Engineering Simplicity

Junos® OS

Junos CLI Reference

Published
2024-12-18

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Junos® OS Junos CLI Reference
Copyright © 2024 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

Junos CLI Reference | cdlxvii

1 What's New in the CLI

What's New in the CLI | 2

2 Configuration Statements by Hierarchy

[edit] | 17

[edit access] | 19

[aaa-options] | 22

[accounting-backup-options] | 22

[address] | 23

[address-assignment] | 23

[diameter] | 26

[domain] | 27

[firewall-authentication] | 28

[group-profile] | 29

[gx-plus] | 30

[ldap-options] | 31

[ldap-server] | 32

[nasreq] | 32

[ocs] | 32

[pcrf] | 34

[profile] | 35

[profileldap-options] | 41

[protocol-attributes] | 41

[radius] | 42

[radius-disconnect] | 43

[radius-options] | 44

[radius-server] | 44

[radsec] | 44

[s6a] | 45

[saml] | 45

[securid-server] | 45

[tunnel-profile] | 46

[tunnel-switch-profile] | 47

[edit accounting-options] | 47

[class-usage-profile] | 48

[file] | 48

[filter-profile] | 49

[flat-file-profile] | 49

[interface-profile] | 50

[mib-profile] | 50

[routing-engine-profile] | 50

[edit aggregated-devices] | 51

[ethernet] | 51

[edit applications] | 52

[application] | 52

[application-set] | 54

[edit bridge-domains] | 54

[<domain-name>] | 55

[bridge-options] | 55

[forwarding-options] | 56

[interface] | 59

[multicast-snooping-options] | 60

[protocols] | 60

[vxlan] | 63

[edit bridge-options] | 63

[edit chassis] | 63

[(EX Series)] | 68

[(QFX Series)] | 69

[afeb] | 69

[aggregated-devices] | 70

[alarm] | 70

[cb] | 71

[cfeb] | 72

[cluster] | 72

[fabric] | 73

[fabric-card] | 74

[feb] | 74

[forwarding-options] | 75

[fpc] | 76

[high-availability] | 84

[interconnect-device] | 85

[lcc] | 86

[lcd-menu] | 88

[member] | 88

[memory-enhanced] | 89

[multicast-loadbalance] | 89

[network-services] | 90

[network-slices] | 90

[node-group] | 91

[nssu] | 92

[pem] | 92

[pseudowire-service] | 93

[psm] | 93

[redundancy] | 93

[redundancy-group] | 94

[routing-engine] | 95

[satellite-management] | 95

[source] | 97

[synchronization] | 98

[tfeb] | 102

[edit class-of-service] | 103

[adaptive-shapers] | 107

[adjustment-control-profiles] | 107

[application-traffic-control] | 108

[class] | 108

[classifiers] | 108

[code-point-aliases] | 110

[congestion-notification-profile] | 110

[drop-profiles] | 112

[fabric] | 113

[forwarding-class-sets] | 113

[forwarding-classes] | 113

[forwarding-classes-interface-specific] | 114

[forwarding-policy] | 114

[fragmentation-map] | 115

[fragmentation-maps] | 115

[host-outbound-traffic] | 116

[interface] | 116

[interfaces] | 117

[loss-priority-maps] | 122

[loss-priority-rewrites] | 123

[multi-destination] | 123

[queue] | 124

[restricted-queues] | 124

[rewrite-rules] | 124

[routing-instances] | 125

[scheduler-map-forwarding-class-sets] | 126

[scheduler-maps] | 126

[schedulers] | 126

[shared-buffer] | 128

[system-defaults] | 128

[traffic-control-profile] | 129

[traffic-control-profiles] | 129

[virtual-channel-groups] | 130

[edit diameter] | 131

[network-element] | 132

[network-statement] | 132

[peer] | 133

[transport] | 133

[edit dynamic] | 133

[policy-options] | 134

[edit dynamic-profiles] | 134

[access-cac] | 136

[class-of-service] | 137

[firewall] | 140

[interfaces] | 144

[logical-systems] | 160

[multicast] | 163

[name] | 163

[policy-options] | 163

[protocols] | 163

[routing-instances] | 167

[routing-options] | 170

[services] | 171

[telemetry] | 172

[variables] | 172

[edit ethernet-switching] | 173

[options] | 173

[edit ethernet-switching-options] | 174

[analyzer] | 175

[bpdu-block] | 176

[dot1q-tunneling] | 177

[mac-notification] | 177

[port-error-disable] | 177

[redundant-trunk-group] | 177

[secure-access-port] | 178

[static] | 180

[storm-control] | 180

[unknown-unicast-forwarding] | 180

[edit event-options] | 181

[destinations] | 182

[event-script] | 182

[policy] | 183

[edit fabric] | 186

[aliases] | 187

[flow-groups] | 187

[protocols] | 187

[resources] | 191

[routing-options] | 191

[edit family] | 191

[ethernet-switching] | 192

[filter] | 192

[edit fc-fabrics] | 193

[fc2] | 194

[interface] | 194

[protocols] | 194

[proxy] | 195

[edit fc-options] | 195

[edit firewall] | 196

[atm-policer] | 197

[family] | 197

[filter] | 202

[hierarchical-policer] | 202

[policer] | 203

[three-color-policer] | 204

[edit forwarding-options] | 205

[access-security] | 209

[accounting] | 209

[accountingoutput] | 211

[analyzer] | 211

[dhcp-relay] | 213

[dhcp-relaygroup] | 235

[enhanced-hash-key] | 236

[evpn-vxlan] | 237

[family] | 237

[hash-key] | 239

[helpers] | 240

[helperstraceoptions] | 243

[interface] | 243

[load-balance] | 243

[monitoring] | 244

[next-hop] | 245

[next-hop-group] | 245

[pfe-sensor] | 246

[port-mirror] | 246

[port-mirroring] | 246

[sampling] | 250

[satellite] | 257

[sflow] | 258

[storm-control] | 258

[storm-control-profiles] | 258

[vlans] | 259

[vxlan-routing] | 259

[edit fpc] | 259

[pic] | 260

[edit hierarchical-policer] | 260

[aggregate] | 261

[premium] | 261

[edit interface] | 262

[(all | {interface-name})] | 262

[gigether-options] | 263

[unit] | 263

[edit interfaces] | 263

[(fxp0 | em0)] | 275

[(ls-] | 275

[(ml-] | 276

[(QFX Series)] | 276

[(rsp0 | rsp1)] | 276

[advisory-options] | 277

[ae] | 277

[aeaggregated-ether-options] | 277

[aeaggregated-sonet-options] | **279**

[aeesi] | **279**

[aeforwarding-options] | **280**

[aeunit] | **280**

[aex] | **280**

[aggregated-ether-options] | **281**

[aggregated-ethernet-options] | **282**

[asaggregated-sonet-options] | **283**

[asi] | **283**

[asiaggregated-inline-services-options] | **283**

[at-atm-options] | **283**

[at-dsl-options] | **284**

[at-unit] | **284**

[atm-options] | **285**

[auto-configure] | **285**

[br-] | **287**

[br-dialer-options] | **288**

[br-isdn-options] | **288**

[br-unit] | **288**

[br-{pim}/0/{port}isdn-options] | **288**

[ce1-] | **289**

[cellular-options] | **289**

[ct1-] | **290**

[demux0] | **290**

[dlunit] | **291**

[ds0-options] | **291**

[e1-options] | **292**

[e3-options] | **292**

[es-options] | **292**

[es-unit] | **292**

[esi] | **293**

[et-] | **293**

[et-fpc/pic/port] | **293**

[ether-options] | **294**

[family] | **295**

[fastether-options] | **295**

[fibrechannel-options] | **296**

[forwarding-class-accounting] | **297**

[fti] | **297**

[ftiunit] | **298**

[ge-] | **298**

[ge-gigether-options] | **298**

[ge-unit] | **299**

[ge-{pim}/0/0] | **299**

[gigether-options] | **300**

[gr-fpc/pic/port] | **303**

[gr-fpc/pic/port.unit] | **303**

[gr-unit] | **303**

[gre] | **304**

[interface-name] | **304**

[interface-range] | **305**

[interface-set] | **307**

[ip-fpc/pic/port] | **308**

[irb] | **308**

[lacp] | **309**

[layer2-policer] | **309**

[link-degrade-monitor] | **309**

[lo0] | **309**

[load-balancing-options] | **310**

[logical-tunnel-options] | **311**

[ls-] | **311**

[ls-mlfr-uni-nni-bundle-options] | **311**

[ls-unit] | **312**

[lsq-] | **312**

[lsq-lsq-failure-options] | **312**

[lsq-unit] | **313**

[mapping-range] | **313**

[media-type] | **313**

[mif] | **314**

[mlfr-uni-nni-bundle-options] | **314**

[mo-] | **315**

[mo-multiservice-options] | **315**

[multi-chassis-protection] | **315**

[multiservice-options] | **315**

[optics-options] | **316**

[otn-options] | **316**

[pic-set] | **318**

[pp0] | **318**

[ppp-options] | 319

[ps] | 320

[ps0] | 320

[pstargeted-options] | 321

[radius-options] | 321

[redundant-ether-options] | 321

[rlsq] | 322

[rlsqredundancy-options] | 322

[rlt] | 323

[rms] | 323

[rmsredundancy-options] | 323

[rsp] | 323

[rspredundancy-options] | 324

[se-] | 324

[serial-options] | 324

[services-options] | 325

[shdsl-options] | 327

[si] | 327

[si-] | 327

[si-unit] | 328

[so-unit] | 328

[sonet-options] | 328

[sp-unit] | 329

[speed] | 329

[t1-options] | 330

[t3-options] | 330

[tdm-options] | 330

[umd0] | 330

[unit] | 331

[unitradio-router] | 350

[vlan] | 350

[xe-] | 351

[xe-,edit] | 351

[xe-fpc/pic/port] | 351

[xe-unit] | 352

[{interface-name] | 352

[{interface-name}redundant-ether-options] | 352

[{interface-name}unit] | 352

[edit jsrc] | 353

[partition] | 353

[edit logical-routers] | 354

[bridge-domains] | 354

[edit logical-systems] | 354

[...] | 357

[access] | 358

[aeaggregated-ether-options] | 359

[applications] | 359

[bridge-domains] | 359

[class-of-service] | 362

[firewall] | 363

[firewallpolicer] | 366

[forwarding-options] | 366

[interfaces] | 377

[l2circuit] | 392

[logical-system-name] | 392

[name] | 394

[neighbor] | 394

[policer] | 394

[policy-options] | 395

[policy-statement] | 398

[protocols] | 399

[realm] | 473

[routing] | 473

[routing-instances] | 473

[routing-instancesrouting-options] | 584

[routing-options] | 584

[security] | 599

[services] | 609

[statement-name] | 610

[switch-options] | 610

[system] | 611

[tenants] | 622

[three-color-policer] | 642

[transport-ls] | 642

[vlans] | 643

[edit multi-chassis] | 645

[multi-chassis-protection] | 645

[edit multicast-snooping-options] | 646

[graceful-restart] | 647

[edit node-group] | 647

[aggregated-devices] | 647

[node-device] | 648

[edit poe] | 648

[(all | fpc {slot-number})] | 649

[fpc] | 649

[interface] | 650

[notification-control] | 651

[edit policy-options] | 651

[application-maps] | 653

[as-path] | 653

[as-path-group] | 654

[community] | 654

[condition] | 654

[defaults] | 655

[policy-statement] | 655

[policy-statementterm] | 658

[policy-statementthen] | 658

[prefix-list] | 658

[satellite-policies] | 659

[then] | 660

[edit protocols] | 661

[(mstp | rstp| vstp)] | 668

[(bgp | esis | isis | ospf | ospf3 | ldp | pim | rip | ripng | rsvp)] | 669

[(bgp | rip | ripng)] | 669

[(esis | isis | ospf | ospf3 | pim)] | **669**

[(isis | ldp | ospf | ospf3 | rsvp)] | **669**

[(mstp | rstp | vstp | vstp vlan {vlan-id})] | **670**

[(mstp | rstp | vstp)] | **670**

[(mstp | rstp)] | **671**

[(ospf | ospf3)] | **671**

[(rstp | vstp)] | **673**

[amt] | **673**

[ancp] | **674**

[authentication-access-control] | **675**

[bfd] | **676**

[bgp] | **676**

[connections] | **701**

[dcbx] | **702**

[dot1x] | **703**

[dvmrp] | **704**

[edge-virtual-bridging] | **705**

[ethernet] | **705**

[evpn] | **705**

[group] | **707**

[iccp] | **707**

[igmp] | **708**

[igmp-snooping] | **710**

[interface] | **712**

[isis] | **712**

[l2-learning] | **716**

[l2circuit] | **717**

[lacp] | **720**

[layer2-control] | **721**

[ldp] | **721**

[link-management] | **725**

[lldp] | **726**

[lldp-med] | **727**

[mld] | **728**

[mld-snooping] | **729**

[mpls] | **730**

[msdp] | **745**

[mstp] | **747**

[mstpmsti] | **749**

[mvpn] | **749**

[mvrp] | **751**

[neighbor-discovery] | **751**

[oam] | **752**

[openflow] | **761**

[ospf] | **762**

[ospf3] | **765**

[ovsdb] | **767**

[pcep] | **767**

[pim] | **768**

[piminterface] | **776**

[ppp] | **776**

[ppp-service] | **776**

[pppoe] | 777

[protection-group] | 778

[ptp] | 779

[rift] | 782

[rip] | 783

[ripng] | 784

[router-advertisement] | 785

[router-discovery] | 786

[rstp] | 787

[rsvp] | 788

[sap] | 791

[sflow] | 791

[source-packet-routing] | 792

[spring-traffic-engineering] | 794

[stp] | 794

[uplink-failure-detection] | 795

[vlan] | 795

[vpls] | 796

[vrrp] | 796

[vstp] | 796

[edit pvlan] | 798

[edit redundant-power-system] | 798

[member] | 799

[edit routing-instances] | 799

[(Multiple Routing Entities)] | 803

[...] | 803

[access] | **804**

[bridge-domains] | **804**

[class-of-service] | **812**

[edit] | **812**

[firewall] | **813**

[flow] | **813**

[forwarding-options] | **813**

[igmp-snooping-options] | **826**

[instance-name] | **826**

[instance-type] | **827**

[interface] | **827**

[logical-systems] | **828**

[multicast-snooping-options] | **828**

[pbb-options] | **829**

[protocol] | **829**

[protocols] | **829**

[provider-tunnel] | **893**

[realm] | **900**

[routing-instance-name] | **900**

[routing-options] | **900**

[service-groups] | **912**

[statement-name] | **912**

[switch-options] | **913**

[system] | **914**

[vlans] | **925**

[vrf-advertise-selective] | **928**

[vrf-table-label] | 928

[vxlan] | 928

[edit routing-instance-rib-inet6.0] | 929

[flow] | 929

[edit routing-options] | 929

[(aggregate | generate | static)] | 934

[(aggregate | generate)] | 934

[access] | 934

[access-internal] | 935

[autonomous-system] | 935

[backup-selection] | 936

[bmp] | 936

[client] | 937

[dynamic-tunnels] | 937

[fate-sharing] | 938

[flex-algorithm] | 939

[flow] | 939

[forwarding-table] | 939

[graceful-restart] | 941

[host-fast-reroute] | 941

[interface-routes] | 941

[multicast] | 942

[nonstop-routing-options] | 944

[outbound-route-filter] | 945

[ppm] | 945

[resolution] | 945

[rib] | 946

[rib-groups] | 948

[source-packet-routing] | 948

[srlg] | 949

[static] | 949

[topologies] | 950

[transport-class] | 950

[validation] | 951

[edit routing-options-rib-inet6.0] | 952

[flow] | 952

[edit schedulers] | 953

[scheduler] | 953

[edit security] | 955

[address-book] | 960

[advance-policy-based-routing] | 960

[advanced-policy-based-routing] | 961

[alarms] | 961

[alg] | 962

[application-firewall] | 965

[application-screen] | 966

[application-tracking] | 966

[authentication-key-chains] | 966

[auto-re-enrollment] | 967

[casb] | 967

[certificates] | 967

[cloud] | 968

[datapath-debug] | **968**

[dynamic-address] | **969**

[dynamic-application] | **969**

[dynamic-vpn] | **970**

[firewall-authentication] | **970**

[flow] | **970**

[forwarding-options] | **973**

[forwarding-process] | **974**

[group-vpn] | **975**

[gtp] | **977**

[host-vpn] | **978**

[idp] | **979**

[ike] | **989**

[ipsec] | **991**

[log] | **994**

[logical-systems] | **996**

[macsec] | **996**

[nat] | **999**

[pki] | **1006**

[policies] | **1007**

[remote-access] | **1013**

[resource-manager] | **1013**

[screen] | **1014**

[sctp] | **1016**

[sip] | **1017**

[softwires] | **1017**

[tcp-encap] | **1017**

[trusted-channel] | **1017**

[tunnel-inspection] | **1018**

[user-identification] | **1018**

[utm] | **1018**

[zones] | **1044**

[edit service-set] | 1045

[service-set-options] | **1046**

[edit services] | 1046

[(CoS)] | **1053**

[(IPsec VPN)] | **1053**

[(Stateful Firewall)] | **1054**

[aacl] | **1054**

[adaptive-services-pics] | **1056**

[address-book] | **1056**

[advanced-anti-malware] | **1056**

[alg] | **1056**

[analytics] | **1059**

[app-engine] | **1061**

[application-identification] | **1062**

[captive-portal] | **1064**

[captive-portal-content-delivery] | **1065**

[cos] | **1065**

[dynamic-flow-capture] | **1068**

[evpn] | **1069**

[flow-collector] | **1070**

[flow-monitoring] | **1071**

[flow-tap] | **1073**

[hcm] | **1074**

[hosted-services] | **1075**

[icap-redirect] | **1076**

[ids] | **1076**

[ike] | **1079**

[inline-monitoring] | **1079**

[ip-reassembly] | **1080**

[ip-reassemblyrule] | **1080**

[ipsec-vpn] | **1081**

[jdaf] | **1084**

[jflow-log] | **1085**

[l2tp] | **1085**

[logging] | **1089**

[lrf] | **1089**

[mobile-edge] | **1091**

[monitoring] | **1092**

[nat] | **1093**

[network-monitoring] | **1100**

[paa] | **1100**

[pcef] | **1101**

[pcp] | **1103**

[pfe] | **1104**

[plmn] | **1104**

[policies] | **1105**

[proxy] | **1105**

[radius-flow-tap] | **1105**

[radius-flow-tappolicy{] | **1107**

[redundancy-set] | **1107**

[rpm] | **1107**

[rtlog] | **1111**

[screen] | **1112**

[security-metadata-streaming] | **1114**

[service-device-pools] | **1114**

[service-interface-pools] | **1114**

[service-set] | **1115**

[service-set-options] | **1120**

[services-set] | **1120**

[softwire] | **1121**

[softwires] | **1121**

[source] | **1122**

[ssl] | **1122**

[stateful-firewall] | **1124**

[traffic-load-balance] | **1125**

[unified-access-control] | **1126**

[url-filter] | **1127**

[user-identification] | **1127**

[video-monitoring] | **1129**

[web-filter] | **1129**

[web-proxy] | **1131**

[edit smtp] | 1131

[traceoptions] | **1132**

[edit snmp] | 1132

[alarm-management] | **1134**

[community] | **1134**

[communitycommunity] | **1135**

[health-monitor] | **1135**

[nonvolatile] | **1136**

[rmon] | **1136**

[routing-instance-access] | **1137**

[trap-group] | **1137**

[trap-options] | **1138**

[v3] | **1138**

[v3target-address] | **1143**

[view] | **1143**

[edit switch-options] | 1144

[interface] | **1145**

[interface-mac-limit] | **1146**

[mac-notification] | **1146**

[mac-table-size] | **1146**

[redundant-trunk-group] | **1146**

[voip] | **1147**

[edit system] | 1147

[access-line] | **1152**

[accounting] | **1153**

[arp] | **1154**

[autoinstallation] | **1154**

[commit] | **1154**

[configuration] | **1155**

[configuration-database] | **1155**

[ddos-protection] | **1156**

[demux-options] | **1158**

[dynamic-profile-options] | **1158**

[extension] | **1159**

[extensions] | **1159**

[fib-streaming] | **1160**

[firewall-profile] | **1160**

[internet-options] | **1160**

[license] | **1160**

[login] | **1161**

[management-traffic] | **1161**

[monitor] | **1162**

[ntp] | **1162**

[packet-forwarding-options] | **1163**

[ports] | **1164**

[processes] | **1165**

[radius-server] | **1168**

[root-authentication] | **1168**

[scripts] | **1168**

[security-profile] | **1172**

[services] | **1173**

[snapshot] | **1201**

[syslog] | **1201**

[system-profile] | **1203**

[tacplus-server] | **1203**

[tracing] | **1204**

[edit tenants] | 1204

[<tenant-name>] | **1205**

[access] | **1206**

[applications] | **1206**

[class-of-service] | **1206**

[interfaces] | **1207**

[name] | **1208**

[nat] | **1208**

[routing-instances] | **1208**

[security] | **1219**

[security-profile] | **1230**

[securitypolicies] | **1230**

[services] | **1231**

[system] | **1231**

[edit unified-edge] | 1231

[aaa] | **1232**

[diameter-profiles] | **1233**

[gateways] | **1234**

[pcef] | **1241**

[tdf] | **1244**

[edit virtual-chassis] | 1245

[member] | **1246**

[edit virtual-network-functions] | 1246

[edit vlans] | 1247

[dot1q-tunneling] | **1248**

[forwarding-options] | **1249**

[interface] | **1253**

[switch-options] | **1253**

[vlan-id] | **1254**

[vxlan] | **1255**

[{vlan-name] | **1255**

[edit vmhost] | 1255

[edit wlan] | 1256

[access-point] | **1256**

3 **Configuration Statements**

Configuration Statements | 1258

4 **CLI Commands**

CLI Commands | 1487

5 **CLI Reference**

0-9 | 1578

3gpp-imsi | **1578**

802.3ad (EX and QFX Series) | **1579**

802.3ad | **1582**

A | 1584

aaa | **1605**

aaa-access-profile (L2TP LNS) | **1607**

aaa clients (TDF) | **1608**

aaa-context (AAA Options) | **1610**

aaa-logical-system (Domain Map) | **1611**

aaa-options (Access Profile) | **1613**

aaa-options (PPP Profile) | **1614**

aaa-policy-control (PCEF Profile) | **1616**

aaa-profile (PCEF Profile) | **1618**

aaa-routing-instance (Domain Map) | **1619**

aacl-fields | **1621**

aacl-statistics-profile | **1623**

abstract-hop | **1625**

accept (Router Advertisement Guard Policy) | **1627**

accept (Dynamic Profile) | **1629**

accept-data | **1631**

accept-join-always-from | **1632**

accept-max-tcp-connections (System Process) | **1634**

accept-out-of-band | **1636**

accept-own | **1637**

accept-prpd-connection | **1639**

accept-remote-nexthop | **1640**

accept-remote-source | **1642**

accept-sdr (PCRF Partition) | **1644**

accept-source-mac (SRX) | **1645**

accept-source-mac | **1647**

accept-srv6-service | **1649**

accepted-prefix-limit | **1651**

access-concentrator | **1655**

access-domain (Remote Device Management) | **1657**

access (Dynamic Access Routes) | **1659**

access-identifier | **1661**

access-interfaces (IFL Subscriber) | **1662**

access-interfaces (IP Subscriber) | **1664**

access-internal (Dynamic Access-Internal Routes) | **1666**

access-internal (Static Access-Internal Routes) | **1668**

access-line (Access-Line Rate Adjustment) | **1669**

access-line-information (L2TP) | **1684**

access-list | **1686**

access-point name | **1687**

access-profile | **1688**

access-profile (AAA Options) | **1690**

access-profile (Domain Map) | **1691**

access-profile (Dynamic Stacked VLAN) | **1693**

access-profile (Dynamic VLAN) | **1694**

access-profile (Extensible Subscriber Services Manager) | **1695**

access-profile-name (Duplicate Accounting) | **1697**

access-profile (Static Subscribers) | **1698**

access-security | **1700**

access (SNMP) | **1702**

access (Static Framed Routes) | **1703**

access-trunk | **1705**

account-layer2-overhead (PIC Level) | **1706**

accounting (Forwarding Options) | **1707**

accounting (interfaces) | **1709**

accounting (AAA Profile) | **1710**

accounting (Access Profile) | **1711**

accounting-backup-options (Access Profile) | **1713**

accounting (Dynamic IGMP Interface) | **1714**

accounting (Dynamic MLD Interface) | **1716**

accounting-options | **1717**

accounting-order | **1721**

accounting-order (Service Accounting) | **1722**

accounting-port (RADIUS Server) | **1724**

accounting-profile (firewall) | **1726**

accounting-profile | **1727**

accounting (Protocols AMT Interface) | **1728**

accounting (Protocols IGMP) | **1730**

accounting (Protocols IGMP AMT Interface) | **1731**

accounting (Protocols IGMP Interface) | **1732**

accounting (Protocols MLD) | **1733**

accounting (Protocols MLD Interface) | **1735**

accounting (RADIUS Client) | **1736**

accounting-secret (RADIUS Server) | **1737**

accounting-server | **1739**

accounting-server-name | **1740**

accounting (Service Accounting) | **1742**

accounting-stop-on-access-deny | **1743**

accounting-stop-on-failure | **1744**

accounting (System) | **1746**

acfc | **1748**

ack-number | **1750**

acknowledge-retries | **1751**

acknowledge-timer | **1752**

acquiring-state-announce-grandmaster | **1754**

action (QFX) | **1755**

action (chassis error) | **1756**

action | **1759**

action (APN GTP) | **1760**

action (OAM) | **1762**

action (OAM LFM) | **1764**

action (Policer) | **1765**

action-priority (Routing Instance) | **1766**

action-priority | **1768**

action-priority (bridge domain) | **1769**

action-profile | **1771**

action-profile (Network Management) | **1773**

action-profile (Applying to CFM) | **1775**

action-profile (Applying to OAM CFM, for EX Series Switch Only) | **1776**

action-profile (Defining for CFM) | **1778**

action-profile (MEP) | **1780**

action-red-differential-delay | **1782**

action (Security Content Security Web Filtering) | **1783**

action (Security Rulebase IPS) | **1785**

action-shutdown | **1787**

action (Uplink Failure Detection) | **1789**

actions | **1790**

actions (Services SSL Initiation) | **1793**

activate | **1795**

activation-attribute (AAA Profile) | **1796**

activation-delay | **1797**

activation-priority | **1799**

active | **1800**

active-directory-access | **1802**

active-interface (VPLS Multihoming) | **1805**

active-leasequery (DHCP Local Server) | **1806**

active-leasequery (DHCP Relay Agent) | **1808**

active-policy | **1811**

active-probe-params | **1813**

active-server-group | **1817**

active-source-limit | **1819**

activeness-probe | **1821**

actual-transit-statistics (Dynamic Profiles) | **1823**

adaptive (protocols mpls) | **1825**

adaptive | **1826**

adaptive (RSVP) | **1828**

adaptive-sample-rate | **1830**

adaptive-services | **1832**

adaptive-services-pics | **1834**

adaptive-shaper | **1835**

adaptive-shapers | **1836**

add-attribute-length-in-pdu | **1838**

add-path | **1839**

add-path-display-ipv4-address | **1842**

address (tunnel-encapsulation) | **1843**

address (SNMP) | **1844**

address | **1846**

address (IP Address) | **1849**

address (Address Book Next Gen Services) | **1852**

address-allocation | **1853**

address (Analytics Collector) | **1855**

address (Anycast RPs) | **1856**

address-assignment (Access) | **1858**

address-assignment (Address-Assignment Pools) | **1862**

address (Bidirectional Rendezvous Points) | **1865**

address-blacklist | **1867**

address-book | **1868**

address-book (System) | **1870**

address-change-immediate-update | **1872**

address (Diameter Peer) | **1874**

address (Diameter Transport) | **1875**

address-family | **1876**

address (Interfaces) | **1878**

address (Interfaces, Flow Monitoring) | **1879**

address (L2TP Destination) | **1880**

address (L2TP Tunnel Destination) | **1882**

address (LNS Local Gateway) | **1883**

address (Local RPs) | **1885**

address (LRF Profile) | **1886**

address-mapping | **1887**

address-mapping (Application Identification) | **1889**

address-mask | **1891**

address-name | **1892**

address (NAT Pool Next Gen Services) | **1894**

address (Peer) | **1895**

address-persistent (Security Source NAT) | **1896**

address-persistent (Security Source NAT Pool) | **1897**

address-pool | **1899**

address-pool (Access) | **1901**

address-pool (Domain Map) | **1902**

address-pooling | **1904**

address-pooling (Security Source NAT) | **1905**

address-pooling (Source NAT Next Gen Services) | **1907**

address-pools | **1908**

address-protection | **1910**

address (Protocols OpenFlow) | **1912**

address (Protocols Router Discovery) | **1913**

address (RADIUS Clients) | **1915**

address (RADIUS Server) | **1916**

address-range | **1917**

address-ranges (Demux) | **1919**

address (Security Address Book) | **1920**

address (Security ARP Proxy) | **1922**

address (Security Destination NAT) | **1924**

address (Security NDP Proxy) | **1925**

address (Services) | **1926**

address (Services Dynamic Flow Capture) | **1928**

address (Services NAT Pool) | **1929**

address-set | **1930**

address-shared (Security Source NAT) | **1932**

address (Static RPs) | **1933**

address (Tunnel Profile Remote Gateway) | **1935**

address (Tunnel Profile Source Gateway) | **1936**

address-whitelist | **1937**

addresses (REST API) | **1938**

adf (Dynamic Firewalls) | **1939**

adjacency-loss | **1941**

adjacency-loss-hold-time (ANCP) | **1943**

adjacency-timer | **1944**

adjust-interval | **1946**

adjust-minimum | **1947**

adjust-minimum (Dynamic Shaping and Scheduling) | **1949**

adjust-percent | **1950**

adjust-percent (Dynamic Schedulers) | **1952**

adjust-threshold | **1953**

adjust-threshold-absolute | **1954**

adjust-threshold-activate-bandwidth | **1956**

adjust-threshold-overflow-limit | **1958**

adjust-threshold-underflow-limit | **1959**

adjustment-control-profiles | **1961**

admin-down | **1962**

admin-email | **1963**

admin-group (Backup Selection) | **1965**

admin-group | **1967**

admin-group-extended | **1968**

admin-group (for Interfaces) | **1971**

admin-group (for LSPs) | **1972**

admin-groups | **1973**

admin-groups-extended | **1975**

admin-groups-extended-range | **1977**

admin-search | **1978**

administrator-email (Security Fallback Block) | **1980**

administrator-email (Security Virus Detection) | **1981**

adsl-bytes | **1982**

adsl2-bytes | **1984**

adsl2-plus-bytes | **1986**

advance-policy-based-routing | **1988**

advance-policy-based-routing-profile (Security Zones) | **1993**

advanced-anti-malware connection | **1994**

advanced-anti-malware policy | **1997**

advanced-anti-malware traceoptions | **2005**

advertise | **2007**

advertise-bgp-static | **2009**

advertise-external | **2011**

advertise-from-main-vpn-tables | **2013**

advertise-inactive | **2015**

advertise-interval | **2017**

advertise-mode (MPLS) | **2018**

advertise-peer-as | **2020**

advertise-prefix-sid | **2022**

advertise-srv6-service | **2024**

advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper) | **2026**

advertise-zero-ethernet-tag | **2028**

advertisement | **2030**

advertisement (Delay Measurement) | **2031**

advertisement-hold-time | **2033**

advisory-options (Traffic Shaping) | **2035**

advpn | **2036**

af-mode | **2039**

af-name | **2040**

age | **2042**

age-of-attack | **2043**

agent-address | **2044**

agent (Analytics) | **2046**

agent-circuit-identifier (Dynamic ACI VLANs) | **2048**

agent-id | **2050**

agent-specifier | **2051**

agf | **2053**

aggregate-clients (DHCP Local Server) | **2056**

aggregate-clients (DHCP Relay Agent) | **2057**

aggregate-clients (Static Subscribers) | **2060**

aggregate-export-interval | **2062**

aggregate (Gigabit Ethernet CoS Policer) | **2063**

aggregate (Hierarchical Policer) | **2064**

aggregate-label | **2066**

aggregate (Protocols RSVP) | **2068**

aggregate (Routing) | **2070**

aggregate-wait-time | **2073**

aggregated-devices | **2074**

aggregated-ether | **2077**

aggregated-ether-options | **2078**

aggregated-inline-services-options (Aggregated Inline Services) | **2082**

aggregation (Forwarding Options) | **2084**

aggregation (IDS) | **2086**

aggregations (IDS Screen Next Gen Services) | **2088**

aging | **2089**

aging-timer (Junos Fusion) | **2091**

aigp | **2092**

aigp-adjust (Policy Action) | **2095**

aigp-originate | **2097**

ais-trigger-condition | **2099**

alarm (chassis) | **2100**

alarm-id | **2102**

alarm-list-name | **2103**

alarm low-light-alarm | **2105**

alarm-management | **2106**

alarm-mode | **2108**

alarm (optics-options) | **2110**

alarm-port | **2111**

alarm (Satellite Policies) | **2114**

alarm (SNMP RMON) | **2115**

alarm-state | **2117**

alarm-threshold | **2119**

alarm-threshold (Security GPRS) | **2120**

alarm-without-drop | **2122**

alarm-without-drop (IDS Screen Next Gen Services) | **2123**

alarms | **2124**

alarms (Video Monitoring) | **2125**

alarms (Security) | **2128**

alg | **2130**

alg (Applications) | **2138**

alg-manager | **2141**

alg-support-lib | **2142**

algorithm | **2144**

algorithm (Junos FIPS) | **2145**

alias (Interfaces) | **2146**

alias (Junos Fusion) | **2148**

alias (Junos Fusion Satellite Device Clustering) | **2150**

aliases | **2151**

aliases (Virtual Chassis) | **2153**

all-defects | **2155**

all-regions | **2157**

all-tcp | **2159**

allow | **2160**

allow-6pe-traceroute | **2162**

allow-active-leasequery (DHCP Local Server) | **2163**

allow-bulk-leasequery (DHCP Local Server) | **2165**

allow-clear (Accounting Options) | **2167**

allow-commands (Op Scripts) | **2169**

allow-dns-reply | **2170**

allow-duplicates | **2172**

allow-dynamic-requests (RADIUS Server) | **2173**

allow-email (Security Fallback Block) | **2175**

allow-email (Security Virus Detection) | **2176**

allow-embedded-icmp | **2177**

allow-fragmentation (protocols mpls) | **2178**

allow-fragmentation | **2179**

allow-fragmentation (fti) | **2181**

allow-icmp-without-flow | **2182**

allow-ip-options (IDS MS-MPC) | **2183**

allow-ip-options (Services Stateful Firewall) | **2186**

allow-ipv6-extension-header (IDS MS-MPC) | **2188**

allow-l3vpn-traceroute-src-select | **2190**

allow-leasequery (DHCP Local Server) | **2191**

allow-maximum (Multicast) | **2193**

allow-multicast | **2195**

allow-no-end-option (DHCP Relay Agent) | **2196**

allow-overlapping-nat-pools | **2198**

allow-overlapping-pools (NAT Next Gen Services) | **2199**

allow-remote-loopback | **2200**

allow-reverse-ecmp | **2201**

allow-route-leaking | **2202**

allow-snooped-clients | **2204**

allow-sram-parity-errors | **2206**

allow-subnet-mismatch | **2207**

allow-transients | **2208**

allow-url-for-python (Op Scripts) | **2209**

allow-v4mapped-packets | **2211**

allowed-destinations | **2212**

allowed-mac | **2214**

allowed-sources (REST API) | **2216**

alt-name (Application Identification) | **2217**

alternative-partition-name (OCS Partition) | **2218**

always-mark-connection-protection-tlv | **2219**

always-write-giaddr | **2221**

always-write-option-82 | **2222**

ambient-temperature | **2224**

amf | **2226**

amt (IGMP) | **2228**

amt (Protocols) | **2230**

analytics | **2231**

analyzer | **2237**

analyzer-address | **2241**

analyzer-id | **2243**

anchor-point (Pseudowire Subscriber Interfaces) | **2244**

ancp | **2247**

ancp (Adjustment Control Profiles) | **2249**

ancp-speed-change-immediate-update (ANCP) | **2251**

annotate | **2252**

anomaly | **2254**

anti-replay-window-size | **2255**

anti-replay-window-size (Services IPsec VPN) | **2256**

anti-replay-window-size (Services Service Set) | **2258**

anti-spam (Logical Systems) | **2259**

anti-spam | **2262**

anti-spam (Security Content Security Policy) | **2265**

anti-virus (Logical Systems) | **2266**

anti-virus (Content Security) | **2269**

anti-virus (Content Security Policy) | **2273**

anti-virus (Services) | **2275**

any-sender | **2281**

any (VPLS Multihoming) | **2282**

anycast-pim | **2283**

anycast-prefix | **2285**

apn | **2286**

apn-control-group (Security GTP) | **2288**

apn-control (Security GTP) | **2289**

apn-services (control plane services) | **2292**

app-mapping-timeout | **2293**

appfw-profile (System) | **2294**

appfw-rule | **2296**

appfw-rule-set | **2298**

application | **2299**

application (Adjustment Control Profiles) | **2301**

application (Application Identification) | **2303**

application (Application Identification) | **2306**

application (Application Maps) | **2308**

application (Applications) | **2310**

application-bypass (Juniper Secure Connect) | **2314**

application (Extensions) | **2316**

application-firewall | **2318**

application-firewall (Application Services) | **2320**

application-group | **2322**

application-group-any (AACL) | **2323**

application-group (Services) | **2325**

application-groups (PCC Rules) | **2326**

application-groups (Services AACL) | **2328**

application-identification (Services) | **2330**

application-identification (IDP Sensor Configuration) | **2336**

application-identification (Application Identification) | **2337**

application-identification-profile (Service Set) | **2343**

application-map | **2345**

application-maps | **2346**

application (NAT Next Gen Services) | **2348**

application-profile | **2349**

application-profile (Services CoS Next Gen Services) | **2351**

application-protocol | **2352**

application-protocol (Applications) | **2355**

application-proxy (Security Content Security) | **2358**

application-screen (Security H323) | **2359**

application-screen (Security MGCP) | **2361**

application-screen (Security SCCP) | **2362**

application-screen (Security SIP) | **2364**

application (Security Alarms) | **2366**

application (Security Custom Attack) | **2367**

application (Security Destination NAT) | **2369**

application (Security IDP) | **2370**

application (Security Policies) | **2371**

application (Security Source NAT) | **2373**

application-services (Security Forwarding Process) | **2375**

application-services (security-metadata-streaming) | **2377**

application-services (Security Policies) | **2378**

application-set | **2382**

application-sets (IDS MS-DPC) | **2383**

application-sets (PCP) | **2384**

application-sets (Services CoS) | **2386**

application-sets (Services NAT) | **2387**

application-sets (Services Stateful Firewall) | **2388**

application-specific (Protocols OSPF) | **2389**

application-status | **2392**

application-system-cache | **2393**

application-system-cache-timeout (Services) | **2395**

application-tracking | **2397**

application-tracking (Security Zones) | **2399**

application-traffic-control | **2400**

application-traffic-control (Application Services) | **2402**

application-unknown | **2404**

applications (Applications) | **2405**

applications (DCBX) | **2406**

applications (Diameter) | **2408**

applications (IDS MS-DPC) | **2409**

applications (PCC Rules) | **2410**

applications (PCP) | **2412**

applications (Services AACL) | **2414**

applications (Services ALGs) | **2415**

applications (Services Application Identification) | **2416**

applications (Services CoS) | **2418**

applications (Services NAT) | **2419**

applications (Services Stateful Firewall) | **2420**

apply-action-profile | **2421**

apply-groups | **2423**

apply-groups (Chassis Cluster) | **2424**

apply-groups-except | **2426**

apply-groups-except (Subscriber Secure Policy) | **2427**

apply-groups (Subscriber Secure Policy) | **2428**

apply-macro | **2430**

apply-path | **2431**

apply-secondary-as-giaddr | **2433**

apply-to-half-close-state | **2435**

apts | **2436**

archival | **2438**

archive (All System Log Files) | **2440**

archive (Individual System Log File) | **2442**

archive (QFabric System) | **2445**

archive-sites | **2446**

archive-sites (Flow Monitoring) | **2448**

archive-sites (Event Policy) | **2449**

area | **2451**

area-range | **2453**

arguments (Event Policy) | **2455**

arguments (Op Scripts) | **2457**

arp | **2458**

arp-detect | **2463**

arp-enhanced-scale | **2465**

arp-inspection | **2466**

arp-inspection (MX Series) | **2469**

arp-l2-validate | **2470**

arp-max-cache | **2471**

arp-new-hold-limit | **2473**

arp policer | **2475**

arp-prefix-limit (Host Fast Reroute) | **2477**

arp-resp | **2479**

arp-system-cache-limit | **2480**

arp-throttle | **2482**

arp-nd-probe-disable | **2484**

arp-nd-probe-failed-log | **2485**

as-external | **2486**

as-list | **2487**

as-override | **2489**

as-path-compare | **2490**

as-path-group | **2492**

as-path (Policy Options) | **2493**

as-path (Routing Options) | **2495**

asm-override-ssm | **2497**

assemble | **2499**

assert-timeout | **2500**

identity-provider attribute-mapping (SAML) | **2501**

assignment-id-format (L2TP LAC) | **2504**

assisted-replication | **2505**

assisted-replication (Deterministic AR Replicator Policy Actions) | **2507**

associate-backup-pe-groups | **2511**

associate-lsp | **2512**

associate-profile | **2513**

association-timeout | **2515**

asymmetric-hold-time | **2516**

asymmetric-hold-time | **2517**

asymmetric-lease-time (DHCP Overrides) | **2519**

asymmetric-prefix-lease-time (DHCP Overrides) | **2521**

asynchronous-notification | **2522**

atm-options | **2524**

attach | **2527**

attack-threshold | **2528**

attack-type (Security Anomaly) | **2529**

attack-type (Security Chain) | **2531**

attack-type (Security IDP) | **2533**

attack-type (Security Signature) | **2540**

attacks (Security Exempt Rulebase) | **2546**

attacks (Security IPS Rulebase) | **2547**

attempts (DHCP Local Server) | **2549**

attribute | **2551**

attributes (Access-Line Rate Adjustment) | **2553**

attributes (Diameter Gx Profiles) | **2555**

attributes (JSRC Attributes) | **2557**

attributes-match (Event Policy) | **2558**

attributes (RADIUS Attributes) | **2560**

audible (Security Alarms) | **2563**

auth-entry | **2564**

auth-only-browser | **2566**

auth-user-agent | **2568**

authentication (Protocols OSPF) | **2570**

authentication (services) | **2572**

authentication | **2574**

authentication (AAA Profile) | **2576**

authentication-algorithm | **2577**

authentication-algorithm (Security IKE) | **2580**

authentication-algorithm (Security IPsec) | **2582**

authentication-algorithm (Services IKE) | **2584**

authentication-algorithm (Services IPsec) | **2585**

authentication (BGP BFD Liveness Detection) | **2588**

authentication (Demux) | **2592**

authentication (DHCP Local Server) | **2593**

authentication (DHCP Relay Agent) | **2595**

authentication-entry-timeout (Identity Management Advanced Query) | **2597**

authentication-entry-timeout (Services User Identification) | **2599**

authentication-key (Protocols MSDP) | **2600**

authentication-key (SNMP) | **2602**

authentication-key | **2603**

authentication-key-chain (Protocols BGP and BMP) | **2605**

authentication-key-chain (Protocols LDP) | **2607**

authentication-key-chain (TWAMP) | **2608**

authentication-key-chains | **2610**

authentication-key-chains (TCP-AO) | **2612**

authentication-key (ICCP) | **2614**

authentication-key (Protocols BGP and BMP) | **2616**

authentication-key (Protocols LDP) | **2618**

authentication-key (Protocols RIP) | **2619**

authentication-key (Protocols RSVP) | **2620**

authentication (LAG) | **2622**

authentication-md5 | **2624**

authentication-method (services) | **2626**

authentication-method | **2627**

authentication-method (Interfaces) | **2629**

authentication-mode | **2630**

authentication-none | **2632**

authentication-order | **2633**

authentication-order (Access Profile) | **2635**

authentication-order (Authenticator) | **2638**

authentication-order (System) | **2641**

authentication-password | **2643**

authentication-protocol | **2645**

authentication (Protocols PIM) | **2647**

authentication (Security Alarms) | **2648**

authentication (Security IPsec) | **2650**

authentication-server | **2652**

authentication-sha | **2653**

authentication-sha224 | **2654**

authentication-sha256 | **2656**

authentication-sha348 | **2657**

authentication-sha512 | **2658**

authentication-source (Services User Identification ClearPass) | **2660**

authentication-source (Services User Identification Device Identity) | **2662**

authentication (Static and Dynamic PPP) | **2664**

authentication (Static Subscribers) | **2665**

authentication-type | **2667**

authentication-type (Protocols RIP) | **2669**

authentication-whitelist | **2671**

authenticator | **2673**

authorization (SNMP) | **2677**

authorization (icap-redirect profile) | **2678**

authorization-order | **2680**

auto-bandwidth | **2682**

auto-bandwidth (MPLS Statistics) | **2684**

auto-bandwidth (MPLS Tunnel) | **2686**

auto-configuration | **2688**

auto-configuration (System) | **2689**

auto-configure | **2692**

auto-configure (Demux) | **2694**

auto-configure (Dynamic VLAN Interface Sets) | **2696**

auto-configure (IPv4) | **2698**

auto-configure (IPv6) | **2700**

auto-configure-trigger interface (ANCP) | **2702**

auto-conversion | **2704**

auto-dad (SLAAC Snooping) | **2705**

auto-derive | **2707**

auto-discovery | **2709**

auto-discovery-only | **2711**

auto-export | **2712**

auto-image-upgrade | **2715**

auto-load-rebalance | **2717**

auto-negotiation (QFX Series) | **2718**

auto-negotiation | **2719**

auto-policing | **2724**

auto-provisioned | **2725**

auto-re-enrollment | **2727**

auto-re-enrollment (Security) | **2728**

auto-reconnect | **2732**

auto-recovery | **2733**

auto-rp | **2735**

auto-satellite-conversion (Junos Fusion) | **2736**

auto-sense | **2738**

auto-service-id | **2739**

auto-snapshot | **2741**

auto-sw-sync | **2743**

auto-sw-update | **2745**

auto-targeted-session | **2748**

autodiscovery | **2750**

autodiscovery-only | **2752**

autoinstallation | **2753**

autoinstallation (JNU Satellite Devices) | **2755**

automatic (Security) | **2757**

automatic-site-id | **2758**

automatic (Source NAT Next Gen Services) | **2760**

autonomous | **2761**

autonomous (Dynamic Router Advertisement) | **2763**

autonomous-system | **2764**

autonomous-system-type (Flow Monitoring) | **2767**

autostate-exclude | **2769**

auxiliary | **2771**

auxiliary-spi | **2773**

auxiliary-spi (Security IPsec) | **2774**

avg-fd-twoway-threshold | **2775**

avg-flr-backward-threshold | **2777**

avg-flr-forward-threshold | **2779**

avg-ifdv-twoway-threshold | **2781**

avira-engine | **2783**

avp (L2TP Tunnel Switching) | **2784**

B | 2786

backoff-period | **2791**

backup-bridge-priority | **2793**

backup-destination | **2795**

backup-interface | **2796**

backup-interface (Layer 2 Circuits) | **2797**

backup-liveness-detection | **2799**

backup (MBGP MVPN) | **2800**

backup-neighbor | **2802**

backup (OCS Partition) | **2804**

backup-on-failure (Accounting Options) | **2806**

backup-options | **2807**

backup-pe-group | **2808**

backup-peer-ip | **2810**

backup-remote-gateway | **2811**

backup-router | **2812**

backup-selection (Protocols IS-IS) | **2814**

backup-selection (Protocols OSPF) | **2819**

backup-spf-options (Protocols IS-IS) | **2821**

backup-spf-options (Protocols OSPF) | **2823**

backups | **2826**

backward-frr-enable | **2827**

bad-inner-header | **2829**

bad-option (IDS Screen Next Gen Services) | **2830**

bandwidth (multicast) | **2831**

bandwidth (storm control) | **2833**

bandwidth-based-metrics (Protocols OSPF) | **2835**

bandwidth-based-metrics | **2837**

bandwidth (DDoS) | **2839**

bandwidth-degradation | **2841**

bandwidth (Fast Reroute, Signaled, and Multiclass LSPs) | **2843**

bandwidth (Inline Services) | **2845**

bandwidth (Interfaces) | **2847**

bandwidth (Interfaces) | **2849**

bandwidth-kbps | **2850**

bandwidth-level | **2851**

bandwidth-limit | **2853**

bandwidth-limit (Hierarchical Policer) | **2855**

bandwidth-limit (Policer) | **2857**

bandwidth-limit (Policer for Gigabit Ethernet Interfaces) | **2860**

bandwidth-model | **2861**

bandwidth (MPC Bandwidth) | **2863**

bandwidth-percent (protocols mpls) | **2864**

bandwidth-percent | **2866**

bandwidth-percentage | **2869**

bandwidth (Protocols Layer 2 Circuit) | **2870**

bandwidth (Protocols RSVP) | **2872**

bandwidth (RFC 2544 Benchmarking) | **2874**

bandwidth-scale (DDoS) | **2875**

bandwidth (Static LSP) | **2876**

bandwidth-threshold | **2878**

bandwidth (Tunnel Services) | **2880**

banner (Access FTP HTTP Telnet Authentication) | **2882**

banner (Access Web Authentication) | **2884**

base-distinguished-name | **2885**

batch query | **2886**

bb-sc-n | **2889**

beacon-period | **2890**

bearer-type (L2TP Tunnel Switching) | **2892**

ber-threshold-clear | **2894**

ber-threshold-signal-degrade | **2897**

bert-algorithm | **2900**

bert-error-rate | **2903**

bert-period | **2905**

best-site | **2908**

bfd | **2909**

bfd (protocols) | **2911**

bfd-liveliness | **2914**

bfd-liveliness (Activeness Probe) | **2917**

bfd-liveness-detection (BGP) | **2918**

bfd-liveness-detection (EVPN Node Detection) | **2923**

bfd-liveness-detection (LAG) | **2926**

bfd-liveness-detection (Layer 2 VPN and VPLS) | **2929**

bfd-liveness-detection (LSP) | **2931**

bfd-liveness-detection (Protocols IS-IS) | **2933**

bfd-liveness-detection (Protocols LDP) | **2937**

bfd-liveness-detection (Protocols MPLS) | **2939**

bfd-liveness-detection (Protocols OSPF) | **2941**

bfd-liveness-detection (Protocols PIM) | **2945**

bfd-liveness-detection (Protocols RIP) | **2947**

bfd-liveness-detection (Routing Options Static Route) | **2950**

bfd-liveness-detection (Source Packet Routing Template) | **2955**

bgp | **2956**

bgp (RPM) | **2958**

bgp-error-tolerance (Protocols BGP) | **2959**

bgp-identifier | **2962**

bgp-orf-cisco-mode | **2963**

bgp-peer | **2965**

bgp-static | **2966**

bgpmcast | **2968**

bidirectional (Interface) | **2969**

bidirectional (RP) | **2971**

bier | **2972**

bier-sub-domain (Protocols IS-IS) | **2974**

bits (Chassis Synchronization) | **2975**

black-list | **2978**

blackhole-action | **2979**

block | **2980**

block-allocation (Source NAT Next Gen Services) | **2983**

block-command | **2985**

block-content-type | **2986**

block-extension | **2987**

block-frag (IDS Screen Next Gen Services) | **2988**

block-message (Application Firewall) | **2989**

block-message (Security Content Security) | **2993**

block-mime | **2994**

block (Spanning Trees) | **2995**

blocklists (Security IKE) | **2997**

bmp | **3000**

boot-loader-authentication | **3008**

boot-server (DHCP) | **3010**

bootp | **3012**

bootp (System Processes) | **3013**

bootstrap | **3015**

bootstrap-export | **3016**

bootstrap-import | **3018**

bootstrap-priority | **3019**

bounded-delay | **3021**

bpdu-block | **3022**

bpdu-block-on-edge | **3024**

bpdu-destination-mac-address | **3026**

bpdu-destination-mac-address (Spanning Tree) | **3027**

bridge-domain | **3029**

bridge-domain (Analyzer) | **3031**

bridge-domains | **3032**

bridge-options | **3034**

bridge-priority | **3036**

bridge-template | **3038**

brief | **3039**

broadcast (Interfaces) | **3041**

broadcast (Router Discovery) | **3043**

broadcast (Forwarding class) | **3044**

bucket-size | **3046**

buffer-dynamic-threshold | **3047**

buffer-partition (Egress) | **3049**

buffer-partition (Ingress) | **3052**

buffer-size | **3055**

buffer-size (Dynamic Scheduling) | **3061**

buffer-size (Schedulers) | **3063**

buffer-rate (Schedulers) | **3065**

bulk-call-mode | **3067**

bulk-leasequery (DHCP Relay Agent) | **3068**

bum-hashing | **3070**

bundle | **3072**

bundle (multilink) | **3073**

bundle (Interfaces) | **3074**

burst (DDoS) | **3075**

burst-scale (DDoS) | **3077**

burst-size | **3078**

burst-size (Default Local Policy) | **3081**

burst-size-limit | **3083**

burst-size-limit (Hierarchical Policer) | **3084**

burst-size-limit (Policer) | **3086**

burst-size-limit (Policer for Gigabit Ethernet Interfaces) | **3089**

burst-size (TDF Domain) | **3090**

by-destination (IDS MS-DPC) | **3092**

by-destination (IDS MS-MPC) | **3093**

by-destination (IDS Screen Next Gen Services) | **3095**

by-pair (IDS MS-DPC) | **3098**

by-protocol (IDS MS-MPC) | **3099**

by-protocol (IDS Screen Next Gen Services) | **3102**

by-source (IDS MS-DPC) | **3105**

by-source (IDS MS-MPC) | **3106**

by-source (IDS Screen Next Gen Services) | **3108**

bypass | **3111**

bypass-aggregate (DDoS) | **3112**

bypass-queuing-chip | **3114**

bypass (Signaled LSP) | **3115**

bypass (Static LSP) | **3117**

bypass-traffic-on-exceeding-flow-limits | **3119**

bypass-traffic-on-pic-failure | **3120**

bytes (Dynamic Traffic Shaping) | **3121**

bytes (otn-options) | **3123**

C | 3124

c-timeout | **3152**

ca-chain (REST API) | **3154**

ca-identity | **3155**

ca-name | **3156**

ca-profile (Security PKI) | **3157**

ca-profile (Services) | **3160**

ca-type | **3161**

cable-length (Congestion Notification) | **3164**

cac (TDF Gateway) | **3165**

cache | **3167**

cache (Security Log) | **3168**

cache-size | **3170**

cache-timeout-negative | **3171**

cacheable (Application Identification) | **3173**

cak | **3174**

cak (MX Series) | **3176**

calculation-weight | **3177**

call-flood | **3179**

call-rate-statistics | **3180**

callback | **3182**

callback-wait-period | **3183**

called-station-id | **3185**

called-station-id (OCS Partition) | **3186**

caller | **3188**

calling-number | **3189**

calling-number (L2TP Tunnel Switching) | **3190**

calling-station-id | **3192**

calling-station-id-format (Subscriber Management) | **3193**

candidate-uplink-port-policy (satellite-policies) | **3196**

capabilities (Protocols RIFT) | **3197**

captive-portal | **3199**

captive-portal-content-delivery-profile (Services) | **3202**

captive-portal (Services UAC) | **3203**

captive-portal (Services UAC Policy) | **3205**

capture-file (Security) | **3206**

capture-group | **3208**

casb | **3210**

casb-policy | **3212**

cascade-port | **3214**

cascade-ports | **3216**

cascade-ports (Junos Fusion Satellite Device Cluster) | **3218**

categories | **3219**

category | **3221**

category (Security Dynamic Attack Group) | **3223**

category (Security Logging) | **3224**

category (Security Web Filtering) | **3226**

category (System Logging) | **3236**

cbr | **3238**

cbr rate | **3240**

ccc | **3241**

ccc (Routing Policy Condition) | **3242**

cdvt | **3244**

cell-mode (Dynamic Traffic Shaping) | **3245**

cell-overhead | **3248**

cellular-options | **3249**

centralized | **3250**

cert-file (Scripts) | **3252**

certificate | **3254**

certificate-authority (REST API) | **3256**

certificate-id | **3257**

certificate (Juniper Secure Connect) | **3259**

certificate-verification | **3260**

certificates | **3262**

certification-authority | **3263**

ces-psn-channel (tdm-options) | **3265**

cfeb | **3267**

cflowd (Discard Accounting) | **3268**

cflowd (Flow Monitoring) | **3270**

cfp-to-et | **3272**

cgn-pic | **3273**

chain | **3274**

chain-order (Application Identification) | **3276**

chained-composite-next-hop | **3277**

challenge-length (Static and Dynamic PPP) | **3281**

challenge-password | **3283**

change-configuration (Event Policy) | **3284**

channel | **3286**

channel-speed | **3287**

chap | **3289**

chap (Dynamic PPP) | **3291**

chap (L2TP) | **3292**

charging-id (OCS Partition) | **3293**

charging-service-list | **3295**

chassis | **3296**

chassis-id (LLDP) | **3301**

chassis-id (Aggregated Ethernet Interface) | **3302**

chassis (QFabric System) | **3303**

chassis (Subscriber Limits) | **3306**

check-bytes (Application Identification) | **3308**

check-zero | **3309**

checksum (Scripts) | **3311**

checksum-validate | **3313**

child-inactivity-timeout | **3314**

children | **3315**

cipher-list (REST API) | **3318**

cipher-suite (MACsec) | **3320**

circuit-id | **3322**

circuit-id (DHCP Relay Agent) | **3324**

circuit-id (VLAN Authentication Username) | **3328**

circuit-type | **3329**

circuit-type (DHCP Local Server) | **3330**

circuit-type (DHCP Relay Agent) | **3332**

cisco-interoperability | **3334**

cisco-nas-port-info (L2TP Tunnel Switching) | **3335**

civic-based | **3337**

ckn | **3339**

ckn (MX Series) | **3341**

class | **3342**

class (Unified Edge) | **3345**

class (CoS-Based Forwarding) | **3346**

class (Defining Login Classes) | **3347**

class (Forwarding Class Sets) | **3361**

class (Forwarding classes) | **3362**

class (Forwarding Classes) | **3364**

class-of-service (Dynamic Profiles) | **3366**

class-of-service (Protocols MPLS) | **3369**

class-of-service (Protocols RSVP) | **3370**

class-of-service (saegw system) | **3372**

class-usage-profile | **3373**

classification-key (DHCP Local Server) | **3375**

classification-key (DHCP Relay Agent) | **3376**

classification-override | **3379**

classifier | **3380**

classifiers (VRF) | **3382**

classifiers (Class of Service) | **3383**

classifiers (Definition) | **3386**

classifiers (Definition) | **3388**

classifiers (Dynamic CoS Application) | **3390**

classifiers (Interfaces) | **3391**

classifiers (Physical Interface) | **3393**

classifiers (Routing Instance) | **3395**

clat-ipv6-prefix-length | **3396**

clat-prefix | **3398**

clat-prefix (Source NAT Next Gen Services) | **3399**

cleanup-interval (Accounting Options) | **3400**

clear | **3402**

clear access-security router-advertisement statistics | **3403**

clear access-security slaac-snooping binding | 3404

clear access-security slaac-snooping statistics | 3406

clear accounting server statistics archival-transfer | 3407

clear-action (CFM) | 3409

clear amt statistics | 3410

clear amt tunnel | 3412

clear ancp access-loop | 3414

clear ancp neighbor | 3416

clear ancp statistics | 3418

clear ancp subscriber | 3421

clear arp | 3424

clear arp inspection statistics | 3426

clear auto-configuration interfaces | 3428

clear auto-configuration interfaces interface-set | 3430

clear bfd adaptation | 3432

clear bfd session | 3434

clear bgp damping | 3436

clear bgp neighbor | 3437

clear bgp table | 3441

clear bridge mac-ip-table | 3443

clear bridge mac-table | 3444

clear bridge recovery-timeout | 3446

clear bridge statistics | 3447

clear captive-portal | 3450

clear chassis alarms fabric degraded | 3453

clear chassis cluster control-plane statistics | 3455

clear chassis cluster data-plane statistics | 3457

clear chassis cluster failover-count | 3458

clear chassis cluster ip-monitoring failure-count | 3460

clear chassis cluster ip monitoring failure count ip address | 3462

clear chassis cluster statistics | 3463

clear chassis display message | 3465

clear chassis fpc errors | 3468

clear chassis hardware history | 3470

clear chassis high-availability data-plane statistics | 3472

clear chassis high-availability information | 3473

clear chassis sib errors | 3474

clear class-of-service application-traffic-control counter | 3476

clear class-of-service application-traffic-control rate-limiters | 3478

clear class-of-service application-traffic-control statistics rule | 3479

clear cli logical-system | 3481

clear ddos-protection protocols | 3482

clear dhcp client binding | 3484

clear dhcp client statistics | 3486

clear dhcp relay active-leasequery statistics | 3488

clear dhcp relay binding (Security) | 3491

clear dhcp relay binding | 3493

clear dhcp relay lockout-entries | 3496

clear dhcp relay statistics (Security) | 3499

clear dhcp relay statistics | 3500

clear dhcp-security binding | 3503

clear dhcp-security ipv6 binding | 3505

clear dhcp server active-leasequery statistics | 3506

clear dhcp server binding | 3508

clear dhcp server lockout-entries | 3512

clear dhcp server statistics | 3515

clear dhcp snooping binding | 3517

clear dhcp snooping statistics | 3519

clear dhcpv6 client binding | 3521

clear dhcpv6 client statistics | 3522

clear dhcpv6 relay active-leasequery statistics | 3524

clear dhcpv6 relay binding | 3527

clear dhcpv6 relay lockout-entries | 3531

clear dhcpv6 relay statistics | 3534

clear dhcpv6 server active-leasequery statistics | 3536

clear dhcpv6 server binding | 3538

clear dhcpv6 server binding (Local Server) | 3542

clear dhcpv6 server lockout-entries | 3543

clear dhcpv6 server statistics | 3546

clear dhcpv6 server statistics (Local Server) | 3548

clear dhcpv6 snooping binding | 3549

clear dhcpv6 snooping statistics | 3551

clear diameter function statistics | 3553

clear diameter peer | 3554

clear-dont-fragment-bit (fti) | 3556

clear-dont-fragment-bit (Interfaces GRE Tunnels) | 3557

clear-dont-fragment-bit (NAT Next Gen Services) | 3559

clear-dont-fragment-bit (Services IPsec VPN) | 3560

clear-dont-fragment-bit (Services NAT Options) | 3561

clear-dont-fragment-bit (Services Service Set) | 3562

clear dot1x | 3564

clear edge-virtual-bridging | 3566

clear error bpdu interface | 3567

clear error mac-rewrite | 3568

clear ethernet-switching bpdu-error interface | 3570

clear ethernet-switching evpn arp-table | 3572

clear ethernet-switching evpn nd-statistics | 3573

clear ethernet-switching evpn nd-table | 3574

clear ethernet-switching layer2-protocol-tunneling error | 3576

clear ethernet-switching layer2-protocol-tunneling statistics | 3577

clear ethernet-switching mac-ip-table | 3579

clear ethernet-switching port-error | 3580

clear ethernet-switching recovery-timeout | 3582

clear ethernet-switching recovery-timeout (SRX) | 3583

clear ethernet-switching statistics mac-learning | 3584

clear ethernet-switching table persistent-learning | 3586

clear evpn duplicate-mac-suppression | 3588

clear evpn mac-ip-table | 3589

clear evpn nd-table | 3590

clear evpn statistics | 3592

clear extensible-subscriber-services counters | 3593

clear extensible-subscriber-services sessions | 3594

clear fibre-channel fc2 statistics | 3596

clear fibre-channel fip enode | 3597

clear fibre-channel fip statistics | 3599

clear fibre-channel fip vn-port | 3600

clear fibre-channel flogi statistics | 3602

clear fibre-channel proxy statistics | 3603

clear fip snooping enode | 3604

clear fip snooping statistics | 3606

clear fip snooping vlan | 3608

clear fip vlan-discovery statistics | 3610

clear firewall (QFX) | 3611

clear firewall | 3613

clear forwarding-options port-mirroring instance family | 3616

clear gtp tunnels | 3618

clear igmp membership | 3619

clear igmp snooping membership | 3623

clear igmp snooping statistics | 3625

clear igmp statistics | 3627

clear-ike-sas-on-pic-restart | 3631

clear ike security-associations | 3632

clear ilmi statistics | 3634

clear inband-flow-telemetry stats | 3635

clear interfaces aeX forwarding-options load-balance state | 3637

clear interfaces aggregate forwarding-options load-balance state | 3638

clear interfaces interface-set statistics | 3640

clear interfaces interval | 3641

clear interfaces lb-stats af | 3643

clear interfaces mac-database | 3648

clear interfaces mac-database statistics | **3649**

clear interfaces statistics | **3651**

clear interfaces statistics swfabx | **3653**

clear interfaces transport pm | **3654**

clear-ipsec-sas-on-pic-restart | **3656**

clear ipsec security-associations | **3657**

clear ipv6 neighbors | **3659**

clear ipv6 neighbors (SRX Series) | **3661**

clear ipv6 router-advertisement | **3663**

clear isis adjacency | **3664**

clear isis database | **3667**

clear isis overload | **3670**

clear isis spring traffic-statistics | **3673**

clear isis statistics | **3674**

clear l2-learning mac-move-buffer | **3677**

clear lacp statistics interfaces | **3679**

clear ldp neighbor | **3680**

clear ldp session | **3682**

clear ldp statistics | **3683**

clear lldp neighbors | **3685**

clear lldp statistics | **3686**

clear log | **3688**

clear log (JDM) | **3690**

clear log (Next Gen Services) | **3692**

clear loop-detect enhanced interface | **3693**

clear mac-vrf forwarding mac-ip-table | **3695**

clear mac-vrf forwarding mac-learning-log | 3697

clear mac-vrf forwarding mac-table | 3698

clear mac-vrf forwarding recovery-timeout | 3700

clear mac-vrf forwarding redundancy-group | 3701

clear mld membership | 3702

clear mld snooping membership | 3704

clear mld snooping statistics | 3705

clear mld statistics | 3707

clear monitor security flow filter | 3709

clear mpls container-lsp | 3711

clear mpls lsp | 3713

clear mpls tunnel manager statistics | 3716

clear msdp cache | 3717

clear msdp statistics | 3719

clear multicast bandwidth-admission | 3720

clear multicast forwarding-cache | 3723

clear multicast scope | 3724

clear multicast sessions | 3726

clear multicast snooping statistics | 3728

clear multicast statistics | 3730

clear mvrp statistics | 3732

clear neighbor-discovery-inspection statistics | 3734

clear network-access aaa statistics | 3736

clear network-access aaa subscriber | 3741

clear network-access gx-plus replay | 3743

clear network-access gx-plus statistics | 3745

clear network access nas-5g statistics | 3746

clear network-access ocs statistics | 3747

clear network-access pcrf | 3749

clear network-access radsec state | 3751

clear network-access radsec statistics | 3752

clear network-access requests pending | 3753

clear network-access requests statistics | 3755

clear network-access securid-node-secret-file | 3756

clear network-access aaa saml assertion-cache | 3757

clear network-access aaa saml idp-metadata | 3759

clear node reboot | 3760

clear oam ethernet connectivity-fault-management continuity-measurement | 3761

clear oam ethernet connectivity-fault-management delay-statistics | 3763

clear oam ethernet connectivity-fault-management event | 3765

clear oam ethernet connectivity-fault-management linktrace path-database | 3767

clear oam ethernet connectivity-fault-management loss-statistics | 3768

clear oam ethernet connectivity-fault-management path-database | 3770

clear oam ethernet connectivity-fault-management policer | 3772

clear oam ethernet connectivity-fault-management sla-iterator-statistics | 3774

clear oam ethernet connectivity-fault-management statistics | 3775

clear oam ethernet connectivity-fault-management synthetic-loss-measurement | 3777

clear oam ethernet link-fault-management state | 3779

clear oam ethernet link-fault-management statistics | 3781

clear-on-abort (DHCP Local Server) | 3782

clear (ospf | ospf3) database | 3784

clear (ospf | ospf3) database-protection | 3788

clear (ospf | ospf3) io-statistics | 3789

clear (ospf | ospf3) neighbor | 3791

clear (ospf | ospf3) overload | 3793

clear (ospf | ospf3) statistics | 3795

clear ovsdb commit failures | 3798

clear passive-monitoring statistics | 3800

clear path-computation-client statistics | 3802

clear performance-monitoring mpls lsp | 3803

clear pfe sflow egress-mcast-stats | 3805

clear pim join | 3806

clear pim join-distribution | 3809

clear pim register | 3812

clear pim snooping join | 3814

clear pim snooping statistics | 3817

clear pim statistics | 3819

clear policy statistics | 3823

clear pppoe lockout | 3824

clear pppoe lockout atm-identifier | 3827

clear pppoe lockout vlan-identifier | 3829

clear pppoe statistics | 3832

clear protection-group ethernet-ring statistics | 3834

clear ptp statistics stream | 3836

clear remote-device-management statistics | 3837

clear rip general-statistics | 3838

clear rip statistics | 3840

clear ripng general-statistics | 3842

clear ripng statistics | **3844**

clear rsvp session | **3846**

clear rsvp statistics | **3848**

clear security advance-policy-based-routing sla statistics | **3850**

clear security alarms | **3851**

clear security alg h323 counters | **3855**

clear security alg ike-esp-nat | **3857**

clear security alg sccp calls | **3858**

clear security alg sccp counters | **3860**

clear security alg sip calls | **3862**

clear security alg sip counters | **3863**

clear security application-firewall rule-set statistics | **3865**

clear security application-firewall rule-set statistics logical-system | **3866**

clear security datapath-debug counters | **3868**

clear security dns-cache | **3869**

clear security dynamic-vpn all | **3871**

clear security dynamic-vpn user | **3872**

clear security firewall-authentication history | **3874**

clear security firewall-authentication history address | **3876**

clear security firewall-authentication history identifier | **3878**

clear security firewall-authentication users | **3881**

clear security firewall-authentication users address | **3883**

clear security firewall-authentication users identifier | **3885**

clear security flow ip-action | **3887**

clear security flow session | **3890**

clear security flow session application | **3893**

clear security flow session application-traffic-control | **3896**

clear security flow session conn-tag | **3899**

clear security flow session destination-port | **3901**

clear security flow session destination-prefix | **3903**

clear security flow session family | **3905**

clear security flow session IDP | **3907**

clear security flow session interface | **3910**

clear security flow session nat | **3912**

clear security flow session protocol | **3915**

clear security flow session resource-manager | **3918**

clear security flow session services-offload | **3920**

clear security flow session session-identifier | **3924**

clear security flow session source-port | **3926**

clear security flow session source-prefix | **3928**

clear security flow session tenant | **3931**

clear security flow session tunnel | **3933**

clear security forward-options mirror filter | **3935**

clear security group-vpn member group | **3937**

clear security group-vpn member ike security-associations | **3938**

clear security group-vpn member ipsec security-associations | **3939**

clear security group-vpn member ipsec security-associations statistics | **3941**

clear security group-vpn member ipsec statistics | **3942**

clear security group-vpn member kek security-associations | **3943**

clear security group-vpn server | **3944**

clear security group-vpn server server-cluster statistics | **3946**

clear security group-vpn server statistics | **3947**

clear security gtp counters | 3949

clear security gtp sctp association | 3952

clear security gtp sctp counters | 3954

clear security host-vpn security-associations | 3956

clear security idp | 3957

clear security idp attack table | 3959

clear security idp counters action | 3961

clear security idp counters application-identification | 3962

clear security idp counters dfa | 3964

clear security idp counters flow | 3965

clear security idp counters http-decoder | 3966

clear security idp counters ips | 3968

clear security idp counters log | 3969

clear security idp counters memory | 3971

clear security idp counters packet | 3972

clear security idp counters packet-log | 3974

clear security idp counters pdf-decoder | 3975

clear security idp counters policy-manager | 3977

clear security idp counters ssl-inspection | 3978

clear security idp counters tcp-reassembler | 3980

clear security idp ssl-inspection session-id-cache | 3981

clear security ike active-peer aaa-username | 3983

clear security ike peers | 3984

clear security ike peers statistics | 3985

clear security ike respond-bad-spi-count | 3987

clear security ike security-associations | 3988

clear security ike stats | **3990**

clear security ipsec security-associations | **3993**

clear security ipsec statistics | **3995**

clear security ipsec tunnel-events-statistics | **3997**

clear security log | **3998**

clear security log file | **4001**

clear security log stream file | **4002**

clear security mka statistics | **4004**

clear security mka statistics (MX Series) | **4006**

clear security nat incoming-table | **4007**

clear security nat source persistent-nat-table | **4009**

clear security nat statistics destination pool | **4010**

clear security nat statistics destination rule | **4011**

clear security nat statistics source pool | **4013**

clear security nat statistics source rule | **4014**

clear security nat statistics static rule | **4015**

clear security pki ca-certificate | **4017**

clear security pki certificate-request | **4018**

clear security pki crl | **4020**

clear security pki key-pair | **4021**

clear security pki key-pair (Local Certificate) | **4023**

clear security pki local-certificate | **4024**

clear security pki local-certificate | **4026**

clear security pki local-certificate (Device) | **4027**

clear security pki node-local key-pair | **4030**

clear security pki node-local local-certificate | **4031**

clear security pki statistics | 4033

clear security policies file-serialization counters | 4034

clear security policies hit-count | 4035

clear security policies statistics | 4037

clear security screen statistics | 4038

clear security screen statistics interface | 4039

clear security screen statistics zone | 4042

clear security ssh key-pair-identity | 4044

clear security tcp-encap statistics | 4046

clear security user-identification local-authentication-table | 4047

clear security utm anti-spam statistics | 4048

clear security utm antivirus statistics | 4051

clear security utm content-filtering statistics | 4054

clear security utm session | 4057

clear security utm web-filtering statistics | 4058

clear service l2tp session routing-instance | 4062

clear service user-identification identity-management counter | 4063

clear services accounting statistics inline-jflow | 4064

clear services advanced-anti-malware statistics | 4066

clear services agf nas statistics | 4068

clear services agf ngap amf | 4069

clear services alg statistics | 4070

clear services application-aware-access-list statistics | 4071

clear services application-identification application-statistics | 4072

clear services application-identification application-statistics cumulative | 4073

clear services application-identification application-statistics interval | 4074

clear services application-identification application-system-cache | **4075**

clear services application-identification application-system-cache (Junos OS) | **4076**

clear services application-identification counter (Values) | **4078**

clear services application-identification packet-capture counters | **4081**

clear services application-identification statistics | **4082**

clear services captive-portal-content-delivery statistics | **4085**

clear services cos statistics | **4086**

clear services crtp statistics | **4088**

clear services dns-filtering cache | **4089**

clear services dynamic-flow-capture | **4093**

clear services flow-collector statistics | **4094**

clear services icap-redirect statistic | **4096**

clear services ids | **4098**

clear services ids destination-table | **4100**

clear services ids pair-table | **4101**

clear services ids source-table | **4103**

clear services inline-monitoring statistics | **4105**

clear services inline nat pool | **4106**

clear services inline nat statistics | **4107**

clear services inline softwire statistics | **4109**

clear services ipsec-vpn certificates | **4110**

clear services ipsec-vpn ike security-associations | **4112**

clear services ipsec-vpn ipsec security-associations | **4113**

clear services ipsec-vpn ipsec statistics | **4115**

clear services l2tp destination | **4117**

clear services l2tp destination lockout | **4119**

clear services l2tp destination statistics | **4121**

clear services l2tp multilink | **4123**

clear services l2tp session | **4125**

clear services l2tp session statistics | **4129**

clear services l2tp tunnel | **4131**

clear services l2tp tunnel statistics | **4134**

clear services lrf collector statistics | **4136**

clear services lrf statistics | **4138**

clear services mobile-edge sessions | **4139**

clear services monitoring rfc2544 | **4140**

clear services monitoring twamp server control-connection | **4141**

clear services nat flows | **4142**

clear services nat mappings | **4144**

clear services nat mappings app | **4147**

clear services nat mappings eim | **4149**

clear services nat mappings pcp | **4151**

clear services nat source mappings | **4154**

clear services redundancy-set last-saved-state id | **4157**

clear services rpm rfc2544-benchmarking | **4158**

clear services rpm twamp server connection | **4159**

clear services security-intelligence dns-statistics | **4160**

clear services security-metadata-streaming | **4161**

clear services service-set statistics ids drops | **4164**

clear services service-sets statistics ids session-limits counters | **4165**

clear services service-sets statistics integrity-drops | **4166**

clear services service-sets statistics jflow-log | **4168**

clear services service-sets statistics packet-drops | **4170**

clear services service-sets statistics syslog | **4171**

clear services sessions | **4173**

clear services sessions analysis | **4178**

clear services ssl proxy statistics | **4179**

clear services stateful-firewall flows | **4180**

clear services stateful-firewall sip-call | **4183**

clear services stateful-firewall sip-register | **4187**

clear services stateful-firewall statistics | **4191**

clear services subscriber analysis | **4193**

clear services user-identification active-directory-access | **4194**

clear services user-identification authentication-table | **4195**

clear services user-identification logical-domain-identity-management counters | **4197**

clear services video-monitoring mdi errors fpc-slot | **4199**

clear services video-monitoring mdi statistics fpc-slot | **4200**

clear-services-web-filter-dns-feed-statistics-feed-name | **4201**

clear services web-filter statistics profile | **4202**

clear sflow collector statistics (QFX Series) | **4204**

clear sflow collectors statistics | **4206**

clear snmp history | **4207**

clear snmp statistics | **4208**

clear spanning-tree protocol-migration | **4210**

clear spanning-tree statistics | **4212**

clear spanning-tree statistics bridge | **4214**

clear spanning-tree stp-buffer | **4215**

clear synchronous-ethernet esmc statistics | **4216**

clear synchronous-ethernet wait-to-restore | **4217**

clear system commit | **4219**

clear system commit prepared | **4221**

clear system errors | **4223**

clear system login lockout | **4225**

clear system reboot | **4226**

clear system services dhcp binding | **4230**

clear system services dhcp conflict | **4232**

clear system services dhcp statistics | **4233**

clear system services dns dns-proxy | **4235**

clear system subscriber-management statistics | **4236**

clear tcp-forwarding connections | **4237**

clear tcp-forwarding statistics | **4241**

clear-threshold | **4245**

clear trace | **4246**

clear unified-edge tdf aaa radius client statistics | **4248**

clear unified-edge tdf aaa radius network-element statistics | **4250**

clear unified-edge tdf aaa radius server statistics | **4252**

clear unified-edge tdf aaa radius snoop-segment statistics | **4253**

clear unified-edge tdf aaa statistics | **4255**

clear unified-edge tdf address-assignment pool | **4257**

clear unified-edge tdf address-assignment statistics | **4259**

clear unified-edge tdf call-admission-control statistics | **4261**

clear unified-edge tdf diameter network-element statistics | **4262**

clear unified-edge tdf diameter pcc-gx statistics | **4264**

clear unified-edge tdf diameter peer statistics | **4266**

clear unified-edge tdf statistics | **4268**

clear unified-edge tdf subscribers | **4269**

clear unified-edge tdf subscribers peer | **4271**

clear validation database | **4273**

clear validation session | **4275**

clear validation statistics | **4277**

clear virtual-chassis heartbeat (MX Series Virtual Chassis) | **4278**

clear virtual-chassis protocol statistics | **4281**

clear virtual-chassis vc-port statistics | **4283**

clear vpls mac-address | **4285**

clear vpls mac-move-action | **4287**

clear vpls mac-table | **4288**

clear vrrp | **4290**

client | **4292**

client (TWAMP) | **4295**

client (interfaces) | **4299**

client (Unified Edge) | **4300**

client-address | **4302**

client-config (Juniper Secure Connect) | **4303**

client-delegate-probes | **4307**

client-discover-match (DHCP Local Server) | **4308**

client-discover-match (DHCP Relay Agent) | **4310**

client-group | **4312**

client-ia-type | **4313**

client-id (DHCP Local Server) | **4315**

client-id (DHCP Relay Agent) | **4317**

client-identifier (DHCP Client) | **4319**

client-identifier (DHCPv6 Client) | **4321**

client-identifier (Interfaces) | **4322**

client-identity (NETCONF TLS) | **4324**

client-idle-timeout (Access Profile) | **4326**

client-list | **4328**

client-list (SNMP) | **4329**

client-list-name | **4331**

client-name-filter | **4332**

client-negotiation-match (DHCPv6 Local Server) | **4333**

client-negotiation-match (DHCPv6 Relay Agent) | **4335**

client-response-ttl | **4336**

client-session-timeout (Access Profile) | **4338**

client (System Services) | **4339**

client-type | **4340**

client-type (Subscriber Limits) | **4342**

clients (SNMP Community) | **4344**

clients | **4345**

clients (Security) | **4347**

clns-routing | **4349**

clns-updown-compatibility | **4350**

clock-class | **4352**

clock-class-to-quality-level-mapping | **4353**

clock-ipv6-client | **4355**

clock-ipv6-source | **4356**

clock-mode | **4357**

clock-mode (Clock Synchronization) | **4359**

clock-rate | **4360**

clock-source (hybrid) | **4362**

clocking-mode | **4364**

close-timeout | **4365**

cluster | **4367**

cluster (Chassis) | **4369**

cluster-id (Junos Fusion Satellite Device Cluster) | **4373**

cluster (Junos Fusion) | **4375**

cmcast-joins-limit-inet (MVPN Selective Tunnels) | **4376**

cmcast-joins-limit-inet6 (MVPN Selective Tunnels) | **4379**

coa-accounting (AAA Profile) | **4381**

coa-immediate-update | **4382**

coa-no-override service-class-attribute | **4384**

code | **4385**

code (IDP) | **4386**

code (AAA Profile) | **4388**

code (Application Identification) | **4389**

code-point | **4390**

code-point-aliases | **4392**

code-point (Congestion Notification) | **4393**

code-point (Fibre Channel Interfaces) | **4395**

code-point (Input Congestion Notification) | **4397**

code-point (Output Congestion Notification) | **4398**

code-point (Rewrite Rules) | **4400**

code-points (Application Maps) | **4402**

code-points (Classifiers) | **4403**

code-points (Classifiers) | **4405**

collapsed-forward | **4406**

collector | **4408**

collector (Flow Collection) | **4410**

collector (Analytics) | **4412**

collector (Flow Monitoring Logs for NAT) | **4413**

collector (Flow Template Profiles for NAT) | **4415**

collector-group (Flow Monitoring Logs for NAT) | **4416**

collector-group (Flow Template Profiles for NAT) | **4418**

collector (Inline Monitoring) | **4420**

collector (LRF Profile) | **4422**

collector (LRF Rule) | **4424**

color | **4425**

color-aware | **4428**

color-blind | **4429**

command (Op Scripts) | **4431**

commands (Event Policy Change Configuration) | **4432**

commands (Event Policy Execute Commands) | **4434**

commit | **4436**

commit activate | **4442**

commit (Commit Scripts) | **4444**

commit-delay | **4446**

commit-interval | **4447**

commit-options | **4449**

commit prepare | **4451**

commit-synchronize-server | **4452**

commit (System) | **4455**

committed-burst-size | **4458**

committed-information-rate | **4460**

community | **4462**

community-name | **4464**

community (Policy Options) | **4465**

community (Protocols Layer 2 Circuit) | **4469**

community (RMON) | **4471**

community (Routing Options) | **4472**

community (SNMP) | **4475**

community-vlan | **4477**

community-vlans (MX Series) | **4478**

compatibility (Application Identification) | **4480**

compatibility-version | **4481**

compliance (Juniper Secure Connect) | **4483**

compress (Accounting Options) | **4488**

compress-configuration-files (System) | **4489**

compression | **4491**

compression-device (Interfaces) | **4492**

compression-device (Interfaces) | **4493**

compression (PPP Properties) | **4494**

compute-options | **4496**

compute-profile | **4497**

concurrent-data-sessions | **4500**

condition | **4502**

conditional-match | **4504**

conditional-metric | **4506**

confederation | **4508**

configuration-database | **4510**

configuration-database (Enhanced Subscriber Management) | **4513**

configuration-file | **4514**

configuration-name | **4516**

configuration-servers | **4517**

configuration-synchronize (Chassis Cluster) | **4519**

configure | **4520**

configured-flow-control | **4523**

congestion-notification-profile | **4525**

connect-actively | **4529**

connection-flood | **4531**

connection (Identity Management Advanced Query) | **4533**

connection-limit | **4537**

connection-limit (REST API) | **4538**

connection-protection | **4540**

connection-protection-tlv | **4541**

connections (Circuits) | **4542**

connections (Host VPN) | **4544**

connections (MPLS) | **4547**

connectivity-association | **4548**

connectivity-association (MACsec Interfaces) | **4551**

connectivity-association (MACsec Interfaces for MX Series) | **4553**

connectivity-association (MX Series) | **4554**

connectivity-fault-management | **4556**

connectivity-fault-management-switches | **4560**

connectivity-type | **4563**

console (System Logging) | **4565**

console (System Ports) | **4567**

constant | **4570**

constituent-list | **4571**

cont-stats-collection-interval | **4573**

contact (SNMP) | **4575**

container-devices | **4576**

container-label-switched-path | **4577**

content-destination | **4579**

content-filtering (Logical System Security Feature Profile) | **4580**

content-filtering (Security Content Security Policy) | **4583**

content-filtering (Security Feature Profile) | **4586**

Content Security | **4589**

Content Security-policy (Application Services) | **4600**

content-size | **4601**

content-size-limit | **4603**

content-size (Security Antivirus Sophos Engine) | **4604**

context (Application Identification) | **4606**

context-identifier (Protocols IS-IS) | **4611**

context-identifier (Protocols OSPF) | **4613**

context (Security Custom Attack) | **4614**

context (SNMPv3) | **4615**

continuity-check | **4616**

continuity-check (EX Series Switch Only) | **4618**

control-channel | **4620**

control-channel (Protocols Link Management Peer) | **4621**

control-channel (Protocols OAM) | **4623**

control-connection (Junos OS) | **4624**

control-connection (Junos OS Evolved) | **4627**

control-cores | **4631**

control-link-recovery | **4633**

control-plane-bandwidth-percent (Node Slicing) | **4634**

control-polarity | **4636**

control-ports | **4637**

control (REST API) | **4639**

control-signal | **4640**

control-source | **4641**

control-vlan | **4643**

control-word | **4644**

control-word (Protocols Layer 2 Circuit Neighbor) | **4647**

control-word (Protocols Layer 2 VPN) | **4648**

controller | **4650**

controller (OVSDB) | **4652**

controller (Protocols OpenFlow) | **4654**

convert-clock-class-to-quality-level | **4656**

convert-sa-to-msdp | **4657**

convey-loss-threshold | **4658**

copy | **4660**

copy-dont-fragment-bit (Services IPsec VPN) | **4661**

copy-dont-fragment-bit (Services Set) | **4662**

copy-plp-all | **4664**

copy-tos-to-outer | **4665**

copy-tos-to-outer-ip-header | **4666**

copy-tos-to-outer-ip-header | **4668**

copy-tos-to-outer-ip-header-transit | **4669**

core-dump | **4670**

core-facing | **4672**

corouted-bidirectional | **4673**

corouted-bidirectional-passive | **4674**

corrupt-file | **4676**

cos-rule-sets (Service Set Next Gen Services) | **4677**

cos-rules (Service Set) | **4679**

cos-rules (Service Set Next Gen Services) | **4680**

cost | **4681**

count (Protocols PIM) | **4684**

count | **4685**

count (HTTP Header Enrichment) | **4686**

count (Security Custom Attack) | **4688**

count (Security Policies) | **4689**

counter | **4690**

counters | **4691**

cpu | **4693**

cpu-factor (Dampen Options) | **4694**

cpu-load-threshold | **4696**

cpu (resource-manager) | **4697**

cpu (TDF Gateway) | **4699**

cpu-throttle (Next Gen Services) | **4701**

craft-lockout | **4703**

crb-proxy-mac | **4704**

create-new-ucast-tunnel | **4707**

create-req | **4708**

credibility | **4710**

credit (Interfaces) | **4712**

critical-event (Process) | **4713**

critical-event (System) | **4715**

crl | **4717**

crl (Adaptive Services Interface) | **4718**

crl (Encryption Interface) | **4720**

crl (Security) | **4721**

cross-connect | **4724**

cross-connect-ccm | **4726**

cryptographic-address | **4728**

csds | **4729**

cspf-backoff-time | **4731**

cts | **4732**

cts-polarity | **4734**

current-hop-limit | **4735**

current-hop-limit (Dynamic Router Advertisement) | **4737**

custom-attack | **4738**

custom-attack-group | **4746**

custom-attack-groups (Security IDP) | **4747**

custom-attacks | **4748**

custom-block-message | **4750**

custom-ciphers | **4751**

custom-message (Security Content Filtering) | **4758**

custom-message (Security Email Notify) | **4759**

custom-message (Security Fallback Block) | **4761**

custom-message (Security Fallback Non-Block) | **4762**

custom-message (Security Virus Detection) | **4763**

custom-message (Security Web Filtering) | **4764**

custom-message-subject (Security Email Notify) | **4766**

custom-message-subject (Security Fallback Block) | **4767**

custom-message-subject (Security Fallback Non-Block) | **4769**

custom-message-subject (Security Virus Detection) | **4770**

custom-objects | **4771**

custom-options | **4773**

custom-page | **4777**

custom-page-file | **4778**

custom-tag-string | **4780**

custom-url-category | **4781**

customer-vlans | **4783**

customization (SNMP) | **4784**

cut-through | **4786**

cvss-score | **4787**

cycle-time | **4788**

D | 4790

dad-proxy | **4809**

dampen | **4811**

dampen (Protocols MVPN) | **4813**

dampen-options | **4814**

damping (BFD Liveness Detection) | **4815**

damping (Interfaces) | **4818**

damping (Policy Options) | **4820**

damping (Protocols BGP) | **4822**

data-channel | **4825**

data-cores | **4827**

data-encapsulation | **4828**

data-fill | **4830**

data-fill-with zeros | **4831**

data-flow-affinity | **4832**

data-format | **4834**

data-forwarding | **4835**

data (FTP) | **4837**

data-length | **4838**

data-mdt-reuse | **4840**

data-rate | **4841**

data-record-fields | **4842**

data-size | **4847**

data-tlv-size | **4849**

database | **4850**

database-protection | **4852**

database-replication (Subscriber Session Database) | **4854**

datapath-debug | **4855**

dcbx | **4858**

dcbx-version | **4860**

dcd | **4861**

dcd-polarity | **4863**

dce-options | **4864**

ddos-protection | **4865**

ddos-protection (DDoS) | **5137**

deactivate | **5142**

deactivation-attribute (AAA Profile) | **5143**

deactivation-delay | **5144**

dead-criteria-retries (RADIUS Server) | **5146**

dead-peer-detection | **5147**

dead-peer-detection (Security Group VPN Server) | **5149**

dead-peer-detection (Services IPsec VPN) | **5151**

deaggregate | **5153**

debounce-interval | **5154**

debug-level (System Services) | **5155**

debug-log (System Services) | **5157**

decapsulate-accept-inner-vlan | **5158**

decapsulate (Firewall Filter) | **5160**

decapsulate-inner-vlan | **5163**

decompress-layer | **5165**

decompress-layer-limit | **5166**

decryption-failures | **5168**

dedicated-buffer | **5170**

dedicated-buffer-profile | **5172**

dedicated-ukern-cpu (BFD) | **5174**

deep-vlan-qualified-learning | **5176**

default-action (DHCP Relay Agent Option) | **5177**

default-action (Protocols OpenFlow) | **5179**

default-actions | **5180**

default-address-selection | **5182**

default-chap-secret | **5184**

default-client-identity (NETCONF TLS) | **5185**

default (CoS) | **5187**

default (CoS Host Outbound Traffic) | **5189**

default-gateway | **5190**

default-isid | **5191**

default-lifetime | **5193**

default-lifetime (Dynamic Router Advertisement) | **5194**

default-local-policy | **5195**

default-lsa | **5197**

default-pap-password | **5199**

default-peer | **5200**

default-policy | **5202**

default-policy (IDP) | **5203**

default-pool (Address Pools) | **5205**

default-prefixes (Protocols RIFT) | **5206**

default-profile | **5208**

default-profile (Juniper Secure Connect) | **5209**

default-rule | **5211**

default (Security Antivirus) | **5213**

default (Security Antivirus Sophos Engine) | **5214**

default (Security Content Security) | **5216**

default (Security Web Filtering) | **5217**

default-trusted-ca-certs (Security) | **5218**

default-value | **5221**

default-vpn-source | **5222**

defaults | **5224**

defaults (ACX series, JRR200, MX204, PTX5000, and VRR) | **5225**

defaults (early-aging) | **5226**

defaults (Policy Options) | **5228**

defer-initial-multipath-build | **5229**

definition | **5231**

definition (Protocols OSPF) | **5234**

degraded | **5236**

delay | **5238**

delay-advertise (DHCPv6) | **5239**

delay-authentication (DHCP Relay Agent) | **5242**

delay-buffer-rate | **5244**

delay-buffer-rate (Dynamic Traffic Shaping) | **5246**

delay-factor | **5248**

delay-measurement (Protocols IS-IS) | **5249**

delay-measurement (Protocols OSPF) | **5251**

delay-metric | **5252**

delay-metric (Protocols OSPF) | **5253**

delay-offer (DHCPv4) | **5254**

delay (PPPoE Service Name Tables) | **5257**

delay (querier) | **5259**

delay (responder) | **5260**

delay-route-advertisements | **5262**

delay-variation | **5265**

delegate-probes | **5266**

delegate-processing (VRRP) | **5268**

delegate-server-processing | **5270**

delegated-pool (DHCP Local Server) | **5271**

delegation-cleanup-timeout | **5272**

delegation-priority | **5274**

delete | **5276**

delete-after-commit (JNU Satellites) | **5278**

delete-binding-on-renegotiation (DHCP Local Server and Relay Agent) | **5279**

delete chassis fpc | **5281**

delete-req | **5282**

delete services jinsightd subscribe health-monitor | **5284**

delimiter | **5285**

delimiter (Access Profile) | **5286**

delimiter (DHCP Local Server) | **5288**

delimiter (DHCP Relay Agent) | **5290**

delimiter (Domain Map) | **5293**

demand-circuit (Protocols RIP) | **5294**

demux-destination (Demux Interface) | **5296**

demux-destination (Underlying Interface) | **5297**

demux (Interfaces) | **5299**

demux-options (All Demux Interfaces) | **5301**

demux-options (Dynamic Interface) | **5302**

demux-options (Static Interface) | **5304**

demux-source (Demux Interface) | **5305**

demux-source (Dynamic IP Demux Interface) | **5306**

demux-source (Dynamic Underlying Interface) | **5308**

demux-source (Underlying Interface) | **5309**

demux0 (Dynamic Interface) | **5311**

demux0 (Static Interface) | **5313**

dense-groups | **5315**

deny (Security Policies) | **5316**

deny (Security SIP) | **5317**

depth-threshold | **5319**

description (RMON) | **5321**

description (SNMP) | **5322**

description (Abstracted Fabric) | **5324**

description (Application Identification) | **5325**

description (Applications) | **5326**

description (Chassis Redundancy) | **5328**

description (Fibre Channel Fabrics) | **5329**

description (Forwarding Options) | **5330**

description (GNF) | **5332**

description (IKE policy) | **5333**

description (Interfaces) | **5334**

description (Junos Fusion) | **5336**

description (Junos Fusion Satellite Device Cluster) | **5337**

description (Op Scripts) | **5339**

description (Protocols BGP) | **5340**

description (Protocols Layer 2 Circuit Neighbor) | **5342**

description (Protocols Layer 2 VPN) | **5343**

description (Protocols MPLS) | **5345**

description (Routing Instances) | **5346**

description (Security Address Book) | **5347**

description (Security IDP Policy) | **5349**

description (Security NAT Pool) | **5350**

description (Security NAT Rule) | **5351**

description (Security NAT Rule Set) | **5353**

description (Security Policies Next Gen Services) | **5354**

description (Security Screen) | **5355**

description (Security Zone) | **5357**

description (Services IPsec VPN) | **5358**

description (VLAN) | **5360**

deselect-on-bandwidth-failure | **5361**

designated-forwarder-election-hold-time | **5363**

destination | **5364**

destination (Accounting) | **5366**

destination-address | **5368**

destination-address (IPsec-VPN) | **5369**

destination-address (Firewall) | **5371**

destination-address-excluded | **5372**

destination-address (Flow Monitoring Logs for NAT) | **5373**

destination-address (HTTP Header Enrichment) | **5375**

destination-address (IDS MS-DPC) | **5377**

destination-address-name (NAT Next Gen Services) | **5378**

destination-address-name (Security Destination NAT) | **5379**

destination-address-name (Security Source NAT) | **5380**

destination-address-name (Security Static NAT) | **5382**

destination-address (NAT Next Gen Services) | **5383**

destination-address (PCP) | **5384**

destination-address-range | **5386**

destination-address-range (HTTP Header Enrichment) | **5387**

destination-address-range (IDS MS-DPC)) | **5389**

destination-address-range (PCP) | **5390**

destination-address-range (Services NAT) | **5391**

destination-address-range (Services Stateful Firewall) | **5393**

destination-address (Security Destination NAT) | **5394**

destination-address (Security IDP Policy) | **5396**

destination-address (Security Policies) | **5397**

destination-address (Security Policies Flag) | **5399**

destination-address (Security Source NAT) | **5400**

destination-address (Security Static NAT) | **5401**

destination-address (Services CoS) | **5403**

destination-address (Services NAT) | **5404**

destination-address (Services Stateful Firewall) | **5405**

destination-address (Subscriber Secure Policy) | **5407**

destination (Application Identification) | **5408**

destination (Chassis) | **5410**

destination-class-usage | **5411**

destination-classes | **5412**

destination (Diameter Network Element) | **5414**

destination-equal-load-balancing (L2TP LAC) | **5415**

destination (Event Policy) | **5416**

destination-except | **5418**

destination-header | **5419**

destination-host | **5421**

destination-host (Gx-Plus) | **5422**

destination-host (OCS Partition) | **5423**

destination-host (PCRF Partition) | **5424**

destination-hostname | **5426**

destination-identity-context | **5427**

destination-identity-context-profile | **5429**

destination-interface | **5430**

destination (Interfaces) | **5433**

destination-ip-address (RADIUS Snoop Segment) | **5434**

destination-ip-based | **5435**

destination-ip (Security Alarms) | **5437**

destination (IPCP) | **5439**

destination-ipv4-address | **5440**

destination-ipv4-address (RFC 2544 Benchmarking Tests) | **5442**

destination-ipv6-address | **5443**

destination (L2TP) | **5445**

destination (LRF Profile) | **5446**

destination-mac-address | **5448**

destination-nat | **5449**

destination-networks | **5451**

destination-option | 5452

destination-override | 5454

destination-path-group | 5455

destination-pool | 5457

destination-port (firewall) | 5458

destination-port (protocols pcep) | 5460

destination-port (RPM and TWAMP) | 5461

destination-port (applications) | 5464

destination-port (Applications) | 5465

destination-port (Flow Monitoring Logs for NAT) | 5471

destination-port (PCP) | 5472

destination-port (RADIUS Snoop Segment) | 5474

destination-port range | 5475

destination-port-range (HTTP Header Enrichment) | 5476

destination-port (Security Destination NAT) | 5478

destination-port (Security Forwarding Options) | 5479

destination-port (Security Signature Attack) | 5485

destination-port (Security Source NAT) | 5486

destination-port (Security Static NAT) | 5487

destination-port (SNMP) | 5488

destination-port (Subscriber Secure Policy) | 5490

destination-ports (HTTP Header Enrichment) | 5491

destination-prefix (Destination NAT Next Gen Services) | 5493

destination-prefix (IDS) | 5494

destination-prefix-ipv6 (IDS) | 5495

destination-prefix-list (HTTP Header Enrichment) | 5497

destination-prefix-list (PCP) | **5498**

destination-prefix-list (Services AACL) | **5500**

destination-prefix-list (Services CoS) | **5501**

destination-prefix-list (Services IDS) | **5502**

destination-prefix-list (Services NAT) | **5504**

destination-prefix-list (Services Stateful Firewall) | **5505**

destination-prefix (Security Forwarding Options) | **5507**

destination-prefix (Services NAT) | **5508**

destination-profile | **5509**

destination (Protocols OSPF) | **5511**

destination-realm (Gx-Plus) | **5512**

destination-realm (JSRC) | **5514**

destination-realm (OCS Partition) | **5515**

destination-realm (PCRF Partition) | **5517**

destination (Routing Instance) | **5518**

destination (saegw system) | **5520**

destination (Security Destination NAT) | **5521**

destination (Security IP Headers Attack) | **5524**

destination-threshold | **5525**

destination (Tunnel Remote End) | **5526**

destination (Tunnels) | **5528**

destination-udp-port | **5529**

destination-udp-port (FTI) | **5531**

destinations | **5532**

destinations (Event Policy) | **5533**

destined-port | **5535**

destruct-timeout (L2TP) | **5536**

detect-loc | **5538**

detection (Network Isolation Groups) | **5540**

detection-time | **5542**

detection-time (BFD for PIM) | **5543**

detection-time (BFD Liveness Detection) | **5545**

detection-time (LAG) | **5549**

detection-time (Liveness Detection) | **5550**

detector | **5551**

deterministic-ar-policy | **5553**

deterministic-nat-configuration-log-interval (Source NAT Next Gen Services) | **5555**

deterministic-port-block-allocation | **5557**

deterministic (Source NAT Next Gen Services) | **5559**

device-attribute (Easy EVPN LAG Configuration) | **5560**

device-authentication | **5563**

device-count | **5565**

device-count (MX Series) | **5567**

device-count (Pseudowire Subscriber Interfaces) | **5569**

devices | **5571**

df-election | **5572**

df-election-granularity | **5574**

df-election-type | **5576**

dh-group (services) | **5578**

dh-group | **5579**

dh-group (Security IKE) | **5581**

dhcp | **5583**

dhcp-attributes (Access IPv4 Address Pools) | **5585**

dhcp-attributes (Access IPv6 Address Pools) | **5587**

dhcp-attributes (Address-Assignment Pools) | **5589**

dhcp-client | **5596**

dhcp (DHCP Client) | **5598**

dhcp-local-server | **5601**

dhcp-local-server (System Services) | **5613**

dhcp-option82 | **5619**

dhcp-relay | **5620**

dhcp-relay (DHCP Spoofing Prevention) | **5637**

dhcp-relay (Easy EVPN LAG Configuration) | **5638**

dhcp-security | **5640**

dhcp-security (MX Series) | **5644**

dhcp-service | **5646**

dhcp-snooping-file (EX) | **5649**

dhcp-snooping-file | **5650**

dhcp-tags (Adjustment Control Profiles) | **5652**

dhcp-trusted | **5654**

dhcpv6-client | **5655**

dhcpv6 (DHCP Local Server) | **5658**

dhcpv6 (DHCP Relay Agent) | **5665**

dhcpv6-options | **5673**

dhcpv6-snooping-file | **5675**

dhcpv6 (System Services) | **5677**

diag-port-authentication | **5682**

dial-options | **5684**

dial-options (Dynamic Profiles) | **5685**

dial-string | **5687**

dialer | **5688**

dialer-options | **5689**

dialin | **5691**

diameter | **5693**

diameter-instance (Diameter Applications) | **5695**

diameter-instance (JSRC) | **5696**

diameter-profile (PCEF Profile) | **5698**

diameter (Subscriber Aware Policy Control) | **5699**

diameter (TDF Gateway) | **5702**

dictionary | **5703**

diffserv-te | **5705**

direct-access (Commit Scripts) | **5706**

direct-connect | **5708**

direction (RFC 2544 Benchmarking Tests) | **5709**

direction (services) | **5711**

direction (Secure-channel) | **5713**

direction (Application Identification) | **5715**

direction (forwarding-class-accounting) | **5716**

direction (Junos-FIPS Software) | **5718**

direction (Junos OS) | **5719**

direction (MEP) | **5721**

direction (MX Series) | **5722**

direction (Security Custom Attack) | **5724**

direction (Security Dynamic Attack Group) | **5726**

direction (Service Data Flow Filters) | **5727**

director-device (Aliases) | **5729**

disable (Protocols) | **5731**

disable | **5736**

disable (Routing Graceful Restart) | **5738**

disable (STP) | **5739**

disable-4byte-as | **5742**

disable-action | **5743**

disable-all-instances | **5744**

disable (Application Tracking) | **5745**

disable (BGP Graceful Restart) | **5747**

disable-calling-number-avp (L2TP LAC) | **5748**

disable (DCBX) | **5749**

disable (Dynamic IGMP) | **5751**

disable (Dynamic MLD) | **5752**

disable (Extensible Subscriber Services Manager) | **5753**

disable-failover-protocol (L2TP) | **5755**

disable (Forwarding Options) | **5756**

disable-fpc (DDoS) | **5758**

disable-global-timeout-override | **5760**

disable (GMPLS) | **5761**

disable-grant-bypass | **5762**

disable (IGMP Snooping) | **5764**

disable (Interface) | **5765**

disable (Link Protection) | **5767**

disable-logging (DDoS) | **5769**

disable (Long-Lived Graceful Restart for BGP Restarter) | **5770**

disable-mlppp-inner-ppp-pfc | **5772**

disable (Multicast Load Balancing) | **5773**

disable (Multicast Snooping) | **5775**

disable (MVRP) | **5776**

disable-natt (Services IPsec VPN) | **5777**

disable-notification-extensions (BGP Graceful Restart) | **5779**

disable-notification-flag (BGP Graceful Restart) | **5780**

disable (PIM) | **5782**

disable (PoE) | **5784**

disable (Power over Ethernet) | **5785**

disable-preceding-key | **5787**

disable (Protocols BGP) | **5788**

disable (Protocols LDP) | **5789**

disable (Protocols MLD) | **5791**

disable (Protocols MLD Snooping) | **5792**

disable (Protocols MPLS) | **5793**

disable (Protocols MSDP) | **5795**

disable (Protocols Router Discovery) | **5797**

disable (Protocols RSVP) | **5798**

disable (Protocols SAP) | **5800**

disable-relay | **5801**

disable-routing-engine (DDoS) | **5802**

disable (Routing Options) | **5804**

disable-signature-check | **5805**

disable-timeout | **5806**

disable-timeout (BPDU) | **5808**

disable-timeout (Port Error Disable) | **5810**

disable-https-filtering | **5812**

disable-packet-register | **5813**

discard (forwarding-options) | **5814**

discard (routing-options) | **5816**

discard-action-for-unresolved-redir-addr | **5818**

discard (Forwarding Class) | **5819**

disconnect-peer-timeout | **5820**

disk-failure-action | **5821**

disk-failure-action | **5823**

disk-partition | **5824**

display-host (Security Fallback Block) | **5826**

display-host (Security Virus Detection) | **5827**

distinguished-name | **5828**

distinguished-name (Access) | **5830**

distinguished-name (Security) | **5831**

distributed-dr | **5833**

distributed (IGMP) | **5834**

distribution-ip (Security SIP) | **5836**

distribution-profile | **5838**

dlb | **5841**

dlci | **5843**

dlv | **5844**

dne-origin (Diameter Network Element) | **5846**

dns | **5847**

dns-alg-pool | **5849**

dns-alg-prefix | **5850**

dns-cache | **5852**

dns-filter | **5853**

dns-filter-template | **5856**

dns-filtering | **5859**

dns-proxy | **5862**

dns-server-address | **5864**

dns-server-address (Dynamic Profiles) | **5865**

dns (System Services) | **5867**

dnssec | **5869**

do-not-fragment | **5871**

dod-request-policy | **5873**

domain (DNS Forwarding Options) | **5874**

domain (Precision Time Protocol) | **5875**

domain (Domain Map) | **5877**

domain-name | **5879**

domain-name (System) | **5880**

domain-name (Access Profile) | **5881**

domain-name (DHCP) | **5882**

domain-name (DHCP Local Server) | **5884**

domain-name (DHCP Relay Agent) | **5886**

domain-name-server-inet (Routing Instances and Access Profiles) | **5888**

domain-name-server-inet6 (Routing Instances and Access Profiles) | **5890**

domain-name-server (Routing Instances and Access Profiles) | **5891**

domain-name (Static Subscribers) | **5892**

domain-path-id | **5894**

domain-path-id (Protocols BGP) | **5898**

domain-search | **5899**

domain-selection | **5901**

domain (TDF Domain Selection) | **5905**

domain-type (Bridge Domains) | **5906**

domains | **5908**

dont-help-shared-fate-bfd-down | **5912**

dot1p-priority | **5913**

dot1q-tunneling | **5915**

dot1x | **5917**

dot1x (MX Series in Enhanced LAN Mode) | **5920**

downlink-dscp-to-egress-forwarding-class | **5922**

downlink-dscp-to-forwarding-class | **5923**

ignore-server-validation (Application Identification) | **5924**

download-profile (Security Antivirus FTP) | **5926**

download-profile (Security Content Filtering FTP) | **5927**

download (Services) | **5928**

download-timeout | **5930**

downstream | **5932**

downstream-mode | **5933**

downstream-on-demand | **5934**

downstream-rate (Traffic Shaping) | **5935**

dr-election-on-p2p | **5937**

dr-register-policy | **5938**

drain | **5940**

draining (Diameter Applications) | **5941**

draining-response-timeout (Diameter Applications) | **5943**

drop (BPDU Block) | **5944**

drop (DHCP Relay Agent Option) | **5946**

drop-flow | **5947**

drop-member-traffic (Aggregated Multiservices) | **5949**

drop-policy (Subscriber Secure Policy) | **5950**

drop (PPPoE Service Name Tables) | **5952**

drop-probability | **5953**

drop-probability (Fill Level) | **5955**

drop-probability (Interpolated Value) | **5956**

drop-probability (Percentage) | **5958**

drop-profile (Schedulers) | **5959**

drop-profile (Dynamic Schedulers) | **5960**

drop-profile-map (Dynamic Schedulers) | **5962**

drop-profile-map (Schedulers) | **5964**

drop-profile-map (Schedulers) | **5965**

drop-profile (Schedulers) | **5966**

drop-profiles (Definition) | **5968**

drop (Security GTP) | **5969**

drop (Security SCTP) | **5974**

drop-threshold | **5978**

drop-threshold (Security GPRS) | **5980**

drop-timeout | **5981**

drop-timeout (Forwarding Class) | **5983**

drop-unknown-mac-ip | **5984**

ds-lite | **5986**

dscp (Class of Service) | **5988**

dscp-code-point (CoS Host Outbound Traffic) | **5991**

dscp-code-points (RFC 2544 Benchmarking) | **5993**

dscp-code-points (RPM) | **5995**

dscp (CoS Classifiers) | **5998**

dscp (CoS Interfaces) | **5999**

dscp (Dynamic Classifiers) | **6000**

dscp (Dynamic Rewrite Rules) | **6002**

dscp (flow-server) | **6003**

dscp (Input Congestion Notification) | **6005**

dscp-ipv6 (Class of Service) | **6007**

dscp-ipv6 (Class of Service) | **6009**

dscp-ipv6 (Dynamic Classifiers) | **6013**

dscp-ipv6 (Dynamic Rewrite Rules) | **6014**

dscp-mapping-classifier | **6016**

dscp (Multifield Classifier) | **6017**

dscp-rewrite | **6019**

dscp (Rewrite Rules on Interfaces) | **6020**

dscp (Services CoS) | **6022**

dscp (Subscriber Secure Policy) | **6023**

dsl (Access-Line Rate Adjustment) | **6025**

dsl-sfp-options | **6030**

dslite-softwire-initiator | **6033**

dsr | **6035**

dsr-polarity | **6037**

dtcp-only (System Services) | **6038**

dte-options | **6039**

dtr | **6041**

dtr-circuit | **6042**

dtr-polarity | **6044**

dual-dr | **6045**

dual-phase-bootup | **6047**

dual-stack (DHCP Local Server Overrides) | **6049**

dual-stack (DHCP Relay Agent Overrides) | **6050**

dual-stack-group (DHCP Local Server) | **6052**

dual-stack-group (DHCP Relay Agent) | **6055**

dual-stack-interface-client-limit (DHCP Local Server and Relay Agent) | **6058**

dual-transport | **6059**

dualstack-support (JSRC) | **6061**

dump-on-flow-control | **6063**

duplicate-clients (DHCPv6 Local Server and Relay Agent) | **6064**

duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent) | **6066**

duplicate-mac-detection | **6068**

duplicate-protection (Dynamic PPPoE) | **6071**

duplicates-dropped-periodicity | **6072**

duplication (Access Profile) | **6074**

duplication-filter (Access Profile) | **6075**

duplication-vrf (Duplicate Accounting) | **6077**

duration | **6078**

duration (PoE) | **6080**

dvmrp | **6081**

dynamic (services) | **6083**

dynamic | **6084**

dynamic-address | **6086**

dynamic-address | **6088**

dynamic-application | **6090**

dynamic-application-group | **6092**

dynamic-application (Security) | **6093**

dynamic-application (Security Policies) | **6095**

dynamic-attack-group | **6097**

dynamic-attack-groups (Security IDP) | **6100**

dynamic-bidirectional-transport | **6101**

dynamic-class-of-service-options (Dynamic Traffic Shaping) | **6102**

dynamic-db | **6105**

dynamic-dns | **6106**

dynamic-filter | **6108**

dynamic-flow-capture | **6109**

dynamic-peers | **6111**

dynamic-policy-control | **6112**

dynamic-pool | **6114**

dynamic-profile (Demux) | **6115**

dynamic-profile (DHCP Local Server) | **6117**

dynamic-profile (DHCP Relay Agent) | **6119**

dynamic-profile (Dynamic Access-Line-Identifier VLANs) | **6121**

dynamic-profile (Dynamic ACI VLANs) | **6122**

dynamic-profile (Dynamic PPPoE) | **6124**

dynamic-profile (L2TP) | **6126**

dynamic-profile-options | **6127**

dynamic-profile (PPP) | **6128**

dynamic-profile (PPPoE Service Name Tables) | **6129**

dynamic-profile (Stacked VLAN) | **6131**

dynamic-profile (VLAN) | **6133**

dynamic-profiles | **6134**

dynamic-requests | **6148**

dynamic-requests-secret (RADIUS Server) | **6149**

dynamic (Security) | **6151**

dynamic-server | **6154**

dynamic (Source Packet Routing) | **6156**

dynamic-tunnel-gre-key | **6158**

dynamic-tunnel-reassembly | **6159**

dynamic-tunnels | **6161**

dynamic-vpn | **6163**

E | 6166

enhanced-scale (IRB Interface) | **6175**

e1-options (BITS Interfaces Signal Type) | **6176**

e1-options (Clock Synchronization) | **6178**

e2e-transparent | **6179**

eapol-address (MACSec) | **6181**

eapol-block | **6183**

early-ageout | **6185**

east-interface | **6186**

echo-req | **6188**

ecmp | **6190**

ecmp-alb | **6191**

ecmp-fast-reroute | **6193**

ecmp-local-bias | **6195**

ecmp-resilient-hash | **6196**

edge | **6197**

edge-virtual-bridging | **6200**

edit | **6201**

effective-aigp | **6203**

effective-shaping-rate | **6204**

egress | **6205**

egress (Analyzer) | **6207**

egress (Buffer Configuration) | **6208**

egress-chaining | **6210**

egress-l2-extended-match | **6212**

egress-policer-overhead | **6213**

egress-policy | **6215**

egress-protection (BGP) | **6216**

egress-protection (Layer 2 circuit) | **6218**

egress-protection (MPLS) | **6220**

egress-protection (Routing Instances) | **6221**

egress-shaping-overhead | **6223**

egress-stats (Flat-File Accounting Options) | **6225**

egress-te | **6227**

egress-te-adj-segment | **6229**

egress-te-backup-paths | **6232**

egress-te-node-segment | **6234**

egress-te-set-segment | **6236**

egress-quantization | **6238**

ei-mapping-timeout | **6241**

ei-mapping-timeout (Source NAT Next Gen Services) | **6242**

eif-flow-limit | **6244**

em-hw-profile | **6245**

email-notify | **6247**

embedded-rp | **6248**

enable | **6249**

enable | **6251**

enable-all-ifl | **6252**

enable-all-qmodules | **6253**

enable-asymmetric-traffic-processing (Service Set Next Gen Services) | **6254**

enable-call-distribution | **6255**

enable-descriptive-session-syslog | **6256**

enable (Enhanced Subscriber Management) | **6257**

enable-explorer (REST API) | **6259**

enable-extended-ql-tlv | **6260**

enable-flow-tracing (Services) | **6263**

enable-gtpu-distribution | **6264**

enable-ipv6-services-for-lac (L2TP) | **6265**

enable-mac-move-action | **6267**

enable-packet-pool | **6268**

enable-performance-mode | **6269**

enable-performance-mode (Application Identification) | **6271**

enable-rejoin (Aggregated Multiservices) | **6272**

enable-reroute-uniform-link-check | **6273**

enable-reverse-reroute | **6275**

enable-session-cache | **6276**

enable-snmp-tunnel-statistics (L2TP) | **6277**

enable-subscriber-analysis (Services Options VMS Interfaces) | **6279**

encap-type | **6280**

encapsulate-inner-vlan | **6282**

encapsulation (protocols evpn) | **6284**

encapsulation (logical-link-layer) | **6286**

encapsulation (interfaces-fti) | **6288**

encapsulation (interfaces) | **6289**

encapsulation (Dynamic Interfaces) | **6297**

encapsulation (Logical Interface) | **6300**

encapsulation-overhead | **6305**

encapsulation (Physical Interface) | **6307**

encapsulation-type (Layer 2 Circuits) | **6308**

encapsulation-type (Layer 2 VPNs) | **6310**

encapsulation-udp-fti | **6313**

encoding (QFX Series) | **6314**

encoding (security) | **6315**

encoding | **6317**

encoding-type | **6318**

encrypt (HTTP Header Enrichment) | **6320**

encryption | **6322**

encryption-algorithm | **6324**

encryption-algorithm (Security) | **6326**

encryption-algorithm (Security IKE) | **6327**

encryption-failures | **6330**

encryption (Junos-FIPS Software) | **6331**

encryption (Junos OS) | **6332**

encryption (MACsec) | **6334**

encryption (MACsec for MX Series) | **6336**

end-interface | **6338**

end-user-profile | **6339**

endpoint-registration-timeout | **6342**

enforce-first-as | **6343**

enforce-strict-scale-limit-license (Subscriber Management) | **6345**

engine-id | **6346**

engine-id (Forwarding Options) | **6348**

engine-not-ready | **6350**

engine-not-ready (Security Antivirus Sophos Engine) | **6351**

engine-type | **6352**

enhanced | **6354**

enhanced-cfm-mode | **6355**

enhanced-convergence | **6357**

enhanced (forwarding-class-accounting) | **6358**

enhanced-hash-key | **6362**

enhanced-hash-key (EX) | **6378**

enhanced-hash-key profile1/profile2 | **6389**

enhanced-hierarchical-policer | **6394**

enhanced-mode | **6396**

enhanced-mode (Network Services) | **6398**

enhanced-mode-override | **6401**

enhanced-policer | **6403**

enhanced-priority-mode | **6404**

enhanced-transmission-selection | **6406**

enrollment | **6408**

enrollment-retry | **6409**

enrollment-url | **6411**

enterprise-oid | **6412**

entropy-label (protocols bgp) | **6413**

entropy-label (protocols mpls) | **6416**

environment-monitoring-policy (satellite-management) | **6417**

environment-monitoring-policy (satellite-policies) | **6419**

epacl-firewall-optimization | **6421**

ephemeral | **6423**

equals | **6426**

equals (DHCP Relay Agent) | **6428**

equals (Dynamic Profile) | **6433**

equals (Event Policy) | **6434**

eracl-ip6-match (packet-forwarding-options) | **6435**

eracl-profile (packet-forwarding-options) | **6437**

erroneous-ccm | **6439**

error | **6441**

error (Security) | **6446**

error-id | **6449**

es-options | **6451**

esi | **6453**

esi-resolution-per-bridge-domain | **6456**

esmc-transmit (Chassis Synchronization) | **6458**

establish-tunnels | **6460**

ether-options | **6461**

ether-pseudowire | **6471**

ether-type (Dynamic Load Balancing) | **6473**

ether-type (VLAN) | **6474**

ether-type | **6476**

ether-type (Chassis FPC) | **6477**

ethernet | **6480**

ethernet (Aggregated Devices) | **6481**

ethernet (Alarm) | **6482**

ethernet (Chassis) | **6484**

ethernet (CoS for Multidestination Traffic) | **6485**

ethernet-frame-size | **6486**

ethernet-policer-profile | **6488**

ethernet (Protocols OAM) | **6490**

ethernet-ring | **6497**

ethernet-switch-profile | **6499**

ethernet-switching | **6501**

ethernet-switching | **6503**

ethernet-switching-options | **6506**

ethernet-switching (Port Mirroring) | **6513**

ethernet-vlan (Protocols Link Management) | **6514**

eui-64 | **6515**

evc-protocol cfm | **6517**

evcs | **6518**

event | **6520**

event (CFM) | **6521**

event (Enhanced Subscriber Management) | **6523**

event (LFM) | **6525**

event-options | **6527**

event-rate | **6531**

event-rate (Next Gen Services Service-Set Local System Logging) | **6532**

event-script (Event Options) | **6533**

event-script (Event Policy) | **6536**

event-thresholds | **6537**

events (Event Policy Correlated Events) | **6538**

events (Event Policy Trigger Events) | **6540**

evpn | **6541**

evpn-aliasing-optimize | **6547**

evpn (Easy EVPN LAG Configuration) | **6548**

evpn-ssm-reports-only | **6553**

evpn-vxlan | **6555**

evpn-vxlan (Easy EVPN LAG Configuration) | **6559**

examine-dhcp | **6561**

examine-dhcpv6 | **6563**

examine-fip | **6565**

examine-vn2vf | **6566**

examine-vn2vn | **6568**

exceed-action | **6570**

exception | **6571**

exception-reporting | **6572**

exception (Security Content Filtering) | **6573**

exceptions | **6575**

excess-bandwidth-share | **6577**

excess-burst-size (QFX) | **6579**

excess-burst-size | **6580**

excess-priority | **6582**

excess-priority (Dynamic Schedulers) | **6584**

excess-rate (Schedulers) | **6585**

excess-rate (Dynamic Schedulers) | **6588**

excess-rate (Dynamic Traffic Shaping) | **6589**

excess-rate-high | **6591**

excess-rate-high (Dynamic Traffic Shaping) | **6593**

excess-rate-low | **6594**

excess-rate-low (Dynamic Traffic Shaping) | **6596**

excess-rate-medium-high | **6597**

excess-rate-medium-low | **6599**

excess-rate (Schedulers) | **6601**

exclude (Diameter Gx Profiles) | **6602**

exclude (Dynamic MLD Interface) | **6604**

exclude (for Administrative Groups) | **6605**

exclude (for Fast Reroute) | **6606**

exclude-hostname | **6608**

exclude (JSRC Attributes) | **6609**

exclude-protocol | **6610**

exclude-protocol (MX Series) | **6612**

exclude (Protocols IGMP) | **6614**

exclude (Protocols MLD) | **6615**

exclude-queue-overhead-bytes | **6616**

exclude (RADIUS Attributes) | **6617**

exclude (Schedulers) | **6626**

exclude (Security Log) | **6627**

exclude-srlg | **6629**

excluded-address (Address-Assignment Pools) | **6631**

excluded-range (Address-Assignment Pools) | **6632**

exclusive-mac | **6634**

execute-commands (Event Policy) | **6636**

exercise | **6637**

exit | **6638**

exp (Class of Service) | **6639**

exp-push-push-push | **6642**

exp (Rewrite rules) | **6643**

exp-swap-push-push | **6646**

expand-loose-hop | **6647**

explicit-congestion-notification | **6649**

explicit-null-cos | **6650**

explicit-null (Protocols BGP) | **6652**

explicit-null (Protocols LDP) | **6654**

explicit-null (Protocols MPLS) | **6655**

explicit-priority | **6657**

export (ISIS) | **6658**

export (OSPF) | **6660**

export (RIP) | **6661**

export (Bootstrap) | **6663**

export-format (JSON) | **6665**

export-format (Flow Monitoring) | **6667**

export (MPLS Traffic engineering database) | **6668**

export-profile (Junos Telemetry Interface) | **6669**

export-profiles | **6674**

export (Protocols BGP) | **6677**

export (Protocols DVMRP) | **6678**

export (Protocols LDP) | **6680**

export (Protocols MSDP) | **6681**

export (Protocols PIM) | **6683**

export (Protocols RIFT) | **6685**

export (Protocols RIPng) | **6686**

export-rib | **6688**

export (Routing Options) | **6691**

export-target | **6693**

expression | **6694**

extend-secondary-vlan-id | **6696**

extended-echo | **6697**

extended-isid-list | **6698**

extended-nexthop | **6700**

extended-nexthop-tunnel | **6701**

extended-space | **6703**

extended-statistics | **6704**

extended-system-id | **6706**

extended-vlan-list | **6707**

extended-vni-all | **6709**

extended-vni-list | **6710**

Extensible Subscriber Services | **6712**

extension-header | **6714**

extension-provider | **6716**

extension-service (System Extensions) | **6717**

extension-service (System Services) | **6719**

extensions | **6722**

external-assigned (Address Pools) | **6725**

external-authority | **6726**

F | 6727

f-max-period | **6742**

f-max-period (rtp) | **6744**

fabric | **6745**

fabric (Class-of-Service) | **6747**

fabric-control | **6749**

fabric-id | **6750**

fabric-interfaces | **6751**

fabric-load-balance | **6753**

fabric loopback wan | **6755**

fabric-maintenance-associations | **6756**

fabric-maintenance-end-points | **6759**

fabric (OAM) | **6761**

fabric-options | **6763**

fabric-tree-root | **6766**

fabric-type | **6767**

facility | **6769**

facility-override | **6770**

facility-override (Security) | **6771**

facility-override (Service Sets) | **6772**

facility-override (System Log Reporting) | **6774**

fail | **6775**

fail-filter (Dynamic Profiles) | **6777**

failover (Chassis) | **6778**

failover-delay | **6780**

failover-delay | **6781**

failover-resync | **6783**

failover (System Process) | **6784**

failover-within-preference (L2TP LAC) | **6786**

failure-action | **6787**

failure-action (Protocols LDP) | **6789**

failure-action (Protocols MPLS) | **6790**

fallback-block (Security Antivirus) | **6792**

fallback-key | **6794**

fallback-non-block (Security Antivirus) | **6795**

fallback-option (ICAP Redirect Service) | **6797**

fallback-options (Security Antivirus Juniper Express Engine) | **6799**

fallback-options (Security Antivirus Kaspersky Lab Engine) | **6800**

fallback-options (Security Antivirus Sophos Engine) | **6802**

fallback-settings (Security Web Filtering) | **6804**

fallback-settings (Security Web Filtering Juniper Local) | **6805**

fallback-settings (Security Web Filtering Websense Redirect) | **6806**

falling-event-index | **6808**

falling-event-index (RMON) | **6809**

falling-threshold (Health Monitor) | **6811**

falling-threshold (integer) | **6812**

falling-threshold-interval | **6814**

falling-threshold (percentage) | **6815**

falling-threshold (RMON) | **6817**

false-positives | **6818**

family (RFC 2544 Benchmarking Tests) | **6820**

family (logical-interface) | **6822**

family (ELS) | **6824**

family (QFX) | **6829**

family (protocols ldp) | **6831**

family (interfaces) | **6832**

family (Link Interfaces) | **6839**

family (Address-Assignment Pools) | **6841**

family (Address Pools) | **6843**

family (Aggregated Multiservices) | **6844**

family (Bootstrap) | **6846**

family bridge | **6847**

family bridge (GRE Interfaces) | **6849**

family (chassis) | **6851**

family (Dynamic Demux Interface) | **6853**

family (Dynamic Firewalls) | **6855**

family (Dynamic PPPoE) | **6857**

family (Dynamic Standard Interface) | **6859**

family (Exclude Prefix) | **6862**

family fcoe | **6863**

family (Filtering) | **6865**

family (Firewall) | **6867**

family (Firewall Filter) | **6869**

family inet | **6871**

family inet (Interfaces) | **6874**

family inet6 | **6878**

family (Interfaces) | **6881**

family (Local RP) | **6883**

family (Monitoring) | **6885**

family mpls | **6886**

family (Multi-destination) | **6890**

family (Multifield Classifier) | **6891**

family multiservice | **6893**

family (Port Mirroring) | **6896**

family (Protocols AMT Relay) | **6898**

family (Protocols BGP) | **6900**

family (Protocols IS-IS) | **6905**

family (Protocols PIM) | **6907**

family (Protocols PIM Interface) | **6909**

family route-target | **6911**

family (Sampling) | **6913**

family-state-change-immediate-update | **6916**

family (TDF Interface) | **6917**

family (Voice Services) | **6918**

family vpls (Layer 2 Pseudowires) | **6920**

family (VRF Advertisement) | **6921**

fast-aps-switch | **6922**

fast-failover | **6924**

fast-interval | **6926**

fast-lookup-filter | **6928**

fast-lookup-filter (PTX) | **6930**

fast-lookup-tuple | **6935**

fast-lookup-tuple-list | **6937**

fast-lookup-tuple-list (policy-options) | **6938**

fast-poe | **6939**

fast-reroute-priority | **6941**

fast-reroute (Protocols MPLS) | **6943**

fast-reroute (Protocols RSVP) | **6944**

fast-update-filter (Dynamic Firewalls) | **6946**

fast-lookup-filter (Protocols BGP) | **6947**

fastether-options | **6949**

fatal | **6950**

fate-sharing | **6952**

fc-fabrics | **6954**

fc-map | **6957**

fc-options | **6959**

fc2 | **6960**

fcoe | **6962**

fcoe-lag | **6963**

fcoe-trusted | **6965**

feature-profile | **6967**

features (VNF) | **6978**

features | **6979**

feb (Assigning a FEB to a Redundancy Group) | **6981**

feb (chassis) | **6982**

feb (Creating a Redundancy Group) | **6984**

fec | **6985**

fec (MX Series) | **6987**

fec (protocols ldp) | **6990**

fec (ether) | **6992**

fec (gigether) | **6995**

feed-name | **6997**

feed-server | **6999**

feeds (T640, T1600, and T4000 Routers with Six-Input DC Power Supply) | **7001**

fib-install-priority | **7003**

fib-local | **7005**

fib-next-hop-split | **7006**

fib-prioritization | **7007**

fib-remote | **7009**

fibre-channel (Alarm) | **7010**

fibre-channel (Family Interfaces) | **7012**

fibre-channel (Port) | **7013**

fibrechannel-options | **7015**

fields (Flat-File Accounting Options) | **7016**

fields (for Interface Profiles) | **7019**

fields (for Routing Engine Profiles) | **7021**

file | **7023**

file | **7025**

file (Security Certificates) | **7027**

file (Analytics) | **7028**

file (App Engine Virtual Machine Management Service) | **7030**

file archive | **7032**

file (Associating with a Profile) | **7036**

file checksum md5 | **7037**

file checksum sha-256 | **7039**

file checksum sha1 | **7041**

file (Commit Scripts) | **7042**

file compare | **7044**

file (Configuring a Log File) | **7047**

file copy | **7049**

file delete | **7052**

file (Event Scripts) | **7054**

file (Extended DHCP Relay Agent and Helpers Trace Options) | **7056**

file (Flat-File Accounting Options) | **7058**

file (JET) | **7059**

file list | **7062**

file (Next Gen Services Global System Logging) | **7065**

file (Op Scripts) | **7067**

file (Packet Forwarding Engine Traffic Traceoptions) | **7069**

file (QFabric System) | **7071**

file rename | **7072**

file (Sampling) | **7075**

file (Security Log) | **7076**

file show | **7078**

file (SNMP Scripts) | **7080**

file-specification (File Format) | **7082**

file-specification (Interface Mapping) | **7083**

file (System Logging) | **7085**

file (System Logging) | **7086**

file (Trace Options) | **7090**

file (Tracing for Origin AS Validation) | **7091**

file-type | **7093**

file-serialization | **7094**

filename | **7096**

filename-extension | **7097**

filename (Next Gen Services Global System Logging) | **7099**

filename-prefix | **7100**

filename (Sampling) | **7101**

files | **7103**

files (accounting) | **7104**

files (Next Gen Services Global System Logging) | **7105**

files (Sampling and Traceoptions) | **7107**

files (syslog) | **7108**

fill-level | **7110**

fill-level (Drop Profiles) | **7112**

fill-level (Interpolated Value) | **7113**

filter (Flow Monitoring) | **7115**

filter (interfaces) | **7116**

filter (Interfaces) | **7117**

filter | **7119**

filter (routing-options) | **7122**

filter (Applying to a Logical Interface) | **7123**

filter (Applying to an Interface) | **7125**

filter (Configuring) | **7127**

filter-duplicates | **7129**

filter (Dynamic Interface Unit) | **7130**

filter (Dynamic Profiles Filter Attachment) | **7132**

filter (Dynamic Profiles Filter Creation) | **7134**

filter (export routing policy action) | **7135**

filter (Firewall Filters) | **7136**

filter (Identity Management Advanced Query) | **7138**

filter-interfaces | **7142**

filter (IPv4, IPv6, and MPLS) | **7144**

filter (Layer 2 and Layer 3 Interfaces) | **7145**

filter (Layer 3 Interfaces) | **7146**

filter-list-template | **7148**

filter (memory-enhanced) | **7150**

filter-profile | **7151**

filter (Security) | **7153**

filter-specific (QFX) | **7154**

filter-specific | **7155**

filter (VLANs) | **7157**

filter (VPLS) | **7158**

filtering-type | **7160**

filtering-type (Source NAT Next Gen Services) | **7161**

filters | **7162**

fin-invalidate-session | **7165**

fin-no-ack | **7167**

fin-no-ack (IDS Screen Next Gen Services) | **7168**

final-response-timeout (OCS Partition) | **7170**

finger | **7171**

fip | **7173**

fip-security | **7175**

firewall | **7176**

firewall-authentication | **7179**

firewall-authentication-forced-timeout | **7181**

firewall-authentication (Security) | **7184**

firewall-authentication (Security Policies) | **7186**

firewall-authentication-service | **7188**

firewall-authentication (tenants) | **7189**

firewall (Dynamic Firewalls) | **7191**

firewall-install-disable | **7194**

firewall-user | **7196**

firmware-revision | **7197**

fiveqi-map | **7199**

fixed-stuff-bytes | **7201**

fka-adv-period | **7202**

flag (App Engine Virtual Machine Management Service) | **7204**

flag (Interfaces) | **7206**

flag (Next Gen Services Global System Logging) | **7208**

flag (Services) | **7209**

flag (SMTP) | **7211**

flag (Tracing for Origin AS Validation) | **7212**

flap-trap-monitor | **7215**

flat-file-profile (Accounting Options) | **7217**

flat-file-profile (Extensible Subscriber Services) | **7220**

flex-algorithm | **7221**

flex-algorithm (Protocols OSPF) | **7223**

flex-flow-sizing | **7225**

flex-hashing | **7227**

flexible-cross-connect-vlan-aware | **7230**

flexible-cross-connect-vlan-unaware | **7232**

flexible-queuing-mode | **7233**

flexible-vlan-tagging | **7235**

flexible-vlan-tagging (Interfaces) | **7237**

flood | **7238**

flood-as-unknown-unicast | **7239**

flood-groups | **7241**

flood-reflector (Protocols IS-IS) | **7242**

flood (Security ICMP) | **7245**

flood (Security UDP) | **7246**

flood (VLANs) | **7248**

flow (protocols bgp) | **7250**

flow | **7252**

flow-action | **7259**

flow-active-timeout (Flow Monitoring) | **7261**

flow-collector (Flow Collection) | **7263**

flow-control (ACX Series) | **7265**

flow-control | **7267**

flow-control (Interfaces) | **7269**

flow-control-options | **7270**

flow-control-queue (Output Congestion Notification) | **7272**

flow-descriptions | **7274**

flow-detect-time (DDoS Flow Detection) | **7276**

flow-detection (DDoS Flow Detection) | **7278**

flow-detection (DDoS Packet Level) | **7279**

flow-detection-mode (DDoS Flow Detection) | **7281**

flow-detection-mode (DDoS Global Flow Detection) | **7283**

flow-export-destination (Flow Collection) | **7285**

flow-export-rate | **7286**

flow-export-timer | **7288**

flow-gate | **7289**

flow-groups | **7291**

flow-inactive-timeout (Flow Monitoring) | **7296**

flow (IPv6) | **7298**

flow-key (Flow Monitoring) | **7304**

flow-label | **7306**

flow-label-receive | **7307**

flow-label-receive-static | **7309**

flow-label-static | **7311**

flow-label-transmit | **7312**

flow-label-transmit-static | **7314**

flow-level-bandwidth (DDoS Flow Detection) | **7316**

flow-level-control (DDoS Flow Detection) | **7317**

flow-level-control (DDoS Global Flow Detection) | **7319**

flow-level-detection (DDoS Flow Detection) | **7320**

flow-map | **7322**

flow-monitoring | **7323**

flow-monitoring (Inline Monitoring Services) | **7326**

flow-monitoring (Services) | **7328**

flow-recover-time (DDoS Flow Detection) | **7330**

flow-report-rate (DDoS Flow Detection) | **7331**

flow (Security Flow) | **7333**

flow (Security IDP) | **7339**

flow-server | **7344**

flow-session | **7347**

flow-specs | **7349**

flow-table-size | **7351**

flow-table-size (Chassis) | **7354**

flow-tap | **7355**

flow-tap-dtcp | **7357**

flow-timeout-time (DDoS Flow Detection) | **7359**

flow-sample | **7360**

flows (PCC Rules) | **7362**

fnp | **7364**

force-continue (OCS Partition) | **7365**

force-control-packets-on-transit-path | **7367**

force-discover (DHCP Client) | **7369**

force-entry (IDS MS-DPC) | **7370**

force-ip-reassembly | **7371**

force-premium (Firewall Filter Action) | **7373**

force switch | **7374**

force-switch | **7375**

force-up (EX and QFX Series) | **7376**

force-up | **7378**

force-version (IEEE 802.1D STP) | **7379**

force-fbf-terms | **7381**

format (Flat-File Accounting Options) | **7382**

format (LRF Profile) | **7383**

format (Next Gen Services Service-Set Remote System Logging) | **7385**

format (Security Log) | **7386**

format (Security Log Stream) | **7388**

format (Unified Edge Gateways) | **7389**

forward-and-send-to-re | **7392**

forward-delay | **7393**

forward-only | **7395**

forward-only (DHCP Relay Agent) | **7396**

forward-only (DHCP Relay Agent Option) | **7398**

forward-only (DNS) | **7400**

forward-only-replies (DHCP Relay Agent) | **7401**

forward-policy-mismatch (Security Group VPN Member) | **7402**

forward-snooped-clients (DHCP Local Server) | **7404**

forward-snooped-clients (DHCP Relay Agent) | **7405**

forwarding-cache (Bridge Domains) | **7407**

forwarding-cache (Flow Maps) | **7408**

forwarding-class | **7410**

forwarding-class-accounting | **7413**

forwarding-class (BA Classifiers) | **7415**

forwarding-class (CoS Host Outbound Traffic) | **7417**

forwarding-class-default (Forwarding Policy) | **7418**

forwarding-class (Dynamic Scheduler Maps) | **7420**

forwarding-class (Fibre Channel Interfaces) | **7421**

forwarding-class (Firewall Filter Action) | **7423**

forwarding-class (for DHCP Snooping or DAI Packets) | **7425**

forwarding-class (Forwarding Policy) | **7426**

forwarding-class (Fragmentation) | **7428**

forwarding-class (Gigabit Ethernet IQ Classifier) | **7430**

forwarding-class (Interfaces) | **7432**

forwarding-class (Multifield Classifiers) | **7433**

forwarding-class (PCC Action Profiles) | **7434**

forwarding-class (Restricted Queues) | **7436**

forwarding-class (RFC 2544 Benchmarking) | **7437**

forwarding-class (Sampling) | **7439**

forwarding-class (Services CoS Fragmentation Properties) | **7440**

forwarding-class (Services PIC Classifiers) | **7442**

forwarding-class-set | **7443**

forwarding-class-set (Interconnect Device) | **7444**

forwarding-class-sets | **7446**

forwarding-classes | **7447**

forwarding-context (Protocols BGP) | **7453**

forwarding-db-size | **7454**

forwarding (Diameter Network Element) | **7456**

forwarding-instance identifier | **7457**

forwarding-mode (100-Gigabit Ethernet) | **7459**

forwarding-mode (PTX Series Packet Transport Routers) | **7460**

forwarding-mode (Security Content Security Policy) | **7461**

forwarding-options (Routing Instance) | **7463**

forwarding-options (Bridge Domains or VLANs) | **7467**

forwarding-options (Security) | **7473**

forwarding-packages | **7478**

forwarding-policy | **7479**

forwarding-process | **7481**

forwarding-rib | **7483**

forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter) | **7485**

forwarding-state-bit (Per Family for BGP Graceful Restart) | **7488**

forwarding-table | **7490**

forwarding-options (Unified Forwarding Table Profile) | **7493**

fpc (QFX Series) | **7499**

fpc (Chassis) | **7501**

fpc (DDoS) | **7505**

fpc error | **7507**

fpc-feb-connectivity | **7509**

fpc (Interfaces) | **7511**

fpc (Junos Fusion) | **7512**

fpc (Junos Fusion Satellite Device Cluster) | **7514**

fpc (M320, T320 and T640 Devices) | **7516**

fpc (MX Series 5G Universal Routing Platforms) | **7519**

fpc-nmi-volt-fail-knob | **7521**

fpc (Notification Control) | **7523**

fpc-offline-on-blackholing | **7524**

fpc (PTX Series Devices) | **7525**

fpc-restart | **7528**

fpc-resync | **7529**

fpc-slice | **7530**

fpc (Subscriber Limits) | **7533**

fpc (TX Matrix and TX Matrix Plus Routers) | **7534**

fpcs (Node Slicing) | **7536**

fpcs (NSSU Upgrade Groups) | **7538**

fragment (IDS Screen Next Gen Services) | **7540**

fragment-limit | **7541**

fragment-threshold | **7542**

fragment-threshold (Forwarding Class Maps) | **7544**

fragment-threshold (Interfaces LSQ) | **7545**

fragmentation-map | **7547**

fragmentation-maps | **7548**

frame-error | **7550**

frame-mode (Dynamic Traffic Shaping) | **7552**

frame-period | **7554**

frame-period-summary | **7555**

frame-relay-de (CoS Interfaces) | **7557**

frame-relay-de (CoS Loss Priority) | **7558**

frame-relay-de (CoS Rewrite Rule) | **7560**

frame-relay-de (Defining Loss Priority Maps) | **7561**

frame-relay-de (Defining Loss Priority Rewrites) | **7563**

framed-ip-address | **7564**

framed-ipv6-prefix | **7566**

framed-pool | **7567**

framing (10-Gigabit Ethernet Interfaces) | **7568**

framing (E1 Options) | **7571**

framing (E1 Options for BITS Interfaces) | **7572**

framing (T1 Options) | **7573**

framing (T1 Options for BITS Interfaces) | **7575**

free-fw-memory-watermark (Resource Monitor) | **7576**

free-heap-memory-watermark (Resource Monitor) | **7578**

free-nh-memory-watermark (Resource Monitor) | **7579**

frequency-lock-threshold | **7580**

from (Application Aware Access List) | **7582**

from (services-ipsec-vpn) | **7584**

from (Firewall) | **7585**

from (Captive Portal Content Delivery Tags) | **7587**

from (HTTP Header Enrichment) | **7588**

from (IDS MS-DPC) | **7590**

from (PCC Rules) | **7591**

from (PCP) | **7593**

from (Protocols MPLS) | **7595**

from (Security NAT) | **7596**

from (Services CoS) | **7597**

from (Services NAT) | **7599**

from (Services Stateful Firewall) | **7600**

from (Subscriber Secure Policy) | **7602**

from (TDF Domain Selection) | **7603**

from-zone (Security IDP Policy) | **7607**

from-zone (Security Policies) | **7608**

from-zone (Security Policies Global) | **7612**

fru-poweron-sequence | **7613**

fte (Port) | **7615**

ftp | **7617**

ftp (Access) | **7619**

ftp (Content Security Policy Anti-Virus) | **7621**

ftp (Content Security Policy Content Filtering) | **7622**

ftp (Flow Collector Files) | **7624**

ftp (Security ALG) | **7626**

ftp (Services CoS) | **7628**

ftp (Services CoS Next Gen Services) | **7629**

ftp (Transfer Log Files) | **7631**

full | **7632**

function (Diameter Network Element for Subscriber Management) | **7633**

function (Diameter Network Element) | **7634**

function (Diameter Route) | **7636**

functional-zone | **7638**

G | 7640

g-duplicates-dropped-periodicity | **7645**

g-max-duplicates | **7646**

garp-na-enable | **7648**

gate-status | **7649**

gate-timeout | **7651**

gatekeeper | **7653**

gateway-name (LNS Local Gateway) | **7654**

gateway-name (Tunnel Profile Remote Gateway) | **7655**

gateway-name (Tunnel Profile Source Gateway) | **7657**

gateway (Security Group VPN Member IKE) | **7658**

gateway (Security Group VPN Server IKE) | **7660**

gateway (Security IKE) | **7663**

general-ikeid | **7669**

general-param (Flat-File Accounting Options) | **7670**

generate | **7672**

generate-event | **7674**

generate-snmp-traps | **7676**

geo-ip (global-filter) | **7678**

geo-ip (group-filter) | **7679**

get-route-range | **7681**

ggsn-address (OCS Partition) | **7682**

ggsn-mcc-mnc (OCS Partition) | **7684**

gigether-options | **7685**

gigether-options (Chassis Cluster) | **7694**

global-advertisements-threshold | **7696**

global-arp-prefix-limit (Host Fast Reroute) | **7698**

global-config (Services) | **7700**

global (DDoS) | **7703**

global-dns-stats-log-timer | **7705**

global (Gx-Plus) | **7706**

global-limit | **7708**

global-mac-ip-limit | **7709**

global-mac-ip-table-aging-time | **7711**

global-mac-limit | **7712**

global-mac-limit (Protocols) | **7714**

global-mac-move | **7716**

global-mac-statistics | **7717**

global-mac-table-aging-time | **7719**

global-mode (Protocols) | **7721**

global-no-control-mac-aging | **7722**

global-no-mac-learning | **7723**

global (OCS) | **7725**

global-options (Juniper Secure Connect) | **7726**

global-parameters (Easy EVPN LAG Configuration) | **7728**

global (PCRF) | **7731**

global (Security IDP) | **7733**

global (Security Policies) | **7734**

global-supplementary-blackout-timer (Host Fast Reroute) | **7737**

global-threshold | **7740**

global-wait-to-restore | **7741**

global-weight | **7743**

global-load-balancing | **7744**

gnf | **7746**

gpid | **7748**

gprs | **7750**

gprs-gtp-profile | **7755**

gprs-sctp-profile | **7757**

graceful-deletion-timeout | **7758**

graceful-restart (Enabling Globally) | **7760**

graceful-restart (Fabric Control) | **7762**

graceful-restart (Long-Lived for BGP Helper) | **7763**

graceful-restart (Long-Lived for BGP Restarter) | **7765**

graceful-restart (Multicast Snooping) | **7768**

graceful-restart (Protocols BGP) | **7769**

graceful-restart (Protocols IS-IS) | **7771**

graceful-restart (Protocols LDP) | **7773**

graceful-restart (Protocols OSPF) | **7775**

graceful-restart (Protocols PIM) | **7777**

graceful-restart (Protocols RIP) | **7779**

graceful-restart (Protocols RIPng) | **7780**

graceful-restart (Protocols RSVP) | **7782**

graceful-shutdown (Protocols BGP) | **7784**

graceful-switchover (Chassis) | **7786**

gratuitous-arp-count | **7787**

gratuitous-arp-reply | **7789**

gre | **7790**

gre (Ids-option) | **7791**

gre-in | **7793**

gre-out | **7794**

gre-performance-acceleration (Security Flow) | **7796**

gre (Routing Options) | **7797**

greater-than | **7798**

gres-route-flush-delay (Subscriber Management) | **7800**

group | **7801**

group-address (Routing Instances Tunnel Group) | **7802**

group-address (Routing Instances VPN) | **7804**

group (Bridge Domains) | **7806**

group (Configuring Group Name) | **7807**

group-count (Dynamic MLD Interface) | **7809**

group-count (Protocols IGMP) | **7810**

group-count (Protocols MLD) | **7812**

group (Defining Access Privileges for an SNMPv3 Group) | **7813**

group (DHCP Local Server) | **7814**

group (DHCP Relay Agent) | **7820**

group (DHCP Security) | **7826**

group (DHCP Security for MX Series) | **7827**

group (DHCP Spoofing Prevention) | **7829**

group (Distributed IGMP) | **7830**

group (Dynamic IGMP Interface) | **7832**

group (Dynamic MLD Interface) | **7834**

group (IGMP Snooping) | **7835**

group-increment (Dynamic MLD Interface) | **7837**

group-increment (Protocols IGMP) | **7838**

group-increment (Protocols MLD) | **7840**

group-limit (Dynamic IGMP Interface) | **7841**

group-limit (Dynamic MLD Interface) | **7843**

group-limit (IGMP) | **7844**

group-limit (IGMP and MLD Snooping) | **7846**

group-limit (Protocols MLD) | **7847**

group-members | **7849**

group (Origin Validation for BGP) | **7850**

group-policy (Dynamic IGMP Interface) | **7852**

group-policy (Dynamic MLD Interface) | **7854**

group-policy (Protocols IGMP) | **7855**

group-policy (Protocols IGMP AMT Interface) | **7856**

group-policy (Protocols MLD) | **7858**

group-profile (Access) | **7859**

group-profile (Group Profile) | **7861**

group (Protocols BGP) | **7863**

group (Protocols IGMP) | **7869**

group (Protocols MLD) | **7870**

group (Protocols MLD Snooping) | **7872**

group (Protocols MSDP) | **7873**

group (Protocols PIM) | **7876**

group (Protocols RIP) | **7877**

group (Protocols RIPng) | **7880**

group-range (Data MDTs) | **7882**

group-range (MBGP MVPN Tunnel) | **7883**

group-ranges | **7885**

group (Redundant Trunk Groups) | **7887**

group (Routing Instances) | **7889**

group-rp-mapping | **7891**

group (RPF Selection) | **7893**

group (Security Group VPN) | **7894**

group (Static Subscribers) | **7897**

group (System Services DHCP) | **7899**

group-threshold (Protocols IGMP Interface) | **7903**

group-threshold (Protocols MLD Interface) | **7905**

group (Traffic Load Balancer) | **7906**

group-type (Unknown Unicast Forwarding) | **7908**

group (Uplink Failure Detection) | **7910**

group-vpn | **7911**

grouped-ie-profile | **7915**

groups | **7916**

grpc | **7920**

grpc-replay | **7922**

grpc-tunnel | **7923**

gsm-options | **7926**

gsmp-syn-timeout (ANCP) | **7928**

gsmp-syn-wait (ANCP) | **7929**

gtp | **7931**

gtp-header-offset | **7936**

gtp-tunnel-endpoint-identifier | **7938**

guaranteed-rate (Traffic Control Profiles) | **7940**

guaranteed-rate (Dynamic Traffic Shaping) | **7942**

guard-band | **7944**

guard-band (PoE) | **7946**

guard-interval | **7947**

guest-network-functions | **7949**

gw-interface | **7951**

gx-plus (Gx-Plus) | **7952**

gx-profile | **7954**

H | 7956

h323 | **7961**

halt-on-prefix-down (RFC 2544 Benchmarking) | **7963**

handover-default | **7964**

handover-group | **7966**

handshake-timeout | **7967**

hard-limit | **7969**

hard-limit-target | **7970**

hardware-assisted-keepalives | **7971**

hardware-assisted-timestamping | **7973**

hardware-timestamp | **7975**

hardware-upgrade | **7976**

has-prefix (Unified Edge Gateways) | **7977**

has-suffix | **7979**

hash-key (Chassis) | **7980**

hash-key (Chassis LAG) | **7982**

hash-key (Forwarding Options) | **7984**

hash-key (Forwarding Options) | **7986**

hash-keys (Aggregated Multiservices) | **7992**

hash-keys (Interfaces) | **7995**

hash-mode | **7997**

hash-parameters | **8000**

hash-seed | **8002**

hash-seed (profile) | **8003**

hcm (HTTP Header Enrichment) | **8005**

hcm-profile (HTTP Header Enrichment) | **8007**

hcm-profile (PCC Action Profiles) | **8008**

header-integrity-check | **8010**

header-integrity-check (Next Gen Services) | **8012**

header-length | **8014**

header-type | **8015**

health-monitor | **8016**

health-monitor (KHMS) | **8018**

heartbeat-address (MX Series Virtual Chassis) | **8020**

heartbeat-interval | **8022**

heartbeat-threshold | **8024**

heartbeat-timeout (MX Series Virtual Chassis) | **8025**

heartbeat-tos (MX Series Virtual Chassis) | **8027**

hello-acknowledgements | **8029**

hello-dead-interval | **8030**

hello-interval | **8032**

hello-interval (L2TP) | **8033**

hello-interval (LMP) | **8035**

hello-interval (Protocols LDP) | **8036**

hello-interval (Protocols RSVP) | **8038**

hello-message (NETCONF) | **8039**

hello-time | **8041**

hello-timer | **8043**

help | **8045**

helper-disable (LDP) | **8046**

helper-disable (Multiple Protocols) | **8048**

helpers | **8049**

hide-avps | **8053**

hide-via-headers | **8054**

hierarchical-policer | **8055**

hierarchical-policer (aggregate or premium) | **8057**

hierarchical-scheduler | **8060**

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) | **8062**

high-availability (Chassis) | **8064**

high-availability-options (Aggregated Multiservices) | **8068**

high-availability (security cloud) | **8070**

high-availability (Security IDP) | **8073**

high-cos-queue-threshold | **8074**

high-polarization | **8076**

high-power | **8077**

high-threshold (Resource Monitor) | **8078**

high-watermark | **8080**

hint | **8081**

history | **8082**

history-size | **8084**

hold-down | **8086**

hold-down-interval | **8088**

hold-interval | **8089**

hold-interval (Clock Synchronization) | **8091**

hold-interval (OAM) | **8093**

hold-interval (OAM CFM, for EX Series Switch Only) | **8094**

hold-interval (Protection Group) | **8095**

hold off time | **8097**

hold-time (ae interface) | **8098**

hold-time | **8101**

hold-time (OAM) | **8102**

hold-time (Protocols BGP) | **8104**

hold-time (Protocols DVMRP) | **8106**

hold-time (Protocols IS-IS) | **8108**

hold-time (Protocols LDP) | **8110**

hold-time (Protocols MSDP) | **8112**

hold-time (Protocols PIM) | **8114**

hold-time (Redundant Ethernet Interfaces) | **8115**

hold-time up | **8117**

hold-time (VRRP) | **8119**

holddown (candidate-uplink-port-profile) | **8120**

holddown-interval | **8122**

holddown-interval (protocols ldp) | **8124**

holddown (Protocols RIP) | **8125**

holddown (Protocols RIPng) | **8127**

holdover-mode-disable | **8128**

holdover-time-error-budget | **8129**

home-address | **8131**

hop-by-hop-header | **8132**

hop-limit | **8134**

host (Address-Assignment Pools) | **8136**

host-address-base | **8137**

host-address-base (Source NAT Next Gen Services) | **8139**

host (Diameter Origin) | **8140**

host-fast-reroute | **8141**

host-inbound-traffic | **8142**

host (Interfaces) | **8144**

host (L2TP) | **8146**

host-name | **8147**

host-name (System) | **8148**

host-name (DHCP Relay Agent) | **8150**

host-name (DHCPv6 Relay Agent) | **8151**

host (Next Gen Services Service-Set Remote System Logging) | **8152**

host-only-interface | **8154**

host-os forwarding-options analyzer | **8155**

host-outbound media-interface | **8157**

host-outbound-traffic (Class-of-Service) | **8158**

host-outbound-traffic (Class-of-Service) | **8160**

host-outbound-traffic (Multicast Snooping) | **8162**

host-prefix-only | **8163**

host (Security IDP Sensor Configuration) | **8165**

host (Security Log) | **8166**

host (Security Web Filtering) | **8167**

host (service-set) | **8169**

host (System) | **8171**

host-vpn | **8174**

hosted-service-identifier | **8176**

hosted-services | **8178**

hostkey-algorithm | **8179**

hostkey-algorithm-list | **8182**

hot-root-standby (MBGP MVPN) | **8184**

hot-standby (Routing Instance) | **8186**

hot-standby | **8187**

hot-standby (Protocols Layer 2 Circuit) | **8189**

hot-standby-vc-on (Protocols Layer 2 Circuit) | **8190**

http (Access) | **8192**

http (icap-redirect profile) | **8193**

http-log-multiple-transactions (LRF Profile) | **8195**

http-persist | **8196**

http-profile (Security Antivirus) | **8198**

http-profile (Security Content Filtering) | **8199**

http-profile (Security Web Filtering) | **8200**

http-reassemble | **8201**

http (REST API) | **8202**

http (System Services) | **8204**

http (Web Management) | **8205**

https (REST API) | **8207**

https (System Services) | **8208**

https (Web Management) | **8211**

hub-assist | **8213**

hugepages | **8215**

hw-db-profile | **8216**

hybrid | **8220**

hyper-mode (forwarding-options) | **8221**

I | 8223

irb-enhanced-scale (Bridge Domain) | **8243**

icap-redirect | **8245**

iccp (Traceoptions) | **8247**

iccp | **8250**

icl-down-delay | **8252**

icmp | **8253**

icmp-code | **8255**

icmp-code (Applications) | **8256**

icmp (Error Message Rate Limit) | **8257**

icmp-fragment-check (IDS MS-MPC) | **8259**

icmp (IDS Screen Next Gen Services) | **8260**

icmp-large-packet-check (IDS MS-MPC) | **8262**

icmp-mapping (Application Identification) | **8263**

icmp (Security IDP Custom Attack) | **8264**

icmp (Security IDP Signature Attack) | **8265**

icmp (Security Screen) | **8267**

icmp (System) | **8269**

icmp-type | **8270**

icmp-type (Applications) | **8271**

icmp6 (Error Message Rate Limit) | **8273**

icmpv6-malformed | **8274**

icmpv6-malformed (IDS Screen Next Gen Services) | **8276**

icmpv6 (Security IDP) | **8277**

icmpv6 (Security IDP Custom Attack) | **8278**

id | **8279**

id-components | **8281**

id (MACsec for MX Series) | **8283**

id (Protocols OpenFlow) | **8284**

identification (Security ICMP Headers) | **8286**

identification (Tunnel Profile) | **8287**

identifier (VPLS Multihoming for FEC 129) | **8288**

identity | **8290**

identity-context-profile | **8291**

identity-management | **8293**

identity-provider (SAML) | **8298**

idle-after-switch-over | **8300**

idle-standby-path-switchover-delay | **8301**

idle-timeout | **8303**

idle-timeout (Unified Edge) | **8304**

idle-timeout (Access) | **8306**

idle-timeout (L2TP) | **8307**

idle-timeout (System) | **8309**

idle-timeout (System Services) | **8311**

idp (Application Services) | **8312**

idp (logical-systems) | **8313**

idp-policy | **8315**

idp-policy (Application Services) | **8316**

idp-policy (Security) | **8317**

idp (Security) | **8321**

idp (Security Alarms) | **8333**

idp (SNMP) | **8334**

ids-option | **8335**

ids-rule-sets | **8340**

ids-rules | **8341**

ie-set | **8343**

ieee-802.1 (Class of Service) | **8344**

ieee-802.1 (Classifier on Physical Interface) | **8348**

ieee-802.1 (Congestion Notification) | **8350**

ieee-802.1 (Dynamic Classifiers) | **8351**

ieee-802.1 (Dynamic Rewrite Rules) | **8352**

ieee-802.1 (Fibre Channel Interfaces) | **8354**

ieee-802.1 (Host Outbound Traffic) | **8356**

ieee-802.1 (Input Congestion Notification) | **8357**

ieee-802.1 (Output Congestion Notification) | **8358**

ieee-802.1 (Rewrite Rules on Logical Interface) | **8360**

ieee-802.1 (Rewrite Rules on Physical Interface) | **8362**

ieee-802.1ad | **8363**

ieee-802-3az-eee | **8364**

ieee802.1p | **8366**

ietf-mode | **8367**

if-exceeding (QFX) | **8368**

if-exceeding (EX) | **8370**

if-exceeding (Hierarchical Policer) | **8371**

if-exceeding (Policer) | **8373**

if-exceeding-pps (Hierarchical Policer) | **8374**

if-exceeding-pps (Policer) | **8376**

if-route-exists | **8377**

ifl-subscriber | **8379**

igmp | **8381**

igmp (Dynamic Profiles) | **8383**

igmp-querier (QFabric Systems only) | **8385**

igmp-snooping | **8387**

igmp-snooping-options | **8393**

ignore | **8394**

ignore-all | **8394**

ignore-attached-bit | **8396**

ignore-encapsulation-mismatch | **8397**

ignore-entry | **8399**

ignore (Event Policy) | **8399**

ignore-l3-incompletes | **8400**

ignore-lsp-metrics | **8402**

ignore-magic-number-mismatch (Access Group Profile) | **8403**

ignore-magic-number-mismatch (Dynamic Profiles) | **8405**

ignore-memory-overflow | **8407**

ignore-mtu-mismatch | **8408**

ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-overflow | **8409**

ignore-reassembly-overflow | **8411**

ignore-stp-topology-change | **8412**

ignore-server-validation (Application Identification) | **8413**

igp-synchronization | **8414**

ike | **8415**

ike-access-profile | **8418**

ike (Access Profile) | **8419**

ike-esp-nat | **8421**

ike (High Availability) | **8424**

ike-log | **8426**

ike-phase1-failures | **8428**

ike-phase2-failures | **8429**

ike-secrets | **8431**

ike (Security) | **8432**

ike (Security Group VPN Member) | **8438**

ike (Security Group VPN Server) | **8440**

ike (Security IPsec VPN) | **8443**

ike (Security Services) | **8446**

image | **8447**

imap-profile (Security Content Security Policy Antivirus) | **8450**

imap-profile (Security Content Security Policy Content Filtering) | **8451**

immediate-accounting-response | **8452**

immediate-leave | **8454**

immediate-leave (Dynamic IGMP Interface) | **8456**

immediate-leave (Dynamic MLD Interface) | **8458**

immediate-update | **8459**

import (BGP) | **8461**

import (Protocols OSPF) | **8463**

import (Routing Options) | **8465**

import (Classifiers) | **8466**

import (CoS) | **8467**

import-labeled-routes (Routing Instances VPLS) | **8469**

import (MPLS Traffic Engineering Database) | **8471**

import-policy | **8473**

import (Protocols DVMRP) | **8474**

import (Protocols IS-IS) | **8476**

import (Protocols LDP) | **8477**

import (Protocols MSDP) | **8478**

import (Protocols PIM) | **8480**

import (Protocols PIM Bootstrap) | **8482**

import (Protocols RIP) | **8483**

import (Protocols RIPng) | **8485**

import (Rewrite Rules) | **8486**

import-rib | **8488**

import-target | **8489**

imsi-prefix | **8491**

in-place-lsp-bandwidth-update (Protocols MPLS) | **8493**

in-service | **8494**

inactive-media-timeout (Security MGCP) | **8496**

inactive-media-timeout (Security SCCP) | **8497**

inactive-media-timeout (Security SIP) | **8499**

inactivity-asymm-tcp-timeout (Service Set Next Gen Services) | **8500**

inactivity-interval (Fabric Load Balance) | **8501**

inactivity-probe-duration | **8503**

inactivity-timeout | **8505**

inactivity-timeout (services) | **8506**

inactivity-timeout (Applications) | **8508**

inactivity-timeout (Security Persistent NAT) | **8509**

inactivity-timeout (Services RPM) | **8510**

inband-flow-telemetry | **8512**

inband-ztp | **8516**

include-all (for Administrative Groups) | **8518**

include-all (for Fast Reroute) | **8519**

include-any (for Administrative Groups) | **8521**

include-any (for Fast Reroute) | **8522**

include (Diameter Gx Profiles) | **8524**

include (Dynamic Access-Line-Identifier VLANs) | **8525**

include-igp-metric | **8527**

include-ipv6 (Gx-Plus) | **8528**

include-irb-and-l2 | **8529**

include-mp-next-hop | **8532**

include-option-82 (DHCP Local Server) | **8534**

include-sci | **8536**

include-sci (MACsec for MX Series) | **8537**

inclusive | **8539**

incoming-map | **8540**

incoming-queue | **8541**

independent-domain | **8543**

indexed-load-balance | **8545**

indication | **8546**

indication-polarity | **8548**

indirect-next-hop | **8549**

indirect-next-hop-change-acknowledgements | **8551**

ineligible | **8553**

inet | **8554**

inet (AMT Protocol) | **8555**

inet-bgpmcast | **8557**

inet (chassis) | **8558**

inet (CoS) | **8560**

inet (enhanced-hash-key) | **8561**

inet (forwarding options) | **8564**

inet (interfaces (QFX Series)) | **8565**

inet (Interfaces) | **8568**

inet-mdt | **8570**

inet-mdt (Signaling) | **8571**

inet-mvpn (BGP) | **8574**

inet-mvpn (VRF Advertisement) | **8575**

inet (Port Mirroring) | **8576**

inet-precedence | **8578**

inet-precedence (Classifier on Physical Interface) | **8580**

inet-precedence (CoS Classifiers) | **8581**

inet-precedence (CoS Rewrite Rules) | **8582**

inet-precedence (Dynamic Classifiers) | **8584**

inet-precedence (Dynamic Rewrite Rules) | **8585**

inet-precedence (Rewrite Rules on Physical Interface) | **8587**

inet (Routing Policy Condition) | **8588**

inet (Security Static NAT) | **8590**

inet (Subscriber Secure Policy) | **8591**

inet (TDF Subscriber Address) | **8593**

inet (TDF Subscriber Exclude Prefix) | **8594**

inet6 | **8595**

inet6-advertise-interval (VRRP) | **8596**

inet6-backup-router | **8598**

inet6-bgpmcast | **8599**

inet6 (CoS Multidestination) | **8601**

inet6 (enhanced-hash-key) | **8602**

inet6 (forwarding optons) | **8605**

inet6 (interfaces (QFX Series)) | **8607**

inet6 (Interfaces) | **8608**

inet6-mvpn (BGP) | **8610**

inet6-mvpn (VRF Advertisement) | **8612**

inet6-options (Services) | **8613**

inet6-precedence (CoS Rewrite Rules) | **8614**

inet6 (Subscriber Secure Policy) | **8616**

inet6 (TDF Subscriber Address) | **8617**

inet6 (TDF Subscriber Exclude Prefix) | **8619**

inet6-vpn | **8620**

infinity | **8622**

infranet-controller | **8623**

ingress (Analyzer) | **8625**

ingress (Buffer Configuration) | **8626**

ingress (Chained Composite Next Hop) | **8628**

ingress-egress-recycle | **8631**

ingress (LSP) | **8632**

ingress-node-replication | **8634**

ingress-node-replication (EVPN) | **8635**

ingress-only | **8637**

ingress-policer-overhead | **8638**

ingress-policy | **8641**

ingress (Port Mirroring) | **8643**

ingress-queuing-filter | **8644**

ingress-rate-limit | **8646**

ingress-replication | **8647**

ingress-shaping-overhead | **8649**

ingress-stats (Flat-File Accounting Options) | **8651**

ingress (vlans) | **8653**

inherit-advertisement-interval | **8654**

init-command-string | **8655**

init-descriptor | **8657**

initial-route-check | **8658**

initial-tcp-mss | **8659**

initiate-dead-peer-detection | **8661**

initiate-ncp (Dynamic and Static PPP) | **8662**

initiation (Services) | **8664**

inline | **8666**

inline-fpga-crypto | **8668**

inline-jflow | **8669**

inline-jflow (Forwarding Options) | **8671**

inline-monitoring | **8672**

inline-services | **8674**

inline-services (PIC level) | **8676**

inline-tap | **8677**

inner-tag-protocol-id | **8679**

inner-tag-protocol-id (Dynamic VLANs) | **8680**

inner-vlan-id | **8682**

inner-vlan-id (Dynamic VLANs) | **8683**

inner-vlan-id-swap-ranges | **8685**

input | **8687**

input (Analyzer) | **8688**

input-chain | **8690**

input (Congestion Notification) | **8692**

input (Congestion Notification) | **8693**

input-current (T4000 Routers) | **8695**

input (Dynamic Service Sets) | **8697**

input-excess-bandwidth-share | **8698**

input (Fibre Channel Interfaces) | **8700**

input-firewall-optimized-profile | **8701**

input (Forwarding Table) | **8703**

input-hierarchical-policer | **8704**

input-interface-index | **8705**

input (Interfaces) | **8707**

input-list | **8708**

input-native-vlan-push | **8710**

input-packet-rate-threshold | **8711**

input-policer | **8713**

input (Port Mirroring) | **8714**

input-priority-map | **8716**

input (Routing Options RIB) | **8717**

input (Sampling) | **8718**

input-scheduler-map | **8720**

input-shaping-rate (Logical Interface) | **8721**

input-shaping-rate (Physical Interface) | **8723**

input-three-color | **8724**

input-traffic-control-profile | **8726**

input-traffic-control-profile-remaining | **8728**

input-vlan-map | **8729**

input-vlan-map (Aggregated Ethernet) | **8731**

input-vlan-map (Dynamic Interfaces) | **8732**

inputs (Analytics) | **8734**

insert | **8738**

inspection-profile | **8739**

install | **8741**

install-nexthop | **8742**

install (Protocols MPLS) | **8744**

install (Routing Options) | **8747**

instance | **8748**

instance (Inline Monitoring) | **8750**

instance (OAM) | **8753**

instance (CASB) | **8754**

instance (Ephemeral Database) | **8756**

instance-export | **8759**

instance-import | **8760**

instance (Port Mirroring) | **8762**

instance-role | **8764**

instance (Sampling) | **8766**

instance-shared | **8769**

instance (Traffic Load Balancer) | **8770**

instance-type | **8773**

integer | **8777**

intelligent-prescreening | **8779**

intelligent-offload-tunable (IDP) | **8780**

inter-area-prefix-export | **8782**

inter-area-prefix-import | **8783**

inter-as (Routing Instances) | **8785**

inter-domain | **8788**

inter-region | **8789**

inter-region-segmented | **8791**

inter-region-template | **8793**

inter-switch-link | **8795**

interconnect | **8797**

interconnect-device (Aliases) | **8800**

interconnect-device (Chassis) | **8802**

interconnect-device (Flow Groups) | **8804**

interconnect-multihoming-peer-gateways | **8806**

interconnected-vlan-list | **8808**

interconnected-vni-list | **8809**

interface (services) | **8811**

interface (VRRP) | **8812**

interface (Logical Interfaces) | **8814**

interface (Precision Time Protocol) | **8815**

interface (LLDP) | **8817**

interface (802.1X) | **8819**

interface (Access Port Security) | **8829**

interface (Accounting or Sampling) | **8831**

interface (advance-policy-based-routing) | **8833**

interface (Analyzer) | **8835**

interface (Backup Selection OSPF) | **8837**

interface (BOOTP) | **8841**

interface (BPDU Blocking) | **8842**

interface (Bridge Domains) | **8845**

interface (Captive Portal) | **8846**

interface (Chassis Cluster) | **8850**

interface (DCBX) | **8851**

interface-delete (Subscriber Management or DHCP Client Management) | **8853**

interface-description (DHCP Local Server) | **8854**

interface-description (DHCP Relay Agent) | **8856**

interface-description-format | **8858**

interface (DHCP Local Server) | **8860**

interface (DHCP Relay Agent) | **8863**

interface (DHCP Security for MX Series) | **8866**

interface (DHCP Spoofing Prevention) | **8867**

interface (DNS, Port, and TFTP Packet Forwarding or Relay Agent) | **8868**

interface-down | **8870**

interface (Dynamic IGMP) | **8871**

interface (Dynamic Interface Sets) | **8873**

interface (Dynamic MLD) | **8875**

interface (Dynamic Router Advertisement) | **8877**

interface (Dynamic Routing Instances) | **8879**

interface (Dynamic Routing Options) | **8880**

interface (EVPN Routing Instances) | **8882**

interface (Export Profiles) | **8884**

interface (Fibre Channel Fabric) | **8886**

interface (FIP) | **8887**

interface (FIP Snooping) | **8889**

interface-group | **8890**

interface-group (Decapsulate GRE) | **8892**

interface-group-down | **8893**

interface-group (Routing Options) | **8894**

interface (Hierarchical CoS Schedulers) | **8896**

interface (Host Fast Reroute) | **8897**

interface-id | **8898**

interface (IEEE 802.1ag OAM Connectivity-Fault Management) | **8900**

interface (IGMP Snooping) | **8902**

interface-in (Security Forwarding Options) | **8903**

interface (Layer 2 Protocol Tunneling) | **8904**

interface (LLDP-MED) | **8906**

interface-mac-ip-limit | **8908**

interface-mac-limit | **8910**

interface-mac-limit (VPLS) | **8913**

interface-map | **8915**

interface-mib (Enhanced Subscriber Management) | **8916**

interface (MLD Snooping) | **8918**

interface-mode | **8920**

interface-monitor | **8922**

interface (Monitoring) | **8923**

interface (MPLS) | **8925**

interface (Multicast Scoping) | **8927**

interface (Multicast Static Routes) | **8928**

interface (Multichassis Protection) | **8930**

interface (MVRP) | **8931**

interface-name | **8934**

interface-name | **8935**

interface-name (DHCP Local Server) | **8937**

interface-name (DHCP Relay Agent) | **8938**

interface (Next-Hop Group) | **8939**

interface-none | **8941**

interface-notification (programmable-rpd) | **8942**

interface-num | **8943**

interface (OAM CFM, for EX Series Switch Only) | **8945**

interface (OAM Link-Fault Management) | **8947**

interface-out (Security Forwarding Options) | **8948**

interface (PIC Bundle) | **8949**

interface (PoE) | **8951**

interface (Port Mirroring) | **8952**

interface (Port Mirroring) | **8954**

interface (Port Mirroring) | **8956**

interface (Power over Ethernet) | **8957**

interface-profile | **8959**

interface (Protocols DVMRP) | **8960**

interface (Protocols IGMP) | **8962**

interface (Protocols IPv6 Neighbor Discovery) | **8964**

interface (Protocols IS-IS) | **8966**

interface (Protocols Layer 2 Circuit) | **8974**

interface (Protocols Layer 2 VPN) | **8977**

interface (Protocols LDP) | **8978**

interface (Protocols Link Management) | **8980**

interface (Protocols MLD) | **8981**

interface (Protocols MPLS) | **8983**

interface (Protocols OSPF) | **8984**

interface (Protocols PIM) | **8991**

interface (Protocols RIFT) | **8994**

interface (Protocols Router Discovery) | **8998**

interface (Protocols RSVP) | **8999**

interface (RA Guard) | **9002**

interface-range (QFX Series) | **9003**

interface-range | **9006**

interface (Redundant Trunk Groups) | **9008**

interface-routes | **9010**

interface (Routing Instances) | **9012**

interface (Routing Options) | **9014**

interface (Scoping) | **9015**

interface (Secure Access Port) | **9017**

interface (Security NAT ARP Proxy) | **9018**

interface (Security NAT NDP Proxy) | **9020**

interface (Security Source NAT) | **9021**

interface (Security Source NAT Rule Set) | **9022**

interface (Service Interfaces) | **9023**

interface (Service Redundancy Group Monitor) | **9025**

interface-service (Services Interfaces) | **9027**

interface (Services) | **9028**

interface (Services Flow Tap) | **9029**

interface (Services PIC) | **9031**

interface (Session PICs) | **9033**

interface-set | **9034**

interface-set (ANCP) | **9036**

interface-set (Dynamic Profiles) | **9037**

interface-set (Dynamic VLAN Interface Sets Association) | **9040**

interface-set (Dynamic VLAN Interface Sets Definition) | **9042**

interface-set (Ethernet Interfaces) | **9044**

interface-set (Hierarchical Schedulers) | **9045**

interface-set (IP Demux Interfaces) | **9046**

interface-shared | **9048**

interface-shared-with | **9049**

interface-shutdown-action | **9050**

interface (SLAAC Snooping) | **9052**

interface (SNMP) | **9054**

interface (SNMP RMON History) | **9056**

interface (Spanning Tree) | **9057**

interface-specific | **9060**

interface-specific (Dynamic Firewalls) | **9061**

interface-specific (Firewall Filters) | **9063**

interface-state | **9064**

interface (Static MAC Bypass) | **9067**

interface (Static Subscriber Group) | **9068**

interface (Static Subscriber Username) | **9070**

interface-status-send-rdi | **9071**

interface-status-tlv | **9073**

interface (Storm Control) | **9074**

interface-switch | **9076**

interface (Switching Options) | **9078**

interface (System Services DHCP) | **9079**

interface-tag (DHCP Local Server) | **9081**

interface-tag (DHCP Relay Server) | **9083**

interface-traceoptions (DHCP) | **9085**

interface-traceoptions (System Services DHCP) | **9087**

interface-transmit-statistics | **9090**

interface-type (Protocols OSPF) | **9091**

interface-type (Protocols RIP) | **9095**

interface (Unknown Unicast Forwarding) | **9096**

interface (Virtual Tunnel in Routing Instances) | **9097**

interface (VoIP) | **9099**

interface (VPLS Mesh-Group) | **9102**

interface (VPLS Multihoming for FEC 129) | **9103**

interface (VPLS Routing Instances) | **9104**

interface-client-limit (DHCP Local Server) | **9106**

interface-client-limit (DHCP Relay Agent) | **9108**

interfaces (VNF) | **9111**

interfaces | **9113**

interfaces (Aggregated Multiservices) | **9115**

interfaces (Analytics) | **9117**

interfaces (Analytics Resource) | **9119**

interfaces (ANCP) | **9121**

interfaces (Autoinstallation) | **9122**

interfaces (Chassis Synchronization Source) | **9124**

interfaces (Class of Service) | **9126**

interfaces (Class of Service) | **9128**

interfaces (Class of Service) | **9130**

interfaces (Dynamic CoS Definition) | **9133**

interfaces (EX Series switches) | **9135**

interfaces external | **9144**

interfaces (Junos Node Slicing) | **9145**

interfaces (MACsec) | **9147**

interfaces (MACsec for MX Series) | **9149**

interfaces (OVSDB) | **9151**

interfaces (Protocols OpenFlow) | **9153**

interfaces (Q-in-Q Tunneling) | **9154**

interfaces (QFX Series, ACX Series) | **9156**

interfaces (Security Zones) | **9165**

interfaces (Services Dynamic Flow Capture) | **9167**

interfaces (sFlow) | **9168**

interfaces (Static and Dynamic Subscribers) | **9170**

interfaces (Video Monitoring) | **9176**

interfaces (Voice Services) | **9181**

interim-logging-interval | **9182**

interim-rate (Access) | **9183**

interleave-fragments | **9185**

internal | **9186**

internal-forwarding-bandwidth | **9187**

internal-node | **9189**

internal (Security IPsec) | **9190**

internet-options | **9192**

interpolate (Drop Profiles) | **9197**

interval | **9198**

interval (services) | **9201**

interval (Protocols PIM) | **9202**

interval (OAM) | **9203**

interval (RMON) | **9206**

interval (Accounting Options) | **9207**

interval (CFM MEP) | **9209**

interval (EX Series Switch Only) | **9210**

interval (Flat-File Accounting Options) | **9211**

interval (Health Monitor) | **9213**

interval (Interfaces) | **9214**

interval (PoE) | **9216**

interval (Power over Ethernet) | **9217**

interval (Security Antivirus) | **9219**

interval (Security IDP) | **9220**

interval (Services) | **9221**

interval (SNMP RMON) | **9223**

intra-area-prefix | **9224**

intra-as | **9226**

invalid-authentication-entry-timeout (Services User Identification Active Directory and ClearPass) | **9227**

ioc-npc-connectivity | **9230**

ip-action (Security IDP Rulebase IPS) | **9231**

ip-address (Access Profile) | **9234**

ip-address-first | **9235**

ip-address-owner | **9236**

ip-address (Port Mirroring) | **9238**

ip-can-type (PCRF Partition) | **9239**

ip (IDS Screen Next Gen Services) | **9241**

ip-in-udp | **9243**

ip-monitoring | **9244**

ip-monitoring (Services) | **9246**

ip-options-protocol-queue | **9249**

ip-prefix-routes | **9250**

ip-prefix-support | **9255**

ip-protocol-mapping (Application Identification) | **9258**

ip-query (Identity Management Advanced Query) | **9259**

ip-reassembly | **9262**

ip-reassembly (L2TP) | **9263**

ip-reassembly-rules (Service Set) | **9265**

ip (Security IDP Custom Attack) | **9266**

ip (Security Screen) | **9268**

ip-source-guard | **9272**

ip-source-guard (MX Series) | **9274**

ip-subscriber | **9275**

ip-swap | **9277**

ip-sweep | **9279**

ip-tunnel-rpf-check | **9281**

ip-user-mapping | **9283**

ip-version | **9285**

ip-version (Family MPLS) | **9286**

ipc | **9288**

ipcp-suggest-dns-option | **9289**

ipfix-sw-mode | **9291**

ipip | **9292**

ipip (Ids-option) | **9293**

ips | **9295**

ipsec | **9298**

ipsec-group-vpn (Security Policies) | **9300**

ipsec (High Availability) | **9301**

ipsec-inside-interface | **9304**

ipsec-policy | **9306**

ipsec-sa | **9307**

ipsec-sa (Protocols BGP) | **9308**

ipsec-sa (Security Group VPN) | **9310**

ipsec (Security) | **9312**

ipsec (Security Group VPN Member) | **9314**

ipsec (Security Group VPN Server) | **9317**

ipsec (Security Interfaces) | **9319**

ipsec-traffic-selector | **9321**

ipsec-vpn-options | **9323**

ipsec-vpn-rules | **9325**

ipsec-vpn (Security Policies) | **9326**

ipv4-address (Steering Path) | **9327**

ipv4-dscp | **9329**

ipv4 (Family MPLS) | **9330**

ipv4-flow-table-size | **9332**

ipv4-mask (HTTP Header Enrichment) | **9335**

ipv4-or-value (HTTP Header Enrichment) | **9337**

ipv4-prefix | **9338**

ipv4 (Security IDP Signature Attack) | **9340**

ipv4-template | **9344**

ipv4-template (Services) | **9345**

ipv6-address (Steering Path) | **9347**

ipv6-extended | **9348**

ipv6-extended-attrib | **9349**

ipv6-extension-header | **9350**

ipv6-extension-header (IDS Screen Next Gen Services) | **9353**

ipv6-extension-header-limit | **9355**

ipv6 (Family MPLS) | **9357**

ipv6-flow-label | **9359**

ipv6-flow-table-size | **9360**

ipv6-malformed-header | **9362**

ipv6-mask (HTTP Header Enrichment) | **9363**

ipv6-multicast-interfaces | **9364**

ipv6-or-value (HTTP Header Enrichment) | **9366**

ipv6-origin | **9367**

ipv6-prefix | **9368**

ipv6-source-guard | **9370**

ipv6-source-guard-sessions | **9372**

ipv6-target | **9374**

ipv6-template | **9375**

ipv6-template (Services) | **9376**

ipv6-tunneling | **9377**

ipvacl-firewall-l2-profile | **9378**

iq-policing-filter | **9380**

ir-mode | **9381**

iracl-firewall-ipv4-profile | **9383**

irb | **9384**

irb (Easy EVPN LAG Configuration) | **9386**

irb (Interfaces) | **9389**

irb-symmetric-routing | **9394**

is-ma | **9396**

isdn-options | **9397**

isid | **9399**

isid-list | **9400**

isis | **9401**

isis-instance | **9403**

iso-vpn | **9404**

isolated | **9406**

isolated-vlan | **9407**

isolated-vlan (MX Series) | **9408**

isolation-id | **9410**

isolation-vlan-id | **9411**

iteration-count | **9412**

iteration-period | **9414**

ivacl-firewall-no-portrange-profile | **9415**

ivlan-cfi (RFC 2544 Benchmarking) | **9416**

ivlan-id (RFC 2544 Benchmarking) | **9417**

ivlan-priority (RFC 2544 Benchmarking) | **9419**

iwf-params (tdm-options) | **9420**

J | 9422

jdaf | **9422**

jeb | **9423**

jflow-log (Interfaces) | **9426**

jflow-log (Services) | **9427**

jnu-management | **9429**

join-load-balance | **9431**

join-prune-timeout | **9433**

join-timer (MVRP) | **9434**

json | **9436**

jsrc (Access Profile) | **9438**

jsrc (JSRC) | **9439**

jsrc-partition | **9441**

juniper-enhanced | **9442**

juniper-express-engine | **9444**

juniper-local | **9447**

K | 9448

kaspersky-lab-engine | **9450**

keep | **9452**

keep-address-and-control | **9455**

keep-alive (Protocols MSDP) | **9457**

keep-existing-steering | **9459**

keep-incoming-circuit-id (DHCP Relay Agent) | **9460**

keep-incoming-interface-id (DHCP Relay Agent) | **9462**

keep-incoming-remote-id (DHCP Relay Agent) | **9463**

keep-multiplier | **9465**

keepalive | **9466**

keepalive-interval | **9468**

keepalive-retries | **9469**

keepalive-time (Chassis) | **9471**

keepalive-time (Protocols OAM) | **9472**

keepalive-timeout | **9474**

keepalives | **9475**

keepalives (Dynamic Profiles) | **9478**

kernel-replication | **9480**

key | **9481**

key (gre) | **9482**

key (interfaces) | **9484**

key (Authentication Keychain) | **9485**

key-chain (Authentication Keychain) | **9488**

key-chain (Protocols PIM) | **9490**

key-exchange | **9491**

key (Junos FIPS) | **9494**

key-length | **9495**

key (MACsec) | **9497**

key (MACsec for MX Series) | **9498**

key-pair | **9500**

key-server-priority (MACsec) | **9502**

key-server-priority (MACsec for MX Series) | **9503**

krt-nexthop-ack | **9505**

krt-nexthop-ack-timeout | **9506**

L | 9508

l2-channel-error-threshold (chassis) | **9520**

l2-channel-errors (chassis) | **9521**

l2-domain-id-for-l3 | **9523**

l2-learning | **9524**

l2-querier | **9527**

l2-smart-policy | **9528**

l2-stats (Flat-File Accounting Options) | **9530**

l2circuit | **9531**

l2circuit-control-passthrough | **9534**

l2ckt | **9535**

l2pt (Destination Tunnels) | **9537**

l2tp | **9541**

l2tp (DDoS Protection) | **9544**

l2tp-access-profile | **9547**

l2tp (Group Profile) | **9548**

l2tp-interface-id | **9550**

l2tp-maximum-session (Service Interfaces) | **9551**

l2tp (Profile) | **9553**

l2tp-tunnel-session-identifier | **9558**

l2vpn | **9559**

l2vpn-id | **9562**

l2vpn (routing-options) | **9564**

l2vpn-use-bgp-rules-edit-protocols-bgp | **9565**

l3-interface-ingress-counting | **9567**

l3-interface (VLAN) | **9568**

l3vpn (ingress) | **9571**

l3vpn (transit) | **9573**

label | **9574**

label-allocation | **9576**

label-block-size | **9577**

label-position | **9579**

label-switched-path | **9580**

label-switched-path (Protocols IS-IS) | **9582**

label-switched-path (Protocols Link Management) | **9583**

label-switched-path (Protocols MPLS) | **9585**

label-switched-path (Protocols MVPN) | **9590**

label-switched-path (Protocols OSPF) | **9592**

label-switched-path-template | **9593**

label-switched-path-template (Container LSP) | **9595**

label-switched-path-template (Multicast) | **9596**

label-withdrawal-delay | **9598**

labeled-bgp | **9600**

labeled-unicast | **9601**

labeled-unicast (Protocols BGP) | **9603**

lacp | **9605**

lacp (802.3ad) | **9607**

lacp (802.3ad) (MX Series) | **9609**

lacp (Aggregated Ethernet) | **9610**

lacp (Interfaces) | **9614**

lacp (Protocols) | **9615**

land | **9617**

land-attack-check | **9619**

land-attack-check (IDS MS-MPC) | **9620**

land (IDS Screen Next Gen Services) | **9622**

lane | **9623**

lane-all | **9626**

language (Scripts) | **9630**

large | **9632**

large (IDS Screen Next Gen Services) | **9633**

large-community-count | **9634**

laser-enable | **9635**

last-block-recycle-timeout | **9637**

latency (Interfaces) | **9638**

latency-threshold | **9639**

layer2-control | **9641**

layer2 (enhanced-hash-key) | **9644**

layer2-liveness-detection (Receive) | **9647**

layer2-liveness-detection (Send) | **9649**

layer2-map | **9651**

layer2-policer | **9653**

layer2-policer (EX) | **9654**

layer2-policer (Hierarchical Policer) | **9656**

layer2-protocol-tunneling | **9658**

layer2-unicast-replies | **9661**

layer2-policer (firewall) | **9662**

layer3-domain-identifier | **9666**

lcd-menu | **9667**

lcp-max-conf-req | **9669**

lcp-renegotiation | **9671**

lcp-restart-timer | **9672**

ldap-options | **9673**

ldap-server | **9676**

ldap-url | **9678**

ldp | **9679**

ldp-p2mp | **9685**

ldp-stitching (Protocols OSPF) | **9686**

ldp-tunneling | **9687**

leaf | **9689**

leaf-tunnel-limit-inet (MVPN Selective Tunnels) | **9691**

leaf-tunnel-limit-inet6 (MVPN Selective Tunnels) | **9693**

learn-sip-register | **9695**

lease-time | **9696**

lease-time (dhcp-client) | **9698**

lease-time-threshold (DHCP Local Server and DHCP Relay Agent) | **9699**

lease-time-validation (DHCP Local Server and DHCP Relay Agent) | **9701**

leasequery (DHCP Relay Agent) | **9702**

least-fill | **9704**

leave-timer (MVRP) | **9705**

leaveall-timer (MVRP) | **9707**

led-beacon | **9710**

legacy-redirect-ip-action | **9711**

less-than | **9713**

level (Connectivity Fault Management) | **9715**

level (OAM) | **9716**

level (Disk Partition) | **9718**

level (App Engine Virtual Machine Management Service) | **9719**

level (CFM MEP) | **9721**

level (EX Series Switch Only) | **9722**

level (Global IS-IS) | **9724**

level (IS-IS Interfaces) | **9728**

level (Protocols RIFT) | **9734**

level (Services) | **9735**

license | **9737**

license-mode | **9739**

license-server | **9740**

lie-receive-address (Protocols RIFT) | **9742**

lie-transmit-address (Protocols RIFT) | **9744**

lifetime (Router Advertisement) | **9745**

lifetime (Router Discovery) | **9747**

lifetime-seconds | **9749**

lifetime-seconds (Security) | **9750**

light | **9752**

light-weight-dhcpv6-relay | **9755**

limit | **9757**

limit (Content Security Policy) | **9758**

limit-ports-per-address | **9759**

limit (Security Log) | **9760**

limit (Security SCTP) | **9762**

limit-session | **9764**

limit-session (IDS Screen Next Gen Services) | **9765**

limited-ifl-scaling | **9768**

line-encoding (E1 Options) | **9769**

line-encoding (T1 Options) | **9770**

line-identity (Dynamic Access-Line-Identifier VLANs) | **9772**

line-interval (Dampen Options) | **9773**

line-loopback | **9775**

line-protocol | **9776**

line-rate | **9778**

line-rate (Interfaces) | **9779**

linerate-mode | **9780**

link (Access) | **9781**

link-adjacency-loss | **9783**

link-degrade-monitor | **9784**

link-discovery | **9787**

link-down (QFX Series) | **9788**

link-down | **9790**

link-event-rate | **9791**

link-fault-management | **9792**

link-group-protection (Protocols IS-IS) | **9794**

link-layer-broadcast-inet-check | **9796**

link-layer-overhead | **9797**

link-management | **9798**

link-mode | **9800**

link-mtu | **9802**

link-protection | **9803**

link-protection (Chassis Cluster) | **9805**

link-protection (Dynamic LSPs) | **9807**

link-protection (Host Fast Reroute) | **9808**

link-protection (non-LACP) | **9810**

link-protection (Protocols IS-IS) | **9811**

link-protection (Protocols LACP) | **9813**

link-protection (Protocols OSPF) | **9814**

link-protection (RSVP) | **9816**

link-protection (Static LSPs) | **9818**

link-protection-sub-group (802.3ad) | **9820**

link-protection-sub-group (aggregated-ether-options) | **9821**

link-speed | **9823**

link-speed (Aggregated Ethernet) | **9826**

link-speed (Aggregated SONET/SDH) | **9828**

link-speed (Interfaces) | **9830**

link-to-disable | **9831**

link-to-monitor | **9832**

link-tracking (Network Isolation Groups) | **9833**

linkdown (satellite-policies alarm) | **9835**

linked-pool-aggregation (Address-Assignment Pools) | **9837**

linktrace | **9838**

list | **9839**

list (Security Content Filtering Block Mime) | **9841**

listen | **9842**

listening-mode | **9843**

liveness-detection (iccp-peer) | **9845**

liveness-detection | **9846**

liveness-detection (high availability) | **9848**

lldp | **9851**

lldp-med (Ethernet Switching) | **9857**

lldp-priority | **9860**

lmi (Ethernet OAM) | **9861**

lmi-type | **9863**

lmp-control-channel | **9864**

lmp-protocol | **9865**

load | **9867**

load access-security slaac-snooping persistent-file | **9870**

load-balance (ecmp) | **9872**

load-balance | **9874**

load-balance-algorithm | **9876**

load-balance (Forwarding Options) | **9878**

load-balance-group | **9880**

load-balance-label-capability | **9882**

load-balance profile1/profile2 | **9883**

load-balance (Protocols RSVP) | **9885**

load-balance-stateful (Aggregated Ethernet Interfaces) | **9886**

load-balancing-options (Aggregated Multiservices) | **9888**

load-balancing-options (Service Set) | **9890**

load-interval | **9892**

load-scripts-from-flash (Scripts) | **9893**

load-threshold | **9894**

load-type (Aggregated Ethernet Interfaces) | **9896**

loc-rib | **9897**

local | **9899**

local (certificates) | **9900**

local-address (Junos Telemetry Interface) | **9902**

local-address (Protocols AMT) | **9904**

local-address (Protocols BGP) | **9905**

local-address (Protocols Link Management) | **9908**

local-address (Protocols MSDP) | **9910**

local-address (Protocols PIM) | **9911**

local-address (Routing Options) | **9913**

local (Analytics Collector) | **9914**

local-as | **9916**

local-authentication (Dynamic PPP Options) | **9919**

local-bias (ae) | **9921**

local-bias | **9922**

local-bias (ae load-balance) | **9923**

local-bw-override-threshold | **9925**

local-category (Next Gen Services Service-Set Local System Logging) | **9926**

local-certificate | **9929**

local-certificate (System) | **9930**

local-certificate (Security) | **9932**

local-chap | **9933**

local-decision (PCRF Partition) | **9934**

local-dump (Flow Aggregation) | **9938**

local-endpoint | **9939**

local-engine | **9941**

local (Flat-File Access Profile) | **9942**

local-gateway (IPSec) | **9944**

local-gateway (L2TP LNS) | **9945**

local-id | **9947**

local-id (High Availability) | **9948**

local-identity | **9950**

local-interface (IPv6) | **9952**

local-ip-addr (ICCP) | **9953**

local-ipv4-address | **9955**

local-ipv6-address | **9956**

local-key-pair | **9958**

local-log-tag (Next Gen Services Service-Set System Logging) | **9959**

local-loopback | **9960**

local-minimum-links-threshold | **9962**

local-name | **9964**

local-password | **9966**

local-port-range | **9967**

local-ports | **9969**

local-preference | **9971**

local-priority | **9973**

local-replication | **9974**

local-server-group (DHCP Relay Agent Option) | **9976**

local-switching (Layer 2 Circuits) | **9977**

local-switching (VPLS) | **9979**

local-traffic-selector | **9980**

locality-bias (MX Series Virtual Chassis) | **9981**

localized-fib | **9983**

location (DHCP Snooping Database) | **9984**

location (DHCP Relay Agent) | **9985**

location (DHCP Snooping Database) | **9987**

location (DHCPv6 Relay Agent) | **9988**

location (LLDP-MED) | **9990**

location (SNMP) | **9992**

location (System) | **9993**

location (Virtual Chassis) | **9995**

locator | **9997**

lockout | **9999**

lockout-timeout (L2TP Destination Lockout) | **10001**

log | **10002**

log (aacl) | **10005**

log-and-generate-ais | **10006**

log-attacks | **10007**

log-interval (IGMP Interface) | **10009**

log-interval (MLD Interface) | **10011**

log-interval (PIM Entries) | **10012**

log-interval (Protocols MSDP) | **10014**

log (Logical Systems and Tenant Systems) | **10016**

log-prefix (Interfaces) | **10020**

log-prefix (L2TP) | **10021**

log-prefix (Services) | **10022**

log-prefix (System) | **10023**

log-rotate-frequency | **10025**

log (Security) | **10026**

log (Security GTP) | **10032**

log (Security IDP Sensor Configuration) | **10033**

log (Security Policies) | **10035**

log (Security SCTP) | **10037**

log (Services) | **10038**

log (Spanning Trees) | **10040**

log-updown | **10042**

log-updown (Protocols LDP) | **10043**

log-updown (Protocols MPLS) | **10045**

log-warning (Multicast Forwarding Cache) | **10047**

log-warning (Protocols MSDP) | **10049**

log-attacks | **10050**

log-attacks (rulebase-exempt) | **10052**

logging (IDS MS-DPC) | **10053**

logging-rule (PCC Action Profile) | **10054**

logging (Services) | **10056**

logical-bandwidth-policer | **10058**

logical-domain-identity-management | **10059**

logical-interface-aggregate-statistics | **10062**

logical-interface-chassis-redundancy (MX Series Virtual Chassis) | **10063**

logical-interface (DDoS Flow Detection) | **10064**

logical-interface-fpc-redundancy (Aggregated Ethernet Subscriber Interfaces) | **10067**

logical-interface-fpc-redundancy (PS Interface) | **10068**

logical-interface-policer | **10069**

logical-interface-unit-range | **10072**

logical-system | **10073**

logical-system (SNMP) | **10075**

logical-system (Diameter Peer) | **10076**

logical-system (Diameter Transport) | **10078**

logical-system-mux | **10079**

logical-system-name (DHCP Local Server) | **10082**

logical-system-name (DHCP Relay Agent) | **10083**

logical-system-name (Static Subscribers) | **10085**

logical-system (System Security Profile) | **10086**

logical-system-trap-filter | **10088**

logical-system (Tunnel Profile) | **10089**

logical-systems | **10090**

logical-systems (All) | **10092**

logical-tunnel-options | **10093**

login | **10095**

login (Access) | **10100**

logout-response-timeout (PCRF Partition) | **10101**

long-lived (Graceful Restart for BGP Helper) | **10103**

long-lived (Graceful Restart for BGP Restarter) | **10105**

longest-match | **10108**

longest-match (Static Routes) | **10109**

lookup-intact-on-commit | **10112**

loop-detect | **10113**

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) | **10117**

loopback-clear-timer | **10119**

loopback (Fibre Channel Interface) | **10120**

loopback-firewall-optimization | **10121**

loopback-firewallv6-optimization | **10123**

loopback (Local and Remote) | **10125**

loopback-port | **10126**

loopback-remote | **10128**

loopback (Serial) | **10129**

loopback-tracking | **10131**

loops (Autonomous System) | **10132**

loops (BGP Address Family) | **10134**

loose-authentication-check | **10137**

loose-check | **10138**

loose-source-route-option (IDS Screen Next Gen Services) | **10139**

loss-delay (querier) | **10140**

loss-priority | **10142**

loss-priority (Classifiers) | **10143**

loss-priority (Classifiers) | **10145**

loss-priority (Classifiers and Rewrite Rules) | **10146**

loss-priority (Frame Relay Loss Priority) | **10148**

loss-priority (Drop Profiles) | **10150**

loss-priority (Dynamic Schedulers) | **10151**

loss-priority (Firewall Filter) | **10152**

loss-priority (Firewall Filter Action) | **10154**

loss-priority high then discard (Three-Color Policer) | **10155**

loss-priority-maps | **10157**

loss-priority-maps (Assigning to an Interface) | **10159**

loss-priority-maps (CoS) | **10160**

loss-priority-maps (CoS Interfaces) | **10161**

loss-priority (Rewrite Rules) | **10163**

loss-priority (Rewrite Rules) | **10164**

loss-priority-rewrites | **10166**

loss-priority-rewrites (Assigning to an Interface) | **10167**

loss-priority (Scheduler Drop Profiles) | **10169**

loss-priority (Simple Firewall Filter) | **10170**

loss (querier) | **10171**

loss (responder) | **10173**

loss-threshold | **10175**

low-latency | **10176**

low-watermark | **10177**

lowest-priority-defect | **10179**

lrf-profile (Service Set) | **10180**

lsa-refresh-interval | **10182**

lsp-attributes | **10183**

lsp-cleanup-timer | **10185**

lsp-external-controller | **10187**

lsp-lifetime | **10189**

lsp-next-hop (CoS-Based Forwarding) | **10191**

lsp-next-hop (Static Routes) | **10192**

lsp-switch | **10194**

lsping-channel-type | **10195**

lsq-failure-options | **10197**

ltv-syslog-interval (System Process) | **10198**

M | 10199

mac | **10215**

mac (MPLS) | **10217**

mac-address | **10218**

mac-address (Accept Source Mac) | **10219**

mac-address (DHCP Local Server) | **10221**

mac-address (DHCP Relay Agent) | **10223**

mac-address (Dynamic Access-Internal Routes) | **10224**

mac-address (MACsec) | **10226**

mac-address (VLAN and Stacked VLAN Interfaces) | **10227**

mac-flush | **10229**

mac-ip-table-size | **10231**

mac-learn-enable | **10233**

mac-learning-priority | **10234**

mac-limit | **10236**

mac-limit (VLAN) | **10238**

mac-limit (Access Port Security) | **10241**

mac-list | **10243**

mac-lookup-length | **10244**

mac-move-limit | **10246**

mac-notification | **10249**

mac (Option 82) | **10250**

mac-persistence-timer | **10252**

mac-pinning | **10254**

mac-pinning (EVPN Routing Instances) | **10256**

mac-radius | **10257**

mac-rewrite | **10259**

mac (Static MAC-Based VLANs) | **10261**

mac-statistics | **10263**

mac-table-aging-time (MX Series) | **10265**

mac-table-aging-time | **10267**

mac-table-size (VPLS) | **10269**

mac-table-size | **10271**

mac-validate | **10273**

mac-validate (Dynamic IP Demux Interface) | **10275**

mac-vrf | **10276**

macsec | **10278**

macsec (MX Series) | **10281**

maintain-subscriber (Subscriber Management) | **10284**

maintenance-association | **10285**

maintenance-domain | **10287**

major | **10290**

major-event (Process) | **10291**

major-event (System) | **10293**

major-ring-name | **10295**

make-before-break (LDP) | **10296**

managed | **10298**

managed-configuration | **10301**

managed-configuration (Dynamic Router Advertisement) | **10302**

managed-services | **10304**

management | **10305**

management-ethernet (Alarm) | **10307**

management-instance | **10309**

management (PoE) | **10310**

management (Security Zones) | **10311**

mandatory | **10313**

manual | **10315**

manual (Junos-FIPS Software) | **10316**

manual (Junos OS) | **10318**

manual (Security IPsec) | **10320**

manual switch | **10323**

many-to-one (Aggregated Multiservices) | **10324**

map-dest-bmac-to-dest-cmac | **10326**

map (Domain Map) | **10327**

map-e | **10329**

map-entry-timeout | **10333**

map-t | **10334**

mapped-port (Security Static NAT) | **10336**

mapping | **10338**

mapping (VNF) | **10340**

mapping-agent-election | **10342**

mapping-range | **10344**

mapping-refresh | **10345**

mapping-server-entry | **10346**

mapping-timeout | **10348**

mapping-timeout (Source NAT Next Gen Services) | **10349**

mapping-type | **10351**

mapping-type (Source NAT Next Gen Services) | **10352**

mark-interface (RA Guard) | **10353**

mark-local-ip-host-routes | **10355**

martians | **10356**

mask uli | **10358**

master | **10360**

master-only | **10364**

master-password | **10365**

mastership-priority | **10367**

match (Syslog) | **10369**

match-direction | **10371**

match-direction (ipsec-vpn) | **10372**

match-direction (IDS) | **10373**

match-direction (IP Reassembly Rule) | **10375**

match-direction (NAT Next Gen Services) | **10376**

match-direction (PCP) | **10377**

match-direction (Services CoS) | **10378**

match-direction (Services NAT) | **10380**

match-direction (Services Stateful Firewall) | **10381**

match (Forwarding-Options) | **10382**

match-list | **10383**

match (Next Gen Services Global System Logging) | **10386**

match-option | **10387**

match-order (Dynamic Firewalls) | **10389**

match-rules-on-reverse-flow | **10391**

match-rules-on-reverse-flow (Next Gen Services) | **10392**

match (Security Destination NAT) | **10393**

match (Security IDP Policy) | **10395**

match (Security Policies) | **10396**

match (Security Policies Global) | **10398**

match (Security Source NAT) | **10401**

match (Security Static NAT) | **10402**

match (Services CoS Next Gen Services) | **10404**

match (Stateful Firewall Rule Next Gen Services) | **10406**

match-strings | **10407**

matches | **10409**

matches (Event Policy) | **10412**

max-advertisement-interval (Dynamic Router Advertisement) | **10413**

max-advertisement-interval (Protocols IPv6 Neighbor Discovery) | **10415**

max-advertisement-interval (Protocols Router Discovery) | **10416**

max-age | **10418**

max-areas | **10420**

max-bypasses | **10422**

max-configurations-on-flash (System) | **10423**

max-connection-duration | **10424**

max-connections | **10426**

max-data-sessions-per-subscriber | **10427**

max-datasize | **10429**

max-db-size (Enhanced Subscriber Management) | **10432**

max-drop-flows | **10435**

max-duplicates | **10436**

max-failures | **10438**

max-flows | **10439**

max-hops | **10441**

max-login-sessions | **10442**

max-login-sessions-per-node | **10444**

max-lsp-size | **10446**

max-message-length | **10448**

max-outstanding-requests (Diameter Applications) | **10449**

max-packet-memory-ratio | **10451**

max-packets-per-second (Sampling) | **10452**

max-pending-accounting-stops (Access Profile) | **10454**

max-policies (Event Policy) | **10455**

max-power | **10456**

max-queues | **10458**

max-queues-per-interface | **10461**

max-reass-packet-memory-ratio | **10463**

max-retrans-time | **10464**

max-session-number | **10466**

max-session-setup-rate (Service Set) | **10467**

max-sessions (Dynamic PPPoE) | **10469**

max-sessions-per-connection | **10471**

max-sessions-per-enode | **10472**

max-sessions-per-subscriber | **10474**

max-sessions-per-subscriber (Service Set Next Gen Services) | **10475**

max-sessions (PPPoE Service Name Tables) | **10477**

max-sessions (Security Packet Log) | **10478**

max-sessions (Tunnel Profile) | **10480**

max-sessions-vsa-ignore (Static and Dynamic Subscribers) | **10481**

max-snp-size | **10483**

max-tcp-session-packet-memory | **10484**

max-time-report | **10486**

max-transmit-quality-level | **10487**

max-udp-session-packet-memory | **10489**

max-unknown-messages | **10491**

max-unknown-requests | **10492**

max-withhold-time (Access Profile) | **10494**

maximize-idp-sessions | **10495**

maximum | **10497**

maximum-age | **10498**

maximum-allowed-contentions | **10499**

maximum-backoff-duration | **10501**

maximum-bandwidth | **10502**

maximum-bandwidth (Protocols MPLS) | **10504**

maximum-bit-rate (Default Local Policy) | **10505**

maximum-bit-rate (PCC Action Profiles) | **10507**

maximum-bit-rate (TDF Domain) | **10509**

maximum-call-duration (Security) | **10510**

maximum-capture-size (Datapath Debug) | **10512**

maximum-certificates | **10513**

maximum-connections | **10514**

maximum-connections-per-client | **10516**

maximum-contexts | **10517**

maximum-discovery-table-entries | **10519**

maximum-ecmp | **10520**

maximum-helper-recovery-time | **10522**

maximum-helper-restart-time | **10524**

maximum-helper-restart-time (RSVP) | **10525**

maximum-hop-count | **10527**

maximum-labels | **10528**

maximum-lease-time (DHCP) | **10530**

maximum-length (Origin Validation for BGP) | **10531**

maximum-links | **10532**

maximum-message-length | **10534**

maximum (MSDP Active Source Messages) | **10536**

maximum-neighbor-reconnect-time | **10537**

maximum-neighbor-recovery-time | **10538**

maximum-packet-length | **10540**

maximum-paths | **10542**

maximum-pending-reqs-limit | **10544**

maximum-pending-requests (Diameter) | **10546**

maximum (PIM Entries) | **10547**

maximum-power (Interface) | **10549**

maximum-power (PoE) | **10552**

maximum-prefixes | **10554**

maximum-rps | **10556**

maximum-send-window | **10558**

maximum-sessions | **10559**

maximum-sessions (L2TP) | **10560**

maximum-sessions-per-connection | **10562**

maximum-sessions-per-tunnel | **10563**

maximum-sessions (TDF Gateway) | **10565**

maximum-sessions-trap-percentage (TDF Gateway) | **10566**

maximum-subscribers | **10568**

maximum-subscribers | **10569**

maximum-transactions | **10570**

maximum-transmit-rate (Protocols IGMP) | **10572**

maximum-transmit-rate (Protocols MLD) | **10574**

mc-ae | **10575**

mc-ae-id | **10581**

mcast stats-enable | **10582**

mclag | **10584**

mclag-arp-nd-sync | **10585**

mclag-arpreq-sync | **10586**

mdi-mode | **10587**

mdt | **10590**

measurement-interval | **10591**

measurement-type | **10593**

med-igp-update-interval | **10594**

media-loss-rate | **10596**

media-rate-variation | **10597**

media-source-port-any | **10599**

media-type (Dual-Purpose Uplink Ports) | **10600**

media-type (Interfaces) | **10602**

medium (Tunnel Profile) | **10603**

member | **10604**

member (Application Identification) | **10606**

member-failure-options (Aggregated Multiservices) | **10608**

member-id (Junos Fusion Satellite Device Cluster) | **10611**

member-interface (Aggregated Multiservices) | **10612**

member-link-scheduler | **10614**

member (MX Series Virtual Chassis) | **10615**

member (NSSU Upgrade Groups) | **10617**

member-range | **10618**

member (Redundant Power System) | **10620**

member (Security Group VPN) | **10621**

member (Security IDP) | **10624**

members | **10625**

memory | **10628**

memory-enhanced | **10629**

memory (System) | **10630**

memory (TDF Gateway) | **10633**

menu-item | **10634**

mep | **10639**

mesh-group (Protocols VPLS) | **10641**

message-flood (Security H323) | **10644**

message-flood (Security MGCP) | **10645**

message-ie-profile-v1 | **10647**

message-ie-profile-v2 | **10648**

message-list | **10650**

message-processing-model | **10652**

message-rate-limit (interfaces) | **10653**

message-rate-limit | **10655**

message-rate-limit (Flow Monitoring Logs for NAT) | **10656**

message-size | **10658**

message-type | **10659**

saml (Access Profile) | **10662**

method | **10663**

metric | **10665**

metric (Aggregate, Generated, or Static Route) | **10667**

metric (Diameter Route) | **10668**

metric (Dynamic Access-Internal Routes) | **10670**

metric-in (Protocols RIP) | **10671**

metric-in (Protocols RIPng) | **10673**

metric-out (Protocols RIP) | **10674**

metric-out | **10676**

metric-out (Protocols RIPng) | **10679**

metric (Policy Action) | **10681**

metric (Protocols DVMRP) | **10682**

metric (Protocols IS-IS) | **10684**

metric (Protocols MPLS) | **10686**

metric (Protocols OSPF Sham Link) | **10688**

metric (Qualified Next Hop on Static Route) | **10689**

metrics-profile | **10691**

mgcp | **10693**

mib-profile | **10695**

micro-sid | **10697**

microloop-avoidance | **10700**

microloop-avoidance (Protocols OSPF) | **10702**

mif (TDF Interface) | **10703**

mime-pattern | **10705**

mime-whitelist | **10706**

min-advertisement-interval (Dynamic Router Advertisement) | **10708**

min-advertisement-interval (Protocols IPv6 Neighbor Discovery) | **10709**

min-advertisement-interval (Protocols Router Discovery) | **10711**

min-message-length | **10713**

min-rate | **10714**

min-rate (source-active-advertisement) | **10718**

mine-dhcp-client-options | **10719**

mine-dhcpv6-client-options | **10721**

minimum-bandwidth | **10722**

minimum-bandwidth-adjust-interval | **10724**

minimum-bandwidth-adjust-threshold-change | **10725**

minimum-bandwidth-adjust-threshold-value | **10726**

minimum-bandwidth (aggregated Ethernet) | **10728**

minimum-effective-aigp | **10730**

minimum-hold-time | **10731**

minimum-interval | **10733**

minimum-interval (BFD Liveness Detection) | **10735**

minimum-interval (Liveness Detection) | **10737**

minimum-interval (PIM BFD Liveness Detection) | **10738**

minimum-interval (PIM BFD Transmit Interval) | **10739**

minimum-links | **10741**

minimum-links (candidate-uplink-port-profile) | **10744**

minimum-links (Interfaces) | **10746**

minimum-priority | **10747**

minimum-quality | **10748**

minimum-receive-interval (Protocols PIM) | **10751**

minimum-receive-interval | **10752**

minimum-receive-interval (BFD Liveness Detection) | **10754**

minimum-receive-interval (Liveness Detection) | **10756**

minimum-retransmission-timeout (L2TP Tunnel) | **10757**

minimum-wait-time | **10758**

minor | **10760**

minor-event (Process) | **10761**

minor-event (System) | **10763**

mip-half-function | **10765**

mirror-decrypt-traffic | **10766**

mirror-filter (Security Forwarding Options) | **10768**

mirror-flash-on-disk | **10771**

mirror-once | **10772**

mixed-rate-mode | **10774**

mka | **10775**

mka (MX Series) | **10777**

mld | **10778**

mld (Dynamic Profiles) | **10781**

mld-snooping | **10783**

mldp-inband-signalling (Protocols Multipoint LDP) | **10787**

mlfr-uni-nni-bundle-options | **10789**

mlfr-uni-nni-bundles-inline | **10790**

mobile-edge session-scale | **10792**

mobility | **10793**

mode (RFC 2544 Benchmarking Tests) | **10795**

mode (ike) | **10797**

mode | **10798**

mode (Dynamic Profiles) | **10801**

mode (IKE) | **10802**

mode (Interfaces) | **10804**

mode (IPsec) | **10805**

mode (Traffic Manager) | **10807**

mode (Multicast VLAN Registration) | **10809**

mode (Next Gen Services Service-Set System Logging) | **10811**

mode (Protocols DVMRP) | **10813**

mode (Protocols MSDP) | **10814**

mode (Protocols PIM) | **10816**

mode (QFX Series) | **10818**

mode (Security Forwarding Options) | **10819**

mode (Security Group VPN) | **10821**

mode (Security Log) | **10823**

modem-options | **10825**

modulation-format | **10826**

mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain) | **10827**

mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain) | **10829**

mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain) | **10830**

mofrr-primary-path-selection-by-routing (Multicast-Only Fast Reroute) | **10832**

monitor-bandwidth | **10834**

monitor-end-point | **10835**

monitor ethernet delay-measurement | **10836**

monitor ethernet loss-measurement | **10843**

monitor ethernet synthetic-loss-measurement | **10850**

monitor interface | **10856**

monitor interface error | **10874**

monitor label-switched-path | **10877**

monitor list | **10881**

monitor list (JDM) | **10883**

monitor mpls delay rsvp | **10884**

monitor mpls loss-delay rsvp | **10890**

monitor mpls loss rsvp | **10896**

monitor (Multinode High Availability) | **10903**

monitor-object | **10906**

monitor pfe traffic interface | **10908**

monitor-profile | **10912**

monitor security flow file | **10914**

monitor security flow filter | **10916**

monitor security flow start | **10918**

monitor security flow stop | **10920**

monitor security packet-drop | **10921**

monitor-session | **10926**

monitor start | **10927**

monitor start (JDM) | **10930**

monitor stop | **10932**

monitor stop (JDM) | **10934**

monitor traffic | **10935**

monitor-timer | **10953**

monitoring | **10954**

monitoring (Forwarding Options) | **10956**

monitoring-key (PCC Action Profile) | **10957**

monitoring-profile | **10959**

monitoring (Services) | **10961**

most-fill | **10967**

moving-average-size | **10967**

mpls | **10969**

mpls-flow-table-size | **10972**

mpls (security) | **10974**

mpls (forwarding-options) | **10975**

mpls-inet6 | **10977**

mpls-internet-multicast | **10978**

mpls-ipv4-template | **10979**

mpls-ipvx-template | **10981**

mpls (packet-forwarding-options) | **10982**

mpls (Protocols) | **10984**

mpls (Security Forwarding Options) | **10985**

mpls-template | **10986**

mpls-tp-mode | **10988**

mrru | **10989**

mru (Gigabit Ethernet Options) | **10991**

mru (Congestion Notification Profiles) | **10992**

mru (Dynamic and Static PPPoE) | **10994**

msdp | **10996**

msrpc | **10999**

mss (IDS MS-DPC) | **11001**

mss (Security IDP) | **11002**

msti | **11003**

mstp | **11006**

mtrace | **11010**

mtrace from-source | **11013**

mtrace monitor | **11017**

mtrace to-gateway | **11019**

mtu (QFX Series) | **11023**

mtu (Multilink and Link Services Logical Interface) | **11025**

mtu | **11026**

mtu (interfaces) | **11027**

mtu-discovery | **11032**

mtu (Dynamic and Static PPPoE) | **11034**

mtu (Dynamic Profiles) | **11035**

mtu-signaling | **11037**

mtu (TDF Interface) | **11038**

multi-address-embedded-option-response (DHCP Local Server) | **11039**

multi-bgp-path | **11041**

multi-chassis | **11042**

multi-chassis-protection | **11044**

multi-chassis-protection (Interfaces) | **11045**

multi-destination (Class of Service) | **11047**

multi-domain | **11048**

multi-homing (VPLS Multihoming for BGP VPLS) | **11051**

multi-homing (VPLS Multihoming for FEC 129) | **11052**

multi-link-layer-2-inline | **11054**

multi-sa | **11055**

multi-tenant-hash | **11057**

multi-tenant-support | **11059**

multi-bgp-path | **11060**

multicast | **11062**

multicast (routing-options) | **11063**

multicast-dlci | **11066**

multicast (Dynamic Routing Options) | **11067**

multicast-group | **11069**

multicast-ipv4 | **11070**

multicast-loadbalance | **11072**

multicast-mode (EVPN) | **11074**

multicast-only | **11076**

multicast (Protocols Router Discovery) | **11077**

multicast (QFabric Routing Options) | **11078**

multicast-replication | **11080**

multicast-router-interface (IGMP Snooping) | **11084**

multicast-router-interface (MLD Snooping) | **11086**

multicast-snooping-options | **11087**

multicast-statistics | **11089**

multicast-statistics (packet-forwarding-options) | **11091**

multicast (Virtual Tunnel in Routing Instances) | **11092**

multicast-vlan-flood | **11094**

multichassis-lag-replicate-state | **11095**

multichunk-inspection | **11097**

multihop | **11098**

multilink | **11103**

multilink-class | **11104**

multilink-max-classes | **11106**

multipath-build-priority | **11107**

multipath (OSPF SPF Options) | **11109**

multipath (Protocols BGP) | **11110**

multipath (Routing Options) | **11113**

multipath-rule | **11115**

multipath (SPF Options) | **11117**

multiple-triggered-joins | **11119**

multiplier (Protocols PIM) | **11120**

multiplier | **11122**

multiplier (BFD Liveness Detection) | **11123**

multiplier (Liveness Detection) | **11125**

multipoint | **11127**

multiservice | **11128**

multiservice-options (Flow Monitoring) | **11130**

multiservice-options | **11131**

must-secure | **11133**

mutual-authentication (REST API) | **11134**

mvpn | **11135**

mvpn (Draft-Rosen MVPN) | **11138**

mvpn-iana-rt-import | **11140**

mvpn-mode | **11142**

mvpn (NG-MVPN) | **11143**

mvrp | **11145**

N | 11149

n-plus-n (Power Management) | **11160**

n-plus-n (satellite-management) | **11162**

n391 | **11163**

n392 | **11164**

n393 | **11166**

name (SNMP) | **11167**

name (early-aging) | **11168**

name-format (Flow Collection) | **11170**

name-format | **11172**

name (L2TP Destination) | **11173**

name (L2TP Tunnel Destination) | **11175**

name (Next Gen Services Global System Logging) | **11176**

name-server (Access) | **11178**

name-server (System Services) | **11179**

name (Transport-class) | **11181**

nas-ip-address | **11182**

nas-port-extended-format | **11184**

nas-port-extended-format (Interfaces) | **11186**

nas-port-id-format (Subscriber Management) | **11188**

nas-port-method (L2TP LAC) | **11191**

nas-port-method (Tunnel Profile) | **11192**

nas-port-options (RADIUS Options) | **11193**

nas-port-type | **11195**

nas-port-type (RADIUS Options) | **11197**

nas-port-type (Subscriber Management) | **11199**

nasreq (Diameter Application) | **11202**

nat | **11205**

nat-cone-binding | **11210**

nat-destination-pool | **11212**

nat-destination-rule | **11214**

nat-interface-port-ol (System) | **11216**

nat-keepalive (Services IPsec VPN) | **11218**

nat-nopat-address | **11219**

nat-options | **11221**

nat-options (Next Gen Services) | **11222**

nat-pat-address | **11223**

nat-pat-portnum | **11225**

nat-port-ol-ipnumber | **11227**

nat-rule-referenced-prefix (System) | **11229**

nat-rule-sets (Service Set) | **11230**

nat-rule-sets (Service Set Next Gen Services) | **11231**

nat-rules | **11232**

nat-source-pool | **11234**

nat-source-rule | **11236**

nat-static-rule | **11238**

native-vlan-id | **11240**

natt-install-interval | **11242**

ncp-max-conf-req | **11243**

ncp-restart-timer | **11245**

nd-maxucast-retry | **11246**

nd-override-preferred-src | **11248**

nd-retransmit-timer | **11249**

nd-system-cache-limit | **11251**

nd6-max-cache | **11252**

nd6-new-hold-limit | **11254**

nd6-stale-time | **11256**

ndp | **11257**

ndp-proxy | **11259**

negate | **11260**

negotiate-address | **11261**

negotiation-options | **11263**

neighbor | **11264**

neighbor (Define ANCP) | **11266**

neighbor-discovery | **11268**

neighbor-discovery-inspection | **11270**

neighbor-policy | **11272**

neighbor (Protocols BGP) | **11273**

neighbor (Protocols Layer 2 Circuit) | **11278**

neighbor (Protocols LDP) | **11281**

neighbor (Protocols RIPng) | **11282**

neighbor (Protocols VPLS) | **11284**

nested-application (Security IDP) | **11286**

nested-application-unknown | **11287**

nested-applications | **11288**

netconf | **11290**

netconf-monitoring (NETCONF) | **11293**

netconf (Remote Device Management) | **11295**

network | **11298**

network (Access) | **11300**

network (Address Pools) | **11301**

network-domain | **11302**

network-element (AAA Profile) | **11304**

network-element (Diameter Base Protocol) | **11305**

network-element (Diameter Base Protocol) | **11307**

network-element (Subscriber Aware Policy Control) | **11308**

network-elements (RADIUS) | **11310**

network-isolation | **11312**

network-isolation-profile | **11314**

network-management | **11317**

network-option | **11318**

network-services | **11320**

network-slices | **11322**

network-slicing | **11323**

network-summary-export | **11325**

network-summary-import | **11326**

network (TDF Domain) | **11328**

next-hop (Access) | **11329**

next-hop (Access Internal) | **11330**

next-hop-based-tunnel (Routing Options) | **11332**

next-hop (Class-Of-Service) | **11333**

next-hop (Dynamic Access Routes) | **11334**

next-hop (Forwarding Options) | **11336**

next-hop-group | **11337**

next-hop-group (Analyzer) | **11339**

next-hop-group (Forwarding Options) | **11340**

next-hop-group (Port Mirroring) | **11343**

next-hop-group (Unknown Unicast Forwarding) | **11344**

next-hop-map | **11346**

next-hop (PIM RPF Selection) | **11347**

next-hop (Protocols MPLS) | **11349**

next-hop (RPM) | **11350**

next-hop-service | **11352**

next-hop (Static MAC-Based VLANs) | **11353**

next-hop-tunnel | **11355**

next-hop (VXLAN Routing) | **11356**

next-server | **11358**

next-table | **11359**

nextboot | **11361**

nexthop-hold-time | **11362**

nexthop-learning | **11363**

ng-juniper | **11366**

ngfw | **11369**

no-accounting | **11371**

no-adaptation | **11373**

no-adaptation (BFD Liveness Detection) | **11374**

no-adaptation (Liveness Detection) | **11376**

no-adaptation (PIM BFD Liveness Detection) | **11377**

no-adaptive | **11378**

no-adjacency-down-notification (Protocols IS-IS) | **11380**

no-advertise-adjacency-segment (Protocols IS-IS) | **11382**

no-advertise-adjacency-segment (Protocols OSPF) | **11383**

no advertise-peer-as | **11385**

no-aggregate-delegate-processing | **11386**

no-aggregator-id | **11387**

no-allow-link-events | **11389**

no-allow-snooped-clients | **11390**

no-allow-url (Op Scripts) | **11392**

no-allowed-mac-log | **11393**

no-anti-replay | **11395**

no-anti-replay (Services Service Set) | **11396**

no-application-identification (Services) | **11397**

no-application-system-cache | **11399**

no-application-system-cache (Services) | **11400**

no-attribute-length-in-pdu | **11401**

no-authentication-check (Protocols RSVP) | **11403**

no-auto-conversion | **11404**

no-auto-failover | **11406**

no-auto-iccp-provisioning (Junos Fusion Redundancy Group) | **11407**

no-auto-mdix | **11409**

no-auto-virtual-gateway-esi | **11410**

no-autostart | **11412**

no-autoupdate | **11413**

no-bfd-triggered-local-repair | **11414**

no-bidirectional-mode | **11416**

no-bind-on-request (DHCP Relay Agent) | **11418**

no-broadcast | **11420**

no-bundle-flap | **11422**

no-certificate-chain-in-ike | **11423**

no-client-reflect | **11424**

no-control-word | **11426**

no-control-word (Protocols Layer 2 VPN) | **11427**

no-cspf | **11429**

no-cspf (Protocols RSVP) | **11431**

no-decrement-ttl | **11432**

no-default-gateway-ext-comm | **11434**

no-delegate-processing | **11435**

no-dhcp-snooping | **11436**

no-dhcp-trusted | **11438**

no-dhcpv6-options | **11440**

no-dhcpv6-snooping | **11441**

no-dns-propagation | **11442**

no-domain-vpn-tag | **11443**

no-dr-flood (PIM Snooping) | **11445**

no-dynamic-vlan | **11446**

no-encryption (MACsec) | **11447**

no-encryption (MACsec for MX Series) | **11449**

no-enhanced-frr-bypass (Protocols RSVP) | **11451**

no-examine-dhcpv6 | **11452**

no-external-export | **11454**

no-fabric-wwn-verify | **11456**

no-fcoe-lag | **11457**

no-fcoe-trusted | **11458**

no-filter-check | **11460**

no-fip-snooping-scaling | **11461**

no-flow-logging (DDoS Flow Detection) | **11463**

no-forwarding | **11464**

no-fragmentation | **11466**

no-gratuitous-arp-request | **11467**

no-hidden-commands | **11469**

no-hierarchical-ecmp | **11470**

no-icmp-packet-too-big | **11472**

no-incremental-update | **11473**

no-inner-payload | **11474**

no-install | **11475**

no-install-to-address | **11476**

no-intelligent-prescreening | **11477**

no-interface-hello | **11479**

no-ip-tos-rewrite | **11480**

no-ipsec-tunnel-in-traceroute | **11481**

no-ipv4-routing | **11483**

no-ipv6-routing | **11484**

no-irb-layer-2-copy | **11486**

no-issu-timer-negotiation | **11487**

no-keepalives | **11488**

no-keepalives (Dynamic Profiles) | **11490**

no-l2ckt | **11491**

no-l2vpn | **11492**

no-listen | **11494**

no-load-balance-label-capability | **11495**

no-load-throttle (Resource Monitor) | **11496**

no-local-reversion | **11497**

no-local-switching | **11499**

no-local-switching (VLAN) | **11500**

no-local-switching (VPLS) | **11502**

no-logging (Resource Monitor) | **11503**

no-mac-learning | **11504**

no-make-before-break | **11508**

no-management-vlan | **11509**

no-mcast-replication | **11510**

no-multicast | **11512**

no-multicast-echo | **11514**

no-nat-traversal (Services IPsec VPN) | **11515**

no-native-vlan-insert | **11517**

no-neighbor-down-notification | **11519**

no-neighbor-learn | **11520**

no-next-header-to-payload-protocol-mapping | **11522**

no-node-hello-on-bypass | **11523**

no-node-id-subobject | **11524**

no-normalization | **11526**

no-notify-mail-recipient | **11528**

no-notify-mail-sender (Security Content Filtering Notification Options) | **11529**

no-notify-mail-sender (Security Fallback Block) | **11530**

no-notify-mail-sender (Security Virus Detection) | **11531**

no-nssa-abr | **11533**

no-odu-backward-frr-enable | **11534**

no-odu-signal-degrade-monitor-enable | **11536**

no-option16 | **11537**

no-option18 | **11538**

no-option37 | **11540**

no-option82 | **11541**

no-overload-on-prefix-export-limit | **11542**

no-p2mp-sublsp | **11544**

no-per-unit-scheduler | **11545**

no-ping-record-route | **11546**

no-ping-time-stamp | **11547**

no-policy-cold-synchronization | **11548**

no-power-budget | **11550**

no-pre-classifier | **11551**

no-propagate-ttl | **11553**

no-qos-adjust | **11554**

no-qos-adjust (Dynamic Routing Options) | **11556**

no-recommendation-tlv | **11557**

no-recommended | **11558**

no-redirects | **11559**

no-redirects (IPv4 Traffic) | **11560**

no-registered-multicast | **11562**

no-remote-trace (Next Gen Services Global System Logging) | **11564**

no-remote-trace (System) | **11565**

no-remote-trace (Trace Options) | **11566**

no-revert (Local Switching) | **11568**

no-revert (Neighbor Interface) | **11569**

no-rfc-1583 | **11571**

no-root-port | **11572**

no-route-localize | **11574**

no-sbl-default-server | **11575**

no-send-pads-ac-info | **11577**

no-send-pads-error | **11578**

no-send-to-ue | **11579**

no-sequence-check | **11580**

no-snoop (DHCP Local Server and Relay Agent) | **11582**

no-source-packet-routing (Protocols OSPF) | **11583**

no-split-detection | **11585**

no-split-detection (MX Series Virtual Chassis) | **11586**

no-syn-check | **11588**

no-syn-check-in-tunnel | **11589**

no-syslog | **11590**

no-syslog-generation | **11591**

no-tag | **11592**

no-tag | **11594**

no-temperature-monitoring | **11595**

no-termination-request | **11596**

no-throttle (Resource Monitor) | **11598**

no-tls-certificate-check | **11599**

no-transit-statistics | **11600**

no-translation | **11601**

no-translation (Source NAT Next Gen Services) | **11603**

no-trap | **11604**

no-tunnel-services | **11605**

no-unknown-unicast | **11607**

no-unregistered-multicast | **11609**

no-unsolicited-ra (Enhanced Subscriber Management) | **11611**

no-uri-check | **11612**

no-user-query (Services User Identification) | **11614**

no-validate | **11615**

no-vlan-interface-name | **11617**

no-vrf-advertise | **11620**

no-vrf-propagate-ttl | **11621**

no-walkup | **11623**

no-world-readable (Next Gen Services Global System Logging) | **11624**

no-zero-suppression | **11626**

no-dhcp-flood | **11627**

no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping) | **11628**

no-incremental-update | **11630**

node (Chassis Cluster Redundancy Group) | **11631**

node-detection (EVPN-VXLAN) | **11633**

node-device (Aliases) | **11635**

node-device (Chassis) | **11637**

node-device (Flow Groups) | **11638**

node-device (Resources) | **11640**

node-group (Chassis) | **11641**

node-group (OxID Hash Control) | **11643**

node-group (Resources) | **11645**

node-hello | **11646**

node-id | **11648**

node-link-protection (Protocols MPLS) | **11649**

node-protection | **11651**

node-protection (Static LSP) | **11652**

node-segment (Protocols OSPF) | **11654**

node (System) | **11656**

non-active-hold-priority | **11657**

non-lsp-next-hop | **11659**

non-revertive | **11660**

non-revertive (Chassis) | **11661**

non-revertive (Interfaces) | **11662**

non-strict-priority-scheduling | **11664**

non-subscriber-no-reply | **11665**

non-vc-mode | **11666**

nonpersistent | **11668**

nonstop-bridging | **11669**

nonstop-bridging (Ethernet Switching) | **11670**

nonstop-routing | **11671**

nonstop-routing-options | **11673**

nonvolatile | **11675**

normalization | **11676**

normalization (Protocols EVPN) | **11678**

not (Event Policy) | **11680**

not-on-disk-underperform | **11681**

not-present (DHCP Relay Agent) | **11682**

notification | **11686**

notification (Commit) | **11687**

notification-control | **11688**

notification-interval | **11690**

notification (NETCONF) | **11691**

notification-options (Security Antivirus) | **11693**

notification-options (Security Content Filtering) | **11695**

notification (System Services) | **11697**

notification-targets | **11699**

notification-time | **11700**

notify | **11701**

notify-filter (Applying to the Management Target) | **11703**

notify-filter (Configuring the Profile Name) | **11704**

notify-mail-recipient | **11705**

notify-mail-sender (Security Content Filtering Notification Options) | **11707**

notify-mail-sender (Security Fallback Block) | **11708**

notify-mail-sender (Security Virus Detection) | **11709**

notify-view | **11710**

np-cache (Flexible PIC Concentrator) | **11711**

nptv6-prefix | **11713**

nptv6-prefix-name | **11714**

nsr-phantom-holdtime | **11715**

nssa | **11717**

nssu | **11719**

ntp | **11721**

ntp (QFabric) | **11729**

ntp threshold | **11731**

nullpdu | **11733**

num-65-127-prefix | **11734**

number-of-frames | **11736**

number-of-ports | **11737**

number-of-sub-ports (interface) | **11740**

O | 11741

oam (Layer 2 Circuit Protocols) | **11748**

oam (MX Series) | **11749**

oam | **11753**

oam detection-disable | **11756**

oam-on-svlan (Ethernet Interfaces) | **11758**

oam (Protocols LDP) | **11759**

oam (Protocols MPLS) | **11762**

oam runtime-disable | **11764**

object-cache-size | **11765**

object-names | **11767**

observation-domain-id | **11768**

oc192 | **11770**

ocs (Diameter Applications) | **11771**

ocsp (Security PKI) | **11774**

odu-backward-frr-enable | **11776**

odu-delay-management | **11777**

odu-signal-degrade | **11779**

odu-signal-degrade-monitor-enable | **11781**

odu-ttim-action-enable | **11782**

off (Destination NAT Next Gen Services) | **11784**

offer-period | **11785**

offline-on-fabric-bandwidth-reduction | **11786**

offload-type | **11788**

offloads | **11789**

offset | **11791**

offset (MX Series) | **11793**

oid | **11796**

oid (SNMP Scripts) | **11797**

oid (SNMPv3) | **11799**

oif-map (Dynamic IGMP Interface) | **11800**

oif-map (Dynamic MLD Interface) | **11801**

oif-map (IGMP Interface) | **11803**

oif-map (MLD Interface) | **11804**

oism | **11805**

oism (Multicast Snooping Options) | **11807**

omit-no-export (Graceful Restart for BGP Helper) | **11810**

omit-wildcard-address | **11812**

on-demand-address-allocation | **11813**

on-demand-ip-address | **11815**

on-disk-failure | **11817**

on-disk-failure (Chassis Routing Engine) | **11818**

on-error | **11819**

on-link | **11821**

on-link (Dynamic Router Advertisement) | **11823**

on-loss-of-keepalives | **11824**

on-loss-of-keepalives | **11826**

on-re-to-fpc-stale | **11827**

one-to-one (Aggregated Multiservices) | **11829**

one-way-hardware-timestamp | **11830**

onlink-subnet-only | **11832**

op | **11833**

op invoke-debugger cli | **11836**

op (Op Scripts) | **11838**

open-timeout | **11841**

openflow (Protocols OpenFlow) | **11842**

operating-mode | **11844**

operation | **11846**

optics-options | **11847**

optimal-route-reflection | **11850**

optimize-adaptive-teardown | **11852**

optimize-aggregate-frr | **11853**

optimize-aggressive | **11855**

optimize-hold-dead-delay | **11856**

optimize-switchover-delay | **11857**

optimize-timer (Protocols MPLS) | **11859**

optimize-timer (Protocols RSVP) | **11861**

option-15 (DHCP Relay Agent) | **11862**

option-16 (DHCP Relay Agent) | **11865**

option-16 (DHCPv6 Snooping) | **11868**

option-18 (DHCPv6 Snooping) | **11869**

option-18 (Interface-ID for DHCPv6 Autosense VLANs) | **11871**

option-37 (DHCPv6 Snooping) | **11872**

option-37 (Relay Agent Remote-ID for DHCPv6 Autosense VLANs) | **11874**

option-60 (DHCP Local Server) | **11876**

option-60 (DHCP Relay Agent) | **11877**

option-77 (DHCP Relay Agent) | **11880**

option-79 | **11883**

option-82 | **11884**

option-82 (interfaces) | **11886**

option-82 (Address-Assignment Pools) | **11888**

option-82 (DHCP Local Server Authentication) | **11890**

option-82 (DHCP Local Server Pool Matching) | **11892**

option-82 (DHCP Relay Agent) | **11893**

option (DHCP server) | **11895**

option-match | **11896**

option-number (DHCP Relay Agent Option) | **11898**

option-order (DHCP Relay Agent) | **11899**

option-refresh-rate | **11902**

option-refresh-rate (Services) | **11903**

option (Security IDP) | **11905**

option-type | **11906**

optional (Commit Scripts) | **11907**

optional (Event Scripts) | **11909**

optional-parameters | **11910**

options (Access Profile) | **11912**

options (Routing Options) | **11922**

options-template-id | **11924**

order | **11926**

order (Application Identification) | **11927**

order-priority (Application Identification) | **11929**

order (Security IDP) | **11930**

origin-autonomous-system (Origin Validation for BGP) | **11932**

origin (Diameter Base Protocol) | **11934**

origin (Diameter Base Protocol) | **11935**

origin-validation | **11937**

ospf | **11938**

ospf-instance | **11940**

ospf3 | **11942**

oss-map | **11944**

other | **11945**

other-bytes | **11949**

other-overhead-adjust | **11951**

other-stateful-configuration | **11953**

other-stateful-configuration (Dynamic Router Advertisement) | **11954**

otn-options | **11955**

otu-ttim-action-enable | **11958**

otu4 | **11960**

out-delay | **11961**

out-of-resources | **11963**

out-of-resources (Security Antivirus Sophos Engine) | **11965**

outbound-https | **11966**

outbound-route-filter | **11970**

outbound-ssh | **11972**

outer-tag-protocol-id (RFC 2544 Benchmarking) | **11975**

outgoing-queue | **11976**

output (Service Set) | **11978**

output | **11980**

output (Accounting) | **11981**

output-chain | **11983**

output (Congestion Notification) | **11984**

output (Dynamic Service Sets) | **11986**

output-filename (Event Policy) | **11988**

output-format (Event Policy) | **11989**

output-forwarding-class-map | **11991**

output (Forwarding Table) | **11992**

output-interface (CCC) | **11994**

output-interface-index | **11995**

output interfaces (Chassis Synchronization) | **11996**

output (Mirroring) | **11998**

output (Monitoring) | **12000**

output-policer | **12001**

output (Port Mirroring) | **12003**

output (Port Mirroring) | **12004**

output-priority-map | **12007**

output-queue-priority | **12008**

output (Sampling) | **12010**

output (Security Forwarding Options) | **12013**

output-three-color | **12014**

output-traffic-control-profile | **12016**

output-traffic-control-profile (Dynamic CoS Definition) | **12018**

output-traffic-control-profile-remaining | **12019**

output-vlan-map | **12021**

output-vlan-map (Aggregated Ethernet) | **12022**

output-vlan-map (Dynamic Interfaces) | **12023**

outputs (Analytics) | **12025**

over (Application Identification) | **12028**

over-limit | **12031**

overall-packet (Flat-File Accounting Options) | **12032**

overflow-pool | **12035**

overhead-accounting | **12036**

overhead-accounting (ANCP) | **12038**

overhead-accounting (Dynamic Traffic Shaping) | **12039**

overlay-ecmp | **12041**

overlay (Packet Forwarding Options) | **12042**

overlay-path | **12044**

overload-pool | **12046**

overload-prefix | **12047**

overload (Protocols IS-IS) | **12048**

overload (Protocols OSPF) | **12051**

override | **12054**

override-chap-password | **12055**

override-firewall | **12057**

override-interval | **12058**

override-password (Domain Map) | **12060**

override (PIM Static RP) | **12061**

override (RADIUS Options) | **12063**

overrides (DHCP Local Server) | **12065**

overrides (DHCP Relay Agent) | **12068**

overrides (DHCP Security) | **12071**

overrides (DHCP Security for MX Series) | **12072**

overrides (Enhanced Subscriber Management) | **12074**

overrides (System Services DHCP) | **12077**

oversize-message-drop | **12081**

ovlan-cfi (RFC 2544 Benchmarking) | **12082**

ovlan-id (RFC 2544 Benchmarking) | **12084**

ovlan-priority (RFC 2544 Benchmarking) | **12085**

ovsdb | **12086**

ovsdb-managed | **12088**

owner | **12090**

owner (RMON) | **12092**

oxid | **12093**

P | 12095

p2mp-ldp-next-hop | **12119**

p2mp-lsp-next-hop | **12120**

p2mp (Protocols LDP) | **12122**

p2mp (Protocols MPLS) | **12124**

p2mp-receive-switch | **12125**

p2mp-transmit-switch | **12127**

package (Loading on PIC) | **12128**

packet-action | **12129**

packet-burst (Hierarchical Policer) | **12133**

packet-burst (Policer) | **12135**

packet-capture | **12137**

packet-filter | **12141**

packet-format-match | **12143**

packet-forwarding-options | **12145**

packet-log (Security Flow) | **12150**

packet-log (Security IDP Policy) | **12152**

packet-log (Security IDP Sensor Configuration) | **12153**

packet-loss-priority (RFC 2544 Benchmarking) | **12155**

packet-ordering-mode (Application Services) | **12157**

packet-scheduling | **12158**

packet-size (RFC 2544 Benchmarking) | **12160**

packet-timestamp | **12162**

packet-types (Dynamic VLAN Authentication) | **12163**

packet-timestamp | **12165**

pado-advertise | **12166**

pair-policy | **12168**

pap | **12169**

pap (Dynamic PPP) | **12171**

pap (L2TP) | **12172**

parameter-preference | **12173**

parameters | **12175**

parse-direction (Access Profile) | **12176**

parse-direction (Domain Map) | **12178**

parse-order (Domain Map) | **12179**

partition | **12181**

partition (Gx-Plus) | **12182**

partition (NASREQ Diameter Application) | **12183**

partition (OCS) | **12185**

partition (PCRF) | **12187**

partition (s6a) | **12190**

pass-through (interfaces) | **12192**

pass-through (Security Policy) | **12194**

pass-through (Firewall-authentication) | **12196**

passive (CHAP) | **12198**

passive (Dynamic IGMP Interface) | **12199**

passive (Dynamic MLD Interface) | **12201**

passive (IGMP) | **12202**

passive (MLD) | **12204**

passive-mode | **12205**

passive-mode-tunneling | **12207**

passive-mode-tunneling (Security IPsec) | **12208**

passive-monitor-mode | **12212**

passive (MVRP) | **12213**

passive (PAP) | **12214**

passive-probe-params | **12216**

passive (Protocols BGP) | **12218**

passive (Protocols IS-IS) | **12219**

passive (Protocols Link Management) | **12222**

passive (Protocols OSPF) | **12223**

passive (Routing Options) | **12225**

password (Access) | **12225**

password (DHCP Local Server) | **12226**

password (DHCP Relay Agent) | **12229**

password-file | **12231**

password (Flow Collector File Servers) | **12232**

password (Interfaces) | **12233**

password (Login) | **12235**

password (MAC RADIUS Authentication) | **12241**

password-options | **12243**

password (Security Antivirus) | **12245**

password (Services) | **12246**

password (Static Subscribers) | **12247**

password (Transfer Log File Servers) | **12249**

path | **12250**

path-database-size | **12252**

path-database-size (EX Series Switch Only) | **12253**

path-length | **12255**

path-mtu | **12256**

path-mtu-discovery (Tunnel) | **12257**

path (Protocols MPLS) | **12259**

path (Protocols RSVP) | **12261**

path-rate-limit | **12263**

path-selection | **12266**

path (Steering) | **12269**

pattern (Application Identification) | **12271**

pattern (Class Attribute) | **12272**

pattern-pcre (Security IDP) | **12274**

pattern (Security IDP) | **12275**

pattern-update (Security Antivirus) | **12276**

pause-computation-during-churn | **12277**

payload | **12279**

pba-interim-logging-interval | **12280**

pbb-evpn-core | **12282**

pbb-service-options | **12283**

pcc-action-profile (PCC Rules) | **12284**

pcc-action-profiles | **12286**

pcc-context | **12289**

pcc-rule | **12291**

pcc-rulebases (PCEF) | **12293**

pcc-rulebases (PCEF Profile) | **12294**

pcc-rules (PCEF) | **12297**

pcc-rules (PCEF Profile) | **12299**

pcc-time-of-day-profiles | **12301**

pce | **12303**

pce-group (PCE) | **12306**

pce-group (Protocols PCEP) | **12307**

pce-type | **12309**

pcef | **12311**

pcef (Dynamic Profiles) | **12314**

pcef-profile (Service Set) | **12315**

pcef-profile (TDF Domain) | **12317**

pcef-profile (TDF Domain Selection) | **12318**

pcep | **12320**

pci-address | **12322**

pcp-rules | **12323**

pcp-server | **12324**

pcrf (Diameter Applications) | **12325**

pdu-interval | **12328**

pdu-threshold | **12329**

peak-burst-size (QFX) | **12331**

peak-burst-size | **12332**

peak-information-rate (QFX) | **12334**

peak-information-rate | **12336**

peak-rate | **12338**

peer | **12339**

peer-active (VPLS Multihoming for FEC 129) | **12341**

peer-as-billing-template | **12343**

peer-as-list | **12345**

peer-as (Protocols BGP) | **12346**

peer-as (VPLS) | **12348**

peer-certificate-type | **12350**

peer (Diameter Base Protocol) | **12351**

peer (Diameter Base Protocol) | **12353**

peer (Diameter Network Element) | **12354**

peer (Diameter Network Element) | **12356**

peer-forwarding-ip | **12357**

peer-gnf | **12359**

peer-groups (access network peers) | **12361**

peer-groups (control plane peers) | **12362**

peer-groups (core network peers) | **12364**

peer (ICCP) | **12366**

peer-id | **12367**

peer-interface (Protocols OSPF) | **12370**

peer-interface (Protocols RSVP) | **12372**

peer-ip-address-optional | **12373**

peer-limit | **12375**

peer (Multichassis) | **12376**

peer-origin (Diameter Peer) | **12377**

peer (Protocols LMP) | **12379**

peer (Protocols MSDP) | **12380**

peer (Protocols RIP) | **12382**

peer-selection-service | **12384**

peer-subnet (Easy EVPN LAG Configuration) | **12385**

peer-to-peer (Easy EVPN LAG Configuration) | **12388**

peer-unit | **12390**

peer-unit (Routing Policy Condition) | **12391**

peers (Commit) | **12393**

peers-synchronize | **12394**

peg-df-election | **12395**

pem-absence | **12397**

pem (MX Series Devices, ACX7024 Devices) | **12399**

pem (T640, T1600, and T4000 Routers with Six-Input DC Power Supply) | **12400**

pending-queue-watermark | **12402**

pending-queue-watermark-abate | **12403**

pending-sess-queue-length | **12405**

per-esi | **12406**

per-esi-vlan | **12408**

per-evi | **12409**

per-flow | **12413**

per-flow (Aggregated Ethernet Interfaces) | **12414**

per-group-label | **12416**

per-prefix | **12417**

per-prefix-label | **12419**

per-route-accounting | **12420**

per-session-scheduler | **12422**

per-unit-scheduler | **12423**

perfect-forward-secrecy (Security) | **12424**

perfect-forward-secrecy (Services) | **12426**

performance | **12428**

performance-monitor | **12429**

performance-monitoring | **12431**

performance-monitoring (OAM LFM) | **12432**

performance-monitoring (Protocols MPLS) | **12434**

periodic (Chassis) | **12436**

periodic | **12438**

periodic (Interfaces) | **12439**

periodic-traceroute | **12441**

permit-command | **12443**

permit (Security Persistent NAT) | **12444**

permit (Security Policies) | **12446**

permit (Security SCTP) | **12449**

perpetual-poe | **12450**

persistent-learning | **12451**

persistent-nat | **12453**

persistent-results | **12454**

pfc | **12455**

pfc (Input Congestion Notification) | **12456**

pfc-priority | **12458**

pfe | **12460**

pfe power off | **12463**

pgcp | **12465**

pgcp-rules | **12466**

pgcp-service | **12468**

phase-adjust-threshold | **12469**

phase-lock-threshold | **12471**

phy-timestamping | **12472**

physical-interface (DDoS Flow Detection) | **12474**

physical-interface-filter | **12476**

physical-interface-policer | **12479**

pic | **12480**

pic-boot-timeout | **12484**

pic-console-authentication | **12485**

pic (M Series and T Series Routers) | **12487**

pic-memory-threshold | **12489**

pic-mode | **12490**

pic-mode (MX Series) | **12491**

pic-mode (Chassis T1 Mode) | **12494**

pic (MX Series Routers) | **12496**

pic (Port) | **12498**

pic (PTX Series Routers) | **12500**

pic (satellite-policies port-group-alias) | **12502**

pic-services-logging | **12504**

pic (Subscriber Limits) | **12506**

pim | **12507**

pim-asm | **12514**

pim-evpn-gateway | **12515**

pim-snooping | **12517**

pim-ssm (Provider Tunnel) | **12519**

pim-ssm (Selective Tunnel) | **12520**

pim-to-igmp-proxy | **12522**

pim-to-mld-proxy | **12523**

ping | **12525**

ping clns | **12534**

ping-death | **12537**

ping-death (IDS Screen Next Gen Services) | **12538**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

ping-interval | **12545**

ping mpls bgp | **12547**

ping mpls l2circuit | **12550**

ping mpls l2vpn | **12554**

ping mpls l3vpn | **12558**

ping mpls ldp | **12561**

ping mpls lsp-end-point | **12565**

ping mpls rsvp | **12568**

ping mpls segment routing isis | **12575**

ping mpls segment routing ospf | **12579**

ping mpls segment routing spring-te | **12583**

ping mpls segment routing static egress-ip | **12592**

ping overlay | **12596**

ping srv6 | **12603**

ping vpls instance | **12608**

pki | **12610**

pki (Services) | **12612**

plmn | **12615**

poe | **12616**

point-to-point | **12619**

point-to-point (MVRP) | **12620**

policer | **12622**

policer (Applying to a Logical Interface) | **12624**

policer (CFM Firewall) | **12626**

policer (CFM Global) | **12627**

policer (CFM Session) | **12629**

policer (Configuring) | **12630**

policer (CoS) | **12633**

policer-drop-probability-low | **12635**

policer (Firewall Filter Action) | **12636**

policer (Interface) | **12638**

policer (Layer 2 VPN) | **12639**

policer (MAC) | **12641**

policer-overhead | **12643**

policer-overhead-adjustment | **12644**

policer-charge | **12646**

policer-charge | **12648**

policies | **12650**

policing | **12661**

policing (Protocols LDP) | **12662**

policing (Protocols MPLS) | **12664**

policy | **12665**

policy (advanced-policy-based-routing) | **12667**

policy (Aggregate and Generated Routes) | **12670**

policy-based-logging (LRF Profile) | **12672**

policy-db-size | **12674**

policy (Event Policy) | **12675**

policy (Explicit-Proxy) | **12678**

policy (Flow Maps) | **12683**

policy (IPsec) | **12684**

policy-lookup-cache | **12685**

policy-map (Class of Service) | **12686**

policy-match | **12689**

policy (Multicast-Only Fast Reroute) | **12691**

policy-multipath | **12694**

policy-oids | **12696**

policy-options | **12697**

policy-options (QFX) | **12700**

policy-options (Dynamic Profiles) | **12702**

policy (PIM rpf-vector) | **12703**

policy (Protocols Multipoint LDP) | **12706**

policy-rematch | **12707**

policy (Security Alarms) | **12709**

policy (Security Group VPN IKE) | **12711**

policy (Security IKE) | **12713**

policy (Security Interfaces) | **12716**

policy (Security IPsec) | **12718**

policy (Security IPsec Interfaces) | **12722**

policy (Security Policies) | **12724**

policy (Services CoS Next Gen Services) | **12728**

policy (Services IKE) | **12730**

policy-set | **12731**

policy (SSM Maps) | **12733**

policy (Stateful Firewall Rules Next Gen Services) | **12734**

policy-statement | **12736**

policy-stats | **12743**

policy (Subscriber Secure Policy) | **12745**

policy (System Security Profile) | **12747**

policy-with-count | **12749**

policy-map-marking | **12751**

polling-interval (Network Management) | **12752**

polling-interval | **12753**

pon (Access-Line Rate Adjustment) | **12755**

pool (NAT) | **12760**

pool | **12762**

pool (Access) | **12763**

pool (Address-Assignment Pools) | **12766**

pool-default-port-range | **12768**

pool-default-port-range (Source NAT Next Gen Services) | **12770**

pool-default-twin-port-range | **12771**

pool (Destination NAT Next Gen Services) | **12773**

pool (DHCP Local Server Overrides) | **12774**

pool-match-order | **12777**

pool (NAT Rule Next Gen Services) | **12778**

pool (Security Destination NAT) | **12779**

pool (Security Source NAT) | **12781**

pool (Security Source NAT Rule Set) | **12783**

pool (Service Interface) | **12784**

pool (Service Interfaces) | **12786**

pool (Source NAT Next Gen Services) | **12787**

pool (System) | **12789**

pool (TDF Domain) | **12791**

pool-utilization-alarm | **12793**

pool-utilization-alarm (Security Source NAT Pool) | **12794**

pool-utilization-alarm (Source NAT Next Gen Services) | **12796**

pop (protocols mpls) | **12797**

pop | **12799**

pop-all-labels | **12800**

pop-and-forward (Protocols MPLS) | **12802**

pop-and-forward (Protocols RSVP) | **12803**

pop (Dynamic VLANs) | **12804**

pop-pop | **12806**

pop-swap | **12807**

pop3-profile (Security Content Security Policy Antivirus) | **12808**

pop3-profile (Security Content Security Policy Content Filtering) | **12809**

port | **12811**

port (Access LDAP) | **12812**

port (cflowd) | **12813**

port (Chassis) | **12815**

port (Chassis Synchronization) | **12818**

port (Diameter Peer) | **12819**

port-error-disable | **12821**

port (Flow Monitoring) | **12823**

port-forwarding | **12824**

port-forwarding (Destination NAT Next Gen Services) | **12826**

port-forwarding-mappings | **12827**

port-forwarding-mappings (Destination NAT Rule Next Gen Services) | **12828**

port-group-alias (satellite-policies) | **12829**

port-id | **12831**

port-id (MACsec for MX Series) | **12833**

port (LRF Profile) | **12834**

port-mirroring | **12836**

port-mode | **12841**

port-mode (Fibre Channel Interfaces) | **12843**

port (NETCONF) | **12844**

port (Origin Validation for BGP) | **12846**

port-overloading-factor (Security Source NAT Interface) | **12847**

port-overloading-factor (Security Source NAT Pool) | **12849**

port-overloading (Security Source NAT Interface) | **12850**

port (OVSDB) | **12852**

port (Packet Forwarding) | **12853**

port-priority | **12856**

port (Protocols OpenFlow) | **12857**

port (RADIUS Server) | **12859**

port-randomization | **12860**

port-range | **12861**

port-range (Address Mapping) | **12863**

port-range (Application Identification) | **12865**

port (REST API) | **12866**

port-round-robin | **12867**

port-round-robin (Source NAT Next Gen Services) | **12869**

port (RPM) | **12870**

port (satellite-policies port-group-alias) | **12871**

port-scaling-enlargement | **12873**

port-scan | **12874**

port (Security Antivirus) | **12875**

port (Security Source NAT) | **12877**

port (Security Web Filtering Server) | **12879**

port (Services) | **12880**

port (Services NAT) | **12881**

port (Services Voice) | **12884**

port (Source NAT Next Gen Services) | **12885**

port (SRC Server) | **12887**

port-status-tlv | **12888**

port (Subscriber Limits) | **12890**

port (Syslog) | **12891**

port (System Log Messsages) | **12893**

port (TWAMP) | **12894**

ports (System) | **12896**

ports-per-session | **12897**

post-attack | **12898**

post-attack-timeout | **12900**

post-cli-implicit-firewall | **12901**

post-convergence-lfa | **12903**

post-convergence-lfa (Protocols OSPF) | **12904**

post-service-filter | **12907**

post-service-filter (Dynamic Service Sets) | **12908**

post-policy | **12910**

potential-violation | **12912**

power | **12915**

power | **12918**

power-budget-priority | **12919**

power-mode-disable | **12921**

power-mode-ipsec | **12922**

power-mode-ipsec-qat | **12924**

pp0 (Dynamic PPPoE) | **12925**

ppm | **12927**

ppm (Ethernet Switching) | **12930**

ppp | **12931**

ppp-access-profile | **12944**

ppp (Group Profile) | **12945**

ppp-options | **12949**

ppp-options (Dynamic PPP) | **12951**

ppp-options (L2TP) | **12954**

ppp-over-ether | **12956**

ppp (Profile) | **12957**

ppp-subscriber-services | **12959**

pppoe (System Processes) | **12960**

pppoe (DDoS Protection) | **12962**

pppoe-options (SRX Series) | **12972**

pppoe-options | **12974**

pppoe-options (Dynamic PPPoE) | **12976**

pppoe-tags (Adjustment Control Profiles) | **12977**

pppoe-underlying-options (Dynamic VLAN Interface Sets) | **12979**

pppoe-underlying-options (Static and Dynamic Subscribers) | **12980**

pps-limit (Hierarchical Policer) | **12982**

pps-limit (Policer) | **12984**

pptp | **12986**

prbs | **12987**

prbs-test-start | **12989**

prbs-test-stop | **12991**

pre-attack | **12993**

pre-filter-shellcode | **12994**

pre-id-default-policy | **12995**

pre-ietf-mode | **12998**

pre-rewrite-tos | **13000**

pre-shared-key | **13001**

pre-shared-key (MX Series) | **13003**

pre-shared-key (Security) | **13004**

pre-shared-key (Services IKE) | **13006**

preauthentication-order (Access Profile) | **13007**

precedence | **13008**

precision-timers | **13010**

precision-timers-max-period | **13012**

predefined-attack-groups | **13013**

predefined-attacks | **13015**

predefined-variable-defaults (Dynamic Client Profiles) | **13016**

preempt | **13019**

preempt (Chassis Cluster) | **13021**

preempt-cutover-timer | **13023**

preempt (VRRP) | **13025**

preemption | **13026**

preemptive-fast-reroute | **13028**

prefer-framed-ip-address (RADIUS Clients) | **13030**

prefer-framed-ipv6-prefix (RADIUS Clients) | **13031**

preference (Access) | **13032**

preference (Interface-Level Preference for VPLS Multihoming for FEC 129) | **13033**

preference (IPv6 Router Advertisement) | **13035**

preference (PEG DF Election) | **13037**

preference (Protocols BGP) | **13038**

preference (Protocols LDP) | **13040**

preference (Protocols MPLS) | **13042**

preference (Protocols RIP) | **13043**

preference (Protocols RIPng) | **13045**

preference (Routing Options) | **13046**

preference (Site-Level Preference for VPLS Multihoming for FEC 129) | **13049**

preference (Subscriber Management) | **13050**

preference (Tunnel Profile) | **13052**

preferred | **13053**

preferred-ciphers | **13055**

preferred-lifetime | **13056**

preferred-lifetime (Dynamic Router Advertisement) | **13058**

preferred-prefix-length | **13059**

preferred-source-address | **13060**

prefix | **13062**

prefix-65-127-disable | **13064**

prefix (Access IPv6) | **13067**

prefix-action (Configuring) | **13068**

prefix-action (Firewall Filter Action) | **13070**

prefix (Address-Assignment Pools) | **13071**

prefix-attribute-flags | **13072**

prefix (Circuit ID for Option 82) | **13074**

prefix (DHCP Client) | **13076**

prefix (DHCP Relay Agent) | **13077**

prefix (DHCPv6 Options) | **13079**

prefix (Dynamic Router Advertisement) | **13081**

prefix-export-limit (Protocols IS-IS) | **13083**

prefix-export-limit (Protocols OSPF) | **13085**

prefix-limit | **13087**

prefix-list | **13090**

prefix-list-filter | **13092**

prefix-list (Junos Telemetry Interface) | **13093**

prefix-list-name | **13095**

prefix-list (PIM RPF Selection) | **13097**

prefix-name (Security Static NAT) | **13098**

prefix (Protocols IPv6 Neighbor Discovery) | **13100**

prefix (Remote ID for Option 82) | **13101**

prefix (Security Static NAT) | **13103**

prefix-segment | **13105**

prefix-segment-range | **13107**

prefix-segment (Routing Options) | **13109**

premium (Hierarchical Policer) | **13110**

premium (Output Priority Map) | **13112**

premium (Policer) | **13113**

preprovisioned | **13115**

preprovisioned (MX Series Virtual Chassis) | **13116**

preserve-fpc-poweron-sequence | **13118**

preserve-incoming-fragment-size | **13119**

preserve-interface | **13121**

preserve-nexthop-hierarchy (SRv6 TE) | **13122**

preserve-parity (Source NAT Next Gen Services) | **13124**

preserve-range (Source NAT Next Gen Services) | **13125**

primary | **13126**

primary | **13127**

primary (Adaptive Services Interfaces) | **13129**

primary (Address on Interface) | **13130**

primary connection (Identity Management Advanced Query) | **13131**

primary-data-record-fields | **13134**

primary-dns | **13137**

primary-interface (Aggregated Inline Services) | **13138**

primary (Interface for Router) | **13140**

primary (Link Services IQ PIC Interfaces) | **13142**

primary (MBGP MVPN) | **13143**

primary (Protocols MPLS) | **13144**

primary-server | **13146**

primary (Virtual Tunnel in Routing Instances) | **13148**

primary-vlan | **13150**

primary (VPLS Multihoming) | **13151**

primary-wins | **13153**

priority | **13154**

priority (Application Identification With Next Gen Services) | **13156**

priority (Bootstrap) | **13157**

priority (CFM MEP) | **13159**

priority (Chassis Synchronization Source) | **13160**

priority (Clock Synchronization) | **13162**

priority-cost (VRRP) | **13163**

priority (DDoS) | **13165**

priority (Diameter Network Element) | **13166**

priority (Diameter Peer) | **13168**

priority (Dynamic Schedulers) | **13169**

priority (Fabric Priority) | **13171**

priority (Fabric Queues, Schedulers) | **13172**

priority (FIP) | **13174**

priority-flow-control | **13176**

priority-hold-time | **13177**

priority-hold-time | **13179**

priority (Management Traffic) | **13180**

priority (OAM Connectivity-Fault Management) | **13182**

priority-override (Event Policy) | **13183**

priority (PIM Interfaces) | **13184**

priority (PIM RPs) | **13186**

priority (PoE) | **13188**

priority (policy-options) | **13189**

priority (Power over Ethernet) | **13191**

priority (Protocols MPLS) | **13193**

priority (Protocols OAM) | **13194**

priority (Protocols Router Discovery) | **13196**

priority (Protocols RSVP) | **13197**

priority (Protocols STP) | **13199**

priority (Protocols VRRP) | **13201**

priority (RADIUS Network Elements) | **13203**

priority (Redundant Power System) | **13204**

priority (Schedulers) | **13206**

priority1 | **13209**

priority2 | **13211**

privacy-3des | **13212**

privacy-aes128 | **13213**

privacy-des | **13215**

privacy-key | **13216**

privacy-none | **13217**

privacy-password | **13218**

private-vlan | **13220**

proactive-arp-detection | **13221**

probe (ICMP Router Discovery) | **13222**

probe | **13228**

probe-count | **13231**

probe-interval | **13233**

probe-limit | **13234**

probe-server | **13236**

probe-type | **13238**

process-inform | **13240**

process (Memory) | **13242**

process-monitor | **13245**

process-non-null-as-null-register | **13246**

processes | **13248**

product-model (Junos Fusion) | **13250**

product-name | **13252**

products | **13253**

profile (Application Identification) | **13254**

profile | **13256**

profile | **13257**

profile (Access) | **13259**

profile (Access) | **13263**

profile (Application Firewall) | **13271**

profile (Captive Portal Content Delivery) | **13273**

profile(dynamic-application) | **13275**

profile (Explicit Proxy) | **13277**

profile (HTTP Header Enrichment) | **13282**

profile (icap-redirect) | **13283**

profile (Juniper Secure Connect) | **13286**

profile (LRF) | **13289**

profile (Rule Sets) | **13291**

profile (security) | **13292**

profile (Security Antispam SBL) | **13310**

profile (Security Antivirus Juniper Express Engine) | **13312**

profile (Security Antivirus Kaspersky Lab Engine) | **13314**

profile (Security Content Filtering) | **13317**

profile (Security GTP) | **13318**

profile (Security SCTP) | **13323**

profile (Security Sophos Engine Antivirus) | **13326**

profile (Security Web Filtering Juniper Enhanced) | **13328**

profile (Security Web Filtering Juniper Local) | **13330**

profile (Security Web Filtering Surf Control Integrated) | **13332**

profile (Security Web Filtering Websense Redirect) | **13333**

profile (Services Application Identification) | **13335**

profile (Services PCEF) | **13336**

profile (Services Proxy) | **13338**

profile (Services SSL Proxy) | **13339**

profile (SSL Initiation) | **13343**

profile (SSL Termination) | **13347**

profile (Traffic Load Balancer) | **13349**

profile-type (Dynamic Service Profiles) | **13353**

profile-variable-set (Dynamic Profiles) | **13354**

profile-variable-set (Routing Instances) | **13356**

profile (Web Filter) | **13357**

profilerd | **13360**

profiles | **13362**

profiles (AAA) | **13363**

profiles (PCEF) | **13365**

profiles (RFC 2544 Benchmarking) | **13368**

programmable-rpd | **13369**

promiscuous | **13372**

promiscuous-mode (Dynamic IGMP Interface) | **13373**

promiscuous-mode (Interfaces) | **13374**

promiscuous-mode (Protocols IGMP) | **13375**

promote | **13377**

propagate-settings | **13379**

propagate-tc | **13380**

propagation-delay | **13381**

proposal (Security Group VPN Member IKE) | **13383**

proposal (Security Group VPN Server IKE) | **13385**

proposal (Security Group VPN Server IPsec) | **13387**

proposal (Security IKE) | **13389**

proposal (Security IKE Dynamic SA) | **13393**

proposal (Security IPsec) | **13394**

proposal (Services IKE) | **13399**

proposal (Services IPsec VPN) | **13400**

proposals (Services) | **13402**

proposals | **13403**

protect | **13404**

protect (Security) | **13406**

protect core | **13407**

protect-interface | **13409**

protect-maintenance-association (OAM) | **13410**

protected-l2circuit | **13412**

protection-group | **13413**

protection (Protocols BGP) | **13416**

protection (Protocols MPLS) | **13417**

protection-revert-time | **13419**

protector-interface | **13420**

protector-pe | **13421**

protocol (Firewall Filter) | **13423**

protocol | **13424**

protocol (early-aging) | **13428**

protocol (Application Identification) | **13429**

protocol (Applications) | **13431**

protocol-binding | **13433**

protocol-command | **13434**

protocol-down | **13436**

protocol (Drop Profile Map) | **13437**

protocol (Drop Profiles) | **13438**

protocol (Dynamic Schedulers) | **13439**

protocol (Flow Descriptions) | **13441**

protocol (Host Outbound Traffic) | **13443**

protocol (IPsec) | **13444**

protocol (Junos-FIPS Software) | **13445**

protocol (Junos OS) | **13446**

protocol-master | **13448**

protocol-name | **13451**

protocol-origin | **13452**

protocol (OVSDB) | **13454**

protocol (Protocols OpenFlow) | **13455**

protocol (Rewrite Rules) | **13457**

protocol (Schedulers) | **13459**

protocol (Security Destination NAT) | **13460**

protocol (Security Forwarding Options) | **13461**

protocol (Security IDP Signature Attack) | **13463**

protocol (Security Source NAT) | **13469**

protocol (Server MEP) | **13471**

protocol (Services Proxy) | **13472**

protocol (Subscriber Secure Policy) | **13474**

protocol-version | **13475**

protocols | **13477**

protocols | **13748**

protocols (DDoS) | **13751**

protocols (DDoS) (ACX Series, PTX Series, and QFX Series) | **13765**

protocols (Dynamic Profiles) | **13791**

protocols (Fabric) | **13794**

protocols (FIP) | **13795**

protocols rsvp expand-flood-reflector-hop | **13797**

protocols (Security Zones Host Inbound Traffic) | **13798**

protocols (Security Zones Interfaces) | **13800**

protocols (STP Type) | **13802**

protocols (Tenant Systems) | **13804**

provider-tunnel | **13806**

providers | **13812**

provisioning-method (Remote Device Management) | **13814**

provisioning-order (Diameter Applications) | **13815**

proxy (interfaces) | **13817**

proxy (Bridge Domains) | **13819**

proxy-arp | **13820**

proxy-arp (Dynamic Profiles) | **13822**

proxy-arp (Security NAT) | **13823**

proxy (Fibre Channel) | **13825**

proxy-generate | **13826**

proxy (Interfaces) | **13828**

proxy-mac | **13829**

proxy-macip-advertisement | **13831**

proxy-mode | **13833**

proxy (Multicast VLAN Registration) | **13834**

proxy-ndp (Security NAT) | **13836**

proxy-profile | **13838**

proxy (Security Antivirus) | **13839**

proxy (Services) | **13841**

proxy (snmp) | **13843**

proxy (System) | **13845**

ps0 (Pseudowire Subscriber Interfaces) | **13847**

pseudowire-service (Pseudowire Subscriber Interfaces) | **13848**

pseudowire-status-tlv | **13849**

psn-tunnel-endpoint | **13851**

psu | **13852**

psu (satellite-management) | **13853**

ptp | **13855**

ptp-mode | **13865**

ptsp-rules | **13866**

pulse-per-second-enable | **13868**

purge-flow-timer (Protocols OpenFlow) | **13869**

purge-timeout (programmable-rpd) | **13870**

purging | **13872**

push (protocols mpls) | **13873**

push | **13875**

push-backup-to-master (Accounting Options) | **13876**

push (Dynamic VLANs) | **13878**

push-push | **13879**

push-to-identity-management | **13880**

pvlan | **13882**

pvlan-trunk | **13883**

python-script-user (Event and SNMP Scripts) | **13884**

Q | 13886

qos-adjust | **13888**

qos-adjust-adsl | **13890**

qos-adjust-adsl2 | **13892**

qos-adjust-adsl2-plus | **13894**

qos-adjust-hierarchical | **13896**

qos-adjust-other | **13897**

qos-adjust-sdsl | **13899**

qos-adjust-vdsl | **13901**

qos-adjust-vdsl2 | **13903**

qualified-bum-pruning-mode | **13905**

qualified-next-hop (Access) | **13906**

qualified-next-hop (Access-Internal) | **13907**

qualified-next-hop (Dynamic Access-Internal Routes) | **13908**

qualified-next-hop (Static Routes) | **13910**

qualified-vlan | **13912**

quality (Interfaces) | **13914**

quality-level | **13915**

quality-level (Clock Synchronization) | **13917**

quality-level (hybrid) | **13919**

quality-mode-enable (MX) | **13920**

quality-mode-enable (PTX) | **13922**

quarantine-message (Security Content Security) | **13924**

querier (performance-monitoring) | **13925**

query-interval (Bridge Domains) | **13927**

query-interval (Protocols IGMP) | **13929**

query-interval (Protocols IGMP AMT) | **13930**

query-interval (Protocols MLD) | **13932**

query-last-member-interval (Bridge Domains) | **13933**

query-last-member-interval (Protocols IGMP) | **13935**

query-last-member-interval (Protocols MLD) | **13936**

query-response-interval (Bridge Domains) | **13938**

query-response-interval (Protocols IGMP) | **13940**

query-response-interval (Protocols IGMP AMT) | **13941**

query-response-interval (Protocols MLD) | **13943**

queue (Global Queues) | **13944**

queue-monitoring | **13946**

queue-num | **13947**

queue (Restricted Queues) | **13949**

queue-statistics | **13950**

queue-threshold | **13953**

queues | **13955**

quit | **13956**

R | 13958

r2cp | **13990**

radio-router (Interfaces) | **13991**

radius (Access Profile) | **13993**

radius-accounting | **13997**

radius-coa (Adjustment Control Profiles) | **13999**

radius-disconnect | **14001**

radius-disconnect (DHCP Local Server) | **14003**

radius-disconnect-port | **14005**

radius-flow-tap | **14006**

radius-options (Access) | **14010**

radius-options (Interfaces) | **14011**

radius-options (System) | **14013**

radius-reachability | **14015**

radius-realm | **14017**

radius (Security Log) | **14018**

radius-server | **14020**

radius-server (Access) | **14026**

radius-server-name | **14028**

radius-server (System) | **14030**

radius (System) | **14034**

radsec | **14036**

radsec-destination | **14040**

raise-threshold | **14041**

raise-trap (Event Policy) | **14042**

random | **14043**

range (Access) | **14045**

range-address | **14047**

range (Address-Assignment Pools) | **14048**

range (Source NAT Next Gen Services) | **14049**

ranges (Dynamic Stacked VLAN) | **14051**

ranges (Dynamic VLAN) | **14052**

rapid-commit (DHCPv6 Client) | **14054**

rapid-commit (DHCPv6 Local Server) | **14055**

rate | **14057**

rate-cap | **14058**

rate (Forwarding Options) | **14060**

rate (Interface Services) | **14061**

rate-limit | **14063**

rate-limit (Aggregated rate limit) | **14065**

rate-limit (Security GTP) | **14068**

rate-limit (Security Log) | **14069**

rate-limiters | **14071**

rate (Routing Instances) | **14073**

rcp-count | **14075**

rdi | **14077**

rdma-opcode | **14078**

re-assembler | **14079**

re-enroll-trigger-time-percentage | **14081**

re-generate-keypair | **14082**

reachability-fault | **14083**

reachable-time | **14086**

reachable-time (Dynamic Router Advertisement) | **14087**

read-view | **14089**

readvertise | **14090**

real-service (Traffic Load Balancer) | **14092**

realm | **14093**

realm-delimiter (Domain Map) | **14095**

realm (Diameter Origin) | **14096**

realm-parse-direction (Domain Map) | **14098**

realtime-ukern-thread (BFD) | **14099**

reassemble-packets | **14100**

reassembly-timeout | **14102**

reauthenticate (DHCP Local Server) | **14103**

rebalance (Aggregated Ethernet Interfaces) | **14106**

rebalance-periodic (Aggregated Ethernet Subscriber Interfaces) | **14108**

rebalance-periodic (PS Interface) | **14109**

rebalance-subscriber-granularity (PS interface) | **14111**

receive-ais | **14112**

receive-failure-threshold (RFC 2544 Benchmarking) | **14113**

receive-options-packets (Passive Flow Monitoring) | **14114**

receive-options-packets | **14116**

receive (Protocols RIP) | **14117**

receive (Protocols RIPng) | **14119**

receive-ttl-exceeded (Passive Flow Monitoring) | **14120**

receive-ttl-exceeded | **14121**

receive-window | **14123**

receiver | **14124**

receiver (Graceful Restart for BGP Helper) | **14127**

recommendation-tlv | **14129**

recommended | **14130**

recommended-action | **14131**

reconfigure (DHCP Local Server) | **14133**

reconnect-time | **14135**

record | **14136**

record (Origin Validation for BGP) | **14138**

record-route-option (IDS Screen Next Gen Services) | **14139**

recover-time (DDoS) | **14141**

recovered-clock | **14142**

recovery | **14143**

recovery | **14146**

recovery-time (LDP) | **14148**

recovery-timeout | **14149**

recycle-bandwidth-profiles | **14152**

red-buffer-occupancy | **14154**

red-differential-delay | **14156**

redial-delay | **14157**

redirect (PCC Action Profiles) | **14159**

redirect-traffic | **14160**

redirect-url | **14162**

redirect-wx (Application Services) | **14164**

redistribute-all-traffic (Aggregated Multiservices) | **14165**

redundancy | **14166**

redundancy-event (Services Redundancy Daemon) | **14168**

redundancy (Graceful Switchover) | **14170**

redundancy-group | **14171**

redundancy-group (Chassis Cluster) | **14173**

redundancy-group (Chassis - MX Series) | **14176**

redundancy-group-id (Junos Fusion) | **14177**

redundancy-group (Interfaces) | **14179**

redundancy-group (Redundant Tunnel) | **14180**

redundancy-interface-process | **14182**

redundancy (Junos OS) | **14183**

redundancy (Junos OS Evolved) | **14185**

redundancy (M:N Subscriber Redundancy) | **14186**

redundancy-mode | **14191**

redundancy-options (Adaptive Services Interfaces) | **14192**

redundancy-options (Aggregated Multiservices) | **14194**

redundancy-options (Link Services IQ PIC Interfaces) | **14195**

redundancy-options (Stateful Synchronization) | **14197**

redundancy-policy (Interchassis Services Redundancy) | **14199**

redundancy (Power Management) | **14201**

redundancy (satellite-management) | **14203**

redundancy-set | **14204**

redundancy-set-id (Service Set) | **14206**

redundancy-mode | **14207**

redundant-ether-options | **14209**

redundant-logical-tunnel | **14211**

redundant-parent (Interfaces) | **14212**

redundant-parent (Interfaces Fast Ethernet) | **14214**

redundant-parent (Interfaces Gigabit Ethernet) | **14215**

redundant-power-system | **14216**

redundant-pseudo-interface-options | **14218**

redundant-sources | **14219**

redundant-trunk-group | **14221**

redundant-virtual-tunnel | **14222**

reference-bandwidth (Protocols IS-IS) | **14224**

reference-bandwidth (Protocols OSPF) | **14226**

reflect-etype | **14228**

reflect-mode | **14229**

reflective-relay | **14232**

reflexive | revert | reverse | **14233**

refresh (Commit Scripts) | **14234**

refresh (Event Scripts) | **14236**

refresh-from (Commit Scripts) | **14238**

refresh-from (Event Scripts) | **14240**

refresh-from (JET) | **14242**

refresh-from (Op Scripts) | **14244**

refresh-from (SNMP Scripts) | **14246**

refresh-interval | **14248**

refresh (JET) | **14250**

refresh (Op Scripts) | **14251**

refresh-rate (Flow Monitoring Logs for NAT) | **14253**

refresh (SNMP Scripts) | **14255**

refresh-time | **14257**

regex-additive-logic | **14258**

regex (Class Attribute) | **14260**

regexp | **14261**

region | **14262**

register-limit | **14264**

register-probe-time | **14266**

registration | **14268**

reject (Security) | **14270**

reject-unauthorized-ipv6cp | **14271**

rejoin-timeout (Aggregated Multiservices) | **14273**

relay-agent-interface-id (DHCP Local Server) | **14275**

relay-agent-interface-id (DHCPv6 Relay Agent) | **14276**

relay-agent-interface-id (DHCPv6 Relay Agent Username) | **14278**

relay-agent-option | **14279**

relay-agent-option-79 | **14281**

relay-agent-remote-id (DHCP Local Server) | **14282**

relay-agent-remote-id (DHCPv6 Relay Agent) | **14284**

relay-agent-remote-id (DHCPv6 Relay Agent Username) | **14286**

relay-agent-subscriber-id (DHCP Local Server) | **14288**

relay-agent-subscriber-id (DHCPv6 Relay Agent) | **14289**

relay (AMT Protocol) | **14291**

relay (IGMP) | **14292**

relay-option-82 | **14294**

relay-option (DHCP Relay Agent) | **14297**

relay-option-vendor-specific (dhcpv6) | **14299**

relay-server-group (DHCP Relay Agent Option) | **14300**

relay-source | **14302**

reliable | **14304**

remote | **14305**

remote-access (Juniper Secure Connect) | **14307**

remote-address | **14311**

remote-address (for LMP Control Channel) | **14312**

remote-address (for LMP Traffic Engineering) | **14314**

remote-debug-permission | **14315**

remote-device-management | **14317**

remote-engine | **14318**

remote-execution (Event Scripts) | **14320**

remote-gateway | **14322**

remote-gateway (Tunnel Profile) | **14323**

remote (Host VPN) | **14324**

remote-id (Services) | **14326**

remote-id | **14327**

remote-id (protocols link-management) | **14330**

remote-id (DHCP Relay Agent) | **14331**

remote-id-mismatch (DHCP Local Server and DHCP Relay Agent) | **14334**

remote-id (MX Series) | **14337**

remote-id (VLAN Authentication Username) | **14338**

remote-identity | **14340**

remote-interface-switch | **14342**

remote-ip-address | **14344**

remote-ip-host-routes | **14346**

remote-loop-enable | **14348**

remote-loopback | **14350**

remote-maintenance-association (OAM) | **14351**

remote-mep | **14352**

remote-mep (EX Series Switch Only) | **14354**

remote-port-range | **14355**

remote-ports | **14357**

remote-site-id | **14360**

remote-traffic-selector | **14361**

remote-vtep-list | **14362**

remotely-controlled | **14364**

remove-ie | **14366**

remove-private | **14368**

remove-when-no-subscribers | **14370**

rename | **14371**

renegotiation (Services) | **14373**

reopen-time | **14374**

replace | **14376**

replace-ip-source-with (DHCP Relay Agent) | **14377**

replay-attacks | **14379**

replay-protect | **14380**

replay-protect (MX Series) | **14382**

replay-window-size | **14383**

replay-window-size (MX Series) | **14385**

replicator | **14387**

report-ingress-shaping-rate (Dynamic CoS Interfaces) | **14390**

report-interface-descriptions (Access) | **14392**

report-local-rule (PCRF Partition) | **14393**

report (LRF Rule) | **14395**

report-resource-allocation (PCRF Partition) | **14396**

report (Security Log) | **14398**

report-skip | **14403**

report-successful-resource-allocation (PCRF Partition) | **14404**

report-zero-oif-gw-on-discard | **14406**

req-option | **14408**

req-timeout | **14409**

request | **14411**

request access-security router-advertisement-guard-block | **14413**

request access-security router-advertisement-guard-forward | **14414**

request access-security router-advertisement-guard-learn interface | **14415**

request access-security slaac-snooping unblock | **14417**

request ancp oam interface | **14419**

request ancp oam neighbor | **14421**

request ancp oam port-down | **14423**

request ancp oam port-up | **14425**

request app-engine cleanup | **14427**

request app-engine file-copy (crash | log) from-jhost to-vjunos | **14429**

request auto-configuration reconnect-pending | **14431**

request-cache-timeout (RADIUS Snoop Segment) | **14432**

request chassis afeb | **14434**

request chassis beacon | **14435**

request chassis cb (Security) | **14440**

request chassis cb | **14441**

request chassis ccg | **14447**

request chassis cfeb | **14449**

request chassis clock master switch | **14450**

request chassis cluster configuration-synchronize | **14452**

request chassis cluster failover node | **14454**

request chassis cluster failover redundancy-group | **14455**

request chassis cluster failover reset | **14457**

request chassis cluster in-service-upgrade abort (ISSU) | **14459**

request chassis device-mode | **14460**

request chassis device-mode satellite | **14463**

request chassis fabric fpc | **14465**

request chassis fabric pfe | **14467**

request chassis fabric plane | **14468**

request chassis feb | **14472**

request chassis fpc | **14475**

request chassis fpc-control-port | **14486**

request chassis fpc slc | **14488**

request chassis fpm resync | **14489**

request chassis high-availability failover services-redundancy-group | **14492**

request chassis mcs | **14494**

request chassis mic | **14495**

request chassis optics | **14499**

request chassis pcg | **14501**

request chassis pic | **14502**

request chassis pic-mode | **14509**

request chassis port-led | **14512**

request chassis primary-ha-control-port-transition | **14516**

request chassis redundancy feb slot | **14518**

request chassis routing-engine hard-disk-test | **14520**

request chassis routing-engine master | **14529**

request chassis satellite beacon | **14535**

request chassis satellite disable | **14537**

request chassis satellite enable | **14539**

request chassis satellite file-copy | **14541**

request chassis satellite install | **14542**

request chassis satellite interface | **14545**

request chassis satellite login | **14547**

request chassis satellite reboot | **14548**

request chassis satellite restart | **14550**

request chassis satellite shell-command | **14551**

request chassis scg | **14553**

request chassis sfb | **14555**

request chassis sfm | **14558**

request chassis sfm master switch | **14559**

request chassis sib | **14562**

request chassis spmb restart | **14565**

request chassis ssb master switch | **14568**

request chassis synchronization mode | **14570**

request (Chassis Synchronization Source) | **14573**

request chassis synchronization switch | **14574**

request chassis system-mode | **14577**

request (Clock Synchronization) | **14579**

request component login | **14580**

request csds add-vsrx | **14583**

request csds authenticate-host | **14584**

request csds delete-vsrx | **14586**

request csds extract-vsrx-keys | **14587**

request csds jdm | **14589**

request csds sync-controller | **14592**

request debug information | **14594**

request dhcp client renew | **14595**

request dhcp relay bulk-leasequery | **14597**

request dhcp relay leasequery | **14599**

request dhcp server reconfigure | **14601**

request dhcpv6 client renew | **14604**

request dhcpv6 relay bulk-leasequery | **14605**

request dhcpv6 relay leasequery | **14607**

request dhcpv6 server reconfigure | **14609**

request diagnostics tdr | **14611**

request extension-service restart-daemonize-app | **14614**

request extension-service (start | stop) | **14615**

request fabric administration director-group change-master | **14616**

request fabric administration power-off interconnect-device | **14618**

request fabric administration power-off interconnect-device fpc | **14620**

request fabric administration power-off node-device | **14622**

request fabric administration remove | **14624**

request fabric administration system mac-pool add | **14626**

request fabric administration system mac-pool delete | **14628**

request fibre-channel proxy load-rebalance | **14630**

request flight-recorder set high-cpu | **14633**

request forwarding-options port-mirroring instance family self-mirror-start | **14636**

request forwarding-options port-mirroring instance family self-mirror-stop | **14638**

request interface bounce | **14639**

request interface link-degrade-recover | **14642**

request interface load-balancing revert (Aggregated Multiservices) | **14646**

request interface load-balancing switchover (Aggregated Multiservices) | **14647**

request interface mc-ae switchover (Multichassis Link Aggregation) | **14649**

request interface optics-reset | **14651**

request interface rebalance (Aggregated Ethernet for Subscriber Management) | **14653**

request interface revert | **14657**

request interface (revert | switchover) (Adaptive Services) | **14658**

request interface (revert | switchover) (Aggregated Ethernet Link Protection) | **14660**

request interface (revert | switchover) (Aggregated Inline Service Interfaces) | **14662**

request interface switchover | **14664**

request ipsec switch | **14665**

request jdaf clear statistics service | **14667**

request jnu role | **14668**

request jnu satellite sync | **14670**

request jnu schema add | **14671**

request jnu schema delete | **14673**

request jnu schema merge | **14674**

request l2circuit-switchover | **14675**

request lacp link-switchover | **14677**

request-max-tcp-connections (System Process) | **14678**

request message | **14680**

request modem wireless create-profile | **14681**

request modem wireless fota | **14684**

request modem wireless sim-lock | **14685**

request modem wireless sim-unlock | **14687**

request mpls container-lsp | **14689**

request mpls lsp adjust-autobandwidth | **14691**

request network-access aaa accounting | **14693**

request network-access aaa replay pending-accounting-stops | **14695**

request network-access aaa subscriber add session-id | **14696**

request network-access aaa subscriber delete session-id | **14699**

request network-access aaa subscriber modify session-id | **14702**

request network-access aaa subscriber set session-id | **14704**

request network-access aaa saml export-sp-metadata | **14706**

request network-access aaa saml load-idp-metadata | **14707**

request node halt (Junos OS Evolved) | **14708**

request node (offline | online) (Junos OS Evolved) | **14711**

request node power-off (Junos OS Evolved) | **14713**

request node power-on (Junos OS Evolved) | **14715**

request node reboot (re0 | re1) (Junos OS Evolved) | **14717**

request packet-capture start | **14720**

request packet-capture stop | **14726**

request path-computation-client active-pce | **14728**

request pim multicast-tunnel rebalance | **14729**

request pppoe connect | **14731**

request pppoe disconnect | **14732**

request protection-group ethernet-aps | **14734**

request protection-group ethernet-ring | **14736**

request-rate (Access) | **14738**

request redundant-power-system multi-backup | **14739**

request-response | **14741**

request restart packet-forwarding-engine | **14742**

request routing-engine | **14743**

request security certificate enroll (Signed) | **14745**

request security certificate enroll (Unsigned) | **14747**

request security datapath-debug capture start | **14749**

request security datapath-debug capture stop | **14750**

request security idp jist-conversion | **14751**

request security idp security-package download | **14754**

request security idp security-package install | **14757**

request security idp security-package offline-download | **14760**

request security idp security package rollback | **14761**

request security idp ssl-inspection key add | **14763**

request security idp ssl-inspection key delete | **14766**

request security idp storage-cleanup | **14768**

request security ike debug-disable | **14770**

request security ike debug-enable | **14771**

request security integrity | **14773**

request security internal-security-association refresh | **14774**

request security key-pair | **14776**

request security pki ca-certificate ca-profile-group default-trusted-ca-certs | **14778**

request security pki ca-certificate ca-profile-group load | **14781**

request security pki ca-certificate enroll | **14784**

request security pki ca-certificate enroll (Security) | **14786**

request security pki ca-certificate load | **14787**

request security pki ca-certificate load (Security) | **14789**

request security pki ca-certificate verify | **14791**

request security pki ca-certificate verify (Security) | **14792**

request security pki ca-profile-group-status | **14795**

request security pki crl load | **14797**

request security pki crl load (Security) | **14798**

request security pki generate-certificate-request | **14800**

request security pki generate-certificate-request (Security) | **14802**

request security pki generate-key-pair | **14805**

request security pki generate-key-pair | **14806**

request security pki generate-key-pair (Security) | **14808**

request security pki key-pair export | **14810**

request security pki local-certificate enroll | **14812**

request security pki local-certificate enroll cmpv2 | **14814**

request security pki local-certificate enroll scep | **14817**

request security pki local-certificate export | **14822**

request security pki local-certificate generate-self-signed | **14823**

request security pki local-certificate generate-self-signed | **14825**

request security pki local-certificate generate-self-signed (Security) | **14827**

request security pki local-certificate load | **14830**

request security pki local-certificate re-enroll cmpv2 | **14831**

request security pki local-certificate re-enroll scep | **14833**

request security pki local-certificate verify | **14835**

request security pki node-local generate-certificate-request | **14837**

request security pki node-local generate-key-pair | **14840**

request security pki node-local key-pair export | **14842**

request security pki node-local local-certificate enroll | **14843**

request security pki node-local local-certificate export | **14847**

request security pki node-local local-certificate load | **14848**

request security pki node-local local-certificate re-enroll | **14850**

request security pki node-local local-certificate verify | **14852**

request security pki sync-from-peer | **14854**

request security pki verify-integrity-status | **14856**

request security policies check | **14857**

request security policies resync | **14861**

request security re-distribution ipsec-vpn | **14863**

request security ssh key-pair-identity generate | **14865**

request security ssh password-less-authentication (Junos OS Evolved) | **14867**

request security system-keystore | **14869**

request security system-keystore unlink | **14870**

request security tpm master-encryption-password set | **14871**

request security uefi revoke | **14873**

request security user-identification local-authentication-table delete | **14875**

request security user-identification local-authorization-table add | **14877**

request security utm anti-virus avira-engine | **14879**

request security utm anti-virus juniper-express-engine | **14881**

request security utm anti-virus kaspersky-lab-engine | **14883**

request security utm anti-virus sophos-engine | **14885**

request security utm custom-objects url-feed update | **14886**

request security utm custom-objects url-feed update feed-name force | **14888**

request security utm web-filtering categorize | **14889**

request security utm web-filtering category download-install [version] | **14891**

request security utm web-filtering category download [version] | **14892**

request security utm web-filtering category install | **14893**

request security utm web-filtering category migrate-to-ng-juniper | **14895**

request security utm web-filtering category uninstall | **14897**

request security utm web-filtering custom-page reload | **14898**

request security utm web-filtering recategorize | **14899**

request server authenticate-peer-server | **14901**

request services advanced-anti-malware data-connection | **14903**

request services advanced-anti-malware diagnostic | **14905**

request services advanced-anti-malware redirect-file | **14910**

request services anti-virus update | **14912**

request services anti-virus machine-learning-engine | **14914**

request services application-identification application | **14916**

request services application-identification clear packet-capture all | **14918**

request services application-identification download | **14919**

request services application-identification download status | **14923**

request services application-identification group | **14925**

request services application identification install ignore duplicate version check | **14928**

request services application-identification install status | **14929**

request services application identification new to production | **14931**

request services application-identification rollback status | **14933**

request services application-identification install | **14934**

request services application-identification offline-download package-path | **14936**

request services application-identification proto-bundle-status | **14937**

request services application-identification uninstall | **14939**

request services application-identification uninstall status | **14940**

request services extensible-subscriber-services reload-dictionary | **14942**

request services flow-collector change-destination primary interface | **14943**

request services flow-collector change-destination secondary interface | **14945**

request services flow-collector test-file-transfer | **14946**

request services ip-monitoring preempt-restore policy | **14948**

request services ipsec-vpn ipsec switch tunnel | **14950**

request services l2tp destination unlock | **14951**

request services monitoring twamp client | **14953**

request services paa install | **14954**

request services paa uninstall | **14958**

request services redundancy-set trigger | **14960**

request services remote-device-management reconfigure service-device | **14961**

request services remote-device-management reload-dictionary | **14962**

request services rpm twamp | **14964**

request services static-subscribers login group | **14966**

request services static-subscribers login interface | **14968**

request services static-subscribers logout group | **14969**

request services static-subscribers logout interface | **14971**

request services subscribers | **14972**

request services subscribers clear | **14974**

request services url-filter delete gencfg-data | **14975**

request services url-filter force dns-resolution | **14976**

request services url-filter update url-filter-database file | **14979**

request services url-filter validate | **14980**

request services user-identification active-directory-access active-directory-authentication-table delete | **14981**

request services user-identification active-directory-access domain-controller | **14983**

request services user-identification active-directory-access ip-user-probe | **14985**

request services user-identification authentication-source aruba-clearpass user-query | **14987**

request services user-identification authentication-source jims groups domain <domain-name> (force-fetch|status) | **14989**

request services user-identification authentication-source jims validate (user <user-name>| group <group-name>|device <device-name>) domain <domain-name> | **14991**

request services user-identification authentication-table delete | **14994**

request services web-filter delete gencfg-data | **15004**

request services web-filter force dns-resolution | **15005**

request services web-filter update dns-filter-database | **15006**

request services web-filter update url-filter-database file | **15007**

request services web-filter validate dns-filter-file-name | **15009**

request services web-filter validate url-filter-file-name | **15010**

request session member | **15011**

request snmp spoof-trap | **15012**

request snmp utility-mib clear instance | **15021**

request snmp utility-mib set instance | **15023**

request support information | **15024**

request system application (Junos OS Evolved) | **15042**

request system autorecovery state | **15044**

request system boot console | **15047**

request system certificate add | **15048**

request system commit server pause | **15050**

request system commit server queue cleanup | **15051**

request system commit server start | **15053**

request system configuration database resize | **15054**

request system configuration rescue delete | **15057**

request system configuration rescue save | **15059**

request system convert-json-configuration | **15060**

request system debug-info | **15062**

request system decrypt password | **15066**

request system disable unified-services | **15068**

request system download abort | **15069**

request system download clear | **15071**

request system download pause | **15072**

request system download resume | **15074**

request system download start | **15076**

request system enable unified-services | **15078**

request system filesystem encryption enable | **15079**

request system filesystem encryption keys delete | **15082**

request system firmware downgrade optics | **15085**

request system firmware reload (Junos OS Evolved) | **15087**

request system firmware upgrade | **15089**

request system firmware upgrade (Junos OS Evolved) | **15099**

request system firmware upgrade poe | **15102**

request system halt | **15105**

request-system-integrity-attestation | **15112**

request-system-integrity-log-retrieval | **15114**

request system license update | **15115**

request system logout | **15117**

request system malware-scan | **15119**

request system partition abort | **15122**

request system partition compact-flash | **15125**

request system partition hard-disk | **15127**

request system power-off | **15131**

request system power-off (SRX Series) | **15135**

request system process terminate | **15138**

request system reboot (Junos OS) | **15140**

request system reboot (Junos OS Evolved) | **15151**

request system reboot (Junos OS with Upgraded FreeBSD) | **15154**

request system reboot (SRX Series) | **15159**

request system reboot usb (SRX Series) | **15160**

request system recover | **15161**

request system restore state path | **15165**

request system save state path | **15166**

request system scripts add (Chassis Cluster) | **15167**

request system scripts add | **15171**

request system scripts convert | **15172**

request system scripts delete | **15175**

request system scripts event-scripts reload | **15176**

request system scripts refresh-from | **15177**

request system scripts rollback | **15181**

request system scripts synchronize | **15182**

request system services dhcp | **15185**

request system snapshot (Junos OS Evolved) | **15187**

request system snapshot (Junos OS with FreeBSD Prior to Release 10) | **15190**

request system snapshot (Junos OS with Upgraded FreeBSD) | **15203**

request system snapshot (Maintenance) | **15208**

request system software abort | **15212**

request system software abort in-service-upgrade (ICU) | **15215**

request system software add (Junos OS) | **15216**

request system software add (Junos OS Evolved) | **15233**

request system software add (Maintenance) | **15238**

request system software add optional://auto-snapshot | **15240**

request system software add restart | **15242**

request system software configuration-backup | **15245**

request system software configuration-restore | **15247**

request system software delete-backup | **15248**

request system software delete (Junos OS) | **15250**

request system software delete (Junos OS Evolved) | **15254**

request system software download | **15258**

request system software format-qfabric-backup | **15260**

request system software in-service-upgrade | **15262**

request system software in-service-upgrade (Maintenance) | **15282**

request system software in-service-upgrade (MX Series 5G Universal Routing Platforms and EX9200 Switches) | **15288**

request system software nonstop-upgrade (QFabric Systems) | **15311**

request system software nonstop-upgrade | **15319**

request system software recover-from-restore-point | **15332**

request system software recovery-package | **15334**

request system software restore-point | **15338**

request system software rollback (Junos OS) | **15340**

request system software rollback (Junos OS Evolved) | **15345**

request system software sync | **15348**

request system software system-backup | **15353**

request system software validate in-service-upgrade | **15355**

request system software validate (Junos OS) | **15360**

request system software validate (Junos OS Evolved) | **15365**

request system software validate on (Junos OS with Upgraded FreeBSD) | **15368**

request system software validate-restart (Junos OS Evolved) | **15371**

request system storage cleanup (Junos OS) | **15374**

request system storage cleanup (Junos OS Evolved) | **15388**

request system storage cleanup (SRX Series) | **15400**

request system yang add | **15404**

request system yang delete | **15408**

request system yang disable | **15410**

request system yang enable | **15413**

request system yang update | **15414**

request system yang validate | **15417**

request system zeroize | **15419**

request system zeroize (Junos OS) | **15421**

request system zeroize (Junos OS Evolved) | **15425**

request system zeroize ztp-option secure | **15430**

request-timeout | **15432**

request-type | **15433**

request unified-edge tdf call-trace clear | **15435**

request unified-edge tdf call-trace show | **15436**

request unified-edge tdf call-trace start | **15440**

request unified-edge tdf call-trace stop | **15443**

request validation policy | **15445**

request virtual-chassis device-reachability | **15447**

request virtual-chassis member-id delete (MX Series Virtual Chassis) | **15451**

request virtual-chassis member-id set | **15452**

request virtual-chassis mode | **15455**

request virtual-chassis reactivate | **15461**

request virtual-chassis recycle | **15462**

request virtual-chassis renumber | **15464**

request virtual-chassis routing-engine master switch | **15465**

request virtual-chassis vc-port | **15468**

request virtual-chassis vc-port (Dedicated VCP) | **15472**

request virtual-chassis vc-port delete (MX Series Virtual Chassis) | **15474**

request virtual-chassis vc-port diagnostics optics | **15476**

request virtual-chassis vc-port set (MX Series Virtual Chassis) | **15478**

request virtual-network-functions | **15481**

request vmhost cleanup | **15482**

request vmhost cleanup (NFX Series) | **15485**

request vmhost copy jnode-to-vjunos | **15486**

request vmhost copy vjunos-to-jnode | **15487**

request vmhost file-copy | **15488**

request vmhost file-copy (NFX Series) | **15490**

request vmhost halt | **15492**

request vmhost halt (NFX Series) | **15494**

request vmhost hard-disk-test | **15495**

request vmhost jdm add (In-Chassis Model) | **15497**

request vmhost jdm delete (In-Chassis Model) | **15499**

request vmhost jdm login (In-Chassis Model) | **15500**

request vmhost jdm start (In-Chassis Model) | **15501**

request vmhost jdm stop (In-Chassis Model) | **15502**

request vmhost local-file-copy | **15503**

request vmhost mode | **15506**

request-vmhost-power-cycle-cb | **15507**

request-vmhost-power-cycle-re | **15516**

request vmhost power-off | **15525**

request vmhost power-off (NFX Series) | **15526**

request vmhost power-on | **15528**

request vmhost reboot | **15529**

request vmhost reboot (NFX Series) | **15532**

request vmhost remote-file-copy | **15535**

request vmhost snapshot | **15539**

request vmhost software abort in-service-upgrade | **15542**

request vmhost software add | **15543**

request vmhost software add (NFX Series) | **15548**

request vmhost software in-service-upgrade | **15551**

request vmhost software rollback | **15556**

request vmhost software validate | **15559**

request vmhost storage | **15561**

request vmhost zeroize | **15563**

request vmhost zeroize ztp-option secure | **15566**

request vpls-switchover | **15568**

request wlan access-point packet capture | **15570**

requested-ip-network-match (DHCP Local Server) | **15572**

required-depth | **15574**

reroute-address | **15575**

reserve ingress-egress-mcast | **15577**

reserved-bandwidth | **15579**

reserved (Security IDP Custom Attack) | **15580**

reset (Security IDP) | **15582**

reset-tracking-bit | **15583**

reset-count | **15584**

resignal-minimum-bandwidth | **15585**

resiliency | **15586**

resilient-hash | **15589**

resize (Ephemeral Database) | **15590**

resize-ephemeral-default | **15592**

resolution | **15594**

resolution-map | **15597**

resolution-ribs | **15599**

resolve | **15600**

resolve-vpn | **15602**

resource (Analytics) | **15604**

resource-category jtree (Resource Monitor) | **15605**

resource-cleanup | **15607**

resource (Interfaces) | **15608**

resource-list (Network Management) | **15610**

resource-list (Telemetry) | **15612**

resource-manager | **15614**

resource-manager (Forwarding-options) | **15616**

resource-monitor | **15618**

resource-profiles (Analytics) | **15620**

resource-watch | **15622**

resources | **15625**

respond-bad-spi (Services IKE Policy) | **15626**

responder (performance-monitoring) | **15628**

response-cache-timeout (RADIUS Client) | **15629**

rest | **15631**

restart | **15633**

restart chassis-control | **15648**

restart-duration (Protocols) | **15651**

restart-duration (Multicast Snooping) | **15653**

restart-duration (Routing Options) | **15655**

restart extensible-subscriber-services | **15656**

restart (JDM) | **15657**

restart (Junos OS Evolved) | **15660**

restart-path | **15662**

restart-time (BGP Graceful Restart) | **15663**

restart-time (Fabric Control) | **15665**

restarter (Graceful Restart for BGP Restarter) | **15667**

restore-interval | **15669**

restricted-queues | **15671**

retain | **15672**

retain-hold-resource | **15674**

reth-count (Chassis Cluster) | **15675**

retransmission-attempt (DHCP Client) | **15677**

retransmission-attempt (DHCPv6 Client) | **15678**

retransmission-count-established (L2TP) | **15680**

retransmission-count-not-established (L2TP) | **15681**

retransmission-interval | **15683**

retransmission-interval (DHCP) | **15684**

retransmission-interval (DHCP Client) | **15685**

retransmit-interval (Services) | **15687**

retransmit-timer | **15689**

retransmit-timer (Dynamic Router Advertisement) | **15690**

retry (Access LDAP) | **15691**

retry (Access RADIUS) | **15693**

retry (Adaptive Services Interface) | **15694**

retry-count (Security Log) | **15696**

retry-count (SNMP) | **15697**

retry-count (FPC Error Condition) | **15698**

retry-count (Chassis Cluster) | **15700**

retry-count (Event Policy) | **15701**

retry-delay | **15703**

retry (Event Policy) | **15704**

retry-interval | **15706**

retry-interval (Chassis Cluster) | **15707**

retry-limit | **15709**

retry-limit (Protocols Link Management) | **15710**

retry-options | **15712**

retry (RADIUS Server) | **15714**

retry (Services Flow Collector) | **15716**

retry-timer | **15717**

reverse | **15718**

reverse-oif-mapping | **15720**

reverse-route-packet-mode-vr | **15721**

reverse-tcp-mss | **15722**

revert-interval (Access) | **15724**

revert-interval (Access LDAP) | **15725**

revert-interval (Access RADIUS) | **15727**

revert-interval (RADIUS Server) | **15728**

revert-time (Protocols Layer 2 Circuits) | **15729**

revert-timer | **15732**

revertive | **15734**

revision-level | **15735**

revocation-check | **15737**

revocation-check (Security PKI) | **15739**

rewrite-rules (Definition) | **15741**

rewrite-rules (CoS Host Outbound Traffic) | **15743**

rewrite-rules (Definition) | **15744**

rewrite-rules (Definition) | **15746**

rewrite-rules (Dynamic CoS Interfaces) | **15748**

rewrite-rules (Interfaces) | **15749**

rewrite-rules (Interfaces) | **15751**

rewrite-value (Fibre Channel Interfaces) | **15753**

rfc-compliant (NETCONF) | **15755**

rfc2544 | **15757**

rfc2544-benchmarking | **15759**

rfc6514-compliant-safi129 (Protocols BGP) | **15763**

rib (General) | **15764**

rib-group (Protocols BGP) | **15767**

rib-group (Protocols DVMRP) | **15769**

rib-group (Protocols IS-IS) | **15770**

rib-group (Protocols MSDP) | **15772**

rib-group (Protocols OSPF) | **15774**

rib-group (Protocols PIM) | **15775**

rib-group (Protocols RIP) | **15777**

rib-group (Routing Options) | **15778**

rib-group (Source Packet Routing) | **15780**

rib-groups | **15781**

rib-list | **15783**

rib (Multitopology Routing) | **15785**

rib (Protocols BGP) | **15787**

rib (Route Resolution) | **15789**

rib-service (programmable-rpd) | **15790**

rib-sharding | **15792**

rift | **15794**

ring-id | **15797**

ring-protection-link-end | **15799**

ring-protection-link-owner | **15800**

riot-loopback | **15801**

rip | **15803**

ripng | **15804**

rising-event-index | **15805**

rising-threshold (Health Monitor) | **15807**

rising-threshold (RMON) | **15808**

rmon | **15810**

rng | **15812**

roaming-mode | **15813**

robust-count (IGMP Snooping) | **15815**

robust-count (MLD Snooping) | **15816**

robust-count (Protocols IGMP) | **15818**

robust-count (Protocols IGMP AMT) | **15820**

robust-count (Protocols MLD) | **15821**

robustness-count | **15823**

role | **15824**

role (MX Series Virtual Chassis) | **15827**

role (Protocols OpenFlow) | **15829**

rollback | **15830**

root-authentication | **15832**

root-ca (Services) | **15834**

root-logical-system | **15835**

root-login (JDM) | **15837**

root-streaming | **15838**

root-address | **15840**

route (Access) | **15841**

route (Access Internal) | **15843**

route (Access-Internal) | **15845**

route-accounting | **15846**

route-active-on | **15848**

route-attributes | **15850**

route-change-timeout | **15852**

route (chassis) | **15853**

route (Diameter Network Element) | **15854**

route-distinguisher | **15856**

route-distinguisher-id | **15860**

route-filter | **15861**

route-filter-list | **15863**

route (Interfaces) | **15868**

route-localization | **15869**

route-monitoring | **15871**

route-record | **15874**

route-refresh-priority | **15876**

route-suppression (DHCP Local Server and Relay Agent) | **15877**

route-target-filter | **15879**

route-target (Protocols BGP) | **15881**

route-target (Protocols MVPN) | **15883**

route-timeout (Protocols RIP) | **15885**

route-timeout (Protocols RIPng) | **15887**

route-distinguisher-id-use-router-id | **15889**

router-advertisement | **15890**

router-advertisement (Dynamic Profiles) | **15895**

router-advertisement-guard | **15897**

router-discovery | **15900**

router-id | **15902**

routing-engine | **15904**

routing-engine (Chassis) | **15905**

routing-engine (Chassis Redundancy) | **15908**

routing-engine-profile | **15909**

routing-engine-services | **15911**

routing-engine (SNMP Resource Level) | **15912**

routing-header | **15914**

routing-instance (Services) | **15915**

routing-instance (SNMPv3) | **15917**

routing-instance (SNMP) | **15918**

routing-instance (Forwarding-options Analyzer) | **15920**

routing-instance | **15921**

routing-instance (Interfaces) | **15923**

routing-instance-access | **15924**

routing-instance (Access LDAP) | **15925**

routing-instance (Access RADIUS) | **15927**

routing-instance (Advanced Policy-Based Routing) | **15928**

routing-instance (Flow Monitoring) | **15929**

routing-instance (Diameter Peer) | **15931**

routing-instance (Diameter Transport) | **15932**

routing-instance (Event Scripts) | **15934**

routing-instance (JET Scripts) | **15935**

routing-instance (Junos Telemetry Interface) | **15937**

routing-instance (L2TP Destination) | **15938**

routing-instance (L2TP Tunnel Destination) | **15940**

routing-instance-list (TWAMP) | **15941**

routing-instance-name | **15943**

routing-instance-name (circuit-id) | **15944**

routing-instance-name (DHCP Local Server) | **15945**

routing-instance-name (DHCP Relay Agent) | **15947**

routing-instance-name (Static Subscribers) | **15949**

routing-instance (PCC Action Profiles) | **15951**

routing-instance (Port Mirroring) | **15953**

routing-instance (PPPoE Service Name Tables) | **15954**

routing-instance (RPM) | **15956**

routing-instance (Security Content Security) | **15958**

routing-instance (Security Destination NAT) | **15959**

routing-instance (Security Source NAT) | **15960**

routing-instance (Syslog) | **15962**

routing-instance (System Scripts) | **15963**

routing-instance (Tunnel Profile) | **15965**

routing-instances (Logical Systems) | **15966**

routing-instances (JDAF) | **15967**

routing-instances | **15969**

routing-instances (RPM) | **15970**

routing-instances (CoS) | **15971**

routing-instances (Dynamic Profiles) | **15973**

routing-instances (Multiple Routing Entities) | **15976**

routing-interface | **15978**

routing-options | **15980**

routing-options (Dynamic Profiles) | **15981**

routing-options (Junos Node Slicing) | **15983**

routing-options-max-ecmp | **15985**

routing-options (QFabric System) | **15987**

routing-service (Dynamic Profiles) | **15988**

routing-services (Enhanced Subscriber Management) | **15991**

routing (System Processes) | **15993**

routing-engine-mode (TLB) | **15995**

rp | **15996**

rp-register-policy | **16000**

rp-set | **16001**

rpc | **16003**

rpc-program-number | **16004**

rpc-program-number (Applications) | **16005**

rpf-check | **16007**

rpf-check (Dynamic Profiles) | **16009**

rpf-check-policy (Routing Options) | **16011**

rpf-check-policy (Routing Options RPF) | **16012**

rpf-loose-mode-discard | **16014**

rpf-selection | **16015**

rpf-vector (PIM) | **16017**

rpm (Interfaces) | **16019**

rpm-scale | **16020**

rpm (Services) | **16023**

rpm-tracking | **16038**

rpt-spt | **16040**

rsh | **16041**

rst-invalidate-session | **16043**

rst-sequence-check | **16044**

rstp | **16046**

rsvp | **16050**

rsvp-error-hold-time | **16056**

rsvp-te (Protocols MVPN) | **16058**

rsvp-te (Routing Instances Provider Tunnel) | **16059**

rsvp-te (Routing Instances Provider Tunnel Selective) | **16061**

rsvp-te (Routing Options) | **16062**

rtf-prefix-list | **16064**

rtlog (Next Gen Services Global System Logging) | **16066**

rtp | **16068**

rts | **16069**

rts-polarity | **16071**

rtsp | **16072**

rule | **16074**

rule (AACL Rule Set) | **16076**

rule-activation-time | **16078**

rule (Advanced Policy-Based Routing) | **16080**

rule (Application Firewall) | **16083**

rule (Captive Portal Content Delivery) | **16085**

rule (casb-policy) | **16087**

rule-deactivation-time | **16089**

rule (Destination NAT Next Gen Services) | **16091**

rule (HTTP Header Enrichment for Tag Rule Set) | **16093**

rule (IDS MS-DPC) | **16094**

rule (IDS MS-MPC) | **16096**

rule (IP Reassembly) | **16100**

rule (LRF) | **16101**

rule (PCP) | **16103**

rule (Security Destination NAT) | **16105**

rule (Security Exempt Rulebase) | **16107**

rule (Security IPS Rulebase) | **16109**

rule (Security Source NAT) | **16111**

rule (Security Static NAT) | **16113**

rule (Services CoS) | **16115**

rule (Services CoS Next Gen Services) | **16117**

rule (Services NAT) | **16120**

rule (Services Stateful Firewall) | **16122**

rule-session-count-alarm (Security Destination NAT Rule Set) | **16124**

rule-session-count-alarm (Security Source NAT Rule Set) | **16125**

rule-session-count-alarm (Security Static NAT Rule Set) | **16126**

rule-set | **16128**

rule-set (Captive Portal Content Delivery) | **16129**

rule-set (Security Destination NAT) | **16131**

rule-set (Security Source NAT) | **16133**

rule-set (Security Static NAT) | **16135**

rule-set (Services AACL) | **16138**

rule-set (Services CoS) | **16139**

rule-set (Services CoS Next Gen Services) | **16140**

rule-set (Services IDS) | **16142**

rule-set (Services NAT) | **16143**

rule-set (Services Stateful Firewall) | **16144**

rule-set (Softwire) | **16146**

rule-set (Softwires Next Gen Services) | **16147**

rule-sets (CoS AppQoS) | **16149**

rule-sets (Security Application Firewall) | **16151**

rule (Softwire) | **16154**

rule (Source NAT Next Gen Services) | **16156**

rulebase-exempt | **16157**

rulebase-ips | **16159**

run | **16162**

run-length | **16163**

rx-buffers | **16164**

rx-connect-speed-when-equal (L2TP LAC) | **16166**

rx-enable | **16168**

rx-max-duration | **16169**

rx-window-size (L2TP) | **16171**

S | 16172

s-nssai | **16259**

s6a | **16260**

sa-hold-time (Protocols MSDP) | **16262**

sa-multicast (100-Gigabit Ethernet) | **16264**

sa-multicast (PTX Series Packet Transport Routers) | **16266**

sabit | **16267**

saegw | **16269**

saegw access-network-peers | **16270**

saegw control-plane-peers | **16272**

saegw-core-network-peers | **16276**

saegw system | **16278**

sak-rekey-interval | **16280**

saml | **16281**

saml-options | **16284**

sample | **16285**

sample-once | **16287**

sample-rate | **16288**

sample-type | **16290**

sampling | **16291**

sampling (Forwarding Options) | **16293**

sampling-instance | **16297**

sampling (Interfaces) | **16298**

sampling (Protocols MPLS) | **16300**

sanity-poll | **16302**

sap | **16303**

satellite (Junos Fusion Automatic Satellite Conversion) | **16305**

satellite (Junos Fusion Satellite Device Homing) | **16307**

satellite (Junos Fusion Satellite Software Upgrade Groups) | **16308**

satellite-management (Junos Fusion) | **16310**

satellite-policies | **16313**

save | **16315**

save (Configuration Mode) | **16317**

saved-core-context (System) | **16320**

saved-core-files (System) | **16321**

sbfd | **16323**

sbfd (protocols bfd) | **16324**

sbl | **16326**

sbl-default-server | **16327**

scale-mode | **16328**

scale-optimized | **16329**

scale-mode | **16331**

scan-extension | **16332**

scan-mode | **16333**

scan-options (Security Antivirus Avira Engine) | **16335**

scan-options (Security Antivirus Juniper Express Engine) | **16336**

scan-options (Security Antivirus Kaspersky Lab Engine) | **16338**

scan-options (Security Antivirus Sophos Engine) | **16340**

sccp | **16341**

scheduler | **16343**

scheduler (Dynamic Scheduler Maps) | **16344**

scheduler (Fabric Queues) | **16345**

scheduler (Interconnect Device) | **16347**

scheduler-map-chassis | **16348**

scheduler-map (CoS Virtual Channels) | **16350**

scheduler-map (Dynamic Traffic Shaping) | **16351**

scheduler-map (Fabric Queues) | **16353**

scheduler-map-forwarding-class-sets | **16354**

scheduler-map (Interfaces and Traffic-Control Profiles) | **16355**

scheduler-map (Interfaces and Traffic-Control Profiles) | **16357**

scheduler-maps (Definition) | **16358**

scheduler-maps (Dynamic CoS Definition) | **16360**

scheduler-name | **16361**

scheduler (Scheduler Map) | **16362**

scheduler (Security Policies) | **16364**

scheduler (System Security Profile) | **16366**

schedulers (CoS) | **16368**

schedulers (CoS) | **16369**

schedulers (Dynamic CoS Definition) | **16371**

schedulers (Interfaces) | **16373**

schedulers (Security Policies) | **16375**

schema | **16376**

schema-version (Flat-File Accounting Options) | **16378**

scope | **16379**

scope-policy | **16380**

scope (Security IDP Chain Attack) | **16382**

scope (Security IDP Custom Attack) | **16383**

screen | **16384**

screen (Security Zones) | **16390**

scripts | **16392**

sctp | **16396**

sdsl-bytes | **16398**

sdsl-overhead-adjust | **16400**

search | **16402**

search-filter | **16404**

secondary (Label Switched Path) | **16405**

secondary (Precision Time Protocol) | **16406**

secondary (Adaptive Services Interfaces) | **16408**

secondary connection (Identity Management Advanced Query) | **16409**

secondary-dns | **16412**

secondary-interface (Aggregated Inline Services) | **16413**

secondary (Link Services IQ PIC Interfaces) | **16414**

secondary (Protocols MPLS) | **16416**

secondary-server | **16418**

secondary-vtep-address | **16419**

secondary-wins | **16420**

secret (Access Profile) | **16422**

secret-key-timeout | **16423**

secret (RADIUS) | **16424**

secret (RADIUS Client) | **16426**

secret (RADIUS Server) | **16427**

secret (Tunnel Profile) | **16429**

secure | **16430**

secure-access-port | **16432**

secure-channel | **16435**

secure-nat-mapping | **16436**

secure-nat-mapping (Source NAT Next Gen Services) | **16438**

secure-neighbor-discovery | **16439**

secure-proxy | **16441**

secure-wire (System Security Profile) | **16443**

secured-port-block-allocation | **16444**

securid-server | **16447**

security | **16448**

security-association (MACSec) | **16450**

security-association | **16452**

security-association (IPSec) | **16454**

security-association (Junos-FIPS Software) | **16458**

security-association (Junos OS) | **16460**

security-intelligence (Services) | **16462**

security-intelligence | **16464**

security-intelligence (IDP) | **16465**

security-intelligence-policy | **16467**

security-intelligence(services) | **16469**

security-level | **16480**

security-level (Defining Access Privileges) | **16481**

security-level (Generating SNMP Notifications) | **16482**

security-log | **16484**

security-log-percent-full | **16485**

security-metadata-streaming | **16487**

security-mode (SRX) | **16492**

security-mode | **16494**

security-model (Access Privileges) | **16496**

security-model (Group) | **16497**

security-model (SNMP Notifications) | **16499**

security-name (Community String) | **16500**

security-name (Security Group) | **16501**

security-name (SNMP Notifications) | **16503**

security-option (IDS Screen Next Gen Services) | **16504**

security-package | **16505**

security-profile | **16507**

security-profile-resources | **16512**

security-service (Security Forwarding Options) | **16513**

security-to-group | **16515**

security-zone | **16516**

segment | **16519**

segment-list | **16521**

select | **16525**

select-profile | **16527**

selection-mode | **16528**

selective | **16529**

selective-vlan-switching | **16532**

send-acct-status-on-config-change (Access Profile) | **16534**

send (add-path) | **16535**

send-critical-event | **16539**

send-dyn-subscription-indicator (PCRF Partition) | **16540**

send-network-family-indicator (PCRF Partition) | **16542**

send-oam | **16544**

send-origin-state-id (Diameter Applications) | **16545**

send (Protocols RIP) | **16546**

send (Protocols RIPng) | **16548**

send-release-on-delete (DHCP Relay Agent) | **16550**

sender-based-rpf (MBGP MVPN) | **16552**

sendid-tlv | **16554**

sensor-based-stats | **16555**

sensor-based-stats (Junos Telemetry Interface) | **16558**

sensor-configuration | **16560**

sensor (Junos Telemetry Interface) | **16564**

separator | **16593**

seq-number-validated | **16594**

sequence-check-required | **16595**

sequence-number (Security IDP ICMP Headers) | **16597**

sequence-number (Security IDP TCP Headers) | **16598**

serial-number | **16599**

serial-number (Junos Fusion) | **16601**

serial-number (MX Series Virtual Chassis) | **16603**

serial-options | **16604**

server (Junos node slicing) | **16606**

server | **16608**

server-address | **16609**

server-address (dhcp-client) | **16611**

server-address (Hosted Services) | **16612**

server (Batch Commits) | **16613**

server-certificate (REST API) | **16615**

server-certificate (Services) | **16617**

server-certificate-subject | **16618**

server-cluster (Security Group VPN Server) | **16619**

server-connectivity | **16621**

server (DHCP and BOOTP Relay Agent) | **16623**

server (DNS, Port, and TFTP Service) | **16624**

server-duid-type (DHCP Local Server) | **16626**

server (Dynamic PPPoE) | **16627**

server (Easy EVPN LAG Configuration) | **16628**

server-group | **16633**

server (icap-redirect profile) | **16635**

server-id-override | **16637**

server-identifier | **16639**

server-inactivity-timeout | **16641**

server (Junos OS) | **16642**

server (Junos OS Evolved) | **16644**

server-member-communication (Security Group VPN Server) | **16646**

server (NTP) | **16648**

server (pcp) | **16651**

server-profile | **16654**

server-profile (Active Flow Monitoring) | **16655**

server (RADIUS Accounting) | **16656**

server (RADIUS Network Elements) | **16660**

server-reject-bridge-domain | server-reject-vlan | **16661**

server-response-time (DHCP Relay Agent) | **16664**

server (Security Antivirus) | **16665**

server (Security Group VPN) | **16667**

server (Security Sophos Engine Antivirus) | **16670**

server (Security Web Filtering) | **16672**

server (TACACS+ Accounting) | **16674**

servers | **16677**

servers (RADIUS) | **16678**

service | **16680**

service-agents (Analytics) | **16681**

service-context-id (OCS) | **16684**

service-deployment | **16686**

service-device-pool (L2TP) | **16687**

service-device-pools (Service Interfaces) | **16688**

service-device (Remote Device Management) | **16690**

service-domain | **16693**

service (Dynamic Profiles) | **16694**

service (Dynamic Service Sets) | **16695**

service-filter (Dynamic Service Sets) | **16697**

service-filter (Firewall) | **16699**

service-filter (Interfaces) | **16701**

service-groups | **16702**

service-id (Bridge Domains) | **16704**

service-id (Switch Options) | **16705**

service-interface (L2TP Processing) | **16706**

service-interface-pools | **16708**

service-interface (Services Interfaces) | **16709**

service-mode | **16711**

service-mode (mobile-edge SAEGW) | **16713**

service-name | **16714**

service-name-table | **16716**

service-name-tables | **16717**

service-order | **16719**

service-package | **16721**

service-pics | **16723**

service-port (UDP, Dynamic Flow Capture) | **16724**

service-port (Inline Services) | **16725**

service (PPPoE) | **16727**

service-profile (DHCP Local Server) | **16729**

service-profile (DHCP Relay Agent) | **16731**

service-profile (L2TP) | **16733**

service-profile (Static Subscribers) | **16735**

service (Security IDP Anomaly Attack) | **16736**

service (Security IDP Dynamic Attack Group) | **16737**

service (Service Accounting) | **16739**

service-set (Application-Aware Control Policy) | **16740**

service-set (Dynamic Service Sets) | **16742**

service-set (Interfaces) | **16744**

service-set-options | **16746**

service-set-options (Next Gen Services Services) | **16748**

service-set (Services) | **16750**

service-set (Subscriber-Aware) | **16754**

service-set (TDF Interface) | **16756**

service-tracking-action (Network Isolation Groups) | **16758**

service-tracking (Network Isolation Groups) | **16760**

service-type | **16761**

service-type (RFC 2544 Benchmarking Tests) | **16763**

service-type (MAC-VRF Routing Instances) | **16764**

service-provider (SAML) | **16767**

services (Captive Portal Content Delivery) | **16769**

services (CoS) | **16772**

services (Hierarchy) | **16773**

services (IDS) | **16774**

services (Interfaces) | **16776**

services (IPsec VPN) | **16777**

services (L2TP) | **16778**

services (L2TP System Logging) | **16780**

services-memory (resource-manager) | **16781**

services (NAT) | **16783**

services-offload | **16784**

services-offload (Security) | **16786**

services-options | **16787**

services-options (Next Gen Services Interfaces) | **16789**

services-redundancy-group | **16793**

services (Stateful Firewall) | **16799**

services (System Logging) | **16800**

services (system services) | **16802**

session-affinity | **16804**

session-close | **16806**

session-drop-hold-down | **16807**

session-establishment-hold-time | **16808**

session-expiry (MX Series in Enhanced LAN Mode) | **16809**

session-group | **16810**

session-id-change-limiter-indirect | **16813**

session-init | **16815**

session-limit | **16816**

session-limit (IDS MS-DPC) | **16817**

session-limit (IDS MS-MPC) | **16819**

session-limit-per-username (Access Profile) | **16821**

session-limit (Service Set Next Gen Services) | **16823**

session-limit (System Services) | **16824**

session-memory (resource-manager) | **16826**

session-mode | **16828**

session-offload | **16829**

session-options | **16831**

session-options (Access Profile) | **16835**

session (Origin Validation for BGP) | **16836**

session-persistence-scan | **16839**

session-pics | **16840**

session-pics (Diameter) | **16842**

session-protection | **16843**

session (Protocols LDP) | **16845**

session-scan | **16847**

session (Security IKE) | **16848**

session (System Services) | **16853**

session-timeout (Service Set Next Gen Services) | **16854**

session-update-interval | **16856**

session (Web Management) | **16857**

sessions-limit-group (L2TP) | **16858**

set | **16860**

set chassis cluster cluster-id node node-number reboot | **16861**

set chassis cluster disable reboot | **16863**

set chassis display message | **16865**

set chassis virtual-chassis | **16869**

set cli complete-on-space | **16870**

set cli directory | **16872**

set cli idle-timeout | **16874**

set cli logical-system | **16875**

set cli prompt | **16877**

set cli restart-on-upgrade | **16878**

set cli screen-length | **16880**

set cli screen-width | **16881**

set cli terminal | **16882**

set cli timestamp | **16884**

set date | **16885**

set date ntp | **16888**

set-dont-fragment-bit (Services IPsec VPN) | **16890**

set-dont-fragment-bit (Services Set) | **16891**

set services jinsightd subscribe health-monitor | **16892**

set services jinsightd traceoptions | **16894**

identity-provider settings (SAML) | **16895**

setup-protection | **16897**

severity (Event Policy) | **16898**

severity (Next Gen Services Service-Set Remote System Logging) | **16900**

severity (Security IDP Custom Attack) | **16901**

severity (Security IDP Dynamic Attack Group) | **16903**

severity (Security IDP IPS Rulebase) | **16904**

severity (Security Log) | **16906**

sflow | **16907**

sfm (Chassis) | **16912**

sfm (Chassis Redundancy) | **16913**

sfpplus | **16915**

sftp-backup (OCS Partition) | **16916**

sglimit | **16919**

sham-link | **16921**

sham-link-remote | **16923**

shaping | **16925**

shaping-rate (Traffic Control Profiles) | **16926**

shaping-rate (Applying to an Interface) | **16929**

shaping-rate (CoS Adaptive Shapers) | **16932**

shaping-rate (CoS Interfaces) | **16933**

shaping-rate (CoS Virtual Channels) | **16936**

shaping-rate (Dynamic Traffic Shaping and Scheduling) | **16937**

shaping-rate-excess-high | **16939**

shaping-rate-excess-low | **16941**

shaping-rate-excess-medium-high | **16943**

shaping-rate-excess-medium-low | **16945**

shaping-rate (Oversubscribing an Interface) | **16947**

shaping-rate-priority-high | **16950**

shaping-rate-priority-low | **16952**

shaping-rate-priority-medium | **16954**

shaping-rate-priority-medium-low | **16956**

shaping-rate-priority-strict-high | **16958**

shaping-rate (Schedulers) | **16960**

shared-bandwidth-policer (Configuring) | **16962**

shared-buffer | **16964**

shared-instance | **16967**

shared-key | **16968**

shared-name | **16970**

shared-scheduler | **16971**

shared-secret (RADIUS Snoop Segment) | **16973**

shdsl-options | **16974**

shell | **16975**

shellcode | **16977**

shmlog (Shared Memory Log) | **16978**

short-cycle-protection (DHCP Local Server and Relay Agent) | **16982**

short-cycle-protection (Static and Dynamic Subscribers) | **16984**

short-name-format | **16986**

short-reach-mode | **16988**

short-sequence | **16989**

shortcuts (Protocols OSPF) | **16991**

should-secure | **16992**

show | **16994**

show access-cac interface-set | **16995**

show access-security router-advertisement state | **16998**

show access-security router-advertisement statistics | **17001**

show access-security slaac-snooping binding | **17003**

show access-security slaac-snooping state | **17005**

show access-security slaac-snooping statistics | **17007**

show accounting pending-accounting-stops | **17011**

show accounting profile | **17016**

show accounting server statistics archival-transfer | **17023**

show agent sensors | **17024**

show amt statistics | **17033**

show amt summary | **17037**

show amt tunnel | **17039**

show analyzer | **17044**

show ancp cos | **17046**

show ancp neighbor | **17054**

show ancp statistics | **17066**

show ancp subscriber | **17072**

show ancp summary | **17083**

show ancp summary neighbor | **17086**

show ancp summary subscriber | **17089**

show app-engine crash | **17091**

show app-engine logs | **17093**

show arp | **17096**

show arp inspection statistics | **17103**

show as-path | **17105**

show as-path domain | **17111**

show as-path summary | **17113**

show auto-bandwidth | **17116**

show auto-configuration out-of-band | **17120**

show backup-selection | **17126**

show bfd seamless session | **17132**

show bfd session | **17134**

show bfd subscriber session | **17144**

show bgp bmp | **17151**

show bgp group | **17153**

show bgp group output-queues | **17166**

show bgp group traffic-statistics | **17171**

show bgp neighbor | **17175**

show bgp output-scheduler | **17216**

show bgp replication | **17218**

show bgp replication logical-system | **17223**

show bgp summary | **17227**

show bridge domain | **17237**

show bridge flood | **17240**

show bridge flood nexthops satellite | **17249**

show bridge flood satellite | **17253**

show bridge mac-ip-table | **17257**

show bridge mac-table | **17264**

show bridge satellite device | **17272**

show bridge statistics | **17276**

show captive-portal authentication-failed-users | **17281**

show captive-portal firewall | **17283**

show captive-portal interface | **17286**

show chassis adc | **17291**

show chassis adc (GNF) | **17295**

show chassis adc (Node Slicing) | **17297**

show chassis afeb | **17299**

show chassis alarms | **17302**

show chassis alarms | **17304**

show chassis beacon | **17334**

show chassis cfeb | **17337**

show chassis cluster control-interface status | **17341**

show chassis cluster control-plane statistics | **17351**

show chassis cluster data-plane interfaces | **17354**

show chassis cluster data-plane statistics | **17356**

show chassis cluster ethernet-switching interfaces | **17359**

show chassis cluster ethernet-switching status | **17361**

show chassis cluster information | **17364**

show chassis cluster information configuration-synchronization | **17378**

show chassis cluster information issu | **17381**

show chassis cluster interfaces | **17383**

show chassis cluster ip-monitoring status redundancy-group | **17392**

show chassis cluster port-peering | **17396**

show chassis cluster statistics | **17399**

show chassis cluster status | **17405**

show chassis craft-interface | **17410**

show chassis dedicated-ukern-cpu | **17418**

show chassis device-mode | **17419**

show chassis enhanced-temperature-thresholds | **17423**

show chassis environment | **17429**

show chassis environment adc | **17553**

show chassis environment cb (Security) | **17569**

show chassis environment cb | **17573**

show chassis environment ccg | **17613**

show chassis environment fan | **17616**

show chassis environment fpc | **17621**

show chassis environment fpm | **17705**

show chassis environment mcs | **17714**

show chassis environment monitored | **17718**

show chassis environment pcg | **17748**

show chassis environment pdu | **17752**

show chassis environment pem | **17760**

show chassis environment power-supply-unit | **17782**

show chassis environment psm | **17784**

show chassis environment psu | **17798**

show chassis environment routing-engine | **17800**

show chassis environment scg | **17811**

show chassis environment (Security) | **17818**

show chassis environment sfb | **17825**

show chassis environment sfm | **17844**

show chassis environment sib | **17849**

show chassis errors active | **17895**

show chassis ethernet-switch (Security) | **17901**

show chassis ethernet-switch | **17908**

show chassis ethernet-switch interconnect-device cb | **17960**

show chassis ethernet-switch interconnect-device fpc | **17985**

show chassis fabric connectivity | **18022**

show chassis fabric degradation | **18031**

show chassis fabric degradation actions | **18033**

show chassis fabric degraded-fabric-reachability | **18035**

show chassis fabric destinations | **18037**

show chassis fabric device | **18066**

show chassis fabric errors | **18068**

show chassis fabric faults recovery-actions | **18075**

show chassis fabric feb | **18078**

show chassis fabric fpcs | **18079**

show chassis fabric map | **18091**

show chassis fabric mesh | **18101**

show chassis fabric plane (Security) | **18104**

show chassis fabric plane | **18113**

show chassis fabric plane-location (Security) | **18154**

show chassis fabric plane-location | **18156**

show chassis fabric reachability | **18168**

show chassis fabric redundancy-mode | **18181**

show chassis fabric sibs | **18183**

show chassis fabric stream-info | **18203**

show chassis fabric summary (Security) | **18209**

show chassis fabric summary | **18212**

show chassis fabric topology | **18220**

show chassis fabric unreachable-destinations | **18257**

show chassis fan | **18260**

show chassis feb | **18278**

show chassis firmware | **18283**

show chassis forwarding | **18304**

show chassis forwarding-options | **18306**

show chassis fpc | **18311**

show chassis fpc errors | **18387**

show chassis fpc ether-types | **18392**

show chassis fpc-feb-connectivity | **18397**

show chassis fpc global pfe-id | **18400**

show chassis fpc (GNF) | **18404**

show chassis fpc (Node Slicing) | **18408**

show chassis fpc pic-status (GNF) | **18411**

show chassis fpc pic-status (Node Slicing) | **18413**

show chassis fpc (View) | **18415**

show chassis gnf | **18427**

show chassis hardware | **18430**

show chassis hardware (GNF) | **18453**

show chassis hardware (View) | **18457**

show chassis high-availability data-plane statistics | **18476**

show chassis high-availability information | **18482**

show chassis high-availability peer-info | **18492**

show chassis high-availability prefix-srgid-table | **18495**

show chassis high-availability services-redundancy-group | **18497**

show chassis in-service-upgrade | **18507**

show chassis jnu satellite | **18513**

show chassis lcd | **18515**

show chassis led | **18534**

show chassis led satellite | **18551**

show chassis location | **18557**

show chassis mac-addresses | **18563**

show chassis network-services | **18571**

show chassis network-slices | **18575**

show chassis network-slices fpcs | **18579**

show chassis nonstop-upgrade | **18581**

show chassis nonstop-upgrade node-group | **18585**

show chassis oss-map | **18587**

show chassis pic | **18588**

show chassis pic-mode | **18628**

show chassis pic (Security) | **18630**

show chassis power | **18633**

show chassis power (Security) | **18676**

show chassis power-budget-statistics | **18681**

show chassis power sequence | **18688**

show chassis psd | **18690**

show chassis realtime-ukern-thread | **18692**

show chassis redundancy feb | **18694**

show chassis redundant-power-system | **18698**

show chassis reswatch | **18701**

show chassis routing-engine | **18703**

show chassis routing-engine (View) | **18720**

show chassis satellite | **18727**

show chassis satellite-cluster | **18743**

show chassis satellite-cluster route | **18751**

show chassis satellite-cluster statistics | **18754**

show chassis satellite extended-port | **18759**

show chassis satellite interface | **18765**

show chassis satellite neighbor | **18772**

show chassis satellite power-budget-statistics | **18779**

show chassis satellite redundancy-group | **18783**

show chassis satellite redundancy-group devices | **18787**

show chassis satellite software | **18790**

show chassis satellite statistics | **18794**

show chassis satellite unprovision | **18801**

show chassis satellite upgrade-group | **18809**

show chassis scb | **18813**

show chassis scb-ha-port | **18816**

show chassis security uefi variables | **18821**

show chassis sfb | **18823**

show chassis sfb errors | **18828**

show chassis sfm | **18831**

show chassis sibs | **18835**

show chassis spmb | **18853**

show chassis spmb sibs | **18867**

show chassis ssb | **18876**

show chassis synchronization | **18879**

show chassis synchronization (MX Series Routers) | **18894**

show chassis temperature-thresholds (Security) | **18911**

show chassis temperature-thresholds | **18914**

show chassis usb storage | **18945**

show chassis zones | **18947**

show class-of-service | **18959**

show class-of-service adjustment-control-profile | **18965**

show class-of-service application-traffic-control counter | **18968**

show class-of-service application-traffic-control rate-limiters | **18974**

show class-of-service application-traffic-control statistics rate-limiter | **18986**

show class-of-service application-traffic-control statistics rule | **18990**

show class-of-service classifier | **18993**

show class-of-service code-point-aliases | **18999**

show class-of-service congestion-notification | **19001**

show class-of-service dedicated-buffer-profile | **19007**

show class-of-service drop-profile | **19009**

show class-of-service fabric scheduler-map | **19011**

show class-of-service fabric statistics | **19014**

show class-of-service forwarding-class | **19019**

show class-of-service forwarding-class-set | **19024**

show class-of-service forwarding-table | **19026**

show class-of-service forwarding-table classifier | **19031**

show class-of-service forwarding-table classifier mapping | **19033**

show class-of-service forwarding-table drop-profile | **19035**

show class-of-service forwarding-table fabric scheduler-map | **19038**

show class-of-service forwarding-table rewrite-rule | **19040**

show class-of-service forwarding-table rewrite-rule mapping | **19042**

show class-of-service forwarding-table scheduler-map | **19044**

show class-of-service forwarding-table slice | **19047**

show class-of-service forwarding-table slice mapping | **19049**

show class-of-service forwarding-table traffic-class-map | **19051**

show class-of-service fragmentation-map | **19054**

show class-of-service interface | **19056**

show class-of-service interface-set | **19100**

show class-of-service l2tp-session | **19103**

show class-of-service loss-priority-rewrite | **19106**

show class-of-service monitoring-profile | **19108**

show class-of-service multi-destination | **19110**

show class-of-service policy-map | **19112**

show class-of-service rewrite-rule | **19114**

show class-of-service routing-instance | **19117**

show class-of-service scheduler-hierarchy interface | **19119**

show class-of-service scheduler-hierarchy interface-set | **19122**

show class-of-service scheduler-map | **19125**

show class-of-service scheduler-map-forwarding-class-sets | **19131**

show class-of-service scheduler-resources fpc | **19134**

show class-of-service shared-buffer | **19137**

show class-of-service slice <slice_name> interface <interface_name> | **19140**

show class-of-service traffic-class-map | **19142**

show class-of-service traffic-control-profile | **19145**

show class-of-service translation-table | **19151**

show cli | **19157**

show cli authorization | **19159**

show cli directory | **19164**

show cli history | **19166**

show configuration | **19167**

show configuration chassis cluster traceoptions | **19171**

show configuration services analytics sensor ddos | **19173**

show configuration smtp | **19174**

show connections | **19176**

show cos halp flow queue-stats | **19181**

show database-replication statistics | **19183**

show database-replication summary | **19185**

show dcbx | **19189**

show dcbx neighbors | **19191**

show ddos-protection protocols | **19222**

show ddos-protection protocols culprit-flows | **19239**

show ddos-protection protocols flow-detection | **19249**

show ddos-protection protocols isis | **19254**

show ddos-protection protocols parameters | **19265**

show ddos-protection protocols statistics | **19276**

show ddos-protection protocols violations | **19292**

show ddos-protection statistics | **19295**

show ddos-protection version | **19298**

show dhcp client binding | **19300**

show dhcp client statistics | **19305**

show dhcp relay active-leasequery | **19309**

show dhcp relay binding (Security) | **19316**

show dhcp relay binding | **19319**

show dhcp relay lockout-entries | **19328**

show dhcp relay statistics (Security) | **19331**

show dhcp relay statistics | **19334**

show dhcp-security arp inspection statistics | **19342**

show dhcp-security binding | **19344**

show dhcp-security binding ip-source-guard | **19348**

show dhcp-security ipv6 binding | **19351**

show dhcp-security ipv6 statistics | **19355**

show dhcp-security neighbor-discovery-inspection statistics | **19359**

show dhcp server active-leasequery | **19362**

show dhcp server active-leasequery statistics | **19364**

show dhcp server binding | **19367**

show dhcp server lockout-entries | **19377**

show dhcp server statistics | **19381**

show dhcp snooping binding | **19387**

show dhcp snooping statistics | **19390**

show dhcpv6 client binding | **19392**

show dhcpv6 client statistics | **19395**

show dhcpv6 relay active-leasequery | **19398**

show dhcpv6 relay binding | **19404**

show dhcpv6 relay lockout-entries | **19417**

show dhcpv6 relay statistics | **19420**

show dhcpv6 server active-leasequery | **19426**

show dhcpv6 server active-leasequery statistics | **19428**

show dhcpv6 server binding | **19431**

show dhcpv6 server binding (View) | **19440**

show dhcpv6 server lockout-entries | **19446**

show dhcpv6 server statistics | **19450**

show dhcpv6 server statistics (View) | **19456**

show dhcpv6 snooping binding | **19460**

show dhcpv6 snooping statistics | **19463**

show diagnostics tdr | **19465**

show diameter | **19472**

show diameter function | **19482**

show diameter function statistics | **19488**

show diameter instance | **19493**

show diameter network-element | **19495**

show diameter network-element map | **19500**

show diameter peer | **19504**

show diameter peer map | **19510**

show diameter peer statistics | **19514**

show diameter route | **19519**

show | display | **19522**

show | display commit-scripts | **19527**

show | display inheritance | **19529**

show | display json | **19532**

show | display omit | **19535**

show | display rfc5952 | **19537**

show | display set | **19538**

show | display set relative | **19540**

show | display translation-scripts | **19542**

show | display xml | **19546**

show dot1x | **19548**

show dot1x accounting attribute | **19557**

show dot1x authentication-bypassed-users | **19561**

show dot1x authentication-failed-users | **19562**

show dot1x firewall | **19564**

show dot1x interface | **19567**

show dot1x static-mac-address | **19577**

show dot1x statistics | **19580**

show dual-phase-bootup-status | **19583**

show dvmrp interfaces | **19584**

show dvmrp neighbors | **19587**

show dvmrp prefix | **19589**

show dvmrp prunes | **19592**

show dynamic-profile session | **19595**

show dynamic-tunnels database | **19601**

show dynamic-tunnels pfe-tunnel-localization | **19609**

show edge-virtual-bridging | **19616**

show ephemeral-configuration | **19620**

show ethernet-switching evpn nd-statistics | **19623**

show ethernet-switching evpn nd-table | **19626**

show ethernet-switching flood | **19630**

show ethernet-switching interface | **19636**

show ethernet-switching interfaces | **19641**

show ethernet-switching layer2-protocol-tunneling interface | **19652**

show ethernet-switching layer2-protocol-tunneling statistics | **19654**

show ethernet-switching layer2-protocol-tunneling vlan | **19658**

show ethernet-switching mac-ip-table | **19660**

show ethernet-switching mac-learning-log | **19670**

show ethernet-switching mgrp-policy | **19676**

show ethernet-switching redundancy-groups | **19682**

show ethernet-switching statistics | **19688**

show ethernet-switching statistics aging | **19692**

show ethernet-switching statistics mac-learning | **19694**

show ethernet-switching table | **19700**

show ethernet-switching vxlan-tunnel-end-point esi | **19724**

show ethernet-switching vxlan-tunnel-end-point remote | **19728**

show ethernet-switching vxlan-tunnel-end-point source | **19730**

show ethernet-switching vxlan-tunnel-end-point svlbnh | **19733**

show evpn arp-table | **19736**

show evpn database | **19740**

show evpn flood | **19748**

show evpn gbp-src-tag | **19749**

show evpn igmp-snooping database | **19752**

show evpn igmp-snooping proxy | **19756**

show evpn instance | **19761**

show evpn ip-prefix-database | **19779**

show evpn l3-context | **19790**

show evpn mac-ip-table | **19794**

show evpn mac-table | **19798**

show evpn mld-snooping database | **19800**

show evpn multicast-snooping assisted-replication multihomed-peers | **19804**

show evpn multicast-snooping assisted-replication next-hops | **19807**

show evpn multicast-snooping assisted-replication replicators | **19811**

show evpn multicast-snooping status | **19815**

show evpn nd-table | **19818**

show evpn oism | **19822**

show evpn oism peg-df-status | **19826**

show evpn oism spmsi-ad | **19829**

show evpn p2mp | **19833**

show evpn peer-gateway-macs | **19836**

show evpn prefix | **19838**

show evpn vpws-instance | **19840**

show extensible-subscriber-services counters | **19849**

show extensible-subscriber-services debug-information | **19851**

show extensible-subscriber-services dictionary | **19853**

show extensible-subscriber-services dictionary attributes | **19859**

show extensible-subscriber-services dictionary services | **19863**

show extensible-subscriber-services service | **19867**

show extensible-subscriber-services sessions | **19869**

show extension-service request-response clients | **19871**

show extension-service request-response servers | **19874**

show extension-service status | **19876**

show fabric administration inventory | **19878**

show fabric administration inventory director-group status | **19885**

show fabric administration inventory infrastructure | **19893**

show fabric administration inventory interconnect-devices | **19898**

show fabric administration inventory node-devices | **19900**

show fabric administration inventory node-groups | **19903**

show fabric administration system mac-pool | **19906**

show fabric inventory | **19908**

show fabric session-host | **19912**

show fib-local-accounting ip | **19914**

show fib-streaming | **19915**

show fib-streaming route-tables | **19921**

show fib-streaming routes | **19924**

show fib-streaming state | **19927**

show fibre-channel fabric | **19929**

show fibre-channel fc2 sessions | **19932**

show fibre-channel fc2 statistics | **19936**

show fibre-channel fip | **19938**

show fibre-channel fip enode | **19946**

show fibre-channel fip fabric | **19951**

show fibre-channel fip fcf | **19956**

show fibre-channel fip interface | **19961**

show fibre-channel fip statistics | **19966**

show fibre-channel flogi fport | **19971**

show fibre-channel flogi nport | **19973**

show fibre-channel flogi statistics | **19977**

show fibre-channel interfaces | **19980**

show fibre-channel next-hops | **19985**

show fibre-channel proxy fabric-state | **19987**

show fibre-channel proxy login-table | **19993**

show fibre-channel proxy np-port | **19997**

show fibre-channel proxy statistics | **20003**

show fibre-channel routes | **20006**

show fip snooping | **20008**

show fip snooping enode | **20015**

show fip snooping fcf | **20021**

show fip snooping interface | **20024**

show fip snooping statistics | **20029**

show fip snooping vlan | **20034**

show fip vlan-discovery | **20040**

show firewall (QFX) | **20043**

show firewall | **20048**

show firewall filter version | **20060**

show firewall log | **20062**

show firewall policer | **20066**

show firewall prefix-action-stats | **20068**

show firewall templates-in-use | **20071**

show firewall (View) | **20073**

show flexible-tunnels profiles | **20077**

show flight-recorder status | **20083**

show forwarding-options enhanced-hash-key | **20086**

show forwarding-options hash-key | **20095**

show forwarding-options hyper-mode | **20100**

show forwarding-options load-balance | **20102**

show forwarding-options load-balance ecmp|trunk | **20107**

show forwarding-options next-hop-group | **20111**

show forwarding-options port-mirroring | **20116**

show groups junos-defaults | **20120**

show gtp tunnels | **20122**

show hfrr profiles | **20128**

show host | **20131**

show iccp | **20133**

show igmp group | **20135**

show igmp interface | **20142**

show igmp snooping data-forwarding | **20148**

show igmp snooping evpn status | **20153**

show igmp snooping interface | **20157**

show igmp snooping membership | **20165**

show igmp-snooping membership | **20173**

show igmp snooping options | **20179**

show igmp-snooping route | **20181**

show igmp-snooping route (View) | **20185**

show igmp snooping statistics | **20187**

show igmp-snooping statistics | **20196**

show igmp-snooping vlans | **20198**

show igmp statistics | **20203**

show ike security-associations | **20210**

show ilmi statistics | **20217**

show ingress-replication mvpn | **20221**

show interfaces (NFX150) | **20223**

show interfaces | **20227**

show interfaces (Abstracted Fabric) | **20371**

show interfaces (Adaptive Services) | **20401**

show interfaces (Aggregated Ethernet) | **20411**

show interfaces anchor-group (Aggregated Packet Forwarding Engine) | **20437**

show interfaces (ATM) | **20442**

show interfaces demux0 (Demux Interfaces) | **20497**

show interfaces destination-class | **20517**

show interfaces detail | **20521**

show interfaces diagnostics optics | **20526**

show interfaces diagnostics optics | **20535**

show interfaces diagnostics optics (Security) | **20569**

show interfaces diagnostics optics-applications | **20576**

show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port) | **20580**

show interfaces diagnostics optics-profile | **20645**

show interfaces (Discard) | **20651**

show interfaces (Dynamic Flow Capture) | **20660**

show interfaces (Encryption) | **20667**

show interfaces extensive | **20676**

show interfaces extensive satellite-device | **20747**

show interfaces fabric | **20774**

show interfaces (far-end-interval) | **20803**

show interfaces (Fast Ethernet) | **20805**

show interfaces filters | **20832**

show interfaces (Flow Collector) | **20835**

show interfaces (Flow Monitoring) | **20844**

show interfaces flow-statistics | **20852**

show interfaces forwarding-class-counters | **20859**

show interfaces fti | **20866**

show interfaces ge | **20881**

show interfaces (GRE) | **20900**

show interfaces interface-set (Ethernet Interface Set) | **20913**

show interfaces interface-set (PS Interface Set) | **20920**

show interfaces interface-set queue | **20922**

show interfaces interval | **20934**

show interfaces (IP-over-IP) | **20938**

show interfaces irb | **20945**

show interfaces l2-routing-instance | **20957**

show interfaces lb-stats af | **20959**

show interfaces (Link Services) | **20965**

show interfaces (Link Services IQ) | **20987**

show interfaces load-balancing (Aggregated Multiservices) | **21024**

show interfaces (Logical Tunnel) | **21029**

show interfaces (Loopback) | **21040**

show interfaces lsi (Label-Switched Interface) | **21051**

show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet) | **21056**

show interfaces mac-database | **21080**

show interfaces mc-ae (MX Series) | **21090**

show interfaces mc-ae | **21098**

show interfaces me0 | **21101**

show interfaces media | **21113**

show interfaces (Multicast Tunnel) | **21117**

show interfaces (Multilink Services) | **21124**

show interfaces (PIM) | **21137**

show interfaces policers | **21143**

show interfaces port-profile | **21146**

show interfaces (PPPoE) | **21154**

show interfaces prbs-stats | **21170**

show interfaces priority-group | **21172**

show interfaces ps0 (Pseudowire Subscriber Interfaces) | **21175**

show interfaces (PTX Series Packet Transport Routers) | **21182**

show interfaces queue | **21208**

show interfaces queue fabric | **21257**

show interfaces redundancy | **21288**

show interfaces (Redundant Adaptive Services) | **21292**

show interfaces (Redundant Link Services IQ) | **21295**

show interfaces routing | **21316**

show interfaces routing-instance | **21325**

show interfaces routing summary | **21328**

show interfaces satellite-device | **21332**

show interfaces (Serial) | **21339**

show interfaces smart-sfp-defects | **21359**

show interfaces smart-sfp-statistics | **21372**

show interfaces snmp-index | **21389**

show interfaces source-class | **21391**

show interfaces (SRX Series) | **21395**

show interfaces statistics | **21442**

show interfaces statistics fabric | **21461**

show interfaces statistics st0 | **21488**

show interfaces summary | **21490**

show interfaces swfabx | **21492**

show interfaces targeting (Aggregated Ethernet for Subscriber Management) | **21494**

show interfaces terse | **21497**

show interfaces terse satellite-device | **21504**

show interfaces terse zone | **21508**

show interfaces transport pm | **21509**

show interfaces vcp | **21524**

show interfaces (View Aggregated Ethernet) | **21552**

show interfaces (Virtual Loopback Tunnel) | **21567**

show interfaces vlan | **21575**

show interfaces voq | **21593**

show interfaces xe | **21612**

show interfaces xle | **21639**

show ip-source-guard | **21666**

show ipsec certificates | **21668**

show ipsec redundancy | **21672**

show ipsec security-associations | **21675**

show ipv6 neighbors | **21680**

show ipv6 neighbors (SRX Series) | **21683**

show ipv6 router-advertisement | **21685**

show ipv6-source-guard | **21692**

show isis adjacency | **21694**

show isis adjacency holddown | **21702**

show isis authentication | **21706**

show isis backup coverage | **21709**

show isis backup label-switched-path | **21711**

show isis backup spf results | **21714**

show isis context-identifier | **21720**

show isis database | **21722**

show isis flex-algorithm | **21732**

show isis hostname | **21736**

show isis interface | **21738**

show isis interface-group | **21746**

show isis layer2-map | **21748**

show isis overview | **21751**

show isis purge log | **21759**

show isis route | **21762**

show isis route download priority | **21768**

show isis spf | **21772**

show isis spring interface traffic-statistics | **21779**

show isis spring sensor info | **21780**

show isis statistics | **21783**

show l2-learning global-information | **21787**

show l2-learning global-mac-count | **21790**

show l2-learning instance | **21791**

show l2-learning interface | **21794**

show l2-learning mac-move-buffer | **21796**

show l2-learning redundancy-groups | **21798**

show l2circuit connections | **21805**

show l2pt interface | **21816**

show l2vpn connections | **21819**

show lacp interfaces | **21829**

show lacp interfaces (View) | **21836**

show lacp statistics interfaces (View) | **21843**

show lacp timeouts | **21845**

show ldp database | **21848**

show ldp database-label-requests | **21861**

show ldp fec-filters | **21864**

show ldp interface | **21866**

show ldp neighbor | **21868**

show ldp overview | **21871**

show ldp p2mp fec | **21879**

show ldp p2mp path | **21881**

show ldp p2mp tunnel | **21884**

show ldp path | **21886**

show ldp rib-groups | **21888**

show ldp route | **21890**

show ldp session | **21901**

show ldp statistics | **21912**

show ldp traffic-statistics | **21918**

show link-management | **21923**

show link-management peer | **21928**

show link-management routing | **21931**

show link-management statistics | **21935**

show link-management te-link | **21939**

show lldp | **21942**

show lldp local-information | **21952**

show lldp neighbors | **21956**

show lldp neighbors-vlan-name-tlv-list interface | **21973**

show lldp remote-global-statistics | **21974**

show lldp statistics | **21977**

show log | **21980**

show loop-detect enhanced interface | **21987**

show mac-rewrite interface | **21993**

show mac-vrf forwarding flood | **21996**

show mac-vrf forwarding flood-group | **22000**

show mac-vrf forwarding global-information | **22005**

show mac-vrf forwarding global-mac-count | **22008**

show mac-vrf forwarding global-mac-ip-count | **22010**

show mac-vrf forwarding instance | **22011**

show mac-vrf forwarding instance-mapping | **22018**

show mac-vrf forwarding interface | **22020**

show mac-vrf forwarding mac-ip-table | **22025**

show mac-vrf forwarding mac-table | **22031**

show mac-vrf forwarding mgrp-policy | **22035**

show mac-vrf forwarding statistics | **22042**

show mac-vrf forwarding vlans | **22046**

show mac-vrf forwarding vxlan-tunnel-end-point esi | **22050**

show mac-vrf forwarding vxlan-tunnel-end-point remote | **22054**

show mac-vrf forwarding vxlan-tunnel-end-point svlbnh | **22063**

show mld group | **22066**

show mld interface | **22074**

show mld snooping evpn database | **22081**

show mld snooping interface | **22085**

show mld snooping membership | **22090**

show mld-snooping route | **22094**

show mld snooping statistics | **22098**

show mld-snooping vlans | **22100**

show mld statistics | **22104**

show modem wireless firmware | **22111**

show modem wireless network | **22115**

show modem wireless profiles | **22120**

show monitor security flow | **22122**

show mpls abstract-hop-membership | **22125**

show mpls admin-groups | **22128**

show mpls association | **22130**

show mpls call-admission-control | **22132**

show mpls container-lsp | **22136**

show mpls context-identifer | **22147**

show mpls correlation label | **22150**

show mpls correlation nexthop-id | **22151**

show mpls cspf | **22153**

show mpls diffserv-te | **22156**

show mpls egress-protection | **22159**

show mpls interface | **22162**

show mpls label usage | **22165**

show mpls label usage label-range | **22169**

show mpls lsp | **22172**

show mpls lsp abstract-computation | **22203**

show mpls lsp autobandwidth | **22206**

show mpls path | **22210**

show mpls srlg | **22212**

show mpls static-lsp | **22214**

show mpls tunnel manager statistics | **22219**

show msdp | **22222**

show msdp source | 22225

show msdp source-active | 22228

show msdp statistics | 22232

show multi-chassis mc-lag configuration-consistency | 22237

show multi-chassis mc-lag configuration-consistency global-config | 22243

show multi-chassis mc-lag configuration-consistency icl-config | 22246

show multi-chassis mc-lag configuration-consistency list-of-parameters | 22249

show multi-chassis mc-lag configuration-consistency mcae-config | 22262

show multi-chassis mc-lag configuration-consistency vlan-config | 22267

show multi-chassis mc-lag configuration-consistency vrrp-config | 22270

show multicast backup-pe-groups | 22273

show multicast ecid-mapping satellite | 22276

show multicast flow-map | 22280

show multicast forwarding-cache statistics | 22283

show multicast interface | 22286

show multicast mrinfo | 22290

show multicast next-hops | 22292

show multicast next-hops satellite | 22298

show multicast pim-to-igmp-proxy | 22301

show multicast pim-to-mld-proxy | 22304

show multicast route | 22306

show multicast rpf | 22323

show multicast scope | 22328

show multicast sessions | 22331

show multicast snooping next-hops | 22335

show multicast snooping next-hops satellite | 22338

show multicast snooping route | **22344**

show multicast snooping route satellite | **22350**

show multicast statistics | **22356**

show multicast statistics satellite | **22363**

show multicast summary satellite | **22365**

show multicast usage | **22366**

show mvpn c-multicast | **22370**

show mvpn instance | **22375**

show mvpn neighbor | **22381**

show mvpn suppressed | **22388**

show mvrp | **22390**

show mvrp applicant-state | **22394**

show mvrp dynamic-vlan-memberships | **22398**

show mvrp interface | **22401**

show mvrp registration-state | **22403**

show mvrp statistics | **22406**

show neighbor-discovery-inspection statistics | **22413**

show network-access aaa accounting | **22415**

show network-access aaa radius-servers | **22417**

show network-access aaa statistics | **22428**

show network-access aaa statistics accounting | **22447**

show network-access aaa statistics authentication | **22449**

show network-access aaa statistics dynamic-requests | **22453**

show network-access aaa statistics pending-accounting-stops | **22455**

show network-access aaa statistics preauthentication | **22456**

show network-access aaa statistics re-authentication | **22459**

show network-access aaa subscribers | **22461**

show network-access aaa subscribers session-id | **22468**

show network-access aaa terminate-code | **22479**

show network-access address-assignment pool | **22486**

show network-access address-assignment pool (View) | **22488**

show network-access domain-map | **22490**

show network-access gx-plus | **22492**

show network access nas-5g | **22509**

show network-access nasreq statistics | **22514**

show network-access ocs | **22520**

show network-access pcrf | **22523**

show network-access radsec local-certificate | **22528**

show network-access radsec state | **22531**

show network-access radsec statistics | **22535**

show network-access requests pending | **22538**

show network-access requests statistics | **22542**

show network-access s6a | **22544**

show network-access securid-node-secret-file | **22547**

show network-agent statistics | **22549**

show network-access aaa saml assertion-cache | **22560**

show nlsd mode | **22562**

show node reboot | **22564**

show node statistics | **22565**

show nonstop-routing | **22571**

show nsagentd | **22575**

show ntp associations | **22577**

show ntp status | **22581**

show oam ethernet connectivity-fault-management delay-statistics | **22585**

show oam ethernet connectivity-fault-management forwarding-state | **22590**

show oam ethernet connectivity-fault-management interfaces | **22596**

show oam ethernet connectivity-fault-management linktrace path-database | **22611**

show oam ethernet connectivity-fault-management loss-statistics | **22615**

show oam ethernet connectivity-fault-management mep-database | **22620**

show oam ethernet connectivity-fault-management mep-statistics | **22637**

show oam ethernet connectivity-fault-management mip | **22654**

show oam ethernet connectivity-fault-management path-database | **22656**

show oam ethernet connectivity-fault-management policer | **22660**

show oam ethernet connectivity-fault-management sla-iterator-statistics | **22664**

show oam ethernet connectivity-fault-management synthetic-loss-statistics | **22677**

show oam ethernet evc | **22681**

show oam ethernet fnp interface | **22684**

show oam ethernet fnp messages | **22686**

show oam ethernet fnp status | **22688**

show oam ethernet link-fault-management | **22691**

show oam ethernet lmi | **22701**

show oam ethernet lmi statistics | **22704**

show oam fabric flow-specification | **22707**

show oam fabric interfaces | **22711**

show openflow capability | **22713**

show openflow controller | **22722**

show openflow filters | **22726**

show openflow flows | **22730**

show openflow groups | **22736**

show openflow interfaces | **22740**

show openflow statistics flows | **22745**

show openflow statistics groups | **22750**

show openflow statistics interfaces | **22753**

show openflow statistics packet | **22758**

show openflow statistics queue | **22760**

show openflow statistics summary | **22764**

show openflow statistics tables | **22766**

show openflow summary | **22771**

show openflow switch | **22773**

show ospf context-identifier | **22776**

show ospf database | **22780**

show ospf flex-algorithm | **22791**

show (ospf | ospf3) backup coverage | **22800**

show (ospf | ospf3) backup lsp | **22804**

show (ospf | ospf3) backup neighbor | **22807**

show (ospf | ospf3) backup spf | **22809**

show (ospf | ospf3) bgp-orr | **22813**

show (ospf | ospf3) interface | **22817**

show (ospf | ospf3) io-statistics | **22829**

show (ospf | ospf3) log | **22831**

show (ospf | ospf3) neighbor | **22835**

show (ospf | ospf3) overview | **22843**

show (ospf | ospf3) route | **22851**

show (ospf | ospf3) statistics | **22859**

show ospf spring prefix sid map | **22864**

show ospf spring sid database | **22866**

show ospf3 database | **22869**

show ovsdb commit failures | **22881**

show ovsdb controller | **22884**

show ovsdb interface | **22887**

show ovsdb logical-switch | **22890**

show ovsdb mac | **22893**

show ovsdb statistics interface | **22899**

show ovsdb tunnels | **22901**

show ovsdb virtual-tunnel-end-point | **22905**

show passive-monitoring error | **22908**

show passive-monitoring flow | **22911**

show passive-monitoring memory | **22915**

show passive-monitoring status | **22917**

show passive-monitoring usage | **22919**

show path-computation-client active-pce | **22922**

show path-computation-client lsp | **22927**

show path-computation-client statistics | **22936**

show path-computation-client status | **22944**

show performance-monitoring mpls lsp | **22947**

show pfe cfeb | **22955**

show pfe data | **22961**

show pfe feb | **22974**

show pfe filter hw summary | **22983**

show pfe fpc | **22986**

show pfe next-hop | 23013

show pfe route | 23019

show pfe sFlow egress-mcast-stats | 23033

show pfe ssb | 23036

show pfe statistics bridge | 23045

show pfe terse | 23051

show pfe version | 23054

show pim bidirectional df-election | 23056

show pim bidirectional df-election interface | 23060

show pim bootstrap | 23064

show pim interfaces | 23066

show pim join | 23070

show pim mdt | 23093

show pim mdt data-mdt-joins | 23099

show pim mdt data-mdt-limit | 23102

show pim mvpn | 23104

show pim neighbors | 23106

show pim rps | 23113

show pim snooping interfaces | 23124

show pim snooping join | 23128

show pim snooping neighbors | 23133

show pim snooping statistics | 23141

show pim source | 23147

show pim statistics | 23152

show poe controller | 23171

show poe controller (View) | 23175

show poe interface | **23177**

show poe notification-control | **23187**

show poe telemetries | **23189**

show policer (EX) | **23192**

show policer | **23194**

show policy | **23199**

show policy conditions | **23203**

show policy damping | **23206**

show ppp address-pool | **23209**

show ppp interface | **23211**

show ppp statistics | **23229**

show ppp summary | **23239**

show pppoe interfaces (Security) | **23241**

show pppoe interfaces | **23246**

show pppoe lockout | **23252**

show pppoe lockout atm-identifier | **23256**

show pppoe lockout vlan-identifier | **23260**

show pppoe service-name-tables | **23265**

show pppoe sessions | **23269**

show pppoe statistics | **23271**

show pppoe underlying-interfaces | **23275**

show pppoe version | **23285**

show programmable-rpd clients | **23287**

show programmable-rpd clients route-summary | **23289**

show programmable-rpd statistics | **23292**

show protection-group ethernet-ring aps | **23294**

show protection-group ethernet-ring configuration | **23300**

show protection-group ethernet-ring data-channel | **23309**

show protection-group ethernet-ring flush-info | **23312**

show protection-group ethernet-ring interface | **23315**

show protection-group ethernet-ring node-state | **23320**

show protection-group ethernet-ring statistics | **23328**

show protection-group ethernet-ring vlan | **23335**

show ptp aggregated-ethernet interfaces | **23342**

show ptp all master clock | **23344**

show ptp clock | **23348**

show ptp global-information | **23354**

show ptp hybrid | **23369**

show ptp lock-status | **23372**

show ptp master | **23377**

show ptp passive-port-monitor-status | **23383**

show ptp path-trace detail | **23389**

show ptp phy-timestamping-interfaces | **23390**

show ptp port | **23392**

show ptp quality-level-mapping | **23400**

show ptp slave | **23401**

show ptp stateful detail | **23406**

show ptp statistics | **23408**

show redundant-power-system led | **23414**

show redundant-power-system multi-backup | **23418**

show redundant-power-system network | **23419**

show redundant-power-system power-supply | **23420**

show redundant-power-system status | **23423**

show redundant-power-system upgrade | **23426**

show redundant-power-system version | **23427**

show redundant-trunk-group | **23430**

show remote-device-management service-devices | **23432**

show remote-device-management services | **23441**

show remote-device-management statistics | **23445**

show remote-device-management subscribers | **23451**

show remote-device-management summary | **23455**

show resource-manager | **23458**

show resource-manager cpu | **23460**

show resource-manager memory | **23461**

show rift database | **23463**

show rift flood-reduction | **23467**

show rift interface | **23469**

show rift node | **23476**

show rift path-computation | **23480**

show rift routes | **23482**

show rift tie | **23486**

show rift topology | **23488**

show rift versions | **23491**

show rip general-statistics | **23493**

show rip neighbor | **23496**

show rip statistics | **23499**

show ripng general-statistics | **23504**

show ripng neighbor | **23506**

show ripng statistics | **23508**

show route | **23512**

show route active-path | **23526**

show route advertising-protocol | **23533**

show route all | **23542**

show route aspath-regex | **23545**

show route best | **23548**

show route brief | **23553**

show route ccc | **23555**

show route community | **23557**

show route community-name | **23560**

show route cumulative | **23562**

show route damping | **23564**

show route detail | **23573**

show route exact | **23608**

show route export | **23612**

show route export vrf-target | **23615**

show route extensive | **23619**

show route extensive <route> expanded-nh | **23643**

show route flow validation | **23649**

show route forwarding-table | **23652**

show route forwarding-table family fibre-channel | **23673**

show route forwarding-table interface-name | **23676**

show route hidden | **23680**

show route inactive-path | **23685**

show route inactive-prefix | **23690**

show route instance | 23692

show route label | 23698

show route label-switched-path | 23705

show route localization | 23707

show route martians | 23709

show route match-prefix | 23713

show route next-hop | 23715

show route no-community | 23719

show route output | 23723

show route pfe-install-error | 23727

show route protocol | 23730

show route range | 23740

show route receive-protocol | 23746

show route resolution | 23753

show route rpm-tracking | 23759

show route sla-tracking | 23764

show route snooping | 23767

show route source-gateway | 23771

show route summary | 23779

show route table | 23789

show route tenant | 23821

show route terse | 23824

show route vpn-localization | 23829

show routing nhdb pfe-install-error | 23835

show rsvp interface | 23837

show rsvp neighbor | 23847

show rsvp pop-and-forward | **23854**

show rsvp route-session-id | **23857**

show rsvp session | **23860**

show rsvp statistics | **23878**

show rsvp version | **23887**

show sap listen | **23892**

show security advance-policy-based-routing detail | **23894**

show security advance policy based routing multipath policy statistics | **23898**

show security advance-policy-based-routing multipath rule | **23901**

show security advance-policy-based-routing profile | **23913**

show security advance-policy-based-routing sla active-probe-statistics | **23915**

show security advance-policy-based-routing sla profile (Application Name) | **23919**

show security advance-policy-based-routing sla profile (Next-Hop) | **23921**

show security advance-policy-based-routing sla profile (Status) | **23928**

show security advance-policy-based-routing sla statistics | **23931**

show security advance-policy-based-routing sla status | **23934**

show security advance-policy-based-routing sla version | **23935**

show security advance-policy-based-routing statistics | **23936**

show security advance-policy-based-routing status | **23943**

show security advanced-policy-based-routing policy-name | **23945**

show security alarms | **23952**

show security alg h323 counters | **23958**

show security alg ike-esp-nat summary | **23963**

show security alg msrpc | **23964**

show security alg sccp calls | **23966**

show security alg sccp counters | **23969**

show security alg sip calls | **23973**

show security alg sip counters | **23975**

show security alg sip rate | **23996**

show security alg status | **23999**

show security alg status logical-system | **24007**

show security application-firewall rule-set | **24012**

show security application-firewall rule-set logical-system | **24017**

show security application-tracking counters | **24021**

show security casb policies | **24023**

show security datapath-debug capture | **24025**

show security datapath-debug counter | **24027**

show security dns-cache | **24028**

show security dynamic-address | **24031**

show security dynamic-policies | **24045**

show security dynamic-vpn users | **24053**

show security dynamic-vpn users terse | **24055**

show security firewall-authentication history | **24058**

show security firewall-authentication history address | **24061**

show security firewall-authentication history identifier | **24065**

show security firewall-authentication jims | **24069**

show security firewall-authentication users | **24071**

show security firewall-authentication users address | **24075**

show security firewall-authentication users auth-type | **24080**

show security firewall-authentication users identifier | **24082**

show security flow cp-session | **24086**

show security flow cp-session destination-port | **24091**

show security flow cp-session destination-prefix | **24095**

show security flow cp-session family | **24098**

show security flow cp-session protocol | **24102**

show security flow cp-session source-port | **24105**

show security flow cp-session source-prefix | **24109**

show security flow gate | **24112**

show security flow gate brief node | **24118**

show security flow gate destination-port | **24127**

show security flow gate destination-prefix | **24131**

show security flow gate protocol | **24135**

show security flow gate summary node | **24140**

show security flow ip-action | **24147**

show security flow pmi statistics | **24159**

show security flow session | **24162**

show security flow session advanced-anti-malware | **24177**

show security flow session application | **24184**

show security flow session application-firewall | **24190**

show security flow session brief node | **24199**

show security flow session destination-port | **24205**

show security flow session destination-prefix | **24211**

show security flow session extensive node | **24217**

show security flow session family | **24227**

show security flow session idp family | **24235**

show security flow session idp summary | **24237**

show security flow session interface | **24240**

show security flow session nat | **24246**

show security flow session plugins | **24251**

show security flow session policy-id | **24254**

show security flow session pretty | **24259**

show security flow session protocol | **24263**

show security flow session resource-manager | **24270**

show security flow session services-offload | **24276**

show security flow session session-identifier | **24283**

show security flow session session-state | **24291**

show security flow session source-port | **24295**

show security flow session source-prefix | **24300**

show security flow session ssl | **24306**

show security flow session summary family | **24310**

show security flow session summary node | **24313**

show security flow session summary services-offload | **24322**

show security flow session tenant | **24327**

show security flow session tunnel | **24330**

show security flow statistics | **24341**

show security flow status | **24348**

show security flow tunnel inspection statistics | **24358**

show security forward-options resource-manager | **24360**

show security forwarding-options mirror-filter | **24363**

show security group-vpn member ike security-associations | **24366**

show security group-vpn member ipsec inactive-tunnels | **24372**

show security group-vpn member ipsec security-associations | **24375**

show security group-vpn member ipsec statistics | **24381**

show security group-vpn member kek security-associations | **24385**

show security group-vpn member policy | **24391**

show security group-vpn server ike security-associations | **24394**

show security group-vpn server ipsec security-associations | **24401**

show security group-vpn server kek security-associations | **24404**

show security group-vpn server registered-members | **24409**

show security group-vpn server server-cluster | **24412**

show security group-vpn server statistics | **24416**

show security gtp | **24419**

show security gtp counters | **24420**

show security gtp counters path-rate-limit | **24432**

show security gtp grouped-ie-profile | **24434**

show security gtp gsn statistics | **24436**

show security gtp handover-group | **24437**

show security gtp ie-set | **24439**

show security gtp ip-group | **24441**

show security gtp message-ie-profile-v1 | **24443**

show security gtp message-ie-profile-v2 | **24445**

show security gtp message-list | **24447**

show security gtp profile | **24448**

show security gtp rate-limit default | **24460**

show security gtp sctp association | **24461**

show security gtp sctp counters | **24465**

show security host-vpn security-associations | **24473**

show security host-vpn version | **24477**

show security idp active-policy | **24478**

show security idp attack attack-list | **24480**

show security idp attack attack-list policy | **24485**

show security idp attack context | **24490**

show security idp attack deprecated-list | **24491**

show security idp attack description | **24492**

show security idp attack detail | **24494**

show security idp attack group-list | **24498**

show security idp attack table | **24500**

show security idp attacks deprecated-attacks policy policy_name | **24502**

show security idp counters action | **24504**

show security idp counters application-identification | **24507**

show security idp counters dfa | **24514**

show security idp counters flow | **24517**

show security idp counters http-decoder | **24534**

show security idp counters ips | **24537**

show security idp counters log | **24545**

show security idp counters memory | **24552**

show security idp counters packet | **24556**

show security idp counters packet-log | **24567**

show security idp counters pdf-decoder | **24581**

show security idp counters policy-manager | **24585**

show security idp counters ssl-inspection | **24588**

show security idp counters tcp-reassembler | **24592**

show security idp jist-package-version | **24599**

show security idp logical system | **24600**

show security idp logical-system policy-association | **24602**

show security idp memory | **24603**

show security idp policies | **24605**

show security idp policy-commit-status | **24607**

show security idp policy-commit-status clear | **24610**

show security idp policy-templates-list | **24611**

show security idp predefined-attacks | **24613**

show security idp security-package-version | **24615**

show security idp ssl-inspection key | **24618**

show security idp ssl-inspection session-id-cache | **24621**

show security idp status | **24622**

show security idp status detail | **24625**

show security ike active-peer | **24630**

show security ike debug-status | **24638**

show security ike peers | **24640**

show security ike peers statistics | **24644**

show security ike pre-shared-key | **24647**

show security ike security-associations | **24648**

show security ike stats | **24671**

show security ike tunnel-map | **24681**

show security integrity | **24685**

show security internal-security-association | **24686**

show security ipsec control-plane-security-associations | **24688**

show security ipsec inactive-tunnels | **24691**

show security ipsec next-hop-tunnels | **24699**

show security ipsec security-associations | **24701**

show security ipsec statistics | **24744**

show security ipsec traffic-selector | **24753**

show security ipsec tunnel-distribution | **24755**

show security ipsec tunnel-events-statistics | **24761**

show security ipsec vpn vpnname ike idle time | **24763**

show security keychain | **24765**

show security log | **24769**

show security log file | **24773**

show security log profile | **24781**

show security log query | **24791**

show security log report | **24794**

show security log report in-detail | **24796**

show security log report in-interval | **24802**

show security log report summary | **24804**

show security log severity | **24807**

show security macsec connections (SRX) | **24809**

show security macsec connections | **24812**

show security macsec connections (MX Series) | **24815**

show security macsec statistics | **24819**

show security macsec statistics (SRX Series Firewalls) | **24825**

show security match-policies | **24831**

show security mka sessions | **24840**

show security mka sessions (MX Series) | **24844**

show security mka sessions (SRX Series Device | **24849**

show security mka sessions summary | **24852**

show security mka statistics | **24854**

show security mka statistics | **24857**

show security mka statistics (MX Series) | **24860**

show security monitoring | **24864**

show security monitoring fpc fpc-number | **24867**

show security monitoring performance session | **24871**

show security monitoring performance spu | **24873**

show security nat destination pool | **24877**

show security nat destination rule | **24881**

show security nat destination rule-application | **24886**

show security nat destination summary | **24889**

show security nat incoming-table | **24893**

show security nat interface-nat-ports | **24895**

show security nat resource-usage source-pool | **24900**

show security nat source deterministic | **24904**

show security nat source paired-address | **24907**

show security nat source persistent-nat-table | **24910**

show security nat source pool | **24914**

show security nat source port-block | **24927**

show security nat source rule | **24931**

show security nat source rule-application | **24937**

show security nat source summary | **24940**

show security nat static rule | **24948**

show security np-cache summary | **24955**

show security pki ca-certificate | **24957**

show security pki ca-certificate (View) | **24962**

show security pki certificate-request | **24968**

show security pki certificate-request (View) | **24972**

show security pki crl | **24975**

show security pki crl (View) | **24979**

show security pki local-certificate | **24982**

show security pki local-certificate | **24987**

show security pki local-certificate (View) | **24992**

show security pki node-local certificate-request | **24999**

show security pki node-local local-certificate | **25003**

show security pki statistics | **25009**

show security policies | **25028**

show security policies checksum | **25056**

show security policies file-serialization counters | **25059**

show security policies hit-count | **25061**

show security policies information | **25068**

show security policies lookup-intact-on-commit | **25077**

show security policies policy set | **25080**

show security policies unknown-source-identity | **25083**

show security policies zone-context | **25085**

show security policy-report | **25089**

show security re-distribution ipsec-vpn | **25093**

show security resource-manager group active | **25095**

show security resource-manager resource active | **25099**

show security resource-manager settings | **25103**

show security resource-manager summary | **25106**

show security screen ids-option | **25108**

show security screen statistics | **25116**

show security screen status | **25132**

show security screen white-list | **25133**

show security shadow-policies | **25135**

show security softwires | **25138**

show security softwires map-e confidentiality status | **25140**

show security softwires map-e domain | **25142**

show security ssh key-pair-identity | **25144**

show security tcp-encap connection | **25146**

show security tcp-encap statistics | **25149**

show security tpm status | **25150**

show security tunnel inspection | **25153**

show security user-identification device-provision authentication-source active-directory start 1 count 9 (match-string|prefix) | **25156**

show security user-identification device-provision authentication-source jims start 1 count 9 (match-string|prefix) | **25158**

show security user-identification local-authentication-table | **25159**

show security user-identification role-provision all | **25163**

show security user-identification role-provision authentication-source active-directory start 1 count 9 (match-string|prefix) | **25164**

show security user-identification role-provision authentication-source jims start 1 count 9 (match-string|prefix) | **25166**

show security user-identification source-identity-provision all | **25169**

show security user-identification user-provision all | **25171**

show security user-identification user-provision authentication-source active-directory start 1 count 9 (match-string|prefix) | **25172**

show security user-identification user-provision authentication-source jims start 1 count 9 (match-string|prefix) | **25174**

show security utm anti-spam statistics | **25176**

show security utm anti-spam status | **25182**

show security utm anti-virus statistics | **25183**

show security utm anti-virus status | **25191**

show security utm content-filtering statistics | **25195**

show security utm custom-objects url-feed status | **25207**

show security utm session | **25212**

show security utm status | **25213**

show security utm web-filtering category base-filter | **25215**

show security utm web-filtering category category | **25217**

show security utm web-filtering category status | **25219**

show security utm web-filtering statistics | **25220**

show security utm web-filtering status | **25228**

show security zones | **25232**

show security zones type | **25242**

show service l2tp session routing-instance | **25246**

show service user-identification authentication-source aruba-clearpass user-query counters | **25248**

show service user-identification authentication-source aruba-clearpass user-query status | **25250**

show service user-identification identity-management | **25252**

show services accounting | **25256**

show services accounting aggregation | **25260**

show services accounting aggregation template | **25266**

show services accounting aggregation template (View) | **25268**

show services accounting aggregation (View) | **25269**

show services accounting errors | **25270**

show services accounting flow | **25277**

show services accounting flow-detail | **25286**

show services accounting flow-detail (View) | **25293**

show services accounting memory | **25295**

show services accounting packet-size-distribution | **25297**

show services accounting status | **25300**

show services accounting usage | **25305**

show services advanced-anti-malware dynamic-filter status | **25308**

show services advanced-anti-malware machine-learning-scan | **25311**

show services advanced-anti-malware policy | **25313**

show services advanced-anti-malware profile | **25317**

show services advanced-anti-malware statistics | **25321**

show services advanced-anti-malware status | **25328**

show services agf fiveqi remarking | **25331**

show services agf nas | **25334**

show services agf ngap | **25338**

show services alg conversations | **25348**

show services alg statistics | **25356**

show services analytics agent | **25374**

show services anti-virus statistics | **25378**

show services application-aware-access-list flows | **25381**

show services application-aware-access-list statistics | **25386**

show services application-identification application | **25389**

show services application-identification application micro-applications | **25402**

show services application-identification application non-configurable | **25404**

show services application identification application obsolete applications | **25406**

show services application-identification application-system-cache | **25408**

show services application-identification application-system-cache (View) | **25413**

show services application-identification commit-status | **25419**

show services application-identification commit-status (Next Gen Services) | **25420**

show services application-identification counter | **25422**

show services application-identification counter (AppSecure) | **25426**

show services application-identification entries | **25435**

show services application-identification group | **25438**

show services application-identification packet-capture counters | **25443**

show services application-identification statistics application-groups | **25446**

show services application-identification statistics applications | **25451**

show services application-identification status | **25456**

show services application-identification version | **25469**

show services application-identification group obsolete-groups | **25472**

show services application-identification recent-appid-sigpack-versions | **25473**

show services application-identification recent-appid-sigpack-versions | **25474**

show services captive-portal-content-delivery | **25476**

show services cos statistics | **25483**

show services cos statistics (Next Gen Services) | **25487**

show services crtp | **25492**

show services crtp flows | **25496**

show services dns-filtering cache | **25498**

show services dynamic-flow-capture content-destination | **25502**

show services dynamic-flow-capture control-source | **25504**

show services dynamic-flow-capture statistics | **25508**

show services fixed-wireless-access statistics | **25512**

show services flow-collector file interface | **25515**

show services flow-collector input interface | **25519**

show services flow-collector interface | **25521**

show services ha detail | **25532**

show services ha statistics | **25535**

show services hcm pic-statistics | **25542**

show services hcm statistics | **25550**

show services icap-redirect server status | **25553**

show services icap-redirect statistic | **25554**

show services icap-redirect status | **25559**

show services ids | **25562**

show services inband-flow-telemetry | **25574**

show services inline ip-reassembly statistics | **25577**

show services inline-monitoring feature-profile-mapping fpc-slot | **25586**

show services inline-monitoring statistics fpc-slot | **25588**

show services inline nat pool | **25591**

show services inline nat statistics | **25594**

show services inline softwire statistics | **25598**

show services ip-monitoring status | **25604**

show services ipsec-vpn certificates | **25610**

show services ipsec-vpn ike security-associations | **25615**

show services ipsec-vpn ipsec security-associations | **25622**

show services ipsec-vpn ipsec statistics | **25632**

show services l2tp client | **25638**

show services l2tp destination | **25641**

show services l2tp destination lockout | **25647**

show services l2tp multilink | **25649**

show services l2tp radius | **25656**

show services l2tp session | **25661**

show services l2tp session-limit-group | **25674**

show services l2tp summary | **25676**

show services l2tp tunnel | **25685**

show services l2tp tunnel-group | **25694**

show services l2tp tunnel-switch destination | **25696**

show services l2tp tunnel-switch session | **25702**

show services l2tp tunnel-switch summary | **25709**

show services l2tp tunnel-switch tunnel | **25711**

show services l2tp user | **25718**

show services link-services cpu-usage | **25724**

show services lrf collector statistics | **25730**

show services lrf rule statistics | **25732**

show services lrf statistics | **25735**

show services lrf template | **25737**

show services mobile-edge anchor class-of-service | **25740**

show services mobile-edge peers | **25742**

show services mobile-edge resources call-admission-control statistics | **25747**

show services mobile-edge sessions | **25749**

show services mobile-edge summary | **25772**

show services monitoring rfc2544 | **25775**

show services monitoring rfc2544 | **25780**

show services monitoring rpm history-results | **25787**

show services monitoring rpm probe-results | **25794**

show services monitoring twamp client control-info | **25804**

show services monitoring twamp client history-results | **25807**

show services monitoring twamp client probe-results | **25814**

show services monitoring twamp client test-info | **25825**

show services monitoring twamp server control-info | **25828**

show services monitoring twamp server test-info | **25831**

show services nat destination pool | **25834**

show services nat destination rule | **25836**

show services nat destination summary | **25839**

show services nat deterministic-nat internal-host | **25842**

show services nat deterministic-nat nat-port-block | **25844**

show services nat ipv6-multicast-interfaces | **25847**

show services nat pool | **25850**

show services nat resource-usage source-pool | **25858**

show services nat source deterministic | **25861**

show services nat source mappings address-pooling-paired | **25864**

show services nat source mappings endpoint-independent | **25868**

show services nat source mappings pcp | **25871**

show services nat source mappings summary | **25873**

show services nat source pool | **25875**

show services nat source port-block | **25881**

show services nat source rule | **25884**

show services nat source rule-application | **25888**

show services nat source summary | **25891**

show services paa status | **25893**

show services pcef pic | **25899**

show services pcef subscribers | **25901**

show services pcp statistics | **25910**

show services policies | **25915**

show services policies detail | **25917**

show services policies hit-count | **25921**

show services policies interface | **25922**

show services policies service-set | **25924**

show services redundancy-group | **25925**

show services rpm active-servers | **25937**

show services rpm history-results | **25939**

show services rpm probe-results | **25944**

show services rpm rfc2544-benchmarking | **25960**

show services rpm rfc2544-benchmarking test-id | **25969**

show services rpm twamp client connection | **25995**

show services rpm twamp client history-results | **25997**

show services rpm twamp client probe-results | **26002**

show services rpm twamp client session | **26013**

show services rpm twamp server connection | **26016**

show services rpm twamp server session | **26018**

show services screen ids-option | **26021**

show services screen-statistics service-set | **26024**

show services security-intelligence | **26032**

show services security-intelligence category summary | **26034**

show services security-intelligence dns-statistics | **26037**

show services security-intelligence update status | **26041**

show services security-metadata-streaming | **26042**

show services service-redirect statistic | **26048**

show services service-set statistics ids drops | **26049**

show services service-sets cpu-usage | **26058**

show services service-sets memory-usage | **26061**

show services service-sets plug-ins | **26063**

show services service-sets statistic screen-drops (Next Gen Services) | **26065**

show services service-sets statistic screen-session-limit-counters (Next Gen Services) | **26073**

show services service-sets statistics ids session-limits counters | **26084**

show services service-sets statistics integrity-drops | **26095**

show services service-sets statistics jflow-log | **26101**

show services service-sets statistics packet-drops | **26111**

show services service-sets statistics syslog | **26114**

show services service-sets statistics tcp | **26123**

show services service-sets statistics tcp-mss | **26125**

show services service-sets summary | **26127**

show services sessions | **26129**

show services sessions (Aggregated Multiservices) | **26141**

show services sessions analysis | **26151**

show services sessions analysis (USF) | **26156**

show services sessions count | **26162**

show services sessions (Next Gen Services) | **26163**

show services sessions service-set | **26176**

show services sessions softwire | **26178**

show services sessions tcp-log | **26183**

show services sessions utilization | **26185**

show services soft-gre tunnel | **26186**

show services softwire | **26190**

show services softwire flows | **26192**

show services softwire statistics | **26197**

show services ssl certificate | **26209**

show services ssl droplogs | **26215**

show services ssl initiation counters | 26217

show services ssl initiation profile | 26222

show services ssl proxy certificate-cache entries | 26228

show services ssl proxy certificate–cache statistics | 26231

show services ssl proxy counters | 26233

show services ssl proxy profile | 26239

show services ssl proxy session-cache entries | 26243

show services ssl proxy session-cache statistics | 26249

show services ssl proxy statistics | 26252

show services ssl proxy status | 26255

show services ssl session | 26258

show services ssl termination counters | 26262

show services ssl termination profile | 26267

show services stateful-firewall conversations | 26273

show services stateful-firewall flow-analysis | 26279

show services stateful-firewall flows | 26286

show services stateful-firewall sip-call | 26294

show services stateful-firewall sip-register | 26300

show services stateful-firewall statistics | 26305

show services stateful-firewall statistics application-protocol sip | 26318

show services stateful-firewall subscriber-analysis | 26322

show services subscriber analysis | 26327

show services tcp-log | 26330

show services tcp-log connections | 26331

show services traffic-detection-function hcm statistics | 26333

show services traffic-detection-function sessions | 26338

show services traffic-load-balance statistics | **26341**

show services unified-access-control authentication-table | **26358**

show services unified-access-control counters | **26362**

show services unified-access-control policies | **26365**

show services unified-access-control roles | **26367**

show services unified-access-control status | **26369**

show services url-filter dns-resolution profile | **26370**

show services url-filter dns-resolution-statistics profile template | **26374**

show services url-filter statistics profile template | **26381**

show services user-identification active-directory-access domain-controller status | **26385**

show services user-identification active-directory-access statistics | **26389**

show services user-identification active-directory-access user-group-mapping | **26393**

show services user-identification authentication-table | **26397**

show services user-identification device-information table | **26421**

show services user-identification logical-domain-identity-management | **26426**

show services user-identification validate-statistics | **26430**

show services video-monitoring mdi errors fpc-slot | **26431**

show services video-monitoring mdi flows fpc-slot | **26434**

show services video-monitoring mdi stats fpc-slot | **26441**

show services web-filter dns-resolution profile | **26445**

show services web-filter dns-resolution-statistics profile template | **26449**

show services web-filter secintel-policy status | **26455**

show services web-filter statistics dns-filter-template | **26461**

show services web-filter statistics profile | **26464**

show services web-proxy dns forwarding-cache | **26470**

show services web-proxy dns global-cache statistics | **26473**

show services web-proxy session | **26476**

show sflow | **26478**

show sflow collector | **26482**

show sflow interface | **26485**

show snmp health-monitor | **26489**

show snmp inform-statistics | **26499**

show snmp mib | **26502**

show snmp rmon | **26516**

show snmp rmon history | **26523**

show snmp statistics | **26529**

show snmp stats-response-statistics | **26539**

show snmp v3 | **26544**

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

show spanning-tree mstp configuration | **26568**

show spanning-tree statistics | **26572**

show spanning-tree statistics bridge | **26575**

show spanning-tree statistics interface | **26576**

show spanning-tree statistics message-queues | **26578**

show spanning-tree stp-buffer see-all | **26580**

show spring-traffic-engineering | **26587**

show-spring-traffic-statistics | **26600**

show ssl-certificates | **26602**

show static-subscribers sessions | **26605**

show subscribers | **26609**

show subscribers summary | **26663**

show synchronous-ethernet esmc statistics | **26674**

show synchronous-ethernet esmc transmit | **26676**

show synchronous-ethernet global-information | **26679**

show system alarms | **26682**

show system alarms | **26683**

show system anomalies (GNF) | **26689**

show system anomalies gnf-id | **26692**

show system applications (Junos OS Evolved) | **26695**

show system audit | **26709**

show system auto-snapshot | **26720**

show system autoinstallation status | **26723**

show system autorecovery state | **26725**

show system boot console | **26727**

show system boot-messages | **26729**

show system buffers | **26741**

show system certificate | **26751**

show system commit | **26754**

show system commit server queue | **26758**

show system commit server status | **26763**

show system configuration archival | **26765**

show system configuration database status | **26766**

show system configuration database usage | **26768**

show system configuration rescue | **26770**

show system configuration revision | **26772**

show system connections | **26775**

show system core-dumps | **26783**

show system core dumps (Junos OS Evolved) | **26802**

show system cpu (JDM) | **26804**

show system data-models | **26807**

show system directory-usage | **26808**

show system download | **26815**

show system errors | **26817**

show system errors active | **26822**

show system errors count | **26829**

show system errors error-id | **26831**

show system errors fru | **26834**

show system errors history | **26841**

show system errors inactive | **26844**

show system filesystem encryption status | **26852**

show system firmware | **26855**

show system firmware (Junos OS Evolved) | **26865**

show system health-monitor | **26868**

show system information | **26873**

show system inventory hardware cpu | **26874**

show-system-inventory-hardware-memory | **26881**

show-system-inventory-hardware-optics | **26884**

show-system-inventory-hardware-storage | **26886**

show system khms-stats | **26889**

show system license (View) | **26896**

show system login lockout | **26902**

show system mac-addresses (JDM) | **26903**

show system max-smbios-version (JDM) | **26907**

show system memory | **26908**

show system memory (JDM) | **26918**

show system monitor memory actions | **26921**

show system monitor memory events | **26923**

show system monitor memory status | **26926**

show system name-resolution | **26929**

show system network (JDM) | **26931**

show system node-attributes | **26936**

show system nodes | **26939**

show system ntp threshold | **26941**

show system packet-forwarding-options hw-resource-monitor monitor-profile | **26943**

show system packet-forwarding-options hw-resource-monitor resource-list | **26946**

show system packet-forwarding-options hw-resource-monitor utilization-info | **26948**

show system packet-forwarding-options hw-resource-utilization-info | **26952**

show system processes | **26957**

show system queues | **26975**

show system random-mac-prefix (JDM) | **26981**

show system reboot | **26982**

show system resource-monitor | **26987**

show system resource-monitor ifd-cos-queue-mapping fpc | **26992**

show system resource-monitor subscribers-limit | **26995**

show system rollback | **27001**

show system schema | **27004**

show system security-profile | **27008**

show system security-profile scheduler | **27016**

show system security-profile secure-wire | **27020**

show system security-profile security-log-stream-number | **27025**

show system security-profile security-log-stream-number detail | **27028**

show system security-profile security-log-stream-number summary | **27032**

show system services dhcp binding | **27034**

show system services dhcp client | **27038**

show system services dhcp conflict | **27043**

show system services dhcp global | **27045**

show system services dhcp pool | **27047**

show system services dhcp relay-statistics | **27051**

show system services dhcp statistics | **27053**

show system services dns dns-proxy | **27057**

show system services dynamic-dns | **27061**

show system services outbound-ssh jefflee | **27063**

show system services service-deployment | **27067**

show system snapshot (Junos OS) | **27069**

show system snapshot (Junos OS Evolved) | **27073**

show system snapshot (Junos OS with Upgraded FreeBSD) | **27075**

show system snapshot media | **27078**

show system software | **27081**

show system software add-restart (Junos OS Evolved) | **27086**

show system software list | **27088**

show system software recovery-package | **27092**

show system software restore-point-status | **27094**

show system software upgrade status | **27096**

show system software usb-software-version | **27098**

show system statistics | **27101**

show system statistics arp | **27119**

show system statistics icmp | **27129**

show system statistics icmp6 | **27136**

show system statistics igmp | **27146**

show system statistics ip | **27151**

show system statistics ip6 | **27162**

show system statistics netisr per-priority-stats | **27173**

show system statistics netisr qdrops | **27174**

show system statistics netisr socket-drops | **27176**

show system statistics tcp | **27178**

show system storage | **27190**

show system storage partitions | **27198**

show system storage satellite | **27202**

show system subscriber-management redundancy-state dhcp active-leasequery interface | **27205**

show system subscriber-management redundancy-state interface | **27208**

show system subscriber-management resiliency | **27211**

show system subscriber-management route | **27222**

show system subscriber-management statistics | **27229**

show system subscriber-management summary | **27242**

show system switchover | **27246**

show system unified-services status | **27251**

show system uptime | **27253**

show system users | **27260**

show system virtual-memory | **27266**

show system visibility cpu | **27279**

show system visibility host | **27286**

show system visibility jcp | **27296**

show system visibility jdm | **27300**

show system visibility memory | **27304**

show system visibility network | **27308**

show system visibility vnf | **27316**

show system yang package | **27323**

show system ztp | **27325**

show task | **27331**

show task io | **27336**

show task logical-system-mux | **27339**

show task memory | **27342**

show task replication | **27348**

show tcp-forwarding status | **27351**

show ted database | **27357**

show-ted-ipv4-prefix | **27372**

show-ted-ipv6-prefix | **27374**

show ted link | **27377**

show ted protocol | **27384**

show ted spring-te-policy | **27386**

show trace | **27392**

show unified-edge tdf aaa radius client statistics | **27396**

show unified-edge tdf aaa radius client status | **27404**

show unified-edge tdf aaa radius network-element statistics | **27406**

show unified-edge tdf aaa radius server statistics | **27410**

show unified-edge tdf aaa radius server status | **27416**

show unified-edge tdf aaa radius snoop-segment statistics | **27420**

show unified-edge tdf aaa statistics | **27425**

show unified-edge tdf address-assignment pool | **27437**

show unified-edge tdf address-assignment service-mode | **27442**

show unified-edge tdf address-assignment statistics | **27445**

show unified-edge tdf call-admission-control statistics | **27448**

show unified-edge tdf call-rate statistics | **27452**

show unified-edge tdf diameter network-element statistics | **27456**

show unified-edge tdf diameter network-element status | **27459**

show unified-edge tdf diameter pcc-gx statistics | **27462**

show unified-edge tdf diameter peer statistics | **27469**

show unified-edge tdf diameter peer status | **27476**

show unified-edge tdf domain service-mode | **27481**

show unified-edge tdf domain statistics | **27484**

show unified-edge tdf resource-manager clients | **27490**

show unified-edge tdf service-mode | **27493**

show unified-edge tdf statistics | **27496**

show unified-edge tdf status | **27507**

show unified-edge tdf subscribers | **27513**

show unified-edge tdf system interfaces | **27532**

show unified-edge tdf system interfaces service-mode | **27534**

show uplink-failure-detection | **27538**

show uplink-failure-detection | **27541**

show v4ov6-tunnels | **27544**

show validation database | **27547**

show validation group | **27551**

show validation replication database | **27553**

show validation session | **27556**

show validation statistics | **27563**

show version (JDM) | **27567**

show version (Junos OS) | **27569**

show version (Junos OS Evolved) | **27572**

show version vnf | **27575**

show virtual-chassis | **27578**

show virtual-chassis active-topology | **27587**

show virtual-chassis active-topology (MX Series Virtual Chassis) | **27593**

show virtual-chassis device-topology | **27596**

show virtual-chassis device-topology (MX Series Virtual Chassis) | **27605**

show virtual-chassis fast-failover | **27609**

show virtual-chassis heartbeat (MX Series Virtual Chassis) | **27611**

show virtual-chassis login | **27616**

show virtual-chassis mode | **27617**

show virtual-chassis protocol adjacency | **27622**

show virtual-chassis protocol adjacency (MX Series Virtual Chassis) | **27627**

show virtual-chassis protocol database | **27633**

show virtual-chassis protocol database (MX Series Virtual Chassis) | **27638**

show virtual-chassis protocol interface | **27647**

show virtual-chassis protocol interface (MX Series Virtual Chassis) | **27650**

show virtual-chassis protocol route | **27655**

show virtual-chassis protocol route (MX Series Virtual Chassis) | **27659**

show virtual-chassis protocol statistics | **27663**

show virtual-chassis protocol statistics (MX Series Virtual Chassis) | **27668**

show virtual-chassis status (MX Series Virtual Chassis) | **27672**

show virtual-chassis vc-path | **27676**

show virtual-chassis vc-port | **27681**

show virtual-chassis vc-port diagnostics optics | **27688**

show virtual-chassis vc-port (MX Series Virtual Chassis) | **27709**

show virtual-chassis vc-port statistics | **27713**

show virtual-network-functions | **27723**

show virtual-network-functions (Junos Node Slicing) | **27728**

show vlans | **27738**

show vlans (JDM) | **27767**

show vlans evpn nd-table | **27769**

show vlans vlan-groups | **27774**

show vmhost bridge | **27775**

show vmhost connections | **27778**

show vmhost control-plane | **27780**

show vmhost crash | **27782**

show vmhost crash (NFX Series) | **27784**

show vmhost forwarding-options analyzer | **27785**

show vmhost hard-disk-test | **27787**

show vmhost hardware | **27789**

show vmhost information | **27791**

show vmhost jdm status (In-Chassis Model) | **27793**

show vmhost logs | **27795**

show vmhost management-if | **27797**

show vmhost memory | **27798**

show vmhost mode | **27799**

show vmhost netstat | **27811**

show vmhost network external-network | **27813**

show vmhost processes | **27815**

show vmhost resource-usage | **27818**

show vmhost snapshot | **27821**

show vmhost snmp mib walk | **27823**

show vmhost status | **27825**

show vmhost status (NFX Series) | **27827**

show vmhost status (In-Chassis Model) | **27829**

show vmhost storage | **27831**

show vmhost uptime | **27838**

show vmhost uptime (NFX Series) | **27840**

show vmhost version | **27841**

show vmhost version (NFX Series) | **27844**

show vmhost vlans | **27847**

show vpls connections | **27849**

show vpls flood event-queue | **27865**

show vpls flood instance | **27867**

show vpls flood route | **27871**

show vpls mac-move-action | **27874**

show vpls mac-table | **27876**

show vpls statistics | **27884**

show vrrp | **27888**

show vrrp | **27905**

show vrrp track | **27920**

show wlan access-points | **27926**

shutdown (BPDU Block) | **27935**

shutdown (Protocols BGP) | **27936**

shutdown-threshold | **27937**

sib | **27939**

signal-bandwidth | **27941**

signal-degrade | **27942**

signal-degrade-monitor-enable | **27944**

signal-type | **27946**

signal-type (BITS Interfaces) | **27947**

signaling | **27949**

signature | **27950**

signature (Application Identification) | **27952**

Signature Hash Algorithm (Security IKE) | **27953**

signature (Security IDP) | **27955**

simple-filter | **27962**

simple-filter (Applying to an Interface) | **27964**

simple-filter (Interfaces) | **27965**

simple-mail-client-service | **27967**

single-home (Junos Fusion) | **27968**

single-rate | **27969**

sip | **27971**

sip (Application Profile) | **27973**

sip-call-hold-timeout | **27974**

sip-password | **27976**

sip (Security) | **27977**

sip-server | **27979**

sip (Services CoS Next Gen Services) | **27981**

sip-user-id | **27983**

site-identifier (Layer 2 Circuits) | **27984**

site-identifier (VPLS) | **27985**

site (Layer 2 Circuits) | **27987**

site-preference | **27989**

site-range | **27990**

site-reputation-action | **27992**

site (VPLS Multihoming for FEC 128) | **27994**

site (VPLS Multihoming for FEC 129) | **27995**

size | **27997**

size (VNF) | **27999**

size (Interface) | **28000**

size (Next Gen Services Global System Logging) | **28001**

size (Sampling and Traceoptions) | **28003**

size (Security Web Filtering Cache) | **28004**

size (Services) | **28006**

size (System) | **28007**

skew-timer-disable | **28009**

skip-arp-iteration (RFC 2544 Benchmarking) | **28010**

sla-iterator-profile | **28012**

sla-iterator-profiles | **28013**

sla-iterator-profiles (OAM CFM) | **28015**

sla-options | **28017**

sla-rule | **28019**

sla-tracking | **28027**

slaac-snooping | **28030**

slamon-services | **28032**

slave | **28034**

slice (CoS Interfaces) | **28039**

slice (export routing policy action) | **28041**

slice (firewall filter action) | **28042**

slice (firewall filter match condition) | **28043**

slow-pfe-alarm | **28044**

smart-check | **28045**

smart-optimize-timer | **28048**

smg-service (Enhanced Subscriber Management) | **28050**

smtp-profile (Security Content Security Policy Antispam) | **28051**

smtp-profile (Security Content Security Policy Antivirus) | **28052**

smtp-profile (Security Content Security Policy Content Filtering) | **28053**

snapshot | **28055**

snext | **28056**

snmp | **28058**

snmp-command | **28064**

snmp-community | **28065**

snmp-options | **28067**

snmp (SNMP Scripts) | **28068**

snmp-trap-thresholds | **28070**

snmp-trap-thresholds (Next Gen Services) | **28072**

snmp-value-match-msmic (Services NAT Options) | **28074**

snoop-pseudowires | **28075**

snoop-segment (TDF Domain Selection) | **28076**

snoop-segments (RADIUS) | **28078**

snoop-segments (TDF Gateway) | **28080**

snr-margin | **28081**

sockets | **28082**

soft-gre | **28083**

soft-limit | **28085**

soft-limit-clear | **28087**

soft-preemption (Protocols MPLS) | **28088**

soft-preemption (Protocols RSVP) | **28089**

software-upgrade | **28090**

softwire-concentrator | **28092**

softwire-name (Next Gen Services) | **28093**

softwire-options | **28095**

softwire-rules | **28097**

softwire-types (Next Gen Services) | **28098**

softwires | **28102**

softwires (Next Gen Services) | **28104**

softwires-rule-set (Service Set Next Gen Services) | **28105**

solicit-router-advertisement-unicast | **28106**

sonet | **28108**

sophos-engine | **28109**

source (GRE Encapsulation) | **28111**

source (Interfaces) | **28112**

source-active-advertisement | **28114**

source-address (Services IPSec VPN) | **28115**

source-address (Dynamic Tunnel) | **28117**

source-address (IGMP and MLD Snooping) | **28118**

source-address | **28120**

source-address (Content Security) | **28121**

source-address (AACL) | **28123**

source-address (Access LDAP) | **28124**

source-address (Access RADIUS) | **28125**

source-address-excluded | **28127**

source-address-filter | **28128**

source-address-filter (Interfaces) | **28130**

source-address-filter-list | **28131**

source-address (Forwarding Options) | **28133**

source-address-giaddr | **28135**

source-address (IDS MS-DPC) | **28136**

source-address (LRF Profile) | **28138**

source-address-name (NAT Next Gen Services) | **28139**

source-address-name (Security Destination NAT) | **28140**

source-address-name (Security Source NAT) | **28141**

source-address-name (Security Static NAT Rule Set) | **28142**

source-address (NAT Next Gen Services) | **28144**

source-address (Next Gen Services Service-Set Remote System Logging) | **28145**

source-address (PCP) | **28146**

source-address (RADIUS, System Logging, or TACACS+) | **28148**

source-address-range | **28149**

source-address-range (IDS MS-DPC) | **28150**

source-address-range (PCP) | **28152**

source-address-range (Services NAT) | **28153**

source-address-range (Services Stateful Firewall) | **28155**

source-address (Routing Options) | **28156**

source-address (RPM) | **28157**

source-address (Security Destination NAT) | **28160**

source-address (Security IDP) | **28161**

source-address (Security IDP Policy) | **28162**

source-address (Security IDP Sensor Configuration) | **28163**

source-address (Security Policies) | **28164**

source-address (Security Source NAT) | **28166**

source-address (Security Static NAT Rule Set) | **28167**

source-address (Service Sets) | **28168**

source-address (Services CoS) | **28169**

source-address (Services NAT) | **28171**

source-address (Services Stateful Firewall) | **28172**

source-address (SNMP) | **28174**

source-address (SRC Software) | **28175**

source-address (Subscriber Secure Policy) | **28176**

source-address (System Logging) | **28177**

source-address (TWAMP) | **28179**

source-addresses | **28181**

source (Application Identification) | **28182**

source-attachment-identifier (Protocols VPWS) | **28183**

source-bmac | **28185**

source (Bridge Domains) | **28187**

source (Chassis Synchronization) | **28188**

source-checking | **28190**

source-class-usage | **28191**

source-classes | **28193**

source (Commit Scripts) | **28194**

source-count (Dynamic MLD Interface) | **28196**

source-count (Protocols IGMP) | **28197**

source-count (Protocols MLD) | **28199**

source (Distributed IGMP) | **28200**

source (Dynamic IGMP Interface) | **28201**

source (Dynamic MLD Interface) | **28203**

source-end-user-profile | **28204**

source (Event Scripts) | **28206**

source-except | **28207**

source-filtering | **28208**

source-filtering (Interfaces) | **28210**

source-gateway (Tunnel Profile) | **28211**

source-hostname | **28213**

source-id | **28214**

source-identity | **28216**

source-identity-log (Security) | **28218**

source-increment (Dynamic MLD Interface) | **28220**

source-increment (Protocols IGMP) | **28221**

source-increment (Protocols MLD) | **28223**

source-interface (RADIUS Client) | **28224**

source-interface | **28225**

source-interface (RADIUS Server) | **28227**

source-interface (RADIUS Snoop Segment) | **28228**

source-interface-set-at-login | **28229**

source interfaces | **28231**

source-ip | **28233**

source-ip-address-list | **28235**

source-ip-address (RADIUS Snoop Segment) | **28237**

source-ip-based | **28238**

source-ip-change (Forwarding Options) | **28239**

source-ip (Flow Monitoring Logs for NAT) | **28240**

source-ip (Security Alarms) | **28242**

source-ipv4-address (RFC 2544 Benchmarking) | **28244**

source (JET Scripts) | **28245**

source-mac-address | **28247**

source-mac-address-list | **28249**

source-mode | **28250**

source (Multicast VLAN Registration) | **28252**

source-nat | **28254**

source (Op Scripts) | **28255**

source-packet-routing (Segment Routing Traffic Engineering) | **28257**

source-packet-routing (Junos Telemetry Interface) | **28261**

source-packet-routing (routing-options) | **28264**

source-packet-routing (Protocols IS-IS) | **28266**

source-packet-routing (Protocols OSPF) | **28270**

source-packet-routing (Routing Instances Protocols EVPN) | **28273**

source (PIM RPF Selection) | **28275**

source-pool | **28276**

source-port | **28277**

source-port (firewall) | **28278**

source-port (Applications) | **28280**

source-port (Security Forwarding Options) | **28281**

source-port (Security IDP) | **28286**

source-port (Security Source NAT Rule Set) | **28288**

source-port (Security Static NAT Rule Set) | **28289**

source-port (Subscriber Secure Policy) | **28290**

source-prefix (IDS) | **28292**

source-prefix-ipv6 (IDS) | **28293**

source-prefix-list (PCP) | **28295**

source-prefix-list (Services AACL) | **28296**

source-prefix-list (Services CoS) | **28297**

source-prefix-list (Services IDS) | **28299**

source-prefix-list (Services NAT) | **28300**

source-prefix-list (Services Stateful Firewall) | **28301**

source-prefix (Security Forwarding Options) | **28303**

source-prefix (Services NAT) | **28304**

source (Protocols IGMP) | **28305**

source (Protocols MLD) | **28307**

source (Protocols MSDP) | **28308**

source-route-option (IDS Screen Next Gen Services) | **28310**

source-routing | **28311**

source (Routing Instances) | **28312**

source (Routing Instances Provider Tunnel Selective) | **28314**

source-routing-path | **28316**

source-routing-path-template | **28319**

source (Security Source NAT) | **28322**

source (SNMP Scripts) | **28326**

source (Source-Specific Multicast) | **28327**

source-threshold | **28329**

source-tree (MBGP MVPN) | **28330**

source-udp-port (RFC 2544 Benchmarking) | **28332**

source-based-fpolicy-statementorwarding | **28333**

source-lookup | **28338**

source-routing | **28339**

spam-action | **28341**

speed (24-port and 12-port 10 Gigabit Ethernet PIC) | **28342**

speed | **28344**

speed (Chassis Cluster) | **28347**

speed (Ethernet) | **28349**

speed (Fibre Channel Interfaces) | **28370**

speed (Gigabit Ethernet interface) | **28372**

speed (Gigabit Ethernet interface - MX Series) | **28373**

speed (Interfaces) | **28376**

speed (MX Series DPC) | **28377**

spf-options (Protocols IS-IS) | **28379**

spf-options (Protocols OSPF) | **28381**

spi | **28384**

spi (Junos-FIPS Software) | **28386**

spi (Junos OS) | **28387**

spid1 | **28388**

spid2 | **28389**

splitting-merging | **28391**

spring-te | **28393**

spring-te (Dynamic Tunnels) | **28395**

spring-te-lsp-next-hop | **28397**

spring-traffic-engineering-edit-protocols | **28398**

spt-only | **28400**

spt-threshold | **28401**

sql | **28402**

squelch | **28404**

srgb (Protocols BGP) | **28407**

srgb (Protocols IS-IS) | **28408**

srgb-label-range | **28410**

srlg | **28411**

srlg-cost | **28413**

srlg-value | **28414**

srv6 (Routing Options) | **28415**

srv6 (Protocols BGP) | **28417**

srv6-sid | **28419**

ssb | **28421**

ssh | **28422**

ssh-known-hosts | **28426**

ssh (NETCONF) | **28428**

ssh (System Services) | **28430**

ssl | **28439**

ssl-encryption | **28442**

ssl-inspection | **28443**

ssl-proxy (Application Services) | **28446**

ssl (Services) | **28447**

ssl-termination-profile | **28450**

ssm-groups | **28452**

ssm-map (Dynamic IGMP Interface) | **28453**

ssm-map (Dynamic MLD Interface) | **28455**

ssm-map-policy (Dynamic IGMP Interface) | **28458**

ssm-map-policy (Dynamic MLD Interface) | **28459**

ssm-map-policy (IGMP) | **28461**

ssm-map-policy (MLD) | **28462**

ssm-map (Protocols IGMP) | **28464**

ssm-map (Protocols IGMP AMT) | **28465**

ssm-map (Protocols MLD) | **28466**

ssm-map (Routing Options Multicast) | **28468**

stacked-interface-set (Dynamic Profiles) | **28469**

stacked-vlan-ranges | **28471**

stacked-vlan-ranges (RADIUS Options) | **28473**

stacked-vlan-tagging | **28475**

stale-labels-holddown-period | **28476**

stale-routes-time | **28478**

stale-routes-time (Fabric Control) | **28479**

stale-time (Long-Lived Graceful Restart for BGP Restarter) | **28481**

stamp (Sampling) | **28484**

standby | **28485**

standby | **28486**

standby-path-creation-delay | **28488**

standby (Protocols Layer 2 Circuit) | **28489**

standby (Routing Policy Condition) | **28491**

start-date | **28492**

start-log | **28493**

start-measurement | **28494**

start (Scripts) | **28496**

start shell | **28498**

start-time | **28500**

start-time (Accounting) | **28502**

start-time (Schedulers) | **28503**

start-time (Security IDP) | **28504**

starts-with (DHCP Relay Agent Option) | **28506**

starts-with (Event Policy) | **28510**

startup-alarm | **28511**

startup-silent-period | **28513**

stateful (Router Advertisement Guard) | **28514**

stateful (Precision Time Protocol) | **28516**

stateful-firewall-rule-set (Next Gen Services) | **28518**

stateful-firewall-rule-sets (Service Set Next Gen Services) | **28519**

stateful-firewall-rules | **28520**

stateful-firewall-rules (Service Set Next Gen Services) | **28521**

stateful-nat64 | **28523**

stateless | **28524**

static (802.1X) | **28526**

static-binding | **28528**

static (Bridge Domains) | **28529**

static (Distributed IGMP) | **28531**

static (Dynamic IGMP Interface) | **28532**

static (Dynamic MLD Interface) | **28534**

static-host-mapping | **28535**

static (IGMP Snooping) | **28537**

static-interface | **28539**

static-ip | **28540**

static-ip (MX Series) | **28542**

static-ipv6 | **28544**

static-label-switched-path | **28545**

static-lsp | **28547**

static-mac | **28549**

static-nat | **28551**

static (Origin Validation for BGP) | **28553**

static-policy-control | **28555**

static (Protocols IGMP) | **28557**

static (Protocols Layer 2 Circuit) | **28559**

static (Protocols MLD) | **28560**

static (Protocols PIM) | **28562**

static (Protocols VPLS) | **28564**

static-remote-vtep-list | **28566**

static (Routing Options) | **28570**

static (Security NAT) | **28577**

static-subscriber-application | **28579**

static-subscribers | **28581**

static-subscribers (Dynamic Service Provisioning) | **28582**

static-tei-val | **28584**

static-umh (MBGP MVPN) | **28585**

static-mac-ip | **28587**

static-multihoming-peer | **28589**

statistics | **28591**

statistics (Access Profile) | **28592**

statistics (Protocols MPLS) | **28594**

statistics-service | **28596**

statistics (Service Accounting) | **28598**

statistics (Services) | **28599**

stats-aggregation-interval (programmable-rpd) | **28601**

stats-sampling-interval (programmable-rpd) | **28602**

status | **28603**

status-control | **28604**

steering | **28605**

step-percent (RFC 2544 Benchmarking) | **28607**

stickydr | **28609**

stop-date | **28610**

stop-time | **28612**

storage | **28613**

store | **28615**

storm-control (Video Monitoring) | **28619**

storm-control | **28620**

storm-control (switch) | **28622**

storm-control-profiles | **28624**

stp | **28625**

stream (Logical Systems and Tenant Systems) | **28628**

stream (Next Gen Services Service-Set Remote System Logging) | **28631**

stream-option (IDS Screen Next Gen Services) | **28633**

stream-protection (Multicast-Only Fast Reroute) | **28634**

stream (Security Log) | **28635**

streaming-server (Junos Telemetry Interface) | **28640**

streaming-servers | **28642**

strict-asla-based-flex-algorithm | **28645**

strict (DHCP Local Server) | **28646**

strict-dual-isis (Protocols IS-IS) | **28648**

strict-priority-scheduler | **28650**

strict-source-route-option (IDS Screen Next Gen Services) | **28652**

strict-syn-check | **28653**

strict-targeted-hellos | **28654**

string | **28655**

strip-domain (Domain Map) | **28657**

strip-nexthop | **28658**

strip-user-name (Access Profile) | **28659**

strip-username (Domain Map) | **28661**

strip-as-path | **28662**

structured-data | **28664**

stub | **28665**

stub-network | **28667**

sub-domain | **28668**

sub-prefix-length | **28674**

subnet-address (Easy EVPN LAG Configuration) | **28675**

subscriber (Access Profile) | **28680**

subscriber-address | **28683**

subscriber-aware-services | **28684**

subscriber-awareness (Service Set Options) | **28685**

subscriber-context (AAA Options) | **28687**

subscriber (DDoS Flow Detection) | **28688**

subscriber-exclude-prefix | **28691**

subscriber-extension | **28692**

subscriber-leave-timer | **28694**

subscriber-management-helper | **28695**

subscriber-management | **28697**

subscriber-packet-idle-timeout | **28699**

subscriber-profile | **28700**

subscriber-type (TDF Domain) | **28702**

subscriber-info | **28704**

subscriber-pool-utilization-alarm | **28705**

subscribers-limit (Resource Monitor) | **28707**

subscription | **28709**

subscription-id | **28711**

subscription-id-data-include (PCRF Partition) | **28713**

subscription-id-options | **28715**

subscription-id-type (Class Attribute) | **28717**

subscription-id-type (PCRF Partition) | **28719**

success | **28720**

sunrpc | **28722**

supplementary-blackout-timer (Host Fast Reroute) | **28724**

supplicant | **28726**

suppress-host-routes-from-dci-to-dc | **28728**

suppress-interface-leaf | **28730**

suppression | **28731**

surf-control-integrated | **28733**

suspend-for | **28735**

suspend-on-request | **28737**

swap (protocols mpls) | **28738**

swap | **28740**

swap-by-poppush | **28741**

swap (Dynamic VLANs) | **28742**

swap-push | **28744**

swap-swap | **28745**

switch-away-lsps | **28746**

switch-options (Interfaces) | **28748**

switch-options | **28749**

switch-options (VLANs) | **28751**

switch-port | **28753**

switch (Protocols OpenFlow) | **28755**

switch-type | **28757**

switching-type | **28758**

switchover-delay | **28759**

switchover-mode | **28761**

switchover-on-routing-crash | **28763**

sxl-retry | **28764**

sxl-timeout | **28765**

symbol-period | **28767**

symmetric-hash | **28768**

syn-ack-ack-proxy | **28770**

syn-ack-ack-proxy (IDS Screen Next Gen Services) | **28771**

syn-check-required | **28773**

syn-cookie (IDS MS-DPC) | **28774**

syn-fin | **28776**

syn-fin (IDS Screen Next Gen Services) | **28777**

syn-flood | **28778**

syn-flood-protection-mode | **28779**

syn-frag | **28781**

syn-frag (IDS Screen Next Gen Services) | **28782**

sync-active-path-bandwidth | **28783**

sync-reset | **28785**

synchronization (ACX Series) | **28787**

synchronization (M Series and T Series) | **28792**

synchronization (MX Series) | **28795**

synchronization (PTX Series) | **28802**

synchronize | **28804**

synchronize (Scripts) | **28806**

synchronous-ethernet-mapping | **28808**

profile-type (PTP) | **28809**

syntonized-e2e-transparent | **28813**

syslog | **28815**

syslog (Services IPsec VPN) | **28816**

syslog (Chassis) | **28817**

syslog (IDS MS-DPC) | **28819**

syslog (Interfaces) | **28820**

syslog-local-system-timestamp | **28821**

syslog (Next Gen Services Service-Set System Logging) | **28823**

syslog (OAM Action) | **28824**

syslog (QFabric System) | **28825**

syslog (Services CoS) | **28827**

syslog (Services L2TP) | **28829**

syslog (Services NAT) | **28830**

syslog (Services Service Set) | **28831**

syslog (Services Stateful Firewall) | **28833**

syslog-subtag | **28834**

syslog (System) | **28836**

system (Analytics Resource) | **28840**

system-defaults | **28842**

system (Export Profiles) | **28844**

system-id (Interfaces) | **28845**

system-id | **28847**

system-id (Junos Fusion) | **28848**

system-id (Junos Fusion Satellite Device Cluster) | **28850**

system login (Junos Node Slicing) | **28852**

system (Memory) | **28853**

system-priority (chassis) | **28855**

system-priority (interfaces) | **28857**

system-services (Security Zones Host Inbound Traffic) | **28859**

system-services (Security Zones Interfaces) | **28862**

T | 28864

t1-interval | **28885**

t1-options | **28886**

t1-options (BITS Interfaces Signal Type) | **28887**

t310 | **28889**

t391 | **28890**

t392 | **28892**

t4-interval | **28893**

table | **28895**

tacplus | **28896**

tacplus-options | **28898**

tacplus-server | **28902**

tag | **28904**

tag (Access) | **28905**

tag-attribute (HTTP Header Enrichment) | **28907**

tag-attribute (HTTP Header Enrichment Tag Rule) | **28908**

tag (Configuring Notification Targets) | **28910**

tag-group | **28911**

tag-header (HTTP Header Enrichment) | **28913**

tag (HTTP Header Enrichment) | **28914**

tag-list | **28916**

tag-operation (HTTP Header Enrichment) | **28917**

tag-protocol-id (Dynamic VLANs) | **28919**

tag-protocol-id (TPID to Rewrite) | **28920**

tag-protocol-id TPID(TPIDs Expected to Be Sent or Received) | **28922**

tag (Routing Options) | **28924**

tag-rule (HTTP Header Enrichment) | **28926**

tag-rule (Profiles for HTTP Header Enrichment) | **28928**

tag-rule-set (HTTP Header Enrichment) | **28930**

tag-rule-sets (Service Set) | **28931**

tag-rules (Service Set) | **28932**

tag-separator (HTTP Header Enrichment) | **28934**

tag-value (HTTP Header Enrichment) | **28935**

tag2 (Dynamic Access Routes) | **28937**

tags (Application Identification) | **28938**

talk | **28940**

target | **28941**

target-address (TWAMP) | **28944**

target-address (SNMP) | **28946**

target-attachment-identifier (Protocols VPWS) | **28947**

target-logical-system (Domain Map) | **28949**

target-parameters | **28951**

target-routing-instance (Domain Map) | **28952**

target (Routing Instances MVPN) | **28954**

targeted-broadcast | **28956**

targeted-distribution | **28958**

targeted-distribution (Dynamic Demux Interfaces over Aggregated Ethernet) | **28959**

targeted-distribution (Static Interfaces over Aggregated Ethernet) | **28961**

targeted-hello | **28962**

targeted-options (Grouping Subscribers by Bandwidth Usage) | **28964**

targeted-options (Manual Targeting) | **28966**

targeted-options (PS Interface) | **28968**

targets | **28970**

targets (Unified Edge) | **28971**

tca | **28973**

tcp | **28977**

tcp-aggressive-transmission | **28978**

tcp-encap | **28980**

tcp-fast-open | **28982**

tcp-flags | **28983**

tcp-forwarding (Processes) | **28985**

tcp-forwarding (Remote Device Management) | **28987**

tcp-initial-timeout | **28989**

tcp-keepcnt | **28991**

tcp-keepidle | **28992**

tcp-keepintvl | **28994**

tcp-log | **28995**

tcp-mss (Internet Options) | **28996**

tcp-mss (Protocols BGP) | **28998**

tcp-mss (Security Flow) | **29000**

tcp-mss (Services) | **29002**

tcp-no-flag | **29004**

tcp-no-flag (IDS Screen Next Gen Services) | **29005**

tcp-non-syn | **29006**

tcp-options (Security Policies) | **29008**

tcp-rst | **29012**

tcp (Security IDP Protocol Binding) | **29013**

tcp (Security IDP Signature Attack) | **29014**

tcp (Security Screen) | **29017**

tcp-seq-ooo-window (Security Forwarding Options) | **29019**

tcp-session | **29020**

tcp-session-install-interval (Security Forwarding Options) | **29022**

tcp-session (Service Set Next Gen Services) | **29023**

tcp-sweep | **29025**

tcp-syn-defense (IDS MS-MPC) | **29027**

tcp-syn-fragment-check (IDS MS-MPC) | **29028**

tcp-tickles | **29029**

tcp-tickles (Service Set Next Gen Services) | **29030**

tcp-winnuke-check (IDS MS-MPC) | **29031**

tcp-mss (Interfaces) | **29033**

tcpao-auth-mismatch | **29034**

tdf-interface | **29036**

tdf (Unified Edge) | **29037**

tdm-options (Interfaces) | **29039**

te-class-matrix | **29041**

te-link | **29043**

te-policy | **29044**

tear-drop (IDS Screen Next Gen Services) | **29046**

tei-option | **29047**

telemetries | **29048**

telemetries (PoE) | **29050**

telemetry | **29051**

telemetry | **29054**

telnet | **29056**

telnet (Access) | **29059**

telnet (Services) | **29060**

temperature-sensor | **29063**

template | **29071**

template (Services URL Filter) | **29072**

template (Flow Monitoring IPFIX Version) | **29075**

template (Flow Monitoring Version 9) | **29077**

template (Forwarding Options) | **29079**

template (Forwarding Options Version IPFIX) | **29080**

template-id | **29081**

template (Inline Monitoring) | **29084**

template (LRF Profile) | **29087**

template (LRF Rule) | **29089**

template-profile (Flow Monitoring Logs for NAT) | **29090**

template(Protocols MVPN) | **29092**

template-refresh-rate | **29094**

template-refresh-rate (Services) | **29096**

template-tx-interval (LRF Profile) | **29097**

template-type (Flow Monitoring Logs for NAT) | **29098**

template-type (LRF Profile) | **29100**

templates | **29102**

term (Application Aware Access List) | **29106**

term (IP Security VPN) | **29107**

term (Applications) | **29110**

term (candidate-uplink-policy) | **29111**

term (Captive Portal Content Delivery) | **29113**

term (Dynamic Profiles) | **29115**

term (Firewall Filter) | **29117**

term (HTTP Header Enrichment) | **29120**

term (IDS MS-DPC) | **29123**

term (IDS MS-MPC) | **29125**

term (PCP) | **29128**

term (satellite-policies) | **29130**

term (Services CoS) | **29132**

term (Services NAT) | **29134**

term (Services Stateful Firewall) | **29136**

term (Simple Filter) | **29137**

term (TDF Domain Selection) | **29139**

term (URL Filter) | **29143**

terminal | **29146**

terminate-code | **29147**

terminate (PPPoE Service Name Tables) | **29149**

termination (Services) | **29151**

test | **29152**

test aaa authd-lite user | **29156**

test aaa dhcp user | **29161**

test aaa ppp user | **29169**

test access profile | **29178**

test access radius-server | **29184**

test-agent (Paragon Active Assurance) | **29187**

test configuration | **29190**

test-count | **29192**

test-finish-wait-duration (RFC 2544 Benchmarking) | **29193**

test interface restart-auto-negotiation | **29195**

test-interface (RFC 2544 Benchmarking) | **29196**

test-interval | **29198**

test-iterator-duration (RFC 2544 Benchmarking) | **29200**

test-iterator-pass-threshold (RFC 2544 Benchmarking) | **29201**

test msdp | **29202**

test-name | **29204**

test-only-mode | **29208**

test policy | **29209**

test-profile (RFC 2544 Benchmarking) | **29211**

test (Security IDP) | **29214**

test security utm anti-spam | **29215**

test security utm enhanced-web-filtering url-check | **29219**

test security utm web-filtering profile | **29224**

test services l2tp tunnel | **29227**

test services monitoring rfc2544 | **29230**

test services rpm rfc2544-benchmarking test | **29234**

test-session (Junos OS) | **29237**

test-session (Junos OS Evolved) | **29239**

test-type (RFC 2544 Benchmarking) | **29244**

tests | **29247**

tftp | **29251**

tftp (Security ALG) | **29252**

tftp-server | **29254**

then (IP Security VPN) | **29255**

then (Hierarchical Policer) | **29257**

then (Captive Portal Content Delivery) | **29259**

then (Event Policy) | **29261**

then (Filters) | **29264**

then (Firewall Filters) | **29266**

then (HTTP Header Enrichment) | **29267**

then (IDS MS-DPC) | **29269**

then (IDS MS-MPC) | **29271**

then (LRF rule) | **29274**

then (PCC Rules) | **29276**

then (PCP) | **29278**

then (Policer Action) | **29279**

then (Policers) | **29281**

then (Security Application Firewall) | **29282**

then (Security Destination NAT) | **29284**

then (Security IDP Policy) | **29285**

then (Security Policies) | **29287**

threshold-logging-interval | **29290**

then (Security Source NAT) | **29292**

then (Security Static NAT) | **29293**

then (Services CoS) | **29295**

then (Services CoS Next Gen Services) | **29296**

then (Services NAT) | **29299**

then (Services Stateful Firewall) | **29300**

then (Stateful Firewall Rule Next Gen Services) | **29302**

then (TDF Domain Selection) | **29303**

thermal-health-check | **29305**

three-color-policer | **29309**

three-color-policer (Applying) | **29311**

three-color-policer (Configuring) (EX) | **29313**

three-color-policer (Configuring) | **29314**

threshold | **29317**

threshold (Bridge Domains) | **29319**

threshold-count | **29320**

threshold (Detection Time) | **29322**

threshold (detection-time) | **29323**

threshold (Interfaces) | **29325**

threshold-logging-interval | **29326**

threshold (MSDP Active Source Messages) | **29327**

threshold (Multicast Forwarding Cache) | **29329**

threshold (PIM BFD Detection Time) | **29331**

threshold (PIM BFD Transmit Interval) | **29333**

threshold (PIM Entries) | **29334**

threshold-rate | **29336**

threshold (Routing Instances) | **29338**

threshold (Services IPsec) | **29340**

threshold (Services Logging and SYN-Cookie Defenses) | **29341**

threshold-time | **29343**

threshold (transmit-interval) | **29344**

thresholds | **29346**

thresholds (Junos OS) | **29348**

thresholds (Junos OS Evolved) | **29350**

time | **29352**

time-binding | **29354**

time-format | **29356**

time-interval (Dampen Options) | **29358**

time-limit (LRF Rule) | **29359**

time-wait-state | **29361**

time-zone | **29362**

timeout (SNMP) | **29365**

timeout | **29367**

timeout (Access LDAP) | **29368**

timeout (Access RADIUS) | **29370**

timeout-action | **29371**

timeout-active-flows (DDoS Flow Detection) | **29372**

timeout (DHCP Local Server) | **29374**

timeout (Diameter Network Element) | **29376**

timeout (Flow Maps) | **29377**

timeout-grace (Access) | **29379**

timeout (Multicast) | **29380**

timeout (RADIUS Server) | **29382**

timeout (Security Antivirus Fallback Options) | **29383**

timeout (Security Antivirus Fallback Options Sophos Engine) | **29385**

timeout (Security Antivirus Scan Options) | **29386**

timeout (Security GTP) | **29388**

timeout (Security Screen) | **29389**

timeout (Security Web Filtering) | **29391**

timeout (Security Web Filtering Cache) | **29392**

timeout (Security Web Filtering Fallback Settings) | **29393**

timeout (Services) | **29395**

timestamp | **29396**

timestamp-format (RFC 2544 Benchmarking) | **29398**

timestamp-option (IDS Screen Next Gen Services) | **29399**

tls | **29400**

tls-min-version | **29402**

tls (NETCONF) | **29403**

tls-peer-name | **29405**

tls-timeout | **29406**

tls-type | **29407**

tlsdetails | **29409**

tlv-filter | **29410**

tlv-select | **29414**

tm | **29417**

tm-polarity | **29418**

to | **29419**

to (Security Source NAT) | **29421**

to-zone (Security IDP Policy) | **29422**

to-zone (Security Policies) | **29423**

to-zone (Security Policies Global) | **29427**

token (DHCP Local Server) | **29428**

too-many-requests (Security Antivirus Fallback Options) | **29430**

too-many-requests (Security Antivirus Fallback Options Sophos Engine) | **29432**

too-many-requests (Security Web Filtering Fallback Settings) | **29433**

top | **29435**

topologies (Multitopology Routing) | **29436**

topologies (Protocols IS-IS) | **29438**

topology (Filter-Based Forwarding) | **29440**

topology-id | **29441**

topology (Multitopology Routing) | **29443**

topology (OSPF) | **29445**

topology (OSPF Interface) | **29446**

topology (Protocols BGP) | **29449**

tos-reflect (L2TP) | **29451**

total-memory | **29452**

trace | **29454**

trace (DHCP Local Server) | **29456**

trace (DHCP Relay Agent) | **29458**

trace (EVPN-VXLAN) | **29459**

traceoptions (routing-options) | **29462**

traceoptions (Ethernet Switching) | **29465**

traceoptions (High Availability) | **29471**

traceoptions (802.1X) | **29474**

traceoptions (Access) | **29477**

traceoptions (Access Port Security) | **29479**

traceoptions (Active Directory Access) | **29483**

traceoptions (Address-Assignment Pool) | **29485**

traceoptions (advanced policy-based routing) | **29488**

traceoptions (Analytics) | **29491**

traceoptions (Analytics Agent) | **29492**

traceoptions (ANCP) | **29494**

traceoptions (App Engine Virtual Machine Management Service) | **29497**

traceoptions (Batch Commits) | **29499**

traceoptions (Captive Portal Content Delivery) | **29501**

traceoptions (CASB) | **29503**

traceoptions (Chassis Cluster) | **29506**

traceoptions (Class of Service) | **29508**

traceoptions (Commit, Op, and SNMP Scripts) | **29511**

traceoptions (DDoS) | **29515**

traceoptions (DHCP) | **29517**

traceoptions (DHCP Server) | **29521**

traceoptions (Diameter Base Protocol for Subscriber Aware Policy Control) | **29524**

traceoptions (Diameter Base Protocol) | **29527**

traceoptions (DNS, Port, and TFTP Packet Forwarding) | **29530**

traceoptions (dynamic-address) | **29533**

traceoptions (dynamic-application) | **29535**

traceoptions (Dynamic Flow Capture) | **29538**

traceoptions (Egress Protection) | **29539**

traceoptions (Enhanced Subscriber Management) | **29541**

traceoptions (Event Options) | **29544**

traceoptions (Event Scripts) | **29547**

traceoptions (Extensible Subscriber Services Manager) | **29550**

traceoptions (Extensions) | **29552**

traceoptions (FC-2 Fibre Channel) | **29554**

traceoptions (Fibre Channel) | **29557**

traceoptions (FIP Protocol Fibre Channel) | **29560**

traceoptions (flow) | **29563**

traceoptions (Forwarding Options) | **29565**

traceoptions (General Authentication Service) | **29567**

traceoptions (Health Check Monitoring) | **29569**

traceoptions (IGMP Snooping) | **29573**

traceoptions (Individual Interfaces) | **29576**

traceoptions (Inline Monitoring) | **29585**

traceoptions (Interfaces) | **29587**

traceoptions (Juniper Secure Connect) | **29589**

traceoptions (LACP) | **29591**

traceoptions (Layer 2 Learning) | **29593**

traceoptions (LLDP) | **29596**

traceoptions (MACsec) | **29599**

traceoptions (MACsec interfaces) | **29602**

traceoptions (Multicast Snooping Options) | **29604**

traceoptions (MVRP) | **29607**

traceoptions (MX Series Virtual Chassis) | **29609**

traceoptions (NETCONF and Junos XML Protocol) | **29613**

traceoptions (NETCONF TLS) | **29616**

traceoptions (Next Gen Services Global System Logging) | **29619**

traceoptions (Next Gen Services Service-Set Flow) | **29620**

traceoptions (Next Gen Services Softwires) | **29624**

traceoptions (Origin Validation for BGP) | **29626**

traceoptions (Outbound SSH) | **29627**

traceoptions (OVSDB) | **29630**

traceoptions (PCE) | **29633**

traceoptions (PCEF) | **29635**

traceoptions (PIM Snooping) | **29638**

traceoptions (Port Mirroring and Traffic Sampling) | **29640**

traceoptions (PPP Process) | **29641**

traceoptions (PPPoE) | **29645**

traceoptions (Process Monitor) | **29648**

traceoptions (programmable-rpd) | **29651**

traceoptions (Protocols) | **29654**

traceoptions (Protocols AMT) | **29656**

traceoptions (Protocols BFD) | **29659**

traceoptions (Protocols BGP) | **29662**

traceoptions (Protocols BMP) | **29666**

traceoptions (Protocols DVMRP) | **29669**

traceoptions (Protocols EVPN) | **29672**

traceoptions (Protocols IGMP) | **29675**

traceoptions (Protocols IGMP Snooping) | **29678**

traceoptions (Protocols IPv6 Neighbor Discovery) | **29681**

traceoptions (Protocols IS-IS) | **29684**

traceoptions (Protocols Layer 2 Circuit) | **29688**

traceoptions (Protocols Layer 2 VPN) | **29690**

traceoptions (Protocols LDP) | **29693**

traceoptions (Protocols Link Management) | **29697**

traceoptions (Protocols MPLS) | **29699**

traceoptions (Protocols MSDP) | **29703**

traceoptions (Protocols MVPN) | **29706**

traceoptions (Protocols OpenFlow) | **29710**

traceoptions (Protocols OSPF) | **29713**

traceoptions (Protocols PCEP) | **29717**

traceoptions (Protocols PIM) | **29719**

traceoptions (Protocols PPP Service) | **29723**

traceoptions (Protocols RIP) | **29727**

traceoptions (Protocols RIPng) | **29730**

traceoptions (Protocols Router Discovery) | **29733**

traceoptions (Protocols RSVP) | **29736**

traceoptions (Protocols Secure Neighbor Discovery) | **29739**

traceoptions (Protocols Spring-TE) | **29742**

traceoptions (Protocols VPLS) | **29744**

traceoptions (Protocols VRRP) | **29747**

traceoptions (Proxy Fibre Channel) | **29750**

traceoptions (Resiliency) | **29752**

traceoptions (Resource Cleanup) | **29755**

traceoptions (Resource Monitor) | **29758**

traceoptions (REST API) | **29760**

traceoptions (RPM) | **29761**

traceoptions (SBC Configuration Process) | **29764**

traceoptions (Security) | **29767**

traceoptions (Security ALG) | **29770**

traceoptions (Security Antispam) | **29773**

traceoptions (Security Antivirus) | **29774**

traceoptions (Security Application Firewall) | **29776**

traceoptions (Security Application Proxy) | **29778**

traceoptions (Security Content Filtering) | **29781**

traceoptions (Security Content Security) | **29782**

traceoptions (Security Datapath Debug) | **29783**

traceoptions (Security Dynamic VPN) | **29786**

traceoptions (Security Firewall Authentication) | **29788**

traceoptions (Security Flow) | **29789**

traceoptions (Security Group VPN) | **29795**

traceoptions (Security GTP) | **29798**

traceoptions (Security H323 ALG) | **29801**

traceoptions (Security IDP) | **29803**

traceoptions (Security IKE) | **29806**

traceoptions (Security IPsec) | **29809**

traceoptions (Security Log) | **29811**

traceoptions (Security MGCP ALG) | **29814**

traceoptions (Security NAT) | **29815**

traceoptions (Security PKI) | **29818**

traceoptions (Security Policies) | **29821**

traceoptions (Security SCCP ALG) | **29823**

traceoptions (Security Screen) | **29825**

traceoptions (Security SCTP) | **29828**

traceoptions (Security SIP ALG) | **29830**

traceoptions (Security User Identification) | **29832**

traceoptions (Security Web Filtering) | **29835**

traceoptions (Services) | **29836**

traceoptions (Services Application Identification) | **29839**

traceoptions (Services IPsec VPN) | **29842**

traceoptions (Services L2TP) | **29844**

traceoptions (Services Logging) | **29849**

traceoptions (Services Redundancy Daemon) | **29851**

traceoptions (Services SSL) | **29854**

traceoptions (Services UAC) | **29857**

traceoptions (Services User Identification) | **29859**

traceoptions (sFlow Technology) | **29862**

traceoptions (SMTP) | **29865**

traceoptions (SNMP) | **29866**

traceoptions (Spanning Tree) | **29869**

traceoptions (Static Subscribers) | **29873**

traceoptions (Subscriber Management) | **29876**

traceoptions (Subscriber Session Database Replication) | **29878**

traceoptions (System Accounting) | **29881**

traceoptions (System License) | **29883**

traceoptions (System Monitor) | **29885**

traceoptions (System Services DNS) | **29888**

traceoptions (TCP Encapsulation) | **29891**

traceoptions (TCP Port Forwarding) | **29894**

traceoptions (TDF Gateway) | **29896**

traceoptions (Traffic Load Balancer) | **29899**

traceoptions (Uplink Failure Detection) | **29903**

traceoptions (Virtual Chassis) | **29905**

traceroute | **29909**

traceroute clns | **29917**

traceroute ethernet | **29919**

traceroute fabric unicast-flow | **29924**

traceroute monitor | **29927**

traceroute mpls bgp | **29930**

traceroute mpls ldp | **29935**

traceroute mpls rsvp | **29940**

traceroute mpls segment-routing isis | **29946**

traceroute mpls segment-routing ospf | **29953**

traceroute mpls segment-routing spring-te | **29959**

traceroute mpls segment-routing static egress-ip | **29970**

traceroute overlay | **29976**

traceroute srv6 | **29982**

tracing | **29987**

track-igp-metric | **29989**

track-igp-metric (LSP) | **29990**

track (VRRP) | **29992**

tracking-area | **29994**

traffic-class (delay) | **29996**

traffic-class (loss) | **29998**

traffic-class (loss-delay) | **30000**

traffic-class-map | **30003**

traffic-class-map (Apply to Interface) | **30006**

traffic-class (Tunnels) | **30008**

traffic-control-profiles | **30009**

traffic-control-profiles (Dynamic CoS Definition) | **30012**

traffic-engineering | **30014**

traffic-engineering (policy-options) | **30016**

traffic-engineering (OSPF) | **30018**

traffic-engineering (Passive TE Mode) | **30021**

traffic-engineering (Protocols BGP) | **30023**

traffic-engineering (Protocols IS-IS) | **30025**

traffic-engineering (Protocols MPLS) | **30031**

traffic-engineering (SR-TE import policy) | **30032**

traffic-load-balance (Traffic Load Balancer) | **30034**

traffic-manager | **30036**

traffic-options | **30040**

traffic-selector | **30042**

traffic-statistics (Services Analytics) | **30045**

traffic-statistics (Protocols IS-IS) | **30048**

traffic-statistics-labeled-path | **30049**

traffic-statistics (Protocols BGP) | **30051**

traffic-statistics (Protocols LDP) | **30053**

transaction-timeout | **30055**

transfer | **30056**

transfer-delay (Event Policy) | **30058**

transfer-interval | **30059**

transfer-log-archive | **30061**

transient-filter | **30062**

transit | **30064**

transit (Chained Composite Next Hops) | **30066**

transit-lsp-association | **30069**

translated | **30070**

translated-port | **30072**

translation | **30073**

translation-table | **30075**

translation-type | **30077**

translation-vni | **30080**

transmit-clock | **30083**

transmit-failure-threshold (RFC 2544 Benchmarking) | **30084**

transmit-interval | **30085**

transmit-interval (BFD Liveness Detection) | **30087**

transmit-interval (LAG) | **30090**

transmit-interval (Liveness Detection) | **30092**

transmit-interval (MACsec) | **30093**

transmit-interval (MACsec for MX Series) | **30095**

transmit-interval (PIM BFD Liveness Detection) | **30096**

transmit-rate | **30098**

transmit-rate (Dynamic Schedulers) | **30103**

transmit-rate (EX Series Switches) | **30105**

transmit-rate (Schedulers) | **30107**

transparent | **30110**

transport (Services) | **30111**

transport (System Syslog) | **30112**

transport (Protocols BGP) | **30113**

transport-address | **30118**

transport-class | **30120**

transport (Diameter Base Protocol) | **30123**

transport (Diameter Peer) | **30124**

transport-monitoring | **30126**

transport (Next Gen Services Syslog Message Security) | **30127**

transport (Security Log) | **30129**

transport-type | **30131**

transport-class | **30132**

trap | **30133**

trap-group | **30135**

trap-options | **30136**

trap (Security IKE) | **30139**

trap (Security PKI) | **30140**

traps | **30142**

traps (RPM and TWAMP) | **30144**

tri-color | **30147**

tri-color (EX Series Switches) | **30149**

trickling | **30152**

trigger | **30154**

trigger (CoS) | **30159**

trigger (DHCP Local Server) | **30160**

trigger (Event Policy) | **30162**

trigger-link-failure | **30164**

trigger-type (LRF Profile) | **30165**

trio-flow-offload | **30167**

trust-option-82 | **30168**

trusted | **30170**

trusted-ca | **30171**

trusted-ca (Services) | **30172**

trusted (DHCP Security) | **30174**

trusted-keys (DNSSEC) | **30175**

tti | **30176**

ttl | **30178**

ttl (Dynamic Flow Capture) | **30179**

ttl (RPM probe) | **30180**

ttl-threshold | **30183**

tunable-name | **30184**

tunable-value | **30186**

tunables | **30187**

tunnel | **30189**

tunnel | **30190**

tunnel | **30192**

tunnel-attributes | **30194**

tunnel-devices (Protocols AMT) | **30196**

tunnel-devices (Tunnel-Capable PICs) | **30198**

tunnel-end-point | **30200**

tunnel (fti) | **30204**

tunnel-group | **30206**

tunnel-inspection | **30208**

tunnel (L2TP) | **30210**

tunnel-limit (Protocols AMT) | **30212**

tunnel-limit (Routing Instances) | **30213**

tunnel-limit (Routing Instances Provider Tunnel Selective) | **30215**

tunnel-mtu | **30216**

tunnel-mtu (Services IPsec VPN) | **30218**

tunnel-mtu (Services Service Set) | **30219**

tunnel-observation | **30221**

tunnel-port | **30223**

tunnel-port (GRE) | **30224**

tunnel-profile (Domain Map) | **30225**

tunnel-profile (L2TP Tunnel Switching) | **30226**

tunnel-profile (Tunnel Profile) | **30228**

tunnel-queuing | **30230**

tunnel (Security Policies) | **30231**

tunnel (Security Screen) | **30232**

tunnel-services (Chassis) | **30234**

tunnel-services (Routing Instances VPLS) | **30236**

tunnel-services (RSVP) | **30238**

tunnel-source | **30239**

tunnel-source-protocol (Protocols OSPF) | **30240**

tunnel-switch-profile (L2TP Tunnel Switching, Application) | **30242**

tunnel-switch-profile (L2TP Tunnel Switching, Definition) | **30243**

tunnel-termination | **30245**

tunnel-timeout | **30247**

tunnel (Tunnel Profile) | **30248**

twamp | **30250**

twamp (Security ALG) | **30256**

twamp-server | **30258**

two-rate | **30259**

tx-address-change (L2TP LAC) | **30261**

tx-buffers | **30263**

tx-connect-speed-method (L2TP LAC) | **30265**

tx-dnu-to-line-source-enable | **30268**

tx-duration | **30269**

tx-enable | **30271**

tx-power | **30272**

type | **30273**

type (Application Identification) | **30275**

type (ICMP Mapping for Application Identification) | **30276**

type (Protocols BGP) | **30278**

type (RMON) | **30279**

type (RMON Notification) | **30281**

type (Security Antivirus Feature Profile) | **30282**

type (Security Content Filtering Notification Options) | **30284**

type (Security Fallback Block) | **30285**

type (Security IDP Dynamic Attack Group) | **30287**

type (Security IDP ICMP Headers) | **30288**

type (Security Virus Detection) | **30289**

type (Security Web Filtering) | **30291**

type (SNMP v3) | **30293**

type (Tunnel Profile) | **30294**

U | 30295

uac-policy (Application Services) | **30301**

uac-service | **30303**

ucode-imem-remap | **30304**

udp | **30306**

udp-encapsulation | **30307**

udp (FTI) | **30309**

udp-port | **30310**

udp (Security IDP Protocol Binding) | **30312**

udp (Security IDP Signature Attack) | **30313**

udp (Security Screen) | **30314**

udp-sweep | **30316**

udp-tcp-port-swap | **30318**

uid (Dynamic Profiles) | **30319**

uid-reference | **30320**

ukern-trace | **30322**

ultimate-hop-popping | **30323**

ultrahigh-power | **30325**

unattended-boot | **30326**

unconditional-src-learn | **30328**

unconfigured-peer-graceful-restart | **30329**

underlying-interface | **30330**

underlying-interface (ANCP) | **30332**

underlying-interface (demux0) | **30333**

underlying-interface (Dynamic PPPoE) | **30335**

unframed | no-unframed (Interfaces) | **30337**

unicast-ethernet | **30338**

unicast-ethernet-ipv4 | **30340**

unicast-in-lpm | **30342**

unicast-mode | **30344**

unicast-negotiation | **30347**

unicast-reverse-path | **30348**

unicast (Route Target Community) | **30350**

unicast-stream-limit (Protocols AMT) | **30351**

unicast-umh-election | **30353**

unicast (Virtual Tunnel in Routing Instances) | **30354**

unicast-tunnel-name-regular-expression | **30355**

unidirectional | **30356**

unidirectional-session-refreshing | **30358**

unified-access-control (Services) | **30359**

unified-policy-explicit-match | **30361**

unified-policy max-lookups | **30362**

unique-nas-port (Access) | **30364**

unit (Sampling) | **30366**

unit (QFX Series) | **30367**

unit (interfaces) | **30370**

unit (Aggregated Multiservices) | **30380**

unit (CoS Interfaces) | **30382**

unit (CoS Interfaces) | **30384**

unit (Dynamic Demux Interface) | **30385**

unit (Dynamic Interface Sets) | **30388**

unit (Dynamic PPPoE) | **30390**

unit (Dynamic Profiles Standard Interface) | **30392**

unit (Dynamic Traffic Shaping) | **30397**

unit (Interfaces Services) | **30399**

unit (Interfaces) | **30401**

unit (Link-Multilink) | **30402**

unit (Tunnel) | **30404**

unit (TDF Interface) | **30406**

unit (Voice Services) | **30408**

unknown-message (Security H323 ALG) | **30410**

unknown-message (Security MGCP ALG) | **30411**

unknown-message (Security SCCP ALG) | **30413**

unknown-message (Security SIP ALG) | **30415**

unknown-protocol (IDS Screen Next Gen Services) | **30416**

unknown-unicast-forwarding | **30417**

unnumbered-address (Demux) | **30419**

unnumbered-address (Dynamic PPPoE) | **30421**

unnumbered-address (Dynamic Profiles) | **30422**

unnumbered-address (Ethernet) | **30425**

unnumbered-address (PPP) | **30426**

unprotect | **30428**

unreachable-vtep-aging-timer | **30429**

untagged | **30431**

untrusted | **30432**

untrusted (VLAN) | **30433**

Unused | **30435**

up | **30436**

update | **30438**

update-interval | **30439**

update-interval (Protocols RIP) | **30441**

update-interval (Protocols RIPng) | **30443**

update-interval (Service Accounting) | **30444**

update-rate-limit | **30446**

update-response-timeout (PCRF Partition) | **30447**

update-router-advertisement | **30449**

update-server | **30450**

update-server (dhcp-client) | **30452**

update-server (dhcpv6-client) | **30453**

update-threading | **30454**

update-threshold | **30456**

update-threshold-max-reservable | **30458**

upgrade-group | **30460**

upgrade-groups (Junos Fusion) | **30461**

upgrade-mode | **30464**

uplink-failure-detection | **30465**

uplink-failure-detection (Junos Fusion) | **30467**

uplink-port-group (Junos Fusion) | **30469**

upload (Event Policy) | **30471**

upload-profile (Security Antivirus FTP) | **30472**

upload-profile (Security Content Filtering FTP) | **30473**

upstream-interface | **30474**

upstream-label | **30476**

upstream-mode | **30478**

upstream-rate (Traffic Shaping) | **30479**

upstream-rewrite | **30481**

urgent-pointer | **30483**

uri-check | **30484**

url | **30485**

url (Security) | **30487**

url-blacklist | **30488**

url-feed | **30489**

url-filter | **30492**

url-filter-profile | **30495**

url-filter-template | **30496**

url-pattern | **30499**

url (Security) | **30503**

url (Security Antivirus) | **30504**

url (Security IDP) | **30506**

url-whitelist | **30507**

url-whitelist | **30508**

usb | **30509**

usb-control | **30510**

use-active-child-mac-on-reth | **30511**

use-actual-mac-on-physical-interfaces | **30512**

use-class (Class Attribute) | **30514**

use-extended-flow-memory | **30515**

use-for-post-convergence-lfa | **30517**

use-imported-time-zones | **30518**

use-interface | **30519**

use-interface-description (MX) | **30520**

use-interface-description | **30522**

use-interface-index | **30525**

use-interface-name | **30526**

use-option-82 | **30528**

use-p2mp-lsp | **30530**

use-post-convergence-lfa | **30531**

use-post-convergence-lfa (Protocols OSPF) | **30533**

use-primary (DHCP Local Server) | **30535**

use-primary (DHCP Relay Agent) | **30537**

use-string | **30539**

use-transport-class | **30541**

use-underlying-interface-mac | **30542**

use-vlan-id | **30544**

use-vlan-id (DHCP Relay Agent) | **30546**

user (Access) | **30547**

user-defined-option-82 | **30551**

user-firewall | **30553**

user-group-mapping | **30555**

user-id | **30557**

user-identification | **30558**

user-identification (Services) | **30560**

user-name | **30568**

user-name (Event Policy) | **30569**

user-name-include (OCS Partition) | **30571**

user-password (PCEF Profile) | **30574**

user-planes | **30575**

user-prefix | **30577**

user-prefix (DHCP Local Server) | **30578**

user-prefix (DHCP Relay Agent) | **30580**

user-prefix (Static Subscribers) | **30582**

user-query (Services User Identification) | **30584**

user (SNMP v3) | **30588**

user (System Logging) | **30589**

user (System Services) | **30591**

username-include (Demux) | **30593**

username-include (DHCP Local Server) | **30595**

username-include (DHCP Relay Agent) | **30597**

username-include (Interfaces) | **30600**

username-include (Local Authentication) | **30602**

username-include (Static Subscribers) | **30604**

username (Security Antivirus) | **30606**

username (Services) | **30607**

using-user-password | **30609**

usm | **30610**

utc-timestamp | **30612**

utm default-configuration | **30614**

uuid | **30621**

uuid (Applications) | **30623**

V | 30624

v3 | **30631**

v4address | **30635**

v4ov6 | **30636**

v6address | **30638**

v6prefix | **30639**

v6rd | **30641**

vacm | **30642**

valid-lifetime | **30644**

valid-lifetime (Dynamic Router Advertisement) | **30645**

validation (Origin Validation for BGP) | **30647**

validity-period | **30650**

variable | **30651**

variable | **30653**

variables (Dynamic Service Profiles) | **30654**

variant | **30655**

vbr rate | **30657**

vc-port | **30658**

vci | **30660**

vcp-no-hold-time | **30662**

vcp-snmp-statistics | **30664**

vdsl-bytes | **30665**

vdsl-overhead-adjust | **30667**

vdsl-profile | **30669**

vdsl2-bytes | **30671**

vdsl2-overhead-adjust | **30673**

vendor | **30675**

vendor-id (DHCP Client) | **30676**

vendor-id (dhcp-option82) | **30677**

vendor-id (Unified Edge) | **30679**

vendor-id (AAA Profile) | **30681**

vendor-id (Interfaces) | **30682**

vendor-option | **30684**

vendor-specific (DHCP Relay Agent) | **30686**

vendor-support | **30687**

verify-path | **30688**

version-3 | **30691**

version (Flow Aggregation) | **30692**

version-alias (Dynamic Profiles) | **30694**

version (BFD) | **30695**

version (BFD Liveness Detection) | **30697**

version (Dynamic IGMP Interface) | **30699**

version (Dynamic MLD Interface) | **30701**

version (Flow Monitoring Logs for NAT) | **30702**

version (IGMP Snooping) | **30704**

version (IKE) | **30706**

version-ipfix (Forwarding Options) | **30707**

version-ipfix (Services) | **30708**

version (Liveness Detection) | **30710**

version (MLD Snooping) | **30712**

version (Protocols BFD) | **30714**

version (Protocols IGMP) | **30715**

version (Protocols IGMP AMT) | **30716**

version (Protocols MLD) | **30718**

version (Protocols PIM) | **30719**

version (SNMP) | **30721**

version9 (Flow Monitoring) | **30722**

version9 (Forwarding Options) | **30724**

versioning | **30725**

video | **30727**

video (Application Profile) | **30728**

video-monitoring | **30729**

view (Configuring a MIB View) | **30733**

view (SNMP Community) | **30734**

violation-action (DHCP Local Server and DHCP Relay Agent) | **30735**

violation-report-rate (DDoS Flow Detection) | **30737**

virtual-address | **30739**

virtual-channel-group (CoS Interfaces) | **30740**

virtual-channel-groups | **30742**

virtual-channels | **30743**

virtual-chassis | **30745**

virtual-chassis (MX Series Virtual Chassis) | **30747**

virtual-circuit-id | **30749**

virtual-control-channel | **30751**

virtual-cpu | **30752**

virtual-gateway-address | **30754**

virtual-gateway-v4-mac | **30755**

virtual-gateway-v6-mac | **30757**

virtual-inet6-address | **30760**

virtual-ip | **30761**

virtual-link | **30763**

virtual-link-local-address | **30765**

virtual-mac | **30766**

virtual-network-functions | **30768**

virtual-network-functions (Junos node slicing) | **30773**

virtual-service (Traffic Load Balancer) | **30774**

virtual-switch | **30777**

virus-detection (Security Antivirus) | **30778**

vjunos0 | **30780**

vlan | **30782**

vlan (Access Port Security) | **30785**

vlan (Bridge Domains) | **30788**

vlan (DHCP Bindings on Access Ports) | **30790**

vlan-group | **30791**

vlan-groups | **30792**

vlan-id (Fast Ethernet) | **30794**

vlan-id | **30795**

vlan-id (Dynamic Profiles) | **30801**

vlan-id (Dynamic VLANs) | **30802**

vlan-id inner-all | **30804**

vlan-id-list | **30805**

vlan-id-list (Ethernet VLAN Circuit) | **30808**

vlan-id-list (Interface in Bridge Domain) | **30810**

vlan-id-list (Interface in VPLS) | **30812**

vlan-id (Logical Port in Bridge Domain) | **30813**

vlan-id-range | **30815**

vlan-id (routing instance) | **30817**

vlan-id-start | **30818**

vlan-id (VLAN ID to Be Bound to a Logical Interface) | **30820**

vlan (IGMP Snooping) | **30821**

vlan-isolation | **30826**

vlan members (VLANs) | **30827**

vlan (Mirroring) | **30829**

vlan (MLD Snooping) | **30830**

vlan-model | **30833**

vlan (MSTP) | **30835**

vlan (PIM Snooping) | **30837**

vlan (Port Mirroring) | **30839**

vlan (Port Mirroring) | **30841**

vlan-prune | **30842**

vlan-range | **30843**

vlan-ranges | **30845**

vlan-ranges (RADIUS Options) | **30847**

vlan-rewrite | **30849**

vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) | **30850**

vlan (Secure Access Port) | **30852**

vlan (Static IP) | **30854**

vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) | **30855**

vlan-tag | **30857**

vlan-tag (Dynamic Classifiers) | **30858**

vlan-tag (Dynamic Rewrite Rules) | **30860**

vlan-tagging | **30861**

vlan-tagging (Dynamic) | **30864**

vlan-tags | **30865**

vlan-tags (Dynamic Profiles) | **30867**

vlan-tags (Dual-Tagged Logical Interface) | **30869**

vlan-tags-outer | **30872**

vlan-tags (Stacked VLAN Tags) | **30873**

vlan (Unknown Unicast Forwarding) | **30876**

vlan (VSTP) | **30878**

vlans | **30880**

vlans (RA Guard) | **30894**

vmhost | **30895**

vmhost management-if disable | **30899**

vmhost management-if link-mode | **30900**

vmhost management-if speed | **30902**

vmhost resize vjunos compact | **30904**

vnf-license-supplement | **30905**

vnf-name | **30906**

vni | **30909**

vni(Interfaces) | **30911**

vni-options | **30912**

vnid (EVPN) | **30914**

voice | **30915**

voice (Application Profile) | **30917**

voip | **30918**

voip-mac-exclusive | **30920**

volume-limit (LRF Rule) | **30921**

vpi (Define Virtual Path) | **30923**

vpls-flow-table-size | **30924**

vpls-flush-on-topology-change | **30926**

vpls-id | **30927**

vpls-id-list (protocols vpls mesh-group) | **30929**

vpls (Interfaces) | **30931**

vpls-mac-move | **30932**

vpls (Routing Instance) | **30933**

vpls-template | **30937**

vpn-apply-export | **30938**

vpn-group-address | **30939**

vpn-label | **30941**

vpn-localization | **30942**

vpn-monitor | **30943**

vpn (Security) | **30946**

vpn-unequal-cost | **30951**

vpn-global-import | **30953**

vpws-service-id | **30954**

vrf-advertise-selective | **30957**

vrf-export | **30958**

vrf-import | **30960**

vrf-mtu-check | **30962**

vrf-name (Duplicate Accounting) | **30963**

vrf-propagate-ttl | **30965**

vrf-table-label | **30966**

vrf-target | **30968**

vrrp | **30971**

vrrp-group | **30973**

vrrp-inet6-group | **30976**

vrrp-inherit-from | **30978**

vsi-discovery | **30980**

vsi-policy | **30981**

vstp | **30983**

vtep-source-interface | **30988**

vulnerability-type | **30990**

vxlan | **30991**

vxlan-disable-copy-tos-decap | **30993**

vxlan-disable-copy-tos-encap | **30995**

vxlan-extended | **30998**

vxlan-gbp-l2-profile | **30999**

vxlan-gbp-l3-profile | **31001**

vxlan-gbp-profile | **31003**

vxlan-gpe (FTI) | **31005**

vxlan-routing | **31007**

vxlan-stitching | **31009**

vxlan-tos-copy-filter | **31011**

W | 31012

wait-for-acct-on-ack (Access Profile) | **31014**

wait-to-block-interval | **31016**

wait-to-restore | **31017**

walkup | **31019**

wander-filter-disable | **31021**

warm-standby (Interfaces) | **31022**

warm-standby | **31023**

warning | **31024**

watch-list | **31026**

watchdog | **31027**

watchdog (PSM) | **31028**

watchdog-timeout | **31031**

wavelength | **31032**

web-authentication | **31057**

web-authentication (Access) | **31058**

web-authentication (Interfaces) | **31060**

web-filter | **31062**

web-filter-profile | **31065**

web-filtering | **31066**

web-filtering (Logical System Security Feature Profile) | **31072**

web-filtering (Security Content Security Policy) | **31076**

web-management (System Processes) | **31078**

web-management (System Services) | **31079**

web-redirect | **31084**

web-redirect-to-https | **31085**

web-server (Services) | **31086**

webapi-clear-text (Security) | **31089**

webapi-ssl (Security) | **31090**

webapi (System Services) | **31091**

websense-redirect | **31093**

weight | **31095**

weight (Targeted Subscribers) | **31097**

weight (Security) | **31099**

weighted-load-balancing (L2TP LAC) | **31100**

west-interface | **31102**

when | **31104**

white-list | **31106**

white-list (Security Screen) | **31107**

white-list (Security-Zone) | **31109**

whitelist | **31111**

whitelist-url-categories | **31112**

wildcard delete | **31114**

wildcard-group-inet | **31115**

wildcard-group-inet6 | **31117**

wildcard-source (PIM RPF Selection) | **31119**

wildcard-source (Selective Provider Tunnels) | **31120**

window-scale | **31122**

window-size | **31123**

windows-logon (Juniper Secure Connect) | **31124**

winnuke | **31126**

winnuke (IDS Screen Next Gen Services) | **31127**

wins-server (Access) | **31129**

wins-server (System) | **31130**

wired-max-processes | **31131**

wired-process-mem-size | **31133**

withdraw-priority | **31134**

within (Event Policy) | **31136**

wlan | **31137**

world-readable (Forwarding Options) | **31142**

world-readable (System Syslog) | **31144**

world-readable (Next Gen Services Global System Logging) | **31145**

write-interval | **31146**

write-view | **31148**

X | 31150

xauth-attributes | **31150**

xe (Port) | **31152**

xlat-source-rule | **31154**

xle (Port) | **31155**

xnm-clear-text | **31156**

xnm-ssl | **31158**

xon (Input Congestion Notification) | **31161**

Y | 31163

yang-compliant (NETCONF) | **31163**

yang-modules (NETCONF) | **31166**

yellow-differential-delay | **31168**

Z | 31169

zone (System Security Profile) | **31170**

zones | **31172**

Others | 31175

| (pipe) | **31175**

Junos CLI Reference

Overview

Use this guide to learn about the components that make up the Junos CLI commands and configuration statements and the contexts in which you'll use these CLI elements in your network configurations and operations. In this reference, you'll find all Junos CLI commands and statements in one place.

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- "Configuration Statements" on page 1258
- "Operational Commands" on page 1487

1

CHAPTER

What's New in the CLI

What's New in the CLI | 2

What's New in the CLI

IN THIS SECTION

- 24.4 | 2
- 24.2 | 4
- 23.4 | 6
- 23.2 | 9
- 22.4 | 12

Here are the new commands and statements added to the CLI in the last five Junos OS and Junos OS Evolved releases.

24.4

Configuration Statements		
"identity-provider attribute-mapping (SAML)" on page 2501	"buffer-rate (Schedulers)" on page 3065	"chassis" on page 3296
"csds" on page 4729	"disable-https-filtering" on page 5812	"file (Packet Forwarding Engine Traffic Traceoptions)" on page 7069
"file-serialization" on page 7094	"flow-sample" on page 7360	"force-fbf-terms" on page 7381
"high-availability (security cloud)" on page 8070	"identity-provider (SAML)" on page 8298	"intelligent-offload-tunable (IDP)" on page 8780
"jnu-management" on page 9429	"large-community-count" on page 9634	"layer2-policer (firewall)" on page 9662

(Continued)

Configuration Statements		
"log-attacks (rulebase-exempt)" on page 10052	"lookup-intact-on-commit" on page 10112	"saml (Access Profile)" on page 10662
"monitor-timer" on page 10953	"no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping)" on page 11628	"overlay (Packet Forwarding Options)" on page 12042
"policy-map-marking" on page 12751	"redundancy-mode" on page 14207	"reset-count" on page 15584
"root-address" on page 15840	"routing-engine-mode (TLB)" on page 15995	"saml" on page 16281
"saml-options" on page 16284	"service-provider (SAML)" on page 16767	"identity-provider settings (SAML)" on page 16895
"Signature Hash Algorithm (Security IKE)" on page 27953	"source-based-fpolicy-statementorwarding" on page 28333	"source-lookup" on page 28338
"source-routing" on page 28339	"static-multihoming-peer" on page 28589	"subscriber-info" on page 28704
"subscriber-pool-utilization-alarm" on page 28705	"transport-class" on page 30132	"unicast-tunnel-name-regular-expression" on page 30355

Operational Commands		
"clear network-access aaa saml assertion-cache" on page 3757	"clear network-access aaa saml idp-metadata" on page 3759	"clear security policies file-serialization counters" on page 4034
"monitor pfe traffic interface" on page 10908	"request csds add-vsrx" on page 14583	"request csds authenticate-host" on page 14584
"request csds delete-vsrx" on page 14586	"request csds extract-vsrx-keys" on page 14587	"request csds jdm" on page 14589

(Continued)

Operational Commands		
"request csds sync-controller" on page 14592	"request jnu satellite sync" on page 14670	"request network-access aaa saml export-sp-metadata" on page 14706
"request network-access aaa saml load-idp-metadata" on page 14707	"request services application-identification offline-download package-path " on page 14936	"show chassis enhanced-temperature-thresholds" on page 17423
"show chassis jnu satellite" on page 18513	"show network-access aaa saml assertion-cache" on page 22560	"show route pfe-install-error" on page 23727
"show route sla-tracking" on page 23764	"show routing nhdb pfe-install-error" on page 23835	"show security policies file-serialization counters" on page 25059
"show security policies lookup-intact-on-commit" on page 25077	"show services application-identification group obsolete-groups" on page 25472	"show services application-identification recent-appid-sigpack-versions" on page 25473
"show services application-identification recent-appid-sigpack-versions" on page 25473		

24.2

Configuration Statements		
"advertise-zero-ethernet-tag" on page 2028	"apts" on page 2436	"auto-sense" on page 2738
"casb" on page 3210	"casb-policy" on page 3212	"channel" on page 3286
"egress-quantization" on page 6238	"fast-lookup-tuple" on page 6935	"fast-lookup-tuple-list" on page 6937

(Continued)

Configuration Statements		
"fast-lookup-tuple-list (policy-options)" on page 6938	"fast-lookup-filter (Protocols BGP)" on page 6947	"feb (chassis)" on page 6982
"fpc (Chassis)" on page 7501	"inband-ztp" on page 8516	"instance (CASB)" on page 8754
"l2pt (Destination Tunnels)" on page 9537	"local-address (Junos Telemetry Interface)" on page 9902	"mcast stats-enable" on page 10582
"monitor-profile" on page 10912	"multipath (OSPF SPF Options)" on page 11109	"ospf-instance" on page 11940
"per-evi" on page 12409	"polling-interval (Network Management)" on page 12752	"prefix-list (Junos Telemetry Interface)" on page 13093
"radius-accounting" on page 13997	"radius (Security Log)" on page 14018	"rdma-opcode" on page 14078
"recovery" on page 14143	"redundancy-group (Redundant Tunnel)" on page 14180	"resource-list (Network Management)" on page 15610
"retry-count (Security Log)" on page 15696	"route-distinguisher-id-use-router-id" on page 15889	"routing-instance (Junos Telemetry Interface)" on page 15937
"rule (casb-policy)" on page 16087	"scale-mode" on page 16328	"scale-mode" on page 16328
"snapshot" on page 28055	"subscriber-extension" on page 28692	"suppress-host-routes-from-dci-to-dc" on page 28728
"traceoptions (CASB)" on page 29503	"traceoptions (flow)" on page 29563	"transient-filter" on page 30062
"vpn-global-import" on page 30953		

Operational Commands		
"clear ptp statistics stream" on page 3836	"request services anti-virus machine-learning-engine" on page 14914	"request vmhost local-file-copy" on page 15503
"request vmhost remote-file-copy" on page 15535	"show evpn gbp-src-tag" on page 19749	"show fib-streaming state" on page 19927
"show l2pt interface" on page 21816	"show programmable-rpd statistics" on page 23292	"show services advanced-anti-malware machine-learning-scan" on page 25311
"show system data-models" on page 26807	"show system packet-forwarding-options hw-resource-monitor monitor-profile" on page 26943	"show system packet-forwarding-options hw-resource-monitor resource-list" on page 26946
"show system packet-forwarding-options hw-resource-monitor utilization-info" on page 26948	"show system packet-forwarding-options hw-resource-utilization-info" on page 26952	"show vmhost network external-network" on page 27813

23.4

Configuration Statements		
"anti-virus (Services)" on page 2275	"arp-nd-probe-disable" on page 2484	"arp-nd-probe-failed-log" on page 2485
"assisted-replication (Deterministic AR Replicator Policy Actions)" on page 2507	"bandwidth (RFC 2544 Benchmarking)" on page 2874	"bfd-liveliness (Activeness Probe)" on page 2917
"bier" on page 2972	"bier-sub-domain (Protocols IS-IS)" on page 2974	"blocklists (Security IKE)" on page 2997

(Continued)

Configuration Statements		
"chassis-id (LLDP)" on page 3301	"damping (BFD Liveness Detection)" on page 4815	"defaults (early-aging)" on page 5226
"destination-identity-context" on page 5427	"destination-identity-context-profile" on page 5429	"deterministic-ar-policy" on page 5553
"disable-packet-register" on page 5813	"drop-flow" on page 5947	"drop-unknown-mac-ip" on page 5984
"dual-phase-bootup" on page 6047	"egress-l2-extended-match" on page 6212	"enhanced-hierarchical-policer" on page 6394
"frequency-lock-threshold" on page 7580	"garp-na-enable" on page 7648	"global-limit" on page 7708
"hw-db-profile" on page 8216	"identity-context-profile" on page 8291	"map-t" on page 10334
"mark-local-ip-host-routes" on page 10355	"mobile-edge session-scale" on page 10792	"monitor-object" on page 10906
"name (early-aging)" on page 11168	"ng-juniper" on page 11366	"node-detection (EVPN-VXLAN)" on page 11633
"peer-limit" on page 12375	"peg-df-election" on page 12395	"performance-monitor" on page 12429
"phase-adjust-threshold" on page 12469	"phase-lock-threshold" on page 12471	"policer-charge" on page 12646
"policer-charge" on page 12646	"policy (Explicit-Proxy)" on page 12678	"preference (PEG DF Election)" on page 13037

(Continued)

Configuration Statements		
"profile (Explicit Proxy)" on page 13277	"protocol (early-aging)" on page 13428	"recycle-bandwidth-profiles" on page 14152
"redundancy (Junos OS)" on page 14183	"redundancy (Junos OS Evolved)" on page 14185	"resource-list (Telemetry)" on page 15612
"server-reject-bridge-domain | server-reject-vlan" on page 16661	"session (Security IKE)" on page 16848	"source-packet-routing (Routing Instances Protocols EVPN)" on page 28273
"static-mac-ip" on page 28587	"test-agent (Paragon Active Assurance)" on page 29187	"trace (EVPN-VXLAN)" on page 29459
"url-feed" on page 30489	"version (IGMP Snooping)" on page 30704	"version (MLD Snooping)" on page 30712
"vxlan-extended" on page 30998	"white-list (Security-Zone)" on page 31109	

Operational Commands		
"clear chassis hardware history" on page 3470	"clear security ike peers" on page 3984	"clear security ike peers statistics" on page 3985
"clear services mobile-edge sessions" on page 4139	"request interface bounce" on page 14639	"request interface optics-reset" on page 14651
"request security utm custom-objects url-feed update" on page 14886	"request security utm custom-objects url-feed update feed-name force" on page 14888	"request security utm web-filtering categorize" on page 14889
"request security utm web-filtering category migrate-to-ng-juniper" on page 14895	"request security utm web-filtering recategorize" on page 14899	"request services anti-virus update" on page 14912

(Continued)

Operational Commands		
"request vmhost zeroize ztp-option secure" on page 15566	"show chassis cluster control-interface status" on page 17341	"show dual-phase-bootup-status" on page 19583
"show evpn oism spmsi-ad" on page 19829	"show ptp passive-port-monitor-status" on page 23383	"show security casb policies" on page 24023
"show security flow session session-state" on page 24291	"show security ike peers" on page 24640	"show security ike peers statistics" on page 24644
"show security utm custom-objects url-feed status" on page 25207	"show services anti-virus statistics" on page 25378	"show system configuration database status" on page 26766
"show vlans vlan-groups" on page 27774	"vlan-groups" on page 30792	

23.2

Configuration Statements		
"allow-fragmentation (fti)" on page 2181	"application-bypass (Juniper Secure Connect)" on page 2314	"application-status" on page 2392
"auto-recovery" on page 2733	"block" on page 2980	"clear-dont-fragment-bit (fti)" on page 3556
"compliance (Juniper Secure Connect)" on page 4483	"dedicated-buffer" on page 5170	"dedicated-buffer-profile" on page 5172
"default-trusted-ca-certs (Security)" on page 5218	"destination-hostname" on page 5426	"detection (Network Isolation Groups)" on page 5540

(Continued)

Configuration Statements		
"device-attribute (Easy EVPN LAG Configuration)" on page 5560	"dhcp-relay (Easy EVPN LAG Configuration)" on page 5638	"downlink-dscp-to-egress-forwarding-class" on page 5922
"encap-type" on page 6280	"enhanced-hash-key profile1/profile2" on page 6389	"evpn (Easy EVPN LAG Configuration)" on page 6548
"evpn-vxlan (Easy EVPN LAG Configuration)" on page 6559	"fast-lookup-filter (PTX)" on page 6930	"fiveqi-map" on page 7199
"global-parameters (Easy EVPN LAG Configuration)" on page 7728	"hash-seed (profile)" on page 8003	"include-igp-metric" on page 8527
"interface-tag (DHCP Local Server)" on page 9081	"interface-tag (DHCP Relay Server)" on page 9083	"ipv6-extended" on page 9348
"ipv6-origin" on page 9367	"ipv6-target" on page 9374	"irb (Easy EVPN LAG Configuration)" on page 9386
"link-tracking (Network Isolation Groups)" on page 9833	"load-balance profile1/profile2" on page 9883	"loopback-firewallv6-optimization" on page 10123
"mask uli" on page 10358	"micro-sid" on page 10697	"mine-dhcp-client-options" on page 10719
"mine-dhcpv6-client-options" on page 10721	"network-isolation" on page 11312	"network-isolation-profile" on page 11314
"no-external-export" on page 11454	"no-next-header-to-payload-protocol-mapping" on page 11522	"no-overload-on-prefix-export-limit" on page 11542
"optimal-route-reflection" on page 11850	"passive-mode-tunneling (Security IPsec)" on page 12208	"password (MAC RADIUS Authentication)" on page 12241

(Continued)

Configuration Statements		
"peer-subnet (Easy EVPN LAG Configuration)" on page 12385	"peer-to-peer (Easy EVPN LAG Configuration)" on page 12388	"resize (Ephemeral Database)" on page 15590
"resize-ephemeral-default" on page 15592	"resource-manager (Forwarding-options)" on page 15616	"server (Easy EVPN LAG Configuration)" on page 16628
"service-mode (mobile-edge SAEGW)" on page 16713	"service-tracking-action (Network Isolation Groups)" on page 16758	"service-tracking (Network Isolation Groups)" on page 16760
"shell" on page 16975	"source-hostname" on page 28213	"source-interface" on page 28225
"subnet-address (Easy EVPN LAG Configuration)" on page 28675	"suppress-interface-leaf" on page 28730	"tcp-seq-ooo-window (Security Forwarding Options)" on page 29019
"tcp-session-install-interval (Security Forwarding Options)" on page 29022	"upstream-rewrite" on page 30481	"vxlan-gbp-l2-profile" on page 30999
"vxlan-gbp-l3-profile" on page 31001	"vxlan-stitching" on page 31009	

Operational Commands		
"request extension-service restart-daemonize-app" on page 14614	"request restart packet-forwarding-engine" on page 14742	"request security pki ca-certificate ca-profile-group default-trusted-ca-certs" on page 14778
"request system restore state path" on page 15165	"request system save state path" on page 15166	"request system zeroize ztp-option secure" on page 15430
"show chassis scb-ha-port" on page 18816	"show class-of-service dedicated-buffer-profile" on page 19007	"show interfaces diagnostics optics-applications" on page 20576
"show ospf spring prefix sid map" on page 22864	"show ospf spring sid database" on page 22866	"show security log report in-detail" on page 24796

(Continued)

Operational Commands		
"show security log report in-interval" on page 24802	"show security log report summary" on page 24804	"show services agf fiveqi remarking" on page 25331
"show-spring-traffic-statistics" on page 26600		

22.4

Configuration Statements		
"evpn-vxlan" on page 6555	"flood-as-unknown-unicast" on page 7239	"grpc-tunnel " on page 7923
"icmp (System)" on page 8269	"loc-rib" on page 9897	"mac-learning-priority" on page 10234
"managed-services" on page 10304	"no-dhcp-flood" on page 11627	"priority (Management Traffic)" on page 13180
"reachability-fault" on page 14083	"rib-list" on page 15783	"sbfd (protocols bfd)" on page 16324
"start (Scripts)" on page 28496	"traffic-engineering (SR-TE import policy)" on page 30032	"tunnel-source-protocol (Protocols OSPF)" on page 30240

Operational Commands		
"clear service l2tp session routing-instance" on page 4062	"ping mpls segment routing static egress-ip" on page 12592	"request-system-integrity-attestation" on page 15112
"request-system-integrity-log-retrieval" on page 15114	"show bfd seamless session" on page 17132	"show chassis high-availability prefix-srgid-table" on page 18495

(Continued)

Operational Commands		
"show lldp neighbors-vlan-name-tlv-list interface" on page 21973	"show service l2tp session routing-instance" on page 25246	"show system statistics netisr per-priority-stats" on page 27173
"show system statistics netisr qdrops" on page 27174	"show system statistics netisr socket-drops" on page 27176	"show ted spring-te-policy" on page 27386
"traceroute mpls segment-routing static egress-ip" on page 29970		

2

CHAPTER

Configuration Statements by Hierarchy

[edit] | 17

[edit access] | 19

[edit accounting-options] | 47

[edit aggregated-devices] | 51

[edit applications] | 52

[edit bridge-domains] | 54

[edit bridge-options] | 63

[edit chassis] | 63

[edit class-of-service] | 103

[edit diameter] | 131

[edit dynamic] | 133

[edit dynamic-profiles] | 134

[edit ethernet-switching] | 173

[edit ethernet-switching-options] | 174

[edit event-options] | 181

[edit fabric] | 186

[edit family] | 191

[edit fc-fabrics] | 193

[edit fc-options] | 195

[edit firewall] | 196

[edit forwarding-options] | 205

[edit fpc] | 259

[edit hierarchical-policer] | 260

[edit interface] | 262

[edit interfaces] | 263

[edit jsrc] | 353

[edit logical-routers] | 354

[edit logical-systems] | 354

[edit multi-chassis] | 645

[edit multicast-snooping-options] | 646

[edit node-group] | 647

[edit poe] | 648

[edit policy-options] | 651

[edit protocols] | 661

[edit pvlan] | 798

[edit redundant-power-system] | 798

[edit routing-instances] | 799

[edit routing-instance-rib-inet6.0] | 929

[edit routing-options] | 929

[edit routing-options-rib-inet6.0] | 952

[edit schedulers] | 953

[edit security] | 955

[edit service-set] | 1045

[edit services] | 1046

[edit smtp] | 1131

[edit snmp] | 1132

[edit switch-options] | 1144

[edit system] | 1147

[edit tenants] | 1204

[edit unified-edge] | 1231

[edit virtual-chassis] | 1245

[edit virtual-network-functions] | 1246

[edit vlans] | 1247

[edit vmhost] | 1255

[edit wlan] | 1256

[edit]

[edit]		
"accounting-options" on page 1717	"applications (Applications)" on page 2405	"applications (Services ALGs)" on page 2415
"auto-conversion" on page 2704	"auto-provisioned" on page 2725	"bridge-domains" on page 3032
"chassis" on page 3296	"chassis (QFabric System)" on page 3303	"cross-connect" on page 4724
"diameter" on page 5693	"dynamic-profiles" on page 6134	"ethernet-switching-options" on page 6506
"event-options" on page 6527	"fabric" on page 6745	"fc-fabrics" on page 6954
"fc-options" on page 6959	"firewall" on page 7176	"forwarding-options (Bridge Domains or VLANs)" on page 7467
"forwarding-options (Routing Instance)" on page 7463	"groups" on page 7916	"host-os forwarding-options analyzer" on page 8155
"interface-range" on page 9006	"interfaces (Aggregated Multiservices)" on page 9115	"interfaces (EX Series switches)" on page 9135
"interfaces (QFX Series, ACX Series)" on page 9156	"interfaces (Voice Services)" on page 9181	"jsrc (JSRC)" on page 9439
"jsrc-partition" on page 9441	"license-server" on page 9740	"logical-systems" on page 10090
"logical-systems (All)" on page 10092	"mapping-range" on page 10344	"multi-chassis" on page 11042
"nextboot" on page 11361	"no-auto-conversion" on page 11404	"oism (Multicast Snooping Options)" on page 11807

(Continued)

[edit]		
"poe" on page 12616	"policy-options" on page 12697	"policy-options (QFX)" on page 12700
"primary-vlan" on page 13150	"protocols (STP Type)" on page 13802	"pvlan" on page 13882
"redundant-power-system" on page 14216	"routing-instances" on page 15969	"routing-instances (Logical Systems)" on page 15966
"routing-instances (Multiple Routing Entities)" on page 15976	"routing-options" on page 15980	"schedulers (Security Policies)" on page 16375
"security" on page 16448	"services (Captive Portal Content Delivery)" on page 16769	"services (CoS)" on page 16772
"services (Hierarchy)" on page 16773	"services (IDS)" on page 16774	"services (IPsec VPN)" on page 16777
"services (L2TP)" on page 16778	"services (NAT)" on page 16783	"services (Stateful Firewall)" on page 16799
"snmp" on page 28058	"switch-options" on page 28749	"unknown-unicast-forwarding" on page 30417
"virtual-chassis" on page 30745	"virtual-chassis (MX Series Virtual Chassis)" on page 30747	"vlans" on page 30880
"vmhost" on page 30895	"vmhost resize vjunos compact" on page 30904	

[edit access]

IN THIS SECTION

- [aaa-options] | 22
- [accounting-backup-options] | 22
- [address] | 23
- [address-assignment] | 23
- [diameter] | 26
- [domain] | 27
- [firewall-authentication] | 28
- [group-profile] | 29
- [gx-plus] | 30
- [ldap-options] | 31
- [ldap-server] | 32
- [nasreq] | 32
- [ocs] | 32
- [pcrf] | 34
- [profile] | 35
- [profileldap-options] | 41
- [protocol-attributes] | 41
- [radius] | 42
- [radius-disconnect] | 43
- [radius-options] | 44
- [radius-server] | 44
- [radsec] | 44
- [s6a] | 45
- [saml] | 45
- [securid-server] | 45
- [tunnel-profile] | 46
- [tunnel-switch-profile] | 47

Sub Hierarchies

"[edit access aaa-options {aaa-options-name}]" on page 22

"[edit access accounting-backup-options]" on page 22

"[edit access address]" on page 23

"[edit access address-assignment]" on page 23

"[edit access diameter]" on page 26

"[edit access domain]" on page 27

"[edit access firewall-authentication]" on page 28

"[edit access group-profile {profile-name}]" on page 29

"[edit access gx-plus]" on page 30

"[edit access ldap-options]" on page 31

"[edit access ldap-server {server-address}]" on page 32

"[edit access nasreq]" on page 32

"[edit access ocs]" on page 32

"[edit access pcrf]" on page 34

"[edit access profile {profile-name}]" on page 35

"[edit access profile{profile-name} ldap-options]" on page 41

"[edit access protocol-attributes {attribute-set-name}]" on page 41

"[edit access radius]" on page 42

"[edit access radius-disconnect {client-address}]" on page 43

"[edit access radius-options]" on page 44

"[edit access radius-server {server-address}]" on page 44

"[edit access radsec]" on page 44

"[edit access s6a]" on page 45

"[edit access saml]" on page 45

"[edit access securid-server]" on page 45

"[edit access tunnel-profile {profile-name}]" on page 46

"[edit access tunnel-switch-profile {profile-name}]" on page 47

[edit access]		
"aaa-options (Access Profile)" on page 1613	"accounting-backup-options (Access Profile)" on page 1713	"address-assignment (Access)" on page 1858
"address-assignment (Address-Assignment Pools)" on page 1862	"address-pool" on page 1899	"address-pool (Access)" on page 1901
"address-protection" on page 1910	"diameter (Subscriber Aware Policy Control)" on page 5699	"domain (Domain Map)" on page 5877
"domain-name-server (Routing Instances and Access Profiles)" on page 5891	"domain-name-server-inet (Routing Instances and Access Profiles)" on page 5888	"domain-name-server-inet6 (Routing Instances and Access Profiles)" on page 5890
"firewall-authentication" on page 7179	"group-profile (Access)" on page 7859	"group-profile (Group Profile)" on page 7861
"gx-plus (Gx-Plus)" on page 7952	"ldap-options" on page 9673	"ldap-server" on page 9676
"linked-pool-aggregation (Address-Assignment Pools)" on page 9837	"nasreq (Diameter Application)" on page 11202	"ocs (Diameter Applications)" on page 11771
"pcrf (Diameter Applications)" on page 12325	"profile" on page 13256	"profile" on page 13257
"profile (Access)" on page 13259	"profile (Access)" on page 13263	"radius-disconnect" on page 14001
"radius-disconnect-port" on page 14005	"radius-options (Access)" on page 14010	"radius-server" on page 14020
"radius-server (Access)" on page 14026	"radius-server-name" on page 14028	"radsec " on page 14036

(Continued)

[edit access]		
"report-interface-descriptions (Access)" on page 14392	"s6a" on page 16260	"saml" on page 16281
"saml-options" on page 16284	"securid-server" on page 16447	"terminate-code" on page 29147
"tunnel-profile (Tunnel Profile)" on page 30228	"tunnel-switch-profile (L2TP Tunnel Switching, Definition)" on page 30243	"wins-server (Access)" on page 31129

[aaa-options]

[edit access aaa-options {aaa-options-name}]		
"aaa-context (AAA Options)" on page 1610	"access-profile (AAA Options)" on page 1690	"subscriber-context (AAA Options)" on page 28687

[accounting-backup-options]

[edit access accounting-backup-options]		
"max-pending-accounting-stops (Access Profile)" on page 10454	"max-withhold-time (Access Profile)" on page 10494	

[address]

[edit access address assignment-address pool]	
"excluded-address (Address-Assignment Pools)" on page 6631	

[address-assignment]

[edit access address-assignment]		
"address-pools" on page 1908	"pool (Access)" on page 12763	"pool (Address-Assignment Pools)" on page 12766

[edit access address-assignment (Access)]	
"dynamic-pool" on page 6114	

[edit access address-assignment address-pools {name}]		
"default-pool (Address Pools)" on page 5205	"family (Address Pools)" on page 6843	

[edit access address-assignment address-pools {name} family inet]	
"network (Address Pools)" on page 11301	

[edit access address-assignment address-pools {name} family inet network {network-prefix}]	
"external-assigned (Address Pools)" on page 6725	

[edit access address-assignment address-pools {name} family inet6]		
"network (Address Pools)" on page 11301		

[edit access address-assignment address-pools {name} family inet6 network {network-prefix}]		
"external-assigned (Address Pools)" on page 6725		

[edit access address-assignment pool {pool-name}]		
"family (Address-Assignment Pools)" on page 6841	"link (Access)" on page 9781	

[edit access address-assignment pool <name> family (inet | inet6)]		
"network (Access)" on page 11300		

[edit access address-assignment pool <name> family (inet | inet6) dhcp-attributes]		
"name-server (Access)" on page 11178		

[edit access address-assignment pool <name> family (inet | inet6) xauth-attributes]		
"wins-server (Access)" on page 31129		

[edit access address-assignment pool <name> family inet]		
"xauth-attributes" on page 31150		

[edit access address-assignment pool <name> family inet6]		
"xauth-attributes" on page 31150		

[edit access address-assignment pool {pool-name} family {family}]		
"dhcp-attributes (Address-Assignment Pools)" on page 5589		

[edit access address-assignment pool {name} family (inet | inet6)]		
"excluded-range (Address-Assignment Pools)" on page 6632	"host (Address-Assignment Pools)" on page 8136	"range (Address-Assignment Pools)" on page 14048

[edit access address-assignment pool {pool-name} family inet]		
"dhcp-attributes (Access IPv4 Address Pools)" on page 5585	"network" on page 11298	

[edit access address-assignment pool {pool-name} family inet dhcp-attributes]		
"option-match" on page 11896		

[edit access address-assignment pool {pool-name} family inet dhcp-attributes option-match]		
"option-82 (Address-Assignment Pools)" on page 11888		

[edit access address-assignment pool {pool-name} family inet6]		
"dhcp-attributes (Access IPv6 Address Pools)" on page 5587	"prefix (Address-Assignment Pools)" on page 13071	

[edit access address-assignment pool pool-name family inet]		
"range (Access)" on page 14045		

[edit access address-assignment pool pool-name family inet6]		
"prefix (Access IPv6)" on page 13067	"range (Access)" on page 14045	

[diameter]

[edit access diameter]		
"applications (Diameter)" on page 2408	"firmware-revision" on page 7197	"network-element (Diameter Base Protocol)" on page 11305
"origin (Diameter Base Protocol)" on page 11934	"peer (Diameter Base Protocol)" on page 12351	"product-name" on page 13252
"traceoptions (Diameter Base Protocol for Subscriber Aware Policy Control)" on page 29524		

[edit access diameter applications pcc-gx]	
"maximum-pending-requests (Diameter)" on page 10546	

[edit access diameter network-element {element-name}]		
"function (Diameter Network Element for Subscriber Management)" on page 7633	"peer (Diameter Network Element)" on page 12354	

[edit access diameter network-element {element-name} peer {peer-name}]		
"priority (Diameter Network Element)" on page 13166	"timeout (Diameter Network Element)" on page 29376	

[edit access diameter origin]		
"host (Diameter Origin)" on page 8140	"realm (Diameter Origin)" on page 14096	

[edit access diameter peer {peer-name}]		
"address (Diameter Peer)" on page 1874	"connect-actively" on page 4529	"disconnect-peer-timeout" on page 5820
"incoming-queue" on page 8541	"outgoing-queue" on page 11976	"watchdog-timeout" on page 31031

[domain]

[edit access domain]		
"delimiter (Domain Map)" on page 5293	"map (Domain Map)" on page 10327	"parse-direction (Domain Map)" on page 12178
"parse-order (Domain Map)" on page 12179	"realm-delimiter (Domain Map)" on page 14095	"realm-parse-direction (Domain Map)" on page 14098

[edit access domain map {domain-map-name}]		
"aaa-logical-system (Domain Map)" on page 1611	"aaa-routing-instance (Domain Map)" on page 1619	"access-profile (Domain Map)" on page 1691
"address-pool (Domain Map)" on page 1902	"override-chap-password" on page 12055	"override-password (Domain Map)" on page 12060
"strip-domain (Domain Map)" on page 28657	"strip-username (Domain Map)" on page 28661	"sub-domain" on page 28668

(Continued)

[edit access domain map {domain-map-name}]		
"target-logical-system (Domain Map)" on page 28949	"target-routing-instance (Domain Map)" on page 28952	"tunnel-profile (Domain Map)" on page 30225
"tunnel-switch-profile (L2TP Tunnel Switching, Application)" on page 30242	"using-user-password" on page 30609	

[edit access domain map {domain-map-name} aaa-logical-system {logical-system-name}]	
"aaa-routing-instance (Domain Map)" on page 1619	

[edit access domain map {domain-map-name} target-logical-system {logical-system-name}]	
"target-routing-instance (Domain Map)" on page 28952	

[firewall-authentication]

[edit access firewall-authentication]		
"pass-through (Firewall-authentication)" on page 12196	"traceoptions (Access)" on page 29477	"web-authentication (Access)" on page 31058

[edit access firewall-authentication pass-through]		
"default-profile" on page 5208	"ftp (Access)" on page 7619	"http (Access)" on page 8192
"telnet (Access)" on page 29059		

[edit access firewall-authentication pass-through (ftp | http | telnet)]		
"banner (Access FTP HTTP Telnet Authentication)" on page 2882		

[edit access firewall-authentication pass-through default-profile {profile-name} (ftp | http | telnet) banner]		
"fail" on page 6775	"login (Access)" on page 10100	"success" on page 28720

[edit access firewall-authentication web-authentication]		
"banner (Access Web Authentication)" on page 2884	"success" on page 28720	

[group-profile]

[edit access group-profile {profile-name}]		
"l2tp (Group Profile)" on page 9548	"ppp (Group Profile)" on page 12945	

[edit access group-profile {profile-name} client {client-name} ppp]		
"primary-dns" on page 13137	"primary-wins" on page 13153	

[edit access group-profile {profile-name} l2tp]		
"interface-id" on page 8898	"lcp-renegotiation" on page 9671	"local-chap" on page 9933
"maximum-sessions-per-tunnel" on page 10563		

[edit access group-profile {profile-name} ppp]		
"cell-overhead" on page 3248	"encapsulation-overhead" on page 6305	"framed-pool" on page 7567
"interface-id" on page 8898	"keepalive" on page 9466	"ppp-options (L2TP)" on page 12954
"secondary-dns" on page 16412	"secondary-wins" on page 16420	

[edit access group-profile {profile-name} ppp ppp-options]		
"aaa-options (PPP Profile)" on page 1614	"chap (L2TP)" on page 3292	"ignore-magic-number-mismatch (Access Group Profile)" on page 8403
"initiate-ncp (Dynamic and Static PPP)" on page 8662	"ipcp-suggest-dns-option" on page 9289	"mru (Dynamic and Static PPPoE)" on page 10994
"mtu (Dynamic and Static PPPoE)" on page 11034	"pap (L2TP)" on page 12172	"peer-ip-address-optional" on page 12373

[edit access group-profile {profile-name}ppp ppp]	
"idle-timeout (Access)" on page 8306	

[gx-plus]

[edit access gx-plus]		
"global (Gx-Plus)" on page 7706	"partition (Gx-Plus)" on page 12182	

[edit access gx-plus global]		
"include-ipv6 (Gx-Plus)" on page 8528	"max-outstanding-requests (Diameter Applications)" on page 10449	

[edit access gx-plus partition {partition-name}]		
"destination-host (Gx-Plus)" on page 5422	"destination-realm (Gx-Plus)" on page 5512	"diameter-instance (Diameter Applications)" on page 5695

[ldap-options]

[edit access ldap-options]		
"assemble" on page 2499	"base-distinguished-name" on page 2885	"revert-interval (Access LDAP)" on page 15725
"search" on page 16402		

[edit access ldap-options search]		
"admin-search" on page 1978	"search-filter" on page 16404	

[edit access ldap-options search admin-search]		
"distinguished-name (Access)" on page 5830	"password (Access)" on page 12225	

[ldap-server]

[edit access ldap-server {server-address}]		
"port (Access LDAP)" on page 12812	"retry (Access LDAP)" on page 15691	"routing-instance (Access LDAP)" on page 15925
"timeout (Access LDAP)" on page 29368		

[edit access ldap-server server-address]	
"source-address (Access LDAP)" on page 28124	

[nasreq]

[edit access nasreq]	
"partition (NASREQ Diameter Application)" on page 12183	

[ocs]

[edit access ocs]		
"global (OCS)" on page 7725	"partition (OCS)" on page 12185	

[edit access ocs global]		
"service-context-id (OCS)" on page 16684		

[edit access ocs partition {partition-name}]		
"backup (OCS Partition)" on page 2804	"called-station-id (OCS Partition)" on page 3186	"charging-id (OCS Partition)" on page 3293
"destination-host (OCS Partition)" on page 5423	"destination-realm (OCS Partition)" on page 5515	"diameter-instance (Diameter Applications)" on page 5695
"draining (Diameter Applications)" on page 5941	"draining-response-timeout (Diameter Applications)" on page 5943	"final-response-timeout (OCS Partition)" on page 7170
"force-continue (OCS Partition)" on page 7365	"ggsn-address (OCS Partition)" on page 7682	"ggsn-mcc-mnc (OCS Partition)" on page 7684
"max-outstanding-requests (Diameter Applications)" on page 10449	"send-origin-state-id (Diameter Applications)" on page 16545	"sftp-backup (OCS Partition)" on page 16916
"user-name-include (OCS Partition)" on page 30571		

[edit access ocs partition alternative-partition-name {alternative-partition-name}]		
"alternative-partition-name (OCS Partition)" on page 2218		

[pcrf]

[edit access pcrf]		
"global (PCRF)" on page 7731	"partition (PCRF)" on page 12187	

[edit access pcrf partition {partition-name}]		
"accept-sdr (PCRF Partition)" on page 1644	"destination-host (PCRF Partition)" on page 5424	"destination-realm (PCRF Partition)" on page 5517
"diameter-instance (Diameter Applications)" on page 5695	"draining (Diameter Applications)" on page 5941	"draining-response-timeout (Diameter Applications)" on page 5943
"ip-can-type (PCRF Partition)" on page 9239	"local-decision (PCRF Partition)" on page 9934	"logout-response-timeout (PCRF Partition)" on page 10101
"max-outstanding-requests (Diameter Applications)" on page 10449	"report-local-rule (PCRF Partition)" on page 14393	"report-resource-allocation (PCRF Partition)" on page 14396
"report-successful-resource-allocation (PCRF Partition)" on page 14404	"send-network-family-indicator (PCRF Partition)" on page 16542	"send-origin-state-id (Diameter Applications)" on page 16545
"subscription-id-data-include (PCRF Partition)" on page 28713	"subscription-id-type (PCRF Partition)" on page 28719	"update-response-timeout (PCRF Partition)" on page 30447

[edit access pcrf partition {partition-name}],]	
"send-dyn-subscription-indicator (PCRF Partition)" on page 16540	

[profile]

[edit access profile {profile-name}]		
"accounting (Access Profile)" on page 1711	"accounting-order" on page 1721	"authentication-order" on page 2633
"authentication-order (Access Profile)" on page 2635	"authorization-order" on page 2680	"charging-service-list" on page 3295
"client" on page 4292	"client-group" on page 4312	"client-name-filter" on page 4332
"domain-name-server (Routing Instances and Access Profiles)" on page 5891	"domain-name-server-inet (Routing Instances and Access Profiles)" on page 5888	"domain-name-server-inet6 (Routing Instances and Access Profiles)" on page 5890
"jsrc (Access Profile)" on page 9438	"ldap-server" on page 9676	"local (Flat-File Access Profile)" on page 9942
"preauthentication-order (Access Profile)" on page 13007	"provisioning-order (Diameter Applications)" on page 13815	"radius (Access Profile)" on page 13993
"radius-server" on page 14020	"radius-server (Access)" on page 14026	"radius-server-name" on page 14028
"saml (Access Profile)" on page 10662	"service (Service Accounting)" on page 16739	"session-limit-per-username (Access Profile)" on page 16821
"session-options" on page 16831	"session-options (Access Profile)" on page 16835	"subscriber (Access Profile)" on page 28680

[edit access profile {profile-name} accounting]		
"accounting-stop-on-access-deny" on page 1743	"accounting-stop-on-failure" on page 1744	"address-change-immediate-update" on page 1872

(Continued)

[edit access profile {profile-name} accounting]		
"ancp-speed-change-immediate-update (ANCP)" on page 2251	"coa-immediate-update" on page 4382	"coa-no-override service-class-attribute" on page 4384
"duplication (Access Profile)" on page 6074	"duplication-filter (Access Profile)" on page 6075	"duplication-vrf (Duplicate Accounting)" on page 6077
"family-state-change-immediate-update" on page 6916	"immediate-update" on page 8459	"order" on page 11926
"send-acct-status-on-config-change (Access Profile)" on page 16534	"statistics (Access Profile)" on page 28592	"update-interval" on page 30439
"wait-for-acct-on-ack (Access Profile)" on page 31014		

[edit access profile {profile-name} accounting duplication-vrf]		
"access-profile-name (Duplicate Accounting)" on page 1697	"vrf-name (Duplicate Accounting)" on page 30963	

[edit access profile {profile-name} client {client-name}]		
"client-group" on page 4312	"firewall-user" on page 7196	"ike (Access Profile)" on page 8419
"l2tp (Profile)" on page 9553	"maximum-sessions (L2TP)" on page 10560	"ppp (Profile)" on page 12957

[edit access profile {profile-name} client {client-name} ike]	
"interface-id" on page 8898	

[edit access profile {profile-name} client {client-name} l2tp]		
"aaa-access-profile (L2TP LNS)" on page 1607	"interface-id" on page 8898	"lcp-renegotiation" on page 9671
"local-chap" on page 9933	"maximum-sessions-per-tunnel" on page 10563	"multilink" on page 11103
"service-profile (L2TP)" on page 16733		

[edit access profile {profile-name} client {client-name} ppp]		
"cell-overhead" on page 3248	"encapsulation-overhead" on page 6305	"framed-pool" on page 7567
"interface-id" on page 8898	"keepalive" on page 9466	"keepalive-retries" on page 9469
"secondary-dns" on page 16412	"secondary-wins" on page 16420	

[edit access profile {name} client {name} xauth]	
"ip-address (Access Profile)" on page 9234	

[edit access profile {profile-name} client {client-name}ppp]	
"idle-timeout (Access)" on page 8306	

[edit access profile {profile-name} client-name-filter {client-name}]		
"count" on page 4685	"domain-name (Access Profile)" on page 5881	"separator" on page 16593

[edit access profile {profile-name} jsrc]		
"attributes (JSRC Attributes)" on page 2557		

[edit access profile {profile-name} jsrc attributes]		
"exclude (JSRC Attributes)" on page 6609		

[edit access profile {profile-name} ldap-options]		
"assemble" on page 2499	"revert-interval (Access LDAP)" on page 15725	"search" on page 16402

[edit access profile {profile-name} ldap-options search]		
"admin-search" on page 1978	"search-filter" on page 16404	

[edit access profile {profile-name} ldap-options search admin-search]		
"distinguished-name (Access)" on page 5830	"password (Access)" on page 12225	

[edit access profile {profile-name} ldap-server {ip-address}]		
"no-tls-certificate-check" on page 11599	"port (Access LDAP)" on page 12812	"retry (Access LDAP)" on page 15691
"routing-instance (Access LDAP)" on page 15925	"source-address (Access LDAP)" on page 28124	"timeout (Access LDAP)" on page 29368
"tls-min-version" on page 29402	"tls-peer-name" on page 29405	"tls-timeout" on page 29406
"tls-type" on page 29407		

[edit access profile {profile-name} ppp]		
"primary-dns" on page 13137	"primary-wins" on page 13153	

[edit access profile profile-name]	
"radius-options (Access)" on page 14010	

[edit access profile profile-name authentication-order ldap]	
"ldap-options" on page 9673	

[edit access profile {profile-name} radius]		
"accounting-server" on page 1739	"accounting-server-name" on page 1740	"attributes (RADIUS Attributes)" on page 2560
"authentication-server" on page 2652	"options (Access Profile)" on page 11912	

[edit access profile {profile-name} radius attributes]	
"exclude (RADIUS Attributes)" on page 6617	

[edit access profile {profile-name} radius options]		
"calling-station-id-format (Subscriber Management)" on page 3193	"interface-description-format" on page 8858	"nas-port-extended-format" on page 11184
"nas-port-id-format (Subscriber Management)" on page 11188	"nas-port-type" on page 11195	"nas-port-type (Subscriber Management)" on page 11199

(Continued)

[edit access profile {profile-name} radius options]		
"override (RADIUS Options)" on page 12063	"revert-interval (Access)" on page 15724	

[edit access profile {profile-name} radius-server {server-address}]		
"radsec-destination" on page 14040	"retry (Access RADIUS)" on page 15693	"routing-instance (Access RADIUS)" on page 15927
"secret (Access Profile)" on page 16422	"secret (RADIUS)" on page 16424	"source-address (Access RADIUS)" on page 28125
"timeout (Access RADIUS)" on page 29370		

[edit access profile {profile-name} service]		
"accounting (Service Accounting)" on page 1742	"accounting-order (Service Accounting)" on page 1722	

[edit access profile {profile-name} service accounting]		
"statistics (Service Accounting)" on page 28598	"update-interval (Service Accounting)" on page 30444	

[edit access profile {profile-name} session-options]		
"client-idle-timeout (Access Profile)" on page 4326	"client-session-timeout (Access Profile)" on page 4338	"pcc-context" on page 12289
"strip-user-name (Access Profile)" on page 28659		

[edit access profile {profile-name} session-options strip-user-name]		
"delimiter (Access Profile)" on page 5286	"parse-direction (Access Profile)" on page 12176	

[profileldap-options]

[edit access profile{profile-name} ldap-options]	
"base-distinguished-name" on page 2885	

[protocol-attributes]

[edit access protocol-attributes {attribute-set-name}]	
"option-match" on page 11896	

[edit access protocol-attributes {attribute-set-name} option-match]	
"option-82 (Address-Assignment Pools)" on page 11888	

[radius]

[edit access radius]		
"clients" on page 4345	"network-elements (RADIUS)" on page 11310	"servers (RADIUS)" on page 16678
"snoop-segments (RADIUS)" on page 28078		

[edit access radius clients {client-name}]		
"accounting (RADIUS Client)" on page 1736	"address (RADIUS Clients)" on page 1915	"prefer-framed-ip-address (RADIUS Clients)" on page 13030
"prefer-framed-ipv6-prefix (RADIUS Clients)" on page 13031	"source-interface (RADIUS Client)" on page 28224	

[edit access radius clients {client-name} accounting]		
"response-cache-timeout (RADIUS Client)" on page 15629	"secret (RADIUS Client)" on page 16426	

[edit access radius network-element {name}]		
"maximum-pending-reqs-limit" on page 10544	"pending-queue-watermark" on page 12402	"pending-queue-watermark-abate" on page 12403
"server (RADIUS Network Elements)" on page 16660		

[edit access radius network-element {name} server {name}]	
"priority (RADIUS Network Elements)" on page 13203	

[edit access radius servers {name}]		
"accounting-port (RADIUS Server)" on page 1724	"accounting-secret (RADIUS Server)" on page 1737	"address (RADIUS Server)" on page 1916
"allow-dynamic-requests (RADIUS Server)" on page 2173	"dead-criteria-retries (RADIUS Server)" on page 5146	"dynamic-requests-secret (RADIUS Server)" on page 6149
"port (RADIUS Server)" on page 12859	"retry (RADIUS Server)" on page 15714	"revert-interval (RADIUS Server)" on page 15728
"secret (RADIUS Server)" on page 16427	"source-interface (RADIUS Server)" on page 28227	"timeout (RADIUS Server)" on page 29382

[edit access radius snoop-segments {segment-name}]		
"destination-ip-address (RADIUS Snoop Segment)" on page 5434	"destination-port (RADIUS Snoop Segment)" on page 5474	"request-cache-timeout (RADIUS Snoop Segment)" on page 14432
"shared-secret (RADIUS Snoop Segment)" on page 16973	"source-interface (RADIUS Snoop Segment)" on page 28228	"source-ip-address (RADIUS Snoop Segment)" on page 28237

[radius-disconnect]

[edit access radius-disconnect {client-address}]	
"secret (RADIUS)" on page 16424	

[radius-options]

[edit access radius-options]		
"interim-rate (Access)" on page 9183	"request-rate (Access)" on page 14738	"revert-interval (Access RADIUS)" on page 15727
"revert-interval (Access)" on page 15724	"timeout-grace (Access)" on page 29379	"unique-nas-port (Access)" on page 30364

[radius-server]

[edit access radius-server {server-address}]		
"radsec-destination" on page 14040	"retry (Access RADIUS)" on page 15693	"routing-instance (Access RADIUS)" on page 15927
"secret (RADIUS)" on page 16424	"source-address (Access RADIUS)" on page 28125	"timeout (Access RADIUS)" on page 29370

[radsec]

[edit access radsec destination]	
"dynamic-requests" on page 6148	

[s6a]

[edit access s6a]	
"partition (s6a)" on page 12190	

[saml]

[edit access saml]		
"identity-provider (SAML)" on page 8298	"service-provider (SAML)" on page 16767	

[edit access saml identity-provider {identity-provider}]		
"identity-provider attribute-mapping (SAML)" on page 2501	"identity-provider settings (SAML)" on page 16895	

[securid-server]

[edit access securid-server]	
"configuration-file" on page 4514	

[tunnel-profile]

[edit access tunnel-profile {profile-name}]		
"tunnel (Tunnel Profile)" on page 30248		

[edit access tunnel-profile {profile-name} tunnel {tunnel-id}]		
"identification (Tunnel Profile)" on page 8287	"logical-system (Tunnel Profile)" on page 10089	"max-sessions (Tunnel Profile)" on page 10480
"medium (Tunnel Profile)" on page 10603	"nas-port-method (Tunnel Profile)" on page 11192	"preference (Tunnel Profile)" on page 13052
"remote-gateway (Tunnel Profile)" on page 14323	"routing-instance (Tunnel Profile)" on page 15965	"secret (Tunnel Profile)" on page 16429
"source-gateway (Tunnel Profile)" on page 28211	"type (Tunnel Profile)" on page 30294	

[edit access tunnel-profile {profile-name} tunnel {tunnel-id} remote-gateway]		
"address (Tunnel Profile Remote Gateway)" on page 1935	"gateway-name (Tunnel Profile Remote Gateway)" on page 7655	

[edit access tunnel-profile {profile-name} tunnel {tunnel-id} source-gateway]		
"address (Tunnel Profile Source Gateway)" on page 1936	"gateway-name (Tunnel Profile Source Gateway)" on page 7657	

[tunnel-switch-profile]

[edit access tunnel-switch-profile {profile-name}]		
"avp (L2TP Tunnel Switching)" on page 2784	"tunnel-profile (L2TP Tunnel Switching)" on page 30226	

[edit access tunnel-switch-profile {profile-name} avp]		
"bearer-type (L2TP Tunnel Switching)" on page 2892	"calling-number (L2TP Tunnel Switching)" on page 3190	"cisco-nas-port-info (L2TP Tunnel Switching)" on page 3335

[edit accounting-options]

IN THIS SECTION

- [class-usage-profile] | 48
- [file] | 48
- [filter-profile] | 49
- [flat-file-profile] | 49
- [interface-profile] | 50
- [mib-profile] | 50
- [routing-engine-profile] | 50

Sub Hierarchies

"[edit accounting-options class-usage-profile {profile-name}]" on page 48

"[edit accounting-options file {filename}]" on page 48

"[edit accounting-options filter-profile {profile-name}]" on page 49

"[edit accounting-options flat-file-profile {profile-name}]" on page 49

"[edit accounting-options interface-profile {profile-name}]" on page 50

"[edit accounting-options mib-profile {profile-name}]" on page 50

"[edit accounting-options routing-engine-profile {profile-name}]" on page 50

[edit accounting-options]		
"class-usage-profile" on page 3373	"file (Configuring a Log File)" on page 7047	"filter-profile" on page 7151
"interface-profile" on page 8959	"mib-profile" on page 10695	"routing-engine-profile" on page 15909

[class-usage-profile]

[edit accounting-options class-usage-profile {profile-name}]		
"destination-classes" on page 5412	"file (Associating with a Profile)" on page 7036	"interval (Accounting Options)" on page 9207
"source-classes" on page 28193		

[file]

[edit accounting-options file {filename}]		
"archive-sites" on page 2446	"files (accounting)" on page 7104	"nonpersistent" on page 11668

(Continued)

[edit accounting-options file {filename}]		
"size (Interface)" on page 28000	"start-time (Accounting)" on page 28502	"transfer-interval" on page 30059

[filter-profile]

[edit accounting-options filter-profile {profile-name}]		
"counters" on page 4691	"file (Associating with a Profile)" on page 7036	"interval (Accounting Options)" on page 9207

[flat-file-profile]

[edit accounting-options flat-file-profile {profile-name}]		
"fields (Flat-File Accounting Options)" on page 7016	"file (Flat-File Accounting Options)" on page 7058	"format (Flat-File Accounting Options)" on page 7382
"interval (Flat-File Accounting Options)" on page 9211	"schema-version (Flat-File Accounting Options)" on page 16378	

[edit accounting-options flat-file-profile {profile-name} fields]		
"egress-stats (Flat-File Accounting Options)" on page 6225	"general-param (Flat-File Accounting Options)" on page 7670	"ingress-stats (Flat-File Accounting Options)" on page 8651
"l2-stats (Flat-File Accounting Options)" on page 9530	"overall-packet (Flat-File Accounting Options)" on page 12032	

[interface-profile]

[edit accounting-options interface-profile {profile-name}]		
"allow-clear (Accounting Options)" on page 2167	"fields (for Interface Profiles)" on page 7019	"file (Associating with a Profile)" on page 7036
"interval (Accounting Options)" on page 9207		

[mib-profile]

[edit accounting-options mib-profile {profile-name}]		
"file (Associating with a Profile)" on page 7036	"interval (Accounting Options)" on page 9207	"object-names" on page 11767
"operation" on page 11846		

[routing-engine-profile]

[edit accounting-options routing-engine-profile {profile-name}]		
"fields (for Routing Engine Profiles)" on page 7021	"file (Associating with a Profile)" on page 7036	"interval (Accounting Options)" on page 9207

[edit aggregated-devices]

IN THIS SECTION

- [ethernet] | 51

Sub Hierarchies

"[edit aggregated-devices ethernet]" on page 51

[edit aggregated-devices]		
"ethernet" on page 6480	"ethernet (Aggregated Devices)" on page 6481	

[ethernet]

[edit aggregated-devices ethernet]	
"device-count" on page 5565	

[edit aggregated-devices ethernet (Aggregated Devices)]	
"device-count" on page 5565	

[edit applications]

IN THIS SECTION

- [application] | 52
- [application-set] | 54

Sub Hierarchies

"[edit applications application {application-name}]" on page 52

"[edit applications application-set {application-set-name}]" on page 54

[edit applications]		
"application" on page 2299	"application (Applications)" on page 2310	"application-set" on page 2382

[application]

[edit applications application {application-name}]		
"application-protocol" on page 2352	"application-protocol (Applications)" on page 2355	"description (Applications)" on page 5326
"destination-port (Applications)" on page 5465	"destination-port (applications)" on page 5464	"ether-type" on page 6476

(Continued)

[edit applications application {application-name}]		
"icmp-code" on page 8255	"icmp-code (Applications)" on page 8256	"icmp-type" on page 8270
"icmp-type (Applications)" on page 8271	"inactivity-timeout" on page 8505	"inactivity-timeout (Applications)" on page 8508
"learn-sip-register" on page 9695	"protocol (Applications)" on page 13431	"rpc-program-number" on page 16004
"rpc-program-number (Applications)" on page 16005	"sip-call-hold-timeout" on page 27974	"snmp-command" on page 28064
"source-port" on page 28277	"source-port (Applications)" on page 28280	"term (Applications)" on page 29110
"ttl-threshold" on page 30183	"uuid" on page 30621	

[edit applications application ike-esp-nat]		
"child-inactivity-timeout" on page 3314	"gate-timeout" on page 7651	

[edit applications application {application-name} term {term-name}]		
"alg (Applications)" on page 2138	"destination-port (Applications)" on page 5465	"icmp-code (Applications)" on page 8256
"icmp-type (Applications)" on page 8271	"inactivity-timeout (Applications)" on page 8508	"protocol (Applications)" on page 13431
"rpc-program-number (Applications)" on page 16005	"source-port (Applications)" on page 28280	"uuid (Applications)" on page 30623

[application-set]

[edit applications application-set {application-set-name}]		
"application" on page 2299	"description (Applications)" on page 5326	

[edit bridge-domains]

IN THIS SECTION

- [<domain-name>] | 55
- [bridge-options] | 55
- [forwarding-options] | 56
- [interface] | 59
- [multicast-snooping-options] | 60
- [protocols] | 60
- [vxlan] | 63

Sub Hierarchies

"[edit bridge-domains <domain-name>]" on page 55

"[edit bridge-domains {bridge-domain-name} bridge-options]" on page 55

"[edit bridge-domains forwarding-options]" on page 56

"[edit bridge-domains {bridge-domain-name} interface]" on page 59

"[edit bridge-domains {bridge-domain-name} multicast-snooping-options]" on page 60

"[edit bridge-domains {bridge-domain-name} protocols]" on page 60

"[edit bridge-domains {bridge-domain-name} vxlan]" on page 63

[edit bridge-domains {bridge-domain-name}]		
"bridge-options" on page 3034	"community-vlans (MX Series)" on page 4478	"enable-mac-move-action" on page 6267
"forwarding-options (Bridge Domains or VLANs)" on page 7467	"interface (Logical Interfaces)" on page 8814	"isolated-vlan (MX Series)" on page 9408
"no-irb-layer-2-copy" on page 11486	"no-local-switching (VPLS)" on page 11502	"routing-interface" on page 15978
"service-id (Bridge Domains)" on page 16704	"vlan-id-list" on page 30805	"vlan-tags" on page 30865
"vxlan" on page 30991		

[<domain-name>]

[edit bridge-domains <domain-name>]	
"irb-enhanced-scale (Bridge Domain)" on page 8243	

[bridge-options]

[edit bridge-domains {bridge-domain-name} bridge-options]		
"interface-mac-limit" on page 8910	"mac-statistics" on page 10263	"mac-table-size" on page 10271

(Continued)

[edit bridge-domains {bridge-domain-name} bridge-options]		
"network-isolation-profile" on page 11314		

[edit bridge-domains {bridge-domain-name} bridge-options interface {interface-name}]		
"action-priority (bridge domain)" on page 1769	"interface-mac-limit" on page 8910	"mac-pinning" on page 10254
"static-mac" on page 28549		

[edit bridge-domains {bridge-domain-name} bridge-options interface {interface-name} interface-mac-limit {limit}]	
"packet-action" on page 12129	

[edit bridge-domains {bridge-domain-name} bridge-options interface-mac-limit {limit}]	
"packet-action" on page 12129	

[forwarding-options]

[edit bridge-domains {name} forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit bridge-domains {name} forwarding-options dhcp-relay dhcpv6 group {name} relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit bridge-domains {name} forwarding-options dhcp-relay dhcpv6 relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit bridge-domains {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit bridge-domains {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit bridge-domains {name} forwarding-options dhcp-relay relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit bridge-domains {name} forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit bridge-domains {name} forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit bridge-domains {bridge-domain-name} forwarding-options dhcp-security]		
"arp-inspection (MX Series)" on page 2469	"dhcp-security (MX Series)" on page 5644	"group (DHCP Security for MX Series)" on page 7827
"ip-source-guard (MX Series)" on page 9274	"no-dhcp-snooping" on page 11436	"option-82" on page 11884

[edit bridge-domains {bridge-domain-name} forwarding-options dhcp-security group {group-name}]		
"interface (DHCP Security for MX Series)" on page 8866	"overrides (DHCP Security for MX Series)" on page 12072	

[edit bridge-domains {bridge-domain-name} forwarding-options dhcp-security group {group-name} interface {interface-name}]	
"static-ip (MX Series)" on page 28542	

[edit bridge-domains {bridge-domain-name} forwarding-options dhcp-security group {group-name} overrides]		
"trusted" on page 30170	"untrusted" on page 30432	

[edit bridge-domains {bridge-domain-name} forwarding-options dhcp-security option-82]	
"remote-id (MX Series)" on page 14337	

[edit bridge-domains {bridge domain name} forwarding-options dhcp-security option-82 circuit-id prefix]	
"routing-instance-name" on page 15943	

[edit bridge-domains {bridge-domain-name} forwarding-options dhcp-security remote-id option-82]	
"host-name" on page 8147	

[edit bridge-domains {bridge-domain-name} forwarding-options group group {group-name} overrides]	
"no-option82" on page 11541	

[interface]

[edit bridge-domains {bridge-domain-name} interface]	
"protect-interface" on page 13409	

[multicast-snooping-options]

[edit bridge-domains {bridge-domain-name} multicast-snooping-options]		
"flood-groups" on page 7241	"forwarding-cache (Bridge Domains)" on page 7407	"host-outbound-traffic (Multicast Snooping)" on page 8162
"ignore-stp-topology-change" on page 8412		

[edit bridge-domains {bridge-domain-name} multicast-snooping-options forwarding-cache]	
"threshold (Bridge Domains)" on page 29319	

[protocols]

[edit bridge-domains {bridge-domain-name} protocols]		
"igmp-snooping" on page 8387	"mld-snooping" on page 10783	

[edit bridge-domains {bridge-domain-name} protocols igmp-snooping]		
"evpn-ssm-reports-only" on page 6553	"immediate-leave" on page 8454	"interface (Bridge Domains)" on page 8845
"proxy (Bridge Domains)" on page 13819	"traceoptions (Protocols IGMP Snooping)" on page 29678	"version (IGMP Snooping)" on page 30704
"vlan (Bridge Domains)" on page 30788		

[edit bridge-domains {bridge-domain-name} protocols igmp-snooping interface {interface-name}]		
"group-limit (IGMP and MLD Snooping)" on page 7846	"host-only-interface" on page 8154	"immediate-leave" on page 8454
"multicast-router-interface (IGMP Snooping)" on page 11084	"query-interval (Bridge Domains)" on page 13927	"query-last-member-interval (Bridge Domains)" on page 13933
"query-response-interval (Bridge Domains)" on page 13938	"robust-count (IGMP Snooping)" on page 15815	"static (Bridge Domains)" on page 28529

[edit bridge-domains {bridge-domain-name} protocols igmp-snooping interface {interface-name} static]	
"group (Bridge Domains)" on page 7806	

[edit bridge-domains {bridge-domain-name} protocols igmp-snooping interface {interface-name} static group]		
"source (Bridge Domains)" on page 28187	"source (Bridge Domains)" on page 28187	

[edit bridge-domains {bridge-domain-name} protocols igmp-snooping no-flood-to-multicast-router-interfaces]	
"no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping)" on page 11628	

[edit bridge-domains {bridge-domain-name} protocols igmp-snooping proxy]	
"source-address (IGMP and MLD Snooping)" on page 28118	

[edit bridge-domains {bridge-domain-name} protocols igmp-snooping vlan {vlan-id}]		
"interface (Bridge Domains)" on page 8845	"proxy (Bridge Domains)" on page 13819	

[edit bridge-domains {bridge-domain-name} protocols igmp-snooping vlan {vlan-id} interface {interface-name}]		
"group-limit (IGMP and MLD Snooping)" on page 7846	"host-only-interface" on page 8154	"immediate-leave" on page 8454
"multicast-router-interface (IGMP Snooping)" on page 11084	"query-interval (Bridge Domains)" on page 13927	"query-last-member-interval (Bridge Domains)" on page 13933
"query-response-interval (Bridge Domains)" on page 13938	"robust-count (IGMP Snooping)" on page 15815	"static (Bridge Domains)" on page 28529

[edit bridge-domains {bridge-domain-name} protocols igmp-snooping vlan {vlan-id} interface {interface-name} static]	
"group (Bridge Domains)" on page 7806	

[edit bridge-domains {bridge-domain-name} protocols igmp-snooping vlan {vlan-id} proxy]	
"source-address (IGMP and MLD Snooping)" on page 28118	

[edit bridge-domains {bridge-domain-name} protocols mld-snooping]		
"evpn-ssm-reports-only" on page 6553	"query-interval (Bridge Domains)" on page 13927	"query-last-member-interval (Bridge Domains)" on page 13933
"query-response-interval (Bridge Domains)" on page 13938	"version (MLD Snooping)" on page 30712	

[edit bridge-domains {bridge-domain-name} protocols mld-snooping no-flood-to-multicast-router-interfaces]	
"no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping)" on page 11628	

[vxlan]

[edit bridge-domains {bridge-domain-name} vxlan]		
"ingress-node-replication" on page 8634	"ovsdb-managed" on page 12088	"static-remote-vtep-list" on page 28566

[edit bridge-options]

[edit bridge-options]	
"mac-table-aging-time (MX Series)" on page 10265	

[edit chassis]

IN THIS SECTION

- [(EX Series)] | 68
- [(QFX Series)] | 69
- [afeb] | 69
- [aggregated-devices] | 70
- [alarm] | 70
- [cb] | 71
- [cfeb] | 72
- [cluster] | 72
- [fabric] | 73

- [fabric-card] | 74
- [feb] | 74
- [forwarding-options] | 75
- [fpc] | 76
- [high-availability] | 84
- [interconnect-device] | 85
- [lcc] | 86
- [lcd-menu] | 88
- [member] | 88
- [memory-enhanced] | 89
- [multicast-loadbalance] | 89
- [network-services] | 90
- [network-slices] | 90
- [node-group] | 91
- [nssu] | 92
- [pem] | 92
- [pseudowire-service] | 93
- [psm] | 93
- [redundancy] | 93
- [redundancy-group] | 94
- [routing-engine] | 95
- [satellite-management] | 95
- [source] | 97
- [synchronization] | 98
- [tfeb] | 102

Sub Hierarchies

"[edit chassis (EX Series)]" on page 68

"[edit chassis (QFX Series)]" on page 69

"[edit chassis afeb]" on page 69

"[edit chassis aggregated-devices]" on page 70

"[edit chassis alarm]" on page 70

"[edit chassis cb]" on page 71

"[edit chassis cfeb {slot-number}]" on page 72

"[edit chassis cluster]" on page 72

"[edit chassis fabric]" on page 73

"[edit chassis fabric-card]" on page 74

"[edit chassis feb {slot-number}]" on page 74

"[edit chassis forwarding-options]" on page 75

"[edit chassis fpc {slot-number}]" on page 76

"[edit chassis high-availability]" on page 84

"[edit chassis interconnect-device {name}]" on page 85

"[edit chassis lcc {number}]" on page 86

"[edit chassis lcd-menu]" on page 88

"[edit chassis member]" on page 88

"[edit chassis memory-enhanced]" on page 89

"[edit chassis multicast-loadbalance]" on page 89

"[edit chassis network-services]" on page 90

"[edit chassis network-slices]" on page 90

"[edit chassis node-group {name}]" on page 91

"[edit chassis nssu]" on page 92

"[edit chassis pem]" on page 92

"[edit chassis pseudowire-service]" on page 93

"[edit chassis psm]" on page 93

"[edit chassis redundancy]" on page 93

"[edit chassis redundancy-group]" on page 94

"[edit chassis routing-engine]" on page 95

"[edit chassis satellite-management]" on page 95

"[edit chassis source]" on page 97

"[edit chassis synchronization]" on page 98

"[edit chassis tfeb]" on page 102

[edit chassis]		
"alarm (chassis)" on page 2100	"alarm-port" on page 2111	"auto-image-upgrade" on page 2715
"cluster (Chassis)" on page 4369	"container-devices" on page 4576	"copy-tos-to-outer" on page 4665
"craft-lockout" on page 4703	"dedicated-ukern-cpu (BFD)" on page 5174	"ecmp-alb" on page 6191
"effective-shaping-rate" on page 6204	"enhanced-policer" on page 6403	"extended-statistics" on page 6704
"feb (chassis)" on page 6982	"fib-prioritization" on page 7007	"flexible-queuing-mode" on page 7233
"forwarding-options (Unified Forwarding Table Profile)" on page 7493	"fpc (Chassis)" on page 7501	"fpc (M320, T320 and T640 Devices)" on page 7516
"fpc (MX Series 5G Universal Routing Platforms)" on page 7519	"fpc (PTX Series Devices)" on page 7525	"fpc (QFX Series)" on page 7499
"fpc error" on page 7507	"fpc-feb-connectivity" on page 7509	"fpc-nmi-volt-fail-knob" on page 7521
"fpc-resync" on page 7529	"fru-poweron-sequence" on page 7613	"high-availability (Chassis)" on page 8064

(Continued)

[edit chassis]		
"host-outbound media-interface" on page 8157	"icmp (Error Message Rate Limit)" on page 8257	"icmp6 (Error Message Rate Limit)" on page 8273
"input-firewall-optimized-profile" on page 8701	"interconnect-device (Chassis)" on page 8802	"ioc-npc-connectivity" on page 9230
"ipvacl-firewall-l2-profile" on page 9378	"iracl-firewall-ipv4-profile" on page 9383	"ivacl-firewall-no-portrange-profile" on page 9415
"jnu-management" on page 9429	"lcd-menu" on page 9667	"loopback-firewall-optimization" on page 10121
"loopback-firewallv6-optimization" on page 10123	"maximum-ecmp" on page 10520	"memory-enhanced" on page 10629
"multicast-loadbalance" on page 11072	"network-services" on page 11320	"network-slices" on page 11322
"no-power-budget" on page 11550	"node-group (Chassis)" on page 11641	"nssu" on page 11719
"oss-map" on page 11944	"packet-scheduling" on page 12158	"pem (MX Series Devices, ACX7024 Devices)" on page 12399
"pem (T640, T1600, and T4000 Routers with Six-Input DC Power Supply)" on page 12400	"periodic (Chassis)" on page 12436	"policer-drop-probability-low" on page 12635
"ppp-subscriber-services" on page 12959	"pseudowire-service (Pseudowire Subscriber Interfaces)" on page 13843	"realtime-ukern-thread (BFD)" on page 14099
"redundancy" on page 14166	"redundancy-group (Chassis - MX Series)" on page 14176	"resource-watch" on page 15622

(Continued)

[edit chassis]		
"routing-engine" on page 15904	"routing-engine (Chassis)" on page 15905	"satellite-management (Junos Fusion)" on page 16310
"sfm (Chassis)" on page 16912	"sib" on page 27939	"slow-pfe-alarm" on page 28044
"synchronization (ACX Series)" on page 28787	"synchronization (M Series and T Series)" on page 28792	"synchronization (MX Series)" on page 28795
"synchronization (PTX Series)" on page 28802	"thermal-health-check" on page 29305	"usb" on page 30509
"vrf-mtu-check" on page 30962		

[(EX Series)]

[edit chassis (EX Series)]		
"psu" on page 13852	"redundancy (Graceful Switchover)" on page 14170	

[edit chassis (EX Series) psu]	
"redundancy (Power Management)" on page 14201	

[edit chassis (EX Series) psu redundancy]	
"n-plus-n (Power Management)" on page 11160	

[(QFX Series)]

[edit chassis (QFX Series)]	
"forwarding-options (Bridge Domains or VLANs)" on page 7467	

[edit chassis (QFX Series) forwarding-options {profile-name}]	
"num-65-127-prefix" on page 11734	

[edit chassis (QFX Series) forwarding-options lpm-profile]	
"prefix-65-127-disable" on page 13064	

[edit chassis (QFX Series) fpc {fpc-id} pic {pic-id}]	
"fibre-channel (Port)" on page 7013	

[edit chassis (QFX Series) fpc {fpc-id} pic {pic-id} fibre-channel]	
"port-range" on page 12861	

[afeb]

[edit chassis afeb slot {name} pfe core]	
"channel" on page 3286	

[aggregated-devices]

[edit chassis aggregated-devices]		
"ethernet (Chassis)" on page 6484	"maximum-links" on page 10532	"sonet" on page 28108

[edit chassis aggregated-devices ethernet]		
"device-count (MX Series)" on page 5567	"lacp" on page 9605	"lacp (802.3ad)" on page 9607

[edit chassis aggregated-devices ethernet lacp]		
"link-protection (Chassis Cluster)" on page 9805	"link-protection (Protocols LACP)" on page 9813	"system-priority (chassis)" on page 28855

[edit chassis aggregated-devices ethernet lacp link-protection]	
"non-revertive (Chassis)" on page 11661	

[edit chassis aggregated-devices sonet]	
"device-count (MX Series)" on page 5567	

[alarm]

[edit chassis alarm]		
"ethernet (Alarm)" on page 6482	"fibre-channel (Alarm)" on page 7010	"management-ethernet (Alarm)" on page 10307

(Continued)

[edit chassis alarm]		
"pem-absence" on page 12397		

[edit chassis alarm ethernet (Alarm)]		
"link-down (QFX Series)" on page 9788		

[edit chassis alarm fibre-channel]		
"link-down (QFX Series)" on page 9788		

[edit chassis alarm l2-chan-err-ethernet]		
"l2-channel-error-threshold (chassis)" on page 9520	"l2-channel-errors (chassis)" on page 9521	

[cb]

[edit chassis cb 0]		
"temperature-sensor " on page 29063		

[cfeb]

[edit chassis cfeb {slot-number}]		
"on-error" on page 11819	"retry-count (FPC Error Condition)" on page 15698	"sanity-poll" on page 16302

[edit chassis cfeb slot {name} pfe core]	
"channel" on page 3286	

[cluster]

[edit chassis cluster]		
"apply-groups (Chassis Cluster)" on page 2424	"control-link-recovery" on page 4633	"control-ports" on page 4637
"heartbeat-interval" on page 8022	"heartbeat-threshold" on page 8024	"network-management" on page 11317
"redundancy-group (Chassis Cluster)" on page 14173	"redundancy-mode" on page 14207	"reth-count (Chassis Cluster)" on page 15675
"traceoptions (Chassis Cluster)" on page 29506	"use-active-child-mac-on-reth" on page 30511	"use-actual-mac-on-physical-interfaces" on page 30512

[edit chassis cluster redundancy-group {group-number}]		
"hold-down-interval" on page 8088	"interface-monitor" on page 8922	"ip-monitoring" on page 9244
"node (Chassis Cluster Redundancy Group)" on page 11631	"preempt (Chassis Cluster)" on page 13021	

[edit chassis cluster redundancy-group group-number]	
"gratuitous-arp-count" on page 7787	

[edit chassis cluster redundancy-group {group-number} interface-monitor {interface}]	
"weight" on page 31095	

[edit chassis cluster redundancy-group {group-number} ip-monitoring]		
"global-threshold" on page 7740	"retry-count (Chassis Cluster)" on page 15700	"retry-count (FPC Error Condition)" on page 15698
"retry-interval (Chassis Cluster)" on page 15707	"weight" on page 31095	

[edit chassis cluster redundancy-group{group-number} ip-monitoring family {family-name} {IP–address}]	
"interface (Chassis Cluster)" on page 8850	

[edit chassis cluster redundancy-group{group-number}ip-monitoring]	
"global-weight" on page 7743	

[fabric]

[edit chassis fabric]		
"degraded" on page 5236	"disable-grant-bypass" on page 5762	"oam detection-disable" on page 11756

(Continued)

[edit chassis fabric]		
"oam runtime-disable" on page 11764	"redundancy-mode" on page 14191	"upgrade-mode" on page 30464

[edit chassis fabric degraded]		
"fpc-offline-on-blackholing" on page 7524	"fpc-restart" on page 7528	

[edit chassis fabric event]	
"reachability-fault" on page 14083	

[fabric-card]

[edit chassis fabric-card]	
"reset-count" on page 15584	

[feb]

[edit chassis feb {slot-number}]		
"on-error" on page 11819	"retry-count (FPC Error Condition)" on page 15698	"sanity-poll" on page 16302

[edit chassis feb slot {slot-number}]		
"power" on page 12915	"ucode-imem-remap" on page 30304	

[edit chassis feb slot {name} pfe core]	
"channel" on page 3286	

[edit chassis feb slot {slot-number} pfe {pfe-id} core {core-id} channel {channel-id} tunnel-services]	
"bandwidth (Tunnel Services)" on page 2880	

[forwarding-options]

[edit chassis forwarding-options]		
"em-hw-profile" on page 6245	"epacl-firewall-optimization" on page 6421	"vxlan-gbp-l2-profile" on page 30999
"vxlan-gbp-l3-profile" on page 31001	"vxlan-gbp-profile" on page 31003	

[edit chassis forwarding-options lpm-profile]	
"unicast-in-lpm" on page 30342	

[fpc]

[edit chassis fpc {slot-number}]		
"allow-sram-parity-errors" on page 2206	"cfp-to-et" on page 3272	"ether-type (Chassis FPC)" on page 6477
"fabric loopback wan" on page 6755	"flow-table-size" on page 7351	"inline-services" on page 8674
"ir-mode" on page 9381	"license-mode" on page 9739	"max-queues" on page 10458
"np-cache (Flexible PIC Concentrator)" on page 11711	"number-of-ports" on page 11737	"offline-on-fabric-bandwidth-reduction" on page 11786
"pfe" on page 12460	"pfe power off" on page 12463	"pic (M Series and T Series Routers)" on page 12487
"pic (MX Series Routers)" on page 12496	"pic (PTX Series Routers)" on page 12500	"power" on page 12915
"power" on page 12915	"power-budget-priority" on page 12919	"sampling-instance" on page 16297
"sanity-poll" on page 16302	"slamon-services" on page 28032	"temperature-sensor " on page 29063
"traffic-manager" on page 30036	"ukern-trace" on page 30322	"vpn-localization" on page 30942

[edit chassis fpc {slot-number} fabric]	
"blackhole-action" on page 2979	

[edit chassis fpc {slot-number} flow-table-size]	
"ipv4-flow-table-size " on page 9332	

[edit chassis fpc {slot-number} inline-services]		
"flex-flow-sizing" on page 7225	"report-zero-oif-gw-on-discard" on page 14406	"use-extended-flow-memory" on page 30515

[edit chassis fpc {slot-number} inline-services flow-table-size]		
"mpls-flow-table-size " on page 10972	"vpls-flow-table-size" on page 30924	

[edit chassis fpc {slot-number} inline-services ipv6 flow-table-size]		
"ipv6-extended-attrib " on page 9349	"ipv6-flow-table-size " on page 9360	

[edit chassis fpc {slot} inline-video-monitoring]	
"flow-table-size (Chassis)" on page 7354	

[edit chassis fpc {fpc-slot} pfe {pfe-id}]	
"forwarding-packages" on page 7478	

[edit chassis fpc {name} pfe core]	
"channel" on page 3286	

[edit chassis fpc {slot-number} pfe {pfe-id} core {core-id} channel {channel-id} tunnel-services]	
"bandwidth (Tunnel Services)" on page 2880	

[edit chassis fpc {slot-number} pfe {identifier} exception-reporting]	
"category" on page 3221	

[edit chassis fpc {fpc-slot} pfe {pfe-id} forwarding-packages]	
"mobility" on page 10793	

[edit chassis fpc {slot-number} pfe identifier]	
"exception-reporting" on page 6572	

[edit chassis fpc {slot-number} pic {pic-number}]		
"account-layer2-overhead (PIC Level)" on page 1706	"adaptive-services" on page 1832	"egress-policer-overhead" on page 6213
"forwarding-mode (100-Gigabit Ethernet)" on page 7459	"hash-key (Chassis LAG)" on page 7982	"hash-key (Forwarding Options)" on page 7984
"ingress-policer-overhead" on page 8638	"inline-services (PIC level)" on page 8676	"max-queues-per-interface" on page 10461
"mlfr-uni-nni-bundles-inline" on page 10790	"multi-link-layer-2-inline" on page 11054	"no-mcast-replication" on page 11510
"no-pre-classifier" on page 11551	"number-of-ports" on page 11737	"pic-mode (MX Series)" on page 12491
"port (Chassis)" on page 12815	"power" on page 12915	"ptp-mode" on page 13865
"recovered-clock" on page 14142	"red-buffer-occupancy" on page 14154	"services-offload" on page 16784

(Continued)

[edit chassis fpc {slot-number} pic {pic-number}]		
"short-reach-mode" on page 16988	"speed (24-port and 12-port 10 Gigabit Ethernet PIC)" on page 28342	"subscriber-aware-services" on page 28684
"traffic-manager" on page 30036	"tunnel-port" on page 30223	"tunnel-queuing" on page 30230
"tunnel-services (Chassis)" on page 30234		

[edit chassis fpc {slot-number} pic {pic-number} (port {port-number} | port-range {port-range-low} {port-range-high})]	
"channel-speed" on page 3287	

[edit chassis fpc {slot-number} pic {pic-number} adaptive-services]	
"service-package" on page 16721	

[edit chassis fpc {slot-number} pic {pic-number} adaptive-services service-package]	
"extension-provider" on page 6716	

[edit chassis fpc {slot-number} pic {pic-number} adaptive-services service-package extension-provider]		
"control-cores" on page 4631	"data-cores" on page 4827	"data-flow-affinity" on page 4832
"forwarding-db-size" on page 7454	"object-cache-size" on page 11765	"package (Loading on PIC)" on page 12128
"policy-db-size" on page 12674	"session-offload" on page 16829	"syslog (Chassis)" on page 28817
"wired-max-processes" on page 31131	"wired-process-mem-size" on page 31133	

[edit chassis fpc {slot-number} pic {pic-number} adaptive-services service-package extension-provider data-flow-affinity]		
"hash-key (Chassis)" on page 7980		

[edit chassis fpc {slot-number} pic {pic-number} adaptive-services service-package extension-provider syslog {facility}]		
"destination (Chassis)" on page 5410		

[edit chassis fpc {slot-number} pic {pic-number} ethernet]		
"pic-mode (Chassis T1 Mode)" on page 12494		

[edit chassis fpc {slot} pic {slot} forwarding-mode]		
"sa-multicast (100-Gigabit Ethernet)" on page 16264	"vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP)" on page 30855	

[edit chassis fpc {slot} pic {slot} forwarding-mode vlan-steering]		
"vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP)" on page 30850		

[edit chassis fpc {slot-number} pic {pic-number} hash-key]		
"family (chassis)" on page 6851		

[edit chassis fpc {slot-number} pic {pic-number} hash-key family]		
"inet (chassis)" on page 8558	"multiservice" on page 11128	

[edit chassis fpc {slot-number} pic {pic-number} hash-key family multiservice]	
"payload" on page 12279	

[edit chassis fpc {slot-number} pic {number} inline-services]	
"bandwidth (Inline Services)" on page 2845	

[edit chassis fpc {slot-number} pic {pic-number} linerate-mode]	
"linerate-mode" on page 9780	

[edit chassis fpc {slot-number} pic {pic-number} mixed-rate-mode]	
"mixed-rate-mode" on page 10774	

[edit chassis {lcc number} fpc {slot-number} pic {pic-number} mixed-rate-mode](Routing Matrix)]		
"mixed-rate-mode" on page 10774	"speed (24-port and 12-port 10 Gigabit Ethernet PIC)" on page 28342	

[edit chassis fpc {fpc-slot} pic {pic-number} pic-mode {pic-speed}]	
"number-of-ports" on page 11737	

[edit chassis fpc {slot} pic {slot} port {port-number}]		
"forwarding-mode (PTX Series Packet Transport Routers)" on page 7460	"no-temperature-monitoring" on page 11595	"speed" on page 28344
"speed (24-port and 12-port 10 Gigabit Ethernet PIC)" on page 28342	"traffic-manager" on page 30036	"Unused" on page 30435

[edit chassis fpc {slot} pic {slot} port {port-number} forwarding-mode]	
"sa-multicast (PTX Series Packet Transport Routers)" on page 16266	

[edit chassis fpc {slot-number} pic {pic-number} port {port-number} traffic-manager]	
"mode (Traffic Manager)" on page 10807	

[edit chassis fpc {fpc-slot} pic {pic-slot} port-range {port-range-low} {port-range-high}]	
"short-reach-mode" on page 16988	

[edit chassis fpc {fpc-slot-number} pic {pic-slot-number} services-offload]	
"low-latency" on page 10176	

[edit chassis fpc {slot-number} pic {pic-number} traffic-manager]		
"egress-shaping-overhead" on page 6223	"ingress-shaping-overhead" on page 8649	"mode (Traffic Manager)" on page 10807

[edit chassis fpc {slot-number} pic {number} tunnel-services]	
"bandwidth (Tunnel Services)" on page 2880	

[edit chassis fpc {fpc-number} route-localization]		
"fib-local" on page 7005	"fib-remote" on page 7009	

[edit chassis fpc {slot-number} sanity-poll]		
"on-error" on page 11819	"retry-count (FPC Error Condition)" on page 15698	

[edit chassis fpc {slot} slc {slc-id}]		
"power" on page 12915	"power" on page 12918	

[edit chassis fpc slot-number pic number inline-services]	
"service-port (Inline Services) " on page 16725	

[edit chassis fpc slot-number pic number tunnel-services]	
"tunnel-port (GRE)" on page 30224	

[edit chassis fpc slot-number pic slot-number hash-key family inet]	
"symmetric-hash" on page 28768	

[edit chassis fpc slot-number pic slot-number hash-key family multiservice]	
"symmetric-hash" on page 28768	

[edit chassis fpc {slot-number} traffic-manager]		
"enhanced-priority-mode " on page 6404	"packet-timestamp" on page 12162	"queue-threshold" on page 13953

[edit chassis fpc {slot-number}fabric]	
"bandwidth-degradation" on page 2841	

[high-availability]

[edit chassis high-availability]		
"hardware-upgrade" on page 7976	"local-id (High Availability)" on page 9948	"peer-id" on page 12367
"services-redundancy-group" on page 16793	"software-upgrade" on page 28090	"traceoptions (High Availability)" on page 29471

[edit chassis high-availability peer-id]		
"liveness-detection (high availability)" on page 9848	"peer-forwarding-ip" on page 12357	"peer-forwarding-ip" on page 12357

[edit chassis high-availability peer-id peer-forwarding-ip]	
"liveness-detection (high availability)" on page 9848	

[edit chassis high-availability services-redundancy-group]		
"activeness-probe" on page 1821	"managed-services" on page 10304	"monitor (Multinode High Availability)" on page 10903
"virtual-ip" on page 30761		

[edit chassis high-availability services-redundancy-group {name} activeness-probe]		
"bfd-liveliness (Activeness Probe)" on page 2917		

[edit chassis high-availability services-redundancy-group {name} monitor]		
"bfd-liveliness" on page 2914	"interface (Service Redundancy Group Monitor)" on page 9025	"monitor-object" on page 10906

[edit chassis high-availability services-redundancy-group {name} monitor monitor-object {name}]		
"bfd-liveliness" on page 2914	"interface (Service Redundancy Group Monitor)" on page 9025	

[interconnect-device]

[edit chassis interconnect-device {name}]		
"alarm (chassis)" on page 2100	"container-devices" on page 4576	"routing-engine" on page 15904

[edit chassis interconnect-device {name} alarm]		
"ethernet (Alarm)" on page 6482	"fibre-channel (Alarm)" on page 7010	"management-ethernet (Alarm)" on page 10307

[edit chassis interconnect-device {name} alarm ethernet (Alarm)]		
"link-down (QFX Series)" on page 9788		

[edit chassis interconnect-device {name} fpc {slot}]	
"pic (Port)" on page 12498	

[lcc]

[edit chassis lcc {number}]		
"fpc (TX Matrix and TX Matrix Plus Routers)" on page 7534	"pem (T640, T1600, and T4000 Routers with Six-Input DC Power Supply)" on page 12400	

[edit chassis lcc {name} fpc {slot-number}]		
"pfe" on page 12460	"sampling-instance" on page 16297	"sanity-poll" on page 16302
"ukern-trace" on page 30322		

[edit chassis lcc {name} fpc {name} pfe core]	
"channel" on page 3286	

[edit chassis lcc {number} fpc {slot-number} pic {pic-number}]		
"no-mcast-replication" on page 11510	"port (Chassis)" on page 12815	"power" on page 12915
"red-buffer-occupancy" on page 14154	"traffic-manager" on page 30036	

[edit chassis lcc {number} fpc {slot-number} pic {pic-number}] (Routing Matrix)]		
"max-queues-per-interface" on page 10461		

[edit chassis lcc {number} fpc {slot-number} pic {pic-number} linerate-mode]		
"linerate-mode" on page 9780		

[edit chassis lcc {name} fpc {name} pic {name} port]		
"no-temperature-monitoring" on page 11595		

[edit chassis lcc {number} fpc {slot-number} pic {pic-number} traffic-manager]		
"egress-shaping-overhead" on page 6223	"ingress-shaping-overhead" on page 8649	"mode (Traffic Manager)" on page 10807

[edit chassis lcc {number} fpc {number} sanity-poll]		
"on-error" on page 11819	"retry-count (FPC Error Condition)" on page 15698	

[edit chassis lcc {lcc-number} pem]		
"feeds (T640, T1600, and T4000 Routers with Six-Input DC Power Supply)" on page 7001	"redundancy (Junos OS)" on page 14183	

[lcd-menu]

[edit chassis lcd-menu]	
"menu-item" on page 10634	

[edit chassis lcd-menu fpc {slot-number}]	
"menu-item" on page 10634	

[member]

[edit chassis member {name} fabric]	
"upgrade-mode" on page 30464	

[edit chassis member {name} fabric event]	
"reachability-fault" on page 14083	

[edit chassis member {name} fpc {slot-number}]		
"pfe" on page 12460	"ukern-trace" on page 30322	

[edit chassis member {name} fpc {name} pfe core]	
"channel" on page 3286	

[edit chassis member {name} fpc {name} pic]		
"port (Chassis)" on page 12815	"power" on page 12915	

[edit chassis member {name} fpc {name} pic {name} port]	
"no-temperature-monitoring" on page 11595	

[edit chassis member {member-number} fpc slot {slot-number}]	
"sampling-instance" on page 16297	

[memory-enhanced]

[edit chassis memory-enhanced]		
"filter (memory-enhanced)" on page 7150	"route (chassis)" on page 15853	"vpn-label" on page 30941

[multicast-loadbalance]

[edit chassis multicast-loadbalance]	
"disable (Multicast Load Balancing)" on page 5773	

[network-services]

[edit chassis network-services]	
"enhanced-mode (Network Services)" on page 6398	

[edit chassis network-services enhanced-ip]	
"limited-ifl-scaling" on page 9768	

[network-slices]

[edit chassis network-slices]		
"guest-network-functions" on page 7949	"vlan-isolation" on page 30826	

[edit chassis network-slices guest-network-functions]	
"gnf" on page 7746	

[edit chassis network-slices guest-network-functions gnf {id} {af-name}]		
"collapsed-forward" on page 4406	"control-plane-bandwidth-percent (Node Slicing)" on page 4634	"description (GNF)" on page 5332
"fpc-slice" on page 7530	"fpcs (Node Slicing)" on page 7536	

[edit chassis network-slices guest-network-functions gnf af-name]		
"description (Abstracted Fabric)" on page 5324		

[edit chassis network-slices guest-network-functions gnf id]		
"af-name" on page 2040	"peer-gnf" on page 12359	

[node-group]

[edit chassis node-group {name}]		
"alarm (chassis)" on page 2100	"container-devices" on page 4576	"node-device (Chassis)" on page 11637
"routing-engine" on page 15904		

[edit chassis node-group aggregated-devices]		
"ethernet" on page 6480		

[edit chassis node-group {name} alarm]		
"ethernet (Alarm)" on page 6482	"fibre-channel (Alarm)" on page 7010	"management-ethernet (Alarm)" on page 10307

[edit chassis node-group {name} alarm fibre-channel]		
"link-down (QFX Series)" on page 9788		

[nssu]

[edit chassis nssu]		
"rcp-count" on page 14075	"upgrade-group" on page 30460	

[edit chassis nssu upgrade-group {group-name}]		
"fpcs (NSSU Upgrade Groups)" on page 7538	"member (NSSU Upgrade Groups)" on page 10617	

[edit chassis nssu upgrade-group {group-name} member {member-id}]	
"fpcs (NSSU Upgrade Groups)" on page 7538	

[pem]

[edit chassis pem]		
"feeds (T640, T1600, and T4000 Routers with Six-Input DC Power Supply)" on page 7001	"input-current (T4000 Routers)" on page 8695	"redundancy (Junos OS)" on page 14183
"watchdog (PSM)" on page 31028		

[pseudowire-service]

[edit chassis pseudowire-service]	
"device-count (Pseudowire Subscriber Interfaces)" on page 5569	

[psm]

[edit chassis psm]		
"redundancy (Junos OS Evolved)" on page 14185	"watchdog (PSM)" on page 31028	

[redundancy]

[edit chassis redundancy]		
"cfeb" on page 3267	"failover (Chassis)" on page 6778	"feb (Creating a Redundancy Group)" on page 6984
"graceful-switchover (Chassis)" on page 7786	"keepalive-time (Chassis)" on page 9471	"on-disk-failure" on page 11817
"routing-engine (Chassis Redundancy)" on page 15908	"sfm (Chassis Redundancy)" on page 16913	"ssb" on page 28421

[edit chassis redundancy failover]		
"not-on-disk-underperform" on page 11681	"on-loss-of-keepalives" on page 11826	"on-loss-of-keepalives" on page 11824

[edit chassis redundancy feb]	
"redundancy-group" on page 14171	

[edit chassis redundancy feb redundancy-group {group-name}]		
"description (Chassis Redundancy)" on page 5328	"feb (Assigning a FEB to a Redundancy Group)" on page 6981	"no-auto-failover" on page 11406

[edit chassis redundancy on-disk-failure]	
"disk-failure-action" on page 5823	

[redundancy-group]

[edit chassis redundancy-group interface-type]		
"redundant-logical-tunnel" on page 14211	"redundant-virtual-tunnel" on page 14222	

[routing-engine]

[edit chassis routing-engine]		
"on-disk-failure" on page 11817	"on-disk-failure (Chassis Routing Engine)" on page 11818	

[edit chassis routing-engine disk]	
"smart-check" on page 28045	

[edit chassis routing-engine on-disk-failure]		
"disk-failure-action" on page 5823	"disk-failure-action" on page 5821	

[satellite-management]

[edit chassis satellite-management]		
"aging-timer (Junos Fusion)" on page 2091	"auto-satellite-conversion (Junos Fusion)" on page 2736	"cluster (Junos Fusion)" on page 4375
"environment-monitoring-policy (satellite-management)" on page 6417	"fpc (Junos Fusion)" on page 7512	"psu (satellite-management)" on page 13853
"single-home (Junos Fusion)" on page 27968	"upgrade-groups (Junos Fusion)" on page 30461	"uplink-failure-detection (Junos Fusion)" on page 30467

[edit chassis satellite-management auto-satellite-conversion]	
"satellite (Junos Fusion Automatic Satellite Conversion)" on page 16305	

[edit chassis satellite-management cluster {cluster-name}]		
"cascade-ports (Junos Fusion Satellite Device Cluster)" on page 3218	"cluster-id (Junos Fusion Satellite Device Cluster)" on page 4373	"fpc (Junos Fusion Satellite Device Cluster)" on page 7514

[edit chassis satellite-management cluster {cluster-name} fpc {slot-id}]		
"alias (Junos Fusion Satellite Device Clustering)" on page 2150	"description (Junos Fusion Satellite Device Cluster)" on page 5337	"member-id (Junos Fusion Satellite Device Cluster)" on page 10611
"system-id (Junos Fusion Satellite Device Cluster)" on page 28850		

[edit chassis satellite-management fpc {slot-id}]		
"alias (Junos Fusion)" on page 2148	"cascade-ports" on page 3216	"description (Junos Fusion)" on page 5336
"environment-monitoring-policy (satellite-management)" on page 6417	"serial-number (Junos Fusion)" on page 16601	"system-id (Junos Fusion)" on page 28848
"uplink-failure-detection (Junos Fusion)" on page 30467		

[edit chassis satellite-management psu]	
"redundancy (satellite-management)" on page 14203	

[edit chassis satellite-management psu redundancy]	
"n-plus-n (satellite-management)" on page 11162	

[edit chassis satellite-management redundancy-groups {redundancy-group-name}]	
"redundancy-group-id (Junos Fusion)" on page 14177	

[edit chassis satellite-management redundancy-groups {redundancy-group-name} peer-chassis-id {peer-chassis-id-number}]	
"no-auto-iccp-provisioning (Junos Fusion Redundancy Group)" on page 11407	

[edit chassis satellite-management single-home]	
"satellite (Junos Fusion Satellite Device Homing)" on page 16307	

[edit chassis satellite-management upgrade-groups {upgrade-group-name}]	
"satellite (Junos Fusion Satellite Software Upgrade Groups)" on page 16308	

[source]

[edit chassis source interfaces (external | interface {interface-name})]	
"wait-to-restore" on page 31017	

[edit chassis source interfaces (external-0/0 | external-1/0 | interface {interface-name})]	
"wait-to-restore" on page 31017	

[edit chassis source interfaces (external-a | external-b | interface {interface-name})]	
"wait-to-restore" on page 31017	

[synchronization]

[edit chassis synchronization]		
"clock-mode (Clock Synchronization)" on page 4359	"enable-extended-ql-tlv" on page 6260	"esmc-transmit (Chassis Synchronization)" on page 6458
"global-wait-to-restore" on page 7741	"hold-interval (Clock Synchronization)" on page 8091	"interfaces external" on page 9144
"max-transmit-quality-level" on page 10487	"minimum-quality" on page 10748	"network-option" on page 11318
"output interfaces (Chassis Synchronization)" on page 11996	"port (Chassis Synchronization)" on page 12818	"quality-mode-enable (MX)" on page 13920
"selection-mode" on page 16528	"source (Chassis Synchronization)" on page 28188	"source interfaces" on page 28231
"switchover-mode" on page 28761		

[edit chassis synchronization (PTX Series)]	
"quality-mode-enable (PTX)" on page 13922	

[edit chassis synchronization esmc-transmit interfaces]	
"squelch" on page 28404	

[edit chassis synchronization interfaces]		
"bits (Chassis Synchronization)" on page 2975		

[edit chassis synchronization interfaces (external-0/0 | external-1/0)]		
"e1-options (Clock Synchronization)" on page 6178	"pulse-per-second-enable" on page 13868	"signal-type" on page 27946

[edit chassis synchronization interfaces (external-0/0 | external-1/0) e1-options]		
"framing (E1 Options)" on page 7571	"line-encoding (E1 Options)" on page 9769	"sabit" on page 16267
"t1-options" on page 28886		

[edit chassis synchronization interfaces (external-0/0 | external-1/0) t1-options]		
"framing (T1 Options)" on page 7573	"line-encoding (T1 Options)" on page 9770	

[edit chassis synchronization interfaces bits]		
"e1-options (BITS Interfaces Signal Type)" on page 6176	"signal-type (BITS Interfaces)" on page 27947	"source-mode" on page 28250
"t1-options (BITS Interfaces Signal Type)" on page 28887		

[edit chassis synchronization interfaces bits e1-options]		
"framing (E1 Options for BITS Interfaces)" on page 7572	"sabit" on page 16267	

[edit chassis synchronization interfaces bits t1-options]		
"framing (T1 Options for BITS Interfaces)" on page 7575		

[edit chassis synchronization interfaces external]		
"e1-options (Clock Synchronization)" on page 6178	"pulse-per-second-enable" on page 13868	"signal-type" on page 27946
"t1-options" on page 28886		

[edit chassis synchronization interfaces external e1-options]		
"framing (E1 Options)" on page 7571	"line-encoding (E1 Options)" on page 9769	"sabit" on page 16267

[edit chassis synchronization interfaces external t1-options]		
"framing (T1 Options)" on page 7573	"line-encoding (T1 Options)" on page 9770	

[edit chassis synchronization output interfaces]		
"bits (Chassis Synchronization)" on page 2975		

[edit chassis synchronization output interfaces (external-0/0 | external-1/0)]		
"holdover-mode-disable" on page 8128	"minimum-quality" on page 10748	"source-mode" on page 28250
"tx-dnu-to-line-source-enable" on page 30268	"wander-filter-disable" on page 31021	

[edit chassis synchronization output interfaces bits]		
"squelch" on page 28404		

[edit chassis synchronization output interfaces external]		
"holdover-mode-disable" on page 8128	"minimum-quality" on page 10748	"source-mode" on page 28250
"tx-dnu-to-line-source-enable" on page 30268	"wander-filter-disable" on page 31021	

[edit chassis synchronization source]		
"interfaces (Chassis Synchronization Source)" on page 9124		

[edit chassis synchronization source (bits | gps | interfaces {interface-name})]		
"quality-level " on page 13915		

[edit chassis synchronization source (bits | gps)]		
"priority (Chassis Synchronization Source)" on page 13160	"request (Chassis Synchronization Source)" on page 14573	

[edit chassis synchronization source (external | {interface-name})]		
"force-switch" on page 7375		

[edit chassis synchronization source interfaces]		
"aggregated-ether" on page 2077	"priority (Chassis Synchronization Source)" on page 13160	"request (Chassis Synchronization Source)" on page 14573

[edit chassis synchronization source interfaces (external | interface {interface-name})]		
"port (Chassis Synchronization)" on page 12818	"priority (Clock Synchronization)" on page 13162	"quality-level (Clock Synchronization)" on page 13917
"request (Clock Synchronization)" on page 14579		

[edit chassis synchronization source interfaces (external-0/0 | external-1/0 | interface {interface-name})]		
"port (Chassis Synchronization)" on page 12818	"priority (Clock Synchronization)" on page 13162	"quality-level (Clock Synchronization)" on page 13917
"request (Clock Synchronization)" on page 14579		

[edit chassis synchronization source interfaces (external-a | external-b | interface {interface-name})]		
"port (Chassis Synchronization)" on page 12818	"priority (Clock Synchronization)" on page 13162	"quality-level (Clock Synchronization)" on page 13917
"request (Clock Synchronization)" on page 14579		

[tfeb]

[edit chassis tfeb slot {name} pfe core]	
"channel" on page 3286	

[edit class-of-service]

IN THIS SECTION

- [adaptive-shapers] | 107
- [adjustment-control-profiles] | 107
- [application-traffic-control] | 108
- [class] | 108
- [classifiers] | 108
- [code-point-aliases] | 110
- [congestion-notification-profile] | 110
- [drop-profiles] | 112
- [fabric] | 113
- [forwarding-class-sets] | 113
- [forwarding-classes] | 113
- [forwarding-classes-interface-specific] | 114
- [forwarding-policy] | 114
- [fragmentation-map] | 115
- [fragmentation-maps] | 115
- [host-outbound-traffic] | 116
- [interface] | 116
- [interfaces] | 117
- [loss-priority-maps] | 122
- [loss-priority-rewrites] | 123
- [multi-destination] | 123
- [queue] | 124
- [restricted-queues] | 124
- [rewrite-rules] | 124
- [routing-instances] | 125
- [scheduler-map-forwarding-class-sets] | 126
- [scheduler-maps] | 126
- [schedulers] | 126

- [shared-buffer] | **128**
- [system-defaults] | **128**
- [traffic-control-profile] | **129**
- [traffic-control-profiles] | **129**
- [virtual-channel-groups] | **130**

Sub Hierarchies

"[edit class-of-service adaptive-shapers {adaptive-shaper-name}]" on page 107

"[edit class-of-service adjustment-control-profiles {profile-name}]" on page 107

"[edit class-of-service application-traffic-control]" on page 108

"[edit class-of-service class]" on page 108

"[edit class-of-service classifiers]" on page 108

"[edit class-of-service code-point-aliases]" on page 110

"[edit class-of-service congestion-notification-profile {profile-name}]" on page 110

"[edit class-of-service drop-profiles {profile-name}]" on page 112

"[edit class-of-service fabric]" on page 113

"[edit class-of-service forwarding-class-sets {forwarding-class-set-name}]" on page 113

"[edit class-of-service forwarding-classes]" on page 113

"[edit class-of-service forwarding-classes-interface-specific]" on page 114

"[edit class-of-service forwarding-policy]" on page 114

"[edit class-of-service fragmentation-map]" on page 115

"[edit class-of-service fragmentation-maps {map-name}]" on page 115

"[edit class-of-service host-outbound-traffic]" on page 116

"[edit class-of-service interface {interface-name}]" on page 116

"[edit class-of-service interfaces {interface-name}]" on page 117

"[edit class-of-service loss-priority-maps]" on page 122

"[edit class-of-service loss-priority-rewrites]" on page 123

"[edit class-of-service multi-destination]" on page 123

"[edit class-of-service queue {queue-number} {class-name}]" on page 124

"[edit class-of-service restricted-queues]" on page 124

"[edit class-of-service rewrite-rules]" on page 124

"[edit class-of-service routing-instances {routing-instance-name}]" on page 125

"[edit class-of-service scheduler-map-forwarding-class-sets {scheduler-map-fcset-name}]" on page 126

"[edit class-of-service scheduler-maps {map-name}]" on page 126

"[edit class-of-service schedulers {scheduler-name}]" on page 126

"[edit class-of-service shared-buffer]" on page 128

"[edit class-of-service system-defaults]" on page 128

"[edit class-of-service traffic-control-profile {lcp-name}]" on page 129

"[edit class-of-service traffic-control-profiles {traffic-control-profile-name}]" on page 129

"[edit class-of-service virtual-channel-groups {group-name} {virtual-channel-name}]" on page 130

[edit class-of-service]		
"adaptive-shapers" on page 1836	"adjustment-control-profiles" on page 1961	"application-traffic-control" on page 2400
"classifiers (Class of Service)" on page 3383	"classifiers (Definition)" on page 3386	"classifiers (Definition)" on page 3388
"code-point-aliases" on page 4392	"congestion-notification-profile" on page 4525	"copy-plp-all" on page 4664
"dedicated-buffer" on page 5170	"dedicated-buffer-profile" on page 5172	"drop-profiles (Definition)" on page 5968

(Continued)

[edit class-of-service]		
"dscp (CoS Classifiers)" on page 5998	"fabric (Class-of-Service)" on page 6747	"forwarding-class-sets" on page 7446
"forwarding-classes" on page 7447	"forwarding-policy" on page 7479	"fragmentation-maps" on page 7548
"host-outbound-traffic (Class-of-Service)" on page 8158	"host-outbound-traffic (Class-of-Service)" on page 8160	"interfaces (Class of Service)" on page 9126
"interfaces (Class of Service)" on page 9128	"interfaces (Class of Service)" on page 9130	"loss-priority-maps" on page 10157
"loss-priority-maps (CoS)" on page 10160	"loss-priority-rewrites" on page 10166	"monitoring-profile" on page 10959
"multi-destination (Class of Service)" on page 11047	"non-strict-priority-scheduling" on page 11664	"policy-map (Class of Service)" on page 12686
"queue (Global Queues)" on page 13944	"restricted-queues" on page 15671	"rewrite-rules (Definition)" on page 15741
"rewrite-rules (Definition)" on page 15744	"rewrite-rules (Definition)" on page 15746	"routing-instances (CoS)" on page 15971
"scheduler-map-forwarding-class-sets" on page 16354	"scheduler-maps (Definition)" on page 16358	"schedulers (CoS)" on page 16368
"schedulers (CoS)" on page 16369	"shared-buffer" on page 16964	"system-defaults" on page 28842
"traceoptions (Access Port Security)" on page 29479	"traceoptions (Class of Service)" on page 29508	"traffic-class-map" on page 30003
"traffic-control-profiles" on page 30009	"translation-table" on page 30075	"tri-color" on page 30147

(Continued)

[edit class-of-service]		
"tri-color (EX Series Switches)" on page 30149	"virtual-channel-groups" on page 30742	"virtual-channels" on page 30743

[adaptive-shapers]

[edit class-of-service adaptive-shapers {adaptive-shaper-name}]	
"trigger (CoS)" on page 30159	

[edit class-of-service adaptive-shapers {adaptive-shaper-name} trigger {type}]	
"shaping-rate (CoS Adaptive Shapers)" on page 16932	

[adjustment-control-profiles]

[edit class-of-service adjustment-control-profiles {profile-name}]	
"application (Adjustment Control Profiles)" on page 2301	

[edit class-of-service adjustment-control-profiles {profile-name} application]		
"ancp (Adjustment Control Profiles)" on page 2249	"dhcp-tags (Adjustment Control Profiles)" on page 5652	"pppoe-tags (Adjustment Control Profiles)" on page 12977

(Continued)

[edit class-of-service adjustment-control-profiles {profile-name} application]	
"radius-coa (Adjustment Control Profiles)" on page 13999	

[application-traffic-control]

[edit class-of-service application-traffic-control]		
"rate-limiters" on page 14071	"rule-sets (CoS AppQoS)" on page 16149	

[class]

[edit class-of-service class {class-name} queue-num {queue-number}]	
"priority (Fabric Priority)" on page 13171	

[classifiers]

[edit class-of-service classifiers]		
"dscp (Class of Service)" on page 5988	"dscp-ipv6 (Class of Service)" on page 6007	"dscp-ipv6 (Class of Service)" on page 6009

(Continued)

[edit class-of-service classifiers]		
"exp (Class of Service)" on page 6639	"forwarding-class (BA Classifiers)" on page 7415	"ieee-802.1 (Class of Service)" on page 8344
"import (Classifiers)" on page 8466	"inet-precedence" on page 8578	"inet-precedence (CoS Classifiers)" on page 8581

[edit class-of-service classifiers (dscp | dscp-ipv6 | ieee-802.1 | exp) {classifier-name}]	
"import (CoS)" on page 8467	

[edit class-of-service classifiers (dscp | ieee-802.1 | inet-precedence | exp) {classifier-name} forwarding-class {class-name}]	
"loss-priority (Classifiers and Rewrite Rules)" on page 10146	

[edit class-of-service classifiers (dscp | ieee-802.1 | inet-precedence) forwarding-class {class-name} loss-priority {level}]	
"code-points (Classifiers)" on page 4403	

[edit class-of-service classifiers {classifier-type} {classifier-name} forwarding-class {class-name}]		
"loss-priority (Classifiers)" on page 10145	"loss-priority (Classifiers)" on page 10143	

[edit class-of-service classifiers {type} {classifier-name} forwarding-class {class-name} loss-priority {level}]	
"code-points (Classifiers)" on page 4405	

[code-point-aliases]

[edit class-of-service code-point-aliases]		
"dscp (Class of Service)" on page 5988	"dscp-ipv6 (Class of Service)" on page 6007	"dscp-ipv6 (Class of Service)" on page 6009
"exp (Class of Service)" on page 6639	"ieee-802.1 (Class of Service)" on page 8344	"inet-precedence" on page 8578

[congestion-notification-profile]

[edit class-of-service congestion-notification-profile {profile-name}]		
"ieee-802.1 (Congestion Notification)" on page 8350	"input (Congestion Notification)" on page 8693	"input (Congestion Notification)" on page 8692
"output (Congestion Notification)" on page 11984		

[edit class-of-service congestion-notification-profile {profile-name} input]		
"cable-length (Congestion Notification)" on page 3164	"dscp (Input Congestion Notification)" on page 6005	"ieee-802.1 (Input Congestion Notification)" on page 8357

[edit class-of-service congestion-notification-profile {profile-name} input dscp]	
"code-point (Input Congestion Notification)" on page 4397	

[edit class-of-service congestion-notification-profile {profile-name} input dscp code-point {code-point-bits}]		
"pfc (Input Congestion Notification)" on page 12456	"xon (Input Congestion Notification)" on page 31161	

[edit class-of-service congestion-notification-profile {profile-name} input dscp code-point {code-point-bits} pfc]	
"mru (Congestion Notification Profiles)" on page 10992	

[edit class-of-service congestion-notification-profile {profile-name} input ieee-802.1]		
"code-point (Congestion Notification)" on page 4393	"code-point (Input Congestion Notification)" on page 4397	

[edit class-of-service congestion-notification-profile {profile-name} input ieee-802.1 code-point {code-point-bits}]		
"pfc (Input Congestion Notification)" on page 12456	"xon (Input Congestion Notification)" on page 31161	

[edit class-of-service congestion-notification-profile {profile-name} input ieee-802.1 code-point {code-point-bits} pfc]	
"mru (Congestion Notification Profiles)" on page 10992	

[edit class-of-service congestion-notification-profile {profile-name} output]	
"ieee-802.1 (Output Congestion Notification)" on page 8358	

[edit class-of-service congestion-notification-profile {profile-name} output ieee-802.1]	
"code-point (Output Congestion Notification)" on page 4398	

[edit class-of-service congestion-notification-profile {profile-name} output ieee-802.1 code-point {code-point-bits}]	
"flow-control-queue (Output Congestion Notification)" on page 7272	

[drop-profiles]

[edit class-of-service drop-profiles {profile-name}]		
"drop-probability (Percentage)" on page 5958	"fill-level (Drop Profiles)" on page 7112	"interpolate (Drop Profiles)" on page 9197

[edit class-of-service drop-profiles {profile-name} fill-level {percentage}]	
"drop-probability (Fill Level)" on page 5955	

[edit class-of-service drop-profiles {profile-name} interpolate]		
"drop-probability" on page 5953	"drop-probability (Interpolated Value)" on page 5956	"fill-level" on page 7110
"fill-level (Interpolated Value)" on page 7113		

[fabric]

[edit class-of-service fabric]	
"scheduler-map (Fabric Queues)" on page 16353	

[edit class-of-service fabric scheduler-map]	
"priority (Fabric Queues, Schedulers)" on page 13172	

[edit class-of-service fabric scheduler-map priority (high | low)]	
"scheduler (Fabric Queues)" on page 16345	

[forwarding-class-sets]

[edit class-of-service forwarding-class-sets {forwarding-class-set-name}]	
"class (Forwarding Class Sets)" on page 3361	

[forwarding-classes]

[edit class-of-service forwarding-classes]		
"class (Forwarding classes)" on page 3362	"class (Forwarding Classes)" on page 3364	

[edit class-of-service forwarding-classes class {class-name}]		
"pfc-priority" on page 12458	"queue-num" on page 13947	

[forwarding-classes-interface-specific]

[edit class-of-service forwarding-classes-interface-specific]	
"output-forwarding-class-map" on page 11991	

[forwarding-policy]

[edit class-of-service forwarding-policy]		
"class (CoS-Based Forwarding)" on page 3346	"next-hop-map" on page 11346	

[edit class-of-service forwarding-policy class {class-name}]	
"classification-override" on page 3379	

[edit class-of-service forwarding-policy class {class-name} classification-override]	
"forwarding-class (Forwarding Policy)" on page 7426	

[edit class-of-service forwarding-policy next-hop-map {map-name}]		
"forwarding-class (Forwarding Policy)" on page 7426	"forwarding-class-default (Forwarding Policy)" on page 7418	

[edit class-of-service forwarding-policy next-hop-map {map-name} forwarding-class {class-name}]		
"discard (Forwarding Class)" on page 5819	"lsp-next-hop (CoS-Based Forwarding)" on page 10191	"next-hop (Class-Of-Service)" on page 11333
"non-lsp-next-hop" on page 11659		

[fragmentation-map]

[edit class-of-service fragmentation-map {map-name} forwarding-class {class-name}]	
"drop-timeout (Forwarding Class)" on page 5983	

[fragmentation-maps]

[edit class-of-service fragmentation-maps {map-name}]		
"forwarding-class (Fragmentation)" on page 7428	"forwarding-class (Services CoS Fragmentation Properties)" on page 7440	

[edit class-of-service fragmentation-maps forwarding-class {class-name}]		
"fragment-threshold (Forwarding Class Maps)" on page 7544	"multilink-class" on page 11104	"no-fragmentation" on page 11466

[host-outbound-traffic]

[edit class-of-service host-outbound-traffic]		
"dscp-code-point (CoS Host Outbound Traffic)" on page 5991	"forwarding-class (CoS Host Outbound Traffic)" on page 7417	"ieee-802.1 (Host Outbound Traffic)" on page 8356
"override-firewall" on page 12057	"protocol (Host Outbound Traffic)" on page 13443	

[edit class-of-service host-outbound-traffic ieee-802.1]		
"default (CoS Host Outbound Traffic)" on page 5189	"rewrite-rules (CoS Host Outbound Traffic)" on page 15743	

[interface]

[edit class-of-service interface {interface-name}]	
"queue-monitoring" on page 13946	

[interfaces]

[edit class-of-service interfaces {interface-name}]		
"classifiers (Interfaces)" on page 3391	"classifiers (Physical Interface)" on page 3393	"congestion-notification-profile" on page 4525
"exclude-queue-overhead-bytes" on page 6616	"forwarding-class (Fibre Channel Interfaces)" on page 7421	"forwarding-class-set" on page 7443
"input-excess-bandwidth-share" on page 8698	"input-scheduler-map" on page 8720	"input-shaping-rate (Physical Interface)" on page 8723
"input-traffic-control-profile" on page 8726	"input-traffic-control-profile-remaining" on page 8728	"interface-set (Hierarchical Schedulers)" on page 9045
"irb" on page 9384	"logical-interface-aggregate-statistics" on page 10062	"member-link-scheduler" on page 10614
"multi-destination (Class of Service)" on page 11047	"output-traffic-control-profile" on page 12016	"output-traffic-control-profile-remaining" on page 12019
"rewrite-rules (Interfaces)" on page 15751	"rewrite-value (Fibre Channel Interfaces)" on page 15753	"scheduler-map (Interfaces and Traffic-Control Profiles)" on page 16357
"shaping-rate (Applying to an Interface)" on page 16929	"shaping-rate (CoS Interfaces)" on page 16933	"slice (CoS Interfaces)" on page 28039
"traffic-class-map (Apply to Interface)" on page 30006	"unit (CoS Interfaces)" on page 30384	"unit (CoS Interfaces)" on page 30382

[edit class-of-service interfaces {interface-name} classifiers]		
"dscp (CoS Interfaces)" on page 5999	"ieee-802.1 (Classifier on Physical Interface)" on page 8348	"inet-precedence (Classifier on Physical Interface)" on page 8580

[edit class-of-service interfaces {interface-name} congestion-notification-profile {profile-name}]	
"input (Congestion Notification)" on page 8692	

[edit class-of-service interfaces {interface-name} congestion-notification-profile {profile-name} input ieee-802.1]	
"code-point (Congestion Notification)" on page 4393	

[edit class-of-service interfaces {interface-name} forwarding-class-set {forwarding-class-set-name}]	
"output-traffic-control-profile" on page 12016	

[edit class-of-service interfaces interface-name congestion-notification-profile {profile-name}]	
"ieee-802.1 (Congestion Notification)" on page 8350	

[edit class-of-service interfaces interface-set {interface-set-name}]		
"excess-bandwidth-share" on page 6577	"input-excess-bandwidth-share" on page 8698	"input-traffic-control-profile-remaining" on page 8728
"internal-node" on page 9189	"output-traffic-control-profile" on page 12016	"output-traffic-control-profile-remaining" on page 12019

[edit class-of-service interfaces {interface-name} rewrite-rules]		
"dscp (Rewrite Rules on Interfaces)" on page 6020	"ieee-802.1 (Rewrite Rules on Physical Interface)" on page 8362	"inet-precedence (Rewrite Rules on Physical Interface)" on page 8587

[edit class-of-service interfaces {fibre-channel-interface-name} rewrite-value]	
"input (Fibre Channel Interfaces)" on page 8700	

[edit class-of-service interfaces {fibre-channel-interface-name} rewrite-value input]	
"ieee-802.1 (Fibre Channel Interfaces)" on page 8354	

[edit class-of-service interfaces {fibre-channel-interface-name} rewrite-value input ieee-802.1]	
"code-point (Fibre Channel Interfaces)" on page 4395	

[edit class-of-service interfaces {interface-name} slice {slice-name}]	
"output-traffic-control-profile" on page 12016	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number}]		
"adaptive-shaper" on page 1835	"classifiers (Class of Service)" on page 3383	"classifiers (Interfaces)" on page 3391
"forwarding-class (Interfaces)" on page 7432	"fragmentation-map" on page 7547	"input-scheduler-map" on page 8720
"input-shaping-rate (Logical Interface)" on page 8721	"input-traffic-control-profile" on page 8726	"loss-priority-maps (Assigning to an Interface)" on page 10159
"loss-priority-maps (CoS Interfaces)" on page 10161	"loss-priority-rewrites (Assigning to an Interface)" on page 10167	"output-traffic-control-profile" on page 12016
"policy-map-marking" on page 12751	"rewrite-rules (Definition)" on page 15744	"rewrite-rules (Interfaces)" on page 15749

(Continued)

[edit class-of-service interfaces {interface-name} unit {logical-unit-number}]		
"rewrite-rules (Interfaces)" on page 15751	"scheduler-map (Interfaces and Traffic-Control Profiles)" on page 16357	"shaping-rate (Applying to an Interface)" on page 16929
"shaping-rate (CoS Interfaces)" on page 16933	"virtual-channel-group (CoS Interfaces)" on page 30740	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} classifiers]		
"dscp (Class of Service)" on page 5988	"dscp-ipv6 (Class of Service)" on page 6007	"dscp-ipv6 (Class of Service)" on page 6009
"ieee-802.1 (Class of Service)" on page 8344	"inet-precedence" on page 8578	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} classifiers ieee802.1 vlan-tag]	
"transparent" on page 30110	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} input-traffic-control-profile]	
"shared-instance" on page 16967	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} loss-priority-maps]	
"frame-relay-de (CoS Interfaces)" on page 7557	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} output-traffic-control-profile]	
"shared-instance" on page 16967	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} rewrite-rules]		
"dscp (Class of Service)" on page 5988	"dscp (Rewrite Rules on Interfaces)" on page 6020	"dscp-ipv6 (Class of Service)" on page 6007
"dscp-ipv6 (Class of Service)" on page 6009	"exp (Class of Service)" on page 6639	"exp (Rewrite rules)" on page 6643
"exp-push-push-push" on page 6642	"exp-swap-push-push" on page 6646	"frame-relay-de (CoS Interfaces)" on page 7557
"ieee-802.1 (Class of Service)" on page 8344	"ieee-802.1 (Rewrite Rules on Logical Interface)" on page 8360	"ieee-802.1ad" on page 8363
"inet-precedence (CoS Rewrite Rules)" on page 8582	"inet6-precedence (CoS Rewrite Rules)" on page 8614	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} rewrite-rules dscp {rewrite-name}]	
"protocol (Rewrite Rules)" on page 13457	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} rewrite-rules dscp-ipv6 {rewrite-name}]	
"protocol (Rewrite Rules)" on page 13457	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} rewrite-rules exp {rewrite-name}]	
"protocol (Rewrite Rules)" on page 13457	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} rewrite-rules ieee-802.1 ({rewrite-name} | default)]	
"vlan-tag" on page 30857	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} rewrite-rules inet-precedence {rewrite-name}]	
"protocol (Rewrite Rules)" on page 13457	

[edit class-of-service interfaces {interface-name} unit {logical-unit-number} rewrite-rules inet6-precedence {rewrite-name}]	
"protocol (Rewrite Rules)" on page 13457	

[edit class-of-service interfaces {interface-type}-{fpc}/{pic}/*]	
"scheduler-map-chassis" on page 16348	

[loss-priority-maps]

[edit class-of-service loss-priority-maps]		
"frame-relay-de (CoS Loss Priority)" on page 7558	"frame-relay-de (Defining Loss Priority Maps)" on page 7561	

[edit class-of-service loss-priority-maps frame-relay-de {map-name}]	
"loss-priority (Frame Relay Loss Priority)" on page 10148	

[loss-priority-rewrites]

[edit class-of-service loss-priority-rewrites]	
"frame-relay-de (Defining Loss Priority Rewrites)" on page 7563	

[multi-destination]

[edit class-of-service multi-destination]		
"classifiers (Class of Service)" on page 3383	"family (Multi-destination)" on page 6890	"scheduler-map (Interfaces and Traffic-Control Profiles)" on page 16357

[edit class-of-service multi-destination classifiers]		
"dscp (Class of Service)" on page 5988	"ieee-802.1 (Class of Service)" on page 8344	

[edit class-of-service multi-destination family]		
"ethernet (CoS for Multidestination Traffic)" on page 6485	"inet (CoS)" on page 8560	"inet6 (CoS Multidestination)" on page 8601

[edit class-of-service multi-destination family ethernet]	
"broadcast (Forwarding class)" on page 3044	

[queue]

[edit class-of-service queue {queue-number} {class-name}]	
"priority (Fabric Priority)" on page 13171	

[restricted-queues]

[edit class-of-service restricted-queues]	
"forwarding-class (Restricted Queues)" on page 7436	

[edit class-of-service restricted-queues forwarding-class {class-name}]	
"queue (Restricted Queues)" on page 13949	

[rewrite-rules]

[edit class-of-service rewrite-rules]		
"dscp (Class of Service)" on page 5988	"dscp-ipv6 (Class of Service)" on page 6007	"dscp-ipv6 (Class of Service)" on page 6009
"exp (Class of Service)" on page 6639	"frame-relay-de (CoS Rewrite Rule)" on page 7560	"ieee-802.1 (Class of Service)" on page 8344

(Continued)

[edit class-of-service rewrite-rules]		
"import (Rewrite Rules)" on page 8486	"inet-precedence" on page 8578	

[edit class-of-service rewrite-rules (dscp | dscp-ipv6 | ieee-802.1 | exp) {classifier-name}]	
"import (CoS)" on page 8467	

[edit class-of-service rewrite-rules (dscp | ieee-802.1 | inet-precedence | exp) {rewrite-rule-name} forwarding-class {class-name}]	
"loss-priority (Classifiers and Rewrite Rules)" on page 10146	

[edit class-of-service rewrite-rules (dscp | ieee-802.1) forwarding-class {class-name} loss-priority {level}]	
"code-point (Rewrite Rules)" on page 4400	

[edit class-of-service rewrite-rules {type} {rewrite-name} forwarding-class {class-name}]		
"code-point" on page 4390	"loss-priority (Rewrite Rules)" on page 10163	"loss-priority (Rewrite Rules)" on page 10164

[routing-instances]

[edit class-of-service routing-instances {routing-instance-name}]		
"classifiers (Routing Instance)" on page 3395	"classifiers (VRF)" on page 3382	

[scheduler-map-forwarding-class-sets]

[edit class-of-service scheduler-map-forwarding-class-sets {scheduler-map-fcset-name}]	
"forwarding-class-set (Interconnect Device)" on page 7444	

[edit class-of-service scheduler-map-forwarding-class-sets {scheduler-map-fcset-name} forwarding-class-set {fabric-forwarding-class-set-name}]	
"scheduler (Interconnect Device)" on page 16347	

[scheduler-maps]

[edit class-of-service scheduler-maps {map-name}]	
"scheduler (Scheduler Map)" on page 16362	

[edit class-of-service scheduler-maps {map-name} forwarding-class {class-name}]	
"scheduler" on page 16343	

[schedulers]

[edit class-of-service schedulers {scheduler-name}]		
"adjust-minimum" on page 1947	"adjust-percent" on page 1950	"buffer-dynamic-threshold" on page 3047

(Continued)

[edit class-of-service schedulers {scheduler-name}]		
"buffer-rate (Schedulers)" on page 3065	"buffer-size" on page 3055	"buffer-size" on page 3055
"buffer-size (Schedulers)" on page 3063	"drop-profile-map (Schedulers)" on page 5964	"drop-profile-map (Schedulers)" on page 5965
"excess-priority" on page 6582	"excess-rate (Schedulers)" on page 6601	"excess-rate (Schedulers)" on page 6585
"explicit-congestion-notification" on page 6649	"priority (Schedulers)" on page 13206	"protocol (Drop Profiles)" on page 13438
"queue-monitoring" on page 13946	"shaping-rate (Schedulers)" on page 16960	"transmit-rate" on page 30098
"transmit-rate (EX Series Switches)" on page 30105	"transmit-rate (Schedulers)" on page 30107	

[edit class-of-service schedulers {scheduler-name} drop-profile-map]		
"loss-priority (Drop Profiles)" on page 10150	"loss-priority (Scheduler Drop Profiles)" on page 10169	"protocol (Schedulers)" on page 13459

[edit class-of-service schedulers {scheduler-name} drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any | non-tcp | tcp)]	
"drop-profile (Schedulers)" on page 5966	

[edit class-of-service schedulers {scheduler-name} drop-profile-map loss-priority (low | medium-high | high)]	
"protocol (Drop Profile Map)" on page 13437	

[edit class-of-service schedulers {scheduler-name} drop-profile-map loss-priority (low | medium-high | high) protocol {protocol}]	
"drop-profile (Schedulers)" on page 5959	

[shared-buffer]

[edit class-of-service shared-buffer]		
"egress (Buffer Configuration)" on page 6208	"ingress (Buffer Configuration)" on page 8626	

[edit class-of-service shared-buffer egress]	
"buffer-partition (Egress)" on page 3049	

[edit class-of-service shared-buffer ingress]	
"buffer-partition (Ingress)" on page 3052	

[system-defaults]

[edit class-of-service system-defaults]	
"classifiers (Class of Service)" on page 3383	

[edit class-of-service system-defaults classifiers]	
"exp (Class of Service)" on page 6639	

[traffic-control-profile]

[edit class-of-service traffic-control-profile {lcp-name}]	
"queue-monitoring" on page 13946	

[traffic-control-profiles]

[edit class-of-service traffic-control-profiles {traffic-control-profile-name}]		
"adjust-minimum" on page 1947	"delay-buffer-rate" on page 5244	"excess-rate (Schedulers)" on page 6585
"excess-rate-high" on page 6591	"excess-rate-low" on page 6594	"excess-rate-medium-high" on page 6597
"excess-rate-medium-low" on page 6599	"guaranteed-rate (Traffic Control Profiles)" on page 7940	"overhead-accounting" on page 12036
"peak-rate" on page 12338	"scheduler-map (Interfaces and Traffic-Control Profiles)" on page 16357	"shaping-rate (Oversubscribing an Interface)" on page 16947
"shaping-rate (Traffic Control Profiles)" on page 16926	"shaping-rate-excess-high" on page 16939	"shaping-rate-excess-low" on page 16941

(Continued)

[edit class-of-service traffic-control-profiles {traffic-control-profile-name}]		
"shaping-rate-excess-medium-high" on page 16943	"shaping-rate-excess-medium-low" on page 16945	"shaping-rate-priority-high" on page 16950
"shaping-rate-priority-low" on page 16952	"shaping-rate-priority-medium" on page 16954	"shaping-rate-priority-medium-low" on page 16956
"shaping-rate-priority-strict-high" on page 16958	"strict-priority-scheduler" on page 28650	

[edit class-of-service traffic-control-profiles {profile-name} overhead-accounting]	
"bytes (Dynamic Traffic Shaping)" on page 3121	

[edit class-of-service traffic-control-profiles {profile-name} overhead-accounting],]		
"cell-mode (Dynamic Traffic Shaping)" on page 3245	"frame-mode (Dynamic Traffic Shaping)" on page 7552	

[virtual-channel-groups]

[edit class-of-service virtual-channel-groups {group-name} {virtual-channel-name}]		
"default (CoS)" on page 5187	"scheduler-map (CoS Virtual Channels)" on page 16350	"shaping-rate (CoS Virtual Channels)" on page 16936

[edit diameter]

IN THIS SECTION

- [network-element] | 132
- [network-statement] | 132
- [peer] | 133
- [transport] | 133

Sub Hierarchies

"[edit diameter network-element {element-name}]" on page 132

"[edit diameter network-statement]" on page 132

"[edit diameter peer {peer-name}]" on page 133

"[edit diameter transport {transport-name}]" on page 133

[edit diameter]		
"network-element (Diameter Base Protocol)" on page 11307	"origin (Diameter Base Protocol)" on page 11935	"peer (Diameter Base Protocol)" on page 12353
"transport (Diameter Base Protocol)" on page 30123		

[network-element]

[edit diameter network-element {element-name}]		
"dne-origin (Diameter Network Element)" on page 5846	"forwarding (Diameter Network Element)" on page 7456	"function (Diameter Network Element)" on page 7634
"peer (Diameter Network Element)" on page 12356		

[edit diameter network-element {element-name} forwarding]	
"route (Diameter Network Element)" on page 15854	

[edit diameter network-element {element-name} forwarding route {dne-route-name}]		
"destination (Diameter Network Element)" on page 5414	"function (Diameter Route)" on page 7636	

[edit diameter network-element {element-name} peer {peer-name}]	
"priority (Diameter Peer)" on page 13168	

[network-statement]

[edit diameter network-statement {element-name} forwarding route {dne-route-name}]	
"metric (Diameter Route)" on page 10668	

[peer]

[edit diameter peer {peer-name}]		
"address (Diameter Peer)" on page 1874	"logical-system (Diameter Peer)" on page 10076	"peer-origin (Diameter Peer)" on page 12377
"routing-instance (Diameter Peer)" on page 15931	"send-origin-state-id (Diameter Applications)" on page 16545	"transport (Diameter Peer)" on page 30124

[edit diameter peer {peer-name} connect-actively]	
"port (Diameter Peer)" on page 12819	

[transport]

[edit diameter transport {transport-name}]		
"address (Diameter Transport)" on page 1875	"logical-system (Diameter Transport)" on page 10078	"routing-instance (Diameter Transport)" on page 15932

[edit dynamic]

IN THIS SECTION

- [policy-options] | **134**

Sub Hierarchies

"[edit dynamic policy-options]" on page 134

[policy-options]

[edit dynamic policy-options]		
"as-path (Policy Options)" on page 2493	"as-path-group" on page 2492	"community (Policy Options)" on page 4465
"condition" on page 4502	"prefix-list" on page 13090	

[edit dynamic policy-options condition]	
"table" on page 28895	

[edit dynamic-profiles]

IN THIS SECTION

- [access-cac] | **136**
- [class-of-service] | **137**
- [firewall] | **140**
- [interfaces] | **144**
- [logical-systems] | **160**
- [multicast] | **163**
- [name] | **163**
- [policy-options] | **163**

- [protocols] | **163**
- [routing-instances] | **167**
- [routing-options] | **170**
- [services] | **171**
- [telemetry] | **172**
- [variables] | **172**

Sub Hierarchies

"[edit dynamic-profiles access-cac]" on page 136

"[edit dynamic-profiles {profile-name} class-of-service]" on page 137

"[edit dynamic-profiles {profile-name} firewall]" on page 140

"[edit dynamic-profiles {profile-name} interfaces]" on page 144

"[edit dynamic-profiles logical-systems]" on page 160

"[edit dynamic-profiles {name} multicast]" on page 163

"[edit dynamic-profiles name]" on page 163

"[edit dynamic-profiles {profile-name} policy-options]" on page 163

"[edit dynamic-profiles {profile-name} protocols]" on page 163

"[edit dynamic-profiles routing-instances]" on page 167

"[edit dynamic-profiles {profile-name} routing-options]" on page 170

"[edit dynamic-profiles services]" on page 171

"[edit dynamic-profiles {name} telemetry]" on page 172

"[edit dynamic-profiles {profile-name} variables {variable-name}]" on page 172

[edit dynamic-profiles {profile-name}]		
"class-of-service (Dynamic Profiles)" on page 3366	"filter-specific" on page 7155	"firewall (Dynamic Firewalls)" on page 7191
"if-exceeding (Policer)" on page 8373	"if-exceeding-pps (Policer)" on page 8376	"interface (Dynamic Routing Instances)" on page 8879
"interfaces (Static and Dynamic Subscribers)" on page 9170	"policy-options" on page 12697	"policy-options (Dynamic Profiles)" on page 12702
"predefined-variable-defaults (Dynamic Client Profiles)" on page 13016	"profile-type (Dynamic Service Profiles)" on page 13353	"profile-variable-set (Dynamic Profiles)" on page 13354
"protocols (Dynamic Profiles)" on page 13791	"routing-instances (Dynamic Profiles)" on page 15973	"routing-options (Dynamic Profiles)" on page 15981
"services (Captive Portal Content Delivery)" on page 16769	"telemetry" on page 29051	"variables (Dynamic Service Profiles)" on page 30654
"version-alias (Dynamic Profiles)" on page 30694		

[access-cac]

[edit dynamic-profiles {dynamic-profile-name} access-cac interface {$junos-interface-name}]	
"qos-adjust-hierarchical" on page 13896	

[class-of-service]

[edit dynamic-profiles {profile-name} class-of-service]		
"dynamic-class-of-service-options (Dynamic Traffic Shaping)" on page 6102	"interfaces (Dynamic CoS Definition)" on page 9133	"monitoring-profile" on page 10959
"non-strict-priority-scheduling" on page 11664	"scheduler-maps (Dynamic CoS Definition)" on page 16360	"schedulers (Dynamic CoS Definition)" on page 16371
"system-defaults" on page 28842	"traffic-control-profiles (Dynamic CoS Definition)" on page 30012	

[edit dynamic-profiles {name} class-of-service host-outbound-traffic]		
"override-firewall" on page 12057	"protocol (Host Outbound Traffic)" on page 13443	

[edit dynamic-profiles {name} class-of-service interfaces]		
"slice (CoS Interfaces)" on page 28039	"unit (Dynamic Traffic Shaping)" on page 30397	

[edit dynamic-profiles {profile-name} class-of-service interfaces {interface-name} unit {logical-unit-number}]		
"classifiers (Dynamic CoS Application)" on page 3390	"output-traffic-control-profile (Dynamic CoS Definition)" on page 12018	"report-ingress-shaping-rate (Dynamic CoS Interfaces)" on page 14390
"rewrite-rules (Dynamic CoS Interfaces)" on page 15748		

[edit dynamic-profiles {profile-name} class-of-service interfaces {interface-name} unit {logical-unit-number} classifiers]		
"dscp (Dynamic Classifiers)" on page 6000	"dscp-ipv6 (Dynamic Classifiers)" on page 6013	"ieee-802.1 (Dynamic Classifiers)" on page 8351
"inet-precedence (Dynamic Classifiers)" on page 8584		

[edit dynamic-profiles {profile-name} class-of-service interfaces {interface-name} unit {logical-unit-number} classifiers ieee-802.1]	
"vlan-tag (Dynamic Classifiers)" on page 30858	

[edit dynamic-profiles {profile-name} class-of-service interfaces {interface-name} unit {logical-unit-number} rewrite-rules]		
"dscp (Dynamic Rewrite Rules)" on page 6002	"dscp-ipv6 (Dynamic Rewrite Rules)" on page 6014	"ieee-802.1 (Dynamic Rewrite Rules)" on page 8352
"inet-precedence (Dynamic Rewrite Rules)" on page 8585		

[edit dynamic-profiles {profile-name} class-of-service interfaces {interface-name} unit {logical-unit-number} rewrite-rules ieee-802.1]	
"vlan-tag (Dynamic Rewrite Rules)" on page 30860	

[edit dynamic-profiles {profile-name} class-of-service scheduler-maps {map-name}]	
"forwarding-class (Dynamic Scheduler Maps)" on page 7420	

[edit dynamic-profiles {profile-name} class-of-service scheduler-maps {map-name} forwarding-class {class-name}]		
"scheduler (Dynamic Scheduler Maps)" on page 16344		

[edit dynamic-profiles {profile-name} class-of-service schedulers {scheduler-name}]		
"adjust-minimum (Dynamic Shaping and Scheduling)" on page 1949	"adjust-percent (Dynamic Schedulers)" on page 1952	"buffer-size (Dynamic Scheduling)" on page 3061
"drop-profile-map (Dynamic Schedulers)" on page 5962	"excess-priority (Dynamic Schedulers)" on page 6584	"excess-rate (Dynamic Schedulers)" on page 6588
"priority (Dynamic Schedulers)" on page 13169	"shaping-rate (Dynamic Traffic Shaping and Scheduling)" on page 16937	"transmit-rate (Dynamic Schedulers)" on page 30103

[edit dynamic-profiles {profile-name} class-of-service schedulers {scheduler-name} drop-profile-map]		
"loss-priority (Dynamic Schedulers)" on page 10151	"protocol (Dynamic Schedulers)" on page 13439	

[edit dynamic-profiles {profile-name} class-of-service schedulers {scheduler-name} drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any | non-tcp | tcp)]		
"drop-profile (Dynamic Schedulers)" on page 5960		

[edit dynamic-profiles {profile-name} class-of-service traffic-control-profiles {traffic-control-profile-name}]		
"adjust-minimum (Dynamic Shaping and Scheduling)" on page 1949	"delay-buffer-rate (Dynamic Traffic Shaping)" on page 5246	"excess-rate (Dynamic Traffic Shaping)" on page 6589
"excess-rate-high (Dynamic Traffic Shaping)" on page 6593	"excess-rate-low (Dynamic Traffic Shaping)" on page 6596	"guaranteed-rate (Dynamic Traffic Shaping)" on page 7942

(Continued)

[edit dynamic-profiles {profile-name} class-of-service traffic-control-profiles {traffic-control-profile-name}]		
"overhead-accounting (Dynamic Traffic Shaping)" on page 12039	"scheduler-map (Dynamic Traffic Shaping)" on page 16351	"shaping-rate (Dynamic Traffic Shaping and Scheduling)" on page 16937
"shaping-rate-excess-high" on page 16939	"shaping-rate-excess-low" on page 16941	"shaping-rate-excess-medium-high" on page 16943
"shaping-rate-excess-medium-low" on page 16945	"shaping-rate-priority-high" on page 16950	"shaping-rate-priority-low" on page 16952
"shaping-rate-priority-medium" on page 16954	"shaping-rate-priority-medium-low" on page 16956	"shaping-rate-priority-strict-high" on page 16958

[edit dynamic-profiles {profile-name} class-of-service traffic-control-profiles {profile-name} overhead-accounting]		
"bytes (Dynamic Traffic Shaping)" on page 3121	"cell-mode (Dynamic Traffic Shaping)" on page 3245	"frame-mode (Dynamic Traffic Shaping)" on page 7552

[firewall]

[edit dynamic-profiles {profile-name} firewall]		
"family (Dynamic Firewalls)" on page 6855	"hierarchical-policer (aggregate or premium)" on page 8057	"policer (Configuring)" on page 12630
"three-color-policer (Configuring)" on page 29314		

[edit dynamic-profiles {profile-name} firewall family {family}]		
"fast-update-filter (Dynamic Firewalls)" on page 6946	"filter (Dynamic Profiles Filter Creation)" on page 7134	

[edit dynamic-profiles {profile-name} firewall family {family} fast-update-filter {filter-name}]		
"interface-specific (Dynamic Firewalls)" on page 9061	"match-order (Dynamic Firewalls)" on page 10389	"term (Dynamic Profiles)" on page 29115

[edit dynamic-profiles {profile-name} firewall family {family-name} filter {filter-name}]		
"enhanced-mode" on page 6396	"enhanced-mode-override" on page 6401	"interface-shared" on page 9048
"interface-specific (Dynamic Firewalls)" on page 9061	"interface-specific (Firewall Filters)" on page 9063	"term (Dynamic Profiles)" on page 29115
"term (Firewall Filter)" on page 29117		

[edit dynamic-profiles {name} firewall family inet filter {name} term {name} from]	
"loss-priority (Firewall Filter)" on page 10152	

[edit dynamic-profiles {name} firewall filter {name} term {name} from]	
"loss-priority (Firewall Filter)" on page 10152	

[edit dynamic-profiles {profile-name} firewall hierarchical-policer {name}]		
"aggregate (Hierarchical Policer)" on page 2064	"premium (Hierarchical Policer)" on page 13110	

[edit dynamic-profiles {profile-name} firewall hierarchical-policer aggregate]		
"if-exceeding (Hierarchical Policer)" on page 8371	"if-exceeding-pps (Hierarchical Policer)" on page 8374	

[edit dynamic-profiles {profile-name} firewall hierarchical-policer aggregate if-exceeding]		
"bandwidth-limit (Hierarchical Policer)" on page 2855	"burst-size-limit (Hierarchical Policer)" on page 3084	"pps-limit (Hierarchical Policer)" on page 12982

[edit dynamic-profiles {profile-name} firewall hierarchical-policer {hierarchical-policer-name} aggregate if-exceeding-pps]	
"packet-burst (Hierarchical Policer)" on page 12133	

[edit dynamic-profiles {profile-name} firewall hierarchical-policer premium]		
"if-exceeding (Hierarchical Policer)" on page 8371	"if-exceeding-pps (Hierarchical Policer)" on page 8374	

[edit dynamic-profiles {profile-name} firewall hierarchical-policer premium if-exceeding]		
"bandwidth-limit (Hierarchical Policer)" on page 2855	"burst-size-limit (Hierarchical Policer)" on page 3084	"pps-limit (Hierarchical Policer)" on page 12982

[edit dynamic-profiles {profile-name} firewall hierarchical-policer {hierarchical-policer-name} premium if-exceeding-pps]	
"packet-burst (Hierarchical Policer)" on page 12133	

[edit dynamic-profiles {profile-name} firewall policer {policer-name}]		
"logical-bandwidth-policer" on page 10058	"logical-interface-policer" on page 10069	"physical-interface-policer" on page 12479

[edit dynamic-profiles {profile-name} firewall policer {policer-name} if-exceeding]		
"bandwidth-limit (Policer)" on page 2857	"bandwidth-percent" on page 2866	"burst-size-limit (Policer)" on page 3086

[edit dynamic-profiles {profile-name} firewall policer {policer-name} if-exceeding-pps]		
"packet-burst (Policer)" on page 12135	"pps-limit (Policer)" on page 12984	

[edit dynamic-profiles {profile-name} firewall three-color-policer {name}]		
"action" on page 1759	"logical-interface-policer" on page 10069	"single-rate" on page 27969
"two-rate" on page 30259		

[edit dynamic-profiles {profile-name} firewall three-color-policer {name} action]	
"loss-priority high then discard (Three-Color Policer)" on page 10155	

[edit dynamic-profiles {profile-name} firewall three-color-policer {name} single-rate]		
"color-blind" on page 4429	"committed-burst-size" on page 4458	"committed-information-rate" on page 4460
"excess-burst-size" on page 6580		

[edit dynamic-profiles {profile-name} firewall three-color-policer {name} two-rate]		
"color-blind" on page 4429	"committed-burst-size" on page 4458	"committed-information-rate" on page 4460
"peak-burst-size" on page 12332	"peak-information-rate" on page 12336	

[interfaces]

[edit dynamic-profiles {profile-name} interfaces]		
"demux0 (Dynamic Interface)" on page 5311	"hierarchical-scheduler" on page 8060	"interface-mib (Enhanced Subscriber Management)" on page 8916
"interface-set (Dynamic Profiles)" on page 9037	"interface-set (Dynamic VLAN Interface Sets Association)" on page 9040	"interface-set (Dynamic VLAN Interface Sets Definition)" on page 9042
"logical-tunnel-options" on page 10093	"pp0 (Dynamic PPPoE)" on page 12925	"stacked-interface-set (Dynamic Profiles)" on page 28469
"traps" on page 30142	"unit (Dynamic Profiles Standard Interface)" on page 30392	"vlan-tagging (Dynamic)" on page 30864

[edit dynamic-profiles {profile-name} interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]	
"keepalives (Dynamic Profiles)" on page 9478	

[edit dynamic-profiles {profile-name} interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit” ppp-options]		
"chap (Dynamic PPP)" on page 3291	"ignore-magic-number-mismatch (Dynamic Profiles)" on page 8405	"local-name" on page 9964
"mru (Dynamic and Static PPPoE)" on page 10994	"mtu (Dynamic and Static PPPoE)" on page 11034	"pap (Dynamic PPP)" on page 12171

[edit dynamic-profiles {profile-name} interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit” ppp-options chap]	
"challenge-length (Static and Dynamic PPP)" on page 3281	

[edit dynamic-profiles {profile-name} interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]. [edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit”]	
"ppp-options (Dynamic PPP)" on page 12951	

[edit dynamic-profiles {profile-name} interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]. [edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit” ppp-options]	
"on-demand-ip-address" on page 11815	

[edit dynamic-profiles {profile-name} interfaces $junos-interface-ifd-name unit $junos-interface-unit]		
"actual-transit-statistics (Dynamic Profiles)" on page 1823	"advisory-options (Traffic Shaping)" on page 2035	

[edit dynamic-profiles {profile-name} interfaces $junos-interface-ifd-name unit $junos-interface-unit advisory-options]		
"downstream-rate (Traffic Shaping)" on page 5935	"upstream-rate (Traffic Shaping)" on page 30479	

[edit dynamic-profiles {name} interfaces $junos-interface-ifd-name unit $junos-interface-unit ppp-options]	
"local-authentication (Dynamic PPP Options)" on page 9919	

[edit dynamic-profiles {name} interfaces $junos-interface-ifd-name unit $junos-interface-unit ppp-options local-authentication]	
"username-include (Local Authentication)" on page 30602	

[edit dynamic-profiles {name} interfaces {name} aggregated-ether-options]		
"bfd-liveness-detection (LAG)" on page 2926	"load-balance" on page 9874	"optimize-aggregate-frr" on page 11853
"targeted-options (Manual Targeting)" on page 28966		

[edit dynamic-profiles {name} interfaces {name} aggregated-ether-options bfd-liveness-detection]	
"damping (BFD Liveness Detection)" on page 4815	

[edit dynamic-profiles {name} interfaces {name} aggregated-ether-options lacp]	
"sync-reset" on page 28785	

[edit dynamic-profiles {name} interfaces {name} aggregated-ether-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[edit dynamic-profiles {profile-name} interfaces demux0]	
"unit (Dynamic Demux Interface)" on page 30385	

[edit dynamic-profiles {profile-name} interfaces demux0 {interface-name} unit {logical-unit-number}]		
"demux-options (Dynamic Interface)" on page 5302	"family (Dynamic Demux Interface)" on page 6853	"targeted-distribution (Dynamic Demux Interfaces over Aggregated Ethernet)" on page 28959

[edit dynamic-profiles {profile-name} interfaces demux0 unit {logical-unit-number} family {family}]		
"demux-source (Dynamic IP Demux Interface)" on page 5306	"filter (Dynamic Profiles Filter Attachment)" on page 7132	"preferred-source-address" on page 13060
"unnumbered-address (Dynamic Profiles)" on page 30422		

[edit dynamic-profiles {profile-name} interfaces demux0 unit {logical-unit-number} family {family} filter input {filter-name}]	
"precedence" on page 13008	

[edit dynamic-profiles {profile-name} interfaces demux0 unit {logical-unit-number} family {family} filter output {filter-name}]	
"precedence" on page 13008	

[edit dynamic-profiles {profile-name} interfaces demux0 unit {logical-unit-number} family inet]	
"mac-validate (Dynamic IP Demux Interface)" on page 10275	

[edit dynamic-profiles {profile-name} interfaces demux0 unit {logical-unit-number} family pppoe]		
"access-concentrator" on page 1655	"direct-connect" on page 5708	"duplicate-protection (Dynamic PPPoE)" on page 6071
"dynamic-profile (Dynamic PPPoE)" on page 6124	"max-sessions (Dynamic PPPoE)" on page 10469	"max-sessions-vsa-ignore (Static and Dynamic Subscribers)" on page 10481
"service-name-table" on page 16716	"short-cycle-protection (Static and Dynamic Subscribers)" on page 16984	

[edit dynamic-profiles {profile-name} interfaces demux0 unit {logical-unit-number} family {family} rpf-check]	
"fail-filter (Dynamic Profiles)" on page 6777	

[edit dynamic-profiles {profile-name} interfaces demux0 {interface-name} unit unit {logical-unit-number} demux-options]	
"underlying-interface (demux0)" on page 30333	

[edit dynamic-profiles {profile-name} interfaces demux0 unit {logical-unit-number}],]	
"filter (Dynamic Interface Unit)" on page 7130	

[edit dynamic-profiles {name} interfaces {name} esi]	
"df-election-type" on page 5576	

[edit dynamic-profiles {name} interfaces {name} ether-options]		
"fec (ether)" on page 6992		

[edit dynamic-profiles {name} interfaces {name} gigether-options]		
"loopback-remote" on page 10128	"mru (Gigabit Ethernet Options)" on page 10991	

[edit dynamic-profiles {name} interfaces interface-range {name} aggregated-ether-options]		
"load-balance" on page 9874		

[edit dynamic-profiles {name} interfaces interface-range {name} aggregated-ether-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[edit dynamic-profiles {name} interfaces interface-range {name} logical-tunnel-options]		
"load-balance" on page 9874		

[edit dynamic-profiles {name} interfaces interface-range {name} logical-tunnel-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[edit dynamic-profiles {profile-name} interfaces interface-set {interface-set-name}]		
"interface (Dynamic Interface Sets)" on page 8873		

[edit dynamic-profiles {profile-name} interfaces interface-set $junos-interface-set-name interface $junos-interface-ifd-name]		
"advisory-options (Traffic Shaping)" on page 2035		

[edit dynamic-profiles {profile-name} interfaces interface-set $junos-interface-set-name interface $junos-interface-ifd-name advisory-options]		
"downstream-rate (Traffic Shaping)" on page 5935	"upstream-rate (Traffic Shaping)" on page 30479	

[edit dynamic-profiles {profile-name} interfaces interface-set {interface-set-name} interface {interface-name}]		
"unit (Dynamic Interface Sets)" on page 30388	"unit (Dynamic Traffic Shaping)" on page 30397	

[edit dynamic-profiles {profile-name} interfaces interface-set {interface-set-name} pppoe-underlying-options]	
"max-sessions (Dynamic PPPoE)" on page 10469	

[edit dynamic-profiles {profile-name} interfaces interface-set “$junos-interface-set-name”]	
"pppoe-underlying-options (Dynamic VLAN Interface Sets)" on page 12979	

[edit dynamic-profiles {name} interfaces {name} logical-tunnel-options]	
"load-balance" on page 9874	

[edit dynamic-profiles {name} interfaces {name} logical-tunnel-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[edit dynamic-profiles {name} interfaces {name} optics-options]		
"alarm (optics-options)" on page 2110	"is-ma" on page 9396	"is-ma" on page 9396

[edit dynamic-profiles {profile-name} interfaces pp0]	
"unit (Dynamic PPPoE)" on page 30390	

[edit dynamic-profiles {profile-name} interfaces pp0 unit {logical-unit-number}]	
"keepalives (Dynamic Profiles)" on page 9478	

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit”]		
"family (Dynamic PPPoE)" on page 6857	"keepalives (Dynamic Profiles)" on page 9478	"no-keepalives (Dynamic Profiles)" on page 11490
"pppoe-options (Dynamic PPPoE)" on page 12976		

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit” family inet]	
"unnumbered-address (Dynamic PPPoE)" on page 30421	

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit” family {family} service input]	
"service-set (Dynamic Service Sets)" on page 16742	

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit” family {family} service input service-set {service-set-name}]	
"service-filter (Dynamic Service Sets)" on page 16697	

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit” family {family} service output]	
"service-set (Dynamic Service Sets)" on page 16742	

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit” family {family} service output service-set {service-set-name}]	
"service-filter (Dynamic Service Sets)" on page 16697	

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit” ppp-options]		
"aaa-options (PPP Profile)" on page 1614	"authentication (Static and Dynamic PPP)" on page 2664	"ignore-magic-number-mismatch (Dynamic Profiles)" on page 8405
"initiate-ncp (Dynamic and Static PPP)" on page 8662	"ipcp-suggest-dns-option" on page 9289	"local-name" on page 9964
"mru (Dynamic and Static PPPoE)" on page 10994	"mtu (Dynamic and Static PPPoE)" on page 11034	"pap (Dynamic PPP)" on page 12171
"peer-ip-address-optional" on page 12373		

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit” ppp-options chap]	
"challenge-length (Static and Dynamic PPP)" on page 3281	

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit” ppp-options],]		
"chap (Dynamic PPP)" on page 3291		

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos-interface-unit” pppoe-options]		
"server (Dynamic PPPoE)" on page 16627	"underlying-interface (Dynamic PPPoE)" on page 30335	

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos–interface–unit” family {family}]		
"filter (Dynamic Profiles Filter Attachment)" on page 7132	"service (Dynamic Service Sets)" on page 16695	

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos–interface–unit” family {family} filter input {filter-name}]		
"precedence" on page 13008		

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos–interface–unit” family {family} filter output {filter-name}]		
"precedence" on page 13008		

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos–interface–unit” family {family} service]		
"input (Dynamic Service Sets)" on page 8697	"output (Dynamic Service Sets)" on page 11986	

[edit dynamic-profiles {profile-name} interfaces pp0 unit “$junos–interface–unit” family {family} service input]		
"post-service-filter (Dynamic Service Sets)" on page 12908		

[edit dynamic-profiles interfaces {interface-name} unit {logical-unit-number}]		
"demux-source (Dynamic Underlying Interface)" on page 5308	"dial-options (Dynamic Profiles)" on page 5685	"encapsulation (Dynamic Interfaces)" on page 6297
"family (Dynamic Standard Interface)" on page 6859	"family vpls (Layer 2 Pseudowires)" on page 6920	"filter (Dynamic Interface Unit)" on page 7130
"host-prefix-only" on page 8163	"input-vlan-map (Dynamic Interfaces)" on page 8732	"layer2-policer" on page 9653
"no-auto-virtual-gateway-esi" on page 11410	"no-keepalives (Dynamic Profiles)" on page 11490	"output-vlan-map (Dynamic Interfaces)" on page 12023
"proxy-arp (Dynamic Profiles)" on page 13822	"routing-service (Dynamic Profiles)" on page 15988	"routing-services (Enhanced Subscriber Management)" on page 15991
"service (Dynamic Profiles)" on page 16694	"targeted-options (Grouping Subscribers by Bandwidth Usage)" on page 28964	"virtual-gateway-v4-mac" on page 30755
"virtual-gateway-v6-mac" on page 30757	"vlan-id (Dynamic Profiles)" on page 30801	"vlan-tags (Dynamic Profiles)" on page 30867

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family {family}]		
"filter (Dynamic Profiles Filter Attachment)" on page 7132	"rpf-check (Dynamic Profiles)" on page 16009	"service (Dynamic Service Sets)" on page 16695
"unnumbered-address (Dynamic Profiles)" on page 30422		

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family (inet) rpf-check]	
"mode (Dynamic Profiles)" on page 10801	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family {family} filter]		
"adf (Dynamic Firewalls)" on page 1939		

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family {family} filter input {filter-name}]		
"precedence" on page 13008	"shared-name" on page 16970	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family {family} filter output {filter-name}]		
"precedence" on page 13008	"shared-name" on page 16970	

[edit dynamic-profiles {name} interfaces {junos-interface-ifd-name} unit {junos-interface-unit} family inet]		
"auto-configure (IPv4)" on page 2698	"mtu (Dynamic Profiles)" on page 11035	"tcp-mss (Interfaces)" on page 29033

[edit dynamic-profiles {name} interfaces {junos-interface-ifd-name} unit {junos-interface-unit} family inet6]		
"auto-configure (IPv6)" on page 2700	"dad-proxy" on page 4809	"ndp-proxy" on page 11259
"tcp-mss (Interfaces)" on page 29033		

[edit dynamic-profiles {name} interfaces {name} unit {name} family inet6 address]
"ndp" on page 11257

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family pppoe]		
"access-concentrator" on page 1655	"direct-connect" on page 5708	"duplicate-protection (Dynamic PPPoE)" on page 6071

(Continued)

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family pppoe]		
"dynamic-profile (Dynamic PPPoE)" on page 6124	"max-sessions (Dynamic PPPoE)" on page 10469	"max-sessions-vsa-ignore (Static and Dynamic Subscribers)" on page 10481
"service-name-table" on page 16716	"short-cycle-protection (Static and Dynamic Subscribers)" on page 16984	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family {family} rpf-check]	
"fail-filter (Dynamic Profiles)" on page 6777	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family {family} service]		
"input (Dynamic Service Sets)" on page 8697	"output (Dynamic Service Sets)" on page 11986	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family {family} service input]		
"post-service-filter (Dynamic Service Sets)" on page 12908	"service-set (Dynamic Service Sets)" on page 16742	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family {family} service input service-set {service-set-name}]	
"service-filter (Dynamic Service Sets)" on page 16697	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family {family} service output]	
"service-set (Dynamic Service Sets)" on page 16742	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} family {family} service output service-set {service-set-name}]	
"service-filter (Dynamic Service Sets)" on page 16697	

[edit dynamic-profiles interfaces {interface-name} unit {logical-unit-number} family {family} unnumbered-address {interface-name}]	
"preferred-source-address" on page 13060	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} filter input {filter-name}]		
"precedence" on page 13008	"shared-name" on page 16970	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} filter output {filter-name}]		
"precedence" on page 13008	"shared-name" on page 16970	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} input-vlan-map]		
"inner-vlan-id (Dynamic VLANs)" on page 8683	"push (Dynamic VLANs)" on page 13878	"swap (Dynamic VLANs)" on page 28742
"tag-protocol-id (Dynamic VLANs)" on page 28919	"vlan-id (Dynamic VLANs)" on page 30802	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} output-vlan-map]		
"inner-vlan-id (Dynamic VLANs)" on page 8683	"pop (Dynamic VLANs)" on page 12804	"swap (Dynamic VLANs)" on page 28742
"tag-protocol-id (Dynamic VLANs)" on page 28919	"vlan-id (Dynamic VLANs)" on page 30802	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} pppoe-underlying-options]		
"direct-connect" on page 5708	"max-sessions-vsa-ignore (Static and Dynamic Subscribers)" on page 10481	"short-cycle-protection (Static and Dynamic Subscribers)" on page 16984

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number} service]	
"pcef (Dynamic Profiles)" on page 12314	

[edit dynamic-profiles {name} interfaces {name} unit {name} tunnel]	
"ipip" on page 9292	

[edit dynamic-profiles {name} interfaces {name} unit {name} tunnel encapsulation]	
"vxlan-gpe (FTI)" on page 31005	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number}input-vlan-map]	
"inner-tag-protocol-id (Dynamic VLANs)" on page 8680	

[edit dynamic-profiles {profile-name} interfaces {interface-name} unit {logical-unit-number}output-vlan-map]	
"inner-tag-protocol-id (Dynamic VLANs)" on page 8680	

[edit dynamic-profiles {profile-name} interfaces “$junos-interface-ifd-name” unit “$junos-interface-unit”]	
"auto-configure (Dynamic VLAN Interface Sets)" on page 2696	

[edit dynamic-profiles {profile-name} interfaces “$junos-interface-ifd-name” unit “$junos-interface-unit” auto-configure]		
"agent-circuit-identifier (Dynamic ACI VLANs)" on page 2048	"line-identity (Dynamic Access-Line-Identifier VLANs)" on page 9772	

[edit dynamic-profiles {profile-name} interfaces “$junos-interface-ifd-name” unit “$junos-interface-unit” auto-configure agent-circuit-identifier]	
"dynamic-profile (Dynamic ACI VLANs)" on page 6122	

[edit dynamic-profiles {profile-name} interfaces “$junos-interface-ifd-name” unit “$junos-interface-unit” auto-configure line-identity]		
"dynamic-profile (Dynamic Access-Line-Identifier VLANs)" on page 6121	"include (Dynamic Access-Line-Identifier VLANs)" on page 8525	

[edit dynamic-profiles {profile-name} interfaces “$junos-interface-ifd-name” unit “$junos-interface-unit” ppp-options]		
"initiate-ncp (Dynamic and Static PPP)" on page 8662	"ipcp-suggest-dns-option" on page 9289	

[logical-systems]

[edit dynamic-profiles {name} logical-systems {name} interfaces]		
"hierarchical-scheduler" on page 8060	"logical-tunnel-options" on page 10093	

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} aggregated-ether-options]		
"bfd-liveness-detection (LAG)" on page 2926	"load-balance" on page 9874	"optimize-aggregate-frr" on page 11853
"targeted-options (Manual Targeting)" on page 28966		

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} aggregated-ether-options bfd-liveness-detection]	
"damping (BFD Liveness Detection)" on page 4815	

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} aggregated-ether-options lacp]	
"sync-reset" on page 28785	

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} aggregated-ether-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} esi]	
"df-election-type" on page 5576	

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} ether-options]		
"fec (ether)" on page 6992		

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} gigether-options]		
"loopback-remote" on page 10128	"mru (Gigabit Ethernet Options)" on page 10991	

[edit dynamic-profiles {name} logical-systems {name} interfaces interface-range {name} aggregated-ether-options]		
"load-balance" on page 9874		

[edit dynamic-profiles {name} logical-systems {name} interfaces interface-range {name} aggregated-ether-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[edit dynamic-profiles {name} logical-systems {name} interfaces interface-range {name} logical-tunnel-options]		
"load-balance" on page 9874		

[edit dynamic-profiles {name} logical-systems {name} interfaces interface-range {name} logical-tunnel-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} logical-tunnel-options]		
"load-balance" on page 9874		

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} logical-tunnel-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} optics-options]		
"alarm (optics-options)" on page 2110	"is-ma" on page 9396	

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} unit]		
"no-auto-virtual-gateway-esi" on page 11410	"targeted-options (Grouping Subscribers by Bandwidth Usage)" on page 28964	"virtual-gateway-v4-mac" on page 30755
"virtual-gateway-v6-mac" on page 30757		

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} unit {name} family inet]		
"mtu (Dynamic Profiles)" on page 11035	"tcp-mss (Interfaces)" on page 29033	

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} unit {name} family inet6]		
"dad-proxy" on page 4809	"ndp-proxy" on page 11259	"tcp-mss (Interfaces)" on page 29033

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} unit {name} family inet6 address]	
"ndp" on page 11257	

[edit dynamic-profiles {name} logical-systems {name} interfaces {name} unit {name} tunnel]	
"ipip" on page 9292	

[multicast]

[edit dynamic-profiles {name} multicast]	
"omit-wildcard-address" on page 11812	

[name]

[edit dynamic-profiles name interfaces name load-balancing-options]	
"no-bundle-flap" on page 11422	

[policy-options]

[edit dynamic-profiles {profile-name} policy-options]	
"policy-statement" on page 12736	

[protocols]

[edit dynamic-profiles {profile-name} protocols]		
"igmp (Dynamic Profiles)" on page 8383	"mld (Dynamic Profiles)" on page 10781	"router-advertisement (Dynamic Profiles)" on page 15895

[edit dynamic-profiles {profile-name} protocols igmp]		
"interface (Dynamic IGMP)" on page 8871		

[edit dynamic-profiles {profile-name} protocols igmp interface {interface-name}]		
"group-policy (Dynamic IGMP Interface)" on page 7852	"log-interval (IGMP Interface)" on page 10009	"no-accounting" on page 11371
"oif-map (Dynamic IGMP Interface)" on page 11800	"passive (Dynamic IGMP Interface)" on page 12199	"ssm-map (Dynamic IGMP Interface)" on page 28453
"ssm-map-policy (Dynamic IGMP Interface)" on page 28458	"static (Dynamic IGMP Interface)" on page 28532	

[edit dynamic-profiles protocols igmp interface $junos-interface-name]		
"distributed (IGMP)" on page 5834		

[edit dynamic-profiles {profile-name} protocols igmp interface {interface-name} static]		
"source (Dynamic IGMP Interface)" on page 28201		

[edit dynamic-profiles {profile-name} protocols igmp interface {interface-name} static],]		
"group (Dynamic IGMP Interface)" on page 7832		

[edit dynamic-profiles {profile-name} protocols igmp interface {interface-name}],]		
"accounting (Dynamic IGMP Interface)" on page 1714	"disable (Dynamic IGMP)" on page 5751	"group-limit (Dynamic IGMP Interface)" on page 7841

(Continued)

[edit dynamic-profiles {profile-name} protocols igmp interface {interface-name}],]		
"immediate-leave (Dynamic IGMP Interface)" on page 8456	"promiscuous-mode (Dynamic IGMP Interface)" on page 13373	

[edit dynamic-profiles {profile-name} protocols igmpinterface {interface-name}]	
"version (Dynamic IGMP Interface)" on page 30699	

[edit dynamic-profiles {profile-name} protocols mld]	
"interface (Dynamic MLD)" on page 8875	

[edit dynamic-profiles {profile-name} protocols mld interface {interface-name}]		
"accounting (Dynamic MLD Interface)" on page 1716	"disable (Dynamic MLD)" on page 5752	"group-limit (Dynamic MLD Interface)" on page 7843
"group-policy (Dynamic MLD Interface)" on page 7854	"immediate-leave (Dynamic MLD Interface)" on page 8458	"log-interval (MLD Interface)" on page 10011
"oif-map (Dynamic MLD Interface)" on page 11801	"passive (Dynamic MLD Interface)" on page 12201	"ssm-map (Dynamic MLD Interface)" on page 28455
"ssm-map-policy (Dynamic MLD Interface)" on page 28459	"static (Dynamic MLD Interface)" on page 28534	"version (Dynamic MLD Interface)" on page 30701

[edit dynamic-profiles {profile-name} protocols mld interface {interface-name} static]	
"group (Dynamic MLD Interface)" on page 7834	

[edit dynamic-profiles {profile-name} protocols mld interface {interface-name} static group {multicast-group-address}]		
"exclude (Dynamic MLD Interface)" on page 6604	"group-count (Dynamic MLD Interface)" on page 7809	"source (Dynamic MLD Interface)" on page 28203

[edit dynamic-profiles {profile-name} protocols mld interface {interface-name} static group {multicast-group-address} source]		
"group-increment (Dynamic MLD Interface)" on page 7837	"source-count (Dynamic MLD Interface)" on page 28196	"source-increment (Dynamic MLD Interface)" on page 28220

[edit dynamic-profiles {name} protocols pim]		
"disable-packet-register" on page 5813	"rpf-vector (PIM)" on page 16017	"secondary-vtep-address" on page 16419

[edit dynamic-profiles {name} protocols pim interface]	
"dual-dr" on page 6045	

[edit dynamic-profiles {name} protocols pim interface (Protocols PIM) {interface-name}]	
"distributed-dr" on page 5833	

[edit dynamic-profiles {name} protocols pim rp rpf-vector]	
"policy (PIM rpf-vector) " on page 12703	

[edit dynamic-profiles protocols router-advertisement]	
"interface (Dynamic Router Advertisement)" on page 8877	

[edit dynamic-profiles protocols router-advertisement interface {interface-name}]		
"current-hop-limit (Dynamic Router Advertisement)" on page 4737	"dns-server-address (Dynamic Profiles)" on page 5865	"managed-configuration (Dynamic Router Advertisement)" on page 10302
"max-advertisement-interval (Dynamic Router Advertisement)" on page 10413	"min-advertisement-interval (Dynamic Router Advertisement)" on page 10708	"on-link (Dynamic Router Advertisement)" on page 11823
"other-stateful-configuration (Dynamic Router Advertisement)" on page 11954	"preferred-lifetime (Dynamic Router Advertisement)" on page 13058	"prefix (Dynamic Router Advertisement)" on page 13081

[edit dynamic-profiles {profile-name} protocols router-advertisement interface {interface-name} prefix {prefix}]		
"autonomous (Dynamic Router Advertisement)" on page 2763	"valid-lifetime (Dynamic Router Advertisement)" on page 30645	

[routing-instances]

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance]	
"routing-options (Dynamic Profiles)" on page 15981	

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options]		
"access (Dynamic Access Routes)" on page 1659	"access-internal (Dynamic Access-Internal Routes)" on page 1666	

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options access]		
"route (Access)" on page 15841		

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options access route {prefix}]		
"metric (Dynamic Access-Internal Routes)" on page 10670	"next-hop (Dynamic Access Routes)" on page 11334	"preference (Subscriber Management)" on page 13050
"tag (Access)" on page 28905	"tag2 (Dynamic Access Routes)" on page 28937	

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options access-internal]	
"route (Access Internal)" on page 15843	

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options access-internal route {subscriber-ip-address}]	
"qualified-next-hop (Dynamic Access-Internal Routes)" on page 13908	

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options access-internal route {subscriber-ip-address} qualified-next-hop {underlying-interface}]	
"mac-address (Dynamic Access-Internal Routes)" on page 10224	

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options rib {routing-table-name}]		
"access (Dynamic Access Routes)" on page 1659	"access-internal (Dynamic Access-Internal Routes)" on page 1666	

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options rib {routing-table-name} access]		
"route (Access)" on page 15841		

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options rib {routing-table-name} access route {prefix}]		
"metric (Dynamic Access-Internal Routes)" on page 10670	"next-hop (Dynamic Access Routes)" on page 11334	"preference (Subscriber Management)" on page 13050
"tag (Access)" on page 28905		

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options rib {routing-table-name} access-internal]	
"route (Access Internal)" on page 15843	

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options rib {routing-table-name} access-internal route {subscriber-ip-address}]	
"qualified-next-hop (Dynamic Access-Internal Routes)" on page 13908	

[edit dynamic-profiles {profile-name} routing-instances $junos-routing-instance routing-options rib {routing-table-name} access-internal route {subscriber-ip-address} qualified-next-hop {underlying-interface}]	
"mac-address (Dynamic Access-Internal Routes)" on page 10224	

[edit dynamic-profiles {profile-name} routing-instances {routing-instance-name} routing-options]		
"multicast (Dynamic Routing Options)" on page 11067	"multicast (routing-options)" on page 11063	

[edit dynamic-profiles {profile-name} routing-instances {routing-instance-name} routing-options multicast]	
"interface (Dynamic Routing Options)" on page 8880	

[routing-options]

[edit dynamic-profiles {profile-name} routing-options]		
"access (Dynamic Access Routes)" on page 1659	"access-internal (Dynamic Access-Internal Routes)" on page 1666	"multicast (Dynamic Routing Options)" on page 11067
"multicast (routing-options)" on page 11063		

[edit dynamic-profiles {profile-name} routing-options access]	
"route (Access)" on page 15841	

[edit dynamic-profiles {profile-name} routing-options access route {prefix}]		
"metric (Dynamic Access-Internal Routes)" on page 10670	"next-hop (Dynamic Access Routes)" on page 11334	"preference (Subscriber Management)" on page 13050
"tag (Access)" on page 28905	"tag2 (Dynamic Access Routes)" on page 28937	

[edit dynamic-profiles {profile-name} routing-options access-internal]	
"route (Access Internal)" on page 15843	

[edit dynamic-profiles {profile-name} routing-options access-internal route {subscriber-ip-address} qualified-next-hop {underlying-interface}].]	
"qualified-next-hop (Dynamic Access-Internal Routes)" on page 13908	

[edit dynamic-profiles routing-options access-internal route {subscriber-ip-address }qualified-next-hop {underlying-interface}].]	
"mac-address (Dynamic Access-Internal Routes)" on page 10224	

[edit dynamic-profiles {profile-name} routing-options multicast]	
"interface (Dynamic Routing Options)" on page 8880	

[edit dynamic-profiles {profile-name} routing-options multicast interface {interface-name}]	
"no-qos-adjust (Dynamic Routing Options)" on page 11556	

[edit dynamic-profiles {profile-name} routing-options rib {routing-table-name} access]	
"route (Access)" on page 15841	

[services]

[edit dynamic-profiles {profile-name} services captive-portal-content-delivery]	
"rule (Captive Portal Content Delivery)" on page 16085	

[edit dynamic-profiles {profile-name} services captive-portal-content-delivery rule {rule-name}]	
"term (Captive Portal Content Delivery)" on page 29113	

[edit dynamic-profiles {profile-name} services captive-portal-content-delivery rule {rule-name} term {term-name}]	
"then (Captive Portal Content Delivery)" on page 29259	

[telemetry]

[edit dynamic-profiles {name} telemetry]	
"subscriber-info" on page 28704	

[variables]

[edit dynamic-profiles {profile-name} variables {variable-name}]		
"default-value" on page 5221	"equals (Dynamic Profile)" on page 6433	"mandatory" on page 10313
"uid (Dynamic Profiles)" on page 30319	"uid-reference" on page 30320	

[edit ethernet-switching]

IN THIS SECTION

- [options] | 173

Sub Hierarchies

"[edit ethernet-switching options]" on page 173

[options]

[edit ethernet-switching options secure-access-port vlan (all | {vlan-name}) examine-fip]		
"examine-vn2vn" on page 6568	"fc-map" on page 6957	

[edit ethernet-switching options secure-access-port vlan (all | {vlan-name}) examine-fip examine-vn2vn]	
"beacon-period" on page 2890	

[edit ethernet-switching-options]

IN THIS SECTION

- [analyzer] | 175
- [bpdu-block] | 176
- [dot1q-tunneling] | 177
- [mac-notification] | 177
- [port-error-disable] | 177
- [redundant-trunk-group] | 177
- [secure-access-port] | 178
- [static] | 180
- [storm-control] | 180
- [unknown-unicast-forwarding] | 180

Sub Hierarchies

"[edit ethernet-switching-options analyzer {name}]" on page 175

"[edit ethernet-switching-options bpdu-block]" on page 176

"[edit ethernet-switching-options dot1q-tunneling]" on page 177

"[edit ethernet-switching-options mac-notification]" on page 177

"[edit ethernet-switching-options port-error-disable]" on page 177

"[edit ethernet-switching-options redundant-trunk-group]" on page 177

"[edit ethernet-switching-options secure-access-port]" on page 178

"[edit ethernet-switching-options static]" on page 180

"[edit ethernet-switching-options storm-control]" on page 180

"[edit ethernet-switching-options unknown-unicast-forwarding]" on page 180

[edit ethernet-switching-options]		
"analyzer" on page 2237	"authentication-whitelist" on page 2671	"bpdu-block" on page 3022
"interfaces (Q-in-Q Tunneling)" on page 9154	"mac-lookup-length" on page 10244	"mac-notification" on page 10249
"mac-table-aging-time" on page 10267	"nonstop-bridging (Ethernet Switching)" on page 11670	"port-error-disable" on page 12821
"secure-access-port" on page 16432	"storm-control (switch)" on page 28622	"traceoptions (Access Port Security)" on page 29479
"traceoptions (Ethernet Switching)" on page 29465		

[analyzer]

[edit ethernet-switching-options analyzer {name}]		
"input" on page 8687	"output" on page 11980	

[edit ethernet-switching-options analyzer {name} input]		
"egress" on page 6205	"ingress (Port Mirroring)" on page 8643	

[edit ethernet-switching-options analyzer {name} input (egress | ingress)]	
"interface (Port Mirroring)" on page 8952	

[edit ethernet-switching-options analyzer {name} input ingress]	
"vlan (Port Mirroring)" on page 30839	

[edit ethernet-switching-options analyzer {name} output]		
"interface (Port Mirroring)" on page 8952	"vlan (Port Mirroring)" on page 30839	

[edit ethernet-switching-options analyzer {name} output vlan (vlan-id | vlan-name)]	
"no-tag" on page 11594	

[bpdu-block]

[edit ethernet-switching-options bpdu-block]		
"disable-timeout (BPDU)" on page 5808	"interface (BPDU Blocking)" on page 8842	

[edit ethernet-switching-options bpdu-block bpdu-block (all | [{interface-name}])]	
"drop (BPDU Block)" on page 5944	

[edit ethernet-switching-options bpdu-block interface]	
"disable (STP)" on page 5739	

[edit ethernet-switching-options bpdu-block interface (all | [{interface-name}])]	
"shutdown (BPDU Block)" on page 27935	

[dot1q-tunneling]

[edit ethernet-switching-options dot1q-tunneling]	
"ether-type (VLAN)" on page 6474	

[mac-notification]

[edit ethernet-switching-options mac-notification]	
"notification-interval" on page 11690	

[port-error-disable]

[edit ethernet-switching-options port-error-disable]		
"disable-timeout" on page 5806	"disable-timeout (Port Error Disable)" on page 5810	

[redundant-trunk-group]

[edit ethernet-switching-options redundant-trunk-group group {name}]	
"interface (Redundant Trunk Groups)" on page 9008	

[secure-access-port]

[edit ethernet-switching-options secure-access-port]		
"dhcp-snooping-file (EX)" on page 5649	"interface (Access Port Security)" on page 8829	"interface (Secure Access Port)" on page 9017
"ipv6-source-guard-sessions" on page 9372	"vlan (Access Port Security)" on page 30785	"vlan (Secure Access Port)" on page 30852

[edit ethernet-switching-options secure-access-port dhcp-snooping-file]		
"location (DHCP Snooping Database)" on page 9987	"location (DHCP Snooping Database)" on page 9984	"timeout" on page 29367

[edit ethernet-switching-options secure-access-port dhcpv6-snooping-file]		
"location (DHCP Snooping Database)" on page 9987	"timeout" on page 29367	

[edit ethernet-switching-options secure-access-port interface {interface-name}]		
"fcoe-trusted" on page 6965	"no-fcoe-trusted" on page 11458	"static-ipv6" on page 28544

[edit ethernet-switching-options secure-access-port interface (Access Port Security) (all | {interface-name})]	
"no-dhcp-trusted" on page 11438	

[edit ethernet-switching-options secure-access-port interface (all | {interface-name}) static-ip {ip-address}]		
"vlan (DHCP Bindings on Access Ports)" on page 30790	"vlan (Static IP)" on page 30854	

[edit ethernet-switching-options secure-access-port interface (all | {interface-name})]		
"allowed-mac" on page 2214	"dhcp-trusted" on page 5654	"mac-limit" on page 10236
"mac-limit (Access Port Security)" on page 10241	"no-allowed-mac-log" on page 11393	"voip-mac-exclusive" on page 30920

[edit ethernet-switching-options secure-access-port interface {interface-name} vlan {vlan-name}]	
"mac-limit (Access Port Security)" on page 10241	

[edit ethernet-switching-options secure-access-port vlan (all | {vlan-name})]		
"dhcp-option82" on page 5619	"examine-dhcp" on page 6561	"examine-dhcpv6" on page 6563
"examine-fip" on page 6565	"neighbor-discovery-inspection" on page 11270	"no-examine-dhcpv6" on page 11452

[edit ethernet-switching-options secure-access-port vlan (all | {vlan-name}) (examine-dhcp | arp-inspection)]	
"forwarding-class (for DHCP Snooping or DAI Packets)" on page 7425	

[edit ethernet-switching-options secure-access-port vlan (all | {vlan-name}) dhcp-option82 remote-id]	
"prefix (Remote ID for Option 82)" on page 13101	

[edit ethernet-switching-options secure-access-port vlan (all | {vlan-name}) examine-fip]	
"no-fip-snooping-scaling" on page 11461	

[static]

[edit ethernet-switching-options static vlan {vlan-name}]	
"mac (Static MAC-Based VLANs)" on page 10261	

[edit ethernet-switching-options static vlan {vlan-name} mac {mac-address}]	
"next-hop (Static MAC-Based VLANs)" on page 11353	

[storm-control]

[edit ethernet-switching-options storm-control]	
"interface (Storm Control)" on page 9074	

[edit ethernet-switching-options storm-control interface (all | {interface-name})]		
"bandwidth (storm control)" on page 2833	"level (OAM)" on page 9716	

[unknown-unicast-forwarding]

[edit ethernet-switching-options unknown-unicast-forwarding]	
"vlan (Unknown Unicast Forwarding)" on page 30876	

[edit event-options]

IN THIS SECTION

- [destinations] | 182
- [event-script] | 182
- [policy] | 183

Sub Hierarchies

"[edit event-options destinations {destination-name}]" on page 182

"[edit event-options event-script]" on page 182

"[edit event-options policy {policy-name}]" on page 183

[edit event-options]		
"destinations (Event Policy)" on page 5533	"event-script (Event Options)" on page 6533	"generate-event" on page 7674
"max-policies (Event Policy)" on page 10455	"policy (Event Policy)" on page 12675	"redundancy-event (Services Redundancy Daemon)" on page 14168
"traceoptions (Event Options)" on page 29544		

[destinations]

[edit event-options destinations {destination-name}]		
"archive-sites (Event Policy)" on page 2449	"transfer-delay (Event Policy)" on page 30058	

[event-script]

[edit event-options event-script]		
"dampen" on page 4811	"file (Event Scripts)" on page 7054	"max-datasize" on page 10429
"optional (Event Scripts)" on page 11909	"refresh (Event Scripts)" on page 14236	"refresh-from (Event Scripts)" on page 14240
"start (Scripts)" on page 28496	"traceoptions (Event Scripts)" on page 29547	

[edit event-options event-script dampen]	
"dampen-options" on page 4814	

[edit event-options event-script dampen dampen-options]		
"cpu-factor (Dampen Options)" on page 4694	"line-interval (Dampen Options)" on page 9773	"time-interval (Dampen Options)" on page 29358

[edit event-options event-script file {filename}]		
"cert-file (Scripts)" on page 3252	"checksum (Scripts)" on page 3311	"dampen" on page 4811

(Continued)

[edit event-options event-script file {filename}]		
"python-script-user (Event and SNMP Scripts)" on page 13884	"refresh (Event Scripts)" on page 14236	"refresh-from (Event Scripts)" on page 14240
"remote-execution (Event Scripts)" on page 14320	"routing-instance (Event Scripts)" on page 15934	"source (Event Scripts)" on page 28206
"start (Scripts)" on page 28496		

[edit event-options event-script file {filename} dampen]	
"dampen-options" on page 4814	

[edit event-options event-script file {filename} dampen dampen-options]		
"cpu-factor (Dampen Options)" on page 4694	"line-interval (Dampen Options)" on page 9773	"time-interval (Dampen Options)" on page 29358

[policy]

[edit event-options policy {policy-name}]		
"attributes-match (Event Policy)" on page 2558	"events (Event Policy Trigger Events)" on page 6540	"then (Event Policy)" on page 29261
"within (Event Policy)" on page 31136		

[edit event-options policy {policy-name} attributes-match]		
"equals (Event Policy)" on page 6434	"matches (Event Policy)" on page 10412	"starts-with (Event Policy)" on page 28510

[edit event-options policy {policy-name} then]		
"change-configuration (Event Policy)" on page 3284	"event-script (Event Policy)" on page 6536	"execute-commands (Event Policy)" on page 6636
"ignore (Event Policy)" on page 8399	"priority-override (Event Policy)" on page 13183	"raise-trap (Event Policy)" on page 14042
"upload (Event Policy)" on page 30471		

[edit event-options policy {policy-name} then change-configuration]		
"commands (Event Policy Change Configuration)" on page 4432	"commit-options" on page 4449	"retry (Event Policy)" on page 15704
"user-name (Event Policy)" on page 30569		

[edit event-options policy {policy-name} then event-script {filename}]		
"arguments (Event Policy)" on page 2455	"destination (Event Policy)" on page 5416	"output-filename (Event Policy)" on page 11988
"output-format (Event Policy)" on page 11989	"user-name (Event Policy)" on page 30569	

[edit event-options policy {policy-name} then event-script {filename} destination {destination-name}]		
"retry-count (Event Policy)" on page 15701	"transfer-delay (Event Policy)" on page 30058	

[edit event-options policy {policy-name} then execute-commands]		
"commands (Event Policy Execute Commands)" on page 4434	"destination (Event Policy)" on page 5416	"output-filename (Event Policy)" on page 11988
"output-format (Event Policy)" on page 11989	"user-name (Event Policy)" on page 30569	

[edit event-options policy {policy-name} then execute-commands destination {destination-name}]		
"retry-count (Event Policy)" on page 15701	"transfer-delay (Event Policy)" on page 30058	

[edit event-options policy {policy-name} then priority-override]		
"facility" on page 6769	"severity (Event Policy)" on page 16898	

[edit event-options policy {policy-name} then upload filename ({filename} | committed) destination {destination-name}]		
"retry-count (Event Policy)" on page 15701	"transfer-delay (Event Policy)" on page 30058	"user-name (Event Policy)" on page 30569

[edit event-options policy {policy-name} within {seconds}]		
"events (Event Policy Correlated Events)" on page 6538	"not (Event Policy)" on page 11680	"trigger (Event Policy)" on page 30162

[edit fabric]

IN THIS SECTION

- [aliases] | 187
- [flow-groups] | 187
- [protocols] | 187
- [resources] | 191
- [routing-options] | 191

Sub Hierarchies

"[edit fabric aliases]" on page 187

"[edit fabric flow-groups {flow-group-name}]" on page 187

"[edit fabric protocols]" on page 187

"[edit fabric resources]" on page 191

"[edit fabric routing-options]" on page 191

[edit fabric]		
"aliases" on page 2151	"flow-groups" on page 7291	"protocols (Fabric)" on page 13794
"resources" on page 15625	"routing-options (QFabric System)" on page 15987	

[aliases]

[edit fabric aliases]		
"director-device (Aliases)" on page 5729	"interconnect-device (Aliases)" on page 8800	"node-device (Aliases)" on page 11635

[flow-groups]

[edit fabric flow-groups {flow-group-name}]		
"interconnect-device (Flow Groups)" on page 8804	"node-device (Flow Groups)" on page 11638	

[protocols]

[edit fabric protocols]	
"fabric-control" on page 6749	

[edit fabric protocols bgp]		
"advertise-prefix-sid" on page 2022	"defaults (ACX series, JRR200, MX204, PTX5000, and VRR)" on page 5225	"enforce-first-as" on page 6343
"multipath-build-priority" on page 11107	"stale-labels-holddown-period" on page 28476	"tcpao-auth-mismatch" on page 29034

[edit fabric protocols bgp family inet labeled-unicast]		
"entropy-label (protocols bgp)" on page 6413		

[edit fabric protocols bgp family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit fabric protocols bgp family inet6]		
"transport (Protocols BGP)" on page 30113		

[edit fabric protocols bgp family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit fabric protocols bgp group {group_name}]		
"enforce-first-as" on page 6343	"peer-as-list" on page 12345	

[edit fabric protocols bgp group {name} dynamic-neighbor]		
"peer-limit" on page 12375		

[edit fabric protocols bgp group {name} family inet labeled-unicast]		
"entropy-label (protocols bgp)" on page 6413		

[edit fabric protocols bgp group {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit fabric protocols bgp group {name} family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit fabric protocols bgp group {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit fabric protocols bgp group {name} link-bandwidth]	
"auto-sense" on page 2738	

[edit fabric protocols bgp group {name} metric-out]		
"effective-aigp" on page 6203	"minimum-effective-aigp" on page 10730	

[edit fabric protocols bgp group {group_name} neighbor {address}]	
"enforce-first-as" on page 6343	

[edit fabric protocols bgp group {name} neighbor {name} family inet labeled-unicast]	
"entropy-label (protocols bgp)" on page 6413	

[edit fabric protocols bgp group {name} neighbor {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit fabric protocols bgp group {name} neighbor {name} family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit fabric protocols bgp group {name} neighbor {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit fabric protocols bgp multipath]	
"pause-computation-during-churn" on page 12277	

[edit fabric protocols bgp source-packet-routing]		
"srgb (Protocols BGP)" on page 28407	"srv6 (Protocols BGP)" on page 28417	

[edit fabric protocols fabric-control]	
"graceful-restart (Fabric Control)" on page 7762	

[edit fabric protocols fabric-control graceful-restart]		
"restart-time (Fabric Control)" on page 15665	"stale-routes-time (Fabric Control)" on page 28479	

[resources]

[edit fabric resources]	
"node-group (Resources)" on page 11645	

[edit fabric resources node-group {node-group-name}]		
"network-domain" on page 11302	"node-device (Resources)" on page 11640	

[routing-options]

[edit fabric routing-options]	
"multicast (QFabric Routing Options)" on page 11078	

[edit fabric routing-options multicast]	
"no-make-before-break" on page 11508	

[edit family]

IN THIS SECTION

[ethernet-switching] | **192**

Sub Hierarchies

"[edit family ethernet-switching]" on page 192

"[edit family {family-name} filter {filter-name}]" on page 192

[ethernet-switching]

[edit family ethernet-switching filter {filter name} term {term name} from ip-version ipv4]	
"rdma-opcode" on page 14078	

[edit family ethernet-switching filter {filter name} term {term name} from ip-version ipv6]	
"rdma-opcode" on page 14078	

[filter]

[edit family {family-name} filter {filter-name}]	
"interface-specific (Firewall Filters)" on page 9063	

[edit fc-fabrics]

IN THIS SECTION

- [fc2] | 194
- [interface] | 194
- [protocols] | 194
- [proxy] | 195

Sub Hierarchies

"[edit fc-fabrics {fabric-name} fc2]" on page 194

"[edit fc-fabrics {fc-fabric-name} interface {interface-name}]" on page 194

"[edit fc-fabrics {fc-fabric-name} protocols]" on page 194

"[edit fc-fabrics {fabric-name} proxy]" on page 195

[edit fc-fabrics {fabric-name}]		
"description (Fibre Channel Fabrics)" on page 5329	"fabric-id" on page 6750	"fabric-type" on page 6767
"fc2" on page 6960	"interface (Fibre Channel Fabric)" on page 8886	"max-login-sessions" on page 10442
"protocols (FIP)" on page 13795	"proxy (Fibre Channel)" on page 13825	

[fc2]

[edit fc-fabrics {fabric-name} fc2]	
"traceoptions (FC-2 Fibre Channel)" on page 29554	

[interface]

[edit fc-fabrics {fc-fabric-name} interface {interface-name}]	
"max-login-sessions" on page 10442	

[protocols]

[edit fc-fabrics {fc-fabric-name} protocols]	
"fip" on page 7173	

[edit fc-fabrics {fc-fabric-name} protocols fip]		
"fc-map" on page 6957	"fcoe-trusted" on page 6965	"fka-adv-period" on page 7202
"interface (FIP)" on page 8887	"max-sessions-per-enode" on page 10472	"priority (FIP)" on page 13174
"traceoptions (FIP Protocol Fibre Channel)" on page 29560		

[edit fc-fabrics {fc-fabric-name} protocols fip interface {interface-name}]		
"fka-adv-period" on page 7202	"priority (FIP)" on page 13174	

[proxy]

[edit fc-fabrics {fabric-name} proxy]		
"auto-load-rebalance" on page 2717	"load-balance-algorithm" on page 9876	"no-fabric-wwn-verify" on page 11456
"traceoptions (Proxy Fibre Channel)" on page 29750		

[edit fc-options]

[edit fc-options]		
"max-login-sessions-per-node" on page 10444	"no-fip-snooping-scaling" on page 11461	"traceoptions (Fibre Channel)" on page 29557

[edit firewall]

IN THIS SECTION

- [atm-policer] | 197
- [family] | 197
- [filter] | 202
- [hierarchical-policer] | 202
- [policer] | 203
- [three-color-policer] | 204

Sub Hierarchies

"[edit firewall atm-policer {atm-policer-name}]" on page 197

"[edit firewall family {family-name}]" on page 197

"[edit firewall filter {filter-name}]" on page 202

"[edit firewall hierarchical-policer]" on page 202

"[edit firewall policer {policer-name}]" on page 203

"[edit firewall three-color-policer {name}]" on page 204

[edit firewall]		
"arp policer" on page 2475	"enhanced-hierarchical-policer" on page 6394	"family (Firewall Filter)" on page 6869
"family (Firewall)" on page 6867	"family (Multifield Classifier)" on page 6891	"family (QFX)" on page 6829

(Continued)

[edit firewall]		
"hierarchical-policer" on page 8055	"hierarchical-policer (aggregate or premium)" on page 8057	"interface-set" on page 9034
"load-balance-group" on page 9880	"no-next-header-to-payload-protocol-mapping" on page 11522	"policer" on page 12622
"policer (CFM Firewall)" on page 12626	"policer (Configuring)" on page 12630	"three-color-policer" on page 29309
"three-color-policer (Configuring)" on page 29314	"three-color-policer (Configuring) (EX)" on page 29313	"tunnel-end-point " on page 30200

[atm-policer]

[edit firewall atm-policer {atm-policer-name}]		
"cdvt" on page 3244	"logical-interface-policer" on page 10069	

[family]

[edit firewall family {family-name}]		
"filter" on page 7119	"filter (Configuring)" on page 7127	"filter (Firewall Filters)" on page 7136

[edit firewall family (inet | inet6]	
"service-filter (Firewall)" on page 16699	

[edit firewall family (inet | inet6) filter {filter-name}]	
"filter-list-template" on page 7148	

[edit firewall family (inet | inet6) filter {filter-name} term {term-name} then]	
"topology (Filter-Based Forwarding)" on page 29440	

[edit firewall family (inet | inet6) filter {filter-name} term {term-name} then logical-system {logical-system-name}]	
"topology (Filter-Based Forwarding)" on page 29440	

[edit firewall family (inet | inet6) filter {filter-name} term {term-name} then logical-system {logical-system-name} routing-instance {routing-instance-name}]	
"topology (Filter-Based Forwarding)" on page 29440	

[edit firewall family (inet | inet6) filter {filter-name} term {term-name} then routing-instance {routing-instance-name}]	
"topology (Filter-Based Forwarding)" on page 29440	

[edit firewall family family-name filter output filter-name]	
"fast-lookup-filter (PTX)" on page 6930	

[edit firewall family {family-name} filter {filter-name}]		
"accounting-profile (firewall)" on page 1726	"enhanced-mode" on page 6396	"enhanced-mode-override" on page 6401
"fast-lookup-filter" on page 6928	"instance-shared" on page 8769	"interface-shared" on page 9048
"interface-specific" on page 9060	"layer2-policer (firewall)" on page 9662	"no-incremental-update" on page 11473
"physical-interface-filter" on page 12476	"promote" on page 13377	"scale-mode" on page 16328
"scale-optimized" on page 16329	"term (Firewall Filter)" on page 29117	"transient-filter" on page 30062

[edit firewall family {family-name} filter {filter-name} term {term-name}]		
"from (Firewall)" on page 7585	"then (Filters)" on page 29264	"then (Firewall Filters)" on page 29266

[edit firewall family {family-name} filter {filter-name} term {term-name} from]		
"destination-address (Firewall)" on page 5371	"ip-version" on page 9285	"rdma-opcode" on page 14078
"rdma-opcode" on page 14078	"rdma-opcode" on page 14078	"slice (firewall filter match condition)" on page 28043
"source-address " on page 28120		

[edit firewall family {family-name} filter {filter-name} term {term-name} from ip-version {ip-version}]		
"destination-address (Firewall)" on page 5371	"protocol (Firewall Filter)" on page 13423	"source-address " on page 28120

[edit firewall family {family-name} filter {filter-name} term {term-name} from ip-version {ip-version} protocol (tcp| udp)]		
"destination-port (firewall)" on page 5458	"source-port (firewall)" on page 28278	

[edit firewall family {family-name} filter {filter-name} term {term-name} then]		
"decapsulate (Firewall Filter)" on page 5160	"dscp (Multifield Classifier)" on page 6017	"force-premium (Firewall Filter Action)" on page 7373
"forwarding-class (Firewall Filter Action)" on page 7423	"forwarding-class (Multifield Classifiers)" on page 7433	"loss-priority (Firewall Filter Action)" on page 10154
"policer (Firewall Filter Action)" on page 12636	"policer-charge" on page 12646	"slice (firewall filter action)" on page 28042
"three-color-policer (Applying)" on page 29311		

[edit firewall family {family-name} filter {filter-name} term {term-name} then decapsulate gre]	
"interface-group (Decapsulate GRE)" on page 8892	

[edit firewall family inet]		
"prefix-action (Configuring)" on page 13068	"simple-filter" on page 27962	

[edit firewall family inet filter {firewall filter name} term {term name} from]		
"fast-lookup-tuple" on page 6935	"fast-lookup-tuple-list" on page 6937	"loss-priority (Firewall Filter)" on page 10152

[edit firewall family inet filter {filter-name} term {term-name} then]		
"prefix-action (Firewall Filter Action)" on page 13070	"routing-instance" on page 15921	

[edit firewall family inet prefix-action {name}]	
"filter-specific" on page 7155	

[edit firewall family inet simple-filter {filter-name}]	
"term (Simple Filter)" on page 29137	

[edit firewall family mpls filter {name} term {name} from]	
"ip-version (Family MPLS)" on page 9286	

[edit firewall family mpls filter {name} term {name} from ip-version]		
"ipv4 (Family MPLS)" on page 9330	"ipv6 (Family MPLS)" on page 9357	

[edit firewall family {family-name} service-filter {filter-name}]	
"term (Firewall Filter)" on page 29117	

[edit firewall family {family-name} simple-filter {filter-name}]	
"term (Firewall Filter)" on page 29117	

[edit firewall family {family-name} simple-filter {filter-name} term {term-name} then]	
"loss-priority (Simple Firewall Filter)" on page 10170	

[filter]

[edit firewall filter {filter-name}]		
"enhanced-mode" on page 6396	"enhanced-mode-override" on page 6401	

[edit firewall filter {name} term {name} from]	
"loss-priority (Firewall Filter)" on page 10152	

[hierarchical-policer]

[edit firewall hierarchical-policer]		
"aggregate (Hierarchical Policer)" on page 2064	"premium (Hierarchical Policer)" on page 13110	"shared-bandwidth-policer (Configuring)" on page 16962

[edit firewall hierarchical-policer aggregate]		
"if-exceeding (Hierarchical Policer)" on page 8371	"if-exceeding-pps (Hierarchical Policer)" on page 8374	"then (Hierarchical Policer) " on page 29257

[edit firewall hierarchical-policer aggregate if-exceeding]		
"bandwidth-limit (Hierarchical Policer)" on page 2855	"burst-size-limit (Hierarchical Policer)" on page 3084	

[edit firewall hierarchical-policer premium]		
"if-exceeding (Hierarchical Policer)" on page 8371	"if-exceeding-pps (Hierarchical Policer)" on page 8374	"then (Hierarchical Policer) " on page 29257

[edit firewall hierarchical-policer premium if-exceeding]		
"bandwidth-limit (Hierarchical Policer)" on page 2855	"burst-size-limit (Hierarchical Policer)" on page 3084	

[policer]

[edit firewall policer {policer-name}]		
"counter" on page 4690	"filter-specific" on page 7155	"filter-specific (QFX)" on page 7154
"if-exceeding (EX)" on page 8370	"if-exceeding (Policer)" on page 8373	"if-exceeding (QFX)" on page 8368
"if-exceeding-pps (Policer)" on page 8376	"logical-bandwidth-policer" on page 10058	"logical-interface-policer" on page 10069
"logical-interface-policer" on page 10069	"physical-interface-policer" on page 12479	"shared-bandwidth-policer (Configuring)" on page 16962
"then (Policer Action)" on page 29279	"then (Policers)" on page 29281	

[edit firewall policer {policer-name} if-exceeding]		
"bandwidth-limit" on page 2853	"bandwidth-limit (Policer)" on page 2857	"bandwidth-percent" on page 2866
"burst-size-limit" on page 3083	"burst-size-limit (Policer)" on page 3086	

[edit firewall policer {policer-name} if-exceeding-pps]		
"packet-burst (Policer)" on page 12135	"pps-limit (Policer)" on page 12984	

[three-color-policer]

[edit firewall three-color-policer {name}]		
"action" on page 1759	"action (Policer)" on page 1765	"action (QFX)" on page 1755
"logical-interface-policer" on page 10069	"physical-interface-policer" on page 12479	"shared-bandwidth-policer (Configuring)" on page 16962
"single-rate" on page 27969	"two-rate" on page 30259	

[edit firewall three-color-policer {policer-name} action]	
"loss-priority high then discard (Three-Color Policer)" on page 10155	

[edit firewall three-color-policer {policer-name} single-rate]		
"color-aware" on page 4428	"color-blind" on page 4429	"committed-burst-size" on page 4458
"committed-information-rate" on page 4460	"excess-burst-size" on page 6580	"excess-burst-size (QFX)" on page 6579

[edit firewall three-color-policer {policer-name} two-rate]		
"color-aware" on page 4428	"color-blind" on page 4429	"committed-burst-size" on page 4458

(Continued)

[edit firewall three-color-policer {policer-name} two-rate]		
"committed-information-rate" on page 4460	"peak-burst-size" on page 12332	"peak-burst-size (QFX)" on page 12331
"peak-information-rate" on page 12336	"peak-information-rate (QFX)" on page 12334	

[edit forwarding-options]

IN THIS SECTION

- [access-security] | 209
- [accounting] | 209
- [accountingoutput] | 211
- [analyzer] | 211
- [dhcp-relay] | 213
- [dhcp-relaygroup] | 235
- [enhanced-hash-key] | 236
- [evpn-vxlan] | 237
- [family] | 237
- [hash-key] | 239
- [helpers] | 240
- [helperstraceoptions] | 243
- [interface] | 243
- [load-balance] | 243
- [monitoring] | 244
- [next-hop] | 245
- [next-hop-group] | 245

- [pfe-sensor] | **246**
- [port-mirror] | **246**
- [port-mirroring] | **246**
- [sampling] | **250**
- [satellite] | **257**
- [sflow] | **258**
- [storm-control] | **258**
- [storm-control-profiles] | **258**
- [vlans] | **259**
- [vxlan-routing] | **259**

Sub Hierarchies

"[edit forwarding-options access-security {policy-name}]" on page 209

"[edit forwarding-options accounting {name}]" on page 209

"[edit forwarding-options accountingoutput]" on page 211

"[edit forwarding-options analyzer {name}]" on page 211

"[edit forwarding-options dhcp-relay]" on page 213

"[edit forwarding-options dhcp-relaygroup]" on page 235

"[edit forwarding-options enhanced-hash-key]" on page 236

"[edit forwarding-options evpn-vxlan]" on page 237

"[edit forwarding-options family]" on page 237

"[edit forwarding-options hash-key]" on page 239

"[edit forwarding-options helpers]" on page 240

"[edit forwarding-options helperstraceoptions]" on page 243

"[edit forwarding-options interface {interface-name}]" on page 243

"[edit forwarding-options load-balance]" on page 243

"[edit forwarding-options monitoring {name}]" on page 244

"[edit forwarding-options next-hop]" on page 245

"[edit forwarding-options next-hop-group]" on page 245

"[edit forwarding-options pfe-sensor]" on page 246

"[edit forwarding-options port-mirror]" on page 246

"[edit forwarding-options port-mirroring]" on page 246

"[edit forwarding-options sampling]" on page 250

"[edit forwarding-options satellite]" on page 257

"[edit forwarding-options sflow]" on page 258

"[edit forwarding-options storm-control]" on page 258

"[edit forwarding-options storm-control-profiles]" on page 258

"[edit forwarding-options vlans {vlan-name}]" on page 259

"[edit forwarding-options vxlan-routing]" on page 259

[edit forwarding-options]		
"access-security" on page 1700	"accounting (Forwarding Options)" on page 1707	"analyzer" on page 2237
"cut-through" on page 4786	"dhcp-relay" on page 5620	"enhanced-hash-key" on page 6362
"enhanced-hash-key (EX)" on page 6378	"enhanced-hash-key profile1/ profile2" on page 6389	"evpn-vxlan" on page 6555
"explicit-null-cos" on page 6650	"family (Filtering)" on page 6865	"fast-reroute-priority" on page 6941
"hash-key (Forwarding Options)" on page 7984	"hash-key (Forwarding Options)" on page 7986	"helpers" on page 8049
"hyper-mode (forwarding-options)" on page 8221	"interface (RA Guard)" on page 9002	"interface (SLAAC Snooping)" on page 9052

(Continued)

[edit forwarding-options]		
"ip-options-protocol-queue " on page 9249	"l2circuit-control-passthrough" on page 9534	"link-layer-broadcast-inet-check" on page 9796
"load-balance (Forwarding Options)" on page 9878	"load-balance-label-capability" on page 9882	"loopback-port" on page 10126
"monitoring" on page 10954	"monitoring (Forwarding Options)" on page 10956	"multicast-replication" on page 11080
"next-hop-group (Forwarding Options)" on page 11340	"next-hop-group (Unknown Unicast Forwarding)" on page 11344	"no-hierarchical-ecmp" on page 11470
"no-load-balance-label-capability" on page 11495	"port-mirroring" on page 12836	"router-advertisement-guard" on page 15897
"rpf-loose-mode-discard" on page 16014	"sampling (Forwarding Options)" on page 16293	"slaac-snooping" on page 28030
"storm-control-profiles" on page 28624	"suppress-interface-leaf" on page 28730	"tunnel-termination" on page 30245
"vlans (RA Guard)" on page 30894	"vxlan-disable-copy-tos-decap" on page 30993	"vxlan-disable-copy-tos-encap" on page 30995
"vxlan-routing " on page 31007	"vxlan-tos-copy-filter " on page 31011	

[access-security]

[edit forwarding-options access-security {policy-name}]		
"accept (Router Advertisement Guard Policy)" on page 1627	"discard (forwarding-options)" on page 5814	"match-list" on page 10383
"match-option" on page 10387	"policy" on page 12665	"prefix-list-name" on page 13095
"prefix-list-name" on page 13095	"source-ip-address-list" on page 28235	"source-ip-address-list" on page 28235
"source-mac-address-list" on page 28249	"source-mac-address-list" on page 28249	

[edit forwarding-options access-security interface {interface-name}]	
"policy" on page 12665	

[edit forwarding-options access-security vlans ({vlan-name}| all)]	
"policy" on page 12665	

[accounting]

[edit forwarding-options accounting {name}]	
"output (Accounting)" on page 11981	

[edit forwarding-options accounting {name} output]		
"aggregate-export-interval" on page 2062	"cflowd (Discard Accounting)" on page 3268	"flow-active-timeout (Flow Monitoring)" on page 7261
"flow-inactive-timeout (Flow Monitoring)" on page 7296	"interface (Accounting or Sampling)" on page 8831	

[edit forwarding-options accounting output cflowd {hostname}]		
"aggregation (Forwarding Options)" on page 2084	"port (cflowd)" on page 12813	"port (Flow Monitoring)" on page 12823

[edit forwarding-options accounting {name} output flow-server {hostname}]	
"version (Flow Aggregation)" on page 30692	

[edit forwarding-options accounting {name} output interface {interface-name}]		
"engine-id (Forwarding Options)" on page 6348	"engine-type" on page 6352	

[edit forwarding-options accounting {name} outputinterface {interface-name}]	
"source-address (Forwarding Options)" on page 28133	

[accountingoutput]

[edit forwarding-options accountingoutput cflowd {hostname}]	
"autonomous-system-type (Flow Monitoring)" on page 2767	

[analyzer]

[edit forwarding-options analyzer {name}]		
"input" on page 8687	"input (Analyzer)" on page 8688	"output" on page 11980
"output (Mirroring)" on page 11998		

[edit forwarding-options analyzer analyzer-name input]	
"maximum-packet-length" on page 10540	

[edit forwarding-options analyzer {name} input]		
"egress" on page 6205	"egress (Analyzer)" on page 6207	"ingress (Analyzer)" on page 8625
"ingress (Port Mirroring)" on page 8643	"rate (Forwarding Options)" on page 14060	

[edit forwarding-options analyzer {name} input (egress | ingress)]		
"interface (Port Mirroring)" on page 8952	"vlan (Port Mirroring)" on page 30839	

[edit forwarding-options analyzer {analyzer-name} input egress]		
"bridge-domain (Analyzer)" on page 3031	"interface (Analyzer)" on page 8835	"routing-instance (Forwarding-options Analyzer)" on page 15920

[edit forwarding-options analyzer {analyzer-name} input egress routing-instance {instance-name}]	
"bridge-domain (Analyzer)" on page 3031	

[edit forwarding-options analyzer {analyzer-name} input ingress]		
"bridge-domain (Analyzer)" on page 3031	"interface (Analyzer)" on page 8835	"routing-instance (Forwarding-options Analyzer)" on page 15920
"vlan (Mirroring)" on page 30829		

[edit forwarding-options analyzer {analyzer-name} input ingress routing-instance {instance-name}]		
"bridge-domain (Analyzer)" on page 3031	"vlan (Mirroring)" on page 30829	

[edit forwarding-options analyzer {analyzer-name} ouput]	
"next-hop-group (Analyzer)" on page 11339	

[edit forwarding-options analyzer {analyzer-name} output]		
"bridge-domain (Analyzer)" on page 3031	"interface (Analyzer)" on page 8835	"interface (Port Mirroring)" on page 8952
"ip-address (Port Mirroring)" on page 9238	"routing-instance (Forwarding-options Analyzer)" on page 15920	"routing-instance (Port Mirroring)" on page 15953

(Continued)

[edit forwarding-options analyzer {analyzer-name} output]		
"vlan (Mirroring)" on page 30829	"vlan (Port Mirroring)" on page 30839	

[edit forwarding-options analyzer {analyzer-name} output routing-instance {instance-name}]		
"bridge-domain (Analyzer)" on page 3031	"vlan (Mirroring)" on page 30829	

[edit forwarding-options analyzer {analyzer-name} output vlan ({vlan-id} | {vlan-name})]		
"no-tag" on page 11592		

[dhcp-relay]

[edit forwarding-options dhcp-relay]		
"active-leasequery (DHCP Relay Agent)" on page 1808	"active-server-group" on page 1817	"authentication (DHCP Relay Agent)" on page 2595
"bulk-leasequery (DHCP Relay Agent)" on page 3068	"dhcpv6 (DHCP Relay Agent)" on page 5665	"dual-stack-group (DHCP Relay Agent)" on page 6055
"duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)" on page 6066	"dynamic-profile (DHCP Relay Agent)" on page 6119	"forward-only (DHCP Relay Agent)" on page 7396
"forward-only-replies (DHCP Relay Agent)" on page 7401	"forward-snooped-clients (DHCP Relay Agent)" on page 7405	"group (DHCP Relay Agent)" on page 7820

(Continued)

[edit forwarding-options dhcp-relay]		
"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701	"leasequery (DHCP Relay Agent)" on page 9702	"liveness-detection" on page 9846
"no-snoop (DHCP Local Server and Relay Agent)" on page 11582	"overrides (DHCP Relay Agent)" on page 12068	"relay-option (DHCP Relay Agent)" on page 14297
"relay-option-82" on page 14294	"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334	"route-suppression (DHCP Local Server and Relay Agent)" on page 15877
"server-group" on page 16633	"server-response-time (DHCP Relay Agent)" on page 16664	"service-profile (DHCP Relay Agent)" on page 16731
"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982	"source-ip-change (Forwarding Options)" on page 28239	

[edit forwarding-options dhcp-relay ... overrides]	
"delay-authentication (DHCP Relay Agent)" on page 5242	

[edit forwarding-options dhcp-relay authentication]		
"password (DHCP Relay Agent)" on page 12229	"username-include (DHCP Relay Agent)" on page 30597	

[edit forwarding-options dhcp-relay authentication username-include]		
"circuit-type (DHCP Relay Agent)" on page 3332	"client-id (DHCP Relay Agent)" on page 4317	"delimiter (DHCP Relay Agent)" on page 5290

(Continued)

[edit forwarding-options dhcp-relay authentication username-include]		
"domain-name (DHCP Relay Agent)" on page 5886	"interface-description (DHCP Relay Agent)" on page 8856	"interface-name (DHCP Relay Agent)" on page 8938
"logical-system-name (DHCP Relay Agent)" on page 10083	"mac-address (DHCP Relay Agent)" on page 10223	"option-82 (DHCP Relay Agent)" on page 11893
"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580	

[edit forwarding-options dhcp-relay dhcpv6]		
"active-leasequery (DHCP Relay Agent)" on page 1808	"active-server-group" on page 1817	"authentication (DHCP Relay Agent)" on page 2595
"bulk-leasequery (DHCP Relay Agent)" on page 3068	"duplicate-clients (DHCPv6 Local Server and Relay Agent)" on page 6064	"dynamic-profile (DHCP Relay Agent)" on page 6119
"forward-only (DHCP Relay Agent)" on page 7396	"forward-only-replies (DHCP Relay Agent)" on page 7401	"forward-snooped-clients (DHCP Relay Agent)" on page 7405
"group (DHCP Relay Agent)" on page 7820	"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701	"leasequery (DHCP Relay Agent)" on page 9702
"liveness-detection" on page 9846	"no-snoop (DHCP Local Server and Relay Agent)" on page 11582	"overrides (DHCP Relay Agent)" on page 12068
"relay-agent-interface-id (DHCPv6 Relay Agent)" on page 14276	"relay-agent-option-79 " on page 14281	"relay-agent-remote-id (DHCPv6 Relay Agent)" on page 14284
"relay-option (DHCP Relay Agent)" on page 14297	"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334	"route-suppression (DHCP Local Server and Relay Agent)" on page 15877

(Continued)

[edit forwarding-options dhcp-relay dhcpv6]		
"server-group" on page 16633	"server-response-time (DHCP Relay Agent)" on page 16664	"service-profile (DHCP Relay Agent)" on page 16731
"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982	"use-vlan-id (DHCP Relay Agent)" on page 30546	

[edit forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)]		
"include-irb-and-l2" on page 8529	"no-vlan-interface-name" on page 11617	"prefix (DHCP Relay Agent)" on page 13077
"use-interface-description" on page 30522	"use-option-82" on page 30528	

[edit forwarding-options dhcp-relay dhcpv6 ... overrides]	
"delay-authentication (DHCP Relay Agent)" on page 5242	

[edit forwarding-options dhcp-relay dhcpv6 authentication]		
"password (DHCP Relay Agent)" on page 12229	"username-include (DHCP Relay Agent)" on page 30597	

[edit forwarding-options dhcp-relay dhcpv6 authentication username-include]		
"circuit-type (DHCP Relay Agent)" on page 3332	"client-id (DHCP Relay Agent)" on page 4317	"delimiter (DHCP Relay Agent)" on page 5290
"domain-name (DHCP Relay Agent)" on page 5886	"interface-description (DHCP Relay Agent)" on page 8856	"interface-name (DHCP Relay Agent)" on page 8938

(Continued)

[edit forwarding-options dhcp-relay dhcpv6 authentication username-include]		
"logical-system-name (DHCP Relay Agent)" on page 10083	"mac-address (DHCP Relay Agent)" on page 10223	"relay-agent-interface-id (DHCPv6 Relay Agent Username)" on page 14278
"relay-agent-remote-id (DHCPv6 Relay Agent Username)" on page 14286	"relay-agent-subscriber-id (DHCPv6 Relay Agent)" on page 14289	"routing-instance-name (DHCP Relay Agent)" on page 15947
"user-prefix (DHCP Relay Agent)" on page 30580		

[edit forwarding-options dhcp-relay dhcpv6 dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit forwarding-options dhcp-relay dhcpv6 group {group-name}]		
"active-server-group" on page 1817	"authentication (DHCP Relay Agent)" on page 2595	"dynamic-profile (DHCP Relay Agent)" on page 6119
"forward-only (DHCP Relay Agent)" on page 7396	"interface (DHCP Relay Agent)" on page 8863	"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701
"liveness-detection" on page 9846	"overrides (DHCP Relay Agent)" on page 12068	"relay-agent-interface-id (DHCPv6 Relay Agent)" on page 14276
"relay-agent-remote-id (DHCPv6 Relay Agent)" on page 14284	"relay-option (DHCP Relay Agent)" on page 14297	"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334

(Continued)

[edit forwarding-options dhcp-relay dhcpv6 group {group-name}]		
"route-suppression (DHCP Local Server and Relay Agent)" on page 15877	"service-profile (DHCP Relay Agent)" on page 16731	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982
"use-vlan-id (DHCP Relay Agent)" on page 30546		

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} (relay-agent-interface-id | relay-agent-remote-id)]		
"include-irb-and-l2" on page 8529	"no-vlan-interface-name" on page 11617	"prefix (DHCP Relay Agent)" on page 13077
"use-interface-description" on page 30522	"use-option-82" on page 30528	

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} authentication]		
"password (DHCP Relay Agent)" on page 12229	"username-include (DHCP Relay Agent)" on page 30597	

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} authentication username-include]		
"circuit-type (DHCP Relay Agent)" on page 3332	"client-id (DHCP Relay Agent)" on page 4317	"delimiter (DHCP Relay Agent)" on page 5290
"domain-name (DHCP Relay Agent)" on page 5886	"interface-description (DHCP Relay Agent)" on page 8856	"interface-name (DHCP Relay Agent)" on page 8938
"logical-system-name (DHCP Relay Agent)" on page 10083	"mac-address (DHCP Relay Agent)" on page 10223	"relay-agent-interface-id (DHCPv6 Relay Agent Username)" on page 14278

(Continued)

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} authentication username-include]		
"relay-agent-remote-id (DHCPv6 Relay Agent Username)" on page 14286	"relay-agent-subscriber-id (DHCPv6 Relay Agent)" on page 14289	"routing-instance-name (DHCP Relay Agent)" on page 15947
"user-prefix (DHCP Relay Agent)" on page 30580		

[edit forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit forwarding-options dhcp-relay dhcpv6 group { group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} interface {interface-name}]		
"dynamic-profile (DHCP Relay Agent)" on page 6119	"overrides (DHCP Relay Agent)" on page 12068	"service-profile (DHCP Relay Agent)" on page 16731
"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982	"trace (DHCP Relay Agent)" on page 29458	

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} interface {interface-name} overrides]		
"allow-snooped-clients" on page 2204	"asymmetric-lease-time (DHCP Overrides)" on page 2519	"asymmetric-prefix-lease-time (DHCP Overrides)" on page 2521
"client-negotiation-match (DHCPv6 Relay Agent)" on page 4335	"dual-stack (DHCP Relay Agent Overrides)" on page 6050	"no-allow-snooped-clients" on page 11390

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} lease-time-validation]		
"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699	"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} liveness-detection]		
"failure-action" on page 6787	"method" on page 10663	

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} liveness-detection method]		
"bfd" on page 2909	"layer2-liveness-detection (Send)" on page 9649	

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} liveness-detection method bfd]		
"detection-time" on page 5542	"holddown-interval" on page 8122	"minimum-interval" on page 10733
"minimum-receive-interval" on page 10752	"multiplier" on page 11122	"no-adaptation" on page 11373
"session-mode" on page 16828	"transmit-interval" on page 30085	"version (BFD)" on page 30695

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} liveness-detection method bfd detection-time]	
"threshold (detection-time)" on page 29323	

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} liveness-detection method bfd transmit-interval]		
"minimum-interval" on page 10733	"threshold (transmit-interval)" on page 29344	

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} overrides]		
"allow-snooped-clients" on page 2204	"asymmetric-lease-time (DHCP Overrides)" on page 2519	"asymmetric-prefix-lease-time (DHCP Overrides)" on page 2521
"client-negotiation-match (DHCPv6 Relay Agent)" on page 4335	"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279	"dual-stack (DHCP Relay Agent Overrides)" on page 6050
"interface-client-limit (DHCP Relay Agent)" on page 9108	"no-allow-snooped-clients" on page 11390	"no-bind-on-request (DHCP Relay Agent)" on page 11418
"send-release-on-delete (DHCP Relay Agent)" on page 16550		

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} relay-option]		
"default-action (DHCP Relay Agent Option)" on page 5177	"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865
"option-number (DHCP Relay Agent Option)" on page 11898		

[edit forwarding-options dhcp-relay dhcpv6 group {group-name} relay-option (default-action | equals | starts-with)]		
"drop (DHCP Relay Agent Option)" on page 5946	"forward-only (DHCP Relay Agent Option)" on page 7398	

[edit forwarding-options dhcp-relay dhcpv6 lease-time-validation]		
"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699	"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	

[edit forwarding-options dhcp-relay dhcpv6 liveness-detection]		
"failure-action" on page 6787	"method" on page 10663	

[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method]		
"bfd" on page 2909	"layer2-liveness-detection (Send)" on page 9649	

[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd]		
"detection-time" on page 5542	"holddown-interval" on page 8122	"minimum-interval" on page 10733
"minimum-receive-interval" on page 10752	"multiplier" on page 11122	"no-adaptation" on page 11373
"session-mode" on page 16828	"transmit-interval" on page 30085	"version (BFD)" on page 30695

[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd detection-time]	
"threshold (detection-time)" on page 29323	

[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval]		
"minimum-interval" on page 10733	"threshold (transmit-interval)" on page 29344	

[edit forwarding-options dhcp-relay dhcpv6 overrides]		
"allow-snooped-clients" on page 2204	"asymmetric-lease-time (DHCP Overrides)" on page 2519	"asymmetric-prefix-lease-time (DHCP Overrides)" on page 2521

(Continued)

[edit forwarding-options dhcp-relay dhcpv6 overrides]		
"client-negotiation-match (DHCPv6 Relay Agent)" on page 4335	"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279	"dual-stack (DHCP Relay Agent Overrides)" on page 6050
"interface-client-limit (DHCP Relay Agent)" on page 9108	"no-allow-snooped-clients" on page 11390	"no-bind-on-request (DHCP Relay Agent)" on page 11418
"send-release-on-delete (DHCP Relay Agent)" on page 16550		

[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]	
"keep-incoming-interface-id (DHCP Relay Agent)" on page 9462	

[edit forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id]	
"keep-incoming-remote-id (DHCP Relay Agent)" on page 9463	

[edit forwarding-options dhcp-relay dhcpv6 relay-option]		
"default-action (DHCP Relay Agent Option)" on page 5177	"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865
"option-number (DHCP Relay Agent Option)" on page 11898		

[edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action | equals | starts-with)]		
"drop (DHCP Relay Agent Option)" on page 5946	"forward-only (DHCP Relay Agent Option)" on page 7398	

[edit forwarding-options dhcp-relay dhcpv6 relay-option-vendor-specific]		
"host-name (DHCPv6 Relay Agent)" on page 8151	"location (DHCPv6 Relay Agent)" on page 9988	

[edit forwarding-options dhcp-relay dhcpv6],]	
"relay-option-vendor-specific (dhcpv6)" on page 14299	

[edit forwarding-options dhcp-relay dual-stack-grouop {dual-stack-group-name} authentication username-include]	
"user-prefix (DHCP Relay Agent)" on page 30580	

[edit forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"authentication (DHCP Relay Agent)" on page 2595	"classification-key (DHCP Relay Agent)" on page 3376	"dual-stack-interface-client-limit (DHCP Local Server and Relay Agent)" on page 6058
"dynamic-profile (DHCP Relay Agent)" on page 6119	"liveness-detection" on page 9846	"protocol-master" on page 13448
"relay-agent-interface-id (DHCPv6 Relay Agent)" on page 14276	"relay-agent-remote-id (DHCPv6 Relay Agent Username)" on page 14286	"service-profile (DHCP Relay Agent)" on page 16731
"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982	"use-vlan-id (DHCP Relay Agent)" on page 30546	

[edit forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name} authentication]		
"password (DHCP Relay Agent)" on page 12229	"username-include (DHCP Relay Agent)" on page 30597	

[edit forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name} authentication username-include]		
"circuit-type (DHCP Relay Agent)" on page 3332	"client-id (DHCP Relay Agent)" on page 4317	"delimiter (DHCP Relay Agent)" on page 5290
"domain-name (DHCP Relay Agent)" on page 5886	"interface-description (DHCP Relay Agent)" on page 8856	"interface-name (DHCP Relay Agent)" on page 8938
"logical-system-name (DHCP Relay Agent)" on page 10083	"routing-instance-name (DHCP Relay Agent)" on page 15947	

[edit forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name} liveness-detection]	
"method" on page 10663	

[edit forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name} liveness-detection method]	
"layer2-liveness-detection (Send)" on page 9649	

[edit forwarding-options dhcp-relay dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit forwarding-options dhcp-relay group {group-name}]		
"active-server-group" on page 1817	"authentication (DHCP Relay Agent)" on page 2595	"dynamic-profile (DHCP Relay Agent)" on page 6119
"forward-only (DHCP Relay Agent)" on page 7396	"interface (DHCP Relay Agent)" on page 8863	"interface-tag (DHCP Relay Server)" on page 9083
"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701	"liveness-detection" on page 9846	"overrides (DHCP Relay Agent)" on page 12068
"relay-option (DHCP Relay Agent)" on page 14297	"relay-option-82" on page 14294	"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334
"route-suppression (DHCP Local Server and Relay Agent)" on page 15877	"service-profile (DHCP Relay Agent)" on page 16731	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982
"trace (DHCP Relay Agent)" on page 29458		

[edit forwarding-options dhcp-relay group {group-name} authentication]		
"password (DHCP Relay Agent)" on page 12229	"username-include (DHCP Relay Agent)" on page 30597	

[edit forwarding-options dhcp-relay group {group-name} authentication username-include]		
"circuit-type (DHCP Relay Agent)" on page 3332	"client-id (DHCP Relay Agent)" on page 4317	"delimiter (DHCP Relay Agent)" on page 5290
"domain-name (DHCP Relay Agent)" on page 5886	"interface-description (DHCP Relay Agent)" on page 8856	"interface-name (DHCP Relay Agent)" on page 8938

(Continued)

[edit forwarding-options dhcp-relay group {group-name} authentication username-include]		
"logical-system-name (DHCP Relay Agent)" on page 10083	"mac-address (DHCP Relay Agent)" on page 10223	"option-82 (DHCP Relay Agent)" on page 11893
"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580	

[edit forwarding-options dhcp-relay group {group-name} dhcpv6 relay-option-vendor-specific]	
"location (DHCPv6 Relay Agent)" on page 9988	

[edit forwarding-options dhcp-relay group { group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit forwarding-options dhcp-relay group {group-name} interface {interface-name}]		
"dynamic-profile (DHCP Relay Agent)" on page 6119	"overrides (DHCP Relay Agent)" on page 12068	"service-profile (DHCP Relay Agent)" on page 16731
"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982		

[edit forwarding-options dhcp-relay group {group-name} interface {interface-name} overrides]		
"allow-no-end-option (DHCP Relay Agent)" on page 2196	"allow-snooped-clients" on page 2204	"apply-secondary-as-giaddr" on page 2433
"asymmetric-lease-time (DHCP Overrides)" on page 2519	"client-discover-match (DHCP Relay Agent)" on page 4310	"dual-stack (DHCP Relay Agent Overrides)" on page 6050

(Continued)

[edit forwarding-options dhcp-relay group {group-name} interface {interface-name} overrides]		
"no-allow-snooped-clients" on page 11390	"relay-source" on page 14302	

[edit forwarding-options dhcp-relay group {group-name} interface {interface-name} overrides user-defined-option82 {string}]	
"user-defined-option-82" on page 30551	

[edit forwarding-options dhcp-relay group {group-name} lease-time-validation]		
"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699	"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	

[edit forwarding-options dhcp-relay group {group-name} liveness-detection]		
"failure-action" on page 6787	"method" on page 10663	

[edit forwarding-options dhcp-relay group {group-name} liveness-detection method]		
"bfd" on page 2909	"layer2-liveness-detection (Send)" on page 9649	

[edit forwarding-options dhcp-relay group {group-name} liveness-detection method bfd]		
"detection-time" on page 5542	"holddown-interval" on page 8122	"minimum-interval" on page 10733
"minimum-receive-interval" on page 10752	"multiplier" on page 11122	"no-adaptation" on page 11373
"session-mode" on page 16828	"transmit-interval" on page 30085	"version (BFD)" on page 30695

[edit forwarding-options dhcp-relay group {group-name} liveness-detection method bfd detection-time]	
"threshold (detection-time)" on page 29323	

[edit forwarding-options dhcp-relay group {group-name} liveness-detection method bfd transmit-interval]		
"minimum-interval" on page 10733	"threshold (transmit-interval)" on page 29344	

[edit forwarding-options dhcp-relay group {group-name} overrides]		
"allow-no-end-option (DHCP Relay Agent)" on page 2196	"allow-snooped-clients" on page 2204	"always-write-giaddr" on page 2221
"always-write-option-82" on page 2222	"asymmetric-lease-time (DHCP Overrides)" on page 2519	"client-discover-match (DHCP Relay Agent)" on page 4310
"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279	"disable-relay" on page 5801	"dual-stack (DHCP Relay Agent Overrides)" on page 6050
"interface-client-limit (DHCP Relay Agent)" on page 9108	"layer2-unicast-replies" on page 9661	"no-allow-snooped-clients" on page 11390
"no-bind-on-request (DHCP Relay Agent)" on page 11418	"proxy-mode" on page 13833	"relay-source" on page 14302
"replace-ip-source-with (DHCP Relay Agent)" on page 14377	"send-release-on-delete (DHCP Relay Agent)" on page 16550	"trust-option-82" on page 30168

[edit forwarding-options dhcp-relay group {group-name} overrides user-defined-option82 {string}]	
"user-defined-option-82" on page 30551	

[edit forwarding-options dhcp-relay group {group-name} relay-option]		
"default-action (DHCP Relay Agent Option)" on page 5177	"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682
"option-60 (DHCP Relay Agent)" on page 11877	"option-77 (DHCP Relay Agent)" on page 11880	"option-number (DHCP Relay Agent Option)" on page 11898
"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506	

[edit forwarding-options dhcp-relay group {group-name} relay-option (default-action | equals | starts-with)]		
"drop (DHCP Relay Agent Option)" on page 5946	"forward-only (DHCP Relay Agent Option)" on page 7398	"local-server-group (DHCP Relay Agent Option)" on page 9976

[edit forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit forwarding-options dhcp-relay group {group-name} relay-option-60]	
"vendor-option" on page 30684	

[edit forwarding-options dhcp-relay group {group-name} relay-option-82]		
"circuit-id (DHCP Relay Agent)" on page 3324	"host-name (DHCP Relay Agent)" on page 8150	"remote-id (DHCP Relay Agent)" on page 14331

(Continued)

[edit forwarding-options dhcp-relay group {group-name} relay-option-82]		
"server-id-override" on page 16637	"vendor-specific (DHCP Relay Agent)" on page 30686	

[edit forwarding-options dhcp-relay group {group-name} relay-option-82 (circuit-id | remote-id)]		
"include-irb-and-l2" on page 8529	"no-vlan-interface-name" on page 11617	"prefix (DHCP Relay Agent)" on page 13077
"use-interface-description" on page 30522		

[edit forwarding-options dhcp-relay group {group-name} relay-option-82 vendor-specific]	
"location (DHCP Relay Agent)" on page 9985	

[edit forwarding-options dhcp-relay lease-time-validation]		
"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699	"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	

[edit forwarding-options dhcp-relay liveness-detection]		
"failure-action" on page 6787	"method" on page 10663	"session-mode" on page 16828

[edit forwarding-options dhcp-relay liveness-detection method]		
"bfd" on page 2909	"layer2-liveness-detection (Send)" on page 9649	

[edit forwarding-options dhcp-relay liveness-detection method bfd]		
"detection-time" on page 5542	"holddown-interval" on page 8122	"minimum-interval" on page 10733
"minimum-receive-interval" on page 10752	"multiplier" on page 11122	"no-adaptation" on page 11373
"transmit-interval" on page 30085	"version (BFD)" on page 30695	

[edit forwarding-options dhcp-relay liveness-detection method bfd detection-time]	
"threshold (detection-time)" on page 29323	

[edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval]		
"minimum-interval" on page 10733	"threshold (transmit-interval)" on page 29344	

[edit forwarding-options dhcp-relay overrides]		
"allow-no-end-option (DHCP Relay Agent)" on page 2196	"allow-snooped-clients" on page 2204	"always-write-giaddr" on page 2221
"always-write-option-82" on page 2222	"apply-secondary-as-giaddr" on page 2433	"asymmetric-lease-time (DHCP Overrides)" on page 2519
"client-discover-match (DHCP Relay Agent)" on page 4310	"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279	"disable-relay" on page 5801
"dual-stack (DHCP Relay Agent Overrides)" on page 6050	"interface-client-limit (DHCP Relay Agent)" on page 9108	"layer2-unicast-replies" on page 9661
"no-allow-snooped-clients" on page 11390	"no-bind-on-request (DHCP Relay Agent)" on page 11418	"proxy-mode" on page 13833

(Continued)

[edit forwarding-options dhcp-relay overrides]		
"relay-source" on page 14302	"replace-ip-source-with (DHCP Relay Agent)" on page 14377	"send-release-on-delete (DHCP Relay Agent)" on page 16550
"trust-option-82" on page 30168		

[edit forwarding-options dhcp-relay overrides user-defined-option82 {string}]	
"user-defined-option-82" on page 30551	

[edit forwarding-options dhcp-relay relay-option]		
"default-action (DHCP Relay Agent Option)" on page 5177	"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682
"option-60 (DHCP Relay Agent)" on page 11877	"option-77 (DHCP Relay Agent)" on page 11880	"option-number (DHCP Relay Agent Option)" on page 11898
"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506	

[edit forwarding-options dhcp-relay relay-option (default-action | equals | starts-with)]		
"drop (DHCP Relay Agent Option)" on page 5946	"local-server-group (DHCP Relay Agent Option)" on page 9976	

[edit forwarding-options dhcp-relay relay-option (default-action | equals | starts-with), [edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action | equals | starts-with), [edit forwarding-options dhcp-relay group {group-name} relay-option (default-action | equals | starts-with), [edit forwarding-options dhcp-relay dhcpv6 group {group-name} relay-option (default-action | equals | starts-with), [edit logical-systems {logical-system-name} forwarding-options dhcp-relay ...]	
"relay-server-group (DHCP Relay Agent Option)" on page 14300	

[edit forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit forwarding-options dhcp-relay relay-option-60]	
"vendor-option" on page 30684	

[edit forwarding-options dhcp-relay relay-option-82]		
"circuit-id (DHCP Relay Agent)" on page 3324	"host-name (DHCP Relay Agent)" on page 8150	"remote-id (DHCP Relay Agent)" on page 14331
"server-id-override" on page 16637	"vendor-specific (DHCP Relay Agent)" on page 30686	

[edit forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)]		
"include-irb-and-l2" on page 8529	"no-vlan-interface-name" on page 11617	"prefix (DHCP Relay Agent)" on page 13077

(Continued)

[edit forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)]	
"use-interface-description" on page 30522	

[edit forwarding-options dhcp-relay relay-option-82 circuit-id]		
"keep-incoming-circuit-id (DHCP Relay Agent)" on page 9460	"keep-incoming-remote-id (DHCP Relay Agent)" on page 9463	

[edit forwarding-options dhcp-relay relay-option-82 vendor-specific]	
"location (DHCP Relay Agent)" on page 9985	

[edit forwarding-options dhcp-relay traceoptions]	
"file (Extended DHCP Relay Agent and Helpers Trace Options)" on page 7056	

[dhcp-relaygroup]

[edit forwarding-options dhcp-relaygroup {group-name} dhcpv6 relay-option-vendor-specific]	
"host-name (DHCPv6 Relay Agent)" on page 8151	

[enhanced-hash-key]

[edit forwarding-options enhanced-hash-key]		
"conditional-match" on page 4504	"ecmp-resilient-hash" on page 6196	"fabric-load-balance" on page 6753
"flex-hashing" on page 7227	"hash-mode" on page 7997	"hash-parameters" on page 8000
"hash-seed" on page 8002	"hash-seed (profile)" on page 8003	"layer2 (enhanced-hash-key)" on page 9644

[edit forwarding-options enhanced-hash-key ecmp-dlb]		
"egress-quantization" on page 6238	"ether-type (Dynamic Load Balancing)" on page 6473	

[edit forwarding-options enhanced-hash-key fabric-load-balance flowlet]	
"inactivity-interval (Fabric Load Balance)" on page 8501	

[edit forwarding-options enhanced-hash-key family]		
"inet (enhanced-hash-key)" on page 8561	"inet6 (enhanced-hash-key)" on page 8602	

[edit forwarding-options enhanced-hash-key family inet]		
"gtp-tunnel-endpoint-identifier" on page 7938	"gtp-tunnel-endpoint-identifier" on page 7938	

[edit forwarding-options enhanced-hash-key family inet6]	
"gtp-tunnel-endpoint-identifier" on page 7938	

[edit forwarding-options enhanced-hash-key family mpls]	
"ether-pseudowire" on page 6471	

[edit forwarding-options enhanced-hash-key hash-mode]	
"gtp-header-offset" on page 7936	

[edit forwarding-options enhanced-hash-key inet6]	
"ipv6-flow-label" on page 9359	

[edit forwarding-options enhanced-hash-key vxlan]	
"no-inner-payload" on page 11474	

[evpn-vxlan]

[edit forwarding-options evpn-vxlan]	
"reroute-address" on page 15575	

[family]

[edit forwarding-options family (inet | inet6 | mpls | vpls) filter]	
"input (Forwarding Table)" on page 8703	

[edit forwarding-options family (inet | inet6 | mpls)]		
"filter (IPv4, IPv6, and MPLS)" on page 7144		

[edit forwarding-options family (inet | inet6 | mpls) filter]		
"output (Forwarding Table)" on page 11992		

[edit forwarding-options family (inet | inet6)]		
"dscp-mapping-classifier" on page 6016		

[edit forwarding-options family (inet | inet6) output interface {interface-name}]		
"next-hop (Forwarding Options)" on page 11336		

[edit forwarding-options family inet6]		
"route-accounting" on page 15846	"source-checking" on page 28190	

[edit forwarding-options family vpls]		
"filter (VPLS)" on page 7158	"flood" on page 7238	

[hash-key]

[edit forwarding-options hash-key]		
"family fcoe" on page 6863	"family inet" on page 6871	"family mpls" on page 6886
"family multiservice" on page 6893		

[edit forwarding-options hash-key family fcoe]		
"fabric-interfaces" on page 6751	"oxid" on page 12093	

[edit forwarding-options hash-key family fcoe ethernet-interfaces]	
"node-group (OxID Hash Control)" on page 11643	

[edit forwarding-options hash-key family fcoe ethernet-interfaces node-group ({node-group-name} | all)]	
"oxid" on page 12093	

[edit forwarding-options hash-key family fcoe fabric-interfaces]	
"node-group (OxID Hash Control)" on page 11643	

[edit forwarding-options hash-key family fcoe fabric-interfaces node-group ({node-group-name} | all)]	
"oxid" on page 12093	

[edit forwarding-options hash-key family inet]	
"l2tp-tunnel-session-identifier" on page 9558	

[edit forwarding-options hash-key family inet layer-4]	
"gtp-tunnel-endpoint-identifier" on page 7938	

[edit forwarding-options hash-key family inet6 layer-3]	
"ipv6-flow-label" on page 9359	

[edit forwarding-options hash-key family inet6 layer-4]	
"gtp-tunnel-endpoint-identifier" on page 7938	

[helpers]

[edit forwarding-options helpers]		
"bootp" on page 3012	"domain (DNS Forwarding Options)" on page 5874	"port (Packet Forwarding)" on page 12853
"tftp" on page 29251	"traceoptions (DNS, Port, and TFTP Packet Forwarding)" on page 29530	

[edit forwarding-options helpers bootp]		
"apply-secondary-as-giaddr" on page 2433	"client-response-ttl" on page 4336	"description (Forwarding Options)" on page 5330
"dhcp-option82" on page 5619	"interface (BOOTP)" on page 8841	"maximum-hop-count" on page 10527
"minimum-wait-time" on page 10758	"relay-agent-option" on page 14279	"server (DHCP and BOOTP Relay Agent)" on page 16623

(Continued)

[edit forwarding-options helpers bootp]		
"source-address-giaddr" on page 28135		

[edit forwarding-options helpers bootp dhcp-option82 remote-id]		
"prefix (Remote ID for Option 82)" on page 13101		

[edit forwarding-options helpers bootp interface]		
"apply-secondary-as-giaddr" on page 2433	"dhcp-option82" on page 5619	"source-address-giaddr" on page 28135

[edit forwarding-options helpers bootp interface ({interface-name} | {interface-group})]		
"client-response-ttl" on page 4336	"description (Forwarding Options)" on page 5330	"maximum-hop-count" on page 10527
"minimum-wait-time" on page 10758	"no-listen" on page 11494	"server (DHCP and BOOTP Relay Agent)" on page 16623

[edit forwarding-options helpers bootp interface {interface-name} dhcp-option82 remote-id]		
"prefix (Remote ID for Option 82)" on page 13101		

[edit forwarding-options helpers domain]		
"description (Forwarding Options)" on page 5330	"interface (DNS, Port, and TFTP Packet Forwarding or Relay Agent)" on page 8868	"server (DNS, Port, and TFTP Service)" on page 16624

[edit forwarding-options helpers domain interface {interface-name}]		
"description (Forwarding Options)" on page 5330	"no-listen" on page 11494	"server (DNS, Port, and TFTP Service)" on page 16624

[edit forwarding-options helpers port {port-number}]		
"description (Forwarding Options)" on page 5330	"interface (DNS, Port, and TFTP Packet Forwarding or Relay Agent)" on page 8868	"server (DNS, Port, and TFTP Service)" on page 16624

[edit forwarding-options helpers port {port-number} interface {interface-name}]		
"description (Forwarding Options)" on page 5330	"no-listen" on page 11494	"server (DNS, Port, and TFTP Service)" on page 16624

[edit forwarding-options helpers tftp]		
"description (Forwarding Options)" on page 5330	"interface (DNS, Port, and TFTP Packet Forwarding or Relay Agent)" on page 8868	"server (DNS, Port, and TFTP Service)" on page 16624

[edit forwarding-options helpers tftp interface {interface-name}]		
"description (Forwarding Options)" on page 5330	"no-listen" on page 11494	"server (DNS, Port, and TFTP Service)" on page 16624

[edit forwarding-options helpers traceoptions]	
"file (Extended DHCP Relay Agent and Helpers Trace Options)" on page 7056	

[helperstraceoptions]

[edit forwarding-options helperstraceoptions file]		
"files (Sampling and Traceoptions)" on page 7107	"size (Sampling and Traceoptions)" on page 28003	

[interface]

[edit forwarding-options interface {interface-name}]	
"mark-interface (RA Guard)" on page 10353	

[edit forwarding-options interface ({interface-name} | all)]		
"auto-dad (SLAAC Snooping)" on page 2705	"maximum-allowed-contentions " on page 10499	

[edit forwarding-options interface ({interface-name} | {interface-range-name})]		
"stateful (Router Advertisement Guard) " on page 28514	"stateless" on page 28524	

[load-balance]

[edit forwarding-options load-balance]		
"ecmp-local-bias" on page 6195	"indexed-load-balance" on page 8545	"per-flow" on page 12413

(Continued)

[edit forwarding-options load-balance]	
"per-prefix" on page 12417	

[monitoring]

[edit forwarding-options monitoring {name}]		
"family (Monitoring)" on page 6885	"source-address (Forwarding Options)" on page 28133	

[edit forwarding-options monitoring {name} family inet]	
"output (Monitoring)" on page 12000	

[edit forwarding-options monitoring {group-name} family inet output]	
"flow-export-destination (Flow Collection)" on page 7285	

[edit forwarding-options monitoring {name} family inet output cflowd {hostname}]	
"port (Flow Monitoring)" on page 12823	

[edit forwarding-options monitoring {group-name} family inet output flow-server {hostname}]	
"port (cflowd)" on page 12813	

[edit forwarding-options monitoring {name} inet output]		
"cflowd (Flow Monitoring)" on page 3270	"interface (Monitoring)" on page 8923	

[edit forwarding-options monitoring {name} output]		
"export-format (Flow Monitoring)" on page 6667	"flow-active-timeout (Flow Monitoring)" on page 7261	"flow-inactive-timeout (Flow Monitoring)" on page 7296

[edit forwarding-options monitoring {name} output interface {interface-name}]		
"engine-id (Forwarding Options)" on page 6348	"engine-type" on page 6352	"input-interface-index" on page 8705
"output-interface-index" on page 11995		

[next-hop]

[edit forwarding-options next-hop]		
"arp-detect" on page 2463	"arp-throttle" on page 2482	

[next-hop-group]

[edit forwarding-options next-hop-group]		
"group-type (Unknown Unicast Forwarding)" on page 7908	"interface (Next-Hop Group)" on page 8939	

[pfe-sensor]

[edit forwarding-options pfe-sensor npu-memory]	
"resource-list (Telemetry)" on page 15612	

[port-mirror]

[edit forwarding-options port-mirror]		
"disable (Forwarding Options)" on page 5756	"disable-all-instances" on page 5744	

[edit forwarding-options port-mirror instance {instance-name}]	
"disable (Forwarding Options)" on page 5756	

[port-mirroring]

[edit forwarding-options port-mirroring]		
"family (Port Mirroring)" on page 6896	"input (Port Mirroring)" on page 8714	"instance" on page 8748
"instance (Port Mirroring)" on page 8762	"mirror-once" on page 10772	"traceoptions (Forwarding Options)" on page 29565

(Continued)

[edit forwarding-options port-mirroring]	
"traceoptions (Port Mirroring and Traffic Sampling)" on page 29640	

[edit forwarding-options port-mirroring family (any | ccc | inet | inet6 | mpls | vpls)]	
"output (Port Mirroring)" on page 12004	

[edit forwarding-options port-mirroring family (inet | inet6 | ccc | vpls) output]	
"no-filter-check" on page 11460	

[edit forwarding-options port-mirroring family (inet | inet6)]	
"output (Port Mirroring)" on page 12003	

[edit forwarding-options port-mirroring family (inet | inet6) output]		
"interface (Port Mirroring)" on page 8956	"interface (Port Mirroring)" on page 8954	"no-filter-check" on page 11460

[edit forwarding-options port-mirroring family (inet | vpls) output]	
"next-hop-group (Port Mirroring)" on page 11343	

[edit forwarding-options port-mirroring family (inet|inet6) input]	
"run-length" on page 16163	

[edit forwarding-options port-mirroring family ethernet-switching output]		
"vlan (Port Mirroring)" on page 30841		

[edit forwarding-options port-mirroring family ethernet-switching output vlan (vlan-name | vlan-id)]		
"no-tag" on page 11594		

[edit forwarding-options port-mirroring family inet6 output]		
"next-hop-group" on page 11337		

[edit forwarding-options port-mirroring input]		
"maximum-packet-length" on page 10540	"rate (Forwarding Options)" on page 14060	"run-length" on page 16163

[edit forwarding-options port-mirroring instance {instance-name}]		
"input (Port Mirroring)" on page 8714	"output (Port Mirroring)" on page 12003	

[edit forwarding-options port-mirroring [instance {name}] family]		
"ethernet-switching (Port Mirroring)" on page 6513	"inet (Port Mirroring)" on page 8576	

[edit forwarding-options port-mirroring instance {instance-name} family (any | ccc | inet | inet6 | mpls | vpls)]		
"output (Port Mirroring)" on page 12004		

[edit forwarding-options port-mirroring instance {instance-name} family (inet | inet6 | mpls) output]		
"server-profile (Active Flow Monitoring)" on page 16655		

[edit forwarding-options port-mirroring instance {instance-name} family (inet | vpls) output]		
"next-hop-group (Port Mirroring)" on page 11343		

[edit forwarding-options port-mirroring instance {name} family ethernet-switching]		
"output" on page 11980		

[edit forwarding-options port-mirroring [instance {name}] family ethernet-switching output]		
"interface (Port Mirroring)" on page 8952	"vlan (Port Mirroring)" on page 30839	

[edit forwarding-options port-mirroring [instance {name}] family ethernet-switching output interface {name}]		
"ip-address (Port Mirroring)" on page 9238		

[edit forwarding-options port-mirroring [instance {name}] family inet output interface {name}]		
"routing-instance (Port Mirroring)" on page 15953		

[edit forwarding-options port-mirroring instance {instance-name} input]		
"maximum-packet-length" on page 10540	"run-length" on page 16163	

[edit forwarding-options port-mirroring instance instance-name family ethernet-switching output]	
"vlan (Port Mirroring)" on page 30841	

[edit forwarding-options port-mirroring instance instance-name family ethernet-switching output vlan (vlan-name | vlan-id)]	
"no-tag" on page 11594	

[edit forwarding-options port-mirroring output]	
"interface (Port Mirroring)" on page 8956	

[edit forwarding-options port-mirroring traceoptions]		
"file (Trace Options)" on page 7090	"no-remote-trace (Trace Options)" on page 11566	

[edit forwarding-options port-mirroring traceoptions file]		
"files" on page 7103	"files (Sampling and Traceoptions)" on page 7107	"match (Forwarding-Options)" on page 10382
"size" on page 27997	"size (Sampling and Traceoptions)" on page 28003	"world-readable (Forwarding Options)" on page 31142

[sampling]

[edit forwarding-options sampling]		
"disable (Forwarding Options)" on page 5756	"family (Sampling)" on page 6913	"input (Sampling)" on page 8718

(Continued)

[edit forwarding-options sampling]		
"instance (Sampling)" on page 8766	"pre-rewrite-tos" on page 13000	"sample-once" on page 16287
"traceoptions (Forwarding Options)" on page 29565	"traceoptions (Port Mirroring and Traffic Sampling)" on page 29640	

[edit forwarding-options sampling family (inet | inet6 | mpls | vpls | bridge) output]	
"flow-server" on page 7344	

[edit forwarding-options sampling family (inet | inet6 | mpls | vpls)]	
"disable (Forwarding Options)" on page 5756	

[edit forwarding-options sampling family (inet | inet6 | mpls | vpls) output file]	
"disable (Forwarding Options)" on page 5756	

[edit forwarding-options sampling family (inet | inet6 | mpls) output]	
"interface (Accounting or Sampling)" on page 8831	

[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server {hostname}]		
"port (cflowd)" on page 12813	"port (Flow Monitoring)" on page 12823	"routing-instance (Flow Monitoring)" on page 15929

[edit forwarding-options sampling family (inet | inet6 | mpls) output interface {interface-name}]		
"engine-id (Forwarding Options)" on page 6348	"engine-type" on page 6352	

[edit forwarding-options sampling family (inet |inet6 |mpls | vpls)]	
"output (Sampling)" on page 12010	

[edit forwarding-options sampling family (inet |inet6 |mpls |vpls) output]	
"flow-active-timeout (Flow Monitoring)" on page 7261	

[edit forwarding-options sampling family (inet |inet6 |mpls) output]		
"aggregate-export-interval" on page 2062	"flow-inactive-timeout (Flow Monitoring)" on page 7296	

[edit forwarding-options sampling family (inet |inet6 |mpls) output file]		
"filename" on page 7096	"files" on page 7103	"size" on page 27997
"stamp (Sampling)" on page 28484	"world-readable (Forwarding Options)" on page 31142	

[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server {hostname}]		
"aggregation (Forwarding Options)" on page 2084	"autonomous-system-type (Flow Monitoring)" on page 2767	"local-dump (Flow Aggregation)" on page 9938
"version (Flow Aggregation)" on page 30692		

[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server {hostname} version9]		
"template (Forwarding Options)" on page 29079		

[edit forwarding-options sampling family (inet |inet6 |mpls) output interface {interface-name}]		
"source-address (Forwarding Options)" on page 28133		

[edit forwarding-options sampling family (inet |inet6 |mpls| vpls |bridge) output flow-server {hostname}]		
"version9 (Forwarding Options)" on page 30724		

[edit forwarding-options sampling family inet output]		
"file (Sampling)" on page 7075		

[edit forwarding-options sampling family {family-name} output file]		
"filename (Sampling)" on page 7101	"files (Sampling and Traceoptions)" on page 7107	"size (Sampling and Traceoptions)" on page 28003

[edit forwarding-options sampling input]		
"max-packets-per-second (Sampling)" on page 10452	"maximum-packet-length" on page 10540	"rate (Forwarding Options)" on page 14060
"run-length" on page 16163		

[edit forwarding-options sampling instance {instance-name}]		
"disable (Forwarding Options)" on page 5756	"family (Sampling)" on page 6913	"input (Sampling)" on page 8718

[edit forwarding-options sampling instance {instance-name} family (bridge | inet |inet6 |mpls vpls) output flow-server {hostname}]	
"version9 (Forwarding Options)" on page 30724	

[edit forwarding-options sampling instance {instance-name} family (inet | inet6 | bridge)]	
"family (Sampling)" on page 6913	

[edit forwarding-options sampling instance {instance-name} family (inet | inet6 | mpls | vpls | bridge) output]	
"flow-server" on page 7344	

[edit forwarding-options sampling instance {instance-name} family (inet | inet6 | mpls | vpls) output flow-server {hostname}]	
"version-ipfix (Forwarding Options)" on page 30707	

[edit forwarding-options sampling instance family (inet | inet6 | mpls | vpls) output flow-server {hostname} version-ipfix]	
"template (Forwarding Options Version IPFIX)" on page 29080	

[edit forwarding-options sampling instance {instance-name} family (inet | inet6 | mpls) output]	
"interface (Accounting or Sampling)" on page 8831	

[edit forwarding-options sampling instance {sampling-instance-name} family (inet | inet6 | mpls) output flow-server {hostname}]	
"routing-instance (Flow Monitoring)" on page 15929	

[edit forwarding-options sampling instance {instance-name} family (inet | inet6 | mpls) output interface {interface-name}]		
"engine-id (Forwarding Options)" on page 6348	"engine-type" on page 6352	

[edit forwarding-options sampling instance {instance-name} family (inet | inet6) output flow-server {hostname}]		
"dscp (flow-server)" on page 6003	"forwarding-class (Sampling)" on page 7439	

[edit forwarding-options sampling instance {instance-name} family (inet |inet6 |mpls | vpls) output]		
"flow-active-timeout (Flow Monitoring)" on page 7261	"flow-inactive-timeout (Flow Monitoring)" on page 7296	

[edit forwarding-options sampling instance {instance-name} family (inet |inet6 |mpls | vpls) output interface {interface-name}]	
"source-address (Forwarding Options)" on page 28133	

[edit forwarding-options sampling instance {instance-name} family (inet |inet6 |mpls)]	
"output (Sampling)" on page 12010	

[edit forwarding-options sampling instance {instance-name} family (inet |inet6 |mpls) output]		
"aggregate-export-interval" on page 2062		

[edit forwarding-options sampling instance {instance-name} family (inet |inet6 |mpls) output flow-server {hostname}]		
"aggregation (Forwarding Options)" on page 2084	"autonomous-system-type (Flow Monitoring)" on page 2767	"local-dump (Flow Aggregation)" on page 9938
"port (Flow Monitoring)" on page 12823	"version (Flow Aggregation)" on page 30692	

[edit forwarding-options sampling instance {instance-name} family (inet |inet6 |mpls) output flow-server {hostname} version9]		
"template (Forwarding Options)" on page 29079		

[edit forwarding-options sampling instance {instance-name} family inet output]		
"inline-jflow" on page 8669	"inline-jflow (Forwarding Options)" on page 8671	

[edit forwarding-options sampling instance {instance-name} family inet output inline-jflow]		
"flow-export-rate" on page 7286	"source-address (Forwarding Options)" on page 28133	

[edit forwarding-options sampling instance {instance- name} family inet6 output]		
"inline-jflow (Forwarding Options)" on page 8671		

[edit forwarding-options sampling instance {instance-name} input]		
"max-packets-per-second (Sampling)" on page 10452	"maximum-packet-length" on page 10540	"rate (Forwarding Options)" on page 14060
"run-length" on page 16163		

[edit forwarding-options sampling traceoptions]		
"file (Trace Options)" on page 7090	"no-remote-trace (Trace Options)" on page 11566	

[edit forwarding-options sampling traceoptions file]		
"files" on page 7103	"files (Sampling and Traceoptions)" on page 7107	"match (Forwarding-Options)" on page 10382
"size" on page 27997	"size (Sampling and Traceoptions)" on page 28003	

[edit forwarding-options sampling traceoptionsfile]	
"world-readable (Forwarding Options)" on page 31142	

[satellite]

[edit forwarding-options satellite]	
"local-replication" on page 9974	

[edit forwarding-options satellite fpc]	
"selective-vlan-switching" on page 16532	

[sflow]

[edit forwarding-options sflow]	
"flow-sample" on page 7360	

[storm-control]

[edit forwarding-options storm-control]	
"enhanced" on page 6354	

[storm-control-profiles]

[edit forwarding-options storm-control-profiles {profile-name} all]		
"bandwidth-level" on page 2851	"bandwidth-percentage" on page 2869	"burst-size" on page 3078

[vlans]

[edit forwarding-options vlans {vlan-name}]		
"stateful (Router Advertisement Guard) " on page 28514	"stateless" on page 28524	

[vxlan-routing]

[edit forwarding-options vxlan-routing]		
"interface-num" on page 8943	"next-hop (VXLAN Routing)" on page 11356	"overlay-ecmp" on page 12041

[edit fpc]

IN THIS SECTION

- [pic] | 260

Sub Hierarchies

"[edit fpc {slot} pic {pic-number}]" on page 260

[edit fpc {slot}]		
"pic" on page 12480	"pic" on page 12480	

[pic]

[edit fpc {slot} pic {pic-number}]		
"fte (Port)" on page 7615	"sfpplus" on page 16915	"xe (Port)" on page 31152
"xle (Port)" on page 31155		

[edit fpc {slot} pic {pic-number} sfpplus]	
"pic-mode" on page 12490	

[edit hierarchical-policer]

IN THIS SECTION

- [aggregate] | **261**
- [premium] | **261**

Sub Hierarchies

"[edit hierarchical-policer aggregate]" on page 261

"[edit hierarchical-policer premium]" on page 261

[aggregate]

[edit hierarchical-policer {hierarchical-policer-name} aggregate if-exceeding]	
"pps-limit (Hierarchical Policer)" on page 12982	

[edit hierarchical-policer {hierarchical-policer-name} aggregate if-exceeding-pps]	
"packet-burst (Hierarchical Policer)" on page 12133	

[premium]

[edit hierarchical-policer {hierarchical-policer-name} premium if-exceeding]	
"pps-limit (Hierarchical Policer)" on page 12982	

[edit hierarchical-policer {hierarchical-policer-name} premium if-exceeding-pps]	
"packet-burst (Hierarchical Policer)" on page 12133	

[edit interface]

IN THIS SECTION

- [(all | {interface-name})] | 262
- [gigether-options] | 263
- [unit] | 263

Sub Hierarchies

"[edit interface (all | {interface-name})]" on page 262

"[edit interface {interface-name} gigether-options]" on page 263

"[edit interface unit]" on page 263

[edit interface {interface-name}]	
"Unused" on page 30435	

[(all | {interface-name})]

[edit interface (all | {interface-name})]	
"multicast" on page 11062	

[gigether-options]

[edit interface {interface-name} gigether-options]	
"no-auto-mdix" on page 11409	

[unit]

[edit interface {interface-name} unit {logical-unit-number} family inet service]		
"input (Interfaces)" on page 8707	"output (Service Set)" on page 11978	

[edit interface {ppp interface} unit {logical-unit-number} ppp-max-payload {ppp-max-payload}]	
"pppoe-options" on page 12974	

[edit interfaces]

IN THIS SECTION

- [(fxp0 | em0)] | 275
- [(ls-] | 275
- [(ml-] | 276
- [(QFX Series)] | 276
- [(rsp0 | rsp1)] | 276
- [advisory-options] | 277
- [ae] | 277
- [aeaggregated-ether-options] | 277

- [aeaggregated-sonet-options] | **279**
- [aeesi] | **279**
- [aeforwarding-options] | **280**
- [aeunit] | **280**
- [aex] | **280**
- [aggregated-ether-options] | **281**
- [aggregated-ethernet-options] | **282**
- [asaggregated-sonet-options] | **283**
- [asi] | **283**
- [asiaggregated-inline-services-options] | **283**
- [at-atm-options] | **283**
- [at-dsl-options] | **284**
- [at-unit] | **284**
- [atm-options] | **285**
- [auto-configure] | **285**
- [br-] | **287**
- [br-dialer-options] | **288**
- [br-isdn-options] | **288**
- [br-unit] | **288**
- [br-{pim}/0/{port}isdn-options] | **288**
- [ce1-] | **289**
- [cellular-options] | **289**
- [ct1-] | **290**
- [demux0] | **290**
- [dlunit] | **291**
- [ds0-options] | **291**
- [e1-options] | **292**
- [e3-options] | **292**
- [es-options] | **292**
- [es-unit] | **292**
- [esi] | **293**
- [et-] | **293**

- [et-fpc/pic/port] | **293**
- [ether-options] | **294**
- [family] | **295**
- [fastether-options] | **295**
- [fibrechannel-options] | **296**
- [forwarding-class-accounting] | **297**
- [fti] | **297**
- [ftiunit] | **298**
- [ge-] | **298**
- [ge-gigether-options] | **298**
- [ge-unit] | **299**
- [ge-{pim}/0/0] | **299**
- [gigether-options] | **300**
- [gr-fpc/pic/port] | **303**
- [gr-fpc/pic/port.unit] | **303**
- [gr-unit] | **303**
- [gre] | **304**
- [interface-name] | **304**
- [interface-range] | **305**
- [interface-set] | **307**
- [ip-fpc/pic/port] | **308**
- [irb] | **308**
- [lacp] | **309**
- [layer2-policer] | **309**
- [link-degrade-monitor] | **309**
- [lo0] | **309**
- [load-balancing-options] | **310**
- [logical-tunnel-options] | **311**
- [ls-] | **311**
- [ls-mlfr-uni-nni-bundle-options] | **311**
- [ls-unit] | **312**
- [lsq-] | **312**

- [lsq-lsq-failure-options] | **312**
- [lsq-unit] | **313**
- [mapping-range] | **313**
- [media-type] | **313**
- [mif] | **314**
- [mlfr-uni-nni-bundle-options] | **314**
- [mo-] | **315**
- [mo-multiservice-options] | **315**
- [multi-chassis-protection] | **315**
- [multiservice-options] | **315**
- [optics-options] | **316**
- [otn-options] | **316**
- [pic-set] | **318**
- [pp0] | **318**
- [ppp-options] | **319**
- [ps] | **320**
- [ps0] | **320**
- [pstargeted-options] | **321**
- [radius-options] | **321**
- [redundant-ether-options] | **321**
- [rlsq] | **322**
- [rlsqredundancy-options] | **322**
- [rlt] | **323**
- [rms] | **323**
- [rmsredundancy-options] | **323**
- [rsp] | **323**
- [rspredundancy-options] | **324**
- [se-] | **324**
- [serial-options] | **324**
- [services-options] | **325**
- [shdsl-options] | **327**
- [si] | **327**

- [si-] | 327
- [si-unit] | 328
- [so-unit] | 328
- [sonet-options] | 328
- [sp-unit] | 329
- [speed] | 329
- [t1-options] | 330
- [t3-options] | 330
- [tdm-options] | 330
- [umd0] | 330
- [unit] | 331
- [unitradio-router] | 350
- [vlan] | 350
- [xe-] | 351
- [xe-,edit] | 351
- [xe-fpc/pic/port] | 351
- [xe-unit] | 352
- [{interface-name] | 352
- [{interface-name}redundant-ether-options] | 352
- [{interface-name}unit] | 352

Sub Hierarchies

"[edit interfaces (fxp0 | em0)]" on page 275

"[edit interfaces (ls-{fpc}/{pic}/{port} | ml-{fpc}/{pic}/{port}) unit {logical-unit-number}]" on page 275

"[edit interfaces (ml-{fpc}/{pic}/{port}| ls-{fpc}/{pic}/{port}) unit {logical-unit-number}]" on page 276

"[edit interfaces (QFX Series) {interface-name}]" on page 276

"[edit interfaces (rsp0 | rsp1)]" on page 276

"[edit interfaces {interface-name} {logical-unit-number} advisory-options]" on page 277

"[edit interfaces ae{x}]" on page 277

"[edit interfaces ae{X} aggregated-ether-options]" on page 277

"[edit interfaces ae{x} aggregated-sonet-options]" on page 279

"[edit interfaces ae{X} esi]" on page 279

"[edit interfaces ae{X} {unit logical-unit-number} forwarding-options]" on page 280

"[edit interfaces aeunit]" on page 280

"[edit interfaces aex]" on page 280

"[edit interfaces {name} aggregated-ether-options]" on page 281

"[edit interfaces aggregated-ethernet-options]" on page 282

"[edit interfaces as{x} aggregated-sonet-options]" on page 283

"[edit interfaces asi{x}]" on page 283

"[edit interfaces asi{x} aggregated-inline-services-options]" on page 283

"[edit interfaces at-{fpc}/{pic}/{port} atm-options]" on page 283

"[edit interfaces at-{fpc}/{pic}/{port} dsl-options]" on page 284

"[edit interfaces at-{fpc}/{pic}/{port} unit {logical-unit-number}]" on page 284

"[edit interfaces atm-options]" on page 285

"[edit interfaces {interface-name} auto-configure]" on page 285

"[edit interfaces br-{pim}/0/{port}]" on page 287

"[edit interfaces br-{pim}/0/{port} dialer-options]" on page 288

"[edit interfaces br-{pim}/0/{port} isdn-options]" on page 288

"[edit interfaces br-unit]" on page 288

"[edit interfaces br-{pim}/0/{port}isdn-options]" on page 288

"[edit interfaces ce1-{fpc}/{pic}/{port}]" on page 289

"[edit interfaces {interface-name} cellular-options]" on page 289

"[edit interfaces ct1-{fpc}/{pic}/{port}]" on page 290

"[edit interfaces demux0]" on page 290

"[edit interfaces dl{n} unit {logical-unit-number}]" on page 291

"[edit interfaces {interface-name} ds0-options]" on page 291

"[edit interfaces {interface-name} e1-options]" on page 292

"[edit interfaces {interface-name} e3-options]" on page 292

"[edit interfaces {interface-name} es-options]" on page 292

"[edit interfaces es-unit]" on page 292

"[edit interfaces {name} esi]" on page 293

"[edit interfaces et-{fpc}/{pic}/{port}]" on page 293

"[edit interfaces et-fpc/pic/port]" on page 293

"[edit interfaces {interface-name} ether-options]" on page 294

"[edit interfaces family]" on page 295

"[edit interfaces {interface-name} fastether-options]" on page 295

"[edit interfaces {interface-name} fibrechannel-options]" on page 296

"[edit interfaces {interface-name} forwarding-class-accounting]" on page 297

"[edit interfaces fti]" on page 297

"[edit interfaces ftiunit]" on page 298

"[edit interfaces ge-{pim}/0/0 switch-options switch-port {port-number}]" on page 298

"[edit interfaces ge-{fpc}/{pic}/{port} gigether-options]" on page 298

"[edit interfaces ge-{fpc}/{pic}/{slot} unit {logical-unit-number}]" on page 299

"[edit interfaces ge-{pim}/0/0]" on page 299

"[edit interfaces {interface-name} gigether-options]" on page 300

"[edit interfaces gr-fpc/pic/port]" on page 303

"[edit interfaces gr-fpc/pic/port.unit]" on page 303

"[edit interfaces gr-{fpc}/{pic}/{port} unit {logical-unit-number}]" on page 303

"[edit interfaces gre]" on page 304

"[edit interfaces interface-name]" on page 304

"[edit interfaces interface-range {name}]" on page 305

"[edit interfaces interface-set {interface-set-name}]" on page 307

"[edit interfaces ip-fpc/pic/port]" on page 308

"[edit interfaces irb]" on page 308

"[edit interfaces {name} lacp]" on page 309

"[edit interfaces {interface-name} layer2-policer]" on page 309

"[edit interfaces {interfaces-name} link-degrade-monitor]" on page 309

"[edit interfaces lo0]" on page 309

"[edit interfaces {interface-name} load-balancing-options]" on page 310

"[edit interfaces {name} logical-tunnel-options]" on page 311

"[edit interfaces ls-{fpc}/{pic}/{port}:{channel}]" on page 311

"[edit interfaces ls-{fpc}/{pic}/{port}:{channel} mlfr-uni-nni-bundle-options]" on page 311

"[edit interfaces ls-{fpc}/{pic}/{port}:{channel} unit {logical-unit-number}]" on page 312

"[edit interfaces lsq-{fpc}/{pic}/{port}]" on page 312

"[edit interfaces lsq-{fpc}/{pic}/{port} lsq-failure-options]" on page 312

"[edit interfaces lsq-{fpc}/{pic}/{port} unit {logical-unit-number}]" on page 313

"[edit interfaces mapping-range]" on page 313

"[edit interfaces {interface-name} media-type]" on page 313

"[edit interfaces mif]" on page 314

"[edit interfaces {interface-name} mlfr-uni-nni-bundle-options]" on page 314

"[edit interfaces mo-{fpc}/{pic}/{port}]" on page 315

"[edit interfaces mo-{fpc}/{pic}/{port} multiservice-options]" on page 315

"[edit interfaces {name} multi-chassis-protection {IP address of MC-LAG peer}]" on page 315

"[edit interfaces {interface-name} multiservice-options]" on page 315

"[edit interfaces {interface-name} optics-options]" on page 316

"[edit interfaces {interface-name} otn-options]" on page 316

"[edit interfaces pic-set {pic-set-name}]" on page 318

"[edit interfaces pp0]" on page 318

"[edit interfaces {interface-name} ppp-options]" on page 319

"[edit interfaces ps {number}]" on page 320

"[edit interfaces ps0]" on page 320

"[edit interfaces ps{number} targeted-options]" on page 321

"[edit interfaces {interface-name} radius-options]" on page 321

"[edit interfaces {interface-name} redundant-ether-options]" on page 321

"[edit interfaces rlsq {number}:{number}]" on page 322

"[edit interfaces rlsq{number} redundancy-options]" on page 322

"[edit interfaces rlt]" on page 323

"[edit interfaces rms{number}]" on page 323

"[edit interfaces rms{number} redundancy-options]" on page 323

"[edit interfaces rsp{number}]" on page 323

"[edit interfaces rsp{number} redundancy-options]" on page 324

"[edit interfaces se-{pim}/0/{port}]" on page 324

"[edit interfaces {interface-name} serial-options]" on page 324

"[edit interfaces {interface-name} services-options]" on page 325

"[edit interfaces {interface-name} shdsl-options]" on page 327

"[edit interfaces si]" on page 327

"[edit interfaces si-{slot}/{pic}/{port}]" on page 327

"[edit interfaces si-{fpc}/{pic}/{port} unit {logical-unit-number}]" on page 328

"[edit interfaces so-{fpc}/{pic}/{slot} unit {logical-unit-number}]" on page 328

"[edit interfaces sonet-options]" on page 328

"[edit interfaces sp-{fpc}/{pic}/{port} unit {logical-unit-number}]" on page 329

"[edit interfaces {interface-name} speed]" on page 329

"[edit interfaces {interface-name} t1-options]" on page 330

"[edit interfaces {interface-name} t3-options]" on page 330

"[edit interfaces {interface-name} tdm-options]" on page 330

"[edit interfaces umd0]" on page 330

"[edit interfaces {interface-name} unit {logical-unit-number}]" on page 331

"[edit interfaces unitradio-router]" on page 350

"[edit interfaces vlan]" on page 350

"[edit interfaces xe-{fpc}/{pic}/{port}]" on page 351

"[edit interfaces xe-,edit]" on page 351

"[edit interfaces xe-fpc/pic/port]" on page 351

"[edit interfaces xe-{fpc}/{pic}/{slot} unit {logical-unit-number}]" on page 352

"[edit interfaces {interface-name]" on page 352

"[edit interfaces {interface-name}redundant-ether-options]" on page 352

"[edit interfaces {interface-name}unit]" on page 352

[edit interfaces {interface-name}]		
"accounting-profile" on page 1727	"alarm low-light-alarm" on page 2105	"atm-options" on page 2524
"auto-configure" on page 2692	"bypass-queuing-chip" on page 3114	"cascade-port" on page 3214
"cellular-options" on page 3249	"compression" on page 4491	"damping (Interfaces)" on page 4818
"damping (Interfaces)" on page 4818	"demux0 (Static Interface)" on page 5313	"description (Interfaces)" on page 5334
"disable (Interface)" on page 5765	"dsl-sfp-options" on page 6030	"encapsulation (interfaces)" on page 6289
"encapsulation (Physical Interface)" on page 6307	"encoding (QFX Series)" on page 6314	"es-options" on page 6451
"esi" on page 6453	"ether-options" on page 6461	"fabric-options" on page 6763

(Continued)

[edit interfaces {interface-name}]		
"family bridge" on page 6847	"fastether-options" on page 6949	"fibrechannel-options" on page 7015
"flexible-vlan-tagging (Interfaces)" on page 7237	"flow-control" on page 7267	"forwarding-class-accounting" on page 7413
"gigether-options" on page 7685	"gigether-options (Chassis Cluster)" on page 7694	"gratuitous-arp-reply" on page 7789
"hierarchical-scheduler" on page 8060	"hierarchical-scheduler (Subscriber Interfaces on MX Series Routers)" on page 8062	"hold-time (ae interface) " on page 8098
"hold-time (Redundant Ethernet Interfaces)" on page 8115	"input-native-vlan-push" on page 8710	"interface-range (QFX Series)" on page 9003
"interface-set (Ethernet Interfaces)" on page 9044	"interface-set (IP Demux Interfaces)" on page 9046	"interface-transmit-statistics" on page 9090
"irb (Interfaces)" on page 9389	"isolated" on page 9406	"keepalives" on page 9475
"layer2-policer (Hierarchical Policer)" on page 9656	"led-beacon" on page 9710	"link-degrade-monitor" on page 9784
"load-balancing-options (Aggregated Multiservices)" on page 9888	"logical-tunnel-options" on page 10093	"mac (MPLS)" on page 10217
"media-type (Dual-Purpose Uplink Ports)" on page 10600	"mif (TDF Interface)" on page 10703	"modulation-format" on page 10826
"mtu" on page 11026	"mtu (interfaces)" on page 11027	"mtu (Multilink and Link Services Logical Interface)" on page 11025

(Continued)

[edit interfaces {interface-name}]		
"mtu (QFX Series)" on page 11023	"multi-chassis-protection (Interfaces)" on page 11045	"multicast-statistics" on page 11089
"native-vlan-id" on page 11240	"network-isolation-profile" on page 11314	"no-gratuitous-arp-request" on page 11467
"no-keepalives" on page 11488	"no-native-vlan-insert" on page 11517	"no-per-unit-scheduler" on page 11545
"number-of-sub-ports (interface)" on page 11740	"oam-on-svlan (Ethernet Interfaces)" on page 11758	"optics-options" on page 11847
"passive-monitor-mode" on page 12212	"per-unit-scheduler" on page 12423	"ppp-options" on page 12949
"promiscuous" on page 13372	"promiscuous-mode (Interfaces)" on page 13374	"radius-options (Interfaces)" on page 14011
"redundancy-group (Redundant Tunnel)" on page 14180	"redundancy-options (Aggregated Multiservices)" on page 14194	"redundancy-options (Stateful Synchronization)" on page 14197
"redundant-ether-options" on page 14209	"schedulers (Interfaces)" on page 16373	"services-options" on page 16787
"services-options (Next Gen Services Interfaces)" on page 16789	"shared-scheduler" on page 16971	"shdsl-options" on page 16974
"speed (Ethernet)" on page 28349	"speed (Ethernet)" on page 28349	"stacked-vlan-tagging" on page 28475
"targeted-options (PS Interface)" on page 28968	"tdm-options (Interfaces)" on page 29039	"traceoptions (Individual Interfaces)" on page 29576
"traceoptions (Individual Interfaces)" on page 29576	"transport-monitoring" on page 30126	"traps" on page 30142

(Continued)

[edit interfaces {interface-name}]		
"traps" on page 30142	"unidirectional" on page 30356	"unit (Aggregated Multiservices)" on page 30380
"unit (Interfaces Services) " on page 30399	"unit (Interfaces)" on page 30401	"unit (interfaces)" on page 30370
"unit (Link-Multilink)" on page 30402	"unit (Sampling)" on page 30366	"unit (Tunnel)" on page 30404
"unit (Voice Services)" on page 30408	"vlan-tagging" on page 30861	"vlan-tagging (Dynamic)" on page 30864

[(fxp0 | em0)]

[edit interfaces (fxp0 | em0) unit {logical-unit-number} family {family} address]	
"master-only" on page 10364	

[(ls-]

[edit interfaces (ls-{fpc}/{pic}/{port} | ml-{fpc}/{pic}/{port}) unit {logical-unit-number}]		
"drop-timeout" on page 5981	"fragment-threshold" on page 7542	"short-sequence" on page 16989

[(ml-]

[edit interfaces (ml-{fpc}/{pic}/{port}| ls-{fpc}/{pic}/{port}) unit {logical-unit-number}]	
"mrru" on page 10989	

[(QFX Series)]

[edit interfaces (QFX Series) {interface-name}]	
"unit (QFX Series)" on page 30367	

[edit interfaces (QFX Series) interface-range {interface-range-name}]	
"unit (QFX Series)" on page 30367	

[edit interfaces (QFX Series) {interface-name} unit 0 family ethernet-switching vlan]	
"members" on page 10625	

[(rsp0 | rsp1)]

[edit interfaces (rsp0 | rsp1) redundancy-options]		
"primary (Adaptive Services Interfaces)" on page 13129	"secondary (Adaptive Services Interfaces)" on page 16408	

[advisory-options]

[edit interfaces {interface-name} {logical-unit-number} advisory-options]	
"upstream-rate (Traffic Shaping)" on page 30479	

[ae]

[edit interfaces ae{x}]		
"aggregated-ether-options" on page 2078	"flexible-vlan-tagging" on page 7235	"hold-time (ae interface) " on page 8098

[edit interfaces ae {number} aggregated-ether-options]	
"rebalance-periodic (Aggregated Ethernet Subscriber Interfaces)" on page 14108	

[aeaggregated-ether-options]

[edit interfaces ae{X} aggregated-ether-options]		
"lacp (802.3ad)" on page 9607	"lacp (Aggregated Ethernet)" on page 9610	"link-protection" on page 9803
"link-protection" on page 9803	"link-protection (non-LACP)" on page 9810	"link-protection-sub-group (aggregated-ether-options)" on page 9821

(Continued)

[edit interfaces ae{X} aggregated-ether-options]		
"link-speed" on page 9823	"link-speed" on page 9823	"link-speed (Aggregated Ethernet)" on page 9826
"local-bias (ae)" on page 9921	"local-minimum-links-threshold" on page 9962	"logical-interface-chassis-redundancy (MX Series Virtual Chassis)" on page 10063
"logical-interface-fpc-redundancy (Aggregated Ethernet Subscriber Interfaces)" on page 10067	"mc-ae" on page 10575	"minimum-bandwidth (aggregated Ethernet)" on page 10728
"minimum-links" on page 10741	"minimum-links" on page 10741	"resilient-hash" on page 15589

[edit interfaces ae{x} aggregated-ether-options ethernet-switch-profile]	
"mac-learn-enable" on page 10233	

[edit interfaces ae{x} aggregated-ether-options lacp]		
"hold-time up" on page 8117	"periodic" on page 12438	"system-id (Interfaces)" on page 28845
"system-priority (interfaces)" on page 28857		

[edit interfaces ae{X} aggregated-ether-options lacp link-protection]		
"disable (Link Protection)" on page 5767	"non-revertive (Interfaces)" on page 11662	"revertive" on page 15734

[edit interfaces ae{x} aggregated-ether-options load-balance]	
"local-bias (ae load-balance)" on page 9923	

[edit interfaces ae{X} aggregated-ether-options mc-ae]	
"enhanced-convergence" on page 6357	

[aeaggregated-sonet-options]

[edit interfaces ae{x} aggregated-sonet-options]	
"minimum-links" on page 10741	

[aeesi]

[edit interfaces ae{X} esi]	
"auto-derive" on page 2707	

[aeforwarding-options]

[edit interfaces ae{X} {unit logical-unit-number} forwarding-options]	
"load-balance-stateful (Aggregated Ethernet Interfaces)" on page 9886	

[edit interfaces ae{X} {unit logical-unit-number} forwarding-options load-balance-stateful]		
"load-type (Aggregated Ethernet Interfaces)" on page 9896	"per-flow (Aggregated Ethernet Interfaces)" on page 12414	

[edit interfaces ae{X} {unit logical-unit-number} forwarding-options load-balance-stateful per-flow]	
"rebalance (Aggregated Ethernet Interfaces)" on page 14106	

[aeunit]

[edit interfaces ae{X} unit {logical-unit-number} esi]	
"auto-derive" on page 2707	

[aex]

[edit interfaces aex aggregated-ether-options]	
"dlb" on page 5841	

[edit interfaces aex aggregated-ether-options bfd-liveness-detection]		
"authentication (LAG)" on page 2622	"detection-time (LAG)" on page 5549	"transmit-interval (LAG)" on page 30090

[aggregated-ether-options]

[edit interfaces {name} aggregated-ether-options]		
"bfd-liveness-detection (LAG)" on page 2926	"ethernet-switch-profile" on page 6499	"ethernet-switch-profile" on page 6499
"fcoe-lag" on page 6963	"flow-control (ACX Series)" on page 7265	"lacp (Aggregated Ethernet)" on page 9610
"load-balance" on page 9874	"loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)" on page 10117	"loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)" on page 10117
"no-fcoe-lag" on page 11457	"optimize-aggregate-frr" on page 11853	"source-address-filter" on page 28128
"source-filtering" on page 28208	"targeted-options (Manual Targeting)" on page 28966	

[edit interfaces {name} aggregated-ether-options bfd-liveness-detection]	
"damping (BFD Liveness Detection)" on page 4815	

[edit interfaces {interface-name} aggregated-ether-options ethernet-switch-profile]		
"ethernet-policer-profile" on page 6488	"tag-protocol-id TPID(TPIDs Expected to Be Sent or Received)" on page 28922	"tag-protocol-id TPID(TPIDs Expected to Be Sent or Received)" on page 28922

[edit interfaces {name} aggregated-ether-options lacp]	
"sync-reset" on page 28785	

[edit interfaces {name} aggregated-ether-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[edit interfaces aggregated-ether-options mc-ae]		
"chassis-id (Aggregated Ethernet Interface)" on page 3302	"mc-ae-id" on page 10581	"mode (QFX Series)" on page 10818
"status-control" on page 28604		

[aggregated-ethernet-options]

[edit interfaces {interface-name} aggregated-ethernet-options lacp]		
"force-up" on page 7378	"force-up (EX and QFX Series)" on page 7376	

[asaggregated-sonet-options]

[edit interfaces as{x} aggregated-sonet-options]	
"link-speed (Aggregated SONET/SDH)" on page 9828	

[asi]

[edit interfaces asi{x}]		
"aggregated-inline-services-options (Aggregated Inline Services)" on page 2082	"l2tp-maximum-session (Service Interfaces)" on page 9551	

[asiaggregated-inline-services-options]

[edit interfaces asi{x} aggregated-inline-services-options]		
"primary-interface (Aggregated Inline Services)" on page 13138	"secondary-interface (Aggregated Inline Services)" on page 16413	

[at-atm-options]

[edit interfaces at-{fpc}/{pic}/{port} atm-options]	
"vpi (Define Virtual Path)" on page 30923	

[edit interfaces at-{fpc}/{pic}/{port} atm-options vpi {vpi-identifier} shaping]	
"cbr" on page 3238	

[at-dsl-options]

[edit interfaces at-{fpc}/{pic}/{port} dsl-options]	
"operating-mode" on page 11844	

[at-unit]

[edit interfaces at-{fpc}/{pic}/{port} unit {logical-unit-number}]	
"vci" on page 30660	

[edit interfaces at-{fpc}/{pic}/{port} unit {logical-unit-number} address {address} family {family} multipoint-destination {address} shaping]	
"cbr" on page 3238	

[edit interfaces at-{fpc}/{pic}/{port} unit {logical-unit-number} family {family} address {address} multipoint-destination {address}]	
"vci" on page 30660	

[edit interfaces at-{fpc}/{pic}/{port} unit {logical-unit-number} shaping]	
"cbr" on page 3238	

[atm-options]

[edit interfaces {interface-name} atm-options mpls]	
"pop-all-labels" on page 12800	

[edit interfaces {interface-name} atm-options mpls pop-all-labels]	
"required-depth" on page 15574	

[edit interfaces {interface-name} atm-options vpi {vpi-identifier}]	
"shaping" on page 16925	

[auto-configure]

[edit interfaces {interface-name} auto-configure]		
"remove-when-no-subscribers" on page 14370	"stacked-vlan-ranges" on page 28471	"vlan-ranges" on page 30845

[edit interfaces {interface-name} auto-configure stacked-vlan-ranges]		
"authentication" on page 2574	"dynamic-profile (Stacked VLAN)" on page 6131	"override" on page 12054

[edit interfaces {interface-name} auto-configure stacked-vlan-ranges authentication]		
"packet-types (Dynamic VLAN Authentication)" on page 12163	"password (Interfaces)" on page 12233	"username-include (Interfaces)" on page 30600

[edit interfaces {interface-name} auto-configure stacked-vlan-ranges authentication username-include]		
"delimiter" on page 5285	"domain-name" on page 5879	"option-18 (Interface-ID for DHCPv6 Autosense VLANs)" on page 11871
"option-37 (Relay Agent Remote-ID for DHCPv6 Autosense VLANs)" on page 11874	"option-82 (interfaces)" on page 11886	"radius-realm" on page 14017
"user-prefix" on page 30577		

[edit interfaces {interface-name} auto-configure stacked-vlan-ranges authentication username-include],]		
"circuit-type" on page 3329	"interface-name" on page 8935	"mac-address (VLAN and Stacked VLAN Interfaces)" on page 10227

[edit interfaces {interface-name} auto-configure stacked-vlan-ranges dynamic-profile {profile-name}]		
"accept (Dynamic Profile)" on page 1629	"access-profile (Dynamic Stacked VLAN)" on page 1693	"ranges (Dynamic Stacked VLAN)" on page 14051

[edit interfaces {interface-name} auto-configure vlan-ranges]		
"authentication" on page 2574	"dynamic-profile (VLAN)" on page 6133	"override" on page 12054

[edit interfaces {interface-name} auto-configure vlan-ranges authentication]		
"packet-types (Dynamic VLAN Authentication)" on page 12163	"password (Interfaces)" on page 12233	"username-include (Interfaces)" on page 30600

[edit interfaces {interface-name} auto-configure vlan-ranges authentication username-include]		
"circuit-id (VLAN Authentication Username)" on page 3328	"circuit-type" on page 3329	"delimiter" on page 5285

(Continued)

[edit interfaces {interface-name} auto-configure vlan-ranges authentication username-include]		
"domain-name" on page 5879	"interface-name" on page 8935	"mac-address (VLAN and Stacked VLAN Interfaces)" on page 10227
"option-18 (Interface-ID for DHCPv6 Autosense VLANs)" on page 11871	"option-37 (Relay Agent Remote-ID for DHCPv6 Autosense VLANs)" on page 11874	"option-82 (interfaces)" on page 11886
"radius-realm" on page 14017	"remote-id (VLAN Authentication Username)" on page 14338	"user-prefix" on page 30577

[edit interfaces {interface-name} auto-configure vlan-ranges dynamic-profile {profile-name}]		
"accept (Dynamic Profile)" on page 1629	"accept-out-of-band" on page 1636	"access-profile (Dynamic VLAN)" on page 1694
"ranges (Dynamic VLAN)" on page 14052		

[br-]

[edit interfaces br-{pim}/0/{port}]	
"isdn-options" on page 9397	

[br-dialer-options]

[edit interfaces br-{pim}/0/{port} dialer-options]	
"pool" on page 12762	

[br-isdn-options]

[edit interfaces br-{pim}/0/{port} isdn-options]		
"calling-number" on page 3189	"spid1" on page 28388	"spid2" on page 28389
"static-tei-val" on page 28584	"switch-type" on page 28757	"t310" on page 28889

[br-unit]

[edit interfaces br-{pim}/0/{port} unit {logical-unit-number} dialer-options]	
"dial-string" on page 5687	

[br-{pim}/0/{port}isdn-options]

[edit interfaces br-{pim}/0/{port}isdn-options]	
"tei-option" on page 29047	

[ce1-]

[edit interfaces ce1-{fpc}/{pic}/{port}]		
"bert-algorithm" on page 2900	"bert-error-rate" on page 2903	"bert-period" on page 2905
"isdn-options" on page 9397		

[cellular-options]

[edit interfaces {interface-name} cellular-options]		
"gsm-options" on page 7926	"roaming-mode" on page 15813	

[edit interfaces {interface-name} cellular-options gsm-options]		
"profiles" on page 13362	"select-profile" on page 16527	

[edit interfaces {interface-name} cellular-options gsm-options profiles {profile-name}]		
"access-point name" on page 1687	"authentication-method (Interfaces)" on page 2629	"sip-password" on page 27976
"sip-user-id" on page 27983		

[ct1-]

[edit interfaces ct1-{fpc}/{pic}/{port}]		
"bert-algorithm" on page 2900	"bert-error-rate" on page 2903	"bert-period" on page 2905
"isdn-options" on page 9397		

[demux0]

[edit interfaces demux0 unit {logical-unit-number}]		
"advisory-options (Traffic Shaping)" on page 2035	"targeted-distribution (Static Interfaces over Aggregated Ethernet)" on page 28961	

[edit interfaces demux0 unit {logical-unit-number} advisory-options]		
"downstream-rate (Traffic Shaping)" on page 5935	"upstream-rate (Traffic Shaping)" on page 30479	

[edit interfaces demux0 unit {logical-unit-number} demux-options]	
"underlying-interface" on page 30330	

[edit interfaces demux0 unit {logical-unit-number} family pppoe]	
"short-cycle-protection (Static and Dynamic Subscribers)" on page 16984	

[dlunit]

[edit interfaces dl{n} unit {logical-unit-number}]		
"dialer-options" on page 5689		

[edit interfaces dl{n} unit {logical-unit-number} dialer-options]		
"activation-delay" on page 1797	"callback-wait-period" on page 3183	"deactivation-delay" on page 5144
"idle-timeout" on page 8303	"incoming-map" on page 8540	"initial-route-check" on page 8658
"load-interval" on page 9892	"load-threshold" on page 9894	"pool" on page 12762
"redial-delay" on page 14157	"watch-list" on page 31026	

[edit interfaces dl{n} unit {logical-unit-number} dialer-options incoming-map]		
"callback" on page 3182	"caller" on page 3188	

[ds0-options]

[edit interfaces {interface-name} ds0-options]		
"bert-algorithm" on page 2900	"bert-error-rate" on page 2903	"bert-period" on page 2905

[e1-options]

[edit interfaces {interface-name} e1-options]		
"bert-algorithm" on page 2900	"bert-error-rate" on page 2903	"bert-period" on page 2905

[e3-options]

[edit interfaces {interface-name} e3-options]		
"bert-algorithm" on page 2900	"bert-error-rate" on page 2903	"bert-period" on page 2905

[es-options]

[edit interfaces {interface-name} es-options]	
"backup-interface" on page 2796	

[es-unit]

[edit interfaces es-{fpc}/{pic}/{port} unit {logical-unit-number} family inet]	
"ipsec-sa" on page 9307	

[esi]

[edit interfaces {name} esi]		
"df-election-granularity" on page 5574	"df-election-type" on page 5576	"per-esi" on page 12406
"per-esi-vlan" on page 12408		

[edit interfaces {interface-name} esi df-election-granularity]	
"per-evi" on page 12409	

[et-]

[edit interfaces et-{fpc}/{pic}/{port}]	
"flexible-vlan-tagging" on page 7235	

[et-fpc/pic/port]

[edit interfaces et-fpc/pic/port]	
"otn-options" on page 11955	

[ether-options]

[edit interfaces {interface-name} ether-options]		
"802.3ad (EX and QFX Series)" on page 1579	"802.3ad (EX and QFX Series)" on page 1579	"auto-negotiation" on page 2719
"auto-negotiation (QFX Series)" on page 2718	"autostate-exclude" on page 2769	"configured-flow-control" on page 4523
"ethernet-switch-profile" on page 6499	"fec (ether)" on page 6992	"flow-control (ACX Series)" on page 7265
"ieee-802-3az-eee" on page 8364	"link-mode" on page 9800	"loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)" on page 10117
"loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)" on page 10117	"mdi-mode" on page 10587	"speed (Ethernet)" on page 28349

[edit interfaces {interface-name} ether-options 802.3ad]		
"lacp (802.3ad)" on page 9607	"lacp (802.3ad)" on page 9607	"link-protection-sub-group (802.3ad)" on page 9820

[edit interfaces {interface-name} ether-options 802.3ad lacp]		
"force-up (EX and QFX Series)" on page 7376	"force-up (EX and QFX Series)" on page 7376	

[edit interfaces {interface-name} ether-options configured-flow-control]		
"rx-buffers" on page 16164	"tx-buffers" on page 30263	

[edit interfaces {interface-name} ether-options ethernet-switch-profile]		
"storm-control" on page 28620	"tag-protocol-id TPID(TPIDs Expected to Be Sent or Received)" on page 28922	

[edit interfaces {interface-name} ether-options ethernet-switch-profile ethernet-policer-profile]		
"input-priority-map" on page 8716	"output-priority-map" on page 12007	

[edit interfaces {interface-name} ether-options ethernet-switch-profile ethernet-policer-profile input-priority-map]	
"ieee802.1p" on page 8366	

[family]

[edit interfaces {interface-name} family inet]	
"tunnel-termination" on page 30245	

[fastether-options]

[edit interfaces {interface-name} fastether-options]		
"802.3ad" on page 1582	"flow-control (ACX Series)" on page 7265	"flow-control (Interfaces)" on page 7269

(Continued)

[edit interfaces {interface-name} fastether-options]		
"ignore-l3-incompletes" on page 8400	"ingress-rate-limit" on page 8646	"loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)" on page 10117
"redundant-parent (Interfaces Fast Ethernet)" on page 14214	"redundant-parent (Interfaces)" on page 14212	"source-address-filter" on page 28128
"source-filtering" on page 28208		

[edit interfaces {interface-name} fastether-options 802.3ad]	
"lacp (802.3ad) (MX Series)" on page 9609	

[edit interfaces {interface-name} fastether-options mpls]	
"pop-all-labels" on page 12800	

[edit interfaces {interface-name} fastether-options mpls pop-all-labels]	
"required-depth" on page 15574	

[fibrechannel-options]

[edit interfaces {interface-name} fibrechannel-options]		
"bb-sc-n" on page 2889	"loopback (Fibre Channel Interface)" on page 10120	"speed (Fibre Channel Interfaces)" on page 28370

[forwarding-class-accounting]

[edit interfaces {interface-name} forwarding-class-accounting]	
"enhanced (forwarding-class-accounting)" on page 6358	

[fti]

[edit interfaces fti {number} unit {logical-unit-number}]	
"tunnel (fti)" on page 30204	

[edit interfaces fti {number} unit {logical-unit-number} family inet {address}]	
"clear-dont-fragment-bit (fti)" on page 3556	

[edit interfaces fti {number} unit {logical-unit-number} tunnel]		
"encapsulation (interfaces-fti)" on page 6288	"encapsulation-udp-fti" on page 6313	

[edit interfaces fti {number} unit {logical-unit-number} tunnel encapsulation]	
"gre" on page 7790	

[edit interfaces fti {number} unit {logical-unit-number} tunnel encapsulation (gre | ipip | udp)]		
"destination" on page 5364	"source (GRE Encapsulation) " on page 28111	"tunnel-termination" on page 30245

[edit interfaces fti {number} unit {logical-unit-number} tunnel encapsulation (gre|udp|ipip)]	
"allow-fragmentation (fti)" on page 2181	

[ftiunit]

[edit interfaces fti{number} unit {logical-unit-number} tunnelencapsulation (gre | ipip | udp)]	
"key" on page 9481	

[ge-]

[edit interfaces ge-{pim}/0/0 switch-options switch-port {port-number}]		
"auto-negotiation" on page 2719	"flexible-vlan-tagging" on page 7235	"otn-options" on page 11955
"speed (Ethernet)" on page 28349	"speed (MX Series DPC)" on page 28377	

[ge-gigether-options]

[edit interfaces ge-{fpc}/{pic}/{port} gigether-options]	
"asynchronous-notification" on page 2522	

[ge-unit]

[edit interfaces ge-{fpc}/{pic}/{slot} unit {logical-unit-number}]	
"interface-shared-with" on page 9049	

[edit interfaces ge-{fpc}/{slot}/{port} unit {logical-unit-number} family]		
"ccc" on page 3241	"ethernet-switching" on page 6501	

[edit interfaces ge-{chassis}/{slot}/{port} unit {logical-unit-number}] family [edit security flow]	
"ethernet-switching" on page 6503	

[edit interfaces ge-{chassis}/{slot}/{port} unit {logical-unit-number} family ethernet-switching]	
"l3-interface (VLAN)" on page 9568	

[ge-{pim}/0/0]

[edit interfaces ge-{pim}/0/0]	
"switch-options (Interfaces)" on page 28748	

[edit interfaces ge-{pim}/0/0 switch-options]	
"switch-port" on page 28753	

[gigether-options]

[edit interfaces {interface-name} gigether-options]		
"802.3ad" on page 1582	"auto-negotiation" on page 2719	"ethernet-switch-profile" on page 6499
"fec (gigether)" on page 6995	"flow-control (ACX Series)" on page 7265	"flow-control (Interfaces)" on page 7269
"ignore-l3-incompletes" on page 8400	"loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)" on page 10117	"loopback (Local and Remote)" on page 10125
"loopback-remote" on page 10128	"mru (Gigabit Ethernet Options)" on page 10991	"mru (Gigabit Ethernet Options)" on page 10991
"redundant-parent (Interfaces Gigabit Ethernet)" on page 14215	"redundant-parent (Interfaces)" on page 14212	"source-address-filter" on page 28128
"source-filtering" on page 28208	"speed (Gigabit Ethernet interface - MX Series)" on page 28373	"speed (Gigabit Ethernet interface)" on page 28372

[edit interfaces {interface-name} gigether-options 802.3ad]	
"lacp (802.3ad) (MX Series)" on page 9609	

[edit interfaces {interface-name} gigether-options 802.3ad lacp]	
"port-priority" on page 12856	

[edit interfaces {interface-name} gigether-options ethernet-switch-profile]		
"ethernet-policer-profile" on page 6488	"mac-learn-enable" on page 10233	"tag-protocol-id TPID(TPIDs Expected to Be Sent or Received)" on page 28922

[edit interfaces {interface-name} gigether-options ethernet-switch-profile ethernet-policer-profile]		
"input-priority-map" on page 8716	"output-priority-map" on page 12007	"policer (CoS)" on page 12633

[edit interfaces {interface-name} gigether-options ethernet-switch-profile ethernet-policer-profile input-priority-map]	
"ieee802.1p" on page 8366	

[edit interfaces {interface-name} gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map]	
"classifier" on page 3380	

[edit interfaces {interface-name} gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier]	
"premium (Output Priority Map)" on page 13112	

[edit interfaces {interface-name} gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier premium]	
"forwarding-class (Gigabit Ethernet IQ Classifier)" on page 7430	

[edit interfaces {interface-name} gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier premium forwarding-class {class-name}]		
"loss-priority" on page 10142		

[edit interfaces {interface-name} gigether-options ethernet-switch-profile ethernet-policer-profile policer {cos-policer-name}]		
"aggregate (Gigabit Ethernet CoS Policer)" on page 2063	"premium (Policer)" on page 13113	

[edit interfaces {interface-name} gigether-options ethernet-switch-profile ethernet-policer-profile policer {cos-policer-name} aggregate]		
"bandwidth-limit (Policer for Gigabit Ethernet Interfaces)" on page 2860	"burst-size-limit (Policer for Gigabit Ethernet Interfaces)" on page 3089	

[edit interfaces {interface-name} gigether-options ethernet-switch-profile ethernet-policer-profile policer {cos-policer-name} premium]		
"bandwidth-limit (Policer for Gigabit Ethernet Interfaces)" on page 2860	"burst-size-limit (Policer for Gigabit Ethernet Interfaces)" on page 3089	

[edit interfaces {interface-name} gigether-options mpls]	
"pop-all-labels" on page 12800	

[edit interfaces {interface-name} gigether-options mpls pop-all-labels]	
"required-depth" on page 15574	

[gr-fpc/pic/port]

[edit interfaces gr-fpc/pic/port unit logical-unit-number tunnel]	
"path-mtu-discovery (Tunnel)" on page 12257	

[gr-fpc/pic/port.unit]

[edit interfaces gr-fpc/pic/port.unit]	
"family bridge (GRE Interfaces)" on page 6849	

[gr-unit]

[edit interfaces gr-{fpc}/{pic}/{port} unit {logical-unit-number}]		
"clear-dont-fragment-bit (Interfaces GRE Tunnels) " on page 3557	"copy-tos-to-outer-ip-header" on page 4666	"copy-tos-to-outer-ip-header" on page 4668
"copy-tos-to-outer-ip-header-transit" on page 4669	"reassemble-packets" on page 14100	

[edit interfaces gr-{fpc}/{pic}/{port} unit {logical-unit-number} tunnel]		
"allow-fragmentation" on page 2179	"do-not-fragment" on page 5871	

[gre]

[edit interfaces gre unit {logical-unit-number}]		
"copy-tos-to-outer-ip-header" on page 4666	"copy-tos-to-outer-ip-header" on page 4668	

[interface-name]

[edit interfaces interface-name atm-options vpi vpi-identifier shaping]		
"cbr rate" on page 3240	"vbr rate" on page 30657	

[edit interfaces interface-name radio-router]		
"bandwidth (Interfaces)" on page 2849	"credit (Interfaces)" on page 4712	"resource (Interfaces)" on page 15608
"threshold (Interfaces)" on page 29325		

[edit interfaces interface-name traceoptions]		
"flag (Interfaces)" on page 7206	"traceoptions (Interfaces)" on page 29587	

[edit interfaces interface-name unit logical-unit-number]	
"radio-router (Interfaces)" on page 13991	

[edit interfaces interface-name unit logical-unit-number family family dhcpv6-client prefix-delegating]		
"preferred-prefix-length" on page 13059	"sub-prefix-length" on page 28674	

[edit interfaces interface-name unit logical-unit-number radio-router]		
"latency (Interfaces)" on page 9638	"quality (Interfaces)" on page 13914	

[edit interfaces interface-name vdsl-options]		
"vdsl-profile" on page 30669		

[interface-range]

[edit interfaces interface-range {name}]		
"accounting-profile" on page 1727	"ether-options" on page 6461	"gratuitous-arp-reply" on page 7789
"hold-time (ae interface) " on page 8098	"member-range" on page 10618	"mtu (interfaces)" on page 11027
"mtu (QFX Series)" on page 11023	"traps" on page 30142	"traps" on page 30142

[edit interfaces interface-range {name} aggregated-ether-options]		
"flow-control (ACX Series)" on page 7265	"link-speed" on page 9823	"link-speed (Aggregated Ethernet)" on page 9826
"load-balance" on page 9874	"minimum-links" on page 10741	

[edit interfaces interface-range {name} aggregated-ether-options lacp]		
"periodic" on page 12438		

[edit interfaces interface-range {name} aggregated-ether-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[edit interfaces interface-range {name} aggregated-sonet-options]		
"link-speed" on page 9823	"link-speed (Aggregated Ethernet)" on page 9826	"minimum-links" on page 10741

[edit interfaces interface-range {name} ether-options]		
"flow-control (ACX Series)" on page 7265	"loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)" on page 10117	"mdi-mode" on page 10587

[edit interfaces interface-range {name} logical-tunnel-options]		
"load-balance" on page 9874		

[edit interfaces interface-range {name} logical-tunnel-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[edit interfaces interface-range {interface-name} unit {logical-unit-number} family]		
"family (ELS)" on page 6824	"inet (interfaces (QFX Series))" on page 8565	"inet6 (interfaces (QFX Series))" on page 8607

[edit interfaces interface-range {interface-range-name} unit {logical-unit-number} family inet]	
"targeted-broadcast" on page 28956	

[edit interfaces interface-range {name} unit {name} family inet dhcp]	
"client-identifier (DHCP Client)" on page 4319	

[interface-set]

[edit interfaces interface-set {interface-set-name}]	
"interface (Hierarchical CoS Schedulers)" on page 8896	

[edit interfaces interface-set {interface-set-name} interface {interface-name}]		
"unit (interfaces)" on page 30370	"vlan-tags-outer" on page 30872	

[edit interfaces interface-set {interface-set-name} targeted-options]	
"weight (Targeted Subscribers)" on page 31097	

[ip-fpc/pic/port]

[edit interfaces ip-fpc/pic/port unit logical-unit-number tunnel]	
"path-mtu-discovery (Tunnel)" on page 12257	

[irb]

[edit interfaces irb]		
"arp-l2-validate" on page 2470		

[edit interfaces irb unit {unit-number}]		
"enhanced-convergence" on page 6357	"proxy-macip-advertisement" on page 13831	

[edit interfaces irb unit <unit>]	
"enhanced-scale (IRB Interface)" on page 6175	

[edit interfaces irb unit {name} family inet address {<local-ip-address>/<prefix-length>} | arp {<local-ip-address>} mac {<local-ip-address>}]	
"flood-as-unknown-unicast" on page 7239	

[lacp]

[edit interfaces {name} lacp]	
"aggregate-wait-time" on page 2073	

[layer2-policer]

[edit interfaces {interface-name} layer2-policer]	
"input-hierarchical-policer" on page 8704	

[link-degrade-monitor]

[edit interfaces {interfaces-name} link-degrade-monitor]		
"recovery" on page 14146	"thresholds" on page 29346	

[lo0]

[edit interfaces lo0]	
"redundant-pseudo-interface-options" on page 14218	

[load-balancing-options]

[edit interfaces {interface-name} load-balancing-options]		
"high-availability-options (Aggregated Multiservices)" on page 8068	"member-failure-options (Aggregated Multiservices)" on page 10608	"member-interface (Aggregated Multiservices)" on page 10612

[edit interfaces {interface-name} load-balancing-options high-availability-options]		
"many-to-one (Aggregated Multiservices)" on page 10324	"one-to-one (Aggregated Multiservices)" on page 11829	

[edit interfaces {interface-name} load-balancing-options member-failure-options]		
"drop-member-traffic (Aggregated Multiservices)" on page 5949	"redistribute-all-traffic (Aggregated Multiservices)" on page 14165	

[edit interfaces {interface-name} load-balancing-options member-failure-options drop-member-traffic]	
"rejoin-timeout (Aggregated Multiservices)" on page 14273	

[edit interfaces {interface-name} load-balancing-options member-failure-options redistribute-all-traffic]	
"enable-rejoin (Aggregated Multiservices)" on page 6272	

[logical-tunnel-options]

[edit interfaces {name} logical-tunnel-options]	
"load-balance" on page 9874	

[edit interfaces {name} logical-tunnel-options load-balance]		
"adaptive" on page 1826	"no-adaptive" on page 11378	

[ls-]

[edit interfaces ls-{fpc}/{pic}/{port}:{channel}]	
"mlfr-uni-nni-bundle-options" on page 10789	

[ls-mlfr-uni-nni-bundle-options]

[edit interfaces ls-{fpc}/{pic}/{port}:{channel} mlfr-uni-nni-bundle-options]		
"acknowledge-retries" on page 1751	"acknowledge-timer" on page 1752	"action-red-differential-delay" on page 1782
"drop-timeout" on page 5981	"fragment-threshold" on page 7542	"hello-timer" on page 8043
"lmi-type" on page 9863	"mrru" on page 10989	"n391" on page 11163
"n392" on page 11164	"n393" on page 11166	"red-differential-delay" on page 14156

(Continued)

[edit interfaces ls-{fpc}/{pic}/{port}:{channel} mlfr-uni-nni-bundle-options]		
"t391" on page 28890	"t392" on page 28892	"yellow-differential-delay" on page 31168

[ls-unit]

[edit interfaces ls-{fpc}/{pic}/{port}:{channel} unit {logical-unit-number}]	
"interleave-fragments" on page 9185	

[lsq-]

[edit interfaces lsq-{fpc}/{pic}/{port}]	
"lsq-failure-options" on page 10197	

[lsq-lsq-failure-options]

[edit interfaces lsq-{fpc}/{pic}/{port} lsq-failure-options]		
"no-termination-request" on page 11596	"trigger-link-failure" on page 30164	

[lsq-unit]

[edit interfaces lsq-{fpc}/{pic}/{port} unit {logical-unit-number}]	
"fragment-threshold (Interfaces LSQ)" on page 7545	

[edit interfaces lsq-{fpc}/{pic}/{port} unit {logical-unit-number} compression rtp]	
"port (Services Voice) " on page 12884	

[edit interfaces lsq-{fpc}/{pic}/{port} unit {logical-unit-number} family mlppp]	
"bundle" on page 3072	

[mapping-range]

[edit {vlan-name} interfaces {interface-name} mapping-range {C-VLAN-range} (push | swap)]	
"vlan-id-start" on page 30818	

[media-type]

[edit interfaces {interface-name} media-type]	
"media-type (Interfaces)" on page 10602	

[mif]

[edit interfaces mif]		
"mtu (TDF Interface)" on page 11038	"unit (TDF Interface)" on page 30406	

[edit interfaces mif unit {interface-unit-number}]	
"family (TDF Interface)" on page 6917	

[edit interfaces mif unit {number} family inet service input]	
"service-set (TDF Interface)" on page 16756	

[edit interfaces mif unit {number} family inet service output]	
"service-set (TDF Interface)" on page 16756	

[mlfr-uni-nni-bundle-options]

[edit interfaces {interface-name} mlfr-uni-nni-bundle-options]		
"cisco-interoperability" on page 3334	"link-layer-overhead" on page 9797	"minimum-links" on page 10741

[mo-]

[edit interfaces mo-{fpc}/{pic}/{port}]		
"multiservice-options" on page 11131	"multiservice-options (Flow Monitoring)" on page 11130	

[mo-multiservice-options]

[edit interfaces mo-{fpc}/{pic}/{port} multiservice-options]		
"core-dump" on page 4670	"flow-control-options" on page 7270	"syslog" on page 28815

[multi-chassis-protection]

[edit interfaces {name} multi-chassis-protection {IP address of MC-LAG peer}]		
"icl-down-delay" on page 8252	"peer" on page 12339	

[multiservice-options]

[edit interfaces {interface-name} multiservice-options]		
"dump-on-flow-control" on page 6063	"flow-control (ACX Series)" on page 7265	"pic-boot-timeout" on page 12484

[optics-options]

[edit interfaces {interface-name} optics-options]		
"alarm (optics-options)" on page 2110	"alarms" on page 2124	"fec (MX Series)" on page 6987
"high-polarization" on page 8076	"is-ma" on page 9396	"lane" on page 9623
"lane-all" on page 9626	"tca" on page 28973	"tx-power" on page 30272
"warning" on page 31024	"wavelength" on page 31032	

[otn-options]

[edit interfaces {interface-name} otn-options]		
"bytes (otn-options)" on page 3123	"is-ma" on page 9396	"laser-enable" on page 9635
"line-loopback" on page 9775	"local-loopback" on page 9960	"odu-delay-management" on page 11777
"odu-signal-degrade" on page 11779	"odu-ttim-action-enable" on page 11782	"otu-ttim-action-enable" on page 11958
"prbs" on page 12987	"preemptive-fast-reroute" on page 13028	"rate" on page 14057
"signal-degrade" on page 27942	"tca" on page 28973	"trigger" on page 30154
"tti" on page 30176		

[edit interfaces {interface-name} otn-options odu-delay-management]		
"bypass" on page 3111	"monitor-end-point" on page 10835	"number-of-frames" on page 11736
"remote-loop-enable" on page 14348	"start-measurement" on page 28494	

[edit interfaces {interface-name} otn-options odu-signal-degrade]		
"ber-threshold-clear" on page 2894	"ber-threshold-signal-degrade" on page 2897	"interval" on page 9198

[edit interfaces {interface-name} otn-options preemptive-fast-reroute]		
"backward-frr-enable" on page 2827	"no-odu-backward-frr-enable" on page 11534	"no-odu-signal-degrade-monitor-enable" on page 11536
"odu-backward-frr-enable" on page 11776	"odu-signal-degrade-monitor-enable" on page 11781	"signal-degrade-monitor-enable" on page 27944

[edit interfaces {interface-name} otn-options rate]		
"fixed-stuff-bytes" on page 7201	"oc192" on page 11770	"otu4" on page 11960
"pass-through (interfaces)" on page 12192		

[edit interfaces {interface-name} otn-options signal-degrade]		
"ber-threshold-clear" on page 2894	"ber-threshold-signal-degrade" on page 2897	"interval" on page 9198

[pic-set]

[edit interfaces pic-set {pic-set-name}]	
"fpc (Interfaces)" on page 7511	

[edit interfaces pic-set pic-set-name]	
"interface (PIC Bundle)" on page 8949	

[pp0]

[edit interfaces pp0 unit {logical-unit-number}]		
"pppoe-options" on page 12974	"pppoe-options (SRX Series)" on page 12972	"targeted-distribution (Static Interfaces over Aggregated Ethernet)" on page 28961

[edit interfaces pp0 unit {unit-number} ppp-options]		
"authentication (Static and Dynamic PPP)" on page 2664	"initiate-ncp (Dynamic and Static PPP)" on page 8662	"ipcp-suggest-dns-option" on page 9289
"mru (Dynamic and Static PPPoE)" on page 10994	"mtu (Dynamic and Static PPPoE)" on page 11034	"on-demand-ip-address" on page 11815
"peer-ip-address-optional" on page 12373		

[edit interfaces pp0 unit {unit-number} ppp-options chap]		
"challenge-length (Static and Dynamic PPP)" on page 3281		

[edit interfaces pp0 unit {logical-unit-number} pppoe-options]		
"client (interfaces)" on page 4299	"server" on page 16608	"service-name" on page 16714
"underlying-interface" on page 30330		

[ppp-options]

[edit interfaces {interface-name} ppp-options]		
"chap" on page 3289	"compression (PPP Properties)" on page 4494	"no-termination-request" on page 11596
"pap" on page 12169		

[edit interfaces {interface-name} ppp-options chap]		
"access-profile" on page 1688	"default-chap-secret" on page 5184	"local-name" on page 9964
"passive (CHAP)" on page 12198		

[edit interfaces {interface-name} ppp-options compression]		
"acfc" on page 1748	"pfc" on page 12455	

[edit interfaces {interface-name} ppp-options pap]		
"access-profile" on page 1688	"local-name" on page 9964	"local-password" on page 9966
"passive (PAP)" on page 12214		

[ps]

[edit interfaces ps {number}]	
"rebalance-periodic (PS Interface)" on page 14109	

[edit interfaces ps {number} targeted-options]	
"logical-interface-fpc-redundancy (PS Interface)" on page 10068	

[ps0]

[edit interfaces ps0]		
"anchor-point (Pseudowire Subscriber Interfaces)" on page 2244	"flexible-vlan-tagging" on page 7235	"untagged" on page 30431

[pstargeted-options]

[edit interfaces ps{number} targeted-options]	
"rebalance-subscriber-granularity (PS interface)" on page 14111	

[radius-options]

[edit interfaces {interface-name} radius-options]	
"nas-port-options (RADIUS Options)" on page 11193	

[edit interfaces {interface-name} radius-options nas-port-options {nas-port-options-name}]		
"nas-port-extended-format (Interfaces)" on page 11186	"nas-port-type (RADIUS Options)" on page 11197	"stacked-vlan-ranges (RADIUS Options)" on page 28473
"vlan-ranges (RADIUS Options)" on page 30847		

[redundant-ether-options]

[edit interfaces {interface-name} redundant-ether-options]		
"flow-control (Interfaces)" on page 7269	"lacp (Interfaces)" on page 9614	"link-speed (Interfaces)" on page 9830

(Continued)

[edit interfaces {interface-name} redundant-ether-options]		
"loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)" on page 10117	"minimum-links (Interfaces)" on page 10746	"redundancy-group (Interfaces)" on page 14179
"source-address-filter (Interfaces)" on page 28130	"source-filtering (Interfaces)" on page 28210	

[rlsq]

[edit interfaces rlsq {number}:{number}]		
"encapsulation (interfaces)" on page 6289	"encapsulation (Physical Interface)" on page 6307	"redundancy-options (Link Services IQ PIC Interfaces) " on page 14195

[edit interfaces rlsq {number} unit {logical-unit-number}]	
"encapsulation (Logical Interface)" on page 6300	

[rlsqredundancy-options]

[edit interfaces rlsq{number} redundancy-options]		
"hot-standby" on page 8187	"hot-standby" on page 8187	"primary (Link Services IQ PIC Interfaces) " on page 13142
"secondary (Link Services IQ PIC Interfaces) " on page 16414	"warm-standby (Interfaces)" on page 31022	

[rlt]

[edit interfaces rlt {x} logical-tunnel-options load-balance]	
"local-bias" on page 9922	

[rms]

[edit interfaces rms{number}]	
"redundancy-options (Adaptive Services Interfaces) " on page 14192	

[rmsredundancy-options]

[edit interfaces rms{number} redundancy-options]	
"hot-standby" on page 8187	

[rsp]

[edit interfaces rsp{number}]	
"redundancy-options (Adaptive Services Interfaces) " on page 14192	

[rspredundancy-options]

[edit interfaces rsp{number} redundancy-options]	
"hot-standby" on page 8187	

[se-]

[edit interfaces se-{pim}/0/{port}]	
"serial-options" on page 16604	

[serial-options]

[edit interfaces {interface-name} serial-options]		
"clock-rate" on page 4360	"clocking-mode" on page 4364	"control-polarity" on page 4636
"cts-polarity" on page 4734	"dcd-polarity" on page 4863	"dce-options" on page 4864
"dsr-polarity" on page 6037	"dte-options" on page 6039	"dtr-circuit" on page 6042
"dtr-polarity" on page 6044	"encoding" on page 6317	"indication-polarity" on page 8548
"line-protocol" on page 9776	"loopback (Serial)" on page 10129	"rts-polarity" on page 16071
"tm-polarity" on page 29418	"transmit-clock" on page 30083	

[edit interfaces {interface-name} serial-options dce-options]		
"control-signal" on page 4640	"cts" on page 4732	"dcd" on page 4861
"dsr" on page 6035	"dtr" on page 6041	"ignore-all" on page 8394
"indication" on page 8546	"rts" on page 16069	"tm" on page 29417

[edit interfaces {interface-name} serial-options dte-options]		
"control-signal" on page 4640	"cts" on page 4732	"dcd" on page 4861
"dsr" on page 6035	"dtr" on page 6041	"ignore-all" on page 8394
"indication" on page 8546	"rts" on page 16069	"tm" on page 29417

[services-options]

[edit interfaces {interface-name} services-options]		
"cgn-pic" on page 3273	"close-timeout" on page 4365	"disable-global-timeout-override" on page 5760
"enable-subscriber-analysis (Services Options VMS Interfaces)" on page 6279	"fragment-limit" on page 7541	"inactivity-timeout (services)" on page 8506
"jflow-log (Interfaces)" on page 9426	"open-timeout" on page 11841	"pba-interim-logging-interval" on page 12280
"reassembly-timeout" on page 14102	"session-limit" on page 16816	"syslog (Interfaces)" on page 28820

(Continued)

[edit interfaces {interface-name} services-options]		
"syslog-local-system-timestamp" on page 28821	"tcp-tickles" on page 29029	"trio-flow-offload" on page 30167
"utc-timestamp" on page 30612		

[edit interfaces {interface-name} services-options jflow-log]	
"message-rate-limit (Flow Monitoring Logs for NAT)" on page 10656	

[edit interfaces {interface-name} services-options session-limit]		
"cpu-load-threshold" on page 4696	"maximum" on page 10497	"rate (Interface Services)" on page 14061

[edit interfaces {interface-name} services-options syslog]	
"host (Interfaces)" on page 8144	

[edit interfaces {interface-name} services-options syslog host {hostname}]		
"log-prefix (Interfaces)" on page 10020	"port (System Log Messsages)" on page 12893	"services (Interfaces)" on page 16776

[edit interfaces {interface-name} services-options sysloghost {hostname}]	
"facility-override" on page 6770	

[shdsl-options]

[edit interfaces {interface-name} shdsl-options]		
"line-rate" on page 9778	"line-rate (Interfaces)" on page 9779	"snr-margin" on page 28081

[edit interfaces {interface-name} shdsl-options snr-margin]	
"snext" on page 28056	

[si]

[edit interfaces si {interface-id} unit {unit-number} ppp-options]		
"mru (Dynamic and Static PPPoE)" on page 10994	"mtu (Dynamic and Static PPPoE)" on page 11034	

[si-]

[edit interfaces si-{slot}/{pic}/{port}]	
"l2tp-maximum-session (Service Interfaces)" on page 9551	

[si-unit]

[edit interfaces si-{fpc}/{pic}/{port} unit {logical-unit-number}]		
"dial-options" on page 5684		

[edit interfaces si-{fpc}/{pic}/{port} unit {logical-unit-number} ppp-options]		
"initiate-ncp (Dynamic and Static PPP)" on page 8662	"ipcp-suggest-dns-option" on page 9289	

[so-unit]

[edit interfaces so-{fpc}/{pic}/{slot} unit {logical-unit-number}]		
"interface-shared-with" on page 9049		

[edit interfaces so-{fpc}/{pic}/{port} unit {number} ppp-options]		
"lcp-max-conf-req" on page 9669	"ncp-max-conf-req" on page 11243	

[sonet-options]

[edit interfaces {interface-name} sonet-options aps]		
"fast-aps-switch" on page 6922	"preserve-interface" on page 13121	

[edit interfaces {interface-name} sonet-options mpls]	
"pop-all-labels" on page 12800	

[edit interfaces {interface-name} sonet-options mpls pop-all-labels]	
"required-depth" on page 15574	

[sp-unit]

[edit interfaces sp-{fpc}/{pic}/{port} unit {logical-unit-number}]		
"dial-options" on page 5684	"twamp-server" on page 30258	

[edit interfaces sp-{fpc}/{pic}/{port} unit {logical-unit-number} interface]	
"l2tp-interface-id" on page 9550	

[speed]

[edit interfaces {interface-name} speed]	
"speed (Interfaces)" on page 28376	

[t1-options]

[edit interfaces {interface-name} t1-options]		
"bert-algorithm" on page 2900	"bert-error-rate" on page 2903	"bert-period" on page 2905

[t3-options]

[edit interfaces {interface-name} t3-options]		
"bert-algorithm" on page 2900	"bert-error-rate" on page 2903	"bert-period" on page 2905
"unframed | no-unframed (Interfaces)" on page 30337		

[tdm-options]

[edit interfaces {interface-name} tdm-options]		
"ces-psn-channel (tdm-options)" on page 3265	"iwf-params (tdm-options)" on page 9420	

[umd0]

[edit interfaces umd0]		
"dialer-options" on page 5689	"modem-options" on page 10825	

[edit interfaces umd0 dialer-options]		
"pool" on page 12762		

[edit interfaces umd0 modem-options]		
"dialin" on page 5691	"init-command-string" on page 8655	

[unit]

[edit interfaces {interface-name} unit {logical-unit-number}]		
"accept-source-mac" on page 1647	"accounting-profile" on page 1727	"address (Interfaces, Flow Monitoring)" on page 1879
"advisory-options (Traffic Shaping)" on page 2035	"alias (Interfaces)" on page 2146	"arp-resp" on page 2479
"auto-configure (Dynamic VLAN Interface Sets)" on page 2696	"backup-options" on page 2807	"bandwidth (Interfaces)" on page 2847
"client-delegate-probes" on page 4307	"compression-device (Interfaces)" on page 4493	"deep-vlan-qualified-learning" on page 5176
"demux (Interfaces)" on page 5299	"demux-destination (Underlying Interface)" on page 5297	"demux-options (Static Interface)" on page 5304
"demux-source (Underlying Interface)" on page 5309	"description (Interfaces)" on page 5334	"disable (Interface)" on page 5765
"disable-mlppp-inner-ppp-pfc" on page 5772	"dlci" on page 5843	"encapsulation (interfaces)" on page 6289

(Continued)

[edit interfaces {interface-name} unit {logical-unit-number}]		
"encapsulation (Logical Interface)" on page 6300	"encapsulation (logical-link-layer)" on page 6286	"esi" on page 6453
"f-max-period" on page 6742	"family (Aggregated Multiservices)" on page 6844	"family (ELS)" on page 6824
"family (Interfaces)" on page 6881	"family (interfaces)" on page 6832	"family (Link Interfaces)" on page 6839
"family (logical-interface)" on page 6822	"family (Voice Services)" on page 6918	"family inet (Interfaces)" on page 6874
"family inet6" on page 6878	"filter (Applying to a Logical Interface)" on page 7123	"force-control-packets-on-transit-path" on page 7367
"forwarding-class-accounting" on page 7413	"host-prefix-only" on page 8163	"inner-vlan-id-swap-ranges" on page 8685
"input-vlan-map " on page 8729	"input-vlan-map (Aggregated Ethernet)" on page 8731	"interface-state" on page 9064
"ip-address-owner" on page 9236	"keepalives" on page 9475	"layer2-policer (EX)" on page 9654
"layer2-policer (Hierarchical Policer)" on page 9656	"link-layer-overhead" on page 9797	"minimum-links" on page 10741
"multicast-dlci" on page 11066	"multilink-max-classes" on page 11106	"multipoint" on page 11127
"no-auto-virtual-gateway-esi" on page 11410	"no-dhcp-flood" on page 11627	"no-keepalives" on page 11488
"output-vlan-map " on page 12021	"output-vlan-map (Aggregated Ethernet)" on page 12022	"passive-monitor-mode" on page 12212

(Continued)

[edit interfaces {interface-name} unit {logical-unit-number}]		
"peer-unit" on page 12390	"per-session-scheduler" on page 12422	"point-to-point" on page 12619
"policer (Applying to a Logical Interface)" on page 12624	"policer-overhead" on page 12643	"policer-overhead-adjustment" on page 12644
"ppp-options" on page 12949	"ppp-over-ether" on page 12956	"pppoe-underlying-options (Static and Dynamic Subscribers)" on page 12980
"proxy-arp" on page 13820	"rpm (Interfaces) " on page 16019	"shaping" on page 16925
"swap-by-poppush" on page 28741	"targeted-options (Grouping Subscribers by Bandwidth Usage)" on page 28964	"traps" on page 30142
"traps" on page 30142	"tunnel" on page 30189	"tunnel" on page 30190
"tunnel" on page 30192	"virtual-gateway-v4-mac" on page 30755	"virtual-gateway-v6-mac" on page 30757
"vlan-id (Fast Ethernet)" on page 30794	"vlan-id (VLAN ID to Be Bound to a Logical Interface)" on page 30820	"vlan-id-list" on page 30805
"vlan-id-list (Ethernet VLAN Circuit)" on page 30808	"vlan-id-list (Interface in VPLS)" on page 30812	"vlan-id-range" on page 30815
"vlan-tags (Dual-Tagged Logical Interface)" on page 30869	"vlan-tags (Stacked VLAN Tags)" on page 30873	

[edit interfaces {interface-name} unit (Interfaces) {logical-unit-number}]	
"compression-device (Interfaces)" on page 4492	

[edit interfaces {interface-name} unit 0]	
"vlan-id-list" on page 30805	

[edit interfaces {interface-name} unit 0 family bridge]	
"recovery-timeout" on page 14149	

[edit interfaces {interface-name} unit 0 family ethernet-switching]		
"native-vlan-id" on page 11240	"recovery-timeout" on page 14149	

[edit interfaces {interface-name} unit {logical-unit-number} accept-source-mac]	
"mac-address (Accept Source Mac)" on page 10219	

[edit interfaces {interface-name} unit {logical-unit-number} accept-source-mac mac-address {mac-address}]	
"policer (MAC)" on page 12641	

[edit interfaces {interface-name} unit {logical-unit-number} address {address} family {family} multipoint-destination {address}]	
"shaping" on page 16925	

[edit interfaces {interface-name} unit {logical-unit-number} advisory-options]	
"downstream-rate (Traffic Shaping)" on page 5935	

[edit interfaces {interface-name} unit {logical-unit-number} auto-configure]		
"agent-circuit-identifier (Dynamic ACI VLANs)" on page 2048	"line-identity (Dynamic Access-Line-Identifier VLANs)" on page 9772	

[edit interfaces {interface-name} unit {logical-unit-number} auto-configure agent-circuit-identifier]	
"dynamic-profile (Dynamic ACI VLANs)" on page 6122	

[edit interfaces {interface-name} unit {logical-unit-number} auto-configure line-identity]		
"dynamic-profile (Dynamic Access-Line-Identifier VLANs)" on page 6121	"include (Dynamic Access-Line-Identifier VLANs)" on page 8525	

[edit interfaces {interface-name} unit {logical-unit-number} compression]	
"rtp" on page 16068	

[edit interfaces {interface-name} unit {logical-unit-number} compression rtp]		
"f-max-period (rtp)" on page 6744	"maximum-contexts" on page 10517	"queues" on page 13955

[edit interfaces {interface-name} unit {unit-number} demux]		
"inet (Interfaces)" on page 8568	"inet6 (Interfaces)" on page 8608	

[edit interfaces {interface-name} unit {unit-number} demux inet]	
"auto-configure (Demux)" on page 2694	

[edit interfaces {interface-name} unit {unit-number} demux inet auto-configure]		
"address-ranges (Demux)" on page 1919		

[edit interfaces {interface-name} unit {unit-number} demux inet auto-configure address-ranges]		
"authentication (Demux)" on page 2592	"dynamic-profile (Demux)" on page 6115	

[edit interfaces {interface-name} unit {unit-number} demux inet auto-configure address-ranges authentication]		
"username-include (Demux)" on page 30593		

[edit interfaces {interface-name} unit {unit-number} demux inet6]		
"auto-configure (Demux)" on page 2694		

[edit interfaces {interface-name} unit {unit-number} demux inet6 auto-configure]		
"address-ranges (Demux)" on page 1919		

[edit interfaces {interface-name} unit {unit-number} demux inet6 auto-configure address-ranges]		
"authentication (Demux)" on page 2592	"dynamic-profile (Demux)" on page 6115	

[edit interfaces {interface-name} unit {unit-number} demux inet6 auto-configure address-ranges authentication]		
"username-include (Demux)" on page 30593		

[edit interfaces {interface-name} unit {logical-unit-number} dhcp-client]	
"force-discover (DHCP Client)" on page 7369	

[edit interfaces {interface-name} unit {logical-unit-number} dynamic-call-admission-control]	
"activation-priority" on page 1799	

[edit interfaces {name} unit {number} esi]		
"df-election-granularity" on page 5574	"df-election-type" on page 5576	"source-bmac" on page 28185

[edit interfaces {interface-name} unit {logical-unit-number} esi df-election-granularity]	
"per-evi" on page 12409	

[edit interfaces {interface-name} unit {logical-unit-number} family {family}]		
"address" on page 1846	"address (Interfaces)" on page 1878	"core-facing" on page 4672
"demux-destination (Demux Interface)" on page 5296	"demux-source (Demux Interface)" on page 5305	"dhcp-client" on page 5596
"dialer" on page 5688	"fibre-channel (Family Interfaces)" on page 7012	"filter (Applying to a Logical Interface)" on page 7123
"filter (Applying to an Interface)" on page 7125	"filter (Interfaces)" on page 7117	"filter (Layer 2 and Layer 3 Interfaces)" on page 7145
"filter (Layer 3 Interfaces)" on page 7146	"inet (interfaces (QFX Series))" on page 8565	"inet6 (interfaces (QFX Series))" on page 8607

(Continued)

[edit interfaces {interface-name} unit {logical-unit-number} family {family}]		
"ingress-queuing-filter" on page 8644	"iq-policing-filter" on page 9380	"mac-validate" on page 10273
"mtu (interfaces)" on page 11027	"mtu (Multilink and Link Services Logical Interface)" on page 11025	"next-hop-tunnel" on page 11355
"no-dns-propagation" on page 11442	"no-redirects" on page 11559	"no-redirects (IPv4 Traffic)" on page 11560
"policer (Applying to a Logical Interface)" on page 12624	"policer (Interface)" on page 12638	"primary (Interface for Router)" on page 13140
"rpf-check " on page 16007	"rpf-check " on page 16007	"simple-filter (Interfaces)" on page 27965
"source-lookup" on page 28338	"unnumbered-address (Ethernet)" on page 30425	"vpls (Interfaces)" on page 30931
"web-authentication (Interfaces)" on page 31060		

[edit interfaces {interface-name} unit {logical-unit-number} family (ccc | inet | tcc)]	
"policer (Layer 2 VPN)" on page 12639	

[edit interfaces {interface-name} unit {logical-unit-number} family (inet | inet6) rpf-check]	
"mode (Interfaces)" on page 10804	

[edit interfaces {interface-name} unit {logical-unit-number} family {family} address {address}]		
"broadcast (Interfaces)" on page 3041	"preferred" on page 13053	"primary (Address on Interface)" on page 13130
"virtual-gateway-address" on page 30754		

[edit interfaces {interface-name} unit {logical-unit-number} family bridge]		
"interface-mode" on page 8920	"isid-list" on page 9400	"storm-control" on page 28620
"vlan-id (Logical Port in Bridge Domain)" on page 30813	"vlan-id-list (Interface in Bridge Domain)" on page 30810	"vlan-rewrite" on page 30849

[edit interfaces {interface-name} unit {logical-unit-number} family ccc]	
"keep-address-and-control" on page 9455	

[edit interfaces {interface-name} unit {logical-unit-number} family {family-name} dhcp]		
"client-identifier (Interfaces)" on page 4322	"vendor-id (Interfaces)" on page 30682	

[edit interfaces {interface-name} unit {logical-unit-number} family {family} dhcp-client]		
"lease-time (dhcp-client)" on page 9698	"retransmission-attempt (DHCP Client)" on page 15677	"retransmission-interval (DHCP Client)" on page 15685
"server-address (dhcp-client)" on page 16611	"update-server (dhcp-client)" on page 30452	"vendor-id (DHCP Client)" on page 30676

[edit interfaces {interface-name} unit {logical-unit-number} family {family} dhcp-client client-identifier]		
"prefix (DHCP Client)" on page 13076	"use-interface" on page 30519	"user-id" on page 30557

[edit interfaces {interface-name} unit {logical-unit-number} family {family} dhcpv6-client]		
"client-identifier (DHCPv6 Client)" on page 4321	"rapid-commit (DHCPv6 Client)" on page 14054	"req-option" on page 14408
"retransmission-attempt (DHCPv6 Client)" on page 15678	"update-router-advertisement" on page 30449	"update-server (dhcpv6-client)" on page 30453

[edit interfaces {interface-name} unit {logical-unit-number} family ethernet-switching]		
"interface-mode" on page 8920	"port-mode" on page 12841	"reflective-relay" on page 14232
"storm-control" on page 28620	"vlan-rewrite" on page 30849	

[edit interfaces {interface-name} unit {logical-unit-number} family ethernet-switching interface-mode trunk]	
"inter-switch-link" on page 8795	

[edit interfaces {interface-name} unit {logical-unit-number} family fibre-channel]	
"port-mode (Fibre Channel Interfaces)" on page 12843	

[edit interfaces {interface-name} unit {logical-unit-number} family {family} filter]		
"input-chain" on page 8690	"input-list" on page 8708	"output-chain" on page 11983

[edit interfaces {interface-name} unit {logical-unit-number} family inet]		
"accounting (interfaces)" on page 1709	"arp-max-cache" on page 2471	"arp-new-hold-limit" on page 2473
"dhcp (DHCP Client)" on page 5598	"dhcp (DHCP Client)" on page 5598	"filter (Flow Monitoring)" on page 7115
"filter (interfaces)" on page 7116	"mtu (Dynamic Profiles)" on page 11035	"multicast-only" on page 11076
"negotiate-address" on page 11261	"no-neighbor-learn" on page 11520	"receive-options-packets" on page 14116
"receive-options-packets (Passive Flow Monitoring)" on page 14114	"receive-ttl-exceeded" on page 14121	"receive-ttl-exceeded (Passive Flow Monitoring)" on page 14120
"rpf-check " on page 16007	"sampling (Interfaces)" on page 16298	"service" on page 16680
"service-domain" on page 16693	"simple-filter (Applying to an Interface)" on page 27964	"targeted-broadcast" on page 28956
"targeted-broadcast" on page 28956	"targeted-broadcast" on page 28956	"tcp-mss (Interfaces)" on page 29033
"unconditional-src-learn" on page 30328	"unnumbered-address (Demux)" on page 30419	"unnumbered-address (PPP)" on page 30426
"vrrp-inherit-from" on page 30978		

[edit interfaces {interface-name} unit {logical-unit-number} family inet accounting]		
"destination-class-usage" on page 5411	"source-class-usage" on page 28191	

[edit interfaces {interface-name} unit {logical-unit-number} family inet address {address}]		
"arp" on page 2458	"destination (Interfaces)" on page 5433	"destination (Tunnels)" on page 5528
"destination-profile" on page 5509	"policing" on page 12661	"vrrp-group" on page 30973

[edit interfaces {interface-name} unit {logical-unit-number} family inet address {address} vrrp-group {group-id}]		
"accept-data" on page 1631	"advertise-interval" on page 2017	"authentication-key" on page 2603
"authentication-type" on page 2667	"fast-interval" on page 6926	"preempt (VRRP)" on page 13025
"priority" on page 13154	"priority (Protocols VRRP)" on page 13201	"track (VRRP)" on page 29992
"virtual-address" on page 30739		

[edit interfaces {interface-name} unit {logical-unit-number} family inet address {address} vrrp-group {group-id} preempt]	
"hold-time (VRRP)" on page 8119	

[edit interfaces {interface-name} unit {logical-unit-number} family inet address {address} vrrp-group {group-id} track]		
"interface (VRRP)" on page 8812	"priority-hold-time" on page 13179	"route (Interfaces)" on page 15868

[edit interfaces {interface-name} unit {logical-unit-number} family inet address {address} vrrp-group {group-id} track interface {interface-name}]		
"bandwidth-threshold" on page 2878	"priority-cost (VRRP)" on page 13163	

[edit interfaces {name} unit {name} family inet dhcp]		
"client-identifier (DHCP Client)" on page 4319	"lease-time" on page 9696	"retransmission-interval (DHCP)" on page 15684
"update-server" on page 30450		

[edit interfaces {name} unit {name} family inet inet6 address]	
"vrrp-group" on page 30973	

[edit interfaces {interface-name} unit {logical-unit-number} family inet service (input | output)]	
"service-set (Interfaces)" on page 16744	

[edit interfaces {interface-name} unit {logical-unit-number} family inet service (input | output) service-set {service-set-name}]	
"service-filter (Interfaces)" on page 16701	

[edit interfaces {interface-name} unit {logical-unit-number} family inet service input]	
"post-service-filter" on page 12907	

[edit interfaces {interface-name} unit {logical-unit-number} family inet targeted-broadcast]		
"forward-and-send-to-re" on page 7392	"forward-only" on page 7395	

[edit interfaces {interface-name} unit {logical-unit-number} family inet unnumbered-address {interface-name}]		
"destination (IPCP)" on page 5439	"destination (Tunnels)" on page 5528	

[edit interfaces {interface-name} unit {logical-unit-number} family inet unnumbered-address {interface-name} destination {address}]		
"destination-profile" on page 5509		

[edit interfaces {name} unit {name} family inet6]		
"dad-proxy" on page 4809	"dhcpv6-client" on page 5655	"nd6-max-cache" on page 11252
"nd6-new-hold-limit" on page 11254	"nd6-stale-time" on page 11256	"ndp-proxy" on page 11259
"rpf-check " on page 16007	"tcp-mss (Interfaces)" on page 29033	"vrrp-inherit-from" on page 30978

[edit interfaces {interface-name} unit {number} family inet6 address {address}]		
"eui-64" on page 6515	"ndp" on page 11257	"vrrp-inet6-group" on page 30976

[edit interfaces {interface-name} unit {logical-unit-number} family inet6 address {address} vrrp-inet6-group {group-id}]		
"accept-data" on page 1631	"fast-interval" on page 6926	"inet6-advertise-interval (VRRP)" on page 8596
"preempt" on page 13019	"preempt (VRRP)" on page 13025	"priority" on page 13154
"priority (Protocols VRRP)" on page 13201	"track (VRRP)" on page 29992	"virtual-inet6-address" on page 30760
"virtual-link-local-address" on page 30765		

[edit interfaces {interface-name} unit {logical-unit-number} family inet6 address {address} vrrp-inet6-group {group-id} preempt]		
"hold-time" on page 8101	"hold-time (VRRP)" on page 8119	

[edit interfaces {interface-name} unit {logical-unit-number} family inet6 address {address} vrrp-inet6-group {group-id} track]		
"interface (VRRP)" on page 8812	"priority-hold-time" on page 13179	"route (Interfaces)" on page 15868

[edit interfaces {interface-name} unit {logical-unit-number} family inet6 address {address} vrrp-inet6-group {group-id} track interface {interface-name}]		
"bandwidth-threshold" on page 2878	"priority-cost (VRRP)" on page 13163	

[edit interfaces {interface-name} unit {logical-unit-number} family inet6 dhcpv6-client]		
"client-ia-type" on page 4313	"client-type" on page 4340	

[edit interfaces {interface-name} unit {logical-unit-number} family mlfr-end-to-end]	
"bundle (multilink)" on page 3073	

[edit interfaces {interface-name} unit {logical-unit-number} family mlfr-uni-nni]	
"bundle (multilink)" on page 3073	

[edit interfaces {interface-name} unit {logical-unit-number} family mlppp]	
"bundle (Interfaces)" on page 3074	

[edit interfaces {interface-name} unit {logical-unit-number} family mpls]	
"maximum-labels" on page 10528	

[edit interfaces {interface-name} unit {logical-unit-number} family pppoe]		
"access-concentrator" on page 1655	"dynamic-profile (Dynamic PPPoE)" on page 6124	"max-sessions (Dynamic PPPoE)" on page 10469

(Continued)

[edit interfaces {interface-name} unit {logical-unit-number} family pppoe]		
"max-sessions-vsa-ignore (Static and Dynamic Subscribers)" on page 10481	"service-name-table" on page 16716	"short-cycle-protection (Static and Dynamic Subscribers)" on page 16984

[edit interfaces {interface-name} unit {logical-unit-number} family tcc]		
"proxy (Interfaces)" on page 13828	"proxy (interfaces)" on page 13817	"remote" on page 14305

[edit interfaces {interface-name} unit {logical-unit-number} forwarding-class-accounting]		
"direction (forwarding-class-accounting)" on page 5716	"enhanced (forwarding-class-accounting)" on page 6358	

[edit interfaces {interface-name} unit {logical-unit-number} input-vlan-map]		
"inner-tag-protocol-id" on page 8679	"inner-vlan-id" on page 8682	"pop" on page 12799
"pop-pop" on page 12806	"pop-swap" on page 12807	"push" on page 13875
"push-push" on page 13879	"swap" on page 28740	"swap-push" on page 28744
"swap-swap" on page 28745	"tag-protocol-id (TPID to Rewrite)" on page 28920	

[edit interfaces {interface-name} unit {logical-unit-number} layer2-policer]		
"input-hierarchical-policer" on page 8704	"input-policer" on page 8713	"input-three-color" on page 8724
"output-policer" on page 12001	"output-three-color" on page 12014	

[edit interfaces unit {unit-name} load-balancing-options]	
"hash-keys (Interfaces)" on page 7995	

[edit interfaces {interface-name} unit logical-unit-number]	
"accept-source-mac (SRX)" on page 1645	

[edit interfaces {interface-name} unit {logical-unit-number} output-vlan-map]		
"inner-tag-protocol-id" on page 8679	"inner-vlan-id" on page 8682	"pop" on page 12799
"pop-pop" on page 12806	"pop-swap" on page 12807	"push" on page 13875
"push-push" on page 13879	"swap" on page 28740	"swap-push" on page 28744
"swap-swap" on page 28745	"tag-protocol-id (TPID to Rewrite)" on page 28920	

[edit interfaces {interface-name} unit {logical-unit-number} ppp-options]		
"chap" on page 3289	"compression (PPP Properties)" on page 4494	"dynamic-profile (PPP)" on page 6128
"lcp-restart-timer" on page 9672	"loopback-clear-timer" on page 10119	"ncp-restart-timer" on page 11245
"pap" on page 12169		

[edit interfaces {interface-name} unit {logical-unit-number} ppp-options chap]		
"access-profile" on page 1688	"default-chap-secret" on page 5184	"local-name" on page 9964

(Continued)

[edit interfaces {interface-name} unit {logical-unit-number} ppp-options chap]		
"passive (CHAP)" on page 12198		

[edit interfaces {interface-name} unit {logical-unit-number} ppp-options compression]		
"acfc" on page 1748	"pfc" on page 12455	

[edit interfaces {interface-name} unit {logical-unit-number} ppp-options pap]		
"access-profile" on page 1688	"default-pap-password" on page 5199	"local-name" on page 9964
"local-password" on page 9966	"passive (PAP)" on page 12214	

[edit interfaces {interface-name} unit {logical-unit-number} pppoe-options]		
"access-concentrator" on page 1655	"auto-reconnect" on page 2732	

[edit interfaces {interface-name} unit {logical-unit-number} pppoe-underlying-options]		
"access-concentrator" on page 1655	"direct-connect" on page 5708	"duplicate-protection (Dynamic PPPoE)" on page 6071
"dynamic-profile (Dynamic PPPoE)" on page 6124	"max-sessions (Dynamic PPPoE)" on page 10469	"max-sessions-vsa-ignore (Static and Dynamic Subscribers)" on page 10481
"service-name-table" on page 16716	"short-cycle-protection (Static and Dynamic Subscribers)" on page 16984	

[edit interfaces {interface-name} unit {logical-unit-number} radio-router]		
"data-rate" on page 4841	"hub-assist" on page 8213	

[edit interfaces {interface-name} unit {logical-unit-number} rtp]		
"f-max-period" on page 6742		

[edit interfaces {interface-name} unit {logical-unit-number} tunnel]		
"backup-destination" on page 2795	"destination (Interfaces)" on page 5433	"destination (Interfaces)" on page 5433
"destination (Tunnel Remote End)" on page 5526	"destination (Tunnels)" on page 5528	"destination-hostname" on page 5426
"encap-type" on page 6280	"ipip" on page 9292	"key (gre)" on page 9482
"key (interfaces)" on page 9484	"routing-instance (Interfaces)" on page 15923	"source (Interfaces)" on page 28112
"source (Interfaces)" on page 28112	"source-hostname" on page 28213	"source-interface" on page 28225
"traffic-class (Tunnels)" on page 30008	"ttl" on page 30178	

[edit interfaces {name} unit {name} tunnel encapsulation {vxlan-gpe}]		
"destination-udp-port (FTI)" on page 5531	"udp (FTI)" on page 30309	"vni(Interfaces)" on page 30911
"vxlan-gpe (FTI)" on page 31005		

[edit interfaces {name} unit {name} tunnel encapsulation(gre | ipip | udp) destination]	
"address (tunnel-encapsulation)" on page 1843	

[edit interfaces {interface-name} unit {logical-unit-number} tunnel routing-instance]	
"destination (Routing Instance)" on page 5518	

[edit interfaces {interface-name} unit {logical-unit-number}family pppoe]		
"direct-connect" on page 5708	"duplicate-protection (Dynamic PPPoE)" on page 6071	

[unitradio-router]

[edit interfaces {interface-name} unit{logical-unit-number} radio-router credit]	
"interval (Interfaces)" on page 9214	

[vlan]

[edit interfaces vlan unit {logical-unit-number} family]	
"fibre-channel (Family Interfaces)" on page 7012	

[edit interfaces vlan unit {unit} family fibre-channel]	
"port-mode (Fibre Channel Interfaces)" on page 12843	

[xe-]

[edit interfaces xe-{fpc}/{pic}/{port}]	
"flexible-vlan-tagging" on page 7235	

[xe-,edit]

[edit interfaces xe-{fpc}/{pic}/{port}],[edit interfaces et-{fpc}/{pic}/{port}] (PTX Series Packet Transport Routers and MX Series Routers)]	
"framing (10-Gigabit Ethernet Interfaces)" on page 7568	

[xe-fpc/pic/port]

[edit interfaces xe-fpc/pic/port]	
"otn-options" on page 11955	

[xe-unit]

[edit interfaces xe-{fpc}/{pic}/{slot} unit {logical-unit-number}]	
"interface-shared-with" on page 9049	

[{interface-name]

[edit interfaces {interface-name }forwarding-class-accounting]	
"direction (forwarding-class-accounting)" on page 5716	

[{interface-name}redundant-ether-options]

[edit interfaces {interface-name}redundant-ether-options lacp]	
"periodic (Interfaces)" on page 12439	

[{interface-name}unit]

[edit interfaces {interface-name}unit {logical-unit-number}family inet dhcp]	
"server-address" on page 16609	

[edit jsrc]

IN THIS SECTION

- [partition] | 353

Sub Hierarchies

"[edit jsrc partition {partition-name}]" on page 353

[edit jsrc]		
"dualstack-support (JSRC)" on page 6061	"partition" on page 12181	

[partition]

[edit jsrc partition {partition-name}]		
"destination-host" on page 5421	"destination-realm (JSRC)" on page 5514	"diameter-instance (JSRC)" on page 5696

[edit logical-routers]

IN THIS SECTION

- [bridge-domains] | 354

Sub Hierarchies

"[edit logical-routers {logical-router-name} bridge-domains]" on page 354

[bridge-domains]

[edit logical-routers {logical-router-name} bridge-domains]	
"no-irb-layer-2-copy" on page 11486	

[edit logical-systems]

IN THIS SECTION

- [...] | 357
- [access] | 358
- [aeaggregated-ether-options] | 359
- [applications] | 359

- [bridge-domains] | 359
- [class-of-service] | 362
- [firewall] | 363
- [firewallpolicer] | 366
- [forwarding-options] | 366
- [interfaces] | 377
- [l2circuit] | 392
- [logical-system-name] | 392
- [name] | 394
- [neighbor] | 394
- [policer] | 394
- [policy-options] | 395
- [policy-statement] | 398
- [protocols] | 399
- [realm] | 473
- [routing] | 473
- [routing-instances] | 473
- [routing-instancesrouting-options] | 584
- [routing-options] | 584
- [security] | 599
- [services] | 609
- [statement-name] | 610
- [switch-options] | 610
- [system] | 611
- [tenants] | 622
- [three-color-policer] | 642
- [transport-ls] | 642
- [vlans] | 643

Sub Hierarchies

"[edit logical-systems {logical-system-name} ...]" on page 357

"[edit logical-systems access]" on page 358

"[edit logical-systems {logical-system-name interfaces} ae{X} aggregated-ether-options]" on page 359

"[edit logical-systems applications]" on page 359

"[edit logical-systems {logical-system-name} bridge-domains {bridge-domain-name}]" on page 359

"[edit logical-systems class-of-service]" on page 362

"[edit logical-systems {logical-system-name} firewall]" on page 363

"[edit logical-systems firewallpolicer]" on page 366

"[edit logical-systems {logical-system-name} forwarding-options]" on page 366

"[edit logical-systems {logical-system-name} interfaces]" on page 377

"[edit logical-systems l2circuit]" on page 392

"[edit logical-systems logical-system-name]" on page 392

"[edit logical-systems name]" on page 394

"[edit logical-systems neighbor]" on page 394

"[edit logical-systems policer]" on page 394

"[edit logical-systems {name} policy-options]" on page 395

"[edit logical-systems policy-statement]" on page 398

"[edit logical-systems {logical-system-name} protocols]" on page 399

"[edit logical-systems realm]" on page 473

"[edit logical-systems routing]" on page 473

"[edit logical-systems {logical-system-name} routing-instances {routing-instance-name}]" on page 473

"[edit logical-systems routing-instancesrouting-options]" on page 584

"[edit logical-systems {logical-system-name} routing-options]" on page 584

"[edit logical-systems {name} security]" on page 599

"[edit logical-systems {logical-system-name} services]" on page 609

"[edit logical-systems statement-name]" on page 610

"[edit logical-systems {name} switch-options]" on page 610

"[edit logical-systems system]" on page 611

"[edit logical-systems tenants]" on page 622

"[edit logical-systems {logical-system-name} three-color-policer {policer-name}]" on page 642

"[edit logical-systems transport-ls]" on page 642

"[edit logical-systems vlans]" on page 643

[edit logical-systems {name}]		
"enforce-first-as" on page 6343	"protocols (STP Type)" on page 13802	"routing-instances (Dynamic Profiles)" on page 15973
"routing-instances (Logical Systems)" on page 15966	"routing-instances (Multiple Routing Entities)" on page 15976	"routing-options" on page 15980
"switch-options" on page 28749		

[...]

[edit logical-systems {logical-system-name} ...]		
"active-leasequery (DHCP Relay Agent)" on page 1808	"bulk-leasequery (DHCP Relay Agent)" on page 3068	"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279
"duplicate-clients (DHCPv6 Local Server and Relay Agent)" on page 6064	"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699	"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701

(Continued)

[edit logical-systems {logical-system-name} ...]		
"leasequery (DHCP Relay Agent)" on page 9702	"no-snoop (DHCP Local Server and Relay Agent)" on page 11582	"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334
"route-suppression (DHCP Local Server and Relay Agent)" on page 15877	"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	

[edit logical-systems {logical-system-name} ... forwarding-options dhcp-relay ... relay-option-82 (circuit-id | remote-id)]		
"include-irb-and-l2" on page 8529	"no-vlan-interface-name" on page 11617	"prefix (DHCP Relay Agent)" on page 13077

[edit logical-systems {logical-system-name} ... forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)]		
"include-irb-and-l2" on page 8529	"no-vlan-interface-name" on page 11617	"prefix (DHCP Relay Agent)" on page 13077

[access]

[edit logical-systems {name} access address-assignment pool {name} family (inet | inet6)]	
"excluded-range (Address-Assignment Pools)" on page 6632	

[aeaggregated-ether-options]

[edit logical-systems {logical-system-name interfaces} ae{X} aggregated-ether-options]	
"mc-ae" on page 10575	

[applications]

[edit logical-systems {name} applications application]	
"application-protocol (Applications)" on page 2355	

[bridge-domains]

[edit logical-systems {logical-system-name} bridge-domains {bridge-domain-name}]		
"routing-interface" on page 15978	"vlan-id-list" on page 30805	"vlan-tags" on page 30865

[edit logical-systems {logical-system-name} bridge-domains {bridge-domain-name} bridge-options]		
"interface-mac-limit" on page 8910	"mac-statistics" on page 10263	"mac-table-size" on page 10271

[edit logical-systems {logical-system-name} bridge-domains {bridge-domain-name} bridge-options interface {interface-name}]		
"interface-mac-limit" on page 8910	"mac-pinning" on page 10254	"static-mac" on page 28549

[edit logical-systems {logical-system-name} bridge-domains {bridge-domain-name} bridge-options interface {interface-name} interface-mac-limit {limit}]		
"packet-action" on page 12129		

[edit logical-systems {logical-system-name} bridge-domains {bridge-domain-name} bridge-options interface-mac-limit {limit}]		
"packet-action" on page 12129		

[edit logical-systems {name} bridge-domains {name} forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit logical-systems {name} bridge-domains {name} forwarding-options dhcp-relay dhcpv6 group {name} relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {name} bridge-domains {name} forwarding-options dhcp-relay dhcpv6 relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {name} bridge-domains {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit logical-systems {name} bridge-domains {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit logical-systems {name} bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} bridge-domains {name} forwarding-options dhcp-relay relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} bridge-domains {name} forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} bridge-domains {name} forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {logical-system-name} bridge-domains {bridge-domains} interface]	
"protect-interface" on page 13409	

[edit logical-systems {logical-system-name} bridge-domains {domain-name} protocols igmp-snooping]	
"traceoptions (Protocols IGMP Snooping)" on page 29678	

[class-of-service]

[edit logical-systems {logical-system-name} class-of-service application-traffic-control]		
"rate-limiters" on page 14071	"rule-sets (CoS AppQoS)" on page 16149	

[edit logical-systems {logical-system-name} class-of-service interfaces {interface-name}]	
"member-link-scheduler" on page 10614	

[firewall]

[edit logical-systems {logical-system-name} firewall]		
"family (Firewall)" on page 6867	"interface-set" on page 9034	"policer (Configuring)" on page 12630
"three-color-policer" on page 29309	"three-color-policer (Configuring)" on page 29314	

[edit logical-systems {logical-system-name} firewall family {family-name}]	
"filter (Configuring)" on page 7127	

[edit logical-systems {logical-system-name} firewall family (inet | inet6)]	
"service-filter (Firewall)" on page 16699	

[edit logical-systems {logical-system-name} firewall family (inet | inet6) filter {filter-name}]	
"filter-list-template" on page 7148	

[edit logical-systems {logical-system-name} firewall family {family-name} filter {filter-name}]		
"enhanced-mode" on page 6396	"enhanced-mode-override" on page 6401	"fast-lookup-filter" on page 6928
"instance-shared" on page 8769	"interface-specific (Firewall Filters)" on page 9063	"physical-interface-filter" on page 12476
"scale-optimized" on page 16329	"term (Firewall Filter)" on page 29117	

[edit logical-systems {logical-system-name} firewall family {family-name} filter {filter-name} term {term-name} then]		
"force-premium (Firewall Filter Action)" on page 7373	"forwarding-class (Firewall Filter Action)" on page 7423	"loss-priority (Firewall Filter Action)" on page 10154
"policer (Firewall Filter Action)" on page 12636	"three-color-policer (Applying)" on page 29311	

[edit logical-systems {logical-system-name} firewall family inet]		
"prefix-action (Configuring)" on page 13068	"simple-filter" on page 27962	

[edit logical-systems {name} firewall family inet filter {name} term {name} from]	
"loss-priority (Firewall Filter)" on page 10152	

[edit logical-systems {logical-system-name} firewall family inet filter {filter-name} term {term-name} then]	
"prefix-action (Firewall Filter Action)" on page 13070	

[edit logical-systems {logical-system-name} firewall family inet prefix-action {name}]	
"filter-specific" on page 7155	

[edit logical-systems {logical-system-name} firewall family {family-name} service-filter {filter-name}]	
"term (Firewall Filter)" on page 29117	

[edit logical-systems {logical-system-name} firewall family {family-name} simple-filter {filter-name}]	
"term (Firewall Filter)" on page 29117	

[edit logical-systems {logical-system-name} firewall filter {filter-name}]		
"enhanced-mode" on page 6396	"enhanced-mode-override" on page 6401	

[edit logical-systems {name} firewall filter{name} term {name} from]	
"loss-priority (Firewall Filter)" on page 10152	

[edit logical-systems {logical-system-name} firewall policer {policer-name}]		
"filter-specific" on page 7155	"if-exceeding (Policer)" on page 8373	"if-exceeding-pps (Policer)" on page 8376
"logical-bandwidth-policer" on page 10058	"logical-interface-policer" on page 10069	"physical-interface-policer" on page 12479
"then (Policer Action)" on page 29279		

[edit logical-systems {logical-system-name} firewall policer {policer-name} if-exceeding]		
"bandwidth-limit" on page 2853	"burst-size-limit" on page 3083	

[edit logical-systems {logical-system-name} firewall policer {policer-name} if-exceeding-pps]	
"packet-burst (Policer)" on page 12135	

[edit logical-systems {logical-system-name} firewall three-color-policer {name}]		
"action" on page 1759	"logical-interface-policer" on page 10069	"single-rate" on page 27969
"two-rate" on page 30259		

[edit logical-systems {logical-system-name} firewall three-color-policer {policer-name} action]	
"loss-priority high then discard (Three-Color Policer)" on page 10155	

[firewallpolicer]

[edit logical-systems {logical-system-name} firewallpolicer {policer-name} if-exceeding-pps]	
"pps-limit (Policer)" on page 12984	

[forwarding-options]

[edit logical-systems {logical-system-name} forwarding-options]		
"dhcp-relay" on page 5620	"load-balance (Forwarding Options)" on page 9878	"storm-control-profiles" on page 28624

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay]		
"active-server-group" on page 1817	"dhcpv6 (DHCP Relay Agent)" on page 5665	"dual-stack-group (DHCP Relay Agent)" on page 6055
"duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)" on page 6066	"forward-snooped-clients (DHCP Relay Agent)" on page 7405	"relay-option-82" on page 14294
"server-group" on page 16633		

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay ...]		
"allow-snooped-clients" on page 2204	"asymmetric-lease-time (DHCP Overrides)" on page 2519	"authentication (DHCP Relay Agent)" on page 2595
"circuit-type (DHCP Relay Agent)" on page 3332	"default-action (DHCP Relay Agent Option)" on page 5177	"drop (DHCP Relay Agent Option)" on page 5946
"dual-stack (DHCP Relay Agent Overrides)" on page 6050	"dynamic-profile (DHCP Relay Agent)" on page 6119	"forward-only (DHCP Relay Agent Option)" on page 7398
"forward-only (DHCP Relay Agent)" on page 7396	"forward-only-replies (DHCP Relay Agent)" on page 7401	"group (DHCP Relay Agent)" on page 7820
"interface (DHCP Relay Agent)" on page 8863	"interface-description (DHCP Relay Agent)" on page 8856	"interface-name (DHCP Relay Agent)" on page 8938
"local-server-group (DHCP Relay Agent Option)" on page 9976	"mac-address (DHCP Relay Agent)" on page 10223	"no-allow-snooped-clients" on page 11390
"option-number (DHCP Relay Agent Option)" on page 11898	"overrides (DHCP Relay Agent)" on page 12068	"relay-option (DHCP Relay Agent)" on page 14297
"server-response-time (DHCP Relay Agent)" on page 16664	"service-profile (DHCP Relay Agent)" on page 16731	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982
"username-include (DHCP Relay Agent)" on page 30597		

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay ... overrides]		
"client-discover-match (DHCP Relay Agent)" on page 4310	"delay-authentication (DHCP Relay Agent)" on page 5242	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay ... relay-option-82]	
"remote-id (DHCP Relay Agent)" on page 14331	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay authentication]	
"password (DHCP Relay Agent)" on page 12229	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083
"option-82 (DHCP Relay Agent)" on page 11893	"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6]		
"active-server-group" on page 1817	"server-group" on page 16633	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)]	
"use-interface-description" on page 30522	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 ...]		
"client-id (DHCP Relay Agent)" on page 4317	"relay-agent-interface-id (DHCPv6 Relay Agent)" on page 14276	"relay-agent-remote-id (DHCPv6 Relay Agent)" on page 14284
"use-vlan-id (DHCP Relay Agent)" on page 30546		

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 ... (relay-agent-interface-id | relay-agent-remote-id)]	
"use-option-82" on page 30528	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 authentication]	
"password (DHCP Relay Agent)" on page 12229	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083
"relay-agent-interface-id (DHCPv6 Relay Agent Username)" on page 14278	"relay-agent-remote-id (DHCPv6 Relay Agent Username)" on page 14286	"relay-agent-subscriber-id (DHCPv6 Relay Agent)" on page 14289
"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 group {group-name}]	
"active-server-group" on page 1817	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 group {group-name} authentication]		
"password (DHCP Relay Agent)" on page 12229		

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 group {group-name} authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083
"relay-agent-interface-id (DHCPv6 Relay Agent Username)" on page 14278	"relay-agent-remote-id (DHCPv6 Relay Agent Username)" on page 14286	"relay-agent-subscriber-id (DHCPv6 Relay Agent)" on page 14289
"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580	

[edit logical-systems {name} forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 group {group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 group {group-name} interface {interface-name}]	
"trace (DHCP Relay Agent)" on page 29458	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 group {group-name} overrides]		
"interface-client-limit (DHCP Relay Agent)" on page 9108	"no-bind-on-request (DHCP Relay Agent)" on page 11418	"send-release-on-delete (DHCP Relay Agent)" on page 16550

[edit logical-systems {name} forwarding-options dhcp-relay dhcpv6 group {name} relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 overrides]		
"interface-client-limit (DHCP Relay Agent)" on page 9108	"no-bind-on-request (DHCP Relay Agent)" on page 11418	"send-release-on-delete (DHCP Relay Agent)" on page 16550

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]	
"keep-incoming-interface-id (DHCP Relay Agent)" on page 9462	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id]	
"keep-incoming-remote-id (DHCP Relay Agent)" on page 9463	

[edit logical-systems {name} forwarding-options dhcp-relay dhcpv6 relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dhcpv6...]		
"asymmetric-lease-time (DHCP Overrides)" on page 2519	"asymmetric-prefix-lease-time (DHCP Overrides)" on page 2521	

[edit logical-systems {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"dual-stack-interface-client-limit (DHCP Local Server and Relay Agent)" on page 6058	"protocol-master" on page 13448
"relay-agent-interface-id (DHCPv6 Relay Agent)" on page 14276	"use-vlan-id (DHCP Relay Agent)" on page 30546	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit logical-systems {name} forwarding-options dhcp-relay group]		
"interface-tag (DHCP Relay Server)" on page 9083	"relay-option-82" on page 14294	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay group {group-name} authentication]	
"password (DHCP Relay Agent)" on page 12229	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay group {group-name} authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083
"option-82 (DHCP Relay Agent)" on page 11893	"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay group {group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay group {group-name} interface {interface-name}]	
"trace (DHCP Relay Agent)" on page 29458	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay group {group-name} interface {interface-name} overrides]	
"relay-source" on page 14302	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay group {group-name} overrides]		
"always-write-giaddr" on page 2221	"always-write-option-82" on page 2222	"disable-relay" on page 5801
"interface-client-limit (DHCP Relay Agent)" on page 9108	"layer2-unicast-replies" on page 9661	"no-bind-on-request (DHCP Relay Agent)" on page 11418
"proxy-mode" on page 13833	"relay-source" on page 14302	"replace-ip-source-with (DHCP Relay Agent)" on page 14377
"send-release-on-delete (DHCP Relay Agent)" on page 16550	"trust-option-82" on page 30168	

[edit logical-systems {name} forwarding-options dhcp-relay group {name} relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877

(Continued)

[edit logical-systems {name} forwarding-options dhcp-relay group {name} relay-option]		
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay group {group-name} relay-option-60]	
"vendor-option" on page 30684	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay group {group-name} relay-option-82]		
"circuit-id (DHCP Relay Agent)" on page 3324	"server-id-override" on page 16637	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay overrides]		
"always-write-giaddr" on page 2221	"always-write-option-82" on page 2222	"disable-relay" on page 5801
"interface-client-limit (DHCP Relay Agent)" on page 9108	"layer2-unicast-replies" on page 9661	"no-bind-on-request (DHCP Relay Agent)" on page 11418
"proxy-mode" on page 13833	"relay-source" on page 14302	"replace-ip-source-with (DHCP Relay Agent)" on page 14377

(Continued)

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay overrides]		
"send-release-on-delete (DHCP Relay Agent)" on page 16550	"trust-option-82" on page 30168	

[edit logical-systems {name} forwarding-options dhcp-relay relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay relay-option-60]	
"vendor-option" on page 30684	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay relay-option-82]		
"circuit-id (DHCP Relay Agent)" on page 3324	"server-id-override" on page 16637	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)]	
"use-interface-description" on page 30522	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay relay-option-82 circuit-id]	
"keep-incoming-circuit-id (DHCP Relay Agent)" on page 9460	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relay relay-option-82 XXXcircuit-id]	
"keep-incoming-remote-id (DHCP Relay Agent)" on page 9463	

[edit logical-systems {logical-system-name} forwarding-options dhcp-relaygroup {group-name}]	
"active-server-group" on page 1817	

[edit logical-systems {logical-system-name} forwarding-options family vpls]	
"filter (VPLS)" on page 7158	

[edit logical-systems {logical-system-name} forwarding-options load-balance]		
"indexed-load-balance" on page 8545	"per-flow" on page 12413	"per-prefix" on page 12417

[interfaces]

[edit logical-systems {logical-system-name} interfaces]		
"demux0 (Static Interface)" on page 5313	"family bridge" on page 6847	"shdsl-options" on page 16974
"unidirectional" on page 30356	"unit (interfaces)" on page 30370	"unit (Tunnel)" on page 30404
"vlan-tagging" on page 30861		

[edit logical-systems {logical-system-name} interfaces (ls-{fpc}/{pic}/{port}| ml-{fpc}/{pic}/{port}) unit {logical-unit-number}]		
"drop-timeout" on page 5981	"fragment-threshold" on page 7542	"short-sequence" on page 16989

[edit logical-systems {logical-system-name} interfaces (ml-{fpc}/{pic}/{port}| ls-{fpc}/{pic}/{port}) unit {logical-unit-number}]	
"mrru" on page 10989	

[edit logical-systems {logical-system-name} interfaces ae{X} aggregated-ether-options]	
"lacp (Aggregated Ethernet)" on page 9610	

[edit logical-systems {logical-system-name} interfaces at-{fpc}/{pic}/{port} unit {logical-unit-number}]	
"vci" on page 30660	

[edit logical-systems {logical-system-name} interfaces at-{fpc}/{pic}/{port} unit {logical-unit-number} address {address} family {family} multipoint-destination {address} shaping]	
"cbr" on page 3238	

[edit logical-systems {logical-system-name} interfaces at-{fpc}/{pic}/{port} unit {logical-unit-number} family {family} address {address} multipoint-destination {address}]	
"vci" on page 30660	

[edit logical-systems {logical-system-name} interfaces at-{fpc}/{pic}/{port} unit {logical-unit-number} shaping]	
"cbr" on page 3238	

[edit logical-systems {logical-system-name} interfaces br-{pim}/0/{port} unit {logical-unit-number} dialer-options]	
"dial-string" on page 5687	

[edit logical-systems {logical-system-name} interfaces demux0 unit {logical-unit-number} demux-options]	
"underlying-interface" on page 30330	

[edit logical-systems {logical-system-name} interfaces dl{n} unit {logical-unit-number}]	
"dialer-options" on page 5689	

[edit logical-systems {logical-system-name} interfaces dl{n} unit {logical-unit-number} dialer-options]		
"callback-wait-period" on page 3183	"incoming-map" on page 8540	"load-interval" on page 9892
"load-threshold" on page 9894	"pool" on page 12762	"redial-delay" on page 14157

[edit logical-systems {logical-system-name} interfaces dl{n} unit {logical-unit-number} dialer-options incoming-map]		
"callback" on page 3182	"caller" on page 3188	

[edit logical-systems {logical-system-name} interfaces fxp0 unit {logical-unit-number} family {family} address]		
"master-only" on page 10364		

[edit logical-systems {logical-system-name} interfaces gr-{fpc}/{pic}/{port} unit {logical-unit-number}]		
"clear-dont-fragment-bit (Interfaces GRE Tunnels) " on page 3557	"copy-tos-to-outer-ip-header" on page 4666	"copy-tos-to-outer-ip-header" on page 4668
"reassemble-packets" on page 14100		

[edit logical-systems {logical-system-name} interfaces gr-{fpc}/{pic}/{port} unit {logical-unit-number} tunnel]		
"allow-fragmentation" on page 2179	"do-not-fragment" on page 5871	

[edit logical-systems {logical-system-name} interfaces gre unit {logical-unit-number}]		
"copy-tos-to-outer-ip-header" on page 4666	"copy-tos-to-outer-ip-header" on page 4668	

[edit logical-systems {logical-system-name} interfaces interface-name unit {logical-unit-number} family inet address {address}]	
"destination (Tunnels)" on page 5528	

[edit logical-systems {logical-system-name} interfaces ls-{fpc}/{pic}/{port}:{channel} mlfr-uni-nni-bundle-options]		
"drop-timeout" on page 5981	"fragment-threshold" on page 7542	"mrru" on page 10989

[edit logical-systems {logical-system-name} interfaces ls-{fpc}/{pic}/{port} unit {logical-unit-number}]	
"interleave-fragments" on page 9185	

[edit logical-systems {logical-system-name} interfaces lsq-{fpc}/{pic}/{port} unit {logical-unit-number}]	
"fragment-threshold (Interfaces LSQ)" on page 7545	

[edit logical-systems {logical-system-name} interfaces lsq-{fpc}/{pic}/{port} unit {logical-unit-number} compression rtp]	
"port (Services Voice) " on page 12884	

[edit logical-systems {logical-system-name} interfaces pp0 unit {logical-unit-number}]		
"pppoe-options" on page 12974	"pppoe-options (SRX Series)" on page 12972	

[edit logical-systems {logical-system-name} interfaces pp0 unit {logical-unit-number} pppoe-options]		
"client (interfaces)" on page 4299	"server" on page 16608	"service-name" on page 16714
"underlying-interface" on page 30330		

[edit logical-systems {logical-system-name} interfaces {interface-name}unit {logical-unit-number} ppp-options]		
"lcp-restart-timer" on page 9672	"loopback-clear-timer" on page 10119	

[edit logical-systems {logical-system-name} interfaces {interface-name} shdsl-options]		
"line-rate" on page 9778	"snr-margin" on page 28081	

[edit logical-systems {logical-system-name} interfaces {interface-name} shdsl-options snr-margin]	
"snext" on page 28056	

[edit logical-systems {logical-system-name} interfaces si-{fpc}/{pic}/{port} unit {logical-unit-number}]	
"dial-options" on page 5684	

[edit logical-systems {logical-system-name} interfaces sp-{fpc}/{pic}/{port} unit {logical-unit-number}]	
"dial-options" on page 5684	

[edit logical-systems {logical-system-name} interfaces sp-{fpc}/{pic}/{port} unit {logical-unit-number} interface]	
"l2tp-interface-id" on page 9550	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number}]		
"accept-source-mac" on page 1647	"alias (Interfaces)" on page 2146	"backup-options" on page 2807
"bandwidth (Interfaces)" on page 2847	"compression" on page 4491	"compression-device (Interfaces)" on page 4493
"demux-destination (Underlying Interface)" on page 5297	"demux-options (Static Interface)" on page 5304	"demux-source (Underlying Interface)" on page 5309
"description (Interfaces)" on page 5334	"disable (Interface)" on page 5765	"disable-mlppp-inner-ppp-pfc" on page 5772
"dlci" on page 5843	"encapsulation (Logical Interface)" on page 6300	"encapsulation (logical-link-layer)" on page 6286
"family (Interfaces)" on page 6881	"family (interfaces)" on page 6832	"filter (Applying to a Logical Interface)" on page 7123
"host-prefix-only" on page 8163	"input-vlan-map " on page 8729	"input-vlan-map (Aggregated Ethernet)" on page 8731
"keepalives" on page 9475	"layer2-policer (EX)" on page 9654	"link-layer-overhead" on page 9797

(Continued)

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number}]		
"minimum-links" on page 10741	"multicast-dlci" on page 11066	"multilink-max-classes" on page 11106
"multipoint" on page 11127	"no-keepalives" on page 11488	"output-vlan-map " on page 12021
"output-vlan-map (Aggregated Ethernet)" on page 12022	"peer-unit" on page 12390	"point-to-point" on page 12619
"policer (Applying to a Logical Interface)" on page 12624	"ppp-options" on page 12949	"pppoe-underlying-options (Static and Dynamic Subscribers)" on page 12980
"proxy-arp" on page 13820	"shaping" on page 16925	"targeted-distribution" on page 28958
"traps" on page 30142	"tunnel" on page 30192	"vlan-id (Fast Ethernet)" on page 30794
"vlan-id (VLAN ID to Be Bound to a Logical Interface)" on page 30820	"vlan-id-list (Ethernet VLAN Circuit)" on page 30808	"vlan-id-list (Interface in VPLS)" on page 30812
"vlan-id-range" on page 30815	"vlan-tags (Dual-Tagged Logical Interface)" on page 30869	"vlan-tags (Stacked VLAN Tags)" on page 30873

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} accept-source-mac]	
"mac-address (Accept Source Mac)" on page 10219	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} accept-source-mac mac-address {mac-address}]	
"policer (MAC)" on page 12641	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} address {address} family {family} multipoint-destination {address}]	
"shaping" on page 16925	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} compression]	
"rtp" on page 16068	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} compression rtp]		
"f-max-period (rtp)" on page 6744	"maximum-contexts" on page 10517	"queues" on page 13955

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} dynamic-call-admission-control]	
"activation-priority" on page 1799	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family {family}]		
"address" on page 1846	"demux-destination (Demux Interface)" on page 5296	"demux-source (Demux Interface)" on page 5305
"dialer" on page 5688	"filter (Applying to a Logical Interface)" on page 7123	"filter (Applying to an Interface)" on page 7125
"filter (Interfaces)" on page 7117	"ingress-queuing-filter" on page 8644	"mtu (interfaces)" on page 11027
"mtu (Multilink and Link Services Logical Interface)" on page 11025	"policer (Applying to a Logical Interface)" on page 12624	"policer (Interface)" on page 12638
"primary (Interface for Router)" on page 13140	"rpf-check " on page 16007	"rpf-check " on page 16007

(Continued)

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family {family}]		
"unnumbered-address (Ethernet)" on page 30425	"vpls (Interfaces)" on page 30931	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family (ccc | inet | tcc)]	
"policer (Layer 2 VPN)" on page 12639	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family (inet | inet6) rpf-check]	
"mode (Interfaces)" on page 10804	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family {family} address {address}]		
"broadcast (Interfaces)" on page 3041	"preferred" on page 13053	"primary (Address on Interface)" on page 13130

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family bridge]		
"interface-mode" on page 8920	"storm-control" on page 28620	"vlan-id (Logical Port in Bridge Domain)" on page 30813
"vlan-id-list (Interface in Bridge Domain)" on page 30810		

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family ccc]	
"keep-address-and-control" on page 9455	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family {family} filter]		
"input-list" on page 8708		

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet]		
"accounting (interfaces)" on page 1709	"dhcp (DHCP Client)" on page 5598	"multicast-only" on page 11076
"negotiate-address" on page 11261	"receive-options-packets" on page 14116	"receive-ttl-exceeded" on page 14121
"sampling (Interfaces)" on page 16298	"service" on page 16680	"service-domain" on page 16693
"simple-filter (Applying to an Interface)" on page 27964	"targeted-broadcast" on page 28956	"unnumbered-address (Demux)" on page 30419
"unnumbered-address (PPP)" on page 30426		

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet accounting]		
"destination-class-usage" on page 5411	"source-class-usage" on page 28191	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet address {address}]		
"arp" on page 2458	"destination (Interfaces)" on page 5433	"destination-profile" on page 5509
"vrrp-group" on page 30973		

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet address {address} vrrp-group {group-id}]		
"accept-data" on page 1631	"advertise-interval" on page 2017	"authentication-key" on page 2603
"authentication-type" on page 2667	"fast-interval" on page 6926	"preempt (VRRP)" on page 13025
"priority (Protocols VRRP)" on page 13201	"track (VRRP)" on page 29992	"virtual-address" on page 30739

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet address {address} vrrp-group {group-id} preempt]	
"hold-time (VRRP)" on page 8119	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet address {address} vrrp-group {group-id} track]		
"interface (VRRP)" on page 8812	"priority-hold-time" on page 13179	"route (Interfaces)" on page 15868

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet address {address} vrrp-group {group-id} track interface {interface-name}]		
"bandwidth-threshold" on page 2878	"priority-cost (VRRP)" on page 13163	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet service]		
"input (Interfaces)" on page 8707	"output (Service Set)" on page 11978	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet service (input | output)]	
"service-set (Interfaces)" on page 16744	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet service (input | output) service-set {service-set-name}]	
"service-filter (Interfaces)" on page 16701	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet service input]	
"post-service-filter" on page 12907	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet targeted-broadcast]		
"forward-and-send-to-re" on page 7392	"forward-only" on page 7395	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet unnumbered-address {interface-name}]		
"destination (IPCP)" on page 5439	"destination (Tunnels)" on page 5528	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet unnumbered-address {interface-name} destination {address}]	
"destination-profile" on page 5509	

[edit logical-systems {name} interfaces {name} unit {name} family inet6]	
"dhcpv6-client" on page 5655	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet6 address {address}]	
"vrrp-inet6-group" on page 30976	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet6 address {address} vrrp-inet6-group {group-id}]		
"accept-data" on page 1631	"fast-interval" on page 6926	"inet6-advertise-interval (VRRP)" on page 8596
"preempt (VRRP)" on page 13025	"priority (Protocols VRRP)" on page 13201	"track (VRRP)" on page 29992
"virtual-inet6-address" on page 30760	"virtual-link-local-address" on page 30765	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet6 address {address} vrrp-inet6-group {group-id} preempt]	
"hold-time (VRRP)" on page 8119	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet6 address {address} vrrp-inet6-group {group-id} track]		
"interface (VRRP)" on page 8812	"priority-hold-time" on page 13179	"route (Interfaces)" on page 15868

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet6 address {address} vrrp-inet6-group {group-id} track interface {interface-name}]		
"bandwidth-threshold" on page 2878	"priority-cost (VRRP)" on page 13163	

[edit logical-systems {logical -system-name} interfaces {interface-name} unit {logical-unit-number} family inet6 dhcpv6-client]		
"client-ia-type" on page 4313	"client-identifier (DHCPv6 Client)" on page 4321	"client-type" on page 4340
"rapid-commit (DHCPv6 Client)" on page 14054	"req-option" on page 14408	"retransmission-attempt (DHCPv6 Client)" on page 15678

(Continued)

[edit logical-systems {logical -system-name} interfaces {interface-name} unit {logical-unit-number} family inet6 dhcpv6-client]		
"update-router-advertisement" on page 30449		

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family inet6 dhcpv6-client prefix-delegating]		
"preferred-prefix-length" on page 13059	"sub-prefix-length" on page 28674	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family mpls]		
"maximum-labels" on page 10528		

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family pppoe]		
"access-concentrator" on page 1655	"direct-connect" on page 5708	"duplicate-protection (Dynamic PPPoE)" on page 6071
"dynamic-profile (Dynamic PPPoE)" on page 6124	"max-sessions (Dynamic PPPoE)" on page 10469	"max-sessions-vsa-ignore (Static and Dynamic Subscribers)" on page 10481
"service-name-table" on page 16716	"short-cycle-protection (Static and Dynamic Subscribers)" on page 16984	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} family tcc]		
"proxy (Interfaces)" on page 13828	"proxy (interfaces)" on page 13817	"remote" on page 14305

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} input-vlan-map]		
"inner-tag-protocol-id" on page 8679	"inner-vlan-id" on page 8682	"pop" on page 12799
"pop-pop" on page 12806	"pop-swap" on page 12807	"push" on page 13875
"push-push" on page 13879	"swap" on page 28740	"swap-push" on page 28744
"swap-swap" on page 28745	"tag-protocol-id (TPID to Rewrite)" on page 28920	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} layer2-policer]		
"input-policer" on page 8713	"input-three-color" on page 8724	"output-policer" on page 12001
"output-three-color" on page 12014		

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} output-vlan-map]		
"inner-tag-protocol-id" on page 8679	"inner-vlan-id" on page 8682	"pop" on page 12799
"pop-pop" on page 12806	"pop-swap" on page 12807	"push" on page 13875
"push-push" on page 13879	"swap" on page 28740	"swap-push" on page 28744
"swap-swap" on page 28745	"tag-protocol-id (TPID to Rewrite)" on page 28920	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} ppp-options]		
"chap" on page 3289	"compression (PPP Properties)" on page 4494	"ncp-restart-timer" on page 11245
"pap" on page 12169		

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} ppp-options chap]		
"access-profile" on page 1688	"default-chap-secret" on page 5184	"local-name" on page 9964
"passive (CHAP)" on page 12198		

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} ppp-options compression]		
"acfc" on page 1748	"pfc" on page 12455	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} ppp-options pap]		
"access-profile" on page 1688	"default-pap-password" on page 5199	"local-name" on page 9964
"local-password" on page 9966	"passive (PAP)" on page 12214	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} pppoe-options]		
"access-concentrator" on page 1655	"auto-reconnect" on page 2732	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} pppoe-underlying-options]		
"access-concentrator" on page 1655	"direct-connect" on page 5708	"duplicate-protection (Dynamic PPPoE)" on page 6071
"dynamic-profile (Dynamic PPPoE)" on page 6124	"max-sessions (Dynamic PPPoE)" on page 10469	"max-sessions-vsa-ignore (Static and Dynamic Subscribers)" on page 10481
"service-name-table" on page 16716	"short-cycle-protection (Static and Dynamic Subscribers)" on page 16984	

[edit logical-systems {logical-system-name} interfaces {interface-name} unit {logical-unit-number} tunnel]		
"backup-destination" on page 2795	"destination (Tunnel Remote End)" on page 5526	"destination (Tunnels)" on page 5528
"key (gre)" on page 9482	"key (interfaces)" on page 9484	"routing-instance (Interfaces)" on page 15923
"source (Interfaces)" on page 28112	"traffic-class (Tunnels)" on page 30003	

[l2circuit]

[edit logical-systems {logical-system-name} l2circuit neighbor {address} interface {interface-name}]	
"static (Protocols Layer 2 Circuit)" on page 28559	

[edit logical-systems {logical-system-name} l2circuit neighbor {address} interface {interface-name} backup-neighbor {neighbor}]	
"static (Protocols Layer 2 Circuit)" on page 28559	

[logical-system-name]

[edit logical-systems logical-system-name firewall family family-name filter output filter-name]	
"fast-lookup-filter (PTX)" on page 6930	

[edit logical-systems logical-system-name firewall policer policer-name]	
"if-exceeding (EX)" on page 8370	

[edit logical-systems logical-system-name protocols bgp family]	
"inet-mdt (Signaling)" on page 8571	

[edit logical-systems logical-system-name protocols bgp group group-name family]	
"inet-mdt (Signaling)" on page 8571	

[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family]	
"inet-mdt (Signaling)" on page 8571	

[edit logical-systems logical-system-name protocols mpls]	
"exclude-srlg" on page 6629	

[edit logical-systems logical-system-name routing-instances instance-name forwarding-options]	
"enhanced-hash-key" on page 6362	

[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-next-hop transit]	
"l3vpn (transit)" on page 9573	

[name]

[edit logical-systems name routing-options rib {rib name} static route {destination}]	
"spring-te-lsp-next-hop" on page 28397	

[edit logical-systems name routing-options static route {destination}]	
"spring-te-lsp-next-hop" on page 28397	

[edit logical-systems name security nat {destination pool}]	
"address (Security Destination NAT)" on page 1924	

[neighbor]

[edit logical-systems {logical-system-name}protocols bgp group {group-name} neighbor {address} local-as]	
"loops (Autonomous System)" on page 10132	

[policer]

[edit logical-systems {logical-system-name} policer {policer-name} if-exceeding]		
"bandwidth-limit (Policer)" on page 2857	"bandwidth-percent" on page 2866	"burst-size-limit (Policer)" on page 3086

[policy-options]

[edit logical-systems {name} policy-options]		
"as-list" on page 2487	"as-path (Policy Options)" on page 2493	"as-path-group" on page 2492
"community (Policy Options)" on page 4465	"community (Protocols Layer 2 Circuit)" on page 4469	"condition" on page 4502
"damping (Policy Options)" on page 4820	"defaults (Policy Options)" on page 5228	"policy-statement" on page 12736
"prefix-list" on page 13090	"resolution-map" on page 15597	"rib-list" on page 15783
"route-filter-list" on page 15863	"rtf-prefix-list" on page 16064	"source-address-filter-list" on page 28131

[edit logical-systems {logical-system-name} policy-options as-path {path-name}]	
"dynamic-db" on page 6105	

[edit logical-systems {logical-system-name} policy-options as-path-group {group-name}]	
"dynamic-db" on page 6105	

[edit logical-systems {logical-system-name} policy-options community {community-name}]	
"dynamic-db" on page 6105	

[edit logical-systems {logical-system-name} policy-options condition {condition-name}]		
"dynamic-db" on page 6105	"if-route-exists" on page 8377	

[edit logical-systems {logical-system-name} policy-options condition if-route-exists]		
"address-family" on page 1876	"table" on page 28895	

[edit logical-systems {logical-system-name} policy-options condition if-route-exists address-family]		
"ccc (Routing Policy Condition)" on page 3242	"inet (Routing Policy Condition)" on page 8588	

[edit logical-systems {logical-system-name} policy-options condition if-route-exists address-family ccc]		
"peer-unit (Routing Policy Condition)" on page 12391	"standby (Routing Policy Condition)" on page 28491	"table" on page 28895

[edit logical-systems {logical-system-name} policy-options condition if-route-exists address-family inet]	
"table" on page 28895	

[edit logical-systems {logical-system-name} policy-options defaults]	
"route-filter" on page 15861	

[edit logical-systems {logical-system-name} policy-options defaults route-filter]		
"no-walkup" on page 11623	"walkup" on page 31019	

[edit logical-systems {logical-system-name} policy-options policy-statement {policy-statement-name}]		
"defaults (Policy Options)" on page 5228	"dynamic-db" on page 6105	

[edit logical-systems {logical-system-name} policy-options policy-statement {policy-statement-name} defaults]	
"route-filter" on page 15861	

[edit logical-systems {logical-system-name} policy-options policy-statement {policy-statement-name} defaults route-filter]		
"no-walkup" on page 11623	"walkup" on page 31019	

[edit logical-systems {name} policy-options policy-statement {name} from]	
"prefix-segment" on page 13105	

[edit logical-systems {name} policy-options policy-statement {name} from route-filter]	
"get-route-range" on page 7681	

[edit logical-systems {name} policy-options policy-statement {name} from traffic-engineering]		
"ipv6-prefix" on page 9368	"srv6-sid" on page 28419	

[edit logical-systems {logical-system-name} policy-options policy-statement {policy-name} term {term-name}]	
"rtf-prefix-list" on page 16064	

[edit logical-systems {logical-system-name} policy-options policy-statement {policy-statement-name} term {term--name} from]	
"prefix-list-filter" on page 13092	

[edit logical-systems {name} policy-options policy-statement {name} term {name} from route-filter]	
"get-route-range" on page 7681	

[edit logical-systems {logical-system-name} policy-options policy-statement {policy-name} term {term-name} from traffic-engineering]	
"ipv4-prefix" on page 9338	

[edit logical-systems {logical-system-name} policy-options policy-statement {policy-name} term {term-name} then]		
"aigp-originate" on page 2097	"install-nexthop" on page 8742	

[edit logical-systems {logical-system-name} policy-options policy-statement {policy-name} then]		
"aigp-originate" on page 2097	"prefix-segment" on page 13105	

[edit logical-systems {name} policy-options policy-statement{policy-name} term {term-name} then]	
"priority (policy-options)" on page 13189	

[edit logical-systems {name} policy-options policy-statement{policy-name} then]	
"priority (policy-options)" on page 13189	

[edit logical-systems {logical-system-name} policy-options prefix-list {name}]		
"apply-path" on page 2431	"dynamic-db" on page 6105	

[policy-statement]

[edit logical-systems {logical-system-name} policy-statement {policy-statement-name} term {term-name} from]		
"route-filter-list" on page 15863	"source-address-filter-list" on page 28131	

[protocols]

[edit logical-systems {logical-system-name} protocols]		
"amt (Protocols)" on page 2230	"bfd (protocols)" on page 2911	"connections (Circuits)" on page 4542
"dot1x" on page 5917	"dvmrp" on page 6081	"flex-algorithm (Protocols OSPF)" on page 7223
"igmp" on page 8381	"interconnect-multihoming-peer-gateways" on page 8806	"isis" on page 9401
"isis-instance" on page 9403	"l2circuit" on page 9531	"ldp" on page 9679
"link-management" on page 9798	"loop-detect" on page 10113	"mld" on page 10778
"mpls (Protocols)" on page 10984	"msdp" on page 10996	"mstp" on page 11006
"mvpn" on page 11135	"ospf" on page 11938	"ospf-instance" on page 11940
"ospf3" on page 11942	"pim" on page 12507	"ptp" on page 13855
"rip" on page 15803	"ripng" on page 15804	"router-advertisement" on page 15890
"router-discovery" on page 15900	"rstp" on page 16046	"rsvp" on page 16050
"sap" on page 16303	"source-packet-routing (Segment Routing Traffic Engineering)" on page 28257	"static-multihoming-peer" on page 28589
"vstp" on page 30983		

[edit logical-systems {logical-system-name} protocols (bgp | isis | ldp | ospf | ospf3 | pim | rip | ripng | rsvp) graceful-restart]	
"disable (Routing Graceful Restart)" on page 5738	

[edit logical-systems {logical-system-name} protocols (bgp | rip | ripng) graceful-restart (Enabling Globally)]	
"restart-time (BGP Graceful Restart)" on page 15663	

[edit logical-systems {logical-system-name} protocols (isis | ldp | ospf | ospf3 | rsvp) graceful-restart]	
"helper-disable (Multiple Protocols)" on page 8048	

[edit logical-systems {logical-system-name} protocols (isis | ospf | ospf3 | pim) graceful-restart]	
"restart-duration (Protocols)" on page 15651	

[edit logical-systems {logical-system-name} protocols (mstp | rstp | vstp)]		
"bpdu-block-on-edge" on page 3024	"interface (Spanning Tree)" on page 9057	"traceoptions (Spanning Tree)" on page 29869

[edit logical-systems {logical-system-name} protocols (mstp | rstp | vstp) interface {interface-name}]		
"cost" on page 4681	"edge" on page 6197	"mode" on page 10798
"no-root-port" on page 11572	"priority (Protocols STP)" on page 13199	

[edit logical-systems {logical-system-name} protocols (mstp | rstp)]		
"backup-bridge-priority" on page 2793	"bpdu-destination-mac-address (Spanning Tree)" on page 3027	"bridge-priority" on page 3036
"forward-delay" on page 7393	"hello-time" on page 8041	"max-age" on page 10418
"priority-hold-time" on page 13177	"system-id" on page 28847	"vpls-flush-on-topology-change" on page 30926

[edit logical-systems {logical-system-name} protocols (oospf | ospf3)]	
"overload (Protocols OSPF)" on page 12051	

[edit logical-systems {logical-system-name} protocols (ospf | ospf3)]		
"area" on page 2451	"backup-spf-options (Protocols OSPF)" on page 2823	"export (OSPF)" on page 6660
"graceful-restart (Protocols OSPF)" on page 7775	"import (Protocols OSPF)" on page 8463	"lsa-refresh-interval" on page 10182
"no-nssa-abr" on page 11533	"no-rfc-1583" on page 11571	"prefix-export-limit (Protocols OSPF)" on page 13085
"reference-bandwidth (Protocols OSPF)" on page 14226	"rib-group (Protocols OSPF)" on page 15774	"spf-options (Protocols OSPF)" on page 28381
"traceoptions (Protocols OSPF)" on page 29713	"traffic-engineering (OSPF)" on page 30018	

[edit logical-systems {logical-system-name} protocols (ospf | ospf3) area {area-id}]		
"area-range" on page 2453	"interface (Protocols OSPF)" on page 8984	"nssa" on page 11717

(Continued)

[edit logical-systems {logical-system-name} protocols (ospf | ospf3) area {area-id}]		
"stub" on page 28665	"virtual-link" on page 30763	

[edit logical-systems {logical-system-name} protocols (ospf | ospf3) area {area-id} interface {interface-name}]		
"bandwidth-based-metrics (Protocols OSPF)" on page 2835	"bfd-liveness-detection (Protocols OSPF)" on page 2941	"interface-type (Protocols OSPF)" on page 9091
"link-protection (Protocols OSPF)" on page 9814	"passive (Protocols OSPF)" on page 12223	

[edit logical-systems {logical-system-name} protocols (ospf | ospf3) area {area-id} interface {interface-name} passive]	
"traffic-engineering (Passive TE Mode)" on page 30021	

[edit logical-systems {logical-system-name} protocols (ospf | ospf3) area {area-id} nssa]		
"area-range" on page 2453	"default-lsa" on page 5197	

[edit logical-systems {logical-system-name} protocols (ospf | ospf3) traffic-engineering]		
"shortcuts (Protocols OSPF)" on page 16991	"tunnel-source-protocol (Protocols OSPF)" on page 30240	

[edit logical-systems {logical-system-name} protocols (ospf |ospf3) area {area-id}]	
"context-identifier (Protocols OSPF)" on page 4613	

[edit logical-systems {logical-system-name} protocols (rstp | vstp)]	
"force-version (IEEE 802.1D STP)" on page 7379	

[edit logical-systems {logical-system-name} protocols amt]		
"relay (AMT Protocol)" on page 14291	"traceoptions (Protocols AMT)" on page 29656	

[edit logical-systems {logical-system-name} protocols amt relay]		
"accounting (Protocols AMT Interface)" on page 1728	"family (Protocols AMT Relay)" on page 6898	"secret-key-timeout" on page 16423
"tunnel-devices (Protocols AMT)" on page 30196	"tunnel-limit (Protocols AMT)" on page 30212	"unicast-stream-limit (Protocols AMT)" on page 30351

[edit logical-systems {logical-system-name} protocols amt relay family]	
"inet (AMT Protocol)" on page 8555	

[edit logical-systems {logical-system-name} protocols amt relay family inet]		
"anycast-prefix" on page 2285	"local-address (Protocols AMT)" on page 9904	

[edit logical-systems {logical-system-name} protocols bfd]	
"no-issu-timer-negotiation" on page 11487	

[edit logical-systems {logical-system-name} protocols bgp]		
"accept-remote-nexthop" on page 1640	"add-path-display-ipv4-address" on page 1842	"advertise-bgp-static" on page 2009
"advertise-from-main-vpn-tables" on page 2013	"advertise-inactive" on page 2015	"advertise-peer-as" on page 2020
"advertise-prefix-sid" on page 2022	"authentication-algorithm" on page 2577	"authentication-key (Protocols BGP and BMP)" on page 2616
"authentication-key-chain (Protocols BGP and BMP)" on page 2605	"bfd-liveness-detection (BGP)" on page 2918	"bgp" on page 2956
"bgp-error-tolerance (Protocols BGP)" on page 2959	"bmp" on page 3000	"cluster" on page 4367
"damping (Protocols BGP)" on page 4822	"defaults (ACX series, JRR200, MX204, PTX5000, and VRR)" on page 5225	"description (Protocols BGP)" on page 5340
"disable (Protocols BGP)" on page 5788	"egress-te" on page 6227	"egress-te-backup-paths" on page 6232
"egress-te-set-segment" on page 6236	"enforce-first-as" on page 6343	"export (Protocols BGP)" on page 6677
"family (Protocols BGP)" on page 6900	"forwarding-context (Protocols BGP)" on page 7453	"graceful-restart (Long-Lived for BGP Helper)" on page 7763
"graceful-restart (Protocols BGP)" on page 7769	"graceful-shutdown (Protocols BGP)" on page 7784	"group (Protocols BGP)" on page 7863
"hold-time (Protocols BGP)" on page 8104	"idle-after-switch-over" on page 8300	"import (BGP)" on page 8461

(Continued)

[edit logical-systems {logical-system-name} protocols bgp]		
"include-mp-next-hop" on page 8532	"ipsec-sa (Protocols BGP)" on page 9308	"keep" on page 9452
"local-address (Protocols BGP)" on page 9905	"local-as" on page 9916	"local-preference" on page 9971
"log-updown" on page 10042	"metric-out" on page 10676	"mtu-discovery" on page 11032
"multihop" on page 11098	"multipath (Protocols BGP)" on page 11110	"multipath-build-priority" on page 11107
"no advertise-peer-as" on page 11385	"no-aggregator-id" on page 11387	"no-client-reflect" on page 11424
"out-delay" on page 11961	"outbound-route-filter" on page 11970	"output-queue-priority" on page 12008
"passive (Protocols BGP)" on page 12218	"path-selection" on page 12266	"peer-as (Protocols BGP)" on page 12346
"precision-timers" on page 13010	"preference (Protocols BGP)" on page 13038	"remove-private" on page 14368
"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763	"shutdown (Protocols BGP)" on page 27936	"tcp-aggressive-transmission" on page 28978
"tcp-mss (Protocols BGP)" on page 28998	"tcpao-auth-mismatch" on page 29034	"traceoptions (Protocols BGP)" on page 29662
"traffic-statistics-labeled-path" on page 30049	"vpn-apply-export" on page 30938	

[edit logical-systems {name} protocols bgp bfd-liveness-detection]		
"authentication (BGP BFD Liveness Detection)" on page 2588		

[edit logical-systems {logical-system-name} protocols bgp bmp]		
"route-monitoring" on page 15871		

[edit logical-systems {logical-system-name} protocols bgp family {family}]		
"damping (Protocols BGP)" on page 4822	"damping (Protocols BGP)" on page 4822	"delay-route-advertisements" on page 5262
"inet-mvpn (BGP)" on page 8574	"inet6-mvpn (BGP)" on page 8610	"inet6-vpn" on page 8620
"local-ipv4-address" on page 9955	"loops (BGP Address Family)" on page 10134	"output-queue-priority" on page 12008
"route-refresh-priority" on page 15876	"route-target (Protocols BGP)" on page 15881	"traffic-engineering (Protocols BGP)" on page 30023
"withdraw-priority" on page 31134		

[edit logical-systems {logical-system-name} protocols bgp family (inet | inet6)]		
"labeled-unicast (Protocols BGP)" on page 9603		

[edit logical-systems {logical-system-name} protocols bgp family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651	"prefix-limit" on page 13087	

[edit logical-systems {logical-system-name} protocols bgp family (inet | inet6) labeled-unicast]		
"rib (Protocols BGP)" on page 15787	"traffic-statistics (Protocols BGP)" on page 30051	

[edit logical-systems {logical-system-name} protocols bgp family (inet | inet6) unicast]	
"topology (Protocols BGP)" on page 29449	

[edit logical-systems {logical-system-name} protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow)]	
"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	

[edit logical-systems {logical-system-name} protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]	
"restarter (Graceful Restart for BGP Restarter)" on page 15667	

[edit logical-systems {logical-system-name} protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]		
"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481	

[edit logical-systems {logical-system-name} protocols bgp family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit logical-systems {logical-system-name} protocols bgp family inet]	
"labeled-unicast" on page 9601	

[edit logical-systems {logical-system-name} protocols bgp family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit logical-systems {logical-system-name} protocols bgp family inet (labeled-unicast | unicast | multicast) graceful-restart]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit logical-systems {name} protocols bgp family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit logical-systems {logical-system-name} protocols bgp family inet labeled-unicast]		
"aggregate-label" on page 2066	"aigp" on page 2092	"entropy-label (protocols bgp)" on page 6413
"explicit-null (Protocols BGP)" on page 6652	"per-group-label" on page 12416	"per-prefix-label" on page 12419
"resolve-vpn" on page 15602		

[edit logical-systems {name} protocols bgp family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229

(Continued)

[edit logical-systems {name} protocols bgp family inet unicast]	
"extended-nexthop" on page 6700	

[edit logical-systems {logical-system-name} protocols bgp family inet-vpn unicast]	
"aggregate-label" on page 2066	

[edit logical-systems {logical-system-name} protocols bgp family inet-vpn6 unicast]	
"aggregate-label" on page 2066	

[edit logical-systems {name} protocols bgp family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit logical-systems {logical-system-name} protocols bgp family inet6 labeled-unicast]		
"aggregate-label" on page 2066	"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652

[edit logical-systems {name} protocols bgp family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {logical-system-name} protocols bgp family l2vpn]	
"auto-discovery-only" on page 2711	

[edit logical-systems {logical-system-name} protocols bgp family route-target]		
"accepted-prefix-limit" on page 1651		

[edit logical-systems {logical-system-name} protocols bgp graceful-restart]		
"disable (BGP Graceful Restart)" on page 5747	"disable-notification-extensions (BGP Graceful Restart)" on page 5779	"disable-notification-flag (BGP Graceful Restart)" on page 5780
"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488	"long-lived (Graceful Restart for BGP Helper)" on page 10103	"stale-routes-time" on page 28478

[edit logical-systems {logical-system-name} protocols bgp graceful-restart long-lived]		
"advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper)" on page 2026	"receiver (Graceful Restart for BGP Helper)" on page 14127	

[edit logical-systems {logical-system-name} protocols bgp graceful-restart long-lived advertise-to-non-llgr-neighbor]		
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810		

[edit logical-systems {logical-system-name} protocols bgp group {group-name}]		
"accept-remote-nexthop" on page 1640	"add-path-display-ipv4-address" on page 1842	"advertise-bgp-static" on page 2009
"advertise-external" on page 2011	"advertise-inactive" on page 2015	"advertise-peer-as" on page 2020
"allow" on page 2160	"as-override" on page 2489	"authentication-algorithm" on page 2577

(Continued)

[edit logical-systems {logical-system-name} protocols bgp group {group-name}]		
"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605	"bfd-liveness-detection (BGP)" on page 2918
"bgp-error-tolerance (Protocols BGP)" on page 2959	"bmp" on page 3000	"cluster" on page 4367
"damping (Protocols BGP)" on page 4822	"description (Protocols BGP)" on page 5340	"egress-te" on page 6227
"egress-te-backup-paths" on page 6232	"enforce-first-as" on page 6343	"export (Protocols BGP)" on page 6677
"family (Protocols BGP)" on page 6900	"family route-target" on page 6911	"graceful-restart (Long-Lived for BGP Helper)" on page 7763
"graceful-restart (Protocols BGP)" on page 7769	"graceful-shutdown (Protocols BGP)" on page 7784	"hold-time (Protocols BGP)" on page 8104
"idle-after-switch-over" on page 8300	"import (BGP)" on page 8461	"include-mp-next-hop" on page 8532
"ipsec-sa (Protocols BGP)" on page 9308	"keep" on page 9452	"local-address (Protocols BGP)" on page 9905
"local-as" on page 9916	"local-preference" on page 9971	"log-updown" on page 10042
"metric-out" on page 10676	"mtu-discovery" on page 11032	"multihop" on page 11098
"multipath (Protocols BGP)" on page 11110	"mvpn-iana-rt-import" on page 11140	"neighbor (Protocols BGP)" on page 11273
"no advertise-peer-as" on page 11385	"no-aggregator-id" on page 11387	"no-client-reflect" on page 11424

(Continued)

[edit logical-systems {logical-system-name} protocols bgp group {group-name}]		
"out-delay" on page 11961	"outbound-route-filter" on page 11970	"passive (Protocols BGP)" on page 12218
"peer-as (Protocols BGP)" on page 12346	"peer-as-list" on page 12345	"preference (Protocols BGP)" on page 13038
"remove-private" on page 14368	"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763	"shutdown (Protocols BGP)" on page 27936
"tcp-aggressive-transmission" on page 28978	"tcp-mss (Protocols BGP)" on page 28998	"traceoptions (Protocols BGP)" on page 29662
"type (Protocols BGP)" on page 30278	"vpn-apply-export" on page 30938	

[edit logical-systems {name} protocols bgp group {name} bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} bmp]	
"route-monitoring" on page 15871	

[edit logical-systems {name} protocols bgp group {name} dynamic-neighbor {dyn-name}]		
"allow" on page 2160	"peer-limit" on page 12375	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family {family} {family-modifier}]		
"add-path" on page 1839	"damping (Protocols BGP)" on page 4822	"delay-route-advertisements" on page 5262

(Continued)

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family {family} {family-modifier}]		
"inet-mvpn (BGP)" on page 8574	"inet6-mvpn (BGP)" on page 8610	"inet6-vpn" on page 8620
"local-ipv4-address" on page 9955	"loops (BGP Address Family)" on page 10134	"output-queue-priority" on page 12008
"route-refresh-priority" on page 15876	"route-target (Protocols BGP)" on page 15881	"traffic-engineering (Protocols BGP)" on page 30023
"withdraw-priority" on page 31134		

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family (inet | inet6)]	
"labeled-unicast (Protocols BGP)" on page 9603	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651	"prefix-limit" on page 13087	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family (inet | inet6) labeled-unicast]		
"rib (Protocols BGP)" on page 15787	"traffic-statistics (Protocols BGP)" on page 30051	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family (inet | inet6) unicast]	
"topology (Protocols BGP)" on page 29449	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family (l2vpn | route-target | inet) (labeled-unicast | flow)]	
"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]	
"restarter (Graceful Restart for BGP Restarter)" on page 15667	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]		
"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family inet]	
"labeled-unicast" on page 9601	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family inet (labeled-unicast | unicast | multicast) graceful-restart]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit logical-systems {name} protocols bgp group {name} family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family inet labeled-unicast]		
"aigp" on page 2092	"entropy-label (protocols bgp)" on page 6413	"explicit-null (Protocols BGP)" on page 6652
"per-group-label" on page 12416	"per-prefix-label" on page 12419	"protection (Protocols MPLS)" on page 13417
"resolve-vpn" on page 15602		

[edit logical-systems {name} protocols bgp group {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229
"extended-nexthop" on page 6700		

[edit logical-systems {name} protocols bgp group {name} family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family inet6 labeled-unicast]		
"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652	

[edit logical-systems {name} protocols bgp group {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family l2vpn]	
"auto-discovery-only" on page 2711	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} family route-target]		
"accepted-prefix-limit" on page 1651	"proxy-generate" on page 13826	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} graceful-restart]		
"disable (BGP Graceful Restart)" on page 5747	"disable-notification-extensions (BGP Graceful Restart)" on page 5779	"disable-notification-flag (BGP Graceful Restart)" on page 5780
"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488	"long-lived (Graceful Restart for BGP Helper)" on page 10103	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} graceful-restart long-lived]		
"advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper)" on page 2026	"receiver (Graceful Restart for BGP Helper)" on page 14127	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} graceful-restart long-lived advertise-to-non-llgr-neighbor]	
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810	

[edit logical-systems {name} protocols bgp group {name} link-bandwidth]	
"auto-sense" on page 2738	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} local-as]	
"loops (Autonomous System)" on page 10132	

[edit logical-systems {name} protocols bgp group {name} metric-out]		
"effective-aigp" on page 6203	"minimum-effective-aigp" on page 10730	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address}]		
"accept-remote-nexthop" on page 1640	"add-path" on page 1839	"advertise-bgp-static" on page 2009
"advertise-external" on page 2011	"advertise-inactive" on page 2015	"advertise-peer-as" on page 2020
"as-override" on page 2489	"authentication-algorithm" on page 2577	"authentication-key (Protocols BGP and BMP)" on page 2616

(Continued)

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address}]		
"authentication-key-chain (Protocols BGP and BMP)" on page 2605	"bfd-liveness-detection (BGP)" on page 2918	"bgp-error-tolerance (Protocols BGP)" on page 2959
"bmp" on page 3000	"cluster" on page 4367	"damping (Protocols BGP)" on page 4822
"description (Protocols BGP)" on page 5340	"egress-te" on page 6227	"egress-te-adj-segment" on page 6229
"egress-te-node-segment" on page 6234	"enforce-first-as" on page 6343	"export (Protocols BGP)" on page 6677
"family (Protocols BGP)" on page 6900	"family route-target" on page 6911	"graceful-restart (Long-Lived for BGP Helper)" on page 7763
"graceful-restart (Protocols BGP)" on page 7769	"graceful-shutdown (Protocols BGP)" on page 7784	"hold-time (Protocols BGP)" on page 8104
"idle-after-switch-over" on page 8300	"import (BGP)" on page 8461	"include-mp-next-hop" on page 8532
"ipsec-sa (Protocols BGP)" on page 9308	"keep" on page 9452	"local-address (Protocols BGP)" on page 9905
"local-as" on page 9916	"local-interface (IPv6)" on page 9952	"local-preference" on page 9971
"log-updown" on page 10042	"metric-out" on page 10676	"mtu-discovery" on page 11032
"multihop" on page 11098	"multipath (Protocols BGP)" on page 11110	"no advertise-peer-as" on page 11385
"no-aggregator-id" on page 11387	"no-client-reflect" on page 11424	"out-delay" on page 11961

(Continued)

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address}]		
"outbound-route-filter" on page 11970	"passive (Protocols BGP)" on page 12218	"peer-as (Protocols BGP)" on page 12346
"preference (Protocols BGP)" on page 13038	"remove-private" on page 14368	"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763
"shutdown (Protocols BGP)" on page 27936	"tcp-aggressive-transmission" on page 28978	"tcp-mss (Protocols BGP)" on page 28998
"traceoptions (Protocols BGP)" on page 29662	"vpn-apply-export" on page 30938	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit logical-systems {name} protocols bgp group {name} neighbor {address} bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} bmp]	
"route-monitoring" on page 15871	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family {family}]		
"damping (Protocols BGP)" on page 4822	"damping (Protocols BGP)" on page 4822	"delay-route-advertisements" on page 5262

(Continued)

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family {family}]		
"local-ipv4-address" on page 9955	"loops (BGP Address Family)" on page 10134	"route-target (Protocols BGP)" on page 15881
"traffic-engineering (Protocols BGP)" on page 30023		

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family (inet | inet6)]	
"labeled-unicast (Protocols BGP)" on page 9603	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651	"prefix-limit" on page 13087	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family (inet | inet6) labeled-unicast]	
"rib (Protocols BGP)" on page 15787	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family (inet | inet6) unicast]	
"topology (Protocols BGP)" on page 29449	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family (l2vpn | route-target | inet) (labeled-unicast | flow)]	
"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]	
"restarter (Graceful Restart for BGP Restarter)" on page 15667	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]		
"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family inet (labeled-unicast | unicast | multicast) graceful-restart]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit logical-systems {name} protocols bgp group {name} neighbor {name} family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family inet labeled-unicast]		
"aigp" on page 2092	"entropy-label (protocols bgp)" on page 6413	"explicit-null (Protocols BGP)" on page 6652
"resolve-vpn" on page 15602		

[edit logical-systems {name} protocols bgp group {name} neighbor {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229
"extended-nexthop" on page 6700		

[edit logical-systems {name} protocols bgp group {name} neighbor {name} family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family inet6 labeled-unicast]		
"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652	

[edit logical-systems {name} protocols bgp group {name} neighbor {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family l2vpn]	
"auto-discovery-only" on page 2711	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} family route-target]		
"accepted-prefix-limit" on page 1651	"proxy-generate" on page 13826	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} graceful-restart]		
"disable (BGP Graceful Restart)" on page 5747	"disable-notification-extensions (BGP Graceful Restart)" on page 5779	"disable-notification-flag (BGP Graceful Restart)" on page 5780
"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488	"long-lived (Graceful Restart for BGP Helper)" on page 10103	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} graceful-restart long-lived]		
"advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper)" on page 2026	"receiver (Graceful Restart for BGP Helper)" on page 14127	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} graceful-restart long-lived advertise-to-non-llgr-neighbor]	
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} neighbor {address} outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit logical-systems {logical-system-name} protocols bgp group {group-name} outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit logical-systems {logical-system-name} protocols bgp local-as]		
"loops (Autonomous System)" on page 10132		

[edit logical-systems {name} protocols bgp multipath]		
"pause-computation-during-churn" on page 12277		

[edit logical-systems {logical-system-name} protocols bgp multipath allow-external-internal]		
"multipath (Protocols BGP)" on page 11110		

[edit logical-systems {logical-system-name} protocols bgp outbound-route-filter]		
"bgp-orf-cisco-mode" on page 2963		

[edit logical-systems {logical-system-name} protocols bgp path-selection]		
"l2vpn-use-bgp-rules-edit-protocols-bgp" on page 9565		

[edit logical-systems {name} protocols bgp source-packet-routing]		
"srgb (Protocols BGP)" on page 28407	"srv6 (Protocols BGP)" on page 28417	

[edit logical-systems {logical-system-name} protocols connections]		
"interface-switch" on page 9076	"lsp-switch" on page 10194	"p2mp-receive-switch" on page 12125
"p2mp-transmit-switch" on page 12127	"remote-interface-switch" on page 14342	

[edit logical-systems {name} protocols dot1x]		
"authenticator" on page 2673		

[edit logical-systems {name} protocols dot1x authenticator]		
"interface (802.1X)" on page 8819	"radius-reachability" on page 14015	

[edit logical-systems {name} protocols dot1x authenticator interface]		
"authentication-order (Authenticator)" on page 2638	"multi-domain" on page 11048	

[edit logical-systems {name} protocols dot1x authenticator interface (all | [{interface-names}])]	
"server-reject-bridge-domain | server-reject-vlan" on page 16661	

[edit logical-systems {name} protocols dot1x authenticator interface (all | interface-names)]	
"eapol-block" on page 6183	

[edit logical-systems {name} protocols dot1x authenticator interface (all | interface-names) (server-reject-bridge-domain | server-reject-vlan)]	
"eapol-block" on page 6183	

[edit logical-systems protocols dot1x authenticator interface {interface-name} mac-radius]	
"authentication-protocol" on page 2645	

[edit logical-systems {logical-system-name} protocols dvmrp]		
"disable (Protocols)" on page 5731	"export (Protocols DVMRP)" on page 6678	"import (Protocols DVMRP)" on page 8474
"interface (Protocols DVMRP)" on page 8960	"rib-group (Protocols DVMRP)" on page 15769	"traceoptions (Protocols DVMRP)" on page 29669

[edit logical-systems {logical-system-name} protocols dvmrp interface {interface-name}]		
"disable (Protocols)" on page 5731	"hold-time (Protocols DVMRP)" on page 8106	"metric (Protocols DVMRP)" on page 10682
"mode (Protocols DVMRP)" on page 10813		

[edit logical-systems { logical-systems-name} protocols evpn]		
"duplicate-mac-detection" on page 6068	"node-detection (EVPN-VXLAN)" on page 11633	

[edit logical-systems {name} protocols evpn node-detection]	
"bfd-liveness-detection (EVPN Node Detection)" on page 2923	

[edit logical-systems {name} protocols group]	
"detection (Network Isolation Groups)" on page 5540	

[edit logical-systems {logical-system-name} protocols iccp]	
"iccp (Traceoptions)" on page 8247	

[edit logical-systems {logical-system-name} protocols igmp]		
"accounting (Protocols IGMP)" on page 1730	"amt (IGMP)" on page 2228	"interface (Protocols IGMP)" on page 8962
"maximum-transmit-rate (Protocols IGMP)" on page 10572	"query-interval (Protocols IGMP)" on page 13929	"query-last-member-interval (Protocols IGMP)" on page 13935
"query-response-interval (Protocols IGMP)" on page 13940	"robust-count (Protocols IGMP)" on page 15818	"traceoptions (Protocols IGMP)" on page 29675

[edit logical-systems {logical-system-name} protocols igmp amt relay defaults]		
"accounting (Protocols IGMP AMT Interface)" on page 1731	"group-policy (Protocols IGMP AMT Interface)" on page 7856	"query-interval (Protocols IGMP AMT)" on page 13930
"query-response-interval (Protocols IGMP AMT)" on page 13941	"robust-count (Protocols IGMP AMT)" on page 15820	"ssm-map (Protocols IGMP AMT)" on page 28465
"version (Protocols IGMP AMT)" on page 30716		

[edit logical-systems {logical-system-name} protocols igmp interface {interface-name}]		
"accounting (Protocols IGMP Interface)" on page 1732	"disable (Protocols)" on page 5731	"group-limit (IGMP) " on page 7844
"group-policy (Protocols IGMP)" on page 7855	"group-threshold (Protocols IGMP Interface)" on page 7903	"immediate-leave" on page 8454
"log-interval (IGMP Interface)" on page 10009	"oif-map (IGMP Interface) " on page 11803	"passive (IGMP)" on page 12202
"promiscuous-mode (Protocols IGMP)" on page 13375	"ssm-map (Protocols IGMP)" on page 28464	"ssm-map-policy (IGMP)" on page 28461

(Continued)

[edit logical-systems {logical-system-name} protocols igmp interface {interface-name}]		
"static (Protocols IGMP)" on page 28557	"version (Protocols IGMP)" on page 30715	

[edit logical-systems {logical-system-name} protocols igmp interface {interface-name} static]	
"group (Protocols IGMP)" on page 7869	

[edit logical-systems {logical-system-name} protocols igmp interface {interface-name} static group {multicast-group-address}]		
"exclude (Protocols IGMP)" on page 6614	"group-count (Protocols IGMP)" on page 7810	"group-increment (Protocols IGMP)" on page 7838
"source (Protocols IGMP)" on page 28305		

[edit logical-systems {logical-system-name} protocols igmp interface {interface-name} static group {multicast-group-address} source]		
"source-count (Protocols IGMP)" on page 28197	"source-increment (Protocols IGMP)" on page 28221	

[edit logical-systems {logical-system-name} protocols igmp-snooping]	
"version (IGMP Snooping)" on page 30704	

[edit logical-systems {name} protocols igmp-snooping vlan {vlan-name}]	
"data-forwarding" on page 4835	

[edit logical-systems {name} protocols igmp-snooping vlan {vlan-name} data-forwarding]	
"receiver" on page 14124	

[edit logical-systems {logical-system-name} protocols isis]		
"backup-spf-options (Protocols IS-IS)" on page 2821	"clns-updown-compatibility" on page 4350	"context-identifier (Protocols IS-IS)" on page 4611
"export (ISIS)" on page 6658	"graceful-restart (Protocols IS-IS)" on page 7771	"ignore-attached-bit" on page 8396
"import (Protocols IS-IS)" on page 8476	"interface (Protocols IS-IS)" on page 8966	"label-switched-path (Protocols IS-IS)" on page 9582
"layer2-map" on page 9651	"level (Global IS-IS)" on page 9724	"loose-authentication-check" on page 10137
"lsp-lifetime" on page 10189	"max-areas" on page 10420	"max-lsp-size" on page 10446
"max-snp-size" on page 10483	"no-ipv4-routing" on page 11483	"no-ipv6-routing" on page 11484
"overload (Protocols IS-IS)" on page 12048	"reference-bandwidth (Protocols IS-IS)" on page 14224	"rib-group (Protocols IS-IS)" on page 15770
"source-packet-routing (Protocols IS-IS)" on page 28266	"spf-options (Protocols IS-IS)" on page 28379	"strict-dual-isis (Protocols IS-IS)" on page 28648
"topologies (Protocols IS-IS)" on page 29438	"traceoptions (Protocols IS-IS)" on page 29684	"traffic-engineering (Protocols IS-IS)" on page 30025

[edit logical-systems {name} protocols isis backup-spf-options]	
"use-post-convergence-lfa" on page 30531	

[edit logical-systems {name} protocols isis dynamic-overload]		
"auto-recovery" on page 2733	"no-overload-on-prefix-export-limit" on page 11542	

[edit logical-systems {logical-system-name} protocols isis interface {interface-name}]		
"bfd-liveness-detection (Protocols IS-IS)" on page 2933	"layer2-map" on page 9651	"level (IS-IS Interfaces)" on page 9728
"link-protection (Protocols IS-IS)" on page 9811	"no-adjacency-down-notification (Protocols IS-IS)" on page 11380	"no-advertise-adjacency-segment (Protocols IS-IS)" on page 11382
"passive (Protocols IS-IS)" on page 12219		

[edit logical-systems {name} protocols isis interface {name} delay-measurement]	
"advertisement (Delay Measurement)" on page 2031	

[edit logical-systems {logical-system-name} protocols isis interface {interface-name} level {level-number}]		
"flood-reflector (Protocols IS-IS)" on page 7242	"hold-time (Protocols IS-IS)" on page 8108	"metric (Protocols IS-IS)" on page 10684
"passive (Protocols IS-IS)" on page 12219	"post-convergence-lfa" on page 12903	

[edit logical-systems {name} protocols isis interface {name} level {number} post-convergence-lfa]	
"node-protection" on page 11651	

[edit logical-systems {logical-system-name} protocols isis interface-group {group-name}]		
"link-group-protection (Protocols IS-IS)" on page 9794		

[edit logical-systems {name} protocols isis interface-group {name} level {name} topology]		
"bandwidth-based-metrics" on page 2837		

[edit logical-systems {logical-system-name} protocols isis level]		
"max-lsp-size" on page 10446	"prefix-export-limit (Protocols IS-IS)" on page 13083	"source-packet-routing (Protocols IS-IS)" on page 28266

[edit logical-systems {name} protocols isis no-external-export protocol]		
"no-external-export" on page 11454		

[edit logical-systems {name} protocols isis source-packet-routing]		
"srgb (Protocols IS-IS)" on page 28408	"strict-asla-based-flex-algorithm" on page 28645	"traffic-statistics (Protocols IS-IS)" on page 30048

[edit logical-systems {name} protocols isis spf-options]		
"microloop-avoidance" on page 10700		

[edit logical-systems {name} protocols isis traffic-engineering]		
"advertisement" on page 2030	"family (Protocols IS-IS)" on page 6905	

[edit logical-systems {name} protocols isissource-packet-routing]		
"flex-algorithm" on page 7221		

[edit logical-systems {name} protocols l2-learning]		
"global-mac-ip-limit" on page 7709	"global-mac-ip-table-aging-time" on page 7711	"global-no-control-mac-aging" on page 7722
"mclag-arp-nd-sync" on page 10585	"mclag-arpreq-sync" on page 10586	

[edit logical-systems {logical-system-name} protocols l2circuit]		
"local-switching (Layer 2 Circuits)" on page 9977	"neighbor (Protocols Layer 2 Circuit)" on page 11278	"traceoptions (Protocols Layer 2 Circuit)" on page 29688

[edit logical-systems {logical-system-name} protocols l2circuit local-switching]		
"interface (Protocols Layer 2 Circuit)" on page 8974		

[edit logical-systems {logical-system-name} protocols l2circuit local-switching interface {interface-name}]		
"backup-neighbor" on page 2802	"connection-protection" on page 4540	"encapsulation-type (Layer 2 Circuits)" on page 6308
"end-interface" on page 6338	"ignore-encapsulation-mismatch" on page 8397	"ignore-mtu-mismatch" on page 8408
"protect-interface" on page 13409		

[edit logical-systems {logical-system-name} protocols l2circuit local-switching interface {interface-name} backup-neighbor {address}]		
"community (Protocols Layer 2 Circuit)" on page 4469	"mtu (interfaces)" on page 11027	"virtual-circuit-id" on page 30749

[edit logical-systems {logical-system-name} protocols l2circuit local-switching interface {interface-name} end-interface]	
"protect-interface" on page 13409	

[edit logical-systems {logical-system-name} protocols l2circuit local-switching interface {interface-name} end-interface interface {interface-name}]	
"backup-interface (Layer 2 Circuits)" on page 2797	

[edit logical-systems {logical-system-name} protocols l2circuit local-switching interface {interface-name} end-interface interface {interface-name} backup-neighbor {address}]		
"hot-standby (Protocols Layer 2 Circuit)" on page 8189	"standby (Protocols Layer 2 Circuit)" on page 28489	

[edit logical-systems {logical-system-name} protocols l2circuit neighbor {address}]	
"interface (Protocols Layer 2 Circuit)" on page 8974	

[edit logical-systems {logical-system-name} protocols l2circuit neighbor {address} interface {interface-name}]		
"backup-neighbor" on page 2802	"bandwidth (Protocols Layer 2 Circuit)" on page 2870	"community (Protocols Layer 2 Circuit)" on page 4469
"connection-protection" on page 4540	"control-word (Protocols Layer 2 Circuit Neighbor)" on page 4647	"description (Protocols Layer 2 Circuit Neighbor)" on page 5342

(Continued)

[edit logical-systems {logical-system-name} protocols l2circuit neighbor {address} interface {interface-name}]		
"egress-protection (Layer 2 circuit)" on page 6218	"encapsulation-type (Layer 2 Circuits)" on page 6308	"encapsulation-type (Layer 2 VPNs)" on page 6310
"ignore-encapsulation-mismatch" on page 8397	"ignore-mtu-mismatch" on page 8408	"mtu (interfaces)" on page 11027
"protect-interface" on page 13409	"pseudowire-status-tlv" on page 13849	"psn-tunnel-endpoint" on page 13851
"revert-time (Protocols Layer 2 Circuits)" on page 15729	"switchover-delay" on page 28759	"virtual-circuit-id" on page 30749

[edit logical-systems {logical-system-name} protocols l2circuit neighbor {address} interface {interface-name} backup-neighbor {address}]		
"community (Protocols Layer 2 Circuit)" on page 4469	"hot-standby (Protocols Layer 2 Circuit)" on page 8189	"mtu (interfaces)" on page 11027
"psn-tunnel-endpoint" on page 13851	"standby (Protocols Layer 2 Circuit)" on page 28489	"virtual-circuit-id" on page 30749

[edit logical-systems {logical-system-name} protocols l2circuit neighbor {address} interface {interface-name} egress-protection]		
"protected-l2circuit" on page 13412	"protector-interface" on page 13420	"protector-pe" on page 13421

[edit logical-systems {logical-system-name} protocols l2circuit neighbor {address} interface {interface-name} oam]	
"ping-interval" on page 12545	

[edit logical-systems {logical-system-name} protocols l2circuit neighbor {address} interface {interface-name} pseudowire-status-tlv]	
"hot-standby-vc-on (Protocols Layer 2 Circuit)" on page 8190	

[edit logical-systems {logical-system-name} protocols l2circuit neighbor {address} interfaces {interface-name}]	
"no-revert (Neighbor Interface)" on page 11569	

[edit logical-systems {name} protocols layer2-control mac-rewrite]	
"interface (Layer 2 Protocol Tunneling)" on page 8904	

[edit logical-systems {name} protocols layer2-control mac-rewrite interface {interface-name}]	
"protocol" on page 13424	

[edit logical-systems {logical-system-name} protocols ldp]		
"auto-targeted-session" on page 2748	"deaggregate" on page 5153	"dod-request-policy" on page 5873
"egress-policy" on page 6215	"entropy-label (protocols mpls)" on page 6416	"explicit-null (Protocols BGP)" on page 6652
"explicit-null (Protocols LDP)" on page 6654	"export (Protocols LDP)" on page 6680	"graceful-restart (Protocols LDP)" on page 7773
"igp-synchronization" on page 8414	"import (Protocols LDP)" on page 8477	"interface (Protocols LDP)" on page 8978
"keepalive-interval" on page 9468	"keepalive-timeout" on page 9474	"l2-smart-policy" on page 9528

(Continued)

[edit logical-systems {logical-system-name} protocols ldp]		
"label-withdrawal-delay" on page 9598	"log-updown (Protocols LDP)" on page 10043	"no-forwarding" on page 11464
"oam (Protocols LDP)" on page 11759	"p2mp (Protocols LDP)" on page 12122	"policing (Protocols LDP)" on page 12662
"preference (Protocols LDP)" on page 13040	"session (Protocols LDP)" on page 16845	"session-group" on page 16810
"session-protection" on page 16843	"strict-targeted-hellos" on page 28654	"targeted-hello" on page 28962
"traceoptions (Protocols LDP)" on page 29693	"track-igp-metric" on page 29989	"traffic-statistics (Protocols LDP)" on page 30053
"transport-address" on page 30118		

[edit logical-systems {logical-system-name} protocols ldp entropy-label]	
"ingress-policy" on page 8641	

[edit logical-systems {logical-system-name} protocols ldp graceful-restart]		
"disable (Protocols LDP)" on page 5789	"helper-disable (LDP)" on page 8046	"maximum-neighbor-reconnect-time" on page 10537
"maximum-neighbor-recovery-time" on page 10538	"reconnect-time" on page 14135	"recovery-time (LDP)" on page 14148

[edit logical-systems {logical-system-name} protocols ldp interface {interface-name}]		
"allow-subnet-mismatch" on page 2207	"disable (Protocols LDP)" on page 5789	"hello-interval (Protocols LDP)" on page 8036

(Continued)

[edit logical-systems {logical-system-name} protocols ldp interface {interface-name}]		
"hold-time (Protocols LDP)" on page 8110	"transport-address" on page 30118	

[edit logical-systems {logical-system-name} protocols ldp oam]		
"bfd-liveness-detection (Protocols LDP)" on page 2937	"fec (protocols ldp)" on page 6990	"ingress-policy" on page 8641
"periodic-traceroute" on page 12441		

[edit logical-systems {logical-system-name} protocols ldp oam bfd-liveness-detection]		
"ecmp" on page 6190	"version (BFD)" on page 30695	

[edit logical-systems {logical-system-name} protocols ldp oam bfd-livenesss-detection]		
"failure-action (Protocols LDP)" on page 6789	"holddown-interval (protocols ldp)" on page 8124	

[edit logical-systems {logical-system-name} protocols ldp oam fec {address}]		
"bfd-liveness-detection (Protocols LDP)" on page 2937	"periodic-traceroute" on page 12441	

[edit logical-systems {logical-system-name} protocols ldp oam fec {address} bfd-liveness-detection]		
"ecmp" on page 6190	"version (BFD)" on page 30695	

[edit logical-systems {logical-system-name} protocols ldp oam fec {address} bfd-livenesss-detection]		
"failure-action (Protocols LDP)" on page 6789	"holddown-interval (protocols ldp)" on page 8124	

[edit logical-systems {logical-system-name} protocols ldp session {session-address}]		
"authentication-algorithm" on page 2577	"authentication-key (Protocols LDP)" on page 2618	"authentication-key-chain (Protocols LDP)" on page 2607
"downstream-on-demand" on page 5934		

[edit logical-systems {logical-system-name} protocols ldp targeted-hello]		
"hello-interval (Protocols LDP)" on page 8036	"hold-time (Protocols LDP)" on page 8110	

[edit logical-systems {logical-system-name} protocols link-management]		
"peer (Protocols LMP)" on page 12379	"traceoptions (Protocols Link Management)" on page 29697	

[edit logical-systems {logical-system-name} protocols link-management peer {peer-name}]		
"address (Peer)" on page 1895	"control-channel (Protocols Link Management Peer)" on page 4621	"lmp-control-channel" on page 9864
"lmp-protocol" on page 9865		

[edit logical-systems {logical-system-name} protocols link-management peer {peer-name} lmp-control-channel {control-channel-interface}]	
"remote-address (for LMP Control Channel)" on page 14312	

[edit logical-systems {logical-system-name} protocols link-management peer {peer-name} lmp-protocol]		
"hello-dead-interval" on page 8030	"hello-interval (LMP)" on page 8035	"passive (Protocols Link Management)" on page 12222
"retransmission-interval" on page 15683	"retry-limit (Protocols Link Management)" on page 15710	

[edit logical-systems {logical-system-name} protocols link-management te-link {te-link-name}]		
"disable (GMPLS)" on page 5761	"interface (Protocols Link Management)" on page 8980	"label-switched-path (Protocols Link Management)" on page 9583
"local-address (Protocols Link Management)" on page 9908	"remote-address (for LMP Traffic Engineering)" on page 14314	"remote-id (protocols link-management)" on page 14330

[edit logical-systems {logical-system-name} protocols link-management te-link {te-link-name} interface {interface-name}]		
"local-address (Protocols Link Management)" on page 9908	"remote-address (for LMP Traffic Engineering)" on page 14314	"remote-id (protocols link-management)" on page 14330

[edit logical-systems {name} protocols lldp]	
"chassis-id (LLDP)" on page 3301	

[edit logical-systems {logical-system-name} protocols mld]		
"accounting (Protocols MLD)" on page 1733	"interface (Protocols MLD)" on page 8981	"maximum-transmit-rate (Protocols MLD)" on page 10574
"query-interval (Protocols MLD)" on page 13932	"query-last-member-interval (Protocols MLD)" on page 13936	"query-response-interval (Protocols MLD)" on page 13943
"robust-count (Protocols MLD)" on page 15821		

[edit logical-systems {logical-system-name} protocols mld interface {interface-name}]		
"accounting (Protocols MLD Interface)" on page 1735	"disable (Protocols MLD)" on page 5791	"disable (Protocols)" on page 5731
"group-limit (Protocols MLD)" on page 7847	"group-policy (Protocols MLD)" on page 7858	"group-threshold (Protocols MLD Interface)" on page 7905
"immediate-leave" on page 8454	"log-interval (MLD Interface)" on page 10011	"oif-map (MLD Interface) " on page 11804
"passive (MLD)" on page 12204	"ssm-map (Protocols MLD)" on page 28466	"ssm-map-policy (MLD)" on page 28462
"static (Protocols MLD)" on page 28560	"version (Protocols MLD)" on page 30718	

[edit logical-systems {logical-system-name} protocols mld interface {interface-name} static]	
"group (Protocols MLD)" on page 7870	

[edit logical-systems {logical-system-name} protocols mld interface {interface-name} static group {multicast-group-address}]		
"exclude (Protocols MLD)" on page 6615	"group-count (Protocols MLD)" on page 7812	"group-increment (Protocols MLD)" on page 7840
"source (Protocols MLD)" on page 28307		

[edit logical-systems {logical-system-name} protocols mld interface {interface-name} static group {multicast-group-address} source]		
"source-count (Protocols MLD)" on page 28199	"source-increment (Protocols MLD)" on page 28223	

[edit logical-systems {logical-system-name} protocols mld-snooping]	
"version (MLD Snooping)" on page 30712	

[edit logical-systems {logical-systems-name} protocols mpls]		
"abstract-hop" on page 1625	"admin-group (for LSPs)" on page 1972	"admin-group-extended" on page 1968
"admin-groups" on page 1973	"advertisement-hold-time" on page 2033	"auto-policing" on page 2724
"bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)" on page 2843	"class-of-service (Protocols MPLS)" on page 3369	"constituent-list" on page 4571
"cspf-backoff-time" on page 4731	"diffserv-te" on page 5705	"disable (Protocols MPLS)" on page 5793
"egress-protection (MPLS)" on page 6220	"expand-loose-hop" on page 6647	"explicit-null (Protocols BGP)" on page 6652
"explicit-null (Protocols MPLS)" on page 6655	"hop-limit" on page 8134	"interface (Protocols MPLS)" on page 8983
"ipv6-tunneling" on page 9377	"label-switched-path (Protocols MPLS)" on page 9585	"log-updown (Protocols MPLS)" on page 10045
"no-cspf" on page 11429	"no-decrement-ttl" on page 11432	"no-propagate-ttl" on page 11553
"optimize-aggressive" on page 11855	"optimize-hold-dead-delay" on page 11856	"optimize-switchover-delay" on page 11857
"optimize-timer (Protocols MPLS)" on page 11859	"path (Protocols MPLS)" on page 12259	"path-mtu" on page 12256

(Continued)

[edit logical-systems {logical-systems-name} protocols mpls]		
"preference (Protocols MPLS)" on page 13042	"priority (Protocols MPLS)" on page 13193	"record" on page 14136
"revert-timer" on page 15732	"rsvp-error-hold-time" on page 16056	"smart-optimize-timer" on page 28048
"standby" on page 28486	"standby" on page 28485	"static-label-switched-path" on page 28545
"statistics (Protocols MPLS)" on page 28594	"sync-active-path-bandwidth" on page 28783	"traceoptions (Protocols MPLS)" on page 29699
"traffic-engineering (Protocols MPLS)" on page 30031	"ultimate-hop-popping" on page 30323	

[edit logical-systems {logical-system-name} protocols mpls diffserv-te]		
"bandwidth-model" on page 2861	"te-class-matrix" on page 29041	

[edit logical-systems {logical-system-name} protocols mpls egress-protection context-identifier {context-id}]	
"advertise-mode (MPLS)" on page 2018	

[edit logical-systems {logical-system-name} protocols mpls interface {interface-name}]		
"admin-group (for Interfaces)" on page 1971	"admin-groups-extended" on page 1975	"always-mark-connection-protection-tlv" on page 2219
"disable (Protocols MPLS)" on page 5793	"switch-away-lsps" on page 28746	

[edit logical-systems {logical-system-name} protocols mpls interface {interface-name} static]	
"protection-revert-time" on page 13419	

[edit logical-systems {logical-system-name} protocols mpls label-switch-path {lsp-name}]	
"optimize-hold-dead-delay" on page 11856	

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name}]		
"adaptive (protocols mpls)" on page 1825	"admin-down" on page 1962	"admin-group (for LSPs)" on page 1972
"admin-group-extended" on page 1968	"associate-backup-pe-groups" on page 2511	"bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)" on page 2843
"class-of-service (Protocols MPLS)" on page 3369	"conditional-metric" on page 4506	"corouted-bidirectional" on page 4673
"corouted-bidirectional-passive" on page 4674	"cspf-backoff-time" on page 4731	"description (Protocols MPLS)" on page 5345
"disable (Protocols MPLS)" on page 5793	"egress-protection (MPLS)" on page 6220	"exclude-srlg" on page 6629
"fast-reroute (Protocols MPLS)" on page 6943	"from (Protocols MPLS)" on page 7595	"hop-limit" on page 8134
"in-place-lsp-bandwidth-update (Protocols MPLS)" on page 8493	"install (Protocols MPLS)" on page 8744	"inter-domain" on page 8788
"ldp-tunneling" on page 9687	"link-protection (Dynamic LSPs)" on page 9807	"lsp-attributes" on page 10183

(Continued)

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name}]		
"metric (Protocols MPLS)" on page 10686	"no-decrement-ttl" on page 11432	"no-install-to-address" on page 11476
"node-link-protection (Protocols MPLS)" on page 11649	"optimize-timer (Protocols MPLS)" on page 11859	"p2mp (Protocols MPLS)" on page 12124
"policing (Protocols MPLS)" on page 12664	"pop-and-forward (Protocols MPLS)" on page 12802	"preference (Protocols MPLS)" on page 13042
"primary (Protocols MPLS)" on page 13144	"priority (Protocols MPLS)" on page 13193	"random" on page 14043
"record" on page 14136	"retry-limit" on page 15709	"retry-timer" on page 15717
"revert-timer" on page 15732	"secondary (Protocols MPLS)" on page 16416	"soft-preemption (Protocols MPLS)" on page 28088
"standby" on page 28486	"standby" on page 28485	"template" on page 29071
"to" on page 29419	"traceoptions (Protocols MPLS)" on page 29699	"ultimate-hop-popping" on page 30323

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name} (primary | secondary) {path-name}]		
"adaptive (protocols mpls)" on page 1825	"admin-group (for LSPs)" on page 1972	"admin-group-extended" on page 1968
"bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)" on page 2843	"class-of-service (Protocols MPLS)" on page 3369	"hop-limit" on page 8134
"no-cspf" on page 11429	"no-decrement-ttl" on page 11432	"optimize-timer (Protocols MPLS)" on page 11859

(Continued)

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name} (primary | secondary) {path-name}]		
"preference (Protocols MPLS)" on page 13042	"priority (Protocols MPLS)" on page 13193	"record" on page 14136
"select" on page 16525	"standby" on page 28486	"standby" on page 28485
"sync-active-path-bandwidth" on page 28783		

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name} (primary | secondary) {path-name} admin-group]		
"exclude (for Administrative Groups)" on page 6605	"include-all (for Administrative Groups)" on page 8518	"include-any (for Administrative Groups)" on page 8521

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name} admin-group]		
"exclude (for Administrative Groups)" on page 6605	"include-all (for Administrative Groups)" on page 8518	"include-any (for Administrative Groups)" on page 8521

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name} auto-bandwidth]		
"adjust-interval" on page 1946	"adjust-threshold" on page 1953	"adjust-threshold-absolute" on page 1954
"adjust-threshold-activate-bandwidth" on page 1956	"adjust-threshold-overflow-limit" on page 1958	"adjust-threshold-underflow-limit" on page 1959
"disable (Protocols MPLS)" on page 5793	"maximum-bandwidth (Protocols MPLS)" on page 10504	"minimum-bandwidth" on page 10722
"minimum-bandwidth-adjust-interval" on page 10724	"minimum-bandwidth-adjust-threshold-change" on page 10725	"minimum-bandwidth-adjust-threshold-value" on page 10726

(Continued)

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name} auto-bandwidth]		
"monitor-bandwidth" on page 10834	"resignal-minimum-bandwidth" on page 15585	"sync-active-path-bandwidth" on page 28783

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name} egress-protection context-identifier {context-id}]	
"advertise-mode (MPLS)" on page 2018	

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name} fast-reroute]		
"bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)" on page 2843	"bandwidth-percent (protocols mpls)" on page 2864	"exclude (for Fast Reroute)" on page 6606
"hop-limit" on page 8134	"include-all (for Fast Reroute)" on page 8519	"include-any (for Fast Reroute)" on page 8522

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name} lsp-attributes]		
"encoding-type" on page 6318	"gpid" on page 7748	"signal-bandwidth" on page 27941
"switching-type" on page 28758		

[edit logical-systems {logical-system-name} protocols mpls label-switched-path {lsp-name} oam bfd-liveness-detection]	
"failure-action (Protocols MPLS)" on page 6790	

[edit logical-systems {name} protocols mpls label-switched-path {name} secondary {name} non-active-hold-priority]	
"non-active-hold-priority" on page 11657	

[edit logical-systems {logical-system-name} protocols mpls log-updown]	
"no-trap" on page 11604	

[edit logical-systems {logical-system-name} protocols mpls oam bfd-liveness-detection]	
"failure-action (Protocols MPLS)" on page 6790	

[edit logical-systems {logical-system-name} protocols mpls path-mtu]	
"allow-fragmentation (protocols mpls)" on page 2178	

[edit logical-systems {logical-system-name} protocols mpls path-mtu rsvp]	
"mtu-signaling" on page 11037	

[edit logical-systems {logical-system-name} protocols mpls static-label-switched-path {lsp-name}]		
"bypass (Static LSP)" on page 3117	"ingress (LSP)" on page 8632	"segment" on page 16519
"transit" on page 30064		

[edit logical-systems {logical-system-name} protocols mpls static-label-switched-path {lsp-name} bypass]		
"bandwidth (Static LSP)" on page 2876	"description (Protocols MPLS)" on page 5345	"next-hop (Protocols MPLS)" on page 11349

(Continued)

[edit logical-systems {logical-system-name} protocols mpls static-label-switched-path {lsp-name} bypass]		
"next-table" on page 11359	"push (protocols mpls)" on page 13873	"to" on page 29419

[edit logical-systems {logical-system-name} protocols mpls static-label-switched-path {lsp-name} ingress]		
"bandwidth (Static LSP)" on page 2876	"class-of-service (Protocols MPLS)" on page 3369	"class-of-service (Protocols MPLS)" on page 3369
"description (Protocols MPLS)" on page 5345	"install (Protocols MPLS)" on page 8744	"link-protection (Static LSPs)" on page 9818
"metric (Protocols MPLS)" on page 10686	"next-hop (Protocols MPLS)" on page 11349	"no-install-to-address" on page 11476
"node-protection (Static LSP)" on page 11652	"policing (Protocols MPLS)" on page 12664	"preference (Protocols MPLS)" on page 13042
"push (protocols mpls)" on page 13873	"to" on page 29419	

[edit logical-systems {logical-system-name} protocols mpls static-label-switched-path {lsp-name} transit {incoming-label}]		
"bandwidth (Static LSP)" on page 2876	"description (Protocols MPLS)" on page 5345	"link-protection (Static LSPs)" on page 9818
"next-hop (Protocols MPLS)" on page 11349	"node-protection (Static LSP)" on page 11652	"pop (protocols mpls)" on page 12797
"swap (protocols mpls)" on page 28738		

[edit logical-systems {logical-system-name} protocols mpls statistics]		
"auto-bandwidth (MPLS Statistics)" on page 2684	"no-transit-statistics" on page 11600	

[edit logical-systems {logical-system-name} protocols mpls traffic-engineering]	
"database" on page 4850	

[edit logical-systems {logical-system-name} protocols mpls traffic-engineering database]		
"export (MPLS Traffic engineering database)" on page 6668	"import (MPLS Traffic Engineering Database)" on page 8471	

[edit logical-systems {logical-system-name} protocols mpls traffic-engineering database export]	
"credibility" on page 4710	

[edit logical-systems {logical-system-name} protocols mplslabel-switched-path {lsp-name}]		
"no-cspf" on page 11429	"sync-active-path-bandwidth" on page 28783	

[edit logical-systems {logical-system-name} protocols msdp]		
"active-source-limit" on page 1819	"data-encapsulation" on page 4828	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681
"group (Protocols MSDP)" on page 7873	"hold-time (Protocols MSDP)" on page 8112	"import (Protocols MSDP)" on page 8478
"keep-alive (Protocols MSDP)" on page 9457	"local-address (Protocols MSDP) " on page 9910	"peer (Protocols MSDP)" on page 12380

(Continued)

[edit logical-systems {logical-system-name} protocols msdp]		
"rib-group (Protocols MSDP)" on page 15772	"sa-hold-time (Protocols MSDP)" on page 16262	"source (Protocols MSDP) " on page 28308
"traceoptions (Protocols MSDP)" on page 29703		

[edit logical-systems {logical-system-name} protocols msdp active-source-limit]		
"log-interval (Protocols MSDP)" on page 10014	"log-warning (Protocols MSDP)" on page 10049	"maximum (MSDP Active Source Messages) " on page 10536
"threshold (MSDP Active Source Messages) " on page 29327		

[edit logical-systems {logical-system-name} protocols msdp group {group-name}]		
"default-peer" on page 5200	"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731
"export (Protocols MSDP)" on page 6681	"import (Protocols MSDP)" on page 8478	"local-address (Protocols MSDP) " on page 9910
"mode (Protocols MSDP)" on page 10814	"peer (Protocols MSDP)" on page 12380	"traceoptions (Protocols MSDP)" on page 29703

[edit logical-systems {logical-system-name} protocols msdp group {group-name} peer {address}]		
"active-source-limit" on page 1819	"authentication-key (Protocols MSDP)" on page 2600	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681

(Continued)

[edit logical-systems {logical-system-name} protocols msdp group {group-name} peer {address}]		
"import (Protocols MSDP)" on page 8478	"local-address (Protocols MSDP) " on page 9910	"traceoptions (Protocols MSDP)" on page 29703

[edit logical-systems {logical-system-name} protocols msdp group {group-name} peer address]		
"hold-time (Protocols MSDP)" on page 8112	"keep-alive (Protocols MSDP)" on page 9457	"sa-hold-time (Protocols MSDP)" on page 16262

[edit logical-systems {logical-system-name} protocols msdp peer {address}]		
"active-source-limit" on page 1819	"authentication-key (Protocols MSDP)" on page 2600	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681
"hold-time (Protocols MSDP)" on page 8112	"import (Protocols MSDP)" on page 8478	"keep-alive (Protocols MSDP)" on page 9457
"local-address (Protocols MSDP) " on page 9910	"sa-hold-time (Protocols MSDP)" on page 16262	"traceoptions (Protocols MSDP)" on page 29703

[edit logical-systems {logical-system-name} protocols msdp source {ip-address}/{prefix-length}]	
"active-source-limit" on page 1819	

[edit logical-systems {logical-system-name} protocols mstp]		
"configuration-name" on page 4516	"disable (STP)" on page 5739	"max-hops" on page 10441
"msti" on page 11003	"revision-level" on page 15735	

[edit logical-systems {logical-system-name} protocols mstp msti {msti-id}]		
"bridge-priority" on page 3036	"vlan (MSTP)" on page 30835	

[edit logical-systems {logical-system-name} protocols mstp msti {msti-id} interface {interface-name}]		
"cost" on page 4681	"edge" on page 6197	"priority (Protocols STP)" on page 13199

[edit logical-systems {logical-system-name} protocols mvpn]	
"traceoptions (Protocols MVPN)" on page 29706	

[edit logical-systems {logical-system--name} protocols mvpn mvpn-mode spt-only]	
"source-active-advertisement" on page 28114	

[edit logical-systems {logical-system--name} protocols mvpn mvpn-mode spt-only source-active-advertisement]	
"dampen (Protocols MVPN)" on page 4813	

[edit logical-systems {logical-system-name} protocols mvrp]		
"join-timer (MVRP)" on page 9434	"no-dynamic-vlan" on page 11446	"traceoptions (MVRP)" on page 29607

[edit logical-systems {logical-system-name} protocols mvrp interface (all | {interface-name})]	
"point-to-point (MVRP) " on page 12620	

[edit logical-systems {name} protocols oam ethernet connectivity-fault-management]	
"enhanced-cfm-mode" on page 6355	

[edit logical-systems {name} protocols oam ethernet connectivity-fault-management maintenance-domain {name} maintenance-association {name} continuity-check]	
"interface-status-send-rdi" on page 9071	

[edit logical-systems {name} protocols oam ethernet connectivity-fault-management maintenance-domain {name} maintenance-association {name} mep {mep-id}]	
"auto-discovery" on page 2709	

[edit logical-systems {name} protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]	
"flap-trap-monitor" on page 7215	

[edit logical-systems {name} protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles {profile-name} avg-fd-twoway-threshold]	
"flap-trap-monitor" on page 7215	

[edit logical-systems {name} protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles {profile-name} avg-flr-backward-threshold]	
"flap-trap-monitor" on page 7215	

[edit logical-systems {name} protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles {profile-name} avg-flr-forward-threshold]	
"flap-trap-monitor" on page 7215	

[edit logical-systems {name} protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles {profile-name} avg-ifdv-twoway-threshold]		
"flap-trap-monitor" on page 7215		

[edit logical-systems {logical-system-name} protocols ospf]		
"source-packet-routing (Protocols IS-IS)" on page 28266	"source-packet-routing (Protocols IS-IS)" on page 28266	"source-packet-routing (Protocols OSPF)" on page 28270
"topology (OSPF)" on page 29445		

[edit logical-systems {logical-system-name} protocols ospf area {area-id}]		
"label-switched-path (Protocols OSPF)" on page 9592	"network-summary-export" on page 11325	"network-summary-import" on page 11326
"peer-interface (Protocols OSPF)" on page 12370		

[edit logical-systems {logical-system-name} protocols ospf area {area-id} interface {interface-name}]		
"application-specific (Protocols OSPF)" on page 2389	"authentication (Protocols OSPF)" on page 2570	"mtu (interfaces)" on page 11027
"no-advertise-adjacency-segment (Protocols OSPF)" on page 11383	"no-neighbor-down-notification" on page 11519	"post-convergence-lfa (Protocols OSPF)" on page 12904
"topology (OSPF Interface)" on page 29446		

[edit logical-systems {logical-system-name} protocols ospf area {area-id} interface {interface-name} topology {topology-name}]		
"bandwidth-based-metrics (Protocols OSPF)" on page 2835	"bandwidth-based-metrics (Protocols OSPF)" on page 2835	

[edit logical-systems {logical-system-name} protocols ospf area {area-id} interface {interface-name} topology (default | {name})]	
"link-protection (Protocols OSPF)" on page 9814	

[edit logical-systems {logical-system-name} protocols ospf area {area-id} virtual-link]	
"authentication (Protocols OSPF)" on page 2570	

[edit logical-systems {logical-system-name} protocols ospf backup-spf-options]		
"no-source-packet-routing (Protocols OSPF)" on page 11583	"use-post-convergence-lfa (Protocols OSPF)" on page 30533	

[edit logical-systems {name} protocols ospf overload]		
"allow-route-leaking" on page 2202	"as-external" on page 2486	"stub-network" on page 28667

[edit logical-systems {logical-system-name} protocols ospf source-packet-routing]	
"node-segment (Protocols OSPF)" on page 11654	

[edit logical-systems {name} protocols ospf spf-options]	
"microloop-avoidance (Protocols OSPF)" on page 10702	

[edit logical-systems {logical-system-name} protocols ospf topology {name}]	
"topology-id" on page 29441	

[edit logical-systems {logical-system-name} protocols ospf topology (default | ipv4-multicast | {name})]		
"overload (Protocols OSPF)" on page 12051	"prefix-export-limit (Protocols OSPF)" on page 13085	"spf-options (Protocols OSPF)" on page 28381

[edit logical-systems {logical-system-name} protocols ospf topology (default | {name})]	
"backup-spf-options (Protocols OSPF)" on page 2823	

[edit logical-systems {logical-system-name} protocols ospf traffic-engineering shortcuts]	
"ignore-lsp-metrics" on page 8402	

[edit logical-systems {logical-system-name} protocols ospf3]	
"realm" on page 14093	

[edit logical-systems {logical-system-name} protocols ospf3 area {area-id}]		
"inter-area-prefix-export" on page 8782	"inter-area-prefix-import" on page 8783	

[edit logical-systems {name} protocols ospf3 overload]	
"intra-area-prefix" on page 9224	

[edit logical-systems {logical-system-name} protocols ospf3 realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) area {area-id} interface {interface-name}]	
"interface-type (Protocols OSPF)" on page 9091	

[edit logical-systems {logical-system-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)]		
"area" on page 2451	"export (OSPF)" on page 6660	"import (Protocols OSPF)" on page 8463
"lsa-refresh-interval" on page 10182	"no-nssa-abr" on page 11533	"no-rfc-1583" on page 11571
"nssa" on page 11717	"overload (Protocols OSPF)" on page 12051	"prefix-export-limit (Protocols OSPF)" on page 13085
"reference-bandwidth (Protocols OSPF)" on page 14226	"reference-bandwidth (Protocols OSPF)" on page 14226	"rib-group (Protocols OSPF)" on page 15774
"spf-options (Protocols OSPF)" on page 28381	"stub" on page 28665	"traceoptions (Protocols OSPF)" on page 29713

[edit logical-systems {logical-system-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id}]		
"inter-area-prefix-import" on page 8783	"interface (Protocols OSPF)" on page 8984	

[edit logical-systems {logical-system-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} interface {interface-name}]		
"bandwidth-based-metrics (Protocols OSPF)" on page 2835	"bfd-liveness-detection (Protocols OSPF)" on page 2941	"passive (Protocols OSPF)" on page 12223

[edit logical-systems {logical-system-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} interface {interface-name} passive]	
"traffic-engineering (Passive TE Mode)" on page 30021	

[edit logical-systems {logical-system-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} nssa]	
"default-lsa" on page 5197	

[edit logical-systems {logical-system-name} protocols ospf3 realm ipv4-unicast area {area-id}]	
"link-protection (Protocols OSPF)" on page 9814	

[edit logical-systems {name} protocols pcep pce]		
"destination-ipv6-address" on page 5443	"local-ipv6-address" on page 9956	

[edit logical-systems {logical-system-name} protocols pim]		
"assert-timeout" on page 2500	"default-vpn-source" on page 5222	"dense-groups" on page 5315
"disable (PIM)" on page 5782	"disable (Protocols)" on page 5731	"disable-packet-register" on page 5813
"dr-election-on-p2p" on page 5937	"export (Protocols PIM)" on page 6683	"family (Protocols PIM)" on page 6907
"graceful-restart (Protocols PIM)" on page 7777	"idle-standby-path-switchover-delay" on page 8301	"import (Protocols PIM)" on page 8480
"interface (Protocols PIM)" on page 8991	"join-load-balance" on page 9431	"join-prune-timeout" on page 9433
"mldp-inband-signalling (Protocols Multipoint LDP)" on page 10787	"override-interval" on page 12058	"propagation-delay" on page 13381
"reset-tracking-bit" on page 15583	"rib-group (Protocols PIM)" on page 15775	"rp" on page 15996

(Continued)

[edit logical-systems {logical-system-name} protocols pim]		
"rpf-vector (PIM)" on page 16017	"secondary-vtep-address" on page 16419	"sglimit" on page 16919
"spt-threshold" on page 28401	"standby-path-creation-delay" on page 28488	"traceoptions (Protocols PIM)" on page 29719

[edit logical-systems {logical-system-name} protocols pim default-vpn-source]	
"interface-name" on page 8934	

[edit logical-systems {logical-system-name} protocols pim family (inet | inet6)]		
"disable (PIM)" on page 5782	"disable (Protocols)" on page 5731	

[edit logical-systems {logical-system-name} protocols pim graceful-restart]	
"no-bidirectional-mode" on page 11416	

[edit logical-systems {logical-system-name} protocols pim interface {interface-name}]		
"accept-remote-source" on page 1642	"dual-dr" on page 6045	

[edit logical-systems {logical-system-name} protocols pim interface (Protocols PIM) {interface-name}]		
"bidirectional (Interface)" on page 2969	"disable (PIM)" on page 5782	"disable (Protocols)" on page 5731
"distributed-dr" on page 5833	"family (Protocols PIM Interface)" on page 6909	"family (Protocols PIM)" on page 6907

(Continued)

[edit logical-systems {logical-system-name} protocols pim interface (Protocols PIM) {interface-name}]		
"hello-interval" on page 8032	"mode (Protocols PIM)" on page 10816	"neighbor-policy" on page 11272
"override-interval" on page 12058	"priority (PIM Interfaces)" on page 13184	"propagation-delay" on page 13381
"reset-tracking-bit" on page 15583	"version (Protocols PIM)" on page 30719	

[edit logical-systems {logical-system-name} protocols pim interface (Protocols PIM) {interface-name} bidirectional]	
"df-election" on page 5572	

[edit logical-systems {logical-system-name} protocols pim interface (Protocols PIM) {interface-name} bidirectional df-election]		
"backoff-period" on page 2791	"offer-period" on page 11785	"robustness-count" on page 15823

[edit logical-systems {logical-system-name} protocols pim local family (inet | inet6) anycast-pim]	
"rp-set" on page 16001	

[edit logical-systems {logical-system-name} protocols pim mldp-inband-signalling]	
"policy (Protocols Multipoint LDP)" on page 12706	

[edit logical-systems {logical-system-name} protocols pim rp]		
"auto-rp" on page 2735	"bidirectional (RP)" on page 2971	"bootstrap" on page 3015

(Continued)

[edit logical-systems {logical-system-name} protocols pim rp]		
"bootstrap-export" on page 3016	"bootstrap-import" on page 3018	"bootstrap-priority" on page 3019
"dr-register-policy" on page 5938	"embedded-rp" on page 6248	"group-rp-mapping" on page 7891
"local" on page 9899	"register-limit" on page 14264	"rp-register-policy" on page 16000
"static (Protocols PIM)" on page 28562		

[edit logical-systems {logical-system-name} protocols pim rp auto-rp]	
"mapping-agent-election" on page 10342	

[edit logical-systems {logical-system-name} protocols pim rp bidirectional]	
"address (Bidirectional Rendezvous Points)" on page 1865	

[edit logical-systems {logical-system-name} protocols pim rp bidirectional address {address}]		
"group-ranges" on page 7885	"hold-time (Protocols PIM)" on page 8114	"priority (PIM RPs)" on page 13186

[edit logical-systems {logical-system-name} protocols pim rp bootstrap]	
"family (Bootstrap)" on page 6846	

[edit logical-systems {logical-system-name} protocols pim rp bootstrap (inet | inet6)]		
"import (Protocols PIM Bootstrap)" on page 8482	"priority (Bootstrap)" on page 13157	

[edit logical-systems {logical-system-name} protocols pim rp bootstrap family (inet | inet6)]	
"export (Bootstrap)" on page 6663	

[edit logical-systems {logical-system-name} protocols pim rp embedded-rp]		
"group-ranges" on page 7885	"maximum-rps" on page 10556	

[edit logical-systems {logical-system-name} protocols pim rp group-rp-mapping]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit logical-systems {logical-system-name} protocols pim rp local]		
"family (Local RP)" on page 6883	"override (PIM Static RP)" on page 12061	

[edit logical-systems {logical-system-name} protocols pim rp local (inet | inet6) anycast-pim rp-set]	
"address (Anycast RPs)" on page 1856	

[edit logical-systems {logical-system-name} protocols pim rp local family (inet | inet6)]		
"address (Local RPs)" on page 1885	"anycast-pim" on page 2283	"disable (PIM)" on page 5782

(Continued)

[edit logical-systems {logical-system-name} protocols pim rp local family (inet | inet6)]	
"disable (Protocols)" on page 5731	

[edit logical-systems {logical-system-name} protocols pim rp local family (inet | inet6) anycast-pim]	
"local-address (Protocols PIM)" on page 9911	

[edit logical-systems {logical-system-name} protocols pim rp local family inet]	
"override (PIM Static RP)" on page 12061	

[edit logical-systems {logical-system-name} protocols pim rp local family inet6]	
"override (PIM Static RP)" on page 12061	

[edit logical-systems {logical-system-name} protocols pim rp register-limit]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit logical-systems {logical-system-name} protocols pim rp static]	
"address (Static RPs)" on page 1933	

[edit logical-systems {logical-system-name} protocols pim rp static address {address}]		
"override (PIM Static RP)" on page 12061	"version (Protocols PIM)" on page 30719	

[edit logical-systems {name} protocols pim rprpf-vector]	
"policy (PIM rpf-vector) " on page 12703	

[edit logical-systems {logical-system-name} protocols pim sglimit]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit logical-systems {logical-system-name} protocols pim spt-threshold]	
"infinity" on page 8622	

[edit logical-systems {name} protocols ppp-service]	
"max-failures" on page 10438	

[edit logical-systems {name} protocols ptp]		
"acquiring-state-announce-grandmaster" on page 1754	"apts" on page 2436	"holdover-time-error-budget" on page 8129
"performance-monitor" on page 12429		

[edit logical-systems {name} protocols ptp stateful interface {name} unicast-mode]	
"remote-ip-address" on page 14344	

[edit logical-systems {logical-system-name} protocols rip]		
"authentication-key (Protocols RIP)" on page 2619	"authentication-type (Protocols RIP)" on page 2669	"check-zero" on page 3309
"graceful-restart (Protocols RIP)" on page 7779	"group (Protocols RIP)" on page 7877	"holddown (Protocols RIP)" on page 8125
"import (Protocols RIP)" on page 8483	"message-size" on page 10658	"metric-in (Protocols RIP)" on page 10671
"receive (Protocols RIP)" on page 14117	"rib-group (Protocols RIP)" on page 15777	"route-timeout (Protocols RIP)" on page 15885
"send (Protocols RIP)" on page 16546	"traceoptions (Protocols RIP)" on page 29727	"update-interval (Protocols RIP)" on page 30441

[edit logical-systems {logical-system-name} protocols rip group {group-name}]		
"bfd-liveness-detection (Protocols RIP)" on page 2947	"demand-circuit (Protocols RIP)" on page 5294	"export (RIP)" on page 6661
"max-retrans-time" on page 10464	"neighbor" on page 11264	"preference (Protocols RIP)" on page 13043
"route-timeout (Protocols RIP)" on page 15885	"update-interval (Protocols RIP)" on page 30441	

[edit logical-systems {logical-system-name} protocols rip group {group-name} neighbor {neighbor-name}]		
"any-sender" on page 2281	"authentication-key (Protocols RIP)" on page 2619	"authentication-type (Protocols RIP)" on page 2669

(Continued)

[edit logical-systems {logical-system-name} protocols rip group {group-name} neighbor {neighbor-name}]		
"check-zero" on page 3309	"demand-circuit (Protocols RIP)" on page 5294	"dynamic-peers" on page 6111
"import (Protocols RIP)" on page 8483	"interface-type (Protocols RIP)" on page 9095	"max-retrans-time" on page 10464
"message-size" on page 10658	"metric-in (Protocols RIP)" on page 10671	"metric-out (Protocols RIP)" on page 10674
"peer (Protocols RIP)" on page 12382	"receive (Protocols RIP)" on page 14117	"route-timeout (Protocols RIP)" on page 15885
"send (Protocols RIP)" on page 16546	"update-interval (Protocols RIP)" on page 30441	

[edit logical-systems {logical-system-name} protocols ripng]		
"graceful-restart (Protocols RIPng)" on page 7780	"group (Protocols RIPng)" on page 7880	"holddown (Protocols RIPng)" on page 8127
"import (Protocols RIPng)" on page 8485	"metric-in (Protocols RIPng)" on page 10673	"receive (Protocols RIPng)" on page 14119
"route-timeout (Protocols RIPng)" on page 15887	"send (Protocols RIPng)" on page 16548	"traceoptions (Protocols RIPng)" on page 29730
"update-interval (Protocols RIPng)" on page 30443		

[edit logical-systems {logical-system-name} protocols ripng group {group-name}]		
"export (Protocols RIPng)" on page 6686	"neighbor (Protocols RIPng)" on page 11282	"preference (Protocols RIPng)" on page 13045

[edit logical-systems {logical-system-name} protocols ripng group {group-name} neighbor {neighbor-name}]		
"import (Protocols RIPng)" on page 8485	"metric-in (Protocols RIPng)" on page 10673	"metric-out (Protocols RIPng)" on page 10679
"receive (Protocols RIPng)" on page 14119	"send (Protocols RIPng)" on page 16548	

[edit logical-systems {logical-system-name} protocols router-advertisement]		
"interface (Protocols IPv6 Neighbor Discovery)" on page 8964	"traceoptions (Protocols IPv6 Neighbor Discovery)" on page 29681	

[edit logical-systems {logical-system-name} protocols router-advertisement interface {interface-name}]		
"current-hop-limit" on page 4735	"default-lifetime" on page 5193	"dns-server-address" on page 5864
"downstream" on page 5932	"downstream-mode" on page 5933	"link-mtu" on page 9802
"managed-configuration" on page 10301	"max-advertisement-interval (Protocols IPv6 Neighbor Discovery)" on page 10415	"min-advertisement-interval (Protocols IPv6 Neighbor Discovery)" on page 10709
"other-stateful-configuration" on page 11953	"parameter-preference" on page 12173	"passive-mode" on page 12205
"preference (IPv6 Router Advertisement)" on page 13035	"prefix (Protocols IPv6 Neighbor Discovery)" on page 13100	"reachable-time" on page 14086
"retransmit-timer" on page 15689	"upstream-mode" on page 30478	

[edit logical-systems {logical-system-name} protocols router-advertisement interface {interface-name} dns-server-address {address}]	
"lifetime (Router Advertisement)" on page 9745	

[edit logical-systems {logical-system-name} protocols router-advertisement interface {interface-name} prefix {prefix}]		
"autonomous" on page 2761	"on-link" on page 11821	"preferred-lifetime" on page 13056
"valid-lifetime" on page 30644		

[edit logical-systems {logical-system-name} protocols router-discovery]		
"address (Protocols Router Discovery)" on page 1913	"disable (Protocols Router Discovery)" on page 5797	"interface (Protocols Router Discovery)" on page 8998
"traceoptions (Protocols Router Discovery)" on page 29733		

[edit logical-systems {logical-system-name} protocols router-discovery address {address}]		
"advertise" on page 2007	"broadcast (Router Discovery)" on page 3043	"ineligible" on page 8553
"multicast (Protocols Router Discovery)" on page 11077	"priority (Protocols Router Discovery)" on page 13196	

[edit logical-systems {logical-system-name} protocols router-discovery interface {interface-name}]		
"lifetime (Router Discovery)" on page 9747	"max-advertisement-interval (Protocols Router Discovery)" on page 10416	"min-advertisement-interval (Protocols Router Discovery)" on page 10711

[edit logical-systems {logical-system-name} protocols rstp]	
"extended-system-id" on page 6706	

[edit logical-systems {logical-system-name} protocols rsvp]		
"devices" on page 5571	"disable (Protocols RSVP)" on page 5798	"fast-reroute (Protocols RSVP)" on page 6944
"graceful-deletion-timeout" on page 7758	"hello-acknowledgements" on page 8029	"interface (Protocols RSVP)" on page 8999
"keep-multiplier" on page 9465	"load-balance (Protocols RSVP)" on page 9885	"no-interface-hello" on page 11479
"no-local-reversion" on page 11497	"no-node-hello-on-bypass" on page 11523	"no-node-id-subobject" on page 11524
"no-p2mp-sublsp" on page 11544	"node-hello" on page 11646	"pop-and-forward (Protocols RSVP)" on page 12803
"preemption" on page 13026	"refresh-time" on page 14257	"setup-protection" on page 16897
"traceoptions (Protocols RSVP)" on page 29736	"tunnel-services (RSVP)" on page 30238	

[edit logical-systems {logical-system-name} protocols rsvp graceful-restart]		
"disable (Protocols RSVP)" on page 5798	"maximum-helper-recovery-time" on page 10522	"maximum-helper-restart-time (RSVP)" on page 10525

[edit logical-systems {logical-system-name} protocols rsvp interface {interface-name}]		
"authentication-key (Protocols RSVP)" on page 2620	"bandwidth (Protocols RSVP)" on page 2872	"disable (Protocols RSVP)" on page 5798
"hello-interval (Protocols RSVP)" on page 8038	"link-protection (RSVP)" on page 9816	"reliable" on page 14304

(Continued)

[edit logical-systems {logical-system-name} protocols rsvp interface {interface-name}]		
"subscription" on page 28709	"update-threshold" on page 30456	"update-threshold-max-reservable" on page 30458

[edit logical-systems {logical-system-name} protocols rsvp interface {interface-name} link-protection]		
"admin-group" on page 1967	"bandwidth (Protocols RSVP)" on page 2872	"bypass (Signaled LSP)" on page 3115
"class-of-service (Protocols RSVP)" on page 3370	"disable (Protocols RSVP)" on page 5798	"exclude-srlg" on page 6629
"hop-limit" on page 8134	"max-bypasses" on page 10422	"no-cspf (Protocols RSVP)" on page 11431
"optimize-timer (Protocols RSVP)" on page 11861	"path (Protocols RSVP)" on page 12261	"priority (Protocols RSVP)" on page 13197
"subscription" on page 28709		

[edit logical-systems {logical-system-name} protocols rsvp interface {interface-name} link-protection bypass {bypass-name}]		
"admin-group" on page 1967	"bandwidth (Protocols RSVP)" on page 2872	"class-of-service (Protocols RSVP)" on page 3370
"exclude-srlg" on page 6629	"hop-limit" on page 8134	"no-cspf (Protocols RSVP)" on page 11431
"path (Protocols RSVP)" on page 12261	"priority (Protocols RSVP)" on page 13197	

[edit logical-systems {name} protocols rsvp interface {name} non-rsvp-bandwdith]		
"local-bw-override-threshold" on page 9925		

[edit logical-systems {logical-system-name} protocols rsvp interface {interface-name} update-threshold]		
"adaptive (RSVP)" on page 1828		

[edit logical-systems {logical-system-name} protocols rsvp peer-interface {peer-interface-name}]		
"aggregate (Protocols RSVP)" on page 2068	"authentication-key (Protocols RSVP)" on page 2620	"disable (Protocols RSVP)" on page 5798
"hello-interval (Protocols RSVP)" on page 8038	"reliable" on page 14304	

[edit logical-systems {logical-system-name} protocols rsvp preemption]		
"soft-preemption (Protocols RSVP)" on page 28089		

[edit logical-systems {logical-system-name} protocols sap]		
"disable (Protocols SAP)" on page 5800	"disable (Protocols)" on page 5731	"listen" on page 9842

[edit logical-systems {name} protocols source-packet-routing]		
"rib-group (Source Packet Routing)" on page 15780	"segment-list" on page 16521	"source-routing-path" on page 28316
"telemetry" on page 29054	"use-transport-class" on page 30541	

[edit logical-systems {name} protocols source-packet-routing telemetry]	
"statistics" on page 28591	

[edit logical-systems {name} protocols uplink-failure-detection group]	
"debounce-interval" on page 5154	

[edit logical-systems {logical-system-name} protocols vpls static-vpls]	
"no-tunnel-services" on page 11605	

[edit logical-systems {logical-system-name} protocols vstp]		
"access-trunk" on page 1705	"vlan (VSTP)" on page 30878	

[edit logical-systems {logical-system-name} protocols vstp vlan {vlan-id}]		
"backup-bridge-priority" on page 2793	"bridge-priority" on page 3036	"forward-delay" on page 7393
"hello-time" on page 8041	"interface (Spanning Tree)" on page 9057	"max-age" on page 10418
"system-id" on page 28847	"vpls-flush-on-topology-change" on page 30926	

[edit logical-systems {logical-system-name} protocols vstp vlan {vlan-id} interface {interface-name}]		
"cost" on page 4681	"edge" on page 6197	"mode" on page 10798
"no-root-port" on page 11572	"priority (Protocols STP)" on page 13199	

[realm]

[edit logical-systems {logical-system-name} realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id}]	
"area-range" on page 2453	

[routing]

[edit logical-systems {logical-system-name} routing -instances {instance-name} forwarding-options family vpls]	
"filter (VPLS)" on page 7158	

[routing-instances]

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name}]		
"bridge-domains" on page 3032	"description (Routing Instances)" on page 5346	"igmp-snooping-options" on page 8393
"instance-role" on page 8764	"instance-type" on page 8773	"interface (Routing Instances)" on page 9012
"interface (Virtual Tunnel in Routing Instances)" on page 9097	"l2vpn-id" on page 9562	"mac-vrf" on page 10276
"multicast-snooping-options" on page 11087	"no-normalization" on page 11526	"no-vrf-advertise" on page 11620
"no-vrf-propagate-ttl" on page 11621	"protocols" on page 13748	"protocols (STP Type)" on page 13802

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name}]		
"provider-tunnel" on page 13806	"qualified-bum-pruning-mode" on page 13905	"route-distinguisher" on page 15856
"routing-options" on page 15980	"service-type (MAC-VRF Routing Instances)" on page 16764	"switch-options" on page 28749
"vlan-id (routing instance)" on page 30817	"vlan-model" on page 30833	"vrf-advertise-selective" on page 30957
"vrf-export" on page 30958	"vrf-import" on page 30960	"vrf-propagate-ttl" on page 30965
"vrf-table-label" on page 30966	"vrf-target" on page 30968	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} ...]		
"active-leasequery (DHCP Relay Agent)" on page 1808	"bulk-leasequery (DHCP Relay Agent)" on page 3068	"leasequery (DHCP Relay Agent)" on page 9702

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} access]	
"address-protection" on page 1910	

[edit logical-systems {name} routing-instances {name} access address-assignment pool {name} family (inet | inet6)]	
"excluded-range (Address-Assignment Pools)" on page 6632	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} bridge-domains {bridge-domain-name}]		
"bridge-options" on page 3034	"interface (Logical Interfaces)" on page 8814	"routing-interface" on page 15978

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} bridge-domains {bridge-domain-name}]		
"vlan-id-list" on page 30805	"vlan-tags" on page 30865	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} bridge-domains {bridge-domain-name}, [edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name}]		
"community-vlans (MX Series)" on page 4478	"isolated-vlan (MX Series)" on page 9408	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} bridge-options]		
"community-vlans (MX Series)" on page 4478	"drop-unknown-mac-ip" on page 5984	"interface-mac-limit" on page 8910
"isolated-vlan (MX Series)" on page 9408	"mac-statistics" on page 10263	"mac-table-size" on page 10271
"no-normalization" on page 11526		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} bridge-options interface {interface-name}]		
"interface-mac-limit" on page 8910	"mac-pinning" on page 10254	"static-mac" on page 28549
"static-mac-ip" on page 28587		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} bridge-options interface {interface-name} interface-mac-limit {limit}]	
"packet-action" on page 12129	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} bridge-options interface-mac-limit {limit}]	
"packet-action" on page 12129	

[edit logical-systems {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit logical-systems {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay dhcpv6 group {name} relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay dhcpv6 relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit logical-systems {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit logical-systems {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} multicast-snooping-options]		
"flood-groups" on page 7241	"forwarding-cache (Bridge Domains)" on page 7407	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} multicast-snooping-options forwarding-cache]	
"threshold (Bridge Domains)" on page 29319	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} firewall family {family-name} filter {filter-name}]	
"physical-interface-filter" on page 12476	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} firewall policer {policer-name}]	
"physical-interface-policer" on page 12479	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} firewall three-color-policer {policer-name}]	
"physical-interface-policer" on page 12479	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options]		
"dhcp-relay" on page 5620	"evpn-vxlan" on page 6555	"fast-reroute-priority" on page 6941
"load-balance (Forwarding Options)" on page 9878		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay]		
"active-server-group" on page 1817	"dhcpv6 (DHCP Relay Agent)" on page 5665	"dual-stack-group (DHCP Relay Agent)" on page 6055
"duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)" on page 6066	"forward-snooped-clients (DHCP Relay Agent)" on page 7405	"relay-option-82" on page 14294
"server-group" on page 16633		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay ...]		
"allow-snooped-clients" on page 2204	"asymmetric-lease-time (DHCP Overrides)" on page 2519	"authentication (DHCP Relay Agent)" on page 2595
"circuit-type (DHCP Relay Agent)" on page 3332	"default-action (DHCP Relay Agent Option)" on page 5177	"drop (DHCP Relay Agent Option)" on page 5946
"dual-stack (DHCP Relay Agent Overrides)" on page 6050	"dynamic-profile (DHCP Relay Agent)" on page 6119	"forward-only (DHCP Relay Agent Option)" on page 7398
"forward-only (DHCP Relay Agent)" on page 7396	"forward-only-replies (DHCP Relay Agent)" on page 7401	"group (DHCP Relay Agent)" on page 7820
"interface (DHCP Relay Agent)" on page 8863	"interface-description (DHCP Relay Agent)" on page 8856	"interface-name (DHCP Relay Agent)" on page 8938

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay ...]		
"local-server-group (DHCP Relay Agent Option)" on page 9976	"mac-address (DHCP Relay Agent)" on page 10223	"no-allow-snooped-clients" on page 11390
"option-number (DHCP Relay Agent Option)" on page 11898	"overrides (DHCP Relay Agent)" on page 12068	"relay-option (DHCP Relay Agent)" on page 14297
"relay-server-group (DHCP Relay Agent Option)" on page 14300	"server-response-time (DHCP Relay Agent)" on page 16664	"service-profile (DHCP Relay Agent)" on page 16731
"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982	"username-include (DHCP Relay Agent)" on page 30597	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay ... overrides]		
"client-discover-match (DHCP Relay Agent)" on page 4310	"delay-authentication (DHCP Relay Agent)" on page 5242	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay ... relay-option-82]	
"remote-id (DHCP Relay Agent)" on page 14331	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay authentication]	
"password (DHCP Relay Agent)" on page 12229	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083
"option-82 (DHCP Relay Agent)" on page 11893	"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6]		
"active-server-group" on page 1817	"server-group" on page 16633	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 ...]		
"client-id (DHCP Relay Agent)" on page 4317	"relay-agent-interface-id (DHCPv6 Relay Agent)" on page 14276	"relay-agent-remote-id (DHCPv6 Relay Agent)" on page 14284
"use-vlan-id (DHCP Relay Agent)" on page 30546		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 ... (relay-agent-interface-id | relay-agent-remote-id)]	
"use-option-82" on page 30528	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 authentication]	
"password (DHCP Relay Agent)" on page 12229	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083
"relay-agent-interface-id (DHCPv6 Relay Agent Username)" on page 14278	"relay-agent-remote-id (DHCPv6 Relay Agent Username)" on page 14286	"relay-agent-subscriber-id (DHCPv6 Relay Agent)" on page 14289
"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name}]	
"active-server-group" on page 1817	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name} authentication]	
"password (DHCP Relay Agent)" on page 12229	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name} authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name} authentication username-include]		
"relay-agent-interface-id (DHCPv6 Relay Agent Username)" on page 14278	"relay-agent-remote-id (DHCPv6 Relay Agent Username)" on page 14286	"relay-agent-subscriber-id (DHCPv6 Relay Agent)" on page 14289
"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580	

[edit logical-systems {name} routing-instances {name} forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name} interface {interface-name}]	
"trace (DHCP Relay Agent)" on page 29458	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name} overrides]		
"interface-client-limit (DHCP Relay Agent)" on page 9108	"no-bind-on-request (DHCP Relay Agent)" on page 11418	"send-release-on-delete (DHCP Relay Agent)" on page 16550

[edit logical-systems {name} routing-instances {name} forwarding-options dhcp-relay dhcpv6 group {name} relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 overrides]		
"interface-client-limit (DHCP Relay Agent)" on page 9108	"no-bind-on-request (DHCP Relay Agent)" on page 11418	"send-release-on-delete (DHCP Relay Agent)" on page 16550

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]	
"keep-incoming-interface-id (DHCP Relay Agent)" on page 9462	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id]	
"keep-incoming-remote-id (DHCP Relay Agent)" on page 9463	

[edit logical-systems {name} routing-instances {name} forwarding-options dhcp-relay dhcpv6 relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6...]		
"asymmetric-lease-time (DHCP Overrides)" on page 2519	"asymmetric-prefix-lease-time (DHCP Overrides)" on page 2521	

[edit logical-systems {name} routing-instances {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"dual-stack-interface-client-limit (DHCP Local Server and Relay Agent)" on page 6058	"protocol-master" on page 13448
"relay-agent-interface-id (DHCPv6 Relay Agent)" on page 14276	"use-vlan-id (DHCP Relay Agent)" on page 30546	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name}]		
"active-server-group" on page 1817	"interface-tag (DHCP Relay Server)" on page 9083	"relay-option-82" on page 14294

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} authentication]	
"password (DHCP Relay Agent)" on page 12229	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083
"option-82 (DHCP Relay Agent)" on page 11893	"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} interface {interface-name}]	
"trace (DHCP Relay Agent)" on page 29458	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} interface {interface-name} overrides]	
"relay-source" on page 14302	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} overrides]		
"always-write-giaddr" on page 2221	"always-write-option-82" on page 2222	"disable-relay" on page 5801
"interface-client-limit (DHCP Relay Agent)" on page 9108	"layer2-unicast-replies" on page 9661	"no-bind-on-request (DHCP Relay Agent)" on page 11418
"proxy-mode" on page 13833	"relay-source" on page 14302	"replace-ip-source-with (DHCP Relay Agent)" on page 14377
"send-release-on-delete (DHCP Relay Agent)" on page 16550	"trust-option-82" on page 30168	

[edit logical-systems {name} routing-instances {name} forwarding-options dhcp-relay group {name} relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877

(Continued)

[edit logical-systems {name} routing-instances {name} forwarding-options dhcp-relay group {name} relay-option]		
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} relay-option-60]	
"vendor-option" on page 30684	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} relay-option-82]		
"circuit-id (DHCP Relay Agent)" on page 3324	"server-id-override" on page 16637	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay overrides]		
"always-write-giaddr" on page 2221	"always-write-option-82" on page 2222	"disable-relay" on page 5801

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay overrides]		
"interface-client-limit (DHCP Relay Agent)" on page 9108	"layer2-unicast-replies" on page 9661	"no-bind-on-request (DHCP Relay Agent)" on page 11418
"proxy-mode" on page 13833	"relay-source" on page 14302	"replace-ip-source-with (DHCP Relay Agent)" on page 14377
"send-release-on-delete (DHCP Relay Agent)" on page 16550	"trust-option-82" on page 30168	

[edit logical-systems {name} routing-instances {name} forwarding-options dhcp-relay relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay relay-option-60]	
"vendor-option" on page 30684	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay relay-option-82]		
"circuit-id (DHCP Relay Agent)" on page 3324	"server-id-override" on page 16637	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay relay-option-82 circuit-id]		
"keep-incoming-circuit-id (DHCP Relay Agent)" on page 9460	"keep-incoming-remote-id (DHCP Relay Agent)" on page 9463	

[edit logical-systems routing-instances {instance-name} forwarding-options helpers bootp]	
"relay-agent-option" on page 14279	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} forwarding-options load-balance]		
"indexed-load-balance" on page 8545	"per-flow" on page 12413	"per-prefix" on page 12417

[edit logical-systems {name} routing-instances {name} forwarding-options pfe-sensor npu-memory]	
"resource-list (Telemetry)" on page 15612	

[edit logical-systems {name} routing-instances {name} forwarding-options storm-control]	
"enhanced" on page 6354	

[edit logical-systems {logical-system-name} routing-instances instance-name bridge-domains {domain-name} protocols igmp-snooping]		
"traceoptions (Protocols IGMP Snooping)" on page 29678		

[edit logical-systems {logical-system-name} routing-instances {instance-name} instance-type {vpls} protocols]		
"pim-snooping" on page 12517		

[edit logical-systems {logical-system-name} routing-instances {name} interface]		
"protect-interface" on page 13409		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} interface vt-{fpc}/{pic}/{port}.{unit-number}]		
"multicast (Virtual Tunnel in Routing Instances)" on page 11092	"primary (Virtual Tunnel in Routing Instances)" on page 13148	"unicast (Virtual Tunnel in Routing Instances)" on page 30354

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} interfaces demux0 unit {logical-unit-number} demux-options]		
"underlying-interface" on page 30330		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} interfaces pp0 unit {logical-unit-number} pppoe-options]		
"underlying-interface" on page 30330		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} interfaces {interface-name} unit {logical-unit-number}]		
"demux-destination (Underlying Interface)" on page 5297	"demux-source (Underlying Interface)" on page 5309	"host-prefix-only" on page 8163

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} interfaces {interface-name} unit {logical-unit-number} family {family}]		
"demux-destination (Demux Interface)" on page 5296	"demux-source (Demux Interface)" on page 5305	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} multicast-snooping-options]		
"flood-groups" on page 7241	"forwarding-cache (Bridge Domains)" on page 7407	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} multicast-snooping-options forwarding-cache]	
"threshold (Bridge Domains)" on page 29319	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols]		
"amt (Protocols)" on page 2230	"bfd (protocols)" on page 2911	"evpn" on page 6541
"extended-isid-list" on page 6698	"extended-vlan-list" on page 6707	"extended-vni-list" on page 6710
"flex-algorithm (Protocols OSPF)" on page 7223	"interface (EVPN Routing Instances)" on page 8882	"ip-prefix-routes" on page 9250
"isis" on page 9401	"isis-instance" on page 9403	"l2vpn" on page 9559
"ldp" on page 9679	"msdp" on page 10996	"mstp" on page 11006
"multicast-mode (EVPN)" on page 11074	"mvpn" on page 11135	"mvpn (NG-MVPN)" on page 11143
"ospf" on page 11938	"ospf-instance" on page 11940	"ospf3" on page 11942

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols]		
"pbb-evpn-core" on page 12282	"pim" on page 12507	"pim-snooping" on page 12517
"rip" on page 15803	"ripng" on page 15804	"rstp" on page 16046
"rsvp" on page 16050	"traceoptions (Protocols EVPN)" on page 29672	"vpls (Routing Instance)" on page 30933
"vstp" on page 30983		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (bgp | ldp | ospf | ospf3 | pim) graceful-restart]	
"disable (Routing Graceful Restart)" on page 5738	

[edit logical-systems {name} routing-instances {name} protocols (evpn | l2vpn | vpls) mesh-group {mesh-group-name} neighbor {neighbor-id} oam bfd-liveness-detection]	
"minimum-receive-interval (BFD Liveness Detection)" on page 10754	

[edit logical-systems {name} routing-instances {name} protocols (evpn | l2vpn | vpls) neighbor {neighbor-id} oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"minimum-receive-interval (BFD Liveness Detection)" on page 10754	"multiplier (BFD Liveness Detection)" on page 11123
"no-adaptation (BFD Liveness Detection)" on page 11374	"version (BFD Liveness Detection)" on page 30697	

[edit logical-systems {name} routing-instances {name} protocols (evpn | l2vpn | vpls) oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"minimum-receive-interval (BFD Liveness Detection)" on page 10754	"multiplier (BFD Liveness Detection)" on page 11123
"no-adaptation (BFD Liveness Detection)" on page 11374	"version (BFD Liveness Detection)" on page 30697	

[edit logical-systems {name} routing-instances {name} protocols (evpn | l2vpn | vpls)mesh-group {mesh-group-name} neighbor {neighbor-id} oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"multiplier (BFD Liveness Detection)" on page 11123	"no-adaptation (BFD Liveness Detection)" on page 11374
"version (BFD Liveness Detection)" on page 30697		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (ldp | ospf | ospf3) graceful-restart]	
"helper-disable (Multiple Protocols)" on page 8048	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (mstp | rstp | vstp)]		
"bpdu-block-on-edge" on page 3024	"interface (Spanning Tree)" on page 9057	"traceoptions (Spanning Tree)" on page 29869

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (mstp | rstp | vstp) interface {interface-name}]		
"cost" on page 4681	"edge" on page 6197	"mode" on page 10798
"no-root-port" on page 11572	"priority (Protocols STP)" on page 13199	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (mstp | rstp)]		
"backup-bridge-priority" on page 2793	"bpdu-destination-mac-address (Spanning Tree)" on page 3027	"bridge-priority" on page 3036
"forward-delay" on page 7393	"hello-time" on page 8041	"max-age" on page 10418
"priority-hold-time" on page 13177	"system-id" on page 28847	"vpls-flush-on-topology-change" on page 30926

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (ospf | ospf3 | pim) graceful-restart]	
"restart-duration (Protocols)" on page 15651	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (ospf | ospf3)]		
"area" on page 2451	"backup-spf-options (Protocols OSPF)" on page 2823	"database-protection" on page 4852
"export (OSPF)" on page 6660	"graceful-restart (Protocols OSPF)" on page 7775	"import (Protocols OSPF)" on page 8463
"lsa-refresh-interval" on page 10182	"no-domain-vpn-tag" on page 11443	"no-nssa-abr" on page 11533
"no-rfc-1583" on page 11571	"overload (Protocols OSPF)" on page 12051	"prefix-export-limit (Protocols OSPF)" on page 13085
"reference-bandwidth (Protocols OSPF)" on page 14226	"rib-group (Protocols OSPF)" on page 15774	"spf-options (Protocols OSPF)" on page 28381
"traceoptions (Protocols OSPF)" on page 29713		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (ospf | ospf3) area {area-id}]		
"area-range" on page 2453	"interface (Protocols OSPF)" on page 8984	"nssa" on page 11717
"stub" on page 28665		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (ospf | ospf3) area {area-id} interface {interface-name}]		
"bandwidth-based-metrics (Protocols OSPF)" on page 2835	"bfd-liveness-detection (Protocols OSPF)" on page 2941	"interface-type (Protocols OSPF)" on page 9091
"link-protection (Protocols OSPF)" on page 9814	"passive (Protocols OSPF)" on page 12223	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (ospf | ospf3) area {area-id} interface {interface-name} passive]	
"traffic-engineering (Passive TE Mode)" on page 30021	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (ospf | ospf3) area {area-id} nssa]		
"area-range" on page 2453	"default-lsa" on page 5197	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (ospf | ospf3) traffic-engineering]		
"shortcuts (Protocols OSPF)" on page 16991	"tunnel-source-protocol (Protocols OSPF)" on page 30240	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols (rstp | vstp)]		
"force-version (IEEE 802.1D STP)" on page 7379		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols amt]		
"relay (AMT Protocol)" on page 14291	"traceoptions (Protocols AMT)" on page 29656	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols amt relay]		
"accounting (Protocols AMT Interface)" on page 1728	"family (Protocols AMT Relay)" on page 6898	"secret-key-timeout" on page 16423
"tunnel-devices (Protocols AMT)" on page 30196	"tunnel-limit (Protocols AMT)" on page 30212	"unicast-stream-limit (Protocols AMT)" on page 30351

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols amt relay family]		
"inet (AMT Protocol)" on page 8555		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols amt relay family inet]		
"anycast-prefix" on page 2285	"local-address (Protocols AMT)" on page 9904	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp]		
"accept-remote-nexthop" on page 1640	"advertise-from-main-vpn-tables" on page 2013	"advertise-inactive" on page 2015
"advertise-peer-as" on page 2020	"advertise-prefix-sid" on page 2022	"authentication-algorithm" on page 2577

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp]		
"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605	"bfd-liveness-detection (BGP)" on page 2918
"bgp" on page 2956	"cluster" on page 4367	"damping (Protocols BGP)" on page 4822
"defaults (ACX series, JRR200, MX204, PTX5000, and VRR)" on page 5225	"description (Protocols BGP)" on page 5340	"disable (Protocols BGP)" on page 5788
"egress-te" on page 6227	"egress-te-set-segment" on page 6236	"enforce-first-as" on page 6343
"export (Protocols BGP)" on page 6677	"family (Protocols BGP)" on page 6900	"forwarding-context (Protocols BGP)" on page 7453
"graceful-restart (Long-Lived for BGP Helper)" on page 7763	"group (Protocols BGP)" on page 7863	"hold-time (Protocols BGP)" on page 8104
"import (BGP)" on page 8461	"include-mp-next-hop" on page 8532	"ipsec-sa (Protocols BGP)" on page 9308
"keep" on page 9452	"local-address (Protocols BGP)" on page 9905	"local-as" on page 9916
"local-preference" on page 9971	"log-updown" on page 10042	"metric-out" on page 10676
"mtu-discovery" on page 11032	"multihop" on page 11098	"multipath (Protocols BGP)" on page 11110
"multipath-build-priority" on page 11107	"no advertise-peer-as" on page 11385	"no-aggregator-id" on page 11387

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp]		
"no-client-reflect" on page 11424	"out-delay" on page 11961	"outbound-route-filter" on page 11970
"passive (Protocols BGP)" on page 12218	"path-selection" on page 12266	"peer-as (Protocols BGP)" on page 12346
"preference (Protocols BGP)" on page 13038	"remove-private" on page 14368	"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763
"tcp-aggressive-transmission" on page 28978	"tcp-mss (Protocols BGP)" on page 28998	"tcpao-auth-mismatch" on page 29034
"traceoptions (Protocols BGP)" on page 29662	"traffic-statistics-labeled-path" on page 30049	

[edit logical-systems {name} routing-instances {name} protocols bgp bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family {family}]		
"damping (Protocols BGP)" on page 4822	"damping (Protocols BGP)" on page 4822	"delay-route-advertisements" on page 5262
"traffic-engineering (Protocols BGP)" on page 30023		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family (inet | inet6)]	
"labeled-unicast (Protocols BGP)" on page 9603	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651	"prefix-limit" on page 13087	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family (inet | inet6) labeled-unicast]	
"rib (Protocols BGP)" on page 15787	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family (inet | inet6) unicast]	
"topology (Protocols BGP)" on page 29449	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow)]	
"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]	
"restarter (Graceful Restart for BGP Restarter)" on page 15667	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]		
"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family inet (labeled-unicast | unicast | multicast) graceful-restart]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit logical-systems {name} routing-instances {name} protocols bgp family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family inet labeled-unicast]		
"aigp" on page 2092	"entropy-label (protocols bgp)" on page 6413	"explicit-null (Protocols BGP)" on page 6652

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family inet labeled-unicast]		
"resolve-vpn" on page 15602		

[edit logical-systems {name} routing-instances {name} protocols bgp family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229
"extended-nexthop" on page 6700		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family inet-mdt]		
"signaling" on page 27949		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family inet-mvpn]		
"signaling" on page 27949		

[edit logical-systems {name} routing-instances {name} protocols bgp family inet6]		
"transport (Protocols BGP)" on page 30113		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family inet6 labeled-unicast]		
"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652	

[edit logical-systems {name} routing-instances {name} protocols bgp family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols bgp family l2vpn]	
"auto-discovery-only" on page 2711	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp family route-target]	
"accepted-prefix-limit" on page 1651	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols bgp graceful-restart]		
"disable-notification-extensions (BGP Graceful Restart)" on page 5779	"disable-notification-flag (BGP Graceful Restart)" on page 5780	"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488
"long-lived (Graceful Restart for BGP Helper)" on page 10103	"restart-time (BGP Graceful Restart)" on page 15663	"stale-routes-time" on page 28478

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols bgp graceful-restart long-lived]		
"advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper)" on page 2026	"receiver (Graceful Restart for BGP Helper)" on page 14127	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols bgp graceful-restart long-lived advertise-to-non-llgr-neighbor]	
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name}]		
"accept-remote-nexthop" on page 1640	"advertise-external" on page 2011	"advertise-inactive" on page 2015
"advertise-peer-as" on page 2020	"allow" on page 2160	"as-override" on page 2489
"authentication-algorithm" on page 2577	"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605
"bfd-liveness-detection (BGP)" on page 2918	"cluster" on page 4367	"damping (Protocols BGP)" on page 4822
"description (Protocols BGP)" on page 5340	"egress-te" on page 6227	"enforce-first-as" on page 6343
"export (Protocols BGP)" on page 6677	"family (Protocols BGP)" on page 6900	"graceful-restart (Long-Lived for BGP Helper)" on page 7763
"hold-time (Protocols BGP)" on page 8104	"import (BGP)" on page 8461	"include-mp-next-hop" on page 8532
"ipsec-sa (Protocols BGP)" on page 9308	"keep" on page 9452	"local-address (Protocols BGP)" on page 9905
"local-as" on page 9916	"local-preference" on page 9971	"log-updown" on page 10042
"metric-out" on page 10676	"mtu-discovery" on page 11032	"multihop" on page 11098

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name}]		
"multipath (Protocols BGP)" on page 11110	"neighbor (Protocols BGP)" on page 11273	"no advertise-peer-as" on page 11385
"no-aggregator-id" on page 11387	"no-client-reflect" on page 11424	"out-delay" on page 11961
"outbound-route-filter" on page 11970	"passive (Protocols BGP)" on page 12218	"peer-as (Protocols BGP)" on page 12346
"peer-as-list" on page 12345	"preference (Protocols BGP)" on page 13038	"remove-private" on page 14368
"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763	"tcp-aggressive-transmission" on page 28978	"tcp-mss (Protocols BGP)" on page 28998
"traceoptions (Protocols BGP)" on page 29662	"type (Protocols BGP)" on page 30278	

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} dynamic-neighbor {dyn-name}]		
"allow" on page 2160	"peer-limit" on page 12375	

[edit logical-systems {logical-system-name} routing-instances {name} protocols bgp group {group-name} family {family} {family-modifier}]		
"add-path" on page 1839	"damping (Protocols BGP)" on page 4822	"damping (Protocols BGP)" on page 4822

(Continued)

[edit logical-systems {logical-system-name} routing-instances {name} protocols bgp group {group-name} family {family} {family-modifier}]		
"delay-route-advertisements" on page 5262	"route-target (Protocols BGP)" on page 15881	"traffic-engineering (Protocols BGP)" on page 30023

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet6)]	
"labeled-unicast (Protocols BGP)" on page 9603	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651	"prefix-limit" on page 13087	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet6) labeled-unicast]	
"rib (Protocols BGP)" on page 15787	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet6) unicast]	
"topology (Protocols BGP)" on page 29449	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family (l2vpn | route-target | inet) (labeled-unicast | flow)]	
"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]	
"restarter (Graceful Restart for BGP Restarter)" on page 15667	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]		
"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481	

[edit logical-systems {logical-system-name} routing-instances {name} protocols bgp group {group-name} family {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family inet (labeled-unicast | unicast | multicast) graceful-restart]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family inet labeled-unicast]		
"aigp" on page 2092	"entropy-label (protocols bgp)" on page 6413	"explicit-null (Protocols BGP)" on page 6652
"resolve-vpn" on page 15602		

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229
"extended-nexthop" on page 6700		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family inet-mdt]	
"signaling" on page 27949	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family inet-mvpn]	
"signaling" on page 27949	

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family inet6 labeled-unicast]		
"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652	

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols bgp group {group-name} family l2vpn]	
"auto-discovery-only" on page 2711	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} family route-target]	
"accepted-prefix-limit" on page 1651	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols bgp group {group-name} graceful-restart]		
"disable-notification-extensions (BGP Graceful Restart)" on page 5779	"disable-notification-flag (BGP Graceful Restart)" on page 5780	"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488
"long-lived (Graceful Restart for BGP Helper)" on page 10103		

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols bgp group {group-name} graceful-restart long-lived]		
"advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper)" on page 2026	"receiver (Graceful Restart for BGP Helper)" on page 14127	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols bgp group {group-name} graceful-restart long-lived advertise-to-non-llgr-neighbor]	
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810	

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} link-bandwidth]	
"auto-sense" on page 2738	

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} metric-out]		
"effective-aigp" on page 6203	"minimum-effective-aigp" on page 10730	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address}]		
"accept-remote-nexthop" on page 1640	"add-path" on page 1839	"advertise-external" on page 2011
"advertise-inactive" on page 2015	"advertise-peer-as" on page 2020	"as-override" on page 2489
"authentication-algorithm" on page 2577	"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605
"bfd-liveness-detection (BGP)" on page 2918	"cluster" on page 4367	"damping (Protocols BGP)" on page 4822

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address}]		
"description (Protocols BGP)" on page 5340	"egress-te" on page 6227	"egress-te-adj-segment" on page 6229
"egress-te-node-segment" on page 6234	"enforce-first-as" on page 6343	"export (Protocols BGP)" on page 6677
"family (Protocols BGP)" on page 6900	"graceful-restart (Long-Lived for BGP Helper)" on page 7763	"hold-time (Protocols BGP)" on page 8104
"import (BGP)" on page 8461	"include-mp-next-hop" on page 8532	"ipsec-sa (Protocols BGP)" on page 9308
"keep" on page 9452	"local-address (Protocols BGP)" on page 9905	"local-as" on page 9916
"local-interface (IPv6)" on page 9952	"local-preference" on page 9971	"log-updown" on page 10042
"metric-out" on page 10676	"mtu-discovery" on page 11032	"multihop" on page 11098
"multipath (Protocols BGP)" on page 11110	"no advertise-peer-as" on page 11385	"no-aggregator-id" on page 11387
"no-client-reflect" on page 11424	"out-delay" on page 11961	"outbound-route-filter" on page 11970
"passive (Protocols BGP)" on page 12218	"peer-as (Protocols BGP)" on page 12346	"preference (Protocols BGP)" on page 13038
"remove-private" on page 14368	"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763	"tcp-mss (Protocols BGP)" on page 28998

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address}]	
"traceoptions (Protocols BGP)" on page 29662	

[edit logical-systems {logical-system-name} routing-instances {name} protocols bgp group {group-name} neighbor {address} {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} neighbor {address} bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family {family}]		
"damping (Protocols BGP)" on page 4822	"delay-route-advertisements" on page 5262	"route-target (Protocols BGP)" on page 15881
"traffic-engineering (Protocols BGP)" on page 30023		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (inet | inet6)]	
"labeled-unicast (Protocols BGP)" on page 9603	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651	"prefix-limit" on page 13087	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (inet | inet6) labeled-unicast]	
"rib (Protocols BGP)" on page 15787	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (inet | inet6) unicast]	
"topology (Protocols BGP)" on page 29449	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (l2vpn | route-target | inet) (labeled-unicast | flow)]	
"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]	
"restarter (Graceful Restart for BGP Restarter)" on page 15667	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]		
"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family inet (labeled-unicast | unicast | multicast) graceful-restart]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family inet labeled-unicast]		
"aigp" on page 2092	"entropy-label (protocols bgp)" on page 6413	"explicit-null (Protocols BGP)" on page 6652
"resolve-vpn" on page 15602		

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229
"extended-nexthop" on page 6700		

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family inet6 labeled-unicast]		
"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652	

[edit logical-systems {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols bgp group {group-name} neighbor {address} family l2vpn]	
"auto-discovery-only" on page 2711	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family route-target]	
"accepted-prefix-limit" on page 1651	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols bgp group {group-name} neighbor {address} graceful-restart]		
"advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper)" on page 2026	"disable-notification-extensions (BGP Graceful Restart)" on page 5779	"disable-notification-flag (BGP Graceful Restart)" on page 5780
"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488	"long-lived (Graceful Restart for BGP Helper)" on page 10103	"receiver (Graceful Restart for BGP Helper)" on page 14127

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols bgp group {group-name} neighbor {address} graceful-restart long-lived advertise-to-non-llgr-neighbor]	
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp group {group-name} outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit logical-systems {name} routing-instances {name} protocols bgp multipath]	
"pause-computation-during-churn" on page 12277	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols bgp outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit logical-systems {name} routing-instances {name} protocols bgp source-packet-routing]		
"srgb (Protocols BGP)" on page 28407	"srv6 (Protocols BGP)" on page 28417	

[edit logical-systems {name} routing-instances {name} protocols evpn]		
"assisted-replication" on page 2505	"auto-service-id" on page 2739	"default-gateway" on page 5190
"drop-unknown-mac-ip" on page 5984	"duplicate-mac-detection" on page 6068	"flexible-cross-connect-vlan-aware" on page 7230
"flexible-cross-connect-vlan-unaware" on page 7232	"mac-ip-table-size" on page 10231	"no-control-word " on page 11426
"source-packet-routing (Routing Instances Protocols EVPN)" on page 28273		

[edit logical-systems {name} routing-instances {name} protocols evpn assisted-replication]		
"leaf" on page 9689	"replicator" on page 14387	

[edit logical-systems {name} routing-instances {name} protocols evpn interface]	
"static-mac-ip" on page 28587	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols igmp]	
"amt (IGMP)" on page 2228	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols igmp amt relay defaults]		
"accounting (Protocols IGMP AMT Interface)" on page 1731	"group-policy (Protocols IGMP AMT Interface)" on page 7856	"query-interval (Protocols IGMP AMT)" on page 13930
"query-response-interval (Protocols IGMP AMT)" on page 13941	"robust-count (Protocols IGMP AMT)" on page 15820	"ssm-map (Protocols IGMP AMT)" on page 28465
"version (Protocols IGMP AMT)" on page 30716		

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols igmp-snooping]	
"traceoptions (Protocols IGMP Snooping)" on page 29678	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols isis]		
"backup-spf-options (Protocols IS-IS)" on page 2821	"clns-updown-compatibility" on page 4350	"export (ISIS)" on page 6658
"ignore-attached-bit" on page 8396	"import (Protocols IS-IS)" on page 8476	"interface (Protocols IS-IS)" on page 8966
"label-switched-path (Protocols IS-IS)" on page 9582	"layer2-map" on page 9651	"level (Global IS-IS)" on page 9724
"loose-authentication-check" on page 10137	"lsp-lifetime" on page 10189	"max-areas" on page 10420
"max-lsp-size" on page 10446	"max-snp-size" on page 10483	"no-ipv4-routing" on page 11483
"no-ipv6-routing" on page 11484	"overload (Protocols IS-IS)" on page 12048	"reference-bandwidth (Protocols IS-IS)" on page 14224

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols isis]		
"rib-group (Protocols IS-IS)" on page 15770	"source-packet-routing (Protocols IS-IS)" on page 28266	"spf-options (Protocols IS-IS)" on page 28379
"topologies (Protocols IS-IS)" on page 29438	"traceoptions (Protocols IS-IS)" on page 29684	

[edit logical-systems {name} routing-instances {name} protocols isis backup-spf-options]	
"use-post-convergence-lfa" on page 30531	

[edit logical-systems {name} routing-instances {name} protocols isis dynamic-overload]		
"auto-recovery" on page 2733	"no-overload-on-prefix-export-limit" on page 11542	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols isis interface {interface-name}]		
"bfd-liveness-detection (Protocols IS-IS)" on page 2933	"level (IS-IS Interfaces)" on page 9728	"link-protection (Protocols IS-IS)" on page 9811
"no-advertise-adjacency-segment (Protocols IS-IS)" on page 11382	"passive (Protocols IS-IS)" on page 12219	

[edit logical-systems {name} routing-instances {name} protocols isis interface {name} delay-measurement]	
"advertisement (Delay Measurement)" on page 2031	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols isis interface {interface-name} level {level-number}]		
"flood-reflector (Protocols IS-IS)" on page 7242	"hold-time (Protocols IS-IS)" on page 8108	"metric (Protocols IS-IS)" on page 10684
"passive (Protocols IS-IS)" on page 12219	"post-convergence-lfa" on page 12903	

[edit logical-systems {name} routing-instances {name} protocols isis interface {name} level {number} post-convergence-lfa]	
"node-protection" on page 11651	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols isis interface-group {group-name}]	
"link-group-protection (Protocols IS-IS)" on page 9794	

[edit logical-systems {name} routing-instances {name} protocols isis interface-group {name} level {name} topology]	
"bandwidth-based-metrics" on page 2837	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols isis level]		
"max-lsp-size" on page 10446	"prefix-export-limit (Protocols IS-IS)" on page 13083	"source-packet-routing (Protocols IS-IS)" on page 28266

[edit logical-systems {name} routing-instances {name} protocols isis no-external-export protocol]	
"no-external-export" on page 11454	

[edit logical-systems {name} routing-instances {name} protocols isis source-packet-routing]		
"flex-algorithm" on page 7221	"srgb (Protocols IS-IS)" on page 28408	"strict-asla-based-flex-algorithm" on page 28645
"traffic-statistics (Protocols IS-IS)" on page 30048		

[edit logical-systems {name} routing-instances {name} protocols isis spf-options]	
"microloop-avoidance" on page 10700	

[edit logical-systems {name} routing-instances {name} protocols isis traffic-engineering]		
"advertisement" on page 2030	"family (Protocols IS-IS)" on page 6905	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols l2vpn]		
"control-word (Protocols Layer 2 VPN)" on page 4648	"encapsulation-type (Layer 2 VPNs)" on page 6310	"interface (Protocols Layer 2 VPN)" on page 8977
"no-control-word (Protocols Layer 2 VPN)" on page 11427	"site (Layer 2 Circuits)" on page 27987	"traceoptions (Protocols Layer 2 VPN)" on page 29690
"vpls-id" on page 30927		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols l2vpn interface {interface-name}]		
"ignore-mtu-mismatch" on page 8408	"mtu (interfaces)" on page 11027	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols l2vpn mesh-group {mesh-group-name}]		
"route-distinguisher" on page 15856	"vpls-id" on page 30927	"vrf-target" on page 30968

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols l2vpn neighbor {address}]	
"encapsulation-type (Layer 2 VPNs)" on page 6310	

[edit logical-systems {name} routing-instances {instance_name} protocols l2vpn oam]		
"bfd-liveness-detection (Layer 2 VPN and VPLS)" on page 2929	"control-channel (Protocols OAM)" on page 4623	"ping-interval" on page 12545

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols l2vpn site {site-name}]		
"site-identifier (Layer 2 Circuits)" on page 27984	"site-preference" on page 27989	"source-attachment-identifier (Protocols VPWS)" on page 28183

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols l2vpn site {site-name} interface {interface-name}]		
"description (Protocols Layer 2 VPN)" on page 5343	"remote-site-id" on page 14360	"target-attachment-identifier (Protocols VPWS)" on page 28947

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ldp]		
"deaggregate" on page 5153	"egress-policy" on page 6215	"explicit-null (Protocols BGP)" on page 6652
"explicit-null (Protocols LDP)" on page 6654	"export (Protocols LDP)" on page 6680	"graceful-restart (Protocols LDP)" on page 7773

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ldp]		
"igp-synchronization" on page 8414	"import (Protocols LDP)" on page 8477	"interface (Protocols LDP)" on page 8978
"keepalive-interval" on page 9468	"keepalive-timeout" on page 9474	"l2-smart-policy" on page 9528
"label-withdrawal-delay" on page 9598	"log-updown (Protocols LDP)" on page 10043	"no-forwarding" on page 11464
"p2mp (Protocols LDP)" on page 12122	"policing (Protocols LDP)" on page 12662	"preference (Protocols LDP)" on page 13040
"session (Protocols LDP)" on page 16845	"session-group" on page 16810	"session-protection" on page 16843
"strict-targeted-hellos" on page 28654	"targeted-hello" on page 28962	"traceoptions (Protocols LDP)" on page 29693
"track-igp-metric" on page 29989	"traffic-statistics (Protocols LDP)" on page 30053	"transport-address" on page 30118

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ldp graceful-restart]		
"helper-disable (LDP)" on page 8046	"maximum-neighbor-recovery-time" on page 10538	"recovery-time (LDP)" on page 14148

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ldp interface {interface-name}]		
"disable (Protocols LDP)" on page 5789	"hello-interval (Protocols LDP)" on page 8036	"hold-time (Protocols LDP)" on page 8110

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ldp session {session-address}]		
"authentication-algorithm" on page 2577	"authentication-key (Protocols LDP)" on page 2618	"authentication-key-chain (Protocols LDP)" on page 2607

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ldp targeted-hello]		
"hello-interval (Protocols LDP)" on page 8036	"hold-time (Protocols LDP)" on page 8110	

[edit logical-systems {name} routing-instances {name} protocols mpls]	
"cspf-backoff-time" on page 4731	

[edit logical-systems {name} routing-instances {name} protocols mpls label-switched-path {name}]		
"conditional-metric" on page 4506	"cspf-backoff-time" on page 4731	"pop-and-forward (Protocols MPLS)" on page 12802

[edit logical-systems {name} routing-instances {name} protocols mpls label-switched-path {name} auto-bandwidth]	
"adjust-threshold-absolute" on page 1954	

[edit logical-systems {name} routing-instances {name} protocols mpls label-switched-path {name} secondary {name} non-active-hold-priority]	
"non-active-hold-priority" on page 11657	

[edit logical-systems {name} routing-instances {name} protocols mpls static-label-switched-path]		
"segment" on page 16519	"transit" on page 30064	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols msdp]		
"active-source-limit" on page 1819	"data-encapsulation" on page 4828	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681
"group (Protocols MSDP)" on page 7873	"hold-time (Protocols MSDP)" on page 8112	"import (Protocols MSDP)" on page 8478
"keep-alive (Protocols MSDP)" on page 9457	"local-address (Protocols MSDP) " on page 9910	"peer (Protocols MSDP)" on page 12380
"rib-group (Protocols MSDP)" on page 15772	"sa-hold-time (Protocols MSDP)" on page 16262	"source (Protocols MSDP) " on page 28308
"traceoptions (Protocols MSDP)" on page 29703	"traceoptions (Protocols MSDP)" on page 29703	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols msdp active-source-limit]		
"log-interval (Protocols MSDP)" on page 10014	"log-warning (Protocols MSDP)" on page 10049	"maximum (MSDP Active Source Messages) " on page 10536
"threshold (MSDP Active Source Messages) " on page 29327		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols msdp group {group-name}]		
"default-peer" on page 5200	"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731
"export (Protocols MSDP)" on page 6681	"import (Protocols MSDP)" on page 8478	"local-address (Protocols MSDP) " on page 9910

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols msdp group {group-name}]		
"mode (Protocols MSDP)" on page 10814	"peer (Protocols MSDP)" on page 12380	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols msdp group {group-name} peer {address}]		
"active-source-limit" on page 1819	"authentication-key (Protocols MSDP)" on page 2600	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681
"hold-time (Protocols MSDP)" on page 8112	"import (Protocols MSDP)" on page 8478	"keep-alive (Protocols MSDP)" on page 9457
"local-address (Protocols MSDP) " on page 9910	"sa-hold-time (Protocols MSDP)" on page 16262	"traceoptions (Protocols MSDP)" on page 29703

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols msdp peer {address}]		
"active-source-limit" on page 1819	"authentication-key (Protocols MSDP)" on page 2600	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681
"hold-time (Protocols MSDP)" on page 8112	"import (Protocols MSDP)" on page 8478	"keep-alive (Protocols MSDP)" on page 9457
"local-address (Protocols MSDP) " on page 9910	"sa-hold-time (Protocols MSDP)" on page 16262	"traceoptions (Protocols MSDP)" on page 29703

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols msdp source {ip-address}/{prefix-length}]	
"active-source-limit" on page 1819	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mstp]		
"configuration-name" on page 4516	"disable (STP)" on page 5739	"max-hops" on page 10441
"msti" on page 11003	"revision-level" on page 15735	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mstp msti {msti-id}]		
"bridge-priority" on page 3036	"vlan (MSTP)" on page 30835	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mstp msti {msti-id} interface {interface-name}]		
"cost" on page 4681	"edge" on page 6197	"priority (Protocols STP)" on page 13199

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvpn]		
"hot-root-standby (MBGP MVPN)" on page 8184	"mvpn-mode" on page 11142	"route-target (Protocols MVPN)" on page 15883
"sender-based-rpf (MBGP MVPN)" on page 16552	"static-umh (MBGP MVPN)" on page 28585	"traceoptions (Protocols MVPN)" on page 29706

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvpn family {inet | inet6}]	
"autodiscovery-only" on page 2752	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvpn family ({inet | inet6}) autodiscovery-only]	
"intra-as" on page 9226	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvpn family ({inet | inet6}) autodiscovery-only intra-as]	
"inclusive" on page 8539	

[edit logical-systems {profile-name} routing-instances {instance-name} protocols mvpn mvpn-mode]		
"rpt-spt" on page 16040	"spt-only" on page 28400	

[edit logical-systems {logical-system--name} routing-instances {instance-name} protocols mvpn mvpn-mode spt-only]	
"source-active-advertisement" on page 28114	

[edit logical-systems {logical-system--name} routing-instances {instance-name} protocols mvpn mvpn-mode spt-only source-active-advertisement]		
"dampen (Protocols MVPN)" on page 4813	"min-rate (source-active-advertisement)" on page 10718	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvpn route-target]		
"export-target" on page 6693	"import-target" on page 8489	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvpn route-target import-target]		
"target (Routing Instances MVPN)" on page 28954	"unicast (Route Target Community)" on page 30350	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvpn static-umh]		
"backup (MBGP MVPN)" on page 2800	"primary (MBGP MVPN)" on page 13143	"source-tree (MBGP MVPN)" on page 28330

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvrp]	
"no-dynamic-vlan" on page 11446	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvrp] (for virtual switch instance type)]	
"interface (MVRP)" on page 8931	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvrp] (for virtual switch instance type), [edit protocols mvrp]	
"traceoptions (MVRP)" on page 29607	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvrp interface (all | {interface-name})] (for virtual switch instance type), [edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvrp] (for virtual switch instance type), [edit logical-systems {logical-system-name} protocols mvrp interface (all | {interface-name})]	
"join-timer (MVRP)" on page 9434	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvrp interface (all | {interface-name})]	
"point-to-point (MVRP) " on page 12620	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols mvrp interface (all | {interface-name})] (for virtual switch instance type)]	
"registration" on page 14268	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf]		
"sham-link" on page 16921	"source-packet-routing (Protocols IS-IS)" on page 28266	"source-packet-routing (Protocols IS-IS)" on page 28266
"source-packet-routing (Protocols OSPF)" on page 28270	"topology (OSPF)" on page 29445	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf area {area-id}]		
"label-switched-path (Protocols OSPF)" on page 9592	"network-summary-export" on page 11325	"network-summary-import" on page 11326
"sham-link-remote" on page 16923	"virtual-link" on page 30763	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf area {area-id} interface {interface-name}]		
"application-specific (Protocols OSPF)" on page 2389	"authentication (Protocols OSPF)" on page 2570	"mtu (interfaces)" on page 11027
"no-advertise-adjacency-segment (Protocols OSPF)" on page 11383	"post-convergence-lfa (Protocols OSPF)" on page 12904	"topology (OSPF Interface)" on page 29446

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf area {area-id} sham-link-remote {address}]	
"metric (Protocols OSPF Sham Link)" on page 10688	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf area {area-id} virtual-link]	
"authentication (Protocols OSPF)" on page 2570	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf backup-spf-options]		
"no-source-packet-routing (Protocols OSPF)" on page 11583	"use-post-convergence-lfa (Protocols OSPF)" on page 30533	

[edit logical-systems {name} routing-instances {name} protocols ospf overload]		
"allow-route-leaking" on page 2202	"as-external" on page 2486	"stub-network" on page 28667

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf source-packet-routing]	
"node-segment (Protocols OSPF)" on page 11654	

[edit logical-systems {name} routing-instances {name} protocols ospf spf-options]	
"microloop-avoidance (Protocols OSPF)" on page 10702	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf topology {name}]	
"topology-id" on page 29441	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf topology (default | ipv4-multicast | {name})]		
"overload (Protocols OSPF)" on page 12051	"prefix-export-limit (Protocols OSPF)" on page 13085	"spf-options (Protocols OSPF)" on page 28381

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf topology (default | {name})]	
"backup-spf-options (Protocols OSPF)" on page 2823	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf3]	
"realm" on page 14093	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf3 area {area-id}]		
"inter-area-prefix-export" on page 8782	"inter-area-prefix-import" on page 8783	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf3 realm (ip4-unicast | ipv4-multicast | ipv6-multicast) area {area-id}]	
"inter-area-prefix-export" on page 8782	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) area {area-id} interface {interface-name}]	
"interface-type (Protocols OSPF)" on page 9091	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)]		
"area" on page 2451	"export (OSPF)" on page 6660	"import (Protocols OSPF)" on page 8463
"lsa-refresh-interval" on page 10182	"no-nssa-abr" on page 11533	"no-rfc-1583" on page 11571
"nssa" on page 11717	"overload (Protocols OSPF)" on page 12051	"prefix-export-limit (Protocols OSPF)" on page 13085
"rib-group (Protocols OSPF)" on page 15774	"spf-options (Protocols OSPF)" on page 28381	"stub" on page 28665
"traceoptions (Protocols OSPF)" on page 29713		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id}]		
"inter-area-prefix-import" on page 8783	"interface (Protocols OSPF)" on page 8984	

[edit logical-systems {logical-system-name} routing-instances {routing-instances} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} interface {interface-name}]		
"bandwidth-based-metrics (Protocols OSPF)" on page 2835	"bfd-liveness-detection (Protocols OSPF)" on page 2941	"passive (Protocols OSPF)" on page 12223

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} interface {interface-name} passive]	
"traffic-engineering (Passive TE Mode)" on page 30021	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} nssa]	
"default-lsa" on page 5197	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim]		
"assert-timeout" on page 2500	"dense-groups" on page 5315	"disable (PIM)" on page 5782
"disable (Protocols)" on page 5731	"disable-packet-register" on page 5813	"dr-election-on-p2p" on page 5937
"export (Protocols PIM)" on page 6683	"family (Protocols PIM)" on page 6907	"graceful-restart (Protocols PIM)" on page 7777
"idle-standby-path-switchover-delay" on page 8301	"import (Protocols PIM)" on page 8480	"interface (Protocols PIM)" on page 8991
"join-load-balance" on page 9431	"join-prune-timeout" on page 9433	"mdt" on page 10590
"mvpn (Draft-Rosen MVPN)" on page 11138	"rib-group (Protocols PIM)" on page 15775	"rp" on page 15996
"rpf-vector (PIM)" on page 16017	"secondary-vtep-address" on page 16419	"sglimit" on page 16919
"spt-threshold" on page 28401	"standby-path-creation-delay" on page 28488	"traceoptions (Protocols PIM)" on page 29719

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim]		
"tunnel-devices (Tunnel-Capable PICs) " on page 30198	"vpn-group-address" on page 30939	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim graceful-restart]	
"no-bidirectional-mode" on page 11416	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim interface {interface-name}]		
"accept-remote-source" on page 1642	"dual-dr" on page 6045	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim interface (Protocols PIM) {interface-name}]		
"bidirectional (Interface)" on page 2969	"disable (PIM)" on page 5782	"disable (Protocols)" on page 5731
"distributed-dr" on page 5833	"family (Protocols PIM Interface)" on page 6909	"family (Protocols PIM)" on page 6907
"hello-interval" on page 8032	"mode (Protocols PIM)" on page 10816	"neighbor-policy" on page 11272
"override-interval" on page 12058	"priority (PIM Interfaces)" on page 13184	"propagation-delay" on page 13381
"reset-tracking-bit" on page 15583	"version (Protocols PIM)" on page 30719	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim interface (Protocols PIM) {interface-name} bidirectional]		
"df-election" on page 5572		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim interface (Protocols PIM) {interface-name} bidirectional df-election]		
"backoff-period" on page 2791	"offer-period" on page 11785	"robustness-count" on page 15823

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim local family (inet | inet6) anycast-pim]		
"rp-set" on page 16001		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim mdt]		
"threshold (Routing Instances)" on page 29338	"tunnel-limit (Routing Instances)" on page 30213	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim mdt threshold]		
"group (Protocols PIM)" on page 7876		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim mdt threshold group {group-address}]		
"source (Routing Instances)" on page 28312		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim mdt threshold group {group-address} source {source-address}]	
"rate (Routing Instances)" on page 14073	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim mvpn family {inet}]	
"autodiscovery" on page 2750	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim mvpn family ({inet | inet6}) autodiscovery]	
"inet-mdt" on page 8570	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp]		
"auto-rp" on page 2735	"bidirectional (RP)" on page 2971	"bootstrap" on page 3015
"bootstrap-export" on page 3016	"bootstrap-import" on page 3018	"bootstrap-priority" on page 3019
"dr-register-policy" on page 5938	"embedded-rp" on page 6248	"group-rp-mapping" on page 7891
"local" on page 9899	"register-limit" on page 14264	"rp-register-policy" on page 16000
"static (Protocols PIM)" on page 28562		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp auto-rp]	
"mapping-agent-election" on page 10342	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp bidirectional]	
"address (Bidirectional Rendezvous Points)" on page 1865	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols pim rp bidirectional address {address}]		
"group-ranges" on page 7885	"hold-time (Protocols PIM)" on page 8114	"priority (PIM RPs)" on page 13186

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp bootstrap]	
"family (Bootstrap)" on page 6846	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp bootstrap (inet | inet6)]		
"import (Protocols PIM Bootstrap)" on page 8482	"priority (Bootstrap)" on page 13157	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp bootstrap family (inet | inet6)]	
"export (Bootstrap)" on page 6663	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp embedded-rp]		
"group-ranges" on page 7885	"maximum-rps" on page 10556	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp group-rp-mapping]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp local]		
"family (Local RP)" on page 6883	"override (PIM Static RP)" on page 12061	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp local (inet | inet6) anycast-pim rp-set]	
"address (Anycast RPs)" on page 1856	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp local family (inet | inet6)]		
"address (Local RPs)" on page 1885	"anycast-pim" on page 2283	"disable (PIM)" on page 5782
"disable (Protocols)" on page 5731		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp local family (inet | inet6) anycast-pim]	
"local-address (Protocols PIM)" on page 9911	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols pim rp local family inet]	
"override (PIM Static RP)" on page 12061	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols pim rp local family inet6]	
"override (PIM Static RP)" on page 12061	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp register-limit]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit logical-systems {name} routing-instances {name} protocols pim rp rpf-vector]	
"policy (PIM rpf-vector) " on page 12703	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim rp static]	
"address (Static RPs)" on page 1933	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols pim rp static address {address}]		
"override (PIM Static RP)" on page 12061	"version (Protocols PIM)" on page 30719	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim sglimit]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols pim spt-threshold]	
"infinity" on page 8622	

[edit logical-systems {<logical-system-name>} routing-instances {<instance-name>} protocols pim-snooping]		
"traceoptions (PIM Snooping)" on page 29638	"vlan (PIM Snooping)" on page 30837	

[edit logical-systems {<logical-system-name>} routing-instances {<instance-name>} protocols pim-snooping traceoptions]	
"no-dr-flood (PIM Snooping)" on page 11445	

[edit logical-systems {<logical-system-name>} routing-instances {<instance-name>} protocols pim-snooping vlan{<vlan-id>}]	
"no-dr-flood (PIM Snooping)" on page 11445	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols rip]		
"authentication-key (Protocols RIP)" on page 2619	"authentication-type (Protocols RIP)" on page 2669	"check-zero" on page 3309

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols rip]		
"group (Protocols RIP)" on page 7877	"holddown (Protocols RIP)" on page 8125	"import (Protocols RIP)" on page 8483
"message-size" on page 10658	"metric-in (Protocols RIP)" on page 10671	"receive (Protocols RIP)" on page 14117
"rib-group (Protocols RIP)" on page 15777	"route-timeout (Protocols RIP)" on page 15885	"send (Protocols RIP)" on page 16546
"traceoptions (Protocols RIP)" on page 29727	"update-interval (Protocols RIP)" on page 30441	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols rip group {group-name}]		
"export (RIP)" on page 6661	"neighbor" on page 11264	"preference (Protocols RIP)" on page 13043
"route-timeout (Protocols RIP)" on page 15885	"update-interval (Protocols RIP)" on page 30441	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols rip group {group-name} neighbor {neighbor-name}]		
"any-sender" on page 2281	"authentication-key (Protocols RIP)" on page 2619	"authentication-type (Protocols RIP)" on page 2669
"bfd-liveness-detection (Protocols RIP)" on page 2947	"check-zero" on page 3309	"demand-circuit (Protocols RIP)" on page 5294
"import (Protocols RIP)" on page 8483	"max-retrans-time" on page 10464	"message-size" on page 10658

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols rip group {group-name} neighbor {neighbor-name}]		
"metric-in (Protocols RIP)" on page 10671	"metric-out (Protocols RIP)" on page 10674	"receive (Protocols RIP)" on page 14117
"route-timeout (Protocols RIP)" on page 15885	"send (Protocols RIP)" on page 16546	"update-interval (Protocols RIP)" on page 30441

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ripng]		
"graceful-restart (Protocols RIPng)" on page 7780	"group (Protocols RIPng)" on page 7880	"holddown (Protocols RIPng)" on page 8127
"import (Protocols RIPng)" on page 8485	"metric-in (Protocols RIPng)" on page 10673	"receive (Protocols RIPng)" on page 14119
"route-timeout (Protocols RIPng)" on page 15887	"send (Protocols RIPng)" on page 16548	"traceoptions (Protocols RIPng)" on page 29730
"update-interval (Protocols RIPng)" on page 30443		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ripng group {group-name}]		
"export (Protocols RIPng)" on page 6686	"neighbor (Protocols RIPng)" on page 11282	"preference (Protocols RIPng)" on page 13045

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ripng group {group-name} neighbor {neighbor-name}]		
"import (Protocols RIPng)" on page 8485	"metric-in (Protocols RIPng)" on page 10673	"metric-out (Protocols RIPng)" on page 10679

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols ripng group {group-name} neighbor {neighbor-name}]		
"receive (Protocols RIPng)" on page 14119	"send (Protocols RIPng)" on page 16548	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols rstp]	
"extended-system-id" on page 6706	

[edit logical-systems {name} routing-instances {name} protocols rsvp]	
"pop-and-forward (Protocols RSVP)" on page 12803	

[edit logical-systems {name} routing-instances {name} protocols rsvp interface]	
"update-threshold-max-reservable" on page 30458	

[edit logical-systems {name} routing-instances {name} protocols rsvp interface {name} non-rsvp-bandwdith]	
"local-bw-override-threshold" on page 9925	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls]		
"bum-hashing" on page 3070	"connectivity-type" on page 4563	"control-word" on page 4644
"encapsulation-type (Layer 2 VPNs)" on page 6310	"ignore-mtu-mismatch" on page 8408	"import-labeled-routes (Routing Instances VPLS)" on page 8469
"interface (VPLS Routing Instances)" on page 9104	"label-block-size" on page 9577	"mac-flush" on page 10229

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls]		
"mac-table-aging-time (MX Series)" on page 10265	"mac-table-size (VPLS)" on page 10269	"mesh-group (Protocols VPLS)" on page 10641
"multi-homing (VPLS Multihoming for FEC 129)" on page 11052	"neighbor (Protocols VPLS)" on page 11284	"no-control-word " on page 11426
"no-tunnel-services" on page 11605	"site (VPLS Multihoming for FEC 128)" on page 27994	"site-range" on page 27990
"traceoptions (Protocols VPLS)" on page 29744	"tunnel-services (Routing Instances VPLS)" on page 30236	"vpls-id" on page 30927

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls mesh-group {mesh-group-name}]		
"interface (VPLS Mesh-Group)" on page 9102	"local-switching (VPLS)" on page 9979	"mac-flush" on page 10229
"neighbor (Protocols VPLS)" on page 11284	"peer-as (VPLS)" on page 12348	"route-distinguisher" on page 15856
"vpls-id" on page 30927	"vrf-export" on page 30958	"vrf-import" on page 30960
"vrf-target" on page 30968		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls mesh-group {mesh-group-name} neighbor {neighbor-id}]		
"ignore-encapsulation-mismatch" on page 8397	"static (Protocols VPLS)" on page 28564	

[edit logical-systems {name} routing-instances {instance_name} protocols vpls mesh-group {mesh-group-name} neighbor {neighbor-id} oam]		
"bfd-liveness-detection (Layer 2 VPN and VPLS)" on page 2929	"ping-interval" on page 12545	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls mesh-group {mesh-group-name} neighbor {address} pseudowire-status-tlv]	
"hot-standby-vc-on (Protocols Layer 2 Circuit)" on page 8190	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls mesh-group {mesh-group-name} neighbor {address} static]		
"control-word" on page 4644	"no-control-word " on page 11426	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols vpls multi-homing]		
"active-interface (VPLS Multihoming)" on page 1805	"peer-active (VPLS Multihoming for FEC 129)" on page 12341	"site (VPLS Multihoming for FEC 129)" on page 27995

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols vpls multi-homing active-interface]		
"any (VPLS Multihoming)" on page 2282	"primary (VPLS Multihoming)" on page 13151	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols vpls multi-homing site {site-name}]		
"active-interface (VPLS Multihoming)" on page 1805	"identifier (VPLS Multihoming for FEC 129)" on page 8288	"interface (VPLS Multihoming for FEC 129)" on page 9103

(Continued)

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols vpls multi-homing site {site-name}]		
"peer-active (VPLS Multihoming for FEC 129)" on page 12341	"preference (Site-Level Preference for VPLS Multihoming for FEC 129)" on page 13049	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols vpls multi-homing site {site-name} active-interface]		
"any (VPLS Multihoming)" on page 2282	"primary (VPLS Multihoming)" on page 13151	

[edit logical-systems {logical-system-name} routing-instances {instance-name} protocols vpls multi-homing site {site-name} interface {interface-name}]	
"preference (Interface-Level Preference for VPLS Multihoming for FEC 129)" on page 13033	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls neighbor {address}]		
"backup-neighbor" on page 2802	"ignore-encapsulation-mismatch" on page 8397	"revert-time (Protocols Layer 2 Circuits)" on page 15729
"static (Protocols VPLS)" on page 28564	"switchover-delay" on page 28759	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls neighbor {address} backup-neighbor {address}]		
"community (Protocols Layer 2 Circuit)" on page 4469	"hot-standby (Protocols Layer 2 Circuit)" on page 8189	"psn-tunnel-endpoint" on page 13851

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls neighbor {address} backup-neighbor {address}]		
"standby (Protocols Layer 2 Circuit)" on page 28489	"static (Protocols VPLS)" on page 28564	

[edit logical-systems {name} routing-instances {instance_name} protocols vpls neighbor {neighbor-id} oam]		
"bfd-liveness-detection (Layer 2 VPN and VPLS)" on page 2929	"ping-interval" on page 12545	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls neighbor {address} static]		
"control-word" on page 4644	"no-control-word " on page 11426	

[edit logical-systems {name} routing-instances {instance_name} protocols vpls oam]		
"bfd-liveness-detection (Layer 2 VPN and VPLS)" on page 2929	"control-channel (Protocols OAM)" on page 4623	"ping-interval" on page 12545

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls site {site-name}]		
"automatic-site-id" on page 2758	"best-site" on page 2908	"interface (VPLS Routing Instances)" on page 9104
"mesh-group (Protocols VPLS)" on page 10641	"multi-homing (VPLS Multihoming for BGP VPLS)" on page 11051	"site-identifier (VPLS)" on page 27985
"site-preference" on page 27989		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vpls site {site-name} interfaces {interface-name}]	
"interface-mac-limit (VPLS)" on page 8913	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vstp]	
"access-trunk" on page 1705	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vstp vlan {vlan-id}]		
"backup-bridge-priority" on page 2793	"bridge-priority" on page 3036	"forward-delay" on page 7393
"hello-time" on page 8041	"interface (Spanning Tree)" on page 9057	"max-age" on page 10418
"system-id" on page 28847	"vpls-flush-on-topology-change" on page 30926	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} protocols vstp vlan {vlan-id} interface {interface-name}]		
"cost" on page 4681	"edge" on page 6197	"mode" on page 10798
"no-root-port" on page 11572	"priority (Protocols STP)" on page 13199	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel]		
"ingress-replication" on page 8647	"ldp-p2mp" on page 9685	"pim-asm" on page 12514
"rsvp-te (Routing Instances Provider Tunnel Selective)" on page 16061	"rsvp-te (Routing Instances Provider Tunnel)" on page 16059	"selective" on page 16529

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel family {inet | inet6}]		
"mdt" on page 10590	"pim-ssm (Provider Tunnel)" on page 12519	"source (Routing Instances)" on page 28312

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel family {inet | inet6} mdt]	
"data-mdt-reuse" on page 4840	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel family {inet | inet6} pim-ssm]	
"group-address (Routing Instances Tunnel Group)" on page 7802	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel family {inet | inet6}mdt threshold group {group-address} source {source-address}]	
"rate (Routing Instances)" on page 14073	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel family {inet | inet6}mdt]		
"group-range (Data MDTs)" on page 7882	"threshold (Routing Instances)" on page 29338	"tunnel-limit (Routing Instances)" on page 30213

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel family {inet | inet6}pim-ssm]	
"tunnel-source" on page 30239	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel family {intet | inet6}mdt threshold]		
"group (Protocols PIM)" on page 7876		

[edit logical-systems {name} routing-instances {name} provider-tunnel ingress-replication]		
"transport-class" on page 30132	"unicast-tunnel-name-regular-expression" on page 30355	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel ingress-replication label-switched-path]		
"label-switched-path-template (Multicast)" on page 9596		

[edit logical-systems {name} routing-instances {name} provider-tunnel ldp-p2mp]		
"root-address" on page 15840		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel pim mdt]		
"data-mdt-reuse" on page 4840		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel pim-asm]		
"group-address (Routing Instances VPN)" on page 7804		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel pim-asm family inet]		
"group-address (Routing Instances VPN)" on page 7804		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel pim-asm family inet6]	
"group-address (Routing Instances VPN)" on page 7804	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel pim-ssm]	
"group-address (Routing Instances VPN)" on page 7804	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel pim-ssm family inet]	
"group-address (Routing Instances VPN)" on page 7804	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel pim-ssm family inet6]	
"group-address (Routing Instances VPN)" on page 7804	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel rsvp-te]		
"label-switched-path-template (Multicast)" on page 9596	"static-lsp" on page 28547	

[edit logical-systems {logical-system-name} routing-instances {instance-name} provider-tunnel selective]		
"cmcast-joins-limit-inet (MVPN Selective Tunnels)" on page 4376	"cmcast-joins-limit-inet6 (MVPN Selective Tunnels)" on page 4379	"group (Routing Instances)" on page 7889
"leaf-tunnel-limit-inet (MVPN Selective Tunnels)" on page 9691	"leaf-tunnel-limit-inet6 (MVPN Selective Tunnels)" on page 9693	"tunnel-limit (Routing Instances Provider Tunnel Selective)" on page 30215

(Continued)

[edit logical-systems {logical-system-name} routing-instances {instance-name} provider-tunnel selective]		
"wildcard-group-inet" on page 31115	"wildcard-group-inet6" on page 31117	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel selective group {address}]		
"source (Routing Instances Provider Tunnel Selective)" on page 28314	"wildcard-source (Selective Provider Tunnels) " on page 31120	

[edit logical-systems {logical-system-name} routing-instances {instance-name} provider-tunnel selective group {group-prefix} source {source-prefix}]		
"ldp-p2mp" on page 9685	"pim-ssm (Selective Tunnel)" on page 12520	"rsvp-te (Routing Instances Provider Tunnel Selective)" on page 16061
"threshold-rate" on page 29336		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel selective group {group-address} source {source-address} pim-ssm]		
"group-range (MBGP MVPN Tunnel)" on page 7883		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel selective group {address} source {source-address} rsvp-te]		
"label-switched-path-template (Multicast)" on page 9596	"static-lsp" on page 28547	

[edit logical-systems {logical-system-name} routing-instances {instance-name} provider-tunnel selective group {group-prefix} wildcard-source]		
"ldp-p2mp" on page 9685	"pim-ssm (Selective Tunnel)" on page 12520	"rsvp-te (Routing Instances Provider Tunnel Selective)" on page 16061
"threshold-rate" on page 29336		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel selective group {group-address} wildcard-source pim-ssm]	
"group-range (MBGP MVPN Tunnel)" on page 7883	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel selective group {address} wildcard-source rsvp-te]	
"static-lsp" on page 28547	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet]		
"wildcard-source (Selective Provider Tunnels) " on page 31120	"wildcard-source (Selective Provider Tunnels) " on page 31120	

[edit logical-systems {logical-system-name} routing-instances {instance-name} provider-tunnel selective wildcard-group-inet wildcard-source]		
"ldp-p2mp" on page 9685	"pim-ssm (Selective Tunnel)" on page 12520	"rsvp-te (Routing Instances Provider Tunnel Selective)" on page 16061
"threshold-rate" on page 29336		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet wildcard-source pim-ssm]		
"group-range (MBGP MVPN Tunnel)" on page 7883		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet wildcard-source rsvp-te]		
"static-lsp" on page 28547		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet6]		
"wildcard-source (Selective Provider Tunnels) " on page 31120	"wildcard-source (Selective Provider Tunnels) " on page 31120	

[edit logical-systems {logical-system-name} routing-instances {instance-name} provider-tunnel selective wildcard-group-inet6 wildcard-source]		
"ldp-p2mp" on page 9685	"pim-ssm (Selective Tunnel)" on page 12520	"rsvp-te (Routing Instances Provider Tunnel Selective)" on page 16061
"threshold-rate" on page 29336		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet6 wildcard-source pim-ssm]		
"group-range (MBGP MVPN Tunnel)" on page 7883		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet6 wildcard-source rsvp-te]		
"static-lsp" on page 28547		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id}]	
"area-range" on page 2453	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options]		
"access (Static Framed Routes)" on page 1703	"access-internal (Static Access-Internal Routes)" on page 1668	"aggregate (Routing)" on page 2070
"auto-bandwidth" on page 2682	"auto-export" on page 2712	"autonomous-system" on page 2764
"backup-selection (Protocols IS-IS)" on page 2814	"backup-selection (Protocols OSPF)" on page 2819	"dynamic-tunnels" on page 6161
"fate-sharing" on page 6952	"flex-algorithm" on page 7221	"flex-algorithm (Protocols OSPF)" on page 7223
"generate" on page 7672	"graceful-restart (Enabling Globally)" on page 7760	"instance-export" on page 8759
"instance-import" on page 8760	"interface (Host Fast Reroute)" on page 8897	"interface-routes" on page 9010
"label" on page 9574	"martians" on page 10356	"maximum-paths" on page 10542
"maximum-prefixes" on page 10554	"multicast (routing-options)" on page 11063	"multipath (Routing Options)" on page 11113
"options (Routing Options)" on page 11922	"resolution" on page 15594	"rib (General)" on page 15764
"rib (Multitopology Routing)" on page 15785	"route-distinguisher-id-use-router-id" on page 15889	"router-id" on page 15902

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options]		
"routing-options-max-ecmp" on page 15985	"source-packet-routing (routing-options)" on page 28264	"static (Routing Options)" on page 28570
"topologies (Multitopology Routing)" on page 29436	"traceoptions (routing-options)" on page 29462	"validation (Origin Validation for BGP)" on page 30647

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options (aggregate | generate | static) (defaults | route)]		
"active" on page 1800	"as-path (Routing Options)" on page 2495	"color" on page 4425
"community (Routing Options)" on page 4472	"preference (Routing Options)" on page 13046	"tag (Routing Options)" on page 28924

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options (aggregate | generate) (defaults | route)]		
"brief" on page 3039	"discard (routing-options)" on page 5816	"policy (Aggregate and Generated Routes)" on page 12670

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options autonomous-system {autonomous-system}]	
"independent-domain" on page 8543	

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options backup-selection]	
"destination (Protocols OSPF)" on page 5511	

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options backup-selection destination {prefix}]	
"interface (Backup Selection OSPF)" on page 8837	

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options backup-selection destination {prefix} interface {interface-name}]	
"admin-group (Backup Selection)" on page 1965	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options dynamic-tunnels {tunnel-name}]		
"destination-networks" on page 5451	"forwarding-rib" on page 7483	"source-address (Dynamic Tunnel)" on page 28117
"source-address (Routing Options)" on page 28156	"tunnel-attributes" on page 30194	"v4ov6" on page 30636

[edit logical-systems {name} routing-instances {name} routing-options dynamic-tunnels tunnel-attributes]	
"dynamic-tunnel-reassembly" on page 6159	

[edit logical-systems {name} routing-instances {name} routing-options fate-sharing group {group-name}]	
"use-for-post-convergence-lfa" on page 30517	

[edit logical-systems {name} routing-instances {name} routing-options flex-algorithm]		
"definition" on page 5231	"definition (Protocols OSPF)" on page 5234	"use-transport-class" on page 30541

[edit logical-systems {name} routing-instances {name} routing-options flow]		
"fast-lookup-filter (Protocols BGP)" on page 6947	"firewall-install-disable" on page 7194	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options forwarding-table]	
"export (Routing Options)" on page 6691	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options graceful-restart]		
"disable (Protocols LDP)" on page 5789	"disable (Routing Options)" on page 5804	"restart-duration (Routing Options)" on page 15655

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options interface {interface-name}]		
"arp-prefix-limit (Host Fast Reroute)" on page 2477	"link-protection (Host Fast Reroute)" on page 9808	"supplementary-blackout-timer (Host Fast Reroute)" on page 28724

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options interface-routes]	
"rib-group (Routing Options)" on page 15778	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast]		
"asm-override-ssm" on page 2497	"backup-pe-group" on page 2808	"cont-stats-collection-interval" on page 4573
"flow-map" on page 7322	"interface (Multicast Static Routes)" on page 8928	"interface (Routing Options)" on page 9014

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast]		
"pim-to-mld-proxy" on page 12523	"rpf-check-policy (Routing Options RPF)" on page 16012	"scope" on page 16379
"ssm-groups" on page 28452	"ssm-map (Routing Options Multicast)" on page 28468	"stream-protection (Multicast-Only Fast Reroute)" on page 28634
"traceoptions (routing-options)" on page 29462		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast backup-pe-group {group-name}]		
"backups" on page 2826	"local-address (Routing Options)" on page 9913	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast flow-map]		
"bandwidth (multicast)" on page 2831	"forwarding-cache (Flow Maps)" on page 7408	"policy (Flow Maps)" on page 12683
"redundant-sources" on page 14219	"timeout (Flow Maps)" on page 29377	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast forwarding-cache]		
"allow-maximum (Multicast)" on page 2193	"threshold (Multicast Forwarding Cache)" on page 29329	"timeout (Multicast)" on page 29380

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast forwarding-cache family (inet | inet6)]	
"threshold (Multicast Forwarding Cache)" on page 29329	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast forwarding-cache family (inet | inet6)threshold]	
"log-warning (Multicast Forwarding Cache)" on page 10047	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast forwarding-cache threshold]	
"log-warning (Multicast Forwarding Cache)" on page 10047	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast interface {interface-name}]		
"maximum-bandwidth" on page 10502	"no-qos-adjust" on page 11554	"reverse-oif-mapping" on page 15720
"subscriber-leave-timer" on page 28694		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast interface {interface-name} reverse-oif-mapping]	
"no-qos-adjust" on page 11554	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast pim-to-igmp-proxy]	
"upstream-interface" on page 30474	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast pim-to-mld-proxy]		
"upstream-interface" on page 30474		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast scope {scope-name}]		
"interface (Multicast Scoping)" on page 8927	"interface (Scoping)" on page 9015	"prefix" on page 13062

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast ssm-map {ssm-map-name}]		
"policy (SSM Maps)" on page 12733	"source (Source-Specific Multicast)" on page 28327	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multicast stream-protection]		
"mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain)" on page 10827	"mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain)" on page 10829	"mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain)" on page 10830
"mofrr-primary-path-selection-by-routing (Multicast-Only Fast Reroute)" on page 10832	"policy (Multicast-Only Fast Reroute)" on page 12691	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options multipath]		
"as-path-compare" on page 2490	"vpn-unequal-cost" on page 30951	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options outbound-route-filter]		
"bgp-orf-cisco-mode" on page 2963		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options resolution]		
"rib (Route Resolution)" on page 15789		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options resolution rib]		
"import (Routing Options)" on page 8465	"resolution-ribs" on page 15599	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options rib {routing-table-name}]		
"aggregate (Routing)" on page 2070	"filter (routing-options)" on page 7122	"generate" on page 7672
"martians" on page 10356	"multipath (Routing Options)" on page 11113	"policy-multipath" on page 12694
"vpn-global-import" on page 30953		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options rib {routing-table-name} (aggregate | generate | static) (defaults | route)]		
"active" on page 1800	"as-path (Routing Options)" on page 2495	"color" on page 4425
"community (Routing Options)" on page 4472	"preference (Routing Options)" on page 13046	"tag (Routing Options)" on page 28924

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options rib {routing-table-name} (aggregate | generate) (defaults | route)]		
"brief" on page 3039	"discard (routing-options)" on page 5816	"policy (Aggregate and Generated Routes)" on page 12670

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options rib {routing-table-name} multipath]		
"vpn-unequal-cost" on page 30951		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options rib {routing-table-name} static (defaults | route)]		
"install (Routing Options)" on page 8747	"readvertise" on page 14090	"resolve" on page 15600
"retain" on page 15672		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options rib {routing-table-name} static route {destination-prefix}]	
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options rib {routing-table-name} static route {destination-prefix} qualified-next-hop ({interface-name} | {address})]	
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options rib-groups {group-name}]		
"export-rib" on page 6688	"import-policy" on page 8473	"import-rib" on page 8488

[edit logical-systems {name} routing-instances {name} routing-options source-packet-routing]	
"mapping-server-entry" on page 10346	

[edit logical-systems {name} routing-instances {name} routing-options source-packet-routing mapping-server-entry]		
"prefix-segment (Routing Options)" on page 13109	"prefix-segment-range" on page 13107	

[edit logical-systems {name} routing-instances {name} routing-options source-packet-routing srv6]		
"block" on page 2980	"locator" on page 9997	

[edit logical-systems {name} routing-instances {name} routing-options source-packet-routing srv6 locator]	
"micro-sid" on page 10697	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options static (defaults | route)]		
"install (Routing Options)" on page 8747	"readvertise" on page 14090	"resolve" on page 15600
"retain" on page 15672		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options static route {destination-prefix}]		
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	"lsp-next-hop (Static Routes)" on page 10192	"p2mp-ldp-next-hop" on page 12119
"p2mp-lsp-next-hop" on page 12120	"qualified-next-hop (Static Routes)" on page 13910	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options static route {destination-prefix} qualified-next-hop ({interface-name} | {address})]	
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} routing-options topologies family (inet | inet6)]	
"topology (Multitopology Routing)" on page 29443	

[edit logical-systems {name} routing-instances {name} routing-options transport-class]	
"name (Transport-class)" on page 11181	

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options validation]		
"group (Origin Validation for BGP)" on page 7850	"static (Origin Validation for BGP)" on page 28553	"traceoptions (Origin Validation for BGP)" on page 29626

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options validation group {group-name}]	
"session (Origin Validation for BGP)" on page 16836	

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options validation group {group-name} session {server-ip-address}]		
"port (Origin Validation for BGP)" on page 12846	"traceoptions (Origin Validation for BGP)" on page 29626	

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options validation group {group-name} session {address} traceoptions]		
"file (Tracing for Origin AS Validation)" on page 7091	"flag (Tracing for Origin AS Validation)" on page 7212	

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options validation static]	
"record (Origin Validation for BGP)" on page 14138	

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options validation static record {destination}]	
"maximum-length (Origin Validation for BGP)" on page 10531	

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options validation static record {destination} maximum-length {prefix-length}]	
"origin-autonomous-system (Origin Validation for BGP)" on page 11932	

[edit logical-systems {logical-system-name} routing-instances {instance-name} routing-options validation traceoptions]		
"file (Tracing for Origin AS Validation)" on page 7091	"flag (Tracing for Origin AS Validation)" on page 7212	

[edit logical-systems {logical-system-name} routing-instances r{outing-instance-name} protocols bgp group {group-name} neighbor {neighbor-address}]	
"tcp-aggressive-transmission" on page 28978	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} statement-name protocols igmp amt]	
"relay (IGMP)" on page 14292	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} statement-name protocols igmp amt relay]	
"defaults" on page 5224	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} switch-options]		
"interface-mac-limit" on page 8910	"mac-statistics" on page 10263	"mac-table-size" on page 10271

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} switch-options interface {interface-name}]		
"interface-mac-limit" on page 8910	"mac-pinning" on page 10254	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} switch-options interface {interface-name} interface-mac-limit {limit}]	
"packet-action" on page 12129	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} switch-options interface-mac-limit {limit}]	
"packet-action" on page 12129	

[edit logical-systems {logical-system-name} routing-instances {routing-instances-name} system processes static-subscribers]	
"traceoptions (Static Subscribers)" on page 29873	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services]		
"dhcp-local-server" on page 5601	"static-subscribers (Dynamic Service Provisioning)" on page 28582	"tcp-forwarding (Remote Device Management)" on page 28987

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services ...]		
"allow-active-leasequery (DHCP Local Server)" on page 2163	"allow-bulk-leasequery (DHCP Local Server)" on page 2165	"allow-leasequery (DHCP Local Server)" on page 2191

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server]		
"dhcpv6 (DHCP Local Server)" on page 5658	"dual-stack-group (DHCP Local Server)" on page 6052	"duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)" on page 6066
"forward-snooped-clients (DHCP Local Server)" on page 7404	"pool-match-order" on page 12777	"reauthenticate (DHCP Local Server)" on page 14103
"reconfigure (DHCP Local Server)" on page 14133		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server (dhcpv6) ...overrides]	
"include-option-82 (DHCP Local Server)" on page 8534	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server ...]		
"authentication (DHCP Local Server)" on page 2593	"dual-stack (DHCP Local Server Overrides)" on page 6049	"dynamic-profile (DHCP Local Server)" on page 6117

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server ...]		
"group (DHCP Local Server)" on page 7814	"interface (DHCP Local Server)" on page 8860	"interface-description (DHCP Local Server)" on page 8854
"interface-name (DHCP Local Server)" on page 8937	"logical-system-name (DHCP Local Server)" on page 10082	"overrides (DHCP Local Server)" on page 12065
"requested-ip-network-match (DHCP Local Server)" on page 15572	"service-profile (DHCP Local Server)" on page 16729	"username-include (DHCP Local Server)" on page 30595

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server ... overrides]	
"client-discover-match (DHCP Local Server)" on page 4308	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server authentication]		
"password (DHCP Local Server)" on page 12226	"password (DHCP Local Server)" on page 12226	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"circuit-type (DHCP Local Server)" on page 3330	"delimiter (DHCP Local Server)" on page 5288
"delimiter (DHCP Local Server)" on page 5288	"domain-name (DHCP Local Server)" on page 5884	"domain-name (DHCP Local Server)" on page 5884
"mac-address (DHCP Local Server)" on page 10221	"option-60 (DHCP Local Server)" on page 11876	"option-60 (DHCP Local Server)" on page 11876

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server authentication username-include]		
"option-82 (DHCP Local Server Authentication)" on page 11890	"option-82 (DHCP Local Server Authentication)" on page 11890	"routing-instance-name (DHCP Local Server)" on page 15945
"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578	"user-prefix (DHCP Local Server)" on page 30578

[edit logical-systems {name} routing-instances {name} system services dhcp-local-server dhcp-local-server...]	
"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982	

[edit logical-systems {name} routing-instances {name} system services dhcp-local-server dhcpv6]		
"dynamic-server" on page 6154	"reauthenticate (DHCP Local Server)" on page 14103	"reconfigure (DHCP Local Server)" on page 14133

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 ...]		
"delegated-pool (DHCP Local Server)" on page 5271	"multi-address-embedded-option-response (DHCP Local Server)" on page 11039	"rapid-commit (DHCPv6 Local Server)" on page 14055

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 authentication]		
"password (DHCP Local Server)" on page 12226	"password (DHCP Local Server)" on page 12226	"user-prefix (DHCP Local Server)" on page 30578

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"circuit-type (DHCP Local Server)" on page 3330	"client-id (DHCP Local Server)" on page 4315
"client-id (DHCP Local Server)" on page 4315	"delimiter (DHCP Local Server)" on page 5288	"delimiter (DHCP Local Server)" on page 5288
"domain-name (DHCP Local Server)" on page 5884	"domain-name (DHCP Local Server)" on page 5884	"relay-agent-interface-id (DHCP Local Server)" on page 14275
"relay-agent-interface-id (DHCP Local Server)" on page 14275	"relay-agent-remote-id (DHCP Local Server)" on page 14282	"relay-agent-remote-id (DHCP Local Server)" on page 14282
"relay-agent-subscriber-id (DHCP Local Server)" on page 14288	"relay-agent-subscriber-id (DHCP Local Server)" on page 14288	"routing-instance-name (DHCP Local Server)" on page 15945
"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name}]	
"reconfigure (DHCP Local Server)" on page 14133	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} authentication]		
"password (DHCP Local Server)" on page 12226	"password (DHCP Local Server)" on page 12226	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"circuit-type (DHCP Local Server)" on page 3330	"client-id (DHCP Local Server)" on page 4315
"client-id (DHCP Local Server)" on page 4315	"delimiter (DHCP Local Server)" on page 5288	"delimiter (DHCP Local Server)" on page 5288
"domain-name (DHCP Local Server)" on page 5884	"domain-name (DHCP Local Server)" on page 5884	"relay-agent-interface-id (DHCP Local Server)" on page 14275
"relay-agent-interface-id (DHCP Local Server)" on page 14275	"relay-agent-remote-id (DHCP Local Server)" on page 14282	"relay-agent-remote-id (DHCP Local Server)" on page 14282
"relay-agent-subscriber-id (DHCP Local Server)" on page 14288	"relay-agent-subscriber-id (DHCP Local Server)" on page 14288	"routing-instance-name (DHCP Local Server)" on page 15945
"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578	"user-prefix (DHCP Local Server)" on page 30578

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name}]		
"trace (DHCP Local Server)" on page 29456	"trace (DHCP Local Server)" on page 29456	

[edit logical-systems {name} routing-instances {name} system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name} overrides]		
"delay-advertise (DHCPv6)" on page 5239	"process-inform" on page 13240	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit logical-systems {name} routing-instances {name} system services dhcp-local-server dhcpv6 group {group-name} overrides]		
"delay-advertise (DHCPv6)" on page 5239	"interface-client-limit (DHCP Local Server)" on page 9106	"interface-client-limit (DHCP Local Server)" on page 9106
"process-inform" on page 13240		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"strict (DHCP Local Server)" on page 28646
"timeout (DHCP Local Server)" on page 29374	"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} reconfigure trigger]	
"radius-disconnect (DHCP Local Server)" on page 14003	

[edit logical-systems {name} routing-instances {name} system services dhcp-local-server dhcpv6 overrides]		
"delay-advertise (DHCPv6)" on page 5239	"interface-client-limit (DHCP Local Server)" on page 9106	"interface-client-limit (DHCP Local Server)" on page 9106
"process-inform" on page 13240		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"strict (DHCP Local Server)" on page 28646
"timeout (DHCP Local Server)" on page 29374	"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 reconfigure trigger]	
"radius-disconnect (DHCP Local Server)" on page 14003	

[edit logical-systems {name} routing-instances {name} system services dhcp-local-server dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Local Server)" on page 3375	"classification-key (DHCP Relay Agent)" on page 3376	"dual-stack-interface-client-limit (DHCP Local Server and Relay Agent)" on page 6058

(Continued)

[edit logical-systems {name} routing-instances {name} system services dhcp-local-server dual-stack-group {dual-stack-group-name}]		
"on-demand-address-allocation" on page 11813	"protocol-master" on page 13448	"reauthenticate (DHCP Local Server)" on page 14103

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Local Server)" on page 2056	"use-primary (DHCP Local Server)" on page 30535	

[edit logical-systems {name} routing-instances {name} system services dhcp-local-server group]		
"interface-tag (DHCP Local Server)" on page 9081	"reconfigure (DHCP Local Server)" on page 14133	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} authentication]		
"password (DHCP Local Server)" on page 12226	"password (DHCP Local Server)" on page 12226	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"circuit-type (DHCP Local Server)" on page 3330	"delimiter (DHCP Local Server)" on page 5288
"delimiter (DHCP Local Server)" on page 5288	"domain-name (DHCP Local Server)" on page 5884	"domain-name (DHCP Local Server)" on page 5884
"mac-address (DHCP Local Server)" on page 10221	"option-60 (DHCP Local Server)" on page 11876	"option-60 (DHCP Local Server)" on page 11876

(Continued)

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} authentication username-include]		
"option-82 (DHCP Local Server Authentication)" on page 11890	"option-82 (DHCP Local Server Authentication)" on page 11890	"routing-instance-name (DHCP Local Server)" on page 15945
"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578	"user-prefix (DHCP Local Server)" on page 30578

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Local Server)" on page 2056	"use-primary (DHCP Local Server)" on page 30535	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} interface {interface-name}]		
"trace (DHCP Local Server)" on page 29456	"trace (DHCP Local Server)" on page 29456	

[edit logical-systems {name} routing-instances {name} system services dhcp-local-server group {group-name} interface {interface-name} overrides]		
"delay-offer (DHCPv4)" on page 5254	"process-inform" on page 13240	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} interface {interface-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit logical-systems {name} routing-instances {name} system services dhcp-local-server group {group-name} overrides]		
"delay-offer (DHCPv4)" on page 5254	"interface-client-limit (DHCP Local Server)" on page 9106	"interface-client-limit (DHCP Local Server)" on page 9106
"process-inform" on page 13240		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"timeout (DHCP Local Server)" on page 29374
"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} reconfigure trigger]	
"radius-disconnect (DHCP Local Server)" on page 14003	

[edit logical-systems {name} routing-instances {name} system services dhcp-local-server overrides]		
"delay-offer (DHCPv4)" on page 5254	"interface-client-limit (DHCP Local Server)" on page 9106	"interface-client-limit (DHCP Local Server)" on page 9106
"process-inform" on page 13240		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server overrides process-inform]		
"pool (DHCP Local Server Overrides)" on page 12774		

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server pool-match-order]		
"external-authority" on page 6726	"ip-address-first" on page 9235	"option-82 (DHCP Local Server Pool Matching)" on page 11892

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"timeout (DHCP Local Server)" on page 29374
"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} system services dhcp-local-server reconfigure trigger]		
"radius-disconnect (DHCP Local Server)" on page 14003		

[edit logical-systems {logical-system-name} routing-instances {routing-instances-name} system services static-subscribers]		
"access-profile (Static Subscribers)" on page 1698	"authentication (Static Subscribers)" on page 2665	"group (Static Subscribers)" on page 7897

[edit logical-systems {logical-system-name} routing-instances {routing-instances-name} system services static-subscribers authentication]		
"password (Static Subscribers)" on page 12247	"username-include (Static Subscribers)" on page 30604	

[edit logical-systems {logical-system-name} routing-instances {routing-instances-name} system services static-subscribers authentication username-include]		
"domain-name (Static Subscribers)" on page 5892	"interface (Static Subscriber Username)" on page 9070	"logical-system-name (Static Subscribers)" on page 10085
"routing-instance-name (Static Subscribers)" on page 15949	"user-prefix (Static Subscribers)" on page 30582	

[edit logical-systems {logical-system-name} routing-instances {routing-instances-name} system services static-subscribers dynamic-profile {profile-name}]	
"aggregate-clients (Static Subscribers)" on page 2060	

[edit logical-systems {logical-system-name} routing-instances {routing-instances-name} system services static-subscribers group {group-name}]		
"access-profile (Static Subscribers)" on page 1698	"authentication (Static Subscribers)" on page 2665	"interface (Static Subscriber Group)" on page 9068

[edit logical-systems {logical-system-name} routing-instances {routing-instances-name} system services static-subscribers group {group-name} authentication]		
"password (Static Subscribers)" on page 12247	"username-include (Static Subscribers)" on page 30604	

[edit logical-systems {logical-system-name} routing-instances {routing-instances-name} system services static-subscribers group {group-name} authentication username-include]		
"domain-name (Static Subscribers)" on page 5892	"interface (Static Subscriber Username)" on page 9070	"logical-system-name (Static Subscribers)" on page 10085
"routing-instance-name (Static Subscribers)" on page 15949	"user-prefix (Static Subscribers)" on page 30582	

[edit logical-systems {logical-system-name} routing-instances {routing-instances-name} system services static-subscribers group {group-name} dynamic-profile {profile-name}]	
"aggregate-clients (Static Subscribers)" on page 2060	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} vlans {vlan-name}bridge-domains {bridge-name} bridge-options-interface {interface-name}]	
"mac-pinning (EVPN Routing Instances)" on page 10256	

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay dhcpv6 group {name} ddual-stack-group {dual-stack-group-name}]	
"classification-key (DHCP Relay Agent)" on page 3376	

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group {dual-stack-group-name}]	
"protocol-master" on page 13448	

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay dhcpv6 group {name} relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay dhcpv6 relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay group {name} relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} vlans {vlan-name} switch-options]	
"interface-mac-ip-limit" on page 8908	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} vlans {vlan-name}routing-instances {routing-instance-name} switch-options interface {interface-name}]	
"mac-pinning (EVPN Routing Instances)" on page 10256	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} vrf-advertise-selective]	
"family (VRF Advertisement)" on page 6921	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name} vrf-advertise-selective family]		
"inet-mvpn (VRF Advertisement)" on page 8575	"inet6-mvpn (VRF Advertisement)" on page 8612	

[edit logical-systems {logical-system-name} routing-instances {routing-instance-name}...]		
"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279	"duplicate-clients (DHCPv6 Local Server and Relay Agent)" on page 6064	"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699
"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701	"no-snoop (DHCP Local Server and Relay Agent)" on page 11582	"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334
"route-suppression (DHCP Local Server and Relay Agent)" on page 15877	"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	

[edit logical-systems {logical-system -name} routing-instances {routing-instance-name}igmp-snooping-options]	
"snoop-pseudowires" on page 28075	

[routing-instancesrouting-options]

[edit logical-systems {name} routing-instances{name} routing-options bmp route-monitoring]	
"loc-rib" on page 9897	

[routing-options]

[edit logical-systems {logical-system-name} routing-options]		
"access (Static Framed Routes)" on page 1703	"access-internal (Static Access-Internal Routes)" on page 1668	"admin-groups-extended" on page 1975
"admin-groups-extended-range" on page 1977	"aggregate (Routing)" on page 2070	"auto-bandwidth" on page 2682
"auto-export" on page 2712	"autonomous-system" on page 2764	"backup-selection (Protocols IS-IS)" on page 2814
"backup-selection (Protocols OSPF)" on page 2819	"bgp-static" on page 2966	"bmp" on page 3000
"confederation" on page 4508	"dynamic-tunnels" on page 6161	"fate-sharing" on page 6952
"flex-algorithm" on page 7221	"flex-algorithm (Protocols OSPF)" on page 7223	"forwarding-table" on page 7490

(Continued)

[edit logical-systems {logical-system-name} routing-options]		
"graceful-restart (Enabling Globally)" on page 7760	"graceful-restart (Protocols RSVP)" on page 7782	"host-fast-reroute" on page 8141
"instance-export" on page 8759	"instance-import" on page 8760	"interface-routes" on page 9010
"krt-nexthop-ack" on page 9505	"martians" on page 10356	"maximum-paths" on page 10542
"maximum-prefixes" on page 10554	"multicast (routing-options)" on page 11063	"no-bfd-triggered-local-repair" on page 11414
"options (Routing Options)" on page 11922	"programmable-rpd" on page 13369	"resolution" on page 15594
"rib (General)" on page 15764	"rib (Multitopology Routing)" on page 15785	"rib-groups" on page 15781
"route-distinguisher-id" on page 15860	"route-distinguisher-id-use-router-id" on page 15889	"route-record" on page 15874
"router-id" on page 15902	"routing-options-max-ecmp" on page 15985	"source-packet-routing (routing-options)" on page 28264
"srlg" on page 28411	"static (Routing Options)" on page 28570	"topologies (Multitopology Routing)" on page 29436
"traceoptions (routing-options)" on page 29462	"validation (Origin Validation for BGP)" on page 30647	

[edit logical-systems {logical-system-name} routing-options (aggregate | generate | static) (defaults | route)]		
"active" on page 1800	"as-path (Routing Options)" on page 2495	"color" on page 4425

(Continued)

[edit logical-systems {logical-system-name} routing-options (aggregate | generate | static) (defaults | route)]		
"community (Routing Options)" on page 4472	"metric (Aggregate, Generated, or Static Route)" on page 10667	"preference (Routing Options)" on page 13046
"tag (Routing Options)" on page 28924		

[edit logical-systems {logical-system-name} routing-options (aggregate | generate) (defaults | route)]		
"brief" on page 3039	"discard (routing-options)" on page 5816	"policy (Aggregate and Generated Routes)" on page 12670

[edit logical-systems {logical-system-name} routing-options autonomous-system {as-number}]	
"loops (Autonomous System)" on page 10132	

[edit logical-systems {logical-system-name} routing-options backup-selection]	
"destination (Protocols OSPF)" on page 5511	

[edit logical-systems {logical-system-name} routing-options backup-selection destination {prefix}]	
"interface (Backup Selection OSPF)" on page 8837	

[edit logical-systems {logical-system-name} routing-options backup-selection destination {prefix} interface {interface name}]	
"admin-group (Backup Selection)" on page 1965	

[edit logical-systems {logical-system-name} routing-options bmp]		
"authentication-algorithm" on page 2577	"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605
"hold-down" on page 8086	"route-monitoring" on page 15871	"traceoptions (Protocols BMP)" on page 29666

[edit logical-systems {name} routing-options bmp route-monitoring]	
"loc-rib" on page 9897	

[edit logical-systems {logical-system-name} routing-options bmp station {station-name}]		
"authentication-algorithm" on page 2577	"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605
"hold-down" on page 8086	"route-monitoring" on page 15871	"traceoptions (Protocols BMP)" on page 29666

[edit logical-systems {name} routing-options client {id}]	
"interface-notification (programmable-rpd)" on page 8942	

[edit logical-systems {logical-system-name} routing-options dynamic-tunnels {tunnel-name}]		
"destination-networks" on page 5451	"forwarding-rib" on page 7483	"gre (Routing Options)" on page 7797
"rsvp-te (Routing Options)" on page 16062	"source-address (Dynamic Tunnel)" on page 28117	"source-address (Routing Options)" on page 28156
"tunnel-attributes" on page 30194	"v4ov6" on page 30636	

[edit logical-systems {logical-system-name} routing-options dynamic-tunnels gre next-hop-based-tunnel]		
"next-hop-based-tunnel (Routing Options)" on page 11332		

[edit logical-systems {logical-system-name} routing-options dynamic-tunnels {tunnel-name} rsvp-te {entry}]		
"destination-networks" on page 5451	"label-switched-path-template (Multicast)" on page 9596	

[edit logical-systems {name} routing-options dynamic-tunnels tunnel-attributes]		
"dynamic-tunnel-reassembly" on page 6159		

[edit logical-systems {name} routing-options fate-sharing group {group-name}]		
"use-for-post-convergence-lfa" on page 30517		

[edit logical-systems {name} routing-options flex-algorithm]		
"definition" on page 5231	"definition (Protocols OSPF)" on page 5234	"use-transport-class" on page 30541

[edit logical-systems {name} routing-options flow]		
"fast-lookup-filter (Protocols BGP)" on page 6947	"firewall-install-disable" on page 7194	

[edit logical-systems {logical-system-name} routing-options forwarding-table]		
"chained-composite-next-hop" on page 3277	"ecmp-fast-reroute" on page 6193	"evpn-aliasing-optimize" on page 6547

(Continued)

[edit logical-systems {logical-system-name} routing-options forwarding-table]		
"export (Routing Options)" on page 6691	"indirect-next-hop" on page 8549	"indirect-next-hop-change-acknowledgements" on page 8551
"krt-nexthop-ack-timeout" on page 9506	"unicast-reverse-path" on page 30348	

[edit logical-systems {logical-system-name} routing-options forwarding-table chained-composite-next-hop]		
"ingress (Chained Composite Next Hop)" on page 8628	"transit (Chained Composite Next Hops)" on page 30066	

[edit logical-systems {logical-system-name} routing-options forwarding-table chained-composite-next-hop ingress]		
"l2ckt" on page 9535	"l2vpn (routing-options)" on page 9564	"l3vpn (ingress)" on page 9571
"labeled-bgp" on page 9600	"no-l2ckt" on page 11491	"no-l2vpn" on page 11492

[edit logical-systems {logical-system-name} routing-options forwarding-table chained-composite-next-hop ingress l3vpn]	
"extended-space" on page 6703	

[edit logical-systems {logical-system-name} routing-options forwarding-table chained-composite-next-hop ingress labeled-bgp]		
"inet" on page 8554	"inet6" on page 8595	

[edit logical-systems {logical-system-name} routing-options graceful-restart]		
"disable (Routing Options)" on page 5804	"restart-duration (Routing Options)" on page 15655	

[edit logical-systems {logical-system-name} routing-options host-fast-reroute]		
"global-arp-prefix-limit (Host Fast Reroute)" on page 7698	"global-supplementary-blackout-timer (Host Fast Reroute)" on page 7737	

[edit logical-systems {logical-system-name} routing-options interface-routes]	
"rib-group (Routing Options)" on page 15778	

[edit logical-systems {logical-system-name} routing-options multicast]		
"asm-override-ssm" on page 2497	"backup-pe-group" on page 2808	"cont-stats-collection-interval" on page 4573
"flow-map" on page 7322	"interface (Multicast Static Routes)" on page 8928	"interface (Routing Options)" on page 9014
"pim-to-igmp-proxy" on page 12522	"pim-to-mld-proxy" on page 12523	"rpf-check-policy (Routing Options RPF)" on page 16012
"rpf-check-policy (Routing Options)" on page 16011	"scope" on page 16379	"scope-policy" on page 16380
"ssm-groups" on page 28452	"ssm-map (Routing Options Multicast)" on page 28468	"stream-protection (Multicast-Only Fast Reroute)" on page 28634
"traceoptions (routing-options)" on page 29462		

[edit logical-systems {logical-system-name} routing-options multicast backup-pe-group {group-name}]		
"backups" on page 2826	"local-address (Routing Options)" on page 9913	

[edit logical-systems {logical-system-name} routing-options multicast flow-map]		
"bandwidth (multicast)" on page 2831	"forwarding-cache (Flow Maps)" on page 7408	"policy (Flow Maps)" on page 12683
"redundant-sources" on page 14219	"timeout (Flow Maps)" on page 29377	

[edit logical-systems {logical-system-name} routing-options multicast forwarding-cache]		
"allow-maximum (Multicast)" on page 2193	"threshold (Multicast Forwarding Cache)" on page 29329	"timeout (Multicast)" on page 29380

[edit logical-systems {logical-system-name} routing-options multicast forwarding-cache family (inet | inet6)]	
"threshold (Multicast Forwarding Cache)" on page 29329	

[edit logical-systems {logical-system-name} routing-options multicast forwarding-cache family (inet | inet6) threshold]	
"log-warning (Multicast Forwarding Cache)" on page 10047	

[edit logical-systems {logical-system-name} routing-options multicast forwarding-cache threshold]	
"log-warning (Multicast Forwarding Cache)" on page 10047	

[edit logical-systems {logical-system-name} routing-options multicast interface {interface-name}]		
"maximum-bandwidth" on page 10502	"no-qos-adjust" on page 11554	"reverse-oif-mapping" on page 15720
"subscriber-leave-timer" on page 28694		

[edit logical-systems {logical-system-name} routing-options multicast interface {interface-name} reverse-oif-mapping]	
"no-qos-adjust" on page 11554	

[edit logical-systems {logical-system-name} routing-options multicast pim-to-igmp-proxy]	
"upstream-interface" on page 30474	

[edit logical-systems {logical-system-name} routing-options multicast pim-to-mld-proxy]	
"upstream-interface" on page 30474	

[edit logical-systems {logical-system-name} routing-options multicast scope {scope-name}]		
"interface (Multicast Scoping)" on page 8927	"interface (Scoping)" on page 9015	"prefix" on page 13062

[edit logical-systems {logical-system-name} routing-options multicast ssm-map {ssm-map-name}]		
"policy (SSM Maps)" on page 12733	"source (Source-Specific Multicast)" on page 28327	

[edit logical-systems {logical-system-name} routing-options multicast stream-protection]		
"mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain)" on page 10827	"mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain)" on page 10829	"mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain)" on page 10830
"mofrr-primary-path-selection-by-routing (Multicast-Only Fast Reroute)" on page 10832	"policy (Multicast-Only Fast Reroute)" on page 12691	

[edit logical-systems {logical-system-name} routing-options outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit logical-systems {logical-system-name} routing-options resolution]	
"rib (Route Resolution)" on page 15789	

[edit logical-systems {logical-system-name} routing-options resolution rib]		
"import (Routing Options)" on page 8465	"resolution-ribs" on page 15599	

[edit logical-systems {logical-system-name} routing-options rib {routing-table-name}]		
"aggregate (Routing)" on page 2070	"bgp-static" on page 2966	"filter (routing-options)" on page 7122
"martians" on page 10356	"policy-multipath" on page 12694	"static (Routing Options)" on page 28570
"vpn-global-import" on page 30953		

[edit logical-systems {logical-system-name} routing-options rib {routing-table-name} (aggregate | generate | static) (defaults | route)]		
"active" on page 1800	"as-path (Routing Options)" on page 2495	"color" on page 4425
"community (Routing Options)" on page 4472	"preference (Routing Options)" on page 13046	"tag (Routing Options)" on page 28924

[edit logical-systems {logical-system-name} routing-options rib {routing-table-name} (aggregate | generate) (defaults | route)]		
"brief" on page 3039	"discard (routing-options)" on page 5816	"policy (Aggregate and Generated Routes)" on page 12670

[edit logical-systems {logical-system-name} routing-options rib bgp.rtarget.0 static]	
"route-target-filter" on page 15879	

[edit logical-systems {logical-system-name} routing-options rib {routing-table-name} filter]	
"input (Routing Options RIB)" on page 8717	

[edit logical-systems {logical-system-name} routing-options rib inet6.0 static route {destination-prefix}]	
"qualified-next-hop (Static Routes)" on page 13910	

[edit logical-systems {logical-system-name} routing-options rib {routing-table-name} static]	
"rib-group (Routing Options)" on page 15778	

[edit logical-systems {logical-system-name} routing-options rib {routing-table-name} static (defaults | route)]		
"install (Routing Options)" on page 8747	"readvertise" on page 14090	"resolve" on page 15600
"retain" on page 15672		

[edit logical-systems {logical-system-name} routing-options rib {routing-table-name} static route {destination-prefix}]	
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	

[edit logical-systems {logical-system-name} routing-options rib {routing-table-name} static route {destination-prefix} qualified-next-hop ({interface-name} | {address})]	
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	

[edit logical-systems {logical-system-name} routing-options rib-groups {group-name}]		
"export-rib" on page 6688	"import-policy" on page 8473	"import-rib" on page 8488

[edit logical-systems {name} routing-options source-packet-routing]	
"mapping-server-entry" on page 10346	

[edit logical-systems {name} routing-options source-packet-routing mapping-server-entry]		
"prefix-segment (Routing Options)" on page 13109	"prefix-segment-range" on page 13107	

[edit logical-systems {name} routing-options source-packet-routing srv6]		
"block" on page 2980	"locator" on page 9997	

[edit logical-systems {name} routing-options source-packet-routing srv6 locator]	
"micro-sid" on page 10697	

[edit logical-systems {logical-system-name} routing-options srlg]		
"srlg-cost" on page 28413	"srlg-value" on page 28414	

[edit logical-systems {logical-system-name} routing-options static]
"rib-group (Routing Options)" on page 15778

[edit logical-systems {logical-system-name} routing-options static (defaults | route)]		
"install (Routing Options)" on page 8747	"readvertise" on page 14090	"resolve" on page 15600
"retain" on page 15672		

[edit logical-systems {logical-system-name} routing-options static route {destination-prefix}]		
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	"lsp-next-hop (Static Routes)" on page 10192	"p2mp-ldp-next-hop" on page 12119
"p2mp-lsp-next-hop" on page 12120	"qualified-next-hop (Static Routes)" on page 13910	

[edit logical-systems {logical-system-name} routing-options static route {destination-prefix} next-hop {address} resolve]	
"longest-match (Static Routes)" on page 10109	

[edit logical-systems {logical-system-name} routing-options static route {destination-prefix} qualified-next-hop]	
"metric (Qualified Next Hop on Static Route)" on page 10689	

[edit logical-systems {logical-system-name} routing-options static route {destination-prefix} qualified-next-hop ({interface-name} | {address})]	
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	

[edit logical-systems {logical-system-name} routing-options topologies family (inet | inet6)]	
"topology (Multitopology Routing)" on page 29443	

[edit logical-systems {name} routing-options transport-class]	
"name (Transport-class)" on page 11181	

[edit logical-systems {logical-system-name} routing-options validation]		
"group (Origin Validation for BGP)" on page 7850	"static (Origin Validation for BGP)" on page 28553	"traceoptions (Origin Validation for BGP)" on page 29626

[edit logical-systems {logical-system-name} routing-options validation group {group-name}]		
"session (Origin Validation for BGP)" on page 16836		

[edit logical-systems {logical-system-name} routing-options validation group {group-name} session {server-ip-address}]		
"port (Origin Validation for BGP)" on page 12846	"traceoptions (Origin Validation for BGP)" on page 29626	

[edit logical-systems {logical-system-name} routing-options validation group {group-name} session {address} traceoptions]		
"file (Tracing for Origin AS Validation)" on page 7091	"flag (Tracing for Origin AS Validation)" on page 7212	

[edit logical-systems {logical-system-name} routing-options validation static]		
"record (Origin Validation for BGP)" on page 14138		

[edit logical-systems {logical-system-name} routing-options validation static record {destination}]		
"maximum-length (Origin Validation for BGP)" on page 10531		

[edit logical-systems {logical-system-name} routing-options validation static record {destination} maximum-length {prefix-length}]		
"origin-autonomous-system (Origin Validation for BGP)" on page 11932		

[edit logical-systems {logical-system-name} routing-options validation traceoptions]		
"file (Tracing for Origin AS Validation)" on page 7091	"flag (Tracing for Origin AS Validation)" on page 7212	

[security]

[edit logical-systems {name} security]		
"alg" on page 2130	"dynamic-address" on page 6086	"idp (logical-systems)" on page 8313
"log (Logical Systems and Tenant Systems)" on page 10016	"log (Security)" on page 10026	"softwires" on page 28102
"traceoptions (Security Application Proxy)" on page 29778		

[edit logical-systems {name} security alg]		
"alg-support-lib" on page 2142	"dns" on page 5847	"ftp (Security ALG)" on page 7626
"ike-esp-nat" on page 8421	"mgcp" on page 10693	"msrpc" on page 10999
"pptp" on page 12986	"rsh" on page 16041	"rtsp" on page 16072
"sccp" on page 16341	"sip (Security) " on page 27977	"sunrpc" on page 28722
"talk" on page 28940	"tftp (Security ALG)" on page 29252	"traceoptions (Security ALG)" on page 29770

[edit logical-systems {name} security alg dns]		
"maximum-message-length" on page 10534	"oversize-message-drop" on page 12081	

[edit logical-systems {name} security alg h323]		
"application-screen (Security H323)" on page 2359	"endpoint-registration-timeout" on page 6342	"media-source-port-any" on page 10599
"traceoptions (Security H323 ALG)" on page 29801		

[edit logical-systems {name} security alg h323 application-screen]		
"message-flood (Security H323)" on page 10644	"unknown-message (Security H323 ALG)" on page 30410	

[edit logical-systems {name} security alg h323 application-screen message-flood]	
"gatekeeper" on page 7653	

[edit logical-systems {name} security alg mgcp]		
"application-screen (Security MGCP)" on page 2361	"inactive-media-timeout (Security MGCP)" on page 8496	"traceoptions (Security MGCP ALG)" on page 29814
"transaction-timeout" on page 30055		

[edit logical-systems {name} security alg mgcp application-screen]		
"connection-flood" on page 4531	"message-flood (Security MGCP)" on page 10645	"unknown-message (Security MGCP ALG)" on page 30411

[edit logical-systems {name} security alg sccp]		
"application-screen (Security SCCP)" on page 2362	"inactive-media-timeout (Security SCCP)" on page 8497	"traceoptions (Security SCCP ALG)" on page 29823

[edit logical-systems {name} security alg sccp application-screen]		
"call-flood" on page 3179	"unknown-message (Security SCCP ALG)" on page 30413	

[edit logical-systems {name} security alg sip]		
"application-screen (Security SIP)" on page 2364	"c-timeout" on page 3152	"t1-interval" on page 28885
"t4-interval" on page 28893	"traceoptions (Security SIP ALG)" on page 29830	

[edit logical-systems {name} security alg sip (Security)]		
"inactive-media-timeout (Security SIP)" on page 8499	"retain-hold-resource" on page 15674	

[edit logical-systems {name} security alg sip (Security) application-screen (Security SIP)]	
"unknown-message (Security SIP ALG)" on page 30415	

[edit logical-systems {name} security alg sip application-screen protect]	
"deny (Security SIP)" on page 5317	

[edit logical-systems {logical-system-name} security dynamic-address]	
"address-name" on page 1892	

[edit logical-systems {name} security flow]		
"traceoptions (flow)" on page 29563	"traceoptions (Security Flow)" on page 29789	

[edit logical-systems {name} security idp idp-policy {name} rulebase-exempt rule {name} then notification]		
"log-attacks" on page 10050	"log-attacks (rulebase-exempt)" on page 10052	

[edit logical-systems {logical system name} security idp sensor-configuration]	
"packet-log (Security IDP Sensor Configuration)" on page 12153	

[edit logical-systems {name} security idp sensor-configuration global intelligent-offload-tunable]	
"intelligent-offload-tunable (IDP)" on page 8780	

[edit logical-systems {name} security log]		
"event-rate" on page 6531	"facility-override (Security)" on page 6771	"file (Security Log)" on page 7076
"mode (Security Log)" on page 10823	"rate-cap" on page 14058	"report (Security Log)" on page 14398
"retry-count (Security Log)" on page 15696	"root-streaming" on page 15838	"stream (Logical Systems and Tenant Systems)" on page 28628
"stream (Security Log)" on page 28635	"traceoptions (Security Log)" on page 29811	"transport (Security Log)" on page 30129

[edit logical-systems {name} security log cache]		
"exclude (Security Log)" on page 6627	"limit (Security Log)" on page 9760	

[edit logical-systems {name} security log stream {stream-name}]		
"category (Security Logging)" on page 3224	"format (Security Log Stream)" on page 7388	"host (Security Log)" on page 8166
"radius (Security Log)" on page 14018	"rate-limit (Security Log)" on page 14069	"severity (Security Log)" on page 16906

[edit logical-systems {name} security nat]	
"address-persistent (Security Source NAT)" on page 1896	

[edit logical-systems {name} security nat destination rule-set {rule-set-name} rule {rule-name} match]	
"application (Security Destination NAT)" on page 2369	

[edit logical-systems {name} security nat proxy-arp interface]	
"address (Security ARP Proxy)" on page 1922	

[edit logical-systems {name} security nat proxy-ndp interface]	
"address (Security NDP Proxy)" on page 1925	

[edit logical-systems {name} security nat source {rules-set} name {rule} {name} then source-nat]	
"address-mapping" on page 1887	

[edit logical-systems {name} security nat source pool]	
"address-shared (Security Source NAT)" on page 1932	

[edit logical-systems {name} security nat source pool {name} address-persistent subscriber]	
"address-persistent (Security Source NAT Pool)" on page 1897	

[edit logical-systems {name} security nat source pool {name} address-pooling]	
"address-pooling (Security Source NAT)" on page 1905	

[edit logical-systems {name} security nat source pool {name} pool-utilization-alarm]	
"clear-threshold" on page 4245	

[edit logical-systems {name} security nat source pool-utilization-alarm]	
"clear-threshold" on page 4245	

[edit logical-systems {name} security nat source rule-set {rule-name} rule {name} match]	
"application (Security Source NAT)" on page 2373	

[edit logical-systems {logical-system-name} security policies]		
"dns-cache" on page 5852	"unified-policy max-lookups" on page 30362	

[edit logical-systems {logical-systems-name} security policies explicit-proxy profile {name} policy {policy-name} from then permit]	
"tcp-options (Security Policies)" on page 29008	

[edit logical-systems {name} security policies from-zone]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	

[edit logical-systems {logical-system-name} security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name}]	
"report-skip" on page 14403	

[edit logical-systems {logical-system-name} security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit application-services]	
"application-traffic-control (Application Services)" on page 2402	

[edit logical-systems {logical-system-name} security policies global policy {policy-name}]	
"report-skip" on page 14403	

[edit logical-systems {name} security policies global policy {name} then]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	

[edit logical-systems {logical-systems-name} security policies global policy {policy-name} then permit]	
"tcp-options (Security Policies)" on page 29008	

[edit logical-systems {logical-systems-name} security policies policy-set {policy-policy-set-name} policy {policy-name} then permit]		
"tcp-options (Security Policies)" on page 29008	"tcp-options (Security Policies)" on page 29008	

[edit logical-systems {name} security policies stateful-firewall-rule {name} policy {name} then]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	

[edit logical-systems {logical-systems-name} security policies-zone from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit]	
"tcp-options (Security Policies)" on page 29008	

[edit logical-systems {logical-system-name} security screen]	
"white-list (Security Screen)" on page 31107	

[edit logical-systems {logical-system-name} security screen ids-option {screen-name} udp flood]	
"white-list (Security Screen)" on page 31107	

[edit logical-systems {name} security tunnel-inspection]	
"inspection-profile" on page 8739	

[edit logical-systems {logical-system-name} security utm]		
"custom-objects" on page 4771		

[edit logical-systems {name} security utm custom-objects]		
"url-feed" on page 30489		

[edit logical-systems {name} security utm custom-objects custom-message]		
"custom-page-file" on page 4778		

[edit logical-systems {name} security utm default-configuration anti-virus]		
"notification-options (Security Antivirus)" on page 11693	"sophos-engine" on page 28109	

[edit logical-systems {name} security utm default-configuration anti-virus avira-engine pattern-update]		
"proxy-profile" on page 13838		

[edit logical-systems {name} security utm default-configuration anti-virus sophos-engine]		
"notification-options (Security Antivirus)" on page 11693		

[edit logical-systems {name} security utm default-configuration anti-virus sophos-engine pattern-update]		
"proxy-profile" on page 13838		

[edit logical-systems {name} security utm default-configuration url-feed]		
"global-limit" on page 7708		

[edit logical-systems {logical-system-name} security utm feature-profile]		
"anti-spam (Logical Systems)" on page 2259	"anti-virus (Logical Systems)" on page 2266	"content-filtering (Logical System Security Feature Profile)" on page 4580
"web-filtering (Logical System Security Feature Profile)" on page 31072		

[edit logical-systems {name} security utm feature-profile anti-virus profile]	
"notification-options (Security Antivirus)" on page 11693	

[edit logical-systems {name} security utm feature-profile anti-virus sophos-engine]	
"server (Security Sophos Engine Antivirus)" on page 16670	

[edit logical-systems {name} security utm feature-profile anti-virus sophos-engine pattern-update]	
"proxy-profile" on page 13838	

[edit logical-systems {name} security utm feature-profile anti-virus sophos-engine profile]	
"notification-options (Security Antivirus)" on page 11693	

[edit logical-systems {name} security utm feature-profile web-filtering juniper-enhanced profile {profile-name}]	
"custom-message (Security Web Filtering)" on page 4764	

[edit logical-systems {name} security utm feature-profile web-filtering juniper-local profile {profile-name}]		
"custom-message (Security Web Filtering)" on page 4764		

[edit logical-systems {name} security utm feature-profile web-filtering websense-redirect profile {profile-name}]		
"custom-message (Security Web Filtering)" on page 4764		

[edit logical-systems {logical-system-name} security utm utm-policy {policy-name}]		
"anti-spam (Security Content Security Policy)" on page 2265	"anti-virus (Content Security Policy)" on page 2273	"content-filtering (Security Content Security Policy)" on page 4583
"web-filtering (Security Content Security Policy)" on page 31076		

[edit logical-systems {name} security zones security-zone]		
"white-list (Security-Zone)" on page 31109		

[services]

[edit logical-systems {logical-system-name} services]		
"icap-redirect" on page 8245		

[edit logical-systems {logical-system-name} services icap-redirect profile {name}]		
"fallback-option (ICAP Redirect Service)" on page 6797	"http (icap-redirect profile)" on page 8193	"server (icap-redirect profile)" on page 16635

[edit logical-systems {logical-system-name} services ssl proxy]	
"profile (Services SSL Proxy)" on page 13339	

[statement-name]

[edit logical-systems {logical-system-name} statement-name protocols igmp amt]	
"relay (IGMP)" on page 14292	

[edit logical-systems {logical-system-name} statement-name protocols igmp amt relay]	
"defaults" on page 5224	

[switch-options]

[edit logical-systems {name} switch-options]		
"authentication-whitelist" on page 2671	"interface-mac-limit" on page 8910	"mac-statistics" on page 10263
"mac-table-size" on page 10271		

[edit logical-systems {logical-system-name} switch-options interface {interface-name}]	
"interface-mac-limit" on page 8910	

[edit logical-systems {logical-system-name} switch-options interface-mac-limit {limit}]	
"packet-action" on page 12129	

[system]

[edit logical-systems {logical-system name} system processes]	
"routing (System Processes)" on page 15993	

[edit logical-systems {logical-system-name} system processes static-subscribers]	
"traceoptions (Static Subscribers)" on page 29873	

[edit logical-systems {logical -system-name} system services]		
"dhcp" on page 5583	"dhcp-local-server" on page 5601	"static-subscribers (Dynamic Service Provisioning)" on page 28582
"tcp-forwarding (Remote Device Management)" on page 28987		

[edit logical-systems {logical-system-name} system services ...]		
"allow-active-leasequery (DHCP Local Server)" on page 2163	"allow-bulk-leasequery (DHCP Local Server)" on page 2165	"allow-leasequery (DHCP Local Server)" on page 2191

[edit logical-systems {logical-system-name} system services dhcp-local-server]		
"dhcpv6 (DHCP Local Server)" on page 5658	"dual-stack-group (DHCP Local Server)" on page 6052	"duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)" on page 6066
"forward-snooped-clients (DHCP Local Server)" on page 7404	"pool-match-order" on page 12777	"reauthenticate (DHCP Local Server)" on page 14103
"reconfigure (DHCP Local Server)" on page 14133		

[edit logical-systems {logical-system-name} system services dhcp-local-server (dhcpv6) ...overrides]	
"include-option-82 (DHCP Local Server)" on page 8534	

[edit logical-systems {logical-system-name} system services dhcp-local-server ...]		
"authentication (DHCP Local Server)" on page 2593	"dual-stack (DHCP Local Server Overrides)" on page 6049	"dynamic-profile (DHCP Local Server)" on page 6117
"group (DHCP Local Server)" on page 7814	"interface (DHCP Local Server)" on page 8860	"interface-description (DHCP Local Server)" on page 8854
"interface-name (DHCP Local Server)" on page 8937	"logical-system-name (DHCP Local Server)" on page 10082	"overrides (DHCP Local Server)" on page 12065
"requested-ip-network-match (DHCP Local Server)" on page 15572	"service-profile (DHCP Local Server)" on page 16729	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982
"username-include (DHCP Local Server)" on page 30595		

[edit logical-systems {logical-system-name} system services dhcp-local-server ...overrides]	
"client-discover-match (DHCP Local Server)" on page 4308	

[edit logical-systems {logical-system-name} system services dhcp-local-server authentication]	
"password (DHCP Local Server)" on page 12226	

[edit logical-systems {logical-system-name} system services dhcp-local-server authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"delimiter (DHCP Local Server)" on page 5288	"domain-name (DHCP Local Server)" on page 5884
"mac-address (DHCP Local Server)" on page 10221	"option-60 (DHCP Local Server)" on page 11876	"option-82 (DHCP Local Server Authentication)" on page 11890
"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578	

[edit logical-systems {name} system services dhcp-local-server dhcpv6]		
"dynamic-server" on page 6154	"reauthenticate (DHCP Local Server)" on page 14103	"reconfigure (DHCP Local Server)" on page 14133

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 ...]	
"rapid-commit (DHCPv6 Local Server)" on page 14055	

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 authentication]	
"password (DHCP Local Server)" on page 12226	

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"client-id (DHCP Local Server)" on page 4315	"delimiter (DHCP Local Server)" on page 5288
"domain-name (DHCP Local Server)" on page 5884	"relay-agent-interface-id (DHCP Local Server)" on page 14275	"relay-agent-remote-id (DHCP Local Server)" on page 14282
"relay-agent-subscriber-id (DHCP Local Server)" on page 14288	"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 group {group-name}]	
"reconfigure (DHCP Local Server)" on page 14133	

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 group {group-name} authentication]	
"password (DHCP Local Server)" on page 12226	

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 group {group-name} authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"client-id (DHCP Local Server)" on page 4315	"delimiter (DHCP Local Server)" on page 5288
"domain-name (DHCP Local Server)" on page 5884	"relay-agent-interface-id (DHCP Local Server)" on page 14275	"relay-agent-remote-id (DHCP Local Server)" on page 14282
"relay-agent-subscriber-id (DHCP Local Server)" on page 14288	"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name}]		
"trace (DHCP Local Server)" on page 29456		

[edit logical-systems {name} system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name} overrides]		
"delay-advertise (DHCPv6)" on page 5239	"process-inform" on page 13240	

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name} overrides process-inform]		
"pool (DHCP Local Server Overrides)" on page 12774		

[edit logical-systems {name} system services dhcp-local-server dhcpv6 group {group-name} overrides]		
"delay-advertise (DHCPv6)" on page 5239	"interface-client-limit (DHCP Local Server)" on page 9106	"process-inform" on page 13240

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 group {group-name} overrides process-inform]		
"pool (DHCP Local Server Overrides)" on page 12774		

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 group {group-name} reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"strict (DHCP Local Server)" on page 28646

(Continued)

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 group {group-name} reconfigure]		
"timeout (DHCP Local Server)" on page 29374	"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 group {group-name} reconfigure trigger]	
"radius-disconnect (DHCP Local Server)" on page 14003	

[edit logical-systems {name} system services dhcp-local-server dhcpv6 overrides]		
"delay-advertise (DHCPv6)" on page 5239	"interface-client-limit (DHCP Local Server)" on page 9106	"process-inform" on page 13240

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"strict (DHCP Local Server)" on page 28646
"timeout (DHCP Local Server)" on page 29374	"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160

[edit logical-systems {logical-system-name} system services dhcp-local-server dhcpv6 reconfigure trigger]	
"radius-disconnect (DHCP Local Server)" on page 14003	

[edit logical-systems {name} system services dhcp-local-server dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Local Server)" on page 3375	"classification-key (DHCP Relay Agent)" on page 3376	"dual-stack-interface-client-limit (DHCP Local Server and Relay Agent)" on page 6058
"on-demand-address-allocation" on page 11813	"protocol-master" on page 13448	"reauthenticate (DHCP Local Server)" on page 14103

[edit logical-systems {logical-system-name} system services dhcp-local-server dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Local Server)" on page 2056	"use-primary (DHCP Local Server)" on page 30535	

[edit logical-systems {name} system services dhcp-local-server group]		
"interface-tag (DHCP Local Server)" on page 9081	"reconfigure (DHCP Local Server)" on page 14133	

[edit logical-systems {logical-system-name} system services dhcp-local-server group {group-name} authentication]	
"password (DHCP Local Server)" on page 12226	

[edit logical-systems {logical-system-name} system services dhcp-local-server group {group-name} authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"delimiter (DHCP Local Server)" on page 5288	"domain-name (DHCP Local Server)" on page 5884
"mac-address (DHCP Local Server)" on page 10221	"option-60 (DHCP Local Server)" on page 11876	"option-82 (DHCP Local Server Authentication)" on page 11890
"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578	

[edit logical-systems {logical-system-name} system services dhcp-local-server group {group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Local Server)" on page 2056	"use-primary (DHCP Local Server)" on page 30535	

[edit logical-systems {logical-system-name} system services dhcp-local-server group {group-name} interface {interface-name}]	
"trace (DHCP Local Server)" on page 29456	

[edit logical-systems {name} system services dhcp-local-server group {group-name} interface {interface-name} overrides]		
"delay-offer (DHCPv4)" on page 5254	"process-inform" on page 13240	

[edit logical-systems {logical-system-name} system services dhcp-local-server group {group-name} interface {interface-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit logical-systems {name} system services dhcp-local-server group {group-name} overrides]		
"delay-offer (DHCPv4)" on page 5254	"interface-client-limit (DHCP Local Server)" on page 9106	"process-inform" on page 13240

[edit logical-systems {logical-system-name} system services dhcp-local-server group {group-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit logical-systems {logical-system-name} system services dhcp-local-server group {group-name} reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"timeout (DHCP Local Server)" on page 29374
"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160	

[edit logical-systems {logical-system-name} system services dhcp-local-server group {group-name} reconfigure trigger]	
"radius-disconnect (DHCP Local Server)" on page 14003	

[edit logical-systems {name} system services dhcp-local-server overrides]		
"delay-offer (DHCPv4)" on page 5254	"interface-client-limit (DHCP Local Server)" on page 9106	"process-inform" on page 13240

[edit logical-systems {logical-system-name} system services dhcp-local-server overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit logical-systems {logical-system-name} system services dhcp-local-server pool-match-order]		
"external-authority" on page 6726	"ip-address-first" on page 9235	"option-82 (DHCP Local Server Pool Matching)" on page 11892

[edit logical-systems {logical-system-name} system services dhcp-local-server reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"timeout (DHCP Local Server)" on page 29374
"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160	

[edit logical-systems {logical-system-name} system services dhcp-local-server reconfigure trigger]		
"radius-disconnect (DHCP Local Server)" on page 14003		

[edit logical-systems {logical-system-name} system services static-subscribers]		
"access-profile (Static Subscribers)" on page 1698	"authentication (Static Subscribers)" on page 2665	"group (Static Subscribers)" on page 7897

[edit logical-systems {logical-system-name} system services static-subscribers authentication]		
"password (Static Subscribers)" on page 12247	"username-include (Static Subscribers)" on page 30604	

[edit logical-systems {logical-system-name} system services static-subscribers authentication username-include]		
"domain-name (Static Subscribers)" on page 5892	"interface (Static Subscriber Username)" on page 9070	"logical-system-name (Static Subscribers)" on page 10085
"routing-instance-name (Static Subscribers)" on page 15949	"user-prefix (Static Subscribers)" on page 30582	

[edit logical-systems {logical-system-name} system services static-subscribers dynamic-profile {profile-name}]		
"aggregate-clients (Static Subscribers)" on page 2060		

[edit logical-systems {logical-system-name} system services static-subscribers group {group-name}]		
"access-profile (Static Subscribers)" on page 1698	"authentication (Static Subscribers)" on page 2665	"interface (Static Subscriber Group)" on page 9068

[edit logical-systems {logical-system-name} system services static-subscribers group {group-name} authentication]		
"password (Static Subscribers)" on page 12247	"username-include (Static Subscribers)" on page 30604	

[edit logical-systems {logical-system-name} system services static-subscribers group {group-name} authentication username-include]		
"domain-name (Static Subscribers)" on page 5892	"interface (Static Subscriber Username)" on page 9070	"logical-system-name (Static Subscribers)" on page 10085
"routing-instance-name (Static Subscribers)" on page 15949	"user-prefix (Static Subscribers)" on page 30582	

[edit logical-systems {logical-system-name} system services static-subscribers group {group-name} dynamic-profile {profile-name}]	
"aggregate-clients (Static Subscribers)" on page 2060	

[edit logical-systems {logical-system-name} system services system services dhcp-local-server dhcpv6 ...]		
"delegated-pool (DHCP Local Server)" on page 5271	"multi-address-embedded-option-response (DHCP Local Server)" on page 11039	

[edit logical-systems {logical-system-name} system syslog]		
"allow-duplicates" on page 2172	"host (System)" on page 8171	

[edit logical-systems {logical-system-name} system syslog file {file-name}]		
"allow-duplicates" on page 2172	"explicit-priority" on page 6657	"file (System Logging)" on page 7086
"match (Syslog)" on page 10369	"structured-data" on page 28664	

[edit logical-systems {logical-system-name} system syslog host {host-name}]		
"allow-duplicates" on page 2172	"explicit-priority" on page 6657	"file (System Logging)" on page 7086

[edit logical-systems {logical-system-name} system syslog user {user-name}]		
"allow-duplicates" on page 2172	"file (System Logging)" on page 7086	

[edit logical-systems {logical-system-name} system syslog user ({username} | *)]	
"match (Syslog)" on page 10369	

[tenants]

[edit logical-systems {name} tenants {name} access]	
"firewall-authentication (tenants)" on page 7189	

[edit logical-systems {name} tenants {name} access firewall-authentication]		
"pass-through (Firewall-authentication)" on page 12196	"web-authentication (Access)" on page 31058	

[edit logical-systems {name} tenants {name} applications application]	
"application-protocol (Applications)" on page 2355	

[edit logical-systems {name} tenants {name} routing-instances {name}]		
"mac-vrf" on page 10276	"service-type (MAC-VRF Routing Instances)" on page 16764	

[edit logical-systems {name} tenants {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit logical-systems {name} tenants {name} routing-instances {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit logical-systems {name} tenants {name} routing-instances {name} forwarding-options pfe-sensor npu-memory]	
"resource-list (Telemetry)" on page 15612	

[edit logical-systems {name} tenants {name} routing-instances {name} forwarding-options storm-control]	
"enhanced" on page 6354	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols]		
"flex-algorithm (Protocols OSPF)" on page 7223	"rsvp" on page 16050	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols (evpn | l2vpn | vpls)mesh-group {mesh-group-name} neighbor {neighbor-id} oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"multiplier (BFD Liveness Detection)" on page 11123	"no-adaptation (BFD Liveness Detection)" on page 11374
"version (BFD Liveness Detection)" on page 30697		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols (evpn | l2vpn | vpls)neighbor {neighbor-id} oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"multiplier (BFD Liveness Detection)" on page 11123	"no-adaptation (BFD Liveness Detection)" on page 11374
"version (BFD Liveness Detection)" on page 30697		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols (evpn | l2vpn | vpls)oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"multiplier (BFD Liveness Detection)" on page 11123	"no-adaptation (BFD Liveness Detection)" on page 11374
"version (BFD Liveness Detection)" on page 30697		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp]		
"advertise-prefix-sid" on page 2022	"bfd-liveness-detection (BGP)" on page 2918	"defaults (ACX series, JRR200, MX204, PTX5000, and VRR)" on page 5225
"egress-te-set-segment" on page 6236	"forwarding-context (Protocols BGP)" on page 7453	"multipath (Protocols BGP)" on page 11110

(Continued)

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp]		
"tcpao-auth-mismatch" on page 29034		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp bfd-liveness-detection]		
"authentication (BGP BFD Liveness Detection)" on page 2588		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp family inet labeled-unicast]		
"entropy-label (protocols bgp)" on page 6413		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp family inet6]		
"transport (Protocols BGP)" on page 30113		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {group-name}]		
"multipath (Protocols BGP)" on page 11110	"peer-as-list" on page 12345	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} bfd-liveness-detection]		
"authentication (BGP BFD Liveness Detection)" on page 2588		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} dynamic-neighbor {dyn-name}]		
"allow" on page 2160	"peer-limit" on page 12375	

[edit logical-systems {logical-system-name} tenants {name} routing-instances {name} protocols bgp group {group-name} family {family} {family-modifier}]	
"add-path" on page 1839	

[edit logical-systems {logical-system-name} tenants {name} routing-instances {name} protocols bgp group {group-name} family {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} family inet labeled-unicast]	
"entropy-label (protocols bgp)" on page 6413	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} family inet6]		
"transport (Protocols BGP)" on page 30113		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} link-bandwidth]		
"auto-sense" on page 2738		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} metric-out]		
"effective-aigp" on page 6203	"minimum-effective-aigp" on page 10730	

[edit logical-systems {logical-system-name} tenants {name} routing-instances {name} protocols bgp group {group-name} neighbor {address} {family} {family-modifier}]		
"add-path" on page 1839	"egress-te-adj-segment" on page 6229	"egress-te-node-segment" on page 6234
"multipath (Protocols BGP)" on page 11110		

[edit logical-systems {logical-system-name} tenants {name} routing-instances {name} protocols bgp group {group-name} neighbor {address} {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} neighbor {address} bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet labeled-unicast]	
"entropy-label (protocols bgp)" on page 6413	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp multipath]	
"pause-computation-during-churn" on page 12277	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgp source-packet-routing]		
"srgb (Protocols BGP)" on page 28407	"srv6 (Protocols BGP)" on page 28417	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgpgroup {group-name}]	
"bfd-liveness-detection (BGP)" on page 2918	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols bgpgroup {group-name} neighbor {address}]	
"bfd-liveness-detection (BGP)" on page 2918	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols evpn interface]	
"static-mac-ip" on page 28587	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols evpn | l2vpn | vpls) mesh-group {mesh-group-name} neighbor {neighbor-id} oam bfd-liveness-detection]	
"minimum-receive-interval (BFD Liveness Detection)" on page 10754	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols evpn | l2vpn | vpls) neighbor {neighbor-id} oam bfd-liveness-detection]	
"minimum-receive-interval (BFD Liveness Detection)" on page 10754	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols evpn | l2vpn | vpls) oam bfd-liveness-detection]	
"minimum-receive-interval (BFD Liveness Detection)" on page 10754	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols isis dynamic-overload]		
"auto-recovery" on page 2733	"no-overload-on-prefix-export-limit" on page 11542	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols isis interface {name} delay-measurement]	
"advertisement (Delay Measurement)" on page 2031	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols isis interface-group {name} level {name} topology]	
"bandwidth-based-metrics" on page 2837	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols isis no-external-export protocol]	
"no-external-export" on page 11454	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols isis source-packet-routing]	
"strict-asla-based-flex-algorithm" on page 28645	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols isis spf-options]		
"microloop-avoidance" on page 10700		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols isis traffic-engineering]		
"advertisement" on page 2030		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols mpls]		
"cspf-backoff-time" on page 4731		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols mpls label-switched-path {name}]		
"conditional-metric" on page 4506	"cspf-backoff-time" on page 4731	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols mpls label-switched-path {name} secondary {name} non-active-hold-priority]		
"non-active-hold-priority" on page 11657		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols mpls static-label-switched-path {name}]		
"transit" on page 30064		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols ospf spf-options]		
"microloop-avoidance (Protocols OSPF)" on page 10702		

[edit logical-systems {name} tenants {name} routing-instances {name} protocols pim]		
"disable-packet-register" on page 5813	"secondary-vtep-address" on page 16419	

[edit logical-systems {name} tenants {name} routing-instances {name} protocols pim interface]	
"dual-dr" on page 6045	

[edit logical-systems {name} tenants {name} routing-instances {name} provider-tunnel ingress-replication]		
"transport-class" on page 30132	"unicast-tunnel-name-regular-expression" on page 30355	

[edit logical-systems {name} tenants {name} routing-instances {name} provider-tunnel ldp-p2mp]	
"root-address" on page 15840	

[edit logical-systems {name} tenants {name} routing-instances {name} routing-options]		
"flex-algorithm" on page 7221	"flex-algorithm (Protocols OSPF)" on page 7223	"route-distinguisher-id-use-router-id" on page 15889
"routing-options-max-ecmp" on page 15985		

[edit logical-systems {name} tenants {name} routing-instances {name} routing-options bmp route-monitoring]	
"loc-rib" on page 9897	

[edit logical-systems {name} tenants {name} routing-instances {name} routing-options dynamic-tunnels]		
"forwarding-rib" on page 7483	"v4ov6" on page 30636	

[edit logical-systems {name} tenants {name} routing-instances {name} routing-options dynamic-tunnels tunnel-attributes]		
"dynamic-tunnel-reassembly" on page 6159		

[edit logical-systems {name} tenants {name} routing-instances {name} routing-options flex-algorithm]		
"definition" on page 5231	"definition (Protocols OSPF)" on page 5234	"use-transport-class" on page 30541

[edit logical-systems {name} tenants {name} routing-instances {name} routing-options flow]		
"fast-lookup-filter (Protocols BGP)" on page 6947		

[edit logical-systems {name} tenants {name} routing-instances {name} routing-options multicast]		
"cont-stats-collection-interval" on page 4573		

[edit logical-systems {name} tenants {name} routing-instances {name} routing-options rib]		
"policy-multipath" on page 12694	"vpn-global-import" on page 30953	

[edit logical-systems {name} tenants {name} routing-instances {name} routing-options source-packet-routing srv6]		
"block" on page 2980	"locator" on page 9997	

[edit logical-systems {name} tenants {name} routing-instances {name} routing-options source-packet-routing srv6 locator]		
"micro-sid" on page 10697		

[edit logical-systems {name} tenants {name} routing-instances {name} routing-options transport-class]	
"name (Transport-class)" on page 11181	

[edit logical-systems {name} tenants {name} routing-instances {name} system services dhcp-local-server dhcpv6]	
"dynamic-server" on page 6154	

[edit logical-systems {name} tenants {name} routing-instances {name} system services dhcp-local-server group]	
"interface-tag (DHCP Local Server)" on page 9081	

[edit logical-systems {name} tenants {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit logical-systems {name} tenants {name} security alg]		
"alg-support-lib" on page 2142	"dns" on page 5847	"ftp (Security ALG)" on page 7626
"ike-esp-nat" on page 8421	"mgcp" on page 10693	"msrpc" on page 10999
"pptp" on page 12986	"rsh" on page 16041	"rtsp" on page 16072
"sccp" on page 16341	"sip (Security) " on page 27977	"sunrpc" on page 28722
"talk" on page 28940	"tftp (Security ALG)" on page 29252	"traceoptions (Security ALG)" on page 29770

[edit logical-systems {name} tenants {name} security alg dns]		
"maximum-message-length" on page 10534	"oversize-message-drop" on page 12081	

[edit logical-systems {name} tenants {name} security alg h323]		
"application-screen (Security H323)" on page 2359	"endpoint-registration-timeout" on page 6342	"media-source-port-any" on page 10599
"traceoptions (Security H323 ALG)" on page 29801		

[edit logical-systems {name} tenants {name} security alg h323 application-screen]		
"message-flood (Security H323)" on page 10644	"unknown-message (Security H323 ALG)" on page 30410	

[edit logical-systems {name} tenants {name} security alg h323 application-screen message-flood]	
"gatekeeper" on page 7653	

[edit logical-systems {name} tenants {name} security alg mgcp]		
"application-screen (Security MGCP)" on page 2361	"inactive-media-timeout (Security MGCP)" on page 8496	"traceoptions (Security MGCP ALG)" on page 29814
"transaction-timeout" on page 30055		

[edit logical-systems {name} tenants {name} security alg mgcp application-screen]		
"connection-flood" on page 4531	"message-flood (Security MGCP)" on page 10645	"unknown-message (Security MGCP ALG)" on page 30411

[edit logical-systems {name} tenants {name} security alg sccp]		
"application-screen (Security SCCP)" on page 2362	"inactive-media-timeout (Security SCCP)" on page 8497	"traceoptions (Security SCCP ALG)" on page 29823

[edit logical-systems {name} tenants {name} security alg sccp application-screen]		
"call-flood" on page 3179	"unknown-message (Security SCCP ALG)" on page 30413	

[edit logical-systems {name} tenants {name} security alg sip]		
"application-screen (Security SIP)" on page 2364	"c-timeout" on page 3152	"t1-interval" on page 28885
"t4-interval" on page 28893	"traceoptions (Security SIP ALG)" on page 29830	

[edit logical-systems {name} tenants {name} security alg sip (Security)]		
"inactive-media-timeout (Security SIP)" on page 8499	"retain-hold-resource" on page 15674	

[edit logical-systems {name} tenants {name} security alg sip (Security) application-screen (Security SIP)]	
"unknown-message (Security SIP ALG)" on page 30415	

[edit logical-systems {name} tenants {name} security alg sip application-screen protect]	
"deny (Security SIP)" on page 5317	

[edit logical-systems {name} tenants {name} security idp idp-policy {name} rulebase-exempt rule {name} then notification]		
"log-attacks" on page 10050	"log-attacks (rulebase-exempt)" on page 10052	

[edit logical-systems {name} tenants {name} security idp sensor-configuration global intelligent-offload-tunable]	
"intelligent-offload-tunable (IDP)" on page 8780	

[edit logical-systems {name} tenants {name} security log]	
"retry-count (Security Log)" on page 15696	

[edit logical-systems {name} tenants {name} security log stream]	
"radius (Security Log)" on page 14018	

[edit logical-systems {name} tenants {name} security nat]	
"address-persistent (Security Source NAT)" on page 1896	

[edit logical-systems {name} tenants {name} security nat destination rule-set {rule-set-name} rule {rule-name} match]	
"application (Security Destination NAT)" on page 2369	

[edit logical-systems {name} tenants {name} security nat proxy-arp interface]	
"address (Security ARP Proxy)" on page 1922	

[edit logical-systems {name} tenants {name} security nat proxy-ndp interface]	
"address (Security NDP Proxy)" on page 1925	

[edit logical-systems {name} tenants {name} security nat source pool]	
"address-shared (Security Source NAT)" on page 1932	

[edit logical-systems {name} tenants {name} security nat source pool {name} address-persistent subscriber]	
"address-persistent (Security Source NAT Pool)" on page 1897	

[edit logical-systems {name} tenants {name} security nat source pool {name} address-pooling]	
"address-pooling (Security Source NAT)" on page 1905	

[edit logical-systems {name} tenants {name} security nat source pool {name} pool-utilization-alarm]	
"clear-threshold" on page 4245	

[edit logical-systems {name} tenants {name} security nat source pool-utilization-alarm]	
"clear-threshold" on page 4245	

[edit logical-systems {name} tenants {name} security nat source rule-set {name} rule {name} match]	
"application (Security Source NAT)" on page 2373	

[edit logical-systems {name} tenants {name} security nat source {rule} {name} then source-nat]	
"address-mapping" on page 1887	

[edit logical-systems {logical-system-name} tenants {tenant-system-name} security policies]		
"dns-cache" on page 5852	"unified-policy max-lookups" on page 30362	

[edit logical-systems {name} tenants {name} security policies explicit-proxy]	
"profile (Explicit Proxy)" on page 13277	

[edit logical-systems {logical-systems-name} tenants {tenants-tenants-systems-name} security policies explicit-proxy profile {policy-name} policy {policy-name} from then]	
"tcp-options (Security Policies)" on page 29008	

[edit logical-systems {name} tenants {name} security policies from-zone]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	

[edit logical-systems {name} tenants {name} security policies global policy {name} then]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	

[edit logical-systems {logical-systems-name} tenants {tenants-tenants-systems-name} security policies global policy {policy-name} then permit]		
"tcp-options (Security Policies)" on page 29008		

[edit logical-systems {logical-systems-name} tenants {tenants-tenants-systems-name} security policies policy-setset {policy-set-name} policy {policy-name} then permit]		
"tcp-options (Security Policies)" on page 29008		

[edit logical-systems {name} tenants {name} security policies stateful-firewall-rule {name} policy {name} then]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	

[edit logical-systems {logical-systems-name} tenants {tenants-systems-name} security policies-zone from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit]		
"tcp-options (Security Policies)" on page 29008		

[edit logical-systems {name} tenants {name} security tunnel-inspection]		
"inspection-profile" on page 8739		

[edit logical-systems {name} tenants {name} security utm custom-objects]		
"url-feed" on page 30489		

[edit logical-systems {name} tenants {name} security utm default-configuration anti-virus]		
"notification-options (Security Antivirus)" on page 11693	"sophos-engine" on page 28109	

[edit logical-systems {name} tenants {name} security utm default-configuration anti-virus avira-engine pattern-update]	
"proxy-profile" on page 13838	

[edit logical-systems {name} tenants {name} security utm default-configuration anti-virus sophos-engine]	
"notification-options (Security Antivirus)" on page 11693	

[edit logical-systems {name} tenants {name} security utm default-configuration anti-virus sophos-engine pattern-update]	
"proxy-profile" on page 13838	

[edit logical-systems {name} tenants {name} security utm default-configuration url-feed]	
"global-limit" on page 7708	

[edit logical-systems {name} tenants {name} security utm feature-profile anti-virus profile]	
"notification-options (Security Antivirus)" on page 11693	

[edit logical-systems {name} tenants {name} security utm feature-profile anti-virus sophos-engine]	
"server (Security Sophos Engine Antivirus)" on page 16670	

[edit logical-systems {name} tenants {name} security utm feature-profile anti-virus sophos-engine pattern-update]	
"proxy-profile" on page 13838	

[edit logical-systems {name} tenants {name} security utm feature-profile anti-virus sophos-engine profile]	
"notification-options (Security Antivirus)" on page 11693	

[edit logical-systems {name} tenants {name} security zones security-zone]	
"white-list (Security-Zone)" on page 31109	

[three-color-policer]

[edit logical-systems {logical-system-name} three-color-policer {policer-name}]	
"physical-interface-policer" on page 12479	

[transport-ls]

[edit logical-systems transport-ls interfaces]	
"ps0 (Pseudowire Subscriber Interfaces)" on page 13847	

[vlans]

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay dhcpv6 group {name} relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay dhcpv6 relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay group {name} relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877

(Continued)

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay group {name} relay-option]		
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {name} vlans {name} forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit logical-systems {logical-system-name} vlans {vlan-name} switch-options]	
"interface-mac-ip-limit" on page 8908	

[edit multi-chassis]

IN THIS SECTION

- [multi-chassis-protection] | 645

Sub Hierarchies

"[edit multi-chassis multi-chassis-protection {IP address of MC-LAG peer}]" on page 645

[edit multi-chassis]		
"multi-chassis-protection" on page 11044	"peer (Multichassis)" on page 12376	

[multi-chassis-protection]

[edit multi-chassis multi-chassis-protection {IP address of MC-LAG peer}]	
"icl-down-delay" on page 8252	

[edit multi-chassis multi-chassis-protection peer]	
"interface (Multichassis Protection)" on page 8930	

[edit multicast-snooping-options]

IN THIS SECTION

- [graceful-restart] | 647

Sub Hierarchies

"[edit multicast-snooping-options graceful-restart]" on page 647

[edit multicast-snooping-options]		
"graceful-restart (Multicast Snooping)" on page 7768	"host-outbound-traffic (Multicast Snooping)" on page 8162	"multichassis-lag-replicate-state" on page 11095
"traceoptions (Multicast Snooping Options)" on page 29604		

[graceful-restart]

[edit multicast-snooping-options graceful-restart]		
"disable (Multicast Snooping)" on page 5775	"restart-duration (Multicast Snooping)" on page 15653	

[edit node-group]

IN THIS SECTION

- [aggregated-devices] | 647
- [node-device] | 648

Sub Hierarchies

"[edit node-group aggregated-devices]" on page 647

"[edit node-group {name} node-device {name}]" on page 648

[aggregated-devices]

[edit node-group {name} aggregated-devices ethernet]	
"device-count" on page 5565	

[node-device]

[edit node-group {name} node-device {name}]	
"pic" on page 12480	

[edit poe]

IN THIS SECTION

- [(all | fpc {slot-number})] | 649
- [fpc] | 649
- [interface] | 650
- [notification-control] | 651

Sub Hierarchies

"[edit poe (all | fpc {slot-number})]" on page 649

"[edit poe fpc]" on page 649

"[edit poe interface]" on page 650

"[edit poe notification-control]" on page 651

[edit poe]		
"fast-poe" on page 6939	"guard-band" on page 7944	"guard-band (PoE)" on page 7946

(Continued)

[edit poe]		
"interface (PoE)" on page 8951	"interface (Power over Ethernet)" on page 8957	"lldp-priority" on page 9860
"management" on page 10305	"management (PoE)" on page 10310	"notification-control" on page 11688
"perpetual-poe" on page 12450		

[(all | fpc {slot-number})]

[edit poe (all | fpc {slot-number})]		
"guard-band" on page 7944	"management" on page 10305	

[fpc]

[edit poe fpc (all | {slot-number})]	
"lldp-priority" on page 9860	

[interface]

[edit poe interface (all | all-extended | {interface-name})]		
"disable (Power over Ethernet)" on page 5785	"maximum-power (Interface)" on page 10549	

[edit poe interface (all | all-extended | {interface-name}) telemetries]	
"disable (Power over Ethernet)" on page 5785	

[edit poe interface (all | {interface-name})]		
"af-mode" on page 2039	"disable (PoE)" on page 5784	"high-power" on page 8077
"maximum-power (PoE)" on page 10552	"priority (PoE)" on page 13188	"telemetries" on page 29048
"telemetries (PoE)" on page 29050	"ultrahigh-power" on page 30325	

[edit poe interface (all | {interface-name}) telemetries]		
"disable (PoE)" on page 5784	"duration" on page 6078	"duration (PoE)" on page 6080
"interval (PoE)" on page 9216	"interval (Power over Ethernet)" on page 9217	

[edit poe interface ({interface-name} | all | all-extended)]	
"priority (Power over Ethernet)" on page 13191	

[notification-control]

[edit poe notification-control]	
"fpc (Notification Control)" on page 7523	

[edit poe notification-control fpc {slot-number}]	
"disable (Power over Ethernet)" on page 5785	

[edit policy-options]

IN THIS SECTION

- [application-maps] | 653
- [as-path] | 653
- [as-path-group] | 654
- [community] | 654
- [condition] | 654
- [defaults] | 655
- [policy-statement] | 655
- [policy-statementterm] | 658
- [policy-statementthen] | 658
- [prefix-list] | 658
- [satellite-policies] | 659
- [then] | 660

Sub Hierarchies

"[edit policy-options application-maps {application-map-name}]" on page 653

"[edit policy-options as-path {path-name}]" on page 653

"[edit policy-options as-path-group {group-name}]" on page 654

"[edit policy-options community {community-name}]" on page 654

"[edit policy-options condition {condition-name}]" on page 654

"[edit policy-options defaults]" on page 655

"[edit policy-options policy-statement {policy-statement-name}]" on page 655

"[edit policy-options policy-statementterm]" on page 658

"[edit policy-options policy-statement{policy-name} then]" on page 658

"[edit policy-options prefix-list {name}]" on page 658

"[edit policy-options satellite-policies]" on page 659

"[edit policy-options {name} then]" on page 660

[edit policy-options]		
"application-maps" on page 2346	"as-list" on page 2487	"as-path (Policy Options)" on page 2493
"as-path-group" on page 2492	"community (Policy Options)" on page 4465	"community (Protocols Layer 2 Circuit)" on page 4469
"condition" on page 4502	"damping (Policy Options)" on page 4820	"defaults (Policy Options)" on page 5228
"fast-lookup-tuple-list (policy-options)" on page 6938	"mac-list" on page 10243	"policy-statement" on page 12736
"prefix-list" on page 13090	"redundancy-policy (Interchassis Services Redundancy)" on page 14199	"resolution-map" on page 15597

(Continued)

[edit policy-options]		
"rib-list" on page 15783	"route-filter-list" on page 15863	"rtf-prefix-list" on page 16064
"satellite-policies" on page 16313	"source-address-filter-list" on page 28131	"vsi-policy" on page 30981

[application-maps]

[edit policy-options application-maps {application-map-name}]	
"application (Application Maps)" on page 2308	

[edit policy-options application-maps {application-map-name} application {application-name}]	
"code-points (Application Maps)" on page 4402	

[as-path]

[edit policy-options as-path {path-name}]	
"dynamic-db" on page 6105	

[as-path-group]

[edit policy-options as-path-group {group-name}]	
"dynamic-db" on page 6105	

[community]

[edit policy-options community {community-name}]	
"dynamic-db" on page 6105	

[edit policy-options community {community-name} members]		
"ipv6-extended" on page 9348	"ipv6-origin" on page 9367	"ipv6-target" on page 9374

[condition]

[edit policy-options condition {condition-name}]		
"dynamic-db" on page 6105	"if-route-exists" on page 8377	"route-active-on" on page 15848

[edit policy-options condition if-route-exists]		
"address-family" on page 1876	"table" on page 28895	

[edit policy-options condition if-route-exists address-family]		
"ccc (Routing Policy Condition)" on page 3242	"inet (Routing Policy Condition)" on page 8588	

[edit policy-options condition if-route-exists address-family ccc]		
"peer-unit (Routing Policy Condition)" on page 12391	"standby (Routing Policy Condition)" on page 28491	"table" on page 28895

[edit policy-options condition if-route-exists address-family inet]	
"table" on page 28895	

[defaults]

[edit policy-options defaults]	
"route-filter" on page 15861	

[edit policy-options defaults route-filter]	
"walkup" on page 31019	

[policy-statement]

[edit policy-options policy-statement {policy-statement-name}]		
"defaults (Policy Options)" on page 5228	"dynamic-db" on page 6105	"large-community-count" on page 9634

[edit policy-options policy-statement {policy-statement-name} defaults]		
"route-filter" on page 15861		

[edit policy-options policy-statement {policy-statement-name} defaults route-filter]		
"walkup" on page 31019		

[edit policy-options policy-statement {name} from]		
"traffic-engineering (policy-options)" on page 30016		

[edit policy-options policy-statement {name} from route-filter]		
"get-route-range" on page 7681		

[edit policy-options policy-statement {name} from traffic-engineering]		
"ipv6-prefix" on page 9368	"spring-te" on page 28393	"srv6-sid" on page 28419
"te-policy" on page 29044		

[edit policy-options policy-statement {name} from traffic-engineering spring-te]		
"protocol-origin" on page 13452		

[edit policy-options policy-statement policy-name term term-name then]		
"fib-install-priority" on page 7003		

[edit policy-options policy-statement {<statement-name>} term {<term-name>}]		
"large-community-count" on page 9634	"rtf-prefix-list" on page 16064	

[edit policy-options policy-statement {policy-statement-name} term {term--name} from]		
"prefix-list-filter" on page 13092	"prefix-segment" on page 13105	"route-filter-list" on page 15863
"source-address-filter-list" on page 28131		

[edit policy-options policy-statement {name} term {name} from route-filter]	
"get-route-range" on page 7681	

[edit policy-options policy-statement {policy-name} term {term-name} from traffic-engineering]	
"ipv4-prefix" on page 9338	

[edit policy-options policy-statement {policy-name} term {term-name} then]		
"aigp-adjust (Policy Action)" on page 2095	"aigp-originate" on page 2097	"filter (export routing policy action)" on page 7135
"install-nexthop" on page 8742	"metric (Policy Action)" on page 10681	"no-route-localize" on page 11574
"prefix-segment" on page 13105	"slice (export routing policy action)" on page 28041	"source-based-fpolicy-statementorwarding" on page 28333
"strip-as-path" on page 28662		

[edit policy-options policy-statement {policy-name} then]		
"aigp-adjust (Policy Action)" on page 2095	"aigp-originate" on page 2097	"load-balance (ecmp)" on page 9872
"load-balance profile1/profile2" on page 9883	"metric (Policy Action)" on page 10681	"prefix-attribute-flags" on page 13072

[policy-statementterm]

[edit policy-options policy-statement{policy-name} term {term-name} then]	
"priority (policy-options)" on page 13189	

[policy-statementthen]

[edit policy-options policy-statement{policy-name} then]	
"priority (policy-options)" on page 13189	

[prefix-list]

[edit policy-options prefix-list {name}]		
"apply-path" on page 2431	"dynamic-db" on page 6105	

[satellite-policies]

[edit policy-options satellite-policies]		
"candidate-uplink-port-policy (satellite-policies)" on page 3196	"environment-monitoring-policy (satellite-policies)" on page 6419	"port-group-alias (satellite-policies)" on page 12829

[edit policy-options satellite-policies candidate-uplink-policy {policy-name}]	
"term (satellite-policies)" on page 29130	

[edit policy-options satellite-policies candidate-uplink-port-profile{profile-name}]		
"holddown (candidate-uplink-port-profile)" on page 8120	"minimum-links (candidate-uplink-port-profile)" on page 10744	"term (candidate-uplink-policy)" on page 29111
"uplink-port-group (Junos Fusion)" on page 30469		

[edit policy-options satellite-policies candidate-uplink-port-profile {profile-name} term {term-name} from]		
"holddown (candidate-uplink-port-profile)" on page 8120	"minimum-links (candidate-uplink-port-profile)" on page 10744	"uplink-port-group (Junos Fusion)" on page 30469

[edit policy-options satellite-policies candidate-uplink–port-policy {policy-name} term {term-name} from]		
"product-model (Junos Fusion)" on page 13250	"product-model (Junos Fusion)" on page 13250	

[edit policy-options satellite-policies environment-monitoring-policy {policy-name}]		
"alarm (Satellite Policies)" on page 2114	"term (satellite-policies)" on page 29130	

[edit policy-options satellite-policies environment-monitoring-policy {policy-name} alarm]	
"linkdown (satellite-policies alarm)" on page 9835	

[edit policy-options satellite-policies forwarding-policy {policy-name}]	
"term (satellite-policies)" on page 29130	

[edit policy-options satellite-policies forwarding-policy {policy-name} term {term-name} from]	
"product-model (Junos Fusion)" on page 13250	

[edit policy-options satellite-policies port-group-alias {port-group-alias-name}]	
"pic (satellite-policies port-group-alias)" on page 12502	

[edit policy-options satellite-policies port-group-alias {port-group-alias-name} pic {pic-number}]	
"port (satellite-policies port-group-alias)" on page 12871	

[then]

[edit policy-options {name} then]	
"assisted-replication (Deterministic AR Replicator Policy Actions)" on page 2507	

[edit protocols]

IN THIS SECTION

- [(mstp | rstp| vstp)] | 668
- [(bgp | esis | isis | ospf | ospf3 | ldp | pim | rip | ripng | rsvp)] | 669
- [(bgp | rip | ripng)] | 669
- [(esis | isis | ospf | ospf3 | pim)] | 669
- [(isis | ldp | ospf | ospf3 | rsvp)] | 669
- [(mstp | rstp | vstp | vstp vlan {vlan-id})] | 670
- [(mstp | rstp | vstp)] | 670
- [(mstp | rstp)] | 671
- [(ospf | ospf3)] | 671
- [(rstp | vstp)] | 673
- [amt] | 673
- [ancp] | 674
- [authentication-access-control] | 675
- [bfd] | 676
- [bgp] | 676
- [connections] | 701
- [dcbx] | 702
- [dot1x] | 703
- [dvmrp] | 704
- [edge-virtual-bridging] | 705
- [ethernet] | 705
- [evpn] | 705
- [group] | 707
- [iccp] | 707
- [igmp] | 708
- [igmp-snooping] | 710
- [interface] | 712
- [isis] | 712

- [l2-learning] | **716**
- [l2circuit] | **717**
- [lacp] | **720**
- [layer2-control] | **721**
- [ldp] | **721**
- [link-management] | **725**
- [lldp] | **726**
- [lldp-med] | **727**
- [mld] | **728**
- [mld-snooping] | **729**
- [mpls] | **730**
- [msdp] | **745**
- [mstp] | **747**
- [mstpmsti] | **749**
- [mvpn] | **749**
- [mvrp] | **751**
- [neighbor-discovery] | **751**
- [oam] | **752**
- [openflow] | **761**
- [ospf] | **762**
- [ospf3] | **765**
- [ovsdb] | **767**
- [pcep] | **767**
- [pim] | **768**
- [piminterface] | **776**
- [ppp] | **776**
- [ppp-service] | **776**
- [pppoe] | **777**
- [protection-group] | **778**
- [ptp] | **779**
- [rift] | **782**
- [rip] | **783**

- [ripng] | **784**
- [router-advertisement] | **785**
- [router-discovery] | **786**
- [rstp] | **787**
- [rsvp] | **788**
- [sap] | **791**
- [sflow] | **791**
- [source-packet-routing] | **792**
- [spring-traffic-engineering] | **794**
- [stp] | **794**
- [uplink-failure-detection] | **795**
- [vlan] | **795**
- [vpls] | **796**
- [vrrp] | **796**
- [vstp] | **796**

Sub Hierarchies

"[edit protocols (mstp | rstp| vstp)]" on page 668

"[edit protocols (bgp | esis | isis | ospf | ospf3 | ldp | pim | rip | ripng | rsvp)]" on page 669

"[edit protocols (bgp | rip | ripng)]" on page 669

"[edit protocols (esis | isis | ospf | ospf3 | pim)]" on page 669

"[edit protocols (isis | ldp | ospf | ospf3 | rsvp)]" on page 669

"[edit protocols (mstp | rstp | vstp | vstp vlan {vlan-id})]" on page 670

"[edit protocols (mstp | rstp | vstp)]" on page 670

"[edit protocols (mstp | rstp)]" on page 671

"[edit protocols (ospf | ospf3)]" on page 671

"[edit protocols (rstp | vstp)]" on page 673

"[edit protocols amt]" on page 673

"[edit protocols ancp]" on page 674

"[edit protocols authentication-access-control]" on page 675

"[edit protocols bfd]" on page 676

"[edit protocols bgp]" on page 676

"[edit protocols connections]" on page 701

"[edit protocols dcbx]" on page 702

"[edit protocols dot1x]" on page 703

"[edit protocols dvmrp]" on page 704

"[edit protocols edge-virtual-bridging]" on page 705

"[edit protocols ethernet]" on page 705

"[edit protocols evpn]" on page 705

"[edit protocols group]" on page 707

"[edit protocols iccp]" on page 707

"[edit protocols igmp]" on page 708

"[edit protocols igmp-snooping]" on page 710

"[edit protocols interface {name}]" on page 712

"[edit protocols isis]" on page 712

"[edit protocols l2-learning]" on page 716

"[edit protocols l2circuit]" on page 717

"[edit protocols lacp]" on page 720

"[edit protocols layer2-control]" on page 721

"[edit protocols ldp]" on page 721

"[edit protocols link-management]" on page 725

"[edit protocols lldp]" on page 726

"[edit protocols lldp-med]" on page 727

"[edit protocols mld]" on page 728

"[edit protocols mld-snooping]" on page 729

"[edit protocols mpls]" on page 730

"[edit protocols msdp]" on page 745

"[edit protocols mstp]" on page 747

"[edit protocols mstpmsti]" on page 749

"[edit protocols mvpn]" on page 749

"[edit protocols mvrp]" on page 751

"[edit protocols neighbor-discovery]" on page 751

"[edit protocols oam]" on page 752

"[edit protocols openflow]" on page 761

"[edit protocols ospf]" on page 762

"[edit protocols ospf3]" on page 765

"[edit protocols ovsdb]" on page 767

"[edit protocols pcep]" on page 767

"[edit protocols pim]" on page 768

"[edit protocols piminterface]" on page 776

"[edit protocols ppp]" on page 776

"[edit protocols ppp-service]" on page 776

"[edit protocols pppoe]" on page 777

"[edit protocols protection-group]" on page 778

"[edit protocols ptp]" on page 779

"[edit protocols rift]" on page 782

"[edit protocols rip]" on page 783

"[edit protocols ripng]" on page 784

"[edit protocols router-advertisement]" on page 785

"[edit protocols router-discovery]" on page 786

"[edit protocols rstp]" on page 787

"[edit protocols rsvp]" on page 788

"[edit protocols sap]" on page 791

"[edit protocols sflow]" on page 791

"[edit protocols source-packet-routing]" on page 792

"[edit protocols spring-traffic-engineering]" on page 794

"[edit protocols stp]" on page 794

"[edit protocols uplink-failure-detection]" on page 795

"[edit protocols vlan {vlan-name}]" on page 795

"[edit protocols vpls]" on page 796

"[edit protocols vrrp]" on page 796

"[edit protocols vstp]" on page 796

[edit protocols]		
"amt (Protocols)" on page 2230	"ancp" on page 2247	"bfd (protocols)" on page 2911
"bgp" on page 2956	"bgpmcast" on page 2968	"bier" on page 2972
"connections (Circuits)" on page 4542	"connections (MPLS)" on page 4547	"dcbx" on page 4858
"dot1x" on page 5917	"dvmrp" on page 6081	"edge-virtual-bridging" on page 6200
"encapsulation (protocols evpn)" on page 6284	"esi-resolution-per-bridge-domain" on page 6456	"evpn" on page 6541
"flex-algorithm (Protocols OSPF)" on page 7223	"iccp" on page 8250	"igmp" on page 8381

(Continued)

[edit protocols]		
"igmp-snooping" on page 8387	"interconnect" on page 8797	"interconnect-multihoming-peer-gateways" on page 8806
"interconnected-vlan-list" on page 8808	"ipsec-traffic-selector" on page 9321	"isis" on page 9401
"isis-instance" on page 9403	"l2-learning" on page 9524	"l2circuit" on page 9531
"l2pt (Destination Tunnels)" on page 9537	"lacp (Protocols)" on page 9615	"layer2-control" on page 9641
"ldp" on page 9679	"link-management" on page 9798	"lldp" on page 9851
"lldp-med (Ethernet Switching)" on page 9857	"loop-detect" on page 10113	"mld" on page 10778
"mld-snooping" on page 10783	"mpls" on page 10969	"mpls (Protocols)" on page 10984
"msdp" on page 10996	"mstp" on page 11006	"mvpn" on page 11135
"neighbor-discovery" on page 11268	"network-isolation" on page 11312	"oam" on page 11753
"oam (MX Series)" on page 11749	"openflow (Protocols OpenFlow)" on page 11842	"ospf" on page 11938
"ospf-instance" on page 11940	"ospf3" on page 11942	"ovsdb" on page 12086
"pcep" on page 12320	"pim" on page 12507	"protection-group" on page 13413
"ptp" on page 13855	"rift" on page 15794	"rip" on page 15803

(Continued)

[edit protocols]		
"ripng" on page 15804	"router-advertisement" on page 15890	"router-discovery" on page 15900
"rstp" on page 16046	"rsvp" on page 16050	"sap" on page 16303
"sflow" on page 16907	"source-packet-routing (Junos Telemetry Interface)" on page 28261	"source-packet-routing (Segment Routing Traffic Engineering)" on page 28257
"spring-traffic-engineering-edit-protocols" on page 28398	"static-multihoming-peer" on page 28589	"stp" on page 28625
"traffic-engineering (SR-TE import policy)" on page 30032	"uplink-failure-detection" on page 30465	"vstp" on page 30983
"vstp" on page 30983		

[(mstp | rstp| vstp)]

[edit protocols (mstp | rstp| vstp)]	
"bpdu-block-on-edge" on page 3024	

[(bgp | esis | isis | ospf | ospf3 | ldp | pim | rip | ripng | rsvp)]

[edit protocols (bgp | esis | isis | ospf | ospf3 | ldp | pim | rip | ripng | rsvp) graceful-restart]	
"disable (Routing Graceful Restart)" on page 5738	

[(bgp | rip | ripng)]

[edit protocols (bgp | rip | ripng) graceful-restart]	
"restart-time (BGP Graceful Restart)" on page 15663	

[(esis | isis | ospf | ospf3 | pim)]

[edit protocols (esis | isis | ospf | ospf3 | pim) graceful-restart]	
"restart-duration (Protocols)" on page 15651	

[(isis | ldp | ospf | ospf3 | rsvp)]

[edit protocols (isis | ldp | ospf | ospf3 | rsvp) graceful-restart]	
"helper-disable (Multiple Protocols)" on page 8048	

[(mstp | rstp | vstp | vstp vlan {vlan-id})]

[edit protocols (mstp | rstp | vstp | vstp vlan {vlan-id})]	
"traceoptions (Spanning Tree)" on page 29869	

[(mstp | rstp | vstp)]

[edit protocols (mstp | rstp | vstp)]		
"interface (Spanning Tree)" on page 9057	"interface (Spanning Tree)" on page 9057	

[edit protocols (mstp | rstp | vstp) interface {interface-name}]		
"cost" on page 4681	"mode" on page 10798	"no-root-port" on page 11572
"priority (Protocols STP)" on page 13199		

[edit protocols (mstp | rstp | vstp) interface (all | {interface-name})]	
"edge" on page 6197	

[(mstp | rstp)]

[edit protocols (mstp | rstp)]		
"backup-bridge-priority" on page 2793	"bpdu-destination-mac-address (Spanning Tree)" on page 3027	"forward-delay" on page 7393
"hello-time" on page 8041	"priority-hold-time" on page 13177	"system-id" on page 28847
"vpls-flush-on-topology-change" on page 30926		

[(ospf | ospf3)]

[edit protocols (ospf | ospf3)]		
"area" on page 2451	"backup-spf-options (Protocols OSPF)" on page 2823	"database-protection" on page 4852
"export (OSPF)" on page 6660	"graceful-restart (Protocols OSPF)" on page 7775	"import (Protocols OSPF)" on page 8463
"lsa-refresh-interval" on page 10182	"no-nssa-abr" on page 11533	"no-rfc-1583" on page 11571
"overload (Protocols OSPF)" on page 12051	"prefix-export-limit (Protocols OSPF)" on page 13085	"reference-bandwidth (Protocols OSPF)" on page 14226
"rib-group (Protocols OSPF)" on page 15774	"spf-options (Protocols OSPF)" on page 28381	"traceoptions (Protocols OSPF)" on page 29713
"traceoptions (Protocols)" on page 29654	"traffic-engineering (OSPF)" on page 30018	

[edit protocols (ospf | ospf3) area {area-id}]		
"area-range" on page 2453	"context-identifier (Protocols OSPF)" on page 4613	"interface (Protocols OSPF)" on page 8984
"nssa" on page 11717	"stub" on page 28665	"virtual-link" on page 30763

[edit protocols (ospf | ospf3) area {area-id} interface {interface-name}]		
"bandwidth-based-metrics (Protocols OSPF)" on page 2835	"bfd-liveness-detection (Protocols OSPF)" on page 2941	"interface-type (Protocols OSPF)" on page 9091
"link-protection (Protocols OSPF)" on page 9814	"passive (Protocols OSPF)" on page 12223	

[edit protocols (ospf | ospf3) area {area-id} interface {interface-name} passive]	
"traffic-engineering (Passive TE Mode)" on page 30021	

[edit protocols (ospf | ospf3) area {area-id} nssa]		
"area-range" on page 2453	"default-lsa" on page 5197	

[edit protocols (ospf | ospf3) source-packet-routing]	
"sensor-based-stats" on page 16555	

[edit protocols (ospf | ospf3) traffic-engineering]		
"shortcuts (Protocols OSPF)" on page 16991	"tunnel-source-protocol (Protocols OSPF)" on page 30240	

[(rstp | vstp)]

[edit protocols (rstp | vstp)]	
"force-version (IEEE 802.1D STP)" on page 7379	

[amt]

[edit protocols amt]		
"relay (AMT Protocol)" on page 14291	"traceoptions (Protocols AMT)" on page 29656	

[edit protocols amt relay]		
"accounting (Protocols AMT Interface)" on page 1728	"family (Protocols AMT Relay)" on page 6898	"secret-key-timeout" on page 16423
"tunnel-devices (Protocols AMT)" on page 30196	"tunnel-limit (Protocols AMT)" on page 30212	"unicast-stream-limit (Protocols AMT)" on page 30351

[edit protocols amt relay family]	
"inet (AMT Protocol)" on page 8555	

[edit protocols amt relay family inet]		
"anycast-prefix" on page 2285	"local-address (Protocols AMT)" on page 9904	

[ancp]

[edit protocols ancp]		
"adjacency-loss-hold-time (ANCP)" on page 1943	"adjacency-timer" on page 1944	"gsmp-syn-timeout (ANCP)" on page 7928
"gsmp-syn-wait (ANCP)" on page 7929	"interfaces (ANCP)" on page 9121	"maximum-discovery-table-entries" on page 10519
"maximum-helper-restart-time" on page 10524	"neighbor (Define ANCP)" on page 11266	"pre-ietf-mode" on page 12998
"qos-adjust" on page 13888	"qos-adjust-adsl" on page 13890	"qos-adjust-adsl2" on page 13892
"qos-adjust-adsl2-plus" on page 13894	"qos-adjust-other" on page 13897	"qos-adjust-sdsl" on page 13899
"qos-adjust-vdsl" on page 13901	"qos-adjust-vdsl2" on page 13903	"traceoptions (ANCP)" on page 29494

[edit protocols ancp interfaces {interface-name}]		
"access-identifier" on page 1661	"interface-set (ANCP)" on page 9036	"overhead-accounting (ANCP)" on page 12038

[edit protocols ancp interfaces interface-set]		
"access-identifier" on page 1661	"underlying-interface (ANCP)" on page 30332	

[edit protocols ancp neighbor {ip-address}]		
"adjacency-loss-hold-time (ANCP)" on page 1943	"adjacency-timer" on page 1944	"auto-configure-trigger interface (ANCP)" on page 2702

(Continued)

[edit protocols ancp neighbor {ip-address}]		
"ietf-mode" on page 8367	"maximum-discovery-table-entries" on page 10519	"pre-ietf-mode" on page 12998

[edit protocols ancp qos-adjust]		
"adsl-bytes" on page 1982	"adsl2-bytes" on page 1984	"adsl2-plus-bytes" on page 1986
"other-bytes" on page 11949	"other-overhead-adjust" on page 11951	"sdsl-bytes" on page 16398
"sdsl-overhead-adjust" on page 16400	"vdsl-bytes" on page 30665	"vdsl-overhead-adjust" on page 30667
"vdsl2-bytes" on page 30671	"vdsl2-overhead-adjust" on page 30673	

[authentication-access-control]

[edit protocols authentication-access-control]	
"interface (Static MAC Bypass)" on page 9067	

[edit protocols authentication-access-control interface (all | [{interface-names}])]	
"session-expiry (MX Series in Enhanced LAN Mode)" on page 16809	

[edit protocols authentication-access-control interface (all | {interface-names})]	
"dot1x (MX Series in Enhanced LAN Mode)" on page 5920	

[bfd]

[edit protocols bfd]		
"no-issu-timer-negotiation" on page 11487	"traceoptions (Protocols BFD)" on page 29659	

[edit protocols bfd sbfd]	
"sbfd (protocols bfd)" on page 16324	

[bgp]

[edit protocols bgp]		
"accept-prpd-connection" on page 1639	"accept-remote-nexthop" on page 1640	"add-path-display-ipv4-address" on page 1842
"advertise-bgp-static" on page 2009	"advertise-from-main-vpn-tables" on page 2013	"advertise-inactive" on page 2015
"advertise-peer-as" on page 2020	"advertise-prefix-sid" on page 2022	"authentication-algorithm" on page 2577

(Continued)

[edit protocols bgp]		
"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605	"bfd-liveness-detection (BGP)" on page 2918
"bgp-error-tolerance (Protocols BGP)" on page 2959	"bgp-identifier" on page 2962	"bmp" on page 3000
"cluster" on page 4367	"damping (Protocols BGP)" on page 4822	"damping (Protocols BGP)" on page 4822
"defaults (ACX series, JRR200, MX204, PTX5000, and VRR)" on page 5225	"description (Protocols BGP)" on page 5340	"disable (Protocols BGP)" on page 5788
"disable-4byte-as" on page 5742	"domain-path-id (Protocols BGP)" on page 5898	"egress-te" on page 6227
"egress-te-backup-paths" on page 6232	"egress-te-set-segment" on page 6236	"enforce-first-as" on page 6343
"export (Protocols BGP)" on page 6677	"family (Protocols BGP)" on page 6900	"forwarding-context (Protocols BGP)" on page 7453
"graceful-restart (Long-Lived for BGP Helper)" on page 7763	"graceful-restart (Protocols BGP)" on page 7769	"graceful-shutdown (Protocols BGP)" on page 7784
"group (Protocols BGP)" on page 7863	"hold-time (Protocols BGP)" on page 8104	"idle-after-switch-over" on page 8300
"import (BGP)" on page 8461	"include-mp-next-hop" on page 8532	"ipsec-sa (Protocols BGP)" on page 9308
"keep" on page 9452	"local-address (Protocols BGP)" on page 9905	"local-as" on page 9916

(Continued)

[edit protocols bgp]		
"local-preference" on page 9971	"log-updown" on page 10042	"metric-out" on page 10676
"minimum-hold-time" on page 10731	"mtu-discovery" on page 11032	"multihop" on page 11098
"multipath (Protocols BGP)" on page 11110	"multipath-build-priority" on page 11107	"no advertise-peer-as" on page 11385
"no-aggregator-id" on page 11387	"no-client-reflect" on page 11424	"out-delay" on page 11961
"outbound-route-filter" on page 11970	"output-queue-priority" on page 12008	"passive (Protocols BGP)" on page 12218
"path-selection" on page 12266	"peer-as (Protocols BGP)" on page 12346	"precision-timers" on page 13010
"preference (Protocols BGP)" on page 13038	"remove-private" on page 14368	"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763
"shutdown (Protocols BGP)" on page 27936	"snmp-options" on page 28067	"stale-labels-holddown-period" on page 28476
"tcp-aggressive-transmission" on page 28978	"tcp-mss (Protocols BGP)" on page 28998	"tcpao-auth-mismatch" on page 29034
"traceoptions (Protocols BGP)" on page 29662	"traffic-statistics-labeled-path" on page 30049	"unconfigured-peer-graceful-restart" on page 30329
"vpn-apply-export" on page 30938		

[edit protocols bgp bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit protocols bgp bmp]	
"route-monitoring" on page 15871	

[edit protocols bgp bmp route-monitoring]	
"post-policy" on page 12910	

[edit protocols bgp ebgp-community-cleanup]	
"origin-validation" on page 11937	

[edit protocols bgp family{name}]		
"delay-route-advertisements" on page 5262	"inet-bgpmcast" on page 8557	"inet-mdt (Signaling)" on page 8571
"inet-mvpn (BGP)" on page 8574	"inet6-bgpmcast" on page 8599	"inet6-mvpn (BGP)" on page 8610
"inet6-vpn" on page 8620	"iso-vpn" on page 9404	"local-ipv4-address" on page 9955
"loops (BGP Address Family)" on page 10134	"output-queue-priority" on page 12008	"route-refresh-priority" on page 15876
"route-target (Protocols BGP)" on page 15881	"traffic-engineering (Protocols BGP)" on page 30023	"withdraw-priority" on page 31134

[edit protocols bgp family (inet | inet-vpn | inet6 | inet6-vpn) flow]	
"strip-nexthop" on page 28658	

[edit protocols bgp family (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) (any | flow | labeled-unicast | multicast | unicast)]	
"no-install" on page 11475	

[edit protocols bgp family (inet | inet6)]	
"labeled-unicast (Protocols BGP)" on page 9603	

[edit protocols bgp family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651	"prefix-limit" on page 13087	

[edit protocols bgp family (inet | inet6) labeled-unicast]		
"rib (Protocols BGP)" on page 15787	"traffic-statistics (Protocols BGP)" on page 30051	

[edit protocols bgp family (inet | inet6) unicast]	
"topology (Protocols BGP)" on page 29449	

[edit protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]	
"restarter (Graceful Restart for BGP Restarter)" on page 15667	

[edit protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]	
"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481	

[edit protocols bgp family accepted-prefix-limit (inet | inet6)(any | flow | labeled-unicast | multicast | unicast)]	
"accepted-prefix-limit" on page 1651	

[edit protocols bgp family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit protocols bgp family inet]	
"labeled-unicast" on page 9601	

[edit protocols bgp family inet (labeled-unicast | flow) graceful-restart long-lived restarter]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit protocols bgp family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit protocols bgp family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit protocols bgp family inet labeled-unicast]		
"aggregate-label" on page 2066	"aigp" on page 2092	"entropy-label (protocols bgp)" on page 6413

(Continued)

[edit protocols bgp family inet labeled-unicast]		
"explicit-null (Protocols BGP)" on page 6652	"per-group-label" on page 12416	"per-prefix-label" on page 12419
"resolve-vpn" on page 15602		

[edit protocols bgp family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229
"extended-nexthop" on page 6700		

[edit protocols bgp family inet-vpn unicast]		
"accept-own" on page 1637	"aggregate-label" on page 2066	

[edit protocols bgp family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit protocols bgp family inet6 labeled-unicast]		
"aggregate-label" on page 2066	"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652

[edit protocols bgp family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit protocols bgp family inet6-vpn unicast]		
"accept-own" on page 1637	"aggregate-label" on page 2066	

[edit protocols bgp family l2vpn]	
"auto-discovery-only" on page 2711	

[edit protocols bgp family prefix-limit (inet | inet6)(any | flow | labeled-unicast | multicast | unicast)]		
"family (Protocols BGP)" on page 6900	"group (Protocols BGP)" on page 7863	"inet-mdt (Signaling)" on page 8571
"labeled-unicast (Protocols BGP)" on page 9603	"neighbor (Protocols BGP)" on page 11273	"prefix-limit" on page 13087
"route-target (Protocols BGP)" on page 15881		

[edit protocols bgp family route-target]	
"accepted-prefix-limit" on page 1651	

[edit protocols bgp graceful-restart]		
"disable (BGP Graceful Restart)" on page 5747	"disable-notification-extensions (BGP Graceful Restart)" on page 5779	"disable-notification-flag (BGP Graceful Restart)" on page 5780
"dont-help-shared-fate-bfd-down" on page 5912	"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488	"long-lived (Graceful Restart for BGP Helper)" on page 10103
"stale-routes-time" on page 28478		

[edit protocols bgp graceful-restart long-lived]		
"advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper)" on page 2026	"receiver (Graceful Restart for BGP Helper)" on page 14127	

[edit protocols bgp graceful-restart long-lived advertise-to-non-llgr-neighbor]	
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810	

[edit protocols bgp group {group-name}]		
"accept-remote-nexthop" on page 1640	"add-path" on page 1839	"add-path-display-ipv4-address" on page 1842
"advertise-bgp-static" on page 2009	"advertise-external" on page 2011	"advertise-inactive" on page 2015
"advertise-peer-as" on page 2020	"allow" on page 2160	"as-override" on page 2489
"authentication-algorithm" on page 2577	"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605
"bfd-liveness-detection (BGP)" on page 2918	"bgp-error-tolerance (Protocols BGP)" on page 2959	"bgp-identifier" on page 2962
"bmp" on page 3000	"cluster" on page 4367	"damping (Protocols BGP)" on page 4822
"description (Protocols BGP)" on page 5340	"disable-4byte-as" on page 5742	"domain-path-id (Protocols BGP)" on page 5898
"egress-te" on page 6227	"egress-te-backup-paths" on page 6232	"enforce-first-as" on page 6343

(Continued)

[edit protocols bgp group {group-name}]		
"export (Protocols BGP)" on page 6677	"family (Protocols BGP)" on page 6900	"family route-target" on page 6911
"graceful-restart (Long-Lived for BGP Helper)" on page 7763	"graceful-restart (Protocols BGP)" on page 7769	"graceful-shutdown (Protocols BGP)" on page 7784
"hold-time (Protocols BGP)" on page 8104	"idle-after-switch-over" on page 8300	"import (BGP)" on page 8461
"include-mp-next-hop" on page 8532	"ipsec-sa (Protocols BGP)" on page 9308	"keep" on page 9452
"local-address (Protocols BGP)" on page 9905	"local-as" on page 9916	"local-preference" on page 9971
"log-updown" on page 10042	"metric-out" on page 10676	"minimum-hold-time" on page 10731
"mtu-discovery" on page 11032	"multihop" on page 11098	"multipath (Protocols BGP)" on page 11110
"mvpn-iana-rt-import" on page 11140	"neighbor (Protocols BGP)" on page 11273	"no advertise-peer-as" on page 11385
"no-aggregator-id" on page 11387	"no-client-reflect" on page 11424	"optimal-route-reflection" on page 11850
"out-delay" on page 11961	"outbound-route-filter" on page 11970	"passive (Protocols BGP)" on page 12218
"peer-as (Protocols BGP)" on page 12346	"peer-as-list" on page 12345	"preference (Protocols BGP)" on page 13038
"remove-private" on page 14368	"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763	"shutdown (Protocols BGP)" on page 27936

(Continued)

[edit protocols bgp group {group-name}]		
"tcp-aggressive-transmission" on page 28978	"tcp-mss (Protocols BGP)" on page 28998	"traceoptions (Protocols BGP)" on page 29662
"type (Protocols BGP)" on page 30278	"vpn-apply-export" on page 30938	

[edit protocols bgp group <group-name> ebgp-community-cleanup]	
"origin-validation" on page 11937	

[edit protocols bgp group <group-name> neighbor <address> ebgp-community-cleanup]	
"origin-validation" on page 11937	

[edit protocols bgp group {group-name} {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit protocols bgp group {name} bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit protocols bgp group {group-name} bmp]	
"route-monitoring" on page 15871	

[edit protocols bgp group {group-name} bmp route-monitoring]	
"post-policy" on page 12910	

[edit protocols bgp group {name} dynamic-neighbor {dyn-name}]		
"allow" on page 2160	"peer-limit" on page 12375	

[edit protocols bgp group {group-name} family {family}]		
"damping (Protocols BGP)" on page 4822	"delay-route-advertisements" on page 5262	"inet-bgpmcast" on page 8557
"inet-mvpn (BGP)" on page 8574	"inet6-bgpmcast" on page 8599	"inet6-mvpn (BGP)" on page 8610
"inet6-vpn" on page 8620	"iso-vpn" on page 9404	"local-ipv4-address" on page 9955
"loops (BGP Address Family)" on page 10134	"route-refresh-priority" on page 15876	"route-target (Protocols BGP)" on page 15881
"traffic-engineering (Protocols BGP)" on page 30023	"withdraw-priority" on page 31134	

[edit protocols bgp group {group-name} family (inet | inet flow)]	
"no-validate" on page 11615	

[edit protocols bgp group {group-name} family (inet | inet-vpn | inet6 | inet6-vpn)]	
"flow (protocols bgp)" on page 7250	

[edit protocols bgp group {group-name} family (inet | inet-vpn | inet6 | inet6-vpn) flow]	
"strip-nexthop" on page 28658	

[edit protocols bgp group {group-name} family (inet | inet-vpn) flow]		
"legacy-redirect-ip-action" on page 9711		

[edit protocols bgp group {group-name} family (inet | inet6)]		
"labeled-unicast (Protocols BGP)" on page 9603		

[edit protocols bgp group {group-name} family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651		

[edit protocols bgp group {group-name} family (inet | inet6) (any | labeled-unicast | multicast | unicast)]		
"prefix-limit" on page 13087		

[edit protocols bgp group {group-name} family (inet | inet6) labeled-unicast]		
"rib (Protocols BGP)" on page 15787	"traffic-statistics (Protocols BGP)" on page 30051	

[edit protocols bgp group {group-name} family (inet | inet6) unicast]		
"topology (Protocols BGP)" on page 29449		

[edit protocols bgp group {group-name} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]		
"restarter (Graceful Restart for BGP Restarter)" on page 15667		

[edit protocols bgp group {group-name} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]	
"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481	

[edit protocols bgp group {group-name} family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit protocols bgp group {group-name} family inet]	
"labeled-unicast" on page 9601	

[edit protocols bgp group {group-name} family inet (labeled-unicast | flow) graceful-restart long-lived restarter]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit protocols bgp group {group-name} family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit protocols bgp group {name} family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit protocols bgp group {group-name} family inet labeled-unicast]		
"aigp" on page 2092	"egress-protection (BGP)" on page 6216	"entropy-label (protocols bgp)" on page 6413
"explicit-null (Protocols BGP)" on page 6652	"per-group-label" on page 12416	"per-prefix-label" on page 12419
"protection (Protocols MPLS)" on page 13417	"resolve-vpn" on page 15602	

[edit protocols bgp group {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229
"extended-nexthop" on page 6700	"extended-nexthop-tunnel" on page 6701	

[edit protocols bgp group {group-name} family inet-vpn unicast]	
"egress-protection (BGP)" on page 6216	

[edit protocols bgp group {name} family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit protocols bgp group {group-name} family inet6 labeled-unicast]		
"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652	

[edit protocols bgp group {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"extended-nexthop-tunnel" on page 6701

[edit protocols bgp group {group-name} family inet6-vpn unicast]	
"egress-protection (BGP)" on page 6216	

[edit protocols bgp group {group-name} family iso-vpn unicast]	
"egress-protection (BGP)" on page 6216	

[edit protocols bgp group {group-name} family l2vpn]	
"auto-discovery-only" on page 2711	

[edit protocols bgp group {group-name} family route-target]		
"accepted-prefix-limit" on page 1651	"proxy-generate" on page 13826	

[edit protocols bgp group {group-name} graceful-restart]		
"disable (BGP Graceful Restart)" on page 5747	"disable (Routing Graceful Restart)" on page 5738	"disable-notification-extensions (BGP Graceful Restart)" on page 5779
"disable-notification-flag (BGP Graceful Restart)" on page 5780	"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488	"long-lived (Graceful Restart for BGP Helper)" on page 10103

[edit protocols bgp group {group-name} graceful-restart long-lived]		
"advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper)" on page 2026	"receiver (Graceful Restart for BGP Helper)" on page 14127	

[edit protocols bgp group {group-name} graceful-restart long-lived advertise-to-non-llgr-neighbor]	
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810	

[edit protocols bgp group group family inet6-vpn unicast]	
"accept-own" on page 1637	

[edit protocols bgp group group family inetvpn unicast]	
"accept-own" on page 1637	

[edit protocols bgp group group-name family]	
"inet-mdt (Signaling)" on page 8571	

[edit protocols bgp group group-name neighbor address family]	
"inet-mdt (Signaling)" on page 8571	

[edit protocols bgp group {name} link-bandwidth]	
"auto-sense" on page 2738	

[edit protocols bgp group {group-name} local-as]		
"loops (Autonomous System)" on page 10132		

[edit protocols bgp group {name} metric-out]		
"effective-aigp" on page 6203	"minimum-effective-aigp" on page 10730	

[edit protocols bgp group name optimal-route-reflection]		
"export (Protocols BGP)" on page 6677		

[edit protocols bgp group {group-name} neighbor {address}]		
"accept-remote-nexthop" on page 1640	"add-path" on page 1839	"advertise-bgp-static" on page 2009
"advertise-external" on page 2011	"advertise-inactive" on page 2015	"advertise-peer-as" on page 2020
"as-override" on page 2489	"authentication-algorithm" on page 2577	"authentication-key (Protocols BGP and BMP)" on page 2616
"authentication-key-chain (Protocols BGP and BMP)" on page 2605	"bfd-liveness-detection (BGP)" on page 2918	"bgp-error-tolerance (Protocols BGP)" on page 2959
"bgp-identifier" on page 2962	"bmp" on page 3000	"cluster" on page 4367
"damping (Protocols BGP)" on page 4822	"description (Protocols BGP)" on page 5340	"disable-4byte-as" on page 5742
"domain-path-id (Protocols BGP)" on page 5898	"egress-te" on page 6227	"egress-te-adj-segment" on page 6229

(Continued)

[edit protocols bgp group {group-name} neighbor {address}]		
"egress-te-node-segment" on page 6234	"enforce-first-as" on page 6343	"export (Protocols BGP)" on page 6677
"family (Protocols BGP)" on page 6900	"family route-target" on page 6911	"graceful-restart (Long-Lived for BGP Helper)" on page 7763
"graceful-restart (Protocols BGP)" on page 7769	"graceful-shutdown (Protocols BGP)" on page 7784	"hold-time (Protocols BGP)" on page 8104
"idle-after-switch-over" on page 8300	"import (BGP)" on page 8461	"include-mp-next-hop" on page 8532
"ipsec-sa (Protocols BGP)" on page 9308	"keep" on page 9452	"local-address (Protocols BGP)" on page 9905
"local-as" on page 9916	"local-interface (IPv6)" on page 9952	"local-preference" on page 9971
"log-updown" on page 10042	"metric-out" on page 10676	"minimum-hold-time" on page 10731
"mtu-discovery" on page 11032	"multihop" on page 11098	"multipath (Protocols BGP)" on page 11110
"no advertise-peer-as" on page 11385	"no-aggregator-id" on page 11387	"no-client-reflect" on page 11424
"out-delay" on page 11961	"outbound-route-filter" on page 11970	"passive (Protocols BGP)" on page 12218
"peer-as (Protocols BGP)" on page 12346	"preference (Protocols BGP)" on page 13038	"remove-private" on page 14368
"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763	"shutdown (Protocols BGP)" on page 27936	"tcp-aggressive-transmission" on page 28978

(Continued)

[edit protocols bgp group {group-name} neighbor {address}]		
"tcp-mss (Protocols BGP)" on page 28998	"tcpao-auth-mismatch" on page 29034	"traceoptions (Protocols BGP)" on page 29662
"vpn-apply-export" on page 30938		

[edit protocols bgp group {group-name} neighbor {address} {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit protocols bgp group {name} neighbor {address} bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit protocols bgp group {group-name} neighbor {address} bmp]	
"route-monitoring" on page 15871	

[edit protocols bgp group {group-name} neighbor {address} family {family}]		
"damping (Protocols BGP)" on page 4822	"delay-route-advertisements" on page 5262	"local-ipv4-address" on page 9955
"loops (BGP Address Family)" on page 10134	"output-queue-priority" on page 12008	"route-refresh-priority" on page 15876
"route-target (Protocols BGP)" on page 15881	"traffic-engineering (Protocols BGP)" on page 30023	"withdraw-priority" on page 31134

[edit protocols bgp group {group-name} neighbor {address} family (inet | inet flow)]	
"no-validate" on page 11615	

[edit protocols bgp group {group-name} neighbor {address} family (inet | inet-vpn | inet6 | inet6-vpn flow)]	
"strip-nexthop" on page 28658	

[edit protocols bgp group {group-name} neighbor {address} family (inet | inet-vpn | inet6 | inet6-vpn)]	
"flow (protocols bgp)" on page 7250	

[edit protocols bgp group {group-name} neighbor {address} family (inet | inet-vpn) flow]	
"legacy-redirect-ip-action" on page 9711	

[edit protocols bgp group {group-name} neighbor {address} family (inet | inet6)]	
"labeled-unicast (Protocols BGP)" on page 9603	

[edit protocols bgp group {group-name} neighbor {address} family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651	"prefix-limit" on page 13087	

[edit protocols bgp group {group-name} neighbor {address} family (inet | inet6) labeled-unicast]	
"rib (Protocols BGP)" on page 15787	

[edit protocols bgp group {group-name} neighbor {address} family (inet | inet6) unicast]	
"topology (Protocols BGP)" on page 29449	

[edit protocols bgp group {group-name} neighbor {address} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]	
"restarter (Graceful Restart for BGP Restarter)" on page 15667	

[edit protocols bgp group {group-name} neighbor {address} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]	
"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481	

[edit protocols bgp group {group-name} neighbor {address} family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit protocols bgp group {group-name} neighbor {address} family inet (labeled-unicast | flow) graceful-restart long-lived restarter]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit protocols bgp group {group-name} neighbor {address} family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit protocols bgp group {name} neighbor {name} family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit protocols bgp group {group-name} neighbor {address} family inet labeled-unicast]		
"aigp" on page 2092	"entropy-label (protocols bgp)" on page 6413	"explicit-null (Protocols BGP)" on page 6652
"resolve-vpn" on page 15602		

[edit protocols bgp group {name} neighbor {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229
"extended-nexthop" on page 6700		

[edit protocols bgp group {group} neighbor {address} family inet-vpn unicast]	
"accept-own" on page 1637	

[edit protocols bgp group {name} neighbor {name} family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit protocols bgp group {group-name} neighbor {address} family inet6 labeled-unicast]		
"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652	

[edit protocols bgp group {name} neighbor {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit protocols bgp group {group} neighbor {address} family inet6-vpn unicast]	
"accept-own" on page 1637	

[edit protocols bgp group {group-name} neighbor {address} family l2vpn]	
"auto-discovery-only" on page 2711	

[edit protocols bgp group {group-name} neighbor {address} family route-target]		
"accepted-prefix-limit" on page 1651	"proxy-generate" on page 13826	

[edit protocols bgp group {group-name} neighbor {address} graceful-restart]		
"disable (BGP Graceful Restart)" on page 5747	"disable (Routing Graceful Restart)" on page 5738	"disable-notification-extensions (BGP Graceful Restart)" on page 5779
"disable-notification-flag (BGP Graceful Restart)" on page 5780	"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488	"long-lived (Graceful Restart for BGP Helper)" on page 10103

[edit protocols bgp group {group-name} neighbor {address} graceful-restart long-lived]		
"advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper)" on page 2026	"receiver (Graceful Restart for BGP Helper)" on page 14127	

[edit protocols bgp group {group-name} neighbor {address} graceful-restart long-lived advertise-to-non-llgr-neighbor]	
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810	

[edit protocols bgp group {group-name} neighbor {address} local-as]	
"loops (Autonomous System)" on page 10132	

[edit protocols bgp group neighbor {group-name} neighbor {address} bmp route-monitoring]	
"post-policy" on page 12910	

[edit protocols bgp group {group-name} neighbor {address} outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit protocols bgp group {group-name} neighbor {address}family]	
"iso-vpn" on page 9404	

[edit protocols bgp group {group-name} outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit protocols bgp local-as]	
"loops (Autonomous System)" on page 10132	

[edit protocols bgp multipath]		
"pause-computation-during-churn" on page 12277		

[edit protocols bgp neighbor {neighbor-address} family]		
"inet-bgpmcast" on page 8557	"inet6-bgpmcast" on page 8599	

[edit protocols bgp outbound-route-filter]		
"bgp-orf-cisco-mode" on page 2963		

[edit protocols bgp path-selection]		
"l2vpn-use-bgp-rules-edit-protocols-bgp" on page 9565		

[edit protocols bgp source-packet-routing]		
"srgb (Protocols BGP)" on page 28407	"srv6 (Protocols BGP)" on page 28417	

[connections]

[edit protocols connections]		
"interface-switch" on page 9076	"lsp-switch" on page 10194	"p2mp-receive-switch" on page 12125
"p2mp-transmit-switch" on page 12127		

[edit protocols connections (MPLS)]	
"remote-interface-switch" on page 14342	

[edit protocols connections p2mp-transmit-switch {p2mp-transmit-switch-name}]	
"output-interface (CCC)" on page 11994	

[dcbx]

[edit protocols dcbx]		
"disable (DCBX)" on page 5749	"interface (DCBX)" on page 8851	

[edit protocols dcbx interface {interface-name}]		
"application-map" on page 2345	"applications (DCBX)" on page 2406	"disable (DCBX)" on page 5749
"enhanced-transmission-selection" on page 6406		

[edit protocols dcbx interface (all | {interface-name})]		
"dcbx-version" on page 4860	"priority-flow-control" on page 13176	

[edit protocols dcbx interface {interface-name} applications]	
"fcoe" on page 6962	

[edit protocols dcbx interface {interface-name} enhanced-transmission-selection]		
"no-recommendation-tlv" on page 11557	"recommendation-tlv" on page 14129	

[dot1x]

[edit protocols dot1x]		
"authenticator" on page 2673	"supplicant" on page 28726	"traceoptions (802.1X)" on page 29474

[edit protocols dot1x authenticator]		
"interface (802.1X)" on page 8819	"radius-reachability" on page 14015	"static (802.1X)" on page 28526

[edit protocols dot1x authenticator interface {interface-name}]		
"authentication-order (Authenticator)" on page 2638	"mac-radius" on page 10257	"multi-domain" on page 11048

[edit protocols dot1x authenticator interface (all | [{interface-names}])]	
"server-reject-bridge-domain | server-reject-vlan" on page 16661	

[edit protocols dot1x authenticator interface (all | interface-names)]	
"eapol-block" on page 6183	

[edit protocols dot1x authenticator interface (all | interface-names) (server-reject-bridge-domain | server-reject-vlan)]	
"eapol-block" on page 6183	

[edit protocols dot1x authenticator interface {interface-name} mac-radius]	
"authentication-protocol" on page 2645	

[edit protocols dot1x authenticator mac-radius]	
"password (MAC RADIUS Authentication)" on page 12241	

[dvmrp]

[edit protocols dvmrp]		
"disable (Protocols)" on page 5731	"export (Protocols DVMRP)" on page 6678	"import (Protocols DVMRP)" on page 8474
"interface (Protocols DVMRP)" on page 8960	"rib-group (Protocols DVMRP)" on page 15769	"traceoptions (Protocols DVMRP)" on page 29669

[edit protocols dvmrp interface {interface-name}]		
"disable (Protocols)" on page 5731	"hold-time (Protocols DVMRP)" on page 8106	"metric (Protocols DVMRP)" on page 10682
"mode (Protocols DVMRP)" on page 10813		

[edge-virtual-bridging]

[edit protocols edge-virtual-bridging]		
"traceoptions (Ethernet Switching)" on page 29465	"vsi-discovery" on page 30980	

[ethernet]

[edit protocols ethernet action-profile {name} event]	
"symbol-period" on page 28767	

[edit protocols ethernet event]	
"protocol-down" on page 13436	

[edit protocols ethernet interface {interface-name}]		
"remote-loopback" on page 14350	"symbol-period" on page 28767	

[evpn]

[edit protocols evpn]		
"assisted-replication" on page 2505	"default-gateway" on page 5190	"designated-forwarder-election-hold-time" on page 5363
"duplicate-mac-detection" on page 6068	"encapsulation (Logical Interface)" on page 6300	"encapsulation (protocols evpn)" on page 6284

(Continued)

[edit protocols evpn]		
"multicast-mode (EVPN)" on page 11074	"node-detection (EVPN-VXLAN)" on page 11633	"remote-ip-host-routes" on page 14346
"vni-options" on page 30912		

[edit protocols evpn assisted-replication]		
"leaf" on page 9689	"replicator" on page 14387	

[edit protocols evpn interconnect]		
"esi" on page 6453	"interconnected-vni-list" on page 8809	"route-distinguisher" on page 15856
"vrf-export" on page 30958	"vrf-import" on page 30960	"vrf-target" on page 30968

[edit protocols evpn interconnect esi]	
"df-election-type" on page 5576	

[edit protocols evpn node-detection]	
"bfd-liveness-detection (EVPN Node Detection)" on page 2923	

[group]

[edit protocols group]		
"detection (Network Isolation Groups)" on page 5540	"service-tracking-action (Network Isolation Groups)" on page 16758	

[edit protocols group {name} detection]		
"link-tracking (Network Isolation Groups)" on page 9833	"service-tracking (Network Isolation Groups)" on page 16760	

[iccp]

[edit protocols iccp]		
"authentication-key (ICCP)" on page 2614	"iccp (Traceoptions)" on page 8247	"local-ip-addr (ICCP)" on page 9953
"peer (ICCP)" on page 12366	"session-establishment-hold-time" on page 16808	

[edit protocols iccp peer {peer-IP-address}]		
"authentication-key (ICCP)" on page 2614	"backup-liveness-detection" on page 2799	"liveness-detection (iccp-peer)" on page 9845
"local-ip-addr (ICCP)" on page 9953	"session-establishment-hold-time" on page 16808	

[edit protocols iccp peer backup-liveness-detection]	
"backup-peer-ip" on page 2810	

[edit protocols iccp peer liveness-detection]		
"detection-time (Liveness Detection)" on page 5550	"minimum-interval (Liveness Detection)" on page 10737	"minimum-receive-interval (Liveness Detection)" on page 10756
"multiplier (Liveness Detection)" on page 11125	"no-adaptation (Liveness Detection)" on page 11376	"transmit-interval (Liveness Detection)" on page 30092
"version (Liveness Detection)" on page 30710		

[edit protocols iccp peer liveness-detection detection-time]	
"threshold (Detection Time)" on page 29322	

[igmp]

[edit protocols igmp]		
"accounting (Protocols IGMP)" on page 1730	"amt (IGMP)" on page 2228	"interface (Protocols IGMP)" on page 8962
"maximum-transmit-rate (Protocols IGMP)" on page 10572	"query-interval (Protocols IGMP)" on page 13929	"query-last-member-interval (Protocols IGMP)" on page 13935
"query-response-interval (Protocols IGMP)" on page 13940	"robust-count (Protocols IGMP)" on page 15818	"traceoptions (Protocols IGMP)" on page 29675

[edit protocols igmp amt]	
"relay (IGMP)" on page 14292	

[edit protocols igmp amt relay]	
"defaults" on page 5224	

[edit protocols igmp amt relay defaults]		
"accounting (Protocols IGMP AMT Interface)" on page 1731	"group-policy (Protocols IGMP AMT Interface)" on page 7856	"query-interval (Protocols IGMP AMT)" on page 13930
"query-response-interval (Protocols IGMP AMT)" on page 13941	"robust-count (Protocols IGMP AMT)" on page 15820	"ssm-map (Protocols IGMP AMT)" on page 28465
"version (Protocols IGMP AMT)" on page 30716		

[edit protocols igmp interface {interface-name}]		
"accounting (Protocols IGMP Interface)" on page 1732	"disable (Protocols)" on page 5731	"distributed (IGMP)" on page 5834
"group-limit (IGMP) " on page 7844	"group-policy (Protocols IGMP)" on page 7855	"group-threshold (Protocols IGMP Interface)" on page 7903
"immediate-leave" on page 8454	"log-interval (IGMP Interface)" on page 10009	"oif-map (IGMP Interface) " on page 11803
"passive (IGMP)" on page 12202	"promiscuous-mode (Protocols IGMP)" on page 13375	"ssm-map (Protocols IGMP)" on page 28464
"ssm-map-policy (IGMP)" on page 28461	"static (Protocols IGMP)" on page 28557	"version (Protocols IGMP)" on page 30715

[edit protocols igmp interface {interface-name} static]	
"group (Protocols IGMP)" on page 7869	

[edit protocols igmp interface {interface-name} static group {multicast-group-address}]		
"exclude (Protocols IGMP)" on page 6614	"group-count (Protocols IGMP)" on page 7810	"group-increment (Protocols IGMP)" on page 7838
"source (Protocols IGMP)" on page 28305		

[edit protocols igmp interface {interface-name} static group {multicast-group-address} source]		
"source-count (Protocols IGMP)" on page 28197	"source-increment (Protocols IGMP)" on page 28221	

[igmp-snooping]

[edit protocols igmp-snooping]		
"traceoptions (IGMP Snooping)" on page 29573	"version (IGMP Snooping)" on page 30704	"vlan (IGMP Snooping)" on page 30821

[edit protocols igmp-snooping vlan {vlan-name}]		
"data-forwarding" on page 4835	"disable (IGMP Snooping)" on page 5764	"disable (Protocols)" on page 5731
"evpn-ssm-reports-only" on page 6553	"igmp-querier (QFabric Systems only)" on page 8385	"l2-querier" on page 9527
"qualified-vlan " on page 13912	"query-interval (Bridge Domains)" on page 13927	"query-last-member-interval (Bridge Domains)" on page 13933
"query-response-interval (Bridge Domains)" on page 13938	"robust-count (IGMP Snooping)" on page 15815	"traceoptions (Protocols IGMP Snooping)" on page 29678

(Continued)

[edit protocols igmp-snooping vlan {vlan-name}]		
"version (IGMP Snooping)" on page 30704		

[edit protocols igmp-snooping vlan (all | {vlan-name}) interface {interface-name}]		
"immediate-leave" on page 8454	"static (IGMP Snooping)" on page 28537	

[edit protocols igmp-snooping vlan (all | {vlan-name})]		
"immediate-leave" on page 8454	"interface (IGMP Snooping)" on page 8902	"proxy (Multicast VLAN Registration)" on page 13834
"version (Protocols IGMP)" on page 30715		

[edit protocols igmp-snooping vlan (all | {vlan-name}) interface (all | {interface-name})]	
"multicast-router-interface (IGMP Snooping)" on page 11084	

[edit protocols igmp-snooping vlan (all | {vlan-name}) interface (all | {interface-name}) static]	
"group (IGMP Snooping)" on page 7835	

[edit protocols igmp-snooping vlan ({vlan-name}) data-forwarding]	
"receiver" on page 14124	

[edit protocols igmp-snooping vlan {vlan-name} data-forwarding]	
"source (Multicast VLAN Registration)" on page 28252	

[edit protocols igmp-snooping vlan {vlan-name} interface {interface-name}]	
"multicast-router-interface (IGMP Snooping)" on page 11084	

[edit protocols igmp-snooping vlan {vlan-name} no-flood-to-multicast-router-interfaces]	
"no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping)" on page 11628	

[interface]

[edit protocols interface {name}]	
"accept-join-always-from" on page 1632	

[isis]

[edit protocols isis]		
"backup-spf-options (Protocols IS-IS)" on page 2821	"bier-sub-domain (Protocols IS-IS)" on page 2974	"clns-routing" on page 4349

(Continued)

[edit protocols isis]		
"clns-updown-compatibility" on page 4350	"context-identifier (Protocols IS-IS)" on page 4611	"export (ISIS)" on page 6658
"graceful-restart (Protocols IS-IS)" on page 7771	"ignore-attached-bit" on page 8396	"import (Protocols IS-IS)" on page 8476
"interface (Protocols IS-IS)" on page 8966	"label-switched-path (Protocols IS-IS)" on page 9582	"layer2-map" on page 9651
"level (Global IS-IS)" on page 9724	"loose-authentication-check" on page 10137	"lsp-lifetime" on page 10189
"max-areas" on page 10420	"max-lsp-size" on page 10446	"max-snp-size" on page 10483
"no-ipv4-routing" on page 11483	"no-ipv6-routing" on page 11484	"overload (Protocols IS-IS)" on page 12048
"reference-bandwidth (Protocols IS-IS)" on page 14224	"rib-group (Protocols IS-IS)" on page 15770	"source-packet-routing (Protocols IS-IS)" on page 28266
"spf-options (Protocols IS-IS)" on page 28379	"strict-dual-isis (Protocols IS-IS)" on page 28648	"topologies (Protocols IS-IS)" on page 29438
"traceoptions (Protocols IS-IS)" on page 29684	"traceoptions (Protocols)" on page 29654	"traffic-engineering (Protocols IS-IS)" on page 30025

[edit protocols isis backup-spf-options]	
"use-post-convergence-lfa" on page 30531	

[edit protocols isis dynamic-overload]		
"auto-recovery" on page 2733	"no-overload-on-prefix-export-limit" on page 11542	

[edit protocols isis interface {interface-name}]		
"bfd-liveness-detection (Protocols IS-IS)" on page 2933	"delay-measurement (Protocols IS-IS)" on page 5249	"delay-metric" on page 5252
"layer2-map" on page 9651	"level (IS-IS Interfaces)" on page 9728	"link-protection (Protocols IS-IS)" on page 9811
"no-adjacency-down-notification (Protocols IS-IS)" on page 11380	"no-advertise-adjacency-segment (Protocols IS-IS)" on page 11382	"passive (Protocols IS-IS)" on page 12219

[edit protocols isis interface {name} delay-measurement]	
"advertisement (Delay Measurement)" on page 2031	

[edit protocols isis interface {interface-name} family inet|inet6]	
"bfd-liveness-detection (Protocols IS-IS)" on page 2933	

[edit protocols isis interface {interface-name} level {level-number}]		
"flood-reflector (Protocols IS-IS)" on page 7242	"hold-time (Protocols IS-IS)" on page 8108	"metric (Protocols IS-IS)" on page 10684
"passive (Protocols IS-IS)" on page 12219	"post-convergence-lfa" on page 12903	

[edit protocols isis interface {name} level {number} post-convergence-lfa]		
"node-protection" on page 11651		

[edit protocols isis interface-group {group-name}]		
"link-group-protection (Protocols IS-IS)" on page 9794		

[edit protocols isis interface-group {name} level {name} topology]		
"bandwidth-based-metrics" on page 2837		

[edit protocols isis level {level-number}]		
"flood-reflector (Protocols IS-IS)" on page 7242	"prefix-export-limit (Protocols IS-IS)" on page 13083	"source-packet-routing (Protocols IS-IS)" on page 28266

[edit protocols isis no-external-export protocol]		
"no-external-export" on page 11454		

[edit protocols isis source-packet-routing]		
"flex-algorithm" on page 7221	"sensor-based-stats" on page 16555	"srgb (Protocols IS-IS)" on page 28408
"strict-asla-based-flex-algorithm" on page 28645	"traffic-statistics (Protocols IS-IS)" on page 30048	

[edit protocols isis spf-options]		
"microloop-avoidance" on page 10700	"multipath (SPF Options)" on page 11117	

[edit protocols isis traffic-engineering]		
"advertisement" on page 2030	"family (Protocols IS-IS)" on page 6905	

[l2-learning]

[edit protocols l2-learning]		
"arp-nd-probe-disable" on page 2484	"arp-nd-probe-failed-log" on page 2485	"crb-proxy-mac" on page 4704
"decapsulate-accept-inner-vlan" on page 5158	"garp-na-enable" on page 7648	"global-mac-ip-limit" on page 7709
"global-mac-ip-table-aging-time" on page 7711	"global-mac-limit" on page 7712	"global-mac-limit (Protocols)" on page 7714
"global-mac-move" on page 7716	"global-mac-statistics" on page 7717	"global-mac-table-aging-time" on page 7719
"global-mode (Protocols)" on page 7721	"global-no-control-mac-aging" on page 7722	"global-no-mac-learning" on page 7723
"global-no-mac-learning" on page 7723	"mclag-arp-nd-sync" on page 10585	"mclag-arpreq-sync" on page 10586
"traceoptions (Layer 2 Learning)" on page 29593	"vpls-mac-move" on page 30932	

[edit protocols l2-learning global-mac-limit {limit}]	
"packet-action" on page 12129	

[edit protocols l2-learning global-mac-move]		
"disable-action" on page 5743	"exclusive-mac" on page 6634	"notification-time" on page 11700
"reopen-time" on page 14374	"threshold-count" on page 29320	"threshold-time" on page 29343
"virtual-mac" on page 30766		

[l2circuit]

[edit protocols l2circuit]		
"local-switching (Layer 2 Circuits)" on page 9977	"neighbor (Protocols Layer 2 Circuit)" on page 11278	"traceoptions (Protocols Layer 2 Circuit)" on page 29688

[edit protocols l2circuit local-switching]	
"interface (Protocols Layer 2 Circuit)" on page 8974	

[edit protocols l2circuit local-switching interface {interface-name}]		
"backup-neighbor" on page 2802	"connection-protection" on page 4540	"encapsulation-type (Layer 2 Circuits)" on page 6308
"end-interface" on page 6338	"ignore-encapsulation-mismatch" on page 8397	"ignore-mtu-mismatch" on page 8408
"protect-interface" on page 13409		

[edit protocols l2circuit local-switching interface {interface-name} backup-neighbor {address}]		
"community (Protocols Layer 2 Circuit)" on page 4469	"mtu (interfaces)" on page 11027	"virtual-circuit-id" on page 30749

[edit protocols l2circuit local-switching interface {interface-name} end-interface]	
"protect-interface" on page 13409	

[edit protocols l2circuit local-switching interface {interface-name} end-interface interface {interface-name}]	
"backup-interface (Layer 2 Circuits)" on page 2797	

[edit protocols l2circuit local-switching interfaces {interface-name}]	
"no-revert (Local Switching)" on page 11568	

[edit protocols l2circuit local-switching interfaces {interface-name} end-interface {interface-name}]	
"no-revert (Local Switching)" on page 11568	

[edit protocols l2circuit neighbor {address}]	
"interface (Protocols Layer 2 Circuit)" on page 8974	

[edit protocols l2circuit neighbor address interface interface-name static]	
"send-oam" on page 16544	

[edit protocols l2circuit neighbor {address} interface {interface-name}]		
"backup-neighbor" on page 2802	"bandwidth (Protocols Layer 2 Circuit)" on page 2870	"community (Protocols Layer 2 Circuit)" on page 4469
"connection-protection" on page 4540	"control-word (Protocols Layer 2 Circuit Neighbor)" on page 4647	"description (Protocols Layer 2 Circuit Neighbor)" on page 5342
"egress-protection (Layer 2 circuit)" on page 6218	"encapsulation-type (Layer 2 Circuits)" on page 6308	"encapsulation-type (Layer 2 VPNs)" on page 6310
"flow-label-receive" on page 7307	"flow-label-receive-static" on page 7309	"flow-label-transmit" on page 7312
"flow-label-transmit-static" on page 7314	"ignore-encapsulation-mismatch" on page 8397	"ignore-mtu-mismatch" on page 8408
"mtu (interfaces)" on page 11027	"oam (Layer 2 Circuit Protocols)" on page 11748	"protect-interface" on page 13409
"pseudowire-status-tlv" on page 13849	"psn-tunnel-endpoint" on page 13851	"revert-time (Protocols Layer 2 Circuits)" on page 15729
"static (Protocols Layer 2 Circuit)" on page 28559	"switchover-delay" on page 28759	"virtual-circuit-id" on page 30749

[edit protocols l2circuit neighbor {address} interface {interface-name} backup-neighbor {address}]		
"community (Protocols Layer 2 Circuit)" on page 4469	"hot-standby (Protocols Layer 2 Circuit)" on page 8189	"mtu (interfaces)" on page 11027
"psn-tunnel-endpoint" on page 13851	"standby (Protocols Layer 2 Circuit)" on page 28489	"static (Protocols Layer 2 Circuit)" on page 28559
"virtual-circuit-id" on page 30749		

[edit protocols l2circuit neighbor {address} interface {interface-name} egress-protection]		
"protected-l2circuit" on page 13412	"protector-interface" on page 13420	"protector-pe" on page 13421

[edit protocols l2circuit neighbor {address} interface {interface-name} oam]	
"ping-interval" on page 12545	

[edit protocols l2circuit neighbor {address} interface {interface-name} pseudowire-status-tlv]	
"hot-standby-vc-on (Protocols Layer 2 Circuit)" on page 8190	

[edit protocols l2circuit neighbor {address} interfaces {interface-name}]	
"no-revert (Neighbor Interface)" on page 11569	

[lacp]

[edit protocols lacp]		
"ppm (Ethernet Switching)" on page 12930	"traceoptions (LACP)" on page 29591	

[edit protocols lacp ppm]		
"centralized" on page 3250	"inline" on page 8666	

[layer2-control]

[edit protocols layer2-control]		
"bpdu-block" on page 3022	"mac-rewrite" on page 10259	"nonstop-bridging" on page 11669

[edit protocols layer2-control bpdu-block]		
"disable-timeout (BPDU)" on page 5808	"interface (BPDU Blocking)" on page 8842	

[edit protocols layer2-control bpdu-block interface {interface-name}]	
"drop-threshold" on page 5978	

[edit protocols layer2-control mac-rewrite]	
"interface (Layer 2 Protocol Tunneling)" on page 8904	

[edit protocols layer2-control mac-rewrite interface {interface-name}]		
"enable-all-ifl" on page 6252	"protocol" on page 13424	

[ldp]

[edit protocols ldp]		
"auto-targeted-session" on page 2748	"deaggregate" on page 5153	"dod-request-policy" on page 5873

(Continued)

[edit protocols ldp]		
"dual-transport" on page 6059	"egress-policy" on page 6215	"explicit-null (Protocols BGP)" on page 6652
"explicit-null (Protocols LDP)" on page 6654	"export (Protocols LDP)" on page 6680	"family (protocols ldp)" on page 6831
"graceful-restart (Protocols LDP)" on page 7773	"igp-synchronization" on page 8414	"import (Protocols LDP)" on page 8477
"interface (Protocols LDP)" on page 8978	"keepalive-interval" on page 9468	"keepalive-timeout" on page 9474
"l2-smart-policy" on page 9528	"label-withdrawal-delay" on page 9598	"log-updown (Protocols LDP)" on page 10043
"longest-match" on page 10108	"make-before-break (LDP)" on page 10296	"neighbor (Protocols LDP)" on page 11281
"no-forwarding" on page 11464	"oam (Protocols LDP)" on page 11759	"p2mp (Protocols LDP)" on page 12122
"policing (Protocols LDP)" on page 12662	"preference (Protocols LDP)" on page 13040	"session (Protocols LDP)" on page 16845
"session-group" on page 16810	"session-protection" on page 16843	"strict-targeted-hellos" on page 28654
"targeted-hello" on page 28962	"traceoptions (Protocols LDP)" on page 29693	"track-igp-metric" on page 29989
"traffic-statistics (Protocols LDP)" on page 30053	"transport-address" on page 30118	

[edit protocols ldp entropy-label]		
"ingress-policy" on page 8641		

[edit protocols ldp graceful-restart]		
"disable (Protocols LDP)" on page 5789	"helper-disable (LDP)" on page 8046	"maximum-neighbor-reconnect-time" on page 10537
"maximum-neighbor-recovery-time" on page 10538	"reconnect-time" on page 14135	"recovery-time (LDP)" on page 14148

[edit protocols ldp interface {interface-name}]		
"allow-subnet-mismatch" on page 2207	"disable (Protocols LDP)" on page 5789	"hello-interval (Protocols LDP)" on page 8036
"hold-time (Protocols LDP)" on page 8110	"preference (Protocols LDP)" on page 13040	"preference (Protocols LDP)" on page 13040
"transport-address" on page 30118		

[edit protocols ldp oam]		
"bfd-liveness-detection (Protocols LDP)" on page 2937	"fec (protocols ldp)" on page 6990	"ingress-policy" on page 8641
"periodic-traceroute" on page 12441		

[edit protocols ldp oam bfd-liveness-detection]		
"ecmp" on page 6190	"version (BFD)" on page 30695	

[edit protocols ldp oam bfd-livenesss-detection]		
"failure-action (Protocols LDP)" on page 6789	"holddown-interval (protocols ldp)" on page 8124	

[edit protocols ldp oam fec {address}]		
"bfd-liveness-detection (Protocols LDP)" on page 2937	"periodic-traceroute" on page 12441	

[edit protocols ldp oam fec {address} bfd-liveness-detection]		
"ecmp" on page 6190	"version (BFD)" on page 30695	

[edit protocols ldp oam fec {address} bfd-livenesss-detection]		
"failure-action (Protocols LDP)" on page 6789	"holddown-interval (protocols ldp)" on page 8124	

[edit protocols ldp session {session-address}]		
"authentication-algorithm" on page 2577	"authentication-key (Protocols LDP)" on page 2618	"authentication-key-chain (Protocols LDP)" on page 2607
"downstream-on-demand" on page 5934	"tcpao-auth-mismatch" on page 29034	

[edit protocols ldp session-group {address}]	
"authentication-key (Protocols LDP)" on page 2618	

[edit protocols ldp targeted-hello]		
"hello-interval (Protocols LDP)" on page 8036	"hold-time (Protocols LDP)" on page 8110	

[edit protocols ldp traffic-statistics]	
"sensor-based-stats (Junos Telemetry Interface)" on page 16558	

[link-management]

[edit protocols link-management]		
"peer (Protocols LMP)" on page 12379	"te-link" on page 29043	"traceoptions (Protocols Link Management)" on page 29697

[edit protocols link-management peer {peer-name}]		
"address (Peer)" on page 1895	"control-channel (Protocols Link Management Peer)" on page 4621	"lmp-control-channel" on page 9864
"lmp-protocol" on page 9865	"te-link" on page 29043	

[edit protocols link-management peer {peer-name} lmp-control-channel {control-channel-interface}]	
"remote-address (for LMP Control Channel)" on page 14312	

[edit protocols link-management peer {peer-name} lmp-protocol]		
"hello-dead-interval" on page 8030	"hello-interval (LMP)" on page 8035	"passive (Protocols Link Management)" on page 12222

(Continued)

[edit protocols link-management peer {peer-name} lmp-protocol]		
"retransmission-interval" on page 15683	"retry-limit (Protocols Link Management)" on page 15710	

[edit protocols link-management te-link {te-link-name}]		
"disable (GMPLS)" on page 5761	"ethernet-vlan (Protocols Link Management)" on page 6514	"interface (Protocols Link Management)" on page 8980
"label-switched-path (Protocols Link Management)" on page 9583	"local-address (Protocols Link Management)" on page 9908	"remote-address (for LMP Traffic Engineering)" on page 14314
"remote-id (protocols link-management)" on page 14330		

[edit protocols link-management te-link {te-link-name} interface {interface-name}]		
"local-address (Protocols Link Management)" on page 9908	"remote-address (for LMP Traffic Engineering)" on page 14314	"remote-id (protocols link-management)" on page 14330

[lldp]

[edit protocols lldp]		
"chassis-id (LLDP)" on page 3301	"interface (LLDP)" on page 8817	"tlv-filter" on page 29410
"tlv-select" on page 29414	"traceoptions (LLDP)" on page 29596	

[edit protocols lldp interface {interface-name}]		
"tlv-filter" on page 29410	"tlv-select" on page 29414	

[lldp-med]

[edit protocols lldp-med]		
"interface (LLDP-MED)" on page 8906	"tlv-filter" on page 29410	"tlv-select" on page 29414

[edit protocols lldp-med interface {interface-name}]		
"tlv-filter" on page 29410	"tlv-select" on page 29414	

[edit protocols lldp-med interface (all | {interface-name} location civic-based)]	
"ca-type" on page 3161	

[edit protocols lldp-med interface (all | {interface-name})]	
"location (LLDP-MED)" on page 9990	

[edit protocols lldp-med interface (all | {interface-name}) location]	
"civic-based" on page 3337	

[mld]

[edit protocols mld]		
"accounting (Protocols MLD)" on page 1733	"interface (Protocols MLD)" on page 8981	"maximum-transmit-rate (Protocols MLD)" on page 10574
"query-interval (Protocols MLD)" on page 13932	"query-last-member-interval (Protocols MLD)" on page 13936	"query-response-interval (Protocols MLD)" on page 13943
"robust-count (Protocols MLD)" on page 15821		

[edit protocols mld interface {interface-name}]		
"accounting (Protocols MLD Interface)" on page 1735	"disable (Protocols MLD)" on page 5791	"disable (Protocols)" on page 5731
"group-limit (Protocols MLD)" on page 7847	"group-policy (Protocols MLD)" on page 7858	"group-threshold (Protocols MLD Interface)" on page 7905
"immediate-leave" on page 8454	"log-interval (MLD Interface)" on page 10011	"oif-map (MLD Interface) " on page 11804
"passive (MLD)" on page 12204	"ssm-map (Protocols MLD)" on page 28466	"ssm-map-policy (MLD)" on page 28462
"static (Protocols MLD)" on page 28560	"version (Protocols MLD)" on page 30718	

[edit protocols mld interface {interface-name} static]	
"group (Protocols MLD)" on page 7870	

[edit protocols mld interface {interface-name} static group {multicast-group-address}]		
"exclude (Protocols MLD)" on page 6615	"group-count (Protocols MLD)" on page 7812	"group-increment (Protocols MLD)" on page 7840
"source (Protocols MLD)" on page 28307		

[edit protocols mld interface {interface-name} static group {multicast-group-address} source]		
"source-count (Protocols MLD)" on page 28199	"source-increment (Protocols MLD)" on page 28223	

[mld-snooping]

[edit protocols mld-snooping]		
"version (MLD Snooping)" on page 30712	"vlan (MLD Snooping)" on page 30830	

[edit protocols mld-snooping vlan {vlan-name}]		
"evpn-ssm-reports-only" on page 6553	"qualified-vlan " on page 13912	"query-last-member-interval (Protocols MLD)" on page 13936
"query-response-interval (Protocols MLD)" on page 13943	"version (MLD Snooping)" on page 30712	

[edit protocols mld-snooping vlan (all | {vlan-name})]		
"disable (Protocols MLD Snooping)" on page 5792	"disable (Protocols)" on page 5731	"immediate-leave" on page 8454

(Continued)

[edit protocols mld-snooping vlan (all | {vlan-name})]		
"interface (MLD Snooping)" on page 8918	"robust-count (MLD Snooping)" on page 15816	

[edit protocols mld-snooping vlan (all | {vlan-name}) interface (all | {interface-name})]	
"multicast-router-interface (MLD Snooping)" on page 11086	

[edit protocols mld-snooping vlan (all | {vlan-name}) interface (all | {interface-name}) static]	
"group (Protocols MLD Snooping)" on page 7872	

[edit protocols mld-snooping vlan {vlan-name} interface {interface-name}]	
"immediate-leave" on page 8454	

[edit protocols mld-snooping vlan {vlan-name} no-flood-to-multicast-router-interfaces]	
"no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping)" on page 11628	

[mpls]

[edit protocols mpls]		
"abstract-hop" on page 1625	"admin-group (for LSPs)" on page 1972	"admin-group-extended" on page 1968

(Continued)

[edit protocols mpls]		
"admin-groups" on page 1973	"advertisement-hold-time" on page 2033	"auto-policing" on page 2724
"bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)" on page 2843	"class-of-service (Protocols MPLS)" on page 3369	"constituent-list" on page 4571
"container-label-switched-path" on page 4577	"cspf-backoff-time" on page 4731	"deselect-on-bandwidth-failure" on page 5361
"diffserv-te" on page 5705	"disable (Protocols MPLS)" on page 5793	"disable (STP)" on page 5739
"egress-protection (MPLS)" on page 6220	"exclude-srlg" on page 6629	"expand-loose-hop" on page 6647
"explicit-null (Protocols BGP)" on page 6652	"explicit-null (Protocols MPLS)" on page 6655	"hop-limit" on page 8134
"interface (MPLS)" on page 8925	"interface (Protocols MPLS)" on page 8983	"ipv6-tunneling" on page 9377
"label-switched-path" on page 9580	"label-switched-path (Protocols MPLS)" on page 9585	"log-updown (Protocols MPLS)" on page 10045
"lsp-external-controller" on page 10187	"no-cspf" on page 11429	"no-decrement-ttl" on page 11432
"no-propagate-ttl" on page 11553	"oam (Protocols MPLS)" on page 11762	"optimize-adaptive-teardown" on page 11852
"optimize-aggressive" on page 11855	"optimize-hold-dead-delay" on page 11856	"optimize-switchover-delay" on page 11857

(Continued)

[edit protocols mpls]		
"optimize-timer (Protocols MPLS)" on page 11859	"path" on page 12250	"path (Protocols MPLS)" on page 12259
"path-mtu" on page 12256	"preference (Protocols MPLS)" on page 13042	"priority (Protocols MPLS)" on page 13193
"record" on page 14136	"revert-timer" on page 15732	"rsvp-error-hold-time" on page 16056
"sensor-based-stats (Junos Telemetry Interface)" on page 16558	"smart-optimize-timer" on page 28048	"standby" on page 28486
"standby" on page 28485	"static-label-switched-path" on page 28545	"statistics (Protocols MPLS)" on page 28594
"sync-active-path-bandwidth" on page 28783	"traceoptions (Protocols MPLS)" on page 29699	"track-igp-metric (LSP)" on page 29990
"traffic-engineering (Protocols MPLS)" on page 30031	"transit-lsp-association" on page 30069	"ultimate-hop-popping" on page 30323

[edit protocols mpls container-label-switched-path {lsp-name}]		
"label-switched-path-template (Container LSP)" on page 9595	"splitting-merging" on page 28391	

[edit protocols mpls container-label-switched-path {lsp-name} splitting-merging]		
"normalization" on page 11676	"sampling (Protocols MPLS)" on page 16300	

[edit protocols mpls diffserv-te]		
"bandwidth-model" on page 2861	"te-class-matrix" on page 29041	

[edit protocols mpls egress-protection]	
"traceoptions (Egress Protection)" on page 29539	

[edit protocols mpls egress-protection context-identifier {context-id}]	
"advertise-mode (MPLS)" on page 2018	

[edit protocols mpls interface {interface-name}]		
"admin-group (for Interfaces)" on page 1971	"admin-groups-extended" on page 1975	"always-mark-connection-protection-tlv" on page 2219
"disable (Protocols MPLS)" on page 5793	"srlg" on page 28411	"switch-away-lsps" on page 28746

[edit protocols mpls interface (all | {interface-name})]	
"disable (STP)" on page 5739	

[edit protocols mpls interface {interface-name} static]	
"protection-revert-time" on page 13419	

[edit protocols mpls label-range]	
"srgb-label-range" on page 28410	

[edit protocols mpls label-switch-path {lsp-name}]	
"optimize-hold-dead-delay" on page 11856	

[edit protocols mpls label-switched-path {lsp-name}]		
"adaptive (protocols mpls)" on page 1825	"admin-down" on page 1962	"admin-group (for LSPs)" on page 1972
"admin-group-extended" on page 1968	"associate-backup-pe-groups" on page 2511	"auto-bandwidth (MPLS Tunnel)" on page 2686
"bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)" on page 2843	"class-of-service (Protocols MPLS)" on page 3369	"conditional-metric" on page 4506
"corouted-bidirectional" on page 4673	"corouted-bidirectional-passive" on page 4674	"cspf-backoff-time" on page 4731
"description (Protocols MPLS)" on page 5345	"deselect-on-bandwidth-failure" on page 5361	"disable (Protocols MPLS)" on page 5793
"egress-protection (MPLS)" on page 6220	"exclude-srlg" on page 6629	"fast-reroute (Protocols MPLS)" on page 6943
"from (Protocols MPLS)" on page 7595	"hop-limit" on page 8134	"in-place-lsp-bandwidth-update (Protocols MPLS)" on page 8493
"install (Protocols MPLS)" on page 8744	"inter-domain" on page 8788	"ldp-tunneling" on page 9687
"link-protection (Dynamic LSPs)" on page 9807	"lsp-attributes" on page 10183	"lsp-external-controller" on page 10187
"metric (Protocols MPLS)" on page 10686	"no-cspf" on page 11429	"no-decrement-ttl" on page 11432

(Continued)

[edit protocols mpls label-switched-path {lsp-name}]		
"no-install-to-address" on page 11476	"node-link-protection (Protocols MPLS)" on page 11649	"oam (Protocols MPLS)" on page 11762
"optimize-timer (Protocols MPLS)" on page 11859	"p2mp (Protocols MPLS)" on page 12124	"policing" on page 12661
"policing (Protocols MPLS)" on page 12664	"pop-and-forward (Protocols MPLS)" on page 12802	"preference (Protocols MPLS)" on page 13042
"primary" on page 13126	"primary (Protocols MPLS)" on page 13144	"priority (Protocols MPLS)" on page 13193
"random" on page 14043	"record" on page 14136	"retry-limit" on page 15709
"retry-timer" on page 15717	"revert-timer" on page 15732	"secondary (Label Switched Path) " on page 16405
"secondary (Protocols MPLS)" on page 16416	"soft-preemption (Protocols MPLS)" on page 28088	"standby" on page 28486
"standby" on page 28485	"sync-active-path-bandwidth" on page 28783	"template" on page 29071
"to" on page 29419	"traceoptions (Protocols MPLS)" on page 29699	"track-igp-metric (LSP)" on page 29990
"ultimate-hop-popping" on page 30323		

[edit protocols mpls label-switched-path {lsp-name} (primary | secondary) {path-name}]		
"adaptive (protocols mpls)" on page 1825	"admin-group (for LSPs)" on page 1972	"admin-group-extended" on page 1968

(Continued)

[edit protocols mpls label-switched-path {lsp-name} (primary | secondary) {path-name}]		
"bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)" on page 2843	"class-of-service (Protocols MPLS)" on page 3369	"hop-limit" on page 8134
"no-cspf" on page 11429	"no-decrement-ttl" on page 11432	"optimize-timer (Protocols MPLS)" on page 11859
"preference (Protocols MPLS)" on page 13042	"priority (Protocols MPLS)" on page 13193	"record" on page 14136
"select" on page 16525	"standby" on page 28486	"standby" on page 28485
"sync-active-path-bandwidth" on page 28783		

[edit protocols mpls label-switched-path {lsp-name} (primary | secondary) {path-name} admin-group]		
"exclude (for Administrative Groups)" on page 6605	"include-all (for Administrative Groups)" on page 8518	"include-any (for Administrative Groups)" on page 8521

[edit protocols mpls label-switched-path {lsp-name} admin-group]		
"exclude (for Administrative Groups)" on page 6605	"include-all (for Administrative Groups)" on page 8518	"include-any (for Administrative Groups)" on page 8521

[edit protocols mpls label-switched-path {lsp-name} auto-bandwidth]		
"adjust-interval" on page 1946	"adjust-threshold" on page 1953	"adjust-threshold-absolute" on page 1954
"adjust-threshold-activate-bandwidth" on page 1956	"adjust-threshold-overflow-limit" on page 1958	"adjust-threshold-underflow-limit" on page 1959

(Continued)

[edit protocols mpls label-switched-path {lsp-name} auto-bandwidth]		
"maximum-bandwidth (Protocols MPLS)" on page 10504	"minimum-bandwidth" on page 10722	"minimum-bandwidth-adjust-interval" on page 10724
"minimum-bandwidth-adjust-threshold-change" on page 10725	"minimum-bandwidth-adjust-threshold-value" on page 10726	"monitor-bandwidth" on page 10834
"resignal-minimum-bandwidth" on page 15585	"sync-active-path-bandwidth" on page 28783	

[edit protocols mpls label-switched-path {lsp-name} egress-protection context-identifier {context-id}]	
"advertise-mode (MPLS)" on page 2018	

[edit protocols mpls label-switched-path {lsp-name} fast-reroute]		
"bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)" on page 2843	"bandwidth-percent (protocols mpls)" on page 2864	"exclude (for Fast Reroute)" on page 6606
"hop-limit" on page 8134	"include-all (for Fast Reroute)" on page 8519	"include-any (for Fast Reroute)" on page 8522

[edit protocols mpls label-switched-path {lsp-name} lsp-attributes]		
"encoding-type" on page 6318	"gpid" on page 7748	"signal-bandwidth" on page 27941
"switching-type" on page 28758	"upstream-label" on page 30476	

[edit protocols mpls label-switched-path {lsp-name} oam]		
"associate-lsp" on page 2512	"bfd-liveness-detection (Protocols MPLS)" on page 2939	"mpls-tp-mode" on page 10988

(Continued)

[edit protocols mpls label-switched-path {lsp-name} oam]		
"performance-monitoring (Protocols MPLS)" on page 12434		

[edit protocols mpls label-switched-path {lsp-name} oam bfd-liveness-detection]		
"failure-action (Protocols MPLS)" on page 6790		

[edit protocols mpls label-switched-path { lsp-name} oam mpls-tp-mode]		
"lsping-channel-type" on page 10195		

[edit protocols mpls label-switched-path {lsp-name} oam performance-monitoring]		
"querier (performance-monitoring)" on page 13925	"responder (performance-monitoring)" on page 15628	

[edit protocols mpls label-switched-path {lsp-name} oam performance-monitoring querier]		
"delay (querier)" on page 5259	"loss (querier)" on page 10171	"loss-delay (querier)" on page 10140

[edit protocols mpls label-switched-path {lsp-name} oam performance-monitoring querier delay]		
"traffic-class (delay)" on page 29996		

[edit protocols mpls label-switched-path {lsp-name} oam performance-monitoring querier loss]		
"traffic-class (loss)" on page 29998		

[edit protocols mpls label-switched-path {lsp-name} oam performance-monitoring querier loss-delay]		
"traffic-class (loss-delay)" on page 30000		

[edit protocols mpls label-switched-path {lsp-name} oam performance-monitoring responder]		
"delay (responder)" on page 5260	"loss (responder)" on page 10173	

[edit protocols mpls label-switched-path {lsp-name} primary {path-name}]		
"oam (Protocols MPLS)" on page 11762		

[edit protocols mpls label-switched-path {lsp-name} primary {path-name} oam]		
"performance-monitoring (Protocols MPLS)" on page 12434		

[edit protocols mpls label-switched-path {lsp-name} primary {path-name} oam performance-monitoring]		
"querier (performance-monitoring)" on page 13925	"responder (performance-monitoring)" on page 15628	

[edit protocols mpls label-switched-path {lsp-name} primary {path-name} oam performance-monitoring querier]		
"delay (querier)" on page 5259	"loss (querier)" on page 10171	"loss-delay (querier)" on page 10140

[edit protocols mpls label-switched-path {lsp-name} primary {path-name} oam performance-monitoring querier delay]		
"traffic-class (delay)" on page 29996		

[edit protocols mpls label-switched-path {lsp-name} primary {path-name} oam performance-monitoring querier loss]		
"traffic-class (loss)" on page 29998		

[edit protocols mpls label-switched-path {lsp-name} primary {path-name} oam performance-monitoring querier loss-delay]		
"traffic-class (loss-delay)" on page 30000		

[edit protocols mpls label-switched-path {lsp-name} primary {path-name} oam performance-monitoring responder]		
"delay (responder)" on page 5260	"loss (responder)" on page 10173	

[edit protocols mpls label-switched-path {name} secondary {name} non-active-hold-priority]		
"non-active-hold-priority" on page 11657		

[edit protocols mpls label-switched-path {lsp-name} secondary {path-name} oam]		
"performance-monitoring (Protocols MPLS)" on page 12434		

[edit protocols mpls label-switched-path {lsp-name} secondary {path-name} oam performance-monitoring]		
"querier (performance-monitoring)" on page 13925	"responder (performance-monitoring)" on page 15628	

[edit protocols mpls label-switched-path {lsp-name} secondary {path-name} oam performance-monitoring querier]		
"delay (querier)" on page 5259	"loss (querier)" on page 10171	"loss-delay (querier)" on page 10140

[edit protocols mpls label-switched-path {lsp-name} secondary {path-name} oam performance-monitoring querier delay]	
"traffic-class (delay)" on page 29996	

[edit protocols mpls label-switched-path {lsp-name} secondary {path-name} oam performance-monitoring querier loss]	
"traffic-class (loss)" on page 29998	

[edit protocols mpls label-switched-path {lsp-name} secondary {path-name} oam performance-monitoring querier loss-delay]	
"traffic-class (loss-delay)" on page 30000	

[edit protocols mpls label-switched-path {lsp-name} secondary {path-name} oam performance-monitoring responder]		
"delay (responder)" on page 5260	"loss (responder)" on page 10173	

[edit protocols mpls log-updown]	
"no-trap" on page 11604	

[edit protocols mpls lsp-external-controller {lsp-external-controller}]	
"label-switched-path-template" on page 9593	

[edit protocols mpls oam]		
"bfd-liveness-detection (Protocols MPLS)" on page 2939	"mpls-tp-mode" on page 10988	"performance-monitoring (Protocols MPLS)" on page 12434

[edit protocols mpls oam bfd-liveness-detection]		
"failure-action (Protocols MPLS)" on page 6790		

[edit protocols mpls oam mpls-tp-mode]		
"lsping-channel-type" on page 10195		

[edit protocols mpls oam performance-monitoring]		
"querier (performance-monitoring)" on page 13925	"responder (performance-monitoring)" on page 15628	

[edit protocols mpls oam performance-monitoring querier]		
"delay (querier)" on page 5259	"loss (querier)" on page 10171	"loss-delay (querier)" on page 10140

[edit protocols mpls oam performance-monitoring querier delay]		
"traffic-class (delay)" on page 29996		

[edit protocols mpls oam performance-monitoring querier loss]		
"traffic-class (loss)" on page 29998		

[edit protocols mpls oam performance-monitoring querier loss-delay]		
"traffic-class (loss-delay)" on page 30000		

[edit protocols mpls oam performance-monitoring responder]		
"delay (responder)" on page 5260	"loss (responder)" on page 10173	

[edit protocols mpls path-mtu]		
"allow-fragmentation (protocols mpls)" on page 2178		

[edit protocols mpls path-mtu rsvp]		
"mtu-signaling" on page 11037		

[edit protocols mpls static-label-switched-path {lsp-name}]		
"bypass (Static LSP)" on page 3117	"ingress (LSP)" on page 8632	"segment" on page 16519
"transit" on page 30064		

[edit protocols mpls static-label-switched-path {lsp-name} bypass]		
"bandwidth (Static LSP)" on page 2876	"description (Protocols MPLS)" on page 5345	"next-hop (Protocols MPLS)" on page 11349
"next-table" on page 11359	"push (protocols mpls)" on page 13873	"to" on page 29419

[edit protocols mpls static-label-switched-path {lsp-name} ingress]		
"bandwidth (Static LSP)" on page 2876	"description (Protocols MPLS)" on page 5345	"install (Protocols MPLS)" on page 8744
"link-protection (Static LSPs)" on page 9818	"metric (Protocols MPLS)" on page 10686	"next-hop (Protocols MPLS)" on page 11349
"no-install-to-address" on page 11476	"node-protection (Static LSP)" on page 11652	"policing (Protocols MPLS)" on page 12664

(Continued)

[edit protocols mpls static-label-switched-path {lsp-name} ingress]		
"preference (Protocols MPLS)" on page 13042	"push (protocols mpls)" on page 13873	"to" on page 29419

[edit protocols mpls static-label-switched-path {lsp-name} transit {incoming-label}]		
"bandwidth (Static LSP)" on page 2876	"description (Protocols MPLS)" on page 5345	"link-protection (Static LSPs)" on page 9818
"next-hop (Protocols MPLS)" on page 11349	"node-protection (Static LSP)" on page 11652	"pop (protocols mpls)" on page 12797
"swap (protocols mpls)" on page 28738		

[edit protocols mpls statistics]		
"auto-bandwidth (MPLS Statistics)" on page 2684	"no-transit-statistics" on page 11600	

[edit protocols mpls traffic-engineering]	
"database" on page 4850	

[edit protocols mpls traffic-engineering database]		
"export (MPLS Traffic engineering database)" on page 6668	"import (MPLS Traffic Engineering Database)" on page 8471	

[edit protocols mpls traffic-engineering database export]	
"credibility" on page 4710	

[msdp]

[edit protocols msdp]		
"active-source-limit" on page 1819	"data-encapsulation" on page 4828	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681
"group (Protocols MSDP)" on page 7873	"hold-time (Protocols MSDP)" on page 8112	"import (Protocols MSDP)" on page 8478
"keep-alive (Protocols MSDP)" on page 9457	"local-address (Protocols MSDP) " on page 9910	"peer (Protocols MSDP)" on page 12380
"rib-group (Protocols MSDP)" on page 15772	"sa-hold-time (Protocols MSDP)" on page 16262	"source (Protocols MSDP) " on page 28308
"traceoptions (Protocols MSDP)" on page 29703		

[edit protocols msdp active-source-limit]		
"log-interval (Protocols MSDP)" on page 10014	"log-warning (Protocols MSDP)" on page 10049	"maximum (MSDP Active Source Messages) " on page 10536
"threshold (MSDP Active Source Messages) " on page 29327		

[edit protocols msdp group {group-name}]		
"default-peer" on page 5200	"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731
"export (Protocols MSDP)" on page 6681	"import (Protocols MSDP)" on page 8478	"local-address (Protocols MSDP) " on page 9910

(Continued)

[edit protocols msdp group {group-name}]		
"mode (Protocols MSDP)" on page 10814	"peer (Protocols MSDP)" on page 12380	"traceoptions (Protocols MSDP)" on page 29703

[edit protocols msdp group {group-name} peer {address}]		
"active-source-limit" on page 1819	"authentication-key (Protocols MSDP)" on page 2600	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681
"hold-time (Protocols MSDP)" on page 8112	"import (Protocols MSDP)" on page 8478	"keep-alive (Protocols MSDP)" on page 9457
"local-address (Protocols MSDP) " on page 9910	"sa-hold-time (Protocols MSDP)" on page 16262	"traceoptions (Protocols MSDP)" on page 29703

[edit protocols msdp peer {address}]		
"active-source-limit" on page 1819	"authentication-key (Protocols MSDP)" on page 2600	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681
"hold-time (Protocols MSDP)" on page 8112	"import (Protocols MSDP)" on page 8478	"keep-alive (Protocols MSDP)" on page 9457
"local-address (Protocols MSDP) " on page 9910	"sa-hold-time (Protocols MSDP)" on page 16262	"traceoptions (Protocols MSDP)" on page 29703

[edit protocols msdp source {ip-address}/{prefix-length}]	
"active-source-limit" on page 1819	

[mstp]

[edit protocols mstp]		
"bridge-priority" on page 3036	"configuration-name" on page 4516	"disable (STP)" on page 5739
"max-age" on page 10418	"max-hops" on page 10441	"msti" on page 11003
"revision-level" on page 15735		

[edit protocols mstp (all | {interface-name})]	
"disable (STP)" on page 5739	

[edit protocols mstp (all | {interface-name}) arp-on-stp]	
"disable (STP)" on page 5739	

[edit protocols mstp ;interface (all | {interface-name}) arp-on-stp]	
"priority (Protocols STP)" on page 13199	

[edit protocols mstp interface (all | {interface-name}) arp-on-stp]		
"mode" on page 10798	"no-root-port" on page 11572	

[edit protocols mstp interface (all | {interface-name}) bpdu-timeout-action]		
"block (Spanning Trees)" on page 2995	"log (Spanning Trees)" on page 10040	

[edit protocols mstp interface {interface-name}, [edit protocols mstp msti {msti-id} vlan (all {vlan-id} | {vlan-name}) interface {interface-name}]	
"disable (STP)" on page 5739	

[edit protocols mstp msti {msti-id}]		
"bridge-priority" on page 3036	"vlan (MSTP)" on page 30835	"vlan (MSTP)" on page 30835

[edit protocols mstp msti {msti-id} interface {interface-name}]		
"cost" on page 4681	"priority (Protocols STP)" on page 13199	

[edit protocols mstp msti {msti-id} interface (all | {interface-name})]	
"edge" on page 6197	

[edit protocols mstp msti {msti-id} vlan ({vlan-id} | {vlan-name}) interface {interface-name} arp-on-stp]	
"disable (STP)" on page 5739	

[edit protocols mstp msti {msti-id} vlan ({vlan-id} | {vlan-name}) interface (all | {interface-name})]	
"disable (STP)" on page 5739	

[mstpmsti]

[edit protocols mstpmsti {msti-id} interface {interface-name} arp-on-stp]	
"priority (Protocols STP)" on page 13199	

[edit protocols mstpmsti {msti-id} interface {interface-name}) arp-on-stp]	
"mode" on page 10798	

[mvpn]

[edit protocols mvpn]		
"inter-region-template" on page 8793	"traceoptions (Protocols MVPN)" on page 29706	

[edit protocols mvpn inter-region-segmented template {template-name} region {region-name} ingress-replication label-switched-path]	
"label-switched-path-template (Multicast)" on page 9596	

[edit protocols mvpn inter-region-segmented template {template-name} region {region-name} rsvpe-te]	
"label-switched-path-template (Multicast)" on page 9596	

[edit protocols mvpn inter-region-template]		
"template(Protocols MVPN)" on page 29092		

[edit protocols mvpn inter-region-template template {template-name}]		
"all-regions" on page 2157	"region" on page 14262	

[edit protocols mvpn inter-region-template template {template-name} all-regions]		
"ingress-replication" on page 8647	"ldp-p2mp" on page 9685	"rsvp-te (Protocols MVPN)" on page 16058

[edit protocols mvpn inter-region-template template {template-name} all-regions ingress-replication]		
"label-switched-path (Protocols MVPN)" on page 9590		

[edit protocols mvpn inter-region-template template {template-name} all-regions ingress-replication label-switched-path]		
"label-switched-path-template (Multicast)" on page 9596		

[edit protocols mvpn inter-region-template template {template-name} all-regions rsvp-te]		
"label-switched-path-template (Multicast)" on page 9596		

[edit protocols mvpn inter-region-template template {template-name} region {region-name}]		
"ingress-replication" on page 8647	"ldp-p2mp" on page 9685	"rsvp-te (Protocols MVPN)" on page 16058

[edit protocols mvpn inter-region-template template {template-name} region {region-name} ingress-replication]	
"label-switched-path (Protocols MVPN)" on page 9590	

[mvrp]

[edit protocols mvrp]		
"add-attribute-length-in-pdu" on page 1838	"disable (MVRP)" on page 5776	"interface (MVRP)" on page 8931
"no-attribute-length-in-pdu" on page 11401	"no-dynamic-vlan" on page 11446	"passive (MVRP)" on page 12213

[edit protocols mvrp interface (all | {interface-name})]		
"join-timer (MVRP)" on page 9434	"point-to-point (MVRP) " on page 12620	"registration" on page 14268

[edit protocols mvrp interface(all | {interface-name})]		
"disable (MVRP)" on page 5776	"passive (MVRP)" on page 12213	

[neighbor-discovery]

[edit protocols neighbor-discovery]		
"onlink-subnet-only" on page 11832	"secure" on page 16430	

[edit protocols neighbor-discovery secure]		
"cryptographic-address" on page 4728	"security-level" on page 16480	"timestamp" on page 29396
"traceoptions (Protocols Secure Neighbor Discovery)" on page 29739		

[edit protocols neighbor-discovery secure cryptographic-address]		
"key-length" on page 9495	"key-pair" on page 9500	

[oam]

[edit protocols oam]		
"ethernet (Protocols OAM)" on page 6490	"fabric (OAM)" on page 6761	"hold-time (OAM)" on page 8102

[edit protocols oam ethernet]		
"connectivity-fault-management" on page 4556	"connectivity-fault-management-switches " on page 4560	"evcs" on page 6518
"fnp" on page 7364	"link-fault-management" on page 9792	"lmi (Ethernet OAM)" on page 9861

[edit protocols oam ethernet connectivity-fault-management]		
"action-profile (Applying to OAM CFM, for EX Series Switch Only)" on page 1776	"action-profile (Defining for CFM)" on page 1778	"enhanced-cfm-mode" on page 6355

(Continued)

[edit protocols oam ethernet connectivity-fault-management]		
"linktrace" on page 9838	"maintenance-domain" on page 10287	"no-aggregate-delegate-processing" on page 11386
"performance-monitoring" on page 12431	"performance-monitoring (OAM LFM)" on page 12432	"policer (CFM Global)" on page 12627
"protocol (Server MEP)" on page 13471	"sendid-tlv" on page 16554	

[edit protocols oam ethernet connectivity-fault-management action-profile {profile-name}]		
"clear-action (CFM)" on page 3409	"default-actions" on page 5180	"event (CFM)" on page 6521

[edit protocols oam ethernet connectivity-fault-management action-profile {action-profile-name} action]		
"interface-group-down " on page 8893	"log-and-generate-ais" on page 10006	

[edit protocols oam ethernet connectivity-fault-management action-profile {action-profile-name} action log-and-generate-ais]		
"interval (CFM MEP)" on page 9209	"level (CFM MEP)" on page 9721	

[edit protocols oam ethernet connectivity-fault-management action-profile {action-profile-name} action log-generate-ais]	
"priority (CFM MEP)" on page 13159	

[edit protocols oam ethernet connectivity-fault-management action-profile {profile-name} default-actions]	
"interface-down" on page 8870	

[edit protocols oam ethernet connectivity-fault-management action-profile {action-profile-name} event]		
"ais-trigger-condition" on page 2099	"interface-status-tlv" on page 9073	"port-status-tlv" on page 12888
"rdi" on page 14077		

[edit protocols oam ethernet connectivity-fault-management action-profile {action-profile-name} event ais-trigger-condition]		
"adjacency-loss" on page 1941	"all-defects" on page 2155	"cross-connect-ccm" on page 4726
"erroneous-ccm" on page 6439	"receive-ais" on page 14112	

[edit protocols oam ethernet connectivity-fault-management expected-defect]		
"rx-enable" on page 16168	"rx-max-duration" on page 16169	"tx-duration" on page 30269
"tx-enable" on page 30271		

[edit protocols oam ethernet connectivity-fault-management linktrace]		
"age" on page 2042	"path-database-size" on page 12252	"path-database-size (EX Series Switch Only)" on page 12253

[edit protocols oam ethernet connectivity-fault-management maintenance-domain {maintenance-domain-name}]		
"bridge-domain" on page 3029	"instance (OAM)" on page 8753	"level (Connectivity Fault Management)" on page 9715
"level (EX Series Switch Only)" on page 9722	"maintenance-association" on page 10285	"name-format" on page 11172

[edit protocols oam ethernet connectivity-fault-management maintenance-domain {domain-name} default-{x}]	
"virtual-switch" on page 30777	

[edit protocols oam ethernet connectivity-fault-management maintenance-domain {name} level {number} maintenance-association {name}]	
"policer (CFM Session)" on page 12629	

[edit protocols oam ethernet connectivity-fault-management maintenance-domain {domain-name} maintenance-association {ma-name}]		
"continuity-check" on page 4616	"continuity-check (EX Series Switch Only)" on page 4618	"mep" on page 10639
"protect-maintenance-association (OAM)" on page 13410	"remote-maintenance-association (OAM)" on page 14351	"short-name-format" on page 16986

[edit protocols oam ethernet connectivity-fault-management maintenance-domain {domain-name} maintenance-association {ma-name} continuity-check]		
"connection-protection-tlv" on page 4541	"convey-loss-threshold" on page 4658	"hold-interval (OAM CFM, for EX Series Switch Only)" on page 8094
"hold-interval (OAM)" on page 8093	"interface-status-send-rdi" on page 9071	"interface-status-tlv" on page 9073
"interval (EX Series Switch Only)" on page 9210	"interval (OAM)" on page 9203	"loss-threshold" on page 10175
"port-status-tlv" on page 12888	"sendid-tlv" on page 16554	

[edit protocols oam ethernet connectivity-fault-management maintenance-domain {md-name} maintenance-association {ma-name} mep {mep-id}]		
"action-profile (MEP)" on page 1780	"auto-discovery" on page 2709	"direction (MEP)" on page 5721
"interface (IEEE 802.1ag OAM Connectivity-Fault Management)" on page 8900	"interface (OAM CFM, for EX Series Switch Only)" on page 8945	"lowest-priority-defect" on page 10179
"remote-mep" on page 14352	"remote-mep (EX Series Switch Only)" on page 14354	

[edit protocols oam ethernet connectivity-fault-management maintenance-domain {domain-name} maintenance-association {ma-name} mep {mep-id} remote-mep {mep-id}]		
"action-profile (Applying to CFM)" on page 1775	"action-profile (Applying to OAM CFM, for EX Series Switch Only)" on page 1776	"detect-loc" on page 5538
"interface-group" on page 8890	"sla-iterator-profile" on page 28012	

[edit protocols oam ethernet connectivity-fault-management maintenance-domain {md-name} maintenance-association {ma-name} mep {mep-id} remote-mep {remote-mep-id} sla-iterator-profile {profile-name}]		
"data-tlv-size" on page 4849	"iteration-count" on page 9412	"priority (Protocols OAM)" on page 13194

[edit protocols oam ethernet connectivity-fault-management maintenance-domain {md-name}maintenance-association {ma-name} mep {mep-id}]	
"priority (OAM Connectivity-Fault Management)" on page 13182	

[edit protocols oam ethernet connectivity-fault-management maintenance-domain {maintenance-domain-name} virtual-switch {virtual-switch-name}]	
"bridge-domain" on page 3029	

[edit protocols oam ethernet connectivity-fault-management maintenance-domain{domain-name}maintenance-association {ma-name} mep{mep-id}]	
"priority (OAM Connectivity-Fault Management)" on page 13182	

[edit protocols oam ethernet connectivity-fault-management performance-monitoring]		
"delegate-server-processing" on page 5270	"hardware-assisted-keepalives" on page 7971	"hardware-assisted-timestamping" on page 7973
"sla-iterator-profiles" on page 28013	"sla-iterator-profiles (OAM CFM)" on page 28015	

[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles {profile-name}]		
"avg-fd-twoway-threshold" on page 2775	"avg-flr-backward-threshold" on page 2777	"avg-flr-forward-threshold" on page 2779
"avg-ifdv-twoway-threshold" on page 2781	"calculation-weight" on page 3177	"cycle-time" on page 4788
"flap-trap-monitor" on page 7215	"iteration-period" on page 9414	"measurement-type" on page 10593

[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles {profile-name} calculation-weight]		
"delay" on page 5238	"delay-variation" on page 5265	

[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles{profile-name} avg-fd-twoway-threshold]	
"flap-trap-monitor" on page 7215	

[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles{profile-name} avg-ifdv-twoway-threshold]	
"flap-trap-monitor" on page 7215	

[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles{profile-name}avg-flr-backward-threshold]	
"flap-trap-monitor" on page 7215	

[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles{profile-name}avg-flr-forward-threshold]	
"flap-trap-monitor" on page 7215	

[edit protocols oam ethernet connectivity-fault-managementmaintenance-association {ma-name}]	
"mip-half-function" on page 10765	

[edit protocols oam ethernet connectivity-fault-managementmaintenance-domain {md-name}]	
"mip-half-function" on page 10765	

[edit protocols oam ethernet evcs]	
"evc-protocol cfm" on page 6517	

[edit protocols oam ethernet link-fault-management]		
"action (OAM LFM)" on page 1764	"action-profile " on page 1771	"interface (OAM Link-Fault Management)" on page 8947
"loopback-tracking" on page 10131		

[edit protocols oam ethernet link-fault-management action-profile]		
"action (OAM)" on page 1762	"event (LFM)" on page 6525	

[edit protocols oam ethernet link-fault-management action-profile action]		
"link-down" on page 9790	"send-critical-event" on page 16539	"syslog (OAM Action)" on page 28824

[edit protocols oam ethernet link-fault-management action-profile event]		
"link-adjacency-loss" on page 9783	"link-event-rate" on page 9791	

[edit protocols oam ethernet link-fault-management action-profile event link-event-rate]		
"frame-error" on page 7550	"frame-period" on page 7554	"frame-period-summary" on page 7555

[edit protocols oam ethernet link-fault-management interface]		
"apply-action-profile" on page 2421	"detect-loc" on page 5538	"link-discovery" on page 9787
"negotiation-options" on page 11263	"pdu-interval" on page 12328	"pdu-threshold" on page 12329

[edit protocols oam ethernet link-fault-management interface {interface-name} negotiation-options]	
"no-allow-link-events" on page 11389	

[edit protocols oam fabric]		
"fabric-maintenance-associations" on page 6756	"flow-specs" on page 7349	

[edit protocols oam fabric fabric-maintenance-associations {fma-name}]	
"fabric-maintenance-end-points" on page 6759	

[edit protocols oam fabric flow-specs {flow-specification-name}]		
"ethernet-frame-size" on page 6486	"multicast-ipv4" on page 11070	"multicast-vlan-flood" on page 11094
"unicast-ethernet" on page 30338	"unicast-ethernet-ipv4" on page 30340	

[edit protocols oam gre-tunnel interface {interface-name}]		
"hold-time (OAM)" on page 8102	"keepalive-time (Protocols OAM)" on page 9472	"keepalive-time (Protocols OAM)" on page 9472

[edit protocols oam link-fault-management interface {interface-name}]	
"event-thresholds" on page 6537	

[edit protocols oam link-fault-management interface {interface-name} event-thresholds]		
"frame-error" on page 7550	"frame-period" on page 7554	"frame-period-summary" on page 7555

[edit protocols oam link-fault-management interface {interface-name} negotiation-options]	
"allow-remote-loopback" on page 2200	

[openflow]

[edit protocols openflow]		
"switch (Protocols OpenFlow)" on page 28755	"traceoptions (Protocols OpenFlow)" on page 29710	

[edit protocols openflow switch {switch-name}]		
"controller (Protocols OpenFlow)" on page 4654	"default-action (Protocols OpenFlow)" on page 5179	"interfaces (Protocols OpenFlow)" on page 9153
"purge-flow-timer (Protocols OpenFlow)" on page 13869		

[edit protocols openflow switch {switch-name} controller]		
"address (Protocols OpenFlow)" on page 1912	"id (Protocols OpenFlow)" on page 8284	"protocol (Protocols OpenFlow)" on page 13455
"role (Protocols OpenFlow)" on page 15829		

[edit protocols openflow switch {switch-name} controller protocol {protocol}]	
"port (Protocols OpenFlow)" on page 12857	

[ospf]

[edit protocols ospf]		
"source-packet-routing (Protocols IS-IS)" on page 28266	"source-packet-routing (Protocols IS-IS)" on page 28266	"source-packet-routing (Protocols OSPF)" on page 28270
"topology (OSPF)" on page 29445		

[edit protocols ospf area {area-id}]		
"label-switched-path (Protocols OSPF)" on page 9592	"network-summary-export" on page 11325	"network-summary-import" on page 11326
"peer-interface (Protocols OSPF)" on page 12370		

[edit protocols ospf area {area-id} interface {interface-name}]		
"application-specific (Protocols OSPF)" on page 2389	"authentication (Protocols OSPF)" on page 2570	"mtu (interfaces)" on page 11027
"no-advertise-adjacency-segment (Protocols OSPF)" on page 11383	"no-neighbor-down-notification" on page 11519	"post-convergence-lfa (Protocols OSPF)" on page 12904
"topology (OSPF Interface)" on page 29446		

[edit protocols ospf area {area-id} interface {interface-name} topology {topology-name}]	
"bandwidth-based-metrics (Protocols OSPF)" on page 2835	

[edit protocols ospf area {area-id} interface {interface-name} topology (default | {name})]	
"link-protection (Protocols OSPF)" on page 9814	

[edit protocols ospf area {area-id} virtual-link]	
"authentication (Protocols OSPF)" on page 2570	

[edit protocols ospf backup-spf-options]		
"no-source-packet-routing (Protocols OSPF)" on page 11583	"use-post-convergence-lfa (Protocols OSPF)" on page 30533	

[edit protocols ospf interface {name}]		
"delay-measurement (Protocols OSPF)" on page 5251	"delay-metric (Protocols OSPF)" on page 5253	

[edit protocols ospf overload]		
"allow-route-leaking" on page 2202	"as-external" on page 2486	"stub-network" on page 28667

[edit protocols ospf source-packet-routing]		
"ldp-stitching (Protocols OSPF)" on page 9686	"node-segment (Protocols OSPF)" on page 11654	"prefix-segment" on page 13105

[edit protocols ospf source-packet-routing srgb]		
"srgb-label-range" on page 28410		

[edit protocols ospf spf-options]		
"microloop-avoidance (Protocols OSPF)" on page 10702	"multipath (OSPF SPF Options)" on page 11109	

[edit protocols ospf topology {name}]		
"topology-id" on page 29441		

[edit protocols ospf topology (default | ipv4-multicast | {name})]		
"overload (Protocols OSPF)" on page 12051	"prefix-export-limit (Protocols OSPF)" on page 13085	"spf-options (Protocols OSPF)" on page 28381

[edit protocols ospf topology (default | {name})]		
"backup-spf-options (Protocols OSPF)" on page 2823		

[edit protocols ospf traffic-engineering shortcuts]		
"ignore-lsp-metrics" on page 8402		

[edit protocols ospf | isis]		
"traffic-engineering" on page 30014		

[ospf3]

[edit protocols ospf3]		
"realm" on page 14093		

[edit protocols ospf3 area {area-id}]		
"inter-area-prefix-export" on page 8782	"inter-area-prefix-import" on page 8783	

[edit protocols ospf3 overload]		
"allow-route-leaking" on page 2202	"as-external" on page 2486	"intra-area-prefix" on page 9224

[edit protocols ospf3 realm (ip4-unicast | ipv4-multicast | ipv6-multicast)]		
"inter-area-prefix-import" on page 8783		

[edit protocols ospf3 realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) area {area-id} interface {interface-name}]		
"interface-type (Protocols OSPF)" on page 9091		

[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)]		
"area" on page 2451	"export (OSPF)" on page 6660	"import (Protocols OSPF)" on page 8463
"lsa-refresh-interval" on page 10182	"no-nssa-abr" on page 11533	"no-rfc-1583" on page 11571
"nssa" on page 11717	"overload (Protocols OSPF)" on page 12051	"prefix-export-limit (Protocols OSPF)" on page 13085

(Continued)

[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)]		
"reference-bandwidth (Protocols OSPF)" on page 14226	"rib-group (Protocols OSPF)" on page 15774	"spf-options (Protocols OSPF)" on page 28381
"stub" on page 28665	"traceoptions (Protocols OSPF)" on page 29713	

[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id}]		
"area-range" on page 2453	"inter-area-prefix-export" on page 8782	"interface (Protocols OSPF)" on page 8984

[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} interface {interface-name}]		
"bandwidth-based-metrics (Protocols OSPF)" on page 2835	"bfd-liveness-detection (Protocols OSPF)" on page 2941	"passive (Protocols OSPF)" on page 12223

[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} interface {interface-name} passive]	
"traffic-engineering (Passive TE Mode)" on page 30021	

[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} nssa]	
"default-lsa" on page 5197	

[edit protocols ospf3 realm ipv4-unicast area {area-id}]	
"link-protection (Protocols OSPF)" on page 9814	

[ovsdb]

[edit protocols ovsdb]		
"controller (OVSDB)" on page 4652	"interfaces (OVSDB)" on page 9151	"traceoptions (OVSDB)" on page 29630

[edit protocols ovsdb controller]		
"inactivity-probe-duration" on page 8503	"maximum-backoff-duration" on page 10501	"protocol (OVSDB)" on page 13454

[edit protocols ovsdb controller protocol]		
"port (OVSDB)" on page 12852		

[pcep]

[edit protocols pcep]		
"message-rate-limit" on page 10655	"pce" on page 12303	"pce-group (Protocols PCEP)" on page 12307
"traceoptions (Protocols PCEP)" on page 29717	"update-rate-limit" on page 30446	

[edit protocols pcep pce {pce-id}]		
"delegation-cleanup-timeout" on page 5272	"delegation-priority" on page 5274	"destination-ipv4-address" on page 5440
"destination-ipv6-address" on page 5443	"destination-port (protocols pcep)" on page 5460	"local-ipv6-address" on page 9956

(Continued)

[edit protocols pcep pce {pce-id}]		
"lsp-cleanup-timer" on page 10185	"max-unknown-messages" on page 10491	"max-unknown-requests" on page 10492
"pce-group (PCE)" on page 12306	"pce-type" on page 12309	"traceoptions (PCE)" on page 29633

[edit protocols pcep pce-group {group-id}]		
"lsp-cleanup-timer" on page 10185	"max-unknown-requests" on page 10492	

[pim]

[edit protocols pim]		
"assert-timeout" on page 2500	"count (Protocols PIM)" on page 4684	"default-vpn-source" on page 5222
"dense-groups" on page 5315	"disable (PIM)" on page 5782	"disable (Protocols)" on page 5731
"disable-packet-register" on page 5813	"dr-election-on-p2p" on page 5937	"export (Protocols PIM)" on page 6683
"family (Protocols PIM)" on page 6907	"graceful-restart (Protocols PIM)" on page 7777	"idle-standby-path-switchover-delay" on page 8301
"import (Protocols PIM)" on page 8480	"interface (Protocols PIM)" on page 8991	"interval (Protocols PIM)" on page 9202
"join-load-balance" on page 9431	"join-prune-timeout" on page 9433	"mldp-inband-signalling (Protocols Multipoint LDP)" on page 10787

(Continued)

[edit protocols pim]		
"mpls-internet-multicast" on page 10978	"multiple-triggered-joins" on page 11119	"override-interval" on page 12058
"propagation-delay" on page 13381	"reset-tracking-bit" on page 15583	"rib-group (Protocols PIM)" on page 15775
"rp" on page 15996	"rpf-selection" on page 16015	"rpf-vector (PIM)" on page 16017
"secondary-vtep-address" on page 16419	"sglimit" on page 16919	"spt-threshold" on page 28401
"standby-path-creation-delay" on page 28488	"traceoptions (Protocols PIM)" on page 29719	

[edit protocols pim default-vpn-source]	
"interface-name" on page 8934	

[edit protocols pim family (inet | inet6)]		
"disable (PIM)" on page 5782	"disable (Protocols)" on page 5731	

[edit protocols pim graceful-restart]	
"no-bidirectional-mode" on page 11416	

[edit protocols pim interface {interface-name}]		
"accept-remote-source" on page 1642	"dual-dr" on page 6045	"stickydr" on page 28609

[edit protocols pim interface (Protocols PIM) {interface-name}]		
"bidirectional (Interface)" on page 2969	"disable (PIM)" on page 5782	"disable (Protocols)" on page 5731
"distributed-dr" on page 5833	"family (Protocols PIM Interface)" on page 6909	"family (Protocols PIM)" on page 6907
"hello-interval" on page 8032	"mode (Protocols PIM)" on page 10816	"neighbor-policy" on page 11272
"override-interval" on page 12058	"priority (PIM Interfaces)" on page 13184	"propagation-delay" on page 13381
"reset-tracking-bit" on page 15583	"version (Protocols PIM)" on page 30719	

[edit protocols pim interface (Protocols PIM) {interface-name} bfd-liveness-detection]		
"minimum-interval (PIM BFD Liveness Detection)" on page 10738	"minimum-receive-interval (Protocols PIM)" on page 10751	"multiplier (Protocols PIM)" on page 11120

[edit protocols pim interface (Protocols PIM) {interface-name} bidirectional]	
"df-election" on page 5572	

[edit protocols pim interface (Protocols PIM) {interface-name} bidirectional df-election]		
"backoff-period" on page 2791	"offer-period" on page 11785	"robustness-count" on page 15823

[edit protocols pim interface (Protocols PIM) {interface-name} family (inet | inet6)]	
"bfd-liveness-detection (Protocols PIM)" on page 2945	

[edit protocols pim interface {interface-name} bfd-liveness-detection]		
"detection-time (BFD for PIM)" on page 5543	"no-adaptation (PIM BFD Liveness Detection)" on page 11377	"transmit-interval (PIM BFD Liveness Detection)" on page 30096

[edit protocols pim interface {interface-name} bfd-liveness-detection authentication]		
"algorithm" on page 2144	"loose-check" on page 10138	

[edit protocols pim interface {interface-name} bfd-liveness-detection detection-time]	
"threshold (PIM BFD Detection Time)" on page 29331	

[edit protocols pim interface {interface-name} bfd-liveness-detection transmit-interval]		
"minimum-interval (PIM BFD Transmit Interval) " on page 10739	"threshold (PIM BFD Transmit Interval)" on page 29333	

[edit protocols pim interface {interface-name} family (inet | inet6) bfd-liveness-detection]	
"authentication (Protocols PIM)" on page 2647	

[edit protocols pim interface {interface-name} family {inet | inet6} bfd-liveness-detection authentication]	
"key-chain (Protocols PIM)" on page 9490	

[edit protocols pim interface {interface-name} multiple-triggered-joins]	
"interface (Protocols PIM)" on page 8991	

[edit protocols pim local family (inet | inet6) anycast-pim]	
"rp-set" on page 16001	

[edit protocols pim mldp-inband-signalling]	
"policy (Protocols Multipoint LDP)" on page 12706	

[edit protocols pim rp]		
"auto-rp" on page 2735	"bidirectional (RP)" on page 2971	"bootstrap" on page 3015
"bootstrap-export" on page 3016	"bootstrap-import" on page 3018	"bootstrap-priority" on page 3019
"dr-register-policy" on page 5938	"embedded-rp" on page 6248	"group-rp-mapping" on page 7891
"local" on page 9899	"register-limit" on page 14264	"register-probe-time" on page 14266
"rp-register-policy" on page 16000	"static (Protocols PIM)" on page 28562	

[edit protocols pim rp auto-rp]	
"mapping-agent-election" on page 10342	

[edit protocols pim rp bidirectional]	
"address (Bidirectional Rendezvous Points)" on page 1865	

[edit protocols pim rp bidirectional address {address}]		
"group-ranges" on page 7885	"hold-time (Protocols PIM)" on page 8114	"priority (PIM RPs)" on page 13186

[edit protocols pim rp bootstrap]	
"family (Bootstrap)" on page 6846	

[edit protocols pim rp bootstrap (inet | inet6)]		
"import (Protocols PIM Bootstrap)" on page 8482	"priority (Bootstrap)" on page 13157	

[edit protocols pim rp bootstrap family (inet | inet6)]	
"export (Bootstrap)" on page 6663	

[edit protocols pim rp embedded-rp]		
"group-ranges" on page 7885	"maximum-rps" on page 10556	

[edit protocols pim rp group-rp-mapping]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit protocols pim rp local]		
"family (Local RP)" on page 6883	"override (PIM Static RP)" on page 12061	"process-non-null-as-null-register" on page 13246

[edit protocols pim rp local (inet | inet6) anycast-pim rp-set]		
"address (Anycast RPs)" on page 1856		

[edit protocols pim rp local family (inet | inet6)]		
"address (Local RPs)" on page 1885	"anycast-pim" on page 2283	"disable (PIM)" on page 5782
"disable (Protocols)" on page 5731	"group-ranges" on page 7885	"hold-time (Protocols PIM)" on page 8114
"priority (PIM RPs)" on page 13186		

[edit protocols pim rp local family (inet | inet6) anycast-pim]		
"local-address (Protocols PIM)" on page 9911		

[edit protocols pim rp local family inet]		
"override (PIM Static RP)" on page 12061		

[edit protocols pim rp local family inet6]		
"override (PIM Static RP)" on page 12061		

[edit protocols pim rp register-limit]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547

(Continued)

[edit protocols pim rp register-limit]		
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit protocols pim rp rpf-vector]	
"policy (PIM rpf-vector) " on page 12703	

[edit protocols pim rp static address {address}]		
"group-ranges" on page 7885	"override (PIM Static RP)" on page 12061	"version (Protocols PIM)" on page 30719

[edit protocols pim sglimit]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit protocols pim spt-threshold]	
"infinity" on page 8622	

[edit protocols pim static]	
"address (Static RPs)" on page 1933	

[piminterface]

[edit protocols piminterface (Protocols PIM) {interface-name} bfd-liveness-detection]	
"version (Protocols BFD)" on page 30714	

[ppp]

[edit protocols ppp]		
"monitor-session" on page 10926	"traceoptions (PPP Process)" on page 29641	

[ppp-service]

[edit protocols ppp-service]		
"max-failures" on page 10438	"on-demand-ip-address" on page 11815	"reject-unauthorized-ipv6cp" on page 14271
"source-interface-set-at-login" on page 28229	"traceoptions (Protocols PPP Service)" on page 29723	

[pppoe]

[edit protocols pppoe]		
"no-send-pads-ac-info" on page 11577	"no-send-pads-error" on page 11578	"pado-advertise" on page 12166
"service-name-tables" on page 16717	"traceoptions (PPPoE)" on page 29645	

[edit protocols pppoe service-name-tables {table-name}]	
"service (PPPoE)" on page 16727	

[edit protocols pppoe service-name-tables {table-name} service {service-name}]		
"agent-specifier" on page 2051	"delay (PPPoE Service Name Tables)" on page 5257	"drop (PPPoE Service Name Tables)" on page 5952
"dynamic-profile (PPPoE Service Name Tables)" on page 6129	"max-sessions (PPPoE Service Name Tables)" on page 10477	"routing-instance (PPPoE Service Name Tables)" on page 15954
"terminate (PPPoE Service Name Tables)" on page 29149		

[edit protocols pppoe service-name-tables {table-name} service {service-name} agent-specifier aci {circuit-id-string} ari {remote-id-string}]		
"delay (PPPoE Service Name Tables)" on page 5257	"drop (PPPoE Service Name Tables)" on page 5952	"dynamic-profile (PPPoE Service Name Tables)" on page 6129
"routing-instance (PPPoE Service Name Tables)" on page 15954	"static-interface" on page 28539	"terminate (PPPoE Service Name Tables)" on page 29149

[protection-group]

[edit protocols protection-group]		
"ethernet-ring" on page 6497	"traceoptions (Ethernet Switching)" on page 29465	

[edit protocols protection-group ethernet-aps]		
"clear" on page 3402	"exercise" on page 6637	"force switch" on page 7374
"lockout" on page 9999	"manual switch" on page 10323	

[edit protocols protection-group ethernet-ring {ring-name}]		
"compatibility-version" on page 4481	"control-vlan" on page 4643	"data-channel" on page 4825
"dot1p-priority" on page 5913	"east-interface" on page 6186	"guard-interval" on page 7947
"hold-interval (Protection Group)" on page 8095	"major-ring-name" on page 10295	"node-id" on page 11648
"non-revertive" on page 11660	"non-vc-mode" on page 11666	"propagate-tc" on page 13380
"restore-interval" on page 15669	"ring-id" on page 15797	"ring-protection-link-owner" on page 15800
"wait-to-block-interval" on page 31016	"west-interface" on page 31102	

[edit protocols protection-group ethernet-ring {name} (east-interface | west-interface)]		
"control-channel" on page 4620	"ring-protection-link-end" on page 15799	"virtual-control-channel" on page 30751

[edit protocols protection-group ethernet-ring {ring-name} east-interface]	
"interface-none" on page 8941	

[edit protocols protection-group ethernet-ring name (east-interface |west-interface)]	
"control-vlan" on page 4643	

[edit protocols protection-group ethernet-ring {ring-name} west-interface]	
"interface-none" on page 8941	

[ptp]

[edit protocols ptp]		
"acquiring-state-announce-grandmaster" on page 1754	"apts" on page 2436	"clock-mode" on page 4357
"domain (Precision Time Protocol)" on page 5875	"e2e-transparent" on page 6179	"frequency-lock-threshold" on page 7580
"holdover-time-error-budget" on page 8129	"ipv4-dscp" on page 9329	"local-priority" on page 9973
"master" on page 10360	"performance-monitor" on page 12429	"phase-adjust-threshold" on page 12469
"phase-lock-threshold" on page 12471	"phy-timestamping" on page 12472	"priority1" on page 13209
"priority2" on page 13211	"profile-type (PTP)" on page 28809	"slave" on page 28034

(Continued)

[edit protocols ptp]		
"stateful (Precision Time Protocol) " on page 28516	"syntonized-e2e-transparent" on page 28813	"unicast-negotiation" on page 30347

[edit protocols ptp client clock-class-to-quality-level-mapping]	
"clock-class" on page 4352	

[edit protocols ptp client interface {interface-name}]		
"primary" on page 13127	"secondary (Precision Time Protocol) " on page 16406	"unicast-mode " on page 30344

[edit protocols ptp client interface {interface-name} unicast-mode transport]	
"clock-ipv6-source " on page 4356	

[edit protocols ptp master interface {interface-name}]		
"primary" on page 13127	"secondary (Precision Time Protocol) " on page 16406	

[edit protocols ptp master interface interface-name multicast-mode]	
"local-priority" on page 9973	

[edit protocols ptp master interface interface-name unicast-mode]	
"local-priority" on page 9973	

[edit protocols ptp master interface interface-name unicast-mode transport ipv4]		
"clock-ipv6-client" on page 4355		

[edit protocols ptp slave]		
"clock-class-to-quality-level-mapping" on page 4353	"convert-clock-class-to-quality-level" on page 4656	"hybrid" on page 8220
"interface (Precision Time Protocol)" on page 8815		

[edit protocols ptp slave clock-class-to-quality-level-mapping]	
"quality-level (hybrid)" on page 13919	

[edit protocols ptp slave hybrid]	
"synchronous-ethernet-mapping" on page 28808	

[edit protocols ptp slave hybrid synchronous-ethernet-mapping]	
"clock-source (hybrid)" on page 4362	

[edit protocols ptp slave interface interface-name multicast-mode]	
"local-priority" on page 9973	

[edit protocols ptp slave interface interface-name unicast-mode]	
"local-priority" on page 9973	

[edit protocols ptp stateful interface interface-name multicast-mode]	
"local-priority" on page 9973	

[edit protocols ptp stateful interface interface-name unicast-mode]	
"local-priority" on page 9973	

[edit protocols ptp stateful interface {name} unicast-mode]	
"remote-ip-address" on page 14344	

[rift]

[edit protocols rift]		
"capabilities (Protocols RIFT)" on page 3197	"default-prefixes (Protocols RIFT)" on page 5206	"export (Protocols RIFT)" on page 6685
"interface (Protocols RIFT)" on page 8994	"level (Protocols RIFT)" on page 9734	"lie-receive-address (Protocols RIFT)" on page 9742

[edit protocols rift interface {interface-name}]	
"lie-transmit-address (Protocols RIFT)" on page 9744	

[rip]

[edit protocols rip]		
"authentication-key (Protocols RIP)" on page 2619	"authentication-type (Protocols RIP)" on page 2669	"check-zero" on page 3309
"graceful-restart (Protocols RIP)" on page 7779	"group (Protocols RIP)" on page 7877	"holddown (Protocols RIP)" on page 8125
"import (Protocols RIP)" on page 8483	"message-size" on page 10658	"metric-in (Protocols RIP)" on page 10671
"receive (Protocols RIP)" on page 14117	"rib-group (Protocols RIP)" on page 15777	"route-timeout (Protocols RIP)" on page 15885
"send (Protocols RIP)" on page 16546	"traceoptions (Protocols RIP)" on page 29727	"update-interval (Protocols RIP)" on page 30441

[edit protocols rip group {group-name}]		
"bfd-liveness-detection (Protocols RIP)" on page 2947	"demand-circuit (Protocols RIP)" on page 5294	"export (RIP)" on page 6661
"max-retrans-time" on page 10464	"neighbor" on page 11264	"preference (Protocols RIP)" on page 13043
"route-timeout (Protocols RIP)" on page 15885	"update-interval (Protocols RIP)" on page 30441	

[edit protocols rip group {group-name} neighbor {neighbor-name}]		
"any-sender" on page 2281	"authentication-key (Protocols RIP)" on page 2619	"authentication-type (Protocols RIP)" on page 2669

(Continued)

[edit protocols rip group {group-name} neighbor {neighbor-name}]		
"check-zero" on page 3309	"demand-circuit (Protocols RIP)" on page 5294	"dynamic-peers" on page 6111
"import (Protocols RIP)" on page 8483	"interface-type (Protocols RIP)" on page 9095	"max-retrans-time" on page 10464
"message-size" on page 10658	"metric-in (Protocols RIP)" on page 10671	"metric-out (Protocols RIP)" on page 10674
"peer (Protocols RIP)" on page 12382	"receive (Protocols RIP)" on page 14117	"route-timeout (Protocols RIP)" on page 15885
"send (Protocols RIP)" on page 16546	"update-interval (Protocols RIP)" on page 30441	

[ripng]

[edit protocols ripng]		
"graceful-restart (Protocols RIPng)" on page 7780	"group (Protocols RIPng)" on page 7880	"holddown (Protocols RIPng)" on page 8127
"import (Protocols RIPng)" on page 8485	"metric-in (Protocols RIPng)" on page 10673	"receive (Protocols RIPng)" on page 14119
"route-timeout (Protocols RIPng)" on page 15887	"send (Protocols RIPng)" on page 16548	"traceoptions (Protocols RIPng)" on page 29730
"update-interval (Protocols RIPng)" on page 30443		

[edit protocols ripng group {group-name}]		
"export (Protocols RIPng)" on page 6686	"neighbor (Protocols RIPng)" on page 11282	"preference (Protocols RIPng)" on page 13045

[edit protocols ripng group {group-name} neighbor {neighbor-name}]		
"import (Protocols RIPng)" on page 8485	"metric-in (Protocols RIPng)" on page 10673	"metric-out (Protocols RIPng)" on page 10679
"receive (Protocols RIPng)" on page 14119	"send (Protocols RIPng)" on page 16548	

[router-advertisement]

[edit protocols router-advertisement]		
"interface (Protocols IPv6 Neighbor Discovery)" on page 8964	"traceoptions (Protocols IPv6 Neighbor Discovery)" on page 29681	

[edit protocols router-advertisement interface {interface-name}]		
"current-hop-limit" on page 4735	"default-lifetime" on page 5193	"default-lifetime (Dynamic Router Advertisement)" on page 5194
"dns-server-address" on page 5864	"downstream" on page 5932	"downstream-mode" on page 5933
"link-mtu" on page 9802	"managed-configuration" on page 10301	"max-advertisement-interval (Protocols IPv6 Neighbor Discovery)" on page 10415
"min-advertisement-interval (Protocols IPv6 Neighbor Discovery)" on page 10709	"other-stateful-configuration" on page 11953	"parameter-preference" on page 12173

(Continued)

[edit protocols router-advertisement interface {interface-name}]		
"passive-mode" on page 12205	"preference (IPv6 Router Advertisement)" on page 13035	"prefix (Protocols IPv6 Neighbor Discovery)" on page 13100
"reachable-time" on page 14086	"reachable-time (Dynamic Router Advertisement)" on page 14087	"retransmit-timer" on page 15689
"retransmit-timer (Dynamic Router Advertisement)" on page 15690	"solicit-router-advertisement-unicast" on page 28106	"upstream-mode" on page 30478

[edit protocols router-advertisement interface {interface-name} dns-server-address {address}]	
"lifetime (Router Advertisement)" on page 9745	

[edit protocols router-advertisement interface {interface-name} prefix {prefix}]		
"autonomous" on page 2761	"on-link" on page 11821	"preferred-lifetime" on page 13056
"valid-lifetime" on page 30644		

[router-discovery]

[edit protocols router-discovery]		
"address (Protocols Router Discovery)" on page 1913	"disable (Protocols Router Discovery)" on page 5797	"interface (Protocols Router Discovery)" on page 8998
"traceoptions (Protocols Router Discovery)" on page 29733		

[edit protocols router-discovery address {address}]		
"advertise" on page 2007	"broadcast (Router Discovery)" on page 3043	"ineligible" on page 8553
"multicast (Protocols Router Discovery)" on page 11077	"priority (Protocols Router Discovery)" on page 13196	

[edit protocols router-discovery interface {interface-name}]		
"lifetime (Router Discovery)" on page 9747	"max-advertisement-interval (Protocols Router Discovery)" on page 10416	"min-advertisement-interval (Protocols Router Discovery)" on page 10711

[rstp]

[edit protocols rstp]		
"bridge-priority" on page 3036	"disable (STP)" on page 5739	"extended-system-id" on page 6706
"max-age" on page 10418		

[edit protocols rstp force-version]	
"stp" on page 28625	

[edit protocols rstp interface {interface-name}]	
"disable (STP)" on page 5739	

[edit protocols rstp interface (all | {interface-name})]		
"disable (STP)" on page 5739		

[edit protocols rstp interface (all | {interface-name}) arp-on-stp]		
"disable (STP)" on page 5739	"mode" on page 10798	"no-root-port" on page 11572
"priority (Protocols STP)" on page 13199		

[edit protocols rstp interface (all | {interface-name}) bpdu-timeout-action]		
"block (Spanning Trees)" on page 2995	"log (Spanning Trees)" on page 10040	

[rsvp]

[edit protocols rsvp]		
"authentication-key (Protocols RSVP)" on page 2620	"devices" on page 5571	"disable (Protocols RSVP)" on page 5798
"fast-reroute (Protocols RSVP)" on page 6944	"graceful-deletion-timeout" on page 7758	"graceful-restart (Protocols RSVP)" on page 7782
"hello-acknowledgements" on page 8029	"interface (Protocols RSVP)" on page 8999	"keep-multiplier" on page 9465
"load-balance (Protocols RSVP)" on page 9885	"no-authentication-check (Protocols RSVP)" on page 11403	"no-enhanced-frr-bypass (Protocols RSVP)" on page 11451

(Continued)

[edit protocols rsvp]		
"no-interface-hello" on page 11479	"no-local-reversion" on page 11497	"no-node-hello-on-bypass" on page 11523
"no-node-id-subobject" on page 11524	"no-p2mp-sublsp" on page 11544	"node-hello" on page 11646
"peer-interface (Protocols RSVP)" on page 12372	"pop-and-forward (Protocols RSVP)" on page 12803	"preemption" on page 13026
"refresh-time" on page 14257	"setup-protection" on page 16897	"traceoptions (Protocols RSVP)" on page 29736
"tunnel-services (RSVP)" on page 30238		

[edit protocols rsvp graceful-restart]		
"disable (Protocols RSVP)" on page 5798	"maximum-helper-recovery-time" on page 10522	"maximum-helper-restart-time (RSVP)" on page 10525

[edit protocols rsvp interface {interface-name}]		
"authentication-key (Protocols RSVP)" on page 2620	"bandwidth (Protocols RSVP)" on page 2872	"disable (Protocols RSVP)" on page 5798
"hello-interval (Protocols RSVP)" on page 8038	"link-protection (RSVP)" on page 9816	"reliable" on page 14304
"subscription" on page 28709	"update-threshold" on page 30456	"update-threshold-max-reservable" on page 30458

[edit protocols rsvp interface {interface-name} link-protection]		
"admin-group" on page 1967	"bandwidth (Protocols RSVP)" on page 2872	"bypass (Signaled LSP)" on page 3115
"class-of-service (Protocols RSVP)" on page 3370	"disable (Protocols RSVP)" on page 5798	"exclude-srlg" on page 6629
"hop-limit" on page 8134	"max-bypasses" on page 10422	"no-cspf (Protocols RSVP)" on page 11431
"optimize-timer (Protocols RSVP)" on page 11861	"path (Protocols RSVP)" on page 12261	"priority (Protocols RSVP)" on page 13197
"subscription" on page 28709		

[edit protocols rsvp interface {interface-name} link-protection bypass {bypass-name}]		
"admin-group" on page 1967	"bandwidth (Protocols RSVP)" on page 2872	"class-of-service (Protocols RSVP)" on page 3370
"exclude-srlg" on page 6629	"hop-limit" on page 8134	"no-cspf (Protocols RSVP)" on page 11431
"path (Protocols RSVP)" on page 12261	"priority (Protocols RSVP)" on page 13197	

[edit protocols rsvp interface {name} non-rsvp-bandwdith]	
"local-bw-override-threshold" on page 9925	

[edit protocols rsvp interface {interface-name} update-threshold]	
"adaptive (RSVP)" on page 1828	

[edit protocols rsvp peer-interface {peer-interface-name}]		
"aggregate (Protocols RSVP)" on page 2068	"authentication-key (Protocols RSVP)" on page 2620	"disable (Protocols RSVP)" on page 5798
"dynamic-bidirectional-transport" on page 6101	"hello-interval (Protocols RSVP)" on page 8038	"reliable" on page 14304

[edit protocols rsvp preemption]	
"soft-preemption (Protocols RSVP)" on page 28089	

[edit protocols rsvp route]	
"include-igp-metric" on page 8527	

[sap]

[edit protocols sap]		
"disable (Protocols SAP)" on page 5800	"disable (Protocols)" on page 5731	"listen" on page 9842

[sflow]

[edit protocols sflow]		
"adaptive-sample-rate" on page 1830	"agent-id" on page 2050	"collector" on page 4408

(Continued)

[edit protocols sflow]		
"interfaces (sFlow)" on page 9168	"polling-interval" on page 12753	"sample-rate" on page 16288
"source-ip" on page 28233	"traceoptions (sFlow Technology)" on page 29862	

[edit protocols sflow collector]	
"udp-port" on page 30310	

[edit protocols sflow interfaces {interface-name}]		
"polling-interval" on page 12753	"sample-rate" on page 16288	

[source-packet-routing]

[edit protocols source-packet-routing]		
"compute-options" on page 4496	"compute-profile" on page 4497	"preserve-nexthop-hierarchy (SRv6 TE)" on page 13122
"rib-group (Source Packet Routing)" on page 15780	"segment-list" on page 16521	"source-routing-path" on page 28316
"source-routing-path-template" on page 28319	"telemetry" on page 29054	"traceoptions (Protocols Spring-TE)" on page 29742
"use-transport-class" on page 30541		

[edit protocols source-packet-routing segment-list]		
"bfd-liveness-detection (LSP)" on page 2931	"dynamic (Source Packet Routing)" on page 6156	

[edit protocols source-packet-routing source-routing-path {lsp-name}]		
"install (Protocols MPLS)" on page 8744	"ldp-tunneling" on page 9687	

[edit protocols source-packet-routing source-routing-path { lsp-name} primary {segment-list-name}]	
"bfd-liveness-detection (LSP)" on page 2931	

[edit protocols source-packet-routing source-routing-path {name} primary {name} bfd-liveness-detection (LSP)]	
"sbfd" on page 16323	

[edit protocols source-packet-routing source-routing-path { lsp-name} secondary {segment-list-name}]	
"bfd-liveness-detection (LSP)" on page 2931	

[edit protocols source-packet-routing source-routing-path {name} secondary {name} bfd-liveness-detection (LSP)]	
"sbfd" on page 16323	

[edit protocols source-packet-routing source-routing-path-template]	
"bfd-liveness-detection (Source Packet Routing Template)" on page 2955	

[edit protocols source-packet-routing telemetry]	
"statistics" on page 28591	

[spring-traffic-engineering]

[edit protocols spring-traffic-engineering]	
"lsp-external-controller" on page 10187	

[stp]

[edit protocols stp]		
"bridge-priority" on page 3036	"disable (STP)" on page 5739	"max-age" on page 10418

[edit protocols stp interface (all | {interface-name})]		
"disable (STP)" on page 5739	"mode" on page 10798	

[edit protocols stp interface (all | {interface-name}) arp-on-stp]		
"disable (STP)" on page 5739	"mode" on page 10798	"no-root-port" on page 11572
"priority (Protocols STP)" on page 13199		

[edit protocols stp interface (all | {interface-name}) bpdu-timeout-action]		
"block (Spanning Trees)" on page 2995	"log (Spanning Trees)" on page 10040	

[uplink-failure-detection]

[edit protocols uplink-failure-detection]		
"action (Uplink Failure Detection)" on page 1789	"group" on page 7801	"group (Uplink Failure Detection)" on page 7910
"traceoptions (Uplink Failure Detection)" on page 29903		

[edit protocols uplink-failure-detection group]		
"debounce-interval" on page 5154	"link-to-disable" on page 9831	"link-to-monitor" on page 9832

[vlan]

[edit protocols vlan {vlan-name}]	
"mode (Multicast VLAN Registration)" on page 10809	

[vpls]

[edit protocols vpls static-vpls]	
"no-tunnel-services" on page 11605	

[vrrp]

[edit protocols vrrp]		
"asymmetric-hold-time" on page 2517	"asymmetric-hold-time" on page 2516	"delegate-processing (VRRP)" on page 5268
"failover-delay" on page 6781	"failover-delay" on page 6780	"global-advertisements-threshold" on page 7696
"inherit-advertisement-interval" on page 8654	"skew-timer-disable" on page 28009	"startup-silent-period" on page 28513
"traceoptions (Protocols VRRP)" on page 29747	"version-3" on page 30691	

[vstp]

[edit protocols vstp]		
"disable (STP)" on page 5739	"vlan (MSTP)" on page 30835	"vlan (VSTP)" on page 30878
"vlan-group" on page 30791		

[edit protocols vstp {vlan-id} interface {interface-name}]		
"priority (Protocols STP)" on page 13199		

[edit protocols vstp vlan {vlan-id}]		
"backup-bridge-priority" on page 2793	"bridge-priority" on page 3036	"forward-delay" on page 7393
"hello-time" on page 8041	"interface (Spanning Tree)" on page 9057	"interface (Spanning Tree)" on page 9057
"max-age" on page 10418	"system-id" on page 28847	"vpls-flush-on-topology-change" on page 30926

[edit protocols vstp vlan {vlan-identifier} interface {interface-name}]		
"access-trunk" on page 1705	"cost" on page 4681	"disable (STP)" on page 5739
"mode" on page 10798	"no-root-port" on page 11572	

[edit protocols vstp vlan {vlan-id} interface (all | {interface-name})]		
"disable (STP)" on page 5739	"edge" on page 6197	

[edit protocols vstp vlan {vlan-id} interface (all | {interface-name}) arp-on-stp]		
"disable (STP)" on page 5739	"mode" on page 10798	"no-root-port" on page 11572
"priority (Protocols STP)" on page 13199		

[edit protocols vstp vlan {vlan-id} interface (all | {interface-name}) bpdu-timeout-action]		
"block (Spanning Trees)" on page 2995	"log (Spanning Trees)" on page 10040	

[edit protocols vstp vlan-group group {group-name} vlan (vlan-id | vlan-range |open-set-of-values) [edit routing-instances {routing-instance-name} protocols (mstp | rstp | vstp)]	
"interface (Spanning Tree)" on page 9057	

[edit pvlan]

[edit {vlan-name} pvlan]		
"extend-secondary-vlan-id" on page 6696	"isolation-vlan-id" on page 9411	

[edit redundant-power-system]

IN THIS SECTION

- [member] | 799

Sub Hierarchies

"[edit redundant-power-system member]" on page 799

[edit redundant-power-system]	
"member (Redundant Power System)" on page 10620	

[member]

[edit redundant-power-system member]		
"priority (Redundant Power System)" on page 13204	"priority (Redundant Power System)" on page 13204	

[edit routing-instances]

IN THIS SECTION

- [(Multiple Routing Entities)] | 803
- [...] | 803
- [access] | 804
- [bridge-domains] | 804
- [class-of-service] | 812
- [edit] | 812
- [firewall] | 813
- [flow] | 813
- [forwarding-options] | 813

- [igmp-snooping-options] | **826**
- [instance-name] | **826**
- [instance-type] | **827**
- [interface] | **827**
- [logical-systems] | **828**
- [multicast-snooping-options] | **828**
- [pbb-options] | **829**
- [protocol] | **829**
- [protocols] | **829**
- [provider-tunnel] | **893**
- [realm] | **900**
- [routing-instance-name] | **900**
- [routing-options] | **900**
- [service-groups] | **912**
- [statement-name] | **912**
- [switch-options] | **913**
- [system] | **914**
- [vlans] | **925**
- [vrf-advertise-selective] | **928**
- [vrf-table-label] | **928**
- [vxlan] | **928**

Sub Hierarchies

"[edit routing-instances (Multiple Routing Entities) {routing-instance-name}]" on page 803

"[edit routing-instances {routing-instance-name} ...]" on page 803

"[edit routing-instances access]" on page 804

"[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name}]" on page 804

"[edit routing-instances class-of-service]" on page 812

"[edit routing-instances edit]" on page 812

"[edit routing-instances firewall]" on page 813

"[edit routing-instances flow]" on page 813

"[edit routing-instances {routing-instance-name} forwarding-options]" on page 813

"[edit routing-instances {routing-instance-name} igmp-snooping-options]" on page 826

"[edit routing-instances instance-name]" on page 826

"[edit routing-instances {routing-instance-name} instance-type]" on page 827

"[edit routing-instances {routing-instances-name} interface]" on page 827

"[edit routing-instances logical-systems]" on page 828

"[edit routing-instances {routing-instance-name} multicast-snooping-options]" on page 828

"[edit routing-instances {routing-instance-name} pbb-options]" on page 829

"[edit routing-instances protocol]" on page 829

"[edit routing-instances {routing-instance-name} protocols]" on page 829

"[edit routing-instances {routing-instance-name} provider-tunnel]" on page 893

"[edit routing-instances realm]" on page 900

"[edit routing-instances routing-instance-name]" on page 900

"[edit routing-instances {routing-instance-name} routing-options]" on page 900

"[edit routing-instances {routing-instance-name} service-groups {service-group-name}]" on page 912

"[edit routing-instances statement-name]" on page 912

"[edit routing-instances {routing-instance-name} switch-options]" on page 913

"[edit routing-instances system]" on page 914

"[edit routing-instances {routing-instance-name} vlans {vlan-name}]" on page 925

"[edit routing-instances {routing-instance-name} vrf-advertise-selective]" on page 928

"[edit routing-instances {routing-instance-name} vrf-table-label]" on page 928

"[edit routing-instances {routing-instance-name} vxlan]" on page 928

[edit routing-instances {routing-instance-name}]		
"bridge-domains" on page 3032	"description (Routing Instances)" on page 5346	"egress-protection (Routing Instances)" on page 6221
"forwarding-instance identifier" on page 7457	"forwarding-options (Routing Instance)" on page 7463	"igmp-snooping-options" on page 8393
"instance-role" on page 8764	"instance-type" on page 8773	"interface (Routing Instances)" on page 9012
"interface (Virtual Tunnel in Routing Instances)" on page 9097	"l2-domain-id-for-l3" on page 9523	"l2vpn-id" on page 9562
"layer3-domain-identifier" on page 9666	"mac-vrf" on page 10276	"multicast-snooping-options" on page 11087
"no-local-switching" on page 11499	"no-normalization" on page 11526	"no-vrf-advertise" on page 11620
"no-vrf-propagate-ttl" on page 11621	"oism (Multicast Snooping Options)" on page 11807	"protocols" on page 13748
"protocols (STP Type)" on page 13802	"provider-tunnel" on page 13806	"qualified-bum-pruning-mode" on page 13905
"remote-vtep-list" on page 14362	"robust-count (MLD Snooping)" on page 15816	"route-distinguisher" on page 15856
"routing-options" on page 15980	"service-type (MAC-VRF Routing Instances)" on page 16764	"switch-options" on page 28749
"vlan-id (routing instance)" on page 30817	"vlan-id inner-all" on page 30804	"vlan-model" on page 30833
"vlans" on page 30880	"vnid (EVPN)" on page 30914	"vrf-advertise-selective" on page 30957

(Continued)

[edit routing-instances {routing-instance-name}]		
"vrf-export" on page 30958	"vrf-import" on page 30960	"vrf-propagate-ttl" on page 30965
"vrf-table-label" on page 30966	"vtep-source-interface" on page 30988	

[(Multiple Routing Entities)]

[edit routing-instances (Multiple Routing Entities) {routing-instance-name}]		
"service-groups" on page 16702	"service-type" on page 16761	

[...]

[edit routing-instances {routing-instance-name} ...]		
"active-leasequery (DHCP Relay Agent)" on page 1808	"bulk-leasequery (DHCP Relay Agent)" on page 3068	"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279
"duplicate-clients (DHCPv6 Local Server and Relay Agent)" on page 6064	"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699	"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701
"leasequery (DHCP Relay Agent)" on page 9702	"no-snoop (DHCP Local Server and Relay Agent)" on page 11582	"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334

(Continued)

[edit routing-instances {routing-instance-name} ...]		
"route-suppression (DHCP Local Server and Relay Agent)" on page 15877	"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	

[access]

[edit routing-instances {routing-instances-name} access address-assignment]	
"pool (Address-Assignment Pools)" on page 12766	

[edit routing-instances {vrf-name} access address-assignment address-pools {juniper-pool}]	
"service-mode" on page 16711	

[edit routing-instances {name} access address-assignment pool {name} family (inet | inet6)]	
"excluded-range (Address-Assignment Pools)" on page 6632	

[bridge-domains]

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name}]		
"bridge-options" on page 3034	"host-outbound-traffic (Multicast Snooping)" on page 8162	"interface (Logical Interfaces)" on page 8814

(Continued)

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name}]		
"no-irb-layer-2-copy" on page 11486	"proxy-mac" on page 13829	"routing-interface" on page 15978
"vlan-id-list" on page 30805	"vlan-tags" on page 30865	

[edit routing-instances {name} bridge-domains {name} bridge-options]		
"drop-unknown-mac-ip" on page 5984	"interface-mac-ip-limit" on page 8908	"interface-mac-limit" on page 8910
"mac-ip-table-size" on page 10231	"mac-statistics" on page 10263	"mac-table-size" on page 10271
"network-isolation-profile" on page 11314	"no-normalization" on page 11526	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} bridge-options interface {interface-name}]		
"interface-mac-limit" on page 8910	"mac-pinning" on page 10254	"static-mac" on page 28549
"static-mac-ip" on page 28587		

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} bridge-options interface {interface-name} interface-mac-limit {limit}]	
"packet-action" on page 12129	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} bridge-options interface-mac-limit {limit}]	
"packet-action" on page 12129	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} forwarding-options]	
"dhcp-relay (DHCP Spoofing Prevention)" on page 5637	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} forwarding-options dhcp-relay]	
"group (DHCP Spoofing Prevention)" on page 7829	

[edit routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay dhcpv6 group {name} relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay dhcpv6 relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} forwarding-options dhcp-relay group {group-name}]		
"interface (DHCP Spoofing Prevention)" on page 8867	"interface-tag (DHCP Relay Server)" on page 9083	

[edit routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} multicast-snooping-options]		
"flood-groups" on page 7241	"forwarding-cache (Bridge Domains)" on page 7407	"ignore-stp-topology-change" on page 8412

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} multicast-snooping-options forwarding-cache]	
"threshold (Bridge Domains)" on page 29319	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols]		
"igmp-snooping" on page 8387	"mld-snooping" on page 10783	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols igmp-snooping]		
"evpn-ssm-reports-only" on page 6553	"immediate-leave" on page 8454	"interface (Bridge Domains)" on page 8845
"proxy (Bridge Domains)" on page 13819	"traceoptions (Protocols IGMP Snooping)" on page 29678	"version (IGMP Snooping)" on page 30704
"vlan (Bridge Domains)" on page 30788		

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols igmp-snooping interface {interface-name}]		
"group-limit (IGMP and MLD Snooping)" on page 7846	"host-only-interface" on page 8154	"immediate-leave" on page 8454
"multicast-router-interface (IGMP Snooping)" on page 11084	"query-interval (Bridge Domains)" on page 13927	"query-last-member-interval (Bridge Domains)" on page 13933
"query-response-interval (Bridge Domains)" on page 13938	"robust-count (IGMP Snooping)" on page 15815	"static (Bridge Domains)" on page 28529

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols igmp-snooping interface {interface-name} static]	
"group (Bridge Domains)" on page 7806	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols igmp-snooping interface {interface-name} static group]	
"source (Bridge Domains)" on page 28187	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols igmp-snooping no-flood-to-multicast-router-interfaces]	
"no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping)" on page 11628	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols igmp-snooping proxy]	
"source-address (IGMP and MLD Snooping)" on page 28118	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols igmp-snooping vlan {vlan-id}]	
"query-last-member-interval (Bridge Domains)" on page 13933	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols igmp-snooping vlan {vlan-id} interface {interface-name}]		
"query-interval (Bridge Domains)" on page 13927	"robust-count (IGMP Snooping)" on page 15815	"static (Bridge Domains)" on page 28529

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols igmp-snooping vlan {vlan-id} proxy]	
"source-address (IGMP and MLD Snooping)" on page 28118	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols igmp-snoopingvlan {vlan-id}]	
"query-response-interval (Bridge Domains)" on page 13938	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols mld-snooping]		
"evpn-ssm-reports-only" on page 6553	"query-interval (Bridge Domains)" on page 13927	"query-last-member-interval (Bridge Domains)" on page 13933
"version (MLD Snooping)" on page 30712		

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols mld-snooping interface {interface-name}]		
"query-response-interval (Bridge Domains)" on page 13938		

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols mld-snooping no-flood-to-multicast-router-interfaces]		
"no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping)" on page 11628		

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols vlan {vlan-id} igmp-snooping]		
"interface (Bridge Domains)" on page 8845	"proxy (Bridge Domains)" on page 13819	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols vlan {vlan-id} igmp-snooping interface {interface-name}]		
"group-limit (IGMP and MLD Snooping)" on page 7846	"host-only-interface" on page 8154	"immediate-leave" on page 8454
"multicast-router-interface (IGMP Snooping)" on page 11084		

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols vlan {vlan-id} igmp-snooping interface {interface-name} static]		
"group (Bridge Domains)" on page 7806		

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} protocols vlan {vlan-id} igmp-snooping interface {interface-name} static group]	
"source (Bridge Domains)" on page 28187	

[edit routing-instances {routing-instance-name} bridge-domains {bridge-domain-name} vxlan]		
"ingress-node-replication" on page 8634	"ovsdb-managed" on page 12088	"static-remote-vtep-list" on page 28566

[class-of-service]

[edit routing-instances {routing-instance-name} class-of-service interfaces {interface-name}]	
"member-link-scheduler" on page 10614	

[edit]

[edit routing-instances {routing-instance-name} edit protocols pim rpf-selection]	
"group (RPF Selection)" on page 7893	

[firewall]

[edit routing-instances {routing-instance-name} firewall family {family-name} filter {filter-name}]	
"physical-interface-filter" on page 12476	

[edit routing-instances {routing-instance-name} firewall policer {policer-name}]	
"physical-interface-policer" on page 12479	

[edit routing-instances {routing-instance-name} firewall three-color-policer {policer-name}]	
"physical-interface-policer" on page 12479	

[flow]

[edit routing-instances flow]	
"per-route-accounting" on page 12420	

[forwarding-options]

[edit routing-instances {routing-instance-name} forwarding-options]		
"dhcp-relay" on page 5620	"dhcp-relay (DHCP Spoofing Prevention)" on page 5637	"fast-reroute-priority" on page 6941

(Continued)

[edit routing-instances {routing-instance-name} forwarding-options]		
"load-balance (Forwarding Options)" on page 9878		

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay]		
"active-server-group" on page 1817	"dhcpv6 (DHCP Relay Agent)" on page 5665	"dual-stack-group (DHCP Relay Agent)" on page 6055
"duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)" on page 6066	"forward-snooped-clients (DHCP Relay Agent)" on page 7405	"group (DHCP Spoofing Prevention)" on page 7829
"relay-option-82" on page 14294	"server-group" on page 16633	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay ...]		
"allow-no-end-option (DHCP Relay Agent)" on page 2196	"allow-snooped-clients" on page 2204	"asymmetric-lease-time (DHCP Overrides)" on page 2519
"authentication (DHCP Relay Agent)" on page 2595	"circuit-type (DHCP Relay Agent)" on page 3332	"default-action (DHCP Relay Agent Option)" on page 5177
"drop (DHCP Relay Agent Option)" on page 5946	"dual-stack (DHCP Relay Agent Overrides)" on page 6050	"dynamic-profile (DHCP Relay Agent)" on page 6119
"forward-only (DHCP Relay Agent Option)" on page 7398	"forward-only (DHCP Relay Agent)" on page 7396	"forward-only-replies (DHCP Relay Agent)" on page 7401
"group (DHCP Relay Agent)" on page 7820	"interface (DHCP Relay Agent)" on page 8863	"interface-description (DHCP Relay Agent)" on page 8856
"interface-name (DHCP Relay Agent)" on page 8938	"local-server-group (DHCP Relay Agent Option)" on page 9976	"mac-address (DHCP Relay Agent)" on page 10223

(Continued)

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay ...]		
"no-allow-snooped-clients" on page 11390	"option-number (DHCP Relay Agent Option)" on page 11898	"overrides (DHCP Relay Agent)" on page 12068
"relay-option (DHCP Relay Agent)" on page 14297	"relay-server-group (DHCP Relay Agent Option)" on page 14300	"server-response-time (DHCP Relay Agent)" on page 16664
"service-profile (DHCP Relay Agent)" on page 16731	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982	"username-include (DHCP Relay Agent)" on page 30597

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay ... overrides]	
"delay-authentication (DHCP Relay Agent)" on page 5242	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay ... relay-option-82]	
"remote-id (DHCP Relay Agent)" on page 14331	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay ... relay-option-82 (circuit-id | remote-id)]		
"include-irb-and-l2" on page 8529	"no-vlan-interface-name" on page 11617	"prefix (DHCP Relay Agent)" on page 13077

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay authentication]	
"password (DHCP Relay Agent)" on page 12229	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083
"option-82 (DHCP Relay Agent)" on page 11893	"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6]		
"active-server-group" on page 1817	"server-group" on page 16633	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)]		
"include-irb-and-l2" on page 8529	"no-vlan-interface-name" on page 11617	"prefix (DHCP Relay Agent)" on page 13077
"use-interface-description" on page 30522		

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 ...]		
"client-id (DHCP Relay Agent)" on page 4317	"relay-agent-interface-id (DHCPv6 Relay Agent)" on page 14276	"relay-agent-remote-id (DHCPv6 Relay Agent)" on page 14284
"use-vlan-id (DHCP Relay Agent)" on page 30546		

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 ... (relay-agent-interface-id |relay-agent-remote-id)]	
"use-option-82" on page 30528	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 authentication]	
"password (DHCP Relay Agent)" on page 12229	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083
"relay-agent-interface-id (DHCPv6 Relay Agent Username)" on page 14278	"relay-agent-remote-id (DHCPv6 Relay Agent Username)" on page 14286	"relay-agent-subscriber-id (DHCPv6 Relay Agent)" on page 14289
"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name}]	
"active-server-group" on page 1817	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name} authentication]	
"password (DHCP Relay Agent)" on page 12229	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name} authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083
"relay-agent-interface-id (DHCPv6 Relay Agent Username)" on page 14278	"relay-agent-remote-id (DHCPv6 Relay Agent Username)" on page 14286	"relay-agent-subscriber-id (DHCPv6 Relay Agent)" on page 14289
"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580	

[edit routing-instances {name} forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group { group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name} interface {interface-name}]	
"trace (DHCP Relay Agent)" on page 29458	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name} overrides]		
"interface-client-limit (DHCP Relay Agent)" on page 9108	"no-bind-on-request (DHCP Relay Agent)" on page 11418	"send-release-on-delete (DHCP Relay Agent)" on page 16550

[edit routing-instances {name} forwarding-options dhcp-relay dhcpv6 group {name} relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 overrides]		
"interface-client-limit (DHCP Relay Agent)" on page 9108	"no-bind-on-request (DHCP Relay Agent)" on page 11418	"send-release-on-delete (DHCP Relay Agent)" on page 16550

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id],]	
"keep-incoming-interface-id (DHCP Relay Agent)" on page 9462	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id]	
"keep-incoming-remote-id (DHCP Relay Agent)" on page 9463	

[edit routing-instances {name} forwarding-options dhcp-relay dhcpv6 relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6...]		
"asymmetric-lease-time (DHCP Overrides)" on page 2519	"asymmetric-prefix-lease-time (DHCP Overrides)" on page 2521	

[edit routing-instances {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"dual-stack-interface-client-limit (DHCP Local Server and Relay Agent)" on page 6058	"protocol-master" on page 13448

(Continued)

[edit routing-instances {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"use-vlan-id (DHCP Relay Agent)" on page 30546		

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dual-stack-group],]		
"relay-agent-interface-id (DHCPv6 Relay Agent)" on page 14276		

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name}]		
"active-server-group" on page 1817	"interface (DHCP Spoofing Prevention)" on page 8867	"interface-tag (DHCP Relay Server)" on page 9083
"relay-option-82" on page 14294		

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay group ... overrides]		
"client-discover-match (DHCP Relay Agent)" on page 4310		

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} authentication]		
"password (DHCP Relay Agent)" on page 12229		

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} authentication username-include]		
"delimiter (DHCP Relay Agent)" on page 5290	"domain-name (DHCP Relay Agent)" on page 5886	"logical-system-name (DHCP Relay Agent)" on page 10083
"option-82 (DHCP Relay Agent)" on page 11893	"routing-instance-name (DHCP Relay Agent)" on page 15947	"user-prefix (DHCP Relay Agent)" on page 30580

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay group { group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Relay Agent)" on page 2057	"use-primary (DHCP Relay Agent)" on page 30537	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} interface {interface-name}]	
"trace (DHCP Relay Agent)" on page 29458	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} interface {interface-name} overrides]		
"always-write-giaddr" on page 2221	"always-write-option-82" on page 2222	"disable-relay" on page 5801
"interface-client-limit (DHCP Relay Agent)" on page 9108	"layer2-unicast-replies" on page 9661	"no-bind-on-request (DHCP Relay Agent)" on page 11418
"proxy-mode" on page 13833	"relay-source" on page 14302	"replace-ip-source-with (DHCP Relay Agent)" on page 14377
"send-release-on-delete (DHCP Relay Agent)" on page 16550	"trust-option-82" on page 30168	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} overrides]		
"always-write-giaddr" on page 2221	"always-write-option-82" on page 2222	"disable-relay" on page 5801
"interface-client-limit (DHCP Relay Agent)" on page 9108	"layer2-unicast-replies" on page 9661	"no-bind-on-request (DHCP Relay Agent)" on page 11418
"proxy-mode" on page 13833	"relay-source" on page 14302	"replace-ip-source-with (DHCP Relay Agent)" on page 14377
"send-release-on-delete (DHCP Relay Agent)" on page 16550	"trust-option-82" on page 30168	

[edit routing-instances {name} forwarding-options dhcp-relay group {name} relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} relay-option-60]	
"vendor-option" on page 30684	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay group {group-name} relay-option-82]		
"circuit-id (DHCP Relay Agent)" on page 3324	"server-id-override" on page 16637	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay overrides]		
"always-write-giaddr" on page 2221	"always-write-option-82" on page 2222	"disable-relay" on page 5801
"interface-client-limit (DHCP Relay Agent)" on page 9108	"layer2-unicast-replies" on page 9661	"no-bind-on-request (DHCP Relay Agent)" on page 11418
"proxy-mode" on page 13833	"relay-source" on page 14302	"replace-ip-source-with (DHCP Relay Agent)" on page 14377
"send-release-on-delete (DHCP Relay Agent)" on page 16550	"trust-option-82" on page 30168	

[edit routing-instances {name} forwarding-options dhcp-relay relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay relay-option-60]	
"vendor-option" on page 30684	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay relay-option-82]		
"circuit-id (DHCP Relay Agent)" on page 3324	"server-id-override" on page 16637	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)]	
"use-interface-description" on page 30522	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay relay-option-82 circuit-id]	
"keep-incoming-circuit-id (DHCP Relay Agent)" on page 9460	

[edit routing-instances {routing-instance-name} forwarding-options dhcp-relay relay-option-82 circuit-id],]	
"keep-incoming-remote-id (DHCP Relay Agent)" on page 9463	

[edit routing-instances {name} forwarding-options enhanced-hash-key]	
"hash-parameters" on page 8000	

[edit routing-instances {routing-instance-name} forwarding-options family (inet | inet6 | mpls | vpls) filter]	
"input (Forwarding Table)" on page 8703	

[edit routing-instances {routing-instance-name} forwarding-options family (inet | inet6 | mpls) filter]	
"output (Forwarding Table)" on page 11992	

[edit routing-instances {instance-name} forwarding-options family vpls]		
"filter (VPLS)" on page 7158	"flood" on page 7238	

[edit routing-instances {instance-name} forwarding-options helpers bootp]	
"relay-agent-option" on page 14279	

[edit routing-instances {routing-instance-name} forwarding-options load-balance]		
"indexed-load-balance" on page 8545	"per-flow" on page 12413	"per-prefix" on page 12417

[edit routing-instances {name} forwarding-options pfe-sensor npu-memory]	
"resource-list (Telemetry)" on page 15612	

[edit routing-instances {routing-instance-name} forwarding-options port-mirroring]	
"instance" on page 8748	

[edit routing-instances {name} forwarding-options storm-control]	
"enhanced" on page 6354	

[igmp-snooping-options]

[edit routing-instances {routing-instance-name} igmp-snooping-options]		
"snoop-pseudowires" on page 28075	"use-p2mp-lsp" on page 30530	

[instance-name]

[edit routing-instances instance-name forwarding-options]	
"enhanced-hash-key" on page 6362	

[edit routing-instances instance-name protocols bgp family]	
"inet-mdt (Signaling)" on page 8571	

[edit routing-instances instance-name protocols bgp group group-name family]	
"inet-mdt (Signaling)" on page 8571	

[edit routing-instances instance-name protocols bgp group group-name neighbor address family]	
"inet-mdt (Signaling)" on page 8571	

[instance-type]

[edit routing-instances {routing-instance-name} instance-type]		
"mpls-internet-multicast" on page 10978	"vlan-id (routing instance)" on page 30817	

[edit routing-instances {routing-instance-name} instance-type (layer2-control | virtual-switch)]	
"access-trunk" on page 1705	

[edit routing-instances {routing-instance-name} instance-type {name} routing-options forwarding-table]	
"unicast-reverse-path" on page 30348	

[edit routing-instances {routing-instance-name} instance-type virtual-switch protocols evpn]	
"extended-vlan-list" on page 6707	

[interface]

[edit routing-instances {routing-instances-name} interface]	
"protect-interface" on page 13409	

[edit routing-instances {routing-instance-name} interface vt-{fpc}/{pic}/{port}.{unit-number}]		
"multicast (Virtual Tunnel in Routing Instances)" on page 11092	"primary (Virtual Tunnel in Routing Instances)" on page 13148	"unicast (Virtual Tunnel in Routing Instances)" on page 30354

[logical-systems]

[edit routing-instances {instance-name} logical-systems {logical-system-name} protocols bgp family inet labeled-unicast]	
"per-prefix-label" on page 12419	

[edit routing-instances {instance-name} logical-systems {logical-system-name} protocols bgp group {group-name} family inet labeled-unicast]	
"per-prefix-label" on page 12419	

[multicast-snooping-options]

[edit routing-instances {routing-instance-name} multicast-snooping-options]		
"flood-groups" on page 7241	"forwarding-cache (Bridge Domains)" on page 7407	"host-outbound-traffic (Multicast Snooping)" on page 8162
"nexthop-hold-time" on page 11362		

[edit routing-instances {routing-instance-name} multicast-snooping-options forwarding-cache]	
"threshold (Bridge Domains)" on page 29319	

[pbb-options]

[edit routing-instances {routing-instance-name} pbb-options]		
"source-bmac" on page 28185	"source-bmac" on page 28185	

[edit routing-instances {routing-instance-name} pbb-options {service-group-name} isid {isid-name}]	
"source-bmac" on page 28185	

[protocol]

[edit routing-instances {routing-instance-name} protocol evpn]		
"extended-isid-list" on page 6698	"pbb-evpn-core" on page 12282	

[protocols]

[edit routing-instances {routing-instance-name} protocols]		
"amt (Protocols)" on page 2230	"bfd (protocols)" on page 2911	"bgp" on page 2956
"evpn" on page 6541	"flex-algorithm (Protocols OSPF)" on page 7223	"igmp-snooping" on page 8387
"interconnect" on page 8797	"ip-prefix-routes" on page 9250	"irb-symmetric-routing" on page 9394
"irb-symmetric-routing" on page 9394	"isis" on page 9401	"l2vpn" on page 9559

(Continued)

[edit routing-instances {routing-instance-name} protocols]		
"ldp" on page 9679	"lldp" on page 9851	"mld-snooping" on page 10783
"msdp" on page 10996	"mstp" on page 11006	"mvpn" on page 11135
"mvpn (NG-MVPN)" on page 11143	"oism" on page 11805	"ospf" on page 11938
"ospf3" on page 11942	"peg-df-election" on page 12395	"pim" on page 12507
"pim-evpn-gateway" on page 12515	"pim-snooping" on page 12517	"rip" on page 15803
"ripng" on page 15804	"rstp" on page 16046	"rsvp" on page 16050
"vpls (Routing Instance)" on page 30933	"vstp" on page 30983	

[edit routing-instances {routing-instance-name} protocols (bgp | ldp | ospf | ospf3 | pim) graceful-restart]	
"disable (Routing Graceful Restart)" on page 5738	

[edit routing-instances {name} protocols (evpn | l2vpn | vpls) mesh-group {mesh-group-name} neighbor {neighbor-id} oam bfd-liveness-detection]	
"minimum-receive-interval (BFD Liveness Detection)" on page 10754	

[edit routing-instances {name} protocols (evpn | l2vpn | vpls) neighbor {neighbor-id} oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"minimum-receive-interval (BFD Liveness Detection)" on page 10754	"multiplier (BFD Liveness Detection)" on page 11123

(Continued)

[edit routing-instances {name} protocols (evpn | l2vpn | vpls) neighbor {neighbor-id} oam bfd-liveness-detection]		
"no-adaptation (BFD Liveness Detection)" on page 11374	"version (BFD Liveness Detection)" on page 30697	

[edit routing-instances {name} protocols (evpn | l2vpn | vpls) oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"minimum-receive-interval (BFD Liveness Detection)" on page 10754	"multiplier (BFD Liveness Detection)" on page 11123
"no-adaptation (BFD Liveness Detection)" on page 11374	"version (BFD Liveness Detection)" on page 30697	

[edit routing-instances {name} protocols (evpn | l2vpn | vpls)mesh-group {mesh-group-name} neighbor {neighbor-id} oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"multiplier (BFD Liveness Detection)" on page 11123	"no-adaptation (BFD Liveness Detection)" on page 11374
"version (BFD Liveness Detection)" on page 30697		

[edit routing-instances {routing-instance-name} protocols (ldp | ospf | ospf3) graceful-restart]	
"helper-disable (Multiple Protocols)" on page 8048	

[edit routing-instances {routing-instance-name} protocols (mstp | rstp | vstp)]		
"bpdu-block-on-edge" on page 3024	"traceoptions (Spanning Tree)" on page 29869	

[edit routing-instances {routing-instance-name} protocols (mstp | rstp | vstp) interface {interface-name}]		
"cost" on page 4681	"edge" on page 6197	"mode" on page 10798
"no-root-port" on page 11572	"priority (Protocols STP)" on page 13199	

[edit routing-instances {routing-instance-name} protocols (mstp | rstp)]		
"backup-bridge-priority" on page 2793	"bpdu-destination-mac-address (Spanning Tree)" on page 3027	"bridge-priority" on page 3036
"forward-delay" on page 7393	"hello-time" on page 8041	"max-age" on page 10418
"priority-hold-time" on page 13177	"system-id" on page 28847	"vpls-flush-on-topology-change" on page 30926

[edit routing-instances {routing-instance-name} protocols (ospf | ospf3 | pim) graceful-restart]	
"restart-duration (Protocols)" on page 15651	

[edit routing-instances {routing-instance-name} protocols (ospf | ospf3)]		
"area" on page 2451	"backup-spf-options (Protocols OSPF)" on page 2823	"database-protection" on page 4852
"export (OSPF)" on page 6660	"import (Protocols OSPF)" on page 8463	"lsa-refresh-interval" on page 10182
"no-domain-vpn-tag" on page 11443	"no-nssa-abr" on page 11533	"no-rfc-1583" on page 11571
"overload (Protocols OSPF)" on page 12051	"prefix-export-limit (Protocols OSPF)" on page 13085	"reference-bandwidth (Protocols OSPF)" on page 14226

(Continued)

[edit routing-instances {routing-instance-name} protocols (ospf | ospf3)]		
"rib-group (Protocols OSPF)" on page 15774	"spf-options (Protocols OSPF)" on page 28381	"traceoptions (Protocols OSPF)" on page 29713

[edit routing-instances {routing-instance-name} protocols (ospf | ospf3) area {area-id}]		
"area-range" on page 2453	"interface (Protocols OSPF)" on page 8984	"nssa" on page 11717
"stub" on page 28665		

[edit routing-instances {routing-instance-name} protocols (ospf | ospf3) area {area-id} interface {interface-name}]		
"bandwidth-based-metrics (Protocols OSPF)" on page 2835	"bfd-liveness-detection (Protocols OSPF)" on page 2941	"interface-type (Protocols OSPF)" on page 9091
"link-protection (Protocols OSPF)" on page 9814	"passive (Protocols OSPF)" on page 12223	

[edit routing-instances {routing-instance-name} protocols (ospf | ospf3) area {area-id} interface {interface-name} passive]
"traffic-engineering (Passive TE Mode)" on page 30021

[edit routing-instances {routing-instance-name} protocols (ospf | ospf3) area {area-id} nssa]		
"area-range" on page 2453	"default-lsa" on page 5197	

[edit routing-instances {routing-instance-name} protocols (ospf | ospf3)traffic-engineering]		
"shortcuts (Protocols OSPF)" on page 16991	"tunnel-source-protocol (Protocols OSPF)" on page 30240	

[edit routing-instances {routing-instance-name} protocols (rstp | vstp)]	
"force-version (IEEE 802.1D STP)" on page 7379	

[edit routing-instances {routing-instance-name} protocols amt]		
"relay (AMT Protocol)" on page 14291	"traceoptions (Protocols AMT)" on page 29656	

[edit routing-instances {routing-instance-name} protocols amt relay]		
"accounting (Protocols AMT Interface)" on page 1728	"family (Protocols AMT Relay)" on page 6898	"secret-key-timeout" on page 16423
"tunnel-devices (Protocols AMT)" on page 30196	"tunnel-limit (Protocols AMT)" on page 30212	"unicast-stream-limit (Protocols AMT)" on page 30351

[edit routing-instances {routing-instance-name} protocols amt relay family]	
"inet (AMT Protocol)" on page 8555	

[edit routing-instances {routing-instance-name} protocols amt relay family inet]		
"anycast-prefix" on page 2285	"local-address (Protocols AMT)" on page 9904	

[edit routing-instances {routing-instance-name} protocols bfd]	
"no-issu-timer-negotiation" on page 11487	

[edit routing-instances {routing-instance-name} protocols bgp]		
"accept-remote-nexthop" on page 1640	"advertise-from-main-vpn-tables" on page 2013	"advertise-inactive" on page 2015
"advertise-peer-as" on page 2020	"advertise-prefix-sid" on page 2022	"authentication-algorithm" on page 2577
"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605	"bfd-liveness-detection (BGP)" on page 2918
"cluster" on page 4367	"damping (Protocols BGP)" on page 4822	"defaults (ACX series, JRR200, MX204, PTX5000, and VRR)" on page 5225
"description (Protocols BGP)" on page 5340	"disable (Protocols BGP)" on page 5788	"domain-path-id (Protocols BGP)" on page 5898
"egress-te" on page 6227	"egress-te-set-segment" on page 6236	"enforce-first-as" on page 6343
"export (Protocols BGP)" on page 6677	"family (Protocols BGP)" on page 6900	"forwarding-context (Protocols BGP)" on page 7453
"group (Protocols BGP)" on page 7863	"hold-time (Protocols BGP)" on page 8104	"import (BGP)" on page 8461
"include-mp-next-hop" on page 8532	"ipsec-sa (Protocols BGP)" on page 9308	"keep" on page 9452
"local-address (Protocols BGP)" on page 9905	"local-as" on page 9916	"local-preference" on page 9971
"log-updown" on page 10042	"metric-out" on page 10676	"mtu-discovery" on page 11032

(Continued)

[edit routing-instances {routing-instance-name} protocols bgp]		
"multihop" on page 11098	"multipath (Protocols BGP)" on page 11110	"multipath-build-priority" on page 11107
"no advertise-peer-as" on page 11385	"no-aggregator-id" on page 11387	"no-client-reflect" on page 11424
"out-delay" on page 11961	"outbound-route-filter" on page 11970	"passive (Protocols BGP)" on page 12218
"path-selection" on page 12266	"peer-as (Protocols BGP)" on page 12346	"preference (Protocols BGP)" on page 13038
"remove-private" on page 14368	"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763	"stale-labels-holddown-period" on page 28476
"tcp-aggressive-transmission" on page 28978	"tcp-mss (Protocols BGP)" on page 28998	"tcpao-auth-mismatch" on page 29034
"traceoptions (Protocols BGP)" on page 29662	"traffic-statistics-labeled-path" on page 30049	

[edit routing-instances {name} protocols bgp bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit routing-instances {routing-instance-name} protocols bgp family {family}]		
"damping (Protocols BGP)" on page 4822	"delay-route-advertisements" on page 5262	"iso-vpn" on page 9404
"traffic-engineering (Protocols BGP)" on page 30023		

[edit routing-instances {routing-instance-name} protocols bgp family (inet | inet-vpn | inet6 | inet6-vpn) flow]	
"strip-nexthop" on page 28658	

[edit routing-instances {routing-instance-name} protocols bgp family (inet | inet6)]	
"labeled-unicast (Protocols BGP)" on page 9603	

[edit routing-instances {routing-instance-name} protocols bgp family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651	"prefix-limit" on page 13087	

[edit routing-instances {routing-instance-name} protocols bgp family (inet | inet6) labeled-unicast]	
"rib (Protocols BGP)" on page 15787	

[edit routing-instances {routing-instance-name} protocols bgp family (inet | inet6) unicast]	
"topology (Protocols BGP)" on page 29449	

[edit routing-instances {routing-instance-name} protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow)]		
"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	

[edit routing-instances {routing-instance-name} protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]	
"restarter (Graceful Restart for BGP Restarter)" on page 15667	

[edit routing-instances {routing-instance-name} protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]		
"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481

[edit routing-instances {routing-instance-name} protocols bgp family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit routing-instances {routing-instance-name} protocols bgp family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit routing-instances {routing-instance-name} protocols bgp family inet (labeled-unicast | unicast | multicast) graceful-restart]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit routing-instances {name} protocols bgp family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit routing-instances {routing-instance-name} protocols bgp family inet labeled-unicast]		
"aigp" on page 2092	"entropy-label (protocols bgp)" on page 6413	"explicit-null (Protocols BGP)" on page 6652

(Continued)

[edit routing-instances {routing-instance-name} protocols bgp family inet labeled-unicast]		
"per-prefix-label" on page 12419	"protection (Protocols BGP)" on page 13416	"resolve-vpn" on page 15602

[edit routing-instances {name} protocols bgp family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229
"extended-nexthop" on page 6700	"protection (Protocols BGP)" on page 13416	

[edit routing-instances {routing-instance-name} protocols bgp family inet-mdt]	
"signaling" on page 27949	

[edit routing-instances {routing-instance-name} protocols bgp family inet-mvpn]	
"signaling" on page 27949	

[edit routing-instances {name} protocols bgp family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit routing-instances {routing-instance-name} protocols bgp family inet6 labeled-unicast]		
"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652	"protection (Protocols BGP)" on page 13416

[edit routing-instances {name} protocols bgp family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"protection (Protocols BGP)" on page 13416

[edit routing-instances {instance-name} protocols bgp family l2vpn]	
"auto-discovery-only" on page 2711	

[edit routing-instances {routing-instance-name} protocols bgp family route-target]	
"accepted-prefix-limit" on page 1651	

[edit routing-instances {routing-instance-name} protocols bgp graceful-restart]		
"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488	"restart-time (BGP Graceful Restart)" on page 15663	"stale-routes-time" on page 28478

[edit routing-instances {instance-name} protocols bgp graceful-restart long-lived]	
"receiver (Graceful Restart for BGP Helper)" on page 14127	

[edit routing-instances {instance-name} protocols bgp graceful-restart long-lived advertise-to-non-llgr-neighbor]	
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name}]		
"accept-remote-nexthop" on page 1640	"advertise-external" on page 2011	"advertise-inactive" on page 2015

(Continued)

[edit routing-instances {routing-instance-name} protocols bgp group {group-name}]		
"advertise-peer-as" on page 2020	"allow" on page 2160	"as-override" on page 2489
"authentication-algorithm" on page 2577	"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605
"authentication-key-chains" on page 2610	"bfd-liveness-detection (BGP)" on page 2918	"cluster" on page 4367
"damping (Protocols BGP)" on page 4822	"description (Protocols BGP)" on page 5340	"domain-path-id (Protocols BGP)" on page 5898
"egress-te" on page 6227	"enforce-first-as" on page 6343	"export (Protocols BGP)" on page 6677
"family (Protocols BGP)" on page 6900	"hold-time (Protocols BGP)" on page 8104	"import (BGP)" on page 8461
"include-mp-next-hop" on page 8532	"ipsec-sa (Protocols BGP)" on page 9308	"keep" on page 9452
"local-address (Protocols BGP)" on page 9905	"local-as" on page 9916	"local-preference" on page 9971
"log-updown" on page 10042	"metric-out" on page 10676	"mtu-discovery" on page 11032
"multihop" on page 11098	"multipath (Protocols BGP)" on page 11110	"neighbor (Protocols BGP)" on page 11273
"no advertise-peer-as" on page 11385	"no-aggregator-id" on page 11387	"no-client-reflect" on page 11424
"out-delay" on page 11961	"outbound-route-filter" on page 11970	"passive (Protocols BGP)" on page 12218

(Continued)

[edit routing-instances {routing-instance-name} protocols bgp group {group-name}]		
"peer-as (Protocols BGP)" on page 12346	"peer-as-list" on page 12345	"preference (Protocols BGP)" on page 13038
"remove-private" on page 14368	"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763	"tcp-aggressive-transmission" on page 28978
"tcp-mss (Protocols BGP)" on page 28998	"traceoptions (Protocols BGP)" on page 29662	"type (Protocols BGP)" on page 30278

[edit routing-instances {name} protocols bgp group {group-name} {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit routing-instances {name} protocols bgp group {name} bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit routing-instances {name} protocols bgp group {name} dynamic-neighbor {dyn-name}]		
"allow" on page 2160	"peer-limit" on page 12375	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family {family}]		
"damping (Protocols BGP)" on page 4822	"delay-route-advertisements" on page 5262	"iso-vpn" on page 9404
"route-target (Protocols BGP)" on page 15881	"traffic-engineering (Protocols BGP)" on page 30023	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet flow)]	
"no-validate" on page 11615	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet-vpn | inet6 | inet6-vpn)]	
"flow (protocols bgp)" on page 7250	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet-vpn | inet6 | inet6-vpn) flow]	
"strip-nexthop" on page 28658	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet-vpn) flow]	
"legacy-redirect-ip-action" on page 9711	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet6)]	
"labeled-unicast (Protocols BGP)" on page 9603	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651	"prefix-limit" on page 13087	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet6) labeled-unicast]	
"rib (Protocols BGP)" on page 15787	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family (inet | inet6) unicast]	
"topology (Protocols BGP)" on page 29449	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family (l2vpn | route-target | inet) (labeled-unicast | flow)]		
"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]	
"restarter (Graceful Restart for BGP Restarter)" on page 15667	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]		
"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family inet (labeled-unicast | unicast | multicast) graceful-restart]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit routing-instances {name} protocols bgp group {name} family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family inet labeled-unicast]		
"aigp" on page 2092	"entropy-label (protocols bgp)" on page 6413	"explicit-null (Protocols BGP)" on page 6652
"per-prefix-label" on page 12419	"resolve-vpn" on page 15602	

[edit routing-instances {name} protocols bgp group {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229
"extended-nexthop" on page 6700		

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family inet-mdt]	
"signaling" on page 27949	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family inet-mvpn]	
"signaling" on page 27949	

[edit routing-instances {name} protocols bgp group {name} family inet6]		
"transport (Protocols BGP)" on page 30113		

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family inet6 labeled-unicast]		
"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652	

[edit routing-instances {name} protocols bgp group {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit routing-instances {instance-name} protocols bgp group {group-name} family l2vpn]		
"auto-discovery-only" on page 2711		

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} family route-target]		
"accepted-prefix-limit" on page 1651		

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} graceful-restart]		
"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488		

[edit routing-instances {instance-name} protocols bgp group {group-name} graceful-restart long-lived]		
"receiver (Graceful Restart for BGP Helper)" on page 14127		

[edit routing-instances {instance-name} protocols bgp group {group-name} graceful-restart long-lived advertise-to-non-llgr-neighbor]	
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810	

[edit routing-instances {name} protocols bgp group {name} link-bandwidth]	
"auto-sense" on page 2738	

[edit routing-instances {name} protocols bgp group {name} metric-out]		
"effective-aigp" on page 6203	"minimum-effective-aigp" on page 10730	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address}]		
"accept-remote-nexthop" on page 1640	"advertise-external" on page 2011	"advertise-inactive" on page 2015
"advertise-peer-as" on page 2020	"as-override" on page 2489	"authentication-algorithm" on page 2577
"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605	"bfd-liveness-detection (BGP)" on page 2918
"cluster" on page 4367	"damping (Protocols BGP)" on page 4822	"description (Protocols BGP)" on page 5340
"domain-path-id (Protocols BGP)" on page 5898	"egress-te" on page 6227	"egress-te-adj-segment" on page 6229
"egress-te-node-segment" on page 6234	"enforce-first-as" on page 6343	"export (Protocols BGP)" on page 6677

(Continued)

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address}]		
"family (Protocols BGP)" on page 6900	"hold-time (Protocols BGP)" on page 8104	"import (BGP)" on page 8461
"include-mp-next-hop" on page 8532	"ipsec-sa (Protocols BGP)" on page 9308	"keep" on page 9452
"local-address (Protocols BGP)" on page 9905	"local-as" on page 9916	"local-interface (IPv6)" on page 9952
"local-preference" on page 9971	"log-updown" on page 10042	"metric-out" on page 10676
"mtu-discovery" on page 11032	"multihop" on page 11098	"multipath (Protocols BGP)" on page 11110
"no advertise-peer-as" on page 11385	"no-aggregator-id" on page 11387	"no-client-reflect" on page 11424
"out-delay" on page 11961	"outbound-route-filter" on page 11970	"passive (Protocols BGP)" on page 12218
"peer-as (Protocols BGP)" on page 12346	"preference (Protocols BGP)" on page 13038	"remove-private" on page 14368
"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763	"tcp-aggressive-transmission" on page 28978	"tcp-mss (Protocols BGP)" on page 28998
"traceoptions (Protocols BGP)" on page 29662		

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} (inet | inet6) labeled-unicast]	
"rib (Protocols BGP)" on page 15787	

[edit routing-instances {name} protocols bgp group {group-name} neighbor {address} {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit routing-instances {name} protocols bgp group {name} neighbor {address} bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family {family}]		
"damping (Protocols BGP)" on page 4822	"iso-vpn" on page 9404	"route-target (Protocols BGP)" on page 15881
"traffic-engineering (Protocols BGP)" on page 30023		

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (inet | inet flow)]	
"no-validate" on page 11615	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (inet | inet-vpn | inet6 | inet6-vpn)]	
"flow (protocols bgp)" on page 7250	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (inet | inet-vpn | inet6 | inet6-vpn) flow]	
"strip-nexthop" on page 28658	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (inet | inet-vpn) flow]		
"legacy-redirect-ip-action" on page 9711		

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (inet | inet6)]		
"labeled-unicast (Protocols BGP)" on page 9603	"topology (Protocols BGP)" on page 29449	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]		
"accepted-prefix-limit" on page 1651	"prefix-limit" on page 13087	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (l2vpn | route-target | inet) (labeled-unicast | flow)]		
"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived]		
"restarter (Graceful Restart for BGP Restarter)" on page 15667		

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter]		
"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family {address-family} {subsequent-address-family} graceful-restart]	
"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family inet (labeled-unicast | unicast | multicast)]	
"rib-group (Protocols BGP)" on page 15767	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family inet (labeled-unicast | unicast | multicast) graceful-restart]	
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	

[edit routing-instances {name} protocols bgp group {name} neighbor {name} family inet inet6 unicast]		
"defer-initial-multipath-build" on page 5229	"extended-nexthop" on page 6700	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family inet labeled-unicast]		
"aigp" on page 2092	"entropy-label (protocols bgp)" on page 6413	"explicit-null (Protocols BGP)" on page 6652
"resolve-vpn" on page 15602		

[edit routing-instances {name} protocols bgp group {name} neighbor {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	"defer-initial-multipath-build" on page 5229

(Continued)

[edit routing-instances {name} protocols bgp group {name} neighbor {name} family inet unicast]	
"extended-nexthop" on page 6700	

[edit routing-instances {name} protocols bgp group {name} neighbor {name} family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family inet6 labeled-unicast]		
"aigp" on page 2092	"explicit-null (Protocols BGP)" on page 6652	

[edit routing-instances {name} protocols bgp group {name} neighbor {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit routing-instances {instance-name} protocols bgp group {group-name} neighbor {address} family l2vpn]	
"auto-discovery-only" on page 2711	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} family route-target]	
"accepted-prefix-limit" on page 1651	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} graceful-restart]		
"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488	"receiver (Graceful Restart for BGP Helper)" on page 14127	

[edit routing-instances {instance-name} protocols bgp group {group-name} neighbor {address} graceful-restart long-lived advertise-to-non-llgr-neighbor]	
"omit-no-export (Graceful Restart for BGP Helper)" on page 11810	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} neighbor {address} outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit routing-instances {routing-instance-name} protocols bgp group {group-name} outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit routing-instances {name} protocols bgp multipath]	
"pause-computation-during-churn" on page 12277	

[edit routing-instances {name} protocols bgp neighbor {ip-address} family {name}]	
"delay-route-advertisements" on page 5262	

[edit routing-instances {routing-instance-name} protocols bgp outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit routing-instances {name} protocols bgp source-packet-routing]		
"srgb (Protocols BGP)" on page 28407	"srv6 (Protocols BGP)" on page 28417	

[edit routing-instances {name} protocols evpn]		
"advertise-zero-ethernet-tag" on page 2028	"assisted-replication" on page 2505	"auto-service-id" on page 2739
"control-word" on page 4644	"default-gateway" on page 5190	"designated-forwarder-election-hold-time" on page 5363
"drop-unknown-mac-ip" on page 5984	"encapsulation (protocols evpn)" on page 6284	"extended-vni-all" on page 6709
"extended-vni-list" on page 6710	"flexible-cross-connect-vlan-aware" on page 7230	"flexible-cross-connect-vlan-unaware" on page 7232
"flow-label" on page 7306	"flow-label-receive-static" on page 7309	"flow-label-static" on page 7311
"flow-label-transmit-static" on page 7314	"interface (EVPN Routing Instances)" on page 8882	"interface-mac-ip-limit" on page 8908
"interface-mac-limit (VPLS)" on page 8913	"ip-prefix-support" on page 9255	"label-allocation" on page 9576
"mac-ip-table-size" on page 10231	"mac-pinning" on page 10254	"mac-statistics" on page 10263
"mac-table-aging-time (MX Series)" on page 10265	"mac-table-size (VPLS)" on page 10269	"mark-local-ip-host-routes" on page 10355
"mclag" on page 10584	"multicast-mode (EVPN)" on page 11074	"no-control-word " on page 11426

(Continued)

[edit routing-instances {name} protocols evpn]		
"no-default-gateway-ext-comm" on page 11434	"normalization (Protocols EVPN)" on page 11678	"proxy-mac" on page 13829
"remote-ip-host-routes" on page 14346	"source-packet-routing (Routing Instances Protocols EVPN)" on page 28273	"traceoptions (Protocols EVPN)" on page 29672
"vni-options" on page 30912		

[edit routing-instances {name} protocols evpn assisted-replication]		
"deterministic-ar-policy" on page 5553	"leaf" on page 9689	"replicator" on page 14387

[edit routing-instances {name} protocols evpn encapsulation]	
"mpls-inet6" on page 10977	

[edit routing-instances {routing-instance-name} protocols evpn interconnect]		
"encapsulation (protocols evpn)" on page 6284	"esi" on page 6453	"interconnected-vlan-list" on page 8808
"interconnected-vni-list" on page 8809	"route-distinguisher" on page 15856	"suppress-host-routes-from-dci-to-dc" on page 28728
"vrf-export" on page 30958	"vrf-import" on page 30960	"vrf-target" on page 30968

[edit routing-instances {name} protocols evpn interconnect esi]	
"df-election-type" on page 5576	

[edit routing-instances {instance-name} protocols evpn interface {interface-name}]		
"flow-label-receive-static" on page 7309	"flow-label-transmit-static" on page 7314	"ignore-encapsulation-mismatch" on page 8397
"interface-mac-limit (VPLS)" on page 8913	"static-mac" on page 28549	"vpws-service-id" on page 30954

[edit routing-instances {routing-instance-name} protocols evpn interface {interface-name} interface-mac-limit (VPLS)]	
"packet-action" on page 12129	

[edit routing-instances {name} protocols evpn interface],]	
"static-mac-ip" on page 28587	

[edit routing-instances {routing-instance-name} protocols evpn interface-mac-limit (VPLS)]	
"packet-action" on page 12129	

[edit routing-instances {name} protocols evpn ip-prefix-routes]	
"route-attributes" on page 15850	

[edit routing-instances {routing-instance-name} protocols evpn mac-table-size {limit}]	
"packet-action" on page 12129	

[edit routing-instances {name} protocols evpn mclag]	
"bgp-peer" on page 2965	

[edit routing-instances {name} protocols evpn oism pim-evpn-gateway peg-df-election]	
"preference (PEG DF Election)" on page 13037	

[edit routing-instances {routing-instance-name} protocols evpn vni-options vni {vni-value}]	
"vrf-target" on page 30968	

[edit routing-instances { routing-instance-name} protocols evpn],]	
"duplicate-mac-detection" on page 6068	

[edit routing-instances {routing-instance-name} protocols igmp]	
"amt (IGMP)" on page 2228	

[edit routing-instances {routing-instance-name} protocols igmp amt relay defaults]		
"accounting (Protocols IGMP AMT Interface)" on page 1731	"group-policy (Protocols IGMP AMT Interface)" on page 7856	"query-interval (Protocols IGMP AMT)" on page 13930
"query-response-interval (Protocols IGMP AMT)" on page 13941	"robust-count (Protocols IGMP AMT)" on page 15820	"ssm-map (Protocols IGMP AMT)" on page 28465
"version (Protocols IGMP AMT)" on page 30716		

[edit routing-instances {routing-instance-name} protocols igmp-snooping]		
"evpn-ssm-reports-only" on page 6553	"traceoptions (Protocols IGMP Snooping)" on page 29678	"version (IGMP Snooping)" on page 30704

[edit routing-instances {routing-instance-name} protocols igmp-snooping vlan {vlan-id}]	
"version (IGMP Snooping)" on page 30704	

[edit routing-instances {routing-instance-name} protocols igmp-snooping vlan {vlan-name} no-flood-to-multicast-router-interfaces]	
"no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping)" on page 11628	

[edit routing-instances {name} protocols interface {name}]		
"accept-join-always-from" on page 1632	"action-priority (Routing Instance)" on page 1766	"mac-pinning (EVPN Routing Instances)" on page 10256

[edit routing-instances {routing-instance-name} protocols isis]		
"backup-spf-options (Protocols IS-IS)" on page 2821	"clns-routing" on page 4349	"clns-updown-compatibility" on page 4350
"export (ISIS)" on page 6658	"ignore-attached-bit" on page 8396	"import (Protocols IS-IS)" on page 8476
"interface (Protocols IS-IS)" on page 8966	"label-switched-path (Protocols IS-IS)" on page 9582	"layer2-map" on page 9651
"level (Global IS-IS)" on page 9724	"loose-authentication-check" on page 10137	"lsp-lifetime" on page 10189
"max-areas" on page 10420	"max-lsp-size" on page 10446	"max-snp-size" on page 10483
"no-ipv4-routing" on page 11483	"no-ipv6-routing" on page 11484	"overload (Protocols IS-IS)" on page 12048

(Continued)

[edit routing-instances {routing-instance-name} protocols isis]		
"reference-bandwidth (Protocols IS-IS)" on page 14224	"rib-group (Protocols IS-IS)" on page 15770	"source-packet-routing (Protocols IS-IS)" on page 28266
"spf-options (Protocols IS-IS)" on page 28379	"topologies (Protocols IS-IS)" on page 29438	"traceoptions (Protocols IS-IS)" on page 29684

[edit routing-instances {name} protocols isis backup-spf-options]	
"use-post-convergence-lfa" on page 30531	

[edit routing-instances {name} protocols isis dynamic-overload]		
"auto-recovery" on page 2733	"no-overload-on-prefix-export-limit" on page 11542	

[edit routing-instances {routing-instance-name} protocols isis interface {interface-name}]		
"bfd-liveness-detection (Protocols IS-IS)" on page 2933	"level (IS-IS Interfaces)" on page 9728	"link-protection (Protocols IS-IS)" on page 9811
"no-advertise-adjacency-segment (Protocols IS-IS)" on page 11382	"passive (Protocols IS-IS)" on page 12219	

[edit routing-instances {name} protocols isis interface {name} delay-measurement]	
"advertisement (Delay Measurement)" on page 2031	

[edit routing-instances {routing-instance-name} protocols isis interface {interface-name} level {level-number}]		
"flood-reflector (Protocols IS-IS)" on page 7242	"hold-time (Protocols IS-IS)" on page 8108	"metric (Protocols IS-IS)" on page 10684
"passive (Protocols IS-IS)" on page 12219	"post-convergence-lfa" on page 12903	

[edit routing-instances {name} protocols isis interface {name} level { number} post-convergence-lfa]	
"node-protection" on page 11651	

[edit routing-instances {routing-instance-name} protocols isis interface-group {group-name}]	
"link-group-protection (Protocols IS-IS)" on page 9794	

[edit routing-instances {name} protocols isis interface-group {name} level {name} topology]	
"bandwidth-based-metrics" on page 2837	

[edit routing-instances {routing-instance-name} protocols isis level {level-number}]		
"prefix-export-limit (Protocols IS-IS)" on page 13083	"source-packet-routing (Protocols IS-IS)" on page 28266	

[edit routing-instances {name} protocols isis no-external-export protocol]	
"no-external-export" on page 11454	

[edit routing-instances {name} protocols isis source-packet-routing]		
"flex-algorithm" on page 7221	"srgb (Protocols IS-IS)" on page 28408	"strict-asla-based-flex-algorithm" on page 28645

(Continued)

[edit routing-instances {name} protocols isis source-packet-routing]		
"traffic-statistics (Protocols IS-IS)" on page 30048		

[edit routing-instances {name} protocols isis spf-options]		
"microloop-avoidance" on page 10700		

[edit routing-instances {name} protocols isis traffic-engineering]		
"advertisement" on page 2030	"family (Protocols IS-IS)" on page 6905	

[edit routing-instances {routing-instance-name} protocols l2vpn]		
"control-word (Protocols Layer 2 VPN)" on page 4648	"encapsulation-type (Layer 2 VPNs)" on page 6310	"interface (Protocols Layer 2 VPN)" on page 8977
"no-control-word (Protocols Layer 2 VPN)" on page 11427	"oam (Layer 2 Circuit Protocols)" on page 11748	"site (Layer 2 Circuits)" on page 27987
"traceoptions (Protocols Layer 2 VPN)" on page 29690	"vpls-id" on page 30927	

[edit routing-instances {routing-instance-name} protocols l2vpn interface {interface-name}]		
"ignore-mtu-mismatch" on page 8408	"mtu (interfaces)" on page 11027	

[edit routing-instances {routing-instance-name} protocols l2vpn mesh-group {mesh-group-name}]		
"route-distinguisher" on page 15856	"vpls-id" on page 30927	"vrf-target" on page 30968

[edit routing-instances {routing-instance-name} protocols l2vpn neighbor {address}]		
"encapsulation-type (Layer 2 VPNs)" on page 6310		

[edit routing-instances {instance_name} protocols l2vpn oam]		
"bfd-liveness-detection (Layer 2 VPN and VPLS)" on page 2929	"control-channel (Protocols OAM)" on page 4623	"ping-interval" on page 12545

[edit routing-instances {routing-instance-name} protocols l2vpn site {site-name}]		
"hot-standby (Routing Instance)" on page 8186	"site-identifier (Layer 2 Circuits)" on page 27984	"site-preference" on page 27989
"source-attachment-identifier (Protocols VPWS)" on page 28183		

[edit routing-instances {routing-instance-name} protocols l2vpn site {site-name} interface {interface-name}]		
"description (Protocols Layer 2 VPN)" on page 5343	"remote-site-id" on page 14360	"target-attachment-identifier (Protocols VPWS)" on page 28947

[edit routing-instances {routing-instance-name} protocols l2vpn site name]		
"flow-label-receive" on page 7307	"flow-label-transmit" on page 7312	

[edit routing-instances {routing-instance-name} protocols l2vpn site name interface {interface-name}]		
"flow-label-receive" on page 7307	"flow-label-transmit" on page 7312	

[edit routing-instances {routing-instance-name} protocols ldp]		
"deaggregate" on page 5153	"egress-policy" on page 6215	"explicit-null (Protocols BGP)" on page 6652

(Continued)

[edit routing-instances {routing-instance-name} protocols ldp]		
"explicit-null (Protocols LDP)" on page 6654	"export (Protocols LDP)" on page 6680	"graceful-restart (Protocols LDP)" on page 7773
"igp-synchronization" on page 8414	"import (Protocols LDP)" on page 8477	"interface (Protocols LDP)" on page 8978
"keepalive-interval" on page 9468	"keepalive-timeout" on page 9474	"l2-smart-policy" on page 9528
"label-withdrawal-delay" on page 9598	"log-updown (Protocols LDP)" on page 10043	"no-forwarding" on page 11464
"p2mp (Protocols LDP)" on page 12122	"policing (Protocols LDP)" on page 12662	"session (Protocols LDP)" on page 16845
"session-group" on page 16810	"session-protection" on page 16843	"strict-targeted-hellos" on page 28654
"targeted-hello" on page 28962	"traceoptions (Protocols LDP)" on page 29693	"track-igp-metric" on page 29989
"traffic-statistics (Protocols LDP)" on page 30053		

[edit routing-instances {routing-instance-name} protocols ldp graceful-restart]		
"helper-disable (LDP)" on page 8046	"maximum-neighbor-reconnect-time" on page 10537	"maximum-neighbor-recovery-time" on page 10538
"reconnect-time" on page 14135	"recovery-time (LDP)" on page 14148	

[edit routing-instances {routing-instance-name} protocols ldp interface {interface-name}]		
"disable (Protocols LDP)" on page 5789	"hello-interval (Protocols LDP)" on page 8036	"hold-time (Protocols LDP)" on page 8110
"preference (Protocols LDP)" on page 13040	"transport-address" on page 30118	

[edit routing-instances {routing-instance-name} protocols ldp session {session-address}]		
"authentication-algorithm" on page 2577	"authentication-key (Protocols LDP)" on page 2618	"authentication-key-chain (Protocols LDP)" on page 2607

[edit routing-instances {routing-instance-name} protocols ldp targeted-hello]		
"hello-interval (Protocols LDP)" on page 8036	"hold-time (Protocols LDP)" on page 8110	

[edit routing-instances {routing-instance-name} protocols lldp]		
"interface (LLDP)" on page 8817	"traceoptions (LLDP)" on page 29596	

[edit routing-instances {routing-instance-name} protocols mld-snooping]		
"evpn-ssm-reports-only" on page 6553	"query-interval (Bridge Domains)" on page 13927	"query-response-interval (Bridge Domains)" on page 13938
"version (MLD Snooping)" on page 30712	"vlan (MLD Snooping)" on page 30830	

[edit routing-instances {routing-instance-name} protocols mld-snooping interface {interface-name}]	
"query-last-member-interval (Bridge Domains)" on page 13933	

[edit routing-instances {instance-name} protocols mld-snooping vlan {vlan-name}]		
"interface (MLD Snooping)" on page 8918	"qualified-vlan " on page 13912	"query-last-member-interval (Protocols MLD)" on page 13936
"query-response-interval (Protocols MLD)" on page 13943	"version (MLD Snooping)" on page 30712	

[edit routing-instances {instance-name} protocols mld-snooping vlan {vlan-name} interface {interface-name}]		
"immediate-leave" on page 8454	"multicast-router-interface (MLD Snooping)" on page 11086	

[edit routing-instances {instance-name} protocols mld-snooping vlan {vlan-name} interface {interface-name} static]	
"group (Protocols MLD Snooping)" on page 7872	

[edit routing-instances {routing-instance-name} protocols mld-snooping vlan {vlan-name} no-flood-to-multicast-router-interfaces]	
"no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping)" on page 11628	

[edit routing-instances {name} protocols mpls]	
"cspf-backoff-time" on page 4731	

[edit routing-instances {name} protocols mpls label-switched-path {name}]		
"conditional-metric" on page 4506	"cspf-backoff-time" on page 4731	"pop-and-forward (Protocols MPLS)" on page 12802

[edit routing-instances {name} protocols mpls label-switched-path {name} auto-bandwidth]	
"adjust-threshold-absolute" on page 1954	

[edit routing-instances {name} protocols mpls label-switched-path {name} secondary {name} non-active-hold-priority]	
"non-active-hold-priority" on page 11657	

[edit routing-instances {name} protocols mpls static-label-switched-path]		
"segment" on page 16519	"transit" on page 30064	

[edit routing-instances {routing-instance-name} protocols msdp]		
"active-source-limit" on page 1819	"data-encapsulation" on page 4828	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681
"group (Protocols MSDP)" on page 7873	"hold-time (Protocols MSDP)" on page 8112	"import (Protocols MSDP)" on page 8478
"keep-alive (Protocols MSDP)" on page 9457	"local-address (Protocols MSDP) " on page 9910	"peer (Protocols MSDP)" on page 12380
"rib-group (Protocols MSDP)" on page 15772	"sa-hold-time (Protocols MSDP)" on page 16262	"source (Protocols MSDP) " on page 28308
"traceoptions (Protocols MSDP)" on page 29703		

[edit routing-instances {routing-instance-name} protocols msdp active-source-limit]		
"log-interval (Protocols MSDP)" on page 10014	"log-warning (Protocols MSDP)" on page 10049	"maximum (MSDP Active Source Messages) " on page 10536
"threshold (MSDP Active Source Messages) " on page 29327		

[edit routing-instances {routing-instance-name} protocols msdp group {group-name}]		
"default-peer" on page 5200	"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731
"export (Protocols MSDP)" on page 6681	"import (Protocols MSDP)" on page 8478	"local-address (Protocols MSDP) " on page 9910
"mode (Protocols MSDP)" on page 10814	"peer (Protocols MSDP)" on page 12380	"traceoptions (Protocols MSDP)" on page 29703

[edit routing-instances {routing-instance-name} protocols msdp group {group-name} peer {address}]		
"active-source-limit" on page 1819	"authentication-key (Protocols MSDP)" on page 2600	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681
"hold-time (Protocols MSDP)" on page 8112	"import (Protocols MSDP)" on page 8478	"keep-alive (Protocols MSDP)" on page 9457
"local-address (Protocols MSDP) " on page 9910	"sa-hold-time (Protocols MSDP)" on page 16262	"traceoptions (Protocols MSDP)" on page 29703

[edit routing-instances {routing-instance-name} protocols msdp peer {address}]		
"active-source-limit" on page 1819	"authentication-key (Protocols MSDP)" on page 2600	"default-peer" on page 5200
"disable (Protocols MSDP)" on page 5795	"disable (Protocols)" on page 5731	"export (Protocols MSDP)" on page 6681
"hold-time (Protocols MSDP)" on page 8112	"import (Protocols MSDP)" on page 8478	"keep-alive (Protocols MSDP)" on page 9457
"local-address (Protocols MSDP) " on page 9910	"sa-hold-time (Protocols MSDP)" on page 16262	"traceoptions (Protocols MSDP)" on page 29703

[edit routing-instances {routing-instance-name} protocols msdp source {ip-address}/{prefix-length}]	
"active-source-limit" on page 1819	

[edit routing-instances {routing-instance-name} protocols mstp]		
"configuration-name" on page 4516	"disable (STP)" on page 5739	"max-hops" on page 10441
"msti" on page 11003	"revision-level" on page 15735	

[edit routing-instances {routing-instance-name} protocols mstp msti {msti-id}]		
"bridge-priority" on page 3036	"vlan (MSTP)" on page 30835	

[edit routing-instances {routing-instance-name} protocols mstp msti {msti-id} interface {interface-name}]		
"cost" on page 4681	"edge" on page 6197	"priority (Protocols STP)" on page 13199

[edit routing-instances {routing-instance-name} protocols mvpn]		
"hot-root-standby (MBGP MVPN)" on page 8184	"mvpn-mode" on page 11142	"route-target (Protocols MVPN)" on page 15883
"sender-based-rpf (MBGP MVPN)" on page 16552	"static-umh (MBGP MVPN)" on page 28585	"traceoptions (Protocols MVPN)" on page 29706
"unicast-umh-election" on page 30353		

[edit routing-instances {routing-instance-name} protocols mvpn family ({inet | inet6})]	
"autodiscovery-only" on page 2752	

[edit routing-instances {routing-instance-name} protocols mvpn family ({inet | inet6}) autodiscovery-only]	
"intra-as" on page 9226	

[edit routing-instances {routing-instance-name} protocols mvpn family ({inet | inet6}) autodiscovery-only intra-as]	
"inclusive" on page 8539	

[edit routing-instances {routing-instance-name} protocols mvpn hot-root-standby]	
"min-rate" on page 10714	

[edit routing-instances {instance-name} protocols mvpn mvpn-mode]		
"rpt-spt" on page 16040	"spt-only" on page 28400	

[edit routing-instances {instance-name} protocols mvpn mvpn-mode spt-only]		
"convert-sa-to-msdp" on page 4657	"source-active-advertisement" on page 28114	"source-active-advertisement" on page 28114

[edit routing-instances protocols mvpn mvpn-mode spt-only source-active-advertisement]		
"dampen (Protocols MVPN)" on page 4813	"dampen (Protocols MVPN)" on page 4813	"min-rate (source-active-advertisement)" on page 10718
"min-rate (source-active-advertisement)" on page 10718		

[edit routing-instances {routing-instance-name} protocols mvpn route-target]		
"export-target" on page 6693	"import-target" on page 8489	

[edit routing-instances {routing-instance-name} protocols mvpn route-target import-target]		
"target (Routing Instances MVPN)" on page 28954	"unicast (Route Target Community)" on page 30350	

[edit routing-instances {routing-instance-name} protocols mvpn static-umh]		
"backup (MBGP MVPN)" on page 2800	"primary (MBGP MVPN)" on page 13143	"source-tree (MBGP MVPN)" on page 28330

[edit routing-instances {routing-instance-name} protocols mvrp]	
"no-dynamic-vlan" on page 11446	

[edit routing-instances {routing-instance-name} protocols mvrp] (for virtual switch instance type)]	
"traceoptions (MVRP)" on page 29607	

[edit routing-instances {routing-instance-name} protocols mvrp] (for virtual switch instance type) [edit logical-systems {logical-system-name} protocols mvrp]	
"interface (MVRP)" on page 8931	

[edit routing-instances {routing-instance-name} protocols mvrp] (for virtual switch instance type), [edit routing-instances {routing-instance-name} protocols mvrp interface (all | {interface-name})] (for virtual switch instance type)]	
"join-timer (MVRP)" on page 9434	

[edit routing-instances {routing-instance-name} protocols mvrp interface (all | {interface-name})] (for virtual switch instance type)[edit logical-systems {logical-system-name} protocols mvrp interface (all | {interface-name})]	
"registration" on page 14268	

[edit routing-instances {routing-instance-name} protocols mvrp interface (all | {interface-name})]	
"point-to-point (MVRP) " on page 12620	

[edit routing-instances {routing-instance-name} protocols ospf]		
"graceful-restart (Protocols OSPF)" on page 7775	"sham-link" on page 16921	"source-packet-routing (Protocols IS-IS)" on page 28266
"source-packet-routing (Protocols IS-IS)" on page 28266	"source-packet-routing (Protocols OSPF)" on page 28270	"topology (OSPF)" on page 29445

[edit routing-instances {routing-instance-name} protocols ospf area {area-id}]		
"label-switched-path (Protocols OSPF)" on page 9592	"network-summary-export" on page 11325	"network-summary-import" on page 11326
"sham-link-remote" on page 16923	"virtual-link" on page 30763	

[edit routing-instances {routing-instance-name} protocols ospf area {area-id} interface {interface-name}]		
"application-specific (Protocols OSPF)" on page 2389	"authentication (Protocols OSPF)" on page 2570	"mtu (interfaces)" on page 11027
"no-advertise-adjacency-segment (Protocols OSPF)" on page 11383	"post-convergence-lfa (Protocols OSPF)" on page 12904	"topology (OSPF Interface)" on page 29446

[edit routing-instances {routing-instance-name} protocols ospf area {area-id} interface {interface-name} topology {topology-name}]	
"bandwidth-based-metrics (Protocols OSPF)" on page 2835	

[edit routing-instances {routing-instance-name} protocols ospf area {area-id} interface {interface-name} topology (default | {name})]	
"link-protection (Protocols OSPF)" on page 9814	

[edit routing-instances {routing-instance-name} protocols ospf area {area-id} sham-link-remote {address}]	
"metric (Protocols OSPF Sham Link)" on page 10688	

[edit routing-instances {routing-instance-name} protocols ospf area {area-id} virtual-link]		
"application-specific (Protocols OSPF)" on page 2389	"authentication (Protocols OSPF)" on page 2570	

[edit routing-instances {routing-instance-name} protocols ospf backup-spf-options]		
"no-source-packet-routing (Protocols OSPF)" on page 11583	"use-post-convergence-lfa (Protocols OSPF)" on page 30533	

[edit routing-instances {name} protocols ospf overload]		
"allow-route-leaking" on page 2202	"as-external" on page 2486	"stub-network" on page 28667

[edit routing-instances {name} protocols ospf source-packet-routing]		
"ldp-stitching (Protocols OSPF)" on page 9686	"node-segment (Protocols OSPF)" on page 11654	

[edit routing-instances {name} protocols ospf spf-options]	
"microloop-avoidance (Protocols OSPF)" on page 10702	

[edit routing-instances {routing-instance-name} protocols ospf topology {name}]	
"topology-id" on page 29441	

[edit routing-instances {routing-instance-name} protocols ospf topology (default | ipv4-multicast | {name}]	
"overload (Protocols OSPF)" on page 12051	

[edit routing-instances {routing-instance-name} protocols ospf topology (default | ipv4-multicast | {name})]		
"prefix-export-limit (Protocols OSPF)" on page 13085	"spf-options (Protocols OSPF)" on page 28381	

[edit routing-instances {routing-instance-name} protocols ospf topology (default | {name})]	
"backup-spf-options (Protocols OSPF)" on page 2823	

[edit routing-instances {routing-instance-name} protocols ospf3]		
"realm" on page 14093		

[edit routing-instances {routing-instance-name} protocols ospf3 area {area-id}]		
"inter-area-prefix-export" on page 8782	"inter-area-prefix-import" on page 8783	

[edit routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) area {area-id} interface {interface-name}]		
"interface-type (Protocols OSPF)" on page 9091		

[edit routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)]		
"area" on page 2451	"export (OSPF)" on page 6660	"import (Protocols OSPF)" on page 8463
"lsa-refresh-interval" on page 10182	"no-nssa-abr" on page 11533	"no-rfc-1583" on page 11571
"nssa" on page 11717	"overload (Protocols OSPF)" on page 12051	"prefix-export-limit (Protocols OSPF)" on page 13085
"reference-bandwidth (Protocols OSPF)" on page 14226	"rib-group (Protocols OSPF)" on page 15774	"spf-options (Protocols OSPF)" on page 28381
"stub" on page 28665	"traceoptions (Protocols OSPF)" on page 29713	

[edit routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id}]		
"inter-area-prefix-import" on page 8783	"interface (Protocols OSPF)" on page 8984	

[edit routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} interface {interface-name}]		
"bandwidth-based-metrics (Protocols OSPF)" on page 2835	"bfd-liveness-detection (Protocols OSPF)" on page 2941	"passive (Protocols OSPF)" on page 12223

[edit routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} interface {interface-name} passive]	
"traffic-engineering (Passive TE Mode)" on page 30021	

[edit routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id} nssa]	
"default-lsa" on page 5197	

[edit routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast)]	
"database-protection" on page 4852	

[edit routing-instances {routing-instance-name} protocols ospf3 realm (ipv4-unicast | ipv4-muticast | ipv6-multicast) area {area-id}]	
"inter-area-prefix-export" on page 8782	

[edit routing-instances {routing-instance-name} protocols ospf3 realm ipv4-unicast area {area-id}]	
"link-protection (Protocols OSPF)" on page 9814	

[edit routing-instances {routing-instance-name} protocols pim]		
"assert-timeout" on page 2500	"dense-groups" on page 5315	"disable (PIM)" on page 5782
"disable (Protocols)" on page 5731	"disable-packet-register" on page 5813	"dr-election-on-p2p" on page 5937
"export (Protocols PIM)" on page 6683	"family (Protocols PIM)" on page 6907	"graceful-restart (Protocols PIM)" on page 7777
"idle-standby-path-switchover-delay" on page 8301	"import (Protocols PIM)" on page 8480	"interface (Protocols PIM)" on page 8991
"join-load-balance" on page 9431	"join-prune-timeout" on page 9433	"mdt" on page 10590
"mvpn (Draft-Rosen MVPN)" on page 11138	"override-interval" on page 12058	"propagation-delay" on page 13381
"reset-tracking-bit" on page 15583	"rib-group (Protocols PIM)" on page 15775	"rp" on page 15996
"rpf-selection" on page 16015	"rpf-vector (PIM)" on page 16017	"secondary-vtep-address" on page 16419
"sglimit" on page 16919	"spt-threshold" on page 28401	"standby-path-creation-delay" on page 28488
"traceoptions (Protocols PIM)" on page 29719	"tunnel-devices (Tunnel-Capable PICs) " on page 30198	"vpn-group-address" on page 30939

[edit routing-instances {routing-instance-name} protocols pim family (inet | inet6)]		
"disable (PIM)" on page 5782	"disable (Protocols)" on page 5731	

[edit routing-instances {routing-instance-name} protocols pim graceful-restart]	
"no-bidirectional-mode" on page 11416	

[edit routing-instances {routing-instance-name} protocols pim interface {interface-name}]		
"accept-remote-source" on page 1642	"dual-dr" on page 6045	"stickydr" on page 28609

[edit routing-instances {routing-instance-name} protocols pim interface (Protocols PIM) {interface-name}]		
"bidirectional (Interface)" on page 2969	"disable (PIM)" on page 5782	"disable (Protocols)" on page 5731
"distributed-dr" on page 5833	"family (Protocols PIM Interface)" on page 6909	"family (Protocols PIM)" on page 6907
"hello-interval" on page 8032	"mode (Protocols PIM)" on page 10816	"neighbor-policy" on page 11272
"override-interval" on page 12058	"priority (PIM Interfaces)" on page 13184	"propagation-delay" on page 13381
"reset-tracking-bit" on page 15583	"version (Protocols PIM)" on page 30719	

[edit routing-instances {routing-instance-name} protocols pim interface (Protocols PIM) {interface-name} bfd-liveness-detection]		
"minimum-interval (PIM BFD Liveness Detection)" on page 10738	"minimum-receive-interval (Protocols PIM)" on page 10751	"multiplier (Protocols PIM)" on page 11120

(Continued)

[edit routing-instances {routing-instance-name} protocols pim interface (Protocols PIM) {interface-name} bfd-liveness-detection]		
"version (Protocols BFD)" on page 30714		

[edit routing-instances {routing-instance-name} protocols pim interface (Protocols PIM) {interface-name} bidirectional]		
"df-election" on page 5572		

[edit routing-instances {routing-instance-name} protocols pim interface (Protocols PIM) {interface-name} bidirectional df-election]		
"backoff-period" on page 2791	"offer-period" on page 11785	"robustness-count" on page 15823

[edit routing-instances {routing-instance-name} protocols pim interface (Protocols PIM) {interface-name} family (inet | inet6)]		
"bfd-liveness-detection (Protocols PIM)" on page 2945		

[edit routing-instances {routing-instance-name} protocols pim interface {interface-name} bfd-liveness-detection]		
"detection-time (BFD for PIM)" on page 5543	"no-adaptation (PIM BFD Liveness Detection)" on page 11377	"transmit-interval (PIM BFD Liveness Detection)" on page 30096

[edit routing-instances {routing-instance-name} protocols pim interface {interface-name} bfd-liveness-detection authentication]		
"algorithm" on page 2144	"loose-check" on page 10138	

[edit routing-instances {routing-instance-name} protocols pim interface {interface-name} bfd-liveness-detection detection-time]		
"threshold (PIM BFD Detection Time)" on page 29331		

[edit routing-instances {routing-instance-name} protocols pim interface {interface-name} bfd-liveness-detection transmit-interval]		
"minimum-interval (PIM BFD Transmit Interval) " on page 10739	"threshold (PIM BFD Transmit Interval)" on page 29333	

[edit routing-instances {routing-instance-name} protocols pim interface family (inet | inet6) {interface-name} bfd-liveness-detection]	
"authentication (Protocols PIM)" on page 2647	

[edit routing-instances {routing-instance-name} protocols pim interface {interface-name} family {inet | inet6} bfd-liveness-detection authentication]	
"key-chain (Protocols PIM)" on page 9490	

[edit routing-instances {routing-instance-name} protocols pim local family (inet | inet6) anycast-pim]	
"rp-set" on page 16001	

[edit routing-instances {routing-instance-name} protocols pim mdt]		
"threshold (Routing Instances)" on page 29338	"tunnel-limit (Routing Instances)" on page 30213	

[edit routing-instances {routing-instance-name} protocols pim mdt threshold]		
"group (Protocols PIM)" on page 7876		

[edit routing-instances {routing-instance-name} protocols pim mdt threshold group {group-address}]		
"source (Routing Instances)" on page 28312		

[edit routing-instances {routing-instance-name} protocols pim mdt threshold group {group-address} source {source-address}]		
"rate (Routing Instances)" on page 14073		

[edit routing-instances {routing-instance-name} protocols pim mvpn family {inet}]		
"autodiscovery" on page 2750		

[edit routing-instances {routing-instance-name} protocols pim mvpn family (inet | inet6)]		
"disable (PIM)" on page 5782	"disable (Protocols)" on page 5731	

[edit routing-instances {routing-instance-name} protocols pim mvpn family ({inet | inet6}) autodiscovery]		
"inet-mdt" on page 8570		

[edit routing-instances {routing-instance-name} protocols pim rp]		
"auto-rp" on page 2735	"bidirectional (RP)" on page 2971	"bootstrap" on page 3015
"bootstrap-export" on page 3016	"bootstrap-import" on page 3018	"bootstrap-priority" on page 3019

(Continued)

[edit routing-instances {routing-instance-name} protocols pim rp]		
"dr-register-policy" on page 5938	"embedded-rp" on page 6248	"group-rp-mapping" on page 7891
"local" on page 9899	"register-limit" on page 14264	"rp-register-policy" on page 16000
"static (Protocols PIM)" on page 28562		

[edit routing-instances {routing-instance-name} protocols pim rp auto-rp]	
"mapping-agent-election" on page 10342	

[edit routing-instances {routing-instance-name} protocols pim rp bidirectional]	
"address (Bidirectional Rendezvous Points)" on page 1865	

[edit routing-instances {instance-name} protocols pim rp bidirectional address {address}]		
"group-ranges" on page 7885	"hold-time (Protocols PIM)" on page 8114	"priority (PIM RPs)" on page 13186

[edit routing-instances {routing-instance-name} protocols pim rp bootstrap]	
"family (Bootstrap)" on page 6846	

[edit routing-instances {routing-instance-name} protocols pim rp bootstrap (inet | inet6)]		
"import (Protocols PIM Bootstrap)" on page 8482	"priority (Bootstrap)" on page 13157	

[edit routing-instances {routing-instance-name} protocols pim rp bootstrap family (inet | inet6)]		
"export (Bootstrap)" on page 6663		

[edit routing-instances {routing-instance-name} protocols pim rp embedded-rp]		
"group-ranges" on page 7885	"maximum-rps" on page 10556	

[edit routing-instances {routing-instance-name} protocols pim rp group-rp-mapping]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit routing-instances {routing-instance-name} protocols pim rp local]		
"family (Local RP)" on page 6883	"override (PIM Static RP)" on page 12061	

[edit routing-instances {routing-instance-name} protocols pim rp local (inet | inet6) anycast-pim rp-set]		
"address (Anycast RPs)" on page 1856		

[edit routing-instances {routing-instance-name} protocols pim rp local family (inet | inet6)]		
"address (Local RPs)" on page 1885	"anycast-pim" on page 2283	"disable (PIM)" on page 5782
"disable (Protocols)" on page 5731	"group-ranges" on page 7885	"hold-time (Protocols PIM)" on page 8114
"priority (PIM RPs)" on page 13186		

[edit routing-instances {routing-instance-name} protocols pim rp local family (inet | inet6) anycast-pim]	
"local-address (Protocols PIM)" on page 9911	

[edit routing-instances {instance-name} protocols pim rp local family inet]	
"override (PIM Static RP)" on page 12061	

[edit routing-instances {instance-name} protocols pim rp local family inet6]	
"override (PIM Static RP)" on page 12061	

[edit routing-instances {routing-instance-name} protocols pim rp register-limit]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit routing-instances {name} protocols pim rp rpf-vector]	
"policy (PIM rpf-vector) " on page 12703	

[edit routing-instances {routing-instance-name} protocols pim rp static]	
"address (Static RPs)" on page 1933	

[edit routing-instances {routing-instance-name} protocols pim rp static address {address}]		
"group-ranges" on page 7885	"override (PIM Static RP)" on page 12061	"version (Protocols PIM)" on page 30719

[edit routing-instances {routing-instance-name} protocols pim rpf-selection group {group-address}]		
"source (PIM RPF Selection)" on page 28275	"wildcard-source (PIM RPF Selection)" on page 31119	

[edit routing-instances {routing-instance-name} protocols pim rpf-selection group {group-address} source {source-address}]		
"next-hop (PIM RPF Selection)" on page 11347	"prefix-list (PIM RPF Selection)" on page 13097	

[edit routing-instances {routing-instance-name} protocols pim rpf-selection group {group-address} wildcard-source]		
"next-hop (PIM RPF Selection)" on page 11347	"prefix-list (PIM RPF Selection)" on page 13097	

[edit routing-instances {routing-instance-name} protocols pim rpf-selection prefix-list {prefix-list-addresses}]		
"source (PIM RPF Selection)" on page 28275	"wildcard-source (PIM RPF Selection)" on page 31119	

[edit routing-instances {routing-instance-name} protocols pim rpf-selection prefix-list {prefix-list-addresses} source {source-address}]		
"next-hop (PIM RPF Selection)" on page 11347	"prefix-list (PIM RPF Selection)" on page 13097	

[edit routing-instances {routing-instance-name} protocols pim rpf-selection prefix-list {prefix-list-addresses} wildcard-source]		
"next-hop (PIM RPF Selection)" on page 11347	"prefix-list (PIM RPF Selection)" on page 13097	

[edit routing-instances {routing-instance-name} protocols pim sglimit]		
"log-interval (PIM Entries)" on page 10012	"log-interval (PIM Entries)" on page 10012	"maximum (PIM Entries)" on page 10547
"maximum (PIM Entries)" on page 10547	"threshold (PIM Entries)" on page 29334	"threshold (PIM Entries)" on page 29334

[edit routing-instances {routing-instance-name} protocols pim spt-threshold]	
"infinity" on page 8622	

[edit routing-instances {<instance-name>} protocols pim-snooping]		
"traceoptions (PIM Snooping)" on page 29638	"vlan (PIM Snooping)" on page 30837	

[edit routing-instances {<instance-name>} protocols pim-snooping traceoptions]	
"no-dr-flood (PIM Snooping)" on page 11445	

[edit routing-instances {<instance-name>} protocols pim-snooping vlan {<vlan-id>}]	
"no-dr-flood (PIM Snooping)" on page 11445	

[edit routing-instances {routing-instance-name} protocols rip]		
"authentication-key (Protocols RIP)" on page 2619	"authentication-type (Protocols RIP)" on page 2669	"check-zero" on page 3309
"group (Protocols RIP)" on page 7877	"holddown (Protocols RIP)" on page 8125	"import (Protocols RIP)" on page 8483
"message-size" on page 10658	"metric-in (Protocols RIP)" on page 10671	"receive (Protocols RIP)" on page 14117
"rib-group (Protocols RIP)" on page 15777	"route-timeout (Protocols RIP)" on page 15885	"send (Protocols RIP)" on page 16546
"traceoptions (Protocols RIP)" on page 29727	"update-interval (Protocols RIP)" on page 30441	

[edit routing-instances {routing-instance-name} protocols rip group {group-name}]		
"demand-circuit (Protocols RIP)" on page 5294	"export (RIP)" on page 6661	"max-retrans-time" on page 10464
"neighbor" on page 11264	"preference (Protocols RIP)" on page 13043	"route-timeout (Protocols RIP)" on page 15885
"update-interval (Protocols RIP)" on page 30441		

[edit routing-instances {routing-instance-name} protocols rip group {group-name} neighbor {neighbor-name}]		
"any-sender" on page 2281	"authentication-key (Protocols RIP)" on page 2619	"authentication-type (Protocols RIP)" on page 2669
"bfd-liveness-detection (Protocols RIP)" on page 2947	"check-zero" on page 3309	"demand-circuit (Protocols RIP)" on page 5294

(Continued)

[edit routing-instances {routing-instance-name} protocols rip group {group-name} neighbor {neighbor-name}]		
"import (Protocols RIP)" on page 8483	"max-retrans-time" on page 10464	"message-size" on page 10658
"metric-in (Protocols RIP)" on page 10671	"metric-out (Protocols RIP)" on page 10674	"receive (Protocols RIP)" on page 14117
"route-timeout (Protocols RIP)" on page 15885	"send (Protocols RIP)" on page 16546	"update-interval (Protocols RIP)" on page 30441

[edit routing-instances {routing-instance-name} protocols ripng]		
"graceful-restart (Protocols RIPng)" on page 7780	"group (Protocols RIPng)" on page 7880	"holddown (Protocols RIPng)" on page 8127
"import (Protocols RIPng)" on page 8485	"metric-in (Protocols RIPng)" on page 10673	"receive (Protocols RIPng)" on page 14119
"route-timeout (Protocols RIPng)" on page 15887	"send (Protocols RIPng)" on page 16548	"traceoptions (Protocols RIPng)" on page 29730
"update-interval (Protocols RIPng)" on page 30443		

[edit routing-instances {routing-instance-name} protocols ripng group {group-name}]		
"export (Protocols RIPng)" on page 6686	"neighbor (Protocols RIPng)" on page 11282	"preference (Protocols RIPng)" on page 13045

[edit routing-instances {routing-instance-name} protocols ripng group {group-name} neighbor {neighbor-name}]		
"import (Protocols RIPng)" on page 8485	"metric-in (Protocols RIPng)" on page 10673	"metric-out (Protocols RIPng)" on page 10679

(Continued)

[edit routing-instances {routing-instance-name} protocols ripng group {group-name} neighbor {neighbor-name}]		
"receive (Protocols RIPng)" on page 14119	"send (Protocols RIPng)" on page 16548	

[edit routing-instances {routing-instance-name} protocols rstp]	
"extended-system-id" on page 6706	

[edit routing-instances {name} protocols rsvp]	
"pop-and-forward (Protocols RSVP)" on page 12803	

[edit routing-instances {name} protocols rsvp interface]	
"update-threshold-max-reservable" on page 30458	

[edit routing-instances {name} protocols rsvp interface {name} non-rsvp-bandwdith]	
"local-bw-override-threshold" on page 9925	

[edit routing-instances {routing-instance-name} protocols vpls]		
"associate-profile" on page 2513	"bum-hashing" on page 3070	"connectivity-type" on page 4563
"control-word" on page 4644	"encapsulation-type (Layer 2 VPNs)" on page 6310	"flow-label-receive" on page 7307
"flow-label-transmit" on page 7312	"ignore-mtu-mismatch" on page 8408	"import-labeled-routes (Routing Instances VPLS)" on page 8469

(Continued)

[edit routing-instances {routing-instance-name} protocols vpls]		
"interface (VPLS Routing Instances)" on page 9104	"label-block-size" on page 9577	"mac-flush" on page 10229
"mac-table-aging-time (MX Series)" on page 10265	"mac-table-size (VPLS)" on page 10269	"mesh-group (Protocols VPLS)" on page 10641
"multi-homing (VPLS Multihoming for FEC 129)" on page 11052	"neighbor (Protocols VPLS)" on page 11284	"no-control-word " on page 11426
"no-tunnel-services" on page 11605	"oam (Layer 2 Circuit Protocols)" on page 11748	"site (VPLS Multihoming for FEC 128)" on page 27994
"site-range" on page 27990	"traceoptions (Protocols VPLS)" on page 29744	"tunnel-services (Routing Instances VPLS)" on page 30236
"vpls-id" on page 30927		

[edit routing-instances {routing-instance-name} protocols vpls mesh-group {mesh-group-name}]		
"associate-profile" on page 2513	"interface (VPLS Mesh-Group)" on page 9102	"local-switching (VPLS)" on page 9979
"mac-flush" on page 10229	"neighbor (Protocols VPLS)" on page 11284	"peer-as (VPLS)" on page 12348
"route-distinguisher" on page 15856	"vpls-id" on page 30927	"vrf-export" on page 30958
"vrf-import" on page 30960	"vrf-target" on page 30968	

[edit routing-instances {routing-instance-name} protocols vpls mesh-group {mesh-group-name} neighbor {neighbor-id}]		
"ignore-encapsulation-mismatch" on page 8397	"static (Protocols VPLS)" on page 28564	"vpls-id-list (protocols vpls mesh-group)" on page 30929

[edit routing-instances {routing-instance-name} protocols vpls mesh-group {mesh-group-name} neighbor neighbor-id]	
"associate-profile" on page 2513	

[edit routing-instances {instance_name} protocols vpls mesh-group {mesh-group-name} neighbor {neighbor-id} oam]		
"bfd-liveness-detection (Layer 2 VPN and VPLS)" on page 2929	"ping-interval" on page 12545	

[edit routing-instances {routing-instance-name} protocols vpls mesh-group {mesh-group-name} neighbor {address} pseudowire-status-tlv]	
"hot-standby-vc-on (Protocols Layer 2 Circuit)" on page 8190	

[edit routing-instances {routing-instance-name} protocols vpls mesh-group {mesh-group-name} neighbor {address} static]		
"control-word" on page 4644	"no-control-word " on page 11426	

[edit routing-instances {instance-name} protocols vpls multi-homing]		
"active-interface (VPLS Multihoming)" on page 1805	"peer-active (VPLS Multihoming for FEC 129)" on page 12341	"site (VPLS Multihoming for FEC 129)" on page 27995

[edit routing-instances {instance-name} protocols vpls multi-homing active-interface]		
"any (VPLS Multihoming)" on page 2282	"primary (VPLS Multihoming)" on page 13151	

[edit routing-instances {instance-name} protocols vpls multi-homing site {site-name}]		
"active-interface (VPLS Multihoming)" on page 1805	"identifier (VPLS Multihoming for FEC 129)" on page 8288	"interface (VPLS Multihoming for FEC 129)" on page 9103
"peer-active (VPLS Multihoming for FEC 129)" on page 12341	"preference (Site-Level Preference for VPLS Multihoming for FEC 129)" on page 13049	

[edit routing-instances {instance-name} protocols vpls multi-homing site {site-name} active-interface]		
"any (VPLS Multihoming)" on page 2282	"primary (VPLS Multihoming)" on page 13151	

[edit routing-instances {instance-name} protocols vpls multi-homing site {site-name} interface {interface-name}]	
"preference (Interface-Level Preference for VPLS Multihoming for FEC 129)" on page 13033	

[edit routing-instances {routing-instance-name} protocols vpls neighbor {address}]		
"backup-neighbor" on page 2802	"encapsulation-type (Layer 2 VPNs)" on page 6310	"ignore-encapsulation-mismatch" on page 8397
"oam (Layer 2 Circuit Protocols)" on page 11748	"revert-time (Protocols Layer 2 Circuits)" on page 15729	"static (Protocols VPLS)" on page 28564
"switchover-delay" on page 28759		

[edit routing-instances {routing-instance-name} protocols vpls neighbor {address} backup-neighbor {address}]		
"community (Protocols Layer 2 Circuit)" on page 4469	"hot-standby (Protocols Layer 2 Circuit)" on page 8189	"psn-tunnel-endpoint" on page 13851
"standby (Protocols Layer 2 Circuit)" on page 28489	"static (Protocols VPLS)" on page 28564	

[edit routing-instances {instance_name} protocols vpls neighbor {neighbor-id} oam]		
"bfd-liveness-detection (Layer 2 VPN and VPLS)" on page 2929	"ping-interval" on page 12545	

[edit routing-instances {routing-instance-name} protocols vpls neighbor {address} static]		
"control-word" on page 4644	"no-control-word " on page 11426	

[edit routing-instances {instance_name} protocols vpls oam]		
"bfd-liveness-detection (Layer 2 VPN and VPLS)" on page 2929	"control-channel (Protocols OAM)" on page 4623	"ping-interval" on page 12545

[edit routing-instances {routing-instance-name} protocols vpls site {site-name}]		
"automatic-site-id" on page 2758	"best-site" on page 2908	"interface (VPLS Routing Instances)" on page 9104
"mesh-group (Protocols VPLS)" on page 10641	"multi-homing (VPLS Multihoming for BGP VPLS)" on page 11051	"site-identifier (VPLS)" on page 27985
"site-preference" on page 27989		

[edit routing-instances {routing-instance-name} protocols vpls site {site-name} interfaces {interface-name}]	
"interface-mac-limit (VPLS)" on page 8913	

[edit routing-instances {routing-instance-name} protocols vstp vlan {vlan-id}]		
"backup-bridge-priority" on page 2793	"bridge-priority" on page 3036	"forward-delay" on page 7393
"hello-time" on page 8041	"interface (Spanning Tree)" on page 9057	"max-age" on page 10418
"system-id" on page 28847	"vpls-flush-on-topology-change" on page 30926	

[edit routing-instances {routing-instance-name} protocols vstp vlan {vlan-id} interface {interface-name}]		
"cost" on page 4681	"edge" on page 6197	"mode" on page 10798
"no-root-port" on page 11572	"priority (Protocols STP)" on page 13199	

[provider-tunnel]

[edit routing-instances {routing-instance-name} provider-tunnel]		
"ingress-replication" on page 8647	"inter-as (Routing Instances)" on page 8785	"inter-region" on page 8789
"label-switched-path-template (Container LSP)" on page 9595	"label-switched-path-template (Multicast)" on page 9596	"ldp-p2mp" on page 9685

(Continued)

[edit routing-instances {routing-instance-name} provider-tunnel]		
"pim-asm" on page 12514	"rsvp-te (Routing Instances Provider Tunnel Selective)" on page 16061	"rsvp-te (Routing Instances Provider Tunnel)" on page 16059
"rsvp-te (Routing Instances Provider Tunnel)" on page 16059	"selective" on page 16529	

[edit routing-instances {routing-instance-name} provider-tunnel family {inet | inet6}]		
"mdt" on page 10590	"pim-ssm (Provider Tunnel)" on page 12519	"source (Routing Instances)" on page 28312

[edit routing-instances {routing-instance-name} provider-tunnel family {inet | inet6} mdt]		
"data-mdt-reuse" on page 4840	"tunnel-limit (Routing Instances)" on page 30213	

[edit routing-instances {routing-instance-name} provider-tunnel family {inet | inet6} pim-ssm]	
"group-address (Routing Instances Tunnel Group)" on page 7802	

[edit routing-instances {routing-instance-name} provider-tunnel family {inet | inet6}mdt threshold group {group-address} source {source-address}]	
"rate (Routing Instances)" on page 14073	

[edit routing-instances {routing-instance-name} provider-tunnel family {inet | inet6}mdt]		
"group-range (Data MDTs)" on page 7882	"threshold (Routing Instances)" on page 29338	

[edit routing-instances {routing-instance-name} provider-tunnel family {inet | inet6}mdtthreshold]	
"group (Protocols PIM)" on page 7876	

[edit routing-instances {routing-instance-name} provider-tunnel family {inet | inet6}pim-ssm]	
"tunnel-source" on page 30239	

[edit routing-instances {routing-instance-name} provider-tunnel ingress-replication]		
"create-new-ucast-tunnel" on page 4707	"transport-class" on page 30132	"unicast-tunnel-name-regular-expression" on page 30355

[edit routing-instances {routing-instance-name} provider-tunnel ingress-replication label-switched-path]	
"label-switched-path-template (Multicast)" on page 9596	

[edit routing-instances {instance-name} provider-tunnel inter-as]	
"ldp-p2mp" on page 9685	

[edit routing-instances {name} provider-tunnel ldp-p2mp]	
"root-address" on page 15840	

[edit routing-instances {routing-instance-name} provider-tunnel pim-asm]	
"group-address (Routing Instances VPN)" on page 7804	

[edit routing-instances {routing-instance-name} provider-tunnel pim-asm family inet]		
"group-address (Routing Instances VPN)" on page 7804		

[edit routing-instances {routing-instance-name} provider-tunnel pim-asm family inet6]		
"group-address (Routing Instances VPN)" on page 7804		

[edit routing-instances {routing-instance-name} provider-tunnel pim-ssm]		
"group-address (Routing Instances VPN)" on page 7804		

[edit routing-instances {routing-instance-name} provider-tunnel pim-ssm family inet]		
"group-address (Routing Instances VPN)" on page 7804		

[edit routing-instances {routing-instance-name} provider-tunnel pim-ssm family inet6]		
"group-address (Routing Instances VPN)" on page 7804		

[edit routing-instances {routing-instance-name} provider-tunnel rsvp-te]		
"label-switched-path-template (Multicast)" on page 9596	"static-lsp" on page 28547	

[edit routing-instances {instance-name} provider-tunnel selective]		
"cmcast-joins-limit-inet (MVPN Selective Tunnels)" on page 4376	"cmcast-joins-limit-inet6 (MVPN Selective Tunnels)" on page 4379	"group (Routing Instances)" on page 7889

(Continued)

[edit routing-instances {instance-name} provider-tunnel selective]		
"leaf-tunnel-limit-inet (MVPN Selective Tunnels)" on page 9691	"leaf-tunnel-limit-inet6 (MVPN Selective Tunnels)" on page 9693	"tunnel-limit (Routing Instances Provider Tunnel Selective)" on page 30215
"wildcard-group-inet" on page 31115	"wildcard-group-inet6" on page 31117	

[edit routing-instances {routing-instance-name} provider-tunnel selective group {address}]		
"source (Routing Instances Provider Tunnel Selective)" on page 28314	"wildcard-source (Selective Provider Tunnels) " on page 31120	

[edit routing-instances {routing-instance-name} provider-tunnel selective group {address} source {source-address}]		
"ingress-replication" on page 8647	"inter-region-segmented" on page 8791	"ldp-p2mp" on page 9685
"pim-ssm (Selective Tunnel)" on page 12520	"rsvp-te (Routing Instances Provider Tunnel Selective)" on page 16061	"threshold-rate" on page 29336

[edit routing-instances {routing-instance-name} provider-tunnel selective group {address} source {source-address} ingress-replication]	
"create-new-ucast-tunnel" on page 4707	

[edit routing-instances {routing-instance-name} provider-tunnel selective group {group-address} source {source-address} pim-ssm]	
"group-range (MBGP MVPN Tunnel)" on page 7883	

[edit routing-instances {routing-instance-name} provider-tunnel selective group {address} source {source-address} rsvp-te]		
"label-switched-path-template (Multicast)" on page 9596	"static-lsp" on page 28547	

[edit routing-instances {instance-name} provider-tunnel selective group {multicast IP address} wildcard-source]		
"inter-region-segmented" on page 8791	"ldp-p2mp" on page 9685	"pim-ssm (Selective Tunnel)" on page 12520
"rsvp-te (Routing Instances Provider Tunnel Selective)" on page 16061	"threshold-rate" on page 29336	

[edit routing-instances {routing-instance-name} provider-tunnel selective group {group-address} wildcard-source pim-ssm]	
"group-range (MBGP MVPN Tunnel)" on page 7883	

[edit routing-instances {routing-instance-name} provider-tunnel selective group {address} wildcard-source rsvp-te]	
"static-lsp" on page 28547	

[edit routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet]	
"wildcard-source (Selective Provider Tunnels) " on page 31120	

[edit routing-instances {instance-name} provider-tunnel selective wildcard-group-inet wildcard-source]		
"inter-region-segmented" on page 8791	"ldp-p2mp" on page 9685	"pim-ssm (Selective Tunnel)" on page 12520
"rsvp-te (Routing Instances Provider Tunnel Selective)" on page 16061	"threshold-rate" on page 29336	

[edit routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet wildcard-source pim-ssm]	
"group-range (MBGP MVPN Tunnel)" on page 7883	

[edit routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet wildcard-source rsvp-te]	
"static-lsp" on page 28547	

[edit routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet6]	
"wildcard-source (Selective Provider Tunnels) " on page 31120	

[edit routing-instances {instance-name} provider-tunnel selective wildcard-group-inet6 wildcard-source]		
"inter-region-segmented" on page 8791	"ldp-p2mp" on page 9685	"pim-ssm (Selective Tunnel)" on page 12520
"rsvp-te (Routing Instances Provider Tunnel Selective)" on page 16061	"threshold-rate" on page 29336	

[edit routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet6 wildcard-source pim-ssm]	
"group-range (MBGP MVPN Tunnel)" on page 7883	

[edit routing-instances {routing-instance-name} provider-tunnel selective wildcard-group-inet6 wildcard-source rsvp-te]	
"static-lsp" on page 28547	

[realm]

[edit routing-instances {routing-instance-name} realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area {area-id}]	
"area-range" on page 2453	

[routing-instance-name]

[edit routing-instances routing-instance-name protocols vpls associate-profile]	
"profile-variable-set (Routing Instances)" on page 13356	

[routing-options]

[edit routing-instances {routing-instance-name} routing-options]		
"access (Static Framed Routes)" on page 1703	"access-internal (Static Access-Internal Routes)" on page 1668	"aggregate (Routing)" on page 2070
"auto-bandwidth" on page 2682	"auto-export" on page 2712	"autonomous-system" on page 2764
"backup-selection (Protocols IS-IS)" on page 2814	"backup-selection (Protocols OSPF)" on page 2819	"dynamic-tunnels" on page 6161
"fate-sharing" on page 6952	"flex-algorithm" on page 7221	"flex-algorithm (Protocols OSPF)" on page 7223

(Continued)

[edit routing-instances {routing-instance-name} routing-options]		
"flow" on page 7252	"graceful-restart (Enabling Globally)" on page 7760	"instance-export" on page 8759
"instance-import" on page 8760	"interface (Host Fast Reroute)" on page 8897	"interface-routes" on page 9010
"label" on page 9574	"localized-fib" on page 9983	"martians" on page 10356
"maximum-paths" on page 10542	"maximum-prefixes" on page 10554	"multicast (routing-options)" on page 11063
"multipath (Routing Options)" on page 11113	"options (Routing Options)" on page 11922	"protect core" on page 13407
"resolution" on page 15594	"rib (General)" on page 15764	"rib (Multitopology Routing)" on page 15785
"route-distinguisher-id-use-router-id" on page 15889	"router-id" on page 15902	"routing-options-max-ecmp" on page 15985
"rpm-tracking" on page 16038	"source-packet-routing (routing-options)" on page 28264	"static (Routing Options)" on page 28570
"topologies (Multitopology Routing)" on page 29436	"traceoptions (routing-options)" on page 29462	"validation (Origin Validation for BGP)" on page 30647

[edit routing-instances {routing-instance-name} routing-options (aggregate | generate | static) (defaults | route)]		
"active" on page 1800	"as-path (Routing Options)" on page 2495	"color" on page 4425
"community (Routing Options)" on page 4472	"preference (Routing Options)" on page 13046	"tag (Routing Options)" on page 28924

[edit routing-instances {routing-instance-name} routing-options (aggregate | generate) (defaults | route)]		
"brief" on page 3039	"discard (routing-options)" on page 5816	"policy (Aggregate and Generated Routes)" on page 12670

[edit routing-instances {routing-instance-name} routing-options autonomous-system {autonomous-system}]	
"independent-domain" on page 8543	

[edit routing-instances {instance-name} routing-options backup-selection]		
"destination (Protocols OSPF)" on page 5511	"interface (Backup Selection OSPF)" on page 8837	

[edit routing-instances {instance-name} routing-options backup-selection destination {prefix} interface {interface-name}]	
"admin-group (Backup Selection)" on page 1965	

[edit routing-instances {name} routing-options bmp route-monitoring]	
"loc-rib" on page 9897	

[edit routing-instances {routing-instance-name} routing-options dynamic-tunnels {tunnel-name}]		
"destination-networks" on page 5451	"forwarding-rib" on page 7483	"source-address (Dynamic Tunnel)" on page 28117
"source-address (Routing Options)" on page 28156	"tunnel-attributes" on page 30194	"v4ov6" on page 30636

[edit routing-instances {routing-instance-name} routing-options dynamic-tunnels {tunnel-name} rsvp-te {entry}]		
"destination-networks" on page 5451		

[edit routing-instances {name} routing-options dynamic-tunnels tunnel-attributes]		
"dynamic-tunnel-reassembly" on page 6159		

[edit routing-instances {name} routing-options fate-sharing group {group-name}]		
"use-for-post-convergence-lfa" on page 30517		

[edit routing-instances {name} routing-options flex-algorithm]		
"definition" on page 5231	"definition (Protocols OSPF)" on page 5234	"use-transport-class" on page 30541

[edit routing-instances {name} routing-options flow]		
"discard-action-for-unresolved-redir-addr" on page 5818	"fast-lookup-filter (Protocols BGP)" on page 6947	"firewall-install-disable" on page 7194

[edit routing-instances {routing-instance-name} routing-options forwarding-table]		
"export (Routing Options)" on page 6691	"ip-tunnel-rpf-check" on page 9281	

[edit routing-instances {routing-instance-name} routing-options graceful-restart]		
"disable (Protocols LDP)" on page 5789	"disable (Routing Graceful Restart)" on page 5738	"disable (Routing Options)" on page 5804

(Continued)

[edit routing-instances {routing-instance-name} routing-options graceful-restart]		
"restart-duration (Routing Options)" on page 15655		

[edit routing-instances {instance-name} routing-options interface {interface-name}]		
"arp-prefix-limit (Host Fast Reroute)" on page 2477	"link-protection (Host Fast Reroute)" on page 9808	"supplementary-blackout-timer (Host Fast Reroute)" on page 28724

[edit routing-instances {routing-instance-name} routing-options interface-routes]		
"rib-group (Routing Options)" on page 15778		

[edit routing-instances {routing-instance-name} routing-options multicast]		
"asm-override-ssm" on page 2497	"backup-pe-group" on page 2808	"cont-stats-collection-interval" on page 4573
"flow-map" on page 7322	"interface (Multicast Static Routes)" on page 8928	"interface (Routing Options)" on page 9014
"pim-to-mld-proxy" on page 12523	"rpf-check-policy (Routing Options RPF)" on page 16012	"scope" on page 16379
"ssm-groups" on page 28452	"ssm-map (Routing Options Multicast)" on page 28468	"stream-protection (Multicast-Only Fast Reroute)" on page 28634
"traceoptions (routing-options)" on page 29462		

[edit routing-instances {routing-instance-name} routing-options multicast backup-pe-group {group-name}]		
"backups" on page 2826	"local-address (Routing Options)" on page 9913	

[edit routing-instances {routing-instance-name} routing-options multicast flow-map]		
"bandwidth (multicast)" on page 2831	"forwarding-cache (Flow Maps)" on page 7408	"policy (Flow Maps)" on page 12683
"redundant-sources" on page 14219	"timeout (Flow Maps)" on page 29377	

[edit routing-instances {routing-instance-name} routing-options multicast forwarding-cache]		
"allow-maximum (Multicast)" on page 2193	"threshold (Multicast Forwarding Cache)" on page 29329	"timeout (Multicast)" on page 29380

[edit routing-instances {routing-instance-name} routing-options multicast forwarding-cache (inet | inet6)]	
"threshold (Multicast Forwarding Cache)" on page 29329	

[edit routing-instances {routing-instance-name} routing-options multicast forwarding-cache family (inet | inet6) threshold]	
"log-warning (Multicast Forwarding Cache)" on page 10047	

[edit routing-instances {routing-instance-name} routing-options multicast forwarding-cache threshold]	
"log-warning (Multicast Forwarding Cache)" on page 10047	

[edit routing-instances {routing-instance-name} routing-options multicast interface {interface-name}]		
"maximum-bandwidth" on page 10502	"no-qos-adjust" on page 11554	"reverse-oif-mapping" on page 15720
"subscriber-leave-timer" on page 28694		

[edit routing-instances {routing-instance-name} routing-options multicast interface {interface-name} reverse-oif-mapping]	
"no-qos-adjust" on page 11554	

[edit routing-instances {routing-instance-name} routing-options multicast pim-to-igmp-proxy]	
"upstream-interface" on page 30474	

[edit routing-instances {routing-instance-name} routing-options multicast pim-to-mld-proxy]	
"upstream-interface" on page 30474	

[edit routing-instances {routing-instance-name} routing-options multicast scope {scope-name}]		
"interface (Multicast Scoping)" on page 8927	"interface (Scoping)" on page 9015	"prefix" on page 13062

[edit routing-instances {routing-instance-name} routing-options multicast ssm-map {ssm-map-name}]		
"policy (SSM Maps)" on page 12733	"source (Source-Specific Multicast)" on page 28327	

[edit routing-instances {routing-instance-name} routing-options multicast stream-protection]		
"mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain)" on page 10827	"mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain)" on page 10829	"mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain)" on page 10830
"mofrr-primary-path-selection-by-routing (Multicast-Only Fast Reroute)" on page 10832	"policy (Multicast-Only Fast Reroute)" on page 12691	

[edit routing-instances {routing-instance-name} routing-options multipath]		
"as-path-compare" on page 2490	"vpn-unequal-cost" on page 30951	

[edit routing-instances {routing-instance-name} routing-options outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[edit routing-instances {routing-instance-name} routing-options resolution]	
"rib (Route Resolution)" on page 15789	

[edit routing-instances {routing-instance-name} routing-options resolution rib]		
"import (Routing Options)" on page 8465	"resolution-ribs" on page 15599	

[edit routing-instances {routing-instance-name} routing-options rib {routing-table-name}]		
"aggregate (Routing)" on page 2070	"filter (routing-options)" on page 7122	"martians" on page 10356
"multipath (Routing Options)" on page 11113	"policy-multipath" on page 12694	"vpn-global-import" on page 30953

[edit routing-instances {routing-instance-name} routing-options rib {routing-table-name} (aggregate | generate | static) (defaults | route)]		
"active" on page 1800	"as-path (Routing Options)" on page 2495	"color" on page 4425
"community (Routing Options)" on page 4472	"preference (Routing Options)" on page 13046	"tag (Routing Options)" on page 28924

[edit routing-instances {routing-instance-name} routing-options rib {routing-table-name} (aggregate | generate) (defaults | route)]		
"brief" on page 3039	"discard (routing-options)" on page 5816	"policy (Aggregate and Generated Routes)" on page 12670

[edit routing-instances {routing-instance-name} routing-options rib inet6.0]	
"flow (IPv6)" on page 7298	

[edit routing-instances {routing-instance-name} routing-options rib {routing-table-name} multipath]	
"vpn-unequal-cost" on page 30951	

[edit routing-instances {routing-instance-name} routing-options rib {routing-table-name} static (defaults | route)]		
"install (Routing Options)" on page 8747	"readvertise" on page 14090	"resolve" on page 15600
"retain" on page 15672		

[edit routing-instances {routing-instance-name} routing-options rib {routing-table-name} static route {destination-prefix}]	
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	

[edit routing-instances {routing-instance-name} routing-options rib {routing-table-name} static route {destination-prefix} qualified-next-hop ({interface-name} | {address})]		
"bfd-liveness-detection (Routing Options Static Route)" on page 2950		

[edit routing-instances {routing-instance-name} routing-options rib-groups {group-name}]		
"export-rib" on page 6688	"import-policy" on page 8473	"import-rib" on page 8488

[edit routing-instances {name} routing-options source-packet-routing]		
"mapping-server-entry" on page 10346	"srv6 (Routing Options)" on page 28415	

[edit routing-instances {name} routing-options source-packet-routing mapping-server-entry]		
"prefix-segment (Routing Options)" on page 13109	"prefix-segment-range" on page 13107	

[edit routing-instances {name} routing-options source-packet-routing srv6]		
"block" on page 2980	"locator" on page 9997	

[edit routing-instances {name} routing-options source-packet-routing srv6 locator]	
"micro-sid" on page 10697	

[edit routing-instances {routing-instance-name} routing-options static (defaults | route)]		
"install (Routing Options)" on page 8747	"readvertise" on page 14090	"resolve" on page 15600
"retain" on page 15672		

[edit routing-instances {routing-instance-name} routing-options static route {destination-prefix}]		
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	"lsp-next-hop (Static Routes)" on page 10192	"qualified-next-hop (Static Routes)" on page 13910

[edit routing-instances {routing-instance-name} routing-options static route {destination-prefix} next-hop {address} resolve]	
"longest-match (Static Routes)" on page 10109	

[edit routing-instances {routing-instance-name} routing-options static route {destination-prefix} qualified-next-hop ({interface-name} | {address})]	
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	

[edit routing-instances {routing-instance-name} routing-options static route {destination-prefix}]. [edit routing-options static route {destination-prefix}]		
"p2mp-ldp-next-hop" on page 12119	"p2mp-lsp-next-hop" on page 12120	

[edit routing-instances {routing-instance-name} routing-options topologies family (inet | inet6)]	
"topology (Multitopology Routing)" on page 29443	

[edit routing-instances {name} routing-options transport-class]	
"name (Transport-class)" on page 11181	

[edit routing-instances {instance-name} routing-options validation]		
"group (Origin Validation for BGP)" on page 7850	"static (Origin Validation for BGP)" on page 28553	"traceoptions (Origin Validation for BGP)" on page 29626

[edit routing-instances {instance-name} routing-options validation group {group-name}]	
"session (Origin Validation for BGP)" on page 16836	

[edit routing-instances {instance-name} routing-options validation group {group-name} session {server-ip-address}]		
"port (Origin Validation for BGP)" on page 12846	"traceoptions (Origin Validation for BGP)" on page 29626	

[edit routing-instances {instance-name} routing-options validation group {group-name} session {address} traceoptions]		
"file (Tracing for Origin AS Validation)" on page 7091	"flag (Tracing for Origin AS Validation)" on page 7212	

[edit routing-instances {instance-name} routing-options validation static]	
"record (Origin Validation for BGP)" on page 14138	

[edit routing-instances {instance-name} routing-options validation static record {destination}]	
"maximum-length (Origin Validation for BGP)" on page 10531	

[edit routing-instances {instance-name} routing-options validation static record {destination} maximum-length {prefix-length}]	
"origin-autonomous-system (Origin Validation for BGP)" on page 11932	

[edit routing-instances {instance-name} routing-options validation traceoptions]		
"file (Tracing for Origin AS Validation)" on page 7091	"flag (Tracing for Origin AS Validation)" on page 7212	

[service-groups]

[edit routing-instances {routing-instance-name} service-groups {service-group-name}]		
"isid" on page 9399	"map-dest-bmac-to-dest-cmac" on page 10326	"pbb-service-options" on page 12283

[edit routing-instances {routing-instance-name} service-groups {service-group-name} pbb-service-options]	
"default-isid" on page 5191	

[statement-name]

[edit routing-instances {routing-instance-name} statement-name protocols igmp amt]	
"relay (IGMP)" on page 14292	

[edit routing-instances {routing-instance-name} statement-name protocols igmp amt relay]	
"defaults" on page 5224	

[switch-options]

[edit routing-instances {routing-instance-name} switch-options]		
"interface-mac-limit" on page 8910	"mac-statistics" on page 10263	"mac-table-size" on page 10271
"network-isolation-profile" on page 11314	"ovsdb-managed" on page 12088	

[edit routing-instances {routing-instance-name} switch-options interface {interface-name}]		
"interface-mac-limit" on page 8910	"mac-pinning" on page 10254	

[edit routing-instances {routing-instance-name} switch-options interface {interface-name} interface-mac-limit {limit}]	
"packet-action" on page 12129	

[edit routing-instances {routing-instance-name} switch-options interface-mac-limit {limit}]	
"packet-action" on page 12129	

[system]

[edit routing-instances {routing-instances-name} system processes static-subscribers]	
"traceoptions (Static Subscribers)" on page 29873	

[edit routing-instances {routing-instance-name} system services]		
"dhcp-local-server" on page 5601	"remote-device-management" on page 14317	"static-subscribers (Dynamic Service Provisioning)" on page 28582
"tcp-forwarding (Remote Device Management)" on page 28987		

[edit routing-instances {routing-instance-name} system services ...]		
"allow-active-leasequery (DHCP Local Server)" on page 2163	"allow-bulk-leasequery (DHCP Local Server)" on page 2165	"allow-leasequery (DHCP Local Server)" on page 2191

[edit routing-instances {routing-instance-name} system services dhcp-local-server]		
"dhcpv6 (DHCP Local Server)" on page 5658	"dual-stack-group (DHCP Local Server)" on page 6052	"duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)" on page 6066
"forward-snooped-clients (DHCP Local Server)" on page 7404	"interface-traceoptions (System Services DHCP)" on page 9087	"pool-match-order" on page 12777
"reauthenticate (DHCP Local Server)" on page 14103	"reconfigure (DHCP Local Server)" on page 14133	

[edit routing-instances {routing-instance-name} system services dhcp-local-server (dhcpv6) ...overrides]	
"include-option-82 (DHCP Local Server)" on page 8534	

[edit routing-instances {routing-instance-name} system services dhcp-local-server ...]		
"authentication (DHCP Local Server)" on page 2593	"dual-stack (DHCP Local Server Overrides)" on page 6049	"dynamic-profile (DHCP Local Server)" on page 6117
"group (DHCP Local Server)" on page 7814	"interface (DHCP Local Server)" on page 8860	"interface-description (DHCP Local Server)" on page 8854
"interface-name (DHCP Local Server)" on page 8937	"logical-system-name (DHCP Local Server)" on page 10082	"overrides (DHCP Local Server)" on page 12065
"requested-ip-network-match (DHCP Local Server)" on page 15572	"service-profile (DHCP Local Server)" on page 16729	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982
"username-include (DHCP Local Server)" on page 30595		

[edit routing-instances {routing-instance-name} system services dhcp-local-server ...overrides]	
"client-discover-match (DHCP Local Server)" on page 4308	

[edit routing-instances {routing-instance-name} system services dhcp-local-server authentication]	
"password (DHCP Local Server)" on page 12226	

[edit routing-instances {routing-instance-name} system services dhcp-local-server authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"delimiter (DHCP Local Server)" on page 5288	"domain-name (DHCP Local Server)" on page 5884
"mac-address (DHCP Local Server)" on page 10221	"option-60 (DHCP Local Server)" on page 11876	"option-82 (DHCP Local Server Authentication)" on page 11890

(Continued)

[edit routing-instances {routing-instance-name} system services dhcp-local-server authentication username-include]		
"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578	

[edit routing-instances {name} system services dhcp-local-server dhcpv6]		
"dynamic-server" on page 6154	"reauthenticate (DHCP Local Server)" on page 14103	"reconfigure (DHCP Local Server)" on page 14133

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 ...]	
"rapid-commit (DHCPv6 Local Server)" on page 14055	

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 authentication]	
"password (DHCP Local Server)" on page 12226	

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"client-id (DHCP Local Server)" on page 4315	"delimiter (DHCP Local Server)" on page 5288
"domain-name (DHCP Local Server)" on page 5884	"relay-agent-interface-id (DHCP Local Server)" on page 14275	"relay-agent-remote-id (DHCP Local Server)" on page 14282
"relay-agent-subscriber-id (DHCP Local Server)" on page 14288	"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name}]	
"reconfigure (DHCP Local Server)" on page 14133	

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} authentication]	
"password (DHCP Local Server)" on page 12226	

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"client-id (DHCP Local Server)" on page 4315	"delimiter (DHCP Local Server)" on page 5288
"domain-name (DHCP Local Server)" on page 5884	"relay-agent-interface-id (DHCP Local Server)" on page 14275	"relay-agent-remote-id (DHCP Local Server)" on page 14282
"relay-agent-subscriber-id (DHCP Local Server)" on page 14288	"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name}]	
"trace (DHCP Local Server)" on page 29456	

[edit routing-instances {name} system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name} overrides]		
"delay-advertise (DHCPv6)" on page 5239	"process-inform" on page 13240	

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit routing-instances {name} system services dhcp-local-server dhcpv6 group {group-name} overrides]		
"delay-advertise (DHCPv6)" on page 5239	"interface-client-limit (DHCP Local Server)" on page 9106	"process-inform" on page 13240

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"strict (DHCP Local Server)" on page 28646
"timeout (DHCP Local Server)" on page 29374	"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 group {group-name} reconfigure trigger]	
"radius-disconnect (DHCP Local Server)" on page 14003	

[edit routing-instances {name} system services dhcp-local-server dhcpv6 overrides]		
"delay-advertise (DHCPv6)" on page 5239	"interface-client-limit (DHCP Local Server)" on page 9106	"process-inform" on page 13240

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"strict (DHCP Local Server)" on page 28646
"timeout (DHCP Local Server)" on page 29374	"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160

[edit routing-instances {routing-instance-name} system services dhcp-local-server dhcpv6 reconfigure trigger]	
"radius-disconnect (DHCP Local Server)" on page 14003	

[edit routing-instances {name} system services dhcp-local-server dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Local Server)" on page 3375	"classification-key (DHCP Relay Agent)" on page 3376	"dual-stack-interface-client-limit (DHCP Local Server and Relay Agent)" on page 6058
"on-demand-address-allocation" on page 11813	"protocol-master" on page 13448	"reauthenticate (DHCP Local Server)" on page 14103

[edit routing-instances {routing-instance-name} system services dhcp-local-server dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Local Server)" on page 2056	"use-primary (DHCP Local Server)" on page 30535	

[edit routing-instances {name} system services dhcp-local-server group]		
"interface-tag (DHCP Local Server)" on page 9081	"reconfigure (DHCP Local Server)" on page 14133	

[edit routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} authentication]	
"password (DHCP Local Server)" on page 12226	

[edit routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"delimiter (DHCP Local Server)" on page 5288	"domain-name (DHCP Local Server)" on page 5884
"mac-address (DHCP Local Server)" on page 10221	"option-60 (DHCP Local Server)" on page 11876	"option-82 (DHCP Local Server Authentication)" on page 11890
"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578	

[edit routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Local Server)" on page 2056	"use-primary (DHCP Local Server)" on page 30535	

[edit routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} interface {interface-name}]	
"trace (DHCP Local Server)" on page 29456	

[edit routing-instances {name} system services dhcp-local-server group {group-name} interface {interface-name} overrides]		
"delay-offer (DHCPv4)" on page 5254	"process-inform" on page 13240	

[edit routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} interface {interface-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit routing-instances {name} system services dhcp-local-server group {group-name} overrides]		
"delay-offer (DHCPv4)" on page 5254	"interface-client-limit (DHCP Local Server)" on page 9106	"process-inform" on page 13240

[edit routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"timeout (DHCP Local Server)" on page 29374
"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160	

[edit routing-instances {routing-instance-name} system services dhcp-local-server group {group-name} reconfigure trigger]		
"radius-disconnect (DHCP Local Server)" on page 14003		

[edit routing-instances {name} system services dhcp-local-server overrides]		
"delay-offer (DHCPv4)" on page 5254	"interface-client-limit (DHCP Local Server)" on page 9106	"process-inform" on page 13240

[edit routing-instances {routing-instance-name} system services dhcp-local-server overrides process-inform]		
"pool (DHCP Local Server Overrides)" on page 12774		

[edit routing-instances {routing-instance-name} system services dhcp-local-server pool-match-order]		
"external-authority" on page 6726	"ip-address-first" on page 9235	"option-82 (DHCP Local Server Pool Matching)" on page 11892

[edit routing-instances {routing-instance-name} system services dhcp-local-server reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"timeout (DHCP Local Server)" on page 29374
"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160	

[edit routing-instances {routing-instance-name} system services dhcp-local-server reconfigure trigger]		
"radius-disconnect (DHCP Local Server)" on page 14003		

[edit routing-instances {routing-instance-name} system services remote-device-management]		
"service-device (Remote Device Management)" on page 16690		

[edit routing-instances {routing-instance-name} system services remote-device-management service-device {device-name}]		
"access-domain (Remote Device Management)" on page 1657	"provisioning-method (Remote Device Management)" on page 13814	

[edit routing-instances {routing-instance-name} system services remote-device-management service-device {device-name} provisioning-method]		
"netconf (Remote Device Management)" on page 11295		

[edit routing-instances {routing-instances-name} system services static-subscribers]		
"access-profile (Static Subscribers)" on page 1698	"authentication (Static Subscribers)" on page 2665	"group (Static Subscribers)" on page 7897

[edit routing-instances {routing-instances-name} system services static-subscribers authentication]		
"password (Static Subscribers)" on page 12247	"username-include (Static Subscribers)" on page 30604	

[edit routing-instances {routing-instances-name} system services static-subscribers authentication username-include]		
"domain-name (Static Subscribers)" on page 5892	"interface (Static Subscriber Username)" on page 9070	"logical-system-name (Static Subscribers)" on page 10085
"routing-instance-name (Static Subscribers)" on page 15949	"user-prefix (Static Subscribers)" on page 30582	

[edit routing-instances {routing-instances-name} system services static-subscribers dynamic-profile {profile-name}]		
"aggregate-clients (Static Subscribers)" on page 2060		

[edit routing-instances {routing-instances-name} system services static-subscribers group {group-name}]		
"access-profile (Static Subscribers)" on page 1698	"authentication (Static Subscribers)" on page 2665	"interface (Static Subscriber Group)" on page 9068

[edit routing-instances {routing-instances-name} system services static-subscribers group {group-name} authentication]		
"username-include (Static Subscribers)" on page 30604		

[edit routing-instances {routing-instances-name} system services static-subscribers group {group-name} authentication username-include]		
"domain-name (Static Subscribers)" on page 5892	"interface (Static Subscriber Username)" on page 9070	"logical-system-name (Static Subscribers)" on page 10085
"password (Static Subscribers)" on page 12247	"routing-instance-name (Static Subscribers)" on page 15949	"user-prefix (Static Subscribers)" on page 30582

[edit routing-instances {routing-instances-name} system services static-subscribers group {group-name} dynamic-profile {profile-name}]		
"aggregate-clients (Static Subscribers)" on page 2060		

[edit routing-instances {routing-instance-name} system services system services dhcp-local-server dhcpv6 ...]		
"delegated-pool (DHCP Local Server)" on page 5271	"multi-address-embedded-option-response (DHCP Local Server)" on page 11039	

[vlans]

[edit routing-instances {routing-instance-name} vlans {vlan-name}]		
"decapsulate-inner-vlan" on page 5163		

[edit routing-instances {name} vlans {name} forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit routing-instances {name} vlans {name} forwarding-options dhcp-relay dhcpv6 group {name} relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit routing-instances {name} vlans {name} forwarding-options dhcp-relay dhcpv6 relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit routing-instances {name} vlans {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit routing-instances {name} vlans {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit routing-instances {name} vlans {name} forwarding-options dhcp-relay group {name} relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} vlans {name} forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} vlans {name} forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} vlans {name} forwarding-options dhcp-relay relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} vlans {name} forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} vlans {name} forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit routing-instances {name} vlans {name} switch-options]		
"drop-unknown-mac-ip" on page 5984	"interface-mac-ip-limit" on page 8908	"mac-ip-table-size" on page 10231
"network-isolation-profile" on page 11314		

[edit routing-instances {name} vlans {name} switch-options interface]	
"static-mac-ip" on page 28587	

[edit routing-instances {routing-instance-name} vlans {vlan-name} vxlan]		
"decapsulate-accept-inner-vlan" on page 5158	"encapsulate-inner-vlan" on page 6282	"ingress-node-replication" on page 8634
"ingress-node-replication (EVPN)" on page 8635	"ovsdb-managed" on page 12088	"static-remote-vtep-list" on page 28566
"translation-vni" on page 30080	"vni" on page 30909	

[vrf-advertise-selective]

[edit routing-instances {routing-instance-name} vrf-advertise-selective]	
"family (VRF Advertisement)" on page 6921	

[edit routing-instances {routing-instance-name} vrf-advertise-selective family]		
"inet-mvpn (VRF Advertisement)" on page 8575	"inet6-mvpn (VRF Advertisement)" on page 8612	

[vrf-table-label]

[edit routing-instances {routing-instance-name} vrf-table-label]	
"source-class-usage" on page 28191	

[vxlan]

[edit routing-instances {routing-instance-name} vxlan]		
"ingress-node-replication" on page 8634	"ovsdb-managed" on page 12088	"translation-vni" on page 30080

[edit routing-instance-rib-inet6.0]

IN THIS SECTION

- [flow] | 929

Sub Hierarchies

"[edit routing-instance-rib-inet6.0 flow]" on page 929

[flow]

[edit routing-instance-rib-inet6.0 flow]	
"per-route-accounting" on page 12420	

[edit routing-options]

IN THIS SECTION

- [(aggregate | generate | static)] | 934
- [(aggregate | generate)] | 934
- [access] | 934

- [access-internal] | **935**
- [autonomous-system] | **935**
- [backup-selection] | **936**
- [bmp] | **936**
- [client] | **937**
- [dynamic-tunnels] | **937**
- [fate-sharing] | **938**
- [flex-algorithm] | **939**
- [flow] | **939**
- [forwarding-table] | **939**
- [graceful-restart] | **941**
- [host-fast-reroute] | **941**
- [interface-routes] | **941**
- [multicast] | **942**
- [nonstop-routing-options] | **944**
- [outbound-route-filter] | **945**
- [ppm] | **945**
- [resolution] | **945**
- [rib] | **946**
- [rib-groups] | **948**
- [source-packet-routing] | **948**
- [srlg] | **949**
- [static] | **949**
- [topologies] | **950**
- [transport-class] | **950**
- [validation] | **951**

Sub Hierarchies

"[edit routing-options (aggregate | generate | static)]" on page 934

"[edit routing-options (aggregate | generate)]" on page 934

"[edit routing-options access]" on page 934

"[edit routing-options access-internal]" on page 935

"[edit routing-options autonomous-system {as-number}]" on page 935

"[edit routing-options backup-selection]" on page 936

"[edit routing-options bmp]" on page 936

"[edit routing-options client {id}]" on page 937

"[edit routing-options dynamic-tunnels {tunnel-name}]" on page 937

"[edit routing-options fate-sharing]" on page 938

"[edit routing-options flex-algorithm]" on page 939

"[edit routing-options flow]" on page 939

"[edit routing-options forwarding-table]" on page 939

"[edit routing-options graceful-restart]" on page 941

"[edit routing-options host-fast-reroute]" on page 941

"[edit routing-options interface-routes]" on page 941

"[edit routing-options multicast]" on page 942

"[edit routing-options nonstop-routing-options]" on page 944

"[edit routing-options outbound-route-filter]" on page 945

"[edit routing-options ppm]" on page 945

"[edit routing-options resolution]" on page 945

"[edit routing-options rib {routing-table-name}]" on page 946

"[edit routing-options rib-groups {group-name}]" on page 948

"[edit routing-options source-packet-routing]" on page 948

"[edit routing-options srlg]" on page 949

"[edit routing-options static]" on page 949

"[edit routing-options topologies]" on page 950

"[edit routing-options transport-class]" on page 950

"[edit routing-options validation]" on page 951

[edit routing-options]		
"access (Static Framed Routes)" on page 1703	"access-internal (Static Access-Internal Routes)" on page 1668	"admin-groups-extended" on page 1975
"admin-groups-extended-range" on page 1977	"aggregate (Routing)" on page 2070	"auto-bandwidth" on page 2682
"auto-export" on page 2712	"autonomous-system" on page 2764	"backup-selection (Protocols IS-IS)" on page 2814
"backup-selection (Protocols OSPF)" on page 2819	"bgp-static" on page 2966	"bmp" on page 3000
"confederation" on page 4508	"domain-path-id" on page 5894	"dynamic-tunnels" on page 6161
"fate-sharing" on page 6952	"flex-algorithm" on page 7221	"flex-algorithm (Protocols OSPF)" on page 7223
"flow" on page 7252	"forwarding-table" on page 7490	"generate" on page 7672
"graceful-restart (Enabling Globally)" on page 7760	"graceful-restart (Protocols RSVP)" on page 7782	"host-fast-reroute" on page 8141
"instance-export" on page 8759	"instance-import" on page 8760	"interface-routes" on page 9010
"krt-nexthop-ack" on page 9505	"logical-system-mux" on page 10079	"martians" on page 10356
"maximum-paths" on page 10542	"maximum-prefixes" on page 10554	"med-igp-update-interval" on page 10594
"multicast (routing-options)" on page 11063	"no-bfd-triggered-local-repair" on page 11414	"nonstop-routing" on page 11671

(Continued)

[edit routing-options]		
"nonstop-routing-options" on page 11673	"nsr-phantom-holdtime" on page 11715	"options (Routing Options)" on page 11922
"ppm" on page 12927	"programmable-rpd" on page 13369	"purge-timeout (programmable-rpd)" on page 13870
"resolution" on page 15594	"rib (General)" on page 15764	"rib (Multitopology Routing)" on page 15785
"rib-groups" on page 15781	"rib-service (programmable-rpd)" on page 15790	"route-distinguisher-id" on page 15860
"route-distinguisher-id-use-router-id" on page 15889	"route-record" on page 15874	"router-id" on page 15902
"routing-options-max-ecmp" on page 15985	"rpm-tracking" on page 16038	"session-id-change-limiter-indirect" on page 16813
"sla-tracking" on page 28027	"source-packet-routing (routing-options)" on page 28264	"source-routing" on page 28311
"srlg" on page 28411	"static (Routing Options)" on page 28570	"stats-aggregation-interval (programmable-rpd)" on page 28601
"stats-sampling-interval (programmable-rpd)" on page 28602	"topologies (Multitopology Routing)" on page 29436	"traceoptions (programmable-rpd)" on page 29651
"traceoptions (routing-options)" on page 29462	"transport-class" on page 30120	"validation (Origin Validation for BGP)" on page 30647
"warm-standby" on page 31023		

[(aggregate | generate | static)]

[edit routing-options (aggregate | generate | static) (defaults | route)]		
"active" on page 1800	"as-path (Routing Options)" on page 2495	"color" on page 4425
"community (Routing Options)" on page 4472	"metric (Aggregate, Generated, or Static Route)" on page 10667	"preference (Routing Options)" on page 13046
"tag (Routing Options)" on page 28924		

[(aggregate | generate)]

[edit routing-options (aggregate | generate) (defaults | route)]		
"brief" on page 3039	"discard (routing-options)" on page 5816	"policy (Aggregate and Generated Routes)" on page 12670

[access]

[edit routing-options access]	
"route (Access)" on page 15841	

[edit routing-options access route {ip-prefix</prefix-length>}]		
"metric" on page 10665	"next-hop (Access)" on page 11329	"preference (Access)" on page 13032

(Continued)

[edit routing-options access route {ip-prefix</prefix-length>}]		
"qualified-next-hop (Access)" on page 13906	"tag (Access)" on page 28905	

[access-internal]

[edit routing-options access-internal]	
"route (Access-Internal)" on page 15845	

[edit routing-options access-internal route {ip-prefix</prefix-length>}]		
"next-hop (Access Internal)" on page 11330	"qualified-next-hop (Access-Internal)" on page 13907	

[autonomous-system]

[edit routing-options autonomous-system {as-number}]	
"loops (Autonomous System)" on page 10132	

[backup-selection]

[edit routing-options backup-selection]	
"destination (Protocols OSPF)" on page 5511	

[edit routing-options backup-selection destination {prefix}]	
"interface (Backup Selection OSPF)" on page 8837	

[edit routing-options backup-selection destination {prefix} interface {interface name}]	
"admin-group (Backup Selection)" on page 1965	

[bmp]

[edit routing-options bmp]		
"authentication-algorithm" on page 2577	"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605
"hold-down" on page 8086	"route-monitoring" on page 15871	"traceoptions (Protocols BMP)" on page 29666

[edit routing-options bmp route-monitoring]		
"loc-rib" on page 9897	"post-policy" on page 12910	

[edit routing-options bmp station {station-name}]		
"authentication-algorithm" on page 2577	"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key-chain (Protocols BGP and BMP)" on page 2605
"hold-down" on page 8086	"route-monitoring" on page 15871	"traceoptions (Protocols BMP)" on page 29666
"hold-down" on page 8086	"route-monitoring" on page 15871	"traceoptions (Protocols BMP)" on page 29666

[edit routing-options bmp station {station-name} route-monitoring]	
"post-policy" on page 12910	

[client]

[edit routing-options client {id}]	
"interface-notification (programmable-rpd)" on page 8942	

[dynamic-tunnels]

[edit routing-options dynamic-tunnels {tunnel-name}]		
"destination-networks" on page 5451	"forwarding-rib" on page 7483	"gre (Routing Options)" on page 7797

(Continued)

[edit routing-options dynamic-tunnels {tunnel-name}]		
"rsvp-te (Routing Options)" on page 16062	"source-address (Dynamic Tunnel)" on page 28117	"source-address (Routing Options)" on page 28156
"spring-te (Dynamic Tunnels)" on page 28395	"tunnel-attributes" on page 30194	"v4ov6" on page 30636

[edit routing-options dynamic-tunnels gre next-hop-based-tunnel]	
"next-hop-based-tunnel (Routing Options)" on page 11332	

[edit routing-options dynamic-tunnels {tunnel-name} rsvp-te {entry}]		
"destination-networks" on page 5451	"label-switched-path-template (Multicast)" on page 9596	

[edit routing-options dynamic-tunnels tunnel-attributes]		
"dynamic-tunnel-gre-key" on page 6158	"dynamic-tunnel-reassembly" on page 6159	

[fate-sharing]

[edit routing-options fate-sharing group {group-name}]	
"use-for-post-convergence-lfa" on page 30517	

[flex-algorithm]

[edit routing-options flex-algorithm]		
"definition" on page 5231	"definition (Protocols OSPF)" on page 5234	"use-transport-class" on page 30541

[flow]

[edit routing-options flow]		
"discard-action-for-unresolved-redir-addr" on page 5818	"fast-lookup-filter (Protocols BGP)" on page 6947	"firewall-install-disable" on page 7194
"interface-group (Routing Options)" on page 8894	"per-route-accounting" on page 12420	"traceoptions (routing-options)" on page 29462

[forwarding-table]

[edit routing-options forwarding-table]		
"chained-composite-next-hop" on page 3277	"ecmp-fast-reroute" on page 6193	"egress-chaining" on page 6210
"evpn-aliasing-optimize" on page 6547	"export (Routing Options)" on page 6691	"fib-next-hop-split" on page 7006
"indirect-next-hop" on page 8549	"indirect-next-hop-change-acknowledgements" on page 8551	"krt-nexthop-ack-timeout" on page 9506

(Continued)

[edit routing-options forwarding-table]		
"unicast-reverse-path" on page 30348		

[edit routing-options forwarding-table chained-composite-next-hop]		
"ingress (Chained Composite Next Hop)" on page 8628	"transit (Chained Composite Next Hops)" on page 30066	

[edit routing-options forwarding-table chained-composite-next-hop ingress]		
"l2ckt" on page 9535	"l2vpn (routing-options)" on page 9564	"l3vpn (ingress)" on page 9571
"labeled-bgp" on page 9600	"no-l2ckt" on page 11491	"no-l2vpn" on page 11492

[edit routing-options forwarding-table chained-composite-next-hop ingress l3vpn]		
"extended-space" on page 6703		

[edit routing-options forwarding-table chained-composite-next-hop ingress labeled-bgp]		
"inet" on page 8554	"inet6" on page 8595	

[edit routing-options forwarding-table chained-composite-next-hop transit]		
"l3vpn (transit)" on page 9573		

[graceful-restart]

[edit routing-options graceful-restart]		
"disable (Routing Graceful Restart)" on page 5738	"disable (Routing Options)" on page 5804	"restart-duration (Protocols)" on page 15651
"restart-duration (Routing Options)" on page 15655		

[host-fast-reroute]

[edit routing-options host-fast-reroute]		
"global-arp-prefix-limit (Host Fast Reroute)" on page 7698	"global-supplementary-blackout-timer (Host Fast Reroute)" on page 7737	

[interface-routes]

[edit routing-options interface-routes]	
"rib-group (Routing Options)" on page 15778	

[multicast]

[edit routing-options multicast]		
"asm-override-ssm" on page 2497	"backup-pe-group" on page 2808	"cont-stats-collection-interval" on page 4573
"flow-map" on page 7322	"interface (Multicast Static Routes)" on page 8928	"interface (Routing Options)" on page 9014
"pim-to-igmp-proxy" on page 12522	"pim-to-mld-proxy" on page 12523	"rpf-check-policy (Routing Options RPF)" on page 16012
"rpf-check-policy (Routing Options)" on page 16011	"scope" on page 16379	"scope-policy" on page 16380
"ssm-groups" on page 28452	"ssm-map (Routing Options Multicast)" on page 28468	"stream-protection (Multicast-Only Fast Reroute)" on page 28634
"traceoptions (routing-options)" on page 29462		

[edit routing-options multicast backup-pe-group {group-name}]		
"backups" on page 2826	"local-address (Routing Options)" on page 9913	

[edit routing-options multicast flow-map]		
"bandwidth (multicast)" on page 2831	"forwarding-cache (Flow Maps)" on page 7408	"policy (Flow Maps)" on page 12683
"redundant-sources" on page 14219	"timeout (Flow Maps)" on page 29377	

[edit routing-options multicast forwarding-cache]		
"allow-maximum (Multicast)" on page 2193	"threshold (Multicast Forwarding Cache)" on page 29329	"timeout (Multicast)" on page 29380

[edit routing-options multicast forwarding-cache family (inet | inet6)]	
"threshold (Multicast Forwarding Cache)" on page 29329	

[edit routing-options multicast forwarding-cache family (inet | inet6) threshold]	
"log-warning (Multicast Forwarding Cache)" on page 10047	

[edit routing-options multicast forwarding-cache threshold]	
"log-warning (Multicast Forwarding Cache)" on page 10047	

[edit routing-options multicast interface {interface-name}]		
"maximum-bandwidth" on page 10502	"no-qos-adjust" on page 11554	"reverse-oif-mapping" on page 15720
"subscriber-leave-timer" on page 28694		

[edit routing-options multicast interface {interface-name} reverse-oif-mapping]	
"no-qos-adjust" on page 11554	

[edit routing-options multicast pim-to-igmp-proxy]	
"upstream-interface" on page 30474	

[edit routing-options multicast pim-to-mld-proxy]		
"upstream-interface" on page 30474		

[edit routing-options multicast scope {scope-name}]		
"interface (Multicast Scoping)" on page 8927	"interface (Scoping)" on page 9015	"prefix" on page 13062

[edit routing-options multicast ssm-map {ssm-map-name}]		
"policy (SSM Maps)" on page 12733	"source (Source-Specific Multicast)" on page 28327	

[edit routing-options multicast stream-protection]		
"mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain)" on page 10827	"mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain)" on page 10829	"mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain)" on page 10830
"mofrr-primary-path-selection-by-routing (Multicast-Only Fast Reroute)" on page 10832	"policy (Multicast-Only Fast Reroute)" on page 12691	

[nonstop-routing-options]

[edit routing-options nonstop-routing-options]	
"precision-timers-max-period" on page 13012	

[outbound-route-filter]

[edit routing-options outbound-route-filter]	
"bgp-orf-cisco-mode" on page 2963	

[ppm]

[edit routing-options ppm]	
"no-delegate-processing" on page 11435	

[resolution]

[edit routing-options resolution]		
"rib (Route Resolution)" on page 15789		

[edit routing-options resolution rib]		
"import (Routing Options)" on page 8465	"resolution-ribs" on page 15599	

[rib]

[edit routing-options rib {routing-table-name}]		
"aggregate (Routing)" on page 2070	"bgp-static" on page 2966	"filter (routing-options)" on page 7122
"generate" on page 7672	"martians" on page 10356	"policy-multipath" on page 12694
"static (Routing Options)" on page 28570	"vpn-global-import" on page 30953	

[edit routing-options rib {routing-table-name} (aggregate | generate | static) (defaults | route)]		
"active" on page 1800	"as-path (Routing Options)" on page 2495	"color" on page 4425
"community (Routing Options)" on page 4472	"preference (Routing Options)" on page 13046	"tag (Routing Options)" on page 28924

[edit routing-options rib {routing-table-name} (aggregate | generate) (defaults | route)]		
"brief" on page 3039	"discard (routing-options)" on page 5816	"policy (Aggregate and Generated Routes)" on page 12670

[edit routing-options rib bgp.rtarget.0 static]	
"route-target-filter" on page 15879	

[edit routing-options rib {routing-table-name} filter]	
"input (Routing Options RIB)" on page 8717	

[edit routing-options rib inet6.0]	
"flow (IPv6)" on page 7298	

[edit routing-options rib inet6.0 static route {destination-prefix}]	
"qualified-next-hop (Static Routes)" on page 13910	

[edit routing-options rib {routing-table-name} static]	
"rib-group (Routing Options)" on page 15778	

[edit routing-options rib {routing-table-name} static (defaults | route)]		
"install (Routing Options)" on page 8747	"readvertise" on page 14090	"resolve" on page 15600
"retain" on page 15672		

[edit routing-options rib {routing-table-name} static route {destination-prefix}]		
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	"spring-te-lsp-next-hop" on page 28397	

[edit routing-options rib {routing-table-name} static route {destination-prefix} qualified-next-hop ({interface-name} | {address})]	
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	

[rib-groups]

[edit routing-options rib-groups {group-name}]		
"export-rib" on page 6688	"import-policy" on page 8473	"import-rib" on page 8488

[source-packet-routing]

[edit routing-options source-packet-routing]		
"mapping-server-entry" on page 10346	"srv6 (Routing Options)" on page 28415	

[edit routing-options source-packet-routing mapping-server-entry]		
"prefix-segment (Routing Options)" on page 13109	"prefix-segment-range" on page 13107	

[edit routing-options source-packet-routing srv6]		
"block" on page 2980	"locator" on page 9997	

[edit routing-options source-packet-routing srv6 locator]	
"micro-sid" on page 10697	

[srlg]

[edit routing-options srlg]		
"srlg-cost" on page 28413	"srlg-value" on page 28414	

[static]

[edit routing-options static]		
"rib-group (Routing Options)" on page 15778		

[edit routing-options static (defaults | route)]		
"install (Routing Options)" on page 8747	"readvertise" on page 14090	"resolve" on page 15600
"retain" on page 15672		

[edit routing-options static route {destination-prefix}]		
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	"lsp-next-hop (Static Routes)" on page 10192	"qualified-next-hop (Static Routes)" on page 13910
"spring-te-lsp-next-hop" on page 28397		

[edit routing-options static route {destination-prefix} next-hop {address} resolve]	
"longest-match (Static Routes)" on page 10109	

[edit routing-options static route {destination-prefix} qualified-next-hop]	
"metric (Qualified Next Hop on Static Route)" on page 10689	

[edit routing-options static route {destination-prefix} qualified-next-hop ({interface-name} | {address})]	
"bfd-liveness-detection (Routing Options Static Route)" on page 2950	

[topologies]

[edit routing-options topologies family (inet | inet6)]	
"topology (Multitopology Routing)" on page 29443	

[transport-class]

[edit routing-options transport-class]	
"name (Transport-class)" on page 11181	

[validation]

[edit routing-options validation]		
"group (Origin Validation for BGP)" on page 7850	"static (Origin Validation for BGP)" on page 28553	"traceoptions (Origin Validation for BGP)" on page 29626

[edit routing-options validation group {group-name}]	
"session (Origin Validation for BGP)" on page 16836	

[edit routing-options validation group {group-name} session {server-ip-address}]		
"port (Origin Validation for BGP)" on page 12846	"traceoptions (Origin Validation for BGP)" on page 29626	

[edit routing-options validation group {group-name} session {address} traceoptions]		
"file (Tracing for Origin AS Validation)" on page 7091	"flag (Tracing for Origin AS Validation)" on page 7212	

[edit routing-options validation static]	
"record (Origin Validation for BGP)" on page 14138	

[edit routing-options validation static record {destination}]	
"maximum-length (Origin Validation for BGP)" on page 10531	

[edit routing-options validation static record {destination} maximum-length {prefix-length}]	
"origin-autonomous-system (Origin Validation for BGP)" on page 11932	

[edit routing-options validation traceoptions]		
"file (Tracing for Origin AS Validation)" on page 7091	"flag (Tracing for Origin AS Validation)" on page 7212	

[edit routing-options-rib-inet6.0]

IN THIS SECTION

- [flow] | **952**

Sub Hierarchies

"[edit routing-options-rib-inet6.0 flow]" on page 952

[flow]

[edit routing-options-rib-inet6.0 flow]	
"per-route-accounting" on page 12420	

[edit schedulers]

IN THIS SECTION

- [scheduler] | 953

Sub Hierarchies

"[edit schedulers scheduler {scheduler-name}]" on page 953

[edit schedulers]	
"scheduler (Security Policies)" on page 16364	

[scheduler]

[edit schedulers scheduler {scheduler-name}]		
"start-date" on page 28492	"stop-date" on page 28610	

[edit schedulers scheduler {scheduler-name} daily]		
"exclude (Schedulers)" on page 6626	"start-time (Schedulers)" on page 28503	"stop-time" on page 28612

[edit schedulers scheduler {scheduler-name} friday]		
"exclude (Schedulers)" on page 6626	"start-time (Schedulers)" on page 28503	"stop-time" on page 28612

[edit schedulers scheduler {scheduler-name} monday]		
"exclude (Schedulers)" on page 6626	"start-time (Schedulers)" on page 28503	"stop-time" on page 28612

[edit schedulers scheduler {scheduler-name} saturday]		
"exclude (Schedulers)" on page 6626	"start-time (Schedulers)" on page 28503	"stop-time" on page 28612

[edit schedulers scheduler {scheduler-name} sunday]		
"exclude (Schedulers)" on page 6626	"start-time (Schedulers)" on page 28503	"stop-time" on page 28612

[edit schedulers scheduler {scheduler-name} thursday]		
"exclude (Schedulers)" on page 6626	"start-time (Schedulers)" on page 28503	"stop-time" on page 28612

[edit schedulers scheduler {scheduler-name} tuesday]		
"exclude (Schedulers)" on page 6626	"start-time (Schedulers)" on page 28503	"stop-time" on page 28612

[edit schedulers scheduler {scheduler-name} wednesday]		
"exclude (Schedulers)" on page 6626	"start-time (Schedulers)" on page 28503	"stop-time" on page 28612

[edit security]

IN THIS SECTION

- [address-book] | 960
- [advance-policy-based-routing] | 960
- [advanced-policy-based-routing] | 961
- [alarms] | 961
- [alg] | 962
- [application-firewall] | 965
- [application-screen] | 966
- [application-tracking] | 966
- [authentication-key-chains] | 966
- [auto-re-enrollment] | 967
- [casb] | 967
- [certificates] | 967
- [cloud] | 968
- [datapath-debug] | 968
- [dynamic-address] | 969
- [dynamic-application] | 969
- [dynamic-vpn] | 970
- [firewall-authentication] | 970
- [flow] | 970
- [forwarding-options] | 973
- [forwarding-process] | 974
- [group-vpn] | 975
- [gtp] | 977
- [host-vpn] | 978
- [idp] | 979
- [ike] | 989
- [ipsec] | 991
- [log] | 994

- [logical-systems] | **996**
- [macsec] | **996**
- [nat] | **999**
- [pki] | **1006**
- [policies] | **1007**
- [remote-access] | **1013**
- [resource-manager] | **1013**
- [screen] | **1014**
- [sctp] | **1016**
- [sip] | **1017**
- [softwires] | **1017**
- [tcp-encap] | **1017**
- [trusted-channel] | **1017**
- [tunnel-inspection] | **1018**
- [user-identification] | **1018**
- [utm] | **1018**
- [zones] | **1044**

Sub Hierarchies

"[edit security address-book {book-name}]" on page 960

"[edit security advance-policy-based-routing]" on page 960

"[edit security advanced-policy-based-routing]" on page 961

"[edit security alarms]" on page 961

"[edit security alg]" on page 962

"[edit security application-firewall]" on page 965

"[edit security application-screen]" on page 966

"[edit security application-tracking]" on page 966

"[edit security authentication-key-chains]" on page 966

"[edit security auto-re-enrollment]" on page 967

"[edit security casb]" on page 967

"[edit security certificates]" on page 967

"[edit security cloud]" on page 968

"[edit security datapath-debug]" on page 968

"[edit security dynamic-address]" on page 969

"[edit security dynamic-application]" on page 969

"[edit security dynamic-vpn]" on page 970

"[edit security firewall-authentication]" on page 970

"[edit security flow]" on page 970

"[edit security forwarding-options]" on page 973

"[edit security forwarding-process]" on page 974

"[edit security group-vpn]" on page 975

"[edit security gtp]" on page 977

"[edit security host-vpn]" on page 978

"[edit security idp]" on page 979

"[edit security ike]" on page 989

"[edit security ipsec]" on page 991

"[edit security log]" on page 994

"[edit security logical-systems {logical-system-name}]" on page 996

"[edit security macsec]" on page 996

"[edit security nat]" on page 999

"[edit security pki]" on page 1006

"[edit security policies]" on page 1007

"[edit security remote-access]" on page 1013

"[edit security resource-manager]" on page 1013

"[edit security screen]" on page 1014

"[edit security sctp]" on page 1016

"[edit security sip]" on page 1017

"[edit security softwires]" on page 1017

"[edit security tcp-encap]" on page 1017

"[edit security trusted-channel]" on page 1017

"[edit security tunnel-inspection]" on page 1018

"[edit security user-identification]" on page 1018

"[edit security utm]" on page 1018

"[edit security zones]" on page 1044

[edit security]		
"address-book" on page 1868	"alarms (Security)" on page 2128	"alg" on page 2130
"application-firewall" on page 2318	"application-tracking" on page 2397	"authentication-key-chains" on page 2610
"authentication-key-chains (TCP-AO)" on page 2612	"casb" on page 3210	"casb-policy" on page 3212
"certificates" on page 3262	"Content Security" on page 4589	"datapath-debug" on page 4855
"distribution-profile" on page 5838	"dynamic-address" on page 6086	"dynamic-address" on page 6088
"dynamic-application (Security)" on page 6093	"dynamic-vpn" on page 6163	"firewall-authentication (Security)" on page 7184
"flow (Security Flow)" on page 7333	"forwarding-options (Security)" on page 7473	"forwarding-process" on page 7481
"gprs" on page 7750	"group-vpn" on page 7911	"gtp" on page 7931

(Continued)

[edit security]		
"host-vpn" on page 8174	"idp (Security)" on page 8321	"ike (High Availability)" on page 8424
"ike (Security Services)" on page 8446	"ike (Security)" on page 8432	"ipsec (High Availability)" on page 9301
"ipsec (Security Interfaces)" on page 9319	"ipsec (Security)" on page 9312	"ipsec-policy" on page 9306
"log (Security)" on page 10026	"macsec" on page 10278	"macsec (MX Series)" on page 10281
"nat" on page 11205	"ngfw" on page 11369	"pki" on page 12610
"pki (Services)" on page 12612	"policies" on page 12650	"remote-access (Juniper Secure Connect)" on page 14307
"resource-manager" on page 15614	"screen " on page 16384	"sctp" on page 16396
"softwires" on page 28102	"ssh-known-hosts" on page 28426	"tcp-encap" on page 28980
"traceoptions (Security)" on page 29767	"tunnel-inspection" on page 30208	"twamp (Security ALG)" on page 30256
"user-identification" on page 30558	"zones" on page 31172	

[address-book]

[edit security address-book {book-name}]		
"address (Security Address Book)" on page 1920	"address-set" on page 1930	"attach" on page 2527
"description (Security Address Book)" on page 5347		

[edit security address-book address {address-name}]	
"range-address" on page 14047	

[advance-policy-based-routing]

[edit security advance-policy-based-routing]		
"interface (advance-policy-based-routing)" on page 8833	"multipath-rule" on page 11115	

[edit security advance-policy-based-routing from-zone {zone-name} policy {policy-name} match]	
"source-identity" on page 28216	

[advanced-policy-based-routing]

[edit security advanced-policy-based-routing from-zone {name}]	
"policy (advanced-policy-based-routing)" on page 12667	

[alarms]

[edit security alarms]		
"audible (Security Alarms)" on page 2563	"potential-violation" on page 12912	

[edit security alarms potential-violation]		
"authentication (Security Alarms)" on page 2648	"decryption-failures" on page 5168	"encryption-failures" on page 6330
"idp (Security Alarms)" on page 8333	"ike-phase1-failures" on page 8428	"ike-phase2-failures" on page 8429
"policy (Security Alarms)" on page 12709	"replay-attacks" on page 14379	"security-log-percent-full" on page 16485

[edit security alarms potential-violation policy]		
"application (Security Alarms)" on page 2366	"destination-ip (Security Alarms)" on page 5437	"policy-match" on page 12689
"source-ip (Security Alarms)" on page 28242		

[alg]

[edit security alg]		
"alg-manager" on page 2141	"alg-support-lib" on page 2142	"dns" on page 5847
"ftp (Security ALG)" on page 7626	"h323" on page 7961	"ike-esp-nat" on page 8421
"mgcp" on page 10693	"msrpc" on page 10999	"pptp" on page 12986
"protect (Security)" on page 13406	"rsh" on page 16041	"rtsp" on page 16072
"sccp" on page 16341	"sip (Security) " on page 27977	"sql" on page 28402
"sunrpc" on page 28722	"talk" on page 28940	"tftp (Security ALG)" on page 29252
"traceoptions (Security ALG)" on page 29770		

[edit security alg dns]		
"maximum-message-length" on page 10534	"oversize-message-drop" on page 12081	

[edit security alg h323]		
"application-screen (Security H323)" on page 2359	"dscp-rewrite" on page 6019	"endpoint-registration-timeout" on page 6342
"media-source-port-any" on page 10599	"traceoptions (Security H323 ALG)" on page 29801	

[edit security alg h323 application-screen]		
"message-flood (Security H323)" on page 10644	"unknown-message (Security H323 ALG)" on page 30410	

[edit security alg h323 application-screen message-flood]	
"gatekeeper" on page 7653	

[edit security alg mgcp]		
"application-screen (Security MGCP)" on page 2361	"dscp-rewrite" on page 6019	"inactive-media-timeout (Security MGCP)" on page 8496
"maximum-call-duration (Security)" on page 10510	"traceoptions (Security MGCP ALG)" on page 29814	"transaction-timeout" on page 30055

[edit security alg mgcp application-screen]		
"connection-flood" on page 4531	"message-flood (Security MGCP)" on page 10645	"unknown-message (Security MGCP ALG)" on page 30411

[edit security alg msrpc]	
"map-entry-timeout" on page 10333	

[edit security alg sccp]		
"application-screen (Security SCCP)" on page 2362	"dscp-rewrite" on page 6019	"inactive-media-timeout (Security SCCP)" on page 8497
"traceoptions (Security SCCP ALG)" on page 29823		

[edit security alg sccp application-screen]		
"call-flood" on page 3179	"unknown-message (Security SCCP ALG)" on page 30413	

[edit security alg sip]		
"application-screen (Security SIP)" on page 2364	"bulk-call-mode" on page 3067	"c-timeout" on page 3152
"distribution-ip (Security SIP)" on page 5836	"dscp-rewrite" on page 6019	"enable-call-distribution" on page 6255
"hide-via-headers" on page 8054	"maximum-call-duration (Security)" on page 10510	"t1-interval" on page 28885
"t4-interval" on page 28893	"traceoptions (Security SIP ALG)" on page 29830	

[edit security alg sip (Security)]		
"inactive-media-timeout (Security SIP)" on page 8499	"retain-hold-resource" on page 15674	

[edit security alg sip (Security) application-screen (Security SIP)]	
"unknown-message (Security SIP ALG)" on page 30415	

[edit security alg sip application-screen protect]	
"deny (Security SIP)" on page 5317	

[edit security alg sunrpc]	
"map-entry-timeout" on page 10333	

[application-firewall]

[edit security application-firewall]		
"profile (Application Firewall)" on page 13271	"rule-sets (Security Application Firewall)" on page 16151	"traceoptions (Security Application Firewall)" on page 29776

[edit security application-firewall profile {profile-name}]		
"appfw-profile (System)" on page 2294	"block-message (Application Firewall)" on page 2989	

[edit security application-firewall rule-set {rule-set-name} rule {rule-name}]	
"then (Security Application Firewall)" on page 29282	

[edit security application-firewall rule-sets {rule-set-name}]		
"default-rule" on page 5211	"profile (Rule Sets)" on page 13291	"rule (Application Firewall)" on page 16083

[edit security application-firewall rule-sets {rule-set-name} rule {rule-name} match]		
"dynamic-application" on page 6090	"dynamic-application-group" on page 6092	"ssl-encryption" on page 28442

[application-screen]

[edit security application-screen (Security SIP)]	
"protect (Security)" on page 13406	

[application-tracking]

[edit security application-tracking]		
"disable (Application Tracking)" on page 5745	"session-update-interval" on page 16856	

[authentication-key-chains]

[edit security authentication-key-chains]	
"key-chain (Authentication Keychain)" on page 9488	

[edit security authentication-key-chains key-chain {key-chain-name}]	
"key (Authentication Keychain)" on page 9485	

[auto-re-enrollment]

[edit security auto-re-enrollment]	
"certificate-id" on page 3257	

[casb]

[edit security casb]		
"instance (CASB)" on page 8754	"traceoptions (CASB)" on page 29503	

[edit security casb casb-policy {policy-name}]	
"rule (casb-policy)" on page 16087	

[certificates]

[edit security certificates]		
"cache-size" on page 3170	"cache-timeout-negative" on page 3171	"certification-authority" on page 3263
"crl (Encryption Interface)" on page 4720	"enrollment-retry" on page 6409	"local (certificates)" on page 9900
"maximum-certificates" on page 10513	"path-length" on page 12255	

[edit security certificates certification-authority]		
"ca-name" on page 3156	"encoding (security)" on page 6315	"enrollment-url" on page 6411
"file (Security Certificates)" on page 7027	"ldap-url" on page 9678	

[cloud]

[edit security cloud]	
"high-availability (security cloud)" on page 8070	

[datapath-debug]

[edit security datapath-debug]		
"action-profile (Network Management)" on page 1773	"capture-file (Security)" on page 3206	"maximum-capture-size (Datapath Debug)" on page 10512
"packet-filter" on page 12141	"traceoptions (Security Datapath Debug)" on page 29783	

[dynamic-address]

[edit security dynamic-address]		
"address-name" on page 1892	"feed-server" on page 6999	"traceoptions (DHCP)" on page 29517
"traceoptions (dynamic-address)" on page 29533	"traceoptions (dynamic-application)" on page 29535	

[edit security dynamic-address address-name {name}]	
"session-scan" on page 16847	

[edit security dynamic-address feed-server {feed-server-name}]	
"url (Security)" on page 30487	

[edit security dynamic-address session-scan]	
"hold-interval" on page 8089	

[dynamic-application]

[edit security dynamic-application]		
"profile(dynamic-application)" on page 13275	"traceoptions (dynamic-application)" on page 29535	

[dynamic-vpn]

[edit security dynamic-vpn]		
"clients (Security) " on page 4347	"traceoptions (Security Dynamic VPN)" on page 29786	

[edit security dynamic-vpn traceoptions (Security Dynamic VPN)],]	
"file" on page 7023	

[firewall-authentication]

[edit security firewall-authentication]	
"traceoptions (Security Firewall Authentication)" on page 29788	

[flow]

[edit security flow]		
"aging" on page 2089	"allow-dns-reply" on page 2170	"allow-embedded-icmp" on page 2177
"allow-reverse-ecmp" on page 2201	"drop-flow" on page 5947	"force-ip-reassembly" on page 7371
"fragment-limit" on page 7541	"gre-performance-acceleration (Security Flow)" on page 7796	"packet-log (Security Flow)" on page 12150

(Continued)

[edit security flow]		
"pending-sess-queue-length" on page 12405	"power-mode-disable" on page 12921	"power-mode-ipsec-qat" on page 12924
"preserve-incoming-fragment-size" on page 13119	"reassembly-timeout" on page 14102	"route-change-timeout" on page 15852
"syn-flood-protection-mode" on page 28779	"tcp-mss (Security Flow)" on page 29000	"tcp-session" on page 29020
"traceoptions (Security Flow)" on page 29789		

[edit security flow (Security Flow)]	
"power-mode-ipsec" on page 12922	

[edit security flow advanced-options]	
"reverse-route-packet-mode-vr" on page 15721	

[edit security flow aging]		
"early-ageout" on page 6185	"high-watermark" on page 8080	"low-watermark" on page 10177

[edit security flow aging early-aging]		
"defaults (early-aging)" on page 5226	"protocol (early-aging)" on page 13428	

[edit security flow aging early-aging application]	
"name (early-aging)" on page 11168	

[edit security flow load-distribution]	
"session-affinity" on page 16804	

[edit security flow tcp-mss]		
"all-tcp" on page 2159	"gre-in" on page 7793	"gre-out" on page 7794

[edit security flow tcp-session]		
"fin-invalidate-session" on page 7165	"no-sequence-check" on page 11580	"no-syn-check" on page 11588
"no-syn-check-in-tunnel" on page 11589	"rst-invalidate-session" on page 16043	"rst-sequence-check" on page 16044
"strict-syn-check" on page 28653	"tcp-initial-timeout" on page 28989	"time-wait-state" on page 29361

[edit security flow tcp-session time-wait-state]	
"apply-to-half-close-state" on page 2435	

[forwarding-options]

[edit security forwarding-options]		
"mirror-filter (Security Forwarding Options)" on page 10768	"resource-manager (Forwarding-options)" on page 15616	"security-service (Security Forwarding Options)" on page 16513

[edit security forwarding-options family]		
"inet (forwarding options)" on page 8564	"inet6 (forwarding optons)" on page 8605	"mpls (Security Forwarding Options)" on page 10985
"mpls (security)" on page 10974		

[edit security forwarding-options family inet6]	
"mode (Security Forwarding Options)" on page 10819	

[edit security forwarding-options mirror-filter {filter-name}]		
"destination-port (Security Forwarding Options)" on page 5479	"destination-prefix (Security Forwarding Options)" on page 5507	"interface-in (Security Forwarding Options)" on page 8903
"interface-out (Security Forwarding Options)" on page 8948	"output (Security Forwarding Options)" on page 12013	"protocol (Security Forwarding Options)" on page 13461
"source-port (Security Forwarding Options)" on page 28281	"source-prefix (Security Forwarding Options)" on page 28303	

[edit security forwarding-options services-offload]		
"tcp-seq-ooo-window (Security Forwarding Options)" on page 29019	"tcp-session-install-interval (Security Forwarding Options)" on page 29022	

[edit security forwarding-options tunnel udp payload-port-profile {profile-name}]	
"mpls (forwarding-options)" on page 10975	

[forwarding-process]

[edit security forwarding-process]	
"application-services (Security Forwarding Process)" on page 2375	

[edit security forwarding-process application-services]		
"enable-gtpu-distribution" on page 6264	"maximize-idp-sessions" on page 10495	

[edit security forwarding-process application-services (Security Forwarding Process)]	
"inline-fpga-crypto" on page 8668	

[edit security forwarding-process application-services maximize-idp-sessions]		
"inline-tap" on page 8677	"weight (Security) " on page 31099	

[group-vpn]

[edit security group-vpn]		
"member (Security Group VPN)" on page 10621	"server (Security Group VPN)" on page 16667	

[edit security group-vpn member]		
"ike (Security Group VPN Member)" on page 8438	"ipsec (Security Group VPN Member)" on page 9314	

[edit security group-vpn member ike]		
"gateway (Security Group VPN Member IKE)" on page 7658	"policy (Security Group VPN IKE)" on page 12711	"proposal (Security Group VPN Member IKE)" on page 13383
"traceoptions (Security Group VPN)" on page 29795		

[edit security group-vpn member ike policy {policy-name}]	
"mode (Security Group VPN)" on page 10821	

[edit security group-vpn member ipsec]	
"forward-policy-mismatch (Security Group VPN Member)" on page 7402	

[edit security group-vpn server]		
"group (Security Group VPN)" on page 7894	"ike (Security Group VPN Server)" on page 8440	"ipsec (Security Group VPN Server)" on page 9317

(Continued)

[edit security group-vpn server]		
"traceoptions (Security Group VPN)" on page 29795		

[edit security group-vpn server group {name}]		
"ipsec-sa (Security Group VPN)" on page 9310	"server-cluster (Security Group VPN Server)" on page 16619	"server-member-communication (Security Group VPN Server)" on page 16646

[edit security group-vpn server ike]		
"gateway (Security Group VPN Server IKE)" on page 7660	"policy (Security Group VPN IKE)" on page 12711	"proposal (Security Group VPN Server IKE)" on page 13385

[edit security group-vpn server ike gateway {gateway-name}]		
"dead-peer-detection (Security Group VPN Server)" on page 5149		

[edit security group-vpn server ike policy {policy-name}]		
"mode (Security Group VPN)" on page 10821		

[edit security group-vpn server ipsec]		
"proposal (Security Group VPN Server IPsec)" on page 13387		

[gtp]

[edit security gtp]		
"grouped-ie-profile" on page 7915	"handover-default" on page 7964	"ie-set" on page 8343
"message-ie-profile-v1" on page 10647	"message-ie-profile-v2" on page 10648	"message-list" on page 10650
"profile (Security GTP)" on page 13318	"rate-limit (Aggregated rate limit)" on page 14065	"traceoptions (Security GTP)" on page 29798

[edit security gtp profile {profile-name}]		
"apn" on page 2286	"drop (Security GTP)" on page 5969	"handover-group" on page 7966
"listening-mode" on page 9843	"log (Security GTP)" on page 10032	"mask uli" on page 10358
"max-message-length" on page 10448	"min-message-length" on page 10713	"path-rate-limit" on page 12263
"rate-limit (Security GTP)" on page 14068	"remove-ie" on page 14366	"req-timeout" on page 14409
"restart-path" on page 15662	"seq-number-validated" on page 16594	"timeout (Security GTP)" on page 29388

[edit security gtp profile {profile-name} apn {pattern-string}]	
"imsi-prefix" on page 8491	

[edit security gtp profile {profile-name} apn {pattern-string} imsi-prefix{imsi-prefix}]	
"action (APN GTP)" on page 1760	

[edit security gtp profile {profile-name} path-rate-limit]		
"message-type" on page 10659		

[edit security gtp profile {profile-name} path-rate-limit message-type]		
"create-req" on page 4708	"delete-req" on page 5282	"echo-req" on page 6188
"other" on page 11945		

[edit security gtp profile {profile-name} path-rate-limit message-type (create-req | delete-req | echo-req | other)]		
"alarm-threshold (Security GPRS)" on page 2120	"drop-threshold (Security GPRS)" on page 5980	

[host-vpn]

[edit security host-vpn]		
"connections (Host VPN)" on page 4544	"ike-log" on page 8426	"ike-secrets" on page 8431
"remote (Host VPN)" on page 14324		

[edit security host-vpn connections {connection-name}]		
"children" on page 3315		

[edit security host-vpn connections {connection-name} children {child-name}]		
"local-traffic-selector" on page 9980	"remote-traffic-selector" on page 14361	

[idp]

[edit security idp]		
"active-policy" on page 1811	"custom-attack" on page 4738	"custom-attack-group" on page 4746
"default-policy (IDP)" on page 5203	"dynamic-attack-group" on page 6097	"idp-policy (Security) " on page 8317
"security-package" on page 16505	"sensor-configuration" on page 16560	"traceoptions (Security IDP)" on page 29803

[edit security idp (Security) dynamic-attack-group {name} filters]		
"file-type" on page 7093	"vulnerability-type" on page 30990	

[edit security idp custom-attack {attack-name}]		
"attack-type (Security Anomaly)" on page 2529	"attack-type (Security Chain)" on page 2531	"attack-type (Security Signature)" on page 2540
"recommended-action" on page 14131	"severity (Security IDP Custom Attack)" on page 16901	"time-binding" on page 29354

[edit security idp custom-attack {attack-name} attack-type]		
"anomaly" on page 2254	"chain" on page 3274	"signature (Security IDP)" on page 27955

[edit security idp custom-attack {attack-name} attack-type anomaly]		
"direction (Security Custom Attack)" on page 5724	"service (Security IDP Anomaly Attack)" on page 16736	"shellcode" on page 16977

(Continued)

[edit security idp custom-attack {attack-name} attack-type anomaly]		
"test (Security IDP)" on page 29214		

[edit security idp custom-attack {attack-name} attack-type chain]		
"expression" on page 6694	"member (Security IDP)" on page 10624	"order (Security IDP)" on page 11930
"protocol-binding" on page 13433	"reset (Security IDP)" on page 15582	"scope (Security IDP Chain Attack)" on page 16382

[edit security idp custom-attack {attack-name} attack-type chain member {member-name}]		
"attack-type (Security IDP)" on page 2533		

[edit security idp custom-attack {attack-name} attack-type chain protocol-binding]		
"application (Security Custom Attack)" on page 2367	"icmp (Security IDP Custom Attack)" on page 8264	"icmpv6 (Security IDP)" on page 8277
"ip (Security IDP Custom Attack)" on page 9266	"nested-application (Security IDP)" on page 11286	"rpc" on page 16003
"tcp (Security IDP Protocol Binding)" on page 29013	"udp (Security IDP Protocol Binding)" on page 30312	

[edit security idp custom-attack {attack-name} attack-type signature]		
"context (Security Custom Attack)" on page 4614	"direction (Security Custom Attack)" on page 5724	"negate" on page 11260

(Continued)

[edit security idp custom-attack {attack-name} attack-type signature]		
"optional-parameters" on page 11910	"pattern (Security IDP)" on page 12275	"pattern-pcre (Security IDP)" on page 12274
"protocol (Security IDP Signature Attack)" on page 13463	"protocol-binding" on page 13433	"regexp" on page 14261
"shellcode" on page 16977		

[edit security idp custom-attack {attack-name} attack-type signature protocol]		
"icmp (Security IDP Signature Attack)" on page 8265	"icmpv6 (Security IDP Custom Attack)" on page 8278	"ipv4 (Security IDP Signature Attack)" on page 9340
"reserved (Security IDP Custom Attack)" on page 15580	"tcp (Security IDP Signature Attack)" on page 29014	"udp (Security IDP Signature Attack)" on page 30313

[edit security idp custom-attack {attack-name} attack-type signature protocol {ipv6} extension-header destination-option]	
"option-type" on page 11906	

[edit security idp custom-attack {attack-name} attack-type signature protocol icmp]		
"checksum-validate" on page 3313	"code (IDP)" on page 4386	"data-length" on page 4838
"identification (Security ICMP Headers)" on page 8286	"sequence-number (Security IDP ICMP Headers)" on page 16597	"type (Security IDP ICMP Headers)" on page 30288

[edit security idp custom-attack {attack-name} attack-type signature protocol icmpv6]		
"checksum-validate" on page 3313	"code (IDP)" on page 4386	"data-length" on page 4838

(Continued)

[edit security idp custom-attack {attack-name} attack-type signature protocol icmpv6]		
"identification (Security ICMP Headers)" on page 8286	"sequence-number (Security IDP ICMP Headers)" on page 16597	

[edit security idp custom-attack {attack-name} attack-type signature protocol ipv4]		
"checksum-validate" on page 3313	"destination (Security IP Headers Attack)" on page 5524	

[edit security idp custom-attack {attack-name} attack-type signature protocol ipv6]	
"destination (Security IP Headers Attack)" on page 5524	

[edit security idp custom-attack {attack-name} attack-type signature protocol tcp]		
"ack-number" on page 1750	"checksum-validate" on page 3313	"data-length" on page 4838
"destination-port (Security Signature Attack)" on page 5485	"header-length" on page 8014	"mss (Security IDP)" on page 11002
"option (Security IDP)" on page 11905	"source-port (Security IDP)" on page 28286	"tcp-flags" on page 28983
"urgent-pointer" on page 30483	"window-scale" on page 31122	"window-size" on page 31123

[edit security idp custom-attack {attack-name} attack-type signature protocol udp]		
"checksum-validate" on page 3313	"data-length" on page 4838	"destination-port (Security Signature Attack)" on page 5485
"source-port (Security IDP)" on page 28286		

[edit security idp custom-attack {attack-name} attack-type signature protocol-binding]		
"application (Security Custom Attack)" on page 2367	"icmp (Security IDP Custom Attack)" on page 8264	"icmpv6 (Security IDP)" on page 8277
"ip (Security IDP Custom Attack)" on page 9266	"nested-application (Security IDP)" on page 11286	"rpc" on page 16003
"tcp (Security IDP Protocol Binding)" on page 29013	"udp (Security IDP Protocol Binding)" on page 30312	

[edit security idp custom-attack {attack-name} time-binding]		
"count (Security Custom Attack)" on page 4688	"scope (Security IDP Custom Attack)" on page 16383	

[edit security idp custom-attack {attack-name }attack-type signature protocol tcp]	
"sequence-number (Security IDP TCP Headers)" on page 16598	

[edit security idp custom-attack-group {custom-attack-group-name}]	
"group-members" on page 7849	

[edit security idp dynamic-attack-group {dynamic-attack-group-name}]	
"filters" on page 7162	

[edit security idp dynamic-attack-group {name} filters]		
"age-of-attack" on page 2043	"category (Security Dynamic Attack Group)" on page 3223	"cvss-score" on page 4787

(Continued)

[edit security idp dynamic-attack-group {name} filters]		
"direction (Security Dynamic Attack Group)" on page 5726	"false-positives" on page 6818	"no-recommended" on page 11558
"performance" on page 12428	"products" on page 13253	"recommended" on page 14130
"service (Security IDP Dynamic Attack Group)" on page 16737	"severity (Security IDP Dynamic Attack Group)" on page 16903	"type (Security IDP Dynamic Attack Group)" on page 30287
"vendor" on page 30675		

[edit security idp idp-policy {policy-name}]		
"rulebase-exempt" on page 16157	"rulebase-ips" on page 16159	

[edit security idp idp-policy (Security) {name} rulebase-ips {name} rule (Security IPS Rulebase) then (Security IDP Policy) application-services security-intelligence]	
"security-intelligence (IDP)" on page 16465	

[edit security idp idp-policy {policy-name} rulebase-exempt]	
"rule (Security Exempt Rulebase)" on page 16107	

[edit security idp idp-policy {policy-name} rulebase-exempt rule {rule-name}]		
"description (Security IDP Policy)" on page 5349	"match (Security IDP Policy)" on page 10395	

[edit security idp idp-policy {policy-name} rulebase-exempt rule {rule-name} match]		
"attacks (Security Exempt Rulebase)" on page 2546	"destination-address (Security IDP Policy)" on page 5396	"destination-except" on page 5418
"from-zone (Security IDP Policy)" on page 7607	"source-address (Security IDP Policy)" on page 28162	"source-except" on page 28207
"to-zone (Security IDP Policy)" on page 29422		

[edit security idp idp-policy {policy-name} rulebase-exempt rule {rule-name} match attacks]		
"custom-attack-groups (Security IDP)" on page 4747	"custom-attacks" on page 4748	"dynamic-attack-groups (Security IDP)" on page 6100
"predefined-attack-groups" on page 13013	"predefined-attacks" on page 13015	

[edit security idp idp-policy {name} rulebase-exempt rule {name} then notification]		
"log-attacks" on page 10050	"log-attacks (rulebase-exempt)" on page 10052	

[edit security idp idp-policy {policy-name} rulebase-ips]		
"rule (Security IPS Rulebase)" on page 16109		

[edit security idp idp-policy {policy-name} rulebase-ips rule {rule-name}]		
"description (Security IDP Policy)" on page 5349	"match (Security IDP Policy)" on page 10395	"terminal" on page 29146
"then (Security IDP Policy)" on page 29285		

[edit security idp idp-policy {policy-name} rulebase-ips rule {rule-name} match]		
"application (Security IDP)" on page 2370	"attacks (Security IPS Rulebase)" on page 2547	"destination-address (Security IDP Policy)" on page 5396
"destination-except" on page 5418	"from-zone (Security IDP Policy)" on page 7607	"source-address (Security IDP Policy)" on page 28162
"source-except" on page 28207	"to-zone (Security IDP Policy)" on page 29422	

[edit security idp idp-policy {policy-name} rulebase-ips rule {rule-name} match attacks]		
"custom-attack-groups (Security IDP)" on page 4747	"custom-attacks" on page 4748	"dynamic-attack-groups (Security IDP)" on page 6100
"predefined-attack-groups" on page 13013	"predefined-attacks" on page 13015	

[edit security idp idp-policy {policy-name} rulebase-ips rule {rule-name} then]		
"action (Security Rulebase IPS)" on page 1785	"ip-action (Security IDP Rulebase IPS)" on page 9231	"notification" on page 11686
"severity (Security IDP IPS Rulebase)" on page 16904		

[edit security idp idp-policy {policy-name} rulebase-ips rule {rule-name} then notification]		
"log-attacks" on page 10007	"packet-log (Security IDP Policy)" on page 12152	

[edit security idp idp-policy {policy-name} rulebase-ips rule {rule-name} then notification packet-log]		
"post-attack" on page 12898	"post-attack-timeout" on page 12900	"pre-attack" on page 12993

[edit security idp security-package]		
"automatic (Security)" on page 2757	"install" on page 8741	"source-address (Security IDP)" on page 28161
"url (Security IDP)" on page 30506		

[edit security idp security-package automatic]		
"download-timeout" on page 5930	"enable" on page 6251	"interval (Security IDP)" on page 9220
"start-time (Security IDP)" on page 28504		

[edit security idp sensor-configuration]		
"application-identification (IDP Sensor Configuration)" on page 2336	"detector" on page 5551	"flow (Security IDP)" on page 7339
"global (Security IDP)" on page 7733	"high-availability (Security IDP)" on page 8073	"ips" on page 9295
"log (Security IDP Sensor Configuration)" on page 10033	"packet-log (Security IDP Sensor Configuration)" on page 12153	"re-assembler" on page 14079
"ssl-inspection" on page 28443		

[edit security idp sensor-configuration application-identification]		
"max-packet-memory-ratio" on page 10451	"max-reass-packet-memory-ratio" on page 10463	"max-tcp-session-packet-memory" on page 10484
"max-udp-session-packet-memory" on page 10489		

[edit security idp sensor-configuration detector]	
"protocol-name" on page 13451	

[edit security idp sensor-configuration detector protocol-name {protocol-name}]	
"tunable-name" on page 30184	

[edit security idp sensor-configuration detector protocol-name {protocol-name} tunable-name {tunable-name}]	
"tunable-value" on page 30186	

[edit security idp sensor-configuration flow]	
"allow-icmp-without-flow" on page 2182	

[edit security idp sensor-configuration global]		
"enable-all-qmodules" on page 6253	"enable-packet-pool" on page 6268	"policy-lookup-cache" on page 12685

[edit security idp sensor-configuration global intelligent-offload-tunable]	
"intelligent-offload-tunable (IDP)" on page 8780	

[edit security idp sensor-configuration high-availability]	
"no-policy-cold-synchronization" on page 11548	

[edit security idp sensor-configuration ips]		
"pre-filter-shellcode" on page 12994		

[edit security idp sensor-configuration log]		
"suppression" on page 28731		

[edit security idp sensor-configuration log suppression]		
"max-time-report" on page 10486	"start-log" on page 28493	

[edit security idp sensor-configuration packet-log]		
"host (Security IDP Sensor Configuration)" on page 8165	"max-sessions (Security Packet Log)" on page 10478	"source-address (Security IDP Sensor Configuration)" on page 28163
"total-memory" on page 29452		

[edit security idp sensor-configuration re-assembler]		
"ignore-memory-overflow" on page 8407	"ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-overflow" on page 8409	"ignore-reassembly-overflow" on page 8411

[ike]

[edit security ike]		
"blocklists (Security IKE)" on page 2997	"gateway (Security IKE)" on page 7663	"identity" on page 8290

(Continued)

[edit security ike]		
"policy (Security IKE)" on page 12713	"policy (Security Interfaces)" on page 12716	"proposal (Security IKE Dynamic SA)" on page 13393
"proposal (Security IKE)" on page 13389	"session (Security IKE)" on page 16848	"traceoptions (Security IKE)" on page 29806

[edit security ike gateway {gateway-name}]		
"aaa" on page 1605	"advpn" on page 2036	"dead-peer-detection" on page 5147
"dynamic (Security)" on page 6151	"local-identity" on page 9950	"remote-identity" on page 14340

[edit security ike gateway {gateway-name} dynamic]		
"distinguished-name (Security)" on page 5831	"general-ikeid" on page 7669	

[edit security ike policy {policy-name}]		
"certificate" on page 3254	"description (IKE policy)" on page 5333	"description (Security Policies Next Gen Services)" on page 5354
"encoding (security)" on page 6315	"local-certificate (Security)" on page 9932	"local-key-pair" on page 9958
"mode (IKE)" on page 10802	"pre-shared-key (Security)" on page 13004	"proposals" on page 13403

[edit security ike proposal {ike-proposal-name}]		
"authentication-algorithm (Security IKE)" on page 2580	"authentication-method" on page 2627	"description (IKE policy)" on page 5333
"description (Security Policies Next Gen Services)" on page 5354	"dh-group" on page 5579	"dh-group (Security IKE)" on page 5581
"encryption-algorithm (Security IKE)" on page 6327	"encryption-algorithm (Security)" on page 6326	"lifetime-seconds (Security)" on page 9750
"Signature Hash Algorithm (Security IKE)" on page 27953		

[edit security ike trap]	
"trap (Security IKE)" on page 30139	

[ipsec]

[edit security ipsec]		
"anti-replay-window-size" on page 2255	"internal" on page 9186	"internal (Security IPsec)" on page 9190
"policy (Security IPsec Interfaces)" on page 12722	"policy (Security IPsec)" on page 12718	"proposal (Security IPsec)" on page 13394
"security-association (IPSec)" on page 16454	"security-association (Junos OS)" on page 16460	"security-association (Junos-FIPS Software)" on page 16458
"traceoptions (Security IPsec)" on page 29809	"vpn (Security)" on page 30946	

[edit security ipsec internal security-association manual]		
"direction (Junos-FIPS Software)" on page 5718		

[edit security ipsec internal security-association manual direction]		
"protocol (Junos-FIPS Software)" on page 13445	"spi (Junos-FIPS Software)" on page 28386	

[edit security ipsec internal security-association manual direction encryption]		
"algorithm (Junos FIPS)" on page 2145	"key (Junos FIPS)" on page 9494	

[edit security ipsec policy {ipsec-policy-name}]		
"description (IKE policy)" on page 5333	"description (Security Policies Next Gen Services)" on page 5354	"perfect-forward-secrecy (Security)" on page 12424
"proposals" on page 13403		

[edit security ipsec proposal {ipsec-proposal-name}]		
"authentication-algorithm (Security IPsec)" on page 2582	"description (IKE policy)" on page 5333	"description (Security Policies Next Gen Services)" on page 5354
"encryption-algorithm (Security)" on page 6326	"lifetime-seconds (Security)" on page 9750	"protocol (Junos OS)" on page 13446

[edit security ipsec security-association {sa-name}]		
"description (IKE policy)" on page 5333	"dynamic" on page 6084	"manual (Junos OS)" on page 10318

(Continued)

[edit security ipsec security-association {sa-name}]		
"mode (IPsec)" on page 10805		

[edit security ipsec security-association {sa-name} manual]		
"direction (Junos OS)" on page 5719		

[edit security ipsec security-association {sa-name} manual direction (inbound | outbound | bi-directional)]		
"authentication (Security IPsec)" on page 2650	"auxiliary-spi (Security IPsec)" on page 2774	"spi (Junos OS)" on page 28387

[edit security ipsec security-association {sa-name} manual direction (inbound | outbound | bidirectional)]		
"encryption (Junos OS)" on page 6332	"protocol (Junos OS)" on page 13446	

[edit security ipsec vpn {vpn-name}]		
"ike (Security IPsec VPN)" on page 8443	"manual (Security IPsec)" on page 10320	"multi-sa" on page 11055
"traffic-selector" on page 30042	"vpn-monitor" on page 30943	

[edit security ipsec vpn hub-to-spoke-vpn tunnel-mtu {tunnel-mtu}]		
"tunnel-mtu" on page 30216		

[edit security ipsec vpn {hub-to-spoke-vpn} no-icmp-packet-too-big]	
"no-icmp-packet-too-big" on page 11472	

[edit security ipsec vpn {vpn-name} vpn-monitor]	
"verify-path" on page 30688	

[edit security ipsec vpn vpn-name]	
"passive-mode-tunneling (Security IPsec)" on page 12208	

[log]

[edit security log]		
"cache (Security Log)" on page 3168	"event-rate" on page 6531	"facility-override (Security)" on page 6771
"file (Security Log)" on page 7076	"format (Security Log)" on page 7386	"mode (Security Log)" on page 10823
"profile (security)" on page 13292	"rate-cap" on page 14058	"report (Security Log)" on page 14398
"retry-count (Security Log)" on page 15696	"root-streaming" on page 15838	"stream (Security Log)" on page 28635
"time-format" on page 29356	"traceoptions (Security Log)" on page 29811	"transport (Security Log)" on page 30129

[edit security log cache]		
"exclude (Security Log)" on page 6627	"limit (Security Log)" on page 9760	

[edit security log mode event]	
"file (Security Log)" on page 7076	

[edit security log mode stream]	
"file (Security Log)" on page 7076	

[edit security log mode stream-event]	
"file (Security Log)" on page 7076	

[edit security log stream {stream-name}]		
"category (Security Logging)" on page 3224	"filter (Security)" on page 7153	"format (Security Log Stream)" on page 7388
"host (Security Log)" on page 8166	"radius (Security Log)" on page 14018	"rate-limit (Security Log)" on page 14069
"severity (Security Log)" on page 16906		

[logical-systems]

[edit security logical-systems {logical-system-name}]	
"ngfw" on page 11369	

[macsec]

[edit security macsec]		
"connectivity-association" on page 4548	"connectivity-association (MX Series)" on page 4554	"interfaces (MACsec for MX Series)" on page 9149
"interfaces (MACsec)" on page 9147	"traceoptions (MACsec)" on page 29599	

[edit security macsec cluster-data-port]	
"interfaces (MACsec)" on page 9147	

[edit security macsec connectivity-association {connectivity-association-name}]		
"bounded-delay" on page 3021	"cipher-suite (MACsec)" on page 3320	"disable-preceding-key" on page 5787
"exclude-protocol" on page 6610	"exclude-protocol (MX Series)" on page 6612	"fallback-key" on page 6794
"include-sci" on page 8536	"include-sci (MACsec for MX Series)" on page 8537	"mka" on page 10775

(Continued)

[edit security macsec connectivity-association {connectivity-association-name}]		
"mka (MX Series)" on page 10777	"no-encryption (MACsec for MX Series)" on page 11449	"no-encryption (MACsec)" on page 11447
"offset" on page 11791	"offset (MX Series)" on page 11793	"pre-shared-key" on page 13001
"pre-shared-key (MX Series)" on page 13003	"replay-protect" on page 14380	"replay-protect (MX Series)" on page 14382
"secure-channel" on page 16435	"security-mode" on page 16494	"security-mode (SRX)" on page 16492

[edit security macsec connectivity-association {connectivity-association-name} fallback-key]		
"cak" on page 3174	"ckn" on page 3339	

[edit security macsec connectivity-association {connectivity-association-name} mka]		
"eapol-address (MACSec)" on page 6181	"key-server-priority (MACsec for MX Series)" on page 9503	"key-server-priority (MACsec)" on page 9502
"must-secure" on page 11133	"sak-rekey-interval " on page 16280	"should-secure" on page 16992
"suspend-for" on page 28735	"suspend-on-request" on page 28737	"transmit-interval (MACsec for MX Series)" on page 30095
"transmit-interval (MACsec)" on page 30093		

[edit security macsec connectivity-association {connectivity-association-name} pre-shared-key]		
"cak" on page 3174	"cak (MX Series)" on page 3176	"ckn" on page 3339

(Continued)

[edit security macsec connectivity-association {connectivity-association-name} pre-shared-key]		
"ckn (MX Series)" on page 3341		

[edit security macsec connectivity-association {connectivity-association-name} replay-protect]		
"replay-window-size" on page 14383	"replay-window-size (MX Series)" on page 14385	

[edit security macsec connectivity-association {connectivity-association-name} secure-channel {secure-channel-name}]		
"direction (MX Series)" on page 5722	"direction (Secure-channel)" on page 5713	"encryption (MACsec for MX Series)" on page 6336
"encryption (MACsec)" on page 6334	"id (MACsec for MX Series)" on page 8283	"offset" on page 11791
"offset (MX Series)" on page 11793	"security-association" on page 16452	"security-association (MACSec)" on page 16450

[edit security macsec connectivity-association {connectivity-association-name} secure-channel {secure-channel-name} id]		
"mac-address (MACsec)" on page 10226	"port-id" on page 12831	"port-id (MACsec for MX Series)" on page 12833

[edit security macsec connectivity-association {connectivity-association-name} secure-channel {secure-channel-name} security-association {security-association-number}]		
"key (MACsec for MX Series)" on page 9498	"key (MACsec)" on page 9497	

[edit security macsec interfaces {interface-name}]		
"connectivity-association (MACsec Interfaces for MX Series)" on page 4553	"connectivity-association (MACsec Interfaces)" on page 4551	"traceoptions (MACsec interfaces)" on page 29602

[nat]

[edit security nat]		
"destination (Security Destination NAT)" on page 5521	"proxy-arp (Security NAT)" on page 13823	"proxy-ndp (Security NAT)" on page 13836
"static (Security NAT)" on page 28577	"traceoptions (Security NAT)" on page 29815	

[edit security nat destination]		
"pool (Security Destination NAT)" on page 12779	"rule-set (Security Destination NAT)" on page 16131	

[edit security nat destination pool {pool-name}]		
"address (Security Destination NAT)" on page 1924	"description (Security NAT Pool)" on page 5350	"routing-instance (Security Destination NAT)" on page 15959

[edit security nat destination rule-set {rule-set-name}]		
"description (Security NAT Rule Set)" on page 5353	"from (Security NAT)" on page 7596	"rule (Security Destination NAT)" on page 16105

[edit security nat destination rule-set {rule-set-name} rule {rule-name}]		
"description (Security NAT Rule)" on page 5351	"match (Security Destination NAT)" on page 10393	"then (Security Destination NAT)" on page 29284

[edit security nat destination rule-set {rule-set-name} rule {rule-name} match]		
"application (Security Destination NAT)" on page 2369	"destination-address (Security Destination NAT)" on page 5394	"destination-address-name (Security Destination NAT)" on page 5379
"destination-port (Security Destination NAT)" on page 5478	"protocol (Security Destination NAT)" on page 13460	"source-address (Security Destination NAT)" on page 28160
"source-address-name (Security Destination NAT)" on page 28140		

[edit security nat destination rule-set {rule-set-name} rule {rule-name} then]
"destination-nat" on page 5449

[edit security nat destination rule-set {rule-set-name} rule {rule-name} then destination-nat]
"rule-session-count-alarm (Security Destination NAT Rule Set)" on page 16124

[edit security nat proxy-arp]
"interface (Security NAT ARP Proxy)" on page 9018

[edit security nat proxy-arp interface]
"address (Security ARP Proxy)" on page 1922

[edit security nat proxy-ndp]	
"interface (Security NAT NDP Proxy)" on page 9020	

[edit security nat proxy-ndp interface]	
"address (Security NDP Proxy)" on page 1925	

[edit security nat source]		
"address-persistent (Security Source NAT)" on page 1896	"interface (Security Source NAT)" on page 9021	"pool (Security Source NAT)" on page 12781
"pool-default-port-range" on page 12768	"pool-default-twin-port-range" on page 12771	"pool-utilization-alarm" on page 12793
"port-randomization" on page 12860	"port-round-robin" on page 12867	"port-scaling-enlargement" on page 12873
"radius-accounting" on page 13997	"rule-set (Security Source NAT)" on page 16133	"subscriber-extension" on page 28692
"subscriber-pool-utilization-alarm" on page 28705		

[edit security nat source interface]		
"port-overloading (Security Source NAT Interface)" on page 12850	"port-overloading-factor (Security Source NAT Interface)" on page 12847	

[edit security nat source pool]		
"address-shared (Security Source NAT)" on page 1932	"description (Security NAT Pool)" on page 5350	"host-address-base" on page 8137
"overflow-pool" on page 12035	"pool-utilization-alarm (Security Source NAT Pool)" on page 12794	"port (Security Source NAT)" on page 12877
"routing-instance (Security Source NAT)" on page 15960		

[edit security nat source pool {name} address-persistent subscriber]	
"address-persistent (Security Source NAT Pool)" on page 1897	

[edit security nat source pool {name} address-pooling]	
"address-pooling (Security Source NAT)" on page 1905	

[edit security nat source pool pool-name port]	
"source (Security Source NAT)" on page 28322	

[edit security nat source pool {name} pool-utilization-alarm]	
"clear-threshold" on page 4245	

[edit security nat source pool port block-allocation]		
"interim-logging-interval" on page 9182	"last-block-recycle-timeout" on page 9637	

[edit security nat source pool-utilization-alarm]		
"clear-threshold" on page 4245	"raise-threshold" on page 14041	

[edit security nat source rule-set {rule-set-name}]		
"description (Security NAT Rule Set)" on page 5353	"from (Security NAT)" on page 7596	"rule (Security Source NAT)" on page 16111
"to (Security Source NAT)" on page 29421		

[edit security nat source rule-set {rule-set-name} rule {rule-name}]		
"description (Security NAT Rule)" on page 5351	"match (Security Source NAT)" on page 10401	"then (Security Source NAT)" on page 29292

[edit security nat source rule-set {rule-set-name} rule {rule-name} match]		
"application (Security Source NAT)" on page 2373	"destination-address (Security Source NAT)" on page 5400	"destination-address-name (Security Source NAT)" on page 5380
"destination-port (Security Source NAT)" on page 5486	"protocol (Security Source NAT)" on page 13469	"source-address (Security Source NAT)" on page 28166
"source-address-name (Security Source NAT)" on page 28141	"source-port (Security Source NAT Rule Set)" on page 28288	

[edit security nat source rule-set {rule-set-name} rule {rule-name } then]	
"source-nat" on page 28254	

[edit security nat source rule-set {rule-set-name} rule {rule-name} then source-nat]		
"interface (Security Source NAT Rule Set)" on page 9022	"pool (Security Source NAT Rule Set)" on page 12783	"rule-session-count-alarm (Security Source NAT Rule Set)" on page 16125

[edit security nat source rule-set {ruleset} rule {rule} then source-nat interface]	
"persistent-nat" on page 12453	

[edit security nat source rule-set {ruleset} rule {rule} then source-nat interface persistent-nat]		
"address-mapping" on page 1887	"inactivity-timeout (Security Persistent NAT)" on page 8509	"max-session-number" on page 10466
"permit (Security Persistent NAT)" on page 12444		

[edit security nat source rule-set {ruleset} rule {rule} then source-nat pool]	
"persistent-nat" on page 12453	

[edit security nat source rule-set {ruleset} rule {rule} then source-nat pool persistent-nat]		
"inactivity-timeout (Security Persistent NAT)" on page 8509	"max-session-number" on page 10466	"permit (Security Persistent NAT)" on page 12444

[edit security nat static]	
"rule-set (Security Static NAT)" on page 16135	

[edit security nat static rule-set {rule-set-name}]		
"description (Security NAT Rule Set)" on page 5353	"from (Security NAT)" on page 7596	"rule (Security Static NAT)" on page 16113

[edit security nat static rule-set {rule-set-name} rule {rule-name}]		
"description (Security NAT Rule)" on page 5351	"match (Security Static NAT)" on page 10402	"then (Security Static NAT)" on page 29293

[edit security nat static rule-set {rule-set-name} rule {rule-name} match]		
"destination-address (Security Static NAT)" on page 5401	"destination-address-name (Security Static NAT)" on page 5382	"destination-port (Security Static NAT)" on page 5487
"source-address (Security Static NAT Rule Set)" on page 28167	"source-address-name (Security Static NAT Rule Set)" on page 28142	"source-port (Security Static NAT Rule Set)" on page 28289

[edit security nat static rule-set {rule-set-name} rule {rule-name} then]	
"static-nat " on page 28551	

[edit security nat static rule-set {rule-set-name} rule {rule-name} then static-nat]		
"inet (Security Static NAT)" on page 8590	"nptv6-prefix" on page 11713	"nptv6-prefix-name" on page 11714
"prefix (Security Static NAT)" on page 13103	"prefix-name (Security Static NAT)" on page 13098	"rule-session-count-alarm (Security Static NAT Rule Set)" on page 16126

[edit security nat static rule-set {rule-set-name} rule {rule-name} then static-nat prefix]	
"mapped-port (Security Static NAT)" on page 10336	

[edit security nat static rule-set {rule-set-name} rule {rule-name} then static-nat prefix-name]	
"mapped-port (Security Static NAT)" on page 10336	

[pki]

[edit security pki]		
"auto-re-enrollment" on page 2727	"auto-re-enrollment (Security)" on page 2728	"ca-profile (Security PKI)" on page 3157
"default-trusted-ca-certs (Security)" on page 5218	"traceoptions (Security PKI)" on page 29818	"trap (Security PKI)" on page 30140

[edit security pki auto-re-enrollment certificate-id]		
"challenge-password" on page 3283	"re-enroll-trigger-time-percentage" on page 14081	"re-generate-keypair" on page 14082
"validity-period" on page 30650		

[edit security pki ca-profile {ca-profile-name}]		
"ca-identity" on page 3155	"enrollment" on page 6408	"revocation-check" on page 15737
"revocation-check (Security PKI)" on page 15739		

[edit security pki ca-profile {ca-profile-name} enrollment]		
"retry (Adaptive Services Interface)" on page 15694	"retry-interval" on page 15706	"url (Security)" on page 30503

[edit security pki ca-profile {ca-profile-name} revocation-check]		
"crl (Adaptive Services Interface)" on page 4718	"crl (Security)" on page 4721	"ocsp (Security PKI)" on page 11774

[edit security pki ca-profile {ca-profile-name} revocation-check crl]		
"refresh-interval" on page 14248	"url (Security)" on page 30503	

[policies]

[edit security policies]		
"default-policy" on page 5202	"dns-cache" on page 5852	"file-serialization" on page 7094
"from-zone (Security Policies)" on page 7608	"global (Security Policies)" on page 7734	"lookup-intact-on-commit" on page 10112
"policy-rematch" on page 12707	"policy-set " on page 12731	"policy-stats" on page 12743
"pre-id-default-policy" on page 12995	"traceoptions (Security Policies)" on page 29821	"unified-policy max-lookups" on page 30362

[edit security policies explicit-proxy]	
"profile (Explicit Proxy)" on page 13277	

[edit security policies explicit-proxy profile]	
"policy (Explicit-Proxy)" on page 12678	

[edit security policies explicit-proxy profile [name] policy [name] report-skip]	
"report-skip" on page 14403	

[edit security policies explicit-proxy profile {policy-set-policy-name} policy {policy-name} then permit]	
"tcp-options (Security Policies)" on page 29008	

[edit security policies from-zone]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	"to-zone (Security Policies)" on page 29423

[edit security policies from-zone {from-zone} to-zone {to-zone}]		
"application-services (security-metadata-streaming)" on page 2377	"policy (Security Policies)" on page 12724	

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name}]		
"destination-identity-context" on page 5427	"destination-identity-context-profile" on page 5429	"match (Security Policies)" on page 10396
"report-skip" on page 14403	"scheduler-name" on page 16361	"then (Security Policies)" on page 29287

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} match]		
"application (Security Policies)" on page 2371	"destination-address (Security Policies)" on page 5397	"destination-address-excluded" on page 5372
"dynamic-application (Security Policies)" on page 6095	"source-address (Security Policies)" on page 28164	"source-address-excluded" on page 28127

(Continued)

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} match]		
"source-end-user-profile" on page 28204	"source-identity" on page 28216	

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then]		
"count (Security Policies)" on page 4689	"deny (Security Policies)" on page 5316	"log (Security Policies)" on page 10035
"permit (Security Policies)" on page 12446	"reject (Security) " on page 14270	

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then deny application-services]	
"security-intelligence" on page 16464	

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then log]		
"session-close" on page 16806	"session-init" on page 16815	

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit]		
"application-services (Security Policies)" on page 2378	"destination-address (Security Policies Flag)" on page 5399	"firewall-authentication (Security Policies)" on page 7186
"services-offload (Security)" on page 16786	"tcp-options (Security Policies)" on page 29008	"tunnel (Security Policies)" on page 30231

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit application-services]		
"application-firewall (Application Services)" on page 2320	"application-traffic-control (Application Services)" on page 2402	"Content Security-policy (Application Services)" on page 4600
"idp (Application Services)" on page 8312	"idp-policy (Application Services)" on page 8316	"redirect-wx (Application Services)" on page 14164
"security-intelligence" on page 16464	"ssl-proxy (Application Services)" on page 28446	"uac-policy (Application Services)" on page 30301

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit application-services uac-policy]	
"captive-portal (Services UAC Policy)" on page 3205	

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit firewall-authentication]		
"pass-through (Security Policy)" on page 12194	"push-to-identity-management" on page 13880	"user-firewall" on page 30553
"web-authentication" on page 31057		

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit firewall-authentication pass-through]		
"auth-only-browser" on page 2566	"auth-user-agent" on page 2568	"ssl-termination-profile" on page 28450

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit firewall-authentication pass-through auth-only-browser]	
"auth-user-agent" on page 2568	

[edit security policies from-zone {zone-name} to-zone {zone-name } policy {policy-name} then permit firewall-authentication pass-through user-firewall]		
"web-redirect" on page 31084	"web-redirect-to-https" on page 31085	

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit firewall-authentication user-firewall]		
"auth-only-browser" on page 2566	"auth-user-agent" on page 2568	

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit firewall-authentication user-firewall auth-only-browser]	
"auth-user-agent" on page 2568	

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit tcp-options]		
"initial-tcp-mss" on page 8659	"reverse-tcp-mss" on page 15722	"sequence-check-required" on page 16595
"syn-check-required" on page 28773		

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit tunnel]		
"ipsec-group-vpn (Security Policies)" on page 9300	"ipsec-vpn (Security Policies)" on page 9326	"pair-policy" on page 12168

[edit security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then reject application-services]		
"security-intelligence" on page 16464		

[edit security policies global policy {policy-name}]		
"match (Security Policies Global)" on page 10398	"report-skip" on page 14403	

[edit security policies global policy {policy-name} match]		
"application (Security Policies)" on page 2371	"destination-address (Security Policies)" on page 5397	"dynamic-application (Security Policies)" on page 6095
"from-zone (Security Policies Global)" on page 7612	"source-address (Security Policies)" on page 28164	"source-identity" on page 28216
"to-zone (Security Policies Global)" on page 29427		

[edit security policies global policy {name} then]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	

[edit security policies policies policy-set {policy-policy-set-name} policy {policy-name} then permit]
"tcp-options (Security Policies)" on page 29008

[edit security policies policiesglobal from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit then permit]
"tcp-options (Security Policies)" on page 29008

[edit security policies stateful-firewall-rule [name] policy [name]	
"report-skip" on page 14403	

[edit security policies stateful-firewall-rule {name} policy {name} then]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	

[remote-access]

[edit security remote-access]		
"client-config (Juniper Secure Connect)" on page 4303	"compliance (Juniper Secure Connect)" on page 4483	"default-profile (Juniper Secure Connect)" on page 5209
"global-options (Juniper Secure Connect)" on page 7726	"profile (Juniper Secure Connect)" on page 13286	"traceoptions (Juniper Secure Connect)" on page 29589

[edit security remote-access client-config {name}]		
"application-bypass (Juniper Secure Connect)" on page 2314	"certificate (Juniper Secure Connect)" on page 3259	"windows-logon (Juniper Secure Connect)" on page 31124

[resource-manager]

[edit security resource-manager]		
"cpu (resource-manager)" on page 4697	"services-memory (resource-manager)" on page 16781	"session-memory (resource-manager)" on page 16826

[screen]

[edit security screen]		
"ids-option" on page 8335	"traceoptions (Security Screen)" on page 29825	"white-list (Security Screen)" on page 31107

[edit security screen ids-option {screen-name}]		
"alarm-without-drop" on page 2122	"description (Security Screen)" on page 5355	"icmp (Security Screen)" on page 8267
"ip (Security Screen)" on page 9268	"limit-session" on page 9764	"tcp (Security Screen)" on page 29017
"udp (Security Screen)" on page 30314		

[edit security screen ids-option {screen-name} icmp]		
"flood (Security ICMP)" on page 7245	"icmpv6-malformed" on page 8274	"ip-sweep" on page 9279
"large" on page 9632	"ping-death" on page 12537	

[edit security screen ids-option {screen-name} ip]		
"ipv6-extension-header" on page 9350	"ipv6-extension-header-limit" on page 9355	"ipv6-malformed-header" on page 9362
"tunnel (Security Screen)" on page 30232		

[edit security screen ids-option {screen-name} ip ipv6-extension-header]		
"destination-header" on page 5419	"hop-by-hop-header" on page 8132	

[edit security screen ids-option {ids-option-name} ip tunnel]		
"bad-inner-header" on page 2829	"gre (Ids-option)" on page 7791	"ip-in-udp" on page 9243
"ipip (Ids-option)" on page 9293		

[edit security screen ids-option {screen-name} limit-session]		
"destination-ip-based" on page 5435	"source-ip-based" on page 28238	

[edit security screen ids-option {screen-name} tcp]		
"fin-no-ack" on page 7167	"land" on page 9617	"port-scan" on page 12874
"syn-ack-ack-proxy" on page 28770	"syn-fin" on page 28776	"syn-flood" on page 28778
"syn-frag" on page 28781	"tcp-no-flag" on page 29004	"tcp-sweep" on page 29025
"winnuke" on page 31126		

[edit security screen ids-option {screen-name } tcp syn-flood]		
"alarm-threshold" on page 2119	"attack-threshold" on page 2528	"destination-threshold" on page 5525
"source-threshold" on page 28329	"timeout (Security Screen)" on page 29389	

[edit security screen ids-option {screen-name} udp]		
"flood (Security UDP)" on page 7246	"port-scan" on page 12874	"udp-sweep" on page 30316

[edit security screen ids-option {screen-name} udp flood]	
"white-list (Security Screen)" on page 31107	

[edit security screen trap]	
"trap" on page 30133	

[sctp]

[edit security sctp]		
"log (Security SCTP)" on page 10037	"multichunk-inspection" on page 11097	"nullpdu" on page 11733
"profile (Security SCTP)" on page 13323	"traceoptions (Security SCTP)" on page 29828	

[edit security sctp profile {profile-name}]		
"association-timeout" on page 2515	"drop (Security SCTP)" on page 5974	"handshake-timeout" on page 7967
"limit (Security SCTP)" on page 9762	"permit (Security SCTP)" on page 12449	

[sip]

[edit security sip (Security)]		
"protect (Security)" on page 13406		

[softwires]

[edit security softwires]		
"map-e" on page 10329	"traceoptions (Next Gen Services Softwires)" on page 29624	

[tcp-encap]

[edit security tcp-encap]		
"traceoptions (TCP Encapsulation)" on page 29891		

[trusted-channel]

[edit security trusted-channel ipsec security-association manual]		
"direction (Junos-FIPS Software)" on page 5718		

[tunnel-inspection]

[edit security tunnel-inspection]	
"inspection-profile" on page 8739	

[user-identification]

[edit security user-identification]	
"traceoptions (Security User Identification)" on page 29832	

[utm]

[edit security utm]		
"application-proxy (Security Content Security)" on page 2358	"custom-objects" on page 4771	"feature-profile" on page 6967
"ipc" on page 9288	"traceoptions (Security Content Security)" on page 29782	"utm default-configuration" on page 30614

[edit security utm application-proxy]	
"traceoptions (Security Application Proxy)" on page 29778	

[edit security utm custom-objects]		
"custom-url-category" on page 4781	"filename-extension" on page 7097	"mime-pattern" on page 10705
"protocol-command" on page 13434	"url-feed" on page 30489	"url-pattern" on page 30499

[edit security utm custom-objects custom-message]	
"custom-page-file" on page 4778	

[edit security utm custom-objects custom-message {name} type]	
"custom-page" on page 4777	

[edit security utm default-configuration]		
"action (Security Content Security Web Filtering)" on page 1783	"address-blacklist" on page 1867	"address-whitelist" on page 1937
"administrator-email (Security Fallback Block)" on page 1980	"administrator-email (Security Virus Detection)" on page 1981	"allow-email (Security Fallback Block)" on page 2175
"allow-email (Security Virus Detection)" on page 2176	"anti-spam " on page 2262	"anti-spam (Security Content Security Policy)" on page 2265
"anti-virus (Content Security Policy)" on page 2273	"anti-virus (Content Security)" on page 2269	"application-proxy (Security Content Security)" on page 2358
"block-command" on page 2985	"block-content-type" on page 2986	"block-extension" on page 2987
"block-message (Security Content Security)" on page 2993	"block-mime" on page 2994	"cache" on page 3167

(Continued)

[edit security utm default-configuration]		
"category (Security Web Filtering)" on page 3226	"Content Security" on page 4589	"content-filtering (Security Feature Profile)" on page 4586
"content-size" on page 4601	"content-size (Security Antivirus Sophos Engine)" on page 4604	"corrupt-file" on page 4676
"custom-block-message" on page 4750	"custom-message (Security Content Filtering)" on page 4758	"custom-message (Security Fallback Block)" on page 4761
"custom-message (Security Fallback Non-Block)" on page 4762	"custom-message (Security Virus Detection)" on page 4763	"custom-message-subject (Security Fallback Block)" on page 4767
"custom-message-subject (Security Fallback Non-Block)" on page 4769	"custom-message-subject (Security Virus Detection)" on page 4770	"custom-objects" on page 4771
"custom-tag-string" on page 4780	"custom-url-category" on page 4781	"decompress-layer" on page 5165
"default (Security Antivirus Sophos Engine)" on page 5214	"default (Security Antivirus)" on page 5213	"default (Security Content Security)" on page 5216
"default (Security Web Filtering)" on page 5217	"display-host (Security Fallback Block)" on page 5826	"display-host (Security Virus Detection)" on page 5827
"download-profile (Security Antivirus FTP)" on page 5926	"download-profile (Security Content Filtering FTP)" on page 5927	"engine-not-ready" on page 6350
"engine-not-ready (Security Antivirus Sophos Engine)" on page 6351	"exception " on page 6571	"exception (Security Content Filtering)" on page 6573

(Continued)

[edit security utm default-configuration]		
"fallback-block (Security Antivirus)" on page 6792	"fallback-non-block (Security Antivirus)" on page 6795	"fallback-options (Security Antivirus Juniper Express Engine)" on page 6799
"fallback-options (Security Antivirus Kaspersky Lab Engine)" on page 6800	"fallback-settings (Security Web Filtering Juniper Local)" on page 6805	"fallback-settings (Security Web Filtering Websense Redirect)" on page 6806
"fallback-settings (Security Web Filtering)" on page 6804	"feature-profile" on page 6967	"filename-extension" on page 7097
"ftp (Content Security Policy Anti-Virus)" on page 7621	"ftp (Content Security Policy Content Filtering)" on page 7622	"host (Security Web Filtering)" on page 8167
"http-profile (Security Antivirus)" on page 8198	"http-profile (Security Content Filtering)" on page 8199	"http-profile (Security Web Filtering)" on page 8200
"imap-profile (Security Content Security Policy Antivirus)" on page 8450	"imap-profile (Security Content Security Policy Content Filtering)" on page 8451	"intelligent-prescreening" on page 8779
"ipc" on page 9288	"juniper-enhanced" on page 9442	"juniper-express-engine" on page 9444
"juniper-local" on page 9447	"kaspersky-lab-engine" on page 9450	"limit (Content Security Policy)" on page 9758
"list " on page 9839	"list (Security Content Filtering Block Mime)" on page 9841	"mime-pattern" on page 10705
"mime-whitelist" on page 10706	"no-intelligent-prescreening" on page 11477	"no-notify-mail-recipient" on page 11528

(Continued)

[edit security utm default-configuration]		
"no-notify-mail-sender (Security Content Filtering Notification Options)" on page 11529	"no-notify-mail-sender (Security Fallback Block)" on page 11530	"no-notify-mail-sender (Security Virus Detection)" on page 11531
"no-sbl-default-server" on page 11575	"notification-options (Security Content Filtering)" on page 11695	"notify-mail-recipient" on page 11705
"notify-mail-sender (Security Content Filtering Notification Options)" on page 11707	"notify-mail-sender (Security Fallback Block)" on page 11708	"notify-mail-sender (Security Virus Detection)" on page 11709
"out-of-resources" on page 11963	"out-of-resources (Security Antivirus Sophos Engine)" on page 11965	"over-limit" on page 12031
"password (Security Antivirus)" on page 12245	"password-file" on page 12231	"permit-command" on page 12443
"pop3-profile (Security Content Security Policy Antivirus)" on page 12808	"pop3-profile (Security Content Security Policy Content Filtering)" on page 12809	"port (Security Antivirus)" on page 12875
"port (Security Web Filtering Server)" on page 12879	"profile (Security Antispam SBL)" on page 13310	"profile (Security Antivirus Juniper Express Engine)" on page 13312
"profile (Security Antivirus Kaspersky Lab Engine)" on page 13314	"profile (Security Content Filtering)" on page 13317	"profile (Security Sophos Engine Antivirus)" on page 13326
"profile (Security Web Filtering Juniper Enhanced)" on page 13328	"profile (Security Web Filtering Juniper Local)" on page 13330	"profile (Security Web Filtering Surf Control Integrated)" on page 13332
"profile (Security Web Filtering Websense Redirect)" on page 13333	"protocol-command" on page 13434	"proxy (Security Antivirus)" on page 13839

(Continued)

[edit security utm default-configuration]		
"quarantine-message (Security Content Security)" on page 13924	"sbl" on page 16326	"sbl-default-server" on page 16327
"scan-extension" on page 16332	"scan-mode" on page 16333	"scan-options (Security Antivirus Juniper Express Engine)" on page 16336
"scan-options (Security Antivirus Kaspersky Lab Engine)" on page 16338	"server (Security Antivirus)" on page 16665	"server (Security Web Filtering)" on page 16672
"server-connectivity" on page 16621	"site-reputation-action" on page 27992	"size (Security Web Filtering Cache)" on page 28004
"smtp-profile (Security Content Security Policy Antispam)" on page 28051	"smtp-profile (Security Content Security Policy Antivirus)" on page 28052	"smtp-profile (Security Content Security Policy Content Filtering)" on page 28053
"sockets" on page 28082	"spam-action" on page 28341	"surf-control-integrated" on page 28733
"sxl-retry" on page 28764	"sxl-timeout" on page 28765	"timeout (Security Antivirus Fallback Options Sophos Engine)" on page 29385
"timeout (Security Antivirus Fallback Options)" on page 29383	"timeout (Security Web Filtering Cache)" on page 29392	"timeout (Security Web Filtering Fallback Settings)" on page 29393
"timeout (Security Web Filtering)" on page 29391	"too-many-requests (Security Antivirus Fallback Options Sophos Engine)" on page 29432	"too-many-requests (Security Antivirus Fallback Options)" on page 29430
"too-many-requests (Security Web Filtering Fallback Settings)" on page 29433	"traceoptions (Security Antispam)" on page 29773	"traceoptions (Security Antivirus)" on page 29774

(Continued)

[edit security utm default-configuration]		
"traceoptions (Security Application Proxy)" on page 29778	"traceoptions (Security Content Filtering)" on page 29781	"traceoptions (Security Content Security)" on page 29782
"traceoptions (Security Web Filtering)" on page 29835	"trickling" on page 30152	"type (Security Antivirus Feature Profile)" on page 30282
"type (Security Content Filtering Notification Options)" on page 30284	"type (Security Fallback Block)" on page 30285	"type (Security Virus Detection)" on page 30289
"upload-profile (Security Antivirus FTP)" on page 30472	"upload-profile (Security Content Filtering FTP)" on page 30473	"url-blacklist" on page 30488
"url-whitelist " on page 30508	"url-whitelist " on page 30507	"username (Security Antivirus)" on page 30606
"virus-detection (Security Antivirus)" on page 30778	"web-filtering" on page 31066	"web-filtering (Security Content Security Policy)" on page 31076
"websense-redirect" on page 31093		

[edit security utm default-configuration anti-virus]		
"avira-engine" on page 2783	"fallback-options (Security Antivirus Sophos Engine)" on page 6802	"forwarding-mode (Security Content Security Policy)" on page 7461
"notification-options (Security Antivirus)" on page 11693	"scan-options (Security Antivirus Avira Engine)" on page 16335	"sophos-engine" on page 28109

[edit security utm default-configuration anti-virus avira-engine]		
"pattern-update (Security Antivirus)" on page 12276		

[edit security utm default-configuration anti-virus avira-engine pattern-update]		
"email-notify" on page 6247	"interval (Security Antivirus)" on page 9219	"no-autoupdate" on page 11413
"proxy-profile" on page 13838	"routing-instance (Security Content Security)" on page 15958	"start-time" on page 28500
"url (Security Antivirus)" on page 30504		

[edit security utm default-configuration anti-virus avira-engine pattern-update email-notify]		
"admin-email" on page 1963	"custom-message (Security Email Notify)" on page 4759	

[edit security utm default-configuration anti-virus scan-options]		
"content-size-limit" on page 4603	"decompress-layer-limit" on page 5166	"no-uri-check" on page 11612
"timeout (Security Antivirus Scan Options)" on page 29386	"uri-check" on page 30484	

[edit security utm default-configuration anti-virus sophos-engine]	
"notification-options (Security Antivirus)" on page 11693	

[edit security utm default-configuration anti-virus sophos-engine pattern-update]		
"proxy-profile" on page 13838	"start-time" on page 28500	

[edit security utm default-configuration anti-virus sophos-engine server]	
"source-address (Content Security)" on page 28121	

[edit security utm default-configuration antivirus scan-options]	
"scan-options (Security Antivirus Sophos Engine)" on page 16340	

[edit security utm default-configuration avira-engine pattern-update email-notify]	
"custom-message-subject (Security Email Notify)" on page 4766	

[edit security utm default-configuration url-feed]	
"global-limit" on page 7708	

[edit security utm default-configuration web-filtering]		
"http-persist" on page 8196	"http-reassemble" on page 8201	"type (Security Web Filtering)" on page 30291

[edit security utm default-configuration web-filtering juniper-enhanced]	
"custom-message (Security Web Filtering)" on page 4764	

[edit security utm default-configuration web-filtering juniper-enhanced server]	
"source-address (Content Security)" on page 28121	

[edit security utm default-configuration web-filtering juniper-local]	
"custom-message (Security Web Filtering)" on page 4764	

[edit security utm default-configuration web-filtering ng-juniper]	
"ng-juniper" on page 11366	

[edit security utm default-configuration web-filtering websense-redirect]	
"custom-message (Security Web Filtering)" on page 4764	

[edit security utm default-configuration web-filtering websense-redirect server]	
"source-address (Content Security)" on page 28121	

[edit security utm feature-profile]		
"anti-spam " on page 2262	"anti-virus (Content Security)" on page 2269	"content-filtering (Security Feature Profile)" on page 4586
"web-filtering" on page 31066		

[edit security utm feature-profile anti-spam]		
"address-blacklist" on page 1867	"address-whitelist" on page 1937	"sbl" on page 16326

(Continued)

[edit security utm feature-profile anti-spam]		
"traceoptions (Security Antispam)" on page 29773		

[edit security utm feature-profile anti-spam sbl]		
"profile (Security Antispam SBL)" on page 13310		

[edit security utm feature-profile anti-spam sbl profile {profile-name}]		
"custom-tag-string" on page 4780	"no-sbl-default-server" on page 11575	"sbl-default-server" on page 16327
"spam-action" on page 28341		

[edit security utm feature-profile anti-virus]		
"juniper-express-engine" on page 9444	"kaspersky-lab-engine" on page 9450	"mime-whitelist" on page 10706
"traceoptions (Security Antivirus)" on page 29774	"type (Security Antivirus Feature Profile)" on page 30282	"url-whitelist " on page 30508

[edit security utm feature-profile anti-virus avira-engine profile {profile-name} notification-options]		
"virus-detection (Security Antivirus)" on page 30778		

[edit security utm feature-profile anti-virus juniper-express-engine]		
"pattern-update (Security Antivirus)" on page 12276	"profile (Security Antivirus Juniper Express Engine)" on page 13312	

[edit security utm feature-profile anti-virus juniper-express-engine pattern-update]		
"email-notify" on page 6247	"interval (Security Antivirus)" on page 9219	"no-autoupdate" on page 11413
"proxy (Security Antivirus)" on page 13839	"url (Security Antivirus)" on page 30504	

[edit security utm feature-profile anti-virus juniper-express-engine pattern-update email-notify]		
"admin-email" on page 1963	"custom-message (Security Email Notify)" on page 4759	"custom-message-subject (Security Email Notify)" on page 4766

[edit security utm feature-profile anti-virus juniper-express-engine pattern-update proxy]		
"password (Security Antivirus)" on page 12245	"port (Security Antivirus)" on page 12875	"server (Security Antivirus)" on page 16665
"username (Security Antivirus)" on page 30606		

[edit security utm feature-profile anti-virus juniper-express-engine profile {profile-name}]		
"fallback-options (Security Antivirus Juniper Express Engine)" on page 6799	"scan-options (Security Antivirus Juniper Express Engine)" on page 16336	"trickling" on page 30152

[edit security utm feature-profile anti-virus juniper-express-engine profile {profile-name} fallback-options]		
"content-size" on page 4601	"default (Security Antivirus)" on page 5213	"engine-not-ready" on page 6350
"out-of-resources" on page 11963	"timeout (Security Antivirus Fallback Options)" on page 29383	"too-many-requests (Security Antivirus Fallback Options)" on page 29430

[edit security utm feature-profile anti-virus juniper-express-engine profile {profile-name} notification-options]		
"fallback-block (Security Antivirus)" on page 6792	"fallback-non-block (Security Antivirus)" on page 6795	"virus-detection (Security Antivirus)" on page 30778

[edit security utm feature-profile anti-virus juniper-express-engine profile {profile-name} notification-options fallback-block]		
"administrator-email (Security Fallback Block)" on page 1980	"allow-email (Security Fallback Block)" on page 2175	"custom-message (Security Fallback Block)" on page 4761
"custom-message-subject (Security Fallback Block)" on page 4767	"display-host (Security Fallback Block)" on page 5826	"no-notify-mail-sender (Security Fallback Block)" on page 11530
"notify-mail-sender (Security Fallback Block)" on page 11708	"type (Security Fallback Block)" on page 30285	

[edit security utm feature-profile anti-virus juniper-express-engine profile {profile-name} notification-options fallback-non-block]		
"custom-message (Security Fallback Non-Block)" on page 4762	"custom-message-subject (Security Fallback Non-Block)" on page 4769	"no-notify-mail-recipient" on page 11528
"notify-mail-recipient" on page 11705		

[edit security utm feature-profile anti-virus juniper-express-engine profile {profile-name} notification-options virus-detection]		
"custom-message (Security Virus Detection)" on page 4763	"custom-message-subject (Security Virus Detection)" on page 4770	"no-notify-mail-sender (Security Virus Detection)" on page 11531
"notify-mail-sender (Security Virus Detection)" on page 11709	"type (Security Virus Detection)" on page 30289	

[edit security utm feature-profile anti-virus juniper-express-engine profile {profile-name} scan-options]		
"content-size-limit" on page 4603	"intelligent-prescreening" on page 8779	"no-intelligent-prescreening" on page 11477
"timeout (Security Antivirus Scan Options)" on page 29386		

[edit security utm feature-profile anti-virus kaspersky-lab-engine]		
"pattern-update (Security Antivirus)" on page 12276	"profile (Security Antivirus Kaspersky Lab Engine)" on page 13314	

[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update]		
"email-notify" on page 6247	"interval (Security Antivirus)" on page 9219	"no-autoupdate" on page 11413
"proxy (Security Antivirus)" on page 13839	"url (Security Antivirus)" on page 30504	

[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update email-notify]		
"admin-email" on page 1963	"custom-message (Security Email Notify)" on page 4759	"custom-message-subject (Security Email Notify)" on page 4766

[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update proxy]		
"password (Security Antivirus)" on page 12245	"port (Security Antivirus)" on page 12875	"server (Security Antivirus)" on page 16665
"username (Security Antivirus)" on page 30606		

[edit security utm feature-profile anti-virus kaspersky-lab-engine profile {profile-name}]		
"fallback-options (Security Antivirus Kaspersky Lab Engine)" on page 6800	"scan-options (Security Antivirus Kaspersky Lab Engine)" on page 16338	"trickling" on page 30152

[edit security utm feature-profile anti-virus kaspersky-lab-engine profile {profile-name} fallback-options]		
"content-size" on page 4601	"corrupt-file" on page 4676	"decompress-layer" on page 5165
"default (Security Antivirus)" on page 5213	"engine-not-ready" on page 6350	"out-of-resources" on page 11963
"password-file" on page 12231	"timeout (Security Antivirus Fallback Options)" on page 29383	"too-many-requests (Security Antivirus Fallback Options)" on page 29430

[edit security utm feature-profile anti-virus kaspersky-lab-engine profile {profile-name} notification-options]		
"fallback-block (Security Antivirus)" on page 6792	"fallback-non-block (Security Antivirus)" on page 6795	

[edit security utm feature-profile anti-virus kaspersky-lab-engine profile {profile-name} notification-options fallback-block]		
"administrator-email (Security Fallback Block)" on page 1980	"allow-email (Security Fallback Block)" on page 2175	"custom-message (Security Fallback Block)" on page 4761
"custom-message-subject (Security Fallback Block)" on page 4767	"display-host (Security Fallback Block)" on page 5826	"no-notify-mail-sender (Security Fallback Block)" on page 11530
"notify-mail-sender (Security Fallback Block)" on page 11708	"type (Security Fallback Block)" on page 30285	

[edit security utm feature-profile anti-virus kaspersky-lab-engine profile {profile-name} notification-options fallback-non-block]		
"custom-message (Security Fallback Non-Block)" on page 4762	"custom-message-subject (Security Fallback Non-Block)" on page 4769	"no-notify-mail-recipient" on page 11528
"notify-mail-recipient" on page 11705		

[edit security utm feature-profile anti-virus kaspersky-lab-engine profile {profile-name} notification-options virus-detection]		
"custom-message (Security Virus Detection)" on page 4763	"custom-message-subject (Security Virus Detection)" on page 4770	"no-notify-mail-sender (Security Virus Detection)" on page 11531
"notify-mail-sender (Security Virus Detection)" on page 11709	"type (Security Virus Detection)" on page 30289	

[edit security utm feature-profile anti-virus kaspersky-lab-engine profile {profile-name} scan-options]		
"content-size-limit" on page 4603	"intelligent-prescreening" on page 8779	"no-intelligent-prescreening" on page 11477
"scan-extension" on page 16332	"scan-mode" on page 16333	"timeout (Security Antivirus Scan Options)" on page 29386

[edit security utm feature-profile anti-virus mime-whitelist]		
"exception " on page 6571	"list " on page 9839	

[edit security utm feature-profile anti-virus mime-whitelist list {listname}]	
"exception " on page 6571	

[edit security utm feature-profile anti-virus profile]		
"notification-options (Security Antivirus)" on page 11693		

[edit security utm feature-profile anti-virus profile notification-options virus-detect]		
"allow-email (Security Virus Detection)" on page 2176		

[edit security utm feature-profile anti-virus profile {profile name} notification-options virus-detection]		
"display-host (Security Virus Detection)" on page 5827		

[edit security utm feature-profile anti-virus sophos-engine]		
"profile (Security Sophos Engine Antivirus)" on page 13326	"server (Security Sophos Engine Antivirus)" on page 16670	"sxl-retry" on page 28764
"sxl-timeout" on page 28765		

[edit security utm feature-profile anti-virus sophos-engine pattern-update]		
"email-notify" on page 6247	"interval (Security Antivirus)" on page 9219	"no-autoupdate" on page 11413
"proxy (Security Antivirus)" on page 13839	"proxy-profile" on page 13838	"routing-instance (Security Content Security)" on page 15958
"url (Security Antivirus)" on page 30504		

[edit security utm feature-profile anti-virus sophos-engine pattern-update email-notify]		
"admin-email" on page 1963	"custom-message (Security Email Notify)" on page 4759	"custom-message-subject (Security Email Notify)" on page 4766

[edit security utm feature-profile anti-virus sophos-engine pattern-update proxy]		
"password (Security Antivirus)" on page 12245	"port (Security Antivirus)" on page 12875	"server (Security Antivirus)" on page 16665
"username (Security Antivirus)" on page 30606		

[edit security utm feature-profile anti-virus sophos-engine profile {profile-name}]		
"fallback-options (Security Antivirus Sophos Engine)" on page 6802	"notification-options (Security Antivirus)" on page 11693	"trickling" on page 30152

[edit security utm feature-profile anti-virus sophos-engine profile {profile-name} fallback-options]		
"content-size (Security Antivirus Sophos Engine)" on page 4604	"default (Security Antivirus Sophos Engine)" on page 5214	"engine-not-ready (Security Antivirus Sophos Engine)" on page 6351
"out-of-resources (Security Antivirus Sophos Engine)" on page 11965	"timeout (Security Antivirus Fallback Options Sophos Engine)" on page 29385	"too-many-requests (Security Antivirus Fallback Options Sophos Engine)" on page 29432

[edit security utm feature-profile anti-virus sophos-engine profile {profile-name} notification-options]		
"fallback-block (Security Antivirus)" on page 6792	"fallback-non-block (Security Antivirus)" on page 6795	"virus-detection (Security Antivirus)" on page 30778

[edit security utm feature-profile anti-virus sophos-engine profile {profile-name} notification-options fallback-block]		
"administrator-email (Security Fallback Block)" on page 1980	"allow-email (Security Fallback Block)" on page 2175	"custom-message (Security Fallback Block)" on page 4761
"custom-message-subject (Security Fallback Block)" on page 4767	"display-host (Security Fallback Block)" on page 5826	"no-notify-mail-sender (Security Fallback Block)" on page 11530
"notify-mail-sender (Security Fallback Block)" on page 11708	"type (Security Fallback Block)" on page 30285	

[edit security utm feature-profile anti-virus sophos-engine profile {profile-name} notification-options fallback-non-block]		
"custom-message (Security Fallback Non-Block)" on page 4762	"custom-message-subject (Security Fallback Non-Block)" on page 4769	"no-notify-mail-recipient" on page 11528
"notify-mail-recipient" on page 11705		

[edit security utm feature-profile anti-virus sophos-engine profile {profile name} notification-options virus-detection]		
"administrator-email (Security Virus Detection)" on page 1981	"custom-message (Security Virus Detection)" on page 4763	"custom-message-subject (Security Virus Detection)" on page 4770
"no-notify-mail-sender (Security Virus Detection)" on page 11531	"notify-mail-sender (Security Virus Detection)" on page 11709	"type (Security Virus Detection)" on page 30289

[edit security utm feature-profile anti-virus sophos-engine server]	
"routing-instance (Security Content Security)" on page 15958	

[edit security utm feature-profile content-filtering]		
"profile (Security Content Filtering)" on page 13317	"traceoptions (Security Content Filtering)" on page 29781	

[edit security utm feature-profile content-filtering profile {profile-name}]		
"block-command" on page 2985	"block-content-type" on page 2986	"block-extension" on page 2987
"block-mime" on page 2994	"notification-options (Security Content Filtering)" on page 11695	"permit-command" on page 12443

[edit security utm feature-profile content-filtering profile {profile-name} block-mime]		
"exception (Security Content Filtering)" on page 6573	"list (Security Content Filtering Block Mime)" on page 9841	

[edit security utm feature-profile content-filtering profile {profile-name} notification-options]		
"custom-message (Security Content Filtering)" on page 4758	"no-notify-mail-sender (Security Content Filtering Notification Options)" on page 11529	"notify-mail-sender (Security Content Filtering Notification Options)" on page 11707
"type (Security Content Filtering Notification Options)" on page 30284		

[edit security utm feature-profile web-filtering]		
"http-persist" on page 8196	"http-reassemble" on page 8201	"juniper-enhanced" on page 9442
"juniper-local" on page 9447	"surf-control-integrated" on page 28733	"traceoptions (Security Web Filtering)" on page 29835
"type (Security Web Filtering)" on page 30291	"url-blacklist" on page 30488	"url-whitelist " on page 30507

(Continued)

[edit security utm feature-profile web-filtering]		
"websense-redirect" on page 31093		

[edit security utm feature-profile web-filtering juniper-enhanced]		
"cache" on page 3167	"profile (Security Web Filtering Juniper Enhanced)" on page 13328	"server (Security Web Filtering)" on page 16672

[edit security utm feature-profile web-filtering juniper-enhanced cache]		
"size (Security Web Filtering Cache)" on page 28004	"timeout (Security Web Filtering Cache)" on page 29392	

[edit security utm feature-profile web-filtering juniper-enhanced profile {profile-name}]		
"block-message (Security Content Security)" on page 2993	"category (Security Web Filtering)" on page 3226	"custom-block-message" on page 4750
"custom-message (Security Web Filtering)" on page 4764	"default (Security Content Security)" on page 5216	"fallback-settings (Security Web Filtering)" on page 6804
"quarantine-message (Security Content Security)" on page 13924	"timeout (Security Web Filtering)" on page 29391	

[edit security utm feature-profile web-filtering juniper-enhanced profile {profile-name} category {customurl-last-name}]		
"action (Security Content Security Web Filtering)" on page 1783	"site-reputation-action" on page 27992	

[edit security utm feature-profile web-filtering juniper-enhanced profile {profile-name} fallback-settings]		
"default (Security Web Filtering)" on page 5217	"server-connectivity" on page 16621	"timeout (Security Web Filtering Fallback Settings)" on page 29393
"too-many-requests (Security Web Filtering Fallback Settings)" on page 29433		

[edit security utm feature-profile web-filtering juniper-enhanced profile {profile-name} server]		
"host (Security Web Filtering)" on page 8167	"port (Security Web Filtering Server)" on page 12879	

[edit security utm feature-profile web-filtering juniper-enhanced server]	
"routing-instance (Security Content Security)" on page 15958	

[edit security utm feature-profile web-filtering juniper-local]	
"profile (Security Web Filtering Juniper Local)" on page 13330	

[edit security utm feature-profile web-filtering juniper-local profile {profile-name}]	
"custom-message (Security Web Filtering)" on page 4764	

[edit security utm feature-profile web-filtering juniper–local profile {profile-name}]	
"fallback-settings (Security Web Filtering Juniper Local)" on page 6805	

[edit security utm feature-profile web-filtering juniper–local profile {profile-name} fallback-settings]		
"default (Security Web Filtering)" on page 5217		

[edit security utm feature-profile web-filtering ng-juniper profile {name}]		
"ng-juniper" on page 11366		

[edit security utm feature-profile web-filtering surf-control-integrated]		
"cache" on page 3167	"profile (Security Web Filtering Surf Control Integrated)" on page 13332	"server (Security Web Filtering)" on page 16672

[edit security utm feature-profile web-filtering surf-control-integrated cache]		
"size (Security Web Filtering Cache)" on page 28004	"timeout (Security Web Filtering Cache)" on page 29392	

[edit security utm feature-profile web-filtering surf-control-integrated profile {profile-name}]		
"category (Security Web Filtering)" on page 3226	"custom-block-message" on page 4750	"fallback-settings (Security Web Filtering)" on page 6804

[edit security utm feature-profile web-filtering surf-control-integrated profile {profile-name} category {customurl-last-name}]		
"action (Security Content Security Web Filtering)" on page 1783		

[edit security utm feature-profile web-filtering surf-control-integrated profile {profile-name} fallback-settings]		
"default (Security Web Filtering)" on page 5217	"server-connectivity" on page 16621	"timeout (Security Web Filtering Fallback Settings)" on page 29393

(Continued)

[edit security utm feature-profile web-filtering surf-control-integrated profile {profile-name} fallback-settings]		
"too-many-requests (Security Web Filtering Fallback Settings)" on page 29433		

[edit security utm feature-profile web-filtering surf-control-integrated server]		
"host (Security Web Filtering)" on page 8167	"port (Security Web Filtering Server)" on page 12879	

[edit security utm feature-profile web-filtering websense-redirect]		
"profile (Security Web Filtering Websense Redirect)" on page 13333		

[edit security utm feature-profile web-filtering websense-redirect profile {profile-name}]		
"custom-message (Security Web Filtering)" on page 4764	"fallback-settings (Security Web Filtering Websense Redirect)" on page 6806	"server (Security Web Filtering)" on page 16672
"sockets" on page 28082	"timeout (Security Web Filtering)" on page 29391	

[edit security utm feature-profile web-filtering websense-redirect profile {profile-name} fallback-settings]		
"default (Security Web Filtering)" on page 5217	"server-connectivity" on page 16621	"timeout (Security Web Filtering Fallback Settings)" on page 29393
"too-many-requests (Security Web Filtering Fallback Settings)" on page 29433		

[edit security utm feature-profile web-filtering websense-redirect profile {profile-name} server]		
"host (Security Web Filtering)" on page 8167	"port (Security Web Filtering Server)" on page 12879	"source-address (Content Security)" on page 28121

[edit security utm feature-profile web-filtering websense-redirect profile wr server]	
"routing-instance (Security Content Security)" on page 15958	

[edit security utm utm-policy {policy-name}]		
"anti-spam (Security Content Security Policy)" on page 2265	"anti-virus (Content Security Policy)" on page 2273	"content-filtering (Security Content Security Policy)" on page 4583
"traffic-options" on page 30040	"web-filtering (Security Content Security Policy)" on page 31076	

[edit security utm utm-policy {policy-name} anti-spam]	
"smtp-profile (Security Content Security Policy Antispam)" on page 28051	

[edit security utm utm-policy {policy-name} anti-virus]		
"ftp (Content Security Policy Anti-Virus)" on page 7621	"http-profile (Security Antivirus)" on page 8198	"imap-profile (Security Content Security Policy Antivirus)" on page 8450
"pop3-profile (Security Content Security Policy Antivirus)" on page 12808	"smtp-profile (Security Content Security Policy Antivirus)" on page 28052	

[edit security utm utm-policy {policy-name} anti-virus ftp]		
"download-profile (Security Antivirus FTP)" on page 5926	"upload-profile (Security Antivirus FTP)" on page 30472	

[edit security utm utm-policy {policy-name} content-filtering]		
"ftp (Content Security Policy Content Filtering)" on page 7622	"http-profile (Security Content Filtering)" on page 8199	"imap-profile (Security Content Security Policy Content Filtering)" on page 8451
"pop3-profile (Security Content Security Policy Content Filtering)" on page 12809	"smtp-profile (Security Content Security Policy Content Filtering)" on page 28053	

[edit security utm utm-policy {policy-name} content-filtering ftp]		
"download-profile (Security Content Filtering FTP)" on page 5927	"upload-profile (Security Content Filtering FTP)" on page 30473	

[edit security utm utm-policy {policy-name} traffic-options sessions-per-client]		
"limit (Content Security Policy)" on page 9758	"over-limit" on page 12031	

[edit security utm utm-policy {policy-name} web-filtering]	
"http-profile (Security Web Filtering)" on page 8200	

[zones]

[edit security zones]		
"functional-zone" on page 7638	"security-zone" on page 16516	

[edit security zones functional-zone]	
"management (Security Zones)" on page 10311	

[edit security zones functional-zone management]		
"description (Security Zone)" on page 5357	"host-inbound-traffic" on page 8142	"interfaces (Security Zones)" on page 9165
"screen (Security Zones)" on page 16390		

[edit security zones functional-zone management interfaces {interface-name}]	
"host-inbound-traffic" on page 8142	

[edit security zones security-zone {zone-name}]		
"application-tracking (Security Zones)" on page 2399	"description (Security Zone)" on page 5357	"host-inbound-traffic" on page 8142
"host-inbound-traffic" on page 8142	"interfaces (Security Zones)" on page 9165	"screen (Security Zones)" on page 16390
"source-identity-log (Security)" on page 28218	"tcp-rst" on page 29012	"unidirectional-session-refreshing" on page 30358

(Continued)

[edit security zones security-zone {zone-name}]		
"white-list (Security-Zone)" on page 31109		

[edit security zones security-zone {zone-name} host-inbound-traffic]		
"protocols (Security Zones Host Inbound Traffic)" on page 13798	"system-services (Security Zones Host Inbound Traffic)" on page 28859	

[edit security zones security-zone {zone} host-inbound-traffic system-services]		
"webapi-clear-text (Security)" on page 31089	"webapi-ssl (Security)" on page 31090	

[edit security zones security-zone {zone-name} interfaces {interface-name} host-inbound-traffic]		
"protocols (Security Zones Interfaces)" on page 13800	"system-services (Security Zones Interfaces)" on page 28862	

[edit service-set]

IN THIS SECTION

- [service-set-options] | **1046**

Sub Hierarchies

"[edit service-set service-set-options]" on page 1046

[service-set-options]

[edit service-set {service-set-name} service-set-options tcp-session]	
"tcp-tickles (Service Set Next Gen Services)" on page 29030	

[edit services]

IN THIS SECTION

- [(CoS)] | **1053**
- [(IPsec VPN)] | **1053**
- [(Stateful Firewall)] | **1054**
- [aacl] | **1054**
- [adaptive-services-pics] | **1056**
- [address-book] | **1056**
- [advanced-anti-malware] | **1056**
- [alg] | **1056**
- [analytics] | **1059**
- [app-engine] | **1061**
- [application-identification] | **1062**
- [captive-portal] | **1064**
- [captive-portal-content-delivery] | **1065**
- [cos] | **1065**

- [dynamic-flow-capture] | **1068**
- [evpn] | **1069**
- [flow-collector] | **1070**
- [flow-monitoring] | **1071**
- [flow-tap] | **1073**
- [hcm] | **1074**
- [hosted-services] | **1075**
- [icap-redirect] | **1076**
- [ids] | **1076**
- [ike] | **1079**
- [inline-monitoring] | **1079**
- [ip-reassembly] | **1080**
- [ip-reassemblyrule] | **1080**
- [ipsec-vpn] | **1081**
- [jdaf] | **1084**
- [jflow-log] | **1085**
- [l2tp] | **1085**
- [logging] | **1089**
- [lrf] | **1089**
- [mobile-edge] | **1091**
- [monitoring] | **1092**
- [nat] | **1093**
- [network-monitoring] | **1100**
- [paa] | **1100**
- [pcef] | **1101**
- [pcp] | **1103**
- [pfe] | **1104**
- [plmn] | **1104**
- [policies] | **1105**
- [proxy] | **1105**
- [radius-flow-tap] | **1105**
- [radius-flow-tappolicy{] | **1107**

- [redundancy-set] | **1107**
- [rpm] | **1107**
- [rtlog] | **1111**
- [screen] | **1112**
- [security-metadata-streaming] | **1114**
- [service-device-pools] | **1114**
- [service-interface-pools] | **1114**
- [service-set] | **1115**
- [service-set-options] | **1120**
- [services-set] | **1120**
- [softwire] | **1121**
- [softwires] | **1121**
- [source] | **1122**
- [ssl] | **1122**
- [stateful-firewall] | **1124**
- [traffic-load-balance] | **1125**
- [unified-access-control] | **1126**
- [url-filter] | **1127**
- [user-identification] | **1127**
- [video-monitoring] | **1129**
- [web-filter] | **1129**
- [web-proxy] | **1131**

Sub Hierarchies

"[edit services (cos)]" on page 1053

"[edit services (ipsec vpn)]" on page 1053

"[edit services (stateful firewall)]" on page 1054

"[edit services aacl]" on page 1054

"[edit services adaptive-services-pics]" on page 1056

"[edit services address-book]" on page 1056

"[edit services advanced-anti-malware]" on page 1056

"[edit services alg]" on page 1056

"[edit services analytics]" on page 1059

"[edit services app-engine]" on page 1061

"[edit services application-identification]" on page 1062

"[edit services captive-portal]" on page 1064

"[edit services captive-portal-content-delivery]" on page 1065

"[edit services cos]" on page 1065

"[edit services dynamic-flow-capture]" on page 1068

"[edit services evpn]" on page 1069

"[edit services flow-collector]" on page 1070

"[edit services flow-monitoring]" on page 1071

"[edit services flow-tap]" on page 1073

"[edit services hcm]" on page 1074

"[edit services hosted-services]" on page 1075

"[edit services icap-redirect]" on page 1076

"[edit services ids]" on page 1076

"[edit services ike]" on page 1079

"[edit services inline-monitoring]" on page 1079

"[edit services ip-reassembly]" on page 1080

"[edit services ip-reassemblyrule{rule-name}]" on page 1080

"[edit services ipsec-vpn]" on page 1081

"[edit services jdaf]" on page 1084

"[edit services jflow-log]" on page 1085

"[edit services l2tp]" on page 1085

"[edit services logging]" on page 1089

"[edit services lrf]" on page 1089

"[edit services mobile-edge]" on page 1091

"[edit services monitoring]" on page 1092

"[edit services nat {destination pool}]" on page 1093

"[edit services network-monitoring]" on page 1100

"[edit services paa]" on page 1100

"[edit services pcef]" on page 1101

"[edit services pcp]" on page 1103

"[edit services pfe]" on page 1104

"[edit services {name} plmn]" on page 1104

"[edit services policies]" on page 1105

"[edit services proxy]" on page 1105

"[edit services radius-flow-tap]" on page 1105

"[edit services radius-flow-tappolicy{]" on page 1107

"[edit services redundancy-set]" on page 1107

"[edit services rpm]" on page 1107

"[edit services rtlog]" on page 1111

"[edit services screen]" on page 1112

"[edit services security-metadata-streaming]" on page 1114

"[edit services service-device-pools]" on page 1114

"[edit services service-interface-pools]" on page 1114

"[edit services service-set {service-set-name}]" on page 1115

"[edit services {service-set-name} service-set-options]" on page 1120

"[edit services services-set]" on page 1120

"[edit services softwire]" on page 1121

"[edit services softwires]" on page 1121

"[edit services source]" on page 1122

"[edit services ssl]" on page 1122

"[edit services stateful-firewall]" on page 1124

"[edit services traffic-load-balance]" on page 1125

"[edit services unified-access-control]" on page 1126

"[edit services url-filter]" on page 1127

"[edit services user-identification]" on page 1127

"[edit services video-monitoring]" on page 1129

"[edit services web-filter]" on page 1129

"[edit services web-proxy]" on page 1131

[edit services]		
"access-line-information (L2TP)" on page 1684	"adaptive-services-pics" on page 1834	"agf" on page 2053
"alarm-mode" on page 2108	"alarms (Video Monitoring)" on page 2125	"alg" on page 2130
"amf" on page 2226	"analytics" on page 2231	"anti-virus (Services)" on page 2275
"application-identification (Application Identification)" on page 2337	"application-identification (Services)" on page 2330	"captive-portal" on page 3199
"delay-factor" on page 5248	"dns-filtering" on page 5859	"dynamic-flow-capture" on page 6109
"evpn (Easy EVPN LAG Configuration)" on page 6548	"fiveqi-map" on page 7199	"flow-collector (Flow Collection)" on page 7263
"flow-monitoring" on page 7323	"flow-monitoring (Services)" on page 7328	"flow-tap" on page 7355

(Continued)

[edit services]		
"generate-snmp-traps" on page 7676	"hcm (HTTP Header Enrichment)" on page 8005	"hosted-services" on page 8178
"icap-redirect" on page 8245	"inband-flow-telemetry" on page 8512	"inline-monitoring" on page 8672
"ip-monitoring (Services)" on page 9246	"ip-reassembly" on page 9262	"jdaf" on page 9422
"jflow-log (Services)" on page 9427	"l2tp" on page 9541	"local-endpoint" on page 9939
"logging (Services)" on page 10056	"logical-domain-identity-management" on page 10059	"media-loss-rate" on page 10596
"media-rate-variation" on page 10597	"monitoring (Services)" on page 10961	"network-slicing" on page 11323
"no-syslog-generation" on page 11591	"plmn" on page 12615	"profile (icap-redirect)" on page 13283
"radius-flow-tap" on page 14006	"redundancy-set" on page 14204	"rtlog (Next Gen Services Global System Logging)" on page 16066
"security-intelligence (Services)" on page 16462	"security-intelligence(services)" on page 16469	"security-metadata-streaming" on page 16487
"service-device-pools (Service Interfaces)" on page 16688	"service-interface-pools" on page 16708	"service-set (Application-Aware Control Policy)" on page 16740
"service-set (Services)" on page 16750	"service-set (Subscriber-Aware)" on page 16754	"soft-gre" on page 28083
"softwires" on page 28102	"softwires (Next Gen Services)" on page 28104	"ssl (Services)" on page 28447

(Continued)

[edit services]		
"storm-control (Video Monitoring)" on page 28619	"trace (EVPN-VXLAN)" on page 29459	"tracking-area" on page 29994
"traffic-load-balance (Traffic Load Balancer)" on page 30034	"unified-access-control (Services)" on page 30359	"upstream-rewrite" on page 30481
"url-filter" on page 30492	"user-identification (Services)" on page 30560	"user-planes" on page 30575
"video-monitoring" on page 30729	"web-filter" on page 31062	

[(CoS)]

[edit services (CoS) cos application-profile {profile-name} sip]		
"video" on page 30727	"voice" on page 30915	

[(IPsec VPN)]

[edit services (IPsec VPN) ipsec-vpn ike proposal {proposal-name}]	
"authentication-method (services)" on page 2626	

[edit services (IPsec VPN) ipsec-vpn rule {rule-name} term {term-name} then manual direction {direction}]		
"authentication (services)" on page 2572	"auxiliary-spi" on page 2773	

[(Stateful Firewall)]

[edit services (Stateful Firewall) stateful-firewall rule (Services Stateful Firewall) {rule-name} term (Services Stateful Firewall) {term-name} from (Services Stateful Firewall)]	
"applications (Services Stateful Firewall)" on page 2420	

[edit services (Stateful Firewall) stateful-firewall rule (Services Stateful Firewall) {rule-name} term (Services Stateful Firewall) {term-name} then (Services Stateful Firewall)]	
"allow-ip-options (Services Stateful Firewall)" on page 2186	

[aacl]

[edit services aacl]		
"rule (AACL Rule Set) " on page 16076	"rule-set (Services AACL)" on page 16138	

[edit services aacl rule {rule-name} {term-name}]		
"application-group-any (AACL)" on page 2323	"application-unknown" on page 2404	"log (aacl)" on page 10005

(Continued)

[edit services aacl rule {rule-name} {term-name}]		
"match-direction" on page 10371	"nested-application-unknown" on page 11287	"nested-applications" on page 11288
"term (Application Aware Access List)" on page 29106		

[edit services aacl rule {rule-name} term {term-name}]		
"from (Application Aware Access List)" on page 7582		

[edit services aacl rule {rule-name} term {term-name} from]		
"application-group-any (AACL)" on page 2323	"application-groups (Services AACL)" on page 2328	"applications (Services AACL)" on page 2414
"destination-address" on page 5368	"destination-address-range" on page 5386	"destination-prefix-list (Services AACL)" on page 5500
"source-address (AACL)" on page 28123	"source-address-range" on page 28149	"source-prefix-list (Services AACL)" on page 28296

[edit services aacl rule-set {rule-set-name}]		
"rule (AACL Rule Set) " on page 16076		

[adaptive-services-pics]

[edit services adaptive-services-pics]	
"traceoptions (Services Logging)" on page 29849	

[address-book]

[edit services address-book global]	
"address (Address Book Next Gen Services)" on page 1852	

[advanced-anti-malware]

[edit services advanced-anti-malware]	
"advanced-anti-malware policy" on page 1997	

[alg]

[edit services alg]		
"alg-support-lib" on page 2142	"dns" on page 5847	"ftp (Security ALG)" on page 7626

(Continued)

[edit services alg]		
"ike-esp-nat" on page 8421	"mgcp" on page 10693	"msrpc" on page 10999
"pptp" on page 12986	"rsh" on page 16041	"rtsp" on page 16072
"sccp" on page 16341	"sunrpc" on page 28722	"talk" on page 28940
"tftp (Security ALG)" on page 29252	"traceoptions (Security ALG)" on page 29770	

[edit services alg dns]		
"maximum-message-length" on page 10534	"oversize-message-drop" on page 12081	

[edit services alg h323]		
"application-screen (Security H323)" on page 2359	"endpoint-registration-timeout" on page 6342	"media-source-port-any" on page 10599
"traceoptions (Security H323 ALG)" on page 29801		

[edit services alg h323 application-screen]		
"message-flood (Security H323)" on page 10644	"unknown-message (Security H323 ALG)" on page 30410	

[edit services alg h323 application-screen message-flood]	
"gatekeeper" on page 7653	

[edit services alg mgcp]		
"application-screen (Security MGCP)" on page 2361	"inactive-media-timeout (Security MGCP)" on page 8496	"traceoptions (Security MGCP ALG)" on page 29814
"transaction-timeout" on page 30055		

[edit services alg mgcp application-screen]		
"connection-flood" on page 4531	"message-flood (Security MGCP)" on page 10645	"unknown-message (Security MGCP ALG)" on page 30411

[edit services alg sccp]		
"application-screen (Security SCCP)" on page 2362	"inactive-media-timeout (Security SCCP)" on page 8497	"traceoptions (Security SCCP ALG)" on page 29823

[edit services alg sccp application-screen]		
"call-flood" on page 3179	"unknown-message (Security SCCP ALG)" on page 30413	

[edit services alg sip]		
"application-screen (Security SIP)" on page 2364	"c-timeout" on page 3152	"t1-interval" on page 28885
"t4-interval" on page 28893	"traceoptions (Security SIP ALG)" on page 29830	

[edit services alg sip (Security)]		
"inactive-media-timeout (Security SIP)" on page 8499	"retain-hold-resource" on page 15674	

[edit services alg sip (Security) application-screen (Security SIP)]	
"unknown-message (Security SIP ALG)" on page 30415	

[edit services alg sip application-screen protect]	
"deny (Security SIP)" on page 5317	

[analytics]

[edit services analytics]		
"agent (Analytics)" on page 2046	"collector (Analytics)" on page 4412	"export-profile (Junos Telemetry Interface)" on page 6669
"export-profiles" on page 6674	"interfaces (Analytics)" on page 9117	"queue-statistics" on page 13950
"resource (Analytics)" on page 15604	"resource-profiles (Analytics)" on page 15620	"sensor (Junos Telemetry Interface)" on page 16564
"streaming-server (Junos Telemetry Interface)" on page 28640	"streaming-servers" on page 28642	"traceoptions (Analytics)" on page 29491
"traffic-statistics (Services Analytics)" on page 30045		

[edit services analytics agent]		
"service-agents (Analytics)" on page 16681	"traceoptions (Analytics Agent)" on page 29492	

[edit services analytics agent service-agents {agent-name}]		
"inputs (Analytics)" on page 8734	"outputs (Analytics)" on page 12025	

[edit services analytics collector]		
"address (Analytics Collector)" on page 1855	"local (Analytics Collector)" on page 9914	

[edit services analytics collector local]	
"file (Analytics)" on page 7028	

[edit services analytics export-profile {export-profile-name}]		
"local-address (Junos Telemetry Interface)" on page 9902	"routing-instance (Junos Telemetry Interface)" on page 15937	

[edit services analytics export-profiles]		
"interface (Export Profiles)" on page 8884	"system (Export Profiles)" on page 28844	

[edit services analytics interfaces]		
"depth-threshold" on page 5319	"latency-threshold" on page 9639	

[edit services analytics queue-statistics]	
"file (Analytics)" on page 7028	

[edit services analytics resource]		
"interfaces (Analytics Resource)" on page 9119	"system (Analytics Resource)" on page 28840	

[edit services analytics resource-profiles]		
"depth-threshold" on page 5319	"latency-threshold" on page 9639	

[edit services analytics traffic-statistics]	
"file (Analytics)" on page 7028	

[edit services analytics zero-suppression]	
"no-zero-suppression" on page 11626	

[app-engine]

[edit services app-engine compute-cluster {compute-cluster-name} local-management]	
"routing-instance (Services)" on page 15915	

[edit services app-engine virtual-machines instance {instance-name} local-management]	
"routing-instance (Services)" on page 15915	

[application-identification]

[edit services application-identification]		
"application (Application Identification)" on page 2306	"application (Application Identification)" on page 2303	"application-group" on page 2322
"application-group (Services)" on page 2325	"application-system-cache-timeout (Services)" on page 2395	"download (Services)" on page 5928
"enable-performance-mode" on page 6269	"enable-performance-mode (Application Identification)" on page 6271	"maximum-transactions" on page 10570
"no-application-identification (Services)" on page 11397	"no-application-system-cache" on page 11399	"no-application-system-cache (Services)" on page 11400
"packet-capture" on page 12137	"profile (Application Identification)" on page 13254	"profile (Services Application Identification)" on page 13335
"statistics (Services)" on page 28599	"traceoptions (Services Application Identification)" on page 29839	

[edit services application-identification application {application-name}]		
"alt-name (Application Identification)" on page 2217	"cacheable (Application Identification)" on page 3173	"compatibility (Application Identification)" on page 4480
"description (Application Identification)" on page 5325	"icmp-mapping (Application Identification)" on page 8263	"ip-protocol-mapping (Application Identification)" on page 9258
"order (Application Identification)" on page 11927	"over (Application Identification)" on page 12028	"priority (Application Identification With Next Gen Services)" on page 13156
"tags (Application Identification)" on page 28938	"type (Application Identification)" on page 30275	

[edit services application-identification application {application-name} address-mapping]		
"destination (Application Identification)" on page 5408	"order (Application Identification)" on page 11927	"order-priority (Application Identification)" on page 11929
"source (Application Identification)" on page 28182		

[edit services application-identification application {application-name} icmp-mapping]		
"code (Application Identification)" on page 4389	"order (Application Identification)" on page 11927	"order-priority (Application Identification)" on page 11929
"type (ICMP Mapping for Application Identification)" on page 30276		

[edit services application-identification application {application-name} ip-protocol-mapping]		
"order (Application Identification)" on page 11927	"order-priority (Application Identification)" on page 11929	

[edit services application-identification application {application-name} over {protocol-type}]		
"signature" on page 27950	"signature (Application Identification)" on page 27952	

[edit services application-identification application {application-name} over {protocol-type} signature {l4-l7-signature-name}]		
"chain-order (Application Identification)" on page 3276	"member (Application Identification)" on page 10606	"order (Application Identification)" on page 11927
"order-priority (Application Identification)" on page 11929	"port-range (Application Identification)" on page 12865	"protocol (Application Identification)" on page 13429

[edit services application-identification application {application-name} over {protocol-type} signature {l4-l7-signature-name} member {member-name}]		
"check-bytes (Application Identification)" on page 3308	"pattern (Application Identification)" on page 12271	

[edit services application-identification application-group {group-name}]		
"applications (Services Application Identification)" on page 2416	"tag-group" on page 28911	

[edit services application-identification download]	
"ignore-server-validation (Application Identification)" on page 5924	

[edit services application-identification nested-application-settings]	
"no-application-system-cache" on page 11399	

[captive-portal]

[edit services captive-portal]		
"custom-options" on page 4773	"interface (Captive Portal)" on page 8846	

[captive-portal-content-delivery]

[edit services captive-portal-content-delivery]		
"profile (Captive Portal Content Delivery)" on page 13273	"rule (Captive Portal Content Delivery)" on page 16085	"rule-set (Captive Portal Content Delivery)" on page 16129
"traceoptions (Captive Portal Content Delivery)" on page 29501		

[edit services captive-portal-content-delivery rule {rule-name}]	
"term (Captive Portal Content Delivery)" on page 29113	

[edit services captive-portal-content-delivery rule {rule-name} term {term-name}]		
"from (Captive Portal Content Delivery Tags)" on page 7587	"then (Captive Portal Content Delivery)" on page 29259	

[cos]

[edit services cos]		
"application-profile" on page 2349	"application-profile (Services CoS Next Gen Services)" on page 2351	"rule (Services CoS Next Gen Services)" on page 16117
"rule (Services CoS) " on page 16115	"rule-set (Services CoS Next Gen Services)" on page 16140	"rule-set (Services CoS)" on page 16139

[edit services cos application-profile {profile-name}]		
"ftp (Services CoS Next Gen Services)" on page 7629	"sip" on page 27971	"sip (Application Profile)" on page 27973
"sip (Services CoS Next Gen Services)" on page 27981		

[edit services cos application-profile {profile-name} (ftp | sip) (data | video | voice)]		
"dscp (Services CoS)" on page 6022	"forwarding-class (Services PIC Classifiers)" on page 7442	

[edit services cos application-profile {profile-name} ftp]		
"data (FTP)" on page 4837	"ftp (Services CoS)" on page 7628	

[edit services cos application-profile{profile-name} sip]		
"video (Application Profile)" on page 30728	"voice (Application Profile)" on page 30917	

[edit services cos rule {rule-name}]		
"match-direction (Services CoS) " on page 10378	"policy (Services CoS Next Gen Services)" on page 12728	"term (Services CoS)" on page 29132

[edit services cos rule {rule-name} policy {policy-name}]		
"match (Services CoS Next Gen Services)" on page 10404	"then (Services CoS Next Gen Services)" on page 29296	

[edit services cos rule {rule-name} term {term-name}]		
"from (Services CoS)" on page 7597	"then (Services CoS) " on page 29295	

[edit services cos rule {rule-name} term {term-name} from]		
"application-sets (Services CoS)" on page 2386	"applications (Services CoS)" on page 2418	"destination-address (Services CoS)" on page 5403
"destination-prefix-list (Services CoS)" on page 5501	"source-address (Services CoS)" on page 28169	"source-prefix-list (Services CoS) " on page 28297

[edit services cos rule {rule-name} term {term-name} then]		
"application-profile" on page 2349	"dscp (Services CoS)" on page 6022	"forwarding-class (Services PIC Classifiers)" on page 7442
"reflexive | revert | reverse" on page 14233	"syslog (Services CoS)" on page 28827	

[edit services cos rule {rule-name} term {term-name} then reflexive | revert | reverse]	
"forwarding-class (Services PIC Classifiers)" on page 7442	

[edit services cos rule {rule-name} term {term-name} then reverse]		
"application-profile" on page 2349	"dscp (Services CoS)" on page 6022	"syslog (Services CoS)" on page 28827

[edit services cos rule-set {rule-set-name}]	
"rule (Services CoS) " on page 16115	

[dynamic-flow-capture]

[edit services dynamic-flow-capture]		
"capture-group" on page 3208	"g-duplicates-dropped-periodicity" on page 7645	"g-max-duplicates" on page 7646
"traceoptions (Dynamic Flow Capture)" on page 29538		

[edit services dynamic-flow-capture capture-group {client-name}]		
"content-destination" on page 4579	"control-source" on page 4641	"duplicates-dropped-periodicity" on page 6072
"input-packet-rate-threshold" on page 8711	"interfaces (Services Dynamic Flow Capture)" on page 9167	"max-duplicates" on page 10436
"pic-memory-threshold" on page 12489		

[edit services dynamic-flow-capture capture-group {client-name} content-destination {identifier}]		
"address (Services Dynamic Flow Capture)" on page 1928	"hard-limit" on page 7969	"hard-limit-target" on page 7970
"soft-limit" on page 28085	"soft-limit-clear" on page 28087	"ttl (Dynamic Flow Capture)" on page 30179

[edit services dynamic-flow-capture capture-group {client-name} control-source {identifier}]		
"allowed-destinations" on page 2212	"minimum-priority" on page 10747	"no-syslog" on page 11590
"notification-targets" on page 11699	"service-port (UDP, Dynamic Flow Capture)" on page 16724	"shared-key" on page 16968

(Continued)

[edit services dynamic-flow-capture capture-group {client-name} control-source {identifier}]	
"source-addresses" on page 28181	

[evpn]

[edit services evpn]		
"device-attribute (Easy EVPN LAG Configuration)" on page 5560	"evpn-vxlan (Easy EVPN LAG Configuration)" on page 6559	"global-parameters (Easy EVPN LAG Configuration)" on page 7728

[edit services evpn device-attribute]	
"peer-to-peer (Easy EVPN LAG Configuration)" on page 12388	

[edit services evpn device-attribute peer-to-peer]	
"peer-subnet (Easy EVPN LAG Configuration)" on page 12385	

[edit services evpn evpn-vxlan]		
"dhcp-relay (Easy EVPN LAG Configuration)" on page 5638	"irb (Easy EVPN LAG Configuration)" on page 9386	"server (Easy EVPN LAG Configuration)" on page 16628

[edit services evpn evpn-vxlan irb {irb-instance}]	
"subnet-address (Easy EVPN LAG Configuration)" on page 28675	

[flow-collector]

[edit services flow-collector]		
"analyzer-address" on page 2241	"analyzer-id" on page 2243	"destinations" on page 5532
"file-specification (File Format)" on page 7082	"interface-map" on page 8915	"retry (Services Flow Collector)" on page 15716
"retry-delay" on page 15703	"transfer-log-archive" on page 30061	

[edit services flow-collector destination]	
"ftp (Flow Collector Files)" on page 7624	

[edit services flow-collector destination ftp:{url}]	
"password (Flow Collector File Servers)" on page 12232	

[edit services flow-collector file-specification]	
"variant" on page 30655	

[edit services flow-collector file-specification variant {variant-number}]		
"data-format" on page 4834	"name-format (Flow Collection)" on page 11170	"transfer" on page 30056

[edit services flow-collector interface-map]		
"collector (Flow Collection)" on page 4410	"file-specification (Interface Mapping)" on page 7083	

[edit services flow-collector transfer-log-archive]		
"archive-sites (Flow Monitoring)" on page 2448	"filename-prefix" on page 7100	"maximum-age" on page 10498

[edit services flow-collector transfer-log-archive archive-sites]		
"ftp (Transfer Log Files)" on page 7631	"password (Transfer Log File Servers)" on page 12249	"username (Services)" on page 30607

[flow-monitoring]

[edit services flow-monitoring]		
"version-ipfix (Services)" on page 30708	"version9 (Flow Monitoring)" on page 30722	

[edit services flow-monitoring (version-ipfix |version9) template {template-name}]	
"nexthop-learning" on page 11363	

[edit services flow-monitoring version-ipfix]	
"template (Flow Monitoring IPFIX Version)" on page 29075	

[edit services flow-monitoring version-ipfix template {template-name}]		
"data-record-fields" on page 4842	"flow-active-timeout (Flow Monitoring)" on page 7261	"flow-active-timeout (Flow Monitoring)" on page 7261
"flow-key (Flow Monitoring)" on page 7304	"ipv4-template" on page 9344	"ipv6-template" on page 9375

(Continued)

[edit services flow-monitoring version-ipfix template {template-name}]		
"mpls-ipv4-template" on page 10979	"mpls-ipvx-template" on page 10981	"mpls-template" on page 10986
"multi-bgp-path" on page 11041	"observation-domain-id" on page 11768	"options-template-id" on page 11924
"template-id" on page 29081	"tunnel-observation" on page 30221	"vpls-template" on page 30937

[edit services flow-monitoring version-ipfix template {template-name}],]		
"flow-inactive-timeout (Flow Monitoring)" on page 7296	"option-refresh-rate" on page 11902	"template-refresh-rate" on page 29094

[edit services flow-monitoring version-ipfix |version9 template {template-name}]	
"bridge-template" on page 3038	

[edit services flow-monitoring version9]	
"template (Flow Monitoring Version 9)" on page 29077	

[edit services flow-monitoring version9 template {template-name}]		
"flow-active-timeout (Flow Monitoring)" on page 7261	"flow-active-timeout (Flow Monitoring)" on page 7261	"flow-inactive-timeout (Flow Monitoring)" on page 7296
"flow-key (Flow Monitoring)" on page 7304	"ipv4-template" on page 9344	"ipv4-template (Services)" on page 9345
"ipv6-template" on page 9375	"ipv6-template (Services)" on page 9376	"mpls-ipv4-template" on page 10979

(Continued)

[edit services flow-monitoring version9 template {template-name}]		
"mpls-ipvx-template" on page 10981	"mpls-template" on page 10986	"multi-bgp-path" on page 11041
"option-refresh-rate" on page 11902	"option-refresh-rate (Services)" on page 11903	"options-template-id" on page 11924
"peer-as-billing-template" on page 12343	"source-id" on page 28214	"template-id" on page 29081
"template-refresh-rate" on page 29094	"template-refresh-rate (Services)" on page 29096	"tunnel-observation" on page 30221

[edit services flow-monitoring version9 template {template-name} mpls-ipv4-template]	
"label-position" on page 9579	

[edit services flow-monitoring version9 template {template-name} mpls-template]	
"label-position" on page 9579	

[flow-tap]

[edit services flow-tap]	
"interface (Services Flow Tap)" on page 9029	

[hcm]

[edit services hcm]		
"profile (HTTP Header Enrichment)" on page 13282	"tag-attribute (HTTP Header Enrichment)" on page 28907	"tag-rule (HTTP Header Enrichment)" on page 28926
"tag-rule-set (HTTP Header Enrichment)" on page 28930		

[edit services hcm profile {profile-name}]	
"tag-rule (Profiles for HTTP Header Enrichment)" on page 28928	

[edit services hcm tag-rule {rule-name}]	
"term (HTTP Header Enrichment)" on page 29120	

[edit services hcm tag-rule {rule-name} term {term-number}]		
"from (HTTP Header Enrichment)" on page 7588	"then (HTTP Header Enrichment)" on page 29267	

[edit services hcm tag-rule {rule-name} term {term-number} from]		
"destination-address (HTTP Header Enrichment) " on page 5375	"destination-address-range (HTTP Header Enrichment)" on page 5387	"destination-port-range (HTTP Header Enrichment)" on page 5476
"destination-ports (HTTP Header Enrichment)" on page 5491	"destination-prefix-list (HTTP Header Enrichment)" on page 5497	

[edit services hcm tag-rule {rule-name} term {term-number} then]		
"count (HTTP Header Enrichment)" on page 4686	"tag (HTTP Header Enrichment)" on page 28914	

[edit services hcm tag-rule {rule-name} term {term-number} then tag {tag-name}]		
"encrypt (HTTP Header Enrichment)" on page 6320	"ipv4-mask (HTTP Header Enrichment)" on page 9335	"ipv4-or-value (HTTP Header Enrichment)" on page 9337
"ipv6-mask (HTTP Header Enrichment)" on page 9363	"ipv6-or-value (HTTP Header Enrichment)" on page 9366	"tag-attribute (HTTP Header Enrichment Tag Rule)" on page 28908
"tag-header (HTTP Header Enrichment)" on page 28913	"tag-operation (HTTP Header Enrichment)" on page 28917	"tag-separator (HTTP Header Enrichment)" on page 28934

[edit services hcm tag-rule then tag]	
"tag-value (HTTP Header Enrichment)" on page 28935	

[edit services hcm tag-rule-set]	
"rule (HTTP Header Enrichment for Tag Rule Set)" on page 16093	

[hosted-services]

[edit services hosted-services]	
"server-profile" on page 16654	

[edit services hosted-services server-profile {server-profile-name}]		
"client-address" on page 4302	"hosted-service-identifier" on page 8176	"server-address (Hosted Services)" on page 16612
"transport-type" on page 30131		

[icap-redirect]

[edit services icap-redirect]		
"traceoptions (Services Application Identification)" on page 29839		

[edit services icap-redirect profile {name}]		
"fallback-option (ICAP Redirect Service)" on page 6797	"http (icap-redirect profile)" on page 8193	"server (icap-redirect profile)" on page 16635

[edit services icap-redirect profile {name} server {name}]		
"authorization (icap-redirect profile)" on page 2678		

[ids]

[edit services ids]		
"rule (IDS MS-DPC)" on page 16094	"rule (IDS MS-MPC)" on page 16096	"rule-set (Services IDS)" on page 16142

[edit services ids rule {rule-name}]		
"match-direction (IDS)" on page 10373	"term (IDS MS-DPC)" on page 29123	"term (IDS MS-MPC)" on page 29125

[edit services ids rule {rule-name} term {term-name}]		
"from (IDS MS-DPC)" on page 7590	"then (IDS MS-DPC)" on page 29269	"then (IDS MS-MPC)" on page 29271

[edit services ids rule {rule-name} term {term-name} from]		
"application-sets (IDS MS-DPC)" on page 2383	"applications (IDS MS-DPC)" on page 2409	"destination-address (IDS MS-DPC)" on page 5377
"destination-address-range (IDS MS-DPC))" on page 5389	"destination-prefix-list (Services IDS)" on page 5502	"source-address (IDS MS-DPC)" on page 28136
"source-address-range (IDS MS-DPC)" on page 28150	"source-prefix-list (Services IDS)" on page 28299	

[edit services ids rule {rule-name} term {term-name} then]		
"aggregation (IDS)" on page 2086	"allow-ip-options (IDS MS-MPC)" on page 2183	"allow-ipv6-extension-header (IDS MS-MPC)" on page 2188
"force-entry (IDS MS-DPC)" on page 7370	"icmp-fragment-check (IDS MS-MPC)" on page 8259	"icmp-large-packet-check (IDS MS-MPC)" on page 8262
"land-attack-check (IDS MS-MPC)" on page 9620	"logging (IDS MS-DPC)" on page 10053	"session-limit (IDS MS-DPC)" on page 16817
"session-limit (IDS MS-MPC)" on page 16819	"syn-cookie (IDS MS-DPC)" on page 28774	"tcp-syn-defense (IDS MS-MPC)" on page 29027
"tcp-syn-fragment-check (IDS MS-MPC)" on page 29028	"tcp-winnuke-check (IDS MS-MPC)" on page 29031	

[edit services ids rule {rule-name} term {term-name} then aggregation]		
"destination-prefix (IDS)" on page 5494	"destination-prefix-ipv6 (IDS)" on page 5495	"source-prefix (IDS)" on page 28292
"source-prefix-ipv6 (IDS)" on page 28293		

[edit services ids rule {rule-name} term {term-name} then logging]		
"syslog (IDS MS-DPC)" on page 28819	"threshold (Services Logging and SYN-Cookie Defenses)" on page 29341	

[edit services ids rule {rule-name} term {term-name} then session-limit]		
"by-destination (IDS MS-DPC)" on page 3092	"by-destination (IDS MS-MPC)" on page 3093	"by-pair (IDS MS-DPC)" on page 3098
"by-source (IDS MS-MPC)" on page 3106		

[edit services ids rule {rule-name} term {term-name} then session-limit (IDS MS-DPC)]	
"by-source (IDS MS-DPC)" on page 3105	

[edit services ids rule {rule-name} term {term-name} then session-limit by-destination]	
"by-protocol (IDS MS-MPC)" on page 3099	

[edit services ids rule {rule-name} term {term-name} then session-limit by-source]	
"by-protocol (IDS MS-MPC)" on page 3099	

[edit services ids rule {rule-name} term {term-name} then syn-cookie]		
"mss (IDS MS-DPC)" on page 11001	"threshold (Services Logging and SYN-Cookie Defenses)" on page 29341	

[edit services ids rule-set {rule-set-name}]	
"rule (IDS MS-DPC)" on page 16094	

[ike]

[edit services ike policy {policy-name}]	
"pre-shared-key (Services IKE)" on page 13006	

[inline-monitoring]

[edit services inline-monitoring]		
"instance (Inline Monitoring)" on page 8750	"template (Inline Monitoring)" on page 29084	"traceoptions (Inline Monitoring)" on page 29585

[edit services inline-monitoring feature-profile]	
"features" on page 6979	

[edit services inline-monitoring instance {instance-name}]		
"collector (Inline Monitoring)" on page 4420	"controller" on page 4650	

[edit services inline-monitoring template {template-name}]		
"flow-monitoring (Inline Monitoring Services)" on page 7326	"primary-data-record-fields" on page 13134	

[ip-reassembly]

[edit services ip-reassembly]	
"rule (IP Reassembly)" on page 16100	

[ip-reassemblyrule]

[edit services ip-reassemblyrule{rule-name}]	
"match-direction (IP Reassembly Rule)" on page 10375	

[ipsec-vpn]

[edit services ipsec-vpn]		
"clear-ike-sas-on-pic-restart" on page 3631	"clear-ipsec-sas-on-pic-restart" on page 3656	"disable-natt (Services IPsec VPN)" on page 5777
"establish-tunnels" on page 6460	"ike" on page 8415	"ipsec" on page 9298
"no-ipsec-tunnel-in-traceroute" on page 11481	"rule" on page 16074	"rule-set" on page 16128
"traceoptions (Security)" on page 29767	"traceoptions (Services IPsec VPN)" on page 29842	

[edit services ipsec-vpn ike]		
"policy (Services IKE)" on page 12730	"proposal (Services IKE)" on page 13399	

[edit services ipsec-vpn ike policy {policy-name}]		
"description (Services IPsec VPN)" on page 5358	"local-certificate" on page 9929	"local-id" on page 9947
"mode (ike)" on page 10797	"peer-certificate-type" on page 12350	"proposals (Services)" on page 13402
"respond-bad-spi (Services IKE Policy)" on page 15626	"version (IKE)" on page 30706	

[edit services ipsec-vpn ike policy {policy-name} certificate]	
"policy-oids" on page 12696	

[edit services ipsec-vpn ike proposal {proposal-name}]		
"authentication-algorithm (Services IKE)" on page 2584	"description (Services IPsec VPN)" on page 5358	"dh-group (services)" on page 5578
"encryption-algorithm" on page 6324	"lifetime-seconds" on page 9749	

[edit services ipsec-vpn ikepolicy {policy-name}]	
"remote-id (Services)" on page 14326	

[edit services ipsec-vpn ikepolicy {policy-name} local-id|remote-id]	
"distinguished-name" on page 5828	

[edit services ipsec-vpn ipsec]		
"policy (IPsec)" on page 12684	"proposal (Services IPsec VPN)" on page 13400	

[edit services ipsec-vpn ipsec policy {policy-name}]		
"description (Services IPsec VPN)" on page 5358	"perfect-forward-secrecy (Services)" on page 12426	"proposals (Services)" on page 13402

[edit services ipsec-vpn ipsec proposal {ipsec-proposal-name}]		
"authentication-algorithm (Services IPsec)" on page 2585	"description (Services IPsec VPN)" on page 5358	"encryption-algorithm" on page 6324
"lifetime-seconds" on page 9749	"protocol (IPsec) " on page 13444	

[edit services ipsec-vpn rule {rule-name}]		
"match-direction (ipsec-vpn)" on page 10372	"term (IP Security VPN)" on page 29107	

[edit services ipsec-vpn rule {rule-name} term {term-name}]		
"from (services-ipsec-vpn)" on page 7584	"then (IP Security VPN)" on page 29255	

[edit services ipsec-vpn rule {rule-name} term {term-name} from]		
"destination-address (IPsec-VPN)" on page 5369	"ipsec-inside-interface" on page 9304	"source-address (Services IPSec VPN)" on page 28115

[edit services ipsec-vpn rule {rule-name} term {term-name} then]		
"anti-replay-window-size (Services IPsec VPN)" on page 2256	"backup-remote-gateway" on page 2811	"clear-dont-fragment-bit (Services IPsec VPN)" on page 3560
"copy-dont-fragment-bit (Services IPsec VPN)" on page 4661	"dynamic (services)" on page 6083	"initiate-dead-peer-detection" on page 8661
"manual" on page 10315	"no-anti-replay" on page 11395	"remote-gateway" on page 14322
"set-dont-fragment-bit (Services IPsec VPN)" on page 16890	"syslog (Services IPsec VPN)" on page 28816	"tunnel-mtu (Services IPsec VPN)" on page 30218

[edit services ipsec-vpn rule {rule-name} term {term-name} then dead-peer-detection]		
"interval (services)" on page 9201	"threshold (Services IPsec)" on page 29340	

[edit services ipsec-vpn rule {rule-name} term {term-name} then dynamic]	
"natt-install-interval" on page 11242	

[edit services ipsec-vpn rule {rule-name} term {term-name} then manual]	
"direction (services)" on page 5711	

[edit services ipsec-vpn rule {rule-name} term {term-name} then manual direction {direction}]		
"encryption" on page 6322	"protocol (IPsec) " on page 13444	"spi" on page 28384

[edit services ipsec-vpn rule-set {rule-set-name}]	
"rule" on page 16074	

[edit services ipsec-vpn {rule-name }term {term-name} then]	
"dead-peer-detection (Services IPsec VPN)" on page 5151	

[jdaf]

[edit services jdaf]	
"routing-instances (JDAF)" on page 15967	

[jflow-log]

[edit services jflow-log]		
"collector (Flow Monitoring Logs for NAT)" on page 4413	"collector-group (Flow Monitoring Logs for NAT)" on page 4416	"template-profile (Flow Monitoring Logs for NAT)" on page 29090

[edit services jflow-log collector {collector-name}]		
"destination-address (Flow Monitoring Logs for NAT)" on page 5373	"destination-port (Flow Monitoring Logs for NAT)" on page 5471	"source-ip (Flow Monitoring Logs for NAT)" on page 28240

[edit services jflow-log template-profile {template–profile-name}]		
"collector (Flow Template Profiles for NAT)" on page 4415	"collector-group (Flow Template Profiles for NAT)" on page 4418	"refresh-rate (Flow Monitoring Logs for NAT)" on page 14253
"template-type (Flow Monitoring Logs for NAT)" on page 29098	"version (Flow Monitoring Logs for NAT)" on page 30702	

[l2tp]

[edit services l2tp]		
"destination (L2TP)" on page 5445	"destination-equal-load-balancing (L2TP LAC)" on page 5415	"destruct-timeout (L2TP)" on page 5536
"disable-calling-number-avp (L2TP LAC)" on page 5748	"disable-failover-protocol (L2TP)" on page 5755	"drain" on page 5940
"enable-ipv6-services-for-lac (L2TP)" on page 6265	"enable-snmp-tunnel-statistics (L2TP)" on page 6277	"failover-within-preference (L2TP LAC)" on page 6786

(Continued)

[edit services l2tp]		
"ip-reassembly (L2TP)" on page 9263	"maximum-sessions (L2TP)" on page 10560	"rx-connect-speed-when-equal (L2TP LAC)" on page 16166
"sessions-limit-group (L2TP)" on page 16858	"traceoptions (Services L2TP)" on page 29844	"tunnel (L2TP)" on page 30210
"tunnel-group" on page 30206	"tunnel-switch-profile (L2TP Tunnel Switching, Application)" on page 30242	"tx-connect-speed-method (L2TP LAC)" on page 30265
"weighted-load-balancing (L2TP LAC)" on page 31100		

[edit services l2tp destination {ip-address}]		
"access-line-information (L2TP)" on page 1684	"address (L2TP Destination)" on page 1880	"name (L2TP Destination)" on page 11173

[edit services l2tp destination address {ip-address}]		
"drain" on page 5940	"routing-instance (L2TP Destination)" on page 15938	

[edit services l2tp destination address {ip-address} routing-instance {routing-instance-name}]	
"drain" on page 5940	

[edit services l2tp destination lockout-result-code {lockout-result-code-name} lockout-error-code {lockout-error-code- name}]	
"lockout-timeout (L2TP Destination Lockout)" on page 10001	

[edit services l2tp destination name {destination-name}]	
"drain" on page 5940	

[edit services l2tp sessions-limit-group]	
"maximum-sessions (L2TP)" on page 10560	

[edit services l2tp tunnel]		
"assignment-id-format (L2TP LAC)" on page 2504	"failover-resync" on page 6783	"idle-timeout (L2TP)" on page 8307
"maximum-sessions (L2TP)" on page 10560	"minimum-retransmission-timeout (L2TP Tunnel)" on page 10757	"name (L2TP Tunnel Destination)" on page 11175
"nas-port-method (L2TP LAC)" on page 11191	"retransmission-count-established (L2TP)" on page 15680	"retransmission-count-not-established (L2TP)" on page 15681
"rx-window-size (L2TP)" on page 16171	"tx-address-change (L2TP LAC)" on page 30261	

[edit services l2tp tunnel name {name}]		
"address (L2TP Tunnel Destination)" on page 1882	"drain" on page 5940	

[edit services l2tp tunnel name {name} address {ip-address}]		
"drain" on page 5940	"routing-instance (L2TP Tunnel Destination)" on page 15940	

[edit services l2tp tunnel name {name} address {ip-address} routing-instance {routing-instance-name}]		
"drain" on page 5940		

[edit services l2tp tunnel-group {name}]		
"aaa-access-profile (L2TP LNS)" on page 1607	"dynamic-profile (L2TP)" on page 6126	"hello-interval (L2TP)" on page 8033
"hide-avps" on page 8053	"l2tp-access-profile" on page 9547	"local-gateway (L2TP LNS)" on page 9945
"maximum-send-window" on page 10558	"ppp-access-profile" on page 12944	"receive-window" on page 14123
"retransmit-interval (Services)" on page 15687	"service-device-pool (L2TP)" on page 16687	"service-interface (L2TP Processing) " on page 16706
"service-profile (L2TP)" on page 16733	"syslog (Services L2TP)" on page 28829	"tos-reflect (L2TP)" on page 29451
"tunnel-switch-profile (L2TP Tunnel Switching, Application)" on page 30242	"tunnel-timeout" on page 30247	

[edit services l2tp tunnel-group {name} local-gateway]		
"address (LNS Local Gateway)" on page 1883	"gateway-name (LNS Local Gateway)" on page 7654	

[edit services l2tp tunnel-group {group-name} syslog]	
"host (L2TP)" on page 8146	

[edit services l2tp tunnel-group {group-name} syslog host {hostname}]		
"facility-override (System Log Reporting)" on page 6774	"log-prefix (L2TP)" on page 10021	"services (L2TP System Logging)" on page 16780

[edit services l2tp tunnel-group {group-name}],]	
"maximum-sessions (L2TP)" on page 10560	

[logging]

[edit services logging]	
"traceoptions (Services Logging)" on page 29849	

[lrf]

[edit services lrf]	
"profile (LRF)" on page 13289	

[edit services lrf profile {profile-name}]		
"collector (LRF Profile)" on page 4422	"http-log-multiple-transactions (LRF Profile)" on page 8195	"policy-based-logging (LRF Profile)" on page 12672
"rule (LRF)" on page 16101	"template (LRF Profile)" on page 29087	"vendor-support" on page 30687

[edit services lrf profile {profile-name} collector {collector-name}]		
"destination (LRF Profile)" on page 5446	"source-address (LRF Profile)" on page 28138	

[edit services lrf profile {profile-name} collector {collector-name} destination]		
"address (LRF Profile)" on page 1886	"port (LRF Profile)" on page 12834	

[edit services lrf profile {profile-name} rule {lrf-rule-name}]	
"then (LRF rule)" on page 29274	

[edit services lrf profile {profile-name} rule {lrf-rule-name} then]	
"report (LRF Rule)" on page 14395	

[edit services lrf profile {profile-name} rule {lrf-rule-name} then report]		
"collector (LRF Rule)" on page 4424	"template (LRF Rule)" on page 29089	"time-limit (LRF Rule)" on page 29359
"volume-limit (LRF Rule)" on page 30921		

[edit services lrf profile {profile-name} template {template-name}]		
"format (LRF Profile)" on page 7383	"template-tx-interval (LRF Profile)" on page 29097	"template-type (LRF Profile)" on page 29100
"trigger-type (LRF Profile)" on page 30165		

[mobile-edge]

[edit services mobile-edge gateways]		
"saegw" on page 16269		

[edit services mobile-edge gateways saegw]		
"saegw access-network-peers" on page 16270	"saegw control-plane-peers" on page 16272	"saegw system" on page 16278
"saegw-core-network-peers" on page 16276		

[edit services mobile-edge gateways saegw {name} access-network-peers]		
"peer-groups (access network peers)" on page 12361		

[edit services mobile-edge gateways saegw {name} control-plane-peers]		
"apn-services (control plane services)" on page 2292	"peer-groups (control plane peers)" on page 12362	

[edit services mobile-edge gateways saegw {name} core-network-peers]		
"peer-groups (core network peers)" on page 12364		

[edit services mobile-edge gateways saegw system]		
"class-of-service (saegw system)" on page 3372	"destination (saegw system)" on page 5520	"service-mode (mobile-edge SAEGW)" on page 16713

[edit services mobile-edge gateways saegw system class-of-service]		
"downlink-dscp-to-egress-forwarding-class" on page 5922	"downlink-dscp-to-forwarding-class" on page 5923	"reserved-bandwidth" on page 15579

[edit services mobile-edge session-scale {scale-option}]	
"mobile-edge session-scale" on page 10792	

[monitoring]

[edit services monitoring]		
"rfc2544" on page 15757	"twamp" on page 30250	

[edit services monitoring rpm]	
"owner" on page 12090	

[edit services monitoring rpm owner {name} test {test-name}]		
"offload-type" on page 11788	"thresholds (Junos OS Evolved)" on page 29350	

[edit services monitoring twamp]	
"server (Junos OS Evolved)" on page 16644	

[edit services monitoring twamp client]	
"control-connection (Junos OS Evolved)" on page 4627	

[edit services monitoring twamp client control-connection {name}]	
"test-session (Junos OS Evolved)" on page 29239	

[edit services monitoring twamp client control-connection {name} test-session {name}]		
"source-routing" on page 28339	"thresholds (Junos OS Evolved)" on page 29350	

[edit services monitoring twamp server]	
"managed" on page 10298	

[nat]

[edit services nat {destination pool}]		
"address (Security Destination NAT)" on page 1924	"address-persistent (Security Source NAT)" on page 1896	"allow-overlapping-nat-pools" on page 2198
"allow-overlapping-pools (NAT Next Gen Services)" on page 2199	"ipv6-multicast-interfaces" on page 9364	"pool (NAT)" on page 12760
"port-forwarding" on page 12824	"rule (Services NAT) " on page 16120	"rule-set (Services NAT) " on page 16143

[edit services nat destination]		
"pool (Destination NAT Next Gen Services)" on page 12773	"port-forwarding (Destination NAT Next Gen Services)" on page 12826	

[edit services nat destination pool {nat-pool-name}]	
"address (NAT Pool Next Gen Services)" on page 1894	

[edit services nat destination rule-set {rule-set}]		
"match-direction (NAT Next Gen Services)" on page 10376	"rule (Destination NAT Next Gen Services)" on page 16091	

[edit services nat destination rule-set name rule name then destination-nat]	
"xlat-source-rule" on page 31154	

[edit services nat destination rule-set {rule-set} rule {rule-name} match]		
"application (NAT Next Gen Services)" on page 2348	"application (Security Destination NAT)" on page 2369	"destination-address (NAT Next Gen Services)" on page 5383
"destination-address-name (NAT Next Gen Services)" on page 5378	"source-address (NAT Next Gen Services)" on page 28144	"source-address-name (NAT Next Gen Services)" on page 28139

[edit services nat destination rule-set {rule-set-name} rule {rule-name} then]	
"port-forwarding-mappings (Destination NAT Rule Next Gen Services)" on page 12828	

[edit services nat destination rule-set {rule-set} rule {rule-name} then destination-nat]		
"destination-prefix (Destination NAT Next Gen Services)" on page 5493	"off (Destination NAT Next Gen Services)" on page 11784	

[edit services nat destination rule-set {rule-set} rule {rule-name} then source-nat]	
"pool (NAT Rule Next Gen Services)" on page 12778	

[edit services nat natv6v4]	
"clear-dont-fragment-bit (NAT Next Gen Services)" on page 3559	

[edit services nat pool {nat-pool-name}]		
"address (Services NAT Pool)" on page 1929	"address-allocation" on page 1853	"address-range" on page 1917
"app-mapping-timeout" on page 2293	"ei-mapping-timeout" on page 6241	"limit-ports-per-address" on page 9759
"mapping-timeout" on page 10348	"pgcp" on page 12465	"port (Services NAT) " on page 12881

[edit services nat pool {nat-pool-name} pgcp]		
"hint" on page 8081	"ports-per-session" on page 12897	"remotely-controlled" on page 14364
"transport (Services)" on page 30111		

[edit services nat pool {pool-name} port]		
"deterministic-port-block-allocation" on page 5557	"secured-port-block-allocation" on page 16444	

[edit services nat port-forwarding {map-name}]		
"destined-port" on page 5535	"translated-port" on page 30072	

[edit services nat proxy-arp interface]	
"address (Security ARP Proxy)" on page 1922	

[edit services nat proxy-ndp interface]	
"address (Security NDP Proxy)" on page 1925	

[edit services nat rule {rule-name}]		
"match-direction (Services NAT) " on page 10380	"term (Services NAT) " on page 29134	

[edit services nat rule (Services NAT) {rule-name} term (Services NAT) {term-name} then (Services NAT) translated]		
"address-pooling" on page 1904	"filtering-type" on page 7160	"mapping-type" on page 10351

[edit services nat rule {rule-name} term {term-name}]		
"from (Services NAT)" on page 7599	"then (Services NAT) " on page 29299	

[edit services nat rule {rule-name} term {term-name} from]		
"application-sets (Services NAT)" on page 2387	"applications (Services NAT)" on page 2419	"destination-address (Services NAT)" on page 5404
"destination-address-range (Services NAT)" on page 5391	"destination-port range" on page 5475	"destination-prefix-list (Services NAT) " on page 5504
"source-address (Services NAT)" on page 28171	"source-address-range (Services NAT) " on page 28153	"source-prefix-list (Services NAT) " on page 28300

[edit services nat rule {rule-name} term {term-name} then]		
"no-translation" on page 11601	"port-forwarding-mappings" on page 12827	"syslog (Services NAT) " on page 28830
"translated" on page 30070		

[edit services nat rule {name} term {name} then translated]		
"clat-ipv6-prefix-length" on page 3396	"clat-prefix" on page 3398	"destination-pool" on page 5457
"destination-prefix (Services NAT) " on page 5508	"overload-pool" on page 12046	"overload-prefix" on page 12047
"secure-nat-mapping" on page 16436	"source-pool" on page 28276	"source-prefix (Services NAT) " on page 28304
"translation-type" on page 30077		

[edit services nat rule {rule-name} term {term-name} then translated secure-nat-mapping]		
"eif-flow-limit" on page 6244	"mapping-refresh" on page 10345	

[edit services nat rule-set {rule-set-name}]		
"rule (Services NAT) " on page 16120		

[edit services nat rule{ rule-name} term{ term-name} then translated]		
"dns-alg-pool" on page 5849	"dns-alg-prefix" on page 5850	

[edit services nat source]		
"pool (Source NAT Next Gen Services)" on page 12787	"pool-default-port-range (Source NAT Next Gen Services)" on page 12770	"port-round-robin (Source NAT Next Gen Services)" on page 12869

[edit services nat source pool {nat-pool-name}]		
"address (NAT Pool Next Gen Services)" on page 1894	"address-pooling (Source NAT Next Gen Services)" on page 1907	"address-shared (Security Source NAT)" on page 1932
"ei-mapping-timeout (Source NAT Next Gen Services)" on page 6242	"host-address-base (Source NAT Next Gen Services)" on page 8139	"mapping-timeout (Source NAT Next Gen Services)" on page 10349
"pool-utilization-alarm (Source NAT Next Gen Services)" on page 12796	"port (Source NAT Next Gen Services)" on page 12885	

[edit services nat source pool {name} address-persistent subscriber]	
"address-persistent (Security Source NAT Pool)" on page 1897	

[edit services nat source pool {name} address-pooling]	
"address-pooling (Security Source NAT)" on page 1905	

[edit services nat source pool {name} pool-utilization-alarm]		
"clear-threshold" on page 4245		

[edit services nat source pool {nat-pool-name} port]		
"automatic (Source NAT Next Gen Services)" on page 2760	"block-allocation (Source NAT Next Gen Services)" on page 2983	"deterministic (Source NAT Next Gen Services)" on page 5559
"deterministic-nat-configuration-log-interval (Source NAT Next Gen Services)" on page 5555	"no-translation (Source NAT Next Gen Services)" on page 11603	"preserve-parity (Source NAT Next Gen Services)" on page 13124
"preserve-range (Source NAT Next Gen Services)" on page 13125	"range (Source NAT Next Gen Services)" on page 14049	

[edit services nat source pool-utilization-alarm]		
"clear-threshold" on page 4245		

[edit services nat source rule-set {rule-set}]		
"match-direction (NAT Next Gen Services)" on page 10376	"rule (Source NAT Next Gen Services)" on page 16156	

[edit services nat source rule-set {rule-set} rule {rule-name} match]		
"application (NAT Next Gen Services)" on page 2348	"application (Security Source NAT)" on page 2373	"destination-address (NAT Next Gen Services)" on page 5383
"destination-address-name (NAT Next Gen Services)" on page 5378	"source-address (NAT Next Gen Services)" on page 28144	"source-address-name (NAT Next Gen Services)" on page 28139

[edit services nat source rule-set {rule-set} rule {rule-name} then source-nat]		
"clat-prefix (Source NAT Next Gen Services)" on page 3399	"filtering-type (Source NAT Next Gen Services)" on page 7161	"mapping-type (Source NAT Next Gen Services)" on page 10352
"pool (NAT Rule Next Gen Services)" on page 12778	"secure-nat-mapping (Source NAT Next Gen Services)" on page 16438	

[edit services nat source],]	
"subscriber-pool-utilization-alarm" on page 28705	

[network-monitoring]

[edit services network-monitoring]		
"profile (Traffic Load Balancer)" on page 13349	"traceoptions (Health Check Monitoring)" on page 29569	

[paa]

[edit services paa]	
"test-agent (Paragon Active Assurance)" on page 29187	

[pcef]

[edit services pcef]		
"flow-descriptions" on page 7274	"pcc-action-profiles" on page 12286	"pcc-rulebases (PCEF)" on page 12293
"pcc-rules (PCEF)" on page 12297	"profile (Services PCEF)" on page 13336	"profiles (PCEF)" on page 13365

[edit services pcef flow-description {flow-identifier}]		
"local-ports" on page 9969	"protocol (Flow Descriptions)" on page 13441	"remote-ports" on page 14357

[edit services pcef flow-descriptions {flow-identifier}]		
"direction (Service Data Flow Filters)" on page 5727	"local-port-range" on page 9967	"remote-address" on page 14311
"remote-port-range" on page 14355		

[edit services pcef pcc-action-profiles {profile-name}]		
"forwarding-class (PCC Action Profiles)" on page 7434	"gate-status" on page 7649	"ipv4-address (Steering Path)" on page 9327
"ipv6-address (Steering Path)" on page 9347	"keep-existing-steering" on page 9459	"logging-rule (PCC Action Profile)" on page 10054
"maximum-bit-rate (PCC Action Profiles)" on page 10507	"path (Steering)" on page 12269	"redirect (PCC Action Profiles)" on page 14159
"steering" on page 28605		

[edit services pcef pcc-action-profiles {profile-name} redirect]	
"url" on page 30485	

[edit services pcef pcc-action-profiles {profile-name} steering]	
"routing-instance (PCC Action Profiles)" on page 15951	

[edit services pcef pcc-rule-bases {rulebase-name}]	
"pcc-rule" on page 12291	

[edit services pcef pcc-rules {rule-name}]		
"from (PCC Rules)" on page 7591	"then (PCC Rules)" on page 29276	

[edit services pcef pcc-rules {rule-name} from]		
"application-groups (PCC Rules)" on page 2326	"applications (PCC Rules)" on page 2410	"flows (PCC Rules)" on page 7362

[edit services pcef pcc-rules {rules-name} then]	
"pcc-action-profile (PCC Rules)" on page 12284	

[edit services pcef profiles {profile-name}]	
"static-policy-control" on page 28555	

[edit services pcef profiles {profile-name} dynamic-policy-control]		
"diameter-profile (PCEF Profile)" on page 5698	"pcc-rulebases (PCEF Profile)" on page 12294	"pcc-rules (PCEF Profile)" on page 12299

[edit services pcef profiles {profile-name} static-policy-control]		
"pcc-rulebases (PCEF Profile)" on page 12294	"pcc-rules (PCEF Profile)" on page 12299	

[pcp]

[edit services pcp]		
"rule (PCP)" on page 16103	"server (pcp)" on page 16651	

[edit services pcp rule {rule-name}]		
"match-direction (PCP)" on page 10377	"term (PCP)" on page 29128	

[edit services pcp rule {rule-name} term {term-name}]		
"from (PCP)" on page 7593	"then (PCP)" on page 29278	

[edit services pcp rule {rule-name} term {term-name} from]		
"application-sets (PCP)" on page 2384	"applications (PCP)" on page 2412	"destination-address (PCP)" on page 5384
"destination-address-range (PCP)" on page 5390	"destination-port (PCP)" on page 5472	"destination-prefix-list (PCP)" on page 5498

(Continued)

[edit services pcp rule {rule-name} term {term-name} from]		
"source-address (PCP)" on page 28146	"source-address-range (PCP)" on page 28152	"source-prefix-list (PCP)" on page 28295

[edit services pcp rule {rule-name} term {term-name} then]	
"pcp-server" on page 12324	

[pfe]

[edit services pfe traffic]	
"monitor-timer" on page 10953	

[edit services pfe traffic traceoptions]	
"file (Packet Forwarding Engine Traffic Traceoptions)" on page 7069	

[plmn]

[edit services {name} plmn]	
"s-nssai" on page 16259	

[policies]

[edit services policies]		
"file-serialization" on page 7094	"stateful-firewall-rule-set (Next Gen Services)" on page 28518	

[edit services policies stateful-firewall-rule {rule-name}]	
"policy (Stateful Firewall Rules Next Gen Services)" on page 12734	

[edit services policies stateful-firewall-rule {rule-name} policy {policy-name}]		
"match (Stateful Firewall Rule Next Gen Services)" on page 10406	"then (Stateful Firewall Rule Next Gen Services)" on page 29302	

[proxy]

[edit services proxy]		
"profile (Services Proxy)" on page 13338	"protocol (Services Proxy)" on page 13472	

[radius-flow-tap]

[edit services radius-flow-tap]	
"policy (Subscriber Secure Policy)" on page 12745	

[edit services radius-flow-tap policy {policy-name}]		
"inet (Subscriber Secure Policy)" on page 8591	"inet6 (Subscriber Secure Policy)" on page 8616	

[edit services radius-flow-tap policy{ policy-name }inet6 drop-policy{ rule-name} from]		
"protocol (Subscriber Secure Policy)" on page 13474	"source-port (Subscriber Secure Policy)" on page 28290	

[edit services radius-flow-tap policy{ policy-name } inet drop-policy{ rule-name} from]		
"apply-groups (Subscriber Secure Policy)" on page 2428	"apply-groups-except (Subscriber Secure Policy)" on page 2427	"destination-port (Subscriber Secure Policy)" on page 5490
"source-address (Subscriber Secure Policy)" on page 28176		

[edit services radius-flow-tap policy{ policy-name} inet drop-policy{ rule-name }from]		
"destination-address (Subscriber Secure Policy)" on page 5407	"dscp (Subscriber Secure Policy)" on page 6023	"protocol (Subscriber Secure Policy)" on page 13474

[edit services radius-flow-tap policy{ policy-name} inet6 drop-policy{ rule-name} from]		
"apply-groups (Subscriber Secure Policy)" on page 2428	"apply-groups-except (Subscriber Secure Policy)" on page 2427	"destination-address (Subscriber Secure Policy)" on page 5407
"destination-port (Subscriber Secure Policy)" on page 5490	"dscp (Subscriber Secure Policy)" on page 6023	"source-address (Subscriber Secure Policy)" on page 28176
"source-port (Subscriber Secure Policy)" on page 28290		

[edit services radius-flow-tap policy{ policy-name} inet| inet6]	
"drop-policy (Subscriber Secure Policy)" on page 5950	

[radius-flow-tappolicy{]

[edit services radius-flow-tappolicy{ policy-name }inet| inet6]	
"from (Subscriber Secure Policy)" on page 7602	

[redundancy-set]

[edit services redundancy-set]	
"traceoptions (Services Redundancy Daemon)" on page 29851	

[rpm]

[edit services rpm]		
"probe" on page 13228	"probe-limit" on page 13234	"rfc2544-benchmarking" on page 15759
"traceoptions (RPM)" on page 29761	"twamp" on page 30250	

[edit services rpm bgp]		
"bgp (RPM)" on page 2958	"logical-system" on page 10073	"routing-instances (RPM)" on page 15970

[edit services rpm bgp logical-system {logical-system-name}]	
"routing-instances (RPM)" on page 15970	

[edit services rpm probe {owner}]	
"delegate-probes" on page 5266	

[edit services rpm probe {owner} test {test-name}]		
"destination-interface" on page 5430	"hardware-timestamp" on page 7975	"inet6-options (Services)" on page 8613
"next-hop (RPM)" on page 11350	"one-way-hardware-timestamp" on page 11830	"rpm-scale" on page 16020
"thresholds (Junos OS)" on page 29348		

[edit services rpm probe-server (tcp | udp)]	
"destination-interface" on page 5430	

[edit services rpm rfc2544-benchmarking]	
"profiles (RFC 2544 Benchmarking)" on page 13368	

[edit services rpm rfc2544-benchmarking profilestest-profile {profile-name}]	
"bandwidth-kbps" on page 2850	

[edit services rpm rfc2544-benchmarking tests test-name {test-name}]	
"ivlan-priority (RFC 2544 Benchmarking)" on page 9419	

[edit services rpm rfc2544-benchmarkingtests test-name {test-name}]		
"forwarding-class (RFC 2544 Benchmarking)" on page 7437	"ivlan-cfi (RFC 2544 Benchmarking)" on page 9416	"ivlan-id (RFC 2544 Benchmarking)" on page 9417
"outer-tag-protocol-id (RFC 2544 Benchmarking)" on page 11975	"ovlan-cfi (RFC 2544 Benchmarking)" on page 12082	"ovlan-id (RFC 2544 Benchmarking)" on page 12084
"packet-loss-priority (RFC 2544 Benchmarking)" on page 12155	"receive-failure-threshold (RFC 2544 Benchmarking)" on page 14113	"service-type (RFC 2544 Benchmarking Tests)" on page 16763
"test-finish-wait-duration (RFC 2544 Benchmarking)" on page 29193	"test-iterator-pass-threshold (RFC 2544 Benchmarking)" on page 29201	"timestamp-format (RFC 2544 Benchmarking)" on page 29398
"transmit-failure-threshold (RFC 2544 Benchmarking)" on page 30084		

[edit services rpm twamp]		
"post-cli-implicit-firewall" on page 12901	"server (Junos OS)" on page 16642	

[edit services rpm twamp (client control-connection {control-client-name} |server)]		
"tcp-keepcnt" on page 28991	"tcp-keepidle" on page 28992	"tcp-keepintvl" on page 28994

[edit services rpm twamp client]		
"control-connection (Junos OS)" on page 4624		

[edit services rpm twamp client control-connection {control-client-name}]		
"authentication-mode" on page 2630	"destination-interface" on page 5430	"persistent-results" on page 12454
"test-count" on page 29192	"test-session (Junos OS)" on page 29237	"thresholds (Junos OS)" on page 29348

[edit services rpm twamp client control-connection {control-client-name} test-session {session-name}]		
"data-fill-with zeros" on page 4831	"source-address (TWAMP)" on page 28179	

[edit services rpm twamp server]		
"authentication-key-chain (TWAMP)" on page 2608	"authentication-mode" on page 2630	"inactivity-timeout (Services RPM)" on page 8510
"max-connection-duration" on page 10424	"maximum-connections" on page 10514	"maximum-connections-per-client" on page 10516
"maximum-sessions" on page 10559	"maximum-sessions-per-connection" on page 10562	"routing-instance-list (TWAMP)" on page 15941
"server-inactivity-timeout" on page 16641		

[rtlog]

[edit services rtlog]		
"traceoptions (Next Gen Services Global System Logging)" on page 29619		

[edit services rtlog traceoptions]		
"file (Next Gen Services Global System Logging)" on page 7065	"flag (Next Gen Services Global System Logging)" on page 7208	"no-remote-trace (Next Gen Services Global System Logging)" on page 11564

[edit services rtlog traceoptions file]		
"filename (Next Gen Services Global System Logging)" on page 7099	"match (Next Gen Services Global System Logging)" on page 10386	"no-world-readable (Next Gen Services Global System Logging)" on page 11624
"size (Next Gen Services Global System Logging)" on page 28001	"world-readable (Next Gen Services Global System Logging)" on page 31145	

[edit services rtlog traceoptions file filename]	
"files (Next Gen Services Global System Logging)" on page 7105	

[edit services rtlog traceoptions flag]	
"name (Next Gen Services Global System Logging)" on page 11176	

[screen]

[edit services screen]		
"cpu-throttle (Next Gen Services)" on page 4701		

[edit services screen ids-option {screen-name}]		
"aggregations (IDS Screen Next Gen Services)" on page 2088	"alarm-without-drop (IDS Screen Next Gen Services)" on page 2123	"icmp (IDS Screen Next Gen Services)" on page 8260
"ip (IDS Screen Next Gen Services)" on page 9241	"limit-session (IDS Screen Next Gen Services)" on page 9765	

[edit services screen ids-option {screen-name} icmp]		
"fragment (IDS Screen Next Gen Services)" on page 7540	"icmpv6-malformed (IDS Screen Next Gen Services)" on page 8276	"large (IDS Screen Next Gen Services)" on page 9633
"ping-death (IDS Screen Next Gen Services)" on page 12538		

[edit services screen ids-option {screen-name} ip]		
"bad-option (IDS Screen Next Gen Services)" on page 2830	"block-frag (IDS Screen Next Gen Services)" on page 2988	"ipv6-extension-header (IDS Screen Next Gen Services)" on page 9353
"loose-source-route-option (IDS Screen Next Gen Services)" on page 10139	"record-route-option (IDS Screen Next Gen Services)" on page 14139	"security-option (IDS Screen Next Gen Services)" on page 16504
"source-route-option (IDS Screen Next Gen Services)" on page 28310	"stream-option (IDS Screen Next Gen Services)" on page 28633	"strict-source-route-option (IDS Screen Next Gen Services)" on page 28652

(Continued)

[edit services screen ids-option {screen-name} ip]		
"tear-drop (IDS Screen Next Gen Services)" on page 29046	"timestamp-option (IDS Screen Next Gen Services)" on page 29399	"unknown-protocol (IDS Screen Next Gen Services)" on page 30416

[edit services screen ids-option {screen-name} limit-session]		
"by-destination (IDS Screen Next Gen Services)" on page 3095	"by-source (IDS Screen Next Gen Services)" on page 3108	

[edit services screen ids-option {screen-name} limit-session by-destination]	
"by-protocol (IDS Screen Next Gen Services)" on page 3102	

[edit services screen ids-option {screen-name} limit-session by-source]	
"by-protocol (IDS Screen Next Gen Services)" on page 3102	

[edit services screen ids-option {screen-name} tcp]		
"fin-no-ack (IDS Screen Next Gen Services)" on page 7168	"land (IDS Screen Next Gen Services)" on page 9622	"syn-ack-ack-proxy (IDS Screen Next Gen Services)" on page 28771
"syn-fin (IDS Screen Next Gen Services)" on page 28777	"syn-frag (IDS Screen Next Gen Services)" on page 28782	"tcp-no-flag (IDS Screen Next Gen Services)" on page 29005
"winnuke (IDS Screen Next Gen Services)" on page 31127		

[security-metadata-streaming]

[edit services security-metadata-streaming policy]	
"dynamic-filter" on page 6108	

[service-device-pools]

[edit services service-device-pools]	
"pool (Service Interfaces)" on page 12786	

[edit services service-device-pools pool {pool-name}]	
"interface (Service Interfaces)" on page 9023	

[service-interface-pools]

[edit services service-interface-pools]	
"pool (Service Interface)" on page 12784	

[edit services service-interface-pools pool {pool-name}]	
"interface (services)" on page 8811	

[service-set]

[edit services service-set {service-set-name}]		
"aacl-statistics-profile" on page 1623	"allow-multicast" on page 2195	"application-identification-profile (Service Set)" on page 2343
"captive-portal-content-delivery-profile (Services)" on page 3202	"cos-rule-sets (Service Set Next Gen Services)" on page 4677	"cos-rules (Service Set Next Gen Services)" on page 4680
"cos-rules (Service Set)" on page 4679	"hcm-profile (HTTP Header Enrichment)" on page 8007	"ids-rule-sets" on page 8340
"ids-rules" on page 8341	"interface-service (Services Interfaces)" on page 9027	"ip-reassembly-rules (Service Set)" on page 9265
"ipsec-vpn-options" on page 9323	"ipsec-vpn-rules" on page 9325	"lrf-profile (Service Set)" on page 10180
"max-drop-flows" on page 10435	"max-flows" on page 10439	"max-session-setup-rate (Service Set)" on page 10467
"nat-options" on page 11221	"nat-options (Next Gen Services)" on page 11222	"nat-rule-sets (Service Set Next Gen Services)" on page 11231
"nat-rule-sets (Service Set)" on page 11230	"nat-rules" on page 11232	"next-hop-service" on page 11352
"pcef-profile (Service Set)" on page 12315	"pcp-rules" on page 12323	"pgcp-rules" on page 12466
"ptsp-rules" on page 13866	"redundancy-set-id (Service Set)" on page 14206	"service-order" on page 16719
"service-set-options" on page 16746	"service-set-options (Next Gen Services Services)" on page 16748	"snmp-trap-thresholds" on page 28070

(Continued)

[edit services service-set {service-set-name}]		
"snmp-trap-thresholds (Next Gen Services)" on page 28072	"softwire-options" on page 28095	"softwire-rules" on page 28097
"softwires-rule-set (Service Set Next Gen Services)" on page 28105	"stateful-firewall-rule-sets (Service Set Next Gen Services)" on page 28519	"stateful-firewall-rules" on page 28520
"stateful-firewall-rules (Service Set Next Gen Services)" on page 28521	"subscriber-profile" on page 28700	"syslog (Next Gen Services Service-Set System Logging)" on page 28823
"syslog (Services Service Set)" on page 28831	"tag-rule-sets (Service Set)" on page 28931	"tag-rules (Service Set)" on page 28932
"tcp-log" on page 28995	"tcp-mss (Services)" on page 29002	"url-filter-profile" on page 30495
"web-filter-profile" on page 31065		

[edit services service-set {service-set-name} cos-options]		
"match-rules-on-reverse-flow" on page 10391	"match-rules-on-reverse-flow (Next Gen Services)" on page 10392	

[edit services service-set {service-set-name} interface-service]		
"load-balancing-options (Service Set)" on page 9890	"service-interface (Services Interfaces) " on page 16709	

[edit services service-set {service-set-name} interface-service load-balancing-options]	
"hash-keys (Aggregated Multiservices)" on page 7992	

[edit services service-set {service-set-name} ipsec-vpn-options]		
"anti-replay-window-size (Services Service Set)" on page 2258	"clear-dont-fragment-bit (Services Service Set)" on page 3562	"copy-dont-fragment-bit (Services Set)" on page 4662
"ike-access-profile" on page 8418	"local-gateway (IPSec) " on page 9944	"nat-keepalive (Services IPsec VPN)" on page 11218
"no-anti-replay (Services Service Set)" on page 11396	"no-certificate-chain-in-ike" on page 11423	"no-nat-traversal (Services IPsec VPN)" on page 11515
"passive-mode-tunneling" on page 12207	"set-dont-fragment-bit (Services Set)" on page 16891	"trusted-ca" on page 30171
"tunnel-mtu (Services Service Set)" on page 30219	"udp-encapsulation" on page 30307	

[edit services service-set {service-set-name} ipsec-vpn-options local-gateway {address}]	
"gw-interface" on page 7951	

[edit services service-set {service-set-name} jflow-log]	
"template-profile (Flow Monitoring Logs for NAT)" on page 29090	

[edit services service-set name flow]	
"traceoptions (Next Gen Services Service-Set Flow)" on page 29620	

[edit services service-set {service-set-name} nat-options]		
"land-attack-check" on page 9619	"max-sessions-per-subscriber" on page 10474	"snmp-value-match-msmic (Services NAT Options)" on page 28074

(Continued)

[edit services service-set {service-set-name} nat-options]	
"stateful-nat64" on page 28523	

[edit services service-set {service-set-name} nat-options stateful-nat64]	
"clear-dont-fragment-bit (Services NAT Options)" on page 3561	

[edit services service-set {service-set-name} service-set-options]		
"bypass-traffic-on-exceeding-flow-limits" on page 3119	"bypass-traffic-on-pic-failure" on page 3120	"disable-global-timeout-override" on page 5760
"enable-asymmetric-traffic-processing (Service Set Next Gen Services)" on page 6254	"header-integrity-check" on page 8010	"header-integrity-check (Next Gen Services)" on page 8012
"max-sessions-per-subscriber (Service Set Next Gen Services)" on page 10475	"routing-engine-services" on page 15911	"session-limit (Service Set Next Gen Services)" on page 16823
"session-timeout (Service Set Next Gen Services)" on page 16854	"static-subscriber-application" on page 28579	"subscriber-awareness (Service Set Options)" on page 28685
"tcp-fast-open" on page 28982	"tcp-non-syn" on page 29006	"tcp-session (Service Set Next Gen Services)" on page 29023

[edit services service-set {service-set-name} service-set-options tcp-session]		
"close-timeout" on page 4365	"inactivity-asymm-tcp-timeout (Service Set Next Gen Services)" on page 8500	"open-timeout" on page 11841

[edit services service-set {services-set-name} subscriber-profile {profile-name}]		
"disable" on page 5736	"enable" on page 6249	"max-data-sessions-per-subscriber" on page 10427

[edit services service-set {services-set-name} subscriber-profile {profile-name} enable {service-name}]	
"concurrent-data-sessions" on page 4500	

[edit services service-set {services-set-name} subscriber-profile {profile-name} max-data-sessions-per-subscriber]		
"exceed-action" on page 6570	"limit" on page 9757	

[edit services service-set {service-set-name} syslog]		
"host (service-set)" on page 8169	"local-category (Next Gen Services Service-Set Local System Logging)" on page 9926	"transport (Next Gen Services Syslog Message Security)" on page 30127
"utc-timestamp" on page 30612		

[edit services service-set {name } syslog edit services service-set {name} syslog stream {stream-name}]	
"local-log-tag (Next Gen Services Service-Set System Logging)" on page 9959	

[edit services service-set {service-set-name} syslog host {hostname}]		
"class" on page 3342	"facility-override (Service Sets)" on page 6772	"log-prefix (Services)" on page 10022
"services (System Logging)" on page 16800	"source-address (Service Sets) " on page 28168	

[edit services service-set {service-set-name} syslog stream]	
"category (System Logging)" on page 3236	

[edit services service-set {service-set-name} web-filter]	
"multi-tenant-support" on page 11059	

[service-set-options]

[edit services {service-set-name} service-set-options]	
"inactivity-timeout (services)" on page 8506	

[services-set]

[edit services services-set {name} syslog]		
"event-rate (Next Gen Services Service-Set Local System Logging)" on page 6532	"mode (Next Gen Services Service-Set System Logging)" on page 10811	

[softwire]

[edit services softwire]		
"ipv6-multicast-interfaces" on page 9364	"rule (Softwire)" on page 16154	"rule-set (Softwire)" on page 16146
"softwire-concentrator" on page 28092		

[edit services softwire rule-set {rule-set-name}]	
"rule (Softwire)" on page 16154	

[edit services softwire softwire-concentrator]		
"ds-lite" on page 5986	"map-e" on page 10329	"map-t" on page 10334
"v6rd" on page 30641		

[softwires]

[edit services softwires]		
"rule-set (Softwires Next Gen Services)" on page 16147	"softwire-name (Next Gen Services)" on page 28093	"softwire-types (Next Gen Services)" on page 28098

[edit services softwires softwire-types]		
"ds-lite" on page 5986	"map-e" on page 10329	"v6rd" on page 30641

[source]

[edit services {nat} source {rule-set} {name} {rule} {name} then source-nat]	
"address-mapping" on page 1887	

[ssl]

[edit services ssl]		
"initiation (Services)" on page 8664	"proxy (Services)" on page 13841	"termination (Services)" on page 29151
"traceoptions (Services SSL)" on page 29854		

[edit services ssl initiation]	
"profile (SSL Initiation)" on page 13343	

[edit services ssl initiation profile {profile-name}]		
"actions (Services SSL Initiation)" on page 1793	"custom-ciphers" on page 4751	"enable-flow-tracing (Services)" on page 6263
"enable-session-cache" on page 6276	"preferred-ciphers" on page 13055	"protocol-version" on page 13475
"trusted-ca (Services)" on page 30172		

[edit services ssl initiation profile {profile-name } actions]		
"crl" on page 4717		

[edit services ssl proxy]		
"global-config (Services)" on page 7700	"profile (Services SSL Proxy)" on page 13339	

[edit services ssl proxy (Services) profile (Services SSL Proxy)]		
"actions " on page 1790		

[edit services ssl proxy profile {profile-name}]		
"custom-ciphers" on page 4751	"enable-flow-tracing (Services)" on page 6263	"mirror-decrypt-traffic" on page 10766
"preferred-ciphers" on page 13055	"root-ca (Services)" on page 15834	"trusted-ca (Services)" on page 30172
"whitelist " on page 31111	"whitelist-url-categories" on page 31112	

[edit services ssl proxy profile {profile-name} actions]		
"crl" on page 4717	"log (Services)" on page 10038	"renegotiation (Services)" on page 14373

[edit services ssl termination]		
"profile (SSL Termination)" on page 13347		

[edit services ssl termination profile {profile-name}]		
"custom-ciphers" on page 4751	"enable-flow-tracing (Services)" on page 6263	"enable-session-cache" on page 6276
"preferred-ciphers" on page 13055	"protocol-version" on page 13475	"root-ca (Services)" on page 15834
"server-certificate (Services)" on page 16617	"trusted-ca (Services)" on page 30172	"whitelist " on page 31111

[edit services ssl traceoptions]		
"flag (Services)" on page 7209	"level (Services)" on page 9735	

[edit services ssl traceoptions file {file-name}]	
"size (Services)" on page 28006	

[stateful-firewall]

[edit services stateful-firewall]		
"rule (Services Stateful Firewall) " on page 16122	"rule-set (Services Stateful Firewall) " on page 16144	

[edit services stateful-firewall rule {rule-name}]		
"match-direction (Services Stateful Firewall) " on page 10381	"term (Services Stateful Firewall) " on page 29136	

[edit services stateful-firewall rule {rule-name} term {term-name}]		
"from (Services Stateful Firewall) " on page 7600	"then (Services Stateful Firewall) " on page 29300	

[edit services stateful-firewall rule {rule-name} term {term-name} from]		
"application-sets (Services Stateful Firewall) " on page 2388	"destination-address (Services Stateful Firewall)" on page 5405	"destination-address-range (Services Stateful Firewall)" on page 5393
"destination-prefix-list (Services Stateful Firewall) " on page 5505	"source-address (Services Stateful Firewall) " on page 28172	"source-address-range (Services Stateful Firewall) " on page 28155
"source-prefix-list (Services Stateful Firewall) " on page 28301		

[edit services stateful-firewall rule {rule-name} term {term-name} then]	
"syslog (Services Stateful Firewall) " on page 28833	

[edit services stateful-firewall rule-set {rule-set-name}]	
"rule (Services Stateful Firewall) " on page 16122	

[traffic-load-balance]

[edit services traffic-load-balance]		
"instance (Traffic Load Balancer)" on page 8770	"routing-engine-mode (TLB)" on page 15995	"traceoptions (Traffic Load Balancer)" on page 29899

[edit services traffic-load-balance instance {instance-name}]		
"group (Traffic Load Balancer)" on page 7906	"real-service (Traffic Load Balancer)" on page 14092	"virtual-service (Traffic Load Balancer)" on page 30774

[unified-access-control]

[edit services unified-access-control]		
"captive-portal (Services UAC)" on page 3203	"certificate-verification" on page 3260	"infranet-controller" on page 8623
"interval (Services)" on page 9221	"test-only-mode" on page 29208	"timeout (Services)" on page 29395
"timeout-action" on page 29371	"traceoptions (Services UAC)" on page 29857	

[edit services unified-access-control captive-portal {policy}]		
"redirect-traffic" on page 14160	"redirect-url" on page 14162	

[edit services unified-access-control infranet-controller {hostname}]		
"address (Services)" on page 1926	"ca-profile (Services)" on page 3160	"interface (Services)" on page 9028
"password (Services)" on page 12246	"port (Services)" on page 12880	"server-certificate-subject" on page 16618

[url-filter]

[edit services url-filter]	
"profile (Web Filter)" on page 13357	

[edit services url-filter profile {profile-name}]	
"template (Services URL Filter)" on page 29072	

[edit services url-filter profile {profile-name} template {template-name}]	
"term (URL Filter)" on page 29143	

[user-identification]

[edit services user-identification]		
"active-directory-access" on page 1802	"authentication-source (Services User Identification ClearPass)" on page 2660	"identity-management" on page 8293

[edit services user-identification active-directory domain]	
"ip-user-mapping" on page 9283	

[edit services user-identification active-directory-access]		
"firewall-authentication-forced-timeout" on page 7181	"invalid-authentication-entry-timeout (Services User Identification Active Directory and ClearPass)" on page 9227	"traceoptions (Active Directory Access)" on page 29483

[edit services user-identification active-directory-access domain]	
"user-group-mapping" on page 30555	

[edit services user-identification authentication-source (Services User Identification ClearPass) aruba-clearpass]		
"authentication-entry-timeout (Services User Identification)" on page 2599	"invalid-authentication-entry-timeout (Services User Identification Active Directory and ClearPass)" on page 9227	"no-user-query (Services User Identification)" on page 11614
"traceoptions (Services User Identification)" on page 29859	"user-query (Services User Identification)" on page 30584	

[edit services user-identification authentication-source (Services User Identification ClearPass) aruba-clearpass user-query (Services User Identification)]	
"web-server (Services)" on page 31086	

[edit services user-identification device-information]		
"authentication-source (Services User Identification Device Identity)" on page 2662	"end-user-profile " on page 6339	"identity-context-profile" on page 8291

[edit services user-identification identity-management]		
"authentication-entry-timeout (Identity Management Advanced Query)" on page 2597	"batch query" on page 2886	"connection (Identity Management Advanced Query)" on page 4533
"filter (Identity Management Advanced Query)" on page 7138	"ip-query (Identity Management Advanced Query)" on page 9259	

[edit services user-identification identity-management connection]		
"primary connection (Identity Management Advanced Query) " on page 13131	"secondary connection (Identity Management Advanced Query) " on page 16409	

[video-monitoring]

[edit services video-monitoring]		
"generate-snmp-traps" on page 7676	"interfaces (Video Monitoring)" on page 9176	"templates" on page 29102

[web-filter]

[edit services web-filter]		
"multi-tenant-hash" on page 11057	"profile (Web Filter)" on page 13357	

[edit services web-filter profile {profile-name}]		
"dns-filter" on page 5853	"dns-filter-template" on page 5856	"feed-name" on page 6997
"global-dns-stats-log-timer" on page 7705	"security-intelligence-policy" on page 16467	"url-filter-template" on page 30496

[edit services web-filter profile {profile-name} dns-filter-template {template-name}]		
"dns-filter" on page 5853	"feed-name" on page 6997	

[edit services web-filter profile {profile-name} dns-filter-template {template-name} term {term-name}]	
"feed-name" on page 6997	

[edit services web-filter profile {profile-name} security-intelligence-policy]		
"black-list" on page 2978	"geo-ip (global-filter)" on page 7678	"white-list" on page 31106

[edit services web-filter profile {name} url-filter-template]		
"disable-https-filtering" on page 5812	"security-intelligence-policy" on page 16467	"term (URL Filter)" on page 29143

[edit services web-filter profile {profile-name} url-filter-template {template-name} security-intelligence-policy]	
"geo-ip (group-filter)" on page 7679	

[web-proxy]

[edit services web-proxy]	
"secure-proxy" on page 16441	

[edit smtp]

IN THIS SECTION

- [traceoptions] | **1132**

Sub Hierarchies

"[edit smtp traceoptions]" on page 1132

[edit smtp]		
"primary-server" on page 13146	"secondary-server" on page 16418	"traceoptions (SMTP)" on page 29865

[traceoptions]

[edit smtp traceoptions]	
"flag (SMTP)" on page 7211	

[edit snmp]

IN THIS SECTION

- [alarm-management] | **1134**
- [community] | **1134**
- [communitycommunity] | **1135**
- [health-monitor] | **1135**
- [nonvolatile] | **1136**
- [rmon] | **1136**
- [routing-instance-access] | **1137**
- [trap-group] | **1137**
- [trap-options] | **1138**
- [v3] | **1138**
- [v3target-address] | **1143**
- [view] | **1143**

Sub Hierarchies

"[edit snmp alarm-management]" on page 1134

"[edit snmp community {community-name}]" on page 1134

"[edit snmp communitycommunity {community-name}]" on page 1135

"[edit snmp health-monitor]" on page 1135

"[edit snmp nonvolatile]" on page 1136

"[edit snmp rmon]" on page 1136

"[edit snmp routing-instance-access]" on page 1137

"[edit snmp trap-group {group-name}]" on page 1137

"[edit snmp trap-options]" on page 1138

"[edit snmp v3]" on page 1138

"[edit snmp v3target-address {target-address-name}]" on page 1143

"[edit snmp view {view-name}]" on page 1143

[edit snmp]		
"alarm-management" on page 2106	"client-list (SNMP)" on page 4329	"community (SNMP)" on page 4475
"contact (SNMP)" on page 4575	"customization (SNMP)" on page 4784	"description (SNMP)" on page 5322
"engine-id" on page 6346	"filter-duplicates" on page 7129	"filter-interfaces" on page 7142
"health-monitor" on page 8016	"interface (SNMP)" on page 9054	"location (SNMP)" on page 9992
"logical-system-trap-filter" on page 10088	"name (SNMP)" on page 11167	"nonvolatile" on page 11675
"proxy (snmp)" on page 13843	"rmon" on page 15810	"routing-instance-access" on page 15924
"traceoptions (SNMP)" on page 29866	"trap-group" on page 30135	"trap-options" on page 30136
"v3" on page 30631	"view (Configuring a MIB View)" on page 30733	

[alarm-management]

[edit snmp alarm-management]		
"alarm-list-name" on page 2103		

[edit snmp alarm-management alarm-list-name]		
"alarm-id" on page 2102	"alarm-state" on page 2117	

[community]

[edit snmp community {community-name}]		
"authorization (SNMP)" on page 2677	"client-list-name" on page 4331	"logical-system (SNMP)" on page 10075
"routing-instance (SNMP)" on page 15918	"view (SNMP Community)" on page 30734	

[edit snmp community {community-name} logical-system {logical-system-name}]	
"routing-instance (SNMP)" on page 15918	

[communitycommunity]

[edit snmp communitycommunity {community-name}]	
"clients (SNMP Community)" on page 4344	

[health-monitor]

[edit snmp health-monitor]		
"falling-threshold (Health Monitor)" on page 6811	"falling-threshold (percentage)" on page 6815	"idp (SNMP)" on page 8334
"interval (Health Monitor)" on page 9213	"rising-threshold (Health Monitor)" on page 15807	

[edit snmp health-monitor idp]		
"falling-threshold (percentage)" on page 6815	"interval (Health Monitor)" on page 9213	"rising-threshold (Health Monitor)" on page 15807

[edit snmp health-monitor routing-engine]	
"routing-engine (SNMP Resource Level)" on page 15912	

[nonvolatile]

[edit snmp nonvolatile]		
"commit-delay" on page 4446		

[rmon]

[edit snmp rmon]		
"alarm (SNMP RMON)" on page 2115	"event" on page 6520	"history" on page 8082
"owner (RMON)" on page 12092		

[edit snmp rmon alarm {index}]		
"description (RMON)" on page 5321	"falling-event-index" on page 6808	"falling-event-index (RMON)" on page 6809
"falling-threshold (integer)" on page 6812	"falling-threshold (RMON)" on page 6817	"falling-threshold-interval" on page 6814
"interval (RMON)" on page 9206	"interval (SNMP RMON)" on page 9223	"request-type" on page 15433
"rising-event-index" on page 15805	"rising-threshold (RMON)" on page 15808	"sample-type" on page 16290
"startup-alarm" on page 28511	"syslog-subtag" on page 28834	"variable" on page 30651
"variable" on page 30653		

[edit snmp rmon event {index}]		
"community" on page 4462	"community (RMON)" on page 4471	"description (RMON)" on page 5321
"type (RMON Notification)" on page 30281	"type (RMON)" on page 30279	

[edit snmp rmon history {index}]		
"bucket-size" on page 3046	"interface (SNMP RMON History)" on page 9056	"interval (SNMP RMON)" on page 9223

[routing-instance-access]

[edit snmp routing-instance-access]	
"access-list" on page 1686	

[trap-group]

[edit snmp trap-group {group-name}]		
"categories" on page 3219	"destination-port (SNMP)" on page 5488	"logical-system (SNMP)" on page 10075
"routing-instance (SNMP)" on page 15918	"targets" on page 28970	"version (SNMP)" on page 30721

[trap-options]

[edit snmp trap-options]		
"agent-address" on page 2044	"enterprise-oid" on page 6412	"logical-system (SNMP)" on page 10075
"routing-instance (SNMP)" on page 15918	"source-address (SNMP)" on page 28174	

[v3]

[edit snmp v3]		
"notify" on page 11701	"notify-filter (Configuring the Profile Name)" on page 11704	"security-name (Community String)" on page 16500
"snmp-community" on page 28065	"tag" on page 28904	"tag" on page 28904
"target-address (SNMP)" on page 28946	"target-parameters" on page 28951	"usm" on page 30610
"vacm" on page 30642		

[edit snmp v3 notify {name}]		
"tag (Configuring Notification Targets)" on page 28910	"type (SNMP v3)" on page 30293	

[edit snmp v3 notify-filter {profile-name}]	
"oid (SNMPv3)" on page 11799	

[edit snmp v3 snmp-community {community-index}]		
"community-name" on page 4464	"context (SNMPv3)" on page 4615	

[edit snmp v3 target-address {target-address-name}]		
"address (SNMP)" on page 1844	"address-mask" on page 1891	"port" on page 12811
"retry-count (SNMP)" on page 15697	"routing-instance (SNMPv3)" on page 15917	"tag-list" on page 28916
"target-parameters" on page 28951	"timeout (SNMP)" on page 29365	

[edit snmp v3 target-parameters {target-parameters-name}]		
"notify-filter (Applying to the Management Target)" on page 11703	"parameters" on page 12175	

[edit snmp v3 target-parameters {target-parameter-name} parameters]		
"message-processing-model" on page 10652	"security-level (Generating SNMP Notifications)" on page 16482	"security-model (SNMP Notifications)" on page 16499
"security-name (SNMP Notifications)" on page 16503		

[edit snmp v3 usm]		
"local-engine" on page 9941	"remote-engine" on page 14318	

[edit snmp v3 usm local-engine]		
"user (SNMP v3)" on page 30588		

[edit snmp v3 usm local-engine user {username}]		
"authentication-md5" on page 2624	"authentication-none" on page 2632	"authentication-sha" on page 2653
"authentication-sha224" on page 2654	"authentication-sha256" on page 2656	"authentication-sha348" on page 2657
"authentication-sha512" on page 2658	"privacy-3des" on page 13212	"privacy-aes128" on page 13213
"privacy-des" on page 13215	"privacy-none" on page 13217	

[edit snmp v3 usm local-engine user {username} authentication-md5]		
"authentication-key (SNMP)" on page 2602	"authentication-password" on page 2643	

[edit snmp v3 usm local-engine user {username} authentication-sha]		
"authentication-key (SNMP)" on page 2602	"authentication-password" on page 2643	

[edit snmp v3 usm local-engine user {username} privacy-3des]		
"privacy-key" on page 13216	"privacy-password" on page 13218	

[edit snmp v3 usm local-engine user {username} privacy-aes128]		
"privacy-key" on page 13216	"privacy-password" on page 13218	

[edit snmp v3 usm local-engine user {username} privacy-des]		
"privacy-key" on page 13216	"privacy-password" on page 13218	

[edit snmp v3 usm remote-engine {engine-id}]		
"user (SNMP v3)" on page 30588		

[edit snmp v3 usm remote-engine {engine-id} user {username}]		
"authentication-none" on page 2632	"authentication-sha" on page 2653	"authentication-sha224" on page 2654
"authentication-sha256" on page 2656	"authentication-sha348" on page 2657	"authentication-sha512" on page 2658
"privacy-3des" on page 13212	"privacy-aes128" on page 13213	"privacy-des" on page 13215
"privacy-none" on page 13217		

[edit snmp v3 usm remote-engine {engine-id} user {username} authentication-md5]		
"authentication-key (SNMP)" on page 2602	"authentication-password" on page 2643	

[edit snmp v3 usm remote-engine {engine-id} user {username} authentication-sha]		
"authentication-key (SNMP)" on page 2602	"authentication-password" on page 2643	

[edit snmp v3 usm remote-engine {engine-id} user {username} privacy-3des]		
"privacy-key" on page 13216	"privacy-password" on page 13218	

[edit snmp v3 usm remote-engine {engine-id} user {username} privacy-aes128]		
"privacy-key" on page 13216	"privacy-password" on page 13218	

[edit snmp v3 usm remote-engine {engine-id} user {username} privacy-des]		
"privacy-key" on page 13216	"privacy-password" on page 13218	

[edit snmp v3 usm remote-engine {engine-id} user username]	
"authentication-md5" on page 2624	

[edit snmp v3 vacm]		
"access (SNMP)" on page 1702	"security-model (Group)" on page 16497	"security-to-group" on page 16515

[edit snmp v3 vacm access]	
"group (Configuring Group Name)" on page 7807	

[edit snmp v3 vacm access group {group-name} (default-context-prefix | context-prefix {context-prefix})]	
"security-model (Access Privileges)" on page 16496	

[edit snmp v3 vacm access group {group-name} (default-context-prefix | context-prefix {context-prefix}) security-model (any | usm | v1 | v2c)]	
"security-level (Defining Access Privileges)" on page 16481	

[edit snmp v3 vacm access group {group-name} (default-context-prefix | context-prefix {context-prefix}) security-model (any | usm | v1 | v2c) security-level (authentication | none | privacy)]		
"notify-view" on page 11710	"read-view" on page 14089	"write-view" on page 31148

[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c)]	
"security-name (Security Group)" on page 16501	

[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c) security-name {security-name}]	
"group (Defining Access Privileges for an SNMPv3 Group)" on page 7813	

[v3target-address]

[edit snmp v3target-address {target-address-name}]	
"logical-system (SNMP)" on page 10075	

[view]

[edit snmp view {view-name}]	
"oid" on page 11796	

[edit switch-options]

IN THIS SECTION

- [interface] | 1145
- [interface-mac-limit] | 1146
- [mac-notification] | 1146
- [mac-table-size] | 1146
- [redundant-trunk-group] | 1146
- [voip] | 1147

Sub Hierarchies

"[edit switch-options interface {interface-name}]" on page 1145

"[edit switch-options interface-mac-limit {limit}]" on page 1146

"[edit switch-options mac-notification]" on page 1146

"[edit switch-options mac-table-size {limit}]" on page 1146

"[edit switch-options redundant-trunk-group]" on page 1146

"[edit switch-options voip]" on page 1147

[edit switch-options]		
"authentication-whitelist" on page 2671	"interface-mac-limit" on page 8910	"interface-mac-limit" on page 8910
"interface-shutdown-action" on page 9050	"mac-learning-priority" on page 10234	"mac-notification" on page 10249

(Continued)

[edit switch-options]		
"mac-statistics" on page 10263	"mac-statistics" on page 10263	"mac-table-size" on page 10271
"mac-table-size" on page 10271	"network-isolation-profile" on page 11314	"ovsdb-managed" on page 12088
"remote-vtep-list" on page 14362	"service-id (Switch Options)" on page 16705	"unknown-unicast-forwarding" on page 30417
"vrf-export" on page 30958	"vrf-import" on page 30960	"vrf-target" on page 30968
"vtep-source-interface" on page 30988		

[interface]

[edit switch-options interface {interface-name}]		
"interface-mac-limit" on page 8910	"interface-mac-limit" on page 8910	"mac-pinning" on page 10254

[edit switch-options interface {interface-name} interface-mac-limit {limit}]		
"mac-pinning" on page 10254	"packet-action" on page 12129	

[interface-mac-limit]

[edit switch-options interface-mac-limit {limit}]	
"packet-action" on page 12129	

[mac-notification]

[edit switch-options mac-notification]	
"notification-interval" on page 11690	

[mac-table-size]

[edit switch-options mac-table-size {limit}]	
"packet-action" on page 12129	

[redundant-trunk-group]

[edit switch-options redundant-trunk-group group {name}]	
"interface (Redundant Trunk Groups)" on page 9008	

[voip]

[edit switch-options voip interface (all | {interface-name})]	
"voip-mac-exclusive" on page 30920	

[edit system]

IN THIS SECTION

- [access-line] | 1152
- [accounting] | 1153
- [arp] | 1154
- [autoinstallation] | 1154
- [commit] | 1154
- [configuration] | 1155
- [configuration-database] | 1155
- [ddos-protection] | 1156
- [demux-options] | 1158
- [dynamic-profile-options] | 1158
- [extension] | 1159
- [extensions] | 1159
- [fib-streaming] | 1160
- [firewall-profile] | 1160
- [internet-options] | 1160
- [license] | 1160
- [login] | 1161
- [management-traffic] | 1161
- [monitor] | 1162
- [ntp] | 1162

- [packet-forwarding-options] | **1163**
- [ports] | **1164**
- [processes] | **1165**
- [radius-server] | **1168**
- [root-authentication] | **1168**
- [scripts] | **1168**
- [security-profile] | **1172**
- [services] | **1173**
- [snapshot] | **1201**
- [syslog] | **1201**
- [system-profile] | **1203**
- [tacplus-server] | **1203**
- [tracing] | **1204**

Sub Hierarchies

"[edit system access-line]" on page 1152

"[edit system accounting]" on page 1153

"[edit system arp]" on page 1154

"[edit system autoinstallation]" on page 1154

"[edit system commit]" on page 1154

"[edit system configuration]" on page 1155

"[edit system configuration-database]" on page 1155

"[edit system ddos-protection]" on page 1156

"[edit system demux-options]" on page 1158

"[edit system dynamic-profile-options]" on page 1158

"[edit system extension]" on page 1159

"[edit system extensions]" on page 1159

"[edit system fib-streaming]" on page 1160

"[edit system firewall-profile]" on page 1160

"[edit system internet-options]" on page 1160

"[edit system license]" on page 1160

"[edit system login]" on page 1161

"[edit system management-traffic]" on page 1161

"[edit system monitor]" on page 1162

"[edit system ntp]" on page 1162

"[edit system packet-forwarding-options]" on page 1163

"[edit system ports]" on page 1164

"[edit system processes]" on page 1165

"[edit system radius-server {server-address}]" on page 1168

"[edit system root-authentication]" on page 1168

"[edit system scripts]" on page 1168

"[edit system security-profile {security-profile-name}]" on page 1172

"[edit system services]" on page 1173

"[edit system snapshot]" on page 1201

"[edit system syslog]" on page 1201

"[edit system system-profile]" on page 1203

"[edit system tacplus-server {server-address}]" on page 1203

"[edit system tracing]" on page 1204

[edit system]		
"access-line (Access-Line Rate Adjustment)" on page 1669	"accounting (System)" on page 1746	"allow-6pe-traceroute" on page 2162

(Continued)

[edit system]		
"allow-l3vpn-traceroute-src-select" on page 2190	"allow-v4mapped-packets" on page 2211	"archival" on page 2438
"arp" on page 2458	"arp-enhanced-scale" on page 2465	"arp-system-cache-limit" on page 2480
"authentication-order (System)" on page 2641	"authentication-order (System)" on page 2641	"auto-configuration (System)" on page 2689
"auto-snapshot" on page 2741	"auto-sw-sync" on page 2743	"autoinstallation" on page 2753
"autoinstallation (JNU Satellite Devices)" on page 2755	"backup-router" on page 2812	"boot-loader-authentication" on page 3008
"commit (System)" on page 4455	"commit activate" on page 4442	"commit prepare" on page 4451
"compress-configuration-files (System)" on page 4489	"configuration-database" on page 4510	"configuration-database (Enhanced Subscriber Management)" on page 4513
"csds" on page 4729	"ddos-protection" on page 4865	"ddos-protection (DDoS)" on page 5137
"default-address-selection" on page 5182	"demux-options (All Demux Interfaces)" on page 5301	"device-authentication" on page 5563
"diag-port-authentication" on page 5682	"domain-name (System)" on page 5880	"domain-search" on page 5899
"dual-phase-bootup" on page 6047	"dynamic-profile-options" on page 6127	"exceptions" on page 6575
"export-format (JSON)" on page 6665	"extended-echo" on page 6697	"extensions" on page 6722

(Continued)

[edit system]		
"health-monitor (KHMS)" on page 8018	"host-name (System)" on page 8148	"icmp (System)" on page 8269
"inet6-backup-router" on page 8598	"internet-options" on page 9192	"kernel-replication" on page 9480
"license" on page 9737	"location (System)" on page 9993	"login" on page 10095
"management-instance" on page 10309	"master-password" on page 10365	"max-configurations-on-flash (System)" on page 10423
"mirror-flash-on-disk" on page 10771	"name-server (System Services)" on page 11179	"nd-maxucast-retry" on page 11246
"nd-override-preferred-src" on page 11248	"nd-retransmit-timer" on page 11249	"nd-system-cache-limit" on page 11251
"no-hidden-commands" on page 11469	"no-multicast-echo" on page 11514	"no-ping-record-route" on page 11546
"no-ping-time-stamp" on page 11547	"no-redirects (IPv4 Traffic)" on page 11560	"node (System)" on page 11656
"ntp" on page 11721	"ntp (QFabric)" on page 11729	"packet-forwarding-options" on page 12145
"password-options" on page 12243	"pic-console-authentication" on page 12485	"ports (System)" on page 12896
"processes" on page 13248	"proxy (System)" on page 13845	"radius-options (System)" on page 14013
"radius-server (System)" on page 14030	"regex-additive-logic" on page 14258	"resiliency" on page 15586

(Continued)

[edit system]		
"rng" on page 15812	"root-authentication" on page 15832	"saved-core-context (System)" on page 16320
"saved-core-files (System)" on page 16321	"schema" on page 16376	"scripts" on page 16392
"security-profile" on page 16507	"security-profile-resources" on page 16512	"services (system services)" on page 16802
"snapshot" on page 28055	"static-host-mapping" on page 28535	"store" on page 28615
"switchover-on-routing-crash" on page 28763	"syslog (QFabric System)" on page 28825	"syslog (System)" on page 28836
"tacplus-options" on page 28898	"tacplus-server" on page 28902	"time-zone" on page 29362
"trace" on page 29454	"traceoptions (Resiliency)" on page 29752	"tracing" on page 29987
"unattended-boot" on page 30326	"use-imported-time-zones" on page 30518	

[access-line]

[edit system access-line]		
"attributes (Access-Line Rate Adjustment)" on page 2553	"dsl (Access-Line Rate Adjustment)" on page 6025	"pon (Access-Line Rate Adjustment)" on page 12755

[accounting]

[edit system accounting]		
"destination (Accounting)" on page 5366	"traceoptions (System Accounting)" on page 29881	

[edit system accounting destination]		
"radius (System)" on page 14034	"tacplus" on page 28896	

[edit system accounting destination radius]
"server (RADIUS Accounting)" on page 16656

[edit system accounting destination radius server {server-address}]
"source-address (RADIUS, System Logging, or TACACS+)" on page 28148

[edit system accounting destination tacplus]
"server (TACACS+ Accounting)" on page 16674

[edit system accounting destination tacplus server {server-address}]
"source-address (RADIUS, System Logging, or TACACS+)" on page 28148

[arp]

[edit system arp]		
"non-subscriber-no-reply" on page 11665	"proactive-arp-detection" on page 13221	"purging" on page 13872

[autoinstallation]

[edit system autoinstallation]		
"configuration-servers" on page 4517	"delete-after-commit (JNU Satellites)" on page 5278	"interfaces (Autoinstallation)" on page 9122
"usb" on page 30509		

[commit]

[edit system commit]		
"commit-synchronize-server" on page 4452	"notification (Commit)" on page 11687	"peers (Commit)" on page 12393
"peers-synchronize " on page 12394	"server (Batch Commits)" on page 16613	"synchronize " on page 28804

[edit system commit server]	
"traceoptions (Batch Commits)" on page 29499	

[edit system commit synchronize server]	
"traceoptions (Batch Commits)" on page 29499	

[configuration]

[edit system configuration input format]	
"json" on page 9436	

[configuration-database]

[edit system configuration-database]		
"ephemeral" on page 6423	"max-db-size (Enhanced Subscriber Management)" on page 10432	

[edit system configuration-database ephemeral]		
"instance (Ephemeral Database)" on page 8756	"resize-ephemeral-default" on page 15592	

[edit system configuration-database ephemeral instance {instance-name}]	
"resize (Ephemeral Database)" on page 15590	

[ddos-protection]

[edit system ddos-protection]		
"global (DDoS)" on page 7703	"protocols" on page 13477	"protocols (DDoS)" on page 13751
"protocols (DDoS) (ACX Series, PTX Series, and QFX Series)" on page 13765	"traceoptions (DDoS)" on page 29515	

[edit system ddos-protection global]		
"disable-fpc (DDoS)" on page 5758	"disable-logging (DDoS)" on page 5769	"disable-routing-engine (DDoS)" on page 5802
"flow-detection (DDoS Flow Detection)" on page 7278	"flow-detection-mode (DDoS Global Flow Detection)" on page 7283	"flow-level-control (DDoS Global Flow Detection)" on page 7319
"flow-report-rate (DDoS Flow Detection)" on page 7331	"violation-report-rate (DDoS Flow Detection)" on page 30737	

[edit system ddos-protection protocols {protocol-group} {packet-type}]		
"flow-detection (DDoS Packet Level)" on page 7279	"flow-detection-mode (DDoS Flow Detection)" on page 7281	"flow-level-bandwidth (DDoS Flow Detection)" on page 7316
"flow-level-control (DDoS Flow Detection)" on page 7317	"flow-level-detection (DDoS Flow Detection)" on page 7320	"flow-recover-time (DDoS Flow Detection)" on page 7330
"flow-timeout-time (DDoS Flow Detection)" on page 7359	"l2tp (DDoS Protection) " on page 9544	"no-flow-logging (DDoS Flow Detection)" on page 11463
"ppp" on page 12931	"pppoe (DDoS Protection) " on page 12962	"sample" on page 16285

(Continued)

[edit system ddos-protection protocols {protocol-group} {packet-type}]		
"timeout-active-flows (DDoS Flow Detection)" on page 29372		

[edit system ddos-protection protocols {protocol-group} (aggregate | {packet-type}) fpc {slot-number}]		
"bandwidth-scale (DDoS)" on page 2875	"burst-scale (DDoS)" on page 3077	"disable-fpc (DDoS)" on page 5758

[edit system ddos-protection protocols {protocol-group} (aggregate | {packet-type})]		
"disable-fpc (DDoS)" on page 5758	"disable-logging (DDoS)" on page 5769	"disable-routing-engine (DDoS)" on page 5802
"recover-time (DDoS)" on page 14141		

[edit system ddos-protection protocols {protocol-group} aggregate]		
"disable-routing-engine (DDoS)" on page 5802		

[edit system ddos-protection protocols {protocol-group} {packet-type} flow-detection]		
"flow-detect-time (DDoS Flow Detection)" on page 7276		

[edit system ddos-protection protocols {protocol-group} {packet-type} flow-level-bandwidth]		
"logical-interface (DDoS Flow Detection)" on page 10064	"physical-interface (DDoS Flow Detection)" on page 12474	"subscriber (DDoS Flow Detection)" on page 28688

[edit system ddos-protection protocols {protocol-group} {packet-type} flow-level-control]		
"logical-interface (DDoS Flow Detection)" on page 10064	"physical-interface (DDoS Flow Detection)" on page 12474	"subscriber (DDoS Flow Detection)" on page 28688

[edit system ddos-protection protocols {protocol-group} {packet-type} flow-level-detection]		
"logical-interface (DDoS Flow Detection)" on page 10064	"physical-interface (DDoS Flow Detection)" on page 12474	"subscriber (DDoS Flow Detection)" on page 28688

[demux-options]

[edit system demux-options]	
"use-underlying-interface-mac " on page 30542	

[dynamic-profile-options]

[edit system dynamic-profile-options]	
"versioning" on page 30725	

[extension]

[edit system extension extension-service application]	
"max-datasize" on page 10429	

[extensions]

[edit system extensions]		
"extension-service (System Extensions)" on page 6717	"providers" on page 13812	

[edit system extensions extension-service]	
"application (Extensions)" on page 2316	

[edit system extensions extension-service application]		
"file (JET)" on page 7059	"start (Scripts)" on page 28496	"traceoptions (Extensions)" on page 29552

[edit system extensions extension-service application file {filename}]		
"cert-file (Scripts)" on page 3252	"checksum (Scripts)" on page 3311	"refresh (JET)" on page 14250
"refresh-from (JET)" on page 14242	"routing-instance (JET Scripts)" on page 15935	"source (JET Scripts)" on page 28245
"start (Scripts)" on page 28496	"traceoptions (Extensions)" on page 29552	

[fib-streaming]

[edit system fib-streaming]	
"prefix-list (Junos Telemetry Interface)" on page 13093	

[firewall-profile]

[edit system firewall-profile]	
"egress-l2-extended-match" on page 6212	

[internet-options]

[edit system internet-options]	
"tcp-mss (Internet Options)" on page 28996	

[license]

[edit system license]	
"traceoptions (System License)" on page 29883	

[login]

[edit system login]		
"class (Defining Login Classes)" on page 3347	"idle-timeout (System)" on page 8309	"password (Login)" on page 12235
"retry-options" on page 15712	"user (Access)" on page 30547	

[edit system login class]	
"idle-timeout (System)" on page 8309	

[edit system login user {username} authentication]	
"remote-debug-permission" on page 14315	

[management-traffic]

[edit system management-traffic protocol {protocol-name}]	
"priority (Management Traffic)" on page 13180	

[monitor]

[edit system monitor]		
"memory (System)" on page 10630	"traceoptions (System Monitor)" on page 29885	

[edit system monitor memory]		
"process (Memory)" on page 13242	"system (Memory)" on page 28853	

[edit system monitor memory process]		
"critical-event (Process)" on page 4713	"major-event (Process)" on page 10291	"minor-event (Process)" on page 10761

[edit system monitor memory system]		
"critical-event (System)" on page 4715	"major-event (System)" on page 10293	"minor-event (System)" on page 10763

[ntp]

[edit system ntp]		
"ntp threshold" on page 11731	"server (NTP)" on page 16648	

[edit system ntp nts]		
"local-certificate (System)" on page 9930	"trusted-ca" on page 30171	

[edit system ntp server {server-name} nts]	
"remote-identity" on page 14340	

[packet-forwarding-options]

[edit system packet-forwarding-options]		
"eracl-ip6-match (packet-forwarding-options)" on page 6435	"flow-export-timer" on page 7288	"hw-db-profile" on page 8216
"internal-forwarding-bandwidth" on page 9187	"ipfix-sw-mode" on page 9291	"mcast stats-enable" on page 10582
"mpls (packet-forwarding-options)" on page 10982	"multicast-statistics (packet-forwarding-options)" on page 11091	"no-ip-tos-rewrite" on page 11480
"overlay (Packet Forwarding Options)" on page 12042	"packet-format-match" on page 12143	"recycle-bandwidth-profiles" on page 14152
"reserve ingress-egress-mcast" on page 15577		

[edit system packet-forwarding-options firewall]		
"eracl-profile (packet-forwarding-options)" on page 6437	"force-fbf-terms" on page 7381	

[edit system packet-forwarding-options firewall eracl-profile eracl-scale]	
"eracl-profile (packet-forwarding-options)" on page 6437	

[edit system packet-forwarding-options hw-resource-monitor]		
"monitor-profile" on page 10912	"polling-interval (Network Management)" on page 12752	"resource-list (Network Management)" on page 15610

[edit system packet-forwarding-options mcast-repl-type]		
"ingress-egress-recycle" on page 8631	"ingress-only" on page 8637	

[edit system packet-forwarding-options system-profile]	
"vxlan-stitching" on page 31009	

[ports]

[edit system ports]		
"auxiliary" on page 2771	"console (System Ports)" on page 4567	

[edit system ports console]	
"authentication-order (System)" on page 2641	

[processes]

[edit system processes]		
"auto-configuration" on page 2688	"bootp (System Processes)" on page 3013	"dhcp-service" on page 5646
"failover (System Process)" on page 6784	"firewall-authentication-service" on page 7188	"peer-selection-service" on page 12384
"pgcp-service" on page 12468	"pic-services-logging" on page 12504	"pppoe (System Processes) " on page 12960
"process-monitor" on page 13245	"profilerd" on page 13360	"r2cp" on page 13990
"redundancy-interface-process" on page 14182	"resource-cleanup" on page 15607	"routing (System Processes)" on page 15993
"sampling" on page 16291	"secure-neighbor-discovery" on page 16439	"security-log" on page 16484
"simple-mail-client-service" on page 27967	"smg-service (Enhanced Subscriber Management)" on page 28050	"static-subscribers" on page 28581
"statistics-service" on page 28596	"subscriber-management-helper" on page 28695	"tcp-forwarding (Processes)" on page 28985
"uac-service" on page 30303	"usb-control" on page 30510	"vrrp" on page 30971
"watchdog" on page 31027	"web-management (System Processes)" on page 31078	

[edit system processes app-engine-virtual-machine-management-service]		
"traceoptions (App Engine Virtual Machine Management Service)" on page 29497		

[edit system processes app-engine-virtual-machine-management-service traceoptions]		
"file (App Engine Virtual Machine Management Service)" on page 7030	"flag (App Engine Virtual Machine Management Service)" on page 7204	"level (App Engine Virtual Machine Management Service)" on page 9719

[edit system processes dhcp-service]		
"accept-max-tcp-connections (System Process)" on page 1634	"dhcp-snooping-file" on page 5650	"interface-traceoptions (DHCP)" on page 9085
"log" on page 10002	"ltv-syslog-interval (System Process)" on page 10198	"request-max-tcp-connections (System Process)" on page 14678
"traceoptions (DHCP)" on page 29517		

[edit system processes diameter-service]		
"traceoptions (Diameter Base Protocol)" on page 29527		

[edit system processes general-authentication-service]		
"traceoptions (Address-Assignment Pool)" on page 29485	"traceoptions (General Authentication Service)" on page 29567	

[edit system processes process-monitor]		
"traceoptions (Process Monitor)" on page 29648		

[edit system processes resource-cleanup]		
"traceoptions (Resource Cleanup)" on page 29755		

[edit system processes routing bgp]		
"rib-sharding" on page 15792	"update-threading" on page 30454	

[edit system processes sbc-configuration-process]		
"traceoptions (SBC Configuration Process)" on page 29764		

[edit system processes smg-service]		
"traceoptions (Enhanced Subscriber Management)" on page 29541		

[edit system processes static-subscribers]		
"traceoptions (Static Subscribers)" on page 29873		

[edit system processes tcp-forwarding]		
"traceoptions (TCP Port Forwarding)" on page 29894		

[radius-server]

[edit system radius-server {server-address}]	
"source-address (RADIUS, System Logging, or TACACS+)" on page 28148	

[root-authentication]

[edit system root-authentication]	
"remote-debug-permission" on page 14315	

[scripts]

[edit system scripts]		
"commit (Commit Scripts)" on page 4444	"language (Scripts)" on page 9630	"load-scripts-from-flash (Scripts)" on page 9893
"op (Op Scripts)" on page 11838	"snmp (SNMP Scripts)" on page 28068	"synchronize (Scripts)" on page 28806
"translation" on page 30073		

[edit system scripts commit]		
"allow-transients" on page 2208	"dampen" on page 4811	"direct-access (Commit Scripts)" on page 5706
"file (Commit Scripts)" on page 7042	"max-datasize" on page 10429	"refresh (Commit Scripts)" on page 14234
"refresh-from (Commit Scripts)" on page 14238	"traceoptions (Commit, Op, and SNMP Scripts)" on page 29511	

[edit system scripts commit dampen]	
"dampen-options" on page 4814	

[edit system scripts commit dampen dampen-options]		
"cpu-factor (Dampen Options)" on page 4694	"line-interval (Dampen Options)" on page 9773	"time-interval (Dampen Options)" on page 29358

[edit system scripts commit file {filename}]		
"allow-transients" on page 2208	"cert-file (Scripts)" on page 3252	"checksum (Scripts)" on page 3311
"optional (Commit Scripts)" on page 11907	"refresh (Commit Scripts)" on page 14234	"refresh-from (Commit Scripts)" on page 14238
"routing-instance (System Scripts)" on page 15963	"source (Commit Scripts)" on page 28194	

[edit system scripts commit traceoptions]	
"no-remote-trace (System)" on page 11565	

[edit system scripts op]		
"allow-url-for-python (Op Scripts)" on page 2209	"dampen" on page 4811	"file (Op Scripts)" on page 7067
"max-datasize" on page 10429	"no-allow-url (Op Scripts)" on page 11392	"refresh (Op Scripts)" on page 14251
"refresh-from (Op Scripts)" on page 14244	"start (Scripts)" on page 28496	"traceoptions (Commit, Op, and SNMP Scripts)" on page 29511

[edit system scripts op dampen]	
"dampen-options" on page 4814	

[edit system scripts op dampen dampen-options]		
"cpu-factor (Dampen Options)" on page 4694	"line-interval (Dampen Options)" on page 9773	"time-interval (Dampen Options)" on page 29358

[edit system scripts op file {filename}]		
"allow-commands (Op Scripts)" on page 2169	"arguments (Op Scripts)" on page 2457	"cert-file (Scripts)" on page 3252
"checksum (Scripts)" on page 3311	"command (Op Scripts)" on page 4431	"dampen" on page 4811
"description (Op Scripts)" on page 5339	"refresh (Op Scripts)" on page 14251	"refresh-from (Op Scripts)" on page 14244
"routing-instance (System Scripts)" on page 15963	"source (Op Scripts)" on page 28255	"start (Scripts)" on page 28496

[edit system scripts op file {filename} arguments {argument-name}]		
"description (Op Scripts)" on page 5339		

[edit system scripts op file {filename} dampen]		
"dampen-options" on page 4814		

[edit system scripts op file {filename} dampen dampen-options]		
"cpu-factor (Dampen Options)" on page 4694	"line-interval (Dampen Options)" on page 9773	"time-interval (Dampen Options)" on page 29358

[edit system scripts snmp]		
"file (SNMP Scripts)" on page 7080	"max-datasize" on page 10429	"refresh (SNMP Scripts)" on page 14255
"refresh-from (SNMP Scripts)" on page 14246	"start (Scripts)" on page 28496	"traceoptions (Commit, Op, and SNMP Scripts)" on page 29511

[edit system scripts snmp file {filename}]		
"cert-file (Scripts)" on page 3252	"checksum (Scripts)" on page 3311	"oid (SNMP Scripts)" on page 11797
"python-script-user (Event and SNMP Scripts)" on page 13884	"refresh (SNMP Scripts)" on page 14255	"refresh-from (SNMP Scripts)" on page 14246
"routing-instance (System Scripts)" on page 15963	"source (SNMP Scripts)" on page 28326	"start (Scripts)" on page 28496

[edit system scripts translation]	
"max-datasize" on page 10429	

[security-profile]

[edit system security-profile {security-profile-name}]		
"address-book (System)" on page 1870	"appfw-rule" on page 2296	"appfw-rule-set" on page 2298
"auth-entry" on page 2564	"cpu" on page 4693	"dslite-softwire-initiator" on page 6033
"flow-gate" on page 7289	"flow-session" on page 7347	"idp-policy" on page 8315
"logical-system (System Security Profile)" on page 10086	"nat-cone-binding" on page 11210	"nat-destination-pool" on page 11212
"nat-destination-rule" on page 11214	"nat-interface-port-ol (System)" on page 11216	"nat-nopat-address" on page 11219
"nat-pat-address" on page 11223	"nat-pat-portnum" on page 11225	"nat-port-ol-ipnumber" on page 11227
"nat-rule-referenced-prefix (System)" on page 11229	"nat-source-pool" on page 11234	"nat-source-rule" on page 11236
"nat-static-rule" on page 11238	"policy (System Security Profile)" on page 12747	"policy-with-count" on page 12749
"root-logical-system" on page 15835	"scheduler (System Security Profile)" on page 16366	"secure-wire (System Security Profile)" on page 16443
"zone (System Security Profile)" on page 31170		

[services]

[edit system services]		
"connection-limit" on page 4537	"database-replication (Subscriber Session Database)" on page 4854	"dhcp" on page 5583
"dhcp-local-server" on page 5601	"dhcp-local-server (System Services)" on page 5613	"dhcpv6 (System Services)" on page 5677
"dns (System Services)" on page 5867	"dtcp-only (System Services)" on page 6038	"dynamic-dns" on page 6106
"Extensible Subscriber Services" on page 6712	"extension-service (System Services)" on page 6719	"finger" on page 7171
"flow-tap-dtcp" on page 7357	"ftp" on page 7617	"grpc-tunnel " on page 7923
"inband-ztp" on page 8516	"jeb" on page 9423	"netconf" on page 11290
"outbound-https" on page 11966	"outbound-ssh" on page 11972	"rate-limit" on page 14063
"remote-device-management" on page 14317	"resource-monitor" on page 15618	"rest" on page 15631
"reverse" on page 15718	"service-deployment" on page 16686	"ssh (System Services)" on page 28430
"static-subscribers (Dynamic Service Provisioning)" on page 28582	"subscriber-management" on page 28697	"tcp-forwarding (Remote Device Management)" on page 28987
"telnet (Services)" on page 29060	"tftp-server" on page 29254	"tls (NETCONF)" on page 29403
"web-management (System Services)" on page 31079	"webapi (System Services)" on page 31091	"xnm-clear-text" on page 31156

(Continued)

[edit system services]	
"xnm-ssl" on page 31158	

[edit system services database-replication]	
"traceoptions (Subscriber Session Database Replication)" on page 29878	

[edit system services dhcp]		
"boot-server (DHCP)" on page 3010	"domain-name (DHCP Local Server)" on page 5884	"domain-name (DHCP)" on page 5882
"domain-search" on page 5899	"domain-search" on page 5899	"maximum-lease-time (DHCP)" on page 10530
"name-server (System Services)" on page 11179	"name-server (System Services)" on page 11179	"next-server" on page 11358
"option (DHCP server)" on page 11895	"pool (System)" on page 12789	"propagate-settings" on page 13379
"server-identifier" on page 16639	"sip-server" on page 27979	"sip-server" on page 27979
"static-binding" on page 28528	"static-binding" on page 28528	"traceoptions (DHCP Server)" on page 29521
"wins-server (System)" on page 31130	"wins-server (System)" on page 31130	

[edit system services dhcp dhcp-local-server dual-stack-group {dual-stack-group-name}]	
"classification-key (DHCP Relay Agent)" on page 3376	

[edit system services dhcp dhcp-local-server dual-stack-group {dual-stack-group-name}],]	
"dual-stack-interface-client-limit (DHCP Local Server and Relay Agent)" on page 6058	

[edit system services dhcp pool]		
"boot-server (DHCP)" on page 3010	"domain-name (DHCP)" on page 5882	"domain-search" on page 5899
"name-server (System Services)" on page 11179	"next-server" on page 11358	"option (DHCP server)" on page 11895
"propagate-settings" on page 13379	"server-identifier" on page 16639	"sip-server" on page 27979
"wins-server (System)" on page 31130		

[edit system services dhcp static-binding]		
"boot-server (DHCP)" on page 3010	"domain-name (DHCP)" on page 5882	"domain-search" on page 5899
"name-server (System Services)" on page 11179	"next-server" on page 11358	"option (DHCP server)" on page 11895
"server-identifier" on page 16639	"sip-server" on page 27979	"wins-server (System)" on page 31130

[edit system services dhcp-local-server]		
"active-leasequery (DHCP Local Server)" on page 1806	"allow-active-leasequery (DHCP Local Server)" on page 2163	"allow-bulk-leasequery (DHCP Local Server)" on page 2165
"allow-leasequery (DHCP Local Server)" on page 2191	"authentication (DHCP Local Server)" on page 2593	"dhcpv6 (DHCP Local Server)" on page 5658
"dual-stack-group (DHCP Local Server)" on page 6052	"duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)" on page 6066	"dynamic-profile (DHCP Local Server)" on page 6117
"forward-snooped-clients (DHCP Local Server)" on page 7404	"group (DHCP Local Server)" on page 7814	"interface-traceoptions (System Services DHCP)" on page 9087
"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701	"liveness-detection" on page 9846	"no-snoop (DHCP Local Server and Relay Agent)" on page 11582
"overrides (DHCP Local Server)" on page 12065	"pool-match-order" on page 12777	"reauthenticate (DHCP Local Server)" on page 14103
"reconfigure (DHCP Local Server)" on page 14133	"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334	"requested-ip-network-match (DHCP Local Server)" on page 15572
"route-suppression (DHCP Local Server and Relay Agent)" on page 15877	"service-profile (DHCP Local Server)" on page 16729	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982

[edit system services dhcp-local-server (dhcpv6) group {group-name} interface {interface-name} overrides]	
"include-option-82 (DHCP Local Server)" on page 8534	

[edit system services dhcp-local-server (dhcpv6) group {group-name} overrides]	
"include-option-82 (DHCP Local Server)" on page 8534	

[edit system services dhcp-local-server authentication]		
"password (DHCP Local Server)" on page 12226	"username-include (DHCP Local Server)" on page 30595	

[edit system services dhcp-local-server authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"delimiter (DHCP Local Server)" on page 5288	"domain-name (DHCP Local Server)" on page 5884
"interface-description (DHCP Local Server)" on page 8854	"interface-name (DHCP Local Server)" on page 8937	"logical-system-name (DHCP Local Server)" on page 10082
"mac-address (DHCP Local Server)" on page 10221	"option-60 (DHCP Local Server)" on page 11876	"option-82 (DHCP Local Server Authentication)" on page 11890
"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578	

[edit system services dhcp-local-server dhcp authentication username-include]	
"client-id (DHCP Local Server)" on page 4315	

[edit system services dhcp-local-server dhcp group {group-name} authentication username-include]	
"client-id (DHCP Local Server)" on page 4315	

[edit system services dhcp-local-server dhcpv6]		
"active-leasequery (DHCP Local Server)" on page 1806	"allow-active-leasequery (DHCP Local Server)" on page 2163	"allow-bulk-leasequery (DHCP Local Server)" on page 2165
"allow-leasequery (DHCP Local Server)" on page 2191	"authentication (DHCP Local Server)" on page 2593	"duplicate-clients (DHCPv6 Local Server and Relay Agent)" on page 6064
"dynamic-profile (DHCP Local Server)" on page 6117	"dynamic-server" on page 6154	"group (DHCP Local Server)" on page 7814
"group (System Services DHCP)" on page 7899	"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701	"liveness-detection" on page 9846
"no-snoop (DHCP Local Server and Relay Agent)" on page 11582	"overrides (DHCP Local Server)" on page 12065	"overrides (System Services DHCP)" on page 12077
"password (DHCP Local Server)" on page 12226	"reauthenticate (DHCP Local Server)" on page 14103	"reconfigure (DHCP Local Server)" on page 14133
"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334	"requested-ip-network-match (DHCP Local Server)" on page 15572	"route-suppression (DHCP Local Server and Relay Agent)" on page 15877
"server-duid-type (DHCP Local Server)" on page 16626	"service-profile (DHCP Local Server)" on page 16729	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982

[edit system services dhcp-local-server dhcpv6 authentication]	
"username-include (DHCP Local Server)" on page 30595	

[edit system services dhcp-local-server dhcpv6 authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"client-id (DHCP Local Server)" on page 4315	"delimiter (DHCP Local Server)" on page 5288
"domain-name (DHCP Local Server)" on page 5884	"interface-description (DHCP Local Server)" on page 8854	"interface-name (DHCP Local Server)" on page 8937
"logical-system-name (DHCP Local Server)" on page 10082	"mac-address (DHCP Local Server)" on page 10221	"relay-agent-interface-id (DHCP Local Server)" on page 14275
"relay-agent-remote-id (DHCP Local Server)" on page 14282	"relay-agent-subscriber-id (DHCP Local Server)" on page 14288	"routing-instance-name (DHCP Local Server)" on page 15945
"user-prefix (DHCP Local Server)" on page 30578		

[edit system services dhcp-local-server dhcpv6 group {group-name}]		
"authentication (DHCP Local Server)" on page 2593	"dynamic-profile (DHCP Local Server)" on page 6117	"interface (DHCP Local Server)" on page 8860
"interface (System Services DHCP)" on page 9079	"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701	"liveness-detection" on page 9846
"overrides (DHCP Local Server)" on page 12065	"overrides (System Services DHCP)" on page 12077	"reconfigure (DHCP Local Server)" on page 14133
"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334	"route-suppression (DHCP Local Server and Relay Agent)" on page 15877	"service-profile (DHCP Local Server)" on page 16729
"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982		

[edit system services dhcp-local-server dhcpv6 group {group-name} authentication]		
"password (DHCP Local Server)" on page 12226	"username-include (DHCP Local Server)" on page 30595	

[edit system services dhcp-local-server dhcpv6 group {group-name} authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"client-id (DHCP Local Server)" on page 4315	"delimiter (DHCP Local Server)" on page 5288
"domain-name (DHCP Local Server)" on page 5884	"interface-description (DHCP Local Server)" on page 8854	"interface-name (DHCP Local Server)" on page 8937
"logical-system-name (DHCP Local Server)" on page 10082	"mac-address (DHCP Local Server)" on page 10221	"relay-agent-interface-id (DHCP Local Server)" on page 14275
"relay-agent-remote-id (DHCP Local Server)" on page 14282	"relay-agent-subscriber-id (DHCP Local Server)" on page 14288	"routing-instance-name (DHCP Local Server)" on page 15945
"user-prefix (DHCP Local Server)" on page 30578		

[edit system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name}]		
"dynamic-profile (DHCP Local Server)" on page 6117	"overrides (DHCP Local Server)" on page 12065	"overrides (System Services DHCP)" on page 12077
"service-profile (DHCP Local Server)" on page 16729	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982	"trace (DHCP Local Server)" on page 29456

[edit system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name} overrides]		
"client-negotiation-match (DHCPv6 Local Server)" on page 4333	"delay-advertise (DHCPv6)" on page 5239	"delegated-pool (DHCP Local Server)" on page 5271

(Continued)

[edit system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name} overrides]		
"interface-client-limit (DHCP Local Server)" on page 9106	"multi-address-embedded-option-response (DHCP Local Server)" on page 11039	"process-inform" on page 13240
"rapid-commit (DHCPv6 Local Server)" on page 14055		

[edit system services dhcp-local-server dhcpv6 group {group-name} interface {interface-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit system services dhcp-local-server dhcpv6 group {group-name} lease-time-validation]		
"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699	"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	

[edit system services dhcp-local-server dhcpv6 group {group-name} liveness-detection]		
"failure-action" on page 6787	"method" on page 10663	

[edit system services dhcp-local-server dhcpv6 group {group-name} liveness-detection method]		
"bfd" on page 2909	"layer2-liveness-detection (Send)" on page 9649	

[edit system services dhcp-local-server dhcpv6 group {group-name} liveness-detection method bfd]		
"detection-time" on page 5542	"holddown-interval" on page 8122	"minimum-interval" on page 10733

(Continued)

[edit system services dhcp-local-server dhcpv6 group {group-name} liveness-detection method bfd]		
"minimum-receive-interval" on page 10752	"multiplier" on page 11122	"no-adaptation" on page 11373
"session-mode" on page 16828	"transmit-interval" on page 30085	"version (BFD)" on page 30695

[edit system services dhcp-local-server dhcpv6 group {group-name} liveness-detection method bfd detection-time]	
"threshold (detection-time)" on page 29323	

[edit system services dhcp-local-server dhcpv6 group {group-name} liveness-detection method bfd transmit-interval]		
"minimum-interval" on page 10733	"threshold (transmit-interval)" on page 29344	

[edit system services dhcp-local-server dhcpv6 group {group-name} overrides]		
"client-negotiation-match (DHCPv6 Local Server)" on page 4333	"delay-advertise (DHCPv6)" on page 5239	"delegated-pool (DHCP Local Server)" on page 5271
"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279	"interface-client-limit (DHCP Local Server)" on page 9106	"multi-address-embedded-option-response (DHCP Local Server)" on page 11039
"process-inform" on page 13240	"rapid-commit (DHCPv6 Local Server)" on page 14055	

[edit system services dhcp-local-server dhcpv6 group {group-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit system services dhcp-local-server dhcpv6 group {group-name} reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"strict (DHCP Local Server)" on page 28646
"timeout (DHCP Local Server)" on page 29374	"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160

[edit system services dhcp-local-server dhcpv6 group {group-name} reconfigure trigger]	
"radius-disconnect (DHCP Local Server)" on page 14003	

[edit system services dhcp-local-server dhcpv6 lease-time-validation]		
"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699	"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	

[edit system services dhcp-local-server dhcpv6 liveness-detection]		
"failure-action" on page 6787	"method" on page 10663	

[edit system services dhcp-local-server dhcpv6 liveness-detection method]		
"bfd" on page 2909	"layer2-liveness-detection (Send)" on page 9649	

[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd]		
"detection-time" on page 5542	"holddown-interval" on page 8122	"minimum-interval" on page 10733
"minimum-receive-interval" on page 10752	"multiplier" on page 11122	"no-adaptation" on page 11373

(Continued)

[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd]		
"session-mode" on page 16828	"transmit-interval" on page 30085	"version (BFD)" on page 30695

[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd detection-time]	
"threshold (detection-time)" on page 29323	

[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval]		
"minimum-interval" on page 10733	"threshold (transmit-interval)" on page 29344	

[edit system services dhcp-local-server dhcpv6 overrides]		
"client-negotiation-match (DHCPv6 Local Server)" on page 4333	"delay-advertise (DHCPv6)" on page 5239	"delegated-pool (DHCP Local Server)" on page 5271
"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279	"include-option-82 (DHCP Local Server)" on page 8534	"interface-client-limit (DHCP Local Server)" on page 9106
"multi-address-embedded-option-response (DHCP Local Server)" on page 11039	"process-inform" on page 13240	"rapid-commit (DHCPv6 Local Server)" on page 14055

[edit system services dhcp-local-server dhcpv6 overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit system services dhcp-local-server dhcpv6 reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"strict (DHCP Local Server)" on page 28646
"timeout (DHCP Local Server)" on page 29374	"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160

[edit system services dhcp-local-server dhcpv6 reconfigure trigger]	
"radius-disconnect (DHCP Local Server)" on page 14003	

[edit system services dhcp-local-server dual-stack-group {dual-stack-group-name}]		
"authentication (DHCP Local Server)" on page 2593	"classification-key (DHCP Local Server)" on page 3375	"dynamic-profile (DHCP Local Server)" on page 6117
"liveness-detection" on page 9846	"on-demand-address-allocation" on page 11813	"protocol-master" on page 13448
"reauthenticate (DHCP Local Server)" on page 14103	"service-profile (DHCP Local Server)" on page 16729	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982

[edit system services dhcp-local-server dual-stack-group <group-name> authentication username-include]	
"client-id (DHCP Local Server)" on page 4315	

[edit system services dhcp-local-server dual-stack-group {group-name} authentication]	
"username-include (DHCP Local Server)" on page 30595	

[edit system services dhcp-local-server dual-stack-group {dual-stack-group-name} liveness-detection]	
"method" on page 10663	

[edit system services dhcp-local-server dual-stack-group {dual-stack-group-name} liveness-detection method]	
"layer2-liveness-detection (Send)" on page 9649	

[edit system services dhcp-local-server dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Local Server)" on page 2056	"use-primary (DHCP Local Server)" on page 30535	

[edit system services dhcp-local-server group {group-name}]		
"authentication (DHCP Local Server)" on page 2593	"dynamic-profile (DHCP Local Server)" on page 6117	"interface (DHCP Local Server)" on page 8860
"interface-tag (DHCP Local Server)" on page 9081	"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701	"liveness-detection" on page 9846
"overrides (DHCP Local Server)" on page 12065	"reconfigure (DHCP Local Server)" on page 14133	"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334
"route-suppression (DHCP Local Server and Relay Agent)" on page 15877	"service-profile (DHCP Local Server)" on page 16729	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982

[edit system services dhcp-local-server group {group-name} authentication]		
"password (DHCP Local Server)" on page 12226	"username-include (DHCP Local Server)" on page 30595	

[edit system services dhcp-local-server group {group-name} authentication username-include]		
"circuit-type (DHCP Local Server)" on page 3330	"delimiter (DHCP Local Server)" on page 5288	"domain-name (DHCP Local Server)" on page 5884
"interface-description (DHCP Local Server)" on page 8854	"interface-name (DHCP Local Server)" on page 8937	"logical-system-name (DHCP Local Server)" on page 10082
"mac-address (DHCP Local Server)" on page 10221	"option-60 (DHCP Local Server)" on page 11876	"option-82 (DHCP Local Server Authentication)" on page 11890
"routing-instance-name (DHCP Local Server)" on page 15945	"user-prefix (DHCP Local Server)" on page 30578	

[edit system services dhcp-local-server group {group-name} dynamic-profile {profile-name}]		
"aggregate-clients (DHCP Local Server)" on page 2056	"use-primary (DHCP Local Server)" on page 30535	

[edit system services dhcp-local-server group {group-name} interface {interface-name}]		
"dynamic-profile (DHCP Local Server)" on page 6117	"overrides (DHCP Local Server)" on page 12065	"service-profile (DHCP Local Server)" on page 16729
"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982	"trace (DHCP Local Server)" on page 29456	

[edit system services dhcp-local-server group {group-name} interface {interface-name} overrides]		
"client-discover-match (DHCP Local Server)" on page 4308	"delay-offer (DHCPv4)" on page 5254	"dual-stack (DHCP Local Server Overrides)" on page 6049
"interface-client-limit (DHCP Local Server)" on page 9106	"process-inform" on page 13240	

[edit system services dhcp-local-server group {group-name} interface {interface-name} overrides process-inform]		
"pool (DHCP Local Server Overrides)" on page 12774		

[edit system services dhcp-local-server group {group-name} lease-time-validation]		
"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699	"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	

[edit system services dhcp-local-server group {group-name} liveness-detection]		
"failure-action" on page 6787	"method" on page 10663	

[edit system services dhcp-local-server group {group-name} liveness-detection method]		
"bfd" on page 2909	"layer2-liveness-detection (Send)" on page 9649	

[edit system services dhcp-local-server group {group-name} liveness-detection method bfd]		
"detection-time" on page 5542	"holddown-interval" on page 8122	"minimum-interval" on page 10733
"minimum-receive-interval" on page 10752	"multiplier" on page 11122	"no-adaptation" on page 11373
"session-mode" on page 16828	"transmit-interval" on page 30085	"version (BFD)" on page 30695

[edit system services dhcp-local-server group {group-name} liveness-detection method bfd detection-time]		
"threshold (detection-time)" on page 29323		

[edit system services dhcp-local-server group {group-name} liveness-detection method bfd transmit-interval]		
"minimum-interval" on page 10733	"threshold (transmit-interval)" on page 29344	

[edit system services dhcp-local-server group {group-name} overrides]		
"client-discover-match (DHCP Local Server)" on page 4308	"delay-offer (DHCPv4)" on page 5254	"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279
"dual-stack (DHCP Local Server Overrides)" on page 6049	"interface-client-limit (DHCP Local Server)" on page 9106	"process-inform" on page 13240

[edit system services dhcp-local-server group {group-name} overrides process-inform]	
"pool (DHCP Local Server Overrides)" on page 12774	

[edit system services dhcp-local-server group {group-name} reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"timeout (DHCP Local Server)" on page 29374
"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160	

[edit system services dhcp-local-server group {group-name} reconfigure trigger]	
"radius-disconnect (DHCP Local Server)" on page 14003	

[edit system services dhcp-local-server lease-time-validation]		
"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699	"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	

[edit system services dhcp-local-server liveness-detection]		
"failure-action" on page 6787	"method" on page 10663	

[edit system services dhcp-local-server liveness-detection method]		
"bfd" on page 2909		

[edit system services dhcp-local-server liveness-detection method bfd]		
"detection-time" on page 5542	"holddown-interval" on page 8122	"minimum-interval" on page 10733
"minimum-receive-interval" on page 10752	"multiplier" on page 11122	"no-adaptation" on page 11373
"session-mode" on page 16828	"transmit-interval" on page 30085	"version (BFD)" on page 30695

[edit system services dhcp-local-server liveness-detection method bfd detection-time]		
"threshold (detection-time)" on page 29323		

[edit system services dhcp-local-server liveness-detection method bfd transmit-interval]		
"minimum-interval" on page 10733	"threshold (transmit-interval)" on page 29344	

[edit system services dhcp-local-server liveness-detection method].]		
"layer2-liveness-detection (Send)" on page 9649		

[edit system services dhcp-local-server overrides]		
"client-discover-match (DHCP Local Server)" on page 4308	"delay-offer (DHCPv4)" on page 5254	"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279
"dual-stack (DHCP Local Server Overrides)" on page 6049	"include-option-82 (DHCP Local Server)" on page 8534	"interface-client-limit (DHCP Local Server)" on page 9106
"process-inform" on page 13240		

[edit system services dhcp-local-server overrides process-inform]		
"pool (DHCP Local Server Overrides)" on page 12774		

[edit system services dhcp-local-server pool-match-order]		
"external-authority" on page 6726	"ip-address-first" on page 9235	"option-82 (DHCP Local Server Pool Matching)" on page 11892

[edit system services dhcp-local-server reconfigure]		
"attempts (DHCP Local Server)" on page 2549	"clear-on-abort (DHCP Local Server)" on page 3782	"timeout (DHCP Local Server)" on page 29374
"token (DHCP Local Server)" on page 29428	"trigger (DHCP Local Server)" on page 30160	

[edit system services dhcp-local-server reconfigure trigger]		
"radius-disconnect (DHCP Local Server)" on page 14003		

[edit system services dns]		
"dnssec" on page 5869	"traceoptions (System Services DNS)" on page 29888	

[edit system services dns dns-proxy]		
"dns-proxy" on page 5862		

[edit system services dns dns-proxy view {view-name} domain {domain-name}]		
"forward-only (DNS)" on page 7400		

[edit system services dns dnssec]		
"dlv" on page 5844	"trusted-keys (DNSSEC)" on page 30175	

[edit system services extension-service]		
"notification (System Services)" on page 11697	"request-response" on page 14741	"traceoptions (Services)" on page 29836

[edit system services extension-service request-response]		
"grpc" on page 7920		

[edit system services extension-service request-response grpc]		
"max-connections" on page 10426	"ssl" on page 28439	

[edit system services finger]		
"connection-limit" on page 4537	"rate-limit" on page 14063	

[edit system services ftp]		
"authentication-order (System)" on page 2641	"connection-limit" on page 4537	"rate-limit" on page 14063

[edit system services grpc request-response grpc ssl]	
"local-certificate (System)" on page 9930	

[edit system services local-policy-decision-function statistics]	
"aacl-statistics-profile" on page 1623	

[edit system services local-policy-decision-function statistics aacl-statistics-profile {profile-name}]	
"aacl-fields" on page 1621	

[edit system services netconf]		
"hello-message (NETCONF)" on page 8039	"netconf-monitoring (NETCONF)" on page 11293	"notification (NETCONF)" on page 11691
"rfc-compliant (NETCONF)" on page 15755	"ssh (NETCONF)" on page 28428	"traceoptions (NETCONF and Junos XML Protocol)" on page 29613

(Continued)

[edit system services netconf]		
"yang-compliant (NETCONF)" on page 31163	"yang-modules (NETCONF)" on page 31166	

[edit system services netconf ssh]		
"connection-limit" on page 4537	"port (NETCONF)" on page 12844	"rate-limit" on page 14063

[edit system services netconf tls]		
"client-identity (NETCONF TLS)" on page 4324	"default-client-identity (NETCONF TLS)" on page 5185	"traceoptions (NETCONF TLS)" on page 29616

[edit system services outbound-ssh]	
"traceoptions (Outbound SSH)" on page 29627	

[edit system services packet-triggered-subscribers]	
"subscriber-packet-idle-timeout" on page 28699	

[edit system services remote-device-management]	
"service-device (Remote Device Management)" on page 16690	

[edit system services remote-device-management service-device {device-name}]		
"access-domain (Remote Device Management)" on page 1657	"provisioning-method (Remote Device Management)" on page 13814	

[edit system services remote-device-management service-device {device-name} provisioning-method]		
"netconf (Remote Device Management)" on page 11295		

[edit system services resource-monitor]		
"free-fw-memory-watermark (Resource Monitor)" on page 7576	"free-heap-memory-watermark (Resource Monitor)" on page 7578	"free-nh-memory-watermark (Resource Monitor)" on page 7579
"high-cos-queue-threshold" on page 8074	"high-threshold (Resource Monitor)" on page 8078	"no-load-throttle (Resource Monitor)" on page 11496
"no-logging (Resource Monitor)" on page 11503	"no-throttle (Resource Monitor)" on page 11598	"resource-category jtree (Resource Monitor)" on page 15605
"subscribers-limit (Resource Monitor)" on page 28707	"traceoptions (Resource Monitor)" on page 29758	

[edit system services resource-monitor subscribers-limit]		
"client-type (Subscriber Limits)" on page 4342		

[edit system services resource-monitor subscribers-limit client-type {name}]		
"chassis (Subscriber Limits)" on page 3306	"fpc (Subscriber Limits)" on page 7533	

[edit system services resource-monitor subscribers-limit client-type {name} fpc {slot-number}]		
"pic (Subscriber Limits)" on page 12506		

[edit system services resource-monitor subscribers-limit client-type {name} fpc {slot-number} pic {number}]		
"port (Subscriber Limits)" on page 12890		

[edit system services rest]		
"control (REST API)" on page 4639	"enable-explorer (REST API)" on page 6259	"http (REST API)" on page 8202
"https (REST API)" on page 8207	"traceoptions (REST API)" on page 29760	

[edit system services rest control]		
"allowed-sources (REST API)" on page 2216	"connection-limit (REST API)" on page 4538	

[edit system services rest http]		
"port (REST API)" on page 12866		

[edit system services rest http] , [edit system services rest https]		
"addresses (REST API)" on page 1938		

[edit system services rest https]		
"ca-chain (REST API)" on page 3154	"cipher-list (REST API)" on page 3318	"mutual-authentication (REST API)" on page 11134
"port (REST API)" on page 12866	"server-certificate (REST API)" on page 16615	

[edit system services rest https mutual-authentication]		
"certificate-authority (REST API)" on page 3256		

[edit system services service-deployment]		
"local-certificate (System)" on page 9930	"servers" on page 16677	"source-address (SRC Software)" on page 28175

[edit system services service-deployment servers {server-address}]		
"port (SRC Server)" on page 12887		

[edit system services ssh]		
"authentication-order (System)" on page 2641	"connection-limit" on page 4537	"hostkey-algorithm" on page 8179
"hostkey-algorithm-list" on page 8182	"key-exchange" on page 9491	"max-sessions-per-connection" on page 10471
"rate-limit" on page 14063	"root-login (JDM)" on page 15837	

[edit system services static-subscribers]		
"access-profile (Static Subscribers)" on page 1698	"authentication (Static Subscribers)" on page 2665	"group (Static Subscribers)" on page 7897
"service-profile (Static Subscribers)" on page 16735		

[edit system services static-subscribers authentication]		
"password (Static Subscribers)" on page 12247	"username-include (Static Subscribers)" on page 30604	

[edit system services static-subscribers authentication username-include]		
"domain-name (Static Subscribers)" on page 5892	"interface (Static Subscriber Username)" on page 9070	"logical-system-name (Static Subscribers)" on page 10085
"routing-instance-name (Static Subscribers)" on page 15949	"user-prefix (Static Subscribers)" on page 30582	

[edit system services static-subscribers dynamic-profile {profile-name}]	
"aggregate-clients (Static Subscribers)" on page 2060	

[edit system services static-subscribers group {group-name}]		
"access-profile (Static Subscribers)" on page 1698	"authentication (Static Subscribers)" on page 2665	"interface (Static Subscriber Group)" on page 9068
"service-profile (Static Subscribers)" on page 16735		

[edit system services static-subscribers group {group-name} authentication]		
"password (Static Subscribers)" on page 12247	"username-include (Static Subscribers)" on page 30604	

[edit system services static-subscribers group {group-name} authentication username-include]		
"domain-name (Static Subscribers)" on page 5892	"interface (Static Subscriber Username)" on page 9070	"logical-system-name (Static Subscribers)" on page 10085

(Continued)

[edit system services static-subscribers group {group-name} authentication username-include]		
"routing-instance-name (Static Subscribers)" on page 15949	"user-prefix (Static Subscribers)" on page 30582	

[edit system services static-subscribers group {group-name} dynamic-profile {profile-name}]	
"aggregate-clients (Static Subscribers)" on page 2060	

[edit system services subscriber-management]		
"enable (Enhanced Subscriber Management)" on page 6257	"enforce-strict-scale-limit-license (Subscriber Management)" on page 6345	"gres-route-flush-delay (Subscriber Management)" on page 7800
"maintain-subscriber (Subscriber Management)" on page 10284	"overrides (Enhanced Subscriber Management)" on page 12074	"redundancy (M:N Subscriber Redundancy)" on page 14186
"traceoptions (Subscriber Management)" on page 29876		

[edit system services subscriber-management maintain-subscriber]	
"interface-delete (Subscriber Management or DHCP Client Management)" on page 8853	

[edit system services subscriber-management overrides]		
"event (Enhanced Subscriber Management)" on page 6523	"no-unsolicited-ra (Enhanced Subscriber Management)" on page 11611	"shmlog (Shared Memory Log)" on page 16978

[edit system services subscriber-management overrides interfaces family (inet | inet6)]		
"layer2-liveness-detection (Receive)" on page 9647		

[edit system services telnet]		
"authentication-order (System)" on page 2641	"connection-limit" on page 4537	"rate-limit" on page 14063

[edit system services web-management]		
"http (Web Management)" on page 8205	"https (Web Management)" on page 8211	"session (System Services)" on page 16853
"session (Web Management)" on page 16857		

[edit system services web-management session]		
"idle-timeout (System Services)" on page 8311	"session-limit (System Services)" on page 16824	

[edit system services webapi (System Services)]		
"client (System Services)" on page 4339	"debug-level (System Services)" on page 5155	"debug-log (System Services)" on page 5157
"http (System Services)" on page 8204	"https (System Services)" on page 8208	"user (System Services)" on page 30591

[edit system services xnm-clear-text]		
"connection-limit" on page 4537	"rate-limit" on page 14063	

[edit system services xnm-ssl]		
"connection-limit" on page 4537	"rate-limit" on page 14063	

[snapshot]

[edit system snapshot]	
"recovery" on page 14143	

[syslog]

[edit system syslog]		
"allow-duplicates" on page 2172	"console (System Logging)" on page 4565	"file (QFabric System)" on page 7071
"file (System Logging)" on page 7085	"grpc-replay" on page 7922	"host (System)" on page 8171
"log-rotate-frequency" on page 10025	"routing-instance (Syslog)" on page 15962	"shell" on page 16975
"source-address (RADIUS, System Logging, or TACACS+)" on page 28148	"source-address (System Logging)" on page 28177	"time-format" on page 29356
"user (System Logging)" on page 30589		

[edit system syslog archive]		
"files (syslog)" on page 7108	"size (System)" on page 28007	"world-readable (System Syslog)" on page 31144

[edit system syslog console]	
"application-status" on page 2392	

[edit system syslog file {file-name}]		
"allow-duplicates" on page 2172	"archive (All System Log Files)" on page 2440	"archive (Individual System Log File)" on page 2442
"archive (QFabric System)" on page 2445	"explicit-priority" on page 6657	"file (System Logging)" on page 7086
"match (Syslog)" on page 10369	"match-strings" on page 10407	"structured-data" on page 28664

[edit system syslog file {filename} archive]		
"files (syslog)" on page 7108	"size (System)" on page 28007	"world-readable (System Syslog)" on page 31144

[edit system syslog host {host-name}]		
"allow-duplicates" on page 2172	"exclude-hostname" on page 6608	"explicit-priority" on page 6657
"file (System Logging)" on page 7086	"log-prefix (System)" on page 10023	"match-strings" on page 10407
"routing-instance (Syslog)" on page 15962	"transport (System Syslog)" on page 30112	

[edit system syslog host {hostname} | other-routing-engine| scc-master)]		
"match (Syslog)" on page 10369	"port (Syslog)" on page 12891	

[edit system syslog user {user-name}]		
"allow-duplicates" on page 2172	"file (System Logging)" on page 7086	"match-strings" on page 10407

[edit system syslog user ({username} | *)]	
"match (Syslog)" on page 10369	

[system-profile]

[edit system system-profile]	
"vxlan-extended" on page 30998	

[tacplus-server]

[edit system tacplus-server {server-address}]	
"source-address (RADIUS, System Logging, or TACACS+)" on page 28148	

[tracing]

[edit system tracing]	
"destination-override" on page 5454	

[edit tenants]

IN THIS SECTION

- [<tenant-name>] | 1205
- [access] | 1206
- [applications] | 1206
- [class-of-service] | 1206
- [interfaces] | 1207
- [name] | 1208
- [nat] | 1208
- [routing-instances] | 1208
- [security] | 1219
- [security-profile] | 1230
- [securitypolicies] | 1230
- [services] | 1231
- [system] | 1231

Sub Hierarchies

"[edit tenants <tenant-name>]" on page 1205

"[edit tenants {name} access]" on page 1206

"[edit tenants applications]" on page 1206

"[edit tenants class-of-service]" on page 1206

"[edit tenants interfaces]" on page 1207

"[edit tenants name]" on page 1208

"[edit tenants nat]" on page 1208

"[edit tenants {tenant-name} routing-instances {mac-vrf-routing-instance-name}]" on page 1208

"[edit tenants {tenant-name} security]" on page 1219

"[edit tenants {tenant-name} security-profile]" on page 1230

"[edit tenants securitypolicies]" on page 1230

"[edit tenants {tenants_name} services]" on page 1231

"[edit tenants system]" on page 1231

[edit tenants {tenant-name}]	
"security-profile" on page 16507	

[<tenant-name>]

[edit tenants <tenant-name> schedulers]	
"scheduler (Security Policies)" on page 16364	

[access]

[edit tenants {name} access]		
"firewall-authentication (tenants)" on page 7189		

[edit tenants {name} access firewall-authentication]		
"pass-through (Firewall-authentication)" on page 12196	"web-authentication (Access)" on page 31058	

[applications]

[edit tenants {name} applications application]	
"application-protocol (Applications)" on page 2355	

[class-of-service]

[edit tenants {tenant-name} class-of-service application-traffic-control]		
"rate-limiters" on page 14071	"rule-sets (CoS AppQoS)" on page 16149	

[interfaces]

[edit tenants {tenant-name} interfaces {interface-name} unit {logical-unit-number} family inet]	
"dhcp (DHCP Client)" on page 5598	

[edit tenants {tenant-name} interfaces {name} unit {name} family inet6]	
"dhcpv6-client" on page 5655	

[edit tenants {tenant-name} interfaces {interface-name} unit {logical-unit-number} family inet6 dhcpv6-client]		
"client-ia-type" on page 4313	"client-identifier (DHCPv6 Client)" on page 4321	"client-type" on page 4340
"req-option" on page 14408	"retransmission-attempt (DHCPv6 Client)" on page 15678	"update-router-advertisement" on page 30449

[edit tenants {tenant-name} interfaces {interface-name} unit {logical-unit-number} family inet6 dhcpv6-client prefix-delegating]		
"preferred-prefix-length" on page 13059	"sub-prefix-length" on page 28674	

[edit tenants {tenant-name} interfaces {interface-name} unit {logical-unit-number} family inet6dhcpv6-client]	
"rapid-commit (DHCPv6 Client)" on page 14054	

[name]

[edit tenants name security nat {destination pool}]	
"address (Security Destination NAT)" on page 1924	

[nat]

[edit tenants {name} nat destination rule-set {rule-set-name} rule {rule-name} match]	
"application (Security Destination NAT)" on page 2369	

[routing-instances]

[edit tenants {tenant-name} routing-instances {mac-vrf-routing-instance-name}]		
"forwarding-instance identifier" on page 7457	"mac-vrf" on page 10276	"routing-options" on page 15980
"service-type (MAC-VRF Routing Instances)" on page 16764		

[edit tenants {name} routing-instances {name} bridge-domains {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit tenants {tenant-name} routing-instances {routing-instance-name} forwarding-options dhcp-relay dhcpv6 group {group-name} authentication username-include]	
"routing-instance-name (DHCP Relay Agent)" on page 15947	

[edit tenants {name} routing-instances {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit tenants {name} routing-instances {name} forwarding-options pfe-sensor npu-memory]	
"resource-list (Telemetry)" on page 15612	

[edit tenants {name} routing-instances {name} forwarding-options storm-control]	
"enhanced" on page 6354	

[edit tenants {name} routing-instances {name} protocols]		
"flex-algorithm (Protocols OSPF)" on page 7223	"rsvp" on page 16050	

[edit tenants {name} routing-instances {name} protocols (evpn | l2vpn | vpls) mesh-group {mesh-group-name} neighbor {neighbor-id} oam bfd-liveness-detection]	
"minimum-receive-interval (BFD Liveness Detection)" on page 10754	

[edit tenants {name} routing-instances {name} protocols (evpn | l2vpn | vpls) neighbor {neighbor-id} oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"minimum-receive-interval (BFD Liveness Detection)" on page 10754	"multiplier (BFD Liveness Detection)" on page 11123
"no-adaptation (BFD Liveness Detection)" on page 11374	"version (BFD Liveness Detection)" on page 30697	

[edit tenants {name} routing-instances {name} protocols (evpn | l2vpn | vpls) oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"minimum-receive-interval (BFD Liveness Detection)" on page 10754	"multiplier (BFD Liveness Detection)" on page 11123
"no-adaptation (BFD Liveness Detection)" on page 11374	"version (BFD Liveness Detection)" on page 30697	

[edit tenants {name} routing-instances {name} protocols (evpn | l2vpn | vpls)mesh-group {mesh-group-name} neighbor {neighbor-id} oam bfd-liveness-detection]		
"minimum-interval (BFD Liveness Detection)" on page 10735	"multiplier (BFD Liveness Detection)" on page 11123	"no-adaptation (BFD Liveness Detection)" on page 11374
"version (BFD Liveness Detection)" on page 30697		

[edit tenants {name} routing-instances {name} protocols bgp]		
"advertise-prefix-sid" on page 2022	"bfd-liveness-detection (BGP)" on page 2918	"defaults (ACX series, JRR200, MX204, PTX5000, and VRR)" on page 5225
"egress-te-set-segment" on page 6236	"forwarding-context (Protocols BGP)" on page 7453	"multipath (Protocols BGP)" on page 11110
"tcpao-auth-mismatch" on page 29034		

[edit tenants {name} routing-instances {name} protocols bgp bfd-liveness-detection]		
"authentication (BGP BFD Liveness Detection)" on page 2588		

[edit tenants {name} routing-instances {name} protocols bgp family inet labeled-unicast]		
"entropy-label (protocols bgp)" on page 6413		

[edit tenants {name} routing-instances {name} protocols bgp family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit tenants {name} routing-instances {name} protocols bgp family inet6]		
"transport (Protocols BGP)" on page 30113		

[edit tenants {name} routing-instances {name} protocols bgp family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit tenants {name} routing-instances {name} protocols bgp group {group-name} {family} {family-modifier}]		
"add-path" on page 1839	"bfd-liveness-detection (BGP)" on page 2918	"multipath (Protocols BGP)" on page 11110
"peer-as-list" on page 12345		

[edit tenants {name} routing-instances {name} protocols bgp group {group-name} {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit tenants {name} routing-instances {name} protocols bgp group {name} bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit tenants {name} routing-instances {name} protocols bgp group {name} dynamic-neighbor]	
"peer-limit" on page 12375	

[edit tenants {name} routing-instances {name} protocols bgp group {name} family inet labeled-unicast]	
"entropy-label (protocols bgp)" on page 6413	

[edit tenants {name} routing-instances {name} protocols bgp group {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit tenants {name} routing-instances {name} protocols bgp group {name} family inet6]	
"transport (Protocols BGP)" on page 30113	

[edit tenants {name} routing-instances {name} protocols bgp group {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit tenants {name} routing-instances {name} protocols bgp group {name} metric-out]		
"effective-aigp" on page 6203	"minimum-effective-aigp" on page 10730	

[edit tenants {name} routing-instances {name} protocols bgp group {group-name} neighbor {address} {family} {family-modifier}]		
"add-path" on page 1839	"bfd-liveness-detection (BGP)" on page 2918	"egress-te-adj-segment" on page 6229
"egress-te-node-segment" on page 6234	"multipath (Protocols BGP)" on page 11110	

[edit tenants {name} routing-instances {name} protocols bgp group {group-name} neighbor {address} {family} {family-modifier} add-path]	
"send (add-path)" on page 16535	

[edit tenants {name} routing-instances {name} protocols bgp group {name} neighbor {address} bfd-liveness-detection]	
"authentication (BGP BFD Liveness Detection)" on page 2588	

[edit tenants {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet labeled-unicast]	
"entropy-label (protocols bgp)" on page 6413	

[edit tenants {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit tenants {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet6]		
"transport (Protocols BGP)" on page 30113		

[edit tenants {name} routing-instances {name} protocols bgp group {name} neighbor {name} family inet6 unicast]		
"accept-srv6-service" on page 1649	"advertise-srv6-service" on page 2024	

[edit tenants {name} routing-instances {name} protocols bgp multipath]		
"pause-computation-during-churn" on page 12277		

[edit tenants {name} routing-instances {name} protocols bgp source-packet-routing]		
"srgb (Protocols BGP)" on page 28407	"srv6 (Protocols BGP)" on page 28417	

[edit tenants {name} routing-instances {name} protocols isis dynamic-overload]		
"auto-recovery" on page 2733	"no-overload-on-prefix-export-limit" on page 11542	

[edit tenants {name} routing-instances {name} protocols isis interface {name} delay-measurement]		
"advertisement (Delay Measurement)" on page 2031		

[edit tenants {name} routing-instances {name} protocols isis interface-group {name} level {name} topology]		
"bandwidth-based-metrics" on page 2837		

[edit tenants {name} routing-instances {name} protocols isis no-external-export protocol]		
"no-external-export" on page 11454		

[edit tenants {name} routing-instances {name} protocols isis source-packet-routing]		
"strict-asla-based-flex-algorithm" on page 28645		

[edit tenants {name} routing-instances {name} protocols isis spf-options]		
"microloop-avoidance" on page 10700		

[edit tenants {name} routing-instances {name} protocols isis traffic-engineering]		
"advertisement" on page 2030		

[edit tenants {name} routing-instances {name} protocols mpls]		
"cspf-backoff-time" on page 4731		

[edit tenants {name} routing-instances {name} protocols mpls label-switched-path {name}]		
"conditional-metric" on page 4506	"cspf-backoff-time" on page 4731	

[edit tenants {name} routing-instances {name} protocols mpls label-switched-path {name} secondary {name} non-active-hold-priority]		
"non-active-hold-priority" on page 11657		

[edit tenants {name} routing-instances {name} protocols mpls static-label-switched-path {name}]		
"transit" on page 30064		

[edit tenants {name} routing-instances {name} protocols ospf spf-options]		
"microloop-avoidance (Protocols OSPF)" on page 10702		

[edit tenants {name} routing-instances {name} protocols pim]		
"disable-packet-register" on page 5813	"secondary-vtep-address" on page 16419	

[edit tenants {name} routing-instances {name} protocols pim interface]		
"dual-dr" on page 6045		

[edit tenants {name} routing-instances {name} provider-tunnel ingress-replication]		
"transport-class" on page 30132	"unicast-tunnel-name-regular-expression" on page 30355	

[edit tenants {name} routing-instances {name} provider-tunnel ldp-p2mp]		
"root-address" on page 15840		

[edit tenants {name} routing-instances {name} routing-options]		
"flex-algorithm" on page 7221	"flex-algorithm (Protocols OSPF)" on page 7223	

[edit tenants {name} routing-instances {name} routing-options bmp route-monitoring]		
"loc-rib" on page 9897		

[edit tenants {name} routing-instances {name} routing-options dynamic-tunnels]		
"forwarding-rib" on page 7483	"v4ov6" on page 30636	

[edit tenants {name} routing-instances {name} routing-options dynamic-tunnels tunnel-attributes]		
"dynamic-tunnel-reassembly" on page 6159		

[edit tenants {name} routing-instances {name} routing-options flex-algorithm]		
"definition" on page 5231	"definition (Protocols OSPF)" on page 5234	"use-transport-class" on page 30541

[edit tenants {name} routing-instances {name} routing-options multicast]		
"cont-stats-collection-interval" on page 4573		

[edit tenants {name} routing-instances {name} routing-options rib]		
"policy-multipath" on page 12694		

[edit tenants {name} routing-instances {name} routing-options source-packet-routing]		
"srv6 (Routing Options)" on page 28415		

[edit tenants {name} routing-instances {name} routing-options source-packet-routing srv6]		
"block" on page 2980	"locator" on page 9997	

[edit tenants {name} routing-instances {name} routing-options source-packet-routing srv6 locator]	
"micro-sid" on page 10697	

[edit tenants {name} routing-instances {name} routing-options transport-class]	
"name (Transport-class)" on page 11181	

[edit tenants {tenant-name} routing-instances {routing-instance-name} routing-options].]	
"routing-options-max-ecmp" on page 15985	

[edit tenants {name} routing-instances {name} system services dhcp-local-server dhcpv6]	
"dynamic-server" on page 6154	

[edit tenants {name} routing-instances {name} system services dhcp-local-server group]	
"interface-tag (DHCP Local Server)" on page 9081	

[edit tenants {name} routing-instances {name} vlans {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit tenants {name} routing-instances {name}protocolsisissource-packet-routing]	
"flex-algorithm" on page 7221	

[security]

[edit tenants {tenant-name} security]		
"address-book" on page 1868	"alg" on page 2130	"dynamic-address" on page 6086
"log (Logical Systems and Tenant Systems)" on page 10016	"log (Security)" on page 10026	"nat" on page 11205
"screen " on page 16384		

[edit tenants {name} security alg]		
"alg-support-lib" on page 2142	"dns" on page 5847	"ftp (Security ALG)" on page 7626
"h323" on page 7961	"ike-esp-nat" on page 8421	"mgcp" on page 10693
"msrpc" on page 10999	"pptp" on page 12986	"rsh" on page 16041
"rtsp" on page 16072	"sccp" on page 16341	"sip (Security) " on page 27977
"sql" on page 28402	"sunrpc" on page 28722	"talk" on page 28940
"tftp (Security ALG)" on page 29252	"traceoptions (Security ALG)" on page 29770	

[edit tenants {name} security alg dns]		
"maximum-message-length" on page 10534	"oversize-message-drop" on page 12081	

[edit tenants {name} security alg h323]		
"application-screen (Security H323)" on page 2359	"endpoint-registration-timeout" on page 6342	"media-source-port-any" on page 10599
"traceoptions (Security H323 ALG)" on page 29801		

[edit tenants {name} security alg h323 application-screen]		
"message-flood (Security H323)" on page 10644	"unknown-message (Security H323 ALG)" on page 30410	

[edit tenants {name} security alg h323 application-screen message-flood]	
"gatekeeper" on page 7653	

[edit tenants {name} security alg mgcp]		
"application-screen (Security MGCP)" on page 2361	"inactive-media-timeout (Security MGCP)" on page 8496	"traceoptions (Security MGCP ALG)" on page 29814
"transaction-timeout" on page 30055		

[edit tenants {name} security alg mgcp application-screen]		
"connection-flood" on page 4531	"message-flood (Security MGCP)" on page 10645	"unknown-message (Security MGCP ALG)" on page 30411

[edit tenants {tenant-name} security alg msrpc]		
"map-entry-timeout" on page 10333		

[edit tenants {name} security alg sccp]		
"application-screen (Security SCCP)" on page 2362	"inactive-media-timeout (Security SCCP)" on page 8497	"traceoptions (Security SCCP ALG)" on page 29823

[edit tenants {name} security alg sccp application-screen]		
"call-flood" on page 3179	"unknown-message (Security SCCP ALG)" on page 30413	

[edit tenants {name} security alg sip]		
"application-screen (Security SIP)" on page 2364	"c-timeout" on page 3152	"t1-interval" on page 28885
"t4-interval" on page 28893	"traceoptions (Security SIP ALG)" on page 29830	

[edit tenants {name} security alg sip (Security)]		
"inactive-media-timeout (Security SIP)" on page 8499	"retain-hold-resource" on page 15674	

[edit tenants {name} security alg sip (Security) application-screen (Security SIP)]	
"unknown-message (Security SIP ALG)" on page 30415	

[edit tenants {name} security alg sip application-screen protect]	
"deny (Security SIP)" on page 5317	

[edit tenants {tenant-name} security alg sunrpc]		
"map-entry-timeout" on page 10333		

[edit tenants {tenant-name} security application-firewall]		
"appfw-profile (System)" on page 2294	"appfw-rule" on page 2296	"appfw-rule-set" on page 2298
"traceoptions (Security Application Firewall)" on page 29776		

[edit tenants {tenant-name} security dynamic-address]	
"address-name" on page 1892	

[edit tenants {tenant-name} security flow]	
"traceoptions (Security Flow)" on page 29789	

[edit tenants {tenant-name} security idp]		
"custom-attack" on page 4738	"custom-attack-group" on page 4746	"dynamic-attack-group" on page 6097
"idp-policy (Security) " on page 8317		

[edit tenants {name} security idp idp-policy {name} rulebase-exempt rule {name} then notification]		
"log-attacks" on page 10050	"log-attacks (rulebase-exempt)" on page 10052	

[edit tenants {tenant name} security idp sensor-configuration]	
"packet-log (Security IDP Sensor Configuration)" on page 12153	

[edit tenants {name} security idp sensor-configuration global intelligent-offload-tunable]	
"intelligent-offload-tunable (IDP)" on page 8780	

[edit tenants {name} security ipsec]	
"anti-replay-window-size" on page 2255	

[edit tenants {tenant-name} security log]		
"event-rate" on page 6531	"facility-override (Security)" on page 6771	"file (Security Log)" on page 7076
"mode (Security Log)" on page 10823	"rate-cap" on page 14058	"report (Security Log)" on page 14398
"retry-count (Security Log)" on page 15696	"root-streaming" on page 15838	"stream (Logical Systems and Tenant Systems)" on page 28628
"stream (Security Log)" on page 28635	"traceoptions (Security Log)" on page 29811	"transport (Security Log)" on page 30129

[edit tenants {tenant-name} security log cache]		
"exclude (Security Log)" on page 6627	"limit (Security Log)" on page 9760	

[edit tenants {tenant-name} security log stream {stream-name}]		
"category (Security Logging)" on page 3224	"format (Security Log Stream)" on page 7388	"host (Security Log)" on page 8166
"radius (Security Log)" on page 14018	"rate-limit (Security Log)" on page 14069	"severity (Security Log)" on page 16906

[edit tenants {name} security nat]	
"address-persistent (Security Source NAT)" on page 1896	

[edit tenants {name} security nat proxy-arp interface]	
"address (Security ARP Proxy)" on page 1922	

[edit tenants {name} security nat proxy-ndp interface]	
"address (Security NDP Proxy)" on page 1925	

[edit tenants {name} security nat source pool]	
"address-shared (Security Source NAT)" on page 1932	

[edit tenants {name} security nat source pool {name} address-persistent subscriber]	
"address-persistent (Security Source NAT Pool)" on page 1897	

[edit tenants {name} security nat source pool {name} address-pooling]		
"address-pooling (Security Source NAT)" on page 1905		

[edit tenants {name} security nat source pool {name} pool-utilization-alarm]		
"clear-threshold" on page 4245		

[edit tenants {name} security nat source pool-utilization-alarm]		
"clear-threshold" on page 4245		

[edit tenants {name} security nat source rule-set {rule-set-name} rule {rule-name} match]		
"application (Security Source NAT)" on page 2373		

[edit tenants {name} security nat source {rule-set} {name} {rule} {name} then source nat]		
"address-mapping" on page 1887		

[edit tenants {tenant-system-name} security policies]		
"dns-cache" on page 5852	"unified-policy max-lookups" on page 30362	

[edit tenants {name} security policies explicit-proxy]		
"profile (Explicit Proxy)" on page 13277		

[edit tenants {name} security policies from-zone]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	

[edit tenants {tenant-name} security policies from-zone {zone-name} policy {policy-name}]	
"report-skip" on page 14403	

[edit tenants {tenants-systems-name} security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit]	
"tcp-options (Security Policies)" on page 29008	

[edit tenants {tenant-name} security policies from-zone {zone-name} to-zone {zone-name} policy {policy-name} then permit application-services]	
"application-traffic-control (Application Services)" on page 2402	

[edit tenants {tenant-name} security policies global policy {policy-name}]	
"report-skip" on page 14403	

[edit tenants {name} security policies global policy {name} then]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	

[edit tenants {tenants-systems-name} security policies global policy {policy-name} then permit]	
"tcp-options (Security Policies)" on page 29008	

[edit tenants {tenants-systems-name} security policies policy-set {policy-set-name} policy {policy-name} then permit]	
"tcp-options (Security Policies)" on page 29008	

[edit tenants {name} security policies stateful-firewall-rule {name} policy {name} then]		
"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757	

[edit tenants {tenant-name} security screen]		
"ids-option" on page 8335	"trap" on page 30133	"white-list (Security Screen)" on page 31107

[edit tenants {tenant-name} security screen ids-option {screen-name} udp flood]	
"white-list (Security Screen)" on page 31107	

[edit tenants {name} security tunnel-inspection]	
"inspection-profile" on page 8739	

[edit tenants {tenant-name} security utm]	
"custom-objects" on page 4771	

[edit tenants {name} security utm custom-objects]	
"url-feed" on page 30489	

[edit tenants {name} security utm custom-objects custom-message]		
"custom-page-file" on page 4778		

[edit tenants {name} security utm default-configuration anti-virus]		
"notification-options (Security Antivirus)" on page 11693	"sophos-engine" on page 28109	

[edit tenants {name} security utm default-configuration anti-virus avira-engine pattern-update]		
"proxy-profile" on page 13838		

[edit tenants {name} security utm default-configuration anti-virus sophos-engine]		
"notification-options (Security Antivirus)" on page 11693		

[edit tenants {name} security utm default-configuration anti-virus sophos-engine pattern-update]		
"proxy-profile" on page 13838		

[edit tenants {name} security utm default-configuration url-feed]		
"global-limit" on page 7708		

[edit tenants {tenant-name} security utm feature-profile]		
"anti-spam (Logical Systems)" on page 2259	"anti-virus (Logical Systems)" on page 2266	"content-filtering (Logical System Security Feature Profile)" on page 4580
"web-filtering (Logical System Security Feature Profile)" on page 31072		

[edit tenants {name} security utm feature-profile anti-virus profile]	
"notification-options (Security Antivirus)" on page 11693	

[edit tenants {name} security utm feature-profile anti-virus sophos-engine]	
"server (Security Sophos Engine Antivirus)" on page 16670	

[edit tenants {name} security utm feature-profile anti-virus sophos-engine pattern-update]	
"proxy-profile" on page 13838	

[edit tenants {name} security utm feature-profile anti-virus sophos-engine profile]	
"notification-options (Security Antivirus)" on page 11693	

[edit tenants {name} security utm feature-profile web-filtering juniper-enhanced profile {profile-name}]	
"custom-message (Security Web Filtering)" on page 4764	

[edit tenants {name} security utm feature-profile web-filtering juniper-local profile {profile-name}]	
"custom-message (Security Web Filtering)" on page 4764	

[edit tenants {name} security utm feature-profile web-filtering websense-redirect profile {profile-name}]	
"custom-message (Security Web Filtering)" on page 4764	

[edit tenants {tenant-name} security utm utm-policy {policy-name}]		
"anti-spam (Security Content Security Policy)" on page 2265	"anti-virus (Content Security Policy)" on page 2273	"content-filtering (Security Content Security Policy)" on page 4583
"web-filtering (Security Content Security Policy)" on page 31076		

[edit tenants {name} security zones security-zone]	
"white-list (Security-Zone)" on page 31109	

[security-profile]

[edit tenants {tenant-name} security-profile]	
"nat-destination-pool" on page 11212	

[securitypolicies]

[edit tenants {tenants-systems-name} securitypolicies policies policyexplicit-setproxy profile {policy-set-name} policy {policy-name} then permit]	
"tcp-options (Security Policies)" on page 29008	

[services]

[edit tenants {tenants_name} services]	
"icap-redirect" on page 8245	

[system]

[edit tenants {tenant-name} system security-profile]		
"nat-destination-rule" on page 11214	"nat-source-pool" on page 11234	"nat-source-rule" on page 11236
"nat-static-rule" on page 11238		

[edit unified-edge]

IN THIS SECTION

- [aaa] | **1232**
- [diameter-profiles] | **1233**
- [gateways] | **1234**
- [pcef] | **1241**
- [tdf] | **1244**

Sub Hierarchies

"[edit unified-edge aaa]" on page 1232

"[edit unified-edge diameter-profiles]" on page 1233

"[edit unified-edge gateways]" on page 1234

"[edit unified-edge pcef]" on page 1241

"[edit unified-edge tdf {gateway-name}]" on page 1244

[edit unified-edge]	
"pcef" on page 12311	

[aaa]

[edit unified-edge aaa]	
"profiles (AAA)" on page 13363	

[edit unified-edge aaa profiles {aaa-profile-name} radius]		
"accounting (AAA Profile)" on page 1710	"authentication (AAA Profile)" on page 2576	

[edit unified-edge aaa profiles {aaa-profile-name} radius accounting]	
"network-element (AAA Profile)" on page 11304	

[edit unified-edge aaa profiles {aaa-profile-name} radius authentication]		
"network-element (AAA Profile)" on page 11304		

[edit unified-edge aaa profiles {aaa-profile-name} radius policy]		
"activation-attribute (AAA Profile)" on page 1796	"coa-accounting (AAA Profile)" on page 4381	"deactivation-attribute (AAA Profile)" on page 5143

[edit unified-edge aaa profiles {aaa-profile-name} radius policy activation-attribute]		
"code (AAA Profile)" on page 4388	"vendor-id (AAA Profile)" on page 30681	

[edit unified-edge aaa profiles {aaa-profile-name} radius policy deactivation-attribute]		
"code (AAA Profile)" on page 4388	"vendor-id (AAA Profile)" on page 30681	

[diameter-profiles]

[edit unified-edge diameter-profiles]		
"gx-profile" on page 7954		

[edit unified-edge diameter-profiles gx-profile {profile-name}]		
"attributes (Diameter Gx Profiles)" on page 2555	"request-timeout" on page 15432	"targets (Unified Edge)" on page 28971

[edit unified-edge diameter-profiles gx-profile {profile-name} attributes]		
"exclude (Diameter Gx Profiles)" on page 6602	"include (Diameter Gx Profiles)" on page 8524	

[gateways]

[edit unified-edge gateways]		
"tdf (Unified Edge)" on page 29037		

[edit unified-edge gateways tdf {gateway-name}]		
"aaa clients (TDF)" on page 1608	"cac (TDF Gateway)" on page 3165	"call-rate-statistics" on page 3180
"domain-selection" on page 5901	"domains" on page 5908	"service-mode" on page 16711
"traceoptions (TDF Gateway)" on page 29896		

[edit unified-edge gateways tdf {gateway-name} aaa]		
"snoop-segments (TDF Gateway)" on page 28080		

[edit unified-edge gateways tdf {gateway-name} cac]		
"cpu (TDF Gateway)" on page 4699	"maximum-sessions (TDF Gateway)" on page 10565	"maximum-sessions-trap-percentage (TDF Gateway)" on page 10566
"memory (TDF Gateway)" on page 10633		

[edit unified-edge gateways tdf {gateway-name} domain-selection]		
"term (TDF Domain Selection)" on page 29139		

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name}]		
"from (TDF Domain Selection)" on page 7603	"then (TDF Domain Selection)" on page 29303	

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from]		
"3gpp-imsi" on page 1578	"attribute" on page 2551	"called-station-id" on page 3185
"calling-station-id" on page 3192	"class (Unified Edge)" on page 3345	"client (Unified Edge)" on page 4300
"framed-ip-address" on page 7564	"framed-ipv6-prefix" on page 7566	"nas-ip-address" on page 11182
"snoop-segment (TDF Domain Selection)" on page 28076	"user-name" on page 30568	

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from 3gpp-imsi]		
"equals" on page 6426	"has-prefix (Unified Edge Gateways)" on page 7977	"has-suffix" on page 7979
"matches" on page 10409		

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from attribute {name}]		
"code" on page 4385	"format (Unified Edge Gateways)" on page 7389	"vendor-id (Unified Edge)" on page 30679

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from attribute {name} format]		
"integer" on page 8777	"string" on page 28655	"time" on page 29352
"v4address" on page 30635	"v6address" on page 30638	"v6prefix" on page 30639

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from attribute {name} format integer]		
"equals" on page 6426	"greater-than" on page 7798	"less-than" on page 9713

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from attribute {name} format string]		
"equals" on page 6426	"has-prefix (Unified Edge Gateways)" on page 7977	"has-suffix" on page 7979
"matches" on page 10409		

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from attribute {name} format time]		
"equals" on page 6426	"greater-than" on page 7798	"less-than" on page 9713

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from attribute {name} format v4address]	
"equals" on page 6426	

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from attribute {name} format v6address]	
"equals" on page 6426	

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from attribute {name} format v6prefix]		
"equals" on page 6426		

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from called-station-id]		
"equals" on page 6426	"matches" on page 10409	

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from calling-station-id]		
"equals" on page 6426	"matches" on page 10409	

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from class]		
"equals" on page 6426	"has-prefix (Unified Edge Gateways)" on page 7977	"has-suffix" on page 7979
"matches" on page 10409		

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from framed-ip-address]		
"equals" on page 6426		

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from framed-ipv6-prefix]		
"equals" on page 6426		

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} from nas-ip-address]		
"equals" on page 6426		

[edit unified-edge gateways tdf {gateway-name} domain-selection term {term-name} then]		
"domain (TDF Domain Selection)" on page 5905	"pcef-profile (TDF Domain Selection)" on page 12318	

[edit unified-edge gateways tdf {gateway-name} domains {domain-name}]		
"burst-size (TDF Domain)" on page 3090	"call-rate-statistics" on page 3180	"ifl-subscriber" on page 8379
"ip-subscriber" on page 9275	"maximum-bit-rate (TDF Domain)" on page 10509	"pcef-profile (TDF Domain)" on page 12317
"service-mode" on page 16711	"subscriber-exclude-prefix" on page 28691	"subscriber-type (TDF Domain)" on page 28702
"tdf-interface " on page 29036		

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} ifl-subscriber {subscriber-name}]	
"access-interfaces (IFL Subscriber)" on page 1662	

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} ip-subscriber]		
"access-interfaces (IP Subscriber)" on page 1664	"default-local-policy" on page 5195	"idle-timeout (Unified Edge)" on page 8304
"immediate-accounting-response" on page 8452	"maximum-subscribers" on page 10568	"subscriber-address" on page 28683
"subscription-id" on page 28711		

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} ip-subscriber default-local-policy]		
"burst-size (Default Local Policy)" on page 3081	"flow-action" on page 7259	"maximum-bit-rate (Default Local Policy)" on page 10505

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} ip-subscriber subscriber-address]		
"inet (TDF Subscriber Address)" on page 8593	"inet6 (TDF Subscriber Address)" on page 8617	

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} ip-subscriber subscriber-address (inet | inet6)]	
"pool (TDF Domain)" on page 12791	

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} ip-subscriber subscription-id]		
"constant" on page 4570	"subscription-id-options" on page 28715	

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} ip-subscriber subscription-id subscription-id-options {entry-name}]	
"id-components" on page 8281	

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} subscriber-exclude-prefix]	
"family (Exclude Prefix)" on page 6862	

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} subscriber-exclude-prefix family]		
"inet (TDF Subscriber Exclude Prefix)" on page 8594	"inet6 (TDF Subscriber Exclude Prefix)" on page 8619	

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} subscriber-exclude-prefix family inet]		
"network (TDF Domain)" on page 11328		

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} subscriber-exclude-prefix family inet6]		
"network (TDF Domain)" on page 11328		

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} subscription-id]		
"use-class (Class Attribute)" on page 30514		

[edit unified-edge gateways tdf {gateway-name} domains {domain-name} subscription-id use-class]		
"pattern (Class Attribute)" on page 12272	"regex (Class Attribute)" on page 14260	"subscription-id-type (Class Attribute)" on page 28717

[edit unified-edge gateways tdf {gateway-name} system]		
"service-pics" on page 16723	"session-pics" on page 16840	

[edit unified-edge gateways tdf {gateway-name} system service-pics]		
"interface (Services PIC)" on page 9031		

[edit unified-edge gateways tdf {gateway-name} system session-pics]		
"interface (Session PICs)" on page 9033		

[edit unified-edge gateways tdf {gateway-name} system session-pics interface {interface-name}]	
"service-mode" on page 16711	

[pcef]

[edit unified-edge pcef]		
"flow-descriptions" on page 7274	"pcc-action-profiles" on page 12286	"pcc-rulebases (PCEF)" on page 12293
"pcc-rules (PCEF)" on page 12297	"pcc-time-of-day-profiles" on page 12301	"profiles (PCEF)" on page 13365
"traceoptions (PCEF)" on page 29635		

[edit unified-edge pcef flow-description {flow-identifier}]		
"local-ports" on page 9969	"no-send-to-ue" on page 11579	"protocol (Flow Descriptions)" on page 13441
"remote-ports" on page 14357		

[edit unified-edge pcef flow-descriptions {flow-identifier}]		
"direction (Service Data Flow Filters)" on page 5727	"local-port-range" on page 9967	"remote-address" on page 14311
"remote-port-range" on page 14355		

[edit unified-edge pcef pcc-action-profiles {profile-name}]		
"forwarding-class (PCC Action Profiles)" on page 7434	"gate-status" on page 7649	"hcm-profile (PCC Action Profiles)" on page 8008
"logging-rule (PCC Action Profile)" on page 10054	"maximum-bit-rate (PCC Action Profiles)" on page 10507	"monitoring-key (PCC Action Profile)" on page 10957
"redirect (PCC Action Profiles)" on page 14159	"steering" on page 28605	

[edit unified-edge pcef pcc-action-profiles {profile-name} redirect]	
"url" on page 30485	

[edit unified-edge pcef pcc-action-profiles {profile-name} steering]		
"keep-existing-steering" on page 9459	"path (Steering)" on page 12269	"routing-instance (PCC Action Profiles)" on page 15951

[edit unified-edge pcef pcc-action-profiles {profile-name} steering path]		
"ipv4-address (Steering Path)" on page 9327	"ipv6-address (Steering Path)" on page 9347	

[edit unified-edge pcef pcc-rule-bases {rulebase-name}]	
"pcc-rule" on page 12291	

[edit unified-edge pcef pcc-rules {rule-name}]		
"from (PCC Rules)" on page 7591	"then (PCC Rules)" on page 29276	

[edit unified-edge pcef pcc-rules {rule-name} from]		
"application-groups (PCC Rules)" on page 2326	"applications (PCC Rules)" on page 2410	"flows (PCC Rules)" on page 7362

[edit unified-edge pcef pcc-rules {rules-name} then]	
"pcc-action-profile (PCC Rules)" on page 12284	

[edit unified-edge pcef pcc-time-of-day-profiles {profile-name}]		
"rule-activation-time" on page 16078	"rule-deactivation-time" on page 16089	

[edit unified-edge pcef profiles {profile-name}]		
"aaa-policy-control (PCEF Profile)" on page 1616	"dynamic-policy-control" on page 6112	"static-policy-control" on page 28555

[edit unified-edge pcef profiles {profile-name} aaa-policy-control]		
"aaa-profile (PCEF Profile)" on page 1618	"pcc-rulebases (PCEF Profile)" on page 12294	"pcc-rules (PCEF Profile)" on page 12299
"user-password (PCEF Profile)" on page 30574		

[edit unified-edge pcef profiles {profile-name} dynamic-policy-control]		
"diameter-profile (PCEF Profile)" on page 5698	"pcc-rulebases (PCEF Profile)" on page 12294	"pcc-rules (PCEF Profile)" on page 12299

[edit unified-edge pcef profiles {profile-name} static-policy-control]		
"pcc-rulebases (PCEF Profile)" on page 12294	"pcc-rules (PCEF Profile)" on page 12299	

[tdf]

[edit unified-edge tdf {gateway-name}]	
"diameter (TDF Gateway)" on page 5702	

[edit unified-edge tdf {gateway-name} diameter]	
"network-element (Subscriber Aware Policy Control)" on page 11308	

[edit unified-edge tdf {gateway-name} diameter network-element {element-name}]	
"session-pics (Diameter)" on page 16842	

[edit unified-edge tdf {gateway-name} system interface {interface-name}]	
"service-mode" on page 16711	

[edit virtual-chassis]

IN THIS SECTION

- [member] | 1246

Sub Hierarchies

"[edit virtual-chassis member {member-id}]" on page 1246

[edit virtual-chassis]		
"aliases (Virtual Chassis)" on page 2153	"auto-sw-update" on page 2745	"fast-failover" on page 6924
"heartbeat-address (MX Series Virtual Chassis)" on page 8020	"heartbeat-timeout (MX Series Virtual Chassis)" on page 8025	"heartbeat-tos (MX Series Virtual Chassis)" on page 8027
"id" on page 8279	"locality-bias (MX Series Virtual Chassis)" on page 9981	"mac-persistence-timer" on page 10252
"member" on page 10604	"member (MX Series Virtual Chassis)" on page 10615	"no-split-detection" on page 11585
"no-split-detection (MX Series Virtual Chassis)" on page 11586	"preprovisioned" on page 13115	"preprovisioned (MX Series Virtual Chassis)" on page 13116
"traceoptions (MX Series Virtual Chassis)" on page 29609	"traceoptions (Virtual Chassis)" on page 29905	"vc-port" on page 30658
"vcp-no-hold-time" on page 30662	"vcp-snmp-statistics" on page 30664	

[member]

[edit virtual-chassis member {member-id}]		
"fabric-tree-root" on page 6766	"location (Virtual Chassis)" on page 9995	"mastership-priority" on page 10367
"no-management-vlan" on page 11509	"role" on page 15824	"role (MX Series Virtual Chassis)" on page 15827
"serial-number" on page 16599	"serial-number (MX Series Virtual Chassis)" on page 16603	

[edit virtual-network-functions]

[edit virtual-network-functions]		
"features (VNF)" on page 6978	"hugepages" on page 8215	"image" on page 8447
"init-descriptor" on page 8657	"interfaces (VNF)" on page 9111	"mac-address" on page 10218
"mapping (VNF)" on page 10340	"memory" on page 10628	"no-autostart" on page 11412
"offloads" on page 11789	"pci-address" on page 12322	"size (VNF)" on page 27999
"storage" on page 28613	"type" on page 30273	"virtual-cpu" on page 30752
"vjunos0" on page 30780	"vnf-name" on page 30906	

[edit vlans]

IN THIS SECTION

- [dot1q-tunneling] | **1248**
- [forwarding-options] | **1249**
- [interface] | **1253**
- [switch-options] | **1253**
- [vlan-id] | **1254**
- [vxlan] | **1255**
- [{vlan-name] | **1255**

Sub Hierarchies

"[edit vlans {vlan-name} dot1q-tunneling]" on page 1248

"[edit vlans {vlan-name} forwarding-options]" on page 1249

"[edit vlans {vlan-name} interface {interface-name}]" on page 1253

"[edit vlans {name} switch-options]" on page 1253

"[edit vlans {primary-vlan-name} vlan-id {primary-vlan-vlan-id}]" on page 1254

"[edit vlans {vlan-name} vxlan]" on page 1255

"[edit vlans {vlan-name]" on page 1255

[edit vlans {vlan-name }switch-options interface {interface-name}]		
"action-priority" on page 1768	"decapsulate-inner-vlan" on page 5163	"description (VLAN)" on page 5360

(Continued)

[edit vlans {vlan-name }switch-options interface {interface-name}]		
"filter (VLANs)" on page 7157	"flood (VLANs)" on page 7248	"flood (VLANs)" on page 7248
"forwarding-options (Bridge Domains or VLANs)" on page 7467	"interface (Logical Interfaces)" on page 8814	"l3-interface (VLAN)" on page 9568
"l3-interface (VLAN)" on page 9568	"l3-interface-ingress-counting" on page 9567	"mac-limit (VLAN)" on page 10238
"mac-table-aging-time" on page 10267	"no-local-switching (VLAN)" on page 11500	"private-vlan" on page 13220
"riot-loopback" on page 15801	"service-id (Switch Options)" on page 16705	"vlan members (VLANs)" on page 30827
"vlan-id-list" on page 30805	"vlan-prune" on page 30842	"vlan-range" on page 30843
"vlan-tags" on page 30865	"vxlan" on page 30991	

[dot1q-tunneling]

[edit vlans {vlan-name} dot1q-tunneling]		
"customer-vlans" on page 4783	"layer2-protocol-tunneling" on page 9658	

[edit vlans {vlan-name} dot1q-tunneling layer2-protocol-tunneling (all | protocol-name {vlan-name} dot1q-tunneling layer2-protocol-tunneling (all | {protocol-name})]	
"drop-threshold" on page 5978	

[edit vlans {vlan-name} dot1q-tunneling layer2-protocol-tunneling (all | {protocol-name})]	
"shutdown-threshold" on page 27937	

[forwarding-options]

[edit vlans {vlan-name} forwarding-options]		
"dhcp-security" on page 5640	"filter (VLANs)" on page 7157	"fip-security" on page 7175
"light-weight-dhcpv6-relay" on page 9755	"mine-dhcp-client-options" on page 10719	"mine-dhcpv6-client-options" on page 10721

[edit vlans {name} forwarding-options dhcp-relay dhcpv6 group {name} dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit vlans {name} forwarding-options dhcp-relay dhcpv6 group {name} relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit vlans {name} forwarding-options dhcp-relay dhcpv6 relay-option]		
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	

[edit vlans {name} forwarding-options dhcp-relay dual-stack-group {dual-stack-group-name}]		
"classification-key (DHCP Relay Agent)" on page 3376	"protocol-master" on page 13448	

[edit vlans {name} forwarding-options dhcp-relay group]	
"interface-tag (DHCP Relay Server)" on page 9083	

[edit vlans {name} forwarding-options dhcp-relay group {name} relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit vlans {name} forwarding-options dhcp-relay group {name} relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit vlans {name} forwarding-options dhcp-relay group {name} relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit vlans {name} forwarding-options dhcp-relay relay-option]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-order (DHCP Relay Agent)" on page 11899	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit vlans {name} forwarding-options dhcp-relay relay-option option-60]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit vlans {name} forwarding-options dhcp-relay relay-option option-77]		
"equals (DHCP Relay Agent)" on page 6428	"not-present (DHCP Relay Agent)" on page 11682	"starts-with (DHCP Relay Agent Option)" on page 28506

[edit vlans {vlan-name} forwarding-options dhcp-security]		
"dhcpv6-options" on page 5673	"group (DHCP Security)" on page 7826	"neighbor-discovery-inspection" on page 11270
"no-dhcp-snooping" on page 11436	"no-dhcpv6-snooping" on page 11441	"option-82" on page 11884

[edit vlans {vlan-name} forwarding-options dhcp-security dhcpv6-options]		
"option-16 (DHCPv6 Snooping)" on page 11868	"option-18 (DHCPv6 Snooping)" on page 11869	"option-37 (DHCPv6 Snooping)" on page 11872
"option-79 " on page 11883		

[edit vlans forwarding-options dhcp-security dhcpv6-options option-18]	
"prefix (DHCPv6 Options)" on page 13079	

[edit vlans forwarding-options dhcp-security dhcpv6-options option-37]	
"prefix (DHCPv6 Options)" on page 13079	

[edit vlans {vlan-name} forwarding-options dhcp-security group {group-name}]		
"overrides (DHCP Security)" on page 12071		

[edit vlans {vlan-name} forwarding-options dhcp-security group {group-name} interface {interface-name}]		
"static-ipv6" on page 28544		

[edit vlans {vlan-name} forwarding-options dhcp-security group {group-name} overrides]		
"no-dhcpv6-options" on page 11440	"no-option16" on page 11537	"no-option18" on page 11538
"no-option37" on page 11540	"no-option82" on page 11541	"trusted (DHCP Security)" on page 30174
"untrusted (VLAN)" on page 30433		

[edit vlans forwarding-options dhcp-security option-82 circuit-id prefix]		
"routing-instance-name" on page 15943	"routing-instance-name (circuit-id)" on page 15944	

[edit vlans {vlan-name} forwarding-options dhcp-security option-82 remote-id]		
"host-name" on page 8147	"mac (Option 82)" on page 10250	

[edit vlans {vlan-name} forwarding-options fip-security]		
"beacon-period" on page 2890	"examine-vn2vf" on page 6566	"examine-vn2vn" on page 6568
"fc-map" on page 6957	"interface (FIP Snooping)" on page 8889	

[edit vlans {vlan-name} forwarding-options fip-security interface {interface-name}]		
"fcoe-trusted" on page 6965	"no-fcoe-trusted" on page 11458	

[interface]

[edit vlans {vlan-name} interface {interface-name}]	
"mapping" on page 10338	

[edit vlans {vlan-name} interface {interface-name} egress]	
"mapping" on page 10338	

[edit vlans {vlan-name} interface {interface-name} ingress]	
"mapping" on page 10338	

[switch-options]

[edit vlans {name} switch-options]		
"drop-unknown-mac-ip" on page 5984	"interface (Switching Options)" on page 9078	"interface-mac-ip-limit" on page 8908
"interface-mac-limit" on page 8910	"mac-statistics" on page 10263	"mac-table-size" on page 10271
"network-isolation-profile" on page 11314		

[edit vlans {vlan-name} switch-options interface {interface-name}]		
"interface-mac-limit" on page 8910	"static-mac" on page 28549	"static-mac-ip" on page 28587

[edit vlans {vlan-name} switch-options interface {interface-name} interface-mac-limit {limit}]	
"packet-action" on page 12129	

[edit vlans {vlan-name} switch-options interface-mac-limit {limit}]	
"packet-action" on page 12129	

[edit vlans {vlan-name} switch-options mac-table-size {limit}]	
"packet-action" on page 12129	

[vlan-id]

[edit vlans {primary-vlan-name} vlan-id {primary-vlan-vlan-id}]		
"community-vlan" on page 4477	"isolated-vlan" on page 9407	"isolation-id" on page 9410

[edit vlans {vlan-name} vlan-id {number} interface {interface-name}]		
"ingress (vlans)" on page 8653	"pvlan-trunk" on page 13883	

[vxlan]

[edit vlans {vlan-name} vxlan]		
"encapsulate-inner-vlan" on page 6282	"ingress-node-replication" on page 8634	"ingress-node-replication (EVPN)" on page 8635
"multicast-group" on page 11069	"ovsdb-managed" on page 12088	"static-remote-vtep-list" on page 28566
"translation-vni" on page 30080	"unreachable-vtep-aging-timer" on page 30429	"vni" on page 30909

[{vlan-name]

[edit vlans {vlan-name }forwarding-options dhcp-security dhcpv6-options option-18]		
"include-irb-and-l2" on page 8529	"no-vlan-interface-name" on page 11617	"use-interface-description" on page 30522
"use-interface-index" on page 30525	"use-interface-name" on page 30526	

[edit vlans {vlan-name }forwarding-options dhcp-security dhcpv6-options option-37]		
"include-irb-and-l2" on page 8529	"no-vlan-interface-name" on page 11617	"use-interface-description" on page 30522
"use-interface-index" on page 30525	"use-interface-name" on page 30526	

[edit vmhost]

[edit vmhost]		
"vmhost management-if disable" on page 30899	"vmhost management-if link-mode" on page 30900	"vmhost management-if speed " on page 30902

[edit wlan]

IN THIS SECTION

- [access-point] | **1256**

Sub Hierarchies

"[edit wlan access-point {name}]" on page 1256

[access-point]

[edit wlan access-point {name}]	
"wlan" on page 31137	

3

CHAPTER

Configuration Statements

Configuration Statements | 1258

Configuration Statements

IN THIS SECTION

- 0-9 | 1259
- A | 1259
- B | 1277
- C | 1281
- D | 1293
- E | 1309
- F | 1317
- G | 1329
- H | 1334
- I | 1338
- J | 1356
- K | 1356
- L | 1357
- M | 1368
- N | 1380
- O | 1390
- P | 1396
- Q | 1416
- R | 1417
- S | 1432
- T | 1456
- U | 1473
- V | 1478
- W | 1483
- X | 1485
- Y | 1485
- Z | 1485

0-9

"3gpp-imsi" on page 1578	"802.3ad" on page 1582	"802.3ad (EX and QFX Series)" on page 1579

A

"aaa" on page 1605	"aaa clients (TDF)" on page 1608	"aaa-access-profile (L2TP LNS)" on page 1607
"aaa-context (AAA Options)" on page 1610	"aaa-logical-system (Domain Map)" on page 1611	"aaa-options (Access Profile)" on page 1613
"aaa-options (PPP Profile)" on page 1614	"aaa-policy-control (PCEF Profile)" on page 1616	"aaa-profile (PCEF Profile)" on page 1618
"aaa-routing-instance (Domain Map)" on page 1619	"aacl-fields" on page 1621	"aacl-statistics-profile" on page 1623
"abstract-hop" on page 1625	"accept (Dynamic Profile)" on page 1629	"accept (Router Advertisement Guard Policy)" on page 1627
"accept-data" on page 1631	"accept-join-always-from" on page 1632	"accept-max-tcp-connections (System Process)" on page 1634
"accept-out-of-band" on page 1636	"accept-own" on page 1637	"accept-prpd-connection" on page 1639
"accept-remote-nexthop" on page 1640	"accept-remote-source" on page 1642	"accept-sdr (PCRF Partition)" on page 1644
"accept-source-mac" on page 1647	"accept-source-mac (SRX)" on page 1645	"accept-srv6-service" on page 1649

"accepted-prefix-limit" on page 1651	"access (Dynamic Access Routes)" on page 1659	"access (SNMP)" on page 1702
"access (Static Framed Routes)" on page 1703	"access-concentrator" on page 1655	"access-domain (Remote Device Management)" on page 1657
"access-identifier" on page 1661	"access-interfaces (IFL Subscriber)" on page 1662	"access-interfaces (IP Subscriber)" on page 1664
"access-internal (Dynamic Access-Internal Routes)" on page 1666	"access-internal (Static Access-Internal Routes)" on page 1668	"access-line (Access-Line Rate Adjustment)" on page 1669
"access-line-information (L2TP)" on page 1684	"access-list" on page 1686	"access-point name" on page 1687
"access-profile" on page 1688	"access-profile (AAA Options)" on page 1690	"access-profile (Domain Map)" on page 1691
"access-profile (Dynamic Stacked VLAN)" on page 1693	"access-profile (Dynamic VLAN)" on page 1694	"access-profile (Extensible Subscriber Services Manager)" on page 1695
"access-profile (Static Subscribers)" on page 1698	"access-profile-name (Duplicate Accounting)" on page 1697	"access-security" on page 1700
"access-trunk" on page 1705	"account-layer2-overhead (PIC Level)" on page 1706	"accounting (AAA Profile)" on page 1710
"accounting (Access Profile)" on page 1711	"accounting (Dynamic IGMP Interface)" on page 1714	"accounting (Dynamic MLD Interface)" on page 1716
"accounting (Forwarding Options)" on page 1707	"accounting (interfaces)" on page 1709	"accounting (Protocols AMT Interface)" on page 1728
"accounting (Protocols IGMP)" on page 1730	"accounting (Protocols IGMP AMT Interface)" on page 1731	"accounting (Protocols IGMP Interface)" on page 1732

"accounting (Protocols MLD)" on page 1733	"accounting (Protocols MLD Interface)" on page 1735	"accounting (RADIUS Client)" on page 1736
"accounting (Service Accounting)" on page 1742	"accounting (System)" on page 1746	"accounting-backup-options (Access Profile)" on page 1713
"accounting-options" on page 1717	"accounting-order" on page 1721	"accounting-order (Service Accounting)" on page 1722
"accounting-port (RADIUS Server)" on page 1724	"accounting-profile" on page 1727	"accounting-profile (firewall)" on page 1726
"accounting-secret (RADIUS Server)" on page 1737	"accounting-server" on page 1739	"accounting-server-name" on page 1740
"accounting-stop-on-access-deny" on page 1743	"accounting-stop-on-failure" on page 1744	"acfc" on page 1748
"ack-number" on page 1750	"acknowledge-retries" on page 1751	"acknowledge-timer" on page 1752
"acquiring-state-announce-grandmaster" on page 1754	"action" on page 1759	"action (APN GTP)" on page 1760
"action (chassis error)" on page 1756	"action (OAM)" on page 1762	"action (OAM LFM)" on page 1764
"action (Policer)" on page 1765	"action (QFX)" on page 1755	"action (Security Content Security Web Filtering)" on page 1783
"action (Security Rulebase IPS)" on page 1785	"action (Uplink Failure Detection)" on page 1789	"action-priority" on page 1768
"action-priority (bridge domain)" on page 1769	"action-priority (Routing Instance)" on page 1766	"action-profile " on page 1771

"action-profile (Applying to CFM)" on page 1775	"action-profile (Applying to OAM CFM, for EX Series Switch Only)" on page 1776	"action-profile (Defining for CFM)" on page 1778
"action-profile (MEP)" on page 1780	"action-profile (Network Management)" on page 1773	"action-red-differential-delay" on page 1782
"action-shutdown" on page 1787	"actions " on page 1790	"actions (Services SSL Initiation)" on page 1793
"activation-attribute (AAA Profile)" on page 1796	"activation-delay" on page 1797	"activation-priority" on page 1799
"active" on page 1800	"active-directory-access" on page 1802	"active-interface (VPLS Multihoming)" on page 1805
"active-leasequery (DHCP Local Server)" on page 1806	"active-leasequery (DHCP Relay Agent)" on page 1808	"active-policy" on page 1811
"active-probe-params" on page 1813	"active-server-group" on page 1817	"active-source-limit" on page 1819
"activeness-probe" on page 1821	"actual-transit-statistics (Dynamic Profiles)" on page 1823	"adaptive" on page 1826
"adaptive (protocols mpls)" on page 1825	"adaptive (RSVP)" on page 1828	"adaptive-sample-rate" on page 1830
"adaptive-services" on page 1832	"adaptive-services-pics" on page 1834	"adaptive-shaper" on page 1835
"adaptive-shapers" on page 1836	"add-attribute-length-in-pdu" on page 1838	"add-path" on page 1839
"add-path-display-ipv4-address" on page 1842	"address" on page 1846	"address (Address Book Next Gen Services)" on page 1852

"address (Analytics Collector)" on page 1855	"address (Anycast RPs)" on page 1856	"address (Bidirectional Rendezvous Points)" on page 1865
"address (Diameter Peer)" on page 1874	"address (Diameter Transport)" on page 1875	"address (Interfaces)" on page 1878
"address (Interfaces, Flow Monitoring)" on page 1879	"address (IP Address)" on page 1849	"address (L2TP Destination)" on page 1880
"address (L2TP Tunnel Destination)" on page 1882	"address (LNS Local Gateway)" on page 1883	"address (Local RPs)" on page 1885
"address (LRF Profile)" on page 1886	"address (NAT Pool Next Gen Services)" on page 1894	"address (Peer)" on page 1895
"address (Protocols OpenFlow)" on page 1912	"address (Protocols Router Discovery)" on page 1913	"address (RADIUS Clients)" on page 1915
"address (RADIUS Server)" on page 1916	"address (Security Address Book)" on page 1920	"address (Security ARP Proxy)" on page 1922
"address (Security Destination NAT)" on page 1924	"address (Security NDP Proxy)" on page 1925	"address (Services)" on page 1926
"address (Services Dynamic Flow Capture)" on page 1928	"address (Services NAT Pool)" on page 1929	"address (SNMP)" on page 1844
"address (Static RPs)" on page 1933	"address (Tunnel Profile Remote Gateway)" on page 1935	"address (Tunnel Profile Source Gateway)" on page 1936
"address (tunnel-encapsulation)" on page 1843	"address-allocation" on page 1853	"address-assignment (Access)" on page 1858
"address-assignment (Address-Assignment Pools)" on page 1862	"address-blacklist" on page 1867	"address-book" on page 1868

"address-book (System)" on page 1870	"address-change-immediate-update" on page 1872	"address-family" on page 1876
"address-mapping" on page 1887	"address-mapping (Application Identification)" on page 1889	"address-mask" on page 1891
"address-name" on page 1892	"address-persistent (Security Source NAT)" on page 1896	"address-persistent (Security Source NAT Pool)" on page 1897
"address-pool" on page 1899	"address-pool (Access)" on page 1901	"address-pool (Domain Map)" on page 1902
"address-pooling" on page 1904	"address-pooling (Security Source NAT)" on page 1905	"address-pooling (Source NAT Next Gen Services)" on page 1907
"address-pools" on page 1908	"address-protection" on page 1910	"address-range" on page 1917
"address-ranges (Demux)" on page 1919	"address-set" on page 1930	"address-shared (Security Source NAT)" on page 1932
"address-whitelist" on page 1937	"addresses (REST API)" on page 1938	"adf (Dynamic Firewalls)" on page 1939
"adjacency-loss" on page 1941	"adjacency-loss-hold-time (ANCP)" on page 1943	"adjacency-timer" on page 1944
"adjust-interval" on page 1946	"adjust-minimum" on page 1947	"adjust-minimum (Dynamic Shaping and Scheduling)" on page 1949
"adjust-percent" on page 1950	"adjust-percent (Dynamic Schedulers)" on page 1952	"adjust-threshold" on page 1953
"adjust-threshold-absolute" on page 1954	"adjust-threshold-activate-bandwidth" on page 1956	"adjust-threshold-overflow-limit" on page 1958
"adjust-threshold-underflow-limit" on page 1959	"adjustment-control-profiles" on page 1961	"admin-down" on page 1962

"admin-email" on page 1963	"admin-group" on page 1967	"admin-group (Backup Selection)" on page 1965
"admin-group (for Interfaces)" on page 1971	"admin-group (for LSPs)" on page 1972	"admin-group-extended" on page 1968
"admin-groups" on page 1973	"admin-groups-extended" on page 1975	"admin-groups-extended-range" on page 1977
"admin-search" on page 1978	"administrator-email (Security Fallback Block)" on page 1980	"administrator-email (Security Virus Detection)" on page 1981
"adsl2-bytes" on page 1984	"adsl2-plus-bytes" on page 1986	"adsl-bytes" on page 1982
"advance-policy-based-routing" on page 1988	"advance-policy-based-routing-profile (Security Zones)" on page 1993	"advanced-anti-malware policy" on page 1997
"advertise" on page 2007	"advertise-bgp-static" on page 2009	"advertise-external" on page 2011
"advertise-from-main-vpn-tables" on page 2013	"advertise-inactive" on page 2015	"advertise-interval" on page 2017
"advertise-mode (MPLS)" on page 2018	"advertise-peer-as" on page 2020	"advertise-prefix-sid" on page 2022
"advertise-srv6-service" on page 2024	"advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper)" on page 2026	"advertise-zero-ethernet-tag" on page 2028
"advertisement" on page 2030	"advertisement (Delay Measurement)" on page 2031	"advertisement-hold-time" on page 2033
"advisory-options (Traffic Shaping)" on page 2035	"advpn" on page 2036	"af-mode" on page 2039

"af-name" on page 2040	"age" on page 2042	"age-of-attack" on page 2043
"agent (Analytics)" on page 2046	"agent-address" on page 2044	"agent-circuit-identifier (Dynamic ACI VLANs)" on page 2048
"agent-id" on page 2050	"agent-specifier" on page 2051	"agf" on page 2053
"aggregate (Gigabit Ethernet CoS Policer)" on page 2063	"aggregate (Hierarchical Policer)" on page 2064	"aggregate (Protocols RSVP)" on page 2068
"aggregate (Routing)" on page 2070	"aggregate-clients (DHCP Local Server)" on page 2056	"aggregate-clients (DHCP Relay Agent)" on page 2057
"aggregate-clients (Static Subscribers)" on page 2060	"aggregate-export-interval" on page 2062	"aggregate-label" on page 2066
"aggregate-wait-time" on page 2073	"aggregated-devices" on page 2074	"aggregated-ether" on page 2077
"aggregated-ether-options" on page 2078	"aggregated-inline-services-options (Aggregated Inline Services)" on page 2082	"aggregation (Forwarding Options)" on page 2084
"aggregation (IDS)" on page 2086	"aggregations (IDS Screen Next Gen Services)" on page 2088	"aging" on page 2089
"aging-timer (Junos Fusion)" on page 2091	"aigp" on page 2092	"aigp-adjust (Policy Action)" on page 2095
"aigp-originate" on page 2097	"ais-trigger-condition" on page 2099	"alarm (chassis)" on page 2100
"alarm low-light-alarm" on page 2105	"alarm (optics-options)" on page 2110	"alarm (Satellite Policies)" on page 2114
"alarm (SNMP RMON)" on page 2115	"alarm-id" on page 2102	"alarm-list-name" on page 2103

"alarm-management" on page 2106	"alarm-mode" on page 2108	"alarm-port" on page 2111
"alarm-state" on page 2117	"alarm-threshold" on page 2119	"alarm-threshold (Security GPRS)" on page 2120
"alarm-without-drop" on page 2122	"alarm-without-drop (IDS Screen Next Gen Services)" on page 2123	"alarms" on page 2124
"alarms (Security)" on page 2128	"alarms (Video Monitoring)" on page 2125	"alg" on page 2130
"alg (Applications)" on page 2138	"alg-manager" on page 2141	"alg-support-lib" on page 2142
"algorithm" on page 2144	"algorithm (Junos FIPS)" on page 2145	"alias (Interfaces)" on page 2146
"alias (Junos Fusion)" on page 2148	"alias (Junos Fusion Satellite Device Clustering)" on page 2150	"aliases" on page 2151
"aliases (Virtual Chassis)" on page 2153	"all-defects" on page 2155	"all-regions" on page 2157
"all-tcp" on page 2159	"allow" on page 2160	"allow-6pe-traceroute" on page 2162
"allow-active-leasequery (DHCP Local Server)" on page 2163	"allow-bulk-leasequery (DHCP Local Server)" on page 2165	"allow-clear (Accounting Options)" on page 2167
"allow-commands (Op Scripts)" on page 2169	"allow-dns-reply" on page 2170	"allow-duplicates" on page 2172
"allow-dynamic-requests (RADIUS Server)" on page 2173	"allow-email (Security Fallback Block)" on page 2175	"allow-email (Security Virus Detection)" on page 2176
"allow-embedded-icmp" on page 2177	"allow-fragmentation" on page 2179	"allow-fragmentation (fti)" on page 2181

"allow-fragmentation (protocols mpls)" on page 2178	"allow-icmp-without-flow" on page 2182	"allow-ip-options (IDS MS-MPC)" on page 2183
"allow-ip-options (Services Stateful Firewall)" on page 2186	"allow-ipv6-extension-header (IDS MS-MPC)" on page 2188	"allow-l3vpn-traceroute-src-select" on page 2190
"allow-leasequery (DHCP Local Server)" on page 2191	"allow-maximum (Multicast)" on page 2193	"allow-multicast" on page 2195
"allow-no-end-option (DHCP Relay Agent)" on page 2196	"allow-overlapping-nat-pools" on page 2198	"allow-overlapping-pools (NAT Next Gen Services)" on page 2199
"allow-remote-loopback" on page 2200	"allow-reverse-ecmp" on page 2201	"allow-route-leaking" on page 2202
"allow-snooped-clients" on page 2204	"allow-sram-parity-errors" on page 2206	"allow-subnet-mismatch" on page 2207
"allow-transients" on page 2208	"allow-url-for-python (Op Scripts)" on page 2209	"allow-v4mapped-packets" on page 2211
"allowed-destinations" on page 2212	"allowed-mac" on page 2214	"allowed-sources (REST API)" on page 2216
"alt-name (Application Identification)" on page 2217	"alternative-partition-name (OCS Partition)" on page 2218	"always-mark-connection-protection-tlv" on page 2219
"always-write-giaddr" on page 2221	"always-write-option-82" on page 2222	"ambient-temperature" on page 2224
"amf" on page 2226	"amt (IGMP)" on page 2228	"amt (Protocols)" on page 2230
"analytics" on page 2231	"analyzer" on page 2237	"analyzer-address" on page 2241
"analyzer-id" on page 2243	"anchor-point (Pseudowire Subscriber Interfaces)" on page 2244	"ancp" on page 2247

"ancp (Adjustment Control Profiles)" on page 2249	"ancp-speed-change-immediate-update (ANCP)" on page 2251	"anomaly" on page 2254
"anti-replay-window-size" on page 2255	"anti-replay-window-size (Services IPsec VPN)" on page 2256	"anti-replay-window-size (Services Service Set)" on page 2258
"anti-spam " on page 2262	"anti-spam (Logical Systems)" on page 2259	"anti-spam (Security Content Security Policy)" on page 2265
"anti-virus (Content Security)" on page 2269	"anti-virus (Content Security Policy)" on page 2273	"anti-virus (Logical Systems)" on page 2266
"anti-virus (Services)" on page 2275	"any (VPLS Multihoming)" on page 2282	"any-sender" on page 2281
"anycast-pim" on page 2283	"anycast-prefix" on page 2285	"apn" on page 2286
"apn-control (Security GTP)" on page 2289	"apn-control-group (Security GTP)" on page 2288	"apn-services (control plane services)" on page 2292
"app-mapping-timeout" on page 2293	"appfw-profile (System)" on page 2294	"appfw-rule" on page 2296
"appfw-rule-set" on page 2298	"application" on page 2299	"application (Adjustment Control Profiles)" on page 2301
"application (Application Identification)" on page 2303	"application (Application Identification)" on page 2306	"application (Application Maps)" on page 2308
"application (Applications)" on page 2310	"application (Extensions)" on page 2316	"application (NAT Next Gen Services)" on page 2348
"application (Security Alarms)" on page 2366	"application (Security Custom Attack)" on page 2367	"application (Security Destination NAT)" on page 2369
"application (Security IDP)" on page 2370	"application (Security Policies)" on page 2371	"application (Security Source NAT)" on page 2373

"application-bypass (Juniper Secure Connect)" on page 2314	"application-firewall" on page 2318	"application-firewall (Application Services)" on page 2320
"application-group" on page 2322	"application-group (Services)" on page 2325	"application-group-any (AACL)" on page 2323
"application-groups (PCC Rules)" on page 2326	"application-groups (Services AACL)" on page 2328	"application-identification (Application Identification)" on page 2337
"application-identification (IDP Sensor Configuration)" on page 2336	"application-identification (Services)" on page 2330	"application-identification-profile (Service Set)" on page 2343
"application-map" on page 2345	"application-maps" on page 2346	"application-profile" on page 2349
"application-profile (Services CoS Next Gen Services)" on page 2351	"application-protocol" on page 2352	"application-protocol (Applications)" on page 2355
"application-proxy (Security Content Security)" on page 2358	"application-screen (Security H323)" on page 2359	"application-screen (Security MGCP)" on page 2361
"application-screen (Security SCCP)" on page 2362	"application-screen (Security SIP)" on page 2364	"application-services (Security Forwarding Process)" on page 2375
"application-services (Security Policies)" on page 2378	"application-services (security-metadata-streaming)" on page 2377	"application-set" on page 2382
"application-sets (IDS MS-DPC)" on page 2383	"application-sets (PCP)" on page 2384	"application-sets (Services CoS)" on page 2386
"application-sets (Services NAT)" on page 2387	"application-sets (Services Stateful Firewall) " on page 2388	"application-specific (Protocols OSPF)" on page 2389
"application-status" on page 2392	"application-system-cache" on page 2393	"application-system-cache-timeout (Services)" on page 2395

"application-tracking" on page 2397	"application-tracking (Security Zones)" on page 2399	"application-traffic-control" on page 2400
"application-traffic-control (Application Services)" on page 2402	"application-unknown" on page 2404	"applications (Applications)" on page 2405
"applications (DCBX)" on page 2406	"applications (Diameter)" on page 2408	"applications (IDS MS-DPC)" on page 2409
"applications (PCC Rules)" on page 2410	"applications (PCP)" on page 2412	"applications (Services AACL)" on page 2414
"applications (Services ALGs)" on page 2415	"applications (Services Application Identification)" on page 2416	"applications (Services CoS)" on page 2418
"applications (Services NAT)" on page 2419	"applications (Services Stateful Firewall)" on page 2420	"apply-action-profile" on page 2421
"apply-groups" on page 2423	"apply-groups (Chassis Cluster)" on page 2424	"apply-groups (Subscriber Secure Policy)" on page 2428
"apply-groups-except" on page 2426	"apply-groups-except (Subscriber Secure Policy)" on page 2427	"apply-macro" on page 2430
"apply-path" on page 2431	"apply-secondary-as-giaddr" on page 2433	"apply-to-half-close-state" on page 2435
"apts" on page 2436	"archival" on page 2438	"archive (All System Log Files)" on page 2440
"archive (Individual System Log File)" on page 2442	"archive (QFabric System)" on page 2445	"archive-sites" on page 2446
"archive-sites (Event Policy)" on page 2449	"archive-sites (Flow Monitoring)" on page 2448	"area" on page 2451

"area-range" on page 2453	"arguments (Event Policy)" on page 2455	"arguments (Op Scripts)" on page 2457
"arp" on page 2458	"arp policer" on page 2475	"arp-detect" on page 2463
"arp-enhanced-scale" on page 2465	"arp-inspection" on page 2466	"arp-inspection (MX Series)" on page 2469
"arp-l2-validate" on page 2470	"arp-max-cache" on page 2471	"arp-nd-probe-disable" on page 2484
"arp-nd-probe-failed-log" on page 2485	"arp-new-hold-limit" on page 2473	"arp-prefix-limit (Host Fast Reroute)" on page 2477
"arp-resp" on page 2479	"arp-system-cache-limit" on page 2480	"arp-throttle" on page 2482
"as-external" on page 2486	"as-list" on page 2487	"as-override" on page 2489
"as-path (Policy Options)" on page 2493	"as-path (Routing Options)" on page 2495	"as-path-compare" on page 2490
"as-path-group" on page 2492	"asm-override-ssm" on page 2497	"assemble" on page 2499
"assert-timeout" on page 2500	"identity-provider attribute-mapping (SAML)" on page 2501	"assignment-id-format (L2TP LAC)" on page 2504
"assisted-replication" on page 2505	"assisted-replication (Deterministic AR Replicator Policy Actions)" on page 2507	"associate-backup-pe-groups" on page 2511
"associate-lsp" on page 2512	"associate-profile" on page 2513	"association-timeout" on page 2515
"asymmetric-hold-time" on page 2516	"asymmetric-hold-time" on page 2517	"asymmetric-lease-time (DHCP Overrides)" on page 2519

"asymmetric-prefix-lease-time (DHCP Overrides)" on page 2521	"asynchronous-notification" on page 2522	"atm-options" on page 2524
"attach" on page 2527	"attack-threshold" on page 2528	"attack-type (Security Anomaly)" on page 2529
"attack-type (Security Chain)" on page 2531	"attack-type (Security IDP)" on page 2533	"attack-type (Security Signature)" on page 2540
"attacks (Security Exempt Rulebase)" on page 2546	"attacks (Security IPS Rulebase)" on page 2547	"attempts (DHCP Local Server)" on page 2549
"attribute" on page 2551	"attributes (Access-Line Rate Adjustment)" on page 2553	"attributes (Diameter Gx Profiles)" on page 2555
"attributes (JSRC Attributes)" on page 2557	"attributes (RADIUS Attributes)" on page 2560	"attributes-match (Event Policy)" on page 2558
"audible (Security Alarms)" on page 2563	"auth-entry" on page 2564	"auth-only-browser" on page 2566
"auth-user-agent" on page 2568	"authentication" on page 2574	"authentication (AAA Profile)" on page 2576
"authentication (BGP BFD Liveness Detection)" on page 2588	"authentication (Demux)" on page 2592	"authentication (DHCP Local Server)" on page 2593
"authentication (DHCP Relay Agent)" on page 2595	"authentication (LAG)" on page 2622	"authentication (Protocols OSPF)" on page 2570
"authentication (Protocols PIM)" on page 2647	"authentication (Security Alarms)" on page 2648	"authentication (Security IPsec)" on page 2650
"authentication (services)" on page 2572	"authentication (Static and Dynamic PPP)" on page 2664	"authentication (Static Subscribers)" on page 2665

"authentication-algorithm" on page 2577	"authentication-algorithm (Security IKE)" on page 2580	"authentication-algorithm (Security IPsec)" on page 2582
"authentication-algorithm (Services IKE)" on page 2584	"authentication-algorithm (Services IPsec)" on page 2585	"authentication-entry-timeout (Identity Management Advanced Query)" on page 2597
"authentication-entry-timeout (Services User Identification)" on page 2599	"authentication-key" on page 2603	"authentication-key (ICCP)" on page 2614
"authentication-key (Protocols BGP and BMP)" on page 2616	"authentication-key (Protocols LDP)" on page 2618	"authentication-key (Protocols MSDP)" on page 2600
"authentication-key (Protocols RIP)" on page 2619	"authentication-key (Protocols RSVP)" on page 2620	"authentication-key (SNMP)" on page 2602
"authentication-key-chain (Protocols BGP and BMP)" on page 2605	"authentication-key-chain (Protocols LDP)" on page 2607	"authentication-key-chain (TWAMP)" on page 2608
"authentication-key-chains" on page 2610	"authentication-key-chains (TCP-AO)" on page 2612	"authentication-md5" on page 2624
"authentication-method" on page 2627	"authentication-method (Interfaces)" on page 2629	"authentication-method (services)" on page 2626
"authentication-mode" on page 2630	"authentication-none" on page 2632	"authentication-order" on page 2633
"authentication-order (Access Profile)" on page 2635	"authentication-order (Authenticator)" on page 2638	"authentication-order (System)" on page 2641
"authentication-password" on page 2643	"authentication-protocol" on page 2645	"authentication-server" on page 2652

"authentication-sha" on page 2653	"authentication-sha224" on page 2654	"authentication-sha256" on page 2656
"authentication-sha348" on page 2657	"authentication-sha512" on page 2658	"authentication-source (Services User Identification ClearPass)" on page 2660
"authentication-source (Services User Identification Device Identity)" on page 2662	"authentication-type" on page 2667	"authentication-type (Protocols RIP)" on page 2669
"authentication-whitelist" on page 2671	"authenticator" on page 2673	"authorization (icap-redirect profile)" on page 2678
"authorization (SNMP)" on page 2677	"authorization-order" on page 2680	"auto-sense" on page 2738
"auto-bandwidth" on page 2682	"auto-bandwidth (MPLS Statistics)" on page 2684	"auto-bandwidth (MPLS Tunnel)" on page 2686
"auto-configuration" on page 2688	"auto-configuration (System)" on page 2689	"auto-configure" on page 2692
"auto-configure (Demux)" on page 2694	"auto-configure (Dynamic VLAN Interface Sets)" on page 2696	"auto-configure (IPv4)" on page 2698
"auto-configure (IPv6)" on page 2700	"auto-configure-trigger interface (ANCP)" on page 2702	"auto-conversion" on page 2704
"auto-dad (SLAAC Snooping)" on page 2705	"auto-derive" on page 2707	"auto-discovery" on page 2709
"auto-discovery-only" on page 2711	"auto-export" on page 2712	"auto-image-upgrade" on page 2715
"auto-load-rebalance" on page 2717	"auto-negotiation" on page 2719	"auto-negotiation (QFX Series)" on page 2718

"auto-policing" on page 2724	"auto-provisioned" on page 2725	"auto-re-enrollment" on page 2727
"auto-re-enrollment (Security)" on page 2728	"auto-reconnect" on page 2732	"auto-recovery" on page 2733
"auto-rp" on page 2735	"auto-satellite-conversion (Junos Fusion)" on page 2736	"auto-service-id" on page 2739
"auto-snapshot" on page 2741	"auto-sw-sync" on page 2743	"auto-sw-update" on page 2745
"auto-targeted-session" on page 2748	"autodiscovery" on page 2750	"autodiscovery-only" on page 2752
"autoinstallation" on page 2753	"autoinstallation (JNU Satellite Devices)" on page 2755	"automatic (Security)" on page 2757
"automatic (Source NAT Next Gen Services)" on page 2760	"automatic-site-id" on page 2758	"autonomous" on page 2761
"autonomous (Dynamic Router Advertisement)" on page 2763	"autonomous-system" on page 2764	"autonomous-system-type (Flow Monitoring)" on page 2767
"autostate-exclude" on page 2769	"auxiliary" on page 2771	"auxiliary-spi" on page 2773
"auxiliary-spi (Security IPsec)" on page 2774	"avg-fd-twoway-threshold" on page 2775	"avg-flr-backward-threshold" on page 2777
"avg-flr-forward-threshold" on page 2779	"avg-ifdv-twoway-threshold" on page 2781	"avira-engine" on page 2783
"avp (L2TP Tunnel Switching)" on page 2784		

B

"backoff-period" on page 2791	"backup (MBGP MVPN)" on page 2800	"backup (OCS Partition)" on page 2804
"backup-bridge-priority" on page 2793	"backup-destination" on page 2795	"backup-interface" on page 2796
"backup-interface (Layer 2 Circuits)" on page 2797	"backup-liveness-detection" on page 2799	"backup-neighbor" on page 2802
"backup-on-failure (Accounting Options)" on page 2806	"backup-options" on page 2807	"backup-pe-group" on page 2808
"backup-peer-ip" on page 2810	"backup-remote-gateway" on page 2811	"backup-router" on page 2812
"backup-selection (Protocols IS-IS)" on page 2814	"backup-selection (Protocols OSPF)" on page 2819	"backup-spf-options (Protocols IS-IS)" on page 2821
"backup-spf-options (Protocols OSPF)" on page 2823	"backups" on page 2826	"backward-frr-enable" on page 2827
"bad-inner-header" on page 2829	"bad-option (IDS Screen Next Gen Services)" on page 2830	"bandwidth (DDoS)" on page 2839
"bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)" on page 2843	"bandwidth (Inline Services)" on page 2845	"bandwidth (Interfaces)" on page 2847
"bandwidth (Interfaces)" on page 2849	"bandwidth (MPC Bandwidth)" on page 2863	"bandwidth (multicast)" on page 2831
"bandwidth (Protocols Layer 2 Circuit)" on page 2870	"bandwidth (Protocols RSVP)" on page 2872	"bandwidth (RFC 2544 Benchmarking)" on page 2874

"bandwidth (Static LSP)" on page 2876	"bandwidth (storm control)" on page 2833	"bandwidth (Tunnel Services)" on page 2880
"bandwidth-based-metrics" on page 2837	"bandwidth-based-metrics (Protocols OSPF)" on page 2835	"bandwidth-degradation" on page 2841
"bandwidth-kbps" on page 2850	"bandwidth-level" on page 2851	"bandwidth-limit" on page 2853
"bandwidth-limit (Hierarchical Policer)" on page 2855	"bandwidth-limit (Policer)" on page 2857	"bandwidth-limit (Policer for Gigabit Ethernet Interfaces)" on page 2860
"bandwidth-model" on page 2861	"bandwidth-percent" on page 2866	"bandwidth-percent (protocols mpls)" on page 2864
"bandwidth-percentage" on page 2869	"bandwidth-scale (DDoS)" on page 2875	"bandwidth-threshold" on page 2878
"banner (Access FTP HTTP Telnet Authentication)" on page 2882	"banner (Access Web Authentication)" on page 2884	"base-distinguished-name" on page 2885
"batch query" on page 2886	"bb-sc-n" on page 2889	"beacon-period" on page 2890
"bearer-type (L2TP Tunnel Switching)" on page 2892	"ber-threshold-clear" on page 2894	"ber-threshold-signal-degrade" on page 2897
"bert-algorithm" on page 2900	"bert-error-rate" on page 2903	"bert-period" on page 2905
"best-site" on page 2908	"bfd" on page 2909	"bfd (protocols)" on page 2911
"bfd-liveliness" on page 2914	"bfd-liveliness (Activeness Probe)" on page 2917	"bfd-liveness-detection (BGP)" on page 2918
"bfd-liveness-detection (EVPN Node Detection)" on page 2923	"bfd-liveness-detection (LAG)" on page 2926	"bfd-liveness-detection (Layer 2 VPN and VPLS)" on page 2929

"bfd-liveness-detection (LSP)" on page 2931	"bfd-liveness-detection (Protocols IS-IS)" on page 2933	"bfd-liveness-detection (Protocols LDP)" on page 2937
"bfd-liveness-detection (Protocols MPLS)" on page 2939	"bfd-liveness-detection (Protocols OSPF)" on page 2941	"bfd-liveness-detection (Protocols PIM)" on page 2945
"bfd-liveness-detection (Protocols RIP)" on page 2947	"bfd-liveness-detection (Routing Options Static Route)" on page 2950	"bfd-liveness-detection (Source Packet Routing Template)" on page 2955
"bgp" on page 2956	"bgp (RPM)" on page 2958	"bgp-error-tolerance (Protocols BGP)" on page 2959
"bgp-identifier" on page 2962	"bgp-orf-cisco-mode" on page 2963	"bgp-peer" on page 2965
"bgp-static" on page 2966	"bgpmcast" on page 2968	"bidirectional (Interface)" on page 2969
"bidirectional (RP)" on page 2971	"bier" on page 2972	"bier-sub-domain (Protocols IS-IS)" on page 2974
"bits (Chassis Synchronization)" on page 2975	"black-list" on page 2978	"blackhole-action" on page 2979
"block" on page 2980	"block (Spanning Trees)" on page 2995	"block-allocation (Source NAT Next Gen Services)" on page 2983
"block-command" on page 2985	"block-content-type" on page 2986	"block-extension" on page 2987
"block-frag (IDS Screen Next Gen Services)" on page 2988	"block-message (Application Firewall)" on page 2989	"block-message (Security Content Security)" on page 2993
"block-mime" on page 2994	"blocklists (Security IKE)" on page 2997	"bmp" on page 3000

"boot-loader-authentication" on page 3008	"boot-server (DHCP)" on page 3010	"bootp" on page 3012
"bootp (System Processes)" on page 3013	"bootstrap" on page 3015	"bootstrap-export" on page 3016
"bootstrap-import" on page 3018	"bootstrap-priority" on page 3019	"bounded-delay" on page 3021
"bpdu-block" on page 3022	"bpdu-block-on-edge" on page 3024	"bpdu-destination-mac-address" on page 3026
"bpdu-destination-mac-address (Spanning Tree)" on page 3027	"bridge-domain" on page 3029	"bridge-domain (Analyzer)" on page 3031
"bridge-domains" on page 3032	"bridge-options" on page 3034	"bridge-priority" on page 3036
"bridge-template" on page 3038	"brief" on page 3039	"broadcast (Forwarding class)" on page 3044
"broadcast (Interfaces)" on page 3041	"broadcast (Router Discovery)" on page 3043	"bucket-size" on page 3046
"buffer-dynamic-threshold" on page 3047	"buffer-partition (Egress)" on page 3049	"buffer-partition (Ingress)" on page 3052
"buffer-rate (Schedulers)" on page 3065	"buffer-size" on page 3055	"buffer-size (Dynamic Scheduling)" on page 3061
"buffer-size (Schedulers)" on page 3063	"bulk-call-mode" on page 3067	"bulk-leasequery (DHCP Relay Agent)" on page 3068
"bum-hashing" on page 3070	"bundle" on page 3072	"bundle (Interfaces)" on page 3074
"bundle (multilink)" on page 3073	"burst (DDoS)" on page 3075	"burst-scale (DDoS)" on page 3077

"burst-size" on page 3078	"burst-size (Default Local Policy)" on page 3081	"burst-size (TDF Domain)" on page 3090
"burst-size-limit" on page 3083	"burst-size-limit (Hierarchical Policer)" on page 3084	"burst-size-limit (Policer)" on page 3086
"burst-size-limit (Policer for Gigabit Ethernet Interfaces)" on page 3089	"by-destination (IDS MS-DPC)" on page 3092	"by-destination (IDS MS-MPC)" on page 3093
"by-destination (IDS Screen Next Gen Services)" on page 3095	"by-pair (IDS MS-DPC)" on page 3098	"by-protocol (IDS MS-MPC)" on page 3099
"by-protocol (IDS Screen Next Gen Services)" on page 3102	"by-source (IDS MS-DPC)" on page 3105	"by-source (IDS MS-MPC)" on page 3106
"by-source (IDS Screen Next Gen Services)" on page 3108	"bypass" on page 3111	"bypass (Signaled LSP)" on page 3115
"bypass (Static LSP)" on page 3117	"bypass-aggregate (DDoS)" on page 3112	"bypass-queuing-chip" on page 3114
"bypass-traffic-on-exceeding-flow-limits" on page 3119	"bypass-traffic-on-pic-failure" on page 3120	"bytes (Dynamic Traffic Shaping)" on page 3121
"bytes (otn-options)" on page 3123		

C

"c-timeout" on page 3152	"ca-chain (REST API)" on page 3154	"ca-identity" on page 3155
"ca-name" on page 3156	"ca-profile (Security PKI)" on page 3157	"ca-profile (Services)" on page 3160

"ca-type" on page 3161	"cable-length (Congestion Notification)" on page 3164	"cac (TDF Gateway)" on page 3165
"cache" on page 3167	"cache (Security Log)" on page 3168	"cache-size" on page 3170
"cache-timeout-negative" on page 3171	"cacheable (Application Identification)" on page 3173	"cak" on page 3174
"cak (MX Series)" on page 3176	"calculation-weight" on page 3177	"call-flood" on page 3179
"call-rate-statistics" on page 3180	"callback" on page 3182	"callback-wait-period" on page 3183
"called-station-id" on page 3185	"called-station-id (OCS Partition)" on page 3186	"caller" on page 3188
"calling-number" on page 3189	"calling-number (L2TP Tunnel Switching)" on page 3190	"calling-station-id" on page 3192
"calling-station-id-format (Subscriber Management)" on page 3193	"candidate-uplink-port-policy (satellite-policies)" on page 3196	"capabilities (Protocols RIFT)" on page 3197
"captive-portal" on page 3199	"captive-portal (Services UAC)" on page 3203	"captive-portal (Services UAC Policy)" on page 3205
"captive-portal-content-delivery-profile (Services)" on page 3202	"capture-file (Security)" on page 3206	"capture-group" on page 3208
"casb" on page 3210	"casb-policy" on page 3212	"cascade-port" on page 3214
"cascade-ports" on page 3216	"cascade-ports (Junos Fusion Satellite Device Cluster)" on page 3218	"categories" on page 3219

"category" on page 3221	"category (Security Dynamic Attack Group)" on page 3223	"category (Security Logging)" on page 3224
"category (Security Web Filtering)" on page 3226	"category (System Logging)" on page 3236	"cbr" on page 3238
"cbr rate" on page 3240	"ccc" on page 3241	"ccc (Routing Policy Condition)" on page 3242
"cdvt" on page 3244	"cell-mode (Dynamic Traffic Shaping)" on page 3245	"cell-overhead" on page 3248
"cellular-options" on page 3249	"centralized" on page 3250	"cert-file (Scripts)" on page 3252
"certificate" on page 3254	"certificate (Juniper Secure Connect)" on page 3259	"certificate-authority (REST API)" on page 3256
"certificate-id" on page 3257	"certificate-verification" on page 3260	"certificates" on page 3262
"certification-authority" on page 3263	"ces-psn-channel (tdm-options)" on page 3265	"cfeb" on page 3267
"cflowd (Discard Accounting)" on page 3268	"cflowd (Flow Monitoring)" on page 3270	"cfp-to-et" on page 3272
"cgn-pic" on page 3273	"chain" on page 3274	"chain-order (Application Identification)" on page 3276
"chained-composite-next-hop" on page 3277	"challenge-length (Static and Dynamic PPP)" on page 3281	"challenge-password" on page 3283
"change-configuration (Event Policy)" on page 3284	"channel" on page 3286	"channel-speed" on page 3287
"chap" on page 3289	"chap (Dynamic PPP)" on page 3291	"chap (L2TP)" on page 3292

"charging-id (OCS Partition)" on page 3293	"charging-service-list" on page 3295	"chassis" on page 3296
"chassis (QFabric System)" on page 3303	"chassis (Subscriber Limits)" on page 3306	"chassis-id (Aggregated Ethernet Interface)" on page 3302
"chassis-id (LLDP)" on page 3301	"check-bytes (Application Identification)" on page 3308	"check-zero" on page 3309
"checksum (Scripts)" on page 3311	"checksum-validate" on page 3313	"child-inactivity-timeout" on page 3314
"children" on page 3315	"cipher-list (REST API)" on page 3318	"cipher-suite (MACsec)" on page 3320
"circuit-id" on page 3322	"circuit-id (DHCP Relay Agent)" on page 3324	"circuit-id (VLAN Authentication Username)" on page 3328
"circuit-type" on page 3329	"circuit-type (DHCP Local Server)" on page 3330	"circuit-type (DHCP Relay Agent)" on page 3332
"cisco-interoperability" on page 3334	"cisco-nas-port-info (L2TP Tunnel Switching)" on page 3335	"civic-based" on page 3337
"ckn" on page 3339	"ckn (MX Series)" on page 3341	"class" on page 3342
"class (CoS-Based Forwarding)" on page 3346	"class (Defining Login Classes)" on page 3347	"class (Forwarding Class Sets)" on page 3361
"class (Forwarding classes)" on page 3362	"class (Forwarding Classes)" on page 3364	"class (Unified Edge)" on page 3345
"class-of-service (Dynamic Profiles)" on page 3366	"class-of-service (Protocols MPLS)" on page 3369	"class-of-service (Protocols RSVP)" on page 3370
"class-of-service (saegw system)" on page 3372	"class-usage-profile" on page 3373	"classification-key (DHCP Local Server)" on page 3375

"classification-key (DHCP Relay Agent)" on page 3376	"classification-override" on page 3379	"classifier" on page 3380
"classifiers (Class of Service)" on page 3383	"classifiers (Definition)" on page 3386	"classifiers (Definition)" on page 3388
"classifiers (Dynamic CoS Application)" on page 3390	"classifiers (Interfaces)" on page 3391	"classifiers (Physical Interface)" on page 3393
"classifiers (Routing Instance)" on page 3395	"classifiers (VRF)" on page 3382	"clat-ipv6-prefix-length" on page 3396
"clat-prefix" on page 3398	"clat-prefix (Source NAT Next Gen Services)" on page 3399	"cleanup-interval (Accounting Options)" on page 3400
"clear" on page 3402	"clear-action (CFM)" on page 3409	"clear-dont-fragment-bit (fti)" on page 3556
"clear-dont-fragment-bit (Interfaces GRE Tunnels) " on page 3557	"clear-dont-fragment-bit (NAT Next Gen Services)" on page 3559	"clear-dont-fragment-bit (Services IPsec VPN)" on page 3560
"clear-dont-fragment-bit (Services NAT Options)" on page 3561	"clear-dont-fragment-bit (Services Service Set)" on page 3562	"clear-ike-sas-on-pic-restart" on page 3631
"clear-ipsec-sas-on-pic-restart" on page 3656	"clear-on-abort (DHCP Local Server)" on page 3782	"clear-threshold" on page 4245
"client" on page 4292	"client (interfaces)" on page 4299	"client (System Services)" on page 4339
"client (TWAMP)" on page 4295	"client (Unified Edge)" on page 4300	"client-address" on page 4302
"client-config (Juniper Secure Connect)" on page 4303	"client-delegate-probes" on page 4307	"client-discover-match (DHCP Local Server)" on page 4308

"client-discover-match (DHCP Relay Agent)" on page 4310	"client-group" on page 4312	"client-ia-type" on page 4313
"client-id (DHCP Local Server)" on page 4315	"client-id (DHCP Relay Agent)" on page 4317	"client-identifier (DHCP Client)" on page 4319
"client-identifier (DHCPv6 Client)" on page 4321	"client-identifier (Interfaces)" on page 4322	"client-identity (NETCONF TLS)" on page 4324
"client-idle-timeout (Access Profile)" on page 4326	"client-list" on page 4328	"client-list (SNMP)" on page 4329
"client-list-name" on page 4331	"client-name-filter" on page 4332	"client-negotiation-match (DHCPv6 Local Server)" on page 4333
"client-negotiation-match (DHCPv6 Relay Agent)" on page 4335	"client-response-ttl" on page 4336	"client-session-timeout (Access Profile)" on page 4338
"client-type" on page 4340	"client-type (Subscriber Limits)" on page 4342	"clients" on page 4345
"clients (Security) " on page 4347	"clients (SNMP Community)" on page 4344	"clns-routing" on page 4349
"clns-updown-compatibility" on page 4350	"clock-class" on page 4352	"clock-class-to-quality-level-mapping" on page 4353
"clock-ipv6-client" on page 4355	"clock-ipv6-source " on page 4356	"clock-mode" on page 4357
"clock-mode (Clock Synchronization)" on page 4359	"clock-rate" on page 4360	"clock-source (hybrid)" on page 4362
"clocking-mode" on page 4364	"close-timeout" on page 4365	"cluster" on page 4367
"cluster (Chassis)" on page 4369	"cluster (Junos Fusion)" on page 4375	"cluster-id (Junos Fusion Satellite Device Cluster)" on page 4373

"cmcast-joins-limit-inet (MVPN Selective Tunnels)" on page 4376	"cmcast-joins-limit-inet6 (MVPN Selective Tunnels)" on page 4379	"coa-accounting (AAA Profile)" on page 4381
"coa-immediate-update" on page 4382	"coa-no-override service-class-attribute" on page 4384	"code" on page 4385
"code (AAA Profile)" on page 4388	"code (Application Identification)" on page 4389	"code (IDP)" on page 4386
"code-point" on page 4390	"code-point (Congestion Notification)" on page 4393	"code-point (Fibre Channel Interfaces)" on page 4395
"code-point (Input Congestion Notification)" on page 4397	"code-point (Output Congestion Notification)" on page 4398	"code-point (Rewrite Rules)" on page 4400
"code-point-aliases" on page 4392	"code-points (Application Maps)" on page 4402	"code-points (Classifiers)" on page 4403
"code-points (Classifiers)" on page 4405	"collapsed-forward" on page 4406	"collector" on page 4408
"collector (Analytics)" on page 4412	"collector (Flow Collection)" on page 4410	"collector (Flow Monitoring Logs for NAT)" on page 4413
"collector (Flow Template Profiles for NAT)" on page 4415	"collector (Inline Monitoring)" on page 4420	"collector (LRF Profile)" on page 4422
"collector (LRF Rule)" on page 4424	"collector-group (Flow Monitoring Logs for NAT)" on page 4416	"collector-group (Flow Template Profiles for NAT)" on page 4418
"color" on page 4425	"color-aware" on page 4428	"color-blind" on page 4429
"command (Op Scripts)" on page 4431	"commands (Event Policy Change Configuration)" on page 4432	"commands (Event Policy Execute Commands)" on page 4434
"commit activate" on page 4442	"commit (Commit Scripts)" on page 4444	"commit prepare" on page 4451

"commit (System)" on page 4455	"commit-delay" on page 4446	"commit-interval" on page 4447
"commit-options" on page 4449	"commit-synchronize-server" on page 4452	"committed-burst-size" on page 4458
"committed-information-rate" on page 4460	"community" on page 4462	"community (Policy Options)" on page 4465
"community (Protocols Layer 2 Circuit)" on page 4469	"community (RMON)" on page 4471	"community (Routing Options)" on page 4472
"community (SNMP)" on page 4475	"community-name" on page 4464	"community-vlan" on page 4477
"community-vlans (MX Series)" on page 4478	"compatibility (Application Identification)" on page 4480	"compatibility-version" on page 4481
"compliance (Juniper Secure Connect)" on page 4483	"compress (Accounting Options)" on page 4488	"compress-configuration-files (System)" on page 4489
"compression" on page 4491	"compression (PPP Properties)" on page 4494	"compression-device (Interfaces)" on page 4492
"compression-device (Interfaces)" on page 4493	"compute-options" on page 4496	"compute-profile" on page 4497
"concurrent-data-sessions" on page 4500	"condition" on page 4502	"conditional-match" on page 4504
"conditional-metric" on page 4506	"confederation" on page 4508	"configuration-database" on page 4510
"configuration-database (Enhanced Subscriber Management)" on page 4513	"configuration-file" on page 4514	"configuration-name" on page 4516
"configuration-servers" on page 4517	"configuration-synchronize (Chassis Cluster)" on page 4519	"configured-flow-control" on page 4523

"congestion-notification-profile" on page 4525	"connect-actively" on page 4529	"connection (Identity Management Advanced Query)" on page 4533
"connection-flood" on page 4531	"connection-limit" on page 4537	"connection-limit (REST API)" on page 4538
"connection-protection" on page 4540	"connection-protection-tlv" on page 4541	"connections (Circuits)" on page 4542
"connections (Host VPN)" on page 4544	"connections (MPLS)" on page 4547	"connectivity-association" on page 4548
"connectivity-association (MACsec Interfaces)" on page 4551	"connectivity-association (MACsec Interfaces for MX Series)" on page 4553	"connectivity-association (MX Series)" on page 4554
"connectivity-fault-management" on page 4556	"connectivity-fault-management-switches " on page 4560	"connectivity-type" on page 4563
"console (System Logging)" on page 4565	"console (System Ports)" on page 4567	"constant" on page 4570
"constituent-list" on page 4571	"cont-stats-collection-interval" on page 4573	"contact (SNMP)" on page 4575
"container-devices" on page 4576	"container-label-switched-path" on page 4577	"Content Security" on page 4589
[OBSOLETE] Content Security-policy	"Content Security-policy (Application Services)" on page 4600	"content-destination" on page 4579
"content-filtering (Logical System Security Feature Profile)" on page 4580	"content-filtering (Security Content Security Policy)" on page 4583	"content-filtering (Security Feature Profile)" on page 4586

"content-size" on page 4601	"content-size (Security Antivirus Sophos Engine)" on page 4604	"content-size-limit" on page 4603
"context (Application Identification)" on page 4606	"context (Security Custom Attack)" on page 4614	"context (SNMPv3)" on page 4615
"context-identifier (Protocols IS-IS)" on page 4611	"context-identifier (Protocols OSPF)" on page 4613	"continuity-check" on page 4616
"continuity-check (EX Series Switch Only)" on page 4618	"control (REST API)" on page 4639	"control-channel" on page 4620
"control-channel (Protocols Link Management Peer)" on page 4621	"control-channel (Protocols OAM)" on page 4623	"control-connection (Junos OS)" on page 4624
"control-connection (Junos OS Evolved)" on page 4627	"control-cores" on page 4631	"control-link-recovery" on page 4633
"control-plane-bandwidth-percent (Node Slicing)" on page 4634	"control-polarity" on page 4636	"control-ports" on page 4637
"control-signal" on page 4640	"control-source" on page 4641	"control-vlan" on page 4643
"control-word" on page 4644	"control-word (Protocols Layer 2 Circuit Neighbor)" on page 4647	"control-word (Protocols Layer 2 VPN)" on page 4648
"controller" on page 4650	"controller (OVSDB)" on page 4652	"controller (Protocols OpenFlow)" on page 4654
"convert-clock-class-to-quality-level" on page 4656	"convert-sa-to-msdp" on page 4657	"convey-loss-threshold" on page 4658
"copy-dont-fragment-bit (Services IPsec VPN)" on page 4661	"copy-dont-fragment-bit (Services Set)" on page 4662	"copy-plp-all" on page 4664
"copy-tos-to-outer" on page 4665	"copy-tos-to-outer-ip-header" on page 4666	"copy-tos-to-outer-ip-header" on page 4668

"copy-tos-to-outer-ip-header-transit" on page 4669	"core-dump" on page 4670	"core-facing" on page 4672
"corouted-bidirectional" on page 4673	"corouted-bidirectional-passive" on page 4674	"corrupt-file" on page 4676
"cos-rule-sets (Service Set Next Gen Services)" on page 4677	"cos-rules (Service Set)" on page 4679	"cos-rules (Service Set Next Gen Services)" on page 4680
"cost" on page 4681	"count" on page 4685	"count (HTTP Header Enrichment)" on page 4686
"count (Protocols PIM)" on page 4684	"count (Security Custom Attack)" on page 4688	"count (Security Policies)" on page 4689
"counter" on page 4690	"counters" on page 4691	"cpu" on page 4693
"cpu (resource-manager)" on page 4697	"cpu (TDF Gateway)" on page 4699	"cpu-factor (Dampen Options)" on page 4694
"cpu-load-threshold" on page 4696	"cpu-throttle (Next Gen Services)" on page 4701	"craft-lockout" on page 4703
"crb-proxy-mac" on page 4704	"create-new-ucast-tunnel" on page 4707	"create-req" on page 4708
"credibility" on page 4710	"credit (Interfaces)" on page 4712	"critical-event (Process)" on page 4713
"critical-event (System)" on page 4715	"crl" on page 4717	"crl (Adaptive Services Interface)" on page 4718
"crl (Encryption Interface)" on page 4720	"crl (Security)" on page 4721	"cross-connect" on page 4724
"cross-connect-ccm" on page 4726	"cryptographic-address" on page 4728	"csds" on page 4729

"cspf-backoff-time" on page 4731	"cts" on page 4732	"cts-polarity" on page 4734
"current-hop-limit" on page 4735	"current-hop-limit (Dynamic Router Advertisement)" on page 4737	"custom-attack" on page 4738
"custom-attack-group" on page 4746	"custom-attack-groups (Security IDP)" on page 4747	"custom-attacks" on page 4748
"custom-block-message" on page 4750	"custom-ciphers" on page 4751	"custom-message (Security Content Filtering)" on page 4758
"custom-message (Security Email Notify)" on page 4759	"custom-message (Security Fallback Block)" on page 4761	"custom-message (Security Fallback Non-Block)" on page 4762
"custom-message (Security Virus Detection)" on page 4763	"custom-message (Security Web Filtering)" on page 4764	"custom-message-subject (Security Email Notify)" on page 4766
"custom-message-subject (Security Fallback Block)" on page 4767	"custom-message-subject (Security Fallback Non-Block)" on page 4769	"custom-message-subject (Security Virus Detection)" on page 4770
"custom-objects" on page 4771	"custom-options" on page 4773	"custom-page" on page 4777
"custom-page-file" on page 4778	"custom-tag-string" on page 4780	"custom-url-category" on page 4781
"customer-vlans" on page 4783	"customization (SNMP)" on page 4784	"cut-through" on page 4786
"cvss-score" on page 4787	"cycle-time" on page 4788	

D

"dad-proxy" on page 4809	"dampen" on page 4811	"dampen (Protocols MVPN)" on page 4813
"dampen-options" on page 4814	"damping (BFD Liveness Detection)" on page 4815	"damping (Interfaces)" on page 4818
"damping (Policy Options)" on page 4820	"damping (Protocols BGP)" on page 4822	"data (FTP)" on page 4837
"data-channel" on page 4825	"data-cores" on page 4827	"data-encapsulation" on page 4828
"data-fill" on page 4830	"data-fill-with zeros" on page 4831	"data-flow-affinity" on page 4832
"data-format" on page 4834	"data-forwarding" on page 4835	"data-length" on page 4838
"data-mdt-reuse" on page 4840	"data-rate" on page 4841	"data-record-fields" on page 4842
"data-size" on page 4847	"data-tlv-size" on page 4849	"database" on page 4850
"database-protection" on page 4852	"database-replication (Subscriber Session Database)" on page 4854	"datapath-debug" on page 4855
"dcbx" on page 4858	"dcbx-version" on page 4860	"dcd" on page 4861
"dcd-polarity" on page 4863	"dce-options" on page 4864	"ddos-protection" on page 4865
"ddos-protection (DDoS)" on page 5137	"deactivation-attribute (AAA Profile)" on page 5143	"deactivation-delay" on page 5144
"dead-criteria-retries (RADIUS Server)" on page 5146	"dead-peer-detection" on page 5147	"dead-peer-detection (Security Group VPN Server)" on page 5149

"dead-peer-detection (Services IPsec VPN)" on page 5151	"deaggregate" on page 5153	"debounce-interval" on page 5154
"debug-level (System Services)" on page 5155	"debug-log (System Services)" on page 5157	"decapsulate (Firewall Filter)" on page 5160
"decapsulate-accept-inner-vlan" on page 5158	"decapsulate-inner-vlan" on page 5163	"decompress-layer" on page 5165
"decompress-layer-limit" on page 5166	"decryption-failures" on page 5168	"dedicated-buffer" on page 5170
"dedicated-buffer-profile" on page 5172	"dedicated-ukern-cpu (BFD)" on page 5174	"deep-vlan-qualified-learning" on page 5176
"default (CoS)" on page 5187	"default (CoS Host Outbound Traffic)" on page 5189	"default (Security Antivirus)" on page 5213
"default (Security Antivirus Sophos Engine)" on page 5214	"default (Security Content Security)" on page 5216	"default (Security Web Filtering)" on page 5217
"default-action (DHCP Relay Agent Option)" on page 5177	"default-action (Protocols OpenFlow)" on page 5179	"default-actions" on page 5180
"default-address-selection" on page 5182	"default-chap-secret" on page 5184	"default-client-identity (NETCONF TLS)" on page 5185
"default-gateway" on page 5190	"default-isid" on page 5191	"default-lifetime" on page 5193
"default-lifetime (Dynamic Router Advertisement)" on page 5194	"default-local-policy" on page 5195	"default-lsa" on page 5197
"default-pap-password" on page 5199	"default-peer" on page 5200	"default-policy" on page 5202
"default-policy (IDP)" on page 5203	"default-pool (Address Pools)" on page 5205	"default-prefixes (Protocols RIFT)" on page 5206

"default-profile" on page 5208	"default-profile (Juniper Secure Connect)" on page 5209	"default-rule" on page 5211
"default-trusted-ca-certs (Security)" on page 5218	"default-value" on page 5221	"default-vpn-source" on page 5222
"defaults" on page 5224	"defaults (ACX series, JRR200, MX204, PTX5000, and VRR)" on page 5225	"defaults (early-aging)" on page 5226
"defaults (Policy Options)" on page 5228	"defer-initial-multipath-build" on page 5229	"definition" on page 5231
"definition (Protocols OSPF)" on page 5234	"degraded" on page 5236	"delay" on page 5238
"delay (PPPoE Service Name Tables)" on page 5257	"delay (querier)" on page 5259	"delay (responder)" on page 5260
"delay-advertise (DHCPv6)" on page 5239	"delay-authentication (DHCP Relay Agent)" on page 5242	"delay-buffer-rate" on page 5244
"delay-buffer-rate (Dynamic Traffic Shaping)" on page 5246	"delay-factor" on page 5248	"delay-measurement (Protocols IS-IS)" on page 5249
"delay-measurement (Protocols OSPF)" on page 5251	"delay-metric" on page 5252	"delay-metric (Protocols OSPF)" on page 5253
"delay-offer (DHCPv4)" on page 5254	"delay-route-advertisements" on page 5262	"delay-variation" on page 5265
"delegate-probes" on page 5266	"delegate-processing (VRRP)" on page 5268	"delegate-server-processing" on page 5270
"delegated-pool (DHCP Local Server)" on page 5271	"delegation-cleanup-timeout" on page 5272	"delegation-priority" on page 5274

"delete-after-commit (JNU Satellites)" on page 5278	"delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)" on page 5279	"delete-req" on page 5282
"delimiter" on page 5285	"delimiter (Access Profile)" on page 5286	"delimiter (DHCP Local Server)" on page 5288
"delimiter (DHCP Relay Agent)" on page 5290	"delimiter (Domain Map)" on page 5293	"demand-circuit (Protocols RIP)" on page 5294
"demux (Interfaces)" on page 5299	"demux0 (Dynamic Interface)" on page 5311	"demux0 (Static Interface)" on page 5313
"demux-destination (Demux Interface)" on page 5296	"demux-destination (Underlying Interface)" on page 5297	"demux-options (All Demux Interfaces)" on page 5301
"demux-options (Dynamic Interface)" on page 5302	"demux-options (Static Interface)" on page 5304	"demux-source (Demux Interface)" on page 5305
"demux-source (Dynamic IP Demux Interface)" on page 5306	"demux-source (Dynamic Underlying Interface)" on page 5308	"demux-source (Underlying Interface)" on page 5309
"dense-groups" on page 5315	"deny (Security Policies)" on page 5316	"deny (Security SIP)" on page 5317
"depth-threshold" on page 5319	"description (Abstracted Fabric)" on page 5324	"description (Application Identification)" on page 5325
"description (Applications)" on page 5326	"description (Chassis Redundancy)" on page 5328	"description (Fibre Channel Fabrics)" on page 5329
"description (Forwarding Options)" on page 5330	"description (GNF)" on page 5332	"description (IKE policy)" on page 5333
"description (Interfaces)" on page 5334	"description (Junos Fusion)" on page 5336	"description (Junos Fusion Satellite Device Cluster)" on page 5337

"description (Op Scripts)" on page 5339	"description (Protocols BGP)" on page 5340	"description (Protocols Layer 2 Circuit Neighbor)" on page 5342
"description (Protocols Layer 2 VPN)" on page 5343	"description (Protocols MPLS)" on page 5345	"description (RMON)" on page 5321
"description (Routing Instances)" on page 5346	"description (Security Address Book)" on page 5347	"description (Security IDP Policy)" on page 5349
"description (Security NAT Pool)" on page 5350	"description (Security NAT Rule)" on page 5351	"description (Security NAT Rule Set)" on page 5353
"description (Security Policies Next Gen Services)" on page 5354	"description (Security Screen)" on page 5355	"description (Security Zone)" on page 5357
"description (Services IPsec VPN)" on page 5358	"description (SNMP)" on page 5322	"description (VLAN)" on page 5360
"deselect-on-bandwidth-failure" on page 5361	"designated-forwarder-election-hold-time" on page 5363	"destination" on page 5364
"destination (Accounting)" on page 5366	"destination (Application Identification)" on page 5408	"destination (Chassis)" on page 5410
"destination (Diameter Network Element)" on page 5414	"destination (Event Policy)" on page 5416	"destination (Interfaces)" on page 5433
"destination (IPCP)" on page 5439	"destination (L2TP)" on page 5445	"destination (LRF Profile)" on page 5446
"destination (Protocols OSPF)" on page 5511	"destination (Routing Instance)" on page 5518	"destination (saegw system)" on page 5520
"destination (Security Destination NAT)" on page 5521	"destination (Security IP Headers Attack)" on page 5524	"destination (Tunnel Remote End)" on page 5526

"destination (Tunnels)" on page 5528	"destination-address" on page 5368	"destination-address (Firewall)" on page 5371
"destination-address (Flow Monitoring Logs for NAT)" on page 5373	"destination-address (HTTP Header Enrichment) " on page 5375	"destination-address (IDS MS-DPC)" on page 5377
"destination-address (IPsec-VPN)" on page 5369	"destination-address (NAT Next Gen Services)" on page 5383	"destination-address (PCP)" on page 5384
"destination-address (Security Destination NAT)" on page 5394	"destination-address (Security IDP Policy)" on page 5396	"destination-address (Security Policies)" on page 5397
"destination-address (Security Policies Flag)" on page 5399	"destination-address (Security Source NAT)" on page 5400	"destination-address (Security Static NAT)" on page 5401
"destination-address (Services CoS)" on page 5403	"destination-address (Services NAT)" on page 5404	"destination-address (Services Stateful Firewall)" on page 5405
"destination-address (Subscriber Secure Policy)" on page 5407	"destination-address-excluded" on page 5372	"destination-address-name (NAT Next Gen Services)" on page 5378
"destination-address-name (Security Destination NAT)" on page 5379	"destination-address-name (Security Source NAT)" on page 5380	"destination-address-name (Security Static NAT)" on page 5382
"destination-address-range" on page 5386	"destination-address-range (HTTP Header Enrichment)" on page 5387	"destination-address-range (IDS MS-DPC))" on page 5389
"destination-address-range (PCP)" on page 5390	"destination-address-range (Services NAT)" on page 5391	"destination-address-range (Services Stateful Firewall)" on page 5393
"destination-class-usage" on page 5411	"destination-classes" on page 5412	"destination-equal-load-balancing (L2TP LAC)" on page 5415
"destination-except" on page 5418	"destination-header" on page 5419	"destination-host" on page 5421

"destination-host (Gx-Plus)" on page 5422	"destination-host (OCS Partition)" on page 5423	"destination-host (PCRF Partition)" on page 5424
"destination-hostname" on page 5426	"destination-identity-context" on page 5427	"destination-identity-context-profile" on page 5429
"destination-interface" on page 5430	"destination-ip (Security Alarms)" on page 5437	"destination-ip-address (RADIUS Snoop Segment)" on page 5434
"destination-ip-based" on page 5435	"destination-ipv4-address" on page 5440	"destination-ipv4-address (RFC 2544 Benchmarking Tests)" on page 5442
"destination-ipv6-address" on page 5443	"destination-mac-address" on page 5448	"destination-nat" on page 5449
"destination-networks" on page 5451	"destination-option" on page 5452	"destination-override" on page 5454
"destination-path-group" on page 5455	"destination-pool" on page 5457	"destination-port (applications)" on page 5464
"destination-port (Applications)" on page 5465	"destination-port (firewall)" on page 5458	"destination-port (Flow Monitoring Logs for NAT)" on page 5471
"destination-port (PCP)" on page 5472	"destination-port (protocols pcep)" on page 5460	"destination-port (RADIUS Snoop Segment)" on page 5474
"destination-port range" on page 5475	"destination-port (RPM and TWAMP)" on page 5461	"destination-port (Security Destination NAT)" on page 5478
"destination-port (Security Forwarding Options)" on page 5479	"destination-port (Security Signature Attack)" on page 5485	"destination-port (Security Source NAT)" on page 5486
"destination-port (Security Static NAT)" on page 5487	"destination-port (SNMP)" on page 5488	"destination-port (Subscriber Secure Policy)" on page 5490

"destination-port-range (HTTP Header Enrichment)" on page 5476	"destination-ports (HTTP Header Enrichment)" on page 5491	"destination-prefix (Destination NAT Next Gen Services)" on page 5493
"destination-prefix (IDS)" on page 5494	"destination-prefix (Security Forwarding Options)" on page 5507	"destination-prefix (Services NAT) " on page 5508
"destination-prefix-ipv6 (IDS)" on page 5495	"destination-prefix-list (HTTP Header Enrichment)" on page 5497	"destination-prefix-list (PCP)" on page 5498
"destination-prefix-list (Services AACL)" on page 5500	"destination-prefix-list (Services CoS)" on page 5501	"destination-prefix-list (Services IDS)" on page 5502
"destination-prefix-list (Services NAT) " on page 5504	"destination-prefix-list (Services Stateful Firewall) " on page 5505	"destination-profile" on page 5509
"destination-realm (Gx-Plus)" on page 5512	"destination-realm (JSRC)" on page 5514	"destination-realm (OCS Partition)" on page 5515
"destination-realm (PCRF Partition)" on page 5517	"destination-threshold" on page 5525	"destination-udp-port" on page 5529
"destination-udp-port (FTI)" on page 5531	"destinations" on page 5532	"destinations (Event Policy)" on page 5533
"destined-port" on page 5535	"destruct-timeout (L2TP)" on page 5536	"detect-loc" on page 5538
"detection (Network Isolation Groups)" on page 5540	"detection-time" on page 5542	"detection-time (BFD for PIM)" on page 5543
"detection-time (BFD Liveness Detection)" on page 5545	"detection-time (LAG)" on page 5549	"detection-time (Liveness Detection)" on page 5550
"detector" on page 5551	"deterministic (Source NAT Next Gen Services)" on page 5559	"deterministic-ar-policy" on page 5553

"deterministic-nat-configuration-log-interval (Source NAT Next Gen Services)" on page 5555	"deterministic-port-block-allocation" on page 5557	"device-attribute (Easy EVPN LAG Configuration)" on page 5560
"device-authentication" on page 5563	"device-count" on page 5565	"device-count (MX Series)" on page 5567
"device-count (Pseudowire Subscriber Interfaces)" on page 5569	"devices" on page 5571	"df-election" on page 5572
"df-election-granularity" on page 5574	"df-election-type" on page 5576	"dh-group" on page 5579
"dh-group (Security IKE)" on page 5581	"dh-group (services)" on page 5578	"dhcp" on page 5583
"dhcp (DHCP Client)" on page 5598	"dhcp-attributes (Access IPv4 Address Pools)" on page 5585	"dhcp-attributes (Access IPv6 Address Pools)" on page 5587
"dhcp-attributes (Address-Assignment Pools)" on page 5589	"dhcp-client" on page 5596	"dhcp-local-server" on page 5601
"dhcp-local-server (System Services)" on page 5613	"dhcp-option82" on page 5619	"dhcp-relay" on page 5620
"dhcp-relay (DHCP Spoofing Prevention)" on page 5637	"dhcp-relay (Easy EVPN LAG Configuration)" on page 5638	"dhcp-security" on page 5640
"dhcp-security (MX Series)" on page 5644	"dhcp-service" on page 5646	"dhcp-snooping-file" on page 5650
"dhcp-snooping-file (EX)" on page 5649	"dhcp-tags (Adjustment Control Profiles)" on page 5652	"dhcp-trusted" on page 5654
"dhcpv6 (DHCP Local Server)" on page 5658	"dhcpv6 (DHCP Relay Agent)" on page 5665	"dhcpv6 (System Services)" on page 5677

"dhcpv6-client" on page 5655	"dhcpv6-options" on page 5673	"dhcpv6-snooping-file" on page 5675
"diag-port-authentication" on page 5682	"dial-options" on page 5684	"dial-options (Dynamic Profiles)" on page 5685
"dial-string" on page 5687	"dialer" on page 5688	"dialer-options" on page 5689
"dialin" on page 5691	"diameter" on page 5693	"diameter (Subscriber Aware Policy Control)" on page 5699
"diameter (TDF Gateway)" on page 5702	"diameter-instance (Diameter Applications)" on page 5695	"diameter-instance (JSRC)" on page 5696
"diameter-profile (PCEF Profile)" on page 5698	"dictionary" on page 5703	"diffserv-te" on page 5705
"direct-access (Commit Scripts)" on page 5706	"direct-connect" on page 5708	"direction (Application Identification)" on page 5715
"direction (forwarding-class-accounting)" on page 5716	"direction (Junos OS)" on page 5719	"direction (Junos-FIPS Software)" on page 5718
"direction (MEP)" on page 5721	"direction (MX Series)" on page 5722	"direction (RFC 2544 Benchmarking Tests)" on page 5709
"direction (Secure-channel)" on page 5713	"direction (Security Custom Attack)" on page 5724	"direction (Security Dynamic Attack Group)" on page 5726
"direction (Service Data Flow Filters)" on page 5727	"direction (services)" on page 5711	"director-device (Aliases)" on page 5729
"disable" on page 5736	"disable (Application Tracking)" on page 5745	"disable (BGP Graceful Restart)" on page 5747
"disable (DCBX)" on page 5749	"disable (Dynamic IGMP)" on page 5751	"disable (Dynamic MLD)" on page 5752

"disable (Extensible Subscriber Services Manager)" on page 5753	"disable (Forwarding Options)" on page 5756	"disable (GMPLS)" on page 5761
"disable (IGMP Snooping)" on page 5764	"disable (Interface)" on page 5765	"disable (Link Protection)" on page 5767
"disable (Long-Lived Graceful Restart for BGP Restarter)" on page 5770	"disable (Multicast Load Balancing)" on page 5773	"disable (Multicast Snooping)" on page 5775
"disable (MVRP)" on page 5776	"disable (PIM)" on page 5782	"disable (PoE)" on page 5784
"disable (Power over Ethernet)" on page 5785	"disable (Protocols)" on page 5731	"disable (Protocols BGP)" on page 5788
"disable (Protocols LDP)" on page 5789	"disable (Protocols MLD)" on page 5791	"disable (Protocols MLD Snooping)" on page 5792
"disable (Protocols MPLS)" on page 5793	"disable (Protocols MSDP)" on page 5795	"disable (Protocols Router Discovery)" on page 5797
"disable (Protocols RSVP)" on page 5798	"disable (Protocols SAP)" on page 5800	"disable (Routing Graceful Restart)" on page 5738
"disable (Routing Options)" on page 5804	"disable (STP)" on page 5739	"disable-4byte-as" on page 5742
"disable-action" on page 5743	"disable-all-instances" on page 5744	"disable-calling-number-avp (L2TP LAC)" on page 5748
"disable-failover-protocol (L2TP)" on page 5755	"disable-fpc (DDoS)" on page 5758	"disable-global-timeout-override" on page 5760
"disable-grant-bypass" on page 5762	"disable-https-filtering" on page 5812	"disable-logging (DDoS)" on page 5769

"disable-mlppp-inner-ppp-pfc" on page 5772	"disable-natt (Services IPsec VPN)" on page 5777	"disable-notification-extensions (BGP Graceful Restart)" on page 5779
"disable-notification-flag (BGP Graceful Restart)" on page 5780	"disable-packet-register" on page 5813	"disable-preceding-key" on page 5787
"disable-relay" on page 5801	"disable-routing-engine (DDoS)" on page 5802	"disable-signature-check" on page 5805
"disable-timeout" on page 5806	"disable-timeout (BPDU)" on page 5808	"disable-timeout (Port Error Disable)" on page 5810
"discard (Forwarding Class)" on page 5819	"discard (forwarding-options)" on page 5814	"discard (routing-options)" on page 5816
"discard-action-for-unresolved-redir-addr" on page 5818	"disconnect-peer-timeout" on page 5820	"disk-failure-action" on page 5821
"disk-failure-action" on page 5823	"disk-partition" on page 5824	"display-host (Security Fallback Block)" on page 5826
"display-host (Security Virus Detection)" on page 5827	"distinguished-name" on page 5828	"distinguished-name (Access)" on page 5830
"distinguished-name (Security)" on page 5831	"distributed (IGMP)" on page 5834	"distributed-dr" on page 5833
"distribution-ip (Security SIP)" on page 5836	"distribution-profile" on page 5838	"dlb" on page 5841
"dlci" on page 5843	"dlv" on page 5844	"dne-origin (Diameter Network Element)" on page 5846
"dns" on page 5847	"dns (System Services)" on page 5867	"dns-alg-pool" on page 5849

"dns-alg-prefix" on page 5850	"dns-cache" on page 5852	"dns-filter" on page 5853
"dns-filter-template" on page 5856	"dns-filtering" on page 5859	"dns-proxy" on page 5862
"dns-server-address" on page 5864	"dns-server-address (Dynamic Profiles)" on page 5865	"dnssec" on page 5869
"do-not-fragment" on page 5871	"dod-request-policy" on page 5873	"domain (DNS Forwarding Options)" on page 5874
"domain (Domain Map)" on page 5877	"domain (Precision Time Protocol)" on page 5875	"domain (TDF Domain Selection)" on page 5905
"domain-name" on page 5879	"domain-name (Access Profile)" on page 5881	"domain-name (DHCP)" on page 5882
"domain-name (DHCP Local Server)" on page 5884	"domain-name (DHCP Relay Agent)" on page 5886	"domain-name (Static Subscribers)" on page 5892
"domain-name (System)" on page 5880	"domain-name-server (Routing Instances and Access Profiles)" on page 5891	"domain-name-server-inet (Routing Instances and Access Profiles)" on page 5888
"domain-name-server-inet6 (Routing Instances and Access Profiles)" on page 5890	"domain-path-id" on page 5894	"domain-path-id (Protocols BGP)" on page 5898
"domain-search" on page 5899	"domain-selection" on page 5901	"domain-type (Bridge Domains)" on page 5906
"domains" on page 5908	"dont-help-shared-fate-bfd-down" on page 5912	"dot1p-priority" on page 5913
"dot1q-tunneling " on page 5915	"dot1x" on page 5917	"dot1x (MX Series in Enhanced LAN Mode)" on page 5920

"downlink-dscp-to-egress-forwarding-class" on page 5922	"downlink-dscp-to-forwarding-class" on page 5923	"ignore-server-validation (Application Identification)" on page 5924
"download (Services)" on page 5928	"download-profile (Security Antivirus FTP)" on page 5926	"download-profile (Security Content Filtering FTP)" on page 5927
"download-timeout" on page 5930	"downstream" on page 5932	"downstream-mode" on page 5933
"downstream-on-demand" on page 5934	"downstream-rate (Traffic Shaping)" on page 5935	"dr-election-on-p2p" on page 5937
"dr-register-policy" on page 5938	"drain" on page 5940	"draining (Diameter Applications)" on page 5941
"draining-response-timeout (Diameter Applications)" on page 5943	"drop (BPDU Block)" on page 5944	"drop (DHCP Relay Agent Option)" on page 5946
"drop (PPPoE Service Name Tables)" on page 5952	"drop-profile-map (Schedulers)" on page 5964	"drop (Security GTP)" on page 5969
"drop (Security SCTP)" on page 5974	"drop-flow" on page 5947	"drop-member-traffic (Aggregated Multiservices)" on page 5949
"drop-policy (Subscriber Secure Policy)" on page 5950	"drop-probability" on page 5953	"drop-probability (Fill Level)" on page 5955
"drop-probability (Interpolated Value)" on page 5956	"drop-probability (Percentage)" on page 5958	"drop-profile (Dynamic Schedulers)" on page 5960
"drop-profile (Schedulers)" on page 5959	"drop-profile (Schedulers)" on page 5966	"drop-profile-map (Dynamic Schedulers)" on page 5962
"drop-profile-map (Schedulers)" on page 5965	"drop-profiles (Definition)" on page 5968	"drop-threshold" on page 5978

"drop-threshold (Security GPRS)" on page 5980	"drop-timeout" on page 5981	"drop-timeout (Forwarding Class)" on page 5983
"drop-unknown-mac-ip" on page 5984	"ds-lite" on page 5986	"dscp (Class of Service)" on page 5988
"dscp (CoS Classifiers)" on page 5998	"dscp (CoS Interfaces)" on page 5999	"dscp (Dynamic Classifiers)" on page 6000
"dscp (Dynamic Rewrite Rules)" on page 6002	"dscp (flow-server)" on page 6003	"dscp (Input Congestion Notification)" on page 6005
"dscp-ipv6 (Class of Service)" on page 6007	"dscp (Multifield Classifier)" on page 6017	"dscp (Rewrite Rules on Interfaces)" on page 6020
"dscp (Services CoS)" on page 6022	"dscp (Subscriber Secure Policy)" on page 6023	"dscp-code-point (CoS Host Outbound Traffic)" on page 5991
"dscp-code-points (RFC 2544 Benchmarking)" on page 5993	"dscp-code-points (RPM)" on page 5995	"dscp-ipv6 (Class of Service)" on page 6009
"dscp-ipv6 (Dynamic Classifiers)" on page 6013	"dscp-ipv6 (Dynamic Rewrite Rules)" on page 6014	"dscp-mapping-classifier" on page 6016
"dscp-rewrite" on page 6019	"dsl (Access-Line Rate Adjustment)" on page 6025	"dsl-sfp-options" on page 6030
"dslite-softwire-initiator" on page 6033	"dsr" on page 6035	"dsr-polarity" on page 6037
"dtcp-only (System Services)" on page 6038	"dte-options" on page 6039	"dtr" on page 6041
"dtr-circuit" on page 6042	"dtr-polarity" on page 6044	"dual-dr" on page 6045
"dual-phase-bootup" on page 6047	"dual-stack (DHCP Local Server Overrides)" on page 6049	"dual-stack (DHCP Relay Agent Overrides)" on page 6050

"dual-stack-group (DHCP Local Server)" on page 6052	"dual-stack-group (DHCP Relay Agent)" on page 6055	"dual-stack-interface-client-limit (DHCP Local Server and Relay Agent)" on page 6058
"dual-transport" on page 6059	"dualstack-support (JSRC)" on page 6061	"dump-on-flow-control" on page 6063
"duplicate-clients (DHCPv6 Local Server and Relay Agent)" on page 6064	"duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)" on page 6066	"duplicate-mac-detection" on page 6068
"duplicate-protection (Dynamic PPPoE)" on page 6071	"duplicates-dropped-periodicity" on page 6072	"duplication (Access Profile)" on page 6074
"duplication-filter (Access Profile)" on page 6075	"duplication-vrf (Duplicate Accounting)" on page 6077	"duration" on page 6078
"duration (PoE)" on page 6080	"dvmrp" on page 6081	"dynamic" on page 6084
"dynamic (Security)" on page 6151	"dynamic (services)" on page 6083	"dynamic (Source Packet Routing)" on page 6156
"dynamic-address" on page 6086	"dynamic-address" on page 6088	"dynamic-application" on page 6090
"dynamic-application (Security)" on page 6093	"dynamic-application (Security Policies)" on page 6095	"dynamic-application-group" on page 6092
"dynamic-attack-group" on page 6097	"dynamic-attack-groups (Security IDP)" on page 6100	"dynamic-bidirectional-transport" on page 6101
"dynamic-class-of-service-options (Dynamic Traffic Shaping)" on page 6102	"dynamic-db" on page 6105	"dynamic-dns" on page 6106
"dynamic-filter" on page 6108	"dynamic-flow-capture" on page 6109	"dynamic-peers" on page 6111

"dynamic-policy-control" on page 6112	"dynamic-pool" on page 6114	"dynamic-profile (Demux)" on page 6115
"dynamic-profile (DHCP Local Server)" on page 6117	"dynamic-profile (DHCP Relay Agent)" on page 6119	"dynamic-profile (Dynamic Access-Line-Identifier VLANs)" on page 6121
"dynamic-profile (Dynamic ACI VLANs)" on page 6122	"dynamic-profile (Dynamic PPPoE)" on page 6124	"dynamic-profile (L2TP)" on page 6126
"dynamic-profile (PPP)" on page 6128	"dynamic-profile (PPPoE Service Name Tables)" on page 6129	"dynamic-profile (Stacked VLAN)" on page 6131
"dynamic-profile (VLAN)" on page 6133	"dynamic-profile-options" on page 6127	"dynamic-profiles" on page 6134
"dynamic-requests" on page 6148	"dynamic-requests-secret (RADIUS Server)" on page 6149	"dynamic-server" on page 6154
"dynamic-tunnel-gre-key" on page 6158	"dynamic-tunnel-reassembly" on page 6159	"dynamic-tunnels" on page 6161
"dynamic-vpn" on page 6163		

E

"e1-options (BITS Interfaces Signal Type)" on page 6176	"e1-options (Clock Synchronization)" on page 6178	"e2e-transparent" on page 6179
"eapol-address (MACSec)" on page 6181	"eapol-block" on page 6183	"early-ageout" on page 6185
"east-interface" on page 6186	"echo-req" on page 6188	"ecmp" on page 6190

"ecmp-alb" on page 6191	"ecmp-fast-reroute" on page 6193	"ecmp-local-bias" on page 6195
"ecmp-resilient-hash" on page 6196	"edge" on page 6197	"edge-virtual-bridging" on page 6200
"effective-aigp" on page 6203	"effective-shaping-rate" on page 6204	"egress" on page 6205
"egress (Analyzer)" on page 6207	"egress (Buffer Configuration)" on page 6208	"egress-chaining" on page 6210
"egress-l2-extended-match" on page 6212	"egress-policer-overhead" on page 6213	"egress-policy" on page 6215
"egress-protection (BGP)" on page 6216	"egress-protection (Layer 2 circuit)" on page 6218	"egress-protection (MPLS)" on page 6220
"egress-protection (Routing Instances)" on page 6221	"egress-quantization" on page 6238	"egress-shaping-overhead" on page 6223
"egress-stats (Flat-File Accounting Options)" on page 6225	"egress-te" on page 6227	"egress-te-adj-segment" on page 6229
"egress-te-backup-paths" on page 6232	"egress-te-node-segment" on page 6234	"egress-te-set-segment" on page 6236
"ei-mapping-timeout" on page 6241	"ei-mapping-timeout (Source NAT Next Gen Services)" on page 6242	"eif-flow-limit" on page 6244
"em-hw-profile" on page 6245	"email-notify" on page 6247	"embedded-rp" on page 6248
"enable" on page 6249	"enable" on page 6251	"enable (Enhanced Subscriber Management)" on page 6257
"enable-all-ifl" on page 6252	"enable-all-qmodules" on page 6253	"enable-asymmetric-traffic-processing (Service Set Next Gen Services)" on page 6254

"enable-call-distribution" on page 6255	"enable-descriptive-session-syslog" on page 6256	"enable-explorer (REST API)" on page 6259
"enable-extended-ql-tlv" on page 6260	"enable-flow-tracing (Services)" on page 6263	"enable-gtpu-distribution" on page 6264
"enable-ipv6-services-for-lac (L2TP)" on page 6265	"enable-mac-move-action" on page 6267	"enable-packet-pool" on page 6268
"enable-performance-mode" on page 6269	"enable-performance-mode (Application Identification)" on page 6271	"enable-rejoin (Aggregated Multiservices)" on page 6272
"enable-reroute-uniform-link-check" on page 6273	"enable-reverse-reroute" on page 6275	"enable-session-cache" on page 6276
"enable-snmp-tunnel-statistics (L2TP)" on page 6277	"enable-subscriber-analysis (Services Options VMS Interfaces)" on page 6279	"encap-type" on page 6280
"encapsulate-inner-vlan" on page 6282	"encapsulation (Dynamic Interfaces)" on page 6297	"encapsulation (interfaces)" on page 6289
"encapsulation (interfaces-fti)" on page 6288	"encapsulation (Logical Interface)" on page 6300	"encapsulation (logical-link-layer)" on page 6286
"encapsulation (Physical Interface)" on page 6307	"encapsulation (protocols evpn)" on page 6284	"encapsulation-overhead" on page 6305
"encapsulation-type (Layer 2 Circuits)" on page 6308	"encapsulation-type (Layer 2 VPNs)" on page 6310	"encapsulation-udp-fti" on page 6313
"encoding" on page 6317	"encoding (QFX Series)" on page 6314	"encoding (security)" on page 6315
"encoding-type" on page 6318	"encrypt (HTTP Header Enrichment)" on page 6320	"encryption" on page 6322

"encryption (Junos OS)" on page 6332	"encryption (Junos-FIPS Software)" on page 6331	"encryption (MACsec)" on page 6334
"encryption (MACsec for MX Series)" on page 6336	"encryption-algorithm" on page 6324	"encryption-algorithm (Security)" on page 6326
"encryption-algorithm (Security IKE)" on page 6327	"encryption-failures" on page 6330	"end-interface" on page 6338
"end-user-profile " on page 6339	"endpoint-registration-timeout" on page 6342	"enforce-first-as" on page 6343
"enforce-strict-scale-limit-license (Subscriber Management)" on page 6345	"engine-id" on page 6346	"engine-id (Forwarding Options)" on page 6348
"engine-not-ready" on page 6350	"engine-not-ready (Security Antivirus Sophos Engine)" on page 6351	"engine-type" on page 6352
"enhanced" on page 6354	"enhanced (forwarding-class-accounting)" on page 6358	"enhanced-cfm-mode" on page 6355
"enhanced-convergence" on page 6357	"enhanced-hash-key" on page 6362	"enhanced-hash-key (EX)" on page 6378
"enhanced-hash-key profile1/ profile2" on page 6389	"enhanced-hierarchical-policer" on page 6394	"enhanced-mode" on page 6396
"enhanced-mode (Network Services)" on page 6398	"enhanced-mode-override" on page 6401	"enhanced-policer" on page 6403
"enhanced-priority-mode " on page 6404	"enhanced-transmission-selection" on page 6406	"enrollment" on page 6408
"enrollment-retry" on page 6409	"enrollment-url" on page 6411	"enterprise-oid" on page 6412

"entropy-label (protocols bgp)" on page 6413	"entropy-label (protocols mpls)" on page 6416	"environment-monitoring-policy (satellite-management)" on page 6417
"environment-monitoring-policy (satellite-policies)" on page 6419	"epacl-firewall-optimization" on page 6421	"ephemeral" on page 6423
"equals" on page 6426	"equals (DHCP Relay Agent)" on page 6428	"equals (Dynamic Profile)" on page 6433
"equals (Event Policy)" on page 6434	"eracl-ip6-match (packet-forwarding-options)" on page 6435	"eracl-profile (packet-forwarding-options)" on page 6437
"erroneous-ccm" on page 6439	"error " on page 6441	"error (Security)" on page 6446
"error-id" on page 6449	"es-options" on page 6451	"esi" on page 6453
"esi-resolution-per-bridge-domain" on page 6456	"esmc-transmit (Chassis Synchronization)" on page 6458	"establish-tunnels" on page 6460
"ether-options" on page 6461	"ether-pseudowire" on page 6471	"ether-type" on page 6476
"ether-type (Chassis FPC)" on page 6477	"ether-type (Dynamic Load Balancing)" on page 6473	"ether-type (VLAN)" on page 6474
"ethernet" on page 6480	"ethernet (Aggregated Devices)" on page 6481	"ethernet (Alarm)" on page 6482
"ethernet (Chassis)" on page 6484	"ethernet (CoS for Multidestination Traffic)" on page 6485	"ethernet (Protocols OAM)" on page 6490
"ethernet-frame-size" on page 6486	"ethernet-policer-profile" on page 6488	"ethernet-ring" on page 6497
"ethernet-switch-profile" on page 6499	"ethernet-switching" on page 6501	"ethernet-switching" on page 6503

"ethernet-switching (Port Mirroring)" on page 6513	"ethernet-switching-options" on page 6506	"ethernet-vlan (Protocols Link Management)" on page 6514
"eui-64" on page 6515	"evc-protocol cfm" on page 6517	"evcs" on page 6518
"event" on page 6520	"event (CFM)" on page 6521	"event (Enhanced Subscriber Management)" on page 6523
"event (LFM)" on page 6525	"event-options" on page 6527	"event-rate" on page 6531
"event-rate (Next Gen Services Service-Set Local System Logging)" on page 6532	"event-script (Event Options)" on page 6533	"event-script (Event Policy)" on page 6536
"event-thresholds" on page 6537	"events (Event Policy Correlated Events)" on page 6538	"events (Event Policy Trigger Events)" on page 6540
"evpn" on page 6541	"evpn (Easy EVPN LAG Configuration)" on page 6548	"evpn-aliasing-optimize" on page 6547
"evpn-ssm-reports-only" on page 6553	"evpn-vxlan" on page 6555	"evpn-vxlan (Easy EVPN LAG Configuration)" on page 6559
"examine-dhcp" on page 6561	"examine-dhcpv6" on page 6563	"examine-fip" on page 6565
"examine-vn2vf" on page 6566	"examine-vn2vn" on page 6568	"exceed-action" on page 6570
"exception " on page 6571	"exception (Security Content Filtering)" on page 6573	"exception-reporting" on page 6572
"exceptions" on page 6575	"excess-bandwidth-share" on page 6577	"excess-burst-size" on page 6580
"excess-burst-size (QFX)" on page 6579	"excess-priority" on page 6582	"excess-priority (Dynamic Schedulers)" on page 6584

"excess-rate (Dynamic Schedulers)" on page 6588	"excess-rate (Dynamic Traffic Shaping)" on page 6589	"excess-rate (Schedulers)" on page 6585
"excess-rate (Schedulers)" on page 6601	"excess-rate-high" on page 6591	"excess-rate-high (Dynamic Traffic Shaping)" on page 6593
"excess-rate-low" on page 6594	"excess-rate-low (Dynamic Traffic Shaping)" on page 6596	"excess-rate-medium-high" on page 6597
"excess-rate-medium-low" on page 6599	"exclude (Diameter Gx Profiles)" on page 6602	"exclude (Dynamic MLD Interface)" on page 6604
"exclude (for Administrative Groups)" on page 6605	"exclude (for Fast Reroute)" on page 6606	"exclude (JSRC Attributes)" on page 6609
"exclude-protocol" on page 6610	"exclude (Protocols IGMP)" on page 6614	"exclude (Protocols MLD)" on page 6615
"exclude (RADIUS Attributes)" on page 6617	"exclude (Schedulers)" on page 6626	"exclude (Security Log)" on page 6627
"exclude-hostname" on page 6608	"exclude-protocol (MX Series)" on page 6612	"exclude-queue-overhead-bytes" on page 6616
"exclude-srlg" on page 6629	"excluded-address (Address-Assignment Pools)" on page 6631	"excluded-range (Address-Assignment Pools)" on page 6632
"exclusive-mac" on page 6634	"execute-commands (Event Policy)" on page 6636	"exercise" on page 6637
"exp (Class of Service)" on page 6639	"exp (Rewrite rules)" on page 6643	"exp-push-push-push" on page 6642
"exp-swap-push-push" on page 6646	"expand-loose-hop" on page 6647	"explicit-congestion-notification" on page 6649

"explicit-null (Protocols BGP)" on page 6652	"explicit-null (Protocols LDP)" on page 6654	"explicit-null (Protocols MPLS)" on page 6655
"explicit-null-cos" on page 6650	"explicit-priority" on page 6657	"export (Bootstrap)" on page 6663
"export (ISIS)" on page 6658	"export (MPLS Traffic engineering database)" on page 6668	"export (OSPF)" on page 6660
"export (Protocols BGP)" on page 6677	"export (Protocols DVMRP)" on page 6678	"export (Protocols LDP)" on page 6680
"export (Protocols MSDP)" on page 6681	"export (Protocols PIM)" on page 6683	"export (Protocols RIFT)" on page 6685
"export (Protocols RIPng)" on page 6686	"export (RIP)" on page 6661	"export (Routing Options)" on page 6691
"export-format (Flow Monitoring)" on page 6667	"export-format (JSON)" on page 6665	"export-profile (Junos Telemetry Interface)" on page 6669
"export-profiles" on page 6674	"export-rib" on page 6688	"export-target" on page 6693
"expression" on page 6694	"extend-secondary-vlan-id" on page 6696	"extended-echo" on page 6697
"extended-isid-list" on page 6698	"extended-nexthop" on page 6700	"extended-nexthop-tunnel" on page 6701
"extended-space" on page 6703	"extended-statistics" on page 6704	"extended-system-id" on page 6706
"extended-vlan-list" on page 6707	"extended-vni-all" on page 6709	"extended-vni-list" on page 6710
"Extensible Subscriber Services" on page 6712	"extension-header" on page 6714	"extension-provider" on page 6716

"extension-service (System Extensions)" on page 6717	"extension-service (System Services)" on page 6719	"extensions" on page 6722
"external-assigned (Address Pools)" on page 6725	"external-authority" on page 6726	

F

"f-max-period" on page 6742	"f-max-period (rtp)" on page 6744	"fabric" on page 6745
"fabric (Class-of-Service)" on page 6747	"fabric loopback wan" on page 6755	"fabric (OAM)" on page 6761
"fabric-control" on page 6749	"fabric-id" on page 6750	"fabric-interfaces" on page 6751
"fabric-load-balance" on page 6753	"fabric-maintenance-associations" on page 6756	"fabric-maintenance-end-points" on page 6759
"fabric-options" on page 6763	"fabric-tree-root" on page 6766	"fabric-type" on page 6767
"facility" on page 6769	"facility-override" on page 6770	"facility-override (Security)" on page 6771
"facility-override (Service Sets)" on page 6772	"facility-override (System Log Reporting)" on page 6774	"fail" on page 6775
"fail-filter (Dynamic Profiles)" on page 6777	"failover (Chassis)" on page 6778	"failover (System Process)" on page 6784
"failover-delay" on page 6780	"failover-delay" on page 6781	"failover-resync" on page 6783
"failover-within-preference (L2TP LAC)" on page 6786	"failure-action" on page 6787	"failure-action (Protocols LDP)" on page 6789

"failure-action (Protocols MPLS)" on page 6790	"fallback-block (Security Antivirus)" on page 6792	"fallback-key" on page 6794
"fallback-non-block (Security Antivirus)" on page 6795	"fallback-option (ICAP Redirect Service)" on page 6797	"fallback-options (Security Antivirus Juniper Express Engine)" on page 6799
"fallback-options (Security Antivirus Kaspersky Lab Engine)" on page 6800	"fallback-options (Security Antivirus Sophos Engine)" on page 6802	"fallback-settings (Security Web Filtering)" on page 6804
"fallback-settings (Security Web Filtering Juniper Local)" on page 6805	"fallback-settings (Security Web Filtering Websense Redirect)" on page 6806	"falling-event-index" on page 6808
"falling-event-index (RMON)" on page 6809	"falling-threshold (Health Monitor)" on page 6811	"falling-threshold (integer)" on page 6812
"falling-threshold (percentage)" on page 6815	"falling-threshold (RMON)" on page 6817	"falling-threshold-interval" on page 6814
"false-positives" on page 6818	"family (Address Pools)" on page 6843	"family (Address-Assignment Pools)" on page 6841
"family (Aggregated Multiservices)" on page 6844	"family (Bootstrap)" on page 6846	"family bridge" on page 6847
"family bridge (GRE Interfaces)" on page 6849	"family (chassis)" on page 6851	"family (Dynamic Demux Interface)" on page 6853
"family (Dynamic Firewalls)" on page 6855	"family (Dynamic PPPoE)" on page 6857	"family (Dynamic Standard Interface)" on page 6859
"family (ELS)" on page 6824	"family (Exclude Prefix)" on page 6862	"family fcoe" on page 6863

"family (Filtering)" on page 6865	"family (Firewall)" on page 6867	"family (Firewall Filter)" on page 6869
"family inet" on page 6871	"family inet (Interfaces)" on page 6874	"family inet6" on page 6878
"family (interfaces)" on page 6832	"family (Interfaces)" on page 6881	"family (Link Interfaces)" on page 6839
"family (Local RP)" on page 6883	"family (logical-interface)" on page 6822	"family (Monitoring)" on page 6885
"family mpls" on page 6886	"family (Multi-destination)" on page 6890	"family (Multifield Classifier)" on page 6891
"family multiservice" on page 6893	"family (Port Mirroring)" on page 6896	"family (Protocols AMT Relay)" on page 6898
"family (Protocols BGP)" on page 6900	"family (Protocols IS-IS)" on page 6905	"family (protocols ldp)" on page 6831
"family (Protocols PIM)" on page 6907	"family (Protocols PIM Interface)" on page 6909	"family (QFX)" on page 6829
"family (RFC 2544 Benchmarking Tests)" on page 6820	"family route-target" on page 6911	"family (Sampling)" on page 6913
"family (TDF Interface)" on page 6917	"family (Voice Services)" on page 6918	"family vpls (Layer 2 Pseudowires)" on page 6920
"family (VRF Advertisement)" on page 6921	"family-state-change-immediate-update" on page 6916	"fast-aps-switch" on page 6922
"fast-failover" on page 6924	"fast-interval" on page 6926	"fast-lookup-filter" on page 6928
"fast-lookup-filter (Protocols BGP)" on page 6947	"fast-lookup-filter (PTX)" on page 6930	"fast-lookup-tuple" on page 6935

"fast-lookup-tuple-list" on page 6937	"fast-lookup-tuple-list (policy-options)" on page 6938	"fast-poe" on page 6939
"fast-reroute (Protocols MPLS)" on page 6943	"fast-reroute (Protocols RSVP)" on page 6944	"fast-reroute-priority" on page 6941
"fast-update-filter (Dynamic Firewalls)" on page 6946	"fastether-options" on page 6949	"fatal" on page 6950
"fate-sharing" on page 6952	"fc2" on page 6960	"fc-fabrics" on page 6954
"fc-map" on page 6957	"fc-options" on page 6959	"fcoe" on page 6962
"fcoe-lag" on page 6963	"fcoe-trusted" on page 6965	"feature-profile" on page 6967
"features" on page 6979	"features (VNF)" on page 6978	"feb (Assigning a FEB to a Redundancy Group)" on page 6981
"feb (chassis)" on page 6982	"feb (Creating a Redundancy Group)" on page 6984	"fec" on page 6985
"fec (ether)" on page 6992	"fec (gigether)" on page 6995	"fec (MX Series)" on page 6987
"fec (protocols ldp)" on page 6990	"feed-name" on page 6997	"feed-server" on page 6999
"feeds (T640, T1600, and T4000 Routers with Six-Input DC Power Supply)" on page 7001	"fib-install-priority" on page 7003	"fib-local" on page 7005
"fib-next-hop-split" on page 7006	"fib-prioritization" on page 7007	"fib-remote" on page 7009
"fibre-channel (Alarm)" on page 7010	"fibre-channel (Family Interfaces)" on page 7012	"fibre-channel (Port)" on page 7013
"fibrechannel-options" on page 7015	"fields (Flat-File Accounting Options)" on page 7016	"fields (for Interface Profiles)" on page 7019

"fields (for Routing Engine Profiles)" on page 7021	"file" on page 7023	"file (Analytics)" on page 7028
"file (App Engine Virtual Machine Management Service)" on page 7030	"file (Associating with a Profile)" on page 7036	"file (Commit Scripts)" on page 7042
"file (Configuring a Log File)" on page 7047	"file (Event Scripts)" on page 7054	"file (Extended DHCP Relay Agent and Helpers Trace Options)" on page 7056
"file (Flat-File Accounting Options)" on page 7058	"file (JET)" on page 7059	"file (Next Gen Services Global System Logging)" on page 7065
"file (Op Scripts)" on page 7067	"file (Packet Forwarding Engine Traffic Traceoptions)" on page 7069	"file (QFabric System)" on page 7071
"file (Sampling)" on page 7075	"file (Security Certificates)" on page 7027	"file (Security Log)" on page 7076
"file (SNMP Scripts)" on page 7080	"file (System Logging)" on page 7085	"file (System Logging)" on page 7086
"file (Trace Options)" on page 7090	"file (Tracing for Origin AS Validation)" on page 7091	"file-serialization" on page 7094
"file-specification (File Format)" on page 7082	"file-specification (Interface Mapping)" on page 7083	"file-type" on page 7093
"filename" on page 7096	"filename (Next Gen Services Global System Logging)" on page 7099	"filename (Sampling)" on page 7101
"filename-extension" on page 7097	"filename-prefix" on page 7100	"files" on page 7103
"files (accounting)" on page 7104	"files (Next Gen Services Global System Logging)" on page 7105	"files (Sampling and Traceoptions)" on page 7107

"files (syslog)" on page 7108	"fill-level" on page 7110	"fill-level (Drop Profiles)" on page 7112
"fill-level (Interpolated Value)" on page 7113	"filter" on page 7119	"filter (Applying to a Logical Interface)" on page 7123
"filter (Applying to an Interface)" on page 7125	"filter (Configuring)" on page 7127	"filter (Dynamic Interface Unit)" on page 7130
"filter (Dynamic Profiles Filter Attachment)" on page 7132	"filter (Dynamic Profiles Filter Creation)" on page 7134	"filter (export routing policy action)" on page 7135
"filter (Firewall Filters)" on page 7136	"filter (Flow Monitoring)" on page 7115	"filter (Identity Management Advanced Query)" on page 7138
"filter (interfaces)" on page 7116	"filter (Interfaces)" on page 7117	"filter (IPv4, IPv6, and MPLS)" on page 7144
"filter (Layer 2 and Layer 3 Interfaces)" on page 7145	"filter (Layer 3 Interfaces)" on page 7146	"filter (memory-enhanced)" on page 7150
"filter (routing-options)" on page 7122	"filter (Security)" on page 7153	"filter (VLANs)" on page 7157
"filter (VPLS)" on page 7158	"filter-duplicates" on page 7129	"filter-interfaces" on page 7142
"filter-list-template" on page 7148	"filter-profile" on page 7151	"filter-specific" on page 7155
"filter-specific (QFX)" on page 7154	"filtering-type" on page 7160	"filtering-type (Source NAT Next Gen Services)" on page 7161
"filters" on page 7162	"fin-invalidate-session" on page 7165	"fin-no-ack" on page 7167
"fin-no-ack (IDS Screen Next Gen Services)" on page 7168	"final-response-timeout (OCS Partition)" on page 7170	"finger" on page 7171

"fip" on page 7173	"fip-security" on page 7175	"firewall" on page 7176
"firewall (Dynamic Firewalls)" on page 7191	"firewall-authentication" on page 7179	"firewall-authentication (Security)" on page 7184
"firewall-authentication (Security Policies)" on page 7186	"firewall-authentication (tenants)" on page 7189	"firewall-authentication-forced-timeout" on page 7181
"firewall-authentication-service" on page 7188	"firewall-install-disable" on page 7194	"firewall-user" on page 7196
"firmware-revision" on page 7197	"fiveqi-map" on page 7199	"fixed-stuff-bytes" on page 7201
"fka-adv-period" on page 7202	"flag (App Engine Virtual Machine Management Service)" on page 7204	"flag (Interfaces)" on page 7206
"flag (Next Gen Services Global System Logging)" on page 7208	"flag (Services)" on page 7209	"flag (SMTP)" on page 7211
"flag (Tracing for Origin AS Validation)" on page 7212	"flap-trap-monitor" on page 7215	"flat-file-profile (Accounting Options)" on page 7217
"flat-file-profile (Extensible Subscriber Services)" on page 7220	"flex-algorithm" on page 7221	"flex-algorithm (Protocols OSPF)" on page 7223
"flex-flow-sizing" on page 7225	"flex-hashing" on page 7227	"flexible-cross-connect-vlan-aware" on page 7230
"flexible-cross-connect-vlan-unaware" on page 7232	"flexible-queuing-mode" on page 7233	"flexible-vlan-tagging" on page 7235
"flexible-vlan-tagging (Interfaces)" on page 7237	"flood" on page 7238	"flood (Security ICMP)" on page 7245
"flood (Security UDP)" on page 7246	"flood (VLANs)" on page 7248	"flood-as-unknown-unicast" on page 7239

"flood-groups" on page 7241	"flood-reflector (Protocols IS-IS)" on page 7242	"flow" on page 7252
"flow (IPv6)" on page 7298	"flow (protocols bgp)" on page 7250	"flow (Security Flow)" on page 7333
"flow (Security IDP)" on page 7339	"flow-action" on page 7259	"flow-active-timeout (Flow Monitoring)" on page 7261
"flow-collector (Flow Collection)" on page 7263	"flow-control" on page 7267	"flow-control (ACX Series)" on page 7265
"flow-control (Interfaces)" on page 7269	"flow-control-options" on page 7270	"flow-control-queue (Output Congestion Notification)" on page 7272
"flow-descriptions" on page 7274	"flow-detect-time (DDoS Flow Detection)" on page 7276	"flow-detection (DDoS Flow Detection)" on page 7278
"flow-detection (DDoS Packet Level)" on page 7279	"flow-detection-mode (DDoS Flow Detection)" on page 7281	"flow-detection-mode (DDoS Global Flow Detection)" on page 7283
"flow-export-destination (Flow Collection)" on page 7285	"flow-export-rate" on page 7286	"flow-export-timer" on page 7288
"flow-gate" on page 7289	"flow-groups" on page 7291	"flow-inactive-timeout (Flow Monitoring)" on page 7296
"flow-key (Flow Monitoring)" on page 7304	"flow-label" on page 7306	"flow-label-receive" on page 7307
"flow-label-receive-static" on page 7309	"flow-label-static" on page 7311	"flow-label-transmit" on page 7312
"flow-label-transmit-static" on page 7314	"flow-level-bandwidth (DDoS Flow Detection)" on page 7316	"flow-level-control (DDoS Flow Detection)" on page 7317

"flow-level-control (DDoS Global Flow Detection)" on page 7319	"flow-level-detection (DDoS Flow Detection)" on page 7320	"flow-map" on page 7322
"flow-monitoring" on page 7323	"flow-monitoring (Inline Monitoring Services)" on page 7326	"flow-monitoring (Services)" on page 7328
"flow-recover-time (DDoS Flow Detection)" on page 7330	"flow-report-rate (DDoS Flow Detection)" on page 7331	"flow-sample" on page 7360
"flow-server" on page 7344	"flow-session" on page 7347	"flow-specs" on page 7349
"flow-table-size" on page 7351	"flow-table-size (Chassis)" on page 7354	"flow-tap" on page 7355
"flow-tap-dtcp" on page 7357	"flow-timeout-time (DDoS Flow Detection)" on page 7359	"flows (PCC Rules)" on page 7362
"fnp" on page 7364	"force switch" on page 7374	"force-continue (OCS Partition)" on page 7365
"force-control-packets-on-transit-path" on page 7367	"force-discover (DHCP Client)" on page 7369	"force-entry (IDS MS-DPC)" on page 7370
"force-fbf-terms" on page 7381	"force-ip-reassembly" on page 7371	"force-premium (Firewall Filter Action)" on page 7373
"force-switch" on page 7375	"force-up" on page 7378	"force-up (EX and QFX Series)" on page 7376
"force-version (IEEE 802.1D STP)" on page 7379	"format (Flat-File Accounting Options)" on page 7382	"format (LRF Profile)" on page 7383
"format (Next Gen Services Service-Set Remote System Logging)" on page 7385	"format (Security Log)" on page 7386	"format (Security Log Stream)" on page 7388

"format (Unified Edge Gateways)" on page 7389	"forward-and-send-to-re" on page 7392	"forward-delay" on page 7393
"forward-only" on page 7395	"forward-only (DHCP Relay Agent)" on page 7396	"forward-only (DHCP Relay Agent Option)" on page 7398
"forward-only (DNS)" on page 7400	"forward-only-replies (DHCP Relay Agent)" on page 7401	"forward-policy-mismatch (Security Group VPN Member)" on page 7402
"forward-snooped-clients (DHCP Local Server)" on page 7404	"forward-snooped-clients (DHCP Relay Agent)" on page 7405	"forwarding (Diameter Network Element)" on page 7456
"forwarding-cache (Bridge Domains)" on page 7407	"forwarding-cache (Flow Maps)" on page 7408	"forwarding-class" on page 7410
"forwarding-class (BA Classifiers)" on page 7415	"forwarding-class (CoS Host Outbound Traffic)" on page 7417	"forwarding-class (Dynamic Scheduler Maps)" on page 7420
"forwarding-class (Fibre Channel Interfaces)" on page 7421	"forwarding-class (Firewall Filter Action)" on page 7423	"forwarding-class (for DHCP Snooping or DAI Packets)" on page 7425
"forwarding-class (Forwarding Policy)" on page 7426	"forwarding-class (Fragmentation)" on page 7428	"forwarding-class (Gigabit Ethernet IQ Classifier)" on page 7430
"forwarding-class (Interfaces)" on page 7432	"forwarding-class (Multifield Classifiers)" on page 7433	"forwarding-class (PCC Action Profiles)" on page 7434
"forwarding-class (Restricted Queues)" on page 7436	"forwarding-class (RFC 2544 Benchmarking)" on page 7437	"forwarding-class (Sampling)" on page 7439
"forwarding-class (Services CoS Fragmentation Properties)" on page 7440	"forwarding-class (Services PIC Classifiers)" on page 7442	"forwarding-class-accounting" on page 7413

"forwarding-class-default (Forwarding Policy)" on page 7418	"forwarding-class-set" on page 7443	"forwarding-class-set (Interconnect Device)" on page 7444
"forwarding-class-sets" on page 7446	"forwarding-classes" on page 7447	"forwarding-context (Protocols BGP)" on page 7453
"forwarding-db-size" on page 7454	"forwarding-instance identifier" on page 7457	"forwarding-mode (100-Gigabit Ethernet)" on page 7459
"forwarding-mode (PTX Series Packet Transport Routers)" on page 7460	"forwarding-mode (Security Content Security Policy)" on page 7461	"forwarding-options (Routing Instance)" on page 7463
"forwarding-options (Security)" on page 7473	"forwarding-options (Unified Forwarding Table Profile)" on page 7493	"forwarding-options (Bridge Domains or VLANs)" on page 7467
"forwarding-packages" on page 7478	"forwarding-policy" on page 7479	"forwarding-process" on page 7481
"forwarding-rib" on page 7483	"forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)" on page 7485	"forwarding-state-bit (Per Family for BGP Graceful Restart)" on page 7488
"forwarding-table" on page 7490	"fpc (Chassis)" on page 7501	"fpc (DDoS)" on page 7505
"fpc error" on page 7507	"fpc (Interfaces)" on page 7511	"fpc (Junos Fusion)" on page 7512
"fpc (Junos Fusion Satellite Device Cluster)" on page 7514	"fpc (M320, T320 and T640 Devices)" on page 7516	"fpc (MX Series 5G Universal Routing Platforms)" on page 7519
"fpc (Notification Control)" on page 7523	"fpc (PTX Series Devices)" on page 7525	"fpc (QFX Series)" on page 7499
"fpc (Subscriber Limits)" on page 7533	"fpc (TX Matrix and TX Matrix Plus Routers)" on page 7534	"fpc-feb-connectivity" on page 7509

"fpc-nmi-volt-fail-knob" on page 7521	"fpc-offline-on-blackholing" on page 7524	"fpc-restart" on page 7528
"fpc-resync" on page 7529	"fpc-slice" on page 7530	"fpcs (Node Slicing)" on page 7536
"fpcs (NSSU Upgrade Groups)" on page 7538	"fragment (IDS Screen Next Gen Services)" on page 7540	"fragment-limit" on page 7541
"fragment-threshold" on page 7542	"fragment-threshold (Forwarding Class Maps)" on page 7544	"fragment-threshold (Interfaces LSQ)" on page 7545
"fragmentation-map" on page 7547	"fragmentation-maps" on page 7548	"frame-error" on page 7550
"frame-mode (Dynamic Traffic Shaping)" on page 7552	"frame-period" on page 7554	"frame-period-summary" on page 7555
"frame-relay-de (CoS Interfaces)" on page 7557	"frame-relay-de (CoS Loss Priority)" on page 7558	"frame-relay-de (CoS Rewrite Rule)" on page 7560
"frame-relay-de (Defining Loss Priority Maps)" on page 7561	"frame-relay-de (Defining Loss Priority Rewrites)" on page 7563	"framed-ip-address" on page 7564
"framed-ipv6-prefix" on page 7566	"framed-pool" on page 7567	"framing (10-Gigabit Ethernet Interfaces)" on page 7568
"framing (E1 Options)" on page 7571	"framing (E1 Options for BITS Interfaces)" on page 7572	"framing (T1 Options)" on page 7573
"framing (T1 Options for BITS Interfaces)" on page 7575	"free-fw-memory-watermark (Resource Monitor)" on page 7576	"free-heap-memory-watermark (Resource Monitor)" on page 7578
"free-nh-memory-watermark (Resource Monitor)" on page 7579	"frequency-lock-threshold" on page 7580	"from (Application Aware Access List)" on page 7582
"from (Captive Portal Content Delivery Tags)" on page 7587	"from (Firewall)" on page 7585	"from (HTTP Header Enrichment)" on page 7588

"from (IDS MS-DPC)" on page 7590	"from (PCC Rules)" on page 7591	"from (PCP)" on page 7593
"from (Protocols MPLS)" on page 7595	"from (Security NAT)" on page 7596	"from (Services CoS)" on page 7597
"from (Services NAT)" on page 7599	"from (Services Stateful Firewall) " on page 7600	"from (services-ipsec-vpn)" on page 7584
"from (Subscriber Secure Policy)" on page 7602	"from (TDF Domain Selection)" on page 7603	"from-zone (Security IDP Policy)" on page 7607
"from-zone (Security Policies)" on page 7608	"from-zone (Security Policies Global)" on page 7612	"fru-poweron-sequence" on page 7613
"fte (Port)" on page 7615	"ftp" on page 7617	"ftp (Access)" on page 7619
"ftp (Content Security Policy Anti-Virus)" on page 7621	"ftp (Content Security Policy Content Filtering)" on page 7622	"ftp (Flow Collector Files)" on page 7624
"ftp (Security ALG)" on page 7626	"ftp (Services CoS)" on page 7628	"ftp (Services CoS Next Gen Services)" on page 7629
"ftp (Transfer Log Files)" on page 7631	"full" on page 7632	"function (Diameter Network Element)" on page 7634
"function (Diameter Network Element for Subscriber Management)" on page 7633	"function (Diameter Route)" on page 7636	"functional-zone" on page 7638

G

"g-duplicates-dropped-periodicity" on page 7645	"g-max-duplicates" on page 7646	"garp-na-enable" on page 7648

"gate-status" on page 7649	"gate-timeout" on page 7651	"gatekeeper" on page 7653
"gateway (Security Group VPN Member IKE)" on page 7658	"gateway (Security Group VPN Server IKE)" on page 7660	"gateway (Security IKE)" on page 7663
"gateway-name (LNS Local Gateway)" on page 7654	"gateway-name (Tunnel Profile Remote Gateway)" on page 7655	"gateway-name (Tunnel Profile Source Gateway)" on page 7657
"general-ikeid" on page 7669	"general-param (Flat-File Accounting Options)" on page 7670	"generate" on page 7672
"generate-event" on page 7674	"generate-snmp-traps" on page 7676	"geo-ip (global-filter)" on page 7678
"geo-ip (group-filter)" on page 7679	"get-route-range" on page 7681	"ggsn-address (OCS Partition)" on page 7682
"ggsn-mcc-mnc (OCS Partition)" on page 7684	"gigether-options" on page 7685	"gigether-options (Chassis Cluster)" on page 7694
"global (DDoS)" on page 7703	"global (Gx-Plus)" on page 7706	"global (OCS)" on page 7725
"global (PCRF)" on page 7731	"global (Security IDP)" on page 7733	"global (Security Policies)" on page 7734
"global-advertisements-threshold" on page 7696	"global-arp-prefix-limit (Host Fast Reroute)" on page 7698	"global-config (Services)" on page 7700
"global-dns-stats-log-timer" on page 7705	"global-limit" on page 7708	"global-load-balancing" on page 7744
"global-mac-ip-limit" on page 7709	"global-mac-ip-table-aging-time" on page 7711	"global-mac-limit" on page 7712
"global-mac-limit (Protocols)" on page 7714	"global-mac-move" on page 7716	"global-mac-statistics" on page 7717

"global-mac-table-aging-time" on page 7719	"global-mode (Protocols)" on page 7721	"global-no-control-mac-aging" on page 7722
"global-no-mac-learning" on page 7723	"global-options (Juniper Secure Connect)" on page 7726	"global-parameters (Easy EVPN LAG Configuration)" on page 7728
"global-supplementary-blackout-timer (Host Fast Reroute)" on page 7737	"global-threshold" on page 7740	"global-wait-to-restore" on page 7741
"global-weight" on page 7743	"gnf" on page 7746	"gpid" on page 7748
"gprs" on page 7750	"gprs-gtp-profile" on page 7755	"gprs-sctp-profile" on page 7757
"graceful-deletion-timeout" on page 7758	"graceful-restart (Enabling Globally)" on page 7760	"graceful-restart (Fabric Control)" on page 7762
"graceful-restart (Long-Lived for BGP Helper)" on page 7763	"graceful-restart (Long-Lived for BGP Restarter)" on page 7765	"graceful-restart (Multicast Snooping)" on page 7768
"graceful-restart (Protocols BGP)" on page 7769	"graceful-restart (Protocols IS-IS)" on page 7771	"graceful-restart (Protocols LDP)" on page 7773
"graceful-restart (Protocols OSPF)" on page 7775	"graceful-restart (Protocols PIM)" on page 7777	"graceful-restart (Protocols RIP)" on page 7779
"graceful-restart (Protocols RIPng)" on page 7780	"graceful-restart (Protocols RSVP)" on page 7782	"graceful-shutdown (Protocols BGP)" on page 7784
"graceful-switchover (Chassis)" on page 7786	"gratuitous-arp-count" on page 7787	"gratuitous-arp-reply" on page 7789
"gre" on page 7790	"gre (Ids-option)" on page 7791	"gre (Routing Options)" on page 7797
"gre-in" on page 7793	"gre-out" on page 7794	"gre-performance-acceleration (Security Flow)" on page 7796

"greater-than" on page 7798	"gres-route-flush-delay (Subscriber Management)" on page 7800	"group" on page 7801
"group (Bridge Domains)" on page 7806	"group (Configuring Group Name)" on page 7807	"group (Defining Access Privileges for an SNMPv3 Group)" on page 7813
"group (DHCP Local Server)" on page 7814	"group (DHCP Relay Agent)" on page 7820	"group (DHCP Security)" on page 7826
"group (DHCP Security for MX Series)" on page 7827	"group (DHCP Spoofing Prevention)" on page 7829	"group (Distributed IGMP)" on page 7830
"group (Dynamic IGMP Interface)" on page 7832	"group (Dynamic MLD Interface)" on page 7834	"group (IGMP Snooping)" on page 7835
"group (Origin Validation for BGP)" on page 7850	"group (Protocols BGP)" on page 7863	"group (Protocols IGMP)" on page 7869
"group (Protocols MLD)" on page 7870	"group (Protocols MLD Snooping)" on page 7872	"group (Protocols MSDP)" on page 7873
"group (Protocols PIM)" on page 7876	"group (Protocols RIP)" on page 7877	"group (Protocols RIPng)" on page 7880
"group (Redundant Trunk Groups)" on page 7887	"group (Routing Instances)" on page 7889	"group (RPF Selection)" on page 7893
"group (Security Group VPN)" on page 7894	"group (Static Subscribers)" on page 7897	"group (System Services DHCP)" on page 7899
"group (Traffic Load Balancer)" on page 7906	"group (Uplink Failure Detection)" on page 7910	"group-address (Routing Instances Tunnel Group)" on page 7802
"group-address (Routing Instances VPN)" on page 7804	"group-count (Dynamic MLD Interface)" on page 7809	"group-count (Protocols IGMP)" on page 7810

"group-count (Protocols MLD)" on page 7812	"group-increment (Dynamic MLD Interface)" on page 7837	"group-increment (Protocols IGMP)" on page 7838
"group-increment (Protocols MLD)" on page 7840	"group-limit (Dynamic IGMP Interface)" on page 7841	"group-limit (Dynamic MLD Interface)" on page 7843
"group-limit (IGMP) " on page 7844	"group-limit (IGMP and MLD Snooping)" on page 7846	"group-limit (Protocols MLD)" on page 7847
"group-members" on page 7849	"group-policy (Dynamic IGMP Interface)" on page 7852	"group-policy (Dynamic MLD Interface)" on page 7854
"group-policy (Protocols IGMP)" on page 7855	"group-policy (Protocols IGMP AMT Interface)" on page 7856	"group-policy (Protocols MLD)" on page 7858
"group-profile (Access)" on page 7859	"group-profile (Group Profile)" on page 7861	"group-range (Data MDTs)" on page 7882
"group-range (MBGP MVPN Tunnel)" on page 7883	"group-ranges" on page 7885	"group-rp-mapping" on page 7891
"group-threshold (Protocols IGMP Interface)" on page 7903	"group-threshold (Protocols MLD Interface)" on page 7905	"group-type (Unknown Unicast Forwarding)" on page 7908
"group-vpn" on page 7911	"grouped-ie-profile" on page 7915	"groups" on page 7916
"grpc" on page 7920	"grpc-replay" on page 7922	"grpc-tunnel " on page 7923
"gsm-options" on page 7926	"gsmp-syn-timeout (ANCP)" on page 7928	"gsmp-syn-wait (ANCP)" on page 7929
"gtp" on page 7931	"gtp-header-offset" on page 7936	"gtp-tunnel-endpoint-identifier" on page 7938
"guaranteed-rate (Dynamic Traffic Shaping)" on page 7942	"guaranteed-rate (Traffic Control Profiles)" on page 7940	"guard-band" on page 7944

"guard-band (PoE)" on page 7946	"guard-interval" on page 7947	"guest-network-functions" on page 7949
"gw-interface" on page 7951	"gx-plus (Gx-Plus)" on page 7952	"gx-profile" on page 7954

H

"h323" on page 7961	"halt-on-prefix-down (RFC 2544 Benchmarking)" on page 7963	"handover-default" on page 7964
"handover-group" on page 7966	"handshake-timeout" on page 7967	"hard-limit" on page 7969
"hard-limit-target" on page 7970	"hardware-assisted-keepalives" on page 7971	"hardware-assisted-timestamping" on page 7973
"hardware-timestamp" on page 7975	"hardware-upgrade" on page 7976	"has-prefix (Unified Edge Gateways)" on page 7977
"has-suffix" on page 7979	"hash-key (Chassis)" on page 7980	"hash-key (Chassis LAG)" on page 7982
"hash-key (Forwarding Options)" on page 7984	"hash-key (Forwarding Options)" on page 7986	"hash-keys (Aggregated Multiservices)" on page 7992
"hash-keys (Interfaces)" on page 7995	"hash-mode" on page 7997	"hash-parameters" on page 8000
"hash-seed" on page 8002	"hash-seed (profile)" on page 8003	"hcm (HTTP Header Enrichment)" on page 8005
"hcm-profile (HTTP Header Enrichment)" on page 8007	"hcm-profile (PCC Action Profiles)" on page 8008	"header-integrity-check" on page 8010

"header-integrity-check (Next Gen Services)" on page 8012	"header-length" on page 8014	"header-type" on page 8015
"health-monitor" on page 8016	"health-monitor (KHMS)" on page 8018	"heartbeat-address (MX Series Virtual Chassis)" on page 8020
"heartbeat-interval" on page 8022	"heartbeat-threshold" on page 8024	"heartbeat-timeout (MX Series Virtual Chassis)" on page 8025
"heartbeat-tos (MX Series Virtual Chassis)" on page 8027	"hello-acknowledgements" on page 8029	"hello-dead-interval" on page 8030
"hello-interval" on page 8032	"hello-interval (L2TP)" on page 8033	"hello-interval (LMP)" on page 8035
"hello-interval (Protocols LDP)" on page 8036	"hello-interval (Protocols RSVP)" on page 8038	"hello-message (NETCONF)" on page 8039
"hello-time" on page 8041	"hello-timer" on page 8043	"helper-disable (LDP)" on page 8046
"helper-disable (Multiple Protocols)" on page 8048	"helpers" on page 8049	"hide-avps" on page 8053
"hide-via-headers" on page 8054	"hierarchical-policer" on page 8055	"hierarchical-policer (aggregate or premium)" on page 8057
"hierarchical-scheduler" on page 8060	"hierarchical-scheduler (Subscriber Interfaces on MX Series Routers)" on page 8062	"high-availability (Chassis)" on page 8064
"high-availability (security cloud)" on page 8070	"high-availability (Security IDP)" on page 8073	"high-availability-options (Aggregated Multiservices)" on page 8068
"high-cos-queue-threshold" on page 8074	"high-polarization" on page 8076	"high-power" on page 8077

"high-threshold (Resource Monitor)" on page 8078	"high-watermark" on page 8080	"hint" on page 8081
"history" on page 8082	"history-size" on page 8084	"hold off time" on page 8097
"hold-down" on page 8086	"hold-down-interval" on page 8088	"hold-interval" on page 8089
"hold-interval (Clock Synchronization)" on page 8091	"hold-interval (OAM)" on page 8093	"hold-interval (OAM CFM, for EX Series Switch Only)" on page 8094
"hold-interval (Protection Group)" on page 8095	"hold-time" on page 8101	"hold-time (ae interface) " on page 8098
"hold-time (OAM)" on page 8102	"hold-time (Protocols BGP)" on page 8104	"hold-time (Protocols DVMRP)" on page 8106
"hold-time (Protocols IS-IS)" on page 8108	"hold-time (Protocols LDP)" on page 8110	"hold-time (Protocols MSDP)" on page 8112
"hold-time (Protocols PIM)" on page 8114	"hold-time (Redundant Ethernet Interfaces)" on page 8115	"hold-time up" on page 8117
"hold-time (VRRP)" on page 8119	"holddown (candidate-uplink-port-profile)" on page 8120	"holddown (Protocols RIP)" on page 8125
"holddown (Protocols RIPng)" on page 8127	"holddown-interval" on page 8122	"holddown-interval (protocols ldp)" on page 8124
"holdover-mode-disable" on page 8128	"holdover-time-error-budget" on page 8129	"home-address" on page 8131
"hop-by-hop-header" on page 8132	"hop-limit" on page 8134	"host (Address-Assignment Pools)" on page 8136
"host (Diameter Origin)" on page 8140	"host (Interfaces)" on page 8144	"host (L2TP)" on page 8146

"host (Next Gen Services Service-Set Remote System Logging)" on page 8152	"host (Security IDP Sensor Configuration)" on page 8165	"host (Security Log)" on page 8166
"host (Security Web Filtering)" on page 8167	"host (service-set)" on page 8169	"host (System)" on page 8171
"host-address-base" on page 8137	"host-address-base (Source NAT Next Gen Services)" on page 8139	"host-fast-reroute" on page 8141
"host-inbound-traffic" on page 8142	"host-name" on page 8147	"host-name (DHCP Relay Agent)" on page 8150
"host-name (DHCPv6 Relay Agent)" on page 8151	"host-name (System)" on page 8148	"host-only-interface" on page 8154
"host-os forwarding-options analyzer" on page 8155	"host-outbound media-interface" on page 8157	"host-outbound-traffic (Class-of-Service)" on page 8158
"host-outbound-traffic (Class-of-Service)" on page 8160	"host-outbound-traffic (Multicast Snooping)" on page 8162	"host-prefix-only" on page 8163
"host-vpn" on page 8174	"hosted-service-identifier" on page 8176	"hosted-services" on page 8178
"hostkey-algorithm" on page 8179	"hostkey-algorithm-list" on page 8182	"hot-root-standby (MBGP MVPN)" on page 8184
"hot-standby" on page 8187	"hot-standby (Protocols Layer 2 Circuit)" on page 8189	"hot-standby (Routing Instance)" on page 8186
"hot-standby-vc-on (Protocols Layer 2 Circuit)" on page 8190	"http (Access)" on page 8192	"http (icap-redirect profile)" on page 8193
"http (REST API)" on page 8202	"http (System Services)" on page 8204	"http (Web Management)" on page 8205

"http-log-multiple-transactions (LRF Profile)" on page 8195	"http-persist" on page 8196	"http-profile (Security Antivirus)" on page 8198
"http-profile (Security Content Filtering)" on page 8199	"http-profile (Security Web Filtering)" on page 8200	"http-reassemble" on page 8201
"https (REST API)" on page 8207	"https (System Services)" on page 8208	"https (Web Management)" on page 8211
"hub-assist" on page 8213	"hugepages" on page 8215	"hw-db-profile" on page 8216
"hybrid" on page 8220	"hyper-mode (forwarding-options)" on page 8221	

I

"icap-redirect" on page 8245	"iccp" on page 8250	"iccp (Traceoptions)" on page 8247
"icl-down-delay" on page 8252	"icmp" on page 8253	"icmp (Error Message Rate Limit)" on page 8257
"icmp (IDS Screen Next Gen Services)" on page 8260	"icmp (Security IDP Custom Attack)" on page 8264	"icmp (Security IDP Signature Attack)" on page 8265
"icmp (Security Screen)" on page 8267	"icmp (System)" on page 8269	"icmp6 (Error Message Rate Limit)" on page 8273
"icmp-code" on page 8255	"icmp-code (Applications)" on page 8256	"icmp-fragment-check (IDS MS-MPC)" on page 8259
"icmp-large-packet-check (IDS MS-MPC)" on page 8262	"icmp-mapping (Application Identification)" on page 8263	"icmp-type" on page 8270

"icmp-type (Applications)" on page 8271	"icmpv6 (Security IDP)" on page 8277	"icmpv6 (Security IDP Custom Attack)" on page 8278
"icmpv6-malformed" on page 8274	"icmpv6-malformed (IDS Screen Next Gen Services)" on page 8276	"id" on page 8279
"id (MACsec for MX Series)" on page 8283	"id (Protocols OpenFlow)" on page 8284	"id-components" on page 8281
"identification (Security ICMP Headers)" on page 8286	"identification (Tunnel Profile)" on page 8287	"identifier (VPLS Multihoming for FEC 129)" on page 8288
"identity" on page 8290	"identity-context-profile" on page 8291	"identity-management" on page 8293
"identity-provider (SAML)" on page 8298	"idle-after-switch-over" on page 8300	"idle-standby-path-switchover-delay" on page 8301
"idle-timeout" on page 8303	"idle-timeout (Access)" on page 8306	"idle-timeout (L2TP)" on page 8307
"idle-timeout (System)" on page 8309	"idle-timeout (System Services)" on page 8311	"idle-timeout (Unified Edge)" on page 8304
"idp (Application Services)" on page 8312	"idp (logical-systems)" on page 8313	"idp (Security)" on page 8321
"idp (Security Alarms)" on page 8333	"idp (SNMP)" on page 8334	"idp-policy" on page 8315
"idp-policy (Application Services)" on page 8316	"idp-policy (Security) " on page 8317	"ids-option" on page 8335
"ids-rule-sets" on page 8340	"ids-rules" on page 8341	"ie-set" on page 8343
"ieee802.1p" on page 8366	"ieee-802.1 (Class of Service)" on page 8344	"ieee-802.1 (Classifier on Physical Interface)" on page 8348

"ieee-802.1 (Congestion Notification)" on page 8350	"ieee-802.1 (Dynamic Classifiers)" on page 8351	"ieee-802.1 (Dynamic Rewrite Rules)" on page 8352
"ieee-802.1 (Fibre Channel Interfaces)" on page 8354	"ieee-802.1 (Host Outbound Traffic)" on page 8356	"ieee-802.1 (Input Congestion Notification)" on page 8357
"ieee-802.1 (Output Congestion Notification)" on page 8358	"ieee-802.1 (Rewrite Rules on Logical Interface)" on page 8360	"ieee-802.1 (Rewrite Rules on Physical Interface)" on page 8362
"ieee-802.1ad" on page 8363	"ieee-802-3az-eee" on page 8364	"ietf-mode" on page 8367
"if-exceeding (EX)" on page 8370	"if-exceeding (Hierarchical Policer)" on page 8371	"if-exceeding (Policer)" on page 8373
"if-exceeding (QFX)" on page 8368	"if-exceeding-pps (Hierarchical Policer)" on page 8374	"if-exceeding-pps (Policer)" on page 8376
"if-route-exists" on page 8377	"ifl-subscriber" on page 8379	"igmp" on page 8381
"igmp (Dynamic Profiles)" on page 8383	"igmp-querier (QFabric Systems only)" on page 8385	"igmp-snooping" on page 8387
"igmp-snooping-options" on page 8393	"ignore" on page 8394	"ignore (Event Policy)" on page 8399
"ignore-all" on page 8394	"ignore-attached-bit" on page 8396	"ignore-encapsulation-mismatch" on page 8397
"ignore-entry" on page 8399	"ignore-l3-incompletes" on page 8400	"ignore-lsp-metrics" on page 8402
"ignore-magic-number-mismatch (Access Group Profile)" on page 8403	"ignore-magic-number-mismatch (Dynamic Profiles)" on page 8405	"ignore-memory-overflow" on page 8407

"ignore-mtu-mismatch" on page 8408	"ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-overflow" on page 8409	"ignore-reassembly-overflow" on page 8411
"ignore-server-validation (Application Identification)" on page 5924	"ignore-stp-topology-change" on page 8412	"igp-synchronization" on page 8414
"ike" on page 8415	"ike (Access Profile)" on page 8419	"ike (High Availability)" on page 8424
"ike (Security)" on page 8432	"ike (Security Group VPN Member)" on page 8438	"ike (Security Group VPN Server)" on page 8440
"ike (Security IPsec VPN)" on page 8443	"ike (Security Services)" on page 8446	"ike-access-profile" on page 8418
"ike-esp-nat" on page 8421	"ike-log" on page 8426	"ike-phase1-failures" on page 8428
"ike-phase2-failures" on page 8429	"ike-secrets" on page 8431	"image" on page 8447
"imap-profile (Security Content Security Policy Antivirus)" on page 8450	"imap-profile (Security Content Security Policy Content Filtering)" on page 8451	"immediate-accounting-response" on page 8452
"immediate-leave" on page 8454	"immediate-leave (Dynamic IGMP Interface)" on page 8456	"immediate-leave (Dynamic MLD Interface)" on page 8458
"immediate-update" on page 8459	"import (BGP)" on page 8461	"import (Classifiers)" on page 8466
"import (CoS)" on page 8467	"import (MPLS Traffic Engineering Database)" on page 8471	"import (Protocols DVMRP)" on page 8474
"import (Protocols IS-IS)" on page 8476	"import (Protocols LDP)" on page 8477	"import (Protocols MSDP)" on page 8478

"import (Protocols OSPF)" on page 8463	"import (Protocols PIM)" on page 8480	"import (Protocols PIM Bootstrap)" on page 8482
"import (Protocols RIP)" on page 8483	"import (Protocols RIPng)" on page 8485	"import (Rewrite Rules)" on page 8486
"import (Routing Options)" on page 8465	"import-labeled-routes (Routing Instances VPLS)" on page 8469	"import-policy" on page 8473
"import-rib" on page 8488	"import-target" on page 8489	"imsi-prefix" on page 8491
"in-place-lsp-bandwidth-update (Protocols MPLS)" on page 8493	"in-service" on page 8494	"inactive-media-timeout (Security MGCP)" on page 8496
"inactive-media-timeout (Security SCCP)" on page 8497	"inactive-media-timeout (Security SIP)" on page 8499	"inactivity-asymm-tcp-timeout (Service Set Next Gen Services)" on page 8500
"inactivity-interval (Fabric Load Balance)" on page 8501	"inactivity-probe-duration" on page 8503	"inactivity-timeout" on page 8505
"inactivity-timeout (Applications)" on page 8508	"inactivity-timeout (Security Persistent NAT)" on page 8509	"inactivity-timeout (services)" on page 8506
"inactivity-timeout (Services RPM)" on page 8510	"inband-flow-telemetry" on page 8512	"inband-ztp" on page 8516
"include (Diameter Gx Profiles)" on page 8524	"include (Dynamic Access-Line-Identifier VLANs)" on page 8525	"include-sci" on page 8536
"include-all (for Administrative Groups)" on page 8518	"include-all (for Fast Reroute)" on page 8519	"include-any (for Administrative Groups)" on page 8521
"include-any (for Fast Reroute)" on page 8522	"include-igp-metric" on page 8527	"include-ipv6 (Gx-Plus)" on page 8528

"include-irb-and-l2" on page 8529	"include-mp-next-hop" on page 8532	"include-option-82 (DHCP Local Server)" on page 8534
"include-sci (MACsec for MX Series)" on page 8537	"inclusive" on page 8539	"incoming-map" on page 8540
"incoming-queue" on page 8541	"independent-domain" on page 8543	"indexed-load-balance" on page 8545
"indication" on page 8546	"indication-polarity" on page 8548	"indirect-next-hop" on page 8549
"indirect-next-hop-change-acknowledgements" on page 8551	"ineligible" on page 8553	"inet" on page 8554
"inet (AMT Protocol)" on page 8555	"inet (chassis)" on page 8558	"inet (CoS)" on page 8560
"inet (enhanced-hash-key)" on page 8561	"inet (forwarding options)" on page 8564	"inet (interfaces (QFX Series))" on page 8565
"inet (Interfaces)" on page 8568	"inet (Port Mirroring)" on page 8576	"inet (Routing Policy Condition)" on page 8588
"inet (Security Static NAT)" on page 8590	"inet (Subscriber Secure Policy)" on page 8591	"inet (TDF Subscriber Address)" on page 8593
"inet (TDF Subscriber Exclude Prefix)" on page 8594	"inet6" on page 8595	"inet6 (CoS Multidestination)" on page 8601
"inet6 (enhanced-hash-key)" on page 8602	"inet6 (forwarding optons)" on page 8605	"inet6 (interfaces (QFX Series))" on page 8607
"inet6 (Interfaces)" on page 8608	"inet6 (Subscriber Secure Policy)" on page 8616	"inet6 (TDF Subscriber Address)" on page 8617
"inet6 (TDF Subscriber Exclude Prefix)" on page 8619	"inet6-advertise-interval (VRRP)" on page 8596	"inet6-backup-router" on page 8598

"inet6-bgpmcast" on page 8599	"inet6-mvpn (BGP)" on page 8610	"inet6-mvpn (VRF Advertisement)" on page 8612
"inet6-options (Services)" on page 8613	"inet6-precedence (CoS Rewrite Rules)" on page 8614	"inet6-vpn" on page 8620
"inet-bgpmcast" on page 8557	"inet-mdt" on page 8570	"inet-mdt (Signaling)" on page 8571
"inet-mvpn (BGP)" on page 8574	"inet-mvpn (VRF Advertisement)" on page 8575	"inet-precedence" on page 8578
"inet-precedence (Classifier on Physical Interface)" on page 8580	"inet-precedence (CoS Classifiers)" on page 8581	"inet-precedence (CoS Rewrite Rules)" on page 8582
"inet-precedence (Dynamic Classifiers)" on page 8584	"inet-precedence (Dynamic Rewrite Rules)" on page 8585	"inet-precedence (Rewrite Rules on Physical Interface)" on page 8587
"infinity" on page 8622	"infranet-controller" on page 8623	"ingress (Analyzer)" on page 8625
"ingress (Buffer Configuration)" on page 8626	"ingress (Chained Composite Next Hop)" on page 8628	"ingress (LSP)" on page 8632
"ingress (Port Mirroring)" on page 8643	"ingress (vlans)" on page 8653	"ingress-egress-recycle" on page 8631
"ingress-node-replication" on page 8634	"ingress-node-replication (EVPN)" on page 8635	"ingress-only" on page 8637
"ingress-policer-overhead" on page 8638	"ingress-policy" on page 8641	"ingress-queuing-filter" on page 8644
"ingress-rate-limit" on page 8646	"ingress-replication" on page 8647	"ingress-shaping-overhead" on page 8649
"ingress-stats (Flat-File Accounting Options)" on page 8651	"inherit-advertisement-interval" on page 8654	"init-command-string" on page 8655

"init-descriptor" on page 8657	"initial-route-check" on page 8658	"initial-tcp-mss" on page 8659
"initiate-dead-peer-detection" on page 8661	"initiate-ncp (Dynamic and Static PPP)" on page 8662	"initiation (Services)" on page 8664
"inline" on page 8666	"inline-fpga-crypto" on page 8668	"inline-jflow" on page 8669
"inline-jflow (Forwarding Options)" on page 8671	"inline-monitoring" on page 8672	"inline-services" on page 8674
"inline-services (PIC level)" on page 8676	"inline-tap" on page 8677	"inner-tag-protocol-id" on page 8679
"inner-tag-protocol-id (Dynamic VLANs)" on page 8680	"inner-vlan-id" on page 8682	"inner-vlan-id (Dynamic VLANs)" on page 8683
"inner-vlan-id-swap-ranges" on page 8685	"input" on page 8687	"input (Analyzer)" on page 8688
"input (Congestion Notification)" on page 8692	"input (Congestion Notification)" on page 8693	"input (Dynamic Service Sets)" on page 8697
"input (Fibre Channel Interfaces)" on page 8700	"input (Forwarding Table)" on page 8703	"input (Interfaces)" on page 8707
"input (Port Mirroring)" on page 8714	"input (Routing Options RIB)" on page 8717	"input (Sampling)" on page 8718
"input-chain" on page 8690	"input-current (T4000 Routers)" on page 8695	"input-excess-bandwidth-share" on page 8698
"input-firewall-optimized-profile" on page 8701	"input-hierarchical-policer" on page 8704	"input-interface-index" on page 8705
"input-list" on page 8708	"input-native-vlan-push" on page 8710	"input-packet-rate-threshold" on page 8711

"input-policer" on page 8713	"input-priority-map" on page 8716	"input-scheduler-map" on page 8720
"input-shaping-rate (Logical Interface)" on page 8721	"input-shaping-rate (Physical Interface)" on page 8723	"input-three-color" on page 8724
"input-traffic-control-profile" on page 8726	"input-traffic-control-profile-remaining" on page 8728	"input-vlan-map " on page 8729
"input-vlan-map (Aggregated Ethernet)" on page 8731	"input-vlan-map (Dynamic Interfaces)" on page 8732	"inputs (Analytics)" on page 8734
"inspection-profile" on page 8739	"install" on page 8741	"install (Protocols MPLS)" on page 8744
"install (Routing Options)" on page 8747	"install-nexthop" on page 8742	"instance" on page 8748
"instance (CASB)" on page 8754	"instance (Ephemeral Database)" on page 8756	"instance (Inline Monitoring)" on page 8750
"instance (OAM)" on page 8753	"instance (Port Mirroring)" on page 8762	"instance (Sampling)" on page 8766
"instance (Traffic Load Balancer)" on page 8770	"instance-export" on page 8759	"instance-import" on page 8760
"instance-role" on page 8764	"instance-shared" on page 8769	"instance-type" on page 8773
"integer" on page 8777	"intelligent-offload-tunable (IDP)" on page 8780	"intelligent-prescreening" on page 8779
"inter-area-prefix-export" on page 8782	"inter-area-prefix-import" on page 8783	"inter-as (Routing Instances)" on page 8785
"inter-domain" on page 8788	"inter-region" on page 8789	"inter-region-segmented" on page 8791

"inter-region-template" on page 8793	"inter-switch-link" on page 8795	"interconnect" on page 8797
"interconnect-device (Aliases)" on page 8800	"interconnect-device (Chassis)" on page 8802	"interconnect-device (Flow Groups)" on page 8804
"interconnected-vlan-list" on page 8808	"interconnected-vni-list" on page 8809	"interface (802.1X)" on page 8819
"interface (Access Port Security)" on page 8829	"interface (Accounting or Sampling)" on page 8831	"interface (advance-policy-based-routing)" on page 8833
"interface (Analyzer)" on page 8835	"interface (Backup Selection OSPF)" on page 8837	"interface (BOOTP)" on page 8841
"interface (BPDU Blocking)" on page 8842	"interface (Bridge Domains)" on page 8845	"interface (Captive Portal)" on page 8846
"interface (Chassis Cluster)" on page 8850	"interface (DCBX)" on page 8851	"interface (DHCP Local Server)" on page 8860
"interface (DHCP Relay Agent)" on page 8863	"interface (DHCP Security for MX Series)" on page 8866	"interface (DHCP Spoofing Prevention)" on page 8867
"interface (DNS, Port, and TFTP Packet Forwarding or Relay Agent)" on page 8868	"interface (Dynamic IGMP)" on page 8871	"interface (Dynamic Interface Sets)" on page 8873
"interface (Dynamic MLD)" on page 8875	"interface (Dynamic Router Advertisement)" on page 8877	"interface (Dynamic Routing Instances)" on page 8879
"interface (Dynamic Routing Options)" on page 8880	"interface (EVPN Routing Instances)" on page 8882	"interface (Export Profiles)" on page 8884
"interface (Fibre Channel Fabric)" on page 8886	"interface (FIP)" on page 8887	"interface (FIP Snooping)" on page 8889

"interface (Hierarchical CoS Schedulers)" on page 8896	"interface (Host Fast Reroute)" on page 8897	"interface (IEEE 802.1ag OAM Connectivity-Fault Management)" on page 8900
"interface (IGMP Snooping)" on page 8902	"interface (Layer 2 Protocol Tunneling)" on page 8904	"interface (LLDP)" on page 8817
"interface (LLDP-MED)" on page 8906	"interface (Logical Interfaces)" on page 8814	"interface (MLD Snooping)" on page 8918
"interface (Monitoring)" on page 8923	"interface (MPLS)" on page 8925	"interface (Multicast Scoping)" on page 8927
"interface (Multicast Static Routes)" on page 8928	"interface (Multichassis Protection)" on page 8930	"interface (MVRP)" on page 8931
"interface (Next-Hop Group)" on page 8939	"interface (OAM CFM, for EX Series Switch Only)" on page 8945	"interface (OAM Link-Fault Management)" on page 8947
"interface (PIC Bundle)" on page 8949	"interface (PoE)" on page 8951	"interface (Port Mirroring)" on page 8952
"interface (Port Mirroring)" on page 8954	"interface (Port Mirroring)" on page 8956	"interface (Power over Ethernet)" on page 8957
"interface (Precision Time Protocol)" on page 8815	"interface (Protocols DVMRP)" on page 8960	"interface (Protocols IGMP)" on page 8962
"interface (Protocols IPv6 Neighbor Discovery)" on page 8964	"interface (Protocols IS-IS)" on page 8966	"interface (Protocols Layer 2 Circuit)" on page 8974
"interface (Protocols Layer 2 VPN)" on page 8977	"interface (Protocols LDP)" on page 8978	"interface (Protocols Link Management)" on page 8980
"interface (Protocols MLD)" on page 8981	"interface (Protocols MPLS)" on page 8983	"interface (Protocols OSPF)" on page 8984

"interface (Protocols PIM)" on page 8991	"interface (Protocols RIFT)" on page 8994	"interface (Protocols Router Discovery)" on page 8998
"interface (Protocols RSVP)" on page 8999	"interface (RA Guard)" on page 9002	"interface (Redundant Trunk Groups)" on page 9008
"interface (Routing Instances)" on page 9012	"interface (Routing Options)" on page 9014	"interface (Scoping)" on page 9015
"interface (Secure Access Port)" on page 9017	"interface (Security NAT ARP Proxy)" on page 9018	"interface (Security NAT NDP Proxy)" on page 9020
"interface (Security Source NAT)" on page 9021	"interface (Security Source NAT Rule Set)" on page 9022	"interface (Service Interfaces)" on page 9023
"interface (Service Redundancy Group Monitor)" on page 9025	"interface (services)" on page 8811	"interface (Services)" on page 9028
"interface (Services Flow Tap)" on page 9029	"interface (Services PIC)" on page 9031	"interface (Session PICs)" on page 9033
"interface (SLAAC Snooping)" on page 9052	"interface (SNMP)" on page 9054	"interface (SNMP RMON History)" on page 9056
"interface (Spanning Tree)" on page 9057	"interface (Static MAC Bypass)" on page 9067	"interface (Static Subscriber Group)" on page 9068
"interface (Static Subscriber Username)" on page 9070	"interface (Storm Control)" on page 9074	"interface (Switching Options)" on page 9078
"interface (System Services DHCP)" on page 9079	"interface (Unknown Unicast Forwarding)" on page 9096	"interface (Virtual Tunnel in Routing Instances)" on page 9097
"interface (VoIP)" on page 9099	"interface (VPLS Mesh-Group)" on page 9102	"interface (VPLS Multihoming for FEC 129)" on page 9103

"interface (VPLS Routing Instances)" on page 9104	"interface (VRRP)" on page 8812	"interface-client-limit (DHCP Local Server)" on page 9106
"interface-client-limit (DHCP Relay Agent)" on page 9108	"interface-delete (Subscriber Management or DHCP Client Management)" on page 8853	"interface-description (DHCP Local Server)" on page 8854
"interface-description (DHCP Relay Agent)" on page 8856	"interface-description-format" on page 8858	"interface-down" on page 8870
"interface-group" on page 8890	"interface-group (Decapsulate GRE)" on page 8892	"interface-group (Routing Options)" on page 8894
"interface-group-down " on page 8893	"interface-id" on page 8898	"interface-in (Security Forwarding Options)" on page 8903
"interface-mac-ip-limit" on page 8908	"interface-mac-limit" on page 8910	"interface-mac-limit (VPLS)" on page 8913
"interface-map" on page 8915	"interface-mib (Enhanced Subscriber Management)" on page 8916	"interface-mode" on page 8920
"interface-monitor" on page 8922	"interface-name" on page 8934	"interface-name" on page 8935
"interface-name (DHCP Local Server)" on page 8937	"interface-name (DHCP Relay Agent)" on page 8938	"interface-none" on page 8941
"interface-notification (programmable-rpd)" on page 8942	"interface-num" on page 8943	"interface-out (Security Forwarding Options)" on page 8948
"interface-profile" on page 8959	"interface-range" on page 9006	"interface-range (QFX Series)" on page 9003
"interface-routes" on page 9010	"interface-service (Services Interfaces)" on page 9027	"interface-set" on page 9034

"interface-set (ANCP)" on page 9036	"interface-set (Dynamic Profiles)" on page 9037	"interface-set (Dynamic VLAN Interface Sets Association)" on page 9040
"interface-set (Dynamic VLAN Interface Sets Definition)" on page 9042	"interface-set (Ethernet Interfaces)" on page 9044	"interface-set (Hierarchical Schedulers)" on page 9045
"interface-set (IP Demux Interfaces)" on page 9046	"interface-shared" on page 9048	"interface-shared-with" on page 9049
"interface-shutdown-action" on page 9050	"interface-specific" on page 9060	"interface-specific (Dynamic Firewalls)" on page 9061
"interface-specific (Firewall Filters)" on page 9063	"interface-state" on page 9064	"interface-status-send-rdi" on page 9071
"interface-status-tlv" on page 9073	"interface-switch" on page 9076	"interface-tag (DHCP Local Server)" on page 9081
"interface-tag (DHCP Relay Server)" on page 9083	"interface-traceoptions (DHCP)" on page 9085	"interface-traceoptions (System Services DHCP)" on page 9087
"interface-transmit-statistics" on page 9090	"interface-type (Protocols OSPF)" on page 9091	"interface-type (Protocols RIP)" on page 9095
"interfaces" on page 9113	"interfaces (Aggregated Multiservices)" on page 9115	"interfaces (Analytics)" on page 9117
"interfaces (Analytics Resource)" on page 9119	"interfaces (ANCP)" on page 9121	"interfaces (Autoinstallation)" on page 9122
"interfaces (Chassis Synchronization Source)" on page 9124	"interfaces (Class of Service)" on page 9126	"interfaces (Class of Service)" on page 9128
"interfaces (Class of Service)" on page 9130	"interfaces (Dynamic CoS Definition)" on page 9133	"interfaces (EX Series switches)" on page 9135

"interfaces external" on page 9144	"interfaces (Junos Node Slicing)" on page 9145	"interfaces (MACsec)" on page 9147
"interfaces (MACsec for MX Series)" on page 9149	"interfaces (OVSDB)" on page 9151	"interfaces (Protocols OpenFlow)" on page 9153
"interfaces (Q-in-Q Tunneling)" on page 9154	"interfaces (QFX Series, ACX Series)" on page 9156	"interfaces (Security Zones)" on page 9165
"interfaces (Services Dynamic Flow Capture)" on page 9167	"interfaces (sFlow)" on page 9168	"interfaces (Static and Dynamic Subscribers)" on page 9170
"interfaces (Video Monitoring)" on page 9176	"interfaces (VNF)" on page 9111	"interfaces (Voice Services)" on page 9181
"interim-logging-interval" on page 9182	"interim-rate (Access)" on page 9183	"interleave-fragments" on page 9185
"internal" on page 9186	"internal (Security IPsec)" on page 9190	"internal-forwarding-bandwidth" on page 9187
"internal-node" on page 9189	"internet-options" on page 9192	"interpolate (Drop Profiles)" on page 9197
"interval" on page 9198	"interval (Accounting Options)" on page 9207	"interval (CFM MEP)" on page 9209
"interval (EX Series Switch Only)" on page 9210	"interval (Flat-File Accounting Options)" on page 9211	"interval (Health Monitor)" on page 9213
"interval (Interfaces)" on page 9214	"interval (OAM)" on page 9203	"interval (PoE)" on page 9216
"interval (Power over Ethernet)" on page 9217	"interval (Protocols PIM)" on page 9202	"interval (RMON)" on page 9206
"interval (Security Antivirus)" on page 9219	"interval (Security IDP)" on page 9220	"interval (services)" on page 9201

"interval (Services)" on page 9221	"interval (SNMP RMON)" on page 9223	"intra-area-prefix" on page 9224
"intra-as" on page 9226	"invalid-authentication-entry-timeout (Services User Identification Active Directory and ClearPass)" on page 9227	"ioc-npc-connectivity" on page 9230
"ip (IDS Screen Next Gen Services)" on page 9241	"ip (Security IDP Custom Attack)" on page 9266	"ip (Security Screen)" on page 9268
"ip-action (Security IDP Rulebase IPS)" on page 9231	"ip-address (Access Profile)" on page 9234	"ip-address (Port Mirroring)" on page 9238
"ip-address-first" on page 9235	"ip-address-owner" on page 9236	"ip-can-type (PCRF Partition)" on page 9239
"ip-in-udp" on page 9243	"ip-monitoring" on page 9244	"ip-monitoring (Services)" on page 9246
"ip-options-protocol-queue " on page 9249	"ip-prefix-routes" on page 9250	"ip-prefix-support" on page 9255
"ip-protocol-mapping (Application Identification)" on page 9258	"ip-query (Identity Management Advanced Query)" on page 9259	"ip-reassembly" on page 9262
"ip-reassembly (L2TP)" on page 9263	"ip-reassembly-rules (Service Set)" on page 9265	"ip-source-guard" on page 9272
"ip-source-guard (MX Series)" on page 9274	"ip-subscriber" on page 9275	"ip-swap" on page 9277
"ip-sweep" on page 9279	"ip-tunnel-rpf-check" on page 9281	"ip-user-mapping" on page 9283
"ip-version" on page 9285	"ip-version (Family MPLS)" on page 9286	"ipc" on page 9288

"ipcp-suggest-dns-option" on page 9289	"ipfix-sw-mode" on page 9291	"ipip" on page 9292
"ipip (Ids-option)" on page 9293	"ips" on page 9295	"ipsec" on page 9298
"ipsec (High Availability)" on page 9301	"ipsec (Security)" on page 9312	"ipsec (Security Group VPN Member)" on page 9314
"ipsec (Security Group VPN Server)" on page 9317	"ipsec (Security Interfaces)" on page 9319	"ipsec-group-vpn (Security Policies)" on page 9300
"ipsec-inside-interface" on page 9304	"ipsec-policy" on page 9306	"ipsec-sa" on page 9307
"ipsec-sa (Protocols BGP)" on page 9308	"ipsec-sa (Security Group VPN)" on page 9310	"ipsec-traffic-selector" on page 9321
"ipsec-vpn (Security Policies)" on page 9326	"ipsec-vpn-options" on page 9323	"ipsec-vpn-rules" on page 9325
"ipv4 (Family MPLS)" on page 9330	"ipv4 (Security IDP Signature Attack)" on page 9340	"ipv4-address (Steering Path)" on page 9327
"ipv4-dscp" on page 9329	"ipv4-flow-table-size " on page 9332	"ipv4-mask (HTTP Header Enrichment)" on page 9335
"ipv4-or-value (HTTP Header Enrichment)" on page 9337	"ipv4-prefix" on page 9338	"ipv4-template" on page 9344
"ipv4-template (Services)" on page 9345	"ipv6 (Family MPLS)" on page 9357	"ipv6-address (Steering Path)" on page 9347
"ipv6-extended" on page 9348	"ipv6-extended-attrib " on page 9349	"ipv6-extension-header" on page 9350
"ipv6-extension-header (IDS Screen Next Gen Services)" on page 9353	"ipv6-extension-header-limit" on page 9355	"ipv6-flow-label" on page 9359

"ipv6-flow-table-size " on page 9360	"ipv6-malformed-header" on page 9362	"ipv6-mask (HTTP Header Enrichment)" on page 9363
"ipv6-multicast-interfaces" on page 9364	"ipv6-or-value (HTTP Header Enrichment)" on page 9366	"ipv6-origin" on page 9367
"ipv6-prefix" on page 9368	"ipv6-source-guard" on page 9370	"ipv6-source-guard-sessions" on page 9372
"ipv6-target" on page 9374	"ipv6-template" on page 9375	"ipv6-template (Services)" on page 9376
"ipv6-tunneling" on page 9377	"ipvacl-firewall-l2-profile" on page 9378	"iq-policing-filter" on page 9380
"ir-mode" on page 9381	"iracl-firewall-ipv4-profile" on page 9383	"irb" on page 9384
"irb (Easy EVPN LAG Configuration)" on page 9386	"irb (Interfaces)" on page 9389	"irb-symmetric-routing" on page 9394
"is-ma" on page 9396	"isdn-options" on page 9397	"isid" on page 9399
"isid-list" on page 9400	"isis" on page 9401	"isis-instance" on page 9403
"iso-vpn" on page 9404	"isolated" on page 9406	"isolated-vlan" on page 9407
"isolated-vlan (MX Series)" on page 9408	"isolation-id" on page 9410	"isolation-vlan-id" on page 9411
"iteration-count" on page 9412	"iteration-period" on page 9414	"ivacl-firewall-no-portrange-profile" on page 9415
"ivlan-cfi (RFC 2544 Benchmarking)" on page 9416	"ivlan-id (RFC 2544 Benchmarking)" on page 9417	"ivlan-priority (RFC 2544 Benchmarking)" on page 9419

"iwf-params (tdm-options)" on page 9420	

J

"jdaf" on page 9422	"jeb" on page 9423	"jflow-log (Interfaces)" on page 9426
"jflow-log (Services)" on page 9427	"jnu-management" on page 9429	"join-load-balance" on page 9431
"join-prune-timeout" on page 9433	"join-timer (MVRP)" on page 9434	"json" on page 9436
"jsrc (Access Profile)" on page 9438	"jsrc (JSRC)" on page 9439	"jsrc-partition" on page 9441
"juniper-enhanced" on page 9442	"juniper-express-engine" on page 9444	"juniper-local" on page 9447

K

"kaspersky-lab-engine" on page 9450	"keep" on page 9452	"keep-address-and-control" on page 9455
"keep-alive (Protocols MSDP)" on page 9457	"keep-existing-steering" on page 9459	"keep-incoming-circuit-id (DHCP Relay Agent)" on page 9460
"keep-incoming-interface-id (DHCP Relay Agent)" on page 9462	"keep-incoming-remote-id (DHCP Relay Agent)" on page 9463	"keep-multiplier" on page 9465
"keepalive" on page 9466	"keepalive-interval" on page 9468	"keepalive-retries" on page 9469

"keepalive-time (Chassis)" on page 9471	"keepalive-time (Protocols OAM)" on page 9472	"keepalive-timeout" on page 9474
"keepalives" on page 9475	"keepalives (Dynamic Profiles)" on page 9478	"kernel-replication" on page 9480
"key" on page 9481	"key (Authentication Keychain)" on page 9485	"key (gre)" on page 9482
"key (interfaces)" on page 9484	"key (Junos FIPS)" on page 9494	"key (MACsec)" on page 9497
"key (MACsec for MX Series)" on page 9498	"key-server-priority (MACsec)" on page 9502	"key-chain (Authentication Keychain)" on page 9488
"key-chain (Protocols PIM)" on page 9490	"key-exchange" on page 9491	"key-length" on page 9495
"key-pair" on page 9500	"key-server-priority (MACsec for MX Series)" on page 9503	"krt-nexthop-ack" on page 9505
"krt-nexthop-ack-timeout" on page 9506		

L

"l2-channel-error-threshold (chassis)" on page 9520	"l2-channel-errors (chassis)" on page 9521	"l2-domain-id-for-l3" on page 9523
"l2-learning" on page 9524	"l2-querier" on page 9527	"l2-smart-policy" on page 9528
"l2-stats (Flat-File Accounting Options)" on page 9530	"l2circuit" on page 9531	"l2circuit-control-passthrough" on page 9534

"l2ckt" on page 9535	"l2pt (Destination Tunnels)" on page 9537	"l2tp" on page 9541
"l2tp (DDoS Protection) " on page 9544	"l2tp (Group Profile)" on page 9548	"l2tp (Profile)" on page 9553
"l2tp-access-profile" on page 9547	"l2tp-interface-id" on page 9550	"l2tp-maximum-session (Service Interfaces)" on page 9551
"l2tp-tunnel-session-identifier" on page 9558	"l2vpn" on page 9559	"l2vpn (routing-options)" on page 9564
"l2vpn-id" on page 9562	"l2vpn-use-bgp-rules-edit-protocols-bgp" on page 9565	"l3-interface (VLAN)" on page 9568
"l3-interface-ingress-counting" on page 9567	"l3vpn (ingress)" on page 9571	"l3vpn (transit)" on page 9573
"label" on page 9574	"label-allocation" on page 9576	"label-block-size" on page 9577
"label-position" on page 9579	"label-switched-path" on page 9580	"label-switched-path (Protocols IS-IS)" on page 9582
"label-switched-path (Protocols Link Management)" on page 9583	"label-switched-path (Protocols MPLS)" on page 9585	"label-switched-path (Protocols MVPN)" on page 9590
"label-switched-path (Protocols OSPF)" on page 9592	"label-switched-path-template" on page 9593	"label-switched-path-template (Container LSP)" on page 9595
"label-switched-path-template (Multicast)" on page 9596	"label-withdrawal-delay" on page 9598	"labeled-bgp" on page 9600
"labeled-unicast" on page 9601	"labeled-unicast (Protocols BGP)" on page 9603	"lacp" on page 9605
"lacp (802.3ad)" on page 9607	"lacp (802.3ad) (MX Series)" on page 9609	"lacp (Aggregated Ethernet)" on page 9610

"lacp (Interfaces)" on page 9614	"lacp (Protocols)" on page 9615	"land" on page 9617
"land (IDS Screen Next Gen Services)" on page 9622	"land-attack-check" on page 9619	"land-attack-check (IDS MS-MPC)" on page 9620
"lane" on page 9623	"lane-all" on page 9626	"language (Scripts)" on page 9630
"large" on page 9632	"large (IDS Screen Next Gen Services)" on page 9633	"large-community-count" on page 9634
"laser-enable" on page 9635	"last-block-recycle-timeout" on page 9637	"latency (Interfaces)" on page 9638
"latency-threshold" on page 9639	"layer2 (enhanced-hash-key)" on page 9644	"layer2-control" on page 9641
"layer2-liveness-detection (Receive)" on page 9647	"layer2-liveness-detection (Send)" on page 9649	"layer2-map" on page 9651
"layer2-policer" on page 9653	"layer2-policer (EX)" on page 9654	"layer2-policer (firewall)" on page 9662
"layer2-policer (Hierarchical Policer)" on page 9656	"layer2-protocol-tunneling" on page 9658	"layer2-unicast-replies" on page 9661
"layer3-domain-identifier" on page 9666	"lcd-menu" on page 9667	"lcp-max-conf-req" on page 9669
"lcp-renegotiation" on page 9671	"lcp-restart-timer" on page 9672	"ldap-options" on page 9673
"ldap-server" on page 9676	"ldap-url" on page 9678	"ldp" on page 9679
"ldp-p2mp" on page 9685	"ldp-stitching (Protocols OSPF)" on page 9686	"ldp-tunneling" on page 9687

"leaf" on page 9689	"leaf-tunnel-limit-inet (MVPN Selective Tunnels)" on page 9691	"leaf-tunnel-limit-inet6 (MVPN Selective Tunnels)" on page 9693
"learn-sip-register" on page 9695	"lease-time" on page 9696	"lease-time (dhcp-client)" on page 9698
"lease-time-threshold (DHCP Local Server and DHCP Relay Agent)" on page 9699	"lease-time-validation (DHCP Local Server and DHCP Relay Agent)" on page 9701	"leasequery (DHCP Relay Agent)" on page 9702
"least-fill" on page 9704	"leave-timer (MVRP)" on page 9705	"leaveall-timer (MVRP)" on page 9707
"led-beacon" on page 9710	"legacy-redirect-ip-action" on page 9711	"less-than" on page 9713
"level (App Engine Virtual Machine Management Service)" on page 9719	"level (CFM MEP)" on page 9721	"level (Connectivity Fault Management)" on page 9715
"level (Disk Partition)" on page 9718	"level (EX Series Switch Only)" on page 9722	"level (Global IS-IS)" on page 9724
"level (IS-IS Interfaces)" on page 9728	"level (OAM)" on page 9716	"level (Protocols RIFT)" on page 9734
"level (Services)" on page 9735	"license" on page 9737	"license-mode" on page 9739
"license-server" on page 9740	"lie-receive-address (Protocols RIFT)" on page 9742	"lie-transmit-address (Protocols RIFT)" on page 9744
"lifetime (Router Advertisement)" on page 9745	"lifetime (Router Discovery)" on page 9747	"lifetime-seconds" on page 9749
"lifetime-seconds (Security)" on page 9750	"light" on page 9752	"light-weight-dhcpv6-relay" on page 9755

"limit" on page 9757	"limit (Content Security Policy)" on page 9758	"limit (Security Log)" on page 9760
"limit (Security SCTP)" on page 9762	"limit-ports-per-address" on page 9759	"limit-session" on page 9764
"limit-session (IDS Screen Next Gen Services)" on page 9765	"limited-ifl-scaling" on page 9768	"line-encoding (E1 Options)" on page 9769
"line-encoding (T1 Options)" on page 9770	"line-identity (Dynamic Access-Line-Identifier VLANs)" on page 9772	"line-interval (Dampen Options)" on page 9773
"line-loopback" on page 9775	"line-protocol" on page 9776	"line-rate" on page 9778
"line-rate (Interfaces)" on page 9779	"linerate-mode" on page 9780	"link (Access)" on page 9781
"link-adjacency-loss" on page 9783	"link-degrade-monitor" on page 9784	"link-discovery" on page 9787
"link-down" on page 9790	"link-down (QFX Series)" on page 9788	"link-event-rate" on page 9791
"link-fault-management" on page 9792	"link-group-protection (Protocols IS-IS)" on page 9794	"link-layer-broadcast-inet-check" on page 9796
"link-layer-overhead" on page 9797	"link-management" on page 9798	"link-mode" on page 9800
"link-mtu" on page 9802	"link-protection" on page 9803	"link-protection (Chassis Cluster)" on page 9805
"link-protection (Dynamic LSPs)" on page 9807	"link-protection (Host Fast Reroute)" on page 9808	"link-protection (non-LACP)" on page 9810
"link-protection (Protocols IS-IS)" on page 9811	"link-protection (Protocols LACP)" on page 9813	"link-protection (Protocols OSPF)" on page 9814

"link-protection (RSVP)" on page 9816	"link-protection (Static LSPs)" on page 9818	"link-protection-sub-group (802.3ad)" on page 9820
"link-protection-sub-group (aggregated-ether-options)" on page 9821	"link-speed" on page 9823	"link-speed (Aggregated Ethernet)" on page 9826
"link-speed (Aggregated SONET/SDH)" on page 9828	"link-speed (Interfaces)" on page 9830	"link-to-disable" on page 9831
"link-to-monitor" on page 9832	"link-tracking (Network Isolation Groups)" on page 9833	"linkdown (satellite-policies alarm)" on page 9835
"linked-pool-aggregation (Address-Assignment Pools)" on page 9837	"linktrace" on page 9838	"list " on page 9839
"list (Security Content Filtering Block Mime)" on page 9841	"listen" on page 9842	"listening-mode" on page 9843
"liveness-detection" on page 9846	"liveness-detection (high availability)" on page 9848	"liveness-detection (iccp-peer)" on page 9845
"lldp" on page 9851	"lldp-med (Ethernet Switching)" on page 9857	"lldp-priority" on page 9860
"lmi (Ethernet OAM)" on page 9861	"lmi-type" on page 9863	"lmp-control-channel" on page 9864
"lmp-protocol" on page 9865	"load-balance" on page 9874	"load-balance (ecmp)" on page 9872
"load-balance (Forwarding Options)" on page 9878	"load-balance profile1/profile2" on page 9883	"load-balance (Protocols RSVP)" on page 9885
"load-balance-algorithm" on page 9876	"load-balance-group" on page 9880	"load-balance-label-capability" on page 9882

"load-balance-stateful (Aggregated Ethernet Interfaces)" on page 9886	"load-balancing-options (Aggregated Multiservices)" on page 9888	"load-balancing-options (Service Set)" on page 9890
"load-interval" on page 9892	"load-scripts-from-flash (Scripts)" on page 9893	"load-threshold" on page 9894
"load-type (Aggregated Ethernet Interfaces)" on page 9896	"loc-rib" on page 9897	"local" on page 9899
"local (Analytics Collector)" on page 9914	"local (certificates)" on page 9900	"local (Flat-File Access Profile)" on page 9942
"local-address (Junos Telemetry Interface)" on page 9902	"local-address (Protocols AMT)" on page 9904	"local-address (Protocols BGP)" on page 9905
"local-address (Protocols Link Management)" on page 9908	"local-address (Protocols MSDP) " on page 9910	"local-address (Protocols PIM)" on page 9911
"local-address (Routing Options)" on page 9913	"local-as" on page 9916	"local-authentication (Dynamic PPP Options)" on page 9919
"local-bias" on page 9922	"local-bias (ae)" on page 9921	"local-bias (ae load-balance)" on page 9923
"local-bw-override-threshold" on page 9925	"local-category (Next Gen Services Service-Set Local System Logging)" on page 9926	"local-certificate" on page 9929
"local-certificate (System)" on page 9930	"local-certificate (Security)" on page 9932	"local-chap" on page 9933
"local-decision (PCRF Partition)" on page 9934	"local-dump (Flow Aggregation)" on page 9938	"local-endpoint" on page 9939
"local-engine" on page 9941	"local-gateway (IPSec) " on page 9944	"local-gateway (L2TP LNS)" on page 9945

"local-id" on page 9947	"local-id (High Availability)" on page 9948	"local-identity" on page 9950
"local-interface (IPv6)" on page 9952	"local-ip-addr (ICCP)" on page 9953	"local-ipv4-address" on page 9955
"local-ipv6-address" on page 9956	"local-key-pair" on page 9958	"local-log-tag (Next Gen Services Service-Set System Logging)" on page 9959
"local-loopback" on page 9960	"local-minimum-links-threshold" on page 9962	"local-name" on page 9964
"local-password" on page 9966	"local-port-range" on page 9967	"local-ports" on page 9969
"local-preference" on page 9971	"local-priority" on page 9973	"local-replication" on page 9974
"local-server-group (DHCP Relay Agent Option)" on page 9976	"local-switching (Layer 2 Circuits)" on page 9977	"local-switching (VPLS)" on page 9979
"local-traffic-selector" on page 9980	"locality-bias (MX Series Virtual Chassis)" on page 9981	"localized-fib" on page 9983
"location (DHCP Relay Agent)" on page 9985	"location (DHCP Snooping Database)" on page 9984	"location (DHCP Snooping Database)" on page 9987
"location (DHCPv6 Relay Agent)" on page 9988	"location (LLDP-MED)" on page 9990	"location (SNMP)" on page 9992
"location (System)" on page 9993	"location (Virtual Chassis)" on page 9995	"locator" on page 9997
"lockout" on page 9999	"lockout-timeout (L2TP Destination Lockout)" on page 10001	"log" on page 10002
"log (aacl)" on page 10005	"log (Logical Systems and Tenant Systems)" on page 10016	"log (Security)" on page 10026

"log (Security GTP)" on page 10032	"log (Security IDP Sensor Configuration)" on page 10033	"log (Security Policies)" on page 10035
"log (Security SCTP)" on page 10037	"log (Services)" on page 10038	"log (Spanning Trees)" on page 10040
"log-and-generate-ais" on page 10006	"log-attacks" on page 10050	"log-attacks" on page 10007
"log-attacks (rulebase-exempt)" on page 10052	"log-interval (IGMP Interface)" on page 10009	"log-interval (MLD Interface)" on page 10011
"log-interval (PIM Entries)" on page 10012	"log-interval (Protocols MSDP)" on page 10014	"log-prefix (Interfaces)" on page 10020
"log-prefix (L2TP)" on page 10021	"log-prefix (Services)" on page 10022	"log-prefix (System)" on page 10023
"log-rotate-frequency" on page 10025	"log-updown" on page 10042	"log-updown (Protocols LDP)" on page 10043
"log-updown (Protocols MPLS)" on page 10045	"log-warning (Multicast Forwarding Cache)" on page 10047	"log-warning (Protocols MSDP)" on page 10049
"logging (IDS MS-DPC)" on page 10053	"logging (Services)" on page 10056	"logging-rule (PCC Action Profile)" on page 10054
"logical-bandwidth-policer" on page 10058	"logical-domain-identity-management" on page 10059	"logical-interface (DDoS Flow Detection)" on page 10064
"logical-interface-aggregate-statistics" on page 10062	"logical-interface-chassis-redundancy (MX Series Virtual Chassis)" on page 10063	"logical-interface-fpc-redundancy (Aggregated Ethernet Subscriber Interfaces)" on page 10067
"logical-interface-fpc-redundancy (PS Interface)" on page 10068	"logical-interface-policer" on page 10069	"logical-interface-unit-range" on page 10072
"logical-system" on page 10073	"logical-system (Diameter Peer)" on page 10076	"logical-system (Diameter Transport)" on page 10078

"logical-system (SNMP)" on page 10075	"logical-system (System Security Profile)" on page 10086	"logical-system (Tunnel Profile)" on page 10089
"logical-system-mux" on page 10079	"logical-system-name (DHCP Local Server)" on page 10082	"logical-system-name (DHCP Relay Agent)" on page 10083
"logical-system-name (Static Subscribers)" on page 10085	"logical-system-trap-filter" on page 10088	"logical-systems" on page 10090
"logical-systems (All)" on page 10092	"logical-tunnel-options" on page 10093	"login" on page 10095
"login (Access)" on page 10100	"logout-response-timeout (PCRF Partition)" on page 10101	"long-lived (Graceful Restart for BGP Helper)" on page 10103
"long-lived (Graceful Restart for BGP Restarter)" on page 10105	"longest-match" on page 10108	"longest-match (Static Routes)" on page 10109
"lookup-intact-on-commit" on page 10112	"loop-detect" on page 10113	"loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)" on page 10117
"loopback (Fibre Channel Interface)" on page 10120	"loopback (Local and Remote)" on page 10125	"loopback (Serial)" on page 10129
"loopback-clear-timer" on page 10119	"loopback-firewall-optimization" on page 10121	"loopback-firewallv6-optimization" on page 10123
"loopback-port" on page 10126	"loopback-remote" on page 10128	"loopback-tracking" on page 10131
"loops (Autonomous System)" on page 10132	"loops (BGP Address Family)" on page 10134	"loose-authentication-check" on page 10137
"loose-check" on page 10138	"loose-source-route-option (IDS Screen Next Gen Services)" on page 10139	"loss (querier)" on page 10171

"loss (responder)" on page 10173	"loss-delay (querier)" on page 10140	"loss-priority" on page 10142
"loss-priority (Classifiers)" on page 10143	"loss-priority (Classifiers)" on page 10145	"loss-priority (Classifiers and Rewrite Rules)" on page 10146
"loss-priority (Drop Profiles)" on page 10150	"loss-priority (Dynamic Schedulers)" on page 10151	"loss-priority (Firewall Filter)" on page 10152
"loss-priority (Firewall Filter Action)" on page 10154	"loss-priority (Frame Relay Loss Priority)" on page 10148	"loss-priority high then discard (Three-Color Policer)" on page 10155
"loss-priority (Rewrite Rules)" on page 10163	"loss-priority (Rewrite Rules)" on page 10164	"loss-priority (Scheduler Drop Profiles)" on page 10169
"loss-priority (Simple Firewall Filter)" on page 10170	"loss-priority-maps" on page 10157	"loss-priority-maps (Assigning to an Interface)" on page 10159
"loss-priority-maps (CoS)" on page 10160	"loss-priority-maps (CoS Interfaces)" on page 10161	"loss-priority-rewrites" on page 10166
"loss-priority-rewrites (Assigning to an Interface)" on page 10167	"loss-threshold" on page 10175	"low-latency" on page 10176
"low-watermark" on page 10177	"lowest-priority-defect" on page 10179	"lrf-profile (Service Set)" on page 10180
"lsa-refresh-interval" on page 10182	"lsp-attributes" on page 10183	"lsp-cleanup-timer" on page 10185
"lsp-external-controller" on page 10187	"lsp-lifetime" on page 10189	"lsp-next-hop (CoS-Based Forwarding)" on page 10191
"lsp-next-hop (Static Routes)" on page 10192	"lsp-switch" on page 10194	"lsping-channel-type" on page 10195
"lsq-failure-options" on page 10197	"ltv-syslog-interval (System Process)" on page 10198	

M

"mac" on page 10215	"mac (MPLS)" on page 10217	"mac (Option 82)" on page 10250
"mac (Static MAC-Based VLANs)" on page 10261	"mac-address" on page 10218	"mac-address (Accept Source Mac)" on page 10219
"mac-address (DHCP Local Server)" on page 10221	"mac-address (DHCP Relay Agent)" on page 10223	"mac-address (Dynamic Access-Internal Routes)" on page 10224
"mac-address (MACsec)" on page 10226	"mac-address (VLAN and Stacked VLAN Interfaces)" on page 10227	"mac-flush" on page 10229
"mac-ip-table-size" on page 10231	"mac-learn-enable" on page 10233	"mac-learning-priority" on page 10234
"mac-limit" on page 10236	"mac-limit (Access Port Security)" on page 10241	"mac-limit (VLAN)" on page 10238
"mac-list" on page 10243	"mac-lookup-length" on page 10244	"mac-move-limit" on page 10246
"mac-notification" on page 10249	"mac-persistence-timer" on page 10252	"mac-pinning" on page 10254
"mac-pinning (EVPN Routing Instances)" on page 10256	"mac-radius" on page 10257	"mac-rewrite" on page 10259
"mac-statistics" on page 10263	"mac-table-aging-time" on page 10267	"mac-table-aging-time (MX Series)" on page 10265
"mac-table-size" on page 10271	"mac-table-size (VPLS)" on page 10269	"mac-validate" on page 10273
"mac-validate (Dynamic IP Demux Interface)" on page 10275	"mac-vrf" on page 10276	"macsec" on page 10278

"macsec (MX Series)" on page 10281	"maintain-subscriber (Subscriber Management)" on page 10284	"maintenance-association" on page 10285
"maintenance-domain" on page 10287	"major" on page 10290	"major-event (Process)" on page 10291
"major-event (System)" on page 10293	"major-ring-name" on page 10295	"make-before-break (LDP)" on page 10296
"managed" on page 10298	"managed-configuration" on page 10301	"managed-configuration (Dynamic Router Advertisement)" on page 10302
"managed-services" on page 10304	"management" on page 10305	"management (PoE)" on page 10310
"management (Security Zones)" on page 10311	"management-ethernet (Alarm)" on page 10307	"management-instance" on page 10309
"mandatory" on page 10313	"manual" on page 10315	"manual (Junos OS)" on page 10318
"manual (Junos-FIPS Software)" on page 10316	"manual (Security IPsec)" on page 10320	"manual switch" on page 10323
"many-to-one (Aggregated Multiservices)" on page 10324	"map (Domain Map)" on page 10327	"map-dest-bmac-to-dest-cmac" on page 10326
"map-e" on page 10329	"map-entry-timeout" on page 10333	"map-t" on page 10334
"mapped-port (Security Static NAT)" on page 10336	"mapping" on page 10338	"mapping (VNF)" on page 10340
"mapping-agent-election" on page 10342	"mapping-range" on page 10344	"mapping-refresh" on page 10345
"mapping-server-entry" on page 10346	"mapping-timeout" on page 10348	"mapping-timeout (Source NAT Next Gen Services)" on page 10349

"mapping-type" on page 10351	"mapping-type (Source NAT Next Gen Services)" on page 10352	"mark-interface (RA Guard)" on page 10353
"mark-local-ip-host-routes" on page 10355	"martians" on page 10356	"mask uli" on page 10358
"master" on page 10360	"master-only" on page 10364	"master-password" on page 10365
"mastership-priority" on page 10367	"match (Forwarding-Options)" on page 10382	"match (Next Gen Services Global System Logging)" on page 10386
"match (Security Destination NAT)" on page 10393	"match (Security IDP Policy)" on page 10395	"match (Security Policies)" on page 10396
"match (Security Policies Global)" on page 10398	"match (Security Source NAT)" on page 10401	"match (Security Static NAT)" on page 10402
"match (Services CoS Next Gen Services)" on page 10404	"match (Stateful Firewall Rule Next Gen Services)" on page 10406	"match (Syslog)" on page 10369
"match-direction" on page 10371	"match-direction (IDS)" on page 10373	"match-direction (IP Reassembly Rule)" on page 10375
"match-direction (ipsec-vpn)" on page 10372	"match-direction (NAT Next Gen Services)" on page 10376	"match-direction (PCP)" on page 10377
"match-direction (Services CoS) " on page 10378	"match-direction (Services NAT) " on page 10380	"match-direction (Services Stateful Firewall) " on page 10381
"match-list" on page 10383	"match-option" on page 10387	"match-order (Dynamic Firewalls)" on page 10389
"match-rules-on-reverse-flow" on page 10391	"match-rules-on-reverse-flow (Next Gen Services)" on page 10392	"match-strings" on page 10407

"matches" on page 10409	"matches (Event Policy)" on page 10412	"max-advertisement-interval (Dynamic Router Advertisement)" on page 10413
"max-advertisement-interval (Protocols IPv6 Neighbor Discovery)" on page 10415	"max-advertisement-interval (Protocols Router Discovery)" on page 10416	"max-age" on page 10418
"max-areas" on page 10420	"max-bypasses" on page 10422	"max-configurations-on-flash (System)" on page 10423
"max-connection-duration" on page 10424	"max-connections" on page 10426	"max-data-sessions-per-subscriber" on page 10427
"max-datasize" on page 10429	"max-db-size (Enhanced Subscriber Management)" on page 10432	"max-drop-flows" on page 10435
"max-duplicates" on page 10436	"max-failures" on page 10438	"max-flows" on page 10439
"max-hops" on page 10441	"max-login-sessions" on page 10442	"max-login-sessions-per-node" on page 10444
"max-lsp-size" on page 10446	"max-message-length" on page 10448	"max-outstanding-requests (Diameter Applications)" on page 10449
"max-packet-memory-ratio" on page 10451	"max-packets-per-second (Sampling)" on page 10452	"max-pending-accounting-stops (Access Profile)" on page 10454
"max-policies (Event Policy)" on page 10455	"max-power" on page 10456	"max-queues" on page 10458
"max-queues-per-interface" on page 10461	"max-reass-packet-memory-ratio" on page 10463	"max-retrans-time" on page 10464
"max-session-number" on page 10466	"max-session-setup-rate (Service Set)" on page 10467	"max-sessions (Dynamic PPPoE)" on page 10469

"max-sessions (PPPoE Service Name Tables)" on page 10477	"max-sessions (Security Packet Log)" on page 10478	"max-sessions (Tunnel Profile)" on page 10480
"max-sessions-per-connection" on page 10471	"max-sessions-per-enode" on page 10472	"max-sessions-per-subscriber" on page 10474
"max-sessions-per-subscriber (Service Set Next Gen Services)" on page 10475	"max-sessions-vsa-ignore (Static and Dynamic Subscribers)" on page 10481	"max-snp-size" on page 10483
"max-tcp-session-packet-memory" on page 10484	"max-time-report" on page 10486	"max-transmit-quality-level" on page 10487
"max-udp-session-packet-memory" on page 10489	"max-unknown-messages" on page 10491	"max-unknown-requests" on page 10492
"max-withhold-time (Access Profile)" on page 10494	"maximize-idp-sessions" on page 10495	"maximum" on page 10497
"maximum (MSDP Active Source Messages) " on page 10536	"maximum (PIM Entries)" on page 10547	"maximum-age" on page 10498
"maximum-allowed-contentions " on page 10499	"maximum-backoff-duration" on page 10501	"maximum-bandwidth" on page 10502
"maximum-bandwidth (Protocols MPLS)" on page 10504	"maximum-bit-rate (Default Local Policy)" on page 10505	"maximum-bit-rate (PCC Action Profiles)" on page 10507
"maximum-bit-rate (TDF Domain)" on page 10509	"maximum-call-duration (Security)" on page 10510	"maximum-capture-size (Datapath Debug)" on page 10512
"maximum-certificates" on page 10513	"maximum-connections" on page 10514	"maximum-connections-per-client" on page 10516
"maximum-contexts" on page 10517	"maximum-discovery-table-entries" on page 10519	"maximum-ecmp" on page 10520

"maximum-helper-recovery-time" on page 10522	"maximum-helper-restart-time" on page 10524	"maximum-helper-restart-time (RSVP)" on page 10525
"maximum-hop-count" on page 10527	"maximum-labels" on page 10528	"maximum-lease-time (DHCP)" on page 10530
"maximum-length (Origin Validation for BGP)" on page 10531	"maximum-links" on page 10532	"maximum-message-length" on page 10534
"maximum-neighbor-reconnect-time" on page 10537	"maximum-neighbor-recovery-time" on page 10538	"maximum-packet-length" on page 10540
"maximum-paths" on page 10542	"maximum-pending-reqs-limit" on page 10544	"maximum-pending-requests (Diameter)" on page 10546
"maximum-power (Interface)" on page 10549	"maximum-power (PoE)" on page 10552	"maximum-prefixes" on page 10554
"maximum-rps" on page 10556	"maximum-send-window" on page 10558	"maximum-sessions" on page 10559
"maximum-sessions (L2TP)" on page 10560	"maximum-sessions (TDF Gateway)" on page 10565	"maximum-sessions-per-connection" on page 10562
"maximum-sessions-per-tunnel" on page 10563	"maximum-sessions-trap-percentage (TDF Gateway)" on page 10566	"maximum-subscribers" on page 10568
"maximum-subscribers" on page 10569	"maximum-transactions" on page 10570	"maximum-transmit-rate (Protocols IGMP)" on page 10572
"maximum-transmit-rate (Protocols MLD)" on page 10574	"mc-ae" on page 10575	"mc-ae-id" on page 10581
"mcast stats-enable" on page 10582	"mclag" on page 10584	"mclag-arp-nd-sync" on page 10585
"mclag-arpreq-sync" on page 10586	"mdi-mode" on page 10587	"mdt" on page 10590

"measurement-interval" on page 10591	"measurement-type" on page 10593	"med-igp-update-interval" on page 10594
"media-loss-rate" on page 10596	"media-rate-variation" on page 10597	"media-source-port-any" on page 10599
"media-type (Dual-Purpose Uplink Ports)" on page 10600	"media-type (Interfaces)" on page 10602	"medium (Tunnel Profile)" on page 10603
"member" on page 10604	"member (Application Identification)" on page 10606	"member (MX Series Virtual Chassis)" on page 10615
"member (NSSU Upgrade Groups)" on page 10617	"member (Redundant Power System)" on page 10620	"member (Security Group VPN)" on page 10621
"member (Security IDP)" on page 10624	"member-failure-options (Aggregated Multiservices)" on page 10608	"member-id (Junos Fusion Satellite Device Cluster)" on page 10611
"member-interface (Aggregated Multiservices)" on page 10612	"member-link-scheduler" on page 10614	"member-range" on page 10618
"members" on page 10625	"memory" on page 10628	"memory (System)" on page 10630
"memory (TDF Gateway)" on page 10633	"memory-enhanced" on page 10629	"menu-item" on page 10634
"mep" on page 10639	"mesh-group (Protocols VPLS)" on page 10641	"message-flood (Security H323)" on page 10644
"message-flood (Security MGCP)" on page 10645	"message-ie-profile-v1" on page 10647	"message-ie-profile-v2" on page 10648
"message-list" on page 10650	"message-processing-model" on page 10652	"message-rate-limit" on page 10655

"message-rate-limit (Flow Monitoring Logs for NAT)" on page 10656	"message-rate-limit (interfaces)" on page 10653	"message-size" on page 10658
"message-type" on page 10659	"saml (Access Profile)" on page 10662	"method" on page 10663
"metric" on page 10665	"metric (Aggregate, Generated, or Static Route)" on page 10667	"metric (Diameter Route)" on page 10668
"metric (Dynamic Access-Internal Routes)" on page 10670	"metric (Policy Action)" on page 10681	"metric (Protocols DVMRP)" on page 10682
"metric (Protocols IS-IS)" on page 10684	"metric (Protocols MPLS)" on page 10686	"metric (Protocols OSPF Sham Link)" on page 10688
"metric (Qualified Next Hop on Static Route)" on page 10689	"metric-in (Protocols RIP)" on page 10671	"metric-in (Protocols RIPng)" on page 10673
"metric-out" on page 10676	"metric-out (Protocols RIP)" on page 10674	"metric-out (Protocols RIPng)" on page 10679
"metrics-profile " on page 10691	"mgcp" on page 10693	"mib-profile" on page 10695
"micro-sid" on page 10697	"microloop-avoidance" on page 10700	"microloop-avoidance (Protocols OSPF)" on page 10702
"mif (TDF Interface)" on page 10703	"mime-pattern" on page 10705	"mime-whitelist" on page 10706
"min-advertisement-interval (Dynamic Router Advertisement)" on page 10708	"min-advertisement-interval (Protocols IPv6 Neighbor Discovery)" on page 10709	"min-advertisement-interval (Protocols Router Discovery)" on page 10711
"min-message-length" on page 10713	"min-rate" on page 10714	"min-rate (source-active-advertisement)" on page 10718

"mine-dhcp-client-options" on page 10719	"mine-dhcpv6-client-options" on page 10721	"minimum-bandwidth" on page 10722
"minimum-bandwidth (aggregated Ethernet)" on page 10728	"minimum-bandwidth-adjust-interval" on page 10724	"minimum-bandwidth-adjust-threshold-change" on page 10725
"minimum-bandwidth-adjust-threshold-value" on page 10726	"minimum-effective-aigp" on page 10730	"minimum-hold-time" on page 10731
"minimum-interval" on page 10733	"minimum-interval (BFD Liveness Detection)" on page 10735	"minimum-interval (Liveness Detection)" on page 10737
"minimum-interval (PIM BFD Liveness Detection)" on page 10738	"minimum-interval (PIM BFD Transmit Interval) " on page 10739	"minimum-links" on page 10741
"minimum-links (candidate-uplink-port-profile)" on page 10744	"minimum-links (Interfaces)" on page 10746	"minimum-priority" on page 10747
"minimum-quality" on page 10748	"minimum-receive-interval" on page 10752	"minimum-receive-interval (BFD Liveness Detection)" on page 10754
"minimum-receive-interval (Liveness Detection)" on page 10756	"minimum-receive-interval (Protocols PIM)" on page 10751	"minimum-retransmission-timeout (L2TP Tunnel)" on page 10757
"minimum-wait-time" on page 10758	"minor" on page 10760	"minor-event (Process)" on page 10761
"minor-event (System)" on page 10763	"mip-half-function" on page 10765	"mirror-decrypt-traffic" on page 10766
"mirror-filter (Security Forwarding Options)" on page 10768	"mirror-flash-on-disk" on page 10771	"mirror-once" on page 10772
"mixed-rate-mode" on page 10774	"mka" on page 10775	"mka (MX Series)" on page 10777
"mld" on page 10778	"mld (Dynamic Profiles)" on page 10781	"mld-snooping" on page 10783

"mldp-inband-signalling (Protocols Multipoint LDP)" on page 10787	"mlfr-uni-nni-bundle-options" on page 10789	"mlfr-uni-nni-bundles-inline" on page 10790
"mobile-edge session-scale" on page 10792	"mobility" on page 10793	"mode" on page 10798
"mode (Dynamic Profiles)" on page 10801	"mode (ike)" on page 10797	"mode (IKE)" on page 10802
"mode (Interfaces)" on page 10804	"mode (IPsec)" on page 10805	"mode (Traffic Manager)" on page 10807
"mode (Multicast VLAN Registration)" on page 10809	"mode (Next Gen Services Service-Set System Logging)" on page 10811	"mode (Protocols DVMRP)" on page 10813
"mode (Protocols MSDP)" on page 10814	"mode (Protocols PIM)" on page 10816	"mode (QFX Series)" on page 10818
"mode (RFC 2544 Benchmarking Tests)" on page 10795	"mode (Security Forwarding Options)" on page 10819	"mode (Security Group VPN)" on page 10821
"mode (Security Log)" on page 10823	"modem-options" on page 10825	"modulation-format" on page 10826
"mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain)" on page 10827	"mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain)" on page 10829	"mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain)" on page 10830
"mofrr-primary-path-selection-by-routing (Multicast-Only Fast Reroute)" on page 10832	"monitor (Multinode High Availability)" on page 10903	"monitor-bandwidth" on page 10834
"monitor-end-point" on page 10835	"monitor-object" on page 10906	"monitor-profile" on page 10912
"monitor-session" on page 10926	"monitor-timer" on page 10953	"monitoring" on page 10954
"monitoring (Forwarding Options)" on page 10956	"monitoring (Services)" on page 10961	"monitoring-key (PCC Action Profile)" on page 10957

"monitoring-profile" on page 10959	"most-fill" on page 10967	"moving-average-size" on page 10967
"mpls" on page 10969	"mpls (security)" on page 10974	"mpls (packet-forwarding-options)" on page 10982
"mpls (Protocols)" on page 10984	"mpls (Security Forwarding Options)" on page 10985	"mpls-flow-table-size " on page 10972
"mpls-inet6" on page 10977	"mpls-internet-multicast" on page 10978	"mpls-ipv4-template" on page 10979
"mpls-ipvx-template" on page 10981	"mpls-template" on page 10986	"mpls-tp-mode" on page 10988
"mpls (forwarding-options)" on page 10975	"mrru" on page 10989	"mru (Congestion Notification Profiles)" on page 10992
"mru (Dynamic and Static PPPoE)" on page 10994	"mru (Gigabit Ethernet Options)" on page 10991	"msdp" on page 10996
"msrpc" on page 10999	"mss (IDS MS-DPC)" on page 11001	"mss (Security IDP)" on page 11002
"msti" on page 11003	"mstp" on page 11006	"mtu" on page 11026
"mtu (Dynamic and Static PPPoE)" on page 11034	"mtu (Dynamic Profiles)" on page 11035	"mtu (interfaces)" on page 11027
"mtu (Multilink and Link Services Logical Interface)" on page 11025	"mtu (QFX Series)" on page 11023	"mtu (TDF Interface)" on page 11038
"mtu-discovery" on page 11032	"mtu-signaling" on page 11037	"multi-address-embedded-option-response (DHCP Local Server)" on page 11039
"multi-bgp-path" on page 11041	"multi-bgp-path" on page 11041	"multi-chassis" on page 11042

"multi-chassis-protection" on page 11044	"multi-chassis-protection (Interfaces)" on page 11045	"multi-destination (Class of Service)" on page 11047
"multi-domain" on page 11048	"multi-homing (VPLS Multihoming for BGP VPLS)" on page 11051	"multi-homing (VPLS Multihoming for FEC 129)" on page 11052
"multi-link-layer-2-inline" on page 11054	"multi-sa" on page 11055	"multi-tenant-hash" on page 11057
"multi-tenant-support" on page 11059	"multicast" on page 11062	"multicast (Dynamic Routing Options)" on page 11067
"multicast (Protocols Router Discovery)" on page 11077	"multicast (QFabric Routing Options)" on page 11078	"multicast (routing-options)" on page 11063
"multicast (Virtual Tunnel in Routing Instances)" on page 11092	"multicast-dlci" on page 11066	"multicast-group" on page 11069
"multicast-ipv4" on page 11070	"multicast-loadbalance" on page 11072	"multicast-mode (EVPN)" on page 11074
"multicast-only" on page 11076	"multicast-replication" on page 11080	"multicast-router-interface (IGMP Snooping)" on page 11084
"multicast-router-interface (MLD Snooping)" on page 11086	"multicast-snooping-options" on page 11087	"multicast-statistics" on page 11089
"multicast-statistics (packet-forwarding-options)" on page 11091	"multicast-vlan-flood" on page 11094	"multichassis-lag-replicate-state" on page 11095
"multichunk-inspection" on page 11097	"interconnect-multihoming-peer-gateways" on page 8806	"multihop" on page 11098
"multilink" on page 11103	"multilink-class" on page 11104	"multilink-max-classes" on page 11106

"multipath (OSPF SPF Options)" on page 11109	"multipath (Protocols BGP)" on page 11110	"multipath (Routing Options)" on page 11113
"multipath (SPF Options)" on page 11117	"multipath-build-priority" on page 11107	"multipath-rule" on page 11115
"multiple-triggered-joins" on page 11119	"multiplier" on page 11122	"multiplier (BFD Liveness Detection)" on page 11123
"multiplier (Liveness Detection)" on page 11125	"multiplier (Protocols PIM)" on page 11120	"multipoint" on page 11127
"multiservice" on page 11128	"multiservice-options" on page 11131	"multiservice-options (Flow Monitoring)" on page 11130
"must-secure" on page 11133	"mutual-authentication (REST API)" on page 11134	"mvpn" on page 11135
"mvpn (Draft-Rosen MVPN)" on page 11138	"mvpn (NG-MVPN)" on page 11143	"mvpn-iana-rt-import" on page 11140
"mvpn-mode" on page 11142	"mvrp" on page 11145	

N

"n391" on page 11163	"n392" on page 11164	"n393" on page 11166
"n-plus-n (Power Management)" on page 11160	"n-plus-n (satellite-management)" on page 11162	"name (early-aging)" on page 11168
"name (L2TP Destination)" on page 11173	"name (L2TP Tunnel Destination)" on page 11175	"name (Next Gen Services Global System Logging)" on page 11176

"name (SNMP)" on page 11167	"name (Transport-class)" on page 11181	"name-format" on page 11172
"name-format (Flow Collection)" on page 11170	"name-server (Access)" on page 11178	"name-server (System Services)" on page 11179
"nas-ip-address" on page 11182	"nas-port-extended-format" on page 11184	"nas-port-extended-format (Interfaces)" on page 11186
"nas-port-id-format (Subscriber Management)" on page 11188	"nas-port-method (L2TP LAC)" on page 11191	"nas-port-method (Tunnel Profile)" on page 11192
"nas-port-options (RADIUS Options)" on page 11193	"nas-port-type" on page 11195	"nas-port-type (RADIUS Options)" on page 11197
"nas-port-type (Subscriber Management)" on page 11199	"nasreq (Diameter Application)" on page 11202	"nat" on page 11205
"nat-cone-binding" on page 11210	"nat-destination-pool" on page 11212	"nat-destination-rule" on page 11214
"nat-interface-port-ol (System)" on page 11216	"nat-keepalive (Services IPsec VPN)" on page 11218	"nat-nopat-address" on page 11219
"nat-options" on page 11221	"nat-options (Next Gen Services)" on page 11222	"nat-pat-address" on page 11223
"nat-pat-portnum" on page 11225	"nat-port-ol-ipnumber" on page 11227	"nat-rule-referenced-prefix (System)" on page 11229
"nat-rule-sets (Service Set)" on page 11230	"nat-rule-sets (Service Set Next Gen Services)" on page 11231	"nat-rules" on page 11232
"nat-source-pool" on page 11234	"nat-source-rule" on page 11236	"nat-static-rule" on page 11238
"native-vlan-id" on page 11240	"natt-install-interval" on page 11242	"ncp-max-conf-req" on page 11243

"ncp-restart-timer" on page 11245	"nd6-max-cache" on page 11252	"nd6-new-hold-limit" on page 11254
"nd6-stale-time" on page 11256	"nd-maxucast-retry" on page 11246	"nd-override-preferred-src" on page 11248
"nd-retransmit-timer" on page 11249	"nd-system-cache-limit" on page 11251	"ndp" on page 11257
"ndp-proxy" on page 11259	"negate" on page 11260	"negotiate-address" on page 11261
"negotiation-options" on page 11263	"neighbor" on page 11264	"neighbor (Define ANCP)" on page 11266
"neighbor (Protocols BGP)" on page 11273	"neighbor (Protocols Layer 2 Circuit)" on page 11278	"neighbor (Protocols LDP)" on page 11281
"neighbor (Protocols RIPng)" on page 11282	"neighbor (Protocols VPLS)" on page 11284	"neighbor-discovery" on page 11268
"neighbor-discovery-inspection" on page 11270	"neighbor-policy" on page 11272	"nested-application (Security IDP)" on page 11286
"nested-application-unknown" on page 11287	"nested-applications" on page 11288	"netconf" on page 11290
"netconf (Remote Device Management)" on page 11295	"netconf-monitoring (NETCONF)" on page 11293	"network" on page 11298
"network (Access)" on page 11300	"network (Address Pools)" on page 11301	"network (TDF Domain)" on page 11328
"network-domain" on page 11302	"network-element (AAA Profile)" on page 11304	"network-element (Diameter Base Protocol)" on page 11305
"network-element (Diameter Base Protocol)" on page 11307	"network-element (Subscriber Aware Policy Control)" on page 11308	"network-elements (RADIUS)" on page 11310

"network-isolation" on page 11312	"network-isolation-profile" on page 11314	"network-management" on page 11317
"network-option" on page 11318	"network-services" on page 11320	"network-slices" on page 11322
"network-slicing" on page 11323	"network-summary-export" on page 11325	"network-summary-import" on page 11326
"next-hop (Access)" on page 11329	"next-hop (Access Internal)" on page 11330	"next-hop (Class-Of-Service)" on page 11333
"next-hop (Dynamic Access Routes)" on page 11334	"next-hop (Forwarding Options)" on page 11336	"next-hop (PIM RPF Selection)" on page 11347
"next-hop (Protocols MPLS)" on page 11349	"next-hop (RPM)" on page 11350	"next-hop (Static MAC-Based VLANs)" on page 11353
"next-hop (VXLAN Routing)" on page 11356	"next-hop-based-tunnel (Routing Options)" on page 11332	"next-hop-group" on page 11337
"next-hop-group (Analyzer)" on page 11339	"next-hop-group (Forwarding Options)" on page 11340	"next-hop-group (Port Mirroring)" on page 11343
"next-hop-group (Unknown Unicast Forwarding)" on page 11344	"next-hop-map" on page 11346	"next-hop-service" on page 11352
"next-hop-tunnel" on page 11355	"next-server" on page 11358	"next-table" on page 11359
"nextboot" on page 11361	"nexthop-hold-time" on page 11362	"nexthop-learning" on page 11363
"ng-juniper" on page 11366	"ngfw" on page 11369	"no advertise-peer-as" on page 11385
"no-accounting" on page 11371	"no-adaptation" on page 11373	"no-adaptation (BFD Liveness Detection)" on page 11374

"no-adaptation (Liveness Detection)" on page 11376	"no-adaptation (PIM BFD Liveness Detection)" on page 11377	"no-adaptive" on page 11378
"no-adjacency-down-notification (Protocols IS-IS)" on page 11380	"no-advertise-adjacency-segment (Protocols IS-IS)" on page 11382	"no-advertise-adjacency-segment (Protocols OSPF)" on page 11383
"no-aggregate-delegate-processing" on page 11386	"no-aggregator-id" on page 11387	"no-allow-link-events" on page 11389
"no-allow-snooped-clients" on page 11390	"no-allow-url (Op Scripts)" on page 11392	"no-allowed-mac-log" on page 11393
"no-anti-replay" on page 11395	"no-anti-replay (Services Service Set)" on page 11396	"no-application-identification (Services)" on page 11397
"no-application-system-cache" on page 11399	"no-application-system-cache (Services)" on page 11400	"no-attribute-length-in-pdu" on page 11401
"no-authentication-check (Protocols RSVP)" on page 11403	"no-auto-conversion" on page 11404	"no-auto-failover" on page 11406
"no-auto-iccp-provisioning (Junos Fusion Redundancy Group)" on page 11407	"no-auto-mdix" on page 11409	"no-auto-virtual-gateway-esi" on page 11410
"no-autostart" on page 11412	"no-autoupdate" on page 11413	"no-bfd-triggered-local-repair" on page 11414
"no-bidirectional-mode" on page 11416	"no-bind-on-request (DHCP Relay Agent)" on page 11418	"no-broadcast" on page 11420
"no-bundle-flap" on page 11422	"no-certificate-chain-in-ike" on page 11423	"no-client-reflect" on page 11424
"no-control-word " on page 11426	"no-control-word (Protocols Layer 2 VPN)" on page 11427	"no-cspf" on page 11429

"no-cspf (Protocols RSVP)" on page 11431	"no-decrement-ttl" on page 11432	"no-default-gateway-ext-comm" on page 11434
"no-delegate-processing" on page 11435	"no-dhcp-flood" on page 11627	"no-dhcp-snooping" on page 11436
"no-dhcp-trusted" on page 11438	"no-dhcpv6-options" on page 11440	"no-dhcpv6-snooping" on page 11441
"no-dns-propagation" on page 11442	"no-domain-vpn-tag" on page 11443	"no-dr-flood (PIM Snooping)" on page 11445
"no-dynamic-vlan" on page 11446	"no-encryption (MACsec)" on page 11447	"no-encryption (MACsec for MX Series)" on page 11449
"no-enhanced-frr-bypass (Protocols RSVP)" on page 11451	"no-examine-dhcpv6" on page 11452	"no-external-export" on page 11454
"no-fabric-wwn-verify" on page 11456	"no-fcoe-lag" on page 11457	"no-fcoe-trusted" on page 11458
"no-filter-check" on page 11460	"no-fip-snooping-scaling" on page 11461	"no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping)" on page 11628
"no-flow-logging (DDoS Flow Detection)" on page 11463	"no-forwarding" on page 11464	"no-fragmentation" on page 11466
"no-gratuitous-arp-request" on page 11467	"no-hidden-commands" on page 11469	"no-hierarchical-ecmp" on page 11470
"no-icmp-packet-too-big" on page 11472	"no-incremental-update" on page 11473	"no-incremental-update" on page 11473
"no-inner-payload" on page 11474	"no-install" on page 11475	"no-install-to-address" on page 11476

"no-intelligent-prescreening" on page 11477	"no-interface-hello" on page 11479	"no-ip-tos-rewrite" on page 11480
"no-ipsec-tunnel-in-traceroute" on page 11481	"no-ipv4-routing" on page 11483	"no-ipv6-routing" on page 11484
"no-irb-layer-2-copy" on page 11486	"no-issu-timer-negotiation" on page 11487	"no-keepalives" on page 11488
"no-keepalives (Dynamic Profiles)" on page 11490	"no-l2ckt" on page 11491	"no-l2vpn" on page 11492
"no-listen" on page 11494	"no-load-balance-label-capability" on page 11495	"no-load-throttle (Resource Monitor)" on page 11496
"no-local-reversion" on page 11497	"no-local-switching" on page 11499	"no-local-switching (VLAN)" on page 11500
"no-local-switching (VPLS)" on page 11502	"no-logging (Resource Monitor)" on page 11503	"no-mac-learning" on page 11504
"no-make-before-break" on page 11508	"no-management-vlan" on page 11509	"no-mcast-replication" on page 11510
"no-multicast" on page 11512	"no-multicast-echo" on page 11514	"no-nat-traversal (Services IPsec VPN)" on page 11515
"no-native-vlan-insert" on page 11517	"no-neighbor-down-notification" on page 11519	"no-neighbor-learn" on page 11520
"no-next-header-to-payload-protocol-mapping" on page 11522	"no-node-hello-on-bypass" on page 11523	"no-node-id-subobject" on page 11524
"no-normalization" on page 11526	"no-notify-mail-recipient" on page 11528	"no-notify-mail-sender (Security Content Filtering Notification Options)" on page 11529

"no-notify-mail-sender (Security Fallback Block)" on page 11530	"no-notify-mail-sender (Security Virus Detection)" on page 11531	"no-nssa-abr" on page 11533
"no-odu-backward-frr-enable" on page 11534	"no-odu-signal-degrade-monitor-enable" on page 11536	"no-option16" on page 11537
"no-option18" on page 11538	"no-option37" on page 11540	"no-option82" on page 11541
"no-overload-on-prefix-export-limit" on page 11542	"no-p2mp-sublsp" on page 11544	"no-per-unit-scheduler" on page 11545
"no-ping-record-route" on page 11546	"no-ping-time-stamp" on page 11547	"no-policy-cold-synchronization" on page 11548
"no-power-budget" on page 11550	"no-pre-classifier" on page 11551	"no-propagate-ttl" on page 11553
"no-qos-adjust" on page 11554	"no-qos-adjust (Dynamic Routing Options)" on page 11556	"no-recommendation-tlv" on page 11557
"no-recommended" on page 11558	"no-redirects" on page 11559	"no-redirects (IPv4 Traffic)" on page 11560
"no-registered-multicast" on page 11562	"no-remote-trace (Next Gen Services Global System Logging)" on page 11564	"no-remote-trace (System)" on page 11565
"no-remote-trace (Trace Options)" on page 11566	"no-revert (Local Switching)" on page 11568	"no-revert (Neighbor Interface)" on page 11569
"no-rfc-1583" on page 11571	"no-root-port" on page 11572	"no-route-localize" on page 11574
"no-sbl-default-server" on page 11575	"no-send-pads-ac-info" on page 11577	"no-send-pads-error" on page 11578
"no-send-to-ue" on page 11579	"no-sequence-check" on page 11580	"no-snoop (DHCP Local Server and Relay Agent)" on page 11582

"no-source-packet-routing (Protocols OSPF)" on page 11583	"no-split-detection" on page 11585	"no-split-detection (MX Series Virtual Chassis)" on page 11586
"no-syn-check" on page 11588	"no-syn-check-in-tunnel" on page 11589	"no-syslog" on page 11590
"no-syslog-generation" on page 11591	"no-tag" on page 11592	"no-tag" on page 11594
"no-temperature-monitoring" on page 11595	"no-termination-request" on page 11596	"no-throttle (Resource Monitor)" on page 11598
"no-tls-certificate-check" on page 11599	"no-transit-statistics" on page 11600	"no-translation" on page 11601
"no-translation (Source NAT Next Gen Services)" on page 11603	"no-trap" on page 11604	"no-tunnel-services" on page 11605
"no-unknown-unicast" on page 11607	"no-unregistered-multicast" on page 11609	"no-unsolicited-ra (Enhanced Subscriber Management)" on page 11611
"no-uri-check" on page 11612	"no-user-query (Services User Identification)" on page 11614	"no-validate" on page 11615
"no-vlan-interface-name" on page 11617	"no-vrf-advertise" on page 11620	"no-vrf-propagate-ttl" on page 11621
"no-walkup" on page 11623	"no-world-readable (Next Gen Services Global System Logging)" on page 11624	"no-zero-suppression" on page 11626
"node (Chassis Cluster Redundancy Group)" on page 11631	"node (System)" on page 11656	"node-detection (EVPN-VXLAN)" on page 11633
"node-device (Aliases)" on page 11635	"node-device (Chassis)" on page 11637	"node-device (Flow Groups)" on page 11638

"node-device (Resources)" on page 11640	"node-group (Chassis)" on page 11641	"node-group (OxID Hash Control)" on page 11643
"node-group (Resources)" on page 11645	"node-hello" on page 11646	"node-id" on page 11648
"node-link-protection (Protocols MPLS)" on page 11649	"node-protection" on page 11651	"node-protection (Static LSP)" on page 11652
"node-segment (Protocols OSPF)" on page 11654	"non-active-hold-priority" on page 11657	"non-lsp-next-hop" on page 11659
"non-revertive" on page 11660	"non-revertive (Chassis)" on page 11661	"non-revertive (Interfaces)" on page 11662
"non-strict-priority-scheduling" on page 11664	"non-subscriber-no-reply" on page 11665	"non-vc-mode" on page 11666
"nonpersistent" on page 11668	"nonstop-bridging" on page 11669	"nonstop-bridging (Ethernet Switching)" on page 11670
"nonstop-routing" on page 11671	"nonstop-routing-options" on page 11673	"nonvolatile" on page 11675
"normalization" on page 11676	"normalization (Protocols EVPN)" on page 11678	"not (Event Policy)" on page 11680
"not-on-disk-underperform" on page 11681	"not-present (DHCP Relay Agent)" on page 11682	"notification" on page 11686
"notification (Commit)" on page 11687	"notification (NETCONF)" on page 11691	"notification (System Services)" on page 11697
"notification-control" on page 11688	"notification-interval" on page 11690	"notification-options (Security Antivirus)" on page 11693

"notification-options (Security Content Filtering)" on page 11695	"notification-targets" on page 11699	"notification-time" on page 11700
"notify" on page 11701	"notify-filter (Applying to the Management Target)" on page 11703	"notify-filter (Configuring the Profile Name)" on page 11704
"notify-mail-recipient" on page 11705	"notify-mail-sender (Security Content Filtering Notification Options)" on page 11707	"notify-mail-sender (Security Fallback Block)" on page 11708
"notify-mail-sender (Security Virus Detection)" on page 11709	"notify-view" on page 11710	"np-cache (Flexible PIC Concentrator)" on page 11711
"nptv6-prefix" on page 11713	"nptv6-prefix-name" on page 11714	"nsr-phantom-holdtime" on page 11715
"nssa" on page 11717	"nssu" on page 11719	"ntp" on page 11721
"ntp (QFabric)" on page 11729	"ntp threshold" on page 11731	"nullpdu" on page 11733
"num-65-127-prefix" on page 11734	"number-of-frames" on page 11736	"number-of-ports" on page 11737
"number-of-sub-ports (interface)" on page 11740		

O

"oam" on page 11753	"oam detection-disable" on page 11756	"oam (Layer 2 Circuit Protocols)" on page 11748
"oam (MX Series)" on page 11749	"oam (Protocols LDP)" on page 11759	"oam (Protocols MPLS)" on page 11762

"oam runtime-disable" on page 11764	"oam-on-svlan (Ethernet Interfaces)" on page 11758	"object-cache-size" on page 11765
"object-names" on page 11767	"observation-domain-id" on page 11768	"oc192" on page 11770
"ocs (Diameter Applications)" on page 11771	"ocsp (Security PKI)" on page 11774	"odu-backward-frr-enable" on page 11776
"odu-delay-management" on page 11777	"odu-signal-degrade" on page 11779	"odu-signal-degrade-monitor-enable" on page 11781
"odu-ttim-action-enable" on page 11782	"off (Destination NAT Next Gen Services)" on page 11784	"offer-period" on page 11785
"offline-on-fabric-bandwidth-reduction" on page 11786	"offload-type" on page 11788	"offloads" on page 11789
"offset" on page 11791	"offset (MX Series)" on page 11793	"oid" on page 11796
"oid (SNMP Scripts)" on page 11797	"oid (SNMPv3)" on page 11799	"oif-map (Dynamic IGMP Interface)" on page 11800
"oif-map (Dynamic MLD Interface)" on page 11801	"oif-map (IGMP Interface) " on page 11803	"oif-map (MLD Interface) " on page 11804
"oism" on page 11805	"oism (Multicast Snooping Options)" on page 11807	"omit-no-export (Graceful Restart for BGP Helper)" on page 11810
"omit-wildcard-address" on page 11812	"on-demand-address-allocation" on page 11813	"on-demand-ip-address" on page 11815
"on-disk-failure" on page 11817	"on-disk-failure (Chassis Routing Engine)" on page 11818	"on-error" on page 11819
"on-link" on page 11821	"on-link (Dynamic Router Advertisement)" on page 11823	"on-loss-of-keepalives" on page 11824

"on-loss-of-keepalives" on page 11826	"on-re-to-fpc-stale" on page 11827	"one-to-one (Aggregated Multiservices)" on page 11829
"one-way-hardware-timestamp" on page 11830	"onlink-subnet-only" on page 11832	"op (Op Scripts)" on page 11838
"open-timeout" on page 11841	"openflow (Protocols OpenFlow)" on page 11842	"operating-mode" on page 11844
"operation" on page 11846	"optics-options" on page 11847	"optimal-route-reflection" on page 11850
"optimize-aggregate-frr" on page 11853	"optimize-adaptive-teardown" on page 11852	"optimize-aggressive" on page 11855
"optimize-hold-dead-delay" on page 11856	"optimize-switchover-delay" on page 11857	"optimize-timer (Protocols MPLS)" on page 11859
"optimize-timer (Protocols RSVP)" on page 11861	"option (DHCP server)" on page 11895	"option (Security IDP)" on page 11905
"option-15 (DHCP Relay Agent)" on page 11862	"option-16 (DHCP Relay Agent)" on page 11865	"option-16 (DHCPv6 Snooping)" on page 11868
"option-18 (DHCPv6 Snooping)" on page 11869	"option-18 (Interface-ID for DHCPv6 Autosense VLANs)" on page 11871	"option-37 (DHCPv6 Snooping)" on page 11872
"option-37 (Relay Agent Remote-ID for DHCPv6 Autosense VLANs)" on page 11874	"option-60 (DHCP Local Server)" on page 11876	"option-60 (DHCP Relay Agent)" on page 11877
"option-77 (DHCP Relay Agent)" on page 11880	"option-79 " on page 11883	"option-82" on page 11884
"option-82 (Address-Assignment Pools)" on page 11888	"option-82 (DHCP Local Server Authentication)" on page 11890	"option-82 (DHCP Local Server Pool Matching)" on page 11892

"option-82 (DHCP Relay Agent)" on page 11893	"option-82 (interfaces)" on page 11886	"option-match" on page 11896
"option-number (DHCP Relay Agent Option)" on page 11898	"option-order (DHCP Relay Agent)" on page 11899	"option-refresh-rate" on page 11902
"option-refresh-rate (Services)" on page 11903	"option-type" on page 11906	"optional (Commit Scripts)" on page 11907
"optional (Event Scripts)" on page 11909	"optional-parameters" on page 11910	"options (Access Profile)" on page 11912
"options (Routing Options)" on page 11922	"options-template-id" on page 11924	"order" on page 11926
"order (Application Identification)" on page 11927	"order (Security IDP)" on page 11930	"order-priority (Application Identification)" on page 11929
"origin (Diameter Base Protocol)" on page 11935	"origin (Diameter Base Protocol)" on page 11934	"origin-autonomous-system (Origin Validation for BGP)" on page 11932
"origin-validation" on page 11937	"ospf" on page 11938	"ospf-instance" on page 11940
"ospf3" on page 11942	"oss-map" on page 11944	"other" on page 11945
"other-bytes" on page 11949	"other-overhead-adjust" on page 11951	"other-stateful-configuration" on page 11953
"other-stateful-configuration (Dynamic Router Advertisement)" on page 11954	"otn-options" on page 11955	"otu4" on page 11960
"otu-ttim-action-enable" on page 11958	"out-delay" on page 11961	"out-of-resources" on page 11963

"out-of-resources (Security Antivirus Sophos Engine)" on page 11965	"outbound-https" on page 11966	"outbound-route-filter" on page 11970
"outbound-ssh" on page 11972	"outer-tag-protocol-id (RFC 2544 Benchmarking)" on page 11975	"outgoing-queue" on page 11976
"output" on page 11980	"output (Accounting)" on page 11981	"output (Congestion Notification)" on page 11984
"output (Dynamic Service Sets)" on page 11986	"output (Forwarding Table)" on page 11992	"output interfaces (Chassis Synchronization)" on page 11996
"output (Mirroring)" on page 11998	"output (Monitoring)" on page 12000	"output (Port Mirroring)" on page 12003
"output (Port Mirroring)" on page 12004	"output (Sampling)" on page 12010	"output (Security Forwarding Options)" on page 12013
"output (Service Set)" on page 11978	"output-chain" on page 11983	"output-filename (Event Policy)" on page 11988
"output-format (Event Policy)" on page 11989	"output-forwarding-class-map" on page 11991	"output-interface (CCC)" on page 11994
"output-interface-index" on page 11995	"output-policer" on page 12001	"output-priority-map" on page 12007
"output-queue-priority" on page 12008	"output-three-color" on page 12014	"output-traffic-control-profile" on page 12016
"output-traffic-control-profile (Dynamic CoS Definition)" on page 12018	"output-traffic-control-profile-remaining" on page 12019	"output-vlan-map " on page 12021
"output-vlan-map (Aggregated Ethernet)" on page 12022	"output-vlan-map (Dynamic Interfaces)" on page 12023	"outputs (Analytics)" on page 12025

"over (Application Identification)" on page 12028	"over-limit" on page 12031	"overall-packet (Flat-File Accounting Options)" on page 12032
"overflow-pool" on page 12035	"overhead-accounting" on page 12036	"overhead-accounting (ANCP)" on page 12038
"overhead-accounting (Dynamic Traffic Shaping)" on page 12039	"overlay (Packet Forwarding Options)" on page 12042	"overlay-ecmp" on page 12041
"overlay-path" on page 12044	"overload (Protocols IS-IS)" on page 12048	"overload (Protocols OSPF)" on page 12051
"overload-pool" on page 12046	"overload-prefix" on page 12047	"override" on page 12054
"override (PIM Static RP)" on page 12061	"override (RADIUS Options)" on page 12063	"override-chap-password" on page 12055
"override-firewall" on page 12057	"override-interval" on page 12058	"override-password (Domain Map)" on page 12060
"overrides (DHCP Local Server)" on page 12065	"overrides (DHCP Relay Agent)" on page 12068	"overrides (DHCP Security)" on page 12071
"overrides (DHCP Security for MX Series)" on page 12072	"overrides (Enhanced Subscriber Management)" on page 12074	"overrides (System Services DHCP)" on page 12077
"oversize-message-drop" on page 12081	"ovlan-cfi (RFC 2544 Benchmarking)" on page 12082	"ovlan-id (RFC 2544 Benchmarking)" on page 12084
"ovlan-priority (RFC 2544 Benchmarking)" on page 12085	"ovsdb" on page 12086	"ovsdb-managed" on page 12088
"owner" on page 12090	"owner (RMON)" on page 12092	"oxid" on page 12093

P

"p2mp (Protocols LDP)" on page 12122	"p2mp (Protocols MPLS)" on page 12124	"p2mp-ldp-next-hop" on page 12119
"p2mp-lsp-next-hop" on page 12120	"p2mp-receive-switch" on page 12125	"p2mp-transmit-switch" on page 12127
"package (Loading on PIC)" on page 12128	"packet-action" on page 12129	"packet-burst (Hierarchical Policer)" on page 12133
"packet-burst (Policer)" on page 12135	"packet-capture" on page 12137	"packet-filter" on page 12141
"packet-format-match" on page 12143	"packet-forwarding-options" on page 12145	"packet-log (Security Flow)" on page 12150
"packet-log (Security IDP Policy)" on page 12152	"packet-log (Security IDP Sensor Configuration)" on page 12153	"packet-loss-priority (RFC 2544 Benchmarking)" on page 12155
"packet-ordering-mode (Application Services)" on page 12157	"packet-scheduling" on page 12158	"packet-size (RFC 2544 Benchmarking)" on page 12160
"packet-timestamp" on page 12165	"packet-timestamp" on page 12162	"packet-types (Dynamic VLAN Authentication)" on page 12163
"pado-advertise" on page 12166	"pair-policy" on page 12168	"pap" on page 12169
"pap (Dynamic PPP)" on page 12171	"pap (L2TP)" on page 12172	"parameter-preference" on page 12173
"parameters" on page 12175	"parse-direction (Access Profile)" on page 12176	"parse-direction (Domain Map)" on page 12178
"parse-order (Domain Map)" on page 12179	"partition" on page 12181	"partition (Gx-Plus)" on page 12182

"partition (NASREQ Diameter Application)" on page 12183	"partition (OCS)" on page 12185	"partition (PCRF)" on page 12187
"partition (s6a)" on page 12190	"pass-through (Firewall-authentication)" on page 12196	"pass-through (interfaces)" on page 12192
"pass-through (Security Policy)" on page 12194	"passive (CHAP)" on page 12198	"passive (Dynamic IGMP Interface)" on page 12199
"passive (Dynamic MLD Interface)" on page 12201	"passive (IGMP)" on page 12202	"passive (MLD)" on page 12204
"passive (MVRP)" on page 12213	"passive (PAP)" on page 12214	"passive (Protocols BGP)" on page 12218
"passive (Protocols IS-IS)" on page 12219	"passive (Protocols Link Management)" on page 12222	"passive (Protocols OSPF)" on page 12223
"passive (Routing Options)" on page 12225	"passive-mode" on page 12205	"passive-mode-tunneling" on page 12207
"passive-mode-tunneling (Security IPsec)" on page 12208	"passive-monitor-mode" on page 12212	"passive-probe-params" on page 12216
"password (Access)" on page 12225	"password (DHCP Local Server)" on page 12226	"password (DHCP Relay Agent)" on page 12229
"password (Flow Collector File Servers)" on page 12232	"password (Interfaces)" on page 12233	"password (Login)" on page 12235
"password (MAC RADIUS Authentication)" on page 12241	"password (Security Antivirus)" on page 12245	"password (Services)" on page 12246
"password (Static Subscribers)" on page 12247	"password (Transfer Log File Servers)" on page 12249	"password-file" on page 12231

"password-options" on page 12243	"path" on page 12250	"path (Protocols MPLS)" on page 12259
"path (Protocols RSVP)" on page 12261	"path (Steering)" on page 12269	"path-database-size" on page 12252
"path-database-size (EX Series Switch Only)" on page 12253	"path-length" on page 12255	"path-mtu" on page 12256
"path-mtu-discovery (Tunnel)" on page 12257	"path-rate-limit" on page 12263	"path-selection" on page 12266
"pattern (Application Identification)" on page 12271	"pattern (Class Attribute)" on page 12272	"pattern (Security IDP)" on page 12275
"pattern-pcre (Security IDP)" on page 12274	"pattern-update (Security Antivirus)" on page 12276	"pause-computation-during-churn" on page 12277
"payload" on page 12279	"pba-interim-logging-interval" on page 12280	"pbb-evpn-core" on page 12282
"pbb-service-options" on page 12283	"pcc-action-profile (PCC Rules)" on page 12284	"pcc-action-profiles" on page 12286
"pcc-context" on page 12289	"pcc-rule" on page 12291	"pcc-rulebases (PCEF)" on page 12293
"pcc-rulebases (PCEF Profile)" on page 12294	"pcc-rules (PCEF)" on page 12297	"pcc-rules (PCEF Profile)" on page 12299
"pcc-time-of-day-profiles" on page 12301	"pce" on page 12303	"pce-group (PCE)" on page 12306
"pce-group (Protocols PCEP)" on page 12307	"pce-type" on page 12309	"pcef" on page 12311

"pcef (Dynamic Profiles)" on page 12314	"pcef-profile (Service Set)" on page 12315	"pcef-profile (TDF Domain)" on page 12317
"pcef-profile (TDF Domain Selection)" on page 12318	"pcep" on page 12320	"pci-address" on page 12322
"pcp-rules" on page 12323	"pcp-server" on page 12324	"pcrf (Diameter Applications)" on page 12325
"pdu-interval" on page 12328	"pdu-threshold" on page 12329	"peak-burst-size" on page 12332
"peak-burst-size (QFX)" on page 12331	"peak-information-rate" on page 12336	"peak-information-rate (QFX)" on page 12334
"peak-rate" on page 12338	"peer" on page 12339	"peer (Diameter Base Protocol)" on page 12351
"peer (Diameter Base Protocol)" on page 12353	"peer (Diameter Network Element)" on page 12354	"peer (Diameter Network Element)" on page 12356
"peer (ICCP)" on page 12366	"peer (Multichassis)" on page 12376	"peer (Protocols LMP)" on page 12379
"peer (Protocols MSDP)" on page 12380	"peer (Protocols RIP)" on page 12382	"peer-active (VPLS Multihoming for FEC 129)" on page 12341
"peer-as (Protocols BGP)" on page 12346	"peer-as (VPLS)" on page 12348	"peer-as-billing-template" on page 12343
"peer-as-list" on page 12345	"peer-certificate-type" on page 12350	"peer-forwarding-ip" on page 12357
"peer-gnf" on page 12359	"peer-groups (access network peers)" on page 12361	"peer-groups (control plane peers)" on page 12362
"peer-groups (core network peers)" on page 12364	"peer-id" on page 12367	"peer-interface (Protocols OSPF)" on page 12370

"peer-interface (Protocols RSVP)" on page 12372	"peer-ip-address-optional" on page 12373	"peer-limit" on page 12375
"peer-origin (Diameter Peer)" on page 12377	"peer-selection-service" on page 12384	"peer-subnet (Easy EVPN LAG Configuration)" on page 12385
"peer-to-peer (Easy EVPN LAG Configuration)" on page 12388	"peer-unit" on page 12390	"peer-unit (Routing Policy Condition)" on page 12391
"peers (Commit)" on page 12393	"peers-synchronize " on page 12394	"peg-df-election" on page 12395
"pem (MX Series Devices, ACX7024 Devices)" on page 12399	"pem (T640, T1600, and T4000 Routers with Six-Input DC Power Supply)" on page 12400	"pem-absence" on page 12397
"pending-queue-watermark" on page 12402	"pending-queue-watermark-abate" on page 12403	"pending-sess-queue-length" on page 12405
"per-esi" on page 12406	"per-esi-vlan" on page 12408	"per-evi" on page 12409
"per-flow" on page 12413	"per-flow (Aggregated Ethernet Interfaces)" on page 12414	"per-group-label" on page 12416
"per-prefix" on page 12417	"per-prefix-label" on page 12419	"per-route-accounting" on page 12420
"per-session-scheduler" on page 12422	"per-unit-scheduler" on page 12423	"perfect-forward-secrecy (Security)" on page 12424
"perfect-forward-secrecy (Services)" on page 12426	"performance" on page 12428	"performance-monitor" on page 12429
"performance-monitoring" on page 12431	"performance-monitoring (OAM LFM)" on page 12432	"performance-monitoring (Protocols MPLS)" on page 12434
"periodic (Chassis)" on page 12436	"periodic" on page 12438	"periodic (Interfaces)" on page 12439

"periodic-traceroute" on page 12441	"permit (Security Persistent NAT)" on page 12444	"permit (Security Policies)" on page 12446
"permit (Security SCTP)" on page 12449	"permit-command" on page 12443	"perpetual-poe" on page 12450
"persistent-learning" on page 12451	"persistent-nat" on page 12453	"persistent-results" on page 12454
"pfc" on page 12455	"pfc (Input Congestion Notification)" on page 12456	"pfc-priority" on page 12458
"pfe" on page 12460	"pfe power off" on page 12463	"pgcp" on page 12465
"pgcp-rules" on page 12466	"pgcp-service" on page 12468	"phase-adjust-threshold" on page 12469
"phase-lock-threshold" on page 12471	"phy-timestamping" on page 12472	"physical-interface (DDoS Flow Detection)" on page 12474
"physical-interface-filter" on page 12476	"physical-interface-policer" on page 12479	"pic" on page 12480
"pic (M Series and T Series Routers)" on page 12487	"pic (MX Series Routers)" on page 12496	"pic (Port)" on page 12498
"pic (PTX Series Routers)" on page 12500	"pic (satellite-policies port-group-alias)" on page 12502	"pic (Subscriber Limits)" on page 12506
"pic-boot-timeout" on page 12484	"pic-console-authentication" on page 12485	"pic-memory-threshold" on page 12489
"pic-mode" on page 12490	"pic-mode (Chassis T1 Mode)" on page 12494	"pic-mode (MX Series)" on page 12491
"pic-services-logging" on page 12504	"pim" on page 12507	"pim-asm" on page 12514

"pim-evpn-gateway" on page 12515	"pim-snooping" on page 12517	"pim-ssm (Provider Tunnel)" on page 12519
"pim-ssm (Selective Tunnel)" on page 12520	"pim-to-igmp-proxy" on page 12522	"pim-to-mld-proxy" on page 12523
"ping-death" on page 12537	"ping-death (IDS Screen Next Gen Services)" on page 12538	"ping-interval" on page 12545
"pki" on page 12610	"pki (Services)" on page 12612	"plmn" on page 12615
"poe" on page 12616	"point-to-point" on page 12619	"point-to-point (MVRP) " on page 12620
"policer" on page 12622	"policer (Applying to a Logical Interface)" on page 12624	"policer (CFM Firewall)" on page 12626
"policer (CFM Global)" on page 12627	"policer (CFM Session)" on page 12629	"policer (Configuring)" on page 12630
"policer (CoS)" on page 12633	"policer (Firewall Filter Action)" on page 12636	"policer (Interface)" on page 12638
"policer (Layer 2 VPN)" on page 12639	"policer (MAC)" on page 12641	"policer-charge" on page 12646
"policer-charge" on page 12646	"policer-drop-probability-low" on page 12635	"policer-overhead" on page 12643
"policer-overhead-adjustment" on page 12644	"policies" on page 12650	"policing" on page 12661
"policing (Protocols LDP)" on page 12662	"policing (Protocols MPLS)" on page 12664	"policy" on page 12665
"policy (advanced-policy-based-routing)" on page 12667	"policy (Aggregate and Generated Routes)" on page 12670	"policy (Event Policy)" on page 12675

"policy (Explicit-Proxy)" on page 12678	"policy (Flow Maps)" on page 12683	"policy (IPsec)" on page 12684
"policy (Multicast-Only Fast Reroute)" on page 12691	"policy (PIM rpf-vector) " on page 12703	"policy (Protocols Multipoint LDP)" on page 12706
"policy (Security Alarms)" on page 12709	"policy (Security Group VPN IKE)" on page 12711	"policy (Security IKE)" on page 12713
"policy (Security Interfaces)" on page 12716	"policy (Security IPsec)" on page 12718	"policy (Security IPsec Interfaces)" on page 12722
"policy (Security Policies)" on page 12724	"policy (Services CoS Next Gen Services)" on page 12728	"policy (Services IKE)" on page 12730
"policy (SSM Maps)" on page 12733	"policy (Stateful Firewall Rules Next Gen Services)" on page 12734	"policy (Subscriber Secure Policy)" on page 12745
"policy (System Security Profile)" on page 12747	"policy-based-logging (LRF Profile)" on page 12672	"policy-db-size" on page 12674
"policy-lookup-cache" on page 12685	"policy-map (Class of Service)" on page 12686	"policy-map-marking" on page 12751
"policy-match" on page 12689	"policy-multipath" on page 12694	"policy-oids" on page 12696
"policy-options" on page 12697	"policy-options (Dynamic Profiles)" on page 12702	"policy-options (QFX)" on page 12700
"policy-rematch" on page 12707	"policy-set " on page 12731	"policy-statement" on page 12736
"policy-stats" on page 12743	"policy-with-count" on page 12749	"polling-interval (Network Management)" on page 12752
"polling-interval" on page 12753	"pon (Access-Line Rate Adjustment)" on page 12755	"pool" on page 12762

"pool (Access)" on page 12763	"pool (Address-Assignment Pools)" on page 12766	"pool (Destination NAT Next Gen Services)" on page 12773
"pool (DHCP Local Server Overrides)" on page 12774	"pool (NAT)" on page 12760	"pool (NAT Rule Next Gen Services)" on page 12778
"pool (Security Destination NAT)" on page 12779	"pool (Security Source NAT)" on page 12781	"pool (Security Source NAT Rule Set)" on page 12783
"pool (Service Interface)" on page 12784	"pool (Service Interfaces)" on page 12786	"pool (Source NAT Next Gen Services)" on page 12787
"pool (System)" on page 12789	"pool (TDF Domain)" on page 12791	"pool-default-port-range" on page 12768
"pool-default-port-range (Source NAT Next Gen Services)" on page 12770	"pool-default-twin-port-range" on page 12771	"pool-match-order" on page 12777
"pool-utilization-alarm" on page 12793	"pool-utilization-alarm (Security Source NAT Pool)" on page 12794	"pool-utilization-alarm (Source NAT Next Gen Services)" on page 12796
"pop" on page 12799	"pop (Dynamic VLANs)" on page 12804	"pop (protocols mpls)" on page 12797
"pop3-profile (Security Content Security Policy Antivirus)" on page 12808	"pop3-profile (Security Content Security Policy Content Filtering)" on page 12809	"pop-all-labels" on page 12800
"pop-and-forward (Protocols MPLS)" on page 12802	"pop-and-forward (Protocols RSVP)" on page 12803	"pop-pop" on page 12806
"pop-swap" on page 12807	"port" on page 12811	"port (Access LDAP)" on page 12812
"port (cflowd)" on page 12813	"port (Chassis)" on page 12815	"port (Chassis Synchronization)" on page 12818

"port (Diameter Peer)" on page 12819	"port (Flow Monitoring)" on page 12823	"port (LRF Profile)" on page 12834
"port (NETCONF)" on page 12844	"port (Origin Validation for BGP)" on page 12846	"port (OVSDB)" on page 12852
"port (Packet Forwarding)" on page 12853	"port (Protocols OpenFlow)" on page 12857	"port (RADIUS Server)" on page 12859
"port (REST API)" on page 12866	"port (RPM)" on page 12870	"port (satellite-policies port-group-alias)" on page 12871
"port (Security Antivirus)" on page 12875	"port (Security Source NAT)" on page 12877	"port (Security Web Filtering Server)" on page 12879
"port (Services)" on page 12880	"port (Services NAT) " on page 12881	"port (Services Voice) " on page 12884
"port (Source NAT Next Gen Services)" on page 12885	"port (SRC Server)" on page 12887	"port (Subscriber Limits)" on page 12890
"port (Syslog)" on page 12891	"port (System Log Messsages)" on page 12893	"port (TWAMP)" on page 12894
"port-error-disable" on page 12821	"port-forwarding" on page 12824	"port-forwarding (Destination NAT Next Gen Services)" on page 12826
"port-forwarding-mappings" on page 12827	"port-forwarding-mappings (Destination NAT Rule Next Gen Services)" on page 12828	"port-group-alias (satellite-policies)" on page 12829
"port-id" on page 12831	"port-id (MACsec for MX Series)" on page 12833	"port-mirroring" on page 12836
"port-mode" on page 12841	"port-mode (Fibre Channel Interfaces)" on page 12843	"port-overloading (Security Source NAT Interface)" on page 12850

"port-overloading-factor (Security Source NAT Interface)" on page 12847	"port-overloading-factor (Security Source NAT Pool)" on page 12849	"port-priority" on page 12856
"port-randomization" on page 12860	"port-range" on page 12861	"port-range (Address Mapping)" on page 12863
"port-range (Application Identification)" on page 12865	"port-round-robin" on page 12867	"port-round-robin (Source NAT Next Gen Services)" on page 12869
"port-scaling-enlargement" on page 12873	"port-scan" on page 12874	"port-status-tlv" on page 12888
"ports (System)" on page 12896	"ports-per-session" on page 12897	"post-attack" on page 12898
"post-attack-timeout" on page 12900	"post-cli-implicit-firewall" on page 12901	"post-convergence-lfa" on page 12903
"post-convergence-lfa (Protocols OSPF)" on page 12904	"post-policy" on page 12910	"post-service-filter" on page 12907
"post-service-filter (Dynamic Service Sets)" on page 12908	"potential-violation" on page 12912	"power" on page 12915
"power" on page 12918	"power-budget-priority" on page 12919	"power-mode-disable" on page 12921
"power-mode-ipsec" on page 12922	"power-mode-ipsec-qat" on page 12924	"pp0 (Dynamic PPPoE)" on page 12925
"ppm" on page 12927	"ppm (Ethernet Switching)" on page 12930	"ppp" on page 12931
"ppp (Group Profile)" on page 12945	"ppp (Profile)" on page 12957	"ppp-access-profile" on page 12944
"ppp-options" on page 12949	"ppp-options (Dynamic PPP)" on page 12951	"ppp-options (L2TP)" on page 12954

"ppp-over-ether" on page 12956	"ppp-subscriber-services" on page 12959	"pppoe (DDoS Protection) " on page 12962
"pppoe (System Processes) " on page 12960	"pppoe-options" on page 12974	"pppoe-options (Dynamic PPPoE)" on page 12976
"pppoe-options (SRX Series)" on page 12972	"pppoe-tags (Adjustment Control Profiles)" on page 12977	"pppoe-underlying-options (Dynamic VLAN Interface Sets)" on page 12979
"pppoe-underlying-options (Static and Dynamic Subscribers)" on page 12980	"pps-limit (Hierarchical Policer)" on page 12982	"pps-limit (Policer)" on page 12984
"pptp" on page 12986	"prbs" on page 12987	"pre-shared-key" on page 13001
"pre-attack" on page 12993	"pre-filter-shellcode" on page 12994	"pre-id-default-policy" on page 12995
"pre-ietf-mode" on page 12998	"pre-rewrite-tos" on page 13000	"pre-shared-key (MX Series)" on page 13003
"pre-shared-key (Security)" on page 13004	"pre-shared-key (Services IKE)" on page 13006	"preauthentication-order (Access Profile)" on page 13007
"precedence" on page 13008	"precision-timers" on page 13010	"precision-timers-max-period" on page 13012
"predefined-attack-groups" on page 13013	"predefined-attacks" on page 13015	"predefined-variable-defaults (Dynamic Client Profiles)" on page 13016
"preempt" on page 13019	"preempt (Chassis Cluster)" on page 13021	"preempt (VRRP)" on page 13025
"preempt-cutover-timer" on page 13023	"preemption" on page 13026	"preemptive-fast-reroute" on page 13028

"prefer-framed-ip-address (RADIUS Clients)" on page 13030	"prefer-framed-ipv6-prefix (RADIUS Clients)" on page 13031	"preference (Access)" on page 13032
"preference (Interface-Level Preference for VPLS Multihoming for FEC 129)" on page 13033	"preference (IPv6 Router Advertisement)" on page 13035	"preference (PEG DF Election)" on page 13037
"preference (Protocols BGP)" on page 13038	"preference (Protocols LDP)" on page 13040	"preference (Protocols MPLS)" on page 13042
"preference (Protocols RIP)" on page 13043	"preference (Protocols RIPng)" on page 13045	"preference (Routing Options)" on page 13046
"preference (Site-Level Preference for VPLS Multihoming for FEC 129)" on page 13049	"preference (Subscriber Management)" on page 13050	"preference (Tunnel Profile)" on page 13052
"preferred" on page 13053	"preferred-ciphers" on page 13055	"preferred-lifetime" on page 13056
"preferred-lifetime (Dynamic Router Advertisement)" on page 13058	"preferred-prefix-length" on page 13059	"preferred-source-address" on page 13060
"prefix" on page 13062	"prefix (Access IPv6)" on page 13067	"prefix (Address-Assignment Pools)" on page 13071
"prefix (Circuit ID for Option 82)" on page 13074	"prefix (DHCP Client)" on page 13076	"prefix (DHCP Relay Agent)" on page 13077
"prefix (DHCPv6 Options)" on page 13079	"prefix (Dynamic Router Advertisement)" on page 13081	"prefix (Protocols IPv6 Neighbor Discovery)" on page 13100
"prefix (Remote ID for Option 82)" on page 13101	"prefix (Security Static NAT)" on page 13103	"prefix-65-127-disable" on page 13064
"prefix-action (Configuring)" on page 13068	"prefix-action (Firewall Filter Action)" on page 13070	"prefix-attribute-flags" on page 13072

"prefix-export-limit (Protocols IS-IS)" on page 13083	"prefix-export-limit (Protocols OSPF)" on page 13085	"prefix-limit" on page 13087
"prefix-list (Junos Telemetry Interface)" on page 13093	"prefix-list" on page 13090	"prefix-list (PIM RPF Selection)" on page 13097
"prefix-list-filter" on page 13092	"prefix-list-name" on page 13095	"prefix-name (Security Static NAT)" on page 13098
"prefix-segment" on page 13105	"prefix-segment (Routing Options)" on page 13109	"prefix-segment-range" on page 13107
"premium (Hierarchical Policer)" on page 13110	"premium (Output Priority Map)" on page 13112	"premium (Policer)" on page 13113
"preprovisioned" on page 13115	"preprovisioned (MX Series Virtual Chassis)" on page 13116	"preserve-fpc-poweron-sequence" on page 13118
"preserve-incoming-fragment-size" on page 13119	"preserve-interface" on page 13121	"preserve-nexthop-hierarchy (SRv6 TE)" on page 13122
"preserve-parity (Source NAT Next Gen Services)" on page 13124	"preserve-range (Source NAT Next Gen Services)" on page 13125	"primary" on page 13126
"primary" on page 13127	"primary (Adaptive Services Interfaces)" on page 13129	"primary (Address on Interface)" on page 13130
"primary connection (Identity Management Advanced Query) " on page 13131	"primary (Interface for Router)" on page 13140	"primary (Link Services IQ PIC Interfaces) " on page 13142
"primary (MBGP MVPN)" on page 13143	"primary (Protocols MPLS)" on page 13144	"primary (Virtual Tunnel in Routing Instances)" on page 13148
"primary (VPLS Multihoming)" on page 13151	"primary-data-record-fields" on page 13134	"primary-dns" on page 13137

"primary-interface (Aggregated Inline Services)" on page 13138	"primary-server" on page 13146	"primary-vlan" on page 13150
"primary-wins" on page 13153	"priority" on page 13154	"priority (Application Identification With Next Gen Services)" on page 13156
"priority (Bootstrap)" on page 13157	"priority (CFM MEP)" on page 13159	"priority (Chassis Synchronization Source)" on page 13160
"priority (Clock Synchronization)" on page 13162	"priority (DDoS)" on page 13165	"priority (Diameter Network Element)" on page 13166
"priority (Diameter Peer)" on page 13168	"priority (Dynamic Schedulers)" on page 13169	"priority (Fabric Priority)" on page 13171
"priority (Fabric Queues, Schedulers)" on page 13172	"priority (FIP)" on page 13174	"priority (Management Traffic)" on page 13180
"priority (OAM Connectivity-Fault Management)" on page 13182	"priority (PIM Interfaces)" on page 13184	"priority (PIM RPs)" on page 13186
"priority (PoE)" on page 13188	"priority (policy-options)" on page 13189	"priority (Power over Ethernet)" on page 13191
"priority (Protocols MPLS)" on page 13193	"priority (Protocols OAM)" on page 13194	"priority (Protocols Router Discovery)" on page 13196
"priority (Protocols RSVP)" on page 13197	"priority (Protocols STP)" on page 13199	"priority (Protocols VRRP)" on page 13201
"priority (RADIUS Network Elements)" on page 13203	"priority (Redundant Power System)" on page 13204	"priority (Schedulers)" on page 13206
"priority1" on page 13209	"priority2" on page 13211	"priority-cost (VRRP)" on page 13163

"priority-flow-control" on page 13176	"priority-hold-time" on page 13177	"priority-hold-time" on page 13179
"priority-override (Event Policy)" on page 13183	"privacy-3des" on page 13212	"privacy-aes128" on page 13213
"privacy-des" on page 13215	"privacy-key" on page 13216	"privacy-none" on page 13217
"privacy-password" on page 13218	"private-vlan" on page 13220	"proactive-arp-detection" on page 13221
"probe" on page 13228	"probe-count" on page 13231	"probe-interval" on page 13233
"probe-limit" on page 13234	"probe-server" on page 13236	"probe-type" on page 13238
"process (Memory)" on page 13242	"process-inform" on page 13240	"process-monitor" on page 13245
"process-non-null-as-null-register" on page 13246	"processes" on page 13248	"product-model (Junos Fusion)" on page 13250
"product-name" on page 13252	"products" on page 13253	"profile" on page 13256
"profile" on page 13257	"profile (Access)" on page 13259	"profile (Access)" on page 13263
"profile (Application Firewall)" on page 13271	"profile (Application Identification)" on page 13254	"profile (Captive Portal Content Delivery)" on page 13273
"profile (Explicit Proxy)" on page 13277	"profile (HTTP Header Enrichment)" on page 13282	"profile (icap-redirect)" on page 13283
"profile (Juniper Secure Connect)" on page 13286	"profile (LRF)" on page 13289	"profile (Rule Sets)" on page 13291
"profile (security)" on page 13292	"profile (Security Antispam SBL)" on page 13310	"profile (Security Antivirus Juniper Express Engine)" on page 13312

"profile (Security Antivirus Kaspersky Lab Engine)" on page 13314	"profile (Security Content Filtering)" on page 13317	"profile (Security GTP)" on page 13318
"profile (Security SCTP)" on page 13323	"profile (Security Sophos Engine Antivirus)" on page 13326	"profile (Security Web Filtering Juniper Enhanced)" on page 13328
"profile (Security Web Filtering Juniper Local)" on page 13330	"profile (Security Web Filtering Surf Control Integrated)" on page 13332	"profile (Security Web Filtering Websense Redirect)" on page 13333
"profile (Services Application Identification)" on page 13335	"profile (Services PCEF)" on page 13336	"profile (Services Proxy)" on page 13338
"profile (Services SSL Proxy)" on page 13339	"profile (SSL Initiation)" on page 13343	"profile (SSL Termination)" on page 13347
"profile (Traffic Load Balancer)" on page 13349	"profile (Web Filter)" on page 13357	"profile-type (Dynamic Service Profiles)" on page 13353
"profile-type (PTP)" on page 28809	"profile-variable-set (Dynamic Profiles)" on page 13354	"profile-variable-set (Routing Instances)" on page 13356
"profile(dynamic-application)" on page 13275	"profilerd" on page 13360	"profiles" on page 13362
"profiles (AAA)" on page 13363	"profiles (PCEF)" on page 13365	"profiles (RFC 2544 Benchmarking)" on page 13368
"programmable-rpd" on page 13369	"promiscuous" on page 13372	"promiscuous-mode (Dynamic IGMP Interface)" on page 13373
"promiscuous-mode (Interfaces)" on page 13374	"promiscuous-mode (Protocols IGMP)" on page 13375	"promote" on page 13377
"propagate-settings" on page 13379	"propagate-tc" on page 13380	"propagation-delay" on page 13381

"proposal (Security Group VPN Member IKE)" on page 13383	"proposal (Security Group VPN Server IKE)" on page 13385	"proposal (Security Group VPN Server IPsec)" on page 13387
"proposal (Security IKE)" on page 13389	"proposal (Security IKE Dynamic SA)" on page 13393	"proposal (Security IPsec)" on page 13394
"proposal (Services IKE)" on page 13399	"proposal (Services IPsec VPN)" on page 13400	"proposals" on page 13403
"proposals (Services)" on page 13402	"protect core" on page 13407	"protect (Security)" on page 13406
"protect-interface" on page 13409	"protect-maintenance-association (OAM)" on page 13410	"protected-l2circuit" on page 13412
"protection (Protocols BGP)" on page 13416	"protection (Protocols MPLS)" on page 13417	"protection-group" on page 13413
"protection-revert-time" on page 13419	"protector-interface" on page 13420	"protector-pe" on page 13421
"protocol" on page 13424	"protocol (Application Identification)" on page 13429	"protocol (Applications)" on page 13431
"protocol (Drop Profile Map)" on page 13437	"protocol (Drop Profiles)" on page 13438	"protocol (Dynamic Schedulers)" on page 13439
"protocol (early-aging)" on page 13428	"protocol (Firewall Filter)" on page 13423	"protocol (Flow Descriptions)" on page 13441
"protocol (Host Outbound Traffic)" on page 13443	"protocol (IPsec) " on page 13444	"protocol (Junos OS)" on page 13446
"protocol (Junos-FIPS Software)" on page 13445	"protocol (OVSDB)" on page 13454	"protocol (Protocols OpenFlow)" on page 13455
"protocol (Rewrite Rules)" on page 13457	"protocol (Schedulers)" on page 13459	"protocol (Security Destination NAT)" on page 13460

"protocol (Security Forwarding Options)" on page 13461	"protocol (Security IDP Signature Attack)" on page 13463	"protocol (Security Source NAT)" on page 13469
"protocol (Server MEP)" on page 13471	"protocol (Services Proxy)" on page 13472	"protocol (Subscriber Secure Policy)" on page 13474
"protocol-binding" on page 13433	"protocol-command" on page 13434	"protocol-down" on page 13436
"protocol-master" on page 13448	"protocol-name" on page 13451	"protocol-origin" on page 13452
"protocol-version" on page 13475	"protocols" on page 13477	"protocols" on page 13748
"protocols (DDoS)" on page 13751	"protocols (DDoS) (ACX Series, PTX Series, and QFX Series)" on page 13765	"protocols (Dynamic Profiles)" on page 13791
"protocols (Fabric)" on page 13794	"protocols (FIP)" on page 13795	"protocols rsvp expand-flood-reflector-hop" on page 13797
"protocols (Security Zones Host Inbound Traffic)" on page 13798	"protocols (Security Zones Interfaces)" on page 13800	"protocols (STP Type)" on page 13802
"protocols (Tenant Systems)" on page 13804	"provider-tunnel" on page 13806	"providers" on page 13812
"provisioning-method (Remote Device Management)" on page 13814	"provisioning-order (Diameter Applications)" on page 13815	"proxy (Bridge Domains)" on page 13819
"proxy (Fibre Channel)" on page 13825	"proxy (interfaces)" on page 13817	"proxy (Interfaces)" on page 13828
"proxy (Multicast VLAN Registration)" on page 13834	"proxy (Security Antivirus)" on page 13839	"proxy (Services)" on page 13841
"proxy (snmp)" on page 13843	"proxy (System)" on page 13845	"proxy-arp" on page 13820

"proxy-arp (Dynamic Profiles)" on page 13822	"proxy-arp (Security NAT)" on page 13823	"proxy-generate" on page 13826
"proxy-mac" on page 13829	"proxy-macip-advertisement" on page 13831	"proxy-mode" on page 13833
"proxy-ndp (Security NAT)" on page 13836	"proxy-profile" on page 13838	"ps0 (Pseudowire Subscriber Interfaces)" on page 13847
"pseudowire-service (Pseudowire Subscriber Interfaces)" on page 13848	"pseudowire-status-tlv" on page 13849	"psn-tunnel-endpoint" on page 13851
"psu" on page 13852	"psu (satellite-management)" on page 13853	"ptp" on page 13855
"ptp-mode" on page 13865	"ptsp-rules" on page 13866	"pulse-per-second-enable" on page 13868
"purge-flow-timer (Protocols OpenFlow)" on page 13869	"purge-timeout (programmable-rpd)" on page 13870	"purging" on page 13872
"push" on page 13875	"push (Dynamic VLANs)" on page 13878	"push (protocols mpls)" on page 13873
"push-backup-to-master (Accounting Options)" on page 13876	"push-push" on page 13879	"push-to-identity-management" on page 13880
"pvlan" on page 13882	"pvlan-trunk" on page 13883	"python-script-user (Event and SNMP Scripts)" on page 13884

Q

"qos-adjust" on page 13888	"qos-adjust-adsl" on page 13890	"qos-adjust-adsl2" on page 13892
"qos-adjust-adsl2-plus" on page 13894	"qos-adjust-hierarchical" on page 13896	"qos-adjust-other" on page 13897
"qos-adjust-sdsl" on page 13899	"qos-adjust-vdsl" on page 13901	"qos-adjust-vdsl2" on page 13903
"qualified-bum-pruning-mode" on page 13905	"qualified-next-hop (Access)" on page 13906	"qualified-next-hop (Access-Internal)" on page 13907
"qualified-next-hop (Dynamic Access-Internal Routes)" on page 13908	"qualified-next-hop (Static Routes)" on page 13910	"qualified-vlan " on page 13912
"quality (Interfaces)" on page 13914	"quality-level " on page 13915	"quality-level (Clock Synchronization)" on page 13917
"quality-level (hybrid)" on page 13919	"quality-mode-enable (MX)" on page 13920	"quality-mode-enable (PTX)" on page 13922
"quarantine-message (Security Content Security)" on page 13924	"querier (performance-monitoring)" on page 13925	"query-interval (Bridge Domains)" on page 13927
"query-interval (Protocols IGMP)" on page 13929	"query-interval (Protocols IGMP AMT)" on page 13930	"query-interval (Protocols MLD)" on page 13932
"query-last-member-interval (Bridge Domains)" on page 13933	"query-last-member-interval (Protocols IGMP)" on page 13935	"query-last-member-interval (Protocols MLD)" on page 13936
"query-response-interval (Bridge Domains)" on page 13938	"query-response-interval (Protocols IGMP)" on page 13940	"query-response-interval (Protocols IGMP AMT)" on page 13941
"query-response-interval (Protocols MLD)" on page 13943	"queue (Global Queues)" on page 13944	"queue (Restricted Queues)" on page 13949

"queue-monitoring" on page 13946	"queue-num" on page 13947	"queue-statistics" on page 13950
"queue-threshold" on page 13953	"queues" on page 13955	

R

"r2cp" on page 13990	"radio-router (Interfaces)" on page 13991	"radius (Access Profile)" on page 13993
"radius-accounting" on page 13997	"radius (Security Log)" on page 14018	"radius (System)" on page 14034
"radius-coa (Adjustment Control Profiles)" on page 13999	"radius-disconnect" on page 14001	"radius-disconnect (DHCP Local Server)" on page 14003
"radius-disconnect-port" on page 14005	"radius-flow-tap" on page 14006	"radius-options (Access)" on page 14010
"radius-options (Interfaces)" on page 14011	"radius-options (System)" on page 14013	"radius-reachability" on page 14015
"radius-realm" on page 14017	"radius-server" on page 14020	"radius-server (Access)" on page 14026
"radius-server (System)" on page 14030	"radius-server-name" on page 14028	"radsec " on page 14036
"radsec-destination" on page 14040	"raise-threshold" on page 14041	"raise-trap (Event Policy)" on page 14042
"random" on page 14043	"range (Access)" on page 14045	"range (Address-Assignment Pools)" on page 14048

"range (Source NAT Next Gen Services)" on page 14049	"range-address" on page 14047	"ranges (Dynamic Stacked VLAN)" on page 14051
"ranges (Dynamic VLAN)" on page 14052	"rapid-commit (DHCPv6 Client)" on page 14054	"rapid-commit (DHCPv6 Local Server)" on page 14055
"rate" on page 14057	"rate (Forwarding Options)" on page 14060	"rate (Interface Services)" on page 14061
"rate (Routing Instances)" on page 14073	"rate-cap" on page 14058	"rate-limit" on page 14063
"rate-limit (Aggregated rate limit)" on page 14065	"rate-limit (Security GTP)" on page 14068	"rate-limit (Security Log)" on page 14069
"rate-limiters" on page 14071	"rcp-count" on page 14075	"rdi" on page 14077
"rdma-opcode" on page 14078	"re-assembler" on page 14079	"re-enroll-trigger-time-percentage" on page 14081
"re-generate-keypair" on page 14082	"reachability-fault" on page 14083	"reachable-time" on page 14086
"reachable-time (Dynamic Router Advertisement)" on page 14087	"read-view" on page 14089	"readvertise" on page 14090
"real-service (Traffic Load Balancer)" on page 14092	"realm" on page 14093	"realm (Diameter Origin)" on page 14096
"realm-delimiter (Domain Map)" on page 14095	"realm-parse-direction (Domain Map)" on page 14098	"realtime-ukern-thread (BFD)" on page 14099
"reassemble-packets" on page 14100	"reassembly-timeout" on page 14102	"reauthenticate (DHCP Local Server)" on page 14103
"rebalance (Aggregated Ethernet Interfaces)" on page 14106	"rebalance-periodic (Aggregated Ethernet Subscriber Interfaces)" on page 14108	"rebalance-periodic (PS Interface)" on page 14109

"rebalance-subscriber-granularity (PS interface)" on page 14111	"receive (Protocols RIP)" on page 14117	"receive (Protocols RIPng)" on page 14119
"receive-ais" on page 14112	"receive-failure-threshold (RFC 2544 Benchmarking)" on page 14113	"receive-options-packets" on page 14116
"receive-options-packets (Passive Flow Monitoring)" on page 14114	"receive-ttl-exceeded" on page 14121	"receive-ttl-exceeded (Passive Flow Monitoring)" on page 14120
"receive-window" on page 14123	"receiver" on page 14124	"receiver (Graceful Restart for BGP Helper)" on page 14127
"recommendation-tlv" on page 14129	"recommended" on page 14130	"recommended-action" on page 14131
"reconfigure (DHCP Local Server)" on page 14133	"reconnect-time" on page 14135	"record" on page 14136
"record (Origin Validation for BGP)" on page 14138	"record-route-option (IDS Screen Next Gen Services)" on page 14139	"recover-time (DDoS)" on page 14141
"recovered-clock" on page 14142	"recovery" on page 14143	"recovery" on page 14146
"recovery-time (LDP)" on page 14148	"recovery-timeout" on page 14149	"recycle-bandwidth-profiles" on page 14152
"red-buffer-occupancy" on page 14154	"red-differential-delay" on page 14156	"redial-delay" on page 14157
"redirect (PCC Action Profiles)" on page 14159	"redirect-traffic" on page 14160	"redirect-url" on page 14162
"redirect-wx (Application Services)" on page 14164	"redistribute-all-traffic (Aggregated Multiservices)" on page 14165	"redundancy" on page 14166

"redundancy (Graceful Switchover)" on page 14170	"redundancy (Junos OS)" on page 14183	"redundancy (Junos OS Evolved)" on page 14185
"redundancy (M:N Subscriber Redundancy)" on page 14186	"redundancy (Power Management)" on page 14201	"redundancy (satellite-management)" on page 14203
"redundancy-event (Services Redundancy Daemon)" on page 14168	"redundancy-group" on page 14171	"redundancy-group (Chassis - MX Series)" on page 14176
"redundancy-group (Chassis Cluster)" on page 14173	"redundancy-group (Interfaces)" on page 14179	"redundancy-group (Redundant Tunnel)" on page 14180
"redundancy-group-id (Junos Fusion)" on page 14177	"redundancy-interface-process" on page 14182	"redundancy-mode" on page 14207
"redundancy-mode" on page 14191	"redundancy-options (Adaptive Services Interfaces) " on page 14192	"redundancy-options (Aggregated Multiservices)" on page 14194
"redundancy-options (Link Services IQ PIC Interfaces) " on page 14195	"redundancy-options (Stateful Synchronization)" on page 14197	"redundancy-policy (Interchassis Services Redundancy)" on page 14199
"redundancy-set" on page 14204	"redundancy-set-id (Service Set)" on page 14206	"redundant-ether-options" on page 14209
"redundant-logical-tunnel" on page 14211	"redundant-parent (Interfaces)" on page 14212	"redundant-parent (Interfaces Fast Ethernet)" on page 14214
"redundant-parent (Interfaces Gigabit Ethernet)" on page 14215	"redundant-power-system" on page 14216	"redundant-pseudo-interface-options" on page 14218
"redundant-sources" on page 14219	"redundant-trunk-group" on page 14221	"redundant-virtual-tunnel" on page 14222
"reference-bandwidth (Protocols IS-IS)" on page 14224	"reference-bandwidth (Protocols OSPF)" on page 14226	"reflect-etype" on page 14228

"reflect-mode" on page 14229	"reflective-relay" on page 14232	"reflexive | revert | reverse" on page 14233
"refresh (Commit Scripts)" on page 14234	"refresh (Event Scripts)" on page 14236	"refresh (JET)" on page 14250
"refresh (Op Scripts)" on page 14251	"refresh (SNMP Scripts)" on page 14255	"refresh-from (Commit Scripts)" on page 14238
"refresh-from (Event Scripts)" on page 14240	"refresh-from (JET)" on page 14242	"refresh-from (Op Scripts)" on page 14244
"refresh-from (SNMP Scripts)" on page 14246	"refresh-interval" on page 14248	"refresh-rate (Flow Monitoring Logs for NAT)" on page 14253
"refresh-time" on page 14257	"regex (Class Attribute)" on page 14260	"regex-additive-logic" on page 14258
"regexp" on page 14261	"region" on page 14262	"register-limit" on page 14264
"register-probe-time" on page 14266	"registration" on page 14268	"reject (Security) " on page 14270
"reject-unauthorized-ipv6cp" on page 14271	"rejoin-timeout (Aggregated Multiservices)" on page 14273	"relay (AMT Protocol)" on page 14291
"relay (IGMP)" on page 14292	"relay-agent-interface-id (DHCP Local Server)" on page 14275	"relay-agent-interface-id (DHCPv6 Relay Agent)" on page 14276
"relay-agent-interface-id (DHCPv6 Relay Agent Username)" on page 14278	"relay-agent-option" on page 14279	"relay-agent-option-79 " on page 14281
"relay-agent-remote-id (DHCP Local Server)" on page 14282	"relay-agent-remote-id (DHCPv6 Relay Agent)" on page 14284	"relay-agent-remote-id (DHCPv6 Relay Agent Username)" on page 14286

"relay-agent-subscriber-id (DHCP Local Server)" on page 14288	"relay-agent-subscriber-id (DHCPv6 Relay Agent)" on page 14289	"relay-option (DHCP Relay Agent)" on page 14297
"relay-option-82" on page 14294	"relay-option-vendor-specific (dhcpv6)" on page 14299	"relay-server-group (DHCP Relay Agent Option)" on page 14300
"relay-source" on page 14302	"reliable" on page 14304	"remote" on page 14305
"remote (Host VPN)" on page 14324	"remote-access (Juniper Secure Connect)" on page 14307	"remote-address" on page 14311
"remote-address (for LMP Control Channel)" on page 14312	"remote-address (for LMP Traffic Engineering)" on page 14314	"remote-debug-permission" on page 14315
"remote-device-management" on page 14317	"remote-engine" on page 14318	"remote-execution (Event Scripts)" on page 14320
"remote-gateway" on page 14322	"remote-gateway (Tunnel Profile)" on page 14323	"remote-id" on page 14327
"remote-id (DHCP Relay Agent)" on page 14331	"remote-id (MX Series)" on page 14337	"remote-id (protocols link-management)" on page 14330
"remote-id (Services)" on page 14326	"remote-id (VLAN Authentication Username)" on page 14338	"remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)" on page 14334
"remote-identity" on page 14340	"remote-interface-switch" on page 14342	"remote-ip-address" on page 14344
"remote-ip-host-routes" on page 14346	"remote-loop-enable" on page 14348	"remote-loopback" on page 14350
"remote-maintenance-association (OAM)" on page 14351	"remote-mep" on page 14352	"remote-mep (EX Series Switch Only)" on page 14354

"remote-port-range" on page 14355	"remote-ports" on page 14357	"remote-site-id" on page 14360
"remote-traffic-selector" on page 14361	"remote-vtep-list" on page 14362	"remotely-controlled" on page 14364
"remove-ie" on page 14366	"remove-private" on page 14368	"remove-when-no-subscribers" on page 14370
"renegotiation (Services)" on page 14373	"reopen-time" on page 14374	"replace-ip-source-with (DHCP Relay Agent)" on page 14377
"replay-protect" on page 14380	"replay-window-size" on page 14383	"replay-attacks" on page 14379
"replay-protect (MX Series)" on page 14382	"replay-window-size (MX Series)" on page 14385	"replicator" on page 14387
"report (LRF Rule)" on page 14395	"report (Security Log)" on page 14398	"report-ingress-shaping-rate (Dynamic CoS Interfaces)" on page 14390
"report-interface-descriptions (Access)" on page 14392	"report-local-rule (PCRF Partition)" on page 14393	"report-resource-allocation (PCRF Partition)" on page 14396
"report-skip" on page 14403	"report-successful-resource-allocation (PCRF Partition)" on page 14404	"report-zero-oif-gw-on-discard" on page 14406
"req-option" on page 14408	"req-timeout" on page 14409	"request (Chassis Synchronization Source)" on page 14573
"request (Clock Synchronization)" on page 14579	"request-cache-timeout (RADIUS Snoop Segment)" on page 14432	"request-max-tcp-connections (System Process)" on page 14678
"request-rate (Access)" on page 14738	"request-response" on page 14741	"request-timeout" on page 15432

"request-type" on page 15433	"requested-ip-network-match (DHCP Local Server)" on page 15572	"required-depth" on page 15574
"reroute-address" on page 15575	"reserve ingress-egress-mcast" on page 15577	"reserved (Security IDP Custom Attack)" on page 15580
"reserved-bandwidth" on page 15579	"reset (Security IDP)" on page 15582	"reset-count" on page 15584
"reset-tracking-bit" on page 15583	"resignal-minimum-bandwidth" on page 15585	"resiliency" on page 15586
"resilient-hash" on page 15589	"resize (Ephemeral Database)" on page 15590	"resize-ephemeral-default" on page 15592
"resolution" on page 15594	"resolution-map" on page 15597	"resolution-ribs" on page 15599
"resolve" on page 15600	"resolve-vpn" on page 15602	"resource (Analytics)" on page 15604
"resource (Interfaces)" on page 15608	"resource-category jtree (Resource Monitor)" on page 15605	"resource-cleanup" on page 15607
"resource-list (Network Management)" on page 15610	"resource-list (Telemetry)" on page 15612	"resource-manager" on page 15614
"resource-manager (Forwarding-options)" on page 15616	"resource-monitor" on page 15618	"resource-profiles (Analytics)" on page 15620
"resource-watch" on page 15622	"resources" on page 15625	"respond-bad-spi (Services IKE Policy)" on page 15626
"responder (performance-monitoring)" on page 15628	"response-cache-timeout (RADIUS Client)" on page 15629	"rest" on page 15631
"restart-duration (Multicast Snooping)" on page 15653	"restart-duration (Protocols)" on page 15651	"restart-duration (Routing Options)" on page 15655

"restart-path" on page 15662	"restart-time (BGP Graceful Restart)" on page 15663	"restart-time (Fabric Control)" on page 15665
"restarter (Graceful Restart for BGP Restarter)" on page 15667	"restore-interval" on page 15669	"restricted-queues" on page 15671
"retain" on page 15672	"retain-hold-resource" on page 15674	"reth-count (Chassis Cluster)" on page 15675
"retransmission-attempt (DHCP Client)" on page 15677	"retransmission-attempt (DHCPv6 Client)" on page 15678	"retransmission-count-established (L2TP)" on page 15680
"retransmission-count-not-established (L2TP)" on page 15681	"retransmission-interval" on page 15683	"retransmission-interval (DHCP)" on page 15684
"retransmission-interval (DHCP Client)" on page 15685	"retransmit-interval (Services)" on page 15687	"retransmit-timer" on page 15689
"retransmit-timer (Dynamic Router Advertisement)" on page 15690	"retry (Access LDAP)" on page 15691	"retry (Access RADIUS)" on page 15693
"retry (Adaptive Services Interface)" on page 15694	"retry-count (Security Log)" on page 15696	"retry (Event Policy)" on page 15704
"retry (RADIUS Server)" on page 15714	"retry (Services Flow Collector)" on page 15716	"retry-count (Chassis Cluster)" on page 15700
"retry-count (Event Policy)" on page 15701	"retry-count (FPC Error Condition)" on page 15698	"retry-count (SNMP)" on page 15697
"retry-delay" on page 15703	"retry-interval" on page 15706	"retry-interval (Chassis Cluster)" on page 15707
"retry-limit" on page 15709	"retry-limit (Protocols Link Management)" on page 15710	"retry-options" on page 15712

"retry-timer" on page 15717	"reverse" on page 15718	"reverse-oif-mapping" on page 15720
"reverse-route-packet-mode-vr" on page 15721	"reverse-tcp-mss" on page 15722	"revert-interval (Access)" on page 15724
"revert-interval (Access LDAP)" on page 15725	"revert-interval (Access RADIUS)" on page 15727	"revert-interval (RADIUS Server)" on page 15728
"revert-time (Protocols Layer 2 Circuits)" on page 15729	"revert-timer" on page 15732	"revertive" on page 15734
"revision-level" on page 15735	"revocation-check" on page 15737	"revocation-check (Security PKI)" on page 15739
"rewrite-rules (Definition)" on page 15744	"rewrite-rules (CoS Host Outbound Traffic)" on page 15743	"rewrite-rules (Definition)" on page 15741
"rewrite-rules (Definition)" on page 15746	"rewrite-rules (Dynamic CoS Interfaces)" on page 15748	"rewrite-rules (Interfaces)" on page 15749
"rewrite-rules (Interfaces)" on page 15751	"rewrite-value (Fibre Channel Interfaces)" on page 15753	"rfc2544" on page 15757
"rfc2544-benchmarking" on page 15759	"rfc6514-compliant-safi129 (Protocols BGP)" on page 15763	"rfc-compliant (NETCONF)" on page 15755
"rib (General)" on page 15764	"rib (Multitopology Routing)" on page 15785	"rib (Protocols BGP)" on page 15787
"rib (Route Resolution)" on page 15789	"rib-group (Protocols BGP)" on page 15767	"rib-group (Protocols DVMRP)" on page 15769
"rib-group (Protocols IS-IS)" on page 15770	"rib-group (Protocols MSDP)" on page 15772	"rib-group (Protocols OSPF)" on page 15774

"rib-group (Protocols PIM)" on page 15775	"rib-group (Protocols RIP)" on page 15777	"rib-group (Routing Options)" on page 15778
"rib-group (Source Packet Routing)" on page 15780	"rib-groups" on page 15781	"rib-list" on page 15783
"rib-service (programmable-rpd)" on page 15790	"rib-sharding" on page 15792	"rift" on page 15794
"ring-id" on page 15797	"ring-protection-link-end" on page 15799	"ring-protection-link-owner" on page 15800
"riot-loopback" on page 15801	"rip" on page 15803	"ripng" on page 15804
"rising-event-index" on page 15805	"rising-threshold (Health Monitor)" on page 15807	"rising-threshold (RMON)" on page 15808
"rmon" on page 15810	"rng" on page 15812	"roaming-mode" on page 15813
"robust-count (IGMP Snooping)" on page 15815	"robust-count (MLD Snooping)" on page 15816	"robust-count (Protocols IGMP)" on page 15818
"robust-count (Protocols IGMP AMT)" on page 15820	"robust-count (Protocols MLD)" on page 15821	"robustness-count" on page 15823
"role" on page 15824	"role (MX Series Virtual Chassis)" on page 15827	"role (Protocols OpenFlow)" on page 15829
"root-address" on page 15840	"root-authentication" on page 15832	"root-ca (Services)" on page 15834
"root-logical-system" on page 15835	"root-login (JDM)" on page 15837	"root-streaming" on page 15838
"route (Access)" on page 15841	"route (Access Internal)" on page 15843	"route (Access-Internal)" on page 15845

"route (chassis)" on page 15853	"route (Diameter Network Element)" on page 15854	"route (Interfaces)" on page 15868
"route-accounting" on page 15846	"route-active-on" on page 15848	"route-attributes" on page 15850
"route-change-timeout" on page 15852	"route-distinguisher" on page 15856	"route-distinguisher-id" on page 15860
"route-distinguisher-id-use-router-id" on page 15889	"route-filter" on page 15861	"route-filter-list" on page 15863
"route-localization" on page 15869	"route-monitoring" on page 15871	"route-record" on page 15874
"route-refresh-priority" on page 15876	"route-suppression (DHCP Local Server and Relay Agent)" on page 15877	"route-target (Protocols BGP)" on page 15881
"route-target (Protocols MVPN)" on page 15883	"route-target-filter" on page 15879	"route-timeout (Protocols RIP)" on page 15885
"route-timeout (Protocols RIPng)" on page 15887	"router-advertisement" on page 15890	"router-advertisement (Dynamic Profiles)" on page 15895
"router-advertisement-guard" on page 15897	"router-discovery" on page 15900	"router-id" on page 15902
"routing (System Processes)" on page 15993	"routing-engine" on page 15904	"routing-engine (Chassis)" on page 15905
"routing-engine (Chassis Redundancy)" on page 15908	"routing-engine (SNMP Resource Level)" on page 15912	"routing-engine-mode (TLB)" on page 15995
"routing-engine-profile" on page 15909	"routing-engine-services" on page 15911	"routing-header" on page 15914
"routing-instance" on page 15921	"routing-instance (Access LDAP)" on page 15925	"routing-instance (Access RADIUS)" on page 15927

"routing-instance (Advanced Policy-Based Routing)" on page 15928	"routing-instance (Diameter Peer)" on page 15931	"routing-instance (Diameter Transport)" on page 15932
"routing-instance (Event Scripts)" on page 15934	"routing-instance (Flow Monitoring)" on page 15929	"routing-instance (Forwarding-options Analyzer)" on page 15920
"routing-instance (Interfaces)" on page 15923	"routing-instance (JET Scripts)" on page 15935	"routing-instance (Junos Telemetry Interface)" on page 15937
"routing-instance (L2TP Destination)" on page 15938	"routing-instance (L2TP Tunnel Destination)" on page 15940	"routing-instance (PCC Action Profiles)" on page 15951
"routing-instance (Port Mirroring)" on page 15953	"routing-instance (PPPoE Service Name Tables)" on page 15954	"routing-instance (RPM)" on page 15956
"routing-instance (Security Content Security)" on page 15958	"routing-instance (Security Destination NAT)" on page 15959	"routing-instance (Security Source NAT)" on page 15960
"routing-instance (Services)" on page 15915	"routing-instance (SNMP)" on page 15918	"routing-instance (SNMPv3)" on page 15917
"routing-instance (Syslog)" on page 15962	"routing-instance (System Scripts)" on page 15963	"routing-instance (Tunnel Profile)" on page 15965
"routing-instance-access" on page 15924	"routing-instance-list (TWAMP)" on page 15941	"routing-instance-name" on page 15943
"routing-instance-name (circuit-id)" on page 15944	"routing-instance-name (DHCP Local Server)" on page 15945	"routing-instance-name (DHCP Relay Agent)" on page 15947
"routing-instance-name (Static Subscribers)" on page 15949	"routing-instances" on page 15969	"routing-instances (CoS)" on page 15971
"routing-instances (Dynamic Profiles)" on page 15973	"routing-instances (JDAF)" on page 15967	"routing-instances (Logical Systems)" on page 15966

"routing-instances (Multiple Routing Entities)" on page 15976	"routing-instances (RPM)" on page 15970	"routing-interface" on page 15978
"routing-options" on page 15980	"routing-options (Dynamic Profiles)" on page 15981	"routing-options (Junos Node Slicing)" on page 15983
"routing-options (QFabric System)" on page 15987	"routing-options-max-ecmp" on page 15985	"routing-service (Dynamic Profiles)" on page 15988
"routing-services (Enhanced Subscriber Management)" on page 15991	"rp" on page 15996	"rp-register-policy" on page 16000
"rp-set" on page 16001	"rpc" on page 16003	"rpc-program-number" on page 16004
"rpc-program-number (Applications)" on page 16005	"rpf-check " on page 16007	"rpf-check (Dynamic Profiles)" on page 16009
"rpf-check-policy (Routing Options)" on page 16011	"rpf-check-policy (Routing Options RPF)" on page 16012	"rpf-loose-mode-discard" on page 16014
"rpf-selection" on page 16015	"rpf-vector (PIM)" on page 16017	"rpm (Interfaces) " on page 16019
"rpm (Services)" on page 16023	"rpm-scale" on page 16020	"rpm-tracking" on page 16038
"rpt-spt" on page 16040	"rsh" on page 16041	"rst-invalidate-session" on page 16043
"rst-sequence-check" on page 16044	"rstp" on page 16046	"rsvp" on page 16050
"rsvp-error-hold-time" on page 16056	"rsvp-te (Protocols MVPN)" on page 16058	"rsvp-te (Routing Instances Provider Tunnel)" on page 16059
"rsvp-te (Routing Instances Provider Tunnel Selective)" on page 16061	"rsvp-te (Routing Options)" on page 16062	"rtf-prefix-list" on page 16064

"rtlog (Next Gen Services Global System Logging)" on page 16066	"rtp" on page 16068	"rts" on page 16069
"rts-polarity" on page 16071	"rtsp" on page 16072	"rule" on page 16074
"rule (AACL Rule Set) " on page 16076	"rule (Advanced Policy-Based Routing)" on page 16080	"rule (Application Firewall)" on page 16083
"rule (Captive Portal Content Delivery)" on page 16085	"rule (casb-policy)" on page 16087	"rule (Destination NAT Next Gen Services)" on page 16091
"rule (HTTP Header Enrichment for Tag Rule Set)" on page 16093	"rule (IDS MS-DPC)" on page 16094	"rule (IDS MS-MPC)" on page 16096
"rule (IP Reassembly)" on page 16100	"rule (LRF)" on page 16101	"rule (PCP)" on page 16103
"rule (Security Destination NAT)" on page 16105	"rule (Security Exempt Rulebase)" on page 16107	"rule (Security IPS Rulebase)" on page 16109
"rule (Security Source NAT)" on page 16111	"rule (Security Static NAT)" on page 16113	"rule (Services CoS) " on page 16115
"rule (Services CoS Next Gen Services)" on page 16117	"rule (Services NAT) " on page 16120	"rule (Services Stateful Firewall) " on page 16122
"rule (Softwire)" on page 16154	"rule (Source NAT Next Gen Services)" on page 16156	"rule-activation-time" on page 16078
"rule-deactivation-time" on page 16089	"rule-session-count-alarm (Security Destination NAT Rule Set)" on page 16124	"rule-session-count-alarm (Security Source NAT Rule Set)" on page 16125
"rule-session-count-alarm (Security Static NAT Rule Set)" on page 16126	"rule-set" on page 16128	"rule-set (Captive Portal Content Delivery)" on page 16129

"rule-set (Security Destination NAT)" on page 16131	"rule-set (Security Source NAT)" on page 16133	"rule-set (Security Static NAT)" on page 16135
"rule-set (Services AACL)" on page 16138	"rule-set (Services CoS)" on page 16139	"rule-set (Services CoS Next Gen Services)" on page 16140
"rule-set (Services IDS)" on page 16142	"rule-set (Services NAT) " on page 16143	"rule-set (Services Stateful Firewall) " on page 16144
"rule-set (Softwire)" on page 16146	"rule-set (Softwires Next Gen Services)" on page 16147	"rule-sets (CoS AppQoS)" on page 16149
"rule-sets (Security Application Firewall)" on page 16151	"rulebase-exempt" on page 16157	"rulebase-ips" on page 16159
"run-length" on page 16163	"rx-buffers" on page 16164	"rx-connect-speed-when-equal (L2TP LAC)" on page 16166
"rx-enable" on page 16168	"rx-max-duration" on page 16169	"rx-window-size (L2TP)" on page 16171

S

"s6a" on page 16260	"s-nssai" on page 16259	"sa-hold-time (Protocols MSDP)" on page 16262
"sa-multicast (100-Gigabit Ethernet)" on page 16264	"sa-multicast (PTX Series Packet Transport Routers)" on page 16266	"sabit" on page 16267
"saegw" on page 16269	"saegw access-network-peers" on page 16270	"saegw control-plane-peers" on page 16272
"saegw system" on page 16278	"saegw-core-network-peers" on page 16276	"sak-rekey-interval " on page 16280

"saml" on page 16281	"saml-options" on page 16284	"sample" on page 16285
"sample-once" on page 16287	"sample-rate" on page 16288	"sample-type" on page 16290
"sampling" on page 16291	"sampling (Forwarding Options)" on page 16293	"sampling (Interfaces)" on page 16298
"sampling (Protocols MPLS)" on page 16300	"sampling-instance" on page 16297	"sanity-poll" on page 16302
"sap" on page 16303	"satellite (Junos Fusion Automatic Satellite Conversion)" on page 16305	"satellite (Junos Fusion Satellite Device Homing)" on page 16307
"satellite (Junos Fusion Satellite Software Upgrade Groups)" on page 16308	"satellite-management (Junos Fusion)" on page 16310	"satellite-policies" on page 16313
"save (Configuration Mode)" on page 16317	"saved-core-context (System)" on page 16320	"saved-core-files (System)" on page 16321
"sbfd" on page 16323	"sbfd (protocols bfd)" on page 16324	"sbl" on page 16326
"sbl-default-server" on page 16327	"scale-mode" on page 16328	"scale-mode" on page 16328
"scale-optimized" on page 16329	"scan-extension" on page 16332	"scan-mode" on page 16333
"scan-options (Security Antivirus Avira Engine)" on page 16335	"scan-options (Security Antivirus Juniper Express Engine)" on page 16336	"scan-options (Security Antivirus Kaspersky Lab Engine)" on page 16338
"scan-options (Security Antivirus Sophos Engine)" on page 16340	"sccp" on page 16341	"scheduler" on page 16343
"scheduler (Dynamic Scheduler Maps)" on page 16344	"scheduler (Fabric Queues)" on page 16345	"scheduler (Interconnect Device)" on page 16347

"scheduler-map (Interfaces and Traffic-Control Profiles)" on page 16355	"scheduler (Scheduler Map)" on page 16362	"scheduler (Security Policies)" on page 16364
"scheduler (System Security Profile)" on page 16366	"scheduler-map (CoS Virtual Channels)" on page 16350	"scheduler-map (Dynamic Traffic Shaping)" on page 16351
"scheduler-map (Fabric Queues)" on page 16353	"scheduler-map (Interfaces and Traffic-Control Profiles)" on page 16357	"scheduler-map-chassis" on page 16348
"scheduler-map-forwarding-class-sets" on page 16354	"scheduler-maps (Definition)" on page 16358	"scheduler-maps (Dynamic CoS Definition)" on page 16360
"scheduler-name" on page 16361	"schedulers (CoS)" on page 16368	"schedulers (CoS)" on page 16369
"schedulers (Dynamic CoS Definition)" on page 16371	"schedulers (Interfaces)" on page 16373	"schedulers (Security Policies)" on page 16375
"schema" on page 16376	"schema-version (Flat-File Accounting Options)" on page 16378	"scope" on page 16379
"scope (Security IDP Chain Attack)" on page 16382	"scope (Security IDP Custom Attack)" on page 16383	"scope-policy" on page 16380
"screen " on page 16384	"screen (Security Zones)" on page 16390	"scripts" on page 16392
"sctp" on page 16396	"sdsl-bytes" on page 16398	"sdsl-overhead-adjust" on page 16400
"search" on page 16402	"search-filter" on page 16404	"secondary (Adaptive Services Interfaces)" on page 16408
"secondary connection (Identity Management Advanced Query) " on page 16409	"secondary (Label Switched Path) " on page 16405	"secondary (Link Services IQ PIC Interfaces) " on page 16414

"secondary (Precision Time Protocol) " on page 16406	"secondary (Protocols MPLS)" on page 16416	"secondary-dns" on page 16412
"secondary-interface (Aggregated Inline Services)" on page 16413	"secondary-server" on page 16418	"secondary-vtep-address" on page 16419
"secondary-wins" on page 16420	"secret (Access Profile)" on page 16422	"secret (RADIUS)" on page 16424
"secret (RADIUS Client)" on page 16426	"secret (RADIUS Server)" on page 16427	"secret (Tunnel Profile)" on page 16429
"secret-key-timeout" on page 16423	"secure" on page 16430	"secure-access-port" on page 16432
"secure-channel" on page 16435	"secure-nat-mapping" on page 16436	"secure-nat-mapping (Source NAT Next Gen Services)" on page 16438
"secure-neighbor-discovery" on page 16439	"secure-proxy" on page 16441	"secure-wire (System Security Profile)" on page 16443
"secured-port-block-allocation" on page 16444	"securid-server" on page 16447	"security" on page 16448
"security-association" on page 16452	"security-association (IPSec)" on page 16454	"security-association (Junos OS)" on page 16460
"security-association (Junos-FIPS Software)" on page 16458	"security-association (MACSec)" on page 16450	"security-intelligence" on page 16464
"security-intelligence (IDP)" on page 16465	"security-intelligence (Services)" on page 16462	"security-intelligence-policy" on page 16467
"security-intelligence(services)" on page 16469	"security-level" on page 16480	"security-level (Defining Access Privileges)" on page 16481
"security-level (Generating SNMP Notifications)" on page 16482	"security-log" on page 16484	"security-log-percent-full" on page 16485

"security-metadata-streaming" on page 16487	"security-mode" on page 16494	"security-mode (SRX)" on page 16492
"security-model (Access Privileges)" on page 16496	"security-model (Group)" on page 16497	"security-model (SNMP Notifications)" on page 16499
"security-name (Community String)" on page 16500	"security-name (Security Group)" on page 16501	"security-name (SNMP Notifications)" on page 16503
"security-option (IDS Screen Next Gen Services)" on page 16504	"security-package" on page 16505	"security-profile" on page 16507
"security-profile-resources" on page 16512	"security-service (Security Forwarding Options)" on page 16513	"security-to-group" on page 16515
"security-zone" on page 16516	"segment" on page 16519	"segment-list" on page 16521
"select" on page 16525	"select-profile" on page 16527	"selection-mode" on page 16528
"selective" on page 16529	"selective-vlan-switching" on page 16532	"send (add-path)" on page 16535
"send (Protocols RIP)" on page 16546	"send (Protocols RIPng)" on page 16548	"send-acct-status-on-config-change (Access Profile)" on page 16534
"send-critical-event" on page 16539	"send-dyn-subscription-indicator (PCRF Partition)" on page 16540	"send-network-family-indicator (PCRF Partition)" on page 16542
"send-oam" on page 16544	"send-origin-state-id (Diameter Applications)" on page 16545	"send-release-on-delete (DHCP Relay Agent)" on page 16550
"sender-based-rpf (MBGP MVPN)" on page 16552	"sendid-tlv" on page 16554	"sensor (Junos Telemetry Interface)" on page 16564
"sensor-based-stats" on page 16555	"sensor-based-stats (Junos Telemetry Interface)" on page 16558	"sensor-configuration" on page 16560

"separator" on page 16593	"seq-number-validated" on page 16594	"sequence-check-required" on page 16595
"sequence-number (Security IDP ICMP Headers)" on page 16597	"sequence-number (Security IDP TCP Headers)" on page 16598	"serial-number" on page 16599
"serial-number (Junos Fusion)" on page 16601	"serial-number (MX Series Virtual Chassis)" on page 16603	"serial-options" on page 16604
"server" on page 16608	"server (Batch Commits)" on page 16613	"server (DHCP and BOOTP Relay Agent)" on page 16623
"server (DNS, Port, and TFTP Service)" on page 16624	"server (Dynamic PPPoE)" on page 16627	"server (Easy EVPN LAG Configuration)" on page 16628
"server (icap-redirect profile)" on page 16635	"server (Junos node slicing)" on page 16606	"server (Junos OS)" on page 16642
"server (Junos OS Evolved)" on page 16644	"server (NTP)" on page 16648	"server (pcp)" on page 16651
"server (RADIUS Accounting)" on page 16656	"server (RADIUS Network Elements)" on page 16660	"server (Security Antivirus)" on page 16665
"server (Security Group VPN)" on page 16667	"server (Security Sophos Engine Antivirus)" on page 16670	"server (Security Web Filtering)" on page 16672
"server (TACACS+ Accounting)" on page 16674	"server-address" on page 16609	"server-address (dhcp-client)" on page 16611
"server-address (Hosted Services)" on page 16612	"server-certificate (REST API)" on page 16615	"server-certificate (Services)" on page 16617
"server-certificate-subject" on page 16618	"server-cluster (Security Group VPN Server)" on page 16619	"server-connectivity" on page 16621

"server-duid-type (DHCP Local Server)" on page 16626	"server-group" on page 16633	"server-id-override" on page 16637
"server-identifier" on page 16639	"server-inactivity-timeout" on page 16641	"server-member-communication (Security Group VPN Server)" on page 16646
"server-profile" on page 16654	"server-profile (Active Flow Monitoring)" on page 16655	"server-reject-bridge-domain | server-reject-vlan" on page 16661
"server-response-time (DHCP Relay Agent)" on page 16664	"servers" on page 16677	"servers (RADIUS)" on page 16678
"service" on page 16680	"service (Dynamic Profiles)" on page 16694	"service (Dynamic Service Sets)" on page 16695
"service (PPPoE)" on page 16727	"service (Security IDP Anomaly Attack)" on page 16736	"service (Security IDP Dynamic Attack Group)" on page 16737
"service (Service Accounting)" on page 16739	"service-agents (Analytics)" on page 16681	"service-context-id (OCS)" on page 16684
"service-deployment" on page 16686	"service-device (Remote Device Management)" on page 16690	"service-device-pool (L2TP)" on page 16687
"service-device-pools (Service Interfaces)" on page 16688	"service-domain" on page 16693	"service-filter (Dynamic Service Sets)" on page 16697
"service-filter (Firewall)" on page 16699	"service-filter (Interfaces)" on page 16701	"service-groups" on page 16702
"service-id (Bridge Domains)" on page 16704	"service-id (Switch Options)" on page 16705	"service-interface (L2TP Processing) " on page 16706
"service-interface (Services Interfaces) " on page 16709	"service-interface-pools" on page 16708	"service-mode" on page 16711

"service-mode (mobile-edge SAEGW)" on page 16713	"service-name" on page 16714	"service-name-table" on page 16716
"service-name-tables" on page 16717	"service-order" on page 16719	"service-package" on page 16721
"service-pics" on page 16723	"service-port (Inline Services) " on page 16725	"service-port (UDP, Dynamic Flow Capture)" on page 16724
"service-profile (DHCP Local Server)" on page 16729	"service-profile (DHCP Relay Agent)" on page 16731	"service-profile (L2TP)" on page 16733
"service-profile (Static Subscribers)" on page 16735	"service-provider (SAML)" on page 16767	"service-set (Application-Aware Control Policy)" on page 16740
"service-set (Dynamic Service Sets)" on page 16742	"service-set (Interfaces)" on page 16744	"service-set (Services)" on page 16750
"service-set (Subscriber-Aware)" on page 16754	"service-set (TDF Interface)" on page 16756	"service-set-options" on page 16746
"service-set-options (Next Gen Services Services)" on page 16748	"service-tracking (Network Isolation Groups)" on page 16760	"service-tracking-action (Network Isolation Groups)" on page 16758
"service-type" on page 16761	"service-type (MAC-VRF Routing Instances)" on page 16764	"service-type (RFC 2544 Benchmarking Tests)" on page 16763
"services (Captive Portal Content Delivery)" on page 16769	"services (CoS)" on page 16772	"services (Hierarchy)" on page 16773
"services (IDS)" on page 16774	"services (Interfaces)" on page 16776	"services (IPsec VPN)" on page 16777
"services (L2TP)" on page 16778	"services (L2TP System Logging)" on page 16780	"services (NAT)" on page 16783

"services (Stateful Firewall)" on page 16799	"services (System Logging)" on page 16800	"services (system services)" on page 16802
"services-memory (resource-manager)" on page 16781	"services-offload" on page 16784	"services-offload (Security)" on page 16786
"services-options" on page 16787	"services-options (Next Gen Services Interfaces)" on page 16789	"services-redundancy-group" on page 16793
"session (Origin Validation for BGP)" on page 16836	"session (Protocols LDP)" on page 16845	"session (Security IKE)" on page 16848
"session (System Services)" on page 16853	"session (Web Management)" on page 16857	"session-affinity" on page 16804
"session-close" on page 16806	"session-drop-hold-down" on page 16807	"session-establishment-hold-time" on page 16808
"session-expiry (MX Series in Enhanced LAN Mode)" on page 16809	"session-group" on page 16810	"session-id-change-limiter-indirect" on page 16813
"session-init" on page 16815	"session-limit" on page 16816	"session-limit (IDS MS-DPC)" on page 16817
"session-limit (IDS MS-MPC)" on page 16819	"session-limit (Service Set Next Gen Services)" on page 16823	"session-limit (System Services)" on page 16824
"session-limit-per-username (Access Profile)" on page 16821	"session-memory (resource-manager)" on page 16826	"session-mode" on page 16828
"session-offload" on page 16829	"session-options" on page 16831	"session-options (Access Profile)" on page 16835
"session-persistence-scan" on page 16839	"session-pics" on page 16840	"session-pics (Diameter)" on page 16842

"session-protection" on page 16843	"session-scan" on page 16847	"session-timeout (Service Set Next Gen Services)" on page 16854
"session-update-interval" on page 16856	"sessions-limit-group (L2TP)" on page 16858	"set-dont-fragment-bit (Services IPsec VPN)" on page 16890
"set-dont-fragment-bit (Services Set)" on page 16891	"identity-provider settings (SAML)" on page 16895	"setup-protection" on page 16897
"severity (Event Policy)" on page 16898	"severity (Next Gen Services Service-Set Remote System Logging)" on page 16900	"severity (Security IDP Custom Attack)" on page 16901
"severity (Security IDP Dynamic Attack Group)" on page 16903	"severity (Security IDP IPS Rulebase)" on page 16904	"severity (Security Log)" on page 16906
"sflow" on page 16907	"sfm (Chassis)" on page 16912	"sfm (Chassis Redundancy)" on page 16913
"sfpplus" on page 16915	"sftp-backup (OCS Partition)" on page 16916	"sglimit" on page 16919
"sham-link" on page 16921	"sham-link-remote" on page 16923	"shaping" on page 16925
"shaping-rate (Applying to an Interface)" on page 16929	"shaping-rate (CoS Adaptive Shapers)" on page 16932	"shaping-rate (CoS Interfaces)" on page 16933
"shaping-rate (CoS Virtual Channels)" on page 16936	"shaping-rate (Dynamic Traffic Shaping and Scheduling)" on page 16937	"shaping-rate (Oversubscribing an Interface)" on page 16947
"shaping-rate (Schedulers)" on page 16960	"shaping-rate (Traffic Control Profiles)" on page 16926	"shaping-rate-excess-high" on page 16939
"shaping-rate-excess-low" on page 16941	"shaping-rate-excess-medium-high" on page 16943	"shaping-rate-excess-medium-low" on page 16945

"shaping-rate-priority-high" on page 16950	"shaping-rate-priority-low" on page 16952	"shaping-rate-priority-medium" on page 16954
"shaping-rate-priority-medium-low" on page 16956	"shaping-rate-priority-strict-high" on page 16958	"shared-bandwidth-policer (Configuring)" on page 16962
"shared-buffer" on page 16964	"shared-instance" on page 16967	"shared-key" on page 16968
"shared-name" on page 16970	"shared-scheduler" on page 16971	"shared-secret (RADIUS Snoop Segment)" on page 16973
"shdsl-options" on page 16974	"shell" on page 16975	"shellcode" on page 16977
"shmlog (Shared Memory Log)" on page 16978	"short-cycle-protection (DHCP Local Server and Relay Agent)" on page 16982	"short-cycle-protection (Static and Dynamic Subscribers)" on page 16984
"short-name-format" on page 16986	"short-reach-mode" on page 16988	"short-sequence" on page 16989
"shortcuts (Protocols OSPF)" on page 16991	"should-secure" on page 16992	"shutdown (BPDU Block)" on page 27935
"shutdown (Protocols BGP)" on page 27936	"shutdown-threshold" on page 27937	"sib" on page 27939
"signal-bandwidth" on page 27941	"signal-degrade" on page 27942	"signal-degrade-monitor-enable" on page 27944
"signal-type" on page 27946	"signal-type (BITS Interfaces)" on page 27947	"signaling" on page 27949
"signature" on page 27950	"signature (Application Identification)" on page 27952	"Signature Hash Algorithm (Security IKE)" on page 27953
"signature (Security IDP)" on page 27955	"simple-filter" on page 27962	"simple-filter (Applying to an Interface)" on page 27964

"simple-filter (Interfaces)" on page 27965	"simple-mail-client-service" on page 27967	"single-home (Junos Fusion)" on page 27968
"single-rate" on page 27969	"sip" on page 27971	"sip (Application Profile)" on page 27973
"sip (Security) " on page 27977	"sip (Services CoS Next Gen Services)" on page 27981	"sip-call-hold-timeout" on page 27974
"sip-password" on page 27976	"sip-server" on page 27979	"sip-user-id" on page 27983
"site (Layer 2 Circuits)" on page 27987	"site (VPLS Multihoming for FEC 128)" on page 27994	"site (VPLS Multihoming for FEC 129)" on page 27995
"site-identifier (Layer 2 Circuits)" on page 27984	"site-identifier (VPLS)" on page 27985	"site-preference" on page 27989
"site-range" on page 27990	"site-reputation-action" on page 27992	"size" on page 27997
"size (Interface)" on page 28000	"size (Next Gen Services Global System Logging)" on page 28001	"size (Sampling and Traceoptions)" on page 28003
"size (Security Web Filtering Cache)" on page 28004	"size (Services)" on page 28006	"size (System)" on page 28007
"size (VNF)" on page 27999	"skew-timer-disable" on page 28009	"skip-arp-iteration (RFC 2544 Benchmarking)" on page 28010
"sla-iterator-profile" on page 28012	"sla-iterator-profiles" on page 28013	"sla-iterator-profiles (OAM CFM)" on page 28015
"sla-options" on page 28017	"sla-rule" on page 28019	"sla-tracking" on page 28027
"slaac-snooping" on page 28030	"slamon-services" on page 28032	"slave" on page 28034

"slice (CoS Interfaces)" on page 28039	"slice (export routing policy action)" on page 28041	"slice (firewall filter action)" on page 28042
"slice (firewall filter match condition)" on page 28043	"slow-pfe-alarm" on page 28044	"smart-check" on page 28045
"smart-optimize-timer" on page 28048	"smg-service (Enhanced Subscriber Management)" on page 28050	"smtp-profile (Security Content Security Policy Antispam)" on page 28051
"smtp-profile (Security Content Security Policy Antivirus)" on page 28052	"smtp-profile (Security Content Security Policy Content Filtering)" on page 28053	"snapshot" on page 28055
"snext" on page 28056	"snmp" on page 28058	"snmp (SNMP Scripts)" on page 28068
"snmp-command" on page 28064	"snmp-community" on page 28065	"snmp-options" on page 28067
"snmp-trap-thresholds" on page 28070	"snmp-trap-thresholds (Next Gen Services)" on page 28072	"snmp-value-match-msmic (Services NAT Options)" on page 28074
"snoop-pseudowires" on page 28075	"snoop-segment (TDF Domain Selection)" on page 28076	"snoop-segments (RADIUS)" on page 28078
"snoop-segments (TDF Gateway)" on page 28080	"snr-margin" on page 28081	"sockets" on page 28082
"soft-gre" on page 28083	"soft-limit" on page 28085	"soft-limit-clear" on page 28087
"soft-preemption (Protocols MPLS)" on page 28088	"soft-preemption (Protocols RSVP)" on page 28089	"software-upgrade" on page 28090
"softwire-concentrator" on page 28092	"softwire-name (Next Gen Services)" on page 28093	"softwire-options" on page 28095

"softwire-rules" on page 28097	"softwire-types (Next Gen Services)" on page 28098	"softwires" on page 28102
"softwires (Next Gen Services)" on page 28104	"softwires-rule-set (Service Set Next Gen Services)" on page 28105	"solicit-router-advertisement-unicast" on page 28106
"sonet" on page 28108	"sophos-engine" on page 28109	"source (Application Identification)" on page 28182
"source (Bridge Domains)" on page 28187	"source (Chassis Synchronization)" on page 28188	"source (Commit Scripts)" on page 28194
"source (Distributed IGMP)" on page 28200	"source (Dynamic IGMP Interface)" on page 28201	"source (Dynamic MLD Interface)" on page 28203
"source (Event Scripts)" on page 28206	"source (GRE Encapsulation) " on page 28111	"source (Interfaces)" on page 28112
"source interfaces" on page 28231	"source (JET Scripts)" on page 28245	"source (Multicast VLAN Registration)" on page 28252
"source (Op Scripts)" on page 28255	"source (PIM RPF Selection)" on page 28275	"source (Protocols IGMP)" on page 28305
"source (Protocols MLD)" on page 28307	"source (Protocols MSDP) " on page 28308	"source (Routing Instances)" on page 28312
"source (Routing Instances Provider Tunnel Selective)" on page 28314	"source (Security Source NAT)" on page 28322	"source (SNMP Scripts)" on page 28326
"source (Source-Specific Multicast)" on page 28327	"source-active-advertisement" on page 28114	"source-address " on page 28120
"source-address (AACL)" on page 28123	"source-address (Access LDAP)" on page 28124	"source-address (Access RADIUS)" on page 28125

"source-address (Content Security)" on page 28121	"source-address (Dynamic Tunnel)" on page 28117	"source-address (Forwarding Options)" on page 28133
"source-address (IDS MS-DPC)" on page 28136	"source-address (IGMP and MLD Snooping)" on page 28118	"source-address (LRF Profile)" on page 28138
"source-address (NAT Next Gen Services)" on page 28144	"source-address (Next Gen Services Service-Set Remote System Logging)" on page 28145	"source-address (PCP)" on page 28146
"source-address (RADIUS, System Logging, or TACACS+)" on page 28148	"source-address (Routing Options)" on page 28156	"source-address (RPM)" on page 28157
"source-address (Security Destination NAT)" on page 28160	"source-address (Security IDP)" on page 28161	"source-address (Security IDP Policy)" on page 28162
"source-address (Security IDP Sensor Configuration)" on page 28163	"source-address (Security Policies)" on page 28164	"source-address (Security Source NAT)" on page 28166
"source-address (Security Static NAT Rule Set)" on page 28167	"source-address (Service Sets) " on page 28168	"source-address (Services CoS)" on page 28169
"source-address (Services IPSec VPN)" on page 28115	"source-address (Services NAT)" on page 28171	"source-address (Services Stateful Firewall) " on page 28172
"source-address (SNMP)" on page 28174	"source-address (SRC Software)" on page 28175	"source-address (Subscriber Secure Policy)" on page 28176
"source-address (System Logging)" on page 28177	"source-address (TWAMP)" on page 28179	"source-address-excluded" on page 28127
"source-address-filter" on page 28128	"source-address-filter (Interfaces)" on page 28130	"source-address-filter-list" on page 28131

"source-address-giaddr" on page 28135	"source-address-name (NAT Next Gen Services)" on page 28139	"source-address-name (Security Destination NAT)" on page 28140
"source-address-name (Security Source NAT)" on page 28141	"source-address-name (Security Static NAT Rule Set)" on page 28142	"source-address-range" on page 28149
"source-address-range (IDS MS-DPC)" on page 28150	"source-address-range (PCP)" on page 28152	"source-address-range (Services NAT) " on page 28153
"source-address-range (Services Stateful Firewall) " on page 28155	"source-addresses" on page 28181	"source-attachment-identifier (Protocols VPWS)" on page 28183
"source-based-fpolicy-statementorwarding" on page 28333	"source-bmac" on page 28185	"source-checking" on page 28190
"source-class-usage" on page 28191	"source-classes" on page 28193	"source-count (Dynamic MLD Interface)" on page 28196
"source-count (Protocols IGMP)" on page 28197	"source-count (Protocols MLD)" on page 28199	"source-end-user-profile" on page 28204
"source-except" on page 28207	"source-filtering" on page 28208	"source-filtering (Interfaces)" on page 28210
"source-gateway (Tunnel Profile)" on page 28211	"source-hostname" on page 28213	"source-id" on page 28214
"source-identity" on page 28216	"source-identity-log (Security)" on page 28218	"source-increment (Dynamic MLD Interface)" on page 28220
"source-increment (Protocols IGMP)" on page 28221	"source-increment (Protocols MLD)" on page 28223	"source-interface" on page 28225
"source-interface (RADIUS Client)" on page 28224	"source-interface (RADIUS Server)" on page 28227	"source-interface (RADIUS Snoop Segment)" on page 28228

"source-interface-set-at-login" on page 28229	"source-ip" on page 28233	"source-ip (Flow Monitoring Logs for NAT)" on page 28240
"source-ip (Security Alarms)" on page 28242	"source-ip-address (RADIUS Snoop Segment)" on page 28237	"source-ip-address-list" on page 28235
"source-ip-based" on page 28238	"source-ip-change (Forwarding Options)" on page 28239	"source-ipv4-address (RFC 2544 Benchmarking)" on page 28244
"source-lookup" on page 28338	"source-mac-address" on page 28247	"source-mac-address-list" on page 28249
"source-mode" on page 28250	"source-nat" on page 28254	"source-packet-routing (Junos Telemetry Interface)" on page 28261
"source-packet-routing (Protocols IS-IS)" on page 28266	"source-packet-routing (Protocols OSPF)" on page 28270	"source-packet-routing (Routing Instances Protocols EVPN)" on page 28273
"source-packet-routing (routing-options)" on page 28264	"source-packet-routing (Segment Routing Traffic Engineering)" on page 28257	"source-pool" on page 28276
"source-port" on page 28277	"source-port (Applications)" on page 28280	"source-port (firewall)" on page 28278
"source-port (Security Forwarding Options)" on page 28281	"source-port (Security IDP)" on page 28286	"source-port (Security Source NAT Rule Set)" on page 28288
"source-port (Security Static NAT Rule Set)" on page 28289	"source-port (Subscriber Secure Policy)" on page 28290	"source-prefix (IDS)" on page 28292
"source-prefix (Security Forwarding Options)" on page 28303	"source-prefix (Services NAT) " on page 28304	"source-prefix-ipv6 (IDS)" on page 28293
"source-prefix-list (PCP)" on page 28295	"source-prefix-list (Services AACL)" on page 28296	"source-prefix-list (Services CoS) " on page 28297

"source-prefix-list (Services IDS)" on page 28299	"source-prefix-list (Services NAT) " on page 28300	"source-prefix-list (Services Stateful Firewall) " on page 28301
"source-route-option (IDS Screen Next Gen Services)" on page 28310	"source-routing" on page 28339	"source-routing" on page 28311
"source-routing-path" on page 28316	"source-routing-path-template" on page 28319	"source-threshold" on page 28329
"source-tree (MBGP MVPN)" on page 28330	"source-udp-port (RFC 2544 Benchmarking)" on page 28332	"spam-action" on page 28341
"speed (24-port and 12-port 10 Gigabit Ethernet PIC)" on page 28342	"speed" on page 28344	"speed (Chassis Cluster)" on page 28347
"speed (Ethernet)" on page 28349	"speed (Fibre Channel Interfaces)" on page 28370	"speed (Gigabit Ethernet interface)" on page 28372
"speed (Gigabit Ethernet interface - MX Series)" on page 28373	"speed (Interfaces)" on page 28376	"speed (MX Series DPC)" on page 28377
"spf-options (Protocols IS-IS)" on page 28379	"spf-options (Protocols OSPF)" on page 28381	"spi" on page 28384
"spi (Junos OS)" on page 28387	"spi (Junos-FIPS Software)" on page 28386	"spid1" on page 28388
"spid2" on page 28389	"splitting-merging" on page 28391	"spring-te" on page 28393
"spring-te (Dynamic Tunnels)" on page 28395	"spring-te-lsp-next-hop" on page 28397	"spring-traffic-engineering-edit-protocols" on page 28398
"spt-only" on page 28400	"spt-threshold" on page 28401	"sql" on page 28402
"squelch" on page 28404	"srgb (Protocols BGP)" on page 28407	"srgb (Protocols IS-IS)" on page 28408

"srgb-label-range" on page 28410	"srlg" on page 28411	"srlg-cost" on page 28413
"srlg-value" on page 28414	"srv6 (Protocols BGP)" on page 28417	"srv6 (Routing Options)" on page 28415
"srv6-sid" on page 28419	"ssb" on page 28421	"ssh (NETCONF)" on page 28428
"ssh (System Services)" on page 28430	"ssh-known-hosts" on page 28426	"ssl" on page 28439
"ssl (Services)" on page 28447	"ssl-encryption" on page 28442	"ssl-inspection" on page 28443
"ssl-proxy (Application Services)" on page 28446	"ssl-termination-profile" on page 28450	"ssm-groups" on page 28452
"ssm-map (Dynamic IGMP Interface)" on page 28453	"ssm-map (Dynamic MLD Interface)" on page 28455	"ssm-map (Protocols IGMP)" on page 28464
"ssm-map (Protocols IGMP AMT)" on page 28465	"ssm-map (Protocols MLD)" on page 28466	"ssm-map (Routing Options Multicast)" on page 28468
"ssm-map-policy (Dynamic IGMP Interface)" on page 28458	"ssm-map-policy (Dynamic MLD Interface)" on page 28459	"ssm-map-policy (IGMP)" on page 28461
"ssm-map-policy (MLD)" on page 28462	"stacked-interface-set (Dynamic Profiles)" on page 28469	"stacked-vlan-ranges" on page 28471
"stacked-vlan-ranges (RADIUS Options)" on page 28473	"stacked-vlan-tagging" on page 28475	"stale-labels-holddown-period" on page 28476
"stale-routes-time" on page 28478	"stale-routes-time (Fabric Control)" on page 28479	"stale-time (Long-Lived Graceful Restart for BGP Restarter)" on page 28481
"stamp (Sampling)" on page 28484	"standby" on page 28485	"standby" on page 28486

"standby (Protocols Layer 2 Circuit)" on page 28489	"standby (Routing Policy Condition)" on page 28491	"standby-path-creation-delay" on page 28488
"start (Scripts)" on page 28496	"start-date" on page 28492	"start-log" on page 28493
"start-measurement" on page 28494	"start-time" on page 28500	"start-time (Accounting)" on page 28502
"start-time (Schedulers)" on page 28503	"start-time (Security IDP)" on page 28504	"starts-with (DHCP Relay Agent Option)" on page 28506
"starts-with (Event Policy)" on page 28510	"startup-alarm" on page 28511	"startup-silent-period" on page 28513
"stateful (Precision Time Protocol) " on page 28516	"stateful (Router Advertisement Guard) " on page 28514	"stateful-firewall-rule-set (Next Gen Services)" on page 28518
"stateful-firewall-rule-sets (Service Set Next Gen Services)" on page 28519	"stateful-firewall-rules" on page 28520	"stateful-firewall-rules (Service Set Next Gen Services)" on page 28521
"stateful-nat64" on page 28523	"stateless" on page 28524	"static (802.1X)" on page 28526
"static (Bridge Domains)" on page 28529	"static (Distributed IGMP)" on page 28531	"static (Dynamic IGMP Interface)" on page 28532
"static (Dynamic MLD Interface)" on page 28534	"static (IGMP Snooping)" on page 28537	"static (Origin Validation for BGP)" on page 28553
"static (Protocols IGMP)" on page 28557	"static (Protocols Layer 2 Circuit)" on page 28559	"static (Protocols MLD)" on page 28560
"static (Protocols PIM)" on page 28562	"static (Protocols VPLS)" on page 28564	"static (Routing Options)" on page 28570
"static (Security NAT)" on page 28577	"static-binding" on page 28528	"static-host-mapping" on page 28535

"static-interface" on page 28539	"static-ip" on page 28540	"static-ip (MX Series)" on page 28542
"static-ipv6" on page 28544	"static-label-switched-path" on page 28545	"static-lsp" on page 28547
"static-mac" on page 28549	"static-mac-ip" on page 28587	"static-multihoming-peer" on page 28589
"static-nat " on page 28551	"static-policy-control" on page 28555	"static-remote-vtep-list" on page 28566
"static-subscriber-application" on page 28579	"static-subscribers" on page 28581	"static-subscribers (Dynamic Service Provisioning)" on page 28582
"static-tei-val" on page 28584	"static-umh (MBGP MVPN)" on page 28585	"statistics" on page 28591
"statistics (Access Profile)" on page 28592	"statistics (Protocols MPLS)" on page 28594	"statistics (Service Accounting)" on page 28598
"statistics (Services)" on page 28599	"statistics-service" on page 28596	"stats-aggregation-interval (programmable-rpd)" on page 28601
"stats-sampling-interval (programmable-rpd)" on page 28602	"status-control" on page 28604	"steering" on page 28605
"step-percent (RFC 2544 Benchmarking)" on page 28607	"stickydr" on page 28609	"stop-date" on page 28610
"stop-time" on page 28612	"storage" on page 28613	"store" on page 28615
"storm-control" on page 28620	"storm-control (switch)" on page 28622	"storm-control (Video Monitoring)" on page 28619
"storm-control-profiles" on page 28624	"stp" on page 28625	"stream (Logical Systems and Tenant Systems)" on page 28628

"stream (Next Gen Services Service-Set Remote System Logging)" on page 28631	"stream (Security Log)" on page 28635	"stream-option (IDS Screen Next Gen Services)" on page 28633
"stream-protection (Multicast-Only Fast Reroute)" on page 28634	"streaming-server (Junos Telemetry Interface)" on page 28640	"streaming-servers" on page 28642
"strict (DHCP Local Server)" on page 28646	"strict-asla-based-flex-algorithm" on page 28645	"strict-dual-isis (Protocols IS-IS)" on page 28648
"strict-priority-scheduler" on page 28650	"strict-source-route-option (IDS Screen Next Gen Services)" on page 28652	"strict-syn-check" on page 28653
"strict-targeted-hellos" on page 28654	"string" on page 28655	"strip-as-path" on page 28662
"strip-domain (Domain Map)" on page 28657	"strip-nexthop" on page 28658	"strip-user-name (Access Profile)" on page 28659
"strip-username (Domain Map)" on page 28661	"structured-data" on page 28664	"stub" on page 28665
"stub-network" on page 28667	"sub-domain" on page 28668	"sub-prefix-length" on page 28674
"subnet-address (Easy EVPN LAG Configuration)" on page 28675	"subscriber (Access Profile)" on page 28680	"subscriber (DDoS Flow Detection)" on page 28688
"subscriber-extension" on page 28692	"subscriber-address" on page 28683	"subscriber-aware-services" on page 28684
"subscriber-awareness (Service Set Options)" on page 28685	"subscriber-context (AAA Options)" on page 28687	"subscriber-exclude-prefix" on page 28691
"subscriber-info" on page 28704	"subscriber-leave-timer" on page 28694	"subscriber-management" on page 28697

"subscriber-management-helper" on page 28695	"subscriber-packet-idle-timeout" on page 28699	"subscriber-pool-utilization-alarm" on page 28705
"subscriber-profile" on page 28700	"subscriber-type (TDF Domain)" on page 28702	"subscribers-limit (Resource Monitor)" on page 28707
"subscription" on page 28709	"subscription-id" on page 28711	"subscription-id-data-include (PCRF Partition)" on page 28713
"subscription-id-options" on page 28715	"subscription-id-type (Class Attribute)" on page 28717	"subscription-id-type (PCRF Partition)" on page 28719
"success" on page 28720	"sunrpc" on page 28722	"supplementary-blackout-timer (Host Fast Reroute)" on page 28724
"supplicant" on page 28726	"suppress-host-routes-from-dci-to-dc" on page 28728	"suppress-interface-leaf" on page 28730
"suppression" on page 28731	"surf-control-integrated" on page 28733	"suspend-for" on page 28735
"suspend-on-request" on page 28737	"swap" on page 28740	"swap (Dynamic VLANs)" on page 28742
"swap (protocols mpls)" on page 28738	"swap-by-poppush" on page 28741	"swap-push" on page 28744
"swap-swap" on page 28745	"switch (Protocols OpenFlow)" on page 28755	"switch-away-lsps" on page 28746
"switch-options" on page 28749	"switch-options (Interfaces)" on page 28748	"switch-options (VLANs)" on page 28751
"switch-port" on page 28753	"switch-type" on page 28757	"switching-type" on page 28758
"switchover-delay" on page 28759	"switchover-mode" on page 28761	"switchover-on-routing-crash" on page 28763

"sxl-retry" on page 28764	"sxl-timeout" on page 28765	"symbol-period" on page 28767
"symmetric-hash" on page 28768	"syn-ack-ack-proxy" on page 28770	"syn-ack-ack-proxy (IDS Screen Next Gen Services)" on page 28771
"syn-check-required" on page 28773	"syn-cookie (IDS MS-DPC)" on page 28774	"syn-fin" on page 28776
"syn-fin (IDS Screen Next Gen Services)" on page 28777	"syn-flood" on page 28778	"syn-flood-protection-mode" on page 28779
"syn-frag" on page 28781	"syn-frag (IDS Screen Next Gen Services)" on page 28782	"sync-active-path-bandwidth" on page 28783
"sync-reset" on page 28785	"synchronization (ACX Series)" on page 28787	"synchronization (M Series and T Series)" on page 28792
"synchronization (MX Series)" on page 28795	"synchronization (PTX Series)" on page 28802	"synchronize " on page 28804
"synchronize (Scripts)" on page 28806	"synchronous-ethernet-mapping" on page 28808	"syntonized-e2e-transparent" on page 28813
"syslog" on page 28815	"syslog (Chassis)" on page 28817	"syslog (IDS MS-DPC)" on page 28819
"syslog (Interfaces)" on page 28820	"syslog (Next Gen Services Service-Set System Logging)" on page 28823	"syslog (OAM Action)" on page 28824
"syslog (QFabric System)" on page 28825	"syslog (Services CoS)" on page 28827	"syslog (Services IPsec VPN)" on page 28816
"syslog (Services L2TP)" on page 28829	"syslog (Services NAT) " on page 28830	"syslog (Services Service Set)" on page 28831
"syslog (Services Stateful Firewall) " on page 28833	"syslog (System)" on page 28836	"syslog-local-system-timestamp" on page 28821

"syslog-subtag" on page 28834	"system (Analytics Resource)" on page 28840	"system (Export Profiles)" on page 28844
"system login (Junos Node Slicing)" on page 28852	"system (Memory)" on page 28853	"system-defaults" on page 28842
"system-id" on page 28847	"system-id (Interfaces)" on page 28845	"system-id (Junos Fusion)" on page 28848
"system-id (Junos Fusion Satellite Device Cluster)" on page 28850	"system-priority (chassis)" on page 28855	"system-priority (interfaces)" on page 28857
"system-services (Security Zones Host Inbound Traffic)" on page 28859	"system-services (Security Zones Interfaces)" on page 28862	

T

"t1-interval" on page 28885	"t1-options" on page 28886	"t1-options (BITS Interfaces Signal Type)" on page 28887
"t310" on page 28889	"t391" on page 28890	"t392" on page 28892
"t4-interval" on page 28893	"table" on page 28895	"tacplus" on page 28896
"tacplus-options" on page 28898	"tacplus-server" on page 28902	"tag" on page 28904
"tag (Access)" on page 28905	"tag (Configuring Notification Targets)" on page 28910	"tag (HTTP Header Enrichment)" on page 28914
"tag (Routing Options)" on page 28924	"tag2 (Dynamic Access Routes)" on page 28937	"tag-attribute (HTTP Header Enrichment)" on page 28907

"tag-attribute (HTTP Header Enrichment Tag Rule)" on page 28908	"tag-group" on page 28911	"tag-header (HTTP Header Enrichment)" on page 28913
"tag-list" on page 28916	"tag-operation (HTTP Header Enrichment)" on page 28917	"tag-protocol-id (Dynamic VLANs)" on page 28919
"tag-protocol-id (TPID to Rewrite)" on page 28920	"tag-protocol-id TPID(TPIDs Expected to Be Sent or Received)" on page 28922	"tag-rule (HTTP Header Enrichment)" on page 28926
"tag-rule (Profiles for HTTP Header Enrichment)" on page 28928	"tag-rule-set (HTTP Header Enrichment)" on page 28930	"tag-rule-sets (Service Set)" on page 28931
"tag-rules (Service Set)" on page 28932	"tag-separator (HTTP Header Enrichment)" on page 28934	"tag-value (HTTP Header Enrichment)" on page 28935
"tags (Application Identification)" on page 28938	"talk" on page 28940	"target" on page 28941
"target (Routing Instances MVPN)" on page 28954	"target-address (SNMP)" on page 28946	"target-address (TWAMP)" on page 28944
"target-attachment-identifier (Protocols VPWS)" on page 28947	"target-logical-system (Domain Map)" on page 28949	"target-parameters" on page 28951
"target-routing-instance (Domain Map)" on page 28952	"targeted-broadcast" on page 28956	"targeted-distribution" on page 28958
"targeted-distribution (Dynamic Demux Interfaces over Aggregated Ethernet)" on page 28959	"targeted-distribution (Static Interfaces over Aggregated Ethernet)" on page 28961	"targeted-hello" on page 28962
"targeted-options (Grouping Subscribers by Bandwidth Usage)" on page 28964	"targeted-options (Manual Targeting)" on page 28966	"targeted-options (PS Interface)" on page 28968

"targets" on page 28970	"targets (Unified Edge)" on page 28971	"tca" on page 28973
"tcp" on page 28977	"tcp (Security IDP Protocol Binding)" on page 29013	"tcp (Security IDP Signature Attack)" on page 29014
"tcp (Security Screen)" on page 29017	"tcp-aggressive-transmission" on page 28978	"tcp-encap" on page 28980
"tcp-fast-open" on page 28982	"tcp-flags" on page 28983	"tcp-forwarding (Processes)" on page 28985
"tcp-forwarding (Remote Device Management)" on page 28987	"tcp-initial-timeout" on page 28989	"tcp-keepcnt" on page 28991
"tcp-keepidle" on page 28992	"tcp-keepintvl" on page 28994	"tcp-log" on page 28995
"tcp-mss (Internet Options)" on page 28996	"tcp-mss (Interfaces)" on page 29033	"tcp-mss (Protocols BGP)" on page 28998
"tcp-mss (Security Flow)" on page 29000	"tcp-mss (Services)" on page 29002	"tcp-no-flag" on page 29004
"tcp-no-flag (IDS Screen Next Gen Services)" on page 29005	"tcp-non-syn" on page 29006	"tcp-options (Security Policies)" on page 29008
"tcp-rst" on page 29012	"tcp-seq-ooo-window (Security Forwarding Options)" on page 29019	"tcp-session" on page 29020
"tcp-session (Service Set Next Gen Services)" on page 29023	"tcp-session-install-interval (Security Forwarding Options)" on page 29022	"tcp-sweep" on page 29025
"tcp-syn-defense (IDS MS-MPC)" on page 29027	"tcp-syn-fragment-check (IDS MS-MPC)" on page 29028	"tcp-tickles" on page 29029

"tcp-tickles (Service Set Next Gen Services)" on page 29030	"tcp-winnuke-check (IDS MS-MPC)" on page 29031	"tcpao-auth-mismatch" on page 29034
"tdf (Unified Edge)" on page 29037	"tdf-interface " on page 29036	"tdm-options (Interfaces)" on page 29039
"te-class-matrix" on page 29041	"te-link" on page 29043	"te-policy" on page 29044
"tear-drop (IDS Screen Next Gen Services)" on page 29046	"tei-option" on page 29047	"telemetries" on page 29048
"telemetries (PoE)" on page 29050	"telemetry" on page 29051	"telemetry" on page 29054
"telnet (Access)" on page 29059	"telnet (Services)" on page 29060	"temperature-sensor " on page 29063
"template" on page 29071	"template (Flow Monitoring IPFIX Version)" on page 29075	"template (Flow Monitoring Version 9)" on page 29077
"template (Forwarding Options)" on page 29079	"template (Forwarding Options Version IPFIX)" on page 29080	"template (Inline Monitoring)" on page 29084
"template (LRF Profile)" on page 29087	"template (LRF Rule)" on page 29089	"template (Services URL Filter)" on page 29072
"template-id" on page 29081	"template-profile (Flow Monitoring Logs for NAT)" on page 29090	"template-refresh-rate" on page 29094
"template-refresh-rate (Services)" on page 29096	"template-tx-interval (LRF Profile)" on page 29097	"template-type (Flow Monitoring Logs for NAT)" on page 29098
"template-type (LRF Profile)" on page 29100	"template(Protocols MVPN)" on page 29092	"templates" on page 29102
"term (Application Aware Access List)" on page 29106	"term (Applications)" on page 29110	"term (candidate-uplink-policy)" on page 29111

"term (Captive Portal Content Delivery)" on page 29113	"term (Dynamic Profiles)" on page 29115	"term (Firewall Filter)" on page 29117
"term (HTTP Header Enrichment)" on page 29120	"term (IDS MS-DPC)" on page 29123	"term (IDS MS-MPC)" on page 29125
"term (IP Security VPN)" on page 29107	"term (PCP)" on page 29128	"term (satellite-policies)" on page 29130
"term (Services CoS)" on page 29132	"term (Services NAT) " on page 29134	"term (Services Stateful Firewall) " on page 29136
"term (Simple Filter)" on page 29137	"term (TDF Domain Selection)" on page 29139	"term (URL Filter)" on page 29143
"terminal" on page 29146	"terminate (PPPoE Service Name Tables)" on page 29149	"terminate-code" on page 29147
"termination (Services)" on page 29151	"test" on page 29152	"test (Security IDP)" on page 29214
"test-agent (Paragon Active Assurance)" on page 29187	"test-count" on page 29192	"test-finish-wait-duration (RFC 2544 Benchmarking)" on page 29193
"test-interface (RFC 2544 Benchmarking)" on page 29196	"test-interval" on page 29198	"test-iterator-duration (RFC 2544 Benchmarking)" on page 29200
"test-iterator-pass-threshold (RFC 2544 Benchmarking)" on page 29201	"test-name" on page 29204	"test-only-mode" on page 29208
"test-profile (RFC 2544 Benchmarking)" on page 29211	"test-session (Junos OS)" on page 29237	"test-session (Junos OS Evolved)" on page 29239
"test-type (RFC 2544 Benchmarking)" on page 29244	"tests" on page 29247	"tftp" on page 29251

"tftp (Security ALG)" on page 29252	"tftp-server" on page 29254	"then (Captive Portal Content Delivery)" on page 29259
"then (Event Policy)" on page 29261	"then (Filters)" on page 29264	"then (Firewall Filters)" on page 29266
"then (Hierarchical Policer) " on page 29257	"then (HTTP Header Enrichment)" on page 29267	"then (IDS MS-DPC)" on page 29269
"then (IDS MS-MPC)" on page 29271	"then (IP Security VPN)" on page 29255	"then (LRF rule)" on page 29274
"then (PCC Rules)" on page 29276	"then (PCP)" on page 29278	"then (Policer Action)" on page 29279
"then (Policers)" on page 29281	"then (Security Application Firewall)" on page 29282	"then (Security Destination NAT)" on page 29284
"then (Security IDP Policy)" on page 29285	"then (Security Policies)" on page 29287	"then (Security Source NAT)" on page 29292
"then (Security Static NAT)" on page 29293	"then (Services CoS) " on page 29295	"then (Services CoS Next Gen Services)" on page 29296
"then (Services NAT) " on page 29299	"then (Services Stateful Firewall) " on page 29300	"then (Stateful Firewall Rule Next Gen Services)" on page 29302
"then (TDF Domain Selection)" on page 29303	"thermal-health-check" on page 29305	"three-color-policer" on page 29309
"three-color-policer (Applying)" on page 29311	"three-color-policer (Configuring)" on page 29314	"three-color-policer (Configuring) (EX)" on page 29313
"threshold" on page 29317	"threshold (Bridge Domains)" on page 29319	"threshold (Detection Time)" on page 29322

"threshold (detection-time)" on page 29323	"threshold (Interfaces)" on page 29325	"threshold (MSDP Active Source Messages) " on page 29327
"threshold (Multicast Forwarding Cache)" on page 29329	"threshold (PIM BFD Detection Time)" on page 29331	"threshold (PIM BFD Transmit Interval)" on page 29333
"threshold (PIM Entries)" on page 29334	"threshold (Routing Instances)" on page 29338	"threshold (Services IPsec)" on page 29340
"threshold (Services Logging and SYN-Cookie Defenses)" on page 29341	"threshold (transmit-interval)" on page 29344	"threshold-count" on page 29320
"threshold-logging-interval" on page 29326	"threshold-rate" on page 29336	"threshold-time" on page 29343
"thresholds" on page 29346	"thresholds (Junos OS)" on page 29348	"thresholds (Junos OS Evolved)" on page 29350
"time" on page 29352	"time-binding" on page 29354	"time-format" on page 29356
"time-interval (Dampen Options)" on page 29358	"time-limit (LRF Rule)" on page 29359	"time-wait-state" on page 29361
"time-zone" on page 29362	"timeout" on page 29367	"timeout (Access LDAP)" on page 29368
"timeout (Access RADIUS)" on page 29370	"timeout (DHCP Local Server)" on page 29374	"timeout (Diameter Network Element)" on page 29376
"timeout (Flow Maps)" on page 29377	"timeout (Multicast)" on page 29380	"timeout (RADIUS Server)" on page 29382
"timeout (Security Antivirus Fallback Options)" on page 29383	"timeout (Security Antivirus Fallback Options Sophos Engine)" on page 29385	"timeout (Security Antivirus Scan Options)" on page 29386

"timeout (Security GTP)" on page 29388	"timeout (Security Screen)" on page 29389	"timeout (Security Web Filtering)" on page 29391
"timeout (Security Web Filtering Cache)" on page 29392	"timeout (Security Web Filtering Fallback Settings)" on page 29393	"timeout (Services)" on page 29395
"timeout (SNMP)" on page 29365	"timeout-action" on page 29371	"timeout-active-flows (DDoS Flow Detection)" on page 29372
"timeout-grace (Access)" on page 29379	"timestamp" on page 29396	"timestamp-format (RFC 2544 Benchmarking)" on page 29398
"timestamp-option (IDS Screen Next Gen Services)" on page 29399	"tls" on page 29400	"tls (NETCONF)" on page 29403
"tls-min-version" on page 29402	"tls-peer-name" on page 29405	"tls-timeout" on page 29406
"tls-type" on page 29407	"tlsdetails" on page 29409	"tlv-filter" on page 29410
"tlv-select" on page 29414	"tm" on page 29417	"tm-polarity" on page 29418
"to" on page 29419	"to (Security Source NAT)" on page 29421	"to-zone (Security IDP Policy)" on page 29422
"to-zone (Security Policies)" on page 29423	"to-zone (Security Policies Global)" on page 29427	"token (DHCP Local Server)" on page 29428
"too-many-requests (Security Antivirus Fallback Options)" on page 29430	"too-many-requests (Security Antivirus Fallback Options Sophos Engine)" on page 29432	"too-many-requests (Security Web Filtering Fallback Settings)" on page 29433
"topologies (Multitopology Routing)" on page 29436	"topologies (Protocols IS-IS)" on page 29438	"topology (Filter-Based Forwarding)" on page 29440
"topology (Multitopology Routing)" on page 29443	"topology (OSPF)" on page 29445	"topology (OSPF Interface)" on page 29446

"topology (Protocols BGP)" on page 29449	"topology-id" on page 29441	"tos-reflect (L2TP)" on page 29451
"total-memory" on page 29452	"trace" on page 29454	"trace (DHCP Local Server)" on page 29456
"trace (DHCP Relay Agent)" on page 29458	"trace (EVPN-VXLAN)" on page 29459	"traceoptions (802.1X)" on page 29474
"traceoptions (Access)" on page 29477	"traceoptions (Access Port Security)" on page 29479	"traceoptions (Active Directory Access)" on page 29483
"traceoptions (Address-Assignment Pool)" on page 29485	"traceoptions (advanced policy-based routing)" on page 29488	"traceoptions (Analytics)" on page 29491
"traceoptions (Analytics Agent)" on page 29492	"traceoptions (ANCP)" on page 29494	"traceoptions (App Engine Virtual Machine Management Service)" on page 29497
"traceoptions (Batch Commits)" on page 29499	"traceoptions (Captive Portal Content Delivery)" on page 29501	"traceoptions (CASB)" on page 29503
"traceoptions (Chassis Cluster)" on page 29506	"traceoptions (Class of Service)" on page 29508	"traceoptions (Commit, Op, and SNMP Scripts)" on page 29511
"traceoptions (DDoS)" on page 29515	"traceoptions (DHCP)" on page 29517	"traceoptions (DHCP Server)" on page 29521
"traceoptions (Diameter Base Protocol)" on page 29527	"traceoptions (Diameter Base Protocol for Subscriber Aware Policy Control)" on page 29524	"traceoptions (DNS, Port, and TFTP Packet Forwarding)" on page 29530
"traceoptions (Dynamic Flow Capture)" on page 29538	"traceoptions (dynamic-address)" on page 29533	"traceoptions (dynamic-application)" on page 29535
"traceoptions (Egress Protection)" on page 29539	"traceoptions (Enhanced Subscriber Management)" on page 29541	"traceoptions (Ethernet Switching)" on page 29465

"traceoptions (Event Options)" on page 29544	"traceoptions (Event Scripts)" on page 29547	"traceoptions (Extensible Subscriber Services Manager)" on page 29550
"traceoptions (Extensions)" on page 29552	"traceoptions (FC-2 Fibre Channel)" on page 29554	"traceoptions (Fibre Channel)" on page 29557
"traceoptions (FIP Protocol Fibre Channel)" on page 29560	"traceoptions (flow)" on page 29563	"traceoptions (Forwarding Options)" on page 29565
"traceoptions (General Authentication Service)" on page 29567	"traceoptions (Health Check Monitoring)" on page 29569	"traceoptions (High Availability)" on page 29471
"traceoptions (IGMP Snooping)" on page 29573	"traceoptions (Individual Interfaces)" on page 29576	"traceoptions (Inline Monitoring)" on page 29585
"traceoptions (Interfaces)" on page 29587	"traceoptions (Juniper Secure Connect)" on page 29589	"traceoptions (LACP)" on page 29591
"traceoptions (Layer 2 Learning)" on page 29593	"traceoptions (LLDP)" on page 29596	"traceoptions (MACsec)" on page 29599
"traceoptions (MACsec interfaces)" on page 29602	"traceoptions (Multicast Snooping Options)" on page 29604	"traceoptions (MVRP)" on page 29607
"traceoptions (MX Series Virtual Chassis)" on page 29609	"traceoptions (NETCONF and Junos XML Protocol)" on page 29613	"traceoptions (NETCONF TLS)" on page 29616
"traceoptions (Next Gen Services Global System Logging)" on page 29619	"traceoptions (Next Gen Services Service-Set Flow)" on page 29620	"traceoptions (Next Gen Services Softwires)" on page 29624
"traceoptions (Origin Validation for BGP)" on page 29626	"traceoptions (Outbound SSH)" on page 29627	"traceoptions (OVSDB)" on page 29630

"traceoptions (PCE)" on page 29633	"traceoptions (PCEF)" on page 29635	"traceoptions (PIM Snooping)" on page 29638
"traceoptions (Port Mirroring and Traffic Sampling)" on page 29640	"traceoptions (PPP Process)" on page 29641	"traceoptions (PPPoE)" on page 29645
"traceoptions (Process Monitor)" on page 29648	"traceoptions (programmable-rpd)" on page 29651	"traceoptions (Protocols)" on page 29654
"traceoptions (Protocols AMT)" on page 29656	"traceoptions (Protocols BFD)" on page 29659	"traceoptions (Protocols BGP)" on page 29662
"traceoptions (Protocols BMP)" on page 29666	"traceoptions (Protocols DVMRP)" on page 29669	"traceoptions (Protocols EVPN)" on page 29672
"traceoptions (Protocols IGMP)" on page 29675	"traceoptions (Protocols IGMP Snooping)" on page 29678	"traceoptions (Protocols IPv6 Neighbor Discovery)" on page 29681
"traceoptions (Protocols IS-IS)" on page 29684	"traceoptions (Protocols Layer 2 Circuit)" on page 29688	"traceoptions (Protocols Layer 2 VPN)" on page 29690
"traceoptions (Protocols LDP)" on page 29693	"traceoptions (Protocols Link Management)" on page 29697	"traceoptions (Protocols MPLS)" on page 29699
"traceoptions (Protocols MSDP)" on page 29703	"traceoptions (Protocols MVPN)" on page 29706	"traceoptions (Protocols OpenFlow)" on page 29710
"traceoptions (Protocols OSPF)" on page 29713	"traceoptions (Protocols PCEP)" on page 29717	"traceoptions (Protocols PIM)" on page 29719
"traceoptions (Protocols PPP Service)" on page 29723	"traceoptions (Protocols RIP)" on page 29727	"traceoptions (Protocols RIPng)" on page 29730
"traceoptions (Protocols Router Discovery)" on page 29733	"traceoptions (Protocols RSVP)" on page 29736	"traceoptions (Protocols Secure Neighbor Discovery)" on page 29739

"traceoptions (Protocols Spring-TE)" on page 29742	"traceoptions (Protocols VPLS)" on page 29744	"traceoptions (Protocols VRRP)" on page 29747
"traceoptions (Proxy Fibre Channel)" on page 29750	"traceoptions (Resiliency)" on page 29752	"traceoptions (Resource Cleanup)" on page 29755
"traceoptions (Resource Monitor)" on page 29758	"traceoptions (REST API)" on page 29760	"traceoptions (routing-options)" on page 29462
"traceoptions (RPM)" on page 29761	"traceoptions (SBC Configuration Process)" on page 29764	"traceoptions (Security)" on page 29767
"traceoptions (Security ALG)" on page 29770	"traceoptions (Security Antispam)" on page 29773	"traceoptions (Security Antivirus)" on page 29774
"traceoptions (Security Application Firewall)" on page 29776	"traceoptions (Security Application Proxy)" on page 29778	"traceoptions (Security Content Filtering)" on page 29781
"traceoptions (Security Content Security)" on page 29782	"traceoptions (Security Datapath Debug)" on page 29783	"traceoptions (Security Dynamic VPN)" on page 29786
"traceoptions (Security Firewall Authentication)" on page 29788	"traceoptions (Security Flow)" on page 29789	"traceoptions (Security Group VPN)" on page 29795
"traceoptions (Security GTP)" on page 29798	"traceoptions (Security H323 ALG)" on page 29801	"traceoptions (Security IDP)" on page 29803
"traceoptions (Security IKE)" on page 29806	"traceoptions (Security IPsec)" on page 29809	"traceoptions (Security Log)" on page 29811
"traceoptions (Security MGCP ALG)" on page 29814	"traceoptions (Security NAT)" on page 29815	"traceoptions (Security PKI)" on page 29818
"traceoptions (Security Policies)" on page 29821	"traceoptions (Security SCCP ALG)" on page 29823	"traceoptions (Security Screen)" on page 29825

"traceoptions (Security SCTP)" on page 29828	"traceoptions (Security SIP ALG)" on page 29830	"traceoptions (Security User Identification)" on page 29832
"traceoptions (Security Web Filtering)" on page 29835	"traceoptions (Services)" on page 29836	"traceoptions (Services Application Identification)" on page 29839
"traceoptions (Services IPsec VPN)" on page 29842	"traceoptions (Services L2TP)" on page 29844	"traceoptions (Services Logging)" on page 29849
"traceoptions (Services Redundancy Daemon)" on page 29851	"traceoptions (Services SSL)" on page 29854	"traceoptions (Services UAC)" on page 29857
"traceoptions (Services User Identification)" on page 29859	"traceoptions (sFlow Technology)" on page 29862	"traceoptions (SMTP)" on page 29865
"traceoptions (SNMP)" on page 29866	"traceoptions (Spanning Tree)" on page 29869	"traceoptions (Static Subscribers)" on page 29873
"traceoptions (Subscriber Management)" on page 29876	"traceoptions (Subscriber Session Database Replication)" on page 29878	"traceoptions (System Accounting)" on page 29881
"traceoptions (System License)" on page 29883	"traceoptions (System Monitor)" on page 29885	"traceoptions (System Services DNS)" on page 29888
"traceoptions (TCP Encapsulation)" on page 29891	"traceoptions (TCP Port Forwarding)" on page 29894	"traceoptions (TDF Gateway)" on page 29896
"traceoptions (Traffic Load Balancer)" on page 29899	"traceoptions (Uplink Failure Detection)" on page 29903	"traceoptions (Virtual Chassis)" on page 29905
"tracing" on page 29987	"track (VRRP)" on page 29992	"track-igp-metric" on page 29989
"track-igp-metric (LSP)" on page 29990	"tracking-area" on page 29994	"traffic-class (delay)" on page 29996

"traffic-class (loss)" on page 29998	"traffic-class (loss-delay)" on page 30000	"traffic-class (Tunnels)" on page 30008
"traffic-class-map" on page 30003	"traffic-class-map (Apply to Interface)" on page 30006	"traffic-control-profiles" on page 30009
"traffic-control-profiles (Dynamic CoS Definition)" on page 30012	"traffic-engineering" on page 30014	"traffic-engineering (OSPF)" on page 30018
"traffic-engineering (Passive TE Mode)" on page 30021	"traffic-engineering (policy-options)" on page 30016	"traffic-engineering (Protocols BGP)" on page 30023
"traffic-engineering (Protocols IS-IS)" on page 30025	"traffic-engineering (Protocols MPLS)" on page 30031	"traffic-engineering (SR-TE import policy)" on page 30032
"traffic-load-balance (Traffic Load Balancer)" on page 30034	"traffic-manager" on page 30036	"traffic-options" on page 30040
"traffic-selector" on page 30042	"traffic-statistics (Protocols BGP)" on page 30051	"traffic-statistics (Protocols IS-IS)" on page 30048
"traffic-statistics (Protocols LDP)" on page 30053	"traffic-statistics (Services Analytics)" on page 30045	"traffic-statistics-labeled-path" on page 30049
"transaction-timeout" on page 30055	"transfer" on page 30056	"transfer-delay (Event Policy)" on page 30058
"transfer-interval" on page 30059	"transfer-log-archive" on page 30061	"transient-filter" on page 30062
"transit" on page 30064	"transit (Chained Composite Next Hops)" on page 30066	"transit-lsp-association" on page 30069
"translated" on page 30070	"translated-port" on page 30072	"translation" on page 30073
"translation-table" on page 30075	"translation-type" on page 30077	"translation-vni" on page 30080

"transmit-clock" on page 30083	"transmit-failure-threshold (RFC 2544 Benchmarking)" on page 30084	"transmit-interval" on page 30085
"transmit-interval (BFD Liveness Detection)" on page 30087	"transmit-interval (LAG)" on page 30090	"transmit-interval (Liveness Detection)" on page 30092
"transmit-interval (MACsec)" on page 30093	"transmit-interval (MACsec for MX Series)" on page 30095	"transmit-interval (PIM BFD Liveness Detection)" on page 30096
"transmit-rate" on page 30098	"transmit-rate (Dynamic Schedulers)" on page 30103	"transmit-rate (EX Series Switches)" on page 30105
"transmit-rate (Schedulers)" on page 30107	"transparent" on page 30110	"transport (Diameter Base Protocol)" on page 30123
"transport (Diameter Peer)" on page 30124	"transport (Next Gen Services Syslog Message Security)" on page 30127	"transport (Protocols BGP)" on page 30113
"transport (Security Log)" on page 30129	"transport (Services)" on page 30111	"transport (System Syslog)" on page 30112
"transport-address" on page 30118	"transport-class" on page 30132	"transport-class" on page 30120
"transport-monitoring" on page 30126	"transport-type" on page 30131	"trap" on page 30133
"trap (Security IKE)" on page 30139	"trap (Security PKI)" on page 30140	"trap-group" on page 30135
"trap-options" on page 30136	"traps" on page 30142	"traps (RPM and TWAMP)" on page 30144
"tri-color" on page 30147	"tri-color (EX Series Switches)" on page 30149	"trickling" on page 30152

"trigger" on page 30154	"trigger (CoS)" on page 30159	"trigger (DHCP Local Server)" on page 30160
"trigger (Event Policy)" on page 30162	"trigger-link-failure" on page 30164	"trigger-type (LRF Profile)" on page 30165
"trio-flow-offload" on page 30167	"trust-option-82" on page 30168	"trusted" on page 30170
"trusted (DHCP Security)" on page 30174	"trusted-ca" on page 30171	"trusted-ca (Services)" on page 30172
"trusted-keys (DNSSEC)" on page 30175	"tti" on page 30176	"ttl" on page 30178
"ttl (Dynamic Flow Capture)" on page 30179	"ttl (RPM probe)" on page 30180	"ttl-threshold" on page 30183
"tunable-name" on page 30184	"tunable-value" on page 30186	"tunables" on page 30187
"tunnel" on page 30189	"tunnel" on page 30190	"tunnel" on page 30192
"tunnel (fti)" on page 30204	"tunnel (L2TP)" on page 30210	"tunnel (Security Policies)" on page 30231
"tunnel (Security Screen)" on page 30232	"tunnel (Tunnel Profile)" on page 30248	"tunnel-attributes" on page 30194
"tunnel-devices (Protocols AMT)" on page 30196	"tunnel-devices (Tunnel-Capable PICs) " on page 30198	"tunnel-end-point " on page 30200
"tunnel-group" on page 30206	"tunnel-inspection" on page 30208	"tunnel-limit (Protocols AMT)" on page 30212
"tunnel-limit (Routing Instances)" on page 30213	"tunnel-limit (Routing Instances Provider Tunnel Selective)" on page 30215	"tunnel-mtu" on page 30216

"tunnel-mtu (Services IPsec VPN)" on page 30218	"tunnel-mtu (Services Service Set)" on page 30219	"tunnel-observation" on page 30221
"tunnel-port" on page 30223	"tunnel-port (GRE)" on page 30224	"tunnel-profile (Domain Map)" on page 30225
"tunnel-profile (L2TP Tunnel Switching)" on page 30226	"tunnel-profile (Tunnel Profile)" on page 30228	"tunnel-queuing" on page 30230
"tunnel-services (Chassis)" on page 30234	"tunnel-services (Routing Instances VPLS)" on page 30236	"tunnel-services (RSVP)" on page 30238
"tunnel-source" on page 30239	"tunnel-source-protocol (Protocols OSPF)" on page 30240	"tunnel-switch-profile (L2TP Tunnel Switching, Application)" on page 30242
"tunnel-switch-profile (L2TP Tunnel Switching, Definition)" on page 30243	"tunnel-termination" on page 30245	"tunnel-timeout" on page 30247
"twamp" on page 30250	"twamp (Security ALG)" on page 30256	"twamp-server" on page 30258
"two-rate" on page 30259	"tx-address-change (L2TP LAC)" on page 30261	"tx-buffers" on page 30263
"tx-connect-speed-method (L2TP LAC)" on page 30265	"tx-dnu-to-line-source-enable" on page 30268	"tx-duration" on page 30269
"tx-enable" on page 30271	"tx-power" on page 30272	"type" on page 30273
"type (Application Identification)" on page 30275	"type (ICMP Mapping for Application Identification)" on page 30276	"type (Protocols BGP)" on page 30278
"type (RMON)" on page 30279	"type (RMON Notification)" on page 30281	"type (Security Antivirus Feature Profile)" on page 30282

"type (Security Content Filtering Notification Options)" on page 30284	"type (Security Fallback Block)" on page 30285	"type (Security IDP Dynamic Attack Group)" on page 30287
"type (Security IDP ICMP Headers)" on page 30288	"type (Security Virus Detection)" on page 30289	"type (Security Web Filtering)" on page 30291
"type (SNMP v3)" on page 30293	"type (Tunnel Profile)" on page 30294	

U

"uac-policy (Application Services)" on page 30301	"uac-service" on page 30303	"ucode-imem-remap" on page 30304
"udp" on page 30306	"udp (FTI)" on page 30309	"udp (Security IDP Protocol Binding)" on page 30312
"udp (Security IDP Signature Attack)" on page 30313	"udp (Security Screen)" on page 30314	"udp-encapsulation" on page 30307
"udp-port" on page 30310	"udp-sweep" on page 30316	"udp-tcp-port-swap" on page 30318
"uid (Dynamic Profiles)" on page 30319	"uid-reference" on page 30320	"ukern-trace" on page 30322
"ultimate-hop-popping" on page 30323	"ultrahigh-power" on page 30325	"unattended-boot" on page 30326
"unconditional-src-learn" on page 30328	"unconfigured-peer-graceful-restart" on page 30329	"underlying-interface" on page 30330
"underlying-interface (ANCP)" on page 30332	"underlying-interface (demux0)" on page 30333	"underlying-interface (Dynamic PPPoE)" on page 30335

"unframed | no-unframed (Interfaces)" on page 30337	"unicast (Route Target Community)" on page 30350	"unicast (Virtual Tunnel in Routing Instances)" on page 30354
"unicast-ethernet" on page 30338	"unicast-ethernet-ipv4" on page 30340	"unicast-in-lpm" on page 30342
"unicast-mode " on page 30344	"unicast-negotiation" on page 30347	"unicast-reverse-path" on page 30348
"unicast-stream-limit (Protocols AMT)" on page 30351	"unicast-tunnel-name-regular-expression" on page 30355	"unicast-umh-election" on page 30353
"unidirectional" on page 30356	"unidirectional-session-refreshing" on page 30358	"unified-access-control (Services)" on page 30359
"unified-policy max-lookups" on page 30362	"unified-policy-explicit-match" on page 30361	"unique-nas-port (Access)" on page 30364
"unit (Aggregated Multiservices)" on page 30380	"unit (CoS Interfaces)" on page 30384	"unit (CoS Interfaces)" on page 30382
"unit (Dynamic Demux Interface)" on page 30385	"unit (Dynamic Interface Sets)" on page 30388	"unit (Dynamic PPPoE)" on page 30390
"unit (Dynamic Profiles Standard Interface)" on page 30392	"unit (Dynamic Traffic Shaping)" on page 30397	"unit (interfaces)" on page 30370
"unit (Interfaces)" on page 30401	"unit (Interfaces Services) " on page 30399	"unit (Link-Multilink)" on page 30402
"unit (QFX Series)" on page 30367	"unit (Sampling)" on page 30366	"unit (TDF Interface)" on page 30406
"unit (Tunnel)" on page 30404	"unit (Voice Services)" on page 30408	"unknown-message (Security H323 ALG)" on page 30410
"unknown-message (Security MGCP ALG)" on page 30411	"unknown-message (Security SCCP ALG)" on page 30413	"unknown-message (Security SIP ALG)" on page 30415

"unknown-protocol (IDS Screen Next Gen Services)" on page 30416	"unknown-unicast-forwarding" on page 30417	"unnumbered-address (Demux)" on page 30419
"unnumbered-address (Dynamic PPPoE)" on page 30421	"unnumbered-address (Dynamic Profiles)" on page 30422	"unnumbered-address (Ethernet)" on page 30425
"unnumbered-address (PPP)" on page 30426	"unreachable-vtep-aging-timer" on page 30429	"untagged" on page 30431
"untrusted" on page 30432	"untrusted (VLAN)" on page 30433	"Unused" on page 30435
"update-interval" on page 30439	"update-interval (Protocols RIP)" on page 30441	"update-interval (Protocols RIPng)" on page 30443
"update-interval (Service Accounting)" on page 30444	"update-rate-limit" on page 30446	"update-response-timeout (PCRF Partition)" on page 30447
"update-router-advertisement" on page 30449	"update-server" on page 30450	"update-server (dhcp-client)" on page 30452
"update-server (dhcpv6-client)" on page 30453	"update-threading" on page 30454	"update-threshold" on page 30456
"update-threshold-max-reservable" on page 30458	"upgrade-group" on page 30460	"upgrade-groups (Junos Fusion)" on page 30461
"upgrade-mode" on page 30464	"uplink-failure-detection" on page 30465	"uplink-failure-detection (Junos Fusion)" on page 30467
"uplink-port-group (Junos Fusion)" on page 30469	"upload (Event Policy)" on page 30471	"upload-profile (Security Antivirus FTP)" on page 30472
"upload-profile (Security Content Filtering FTP)" on page 30473	"upstream-interface" on page 30474	"upstream-label" on page 30476
"upstream-mode" on page 30478	"upstream-rate (Traffic Shaping)" on page 30479	"upstream-rewrite" on page 30481

"urgent-pointer" on page 30483	"uri-check" on page 30484	"url" on page 30485
"url (Security)" on page 30487	"url (Security)" on page 30503	"url (Security Antivirus)" on page 30504
"url (Security IDP)" on page 30506	"url-blacklist" on page 30488	"url-feed" on page 30489
"url-filter" on page 30492	"url-filter-profile" on page 30495	"url-filter-template" on page 30496
"url-pattern" on page 30499	"url-whitelist " on page 30507	"url-whitelist " on page 30508
"usb" on page 30509	"usb-control" on page 30510	"use-active-child-mac-on-reth" on page 30511
"use-actual-mac-on-physical-interfaces" on page 30512	"use-class (Class Attribute)" on page 30514	"use-extended-flow-memory" on page 30515
"use-for-post-convergence-lfa" on page 30517	"use-imported-time-zones" on page 30518	"use-interface" on page 30519
"use-interface-description" on page 30522	"use-interface-description (MX)" on page 30520	"use-interface-index" on page 30525
"use-interface-name" on page 30526	"use-option-82" on page 30528	"use-p2mp-lsp" on page 30530
"use-post-convergence-lfa" on page 30531	"use-post-convergence-lfa (Protocols OSPF)" on page 30533	"use-primary (DHCP Local Server)" on page 30535
"use-primary (DHCP Relay Agent)" on page 30537	"use-string" on page 30539	"use-transport-class" on page 30541
"use-underlying-interface-mac " on page 30542	"use-vlan-id" on page 30544	"use-vlan-id (DHCP Relay Agent)" on page 30546
"user (Access)" on page 30547	"user (SNMP v3)" on page 30588	"user (System Logging)" on page 30589

"user (System Services)" on page 30591	"user-defined-option-82" on page 30551	"user-firewall" on page 30553
"user-group-mapping" on page 30555	"user-id" on page 30557	"user-identification" on page 30558
"user-identification (Services)" on page 30560	"user-name" on page 30568	"user-name (Event Policy)" on page 30569
"user-name-include (OCS Partition)" on page 30571	"user-password (PCEF Profile)" on page 30574	"user-planes" on page 30575
"user-prefix" on page 30577	"user-prefix (DHCP Local Server)" on page 30578	"user-prefix (DHCP Relay Agent)" on page 30580
"user-prefix (Static Subscribers)" on page 30582	"user-query (Services User Identification)" on page 30584	"username (Security Antivirus)" on page 30606
"username (Services)" on page 30607	"username-include (Demux)" on page 30593	"username-include (DHCP Local Server)" on page 30595
"username-include (DHCP Relay Agent)" on page 30597	"username-include (Interfaces)" on page 30600	"username-include (Local Authentication)" on page 30602
"username-include (Static Subscribers)" on page 30604	"using-user-password" on page 30609	"usm" on page 30610
"utc-timestamp" on page 30612	"utm default-configuration" on page 30614	"uuid" on page 30621
"uuid (Applications)" on page 30623		

V

"v3" on page 30631	"v4address" on page 30635	"v4ov6" on page 30636
"v6address" on page 30638	"v6prefix" on page 30639	"v6rd" on page 30641
"vacm" on page 30642	"valid-lifetime" on page 30644	"valid-lifetime (Dynamic Router Advertisement)" on page 30645
"validation (Origin Validation for BGP)" on page 30647	"validity-period" on page 30650	"variable" on page 30651
"variable" on page 30653	"variables (Dynamic Service Profiles)" on page 30654	"variant" on page 30655
"vbr rate" on page 30657	"vc-port" on page 30658	"vci" on page 30660
"vcp-no-hold-time" on page 30662	"vcp-snmp-statistics" on page 30664	"vdsl2-bytes" on page 30671
"vdsl2-overhead-adjust" on page 30673	"vdsl-bytes" on page 30665	"vdsl-overhead-adjust" on page 30667
"vdsl-profile" on page 30669	"vendor" on page 30675	"vendor-id (AAA Profile)" on page 30681
"vendor-id (DHCP Client)" on page 30676	"vendor-id (dhcp-option82)" on page 30677	"vendor-id (Interfaces)" on page 30682
"vendor-id (Unified Edge)" on page 30679	"vendor-option" on page 30684	"vendor-specific (DHCP Relay Agent)" on page 30686
"vendor-support" on page 30687	"verify-path" on page 30688	"version (BFD)" on page 30695
"version (BFD Liveness Detection)" on page 30697	"version (Dynamic IGMP Interface)" on page 30699	"version (Dynamic MLD Interface)" on page 30701

"version (Flow Aggregation)" on page 30692	"version (Flow Monitoring Logs for NAT)" on page 30702	"version (IGMP Snooping)" on page 30704
"version (IKE)" on page 30706	"version (Liveness Detection)" on page 30710	"version (MLD Snooping)" on page 30712
"version (Protocols BFD)" on page 30714	"version (Protocols IGMP)" on page 30715	"version (Protocols IGMP AMT)" on page 30716
"version (Protocols MLD)" on page 30718	"version (Protocols PIM)" on page 30719	"version (SNMP)" on page 30721
"version9 (Flow Monitoring)" on page 30722	"version9 (Forwarding Options)" on page 30724	"version-3" on page 30691
"version-alias (Dynamic Profiles)" on page 30694	"version-ipfix (Forwarding Options)" on page 30707	"version-ipfix (Services)" on page 30708
"versioning" on page 30725	"video" on page 30727	"video (Application Profile)" on page 30728
"video-monitoring" on page 30729	"view (Configuring a MIB View)" on page 30733	"view (SNMP Community)" on page 30734
"violation-action (DHCP Local Server and DHCP Relay Agent)" on page 30735	"violation-report-rate (DDoS Flow Detection)" on page 30737	"virtual-address" on page 30739
"virtual-channel-group (CoS Interfaces)" on page 30740	"virtual-channel-groups" on page 30742	"virtual-channels" on page 30743
"virtual-chassis" on page 30745	"virtual-chassis (MX Series Virtual Chassis)" on page 30747	"virtual-circuit-id" on page 30749
"virtual-control-channel" on page 30751	"virtual-cpu" on page 30752	"virtual-gateway-address" on page 30754

"virtual-gateway-v4-mac" on page 30755	"virtual-gateway-v6-mac" on page 30757	"virtual-inet6-address" on page 30760
"virtual-ip" on page 30761	"virtual-link" on page 30763	"virtual-link-local-address" on page 30765
"virtual-mac" on page 30766	"virtual-network-functions" on page 30768	"virtual-network-functions (Junos node slicing)" on page 30773
"virtual-service (Traffic Load Balancer)" on page 30774	"virtual-switch" on page 30777	"virus-detection (Security Antivirus)" on page 30778
"vjunos0" on page 30780	"vlan " on page 30782	"vlan (Access Port Security)" on page 30785
"vlan (Bridge Domains)" on page 30788	"vlan (DHCP Bindings on Access Ports)" on page 30790	"vlan (IGMP Snooping)" on page 30821
"vlan members (VLANs)" on page 30827	"vlan (Mirroring)" on page 30829	"vlan (MLD Snooping)" on page 30830
"vlan (MSTP)" on page 30835	"vlan (PIM Snooping)" on page 30837	"vlan (Port Mirroring)" on page 30841
"vlan (Port Mirroring)" on page 30839	"vlan (Secure Access Port)" on page 30852	"vlan (Static IP)" on page 30854
"vlan (Unknown Unicast Forwarding)" on page 30876	"vlan (VSTP)" on page 30878	"vlan-group" on page 30791
"vlan-id" on page 30795	"vlan-id (Dynamic Profiles)" on page 30801	"vlan-id (Dynamic VLANs)" on page 30802
"vlan-id (Fast Ethernet)" on page 30794	"vlan-id inner-all" on page 30804	"vlan-id (Logical Port in Bridge Domain)" on page 30813

"vlan-id (routing instance)" on page 30817	"vlan-id (VLAN ID to Be Bound to a Logical Interface)" on page 30820	"vlan-id-list" on page 30805
"vlan-id-list (Ethernet VLAN Circuit)" on page 30808	"vlan-id-list (Interface in Bridge Domain)" on page 30810	"vlan-id-list (Interface in VPLS)" on page 30812
"vlan-id-range" on page 30815	"vlan-id-start" on page 30818	"vlan-isolation" on page 30826
"vlan-model" on page 30833	"vlan-prune" on page 30842	"vlan-range" on page 30843
"vlan-ranges" on page 30845	"vlan-ranges (RADIUS Options)" on page 30847	"vlan-rewrite" on page 30849
"vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP)" on page 30850	"vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP)" on page 30855	"vlan-tag" on page 30857
"vlan-tag (Dynamic Classifiers)" on page 30858	"vlan-tag (Dynamic Rewrite Rules)" on page 30860	"vlan-tagging" on page 30861
"vlan-tagging (Dynamic)" on page 30864	"vlan-tags" on page 30865	"vlan-tags (Dual-Tagged Logical Interface)" on page 30869
"vlan-tags (Dynamic Profiles)" on page 30867	"vlan-tags (Stacked VLAN Tags)" on page 30873	"vlan-tags-outer" on page 30872
"vlans" on page 30880	"vlans (RA Guard)" on page 30894	"vmhost" on page 30895
"vmhost management-if disable" on page 30899	"vmhost management-if link-mode" on page 30900	"vmhost management-if speed " on page 30902
"vmhost resize vjunos compact" on page 30904	"vnf-license-supplement" on page 30905	"vnf-name" on page 30906
"vni" on page 30909	"vni-options" on page 30912	"vnid (EVPN)" on page 30914

"vni(Interfaces)" on page 30911	"voice" on page 30915	"voice (Application Profile)" on page 30917
"voip" on page 30918	"voip-mac-exclusive" on page 30920	"volume-limit (LRF Rule)" on page 30921
"vpi (Define Virtual Path)" on page 30923	"vpls (Interfaces)" on page 30931	"vpls (Routing Instance)" on page 30933
"vpls-flow-table-size" on page 30924	"vpls-flush-on-topology-change" on page 30926	"vpls-id" on page 30927
"vpls-id-list (protocols vpls mesh-group)" on page 30929	"vpls-mac-move" on page 30932	"vpls-template" on page 30937
"vpn (Security)" on page 30946	"vpn-apply-export" on page 30938	"vpn-global-import" on page 30953
"vpn-group-address" on page 30939	"vpn-label" on page 30941	"vpn-localization" on page 30942
"vpn-monitor" on page 30943	"vpn-unequal-cost" on page 30951	"vpws-service-id" on page 30954
"vrf-advertise-selective" on page 30957	"vrf-export" on page 30958	"vrf-import" on page 30960
"vrf-mtu-check" on page 30962	"vrf-name (Duplicate Accounting)" on page 30963	"vrf-propagate-ttl" on page 30965
"vrf-table-label" on page 30966	"vrf-target" on page 30968	"vrrp" on page 30971
"vrrp-group" on page 30973	"vrrp-inet6-group" on page 30976	"vrrp-inherit-from" on page 30978
"vsi-discovery" on page 30980	"vsi-policy" on page 30981	"vstp" on page 30983
"vtep-source-interface" on page 30988	"vulnerability-type" on page 30990	"vxlan" on page 30991

"vxlan-stitching" on page 31009	"vxlan-disable-copy-tos-decap" on page 30993	"vxlan-disable-copy-tos-encap" on page 30995
"vxlan-extended" on page 30998	"vxlan-gbp-l2-profile" on page 30999	"vxlan-gbp-l3-profile" on page 31001
"vxlan-gbp-profile" on page 31003	"vxlan-gpe (FTI)" on page 31005	"vxlan-routing " on page 31007
"vxlan-tos-copy-filter " on page 31011		

W

"wait-for-acct-on-ack (Access Profile)" on page 31014	"wait-to-block-interval" on page 31016	"wait-to-restore" on page 31017
"walkup" on page 31019	"wander-filter-disable" on page 31021	"warm-standby" on page 31023
"warm-standby (Interfaces)" on page 31022	"warning" on page 31024	"watch-list" on page 31026
"watchdog" on page 31027	"watchdog (PSM)" on page 31028	"watchdog-timeout" on page 31031
"wavelength" on page 31032	"web-authentication" on page 31057	"web-authentication (Access)" on page 31058
"web-authentication (Interfaces)" on page 31060	"web-filter" on page 31062	"web-filter-profile" on page 31065
"web-filtering" on page 31066	"web-filtering (Logical System Security Feature Profile)" on page 31072	"web-filtering (Security Content Security Policy)" on page 31076

"web-management (System Processes)" on page 31078	"web-management (System Services)" on page 31079	"web-redirect" on page 31084
"web-redirect-to-https" on page 31085	"web-server (Services)" on page 31086	"webapi (System Services)" on page 31091
"webapi-clear-text (Security)" on page 31089	"webapi-ssl (Security)" on page 31090	"websense-redirect" on page 31093
"weight" on page 31095	"weight (Security) " on page 31099	"weight (Targeted Subscribers)" on page 31097
"weighted-load-balancing (L2TP LAC)" on page 31100	"west-interface" on page 31102	"white-list" on page 31106
"white-list (Security Screen)" on page 31107	"white-list (Security-Zone)" on page 31109	"whitelist " on page 31111
"whitelist-url-categories" on page 31112	"wildcard-group-inet" on page 31115	"wildcard-group-inet6" on page 31117
"wildcard-source (PIM RPF Selection)" on page 31119	"wildcard-source (Selective Provider Tunnels) " on page 31120	"window-scale" on page 31122
"window-size" on page 31123	"windows-logon (Juniper Secure Connect)" on page 31124	"winnuke" on page 31126
"winnuke (IDS Screen Next Gen Services)" on page 31127	"wins-server (Access)" on page 31129	"wins-server (System)" on page 31130
"wired-max-processes" on page 31131	"wired-process-mem-size" on page 31133	"withdraw-priority" on page 31134
"within (Event Policy)" on page 31136	"wlan" on page 31137	"world-readable (Forwarding Options)" on page 31142

"world-readable (Next Gen Services Global System Logging)" on page 31145	"world-readable (System Syslog)" on page 31144	"write-interval" on page 31146
"write-view" on page 31148		

X

"xauth-attributes" on page 31150	"xe (Port)" on page 31152	"xlat-source-rule" on page 31154
"xle (Port)" on page 31155	"xnm-clear-text" on page 31156	"xnm-ssl" on page 31158
"xon (Input Congestion Notification)" on page 31161		

Y

"yang-compliant (NETCONF)" on page 31163	"yang-modules (NETCONF)" on page 31166	"yellow-differential-delay" on page 31168

Z

"zone (System Security Profile)" on page 31170	"zones" on page 31172	

4

CHAPTER

CLI Commands

CLI Commands | 1487

CLI Commands

IN THIS SECTION

- A | 1487
- C | 1488
- D | 1502
- E | 1502
- F | 1502
- H | 1503
- I | 1503
- L | 1503
- M | 1503
- O | 1504
- P | 1504
- Q | 1505
- R | 1505
- S | 1519
- T | 1573
- U | 1574
- V | 1575
- W | 1575
- Others | 1575

A

"activate" on page 1795	*[OBSOLETE] ADD (DTCP)*	"advanced-anti-malware connection" on page 1994

"advanced-anti-malware traceoptions" on page 2005	"annotate" on page 2252	

C

"clear access-security router-advertisement statistics" on page 3403	"clear access-security slaac-snooping binding" on page 3404	"clear access-security slaac-snooping statistics" on page 3406
"clear accounting server statistics archival-transfer" on page 3407	"clear amt statistics" on page 3410	"clear amt tunnel" on page 3412
"clear ancp access-loop" on page 3414	"clear ancp neighbor" on page 3416	"clear ancp statistics" on page 3418
"clear ancp subscriber" on page 3421	"clear arp" on page 3424	"clear arp inspection statistics" on page 3426
"clear auto-configuration interfaces" on page 3428	"clear auto-configuration interfaces interface-set" on page 3430	"clear bfd adaptation" on page 3432
"clear bfd session" on page 3434	"clear bgp damping" on page 3436	"clear bgp neighbor" on page 3437
"clear bgp table" on page 3441	"clear bridge mac-ip-table" on page 3443	"clear bridge mac-table" on page 3444
"clear bridge recovery-timeout" on page 3446	"clear bridge statistics" on page 3447	"clear captive-portal" on page 3450
"clear chassis alarms fabric degraded" on page 3453	"clear chassis cluster control-plane statistics" on page 3455	"clear chassis cluster data-plane statistics" on page 3457

"clear chassis cluster failover-count" on page 3458	"clear chassis cluster ip monitoring failure count ip address" on page 3462	"clear chassis cluster ip-monitoring failure-count" on page 3460
"clear chassis cluster statistics" on page 3463	"clear chassis display message" on page 3465	"clear chassis fpc errors" on page 3468
"clear chassis hardware history" on page 3470	"clear chassis high-availability data-plane statistics" on page 3472	"clear chassis high-availability information" on page 3473
"clear chassis sib errors" on page 3474	"clear class-of-service application-traffic-control counter" on page 3476	"clear class-of-service application-traffic-control rate-limiters" on page 3478
"clear class-of-service application-traffic-control statistics rule" on page 3479	"clear cli logical-system" on page 3481	"clear ddos-protection protocols" on page 3482
"clear dhcp client binding " on page 3484	"clear dhcp client statistics" on page 3486	"clear dhcp relay active-leasequery statistics" on page 3488
"clear dhcp relay binding" on page 3493	"clear dhcp relay binding (Security)" on page 3491	"clear dhcp relay lockout-entries" on page 3496
"clear dhcp relay statistics" on page 3500	"clear dhcp relay statistics (Security)" on page 3499	"clear dhcp server active-leasequery statistics" on page 3506
"clear dhcp server binding" on page 3508	"clear dhcp server lockout-entries" on page 3512	"clear dhcp server statistics" on page 3515
"clear dhcp snooping binding" on page 3517	"clear dhcp snooping statistics" on page 3519	"clear dhcp-security binding" on page 3503
"clear dhcp-security ipv6 binding" on page 3505	"clear dhcpv6 client binding" on page 3521	"clear dhcpv6 client statistics" on page 3522

"clear dhcpv6 relay active-leasequery statistics" on page 3524	"clear dhcpv6 relay binding" on page 3527	"clear dhcpv6 relay lockout-entries" on page 3531
"clear dhcpv6 relay statistics" on page 3534	"clear dhcpv6 server active-leasequery statistics" on page 3536	"clear dhcpv6 server binding" on page 3538
"clear dhcpv6 server binding (Local Server)" on page 3542	"clear dhcpv6 server lockout-entries" on page 3543	"clear dhcpv6 server statistics" on page 3546
"clear dhcpv6 server statistics (Local Server)" on page 3548	"clear dhcpv6 snooping binding" on page 3549	"clear dhcpv6 snooping statistics" on page 3551
"clear diameter function statistics" on page 3553	"clear diameter peer" on page 3554	"clear dot1x" on page 3564
"clear edge-virtual-bridging" on page 3566	"clear error bpdu interface" on page 3567	"clear error mac-rewrite" on page 3568
"clear ethernet-switching bpdu-error interface" on page 3570	"clear ethernet-switching evpn arp-table" on page 3572	"clear ethernet-switching evpn nd-statistics" on page 3573
"clear ethernet-switching evpn nd-table" on page 3574	"clear ethernet-switching layer2-protocol-tunneling error" on page 3576	"clear ethernet-switching layer2-protocol-tunneling statistics" on page 3577
"clear ethernet-switching mac-ip-table" on page 3579	"clear ethernet-switching port-error" on page 3580	"clear ethernet-switching recovery-timeout" on page 3582
"clear ethernet-switching recovery-timeout (SRX)" on page 3583	"clear ethernet-switching statistics mac-learning" on page 3584	"clear ethernet-switching table persistent-learning" on page 3586
"clear evpn duplicate-mac-suppression " on page 3588	"clear evpn mac-ip-table" on page 3589	"clear evpn nd-table" on page 3590
"clear evpn statistics" on page 3592	"clear extensible-subscriber-services counters" on page 3593	"clear extensible-subscriber-services sessions" on page 3594

"clear fibre-channel fc2 statistics" on page 3596	"clear fibre-channel fip enode" on page 3597	"clear fibre-channel fip statistics" on page 3599
"clear fibre-channel fip vn-port" on page 3600	"clear fibre-channel flogi statistics" on page 3602	"clear fibre-channel proxy statistics" on page 3603
"clear fip snooping enode" on page 3604	"clear fip snooping statistics" on page 3606	"clear fip snooping vlan" on page 3608
"clear fip vlan-discovery statistics" on page 3610	"clear firewall" on page 3613	"clear firewall (QFX)" on page 3611
"clear forwarding-options port-mirroring instance family" on page 3616	"clear gtp tunnels" on page 3618	"clear igmp membership" on page 3619
"clear igmp snooping membership" on page 3623	"clear igmp snooping statistics" on page 3625	"clear igmp statistics" on page 3627
"clear ike security-associations" on page 3632	"clear ilmi statistics" on page 3634	"clear inband-flow-telemetry stats" on page 3635
"clear interfaces aeX forwarding-options load-balance state" on page 3637	"clear interfaces aggregate forwarding-options load-balance state" on page 3638	"clear interfaces interface-set statistics" on page 3640
"clear interfaces interval" on page 3641	"clear interfaces lb-stats af" on page 3643	"clear interfaces mac-database" on page 3648
"clear interfaces mac-database statistics" on page 3649	"clear interfaces statistics" on page 3651	"clear interfaces statistics swfabx" on page 3653
"clear interfaces transport pm" on page 3654	"clear ipsec security-associations" on page 3657	"clear ipv6 neighbors " on page 3659
"clear ipv6 neighbors (SRX Series)" on page 3661	"clear ipv6 router-advertisement " on page 3663	"clear isis adjacency" on page 3664

"clear isis database" on page 3667	"clear isis overload" on page 3670	"clear isis spring traffic-statistics" on page 3673
"clear isis statistics" on page 3674	"clear l2-learning mac-move-buffer" on page 3677	"clear lacp statistics interfaces" on page 3679
"clear ldp neighbor" on page 3680	"clear ldp session" on page 3682	"clear ldp statistics" on page 3683
"clear lldp neighbors" on page 3685	"clear lldp statistics" on page 3686	"clear log" on page 3688
"clear log (JDM)" on page 3690	"clear log (Next Gen Services)" on page 3692	"clear loop-detect enhanced interface" on page 3693
"clear mac-vrf forwarding mac-ip-table" on page 3695	"clear mac-vrf forwarding mac-learning-log" on page 3697	"clear mac-vrf forwarding mac-table" on page 3698
"clear mac-vrf forwarding recovery-timeout" on page 3700	"clear mac-vrf forwarding redundancy-group" on page 3701	"clear mld membership" on page 3702
"clear mld snooping membership" on page 3704	"clear mld snooping statistics" on page 3705	"clear mld statistics" on page 3707
"clear monitor security flow filter" on page 3709	"clear mpls container-lsp" on page 3711	"clear mpls lsp" on page 3713
"clear mpls tunnel manager statistics" on page 3716	"clear msdp cache" on page 3717	"clear msdp statistics" on page 3719
"clear multicast bandwidth-admission" on page 3720	"clear multicast forwarding-cache" on page 3723	"clear multicast scope" on page 3724
"clear multicast sessions" on page 3726	"clear multicast snooping statistics" on page 3728	"clear multicast statistics" on page 3730
"clear mvrp statistics" on page 3732	"clear neighbor-discovery-inspection statistics" on page 3734	"clear network access nas-5g statistics" on page 3746

"clear network-access aaa saml assertion-cache" on page 3757	"clear network-access aaa saml idp-metadata" on page 3759	"clear network-access aaa statistics" on page 3736
"clear network-access aaa subscriber" on page 3741	"clear network-access gx-plus replay" on page 3743	"clear network-access gx-plus statistics" on page 3745
"clear network-access ocs statistics" on page 3747	"clear network-access pcrf" on page 3749	"clear network-access radsec state" on page 3751
"clear network-access radsec statistics" on page 3752	"clear network-access requests pending" on page 3753	"clear network-access requests statistics" on page 3755
"clear network-access securid-node-secret-file" on page 3756	"clear node reboot" on page 3760	"clear oam ethernet connectivity-fault-management continuity-measurement" on page 3761
"clear oam ethernet connectivity-fault-management delay-statistics" on page 3763	"clear oam ethernet connectivity-fault-management event" on page 3765	"clear oam ethernet connectivity-fault-management linktrace path-database" on page 3767
"clear oam ethernet connectivity-fault-management loss-statistics" on page 3768	"clear oam ethernet connectivity-fault-management path-database" on page 3770	"clear oam ethernet connectivity-fault-management policer" on page 3772
"clear oam ethernet connectivity-fault-management sla-iterator-statistics" on page 3774	"clear oam ethernet connectivity-fault-management statistics" on page 3775	"clear oam ethernet connectivity-fault-management synthetic-loss-measurement" on page 3777
"clear oam ethernet link-fault-management state" on page 3779	"clear oam ethernet link-fault-management statistics" on page 3781	"clear (ospf | ospf3) database" on page 3784
"clear (ospf | ospf3) database" on page 3784	"clear (ospf | ospf3) database" on page 3784	"clear (ospf | ospf3) database-protection" on page 3788
"clear (ospf | ospf3) io-statistics" on page 3789	"clear (ospf | ospf3) neighbor" on page 3791	"clear (ospf | ospf3) overload" on page 3793

"clear (ospf | ospf3) statistics" on page 3795	"clear ovsdb commit failures" on page 3798	"clear passive-monitoring statistics" on page 3800
"clear path-computation-client statistics" on page 3802	"clear performance-monitoring mpls lsp" on page 3803	"clear pfe sflow egress-mcast-stats" on page 3805
"clear pim join" on page 3806	"clear pim join-distribution" on page 3809	"clear pim register" on page 3812
"clear pim snooping join" on page 3814	"clear pim snooping statistics" on page 3817	"clear pim statistics" on page 3819
"clear policy statistics" on page 3823	"clear pppoe lockout" on page 3824	"clear pppoe lockout atm-identifier" on page 3827
"clear pppoe lockout vlan-identifier" on page 3829	"clear pppoe statistics" on page 3832	"clear protection-group ethernet-ring statistics" on page 3834
"clear ptp statistics stream" on page 3836	"clear remote-device-management statistics" on page 3837	"clear rip general-statistics" on page 3838
"clear rip statistics" on page 3840	"clear ripng general-statistics" on page 3842	"clear ripng statistics" on page 3844
"clear rsvp session" on page 3846	"clear rsvp statistics" on page 3848	"clear security advance-policy-based-routing sla statistics " on page 3850
"clear security alarms" on page 3851	"clear security alg h323 counters" on page 3855	"clear security alg ike-esp-nat" on page 3857
"clear security alg sccp calls" on page 3858	"clear security alg sccp counters" on page 3860	"clear security alg sip calls" on page 3862
"clear security alg sip counters" on page 3863	"clear security application-firewall rule-set statistics" on page 3865	"clear security application-firewall rule-set statistics logical-system" on page 3866

"clear security datapath-debug counters" on page 3868	"clear security dns-cache" on page 3869	"clear security dynamic-vpn all" on page 3871
"clear security dynamic-vpn user" on page 3872	"clear security firewall-authentication history" on page 3874	"clear security firewall-authentication history address" on page 3876
"clear security firewall-authentication history identifier" on page 3878	"clear security firewall-authentication users" on page 3881	"clear security firewall-authentication users address" on page 3883
"clear security firewall-authentication users identifier" on page 3885	"clear security flow ip-action " on page 3887	"clear security flow session" on page 3890
"clear security flow session application" on page 3893	"clear security flow session application-traffic-control" on page 3896	"clear security flow session conn-tag" on page 3899
"clear security flow session destination-port" on page 3901	"clear security flow session destination-prefix" on page 3903	"clear security flow session family" on page 3905
"clear security flow session IDP" on page 3907	"clear security flow session interface" on page 3910	"clear security flow session nat" on page 3912
"clear security flow session protocol" on page 3915	"clear security flow session resource-manager" on page 3918	"clear security flow session services-offload" on page 3920
"clear security flow session session-identifier" on page 3924	"clear security flow session source-port" on page 3926	"clear security flow session source-prefix" on page 3928
"clear security flow session tenant" on page 3931	"clear security flow session tunnel" on page 3933	"clear security forward-options mirror filter" on page 3935
"clear security group-vpn member group" on page 3937	"clear security group-vpn member ike security-associations" on page 3938	"clear security group-vpn member ipsec security-associations" on page 3939

"clear security group-vpn member ipsec security-associations statistics" on page 3941	"clear security group-vpn member ipsec statistics" on page 3942	"clear security group-vpn member kek security-associations" on page 3943
"clear security group-vpn server" on page 3944	"clear security group-vpn server server-cluster statistics" on page 3946	"clear security group-vpn server statistics" on page 3947
"clear security gtp counters" on page 3949	"clear security gtp sctp association" on page 3952	"clear security gtp sctp counters" on page 3954
"clear security host-vpn security-associations" on page 3956	"clear security idp" on page 3957	"clear security idp attack table" on page 3959
"clear security idp counters action" on page 3961	"clear security idp counters application-identification" on page 3962	"clear security idp counters dfa" on page 3964
"clear security idp counters flow" on page 3965	"clear security idp counters http-decoder" on page 3966	"clear security idp counters ips" on page 3968
"clear security idp counters log" on page 3969	"clear security idp counters memory" on page 3971	"clear security idp counters packet" on page 3972
"clear security idp counters packet-log" on page 3974	"clear security idp counters pdf-decoder" on page 3975	"clear security idp counters policy-manager" on page 3977
"clear security idp counters ssl-inspection" on page 3978	"clear security idp counters tcp-reassembler" on page 3980	"clear security idp ssl-inspection session-id-cache" on page 3981
"clear security ike active-peer aaa-username" on page 3983	"clear security ike peers" on page 3984	"clear security ike peers statistics" on page 3985
"clear security ike respond-bad-spi-count" on page 3987	"clear security ike security-associations" on page 3988	"clear security ike stats" on page 3990

"clear security ipsec security-associations" on page 3993	"clear security ipsec statistics" on page 3995	"clear security ipsec tunnel-events-statistics" on page 3997
"clear security log" on page 3998	"clear security log file" on page 4001	"clear security log stream file" on page 4002
"clear security mka statistics" on page 4004	"clear security mka statistics (MX Series)" on page 4006	"clear security nat incoming-table" on page 4007
"clear security nat source persistent-nat-table" on page 4009	"clear security nat statistics destination pool" on page 4010	"clear security nat statistics destination rule" on page 4011
"clear security nat statistics source pool" on page 4013	"clear security nat statistics source rule" on page 4014	"clear security nat statistics static rule" on page 4015
"clear security pki ca-certificate" on page 4017	"clear security pki certificate-request" on page 4018	"clear security pki crl" on page 4020
"clear security pki key-pair" on page 4021	"clear security pki key-pair (Local Certificate)" on page 4023	"clear security pki local-certificate" on page 4024
"clear security pki local-certificate" on page 4026	"clear security pki local-certificate (Device)" on page 4027	"clear security pki node-local key-pair" on page 4030
"clear security pki node-local local-certificate" on page 4031	"clear security pki statistics" on page 4033	"clear security policies file-serialization counters" on page 4034
"clear security policies hit-count" on page 4035	"clear security policies statistics" on page 4037	"clear security screen statistics" on page 4038
"clear security screen statistics interface" on page 4039	"clear security screen statistics zone" on page 4042	"clear security ssh key-pair-identity" on page 4044
"clear security tcp-encap statistics" on page 4046	"clear security user-identification local-authentication-table" on page 4047	"clear security utm anti-spam statistics" on page 4048

"clear security utm antivirus statistics" on page 4051	"clear security utm content-filtering statistics" on page 4054	"clear security utm session" on page 4057
"clear security utm web-filtering statistics" on page 4058	"clear service l2tp session routing-instance" on page 4062	"clear service user-identification identity-management counter" on page 4063
"clear services accounting statistics inline-jflow" on page 4064	"clear services advanced-anti-malware statistics" on page 4066	"clear services agf nas statistics" on page 4068
"clear services agf ngap amf" on page 4069	"clear services alg statistics" on page 4070	"clear services application-aware-access-list statistics" on page 4071
"clear services application-identification application-statistics" on page 4072	"clear services application-identification application-statistics cumulative" on page 4073	"clear services application-identification application-statistics interval" on page 4074
"clear services application-identification application-system-cache" on page 4075	"clear services application-identification application-system-cache (Junos OS)" on page 4076	"clear services application-identification counter (Values)" on page 4078
"clear services application-identification packet-capture counters" on page 4081	"clear services application-identification statistics" on page 4082	"clear services captive-portal-content-delivery statistics" on page 4085
"clear services cos statistics" on page 4086	"clear services crtp statistics" on page 4088	"clear services dns-filtering cache" on page 4089
"clear services dynamic-flow-capture" on page 4093	"clear services flow-collector statistics" on page 4094	"clear services icap-redirect statistic" on page 4096
"clear services ids" on page 4098	"clear services ids destination-table" on page 4100	"clear services ids pair-table" on page 4101
"clear services ids source-table" on page 4103	"clear services inline nat pool" on page 4106	"clear services inline nat statistics" on page 4107

"clear services inline softwire statistics" on page 4109	"clear services inline-monitoring statistics" on page 4105	"clear services ipsec-vpn certificates" on page 4110
"clear services ipsec-vpn ike security-associations" on page 4112	"clear services ipsec-vpn ipsec security-associations" on page 4113	"clear services ipsec-vpn ipsec statistics" on page 4115
"clear services l2tp destination" on page 4117	"clear services l2tp destination lockout" on page 4119	"clear services l2tp destination statistics" on page 4121
"clear services l2tp multilink" on page 4123	"clear services l2tp session" on page 4125	"clear services l2tp session statistics" on page 4129
"clear services l2tp tunnel" on page 4131	"clear services l2tp tunnel statistics" on page 4134	"clear services lrf collector statistics" on page 4136
"clear services lrf statistics" on page 4138	"clear services mobile-edge sessions" on page 4139	"clear services monitoring rfc2544" on page 4140
"clear services monitoring twamp server control-connection" on page 4141	"clear services nat flows" on page 4142	"clear services nat mappings" on page 4144
"clear services nat mappings app" on page 4147	"clear services nat mappings eim" on page 4149	"clear services nat mappings pcp" on page 4151
"clear services nat source mappings" on page 4154	"clear services redundancy-set last-saved-state id " on page 4157	"clear services rpm rfc2544-benchmarking" on page 4158
"clear services rpm twamp server connection" on page 4159	"clear services security-intelligence dns-statistics" on page 4160	"clear services security-metadata-streaming" on page 4161
"clear services service-set statistics ids drops" on page 4164	"clear services service-sets statistics ids session-limits counters" on page 4165	"clear services service-sets statistics integrity-drops" on page 4166

"clear services service-sets statistics jflow-log" on page 4168	"clear services service-sets statistics packet-drops" on page 4170	"clear services service-sets statistics syslog" on page 4171
"clear services sessions" on page 4173	"clear services sessions analysis" on page 4178	"clear services ssl proxy statistics" on page 4179
"clear services stateful-firewall flows" on page 4180	"clear services stateful-firewall sip-call" on page 4183	"clear services stateful-firewall sip-register" on page 4187
"clear services stateful-firewall statistics" on page 4191	"clear services subscriber analysis" on page 4193	"clear services user-identification active-directory-access" on page 4194
"clear services user-identification authentication-table" on page 4195	"clear services user-identification logical-domain-identity-management counters" on page 4197	"clear services video-monitoring mdi errors fpc-slot" on page 4199
"clear services video-monitoring mdi statistics fpc-slot" on page 4200	"clear services web-filter statistics profile" on page 4202	"clear sflow collector statistics (QFX Series)" on page 4204
"clear sflow collectors statistics" on page 4206	"clear snmp history" on page 4207	"clear snmp statistics" on page 4208
"clear spanning-tree protocol-migration" on page 4210	"clear spanning-tree statistics" on page 4212	"clear spanning-tree statistics bridge" on page 4214
"clear spanning-tree stp-buffer" on page 4215	"clear synchronous-ethernet esmc statistics" on page 4216	"clear synchronous-ethernet wait-to-restore" on page 4217
"clear system commit" on page 4219	"clear system commit prepared" on page 4221	"clear system errors" on page 4223
"clear system login lockout" on page 4225	"clear system reboot" on page 4226	"clear system services dhcp binding" on page 4230

"clear system services dhcp conflict" on page 4232	"clear system services dhcp statistics" on page 4233	"clear system services dns dns-proxy " on page 4235
"clear system subscriber-management statistics" on page 4236	"clear tcp-forwarding connections" on page 4237	"clear tcp-forwarding statistics" on page 4241
"clear trace" on page 4246	"clear unified-edge tdf aaa radius client statistics" on page 4248	"clear unified-edge tdf aaa radius network-element statistics" on page 4250
"clear unified-edge tdf aaa radius server statistics" on page 4252	"clear unified-edge tdf aaa radius snoop-segment statistics" on page 4253	"clear unified-edge tdf aaa statistics" on page 4255
"clear unified-edge tdf address-assignment pool" on page 4257	"clear unified-edge tdf address-assignment statistics" on page 4259	"clear unified-edge tdf call-admission-control statistics" on page 4261
"clear unified-edge tdf diameter network-element statistics" on page 4262	"clear unified-edge tdf diameter pcc-gx statistics" on page 4264	"clear unified-edge tdf diameter peer statistics" on page 4266
"clear unified-edge tdf statistics" on page 4268	"clear unified-edge tdf subscribers" on page 4269	"clear unified-edge tdf subscribers peer" on page 4271
"clear validation database" on page 4273	"clear validation session" on page 4275	"clear validation statistics" on page 4277
"clear virtual-chassis heartbeat (MX Series Virtual Chassis)" on page 4278	"clear virtual-chassis protocol statistics" on page 4281	"clear virtual-chassis vc-port statistics" on page 4283
"clear vpls mac-address" on page 4285	"clear vpls mac-move-action" on page 4287	"clear vpls mac-table" on page 4288
"clear vrrp" on page 4290	"clear-services-web-filter-dns-feed-statistics-feed-name" on page 4201	"commit" on page 4436

"configure" on page 4520	"copy" on page 4660	

D

"deactivate" on page 5142	"delete" on page 5276	"delete chassis fpc" on page 5281
[OBSOLETE] DELETE (DTCP)	"delete services jinsightd subscribe health-monitor" on page 5284	*[OBSOLETE] DISABLE (DTCP)*

E

"edit" on page 6201	*[OBSOLETE] ENABLE (DTCP)*	"exit" on page 6638

F

"file " on page 7025	"file archive" on page 7032	"file checksum md5" on page 7037
"file checksum sha1" on page 7041	"file checksum sha-256" on page 7039	"file compare" on page 7044
"file copy" on page 7049	"file delete" on page 7052	"file list" on page 7062
"file rename" on page 7072	"file show" on page 7078	

H

"help" on page 8045	

I

"insert" on page 8738	

L

[OBSOLETE] LIST (DTCP)	"load" on page 9867	"load access-security slaac-snooping persistent-file" on page 9870

M

"monitor ethernet delay-measurement" on page 10836	"monitor ethernet loss-measurement" on page 10843	"monitor ethernet synthetic-loss-measurement" on page 10850
"monitor interface" on page 10856	"monitor interface error" on page 10874	"monitor label-switched-path" on page 10877
"monitor list" on page 10881	"monitor list (JDM)" on page 10883	"monitor mpls delay rsvp" on page 10884
"monitor mpls loss rsvp" on page 10896	"monitor mpls loss-delay rsvp" on page 10890	"monitor pfe traffic interface" on page 10908

"monitor security flow file" on page 10914	"monitor security flow filter" on page 10916	"monitor security flow start" on page 10918
"monitor security flow stop" on page 10920	"monitor security packet-drop" on page 10921	"monitor start" on page 10927
"monitor start (JDM)" on page 10930	"monitor stop" on page 10932	"monitor stop (JDM)" on page 10934
"monitor traffic" on page 10935	"mtrace" on page 11010	"mtrace from-source" on page 11013
"mtrace monitor" on page 11017	"mtrace to-gateway" on page 11019	

O

"op" on page 11833	"op invoke-debugger cli" on page 11836	

P

"ping" on page 12525	"ping clns" on page 12534	"ping fabric multicast-flow" on page 12539
"ping fabric unicast-flow" on page 12542	"ping mpls bgp" on page 12547	"ping mpls l2circuit" on page 12550
"ping mpls l2vpn" on page 12554	"ping mpls l3vpn" on page 12558	"ping mpls ldp" on page 12561
"ping mpls lsp-end-point" on page 12565	"ping mpls rsvp" on page 12568	"ping mpls segment routing isis" on page 12575

"ping mpls segment routing ospf" on page 12579	"ping mpls segment routing spring-te" on page 12583	"ping mpls segment routing static egress-ip" on page 12592
"ping overlay" on page 12596	"ping srv6" on page 12603	"ping vpls instance" on page 12608
"prbs-test-start" on page 12989	"prbs-test-stop" on page 12991	"probe (ICMP Router Discovery)" on page 13222
"protect" on page 13404		

Q

"quit" on page 13956		

R

"rename" on page 14371	"replace" on page 14376	"request" on page 14411
"request access-security router-advertisement-guard-block" on page 14413	"request access-security router-advertisement-guard-forward" on page 14414	"request access-security router-advertisement-guard-learn interface" on page 14415
"request access-security slaac-snooping unblock" on page 14417	"request ancp oam interface" on page 14419	"request ancp oam neighbor" on page 14421
"request ancp oam port-down" on page 14423	"request ancp oam port-up" on page 14425	"request app-engine cleanup" on page 14427

"request app-engine file-copy (crash | log) from-jhost to-vjunos" on page 14429	"request auto-configuration reconnect-pending" on page 14431	"request chassis afeb" on page 14434
"request chassis beacon" on page 14435	"request chassis cb" on page 14441	"request chassis cb (Security)" on page 14440
"request chassis ccg" on page 14447	"request chassis cfeb" on page 14449	"request chassis clock master switch" on page 14450
"request chassis cluster configuration-synchronize" on page 14452	"request chassis cluster failover node" on page 14454	"request chassis cluster failover redundancy-group" on page 14455
"request chassis cluster failover reset" on page 14457	"request chassis cluster in-service-upgrade abort (ISSU)" on page 14459	"request chassis device-mode" on page 14460
"request chassis device-mode satellite" on page 14463	"request chassis fabric fpc" on page 14465	"request chassis fabric pfe" on page 14467
"request chassis fabric plane" on page 14468	"request chassis feb" on page 14472	"request chassis fpc" on page 14475
"request chassis fpc slc" on page 14488	"request chassis fpc-control-port" on page 14486	"request chassis fpm resync" on page 14489
"request chassis high-availability failover services-redundancy-group" on page 14492	"request chassis mcs" on page 14494	"request chassis mic" on page 14495
"request chassis optics" on page 14499	"request chassis pcg" on page 14501	"request chassis pic" on page 14502
"request chassis pic-mode" on page 14509	"request chassis port-led" on page 14512	"request chassis primary-ha-control-port-transition" on page 14516

"request chassis redundancy feb slot" on page 14518	"request chassis routing-engine hard-disk-test" on page 14520	"request chassis routing-engine master" on page 14529
"request chassis satellite beacon" on page 14535	"request chassis satellite disable" on page 14537	"request chassis satellite enable" on page 14539
"request chassis satellite file-copy" on page 14541	"request chassis satellite install" on page 14542	"request chassis satellite interface" on page 14545
"request chassis satellite login" on page 14547	"request chassis satellite reboot" on page 14548	"request chassis satellite restart" on page 14550
"request chassis satellite shell-command" on page 14551	"request chassis scg" on page 14553	"request chassis sfb" on page 14555
"request chassis sfm" on page 14558	"request chassis sfm master switch" on page 14559	"request chassis sib" on page 14562
"request chassis spmb restart" on page 14565	"request chassis ssb master switch" on page 14568	"request chassis synchronization mode" on page 14570
"request chassis synchronization switch" on page 14574	"request chassis system-mode" on page 14577	"request component login" on page 14580
"request csds add-vsrx" on page 14583	"request csds authenticate-host" on page 14584	"request csds delete-vsrx" on page 14586
"request csds extract-vsrx-keys" on page 14587	"request csds jdm" on page 14589	"request csds sync-controller" on page 14592
"request debug information" on page 14594	"request dhcp client renew" on page 14595	"request dhcp relay bulk-leasequery" on page 14597
"request dhcp relay leasequery" on page 14599	"request dhcp server reconfigure" on page 14601	"request dhcpv6 client renew" on page 14604

"request dhcpv6 relay bulk-leasequery" on page 14605	"request dhcpv6 relay leasequery" on page 14607	"request dhcpv6 server reconfigure" on page 14609
"request diagnostics tdr" on page 14611	"request extension-service restart-daemonize-app" on page 14614	"request extension-service (start | stop)" on page 14615
"request fabric administration director-group change-master" on page 14616	"request fabric administration power-off interconnect-device" on page 14618	"request fabric administration power-off interconnect-device fpc" on page 14620
"request fabric administration power-off node-device" on page 14622	"request fabric administration remove" on page 14624	"request fabric administration system mac-pool add" on page 14626
"request fabric administration system mac-pool delete" on page 14628	"request fibre-channel proxy load-rebalance" on page 14630	"request flight-recorder set high-cpu" on page 14633
"request forwarding-options port-mirroring instance family self-mirror-start" on page 14636	"request forwarding-options port-mirroring instance family self-mirror-stop" on page 14638	"request interface bounce" on page 14639
"request interface link-degrade-recover" on page 14642	"request interface load-balancing revert (Aggregated Multiservices)" on page 14646	"request interface load-balancing switchover (Aggregated Multiservices)" on page 14647
"request interface mc-ae switchover (Multichassis Link Aggregation)" on page 14649	"request interface optics-reset" on page 14651	"request interface rebalance (Aggregated Ethernet for Subscriber Management)" on page 14653
"request interface revert" on page 14657	"request interface (revert | switchover) (Adaptive Services)" on page 14658	"request interface (revert | switchover) (Aggregated Ethernet Link Protection)" on page 14660
"request interface (revert | switchover) (Aggregated Inline Service Interfaces)" on page 14662	"request interface switchover" on page 14664	"request ipsec switch" on page 14665

"request jdaf clear statistics service" on page 14667	"request jnu role" on page 14668	"request jnu satellite sync" on page 14670
"request jnu schema add" on page 14671	"request jnu schema delete" on page 14673	"request jnu schema merge" on page 14674
"request l2circuit-switchover" on page 14675	"request lacp link-switchover" on page 14677	"request message" on page 14680
"request modem wireless create-profile" on page 14681	"request modem wireless fota" on page 14684	"request modem wireless sim-lock" on page 14685
"request modem wireless sim-unlock" on page 14687	"request mpls container-lsp" on page 14689	"request mpls lsp adjust-autobandwidth" on page 14691
"request network-access aaa accounting" on page 14693	"request network-access aaa replay pending-accounting-stops" on page 14695	"request network-access aaa saml export-sp-metadata" on page 14706
"request network-access aaa saml load-idp-metadata" on page 14707	"request network-access aaa subscriber add session-id" on page 14696	"request network-access aaa subscriber delete session-id " on page 14699
"request network-access aaa subscriber modify session-id" on page 14702	"request network-access aaa subscriber set session-id" on page 14704	"request node halt (Junos OS Evolved)" on page 14708
"request node (offline | online) (Junos OS Evolved)" on page 14711	"request node power-off (Junos OS Evolved)" on page 14713	"request node power-on (Junos OS Evolved)" on page 14715
"request node reboot (re0 | re1) (Junos OS Evolved)" on page 14717	"request packet-capture start" on page 14720	"request packet-capture stop" on page 14726
"request path-computation-client active-pce" on page 14728	"request pim multicast-tunnel rebalance" on page 14729	"request pppoe connect" on page 14731

"request pppoe disconnect " on page 14732	"request protection-group ethernet-aps" on page 14734	"request protection-group ethernet-ring" on page 14736
"request redundant-power-system multi-backup" on page 14739	"request restart packet-forwarding-engine" on page 14742	"request routing-engine" on page 14743
"request security certificate enroll (Signed)" on page 14745	"request security certificate enroll (Unsigned)" on page 14747	"request security datapath-debug capture start" on page 14749
"request security datapath-debug capture stop" on page 14750	"request security idp jist-conversion" on page 14751	"request security idp security package rollback" on page 14761
"request security idp security-package download" on page 14754	"request security idp security-package install" on page 14757	"request security idp security-package offline-download" on page 14760
"request security idp ssl-inspection key add" on page 14763	"request security idp ssl-inspection key delete" on page 14766	"request security idp storage-cleanup " on page 14768
"request security ike debug-disable" on page 14770	"request security ike debug-enable" on page 14771	"request security integrity" on page 14773
"request security internal-security-association refresh" on page 14774	"request security key-pair" on page 14776	"request security pki ca-certificate ca-profile-group default-trusted-ca-certs" on page 14778
"request security pki ca-certificate ca-profile-group load" on page 14781	"request security pki ca-certificate enroll" on page 14784	"request security pki ca-certificate enroll (Security)" on page 14786
"request security pki ca-certificate load" on page 14787	"request security pki ca-certificate load (Security)" on page 14789	"request security pki ca-certificate verify" on page 14791
"request security pki ca-certificate verify (Security)" on page 14792	"request security pki ca-profile-group-status" on page 14795	"request security pki crl load" on page 14797

"request security pki crl load (Security)" on page 14798	"request security pki generate-certificate-request" on page 14800	"request security pki generate-certificate-request (Security)" on page 14802
"request security pki generate-key-pair" on page 14805	"request security pki generate-key-pair" on page 14806	"request security pki generate-key-pair (Security)" on page 14808
"request security pki key-pair export" on page 14810	"request security pki local-certificate enroll" on page 14812	"request security pki local-certificate enroll cmpv2" on page 14814
"request security pki local-certificate enroll scep" on page 14817	"request security pki local-certificate export" on page 14822	"request security pki local-certificate generate-self-signed" on page 14823
"request security pki local-certificate generate-self-signed" on page 14825	"request security pki local-certificate generate-self-signed (Security)" on page 14827	"request security pki local-certificate load" on page 14830
"request security pki local-certificate re-enroll cmpv2" on page 14831	"request security pki local-certificate re-enroll scep" on page 14833	"request security pki local-certificate verify" on page 14835
"request security pki node-local generate-certificate-request" on page 14837	"request security pki node-local generate-key-pair" on page 14840	"request security pki node-local key-pair export" on page 14842
"request security pki node-local local-certificate enroll" on page 14843	"request security pki node-local local-certificate export" on page 14847	"request security pki node-local local-certificate load" on page 14848
"request security pki node-local local-certificate re-enroll" on page 14850	"request security pki node-local local-certificate verify" on page 14852	"request security pki sync-from-peer" on page 14854
"request security pki verify-integrity-status" on page 14856	"request security policies check" on page 14857	"request security policies resync" on page 14861

"request security re-distribution ipsec-vpn" on page 14863	"request security ssh key-pair-identity generate" on page 14865	"request security ssh password-less-authentication (Junos OS Evolved)" on page 14867
"request security system-keystore" on page 14869	"request security system-keystore unlink" on page 14870	"request security tpm master-encryption-password set" on page 14871
"request security uefi revoke" on page 14873	"request security user-identification local-authentication-table delete" on page 14875	"request security user-identification local-authorization-table add" on page 14877
"request security utm anti-virus avira-engine" on page 14879	"request security utm anti-virus juniper-express-engine" on page 14881	"request security utm anti-virus kaspersky-lab-engine" on page 14883
"request security utm anti-virus sophos-engine" on page 14885	"request security utm custom-objects url-feed update" on page 14886	"request security utm custom-objects url-feed update feed-name force" on page 14888
"request security utm web-filtering categorize" on page 14889	"request security utm web-filtering category download [version]" on page 14892	"request security utm web-filtering category download-install [version]" on page 14891
"request security utm web-filtering category install" on page 14893	"request security utm web-filtering category migrate-to-ng-juniper" on page 14895	"request security utm web-filtering category uninstall" on page 14897
"request security utm web-filtering custom-page reload" on page 14898	"request security utm web-filtering recategorize" on page 14899	"request server authenticate-peer-server" on page 14901
"request services advanced-anti-malware data-connection" on page 14903	"request services advanced-anti-malware diagnostic" on page 14905	"request services advanced-anti-malware redirect-file" on page 14910
"request services anti-virus machine-learning-engine" on page 14914	"request services anti-virus update" on page 14912	"request services application identification install ignore duplicate version check" on page 14928

"request services application identification new to production" on page 14931	"request services application-identification application" on page 14916	"request services application-identification clear packet-capture all" on page 14918
"request services application-identification download " on page 14919	"request services application-identification download status" on page 14923	"request services application-identification group" on page 14925
"request services application-identification install" on page 14934	"request services application-identification install status" on page 14929	"request services application-identification offline-download package-path " on page 14936
"request services application-identification proto-bundle-status" on page 14937	"request services application-identification rollback status" on page 14933	"request services application-identification uninstall" on page 14939
"request services application-identification uninstall status" on page 14940	"request services extensible-subscriber-services reload-dictionary" on page 14942	"request services flow-collector change-destination primary interface" on page 14943
"request services flow-collector change-destination secondary interface" on page 14945	"request services flow-collector test-file-transfer" on page 14946	"request services ip-monitoring preempt-restore policy " on page 14948
"request services ipsec-vpn ipsec switch tunnel" on page 14950	"request services l2tp destination unlock" on page 14951	"request services monitoring twamp client" on page 14953
"request services paa install" on page 14954	"request services paa uninstall" on page 14958	"request services redundancy-set trigger" on page 14960
"request services remote-device-management reconfigure service-device" on page 14961	"request services remote-device-management reload-dictionary" on page 14962	"request services rpm twamp" on page 14964
"request services static-subscribers login group" on page 14966	"request services static-subscribers login interface" on page 14968	"request services static-subscribers logout group" on page 14969

"request services static-subscribers logout interface" on page 14971	"request services subscribers" on page 14972	"request services subscribers clear" on page 14974
"request services url-filter delete gencfg-data" on page 14975	"request services url-filter force dns-resolution" on page 14976	"request services url-filter update url-filter-database file" on page 14979
"request services url-filter validate" on page 14980	"request services user-identification active-directory-access active-directory-authentication-table delete" on page 14981	"request services user-identification active-directory-access domain-controller" on page 14983
"request services user-identification active-directory-access ip-user-probe" on page 14985	"request services user-identification authentication-source aruba-clearpass user-query" on page 14987	"request services user-identification authentication-source jims groups domain <domain-name> (force-fetch|status)" on page 14989
"request services user-identification authentication-source jims validate (user <user-name>|group <group-name>|device <device-name>) domain <domain-name>" on page 14991	"request services user-identification authentication-table delete " on page 14994	"request services web-filter delete gencfg-data" on page 15004
"request services web-filter force dns-resolution" on page 15005	"request services web-filter update dns-filter-database" on page 15006	"request services web-filter update url-filter-database file" on page 15007
"request services web-filter validate dns-filter-file-name" on page 15009	"request services web-filter validate url-filter-file-name" on page 15010	"request session member" on page 15011
"request snmp spoof-trap" on page 15012	"request snmp utility-mib clear instance" on page 15021	"request snmp utility-mib set instance" on page 15023
"request support information" on page 15024	"request system application (Junos OS Evolved)" on page 15042	"request system autorecovery state" on page 15044

"request system boot console" on page 15047	"request system certificate add" on page 15048	"request system commit server pause" on page 15050
"request system commit server queue cleanup" on page 15051	"request system commit server start" on page 15053	"request system configuration database resize" on page 15054
"request system configuration rescue delete" on page 15057	"request system configuration rescue save" on page 15059	"request system convert-json-configuration" on page 15060
"request system debug-info" on page 15062	"request system decrypt password" on page 15066	"request system disable unified-services" on page 15068
"request system download abort" on page 15069	"request system download clear" on page 15071	"request system download pause" on page 15072
"request system download resume" on page 15074	"request system download start" on page 15076	"request system enable unified-services" on page 15078
"request system filesystem encryption enable" on page 15079	"request system filesystem encryption keys delete" on page 15082	"request system firmware downgrade optics" on page 15085
"request system firmware reload (Junos OS Evolved)" on page 15087	"request system firmware upgrade" on page 15089	"request system firmware upgrade (Junos OS Evolved)" on page 15099
"request system firmware upgrade poe" on page 15102	"request system halt" on page 15105	"request system license update" on page 15115
"request system logout" on page 15117	"request system malware-scan" on page 15119	"request system partition abort" on page 15122
"request system partition compact-flash" on page 15125	"request system partition hard-disk" on page 15127	"request system power-off" on page 15131
"request system power-off (SRX Series)" on page 15135	"request system process terminate" on page 15138	"request system reboot (Junos OS)" on page 15140

"request system reboot (Junos OS Evolved)" on page 15151	"request system reboot (Junos OS with Upgraded FreeBSD)" on page 15154	"request system reboot (SRX Series)" on page 15159
"request system reboot usb (SRX Series)" on page 15160	"request system recover" on page 15161	"request system restore state path" on page 15165
"request system save state path" on page 15166	"request system scripts add" on page 15171	"request system scripts add (Chassis Cluster)" on page 15167
"request system scripts convert" on page 15172	"request system scripts delete" on page 15175	"request system scripts event-scripts reload" on page 15176
"request system scripts refresh-from" on page 15177	"request system scripts rollback" on page 15181	"request system scripts synchronize" on page 15182
"request system services dhcp" on page 15185	"request system snapshot (Junos OS Evolved)" on page 15187	"request system snapshot (Junos OS with FreeBSD Prior to Release 10)" on page 15190
"request system snapshot (Junos OS with Upgraded FreeBSD)" on page 15203	"request system snapshot (Maintenance)" on page 15208	"request system software abort" on page 15212
"request system software abort in-service-upgrade (ICU)" on page 15215	"request system software add (Junos OS)" on page 15216	"request system software add (Junos OS Evolved)" on page 15233
"request system software add (Maintenance)" on page 15238	"request system software add optional://auto-snapshot" on page 15240	"request system software add restart" on page 15242
"request system software configuration-backup" on page 15245	"request system software configuration-restore" on page 15247	"request system software delete (Junos OS)" on page 15250
"request system software delete (Junos OS Evolved)" on page 15254	"request system software delete-backup" on page 15248	"request system software download" on page 15258

"request system software format-qfabric-backup" on page 15260	"request system software in-service-upgrade" on page 15262	"request system software in-service-upgrade (Maintenance)" on page 15282
"request system software in-service-upgrade (MX Series 5G Universal Routing Platforms and EX9200 Switches)" on page 15288	"request system software nonstop-upgrade" on page 15319	"request system software nonstop-upgrade (QFabric Systems)" on page 15311
"request system software recover-from-restore-point" on page 15332	"request system software recovery-package" on page 15334	"request system software restore-point" on page 15338
"request system software rollback (Junos OS)" on page 15340	"request system software rollback (Junos OS Evolved)" on page 15345	"request system software sync" on page 15348
"request system software system-backup" on page 15353	"request system software validate in-service-upgrade" on page 15355	"request system software validate (Junos OS)" on page 15360
"request system software validate (Junos OS Evolved)" on page 15365	"request system software validate on (Junos OS with Upgraded FreeBSD)" on page 15368	"request system software validate-restart (Junos OS Evolved)" on page 15371
"request system storage cleanup (Junos OS)" on page 15374	"request system storage cleanup (Junos OS Evolved)" on page 15388	"request system storage cleanup (SRX Series)" on page 15400
"request system yang add" on page 15404	"request system yang delete" on page 15408	"request system yang disable" on page 15410
"request system yang enable" on page 15413	"request system yang update" on page 15414	"request system yang validate" on page 15417
"request system zeroize" on page 15419	"request system zeroize (Junos OS)" on page 15421	"request system zeroize (Junos OS Evolved)" on page 15425
"request system zeroize ztp-option secure" on page 15430	"request unified-edge tdf call-trace clear" on page 15435	"request unified-edge tdf call-trace show" on page 15436

"request unified-edge tdf call-trace start" on page 15440	"request unified-edge tdf call-trace stop" on page 15443	"request validation policy" on page 15445
"request virtual-chassis device-reachability" on page 15447	"request virtual-chassis member-id delete (MX Series Virtual Chassis)" on page 15451	"request virtual-chassis member-id set" on page 15452
"request virtual-chassis mode" on page 15455	"request virtual-chassis reactivate" on page 15461	"request virtual-chassis recycle" on page 15462
"request virtual-chassis renumber" on page 15464	"request virtual-chassis routing-engine master switch" on page 15465	"request virtual-chassis vc-port" on page 15468
"request virtual-chassis vc-port (Dedicated VCP)" on page 15472	"request virtual-chassis vc-port delete (MX Series Virtual Chassis)" on page 15474	"request virtual-chassis vc-port diagnostics optics" on page 15476
"request virtual-chassis vc-port set (MX Series Virtual Chassis)" on page 15478	"request virtual-network-functions" on page 15481	"request vmhost cleanup" on page 15482
"request vmhost cleanup (NFX Series)" on page 15485	"request vmhost copy jnode-to-vjunos" on page 15486	"request vmhost copy vjunos-to-jnode" on page 15487
"request vmhost file-copy" on page 15488	"request vmhost file-copy (NFX Series)" on page 15490	"request vmhost halt" on page 15492
"request vmhost halt (NFX Series)" on page 15494	"request vmhost hard-disk-test" on page 15495	"request vmhost jdm add (In-Chassis Model)" on page 15497
"request vmhost jdm delete (In-Chassis Model)" on page 15499	"request vmhost jdm login (In-Chassis Model)" on page 15500	"request vmhost jdm start (In-Chassis Model)" on page 15501
"request vmhost jdm stop (In-Chassis Model)" on page 15502	"request vmhost local-file-copy" on page 15503	"request vmhost mode" on page 15506

"request vmhost power-off" on page 15525	"request vmhost power-off (NFX Series)" on page 15526	"request vmhost power-on" on page 15528
"request vmhost reboot" on page 15529	"request vmhost reboot (NFX Series)" on page 15532	"request vmhost remote-file-copy" on page 15535
"request vmhost snapshot" on page 15539	"request vmhost software abort in-service-upgrade" on page 15542	"request vmhost software add" on page 15543
"request vmhost software add (NFX Series)" on page 15548	"request vmhost software in-service-upgrade" on page 15551	"request vmhost software rollback" on page 15556
"request vmhost software validate" on page 15559	"request vmhost storage" on page 15561	"request vmhost zeroize" on page 15563
"request vmhost zeroize ztp-option secure" on page 15566	"request vpls-switchover" on page 15568	"request wlan access-point packet capture" on page 15570
"request-system-integrity-attestation" on page 15112	"request-system-integrity-log-retrieval" on page 15114	"request-vmhost-power-cycle-cb" on page 15507
"request-vmhost-power-cycle-re" on page 15516	"restart" on page 15633	"restart chassis-control" on page 15648
"restart extensible-subscriber-services" on page 15656	"restart (JDM)" on page 15657	"restart (Junos OS Evolved)" on page 15660
"rollback" on page 15830	"run" on page 16162	

S

"save" on page 16315	"set" on page 16860	"set chassis cluster cluster-id node node-number reboot" on page 16861

"set chassis cluster disable reboot" on page 16863	"set chassis display message" on page 16865	"set chassis virtual-chassis" on page 16869
"set cli complete-on-space" on page 16870	"set cli directory" on page 16872	"set cli idle-timeout" on page 16874
"set cli logical-system" on page 16875	"set cli prompt" on page 16877	"set cli restart-on-upgrade" on page 16878
"set cli screen-length" on page 16880	"set cli screen-width" on page 16881	"set cli terminal" on page 16882
"set cli timestamp" on page 16884	"set date" on page 16885	"set date ntp" on page 16888
"set services jinsightd subscribe health-monitor" on page 16892	"set services jinsightd traceoptions" on page 16894	"show" on page 16994
"show access-cac interface-set" on page 16995	"show access-security router-advertisement state" on page 16998	"show access-security router-advertisement statistics" on page 17001
"show access-security slaac-snooping binding" on page 17003	"show access-security slaac-snooping state" on page 17005	"show access-security slaac-snooping statistics" on page 17007
"show accounting pending-accounting-stops" on page 17011	"show accounting profile" on page 17016	"show accounting server statistics archival-transfer" on page 17023
"show agent sensors" on page 17024	"show amt statistics" on page 17033	"show amt summary" on page 17037
"show amt tunnel" on page 17039	"show analyzer" on page 17044	"show ancp cos" on page 17046
"show ancp neighbor" on page 17054	"show ancp statistics" on page 17066	"show ancp subscriber" on page 17072
"show ancp summary" on page 17083	"show ancp summary neighbor" on page 17086	"show ancp summary subscriber" on page 17089

"show app-engine crash" on page 17091	"show app-engine logs" on page 17093	"show arp" on page 17096
"show arp inspection statistics" on page 17103	"show as-path" on page 17105	"show as-path domain" on page 17111
"show as-path summary" on page 17113	"show auto-bandwidth" on page 17116	"show auto-configuration out-of-band" on page 17120
"show backup-selection" on page 17126	"show bfd seamless session" on page 17132	"show bfd session" on page 17134
"show bfd subscriber session" on page 17144	"show bgp bmp" on page 17151	"show bgp group" on page 17153
"show bgp group output-queues" on page 17166	"show bgp group traffic-statistics" on page 17171	"show bgp neighbor" on page 17175
"show bgp output-scheduler" on page 17216	"show bgp replication" on page 17218	"show bgp replication logical-system" on page 17223
"show bgp summary" on page 17227	"show bridge domain" on page 17237	"show bridge flood" on page 17240
"show bridge flood nexthops satellite" on page 17249	"show bridge flood satellite" on page 17253	"show bridge mac-ip-table" on page 17257
"show bridge mac-table" on page 17264	"show bridge satellite device " on page 17272	"show bridge statistics" on page 17276
"show captive-portal authentication-failed-users" on page 17281	"show captive-portal firewall" on page 17283	"show captive-portal interface" on page 17286
"show chassis adc" on page 17291	"show chassis adc (GNF)" on page 17295	"show chassis adc (Node Slicing)" on page 17297

"show chassis afeb" on page 17299	"show chassis alarms" on page 17302	"show chassis alarms" on page 17304
"show chassis beacon" on page 17334	"show chassis cfeb" on page 17337	"show chassis cluster control-interface status" on page 17341
"show chassis cluster control-plane statistics" on page 17351	"show chassis cluster data-plane interfaces" on page 17354	"show chassis cluster data-plane statistics" on page 17356
"show chassis cluster ethernet-switching interfaces" on page 17359	"show chassis cluster ethernet-switching status" on page 17361	"show chassis cluster information" on page 17364
"show chassis cluster information configuration-synchronization" on page 17378	"show chassis cluster information issu" on page 17381	"show chassis cluster interfaces" on page 17383
"show chassis cluster ip-monitoring status redundancy-group" on page 17392	"show chassis cluster port-peering" on page 17396	"show chassis cluster statistics" on page 17399
"show chassis cluster status" on page 17405	"show chassis craft-interface" on page 17410	"show chassis dedicated-ukern-cpu" on page 17418
"show chassis device-mode" on page 17419	"show chassis enhanced-temperature-thresholds" on page 17423	"show chassis environment" on page 17429
"show chassis environment adc" on page 17553	"show chassis environment cb" on page 17573	"show chassis environment cb (Security)" on page 17569
"show chassis environment ccg" on page 17613	"show chassis environment fan" on page 17616	"show chassis environment fpc" on page 17621
"show chassis environment fpm" on page 17705	"show chassis environment mcs" on page 17714	"show chassis environment monitored" on page 17718

"show chassis environment pcg" on page 17748	"show chassis environment pdu" on page 17752	"show chassis environment pem" on page 17760
"show chassis environment power-supply-unit" on page 17782	"show chassis environment psm" on page 17784	"show chassis environment psu" on page 17798
"show chassis environment routing-engine" on page 17800	"show chassis environment scg" on page 17811	"show chassis environment (Security)" on page 17818
"show chassis environment sfb" on page 17825	"show chassis environment sfm" on page 17844	"show chassis environment sib" on page 17849
"show chassis errors active" on page 17895	"show chassis ethernet-switch" on page 17908	"show chassis ethernet-switch interconnect-device cb" on page 17960
"show chassis ethernet-switch interconnect-device fpc" on page 17985	"show chassis ethernet-switch (Security)" on page 17901	"show chassis fabric connectivity" on page 18022
"show chassis fabric degradation" on page 18031	"show chassis fabric degradation actions" on page 18033	"show chassis fabric degraded-fabric-reachability" on page 18035
"show chassis fabric destinations" on page 18037	"show chassis fabric device" on page 18066	"show chassis fabric errors" on page 18068
"show chassis fabric faults recovery-actions" on page 18075	"show chassis fabric feb" on page 18078	"show chassis fabric fpcs" on page 18079
"show chassis fabric map" on page 18091	"show chassis fabric mesh" on page 18101	"show chassis fabric plane" on page 18113
"show chassis fabric plane (Security)" on page 18104	"show chassis fabric plane-location" on page 18156	"show chassis fabric plane-location (Security)" on page 18154
"show chassis fabric reachability" on page 18168	"show chassis fabric redundancy-mode" on page 18181	"show chassis fabric sibs" on page 18183

"show chassis fabric stream-info" on page 18203	"show chassis fabric summary" on page 18212	"show chassis fabric summary (Security)" on page 18209
"show chassis fabric topology" on page 18220	"show chassis fabric unreachable-destinations" on page 18257	"show chassis fan" on page 18260
"show chassis feb" on page 18278	"show chassis firmware" on page 18283	"show chassis forwarding" on page 18304
"show chassis forwarding-options" on page 18306	"show chassis fpc" on page 18311	"show chassis fpc errors" on page 18387
"show chassis fpc ether-types" on page 18392	"show chassis fpc global pfe-id" on page 18400	"show chassis fpc (GNF)" on page 18404
"show chassis fpc (Node Slicing)" on page 18408	"show chassis fpc pic-status (GNF)" on page 18411	"show chassis fpc pic-status (Node Slicing)" on page 18413
"show chassis fpc (View)" on page 18415	"show chassis fpc-feb-connectivity" on page 18397	"show chassis gnf" on page 18427
"show chassis hardware" on page 18430	"show chassis hardware (GNF)" on page 18453	"show chassis hardware (View)" on page 18457
"show chassis high-availability data-plane statistics" on page 18476	"show chassis high-availability information" on page 18482	"show chassis high-availability peer-info" on page 18492
"show chassis high-availability prefix-srgid-table" on page 18495	"show chassis high-availability services-redundancy-group" on page 18497	"show chassis in-service-upgrade" on page 18507
"show chassis jnu satellite" on page 18513	"show chassis lcd" on page 18515	"show chassis led" on page 18534
"show chassis led satellite" on page 18551	"show chassis location" on page 18557	"show chassis mac-addresses" on page 18563

"show chassis network-services" on page 18571	"show chassis network-slices" on page 18575	"show chassis network-slices fpcs" on page 18579
"show chassis nonstop-upgrade" on page 18581	"show chassis nonstop-upgrade node-group" on page 18585	"show chassis oss-map" on page 18587
"show chassis pic" on page 18588	"show chassis pic (Security)" on page 18630	"show chassis pic-mode" on page 18628
"show chassis power" on page 18633	"show chassis power (Security)" on page 18676	"show chassis power sequence" on page 18688
"show chassis power-budget-statistics" on page 18681	"show chassis psd" on page 18690	"show chassis realtime-ukern-thread" on page 18692
"show chassis redundancy feb" on page 18694	"show chassis redundant-power-system" on page 18698	"show chassis reswatch " on page 18701
"show chassis routing-engine" on page 18703	"show chassis routing-engine (View)" on page 18720	"show chassis satellite" on page 18727
"show chassis satellite extended-port" on page 18759	"show chassis satellite interface" on page 18765	"show chassis satellite neighbor" on page 18772
"show chassis satellite power-budget-statistics" on page 18779	"show chassis satellite redundancy-group" on page 18783	"show chassis satellite redundancy-group devices" on page 18787
"show chassis satellite software" on page 18790	"show chassis satellite statistics" on page 18794	"show chassis satellite unprovision" on page 18801
"show chassis satellite upgrade-group" on page 18809	"show chassis satellite-cluster" on page 18743	"show chassis satellite-cluster route" on page 18751
"show chassis satellite-cluster statistics" on page 18754	"show chassis scb" on page 18813	"show chassis scb-ha-port" on page 18816

"show chassis security uefi variables" on page 18821	"show chassis sfb" on page 18823	"show chassis sfb errors" on page 18828
"show chassis sfm" on page 18831	"show chassis sibs" on page 18835	"show chassis spmb" on page 18853
"show chassis spmb sibs" on page 18867	"show chassis ssb" on page 18876	"show chassis synchronization" on page 18879
"show chassis synchronization (MX Series Routers)" on page 18894	*[OBSOLETE] Show chassis system-mode*	"show chassis temperature-thresholds" on page 18914
"show chassis temperature-thresholds (Security)" on page 18911	"show chassis usb storage" on page 18945	"show chassis zones" on page 18947
"show class-of-service" on page 18959	"show class-of-service adjustment-control-profile" on page 18965	"show class-of-service application-traffic-control counter" on page 18968
"show class-of-service application-traffic-control rate-limiters" on page 18974	"show class-of-service application-traffic-control statistics rate-limiter" on page 18986	"show class-of-service application-traffic-control statistics rule" on page 18990
"show class-of-service classifier" on page 18993	"show class-of-service code-point-aliases" on page 18999	"show class-of-service congestion-notification" on page 19001
"show class-of-service dedicated-buffer-profile" on page 19007	"show class-of-service drop-profile" on page 19009	"show class-of-service fabric scheduler-map" on page 19011
"show class-of-service fabric statistics" on page 19014	"show class-of-service forwarding-class" on page 19019	"show class-of-service forwarding-class-set" on page 19024
"show class-of-service forwarding-table" on page 19026	"show class-of-service forwarding-table classifier" on page 19031	"show class-of-service forwarding-table classifier mapping" on page 19033

"show class-of-service forwarding-table drop-profile" on page 19035	"show class-of-service forwarding-table fabric scheduler-map" on page 19038	"show class-of-service forwarding-table rewrite-rule" on page 19040
"show class-of-service forwarding-table rewrite-rule mapping" on page 19042	"show class-of-service forwarding-table scheduler-map" on page 19044	"show class-of-service forwarding-table slice" on page 19047
"show class-of-service forwarding-table slice mapping" on page 19049	"show class-of-service forwarding-table traffic-class-map" on page 19051	"show class-of-service fragmentation-map" on page 19054
"show class-of-service interface" on page 19056	"show class-of-service interface-set" on page 19100	"show class-of-service l2tp-session" on page 19103
"show class-of-service loss-priority-rewrite" on page 19106	"show class-of-service monitoring-profile" on page 19108	"show class-of-service multi-destination" on page 19110
"show class-of-service policy-map" on page 19112	"show class-of-service rewrite-rule" on page 19114	"show class-of-service routing-instance" on page 19117
"show class-of-service scheduler-hierarchy interface" on page 19119	"show class-of-service scheduler-hierarchy interface-set" on page 19122	"show class-of-service scheduler-map" on page 19125
"show class-of-service scheduler-map-forwarding-class-sets" on page 19131	"show class-of-service scheduler-resources fpc" on page 19134	"show class-of-service shared-buffer" on page 19137
"show class-of-service slice <slice_name> interface <interface_name>" on page 19140	"show class-of-service traffic-class-map" on page 19142	"show class-of-service traffic-control-profile" on page 19145
"show class-of-service translation-table" on page 19151	"show cli" on page 19157	"show cli authorization" on page 19159
"show cli directory" on page 19164	"show cli history" on page 19166	"show configuration" on page 19167

"show configuration chassis cluster traceoptions" on page 19171	"show configuration services analytics sensor ddos" on page 19173	"show configuration smtp " on page 19174
"show connections" on page 19176	"show cos halp flow queue-stats" on page 19181	"show database-replication statistics" on page 19183
"show database-replication summary" on page 19185	"show dcbx" on page 19189	"show dcbx neighbors" on page 19191
"show ddos-protection protocols" on page 19222	"show ddos-protection protocols culprit-flows" on page 19239	"show ddos-protection protocols flow-detection" on page 19249
"show ddos-protection protocols isis" on page 19254	"show ddos-protection protocols parameters" on page 19265	"show ddos-protection protocols statistics" on page 19276
"show ddos-protection protocols violations" on page 19292	"show ddos-protection statistics" on page 19295	"show ddos-protection version" on page 19298
"show dhcp client binding" on page 19300	"show dhcp client statistics" on page 19305	"show dhcp relay active-leasequery" on page 19309
"show dhcp relay binding" on page 19319	"show dhcp relay binding (Security)" on page 19316	"show dhcp relay lockout-entries" on page 19328
"show dhcp relay statistics" on page 19334	"show dhcp relay statistics (Security)" on page 19331	"show dhcp server active-leasequery" on page 19362
"show dhcp server active-leasequery statistics" on page 19364	"show dhcp server binding" on page 19367	"show dhcp server lockout-entries" on page 19377
"show dhcp server statistics" on page 19381	"show dhcp snooping binding" on page 19387	"show dhcp snooping statistics" on page 19390
"show dhcp-security arp inspection statistics" on page 19342	"show dhcp-security binding" on page 19344	"show dhcp-security binding ip-source-guard" on page 19348

"show dhcp-security ipv6 binding" on page 19351	"show dhcp-security ipv6 statistics" on page 19355	"show dhcp-security neighbor-discovery-inspection statistics" on page 19359
"show dhcpv6 client binding" on page 19392	"show dhcpv6 client statistics" on page 19395	"show dhcpv6 relay active-leasequery" on page 19398
"show dhcpv6 relay binding" on page 19404	"show dhcpv6 relay lockout-entries" on page 19417	"show dhcpv6 relay statistics" on page 19420
"show dhcpv6 server active-leasequery" on page 19426	"show dhcpv6 server active-leasequery statistics" on page 19428	"show dhcpv6 server binding" on page 19431
"show dhcpv6 server binding (View)" on page 19440	"show dhcpv6 server lockout-entries" on page 19446	"show dhcpv6 server statistics" on page 19450
"show dhcpv6 server statistics (View)" on page 19456	"show dhcpv6 snooping binding" on page 19460	"show dhcpv6 snooping statistics" on page 19463
"show diagnostics tdr" on page 19465	"show diameter" on page 19472	"show diameter function" on page 19482
"show diameter function statistics" on page 19488	"show diameter instance" on page 19493	"show diameter network-element" on page 19495
"show diameter network-element map" on page 19500	"show diameter peer" on page 19504	"show diameter peer map" on page 19510
"show diameter peer statistics" on page 19514	"show diameter route" on page 19519	"show | display" on page 19522
"show | display commit-scripts" on page 19527	"show | display inheritance " on page 19529	"show | display json" on page 19532
"show | display omit" on page 19535	"show | display rfc5952" on page 19537	"show | display set" on page 19538

"show | display set relative" on page 19540	"show | display translation-scripts" on page 19542	"show | display xml" on page 19546
"show dot1x" on page 19548	"show dot1x accounting attribute" on page 19557	"show dot1x authentication-bypassed-users" on page 19561
"show dot1x authentication-failed-users" on page 19562	"show dot1x firewall" on page 19564	"show dot1x interface" on page 19567
"show dot1x static-mac-address" on page 19577	"show dot1x statistics" on page 19580	"show dual-phase-bootup-status" on page 19583
"show dvmrp interfaces" on page 19584	"show dvmrp neighbors" on page 19587	"show dvmrp prefix" on page 19589
"show dvmrp prunes" on page 19592	"show dynamic-profile session" on page 19595	"show dynamic-tunnels database" on page 19601
"show dynamic-tunnels pfe-tunnel-localization" on page 19609	"show edge-virtual-bridging" on page 19616	"show ephemeral-configuration" on page 19620
"show ethernet-switching evpn nd-statistics" on page 19623	"show ethernet-switching evpn nd-table" on page 19626	"show ethernet-switching flood" on page 19630
"show ethernet-switching interface" on page 19636	"show ethernet-switching interfaces" on page 19641	"show ethernet-switching layer2-protocol-tunneling interface" on page 19652
"show ethernet-switching layer2-protocol-tunneling statistics" on page 19654	"show ethernet-switching layer2-protocol-tunneling vlan" on page 19658	"show ethernet-switching mac-ip-table" on page 19660
"show ethernet-switching mac-learning-log" on page 19670	"show ethernet-switching mgrp-policy" on page 19676	"show ethernet-switching redundancy-groups" on page 19682
"show ethernet-switching statistics" on page 19688	"show ethernet-switching statistics aging" on page 19692	"show ethernet-switching statistics mac-learning" on page 19694

"show ethernet-switching table" on page 19700	"show ethernet-switching vxlan-tunnel-end-point esi" on page 19724	"show ethernet-switching vxlan-tunnel-end-point remote" on page 19728
"show ethernet-switching vxlan-tunnel-end-point source" on page 19730	"show ethernet-switching vxlan-tunnel-end-point svlbnh" on page 19733	"show evpn arp-table" on page 19736
"show evpn database" on page 19740	"show evpn flood " on page 19748	"show evpn igmp-snooping database" on page 19752
"show evpn igmp-snooping proxy" on page 19756	"show evpn instance" on page 19761	"show evpn ip-prefix-database " on page 19779
"show evpn l3-context" on page 19790	"show evpn mac-ip-table" on page 19794	"show evpn mac-table" on page 19798
"show evpn mld-snooping database" on page 19800	"show evpn multicast-snooping assisted-replication multihomed-peers" on page 19804	"show evpn multicast-snooping assisted-replication next-hops" on page 19807
"show evpn multicast-snooping assisted-replication replicators" on page 19811	"show evpn multicast-snooping status" on page 19815	"show evpn nd-table" on page 19818
"show evpn oism" on page 19822	"show evpn oism peg-df-status" on page 19826	"show evpn oism spmsi-ad" on page 19829
"show evpn p2mp" on page 19833	"show evpn peer-gateway-macs" on page 19836	"show evpn prefix" on page 19838
"show evpn vpws-instance" on page 19840	"show extensible-subscriber-services counters" on page 19849	"show extensible-subscriber-services debug-information" on page 19851
"show extensible-subscriber-services dictionary" on page 19853	"show extensible-subscriber-services dictionary attributes" on page 19859	"show extensible-subscriber-services dictionary services" on page 19863

"show extensible-subscriber-services service" on page 19867	"show extensible-subscriber-services sessions" on page 19869	"show extension-service request-response clients" on page 19871
"show extension-service request-response servers" on page 19874	"show extension-service status" on page 19876	"show fabric administration inventory" on page 19878
"show fabric administration inventory director-group status" on page 19885	"show fabric administration inventory infrastructure" on page 19893	"show fabric administration inventory interconnect-devices" on page 19898
"show fabric administration inventory node-devices" on page 19900	"show fabric administration inventory node-groups" on page 19903	"show fabric administration system mac-pool" on page 19906
"show fabric inventory" on page 19908	"show fabric session-host" on page 19912	"show fib-streaming state" on page 19927
"show fib-local-accounting ip" on page 19914	"show fib-streaming" on page 19915	"show fib-streaming route-tables" on page 19921
"show fib-streaming routes" on page 19924	"show fibre-channel fabric" on page 19929	"show fibre-channel fc2 sessions" on page 19932
"show fibre-channel fc2 statistics" on page 19936	"show fibre-channel fip" on page 19938	"show fibre-channel fip enode" on page 19946
"show fibre-channel fip fabric" on page 19951	"show fibre-channel fip fcf" on page 19956	"show fibre-channel fip interface" on page 19961
"show fibre-channel fip statistics" on page 19966	"show fibre-channel flogi fport" on page 19971	"show fibre-channel flogi nport" on page 19973
"show fibre-channel flogi statistics" on page 19977	"show fibre-channel interfaces" on page 19980	"show fibre-channel next-hops" on page 19985
"show fibre-channel proxy fabric-state" on page 19987	"show fibre-channel proxy login-table" on page 19993	"show fibre-channel proxy np-port" on page 19997

"show fibre-channel proxy statistics" on page 20003	"show fibre-channel routes" on page 20006	"show fip snooping" on page 20008
"show fip snooping enode" on page 20015	"show fip snooping fcf" on page 20021	"show fip snooping interface" on page 20024
"show fip snooping statistics" on page 20029	"show fip snooping vlan" on page 20034	"show fip vlan-discovery" on page 20040
"show firewall" on page 20048	"show firewall filter version" on page 20060	"show firewall log" on page 20062
"show firewall policer" on page 20066	"show firewall prefix-action-stats" on page 20068	"show firewall (QFX)" on page 20043
"show firewall templates-in-use" on page 20071	"show firewall (View)" on page 20073	"show flexible-tunnels profiles" on page 20077
"show flight-recorder status" on page 20083	"show forwarding-options enhanced-hash-key" on page 20086	"show forwarding-options hash-key" on page 20095
"show forwarding-options hyper-mode" on page 20100	"show forwarding-options load-balance" on page 20102	"show forwarding-options load-balance ecmp|trunk" on page 20107
"show forwarding-options next-hop-group" on page 20111	"show forwarding-options port-mirroring" on page 20116	"show groups junos-defaults" on page 20120
"show gtp tunnels" on page 20122	"show hfrr profiles" on page 20128	"show host" on page 20131
"show iccp" on page 20133	"show igmp group" on page 20135	"show igmp interface" on page 20142
"show igmp snooping data-forwarding" on page 20148	"show igmp snooping evpn status" on page 20153	"show igmp snooping interface" on page 20157
"show igmp snooping membership" on page 20165	"show igmp snooping options" on page 20179	"show igmp snooping statistics" on page 20187

"show igmp statistics" on page 20203	"show igmp-snooping membership" on page 20173	"show igmp-snooping route" on page 20181
"show igmp-snooping route (View)" on page 20185	"show igmp-snooping statistics" on page 20196	"show igmp-snooping vlans" on page 20198
"show ike security-associations" on page 20210	"show ilmi statistics" on page 20217	"show ingress-replication mvpn" on page 20221
"show interfaces " on page 20227	"show interfaces (Abstracted Fabric)" on page 20371	"show interfaces (Adaptive Services)" on page 20401
"show interfaces (Aggregated Ethernet)" on page 20411	"show interfaces anchor-group (Aggregated Packet Forwarding Engine)" on page 20437	"show interfaces (ATM)" on page 20442
"show interfaces demux0 (Demux Interfaces)" on page 20497	"show interfaces destination-class" on page 20517	"show interfaces detail" on page 20521
"show interfaces diagnostics optics" on page 20526	"show interfaces diagnostics optics" on page 20535	"show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)" on page 20580
"show interfaces diagnostics optics (Security)" on page 20569	"show interfaces diagnostics optics-applications" on page 20576	"show interfaces diagnostics optics-profile" on page 20645
"show interfaces (Discard) " on page 20651	"show interfaces (Dynamic Flow Capture)" on page 20660	"show interfaces (Encryption)" on page 20667
"show interfaces extensive" on page 20676	"show interfaces extensive satellite-device" on page 20747	"show interfaces fabric" on page 20774
"show interfaces (far-end-interval)" on page 20803	"show interfaces (Fast Ethernet)" on page 20805	"show interfaces filters" on page 20832

"show interfaces (Flow Collector)" on page 20835	"show interfaces (Flow Monitoring)" on page 20844	"show interfaces flow-statistics" on page 20852
"show interfaces forwarding-class-counters" on page 20859	"show interfaces fti" on page 20866	"show interfaces ge" on page 20881
"show interfaces (GRE)" on page 20900	"show interfaces interface-set (Ethernet Interface Set)" on page 20913	"show interfaces interface-set (PS Interface Set)" on page 20920
"show interfaces interface-set queue" on page 20922	"show interfaces interval" on page 20934	"show interfaces (IP-over-IP)" on page 20938
"show interfaces irb " on page 20945	"show interfaces l2-routing-instance" on page 20957	"show interfaces lb-stats af" on page 20959
"show interfaces (Link Services)" on page 20965	"show interfaces (Link Services IQ)" on page 20987	"show interfaces load-balancing (Aggregated Multiservices)" on page 21024
"show interfaces (Logical Tunnel)" on page 21029	"show interfaces (Loopback)" on page 21040	"show interfaces lsi (Label-Switched Interface)" on page 21051
"show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet)" on page 21056	"show interfaces mac-database" on page 21080	"show interfaces mc-ae" on page 21098
"show interfaces mc-ae (MX Series)" on page 21090	"show interfaces me0" on page 21101	"show interfaces media" on page 21113
"show interfaces (Multicast Tunnel)" on page 21117	"show interfaces (Multilink Services)" on page 21124	"show interfaces (NFX150)" on page 20223
"show interfaces (PIM)" on page 21137	"show interfaces policers" on page 21143	"show interfaces port-profile" on page 21146

"show interfaces (PPPoE)" on page 21154	"show interfaces prbs-stats" on page 21170	"show interfaces priority-group" on page 21172
"show interfaces ps0 (Pseudowire Subscriber Interfaces)" on page 21175	"show interfaces (PTX Series Packet Transport Routers)" on page 21182	"show interfaces queue" on page 21208
"show interfaces queue fabric" on page 21257	"show interfaces redundancy" on page 21288	"show interfaces (Redundant Adaptive Services)" on page 21292
"show interfaces (Redundant Link Services IQ)" on page 21295	"show interfaces routing" on page 21316	"show interfaces routing summary" on page 21328
"show interfaces routing-instance" on page 21325	"show interfaces satellite-device" on page 21332	"show interfaces (Serial)" on page 21339
"show interfaces smart-sfp-defects" on page 21359	"show interfaces smart-sfp-statistics" on page 21372	"show interfaces snmp-index" on page 21389
"show interfaces source-class" on page 21391	"show interfaces (SRX Series)" on page 21395	"show interfaces statistics" on page 21442
"show interfaces statistics fabric" on page 21461	"show interfaces statistics st0" on page 21488	"show interfaces summary" on page 21490
"show interfaces swfabx" on page 21492	"show interfaces targeting (Aggregated Ethernet for Subscriber Management)" on page 21494	"show interfaces terse" on page 21497
"show interfaces terse satellite-device" on page 21504	"show interfaces terse zone" on page 21508	"show interfaces transport pm" on page 21509
"show interfaces vcp" on page 21524	"show interfaces (View Aggregated Ethernet)" on page 21552	"show interfaces (Virtual Loopback Tunnel) " on page 21567

"show interfaces vlan" on page 21575	"show interfaces voq" on page 21593	"show interfaces xe" on page 21612
"show interfaces xle" on page 21639	"show ip-source-guard" on page 21666	"show ipsec certificates" on page 21668
"show ipsec redundancy" on page 21672	"show ipsec security-associations" on page 21675	"show ipv6 neighbors" on page 21680
"show ipv6 neighbors (SRX Series)" on page 21683	"show ipv6 router-advertisement" on page 21685	"show ipv6-source-guard" on page 21692
"show isis adjacency" on page 21694	"show isis adjacency holddown" on page 21702	"show isis authentication" on page 21706
"show isis backup coverage " on page 21709	"show isis backup label-switched-path" on page 21711	"show isis backup spf results" on page 21714
"show isis context-identifier" on page 21720	"show isis database" on page 21722	"show isis flex-algorithm" on page 21732
"show isis hostname" on page 21736	"show isis interface" on page 21738	"show isis interface-group" on page 21746
"show isis layer2-map" on page 21748	"show isis overview" on page 21751	"show isis purge log" on page 21759
"show isis route" on page 21762	"show isis route download priority" on page 21768	"show isis spf" on page 21772
"show isis spring interface traffic-statistics" on page 21779	"show isis spring sensor info" on page 21780	"show isis statistics" on page 21783
"show l2-learning global-information" on page 21787	"show l2-learning global-mac-count" on page 21790	"show l2-learning instance" on page 21791

"show l2-learning interface" on page 21794	"show l2-learning mac-move-buffer" on page 21796	"show l2-learning redundancy-groups" on page 21798
"show l2circuit connections" on page 21805	"show l2pt interface" on page 21816	"show l2vpn connections" on page 21819
"show lacp interfaces" on page 21829	"show lacp interfaces (View)" on page 21836	"show lacp statistics interfaces (View)" on page 21843
"show lacp timeouts" on page 21845	"show ldp database" on page 21848	"show ldp database-label-requests" on page 21861
"show ldp fec-filters" on page 21864	"show ldp interface" on page 21866	"show ldp neighbor" on page 21868
"show ldp overview" on page 21871	"show ldp p2mp fec" on page 21879	"show ldp p2mp path" on page 21881
"show ldp p2mp tunnel" on page 21884	"show ldp path" on page 21886	"show ldp rib-groups" on page 21888
"show ldp route" on page 21890	"show ldp session" on page 21901	"show ldp statistics" on page 21912
"show ldp traffic-statistics" on page 21918	"show link-management" on page 21923	"show link-management peer" on page 21928
"show link-management routing" on page 21931	"show link-management statistics" on page 21935	"show link-management te-link" on page 21939
"show lldp" on page 21942	"show lldp local-information" on page 21952	"show lldp neighbors" on page 21956
"show lldp neighbors-vlan-name-tlv-list interface" on page 21973	"show lldp remote-global-statistics" on page 21974	"show lldp statistics" on page 21977
"show log" on page 21980	"show loop-detect enhanced interface" on page 21987	"show mac-rewrite interface" on page 21993

"show mac-vrf forwarding flood" on page 21996	"show mac-vrf forwarding flood-group" on page 22000	"show mac-vrf forwarding global-information" on page 22005
"show mac-vrf forwarding global-mac-count" on page 22008	"show mac-vrf forwarding global-mac-ip-count" on page 22010	"show mac-vrf forwarding instance" on page 22011
"show mac-vrf forwarding instance-mapping" on page 22018	"show mac-vrf forwarding interface" on page 22020	"show mac-vrf forwarding mac-ip-table" on page 22025
"show mac-vrf forwarding mac-table" on page 22031	"show mac-vrf forwarding mgrp-policy" on page 22035	"show mac-vrf forwarding statistics" on page 22042
"show mac-vrf forwarding vlans" on page 22046	"show mac-vrf forwarding vxlan-tunnel-end-point esi" on page 22050	"show mac-vrf forwarding vxlan-tunnel-end-point remote" on page 22054
"show mac-vrf forwarding vxlan-tunnel-end-point svlbnh" on page 22063	"show mld group" on page 22066	"show mld interface" on page 22074
"show mld snooping evpn database" on page 22081	"show mld snooping interface" on page 22085	"show mld snooping membership" on page 22090
"show mld snooping statistics" on page 22098	"show mld statistics" on page 22104	"show mld-snooping route" on page 22094
"show mld-snooping vlans" on page 22100	"show modem wireless firmware" on page 22111	"show modem wireless network" on page 22115
"show modem wireless profiles" on page 22120	"show monitor security flow" on page 22122	"show mpls abstract-hop-membership" on page 22125
"show mpls admin-groups" on page 22128	"show mpls association" on page 22130	"show mpls call-admission-control" on page 22132
"show mpls container-lsp" on page 22136	"show mpls context-identifer" on page 22147	"show mpls correlation label" on page 22150

"show mpls correlation nexthop-id" on page 22151	"show mpls cspf" on page 22153	"show mpls diffserv-te" on page 22156
"show mpls egress-protection" on page 22159	"show mpls interface" on page 22162	"show mpls label usage" on page 22165
"show mpls label usage label-range" on page 22169	"show mpls lsp" on page 22172	"show mpls lsp abstract-computation" on page 22203
"show mpls lsp autobandwidth" on page 22206	"show mpls path" on page 22210	"show mpls srlg" on page 22212
"show mpls static-lsp" on page 22214	"show mpls tunnel manager statistics" on page 22219	"show msdp" on page 22222
"show msdp source" on page 22225	"show msdp source-active" on page 22228	"show msdp statistics" on page 22232
"show multi-chassis mc-lag configuration-consistency" on page 22237	"show multi-chassis mc-lag configuration-consistency global-config" on page 22243	"show multi-chassis mc-lag configuration-consistency icl-config" on page 22246
"show multi-chassis mc-lag configuration-consistency list-of-parameters" on page 22249	"show multi-chassis mc-lag configuration-consistency mcae-config" on page 22262	"show multi-chassis mc-lag configuration-consistency vlan-config" on page 22267
"show multi-chassis mc-lag configuration-consistency vrrp-config" on page 22270	"show multicast backup-pe-groups" on page 22273	"show multicast ecid-mapping satellite" on page 22276
"show multicast flow-map" on page 22280	"show multicast forwarding-cache statistics" on page 22283	"show multicast interface" on page 22286
"show multicast mrinfo" on page 22290	"show multicast next-hops" on page 22292	"show multicast next-hops satellite" on page 22298

"show multicast pim-to-igmp-proxy" on page 22301	"show multicast pim-to-mld-proxy" on page 22304	"show multicast route" on page 22306
"show multicast rpf" on page 22323	"show multicast scope" on page 22328	"show multicast sessions" on page 22331
"show multicast snooping next-hops" on page 22335	"show multicast snooping next-hops satellite" on page 22338	"show multicast snooping route" on page 22344
"show multicast snooping route satellite" on page 22350	"show multicast statistics" on page 22356	"show multicast statistics satellite " on page 22363
"show multicast summary satellite" on page 22365	"show multicast usage" on page 22366	"show mvpn c-multicast" on page 22370
"show mvpn instance" on page 22375	"show mvpn neighbor" on page 22381	"show mvpn suppressed" on page 22388
"show mvrp" on page 22390	"show mvrp applicant-state" on page 22394	"show mvrp dynamic-vlan-memberships" on page 22398
"show mvrp interface" on page 22401	"show mvrp registration-state" on page 22403	"show mvrp statistics" on page 22406
"show neighbor-discovery-inspection statistics" on page 22413	"show network access nas-5g " on page 22509	"show network-access aaa accounting" on page 22415
"show network-access aaa radius-servers" on page 22417	"show network-access aaa saml assertion-cache" on page 22560	"show network-access aaa statistics" on page 22428
"show network-access aaa statistics accounting" on page 22447	"show network-access aaa statistics authentication" on page 22449	"show network-access aaa statistics dynamic-requests" on page 22453
"show network-access aaa statistics pending-accounting-stops" on page 22455	"show network-access aaa statistics preauthentication" on page 22456	"show network-access aaa statistics re-authentication" on page 22459

"show network-access aaa subscribers" on page 22461	"show network-access aaa subscribers session-id" on page 22468	"show network-access aaa terminate-code" on page 22479
"show network-access address-assignment pool" on page 22486	"show network-access address-assignment pool (View)" on page 22488	"show network-access domain-map" on page 22490
"show network-access gx-plus" on page 22492	"show network-access nasreq statistics" on page 22514	"show network-access ocs" on page 22520
"show network-access pcrf" on page 22523	"show network-access radsec local-certificate" on page 22528	"show network-access radsec state" on page 22531
"show network-access radsec statistics" on page 22535	"show network-access requests pending" on page 22538	"show network-access requests statistics" on page 22542
"show network-access s6a" on page 22544	"show network-access securid-node-secret-file" on page 22547	"show network-agent statistics" on page 22549
"show nlsd mode" on page 22562	"show node reboot" on page 22564	"show node statistics" on page 22565
"show nonstop-routing" on page 22571	"show nsagentd" on page 22575	"show ntp associations" on page 22577
"show ntp status" on page 22581	"show oam ethernet connectivity-fault-management delay-statistics" on page 22585	"show oam ethernet connectivity-fault-management forwarding-state" on page 22590
"show oam ethernet connectivity-fault-management interfaces" on page 22596	"show oam ethernet connectivity-fault-management linktrace path-database" on page 22611	"show oam ethernet connectivity-fault-management loss-statistics" on page 22615
"show oam ethernet connectivity-fault-management mep-database" on page 22620	"show oam ethernet connectivity-fault-management mep-statistics" on page 22637	"show oam ethernet connectivity-fault-management mip" on page 22654

"show oam ethernet connectivity-fault-management path-database" on page 22656	"show oam ethernet connectivity-fault-management policer" on page 22660	"show oam ethernet connectivity-fault-management sla-iterator-statistics" on page 22664
"show oam ethernet connectivity-fault-management synthetic-loss-statistics" on page 22677	"show oam ethernet evc" on page 22681	"show oam ethernet fnp interface" on page 22684
"show oam ethernet fnp messages" on page 22686	"show oam ethernet fnp status" on page 22688	"show oam ethernet link-fault-management" on page 22691
"show oam ethernet lmi" on page 22701	"show oam ethernet lmi statistics" on page 22704	"show oam fabric flow-specification" on page 22707
"show oam fabric interfaces" on page 22711	"show openflow capability" on page 22713	"show openflow controller" on page 22722
"show openflow filters" on page 22726	"show openflow flows" on page 22730	"show openflow groups" on page 22736
"show openflow interfaces" on page 22740	"show openflow statistics flows" on page 22745	"show openflow statistics groups" on page 22750
"show openflow statistics interfaces" on page 22753	"show openflow statistics packet" on page 22758	"show openflow statistics queue" on page 22760
"show openflow statistics summary" on page 22764	"show openflow statistics tables" on page 22766	"show openflow summary" on page 22771
"show openflow switch" on page 22773	"show ospf context-identifier" on page 22776	"show ospf database" on page 22780
"show ospf flex-algorithm" on page 22791	"show (ospf | ospf3) backup coverage" on page 22800	"show (ospf | ospf3) backup lsp" on page 22804
"show (ospf | ospf3) backup neighbor" on page 22807	"show (ospf | ospf3) backup spf" on page 22809	"show (ospf | ospf3) bgp-orr" on page 22813

"show (ospf | ospf3) interface" on page 22817	"show (ospf | ospf3) io-statistics" on page 22829	"show (ospf | ospf3) log" on page 22831
"show (ospf | ospf3) neighbor" on page 22835	"show (ospf | ospf3) overview" on page 22843	"show (ospf | ospf3) route" on page 22851
"show (ospf | ospf3) statistics" on page 22859	"show ospf spring prefix sid map" on page 22864	"show ospf spring sid database" on page 22866
"show ospf3 database" on page 22869	"show ovsdb commit failures" on page 22881	"show ovsdb controller " on page 22884
"show ovsdb interface" on page 22887	"show ovsdb logical-switch" on page 22890	"show ovsdb mac" on page 22893
"show ovsdb statistics interface" on page 22899	"show ovsdb tunnels" on page 22901	"show ovsdb virtual-tunnel-end-point" on page 22905
"show passive-monitoring error" on page 22908	"show passive-monitoring flow" on page 22911	"show passive-monitoring memory" on page 22915
"show passive-monitoring status" on page 22917	"show passive-monitoring usage" on page 22919	"show path-computation-client active-pce" on page 22922
"show path-computation-client lsp" on page 22927	"show path-computation-client statistics" on page 22936	"show path-computation-client status" on page 22944
"show performance-monitoring mpls lsp" on page 22947	"show pfe cfeb" on page 22955	"show pfe data" on page 22961
"show pfe feb" on page 22974	"show pfe filter hw summary" on page 22983	"show pfe fpc" on page 22986
"show pfe next-hop" on page 23013	"show pfe route" on page 23019	"show pfe sFlow egress-mcast-stats" on page 23033

"show pfe ssb" on page 23036	"show pfe statistics bridge" on page 23045	"show pfe terse" on page 23051
"show pfe version" on page 23054	"show pim bidirectional df-election" on page 23056	"show pim bidirectional df-election interface " on page 23060
"show pim bootstrap" on page 23064	"show pim interfaces" on page 23066	"show pim join" on page 23070
"show pim mdt" on page 23093	"show pim mdt data-mdt-joins" on page 23099	"show pim mdt data-mdt-limit" on page 23102
"show pim mvpn" on page 23104	"show pim neighbors" on page 23106	"show pim rps" on page 23113
"show pim snooping interfaces" on page 23124	"show pim snooping join" on page 23128	"show pim snooping neighbors" on page 23133
"show pim snooping statistics" on page 23141	"show pim source" on page 23147	"show pim statistics" on page 23152
"show poe controller" on page 23171	"show poe controller (View)" on page 23175	"show poe interface" on page 23177
"show poe notification-control" on page 23187	"show poe telemetries" on page 23189	"show policer" on page 23194
"show policer (EX)" on page 23192	"show policy" on page 23199	"show policy conditions" on page 23203
"show policy damping" on page 23206	"show ppp address-pool" on page 23209	"show ppp interface" on page 23211
"show ppp statistics" on page 23229	"show ppp summary" on page 23239	"show pppoe interfaces" on page 23246

"show pppoe interfaces (Security)" on page 23241	"show pppoe lockout" on page 23252	"show pppoe lockout atm-identifier" on page 23256
"show pppoe lockout vlan-identifier" on page 23260	"show pppoe service-name-tables" on page 23265	"show pppoe sessions" on page 23269
"show pppoe statistics" on page 23271	"show pppoe underlying-interfaces" on page 23275	"show pppoe version" on page 23285
"show programmable-rpd clients" on page 23287	"show programmable-rpd clients route-summary" on page 23289	"show programmable-rpd statistics" on page 23292
"show protection-group ethernet-ring aps" on page 23294	"show protection-group ethernet-ring configuration" on page 23300	"show protection-group ethernet-ring data-channel" on page 23309
"show protection-group ethernet-ring flush-info" on page 23312	"show protection-group ethernet-ring interface" on page 23315	"show protection-group ethernet-ring node-state" on page 23320
"show protection-group ethernet-ring statistics" on page 23328	"show protection-group ethernet-ring vlan" on page 23335	"show ptp aggregated-ethernet interfaces" on page 23342
"show ptp all master clock" on page 23344	"show ptp clock" on page 23348	"show ptp global-information" on page 23354
"show ptp hybrid" on page 23369	"show ptp lock-status" on page 23372	"show ptp master" on page 23377
"show ptp passive-port-monitor-status" on page 23383	"show ptp path-trace detail" on page 23389	"show ptp phy-timestamping-interfaces" on page 23390
"show ptp port" on page 23392	"show ptp quality-level-mapping" on page 23400	"show ptp slave" on page 23401
"show ptp stateful detail" on page 23406	"show ptp statistics" on page 23408	"show redundant-power-system led" on page 23414

"show redundant-power-system multi-backup" on page 23418	"show redundant-power-system network" on page 23419	"show redundant-power-system power-supply" on page 23420
"show redundant-power-system status" on page 23423	"show redundant-power-system upgrade" on page 23426	"show redundant-power-system version" on page 23427
"show redundant-trunk-group" on page 23430	"show remote-device-management service-devices" on page 23432	"show remote-device-management services" on page 23441
"show remote-device-management statistics" on page 23445	"show remote-device-management subscribers" on page 23451	"show remote-device-management summary" on page 23455
"show resource-manager" on page 23458	"show resource-manager cpu" on page 23460	"show resource-manager memory" on page 23461
"show rift database" on page 23463	"show rift flood-reduction" on page 23467	"show rift interface" on page 23469
"show rift node" on page 23476	"show rift path-computation" on page 23480	"show rift routes" on page 23482
"show rift tie" on page 23486	"show rift topology" on page 23488	"show rift versions" on page 23491
"show rip general-statistics" on page 23493	"show rip neighbor" on page 23496	"show rip statistics" on page 23499
"show ripng general-statistics" on page 23504	"show ripng neighbor" on page 23506	"show ripng statistics" on page 23508
"show route" on page 23512	"show route active-path" on page 23526	"show route advertising-protocol" on page 23533
"show route all" on page 23542	"show route aspath-regex" on page 23545	"show route best" on page 23548
"show route brief" on page 23553	"show route ccc" on page 23555	"show route community" on page 23557

"show route community-name" on page 23560	"show route cumulative" on page 23562	"show route damping" on page 23564
"show route detail" on page 23573	"show route exact" on page 23608	"show route export" on page 23612
"show route export vrf-target" on page 23615	"show route extensive" on page 23619	"show route extensive <route> expanded-nh" on page 23643
"show route flow validation" on page 23649	"show route forwarding-table" on page 23652	"show route forwarding-table family fibre-channel" on page 23673
"show route forwarding-table interface-name" on page 23676	"show route hidden" on page 23680	"show route inactive-path" on page 23685
"show route inactive-prefix" on page 23690	"show route instance" on page 23692	"show route label" on page 23698
"show route label-switched-path" on page 23705	"show route localization" on page 23707	"show route martians" on page 23709
"show route match-prefix" on page 23713	"show route next-hop" on page 23715	"show route no-community" on page 23719
"show route output" on page 23723	"show route pfe-install-error" on page 23727	"show route protocol" on page 23730
"show route range" on page 23740	"show route receive-protocol" on page 23746	"show route resolution" on page 23753
"show route rpm-tracking" on page 23759	"show route sla-tracking" on page 23764	"show route snooping" on page 23767
"show route source-gateway" on page 23771	"show route summary " on page 23779	"show route table" on page 23789
"show route tenant" on page 23821	"show route terse" on page 23824	"show route vpn-localization" on page 23829

"show routing nhdb pfe-install-error" on page 23835	"show rsvp interface" on page 23837	"show rsvp neighbor" on page 23847
"show rsvp pop-and-forward" on page 23854	"show rsvp route-session-id " on page 23857	"show rsvp session" on page 23860
"show rsvp statistics" on page 23878	"show rsvp version" on page 23887	"show sap listen " on page 23892
"show security advance policy based routing multipath policy statistics" on page 23898	"show security advance-policy-based-routing detail" on page 23894	"show security advance-policy-based-routing multipath rule" on page 23901
"show security advance-policy-based-routing profile" on page 23913	"show security advance-policy-based-routing sla active-probe-statistics" on page 23915	"show security advance-policy-based-routing sla profile (Application Name)" on page 23919
"show security advance-policy-based-routing sla profile (Next-Hop)" on page 23921	"show security advance-policy-based-routing sla profile (Status)" on page 23928	"show security advance-policy-based-routing sla statistics" on page 23931
"show security advance-policy-based-routing sla status" on page 23934	" show security advance-policy-based-routing sla version " on page 23935	"show security advance-policy-based-routing statistics" on page 23936
"show security advance-policy-based-routing status" on page 23943	"show security advanced-policy-based-routing policy-name" on page 23945	"show security alarms" on page 23952
"show security alg h323 counters" on page 23958	"show security alg ike-esp-nat summary" on page 23963	"show security alg msrpc" on page 23964
"show security alg sccp calls" on page 23966	"show security alg sccp counters" on page 23969	"show security alg sip calls" on page 23973
"show security alg sip counters" on page 23975	"show security alg sip rate" on page 23996	"show security alg status" on page 23999

"show security alg status logical-system" on page 24007	"show security application-firewall rule-set" on page 24012	"show security application-firewall rule-set logical-system" on page 24017
"show security application-tracking counters" on page 24021	"show security casb policies" on page 24023	"show security datapath-debug capture" on page 24025
"show security datapath-debug counter" on page 24027	"show security dns-cache " on page 24028	"show security dynamic-address" on page 24031
"show security dynamic-policies" on page 24045	"show security dynamic-vpn users" on page 24053	"show security dynamic-vpn users terse" on page 24055
"show security firewall-authentication history" on page 24058	"show security firewall-authentication history address" on page 24061	"show security firewall-authentication history identifier" on page 24065
"show security firewall-authentication jims" on page 24069	"show security firewall-authentication users" on page 24071	"show security firewall-authentication users address" on page 24075
"show security firewall-authentication users auth-type" on page 24080	"show security firewall-authentication users identifier" on page 24082	"show security flow cp-session" on page 24086
"show security flow cp-session destination-port" on page 24091	"show security flow cp-session destination-prefix" on page 24095	"show security flow cp-session family" on page 24098
"show security flow cp-session protocol" on page 24102	"show security flow cp-session source-port" on page 24105	"show security flow cp-session source-prefix" on page 24109
"show security flow gate" on page 24112	"show security flow gate brief node " on page 24118	"show security flow gate destination-port" on page 24127
"show security flow gate destination-prefix" on page 24131	"show security flow gate protocol" on page 24135	"show security flow gate summary node " on page 24140

"show security flow ip-action" on page 24147	"show security flow pmi statistics" on page 24159	"show security flow session" on page 24162
"show security flow session advanced-anti-malware" on page 24177	"show security flow session application" on page 24184	"show security flow session application-firewall" on page 24190
"show security flow session brief node " on page 24199	"show security flow session destination-port" on page 24205	"show security flow session destination-prefix" on page 24211
"show security flow session extensive node " on page 24217	"show security flow session family" on page 24227	"show security flow session idp family" on page 24235
"show security flow session idp summary" on page 24237	"show security flow session interface" on page 24240	"show security flow session nat" on page 24246
"show security flow session plugins" on page 24251	"show security flow session policy-id" on page 24254	"show security flow session pretty" on page 24259
"show security flow session protocol" on page 24263	"show security flow session resource-manager" on page 24270	"show security flow session services-offload" on page 24276
"show security flow session session-identifier" on page 24283	"show security flow session session-state" on page 24291	"show security flow session source-port" on page 24295
"show security flow session source-prefix" on page 24300	"show security flow session ssl" on page 24306	"show security flow session summary family" on page 24310
"show security flow session summary node " on page 24313	"show security flow session summary services-offload" on page 24322	"show security flow session tenant" on page 24327
"show security flow session tunnel" on page 24330	"show security flow statistics" on page 24341	"show security flow status" on page 24348
"show security flow tunnel inspection statistics" on page 24358	"show security forward-options resource-manager" on page 24360	"show security forwarding-options mirror-filter" on page 24363

"show security group-vpn member ike security-associations" on page 24366	"show security group-vpn member ipsec inactive-tunnels" on page 24372	"show security group-vpn member ipsec security-associations" on page 24375
"show security group-vpn member ipsec statistics" on page 24381	"show security group-vpn member kek security-associations" on page 24385	"show security group-vpn member policy" on page 24391
"show security group-vpn server ike security-associations" on page 24394	"show security group-vpn server ipsec security-associations" on page 24401	"show security group-vpn server kek security-associations" on page 24404
"show security group-vpn server registered-members" on page 24409	"show security group-vpn server server-cluster" on page 24412	"show security group-vpn server statistics" on page 24416
"show security gtp" on page 24419	"show security gtp counters" on page 24420	"show security gtp counters path-rate-limit" on page 24432
"show security gtp grouped-ie-profile" on page 24434	"show security gtp gsn statistics" on page 24436	"show security gtp handover-group" on page 24437
"show security gtp ie-set" on page 24439	"show security gtp ip-group" on page 24441	"show security gtp message-ie-profile-v1" on page 24443
"show security gtp message-ie-profile-v2" on page 24445	"show security gtp message-list" on page 24447	"show security gtp profile" on page 24448
"show security gtp rate-limit default" on page 24460	"show security gtp sctp association" on page 24461	"show security gtp sctp counters" on page 24465
"show security host-vpn security-associations" on page 24473	"show security host-vpn version" on page 24477	"show security idp active-policy" on page 24478
"show security idp attack attack-list" on page 24480	"show security idp attack attack-list policy" on page 24485	"show security idp attack context" on page 24490

"show security idp attack deprecated-list" on page 24491	"show security idp attack description" on page 24492	"show security idp attack detail" on page 24494
"show security idp attack group-list" on page 24498	"show security idp attack table" on page 24500	"show security idp attacks deprecated-attacks policy policy_name" on page 24502
"show security idp counters action" on page 24504	"show security idp counters application-identification" on page 24507	"show security idp counters dfa" on page 24514
"show security idp counters flow" on page 24517	"show security idp counters http-decoder" on page 24534	"show security idp counters ips" on page 24537
"show security idp counters log" on page 24545	"show security idp counters memory" on page 24552	"show security idp counters packet" on page 24556
"show security idp counters packet-log" on page 24567	"show security idp counters pdf-decoder" on page 24581	"show security idp counters policy-manager" on page 24585
"show security idp counters ssl-inspection" on page 24588	"show security idp counters tcp-reassembler" on page 24592	"show security idp jist-package-version" on page 24599
"show security idp logical system" on page 24600	"show security idp logical-system policy-association" on page 24602	"show security idp memory" on page 24603
"show security idp policies" on page 24605	"show security idp policy-commit-status" on page 24607	"show security idp policy-commit-status clear " on page 24610
"show security idp policy-templates-list" on page 24611	"show security idp predefined-attacks" on page 24613	"show security idp security-package-version" on page 24615
"show security idp ssl-inspection key" on page 24618	"show security idp ssl-inspection session-id-cache" on page 24621	"show security idp status" on page 24622
"show security idp status detail" on page 24625	"show security ike active-peer" on page 24630	"show security ike debug-status" on page 24638

"show security ike peers" on page 24640	"show security ike peers statistics" on page 24644	"show security ike pre-shared-key" on page 24647
"show security ike security-associations" on page 24648	"show security ike stats" on page 24671	"show security ike tunnel-map" on page 24681
"show security integrity" on page 24685	"show security internal-security-association" on page 24686	"show security ipsec control-plane-security-associations" on page 24688
"show security ipsec inactive-tunnels" on page 24691	"show security ipsec next-hop-tunnels" on page 24699	"show security ipsec security-associations" on page 24701
"show security ipsec statistics" on page 24744	"show security ipsec traffic-selector" on page 24753	"show security ipsec tunnel-distribution" on page 24755
"show security ipsec tunnel-events-statistics" on page 24761	"show security ipsec vpn vpnname ike idle time" on page 24763	"show security keychain" on page 24765
"show security log" on page 24769	"show security log file" on page 24773	"show security log profile" on page 24781
"show security log query" on page 24791	"show security log report" on page 24794	"show security log report in-detail" on page 24796
"show security log report in-interval" on page 24802	"show security log report summary" on page 24804	"show security log severity" on page 24807
"show security macsec connections" on page 24812	"show security macsec connections (MX Series)" on page 24815	"show security macsec connections (SRX)" on page 24809
"show security macsec statistics" on page 24819	"show security macsec statistics (SRX Series Firewalls)" on page 24825	"show security match-policies" on page 24831
"show security mka sessions" on page 24840	"show security mka sessions (MX Series)" on page 24844	"show security mka sessions (SRX Series Device" on page 24849

"show security mka sessions summary" on page 24852	"show security mka statistics" on page 24854	"show security mka statistics" on page 24857
"show security mka statistics (MX Series)" on page 24860	"show security monitoring" on page 24864	"show security monitoring fpc fpc-number" on page 24867
"show security monitoring performance session" on page 24871	"show security monitoring performance spu" on page 24873	"show security nat destination pool" on page 24877
"show security nat destination rule" on page 24881	"show security nat destination rule-application" on page 24886	"show security nat destination summary" on page 24889
"show security nat incoming-table" on page 24893	"show security nat interface-nat-ports" on page 24895	"show security nat resource-usage source-pool" on page 24900
"show security nat source deterministic" on page 24904	"show security nat source paired-address" on page 24907	"show security nat source persistent-nat-table" on page 24910
"show security nat source pool" on page 24914	"show security nat source port-block" on page 24927	"show security nat source rule" on page 24931
"show security nat source rule-application" on page 24937	"show security nat source summary" on page 24940	"show security nat static rule" on page 24948
"show security np-cache summary" on page 24955	"show security pki ca-certificate" on page 24957	"show security pki ca-certificate (View)" on page 24962
"show security pki certificate-request" on page 24968	"show security pki certificate-request (View)" on page 24972	"show security pki crl" on page 24975
"show security pki crl (View)" on page 24979	"show security pki local-certificate" on page 24982	"show security pki local-certificate" on page 24987
"show security pki local-certificate (View)" on page 24992	"show security pki node-local certificate-request" on page 24999	"show security pki node-local local-certificate" on page 25003

"show security pki statistics" on page 25009	"show security policies" on page 25028	"show security policies checksum" on page 25056
"show security policies file-serialization counters" on page 25059	"show security policies hit-count" on page 25061	"show security policies information" on page 25068
"show security policies lookup-intact-on-commit" on page 25077	"show security policies policy set" on page 25080	"show security policies unknown-source-identity" on page 25083
"show security policies zone-context" on page 25085	"show security policy-report" on page 25089	"show security re-distribution ipsec-vpn" on page 25093
"show security resource-manager group active" on page 25095	"show security resource-manager resource active" on page 25099	"show security resource-manager settings" on page 25103
"show security resource-manager summary" on page 25106	"show security screen ids-option" on page 25108	"show security screen statistics" on page 25116
"show security screen status" on page 25132	"show security screen white-list" on page 25133	"show security shadow-policies" on page 25135
"show security softwires" on page 25138	"show security softwires map-e confidentiality status" on page 25140	"show security softwires map-e domain" on page 25142
"show security ssh key-pair-identity" on page 25144	"show security tcp-encap connection" on page 25146	"show security tcp-encap statistics" on page 25149
"show security tpm status" on page 25150	"show security tunnel inspection" on page 25153	" show security user-identification device-provision authentication-source active-directory start 1 count 9 (match-string|prefix)" on page 25156

"show security user-identification device-provision authentication-source jims start 1 count 9 (match-string|prefix)" on page 25158	"show security user-identification local-authentication-table" on page 25159	"show security user-identification role-provision all" on page 25163
"show security user-identification role-provision authentication-source active-directory start 1 count 9 (match-string|prefix)" on page 25164	"show security user-identification role-provision authentication-source jims start 1 count 9 (match-string|prefix)" on page 25166	"show security user-identification source-identity-provision all" on page 25169
"show security user-identification user-provision all" on page 25171	"show security user-identification user-provision authentication-source active-directory start 1 count 9 (match-string|prefix)" on page 25172	" show security user-identification user-provision authentication-source jims start 1 count 9 (match-string|prefix)" on page 25174
"show security utm anti-spam statistics" on page 25176	"show security utm anti-spam status" on page 25182	"show security utm anti-virus statistics" on page 25183
"show security utm anti-virus status" on page 25191	"show security utm content-filtering statistics" on page 25195	"show security utm custom-objects url-feed status" on page 25207
"show security utm session" on page 25212	"show security utm status" on page 25213	"show security utm web-filtering category base-filter" on page 25215
"show security utm web-filtering category category" on page 25217	"show security utm web-filtering category status" on page 25219	"show security utm web-filtering statistics" on page 25220
"show security utm web-filtering status" on page 25228	"show security zones" on page 25232	"show security zones type" on page 25242
"show service l2tp session routing-instance" on page 25246	"show service user-identification authentication-source aruba-clearpass user-query counters" on page 25248	"show service user-identification authentication-source aruba-clearpass user-query status" on page 25250

"show service user-identification identity-management " on page 25252	"show services accounting" on page 25256	"show services accounting aggregation" on page 25260
"show services accounting aggregation template" on page 25266	"show services accounting aggregation template (View)" on page 25268	"show services accounting aggregation (View)" on page 25269
"show services accounting errors" on page 25270	"show services accounting flow" on page 25277	"show services accounting flow-detail" on page 25286
"show services accounting flow-detail (View)" on page 25293	"show services accounting memory" on page 25295	"show services accounting packet-size-distribution" on page 25297
"show services accounting status" on page 25300	"show services accounting usage" on page 25305	"show services advanced-anti-malware dynamic-filter status" on page 25308
"show services advanced-anti-malware machine-learning-scan" on page 25311	"show services advanced-anti-malware policy" on page 25313	"show services advanced-anti-malware profile " on page 25317
"show services advanced-anti-malware statistics" on page 25321	"show services advanced-anti-malware status" on page 25328	"show services agf fiveqi remarking" on page 25331
"show services agf nas" on page 25334	"show services agf ngap" on page 25338	"show services alg conversations" on page 25348
"show services alg statistics" on page 25356	"show services analytics agent" on page 25374	"show services anti-virus statistics" on page 25378
"show services application identification application obsolete applications" on page 25406	"show services application-aware-access-list flows" on page 25381	"show services application-aware-access-list statistics" on page 25386
"show services application-identification application" on page 25389	"show services application-identification application micro-applications" on page 25402	"show services application-identification application non-configurable" on page 25404

"show services application-identification application-system-cache" on page 25408	"show services application-identification application-system-cache (View)" on page 25413	"show services application-identification commit-status" on page 25419
"show services application-identification commit-status (Next Gen Services)" on page 25420	"show services application-identification counter" on page 25422	"show services application-identification counter (AppSecure)" on page 25426
"show services application-identification entries" on page 25435	"show services application-identification group" on page 25438	"show services application-identification group obsolete-groups" on page 25472
"show services application-identification packet-capture counters" on page 25443	"show services application-identification recent-appid-sigpack-versions" on page 25473	"show services application-identification recent-appid-sigpack-versions" on page 25473
"show services application-identification statistics application-groups" on page 25446	"show services application-identification statistics applications" on page 25451	"show services application-identification status" on page 25456
"show services application-identification version" on page 25469	"show services captive-portal-content-delivery" on page 25476	"show services cos statistics" on page 25483
"show services cos statistics (Next Gen Services)" on page 25487	"show services crtp" on page 25492	"show services crtp flows" on page 25496
"show services dns-filtering cache" on page 25498	"show services dynamic-flow-capture content-destination" on page 25502	"show services dynamic-flow-capture control-source" on page 25504
"show services dynamic-flow-capture statistics " on page 25508	"show services fixed-wireless-access statistics" on page 25512	"show services flow-collector file interface" on page 25515
"show services flow-collector input interface" on page 25519	"show services flow-collector interface" on page 25521	"show services ha detail" on page 25532

"show services ha statistics" on page 25535	"show services hcm pic-statistics" on page 25542	"show services hcm statistics" on page 25550
"show services icap-redirect server status" on page 25553	"show services icap-redirect statistic" on page 25554	"show services icap-redirect status" on page 25559
"show services ids" on page 25562	"show services inband-flow-telemetry" on page 25574	"show services inline ip-reassembly statistics" on page 25577
"show services inline nat pool" on page 25591	"show services inline nat statistics" on page 25594	"show services inline softwire statistics" on page 25598
"show services inline-monitoring feature-profile-mapping fpc-slot" on page 25586	"show services inline-monitoring statistics fpc-slot" on page 25588	"show services ip-monitoring status" on page 25604
"show services ipsec-vpn certificates" on page 25610	"show services ipsec-vpn ike security-associations" on page 25615	"show services ipsec-vpn ipsec security-associations" on page 25622
"show services ipsec-vpn ipsec statistics" on page 25632	"show services l2tp client" on page 25638	"show services l2tp destination" on page 25641
"show services l2tp destination lockout" on page 25647	"show services l2tp multilink" on page 25649	"show services l2tp radius" on page 25656
"show services l2tp session" on page 25661	"show services l2tp session-limit-group" on page 25674	"show services l2tp summary" on page 25676
"show services l2tp tunnel" on page 25685	"show services l2tp tunnel-group" on page 25694	"show services l2tp tunnel-switch destination" on page 25696
"show services l2tp tunnel-switch session" on page 25702	"show services l2tp tunnel-switch summary" on page 25709	"show services l2tp tunnel-switch tunnel" on page 25711
"show services l2tp user" on page 25718	"show services link-services cpu-usage" on page 25724	"show services lrf collector statistics" on page 25730

"show services lrf rule statistics" on page 25732	"show services lrf statistics" on page 25735	"show services lrf template" on page 25737
"show services mobile-edge anchor class-of-service" on page 25740	"show services mobile-edge peers" on page 25742	"show services mobile-edge resources call-admission-control statistics" on page 25747
"show services mobile-edge sessions" on page 25749	"show services mobile-edge summary" on page 25772	"show services monitoring rfc2544" on page 25775
"show services monitoring rfc2544" on page 25780	"show services monitoring rpm history-results" on page 25787	"show services monitoring rpm probe-results" on page 25794
"show services monitoring twamp client control-info" on page 25804	"show services monitoring twamp client history-results" on page 25807	"show services monitoring twamp client probe-results" on page 25814
"show services monitoring twamp client test-info" on page 25825	"show services monitoring twamp server control-info" on page 25828	"show services monitoring twamp server test-info" on page 25831
"show services nat destination pool" on page 25834	"show services nat destination rule" on page 25836	"show services nat destination summary" on page 25839
"show services nat deterministic-nat internal-host" on page 25842	"show services nat deterministic-nat nat-port-block" on page 25844	"show services nat ipv6-multicast-interfaces" on page 25847
"show services nat pool" on page 25850	"show services nat resource-usage source-pool" on page 25858	"show services nat source deterministic" on page 25861
"show services nat source mappings address-pooling-paired" on page 25864	"show services nat source mappings endpoint-independent" on page 25868	"show services nat source mappings pcp" on page 25871
"show services nat source mappings summary" on page 25873	"show services nat source pool" on page 25875	"show services nat source port-block" on page 25881
"show services nat source rule" on page 25884	"show services nat source rule-application" on page 25888	"show services nat source summary" on page 25891

"show services paa status" on page 25893	"show services pcef pic" on page 25899	"show services pcef subscribers" on page 25901
"show services pcp statistics" on page 25910	"show services policies" on page 25915	"show services policies detail" on page 25917
"show services policies hit-count" on page 25921	"show services policies interface" on page 25922	"show services policies service-set" on page 25924
"show services redundancy-group" on page 25925	"show services rpm active-servers" on page 25937	"show services rpm history-results" on page 25939
"show services rpm probe-results" on page 25944	"show services rpm rfc2544-benchmarking " on page 25960	"show services rpm rfc2544-benchmarking test-id" on page 25969
"show services rpm twamp client connection" on page 25995	"show services rpm twamp client history-results" on page 25997	"show services rpm twamp client probe-results" on page 26002
"show services rpm twamp client session" on page 26013	"show services rpm twamp server connection" on page 26016	"show services rpm twamp server session" on page 26018
"show services screen ids-option" on page 26021	"show services screen-statistics service-set" on page 26024	"show services security-intelligence" on page 26032
"show services security-intelligence category summary" on page 26034	"show services security-intelligence dns-statistics" on page 26037	"show services security-intelligence update status " on page 26041
"show services security-metadata-streaming" on page 26042	"show services service-redirect statistic" on page 26048	"show services service-set statistics ids drops" on page 26049
"show services service-sets cpu-usage" on page 26058	"show services service-sets memory-usage" on page 26061	"show services service-sets plug-ins" on page 26063
"show services service-sets statistic screen-drops (Next Gen Services)" on page 26065	"show services service-sets statistic screen-session-limit-counters (Next Gen Services)" on page 26073	"show services service-sets statistics ids session-limits counters" on page 26084

"show services service-sets statistics integrity-drops" on page 26095	"show services service-sets statistics jflow-log" on page 26101	"show services service-sets statistics packet-drops" on page 26111
"show services service-sets statistics syslog" on page 26114	"show services service-sets statistics tcp" on page 26123	"show services service-sets statistics tcp-mss" on page 26125
"show services service-sets summary" on page 26127	"show services sessions" on page 26129	"show services sessions (Aggregated Multiservices)" on page 26141
"show services sessions analysis" on page 26151	"show services sessions analysis (USF)" on page 26156	"show services sessions count" on page 26162
"show services sessions (Next Gen Services)" on page 26163	"show services sessions service-set" on page 26176	"show services sessions softwire" on page 26178
"show services sessions tcp-log " on page 26183	"show services sessions utilization" on page 26185	"show services soft-gre tunnel" on page 26186
"show services softwire " on page 26190	"show services softwire flows" on page 26192	"show services softwire statistics" on page 26197
"show services ssl certificate" on page 26209	"show services ssl droplogs" on page 26215	"show services ssl initiation counters" on page 26217
"show services ssl initiation profile " on page 26222	"show services ssl proxy certificate–cache statistics" on page 26231	"show services ssl proxy certificate-cache entries" on page 26228
"show services ssl proxy counters " on page 26233	"show services ssl proxy profile " on page 26239	"show services ssl proxy session-cache entries" on page 26243
"show services ssl proxy session-cache statistics" on page 26249	"show services ssl proxy statistics" on page 26252	"show services ssl proxy status" on page 26255

"show services ssl session" on page 26258	"show services ssl termination counters" on page 26262	"show services ssl termination profile " on page 26267
"show services stateful-firewall conversations" on page 26273	"show services stateful-firewall flow-analysis" on page 26279	"show services stateful-firewall flows" on page 26286
"show services stateful-firewall sip-call" on page 26294	"show services stateful-firewall sip-register" on page 26300	"show services stateful-firewall statistics" on page 26305
"show services stateful-firewall statistics application-protocol sip" on page 26318	"show services stateful-firewall subscriber-analysis" on page 26322	"show services subscriber analysis" on page 26327
"show services tcp-log" on page 26330	"show services tcp-log connections" on page 26331	"show services traffic-detection-function hcm statistics" on page 26333
"show services traffic-detection-function sessions" on page 26338	"show services traffic-load-balance statistics" on page 26341	"show services unified-access-control authentication-table" on page 26358
"show services unified-access-control counters" on page 26362	"show services unified-access-control policies" on page 26365	"show services unified-access-control roles" on page 26367
"show services unified-access-control status" on page 26369	"show services url-filter dns-resolution profile" on page 26370	"show services url-filter dns-resolution-statistics profile template" on page 26374
"show services url-filter statistics profile template" on page 26381	"show services user-identification active-directory-access domain-controller status" on page 26385	"show services user-identification active-directory-access statistics" on page 26389
"show services user-identification active-directory-access user-group-mapping" on page 26393	"show services user-identification authentication-table" on page 26397	"show services user-identification device-information table" on page 26421

"show services user-identification logical-domain-identity-management" on page 26426	"show services user-identification validate-statistics" on page 26430	"show services video-monitoring mdi errors fpc-slot" on page 26431
"show services video-monitoring mdi flows fpc-slot" on page 26434	"show services video-monitoring mdi stats fpc-slot" on page 26441	"show services web-filter dns-resolution profile" on page 26445
"show services web-filter dns-resolution-statistics profile template" on page 26449	"show services web-filter secintel-policy status " on page 26455	"show services web-filter statistics dns-filter-template" on page 26461
"show services web-filter statistics profile" on page 26464	"show services web-proxy dns forwarding-cache " on page 26470	"show services web-proxy dns global-cache statistics" on page 26473
"show services web-proxy session" on page 26476	"show sflow" on page 26478	"show sflow collector" on page 26482
"show sflow interface" on page 26485	"show snmp health-monitor" on page 26489	"show snmp inform-statistics" on page 26499
"show snmp mib" on page 26502	"show snmp rmon" on page 26516	"show snmp rmon history" on page 26523
"show snmp statistics" on page 26529	"show snmp stats-response-statistics" on page 26539	"show snmp v3" on page 26544
"show spanning-tree bridge" on page 26548	"show spanning-tree interface" on page 26555	"show spanning-tree mstp configuration" on page 26568
"show spanning-tree statistics" on page 26572	"show spanning-tree statistics bridge" on page 26575	"show spanning-tree statistics interface" on page 26576
"show spanning-tree statistics message-queues" on page 26578	"show spanning-tree stp-buffer see-all" on page 26580	"show spring-traffic-engineering" on page 26587

"show ssl-certificates" on page 26602	"show static-subscribers sessions" on page 26605	"show subscribers" on page 26609
"show subscribers summary" on page 26663	"show synchronous-ethernet esmc statistics" on page 26674	"show synchronous-ethernet esmc transmit" on page 26676
"show synchronous-ethernet global-information" on page 26679	"show system alarms" on page 26682	"show system alarms" on page 26683
"show system anomalies (GNF)" on page 26689	"show system anomalies gnf-id" on page 26692	"show system applications (Junos OS Evolved)" on page 26695
"show system audit" on page 26709	"show system auto-snapshot" on page 26720	"show system autoinstallation status" on page 26723
"show system autorecovery state" on page 26725	"show system boot console" on page 26727	"show system boot-messages" on page 26729
"show system buffers" on page 26741	"show system certificate" on page 26751	"show system commit" on page 26754
"show system commit server queue" on page 26758	"show system commit server status" on page 26763	"show system configuration archival" on page 26765
"show system configuration database status" on page 26766	"show system configuration database usage" on page 26768	"show system configuration rescue" on page 26770
"show system configuration revision" on page 26772	"show system connections" on page 26775	"show system core dumps (Junos OS Evolved)" on page 26802
"show system core-dumps" on page 26783	"show system cpu (JDM)" on page 26804	"show system data-models" on page 26807
"show system directory-usage" on page 26808	"show system download" on page 26815	"show system errors" on page 26817

"show system errors active" on page 26822	"show system errors count" on page 26829	"show system errors error-id" on page 26831
"show system errors fru" on page 26834	"show system errors history" on page 26841	"show system errors inactive" on page 26844
"show system filesystem encryption status" on page 26852	"show system firmware" on page 26855	"show system firmware (Junos OS Evolved)" on page 26865
"show system health-monitor" on page 26868	"show system information" on page 26873	"show system inventory hardware cpu" on page 26874
"show system khms-stats" on page 26889	"show system license (View)" on page 26896	"show system login lockout" on page 26902
"show system mac-addresses (JDM)" on page 26903	"show system max-smbios-version (JDM)" on page 26907	"show system memory" on page 26908
"show system memory (JDM)" on page 26918	"show system monitor memory actions" on page 26921	"show system monitor memory events" on page 26923
"show system monitor memory status" on page 26926	"show system name-resolution" on page 26929	"show system network (JDM)" on page 26931
"show system node-attributes" on page 26936	"show system nodes" on page 26939	"show system ntp threshold " on page 26941
"show system packet-forwarding-options hw-resource-monitor monitor-profile" on page 26943	"show system packet-forwarding-options hw-resource-monitor resource-list" on page 26946	"show system packet-forwarding-options hw-resource-monitor utilization-info" on page 26948
"show system packet-forwarding-options hw-resource-utilization-info" on page 26952	"show system processes" on page 26957	"show system queues" on page 26975
"show system random-mac-prefix (JDM)" on page 26981	"show system reboot" on page 26982	"show system resource-monitor " on page 26987

"show system resource-monitor ifd-cos-queue-mapping fpc" on page 26992	"show system resource-monitor subscribers-limit" on page 26995	"show system rollback" on page 27001
"show system schema" on page 27004	"show system security-profile" on page 27008	"show system security-profile scheduler" on page 27016
"show system security-profile secure-wire" on page 27020	"show system security-profile security-log-stream-number" on page 27025	"show system security-profile security-log-stream-number detail" on page 27028
"show system security-profile security-log-stream-number summary" on page 27032	"show system services dhcp binding" on page 27034	"show system services dhcp client" on page 27038
"show system services dhcp conflict" on page 27043	"show system services dhcp global" on page 27045	"show system services dhcp pool" on page 27047
"show system services dhcp relay-statistics" on page 27051	"show system services dhcp statistics" on page 27053	"show system services dns dns-proxy" on page 27057
"show system services dynamic-dns" on page 27061	"show system services outbound-ssh jefflee" on page 27063	"show system services service-deployment" on page 27067
"show system snapshot (Junos OS)" on page 27069	"show system snapshot (Junos OS Evolved)" on page 27073	"show system snapshot (Junos OS with Upgraded FreeBSD)" on page 27075
"show system snapshot media" on page 27078	"show system software" on page 27081	"show system software add-restart (Junos OS Evolved)" on page 27086
"show system software list" on page 27088	"show system software recovery-package" on page 27092	"show system software restore-point-status" on page 27094
"show system software upgrade status" on page 27096	"show system software usb-software-version" on page 27098	"show system statistics" on page 27101

"show system statistics arp" on page 27119	"show system statistics icmp" on page 27129	"show system statistics icmp6" on page 27136
"show system statistics igmp" on page 27146	"show system statistics ip" on page 27151	"show system statistics ip6" on page 27162
"show system statistics netisr per-priority-stats" on page 27173	"show system statistics netisr qdrops" on page 27174	"show system statistics netisr socket-drops" on page 27176
"show system statistics tcp" on page 27178	"show system storage" on page 27190	"show system storage partitions" on page 27198
"show system storage satellite " on page 27202	"show system subscriber-management redundancy-state dhcp active-leasequery interface" on page 27205	"show system subscriber-management redundancy-state interface" on page 27208
"show system subscriber-management resiliency" on page 27211	"show system subscriber-management route" on page 27222	"show system subscriber-management statistics" on page 27229
"show system subscriber-management summary" on page 27242	"show system switchover" on page 27246	"show system unified-services status" on page 27251
"show system uptime" on page 27253	"show system users" on page 27260	"show system virtual-memory" on page 27266
"show system visibility cpu" on page 27279	"show system visibility host" on page 27286	"show system visibility jcp" on page 27296
"show system visibility jdm" on page 27300	"show system visibility memory" on page 27304	"show system visibility network" on page 27308
"show system visibility vnf" on page 27316	"show system yang package" on page 27323	"show system ztp" on page 27325

"show task" on page 27331	"show task io" on page 27336	"show task logical-system-mux" on page 27339
"show task memory" on page 27342	"show task replication" on page 27348	"show tcp-forwarding status" on page 27351
"show ted database" on page 27357	"show ted link" on page 27377	"show ted protocol" on page 27384
"show ted spring-te-policy" on page 27386	"show trace" on page 27392	"show unified-edge tdf aaa radius client statistics" on page 27396
"show unified-edge tdf aaa radius client status" on page 27404	"show unified-edge tdf aaa radius network-element statistics" on page 27406	"show unified-edge tdf aaa radius server statistics" on page 27410
"show unified-edge tdf aaa radius server status" on page 27416	"show unified-edge tdf aaa radius snoop-segment statistics" on page 27420	"show unified-edge tdf aaa statistics" on page 27425
"show unified-edge tdf address-assignment pool" on page 27437	"show unified-edge tdf address-assignment service-mode" on page 27442	"show unified-edge tdf address-assignment statistics" on page 27445
"show unified-edge tdf call-admission-control statistics" on page 27448	"show unified-edge tdf call-rate statistics " on page 27452	"show unified-edge tdf diameter network-element statistics" on page 27456
"show unified-edge tdf diameter network-element status" on page 27459	"show unified-edge tdf diameter pcc-gx statistics" on page 27462	"show unified-edge tdf diameter peer statistics" on page 27469
"show unified-edge tdf diameter peer status" on page 27476	"show unified-edge tdf domain service-mode" on page 27481	"show unified-edge tdf domain statistics" on page 27484
"show unified-edge tdf resource-manager clients" on page 27490	"show unified-edge tdf service-mode" on page 27493	"show unified-edge tdf statistics" on page 27496

"show unified-edge tdf status" on page 27507	"show unified-edge tdf subscribers" on page 27513	"show unified-edge tdf system interfaces" on page 27532
"show unified-edge tdf system interfaces service-mode" on page 27534	"show uplink-failure-detection" on page 27538	"show uplink-failure-detection" on page 27541
"show v4ov6-tunnels" on page 27544	"show validation database" on page 27547	"show validation group" on page 27551
"show validation replication database" on page 27553	"show validation session" on page 27556	"show validation statistics" on page 27563
"show version (JDM)" on page 27567	"show version (Junos OS)" on page 27569	"show version (Junos OS Evolved)" on page 27572
"show version vnf" on page 27575	"show virtual-chassis" on page 27578	"show virtual-chassis active-topology" on page 27587
"show virtual-chassis active-topology (MX Series Virtual Chassis)" on page 27593	"show virtual-chassis device-topology" on page 27596	"show virtual-chassis device-topology (MX Series Virtual Chassis)" on page 27605
"show virtual-chassis fast-failover" on page 27609	"show virtual-chassis heartbeat (MX Series Virtual Chassis)" on page 27611	"show virtual-chassis login" on page 27616
"show virtual-chassis mode" on page 27617	"show virtual-chassis protocol adjacency" on page 27622	"show virtual-chassis protocol adjacency (MX Series Virtual Chassis)" on page 27627
"show virtual-chassis protocol database" on page 27633	"show virtual-chassis protocol database (MX Series Virtual Chassis)" on page 27638	"show virtual-chassis protocol interface" on page 27647
"show virtual-chassis protocol interface (MX Series Virtual Chassis)" on page 27650	"show virtual-chassis protocol route" on page 27655	"show virtual-chassis protocol route (MX Series Virtual Chassis)" on page 27659

"show virtual-chassis protocol statistics" on page 27663	"show virtual-chassis protocol statistics (MX Series Virtual Chassis)" on page 27668	"show virtual-chassis status (MX Series Virtual Chassis)" on page 27672
"show virtual-chassis vc-path" on page 27676	"show virtual-chassis vc-port" on page 27681	"show virtual-chassis vc-port diagnostics optics" on page 27688
"show virtual-chassis vc-port (MX Series Virtual Chassis)" on page 27709	"show virtual-chassis vc-port statistics" on page 27713	"show virtual-network-functions" on page 27723
"show virtual-network-functions (Junos Node Slicing)" on page 27728	"show vlans" on page 27738	"show vlans evpn nd-table" on page 27769
"show vlans (JDM)" on page 27767	"show vlans vlan-groups" on page 27774	"show vmhost bridge" on page 27775
"show vmhost connections" on page 27778	"show vmhost control-plane" on page 27780	"show vmhost crash" on page 27782
"show vmhost crash (NFX Series)" on page 27784	"show vmhost forwarding-options analyzer" on page 27785	"show vmhost hard-disk-test " on page 27787
"show vmhost hardware" on page 27789	"show vmhost information" on page 27791	"show vmhost jdm status (In-Chassis Model)" on page 27793
"show vmhost logs" on page 27795	"show vmhost management-if" on page 27797	"show vmhost memory" on page 27798
"show vmhost mode" on page 27799	"show vmhost netstat" on page 27811	"show vmhost network external-network" on page 27813
"show vmhost processes" on page 27815	"show vmhost resource-usage" on page 27818	"show vmhost snapshot" on page 27821
"show vmhost snmp mib walk" on page 27823	"show vmhost status" on page 27825	"show vmhost status (In-Chassis Model)" on page 27829

"show vmhost status (NFX Series)" on page 27827	"show vmhost storage" on page 27831	"show vmhost uptime" on page 27838
"show vmhost uptime (NFX Series)" on page 27840	"show vmhost version" on page 27841	"show vmhost version (NFX Series)" on page 27844
"show vmhost vlans" on page 27847	"show vpls connections" on page 27849	"show vpls flood event-queue" on page 27865
"show vpls flood instance" on page 27867	"show vpls flood route" on page 27871	"show vpls mac-move-action" on page 27874
"show vpls mac-table" on page 27876	"show vpls statistics" on page 27884	"show vrrp" on page 27888
"show vrrp" on page 27905	"show vrrp track" on page 27920	"show wlan access-points" on page 27926
"show evpn gbp-src-tag" on page 19749	"show-spring-traffic-statistics" on page 26600	"show-system-inventory-hardware-memory" on page 26881
"show-system-inventory-hardware-optics" on page 26884	"show-system-inventory-hardware-storage" on page 26886	"show-ted-ipv4-prefix" on page 27372
"show-ted-ipv6-prefix" on page 27374	"ssh" on page 28422	"start shell" on page 28498
"status" on page 28603		

T

"telnet" on page 29056	"test aaa authd-lite user" on page 29156	"test aaa dhcp user" on page 29161

"test aaa ppp user" on page 29169	"test access profile" on page 29178	"test access radius-server" on page 29184
"test configuration" on page 29190	"test interface restart-auto-negotiation" on page 29195	"test msdp" on page 29202
"test policy" on page 29209	"test security utm anti-spam" on page 29215	"test security utm enhanced-web-filtering url-check" on page 29219
"test security utm web-filtering profile" on page 29224	"test services l2tp tunnel" on page 29227	"test services monitoring rfc2544" on page 29230
"test services rpm rfc2544-benchmarking test" on page 29234	"top" on page 29435	"traceroute" on page 29909
"traceroute clns" on page 29917	"traceroute ethernet" on page 29919	"traceroute fabric unicast-flow" on page 29924
"traceroute monitor" on page 29927	"traceroute mpls bgp" on page 29930	"traceroute mpls ldp" on page 29935
"traceroute mpls rsvp" on page 29940	"traceroute mpls segment-routing isis" on page 29946	"traceroute mpls segment-routing ospf" on page 29953
"traceroute mpls segment-routing spring-te" on page 29959	"traceroute mpls segment-routing static egress-ip" on page 29970	"traceroute overlay" on page 29976
"traceroute srv6" on page 29982		

U

"unprotect" on page 30428	"up" on page 30436	"update" on page 30438

V

"vlan-groups" on page 30792	

W

"when" on page 31104	"wildcard delete" on page 31114	

Others

"| (pipe)" on page 31175	

5

CHAPTER

CLI Reference

0-9 | 1578

A | 1584

B | 2786

C | 3124

D | 4790

E | 6166

F | 6727

G | 7640

H | 7956

I | 8223

J | 9422

K | 9448

L | 9508

M | 10199

N | 11149

O | 11741

P | 12095

Q | 13886

R | 13958

S | 16172

T | 28864

U | 30295

V | 30624

W | 31012

X | 31150

Y | 31163

Z | 31169

Others | 31175

0-9

IN THIS SECTION

- 3gpp-imsi | 1578
- 802.3ad (EX and QFX Series) | 1579
- 802.3ad | 1582

3gpp-imsi

IN THIS SECTION

- Syntax | 1578
- Hierarchy Level | 1579
- Description | 1579
- Required Privilege Level | 1579
- Release Information | 1579

Syntax

```
3gpp-imsi {
    equals value;
    has-prefix value;
    has-suffix value;
    matches value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from]
```

Description

Specify the RADIUS AVP 3GPP-IMSI for the incoming RADIUS request from the subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers*
- *Understanding Selection of Properties for an IP-Based TDF Subscriber*
- *Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber*
- *IP-Based Subscriber Setup Overview*

802.3ad (EX and QFX Series)

IN THIS SECTION

Syntax | 1580

- Hierarchy Level (EX Series) | 1580
- Hierarchy Level (NFX, OCX, and QFX Series) | 1581
- Description | 1581
- Options | 1581
- Required Privilege Level | 1581
- Release Information | 1581

Syntax

Syntax (EX Series)

```
802.3ad {
    aex;
    (backup | primary);
    lacp {
        force-up;
        port-priority
    }
}
```

Syntax (NFX, OCX, and QFX Series)

```
802.3ad ae{x;
    lacp {
        force-up;
        (primary | backup);
    }
    port-priority;
}
```

Hierarchy Level (EX Series)

```
[edit interfaces interface-name ether-options]
```

Hierarchy Level (NFX, OCX, and QFX Series)

```
[edit interfaces interface-name ether-options]
```

Description

Specify the aggregated Ethernet logical interface number.

NOTE:

- The "port-priority" on page 12856 statement is not supported on QFabric systems.
- The "force-up" on page 7376 statement is not supported on QFX10002 switches.

Options

- ae*x*—Name of the LAG. Aggregated Ethernet logical interface number.
- backup—Designate the interface as the backup interface for link-protection mode.
- primary—Designate the interface as the primary interface for link-protection mode.

The remaining statements are described separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches

Configuring Aggregated Ethernet Links (CLI Procedure)

Configuring Aggregated Ethernet LACP (CLI Procedure)

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

Configuring Link Aggregation

Configuring Aggregated Ethernet LACP (CLI Procedure)

Understanding Aggregated Ethernet Interfaces and LACP for Switches

Troubleshooting an Aggregated Ethernet Interface

Junos OS Network Interfaces Library for Routing Devices

802.3ad

IN THIS SECTION

- Syntax | **1582**
- Hierarchy Level | **1583**
- Description | **1583**
- Options | **1583**
- Required Privilege Level | **1583**
- Release Information | **1584**

Syntax

```
802.3ad {
    primary | backup;
    ae interface-number ;
    lacp {
        port-priority priority-number;
    }
    link-index index-number
```

```
    distribution-list distribution-list-number
}
```

Hierarchy Level

```
[edit interfaces interface-name fastether-options],
[edit interfaces interface-name gigether-options]
```

Description

Specify aggregated Ethernet logical interface number.

Options

bundle Join an aggregated Ethernet interface.

ae *interface-number* Aggregated Ethernet logical interface number. For MX Series routers running Junos release 14.2R3 and later you can configure a maximum of 1000 aggregated interfaces. On MX2010 and MX2020 routers you can configure a maximum of 800 aggregated interfaces.

primary | backup For link protection configurations, specify the link as primary link or backup link for egress traffic.

lacp Configure Link Aggregation Control Protocol. Specify the port priority in the range 0 through 65535. Default port-priority is 127.

link-index Specify the desired child link index within the aggregated Ethernet Interface. Index number of the logical interface reflects its initialization sequence.

distribution-list For targeted distribution, specify the distribution list to which the interface belongs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`primary` and `backup` options added in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring an Aggregated Ethernet Interface

Configuring Aggregated Ethernet Link Protection

A

IN THIS SECTION

- aaa | **1605**
- aaa-access-profile (L2TP LNS) | **1607**
- aaa clients (TDF) | **1608**
- aaa-context (AAA Options) | **1610**
- aaa-logical-system (Domain Map) | **1611**
- aaa-options (Access Profile) | **1613**
- aaa-options (PPP Profile) | **1614**
- aaa-policy-control (PCEF Profile) | **1616**
- aaa-profile (PCEF Profile) | **1618**
- aaa-routing-instance (Domain Map) | **1619**
- aacl-fields | **1621**
- aacl-statistics-profile | **1623**
- abstract-hop | **1625**
- accept (Router Advertisement Guard Policy) | **1627**
- accept (Dynamic Profile) | **1629**
- accept-data | **1631**
- accept-join-always-from | **1632**

- accept-max-tcp-connections (System Process) | **1634**
- accept-out-of-band | **1636**
- accept-own | **1637**
- accept-prpd-connection | **1639**
- accept-remote-nexthop | **1640**
- accept-remote-source | **1642**
- accept-sdr (PCRF Partition) | **1644**
- accept-source-mac (SRX) | **1645**
- accept-source-mac | **1647**
- accept-srv6-service | **1649**
- accepted-prefix-limit | **1651**
- access-concentrator | **1655**
- access-domain (Remote Device Management) | **1657**
- access (Dynamic Access Routes) | **1659**
- access-identifier | **1661**
- access-interfaces (IFL Subscriber) | **1662**
- access-interfaces (IP Subscriber) | **1664**
- access-internal (Dynamic Access-Internal Routes) | **1666**
- access-internal (Static Access-Internal Routes) | **1668**
- access-line (Access-Line Rate Adjustment) | **1669**
- access-line-information (L2TP) | **1684**
- access-list | **1686**
- access-point name | **1687**
- access-profile | **1688**
- access-profile (AAA Options) | **1690**
- access-profile (Domain Map) | **1691**
- access-profile (Dynamic Stacked VLAN) | **1693**
- access-profile (Dynamic VLAN) | **1694**
- access-profile (Extensible Subscriber Services Manager) | **1695**
- access-profile-name (Duplicate Accounting) | **1697**
- access-profile (Static Subscribers) | **1698**
- access-security | **1700**

- access (SNMP) | **1702**
- access (Static Framed Routes) | **1703**
- access-trunk | **1705**
- account-layer2-overhead (PIC Level) | **1706**
- accounting (Forwarding Options) | **1707**
- accounting (interfaces) | **1709**
- accounting (AAA Profile) | **1710**
- accounting (Access Profile) | **1711**
- accounting-backup-options (Access Profile) | **1713**
- accounting (Dynamic IGMP Interface) | **1714**
- accounting (Dynamic MLD Interface) | **1716**
- accounting-options | **1717**
- accounting-order | **1721**
- accounting-order (Service Accounting) | **1722**
- accounting-port (RADIUS Server) | **1724**
- accounting-profile (firewall) | **1726**
- accounting-profile | **1727**
- accounting (Protocols AMT Interface) | **1728**
- accounting (Protocols IGMP) | **1730**
- accounting (Protocols IGMP AMT Interface) | **1731**
- accounting (Protocols IGMP Interface) | **1732**
- accounting (Protocols MLD) | **1733**
- accounting (Protocols MLD Interface) | **1735**
- accounting (RADIUS Client) | **1736**
- accounting-secret (RADIUS Server) | **1737**
- accounting-server | **1739**
- accounting-server-name | **1740**
- accounting (Service Accounting) | **1742**
- accounting-stop-on-access-deny | **1743**
- accounting-stop-on-failure | **1744**
- accounting (System) | **1746**
- acfc | **1748**

- ack-number | **1750**
- acknowledge-retries | **1751**
- acknowledge-timer | **1752**
- acquiring-state-announce-grandmaster | **1754**
- action (QFX) | **1755**
- action (chassis error) | **1756**
- action | **1759**
- action (APN GTP) | **1760**
- action (OAM) | **1762**
- action (OAM LFM) | **1764**
- action (Policer) | **1765**
- action-priority (Routing Instance) | **1766**
- action-priority | **1768**
- action-priority (bridge domain) | **1769**
- action-profile | **1771**
- action-profile (Network Management) | **1773**
- action-profile (Applying to CFM) | **1775**
- action-profile (Applying to OAM CFM, for EX Series Switch Only) | **1776**
- action-profile (Defining for CFM) | **1778**
- action-profile (MEP) | **1780**
- action-red-differential-delay | **1782**
- action (Security Content Security Web Filtering) | **1783**
- action (Security Rulebase IPS) | **1785**
- action-shutdown | **1787**
- action (Uplink Failure Detection) | **1789**
- actions | **1790**
- actions (Services SSL Initiation) | **1793**
- activate | **1795**
- activation-attribute (AAA Profile) | **1796**
- activation-delay | **1797**
- activation-priority | **1799**
- active | **1800**

- active-directory-access | **1802**
- active-interface (VPLS Multihoming) | **1805**
- active-leasequery (DHCP Local Server) | **1806**
- active-leasequery (DHCP Relay Agent) | **1808**
- active-policy | **1811**
- active-probe-params | **1813**
- active-server-group | **1817**
- active-source-limit | **1819**
- activeness-probe | **1821**
- actual-transit-statistics (Dynamic Profiles) | **1823**
- adaptive (protocols mpls) | **1825**
- adaptive | **1826**
- adaptive (RSVP) | **1828**
- adaptive-sample-rate | **1830**
- adaptive-services | **1832**
- adaptive-services-pics | **1834**
- adaptive-shaper | **1835**
- adaptive-shapers | **1836**
- add-attribute-length-in-pdu | **1838**
- add-path | **1839**
- add-path-display-ipv4-address | **1842**
- address (tunnel-encapsulation) | **1843**
- address (SNMP) | **1844**
- address | **1846**
- address (IP Address) | **1849**
- address (Address Book Next Gen Services) | **1852**
- address-allocation | **1853**
- address (Analytics Collector) | **1855**
- address (Anycast RPs) | **1856**
- address-assignment (Access) | **1858**
- address-assignment (Address-Assignment Pools) | **1862**
- address (Bidirectional Rendezvous Points) | **1865**

- address-blacklist | **1867**
- address-book | **1868**
- address-book (System) | **1870**
- address-change-immediate-update | **1872**
- address (Diameter Peer) | **1874**
- address (Diameter Transport) | **1875**
- address-family | **1876**
- address (Interfaces) | **1878**
- address (Interfaces, Flow Monitoring) | **1879**
- address (L2TP Destination) | **1880**
- address (L2TP Tunnel Destination) | **1882**
- address (LNS Local Gateway) | **1883**
- address (Local RPs) | **1885**
- address (LRF Profile) | **1886**
- address-mapping | **1887**
- address-mapping (Application Identification) | **1889**
- address-mask | **1891**
- address-name | **1892**
- address (NAT Pool Next Gen Services) | **1894**
- address (Peer) | **1895**
- address-persistent (Security Source NAT) | **1896**
- address-persistent (Security Source NAT Pool) | **1897**
- address-pool | **1899**
- address-pool (Access) | **1901**
- address-pool (Domain Map) | **1902**
- address-pooling | **1904**
- address-pooling (Security Source NAT) | **1905**
- address-pooling (Source NAT Next Gen Services) | **1907**
- address-pools | **1908**
- address-protection | **1910**
- address (Protocols OpenFlow) | **1912**
- address (Protocols Router Discovery) | **1913**

- address (RADIUS Clients) | **1915**
- address (RADIUS Server) | **1916**
- address-range | **1917**
- address-ranges (Demux) | **1919**
- address (Security Address Book) | **1920**
- address (Security ARP Proxy) | **1922**
- address (Security Destination NAT) | **1924**
- address (Security NDP Proxy) | **1925**
- address (Services) | **1926**
- address (Services Dynamic Flow Capture) | **1928**
- address (Services NAT Pool) | **1929**
- address-set | **1930**
- address-shared (Security Source NAT) | **1932**
- address (Static RPs) | **1933**
- address (Tunnel Profile Remote Gateway) | **1935**
- address (Tunnel Profile Source Gateway) | **1936**
- address-whitelist | **1937**
- addresses (REST API) | **1938**
- adf (Dynamic Firewalls) | **1939**
- adjacency-loss | **1941**
- adjacency-loss-hold-time (ANCP) | **1943**
- adjacency-timer | **1944**
- adjust-interval | **1946**
- adjust-minimum | **1947**
- adjust-minimum (Dynamic Shaping and Scheduling) | **1949**
- adjust-percent | **1950**
- adjust-percent (Dynamic Schedulers) | **1952**
- adjust-threshold | **1953**
- adjust-threshold-absolute | **1954**
- adjust-threshold-activate-bandwidth | **1956**
- adjust-threshold-overflow-limit | **1958**
- adjust-threshold-underflow-limit | **1959**

- adjustment-control-profiles | **1961**
- admin-down | **1962**
- admin-email | **1963**
- admin-group (Backup Selection) | **1965**
- admin-group | **1967**
- admin-group-extended | **1968**
- admin-group (for Interfaces) | **1971**
- admin-group (for LSPs) | **1972**
- admin-groups | **1973**
- admin-groups-extended | **1975**
- admin-groups-extended-range | **1977**
- admin-search | **1978**
- administrator-email (Security Fallback Block) | **1980**
- administrator-email (Security Virus Detection) | **1981**
- adsl-bytes | **1982**
- adsl2-bytes | **1984**
- adsl2-plus-bytes | **1986**
- advance-policy-based-routing | **1988**
- advance-policy-based-routing-profile (Security Zones) | **1993**
- advanced-anti-malware connection | **1994**
- advanced-anti-malware policy | **1997**
- advanced-anti-malware traceoptions | **2005**
- advertise | **2007**
- advertise-bgp-static | **2009**
- advertise-external | **2011**
- advertise-from-main-vpn-tables | **2013**
- advertise-inactive | **2015**
- advertise-interval | **2017**
- advertise-mode (MPLS) | **2018**
- advertise-peer-as | **2020**
- advertise-prefix-sid | **2022**
- advertise-srv6-service | **2024**

- advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper) | **2026**
- advertise-zero-ethernet-tag | **2028**
- advertisement | **2030**
- advertisement (Delay Measurement) | **2031**
- advertisement-hold-time | **2033**
- advisory-options (Traffic Shaping) | **2035**
- advpn | **2036**
- af-mode | **2039**
- af-name | **2040**
- age | **2042**
- age-of-attack | **2043**
- agent-address | **2044**
- agent (Analytics) | **2046**
- agent-circuit-identifier (Dynamic ACI VLANs) | **2048**
- agent-id | **2050**
- agent-specifier | **2051**
- agf | **2053**
- aggregate-clients (DHCP Local Server) | **2056**
- aggregate-clients (DHCP Relay Agent) | **2057**
- aggregate-clients (Static Subscribers) | **2060**
- aggregate-export-interval | **2062**
- aggregate (Gigabit Ethernet CoS Policer) | **2063**
- aggregate (Hierarchical Policer) | **2064**
- aggregate-label | **2066**
- aggregate (Protocols RSVP) | **2068**
- aggregate (Routing) | **2070**
- aggregate-wait-time | **2073**
- aggregated-devices | **2074**
- aggregated-ether | **2077**
- aggregated-ether-options | **2078**
- aggregated-inline-services-options (Aggregated Inline Services) | **2082**
- aggregation (Forwarding Options) | **2084**

- aggregation (IDS) | **2086**
- aggregations (IDS Screen Next Gen Services) | **2088**
- aging | **2089**
- aging-timer (Junos Fusion) | **2091**
- aigp | **2092**
- aigp-adjust (Policy Action) | **2095**
- aigp-originate | **2097**
- ais-trigger-condition | **2099**
- alarm (chassis) | **2100**
- alarm-id | **2102**
- alarm-list-name | **2103**
- alarm low-light-alarm | **2105**
- alarm-management | **2106**
- alarm-mode | **2108**
- alarm (optics-options) | **2110**
- alarm-port | **2111**
- alarm (Satellite Policies) | **2114**
- alarm (SNMP RMON) | **2115**
- alarm-state | **2117**
- alarm-threshold | **2119**
- alarm-threshold (Security GPRS) | **2120**
- alarm-without-drop | **2122**
- alarm-without-drop (IDS Screen Next Gen Services) | **2123**
- alarms | **2124**
- alarms (Video Monitoring) | **2125**
- alarms (Security) | **2128**
- alg | **2130**
- alg (Applications) | **2138**
- alg-manager | **2141**
- alg-support-lib | **2142**
- algorithm | **2144**
- algorithm (Junos FIPS) | **2145**

- alias (Interfaces) | **2146**
- alias (Junos Fusion) | **2148**
- alias (Junos Fusion Satellite Device Clustering) | **2150**
- aliases | **2151**
- aliases (Virtual Chassis) | **2153**
- all-defects | **2155**
- all-regions | **2157**
- all-tcp | **2159**
- allow | **2160**
- allow-6pe-traceroute | **2162**
- allow-active-leasequery (DHCP Local Server) | **2163**
- allow-bulk-leasequery (DHCP Local Server) | **2165**
- allow-clear (Accounting Options) | **2167**
- allow-commands (Op Scripts) | **2169**
- allow-dns-reply | **2170**
- allow-duplicates | **2172**
- allow-dynamic-requests (RADIUS Server) | **2173**
- allow-email (Security Fallback Block) | **2175**
- allow-email (Security Virus Detection) | **2176**
- allow-embedded-icmp | **2177**
- allow-fragmentation (protocols mpls) | **2178**
- allow-fragmentation | **2179**
- allow-fragmentation (fti) | **2181**
- allow-icmp-without-flow | **2182**
- allow-ip-options (IDS MS-MPC) | **2183**
- allow-ip-options (Services Stateful Firewall) | **2186**
- allow-ipv6-extension-header (IDS MS-MPC) | **2188**
- allow-l3vpn-traceroute-src-select | **2190**
- allow-leasequery (DHCP Local Server) | **2191**
- allow-maximum (Multicast) | **2193**
- allow-multicast | **2195**
- allow-no-end-option (DHCP Relay Agent) | **2196**

- allow-overlapping-nat-pools | **2198**
- allow-overlapping-pools (NAT Next Gen Services) | **2199**
- allow-remote-loopback | **2200**
- allow-reverse-ecmp | **2201**
- allow-route-leaking | **2202**
- allow-snooped-clients | **2204**
- allow-sram-parity-errors | **2206**
- allow-subnet-mismatch | **2207**
- allow-transients | **2208**
- allow-url-for-python (Op Scripts) | **2209**
- allow-v4mapped-packets | **2211**
- allowed-destinations | **2212**
- allowed-mac | **2214**
- allowed-sources (REST API) | **2216**
- alt-name (Application Identification) | **2217**
- alternative-partition-name (OCS Partition) | **2218**
- always-mark-connection-protection-tlv | **2219**
- always-write-giaddr | **2221**
- always-write-option-82 | **2222**
- ambient-temperature | **2224**
- amf | **2226**
- amt (IGMP) | **2228**
- amt (Protocols) | **2230**
- analytics | **2231**
- analyzer | **2237**
- analyzer-address | **2241**
- analyzer-id | **2243**
- anchor-point (Pseudowire Subscriber Interfaces) | **2244**
- ancp | **2247**
- ancp (Adjustment Control Profiles) | **2249**
- ancp-speed-change-immediate-update (ANCP) | **2251**
- annotate | **2252**

- anomaly | **2254**
- anti-replay-window-size | **2255**
- anti-replay-window-size (Services IPsec VPN) | **2256**
- anti-replay-window-size (Services Service Set) | **2258**
- anti-spam (Logical Systems) | **2259**
- anti-spam | **2262**
- anti-spam (Security Content Security Policy) | **2265**
- anti-virus (Logical Systems) | **2266**
- anti-virus (Content Security) | **2269**
- anti-virus (Content Security Policy) | **2273**
- anti-virus (Services) | **2275**
- any-sender | **2281**
- any (VPLS Multihoming) | **2282**
- anycast-pim | **2283**
- anycast-prefix | **2285**
- apn | **2286**
- apn-control-group (Security GTP) | **2288**
- apn-control (Security GTP) | **2289**
- apn-services (control plane services) | **2292**
- app-mapping-timeout | **2293**
- appfw-profile (System) | **2294**
- appfw-rule | **2296**
- appfw-rule-set | **2298**
- application | **2299**
- application (Adjustment Control Profiles) | **2301**
- application (Application Identification) | **2303**
- application (Application Identification) | **2306**
- application (Application Maps) | **2308**
- application (Applications) | **2310**
- application-bypass (Juniper Secure Connect) | **2314**
- application (Extensions) | **2316**
- application-firewall | **2318**

- application-firewall (Application Services) | **2320**
- application-group | **2322**
- application-group-any (AACL) | **2323**
- application-group (Services) | **2325**
- application-groups (PCC Rules) | **2326**
- application-groups (Services AACL) | **2328**
- application-identification (Services) | **2330**
- application-identification (IDP Sensor Configuration) | **2336**
- application-identification (Application Identification) | **2337**
- application-identification-profile (Service Set) | **2343**
- application-map | **2345**
- application-maps | **2346**
- application (NAT Next Gen Services) | **2348**
- application-profile | **2349**
- application-profile (Services CoS Next Gen Services) | **2351**
- application-protocol | **2352**
- application-protocol (Applications) | **2355**
- application-proxy (Security Content Security) | **2358**
- application-screen (Security H323) | **2359**
- application-screen (Security MGCP) | **2361**
- application-screen (Security SCCP) | **2362**
- application-screen (Security SIP) | **2364**
- application (Security Alarms) | **2366**
- application (Security Custom Attack) | **2367**
- application (Security Destination NAT) | **2369**
- application (Security IDP) | **2370**
- application (Security Policies) | **2371**
- application (Security Source NAT) | **2373**
- application-services (Security Forwarding Process) | **2375**
- application-services (security-metadata-streaming) | **2377**
- application-services (Security Policies) | **2378**
- application-set | **2382**

- application-sets (IDS MS-DPC) | **2383**
- application-sets (PCP) | **2384**
- application-sets (Services CoS) | **2386**
- application-sets (Services NAT) | **2387**
- application-sets (Services Stateful Firewall) | **2388**
- application-specific (Protocols OSPF) | **2389**
- application-status | **2392**
- application-system-cache | **2393**
- application-system-cache-timeout (Services) | **2395**
- application-tracking | **2397**
- application-tracking (Security Zones) | **2399**
- application-traffic-control | **2400**
- application-traffic-control (Application Services) | **2402**
- application-unknown | **2404**
- applications (Applications) | **2405**
- applications (DCBX) | **2406**
- applications (Diameter) | **2408**
- applications (IDS MS-DPC) | **2409**
- applications (PCC Rules) | **2410**
- applications (PCP) | **2412**
- applications (Services AACL) | **2414**
- applications (Services ALGs) | **2415**
- applications (Services Application Identification) | **2416**
- applications (Services CoS) | **2418**
- applications (Services NAT) | **2419**
- applications (Services Stateful Firewall) | **2420**
- apply-action-profile | **2421**
- apply-groups | **2423**
- apply-groups (Chassis Cluster) | **2424**
- apply-groups-except | **2426**
- apply-groups-except (Subscriber Secure Policy) | **2427**
- apply-groups (Subscriber Secure Policy) | **2428**

- apply-macro | 2430
- apply-path | 2431
- apply-secondary-as-giaddr | 2433
- apply-to-half-close-state | 2435
- apts | 2436
- archival | 2438
- archive (All System Log Files) | 2440
- archive (Individual System Log File) | 2442
- archive (QFabric System) | 2445
- archive-sites | 2446
- archive-sites (Flow Monitoring) | 2448
- archive-sites (Event Policy) | 2449
- area | 2451
- area-range | 2453
- arguments (Event Policy) | 2455
- arguments (Op Scripts) | 2457
- arp | 2458
- arp-detect | 2463
- arp-enhanced-scale | 2465
- arp-inspection | 2466
- arp-inspection (MX Series) | 2469
- arp-l2-validate | 2470
- arp-max-cache | 2471
- arp-new-hold-limit | 2473
- arp policer | 2475
- arp-prefix-limit (Host Fast Reroute) | 2477
- arp-resp | 2479
- arp-system-cache-limit | 2480
- arp-throttle | 2482
- arp-nd-probe-disable | 2484
- arp-nd-probe-failed-log | 2485
- as-external | 2486

- as-list | **2487**
- as-override | **2489**
- as-path-compare | **2490**
- as-path-group | **2492**
- as-path (Policy Options) | **2493**
- as-path (Routing Options) | **2495**
- asm-override-ssm | **2497**
- assemble | **2499**
- assert-timeout | **2500**
- identity-provider attribute-mapping (SAML) | **2501**
- assignment-id-format (L2TP LAC) | **2504**
- assisted-replication | **2505**
- assisted-replication (Deterministic AR Replicator Policy Actions) | **2507**
- associate-backup-pe-groups | **2511**
- associate-lsp | **2512**
- associate-profile | **2513**
- association-timeout | **2515**
- asymmetric-hold-time | **2516**
- asymmetric-hold-time | **2517**
- asymmetric-lease-time (DHCP Overrides) | **2519**
- asymmetric-prefix-lease-time (DHCP Overrides) | **2521**
- asynchronous-notification | **2522**
- atm-options | **2524**
- attach | **2527**
- attack-threshold | **2528**
- attack-type (Security Anomaly) | **2529**
- attack-type (Security Chain) | **2531**
- attack-type (Security IDP) | **2533**
- attack-type (Security Signature) | **2540**
- attacks (Security Exempt Rulebase) | **2546**
- attacks (Security IPS Rulebase) | **2547**
- attempts (DHCP Local Server) | **2549**

- attribute | **2551**
- attributes (Access-Line Rate Adjustment) | **2553**
- attributes (Diameter Gx Profiles) | **2555**
- attributes (JSRC Attributes) | **2557**
- attributes-match (Event Policy) | **2558**
- attributes (RADIUS Attributes) | **2560**
- audible (Security Alarms) | **2563**
- auth-entry | **2564**
- auth-only-browser | **2566**
- auth-user-agent | **2568**
- authentication (Protocols OSPF) | **2570**
- authentication (services) | **2572**
- authentication | **2574**
- authentication (AAA Profile) | **2576**
- authentication-algorithm | **2577**
- authentication-algorithm (Security IKE) | **2580**
- authentication-algorithm (Security IPsec) | **2582**
- authentication-algorithm (Services IKE) | **2584**
- authentication-algorithm (Services IPsec) | **2585**
- authentication (BGP BFD Liveness Detection) | **2588**
- authentication (Demux) | **2592**
- authentication (DHCP Local Server) | **2593**
- authentication (DHCP Relay Agent) | **2595**
- authentication-entry-timeout (Identity Management Advanced Query) | **2597**
- authentication-entry-timeout (Services User Identification) | **2599**
- authentication-key (Protocols MSDP) | **2600**
- authentication-key (SNMP) | **2602**
- authentication-key | **2603**
- authentication-key-chain (Protocols BGP and BMP) | **2605**
- authentication-key-chain (Protocols LDP) | **2607**
- authentication-key-chain (TWAMP) | **2608**
- authentication-key-chains | **2610**

- authentication-key-chains (TCP-AO) | **2612**
- authentication-key (ICCP) | **2614**
- authentication-key (Protocols BGP and BMP) | **2616**
- authentication-key (Protocols LDP) | **2618**
- authentication-key (Protocols RIP) | **2619**
- authentication-key (Protocols RSVP) | **2620**
- authentication (LAG) | **2622**
- authentication-md5 | **2624**
- authentication-method (services) | **2626**
- authentication-method | **2627**
- authentication-method (Interfaces) | **2629**
- authentication-mode | **2630**
- authentication-none | **2632**
- authentication-order | **2633**
- authentication-order (Access Profile) | **2635**
- authentication-order (Authenticator) | **2638**
- authentication-order (System) | **2641**
- authentication-password | **2643**
- authentication-protocol | **2645**
- authentication (Protocols PIM) | **2647**
- authentication (Security Alarms) | **2648**
- authentication (Security IPsec) | **2650**
- authentication-server | **2652**
- authentication-sha | **2653**
- authentication-sha224 | **2654**
- authentication-sha256 | **2656**
- authentication-sha348 | **2657**
- authentication-sha512 | **2658**
- authentication-source (Services User Identification ClearPass) | **2660**
- authentication-source (Services User Identification Device Identity) | **2662**
- authentication (Static and Dynamic PPP) | **2664**
- authentication (Static Subscribers) | **2665**

- authentication-type | **2667**
- authentication-type (Protocols RIP) | **2669**
- authentication-whitelist | **2671**
- authenticator | **2673**
- authorization (SNMP) | **2677**
- authorization (icap-redirect profile) | **2678**
- authorization-order | **2680**
- auto-bandwidth | **2682**
- auto-bandwidth (MPLS Statistics) | **2684**
- auto-bandwidth (MPLS Tunnel) | **2686**
- auto-configuration | **2688**
- auto-configuration (System) | **2689**
- auto-configure | **2692**
- auto-configure (Demux) | **2694**
- auto-configure (Dynamic VLAN Interface Sets) | **2696**
- auto-configure (IPv4) | **2698**
- auto-configure (IPv6) | **2700**
- auto-configure-trigger interface (ANCP) | **2702**
- auto-conversion | **2704**
- auto-dad (SLAAC Snooping) | **2705**
- auto-derive | **2707**
- auto-discovery | **2709**
- auto-discovery-only | **2711**
- auto-export | **2712**
- auto-image-upgrade | **2715**
- auto-load-rebalance | **2717**
- auto-negotiation (QFX Series) | **2718**
- auto-negotiation | **2719**
- auto-policing | **2724**
- auto-provisioned | **2725**
- auto-re-enrollment | **2727**
- auto-re-enrollment (Security) | **2728**

- auto-reconnect | 2732
- auto-recovery | 2733
- auto-rp | 2735
- auto-satellite-conversion (Junos Fusion) | 2736
- auto-sense | 2738
- auto-service-id | 2739
- auto-snapshot | 2741
- auto-sw-sync | 2743
- auto-sw-update | 2745
- auto-targeted-session | 2748
- autodiscovery | 2750
- autodiscovery-only | 2752
- autoinstallation | 2753
- autoinstallation (JNU Satellite Devices) | 2755
- automatic (Security) | 2757
- automatic-site-id | 2758
- automatic (Source NAT Next Gen Services) | 2760
- autonomous | 2761
- autonomous (Dynamic Router Advertisement) | 2763
- autonomous-system | 2764
- autonomous-system-type (Flow Monitoring) | 2767
- autostate-exclude | 2769
- auxiliary | 2771
- auxiliary-spi | 2773
- auxiliary-spi (Security IPsec) | 2774
- avg-fd-twoway-threshold | 2775
- avg-flr-backward-threshold | 2777
- avg-flr-forward-threshold | 2779
- avg-ifdv-twoway-threshold | 2781
- avira-engine | 2783
- avp (L2TP Tunnel Switching) | 2784

aaa

IN THIS SECTION

- Syntax | 1605
- Hierarchy Level | 1605
- Description | 1605
- Options | 1606
- Required Privilege Level | 1606
- Release Information | 1606

Syntax

```
aaa {
    access-profile access-profile {
        config-payload-password config-payload-password;
    }
    client {
        password;
        username;
    }
    use-routing-instance-address;
    no-push-to-identity-management;
}
```

Hierarchy Level

```
[edit security ike gateway gateway-name]
```

Description

Specify that extended authentication is performed in addition to IKE Phase 1 authentication for remote users trying to access a VPN tunnel. This authentication can be through Extended Authentication (XAuth) or Extensible Authentication Protocol (EAP). Include a previously created access profile,

configured with the `edit access profile` statement, to specify the access profile to be used for authentication information.

Options

access-profile *profile-name*	Name of the previously created access profile to use for extended authentication for remote users trying to access a VPN.
config-payload-password	Specify common client password for IKEv2 configuration payload with 1 to 128 characters.
client	Specify an AAA client uernamе and password for each configured authenticator that is allowed to request authentications for supplicants. • `password`—AAA client password with 1 to 128 characters. • `username`—AAA client username with 1 to 128 characters.
no-push-to-identity-management	(Optional) Specify the option to prevent pushing of the connection states to the identity management server.
use-routing-instance-address	Specify an AAA address from pool configured in routing-instance of type virtual-router.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D80.

`config-payload-password` option added in Junos OS Release 20.1R1.

`no-push-to-identity-management` option added in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

IPsec Overview

aaa-access-profile (L2TP LNS)

IN THIS SECTION

- Syntax | 1607
- Hierarchy Level | 1607
- Description | 1607
- Options | 1607
- Required Privilege Level | 1608
- Release Information | 1608

Syntax

```
aaa-access-profile profile-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name],
[edit access profile profile-name client client-name l2tp]
```

Description

Specify a AAA access profile that overrides the AAA access profile configured for the routing instance with the `access-profile` statement. You can configure a profile to specify the RADIUS server settings for a tunnel group or for a LAC client, or both. The AAA access profile configured for the client takes precedence over the AAA access profile configured for the tunnel group, which takes precedence over the access profile configured for the routing instance.

Options

profile-name Name of the local access profile for the tunnel group or client.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the `[edit access profile profile-name client client-name l2tp]` hierarchy level introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

aaa clients (TDF)

IN THIS SECTION

- Syntax | 1608
- Hierarchy Level | 1609
- Description | 1609
- Options | 1609
- Required Privilege Level | 1609
- Release Information | 1609

Syntax

```
aaa {
    clients client-name;
```

```
    apply-groups [group-names];
    apply-groups-except [group-names];
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name]
```

Description

Specify the GGSN, PGW, or BNG RADIUS clients that can send RADIUS requests to a TDF gateway.

Options

client-name RADIUS client name that was previously configured at the `[edit access radius clients]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers*
>
> *Configuring a TDF Gateway*

aaa-context (AAA Options)

IN THIS SECTION

- Syntax | 1610
- Hierarchy Level | 1610
- Description | 1610
- Options | 1610
- Required Privilege Level | 1611
- Release Information | 1611

Syntax

```
aaa-context aaa-context-name;
```

Hierarchy Level

```
[edit access aaa-options aaa-options-name]
```

Description

Specify the logical-system:routing-instance (LS:RI) that the subscriber session uses for AAA (RADIUS) interactions like authenticating and accounting. For example, this may correspond to the LS:RI for a retail ISP that provides services to the subscriber.

NOTE: Only the default logical system is supported.

Options

aaa-context-name	Name of the logical-system:routing-instance.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access

Configuring Username Modification for Subscriber Sessions

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

aaa-logical-system (Domain Map)

IN THIS SECTION

- Syntax | 1612
- Hierarchy Level | 1612
- Description | 1612
- Default | 1612
- Options | 1612
- Required Privilege Level | 1612
- Release Information | 1613

Syntax

```
aaa-logical-system logical-system-name {
     aaa-routing-instance routing-instance-name;
}
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Configure a non-default logical system in which the `authd` daemon sends AAA requests for the domain map.

> **NOTE**: Subscriber management is supported in the default logical system only. The `aaa-logical-system` statement is for future extensions of subscriber management and is not supported in current Junos OS releases.

Default

Default logical system for the subscriber.

Options

logical-system-name—Name of the logical system.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Specifying an AAA Logical System/Routing Instance in a Domain Map

aaa-options (Access Profile)

IN THIS SECTION

- Syntax | 1613
- Hierarchy Level | 1613
- Description | 1614
- Options | 1614
- Required Privilege Level | 1614
- Release Information | 1614

Syntax

```
aaa-options aaa-options-name {
    aaa-context aaa-context-name;
    access-profile profile-name;
    subscriber-context subscriber-context-name
}
```

Hierarchy Level

```
[edit access]
```

Description

Define a set of AAA options for authorizing and configuring a subscriber or set of subscribers with a subscriber access profile.

Options

aaa-options-name	Name of the set of options.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- *Understanding Session Options for Subscriber Access*
- *Configuring Username Modification for Subscriber Sessions*
- *L2TP LNS Inline Service Interfaces*
- *L2TP LNS Inline Service Interfaces*

aaa-options (PPP Profile)

IN THIS SECTION

- Syntax | **1615**
- Hierarchy Level | **1615**

- Description | 1615
- Options | 1616
- Required Privilege Level | 1616
- Release Information | 1616

Syntax

```
aaa-options aaa-options-name;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp ppp-options],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" ppp-options]
```

Description

Specify a set of AAA options that is used for authentication of PPP subscribers. The set of options is defined globally with the `aaa-options` *`aaa-options-name`* statement at the `[edit access]` hierarchy level.

You can specify the option set in a dynamic PPP profile or in a group profile.

- In a dynamic PPP profile—In this case, usernames are examined and modified for dynamic PPP subscribers logging in by means of the subscriber and AAA contexts that are specified in the AAA options set. The option set must include the `access-profile` *`profile-name`* statement to specify the name of a subscriber access profile.

- In a group profile—In this case, usernames are examined and modified for tunneled PPP subscribers on the LNS logging in by means of the subscriber and AAA contexts that are specified in the AAA options set.

> **NOTE**: When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of

options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

This meant that `aaa-options` configured in a group profile is not applied when the dynamic profile includes any PPP-option, even when the dynamic profile does not include `aaa-options`.

Options

aaa-options-name	Name of the set of options.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- *Understanding Session Options for Subscriber Access*
- *Configuring Username Modification for Subscriber Sessions*
- *L2TP LNS Inline Service Interfaces*
- *L2TP LNS Inline Service Interfaces*

aaa-policy-control (PCEF Profile)

IN THIS SECTION

- Syntax | **1617**
- Hierarchy Level | **1617**

- Description | 1617
- Required Privilege Level | 1617
- Release Information | 1617

Syntax

```
aaa-policy-control {
    aaa-profile aaa-profile-name;
    pcc-rulebases [rulebase-name];
    user-password password;
}
```

Hierarchy Level

```
[edit unified-edge pcef profiles profile-name]
```

Description

Configure RADIUS-server-controlled policy management for a policy and charging enforcement function (PCEF) profile.

The remaining statements are explained separately.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Policies That a RADIUS Server Controls

Configuring an AAA Profile

aaa-profile (PCEF Profile)

IN THIS SECTION

- Syntax | 1618
- Hierarchy Level | 1618
- Description | 1618
- Options | 1619
- Required Privilege Level | 1619
- Release Information | 1619

Syntax

```
aaa-profile aaa-profile-name;
```

Hierarchy Level

```
[edit unified-edge pcef profiles profile-name aaa-policy-control]
```

Description

Specify the AAA profile that identifies the RADIUS server policy control parameters for the policy and charging enforcement function (PCEF) profile. The AAA profile must already be defined at the `[edit unified-edge aaa]` hierarchy level.

Options

aaa-profile-name	Name of the AAA profile.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Policies That a RADIUS Server Controls

Configuring an AAA Profile

aaa-routing-instance (Domain Map)

IN THIS SECTION

- Syntax | 1620
- Hierarchy Level | 1620
- Description | 1620
- Default | 1620
- Options | 1620
- Required Privilege Level | 1620
- Release Information | 1620

Syntax

```
aaa-routing-instance (routing-instance-name | default);
```

Hierarchy Level

```
[edit access domain map domain-map-name],
[edit access domain map domain-map-name aaa-logical-system logical-system-name]
```

Description

Configure the routing instance in which the `authd` daemon sends AAA requests for the domain map.

> **NOTE**: Subscriber management is supported in the default logical system only. The `aaa-logical-system` statement, which appears in the CLI, is not supported in current Junos OS releases.

Default

Routing instance used for the subscriber context.

Options

routing-instance-name—Name of the routing instance.

`default`—The default routing instance.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`default` option added in Junos OS Release 13.3.

RELATED DOCUMENTATION

Specifying an AAA Logical System/Routing Instance in a Domain Map

Domain Mapping Overview

aacl-fields

IN THIS SECTION

- Syntax | 1621
- Hierarchy Level | 1621
- Description | 1621
- Options | 1622
- Required Privilege Level | 1622
- Release Information | 1622

Syntax

```
aacl-fields {
    field-name;
}
```

Hierarchy Level

```
[edit system services local-policy-decision-function statistics aacl-statistics-profile profile-name]
```

Description

Define the statistics to collect in a data log file.

Options

field-name—Name of the field:

- address—IPv4 address
- all-fields—All available fields
- application—Application name
- application-group—Application group name
- input-bytes—Number of input bytes
- input-interface—Input interface name
- input-packets—Number of input packets
- ipv6-address—IPv6 address
- ipv6-prefix-length—Prefix length associated with the displayed IPv6 address
- mask—Netmask
- output-bytes—Number of output bytes
- output-packets—Number of output packets
- subscriber-name—Subscriber name
- timestamp—Timestamp
- vrf-name—VPN routing and forwarding (VRF) name

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

IPv6 support introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Statistics Profiles

aacl-statistics-profile

IN THIS SECTION

- Syntax | 1623
- Hierarchy Level | 1623
- Description | 1624
- Options | 1624
- Required Privilege Level | 1624
- Release Information | 1624

Syntax

```
aacl-statistics-profile profile-name {
    aacl-fields {
        field-name;
    }
    file filename;
    record-mode  (interim-active-only | interim-full);
    report-interval minutes;
}
```

Hierarchy Level

```
[edit services service-set service-set-name],
[edit system services local-policy-decision-function statistics]
```

Description

Create an AACL statistics profile, which configures the files to which statistics records are exported and the format that is exported.

Options

`file` *`filename`*—Name of the file to receive the statistics data output. Enclose the name within quotation marks. All files are placed in the directory **/var/stats/aacl**.

`profile-name`—Identifier for the profile.

`record-mode`—Record mode for the reporting interval; possible values are `interim-active-only`, which reports only statistics that have changed, or `interim-full`, which reports all available statistics.

`report-interval` *`minutes`*—Frequency at which statistics are recorded, in minutes.

- **Default:** 15 minutes
- **Range:** 5 through 1440 minutes

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

`record-mode` option introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Network Management and Monitoring Guide

Configuring Statistics Profiles

abstract-hop

IN THIS SECTION

- Syntax | 1625
- Hierarchy Level | 1625
- Description | 1625
- Options | 1626
- Required Privilege Level | 1626
- Release Information | 1626

Syntax

```
abstract-hop abstract-hop-name {
    constituent-list constituent-list-name (include-any-list | include-all-list | exclude-all-
list | exclude-any-list);
    operators (AND | OR);
}
```

Hierarchy Level

```
[edit logical-systems logical-systems-name protocols mpls],
[edit protocols mpls]
```

Description

Define router clusters or groups, similar to the sequence of real-hop constraints (strict or loose), as a sequence of abstract hops for setting up a label-switched path (LSP).

An abstract hop is a logical combination of the existing traffic engineering constraints, such as administrative groups, extended administrative groups, and Shared Risk Link Groups (SRLGs), along with the ordering property of real hops. As a result, when a sequence of abstract hops is used in a path constraint, ordering is achieved among the groups of routers that meet a logical combination of link or

node attributes called constituent attributes. A path can use a combination of real and abstract hops as constraints.

Options

abstract-hop-name	Name of the abstract hop that is a logical combination of the existing traffic engineering constraints, such as administrative groups, extended administrative groups, and SRLGs, along with the ordering property of real hops.
constituent-list *constituent-list-name*	Name of the predefined constituent list to be included in defining the abstract hop. A constituent list enables you to define a set of constituent attributes that is identified with a user-defined name.
include-any-list	Satisfy any one of the attributes specified in the constituent list.
include-all-list	Satisfy all of the attributes specified in the constituent list.
exclude-all-list	Satisfy none of the attributes specified in the constituent list.
exclude-any-list	Fail to satisfy any one of the attributes specified in the constituent list.
operators	Specify the operation between constituent lists when more than one constituent list is included in the abstract hop definition.

AND Satisfy all the constituent lists referenced in the abstract hop definition for the attached node to be a member of the abstract hop.

OR Satisy at least one of the constituent lists referenced in the abstract hop definition for the attached node to be a member of the abstract hop.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Example: Configuring Abstract Hops for MPLS LSPs

constituent-list | 4571

show mpls abstract-hop-membership | 22125

show mpls lsp abstract-computation | 22203

accept (Router Advertisement Guard Policy)

IN THIS SECTION

- Syntax | 1627
- Hierarchy Level | 1628
- Description | 1628
- Required Privilege Level | 1628
- Release Information | 1628

Syntax

```
accept {
    match-list {
        match-criteria {
            (match-all | match-any);
        }
        prefix-list-name prefix-list-name;
        source-ip-address-list address-list-name;
        source-mac-address-list address-list-name;
    }
    match-option {
        hop-limit {
            (maximum | minimum) value;
        }
        managed-config-flag;
        other-config-flag;
        router-preference (high | low | medium);
```

```
    }
}
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard policy policy-name]
```

Description

Configure the accept policy for an IPv6 Router Advertisement (RA) guard policy. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.

When RA guard is enabled, the switch compares the information contained in the attributes of RA messages to the information contained in the policy. If RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped.

The criteria are configured either as one or more lists of source address or address prefixes, which are associated with the accept policy by using the `match-list` statement, or match condition parameters, which are associated with the accept policy by using the `match-option` statement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

accept (Dynamic Profile)

IN THIS SECTION

- Syntax | 1629
- Hierarchy Level | 1629
- Description | 1629
- Options | 1629
- Required Privilege Level | 1630
- Release Information | 1630

Syntax

```
accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure stacked-vlan-ranges dynamic-profile profile-name],
[edit interfaces interface-name auto-configure vlan-ranges dynamic-profile profile-name]
```

Description

Specify the type of VLAN Ethernet packet accepted by an interface that is associated with a VLAN dynamic profile or stacked VLAN dynamic profile.

Options

`any`—Any packet type. Specifies that any incoming packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes. This option is used when configuring wholesaling in a Layer 2 network.

`dhcp-v4`—IPv4 DHCP packet type. Specifies that incoming IPv4 DHCP discover packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes

> **NOTE**: The DHCP-specific `mac-address` and `option-82` options are rejected if the `accept` statement is not set to `dhcp-v4`.

`dhcp-v6`—IPv6 DHCP packet type. Specifies that incoming IPv6 DHCP discover packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes.

`inet`—IPv4 Ethernet and ARP packet type.

`inet6`—IPv6 Ethernet packet type.

`pppoe`—Point-to-Point Protocol over Ethernet packet type.

> **NOTE**: The `pppoe` VLAN Ethernet packet type option is supported only for MPC/MIC interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

`dhcp-v4` option added in Junos OS Release 10.0.

`dhcp-v6`, `inet6` and `pppoe` options added in Junos OS Release 10.2.

`any` option added in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

Configuring VLAN Interfaces for the Layer 2 Wholesale Solution

Configuring Subscriber Packet Types to Trigger VLAN Authentication

accept-data

IN THIS SECTION

- Syntax | 1631
- Hierarchy Level | 1631
- Description | 1631
- Default | 1632
- Required Privilege Level | 1632
- Release Information | 1632

Syntax

```
(accept-data | no-accept-data);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
```

Description

In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a router that is acting as the primary router accepts all packets destined for the virtual IP address.

- `accept-data`—Enable the primary router to accept all packets destined for the virtual IP address.
- `no-accept-data`—Prevent the primary router from accepting packets other than the ARP packets destined for the virtual IP address.

Default

If the router acting as the primary router is the IP address owner or has its priority set to 255, the primary router, by default, responds to all packets sent to the virtual IP address. However, if the router acting as the primary router does not own the IP address or has its priority set to a value less than 255, the primary router responds only to ARP requests.

NOTE:

- If you want to restrict the incoming IP packets to ICMP packets only, you must configure firewall filters to accept only ICMP packets.
- If you include the `accept-data` statement, your routing platform configuration does not comply with RFC 3768 (see section 6.4.3 of RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group

accept-join-always-from

IN THIS SECTION

Syntax | 1633
Hierarchy Level | 1633

- Description | 1633
- Required Privilege Level | 1634
- Release Information | 1634

Syntax

```
accept-join-always-from [policy-name ... ];
```

Hierarchy Level

```
[edit protocols interface name],
[edit routing-instances name protocols pim interface name]
```

Description

Configure an interface to accept PIM join or prune messages based on a specified policy.

A device might not accept PIM joins if the upstream neighbor address in the join packet is not that device's address. This option enables the device to accept and to install PIM (S,G) join states that the device would not normally accept.

Configure one or more routing policies that specify the upstream neighbor addresses for which the interface should always accept PIM joins. Then you can configure this statement with those policy options. Use the `prefix-list` and `policy-statement` options at the `[edit policy-options]` hierarchy to configure a routing policy.

For example, the following configuration creates a policy named `ajaf-101` that always accepts join messages from upstream neighbors in a prefix list named `rtr-101`:

```
set routing-instances RI protocols pim interface irb.101 accept-join-always-from ajaf-101
set policy-options prefix-list rtr-101 192.168.101.1/32
```

```
set policy-options policy-statement ajaf-101 term term1 from prefix-list rtr-101
set policy-options policy-statement ajaf-101 term term1 then accept
set policy-options policy-statement ajaf-101 then reject
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Optimized Inter-Subnet Multicast in EVPN Networks

accept-max-tcp-connections (System Process)

IN THIS SECTION

- Syntax | 1634
- Hierarchy Level | 1635
- Description | 1635
- Options | 1635
- Required Privilege Level | 1635
- Release Information | 1635

Syntax

```
accept-max-tcp-connections max-tcp-connections;
```

Hierarchy Level

```
[edit system processes dhcp-service]
```

Description

Specify the number of simultaneous TCP connections that DHCP can use for bulk leasequery operations. This chassis-wide to enables the jdhcpd process to avoid exhaustion of TCP resources across the chassis.

> **NOTE**: Use the `max-connections` option of the "allow-bulk-leasequery" on page 2165 statement to set the number of TCP connections allowed in the local server's logical system/routing instance.

Options

max-tcp-connections	Number of connections. • **Range:** 1 through 10 • **Default:** 5

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Configuring and Using DHCP Individual Leasequery*
>
> *Configuring and Using DHCP Bulk Leasequery*

accept-out-of-band

IN THIS SECTION

- Syntax | 1636
- Hierarchy Level | 1636
- Description | 1636
- Options | 1636
- Required Privilege Level | 1637
- Release Information | 1637

Syntax

```
accept-out-of-band protocol;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges dynamic-profile profile-name]
```

Description

Configure the protocol for which packets are accepted as out-of-band traffic to trigger instantiation or deletion of autosensed dynamic VLANs.

NOTE: A given physical interface can support VLANs created by either conventional packet-triggering or out-of-band triggering, but not both at the same time.

Options

protocol Out-of-band protocol. The following out-of-band protocol is supported:

- `ancp`—ANCP Port Up and Port Down messages trigger instantiation and deletion, respectively, of autosensed, dynamic VLAN.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

accept-own

IN THIS SECTION

- Syntax | 1637
- Hierarchy Level | 1638
- Description | 1638
- Required Privilege Level | 1638
- Release Information | 1638

Syntax

```
accept-own;
```

Hierarchy Level

```
[edit protocols bgp family inet-vpn unicast]
```

```
[edit protocols bgp family inet6-vpn unicast]
```

```
[edit protocols bgp group group family inetvpn unicast]
```

```
[edit protocols bgp group group family inet6-vpn unicast]
```

```
[edit protocols bgp group group neighbor address family inet-vpn unicast]
```

```
[edit protocols bgp group group neighbor address family inet6-vpn unicast]
```

Description

Enable processing of routes with own originator-id or nexthop.

NOTE: We support the `accept-own` configuration only for `inet-vpn unicast` and `inet6-vpn unicast` address families.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1

accept-prpd-connection

IN THIS SECTION

- Syntax | 1639
- Hierarchy Level | 1639
- Description | 1639
- Required Privilege Level | 1639
- Release Information | 1639

Syntax

```
accept-prpd-connection;
```

Hierarchy Level

```
[edit protocols bgp]
```

Description

Used to initialize BGP to accept PRPD connections even if there are no BGP peers configured.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1.

RELATED DOCUMENTATION

Controller-Based BGP Multicast Signaling

bgpmcast | 2968

inet-bgpmcast | 8557

inet6-bgpmcast | 8599

show route table | 23789

show multicast route | 22306

accept-remote-nexthop

IN THIS SECTION

- Syntax | 1640
- Hierarchy Level | 1640
- Description | 1641
- Required Privilege Level | 1641
- Release Information | 1641

Syntax

```
accept-remote-nexthop;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
```

```
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Specify that a single-hop EBGP peer accepts a remote next hop with which it does not share a common subnet. Configure a separate import policy on the EBGP peer to specify the remote next hop.

For Junos OS Release 13.3 and later releases, specify that a multihop EBGP peer accepts a remote next hop with which it does not share a common subnet. This allows working around current resolver limitations to realize multipath forwarding in recursive next-hop resolution scenarios.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring Single-Hop EBGP Peers to Accept Remote Next Hops

Configuring Routing Policies to Control BGP Route Advertisements

multipath (Protocols BGP) | **11110**

accept-remote-source

IN THIS SECTION

- Syntax | 1642
- Hierarchy Level | 1642
- Description | 1642
- Required Privilege Level | 1643
- Release Information | 1644

Syntax

```
accept-remote-source;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface interface-name],
[edit protocols pim interface interface-name],
[edit routing-instances routing-instance-name protocols pim interface interface-name]
```

Description

You can configure an incoming interface to accept multicast traffic from a remote source. A remote source is a source that is not on the same subnet as the incoming interface. Figure 1 on page 1643 shows just such a topology – R2 connects to the R1 source on one subnet, and to the incoming interface on R3 on another subnet (ge-1/3/0.0 in the figure).

Figure 1: Accepting Multicast Traffic from a Remote Source



In this topology R2 is a pass-through device not running PIM, so R3 is the first hop router for multicast packets sent from R1. Because R1 and R3 are in different subnets, the default behavior of R3 is to disregard R1 as a remote source. You can have R3 accept multicast traffic from R1, however, by enabling `accept-remote-source` on the target interface.

```
[edit protocols pim interface ge-1/3/0.0]
user@host# set accept-remote-source
```

NOTE: If the interface you identified is not the only path from the remote source, be sure it is the best path. For example you can configure a static route on the receiver side PE router to the source, or you can prepend the AS path on the other possible routes. That said, do not use `accept-remote-source` to receive multicast traffic over multiple upstream interfaces, as this use case for the command is not supported.

```
[edit policy-options policy-statement as-path-prepend term prepend]
user@host# set from route-filter 192.168.0.0/16 orlonger
user@host# set from route-filter 172.16.0.0/16 orlonger
user@host# set then as-path-prepend "1 1 1 1"
```

Commit the configuration changes, and then to confirm that the interface you configured is accepting traffic from the remote source, run the following command:

```
user@host# show pim statistics
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Allowing MBGP MVPN Remote Sources

Example: Allowing MBGP MVPN Remote Sources

Understanding Prepending AS Numbers to BGP AS Paths

accept-sdr (PCRF Partition)

IN THIS SECTION

- Syntax | 1644
- Hierarchy Level | 1644
- Description | 1645
- Required Privilege Level | 1645
- Release Information | 1645

Syntax

```
accept-sdr;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the service discovery requests for the vendor. If `accept-sdr` is configured for PCRF partition, the application and the vendor ID is advertised in the diameter CER request.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

- *Configuring the PCRF Partition*
- *Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers*
- *Understanding Gx Interactions Between the Router and the PCRF*
- *Understanding Interactions Between the PCRF, PCEF, and OCS*
- *Understanding Upstream and Downstream Messages for the PCRF*
- *3GPP Policy and Charging Control for Wireline Provisioning and Accounting*

accept-source-mac (SRX)

IN THIS SECTION

- Syntax | 1646
- Hierarchy Level | 1646
- Description | 1646
- Options | 1646
- Required Privilege Level | 1646
- Release Information | 1646

Syntax

```
accept-source-mac {
    mac-address mac-address;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

For Gigabit Ethernet (GE), Fast Ethernet (FE), or 10 Gigabit Ethernet (XE) interfaces, specify the MAC addresses from which the interface can receive packets. Ensure that you update the MAC address if the remote Ethernet card is replaced. Replacing the interface card changes the MAC address. If you do not update the MAC address, the interface cannot receive packets from the new card.

NOTE:

- Software-based MAC limiting is supported on SRX300, SRX320, and SRX340 devices. A maximum of 32 MAC addresses is supported per device.

Options

mac-address —MAC address filter. You can specify the MAC address as six hexadecimal bytes in one of the following formats: *nn:nn:nn:nn:nn:nn* (for example, `00:11:22:33:44:55`) or *nnnn:nnnn:nnnn* (for example, `0011.2233.4455`). You can configure up to 32 source addresses. To specify more than one address, include multiple *mac-addresses* in the `source-address-filter` statement.

Required Privilege Level

interface—To view this statement in the configuration..

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

accept-source-mac

IN THIS SECTION

- Syntax | 1647
- Hierarchy Level | 1647
- Description | 1648
- Required Privilege Level | 1648
- Release Information | 1648

Syntax

```
accept-source-mac {
    mac-address mac-address {
        policer {
            input cos-policer-name;
            output cos-policer-name;
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For Gigabit Ethernet intelligent queuing (IQ) interfaces only or for Gigabit Ethernet and aggregated Ethernet interfaces on switches, accept traffic from and to the specified remote media access control (MAC) address.

The `accept-source-mac` statement is equivalent to the `source-address-filter` statement, which is valid for Fast Ethernet, and Gigabit Ethernet interfaces only. To allow the interface to receive packets from specific MAC addresses, include the `accept-source-mac` statement.

On untagged Gigabit Ethernet interfaces, you should not configure the `source-address-filter` statement and the `accept-source-mac` statement simultaneously. On tagged Gigabit Ethernet interfaces, you should not configure the `source-address-filter` statement and the `accept-source-mac` statement with an identical MAC address specified in both filters.

The remaining statements are explained separately. See CLI Explorer.

NOTE: The `policer` statement is not supported on PTX Series Packet Transport Routers.

NOTE: If you configure MAC filtering on the AE interface, you must configure the interface with `family ethernet-switching`.

NOTE: On QFX platforms, if you configure source MAC addresses for an interface using the "static-mac" on page 28549 or "persistent-learning" on page 12451 statements and later configure a different MAC address for the same interface using the `accept-source-mac` statement, the MAC addresses that you previously configured for the interface remain in the ethernet-switching table and can still be used to send packets to the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

Configuring MAC Address Filtering on PTX Series Packet Transport Routers

source-filtering | **28208**

accept-srv6-service

IN THIS SECTION

- Syntax | **1649**
- Hierarchy Level | **1649**
- Description | **1651**
- Required Privilege Level | **1651**
- Release Information | **1651**

Syntax

```
accept-srv6-service;
```

Hierarchy Level

```
[edit fabric protocols bgp family inet unicast],
[edit fabric protocols bgp family inet6 unicast],
[edit fabric protocols bgp group name family inet unicast],
[edit fabric protocols bgp group name family inet6 unicast],
[edit fabric protocols bgp group name neighbor name family inet unicast],
[edit fabric protocols bgp group name neighbor name family inet6 unicast],
[edit logical-systems name protocols bgp family inet unicast],
[edit logical-systems name protocols bgp family inet6 unicast],
[edit logical-systems name protocols bgp group name family inet unicast],
[edit logical-systems name protocols bgp group name family inet6 unicast],
[edit logical-systems name protocols bgp group name neighbor name family inet unicast],
[edit logical-systems name protocols bgp group name neighbor name family inet6 unicast],
```

```
[edit logical-systems name routing-instances name protocols bgp family inet unicast],
[edit logical-systems name routing-instances name protocols bgp family inet6 unicast],
[edit logical-systems name routing-instances name protocols bgp group name family inet unicast],
[edit logical-systems name routing-instances name protocols bgp group name family inet6 unicast],
[edit logical-systems name routing-instances name protocols bgp group name neighbor name family inet
unicast],
[edit logical-systems name routing-instances name protocols bgp group name neighbor name family inet6
unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp family inet unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp family inet6 unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp group name family
inet unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp group name family
inet6 unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp group name neighbor
name family inet unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp group name neighbor
name family inet6 unicast],
[edit protocols bgp family inet unicast],
[edit protocols bgp family inet6 unicast],
[edit protocols bgp group name family inet unicast],
[edit protocols bgp group name family inet6 unicast],
[edit protocols bgp group name neighbor name family inet unicast],
[edit protocols bgp group name neighbor name family inet6 unicast],
[edit routing-instances name protocols bgp family inet unicast],
[edit routing-instances name protocols bgp family inet6 unicast],
[edit routing-instances name protocols bgp group name family inet unicast],
[edit routing-instances name protocols bgp group name family inet6 unicast],
[edit routing-instances name protocols bgp group name neighbor name family inet unicast],
[edit routing-instances name protocols bgp group name neighbor name family inet6 unicast],
[edit tenants name routing-instances name protocols bgp family inet unicast],
[edit tenants name routing-instances name protocols bgp family inet6 unicast],
[edit tenants name routing-instances name protocols bgp group name family inet unicast],
[edit tenants name routing-instances name protocols bgp group name family inet6 unicast],
[edit tenants name routing-instances name protocols bgp group name neighbor name family inet
unicast],
[edit tenants name routing-instances name protocols bgp group name neighbor name family inet6
unicast]
```

Description

Configure the BGP ingress node to accept SRv6 service when you enable BGP services over the SRv6 data plane. Egress provider edge (PE) devices that support SRv6-based Layer 3 services advertise overlay service prefixes along with a service SID. BGP ingress nodes receive the advertisements from the egress provider edge (PE) device and then add the prefix to the corresponding VRF RIB also known as routing information base.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

srv6 (Protocols BGP) | 28417

advertise-srv6-service | 2024

Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP

accepted-prefix-limit

IN THIS SECTION

- Syntax | 1652
- Hierarchy Level | 1652
- Description | 1653
- Options | 1653
- Required Privilege Level | 1654
- Release Information | 1654

Syntax

```
accepted-prefix-limit {
    maximum number;
    teardown <percentage-threshold> idle-timeout (forever | minutes);
    drop-excess <percentage>;
    hide-excess <percentage>;
  }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family (inet | inet6) (any | flow |
labeled-unicast | multicast | unicast)],
[edit logical-systems logical-system-name protocols bgp family route-target],
[edit logical-systems logical-system-name protocols bgp group group-name family (inet | inet6)
(any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems logical-system-name protocols bgp group group-name family route-target],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
(inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
route-target],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family route-target],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family route-target],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family (inet | inet6) (any | flow | labeled-unicast |
multicast | unicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family route-target],
[edit protocols bgp family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit protocols bgp family accepted-prefix-limit (inet | inet6)(any | flow | labeled-unicast |
multicast | unicast)],
[edit protocols bgp family route-target],
[edit protocols bgp group group-name family (inet | inet6) (any | flow | labeled-unicast |
multicast | unicast)],
```

```
[edit protocols bgp group group-name family route-target],
[edit protocols bgp group group-name neighbor address family (inet | inet6) (any | flow |
labeled-unicast | multicast | unicast)],
[edit protocols bgp group group-name neighbor address family route-target],
[edit routing-instances routing-instance-name protocols bgp family (inet | inet6) (any | flow |
labeled-unicast | multicast | unicast)],
[edit routing-instances routing-instance-name protocols bgp family route-target],
[edit routing-instances routing-instance-name protocols bgp group group-name family (inet |
inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit routing-instances routing-instance-name protocols bgp group group-name family route-
target],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family route-target]
```

Description

Configure a limit to the number of prefixes that can be accepted in a BGP peer session. When that limit is exceeded, a system log message is sent.

This statement provides the ability to log a message, reset the BGP session, or do both when the number of prefixes received from the peer and accepted by policy exceeds a preset limit. This functionality is identical to the `prefix-limit` functionality except that it operates against accepted prefixes rather than received prefixes.

Options

maximum *number*

When you set the maximum number of prefixes, a message is logged when that number is exceeded.

- **Range:** 1 through 4,294,967,295 ($2^{32} - 1$)

teardown <*percentage*>

(Optional) If you include the `teardown` statement, the session is torn down when the maximum number of prefixes is reached. If you specify a percentage, messages are logged when the number of prefixes exceeds that percentage. After the session is torn down, it is reestablished in a short time unless you include the `idle-timeout` statement. Then the session can be kept down for a specified amount of time, or forever. If you specify `forever`, the session is reestablished only after you issue a `clear bgp neighbor` command.

If the `teardown` statement is not configured, a message is logged when the number of prefixes exceeds the value configured for the `maximum` option.

- **Range:** 1 through 100

`idle-timeout (forever | timeout-in-minutes)`

(Optional) If you include the `idle-timeout` statement, the session is torn down for a specified amount of time, or forever. If you specify a period of time, the session is allowed to reestablish after this timeout period. If you specify `forever`, the session is reestablished only after you intervene with a `clear bgp neighbor` command.

- **Range:** 1 through 2400

`drop-excess <percentage>`

(Optional) If you include the `drop-excess <percentage>` option, the excess routes are dropped when the maximum number of prefixes is reached. If you specify a percentage, the routes are logged when the number of prefixes exceeds that percentage value of the maximum number.

`hide-excess <percentage>`

(Optional) If you include the `hide-excess <percentage>` option, the excess routes are hidden when the maximum number of prefixes is reached. If you specify a percentage, the routes are logged when the number of prefixes exceeds that percentage value of the maximum number. If the percentage is modified, the routes are re-evaluated automatically. If the active routes drop below the specified percentage, those routes are kept as hidden.

To recover the routes from the hidden state, run the following commands to refresh the routes and get a complete re-evaluation of the routes:

`clear bgp table <table-name>` and `clear bgp neighbor <neighbor address> soft-inbound`

> **NOTE**: You can include either the `teardown <percentage>`, `drop-excess <percentage>`, or `hide-excess <percentage>` statement option one at a time.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

`drop-excess <percentage>` and `hide-excess <percentage>` options are introduced in Junos OS Release 21.2.

`drop-excess` <*percentage*> and `hide-excess` <*percentage*> options are introduced in Junos OS Evolved Release 21.3R1.

RELATED DOCUMENTATION

prefix-limit | 13087

Understanding Multiprotocol BGP

show bgp neighbor | 17175

access-concentrator

IN THIS SECTION

- Syntax | 1655
- Hierarchy Level | 1655
- Description | 1656
- Options | 1656
- Required Privilege Level | 1656
- Release Information | 1656

Syntax

```
access-concentrator name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
pppoe],
[edit interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-options],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options],
```

```
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
pppoe-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
pppoe-underlying-options]
```

Description

Configure an alternative access concentrator name in the AC-NAME tag in a PPPoE control packet for use with a dynamic PPPoE subscriber interface. If you do not configure the access concentrator name, the AC-NAME tag contains the system name.

NOTE: The `[edit ... family pppoe]` hierarchies are supported only on MX Series routers with MPCs.

Options

name—Name of the access concentrator.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options]` and `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number pppoe-underlying-options]` hierarchy levels introduced in Junos OS Release 10.1.

Support at the `[edit ... family pppoe]` hierarchies introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Identifying the Access Concentrator

Configuring the PPPoE Family for an Underlying Interface

Configuring Dynamic PPPoE Subscriber Interfaces

PPPoE Overview

access-domain (Remote Device Management)

IN THIS SECTION

- Syntax | 1657
- Hierarchy Level | 1657
- Description | 1657
- Options | 1658
- Required Privilege Level | 1659
- Release Information | 1659

Syntax

```
access-domain {
    vlan-id-list [vlan-id-low-vlan-id-high vlan-id]
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name system services remote-device-management service-device device-name],
[edit system services remote-device-management service-device device-name]
```

Description

Specify the Layer 2 domain that is served by the service device. The access domain represents the set of subscriber-facing Layer 2 locations that map to the device. You define the domain when you map subscribers to the device by configuring one or more non-overlapping outer VLAN ID ranges or discrete VLAN IDs that are served by the remote device. All BNG subscriber sessions with outer VLAN tags

matching this list are connected to this remote device in the access network. During subscriber negotiation and provisioning, the BNG provisions remote services sourced by the policy server (PCRF) on this remote device.

You can configure the VLAN IDs as discrete values, even if they are consecutive; however, it is more practical to configure ranges of IDs with additional discrete VLAN IDs as needed. You can configure any number of ranges or discrete IDs; no limit is enforced.

The list of VLAN IDs can overlap across remote devices to support redundancy in the access network.

NOTE: To specify more than one range or address in the list, place the values inside square brackets, separated by spaces. For example:

```
[100 102 110 120-130]
```

NOTE: You do not have to specify an access domain (list of VLAN IDs) when you configure a service device. This means that a service device without the access domain can be connected and dynamically assigned an IP address. Then at some later point when the device needs to support subscriber connections, you can configure one or more VLAN IDs or ranges with the `vlan-id-list` statement.

NOTE: You can add new VLAN Ids or ranges to the access domain at any time, even when the device has active subscriber services mapped to it. You cannot delete IDs or ranges when at least one active subscriber service is configured for that discrete VLAN ID or any VLAN ID in the range being deleted.

You can use the "request services remote-device-management reconfigure service-device" on page 14961 command to reconfigure a remote device to provision all active subscriber services matching the access domain.

Options

vlan-id-low	(Optional) Lowest VLAN ID in the range,
vlan-id-high	(Optional) Highest VLAN ID in the range.
vlan-id	(Optional) Discrete VLAN ID that is not part of a range.

Required Privilege Level

system—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Configuring Remote Device Management for Service Provisioning

Remote Device Services Manager (RDSM) Overview

access (Dynamic Access Routes)

IN THIS SECTION

- Syntax | 1659
- Hierarchy Level | 1660
- Description | 1660
- Required Privilege Level | 1660
- Release Information | 1660

Syntax

```
access {
    route prefix {
        next-hop next-hop;
        metric route-cost;
        preference route-distance;
        tag route-tag;
        tag2 route-tag2;
    }
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options],
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
rib routing-table-name],
[edit dynamic-profiles  profile-name routing-options]
```

Description

Dynamically configure access routes in a dynamic client profile.

> **NOTE**: Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support. We recommend that you do not use access-internal routes. If you configure the `access-internal` statement in the dynamic profile, it is ignored. The subscriber's address is stored in the session database entry before the dynamic profile installs the framed route, enabling the next-hop address to be resolved when it is not explicitly specified in the Framed-Route RADIUS attribute [22] or Framed-IPv6-Route attribute [99].

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options]` and `[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options rib routing-table-name]` hierarchy levels introduced in Junos OS Release 10.1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1	Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support.

RELATED DOCUMENTATION

Configuring Dynamic Access Routes for Subscriber Management

access-identifier

IN THIS SECTION

- Syntax | 1661
- Hierarchy Level | 1661
- Description | 1662
- Options | 1662
- Required Privilege Level | 1662
- Release Information | 1662

Syntax

```
access-identifier identifier-string;
```

Hierarchy Level

```
[edit protocols ancp interfaces interface-name],
[edit protocols ancp interfaces interface-set]
```

Description

Associate an access-loop circuit identifier (ACI) with the VLAN or set of VLANs that carry traffic to the subscriber using that access loop; identify a particular subscriber. This statement requires that the name of the interface or interface set is statically configured or deterministic. This means that it can be used with dynamic or static interface sets, VLAN-tagged interface sets, or static VLAN/VLAN demux interfaces.

Options

identifier-string—Unique identifier string for the access loop circuit; also configured on the access node.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent

Associating an Access Node with Subscribers for ANCP Agent Operations

access-interfaces (IFL Subscriber)

IN THIS SECTION

- Syntax | 1663
- Hierarchy Level | 1663
- Description | 1663
- Options | 1663

Required Privilege Level | 1664
Release Information | 1664

Syntax

```
access-interfaces [interface-name];
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ifl-subscriber subscriber-name]
```

Description

Specify one or more interfaces that carry traffic for the subscriber.

Options

interface-name Name of the interface. You can assign only one IFL-based subscriber to an interface. You can specify the following types of interfaces:

- Physical Layer 3 Ethernet interface
- Layer 3 Aggregated Ethernet interface
- IRB interface
- IRB that contains Ether-channel and physical interface members
- Logical Tunnel interface

NOTE: The interfaces and the TDF interface (mif) in the TDF domain must be included in the same VRF routing instance.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain

Understanding the Definition of a Set of IFL-Based Subscriber Properties with a TDF Domain

access-interfaces (IP Subscriber)

IN THIS SECTION

- Syntax | 1664
- Hierarchy Level | 1665
- Description | 1665
- Options | 1665
- Required Privilege Level | 1665
- Release Information | 1665

Syntax

```
access-interfaces [interface-name];
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
```

Description

Specify at least one interface that faces the access network and that carries traffic for the TDF domain for IP-based subscribers. You can specify multiple interfaces by including the `access-interfaces` statement multiple times.

Options

interface-name	Name of the interface.

> **NOTE**: The access-facing interface and the TDF interface (mif) in the TDF domain must be included in the same VRF routing instance.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain*
- *Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain*
- *IP-Based Subscriber Setup Overview*

access-internal (Dynamic Access-Internal Routes)

IN THIS SECTION

- Syntax | 1666
- Hierarchy Level | 1666
- Description | 1666
- Required Privilege Level | 1667
- Release Information | 1667

Syntax

```
access-internal {
    route subscriber-ip-address {
        qualified-next-hop underlying-interface {
            mac-address address;
        }
    }
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options],
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
rib routing-table-name],
[edit dynamic-profiles routing-options]
```

Description

(Releases earlier than Junos OS Release 15.1) Dynamically configure access-internal routes in a dynamic client profile. Access-internal routes are optional, but are used instead of access routes if the next-hop address is not specified in the Framed-Route Attribute [22] for IPv4 or the Framed-IPv6-Route attribute [99] for IPv6.

> **NOTE**: Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support. We recommend that you do not use access-internal routes. If you configure the `access-internal` statement in the dynamic profile, it is ignored. The subscriber's address is stored in the session database entry before the dynamic profile installs the framed route, enabling the next-hop address to be resolved when it is not explicitly specified in the Framed-Route RADIUS attribute (22) or Framed-IPv6-Route attribute [99].

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles` *`profile-name`* `routing-instances $junos-routing-instance routing-options]` and `[edit dynamic-profiles` *`profile-name`* `routing-instances $junos-routing-instance routing-options rib` *`routing-table-name`*`]` hierarchy levels introduced in Junos OS Release 10.1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1	Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support.

RELATED DOCUMENTATION

Configuring Dynamic Access-Internal Routes for DHCP and PPP Subscribers

access-internal (Static Access-Internal Routes)

IN THIS SECTION

- Syntax | 1668
- Hierarchy Level | 1668
- Description | 1668
- Required Privilege Level | 1669
- Release Information | 1669

Syntax

```
access-internal {
    route  ip-prefix</prefix-length> {
        next-hop next-hop;
        qualified-next-hop next-hop
    }
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Configure parameters for internal access routes.

The remaining statements are explained separately.

NOTE: Starting in Junos OS Release 15.1R4, the router no longer supports a configuration where a static route points to a next hop that is tied to a subscriber. Typically, this might occur when RADIUS assigns the next hop with the Framed-IP-Address attribute. An alternative to this misconfiguration is to have the RADIUS server provide a Framed-Route attribute that matches the static route.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring Dynamic Access-Internal Routes for DHCP and PPP Subscribers

access-line (Access-Line Rate Adjustment)

IN THIS SECTION

- Syntax for Releases Earlier than Junos OS Release 19.3R1 | **1670**
- Syntax for Junos OS Release 19.3R1 and Higher Releases. | **1671**
- Hierarchy Level | **1673**
- Description | **1673**
- Options | **1674**
- Required Privilege Level | **1683**
- Release Information | **1683**

Syntax for Releases Earlier than Junos OS Release 19.3R1

```
access-line {
    adsl-overhead-bytes bytes;
    adsl-total-adjust percentage;
    adsl2-overhead-bytes bytes;
    adsl2-total-adjust percentage;
    adsl2-plus-overhead-bytes bytes;
    adsl2-plus-total-adjust percentage;
    gfast-bonded-overhead-adjust percentage
    gfast-bonded-overhead-bytes bytes
    gfast-bonded-total-adjust percentage
    gfast-overhead-adjust percentage
    gfast-overhead-bytes bytes
    gfast-total-adjust percentage
    hierarchical-access-network-detection;
    other-overhead-adjust percentage;
    other-overhead-bytes bytes;
    other-total-adjust percentage;
    sdsl-bonded-overhead-bytes  bytes
    sdsl-bonded-overhead-adjust  percentage
    sdsl-bonded-total-adjust  percentage
    sdsl-overhead-adjust percentage;
    sdsl-overhead-bytes bytes;
    sdsl-total-adjust percentage;
    vdsl-overhead-adjust percentage;
    vdsl-overhead-bytes bytes;
    vdsl-total-adjust percentage;
    vdsl2-annex-q-bonded-overhead-adjust  percentage
    vdsl2-annex-q-bonded-overhead-bytes bytes
    vdsl2-annex-q-bonded-total-adjust  percentage
    vdsl2-annex-q-overhead-adjust  percentage
    vdsl2-annex-q-overhead-bytes  bytes
    vdsl2-annex-q-total-adjust  percentage
    vdsl2-bonded-overhead-adjust  percentage
    vdsl2-bonded-overhead-bytes  bytes
    vdsl2-bonded-total-adjust  percentage
    vdsl2-overhead-adjust percentage;
    vdsl2-overhead-bytes bytes;
    vdsl2-total-adjust percentage;
}
```

Syntax for Junos OS Release 19.3R1 and Higher Releases.

```
access-line {
    attributes {
        preference (dsl | pon);
    }
    dsl {
        adsl {
            overhead-bytes bytes;
            total-adjust percent;
        }
        adsl2 {
            overhead-bytes bytes;
            total-adjust percent;
        }
        adsl2-plus {
            overhead-bytes bytes;
            total-adjust percent;
        }
        gfast {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        gfast-bonded {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        other {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        sdsl {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        sdsl-bonded {
            overhead-adjust percent;
            overhead-bytes bytes;
```

```
            total-adjust percent;
        }
        type tlv-value {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        vdsl {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        vdsl2 {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        vdsl2-annex-q {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        vdsl2-annex-q-bonded {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
    }
    hierarchical-access-network-detection;
    pon {
        gpon {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        other {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        twdm-pon {
            overhead-adjust percent;
            overhead-bytes bytes;
```

```
            total-adjust percent;
        }
        type tlv-value {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        wdm-pon {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        xg-pon1 {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
        xgs-pon {
            overhead-adjust percent;
            overhead-bytes bytes;
            total-adjust percent;
        }
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure values to adjust data rates by a percentage of the actual data rate, or adjust encapsulation overhead by adding to or subtracting from the total cell or frame bytes a specified number of bytes. Depending on the value, it may be reported to AAA, CoS, or both.

The actual (unadjusted) downstream and upstream data rates, line type, and encapsulation mode are received from the access node by the ANCP agent in ANCP port messages, or by the PPPoE daemon from the PPPoE intermediate agent (PPPoE-IA) in PADI or PADR messages. The ANCP agent or PPPoE daemon subsequently adjusts rates and bytes based on the configuration.

Adjustments are applied to all subscribers using access lines of the specific subscriber access line type:

- Adjusted and unadjusted downstream and upstream rates are always reported to AAA in response to an AAA request.
- Adjusted and unadjusted downstream rates and overhead byte adjustments are reported to CoS, but only when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.
- Overhead byte adjustments are not reported to AAA.

AAA reports the adjusted values to the RADIUS server in the Access-Request and Accounting-Request messages through Juniper Networks VSAs 26-141, Downstream-Calculated-Qos-Rate Rate, and 26-142, Upstream-Calculated-Qos-Rate.

The ANCP agent reports these values to the LAC in an L2TP network. The LAC passes the rates to the LNS in the following messages and AVPs:

- AVP 24, Tx Connect Speed (ICCN message)
- AVP 38, Rx Connect Speed (ICCN message)
- AVP 97, Connect Speed Update (CSUN message)

NOTE: Starting in Junos OS Release 19.3R1, the pre-existing adjustment options were renamed and placed in the new `dsl` stanza. The old DSL options are deprecated, but they redirect to the new location. The `hierarchical-access-network-detection` option is unchanged.

BEST PRACTICE: We recommend that you update your scripts to use the "dsl" on page 6025 statement when you upgrade to Junos OS Release 19.3R1 or higher releases. The redirect function will be supported for only a limited time.

Options

adsl-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on an ADSL access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

NOTE: This option replaces the `adsl-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

adsl-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an ADSL access line. The adjusted rate is reported only to AAA.

> **NOTE**: This option replaces the `qos-adjust-adsl` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

adsl2-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on an ADSL2 access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `adsl2-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

adsl2-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an ADSL2 access line. The adjusted rate is reported only to AAA.

> **NOTE**: This option replaces the `qos-adjust-adsl2` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

adsl2-plus-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on an ADSL2+ access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `adsl2-plus-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

adsl2-plus-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an ADSL2+ access line. The adjusted rate is reported only to AAA.

> **NOTE**: This option replaces the `qos-adjust-adsl2-plus` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

gfast-bonded-overhead-adjust *percentage*

Adjustment factor in percent that is applied to the downstream and upstream bonded data overhead rates for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

gfast-bonded-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream cell bonded overhead for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure. Specify G.fast bonded value in bytes.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

gfast-bonded-total-adjust *percentage*

Adjustment factor in percent that is globally applied to the downstream and upstream bonded data rates for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

gfast-overhead-adjust *percentage*

Adjustment factor in percent that is applied to the downstream and upstream data overhead rates for all subscribers on a G.fast high speed DSL line connected to a PON tree infrastructure.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

gfast-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream cell overhead for all subscribers on a G.fast high speed DSL line connected to a PON tree infrastructure.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

gfast-total-adjust *percentage*

Adjustment factor in percent that is globally applied to the downstream and upstream data rates for all subscribers on a G.fast high speed bonded DSL line connected to a PON tree infrastructure.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

hierarchical-access-network-detection

Enable parsing of ANCP subscriber access loop attributes (TLVs) for backhaul line identifiers, to recognize when the access node references a logical interface set rather than an individual subscriber. If the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x03) string begins with a # sign, then the remainder of the string represents a logical intermediate node (DPU-C or PON tree) in the access network to which the subscriber is attached. The string is used as the name of a CoS Level 2 interface set that groups subscribers.

> **NOTE**: The Access-Loop-Remote-ID (TLV (0x02) is similarly parsed for the # character, but the resulting string is not used in the current release.

These TLVs can be parsed in ANCP messages or PPPoE IA tags in PADR messages.

- **Default:** Disabled, in case some users include an initial # character for some other purpose.

other-overhead-adjust *percentage*

Adjust the actual downstream rate for an access line of DSL type OTHER by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

The router reports some access technology types as DSL type OTHER. For example, when an OLT sends PON rates in DSL TLVs, the DSL type is set to OTHER.

> **NOTE**: This option replaces the `other-overhead-adjust` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

other-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on an access line of DSL type OTHER to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

The router reports some access technology types as DSL type OTHER. For example, when an OLT sends PON rates in DSL TLVs, the DSL type is set to OTHER.

> **NOTE**: This option replaces the `other-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

other-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an access line of DSL type OTHER. The adjusted rate is reported only to AAA.

The router reports some access technology types as DSL type OTHER. For example, when an OLT sends PON rates in DSL TLVs, the DSL type is set to OTHER.

> **NOTE**: This option replaces the `qos-adjust-other` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

sdsl-bonded-overhead-adjust *percentage*

Adjust the actual downstream rate for an SDSL bonded access line by multiplying the rate by the specified percentage.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

sdsl-bonded-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on an SDSL bonded access line to account for the traffic encapsulation overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

sdsl-bonded-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an SDSL bonded access line. This value is reported to AAA.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

sdsl-overhead-adjust *percentage*

Adjust the actual downstream rate for an SDSL access line by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `sdsl-overhead-adjust` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

sdsl-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on an SDSL access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `sdsl-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

sdsl-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on an SDSL access line. The adjusted rate is reported only to AAA.

> **NOTE**: This option replaces the `qos-adjust-sdsl` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent

- **Default:** 100 percent

vdsl-overhead-adjust *percentage*

Adjust the actual downstream rate for a VDSL access line by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `vdsl-overhead-adjust` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

vdsl-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `vdsl-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

vdsl-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL access line. The adjusted rate is reported only to AAA.

> **NOTE**: This option replaces the `qos-adjust-vdsl` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

vdsl2-annex-q-bonded-overhead-adjust *percentage*

Adjust the actual downstream rate for a VDSL2 annex q bonded access line by multiplying the rate by the specified percentage. The adjusted rate is reported to AAA.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

vdsl2-annex-q-bonded-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 annex q bonded access line to account for the traffic encapsulation overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

vdsl2-annex-q-bonded-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 annex q bonded access line. The adjusted rate is reported to AAA.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

vdsl2-annex-q-overhead-adjust *percentage*

Adjust the actual downstream rate for a VDSL2 annex q access line by multiplying the rate by the specified percentage. The adjusted rate is reported to AAA.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

vdsl2-annex-q-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 annex q access line to account for the traffic encapsulation overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

vdsl2-annex-q-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 annex q access line. The adjusted rate is reported to AAA.

vdsl2-bonded-overhead-adjust *percentage*

Adjust the actual downstream rate for a VDSL2 bonded access line by multiplying the rate by the specified percentage. The adjusted rate is reported to AAA.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

vdsl2-bonded-overhead-bytes*bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 bonded access line to account for the traffic encapsulation overhead.

- **Range:** -100 through 100 bytes

- **Default:** 0 bytes

vdsl2-bonded-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 bonded access line. The adjusted rate is reported to AAA.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

vdsl2-overhead-adjust *percentage*

Adjust the actual downstream rate for a VDSL2 access line by multiplying the rate by the specified percentage. The adjusted rate is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

> **NOTE**: This option replaces the `vdsl2-overhead-adjust` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

vdsl2-overhead-bytes *bytes*

Number of bytes added to or subtracted from the actual downstream frame overhead for all subscribers on a VDSL2 access line to account for the traffic encapsulation overhead. The adjusted value is reported to CoS when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.

Number of bytes added to or subtracted from the actual downstream frame overhead.

> **NOTE**: This option replaces the `vdsl2-bytes` statement at the `[edit protocols ancp qos-adjust]` hierarchy level.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

vdsl2-total-adjust *percentage*

Adjustment factor applied globally to the downstream and upstream data rates for all subscribers on a VDSL2 access line. The adjusted rate is reported only to AAA.

> **NOTE**: This option replaces the `qos-adjust-vdsl2` statement at the `[edit protocols ancp]` hierarchy level.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

The following options added in Junos OS 18.2R1:

`gfast-bonded-overhead-adjust`	`vdsl2-annex-q-bonded-overhead-adjust`
`gfast-bonded-overhead-bytes`	`vdsl2-annex-q-bonded-overhead-bytes`
`gfast-bonded-total-adjust`	`vdsl2-annex-q-bonded-total-adjust`
`gfast-overhead-adjust`	`vdsl2-annex-q-overhead-adjust`
`gfast-overhead-bytes`	`vdsl2-annex-q-overhead-bytes`
`gfast-total-adjust`	`vdsl2-annex-q-total-adjust`
`sdsl-bonded-overhead-bytes`	`vdsl2-bonded-overhead-bytes`
`sdsl-bonded-overhead-adjust`	`vdsl2-bonded-total-adjust`
`sdsl-bonded-total-adjust`	`vdsl2-bonded-total-adjust`

`hierarchical-access-network-detection` option added in Junos OS 18.4R1.

`attributes`, `dsl`, and `pon` statements added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS

Traffic Rate Reporting and Adjustment by the ANCP Agent

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates

Setting a Global Adjustment Factor per PON Subscriber Line for ANCP Agent-Reported Traffic Rates

Configuring the ANCP Agent

access-line-information (L2TP)

IN THIS SECTION

- Syntax | 1684
- Hierarchy Level | 1684
- Description | 1685
- Options | 1685
- Required Privilege Level | 1685
- Release Information | 1685

Syntax

```
access-line-information <connection-speed-update>;
```

Hierarchy Level

```
[edit services l2tp],
[edit services l2tp destination ip-address]
```

Description

Configure a LAC to forward subscriber line identification and other DSL attributes in ICRQ messages to the LNS by means of L2TP AVPs for all tunnels to all LNSs or for only tunnels with the specified endpoint for a particular LNS. Optionally, configure the LAC to send initial line rates in ICCN messages and subsequent rate updates in CSUN messages.

Configure an LNS to process such line information for all tunnels from all LACs or for only tunnels with the specified endpoint for a particular LAC. Optionally, configure the LNS to process rate updates received in CSUN messages from the LAC.

Including this statement at the `[edit services l2tp destination ip-address]` hierarchy level is useful when you know that some endpoints in the network do not support this feature or have an incorrect implementation. Configuring at this level enables you to restrict the transmission or processing of this information to only LACs and LNSs, respectively, that you know support the feature.

This statement has no effect on existing subscribers; it applies only to new subscribers.

Options

connection-speed-update (Optional) On the LAC, include the Connect Speed Update Enable AVP (98) in ICCN messages from the LAC to alert the LNS that the LAC might send CSUN messages that report speed changes originating with the ANCP agent.

On the LNS, enable processing of CSUN updates. If this option is not configured on the LNS, CSUN updates cannot be processed even when the Connect Speed Update Enable AVP (98) is received from the LAC. In that case, only rates received in AVP 24 (Tx speed) and AVP 38 (Rx speed) in ICCN messages can be applied.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Support at the `[edit services l2tp]` hierarchy level introduced in Junos OS Release 14.2.

Support for the LNS introduced in Junos OS Release 17.4R1 on MX Series routers.

RELATED DOCUMENTATION

L2TP Subscriber Access Lines and Connection Speeds

access-list

IN THIS SECTION

- Syntax | 1686
- Hierarchy Level | 1686
- Description | 1686
- Required Privilege Level | 1687
- Release Information | 1687

Syntax

```
[edit snmp]
    routing-instance-access {
        access-list {
            routing-instance;
            routing-instance restrict;
        }
    }
```

Hierarchy Level

```
[edit snmp routing-instance-access]
```

Description

Create access lists to control SNMP agents in routing instances from accessing SNMP information. To enable the SNMP agent on a routing instance to access SNMP information, specify the routing instance

name. To disable the SNMP agent on a routing instance from accessing SNMP information, include the routing-instance name followed by the **restrict** keyword.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

routing-instance-access | 15924

access-point name

IN THIS SECTION

- Syntax | 1687
- Hierarchy Level | 1688
- Description | 1688
- Options | 1688
- Required Privilege Level | 1688
- Release Information | 1688

Syntax

```
access-point-name apn;
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options profiles profile-name]
```

Description

Configure the access point name (APN) provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.

Options

apn—Access point name.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

access-profile

IN THIS SECTION

- Syntax | 1689
- Hierarchy Level | 1689
- Description | 1689
- Options | 1689
- Required Privilege Level | 1689
- Release Information | 1690

Syntax

```
access-profile name;
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options chap],
[edit interfaces interface-name ppp-options pap],
[edit interfaces interface-name unit logical-unit-number ppp-options chap],
[edit interfaces interface-name unit logical-unit-number ppp-options pap],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options chap],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options pap]
```

Description

For CHAP authentication, the mapping between peer names (or “clients”) and the secrets associated with their respective links. For PAP authentication, the peer's username and password.

For Asynchronous Transfer Mode 2 (ATM2) IQ interfaces only, you can configure a Challenge Handshake Authentication Protocol (CHAP) access profile on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- `atm-ppp-llc`—PPP over AAL5 logical link control (LLC) encapsulation.
- `atm-ppp-vc-mux`—PPP over AAL5 multiplex encapsulation.

Options

name—Name of the access profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for PAP added in Junos OS Release 8.3.

Support for VLAN and stacked VLAN ranges added in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring the PPP Challenge Handshake Authentication Protocol

Configuring the PPP Password Authentication Protocol On a Physical Interface

access-profile (AAA Options)

IN THIS SECTION

- Syntax | 1690
- Hierarchy Level | 1690
- Description | 1691
- Options | 1691
- Required Privilege Level | 1691
- Release Information | 1691

Syntax

```
access-profile profile-name;
```

Hierarchy Level

```
[edit access aaa-options aaa-options-name]
```

Description

Specify the name of the access profile that includes the username stripping configuration.

Options

profile-name Name of the subscriber access profile that includes a subscriber username stripping configuration.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access

Configuring Subscriber Session Timeout Options

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

access-profile (Domain Map)

IN THIS SECTION

- Syntax | 1692
- Hierarchy Level | 1692
- Description | 1692
- Options | 1692

Required Privilege Level | 1692
Release Information | 1692

Syntax

```
access-profile profile-name;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Access profile that defines the AAA services and options for subscribers associated with the domain map.

Options

profile-name—Name of access profile.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Specifying an Access Profile in a Domain Map

access-profile (Dynamic Stacked VLAN)

IN THIS SECTION

- Syntax | 1693
- Hierarchy Level | 1693
- Description | 1693
- Required Privilege Level | 1694
- Release Information | 1694

Syntax

```
access-profile svlan-access-profile-name;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure stacked-vlan-ranges dynamic-profile profile-name]
```

Description

Access profiles contain subscriber access authentication, authorization and accounting (AAA) configuration parameters. You can create an access profiles and then attach it at various configuration levels. When you attach an access profile to an interface configured for dynamic VLAN or stacked VLAN, all the VLANs and stacked VLANs use the same set of AAA parameters configured in that access profile. The different access profiles can have different authentication/authorization settings so you can, for example, have authentication on some VLAN and stacked VLAN ranges but no authentication on other ranges.

You can assign different access profiles to different dynamic profiles on the same interface. If you assign an access profile at the global level, but a different access profile is assigned at the interface level, the access profile at the interface level authenticates all dynamic VLANs and stacked VLANs on the interface. Access profiles can be assigned at various levels, but the most specific access profile takes precedence over any other profile assignments..

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

access-profile (Dynamic VLAN)

IN THIS SECTION

- Syntax | 1694
- Hierarchy Level | 1694
- Description | 1695
- Required Privilege Level | 1695
- Release Information | 1695

Syntax

```
access-profile vlan-access-profile-name;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges dynamic-profile profile-name]
```

Description

Access profiles contain subscriber access authentication, authorization and accounting (AAA) configuration parameters. You can create an access profiles and then attach it at various configuration levels. When you attach an access profile to an interface configured for dynamic VLAN or stacked VLAN, all the VLANs and stacked VLANs use the same set of AAA parameters configured in that access profile. The different access profiles can have different authentication/authorization settings so you can, for example, have authentication on some VLAN or stacked VLAN ranges but no authentication on other ranges.

You can assign different access profiles to different dynamic profiles on the same interface. If you assign an access profile at the global level, but a different access profile is assigned at the interface level, the access profile at the interface level authenticates all dynamic VLANs and stacked VLANs on the interface. Access profiles can be assigned at various levels, but the most specific access profile takes precedence over any other profile assignments.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

access-profile (Extensible Subscriber Services Manager)

IN THIS SECTION

- Syntax | **1696**
- Hierarchy Level | **1696**

- Description | 1696
- Options | 1696
- Required Privilege Level | 1696
- Release Information | 1696

Syntax

```
access-profile access-profile-name;
```

Hierarchy Level

```
[edit system services extensible-subscriber-services]
```

Description

Define the access profile name for time accounting. In most cases, the information about the access profile on the RADIUS server is the same as the information about the access profile of a control session. This configuration is mandatory.

Options

access-profile-name	Name of the access profile.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

access-profile-name (Duplicate Accounting)

IN THIS SECTION

- Syntax | 1697
- Hierarchy Level | 1697
- Description | 1697
- Options | 1697
- Required Privilege Level | 1697
- Release Information | 1698

Syntax

```
access-profile-name [profile-name];
```

Hierarchy Level

```
[edit access profile profile-name accounting duplication-vrf]
```

Description

Specify up to five access profiles, all in the same nondefault VRF (LS:RI combination), each of which lists one or more RADIUS accounting servers to which duplication accounting information is sent.

Options

profile-name Name of an access profile that lists RADIUS accounting servers for duplicate reporting.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.

Statement supported in Junos OS Release 13.2 and later releases.

RELATED DOCUMENTATION

Understanding RADIUS Accounting Duplicate Reporting

Configuring Authentication and Accounting Parameters for Subscriber Access

access-profile (Static Subscribers)

IN THIS SECTION

- Syntax | 1698
- Hierarchy Level | 1698
- Description | 1699
- Options | 1699
- Required Privilege Level | 1699
- Release Information | 1699

Syntax

```
access-profile profile-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name],
```

```
[edit logical-systems logical-system-name system services static-subscribers],
[edit logical-systems logical-system-name system services static-subscribers group group-name],
[edit routing-instances routing-instances-name system services static-subscribers],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name],
[edit system services static-subscribers],
[edit system services static-subscribers group group-name]
```

Description

Specify the access profile that triggers AAA services for all static subscribers on interfaces configured at the `[edit system services static-subscribers interface]` hierarchy level or for the static subscribers in a specific group. The group version of this statement overrides the global configuration.

Options

profile-name—Name of the static subscriber access profile.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Specifying the Static Subscriber Global Access Profile

Specifying the Static Subscriber Group Access Profile

access-security

IN THIS SECTION

- Syntax | 1700
- Hierarchy Level | 1701
- Description | 1701
- Required Privilege Level | 1702
- Release Information | 1702

Syntax

```
access-security {
    router-advertisement-guard {
        interface interface-name {
            mark-interface (trusted | block);
            policy policy-name (stateful | stateless);
        }
        vlans (vlan-name | all) {
            policy policy-name (stateful | stateless);
        }
        policy policy-name {
            accept {
                match-list {
                    match-criteria {
                        (match-all | match-any);
                    }
                    prefix-list-name prefix-list-name;
                    source-ip-address-list address-list-name;
                    source-mac-address-list address-list-name;
                }
                match-option {
                    hop-limit {
                        (maximum | minimum) value;
                    }
                    managed-config-flag;
                    other-config-flag;
```

```
                    router-preference (high | low | medium);
                }
            }
            discard {
                prefix-list-name prefix-list-name;
                source-ip-address-list address-list-name;
                source-mac-address-list address-list-name;
            }
        }
    }
    slaac-snooping {
        interface  (interface-name | all) {
            auto-dad {
                retries retries;
                retrans-interval seconds;
            }
            mark-interface {
                trusted;
            }
            max-allowed-contentions count {
                duration seconds;
            }
        }
        link-local {
            expiry interval seconds;
        }
        vlans (vlan-name | all);
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure IPv6 access security options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

access (SNMP)

IN THIS SECTION

- Syntax | 1702
- Hierarchy Level | 1703
- Description | 1703
- Required Privilege Level | 1703
- Release Information | 1703

Syntax

```
access {
    group group-name {
        (default-context-prefix | context-prefix context-prefix) {
            security-model (any | usm | v1 | v2c) {
                security-level (authentication | none | privacy) {
                    notify-view view-name;
                    read-view view-name;
```

```
                    write-view view-name;
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit snmp v3 vacm]
```

Description

Set SNMP access limits.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

access (Static Framed Routes)

IN THIS SECTION

- Syntax | 1704
- Hierarchy Level | 1704
- Description | 1704
- Required Privilege Level | 1705

Release Information | 1705

Syntax

```
access {
    route  ip-prefix</prefix-length> {
        metric route-cost;
        next-hop next-hop;
        preference route-distance;
        qualified-next-hop next-hop;
        tag tag-number
    }
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Configure access routes under both non-default and default routing instances, Configure access route to configure static framed route on Border Network Gateway.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

NOTE: You can configure access route to configure static framed route on Border Network Gateway only under `[edit logical-systems logical-system-name routing-options]` hierarchy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

access-trunk

IN THIS SECTION

- Syntax | 1705
- Hierarchy Level | 1705
- Description | 1705
- Required Privilege Level | 1706

Syntax

access-trunk;

Hierarchy Level

```
[edit logical-systems logical-system-name protocols vstp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vstp],
[edit protocols vstp vlan vlan-identifier interface interface-name],
[edit routing-instances routing-instance-name instance-type (layer2-control | virtual-switch)]
```

Description

Enable untagged RTSP BDPUs to be sent and received on the interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

account-layer2-overhead (PIC Level)

IN THIS SECTION

- Syntax | 1706
- Hierarchy Level | 1706
- Description | 1706
- Required Privilege Level | 1706
- Release Information | 1707

Syntax

```
account-layer2-overhead;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number]
```

Description

Enable the automatic adjustment of Layer 2 overhead in bytes, which is the octet adjustment per packet, based on the encapsulation on the logical interface for the total octet count for ingress and egress traffic on all the interfaces in the PIC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Accounting of the Layer 2 Overhead Attribute in Interface Statistics

Configuring Layer 2 Overhead Accounting in Interface Statistics

Verifying the Accounting of Layer 2 Overhead in Interface Statistics

accounting (Forwarding Options)

IN THIS SECTION

- Syntax | 1707
- Hierarchy Level | 1708
- Description | 1708
- Required Privilege Level | 1708
- Release Information | 1708

Syntax

```
accounting name {
    output {
        aggregate-export-interval seconds;
        cflowd hostname {
            aggregation {
                autonomous-system;
                destination-prefix;
                protocol-port;
                source-destination-prefix {
```

```
                    caida-compliant;
                }
                source-prefix;
            }
            autonomous-system-type (origin | peer);
            port port-number;
            version format;
        }
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        interface interface-name {
            engine-id number;
            engine-type number;
            source-address address;
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Specify the discard accounting instance name and options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Discard Accounting

Configuring Discard Accounting

accounting (interfaces)

IN THIS SECTION

- Syntax | 1709
- Hierarchy Level | 1709
- Description | 1709
- Required Privilege Level | 1710
- Release Information | 1710

Syntax

```
accounting {
    destination-class-usage;
    source-class-usage {
        direction;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet]
```

Description

Enable IP packet counters on an interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Source Class and Destination Class Usage

accounting (AAA Profile)

IN THIS SECTION

- Syntax | 1710
- Hierarchy Level | 1711
- Description | 1711
- Required Privilege Level | 1711
- Release Information | 1711

Syntax

```
accounting {
    network-element network-element-name;
}
```

Hierarchy Level

```
[edit unified-edge aaa profiles aaa-profile-name radius]
```

Description

Specify the network element providing policy management for TDF subscribers. The network element must already be defined at the `[edit access radius]` hierarchy level. This statement is required if the RADIUS servers cannot initiate a CoA request without an accounting record.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring an AAA Profile

Understanding AAA Profiles

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

accounting (Access Profile)

IN THIS SECTION

- Syntax | 1712
- Hierarchy Level | 1712
- Description | 1712

Required Privilege Level | 1713
Release Information | 1713

Syntax

```
accounting {
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    address-change-immediate-update;
    ancp-speed-change-immediate-update;
    coa-immediate-update;
    coa-no-override service-class-attribute;
    duplication;
    duplication-filter;
    duplication-vrf {
        access-profile-name profile-name;
        vrf-name vrf-name;
    }
    immediate-update;
    order [accounting-method];
    send-acct-status-on-config-change
    statistics (time | volume-time);
    update-interval minutes;
    wait-for-acct-on-ack;
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure RADIUS accounting parameters and enable RADIUS accounting for an access profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Accounting for Subscriber Access

RADIUS Authentication and Accounting Basic Configuration

accounting-backup-options (Access Profile)

IN THIS SECTION

- Syntax | 1713
- Hierarchy Level | 1714
- Description | 1714
- Required Privilege Level | 1714
- Release Information | 1714

Syntax

```
accounting-backup-options {
    max-pending-accounting-stops number;
    max-withhold-time hold-time;
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure options for backing up RADIUS accounting stop requests when all RADIUS accounting servers in the profile are offline.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Configuring Back-up Options for RADIUS Accounting

accounting (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | 1715
- Hierarchy Level | 1715
- Description | 1715
- Required Privilege Level | 1715
- Release Information | 1715

Syntax

```
(accounting | no-accounting);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name],
```

Description

Enable or disable the collection of IGMP join and leave event statistics for dynamically created IGMP interfaces.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Recording IGMP Join and Leave Events

accounting (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 1716
- Hierarchy Level | 1716
- Description | 1716
- Required Privilege Level | 1716
- Release Information | 1716

Syntax

```
(accounting | no-accounting);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Description

Enable or disable the collection of MLD join and leave event statistics for a dynamic interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Dynamic IGMP Configuration Overview

Example: Recording MLD Join and Leave Events

accounting-options

IN THIS SECTION

- Syntax | 1717
- Hierarchy Level | 1720
- Description | 1720
- Required Privilege Level | 1720
- Release Information | 1721

Syntax

```
accounting-options {
    class-usage-profile profile-name {
        file filename;
        interval minutes;
        destination-classes {
            destination-class-name;
        }
        source-classes {
            source-class-name;
        }
    }
    cleanup-interval days;
    file filename {
        archive-sites {
            site-name;
        }
        backup-on-failure (master-and-slave | master-only);
```

```
        compress;
        files number;
        nonpersistent;
        push-backup-to-master;
        size bytes;
        start-time time;
        transfer-interval minutes;
    }
    filter-profile profile-name {
        counters {
            counter-name;
        }
        file filename;
        interval minutes;
    }
    flat-file-profile profile-name{
        fields {
            all-fields;
            egress-stats {
                all-fields;
                input-bytes;
                input-packets;
                output-bytes;
                output-packets;
                queue-id;
                red-drop-bytes;
                red-drop-packets;
                tail-drop-packets;
                total-drop-packets;
            }
            general-param {
                all-fields;
                accounting-type;
                descr;
                line-id;
                logical-interface;
                nas-port-id;
                physical-interface;
                routing-instance;
                timestamp;
                user-name;
                vlan-id;
            }
```

```
            ingress-stats {
                all-fields;
                drop-packets;
                input-bytes;
                input-packets;
                output-bytes;
                output-packets;
                queue-id;
            }
            l2-stats {
                all-fields;
                input-mcast-bytes;
                input-mcast-packets;
            }
            overall-packet {
                all-fields;
                input-bytes;
                input-discards;
                input-errors;
                input-packets;
                inputv6-bytes;
                inputv6-packets;
                output-bytes;
                output-errors;
                output-packets;
                outputv6-bytes;
                outputv6-packets;
                input-v4-bytes;
                input-v4-packets;
                output-v4-bytes;
                output-v4-packets;
                input-bytes-per-sec;
                input-packets-per-sec;
            }
            service-accounting;
        }
        file filename;
        format (csv | ipdr)
        interval minutes;
         schema-version schema-name;
    }
    interface-profile profile-name {
        allow-clear;
```

```
        fields {
            field-name;
        }
        file filename;
        interval minutes;
    }
    mib-profile profile-name {
        file filename;
        interval seconds;
        object-names {
            mib-object-name;
        }
        operation operation-name;
    }
    routing-engine-profile profile-name {
        fields {
            field-name;
        }
        file filename;
        interval minutes;
    }
    detailed-transit-interface-stats enable;
}
```

Hierarchy Level

```
[edit]
```

Description

Configure options for accounting statistics collection.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuration Statements at the [edit accounting-options] Hierarchy Level

Accounting Options Configuration

accounting-order

IN THIS SECTION

- Syntax | 1721
- Hierarchy Level | 1721
- Description | 1721
- Options | 1722
- Required Privilege Level | 1722
- Release Information | 1722

Syntax

```
accounting-order (radius | [accounting-order-data-list]);
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Specify the order in which accounting methods are used.

Options

`radius`—Use the RADIUS accounting method.

[*accounting-order-data-list*]—Set of data listing the accounting order to be used, enclosed in brackets. This can be any combination of accounting methods, up to and including a list of the entire accounting order.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring RADIUS Authentication for L2TP

accounting (System) | **1746**

accounting-order (Service Accounting)

IN THIS SECTION

- Syntax | **1723**
- Hierarchy Level | **1723**
- Description | **1723**
- Default | **1723**
- Options | **1723**
- Required Privilege Level | **1724**
- Release Information | **1724**

Syntax

```
accounting-order (activation-protocol | local | radius);
```

Hierarchy Level

```
[edit access profile profile-name service]
```

Description

Specify the method that is used for reporting subscriber service accounting.

When you specify the `activation-protocol` method, the service accounting statistics are reported differently depending on how the service is activated:

- When the service is activated by the RADIUS server, the statistics are reported to the RADIUS server.
- When the service is activated by JSRC, the statistics are reported to JSRC.
- When the service is activated by the CLI configuration or a command—for example, through DHCP—the statistics are reported as determined by the flat-file profile configured at the `[edit access profile profile-name local]` hierarchy level. In this case, statistics are not reported if no flat-file profile is configured.

Default

`activation-protocol`

Options

activation-protocol Send service accounting reports by means of the application that activates services, such as JSRC, RADIUS, or the CLI.

local Send service accounting information to a flat file stored locally on the router. This method requires you to also configure the name of a flat-file profile at the `[edit access profile profile-name local]` hierarchy level. This profile determines which statistics and which nonstatistical parameters are collected for the service.

> **NOTE**: When you configure the `local` option, both volume and time statistics are collected for the service accounting sessions. In this case, you must not configure the `volume-time` option at the `[edit access profile profile-name service accounting statistics]` hierarchy level, or an error is generated when you commit the configuration.

radius Send service accounting reports to the RADIUS server by means of the RADIUS protocol.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

`local` option added in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Service Accounting with JSRC

Service Accounting with JSRC

Configuring Service Accounting in Local Flat Files

accounting-port (RADIUS Server)

IN THIS SECTION

- Syntax | 1725
- Hierarchy Level | 1725
- Description | 1725
- Options | 1725

Required Privilege Level | 1725
Release Information | 1725

Syntax

```
accounting-port port-number;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Specify the RADIUS server port number to which the MX Series router sends RADIUS accounting-start and accounting-stop requests. RADIUS accounting-start and accounting-stop requests are used when the RADIUS server is not able to initiate a change of authorization request without an accounting record.

Options

port-number Port number to which the RADIUS requests are sent.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

accounting-profile (firewall)

IN THIS SECTION

- Syntax | 1726
- Hierarchy Level | 1726
- Description | 1726
- Options | 1726
- Required Privilege Level | 1727
- Release Information | 1727

Syntax

```
accounting-profile name;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name]
```

Description

Enable collection of accounting data for the specified filter.

Options

name—Name assigned to the accounting profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Accounting for Firewall Filters Overview

accounting-profile

IN THIS SECTION

- Syntax | 1727
- Hierarchy Level | 1728
- Description | 1728
- Options | 1728
- Required Privilege Level | 1728
- Release Information | 1728

Syntax

```
accounting-profile name;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-range name]
```

Description

Enable collection of accounting data for the specified physical or logical interface or interface range.

Options

name—Name of the accounting profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying an Accounting Profile to the Physical Interface

Applying an Accounting Profile to the Logical Interface

accounting (Protocols AMT Interface)

IN THIS SECTION

- Syntax | 1729
- Hierarchy Level | 1729

- Description | 1729
- Default | 1729
- Required Privilege Level | 1729
- Release Information | 1729

Syntax

```
accounting;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols amt relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols amt
relay],
[edit protocols amt relay],
[edit routing-instances routing-instance-name protocols amt relay]
```

Description

Enable the collection of statistics for an Automatic Multicast Tunneling (AMT) interface.

Default

Disabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

accounting (Protocols IGMP)

IN THIS SECTION

- Syntax | 1730
- Hierarchy Level | 1730
- Description | 1730
- Required Privilege Level | 1730
- Release Information | 1731

Syntax

```
accounting;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp],
[edit protocols igmp]
```

Description

Enable the collection of IGMP join and leave event statistics on the system.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Recording IGMP Join and Leave Events

accounting (Protocols IGMP AMT Interface)

IN THIS SECTION

- Syntax | 1731
- Hierarchy Level | 1731
- Description | 1732
- Default | 1732
- Required Privilege Level | 1732
- Release Information | 1732

Syntax

```
(accounting | no-accounting);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp amt relay defaults],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols igmp
amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances routing-instance-name protocols igmp amt relay defaults]
```

Description

Enable or disable the collection of IGMP join and leave event statistics for an Automatic Multicast Tunneling (AMT) interface.

Default

Disabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Default IGMP Parameters for AMT Interfaces

accounting (Protocols IGMP Interface)

IN THIS SECTION

- Syntax | 1733
- Hierarchy Level | 1733
- Description | 1733
- Required Privilege Level | 1733
- Release Information | 1733

Syntax

```
(accounting | no-accounting);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Description

Enable or disable the collection of IGMP join and leave event statistics for an interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Recording IGMP Join and Leave Events

accounting (Protocols MLD)

IN THIS SECTION

- Syntax | 1734
- Hierarchy Level | 1734

Description | 1734
Required Privilege Level | 1734
Release Information | 1734

Syntax

```
accounting;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld],
[edit protocols mld]
```

Description

Enable the collection of MLD join and leave event statistics on the system.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Example: Recording MLD Join and Leave Events

accounting (Protocols MLD Interface)

IN THIS SECTION

- Syntax | 1735
- Hierarchy Level | 1735
- Description | 1735
- Required Privilege Level | 1735
- Release Information | 1735

Syntax

```
(accounting | no-accounting);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

Enable or disable the collection of MLD join and leave event statistics for an interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Example: Recording MLD Join and Leave Events

accounting (RADIUS Client)

IN THIS SECTION

- Syntax | 1736
- Hierarchy Level | 1736
- Description | 1736
- Required Privilege Level | 1737
- Release Information | 1737

Syntax

```
accounting {
    secret password;
    response-cache-timeout seconds;
}
```

Hierarchy Level

```
[edit access radius clients client-name]
```

Description

Specify a shared secret and response cache timeout to be used by the MX Series router and the RADIUS client for accounting.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers

accounting-secret (RADIUS Server)

IN THIS SECTION

- Syntax | 1737
- Hierarchy Level | 1738
- Description | 1738
- Default | 1738
- Options | 1738
- Required Privilege Level | 1738
- Release Information | 1738

Syntax

```
accounting-secret password;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure the secret password to be used when sending accounting-start requests to the RADIUS server if the accounting secret password is different from the authentication secret password. RADIUS accounting-start requests are used when the RADIUS server is not able to initiate a change of authorization request without an accounting record.

Default

Use the same password that is used for authentication requests.

Options

password Password for accounting requests.

- **Range:** 1 through 64 characters

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

accounting-server

IN THIS SECTION

- Syntax | 1739
- Hierarchy Level | 1739
- Description | 1739
- Default | 1739
- Options | 1740
- Required Privilege Level | 1740
- Release Information | 1740

Syntax

```
accounting-server[server-addresses];
```

Hierarchy Level

```
[edit access                profile                profile-name                radius]
```

Description

Configure the Remote Authentication Dial-In User Service (RADIUS) server for authentication. To configure multiple RADIUS servers, include multiple server addresses. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Default

Not enabled

Options

server-addresses—One or more addresses of RADIUS authentication servers.

> **NOTE**: The `[edit access]` hierarchy is not available on QFabric systems.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

RELATED DOCUMENTATION

- *show network-access aaa statistics authentication*
- *Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch*
- *Understanding 802.1X and RADIUS Accounting on Switches*
- *Understanding RADIUS Accounting*

accounting-server-name

IN THIS SECTION

- Syntax | 1741
- Hierarchy Level | 1741
- Description | 1741
- Options | 1741
- Required Privilege Level | 1741
- Release Information | 1741

Syntax

```
accounting-server-name hostname;
```

Hierarchy Level

```
[edit access profile profile-name radius]
```

Description

Configure the RADIUS accounting server for an access profile using a fully-qualified domain name (FQDN). Use an FQDN if the server name can translate to more than one address.

The distribution of RADIUS requests among the resolved addresses is based on the configured client-authentication-algorithm. You can configure the direct or round-robin algorithm at the [`edit access profile profile-name radius options client-authentication-algorithm`] hierarchy level.

Options

hostname—Configure the FQDN to identify the RADIUS server.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Specifying RADIUS Server Connections on Switches (CLI Procedure)

accounting (Service Accounting)

IN THIS SECTION

- Syntax | 1742
- Hierarchy Level | 1742
- Description | 1742
- Required Privilege Level | 1742
- Release Information | 1743

Syntax

```
accounting {
    statistics (time | volume-time);
    update-interval minutes;
}
```

Hierarchy Level

```
[edit access profile profile-name service]
```

Description

Define the service accounting configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R1.

RELATED DOCUMENTATION

Configuring Service Accounting

Configuring Per-Subscriber Session Accounting

Processing Cisco VSAs in RADIUS Messages for Service Provisioning

accounting-stop-on-access-deny

IN THIS SECTION

- Syntax | 1743
- Hierarchy Level | 1743
- Description | 1744
- Required Privilege Level | 1744
- Release Information | 1744

Syntax

```
accounting-stop-on-access-deny;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure RADIUS accounting to send an Acct-Stop message when the AAA server refuses a client request for access.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Authentication and Accounting Basic Configuration

accounting-stop-on-failure

IN THIS SECTION

- Syntax | 1744
- Hierarchy Level | 1745
- Description | 1745
- Required Privilege Level | 1745
- Release Information | 1745

Syntax

```
accounting-stop-on-failure;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure RADIUS accounting to send an Acct-Stop message when a subscriber session has been successfully authenticated and authorized, but then fails before an Acct-Start message is sent. By default, an Acct-Stop message is sent only if an Acct-Start message has been exchanged with the accounting server.

Consider a situation where RADIUS address pools are used to assign IP/IPv6 addresses. After a subscriber session is successfully authenticated, the RADIUS server authorizes the session by assigning an IP address from the RADIUS address pool and conveying that address in the Framed-IP-Address attribute. If a negotiation failure occurs at this point, the session is terminated before activating. The Acct-Start message is never sent because it is initiated by session activation. By default, an Acct-Stop message cannot be sent because the Acct-Start is never sent. However, if the `acct-stop-on-failure` statement is configured, the negotiation failure causes the Acct-Stop message to be sent, which explicitly notifies the RADIUS server that the session is disconnected and that it can free the allocated IP address back to the pool.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Authentication and Accounting Basic Configuration

accounting (System)

IN THIS SECTION

- Syntax | 1746
- Hierarchy Level | 1747
- Description | 1747
- Options | 1747
- Required Privilege Level | 1748
- Release Information | 1748

Syntax

```
accounting {
    destination {
        radius {
            server {
                server-address {
                    accounting-port port-number;
                    accounting-retry number;
                    accounting-timeout seconds;
                    dynamic-request-port number;
                    max-outstanding-requests value;
                    port number;
                    preauthentication-port number;
                    preauthentication-secret secret;
                    retry number;
                    routing-instance routing-instance-name;
                    secret password;
                    source-addresssource-address;
                    timeout seconds;
                }
        }
        tacplus {
            server {
                server-address {
                    port port-number;
```

```
                    routing-instance routing-instance;
                    secret password;
                    single-connection;
                    source-address address
                    timeout seconds;
                }
            }
        }
    }
    enhanced-avs-max <number>;
    events (change-log | interactive-commands | login );
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure an audit of TACACS+ or RADIUS authentication events, configuration changes, and interactive commands. Auditing these factors helps you track network usage for auditing and billing purposes.

Options

enhanced-avs-max

Configure the number of attribute-value pairs to be displayed, each of which can store a maximum of 250 bytes.

- **Default:** 7 pairs
- **Range:** 7 through 15 pairs

events

(Required) Configure the types of events to log.

- Values:
  - change-log—Configuration changes
  - interactive-commands—Interactive commands (any command-line input)
  - login—Login sessions

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`enhanced-avs-max` statement introduced in Junos OS Release 14.1.

`routing-instance` introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Configuring RADIUS System Accounting

Configuring TACACS+ System Accounting

acfc

IN THIS SECTION

- Syntax | 1749
- Hierarchy Level | 1749
- Description | 1749
- Required Privilege Level | 1749
- Release Information | 1749

Syntax

```
acfc;
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options compression],
[edit interfaces interface-name unit logical-unit-number ppp-options compression],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options compression]
```

Description

For interfaces with PPP encapsulation, configure compression of the Data Link Layer address and control fields. The `acfc` option is not supported with `frame-relay-ppp` encapsulation.

On M320, M120, and T Series routers, address and control field compression (ACFC) is not supported for any ISO family protocols. Do not include the `acfc` statement at the `[edit interfaces interface-name ppp-options compression]` hierarchy level when you include the `family iso` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring PPP

ack-number

IN THIS SECTION

- Syntax | 1750
- Hierarchy Level | 1750
- Description | 1750
- Options | 1750
- Required Privilege Level | 1751
- Release Information | 1751

Syntax

```
ack-number {
    match (Security IDP Policy) (equal | greater-than | less-than | not-equal);
    value acknowledgement-number;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify the ACK number of the packet. This number identifies the next sequence number; the ACK flag must be set to activate this field.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`acknowledgement-number`*—Match the ACK number of the packet.
- **Range:** 0 through 4,294,967,295

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

acknowledge-retries

IN THIS SECTION

- Syntax | 1751
- Hierarchy Level | 1751
- Description | 1752
- Options | 1752
- Required Privilege Level | 1752
- Release Information | 1752

Syntax

```
acknowledge-retries number;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

For link services interfaces only, configure the number of retransmission attempts to be made for consecutive hello or remove link messages following the expiration of the acknowledgment timer.

Options

number—Number of retransmission attempts to be made following the expiration of the acknowledgment timer.

- **Range:** 1 through 5
- **Default:** 2

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

action-red-differential-delay | **1782**

hello-timer | **8043**

Configuring Acknowledgment Timers on Link Services Physical Interfaces

acknowledge-timer

IN THIS SECTION

- Syntax | **1753**
- Hierarchy Level | **1753**

- Description | 1753
- Options | 1753
- Required Privilege Level | 1753
- Release Information | 1754

Syntax

```
acknowledge-timer milliseconds;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

For link services interfaces only, configure the maximum time, in milliseconds, to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment message.

Options

milliseconds—Time to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment message.

- **Range:** 1 through 10 milliseconds
- **Default:** 4 milliseconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

address (Interfaces) | 1878

hello-timer | 8043

Configuring Acknowledgment Timers on Link Services Physical Interfaces

acquiring-state-announce-grandmaster

IN THIS SECTION

- Syntax | 1754
- Hierarchy Level | 1754
- Description | 1755
- Options | 1755
- Required Privilege Level | 1755
- Release Information | 1755

Syntax

```
acquiring-state-announce-grandmaster {
    default-ts;
}
```

Hierarchy Level

```
[edit logical-systems name   protocols ptp]
[edit protocols ptp]
```

Description

Use grandmaster data set for announce messages when in acquiring state.

Options

default-ts Local data set for announce messages when in acquiring state.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

IEEE 1588v2 Precision Timing Protocol (PTP)

Precision Time Protocol Overview

action (QFX)

IN THIS SECTION

- Syntax | 1756
- Hierarchy Level | 1756
- Description | 1756
- Options | 1756
- Required Privilege Level | 1756
- Release Information | 1756

Syntax

```
action {
    loss-priority high then discard;
}
```

Hierarchy Level

```
[edit firewall three-color-policer name]
```

Description

Discard traffic on a logical interface using tricolor marking policing.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

action (chassis error)

IN THIS SECTION

- Syntax | 1757
- Hierarchy Level | 1757

- Description | 1757
- Options | 1758
- Usage | 1758
- Required Privilege Level | 1758
- Release Information | 1758

Syntax

```
action {
    alarm;
    get-state;
    log;
    trap;
    (reset | offline | reset-pfe | disable-pfe);
}
```

Hierarchy Level

[edit chassis "fpc" on page 7516 *slot-number* "error " on page 6441 "fatal" on page 6950]

[edit chassis "fpc" on page 7516 *slot-number* "error " on page 6441 "major" on page 10290]

[edit chassis "fpc" on page 7516 *slot-number* "error " on page 6441 "minor" on page 10760]

[edit chassis "error " on page 6441 "fatal" on page 6950]

[edit chassis "error " on page 6441 "major" on page 10290]

[edit chassis "error " on page 6441 "minor" on page 10760]

Description

Depending on the severity level of an error, you can configure one or more actions for the router to perform. For example, if the severity level of an error is major, you can configure the system to raise an alarm and restart the FPC. You can use `offline-pfe` and `reset-pfe` options to map errors belonging to `pfe` category.

Options

alarm Raise an FPC alarm.

disable-pfe (PTX routers) Disable Packet Forwarding Engine interfaces on an FPC. (MX240, MX480, MX960, MX2020 PTX10008 routers). Disables the interface associated with the ASIC that raised the error. Disables all PFE interfaces, alarms and logs. On MX Series routers, an MPC can have multiple ASICs handled by a single PFE. Disabling the PFE in such cases, is not the optimal way to handle the error.

get-state The system starts collecting statistics counters and other data from the affected FPC. The data is written and saved to a file under **/var** on the routing engine.

offline Disables the alarms, logs, and shuts down the affected FPC, allowing traffic to be re-routed through interfaces on other FPC in the device.

reset The system takes the affected FPC offline and resets to online, enables the alarms and logs.

offline-pfe Disables the alarms, logs, and shuts down the affected PFE.

reset-pfe Powers-off the PFE, disables the alarms and logs, then, powers-on the PFE, enables the alarms and logs.

log Log occurrences to the system log file.

trap The system raises traps for the FPC errors.

Usage

The usage guidelines are as follows:

- The specified PFE is disabled automatically if `offline-pfe` or `reset-pfe` is configured. For details, see *Configuring FPC Error Levels and Actions*.

Required Privilege Level

interface To view this statement in the configuration.

interface-control To add this statement to the configuration.

Release Information

Statement introduced for PTX Series routers in Junos OS Release 13.3.

Statement introduced for MX240, MX480, MX960 routers in Junos OS Release 14.2.

RELATED DOCUMENTATION

show chassis fpc errors | 18387

Configuring FPC Error Levels and Actions

action

IN THIS SECTION

- Syntax | 1759
- Hierarchy Level | 1759
- Description | 1759
- Required Privilege Level | 1760
- Release Information | 1760

Syntax

```
action {
    loss-priority high then discard;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name],
[edit firewall three-color-policer name],
[edit logical-systems logical-system-name firewall three-color-policer name]
```

Description

Discard traffic on a logical interface using tricolor marking policing.

NOTE: This statement is supported only on IQ2 interfaces.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Logical systems support introduced in Junos OS Release 9.3.

Support at the `[edit dynamic-profiles ... three-color-policer]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Three-Color Policer Configuration Overview
- Basic Single-Rate Three-Color Policers
- Basic Two-Rate Three-Color Policers
- Two-Color and Three-Color Logical Interface Policers
- Two-Color and Three-Color Physical Interface Policers
- Two-Color and Three-Color Policers at Layer 2
- loss-priority high then discard (Three-Color Policer) | **10155**

action (APN GTP)

IN THIS SECTION

- Syntax | **1761**

- Hierarchy Level | 1761
- Description | 1761
- Options | 1761
- Required Privilege Level | 1762
- Release Information | 1762

Syntax

```
action {
    drop;
    pass;
    selection (ms|net|vrf);
}
```

Hierarchy Level

```
[edit security gtp profile profile-name apn pattern-string imsi-prefiximsi-prefix]
```

Description

Use the GPRS tunneling protocol (GTP) profile access point name (APN) action (pass |drop |selection) to filter packets that helps you to allow or deny access to specific access points.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name apn pattern-string imsi-prefix` hierarchy-level is replaced by `[edit security gtp profile profile-name apn pattern-string imsi-prefix`.

Options

- **drop**—Use drop to specify to deny GTP packets from all selection modes for the specified access points.
- **pass**—Use pass to specify to permit GTP packets from all selection modes for the access points.
- **selection**—Use selection to specify one of the following selection modes for the access points:

- **ms**—The access point name is provided by a mobile station, and the user-subscription is not verified.
- **net**—The access point name is provided by a network, and the user subscription is not verified.
- **vrf**— The access point name is provided by a network or an MS, and the user-subscription is verified.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Understanding GTP APN Filtering

Understanding GTPv2 IMSI Prefix and APN Filtering

Example: Setting a GTP APN and a Selection Mode

action (OAM)

IN THIS SECTION

- Syntax | **1763**
- Hierarchy Level | **1763**
- Description | **1763**
- Required Privilege Level | **1763**
- Release Information | **1763**

Syntax

```
action {
    link-down;
    send-critical-event;
    syslog;
}
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile]
```

Description

Define the action or actions to be taken when the OAM fault event occurs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Specifying the Actions to Be Taken for Link-Fault Management Events*

action (OAM LFM)

IN THIS SECTION

- Syntax | 1764
- Hierarchy Level | 1764
- Description | 1764
- Required Privilege Level | 1764
- Release Information | 1765

Syntax

```
action {
                              syslog;
                              link-down;
 }
```

Hierarchy Level

```
[edit protocols oam                ethernet                link-fault-management]
```

Description

Define the action or actions to be taken when the OAM link fault management (LFM) fault event occurs.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Ethernet OAM Link Fault Management

action (Policer)

IN THIS SECTION

- Syntax | 1765
- Hierarchy Level | 1765
- Description | 1765
- Required Privilege Level | 1766
- Release Information | 1766

Syntax

```
action {
    loss-priority high then discard;
}
```

Hierarchy Level

```
[edit firewall three-color-policer policer-name]
```

Description

This statement discards high loss priority traffic as part of a configuration using tricolor marking on a logical interface.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Junos OS Class of Service User Guide for Routing Devices

logical-interface-policer | 10069

action-priority (Routing Instance)

IN THIS SECTION

- Syntax | 1766
- Hierarchy Level | 1767
- Description | 1767
- Options | 1767
- Required Privilege Level | 1767
- Release Information | 1768

Syntax

```
action-priority value;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols protocol-name interface interface-name]
```

Description

Configure the priority for an interface on which the MAC move is applied. When a MAC move is configured with a corresponding action, and that limit is exceeded, the action is applied to one of the interfaces associated with that MAC address having the highest action priority. If the action configured for the MAC move is `shutdown`, you can use the action priority to determine which interface is disabled. If an interface is trusted, assign a low action priority to decrease the likelihood that the interface shuts down if the configured action is `shutdown` or `vlan-member-shutdown`.

If no action priority is configured, or if the interfaces have the same action priority, then the action is applied to the interface to which the MAC address moved last.

Options

value Value assigned to an interface associated with a MAC address move. The value determines the priority with which an action is applied to the interface if the MAC address move limit is reached. The interface with the highest priority is the interface with the lowest value configured for `action-priority`.

A higher value means lower priority. For example, if a MAC address move occurs between two interfaces with the action priority value set to 5 and 6, the interface with value 5 as the action priority value is blocked.

If a VPLS instance has MAC move action enabled, set the `action-priority` to 8 for the router to detect an interface as a core interface by the VPLS customer edge loop prevention feature. By default, the loop prevention feature supports lsi and vt-type interfaces as core interfaces; however, you can use this special action-priority value to designate any interface as a core interface.

- **Default:** 4
- **Range:** 0 through 8

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring Loop Prevention in VPLS Network Due to MAC Moves

action-priority

IN THIS SECTION

- Syntax | 1768
- Hierarchy Level | 1768
- Description | 1768
- Default | 1769
- Required Privilege Level | 1769
- Release Information | 1769

Syntax

```
action-priority value;
```

Hierarchy Level

```
[edit vlans vlan-name switch-options mac-move-limit interface interface-name]
```

Description

Configure a priority for an interface on which the MAC move limit action will be applied. When a MAC move limit is configured, and a MAC address moves to a new interface more times than is allowed by the

limit, the configured action will be applied to the interface associated with that MAC address having the highest priority. The interface with the highest priority is the interface with the lowest value configured for `action-priority`. The default value for `action-priority` on an interface is 4.

If no action priority is configured, or if the interfaces have the same action priority, then the action will be applied to the interface to which the MAC address moved last.

Default

The default value for `action-priority` on an interface is 4.

Required Privilege Level

system—To view this statement in the configuration.
system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Overview of Port Security

Configuring MAC Move Limiting

Configuring Autorecovery for Port Security Events

action-priority (bridge domain)

IN THIS SECTION

- Syntax | 1770
- Hierarchy Level | 1770
- Description | 1770
- Default | 1770
- Required Privilege Level | 1770

Release Information | 1770

Syntax

```
action-priority number;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name bridge-options interface interface-name]
```

Description

Configure the action priority value for an interface in a bridge domain on MAC move detection. This priority value is used to determine which interface should be blocked when a throttled MAC move is detected between two interfaces. The priority value can be between 0 and 7 inclusive. A higher value means lower priority. For example, if a MAC address move occurs between two interfaces with the action priority value set to 5 and 6, the interface with value 5 as the action priority value is blocked.

Default

4

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring a Bridge Domain

action-profile

IN THIS SECTION

- Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches | 1771
- Syntax: EX Series Switches and NFX Series Devices | 1772
- Hierarchy Level | 1772
- Description | 1772
- Options | 1772
- Required Privilege Level | 1772
- Release Information | 1773

Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches

```
action-profile profile-name {
    action {
        link-down;
        send-critical-event;
        syslog;
    }
    event {
        link-adjacency-loss;
        link-event-rate {
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
        }
        protocol-down;
    }
}
```

Syntax: EX Series Switches and NFX Series Devices

```
action-profile profile-name;
    action {
        syslog;
        link-down;
    }
    event {
        link-adjacency-loss;
        link-event-rate {
            frame-error  count;
            frame-period count;
            frame-period-summary  count;
            symbol-period count;
        }
    }
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management]
```

Description

Configure an Ethernet OAM link fault management (LFM) action profile by specifying a profile name.

The remaining statements are explained separately. See CLI Explorer.

Options

profile-name—Name of the action profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring an OAM Action Profile

Configuring Ethernet OAM Link Fault Management

action-profile (Network Management)

IN THIS SECTION

- Syntax | 1773
- Hierarchy Level | 1774
- Description | 1774
- Options | 1774
- Required Privilege Level | 1775
- Release Information | 1775

Syntax

```
action-profile profile-name {
    event (jexec | lbt | lt-enter |  lt-leave | mac-egress | mac-ingress | np-egress | np-
ingress | pot) {
        count;
        packet-dump;
        packet-summary;
        trace;
    }
    module {
```

```
        flow {
            flag {
                all;
            }
        }
    }
    preserve-trace-order;
    record-pic-history;
}
```

Hierarchy Level

```
[edit security datapath-debug]
```

Description

Configure the action profile options for data path debugging.

Options

- *action-profile name* — Name of the action profile.
- event—Enable the events to trace the packet when the packet hit the events (jexec, lbt, lt-enter, lt-leave, mac-egress, mac-ingress, np-egress, np-ingress, pot)
  - count—Number of times a packet hits the specified event.
  - packet-dump—Capture the packet that hits the specified event.
  - packet-summary—Print the source/destination IP address details with protocol number and IP length details along with trace message for the specified event.
  - trace—Print the standard trace message when the packet hits the specified event.
- module—Turn on the flow session related trace messages.
  - flow—Trace flow session related messages.
  - flag—Specify which flow message needs to be traced.
  - all—Trace all possible flow trace messages.
  - trace—Print the standard trace message when the packet hits the specified event.

- `preserve-trace-order`—Preserve trace order.
- `record-pic-history`—Record the PICs in which the packet has been processed.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

> *Example: Configuring Packet Capture for Datapath Debugging*

action-profile (Applying to CFM)

IN THIS SECTION

- Syntax | 1775
- Hierarchy Level | 1776
- Description | 1776
- Options | 1776
- Required Privilege Level | 1776
- Release Information | 1776

Syntax

```
action-profile profile-name;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id remote-mep mep-id]
```

Description

Identify the action profile to use.

Options

profile-name—Name of the action profile to use.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring a MEP to Generate and Respond to CFM Protocol Messages

action-profile (Applying to OAM CFM, for EX Series Switch Only)

IN THIS SECTION

- Syntax | 1777
- Hierarchy Level | 1777
- Description | 1777

- Options | 1777
- Required Privilege Level | 1778
- Release Information | 1778

Syntax

```
action-profile profile-name {
    action {
        interface-down;
    }
    default-actions {
        interface-down;
    }
    event {
        adjacency-loss;
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id remote-mep mep-id]
```

Description

Configure a name and default action for an action profile.

Options

profile-name—Name of the action profile.

action—Defines the action to be taken when connectivity to the remote MEP is lost.

interface-down—Brings the interface down when a remote MEP connectivity failure. is detected.

default-actions—Defines the default action to be taken when connectivity to the remote MEP is lost.

interface-down—Brings the interface down when a remote MEP connectivity failure is detected.

event—Defines the event to be monitored when a remote MEP connectivity failure is detected.

adjacency-loss—Defines the connectivity loss to the remote MEP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)

Junos OS Network Interfaces Configuration Guide

action-profile (Defining for CFM)

IN THIS SECTION

- Syntax | 1779
- Hierarchy Level | 1779
- Description | 1779
- Options | 1779
- Required Privilege Level | 1780
- Release Information | 1780

Syntax

```
action-profile profile-name {
    event {
        ais-trigger-condition {
            adjacency-loss;
            all-defects;
            cross-connect-ccm;
            erroneous-ccm;
            receive-ais;
        }
        interface-status-tlv (down | lower-layer-down);
        port-status-tlv blocked;
        rdi;
    }
    action {
        interface-down;
        log-and-generate-ais {
            interval(1m | 1s);
            level value;
            priority value;
        }
    }
    default-actions {
        interface-down;
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

Description

Configure a name and default action for an action profile.

Options

profile-name—Name of the action profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring a CFM Action Profile to Specify CFM Actions for CFM Events

default-actions | **5180**

event (CFM) | **6521**

interface-down | **8870**

action-profile (MEP)

IN THIS SECTION

- Syntax | **1781**
- Hierarchy Level | **1781**
- Description | **1781**
- Options | **1781**
- Required Privilege Level | **1781**
- Release Information | **1781**

Syntax

```
action-profile action-profile-name;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id]
```

Description

Attach the configured action profile to the MEP depending on the hierarchy level.

Options

action-profile-name	Name of the action profile that is configured for the CFM MEP and the server MEP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

> *Ethernet Alarm Indication Signal (ETH-AIS) Function Overview*
>
> *Configuring ETH-AIS on a CFM MEP*

action-red-differential-delay

IN THIS SECTION

- Syntax | 1782
- Hierarchy Level | 1782
- Description | 1782
- Options | 1782
- Required Privilege Level | 1783
- Release Information | 1783

Syntax

```
action-red-differential-delay (disable-tx | remove-link);
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

For link services interfaces only, configure the action to be taken when the differential delay exceeds the red limit.

Options

`disable-tx`—Disable transmission on the bundle link.

`remove-link`—Remove the bundle link from service.

- **Default:** `remove-link`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

yellow-differential-delay | 31168

Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16

action (Security Content Security Web Filtering)

IN THIS SECTION

- Syntax | 1783
- Hierarchy Level | 1784
- Description | 1784
- Options | 1784
- Required Privilege Level | 1784
- Release Information | 1784

Syntax

```
action  (block | log-and-permit | permit | quarantine);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated profile profile-name
category customurl-last-name]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name category
customurl-last-name]
```

Description

Enter an action to go with the customurl-list filter.

Options

- block—Log the error and deny the traffic.
- log-and-permit—Log the error and permit the traffic.
- permit—Permit the traffic.
- quarantine—Show the warning message and permit/block the traffic based on user input.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Surf-Control feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

action (Security Rulebase IPS)

IN THIS SECTION

- Syntax | 1785
- Hierarchy Level | 1785
- Description | 1785
- Options | 1785
- Required Privilege Level | 1786
- Release Information | 1786

Syntax

```
action {
    class-of-service {
        dscp-code-point number;
         forwarding-class forwarding-class;
    }
    (close-client | close-client-and-server | close-server |drop-connection | drop-packet |
ignore-connection | mark-diffserv value | no-action | recommended);
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name then]
```

Description

Specify the actions you want IDP to take when the monitored traffic matches the attack objects specified in the rules.

Options

- `no-action`—No action is taken. Use this action when you want to only generate logs for some traffic.

- `ignore-connection`—Stops scanning traffic for the rest of the connection if an attack match is found. IDP disables the rulebase for the specific connection.
- `mark-diffserv` *value*—Assigns the indicated service-differentiation value to the packet in an attack, then passes them on normally.
- `class-of-service`—Associates a class-of-service forwarding class as an action to the IDP policy; also sets the value of the DSCP code point. You can use the default forwarding class names or define new ones. Forwarding-class and dscp-code-point are optional, but one must be set.
- `drop-packet`—Drops a matching packet before it can reach its destination but does not close the connection. Use this action to drop packets for attacks in traffic that is prone to spoofing, such as UDP traffic. Dropping a connection for such traffic could result in a denial of service that prevents you from receiving traffic from a legitimate source-IP address.
- `drop-connection`—Drops all packets associated with the connection, preventing traffic for the connection from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing.
- `close-client`—Closes the connection and sends an RST packet to the client but not to the server.
- `close-server`—Closes the connection and sends an RST packet to the server but not to the client.
- `close-client-and-server`—Closes the connection and sends an RST packet to both the client and the server.
- `recommended`—All predefined attack objects have a default action associated with them. This is the action that Juniper Networks recommends when that attack is detected.

> **NOTE**: The actions are listed in the ascending order of severity from low to high. The most severe action is used when there are multiple rule hits for a single session.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

action-shutdown

IN THIS SECTION

- Syntax | 1787
- Hierarchy Level | 1787
- Description | 1787
- Default | 1788
- Required Privilege Level | 1788
- Release Information | 1788

Syntax

```
action-shutdown;
```

Hierarchy Level

- For platforms with Enhanced Layer 2 Software (ELS):

```
[edit forwarding-options storm-control-profiles profile-name]
```

- For platforms without ELS:

```
[edit ethernet-switching-options storm-control]
```

Description

Logically shut down or temporarily disable interfaces when the storm control level is exceeded.

To configure the shutdown action so that the interfaces are disabled temporarily, and recover automatically after a specified period of time:

- (MX Series, QFX Series, and EX switches that support ELS) Configure both the `action-shutdown` and the `recovery-timeout` statements. The interfaces recover automatically when the recovery timeout expires.

- (EX switches that do not support ELS) Configure both the `action-shutdown` and the `port-error-disable` statements. The interfaces recover automatically when the disable timeout expires. (The `port-error-disable` statement is not supported on QFX Series switches or MX Series routers.)

If you configure the `action-shutdown` statement without configuring either the `port-error-disable` or `recovery-timeout` statement, the interfaces that are enabled for storm control are shut down when the storm control level is exceeded and they do not recover automatically from that port-error condition.

If you do not configure the switch for autorecovery from the disabled condition, you can bring up the disabled interfaces by running one of the following commands:

- (MX Series)—Issue the `clear bridge recovery-timeout`
- (QFX Series)—Issue the `clear ethernet-switching recovery-timeout`
- (EX Series switches that support ELS)—Issue the `clear ethernet-switching recovery-timeout`
- (EX Series switches that do not support ELS)—Issue the `clear ethernet-switching port-error`

> **NOTE**: On EX4300 switches, `action-shutdown` causes an interface to stop learning MAC addresses and it also drops all incoming packets, but does not disable the physical interface.

Default

The `action-shutdown` option is not enabled by default. The switching device drops packets for the controlled traffic types if the ingress rate of the combined traffic streams exceeds the specified storm control level. Depending upon the configuration, applicable traffic could include broadcast, unknown unicast, and multicast traffic.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Hierarchy level `[edit forwarding-options]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

recovery-timeout | 14149

clear bridge recovery-timeout | 3446

clear ethernet-switching recovery-timeout | 3582

Example: Using Storm Control to Prevent Network (MX Routers)

Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches

Configuring Autorecovery for Port Security Events

action (Uplink Failure Detection)

IN THIS SECTION

- Syntax | 1789
- Hierarchy Level | 1789
- Description | 1790
- Options | 1790
- Required Privilege Level | 1790
- Release Information | 1790

Syntax

```
action {
     log;
}
```

Hierarchy Level

```
[edit protocols uplink-failure-detection]
```

Description

Define an action on uplink-failure-detection group state change.

Options

log Generate a system log message.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

actions

IN THIS SECTION

- Syntax | 1790
- Hierarchy Level | 1791
- Description | 1791
- Options | 1791
- Required Privilege Level | 1792
- Release Information | 1793

Syntax

```
actions {
    allow-strong-certificate;
    crl {
```

```
        disable;
        if-not-present (allow | drop);
        ignore-hold-instruction-code;
    }
    disable-session-resumption;
    ignore-server-auth-failure;
    log {
        all;
        errors;
        info;
        sessions-allowed;
        sessions-dropped;
        sessions-ignored;
        sessions-whitelisted;
        warning;
    }
    renegotiation {
        (allow | allow-secure | drop);
    }
}
```

Hierarchy Level

```
[edit services ssl proxy (Services) profile (Services SSL Proxy)]
```

Description

Specify the logging and traffic related actions for a SSL proxy profile.

An SSL proxy profile is required to configure SSL proxy on your SRX Series Firewall. As a part of the proxy profile configuration, you can configure– actions related to certification revocations checks, options to specify if a change in SSL parameters requires renegotiation for a session, option to disable session resumption, option to ignore certificate validation, root CA expiration dates, and other such issues based on your requirements.

Options

- `allow-strong-certificate`—Enable devices to use the RSA certificates with key size 4,096 bits. By default, this option is disabled. Option is available on SRX300, SRX320, and SRX380 devices in standalone mode.

Default - Not configured.

- `crl`—Specify the certificate revocation actions.
  - `disable`—Disable CRL verification.
  - `if-not-present`—Specify actions for sessions.
    - `allow`—Allow sessions when CRL information is not available.
    - `drop`—Drop sessions when CRL information is not available.
  - `ignore-hold-instruction-code`—Ignore the unconfirmed (on hold) revocation status, and accept a certificate.
- `disable-session-resumption`—Disable session resumption.
- `ignore-server-auth-failure`—Ignore server authentication failure.
- `log`—Specify the logging actions.
  - `all`—Log all events.
  - `errors`—Log all error events.
  - `info`—Log all information events.
  - `sessions-allowed`—Log SSL session allowed events after an error.
  - `sessions-dropped`—Log only SSL session dropped events.
  - `sessions-ignored`—Log session ignored events.
  - `sessions-whitelisted`—Log SSL session allowlisted events.
  - `warning`—Log all warning events.
- `renegotiation`—Specify the renegotiation options.
  - `allow`—Allow secure and nonsecure renegotiation.
  - `allow-secure`—Allow secure negotiation only.
  - `drop`—Drop session on renegotiation request.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. The `crl` statement is supported from Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

SSL Proxy Overview

Configuring SSL Forward Proxy

Enabling Debugging and Tracing for SSL Proxy

actions (Services SSL Initiation)

IN THIS SECTION

- Syntax | 1793
- Hierarchy Level | 1794
- Description | 1794
- Options | 1794
- Required Privilege Level | 1794
- Release Information | 1794

Syntax

```
actions {
    crl {
        disable;
        if-not-present (allow | drop);
        ignore-hold-instruction-code;
    }
    ignore-server-auth-failure;
}
```

Hierarchy Level

```
[edit services ssl initiation profile profile-name]
```

Description

Specify the certification revocation checks and traffic related actions for configuring SSL initiation support service. As a part of SSL initiation profile, you can specify actions related to certification revocations checks and chose an option to ignore certificate validation, root CA expiration dates, and other such issues based on your requirements. Commonly ignored errors include the inability to verify CA signature, incorrect certificate expiration dates, and so forth. We do not recommend using this option for authentication because configuring it results in websites not being authenticated at all.

Options

- `crl`—Specify the certificate revocation actions.
  - `disable`—Disable CRL verification.
  - `if-not-present`—Specify actions for sessions.
    - `allow`—Allow sessions when CRL information is not available.
    - `drop`—Drop sessions when CRL information is not available.
  - `ignore-hold-instruction-code`—Ignore the unconfirmed (on hold) revocation status, and accept a certificate.
- `ignore-server-auth-failure`—Ignore server authentication failure.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

SSL Proxy Overview

Configuring SSL Forward Proxy

Enabling Debugging and Tracing for SSL Proxy

activate

IN THIS SECTION

- Syntax | 1795
- Description | 1795
- Options | 1795
- Required Privilege Level | 1796
- Release Information | 1796

Syntax

```
activate <statement | identifier >
```

Description

Remove the `inactive:` tag from a statement, effectively adding the statement or identifier back to the configuration. Statements or identifiers that have been activated take effect when you next issue the `commit` command.

Options

identifier Identifier from which you are removing the inactive tag. It must be an identifier at the current hierarchy level.

statement Statement from which you are removing the inactive tag. It must be a statement at the current hierarchy level.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

deactivate | 5142

Deactivating and Reactivating Statements and Identifiers in a Device Configuration

activation-attribute (AAA Profile)

IN THIS SECTION

- Syntax | 1796
- Hierarchy Level | 1797
- Description | 1797
- Required Privilege Level | 1797
- Release Information | 1797

Syntax

```
activation-attribute {
    <code numeric-code;>
    <vendor-id vendor-id;>
}
```

Hierarchy Level

```
[edit unified-edge aaa profiles aaa-profile-name radius policy]
```

Description

Configure the RADIUS attribute that you want to carry the PCC rulebase name for rulebase activations from the RADIUS policy server to the MX Series router. By default, the rulebase name is carried in the ERX-Service-Activate Juniper vendor-specific attribute.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring an AAA Profile

Understanding AAA Profiles

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

activation-delay

IN THIS SECTION

- Syntax | 1798
- Hierarchy Level | 1798
- Description | 1798

- Options | 1798
- Required Privilege Level | 1798
- Release Information | 1798

Syntax

```
activation-delay seconds;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options]
```

Description

(J Series Services Routers) For ISDN interfaces, configure the ISDN dialer activation delay. Used only for dialer backup and dialer watch cases.

Options

seconds—Interval before the backup interface is activated after the primary interface has gone down.

- **Range:** 1 through 4,294,967,295 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

activation-priority

IN THIS SECTION

- Syntax | 1799
- Hierarchy Level | 1799
- Description | 1799
- Options | 1799
- Required Privilege Level | 1800
- Release Information | 1800

Syntax

```
activation-priority priority;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number dynamic-call-admission-control],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
dynamic-call-admission-control]
```

Description

(J4350 and J6350 Services Routers supporting voice over IP with the TGM550 media gateway module) For Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation, configure the dynamic call admission control (dynamic CAC) activation priority value.

Options

priority—The activation priority in which the interface is used for providing call bandwidth. The interface with the highest activation priority value is used as the primary link for providing call bandwidth. If the primary link becomes unavailable, the TGM550 switches over to the next active interface with the highest activation priority value, and so on.

- **Range:** 0 through 255
- **Default:** 50

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Junos OS Services Interfaces Library for Routing Devices

active

IN THIS SECTION

- Syntax | 1800
- Hierarchy Level | 1801
- Description | 1801
- Default | 1801
- Required Privilege Level | 1801
- Release Information | 1802

Syntax

```
(active | passive);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options (aggregate | generate | static) (defaults | route)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name (aggregate | generate | static) (defaults | route)],
[edit logical-systems logical-system-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-instances routing-instance-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib routing-table-name (aggregate | generate | static) (defaults | route)]
```

Description

Determine whether static, aggregate, or generated routes are removed from the routing and forwarding tables when they become inactive. Static routes are only removed from the routing table if the next hop becomes unreachable. This can occur if the local or neighbor interface goes down. Routes that have been configured to remain continually installed in the routing and forwarding tables are marked with **reject** next hops when they are inactive.

- **active**—Remove a route from the routing and forwarding tables when it becomes inactive.
- **passive**—Have a route remain continually installed in the routing and forwarding tables even when it becomes inactive.

Include the `active` statement when configuring an individual route in the **route** portion of the `static` statement to override a **passive** option specified in the **defaults** portion of the statement.

Default

active

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation

Example: Configuring a Conditional Default Route Policy

active-directory-access

IN THIS SECTION

- Syntax | 1802
- Hierarchy Level | 1803
- Description | 1804
- Options | 1804
- Required Privilege Level | 1804
- Release Information | 1804

Syntax

```
active-directory-access {
    authentication-entry-timeout (Services User Identification) minutes;
    domain name {
        domain-controller domain-controller-name {
            address domain-controller-address;
        }
        ip-user-mapping {
            discovery-method {
                wmi {
                    event-log-scanning-interval seconds;
                    initial-event-log-timespan hours;
                }
```

```
            }
        }
        user (System Services){
            user-name;
            password password;
        }
        user-group-mapping {
            ldap {
                address name {
                    port port;
                }
                authentication-algorithm simple;
                base base;
                ssl;
                user {
                    user-name;
                    password password;
                }
            }
        }
    }
    filter {
        exclude name;
        include name;
    }
    firewall-authentication-forced-timeout minutes;
    invalid-authentication-entry-timeout minutes;
    no-on-demand-probe;
    traceoptions (Active Directory Access) {
        file filename files files match match size size (world-readable | no-world-readable);
        flag name;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
    wmi-timeout seconds;
}
```

Hierarchy Level

```
[edit services user-identification]
```

Description

Identify the domain and domain controllers where the integrated user firewall feature is implemented; configure the IP address-to-user mapping information and the user-to-group mapping information for accessing the LDAP server.

Options

authentication-entry-timeout

Authentication entry timeout number.

- **Range:** 10 through 1440 minutes

firewall-authentication-forced-timeout

Firewall authentication fallback authentication entry forced timeout number.

- **Range:** 10 through 1440 minutes

invalid-authentication-entry-timeout

Invalid authentication entry timeout number.

- **Range:** 10 through 1440 minutes

no-on-demand-probe

Disable on-demand probe.

wmi-timeout

Windows Management Instrumentation (wmi) timeout number.

- **Range:** 3 through 120 seconds

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

user-identification (Services) | **30560**

LDAP Functionality in Integrated User Firewall

active-interface (VPLS Multihoming)

IN THIS SECTION

- Syntax | 1805
- Hierarchy Level | 1805
- Description | 1805
- Required Privilege Level | 1806
- Release Information | 1806

Syntax

```
active-interface {
    any;
    primary interface-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-homing],
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-homing site site-name],
[edit routing-instances instance-name protocols vpls multi-homing],
[edit routing-instances instance-name protocols vpls multi-homing site site-name]
```

Description

Specify a multihomed interface as the primary interface for the VPLS site. If there are multiple interfaces, the remaining interfaces are activated only when the primary interface goes down. If no active interfaces are configured at the site level, it is assumed that all traffic for a VPLS site travels through a single, nonmultihomed PE router.

> **NOTE**: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The remaining statements are explained separately. See CLI Explorer.

For FEC 128, use the `[edit routing-instances instance-name protocols vpls multi-homing]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Specifying an Interface as the Active Interface

active-leasequery (DHCP Local Server)

IN THIS SECTION

- Syntax | 1807
- Hierarchy Level | 1807
- Description | 1807
- Options | 1807
- Required Privilege Level | 1808
- Release Information | 1808

Syntax

```
active-leasequery {
             topology-discovery;
             peer-address {
                address;
               }
             timeout timeout;
             idle-timeout idle-timeout;
             exclude-interface interface-name;
       }
```

Hierarchy Level

```
[edit system services dhcp-local-server]
```

```
[edit system services dhcp-local-server dhcpv6]
```

Description

Subscriber redundancy on DHCP server focuses on subscriber synchronization between the peer servers using active leasequery. If you configure active leasequery for the DHCP server, the peer servers establishes TCP connection and sends active leasequery.

Options

idle-timeout *idle-timeout*

Configure the idle time out value in seconds for the connection between the peer servers.

- Range: 10 thru 3600 seconds
- Default: 60 seconds

peer-address *address*

Configure IPv4 or IPv6 address of the peer DHCP servers. You can configure maximum five IP addresses.

timeout *timeout*

Configure the TCP read/write block timeout value to terminate and restart the TCP connection.

- Range: 10 thru 3600 seconds
- Default: 60 seconds

topology-discovery Enable the topology discovery option.

exclude-interface *interface-name* (Optional) Exclude interface participating in active leasequery.

> **NOTE:**
> `exclude-interface` *`interface-name`*

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

active-leasequery (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 1809
- Hierarchy Level | 1809
- Description | 1809
- Options | 1809
- Required Privilege Level | 1811
- Release Information | 1811

Syntax

```
active-leasequery {
    idle-timeout seconds;
    peer-address address;
    timeout seconds;
    topology-discover;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name ...],
[edit routing-instances routing-instance-name ...]
```

Description

Enable support for active leasequery on a DHCPv4 or DHCPv6 relay agent. You can also configure parameters that the DHCP relay agent uses when sending DHCP active lease query messages to obtain lease information from the DHCP servers in the logical system/routing instance.

> **NOTE**: Stale-timer configuration is required to support an Active-Active lease query. This configuration optimizes the synchronization time when both the peers get the solicit packets at the same time.

Options

idle-timeout *seconds*

(Optional) Specify the number of seconds that pass with no activity before the TCP connection is terminated and restarted.

During active leasequery operations, binding information is updated only when it changes. Consequently, there are periods during which the server or peer relay agent sends no information; the connection is idle. If this period is longer than the idle-timeout, the connection is dropped. To avoid inappropriate connection drops, the server or peer relay agent sends DHCPLEASEACTIVE (DHCPv4) or LEASEQUERY-DATA (DHCPv6) messages

at intervals equal to one-half of the idle timeout period. These messages contain no binding information because they are sent when no updates are available. These messages keep the connection alive by serving as hello or keepalive messages signaling that the lack of activity is not a problem.

- **Range:** 10 through 3600
- **Default:** 60

peer-address *address*

(Optional) Specify the address of a peer relay agent to synchronize binding information. You can configure up to five peer addresses. If the subscriber configuration and interface names are identical on peer relay agents, then active leasequery provides 1:1, chassis-level redundancy between peers for binding information.

> **NOTE**: You can also achieve 1:1 chassis redundancy with an M:N redundancy configuration where:
>
> - You configure redundancy for all DHCP subscribers on the chassis
> - You configure a single chassis to back up all of the subscriber redundancy groups

timeout *seconds*

(Optional) Specify the number of seconds that TCP read/write operations can be blocked before the TCP connection is terminated and restarted.

- **Range:** 10 through 3600
- **Default:** 120

topology-discover

(Optional) Activate the topology discovery mechanism to discover peer access links for a redundancy group and their connections to the local access link. You configure this option on peer relay agents hosted on BNGs that are configured for M:N subscriber redundancy.

> **NOTE**: Configure interfaces participating in redundancy for a subset of subscribers and not for the entire chassis.
>
> Topology discovery is needed only for M:N subscriber redundancy. Topology discovery synchronizes only the subscribers stacked over configured interfaces.
>
> Chassis-to-chassis based redundancy provides the following limitations and advantages:

- Topology discovery cannot be configured.
- Supports both single-stack or dual-stack subscribers.
- Synchronizes all subscribers in the routing instance.
- It is recommended to configure chassis-to-chassis based redundancy for scaling.

1. Each peer relay agent sends topology discovery request messages over the TCP connection to each peer relay agent.
2. The receiving peer returns a topology discovery response message with information about its corresponding access interface.
3. The querying relay agent then updates its translation table to map the remote (peer) access interface with its local access interface.

Each relay agent peer initiates topology discovery for all of its access interfaces. Each peer then updates its own translation table according the information received in topology discovery responses.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

`topology-discover` option added in Junos OS Release 19.2R1.

active-policy

IN THIS SECTION

Syntax | 1812

- Hierarchy Level | 1812
- Description | 1812
- Options | 1812
- Required Privilege Level | 1812
- Release Information | 1812

Syntax

```
active-policy policy-name;
```

Hierarchy Level

```
[edit security idp]
```

Description

Specify which policy among the configured policies to activate.

Options

policy-name—Name of the active policy.

> **NOTE**: You need to make sure the active policy is enforced in the data plane.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Starting with Junos OS Release 18.2R1, IDP policy is directly assigned in the security policy rule. This is to simplify IDP policy usage and to provide flexibility to have multiple policies active at the same time. As a part of session interest check IDP will enabled if IDP policy is present in any of the matched rules. IDP policy is activated in security policies, by permitting the IDP policy within the application services using the `set security policies from-zone zone-name to-zone zone-name policy policy-name then permit application-services idp-policy idp-policy-name` command. Since IDP policy name is directly use in the security policy rule, the `[edit security idp active-policy policy-name]` statement is deprecated.

active-probe-params

IN THIS SECTION

- Syntax | 1813
- Hierarchy Level | 1814
- Description | 1814
- Options | 1815
- Required Privilege Level | 1816
- Release Information | 1816

Syntax

```
active-probe-params probe-name {
    settings {
        burst-size {
            size;
        }
        data-fill {
            string;
        }
        data-size {
            size;
        }
        dscp-code-points {
            dscp;
        }
```

```
        enable-sla-export {
            interval;
        }
        forwarding-class {
            forwarding-class-name;
        }
        loss-priority (low | high | medium-high| medium-low) {
        }
        per-packet-loss-timeout  {
            interval;
        }
        probe-count {
            count;
        }
        probe-interval {
            interval;
        }
    }
}
```

Hierarchy Level

[edit security advance-policy-based-routing]

Description

Specify parameter settings for an active probe.

Application performance is monitored and measured using active probes.

In active probing, custom packets are sent between a spoke device and a hub device on multiple routes to measure RTT, jitter, and packet loss between the book-ended points. You can configure to send active probes periodically on all the active and passive links.

Active probing starts after the configuration is committed. A configured number of samples are collected and used for measuring the SLA. If there is a violation detected for any application, the probe metrics are evaluated to determine the best possible link for that application traffic in order to meet performance requirements as in the SLA.

Consider the example, where you configure the probe count as 1000, probe interval as 10 seconds, and burst size as 100. Burst count is calculated as probe count/burst size (1000/100 = 10). Burst-count is 10. So, probes are sent in sets of 10 bursts each containing 100 packets.

Burst interval is calculated as probe interval/burst-count (10/10 =1). Burst interval is 1 second. So, a burst is sent every 1 second. The active probe is initiated from the spoke device to the hub device on each of the overlay path.

SaaS application types do not support ingress, egress jitter types.

The SaaS probe packets are HTTP head packets, sent between the spoke device and the SaaS server.

Options

probe-name Active probe identifier.

burst-size Number of probes sent as a burst. This value should be less than or equal to probe-count. The burst-size configuration is ignored for SaaS probing and always burst size will be used as one.

- **Range:** 1—100
- **Default:** 10

data-fill *string* Data payload for a probe packet. This is a hexadecimal string, which is used the payload for probe. Not supported for SaaS probing.

data-size *size* Size of the data portion. Not supported for SaaS probing.

dscp-code-points *dscp* DiffServ code point (DSCP) bits value.

enable-sla-export Time Interval (in seconds) at which the active probe data to be exported to controller. This option is disabled by default.

- **Range:** 60-600

forwarding-class *forwarding-class-name* Name of the forwarding class

- **Default:** network-control

loss-priority *level* Map packet values to a loss priority. Loss priority allows you to set the priority for dropping packets. Typically, you mark packets exceeding some service level with a high loss priority—that is, a greater likelihood of being dropped. Level can be one of the following:

- high—Packet has high loss priority.
- medium-high—Packet has medium-high loss priority.

- medium-low—Packet has medium-low loss priority.
- low—Packet has low loss priority.

per-packet-loss-timeout

Time interval between two consecutive SaaS probes.

- **Range:** 100 through 10000 milliseconds.
- **Default:** 1000 milliseconds.

probe-count *count*

Number of samples required to be collected for an SLA measurement. For SaaS applications, probing is to the actual server, we recommend to use three to four samples per probe window.

- **Range:** 1 through 1000
- **Default:** 5

probe-interval *interval*

Time interval between successive probes. Starting in Junos OS release 20.4R1, the upper limit of the probe interval is changed to 60 seconds. For SaaS applications, we recommend to configure 60 seconds interval to avoid aggressive probing to an actual SaaS server.

- **Range:** 1 through 60 seconds
- **Default:** 10 seconds

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1. The options `forwarding-class` and `loss-priority` are introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Application Quality of Experience

Advanced Policy-Based Routing

active-server-group

IN THIS SECTION

- Syntax | 1817
- Hierarchy Level | 1817
- Description | 1818
- Options | 1818
- Required Privilege Level | 1818
- Release Information | 1818

Syntax

```
active-server-group server-group-name <allow-server-change>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relaygroup group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay]
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6],
```

```
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name]
```

Description

Apply a DHCP relay agent configuration to the named group of DHCP server addresses. The server group itself is configured with the `server-group` statement. You can apply an active server group globally or for specific groups of interfaces, configured with the `group` statement. An active server group applied to an interface group overrides a global configuration.

Options

allow-server-change (Optional) (DHCPv4 only) Enable the relay agent to accept and forward a DHCP request (renew or rebind) ACK message to the client from any DHCP local server in the active server group. Starting in Junos OS Release 18.4R1, this option also applies to DHCP information request (DHCPINFORM) ACK messages.

- **Default:** Forward ACK messages from only the original binding server.

server-group-name Name of the group of DHCP or DHCPv6 server addresses to which the DHCP or DHCPv6 relay agent configuration applies.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

`allow-server-change` option added in Junos OS Release 16.2R1.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups

Configuring Group-Specific DHCP Relay Options

active-source-limit

IN THIS SECTION

- Syntax | 1819
- Hierarchy Level | 1819
- Description | 1820
- Default | 1820
- Options | 1820
- Required Privilege Level | 1820
- Release Information | 1820

Syntax

```
active-source-limit {
    log-interval seconds;
    log-warning value;
    maximum number;
    threshold number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name protocols msdp source ip-address/prefix-length],
[edit logical-systems logical-system-name routing-instances instance-name protocols msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
```

```
group group-name peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
source ip-address/prefix-length],
[edit protocols msdp],
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit protocols msdp source ip-address/prefix-length],
[edit routing-instances routing-instance-name protocols msdp],
[edit routing-instances routing-instance-name protocols msdp group group-name peer address],
[edit routing-instances routing-instance-name protocols msdp peer address],
[edit routing-instances routing-instance-name protocols msdp source ip-address/prefix-length]
```

Description

Limit the number of active source messages the routing device accepts.

Default

If you do not include this statement, the router accepts any number of MSDP active source messages.

Options

The options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP with Active Source Limits and Mesh Groups

activeness-probe

IN THIS SECTION

- Syntax | 1821
- Hierarchy Level | 1822
- Description | 1822
- Options | 1822
- Required Privilege Level | 1822
- Release Information | 1823

Syntax

```
activeness-probe {
 bfd-liveliness {
    destination-ip destination-ip;
    interface interface;
    minimum-interval milliseconds;
    multiplier multiplier;
    routing-instance routing-instance;
    source-ip source-ip;
 }
 dest-ip {
    ip-address;
    routing-instance routing-instance;
    src-ip src-ip;
 }
 minimum-interval minimum-interval;
 multiplier multiplier;
  }
```

Hierarchy Level

```
[edit chassis high-availability services-redundancy-group]
```

Description

Specify the probes for split-brain detection.

Options

bfd-liveliness— BFD-based probes to prevent split-brain scenarios.

dest-ip *ip-address*— destination IP address to send the probe requests. Use this option for ICMP based probes for split-brain detection.

- routing-instance *name*—Routing instance used by the probe requests.
- src-ip *ip-address*—Source IP address to send the probe requests.

> **NOTE**: ICMP-based split-brain probing and BFD-based split-brain probing are mutually exclusive.

minimal-interval *milliseconds*— (Optional) Time period between the probes sent to the destination IP address. You can configure 1000 milliseconds for minimal interval activeness probes. (supported in default gateway [switching] and hybrid deployments).

multiplier *number*— (Optional) Time period, after which the backup node transitions to the active state, if the backup node fails to receive response to the activeness-probes from the peer node. (supported in default gateway [switching] and hybrid deployments).

- **Range:** 2 through 15

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3R1

RELATED DOCUMENTATION

> *Example: Configure Multinode High Availability in a Layer 3 Network*
>
> No Link Title

actual-transit-statistics (Dynamic Profiles)

IN THIS SECTION

- Syntax | 1823
- Hierarchy Level | 1823
- Description | 1824
- Default | 1824
- Required Privilege Level | 1824
- Release Information | 1824

Syntax

```
actual-transit-statistics;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-interface-
unit]
```

Description

Enable `actual-transit-statistics` to enable the collection and reporting of accurate accounting statistics for subscribers. The accurate statistics provide packet-level and byte-level counts for all transmit and receive traffic forwarded across the specified interface. The counts do not include overhead bytes, filtered traffic, drops, discards, or control traffic. These accurate statistics match the accounting statistics used by RADIUS. You can display these statistics with the `show subscribers id` *`session-id`* `accounting-statistics` and `show subscribers interface` *`interface`* `accounting-statistics` commands.

> **NOTE**: Starting in Junos OS Release 18.4R1, you must enable `actual-transit-statistics` to collect subscriber statistics. If you do not configure this statement, subscriber statistics are not collected; the `show subscribers accounting-statistics` command displays a value of 0 for subscriber statistics; and the subscriber statistics are reported to RADIUS with values of zero.

Default

Disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

- *Enabling the Reporting of Accurate Subscriber Accounting Statistics to the CLI*
- show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet)
- *Dynamic Profiles Overview*

adaptive (protocols mpls)

IN THIS SECTION

- Syntax | 1825
- Hierarchy Level | 1825
- Description | 1825
- Default | 1825
- Required Privilege Level | 1826
- Release Information | 1826

Syntax

```
adaptive;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

During reroute, do not double-count bandwidth on links shared by the old and new paths. Including this statement causes RSVP to use shared explicit (SE) reservation styles and assists in smooth transition during rerouting.

Default

The configured object is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Adaptive LSP Configuration

adaptive

IN THIS SECTION

- Syntax | 1826
- Hierarchy Level | 1827
- Description | 1827
- Options | 1827
- Required Privilege Level | 1828
- Release Information | 1828

Syntax

```
adaptive {
    pps;
    scan-interval multiple;
    tolerance tolerance-percentage;
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options load-balance],
[edit dynamic-profiles name interfaces name logical-tunnel-options load-balance],
[edit dynamic-profiles name interfaces interface-range name aggregated-ether-options load-
balance],
[edit dynamic-profiles name interfaces interface-range name logical-tunnel-options load-balance],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-ether-options load-
balance],
[edit dynamic-profiles name logical-systems name interfaces name logical-tunnel-options load-
balance],
[edit dynamic-profiles name logical-systems name interfaces interface-range name aggregated-
ether-options load-balance],
[edit dynamic-profiles name logical-systems name interfaces interface-range name logical-tunnel-
options load-balance],
[edit interfaces name aggregated-ether-options load-balance],
[edit interfaces name logical-tunnel-options load-balance],
[edit interfaces interface-range              name aggregated-ether-options load-balance],
[edit interfaces interface-range              name logical-tunnel-options load-balance]
```

Description

Correct a genuine traffic imbalance by using a feedback mechanism to distribute the traffic across the links of an aggregated Ethernet bundle.

Options

pps

(PTX Series only) The type of traffic rate among the members of the AE bundle is measured packets per second. The default rate type is bytes per second.

scan-interval *multiple*

(PTX Series only) Scan interval, as a multiple of a 30-second interval.

- **Range:** 1 through 5
- **Default:** 1

tolerance *tolerance-percentage*

(MX Series and PTX Series) Limit to the variance in the packet traffic flow to the aggregated Ethernet links in a percentage.

- **Range:** 1 through 100 percent
- **Default:** 20 percent

Required Privilege Level

interface - To view this statement in the configuration.

interface-control - To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R3.

RELATED DOCUMENTATION

Understanding Aggregated Ethernet Load Balancing

Example: Configuring Aggregated Ethernet Load Balancing

adaptive (RSVP)

IN THIS SECTION

- Syntax | 1828
- Hierarchy Level | 1829
- Description | 1829
- Options | 1829
- Required Privilege Level | 1829
- Release Information | 1829

Syntax

```
adaptive {
limit threshold-limit {
threshold-percent threshold-percent;
threshold-value threshold-value;
```

```
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name update-
threshold],
[edit protocols rsvp interface interface-name update-threshold]
```

Description

Tune update-threshold dynamically. Use the `limit` option to set the upper limit of available bandwidth for a given range. Use the `threshold-percent` statement to set the percentage at which a change in reserved bandwidth triggers an IGP update. Use the `threshold-value` statement to set the change in reserved bandwidth that can trigger an IGP update.

Options

limit ***threshold-limit***	Specify the upper limit of available bandwidth for the range.
threshold-percent ***threshold-percent***	Specify the percentage change in reserved bandwidth to trigger IGP update.
threshold-value ***threshold-value***	Specify the change in reserved bandwidth to trigger IGP update. (is capped at 20% of link bandwidth).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

`adaptive` option introduced for RSVP protocols in Junos OS Release 22.1R1.

RELATED DOCUMENTATION

update-threshold | 30456

adaptive-sample-rate

IN THIS SECTION

- Syntax | 1830
- Hierarchy Level | 1830
- Description | 1830
- Default | 1831
- Options | 1831
- Required Privilege Level | 1831
- Release Information | 1831

Syntax

```
adaptive-sample-rate rate {
    fallback;
    sample-limit-threshold sample-limit-threshold;
}
```

Hierarchy Level

```
[edit protocols sflow]
```

Description

Configure and back up the adaptive sampling rate in sFlow monitoring technology configurations.

Routers and switches use adaptive sampling to ensure both sampling accuracy and efficiency. Adaptive sampling is a process of monitoring the overall incoming traffic rate on the network device and providing intelligent feedback to interfaces to dynamically adapt the sampling rates on interfaces on the basis of traffic conditions.

Default

Adaptive sampling is enabled by default.

You can use this `adaptive-sample-rate` configuration statement to change the default adaptive sampling from 300 pps to any other value.

Options

rate Configure the maximum number of samples that should be generated.

- **Range:** 300 through 9500 pps

> **NOTE**: You can configure the adaptive sampling rate on QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD devices to a maximum number of 30000 pps.

fallback Enable adaptive sampling fallback (also called reverse adaptive sampling).

sample-limit-threshold Explicitly specify the threshold value in packets per second (pps) for each line card. You must configure the `fallback` statement to be able to configure the `sample-limit-threshold` statement.

- **Range:** 0 through 4750 pps.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R2.

`adaptive-sample-rate` limit is increased in Junos OS Release 19.1R1.

`adaptive-sample-rate` limit is increased in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

sflow | **16907**

show sflow interface | 26485

Understanding How to Use sFlow Technology for Network Monitoring

Example: Monitoring Network Traffic Using sFlow Technology

adaptive-services

IN THIS SECTION

- Syntax | 1832
- Hierarchy Level | 1833
- Description | 1833
- Options | 1833
- Required Privilege Level | 1833
- Release Information | 1833

Syntax

```
adaptive-services {
    service-package {
        extension-provider {
            control-cores control-number;
            data-cores data-number;
            data-flow-affinity {
                hash-key (layer-3 | layer-4);
            }
            forwarding-db-size size;
            object-cache-size size;
            package package-name;
            policy-db-size size;
            syslog {
                facility {
                    severity;
                    destination destination;
                }
```

```
            }
            wired-max-processes num-procs;
            wired-process-mem-size mem-size;
        }
        layer-2;
        layer-3;
    }
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number]
```

Description

Enable a service package on adaptive services interfaces.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Service Packages

Multiservices MIC and Multiservices MPC (MS-MIC and MS-MPC) Overview

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

adaptive-services-pics

IN THIS SECTION

- Syntax | 1834
- Hierarchy Level | 1834
- Description | 1834
- Options | 1834
- Required Privilege Level | 1835
- Release Information | 1835

Syntax

```
adaptive-services-pics {
    traceoptions {
        file filename <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Define global services properties.

Options

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4. The `file` option was added in Release 8.0.

RELATED DOCUMENTATION

Tracing Services PIC Operations

adaptive-shaper

IN THIS SECTION

- Syntax | 1835
- Hierarchy Level | 1835
- Description | 1836
- Options | 1836
- Required Privilege Level | 1836
- Release Information | 1836

Syntax

```
adaptive-shaper adaptive-shaper-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Assign an adaptive shaper to an interface.

Adaptive shapers enable bandwidth limits on Frame Relay interfaces when the device receives frames containing the backward explicit congestion notification (BECN) bit.

Options

adaptive-shaper-name—Name of the adaptive shaper.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

adaptive-shapers | 1836

adaptive-shapers

IN THIS SECTION

- Syntax | 1837
- Hierarchy Level | 1837
- Description | 1837
- Options | 1837
- Required Privilege Level | 1837
- Release Information | 1837

Syntax

```
adaptive-shapers {
    adaptive-shaper-name {
        trigger type shaping-rate (percent percentage | rate);
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Define trigger types and associated rates. Adaptive shapers enable bandwidth limits on Frame Relay interfaces when the Services Router receives frames containing the backward explicit congestion notification (BECN) bit.

Options

adaptive-shaper-name—Name of the adaptive shaper.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface— view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

adaptive-shaper | **1835**

add-attribute-length-in-pdu

IN THIS SECTION

- Syntax | 1838
- Hierarchy Level | 1838
- Description | 1838
- Required Privilege Level | 1839
- Release Information | 1839

Syntax

```
add-attribute-length-in-pdu;
```

Hierarchy Level

```
[edit protocols mvrp]
```

Description

Add an extra byte in protocol data units (PDUs) sent by the Multiple VLAN Registration Protocol (MVRP) in Junos OS Releases 11.3 and later for EX Series switches that do not support the Enhanced Layer 2 Software (ELS). By default, this MVRP version does not include the extra byte. You can add the extra byte in this MVRP version to address an incompatibility issue with the following MVRP versions:

- MVRP in Junos OS Releases 11.2 and earlier, which includes the extra byte.
- MVRP in Junos OS Releases 13.2 and later for EX Series switches with support for ELS, which includes the extra byte.

If this incompatibility issue arises, the MVRP versions that include the extra byte do not recognize PDUs that do not include the extra byte.

You can recognize an MVRP version problem by looking at a switch running an MVRP version that includes the extra byte. Because a switch running an MVRP version that includes the extra byte cannot

interpret an unmodified PDU from an MVRP version that does not include the extra byte, the switch will not add VLANs from the MVRP version that does not include the extra byte. When you execute the command `show mvrp statistics` on the MVRP version that includes the extra byte, the values for *Join Empty received* and *Join In received* will incorrectly display zero, even though the value for *MRPDU received* has been increased. Another indication that MVRP is having a version problem is that unexpected VLAN activity, such as multiple VLAN creation, takes place on the switch running the MVRP version that includes the extra byte.

Required Privilege Level

routing—To view this statement in the configuration.
routing control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Multiple VLAN Registration Protocol (MVRP) on Switches

Understanding Multiple VLAN Registration Protocol (MVRP)

add-path

IN THIS SECTION

- Syntax | 1840
- Hierarchy Level | 1840
- Description | 1841
- Options | 1841
- Required Privilege Level | 1841
- Release Information | 1841

Syntax

```
add-path {
    receive;
    send {
        include-backup-path backup_path_number;
        multipath;
        path-count number;
        path-selection-mode {
            (all-paths | equal-cost-paths);
        }
        prefix-policy [ policy-names ];
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp group group-name family family family-
modifier],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
family-modifier],
[edit logical-systems logical-system-name routing-instances name protocols bgp group group-name
family family family-modifier],
[edit logical-systems logical-system-name routing-instances name protocols bgp group group-name
neighbor address family family-modifier],
[edit logical-systems logical-system-name tenants name routing-instances name protocols bgp
group group-name family family family-modifier],
[edit logical-systems logical-system-name tenants name routing-instances name protocols bgp
group group-name neighbor address family family-modifier],
[edit protocols bgp group group-name family family-modifier],
[edit protocols bgp group group-name neighbor address family family-modifier],
[edit tenants name routing-instances name protocols bgp group group-name family family-
modifier ],
[edit tenants name routing-instances name protocols bgp group group-name neighbor address family
family-modifier]
```

Description

Enable advertisement and/or reception of multiple paths to a destination to/from the same BGP peer, instead of advertising/receiving only the active path to/from the same BGP peer.

> **NOTE**: The minimum configuration for the send side is `add-path send path-count <n>`. The configuration for the receive side is `add-path send receive`.
>
> You cannot just configure `add-path`. You must configure at least one `send/receive`. (If you attempt to configure only `add-path`, it will fail.) If only `send/receive` is configured, the corresponding support is enabled: `send` for advertisement of add-path routes and `receive` for reception of add-path routes.

> **NOTE**: The effective `receive/send add-path` state of a BGP session depends on the `add-path` capability advertisement from both the ends. That is, if R1 configures `send`, but R2 doesn't configure `receive`, then R1 will not send `add-path` to R2.

> **NOTE**: On PTX1000 devices, the `add-path` configuration statement is not supported under routing instances.

Options

receive Enable the router to receive multiple paths to a destination. You can enable the router to receive multiple paths from specified neighbors or from all neighbors.

The remaining statement is explained separately. See CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

Support for range from 2 through 20 (for BGP) introduced in Junos OS Release 14.1.

`multipath` option introduced in Junos OS Release 16.1R2.

Support for 64 BGP add-paths introduced in Junos OS Release 18.4R1 for the MX Series.

`include-backup-path` and `path-selection-mode` options introduced in Junos OS Release 18.4R1 for the MX Series and PTX Series.

RELATED DOCUMENTATION

Understanding the Advertisement of Multiple Paths to a Single Destination in BGP

Example: Advertising Multiple Paths in BGP

multipath (Protocols BGP) | **11110**

add-path-display-ipv4-address

IN THIS SECTION

- Syntax | **1842**
- Hierarchy Level | **1842**
- Description | **1843**
- Required Privilege Level | **1843**
- Release Information | **1843**

Syntax

```
add-path-display-ipv4-address { ... }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit protocols bgp],
[edit protocols bgp group group-name]
```

Description

Enable the display of external BGP (EBGP) *path-id* in IPv4 Address format.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding the Advertisement of Multiple Paths to a Single Destination in BGP

Example: Advertising Multiple Paths in BGP

Example: Configuring EBGP Multihop

address (tunnel-encapsulation)

IN THIS SECTION

- Syntax | 1844
- Hierarchy Level | 1844
- Description | 1844
- Required Privilege Level | 1844
- Release Information | 1844

Syntax

```
address address;
```

Hierarchy Level

```
[edit interfaces name unit name tunnel encapsulation( gre | ipip | udp ) destination]
```

Description

Interface address prefix for tunnel destination.

Required Privilege Level

interface

Release Information

Statement introduced for GRE tunnels in Junos OS Evolved Release 20.2R1.

Statement introduced for IPIP tunnels in Junos OS Evolved Release 20.4R1.

Statement introduced for UDP tunnels in Junos OS Evolved Release 21.1R1.

address (SNMP)

IN THIS SECTION

- Syntax | 1845
- Hierarchy Level | 1845
- Description | 1845
- Options | 1845
- Required Privilege Level | 1845
- Release Information | 1845

Syntax

```
address address;
```

Hierarchy Level

```
[edit snmp v3 target-address target-address-name]
```

Description

Specify the SNMP target address for receiving traps or informs.

Options

address—IPv4 address of the system to receive traps or informs. You must specify an address, not a hostname.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Address

address

IN THIS SECTION

- Syntax | 1846
- Hierarchy Level | 1847
- Description | 1847
- Options | 1848
- Required Privilege Level | 1849
- Release Information | 1849

Syntax

```
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    subnet-router-anycast;
    ndp;
    primary-only;
    multipoint-destination address dlci dlci-identifier;
    multipoint-destination address {
        epd-threshold cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (disable | seconds);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
rate burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
```

```
    }
    primary;
    preferred;
    virtual-gateway-address address;
    (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority-number number;
        track {
            priority-cost seconds;
            priority-hold-time interface-name {
                interface priority;
                bandwidth-threshold bits-per-second {
                    priority;
                }
            }
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family  family]
```

Description

Configure the interface address.

NOTE: If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration, and the remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/8;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
     family inet {
       address 192.168.1.1/8;
     }
  }
}
```

For more information on configuring the same address on multiple interfaces, see Configuring the Interface Address.

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.

Options

address—Address of the interface.

subnet-router-anycast—IPv6 host address to communicate with any of the routers present on the link.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

NOTE: The `edit logical-systems` hierarchy is not available on QFabric systems.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring the Protocol Family
- family
- negotiate-address
- unnumbered-address (Ethernet) | **30425**

address (IP Address)

IN THIS SECTION

- Syntax MX Series and EX Series (dynamic-profiles) | **1850**
- Syntax QFX Series and QFabric (interfaces) | **1850**
- MX Series and EX Series (dynamic-profiles) | **1851**
- QFX Series and QFabric (interfaces) | **1851**
- Description | **1851**
- Options | **1851**
- Required Privilege Level | **1852**
- Release Information | **1852**

Syntax MX Series and EX Series (dynamic-profiles)

```
address (ip-address | ipv6-address);
```

Syntax QFX Series and QFabric (interfaces)

```
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    reui-64;
    master-only;
    multipoint-destination addressdlci dlci-identifier;
    multipoint-destination address {
        epd-threshold cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (disable | seconds);
        shaping {
            (cbr rate |rtvbr peak rate sustained rate burst length |vbr peak rate sustained rate
burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
    primary;
    preferred;
    (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
```

```
        priority-number number;
        track {
            priority-cost seconds;
            priority-hold-time interface-name {
                interface priority;
                bandwidth-threshold bits-per-second {
                    priority;
                }
            }
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
    }
}
```

MX Series and EX Series (dynamic-profiles)

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family],
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family],
[edit dynamic-profiles profile-name interfaces pp0 unit  "$junos-interface-unit" family family],
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family]
```

QFX Series and QFabric (interfaces)

```
[edit interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family  family]
```

Description

Configure the interface address.

Options

ip-address—IPv4 address of the interface.

ipv6-address—IPv6 address of the interface. When configuring an IPv6 address on a dynamically created interface, use the *$junos-ipv6-address* dynamic variable.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the `[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family]` hierarchy level introduced in Junos OS Release 10.1.

Support at the `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy level introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

RELATED DOCUMENTATION

- Configuring the Protocol Family
- Format for Specifying IP Addresses, Network Masks, and Prefixes in Junos OS Configuration Statements

address (Address Book Next Gen Services)

IN THIS SECTION

- Syntax | 1853
- Hierarchy Level | 1853
- Description | 1853
- Options | 1853
- Required Privilege Level | 1853
- Release Information | 1853

Syntax

```
address address-name range-address lower-limit to upper-limit
```

Hierarchy Level

```
[edit services address-book global]
```

Description

Configure a range of addresses that can be referenced in the `match` stanza of a NAT rule.

Options

lower-limit	The lower end of the address range.
upper-limit	The upper end of the address range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

address-allocation

IN THIS SECTION

- Syntax | 1854
- Hierarchy Level | 1854

Description | 1854
Required Privilege Level | 1854
Release Information | 1854

Syntax

```
address-allocation round-robin;
```

Hierarchy Level

```
[edit services nat pool pool-name]
```

Description

When you use round-robin allocation, one port is allocated from each address in a range before repeating the process for each address in the next range. After ports have been allocated for all addresses in the last range, the allocation process wraps around and allocates the next unused port for addresses in the first range.

Regardless of whether the round-robin method of allocation is addresses is enabled by using the `address-allocation round-robin` statement, round-robin allocation is enabled by default on MS-MICs and MS-MPCs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Pools of Addresses and Ports for Network Address Translation Overview

address (Analytics Collector)

IN THIS SECTION

- Syntax | 1855
- Hierarchy Level | 1855
- Description | 1855
- Options | 1856
- Required Privilege Level | 1856
- Release Information | 1856

Syntax

```
address ip-address {
    port number {
        transport protocol {
            export-profile profile-name;
        }
    }
}
```

Hierarchy Level

```
[edit services analytics collector]
```

Description

Configure the address of a remote server to receive streamed analytics (queue and traffic statistics) data.

NOTE: The `address` statement is not available in Junos OS Releases prior to 13.2X51-D15.

Options

ip-address	IP address of the remote server receiving the streamed data.
port *number*	Port number of the remote server receiving the streaming data.
export-profile *profile-name*	Name of the export profile containing the parameters for the analytics data being streamed.

transport *protocol*	A transport protocol used to stream data to the port.
	• Values:
	  • **tcp**—Transmission Control Procol (TCP)
	  • **udp**—User Datagram Protocol (UDP)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

address (Anycast RPs)

IN THIS SECTION

- Syntax | 1857
- Hierarchy Level | 1857

- Description | 1857
- Options | 1857
- Required Privilege Level | 1857
- Release Information | 1858

Syntax

```
address address <forward-msdp-sa>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp local (inet | inet6) anycast-pim rp-set],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp local (inet | inet6) anycast-pim rp-set],
[edit protocols pim rp local (inet | inet6) anycast-pim rp-set],
[edit routing-instances routing-instance-name protocols pim rp local (inet | inet6) anycast-pim rp-
set]
```

Description

Configure the anycast rendezvous point (RP) addresses in the RP set. Multiple addresses can be configured in an RP set. If the RP has peer Multicast Source Discovery Protocol (MSDP) connections, then the RP must forward MSDP source active (SA) messages.

Options

address—RP address in an RP set.

forward-msdp-sa—(Optional) Forward MSDP SAs to this address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

address-assignment (Access)

IN THIS SECTION

- Syntax | 1858
- Hierarchy Level | 1861
- Description | 1861
- Options | 1862
- Required Privilege Level | 1862
- Release Information | 1862

Syntax

```
address-assignment {
    abated-utilization percentage;
    abated-utilization-v6 percentage;
    high-utilization percentage;
    high-utilization-v6 percentage;
    neighbor-discovery-router-advertisement ndra-name;
    pool pool-name {
        family {
            inet {
                dhcp-attributes {
                    boot-file boot-file-name;
                    boot-server boot-server-name;
                    domain-name domain-name;
                    grace-period seconds;
                    maximum-lease-time (seconds | infinite);
                    name-server ipv4-address;
                    netbios-node-type (b-node | h-node | m-node | p-node);
                    next-server next-server-name;
```

```
            option dhcp-option-identifier-code {
                array {
                    byte [8-bit-value];
                    flag [ false| off |on |true];
                    integer [32-bit-numeric-values];
                    ip-address [ip-address];
                    short [signed-16-bit-numeric-value];
                    string [character string value];
                    unsigned-integer [unsigned-32-bit-numeric-value];
                    unsigned-short [16-bit-numeric-value];
                }
                byte 8-bit-value;
                flag (false | off | on | true);
                integer 32-bit-numeric-values;
                ip-address ip-address;
                short signed-16-bit-numeric-value;
                string character string value;
                unsigned-integer unsigned-32-bit-numeric-value;
                unsigned-short 16-bit-numeric-value;
            }
            option-match {
                option-82 {
                    circuit-id match-value {
                        range range-name;
                    }
                    remote-id match-value;
                        range range-name;
                    }
                }
            }
            propagate-ppp-settings [interface-name];
            propagate-settings interface-name;
            router ipv4-address;
            server-identifier ip-address;
            sip-server {
                ip-address ipv4-address;
                name sip-server-name;
            }
            tftp-server server-name;
            wins-server ipv4-address;
        }
        excluded-address;
        excluded-range range-name
```

```
                high upper-limit;
                low lower-limit;
            }
            range range-name {
                high upper-limit;
                low lower-limit;
            }
            host hostname {
                hardware-address mac-address;
                ip-address reserved-address;
                user-name;
            }
            network network address;
            xauth-attributes {
                primary-dns ip-address;
                primary-wins ip-address;
                secondary-dns ip-address;
                secondary-wins ip-address;
            }
        }
        inet6 {
            dhcp-attributes {
                dns-server ipv6-address;
                grace-period seconds;
                maximum-lease-time (seconds | infinite);
                option dhcp-option-identifier-code {
                    array {
                        byte [8-bit-value];
                        flag [ false| off |on |true];
                        integer [32-bit-numeric-values];
                        ip-address [ip-address];
                        short [signed-16-bit-numeric-value];
                        string [character string value];
                        unsigned-integer [unsigned-32-bit-numeric-value];
                        unsigned-short [16-bit-numeric-value];
                    }
                    byte 8-bit-value;
                    flag (false | off | on | true);
                    integer 32-bit-numeric-values;
                    ip-address ip-address;
                    short signed-16-bit-numeric-value;
                    string character string value;
                    unsigned-integer unsigned-32-bit-numeric-value;
```

```
                        unsigned-short 16-bit-numeric-value;
                    }
                    propagate-ppp-settings [interface-name];
                    sip-server-address ipv6-address;
                    sip-server-domain-name domain-name;
                }
                excluded-address;
                excluded-range range-name
                    high upper-limit;
                    low lower-limit;
                }
                host hostname {
                    hardware-address mac-address;
                    ip-address reserved-address;
                    user-name;
                }
                prefix ipv6-network-prefix;
                range range-name {
                    high upper-limit;
                    low lower-limit;
                    prefix-length delegated-prefix-length;
                }
                xauth-attributes {
                    primary-dns-ipv6;
                    secondary-dns-ipv6;
                }
            }
        link pool-name;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

The address-assignment pool feature enables you to create different pools with different attributes. For example, multiple client applications, such as DHCPv4 or DHCPv6, can use an address-assignment pool to provide addresses for their particular clients.

Options

- `host` *`hostname`*—Name by which a network-attached device is known on a network.
- `hardware-address` *`mac-address`*—Specify the MAC address of the client. This is the hardware address that identifies the client on the network.
- `ip-address` *`reserved-address`*—Specify the reserved IP address.
- `user-name`—Specify username or IKE ID.
- `xauth-attributes`—Specify XAuth attributes to use in XAuth authentication.
- `primary-dns-ipv6`—Specify the primary-dns IPv6 address.
- `secondary-dns-ipv6`—Specify the secondary-dns IPv6 address.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`xauth-attributes` option under `inet6` is introduced in Junos OS Release 20.3R1.

`user-name` option under `inet host` *`host-name`* is introduced in Junos OS Release 20.3R1.

`host` option under `inet6` is introduced in Junos OS Release 20.3R1.

address-assignment (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 1863
- Hierarchy Level | 1864
- Description | 1864

- Options | 1864
- Required Privilege Level | 1865
- Release Information | 1865

Syntax

```
address-assignment {
    abated-utilization percentage;
    abated-utilization-v6 percentage;
    high-utilization percentage;
    high-utilization-v6 percentage;
    neighbor-discovery-router-advertisement ndra-pool-name;
    pool pool-name {
        active-drain;
        family family {
            dhcp-attributes {
                protocol-specific attributes;
            }
            excluded-address ip-address;
            excluded-range name low minimum-value high maximum-value;
            host hostname {
                hardware-address mac-address;
                ip-address ip-address;
            }
            network ip-prefix/<prefix-length>;
            prefix ipv6-prefix;
            range range-name {
                high upper-limit;
                low lower-limit;
                prefix-length prefix-length;
            }
        }
        hold-down;
        link pool-name;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure address-assignment pools that can be used by different client applications.

> **NOTE**: Support for subordinate statements is platform-specific. See individual statement topics for support information.

Options

abated-utilization

Generate SNMP traps for DHCP address pools or linked set of address pools. No SNMP traps are generated unless a value is configured. Default: Abated utilization is not set. Delete the abated-utilization value to unset.

- **Values:** *percentage*—Threshold below which an SNMP trap clear is generated. Range: 1 through 98.

abated-utilization-v6

Generate SNMP traps for DHCPv6 address pools or linked set of address pools. No SNMP traps are generated unless a value is configured. Default: Abated utilization is not set. Delete the abated-utilization value to unset.

- **Values:** *percentage*—Threshold below which an SNMP trap clear is generated. Range: 1 through 98.

high-utilization

Generate an SNMP trap when the DHCP address pool or linked set of address pools use surpasses the specified percentage. Default: High utilization is not set. Delete the high-utilization value to unset.

- **Values:** *percentage*—Percentage used to generate a trap. Range: 2 through 99.

high-utilization-v6

Generate an SNMP trap when the DHCPv6 address pool or linked set of address pools use surpasses the specified percentage. Default: High utilization is not set. Delete the high-utilization value to unset.

- **Values:** *percentage*—Percentage used to generate a trap. Range: 2 through 99.

neighbor-discovery-router-advertisement

Configure the name of the address-assignment pool used to assign the router advertisement prefix.

- **Values:** *ndra-pool-name*—Name of the address-assignment pool.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management

L2TP LNS Inline Service Interfaces

Configuring an Address-Assignment Pool Used for Router Advertisements

address (Bidirectional Rendezvous Points)

IN THIS SECTION

- Syntax | 1866
- Hierarchy Level | 1866
- Description | 1866
- Options | 1866
- Required Privilege Level | 1866
- Release Information | 1867

Syntax

```
address address {
    group-ranges {
        destination-ip-prefix</prefix-length>;
    }
    hold-time seconds;
    priority number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp bidirectional],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp bidirectional],
[edit protocols pim rp bidirectional],
[edit routing-instances routing-instance-name protocols pim rp bidirectional]
```

Description

Configure bidirectional rendezvous point (RP) addresses. The address can be a loopback interface address, an address of a link interface, or an address that is not assigned to an interface but belongs to a subnet that is reachable by the bidirectional PIM routers in the network.

Options

address—Bidirectional RP address.

- **Default:** 232.0.0.0/8

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Bidirectional PIM

Example: Configuring Bidirectional PIM

address-blacklist

IN THIS SECTION

- Syntax | 1867
- Hierarchy Level | 1867
- Description | 1867
- Required Privilege Level | 1868
- Release Information | 1868

Syntax

```
address-blacklist list-name;
```

Hierarchy Level

```
[edit security utm feature-profile anti-spam]
[edit security utm default-configuration]
```

Description

Enter an address blocklist (or allowlist) custom object for local list spam filtering.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

address-book

IN THIS SECTION

- Syntax | 1868
- Hierarchy Level | 1869
- Description | 1869
- Options | 1869
- Required Privilege Level | 1870
- Release Information | 1870

Syntax

```
address-book (book-name | global) {
    address (Security Address Book) address-name {
        ip-prefix {
            description text;
        }
        description text;
        dns-name domain-name {
            ipv4-only;
            ipv6-only;
        }
```

```
        range-address lower-limit to upper-limit;
        wildcard-address ipv4-address/wildcard-mask;
    }
    address-set address-set-name {
        address address-name;
        address-set address-set-name;
        description text;
    }
    attach {
        zone zone-name;
    }
    description (Security Address Book) text;
}
```

Hierarchy Level

```
[edit security]
[edit tenants tenant-name security]
```

Description

Define entries in the address book. Address book entries can include any combination of IPv4 addresses, IPv6 addresses, DNS names, wildcard addresses, and address range. You define addresses and address sets in an address book and then use them when configuring different features, such as security policies and NAT.

NOTE: IPv6 wildcard address configuration is not supported in this release.

Options

- `address-book` *`book-name`*–Name of the address book.
- `global`–An address book that is available by default. You can add any combination of IPv4 addresses, IPv6 addresses, wildcard addresses, DNS names, or address range to the global address book. You do not need to attach the global address book to a security zone; entries in the global address book are available to all security zones that are not attached to address books.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

tenants—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 18.3R1 for tenant systems. Support for wildcard addresses added in Junos OS Release 11.1. Statement moved under the security hierarchy in Junos OS Release 11.2. Support for address range added in Junos OS Release 12.1. The `description` option added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Address Books

Understanding Address Sets

address-book (System)

IN THIS SECTION

- Syntax | **1871**
- Hierarchy Level | **1871**
- Description | **1871**
- Options | **1871**
- Required Privilege Level | **1872**
- Release Information | **1872**

Syntax

```
address-book {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile security-profile-name]
```

Description

Address book entries include any combination of IPv4 addresses, IPv6 addresses, DNS names, wildcard addresses, and address range. You define addresses and address sets in an address book and then use those addresses when configuring different features, such as security policies and NAT.

Specify the number of address books that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

- `maximum` *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- `reserved` *`amount`*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Address Books

address-change-immediate-update

IN THIS SECTION

- Syntax | 1872
- Hierarchy Level | 1873
- Description | 1873
- Default | 1873
- Required Privilege Level | 1873
- Release Information | 1873

Syntax

```
address-change-immediate-update;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router to send an Interim-Accounting message to the RADIUS server immediately after on-demand IPv4 allocation and de-allocation.

Changes to this setting take effect for new subscriber logins. Existing subscribers are not impacted by this change except when the AAA daemon restarts.

Default

This functionality is disabled by default.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Enabling Immediate Interim Accounting Messages for On-Demand IPv4 Address Changes

Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation

address (Diameter Peer)

IN THIS SECTION

- Syntax | 1874
- Hierarchy Level | 1874
- Description | 1874
- Options | 1874
- Required Privilege Level | 1874
- Release Information | 1875

Syntax

```
address ip-address;
```

Hierarchy Level

```
[edit access diameter peer peer-name]
[edit diameter peer peer-name]
```

Description

Configure the IP address for the Diameter remote peer.

Options

address—IP address for the Diameter peer.

ip-address—IPv4 or IPv6 address of remote Diameter peer. The address family must match that for the local Diameter transport connection.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for IPv6 addresses added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

address (Diameter Transport)

IN THIS SECTION

- Syntax | 1875
- Hierarchy Level | 1875
- Description | 1876
- Options | 1876
- Required Privilege Level | 1876
- Release Information | 1876

Syntax

```
address ip-address;
```

Hierarchy Level

```
[edit diameter transport transport-name]
```

Description

Configure the source (local) IP address for the Diameter local transport connection.

Options

ip-address—IPv4 or IPv6 address of remote Diameter peer. The address family must match that for the remote Diameter peer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support for IPv6 addresses added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring Diameter

address-family

IN THIS SECTION

- Syntax | 1877
- Hierarchy Level | 1877
- Description | 1877
- Required Privilege Level | 1877
- Release Information | 1877

Syntax

```
address-family {
    inet {
        address;
        table table-name;
    }
    ccc {
        interface-name;
        standby;
        peer-unit unit-number;
        table table-name;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options condition if-route-exists],
[edit policy-options condition if-route-exists]
```

Description

Specify that the route must correspond to certain prefix type to be considered a match.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring Pseudowire Redundancy in a Mobile Backhaul Scenario

address (Interfaces)

IN THIS SECTION

- Syntax | 1878
- Hierarchy Level | 1878
- Description | 1878
- Options | 1878
- Required Privilege Level | 1878
- Release Information | 1879

Syntax

```
address address {
    destination address;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family]
```

Description

Configure the interface address.

Options

address—Address of the interface.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Encryption Interfaces

address (Interfaces, Flow Monitoring)

IN THIS SECTION

- Syntax | 1879
- Hierarchy Level | 1879
- Description | 1880
- Options | 1880
- Required Privilege Level | 1880
- Release Information | 1880

Syntax

```
address address {
    destination address;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-numberfamily family]
```

Description

Configure the interface address.

Options

address—Address of the interface.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Flow Monitoring

Configuring Traffic Sampling on MX, M and T Series Routers

address (L2TP Destination)

IN THIS SECTION

- Syntax | 1881
- Hierarchy Level | 1881
- Description | 1881
- Options | 1881
- Required Privilege Level | 1881
- Release Information | 1881

Syntax

```
address ip-address {
    access-line-information <connection-speed-update>;
    drain;
    routing-instance routing-instance-name {
        drain;
    }
}
```

Hierarchy Level

```
[edit services l2tp destination]
```

Description

Specify the IP address and other attributes for the L2TP destination.

Options

ip-address—IP address of the destination; corresponds to the IP address that is used by LACs to identify the LNS.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

address (L2TP Tunnel Destination)

IN THIS SECTION

- Syntax | 1882
- Hierarchy Level | 1882
- Description | 1882
- Options | 1883
- Required Privilege Level | 1883
- Release Information | 1883

Syntax

```
address ip-address; {
    drain;
    routing-instance routing-instance-name {
        drain;
    }
}
```

Hierarchy Level

```
[edit services l2tp tunnel name name ]
```

Description

Specify the IP address for the L2TP tunnel destination when the `name` statement at the `[edit services l2tp tunnel]` hierarchy level specifies only the name of the tunnel rather than both the name and the

destination address. Do not include the `address` statement when the `name` statement provides both the tunnel name and the destination address.

Options

`ip-address`—IP address of the tunnel destination.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

address (LNS Local Gateway)

IN THIS SECTION

- Syntax | 1884
- Hierarchy Level | 1884
- Description | 1884
- Options | 1884
- Required Privilege Level | 1884
- Release Information | 1884

Syntax

```
address address;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name local-gateway]
```

Description

Specify the local (LNS) IP address for L2TP tunnel.

Options

address—Local IP address; corresponds to the IP address that is used by LACs to identify the LNS. When the LAC is an MX Series router, this address matches the remote gateway address configured in the LAC tunnel profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring L2TP Tunnel Groups
- Configuring L2TP Tunnel Groups
- *L2TP LNS Inline Service Interfaces*

address (Local RPs)

IN THIS SECTION

- Syntax | 1885
- Hierarchy Level | 1885
- Description | 1885
- Options | 1885
- Required Privilege Level | 1886
- Release Information | 1886

Syntax

```
address address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp local family (inet | inet6)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp local family (inet | inet6)],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim rp local family (inet | inet6)]
```

Description

Configure the local rendezvous point (RP) address.

Options

address—Local RP address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Local PIM RPs

address (LRF Profile)

IN THIS SECTION

- Syntax | 1886
- Hierarchy Level | 1886
- Description | 1887
- Options | 1887
- Required Privilege Level | 1887
- Release Information | 1887

Syntax

```
address collector-address;
```

Hierarchy Level

```
[edit services lrf profile profile-name collector collector-name destination]
```

Description

Specify the destination IP address of the collector.

Options

collector-address	IP address of the collector.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Junos OS Subscriber Aware

address-mapping

IN THIS SECTION

- Syntax | 1888
- Hierarchy Level | 1888
- Description | 1888
- Required Privilege Level | 1888
- Release Information | 1888

Syntax

```
address-mapping;
```

Hierarchy Level

```
[edit logical-systems name security nat source rules-set name rule name then source-nat],
[edit logical-systems name tenants name security nat source rule name then source-nat],
[edit security nat source rule-set ruleset rule rule then source-nat interface persistent-nat],
[edit services nat source rule-set name rule name then source-nat],
[edit tenants name security nat source rule-set name rule name then source nat]
```

Description

Allows requests from a specific internal IP address to be mapped to the same reflexive IP address (the public IP address created by the NAT device closest to the STUN server); internal and external ports can be any ports. An external host using any port can send a packet to the internal host by sending the packet to the reflexive IP address (with a configured incoming policy that allows external to internal traffic). If this option is not configured, the persistent NAT binding is for specific internal and reflexive transport addresses.

You can only specify this option when the persistent NAT type is `any-remote-host` and the source NAT rule action is one of the following:

- Source NAT pool with IP address shifting
- Source NAT pool with no port translation and no overflow pool

Required Privilege Level

`security`—To view this statement in the configuration.

`security-control`—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 10.2.

address-mapping (Application Identification)

IN THIS SECTION

- Syntax | 1889
- Hierarchy Level | 1889
- Description | 1889
- Options | 1890
- Required Privilege Level | 1890
- Release Information | 1890

Syntax

```
address-mapping address-name {
    filter {
        ip ip-address-and-prefix-length;
        port-range {
            tcp [port];
            udp [port];
        }
    }
}
```

Hierarchy Level

[edit services application-identification application *application-name*]

Description

Match the specified IP address.

Layer 3 and Layer 4 address mapping defines an application by the IP address and optional port range of the traffic. You can use the address mapping option to configure custom applications signatures when the configuration of your private network predicts application traffic to or from trusted servers.

Address mapping provides efficiency and accuracy in handling traffic from a known application.

Options

name Address mapping name.

filter Specify the application matching criteria by the IP address of the application or the port range to match TCP or UDP destination port.

- ip—IP address and prefix-length.
- port-range—Port range to match a TCP or UDP destination port.
  - **tcp** [*port*]—Define the TCP port range for the application.
  - **udp** [*port*]—Define the UDP port range for the application.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

Understanding Junos OS Application Identification Custom Application Signatures

address-mask

IN THIS SECTION

- Syntax | 1891
- Hierarchy Level | 1891
- Description | 1891
- Options | 1891
- Required Privilege Level | 1891
- Release Information | 1892

Syntax

```
address-mask address-mask;
```

Hierarchy Level

```
[edit snmp v3 target-address target-address-name]
```

Description

Verify the source addresses for a group of target addresses.

Options

address-mask combined with the address defines a range of addresses.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Address Mask

address-name

IN THIS SECTION

- Syntax | 1892
- Hierarchy Level | 1893
- Description | 1893
- Options | 1893
- Required Privilege Level | 1893
- Release Information | 1893

Syntax

```
address-name name {
    description description;
    profile {
        category name {
            feed feed;
            property name {
                string name;
            }
        }
        feed-name name;
    }
}
```

Hierarchy Level

```
[edit security dynamic-address]
[edit logical-systems logical-system-name security dynamic-address]
[edit tenants tenant-name security dynamic-address]
```

Description

Specify the security dynamic address name for IPv4 and IPv6 networks within a logical system and tenant system.

Options

- `profile`—Information to categorize feed data into this dynamic address.
  - `feed-name`— Name of feed in feed-server for the dynamic address.
    - `source-feed-name`— Name of feed in feed-server which is mapped to this dynamic address.
  - `category`—Name of category for the dynamic address.
    - `GeoIP` —A list giving you the ability to filter traffic to and from specific geographies in the world.
      - `feed`—Feed to match.
        - `feed-name`—Name of feed (fd1).
    - `ip-filter` —A list of addresses and ranges of malicious sites that can send junk data.
      - `feed`—Feed to match.
        - `feed-name`—Name of feed (fd1).

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

address (NAT Pool Next Gen Services)

IN THIS SECTION

- Syntax | 1894
- Hierarchy Level | 1894
- Description | 1894
- Options | 1894
- Required Privilege Level | 1895
- Release Information | 1895

Syntax

```
address address-prefix | address address-prefix to address address-prefix;
```

Hierarchy Level

```
[edit services nat destination pool nat-pool-name],
[edit services nat source pool nat-pool-name]
```

Description

Define the addresses or subnets to which source addresses or destination addresses are translated. You can configure a single address, an address range, a single subnet, or a subnet range.

Options

address *address-prefix*	A single address or subnet.
address *address-prefix* to address *address-prefix*	An address range or a subnet range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

address (Peer)

IN THIS SECTION

- Syntax | 1895
- Hierarchy Level | 1895
- Description | 1896
- Default | 1896
- Options | 1896
- Required Privilege Level | 1896
- Release Information | 1896

Syntax

```
address ip-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name],
[edit protocols link-management peer peer-name]
```

Description

Specify the ID of the peer.

Default

The loopback address is advertised.

Options

ip-address—IP address of the peer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the ID for LMP Peers

address-persistent (Security Source NAT)

IN THIS SECTION

- Syntax | 1897
- Hierarchy Level | 1897
- Description | 1897
- Required Privilege Level | 1897
- Release Information | 1897

Syntax

```
address-persistent;
```

Hierarchy Level

```
[edit logical-systems name security nat],
[edit logical-systems name tenants name security nat],
[edit security nat source]
[edit services nat],
[edit tenants name security nat]
```

Description

Enable the device to assign the same, statically chosen, IP address from a source pool to a host for multiple sessions that require the same source IP address for each session. This option is a global configuration and is applied to all source pools. After a session is established from a host and NAT is performed, the subsequent session from the same host will always use the same translated address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

address-persistent (Security Source NAT Pool)

IN THIS SECTION

Syntax | 1898

- Hierarchy Level | 1898
- Description | 1898
- Options | 1898
- Required Privilege Level | 1898
- Release Information | 1899

Syntax

```
address-persistent subscriber ipv6-prefix-length prefix-length;
```

Hierarchy Level

```
[edit logical-systems name security nat source pool name address-persistent subscriber],
    [edit logical-systems name tenants name security nat source pool name address-persistent
subscriber],
    [edit security nat source pool name address-persistent subscriber],
    [edit services nat source pool name address-persistent subscriber],
    [edit tenants name security nat source pool name address-persistent subscriber]
```

Description

Enable the device to translate an IPv6 address, with a consistent IPv6 prefix, to the same IPv4 address to ensure that IPv4 services can be used over IPv6-only networks.

Options

ipv6-prefix-length *prefix-length*

Specify the subscriber IPv6 prefix length.

- **Range:** 8 through 128.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D10.

RELATED DOCUMENTATION

Understanding NAT64 IPv6 Prefix to IPv4 Address-Persistent Translation

address-pool

IN THIS SECTION

- Syntax | 1899
- Hierarchy Level | 1899
- Description | 1900
- Options | 1900
- Required Privilege Level | 1900
- Release Information | 1900

Syntax

```
address-pool pool-name {
    address address-or-prefix;
    address-range <low lower-limit> <high upper-limit>;
}
```

Hierarchy Level

```
[edit access]
```

Description

Allocate IP addresses for clients.

> **NOTE**: This statement is not supported for L2TP LNS on MX Series routers.

Options

pool-name Name assigned to an address pool.

address (EX Series, M Series, PTX Series, T Series only) Configure the IP address or prefix value for clients.

- **Values:** *address-or-prefix*—An address or prefix value.

address-range Configure the address range.

- Values:
  - high *upper-limit*—Upper limit of an address range.
  - low *lower-limit*—Lower limit of an address range.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Address Pool for L2TP Network Server IP Address Allocation

address-pool (Access)

IN THIS SECTION

- Syntax | 1901
- Hierarchy Level | 1901
- Description | 1901
- Options | 1902
- Required Privilege Level | 1902
- Release Information | 1902

Syntax

```
address-pool pool-name {
    address address or address prefix;
    address-range {
        high upper-limit;
        low lower-limit;
        mask network-mask;
    }
    primary-dns IP address;
    primary-wins IP address;
    secondary-dns IP address;
    secondary-wins IP address;
}
```

Hierarchy Level

```
[edit access]
```

Description

Create an address-pool for L2TP clients.

Options

- `pool-name`—Name assigned to the address-pool.
- `address`—Configure subnet information for the address-pool.
- `address-range`—Defines the address range available for clients.
- `primary-dns`—Specify the primary-dns IP address.
- `secondary-dns`—Specify the secondary-dns IP address.
- `primary-wins`—Specify the primary-wins IP address.
- `secondary-wins`—Specify the secondary-wins IP address.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

access-control

address-pool (Domain Map)

IN THIS SECTION

- Syntax | 1903
- Hierarchy Level | 1903
- Description | 1903
- Options | 1903

Required Privilege Level | 1903
Release Information | 1903

Syntax

```
address-pool pool-name;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Specify the address pool used to assign addresses to subscribers associated with the domain map.

Options

pool-name—Name of address pool.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Specifying an Address Pool in a Domain Map

address-pooling

IN THIS SECTION

- Syntax | 1904
- Hierarchy Level | 1904
- Description | 1904
- Options | 1904
- Required Privilege Level | 1904
- Release Information | 1905

Syntax

```
address-pooling paired;
```

Hierarchy Level

```
[edit services nat rule (Services NAT) rule-name term (Services NAT) term-name then (Services NAT)
translated]
```

Description

Specify the NAT address pooling behavior.

Options

paired—Currently, the only valid setting specifies paired address pooling behavior.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.1.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

address-pooling (Security Source NAT)

IN THIS SECTION

- Syntax | 1905
- Hierarchy Level | 1905
- Description | 1906
- Options | 1906
- Required Privilege Level | 1906
- Release Information | 1906

Syntax

```
address-pooling (paired | no-paired);
```

Hierarchy Level

```
[edit logical-systems name security nat source pool name address-pooling],
 [edit logical-systems name tenants name security nat source pool name address-pooling],
 [edit security nat source pool name address-pooling],
 [edit services nat source pool name address-pooling],
 [edit tenants name security nat source pool name address-pooling]
```

Description

The `address-pooling paired` and `address-pooling no-paired` options in a source NAT pool enable you to override the global address-persistent configuration and to control the IP addressing in the pool. When either `address pooling-paired` or `address-pooling no-paired` is configured in a NAT source pool, the address-persistent configuration is disabled for that pool.

Use the `address-pooling paired` option in source NAT pools with port translation for applications that require all sessions associated with one internal IP address to be translated to the same external IP address for multiple sessions. (The default behavior for a source NAT pool with port translation pools is address-pooling no-paired.)

Use the `address-pooling no-paired` option in source NAT pools without port translation for assigning IP addresses using a round-robin fashion. (The default behavior for a source NAT pool without port translation is address-pooling paired.)

Options

no-paired	Allow address-pooling no-paired
paired	Allow address-pooling paired

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding Source NAT Pools with Address Pooling

address-pooling (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 1907
- Hierarchy Level | 1907
- Description | 1907
- Options | 1907
- Required Privilege Level | 1907
- Release Information | 1908

Syntax

```
address-pooling {
    no-paired;
}
```

Hierarchy Level

```
[edit services nat source pool pool-name]
```

Description

Allow address-pooling no-paired for a source pool without port translation

Options

no-paired Allow address-pooling no-paired for a source pool without port translation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

address-pools

IN THIS SECTION

- Syntax | 1908
- Hierarchy Level | 1909
- Description | 1909
- Options | 1909
- Required Privilege Level | 1909
- Release Information | 1909

Syntax

```
address-pools {
    name {
        default-pool;
        family (inet | inet6) {
            network {
                [network-prefix] {
                    external-assigned;
                }
            }
        }
        service-mode service-mode-options;
    }
}
```

Hierarchy Level

```
[edit access address-assignment]
```

Description

Configure the address pools that the TDF domains use to specify the source IP addresses of packets to undergo TDF processing.

Options

name Name of the address pool.

- **Range:** 1 through 63 alphanumeric characters

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Understanding Source IP Filtering with Address Pools in TDF Domains for IP-Based Subscribers

Configuring Address Pools for Source-IP Filtering of IP-Based Subscribers

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

address-protection

IN THIS SECTION

- Syntax | 1910
- Hierarchy Level | 1910
- Description | 1910
- Options | 1911
- Required Privilege Level | 1911
- Release Information | 1912

Syntax

```
address-protection {
    reassign-on-match;
}
```

Hierarchy Level

```
[edit access],
[edit logical-systems logical-system-name routing-instances routing-instance-name access]
```

Description

Prevent IPv4 addresses and IPv6 prefixes from being assigned to more than one subscriber session when you use AAA to supply IPv4 addresses.

For IPv4:

If enabled, the router checks the following attributes received from external servers:

- *Framed-IP-Address*
- *Framed-Pool*

The router then takes one of the following actions:

- If an address matches an address in an address pool, the address is taken from the pool, provided it is available.
- If the address is already in use, it is rejected as unavailable.

For IPv6:

If enabled, the router checks the following attributes received from external servers:

- *Framed-IPv6-Prefix*
- *Framed-IPv6-Pool*

The router then takes one of the following actions:

- If a prefix matches a prefix in an address pool, the prefix is taken from the pool, provided it is available.
- If the prefix is already in use, it is rejected as unavailable.
- If the prefix length requested from the external server does not exactly match the pool's prefix length, the authentication request is denied. If configured, the Acct-Stop message includes the cause for termination.

Options

reassign-on-match

Enable reassignment of an address from an existing subscriber to a new subscriber requesting that address. The address in use must not be part of a locally configured pool and address protection must be enabled. The request from the new subscriber is still rejected, but the existing subscriber is sent a disconnect request to begin the logout process. This enables the new subscriber to renegotiate and be assigned that IP address.

If the requested address is in a locally configured pool, the existing subscriber is not disconnected.

> **NOTE**: This option is not supported for IPv6.

- **Default:** Rejects the address request from the new subscriber; the existing subscriber remains intact with the IP address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

`reassign-on-match` option added in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Configuring Duplicate IPv4 Address Protection for AAA

Configuring Duplicate IPv6 Prefix Protection for Router Advertisement

address (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 1912
- Hierarchy Level | 1912
- Description | 1913
- Options | 1913
- Required Privilege Level | 1913
- Release Information | 1913

Syntax

```
address address;
```

Hierarchy Level

```
[edit protocols openflow switch switch-name controller]
```

Description

Specify the IPv4 address of the OpenFlow controller that will manage OpenFlow on that virtual switch. The switch establishes a connection to the controller using this address.

Options

address	IPv4 address of the OpenFlow controller.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Support for OpenFlow on Devices Running Junos OS

controller (Protocols OpenFlow) | **4654**

protocol (Protocols OpenFlow) | **13455**

switch (Protocols OpenFlow) | **28755**

address (Protocols Router Discovery)

IN THIS SECTION

- Syntax | **1914**
- Hierarchy Level | **1914**
- Description | **1914**
- Options | **1914**

Required Privilege Level | 1914
Release Information | 1914

Syntax

```
address (Protocols Router Discovery) address {
    (advertise | ignore);
    (broadcast | multicast);
    (ineligible | priority number);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery],
[edit protocols router-discovery]
```

Description

Specify the IP addresses to include in router advertisement packets.

Options

address—IP address. To specify more than one address, specify multiple addresses or include multiple `address` statements.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

address (RADIUS Clients)

IN THIS SECTION

- Syntax | 1915
- Hierarchy Level | 1915
- Description | 1915
- Options | 1915
- Required Privilege Level | 1916
- Release Information | 1916

Syntax

```
address client-address;
```

Hierarchy Level

```
[edit access radius clients client-name]
```

Description

Specify the address from which the GGSN, PGW, or BNG RADIUS client sends the RADIUS requests.

Options

client-address	IP address of the PGW client.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers

address (RADIUS Server)

IN THIS SECTION

- Syntax | 1916
- Hierarchy Level | 1916
- Description | 1917
- Options | 1917
- Required Privilege Level | 1917
- Release Information | 1917

Syntax

```
address server-address;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure the address of the RADIUS server.

Options

server-address	IP address for the RADIUS server.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

address-range

IN THIS SECTION

- Syntax | 1918
- Hierarchy Level | 1918
- Description | 1918
- Options | 1918
- Required Privilege Level | 1918
- Release Information | 1918

Syntax

```
address-range low minimum-value high maximum-value;
```

Hierarchy Level

```
[edit services nat pool nat-pool-name]
```

Description

Specify the NAT pool address range.

Options

minimum-value—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

minimum-value and ***maximum-value*** options enhanced to support IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Configuring Source and Destination Addresses Network Address Translation Overview*

address-ranges (Demux)

IN THIS SECTION

- Syntax | 1919
- Hierarchy Level | 1920
- Description | 1920
- Required Privilege Level | 1920
- Release Information | 1920

Syntax

```
address-ranges {
    authentication {
        password password-string;
        username-include {
            auth-server-realm realm-string;
            delimiter delimiter-character;
            domain-name domain-name;
            interface-name;
            source-address;
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        network ip-address {
            range name {
                low lower-limit;
                high upper-limit;
            }
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet auto-configure]
[edit interfaces interface-name unit unit-number demux inet6 auto-configure]
```

Description

Specify the address range for the demultiplexing (demux) interface options. The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles

address (Security Address Book)

IN THIS SECTION

- Syntax | 1921
- Hierarchy Level | 1921
- Description | 1921
- Options | 1921
- Required Privilege Level | 1922

Release Information | 1922

Syntax

```
address address-name {
    ip-prefix {
        description text;
    }
    description text;
    dns-name domain-name {
        ipv4-only;
        ipv6-only;
    }
    range-address lower-limit to upper-limit;
    wildcard-address ipv4-address/wildcard-mask;
}
```

Hierarchy Level

```
[edit security address-book book-name]
```

Description

Add an entry containing an IP address or DNS hostname, or wildcard address to the address book. An address book contains entries for addressable entities in security zones, policies, and NAT rules. Address book entries can include any combination of IPv4 addresses, IPv6 addresses, and DNS names.

Options

- `address` *`address-name`*–Name of an address entry.
- `description` *`text`*–Descriptive text about an address entry.
- dns-address *`domain-name`*–DNS address name.
- *`ip-prefix`*–IP address with prefix.

- `range-address` *`lower-limit`* `to` *`upper-limit`*–Address range for an address book.
- `wildcard-address` *`ipv4-address/wildcard-mask`*–IPv4 wildcard address in the form of A.B.C.D/wildcard-mask.

> **NOTE**: IPv6 wildcard address configuration is not supported in this release.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1. Support for address range added in Junos OS Release 12.1. The `description` option added in Junos OS Release 12.1.

address (Security ARP Proxy)

IN THIS SECTION

- Syntax | **1923**
- Hierarchy Level | **1923**
- Description | **1923**
- Options | **1923**
- Required Privilege Level | **1923**
- Release Information | **1923**

Syntax

```
address ip-address to <ip-address>;
```

Hierarchy Level

```
[edit logical-systems name security nat proxy-arp interface],
    [edit logical-systems name tenants name security nat proxy-arp interface],
    [edit security nat proxy-arp interface],
    [edit services nat proxy-arp interface],
    [edit tenants name security nat proxy-arp interface]
```

Description

Specify a single address or an address range of ARP proxy.

Options

to	Specify the upper limit of the address range.
ip-address	IP address of an ARP proxy.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

address (Security Destination NAT)

IN THIS SECTION

- Syntax | 1924
- Hierarchy Level | 1924
- Description | 1924
- Options | 1924
- Required Privilege Level | 1925
- Release Information | 1925

Syntax

```
address <ip-address> {
    (port port-number | to ip-address);
}
```

Hierarchy Level

```
[edit security nat destination pool pool-name],
[edit logical-systems name security nat destination pool],
[edit services nat destination pool],
[edit tenants name security nat destination pool]
```

Description

Specify a single address or an address range of the destination NAT pool.

Options

- *ip-address* —IP address of a pool.
- `port` *port-number*—Specify the port number.

- `to`—Specify the upper limit of the address range.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

address (Security NDP Proxy)

IN THIS SECTION

- Syntax | 1925
- Hierarchy Level | 1925
- Description | 1926
- Options | 1926
- Required Privilege Level | 1926
- Release Information | 1926

Syntax

```
address ip-address {
    to ip-address;
}
```

Hierarchy Level

```
[edit logical-systems name security nat proxy-ndp interface],
    [edit logical-systems name tenants name security nat proxy-ndp interface],
```

```
[edit security nat proxy-ndp interface],
[edit services nat proxy-ndp interface],
[edit tenants name security nat proxy-ndp interface]
```

Description

Specify a single address or an address range of NDP proxy. IPv6 Neighbor Discovery Protocol (NDP) is to resolve network layer (IP) addresses to link layer addresses, such as Ethernet. Address Resolution Protocol (ARP) performs this function in IPv4.

Options

- *ip-address*—IP address of an NDP proxy.
- `to`—Specify the upper limit of the address range.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

address (Services)

IN THIS SECTION

- Syntax | 1927
- Hierarchy Level | 1927
- Description | 1927
- Required Privilege Level | 1927
- Release Information | 1927

Syntax

```
address ip-address;
```

Hierarchy Level

```
[edit services unified-access-control infranet-controller hostname]
```

Description

Specify the IP address of the IC Series device with which the SRX Series Firewalls should communicate.

This statement is required when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Understanding Unified Access Control

Acquiring User Role Information from an Active Directory Authentication Server

address (Services Dynamic Flow Capture)

IN THIS SECTION

- Syntax | 1928
- Hierarchy Level | 1928
- Description | 1928
- Options | 1928
- Required Privilege Level | 1928
- Release Information | 1929

Syntax

```
address address;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Configure an IP address for the flow capture destination.

Options

address—IP address for the content destination.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Content Destination

address (Services NAT Pool)

IN THIS SECTION

- Syntax | 1929
- Hierarchy Level | 1929
- Description | 1929
- Options | 1930
- Required Privilege Level | 1930
- Release Information | 1930

Syntax

```
address ip-prefix</prefix-length>;
```

Hierarchy Level

```
[edit services nat pool nat-pool-name]
```

Description

Specify the NAT pool prefix value.

The subnet and broadcast addresses are not included in the list of usable IP addresses. For example, if you use 10.11.12.0/28 for the NAT pool prefix value, the addresses 10.11.12.0 (subnet address) and 10.11.12.15 (broadcast address) are not available.

Options

ip-prefix—Specify an IPv4 or IPv6 prefix value.

/prefix-length—(Optional) Specify an IPv4 or IPv6 prefix length.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The *ip-prefix* and */prefix-length* options were enhanced to support IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Source and Destination Addresses Network Address Translation Overview

address-set

IN THIS SECTION

- Syntax | 1931
- Hierarchy Level | 1931
- Description | 1931
- Options | 1931
- Required Privilege Level | 1931

Release Information | 1932

Syntax

```
address-set address-set-name {
    address address-name;
    address-set address-set-name;
    description text;
}
```

Hierarchy Level

```
[edit security address-book book-name]
```

Description

Specify a collection of addresses, as defined in the `address (Address Book)` statement. Using address sets, you can organize addresses in logical groups and use them to easily configure other features, such as policies and NAT rules. Using this statement, you can also include a description for an address set.

You can also define address sets within address sets.

Options

address-set-name–Name of the address set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for nested address sets introduced in Junos OS Release 11.2. The `description` option added in Junos OS Release 12.1.

address-shared (Security Source NAT)

IN THIS SECTION

- Syntax | 1932
- Hierarchy Level | 1932
- Description | 1932
- Required Privilege Level | 1933
- Release Information | 1933

Syntax

```
address-shared;
```

Hierarchy Level

```
[edit logical-systems name security nat source pool],
    [edit logical-systems name tenants name security nat source pool],
    [edit security nat source pool],
    [edit services nat source pool],
    [edit tenants name security nat source pool]
```

Description

Specifies that multiple internal IP addresses can be mapped to the same external IP address. Use this option only when the source NAT pool is configured with no port translation.

When a source NAT pool configured with no port translation has few external IP addresses available, or only one external IP address, the `address-shared` option, with a many-to-one mapping, increases NAT resources and improves traffic.

Required Privilege Level

security—To view this statement in the configuration

security-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding Shared Addresses in Source NAT Pools without PAT

Example: Configuring a Single IP Address in a Source NAT Pool Without PAT

address (Static RPs)

IN THIS SECTION

- Syntax | 1933
- Hierarchy Level | 1934
- Description | 1934
- Options | 1934
- Required Privilege Level | 1934
- Release Information | 1934

Syntax

```
address address {
    group-ranges {
```

```
        destination-ip-prefix</prefix-length>;
    }
    override;
    version version;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp static],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp static],
[edit protocols pim static],
[edit routing-instances routing-instance-name protocols pim rp static]
```

Description

Configure static rendezvous point (RP) addresses. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.

For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.

Options

address—Static RP address.

- **Default:** 224.0.0.0/4

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Static PIM RP Address on the Non-RP Routing Device

address (Tunnel Profile Remote Gateway)

IN THIS SECTION

- Syntax | 1935
- Hierarchy Level | 1935
- Description | 1935
- Options | 1935
- Required Privilege Level | 1936
- Release Information | 1936

Syntax

```
address server-ip-address;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id remote-gateway]
```

Description

Specify the IP address of the remote gateway device at the L2TP tunnel endpoint, the LNS.

Options

server-ip-address—IP address of the remote gateway device.

- **Default:** 0.0.0.0.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

address (Tunnel Profile Source Gateway)

IN THIS SECTION

- Syntax | 1936
- Hierarchy Level | 1936
- Description | 1937
- Options | 1937
- Required Privilege Level | 1937
- Release Information | 1937

Syntax

```
address client-ip-address;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id source-gateway]
```

Description

Specify the IP address of the source gateway device at the local L2TP tunnel endpoint, the LAC. This value overrides the default address for the logical system or routing instance.

Options

client-ip-address—IP address of the source gateway device.

- **Default:** 0.0.0.0.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

address-whitelist

IN THIS SECTION

- Syntax | 1938
- Hierarchy Level | 1938
- Description | 1938
- Required Privilege Level | 1938
- Release Information | 1938

Syntax

```
address-whitelist list-name;
```

Hierarchy Level

```
[edit security utm feature-profile anti-spam]
[edit security utm default-configuration]
```

Description

Enter an address-allowlist (or blocklist) custom-object for local list spam filtering.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

addresses (REST API)

IN THIS SECTION

- Syntax | 1939
- Hierarchy Level | 1939
- Description | 1939
- Required Privilege Level | 1939
- Release Information | 1939

Syntax

```
addresses [ address-list ];
```

Hierarchy Level

```
[edit system services rest http] ,
[edit system services rest https]
```

Description

Specify IP addresses for incoming connections.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

adf (Dynamic Firewalls)

IN THIS SECTION

- Syntax | 1940
- Hierarchy Level | 1940
- Description | 1940
- Options | 1940
- Required Privilege Level | 1941
- Release Information | 1941

Syntax

```
adf {
    counter;
    input-precedence precedence;
    not-mandatory;
    output-precedence precedence;
    rule rule-value;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family filter]
```

Description

Configure an Ascend-Data-Filter that the dynamic profile applies to a subscriber session.

Options

`counter`—Enable a counter that increments each time the Ascend-Data-Filter rule is used. Typically used for testing purposes.

`not-mandatory`—Suppress router from reporting an error when the RADIUS reply message does not include the $junos-adf-rule-v4 or $junos-adf-rule-v6 variable that is configured for the Ascend-Data-Filter in the dynamic profile. In this circumstance, the Ascend-Data-Filter is not created.

`precedence`—Precedence value that sets the order in which dynamic service filters are applied on the interface. The lower the precedence value, the higher the precedence that is given. The precedence setting is used in conjunction with the precedence settings of all dynamic service filters configured (not only Ascend-Data-Filters) on the same interface to establish the order. For example, the order also includes any configured `input` *`filter-name`* `precedence` *`precedence`* and `output` *`filter-name`* `precedence` *`precedence`* statements.

- **Range:** 0 through 255
- **Default:** 0

`rule-value`—Ascend-Data-Filter rule. You can specify either a Junos predefined variable that maps the Ascend-Data-Filter actions to Junos filter functionality or you can manually configure the Ascend-Data-

Filter rule. The router supports two predefined variables depending on family type: `$junos-adf-rule-v4` for family `inet` and `$junos-adf-rule-v6` for family `inet6`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Option `not-mandatory` introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- *Understanding Dynamic Firewall Filters*
- *Classic Filters Overview*
- *Basic Classic Filter Syntax*
- Guidelines for Configuring Service Filters

adjacency-loss

IN THIS SECTION

- Syntax | **1942**
- Hierarchy Level | **1942**
- Description | **1942**
- Required Privilege Level | **1942**
- Release Information | **1942**

Syntax

```
adjacency-loss;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
event ais-trigger-condition]
```

Description

Configure the defect condition that raises an alarm indication signal when physical connectivity is lost between peer MEPs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Alarm Indication Signal (ETH-AIS) Function Overview

Configuring ETH-AIS on a CFM MEP

adjacency-loss-hold-time (ANCP)

IN THIS SECTION

- Syntax | 1943
- Hierarchy Level | 1943
- Description | 1943
- Options | 1943
- Required Privilege Level | 1944
- Release Information | 1944

Syntax

```
adjacency-loss-hold-time seconds;
```

Hierarchy Level

```
[edit protocols ancp],
[edit protocols ancp neighbor ip-address]
```

Description

Configure the ANCP agent to monitor, either globally or for a specified neighbor, how long an ANCP adjacency is down and to trigger a state change for the subscriber access line if the hold timer expires before the adjacency comes back up, as indicated by a Port Up message on the access line. By default, there is no delay between detecting an adjacency loss and triggering the state change.

Options

seconds Duration of period that the ANCP agent monitors loss of adjacency.

- **Default:** 0 seconds
- **Range:** 0 through 1800 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses

Configuring ANCP Neighbors

Configuring the ANCP Agent

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

adjacency-timer

IN THIS SECTION

- Syntax | 1944
- Hierarchy Level | 1945
- Description | 1945
- Options | 1945
- Required Privilege Level | 1945
- Release Information | 1945

Syntax

```
adjacency-timer seconds;
```

Hierarchy Level

```
[edit protocols ancp],
[edit protocols ancp neighbor ip-address]
```

Description

Specify a value for the interval that the ANCP agent proposes during negotiation to establish an adjacency, for all neighbors or a specific neighbor. The larger of the values proposed by the agent and the neighbor is selected for the interval between subsequent adjacency messages exchanged by the agent and the neighbor.

Options

seconds—Number of seconds between adjacency messages.

- **Range:** 1 through 25 seconds
- **Default:** 10 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

> *Configuring the ANCP Agent*
>
> *Specifying the Interval Between ANCP Adjacency Messages*
>
> *Configuring ANCP Neighbors*

adjust-interval

IN THIS SECTION

- Syntax | 1946
- Hierarchy Level | 1946
- Description | 1946
- Options | 1946
- Required Privilege Level | 1947
- Release Information | 1947

Syntax

```
adjust-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Specify the bandwidth reallocation interval.

Options

seconds—Bandwidth reallocation interval, in seconds.

- **Range:** 300 through 315,360,000 seconds
- **Default:** 86,400 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Automatic Bandwidth Allocation Interval

adjust-minimum

IN THIS SECTION

- Syntax | 1947
- Hierarchy Level | 1948
- Description | 1948
- Options | 1948
- Required Privilege Level | 1948
- Release Information | 1948

Syntax

```
adjust-minimum rate;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name],
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
```

Description

For adjustments performed by the ANCP or multicast applications on EQ DPC, MIC, or MPC interfaces, specify the minimum shaping rate for an adjusted scheduler node. The node is associated with a traffic-control profile.

For adjustments performed by the multicast application on MIC or MPC interfaces, specify the minimum shaping rate for an adjusted queue. The queue is associated with a scheduler.

Options

rate—Minimum shaping rate for a node or a queue, in Mbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers

adjust-minimum (Dynamic Shaping and Scheduling)

IN THIS SECTION

- Syntax | 1949
- Hierarchy Level | 1949
- Description | 1949
- Options | 1949
- Required Privilege Level | 1950
- Release Information | 1950

Syntax

```
adjust-minimum (rate | $junos-cos-adjust-minimum);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name],
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles traffic-control-
profile-name]
```

Description

For adjustments performed by the ANCP or multicast applications on EQ DPCs and MPC/MIC interfaces, specify the minimum shaping rate for an adjusted scheduler node. The node is associated with a traffic-control profile.

For adjustments performed by the multicast application on MPC/MIC interfaces, specify the minimum shaping rate for an adjusted queue. The queue is associated with a scheduler.

Options

rate—Minimum shaping rate for a node or a queue, in Mbps

`$junos-cos-adjust-minimum`—Junos OS predefined variable that is replaced with the minimum shaping rate for a node that is obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached. Use this variable at the `[edit dynamic-profiles` *`profile-name`* `class-of-service traffic-control-profiles]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Configuring a Dynamic Minimum Adjusted Shaping Rate on Scheduler Nodes
- Configuring a Dynamic Shaping-Rate Adjustment for Queues

adjust-percent

IN THIS SECTION

- Syntax | **1951**
- Hierarchy Level | **1951**
- Description | **1951**
- Options | **1951**
- Required Privilege Level | **1951**
- Release Information | **1951**

Syntax

```
adjust-percent percentage;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

For a MIC or MPC interface, determine the percentage of adjustment for the shaping rate of a queue.

Options

percentage Percentage of the shaping rate to adjust.

- **Range:** 0 through 100 percent

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Shaping-Rate Adjustments on Queues

adjust-percent (Dynamic Schedulers)

IN THIS SECTION

- Syntax | 1952
- Hierarchy Level | 1952
- Description | 1952
- Options | 1952
- Required Privilege Level | 1952
- Release Information | 1953

Syntax

```
adjust-percent percentage;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name  class-of-service schedulers scheduler-name]
```

Description

For a MPC/MIC interface, determine the percentage of adjustment for the shaping rate of a queue.

Options

percentage—Percentage of the shaping rate to adjust.

- **Range:** 0 through 100 percent

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring a Dynamic Shaping-Rate Adjustment for Queues

adjust-threshold

IN THIS SECTION

- Syntax | 1953
- Hierarchy Level | 1953
- Description | 1954
- Options | 1954
- Required Privilege Level | 1954
- Release Information | 1954

Syntax

```
adjust-threshold percent;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Specify in percentage how sensitive the automatic bandwidth adjustment for a label-switched path (LSP) is to changes in bandwidth utilization.

To specify the changes in the automatic bandwidth adjustment for a LSP in absolute value, use the "adjust-threshold-absolute" on page 1954 statement instead.

Options

percent—Bandwidth demand for the current bandwidth adjustment interval is determined and compared to the LSP's current bandwidth allocation. If the percentage difference in bandwidth is greater than or equal to the percentage specified by this statement, the LSP's bandwidth is adjusted to the current bandwidth demand.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Automatic Bandwidth Adjustment Threshold

adjust-threshold-absolute

IN THIS SECTION

- Syntax | 1955
- Hierarchy Level | 1955
- Description | 1955

- Options | 1955
- Required Privilege Level | 1956
- Release Information | 1956

Syntax

```
adjust-threshold-absolute bps;
```

Hierarchy Level

```
[edit logical-systems name protocols mpls label-switched-path name auto-bandwidth],
[edit logical-systems name routing-instances name protocols mpls label-switched-path name auto-
bandwidth],
[edit protocols mpls label-switched-path name auto-bandwidth],
[edit routing-instances name protocols mpls label-switched-path name auto-bandwidth]
```

Description

Specify in bits per second how sensitive the automatic bandwidth adjustment for a label-switched path (LSP) is to changes in the average LSP utilization.

The `adjust-threshold-absolute` statement works in conjunction with the `adjust-threshold` statement, which specifies the change in automatic bandwidth adjustment for an LSP as a percentage.

By triggering automatic bandwidth LSP resignaling based on absolute change in bandwidth instead of percentage bandwidth change, LSP resignaling can be optimized for both big and small LSPs at the same time.

Options

bps Change in average LSP utilization to trigger automatic bandwidth adjustment in bits per second.

- **Default:** 0 bps

Required Privilege Level

routing

Release Information

Statement introduced before Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Configuring the Automatic Bandwidth Adjustment Threshold

adjust-threshold-activate-bandwidth

IN THIS SECTION

- Syntax | 1956
- Hierarchy Level | 1957
- Description | 1957
- Options | 1957
- Required Privilege Level | 1957
- Release Information | 1957

Syntax

```
adjust-threshold-activate-bandwidth bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Specify an absolute value to prevent automatic adjustment of signaled bandwidth and aggressive re-signaling of a label-switched path (LSP) when the actual bandwidth over the LSP is below the configured threshold, although the adjust-threshold percentage condition is satisfied.

Options

bps—Amount of bandwidth that is compared with the maximum of all traffic samples during an adjustment interval. If the maximum average bandwidth is less than this configured value, automatic bandwidth adjustment or re-signaling does not happen, even if the adjust-threshold percentage condition is satisfied.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring the Automatic Bandwidth Adjustment Threshold

adjust-threshold-overflow-limit

IN THIS SECTION

- Syntax | 1958
- Hierarchy Level | 1958
- Description | 1958
- Options | 1958
- Required Privilege Level | 1959
- Release Information | 1959

Syntax

```
adjust-threshold-overflow-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Specify the number of consecutive bandwidth overflow samples before triggering a bandwidth adjustment.

Options

number—Number of consecutive bandwidth overflow samples.

- **Range:** 1 through 65,535
- **Default:** This feature is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configuring a Limit on Bandwidth Overflow and Underflow Samples

adjust-threshold-underflow-limit

IN THIS SECTION

- Syntax | 1959
- Hierarchy Level | 1960
- Description | 1960
- Options | 1960
- Required Privilege Level | 1960
- Release Information | 1960

Syntax

```
adjust-threshold-underflow-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Specify the number of consecutive bandwidth underflow samples before triggering a bandwidth adjustment.

Options

number—Number of consecutive bandwidth underflow samples.

- **Range:** 1 through 65,535
- **Default:** This feature is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring a Limit on Bandwidth Overflow and Underflow Samples

adjustment-control-profiles

IN THIS SECTION

- Syntax | 1961
- Hierarchy Level | 1961
- Description | 1961
- Options | 1962
- Required Privilege Level | 1962
- Release Information | 1962

Syntax

```
adjustment-control-profiles {
    profile-name {
        application {
            (ancp | dhcp-tags | pppoe-tags | radius-coa)
                priority priority;
                algorithm algorithm;
            }
        }
    }
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Configure the CoS adjustment control profile.

Options

profile-name Name of the adjustment control profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

CoS Adjustment Control Profiles Overview

Configuring CoS Adjustment Control Profiles

Verifying the CoS Adjustment Control Profile Configuration

admin-down

IN THIS SECTION

- Syntax | 1963
- Hierarchy Level | 1963
- Description | 1963
- Required Privilege Level | 1963
- Release Information | 1963

Syntax

```
admin-down;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Set a nonpacket GMPLS LSP to the administrative down state. This statement does not affect control path setup or data forwarding for packet LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Allowing Nonpacket GMPLS LSPs to Establish Paths Through Routers Running Junos OS

admin-email

IN THIS SECTION

Syntax | 1964

- Hierarchy Level | 1964
- Description | 1964
- Required Privilege Level | 1964
- Release Information | 1964

Syntax

```
admin-email email-address;
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus sophos-engine pattern-update email-notify]
[edit security utm default-configuration anti-virus avira-engine pattern-update email-notify]
```

Description

You can configure the device to notify a specified administrator when patterns are updated. This is an e-mail notification with a custom message and a custom subject line.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for Avira engine added in Junos OS Release 18.4R1.

admin-group (Backup Selection)

IN THIS SECTION

- Syntax | 1965
- Hierarchy Level | 1965
- Description | 1965
- Options | 1966
- Required Privilege Level | 1966
- Release Information | 1966

Syntax

```
admin-group {
    exclude [ group-name ];
    include-all [ group-name ];
    include-any [ group-name ];
    preference [ group-name ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options backup-selection destination prefix
interface interface name],
[edit logical-systems logical-system-name routing-instances instance-name routing-options backup-
selection destination prefix interface interface-name],
[edit routing-instances instance-name routing-options backup-selection destination prefix
interface interface-name],
[edit routing-options backup-selection destination prefix interface interface name]
```

Description

Define the administrative groups criteria for the selection of the backup path.

NOTE: Configure group names of admin-group under the [`edit protocols mpls`] hierarchy level.

Options

exclude [*group-name*]

Specify the administrative groups to be excluded. The backup path is not selected as the loop-free alternate (LFA) or backup next hop if any of the links in the path have any one of the listed administrative groups.

group-name Name of one or more admin-group defined under the [`edit protocols mpls`] hierarchy level.

include-all [*group-name*]

Require each link in the backup path to have all the listed administrative groups in order to accept the path.

group-name Name of one or more admin-group defined under the [`edit protocols mpls`] hierarchy level.

include-any [*group-name*]

Require each link in the backup path to have at least one of the listed administrative groups in order to select the path.

group-name Name of one or more admin-group defined under the [`edit protocols mpls`] hierarchy level.

preference [*group-name*]

Define an ordered set of administrative groups that specifies the preference of the backup path. The leftmost element in the set is given the highest preference.

group-name Name of one or more admin-group defined under the [`edit protocols mpls`] hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

- Understanding Backup Selection Policy for IS-IS Protocol
- Configuring Backup Selection Policy for IS-IS Protocol
- *Understanding Backup Selection Policy for OSPF Protocol*
- backup-selection (Protocols IS-IS)

admin-group

IN THIS SECTION

- Syntax | 1967
- Hierarchy Level | 1967
- Description | 1968
- Options | 1968
- Required Privilege Level | 1968
- Release Information | 1968

Syntax

```
admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection bypass bypass-name],
```

```
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Description

Enable you to configure administrative groups for bypass label-switched paths (LSPs). You can configure administrative groups either globally for all bypass LSPs traversing an interface or for just a specific bypass LSP.

Options

`exclude` *`group-names`*—Specify the administrative groups to exclude for a bypass LSP.

`include-all` *`group-names`*—Specify the administrative groups whose links the bypass LSP must traverse.

`include-any` *`group-names`*—Specify the administrative groups whose links the bypass LSP can traverse.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Administrative Groups for Bypass LSPs

admin-group-extended

IN THIS SECTION

Syntax | 1969

- Hierarchy Level | 1969
- Description | 1969
- Options | 1970
- Required Privilege Level | 1970
- Release Information | 1970

Syntax

```
admin-group-extended {
    apply-groups group-value;
    apply-groups-except group-value;
    exclude [ group-values ];
    include-all [ group-values ];
    include-any [ group-values ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name  (primary
| secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

Specifies the group name and group identifier for an administrative group. The group identifier must be within the range of values specified by the `admin-groups-extended-range` statement. The extended administrative group values are global and must be identically configured on all the supported routers participating in the network. The domain-wide extended administrative groups database, learned from other routers through IGP flooding, is used by CSPF for path computation.

Options

apply-groups Apply the specified administrative groups for the LSP or for the primary and secondary paths.

apply-groups-except Exclude the specified administrative groups from the LSP or from the primary and secondary paths.

exclude Define the administrative groups to exclude from an LSP or from the primary and secondary paths.

include-all Require the LSP to traverse links that include all of the defined administrative groups.

include-any Define the administrative groups to include for an LSP for the primary and secondary paths.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Extended Administrative Groups for LSPs

Configuring Administrative Groups for LSPs

admin-groups-extended | **1975**

admin-groups-extended-range | **1977**

admin-group (for Interfaces)

IN THIS SECTION

- Syntax | 1971
- Hierarchy Level | 1971
- Description | 1971
- Options | 1971
- Required Privilege Level | 1971
- Release Information | 1972

Syntax

```
admin-group [ group-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls interface interface-name],
[edit protocols mpls interface interface-name]
```

Description

Define administrative groups for an interface.

Options

group-names—One or more names of groups defined with the "admin-groups" on page 1973 statement at the `[edit protocols mpls]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Administrative Groups for LSPs

admin-groups | 1973

admin-group (for LSPs)

IN THIS SECTION

- Syntax | 1972
- Hierarchy Level | 1972
- Description | 1973
- Options | 1973
- Required Privilege Level | 1973
- Release Information | 1973

Syntax

```
admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
```

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

Define the administrative groups to include or exclude an LSP and a path's primary and secondary paths.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Administrative Groups for LSPs

admin-groups

IN THIS SECTION

- Syntax | 1974
- Hierarchy Level | 1974
- Description | 1974

- Options | 1974
- Required Privilege Level | 1974
- Release Information | 1975

Syntax

```
admin-groups {
    group-name group-value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Configure administrative groups to implement link coloring of resource classes.

Options

group-name—Name of the group. You can assign up to 32 names. The names and their corresponding values must be identical across all routers within a single domain.

group-value—Value assigned to the group. The names and their corresponding values must be identical across all routers within a single domain.

- **Range:** 0 through 31

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Administrative Groups for LSPs

admin-group (for Interfaces) | 1971

admin-groups-extended

IN THIS SECTION

- Syntax | 1975
- Hierarchy Level | 1975
- Description | 1976
- Options | 1976
- Required Privilege Level | 1976
- Release Information | 1976

Syntax

```
admin-groups-extended group-name {
    group-value group-identifier;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls interface interface-name],
[edit logical-systems logical-system-name routing-options],
```

```
[edit protocols mpls interface interface-name],
[edit routing-options]
```

Description

Specifies the group name and group identifier for an administrative group. The group identifier must be within the range of values specified by the `admin-groups-extended-range` statement. The extended administrative group values are global and must be identically configured on all the supported routers participating in the network. The domain-wide extended administrative groups database, learned from other routers through IGP flooding, is used by CSPF for path computation.

Options

group-name	The range of configurable values is between 32 and 4,294,967,295. The range maximum must be greater than the range minimum.
group-value ***group-identifier***	The group identifier must be within the range of configurable values, 32 and 4,294,967,295.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Extended Administrative Groups for LSPs

Configuring Administrative Groups for LSPs

admin-group-extended | **1968**

admin-groups-extended-range | **1977**

admin-groups-extended-range

IN THIS SECTION

- Syntax | 1977
- Hierarchy Level | 1977
- Description | 1977
- Options | 1978
- Required Privilege Level | 1978
- Release Information | 1978

Syntax

```
admin-groups-extended-range {
    maximum maximum-number;
    mininum minimum-number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit routing-options]
```

Description

Enables you to configure extended administrative groups, represented by a 32-bit value, expanding the number of administrative groups supported in the network beyond just 32. In MPLS traffic engineering, a link can be configured with a set of administrative groups (also known as colors or resource classes). Administrative groups are carried in IGPs (OSPFv2 and IS-IS) as a 32-bit value assigned to each link. By default, Juniper Networks routers interpret this 32-bit value as a bit mask with each bit representing a group. This normally limits each network to a total of 32 distinct administrative groups (value range 0 through 31).

The extended administrative groups configuration accepts a set of interfaces with a corresponding set of extended administrative group names. It converts the names into a set of 32-bit values and propagates this information into the IGP. The extended administrative group values are global and must be identically configured on all the supported routers participating in the network. The domain-wide extended administrative groups database, learned from other routers through IGP flooding, is used by CSPF for path computation.

Options

maximum *maximum-number*	The range of configurable values is between 32 and 4,294,967,295. The range maximum must be greater than the range minimum.
minimum *minimum-number*	The range of configurable values is between 32 and 4,294,967,295. The range maximum must be greater than the range minimum.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Extended Administrative Groups for LSPs

Configuring Administrative Groups for LSPs

admin-group-extended | **1968**

admin-search

IN THIS SECTION

Syntax | **1979**

- Hierarchy Level | 1979
- Description | 1979
- Options | 1979
- Required Privilege Level | 1979
- Release Information | 1980

Syntax

```
admin-search {
    distinguished-name distinguished-name;
    password password;
}
```

Hierarchy Level

```
[edit access ldap-options search],
[edit access profile          profile-name        ldap-options search]
```

Description

Specify that a Lightweight Directory Access Protocol (LDAP) administrator search is performed. To perform an administrator search, you must specify administrator credentials, which are used in the bind as part of performing the search.

Options

The remaining statements are explained separately.

- **Default:** Anonymous search. To perform an administrator search, you must specify administrator credentials, which are used in the bind as part of performing the search.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

administrator-email (Security Fallback Block)

IN THIS SECTION

- Syntax | 1980
- Hierarchy Level | 1980
- Description | 1981
- Required Privilege Level | 1981
- Release Information | 1981

Syntax

```
administrator-email email-address;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-block]
```

Description

Configure the administrator e-mail address that will be notified when a fallback-block occurs. This is an e-mail notification with a custom message and a custom subject line.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

administrator-email (Security Virus Detection)

IN THIS SECTION

- Syntax | 1981
- Hierarchy Level | 1982
- Description | 1982
- Required Privilege Level | 1982
- Release Information | 1982

Syntax

```
administrator-email email address;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus sophos-engine profile profile name notification-
options virus-detection]
```

Description

Configure the administrator e-mail address that will be notified when a virus is detected by Sophos antivirus. This is an e-mail notification with a custom message and a custom subject line.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

adsl-bytes

IN THIS SECTION

- Syntax | 1983
- Hierarchy Level | 1983
- Description | 1983
- Options | 1983
- Required Privilege Level | 1983
- Release Information | 1983

Syntax

```
adsl-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of cell overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for an ADSL access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `adsl-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream cell overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS

Traffic Rate Reporting and Adjustment by the ANCP Agent

Configuring the ANCP Agent

adsl2-bytes

IN THIS SECTION

- Syntax | 1984
- Hierarchy Level | 1985
- Description | 1985
- Options | 1985
- Required Privilege Level | 1985
- Release Information | 1985

Syntax

```
adsl2-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of cell overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for an ADSL2 access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

> **NOTE**: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `adsl2-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream cell overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS

Traffic Rate Reporting and Adjustment by the ANCP Agent

Configuring the ANCP Agent

adsl2-plus-bytes

IN THIS SECTION

- Syntax | 1986
- Hierarchy Level | 1987
- Description | 1987
- Options | 1987
- Required Privilege Level | 1987
- Release Information | 1987

Syntax

```
adsl2-plus-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of cell overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for an ADSL2+ access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

> **NOTE**: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-plus-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `adsl2-plus-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream cell overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-plus-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS

Traffic Rate Reporting and Adjustment by the ANCP Agent

Configuring the ANCP Agent

advance-policy-based-routing

IN THIS SECTION

- Syntax | **1988**
- Hierarchy Level | **1992**
- Description | **1992**
- Options | **1992**
- Required Privilege Level | **1993**
- Release Information | **1993**

Syntax

```
advance-policy-based-routing {
    active-probe-params name {
        settings {
            burst-size <size>;
            data-fill fill;
```

```
            data-size <size>;
            dscp-code-points dscp;
            forwarding-class fc-name;
            loss-priority {
            priority (high | low | medium-high | medium-low);
            }
            per-packet-loss-timeout <timeout>;
            probe-count <count>;
            probe-interval <interval>;
            sla-export-interval interval;
        }
     }
destination-path-group name {
    active-probe-properties {
        active-probe-only;
    }
    inline-gre-encap;
    overlay-path name;
    probe-routing-instance <routing-instance-name>;
    }
from-zone name {
    policy name {
        description description;
        match {
            source-address;
            destination-address;
            application [ application ... ];
            destination-address-excluded;
            source-address-excluded
            source-identity;
        }
        scheduler-name;
        then {
            application-services {
                advance-policy-based-routing-profile advance-policy-based-routing-profile;
            }
        }
    }
}
interface name {
    unit number {
        link-tag link-tag;
        multipath-bandwidth bandwidth-value
```

```
        priority priority-number;
        weight weight-number
    }
}
metrics-profile name {
    sla-threshold {
        delay-round-trip <delay-val>;
        jitter <jitter-val>;
        jitter-type (egress-jitter | ingress-jitter | two-way-jitter);
        match (all | any-one);
        packet-loss <loss>;
    }
}
multipath-rule name {
      application application-name;
      application-group application-group-name;
      bandwidth-limit bandwidth-limit;
      enable-reverse-wing;
      link-preferences;
      link-type (IP | MPLS);
      number-of-paths number-of-paths
}
overlay-path name {
    probe-path {
        local {
            ip-address ip_address;
        }
        remote {
            ip-address ip_address;
        }
    }
    tunnel-path {
        local {
            ip-address ip_address;
            }
        remote {
            ip-address ip_address;
            }
        }
    }
profile name {
    rule name {
        disable-midstream-routing;
```

```
        match {
            category;
            dscp [ dscp ... ];
            dynamic-application;
            dynamic-application-group;
        }
        then {
            application-services-bypass;
            routing-instance routing-instance-name;
            sla-rule <sla-rule-name>;
            }
       }
}
sla-options {
    log (disabled | syslog);
    max-passive-probe-limit {
        interval <refresh-interval>;
        number-of-probes <probes>;
    }
}
sla-rule name {
    active-probe-params <probe-params-name>;
    link-preferences {
        affinity strict;
        preferred-tag preferred-tag;
    }
    metrics-profile <metrics-profile-name>;
    multipath-rule multipath-rule-name;
    passive-probe-params {
        sampling-frequency {
            interval <sampling-interval>;
            ratio <sampling-ratio>;
        }
        sampling-percentage <percentage>;
        sampling-period <period>;
        sla-export-factor <factor>;
        type book-ended;
        }
    switch-idle-time <timeout>;
    type saas;
    violation-count <count>;
    }
traceoptions {
```

```
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
    flag name;
    no-remote-trace;
    }
tunables {
    drop-on-zone-mismatch;
    enable-logging;
    max-route-change max-route-change;
    no-l3l4-app-lookup;
    }
}
```

Hierarchy Level

[edit security]

Description

Configure an advanced policy-based routing.

You can create an advanced policy-based routing (APBR) profile (application profile) to match applications and application groups and redirect those matching traffic to the specified routing instance for the route lookup. The profile includes multiple rules. Each rule can contain multiple applications or application groups. If the application matches any of the application or application groups of a rule in a profile, the application profile rule is considered to be a match.

The APBR profile evaluates the application-aware traffic and permits or denies traffic based on the applications and application groups.

The application profile can be attached to a security zone or it can be attached to a specific logical or physical interface associated with the security zone.

Options

profile *profile-name*	Name of the profile. Must be a unique name with a maximum length of 63 characters.
from-zone	Specify a source zone to be associated with the APBR policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Application Quality of Experience

Understanding Advanced Policy-Based Routing

advance-policy-based-routing-profile (Security Zones)

IN THIS SECTION

- Syntax | 1993
- Hierarchy Level | 1993
- Description | 1994
- Required Privilege Level | 1994
- Release Information | 1994

Syntax

```
advance-policy-based-routing-profile profile;
```

Hierarchy Level

[edit security zones security-zone *zone-name*]

Description

Enable Advance Policy Based Routing on the specified security zone.

To classify and redirect the traffic, the APBR profile matches applications and application groups and if the matching rule is found, the packets are routed to the routing instance that sends the traffic to a different interface as specified in the next-hop IP address. So, you must associate the application profile to the ingress traffic—that is, attach the application profile to a security zone.

When the application profile is applied to a security zone, then all interfaces belonging to that zone are attached to the application profile by default unless there is a specific configuration for an interface belonging to that zone.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Example: Configuring Advanced Policy-Based Routing for Application-Aware Traffic Management Solution

Understanding Advanced Policy-Based Routing

advanced-anti-malware connection

IN THIS SECTION

- Syntax | **1995**
- Description | **1995**
- Options | **1996**
- Required Privilege Level | **1996**

- Output Fields | 1996
- Sample Output | 1996
- Release Information | 1996

Syntax

```
set services advanced-anti-malware connection (authentication | source-address | source-
interface | url)
```

Description

Configures the cloud connection, including the hostname, port number and authentication information. Once configured, a certificate-based mutual authentication and TLS connection is established between the SRX Series Firewall and Juniper ATP Cloud. The SRX Series Firewall receives configurations from the cloud and sends health data back to the cloud.

You can also use the op script to perform the same operations as the `set services advanced-anti-malware connection` command. To download the Junos op script, see Download and Run the Juniper ATP Cloud Script.

For SRX300, SRX320, SRX340, SRX345 and SRX500 Series Firewalls, you must run the `set security forwarding-process enhanced-services-mode` command before running the op script or before running the `set services advanced-anti-malware connection` command. For example:

```
user@host> set security forwarding-process enhanced-services-mode
user@host> set services advanced-anti-malware connection url https://xxx.xxxx.junipersecurity.net
user@host> set services advanced-anti-malware connection authentication tls-profile aamw-ssl
```

We recommend that you use the op script instead of using the `set services advanced-anti-malware connection` command.

Use the `show services advanced-anti-malware status` CLI command to verify that connection is made to the cloud server from the SRX Series Firewall. If you do not see entries, we recommend you rerun the op script.

Options

authentication tls-profile ***profile-name***	Name of the TLS profile that contains settings for the TLS-secured connection.
source-address ***address***	The source IP address to send files to the cloud.
source-interface ***interface***	The source interface to send files to the cloud. If you configure the `source-interface` but not the `source-address`, the SRX Series Firewall uses the IP address from the specified interface for connections.
url ***url***	The URL of the cloud. You can optionally specify a port if needed. For example, `https://www.juniper-threat-prevent-cloud.net:443`.

Required Privilege Level

View

Output Fields

This command produces no output.

Sample Output

set services advanced-anti-malware connection url

```
user@host# set services advanced-anti-malware connection url https://sky.junipersecurity.net
```

set services advanced-anti-malware connection authentication

```
user@host# set services advanced-anti-malware connection authentication tls-profile aamw-ssl
```

Release Information

Command introduced in Junos OS Release 15.1X49-D33.

RELATED DOCUMENTATION

show services advanced-anti-malware status

advanced-anti-malware policy

IN THIS SECTION

- Syntax | 1997
- Hierarchy Level | 1999
- Description | 1999
- Options | 2000
- Required Privilege Level | 2004
- Release Information | 2004

Syntax

```
policy <policy-name>{
    blacklist-notification {
        log;
    }
    default-notification {
        log;
    }
    fallback-options {
        action [block | permit];
        invalid-content-size {
            action [block | permit];
            notification {
                log;
            }
        }
        notification {
            log;
        }
        out-of-resources {
```

```
            action [block | permit];
            notification {
                log;
            }
        }
        service-not-ready {
            action [block | permit];
            notification {
                log;
            }
        }
        submission-timeout {
            action [block | permit];
            notification {
                log;
            }
        }
        unknown-file {
            action [block | permit];
            notification {
                log;
            }
        }
        verdict-timeout {
            action [block | permit];
            notification {
                log;
            }
        }
    }
    http {
            action [block | permit];
            client-notify {
            file file-name;
            message block-message-to-client;
            redirect-url redirect-url;
        }
        file-verdict-unknown (block | permit);
        inspection-profile inspection-profile-name;
        notification {
            log;
        }
    }
```

```
    imap {
        inspection-profile inspection-profile-name;
        notification {
            log;
        }
    }
    smb {
        inspection-profile inspection-profile-name;
        notification {
            log;
        }
    }
    smtp {
        inspection-profile inspection-profile-name;
        notification {
            log;
        }
    }
    whitelist-notification {
        log;
    }
```

Hierarchy Level

```
[edit services advanced-anti-malware]
```

Description

The connection to the Juniper Advanced Threat Prevention Cloud is launched on-demand. It is established only when a condition is met and a file or URL must be sent to the cloud. The cloud inspects the file and returns a verdict number (1 through 10). A verdict number is a score or threat level. The higher the number, the higher the malware threat. The SRX Series Firewall compares this verdict number to the Juniper Advanced Threat Prevention Cloud policy settings and either permits or denies the session. If the session is denied, a reset packet is sent to the client and the packets are dropped from the server.

Juniper Advanced Threat Prevention Cloud policies append to the Junos OS security policies by defining the actions to take when a file is considered malware or when an attempt is made to download a file from a location that's on a custom blocklist or allowlist.

Use this command to configure the Juniper Advanced Threat Prevention Cloud policy.

Options

policy-name Name of the Juniper Advanced Threat Prevention Cloud policy.

> **NOTE**: Starting in Junos OS Release 18.2R1, for unified policies, a `default-policy` can be used for anti-malware and security-intelligence policies. The commands are: `set services security-intelligence default-policy` and `set services advanced-anti-malware default-policy`. During the initial policy lookup phase, which occurs prior to a dynamic application being identified, if there are multiple policies present in the potential policy list which contains different security intelligence or anti-malware policies, the SRX Series Firewall applies the default policy until a more explicit match has occurred. See the *Juniper ATP Cloud Administration Guide* and your SRX Series documentation for more details on unified policies.

blacklist-notification (Optional) Create a system log entry when an attempt to access a website listed in the blocklist file is made. Use `blacklist-notification log` to create a log entry. If you do not want to create a log entry, do not specify the `blacklist-notification` option.

default-notification Create a system log entry if the cloud returns a verdict number less than the *verdict-threshold*. Use `default-notification log` to create a log entry. If you do not want to create a log entry, do not specify the `default-notification` option.

fallback-options (action block | action permit) The action to take when the SRX Series Firewall runs out of resources or the connection to the cloud is lost. The default is `action permit`.

fallback-options (Optional) Create a system log entry when fallback occurs. Use `fallback-options notification log` to create a log entry. If you do not want to create a log entry, do not specify the `fallback-options notification` option.

http This command allows you to inspect advanced anti-malware (AAMW) files downloaded by hosts through Hypertext Transfer Protocol (HTTP) protocol. The AAMW files are then submitted to Juniper ATP Cloud for malware screening.

http action (permit | block) This command allows you to permit or block malware based on the detected malware.

Starting in Junos OS release 21.3R1, when the http action is set as block, the detected malware file will be blocked immediately after signature match. Juniper ATP cloud will not get the full file for analysis, instead it will receive a notification of the malware hit event. You can view the malware information Juniper ATP Cloud Portal. The Partial file

tab displays the malware hit event information for all blocked signature match detections.

http(s) client-notify (message | file | redirect-url)

(Starting in Junos OS release 19.3R1) This command allows you to configure HTTP URL redirection for a customized client notification based on detected malware with the block action.

http(s) file-verdict-unknown (permit | block)

(Starting in Junos 19.3R1) This command allows you to permit or block malware based on the detected malware having a verdict of “unknown.” By default, “unknown” malware is permitted.

smb

(Starting in Junos 21.1R1) This command allows you to inspect advanced anti-malware (AAMW) files downloaded by hosts through Server Message Block (SMB) protocol. The AAMW files are then submitted to Juniper ATP Cloud for malware screening.

> **NOTE**:
>
> - Starting in Junos OS Release 21.3R1, SMB supports block mode in addition to permit mode.
> - SMBv3 encryption is not supported.
> - SMB multichannel is not supported.

inspection-profile

Name of the Juniper Advanced Threat Prevention Cloud profile. This profile defines what file types or file categories are to be sent to the cloud for inspection.

match *verdict-threshold*

The *verdict-threshold* defines the number at which you want to label a file as malware. For example, if you set *verdict-threshold* to 7 and the cloud returns a verdict number of 7 or greater, then that file is considered malware. *verdict-threshold* can be any number between 1 and 10, inclusive.

then notification

(Optional) Create a system log entry if the cloud returns a verdict number equal to or greater than the *verdict-threshold*. Use `then notification log` to create a log entry. If you do not want to create a log entry, do not specify the `then notification` option.

whitelist-notification

(Optional) Create a system log entry when an attempt to access a website listed in the allowlist file is made. Use `whitelist-notification log` to create a log entry. If you do not want to create a log entry, do not specify the `whitelist-notification` option.

Table 1 on page 2002 shows examples of using the Juniper Advanced Threat Prevention Cloud policy options.

Table 1: Juniper Advanced Threat Prevention Cloud Security Policy Additions

Addition	Description

Action and notification based on the verdict number and threshold	Defines the threshold value and what to do when the verdict number is greater than or equal to the threshold. For example, if the threshold is 7 and Juniper Advanced Threat Prevention Cloud returns a verdict number of 9 for a file, then that file is blocked from being downloaded and a log entry is created.
	
	`set services advanced-anti-malware policy aamwpol1 match verdict-threshold 7`
	`set services advanced-anti-malware policy aamwpol1 then action block`
	`set services advanced-anti-malware policy aamwpol1 then notification log`

Default action and notification	Defines what to do when the verdict number is less than the threshold. For example, if the threshold is 7 and Juniper Advanced Threat Prevention Cloud returns a verdict number of 3 for a file, then that file is allowed to be downloaded and create a log entry. `set services advanced-anti-malware policy aamwpol1 default-notification log`
Name of the inspection profile	Name of the Juniper Advanced Threat Prevention Cloud profile that defines the types of file to scan. `set services advanced-anti-malware policy aamwpol1  inspection-profile profile1`

Table 1: Juniper Advanced Threat Prevention Cloud Security Policy Additions *(Continued)*

Addition	Description

Fallback options	Defines what to do when error conditions occur or when there is a lack of resources. The following fallback options are available:
	• action—Permit or block the file regardless of its threat level.
	• notification—Add or do not add this event to the log file.
	
	`set services advanced-anti-malware policy aamwpol1 fallback-options action block`
	`set services advanced-anti-malware policy aamwpol1 fallback-options notification log`
	
	The following fallback options are available:
	• invalid-content-size—Content size exceed supported range
	• out-of-resources—Service is out of resources
	• service-not-ready—Service is not yet ready
	• submission-timeout—Submission is timed out
	• unknown-file—File type is unknown
	• verdict-timeout—Verdict is timed out
	
	If none of listed fallback condition is hit, then the default fallback option is applied.

Blocklist notification	Defines whether to create a log entry when attempting to download a file from a site listed in the blocklist file. `set services advanced-anti-malware policy aamwpol1 blacklist-notification log`
Whitelist notification	Defines whether to create a log entry when attempting to download a file from a site listed in the allowlist file. `set services advanced-anti-malware policy aamwpol1 whitelist-notification log`

Table 1: Juniper Advanced Threat Prevention Cloud Security Policy Additions *(Continued)*

Addition	Description

User notification of malware on block action	(Starting in Junos 19.3R1) This command allows you to configure HTTP and HTTPS URL redirection for a customized client notification based on detected malware with the block action. A block message can only be sent when a block action is configured.
	
	NOTE: See "request services advanced-anti-malware redirect-file" on page 14910 for details on adding a custom file.
	
	`set services advanced-anti-malware policy p1 http client-notify message`
	`set services advanced-anti-malware policy p1 http client-notify file`
	`set services advanced-anti-malware policy p1 http client-notify redirect-url <enter URL>`

Block or permit malware when file verdict is “unknown”	(Starting in Junos 19.3R1) This command allows you to permit or block malware based on the detected file having a verdict of “unknown.” By default, an “unknown” file verdict is permitted. (Note this only applies to HTTP and HTTPS traffic.) `set services advanced-anti-malware policy p1 http file-verdict-unknown <block|permit>`

Required Privilege Level

View

Release Information

Command introduced in Junos OS Release 15.1X49-D33.

The `SMB` option is introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

show services advanced-anti-malware policy | **25313**

advanced-anti-malware traceoptions

IN THIS SECTION

- Syntax | 2005
- Description | 2005
- Options | 2005
- Required Privilege Level | 2007
- Output Fields | 2007
- Release Information | 2007

Syntax

```
set services advanced-anti-malware traceoptions
    <apply-groups name>
     <apply-groups-except name>
    <file (trace-file-name | files | match | no-world-readable | size | world-readable)>
    <flag (all | connection | content | daemon | identification | parser | plugin | policy)>
    <level level>
    <no-remote-trace>
```

Description

In Junos, the concept of setting traceoptions requires setting the trace in the configuration itself, rather than as an operational mode command. When the trace is set in the configuration, the “flags” are defined for the actual debug that you want to perform. This command allows you to trace the Juniper Advanced Threat Prevention Cloud configuration and is typically used for troubleshooting.

Options

apply-groups *name*	Groups that these trace option settings apply.
apply-groups-	Groups that these trace option settings do not apply.

except *name*

file Defines the trace file characteristics. Valid options are:

- *trace-file-name*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. Valid values range from 1 and 1024 characters. The name cannot include spaces, /, or % characters. The default filename is security.
- files—Maximum number of trace files that can accumulate. Valid values range from 2 to 1000. The default value is 3.
- match—The criteria that you want the system to use when logging information to the file. You can enter a regular expression. Wildcard (*) characters are accepted.
- no-world-readable—Only the system administrator can read the trace file.
- size—Maximum size to which the trace file can grow. Once the file reaches the specified size, it is compressed and renamed *filename*0.gz, the next file is named *filename*1.gz, and so on. Valid values range from 10240 to 1,073,741,824.
- world-readable—Any user can read the trace file.

flag *flag* Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- all—Trace everything.
- connection—Trace connection to the Juniper Advanced Threat Prevention Cloud server.
- content—Trace the content the SRX is buffering to the client and to the cloud.
- daemon—Trace the Juniper Advanced Threat Prevention Cloud daemon.
- identification—Trace the file type identification. Examples of file types are .exes, .java, .tar and so forth.
- parser—Trace the interface between the Juniper Advanced Threat Prevention Cloud daemon and the module that parses the HTTP protocol and extracts the file content..
- plugin—Trace the Juniper Advanced Threat Prevention Cloud plugin to view session checks, packet processing and actions taken against a file or URL.
- policy—Trace the Juniper Advanced Threat Prevention Cloud security policy.

level *level* Specify level of tracing to perform. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- `all`—Match messages of all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure `verbose`, messages at all higher levels are traced. Therefore, the result is the same as when you configure `all`.
- `warning`—Match warning messages.

`no-remote-trace` Disable remote tracing.

Required Privilege Level

`trace`—To view this statement in the configuration.

`trace-control`—To add this statement to the configuration.

Output Fields

This command produces no output. Use the `show log` *`trace-file-name`* command to view the debug file.

Release Information

Command introduced in Junos OS Release 15.1X49-D33.

advertise

IN THIS SECTION

- Syntax | 2008
- Hierarchy Level | 2008
- Description | 2008
- Default | 2008

Required Privilege Level | 2008
Release Information | 2008

Syntax

```
(advertise | ignore);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery address address],
[edit protocols router-discovery address address]
```

Description

Specify whether the server should advertise the IP address in its router advertisement packets:

- **advertise**—Advertise the IP address in its router advertisement packets.
- **ignore**—Do not advertise the IP addresses in router advertisement packets.

Default

advertise

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

advertise-bgp-static

IN THIS SECTION

- Syntax | 2009
- Hierarchy Level | 2009
- Description | 2010
- Options | 2010
- Required Privilege Level | 2010
- Release Information | 2010

Syntax

```
advertise-bgp-static {
    policy policy-expression;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address]
```

Description

Include this statement to always advertise a BGP-static route, even if it is not the active route for a prefix. You can configure this statement globally to advertise the BGP-static routes to all neighbors. You can also configure this statement to advertise BGP-static routes in a BGP group or to a specific neighbor in a BGP group.

Options

`policy` ***policy-expression*** Specify an additional export policy to control whether or not a given BGP-static route is to be advertised in preference to the active route for a prefix. The policy is applied to the BGP-static route and not to the active route. Only the accept or reject result of the policy expression is observed, and any side-effects, such as, modifying communities, are ignored.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

bgp-static | **2966**

Configuring BGP-Static Routes for Preventing Route Flaps

Example: Configuring BGP-Static Routes to Prevent Route Flaps

Understanding BGP-Static Routes for Preventing Route Flaps

advertise-external

IN THIS SECTION

- Syntax | 2011
- Hierarchy Level | 2011
- Description | 2011
- Default | 2012
- Options | 2012
- Required Privilege Level | 2013
- Release Information | 2013

Syntax

```
advertise-external {conditional};
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name  routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor neighbor-
address]
```

Description

Specify BGP to advertise the best external route into an IBGP mesh group, a route reflector cluster, or an AS confederation even if the best route is an internal route.

In general, deployed BGP implementations do not advertise the external route with the highest local preference value to internal peers unless it is the best route. Although this behavior was required by an earlier version of the BGP version 4 specification, RFC 1771, it was typically not followed in order to minimize the amount of advertised information and to prevent routing loops. However, there are scenarios in which advertising the best external route is beneficial, in particular, situations that can result in IBGP route oscillation.

The `advertise-external` statement is supported at both the group and neighbor level. If you configure the statement at the neighbor level, you must configure it for all neighbors in a group. Otherwise, the group is automatically split into different groups.

In a confederation, when advertising a route to a confederation border router, any route from a different confederation sub-AS is considered external. When configuring the `advertise-external` statement for an AS confederation, it is recommended that EBGP peers belonging to different autonomous systems are configured in a separate EBGP peer group. This ensures consistency while BGP sends the best external route to peers in the configured peer group.

To configure the `advertise-external` statement on a route reflector, you must disable intracluster reflection with the `no-client-reflect` statement.

When a routing device is configured as a route reflector for a cluster, a route advertised by the route reflector is considered internal if it is received from an internal peer with the same cluster identifier or if both peers have no cluster identifier configured. A route received from an internal peer that belongs to another cluster, that is, with a different cluster identifier, is considered external.

The `conditional` option causes BGP to advertise the external route only if the route selection process reaches the point where the multiple exit discriminator (MED) metric is evaluated. As a result, an external route with an AS path longer than that of the active path is not advertised.

Junos OS also provides support for configuring a BGP export policy that matches on the state of an advertised route. You can match on either active or inactive routes.

Default

BGP does not advertise the external route with the highest local preference value to internal peers unless it is the best route.

Options

`conditional`—(Optional) Advertise the best external path only if the route selection process reaches the point at which the multiple exit discriminator (MED) metric is evaluated. The `conditional` option restricts advertisement to when the best external path and the active path are equal until the MED step of the route selection process. This implies that external routes with a longer AS path length than the active

path, for instance, are not advertised. The criteria used for selecting the best external path is the same whether or not the `conditional` option is configured.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Example: Configuring a Routing Policy to Advertise the Best External Route to Internal Peers

advertise-inactive | 2015

advertise-from-main-vpn-tables

IN THIS SECTION

- Syntax | 2013
- Hierarchy Level | 2014
- Description | 2014
- Default | 2014
- Required Privilege Level | 2014
- Release Information | 2014

Syntax

```
advertise-from-main-vpn-tables;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit protocols bgp],
[edit routing-instances routing-instance-name protocols bgp]
```

Description

Advertise VPN routes from the main VPN tables in the master routing instance (for example, bgp.l3vpn.0, bgp.mvpn.0) instead of advertising VPN routes from the tables in the VPN routing instances (for example, *instance-name*.inet.0, *instance-name*.mvpn.0). Enable nonstop active routing (NSR) support for BGP multicast VPN (MVPN).

When this statement is enabled, before advertising a route for a VPN prefix, the path selection algorithm is run on all routes (local and received) that have the same route distinguisher (RD).

> **NOTE**: Adding or removing this statement causes all BGP sessions that have VPN address families to be removed and then added again. On the other hand, having this statement in the configuration prevents BGP sessions from going down when route reflector (RR) or autonomous system border router (ASBR) functionality is enabled or disabled on a routing device that has VPN address families configured.

Default

If you do not include this statement, VPN routes are advertised from the tables in the VPN routing instances.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Understanding Junos OS Routing Tables

Types of VPNs

advertise-inactive

IN THIS SECTION

- Syntax | 2015
- Hierarchy Level | 2015
- Description | 2016
- Default | 2016
- Required Privilege Level | 2016
- Release Information | 2016

Syntax

```
advertise-inactive;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
```

```
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Configure the routing table to export to BGP the best route learned by BGP even if Junos OS did not select this route to be an active route.

One way to achieve multivendor compatibility is to include the `advertise-inactive` statement in the external BGP (EBGP) configuration. By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. The `advertise-inactive` statement causes Junos OS to advertise the best BGP route that is inactive because of IGP preference. When you use the `advertise-inactive` statement, the Junos OS device uses, for example, the OSPF route for forwarding, and the other vendor's device uses the EBGP route for forwarding. However, from the perspective of an EBGP peer in a neighboring AS, both vendors' devices appear to behave the same way.

NOTE: When BGP advertises a network layer reachability information (NLRI) with a label, and the advertised route resides in xxx.xxx.3 routing table such as inet.3, Junos OS automatically advertises such inactive routes even if you have not configured the `advertise-inactive` statement.

NOTE: The `advertise-inactive` statement does not help to advertise the inactive route from the VRF when the router is configured as a route reflector.

Default

By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring BGP to Advertise Inactive Routes

Example: Configuring the Preference Value for BGP Routes

Example: Configuring BGP Route Preference (Administrative Distance)

advertise-external | **2011**

advertise-interval

IN THIS SECTION

- Syntax | **2017**
- Hierarchy Level | **2017**
- Description | **2017**
- Options | **2018**
- Required Privilege Level | **2018**
- Release Information | **2018**

Syntax

```
advertise-interval seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]
```

Description

Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv4 advertisement packets.

All routers in the VRRP group must use the same advertisement interval.

> **NOTE**: When VRRPv3 is enabled, the `advertise-interval` statement cannot be used to configure advertisement intervals. Instead, use the `fast-interval` statement to configure advertisement intervals.

Options

seconds—Interval between advertisement packets.

- **Range:** 1 through 255 seconds
- **Default:** 1 second

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Advertisement Interval for the VRRP Primary Router

fast-interval | **6926**

inet6-advertise-interval (VRRP) | **8596**

version-3 | **30691**

advertise-mode (MPLS)

IN THIS SECTION

Syntax | **2019**

- Hierarchy Level | 2019
- Description | 2019
- Options | 2019
- Required Privilege Level | 2020
- Release Information | 2020

Syntax

```
advertise-mode (stub-alias | stub-proxy);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls egress-protection context-identifier
context-id],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name egress-
protection context-identifier context-id],
[edit protocols mpls egress-protection context-identifier context-id],
[edit protocols mpls label-switched-path lsp-name egress-protection context-identifier context-
id]
```

Description

Configure the method for the interior gateway protocol (IGP) to advertise egress protection availability.

Egress protection availability is advertised in the IGP. Label protocols along with CSPF use this information to do egress protection.

Options

stub-alias Context identifier has an alias.

In the alias method, the LSP end-point address has an explicit backup egress node where the backup node can be learned or configured on the penultimate hop node (PHN) of a protected LSP. With this model, the PHN of a protected LSP sets up the bypass LSP tunnel to back up the

egress node by avoiding the primary egress node. This model requires a Junos OS upgrade in core nodes, but is flexible enough to support all traffic engineering constraints.

stub-proxy Context-identifier has a stub proxy node.

A stub node is one that only appears at the end of an AS path, which means it does not provide transit service. In this mode, known as the virtual or proxy mode, the LSP end-point address is represented as a node with bidirectional links, with the LSP's primary egress node and backup egress node. With this representation, the penultimate hop of the LSP primary egress point can behave like a PLR in setting up a bypass tunnel to back up the egress by avoiding the primary egress node. This model has the advantage that you do not need to upgrade Junos OS on core nodes and thereby helps operators to deploy this technology.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Egress Protection for Layer 3 VPN Edge Protection Overview

Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP

advertise-peer-as

IN THIS SECTION

- Syntax | **2021**
- Hierarchy Level | **2021**
- Description | **2021**
- Default | **2022**

- Required Privilege Level | **2022**
- Release Information | **2022**

Syntax

```
advertise-peer-as;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Disable the default behavior of suppressing AS routes.

If you include the `advertise-peer-as` statement in the configuration, BGP advertises routes learned from one external BGP (EBGP) peer back to another EBGP peer in the same autonomous system (AS) but not back to the originating peer.

Another way to disable the route suppression default behavior is with the `as-override` statement. If you include both the `as-override` and `no-advertise-peer-as` statements in the configuration, the `no-advertise-peer-as` statement is ignored.

Default

By default, Junos OS does not advertise the routes learned from one EBGP peer back to the same external BGP (EBGP) peer. In addition, the software does not advertise those routes back to any EBGP peers that are in the same AS as the originating peer, regardless of the routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Enabling BGP Route Advertisements

Example: Configuring a Layer 3 VPN with Route Reflection and AS Override

no advertise-peer-as | **11385**

advertise-prefix-sid

IN THIS SECTION

- Syntax | **2023**
- Hierarchy Level | **2023**
- Description | **2023**
- Required Privilege Level | **2023**
- Release Information | **2023**

Syntax

```
advertise-prefix-sid;
```

Hierarchy Level

```
[edit fabric protocols bgp],
[edit logical-systems name protocols bgp],
[edit logical-systems name routing-instances name protocols bgp],
[edit logical-systems name tenants name routing-instances name protocols bgp],
[edit protocols bgp],
[edit routing-instances name protocols bgp],
[edit tenants name routing-instances name protocols bgp]
```

Description

Advertise prefix sid to E-BGP peers

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

srgb (Protocols BGP) | **28407**

advertise-srv6-service

IN THIS SECTION

- Syntax | 2024
- Hierarchy Level | 2024
- Description | 2025
- Required Privilege Level | 2025
- Release Information | 2026

Syntax

```
advertise-srv6-service;
```

Hierarchy Level

```
[edit fabric protocols bgp family inet unicast],
[edit fabric protocols bgp family inet6 unicast],
[edit fabric protocols bgp group name family inet unicast],
[edit fabric protocols bgp group name family inet6 unicast],
[edit fabric protocols bgp group name neighbor name family inet unicast],
[edit fabric protocols bgp group name neighbor name family inet6 unicast],
[edit logical-systems name protocols bgp family inet unicast],
[edit logical-systems name protocols bgp family inet6 unicast],
[edit logical-systems name protocols bgp group name family inet unicast],
[edit logical-systems name protocols bgp group name family inet6 unicast],
[edit logical-systems name protocols bgp group name neighbor name family inet unicast],
[edit logical-systems name protocols bgp group name neighbor name family inet6 unicast],
[edit logical-systems name routing-instances name protocols bgp family inet unicast],
[edit logical-systems name routing-instances name protocols bgp family inet6 unicast],
[edit logical-systems name routing-instances name protocols bgp group name family inet unicast],
[edit logical-systems name routing-instances name protocols bgp group name family inet6 unicast],
[edit logical-systems name routing-instances name protocols bgp group name neighbor name family inet
unicast],
[edit logical-systems name routing-instances name protocols bgp group name neighbor name family inet6
```

```
unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp family inet unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp family inet6 unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp group name family
inet unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp group name family
inet6 unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp group name neighbor
name family inet unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp group name neighbor
name family inet6 unicast],
[edit protocols bgp family inet unicast],
[edit protocols bgp family inet6 unicast],
[edit protocols bgp group name family inet unicast],
[edit protocols bgp group name family inet6 unicast],
[edit protocols bgp group name neighbor name family inet unicast],
[edit protocols bgp group name neighbor name family inet6 unicast],
[edit routing-instances name protocols bgp family inet unicast],
[edit routing-instances name protocols bgp family inet6 unicast],
[edit routing-instances name protocols bgp group name family inet unicast],
[edit routing-instances name protocols bgp group name family inet6 unicast],
[edit routing-instances name protocols bgp group name neighbor name family inet unicast],
[edit routing-instances name protocols bgp group name neighbor name family inet6 unicast],
[edit tenants name routing-instances name protocols bgp family inet unicast],
[edit tenants name routing-instances name protocols bgp family inet6 unicast],
[edit tenants name routing-instances name protocols bgp group name family inet unicast],
[edit tenants name routing-instances name protocols bgp group name family inet6 unicast],
[edit tenants name routing-instances name protocols bgp group name neighbor name family inet
unicast],
[edit tenants name routing-instances name protocols bgp group name neighbor name family inet6
unicast]
```

Description

Enable the egress BGP peer to advertise SRv6 service to a remote provider edge (PE) device. BGP advertises the reachability of prefixes of a particular service from an egress PE device to ingress PE nodes through BGP messages. BGP egress nodes advertise a set of reachable prefixes when you enable BGP Layer 3 services over the SRv6 data plane.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

srv6 (Protocols BGP) | 28417

accept-srv6-service | 1649

Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP

advertise-to-non-llgr-neighbor (Graceful Restart for BGP Helper)

IN THIS SECTION

- Syntax | 2026
- Hierarchy Level | 2026
- Description | 2027
- Options | 2027
- Required Privilege Level | 2027
- Release Information | 2027

Syntax

```
advertise-to-non-llgr-neighbor {
    omit-no-export;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp graceful-restart long-lived],
[edit logical-systems logical-system-name protocols bgp group group-name graceful-restart long-
lived],
```

```
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
graceful-restart long-lived],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp graceful-
restart long-lived],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name graceful-restart long-lived],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name neighbor address graceful-restart],
[edit protocols bgp graceful-restart long-lived],
[edit protocols bgp group group-name graceful-restart long-lived],
[edit protocols bgp group group-name neighbor address graceful-restart long-lived]
```

Description

Enable the BGP long-lived graceful restart (LLGR) stale routes to be advertised to neighbors that do not advertise the LLGR capability. This setting applies to both routes that were marked LLGR-stale by this router, and LLGR-stale routes received from neighbors. Ideally, all routers in an autonomous system support the IETF draft specification before it was enabled. However, to facilitate incremental deployment, stale routes might be required to be advertised to neighbors that have not advertised the long-lived graceful restart capability under the following conditions: The neighbors must be internal (IBGP or Confederation) neighbors. The NO_EXPORT community must be attached to the stale routes. The stale routes must have their LOCAL_PREF attribute set to zero. If this technique for partial deployment is used, you must set LOCAL_PREF to zero for all LLGR routes throughout the autonomous system. This configuration trades off a small reduction in flexibility (ordering may not be preserved between competing LLGR routes) for consistency between routers that support and do not support this specification. Because consistency of route selection can be important for preventing forwarding loops, the latter consideration of routers that do not support this specification precedes.

Options

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1 .

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

advertise-zero-ethernet-tag

IN THIS SECTION

- Syntax | 2028
- Hierarchy Level | 2028
- Description | 2028
- Required Privilege Level | 2029
- Release Information | 2029

Syntax

```
advertise-zero-ethernet-tag;
```

Hierarchy Level

```
[edit routing-instances name protocols evpn]
```

Description

Advertise EVPN Type 2 (MAC/IP Advertisement) and EVPN Type 3 (Inclusive Multicast Ethernet Tag) routes with an Ethernet Tag ID value of 0. This statement only applies to `instance-type` EVPN routing instances and must have a valid `vlan-id` and `no-normalization` configured.

You use the `advertise-zero-ethernet-tag` configuration statement when you need a vlan-based service with a valid `vlan-id` providing Layer 3 gateway functionality while also providing RFC 7432 compliance for Layer 2 gateway functionality with control plane EVPN routes advertised with an Ethernet Tag ID value of 0.

The following is an excerpt from a sample configuration:

```
user@mx960> show configuration routing-instances evpna
instance-type evpn;
protocols {
    evpn {
            advertise-zero-ethernet-tag;
    }
}
vlan-id 600;
routing-interface irb.600;
no-normalization;
interface ge-0/0/1.600;
route-distinguisher 1:1;
vrf-target target:1:1;
```

NOTE: The `mod` based `df-election-type` uses the `vlan-id` in the designated forwarder (DF) calculation. This can cause an inconsistency in the DF election when one provider edge (PE) device is configured with `vlan-id none` and the other is configured with a valid `vlan-id`. Therefore, we recommend having the same `vlan-id` configuration under the EVPN routing instance to ensure a consistent DF election.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1

RELATED DOCUMENTATION

VLAN-Based Service for EVPN

advertisement

IN THIS SECTION

- Syntax | 2030
- Hierarchy Level | 2030
- Description | 2031
- Options | 2031
- Required Privilege Level | 2031

Syntax

```
advertisement {
          always;
          application-specific {
             all-applications {
                 legacy;
              }
          }
      }
```

Hierarchy Level

```
[edit logical-systems name protocols isis traffic-engineering],
[edit logical-systems name routing-instances name protocols isis traffic-engineering],
[edit logical-systems name tenants name routing-instances name protocols isis traffic-engineering],
[edit protocols isis traffic-engineering],
[edit routing-instances name protocols isis traffic-engineering],
[edit tenants name routing-instances name protocols isis traffic-engineering]
```

Description

Advertise TE parameters even if RSVP is not turned on.

Options

always Advertise TE parameters in TE LSAs.

application-specific all-applications Use this option to advertise ASLA subtlv. This causes ISIS to advertise all legacy TE attributes inside ASLA sub-TLV as ASLA sub-sub-TLVs.

application-specific all-applications legacy Use this option to advertise ASLA subtlv with L flag set. This would mean an empty ASLA sub-TLV will be advertised along with legacy TE-Attribute sub-TLVs.

Required Privilege Level

routing

RELATED DOCUMENTATION

traffic-engineering (Protocols IS-IS)

advertisement (Delay Measurement)

IN THIS SECTION

- Syntax | 2032
- Hierarchy Level | 2032
- Description | 2032
- Options | 2032
- Required Privilege Level | 2033

Syntax

```
advertisement {
 accelerated {
  threshold percentage;
 }
 periodic {
  interval seconds;
  threshold percentage;
 }
}
```

Hierarchy Level

```
[edit logical-systems name protocols isis interface name delay-measurement],
[edit logical-systems name routing-instances name protocols isis interface name delay-
measurement],
[edit logical-systems name tenants name routing-instances name protocols isis interface name
delay-measurement],
[edit protocols isis interface name delay-measurement],
[edit routing-instances name protocols isis interface name delay-measurement],
[edit tenants name routing-instances name protocols isis interface name delay-measurement]
```

Description

Configure link delay advertisement settings.

Options

accelerated Notify the routing process of a change in the link delay measurement parameters for every probe exceeding the configured threshold.

NOTE: Configuring the accelerated advertisement option for link delay measurement enables aggressive notifications that result in excessive network churn.

(Accelerated advertisement) threshold

Specify the accelerated threshold used for the probes.

- Range: 0 through 100 percent
- Default: 20 percent

periodic

Notify the routing process of a change in the link delay measurement parameters only after crossing the specified interval, and only when the configured threshold is exceeded.

- Default: By default, periodic advertisement is enabled.

interval

Specify the time to wait between the periodic advertisement of link delay measurements.

- Range: 30 through 3600 seconds
- Default: 120 seconds

(Periodic advertisement) threshold

Specify the periodic threshold used for the probes.

- Range: 0 through 100 percent
- Default: 10 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Link Delay Measurement and Advertising in IS-IS

advertisement-hold-time

IN THIS SECTION

Syntax | 2034

Hierarchy Level | 2034

- Description | 2034
- Options | 2034
- Required Privilege Level | 2034
- Release Information | 2034

Syntax

```
advertisement-hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Do not advertise when the LSP goes from up to down, for a certain period of time known as the hold time.

Options

seconds—Hold time, in seconds.

- **Range:** 0 through 65,535 seconds
- **Default:** 5 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Damping Advertisement of LSP State Changes

advisory-options (Traffic Shaping)

IN THIS SECTION

- Syntax | 2035
- Hierarchy Level | 2035
- Description | 2035
- Required Privilege Level | 2036
- Release Information | 2036

Syntax

```
advisory-options {
    downstream-rate rate;
    upstream-rate rate;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-interface-
unit],
[edit dynamic-profiles profile-name interfaces interface-set $junos-interface-set-name interface
$junos-interface-ifd-name],
[edit interfaces demux0 unit logical-unit-number],
[edit interfaces interface-name unit logical-unit-number]
```

Description

Specify a recommended shaping rate to be applied to downstream or upstream traffic on an interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the `[edit interfaces demux0 ...]` hierarchy level introduced in Junos OS Release 12.2.

Support at the `[edit dynamic-profiles ...]` hierarchy level introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces

Configuring the ANCP Agent

Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS

advpn

IN THIS SECTION

- Syntax | 2037
- Hierarchy Level | 2037
- Description | 2037
- Options | 2037
- Required Privilege Level | 2038
- Release Information | 2038

Syntax

```
advpn {
    suggester {
        disable;
    }
    partner {
        connection-limit number;
        idle-threshold   packets/sec;
        idle-time        seconds;
        disable;
    }
}
```

Hierarchy Level

```
[edit security ike gateway gateway-name]
```

Description

Enable Auto Discovery VPN (ADVPN) protocol on the specified gateway. ADVPN dynamically establishes VPN tunnels between spokes to avoid routing traffic through the Hub. When you enable ADVPN, by default, the Junos OS enables both the suggester and partner roles on the device. You cannot use the same device with both the functions together. To use a specific role, you must explicitly disable the other.

Options

suggester VPN peer that can initiate a shortcut exchange to allow shortcut partners to establish dynamic security associations (SAs) with each other. Specify `disable` to disable this role on the gateway.

Both suggester and partner roles are enabled if `advpn` is configured without explicitly configuring `suggester` or `partner` keywords. We do not support suggester and partner roles on the same gateway. You must explicitly configure `disable` with the `suggester` or `partner` keyword to disable that particular role. You cannot disable both suggester and partner roles on the same gateway.

partner VPN peer that can receive a shortcut exchange suggesting that it should establish dynamic SAs with another peer. Specify `disable` to disable this role on the gateway.

The following options can be configured for the partner role:

connection-limit Maximum number of shortcut tunnels that can be created with different shortcut partners using a particular gateway. The maximum number, which is also the default, is platform-dependent.

Reducing the configured `connection-limit` value causes all active shortcut tunnels to be brought down. For example, if `connection-limit` is configured as 100 and you later reconfigure the number to 80, all active shortcut tunnels are brought down. Increasing the configured `connection-limit` value does not cause shortcut tunnels to go down.

idle-threshold Rate, in packets per second, below which the shortcut is brought down.

- **Range:** 3 through 5,000 packets per second.
- **Default:** 5 packets per second.

idle-time Duration, in seconds, after which the shortcut is deleted if the traffic remains below the `idle-threshold` value.

- **Range:** 60 seconds through 86,400 seconds.
- **Default:** 300 seconds.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D10. The range for the `idle-threshold` option and the range and default value for the `idle-time` option revised in Junos OS Release 12.3X48-D20.

Support for ADVPN with the iked process is added in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Understanding Auto Discovery VPN

af-mode

IN THIS SECTION

- Syntax | 2039
- Hierarchy Level | 2039
- Description | 2039
- Default | 2039
- Required Privilege Level | 2040
- Release Information | 2040

Syntax

```
af-mode;
```

Hierarchy Level

```
[edit poe interface (all | interface-name)]
```

Description

Configure a PoE port on an EX6200 switch to support IEEE 802.3af only. The maximum power the port can deliver in either class or static mode is 15.4 W.

Default

PoE ports on an EX6200 switch support IEEE 802.3at (PoE+) by default.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring PoE Interfaces on EX Series Switches

af-name

IN THIS SECTION

- Syntax | 2040
- Hierarchy Level | 2041
- Description | 2041
- Options | 2041
- Required Privilege Level | 2041
- Release Information | 2041

Syntax

```
af-name {
    peer-gnf {
        id peer-gnf-id;
        remote-af-name;
    }
    description af-description;
}
```

Hierarchy Level

```
[edit chassis network-slices guest-network-functions gnf id]
```

Description

Configure abstracted fabric (af) interface between a pair of guest network functions (GNFs). Abstracted fabric interface is a pseudo interface that represents a first class Ethernet interface behavior. An af interface is created on a GNF to communicate with the peer GNF when the two GNFs are connected to each other through the CLI.

NOTE: Only one af interface can be configured between a pair of GNFs.

Options

af-name	Name of the af interface being created on the GNF. • **Range:** af0 through af9
id *peer-gnf -id*	Name of the GNF peer connected using the af interface. • **Range:** 1 through 10
remote-af-name	Name of the af interface on the peer GNF.
description *af-description*	A description for the af interface.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Setting Up Junos Node Slicing

age

IN THIS SECTION

- Syntax | 2042
- Hierarchy Level | 2042
- Description | 2042
- Default | 2042
- Required Privilege Level | 2042
- Release Information | 2043

Syntax

```
age (30m | 10m | 1m | 30s | 10s);
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management linktrace]
```

Description

Time to wait (in minutes or seconds) for a response. If no response is received, the request and response entry is deleted from the linktrace database.

Default

10 minutes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Linktrace Protocol in CFM

age-of-attack

IN THIS SECTION

- Syntax | 2043
- Hierarchy Level | 2043
- Description | 2044
- Options | 2044
- Required Privilege Level | 2044
- Release Information | 2044

Syntax

```
age-of-attack
    {
    greater-than value;
    less-than value;
}
```

Hierarchy Level

```
[edit security idp  dynamic-attack-group              name              filters]
```

Description

Age of an Attack.

Options

greater-than *value*

Match when age of attack in terms of years is greater than the value (years) specified.

> **NOTE**: The first attack was added in the year 2003. So, configuring age greater than 18 will not result in any attacks.

- **Range:** 1 year through 100 years

less-than *value*

Match when age of attack in terms of years is less than the value (years) specified.

- **Range:** 1 year through 100 years

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 18.2R1.

agent-address

IN THIS SECTION

- Syntax | 2045
- Hierarchy Level | 2045
- Description | 2045
- Options | 2045
- Required Privilege Level | 2045
- Release Information | 2045

Syntax

```
agent-address outgoing-interface;
```

Hierarchy Level

```
[edit snmp trap-options]
```

Description

Set the agent address of all SNMPv1 traps generated by this router or switch. Currently, the only option is **outgoing-interface**, which sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.

Options

outgoing-interface—Value of the agent address of all SNMPv1 traps generated by this router or switch. The **outgoing-interface** option sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.

- **Default:** disabled (the agent address is not specified in SNMPv1 traps).

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Agent Address for SNMP Traps

agent (Analytics)

IN THIS SECTION

- Syntax | 2046
- Hierarchy Level | 2047
- Description | 2048
- Required Privilege Level | 2048
- Release Information | 2048

Syntax

```
agent {
    service-agents {
        agent-name {
            inputs {
                analytics {
                    parameters {
                        generate-tags value;
                        sample-frequency value;
                        sensors file-path;
                    }
                }
                input-ipfix {
                    parameters {
                        maximum-connections number;
                        tcp-port port-number;
                        vrf-name name;
                    }
                }
                input-jti-ipfix {
                    parameters {
                        record-group group-name {
                            record ipfix-record-name;
                            reporting-interval seconds;
                        }
                    }
```

```
                }
            }
            outputs {
                file {
                    parameters {
                        path file-path;
                    }
                }
                kafka {
                    parameters {
                        server ip-address;
                        topic topic-name;
                        encoding encoding-type;
                    }
                }
                output-ipfix {
                    parameters {
                        collector-address ip-address;
                        collector-ca-certificate file-path;
                        collector-certificate file-path;
                        collector-certificate-key file-path;
                        collector-connection-retry-interval seconds;
                        collector-tcp-port port-number;
                        collector-vrf-name vrf-name;
                    }
                }
            }
        }
    }
    traceoptions {
        filename filename;
        flag (debug | error | info | trace);
    }
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Configure the Network Telemetry Framework (NTF) agent and corresponding service agents that use input and output plug-ins to collect, transform, and forward network telemetry data.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.3R1.

Statement is not supported on Junos OS Evolved.

RELATED DOCUMENTATION

- *Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data*
- *Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator*
- Configuring NTF Agent
- *IPFIX Mediation on the BNG*
- *Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector*

agent-circuit-identifier (Dynamic ACI VLANs)

IN THIS SECTION

- Syntax | 2049
- Hierarchy Level | 2049
- Description | 2049
- Required Privilege Level | 2049
- Release Information | 2049

Syntax

```
agent-circuit-identifier {
    dynamic-profile profile-name;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces “$junos-interface-ifd-name” unit “$junos-
interface-unit” auto-configure],
[edit interfaces interface-name unit logical-unit-number auto-configure]
```

Description

Configure a static or dynamic underlying VLAN interface to enable dynamic VLAN subscriber interface creation based on agent circuit identifier information.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

> *Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*
>
> *Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*

agent-id

IN THIS SECTION

- Syntax (EX Series) | 2050
- Syntax (MX Series, QFX Series, OCX Series, and PTX Series) | 2050
- Hierarchy Level | 2050
- Description | 2050
- Required Privilege Level | 2051
- Release Information | 2051

Syntax (EX Series)

```
agent-id ip-address;
```

Syntax (MX Series, QFX Series, OCX Series, and PTX Series)

```
agent-id {
    ipv4-address;
    inet6 ipv6-address;
}
```

Hierarchy Level

```
[edit protocols sflow]
```

Description

Configure the IP address to be assigned as the agent ID for the sFlow agent. By assigning an IP address, you ensure that the IP address is not dynamic. You can specify an IPv4 or IPv6 address or both.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Support for IPv6 addresses introduced in Junos OS Release 19.1R1 for MX Series, PTX Series, and QFX Series devices.

RELATED DOCUMENTATION

Configuring sFlow Technology for Network Monitoring (CLI Procedure)

agent-specifier

IN THIS SECTION

- Syntax | 2051
- Hierarchy Level | 2052
- Description | 2052
- Default | 2052
- Options | 2052
- Required Privilege Level | 2053
- Release Information | 2053

Syntax

```
agent-specifier {
    aci circuit-id-string ari remote-id-string {
        drop;
        delay seconds;
```

```
        terminate;
        dynamic-profile profile-name;
        routing-instance routing-instance-name;
        static-interface interface-name;
    }
}
```

Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name]
```

Description

Specify the action taken by the interface for the specified agent circuit identifier/agent remote identifier (ACI/ARI) pair when the interface receives a PPPoE Active Discovery Initiation (PADI) control packet that includes the vendor-specific tag with ACI/ARI pair information. You can configure an ACI/ARI pair for a named service, `empty` service, or `any` service in a PPPoE service name table. A maximum of 8000 ACI/ARI pairs are supported per PPPoE service name table. You can distribute the ACI/ARI pairs in any combination among the named, `empty`, and `any` service entries in the service name table.

You can use an asterisk (*) as a wildcard character to match ACI/ARI pairs, the ACI alone, or the ARI alone. The asterisk can be placed only at the beginning, the end, or both the beginning and end of the identifier string. You can also specify an asterisk alone for either the ACI or the ARI. You cannot specify only an asterisk for both the ACI and the ARI. When you specify a single asterisk as the identifier, that identifier is ignored in the PADI packet.

For example, suppose you care about matching only the ACI and do not care what value the ARI has in the PADI packet, or even whether the packet contains an ARI value. In this case you can set the *remote-id-string* to a single asterisk. Then the interface ignores the ARI received in the packet and the interface takes action based only on matching the specified ACI.

Default

The default action is terminate.

Options

`aci` *circuit-id-string*—Identifier for the agent circuit ID that corresponds to the DSLAM interface that initiated the service request. This is a string of up to 63 characters.

`ari` *`remote-id-string`*—Identifier for the subscriber associated with the DSLAM interface that initiated the service request. This is a string of up to 63 characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

`drop`, `delay`, `terminate`, `dynamic-profile`, `routing-instance`, and `static-interface` options introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information

agf

IN THIS SECTION

- Syntax | 2054
- Hierarchy Level | 2055
- Description | 2055
- Options | 2055
- Required Privilege Level | 2055
- Release Information | 2055

Syntax

```
agf {
 amfname {
 default-amf;
 ip-address ip-address;
 local-endpoint {
  initial-tnla-weight-factor initial-tnla-weight-factor;
  ip-address name {
   primary;
  }
 }
 node-id node-id;
 offline;
 port port;
 routing-instance routing-instance;
 tracking-area name;
}


node-id node-id;
node-name node-name;
plmn name {
 mcc mcc;
 mnc mnc;
}
tracking-area code {
 plmn name {
  s-nssai name {
   sd sd;
   sst(embb | miot | urllc | v2x);
  }
 }
}
user-planesname {
 colocated-user-plane {
  ip-endpoint-address ip-endpoint-address;
 }
 data-network-name data-network-name;
 ip-address ip-address;
 routing-instance routing-instance;
```

```
}
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the Access Gateway Function.

Options

node-id Specifies the AGF node identifier.

- **Range:** 0 through 255

node-name AGF node name

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3R1.

aggregate-clients (DHCP Local Server)

IN THIS SECTION

- Syntax | 2056
- Hierarchy Level | 2056
- Description | 2057
- Options | 2057
- Required Privilege Level | 2057
- Release Information | 2057

Syntax

```
aggregate-clients (merge | replace);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name system services dhcp-local-server dynamic-profile
profile-name],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
dynamic-profile profile-name],
[edit routing-instances routing-instance-name system services dhcp-local-server dynamic-profile
profile-name],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
dynamic-profile profile-name],
[edit system services dhcp-local-server dynamic-profile profile-name],
[edit system services dhcp-local-server group group-name dynamic-profile profile-name]
```

Description

Specify that the router merge (chain) client attributes such as firewall filters and CoS attributes or replace them when multiple client sessions exist on the same underlying VLAN.

Not supported for IP demux subscriber interfaces.

Options

`merge`—Aggregate multiple clients attributes for the same subscriber (logical interface)

`replace`—Replace the entire logical interface whenever a new client logs in to the network using the same VLAN logical interface

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Options `merge` and `replace` introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

aggregate-clients (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2058
- Hierarchy Level | 2058
- Description | 2059

- Options | 2059
- Required Privilege Level | 2059
- Release Information | 2059

Syntax

```
aggregate-clients (merge | replace);
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dhcpv6 group  group-name dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name dynamic-profile profile-
name],
[edit forwarding-options dhcp-relay group  group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 dynamic-profile
profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dynamic-profile profile-
name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
dynamic-profile profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name dynamic-
profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name dynamic-profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 dynamic-
profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dynamic-profile
profile-name],
```

```
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group  group-
name dynamic-profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group  group-name
dynamic-profile profile-name]
```

Description

Specify that the router merge (chain) client attributes such as firewall filters and CoS attributes or replace them when multiple client sessions exist on the same underlying VLAN. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

Not supported for IP demux subscriber interfaces.

Options

`merge`—Aggregate multiple client attributes for the same subscriber (logical interface)

`replace`—Replace the entire logical interface whenever a new client logs in to the network using the same VLAN logical interface

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Options `merge` and `replace` introduced in Junos OS Release 9.5.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

dhcp-relay | 5620

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

Configuring Group-Specific DHCP Relay Options

aggregate-clients (Static Subscribers)

IN THIS SECTION

- Syntax | 2060
- Hierarchy Level | 2060
- Description | 2061
- Default | 2061
- Options | 2061
- Required Privilege Level | 2061
- Release Information | 2061

Syntax

```
aggregate-clients (merge | replace);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name system services static-subscribers dynamic-profile
profile-name],
[edit logical-systems logical-system-name system services static-subscribers group group-name
dynamic-profile profile-name],
[edit routing-instances routing-instances-name system services static-subscribers dynamic-profile
profile-name],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name dynamic-profile profile-name],
```

```
[edit system services static-subscribers dynamic-profile profile-name],
[edit system services static-subscribers group group-name dynamic-profile profile-name]
```

Description

Specify for all static subscribers or for a group of static subscribers that the router merge (chain) subscriber (client) attributes such as firewall filters and CoS attributes or replace them when multiple subscriber sessions exist on the same underlying VLAN. The group version of this statement overrides the global version.

NOTE: This statement is not supported for IP demux subscriber interfaces.

NOTE: This statement is not supported for enhanced subscriber management.

Default

By default, multiple subscribers cannot be on the same logical interface.

Options

`merge`—Aggregate the attributes of multiple subscribers for the logical interface.

`replace`—Replace the entire logical interface whenever a new client logs in to the network using the same VLAN logical interface.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers

aggregate-export-interval

IN THIS SECTION

- Syntax | 2062
- Hierarchy Level | 2062
- Description | 2062
- Options | 2062
- Required Privilege Level | 2063
- Release Information | 2063

Syntax

```
aggregate-export-interval seconds;
```

Hierarchy Level

```
[edit forwarding-options accounting name output],
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output],
[edit forwarding-options sampling family (inet |inet6 |mpls) output]
```

Description

Specify the duration, in seconds, of the interval for exporting aggregate accounting information.

Options

seconds—Duration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Discard Accounting

aggregate (Gigabit Ethernet CoS Policer)

IN THIS SECTION

- Syntax | 2063
- Hierarchy Level | 2064
- Description | 2064
- Required Privilege Level | 2064
- Release Information | 2064

Syntax

```
aggregate {
    bandwidth-limit bps;
    burst-size-limit bytes;
}
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile
policer cos-policer-name]
```

Description

Define a policer to apply to nonpremium traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

premium (Hierarchical Policer) | **13110**

ieee802.1p

aggregate (Hierarchical Policer)

IN THIS SECTION

- Syntax | **2065**
- Hierarchy Level | **2065**
- Description | **2065**

Required Privilege Level | 2065
Release Information | 2066

Syntax

```
aggregate {
    if-exceeding {
        bandwidth-limit bandwidth;
        burst-size-limit burst;
    }
    then {
        discard;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer name],
[edit firewall hierarchical-policer]
```

Description

On M40e, M120, and M320 edge routers with Flexible PIC Concentrator (FPC) input as FFPC and FPC output as SFPC, and on MX Series, T320, T640, and T1600 edge routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure an aggregate hierarchical policer.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles ... hierarchical-policer name]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Hierarchical Policer Configuration Overview
- Hierarchical Policers
- bandwidth-limit (Hierarchical Policer)
- burst-size-limit (Hierarchical Policer) | **3084**
- hierarchical-policer (aggregate or premium) | **8057**
- if-exceeding (Hierarchical Policer) | **8371**
- premium (Hierarchical Policer) | **13110**

aggregate-label

IN THIS SECTION

- Syntax | **2067**
- Hierarchy Level | **2067**
- Description | **2067**
- Options | **2067**
- Required Privilege Level | **2067**
- Release Information | **2067**

Syntax

```
aggregate-label {
    community community-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family inet labeled-unicast],
[edit logical-systems logical-system-name protocols bgp family inet6 labeled-unicast],
[edit logical-systems logical-system-name protocols bgp family inet-vpn unicast],
[edit logical-systems logical-system-name protocols bgp family inet-vpn6 unicast],
[edit protocols bgp family inet labeled-unicast],
[edit protocols bgp family inet6 labeled-unicast],
[edit protocols bgp family inet-vpn unicast],
[edit protocols bgp family inet6-vpn unicast]
```

Description

Specify matching criteria (in the form of a community) such that all routes which match are assigned the same VPN label, selected from one of the several routes in the set defined by this criteria. This reduces the number of VPN labels that the router must consider, and aggregates the received labels.

Options

`community` *`community-name`*—Specify the name of the community to which to apply the aggregate label.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Aggregate Labels for VPNs

aggregate (Protocols RSVP)

IN THIS SECTION

- Syntax | 2068
- Hierarchy Level | 2068
- Description | 2068
- Default | 2069
- Required Privilege Level | 2069
- Release Information | 2069

Syntax

```
(aggregate | no-aggregate);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp peer-interface peer-interface-name],
[edit protocols rsvp peer-interface peer-interface-name]
```

Description

Control the use of RSVP aggregate messages on an interface or peer interface, as described below.

NOTE: Starting in Junos OS Release 15.2, the `aggregate` statement is deprecated at the `[edit protocols rsvp interface interface-name]` and `[edit logical-systems logical-system-name`

protocols rsvp interface *interface-name*] hierarchy levels, as RSVP message aggregation is enabled by default.

- aggregate—Use RSVP aggregate messages.
- no-aggregate—Do not use RSVP aggregate messages.

  Aggregate messages can pack multiple RSVP messages into a single transmission, thereby reducing network overhead and enhancing efficiency. The number of supportable sessions and processing overhead are significantly improved when aggregation is enabled.

  Not all routers connected to a subnet need to support aggregation simultaneously. Each RSVP router negotiates its intention to use aggregate messages on a per-neighbor basis. Only when both routers agree are aggregate messages sent.

To have refresh reduction and reliable delivery, you must include the aggregate and reliable statements.

Default

Aggregation is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

RSVP Refresh Reduction

Configuring RSVP Refresh Reduction

reliable | **14304**

aggregate (Routing)

IN THIS SECTION

- Syntax | 2070
- Hierarchy Level | 2070
- Description | 2071
- Options | 2071
- Required Privilege Level | 2072
- Release Information | 2072

Syntax

```
aggregate {
    defaults {
        ... aggregate-options ...
    }
    route destination-prefix {
        policy policy-name;
        ... aggregate-options ...
            algorithm number
                description description;
            }
         }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name],
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-options rib routing-table-name],
```

```
[edit routing-instances routing-instance-name routing-options],
[edit routing-instances routing-instance-name routing-options rib routing-table-name],
[edit routing-options],
[edit routing-options rib routing-table-name]
```

Description

Configure aggregate routes.

Options

aggregate-options—Additional information about aggregate routes that is included with the route when it is installed in the routing table. Specify zero or more of the following options in ***aggregate-options***. Each option is explained separately.

- ***active***— Removes inactive routes from the forwarding table.
- ***passive***— Retains inactive routes in the forwarding table.
- **as-path <*as-path*> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator *as-number ip-address*>**;

- **(brief | full)**;
- **community [*community-ids*]**;
- **discard**;
- **(metric | metric2 | metric3 | metric4) *value* <type *type*>**;
- **(preference | preference2 | color | color2) *preference* <type *type*>**;
- **tag *metric* type *number***;

algorithm *number*—Specify an IGP algorithm for the aggregate route. IS-IS can advertise aggregate routes as SRv6 locators for locator summarization, In addition to the IPv6 aggregate prefix advertisement, IS-IS advertises the summary locator only if the node participates in the specified flexible algorithm.

- **Range**—Specify an algorithm from 0 through 255.

> **NOTE:**
> *IGP Algorithm Types*

defaults—Specify global aggregate route options. These options only set default attributes inherited by all newly created aggregate routes. These are treated as global defaults and apply to all the aggregate routes you configure in the `aggregate` statement. This part of the `aggregate` statement is optional.

route *destination-prefix*—Configure a nondefault aggregate route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/ prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.
- **next-table *next-table***—Specify the name of the next routing table to the destination. Forwarding for the aggregate prefix is done using this table. When you configure **next-table** to **next-hop**, route lookup is redirected to the route table that the next hop points to.

description *description*—Configure a description for the aggregate aggregate route. The configured description will be stored in backend data-structures and is used as an aggregate description in telemetry streaming.

The `policy` statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation

How to Enable SRv6 Network Programming in IS-IS Networks

aggregate-wait-time

IN THIS SECTION

- Syntax | 2073
- Hierarchy Level | 2073
- Description | 2073
- Default | 2073
- Required Privilege Level | 2074
- Release Information | 2074

Syntax

```
aggregate-wait-time seconds;
```

Hierarchy Level

```
[edit interfaces name aggregated-ether-options lacp]
```

Description

Set the aggregate wait time in seconds for an AE interface. This wait time delays the up event to allow for the interface state to be programmed before it is brought online. The value can be configured within a range from 2 to 1800.

> **NOTE**: Starting from Junos OS Release 22.2R1, the range of the `aggregate-wait-time` statement at the [`edit interfaces name aggregated-ether-options lacp`] hierarchy level has changed from 2 to 1800 seconds to 0 to 1800 seconds.

Default

2

Required Privilege Level

interface—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

aggregated-devices

IN THIS SECTION

- Syntax (EX Series) | 2074
- Syntax (QFX Series, EX4600, OCX1100, NFX Series) | 2074
- Hierarchy Level (EX Series, QFX Series, SRX Series) | 2075
- Description | 2075
- Default | 2076
- Required Privilege Level | 2076
- Release Information | 2076

Syntax (EX Series)

```
aggregated-devices {
    ethernet (Aggregated Devices) {
        device-count number;
        lacp
    }
}
```

Syntax (QFX Series, EX4600, OCX1100, NFX Series)

```
aggregated-devices {
    ethernet {
```

```
        device-count number;
    }
}
```

Hierarchy Level (EX Series, QFX Series, SRX Series)

```
[edit chassis]
```

```
aggregated-devices {
    ethernet {
        device-count number;
        lacp {
            link-protection {
                non-revertive;
            }
            system-priority;
        }
    }
    sonet {
        device-count number;
        }
    maximum-links maximum-links-limit;
}
```

Description

Configure properties for aggregated devices on the switch.

Configure properties for aggregated devices on the router. Aggregate Ethernet links are logical interfaces defined on the device that bundle together multiple physical interfaces into a single interface for the use of redundancy and bandwidth aggregation. When interconnecting devices you can create aggregate ethernet interfaces to bundle together multiple physical ethernet links to increase bandwidth and redundancy between devices.

Link aggregation enables you to group Ethernet interfaces to form a single link layer interface. Link Aggregation Control Protocol (LACP) is supported in chassis cluster deployments, where aggregated Ethernet interfaces and redundant Ethernet interfaces are supported simultaneously.

You must first configure the system to enable configuring the Aggregated Ethernet (ae) Interfaces. By default, Juniper devices do not have any aggregated ethernet interfaces created. To configure the device to support a given number of ae interfaces, you must define it on a per chassis basis using the `set chassis aggregated-devices devices {1-32}` in configuration mode. The number of devices you define will be the number of aggregated ethernet interfaces that the system will create which can be configured just like any other ethernet interface. Also you can view the interfaces created by using the `show interface terse` command. Once you have defined the number of aggregated ethernet devices on the chassis you can then continue to configure the LAG members on a per ethernet interface basis.

The remaining statements are explained separately. See CLI Explorer.

Default

Aggregated devices are disabled.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced for SRX devices in Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0.

Support for LACP link protection and system priority introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*
- *Configuring Link Aggregation*
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*
- Junos OS Network Interfaces Library for Routing Devices
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- Configuring Aggregated Ethernet Links (CLI Procedure)
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*
- Junos OS Ethernet Interfaces Configuration Guide
- Configuring Junos OS for Supporting Aggregated Devices

aggregated-ether

IN THIS SECTION

- Syntax | 2077
- Hierarchy Level | 2077
- Description | 2077
- Options | 2077
- Required Privilege Level | 2077
- Release Information | 2078

Syntax

```
aggregated ether aggregated-ether-group-name;
```

Hierarchy Level

```
[edit chassis synchronization source interfaces ]
```

Description

Specify an aggregated Ethernet group while configuring a member of the group as the clock source. This is an optional setting. Any number of interfaces from an aggregated Ethernet group can be configured as the clock source. If a member from the aggregated Ethernet group is selected as the locked clock source, then all the members of the group will transmit the Do not Use (DNU) ESMC quality level.

Options

aggregated-ether-group-name	aggregated ethernet group name

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Synchronous Ethernet Overview

show chassis synchronization | 18879

aggregated-ether-options

IN THIS SECTION

- Syntax | 2078
- Syntax (NFX, QFX Series, EX4600, OCX1100, QFabric) | 2080
- Hierarchy Level | 2081
- Description | 2081
- Required Privilege Level | 2081
- Release Information | 2082

Syntax

```
aggregated-ether-options {
    ethernet-switch-profile {
        ethernet-policer-profile {
            input-priority-map {
                ieee802.1p premium [ values ];
            }
            output-priority-map {
                classifier {
                    premium {
```

```
                        forwarding-class class-name {
                            loss-priority (high | low);
                        }
                    }
                }
            }
            policer cos-policer-name {
                aggregate {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
                premium {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
            }
        }
        (mac-learn-enable | no-mac-learn-enable);
    }
    (flow-control | no-flow-control);
    lacp {
        (active | passive);
        link-protection {
            disable;
        (revertive | non-revertive);
        periodic interval;
        sync-reset
        system-priority priority;
        system-id system-id;
    }
    load-balance {
        local-bias;
        no-adaptive;
        per-packet;
    }
    local-bias;
    link-speed speed;
    logical-interface-chassis-redundancy;
    logical-interface-fpc-redundancy;
    (loopback | no-loopback);
    minimum-links number;
    rebalance-periodic time hour:minute <interval hours>;
    source-address-filter {
```

```
        mac-address;
        (source-filtering | no-source-filtering);
    }
}
```

Syntax (NFX, QFX Series, EX4600, OCX1100, QFabric)

The `fcoe-lag` and `mc-ae` statements are not supported on OCX Series switches.

```
aggregated-ether-options {
    configured-flow-control {
        rx-buffers (on | off);
        tx-buffers (on | off);
    }
    ethernet-switch-profile {
    tag-protocol-id;
    (fcoe-lag | no-fcoe-lag);
    (flow-control | no-flow-control);
    lacp mode {
        admin-key key;
        periodic interval;
        system-id mac-address;
        force-up;
        }
    }
    (link-protection | no-link-protection);
    link-speed speed;
    local-bias;
    local-minimum-links-threshold threshold-value;
     (loopback | no-loopback);
    mc-ae {
        chassis-id chassis-id;
        enhanced-convergence;
        events {
            iccp-peer-down {
                force-icl-down;
                prefer-status-control-active;
            }
        }
        init-delay-time seconds;
        mc-ae-id mc-ae-id;
```

```
        mode (active-active);
        recovery-delay-time seconds;
        redundancy-group group-id;
        revert-time revert-time;
        status-control (active | standby);
        switchover-mode (non-revertive | revertive);
    }
    minimum-links number;
    rebalance-periodic;
    resilient-hash;
    source-address-filter filter;
    (source-filtering | no-source-filtering);
}
```

Hierarchy Level

```
[edit interfaces aex]
```

Description

Configure aggregated Ethernet-specific interface properties.

NOTE:

- The `fcoe-lag` and `mc-ae` statements are not supported on OCX Series switches.
- The `force-up` statement is not supported on QFX10002 switches.
- The `resilient-hash` statement is not supported on QFX5200, QFX5210, or QFX10002 switches.
- On QFX switches, the `enhanced-convergence` statement is supported on QFX10000 line of switches and `recovery-delay-time` statement is supported on QFX5000 and QFX10000 line of switches.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statements `fcoe-lag` and `no-fcoe-lag` introduced in Junos OS Release 13.2X52-D10 for the QFX Series.

Statements `force-up`, `lacp`, and `resilient-hash` introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Statement `local-minimum-links-threshold` introduced in Junos OS Release 14.1X53-D40 for the QFX Series.

RELATED DOCUMENTATION

- *Ethernet Interfaces Overview*
- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch*
- Junos OS Network Interfaces Library for Routing Devices
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- Configuring Aggregated Ethernet Links (CLI Procedure)
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*
- Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support
- Junos OS Ethernet Interfaces Configuration Guide

aggregated-inline-services-options (Aggregated Inline Services)

IN THIS SECTION

Syntax | **2083**

- Hierarchy Level | 2083
- Description | 2083
- Required Privilege Level | 2084
- Release Information | 2084

Syntax

```
aggregated-inline-services-options {
    primary-interface interface-name;
    secondary-interface interface-name;
}
```

Hierarchy Level

```
[edit interfaces asix]
```

Description

Configure the members of an aggregated inline service interface bundle to provide 1:1 stateful LNS redundancy for an LNS sessions in a tunnel group.

BEST PRACTICE: Follow these guidelines:

- You must configure `unit 0 family inet` for each bundle; otherwise, the session fails to come up.
- The primary (active) and secondary (backup) interfaces must be on different MPCs. If you configure both interfaces on the same MPC, the subsequent configuration commit fails.
- The bandwidth configured at the `[edit chassis fpc slot pic number inline-services bandwidth]` hierarchy level must be the same for both member links.

- An si interface configured as a member of an aggregated inline service bundle cannot be configured as a member of another bundle group.
- An si interface configured as a member of an aggregated inline service bundle cannot also be used for any function that is not related to aggregated services; for example, it cannot be used for inline IP reassembly.
- When you configure an si interface as a member of an aggregated inline services bundle, you can no longer configure that si interface independently. You can configure only the parent bundle; the bundle's configuration is applied immediately to all member interfaces.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

aggregation (Forwarding Options)

IN THIS SECTION

Syntax | 2085

Hierarchy Level | 2085

- Description | 2085
- Options | 2085
- Required Privilege Level | 2086
- Release Information | 2086

Syntax

```
aggregation {
    autonomous-system;
    destination-prefix;
    protocol-port;
    source-destination-prefix {
        caida-compliant;
    }
    source-prefix;
}
```

Hierarchy Level

```
[edit forwarding-options accounting output cflowd hostname],
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-
server hostname],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server hostname]
```

Description

For cflowd version 8 only, specify the type of data to be aggregated; cflowd records and sends only those flows that match the specified criteria.

Options

`autonomous-system`—Aggregate by autonomous system (AS) number.

`caida-compliant`—Record source and destination mask-length values in compliance with the Version 2.1b1 release of CAIDA's cflowd application. If this statement is not configured, the Junos OS records source

and destination mask length values in compliance with the *cflowd Configuration Guide*, dated August 30, 1999.

`destination-prefix`—Aggregate by destination prefix.

`protocol-port`—Aggregate by protocol and port number.

`source-destination-prefix`—Aggregate by source and destination prefix.

`source-prefix`—Aggregate by source prefix.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Flow Aggregation

Directing Traffic Sampling Output to a Server Running the cflowd Application

aggregation (IDS)

IN THIS SECTION

- Syntax | 2087
- Hierarchy Level | 2087
- Description | 2087
- Options | 2087
- Required Privilege Level | 2087
- Release Information | 2087

Syntax

```
aggregation {
    destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
    source-prefix prefix-value | source-prefix-ipv6 prefix-value;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Configure the IDS rule session limits for individual destination or source subnets rather than individual addresses. This applies session limits to an aggregation of all attacks from or to a subnet of the specified length.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IDS Rules on an MS-DPC

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

aggregations (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 2088
- Hierarchy Level | 2088
- Description | 2088
- Options | 2089
- Required Privilege Level | 2089
- Release Information | 2089

Syntax

```
aggregations {
    destination-prefix-ipv6-mask prefix-length;
    destination-prefix-mask prefix-length;
    source-prefix-ipv6-mask prefix-length;
    source-prefix-mask prefix-length;
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name]
```

Description

Configure intrusion detection service session limits for individual destination subnets or source subnets rather than individual addresses. This applies session limits to an aggregation of all sessions from or to an individual subnet of the specified length.

For example, if you configure a value of 24 for `destination-prefix-mask`, then sessions to 10.1.1.2 and 10.1.1.3 are counted as sessions to the 10.1.1/24 subnet.

Options

destination-prefix-ipv6-mask ***prefix-length***

Prefix length for destination IPv6 address subnets.

- **Range:** 0 through 128

destination-prefix-mask ***prefix-length***

Prefix length for destination IPv4 address subnets.

- **Range:** 0 through 32

source-prefix-ipv6-mask ***prefix-length***

Prefix length for source IPv6 address subnets.

- **Range:** 0 through 128

source-prefix-mask ***prefix-length***

Prefix length for source IPv4 address subnets.

- **Range:** 0 through 32

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

aging

IN THIS SECTION

- Syntax | **2090**
- Hierarchy Level | **2090**

- Description | 2090
- Options | 2090
- Required Privilege Level | 2090
- Release Information | 2091

Syntax

```
aging {
    early-ageout seconds;
    high-watermark percent;
    low-watermark percent;
}
```

Hierarchy Level

```
[edit security flow]
```

Description

The session aging mechanism accelerates the session time-out process when the percentage of entries in the session table exceeds the high-watermark setting, and then stops when the percentage of entries in the session table falls below the low-watermark setting.

Options

The remaining statements are explained separately. See the CLI Explorer.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

aging-timer (Junos Fusion)

IN THIS SECTION

- Syntax | 2091
- Hierarchy Level | 2091
- Description | 2091
- Default | 2092
- Options | 2092
- Required Privilege Level | 2092
- Release Information | 2092

Syntax

```
aging-timer aging-timer;
```

Hierarchy Level

```
[edit chassis satellite-management]
```

Description

Configure the aging timer on the aggregation device in a Junos Fusion.

The aging timer is used on the aggregation device to specify the amount of time, in minutes, to maintain the device state of an unreachable satellite device before deleting the satellite device from the Junos Fusion.

If the unreachable satellite device is discovered before the aging timer expires, the satellite device is reactivated in the Junos Fusion without having to restore its device state.

Default

The default aging time is 10 minutes.

Options

The remaining statements are explained separately.

- **Range:** 2 through 60,000 minutes

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

aigp

IN THIS SECTION

Syntax | 2093

- Hierarchy Level | 2093
- Description | 2094
- Options | 2094
- Required Privilege Level | 2094
- Release Information | 2095

Syntax

```
aigp [disable];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family inet labeled-unicast],
[edit logical-systems logical-system-name protocols bgp family inet6 labeled-unicast],
[edit logical-systems logical-system-name protocols bgp group group-name family inet labeled-
unicast],
[edit logical-systems logical-system-name protocols bgp group group-name family inet6 labeled-
unicast],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
inet labeled-unicast],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
inet6 labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family inet labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family inet6 labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family inet labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family inet6 labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family inet labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family inet6 labeled-unicast],
[edit protocols bgp family inet labeled-unicast],
[edit protocols bgp family inet6 labeled-unicast],
```

```
[edit protocols bgp group group-name family inet labeled-unicast],
[edit protocols bgp group group-name family inet6 labeled-unicast],
[edit protocols bgp group group-name neighbor address family inet labeled-unicast],
[edit protocols bgp group group-name neighbor address family inet6 labeled-unicast],
[edit routing-instances routing-instance-name protocols bgp family inet labeled-unicast],
[edit routing-instances routing-instance-name protocols bgp family inet6 labeled-unicast],
[edit routing-instances routing-instance-name protocols bgp group group-name family inet labeled-
unicast],
[edit routing-instances routing-instance-name protocols bgp group group-name family inet6
labeled-unicast],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family inet labeled-unicast],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family inet6 labeled-unicast]
```

Description

Enable the accumulated interior gateway protocol (AIGP) BGP attribute on a protocol family. Configuring AIGP on a particular family enables sending and receiving of the AIGP attribute on that family.

The AIGP attribute enables deployments in which a single administration can run several contiguous BGP autonomous systems (ASs). Such deployments allow BGP to make routing decisions based on the IGP metric. With AIGP enabled, BGP can select paths based on IGP metrics. This enables BGP to choose the shortest path between two nodes, even though the nodes might be in different ASs. The AIGP attribute is particularly useful in networks that use tunneling to deliver a packet to its BGP next hop. Such is the case with MPLS label-switched paths.

Options

disable—Explicitly disables AIGP.

- **Default:** Disabled, meaning that the device does not send an AIGP attribute and silently discards a received AIGP attribute.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring the Accumulated IGP Attribute for BGP

aigp-originate | 2097

aigp-adjust (Policy Action)

IN THIS SECTION

- Syntax | 2095
- Hierarchy Level | 2095
- Description | 2096
- Options | 2096
- Required Privilege Level | 2096
- Release Information | 2096

Syntax

```
aigp-adjust (add|divide|multiply|subtract) (user-value | distance-to-protocol-nexthop);
```

Hierarchy Level

```
[edit policy-options policy-statement policy-name term term-name then],
[edit policy-options policy-statement policy-name then]
```

Description

Specify this CLI policy action in an import or export policy to modify the accumulated interior gateway protocol (AIGP) BGP attribute. You can modify the AIGP attribute for one of the following reasons:

- Modify the AIGP metric on the route if it exists on the external BGP sessions between two ASBR routers where there is no interior gateway protocol (IGP)
- Scale up or down while transitioning from one IGP domain, such as OSPF to another IGP domain, such as IS-IS
- Make very minor adjustments on the AIGP from another AS domain or another vendor's routers

> **CAUTION**: AIGP is rated very high in the best route decision and comes only after the BGP local preference rule. Therefore, use the AIGP adjustment option with caution as it can have a huge impact on the network.

Options

add | divide | multiply | subtract	Specify the mathematical operation that needs to be performed on the original AIGP attribute for adjustment.
user-value	Specify an unsigned 64-bit value in decimal.
distance-to-protocol-nexthop	Use the current metric2 value in the routing table as specified in the routing policy. Configure this option in the export policy if you want full control of the advertised AIGP value and do not want BGP to add up the distance to the protocol nexthop.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

metric (Policy Action) | 10681

show policy | 23199

aigp-originate

IN THIS SECTION

- Syntax | 2097
- Hierarchy Level | 2097
- Description | 2098
- Default | 2098
- Options | 2098
- Required Privilege Level | 2098
- Release Information | 2098

Syntax

```
aigp-originate distance;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options policy-statement policy-name term term-
name then],
[edit logical-systems logical-system-name policy-options policy-statement policy-name
then],
[edit policy-options policy-statement policy-name term term-name then],
[edit policy-options policy-statement policy-name then]
```

Description

Originate an accumulated interior gateway protocol (AIGP) BGP attribute for a given prefix by export policy, using the `aigp-originate` policy action.

To originate an AIGP attribure, you need configure the policy action on only one node. The AIGP attribute is readvertised if the neighbors are AIGP enabled with the `aigp` statement in the BGP configuration.

Default

If you omit the `aigp-originate` policy action, the node still readvertises the AIGP BGP attribute if AIGP is enabled in the BGP configuration. However, the node does not originate its own AIGP attribute for local prefixes.

As the route is readvertised by downstream nodes, the cost of the AIGP attribute reflects the IGP distance to the prefix (zero + IGP distance or configured distance + IGP distance).

Options

`distance`—(Optional) Associate an initial cost when advertising a local prefix with the AIGP BGP attribute.

- **Range:** 0 through 4,294,967,295
- **Default:** The initial cost associated with the AIGP attribute for a local prefix is zero. The *`distance`* option overrides the default zero value for the initial cost.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring the Accumulated IGP Attribute for BGP

aigp

ais-trigger-condition

IN THIS SECTION

- Syntax | 2099
- Hierarchy Level | 2099
- Description | 2099
- Options | 2099
- Required Privilege Level | 2100
- Release Information | 2100

Syntax

```
ais-trigger-condition {
    adjacency-loss;
    all-defects;
    cross-connect-ccm;
    erroneous-ccm;
    receive-ais;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
event]
```

Description

Configure the defect conditions that generate an alarm indication signal (AIS).

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Alarm Indication Signal (ETH-AIS) Function Overview

Configuring ETH-AIS on a CFM MEP

alarm (chassis)

IN THIS SECTION

- Syntax | 2100
- Hierarchy Level | 2101
- Description | 2101
- Options | 2101
- Required Privilege Level | 2101
- Release Information | 2101

Syntax

```
alarm {
    interface-type {
        alarm-name (ignore |red | yellow);
    }
}
```

Hierarchy Level

```
[edit chassis],
[edit chassis interconnect-device name],
[edit chassis node-group name]
```

Description

Configure the chassis alarms and whether they trigger a red or yellow alarm, or whether they are ignored. Red alarm conditions light the `RED ALARM` LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the contact on the craft interface or LCD screen. Yellow alarm conditions light the `YELLOW ALARM` LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the craft interface or LCD screen.

To configure more than one alarm, include multiple *`alarm-name`* lines.

Options

`alarm-name`—Alarm condition. For example, `input1`, `input2`

`ignore`—The specified alarm condition does not set off any alarm.

`interface-type`—Type of interface on which you are configuring the alarm: `psm`, `fpc`, `atm`, `ethernet`, `sonet`, or `t3`.

`red`—The specified alarm condition sets off a red alarm.

`yellow`—The specified alarm condition sets off a yellow alarm.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Alarm Types and Severity Levels

Chassis Conditions That Trigger Alarms

Chassis Alarm Messages on a QFX3500 Device

Interface Alarm Messages

alarm-id

IN THIS SECTION

- Syntax | 2102
- Hierarchy Level | 2102
- Description | 2103
- Required Privilege Level | 2103
- Release Information | 2103

Syntax

```
alarm-id id {
    alarm-state state {
        description alarm-description;
        notification-id notification-id-of-alarm;
        resource-prefix alarm-resource-prefix;
        varbind-index varbind-index-in-alarm-varbind-list;
        varbind-subtree alarm-varbind-subtree;
        varbind-value alarm-varbind-value;
    }
}
```

Hierarchy Level

```
[edit snmp alarm-management alarm-list-name]
```

Description

Specify the identifier of the alarm that you need to configure.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

alarm-list-name | 2103
alarm-management | 2106
alarm-state | 2117

alarm-list-name

IN THIS SECTION

- Syntax | 2104
- Hierarchy Level | 2104
- Description | 2104
- Required Privilege Level | 2104
- Release Information | 2104

Syntax

```
alarm-list-name list-name {
    alarm-id  id {
        alarm-state state {
            description alarm-description;
            notification-id notification-id-of-alarm;
            resource-prefix alarm-resource-prefix;
            varbind-index varbind-index-in-alarm-varbind-list;
            varbind-subtree alarm-varbind-subtree;
            varbind-value alarm-varbind-value;
        }
    }
}
```

Hierarchy Level

```
[edit snmp alarm-management]
```

Description

Specify the name of the alarm list that you need to configure.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

alarm-id | 2102

alarm-management | 2106
alarm-state | 2117

alarm low-light-alarm

IN THIS SECTION

- Syntax | 2105
- Hierarchy Level | 2105
- Description | 2105
- Options | 2105
- Required Privilege Level | 2106
- Release Information | 2106

Syntax

```
alarm low-light-alarm {
    (link-down | syslog);
}
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Description

Specify the action to take if the receiving optics signal is below the optics low-light alarm threshold.

Options

`link-down`—Drop the 10-Gigabit Ethernet link and marks link as down.

`syslog`—Write the optics information to the system log.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

Statement introduced in Junos OS Release 18.2R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

RELATED DOCUMENTATION

Configuring Link Down Notification for Optics Options Alarm or Warning

100-Gigabit Ethernet OTN Options Configuration Overview

alarm-management

IN THIS SECTION

- Syntax | 2106
- Hierarchy Level | 2107
- Description | 2107
- Required Privilege Level | 2107
- Release Information | 2107

Syntax

```
alarm-management {
    alarm-list-name list-name {
```

```
        alarm-id id {
            alarm-state state {
                description alarm-description;
                notification-id notification-id-of-alarm;
                resource-prefix alarm-resource-prefix;
                varbind-index varbind-index-in-alarm-varbind-list;
                varbind-subtree alarm-varbind-subtree;
                varbind-value alarm-varbind-value;
            }
        }
    }
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Configure the alarm management system to monitor and report active alarms as well as the history of alarms through the SNMP MIB tables supported by the *Alarm MIB*.

> **NOTE**: You cannot configure alarms without notifications. It is mandatory to include the notification identifier in the configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

alarm-id | 2102

alarm-list-name | 2103

alarm-state | 2117

alarm-mode

IN THIS SECTION

- Syntax | 2108
- Hierarchy Level | 2108
- Description | 2109
- Default | 2109
- Options | 2109
- Required Privilege Level | 2109
- Release Information | 2109

Syntax

```
alarm-mode {
    mdi-records-count number;
    average;
}
```

Hierarchy Level

```
[edit services]
```

Description

If this statement is configured you can set the alarm as immediate or average mode. If immediate alarm is configured, an immediate trap is raised at the end of interval duration when the metric value exceeds the configured range. If average alarm is configured, a trap is generated based on average value for the specified number of interval duration.

Default

The default alarm mode is immediate mode.

Options

mdi-records-count *number*	Use the specified media delivery index record count number for immediate alarm mode.
average	Generate traps for average values that are not within the configured range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Inline Video Monitoring on MX Series Routers

alarms (Video Monitoring) | **2125**

alarm (optics-options)

IN THIS SECTION

- Syntax | 2110
- Hierarchy Level | 2110
- Description | 2110
- Options | 2110
- Required Privilege Level | 2111
- Release Information | 2111

Syntax

```
alarm low-light-alarm {
    (link-down | syslog);
}
```

Hierarchy Level

```
[edit interfaces interface-name optics-options],
[edit dynamic-profiles name interfaces name              optics-options],
[edit dynamic-profiles name logical-systems name interfaces name              optics-options]
```

Description

Specify the action to take if the receiving optics signal is below the optics low-light alarm threshold. Starting in Junos OS Release 15.1, for all QSFP-based interfaces, you need not explicitly configure the `syslog` option. The `syslog` option is enabled by default.

Options

`low-light-alarm`—Enable the low light alarm.

`link-down`—Drop the 10-Gigabit Ethernet link and marks link as down.

`syslog`—Write the optics information to the system log.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring Link Down Notification for Optics Options Alarm or Warning

100-Gigabit Ethernet OTN Options Configuration Overview

alarm-port

IN THIS SECTION

- Syntax | 2111
- Hierarchy Level | 2112
- Description | 2112
- Options | 2112
- Required Privilege Level | 2113
- Release Information | 2114

Syntax

```
alarm-port {
    input port port-number {
        active {
            (high | low);
```

```
        }
        admin-state {
            (disabled | enabled);
        }
        description description;
    }
        severity {
            (critical | major | minor | warning)
        }
    output port port-number {
        admin-state {
            (disabled | enabled);
        }
        description description;
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure an alarm input or output on the dedicated alarm port on the router. You can use the alarm input configuration to receive alarm inputs from the external devices such as sensors. You can use the alarm output configuration to relay the alarms in the router to external alarm devices. You can configure up to three alarm inputs and one alarm output.

Options

port-number Specify the port number for the alarm input or output. The system supports up to three alarm input ports (port number: 1 to 3). The system supports only one alarm output port (port number: 1)

active Applicable only to alarm inputs.

Configure a signal polarity for the alarm input based on the user environment. The alarm relay becomes operational when the input signal polarity is high or low based on this configuration. The following options are available.

- high
- low
- **Default:** low

admin-state Specify the administrative state of the alarms.

- disabled—Disable the alarm configuration. If you choose this option, the router will not trigger the alarm configured on this port. Disabling an alarm configuration does not delete the alarm. You can enable it when required.
- enabled—Enable the alarm configuration.
- **Default:** disabled

description *description* Provide a description for the alarm input or output. For example, you can provide a description to identify the device connected to the alarm port.

severity Applicable only to alarm inputs.

Configure an alarm severity. The system supports the following alarm severities.

- critical
- major
- minor
- warning

> **NOTE**: The system considers the alarm inputs configured with the critical or major severity as major alarms. The alarm inputs configured with the minor or warning severity are considered minor alarms.

- **Default:** critical

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Configuring Chassis Alarm Input and Output (ACX710 Routers)

alarm (Satellite Policies)

IN THIS SECTION

- Syntax | 2114
- Hierarchy Level | 2114
- Description | 2115
- Default | 2115
- Options | 2115
- Required Privilege Level | 2115
- Release Information | 2115

Syntax

```
alarm {
    linkdown [ignore | red | yellow]
}
```

Hierarchy Level

```
[edit policy-options satellite-policies environment-monitoring-policy policy-name]
```

Description

Configure the link down alarm that is sent within the Junos Fusion whenever a satellite device experiences a link-down error.

Default

Link-down alarms are not sent on satellite devices in a Junos Fusion until an environment monitoring policy is configured.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Satellite Device Alarm Handling Using an Environment Monitoring Satellite Policy in a Junos Fusion

Understanding Satellite Policies in a Junos Fusion

alarm (SNMP RMON)

IN THIS SECTION

- Syntax | 2116
- Hierarchy Level | 2116

- Description | 2116
- Options | 2116
- Required Privilege Level | 2117
- Release Information | 2117

Syntax

```
alarm index {
    description description;
    falling-event-index index;
    falling-threshold integer;
    falling-threshold-interval seconds;
    interval seconds;
    request-type (get-next-request | get-request | walk-request);
    rising-event-index index;
    rising-threshold integer;
    sample-type (absolute-value | delta-value);
    startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
    syslog-subtag syslog-subtag;
    variable oid-variable;
}
```

Hierarchy Level

```
[edit snmp rmon]
```

Description

Configure RMON alarm entries.

Options

index—Identifies this alarm entry as an integer.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an RMON Alarm Entry and Its Attributes

event | 6520

RMON MIB Event, Alarm, Log, and History Control Tables

Monitoring RMON MIB Tables

Understanding RMON

alarm-state

IN THIS SECTION

- Syntax | 2117
- Hierarchy Level | 2118
- Description | 2118
- Options | 2118
- Required Privilege Level | 2118
- Release Information | 2119

Syntax

```
alarm-state state {
    description alarm-description;
```

```
    notification-id notification-id-of-alarm;
    resource-prefix alarm-resource-prefix;
    varbind-index varbind-index-in-alarm-varbind-list;
    varbind-subtree alarm-varbind-subtree;
    varbind-value alarm-varbind-value;
}
```

Hierarchy Level

```
[edit snmp alarm-management alarm-list-name]
```

Description

Specify the state of the alarm and the other parameters that you need to monitor.

Options

description *alarm-description*	Include a brief description of the alarm.
notification-id *notification-id-of-alarm*	Specify the identifier of the notification associated with the alarm.
resource-prefix *alarm-resource-prefix*	Specify the resource prefix of the alarm.
varbind-index *varbind-index-in-alarm-varbind-list*	Specify the varbind index in the alarm varbind list. • **Range:** 0 through 4294967295
varbind-subtree *alarm-varbind-subtree*	Specify the subtree of the varbind.
varbind-value *alarm-varbind-value*	Specify the varbind value of the alarm. • **Range:** 0 through 2147483647

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

alarm-id | 2102
alarm-list-name | 2103
alarm-management | 2106

alarm-threshold

IN THIS SECTION

- Syntax | 2119
- Hierarchy Level | 2119
- Description | 2120
- Options | 2120
- Required Privilege Level | 2120
- Release Information | 2120

Syntax

```
alarm-threshold number;
```

Hierarchy Level

```
[edit security screen ids-option screen-name  tcp syn-flood]
```

Description

Define the number of half-complete proxy connections per second at which theSRX300, SRX320, SRX340, SRX345, SRX380, or SRX550M device makes entries in the event alarm log.

Options

number —Threshold value.

- **Range:** 1 through 500,000 per second
- **Default:** 1024 per second

> **NOTE**: For SRX Series Firewalls the applicable range is 1 through 1,000,000 per second.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Security Policies Overview

alarm-threshold (Security GPRS)

IN THIS SECTION

- Syntax | 2121
- Hierarchy Level | 2121
- Description | 2121

- Options | 2121
- Required Privilege Level | 2121
- Release Information | 2122

Syntax

```
alarm-threshold {
    forward number;
    reverse number;
}
```

Hierarchy Level

```
[edit security gtp profile profile-name path-rate-limit message-type (create-req | delete-req |
echo-req | other)]
```

Description

Use the alarm-threshold parameter to configure the device to raise an alarm, when the GPRS tunneling protocol (GTP) control messages on a path have reached the configured limit.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name path-rate-limit message-type (create-req | delete-req | echo-req | other)]` hierarchy-level is replaced by `[edit security gtp profile profile-name path-rate-limit message-type (create-req | delete-req | echo-req | other)]`.

Options

number—To limit messages in forward or reverse direction.

- **Range:** For create request, delete request and other GTP control messages, the applicable range is 1 through 10,000 packets per second. For echo request, the applicable range is 1 through 10,000 packets per minute.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Example: Limiting the Message Rate and Path Rate for GTP Control Messages

alarm-without-drop

IN THIS SECTION

- Syntax | 2122
- Hierarchy Level | 2122
- Description | 2123
- Required Privilege Level | 2123
- Release Information | 2123

Syntax

```
alarm-without-drop;
```

Hierarchy Level

```
[edit security screen ids-option screen-name]
```

Description

Direct the SRX device to generate an alarm when detecting an attack but not block the attack. Use this statement to allow an attack to enter a segment of your network that you have previously prepared to receive it (for example, a honeynet, which is a decoy network with extensive monitoring capabilities).

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

alarm-without-drop (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 2123
- Hierarchy Level | 2124
- Description | 2124
- Required Privilege Level | 2124
- Release Information | 2124

Syntax

```
alarm-without-drop;
```

Hierarchy Level

```
[edit services screen ids-option screen-name]
```

Description

Configure the IDS screen to log an alarm for an offending packet, but not drop the packet. The screen skips the rest of the screen checks. The packet is not counted as a dropped packet.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

alarms

IN THIS SECTION

- Syntax | 2125
- Hierarchy Level | 2125
- Description | 2125
- Required Privilege Level | 2125
- Release Information | 2125

Syntax

```
alarms;
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Description

For 10-Gigabit Ethernet DPCs, configure the DPC to drop the interface link when the receive power falls below the alarm threshold.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.1.

RELATED DOCUMENTATION

Ethernet DWDM Interface Wavelength Overview

alarms (Video Monitoring)

IN THIS SECTION

- Syntax | 2126
- Hierarchy Level | 2127

- Description | 2127
- Required Privilege Level | 2127
- Release Information | 2127

Syntax

```
alarms {
    delay-factor {
        no-syslog-generation;
        generate-snmp-traps;
        storm-control {
            count number;
            interval number;
        }
        alarm-mode {
            mdi-records-count number;
            average;
        }
    }
    media-rate-variation {
        no-syslog-generation;
        generate-snmp-traps;
        storm-control {
            count number;
            interval number;
        }
        alarm-mode {
            mdi-records-count number;
            average;
        }
    }
    media-loss-rate {
        no-syslog-generation;
        generate-snmp-traps;
        storm-control {
            count number;
            interval number;
        }
```

```
        alarm-mode {
            immediate;
        }
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the alarm to monitor and report active alarms. SNMP is used to monitor alarms raised by the inline video monitoring feature. The alarms are monitored in the network management system either to troubleshoot the problem or to diagnose degradation in video quality.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Inline Video Monitoring on MX Series Routers

delay-factor | **5248**

no-syslog-generation | **11591**

generate-snmp-traps | **7676**

storm-control (Video Monitoring) | **28619**

alarm-mode | **2108**

media-rate-variation | **10597**

media-loss-rate | **10596**

alarms (Security)

IN THIS SECTION

- Syntax | 2128
- Hierarchy Level | 2129
- Description | 2129
- Options | 2129
- Required Privilege Level | 2130
- Release Information | 2130

Syntax

```
alarms {
    audible {
        continuous;
    }
    potential-violation {
        authentication failures;
        cryptographic-self-test;
        decryption-failures {
            threshold value;
        }
        encryption-failures {
            threshold value;
        }
        idp;
        ike-phase1-failures {
            threshold value;
        }
        ike-phase2-failures {
            threshold value;
        }
        key-generation-self-test;
        non-cryptographic-self-test;
        policy {
            application {
```

```
                duration interval;
                size count;
                threshold value;
            }
            destination-ip {
                duration interval;
                size count;
                threshold value;
            }
            policy match {
                duration interval;
                size count;
                threshold value;
            }
            source-ip {
                duration interval;
                size count;
                threshold value;
            }
        }
        replay-attacks {
             threshold value;
        }
        security-log-percent-full percentage;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configures security alarms.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Security Policies Overview

alg

IN THIS SECTION

- Syntax | 2130
- Hierarchy Level | 2136
- Description | 2136
- Options | 2136
- Required Privilege Level | 2138
- Release Information | 2138

Syntax

```
alg {
    alg-manager {
        traceoptions {
            flag {
                all <extensive>;
            }
        }
```

```
    }
    alg-support-lib {
        traceoptions {
            flag {
                all <extensive>;
            }
        }
    }
    dns {
        disable;
        doctoring (none | sanity-check);
        maximum-message-length bytes;
        traceoptions {
            flag {
                all <extensive>;
            }
        }
    }
    ftp {
        allow-mismatch-ip-address;
        disable;
        ftps-extension;
        line-break-extension;
        traceoptions {
            flag {
                all <extensive>;
            }
        }
    }
    h323 {
        application-screen {
            message-flood {
                gatekeeper {
                    threshold rate;
                }
            }
            unknown-message {
                permit-nat-applied;
                permit-routed;
            }
        }
        disable;
        dscp-rewrite {
```

```
            code-point string;
        }
        endpoint-registration-timeout value-in-seconds;
        media-source-port-any;
        traceoptions {
            flag flag <detail | extensive | terse>;
        }
    }
    ike-esp-nat {
        enable;
        esp-gate-timeout value-in-seconds;
        esp-session-timeout value-in-seconds;
        state-timeout value-in-seconds;
        traceoptions {
            flag {
                all <extensive>;
            }
        }
    }
    mgcp {
        application-screen {
            connection-flood {
                threshold rate;
            }
            message-flood {
                threshold rate;
            }
            unknown-message {
                permit-nat-applied;
                permit-routed;
            }
        }
        disable;
        dscp-rewrite {
            code-point string;
        }
        inactive-media-timeout value-in-seconds;
        maximum-call-duration value-in-minutes;
        traceoptions {
            flag flag <extensive>;
        }
        transaction-timeout value-in-seconds;
    }
```

```
    msrpc {
        disable;
        map-entry-timeout;
        traceoptions {
            flag {
                all <extensive>;
            }
        }
    }
    pptp {
        disable;
        traceoptions {
            flag {
                all <extensive>;
            }
        }
    }
    rsh {
        disable;
        traceoptions {
            flag {
                all <extensive>;
            }
        }
    }
    rtsp {
        disable;
        traceoptions {
            flag {
                all <extensive>;
            }
        }
    }
    sccp {
        application-screen {
            call-flood {
                threshold rate;
            }
            unknown-message {
                permit-nat-applied;
                permit-routed;
            }
        }
```

```
        disable;
        dscp-rewrite {
            code-point string;
        }
        inactive-media-timeout value-in-seconds;
        traceoptions {
            flag flag <extensive>;
        }
    }
    sip {
        application-screen {
            protect {
                deny {
                    all {
                        timeout value-in-seconds;
                    }
                    destination-ip address;
                    timeout value-in-seconds;
                }
            }
            unknown-message {
                permit-nat-applied;
                permit-routed;
            }
        }
        c-timeout value-in-minutes;
        disable;
        dscp-rewrite {
            code-point string;
        }
        inactive-media-timeout value-in-seconds;
        maximum-call-duration value-in-minutes;
        retain-hold-resource;
        t1-interval value-in-milliseconds;
        t4-interval value-in-seconds;
        traceoptions {
            flag flag <detail | extensive | terse>;
        }
    }
    sql {
        disable;
        traceoptions {
            flag {
```

```
                all <extensive>;
            }
        }
    }
    sunrpc {
        disable;
        map-entry-timeout;
        traceoptions {
            flag {
                all <extensive>;
            }
        }
    }
    talk {
        disable;
        traceoptions {
            flag {
                all <extensive>;
            }
        }
    }
    tftp {
        disable;
        traceoptions {
            flag {
                all <extensive>;
            }
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            (no-world-readable | world-readable);
            size maximum-file-size;
        }
        level (brief | detail | extensive | verbose);
        no-remote-trace;
    }
    twamp {
        traceoptions {
            flag {
```

```
                all <extensive>;
            }
        }
    }
}
```

Hierarchy Level

```
[edit security],
[edit tenants tenant-name security],
[edit services],
[edit logical-systems name security]
```

Description

Configure an Application Layer Gateway (ALG) on the device. An ALG runs as a service and can be associated in policies with specified types of traffic. ALGs are enabled by default.

Options

`bytes`—Maximum length in bytes of a single DNS message.

- **Range:** 512 through 8192 bytes
- **Default:** 512 bytes

`c-timeout` *`value-in-minutes`*—Specifies the timeout interval for Session Initiation Protocol (SIP) transactions in minutes.

- **Range:** 3 through 10 minutes
- **Default:** 3 minutes

`endpoint-registration-timeout` *`value-in-seconds`*—Specifies the timeout value in seconds for entries in the NAT table.

- **Range:** 10 through 50,000 seconds
- **Default:** 3600 seconds

`inactive-media-timeout` *`value-in-seconds`*—(MGCP) Specifies the maximum time duration that the temporary openings in the firewall (pinholes) remain open for media if no activity is detected.

- **Range:** 10 through 2550 seconds
- **Default:** 120 seconds

`inactive-media-timeout` *`value-in-seconds`*—(SCCP) Specifies the maximum time duration that the temporary openings in the firewall (pinholes) remain open for media if no activity is detected.

- **Range:** 10 through 600 seconds
- **Default:** 120 seconds

`inactive-media-timeout` *`value-in-seconds`*—(SIP) Specifies the maximum time duration that the temporary openings in the firewall (pinholes) remain open for media if no activity is detected.

- **Range:** 0 through 2550 seconds
- **Default:** 120 seconds

maximum-call-duration *`value-in-minutes`*—Specifies the maximum time duration at which the call ends and releases the media sessions.

- **Range:** 3 through 720 minutes
- **Default:** 720 minutes

`media-source-port-any`—Allow media traffic from any port number. By default, this feature is disabled, which allows a temporary opening in the firewall (pinhole) for media traffic to be opened.

`retain-hold-resource`—Enable the device to not free media resources for a Session Initiation Protocol (SIP) Application Layer Gateway (ALG), even when a media stream is placed on hold. By default, media stream resources are released when the media stream is held.

`transaction-timeout` *`value-in-seconds`*—Specifies timeout value for Media Gateway Control Protocol (MGCP) transactions. If the timeout value exceeds the limit, the transaction removed by MGCP transactions ager out processing.

- **Range:** 3 through 50 seconds
- **Default:** 30 seconds

`t1-interval` *`value-in-milliseconds`*—Specifies the maximum round-trip time (RTT) (in milliseconds) allowed for Session Initiation Protocol (SIP) transactions.

- **Range:** 500 through 5000 milliseconds
- **Default:** 500 milliseconds

`t4-interval` *`value-in-seconds`*—Specifies the maximum length of time (in seconds) that the network can take to clear messages between client and server Session Initiation Protocol (SIP) transactions.

- **Range:** 5 through 10 seconds
- **Default:** 5 seconds

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Statement supported in Junos OS Release 18.3R1 for tenant systems.

`maximum-message-length` *bytes* option introduced in Junos OS Release 10.1.

`twamp` option introduced in Junos OS Release 18.2R1.

alg (Applications)

IN THIS SECTION

- Syntax | 2138
- Hierarchy Level | 2139
- Description | 2139
- Options | 2139
- Required Privilege Level | 2140
- Release Information | 2140

Syntax

```
alg (dns |ftp | ftp-data | gprs-gtp-c | gprs-gtp-u | gprs-gtp-v0 | gprs-sctp | http | https |
ignore | ike-esp-nat | imap | login | mgcp-ca | mgcp-ua | ms-rpc | none | pptp | q931 | ras |
```

```
realaudio | rsh | rtsp | sccp | sip | smtp | sqlnet-v2 | ssh | sun-rpc | talk | telnet | tftp |
twamp);
```

Hierarchy Level

```
[edit applications application application-name term term-name]
```

Description

Define individual Application Layer Gateway (ALG).

The `twamp` application protocol is introduced Junos OS Release 18.2R1.

Options

ALG supports the following application protocols:

- `dns` —Domain Name Service
- `ftp` —File Transfer Protocol
- `ftp-data` —Trace FTP Data Session
- `gprs-gtp-c` —GPRS Tunneling Control Plane
- `gprs-gtp-u` —GPRS Tunneling User Plane
- `gprs-gtp-v0` —GPRS Tunneling Version 0
- `gprs-sctp` —GPRS Stream Control Protocol
- `http` —Hypertext Transfer Protocol
- `https` —Hypertext Transfer Protocol
- `ignore`—Ignore application type
- `ike-esp-nat` —IKE ESP NAT application protocol
- `mgcp-ca` —Media Gateway Control Protocol with Call Agent
- `mgcp-ua` —MGCP with User Agent
- `ms-rpc` —Microsoft RPC

- `none` —Bypass all ALG application protocols.
- `pptp` —Point-to-Point Tunneling Protocol
- `q931` —ISDN connection control protocol (Q.931)
- `ras` —Remote Access Service
- `realaudio` —RealAudio
- `rsh` —UNIX remote shell services
- `rtsp` —Real-Time Streaming Protocol
- `sccp` —Skinny Client Control Protocol
- `sip`—Session Initiation Protocol
- `smtp` —Simple Mail Transfer Protocol
- `sqlnet-v2` —Oracle SQLNET v2
- `ssh` —Secure Shell Protocol
- `sun-rpc` —Sun Microsystems RPC
- `talk`—TALK program
- `telnet` —Telnet Protocol
- `tftp` —Trivial File Transfer Protocol
- `twamp`—Two-Way Active Measurement Protocol

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

term (Applications) | **29110**

alg-manager

IN THIS SECTION

- Syntax | 2141
- Hierarchy Level | 2141
- Description | 2141
- Options | 2142
- Required Privilege Level | 2142

Syntax

```
alg-manager {
    clear-dont-frag-bit {
        traceoptions {
            flag {
                all {
                    extensive
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit security alg]
```

Description

Configure the Application Layer Gateway (ALG) manager.

Options

clear-dont-frag-bit Starting in Junos OS Release 15.1X49-D200, we've introduced the `clear-dont-frag-bit` option to disable the do not fragment flag from the packet IP header, which allows to split the packet after the NAT translation.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

RELATED DOCUMENTATION

traceoptions (Security ALG) | 29770

alg-support-lib

IN THIS SECTION

- Syntax | 2142
- Hierarchy Level | 2143
- Description | 2143
- Options | 2143
- Required Privilege Level | 2143
- Release Information | 2143

Syntax

```
alg-support-lib {
    traceoptions {
        flag {
```

```
            all {
                extensive
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Configures the Application Layer Gateway (ALG) support library.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

algorithm

IN THIS SECTION

- Syntax | 2144
- Hierarchy Level | 2144
- Description | 2144
- Options | 2144
- Required Privilege Level | 2145
- Release Information | 2145

Syntax

```
algorithm algorithm-name;
```

Hierarchy Level

```
[edit protocols pim interface interface-name bfd-liveness-detection authentication],
[edit routing-instances routing-instance-name protocols pim interface interface-name bfd-
liveness-detection authentication]
```

Description

Specify the algorithm to use for BFD authentication.

Options

algorithm-name—Name of algorithm to use for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text. One or more passwords can be configured.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive rates greater than 100 ms.

- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive rates greater than 100 ms.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Bidirectional Forwarding Detection Authentication for PIM

Configuring BFD Authentication for PIM

authentication (Protocols PIM) | **2647**

algorithm (Junos FIPS)

IN THIS SECTION

- Syntax | **2146**
- Hierarchy Level | **2146**
- Description | **2146**
- Options | **2146**
- Required Privilege Level | **2146**
- Release Information | **2146**

Syntax

```
algorithm 3des-cbc;
```

Hierarchy Level

```
[edit security ipsec internal security-association manual direction encryption]
```

Description

Select the encryption algorithm for the internal Routing-Engine–to–Routing-Engine IPsec security association (SA) configuration.

Options

Only **3des-cbc** is supported.

Required Privilege Level

Crypto Officer—To add and view this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

alias (Interfaces)

IN THIS SECTION

- Syntax | 2147
- Hierarchy Level | 2147
- Description | 2147
- Options | 2147

Required Privilege Level | 2147
Release Information | 2148

Syntax

```
alias alias-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure a textual description of a physical interface or the logical unit of an interface to be the alias of an interface name. The alias name can be a single line of text. If the text contains spaces, enclose it in quotation marks. If you configure an alias name, the alias name is displayed instead of the interface name in the output of all `show`, `show interfaces`, and other operational mode commands. In Junos OS Release 12.3R8 and later, display of the alias can be suppressed in favor of the actual interface name by using the `display no-interface-alias` parameter along with the show command.

Options

alias-name Text to denote an easily identifiable, meaningful alias name for the interface. If the text includes spaces, enclose the entire text in quotation marks.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Example: Adding an Interface Alias Name

Junos OS Network Interfaces Library for Routing Devices

alias (Junos Fusion)

IN THIS SECTION

- Syntax | 2148
- Hierarchy Level | 2148
- Description | 2149
- Default | 2149
- Options | 2149
- Required Privilege Level | 2149
- Release Information | 2149

Syntax

```
alias alias;
```

Hierarchy Level

```
[edit chassis satellite-management fpc slot-id]
```

Description

Configure an alias to label a satellite device.

Satellite device alias configuration is optional, but recommended. In a Junos Fusion, satellite device aliases assist with administration tasks, such as monitoring satellite devices using **show** command outputs, as well as with some configuration tasks that provide an option to identify a satellite device by its alias.

Default

Satellite devices are not assigned an alias, by default.

Options

alias The user-defined text name of the alias.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

alias (Junos Fusion Satellite Device Clustering)

IN THIS SECTION

- Syntax | 2150
- Hierarchy Level | 2150
- Description | 2150
- Default | 2150
- Options | 2151
- Required Privilege Level | 2151
- Release Information | 2151

Syntax

```
alias alias;
```

Hierarchy Level

```
[edit chassis satellite-management cluster cluster-name fpc slot-id ]
```

Description

Configure an alias to label a satellite device in a satellite device cluster.

Satellite device alias configuration is optional, but recommended. In a Junos Fusion, satellite device aliases assist with administration tasks, such as monitoring satellite devices using **show** command outputs, as well as with some configuration tasks that provide an option to identify a satellite device by its alias.

Default

Satellite devices in a satellite device cluster are not assigned an alias, by default.

Options

alias The user-defined text name of the alias.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

aliases

IN THIS SECTION

- Syntax | 2151
- Hierarchy Level | 2152
- Description | 2152
- Required Privilege Level | 2152
- Release Information | 2152

Syntax

```
aliases {
    director-device director-device-name {
```

```
        assigned-director-device-name;
    }
    interconnect-device interconnect-device-name {
        assigned-interconnect-device-name;
    }
    node-device node-device-name {
        assigned-node-device-name;
    }
```

Hierarchy Level

```
[edit fabric]
```

Description

(QFabric systems only) Specify the mapping of user defined names to QFabric system components. You can reassign names for Director devices, Interconnect devices, and Node devices.

The remaining statements are explained separately. See CLI Explorer.

NOTE: The following rules apply to QFabric component alias naming:

- Alias names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
- The maximum length of an alias name is 30 characters.
- Alias names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.
- You cannot use the reserved names `all`, `fabric`, or `director-group` as an alias name.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

Options `director-group` and `interconnect-device` introduced in Junos OS Release 12.2 for the QFX Series.

RELATED DOCUMENTATION

Configuring Aliases for the QFabric System

Understanding Node Devices

aliases (Virtual Chassis)

IN THIS SECTION

- Syntax | 2153
- Hierarchy Level | 2153
- Description | 2154
- Options | 2154
- Required Privilege Level | 2155
- Release Information | 2155

Syntax

```
aliases {
    serial-number serial-number {
        alias-name alias-name;
    }
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Optionally create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). You can more clearly identify the member switches in your Virtual Chassis or VCF by assigning text labels (alias names) to them, associating the alias name with member switch serial numbers. Aliases are for administrative purposes only and have no effect on member switch operations.

The `show virtual-chassis` command displays assigned alias names in the `Alias-Name` output field.

Options

serial-number The permanent serial number for the member switch in the Virtual Chassis or VCF that you want to label with an alias name. Serial number values are case-sensitive.

The `show virtual-chassis` command displays the serial number for all members in the Virtual Chassis or VCF in the `Serial No` output field.

alias-name The text label you want to assign to the member switch.

For example, if you configure alias name `dc-floor-1` in a VCF for member switch 0 with serial number AB012345678 as follows:

```
[edit virtual-chassis aliases]
user@switch# set serial-number AB0123456789 alias-name dc-floor-1
```

then you'll see the alias **dc-floor-1** in the Alias-Name column in the output from the `show virtual-chassis` command:

show virtual-chassis

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 9d5d.5556.919a
Fabric Mode: Enabled
                                                                  Mstr
Member ID  Status   Serial No      Alias-Name     Model            prio    Role
0 (FPC 0)  Prsnt    AB0123456789   dc-floor-1     qfx5100-48s-6q   129     Master
.
```

```
.
.
```

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- Configuring an EX4650 or a QFX Series Virtual Chassis
- *Understanding Virtual Chassis Fabric Components*
- Understanding Virtual Chassis Components

all-defects

IN THIS SECTION

- Syntax | **2156**
- Hierarchy Level | **2156**
- Description | **2156**
- Required Privilege Level | **2156**
- Release Information | **2156**

Syntax

```
all-defects;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
event ais-trigger-condition]
```

Description

Configure the defect condition that raises an alarm indication signal when any or all possible defects occur in the maintenance domain level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

> *Ethernet Alarm Indication Signal (ETH-AIS) Function Overview*
>
> *Configuring ETH-AIS on a CFM MEP*

all-regions

IN THIS SECTION

- Syntax | 2157
- Hierarchy Level | 2158
- Description | 2158
- Options | 2158
- Required Privilege Level | 2158
- Release Information | 2158

Syntax

```
all-regions {
    incoming;
    ingress-replication {
        create-new-ucast-tunnel;
        label-switched-path {
            label-switched-path-template (Multicast) {
                (default-template | lsp-template-name);
            }
        }
    }
    ldp-p2mp;
    rsvp-te {
        label-switched-path-template (Multicast) {
            (default-template | lsp-template-name);
        }
        static-lsp static-lsp;
    }
}
```

Hierarchy Level

```
[edit protocols mvpn inter-region-template template template-name]
```

Description

Specify the tunnel support for all regions that are not defined in the more specific regions. The `all-regions` statement supports different tunnel types such as incoming, ingress replication, ldp-p2mp, or rsvp-te. The incoming tunnel indicates that the tunnel type across the area boundary router (ABR) is not changed.

Options

incoming Specify the tunnel type to be the same across the ABR.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview

Example: Configuring Segmented Inter-Area P2MP LSP

Configuring Segmented Inter-Area P2MP LSP

inter-region | **8789**

inter-region-segmented | **8791**

inter-region-template | **8793**

region | **14262**

template(Protocols MVPN) | **29092**

all-tcp

IN THIS SECTION

- Syntax | 2159
- Hierarchy Level | 2159
- Description | 2159
- Options | 2159
- Required Privilege Level | 2159
- Release Information | 2160

Syntax

```
all-tcp mss value;
```

Hierarchy Level

```
[edit security flow tcp-mss]
```

Description

Set the TCP maximum segment size (MSS) value to enable MSS override for all TCP packets in network traffic.

Options

`mss` *`value`*—TCP MSS value.

- **Range:** 64 through 65,535 bytes

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

tcp-mss (Security Flow) | 29000

Understanding Traffic Processing on Security Devices

allow

IN THIS SECTION

- Syntax | 2160
- Hierarchy Level | 2160
- Description | 2161
- Options | 2161
- Required Privilege Level | 2161
- Release Information | 2161

Syntax

```
allow (all | [ network/mask-length ]);
```

Hierarchy Level

```
[edit logical-systems name protocols bgp group name],
[edit logical-systems name routing-instances name protocols bgp group name],
[edit protocols bgp group name],
[edit routing-instances name protocols bgp group name],
```

```
[edit logical-systems name protocols bgp group name dynamic-neighbor dyn-name],
[edit logical-systems name routing-instances name protocols bgp group name dynamic-neighbor dyn-
name],
[edit logical-systems name tenants name routing-instances name protocols bgp group name dynamic-
neighbor dyn-name],
[edit protocols bgp group name dynamic-neighbor dyn-name],
[edit routing-instances name protocols bgp group name dynamic-neighbor dyn-name]
```

Description

Implicitly configure BGP peers, allowing peer connections from any of the specified networks or hosts. To configure multiple BGP peers, configure one or more networks and hosts within a single `allow` statement or include multiple `allow` statements.

You can configure authentication for all implicitly configured peers at `[edit protocols bgp group]` level and to configure different authentication values for each prefix, you must configure `allow` under `dynamic-neighbor` *`dyn-name`* hierarchy.

You can limit the number of BGP sessions belonging to a given subnet using the `peer-limit` statement.

> **NOTE**: When `set protocols bgp group` *`group-name`* `allow network` is configured to accept dynamic BGP sessions, `unconfigured-peer-graceful-restart` statement should be configured to avoid traffic drop during graceful restart or graceful Routing Engine switchover.

Options

all—Allow all addresses, which is equivalent to **0.0.0.0/0** (or **::/0**).

network/ mask-length—IPv6 or IPv4 network number of a single address or a range of allowable addresses for BGP peers, followed by the number of significant bits in the subnet mask.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced under `[edit protocols bgp group group-name dynamic-neighbor dyn-name]` hierarchy in Junos OS Release 19.1.

RELATED DOCUMENTATION

neighbor (Protocols BGP) | 11273

allow-6pe-traceroute

IN THIS SECTION

- Syntax | 2162
- Hierarchy Level | 2162
- Description | 2162
- Required Privilege Level | 2163

Syntax

```
allow-6pe-traceroute;
```

Hierarchy Level

```
[edit system]
```

Description

Allow IPv4-mapped IPv6 source addresses in an ICMPv6 echo request TTL expired packets.

In a dual-stack IPv6 network connected over an IPv4 MPLS network, the P routers in the IPv4 MPLS backbone do not have an IPv6 family. Consequently, the transit P routers are not shown in the output when you do an IPv6 traceroute. To generate an ICMPv6 echo request and a TTL expired response

packet to and from the intermediate transit routers in the 6PE network, you must configure this statement along with the `allow-v4-mapped` statement.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configure IPv6 Features

allow-v4mapped-packets | 2211

allow-active-leasequery (DHCP Local Server)

IN THIS SECTION

- Syntax | 2163
- Hierarchy Level | 2164
- Description | 2164
- Options | 2164
- Required Privilege Level | 2164
- Release Information | 2165

Syntax

```
allow-active-leasequery {
    idle-timeout seconds;
    timeout seconds;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services ...],
[edit logical-systems logical-system-name system services ...],
[edit routing-instances routing-instance-name system services ...]
```

Description

Enable a DHCPv4 or DHCPv6 local server to listen for, process, and respond to active leasequery requests received on TCP connections on TCP port 67 for DHCPv4 and on TCP port 547 for DHCPv6.

NOTE: Because active leasequery is an extension of bulk leasequery, you must also configure bulk leasequery before you configure active leasequery.

Options

idle-timeout *seconds*

(Optional) Specify the number of seconds that pass with no activity before the TCP connection is terminated and restarted.

- **Range:** 10 through 3600
- **Default:** 60

timeout *seconds*

(Optional) Specify the number of seconds that TCP read/write operations can be blocked before the TCP connection is terminated and restarted.

- **Range:** 10 through 3600
- **Default:** 120

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

DHCP Leasequery Methods

allow-bulk-leasequery (DHCP Local Server)

IN THIS SECTION

- Syntax | 2165
- Hierarchy Level | 2165
- Description | 2166
- Options | 2166
- Required Privilege Level | 2167
- Release Information | 2167

Syntax

```
allow-bulk-leasequery {
    max-connections number-of-connections;
    max-empty-replies number-of-replies;
    restricted-requestor;
    timeout seconds;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services ...],
[edit logical-systems logical-system-name system services ...],
[edit routing-instances routing-instance-name system services ...]
```

Description

Enable a DHCPv4 or DHCPv6 local server to listen for, process, and respond to bulk leasequery requests received on TCP connections on TCP port 67 for DHCPv4 and on TCP port 547 for DHCPv6.

Options

max-connections *number-of-connections*

Specify the maximum number of concurrent TCP connections allowed in the logical system/routing instance. This setting helps you manage the resources that the `jdhcpd` daemon uses for bulk leasequery operations in the logical system/routing instance. The number of connections you specify for a logical system/routing instances must be less than the number of connections you specify for the global `accept-max-tcp-connections` statement.

- **Range:** 1 through 10
- **Default:** 3

max-empty-replies *number-of-replies*

Specify the maximum number of empty replies that the DHCP local server sends to a specific requestor. When the maximum number is reached, the DHCP server closes the connection.

An empty reply is a response sent from the DHCP local server that contains no bindings or has an option status code error. Empty replies are often the response to an unauthorized requestor that has sent an invalid or incorrect query that produces no binding information. By limiting the number of empty replies that the DHCP local server can send, you prevent the connection from being consumed by an unauthorized or malicious requester, and free up the DHCP local server to support legitimate requestors.

- **Range:** 1 through 100
- **Default:** 5

restricted-requestor

Specify that the DHCP local server responds to a bulk leasequery request by sending the binding information to restricted requestors only. This ensures that the requestor is the originator of the binding.

Restricted requestors are determined as follows:

- For DHCPv4 bulk leasequery requests, the giaddr of the requestor must match the giaddr of the client.
- For DHCPv6 bulk leasequery requests, the requestor's client ID in the bulk leasequery message must match the relay ID sent during binding creation.

timeout *seconds*

Specify the number of seconds that a connection on the TCP socket can be idle before the DHCP local server closes the connection. Closing inactive connections enables the DHCP local server to more efficiently apply resources in support of active TCP connections.

- **Range:** 1 through 1000
- **Default:** 120

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring and Using DHCP Bulk Leasequery

allow-clear (Accounting Options)

IN THIS SECTION

- Syntax | 2168
- Hierarchy Level | 2168
- Description | 2168
- Required Privilege Level | 2168

Release Information | 2168

Syntax

```
allow-clear;
```

Hierarchy Level

```
[edit accounting-options interface-profile profile-name]
```

Description

Reports accounting statistics values cleared by the `clear interfaces statistics` command on a logical interface configured with accounting options to the accounting options flat file associated with the interface. By default, this statement is disabled and the cleared statistics are not reported to the flat file, although the counters do show as cleared in the CLI.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configure the Interface Profile

Accounting Options Configuration

allow-commands (Op Scripts)

IN THIS SECTION

- Syntax | 2169
- Hierarchy Level | 2169
- Description | 2169
- Required Privilege Level | 2170
- Release Information | 2170

Syntax

```
allow-commands "regular-expression";
```

Hierarchy Level

```
[edit system scripts op file filename]
```

Description

Specify a regular expression that defines which operational mode commands to explicitly allow during execution of a local op script.

By default, when a user executes a script, the system does not permit the user to execute operational mode commands within a script for which their login class does not normally have permission to execute. You can explicitly enable the script to execute certain operational mode commands regardless of the permissions of the user invoking the script.

NOTE: The `allow-commands` statement does not support executing configuration mode commands.

> **NOTE**: Remote op scripts that are executed using the `op url` command do not support executing unauthorized operational mode commands even when you configure the `allow-commands` statement.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

- *Define Operational Mode Commands to Allow in an Op Script*
- *Enable an Op Script and Define a Script Alias*
- *Disable an Op Script*
- Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands, Configuration Statements, and Hierarchies

allow-dns-reply

IN THIS SECTION

- Syntax | 2171
- Hierarchy Level | 2171
- Description | 2171
- Required Privilege Level | 2171
- Release Information | 2171

Syntax

```
allow-dns-reply;
```

Hierarchy Level

```
[edit security flow]
```

Description

Allow an incoming Domain Name Service (DNS) reply packet without a matched request. By default, if an incoming UDP first-packet has dst-port 53, the device checks the DNS message packet header to verify that the query bit (QR) is 0, which denotes a query message. If the QR bit is 1, which denotes a response message, the device drops the packet, does not create a session, and increments the illegal packet flow counter for the interface. Using the `allow-dns-reply` directs the device to skip the check.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

allow-duplicates

IN THIS SECTION

- Syntax | 2172
- Hierarchy Level | 2172
- Description | 2172
- Options | 2173
- Required Privilege Level | 2173
- Release Information | 2173

Syntax

```
allow-duplicates;
```

Hierarchy Level

```
[edit logical-systems logical-system-name system syslog],
[edit logical-systems logical-system-name system syslog file file-name],
[edit logical-systems logical-system-name system syslog host host-name],
[edit logical-systems logical-system-name system syslog user user-name],
[edit system syslog],
[edit system syslog file file-name],
[edit system syslog host host-name],
[edit system syslog user user-name]
```

Description

Specify whether to allow the repeated messages in the system log output files. This can be set either at global configuration level or for individual file, host, or user. By default, this parameter is set to disable.

Options

file Name of the file to log messages

host Host to receive the messages

user User to receive the notification of the event

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Logical systems support introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

syslog (System) | 28836

allow-dynamic-requests (RADIUS Server)

IN THIS SECTION

- Syntax | 2174
- Hierarchy Level | 2174
- Description | 2174
- Required Privilege Level | 2174
- Release Information | 2174

Syntax

```
allow-dynamic-requests;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Allow dynamic requests from the RADIUS server so that change of authorization requests can be received.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

allow-email (Security Fallback Block)

IN THIS SECTION

- Syntax | 2175
- Hierarchy Level | 2175
- Description | 2175
- Required Privilege Level | 2175
- Release Information | 2176

Syntax

```
allow-email;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-block]
```

Description

Enable e-mail notification to notify a specified administrator when a fallback-block occurs.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

allow-email (Security Virus Detection)

IN THIS SECTION

- Syntax | 2176
- Hierarchy Level | 2176
- Description | 2176
- Required Privilege Level | 2177
- Release Information | 2177

Syntax

```
allow-email;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus profile notification-options virus-detect]
```

Description

Enable e-mail notification to notify a specified administrator when a virus is detected.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

allow-embedded-icmp

IN THIS SECTION

- Syntax | 2177
- Hierarchy Level | 2177
- Description | 2178
- Required Privilege Level | 2178
- Release Information | 2178

Syntax

```
allow-embedded-icmp;
```

Hierarchy Level

```
[edit security flow]
```

Description

Allow ICMP error packets to pass through the device even when there is no session match for the embedded packet. Once enabled, all packets encapsulated in ICMP pass through and no policy affects this behavior. This feature is useful when you have asymmetric routing in your network and you want to use traceroute and other ICMP applications on your device.

The default behavior is to inspect the IP packet which is embedded in the ICMP error packet. If the embedded packet is part of an existing session, the ICMP packet is allowed to pass through. If there is no match, it is dropped. Without the `allow-embedded-icmp` option configured, the default behavior applies.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D10.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

allow-fragmentation (protocols mpls)

IN THIS SECTION

- Syntax | 2179
- Hierarchy Level | 2179
- Description | 2179
- Required Privilege Level | 2179
- Release Information | 2179

Syntax

```
allow-fragmentation;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls path-mtu],
[edit protocols mpls path-mtu]
```

Description

Allow IP packets to be fragmented before they are encapsulated in MPLS.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Packet Fragmentation

allow-fragmentation

IN THIS SECTION

Syntax | 2180
Hierarchy Level | 2180

- Description | 2180
- Default | 2180
- Required Privilege Level | 2181
- Release Information | 2181

Syntax

```
allow-fragmentation;
```

Hierarchy Level

```
[edit interfaces gr-fpc/pic/port unit logical-unit-number tunnel],
[edit logical-systems logical-system-name interfaces gr-fpc/pic/port unit logical-unit-number
tunnel]
```

Description

For a generic routing encapsulation (GRE) tunnel, enable fragmentation of GRE-encapsulated packets whose size exceeds the maximum transmission unit (MTU) value of a link that the packet passes through. The don't-fragment (DF) bit is not set in the outer IP header of GRE-encapsulated packets.

To enable the reassembly of fragmented GRE-encapsulated packets on GRE tunnel interfaces at the endpoint of the GRE tunnel, include the "reassemble-packets" on page 14100 statement for the interface.

> **NOTE**: The "reassemble-packets" on page 14100 statement is not supported on MPC10E line card in Junos OS Release 19.3.

Default

If you do not include the `allow-fragmentation` statement, fragmentation of GRE-encapsulated packets is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Statement introduced in Junos OS Release 19.3 for MPC10E line card.

RELATED DOCUMENTATION

reassemble-packets | 14100

Enabling Fragmentation and Reassembly on Packets After GRE-Encapsulation

Junos OS Services Interfaces Library for Routing Devices

allow-fragmentation (fti)

IN THIS SECTION

- Syntax | 2181
- Hierarchy Level | 2182
- Description | 2182
- Default | 2182
- Required Privilege Level | 2182
- Release Information | 2182

Syntax

```
allow-fragmentaion;
```

Hierarchy Level

```
[edit interfaces fti number unit logical-unit-number tunnel encapsulation (gre|udp|ipip)]
```

Description

For an encapsulated (UDP/GRE/IP-in-IP) tunnel, enable fragmentation of GRE-encapsulated packets whose size exceeds the maximum transmission unit (MTU) value of a link that the packet passes through. The don't-fragment (DF) bit is not set in the outer IP header of encapsulated packets.

Default

Don't Frafgment (DF) bit is set by default for the outer encapsulated IPv4 header.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

encapsulation

allow-icmp-without-flow

IN THIS SECTION

- Syntax | 2183
- Hierarchy Level | 2183
- Description | 2183
- Required Privilege Level | 2183

- Release Information | 2183

Syntax

```
(allow-icmp-without-flow | no-allow-icmp-without-flow);
```

Hierarchy Level

```
[edit security idp sensor-configuration flow]
```

Description

Allow an ICMP packet without matched request. By default the ICMP flow is enabled.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

allow-ip-options (IDS MS-MPC)

IN THIS SECTION

- Syntax | 2184
- Hierarchy Level | 2184
- Description | 2184

- Default | 2184
- Options | 2185
- Required Privilege Level | 2185
- Release Information | 2185

Syntax

```
allow-ip-options {
    any;
    loose-source-route;
    route-alert;
    route-record;
    security;
    stream-id;
    strict-source-route;
    timestamp;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Specify the type of IPv4 options that the IDS rule permits in the IP header of a packet. If the packet includes an option that is not configured, the packet is blocked. If the packet includes a configured option whose length is an illegal value, the packet is dropped. This IDS rule can only be assigned to a service set on an MS-MPC.

Default

If you do not include the `allow-ip-options` statement, packets with any type of IPv4 option are blocked.

Options

any	Allow all IPv4 options.
loose-source-route	Allow the Loose Source Route option.
route-alert	Allow the Router Alert option.
route-record	Allow the Record Route option.
security	Allow the Security option.
stream-id	Allow the Stream ID option.
strict-source-route	Allow the Strict Source Route option.
timestamp	Allow the Time Stamp option.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

allow-ip-options (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 2186
- Hierarchy Level | 2186
- Description | 2186
- Options | 2186
- Required Privilege Level | 2187
- Release Information | 2187

Syntax

```
allow-ip-options values;
```

Hierarchy Level

```
[edit services (Stateful Firewall) stateful-firewall rule (Services Stateful Firewall)  rule-name term
(Services Stateful Firewall)  term-name then (Services Stateful Firewall) ]
```

Description

Configure how the stateful firewall handles IP header information. This statement is optional.

Options

value—Can be a set or range of numeric values, or one or more of the following predefined option types. You can enter either the option name or its numeric equivalent.

Option Name	Numeric Value
any	0

(Continued)

Option Name	Numeric Value
ip-security	130
ip-stream	8
loose-source-route	3
route-record	7
router-alert	148
strict-source-route	9
timestamp	4

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Stateful Firewall Rules

allow-ipv6-extension-header (IDS MS-MPC)

IN THIS SECTION

- Syntax | 2188
- Hierarchy Level | 2188
- Description | 2188
- Default | 2189
- Options | 2189
- Required Privilege Level | 2189
- Release Information | 2189

Syntax

```
allow-ipv6-extension-header {
    ah;
    any;
    dstopts;
    esp;
    fragment;
    hop-by-hop;
    mobility;
    routing;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Specify the type of IPv6 extension headers that the IDS rule permits in a packet. If the packet includes an extension header that is not configured, the packet is blocked. If the packet includes configured

extension headers that are incorrect, the packet is dropped. This IDS rule can only be assigned to a service set on an MS-MPC.

Default

If you do not include the `allow-ipv6-extension-header` statement, packets with any type of extension header are blocked.

Options

ah	Allow Authentication Header extension header.
any	Allow all IPv6 extension headers.
dstopts	Allow Destination Options extension header.
esp	Allow Encapsulating Security Payload extension header.
fragment	Allow Fragment Header extension header.
hop-by-hop	Allow Hop-by-Hop Option extension header.
mobility	Allow Mobility Header extension header.
routing	Allow Routing Header extension header.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

allow-l3vpn-traceroute-src-select

IN THIS SECTION

- Syntax | 2190
- Hierarchy Level | 2190
- Description | 2190
- Required Privilege Level | 2190
- Release Information | 2191

Syntax

```
allow-l3vpn-traceroute-src-select;
```

Hierarchy Level

```
[edit system]
```

Description

Select best source address for ICMP time-to-live (TTL) expiry error in case of Layer 3 VPNs. When a traceroute occurs over Layer 3 VPN and the TTL expired traffic is generated, the traceroute returns to the source through the routing instance. `allow-l3vpn-traceroute-src-select` determines the correct IP source address by reviewing the destination routing instance and destination IP address for Layer 3 VPN provider edge (PE) routers.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

vrf-table-label | 30966

Pinging a Layer 3 VPN

allow-leasequery (DHCP Local Server)

IN THIS SECTION

- Syntax | 2191
- Hierarchy Level | 2191
- Description | 2192
- Options | 2192
- Required Privilege Level | 2192
- Release Information | 2192

Syntax

```
allow-leasequery {
    restricted-requestor;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services ...],
```

```
[edit logical-systems logical-system-name system services ...],
[edit routing-instances routing-instance-name system services ...]
```

Description

Enable a DHCPv4 or DHCPv6 local server to support individual leasequery operations by listening on the UDP socket and responding to leasequery requests received on UDP connections.

Options

restricted-requestor Specify that the DHCP local server responds to an individual leasequery request by sending the binding information to restricted requestors only. This ensures that the requestor is the originator of the binding.

Restricted requestors are determined as follows:

- For DHCPv4 individual leasequery requests, the giaddr of the requestor must match the giaddr of the client.
- For DHCPv6 individual leasequery requests, the requestor's client ID in the bulk leasequery message must match the relay ID sent during binding creation.

> **NOTE**: The `restricted-requestor` statement is not supported for the DHCPv6 `allow-leasequery` statement.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring and Using DHCP Individual Leasequery

Configuring and Using DHCP Bulk Leasequery

allow-maximum (Multicast)

IN THIS SECTION

- Syntax | 2193
- Hierarchy Level | 2193
- Description | 2193
- Default | 2194
- Required Privilege Level | 2195
- Release Information | 2195

Syntax

```
allow-maximum;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast forwarding-cache],
[edit logical-systems logical-system-name routing-options multicast forwarding-cache],
[edit routing-instances routing-instance-name routing-options multicast forwarding-cache],
[edit routing-options multicast forwarding-cache]
```

Description

Allow the larger of global and family-level threshold values to take effect.

This statement is optional when you configure a forwarding cache or PIM state limits. When this statement is included in the configuration and both a family-specific and a global configuration are present, the higher limits take precedence.

For example:

```
[edit routing-options multicast forwarding-cache]
allow-maximum;
family inet {
    threshold {
        suppress 100;
        reuse 75;
    }
}
family inet6 {
    threshold {
        suppress 600;
        reuse 500;
    }
}
threshold {
    suppress 400;
    reuse 450;
}
```

```
user@host# show multicast forwarding-cache statistics

Instance: master Family: INET
Suppress Threshold                      400
Reuse Value                             400
Currently Used Entries                    0

Instance: master Family: INET6
Suppress Threshold                      600
Reuse Value                             500
Currently Used Entries                    0
```

This statement can be useful in single-stack devices on which IPv4 traffic is expected or IPv6 traffic is expected, but not both.

Default

By default, this statement is disabled.

When this statement is omitted from the configuration, a family-specific forwarding cache configuration and a global forwarding cache configuration cannot be configured together. Either the global-specific configuration or the family-specific configuration is allowed, but not both.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring the Multicast Forwarding Cache

Example: Configuring PIM State Limits

allow-multicast

IN THIS SECTION

- Syntax | 2195
- Hierarchy Level | 2196
- Description | 2196
- Required Privilege Level | 2196
- Release Information | 2196

Syntax

```
allow-multicast;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Allow multicast traffic to be sent to the Adaptive Services or Multiservices PIC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Enabling Services PICs to Accept Multicast Traffic

allow-no-end-option (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2197
- Hierarchy Level | 2197
- Description | 2197
- Required Privilege Level | 2197
- Release Information | 2197

Syntax

```
allow-no-end-option;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Enable a DHCP relay agent to process packets sent from clients without DHCP Option-255 (end-of-options).

The default behavior in Junos OS is to drop packets that do not include Option 255. To override this default behavior, configure the `allow-no-end-option` CLI statement at the `[edit forwarding-options dhcp-relay overrides]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Overriding the Default DHCP Relay Configuration Settings

DHCP Snooping Support

allow-overlapping-nat-pools

IN THIS SECTION

- Syntax | 2198
- Hierarchy Level | 2198
- Description | 2198
- Required Privilege Level | 2198
- Release Information | 2198

Syntax

```
allow-overlapping-nat-pools;
```

Hierarchy Level

```
[edit services nat]
```

Description

Specify that NAT source or destination pools can be shared between multiple service sets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced with Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Service Sets for Network Address Translation

allow-overlapping-pools (NAT Next Gen Services)

IN THIS SECTION

- Syntax | 2199
- Hierarchy Level | 2199
- Description | 2199
- Required Privilege Level | 2199
- Release Information | 2200

Syntax

```
allow-overlapping-pools;
```

Hierarchy Level

```
[edit services nat]
```

Description

Specify that NAT source or destination pools can have IP addresses that overlap with IP addresses in pools used in other service sets. However, pools that configure port-block allocation must not overlap with other pools.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

allow-remote-loopback

IN THIS SECTION

- Syntax | 2200
- Hierarchy Level | 2200
- Description | 2200
- Required Privilege Level | 2201
- Release Information | 2201

Syntax

```
allow-remote-loopback;
```

Hierarchy Level

```
[edit protocols oam link-fault-management interface interface-name negotiation-options]
```

Description

Enable the remote loopback on IQ2 and IQ2-E Gigabit Ethernet interfaces, and Ethernet interfaces on the MX Series routers and EX Series switches. Advertise that the interface is capable of getting into loopback mode. Enable remote loopback in Ethernet OAM link fault management (LFM) on all Ethernet interfaces or the specified interface on the EX Series switch.

WARNING: If you disable this statement on a peer interface, LFM loopback enable and disable commands will not work. Before disabling this configuration, please make sure the remote-loopback interface is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Enabling Remote Loopback Support on the Local Interface

allow-reverse-ecmp

IN THIS SECTION

- Syntax | 2201
- Hierarchy Level | 2202
- Description | 2202
- Required Privilege Level | 2202
- Release Information | 2202

Syntax

```
allow-reverse-ecmp
```

Hierarchy Level

```
[edit security flow]
```

Description

Enable ECMP support for reverse traffic. In this case, Junos OS for SRX Series Firewalls and vSRX Virtual Firewall instances use a hash algorithm to determine the interface to use for reverse traffic in a flow. This process is similar to asymmetric routing in which a packet traverses from a source to a destination in one path and takes a different path when it returns to the source.

If you do not enable this feature, the software selects a route in the ECMP set to the incoming interface for reverse traffic, which is the default behavior.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

Understanding ECMP Flow-Based Forwarding

Understanding ECMP Flow-Based Forwarding for Reverse Traffic on SRX Series Devices and vSRX Instances

allow-route-leaking

IN THIS SECTION

Syntax | 2203

- Hierarchy Level | 2203
- Description | 2203
- Required Privilege Level | 2203
- Release Information | 2203

Syntax

```
allow-route-leaking;
```

Hierarchy Level

```
[edit logical-systems name protocols ospf overload],
[edit logical-systems name routing-instances name protocols ospf overload],
[edit protocols ospf overload],
[edit routing-instances name protocols ospf overload],
[edit protocols ospf3 overload]
```

Description

Allow routes to be leaked when OSPF overload is configured and advertise the external prefixes with normal cost.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2.

RELATED DOCUMENTATION

Understanding OSPF Overload Function

stub-network | 28667

intra-area-prefix | 9224

as-external | 2486

allow-snooped-clients

IN THIS SECTION

- Syntax | 2204
- Hierarchy Level | 2204
- Description | 2205
- Default | 2205
- Required Privilege Level | 2205
- Release Information | 2205

Syntax

```
allow-snooped-clients;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
```

```
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Explicitly enable DHCP snooping support on the DHCP relay agent.

Use the statement at the `[edit ... dhcpv6]` hierarchy levels to explicitly enable snooping support on the router for DHCPv6 relay agent.

Default

DHCP snooping is disabled by default.

> **NOTE**: On EX4300 and EX9200 switches, the `allow-snooped-clients` statement is enabled by default at the `[edit forwarding-options dhcp-relay overrides]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- *Extended DHCP Relay Agent Overview*
- *Overriding the Default DHCP Relay Configuration Settings*
- DHCP Snooping Support
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent

allow-sram-parity-errors

IN THIS SECTION

- Syntax | 2206
- Hierarchy Level | 2206
- Description | 2206
- Required Privilege Level | 2206
- Release Information | 2206

Syntax

```
allow-sram-parity-errors;
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

(T Series routers only) Allow SRAM parity errors to occur without restarting the FPC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

allow-subnet-mismatch

IN THIS SECTION

- Syntax | 2207
- Hierarchy Level | 2207
- Description | 2207
- Default | 2207
- Required Privilege Level | 2207
- Release Information | 2208

Syntax

```
allow-subnet-mismatch;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp interface interface-name],
[edit protocols ldp interface interface-name]
```

Description

Ignore the LDP subnet check. For Junos OS Release 8.4 and later releases, an LDP source address subnet check was added for the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address.

Default

The source address in the LDP link hello packet is matched against the interface address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Ignoring the LDP Subnet Check

allow-transients

IN THIS SECTION

- Syntax | 2208
- Hierarchy Level | 2208
- Description | 2209
- Default | 2209
- Required Privilege Level | 2209
- Release Information | 2209

Syntax

```
allow-transients;
```

Hierarchy Level

```
[edit system scripts commit],
[edit system scripts commit file filename]
```

Description

Enable commit scripts to commit transient configuration changes. Transient changes are made in the checkout configuration but not in the candidate configuration.

Default

Transient changes are disabled by default. If you do not include the `allow-transients` statement, and an enabled commit script generates transient changes, the command-line interface (CLI) generates an error message and the commit operation fails.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

Support at the [`edit system scripts commit file` *`filename`*] hierarchy level introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Generating a Persistent or Transient Configuration Change in SLAX and XSLT Commit Scripts

Creating a Commit Script Macro to Read the Custom Syntax and Generate Related Configuration Statements

allow-url-for-python (Op Scripts)

IN THIS SECTION

- Syntax | 2210
- Hierarchy Level | 2210
- Description | 2210
- Default | 2210

Required Privilege Level | 2211
Release Information | 2211

Syntax

```
allow-url-for-python;
```

Hierarchy Level

```
[edit system scripts op]
```

Description

Enable the execution of Python op scripts from a remote site.

When you configure the `language python` or `language python3` statement at the `[edit system scripts]` hierarchy level, you can execute local Python scripts on a device running Junos OS. In order to execute Python op scripts from a remote site, you must also configure the `allow-url-for-python` statement.

To verify the integrity of the remote op script before it runs on the device, you can optionally confirm the checksum hash for the script. To determine the checksum, use the `file checksum` operational mode command. When you run the `op url` command you can provide the hash value for that script by using the `key` option. If the calculated and specified values do not match, script execution fails.

> **NOTE**: Starting in Junos OS Release 18.2R2 and 18.3R1, Junos OS supports only the SHA-256 hash function for script checksum hashes. Earlier releases support the *MD5*, *SHA-1*, and SHA-256 hash functions.

The `no-allow-url` statement takes precedence over the `allow-url-for-python` statement; if both statements are configured, you cannot execute Python op scripts from a remote site on the device.

Default

If you do not include the `allow-url-for-python` statement, you cannot execute Python op scripts from a remote site, and the `op url` command generates an error.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

> *Execute an Op Script from a Remote Site*
>
> no-allow-url (Op Scripts) | **11392**

allow-v4mapped-packets

IN THIS SECTION

- Syntax | **2211**
- Hierarchy Level | **2211**
- Description | **2212**
- Options | **2212**
- Required Privilege Level | **2212**
- Release Information | **2212**

Syntax

```
allow-v4mapped-packets;
```

Hierarchy Level

```
[edit system]
```

Description

Enable the processing of IPv4-mapped IPv6 packets.

Options

None

- **Default:** IPv4-mapped IPv6 address processing is disabled.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configure IPv6 Features

allow-6pe-traceroute | 2162

allowed-destinations

IN THIS SECTION

- Syntax | 2213
- Hierarchy Level | 2213
- Description | 2213
- Options | 2213
- Required Privilege Level | 2213
- Release Information | 2213

Syntax

```
allowed-destinations [ destinations ];
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

Identify flow capture destinations that are allowed in messages sent from this control source.

Options

destinations—Allowed content destination name.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Control Source

allowed-mac

IN THIS SECTION

- Syntax | 2214
- Hierarchy Level | 2214
- Description | 2214
- Default | 2215
- Options | 2215
- Required Privilege Level | 2215
- Release Information | 2215

Syntax

```
allowed-mac {
                        mac-address-list;
}
```

Hierarchy Level

```
[edit ethernet-switching-options                secure-access-port                interface (all |
interface-name)]
```

Description

Specify particular MAC addresses to be added to the MAC address cache.

NOTE: Although this configuration restricts the addresses that can be added to the MAC address cache, it does not block the switch from receiving Layer 2 control packets—such as Link Layer Discovery Protocol (LLDP) packets—transmitted from MAC addresses that are not specified in the list of allowed MAC addresses. Control packets do not undergo the MAC address check and they are therefore included in the statistics of packets

received. However, they are not forwarded to another destination. They are trapped within the switch.

Default

Allowed MAC addresses take precedence over dynamic MAC values that have been applied with the `mac-limit` statement.

Options

mac-address-list—One or more MAC addresses configured as allowed MAC addresses for a specified interface or all interfaces.

Required Privilege Level

system—To view this statement in the configuration.
system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

mac-limit (Access Port Security) | **10241**

Example: Configuring Port Security (non-ELS)

Example: Protecting Against DHCP Snooping Database Attacks

Example: Protecting against Ethernet Switching Table Overflow Attacks

Example: Protecting against DHCP Starvation Attacks

Configuring MAC Limiting (non-ELS)

Configuring MAC Limiting (J-Web Procedure)

allowed-sources (REST API)

IN THIS SECTION

- Syntax | 2216
- Hierarchy Level | 2216
- Description | 2216
- Required Privilege Level | 2216
- Release Information | 2216

Syntax

```
allowed-sources [ value-list ];
```

Hierarchy Level

```
[edit system services rest control]
```

Description

Specify the allowed source IP addresses for the REST API process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

alt-name (Application Identification)

IN THIS SECTION

- Syntax | 2217
- Hierarchy Level | 2217
- Description | 2217
- Options | 2217
- Required Privilege Level | 2217
- Release Information | 2218

Syntax

```
alt-name alt-name
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Description

Provide an alternate name for the application.

Options

alt-name Alternate name for the application.

- **Range:** 1 through 255 characters

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

alternative-partition-name (OCS Partition)

IN THIS SECTION

- Syntax | 2218
- Hierarchy Level | 2218
- Description | 2219
- Required Privilege Level | 2219
- Release Information | 2219

Syntax

```
alternative-partition-name;
```

Hierarchy Level

```
[edit access ocs partition alternative-partition-name alternative-partition-name]
```

Description

Configure alternative partition name for OCS/GY. When configured, the specific partition name is used to forward diameter data along alternative data path in the specific cases.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

Configuring the OCS Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

always-mark-connection-protection-tlv

IN THIS SECTION

- Syntax | 2220
- Hierarchy Level | 2220
- Description | 2220
- Required Privilege Level | 2220
- Release Information | 2220

Syntax

```
always-mark-connection-protection-tlv;
```

Hierarchy Level

```
[edit logical-systems logical-systems-name protocols mpls interface interface-name],
[edit protocols mpls interface interface-name]
```

Description

(MX Series routers only) Enable you to switch an LSP away from a network node using a bypass LSP. This feature could be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.

This statement marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, you then need to configure the `switch-away-lsps` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Switching LSPs Away from a Network Node

always-write-giaddr

IN THIS SECTION

- Syntax | 2221
- Hierarchy Level | 2221
- Description | 2222
- Required Privilege Level | 2222
- Release Information | 2222

Syntax

```
always-write-giaddr;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Overwrite the gateway IP address (giaddr) of every DHCP packet with the giaddr of the DHCP relay agent before forwarding the packet to the DHCP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

dhcp-relay | 5620

always-write-option-82

IN THIS SECTION

- Syntax | 2222
- Hierarchy Level | 2223
- Description | 2223
- Required Privilege Level | 2223
- Release Information | 2223

Syntax

```
always-write-option-82;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Override the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server. The use of this option causes the DHCP relay agent to perform one of the following actions, depending on how it is configured:

- If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server.
- If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information

Extended DHCP Relay Agent Overview

ambient-temperature

IN THIS SECTION

- Syntax | 2224
- Hierarchy Level | 2224
- Description | 2224
- Default | 2225
- Options | 2225
- Required Privilege Level | 2226
- Release Information | 2226

Syntax

```
chassis ambient-temperature (25C|40C|55C);
```

Hierarchy Level

```
[edit chassis]
```

Description

Set the chassis ambient temperature and instruct the power manager to allocate power to the line cards according to the ambient temperature value.

On system initialization in a PTX Series router, the power manager reads the ambient temperature and allocates power to the line cards according to the power budget policy at that temperature. If the actual power consumption of any line card exceeds the configured value for more than three minutes, the power manager overrides the configured ambient temperature setting of that line card, and resets its

ambient temperature to the next higher level and reallocates power according to the new temperature setting. All the overshooting line cards remain in the dynamic ambient temperature mode until the next reboot, or until you override the configured ambient temperature with a CLI command. The power manager then resets the power budget of the FRUs according to the configured ambient temperature setting.

> **NOTE**: The `set chassis ambient-temperature` *`temperature`* and `show chassis ambient-temperature` commands are not supported on PTX10004, PTX10008, and PTX10016 routers because their power allocation is based on the default ambient temperature.

When an MX Series router or an EX9200 switch restarts, the system adjusts the power allocation or the provisioned power for the line cards on the basis of the configured ambient temperature. If enough power is not available, a minor chassis alarm is raised. However, the chassis continues to run with the configured ambient temperature. You can configure a new higher ambient temperature only after you make more power available by adding new power supply modules or by taking a few line cards offline. By using the provisioned power that is saved by configuring a lower ambient temperature, you can bring more hardware components online.

For MX960 Universal Routing Platform and the MX2000 line of routers, a maximum operating temperature of 40°C is recommended at an altitude of up to 6000 feet above sea level. The router can operate up to a maximum temperature of 46°C at sea level and can be set accordingly.

> **NOTE**: If ambient temperature is not set, then the line cards are allocated power according to the default ambient temperature.

Default

- MX Series routers—40°C
- PTX Series routers—40°C
- EX9200 switches—40°C
- EX9251 switches—40°C
- EX9253 switches—40°C

Options

25C Set the ambient temperature of the chassis to 25°C

40C Set the ambient temperature of the chassis to 40°C

55C Set the ambient temperature of the chassis to 55°C

> **NOTE**: Starting in Junos OS Evolved Release 19.1R1, the PTX10003-80C routers do not support configuration of the ambient temperature value 55C.

Required Privilege Level

chassis—To view this statement in the configuration.

chassis-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

- *Managing Power Allocated to PTX5000 FPCs on the Basis of Chassis Ambient Temperature Configuration*
- *Monitoring the Power Consumption of PTX5000 FPCs by Configuring the Ambient Temperature*
- *Understanding How Configuring Ambient Temperature Helps Optimize Power Utilization*

amf

IN THIS SECTION

- Syntax | 2227
- Hierarchy Level | 2227
- Description | 2227
- Options | 2227
- Required Privilege Level | 2228
- Release Information | 2228

Syntax

```
amf name {
<default-amf>;
 ip-address ip-address;
 local-endpoint {
  <initial-tnla-weight-factor initial-tnla-weight-factor>;
  ip-address address;
  <port port-number>;
  }
 }
 node-id node-id;
 <offline>;
 <port port-number>;
 <routing-instance routing-instance>;
 tracking-area name;
}
```

Hierarchy Level

```
[edit services agf]
```

Description

Configure the parameters that the AGF uses to connect to the Access and Mobility Management Function (AMF).

Options

name	Specifies the name of the Access and Mobility Management Function (AMF).
default-amf	(Optional) Identifies the AMF as the default AMF.
ip-address *address*	Specifies the AMF IP address that AGF uses to establish an initial TNLA.
node-id	Specifies the AMF node identifier.

- **Range:** 0 through 255

`offline` (Optional) Disable communication with this AMF.

`port` *`port number`* (Optional) Specifies the port number on the SCTP endpoint. You can specify a port from 0 through 65535. The default is 38412.

- **Default:** 38412
- **Range:** 0 through 65535

`routing-instance` *`routing-instance-name`* (Optional) Binds the SCTP endpoint routing instance to the AMF connectivity.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3R1.

amt (IGMP)

IN THIS SECTION

- Syntax | 2229
- Hierarchy Level | 2229
- Description | 2229
- Required Privilege Level | 2229
- Release Information | 2229

Syntax

```
amt {
    relay {
        defaults {
            (accounting | no-accounting);
            group-policy [ policy-names ];
            query-interval seconds;
            query-response-interval seconds;
            robust-count number;
            ssm-map ssm-map-name;
            version version;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
igmp],
[edit protocols igmp],
[edit routing-instances routing-instance-name protocols igmp]
```

Description

Configure Automatic Multicast Tunneling (AMT) relay attributes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Default IGMP Parameters for AMT Interfaces

amt (Protocols)

IN THIS SECTION

- Syntax | 2230
- Hierarchy Level | 2231
- Description | 2231
- Required Privilege Level | 2231
- Release Information | 2231

Syntax

```
amt {
    relay {
        accounting;
        family {
            inet {
                anycast-prefix ip-prefix</prefix-length>;
                local-address ip-address;
            }
        }

        secret-key-timeout minutes;
        tunnel-limit number;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable Automatic Multicast Tunneling (AMT) on the router or switch. You must also configure the local address and anycast prefix for AMT to function.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

analytics

IN THIS SECTION

- Syntax (EX Series and QFX Series) | 2232
- Syntax (MX Series & PTX Series) | 2235
- Hierarchy Level | 2236

- Description | 2237
- Required Privilege Level | 2237
- Release Information | 2237

Syntax (EX Series and QFX Series)

Junos OS Release 13.2X51-D15 and later:

```
analytics {
    collector {
        local {
            file filename {
                size size;
                files number;
            }
        }
        address ip-address {
            port number {
                transport protocol {
                    export-profile profile-name;
                }
            }
        }
    }
    export-profiles {
        profile-name {
            interface {
                information;
                statistics {
                    queue;
                    traffic;
                }
                status {
                    link;
                    queue;
                    traffic;
                }
            }
```

```
            stream-format format;
            system {
                information;
                status {
                    queue;
                    traffic;
                }
            }
        }
    }
    resource {
        interfaces {
            interface-name {
                resource-profile name;
            }
        }
        system {
            polling-interval {
                queue-monitoring interval;
                traffic-monitoring interval;
            }
            resource-profile name;
        }
    }
    resource-profiles {
        profile-name {
            depth-threshold {
                high number;
                low number;
            }
            latency-threshold {
                high number;
                low number;
            }
            no-queue-monitoring;
            no-traffic-monitoring;
            queue-monitoring;
            traffic-monitoring;
        }
    }
    traceoptions {
        file filename {
            files number;
```

```
            size size;
        }
    }
}
```

Junos OS Release 13.2X50-D15 and 13.2X51-D10 only:

```
analytics {
    interfaces {
        all {
            depth-threshold high number low number;
            latency-threshold high number low number;
            queue-statistics;
            no-queue-statistics;
            traffic-statistics;
            no-traffic-statistics;
        }
        interface-name {
            depth-threshold high number low number;
            latency-threshold high number low number;
            queue-statistics;
            no-queue-statistics;
            traffic-statistics;
            no-traffic-statistics;
        }
    }
    queue-statistics {
        file filename {
            files number-of-files;
            size size;
        }
        interval interval;
    }
    streaming-servers {
        address ip-address {
            port number {
                stream-format format;
                stream-type type
            }
        }
    }
    traceoptions {
```

```
        file filename {
            files number;
            size size;
        }
    }
    traffic-statistics {
        file filename {
            files number-of-files;
            size size;
        }
        interval interval;
    }
}
```

Syntax (MX Series & PTX Series)

```
analytics {
    agent {
        service-agents {
            agent-name {
                inputs {
                    analytics {
                        parameters {
                            generate-tags value;
                            sample-frequency value;
                            sensors file-path;
                        }
                    }
                    input-ipfix {
                        parameters {
                            maximum-connections number;
                            tcp-port port-number;
                            vrf-name name;
                        }
                    }
                    input-jti-ipfix {
                        parameters {
                            record-group group-name {
                                record ipfix-record-name;
                                reporting-interval seconds;
                            }
```

```
                        }
                    }
                }
                outputs {
                    file {
                        parameters {
                            path file-path;
                        }
                    }
                    kafka {
                        parameters {
                            server ip-address;
                            topic topic-name;
                            encoding encoding-type;
                        }
                    }
                    output-ipfix {
                        parameters {
                            collector-address ip-address;
                            collector-ca-certificate file-path;
                            collector-certificate file-path;
                            collector-certificate-key file-path;
                            collector-connection-retry-interval seconds;
                            collector-tcp-port port-number;
                            collector-vrf-name vrf-name;
                        }
                    }
                }
            }
        }
        traceoptions {
            filename filename;
            flag (debug | error | info | trace);
        }
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the network analytics feature that includes monitoring for traffic and queue statistics. The network analytics processes running on the Packet Forwarding Engine and Routing Engine collect and analyze the data, and generate reports that may be saved in log files or sent as streaming data to remote servers.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

- Network Analytics Overview
- *Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data*
- Configuring NTF Agent
- *IPFIX Mediation on the BNG*
- *Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator*
- *Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector*

analyzer

IN THIS SECTION

- Syntax | 2238
- Hierarchy Level | 2240
- Description | 2240

- Default | 2240
- Options | 2240
- Required Privilege Level | 2241
- Release Information | 2241

Syntax

Routing

```
analyzer {
     analyzer-name {
         input {
             egress {
                interface (all | interface-name);
                 vlan (vlan-id | vlan-name);
            }
             ingress {
                 interface (all | interface-name);
                 vlan (vlan-id | vlan-name);
            }
         output {
             interface interface-name;
             ip-address ip-address;
            routing-instance
            vlan (vlan-id | vlan-name);
        }
    }
}
```

Interface

```
analyzer  {
    analyzer-name {
        input  {
            egress {
                bridge-domain bridge-domain-name;
                interface (all | interface-name);
```

```
                routing-instance {
                    instance-name {
                        bridge-domain bridge-domain-name;
                    }
                }
            }
            ingress {
                bridge-domain bridge-domain-name;
                interface  (all | interface-name);
                routing-instance {
                    instance-name {
                        bridge-domain bridge-domain-name;
                    }
                    vlan (vlan-id | vlan-name);
                }
                vlan (vlan-id | vlan-name);
            }
            maximum-packet-length bytes;
            rate number;
        }
        output {
            bridge-domain bridge-domain-name;
            interface interface-name;
            next-hop-group next-hop-group-name;
            routing-instance {
                instance-name {
                    bridge-domain {
                        bridge-domain-name;
                    }
                }
                vlan (vlan-id | vlan-name);
            }
            vlan (vlan-id | vlan-name);
        }
    }
}
```

Hierarchy Level

For platforms without ELS:

```
[edit ethernet-switching-options]
```

For platforms with ELS:

```
[edit forwarding-options]
```

Description

Configure port mirroring. You can create a total of four port-mirroring configurations on the QFX Series, subject to the following limits:

- There can be no more than two configurations that mirror ingress traffic.
- There can be no more than two configurations that mirror egress traffic.

Default

Port mirroring is disabled and Junos OS creates no default analyzers.

Options

`all`—Mirror all the access interfaces. Using this option does not cause the QSFP+ or management interfaces to be mirrored.

> **CAUTION**: Configuring the `all` option in a QFabric system causes all the access interfaces on all the nodes to be mirrored. Be cautious about using this option on a QFabric system.

`analyzer-name`—Name of the analyzer. The name can include as many as 125 characters; must begin with a letter; and can include uppercase letters, lowercase letters, numbers, dashes, and underscores. No other special characters are allowed.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Option `output vlan` added in Junos OS Release 12.1 for the QFX Series.

Option `output ip-address` added in Junos OS Release 12.3 for the QFX Series for non-ELS platforms and added in 14.1X53-D10 for ELS platforms.

Statement introduced for the “Interface Privilege Level” in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Understanding Port Mirroring and Analyzers

Configuring Port Mirroring

Examples: Configuring Port Mirroring for Local Analysis

Example: Configuring Port Mirroring Analyzers for Local Monitoring of Employee Resource Use

Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use

analyzer-address

IN THIS SECTION

- Syntax | 2242
- Hierarchy Level | 2242
- Description | 2242
- Options | 2242
- Required Privilege Level | 2242

Release Information | 2242

Syntax

```
analyzer-address address;
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure an IP address for the packet analyzer that overrides the default value.

Options

address—IP address for packet analyzer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Packet Analyzer

analyzer-id

IN THIS SECTION

- Syntax | 2243
- Hierarchy Level | 2243
- Description | 2243
- Options | 2243
- Required Privilege Level | 2243
- Release Information | 2244

Syntax

```
analyzer-id name;
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure an identifier for the packet analyzer that overrides the default value.

Options

name—Identifier for packet analyzer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Packet Analyzer

anchor-point (Pseudowire Subscriber Interfaces)

IN THIS SECTION

- Syntax | 2244
- Hierarchy Level | 2244
- Description | 2244
- Options | 2246
- Required Privilege Level | 2246
- Release Information | 2246

Syntax

```
anchor-point lt-device;
```

Hierarchy Level

```
[edit interfaces ps0]
```

Description

Specify the anchor-point, a logical tunnel (lt) interface that identifies the logical tunnel interface that terminates the MPLS pseudowire tunnel at the access node. The other end of the tunnel terminates on

the pseudowire subscriber logical interface, which is configured on an MX Series router that hosts subscriber management and enables you to perform subscriber management services at the interface.

The pseudowire is a virtual device that is stacked above the logical tunnel anchor point on the physical interface (the IFD), and supports a circuit-oriented Layer 2 protocol (either Layer 2 VPN or Layer 2 circuit). The Layer 2 protocol provides the transport and service logical interfaces, and supports the protocol family (IPv4, IPv6, or PPPoE).

NOTE: You cannot dynamically change an anchor point that has active pseudowire devices stacked above it. If you need to change such an anchor point, you must perform the following steps:

1. Deactivate the stacked pseudowires and commit. This may require bringing down any subscribers using the pseudowires.

```
[edit interfaces]
user@host# deactivate psnumber
user@host# commit
```

2. Change the anchor on the deactivated pseudowire and commit.

```
[edit interfaces]
user@host# set psnumber anchor-point lt-new-lt-interface-number
user@host# commit
```

3. Reactivate the stacked pseudowires and commit.

```
[edit interfaces]
user@host# activate psnumber
user@host# commit
```

NOTE: You cannot disable the underlying logical tunnel (lt) interface or redundant logical tunnel (rlt) interface when a pseudowire is anchored on that interface. If you want to disable the underlying interface, you must first deactivate the pseudowire.

1. Deactivate the stacked pseudowires and commit. This may require bringing down any subscribers using the pseudowires.

```
[edit interfaces]
user@host# deactivate psnumber
user@host# commit
```

2. Disable the underlying interface and commit.

```
[edit interfaces]
user@host# set interfaces underlying-interface-name disable
user@host# commit
```

Options

lt-device An lt device in the format `lt-`*fpc/pic/port*

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

ancp

IN THIS SECTION

- Syntax | 2247
- Hierarchy Level | 2248
- Description | 2248
- Required Privilege Level | 2249
- Release Information | 2249

Syntax

```
ancp {
    adjacency-loss-hold-time seconds;
    adjacency-timer seconds;
    gsmp-syn-timeout seconds;
    gsmp-syn-wait;
    interfaces {
        interface-set interface-set-name {
            access-identifier identifier-string;
            underlying-interface underlying-interface-name;
        }
        interface-name {
            access-identifier identifier-string;
        }
    }
    maximum-discovery-table-entries entry-number;
    maximum-helper-restart-time;
    neighbor ip-address {
        adjacency-loss-hold-time seconds;
        adjacency-timer;
        auto-configure-trigger interface interface-name;
        ietf-mode;
        maximum-discovery-table-entries entry-number;
        pre-ietf-mode;
    }
    pre-ietf-mode;
```

```
    qos-adjust {
        adsl-bytes bytes;
        adsl2-bytes bytes;
        adsl2-plus-bytes bytes;
        other-bytes bytes;
        other-overhead-adjust percentage;
        sdsl-bytes bytes;
        sdsl-overhead-adjust percentage;
        vdsl-bytes bytes;
        vdsl-overhead-adjust percentage;
        vdsl2-bytes bytes;
        vdsl2-overhead-adjust percentage;
    }
    qos-adjust-adsl adjustment-factor;
    qos-adjust-adsl2 adjustment-factor;
    qos-adjust-adsl2-plus adjustment-factor;
    qos-adjust-other adjustment-factor;
    qos-adjust-sdsl adjustment-factor;
    qos-adjust-vdsl adjustment-factor;
    qos-adjust-vdsl2 adjustment-factor;
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure Junos OS ANCP agent features.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

> **NOTE**: When you deactivate the ANCP protocol, the router does not perform a commit check to determine whether any ANCP or L2-BSA subscribers are present (active or inactive). Any subscribers that are active at the time of deactivation remain active.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent

ancp (Adjustment Control Profiles)

IN THIS SECTION

- Syntax | 2250
- Hierarchy Level | 2250
- Description | 2250
- Options | 2250
- Required Privilege Level | 2251
- Release Information | 2251

Syntax

```
ancp {
    priority priority;
    algorithm algorithm;
}
```

Hierarchy Level

```
[edit class-of-service adjustment-control-profiles profile-name application]
```

Description

Configure the shaping rate adjustment controls for the ANCP application.

Options

priority Priority of the ANCP application in the adjustment control profile.

- **Range:** 1 through 10; 1 being the highest priority.
- **Default:** 1

algorithm Rate adjustment algorithm used by the ANCP application.

- Values:
  - adjust-never—Do not perform rate adjustments.
  - adjust-always—Adjust the shaping rate unconditionally.
  - adjust-less—Adjust the shaping rate if it is less than the configured value.
  - adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
  - adjust-greater—Adjust the shaping rate if it is greater than the configured value.
  - adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.
- **Default:** adjust-always

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

- *CoS Adjustment Control Profiles Overview*
- *Configuring CoS Adjustment Control Profiles*
- *Verifying the CoS Adjustment Control Profile Configuration*
- adjustment-control-profiles | **1961**
- application (Adjustment Control Profiles) | **2301**

ancp-speed-change-immediate-update (ANCP)

IN THIS SECTION

- Syntax | **2251**
- Hierarchy Level | **2252**
- Description | **2252**
- Required Privilege Level | **2252**
- Release Information | **2252**

Syntax

```
ancp-speed-change-immediate-update;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure AAA to generate immediate interim accounting updates to the RADIUS server in response to ANCP agent notifications of rate changes on subscriber access lines.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications

Configuring Per-Subscriber Session Accounting

Configuring the ANCP Agent

annotate

IN THIS SECTION

- Syntax | 2253
- Description | 2253
- Options | 2253
- Required Privilege Level | 2253
- Release Information | 2254

Syntax

```
annotate <statement> <comment string>
```

Description

Add comments to a configuration. You can add comments only at the current hierarchy level.

Any comments you add appear only when you view the configuration by entering the "show" on page 16994 command in configuration mode or the `show configuration` command in operational mode.

NOTE: The software supports annotation up to the last level in the configuration hierarchy, including oneliners. However, annotation of parts (child statements or identifiers within a oneliner) of the oneliner is not supported. For example, in the following sample configuration hierarchy, annotation is supported up to the oneliner level 1 , but not supported for the metric child statement and its attribute *10*:

```
[edit protocols]
    isis {
        interface ge-0/0/0.0 {
            level 1 metric 10;
        }
    }
}
```

Options

statement Statement to which you are attaching the comment.

comment-string Text of the comment. You must enclose it in quotation marks. In the comment string, you can include the comment delimiters `/* */` or `#`. If you do not specify any, the comment string is enclosed with the `/* */` comment delimiters. If a comment for the specified *statement* already exists, it is deleted and replaced with the new comment.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Adding Comments in a Device Configuration

anomaly

IN THIS SECTION

- Syntax | 2254
- Hierarchy Level | 2254
- Description | 2255
- Options | 2255
- Required Privilege Level | 2255
- Release Information | 2255

Syntax

```
anomaly {
    direction (Security Custom Attack) (any | client-to-server | server-to-client);
    service service-name;
    shellcode (all | intel | no-shellcode | sparc);
    test test-condition;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type]
```

Description

Protocol anomaly attack objects detect abnormal or ambiguous messages within a connection according to the set of rules for the particular protocol being used.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

anti-replay-window-size

IN THIS SECTION

- Syntax | 2255
- Hierarchy Level | 2256
- Description | 2256
- Required Privilege Level | 2256
- Release Information | 2256

Syntax

```
anti-replay-window-size anti-replay-window-size;
```

Hierarchy Level

```
[edit security ipsec],
[edit tenants name security ipsec]
```

Description

To enable the `anti-replay-window-size` option, you first need to configure the option for each VPN object or at the global level. You can configure the anti-replay window size in the range of 64 to 8192 (power of 2). If the anti-replay window size is not configured, the window size is 64 by default. If `anti-replay-window-size` command is configured at both the global and VPN object levels, the configuration on VPN object takes precedence over global configuration.

Note that you can increase the anti-replay window size if your device runs the IPsec VPN service using the IKED process. Devices running the KMD process do not support this feature.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Anti-Replay Window

anti-replay-window-size (Services IPsec VPN)

IN THIS SECTION

Syntax | 2257

- Hierarchy Level | 2257
- Description | 2257
- Options | 2257
- Required Privilege Level | 2257
- Release Information | 2257

Syntax

```
anti-replay-window-size bits;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Specify the size of the IPsec antireplay window.

Options

bits—Size of the antireplay window, in bits.

- **Default:** 64 bits (AS PICs), 128 bits (Multiservices PICs and DPCs)
- **Range:** 64 through 4096 bits

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

anti-replay-window-size (Services Service Set)

IN THIS SECTION

- Syntax | 2258
- Hierarchy Level | 2258
- Description | 2258
- Options | 2259
- Required Privilege Level | 2259
- Release Information | 2259

Syntax

```
anti-replay-window-size bits;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Specify the size of the IPsec antireplay window. This statement is useful for dynamic endpoint tunnels for which you cannot configure the `anti-replay-window-size` statement at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level.

For static IPsec tunnels, this statement sets the antireplay window size for all the static tunnels within this service set. If a particular tunnel needs a specific value for antireplay window size, set the `anti-replay-window-size` statement at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level. If antireplay check has to be disabled for a particular tunnel in this service set, set the `no-anti-replay` statement at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level.

> **NOTE**: The `anti-replay-window-size` and `no-anti-replay` settings at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level override the settings specified at the `[edit services service-set service-set-name ipsec-vpn-options]` hierarchy level.

Options

bits—Size of the antireplay window, in bits.

- **Default:** 64 bits (AS PICs), 128 bits (Multiservices PICs and DPCs)
- **Range:** 64 through 4096 bits

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring IPsec Service Sets

Configuring IPsec Rules

anti-spam (Logical Systems)

IN THIS SECTION

- Syntax | 2260
- Hierarchy Level | 2260
- Description | 2260

- Options | 2261
- Required Privilege Level | 2261
- Release Information | 2261

Syntax

```
anti-spam {
    sbl {
        profile (Security Antispam SBL) name {
            address-blacklist address-blacklist;
            address-whitelist address-whitelist;
            custom-tag-string custom-tag-string;
            (sbl-default-server | no-sbl-default-server);
            spam-action (block | tag-header | tag-subject);
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name security utm feature-profile]
[edit tenants tenant-name security utm feature-profile]
```

Description

Configures the Content Security antispam feature for logical systems. The antispam feature examines transmitted e-mail messages to identify e-mail spam. When the device detects a message deemed to be spam, it blocks the e-mail message or tags the e-mail message header or subject with a preprogrammed string. Antispam filtering uses both a third-party server-based Spam Block List (SBL), and optionally created local allowlists (benign) and blocklists (malicious) for filtering against e-mail messages.

You can also configure the default Content Security configuration for antispam feature profile. If you do not configure any option in the antispam feature profile, the values configured in the default Content Security configuration are applied. In the default Content Security profile, the antispam type is configured as SBL instead of none. This configuration enables the SBL option. However, to use this

feature, enable the SBL server using the `[edit security utm default-configuration anti-spam sbl sbl-default-server]` command.

> **NOTE**: A license check for the antispam configuration is performed at the time of commit and provides a warning if a valid license is not installed on the device. Once a valid license is installed on the device then the custom antispam profile or the default antispam profile is able to process the traffic. If a license is expired or is not installed, the antispam service does not process the traffic.

Options

anti-spam	Configures the Content Security antispam feature for logical system.
address-blacklist	Enter an address blocklist custom object for local list spam filtering.
address-whitelist	Enter an address allowlist custom object for local list spam filtering.
sbl	Antispam filtering allows you to use both a third-party server-based SBL, and optionally created local allowlists and blocklists for filtering against e-mail messages.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Antispam Filtering Overview

utm default-configuration | **30614**

anti-spam

IN THIS SECTION

- Syntax | 2262
- Hierarchy Level | 2263
- Description | 2263
- Options | 2264
- Required Privilege Level | 2264
- Release Information | 2264

Syntax

```
anti-spam {
 address-blacklist address-blacklist;
 address-whitelist address-whitelist;
 sbl {
  custom-tag-string custom-tag-string;
  (sbl-default-server | no-sbl-default-server);
  server {
    ip;
   routing-instance routing-instance;
   source-address source-address;
  }
  spam-action(block | tag-header | tag-subject);
 }
 traceoptions {
  flag name;
 }
 type(anti-spam-none | sbl);
}
```

Hierarchy Level

```
[edit security utm feature-profile]
[edit security utm default-configuration]
```

Description

Configure Content Security antispam features. You can also configure the default Content Security configuration for antispam feature profile. If you do not configure any option in the antispam feature profile, the values configured in the default Content Security configuration are applied.

The antispam feature examines transmitted e-mail messages to identify e-mail spam. When the device detects a message deemed to be spam, it blocks the e-mail message or tags the e-mail message header or subject with a preprogrammed string. Antispam filtering uses both a third-party server-based Spam Block List (SBL) and optionally created local allowlists (benign) and blocklists (malicious) for filtering against e-mail messages.

> **NOTE**: A license check for the antispam configuration is performed at the time of a commit and will provide a warning if a valid license is not installed on the device. Once a valid license is installed on the device then a custom antispam profile or the default profile will be able to process traffic. If a license is expired or is not installed, the antivirus service will not process traffic.

In the default Content Security profile, the antispam type is configured as SBL instead of none. This configuration enables SBL. However, to use this feature, you must enable the SBL server using the `[edit security utm default-configuration anti-spam sbl sbl-default-server]` command.

Starting in Junos OS Release 23.1R1, we've introduced the `server` statement at the `[edit security utm default-configuration anti-spam sbl]` hierarchy level to configure the SBL server details for antispam. Earlier to Junos OS Release 23.1R1, the antispam and antivirus used the same `server` statement at [edit security Content Security default-configuration anti-virus sophos-engine] hierarchy level for the SBL server configuration. For antispam, it is recommended to configure SBL server at the `[edit security utm default-configuration anti-spam sbl server]` hierarchy level rather than at the `[edit security utm default-configuration anti-virus sophos-engine server]` hierarchy level.

Starting in Junos OS Release 23.1R1, content security supports antivirus Sophos Live Protection version 2.0. The new version of Sophos antivirus uses HTTPS connection instead of DNS transport channel for device to server communication. Hence, we've deprecated the `ip` statement at the `[edit security utm default-configuration anti-virus sophos-engine server]` hierarchy level.

Options

anti-spam	Configure antispam feature.
address-blacklist	Enter an address blocklist custom object for local list spam filtering.
address-whitelist	Enter an address-allowlist custom-object for local list spam filtering.
ip	Specify the antispam first-hop DNS server IP address.
routing-instance	Specify the name of the routing instance.
sbl	Antispam filtering allows you to use both a third-party server-based spam block list (SBL) and to optionally create your own local allowlists and blocklists for filtering against e-mail messages.
traceoptions	Defines tracing operations for Content Security antispam features.
source-address	Specify the source IP address used to connect SBL server.
type	Antispam type.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

The `server` option introduced in Junos OS Release 23.1R1.

anti-spam (Security Content Security Policy)

IN THIS SECTION

- Syntax | 2265
- Hierarchy Level | 2265
- Description | 2265
- Options | 2266
- Required Privilege Level | 2266
- Release Information | 2266

Syntax

```
anti-spam {
    smtp-profile profile-name;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name]
[edit logical-systems logical-system-name security utm utm-policy policy-name]
[edit tenants tenant-name security utm utm-policy policy-name]
```

Description

Configures a Content Security policy for the antispam SMTP protocol and attach this policy to a security profile to implement it. The device can block and drop detected spam at either the connection level or the e-mail level. When the SMTP sender is identified as a spam sender based on its IP address, the SMTP connection is rejected and dropped. When a particular e-mail sender is identified as spam sender based on its sender address, the e-mail is rejected and dropped.

Options

The statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for configuration in logical systems introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Antispam Filtering Overview

anti-virus (Logical Systems)

IN THIS SECTION

- Syntax | 2267
- Hierarchy Level | 2268
- Description | 2268
- Options | 2268
- Required Privilege Level | 2269
- Release Information | 2269

Syntax

```
anti-virus {
    profile name {
        fallback-options (Security Antivirus Sophos Engine) {
            content-size (Security Antivirus Sophos Engine) (block | log-and-permit | permit);
            decompress-err (block | log-and-permit | permit);
            default (Security Antivirus Sophos Engine) (block | log-and-permit | permit);
            engine-not-ready (Security Antivirus Sophos Engine) (block | log-and-permit | permit);
            out-of-resources (Security Antivirus Sophos Engine) (block | log-and-permit | permit);
            timeout (Security Antivirus Fallback Options Sophos Engine) (block | log-and-permit |
permit);
            too-many-requests (Security Antivirus Fallback Options Sophos Engine) (block | log-and-
permit | permit);
        }
        mime-whitelist {
            exception exception;
            list list;
        }
        notification-options (Security Antivirus) {
            fallback-block (Security Antivirus) {
                custom-message (Security Fallback Block) custom-message;
                custom-message-subject (Security Fallback Block) custom-message-subject;
                (notify-mail-sender (Security Fallback Block) | no-notify-mail-sender (Security Fallback
Block));
                type (Security Fallback Block) (message | protocol-only);
            }
            fallback-non-block (Security Antivirus) {
                custom-message (Security Fallback Non-Block) custom-message;
                custom-message-subject (Security Fallback Non-Block) custom-message-subject;
                (notify-mail-recipient | no-notify-mail-recipient);
            }
            virus-detection (Security Antivirus) {
                custom-message (Security Virus Detection) custom-message;
                custom-message-subject (Security Virus Detection) custom-message-subject;
                (notify-mail-sender (Security Virus Detection) | no-notify-mail-sender (Security Virus
Detection));
                type (Security Virus Detection) (message | protocol-only);
            }
        }
        url-whitelist url-whitelist;
```

```
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-systems-name security utm feature-profile]
[edit tenants tenant-name security utm feature-profile]
```

Description

Configures the Content Security Sophos Antivirus feature for logical systems. You can also configure the default Content Security configuration for antivirus feature profile. If you do not configure any option in the antivirus feature profile, the values configured in the default Content Security configuration are applied.

> **NOTE**: A license check for the antivirus configuration is performed at the time of a commit and will provide a warning if a valid license is not installed on the device. Once a valid license is installed on the device then a custom antivirus profile or the antivirus default profile is able to process traffic. If a license is expired or is not installed, the antivirus service does not process the traffic.

Options

anti-virus	Configures the Content Security antivirus feature for logical systems.
mime-whitelist	This is the comprehensive list for those MIME types that can bypass antivirus scanning.
sophos-engine	The antivirus engine that is used on the device. You can only have one engine type running and you must restart the device if you change engines.
fallback-options	Fallback options helps the system how to handle the errors.
notification-options	There are multiple notification options you can configure to trigger when a virus is detected.
fallback-non-block	Notifications for fallback nonblocking actions.
virus-detection	Notifications to send when a virus is detected.

`scan-options`	Antivirus sophos-engine scan options.
`trickling`	HTTP trickling is a mechanism used to prevent the HTTP client or server from timing-out during a file transfer or during antivirus scanning.
`url-whitelist`	Antivirus URL allowlist is a unique custom list that includes the URLs or IP addresses category to bypass the antivirus scanning.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Unified Threat Management Overview

utm default-configuration | 30614

anti-virus (Content Security)

IN THIS SECTION

- Syntax | 2270
- Hierarchy Level | 2271
- Description | 2272
- Options | 2272
- Required Privilege Level | 2273

Release Information | 2273

Syntax

```
anti-virus {
    mime-whitelist {
        exception;
        list;
    }
    sophos-engine {
         fallback-options {
            content-size (block | log-and-permit | permit);
            default (block | log-and-permit | permit);
            engine-not-ready (block | log-and-permit | permit);
            out-of-resources (block | log-and-permit | permit);
            timeout (block | log-and-permit | permit);
            too-many-requests (block | log-and-permit | permit);
        }
        notification-options {
            fallback-block {
                custom-message;
                custom-message-subject;
                (notify-mail-sender | no-notify-mail-sender);
                type (message | protocol-only);
            }
            fallback-non-block {
                custom-message;
                custom-message-subject;
                (notify-mail-recipient | no-notify-mail-recipient);
            }
            virus-detection {
                custom-message;
                custom-message-subject;
                (notify-mail-sender | no-notify-mail-sender);
                type (message | protocol-only);
            }
        }
        pattern-update {
            email-notify {
```

```
                admin-email;
                custom-message;
                custom-message-subject;
            }
            interval;
            no-autoupdate;
            proxy {
                password;
                 port;
                server;
                username;
            }
            routing-instance;
            url;
        }
        scan-options {
            content-size-limit;
            timeout seconds;
            (uri-check | no-uri-check);
        }
        server {
            ip;
            routing-instance;
        }
        sxl-retry;
        sxl-timeout seconds;
        trickling timeout;
    }
    traceoptions {
        flag name;
    }
    url-whitelist;
}
```

Hierarchy Level

```
[edit security utm feature-profile]
[edit security utm default-configuration]
```

Description

Configure Content Security Sophos antivirus features. You can also configure the default Content Security configuration for antivirus feature profile. If you do not configure any option in the antivirus feature profile, the values configured in the default Content Security configuration are applied. Antivirus, one of several features including content filtering, antispam, and Web filtering, makes up Juniper's Content Security suite, provides the ability to prevent threats at the gateway before they enter the network.

> **NOTE**: A license check for the antivirus configuration is performed at the time of a commit and will provide a warning if a valid license is not installed on the device. Once a valid license is installed on the device then a custom antivirus profile or the default profile will be able to process traffic. If a license is expired or is not installed, the antivirus service will not process traffic.

Options

`anti-virus`	Configure antivirus feature.
`mime-whitelist`	This is the comprehensive list for those MIME types that can bypass antivirus scanning.
`sophos-engine`	The antivirus engine that is used on the device. You can only have one engine type running and you must restart the device if you change engines.
`fallback-options`	Fallback options tell the system how to handle the errors.
`notification-options`	There are multiple notification options you can configure to trigger when a virus is detected.
`fallback-non-block`	Notifications for fallback nonblocking actions.
`virus-detection`	Notification to send when a virus is detected.
`pattern-update`	You can configure the security device to regularly update the pattern file automatically, or you can update the file manually.
`scan-options`	Antivirus sophos-engine scan options.
`server`	Sophos Antivirus (SAV) and antispam first hop DNS server.
`sxl-retry`	Number of retry attempts to the remote Sophos Extensible List (SXL) server when a request timeout occurs.

- **Range:** 0 through 5

`sxl-timeout` Timeout value for responses to a Sophos checksum or URI query.

- **Range:** 1 through 5

`trickling` HTTP trickling is a mechanism used to prevent the HTTP client or server from timing-out during a file transfer or during antivirus scanning.

`traceoptions` Define tracing operations for Content Security antivirus features.

`url-whitelist` Antivirus URL allowlist. A URL allowlist is a unique custom list that you define in which all the URLs or IP addresses in that list for a specified category are always bypassed for scanning.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

anti-virus (Content Security Policy)

IN THIS SECTION

- Syntax | 2274
- Hierarchy Level | 2274
- Description | 2274
- Options | 2274

Required Privilege Level | 2275
Release Information | 2275

Syntax

```
anti-virus {
    ftp {
        download-profile profile-name;
        upload-profile profile-name;
    }
    http-profile profile-name;
    imap-profile profile-name;
    pop3-profile profile-name;
    smtp-profile profile-name;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name]
[edit logical-systems logical-systems-name security utm utm-policy policy-name]
[edit tenants tenant-name security utm utm-policy policy-name]
```

Description

Configures a Content Security policy for the antivirus protocols and attaches this policy to a security profile to implement it. The internal antivirus scan engine supports scanning for specific Application Layer transactions allowing you to select the content (HTTP, FTP, SMTP, POP3, or IMAP traffic) to scan.

Options

The statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for configuration in logical systems introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Unified Threat Management Overview

anti-virus (Services)

IN THIS SECTION

- Syntax | 2275
- Hierarchy Level | 2277
- Description | 2277
- Options | 2277
- Required Privilege Level | 2280
- Release Information | 2280

Syntax

```
anti-virus {
 policy policy-name {
```

```
 action(block | permit);
 default-notification {
  log;
 }
 fallback-options {
  action(block | permit);
  invalid-content-size {
   action(block | permit);
   notification {
    log;
   }
  }
  notification {
   log;
  }
  out-of-resources {
   action(block | permit);
   notification {
    log;
   }
  }
  service-not-ready {
   action(block | permit);
   notification {
    log;
   }
  }
 }
 http-client-notify {
  (file file-name | message message | redirect-url redirect-URL);
 }
 notification {
  log;
 }
 machine-learning-scan {
  action (block | permit);
  notification {
    log;
  }
 }
 verdict-threshold(1 | 10 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9);
}
traceoptions {
```

```
  (file file-name | files files | match match | size size |(world-readable | no-world-readable));
  flag name;

  level(all | error | info | notice | verbose | warning);
  no-remote-trace;
 }
 update {
  automatic {
   interval interval;
   off;
  }
  ignore-server-validation;
  proxy-profile proxy-profile;
  url URL;
 }
 }
```

Hierarchy Level

```
[edit services]
```

Description

Use this command to configure flow-based antivirus policy and machine learning scan. After configuring the antivirus policy, you must apply it to the network firewall policy using the `set security policies from-zone from-zone to-zone to-zone policy policy-name then permit application-services anti-virus-policy av-policy` command.

Options

policy Configure the antivirus policy details.

- policy-name—Name of the antivirus policy.
- action—Action taken when file meets verdict threshold.
- default-notification—Notification log when file does not meet the verdict threshold.

fallback-options Defines what to do when error conditions occur or when there is a lack of resources. The following fallback options are available:

- action—Permit or block the file regardless of its threat level.
- invalid-content-size—Permit or block the file when content size exceed supported range
- out-of-resources—Permit or block the file when service is out of resources
- service-not-ready—Permit or block the file when service is not yet ready
- notification—Add or do not add this event to the log file.

http-client-notify

(message | file | redirect-url) Notification action taken for contents with verdict meet threshold. This command allows you to configure HTTP URL redirection for a customized client notification based on detected virus with the block action.

machine-learning-scan

Configure the machine learning scan options.

- action—Permit or block the file regardless of its threat level.
- apply-groups—Groups that these machine learning scan settings apply.
- apply-groups-except—Groups that these machine learning scan settings do not apply.
- default-notification—Notification log when file does not meet the verdict threshold.
- notification—Add or do not add this event to the log file.

verdict-threshold

The verdict-threshold defines the number at which you want to label a file as virus. For example, if you set verdict-threshold to 7, then any file with verdict number of 7 or greater is considered virus. verdict-threshold can be any number between 1 and 10, inclusive.

traceoptions

Configure anti-virus trace options. When the trace is set in the configuration, the “flags” are defined for the actual debug that you want to perform. This command allows you to trace the antivirus configuration and is typically used for troubleshooting.

- apply-groups name—Groups that these trace option settings apply.
- apply-groups-except name—Groups that these trace option settings do not apply.
- file— Defines the trace file characteristics. Valid options are:
  - trace-file-name—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. Valid values range from 1 and 1024 characters. The name cannot include spaces, /, or % characters. The default filename is security.

  - files—Maximum number of trace files that can accumulate. Valid values range from 2 to 1000. The default value is 3.
  - match—The criteria that you want the system to use when logging information to the file. You can enter a regular expression. Wildcard (*) characters are accepted.
  - no-world-readable—Only the system administrator can read the trace file.
  - size—Maximum size to which the trace file can grow. Once the file reaches the specified size, it is compressed and renamed filename0.gz, the next file is named filename1.gz, and so on. Valid values range from 10240 to 1,073,741,824.
  - world-readable—Any user can read the trace file.
- flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:
  - all—Trace everything.
  - content—Trace the content the SRX is buffering to the client and to the cloud.
  - daemon—Trace the Juniper Advanced Threat Prevention Cloud (ATP Cloud) daemon.
  - http—Trace HTTP protocol operations
  - identification—Trace the file type identification. Examples of file types are .exes, .java, .tar and so forth.
  - imap—Trace IMAP protocol operations
  - parser—Trace the interface between the Juniper ATP Cloud daemon and the module that parses the HTTP protocol and extracts the file content.
  - plugin—Trace the Juniper ATP Cloud plugin to view session checks, packet processing and actions taken against a file or URL.
  - policy—Trace the Juniper ATP Cloud security policy.
  - smb—Trace SMB protocol operations
  - smtp—Trace SMTP protocol operations
- level—Specify level of tracing to perform. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- all—Match messages of all levels.
- error—Match error conditions.
- info—Match informational messages.
- notice—Match notice messages about conditions requiring special handling.
- verbose—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure verbose, messages at all higher levels are traced. Therefore, the result is the same as when you configure all.
- warning—Match warning messages.

- no-remote-trace—Disable remote tracing.

update Configure antivirus package download and install options.

- automatic— Scheduled download and install.
- interval—Schedule automatic download interval in minutes . Range is 5 to 60 minutes.
- off—Disable automatic antivirus download and install.
- ignore-server-validation—Do not check server certificate.
- proxy profile—Configure proxy profile.
- url—HTTP/HTTPS URL of antivirus package download.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

The `machine-learning-scan` option is introduced in Junos OS Release 24.2R1.

any-sender

IN THIS SECTION

- Syntax | 2281
- Hierarchy Level | 2281
- Description | 2281
- Required Privilege Level | 2281
- Release Information | 2282

Syntax

```
any-sender;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

Disable strict sender address checks.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

any (VPLS Multihoming)

IN THIS SECTION

- Syntax | 2282
- Hierarchy Level | 2282
- Description | 2283
- Default | 2283
- Required Privilege Level | 2283
- Release Information | 2283

Syntax

```
any;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-
homing active-interface],
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-
homing site site-name active-interface],
[edit routing-instances instance-name protocols vpls multi-homing active-interface],
[edit routing-instances instance-name protocols vpls multi-homing site site-name active-
interface]
```

Description

Specify that any multihomed interface can be used as the primary interface by the VPLS site. Depending on the order in which interfaces are listed in the PE router's configuration, the first operational interface in the set of configured interfaces is chosen to be the primary interface.

> **NOTE**: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

For FEC 128, use the `[edit routing-instances instance-name protocols vpls multi-homing active-interface]` hierarchy level.

Default

This is the default behavior.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Specifying an Interface as the Active Interface

primary (VPLS Multihoming) | **13151**

anycast-pim

IN THIS SECTION

Syntax | **2284**

- Hierarchy Level | 2284
- Description | 2284
- Required Privilege Level | 2284
- Release Information | 2284

Syntax

```
anycast-pim {
    rp-set {
        address address <forward-msdp-sa>;
    }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp local family (inet | inet6)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp local family (inet | inet6)],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim rp local family (inet | inet6)]
```

Description

Configure properties for anycast RP using PIM.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring PIM Anycast With or Without MSDP

anycast-prefix

IN THIS SECTION

- Syntax | 2285
- Hierarchy Level | 2285
- Description | 2285
- Default | 2286
- Options | 2286
- Required Privilege Level | 2286
- Release Information | 2286

Syntax

```
anycast-prefix ip-prefix/<prefix-length>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols amt relay family inet],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols amt
relay family inet],
[edit protocols amt relay family inet],
[edit routing-instances routing-instance-name protocols amt relay family inet]
```

Description

Specify an IP address prefix to use for the Automatic Multicast Tunneling (AMT) relay anycast address. The prefix is advertised by unicast routing protocols to route AMT discovery messages to the router from nearby AMT gateways. The IP address that the prefix is derived from can be configured on any

interface in the system. Typically, the router's **lo0.0** loopback address prefix is used for configuring the AMT anycast prefix in the default routing instance, and the router's **lo0.*n*** loopback address prefix is used for configuring the AMT anycast prefix in VPN routing instances. However, the anycast address can be either the primary or secondary **lo0.0** loopback address.

Default

None. The anycast prefix must be configured.

Options

ip-prefix/<prefix-length>—IP address prefix.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

apn

IN THIS SECTION

- Syntax | 2287
- Hierarchy Level | 2287
- Description | 2287
- Options | 2287
- Required Privilege Level | 2288

Release Information | 2288

Syntax

```
apn pattern-string {
    imsi-prefix imsi-prefix-digits {
        action (APN GTP) {
            drop;
            pass;
            selection (ms|net|vrf);
        }
    }
}
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

Use the APN function to allow or deny access to specific access point names (APNs). By default, the device permits all APNs. To enable APN filtering, you must specify one or more APNs. To specify an APN, you need to know the domain name of the network (for example, example.com) and, the operator ID. Because the domain name (network ID) portion of an APN can potentially be very long and contain many characters, you can use the wildcard (*) as the first character of the APN string.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name]` hierarchy level is replaced by `[edit security gtp profile profile-name]`.

Options

- *pattern-string*—To specify APN pattern string, such as “example.net.mcc123.mnc456.gprs”.
- *imsi-prefix-digits*—To specify an IMSI prefix.

  The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. Support.

RELATED DOCUMENTATION

Understanding GTP APN Filtering

Example: Setting a GTP APN and a Selection Mode

Understanding GTPv2 IMSI Prefix and APN Filtering

apn-control-group (Security GTP)

IN THIS SECTION

- Syntax | 2288
- Hierarchy Level | 2288
- Description | 2289
- Release Information | 2289

Syntax

```
apn-control-group <apn-ctrl-grp-name>
```

Hierarchy Level

```
[edit security gtp]
```

Description

Use this configuration to group all the APNs with the same rate limiting requirements and apply the group to the GTP profile.

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Understanding GTP APN Filtering

apn-control (Security GTP)

IN THIS SECTION

- Syntax | 2289
- Hierarchy Level | 2290
- Description | 2290
- Options | 2290
- Release Information | 2291

Syntax

```
apn-control  apn-ctrl-name
{ apn apn-string}
{rate-limit
  {alarm-threshold threshold}
   {drop-threshold threshold}
    { message <v0msg-list-name|v1msg-list-name|v3msg-list-name>}
}
```

Hierarchy Level

```
[edit security gtp]
```

Description

Use this command to configure APN control profile using APN string. You can apply the APN control profile to GTP profile for ARL.

Options

- `Apn` *apn-string*—You can specify APN using APN string. APN string consists of Network Identifier (NI) and Operator Identifier (OI). You can specify APN in the below formats:
  - <NI>: for example internet, internet.public
  - <NI>.<OI>: for example internet.mnc005.mcc244.gprs, internet.public.mnc005.mcc244.gprs

  APN string configuration supports wild card "*" in place of OI. You can specify APN as <NI>."*" and configure the matching APN string. You don't need to specify the OI when you configure APN using wild card "*".

  Range–The length range of the apn-string is 1 through 100.

- `rate-limit`—Use the rate-limit option to limit the GTP messages per second. Once the threshold is reached, this configuration either drops the messages or raises an alarm.
  - `alarm-threshold` *threshold*–This configuration raises an alarm when the apn-messages per second reaches the threshold.

    Range:–1 through 80000

  - `drop-threshold`–This configuration drops the apn-messages when the apn-messages per second reaches the threshold.

    Range:–1 through 80000

  - `message <v0|v1|v2>`–This configuration rate-limits the messages using message type specified in the message-list. You can configure any message type in a message list. While applying the message-list to APN control profile, make sure that the message type configured in the list is a GTP-C message which is related to a GTP session context.

    If the message-lists are not configured, aggregated rate limit uses the default messages.

Table 2: Default aggregated rate limit messages for various GTP versions.

GTP version	Default aggregated rate limit messages
GTPv0	Create PDP Request
	Create AA PDP Request
	Update PDP Request
	Delete PDP Request
	Delete AA PDP Request
GTPv1	Create PDP Request
	Update PDP Request
	Delete PDP Request
GTPv2	Create Session Request
	Create Bearer Request
	Modify Bearer Request
	Delete Session Request
	Delete Bearer Request

Release Information

Statement introduced in Junos OS Release 21.2R1.

apn-services (control plane services)

IN THIS SECTION

- Syntax | 2292
- Hierarchy Level | 2292
- Description | 2292
- Options | 2292
- Required Privilege Level | 2293
- Release Information | 2293

Syntax

```
apn-services {
    apns name {
        mobile-interface mobile-interface;
    }
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw name control-plane-peers]
```

Description

The access point name (APN) is sent by the control plane over PFCP to place a subscriber in a specific network instance.

Options

apn name At least one access point name.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.4R1.

app-mapping-timeout

IN THIS SECTION

- Syntax | 2293
- Hierarchy Level | 2293
- Description | 2293
- Options | 2294
- Required Privilege Level | 2294
- Release Information | 2294

Syntax

```
app-mapping-timeout app-mapping-timeout;
```

Hierarchy Level

```
[edit services nat pool nat-pool-name]
```

Description

Specify the duration for address pooling paired (AP-P) mappings that use the specified NAT pool. If this option is not configured and a timeout value is configured for "mapping-timeout" on page 10348, the

timeout value configured for "mapping-timeout" on page 10348 is used. If neither option is specified, the default value of 300 seconds is used.

Options

app-mapping-timeout—Lifetime of AP-P mappings in seconds.

- **Default:** 300
- **Range:** 120 through 864,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

`mapping-timeout` statement introduced in JUNOS Release 12.3.

RELATED DOCUMENTATION

Configuring Source and Destination Addresses Network Address Translation Overview

appfw-profile (System)

IN THIS SECTION

- Syntax | 2295
- Hierarchy Level | 2295
- Description | 2295
- Options | 2295
- Required Privilege Level | 2295
- Release Information | 2296

Syntax

```
appfw-profile {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level

```
[edit security application-firewall profile profile-name]
[edit tenants tenant-name security application-firewall]
```

Description

Specify the application firewall profile quota of a logical system and tenant systems.

As a primary administrator, you can create a security profile and specify the kinds and amounts of resources that are to be allocated to a logical system to which the security profile is bound. A security profile is used for share the device resources, including policies, zones, addresses and address books, flow sessions, and various forms of NAT, among all the logical systems appropriately. You can dedicate various amounts of a resource to the logical systems and also allow the logical systems to compete for use of the free resources.

Options

- `maximum` *`amount`*—Specify the maximum allowed quota value.

  `Range`: 0 through 1024

- `reserved` *`amount`*—Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

The `edit tenant tenant-name security application-firewall` level is introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Application Firewall Overview

appfw-rule

IN THIS SECTION

- Syntax | 2296
- Hierarchy Level | 2296
- Description | 2297
- Options | 2297
- Required Privilege Level | 2297
- Release Information | 2297

Syntax

```
appfw-rule {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile security-profile-name ]
[edit tenants tenant-name security application-firewall ]
```

Description

Specify the number of application firewall rule configurations that a primary administrator can configure for a primary logical system or user logical system, when the security profile is bound to the logical systems and tenant systems.

Tasks performed by the primary administrator are:

- Uses security profiles to provision logical systems with resources.
- Binds security profiles to the primary logical system and the user logical systems.
- Configures more than one security profile, and allocating different numbers of resources in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

- `maximum` *`amount`*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can use resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- `reserved` *`amount`*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

The `edit tenant` *`tenant-name`* `security application-firewall` level is introduced in Junos OS Release 18.4R1.

appfw-rule-set

IN THIS SECTION

- Syntax | 2298
- Hierarchy Level | 2298
- Description | 2298
- Options | 2299
- Required Privilege Level | 2299
- Release Information | 2299

Syntax

```
appfw-rule-set {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile security-profile-name ]
[edit tenants tenant-name security application-firewall]
```

Description

Specify the number of application firewall rule set configurations that a primary administrator can configure for a primary logical system or user logical system when the security profile is bound to the logical systems and tenant systems.

The primary administrator:

- Uses security profiles to provision logical systems with resources
- Binds security profiles to the primary logical system and the user logical systems

- Can configure more than one security profile, allocating different numbers of resources in various profiles

Only the primary administrator can create security profiles and bind them to logical systems.

Options

- `maximum` *`amount`*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can use resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- `reserved` *`amount`*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

The `edit tenant` *`tenant-name`* `security application-firewall` level is introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Application Firewall Overview

application

IN THIS SECTION

Syntax | 2300

- Hierarchy Level | 2300
- Description | 2300
- Options | 2301
- Required Privilege Level | 2301
- Release Information | 2301

Syntax

```
application application-name {
    application-protocol protocol-name;
    child-inactivity-timeout seconds;
    destination-port port-number;
    gate-timeout seconds;
    icmp-code value;
    icmp-type value;
    inactivity-timeout ( seconds | never );
    protocol type;
    rpc-program-number number;
    snmp-command command;
    source-port port-number;
    ttl-threshold number;
    uuid hex-value;
}
```

Hierarchy Level

```
[edit applications],
[edit applications application-set application-set-name]
```

Description

Configure properties of an application and whether to include it in an application set.

Options

application-name—Identifier of the application.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *ALG Descriptions*
- *Configuring Application Sets*
- *Configuring Application Properties*
- *Examples: Configuring Application Protocols*
- *Verifying the Output of ALG Sessions*

application (Adjustment Control Profiles)

IN THIS SECTION

- Syntax | 2302
- Hierarchy Level | 2302
- Description | 2302
- Required Privilege Level | 2302
- Release Information | 2302

Syntax

```
application {
    (ancp | dhcp-tags | pppoe-tags | radius-coa)
        priority priority;
        algorithm algorithm;
    }
}
```

Hierarchy Level

```
[edit class-of-service adjustment-control-profiles profile-name]
```

Description

Configure which applications in the adjustment control profile can make shaping rate adjustments.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

CoS Adjustment Control Profiles Overview

Configuring CoS Adjustment Control Profiles

Verifying the CoS Adjustment Control Profile Configuration

application (Application Identification)

IN THIS SECTION

- Syntax | 2303
- Hierarchy Level | 2304
- Description | 2304
- Options | 2304
- Required Privilege Level | 2305
- Release Information | 2305

Syntax

```
application application-name {
    address-mappingaddress-name {
        filter {
            ip ip-address-and-prefix-length;
            port-range {
                tcp [port];
                udp [port];
            }
        }
    }
    cacheable;
    description;
    icmp-mapping {
        code number;
        type number;
    }
    ip-protocol-mapping {
        protocol number;
    }
    priority high;
    order;
    over protocol-type  {
        signature name  {
            member name  {
```

```
                context {
                    context;
                }
                direction {
                    any;
                    client-to-server;
                    server-to-client;
                }
                pattern pattern;
                depth byte-number;
            }
            port-range value;
    priority [high | low];
    type;
    risk;
}
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Configure application definition.

You can create custom application signatures by specifying a name, protocol, port where the application runs, and match criteria. You can create ICMP-based, address-based, IP protocol-based, and Layer 7-based custom application signatures. Custom applications are created to to identify applications over Layer 7 and transiting or temporary applications, and to achieve further granularity of known applications.

Custom application definitions can be used for applications that are not part of the Juniper Networks predefined application database.

Options

application *application-name*

Name of the custom application signature. Must be a unique name with a maximum length of 63 characters.

> **NOTE**: Application names are case insensitive.

cacheable Enable caching of application identification results. By enabling this option, you can cache the application detection result in an ASC table. If there is an entry in the ASC table, based on the destination IP address, protocol, and the port, we can identify AppID without again sending packet to engine. This option is not supported for address-based, IP protocol-based, and ICMP-based custom application signatures.

description Description of the application.

priority Priority of custom applications over the predefined applications.

- **Values:** high

order *number* Specify the order for the custom application. Lower order has higher priority. This option is used when multiple custom applications of the same type match the same traffic. However, you cannot use this option to prioritize among different type of applications such as TCP stream-based applications against TCP port-based applications or IP address-based applications against port-based applications.

priority [high | low] Specify the priority over other signature applications.

type Specify if application is a well-known application such as HTTP and FTP.

risk Custom application risk value should range from 1 to 5 to keep in sync with the predefined applications. The default value is 1 when the risk is not configured. Configuring risk value for custom application signatures is not supported.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D40.

`Risk` option introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Example: Configuring Junos OS Application Identification Custom Application Signatures

address-mapping (Application Identification) | **1889**

icmp-mapping (Application Identification) | **8263**

ip-protocol-mapping (Application Identification) | **9258**

over (Application Identification) | **12028**

application (Application Identification)

IN THIS SECTION

- Syntax | **2306**
- Hierarchy Level | **2307**
- Description | **2308**
- Options | **2308**
- Required Privilege Level | **2308**
- Release Information | **2308**

Syntax

```
application application-name <description description> {
   address-mapping name {
       destination {
           ip ip-address-prefix;
       }
       source {
           ip ip-address-prefix;
       }
       order order;
       order-priority (high | low);
       }
   }
   alt-name alt-name;
```

```
    cacheable;
    compatibility junos-compatibility-version;
    description description;
    icmp-mapping {
        code icmp-code;
        order order;
        order-priority (high | low);
        type icmp-type;
    }
    ip-protocol-mapping {
        order order;
        order-priority (high | low);
        protocol protocol-number
    }
    order order;
    over protocol-type {
        signature l4-l7-signature-name {
            chain-order
            member member-name {
                check-bytes  max-bytes-to-check;
                context context;
                pattern pattern;
                direction direction;
            }
            order order;
            order-priority (high | low);
            port-range {
                tcp [port-range];
                udp [port-range];
            }
            protocol (http | ssl | tcp | udp);
        ]
    priority;
    tags tag-value;
    type type;
}
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Configure identification of an application for which one or more custom signatures are defined.

Options

application-name Name of the application for which one or more custom signatures has been defined.

description (Optional) Textual description of the application for which mappings are provided.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

application (Application Maps)

IN THIS SECTION

Syntax | 2309

Hierarchy Level | 2309

- Description | 2309
- Options | 2309
- Required Privilege Level | 2309
- Release Information | 2309

Syntax

```
application application-name {
    code-points [ aliases ] [ bit-patterns ];
}
```

Hierarchy Level

```
[edit policy-options application-maps application-map-name]
```

Description

Add an application to an application map and define the application's code points.

Options

application-name—Name of the application.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring an Application Map for DCBX Application Protocol TLV Exchange

Example: Configuring DCBX Application Protocol TLV Exchange

Example: Configuring DCBX to Support an iSCSI Application

Understanding DCBX Application Protocol TLV Exchange

Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

application (Applications)

IN THIS SECTION

- Syntax | 2310
- Hierarchy Level | 2311
- Description | 2311
- Options | 2311
- Required Privilege Level | 2313
- Release Information | 2313

Syntax

```
application application-name {
    application-protocol (dns | ftp |gprs-gtp-c| gprs-gtp-u | gprs-gtp-v0 |gprs-sctp | http |
ignore | ike-esp-nat |  mgcp-ca |  mgcp-ua |  ms-rpc |   q931 |  ras |  realaudio | rsh |   rtsp
|  sccp | sip  |  sqlnet-v2 | sun-rpc |  talk |  tftp);
    description text;
    destination-port port-identifier;
    do-not-translate-A-query-to-AAAA-query;
    do-not-translate-AAAA-query-to-A-query;
    ether-type hex-value;
    icmp-code value;
    icmp-type value;
    icmp6-code value;
    icmp6-type value;
```

```
    inactivity-timeout (seconds | never);
    protocol number;
    rpc-program-number number;
    source-port port-number;
     term term-name {
        alg application;
        destination-port port-identifier;
        icmp-code value;
        icmp-type value;
        icmp6-code value;
        icmp6-type value;
        inactivity-timeout (seconds | never);
        protocol number;
        rpc-program-number number;
        source-port port-number;
        uuid hex-value;
    }
    uuid hex-value;
}
```

Hierarchy Level

```
[edit applications]
```

Description

Configure application properties at the [applications] hierarchy level.

Options

- `application-protocol` *`protocol-name`*—Specify the name of the application protocol.
- `description` *`text`*—Describe the application.
- `destination-port` *`port-identifier`*—Specify a TCP or UDP destination port number.
- `do-not-translate-A-query-to-AAAA-query`—Set the statement to control the translation of A query to AAAA query.
- `do-not-translate-AAAA-query-to-A-query`—Set the statement to control the translation of AAAA query to A query.

- `ether-type` *`value`*—Specify the Ethernet packet type value.
- `icmp-code` *`value`*—Specify the Internet Control Message Protocol (ICMP) code value.

  Range: 0 through 255.
- `icmp-type` *`value`*—Specify the ICMP packet type value.

  Range: 0 through 255.
- `icmp6-code` *`value`*—Specify the ICMP6 code value.

  Range: 0 through 255.
- `icmp6-type` *`value`*—Specify the ICMP6 packet type value.

  Range: 0 through 255.
- `inactivity-timeout (`*`seconds`*` | never)`—Specify the amount of time the application is inactive before it times out in seconds.

  Range: 4 through 129,600 seconds.

  Default: For TCP, 1800 seconds; for UDP, 60 seconds.
- `protocol` *`value`*—Specify the networking protocol name or number.
- `rpc-program-number` *`value`*—Specify the remote procedure call (RPC) or Distributed Computing Equipment (DCE) value.
- `source-port` *`port-number`*—Specify a TCP or UDP source port number.
- `term` *`term-name`*—Specify the individual application protocol.
  - `alg` *`application`*—Specify the name of the application protocol.
  - `destination-port` *`port-identifier`*—Specify a TCP or UDP destination port number.
  - `icmp-code` *`value`*—Specify the ICMP code value.

    Range: 0 through 255.
  - `icmp-type` *`value`*—Specify the ICMP packet type value.

    Range: 0 through 255.
  - `icmp6-code` *`value`*—Specify the ICMP6 code value.

    Range: 0 through 255.
  - `icmp6-type` *`value`*—Specify the ICMP6 packet type value.

Range: 0 through 255.

- `inactivity-timeout (seconds | never)`—Specify the amount of time the application is inactive before it times out in seconds.

  Range: 4 through 129,600 seconds.

  Default: For TCP, 1800 seconds; for UDP, 60 seconds.

- `protocol number`—Specify the networking protocol name or number.
- `rpc-program-number number`—Specify the RPC or DCE value.
- `source-port port-number`—Specify a TCP or UDP source port number.
- `uuid hex-vale`—Specify the universal unique identifier (UUID) for objects.

- `uuid hex-value`—Specify the UUID for objects.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

alg (Applications) | 2138

application-protocol (Applications) | 2355

destination-port (Applications) | 5465

application-bypass (Juniper Secure Connect)

IN THIS SECTION

- Syntax | 2314
- Hierarchy Level | 2314
- Description | 2314
- Options | 2315
- Required Privilege Level | 2315
- Release Information | 2315

Syntax

```
application-bypass {
    term name {
        description description;
        protocol protocol;
        domain-name domain-name
     }
}
```

Hierarchy Level

```
[edit security remote-access client-config name]
```

Description

Define Juniper Secure Connect remote client configuration parameters for bypassing certain applications based on domain names and protocols without passing through the remote access VPN tunnel. Administrator configures these parameters on the SRX Series Firewall which are pushed to client application after its successful authentication.

Options

application-bypass Define application-bypass configuration.

term Define an application-bypass entry; specify a term *name*.

- Values:
  - Format: Must be a string beginning with a number or letter and consisting of letters, numbers, dashes and underscores. Supports upto 50 term entries.

description Text description of remote-access application-bypass profile

domain-name Specify domain-name to bypass VPN tunnel.

- Values:
  - `contains`—Match any domain name containing the suffix value. Example: *example.com*.
  - `fqdn`—Match fully qualified domain name. Example: *hr.example.com*.
  - `wildcard`—Match any subdomain. Example: *.example.com*.

protocol Specify protocol to bypass VPN tunnel

- Values:
  - `all`—Bypass both TCP and UDP traffic.
  - `tcp`—Bypass only TCP traffic.
  - `udp`—Bypass only UDP traffic.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Juniper Secure Connect Administration Guide

application (Extensions)

IN THIS SECTION

- Syntax | 2316
- Hierarchy Level | 2317
- Description | 2317
- Required Privilege Level | 2317
- Release Information | 2317

Syntax

```
application {
    file script-name {
        arguments arguments;
        cert-file path;
        checksum hash-algorithm hash-value;
        daemonize;
        interpreter (bash | python | python3);
        optional;
        refresh;
        refresh-from;
        respawn-on-normal-exit;
        routing-instance;
        source;
        start {
            start-options {
                mem-factor percentage;
            }
        }
        traceoptions {
            file <filename> <files number> <match regex> <size size> <world-readable | no-world-
```

```
readable>;
            flag flag;
            no-remote-trace;
        }
        username username;
    }
    max-datasize max-datasize;
    start {
        start-options {
            mem-factor percentage;
        }
    }
    traceoptions {
        file <filename> <files number> <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system extensions extension-service]
```

Description

Configure the Junos OS extension service application.

> **NOTE**: Global traceoptions for daemonized applications do not take effect if the daemonized application and global traceoptions are committed separately.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

extension-service (System Extensions) | 6717

application-firewall

IN THIS SECTION

- Syntax | 2318
- Hierarchy Level | 2319
- Description | 2319
- Options | 2320
- Required Privilege Level | 2320
- Release Information | 2320

Syntax

```
application-firewall {
    profile profile-name {
        block-message type {
            custom-text content custom-html-text;
            custom-redirect-url content custom-redirect-url;
        }
    }
    rule-sets rule-set-name {
        default-rule {
            (deny [block-message] | permit | reject [block-message]);
        }
        profile profile-name;
        rule rule-name {
            match {
                dynamic-application [system-application];
                dynamic-application-groups [system-application-group];
                ssl-encryption (any | yes | no);
            }
            then  {
```

```
                (deny [block-message] | permit | reject [block-message]);
            }
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            (world-readable | no-world-readable);
            size maximum-file-size;
        }
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Specify the profile options, rule set and rule specifications, and trace options to be used for application firewall implementations.

You can configure the application firewall by defining a collection of rule sets. These rule sets can be defined independently and shared across network security policies. A rule set defines the rules that match the application ID detected, based on the application signature.

The application firewall support in the security policies provides additional security control for dynamic applications.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1. Updated with the `ssl-encryption` and `reject` options in Junos OS Release 12.1X44-D10. Updated with the `block-message` option in Junos OS Release 12.1X45-D10.

The `tenant` option is introduced for Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Application Firewall Overview

application-firewall (Application Services)

IN THIS SECTION

- Syntax | 2321
- Hierarchy Level | 2321
- Description | 2321
- Options | 2321
- Required Privilege Level | 2321
- Release Information | 2322

Syntax

```
application-firewall {
    rule-set rule-set-name;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
application-services]
```

Description

Specify the rule sets configured as part of application firewall to be applied to permitted traffic in a security policy.

The application firewall is defined by a collection of rule sets. You can implement an application firewall by defining one or more application firewall rule sets and creating rules for each rule set that permit, reject, or deny traffic based on the application ID. These rule sets can be defined independently and shared across network security policies. Then you configure a security policy to invoke the application firewall service and specify the rule set to be applied to permitted traffic.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Options

`rule-set` *`rule-set-name`*—Name of the rule set that contains application firewall specification rules.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Application Firewall Overview

rule-sets (Security Application Firewall) | **16151**

application-group

IN THIS SECTION

- Syntax | **2322**
- Hierarchy Level | **2323**
- Description | **2323**
- Options | **2323**
- Required Privilege Level | **2323**
- Release Information | **2323**

Syntax

```
application-group group-name {
    disable;
    application-groups {
        application-group-name;
    }
    applications {
        application-name;
    }
    index number;
}
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Define the properties and contents of the application group.

Options

group-name—Unique identifier for the group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support added in Junos OS release 19.3R2 and 19.4r1 for Next Gen Services on MX240, MX480, and MX960.

NOTE: The `disable` and `index` options are not supported for Next Gen Services.

application-group-any (AACL)

IN THIS SECTION

- Syntax | 2324
- Hierarchy Level | 2324

- Description | 2324
- Required Privilege Level | 2324
- Release Information | 2324

Syntax

```
application-group-any;
```

Hierarchy Level

```
[edit services aacl rule rule-name term term-name from]
[edit services aacl rule rule-name No Link Title term-name No Link Title]
```

Description

Match any application group defined in the database.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Match Conditions in AACL Rules

application-group (Services)

IN THIS SECTION

- Syntax | 2325
- Hierarchy Level | 2325
- Description | 2325
- Options | 2326
- Required Privilege Level | 2326
- Release Information | 2326

Syntax

```
application-group group-name {
    application-groups application-group-name;
    applications application-name;
}
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Configure a custom application group for application identification.

Applications can be grouped under predefined and custom application groups. You can add number of applications or application groups that you want to include in your custom application group.

You can configure an application group to associates related applications under a single name for simplified, consistent reuse in configuring application-based policies.

Options

group-name	Name of the group. This name is used in policy configuration statements in place of multiple predefined applications, user-defined applications, or other groups.
application-groups *application-group-name*	Name of an application group to be assigned to this group. There is no maximum number of groups that can be assigned to a group. Use multiple commands to assign multiple groups.
applications *application-name*	Name of an application to be assigned to this group. An application can remain unassigned or be assigned to a group, but it cannot be assigned to more than one group. There is no maximum number of applications that can be assigned to a group. Use multiple commands to assign multiple groups.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Example: Configuring a Custom Application Group for Junos OS Application Identification for Simplified Management

application-groups (PCC Rules)

IN THIS SECTION

- Syntax | 2327
- Hierarchy Level | 2327
- Description | 2327

- Options | 2327
- Required Privilege Level | 2328
- Release Information | 2328

Syntax

```
application-groups [application-group-name];
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rule-name from],
[edit services pcef pcc-rules rule-name from]
```

Description

Specify one or more application groups to define the match criteria for the policy and charging control (PCC) rule. You can specify a maximum of 10 application groups in a PCC rule.

> **NOTE**: You must also include the `flows` statement. If you do not want to filter subscriber traffic based on service data flow filters, use `flows any`.

If you are using Junos OS Subscriber Aware, specify the name of the application group at the `[edit unified-edge pcef pcc-rules rule-name from]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the application group at the `[edit services pcef pcc-rules rule-name from]` hierarchy level.

Options

application-group-name Name of an application group that is used to detect IP packet flows.

- **Range:** 1 through 63 characters.

NOTE: The referenced application groups must have been previously configured in the `[edit services application-identification]` hierarchy level.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-rules rule-name from]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Rules

application-groups (Services AACL)

IN THIS SECTION

- Syntax | 2329
- Hierarchy Level | 2329

- Description | 2329
- Options | 2329
- Required Privilege Level | 2329
- Release Information | 2329

Syntax

```
application-groups [ application-group-names ];
```

Hierarchy Level

```
[edit services aacl rule rule-name term term-name from]
```

Description

Identify one or more application groups defined in the application identification configuration for inclusion as a match condition.

Options

application-group-names—Identifiers of the application groups.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Match Conditions in AACL Rules

application-identification (Services)

IN THIS SECTION

- Syntax | 2330
- Hierarchy Level | 2333
- Description | 2333
- Options | 2333
- Required Privilege Level | 2335
- Release Information | 2335

Syntax

```
application-identification {
    application application-name {
        address-mapping address-name {
            filter {
                ip ip-address-and-prefix-length;
                port-range {
                    tcp [port];
                    udp [port];
                }
            }
        }
        cacheable;
        description;
        icmp-mapping {
            code number;
            type number;
        }
        ip-protocol-mapping {
            protocol number;
```

```
        }
        order;
        over protocol-type  {
            signature name  {
                member name  {
                    context  {
                         http-get-url-parsed-param-parsed;
                        http-header-content-type;
                        http-header-cookie;
                        http-header-host;
                        http-header-user-agent;
                        http-post-url-parsed-param-parsed;
                        http-post-variable-parsed ;
                        http-url-parsed;
                         http-url-parsed-param-parsed;
                        ssl-server-name;
                         stream;
                    }
                    direction {
                        any;
                        client-to-server;
                        server-to-client;
                    }
                    pattern pattern;
                }
                port-range value;
        priority [high | low];
        type;
        risk;
    }
    application-group group-name {
        application-groups application-group-name;
        applications application-name;
    }
    application-system-cache-timeout value;
    download (Services) {
        automatic {
            interval hours;
            start-time  MM-DD.hh:mm;
        }
        url url;
    }
    enable-cdn-application-detection
```

```
    enable-performance-mode max-packet-threshold number;
    global-offload-byte-limit  byte-limit-number;
    imap-cache-size number;
    imap-cache-timeout  number;
    inspection-limit {
        tcp {
            byte-limit byte-limit-number;
            packet-limit packet-limit-number;
        }
        udp {
            byte-limit byte-limit-number;
            packet-limit packet-limit-number;
        }
    }
    max-memory memory-value
    maximum-transactions transactions-number;
    micro-apps;
    no-application-identification;
    no-application-system-cache;
    packet-capture {
        aggressive-mode;
        buffer-packets-limit bytes;
        capture-interval capture-interval;
        capture-limit capture-limit;
        global;
        max-bytes bytes;
        max-files max-files;
        max-packets max-packets;
        no-inconclusive;
        storage-limit bytes;
    }
    statistics {
        interval minutes;
    }
    traceoptions {
        file {
            filename ;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
```

```
        level [all | error | info | notice | verbose | warning]
        no-remote-trace;
{
no-application-statistics;
}
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure application identification to identify applications regardless of the application port or protocol that is used to transmit the application.

Use this option to configure various options for the application identification such as application signatures, application groups, signature package download option, enable and deactivating application system cache, application traffic throughput, micro applications, application identification inspection limit, trace options and so on to use the application identification functionality.

Once the application is determined, other AppSecure service modules are configured to monitor and control traffic for tracking, prioritization, access control, detection, and prevention based on the application ID of the traffic.

Options

application *application-name*	Configure application definition. You can create custom application signatures by specifying a name, protocol, port where the application runs, and match criteria.
application-group *group-name*	Configure a custom application group for application identification.
application-system-cache-timeout *value*	Specify the timeout value in seconds for the application system cache (ASC) entries.
download	Configure automatic download for the application identification services application package.

enable-cdn-application-detection

Enable application identification (AppID) to classify a web application hosted on a content delivery network (CDN).

enable-performance-mode max-packet-threshold *number*

Set the deep packet inspection (DPI) in performance mode for application identification.

global-offload-byte-limit *byte-limit-number*

Specify the maximum number of byte limit before concluding the classification for identifying an application.

> **NOTE**: The byte limit excludes the IP header and the TCP/UDP header lengths.

- **Range:** 0 through 4294967295
- **Default:** 10000

imap-cache-size *number*

Configure to limit the maximum number of entries in the IMAP cache.

imap-cache-timeout *time-period*

Specify the timeout value for the entries in the IMAP cache cache.

inspection-limit

Specify the maximum number of byte limit before concluding the classification for identifying an application in TCP and UDP sessions.

> **NOTE**: The byte limit excludes the IP header and the TCP/UDP header lengths.

tcp byte-limit *byte-limit-number*

Specify the byte limit.

- **Range:** 0 through 4294967295
- **Default:** 6000
- **Default:** For Junos OS Release 15.1X49-D200, the default value is 10000.

tcp packet-limit *packet-limit-number*

Specify the packet limit.

- **Range:** 0 through 4294967295
- **Default:** 0

udp byte-limit *byte-limit-number*

Specify the byte limit.

- **Range:** 0 through 4294967295
- **Default:** 0

udp packet-limit *packet-limit-number*

Specify the packet limit.

- **Range:** 0 through 4294967295
- **Default:** 10
- **Default:** For Junos OS Release 15.1X49-D200, the default value is 20.

max-memory *value*

Specify maximum memory limit for the deep packet inspection (DPI).

- **Range:** 1 through 200000 MB

micro-apps

Enable micro-application detection with application identification feature.

no-application-identification;

Disable the application identification of applications running on nonstandard ports. By default, application identification is enabled on the device.

no-application-system-cache

Disable application system cache. ASC is enabled by default when a session is created

interval *interval-number;*

Specify the interval, in minutes, for statistics collection.

traceoptions

pecify the trace file information.

no-application-statistics

Configure this configuration statement to disable the application statistics in the AppTrack session.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Custom application definition option introduced in Junos OS Release 15.1X49-D40.

`Risk` option introduced in Junos OS Release 19.1R1.

`micro-app` option introduced in Junos OS Release 19.2R1.

`global-offload-byte-limit` and `inspection-limit` options are introduced in Junos OS Release 19.4R1 and 15.1X49-D200.

Configuration statement `no-application-statistics` is added in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Understanding Application Identification Techniques

application-identification (IDP Sensor Configuration)

IN THIS SECTION

- Syntax | 2336
- Hierarchy Level | 2337
- Description | 2337
- Options | 2337
- Required Privilege Level | 2337
- Release Information | 2337

Syntax

```
application-identification {
    max-packet-memory-ratio percentage-value;
    max-reass-packet-memory-ratio percentage-value;
    max-tcp-session-packet-memory value;
    max-udp-session-packet-memory value;
}
```

Hierarchy Level

```
[edit security idp sensor-configuration]
```

Description

Enable to identify the TCP/UDP application session running on nonstandard ports to match the application properties of transiting network traffic.

Options define the allocation of IDP memory to application identification for packet and reassembler use.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. Packet memory percentages added in Junos OS Release 12.1X44-D20.

application-identification (Application Identification)

IN THIS SECTION

- Syntax | 2338
- Hierarchy Level | 2340
- Description | 2340
- Options | 2341

- Required Privilege Level | 2343
- Release Information | 2343

Syntax

```
application-identification {
    application application-name <description description> {
        address-mapping name {
            destination {
                ip ip-address-prefix;
            }
            source {
                ip ip-address-prefix;
            }
            order order;
            order-priority (high | low);
            }
        }
        alt-name alt-name;
        cacheable;
        compatibility junos-compatibility-version;
        description description;
        icmp-mapping {
            code icmp-code;
            order order;
            order-priority (high | low);
            type icmp-type;
        }
        ip-protocol-mapping {
            order order;
            order-priority (high | low);
            protocol protocol-number
        }
        order
        over protocol-type {
            signature l4-l7-signature-name {
                chain-order
                member member-name {
```

```
                    check-bytes  max-bytes-to-check;
                    context context;
                    pattern pattern;
                    direction direction;
                }
                order order;
                order-priority (high | low);
                port-range {
                    tcp [port-range];
                    udp [port-range];
                }
                protocol (http | ssl | tcp | udp);
            ]
        }
        priority;
        tags tag-value;
        type type;
    }
    application-group group-name {
        disable;
        application-groups {
            application-group-name;
        }
        applications {
            application-name;
        }
        index number;
    }
    application-system-cache-timeout;
    download  {
    }
    inspection-limit {
        tcp {
            byte-limit byte-limit-number;
            packet-limit packet-limit-number;
        }
        udp {
            byte-limit byte-limit-number;
            packet-limit packet-limit-number;
        }
    }
    micro-apps;
    micro-app-max-transactions value;
```

```
    no-application-system-cache;
    statistics {
        interval minutes;
    }
    traceoptions {
        file {
            filename ;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        level [all | error | info | notice | verbose | warning]
        no-remote-trace;
    }
    no-application-system-cache;
    packet-capture
    profile profile-name
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure application identification to identify applications regardless of the application port or protocol that is used to transmit the application.

Use this option to configure various options for the application identification such as application signatures, application groups, signature package download option, enable and deactivating application system cache, application traffic throughput, micro applications, application identification inspection limit, trace options and so on to use the application identification functionality.

Once the application is determined, other AppSecure service modules are configured to monitor and control traffic for tracking, prioritization, access control, detection, and prevention based on the application ID of the traffic.

Options

application *application-name*

Configure application definition. You can create custom application signatures by specifying a name, protocol, port where the application runs, and match criteria.

application-group *group-name*

Configure a custom application group for application identification.

application-system-cache-timeout *value*

Specify the timeout value in seconds for the application system cache (ASC) entries.

download

Configure automatic download for the application identification services application package.

enable-cdn-application-detection

Enable application identification (AppID) to classify a web application hosted on a content delivery network (CDN).

enable-performance-mode max-packet-threshold *number*

Set the deep packet inspection (DPI) in performance mode for application identification.

global-offload-byte-limit *byte-limit-number*

Specify the maximum number of byte limit before concluding the classification for identifying an application.

> **NOTE**: The byte limit excludes the IP header and the TCP/UDP header lengths.

- **Range:** 0 through 4294967295
- **Default:** 10000

imap-cache-size *number*

Configure to limit the maximum number of entries in the IMAP cache.

imap-cache-timeout *time-period*

Specify the timeout value for the entries in the IMAP cache cache.

inspection-limit

Specify the maximum number of byte limit before concluding the classification for identifying an application in TCP and UDP sessions.

> **NOTE**: The byte limit excludes the IP header and the TCP/UDP header lengths.

tcp byte-limit *byte-limit-number*

Specify the byte limit.

- **Range:** 0 through 4294967295
- **Default:** 6000
- **Default:** For Junos OS Release 15.1X49-D200, the default value is 10000.

tcp packet-limit *packet-limit-number*

Specify the packet limit.

- **Range:** 0 through 4294967295
- **Default:** 0

udp byte-limit *byte-limit-number*

Specify the byte limit.

- **Range:** 0 through 4294967295
- **Default:** 0

udp packet-limit *packet-limit-number*

Specify the packet limit.

- **Range:** 0 through 4294967295
- **Default:** 10
- **Default:** For Junos OS Release 15.1X49-D200, the default value is 20.

max-memory *value*

Specify maximum memory limit for the deep packet inspection (DPI).

- **Range:** 1 through 200000 MB

micro-apps

Enable micro-application detection with application identification feature.

micro-app-max-transactions

Set the number of micro-applications transaction finals to terminate application classification.

- **Range:** 1 through 25
- **Default:** 5

no-application-identification;

Disable the application identification of applications running on nonstandard ports. By default, application identification is enabled on the device.

no-application-system-cache

Disable application system cache. ASC is enabled by default when a session is created

interval ***interval-number;***	Specify the interval, in minutes, for statistics collection.
traceoptions	Specify the trace file information.
no-application-statistics	Configure this configuration statement to disable the application statistics in the AppTrack session.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2 and 19.4R1 on MX Series routers MX240, MX480 and MX960.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

application-identification-profile (Service Set)

IN THIS SECTION

- Syntax | 2344
- Hierarchy Level | 2344
- Description | 2344
- Options | 2344
- Required Privilege Level | 2344
- Release Information | 2344

Syntax

```
application-identification-profile app-id-profile-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the dummy application identification profile that you configured at the `[edit services application-identification profile]` hierarchy level. This profile is a placeholder profile with no configuration options, but it must be specified to enable application identification functionality on the services plane.

Options

app-id-profile-name	Name of the application identification profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> *Applying Services to Subscriber-Aware Traffic with a Service Set*
>
> *Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control*

application-map

IN THIS SECTION

- Syntax | 2345
- Hierarchy Level | 2345
- Description | 2345
- Options | 2345
- Required Privilege Level | 2345
- Release Information | 2346

Syntax

```
application-map application-map-name;
```

Hierarchy Level

```
[edit protocols dcbx interface interface-name]
```

Description

Specify an application map to apply to an interface.

Options

application-map-name—Name of the application map.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- show dcbx neighbors | 19191
- *Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange*
- *Example: Configuring DCBX Application Protocol TLV Exchange*
- Example: Configuring DCBX to Support an iSCSI Application
- *Understanding DCBX Application Protocol TLV Exchange*
- Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

application-maps

IN THIS SECTION

- Syntax | 2346
- Hierarchy Level | 2347
- Description | 2347
- Options | 2347
- Required Privilege Level | 2347
- Release Information | 2347

Syntax

```
application-maps application-map-name {
    application application-name {
        code-points [ aliases ] [ bit-patterns ];
    }
}
```

Hierarchy Level

```
[edit policy-options]
```

Description

Define an application map by specifying the applications that belong to the application map.

Options

application-map-name—Name of the application map.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- *Configuring an Application Map for DCBX Application Protocol TLV Exchange*
- *Example: Configuring DCBX Application Protocol TLV Exchange*
- Example: Configuring DCBX to Support an iSCSI Application
- *Understanding DCBX Application Protocol TLV Exchange*
- Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

application (NAT Next Gen Services)

IN THIS SECTION

- Syntax | 2348
- Hierarchy Level | 2348
- Description | 2348
- Required Privilege Level | 2348
- Release Information | 2348

Syntax

```
application [application-name]
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name match],
[edit services nat source rule-set rule-set rule rule-name match]
```

Description

Specify one or more application protocols to which the NAT rule applies. The number of applications must not exceed 3072.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

application-profile

IN THIS SECTION

- Syntax | 2349
- Hierarchy Level | 2350
- Description | 2350
- Options | 2350
- Required Privilege Level | 2350
- Release Information | 2350

Syntax

```
application-profile profile-name {
    ftp {
        data {
            dscp (alias | bits);
            forwarding-class class-name;
        }
    }
    sip {
        video {
            dscp (alias | bits);
            forwarding-class class-name;
        }
        voice {
            dscp (alias | bits);
            forwarding-class class-name;
        }
    }
}
```

Hierarchy Level

```
[edit services cos],
[edit services cos rule rule-name term term-name then],
[edit services cos rule rule-name term term-name then reverse]
```

Description

Define or apply a CoS application profile. When you apply a CoS application profile in a CoS rule, terminate the profile name with a semicolon (;).

Options

profile-name—Identifier for the application profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring Application Profiles for Use as CoS Rule Actions

application-profile (Services CoS Next Gen Services)

IN THIS SECTION

- Syntax | 2351
- Hierarchy Level | 2352
- Description | 2352
- Options | 2352
- Required Privilege Level | 2352
- Release Information | 2352

Syntax

```
application-profile name {
                        ftp {
                                data {
                                    dscp                        dscp;
                                    forwarding-class                        forwarding-class;
        }
    }
                        sip {
                                video {
                                    dscp                        dscp;
                                    forwarding-class                        forwarding-class;
        }
                                voice {
                                    dscp                        dscp;
                                    forwarding-class                        forwarding-class;
        }
    }
}
```

Hierarchy Level

```
[edit services cos]
```

Description

Configure CoS actions for FTP and SIP traffic. The application profile can then be used in CoS rule actions. This enables you to apply a certain DSCP, or forwarding-class to a set of L7 flows.

Options

profile-name	Name of the application profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Class of Service for Services PICs (Next Gen Services)

application-protocol

IN THIS SECTION

- Syntax | 2353
- Hierarchy Level | 2353

- Description | 2353
- Options | 2353
- Required Privilege Level | 2354
- Release Information | 2354

Syntax

```
application-protocol protocol-name;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Identify the application protocol name. Application protocols are also called application layer gateways (ALGs).

Options

protocol-name—Name of the protocol. The following protocols are supported:

1. `bootp`—Bootstrap protocol
2. `dce-rpc`—DCE RPC
3. `dce-rpc-portmap`—DCE RPC portmap
4. `dns`—Domain Name Service
5. `exec`—Remote Execution Protocol
6. `ftp`—File Transfer Protocol
7. `h323`—H.323
8. `icmp`—ICMP

9. `iiop`—Internet Inter-ORB Protocol
10. `ike-esp-nat`—IKE ALG
11. `ip`—IP
12. `login`—Login
13. `netbios`—NetBIOS
14. `netshow`—NetShow
15. `pptp`—Point-to-Point Tunneling Protocol
16. `ras`—Gatekeeper RAS for H323
17. `realaudio`—RealAudio
18. `rpc`—RPC
19. `rpc-portmap`—RPC portmap
20. `rtsp`—Real Time Streaming Protocol
21. `shell`—Shell
22. `sip`—Session Initiation Protocol
23. `snmp`—SNMP
24. `sqlnet`—SQLNet
25. `talk`–Talk Program
26. `tftp`–Trivial File Transfer Protocol
27. `traceroute`—Traceroute
28. `winframe`—WinFrame

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`login` options introduced in Junos OS Release 7.4.

`ip` option introduced in Junos OS Release 8.2.

`ike-esp-nat` option introduced in Junos OS Release 17.1.

`ras` option introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

ALG Descriptions

Configuring Application Sets

Configuring Application Properties

Examples: Configuring Application Protocols

Verifying the Output of ALG Sessions

application-protocol (Applications)

IN THIS SECTION

- Syntax | 2355
- Hierarchy Level | 2356
- Description | 2356
- Default | 2357
- Required Privilege Level | 2357
- Release Information | 2357

Syntax

```
application-protocol (bootp | dce-rpc | dce-rpc-portmap | dns | exec | ftp | ftp-data | gprs-gtp-
c | gprs-gtp-u | gprs-gtp-v0 | gprs-sctp | h323 | http | https | icmp | icmpv6 | ignore | iiop |
ike-esp-nat | imap | imaps | ip | login | mgcp-ca | mgcp-ua | ms-rpc | netbios | netshow | none
| pop3 | pop3s | pptp | q931 | ras | realaudio | rpc | rpc-portmap | rsh | rtsp | sccp | shell |
```

```
sip | smtp | smtps | snmp | sqlnet | sqlnet-v2 | ssh | sun-rpc | talk | telnet | tftp |
traceroute | twamp | winframe);
```

Hierarchy Level

```
[edit applications application],
[edit logical-systems name applications application],
[edit logical-systems name tenants name applications application],
[edit tenants name applications application]
```

Description

Identify the application protocol name. Application protocols are also called application layer gateways (ALGs). Define an application layer gateway (ALG) for a particular application. The following ALGs are supported:

protocol-name—Name of the protocol. The following protocols are supported:

1. `bootp`—Bootstrap protocol
2. `dce-rpc`—DCE RPC
3. `dce-rpc-portmap`—DCE RPC portmap
4. `dns`—Domain Name Service
5. `exec`—Remote Execution Protocol
6. `ftp`—File Transfer Protocol
7. `h323`—H.323
8. `icmp`—ICMP
9. `iiop`—Internet Inter-ORB Protocol
10. `ike-esp-nat`—IKE ALG
11. `ip`—IP
12. `login`—Login
13. `netbios`—NetBIOS
14. `netshow`—NetShow

15. `pptp`—Point-to-Point Tunneling Protocol
16. `ras`—Gatekeeper RAS for H323
17. `realaudio`—RealAudio
18. `rpc`—RPC
19. `rpc-portmap`—RPC portmap
20. `rtsp`—Real Time Streaming Protocol
21. `shell`—Shell
22. `sip`—Session Initiation Protocol
23. `snmp`—SNMP
24. `sqlnet`—SQLNet
25. `talk`–Talk Program
26. `tftp`–Trivial File Transfer Protocol
27. `traceroute`—Traceroute
28. `winframe`—WinFrame

Default

If `application-protocol` is not defined, then traffic enters the ALG process as long the status for this ALG is enabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`login` options introduced in Junos OS Release 7.4.

`ip` option introduced in Junos OS Release 8.2.

`ike-esp-nat` option introduced in Junos OS Release 17.1.

`ras` option introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Security Policy Application Sets Overview

application-proxy (Security Content Security)

IN THIS SECTION

- Syntax | 2358
- Hierarchy Level | 2358
- Description | 2359
- Options | 2359
- Required Privilege Level | 2359
- Release Information | 2359

Syntax

```
application-proxy {
    traceoptions {
        flag flag;
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm]
```

Description

Configure trace options for the application proxy.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

application-screen (Security H323)

IN THIS SECTION

- Syntax | 2359
- Hierarchy Level | 2360
- Description | 2360
- Options | 2360
- Required Privilege Level | 2360
- Release Information | 2360

Syntax

```
application-screen {
    message-flood {
```

```
        gatekeeper {
            threshold rate;
        }
    }
    unknown-message {
        permit-nat-applied;
        permit-routed;
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg h323],
[edit logical-systems name tenants name security alg h323],
[edit security alg h323],
[edit services alg h323],
[edit tenants name security alg h323]
```

Description

Configure the security screens for the H.323 protocol Application Layer Gateway (ALG) to handle unknown-messages and protect H.323 server.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

h323 | 7961

application-screen (Security MGCP)

IN THIS SECTION

- Syntax | 2361
- Hierarchy Level | 2361
- Description | 2362
- Options | 2362
- Required Privilege Level | 2362
- Release Information | 2362

Syntax

```
application-screen {
    connection-flood {
        threshold rate;
    }
    message-flood {
        threshold rate;
    }
    unknown-message {
        permit-nat-applied;
        permit-routed;
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg mgcp],
[edit logical-systems name tenants name security alg mgcp],
```

```
[edit security alg mgcp],
[edit services alg mgcp],
[edit tenants name security alg mgcp]
```

Description

Configure the security screens for the Media Gateway Control Protocol (MGCP) Application Layer Gateway (ALG) to handle unknown-messages and protect MGCP server.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

mgcp | 10693

application-screen (Security SCCP)

IN THIS SECTION

- Syntax | 2363
- Hierarchy Level | 2363
- Description | 2363
- Options | 2363

Required Privilege Level | 2363
Release Information | 2364

Syntax

```
application-screen {
    call-flood {
        threshold rate;
    }
    unknown-message {
        permit-nat-applied;
        permit-routed;
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg sccp],
[edit logical-systems name tenants name security alg sccp],
[edit security alg sccp],
[edit services alg sccp],
[edit tenants name security alg sccp]
```

Description

Configure the security screens for the Skinny Client Control Protocol (SCCP) Application Layer Gateway (ALG) to handle unknown-messages and protect SCCP server.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sccp | 16341

application-screen (Security SIP)

IN THIS SECTION

- Syntax | 2364
- Hierarchy Level | 2365
- Description | 2365
- Options | 2365
- Required Privilege Level | 2365
- Release Information | 2365

Syntax

```
application-screen {
    protect {
        deny {
            all {
                timeout value-in-seconds;
            }
            destination-ip address;
            timeout value-in-seconds;
        }
    }
    unknown-message {
```

```
        permit-nat-applied;
        permit-routed;
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg sip],
[edit logical-systems name tenants name security alg sip],
[edit security alg sip],
[edit services alg sip],
[edit tenants name security alg sip]
```

Description

Configure the security screens for the Session Initiation Protocol (SIP) Application Layer Gateway (ALG) to handle unknown-messages and protect SIP server.

IPv6 is supported on the SIP ALG along with Network Address Translation Protocol Translation (NAT-PT) mode and NAT64 address translation.

The type of the <destination-ip-address> is changed from IPv4 address to IP prefix to support all kinds of IP addresses, and correspondingly a prefix is supported to allow multiple IP addresses.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sip (Security) | 27977

application (Security Alarms)

IN THIS SECTION

- Syntax | 2366
- Hierarchy Level | 2366
- Description | 2366
- Options | 2367
- Required Privilege Level | 2367
- Release Information | 2367

Syntax

```
application {
    duration interval;
    size count;
    threshold value;
}
```

Hierarchy Level

```
[edit security alarms potential-violation policy]
```

Description

Configure alarms for a number of policy violations to an application within a specified time period.

Options

- `duration` *`interval`*—Indicate the duration of counters.
- **Range:** 1 through 3600 seconds.
- **Default:** 1 second.
- `size` *`count`*—Indicate the number of applications for which policy violation checks can be done concurrently.
- **Range:** 1 through 10240.
- **Default:** 1024.
- `threshold` *`value`*—Indicate the maximum number of application matches required to raise an alarm.
- **Range:** 1 through 4294967295.
- **Default:** 1000.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Security Policies Overview

application (Security Custom Attack)

IN THIS SECTION

Syntax | 2368

- Hierarchy Level | 2368
- Description | 2368
- Options | 2368
- Required Privilege Level | 2368
- Release Information | 2368

Syntax

```
application application-name;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain protocol-binding]
[edit security idp custom-attack attack-name attack-type signature protocol-binding]
```

Description

Allow IDP to match the attack for a specified application.

Options

application-name—Name of the application.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

application (Security Destination NAT)

IN THIS SECTION

- Syntax | 2369
- Hierarchy Level | 2369
- Description | 2369
- Options | 2370
- Required Privilege Level | 2370
- Release Information | 2370

Syntax

```
application {
    [application];
    any;
}
```

Hierarchy Level

```
[edit logical-systems name security nat destination rule-set rule-set-name rule rule-name match],
[edit logical-systems name tenants name security nat destination rule-set rule-set-name rule
rule-name match],
[edit security nat destination rule-set rule-set-name rule rule-name match],
[edit services nat destination rule-set rule-set-name rule rule-name match],
[edit tenants name nat destination rule-set rule-set-name rule rule-name match]
```

Description

Specify an application name to match the rule. You can specify multiple application names, but the number of application terms must not exceed 3072.

Options

application-name—Name of the application.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

application (Security Policies) | 2371

application (Security IDP)

IN THIS SECTION

- Syntax | 2370
- Hierarchy Level | 2371
- Description | 2371
- Options | 2371
- Required Privilege Level | 2371
- Release Information | 2371

Syntax

```
application application-name;
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match]
```

Description

Specify an application or an application set name to match.

Options

application-name—Name of the application.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

application (Security Policies)

IN THIS SECTION

- Syntax | 2372
- Hierarchy Level | 2372
- Description | 2372
- Options | 2372
- Required Privilege Level | 2372
- Release Information | 2373

Syntax

```
application {
    [application];
    any;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name match]
```

```
[edit security policies global policy policy-name match]
```

Description

Specify the IP or remote procedure call (RPC) application or set of applications to be used as match criteria.

Starting in Junos OS Release 19.1R1, configuring the `application` statement at the `[edit security policies from-zone zone-name to-zone zone-name policy policy-name match]` hierarchy level is optional if the `dynamic-application` statement is configured at the same hierarchy level.

Options

application-name-or-set Name of the predefined or custom application or application set used as match criteria.

any Any predefined or custom applications or application sets.

> **NOTE**: A custom application that does not use a predefined destination port for the application will not be included in the `any` option, and must be named explicitly.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

Configure Applications in Unified Policies

application (Security Source NAT)

IN THIS SECTION

- Syntax | 2373
- Hierarchy Level | 2374
- Description | 2374
- Options | 2374
- Required Privilege Level | 2374
- Release Information | 2374

Syntax

```
application {
    [application];
    any;
}
```

Hierarchy Level

```
[edit logical-systems name security nat source rule-set rule-name rule name match],
[edit logical-systems name tenants name security nat source rule-set name rule name match],
[edit security nat source rule-set rule-set-name rule rule-name match],
[edit services nat source rule-set rule-set-name rule rule-name match],
[edit tenants name security nat source rule-set rule-set-name rule rule-name match]
```

Description

Specify an application name to match the rule. You can specify multiple application names, but the number of application terms must not exceed 3072.

Options

application-name—Name of the application.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

application (Security Policies) | 2371

application-services (Security Forwarding Process)

IN THIS SECTION

- Syntax | 2375
- Hierarchy Level | 2375
- Description | 2375
- Options | 2376
- Required Privilege Level | 2376
- Release Information | 2377

Syntax

```
application-services {
    enable-gtpu-distribution;
    inline-fpga-crypto (disabled | enabled);
    maximize-alg-sessions;
    maximize-idp-sessions {
        weight (firewall | idp);
    }
    packet-ordering-mode (Application Services) {
        (hardware | software);

    }
}
```

Hierarchy Level

```
[edit security forwarding-process]
```

Description

You can configure SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices to switch from an integrated firewall mode to maximize Intrusion Detection and Prevention (IDP) mode to run IDP

processing in tap mode and increase the capacity of processing with the `maximize-idp-sessions` option. Inline tap mode can only be configured if the forwarding process mode is set to `maximize-idp-sessions`, which ensures stability and resiliency for firewall services. You also do not need a separate tap or span port to use inline tap mode. When you maximize IDP, you are decoupling IDP processes from firewall processes, allowing the device to support the same number of firewall and IDP sessions, also run the IDP processing in tap mode.

You can configure maximum Application Layer Gateway (ALG) sessions by using the `maximize-alg-sessions` option. The session capacity number for Real-Time Streaming Protocol (RTSP), FTP, and Trivial File Transfer Protocol (TFTP) ALG varies per flow SPU. For SRX5000 line devices the session capacity is 10,240 per flow SPU. You must reboot the device (and its peer in chassis cluster mode) for the configuration to take effect. The `maximize-alg-sessions` option now enables you to increase defaults as follows:

- TCP proxy connection capacity: 40,000 per flow SPU

  Flow session capacity is reduced to half per flow SPU; therefore the aforementioned capacity numbers will not change on central point flow.

Enable GPRS tunneling protocol. GTP-U session distribution is a UE (User equipment) based distribution, generating tunnel based GTP-U session and distributing them across SPUs on a UE basis.

Before 15.1X49-D40, GTP-U sessions are distributed by GGSN IP address always.

15.1X49-D40 onward, the GTP-U distribution is disabled and fat GTP-U sessions are distributed as normal UDP.

Use the `enable-gtpu-distribution` command to enable GTP-U session distribution.

Use the `inline-fpga-crypto (disabled | enabled)` to enable or disable inline FPGA crypto

Options

The remaining statements are explained separately. See the CLI Explorer.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6. Statement updated in Junos OS Release 10.4. Statement updated in Junos OS Release 15.1X49-D40 with the `enable-gtpu-distribution` option. Statement updated in Junos OS Release 20.4R1 with the `inline-fpga-crypto (disabled | enabled)` option.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

application-services (security-metadata-streaming)

IN THIS SECTION

- Syntax | 2377
- Hierarchy Level | 2377
- Description | 2378
- Default | 2378
- Required Privilege Level | 2378
- Release Information | 2378

Syntax

```
application-services {
                security-metadata-streaming-policy policy-name;
                }
```

Hierarchy Level

```
[edit security policies from-zone from-zone to-zone to-zone]
```

Description

Enable security metadata streaming-policy on SRX Series Firewall.

Default

security-metadata-streaming-policy *policy-name*	Specify the security-metadata-streaming policy that you configured at `[edit services]` hierarchy level.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2 R1.

RELATED DOCUMENTATION

security-metadata-streaming | 16487

application-services (Security Policies)

IN THIS SECTION

- Syntax | 2379
- Hierarchy Level | 2379
- Description | 2380
- Options | 2380
- Required Privilege Level | 2381
- Release Information | 2381

Syntax

```
application-services {
    advanced-anti-malware-policy advanced-anti-malware-policy;
    application-firewall {
        rule-set rule-set;
    }
    application-traffic-control {
        rule-set rule-set;
    }
    casb-policy policy-name;
    gprs-gtp-profile gprs-gtp-profile;
    gprs-sctp-profile gprs-sctp-profile;
    idp idp;
    packet-capture;
    (redirect-wx redirect-wx | reverse-redirect-wx reverse-redirect-wx);
    security-intelligence-policy security-intelligence-policy;
    security-intelligence {
    add-destination-identity-to-feed feed-name;
    add-destination-ip-to-feed feed-name;
    add-source-identity-to-feed feed-name;
    add-source-ip-to-feed feed-name;
    }
    security-metadata-streaming-policy policy-name
    ssl-proxy {
        profile-name profile-name;
    }
    uac-policy {
        captive-portal captive-portal;
    }
    utm-policy utm-policy;
    web-proxy {
        profile-name profile-name;
    }
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit]
```

Description

Enable application services within a security policy. You can enable service such as application firewall, IDP, Content Security, SSL proxy, and so on by specifying them in a security policy permit action, when the traffic matches the policy rule.

Options

advanced-anti-malware-policy	Specify advanced-anti-malware policy name.
application-firewall	Specify the rule sets configured as part of application firewall to be applied to the permitted traffic.
application-traffic-control	Specify the rule sets configured as part of AppQoS, application-aware quality of service, to be applied to the permitted traffic.
casb-policy	Specify CASB policy name.
gprs-gtp-profile	Specify GPRS tunneling protocol profile name.
gprs-sctp-profile	Specify GPRS stream control protocol profile name.
idp	Apply Intrusion detection and prevention (IDP) as application services.
redirect-wx	Specify the WX redirection needed for the packets that arrive from the LAN.
reverse-redirect-wx	Specify the WX redirection needed for the reverse flow of the packets that arrive from the WAN.
security-intelligence-policy	Specify security-intelligence policy name.
security-intelligence	Specify the security intelligence feed post action. The following feeds are supported: • add-destination-identity-to-feed • add-destination-ip-to-feed • add-source-identity-to-feed • add-source-ip-to-feed
security-metadata-streaming-policy	Enable metadata streaming of the traffic permitted by the security policy.

uac-policy

Enable Unified Access Control (UAC) for the security policy. This statement is required when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a UAC deployment.

captive-portal *captive-portal*

Specify the preconfigured security policy for captive portal on the Junos OS Enforcer to enable the captive portal feature. The captive portal policy is configured as part of the UAC policy. By configuring the captive portal feature, you can redirect traffic destined for protected resources to the IC Series device or to the URL you configure on the Junos OS Enforcer.

Content Security-policy *Content Security-policy*

Specify Content Security policy name. The Content Security policy configured for antivirus, antispam, content-filtering, traffic-options, and Web-filtering protocols is attached to the security policy to be applied to the permitted traffic.

web-proxy *profile-name*

Specify secure Web proxy profile name. The secure Web proxy profile is configured with dynamic application and external proxy server details. This profile is attached to the security policy and applied on the permitted traffic.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 11.1.

RELATED DOCUMENTATION

Application Firewall Overview

application-set

IN THIS SECTION

- Syntax | 2382
- Hierarchy Level | 2382
- Description | 2382
- Options | 2382
- Required Privilege Level | 2382
- Release Information | 2383

Syntax

```
application-set application-set-name {
    application application-name;
}
```

Hierarchy Level

```
[edit applications]
```

Description

Configure one or more applications to include in an application set.

Options

application-set-name—Identifier of an application set.

Required Privilege Level

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *ALG Descriptions*
- *Configuring Application Sets*
- *Configuring Application Properties*
- *Examples: Configuring Application Protocols*
- *Verifying the Output of ALG Sessions*

application-sets (IDS MS-DPC)

IN THIS SECTION

- Syntax | 2383
- Hierarchy Level | 2383
- Description | 2384
- Options | 2384
- Required Privilege Level | 2384
- Release Information | 2384

Syntax

```
application-sets set-name;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name from]
```

Description

Define one or more target application sets when using the MS-DPC.

Options

set-name—Name of the target application set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Match Conditions in IDS Rules

application-sets (PCP)

IN THIS SECTION

- Syntax | 2385
- Hierarchy Level | 2385
- Description | 2385
- Options | 2385
- Required Privilege Level | 2385
- Release Information | 2385

Syntax

```
applications-sets set-name;
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name from]
```

Description

Define an application set to which the PCP rule applies. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

set-name—Name of the application set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

application-sets (Services CoS)

IN THIS SECTION

- Syntax | 2386
- Hierarchy Level | 2386
- Description | 2386
- Options | 2386
- Required Privilege Level | 2386
- Release Information | 2387

Syntax

```
applications-sets set-name;
```

Hierarchy Level

```
[edit services cos rule rule-name term term-name from]
```

Description

Define one or more target application sets.

Options

set-name—Name of the target application set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring Match Conditions In CoS Rules

application-sets (Services NAT)

IN THIS SECTION

- Syntax | 2387
- Hierarchy Level | 2387
- Description | 2387
- Options | 2388
- Required Privilege Level | 2388
- Release Information | 2388

Syntax

```
applications-sets set-name;
```

Hierarchy Level

```
[edit  services nat rule rule-name term term-name from]
```

Description

Define one or more target application sets.

Options

set-name—Name of the target application set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

application-sets (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 2388
- Hierarchy Level | 2389
- Description | 2389
- Options | 2389
- Required Privilege Level | 2389
- Release Information | 2389

Syntax

```
applications-sets set-name;
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name term term-name from]
```

Description

Define one or more target application sets.

Options

set-name—Name of the target application set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Match Conditions in Stateful Firewall Rules

application-specific (Protocols OSPF)

IN THIS SECTION

- Syntax | 2390
- Hierarchy Level | 2390
- Description | 2390
- Options | 2391

Required Privilege Level | 2391
Release Information | 2391

Syntax

```
application-specific {
            attribute-group name {
                      admin-group;
                      advertise-interface-delay;
                      application flex-algorithm;
                      delay-metric;
                      te-metric;
            }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area area-id interface interface-name],
[edit protocols ospf area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols ospf area area-id interface interface-
name],
[edit routing-instances routing-instance-name protocols ospf area area-id virtual-link]
```

Description

Configure application (flexible algorithm) specific attributes such as te-attributes: te-metric, delay-metric, or admin-groups.

NOTE: The Junos OS and Junos OS Evolved implementation of OSPF application-specific link attribute supports flexible algorithm application only.

Options

attribute-group *name*

Indicates an application-specific link attribute for the link. The attribute-group enables multiple applications to share the same attribute.

- admin-group—Define an admin-group for the attribute group. You can commit this configuration only if an application is associated with the attribute group.
- advertise-interface-delay—Specify to fetch the delay values from the interface configuration hierarchy, that is, legacy delay values. You can commit this configuration only if the following criteria match:
  - An application is associated with the attribute group.
  - The `delay-metric` option is not configured in the hierarchy. The `delay-metric` and `advertise-interface-delay` configuration options must be mutually exclusive.
  - Either of the following configurations: `[edit protocols ospf interface delay-metric` or `delay-measurement` are present.
- application flex-algorithm—Configure application-specific attributes. For example, flexible algorithm specific te-attributes such as te-metric, delay-metric, or admin-groups.
- delay-metric—Define a delay-metric for the attribute group. You can commit this configuration only if an application is associated with the attribute group.
- te-metric—Define a te-metric for the attribute group. You can commit this configuration only if an application is associated with the attribute group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

Understanding OSPF Flexible Algorithm for Segment Routing

application-status

IN THIS SECTION

- Syntax | 2392
- Hierarchy Level | 2392
- Description | 2392
- Default | 2393
- Options | 2393
- Required Privilege Level | 2393
- Release Information | 2393

Syntax

```
application-status (any | error)
```

Hierarchy Level

```
[edit system syslog console]
```

Description

Enable logging of application start or stop status. During system reboot, by default, the systemd process logs messages for applications to the journal log. By default, application stop and start status are written to the journal log. Only application stop and start errors are logged to the console.

This configuration statement allows each application's stop and start application status logs to be logged to the console. You configure the `any` option to allow all systemd application stop and start messages to be written to the console, not just error messages. You configure the `error` option to return to the default setting.

Default

Write only application stop and start error messages to the console.

Options

(any | error)

- any—Write all application stop and start status messages to the console.
- error—Write only application stop and start error messages to the console.
- Default: error

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

System Log Explorer

application-system-cache

IN THIS SECTION

- Syntax | 2394
- Hierarchy Level | 2394
- Description | 2394
- Options | 2394
- Required Privilege Level | 2395

Release Information | 2395

Syntax

```
application-system-cache;
```

Hierarchy Level

```
application-system-cache {
    no-miscellaneous-services;
    security-services;
}
```

Description

Enable application system cache (ASC) to save the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service.

ASC is enabled by default when a session is created. You can manually turn this caching off using the `set services application-identification no-application-system-cache` command. You can re-enable the ASC by using the `delete services application-identification application-system-cache` command.

You can enable the ASC for faster application identification process and disable it for performance benefits and security.

Note the differences in the default behavior of ASC for services starting from Junos OS Release 18.2R1:

- Security services including security policies, application firewall (AppFW), application tracking (AppTrack), application quality of service (AppQoS), Juniper ATP Cloud, IDP, and Content Security do not use the ASC by default.
- Miscellaneous services including advanced policy-based routing (APBR) use the ASC for application identification by default.

Options

no-miscellaneous-services	Disable the ASC for miscellaneous services such as APBR and AppTrack.

security-services Enable the ASC for security services such as security policies, application firewall (AppFW), Juniper ATP Cloud, IDP, and Content Security.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. The options `no-miscellaneous-services` and `security-services` are introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Understanding the Application System Cache

application-system-cache-timeout (Services)

IN THIS SECTION

- Syntax | 2395
- Hierarchy Level | 2396
- Description | 2396
- Options | 2396
- Required Privilege Level | 2396
- Release Information | 2396

Syntax

```
application-system-cache-timeout value;
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Specify the timeout value in seconds for the application system cache (ASC) entries.

ASC saves the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service. By default, the ASC saves the mapping information for 3600 seconds.

> **NOTE**: On SRX Series Firewalls, when you change the timeout value for the application system cache entries using the command `set services application-identification application-system-cache-timeout`, the cache entries need to be cleared to avoid inconsistency in timeout values of existing entries.

> **NOTE**: ASC is not cleared when the IDP policy is loaded. Users need to manually clear or wait for the cache entries to expire.

Options

value—Timeout value for the application system cache entries.

- **Range:** 0 through 1,000,000 seconds
- **Default:** 3600 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. Support.

RELATED DOCUMENTATION

Understanding the Application System Cache

application-tracking

IN THIS SECTION

- Syntax | 2397
- Hierarchy Level | 2397
- Description | 2398
- Options | 2398
- Required Privilege Level | 2399
- Release Information | 2399

Syntax

```
application-tracking {
    (first-update | first-update-interval minutes);
    disable (Application Tracking);
    session-update-interval minutes;
log-session-create;
log-session-close;
session-update-interval session-update-interval;
no-volume-updates;
}
```

Hierarchy Level

```
[edit security]
```

Description

Enable application tracking (AppTrack).

After application identification identifies the application, AppTrack collects statistics for the application usage on the device, and when the session closes, AppTrack generates a message that provides the byte and packet counts and duration of the session, and sends details to the host device such as Security Threat Response Manager (STRM). STRM retrieves the data and provides flow-based application visibility details.

Options

first-update Generate application tracking initial message when a session is created. This option overrides the `first-update-interval` option if both are specified.

first-update-interval Interval when the first update message is sent (minutes).

> **NOTE**: The `first-update-interval` setting is disregarded if the `first-update` option is set to log the first message at session start.

- ***minutes*** Maximum number of minutes after session start for the first update message to be sent. This value must be smaller than the `session-update-interval` setting.
  - **Default:** 1

disable Disable application tracking.

session-update-interval Frequency in which application tracking update messages are generated (minutes).

log-session-create Use this configuration statement to enable application tracking session.

log-session-close Use this configuration statement to re-enable AppTrack session after closing the session.

session-update interval ***session-update-interval*** Configure the interval between session update messages for long-lived sessions being monitored by AppTrack. Byte count, packet count, and start and end times are updated and logged when the amount of time between session start or the previous update and the current time exceeds the interval.

no-volume-updates Use this configuration statement to disable AppTrack volume update.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2. Support for `disable` added in Junos OS Release 11.4.

Configuration statement `log-session-create`, `log-session-close`, `session-update-interval` *`session-update-interval`*, and `no-volume-updates` are added in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Example: Configuring AppTrack

application-tracking (Security Zones)

IN THIS SECTION

- Syntax | 2399
- Hierarchy Level | 2400
- Description | 2400
- Required Privilege Level | 2400
- Release Information | 2400

Syntax

```
application-tracking;
```

Hierarchy Level

```
[edit security zones security-zone zone-name]
```

Description

Enable application tracking support for the zone.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Example: Configuring Application Tracking

application-traffic-control

IN THIS SECTION

- Syntax | 2401
- Hierarchy Level | 2401
- Description | 2402
- Options | 2402
- Required Privilege Level | 2402
- Release Information | 2402

Syntax

```
application-traffic-control {
    rate-limiters {
        rate-limiter-name {
            bandwidth-limit value-in-kbps;
            burst-size-limit value-in-bytes;
        }
    }
    rule-sets ruleset-name{
         {
            rule rule-name {
                match {
                    application application-name1;
                    application-any;
                    application-group application-group-name;
                    application-known;
                    application-unknown;
                }
                then {
                    dscp-code-point dscp-value;
                    forwarding-class forwarding-class-name;
                    log;
                    loss-priority [ high | medium-high | medium-low | low ];
                    rate-limit {
                        loss-priority-high;
                        client-to-server rate-limiter-name;
                        server-to-client rate-limiter-name;
                    }
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Mark DSCP values for outgoing packets or apply rate limits based on the specified Layer 7 application types.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring AppTrack

application-traffic-control (Application Services)

IN THIS SECTION

- Syntax | 2403
- Hierarchy Level | 2403
- Description | 2403
- Options | 2403
- Required Privilege Level | 2403
- Release Information | 2403

Syntax

```
application-traffic-control {
    rule-set rule-set-name;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
application-services]
[edit logical-systems logical-system-name security policies from-zone zone-name to-zone zone-
name policy policy-name then permit application-services]
[edit tenants tenant-name security policies from-zone zone-name to-zone zone-name policy policy-
name then permit application-services]
```

Description

Enables AppQoS, application-aware quality of service, as specified in the rules of the specified rule set.

Options

- `rule-set` *`rule-set-name`*—Name of the rule set that contains application-aware traffic control specification rules.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the following hierarchy levels introduced in Junos OS Release 19.3R1: `[edit logical-systems` *`logical-system-name`* `security policies from-zone` *`zone-name`* `to-zone` *`zone-name`* `policy` *`policy-name`* `then permit application-services]`, and `[edit tenants` *`tenant-name`* `security policies from-zone` *`zone-name`* `to-zone` *`zone-name`* `policy` *`policy-name`* `then permit application-services]`.

RELATED DOCUMENTATION

Example: Configuring Application Quality of Service

Security Policies Overview

application-unknown

IN THIS SECTION

- Syntax | 2404
- Hierarchy Level | 2404
- Description | 2404
- Required Privilege Level | 2404
- Release Information | 2405

Syntax

```
application-unknown
```

Hierarchy Level

```
[edit services aacl rule rule-name No Link Title term-name No Link Title]
```

Description

Enable AACL logging of flows for unknown applications.

Required Privilege Level

Release Information

Statement introduced in Junos OS Release 11.2.

applications (Applications)

IN THIS SECTION

- Syntax | 2405
- Hierarchy Level | 2405
- Description | 2405
- Options | 2406
- Required Privilege Level | 2406
- Release Information | 2406

Syntax

```
applications {
    application application-name {
        destination-port port-value;
        protocol (tcp | udp);
        ether-type type;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Define applications that DCBX advertises.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- *Defining an Application for DCBX Application Protocol TLV Exchange*
- *Example: Configuring DCBX Application Protocol TLV Exchange*
- Example: Configuring DCBX to Support an iSCSI Application
- *Understanding DCBX Application Protocol TLV Exchange*
- Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

applications (DCBX)

IN THIS SECTION

- Syntax | 2407
- Hierarchy Level | 2407
- Description | 2407
- Options | 2407
- Required Privilege Level | 2407
- Release Information | 2407

Syntax

```
applications {
    fcoe {
        no-auto-negotiation;
    }
}
```

Hierarchy Level

```
[edit protocols dcbx interface interface-name]
```

Description

Configure Data Center Bridging Capability Exchange protocol (DCBX) applications on an interface.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show dcbx neighbors | **19191**

Understanding DCB Features and Requirements

applications (Diameter)

IN THIS SECTION

- Syntax | 2408
- Hierarchy Level | 2408
- Description | 2408
- Options | 2408
- Required Privilege Level | 2409
- Release Information | 2409

Syntax

```
applications {
    pcc-gx {
        maximum-pending-requests requests;
    }
}
```

Hierarchy Level

```
[edit access diameter]
```

Description

Configure the parameters for Diameter applications. Specify the Diameter application for which you are configuring the parameters. The Gx application (`pcc-gx`) is currently supported.

Options

pcc-gx Use the parameters for the Gx application.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

applications (IDS MS-DPC)

IN THIS SECTION

- Syntax | 2409
- Hierarchy Level | 2409
- Description | 2410
- Options | 2410
- Required Privilege Level | 2410
- Release Information | 2410

Syntax

```
applications [ application-name ];
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name from]
```

Description

Define one or more applications to which IDS applies when using the MS-DPC.

Options

application-name—Name of the target application.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Match Conditions in IDS Rules

applications (PCC Rules)

IN THIS SECTION

- Syntax | 2411
- Hierarchy Level | 2411
- Description | 2411
- Options | 2411
- Required Privilege Level | 2411
- Release Information | 2412

Syntax

```
applications [application-name];
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rule-name from],
[edit services pcef pcc-rules rule-name from]
```

Description

Specify one or more applications to define the match criteria for the policy and charging control (PCC) rule. You can specify a maximum of 10 applications in a PCC rule.

> **NOTE**: You must also include the `flows` statement. If you do not want to filter subscriber traffic based on service data flow filters, use `flows any`.

If you are using Junos OS Subscriber Aware, specify the name of the applications at the `[edit unified-edge pcef pcc-rules rule-name from]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the applications at the `[edit services pcef pcc-rules rule-name from]` hierarchy level.

Options

application-name Name of one or more applications that is used to detect IP packet flows.

- **Range:** 1 through 63 characters.

> **NOTE**: The referenced application must have been previously configured in the `[edit services application-identification]` hierarchy level.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-rules rule-name from]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Rules

applications (PCP)

IN THIS SECTION

- Syntax | **2413**
- Hierarchy Level | **2413**
- Description | **2413**
- Options | **2413**
- Required Privilege Level | **2413**
- Release Information | **2413**

Syntax

```
applications [ application-name ];
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name from]
```

Description

Define one or more application protocols to which the PCP rule applies. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

application-name	Name of the application.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

applications (Services AACL)

IN THIS SECTION

- Syntax | 2414
- Hierarchy Level | 2414
- Description | 2414
- Options | 2414
- Required Privilege Level | 2414
- Release Information | 2415

Syntax

```
applications [ application-names ];
```

Hierarchy Level

```
[edit services aacl rule rule-name term term-name from]
```

Description

Identify one or more applications defined in the application identification configuration for inclusion as a match condition.

Options

application-names—Identifiers of the applications.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Match Conditions in AACL Rules

applications (Services ALGs)

IN THIS SECTION

- Syntax | 2415
- Hierarchy Level | 2415
- Description | 2415
- Required Privilege Level | 2416
- Release Information | 2416

Syntax

```
applications { ... }
```

Hierarchy Level

```
[edit]
```

Description

Define the applications used in services.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

ALG Descriptions

Configuring Application Sets

Configuring Application Properties

Examples: Configuring Application Protocols

Verifying the Output of ALG Sessions

applications (Services Application Identification)

IN THIS SECTION

- Syntax | 2417
- Hierarchy Level | 2417
- Description | 2417
- Options | 2417
- Required Privilege Level | 2417
- Release Information | 2417

Syntax

```
applications {
    application-name;
}
```

Hierarchy Level

```
[edit services application-identification application-group group-name]
```

Description

Identify the list of applications for inclusion in the application group.

Options

application-name—Identifier for the application. Maximum length is 32 characters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support added in Junos OS release 19.3R2 and 19.4R1 for Next Gen Services on MX240, MX480, and MX960.

applications (Services CoS)

IN THIS SECTION

- Syntax | 2418
- Hierarchy Level | 2418
- Description | 2418
- Options | 2418
- Required Privilege Level | 2418
- Release Information | 2419

Syntax

```
applications [ application-name ];
```

Hierarchy Level

```
[edit services cos rule rule-name term term-name from]
```

Description

Define one or more applications to which the CoS services apply.

Options

application-name—Name of the target application.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring CoS Rules

Configuring Match Conditions In CoS Rules

applications (Services NAT)

IN THIS SECTION

- Syntax | 2419
- Hierarchy Level | 2419
- Description | 2419
- Options | 2420
- Required Privilege Level | 2420
- Release Information | 2420

Syntax

```
applications [ application-name ];
```

Hierarchy Level

```
[edit  services nat rule rule-name term term-name from]
```

Description

Define one or more application protocols to which the NAT services apply.

Options

application-name—Name of the target application.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

applications (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 2420
- Hierarchy Level | 2421
- Description | 2421
- Options | 2421
- Required Privilege Level | 2421
- Release Information | 2421

Syntax

```
applications application-name;
```

Hierarchy Level

```
[edit services (Stateful Firewall) stateful-firewall rule (Services Stateful Firewall)  rule-name term
(Services Stateful Firewall)  term-name from (Services Stateful Firewall) ]
```

Description

Define one or more applications to which the stateful firewall services apply.

Options

application-name—Name of the target application.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Match Conditions in Stateful Firewall Rules

apply-action-profile

IN THIS SECTION

- Syntax | 2422
- Hierarchy Level | 2422
- Description | 2422

Required Privilege Level | 2422
Release Information | 2422

Syntax

```
apply-action-profile profile-name;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface]
```

Description

Apply the specified action profile to the interface for link-fault management.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Applying an Action Profile

apply-groups

IN THIS SECTION

- Syntax | 2423
- Hierarchy Level | 2423
- Description | 2423
- Options | 2423
- Required Privilege Level | 2424
- Release Information | 2424

Syntax

```
apply-groups [ group-names ];
```

Hierarchy Level

All hierarchy levels

Description

Apply a configuration group to a specific hierarchy level in a configuration, to have a configuration inherit the statements in the configuration group.

You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.

Options

group-names One or more names specified in the `groups` statement.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying a Configuration Group

groups | 7916

apply-groups (Chassis Cluster)

IN THIS SECTION

- Syntax | 2424
- Hierarchy Level | 2425
- Description | 2425
- Options | 2425
- Required Privilege Level | 2425
- Release Information | 2425

Syntax

```
apply-groups [$node]
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Apply node-specific parameters to each node in a chassis cluster.

You can use the `[edit groups]` option to create a unique configuration between the cluster members by defining a unique hostname for each node and assigning a unique IP address for the fxp0 interface on each node.

Each group is named after the node it is applied to (node 0 and node 1) and once you apply the configured groups using the `apply-groups` statement, only the group that matches the node name is applied.

The configuration you specified under group node 0 will be active only on node 0 and the configuration you specified under group node 1 will be active only on node 1.

Options

${node} Each node (node0 or node1) in a chassis cluster.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

cluster (Chassis) | **4369**

apply-groups-except

IN THIS SECTION

- Syntax | 2426
- Hierarchy Level | 2426
- Description | 2426
- Options | 2426
- Required Privilege Level | 2426
- Release Information | 2427

Syntax

```
apply-groups-except [ group-names ];
```

Hierarchy Level

All hierarchy levels except the top level

Description

Disable inheritance of a configuration group.

Options

group-names One or more names specified in the `groups` statement.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

groups | 7916

Disabling Inheritance of a Configuration Group

apply-groups-except (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 2427
- Hierarchy Level | 2427
- Description | 2428
- Options | 2428
- Required Privilege Level | 2428
- Release Information | 2428

Syntax

```
apply-groups-except group-name;
```

Hierarchy Level

```
[edit services radius-flow-tap policy policy-name inet drop-policy rule-name from],
[edit services radius-flow-tap policy policy-name inet6 drop-policy rule-name from]
```

Description

Specify groups from which to inherit configuration data for the radius-flow-tap policy.

Options

group-name— Name of the group that does not inherit the configuration data.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Subscriber Secure Policy Overview

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

apply-groups (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 2429
- Hierarchy Level | 2429
- Description | 2429
- Options | 2429
- Required Privilege Level | 2429
- Release Information | 2429

Syntax

```
apply-groups group-name;
```

Hierarchy Level

```
[edit services radius-flow-tap policy policy-name  inet drop-policy rule-name from],
[edit services radius-flow-tap policy policy-name inet6 drop-policy rule-name from]
```

Description

Specify groups from which to inherit configuration data for the radius-flow-tap policy.

Options

group-name— Name of the group that inherits the configuration data.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

> *Subscriber Secure Policy Overview*
>
> *Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview*

apply-macro

IN THIS SECTION

- Syntax | 2430
- Hierarchy Level | 2430
- Description | 2430
- Options | 2431
- Required Privilege Level | 2431
- Release Information | 2431

Syntax

```
apply-macro apply-macro-name {
    parameter-name parameter-value;
}
```

Hierarchy Level

All hierarchy levels

Description

Use custom syntax in your configuration, which is expanded into standard Junos OS configuration statements during the commit process using the instructions defined in the corresponding commit script macro.

Macros work by locating `apply-macro` statements that you include in the *candidate configuration* and using the values specified in the `apply-macro` statement as parameters to a set of instructions (the macro) defined in a commit script. The commit script alters your configuration from one that contains custom syntax into a full configuration containing standard Junos OS statements.

In effect, your custom configuration syntax serves a dual purpose. The syntax allows you to simplify your configuration tasks, and it provides data (or *hooks*) that are used by commit script macros.

You can include the `apply-macro` statement at any level of the configuration hierarchy. You can include multiple `apply-macro` statements at each level of the configuration hierarchy; however, each must have a unique name.

Options

apply-macro-name Name of the `apply-macro` statement.

parameter-name One or more parameters. Parameters can be any text you want to include in your configuration.

parameter-value A value that corresponds to the parameter name. Parameter values can be any text you want to include in your configuration.

Required Privilege Level

configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Overview of Creating Custom Configuration Syntax with Commit Script Macros

apply-path

IN THIS SECTION

- Syntax | 2432
- Hierarchy Level | 2432
- Description | 2432
- Options | 2432

Required Privilege Level | 2433
Release Information | 2433

Syntax

```
apply-path path;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options prefix-list name],
[edit policy-options prefix-list name]
```

Description

Expand a prefix list to include all prefixes pointed to by a defined path. Using the `apply-path` statement eliminates most of the effort required to maintain a group prefix list. You can use this expanded list in defining policies and firewalls.

You can also use the `apply-path` configuration statement with SNMP, as long as the path uses the following configuration: `apply-path "policy-options prefix-list list-name <*>"`. This is the only `apply-path` configuration originally supported for SNMP. As of Junos OS Release 16.1, support for one other SNMP `apply-path` configuration is added: `apply-path "access radius-server <*>"`.

> **NOTE**: The wildcard regex format of the type `" et<*> "` for `apply-path` configuration statement at the `[edit policy-options prefix-list]` hierarchy is not supported.

Options

path—String of elements composed of identifiers or configuration keywords that points to a set of prefixes. Use a space to delimit multiple elements. As long as the last path element is not a leaf, that is, terminal, statement, you can include wildcards (enclosed in angle brackets) to match more than one identifier. No elements can be added after a leaf statement, including wildcards.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for `apply-path "access radius-server <*>"` configuration for SNMP added in Junos OS Release 16.1.

Support for the `prefix-list` match condition added in Junos OS Evolved Release 20.2R1 for PTX10003.

RELATED DOCUMENTATION

Configuring Prefix Lists

Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List

apply-secondary-as-giaddr

IN THIS SECTION

- Syntax | 2433
- Hierarchy Level | 2434
- Description | 2434
- Required Privilege Level | 2434
- Release Information | 2435

Syntax

```
apply-secondary-as-giaddr;
```

Hierarchy Level

For platforms without ELS:

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface]
```

For platforms with ELS:

```
[edit forwarding-options dhcp-relay overrides]
[edit forwarding-options dhcp-relay group name interface name overrides]
```

Description

Configures the interfaces on a switch that are DHCP relay agents to be enabled for smart DHCP relay. Smart DHCP relay enables you to configure alternative IP addresses for the gateway interface so that if the server fails to reply to the requests sent from the primary gateway address, the switch can resend the requests using the alternative gateway addresses. To use this feature, you must configure an IRB interface or Layer 3 subinterface with multiple IP addresses and configure that interface to be a relay agent. This feature is not supported for EVPN-VXLAN deployments.

Depending on where you configure this statement, it enables smart relay on all the interfaces that are relay agents or on specific interfaces. To enable smart relay on all the interfaces that are relay agents, configure this statement at the following locations:

- Junos OS with ELS: configure it under `edit forwarding-options dhcp-relay overrides`.
- Junos OS without ELS: configure it directly under the **bootp** statement.

To enable smart relay on specific interfaces that are relay agents, configure this statement at the following locations:

- Junos OS with ELS: configure it under `edit forwarding-options dhcp-relay group` *`name`* `interface` *`name`* `overrides`.
- Junos OS without ELS: configure it directly under the **interface** statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D26.

RELATED DOCUMENTATION

Configuring DHCP and BOOTP Relay

apply-to-half-close-state

IN THIS SECTION

- Syntax | 2435
- Hierarchy Level | 2435
- Description | 2435
- Required Privilege Level | 2436
- Release Information | 2436

Syntax

```
apply-to-half-close-state;
```

Hierarchy Level

```
[edit security flow tcp-session time-wait-state]
```

Description

Configure the TCP session timeout in a half-closed state. This enables the system to apply the configured session timeout on receiving only one FIN packet (either client-to-server or server-to-client). When this is not configured, the default behavior takes effect—applying the configured TCP session timeout on receiving both the FIN packets. The default session timeout remains 150 seconds.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

apts

IN THIS SECTION

- Syntax | 2436
- Hierarchy Level | 2437
- Description | 2437
- Options | 2437
- Release Information | 2437

Syntax

```
apts {
    domain domain | apply-groups | apply-groups-except
}
```

Hierarchy Level

```
[edit logical-systems name protocols ptp],
[edit protocols ptp]
```

Description

Configure the PTP assisted partial timing support (APTS) related parameters.

Options

domain

Domain used for APTS profile

- Range: 44 through 63
- Default: 44

> **NOTE**: If the domain is not specified, then the default is applied.

apply-groups

Groups to inherit configuration data from

apply-groups-except

Groups to **not** inherit configuration data from

Release Information

Statement introduced in Junos OS Evolved Release 24.2R2.

RELATED DOCUMENTATION

No Link Title

No Link Title

archival

IN THIS SECTION

- Syntax | 2438
- Hierarchy Level | 2438
- Description | 2439
- Options | 2439
- Required Privilege Level | 2440
- Release Information | 2440

Syntax

```
archival {
    configuration {
        archive-sites {
            file://<path>/<filename>;
            ftp://username@host<:port>url-path password password;
            http://username@host<:port>url-path password password;
            pasvftp://username@host<:port>url-path password password;
            scp://username@host<:port>url-path password password;
        }
        transfer-interval interval;
        transfer-on-commit;
    }
    routing-instance routing-instance;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure copying of the currently active configuration to an archive site. An archive site can be a file, or an FTP, HTTP, passive FTP, or SCP location.

Options

configuration

Configure the router or switch to periodically transfer its currently active configuration (or after each commit). Parameters include `archive-sites`, `transfer-interval`, and `transfer-on-commit`.

> **NOTE**: The `[edit system archival]` hierarchy is not available on QFabric systems.

archive-sites

Specify where to transfer the current configuration files. When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (" ") and enclose the IPv6 host address in brackets ([ ]). For example: **"scp:// *username*<:*password*>@[*ipv6-host-address*]<:*port*>/ *url-path*"**.

If you specify more than one archive site, the router or switch attempts to transfer the configuration files to the first archive site in the list, moving to the next only if the transfer fails. The destination filename is saved in the following format, where *n* corresponds to the number of the compressed configuration rollback file that has been archived:

```
router-name_YYYYMMDD_HHMMSS_juniper.conf.n.gz
```

The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router or switch is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.

transfer-interval

The frequency, in minutes, for transferring the current configuration to an archive site. Valid intervals are 15 to 2880 minutes.

transfer-on-commit

Configure the router or switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration.

routing-instance

Defines the routing instance through which a server is reachable.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Backing Up Configurations to an Archive Site

archive (All System Log Files)

IN THIS SECTION

- Syntax | 2440
- Hierarchy Level | 2441
- Description | 2441
- Options | 2441
- Required Privilege Level | 2442
- Release Information | 2442

Syntax

```
archive <files number> <size size> <start-time time> <transfer-interval interval>
<binary-data | no-binary-data> <world-readable | no-world-readable>;
```

Hierarchy Level

```
[edit system syslog file filename]
```

Description

Configure archiving properties for all system log files.

Options

`files number`—Maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file ***logfile***, it closes the file, compresses it, and renames it ***logfile*.0.gz** (the amount of data is determined by the `size` statement at this hierarchy level). The utility then opens and writes to a new file called ***logfile***. When the new file reaches the maximum size, the ***logfile*.0.gz** file is renamed to ***logfile*.1.gz**, and the new file is closed, compressed, and renamed ***logfile*.0.gz**. By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).

- **Range:** 1 through 1000
- **Default:** 10 files

`size size`—Maximum amount of data that the Junos OS logging utility writes to a log file ***logfile*** before archiving it (closing it, compressing it, and changing its name to ***logfile*.0.gz**). The utility then opens and writes to a new file called ***logfile***.

- **Syntax:** *x* `k` to specify the number of kilobytes, *x* `m` for the number of megabytes, or *x* `g` for the number of gigabytes
- **Range:** 64 KB through 1 GB

Default:

- - 128 KB for EX Series switches
  - 1 MB for M Series, MX Series, and T Series routers, OCX Series, and the QFX3500 switch, and SRX Series Firewalls and vSRX Virtual Firewall.
  - 10 MB for TX Matrix and TX Matrix Plus routers

`binary-data | no-binary-data`—Mark file as containing binary data. This allows proper archiving of binary files, such as WTMP files (login records for UNIX based systems)..

- **Default:** `no-binary-data`

`world-readable | no-world-readable`—Grant all users permission to read archived log files, or restrict the permission only to the `root` user and users who have the Junos OS `maintenance` permission.

- **Default:** `no-world-readable`

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Specifying Log File Size, Number, and Archiving Properties

archive (Individual System Log File)

IN THIS SECTION

- Syntax | 2443
- Hierarchy Level | 2443
- Description | 2443
- Options | 2443
- Required Privilege Level | 2444
- Release Information | 2444

Syntax

```
archive <archive-sites (ftp-url  <password password>)> <binary-data binary-data> <files number>
<no-binary-data  no-binary-data> <size size> <start-time "YYYY-MM-DD.hh:mm"> <transfer-
interval minutes> <world-readable | no-world-readable>;
```

Hierarchy Level

```
[edit system syslog file filename]
```

Description

Configure archiving properties for a specific system log file.

Options

archive-sites ***site-name***—FTP URL representing the destination for the archived log file (for information about how to specify valid FTP URLs, see Format for Specifying Filenames and URLs in Junos OS CLI Commands). If more than one site name is configured, a list of archive sites for the system log files is created. When a file is archived, the router attempts to transfer the file to the first URL in the list, moving to the next site only if the transfer does not succeed. The log file is stored at the archive site with the filename specified at the **[edit system syslog]** hierarchy level.

files ***number***—Maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file ***logfile***, it closes the file, compresses it, and renames it ***logfile*.0.gz** (the amount of data is determined by the **size** statement at this hierarchy level). The utility then opens and writes to a new file called ***logfile***. When the new file reaches the maximum size, the ***logfile*.0.gz** file is renamed to ***logfile*.1.gz**, and the new file is closed, compressed, and renamed ***logfile*.0.gz**. By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).

- **Range:** 1 through 1000
- **Default:** 10 files

password ***password***—Password for authenticating with the site specified by the **archive-sites** statement.

binary-data | no-binary-data—Mark file as containing binary data. This allows proper archiving of binary files, such as WTMP files (login records for UNIX based systems)..

- **Default:** `no-binary-data`

size ***size***—Maximum amount of data that the Junos OS logging utility writes to a log file ***logfile*** before archiving it (closing it, compressing it, and changing its name to ***logfile*.0.gz**). The utility then opens and writes to a new file called ***logfile***.

- **Syntax:** ***x*k** to specify the number of kilobytes, ***x*m** for the number of megabytes, or ***x*g** for the number of gigabytes
- **Range:** 64 KB through 1 GB
- **Default:** 128 KB for J Series routers; 1 MB for M Series, MX Series, and T Series routers, and the QFX3500 switch; 10 MB for TX Matrix and TX Matrix Plus routers

start-time "*YYYY-MM-DD.hh:mm*"—Date and time in the local time zone for a one-time transfer of the active log file to the first reachable site in the list of sites specified by the **archive-sites** statement.

transfer-interval ***interval***—Interval at which to transfer the log file to an archive site.

- **Range: 5** through **2880** minutes

world-readable | no-world-readable—Grant all users permission to read archived log files, or restrict the permission only to the **root** user and users who have the Junos OS **maintenance** permission.

- **Default: no-world-readable**

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

start-time and **transfer-interval** statements introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Specifying Log File Size, Number, and Archiving Properties

archive (QFabric System)

IN THIS SECTION

- Syntax | 2445
- Hierarchy Level | 2445
- Description | 2445
- Options | 2445
- Required Privilege Level | 2446
- Release Information | 2446

Syntax

```
archive {
    size size;
}
```

Hierarchy Level

```
[edit system syslog file filename]
```

Description

Configure the archiving properties for the system message log file.

Options

size *size* Maximum amount of system log message data that the QFabric system stores in the log file.

- **Syntax:** *x*k to specify the number of kilobytes, *x*m for the number of megabytes, or *x*g for the number of gigabytes
- **Range:** 65 KB through 1 GB

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

syslog (QFabric System) | 28825

archive-sites

IN THIS SECTION

- Syntax | 2446
- Hierarchy Level | 2447
- Description | 2447
- Options | 2447
- Required Privilege Level | 2447
- Release Information | 2447

Syntax

```
archive-sites site-name {
    password password;
}
```

Hierarchy Level

```
[edit accounting-options file filename]
```

Description

Configure an archive site. If more than one site name is configured, an ordered list of archive sites for the accounting-data log files is created. When a file is archived, the router or switch attempts to transfer the file to the first URL in the list, moving to the next site only if the transfer does not succeed. The log file is stored at the archive site with a filename of the format ***router-name_log-filename_timestamp***. When you configure file archival by using the `archive-states` statement, the transfer file utility uses the default routing instance to connect to the destination server. If the default routing instance is unable to connect to the destination server, the transfer file utility does not work.

Starting in Junos OS Release 18.4R1, when you configure file archival by using the `archive-sites` statement, the transfer file utility does not work if you have enabled the management instance.

To delete an archive site, use the `delete` command instead of `set`.

Options

site-name—Any valid FTP/HTTP/SCP/SFTP URL to a destination.

password *password*—Password for logging into the archive site.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Archive Sites for the Files

archive-sites (Flow Monitoring)

IN THIS SECTION

- Syntax | 2448
- Hierarchy Level | 2448
- Description | 2448
- Required Privilege Level | 2448
- Release Information | 2449

Syntax

```
archive-sites {
    ftp:url {
        password "password";
        username username;
    }
}
```

Hierarchy Level

```
[edit services flow-collector transfer-log-archive]
```

Description

Specify the destination for transfer logs.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Transfer Logs

archive-sites (Event Policy)

IN THIS SECTION

- Syntax | 2449
- Hierarchy Level | 2449
- Description | 2450
- Options | 2450
- Required Privilege Level | 2450
- Release Information | 2450

Syntax

```
archive-sites {
    url <password password>;
}
```

Hierarchy Level

```
[edit event-options destinations destination-name]
```

Description

For a given destination name, configure archive sites to which an event policy can transfer files. To archive files, you must reference the destination name in the event policy. If you specify more than one archive site for a particular destination name, the device attempts to transfer to the first archive site in the list, moving to the next site only if the transfer fails.

Options

url

The archive destination specified as a file URI, an active or passive FTP URI, a Secure FTP (SFTP) URI, or a Secure Copy (SCP) URI. Local device directories are also supported (for example, **/var/tmp/**).

- file:<//host>/path
- ftp://username@host:<port>url-path
- pasvftp://username@host:<port>url-path
- sftp://username@host:<port>url-path
- scp://username@host:<port>url-path
- <path>/<filename>

password *password*

A plain-text password required for logging into the archive site.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Event Policy File Archiving

Define Destinations for Event Policy File Archiving

Configure an Event Policy to Upload Files

destinations (Event Policy) | **5533**

area

IN THIS SECTION

- Syntax | 2451
- Hierarchy Level | 2451
- Description | 2452
- Options | 2452
- Required Privilege Level | 2452
- Release Information | 2453

Syntax

```
area area-id {
    interface interface-name {
        no-eligible-remote-backup;
        passive;
        topology (ipv4-multicast | name) {
            disable;
        }
    }
    virtual-link neighbor-id router-id transit-area area-id {
        topology (ipv4-multicast | name) {
            disable;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Specify the area identifier for this routing device to use when participating in OSPF routing. All routing devices in an area must use the same area identifier to establish adjacencies.

Specify multiple `area` statements to configure the routing device as an area border router. An area border router does not automatically summarize routes between areas. Use the `area-range` statement to configure route summarization. By definition, an area border router must be connected to the backbone area either through a physical link or through a virtual link. To create a virtual link, include the `virtual-link` statement.

To specify that the routing device is directly connected to the OSPF backbone, include the `area 0.0.0.0` statement.

All routing devices on the backbone must be contiguous. If they are not, use the `virtual-link` statement to create the appearance of connectivity to the backbone.

You can also configure any interface that belongs to one or more topologies to advertise the direct interface addresses without actually running OSPF on that interface. By default, OSPF must be configured on an interface in order for direct interface addresses to be advertised as interior routes.

> **NOTE**: If you configure an interface with the `passive` statement, it applies to all the topologies to which the interface belongs. You cannot configure an interface as passive for only one specific topology and have it remain active for any other topologies to which it belongs.

Options

area-id—Area identifier. The identifier can be up to 32 bits. It is common to specify the area number as a simple integer or an IP address. Area number 0.0.0.0 is reserved for the OSPF backbone area.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

Support for the `no-eligible-remote-backup` statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding OSPF Areas

Understanding Multiple Address Families for OSPFv3

virtual-link | 30763

area-range

IN THIS SECTION

- Syntax | 2453
- Hierarchy Level | 2454
- Description | 2454
- Default | 2454
- Options | 2454
- Required Privilege Level | 2455
- Release Information | 2455

Syntax

```
area-range network/mask-length <exact> <override-metric metric> <restrict>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id],
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id nssa],
[edit logical-systems logical-system-name realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)
area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-id nssa],
[edit logical-systems logical-system-name routing-instances routing-instance-name realm (ipv4-
unicast | ipv4-multicast | ipv6-multicast) area area-id],
[edit protocols (ospf | ospf3) area area-id],
[edit protocols (ospf | ospf3) area area-id nssa],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id nssa],
[edit routing-instances routing-instance-name realm (ipv4-unicast | ipv4-multicast | ipv6-
multicast) area area-id]
```

Description

(Area border routers only) For an area, summarize a range of IP addresses when sending summary link advertisements (within an area). To summarize multiple ranges, include multiple `area-range` statements.

For a not-so-stubby area (NSSA), summarize a range of IP addresses when sending NSSA link-state advertisements. The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas. To specify multiple prefixes, include multiple `area-range` statements. All external routes learned within the area that do not fall into one of the prefixes are advertised individually to other areas.

Default

By default, area border routing devices do not summarize routes being sent from one area to other areas, but rather send all routes explicitly.

Options

exact—(Optional) Summarization of a route is advertised only when an exact match is made with the configured summary range.

mask-length—Number of significant bits in the network mask.

network—IP address. You can specify one or more IP addresses.

override-metric ***metric***—(Optional) Override the metric for the IP address range and configure a specific metric value.

restrict—(Optional) Do not advertise the configured summary. This hides all routes that are contained within the summary, effectively creating a route filter.

- **Range:** 1 through 16,777,215

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements

arguments (Event Policy)

IN THIS SECTION

- Syntax | 2456
- Hierarchy Level | 2456
- Description | 2456
- Options | 2456

Required Privilege Level | 2456
Release Information | 2456

Syntax

```
arguments {
    argument-name argument-value;
}
```

Hierarchy Level

```
[edit event-options policy policy-name then event-script filename]
```

Description

Define the arguments that are passed to an event script when it is invoked from an event policy.

Options

argument-name	Name of the argument.
argument-value	String specifying the value of the argument.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configure an Event Policy to Pass Arguments to an Event Script

Execute Event Scripts in an Event Policy

event-script (Event Policy) | 6536

policy (Event Policy) | 12675

arguments (Op Scripts)

IN THIS SECTION

- Syntax | 2457
- Hierarchy Level | 2457
- Description | 2458
- Options | 2458
- Required Privilege Level | 2458
- Release Information | 2458

Syntax

```
arguments {
    argument-name {
        description descriptive-text;
    }
}
```

Hierarchy Level

```
[edit system scripts op file filename]
```

Description

Define the expected command-line arguments for an op script. Once configured, the list of possible arguments will be included in the context-sensitive help in the CLI.

Options

argument-name The name of a command-line argument to an op script.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Declare and Use Command-Line Arguments in Op Scripts

Configure Help Text for Op Scripts

description (Op Scripts) | **5339**

arp

IN THIS SECTION

- Syntax | **2459**
- Syntax (EX Series) | **2459**
- Hierarchy Level | **2459**
- Description | **2460**
- Options | **2460**

Required Privilege Level | 2462
Release Information | 2463

Syntax

```
arp ip-address (mac | multicast-mac) mac-address publish;
```

```
arp {
    aging-timer minutes;
    arp-retries count;
    gratuitous-arp-delayseconds;
    gratuitous-arp-on-ifup;
    interfaces {
        interface-name {
            aging-timer minutes;
        }
    }
    non-subscriber-no-reply;
    passive-learning;
    purging;
    respond-out-of-subnet
    unicast-mode-on-expire;
}
```

Syntax (EX Series)

```
arp {
    aging-timer minutes;
}
```

Hierarchy Level

```
[edit system]
[edit interfaces interface-name unit logical-unit-number family inet address address]
```

```
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet address address]
```

NOTE: The `edit logical-systems` hierarchy is not available on QFabric systems.

Description

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries mapping IP addresses to MAC addresses. IPv4 networks use ARP to map IP network addresses to physical (MAC) addresses. An address is resolved when the host device receives a proper ARP reply in response to the ARP request that it sent to a broadcast Ethernet address. The resolved addresses are stored in the ARP table for a configurable period of time. When an entry is close to expiration, it triggers the host to broadcast another ARP request to update the entry for that address. Only the intended receiver responds to the broadcast request; other recipients silently drop the request packet.

You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.

NOTE: By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the `family inet` statement. By including the `arp` statement at the `[edit interfaces interface-name unit logical-unit-number family inet policer]` hierarchy level, you can apply a specific ARP-packet policer to an interface. This feature is not available on EX Series switches.

When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the `unnumbered-address` statement at the `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy level.

NOTE: For EX-Series switches, set only the time interval between ARP updates.

Options

aging-timer Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, on routers only, metro Ethernet environments), increasing the time between updates can improve system performance.

- **Range:** 1 through 240

arp-retries *count*

Configures the maximum number of retries the device attempts for ARP. The default retry attempts for ARP is 4.

ip-address

IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing `address` statement.

mac *mac-address*

MAC address to map to the IP address. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn*.*nnnn*.*nnnn* or *nn*:*nn*:*nn*:*nn*:*nn*:*nn*. For example, `0000.5e00.5355` or `00:00:5e:00:53:55`.

multicast-mac *mac-address*

Multicast MAC address to map to the IP address. Specify the multicast MAC address as six hexadecimal bytes in one of the following formats: *nnnn*.*nnnn*.*nnnn* or *nn*:*nn*:*nn*:*nn*:*nn*:*nn*. For example, `0000.5e00.5355` or `00:00:5e:00:53:55`.

publish

(Optional) Have the router or switch reply to ARP requests for the specified IP address. If you omit this option, the router or switch uses the entry to reach the destination but does not reply to ARP requests.

> **NOTE**: For unicast MAC addresses only, if you include the `publish` option, the router or switch replies to proxy ARP requests.

gratuitous-arp-delay

Configure a delay for gratuitous ARP requests at the system level. By default, Junos OS sends gratuitous ARP requests immediately after network-related configuration changes are made on an interface (for example, a VLAN ID, MAC address, IP address change, or Aggregated Ethernet deployment). This might lead to the Packet Forwarding Engine dropping some initial request packets if the configuration updates have not been fully processed. To avoid such request packets being dropped, you can configure a delay in gratuitous ARP requests.

- Values:
  - *seconds*—Configure the ARP request delay in seconds. We recommend configuring a value in the range of 3 through 6 seconds.

gratuitous-arp-on-ifup

Add this statement to the `[edit system arp]` hierarchy to configure Junos OS to automatically issue a gratuitous ARP announcement when an interface is online.

interfaces aging-timer *minutes*

Specify the ARP aging timer in minutes for a logical interface of family type `inet`. The ARP aging timer is the time between ARP updates.

- **Default:** 20
- **Range:** 1 through 600,000

non-subscriber-no-reply Configure the device to reply to ARP requests from subscribers only. Do not reply to ARP requests from non-subscribers.

passive-learning Configure backup VRRP routers or switches to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests. By default, the backup VRRP router drops these requests; therefore, if the primary router fails, the backup router must learn all entries present in the ARP cache of the primary router. Configuring passive learning reduces transition delay when the backup router is activated. Learning of ARP mappings (IP-to-MAC address) by backup VRRP routers or switches for hosts sending the requests is disabled unless this statement is configured.

purging Purge obsolete ARP entries from the cache when an interface or link goes offline.

respond-out-of-subnet Respond to ARP requests from a source that is not directly attached to the local subnet.

> **NOTE**: See PR 1710699 for details on support for the `respond-out-of-subnet` option.

unicast-mode-on-expire Send a unicast ARP request instead of the default broadcast request when an ARP cache entry ages out. When you include this option, the host device sends the requests only to the expected (currently cached) address. The ARP retry requests are unicast at intervals of 5 seconds. When you do not configure the `unicast-mode-on-expire` option, ARP retries are broadcast at intervals of 800 milliseconds.

This option reduces the amount of broadcast traffic normally sent to resolve expiring addresses. It also supports a special use case where access nodes are configured not to forward broadcast ARP requests towards customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests.

> **NOTE**: This option affects only the update requests. Initial ARP requests are broadcast as usual.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 11.1.

Statement introduced in Junos OS Release 14.1X53-D20.

`arp-retries` option introduced in Junos OS Evolved Release 22.1R1.

`arp-retries` option introduced in Junos OS Release 22.1R2.

RELATED DOCUMENTATION

Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses

Configure ARP Learning and Aging Options

Adjusting the ARP Aging Timer

arp-detect

IN THIS SECTION

- Syntax | 2463
- Hierarchy Level | 2464
- Description | 2464
- Options | 2464
- Required Privilege Level | 2464
- Release Information | 2465

Syntax

```
arp-detect milliseconds;
```

Hierarchy Level

```
[edit forwarding-options next-hop]
```

Description

Define the length of time (in milliseconds) for an SPU to wait for an acknowledgement from the Routing Engine that an ARP request has been received from the SPU. If the Routing Engine fails to respond within the specific time interval, the SPU considers that the Routing Engine CPU utilization is high at that moment, and initiates the ARP throttling. ARP throttling is initiated on the logical interface, where the incoming packet had triggered the ARP request.

Configuring a shorter ARP detect time interval results in triggering of ARP throttling more frequently. Frequent ARP throttling is useful for lowering Routing Engine CPU utilization caused by excessive ARP requests.

For example, when you configure the `set forwarding-options nexthop arp-detect 300` option, the nexthop resolution request must be acknowledged by the Routing Engine within 300 milliseconds. If the SPU does not get an acknowledgment from the Routing Engine in 300 milliseconds, the logical interface which had received the packet that triggered the nexthop request, changes into ARP throttle state. While the ARP throttle state is active for that interface, traffic entering into that interface does not trigger new nexthop resolution requests.

> **CAUTION**: We recommend that only advanced Junos OS users attempt to configure ARP throttle and ARP detect feature. Because, improper configuration might result in high CPU utilization of Routing Engine affecting other processes on your device.

Options

milliseconds—Number of seconds the SPU waits before receiving a response from Routing Engine.

- **Range:** 1 through 10000 milliseconds
- **Default:** 10000 milliseconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D65 and Junos OS Release 15.1X49-D130.

RELATED DOCUMENTATION

arp-throttle | 2482

arp-enhanced-scale

IN THIS SECTION

- Syntax | 2465
- Hierarchy Level | 2465
- Description | 2465
- Required Privilege Level | 2466
- Release Information | 2466

Syntax

```
arp-enhanced-scale;
```

Hierarchy Level

```
[edit system]
```

Description

Increases the number of ARP and neighbor discovery entries for MC-LAG configured with enhanced convergence and Layer 3 VXLAN deployments.

NOTE: To increase the ARP and discovery entries for MC-LAG with enhanced convergence, you also need to enable the `enhanced-convergence` statement at the `[edit system]` hierarchy. For information on how to configure enhanced convergence, see *Understanding Multichassis Link Aggregation Groups*.

NOTE: This statement is enabled by default on QFX10002-60C devices.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

enhanced-convergence | 6357

Understanding Multichassis Link Aggregation Groups

arp-inspection

IN THIS SECTION

- Syntax | 2467
- Hierarchy Level | 2467
- Description | 2467
- Default | 2468
- Required Privilege Level | 2468
- Release Information | 2468

Syntax

```
arp-inspection {
    forwarding-class class-name;
}
```

Hierarchy Level

- For platforms with ELS:

```
[edit vlans vlan-name forwarding-options dhcp-security],
[edit forwarding-options dhcp-relay ]
```

- For platforms without ELS:

```
[edit ethernet-switching-options   secure-access-port vlan (all | vlan-name)],
[edit forwarding-options dhcp-relay ]
```

Description

Perform dynamic ARP inspection (DAI) on all VLANs or on the specified VLAN.

When DAI is enabled, the switch logs invalid ARP packets that it receives on each interface, along with the sender's IP and MAC addresses. ARP probe packets, which have the sender IP address 0.0.0.0, are validated by DAI.

NOTE: If you configure DAI at the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level:

- DAI can be configured only for a specific VLAN, not for a list or a range of VLAN IDs.
- DHCP snooping is automatically enabled on the specified VLAN.
- The `forwarding-class` statement is not available at the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level.

See *Enabling Dynamic ARP Inspection (ELS)* for more information about this configuration.

> **NOTE**: On EX9200 switches, DAI is not supported in an MC-LAG scenario.

The remaining statement is explained separately.

Default

Disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

- *Example: Protecting Against Address Spoofing and Layer 2 DoS Attacks*
- *Example: Protecting Against ARP Spoofing Attacks*
- *Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing*
- *Example: Prioritizing Snooped and Inspected Packet*
- *Enabling Dynamic ARP Inspection (non-ELS)*
- Enabling Dynamic ARP Inspection (J-Web Procedure)

arp-inspection (MX Series)

IN THIS SECTION

- Syntax | 2469
- Hierarchy Level | 2469
- Description | 2469
- Default | 2469
- Required Privilege Level | 2469
- Release Information | 2470

Syntax

```
arp-inspection;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]
```

Description

Perform dynamic ARP inspection (DAI).

DAI can only be configured for a specific bridge domain, not for a list or a range of bridge domain names.

DHCP snooping is automatically enabled on the specified VLAN or bridge domain.

Default

Disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Applying CoS Forwarding Classes to Prioritize Inspected Packets

Example: Configuring IP Source Guard and Dynamic ARP Inspection on a Specified Bridge Domain to Protect the Devices Against Attacks

arp-l2-validate

IN THIS SECTION

- Syntax | 2470
- Hierarchy Level | 2470
- Description | 2471
- Required Privilege Level | 2471
- Release Information | 2471

Syntax

```
arp-l2-validate
```

Hierarchy Level

```
[edit interfaces irb]
```

Description

Enables periodic checking of ARP Layer 3 addressing and MAC Layer 2 addressing tables, and fixes entries if they become out of sync.

Normally, the ARP and MAC address tables stay synchronized. However, you can configure this option on the `irb` interface of the switch to help avoid traffic loss in network conditions that might cause unresolved inconsistencies to occur between the ARP and MAC address tables, such as:

- When link flapping occurs in a multichassis link aggregation (MC-LAG) group, and the network is attempting to achieve convergence. In this case, frequent MAC table updates are happening, and occasionally a corresponding ARP table update might be lost.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R4.

arp-max-cache

IN THIS SECTION

- Syntax | 2472
- Hierarchy Level | 2472
- Description | 2472
- Options | 2472
- Required Privilege Level | 2472
- Release Information | 2473

Syntax

```
arp-max-cache arp-max-cache;
```

Hierarchy Level

```
[edit interfaces name unit name family inet]
```

Description

The ARP cache limit for resolved next hops can be configured at an interface level. The benefit of configuring the ARP cache limit is to protect the device from DoS attacks.

> **NOTE**: After modifying the default ARP next-hop cache limit on an interface, you must deactivate and then reactivate the interface for the newly configured values to take effect.

Options

arp-max-cache Indicates the maximum number of routes to be held in the ARP cache.

- **Default:**
  - 20,000 (ACX Series routers, EX2200, EX2200-C, EX3200, and EX3300 switches, SRX Series Firewalls)
  - 75,000 (EX4200, EX4300, EX4500, EX4550, and EX6210 switches, MX Series routers)
  - 100,000 (Other platforms)
- **Range:** 1 through 2,000,000

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Example: Configuring ARP Cache Protection

arp | 2458

arp | 2458

arp-inspection (MX Series)

arp-new-hold-limit | 2473

arp-system-cache-limit | 2480

arp-new-hold-limit

IN THIS SECTION

- Syntax | 2473
- Hierarchy Level | 2474
- Description | 2474
- Options | 2474
- Required Privilege Level | 2474
- Release Information | 2474

Syntax

```
arp-new-hold-limit number;
```

Hierarchy Level

```
[edit interfaces interface-name unit interface-unit-number family inet]
```

Description

The ARP cache limit for unresolved next hops can be configured at an interface level. The benefit of configuring the ARP cache limit is to protect the device from DoS attacks.

Options

number Indicates the new or maximum routes to be held, before getting added to the ARP cache.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Example: Configuring ARP Cache Protection

arp | **2458**

arp-inspection (MX Series)

arp-system-cache-limit | **2480**

arp-max-cache | **2471**

arp policer

IN THIS SECTION

- Syntax | 2475
- Hierarchy Level | 2475
- Description | 2475
- Options | 2476
- Required Privilege Level | 2476
- Release Information | 2476

Syntax

```
policer name {
    filter-specific;
        logical-interface-policer;
        physical-interface-policer;
        if-exceeding {
        (bandwidth-limit bits per second | bandwidth-percent percent);
        burst-size-limit bytes;
    then <discard> <forwarding-class forwarding-class> <loss-priority (high | low | medium-high
| medium-low)> <out-of-profile>;
}
```

Hierarchy Level

```
[edit firewall]
```

Description

Apply a policer to an interface. Configure policer rate limits and actions. The traffic to the Routing Engine is controlled by applying the policer on ARP. This prevents network congestion caused by broadcast storms.

A policer applies two types of rate limits on traffic:

- Bandwidth—The number of bits per second permitted, on average
- Maximum burst size—The maximum size permitted for bursts of data that exceed the given bandwidth limit

Options

name	Policer name.
filter-specific	Policer is filter-specific.
logical-interface-policer	Policer is a logical interface policer.
physical-interface-policer	Policer is a physical interface policer.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

- bandwidth-limit (Hierarchical Policer)
- burst-size-limit (Hierarchical Policer)

arp-prefix-limit (Host Fast Reroute)

IN THIS SECTION

- Syntax | 2477
- Hierarchy Level | 2477
- Description | 2477
- Default | 2478
- Options | 2478
- Required Privilege Level | 2478
- Release Information | 2479

Syntax

```
arp-prefix-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
interface interface-name],
[edit routing-instances instance-name routing-options interface interface-name]
```

Description

Override the ARP prefix limit for the specified host fast-reroute (HFRR) profile.

When you configure HFRR, an optional ARP prefix limit sets a maximum for the number of ARP routes and, therefore FRR routes created for each HFRR profile in the routing table. This limit prevents ARP attacks from exhausting the virtual memory on the routing devices.

There are two configuration statements (`global-arp-prefix-limit` and `arp-prefix-limit`) that set the ARP prefix limit, one at the global `[edit routing-options host-fast-reroute]` hierarchy level and the other at the `[edit routing-instances instance-name routing-options interface interface-name]` hierarchy level, respectively. The global `global-arp-prefix-limit` statement sets a default ARP prefix limit for all HFRR profiles

configured on the routing device. The `arp-prefix-limit` statement overrides the `global-arp-prefix-limit` for that HFRR profile for that protected interface.

Warning system log messages begin when the ARP routes in an HFRR profile reaches 80% of the configured limit. When the number crosses the 100% threshold, the HFRR profile is deactivated. When this happens, all ARP/FRR routes are deleted from the routing table. FRR routes are deleted from forwarding table as well.

After the HFRR profile is deactivated, a blackout timer is started. The timeout value of this timer is the ARP cache timeout (kernel timeout) + the supplementary blackout timer.

There are global and per-HFRR CLI statements (`global-supplementary-blackout-timer` and `supplementary-blackout-timer`) to configure the supplementary blackout timer. The global value is at the `[edit routing-options host-fast-reroute]` hierarchy level and applies to all HFRR profiles on the routing device. The value for the routing-instance interface is at the `[edit routing-instances` *`instance-name`* `routing-options interface` *`interface-name`*`]` hierarchy level, and overrides the global value for that HFRR profile only.

When the blackout timer expires, the HFRR profile is reactivated, and the Junos OS relearns the ARP routes and re-creates the HFRR routes. If the ARP prefix limit is not exceeded again, the HFRR routes will be up.

If an HFRR profile is in the deactivated state, a reevaluation of the ARP state is preformed during every commit operation or whenever the routing process (rpd) is restarted with the `restart routing` command.

Default

If you omit the-`arp-prefix-limit` statement, the `global-arp-prefix-limit` takes effect for all HFRR profiles on the device. If you omit both of these statements, there is no ARP prefix limit for host fast reroute.

Options

number Maximum number of ARP HFRR routes allowed.

- **Range:** 1 through 10,000 HFRR routes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring Link Protection with Host Fast Reroute

arp-resp

IN THIS SECTION

- Syntax | 2479
- Hierarchy Level | 2479
- Description | 2479
- Options | 2480
- Required Privilege Level | 2480
- Release Information | 2480

Syntax

```
arp-resp (restricted|unrestricted);
```

Hierarchy Level

```
[edit interfaces interfaces-name unit logical-unit-number ]
```

Description

Configure Address Resolution Protocol (ARP) response on the interface.

Options

- `restricted`—Enable restricted proxy ARP response on the interface. This is the default.
- `unrestricted`—Enable unrestricted ARP response on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

arp-system-cache-limit

IN THIS SECTION

- Syntax | 2480
- Hierarchy Level | 2481
- Description | 2481
- Options | 2481
- Required Privilege Level | 2481
- Release Information | 2481

Syntax

```
arp-system-cache-limit number;
```

Hierarchy Level

```
[edit system]
```

Description

Specify the ARP cache next-hop limit at the system (global) level to restrict the number of next-hop routes. To configure the limit at the interface level, use the "arp-max-cache" on page 2471 statement.

The default behavior of ARP is to remove the cache entry if it is not used within a certain period of time, and not to allow the cache to grow too large. You can also manage the number of ARP cache next-hop entries by configuring a limit to the maximum number of next hops that can be created.

Options

number Indicates the maximum number of routes to be held in the ARP cache.

- **Default:**
  - 20,000 (ACX Series routers, EX2200, EX2200-C, EX3200, and EX3300 switches, SRX Series Firewalls)
  - 75,000 (EX4200, EX4300, EX4500, EX4550, and EX6210 switches, MX Series routers)
  - 100,000 (Other platforms)
- **Range:** 1 through 2,000,000

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Example: Configuring ARP Cache Protection

arp | 2458
arp-inspection (MX Series)
arp-max-cache | 2471
arp-new-hold-limit | 2473

arp-throttle

IN THIS SECTION

- Syntax | 2482
- Hierarchy Level | 2482
- Description | 2482
- Options | 2483
- Required Privilege Level | 2483
- Release Information | 2483

Syntax

```
arp-throttle seconds;
```

Hierarchy Level

```
[edit forwarding-options next-hop]
```

Description

Define the time duration (in seconds) for Address Resolution Protocol (ARP) request throttling to remain active when it is triggered.

When ARP throttling is triggered, it is active for a given logical interface. For the configured duration of time, the ARP throttling remains active, and the traffic entering into the specific interface does not

trigger ARP nexthop resolution requests, which are being sent to the Routing Engine from the specific SPU.

When you configure a longer time duration, the ARP throttling can protect the Routing Engine by preventing too many ARP requests being triggered by incoming traffic.

For example, if there is a large amount of traffic destined to a directly connected, unresolved IP address, chances of getting frequent ARP requests is very high, which eventually results into a high CPU load on the Routing Engine. By setting a longer time interval of the ARP throttle, the Routing Engine is protected from numerous ARP requests.

For example, when you configure the `set forwarding-options nexthop arp-throttle 15` option, and the interface state changes to throttle state, the nexthop requests triggered by incoming traffic into this interface, will not be sent. After 15 seconds, interface changes back from throttle to normal state. Then the nexthop requests triggered by the incoming traffic into this interface are sent to the Routing Engine again.

Options

seconds—Time interval (in seconds) for Address Resolution Protocol (ARP) request throttling to remain active when it is triggered.

- **Range:** 10 through 100 seconds
- **Default:** 10 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D65, Junos OS Release 15.1X49-D60, and Junos OS Release 15.1X49-D130.

RELATED DOCUMENTATION

arp-detect | **2463**

arp-nd-probe-disable

IN THIS SECTION

- Syntax | 2484
- Hierarchy Level | 2484
- Description | 2484
- Required Privilege Level | 2484
- Release Information | 2485

Syntax

```
arp-nd-probe-disable;
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

Disable the default probing that happens when you provision static MAC-IP entries.

This statement ensures that probing is disabled, and the PE router acquires both the MAC address and MAC-IP binding information from the network traffic and communicates the information using EVPN.

Required Privilege Level

[none specified]

Release Information

Statement introduced in Junos OS Release 23.4R1.

arp-nd-probe-failed-log

IN THIS SECTION

- Syntax | 2485
- Hierarchy Level | 2485
- Description | 2485
- Required Privilege Level | 2485
- Release Information | 2486

Syntax

```
arp-nd-probe-failed-log;
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

Enable syslog when the IP address probe fails.

Required Privilege Level

[none specified]

Release Information

Statement introduced in Junos OS Release 23.4R1.

as-external

IN THIS SECTION

- Syntax | 2486
- Hierarchy Level | 2486
- Description | 2486
- Required Privilege Level | 2487
- Release Information | 2487

Syntax

```
as-external;
```

Hierarchy Level

```
[edit logical-systems name protocols ospf overload],
[edit logical-systems name routing-instances name protocols ospf overload],
[edit protocols ospf overload],
[edit routing-instances name protocols ospf overload],
[edit protocols ospf3 overload]
```

Description

Advertise OSPF AS external prefixes with maximum usable metric.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2.

RELATED DOCUMENTATION

Understanding OSPF Overload Function

allow-route-leaking | 2202

stub-network | 28667

intra-area-prefix | 9224

as-list

IN THIS SECTION

- Syntax | 2487
- Hierarchy Level | 2488
- Description | 2488
- Options | 2488
- Required Privilege Level | 2488
- Release Information | 2488

Syntax

```
as-list name {
    members [ members ... ];
}
```

Hierarchy Level

```
[edit logical-systems name policy-options],
[edit policy-options]
```

Description

Define a policy to specify a list of BGP autonomous systems (AS) to configure acceptable AS ranges for EBGP groups that can be used for bringing up BGP peers while establishing a BGP session. BGP accepts a peer request based on the specified AS range and rejects a peer request if the AS does not fall into the specified range. This allows you to control BGP peering when the neighbor's exact IP address is not known.

Options

name Specify a name to identify a BGP as-list.

members Specify a single autonomous system number or a range of autonomous system numbers in plain numbers.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

peer-as-list | 12345

as-override

IN THIS SECTION

- Syntax | 2489
- Hierarchy Level | 2489
- Description | 2489
- Required Privilege Level | 2490
- Release Information | 2490

Syntax

```
as-override;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Compare the AS path of an incoming advertised route with the AS number of the BGP peer under the group and replace all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer.

> **NOTE**: The `as-override` statement is specific to a particular BGP group. This statement does not affect peers from the same remote AS configured in different groups.

Enabling the AS override feature allows routes originating from an AS to be accepted by a router residing in the same AS. Without AS override enabled, the routing device refuses the route advertisement once the AS path shows that the route originated from its own AS. This is done by default to prevent route loops. The `as-override` statement overrides this default behavior.

Note that enabling the AS override feature may result in routing loops. Use this feature only for specific applications that require this type of behavior, and in situations with strict network control. One application is the IGP protocol between the provider edge routing device and the customer edge routing device in a virtual private network.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring a Layer 3 VPN with Route Reflection and AS Override

Junos OS VPNs Library for Routing Devices

as-path-compare

IN THIS SECTION

- Syntax | **2491**
- Hierarchy Level | **2491**
- Description | **2491**

Required Privilege Level | 2491
Release Information | 2491

Syntax

```
as-path-compare;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multipath],
[edit routing-instances routing-instance-name routing-options multipath]
```

Description

Specify to have the algorithm that is used to determine the active path compare the AS numbers in the AS path. In a VPN scenario with multiple BGP paths, the algorithm selects as the active path the route whose AS numbers match. By default, the algorithm evaluates only the length and not the contents of the AS path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *Configuring the Algorithm That Determines the Active Route to Evaluate AS Numbers in AS Paths for VPN Routes*

as-path-group

IN THIS SECTION

- Syntax | 2492
- Hierarchy Level | 2492
- Description | 2492
- Options | 2492
- Required Privilege Level | 2493
- Release Information | 2493

Syntax

```
as-path-group group-name {
    as-path name regular-expression;
}
```

Hierarchy Level

```
[edit dynamic policy-options],
[edit logical-systems logical-system-name policy-options],
[edit policy-options]
```

Description

Define a group containing multiple AS path regular expressions for use in a routing policy match condition.

Options

group-name—Name that identifies the AS path group. One or more AS path regular expressions must be listed below the `as-path-group` hierarchy.

name—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

regular-expression—One or more regular expressions used to match the AS path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for dynamic database configuration introduced in Junos OS Release 9.5.

Support for dynamic database configuration introduced in Junos OS Release 9.5 for EX Series switches.

RELATED DOCUMENTATION

Understanding AS Path Regular Expressions for Use as Routing Policy Match Conditions

dynamic-db | 6105

as-path (Policy Options)

IN THIS SECTION

- Syntax | 2494
- Hierarchy Level | 2494
- Description | 2494
- Options | 2494
- Required Privilege Level | 2494
- Release Information | 2494

Syntax

```
as-path name regular-expression;
```

Hierarchy Level

```
[edit dynamic policy-options],
[edit logical-systems logical-system-name policy-options],
[edit policy-options]
```

Description

Define an autonomous system (AS) path regular expression for use in a routing policy match condition.

Options

name—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 65,536 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

regular-expression—One or more regular expressions used to match the AS path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.

RELATED DOCUMENTATION

Understanding AS Path Regular Expressions for Use as Routing Policy Match Conditions

dynamic-db | 6105

as-path (Routing Options)

IN THIS SECTION

- Syntax | 2495
- Hierarchy Level | 2495
- Description | 2496
- Default | 2496
- Options | 2496
- Required Privilege Level | 2497
- Release Information | 2497

Syntax

```
as-path <as-path> <aggregator as-number ip-address> <atomic-aggregate> <origin (egp | igp |
incomplete)>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options (aggregate | generate | static) (defaults | route)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name (aggregate | generate | static) (defaults | route)],
[edit logical-systems logical-system-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-instances routing-instance-name routing-options (aggregate | generate | static)
```

```
(defaults | route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib routing-table-name (aggregate | generate | static) (defaults | route)]
```

Description

Associate BGP autonomous system (AS) path information with a static, aggregate, or generated route.

In Junos OS Release 9.1 and later, the numeric range for the AS number is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.

In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <*16-bit high-order value in decimal*>.<*16-bit low-order value in decimal*>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value in the range from 0.0 through 65535.65535 in AS-dot notation format.

Default

No AS path information is associated with static routes.

Options

aggregator—(Optional) Attach the BGP **aggregator** path attribute to the aggregate route. You must specify the last AS number that formed the aggregate route (encoded as two octets) for ***as-number***, followed by the IP address of the BGP system that formed the aggregate route for ***ip-address***.

as-path—(Optional) AS path to include with the route. It can include a combination of individual AS path numbers and AS sets. Enclose sets in brackets ([ ]). The first AS number in the path represents the AS immediately adjacent to the local AS. Each subsequent number represents an AS that is progressively farther from the local AS, heading toward the origin of the path. You cannot specify a regular expression for ***as-path***. You must use a complete, valid AS path.

atomic-aggregate—(Optional) Attach the BGP **atomic-aggregate** path attribute to the aggregate route. This path attribute indicates that the local system selected a less specific route instead of a more specific route.

origin egp—(Optional) BGP origin attribute that indicates that the path information originated in another AS.

origin igp—(Optional) BGP origin attribute that indicates that the path information originated within the local AS.

origin incomplete—(Optional) BGP origin attribute that indicates that the path information was learned by some other means.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation

asm-override-ssm

IN THIS SECTION

- Syntax | 2498
- Hierarchy Level | 2498
- Description | 2498
- Required Privilege Level | 2498
- Release Information | 2498

Syntax

```
asm-override-ssm;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Description

Enable the routing device to accept any-source multicast join messages (*,G) for group addresses that are within the default or configured range of source-specific multicast groups.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

> *Example: Configuring Source-Specific Multicast Groups with Any-Source Override*

assemble

IN THIS SECTION

- Syntax | 2499
- Hierarchy Level | 2499
- Description | 2499
- Options | 2499
- Required Privilege Level | 2500
- Release Information | 2500

Syntax

```
assemble {
    common-name common-name;
}
```

Hierarchy Level

```
[edit access ldap-options],
[edit access profile profile-name ldap-options]
```

Description

Specify that a user's LDAP distinguished name (DN) is assembled through the use of a common name identifier, the username, and base distinguished name.

Options

`common-name` *`common-name`*—Common name identifier used as a prefix for the username during the assembly of the user's distinguished name. For example, `uid` specifies " user id," and `cn` specifies "common name."

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

assert-timeout

IN THIS SECTION

- Syntax | 2500
- Hierarchy Level | 2500
- Description | 2501
- Options | 2501
- Required Privilege Level | 2501
- Release Information | 2501

Syntax

```
assert-timeout seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Multicast routing devices running PIM sparse mode often forward the same stream of multicast packets onto the same LAN through the rendezvous-point tree (RPT) and shortest-path tree (SPT). PIM assert messages help routing devices determine which routing device forwards the traffic and prunes the RPT for this group. By default, routing devices enter an assert cycle every 180 seconds. You can configure this assert timeout to be between 5 and 210 seconds.

Options

seconds—Time for routing device to wait before another assert message cycle.

- **Range:** 5 through 210 seconds
- **Default:** 180 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring the PIM Assert Timeout

identity-provider attribute-mapping (SAML)

IN THIS SECTION

- Syntax | 2502
- Hierarchy Level | 2502
- Description | 2502

- Options | 2502
- Required Privilege Level | 2503
- Release Information | 2503

Syntax

```
attribute-mapping {
    address-pool assertion-attribute <mandatory | optional>;
    firstname assertion-attribute <mandatory | optional>;
    group assertion-attribute <mandatory | optional>;
    lastname assertion-attribute <mandatory | optional>;
    attribute-name assertion-attribute <mandatory | optional>;
    username assertion-attribute <mandatory | optional>;
}
```

Hierarchy Level

```
[edit access saml identity-provider identity-provider]
```

Description

Map assertion response attributes to local attributes.

Options

address-pool *assertion-attribute*

Specify the name of the address pool.

- mandatory (default)
- optional

firstname *assertion-attribute*

Specify the user's firstname.

- mandatory (default)

- optional

group ***assertion-attribute***

Specify the user's group attribute details.

- mandatory (default)
- optional

lastname ***assertion-attribute***

Specify the user's lastname.

- mandatory (default)
- optional

attribute-name assertion-attribute

Define and specify custom local attribute name. You can define multiple custom attributes.

- mandatory (default)
- optional

username ***assertion-attribute***

Specify the user's username attribute.

- mandatory (default)
- optional

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

saml | **16281**

assignment-id-format (L2TP LAC)

IN THIS SECTION

- Syntax | 2504
- Hierarchy Level | 2504
- Description | 2504
- Default | 2504
- Options | 2505
- Required Privilege Level | 2505
- Release Information | 2505

Syntax

```
assignment-id-format (assignment-id | client-server-id);
```

Hierarchy Level

```
[edit services l2tp tunnel]
```

Description

Set the format for the name used for a tunnel, the tunnel assignment ID.

NOTE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing `no services l2tp tunnel assignment-id-format`.

Default

`assignment-Id`

Options

assignment-Id The tunnel name corresponds to RADIUS attribute Tunnel-Assignment-Id [82].

client-server-id The tunnel name is a combination of RADIUS attributes Tunnel-Client-Auth-Id [90], Tunnel-Server-Auth-Id [91], and Tunnel-Assignment-Id [82].

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

assisted-replication

IN THIS SECTION

- Syntax | 2506
- Hierarchy Level | 2506
- Description | 2506
- Options | 2507
- Required Privilege Level | 2507
- Release Information | 2507

Syntax

```
assisted-replication {
    leaf {
        replicator-activation-delay seconds;
    }
    replicator {
        inet ip;
        vxlan-encapsulation-source-ip (assisted-replicator-ip | ingress-replication-ip | retain-
incoming-source-ip);
    }
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols evpn],
[edit routing-instances name protocols evpn],
[edit protocols evpn]
```

Description

Enable assisted replication (AR) in an EVPN network.

AR helps to optimize broadcast, unknown unicast, and multicast (BUM) traffic in EVPN networks. To enable AR, you configure devices in the EVPN network to operate as AR replicator and AR leaf devices. Use the statements in this hierarchy to define and set up AR replicators and AR leaf devices.

AR replicator devices help perform BUM traffic replication and forwarding tasks for AR leaf devices. When an AR leaf device receives BUM traffic, instead of always using ingress replication to the EVPN core, it forwards the traffic on an AR overlay tunnel to an AR replicator device that can better handle the replication load. The AR replicator device then replicates and forwards the traffic to other ingress replication overlay tunnels. You configure a loopback interface IP address on each AR replicator device for the AR overlay tunnel.

AR devices can interoperate with devices in the EVPN network that don’t support or are not configured into an AR device role.

Options

leaf Configure a device as an AR leaf device.

replicator Configure a device as an AR replicator device.

These statements and their options are explained in more detail separately.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.4R2 and 19.4R1.

RELATED DOCUMENTATION

Assisted Replication Multicast Optimization in EVPN Networks

assisted-replication (Deterministic AR Replicator Policy Actions)

IN THIS SECTION

- Syntax | 2508
- Hierarchy Level | 2508
- Description | 2508
- Options | 2509
- Required Privilege Level | 2510
- Release Information | 2510

Syntax

```
assisted-replication {
    (fallback-replicator-ip fallback-replicator-ip | strict);
    replicator-ip replicator-ip;;
}
```

Hierarchy Level

```
[edit policy-options policy-statement name then]
```

Description

Configure routing policy actions in a routing policy for an assisted replication (AR) leaf device to deterministically steer multicast flows to specific AR replicator devices in an EVPN network.

Assign the routing policy to an EVPN instance on the AR leaf device using the `deterministic-ar-policy` statement at the `[edit routing-instances name protocols evpn assisted-replication leaf]` hierarchy level.

> **NOTE**: AR is a multicast traffic flow optimization feature where an ingress leaf device sends one copy of a multicast stream to another device called an AR replicator in the EVPN network. AR replicators are devices that can efficiently handle replicating and forwarding multicast streams to the other multicast receivers in the network. See *Assisted Replication Multicast Optimization in EVPN Networks* for details on how AR works and how to configure AR leaf and AR replicator devices.

When you assign routing policies with these actions on an AR leaf device, you control how the device load-balances multicast flows among the available AR replicator devices. The AR leaf device directs flows that match the routing policy's `from` conditions to the AR replicator device or devices you specify with this statement in the policy's `then` clause.

In the policy `from` clause, you can specify to match an IGMPv2 or MLDv1 multicast flow with either or both of the following parameters:

- A multicast group address.
- A bridge domain or VLAN-mapped virtual network identifier (VNI).

You can set match conditions for a source-specific multicast (SSM) IGMPv3 or MLDv2 multicast flow with any combination of the following parameters:

- A multicast group address.
- A bridge domain or VLAN-mapped virtual network identifier (VNI).
- A multicast source device address.

When you configure the `then` clause of the policy with the `assisted-replication` action statement we describe on this page, you must specify the `replicator-ip` option with the address of an AR replicator in the network.

You can configure the policy to direct matching flows strictly to that preferred AR replicator by including the `strict` option, which we call **strict mode**. In strict mode, if the preferred AR replicator goes down, the leaf device drops the matching traffic (even if the network has other available AR replicators).

By default, if you don't configure the `strict` option, the AR leaf device uses **loose mode**. In loose mode, you can include the `fallback-replicator-ip` option to configure the AR leaf device to use a fallback AR replicator if the preferred AR replicator goes down.

Also, consider the following cases in loose mode:

- You specify a fallback replicator, but both the preferred AR replicator and the fallback replicator go down.
- You don't specify a fallback replicator, and the preferred AR replicator goes down.

In those cases, if the AR leaf device detects other available AR replicators in the network, the device uses the default AR replicator load balancing method to distribute multicast traffic among the remaining available AR replicators.

Otherwise, if the AR leaf device doesn't detect any other AR replicators in the network, the device defaults to ingress replication behavior (replicates and forwards the multicast traffic itself to any interested receivers in the EVPN core).

Options

replicator-ip *replicator-ip*

(Required) Set a preferred AR replicator to which the AR leaf device directs the multicast flows that match the conditions in the policy statement's `from` clause.

The *`replicator-ip`* parameter is the secondary loopback address you assign for the AR replicator function when you configure a device as an AR replicator. For example:

```
[edit]
user@ar-replicator1# set interfaces lo0 unit 0 family inet address AR-IP-address

[edit]
```

```
user@ar-replicator1# set protocols evpn assisted-replication replicator inet AR-IP-address
```

See the "replicator" on page 14387 statement and the instructions in *Configure an AR Replicator Device*.

fallback-replicator-ip *fallback-replicator-ip*

(Optional) Set a fallback AR replicator to which the AR leaf device redirects the matching flows if the preferred AR replicator goes down.

The *fallback-replicator-ip* parameter is the AR replicator secondary loopback address of the device you want to use as the fallback device, similar to the address you specify with the `replicator-ip` option explained above.

strict

(Optional) Strictly use only the AR replicator(s) you specify with the `replicator-ip` option. The device doesn't redirect any flows to any other AR replicator if the specified AR replicator goes down; it drops any traffic the policy matches.

> **NOTE**: The `strict` and `fallback-replicator-ip` options are mutually exclusive.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.4R2.

RELATED DOCUMENTATION

Deterministic Load Balancing and Traffic Steering to AR Replicators

Routing Policy Match Conditions

Actions in Routing Policy Terms

deterministic-ar-policy | 5553

associate-backup-pe-groups

IN THIS SECTION

- Syntax | 2511
- Hierarchy Level | 2511
- Description | 2511
- Required Privilege Level | 2511
- Release Information | 2512

Syntax

```
associate-backup-pe-groups;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Enable an LSP to monitor the status of its destination PE router. You can configure multiple backup PE router groups using the same router's address. Backup PE router groups provide ingress PE router redundancy when point-to-multipoint LSPs are configured for multicast distribution. A failure of this LSP indicates to all of the backup PE router groups that the destination PE router is down. This statement is not tied to a specific backup PE router group. It applies to all groups that are interested in the status of the LSP to the destination address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Enabling Point-to-Point LSPs to Monitor Egress PE Routers

associate-lsp

IN THIS SECTION

- Syntax | 2512
- Hierarchy Level | 2512
- Description | 2513
- Options | 2513
- Required Privilege Level | 2513
- Release Information | 2513

Syntax

```
associate-lsp lsp-name {
    from from-ip-address;
}
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name oam]
```

Description

Configure associated bidirectional label-switched paths (LSPs) on the two ends of an LSP for sending and receiving GAL and G-Ach OAM messages.

Options

from *from-ip-address* (Optional) Source address for the associated LSP configuration.

If omitted, this is derived from the `to` address of the ingress LSP configuration.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring the MPLS Transport Profile for OAM

associate-profile

IN THIS SECTION

- Syntax | 2514
- Hierarchy Level | 2514
- Description | 2514
- Options | 2514
- Required Privilege Level | 2514
- Release Information | 2514

Syntax

```
associate-profile {
    dynamic-profile-name;
    profile-variable-set profile-variable-set-name;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols vpls],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name neighbor
neighbor-id]
```

Description

Associate a dynamic profile or a profile variable set with a VPLS instance.

Options

dynamic-profile-name—Name of the dynamic profile to attach to this routing instance.

The remaining option is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

> *Attaching Dynamic Profiles to Routing Instances*
>
> *Attaching Dynamic Profiles to Routing Instances for the Complex Configuration*

association-timeout

IN THIS SECTION

- Syntax | 2515
- Hierarchy Level | 2515
- Description | 2515
- Options | 2515
- Required Privilege Level | 2516
- Release Information | 2516

Syntax

```
association-timeout time-in-minutes;
```

Hierarchy Level

```
[edit security sctp profile profile-name]
```

Description

Use the association-timeout parameter to set the association timeout for Stream Control Transmission Protocol (SCTP).

From Junos OS Release 20.4R1 onwards, the `[edit security gprs sctp profile profile-name]` hierarchy level is replaced by `[edit security sctp profile profile-name]`.

Options

time-in-minutes Number of minutes of association time that elapse before the session is terminated.

- **Range:** 10 through 6000 (100 hours).
- **Default:** 300 minutes (5 hours).

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2. The association timeout range increased in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding Stream Control Transmission Protocol

Understanding SCTP Behavior in Chassis Cluster

asymmetric-hold-time

IN THIS SECTION

- Syntax | 2516
- Hierarchy Level | 2517
- Description | 2517
- Default | 2517
- Required Privilege Level | 2517
- Release Information | 2517

Syntax

```
asymmetric-hold-time;
```

Hierarchy Level

```
[edit protocols vrrp]
```

Description

Enable the VRRP primary router to switch over to the backup router immediately, without waiting for the priority hold time to expire, when a tracked route or interface goes down. When the route or interface comes back online, the original primary router that is now acting as the backup router waits for the priority hold time to expire before it reasserts primary role.

Default

`asymmetric-hold-time` is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring the Asymmetric Hold Time for VRRP Routers

asymmetric-hold-time

IN THIS SECTION

- Syntax | 2518
- Hierarchy Level | 2518

- Description | 2518
- Default | 2518
- Required Privilege Level | 2518
- Release Information | 2518

Syntax

```
asymmetric-hold-time;
```

Hierarchy Level

```
[edit protocols vrrp]
```

Description

Configure a VRRP primary to fail over to a backup immediately—without waiting for the priority hold time to expire—when a tracked route goes down. Otherwise, the primary waits for the hold time to expire before it initiates a failover when a tracked interface or route goes down.

When the tracked interface or route comes up again, the new backup (original primary) router waits for the priority hold time to expire before it reasserts primary role.

Default

asymmetric-hold-time is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 11.3.

RELATED DOCUMENTATION

Configuring VRRP Preemption and Hold Time

asymmetric-lease-time (DHCP Overrides)

IN THIS SECTION

- Syntax | 2519
- Hierarchy Level | 2519
- Description | 2520
- Options | 2520
- Required Privilege Level | 2520
- Release Information | 2520

Syntax

```
asymmetric-lease-time seconds;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
```

```
options dhcp-relay dhcpv6...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6...]
```

Description

Set an asymmetric lease time for an address that is sent to the DHCP client instead of the actual lease time granted by the DHCP local server. The asymmetric lease time is shorter than the granted lease. Asymmetric leases require the DHCP client to renew the lease more frequently, at shorter intervals. If the client does not renew the short lease before it expires, the address is released sooner than if the original lease had been sent to the client, conserving address resources. Although you can configure the asymmetric lease on either the DHCP relay agent or the DHCP local server, in most situations, configuring the asymmetric lease on the relay agent is more useful.

Options

seconds Length of the short lease duration for the DHCPv4 client.

- **Range:** 600 through 65,534

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> *DHCP Asymmetric Leasing Overview*
>
> *Configuring DHCP Asymmetric Leasing*

asymmetric-prefix-lease-time (DHCP Overrides)

IN THIS SECTION

- Syntax | 2521
- Hierarchy Level | 2521
- Description | 2521
- Options | 2522
- Required Privilege Level | 2522
- Release Information | 2522

Syntax

```
asymmetric-prefix-lease-time seconds;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6...]
```

Description

Set an asymmetric lease time for an address that is sent to the DHCPv6 client instead of the actual lease time granted by the DHCP local server. The asymmetric lease time is shorter than the granted lease. Asymmetric leases require the DHCP client to renew the lease more frequently, at shorter intervals. If the client does not renew the short lease before it expires, the address is released sooner than if the original lease had been sent to the client, conserving address resources. Although you can configure the asymmetric lease on either the DHCP relay agent or the DHCP local server, in most situations, configuring the asymmetric lease on the relay agent is more useful.

Options

seconds Length of the short lease duration for the DHCPv6 client.

- **Range:** 600 through 65,534

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

DHCP Asymmetric Leasing Overview

Configuring DHCP Asymmetric Leasing

asynchronous-notification

IN THIS SECTION

- Syntax | 2523
- Hierarchy Level | 2523
- Description | 2523
- Default | 2524
- Required Privilege Level | 2524
- Release Information | 2524

Syntax

```
(asynchronous-notification | no-asynchronous-notification);
```

Hierarchy Level

```
[edit interfaces ge-fpc/pic/port gigether-options]
```

Description

The **asynchronous-notification** feature allows supported interfaces to generate and transfer link down alarms. This feature triggers the laser off (LOS) signal on an interface based on specific events and conditions and sends the LOS alarm to its peer (CE).

When you configure the **asynchronous-notification** knob on an interface, the LOS signal will be triggered under the following conditions:

- **RPD CCC Down Event**: If the RPD (Routing and Policy Database) detects a CCC (Cross-Connect Circuit) down event, indicating a loss of connectivity for the associated service.
- **Interface Encapsulation:** If the interface encapsulation is set to ethernet-ccc, ethernet-vpls, or vlan-ccc.

The LOS alarm will be sent to the peer CE (Customer Edge) device in these cases. This functionality is particularly useful for port-based services.

> **NOTE**: You can only configure the "asynchronous-notification" knob at the Interface Device (IFD) level. The RPD determines the CCC down event on a per Interface Logical (IFL) basis, while the Packet Forwarding Engine (PFE) sets the LOS alarm on the port. Also, This feature is not recommended for IFL-based services.
>
> Also, the asynchronous-notification cannot be configured on aggregated Ethernet (ae) interfaces or their member links. Therefore, the LOS functionality triggered by the knob will not be available on ae interfaces.

By enabling the **asynchronous-notification** feature, you will receive timely notifications about link down events, allowing you to take appropriate actions or perform troubleshooting to ensure seamless network connectivity and service availability.

Default

Support for notification of link down alarm generation and transfer is not enabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

> *Gigabit Ethernet Notification of Link Down Alarm Overview*
>
> *Configuring Gigabit Ethernet Notification of Link Down Alarm*

atm-options

IN THIS SECTION

- Syntax | 2524
- Hierarchy Level | 2526
- Description | 2526
- Required Privilege Level | 2526
- Release Information | 2526

Syntax

```
atm-options {
    cell-bundle-size cells;
    ilmi;
```

```
    linear-red-profiles profile-name {
        high-plp-max-threshold percent;
        low-plp-max-threshold percent;
        queue-depth cells high-plp-threshold percent low-plp-threshold percent;
    }
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
    pic-type (atm1 | atm2);
    plp-to-clp;
    promiscuous-mode {
        vpi vpi-identifier;
    }
    scheduler-maps map-name {
        forwarding-class class-name {
            epd-threshold cells plp1 cells;
            linear-red-profile profile-name;
            priority (high | low);
            transmit-weight (cells number | percent number);
        }
        vc-cos-mode (alternate | strict);
    }
    use-null-cw;
    vpi vpi-identifier {
        maximum-vcs maximum-vcs;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (disable | seconds);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
rate burst length);
            queue-length number;
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure ATM-specific physical interface properties.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

NOTE: Certain options apply only to specific platforms.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Interface Encapsulations Overview
- multipoint-destination
- shaping
- vci | 30660

attach

IN THIS SECTION

- Syntax | 2527
- Hierarchy Level | 2527
- Description | 2527
- Options | 2527
- Required Privilege Level | 2527
- Release Information | 2528

Syntax

```
attach {
    zone zone-name;
}
```

Hierarchy Level

```
[edit security address-book book-name]
```

Description

Attach a security zone to an address book. You do not need to attach a security zone to the global address book. The global address book is available by default.

Options

`zone` *`zone-name`*–Name of a security zone to be attached to the address book.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Address Books

Understanding Address Sets

attack-threshold

IN THIS SECTION

- Syntax | 2528
- Hierarchy Level | 2528
- Description | 2529
- Options | 2529
- Required Privilege Level | 2529
- Release Information | 2529

Syntax

```
attack-threshold number;
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp syn-flood]
```

Description

Define the number of SYN packets per second required to trigger the SYN proxy response.

Options

number —Number of SYN packets per second required to trigger the SYN proxy response.

- **Range:** 1 through 500,000 per second
- **Default:** 200 per second

> **NOTE**: For SRX Series Firewalls, the applicable range is 1 through 1.000,000 per second.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

destination-threshold | 5525

attack-type (Security Anomaly)

IN THIS SECTION

- Syntax | 2530
- Hierarchy Level | 2530

- Description | 2530
- Options | 2530
- Required Privilege Level | 2530
- Release Information | 2531

Syntax

```
attack-type {
    anomaly {
        direction (Security Custom Attack) (any | client-to-server | server-to-client);
        service service-name;
        shellcode (all | intel | no-shellcode | sparc);
        test test-condition;
    }
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name]
```

Description

Specify the type of attack.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

attack-type (Security Chain)

IN THIS SECTION

- Syntax | 2531
- Hierarchy Level | 2532
- Description | 2532
- Options | 2532
- Required Privilege Level | 2532
- Release Information | 2533

Syntax

```
attack-type {
    chain {
        expression boolean-expression;
        member (Security IDP) member-name {
            attack-type {
                (anomaly ...same statements as in [edit security idp custom-attack attack-name
attack-type anomaly] hierarchy level | signature ...same statements as in [edit security idp
custom-attack attack-name attack-type signature] hierarchy level);
            }
        }
        order;
        protocol-binding {
            application application-name;
            icmp;
            icmpv6;
            ip (Security IDP Custom Attack) {
                protocol-number transport-layer-protocol-number;
            }
```

```
            ipv6 {
                protocol-number transport-layer-protocol-number;
            }
            rpc {
                program-number rpc-program-number;
            }
            tcp (Security IDP Protocol Binding) {
                minimum-port port-number <maximum-port port-number>;
            }
            udp (Security IDP Protocol Binding) {
                minimum-port port-number <maximum-port port-number>;
            }
        }
        reset;
        scope (session | transaction);
    }
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name]
```

Description

Specify the type of attack.

> **NOTE**: In a chain attack, you can configure multiple member attacks.
> In an attack, under protocol binding TCP/UDP, you can specify multiple ranges of ports.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

attack-type (Security IDP)

IN THIS SECTION

- Syntax | 2533
- Hierarchy Level | 2539
- Description | 2539
- Options | 2539
- Required Privilege Level | 2539
- Release Information | 2540

Syntax

```
attack-type {
anomaly {
    direction (Security Custom Attack) (any | client-to-server | server-to-client);
    shellcode (all | intel |no-shellcode | sparc);
    test-condition condition-name;
}
signature (Security IDP) {
    context context-name;
    direction (Security Custom Attack) (any | client-to-server | server-to-client);
    negate;
    pattern signature-pattern;
    pattern-pcre signature-pattern-pcre;
    protocol (Security IDP Signature Attack) {
        icmp (Security IDP Custom Attack) {
            checksum-validate {
                match (Security IDP Policy) (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
```

```
            code {
                match (equal | greater-than | less-than | not-equal);
                value code-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
            }
            identification (Security ICMP Headers) {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            sequence-number (Security IDP ICMP Headers) {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
            type (Security IDP ICMP Headers) {
                match (Security IDP Policy) (equal | greater-than | less-than | not-equal);
                value type-value;
            }
        }
        icmpv6 (Security IDP Custom Attack) {
            checksum-validate {
                match (Security IDP Policy) (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
            code {
                match (equal | greater-than | less-than | not-equal);
                value code-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
            }
            identification (Security ICMP Headers) {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            sequence-number (Security IDP ICMP Headers) {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
            type (Security IDP ICMP Headers) {
```

```
                match (equal | greater-than | less-than | not-equal);
                value type-value;
            }
        }
        ipv4 (Security IDP Signature Attack) {
            checksum-validate {
                match (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
            destination (Security IP Headers Attack) {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            identification (Security ICMP Headers) {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            ip-flags {
                (df | no-df);
                (mf | no-mf);
                (rb | no-rb);
            }
            protocol (Security IDP Signature Attack) {
                match (equal | greater-than | less-than | not-equal);
                value transport-layer-protocol-id;
            }
            source {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            tos {
                match (equal | greater-than | less-than | not-equal);
                value type-of-service-in-decimal;
            }
            total-length {
                match (equal | greater-than | less-than | not-equal);
                value total-length-of-ip-datagram;
            }
            ttl {
                match (equal | greater-than | less-than | not-equal);
                value time-to-live;
            }
        }
```

```
        ipv6 {
            destination {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            extension-header {
                destination-option {
                    home-address {
                        match (equal | greater-than | less-than | not-equal);
                        value header-value;
                    }
                        option-type {
                            match (equal | greater-than | less-than | not-equal);
                            value header-value;
                        }
                }
                routing-header {
                    header-type {
                        match (equal | greater-than | less-than | not-equal);
                        value header-value;
                        }
                    }
            }
            flow-label {
                match (equal | greater-than | less-than | not-equal);
                value flow-label-value;
            }
            hop-limit {
                match (equal | greater-than | less-than | not-equal);
                value hop-limit-value;
            }
            next-header {
                match (equal | greater-than | less-than | not-equal);
                value next-header-value;
            }
            payload-length {
                match (equal | greater-than | less-than | not-equal);
                value payload-length-value;
            }
            source {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
```

```
            traffic-class {
                match (equal | greater-than | less-than | not-equal);
                value traffic-class-value;
            }
        tcp (Security IDP Signature Attack) {
            ack-number {
                match (equal | greater-than | less-than | not-equal);
                value acknowledgement-number;
            }
            checksum-validate {
                match (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value tcp-data-length;
            }
            destination-port {
                match (equal | greater-than | less-than | not-equal);
                value destination-port;
            }
            header-length {
                match (equal | greater-than | less-than | not-equal);
                value header-length;
            }
            mss {
                match (equal | greater-than | less-than | not-equal);
                value maximum-segment-size;
            }
            option {
                match (equal | greater-than | less-than | not-equal);
                value tcp-option;
            }
            reserved {
                match (equal | greater-than | less-than | not-equal);
                value reserved-value;
            }
            sequence-number {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
            source-port {
                match (equal | greater-than | less-than | not-equal);
```

```
                value source-port;
            }
            tcp-flags {
                (ack | no-ack);
                (fin | no-fin);
                (psh | no-psh);
                (r1 | no-r1);
                (r2 | no-r2);
                (rst | no-rst);
                (syn | no-syn);
                (urg | no-urg);
            }
            urgent-pointer {
                match (equal | greater-than | less-than | not-equal);
                value urgent-pointer;
            }
            window-scale {
                match (equal | greater-than | less-than | not-equal);
                value window-scale-factor;
            }
            window-size {
                match (equal | greater-than | less-than | not-equal);
                value window-size;
            }
        }
        udp (Security IDP Signature Attack){
            checksum-validate {
                match (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
            }
            destination-port (Security Signature Attack) {
                match (equal | greater-than | less-than | not-equal);
                value destination-port;
            }
            source-port {
                match (equal | greater-than | less-than | not-equal);
                value source-port;
            }
        }
```

```
    }
    protocol-binding {
        application application-name;
        icmp;
        icmpv6;
        ip {
            protocol-number transport-layer-protocol-number;
        }
        ipv6 {
            protocol-number transport-layer-protocol-number;
        }
        rpc {
            program-number rpc-program-number;
        }
        tcp {
            minimum-port port-number <maximum-port port-number>;
        }
        udp {
            minimum-port port-number <maximum-port port-number>;
        }
    }
    regexp regular-expression;
    shellcode (all | intel | no-shellcode | sparc);
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain member member-name]
```

Description

Specify the type of attack.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

attack-type (Security Signature)

IN THIS SECTION

- Syntax | 2540
- Hierarchy Level | 2545
- Description | 2545
- Options | 2545
- Required Privilege Level | 2546
- Release Information | 2546

Syntax

```
attack-type {
    signature (Security IDP) {
        context context-name;
        direction (any | client-to-server | server-to-client);
        negate;
        pattern signature-pattern;
        pattern-pcre signature-pattern-pcre;
        protocol (Security IDP Signature Attack) {
            icmp {
                code {
                    match (equal | greater-than | less-than | not-equal);
                    value code-value;
                }
                data-length {
                    match (equal | greater-than | less-than | not-equal);
                    value data-length;
```

```
            }
            identification {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            sequence-number {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
            type {
                match (equal | greater-than | less-than | not-equal);
                value type-value;
            }
        }
        icmpv6 (Security IDP Custom Attack) {
            code {
                match (equal | greater-than | less-than | not-equal);
                value code-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
            }
            identification {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            sequence-number {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
            type {
                match (equal | greater-than | less-than | not-equal);
                value type-value;
            }
        }
        ipv4 (Security IDP Signature Attack) {
            destination {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            identification {
                match (equal | greater-than | less-than | not-equal);
```

```
                value identification-value;
            }
            ihl {
                match (equal | greater-than | less-than | not-equal);
                value ihl-value;
            }
            ip-flags {
                (df | no-df);
                (mf | no-mf);
                (rb | no-rb);
            }
            protocol {
                match (equal | greater-than | less-than | not-equal);
                value transport-layer-protocol-id;
            }
            source {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            tos {
                match (equal | greater-than | less-than | not-equal);
                value type-of-service-in-decimal;
            }
            total-length {
                match (equal | greater-than | less-than | not-equal);
                value total-length-of-ip-datagram;
            }
            ttl {
                match (equal | greater-than | less-than | not-equal);
                value time-to-live;
            }
        }
        ipv6 {
            destination {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            flow-label {
                match (equal | greater-than | less-than | not-equal);
                value flow-label-value;
            }
            hop-limit {
                match (equal | greater-than | less-than | not-equal);
```

```
                value hop-limit-value;
            }
            next-header {
                match (equal | greater-than | less-than | not-equal);
                value next-header-value;
            }
            payload-length {
                match (equal | greater-than | less-than | not-equal);
                value payload-length-value;
            }
            source {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            traffic-class {
                match (equal | greater-than | less-than | not-equal);
                value traffic-class-value;
            }
        tcp {
            ack-number {
                match (equal | greater-than | less-than | not-equal);
                value acknowledgement-number;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value tcp-data-length;
            }
            destination-port {
                match (equal | greater-than | less-than | not-equal);
                value destination-port;
            }
            header-length {
                match (equal | greater-than | less-than | not-equal);
                value header-length;
            }
            mss {
                match (equal | greater-than | less-than | not-equal);
                value maximum-segment-size;
            }
            option {
                match (equal | greater-than | less-than | not-equal);
                value tcp-option;
            }
```

```
            sequence-number {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
            source-port {
                match (equal | greater-than | less-than | not-equal);
                value source-port;
            }
            tcp-flags {
                (ack | no-ack);
                (fin | no-fin);
                (psh | no-psh);
                (r1 | no-r1);
                (r2 | no-r2);
                (rst | no-rst);
                (syn | no-syn);
                (urg | no-urg);
            }
            urgent-pointer {
                match (equal | greater-than | less-than | not-equal);
                value urgent-pointer;
            }
            window-scale {
                match (equal | greater-than | less-than | not-equal);
                value window-scale-factor;
            }
            window-size {
                match (equal | greater-than | less-than | not-equal);
                value window-size;
            }
        }
        udp (Security IDP Signature Attack) {
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
            }
            destination-port {
                match (equal | greater-than | less-than | not-equal);
                value destination-port;
            }
            source-port {
                match (equal | greater-than | less-than | not-equal);
                value source-port;
```

```
                }
            }
        }
        protocol-binding {
            application application-name;
            icmp;
            icmpv6;
            ip {
                protocol-number transport-layer-protocol-number;
            }
            ipv6 {
                protocol-number transport-layer-protocol-number;
            }
            rpc {
                program-number rpc-program-number;
            }
            tcp {
                minimum-port port-number <maximum-port port-number>;
            }
            udp {
                minimum-port port-number <maximum-port port-number>;
            }
        }
        regexp regular-expression;
        shellcode (all | intel | no-shellcode | sparc);
    }
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name]
```

Description

Specify the type of attack.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

attacks (Security Exempt Rulebase)

IN THIS SECTION

- Syntax | 2546
- Hierarchy Level | 2547
- Description | 2547
- Options | 2547
- Required Privilege Level | 2547
- Release Information | 2547

Syntax

```
attacks {
    custom-attack-group [attack-group-name];
    custom-attack [attack-name];
    dynamic-attack-group [attack-group-name];
    predefined-attack-groups [attack-group-name];
    predefined-attacks [attack-name];
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match]
```

Description

Specify the attacks that you do not want the device to match in the monitored network traffic. Each attack is defined as an attack object, which represents a known pattern of attack.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

attacks (Security IPS Rulebase)

IN THIS SECTION

- Syntax | 2548
- Hierarchy Level | 2548
- Description | 2548
- Options | 2548
- Required Privilege Level | 2548
- Release Information | 2548

Syntax

```
attacks {
    custom-attack-group [attack-group-name];
    custom-attacks [attack-name];
    dynamic-attack-group [attack-group-name];
    predefined-attack-groups [attack-group-name];
    predefined-attacks [attack-name];
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match]
```

Description

Specify the attacks you want the device to match in the monitored network traffic. Each attack is defined as an attack object, which represents a known pattern of attack.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

attempts (DHCP Local Server)

IN THIS SECTION

- Syntax | 2549
- Hierarchy Level | 2549
- Description | 2550
- Options | 2550
- Required Privilege Level | 2550
- Release Information | 2550

Syntax

```
attempts attempt-count;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
```

```
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure],
[edit system services dhcp-local-server reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server group group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Configure how many attempts are made to reconfigure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces before reconfiguration is considered to have failed. A group configuration takes precedence over a DHCP local server configuration.

Options

attempt-count—Maximum number of attempts.

- **Range:** 1 through 10
- **Default:** 8

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

> *Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview*
>
> *Configuring Dynamic Reconfiguration Attempts for DHCP Clients*

attribute

IN THIS SECTION

- Syntax | 2551
- Hierarchy Level | 2552
- Description | 2552
- Options | 2553
- Required Privilege Level | 2553
- Release Information | 2553

Syntax

```
attribute name {
    code numeric-code;
    vendor-id vendor-id;
    format {
        integer {
            equals {
                value;
            }
            greater-than value;
            less-than value;
        }
        string {
            equals {
                value;
            }
            has-prefix{
                value;
            }
            has-suffix {
                value;
            }
            matches {
                value;
            }
```

```
        }
        time {
            equals {
                value;
            }
            greater-than value;
            less-than value;
        }
        v4address {
            equals {
                value;
            }
        }
        v6address {
            equals {
                value;
            }
        }
        v6prefix {
            equals {
                value;
            }
        }
    }
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from]
```

Description

Specify a custom RADIUS attribute for the incoming RADIUS request from the subscriber. You can configure up to five attributes.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

Options

name Name for the attribute.

- **Range:** 1 through 50 alphanumeric characters. Allowed characters are [a-z, A-Z, 0-9]

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

attributes (Access-Line Rate Adjustment)

IN THIS SECTION

- Syntax | 2554
- Hierarchy Level | 2554
- Description | 2554
- Options | 2554
- Required Privilege Level | 2555

Release Information | 2555

Syntax

```
attributes {
    preference (dsl | pon);
}
```

Hierarchy Level

```
[edit system access-line]
```

Description

Selects the set of access line attributes (TLVs) that is saved and processed when the router receives both DSL TLVs and PON TLVs in ANCP port status messages or in PPPoE-IA tags. The default line type is PON. You select the `dsl` options to prefer the DSL TLVs. For example, you might do this when the PON TLVs are unreliable, perhaps because of some issue with the OLT. The `preference` option has no effect when the BNG receives attributes of only one access line type from the OLT.

NOTE: The OLT might redundantly report the PON access line attributes both in PON TLVs and by overloading DSL TLVs. In this case, the DSL-Type TLV (0x91) is set to OTHER and PON rates for the access line are presented in the DSL TLVs, Actual-Net-Data-Rate-Upstream (0x81) and Actual-Net-Data-Rate-Downstream (0x82).

Options

preference dsl Specifies that the router saves and processes DSL TLVs from the DSL-Line-Attributes TLV 0x04. The router discards the PON-Access-Line-Attributes TLV (0x12) and the PON TLVs that it contains.

preference pon Specifies that the router saves and processes PON TLVs from the PON-Line-Attributes TLV 0x12. The router discards the DSL-Access-Line-Attributes TLV (0x04) and the DSL TLVs that it contains.

- **Default:** pon

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Neighbors and Operations

attributes (Diameter Gx Profiles)

IN THIS SECTION

- Syntax | 2556
- Hierarchy Level | 2556
- Description | 2556
- Required Privilege Level | 2556
- Release Information | 2556

Syntax

```
attributes {
    exclude {
        an-gw-address;
        default-eps-bearer-qos;
        packet-filter-information;
        packet-filter-operation;
        rat-type;
    }
    include {
        gx-capability-list;
        rule-suggestion;
    }
}
```

Hierarchy Level

```
[edit unified-edge diameter-profiles gx-profile profile-name]
```

Description

Configure attribute-value pairs (AVPs) that are excluded from or included in the credit control request (CCR) messages between the MX Series router and the policy and charging enforcement function (PCEF).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

gx-profile | 7954

attributes (JSRC Attributes)

IN THIS SECTION

- Syntax | 2557
- Hierarchy Level | 2557
- Description | 2557
- Required Privilege Level | 2558
- Release Information | 2558

Syntax

```
attributes {
    exclude {
        user-name [ authorization-request | provisioning-request ];
        }
    }
}
```

Hierarchy Level

```
[edit access profile profile-name jsrc]
```

Description

Specify how the router processes Diameter attributes.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 14.2.

RELATED DOCUMENTATION

> *Excluding AVPs from Diameter Messages for JSRC*
>
> *Understanding JSRC-SAE Interactions*

attributes-match (Event Policy)

IN THIS SECTION

- Syntax | 2558
- Hierarchy Level | 2559
- Description | 2559
- Required Privilege Level | 2559
- Release Information | 2559

Syntax

```
attributes-match {
    event1.attribute-name equals event2.attribute-name;
    event.attribute-name matches regular-expression;
    event1.attribute-name starts-with event2.attribute-name;
}
```

Hierarchy Level

```
[edit event-options policy policy-name]
```

Description

Execute an event policy only if the attributes of two events are correlated or if the attribute of one event matches a regular expression.

If the `attributes-match` statement includes the `equals` or `starts-with` options, or if it includes a `matches` option that includes a clause for an event that is not specified at the `[edit event-options policy policy-name events]` hierarchy level, you must include one or more `within` statements in the same policy configuration.

You can include event policy variables within the statement to differentiate between a trigger event attribute and a correlated event attribute. You can use variables of the following forms:

- {$$.*attribute-name*}—The double dollar sign ($$) notation represents the event that is triggering a policy. When combined with an attribute name, the variable resolves to the value of the attribute associated with the triggering event.
- {$*event*.*attribute-name*}—The single dollar sign with the event name ($*event*) notation represents the most recent event that matches *event*. When combined with an attribute name, the variable resolves to the value of the attribute associated with that event.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Use Correlated Events to Trigger an Event Policy

Use Regular Expressions to Refine the Set of Events That Trigger a Policy

equals (Event Policy) | **6434**

matches (Event Policy) | **10412**

starts-with (Event Policy) | 28510
within (Event Policy) | 31136

attributes (RADIUS Attributes)

IN THIS SECTION

- Syntax | 2560
- Hierarchy Level | 2561
- Description | 2561
- Options | 2561
- Required Privilege Level | 2562
- Release Information | 2563

Syntax

```
attributes {
    exclude {
        attribute-name packet-type;
        standard-attribute number {
            packet-type [ access-request | accounting-off | accounting-on | accounting-start |
accounting-stop ];
        }
        vendor-id id-number {
            vendor-attribute vsa-number {
                packet-type [ access-request | accounting-off | accounting-on | accounting-start
| accounting-stop ];
            }
        }
    }
    ignore {
        dynamic-iflset-name;
        framed-ip-netmask;
        idle-timeout;
        input-filter;
```

```
        logical-system-routing-instance;
        output-filter;
        session-timeout;
        standard-attribute number;
        vendor-id id-number {
            vendor-attribute vsa-number;
        }
    }
}
```

Hierarchy Level

```
[edit access profile profile-name radius]
```

Description

Specify how the router or switch processes RADIUS attributes.

Options

"exclude" on page 6617

Configure the router or switch to exclude the specified attributes from being sent in the specified type of RADIUS message. Exclusion can be useful, for example, for attributes that do not change values over the lifetime of a subscriber. By not sending these attributes, you reduce the packet size without losing information. Contrast this behavior with that provided by the `ignore` statement.

The options for this statement are explained separately. Click the linked statement for details.

ignore

Configure the router or switch to ignore the specified attributes in RADIUS Access-Accept messages. Standard attributes and VSAs received in RADIUS messages take precedence over internally provisioned attribute values. Ignoring the attributes enables your internally provisioned values to be used instead. Contrast this behavior with that provided by the `exclude` statement.

Starting in Junos OS Release 18.1R1, you can specify RADIUS standard attributes with the attribute number. You can specify vendor-specific attributes (VSAs) with the IANA-assigned vendor ID and the VSA number. With this flexible configuration method, you can configure any standard attribute and VSA supported by your platform to be ignored. The

configuration has no effect if you can configure unsupported attributes, vendors, and VSAs.

The legacy method allows you to configure only those attributes and VSAs for which the statement syntax includes a specific option. Consequently, you can use the legacy method to ignore only a subset of all attributes that can be received in Access-Accept messages.

- Values:
  - dynamic-iflset-name—Ignore Juniper Networks VSA 26-130, Qos-Set-Name.
  - framed-ip-netmask—Ignore RADIUS attribute 9, Framed-IP-Netmask.
  - idle-timeout—Ignore RADIUS attribute 28, Idle-Timeout.
  - input-filter—Ignore Juniper Networks VSA 26-10, Ingress-Policy-Name.
  - logical-system-routing-instance—Ignore Juniper Networks VSA 26-1.
  - output-filter—Ignore Juniper Networks VSA 26-11, Egress-Policy-Name.
  - session-timeout—Ignore RADIUS attribute 27, Session-Timeout.
  - standard-attribute *number*—RADIUS standard attribute number supported by your platform. You can enclose multiple values in square brackets to specify a list of attributes. If you configure an unsupported attribute, that configuration has no effect. Range: 1 through 255.
  - vendor-attribute *vsa-number*—Number identifying a VSA belonging to the specified vendor; both must be supported by your platform. You can enclose multiple values in square brackets to specify a list of VSAs. If you configure an unsupported VSA, that configuration has no effect. Range: 1 through 255.
  - vendor-id *id-number*—IANA vendor ID supported by your platform. If you configure an unsupported vendor ID, that configuration has no effect. Range: 1 through 16777215.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

Standard and Vendor-Specific RADIUS Attributes

audible (Security Alarms)

IN THIS SECTION

- Syntax | 2563
- Hierarchy Level | 2563
- Description | 2564
- Options | 2564
- Required Privilege Level | 2564
- Release Information | 2564

Syntax

```
audible {
    continuous;
}
```

Hierarchy Level

```
[edit security alarms]
```

Description

Configure alarm to beep when a new security alarm is enabled.

Options

`continuous`—Specify alarm to keep beeping until all security alarms are cleared.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 11.2.

RELATED DOCUMENTATION

Security Policies Overview

auth-entry

IN THIS SECTION

- Syntax | 2565
- Hierarchy Level | 2565
- Description | 2565
- Options | 2565
- Required Privilege Level | 2566
- Release Information | 2566

Syntax

```
auth-entry {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of firewall authentication entries that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

- `maximum` *`amount`*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- `reserved` *`amount`*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only)

auth-only-browser

IN THIS SECTION

- Syntax | 2566
- Hierarchy Level | 2567
- Description | 2567
- Options | 2567
- Required Privilege Level | 2567
- Release Information | 2567

Syntax

```
auth-only-browser <auth-user-agent [user-agent] >;
auth-only-browser;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication pass-through]
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication user-firewall]
```

Description

Configure firewall authentication to ignore non-browser HTTP/HTTPS traffic. This feature allows you to ensure that unauthenticated users issuing access requests through HTTP/HTTPS browsers are presented with a captive portal interface to allow them to authenticate. By default, firewall authentication responds to all HTTP/HTTPS traffic.

It can happen that non-browser HTTP/HTTPS services running in the background can trigger captive portal authentication, creating a race condition that suppresses presentation of the captive portal interface to the HTTP/HTTPS browser user.

When **auth-only-browser** is configured, non-browser HTTP traffic is dropped to allow for captive portal to be presented to unauthenticated users who request access using a browser.

Options

auth-user-agent ***user-agent***

Allow the SRX Series device to use the user-agent strings that you specify to verify that the browser traffic is HTTP/HTTPS traffic. Firewall authentication checks the strings against the User-Agent field in the browser header. You can specify only one value for this parameter. It must not contain spaces and it does not need to be enclosed in parenthesis. For example, **auth-user-agent** might specify Opera1 as one of its values.

You can use the **auth-user-agent** parameter alone for pass-through or user-firewall authentication or in conjunction with **auth-only-browser**.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D90.

RELATED DOCUMENTATION

auth-user-agent | 2568

Understanding SRX Series Assured Captive Portal Support for Unauthenticated Browser Users

auth-user-agent

IN THIS SECTION

- Syntax | 2568
- Hierarchy Level | 2568
- Description | 2569
- Options | 2569
- Required Privilege Level | 2569

Syntax

```
auth-user-agent [user-agent];
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication pass-through]
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication pass-through auth-only-browser]
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication user-firewall]]
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication user-firewall auth-only-browser]
```

Description

Specify a user-agent value to be used to verify that the user's browser traffic is HTTP/HTTPS traffic. Firewall authentication checks the value against the User-Agent field in the browser header. For example, the **auth-user-agent** parameter might specify Opera1 to be verified against the brower's User-Agent field for a match.

You can use the **auth-user-agent** parameter alone for pass-through or user-firewall authentication or in conjunction with **auth-only-browser**.

The **auth-only-browser** directs firewall authentication to ignore non-browser HTTP/HTTPS traffic to ensure that unauthenticated users using an HTTP/HTTPS browser are authenticated by captive portal before they are granted access to protected resources. It can happen that non-browser HTTP/HTTPS services running in the background can trigger captive portal authentication creating a race condition that suppresses presentation of the captive portal interface to the HTTP/HTTPS browser user.

Options

user-agent A string to be matched against values specified in the browser's User-Agent header field that identifies the traffic as HTTP/HTTPS browser traffic. You can specify only one user-agent value for a security policy configuration. The value must not contain spaces. You do not need to enclose the string in parenthesis. The length of a string must be 17 characters or less.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

RELATED DOCUMENTATION

auth-only-browser | **2566**

Understanding SRX Series Assured Captive Portal Support for Unauthenticated Browser Users

authentication (Protocols OSPF)

IN THIS SECTION

- Syntax | 2570
- Hierarchy Level | 2570
- Description | 2571
- Options | 2571
- Required Privilege Level | 2572
- Release Information | 2572

Syntax

```
authentication {
    md5 key-identifier {
        key key-value;
        start-time YYYY-MM-DD.hh:mm;
    }
    simple-password key;
    multi-active-md5 key-identifier {
        key key-value delete-if-not-inuse;
    }
    keychain keychain;}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id interface interface-name],
[edit logical-systems logical-system-name protocols ospf area area-id virtual-link],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area area-id virtual-link],
[edit protocols ospf area area-id interface interface-name],
[edit protocols ospf area area-id virtual-link],
[edit routing-instances routing-instance-name protocols ospf area area-id interface interface-
```

```
name],
[edit routing-instances routing-instance-name protocols ospf area area-id virtual-link]
```

Description

Configure an authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface.

All routers that are connected to the same IP subnet must use the same authentication scheme and password.

The remaining statements are explained separately. See CLI Explorer.

Options

simple-password *key*

Configures a simple authentication password string, defined by *key*.

md5

Configure an MD5 password.

- *key-identifier*—MD5 key identifier. Range is from 0 through 255. Default is 0.
- `key` *key-values*—One or more MD5 key strings. The MD5 key values can be from 1 through 16 characters long. You can specify more than one key value within the list. Characters can include ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).
- `start-time` *time*—MD5 start date and time, in the format *YYYY-MM-DD.hh:mm*.

multi-active-md5

Configure authentication type MD5 with multi active keys

- *key-identifier*—multi active MD5 key identifier. Range is from 0 through 255. Default is 0.
- `key` *key-values*One or more MD5 key strings. The MD5 key values can be from 1 through 16 characters long. You can specify more than one key value within the list. Characters can include ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).
- `delete-if-not-inuse`—Delete this key if not in use. This option added under `multi-active-md5` for an OSPF interface marks the key as old. If a new key is available on an OSPFv2 interface and once all the neighbors on the OSPFv2 interface use the new key, then OSPFv2 stops sending OSPFv2 packets with key that is marked as `delete-if-not-inuse`.

keychain *keychain*

For each keychain, you must specify a name. For example, ospf_auth. You can have multiple keychains on a routing device. For example, you can have a keychain for OSPF, a keychain for BGP, and a keychain for LDP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The option `keychain` *`keychain`* `introduced in Junos OS and Junos OS Evolved Release 23.4R1.`

RELATED DOCUMENTATION

Understanding OSPFv2 Authentication

Example: Configuring MD5 Authentication for OSPFv2 Exchanges

Example: Configuring a Transition of MD5 Keys on an OSPFv2 Interface

Example: Configuring Simple Authentication for OSPFv2 Exchanges

authentication (services)

IN THIS SECTION

- Syntax | 2573
- Hierarchy Level | 2573
- Description | 2573
- Options | 2573
- Required Privilege Level | 2573
- Release Information | 2574

Syntax

```
authentication {
    algorithm (hmac-md5-96 | hmac-sha1-96 | hmac-sha-256-128);
    key (ascii-text key | hexadecimal key);
}
```

Hierarchy Level

```
[edit services (IPsec VPN) ipsec-vpn rule rule-name term term-name then manual direction direction]
```

Description

Configure IPsec authentication parameters for a manual security association (SA).

Options

algorithm Hash algorithm that authenticates packet data. The algorithm can be one of the following:

- **hmac-md5-96** Produces a 128-bit digest.
- **hmac-sha1-96** Produces a 160-bit digest.
- **hmac-sha-256-128** Produces a 256-bit digest, truncated to 128 bits.

key Type of authentication key. The key can be one of the following:

- **ascii-text key** ASCII text key. For `hmac-md5-96`, the key is 16 ASCII characters; for `hmac-sha1-96`, the key is 20 ASCII characters.
- **hexadecimal key** Hexadecimal key. For `hmac-md5-96`, the key is 32 hexadecimal characters; for `hmac-sha1-96`, the key is 40 hexadecimal characters.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Security Associations

authentication

IN THIS SECTION

- Syntax | 2574
- Hierarchy Level | 2575
- Description | 2575
- Required Privilege Level | 2575
- Release Information | 2575

Syntax

```
authentication {
    packet-types  [packet-types];
    password password-string;
    username-include {
        circuit-id;
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-18;
        option-37;
        option-82 <circuit-id> <remote-id>;
        radius-realm radius-realm-string;
```

```
        remote-id;
        user-prefix user-prefix-string;
        vlan-tags;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges],
[edit interfaces interface-name auto-configure stacked-vlan-ranges]
```

Description

Specify the authentication parameters that trigger the Access-Request message to AAA for the interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Configuring the Static Subscriber Global Authentication Password

Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

authentication (AAA Profile)

IN THIS SECTION

- Syntax | 2576
- Hierarchy Level | 2576
- Description | 2576
- Required Privilege Level | 2576
- Release Information | 2577

Syntax

```
authentication {
    network-element network-element-name;
}
```

Hierarchy Level

```
[edit unified-edge aaa profiles aaa-profile-name radius]
```

Description

Specify the network element providing policy management for TDF subscribers. The network element must already be defined at the `[edit access radius]` hierarchy level.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring an AAA Profile

Understanding AAA Profiles

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

authentication-algorithm

IN THIS SECTION

- Syntax | 2577
- Hierarchy Level | 2577
- Description | 2578
- Options | 2579
- Required Privilege Level | 2580
- Release Information | 2580

Syntax

```
authentication-algorithm algorithm;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name protocols ldp session session-address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit logical-systems logical-system-name  routing-instances routing-instance-name protocols ldp
session session-address],
[edit logical-systems logical-system-name routing-options bmp],
[edit logical-systems logical-system-name routing-options bmp station station-name],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit protocols ldp session session-address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols ldp session session-address],
[edit routing-options bmp],
[edit routing-options bmp station station-name]
```

Description

Configure an authentication algorithm type.

NOTE: Keep the following points in mind when you configure the authentication algorithm in an IPsec proposal:

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, an error occurs and the tunnel is not established in this scenario. For example, if one end of the tunnel contains router 1 configured with the authentication algorithm as hmac-sha- 256-128 and the other end of the tunnel contains router 2 configured with the authentication algorithm as hmac-md5-96, the VPN tunnel is not established.

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, and when one end of the tunnel contains two IPsec proposals to check whether a less secure algorithm is selected or not, an error occurs and the tunnel is not established. For example, if you configure two authentication algorithms for an IPsec proposal as hmac-sha-256-128 and hmac-md5-96 on one end of the tunnel, router 1, and if you configure the algorithm for an IPsec proposal as hmac-

md5-96 on the other end of the tunnel, router 2, the tunnel is not established and the number of proposals mismatch.

- When you configure two IPsec proposals at both ends of a tunnel, such as the `authentication-algorithm hmac-sha-256-128` and `authentication- algorithm hmac-md5-96` statements at the [`edit services ipsec-vpn ipsec proposal` *`proposal-name`*] hierarchy level on one of the tunnel, router 1 (with the algorithms in two successive statements to specify the order), and the `authentication-algorithm hmac-md5-96` and `authentication-algorithm hmac-sha-256-128` statements at the [`edit services ipsec-vpn ipsec proposal` *`proposal-name`*] hierarchy level on one of the tunnel, router 2 (with the algorithms in two successive statements to specify the order, which is the reverse order of router 1), the tunnel is established in this combination as expected because the number of proposals is the same on both ends and they contain the same set of algorithms. However, the authentication algorithm selected is hmac-md5-96 and not the stronger algorithm of hmac-sha-256-128. This method of selection of the algorithm occurs because the first matching proposal is selected. Also, for a default proposal, regardless of whether the router supports the Advanced Encryption Standard (AES) encryption algorithm, the 3des-cbc algorithm is chosen and not the aes-cfb algorithm, which is because of the first algorithm in the default proposal being selected. In the sample scenario described here, on router 2, if you reverse the order of the algorithm configuration in the proposal so that it is the same order as the one specified on router 1, hmac-sha-256-128 is selected as the authentication method.

- You must be aware of the order of proposals in an IPsec policy at the time of configuration if you want the matching of proposals to happen in a certain order of preference, such as the strongest algorithm to be considered first when a match is made when both policies from the two peers have a proposal.

Options

`algorithm`—Specify one of the following types of authentication algorithms:

- `aes-128-cmac-96`—Cipher-based message authentication code (AES128, 96 bits).
- `hmac-sha-1-96`—Hash-based message authentication code (SHA1, 96 bits).
- `md5`—Message digest 5.
- **Default:** `hmac-sha-1-96`

> **NOTE**: The default is not displayed in the output of the `show bgp bmp` command unless a key or key-chain is also configured.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Statement introduced for BGP in Junos OS Release 8.0.

Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced for BMP in Junos OS Release 13.3.

RELATED DOCUMENTATION

Example: Configuring Router Authentication for BGP

Configuring BGP Monitoring Protocol Version 3

authentication-algorithm (Security IKE)

IN THIS SECTION

- Syntax | 2581
- Hierarchy Level | 2581
- Description | 2581
- Options | 2581
- Required Privilege Level | 2581
- Release Information | 2581

Syntax

```
authentication-algorithm (md5 | sha1 | sha-256 | sha-384);
```

Hierarchy Level

```
[edit security ike proposal ike-proposal-name]
```

Description

Configure the Internet Key Exchange (IKE) authentication algorithm.

Options

`authentication-algorithm`—Hash algorithm that authenticates packet data. It can be one of two algorithms:

- `md5`—Produces a 128-bit digest.
- `sha1`—Produces a 160-bit digest.
- `sha256`—Produces a 256-bit digest.
- `sha384`—Produces a 384-bit digest.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Authentication Algorithm for an IKE Proposal

authentication-algorithm (Security IPsec)

IN THIS SECTION

- Syntax | 2582
- Hierarchy Level | 2582
- Description | 2582
- Options | 2583
- Required Privilege Level | 2584
- Release Information | 2584

Syntax

```
authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
```

Hierarchy Level

```
[edit security ipsec proposal ipsec-proposal-name]
```

Description

Configure the IPsec authentication algorithm.

NOTE: Keep the following points in mind when you configure the authentication algorithm in an IPsec proposal:

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, an error occurs and the tunnel is not established in this scenario. For example, if one end of the tunnel contains router 1 configured with the authentication algorithm as hmac-sha- 256-128 and the other end of the tunnel contains router 2 configured with the authentication algorithm as hmac-md5-96, the VPN tunnel is not established.

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, and when one end of the tunnel contains two IPsec proposals to check whether a less secure algorithm is selected or not, an error occurs and the tunnel is not established. For example, if you configure two authentication algorithms for an IPsec proposal as hmac-sha-256-128 and hmac-md5-96 on one end of the tunnel, router 1, and if you configure the algorithm for an IPsec proposal as hmac-md5-96 on the other end of the tunnel, router 2, the tunnel is not established and the number of proposals mismatch.

- When you configure two IPsec proposals at both ends of a tunnel, such as the `authentication-algorithm hmac-sha-256-128` and `authentication- algorithm hmac-md5-96` statements at the `[edit services ipsec-vpn ipsec proposal` *`proposal-name`*`]` hierarchy level on one of the tunnel, router 1 (with the algorithms in two successive statements to specify the order), and the `authentication-algorithm hmac-md5-96` and `authentication-algorithm hmac-sha-256-128` statements at the `[edit services ipsec-vpn ipsec proposal` *`proposal-name`*`]` hierarchy level on one of the tunnel, router 2 (with the algorithms in two successive statements to specify the order, which is the reverse order of router 1), the tunnel is established in this combination as expected because the number of proposals is the same on both ends and they contain the same set of algorithms. However, the authentication algorithm selected is hmac-md5-96 and not the stronger algorithm of hmac-sha-256-128. This method of selection of the algorithm occurs because the first matching proposal is selected. Also, for a default proposal, regardless of whether the router supports the Advanced Encryption Standard (AES) encryption algorithm, the 3des-cbc algorithm is chosen and not the aes-cfb algorithm, which is because of the first algorithm in the default proposal being selected. In the sample scenario described here, on router 2, if you reverse the order of the algorithm configuration in the proposal so that it is the same order as the one specified on router 1, hmac-sha-256-128 is selected as the authentication method.

- You must be aware of the order of proposals in an IPsec policy at the time of configuration if you want the matching of proposals to happen in a certain order of preference, such as the strongest algorithm to be considered first when a match is made when both policies from the two peers have a proposal.

Options

`authentication-algorithm`—Hash algorithm that authenticates packet data. It can be one of two algorithms:

- `hmac-md5-96`—Produces a 128-bit digest.
- `hmac-sha1-96`—Produces a 160-bit digest.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Authentication Algorithm for an IPsec Proposal

authentication-algorithm (Services IKE)

IN THIS SECTION

- Syntax | 2584
- Hierarchy Level | 2584
- Description | 2585
- Options | 2585
- Required Privilege Level | 2585
- Release Information | 2585

Syntax

```
authentication-algorithm (md5 | sha1 | sha-256);
```

Hierarchy Level

```
[edit services ipsec-vpn ike proposal proposal-name]
```

Description

Configure the Internet Key Exchange (IKE) hash algorithm that authenticates packet data.

Options

`md5`—Produces a 128-bit digest.

`sha1`—Produces a 160-bit digest.

`sha-256`—Produces a 256-bit digest.

`sha-384`—Produces a 384-bit digest.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

sha-256 option added in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring IKE Proposals

authentication-algorithm (Services IPsec)

IN THIS SECTION

- Syntax | 2586
- Hierarchy Level | 2586
- Description | 2586

- Options | 2587
- Required Privilege Level | 2587
- Release Information | 2588

Syntax

```
authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
```

Hierarchy Level

```
[edit services ipsec-vpn ipsec proposal ipsec-proposal-name]
```

Description

Configure the IPsec hash algorithm that authenticates packet data.

NOTE: Keep the following points in mind when you configure the authentication algorithm in an IPsec proposal:

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, an error occurs and the tunnel is not established in this scenario. For example, if one end of the tunnel contains router 1 configured with the authentication algorithm as hmac-sha- 256-128 and the other end of the tunnel contains router 2 configured with the authentication algorithm as hmac-md5-96, the VPN tunnel is not established.

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, and when one end of the tunnel contains two IPsec proposals to check whether a less secure algorithm is selected or not, an error occurs and the tunnel is not established. For example, if you configure two authentication algorithms for an IPsec proposal as hmac-sha-256-128 and hmac-md5-96 on one end of the tunnel, router 1, and if you configure the algorithm for an IPsec proposal as hmac-

md5-96 on the other end of the tunnel, router 2, the tunnel is not established and the number of proposals mismatch.

- When you configure two IPsec proposals at both ends of a tunnel, such as the `authentication-algorithm hmac-sha-256-128` and `authentication- algorithm hmac-md5-96` statements at the [`edit services ipsec-vpn ipsec proposal` *`proposal-name`*] hierarchy level on one of the tunnel, router 1 (with the algorithms in two successive statements to specify the order), and the `authentication-algorithm hmac-md5-96` and `authentication-algorithm hmac-sha-256-128` statements at the [`edit services ipsec-vpn ipsec proposal` *`proposal-name`*] hierarchy level on one of the tunnel, router 2 (with the algorithms in two successive statements to specify the order, which is the reverse order of router 1), the tunnel is established in this combination as expected because the number of proposals is the same on both ends and they contain the same set of algorithms. However, the authentication algorithm selected is hmac-md5-96 and not the stronger algorithm of hmac-sha-256-128. This method of selection of the algorithm occurs because the first matching proposal is selected. Also, for a default proposal, regardless of whether the router supports the Advanced Encryption Standard (AES) encryption algorithm, the 3des-cbc algorithm is chosen and not the aes-cfb algorithm, which is because of the first algorithm in the default proposal being selected. In the sample scenario described here, on router 2, if you reverse the order of the algorithm configuration in the proposal so that it is the same order as the one specified on router 1, hmac-sha-256-128 is selected as the authentication method.

- You must be aware of the order of proposals in an IPsec policy at the time of configuration if you want the matching of proposals to happen in a certain order of preference, such as the strongest algorithm to be considered first when a match is made when both policies from the two peers have a proposal.

Options

`hmac-md5-96`—Produces a 128-bit digest.

`hmac-sha-256-128`—Produces a 256-bit digest.

`hmac-sha1-96`—Produces a 160-bit digest.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Proposals

authentication (BGP BFD Liveness Detection)

IN THIS SECTION

- Syntax | 2588
- Hierarchy Level | 2588
- Description | 2589
- Options | 2589
- Required Privilege Level | 2591
- Release Information | 2591

Syntax

```
authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
}
```

Hierarchy Level

```
[edit logical-systems name protocols bgp bfd-liveness-detection],
[edit logical-systems name protocols bgp group name bfd-liveness-
detection],
```

```
[edit logical-systems name protocols bgp group name neighbor
address bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols bgp bfd-
liveness-detection],
[edit logical-systems name routing-instances name protocols
bgp group name bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols
bgp group name neighbor address bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols
bgp bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols
bgp group name bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols
bgp group name neighbor address bfd-liveness-detection],
[edit protocols bgp bfd-liveness-detection],
[edit protocols bgp group name bfd-liveness-detection],
[edit protocols bgp group name neighbor address bfd-liveness-detection],
[edit routing-instances name protocols bgp bfd-liveness-detection],
[edit routing-instances name protocols bgp group name bfd-liveness-
detection],
[edit routing-instances name protocols bgp group name neighbor
address bfd-liveness-detection],
[edit tenants name routing-instances name protocols bgp bfd-liveness-detection]
[edit tenants name routing-instances name protocols bgp group name bfd-liveness-
detection]
[edit tenants name routing-instances name protocols bgp group name neighbor
address bfd-liveness-detection]
```

Description

Specify the router and route authentication to mitigate the risk of being attacked by a machine or router that has been configured to share incorrect routing information with another router. Router and route authentication enables routers to share information only if they can verify that they are talking to a trusted source, based on a password (key). In this method, a hashed key is sent along with the route being sent to another router. The receiving router compares the sent key to its own configured key. If they are the same, the receiving router accepts the route.

Options

authentication *algorithm-name* Configure the algorithm used to authenticate the specified BFD session.

- **Values:** Specify one of these algorithm names:
  - simple-password—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
  - keyed-md5—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
  - meticulous-keyed-md5—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method can take additional time to authenticate the session.
  - keyed-sha-1—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
  - meticulous-keyed-sha-1—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method can take additional time to authenticate the session.

key-chain *key-chain-name*

Specify the name of an authentication keychain. The keychain name must match one of the keychains configured with the `key-chain` *`key-chain-name`* statement at the `[edit security authentication-key-chain]` hierarchy level.

The authentication keychain associates a security key with the specified BFD session. Each key has a unique start time within the keychain. Keychain authentication allows you to change the password information periodically without bringing down peering sessions. This keychain authentication method is referred to as *hitless* because the

keys roll over from one to the next without resetting any peering sessions or interrupting the routing protocol.

loose-check Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from non-authenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.

- **Default:** Strict authentication is enabled. Specify loose-check to disable it.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

Support for BFD authentication introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

Example: Configuring BFD Authentication for Securing Static Routes

Example: Configuring Router Authentication for BGP

Example: Configuring BFD on Internal BGP Peer Sessions

Example: Configuring EBGP Multihop Sessions

bfd-liveness-detection (BGP) | **2918**

authentication (Demux)

IN THIS SECTION

- Syntax | 2592
- Hierarchy Level | 2592
- Description | 2592
- Required Privilege Level | 2593
- Release Information | 2593

Syntax

```
authentication {
    password password-string;
    username-include {
        auth-server-realm realm-string;
        delimiter delimiter-character;
        domain-name domain-name;
        interface-name;
        source-address;
        user-prefix user-prefix-string;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet auto-configure address-ranges]
[edit interfaces interface-name unit unit-number demux inet6 auto-configure address-ranges]
```

Description

Specify the authentication parameters for the demultiplexing (demux) interface options. The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

> *IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview*
>
> *Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles*

authentication (DHCP Local Server)

IN THIS SECTION

- Syntax | 2593
- Hierarchy Level | 2594
- Description | 2594
- Required Privilege Level | 2594
- Release Information | 2595

Syntax

```
authentication {
    password password-string;
    username-include {
        circuit-type;
        client-id;
        delimiter delimiter-character;
        domain-name domain-name-string;
```

```
        interface-description  (device-interface | logical-interface);
        interface-name ;
        logical-system-name;
        mac-address;
        option-60;
        option-82 <circuit-id> <remote-id>;
        relay-agent-interface-id;
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix-string;
        vlan-tags;
    }
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay or DHCP local server configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

authentication (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 2595
- Hierarchy Level | 2596
- Description | 2596
- Required Privilege Level | 2596
- Release Information | 2596

Syntax

```
authentication {
    password password-string;
    username-include {
        circuit-type;
        client-id;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-description  (device-interface | logical-interface);
        interface-name;
        logical-system-name;
        mac-address;
        option-60;
        option-82 <circuit-id> <remote-id>;
        relay-agent-interface-id;
```

```
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix-string;
        vlan-tags;
    }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay configuration. Use the statement at the `[edit...dhcpv6]` hierarchy levels to configure DHCPv6 support.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the [edit ... dual-stack-group *dual-stack-group-name*] hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

dhcp-relay | 5620

Specifying Authentication Support

authentication-entry-timeout (Identity Management Advanced Query)

IN THIS SECTION

- Syntax | 2597
- Hierarchy Level | 2597
- Description | 2597
- Options | 2598
- Required Privilege Level | 2598
- Release Information | 2598

Syntax

```
authentication-entry-timeout time-out-in-minutes;
```

Hierarchy Level

```
[edit services user-identification identity-management]
```

Description

Configure the time-out for the user identity authentication entries. You configure this parameter as part of the advanced user identity query feature for SRX Series Firewalls.

The advanced user identity query feature for SRX Series Firewalls relies on the Juniper Identity Management Service (JIMS), a centralized identity collection (CIC) system from which the SRX Series Firewall obtains the user identity information. It provides a global, end-to-end user identity management solution that allows you to provision users locally and have their authentication information made available to other sites in your network for policy enforcement and reporting.

WARNING: Before you use this feature, you must disable active-directory-access and authentication-source options under the user-identification hierarchy. You cannot commit this configuration if active directory authentication or the ClearPass query and webapi functions are configured and committed.

Options

time-out-in-minutes The amount of time after which a user identity authentication entry expires.

- **Range:** 0 or 10 through 1440 minutes. Specification of 0 indicates no time-out.
- **Default:** 60 minutes

Required Privilege Level

1. services—To view this statement in the configuration.
2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Understanding Advanced Query Feature for Obtaining User Identity Information from JIMS

batch query | **2886**

filter (Identity Management Advanced Query) | **7138**

secondary connection (Identity Management Advanced Query) | **16409**

authentication-entry-timeout (Services User Identification)

IN THIS SECTION

- Syntax | 2599
- Hierarchy Level | 2599
- Description | 2599
- Options | 2599
- Required Privilege Level | 2600
- Release Information | 2600

Syntax

```
authentication-entry-timeout minutes;
```

Hierarchy Level

```
[edit services user-identification authentication-source (Services User Identification ClearPass)
aruba-clearpass]
```

Description

Configure for the integrated ClearPass authentication and enforcement feature the timeout interval after which idle entries in the ClearPass authentication table expire.

Options

minutes Timeout interval. The timeout interval begins from when the user authentication entry is added to the ClearPass authentication table. If a value of 0 is specified, the entries will never expire.

- **Range:** 10 through 1440 minutes
- **Default:** 30 minutes

Required Privilege Level

1. services—To view this statement in the configuration
2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

authentication-key (Protocols MSDP)

IN THIS SECTION

- Syntax | 2600
- Hierarchy Level | 2600
- Description | 2601
- Default | 2601
- Options | 2601
- Required Privilege Level | 2601
- Release Information | 2601

Syntax

```
authentication-key peer-key;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
```

```
peer address],
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit routing-instances routing-instance-name protocols msdp group group-name peer address],
[edit routing-instances routing-instance-name protocols msdp peer address]
```

Description

Associate a Message Digest 5 (MD5) signature option authentication key with an MSDP peering session.

Default

If you do not include this statement, the routing device accepts any valid MSDP messages from the peer address.

Options

peer-key—MD5 authentication key. The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of (,), &, and [. If you include spaces in an MSDP authentication key, enclose all characters in quotation marks (“ ”).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP in a Routing Instance

authentication-key (SNMP)

IN THIS SECTION

- Syntax | 2602
- Hierarchy Level | 2602
- Description | 2602
- Options | 2602
- Required Privilege Level | 2603
- Release Information | 2603

Syntax

```
authentication-key authentication-key;
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username authentication-md5],
[edit snmp v3 usm local-engine user username authentication-sha],
[edit snmp v3 usm remote-engine engine-id user username authentication-md5],
[edit snmp v3 usm remote-engine engine-id user username authentication-sha]
```

Description

Configure the key for user authentication, if you know it, instead of entering a cleartext password. The encrypted version of a password is known as a key.

Options

authentication-key	The encrypted version from the cleartext password.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

authentication-key

IN THIS SECTION

- Syntax | 2603
- Hierarchy Level | 2603
- Description | 2604
- Options | 2604
- Required Privilege Level | 2604
- Release Information | 2604

Syntax

```
authentication-key key;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-
group group-id],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet address address vrrp-group group-id]
```

Description

Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the `authentication-type` statement.

All devices in the VRRP group must use the same authentication scheme and password.

> **NOTE**: When VRRPv3 is enabled, the `authentication-type` and `authentication-key` statements cannot be configured for any VRRP groups.

Options

key—Authentication password. For simple authentication, it can be 1 through 8 characters long. For Message Digest 5 (MD5) authentication, it can be 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (“ ”).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring VRRP Authentication (IPv4 Only)

authentication-type | **2667**

version-3

Understanding VRRP on SRX Series Devices

Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces

authentication-key-chain (Protocols BGP and BMP)

IN THIS SECTION

- Syntax | 2605
- Hierarchy Level | 2605
- Description | 2606
- Options | 2606
- Required Privilege Level | 2606
- Release Information | 2606

Syntax

```
authentication-key-chain key-chain;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit logical-systems logical-system-name routing-options bmp],
[edit logical-systems logical-system-name routing-options bmp station station-name],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address],
```

```
[edit routing-options bmp],
[edit routing-options bmp station station-name]
```

Description

Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update feature for BGP, you cannot commit the `0.0.0.0/allow` statement with authentication keys or key chains. The CLI issues a warning and fails to commit the configuration.

Options

key-chain—Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

> **NOTE**: For BGP, you must also configure an authentication algorithm by including the `authentication-algorithm` *algorithm* `statement`.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced for BMP in Junos OS Release 13.3.

RELATED DOCUMENTATION

Example: Configuring Router Authentication for BGP

Example: Configuring BFD Authentication for Securing Static Routes

Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols

Configuring BGP Monitoring Protocol Version 3

authentication-algorithm | **2577**

authentication-key-chain (Protocols LDP)

IN THIS SECTION

- Syntax | 2607
- Hierarchy Level | 2607
- Description | 2607
- Options | 2608
- Required Privilege Level | 2608
- Release Information | 2608

Syntax

```
authentication-key-chain key-chain;
```

Hierarchy Level

```
[edit logical-systems name protocols ldp session address],
[edit logical-systems name routing-instances instance-name protocols ldp session address],
[edit protocols ldp session address],
[edit routing-instances instance-name protocols ldp session address]
```

Description

Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for LDP, you cannot commit the `0.0.0.0/allow` statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.

NOTE: You must also configure an authentication algorithm using the `authentication-algorithm` statement.

Options

key-chain—Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols

Configuring Miscellaneous LDP Properties

authentication-algorithm | 2577

authentication-key-chain (TWAMP)

IN THIS SECTION

- Syntax | 2609
- Hierarchy Level | 2609
- Description | 2609
- Options | 2609
- Required Privilege Level | 2609
- Release Information | 2609

Syntax

```
authentication-key-chain identifier {
    key-id identifier {
        secret password-string;
    }
}
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Apply and enable an authentication key chain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for TWAMP, you cannot commit the `0.0.0.0/allow` statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.

Options

identifier—Authentication key chain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

password-string—Authentication key, consisting of 1 through 8 ASCII characters. If the key contains spaces, enclose it in quotation marks.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

authentication-key-chains

IN THIS SECTION

- Syntax | 2610
- Hierarchy Level | 2611
- Description | 2611
- Required Privilege Level | 2611
- Release Information | 2611

Syntax

```
authentication-key-chains {
    key-chain key-chain-name {
        description text-string;
        key key {
            algorithm (md5 | hmac-sha-1);
            key-name authentication-key-name
            algorithm (hmac-sha2-224 | hmac-sha2-256| hmac-sha2-384 | hmac-sha2-512);
            options (basic | isis-enhanced);
            secret secret-data;
            start-time yyyy-mm-dd.hh:mm:ss;
        }
        tolerance seconds;
    }
}
```

Hierarchy Level

```
[edit security]
[edit routing-instances vrf-name protocols bgp group group-name]
```

Description

Configure authentication key updates for the Border Gateway Protocol (BGP) and the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, the Intermediate System-to-Intermediate System (IS-IS) protocol, and the Open Shortest Path First (OSPF) protocol. When the `authentication-key-chains` statement is configured at the `[edit security]` hierarchy level, and is associated with the BGP, LDP, IS-IS protocol, or OSPF protocol at the `[edit protocols]` hierarchy level or with the BFD protocol using the `bfd-liveness-detection` statement, authentication key updates can occur without interrupting signaling protocols such as Resource Reservation Setup Protocol (RSVP).

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Support for the BFD protocol introduced in Junos OS Release 9.6.

Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.

Support for IS-IS introduced in JUNOS OS Release 11.2.

Support for HMAC-SHA2 authentication algorithm that includes HMAC-SHA2-256, HMAC-SHA2-384 and HMAC-SHA2-512 for OSPFv2 protocol introduced in Junos OS and Junos OS Evolved Release 24.2R1.

Support for HMAC-SHA2-224 authentication for OSPFv2 protocol introduced in Junos OS and Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

BGP Route Authentication

Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols

Example: Configuring BFD Authentication for Securing Static Routes

Example: Configuring Hitless Authentication Key Rollover for IS-IS

Configuring Media Access Control Security (MACsec) on Routers

authentication-key-chains (TCP-AO)

IN THIS SECTION

- Syntax | **2612**
- Hierarchy Level | **2613**
- Description | **2613**
- Required Privilege Level | **2614**
- Release Information | **2614**

Syntax

```
authentication-key-chains {
    key-chain keychain-name {
        key key ID {
            secret secretpassword;
            start-time yyyy-mm-dd.hh:mm:ss;
            algorithm ao;
            ao-attribute {
                send-id send ID;
                    recv-id receive ID;
                    cryptographic-algorithm (aes-128-cmac-96 | hmac-sha-1-96);
                tcp-ao-option (disabled | enabled);
            }
        }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure authentication keychains for TCP Authentication Option (TCP-AO) .

Table 3: Options

Option	Description
`key-chain`	Enter a unique name for the keychain. For example, `new-auth-key`.
`key`	Enter a unique key ID for each key. In a key-chain, keys are numbered sequentially, from `key` *0* through `key` *63*.
`secret`	Enter a unique secret key or password for each key. Use any alphanumeric characters without any space. Once configured, it will appear in an encrypted format.
`start-time`	Enter a time in *YYYY-MM-DD.HH:MM* format to specify the time when the control gets passed on from one key to the next. When a configured start time arrives (based on the device's clock), the key with that start time becomes active.
`algorithm`	Enter `ao` to indicate the authentication option.
`send-id`	Enter any two numbers between 0 and 255. You can also use the same number as the `send-id` and the `recv-id` of the same key. You must not use this numbers for any other key inside that key chain. You can configure up to 64 keys in a key-chain.
`recv-id`	Enter any two numbers between 0 and 255. You must not use this numbers for any other key inside that key chain. Reverse the `send-id` and `recv-id` values when you configure the same key in the receiving device.
`cryptographic-algorithm`	Choose a cryptographic algorithm. Starting in Junos OS Release 20.3R1, to be compliant to *RFC5925, The TCP Authentication Option*, we are supporting HMAC-SHA1 and AES-128 algorithms.

Table 3: Options *(Continued)*

Option	Description
`tcp-ao-option`	Choose enable to include the TCP-AO option. **NOTE**: The default value is disabled.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

TCP Authentication Option (TCP-AO)

authentication-key (ICCP)

IN THIS SECTION

- Syntax | 2615
- Hierarchy Level | 2615
- Description | 2615
- Options | 2615
- Required Privilege Level | 2615
- Release Information | 2615

Syntax

```
authentication-key key;
```

Hierarchy Level

```
[edit protocols iccp],
[edit protocols iccp peer peer-IP-address]
```

Description

Specify the authentication key (password). The QFX3500 and MX Series devices use this password to verify the authenticity of packets sent from the peers hosting a multichassis link aggregation group (MC-LAG). Peer-level authentication takes precedence over global-level authentication.

Inter-Chassis Control Protocol (ICCP) uses MD5 authentication.

Options

key—Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

authentication-key (Protocols BGP and BMP)

IN THIS SECTION

- Syntax | 2616
- Hierarchy Level | 2616
- Description | 2617
- Options | 2617
- Required Privilege Level | 2617
- Release Information | 2617

Syntax

```
authentication-key key;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit logical-systems logical-system-name routing-options bmp],
[edit logical-systems logical-system-name routing-options bmp station station-name],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address],
```

```
[edit routing-options bmp],
[edit routing-options bmp station station-name]
```

Description

Configure an MD5 authentication key (password). Neighboring routing devices use the same password to verify the authenticity of BGP packets sent from this system.

Options

key—Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced for BMP version 3 in Junos OS Release 13.3.

RELATED DOCUMENTATION

Example: Configuring Router Authentication for BGP

Configuring BGP Monitoring Protocol Version 3

authentication-key (Protocols LDP)

IN THIS SECTION

- Syntax | 2618
- Hierarchy Level | 2618
- Description | 2618
- Required Privilege Level | 2618
- Release Information | 2619

Syntax

```
authentication-key md5-authentication-key;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp session address]
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp
session address]
[edit protocols ldp session address]
[edit protocols ldp session-group address]
[edit routing-instances routing-instance-name protocols ldp session address]
```

Description

Configure the MD5 authentication signature. The maximum length of the authentication signature is 69 characters.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the TCP MD5 Signature for LDP Sessions

authentication-key (Protocols RIP)

IN THIS SECTION

- Syntax | 2619
- Hierarchy Level | 2619
- Description | 2620
- Options | 2620
- Required Privilege Level | 2620
- Release Information | 2620

Syntax

```
authentication-key password;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip],
[edit protocols rip group group-name neighbor neighbor-name],
```

```
[edit routing-instances routing-instance-name protocols rip],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-name]
```

Description

Require authentication for RIP route queries received on an interface.

Options

password—Authentication password. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring Route Authentication for RIP using single MD5 key

authentication-key (Protocols RSVP)

IN THIS SECTION

- Syntax | 2621
- Hierarchy Level | 2621
- Description | 2621
- Options | 2622

Required Privilege Level | 2622
Release Information | 2622

Syntax

```
authentication-key key;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit logical-systems logical-system-name protocols rsvp peer-interface peer-interface-name],
[edit protocols rsvp],
[edit protocols rsvp interface interface-name],
[edit protocols rsvp peer-interface peer-interface-name]
```

Description

Authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface or peer interface. To authenticate node hellos or remote messages between the Point of Local Repair (PLR) to the Merge Point (MP), enable `authentication-key` at the `[edit protocols rsvp]` hierarchy level.

It is recommended to use the `authentication-key` configuration at the `[edit protocols rsvp]` hierarchy level for the RSVP node-neighbor. Because non-RSVP interfaces do not have RSVP authentication key, the `authentication-key` configuration allows the packets to arrive at any interface regardless of RSVP interfaces configuration being enabled or not.

NOTE: To avoid fxp inheriting the `authentication-key` configuration statement, use specific physical interfaces such as <ge-*> and <ae*> in the group configuration. We don't recommend to use wildcard pattern <*>.

RSVP uses HMAC-MD5 authentication, which is defined in RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*.

All routers that are connected to the same IP subnet must use the same authentication scheme and password.

Options

key—Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring RSVP Authentication

authentication (LAG)

IN THIS SECTION

- Syntax | 2623
- Hierarchy Level | 2623
- Description | 2623
- Options | 2623
- Required Privilege Level | 2624
- Release Information | 2624

Syntax

```
authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
}
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options bfd-liveness-detection]
```

Description

Configure the authentication criteria of the BFD session for aggregated Ethernet interfaces.

Options

algorithm *algorithm-name*

Specify the algorithm to be used to authenticate the BFD session. You can use one of the following algorithms for authentication:

- keyed-md5
- keyed-sha-1
- meticulous-keyed-md5
- meticulous-keyed-sha-1
- simple-password

key-chain *key-chain-name*

Specify the name that is associated with the security key for the BFD session. The name you specify must match one of the keychains configured in the `authentication-key-chains key-chain` statement at the `[edit security]` hierarchy level.

loose-check

(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

- bfd-liveness-detection (LAG) | 2926
- detection-time (LAG) | 5549
- transmit-interval (LAG) | 30090
- *Configuring Micro BFD Sessions for LAG*
- Example: Configuring Independent Micro BFD Sessions for LAG
- Understanding Independent Micro BFD Sessions for LAG

authentication-md5

IN THIS SECTION

- Syntax | 2625
- Hierarchy Level | 2625
- Description | 2625
- Required Privilege Level | 2625
- Release Information | 2625

Syntax

```
authentication-md5 {
    (authentication-key authentication-key | authentication-password authentication-password);
}
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]
```

Description

Configure MD5 as the authentication type for the SNMPv3 user.

NOTE: You can only configure one authentication type for each SNMPv3 user.

For authentication, you can either enter a password in clear text, which is immediately encrypted, or, if you already have the encrypted version (known as a *key*), enter the key directly.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring MD5 Authentication

authentication-method (services)

IN THIS SECTION

- Syntax | 2626
- Hierarchy Level | 2626
- Description | 2626
- Options | 2626
- Required Privilege Level | 2627
- Release Information | 2627

Syntax

```
authentication-method (ecdsa-signatures-256 | ecdsa-signatures-384 | pre-shared-keys | rsa-
signatures);
```

Hierarchy Level

```
[edit services (IPsec VPN) ipsec-vpn ike proposal proposal-name]
```

Description

Configure an IKE authentication method.

Options

NOTE: In Junos FIPS mode, ECDSA is not supported in Junos OS Release 17.3R1. Starting in Junos OS Release 17.4R1, ECDSA is supported in Junos FIPS mode.

ecdsa-signatures-256 Elliptic Curve Digital Signature Algorithm (ECDSA) for 256-bit moduli. This can only be used on an MS-MPC or MS-MIC.

ecdsa-signatures-384 ECDSA for 384-bit moduli. This can only be used on an MS-MPC or MS-MIC.

pre-shared-keys A key derived from an out-of-band mechanism; the key authenticates the exchange.

rsa-signatures Public key algorithm (supports encryption and digital signatures).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`ecdsa-signatures-256` and `ecdsa-signatures-384` options added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring IKE Proposals

authentication-method

IN THIS SECTION

- Syntax | 2628
- Hierarchy Level | 2628
- Description | 2628
- Options | 2628
- Required Privilege Level | 2628
- Release Information | 2628

Syntax

```
authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
```

Hierarchy Level

```
[edit security ike proposal ike-proposal-name]
```

Description

Configure the IKE authentication method.

Options

`dsa-signatures`—Digital Signature Algorithm (DSA)

`rsa-signatures`—A public key algorithm, which supports encryption and digital signatures

`pre-shared-keys`—A key derived from an out-of-band mechanism; the key authenticates the exchange

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Manual IPsec Security Associations for an ES PIC

authentication-method (Interfaces)

IN THIS SECTION

- Syntax | 2629
- Hierarchy Level | 2629
- Description | 2629
- Options | 2629
- Required Privilege Level | 2630
- Release Information | 2630

Syntax

```
authentication-method (pap | chap | none);
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options profiles profile-name]
```

Description

Specify the authentication method for connection to a Global System for Mobile Communications (GSM) cellular network.

Options

- `pap`—Password Authentication Protocol.
- `chap`—Challenge Handshake Authentication Protocol.
- `none`—No authentication method is used.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

authentication-mode

IN THIS SECTION

- Syntax | 2630
- Hierarchy Level | 2630
- Description | 2631
- Options | 2631
- Required Privilege Level | 2631
- Release Information | 2631

Syntax

```
authentication-mode (authenticated | encrypted | none);
```

Hierarchy Level

```
[edit services rpm twamp server],
[edit services rpm twamp client control-connection control-client-name]
```

Description

Specify the authentication or encryption mode support for the TWAMP test protocol. This statement is required in the configuration; if no authentication or encryption is specified, you must set the value to `none`.

Options

`authenticated` Authenticate all TWAMP packets.

> **NOTE**: This mode is supported only on TWAMP servers.

`encrypted` Encrypt all TWAMP packets.

> **NOTE**: This mode is supported only on TWAMP servers.

`none` Do not authenticate or encyrpt packets.

> **NOTE**: This mode is supported on both TWAMP servers and clients.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the [`edit services rpm twamp client control-connection` *`control-client-name`*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

RELATED DOCUMENTATION

> *Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches*
>
> *Understanding Two-Way Active Measurement Protocol on Routers*

authentication-none

IN THIS SECTION

- Syntax | 2632
- Hierarchy Level | 2632
- Description | 2632
- Required Privilege Level | 2632
- Release Information | 2633

Syntax

```
authentication-none;
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]
```

Description

Configure that there should be no authentication for the SNMPv3 user.

NOTE: You can configure only one authentication type for each SNMPv3 user.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring No Authentication

authentication-order

IN THIS SECTION

- Syntax | 2633
- Hierarchy Level | 2633
- Description | 2634
- Options | 2634
- Required Privilege Level | 2635
- Release Information | 2635

Syntax

```
authentication-order [ authentication-methods ];
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Set the order in which AAA tries different authentication methods when verifying that a client can access the router or switch. For each login attempt, AAA tries the authentication methods in order, from first to last.

A given subscriber does not undergo both authentication and authorization as separate steps. When both `authentication-order` and `authorization-order` are specified, DHCP subscribers honor the configured authorization order, all other subscribers use the configured authentication-order.

Starting in Junos OS Release 18.2R1, the `password` option can also be used to specify that local authentication and local authorization is attempted for individual subscribers that are configured with the `subscriber` statement at the [`edit access profile` *`profile-name`*] hierarchy level.

Options

authentication-methods

Ordered list of methods to use for authentication attempts. The list includes one or more of the following methods in any combination:

- `nasreq`—Verify subscribers using the Diameter-based Network Access Server Requirements (NASREQ) protocol.

- `none`—No authentication is performed. Grants authentication without examining the client credentials. Can be used, for example, when the Diameter function Gx-Plus is employed for notification during subscriber provisioning.

  > **NOTE**: Subscriber access management does not support the `none` option; authentication fails when this option is specified.

- `password`—Verify the client using the information configured at the [`edit access profile` *`profile-name`* `client` *`client-name`*] hierarchy level.

  Subscriber access management does not support the `password` option until Junos OS Release 18.2R1. Starting in Junos OS Release 18.2R1, this option is used to enable local authentication and optionally local authorization for individual subscribers. Local authentication is typically used when you do not have external authentication and authorization servers. The password itself must be configured with the `subscriber` statement in the same access profile. Local authentication is performed when a subscriber logs in with a matching username; it succeeds if the subscribers login password matches the password in the profile.

If you have external authentication and authorization servers, you can use local authentication as a backup authentication method. In this case, configure password other than first in the list of methods.

- radius—Verify the client using RADIUS authentication services.
- s6a—Verify subscribers using the Diameter-based s6a protocol.
- **Default:** password

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

none option added in Junos OS Release 11.2.

nasreq option added in Junos OS Release 16.1.

s6a option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Example: Configuring CHAP Authentication with RADIUS

RADIUS Authentication and Accounting Basic Configuration

Example: Configure S6a Application

authentication-order (Access Profile)

IN THIS SECTION

Syntax | 2636

- Hierarchy Level | 2636
- Description | 2636
- Default | 2636
- Options | 2636
- Required Privilege Level | 2637
- Release Information | 2637

Syntax

```
authentication-order [(none | ldap | password | radius | s6a | saml | secureid)];
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure the order of authentication, authorization, and accounting (AAA) methods to use while sending authentication messages.

Default

Not enabled

Options

`none`—No authentication for specified users.

When you enable `none` authentication option, the SRX Series Firewall no longer requires the RADIUS server to authenticate the initiator again with the common shared password used for IKEv2 configuration payload. This is because, the SRX Series Firewall already authenticates the remote peer using a certificated-based authentication. You can use this AAA profile in different combinations, but ensure that it is not used where you do not use a pre-authentication.

For example: Consider a scenario, where to establish a connection from a client to secure gateway IPsec tunnels, client is authenticated using certificates method as per IKE protocol. For simplicity, if you do not prefer dependency on RADIUS server and use local pool for address acquisition without any additional authentication, you can configure the “aaa” profile in IKE gateway hierarchy and set the `authentication-order` value as `none` in the access profile, as follows:

```
set access profile profile-name authentication-order none
set security ike gateway gateway name aaa access-profile profile-name
```

To start the Radius Accounting, use the following commands:

```
set access profile profile-anme radius-server 172.16.0.0 secret "$ABC123"
set access profile profile-name radius accounting-server 172.16.0.0
set access profile profile-name accounting order radius
```

`ldap`—Light weight directory access protocol.

`password`—Locally configured password in access profile.

`radius`—RADIUS authentication.

`s6a`—S6a authentication

`saml`—Security Assertion Markup Language

`securid`—RSA Secure ID authentication.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`none` option introduced in Junos OS Release 20.3R1.

`saml` option added in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

- Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch
- Configuring 802.1X RADIUS Accounting (CLI Procedure)
- Example: Specifying RADIUS Server Connections on a Security Device
- Ethernet Switching User Guide
- saml | **16281**

authentication-order (Authenticator)

IN THIS SECTION

- Syntax | **2638**
- Hierarchy Level | **2638**
- Description | **2639**
- Default | **2640**
- Options | **2640**
- Required Privilege Level | **2641**
- Release Information | **2641**

Syntax

```
authentication-order [dot1x | mac-radius | captive-portal];
```

Hierarchy Level

```
[edit logical-systems name protocols dot1x authenticator
interface],
[edit protocols dot1x authenticator interface interface-name]
```

Description

Configure the preferred order of authentication methods that the device will use when attempting to authenticate a client. If multiple authentication methods are configured on a single interface, when one authentication method fails, the device falls back to another method. You can configure the `authentication-order` statement to specify whether 802.1X authentication or MAC RADIUS authentication must be the first authentication method tried.

By default, the device attempts to authenticate a client by using 802.1X authentication first. If 802.1X authentication fails because there is no response from the client, and MAC RADIUS authentication is configured on the interface, the device falls back to MAC RADIUS authentication. If MAC RADIUS fails, and captive portal is configured on the device, the device falls back to captive portal.

Configuring MAC RADIUS authentication as the first method can help prevent the fallback timeout period which occurs after an 802.1X authentication attempt is made for a host that does not support 802.1X authentication. If MAC RADIUS authentication is configured as the first authentication method on an interface, then on receiving data from any client on that interface, the device attempts to authenticate the client by using MAC RADIUS authentication. If MAC RADIUS authentication fails, then the device falls back to 802.1X authentication. If 802.1X authentication fails, and captive portal is configured on the interface, the device falls back to captive portal.

802.1X authentication always has the highest priority, even if a client has been authenticated using another method. If the device receives an EAP packet from a client that has been authenticated using MAC RADIUS authentication, the device acknowledges the EAP packet and upgrades the authentication using 802.1X authentication credentials. Similarly, if a client has been authenticated through fallback to captive portal, and the device receives an EAP packet from that client, the device attempts to authenticate the client by using 802.1X authentication.

The device attempts authentication using only methods that are configured on the interface. If an authentication method is included in the authentication order, but is not configured on the interface, the device ignores that method and attempts authentication using the next method in the order that is enabled. However, if a method is enabled on the interface, but is not included in the authentication order, the device does not attempt using that method. For example, if captive portal is enabled for an interface, but the authentication order is configured as `[mac-radius dot1x]`, the authentication method for that interface does not fall back to captive portal.

The authentication order can be configured for all interfaces by using the `interface all` option. If the authentication order is configured for an individual interface, and there is also an authentication order configured for all interfaces, then the order for the individual interface is followed. If there is no authentication order configured for an individual interface, and there is an authentication order configured for all interfaces, then the configuration for all interfaces is followed.

Use the following guidelines when configuring the `authentication-order` statement:

- The authentication order must include at least two methods of authentication.

- 802.1X authentication must be one of the methods included in the authentication order.
- If captive portal is included in the authentication order, it must be the last method in the order.
- If `mac-radius-restrict` is configured on an interface, then the authentication order cannot be configured.

The valid combinations for `authentication-order` are as follows:

- `[dot1x mac-radius captive-portal]`
- `[dot1x captive-portal]`
- `[dot1x mac-radius]`
- `[mac-radius dot1x captive-portal]`

Default

If `authentication-order` is not configured, the device attempts to authenticate the client by using 802.1X authentication first, followed by MAC RADIUS authentication, and then captive portal, as follows:

1. 802.1X authentication—If 802.1X is configured on the interface, the device sends EAPoL requests to the end device and attempts to authenticate the end device through 802.1X authentication. If the end device does not respond to the EAP requests, the device checks whether MAC RADIUS authentication is configured on the interface.
2. MAC RADIUS authentication—If MAC RADIUS authentication is configured on the interface, the device sends the MAC RADIUS address of the end device to the authentication server. If MAC RADIUS authentication is not configured, the device checks whether captive portal is configured on the interface.
3. Captive portal authentication—If captive portal is configured on the interface, the device attempts to authenticate the end device by using this method after attempting any other configured authentication methods.

Options

`captive-portal`—Configure captive portal authentication in the order of authentication methods on the interface.

`dot1x`—Configure 802.1X authentication in the order of authentication methods on the interface.

`mac-radius`—Configure MAC RADIUS authentication in the order of authentication methods on the interface.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

Understanding Authentication on Switches

Configuring Flexible Authentication Order

authentication-order (System)

IN THIS SECTION

- Syntax | 2641
- Hierarchy Level (Junos OS) | 2642
- Hierarchy Level (Junos OS Evolved) | 2642
- Description | 2642
- Default | 2642
- Options | 2642
- Required Privilege Level | 2643
- Release Information | 2643

Syntax

```
authentication-order [method1 method2...];
```

Hierarchy Level (Junos OS)

```
[edit system]
[edit system services ftp]
[edit system services ssh]
[edit system services telnet]
[edit system ports console]
```

Hierarchy Level (Junos OS Evolved)

```
[edit system]
```

Description

Configure the order in which the software tries different user authentication methods when attempting to authenticate a user. For each login attempt, the software tries the authentication methods in order, starting with the first one, until the password matches.

Default

If you do not include the `authentication-order` statement, users are verified based on their configured passwords.

Options

authentication-order [*method1 method2...*]

Specify the order in which the software tries different authentication methods when attempting to authenticate a user.

- **Values:** One or more of the following authentication methods listed in the order in which they must be tried:
  - `password`—Use the password configured for the user with the `authentication` statement at the `[edit system login user]` hierarchy level.
  - `radius`—Use RADIUS authentication services.
  - `tacplus`—Use TACACS+ authentication services.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS User Authentication Methods

authentication-password

IN THIS SECTION

- Syntax | 2643
- Hierarchy Level | 2644
- Description | 2644
- Options | 2644
- Required Privilege Level | 2644
- Release Information | 2644

Syntax

```
authentication-password authentication-password;
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username authentication-md5],
[edit snmp v3 usm local-engine user username authentication-sha],
[edit snmp v3 usm remote-engine engine-id user username authentication-md5],
[edit snmp v3 usm remote-engine engine-id user username authentication-sha]
```

Description

Configure the password for user authentication.

Options

authentication-password—Password that a user enters. The password is then converted into a key that is used for authentication.

SNMPv3 has special requirements when you create plain-text passwords on a router or switch:

- The password must be at least eight characters long.
- The password can include alphabetic, numeric, and special characters, but it cannot include control characters.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring MD5 Authentication

Configuring SHA Authentication

authentication-protocol

IN THIS SECTION

- Syntax | 2645
- Hierarchy Level | 2645
- Description | 2645
- Default | 2646
- Options | 2646
- Required Privilege Level | 2646
- Release Information | 2646

Syntax

```
authentication-protocol {
    eap-md5;
    eap-peap {
        resume;
    }
    pap;
}
```

Hierarchy Level

```
[edit logical-systems protocols dot1x authenticator interface interface-name mac-radius]
[edit protocols dot1x authenticator interface interface-name mac-radius]
```

Description

Specify the protocol to be used by a supplicant to provide authentication credentials for MAC RADIUS authentication. The protocols supported for MAC RADIUS authentication are EAP-MD5, which is the default, Protected Extensible Authentication Protocol (EAP-PEAP), and Password Authentication Protocol (PAP).

Default

If `authentication-protocol` is not configured, the EAP-MD5 authentication protocol is used for MAC RADIUS authentication.

Options

eap-md5 Use the EAP-MD5 protocol for MAC RADIUS authentication. EAP-MD5 is an authentication method belonging to the Extensible Authentication Protocol (EAP) authentication framework. EAP-MD5 uses MD5 to hash the username and password. EAP-MD5 provides for a one-way client authentication. The server sends the client a random request for which the client must provide a response containing an encryption of the request and its password for establishing its identity.

eap-peap <resume> Use the EAP-PEAP protocol, also known as Protected EAP or PEAP, for MAC RADIUS authentication. EAP-PEAP is a protocol that encapsulates EAP within a potentially encrypted and authenticated Transport Layer Security (TLS) tunnel. By encapsulating the authentication process in a TLS tunnel, PEAP addresses the vulnerabilities of an EAP like EAP-MD5.

- **Syntax:** resume—(Optional) Enable faster authentication when reconnecting by resuming the TLS session.

pap Use the PAP authentication protocol for MAC RADIUS authentication. PAP provides a simple password-based authentication for users to establish their identity by using a two-way handshake. PAP transmits plaintext passwords over the network without encryption. PAP must be configured if the Lightweight Directory Access Protocol (LDAP), which supports only plaintext passwords for client authentication, is used for RADIUS authentication.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

`eap-peap` introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Understanding Authentication on Switches

Configuring 802.1X RADIUS Accounting (CLI Procedure)

authentication (Protocols PIM)

IN THIS SECTION

- Syntax | 2647
- Hierarchy Level | 2647
- Description | 2648
- Options | 2648
- Required Privilege Level | 2648
- Release Information | 2648

Syntax

```
authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
}
```

Hierarchy Level

```
[edit protocols pim interface interface-name family (inet | inet6) bfd-liveness-detection],
[edit routing-instances routing-instance-name protocols pim interface family (inet | inet6)
interface-name bfd-liveness-detection]
```

Description

Configure the algorithm, security keychain, and level of authentication for BFD sessions running on PIM interfaces.

The remaining statements are explained separately. See CLI Explorer.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD Authentication for PIM

Configuring BFD for PIM

Understanding Bidirectional Forwarding Detection Authentication for PIM

bfd-liveness-detection (Protocols PIM) | **2945**

key-chain (Protocols PIM) | **9490**

loose-check | **10138**

authentication (Security Alarms)

IN THIS SECTION

- Syntax | **2649**
- Hierarchy Level | **2649**

- Description | 2649
- Default | 2649
- Options | 2649
- Required Privilege Level | 2650
- Release Information | 2650

Syntax

```
authentication failures;
```

Hierarchy Level

```
[edit security alarms potential-violation]
```

Description

Raise a security alarm when the device or switch detects a specified number of authentication failures (bad password attempts) before an alarm is raised.

This value must be equal to or less than the `tries-before-disconnect` setting at the `[edit system login retry-options]` hierarchy level; otherwise, the login session ends before the user reaches the alarmable threshold.

Default

Multiple authentication failures do not cause an alarm to be raised.

Options

failures—Number of authentication failures that causes an alarm to be raised.

- **Range:** 1 through 10.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Security Policies Overview

authentication (Security IPsec)

IN THIS SECTION

- Syntax | 2650
- Hierarchy Level | 2651
- Description | 2651
- Options | 2651
- Required Privilege Level | 2651
- Release Information | 2651

Syntax

```
authentication {
    algorithm (hmac-sha1-96 | hmac-sha2-256);
    key (ascii-text key | hexadecimal key);
}
```

Hierarchy Level

```
[edit security ipsec security-association sa-name manual direction (inbound | outbound | bi-
directional)]
```

Description

Configure IP Security (IPsec) authentication parameters for manual security association (SA).

> **NOTE**: The hexadecimal format must be used for the encryption key to be FIPS compliant. The hexadecimal keys provide maximum key strength.

Options

`algorithm`—Hash algorithm that authenticates packet data. It can be one of the following:

- `hmac-md5-96`—Produces a 128-bit digest.
- `hmac-sha1-96`—Produces a 160-bit digest.

`key`—Type of authentication key. It can be one of the following:

- `ascii-text` *key*—ASCII text key. For `hmac-md5-96`, the key is 16 ASCII characters; for `hmac-sha1-96`, the key is 20 ASCII characters.
- `hexadecimal` *key*—Hexadecimal key. For `hmac-md5-96`, the key is 32 hexadecimal characters; for `hmac-sha1-96`, the key is 40 hexadecimal characters.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Manual IPsec Security Associations for an ES PIC

authentication-server

IN THIS SECTION

- Syntax | 2652
- Hierarchy Level | 2652
- Description | 2652
- Options | 2652
- Required Privilege Level | 2653
- Release Information | 2653

Syntax

```
authentication-server [server-addresses];
```

Hierarchy Level

```
[edit access profile profile-name radius]
```

Description

Configure the RADIUS server for authentication. To configure multiple RADIUS servers, include multiple server addresses. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options

server-addresses—Configure one or more RADIUS server addresses.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

RELATED DOCUMENTATION

- Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch
- show network-access aaa statistics authentication

authentication-sha

IN THIS SECTION

- Syntax | 2653
- Hierarchy Level | 2654
- Description | 2654
- Required Privilege Level | 2654
- Release Information | 2654

Syntax

```
authentication-sha {
    (authentication-key authentication-key | authentication-password authentication-password);
}
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]
```

Description

Configure the secure hash algorithm (SHA) as the authentication type for the SNMPv3 user.

NOTE: You can configure only one authentication type for each SNMPv3 user.

For authentication, you can either enter a password in clear text, which is immediately encrypted, or, if you already have the encrypted version (known as a *key*), enter the key directly.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring SHA Authentication

authentication-sha224

IN THIS SECTION

Syntax | 2655

- Hierarchy Level | 2655
- Description | 2655
- Required Privilege Level | 2655
- Release Information | 2655

Syntax

```
authentication-sha224 {
            authentication-key authentication-key;
            authentication-password authentication-password;
         }
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]
```

Description

Configure the secure hash algorithm (SHA)224 as the authentication type for the SNMPv3 user. You can use this configuration statement for local and remote user.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

Configuring SHA Authentication

authentication-sha256

IN THIS SECTION

- Syntax | 2656
- Hierarchy Level | 2656
- Description | 2656
- Required Privilege Level | 2657
- Release Information | 2657

Syntax

```
authentication-sha256 {
            authentication-key authentication-key;
            authentication-password authentication-password;
         }
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]
```

Description

Configure the secure hash algorithm (SHA)256 as the authentication type for the SNMPv3 user. You can use this configuration statement for local and remote user.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

Configuring SHA Authentication

authentication-sha348

IN THIS SECTION

- Syntax | 2657
- Hierarchy Level | 2658
- Description | 2658
- Required Privilege Level | 2658
- Release Information | 2658

Syntax

```
authentication-sha348 {
            authentication-key authentication-key;
            authentication-password authentication-password;
         }
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]
```

Description

Configure the secure hash algorithm (SHA)348 as the authentication type for the SNMPv3 user. You can use this configuration statement for local and remote user.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Configuring SHA Authentication

authentication-sha512

IN THIS SECTION

- Syntax | 2659
- Hierarchy Level | 2659
- Description | 2659
- Required Privilege Level | 2659
- Release Information | 2659

Syntax

```
authentication-sha512 {
            authentication-key authentication-key;
            authentication-password authentication-password;
         }
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]
```

Description

Configure the secure hash algorithm (SHA)512 as the authentication type for the SNMPv3 user. You can use this configuration statement for local and remote user.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Configuring SHA Authentication

authentication-source (Services User Identification ClearPass)

IN THIS SECTION

- Syntax | 2660
- Hierarchy Level | 2661
- Description | 2661
- Options | 2661
- Required Privilege Level | 2661
- Release Information | 2662

Syntax

```
authentication-source aruba-clearpass {
    authentication-entry-timeout (Services User Identification) minutes;
    invalid-authentication-entry-timeoutminutes;
    no-user-query (Services User Identification);
    traceoptions (Services User Identification) {
        file filename files files match match size size(world-readable | no-world-readable);
        flag name;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
    user-query (Services User Identification) {
        ca-certificate ca-certificate;
        client-id client-id;
        client-secret client-secret;
        delay-query-time seconds;
        query-api query-api;
        token-api token-api;
        web-server {
            server-name;
            address address;
            connect-method (http | https);
            port port;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit services user-identification]
```

Description

Configure ClearPass as the authentication source for the integrated ClearPass authentication and enforcement feature.

The ClearPass Policy Manager (CPPM), as the authentication source and client of the SRX Series Firewall HTTP server, initiates a connection to the SRX Series Firewall using the Web API that the SRX Series Firewall exposes to it. The CPPM sends user authentication and identity information to the SRX Series Firewall across this connection using HTTP or HTTPS POST request messages.

`set authentication-source aruba-clearpass` command can be used to configure the Juniper Identity Management Service as the authentication-source.

Options

name	Aruba ClearPass authentication source name.
authentication-entry-timeout	Aruba ClearPass authentication entry timeout number. • **Range:** 10 through 1440 minutes
invalid-authentication-entry-timeout	Invalid authentication entry timeout number. • **Range:** 10 through 1440 minutes
no-user-query	Disable user query from ClearPass.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

1. services—To view this statement in the configuration.
2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

authentication-source (Services User Identification Device Identity)

IN THIS SECTION

- Syntax | 2662
- Hierarchy Level | 2662
- Description | 2662
- Options | 2663
- Required Privilege Level | 2663
- Release Information | 2663

Syntax

```
authentication-source (active-directory | network-access-controller)
```

Hierarchy Level

```
[edit services user-identification device-information]
```

Description

Specify the device identity authentication source. The integrated user firewall device identity authentication feature enables you to control access to resources based on the identity of the device and not that of the user of the device. Supported authentication sources include Active Directory and third-party network access systems.

The SRX Series Firewall obtains the device identity information for authenticated devices from the authentication source. After the SRX Series Firewall obtains the device information, it creates a device identity authentication table to use to store device identity entries.

The SRX Series Firewall searches the device identity authentication table for a device match when traffic issuing from a user's device arrives at the SRX Series Firewall. If it finds a match, the SRX Series Firewall searches for a matching security policy. If it finds a matching security policy, the security policy's action is applied to the traffic.

Options

active-directory

Specifies Microsoft Active Directory as the authentication source.

The SRX Series Firewall obtains the device identity information for authenticated devices from Active Directory. It reads the Active Directory domain controller event logs to obtain the IP addresses of devices logged into the domain and authenticated by Windows. Then, for each authenticated device, it obtains from the Active Directory LDAP server the names of the groups to which the device belongs, based on the IP addresses of the devices.

network-access-controller

Specifies the authentication source as that of a third-party network access controller (NAC) system. If your network environment is configured for a NAC solution and you decide to take this approach, the NAC system sends the device identity information to the SRX Series Firewall. The SRX Series Firewall exposes a RESTful Web services API implementation that enables you to send the device identity information to the SRX Series Firewall in a formal XML structure. If you take this approach, you must verify that your NAC solution works with the SRX Series Firewall.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

Understanding Access Control to Network Resources Based on Device Identity Information

Understanding the Device Identity Attributes and Profiles for the Integrated User Firewall Device Identity Authentication Feature

Understanding the Device Identity Authentication Table and Its Entries

authentication (Static and Dynamic PPP)

IN THIS SECTION

- Syntax | 2664
- Hierarchy Level | 2664
- Description | 2664
- Options | 2664
- Required Privilege Level | 2665
- Release Information | 2665

Syntax

```
authentication [ authentication-protocols ];
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" ppp-options],
[edit interfaces pp0 unit unit-number ppp-options]
```

Description

Specify the order in which the router tries to negotiate PPP authentication protocols when verifying that a PPP client can access the network. By default, the router tries to negotiate Challenge Handshake Authentication Protocol (CHAP) authentication first, and then tries Password Authentication Protocol (PAP) authentication if the attempt to negotiate CHAP authentication is unsuccessful.

You can specify one or both authentication protocols. If you specify both CHAP and PAP in either order, you must enclose the set of protocol names within square brackets ([ ]).

Options

authentication-protocols One or both of the following PPP authentication protocols:

- chap—Challenge Handshake Authentication Protocol
- pap—Password Authentication Protocol

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

PPP Network Control Protocol Negotiation

authentication (Static Subscribers)

IN THIS SECTION

- Syntax | 2665
- Hierarchy Level | 2666
- Description | 2666
- Required Privilege Level | 2666
- Release Information | 2667

Syntax

```
authentication {
    password password-string;
    username-include {
```

```
        delimiter delimiter-character;
        domain-name domain-name;
        interface;
        logical-system-name;
        routing-instance-name;
        user-prefix user-prefix-string;
        vlan-tags;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name],
[edit logical-systems logical-system-name system services static-subscribers],
[edit logical-systems logical-system-name system services static-subscribers group group-name],
[edit routing-instances routing-instances-name system services static-subscribers],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name],
[edit system services static-subscribers],
[edit system services static-subscribers group group-name]
```

Description

Specify the authentication parameters that trigger the Access-Request message to AAA for all static subscribers on interfaces configured at the `[edit system services static-subscribers interface]` hierarchy level, or for the static subscribers in a specific group. The group version of this statement overrides the global configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Configuring the Static Subscriber Global Authentication Password

Configuring the Static Subscriber Group Authentication Password

authentication-type

IN THIS SECTION

- Syntax | 2667
- Hierarchy Level | 2667
- Description | 2668
- Options | 2668
- Required Privilege Level | 2668
- Release Information | 2668

Syntax

```
authentication-type (md5 | simple);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]
```

Description

Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group. If you enable authentication, you must specify a password by including the `authentication-key` statement. The specific type of authentication used by OSPF is encoded in this field.

All devices in the VRRP group must use the same authentication scheme and password.

> **NOTE**: When VRRPv3 is enabled, the `authentication-type` and `authentication-key` statements cannot be configured for any VRRP groups.

Options

authentication Authentication scheme:

- **simple** Use a simple password. The password is included in the transmitted packet, so this method of authentication is relatively insecure.
- **md5** Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing platform uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.
- **Default:** none (no authentication is performed).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring VRRP Authentication (IPv4 Only)

authentication-key | 2603

version-3

Understanding VRRP on SRX Series Devices

Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces

authentication-type (Protocols RIP)

IN THIS SECTION

- Syntax | 2669
- Hierarchy Level | 2669
- Description | 2670
- Default | 2670
- Options | 2670
- Required Privilege Level | 2670
- Release Information | 2670

Syntax

```
authentication-type type;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip],
[edit protocols rip group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

Configure the type of authentication for RIP route queries received on an interface.

Default

If you do not include this statement and the `authentication-key` statement, RIP authentication is disabled.

Options

type—Authentication type:

- `md5`—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing device uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.
- `none`—Disable authentication. If `none` is configured, the configured authentication key is ignored.
- `simple`—Use a simple password. The password is included in the transmitted packet, which makes this method of authentication relatively insecure. The password can be from 1 through 16 contiguous letters or digits long.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring Route Authentication for RIP using single MD5 key

authentication-key (Protocols RIP) | **2619**

authentication-whitelist

IN THIS SECTION

- Syntax | 2671
- Hierarchy Level | 2671
- Description | 2671
- Options | 2672
- Required Privilege Level | 2672
- Release Information | 2672

Syntax

```
authentication-whitelist {
    mac-address {
    bridge-domain-assignment bridge-domain-assignment;
    interface interface-name;
    vlan-assignment  ( vlan-id |vlan-name);
}
```

Hierarchy Level

```
[edit ethernet-switching-options],
[edit logical-systems name switch-options],
[edit switch-options]
```

Description

Configure MAC addresses to exclude from RADIUS authentication. The authentication allowlist provides an authentication bypass mechanism for supplicants connecting to a port, permitting devices, such as printers, to be connected to the network without going through the authentication process.

Options

mac-address	The MAC address of the device for which RADIUS authentication should be bypassed and the device permitted access to the port.
bridge-domain-assignment *bridge-domain-assignment*	(MX Series only) Specify the bridge-domain name or 802.1q tag identifier for the MAC address that should be allowed to bypass RADIUS authentication.
interface [*interface-names***]**	Specify a list of interfaces on which the specified MAC addresses are allowed to bypass RADIUS authentication and allowed to connect to the LAN without authentication.
vlan-assignment (*vlan-id* **|** *vlan-name*	(EX, QFX, and SRX Series only) Specify the VLAN 802.1q tag identifier or name associated with the list of MAC addresses that should be allowed to bypass RADIUS authentication.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

The `[edit switch-options]` hierarchy level was introduced in Junos OS Release 13.2X50-D10 for EX Series switches (ELS).

RELATED DOCUMENTATION

Example: Setting Up Captive Portal Authentication on an EX Series Switch

Example: Setting Up Captive Portal Authentication on an EX Series Switch with ELS Support

Configuring Captive Portal Authentication (CLI Procedure)

Configuring Captive Portal Authentication (CLI Procedure) on an EX Series Switche with ELS Support

authenticator

IN THIS SECTION

- Syntax | 2673
- Hierarchy Level | 2675
- Description | 2675
- Default | 2675
- Options | 2675
- Required Privilege Level | 2676
- Release Information | 2676

Syntax

```
authenticator {
    authentication-profile-name access-profile-name;
    interface (all | [ interface-names ]) {
        authentication-order (captive-portal | dot1x | mac-radius);
        disable;
        guest-bridge-domain guest-bridge-domain;
        guest-vlan guest-vlan;
        ignore-port-bounce;
        mac-radius {
            authentication-protocol {
                eap-md5;
                eap-peap {
                    resume;
                }
                pap;
            }
            flap-on-disconnect;
            restrict;
        }
        maximum-requests number;
        multi-domain {
            max-data-session max-data-session;
```

```
            packet-action (drop-and-log | shutdown);
            recovery-timeout seconds;
        }
        (no-reauthentication | reauthentication interval );
        no-tagged-mac-authentication;
        quiet-period seconds;
        redirect-url  redirect-url;
        retries (802.1X) number;
        server-fail (bridge-domain bridge-domain | deny | permit | use-cache | vlan-name vlan-
name);
        server-fail-voip (deny | permit | use-cache | vlan-name vlan-name);
        server-reject-bridge-domain bridge-domain {
            block-interval seconds;
            eapol-block;
        }
        server-reject-vlan (vlan-id | vlan-name)  {
            block-interval block-interval;
            eapol-block;
        }
        server-timeout seconds;
        supplicant (single | single-secure | multiple);
        supplicant-timeout seconds;
        transmit-period seconds;
    }
    ip-mac-session-binding;
    no-mac-table-binding;
    radius-options {
        add-interface-text-description;
        use-vlan-id;
        use-vlan-name;
    }
    static mac-address {
        bridge-domain-assignment bridge-domain-assignment;
        interface interface;
        vlan-assignment vlan-identifier;
    }
}
```

Hierarchy Level

```
[edit logical-systems name protocols              dot1x],
[edit protocols dot1x]
```

Description

Specify the group of servers to be used for IEEE 802.1X or MAC RADIUS authentication for Port-Based Network Access Control, configure interfaces for 802.1x authentication, and configure static MAC bypass for 802.1x and MAC RADIUS authentication. 802.1X authentication is supported on interfaces that are members of private VLANs (PVLANs).

> **NOTE**: You cannot configure 802.1X user authentication on interfaces that have been enabled for Q-in-Q tunneling.

Default

802.1X authentication is disabled.

Options

authentication-profile-name *access-profile-name*

Specify the name of the access profile to be used for 802.1X or MAC RADIUS user authentication. The access profile is configured at the [edit access profile] hierarchy level and contains the RADIUS server IP address and other information used for authentication.

> **NOTE**: Access profile configuration is required only for 802.1X clients, not for static MAC clients.

- **Default:** No access profile is specified.

ip-mac-session-binding

Configure the switching device to check for an IP-MAC address binding in the DHCP, DHCPv6, or SLAAC snooping table before terminating the authentication session when the MAC address ages out. If the MAC address for the end device is bound to an IP address, then it will be retained in the Ethernet switching table, and the authentication session will remain active.

To configure this feature, you must also disassociate the authentication session table from the Ethernet switching table using the `no-mac-table-binding` statement. This extends the authentication session until the next re-authentication period.

> **NOTE**: This feature requires DHCP, DHCPv6, or SLAAC snooping to be enabled on the device.

- **Default:** Not enabled

no-mac-table-binding

Specify that the device not remove the session from the authentication session table when the MAC address ages out of the Ethernet switching table.

- **Default:** Not enabled

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`no-mac-table-binding` introduced in Junos OS Release 11.1.

`radius-options` introduced in Junos OS Release 12.1.

`add-interface-text-description` introduced in Junos OS Release 18.4.

`ip-mac-session-binding` introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch

Example: Configuring MAC RADIUS Authentication on an EX Series Switch

Configuring 802.1X Interface Settings (CLI Procedure)

Specifying RADIUS Server Connections on Switches (CLI Procedure)

Example: Configuring Static MAC Bypass of 802.1X and MAC RADIUS Authentication on an EX Series Switch

Access Control and Authentication on Switching Devices

Retaining the Authentication Session Based on IP-MAC Address Bindings

no-mac-table-binding (802.1X for MX Series in Enhanced LAN Mode)

authorization (SNMP)

IN THIS SECTION

- Syntax | 2677
- Hierarchy Level | 2677
- Description | 2677
- Options | 2678
- Required Privilege Level | 2678
- Release Information | 2678

Syntax

```
authorization authorization;
```

Hierarchy Level

```
[edit snmp community community-name]
```

Description

Set the access authorization for SNMP **Get**, **GetBulk**, **GetNext**, and **Set** requests.

Options

authorization—Access authorization level:

- **read-only**—Enable **Get**, **GetNext**, and **GetBulk** requests.
- **read-write**—Enable all requests, including **Set** requests. You must configure a view to enable **Set** requests.

> **NOTE**: The **read-write** option is not supported on the QFX3000 QFabric system.

- **Default: read-only**

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring SNMP Communities

authorization (icap-redirect profile)

IN THIS SECTION

- Syntax | 2679
- Hierarchy Level | 2679
- Description | 2679
- Options | 2679
- Required Privilege Level | 2679

Release Information | 2680

Syntax

```
authorization {
    authorization-type authorization-type;
    credentials (ascii ascii | base64 base64);
}
```

Hierarchy Level

```
[edit services icap-redirect profile name server name]
```

Description

User authentication for the ICAP server if the request needs to be authorized.

Options

authorization-type Authentication type for the ICAP server. Authorization type is basic by default.

credentials Credentials (user name and password) for authentication to ICAP server.

- Values:

  ascii *ascii* ASCII string.

  base64 *base64* bBase64 encoded string.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Example: Configuring ICAP Redirect Service on SRX Series Firewalls

authorization-order

IN THIS SECTION

- Syntax | 2680
- Hierarchy Level | 2680
- Description | 2680
- Options | 2681
- Required Privilege Level | 2681
- Release Information | 2681

Syntax

```
authorization-order [ authorization-methods ];
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Set the order in which AAA tries different methods to verify that a client is authorized access to the router or switch. For each login attempt, AAA tries the authorization methods in order, from first to last.

A given subscriber does not undergo both authentication and authorization as separate steps. When both `authentication-order` and `authorization-order` are specified, DHCP subscribers honor the configured authorization order, all other subscribers use the configured authentication-order.

Options

authorization-methods Ordered list of methods to use for authorization attempts. The list includes one or more of the following methods in any combination:

- `jsrc`—Use the JSRC application in an SRC environment to request authorization from the SAE when verifying that a subscriber can access the router or switch.
  - When you configure both this option and `authentication-order`, AAA ignores the authentication order setting for DHCP subscribers. For non-DHCP subscribers, AAA ignores the authorization order and applies only the authentication order.
  - When you configure only this option, AAA applies the authorization order to both DHCP and non-DHCP subscribers.
- `nasreq`—Use the NASREQ application to communicate with a NASREQ server for authorization of any subscriber type as an alternative to RADIUS authorization.
- `none`—No authorization is performed. Can be used, for example, when the Diameter function Gx-Plus is employed for notification during subscriber provisioning.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

`nasreq` option added in Junos OS Release 16.1.

RELATED DOCUMENTATION

JSRC Configuration Overview

Authorizing Subscribers with JSRC

auto-bandwidth

IN THIS SECTION

- Syntax | 2682
- Hierarchy Level | 2682
- Description | 2683
- Options | 2683
- Required Privilege Level | 2684
- Release Information | 2684

Syntax

```
auto-bandwidth {
    template name {
        adjust-interval adjust-interval;
        adjust-threshold adjust-threshold;
        auto-bandwidth-subscription auto-bandwidth-subscription;
        statistic-collection-interval statistic-collection-interval;
    }
    traceoptions {
        file filename <files files>     <size size> <(world-readable | no-world-readable)>;
        flag (all | state | timer) {
            detail detail;
            disable disable;
            receive receive;
            send send;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options],
[edit logical-systems name routing-options],
```

```
[edit routing-instances name routing-options],
[edit routing-options]
```

Description

Configure an auto-bandwidth template to define parameters that measure the available bandwidth and periodically trigger updates to RSVP and IGP. You can configure this option both at the IS-IS protocol level or at the IS-IS interface level.

Traffic samples are collected periodically at each statistics interval and the average SPRING bandwidth utilization is calculated. If this average value increases beyond the configured adjust threshold percentage, that is the last reported SPRING bandwidth, then RSVP adjusts the subscription percentage accordingly. which might cause RSVP LSP preemption, IGP updates, and further re-routing of RSVP LSPs. However, if the SPRING bandwidth utilization drops below the configured threshold, the RSVP subscription percentage is modified to reflect the increase. You can configure parameters such as the threshold, the collection interval, and the adjust interval as per your requirements.

CAUTION: Configure the `adjust-interval` and `statistics-interval` options appropriately to avoid scaling issues. You can fine tune these values based on route and label-switched path (LSP) scale.

Options

template *name*

Specify an auto-bandwidth template name with a maximum length of 64 characters.

adjust-interval

Specify the time interval after which the average bandwidth utilization is measured.

- **Range:** `30` through `3600` seconds
- **Default:** 900 seconds

adjust-threshold

Specify the percentage change threshold after which an action is triggered.

- **Range:** `0` through `100` percent
- **Default:** 10 percent

auto-bandwidth-subscription

Specify the threshold for RSVP subscription in percentage.

- **Range:** 0 through 200 percent
- **Default:** 100 percent

statistic-collection-interval Specify the time interval after which the traffic statistics is collected from the line cards.

- **Range:** 10 through 300 seconds
- **Default:** 60 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

traffic-statistics (Protocols IS-IS) | 30048

update-threshold-max-reservable | 30458

auto-bandwidth (MPLS Statistics)

IN THIS SECTION

- Syntax | 2685
- Hierarchy Level | 2685
- Description | 2685
- Required Privilege Level | 2685
- Release Information | 2685

Syntax

```
auto-bandwidth;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls statistics],
[edit protocols mpls statistics]
```

Description

Collect statistics related to automatic bandwidth.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Automatic Bandwidth Allocation for LSPs

Configuring MPLS to Gather Statistics

statistics (Protocols MPLS) | **28594**

auto-bandwidth (MPLS Tunnel)

IN THIS SECTION

- Syntax | 2686
- Hierarchy Level | 2686
- Description | 2687
- Options | 2687
- Required Privilege Level | 2687
- Release Information | 2687

Syntax

```
auto-bandwidth {
    adjust-interval seconds;
    adjust-threshold percent;
    adjust-threshold-absolute;
    adjust-threshold-activate-bandwidthbps
    adjust-threshold-overflow-limit number;
    adjust-threshold-underflow-limit number;
    maximum-bandwidth bps;
    minimum-bandwidth bps;
    minimum-bandwidth-adjust-interval
    minimum-bandwidth-adjust-threshold-change
    minimum-bandwidth-adjust-threshold-value
    monitor-bandwidth;
    resignal-minimum-bandwidth
    sync-active-path-bandwidth
}
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name]
```

Description

Allow an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel.

> **NOTE**: In calculating the value for `Max AvgBW` (relative to the ingress LSP), the sample collected during make before break (MBB) is ignored to prevent inaccurate results. The first sample after a bandwidth adjustment, or after a change in the LSP ID (regardless of path change), is also ignored.

Options

adjust-threshold-absolute ***adjust-threshold-absolute-value***	Configure an absolute value based threshold along with the percentage based threshold that helps avoid the bandwidth getting triggered for LSPs of both small and large bandwidth reservations.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Automatic Bandwidth Allocation for LSPs

request mpls lsp adjust-autobandwidth | **14691**

show mpls lsp autobandwidth | **22206**

auto-configuration

IN THIS SECTION

- Syntax | 2688
- Hierarchy Level | 2688
- Description | 2688
- Options | 2688
- Required Privilege Level | 2689
- Release Information | 2689

Syntax

```
auto-configuration {
    command binary-file-path;
    disable;
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Configure the autoconfiguration process.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the autoconfiguration process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Autoinstallation Overview

Configuring Autoinstallation of Configuration Files

auto-configuration (System)

IN THIS SECTION

- Syntax | 2689
- Hierarchy Level | 2690
- Description | 2690
- Options | 2690
- Required Privilege Level | 2691
- Release Information | 2691

Syntax

```
auto-configuration {
    traceoptions {
        file {
        filename;
            files number;
            match regular-expression;
```

```
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the autoconfiguration process.

Options

`traceoptions`—Set the trace options.

- `file`—Configure the trace file information.
  - `filename`—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *`number`*—Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed to `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    `Range:` 2 through 1000 files

    `Default:` 10 files
  - `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
  - `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named `trace-file` reaches this size, it is renamed `trace-file.0`. When `trace-file` again reaches its maximum size, `trace-file.0` is renamed `trace-file.1` and `trace-file`

is renamed `trace-file.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

`Syntax:` x K to specify KB, x m to specify MB, or x g to specify GB

`Range:` 10 KB through 1 GB

`Default:` 128 KB

- `world-readable | no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Specify the tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags.
  - `all`—Trace all events.
  - `auth`—Trace VLAN authentication.
  - `configuration`—Trace configurations.
  - `interfaces`—Trace interface operations.
  - `io`—Trace I/O operations.
  - `rtsock`—Trace routing socket operations.
  - `ui`—Trace user interface operations.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Autoinstallation Overview

auto-configure

IN THIS SECTION

- Syntax | 2692
- Hierarchy Level | 2693
- Description | 2694
- Required Privilege Level | 2694
- Release Information | 2694

Syntax

```
auto-configure {
    vlan-ranges {
        access-profile profile-name;
        authentication {
            packet-types  [packet-types];
            password password-string;
            username-include{
                circuit-id;
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-18;
                option-37;
                option-82 <circuit-id> <remote-id>;
                radius-realm radius-realm-string;
                remote-id;
                user-prefix user-prefix-string;
                vlan-tags;
            }
```

```
        }
        dynamic-profile profile-name {
            accept  (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
            accept-out-of-band protocol;
            ranges (any | low-tag)-(any | high-tag);
        }
        override;
    }
    stacked-vlan-ranges {
        access-profile profile-name;
        authentication {
            packet-types  [packet-types];
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-18;
                option-37;
                option-82 <circuit-id> <remote-id>;
                radius-realm radius-realm-string;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        dynamic-profile profile-name {
            accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
            ranges (any | low-tag-high-tag),(any | low-tag-high-tag);
        }
        override;
    }
    remove-when-no-subscribers;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Enable the configuration of dynamic, auto-sensed VLANs.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

auto-configure (Demux)

IN THIS SECTION

- Syntax | 2695
- Hierarchy Level | 2695
- Description | 2695
- Required Privilege Level | 2696
- Release Information | 2696

Syntax

```
auto-configure {
    address-ranges {
        authentication {
            password password-string;
            username-include {
                auth-server-realm realm-string;
                delimiter delimiter-character;
                domain-name domain-name;
                interface-name;
                source-address;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            network ip-address {
                range name {
                    low lower-limit;
                    high upper-limit;
                }
            }
        }
    }

}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet]
[edit interfaces interface-name unit unit-number demux inet6]
```

Description

Enable the configuration of dynamic, auto-sensed subscriber interfaces for the demultiplexing (demux) interface options. The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

> *IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview*
>
> *Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles*

auto-configure (Dynamic VLAN Interface Sets)

IN THIS SECTION

- Syntax | 2696
- Hierarchy Level | 2697
- Description | 2697
- Required Privilege Level | 2697
- Release Information | 2697

Syntax

```
auto-configure {
    agent-circuit-identifier {
        dynamic-profile profile-name;
    }
    line-identity {
        include {
            accept-no-ids;
```

```
            circuit-id;
            remote-id;
        }
        dynamic-profile profile-name;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit"],
[edit interfaces interface-name unit logical-unit-number]
```

Description

Enable the configuration of dynamic, auto-sensed VLAN subscriber interfaces on a static or dynamic underlying VLAN interface. Use the `agent-circuit-identifier` statement to configure dynamic VLANs based only on the ACI. Use the `line-identity` statement to configure dynamic VLANs that can be initiated by receipt of ACI, ARI, both ACI and ARI, or neither.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information

Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information

Configuring Dynamic Underlying VLAN Interfaces to Use Access-Line Identifiers

Configuring Dynamic VLAN Subscriber Interfaces Based on Access-Line Identifiers

auto-configure (IPv4)

IN THIS SECTION

- Syntax | 2698
- Hierarchy Level | 2699
- Description | 2699
- Options | 2699
- Required Privilege Level | 2699
- Release Information | 2700

Syntax

```
auto-configure {
                    address-ranges {
                        session-timeout seconds;
                        dynamic-profile dynamic-profile {
                            network network-address;
                        }
                        authentication {
                            username-include {
                                delimiter;
                                domain-name;
                                user-prefix;
                                auth-server-realm;
                                interface-name;
                                source-address;
                            }
                        }
                    }
                }
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces junos-interface-ifd-name unit junos-interface-unit family
inet
```

Description

Defines the parameters for creating dynamic VLAN (DVLAN) interface on receiving the first VLAN packet from the client.

Options

dynamic-profile *dynamic-profile*
Name of the dynamic profile.

session-timeout *seconds*
Duration of the active dynamic IP subscriber in seconds. If you do not configure the `session-timeout` value, the system does not delete the dynamic IP subscriber unless an active DHCP subscriber is available.

- Range: 600 thru 6000

network *network-address*
IPv4 network address.

username-include
As a part of subscriber authentication, you can define the username with the following strings:

- `delimiter`-Delimiter or separator character
- `domain-name`-Domain name
- `user-prefix`-User defined prefix
- `auth-server-realm`-Authentication server realm name
- `interface-name`-Interface name
- `source-address`-Source address

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery

auto-configure (IPv6) | 2700

auto-configure (IPv6)

IN THIS SECTION

- Syntax | 2700
- Hierarchy Level | 2701
- Description | 2701
- Options | 2701
- Required Privilege Level | 2702
- Release Information | 2702

Syntax

```
auto-configure {
                    address-ranges {
                        session-timeout seconds;
                        dynamic-profile dynamic-profile {
                            prefix prefix;
                        }
                        authentication {
                            username-include {
                                delimiter;
                                domain-name;
                                user-prefix;
                                auth-server-realm;
```

```
                                interface-name;
                                source-address;
                            }
                        }
                    }
                }
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces junos-interface-ifd-name unit junos-interface-unit family
inet6
```

Description

Defines the parameters for creating dynamic VLAN (DVLAN) interface on receiving the first VLAN packet from the client.

Options

dynamic-profile *dynamic-profile*
Name of the dynamic profile.

session-timeout *seconds*
Duration of the active dynamic IP subscriber in seconds. If you do not configure the `session-timeout` value, the system does not delete the dynamic IP subscriber unless an active DHCP subscriber is available.

- Range: 600 thru 6000

prefix *prefix*
IPv6 address prefix.

username-include
As a part of subscriber authentication, you can define the username with the following strings:

- `delimiter`-Delimiter or separator character
- `domain-name`-Domain name
- `user-prefix`-User defined prefix
- `auth-server-realm`-Authentication server realm name
- `interface-name`-Interface name
- `source-address`-Source address

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery

auto-configure (IPv4) | 2698

auto-configure-trigger interface (ANCP)

IN THIS SECTION

- Syntax | 2702
- Hierarchy Level | 2703
- Description | 2703
- Options | 2703
- Required Privilege Level | 2703
- Release Information | 2703

Syntax

```
auto-configure-trigger interface interface-name;
```

Hierarchy Level

```
[edit protocols ancp neighbor ip-address]
```

Description

Map an ANCP neighbor to a subscriber-facing physical interface on the router, so that ANCP Port Up and Port Down messages trigger notifications to the auto-configuration daemon (autoconfd) to initiate VLAN creation (Port Up) or removal (Port Down).

Options

interface-name Name of the physical interface. Can be any of the following interface types: ae, ge, xe, et, demux, ps.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

`ps` interface type support added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Configuring ANCP Neighbors

Configuring the ANCP Agent

auto-conversion

IN THIS SECTION

- Syntax | 2704
- Hierarchy Level | 2704
- Description | 2704
- Required Privilege Level | 2705
- Release Information | 2705

Syntax

```
auto-conversion;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Enable automatic Virtual Chassis port (VCP) conversion (also called VCP autoprovisioning) in a Virtual Chassis.

VCP links connect members in a Virtual Chassis using supported VCP ports on the member switches. Ports that are supported as VCPs and are not dedicated VCPs or configured as VCPs by default must be converted into VCPs to use them to interconnect Virtual Chassis members. You can manually convert a supported port into a VCP port using the `request virtual-chassis vc-port` command, or when automatic VCP conversion is enabled, a port will be automatically converted into a VCP under the following conditions:

- LLDP is enabled on the interfaces for the members on both sides of the link. The two interfaces exchange LLDP packets to accomplish the port conversion.
- The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the `set virtual-chassis member` command.

- The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

You might want to invoke automatic VCP conversion when adding a new switch to an existing Virtual Chassis or adding a redundant VCP link between two existing members of a Virtual Chassis.

> **NOTE**: When automatic VCP conversion is enabled in a Virtual Chassis with switches that have dedicated VCPs (EX4200, EX4500, or EX4550 Virtual Chassis), if network or uplink ports are automatically converted into VCPs to create a redundant link with a dedicated VCP connection, you must reboot the Virtual Chassis to avoid creating a traffic loop within the Virtual Chassis. See *Troubleshooting an EX Series Virtual Chassis* for more information on this potential issue.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 14.1X53-D47 and 15.1R7.

RELATED DOCUMENTATION

Virtual Chassis Overview for Switches

Understanding Virtual Chassis Components

Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port

auto-dad (SLAAC Snooping)

IN THIS SECTION

- Syntax | 2706
- Hierarchy Level | 2706
- Description | 2706
- Options | 2706

Required Privilege Level | 2707
Release Information | 2707

Syntax

```
auto-dad {
    retries retry-count;
    retrans-interval seconds;
}
```

Hierarchy Level

```
[edit forwarding-options access-security slaac-snooping interface  (interface-name | all)]
```

Description

Configure parameters for Duplicate Address Detection (DAD) to be performed by the SLAAC snooping device. DAD is used by IPv6 clients to verify the uniqueness of addresses obtained through stateless address auto-configuration (SLAAC). DAD sends a Neighbor Solicitation message destined to the new address. If there is a reply, then the address is a duplicate; if the address is unique, it is assigned to the interface.

If DAD is disabled on the client side, or if DAD packets are dropped due to traffic congestion, the SLAAC snooping device can perform auto-DAD on the client's behalf. You can configure the number of times that DAD will attempt to verify the address and the length of time that DAD waits for a response before retransmission. The client address binding remains in a hold state until all transmissions are completed.

Options

retries ***retry-count***	Configure the number of times that auto-DAD sends a Neighbor Solicitation message to verify the uniqueness of an address obtained through SLAAC.
retrans-interval ***seconds***	Configure the interval between retransmissions of a Neighbor Solicitation message for auto-DAD.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

IPv6 Stateless Address Auto-configuration (SLAAC) Snooping

auto-derive

IN THIS SECTION

- Syntax | 2707
- Hierarchy Level | 2708
- Description | 2708
- Options | 2708
- Required Privilege Level | 2709
- Release Information | 2709

Syntax

```
auto-derive {
    lacp-pe-system-id-and-admin-key;
    type-1-lacp;
    type-3-system-mac {
        local-discriminator number;
        mac mac-address;
```

```
    }
}
```

Hierarchy Level

```
[edit interfaces aeX esi]
[edit interfaces aeX unit logical-unit-number esi]
```

Description

Configure aggregated Ethernet interfaces and aggregated Ethernet logical interfaces to automatically derive Ethernet segment identifiers (ESIs) from the Link Aggregation Control Protocol (LACP) configuration.

We support the `lacp-pe-system-id-and-admin-key` option in the following environments:

- On Juniper Networks devices that support the automatic ESI feature and are multihomed in active-active mode in an EVPN-VXLAN overlay network.
- On Juniper Networks devices that support the automatic ESI feature and are multihomed in active-standby or active-active mode in an EVPN-MPLS overlay network.

We support configuring `type-1-lacp` and `type-3-system-mac` options in all-active multihome devices in EVPN-VXLAN and EVPN-MPLS overlay networks.

For more guidelines on configuring automatic ESI and information about how the ESI is derived, see *Understanding Automatically Generated ESIs in EVPN Networks*.

Options

lacp-pe-system-id-and-admin-key

Specify that the ESI is automatically derived from the system id and admin key in the LACP configuration for the aggregated Ethernet interface on the local device (actor).

> **NOTE**: In Junos OS 22.2R1,`lacp` was rename to `lacp-pe-system-id-and-admin-key`.

type-1-lacp

Specify that the ESI is automatically derived from the system id and admin key in the LACP configuration for the aggregated Ethernet interface on the remote device (partner).

type-3-system-mac

Specify that the ESI is automatically derived from the configured local-discriminator and mac options.

- local-discriminator *number*—Specify a unique number that will be used in conjunction with the mac option to create the ESI.

  Range: 0 through 16777215

- mac *mac-address*—Specify a system identfier in the form of a MAC address that will be used in conjunction with the local-discriminator to create the ESI.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced on MX Series routers, QFX Series switches, and Junos on White Box in Junos OS Release 18.4R1.

`type-1-lacp` and `type-3-system-mac` options introduced in Junos OS Release 22.2R1.

`lacp` option renamed and aliased to `lacp-pe-system-id-and-admin-key` in Junos OS Release 22.2R1.

Support for EVPN-MPLS introduced in Junos OS Release 23.1R1 on MX series and in Junos OS Evolved Release 23.1R1 on ACX series and PTX routers.

auto-discovery

IN THIS SECTION

- Syntax | 2710
- Hierarchy Level | 2710
- Description | 2710
- Required Privilege Level | 2710
- Release Information | 2710

Syntax

```
auto-discovery;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id]
[edit logical-systems name protocols oam ethernet connectivity-fault-management maintenance-
domain name maintenance-association name mep mep-id]
```

Description

Enable the MEP to accept continuity check messages from all remote MEPs.

The `auto-discovery` option is not available on SRX1500, SRX4100, SRX4200, and SRX4600 devices.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement introduced in Junos OS Release 10.2 for EX Series Switches.

RELATED DOCUMENTATION

Configuring a MEP to Generate and Respond to CFM Protocol Messages

auto-discovery-only

IN THIS SECTION

- Syntax | 2711
- Hierarchy Level | 2711
- Description | 2712
- Required Privilege Level | 2712
- Release Information | 2712

Syntax

```
auto-discovery-only;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family l2vpn],
[edit logical-systems logical-system-name protocols bgp group group-name family l2vpn],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
l2vpn],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp family
l2vpn],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name family l2vpn],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name neighbor address family l2vpn],
[edit protocols bgp family l2vpn],
[edit protocols bgp group group-name family l2vpn],
[edit protocols bgp group group-name neighbor address family l2vpn],
[edit routing-instances instance-name protocols bgp family l2vpn],
[edit routing-instances instance-name protocols bgp group group-name family l2vpn],
[edit routing-instances instance-name protocols bgp group group-name neighbor address family
l2vpn]
```

Description

Enable the router to process only the autodiscovery network layer reachability information (NLRI) update messages for VPWS and LDP-based Layer 2 VPN and VPLS update messages (BGP_L2VPN_AD_NLRI) (FEC 129).

Specifically, the `auto-discovery-only` statement notifies the routing process (rpd) to expect autodiscovery-related NLRI messages so that information can be deciphered and used by LDP, VPLS, and VPWS.

The `auto-discovery-only` statement must be configured on all provider edge (PE) routers in a VPLS or in a VPWS. If you configure route reflection, the `auto-discovery-only` statement is also required on provider (P) routers that act as the route reflector in supporting FEC 129-related updates.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4R2.

RELATED DOCUMENTATION

Example: Configuring BGP Autodiscovery for LDP VPLS

Example: Configuring BGP Autodiscovery for LDP VPLS with User-Defined Mesh Groups

Example: Configuring FEC 129 BGP Autodiscovery for VPWS

auto-export

IN THIS SECTION

- Syntax | 2713
- Hierarchy Level | 2714
- Description | 2714

- Default | 2714
- Options | 2715
- Required Privilege Level | 2715
- Release Information | 2715

Syntax

```
auto-export {
    disable;
    family inet {
        disable;
        flow {
            disable;
            rib-group rib-group;
        }
        multicast {
            disable;
            rib-group rib-group;
        }
        unicast {
            disable;
            rib-group rib-group;
        }
    }
    family inet6 {
        disable;
        multicast {
            disable;
            rib-group rib-group;
        }
        unicast {
            disable;
            rib-group rib-group;
        }
    }
    family iso {
        disable;
        unicast {
```

```
            disable;
            rib-group rib-group;
        }
    }
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable | no-world-
readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Export routes between routing instances.

This statement enables you to leak routes between VPN routing and forwarding (VRF) instances that are locally configured on a provider edge (PE) router. Auto export is always applied on the local PE router, because it applies to only local prefix leaking by evaluating the export policy of each VRF and determining which route targets can be leaked. The standard VRF import and export policies affect remote PE prefix leaking.

You can use this statement as an alternative to using the VRF import and export policies.

Default

`auto-export` is disabled by default. If you want to export routes between routing instances, include the `auto-export` CLI statement at the [`edit routing-options`] or at the [edit routing-instances *routing-instance-name* routing-options] hierarchy levels. You can use the `show route export` CLI command to verify the route export information.

Options

family—Address family.

inet—IP version 4 (IPv4) address family.

multicast—Multicast routing information.

unicast—Unicast routing information.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Technology Overview: Understanding the Auto Export Feature

auto-image-upgrade

IN THIS SECTION

- Syntax | 2716
- Hierarchy Level | 2716
- Description | 2716
- Default | 2716
- Required Privilege Level | 2716
- Release Information | 2716

Syntax

```
auto-image-upgrade;
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable automatic software download using Zero Touch Provisioning (ZTP).

If the `auto-image-upgrade` statement is included in the factory default configuration, then ZTP will start automatically when the device is booted up. Otherwise, ZTP will start with the first power cycle after the `auto-image-update` statement is configured.

The DHCP client compares the software package name in the DHCP server message to the name of the software package that booted the switch. If the software packages are different, the DHCP client downloads and installs the software package specified in the DHCP server message.

Before you upgrade software using automatic software download, ensure that you have configured DHCP services, including configuring a path to a boot server and a boot file. See the *Junos OS System Basics Configuration Guide* for information about using the CLI to configure DHCP services and settings.

Default

Automatic software download is disabled.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Upgrading Software by Using Automatic Software Download for Switches

Understanding Software Installation on EX Series Switches

auto-load-rebalance

IN THIS SECTION

- Syntax | 2717
- Hierarchy Level | 2717
- Description | 2717
- Required Privilege Level | 2717
- Release Information | 2718

Syntax

```
auto-load-rebalance;
```

Hierarchy Level

```
[edit fc-fabrics fabric-name proxy]
```

Description

Configure the system to rebalance NP_Port link loads automatically on an FCoE-FC gateway proxy fabric if the link loads become unbalanced. Load rebalancing is a disruptive action that forces some or all sessions (depending on the configured load-balancing algorithm) to log out and then log in again. When sessions log in again, they are placed on NP_Port interfaces so that the link loads are balanced.

Required Privilege Level

storage—To view this statement in the configuration.

storage-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Defining the Proxy Load-Balancing Algorithm

Example: Configuring Automated Fibre Channel Interface Load Rebalancing

Simulating On-Demand Fibre Channel Link Load Rebalancing (Dry Run Test)

Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric

Monitoring Fibre Channel Interface Load Balancing

auto-negotiation (QFX Series)

IN THIS SECTION

- Syntax | 2718
- Hierarchy Level | 2718
- Description | 2719
- Required Privilege Level | 2719
- Release Information | 2719

Syntax

```
(auto-negotiation | no-auto-negotiation);
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
```

Description

Explicitly enable or disable autonegotiation. Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 1G and 10G.

- **auto-negotiation**—Enable autonegotiation.
- **no-auto-negotiation**—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure link mode and speed options.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

speed (Ethernet) | 28349

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

Junos OS Network Interfaces Library for Routing Devices

auto-negotiation

IN THIS SECTION

- Syntax | 2720
- Hierarchy Level | 2720
- Description | 2720
- Default | 2723
- Options | 2723
- Required Privilege Level | 2723

Release Information | 2723

Syntax

```
(auto-negotiation | no-auto-negotiation) <remote-fault (local-interface-online | local-interface-
offline)>;
```

Hierarchy Level

```
[edit interfaces interface-name ether-options],
[edit interfaces interface-name gigether-options],
[edit interfaces ge-pim/0/0 switch-options switch-port port-number]
```

Description

For Gigabit Ethernet interfaces on EX Series devices enable auto-negotiation. Auto-negotiation is supported in some ACX devices. In ACX5448, ACX5448-D, ACX5448-M and ACX710, auto-negotiation is supported on 1G interface. Beginning with Junos 21.1R1 and beyond, no-auto-negotiation configuration through CLI is supported on ports 24 onwards on ACX5448 variants and port 16 onwards on ACX710 due to hardware limitation for SerDes lanes.

- `auto-negotiation`—Enables autonegotiation. This is the default.
- `no-auto-negotiation`—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure the link mode and speed.

When you configure Tri-Rate Ethernet copper interfaces to operate at 1 Gbps, autonegotiation must be enabled.

> **NOTE**: On EX Series switches, an interface configuration that disables autonegotiation and manually sets the link speed to 1 Gbps is accepted when you commit the configuration; however, if the interface you are configuring is a Tri-Rate Ethernet copper interface, the configuration is ignored as invalid and autonegotiation is enabled by default.

To correct the invalid configuration and disable autonegotiation:

1. Delete the `no-auto-negotiation` statement and commit the configuration.
2. Set the link speed to 10 or 100 Mbps, set `no-auto-negotiation`, and commit the configuration.

On EX Series switches, if the link speed and duplex mode are also configured, the interfaces use the values configured as the desired values in the negotiation. If autonegotiation is disabled, the link speed and link mode must be configured.

NOTE: 1G optical transceivers do not support auto-negotiation on EX4100 1G/10G uplink module (PIC 1/PIC 2) and EX4400 1G/10G interface. 1G-T supports auto-negotiation on these ports.

NOTE: On T4000 routers, the `auto-negotiation` command is ignored for interfaces other than Gigabit Ethernet.

NOTE: On ACX Series routers, when you configure fiber interfaces (fiber media mode) to operate at 1 Gbps, autonegotiation is enabled by default to negotiate the speed and duplex settings. You can disable autonegotiation on ACX Series routers, except ACX5448 routers, in the fiber media mode, by using the (`no-auto-negotiation`) statement, and commit the configuration. In copper interfaces (copper media mode), autonegotiation is enabled by default. To disable autonegotiation, you need to explicitly configure the link speed to 10 or 100 Mbps, set `no-auto-negotiation`, and commit the configuration.

NOTE: Auto-negotiation is not supported on ACX710 and ACX5448 at 10G speed.
1G mode is not supported on pic level port profile and port level port profile.

1G is supported in *no-autonegotiation* mode. This indicates link partner must disable autonegotiation.

Timing is not supported.

NOTE: The auto-negotiation parameter *link partner speed* is displayed based on local speed even when auto-negotiation is disabled on ACX710 and ACX5448.

> **NOTE**: For Gigabit Ethernet interfaces on NFX350 devices, `auto-negotiation` is enabled by default. The `no-auto-negotiation` command is ignored for Gigabit Ethernet interfaces on NFX350 devices.

For SRX Series Firewalls, when autonegotiatiation is disabled, you can set the `mdi-mode` to enable it in case of non-cross table.

From Junos OS Evolved Release 21.2R1, we support autonegotiation and link training using 400G DAC cable for PTX10001-36MR. This is not a default behavior. Use `auto-negotiation` under `[edit interfaces interface-name ether-options]` hierarchy level to enable autonegotiation and link training. Enabling autonegotiation automatically enables link training. This advertises configured speed on the interface hierarchy. It does not upgrade or downgrade speed based on peers advertised capability.

Autonegotiation and link training is supported only on one 400-Gigabit Ethernet interfaces (1x400G), two 200-Gigabit Ethernet interfaces (2x200G), and four 100-Gigabit Ethernet interfaces (4x100G) channelization.

Starting in Junos OS Release 22.1R1, we support auto-negotiation for 1G interface on ACX5448, ACX5448-D, ACX5448-M and ACX710. Use the existing `auto-negotiation | no-auto-negotiation` statement at the `[edit interfaces interface-name gigether-options]` hierarchy level to enable and disable auto-negotiation for 1G interfaces.

No auto-negotiation configuration through CLI will be supported from port 24 onwards on ACX5448 variants and port 16 onwards on ACX710 due to hardware limitation for serdes lanes. No-auto-negotiation configuration on unsupported ports will be ignored and alerted to the user with a PFE log message.

Table 4 on page 2722 shows the behaviour of auto-negotiation in different combinations.

Table 4: Behavior of auto-negotiation

Device (ACX5448/ACX710)	Peer device	Link status
auto-negotiation ON	auto-negotiation ON	UP
auto-negotiation ON	auto-negotiation OFF	Results in anomalous behaviour due to duplex mismatch
auto-negotiation OFF	auto-negotiation ON	Results in anomalous behaviour due to duplex mismatch

Table 4: Behavior of auto-negotiation *(Continued)*

Device (ACX5448/ACX710)	Peer device	Link status
auto-negotiation OFF	auto-negotiation OFF	UP (If the speed and duplex mode matches on both the ends) DOWN (If the speed or duplex mode does not match on both the ends)

Default

Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.

Options

`remote-fault (local-interface-online | local-interface-offline)`—(Optional) For M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers only, manually configure remote fault on an interface.

- **Default:** `local-interface-online`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Gigabit Ethernet Autonegotiation Overview

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support

auto-policing

IN THIS SECTION

- Syntax | 2724
- Hierarchy Level | 2724
- Description | 2724
- Options | 2724
- Required Privilege Level | 2725
- Release Information | 2725

Syntax

```
auto-policing {
    class all (drop | loss-priority-high | loss-priority-low);
    class ctnumber (drop | loss-priority-high | loss-priority-low);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Enable the automatic policing of all the MPLS LSPs on the router or logical system.

Options

`class all`—Apply the same policer action to all the class types (`ct0`, `ct1`, `ct2`, and `ct3`).

`class ct`*`number`*—Specific class type (`ct0`, `ct1`, `ct2`, or `ct3`) to which to apply a policer action.

Policer actions—You can specify the following policer actions:

- **Default:** no action
- `drop`—Drop all packets.
- `loss-priority-high`—Set the packet loss priority (PLP) to high.
- `loss-priority-low`—Set the PLP to low.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for QFX10000 Series switches in release 15.1X53-D40.

Statement introduced for MX240, MX480, MX960, MX2010, and MX2020 routers with MPC10E and MPC11E line cards in Junos OS 20.2R1.

RELATED DOCUMENTATION

policing (Protocols MPLS) | **12664**

Configuring Automatic Policers

auto-provisioned

IN THIS SECTION

- Syntax | **2726**
- Hierarchy Level | **2726**
- Description | **2726**
- Required Privilege Level | **2726**
- Release Information | **2726**

Syntax

```
auto-provisioned;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Enable the auto-provisioned configuration mode for a Virtual Chassis Fabric (VCF).

When a VCF is autoprovisioned, you can plug and play leaf devices into your VCF. You do not need to manually configure the required VCF parameters. When you interconnect a leaf device into the VCF, the leaf device is automatically configured into the linecard role. All other required VCF settings also happen without further user action, such as configuring the member ID, setting fabric mode or mixed mode, and setting up the Virtual Chassis ports (VCPs).

For autoprovisioning to work:

- The leaf device must be zeroized or no VCF parameters have been previously set or configured on the device.
- You must interconnect the leaf device to a spine device using a 10-Gbps SFP+ or 40-Gbps QSFP+ link that is supported as a VCP and is not already converted into a VCP on *both* sides of the link. See for details.

If the VCF automatically sets fabric or mixed mode on a spine or leaf device during autoprovisioning, the VCF then automatically reboots the device to complete the configuration. The device must reboot to apply the mode settings and properly join the VCF. If you want to more deterministically control when a device reboots upon adding it to an autoprovisioned VCF, you can alternatively configure fabric or mixed mode settings manually and reboot the device yourself *before* interconnecting it into the VCF.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D20.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric

Understanding Virtual Chassis Fabric Configuration

auto-re-enrollment

IN THIS SECTION

- Syntax | 2727
- Hierarchy Level | 2727
- Description | 2728
- Options | 2728
- Required Privilege Level | 2728
- Release Information | 2728

Syntax

```
auto-re-enrollment {
    certificate-id {
        ca-profile ca-profile-name;
        challenge-password password;
        re-enroll-trigger-time-percentage percentage;
        re-generate-keypair;
        validity-period days;
    }
}
```

Hierarchy Level

```
[edit security pki]
```

Description

Specify auto-reenrollment parameters for a certificate authority (CA) issued router certificate. Auto-reenrollment requests that the issuing CA replace a router certificate before its specified expiration date.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Auto-Reenrollment of a Router Certificate

Configuring Digital Certificates for Adaptive Services Interfaces

auto-re-enrollment (Security)

IN THIS SECTION

- Syntax | 2729
- Hierarchy Level | 2730
- Description | 2730
- Options | 2730
- Required Privilege Level | 2731
- Release Information | 2731

Syntax

```
auto-re-enrollment {
 acme {
        certificate-id certificate-id-name {
            acme-key-id acme-key-id;
            ca-profile-name ca-profile-name;
            re-enroll-time (days value| hours value| percentage value);
            re-enroll-trigger-time-percentage percentage;
            re-generate-keypair;
        }
    }
    cmpv2 {
        certificate-id certificate-id-name {
            ca-profile-name ca-profile-name;
            challenge-password password;
            re-enroll-trigger-time-percentage percentage;
            re-enroll-time (days value| hours value| percentage value);
            re-generate-keypair;
        }
    }
    scep {
        certificate-id certificate-id-name {
            ca-profile-name ca-profile-name;
            challenge-password password;
            re-enroll-trigger-time-percentage percentage;
            re-enroll-time (days value| hours value| percentage value);
            re-generate-keypair;
            scep-digest-algorithm {
            (md5 | sha1);
             }
            scep-encryption-algorithm {
            (des | des3);
             }
        }
    }
}
```

Hierarchy Level

```
[edit security pki]
```

Description

Configure the automatic reenrollment of a local end-entity (EE) certificate. Auto-reenrollment requests that the issuing CA replace a device certificate before its specified expiration date.

Options

certificate-id Auto reenrollment configuration for certificate ID.

acme-key-id Specify the ACME account key identifier.

ca-profile-name Specify the name of the certificate authority (CA) profile to be used for automatic reenrollment. The CA certificate must be present to initiate reenrollment.

challenge-password Specify the password used by the certificate authority (CA) for enrollment and revocation. If the CA does not provide the challenge password, choose your own password.

re-enroll-trigger-time-percentage Specify the certificate reenrollment trigger as a percentage of the end-entity (EE) certificate's lifetime that remains before certificate reenrollment is initiated. For example, if the renewal request is to be sent when the certificate's remaining lifetime is 10 percent, then configure 10 for `re-enroll-trigger-time-percentage` value. The time at which the certificate reenrollment is initiated is based on the certificate expiry date.

- **Range:** 1 through 99

re-enroll-time This option allows you to trigger auto-re-enrollment ahead of the certificate expiration. You can configure the re-enrollment trigger time in days, or hours, or percentage.

- days *value*—Specify when to trigger re-enrollment in days.
- hours *value*—Specify when to trigger re-enrollment in hours.
- percentage *value*—Specify when to trigger re-enrollment in percentage. **Range**: 1 to 99.

If you configure both `re-enroll-trigger-time-percentage` and `re-enroll-time` options, then `re-enroll-time` configuration take precedence.

Starting Junos OS Release 23.1R1, you must configure either `re-enroll-trigger-time-percentage` or `re-enroll-time` for the `commit check` to be successful.

re-generate-keypair Specify new key pair generation for automatic certificate reenrollment. If this statement is not configured, the current key pair is used. If the key pair does not change, the CA does not issue new certificates. We recommend that a new key pair be generated during reenrollment as it provides better security.

scep-digest-algorithm SCEP digest algorithm.

- Values:
  - md5—Use MD5 as SCEP digest algorithm
  - sha1—Use SHA1 as SCEP digest algorithm

scep-encryption-algorithm SCEP encryption algorithm.

- Values:
  - des—Use DES as SCEP encryption algorithm
  - des3—Use DES3 as SCEP encryption algorithm

cmpv2 Configure automatic reenrollment of a local certificate using CMPv2.

scep Configure automatic reenrollment of a local certificate using Simple Certificate Enrollment Protocol (SCEP).

acme You can configure the ACME auto re-enrollment.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.0. `cmpv2` and `scep` options added in Junos OS Release 15.1X49-D40.

Support for `re-enroll-time (days value| hours value| percentage value)` option added in Junos OS Release 21.4R1.

The acme optioin added in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

Basic Elements of PKI in Junos OS

show security ipsec statistics (MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall) | **25023**

auto-reconnect

IN THIS SECTION

- Syntax | **2732**
- Hierarchy Level | **2732**
- Description | **2733**
- Options | **2733**
- Required Privilege Level | **2733**
- Release Information | **2733**

Syntax

```
auto-reconnect seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number pppoe-options],
[edit logical-systems logical-system-name interfaces interface-name
unit logical-unit-number pppoe-options]
```

Description

PPP over Ethernet interfaces, configure the amount of time to wait before reconnecting after a session has terminated.

Options

seconds—Time to wait before reconnecting after a session has terminated.

- **Range:** 0 through 4,294,967,295 seconds
- **Default:** 0 (never)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the PPPoE Automatic Reconnect Wait Timer

auto-recovery

IN THIS SECTION

- Syntax | 2734
- Hierarchy Level | 2734
- Description | 2734
- Required Privilege Level | 2734
- Release Information | 2734

Syntax

```
auto-recovery;
```

Hierarchy Level

```
[edit logical-systems name protocols isis dynamic-overload],
[edit logical-systems name routing-instances name protocols isis dynamic-overload],
[edit logical-systems name tenants name routing-instances name protocols isis dynamic-overload],
[edit protocols isis dynamic-overload],
[edit routing-instances name protocols isis dynamic-overload],
[edit tenants name routing-instances name protocols isis dynamic-overload]
```

Description

Enable Auto-recovery to recover from dynamic overload state.

IS-IS enters a dynamic overload state when the prefix limit is reached or when the LSP fragment space is exhausted. Enabling auto-recovery prevents ISIS from flushing of all the LSP fragments. When the fragment space becomes available, IS-IS exits from the dynamic overload state.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.1

RELATED DOCUMENTATION

overload (Protocols IS-IS) | **12048**

auto-rp

IN THIS SECTION

- Syntax | 2735
- Hierarchy Level | 2735
- Description | 2735
- Options | 2735
- Required Privilege Level | 2736
- Release Information | 2736

Syntax

```
auto-rp {
    (announce | discovery | mapping);
    (mapping-agent-election | no-mapping-agent-election);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Configure automatic RP announcement and discovery.

Options

announce—Configure the routing device to listen only for mapping packets and also to advertise itself if it is an RP.

discovery—Configure the routing device to listen only for mapping packets.

mapping—Configures the routing device to announce, listen for and generate mapping packets, and announce that the routing device is eligible to be an RP.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

The auto-rp options `announce` and `mapping` are not supported on QFX5220-32CD devices running Junos OS Evolved Release 19.3R1, 19.4R1, or 20.1R1.

RELATED DOCUMENTATION

Configuring PIM Auto-RP

auto-satellite-conversion (Junos Fusion)

IN THIS SECTION

- Syntax | 2737
- Hierarchy Level | 2737
- Description | 2737
- Options | 2737
- Required Privilege Level | 2737
- Release Information | 2737

Syntax

```
auto-satellite-conversion {
    satellite [slot-id | range | all];
}
```

Hierarchy Level

```
[edit chassis satellite-management]
```

Description

Enable automatic satellite conversion in a Junos Fusion.

Automatic satellite conversion automatically configures a switch into a satellite device when it is cabled into the aggregation device.

Additional configuration steps are required to add satellite devices to a Junos Fusion using automatic satellite conversion. See *Configuring Junos Fusion Provider Edge* or Configuring or Expanding a Junos Fusion Enterprise.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

auto-sense

IN THIS SECTION

- Syntax | 2738
- Hierarchy Level | 2738
- Description | 2739
- Options | 2739
- Required Privilege Level | 2739
- Release Information | 2739

Syntax

```
auto-sense {
    hold-down hold-down;
}
```

Hierarchy Level

```
[edit fabric protocols bgp group name link-bandwidth],
[edit logical-systems name protocols bgp group name link-bandwidth],
[edit logical-systems name routing-instances name protocols bgp group name link-bandwidth],
[edit logical-systems name tenants name routing-instances name protocols bgp group name link-
bandwidth],
[edit protocols bgp group name link-bandwidth],
[edit routing-instances name protocols bgp group name link-bandwidth]
```

Description

Enable auto-sensing of link-bandwidth and attaching the link-bandwidth community on a route at import. BGP can communicate link speeds to remote peers, enabling better optimization of traffic distribution for load balancing. This facilitates the integration of devices with different transmission speeds within the network, enabling efficient traffic distribution based on link speed.

Note that when more than one link-bandwidth community is attached to the route, BGP ignores transitivity and installs the lowest link-bandwidth value to the FIB, also known as the forwarding table

Options

hold-down Hold down in secs for link bandwidth auto-sense

- **Range:** 0 through 3600

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

community (Policy Options) | **4465**

group (Protocols BGP) | **7863**

auto-service-id

IN THIS SECTION

Syntax | **2740**

- Hierarchy Level | 2740
- Description | 2740
- Required Privilege Level | 2740
- Release Information | 2740

Syntax

```
auto-service-id;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols evpn],
[edit routing-instances name protocols evpn]
```

Description

Enable auto-derivation of EVPN-VPWS service instance identifier. When you enable `auto-service-id`, the device derives the VPWS service ID from the normalized VLAN ID.

NOTE: This option only applies to the EVPN-VPWS flexible cross-connect (FXC) VLAN aware service (VLAN signaled flexible cross connect).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

flexible-cross-connect-vlan-aware | 7230

auto-snapshot

IN THIS SECTION

- Syntax | 2741
- Hierarchy Level | 2741
- Description | 2741
- Default | 2742
- Required Privilege Level | 2742
- Release Information | 2742

Syntax

```
auto-snapshot;
```

Hierarchy Level

```
[edit system]
```

Description

Enable the automatic snapshot feature, which allows the device to automatically fix a corrupt Junos OS file in the primary root partition. If the automatic snapshot feature is enabled, the device automatically takes a snapshot of the Junos OS root file system in the alternate root partition and copies it onto the primary root partition, thereby repairing the corrupt file in the primary root partition. The automatic snapshot procedure takes place whenever the system reboots from the alternate root partition, regardless of whether the reboot is due to a command or due to corruption of the primary root partition.

NOTE: EX9200 devices do not support the automatic snapshot feature.

Default

- The automatic snapshot feature is enabled by default on the following EX Series devices:
  - EX4550 devices
  - EX Series devices that ship with Junos OS Release 12.3R1 or later
- The automatic snapshot feature is disabled by default on EX Series devices (except the EX4550 devices) running Junos OS Release 12.2 or earlier.
- If the automatic snapshot feature was disabled by default before the device was upgraded to Junos OS Release 12.3R1 or later, the feature remains disabled (for backward compatibility) by default after the upgrade.
- By default, the auto-snapshot feature is disabled on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550M devices.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring Dual-Root Partitions

show system auto-snapshot | **26720**

auto-sw-sync

IN THIS SECTION

- Syntax | 2743
- Hierarchy Level | 2743
- Description | 2743
- Default | 2744
- Options | 2744
- Required Privilege Level | 2744
- Release Information | 2744

Syntax

```
auto-sw-sync node node-name (disable | enable);
```

Hierarchy Level

```
[edit system]
```

Description

(Junos OS Evolved only) When you add a new Routing Engine (RE) to the device and the new RE has a different software version than the rest of the system, by default the RE is kept out of the system. If you want a new RE's software and configuration to be automatically upgraded and synchronized with that of the cluster, configure this statement. Once configured for the device, if an RE fails and you replace the RE, the software and the configuration from the primary RE in the system automatically installs on the new RE.

When you configure this statement, the primary Routing Engine of the system copies over all the images (software and configuration) to the new Routing Engine and reboots the new Routing Engine so it runs the same software version and configuration as the primary Routing Engine. Each software image also contains the configuration running when the software image was last active.

When the chassis first comes up, the Routing Engines elect a "primary" node based on several factors, including which Routing Engine was "primary" last, which Routing Engine is the current hardware primary RE, and the slot position (0 versus 1).

Default

Disabled: When you insert a new Routing Engine with a different software version than the rest of the system and you have not already configured this statement on the system, the Routing Engine is kept out of the system. Thereafter, the newly-inserted RE does not respond to any software event and remains in its original software version.

Options

disable | enable (Required) Specify whether to disable or enable automatic software synchronization from the primary node to the new Routing Engine.

- Default: Disable

node *node-name* Specify the node to be synchronized (`fpc0` | `re0` | `re1`). Deprecated as of Junos OS Evolved Release 20.4R2.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.2R1.

`node` option deprecated as of Junos OS Evolved Release 20.4R2.

RELATED DOCUMENTATION

Replace a Routing Engine in a Dual-Routing Engine System

auto-sw-update

IN THIS SECTION

- Syntax | 2745
- Hierarchy Level | 2745
- Description | 2745
- Default | 2746
- Options | 2747
- Required Privilege Level | 2748
- Release Information | 2748

Syntax

```
auto-sw-update {
    (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
    package-name package-name;
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Enable the automatic software update feature for Virtual Chassis or Virtual Chassis Fabric (VCF) configurations.

You should only use the keywords that specify a device or device family—`ex-4300`, `ex-4600`, `qfx-3`, and `qfx-5`—when configuring automatic software update on a mixed mode Virtual Chassis or Virtual Chassis Fabric (VCF). You can simply specify the *package-name* without specifying individual device keywords in non-mixed Virtual Chassis or VCF topologies.

NOTE: The automatic software update feature is not supported for a mixed mode EX4300 Virtual Chassis with a combination of EX4300 multigigabit model (EX4300-48MP) switches and other EX4300 switches.

You must enter the `auto-sw-update` statement multiple times—once for each device family in your mixed Virtual Chassis or VCF—in most scenarios when enabling the automatic software update for a mixed Virtual Chassis or VCF.

The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. You do not, therefore, need to specify the `ex-4500` keyword when configuring automatic software update for a mixed Virtual Chassis that include EX4500 and EX4550 switches only. You also only have to enter the `ex-4500` keyword once to configure automatic software update for all EX4500 and EX4550 member switches in the s37628)ame mixed Virtual Chassis.

The Junos OS package for a QFX3500 device updates the software for QFX3500 and QFX3600 devices. You do not, therefore, need to specify the `qfx-3` keyword when configuring automatic software update for a Virtual Chassis composed entirely of QFX3500 and QFX3600 devices. You also have to enter the `qfx-3` keyword only once to configure automatic software update for all QFX3500 and QFX3600 member devices in the same mixed Virtual Chassis.

The Junos OS package for a QFX5110 device updates the software for both QFX5110 and QFX5100 devices in a QFX5110 Virtual Chassis or VCF, so you do not need to specify the `qfx-5` keyword when configuring automatic software update for a QFX5110 Virtual Chassis or VCF composed of both types of devices.

CAUTION: A QFX5100 switch running a Junos OS software image with “-qfx-5-” in the package filename *must* first be upgraded to a Junos OS software image with “-qfx-5e-” in the package filename before it can be added to a mixed QFX5110 Virtual Chassis or VCF. The automatic software update process cannot update a switch from a “-qfx-5-” image to a “-qfx-5e-” image. See *Upgrading a QFX5100 Switch with a USB Device to Join a QFX5110 Virtual Chassis or Virtual Chassis Fabric*.

After a QFX5100 switch is installed with a “-qfx-5e-” Junos OS software image, the automatic software update process can successfully update the switch automatically with a different version of a “-qfx-5e-” Junos OS image to match the other members in the Virtual Chassis or VCF.

The remaining statement is explained separately. See CLI Explorer.

Default

The automatic software update feature is disabled.

Options

package-name *package-name*

Specify the software package name or a path to a Junos OS software image to update new members being added to a Virtual Chassis or VCF. .

- If the software package is located on a local directory on the switch, use the following format for *package-name*:

  /pathname/package-name

- If the software package is to be downloaded and installed from a remote location, use one of the following file access and transfer formats:

  `ftp://`*hostname/pathname/package-name*

  `ftp://`*username*`:`*prompt*`@ftp.`*hostname.net/package-name*

  `http://`*hostname/pathname/package-name*

ex-4200

Specify a path to a Junos OS image for an EX4200 switch when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

ex-4300

Specify a path to a Junos OS image for an EX4300 switch when enabling automatic software update for a mixed Virtual Chassis or VCF. Automatic software update is not supported for mixed mode EX4300 Virtual Chassis with a combination of EX4300 multigigabit model (EX4300-48MP) switches and other EX4300 switches, and this option is not applicable in that case.

ex-4500

Specify a path to a Junos OS image for an EX4500 switch, an EX4550 switch, or both types of switches when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. Therefore, you only enter this command once to upgrade the EX4500 and EX4550 member switches in the same mixed Virtual Chassis.

The `ex-4500` keyword also does not need to be specified when configuring automatic software update for a mixed EX4500 and EX4550 Virtual Chassis.

ex-4600

Specify a path to a Junos OS image for an EX4600 switch when enabling automatic software update for a mixed Virtual Chassis.

qfx-3

Specify a path to a Junos OS image for a QFX3500, QFX3600, or both types of devices when enabling automatic software update for a mixed Virtual Chassis or VCF.

qfx-5 Specify a path to a Junos OS image for a QFX5100 device when enabling automatic software update for a mixed QFX5100 Virtual Chassis or VCF. You do not need to use this option for a QFX5110 Virtual Chassis or VCF that contains both QFX5110 and QFX5100 switches because the same software image runs on either type of switch.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statements introduced in Junos OS Release 10.0 for EX Series switches.

The `ex-4200` and `ex-4500` options introduced in Junos OS Release 12.2 for EX Series switches.

Statements introduced in Junos OS Release 13.2X50-D15 for the QFX Series.

The `ex-4300`, `qfx-3`, and `qfx-5` options introduced in Junos OS Release 13.2X51-D20.

Statements introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

The `ex-4600` option introduced in Junos OS Release 13.2X51-D25.

RELATED DOCUMENTATION

- Configuring Automatic Software Update on Virtual Chassis Member Switches
- Understanding Software Upgrades in a Virtual Chassis
- *Understanding Software Upgrades in a Virtual Chassis Fabric*
- Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches

auto-targeted-session

IN THIS SECTION

- Syntax | 2749
- Hierarchy Level | 2749

- Description | 2749
- Options | 2749
- Required Privilege Level | 2750
- Release Information | 2750

Syntax

```
auto-targeted-session {
    maximum-sessions number;
    teardown-delay seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit protocols ldp]
```

Description

Configure session parameters for LDP sessions established with the remote LFA node that are automatically targeted using the loopback addresses. Configure parameters of automatically targeted sessions for remote LFA only.

Options

maximum-sessions *number*

Specify the maximum number of auto-targeted LDP sessions allowed. Include this statement to optimize the use of router memory.

- **Default:** 100
- **Range:** 1 through 1000

teardown-delay *seconds*

Specify the minimum time period for which the auto-targeted session must be alive before tearing down the auto-targeted LDP sessions to the remote LFA node. Include

this statement to prevent rapid route-resolution in case of temporary change in IGP topology.

- **Default:** 90 seconds
- **Range:** 1 through 300 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

autodiscovery

IN THIS SECTION

- Syntax | 2750
- Hierarchy Level | 2751
- Description | 2751
- Options | 2751
- Required Privilege Level | 2751
- Release Information | 2751

Syntax

```
autodiscovery {
    inet-mdt;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
mvpn family inet],
[edit routing-instances routing-instance-name protocols pim mvpn family inet]
```

Description

For draft-rosen 7, enable the PE routers in the VPN to discover one another automatically.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Statement moved to `[..protocols pim mvpn family inet]` from `[.. protocols pim mvpn]` in Junos OS Release 13.3.

RELATED DOCUMENTATION

> *Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs*

autodiscovery-only

IN THIS SECTION

- Syntax | 2752
- Hierarchy Level | 2752
- Description | 2752
- Required Privilege Level | 2752
- Release Information | 2753

Syntax

```
autodiscovery-only {
    intra-as {
        inclusive;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvpn
family inet / inet6],
[edit routing-instances routing-instance-name protocols mvpn family (inet / inet6)]
```

Description

Enable the Rosen multicast VPN to use the MDT-SAFI autodiscovery NLRI.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Statement moved to `[..protocols pim mvpn family inet]` from `[.. protocols mvpn]` in Junos OS Release 13.3.

Support for IPv6 added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs

autoinstallation

IN THIS SECTION

- Syntax | 2753
- Hierarchy Level | 2754
- Description | 2754
- Options | 2754
- Required Privilege Level | 2754
- Release Information | 2754

Syntax

```
autoinstallation;
```

```
autoinstallation {
    configuration-servers {
        url {
             password password;
        }
    }
```

```
    interfaces {
        interface-name {
            bootp;
            rarp;
        }
    }
    usb {
        disable;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Specify the configuration for autoinstallation. Download a configuration file automatically from an FTP, Hypertext Transfer Protocol (HTTP), or Trivial FTP (TFTP) server. When you power on a router or switch configured for autoinstallation, it requests an IP address from a Dynamic Host Configuration Protocol (DHCP) server. Once the router or switch has an address, it sends a request to a configuration server and downloads and installs a configuration.

Options

There are no options.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- ACX Series Autoinstallation Overview
- Before You Begin Autoinstallation on an ACX Series Universal Metro Router
- Autoinstallation Configuration of ACX Series Universal Metro Routers
- USB Autoinstallation on ACX Series Routers
- Verifying Autoinstallation on ACX Series Universal Metro Routers
- show system autoinstallation status
- Upgrading Software by Using Automatic Software Download for Switches
- *Configuring Autoinstallation of Configuration Files*

autoinstallation (JNU Satellite Devices)

IN THIS SECTION

- Syntax | 2755
- Hierarchy Level | 2756
- Description | 2756
- Options | 2756
- Required Privilege Level | 2756
- Release Information | 2756

Syntax

```
autoinstallation {
    delete-after-commit;
    configuration-servers {
        url;
    }
    interfaces {
        interface-name {
            bootp;
            rarp;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

(Satellite devices in a JNU group) Download a configuration file automatically from an FTP or HTTP server. When you power on a router or switch configured for autoinstallation, it requests an IP address from a Dynamic Host Configuration Protocol (DHCP) server. When the router or switch has an address, it sends a request to a configuration server and downloads and installs a configuration.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

- Autoinstallation of Satellite Devices in a Junos Node Unifier Group
- Autoinstallation Process on Satellite Devices in a Junos Node Unifier Group
- *Configuring Autoinstallation of Configuration Files*
- *Configuring Autoinstallation of Configuration Files*
- delete-after-commit (JNU Satellites) | **5278**

automatic (Security)

IN THIS SECTION

- Syntax | 2757
- Hierarchy Level | 2757
- Description | 2757
- Options | 2757
- Required Privilege Level | 2758
- Release Information | 2758

Syntax

```
automatic {
    download-timeout minutes;
    enable;
    interval (Security IDP) hours;
    start-time (Security IDP) start-time;
}
```

Hierarchy Level

```
[edit security idp security-package]
```

Description

Enable the device to automatically download the updated signature database from the specified URL.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

automatic-site-id

IN THIS SECTION

- Syntax | 2758
- Hierarchy Level | 2759
- Description | 2759
- Options | 2759
- Required Privilege Level | 2760
- Release Information | 2760

Syntax

```
automatic-site-id {
    collision-detect-time seconds;
    new-site-wait-time seconds;
    reclaim-wait-time minimum seconds maximum seconds;
    startup-wait-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
site site-name],
[edit routing-instances routing-instance-name protocols vpls site site-name]
```

Description

Enable automatic site identifiers for VPLS routing instances.

When you configure `automatic-site-id` for the first time, you must deactivate and then activate `protocol vpls`. However, if you already have `automatic-site-id` configured, you do not need to deactivate and then activate `protocol vpls`.

Options

`collision-detect-time`—The time in seconds to wait after a claim advertisement is sent to the other routers in a VPLS instance before a PE router can begin using a site identifier. If the PE router receives a competing claim advertisement for the same site identifier during this time period, it initiates the collision resolution procedure for site identifiers.

> **NOTE**: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

new-site-wait-time—The time in seconds to wait to receive VPLS information for a newly configured routing instance or a new site. This time interval is also applied whenever the automatic site identifier feature is activated on a VPLS routing instance other than at startup. Effectively, this timer indicates how long to wait before an attempt is made to allocate a site identifier. This timer is also triggered whenever a VPLS routing instance is enabled.

reclaim-wait-time—The time in seconds to wait to receive VPLS information for a newly configured routing instance or a new site. This time interval is also applied whenever the automatic site identifier feature is activated on a VPLS routing instance other than at startup. Effectively, this timer indicates how long to wait before an attempt is made to allocate a site identifier. This timer is also triggered whenever a VPLS routing instance is enabled. You can configure two values for this option: the **minimum** wait time and the **maximum** wait time.

startup-wait-time—The time in seconds to wait at startup to receive all the VPLS information for the route targets configured on the other PE routers included in the VPLS routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Configuring Automatic Site Identifiers for VPLS

automatic (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 2760
- Hierarchy Level | 2760
- Description | 2761
- Options | 2761
- Required Privilege Level | 2761
- Release Information | 2761

Syntax

```
automatic (random-allocation | round-robin);
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Description

Configure automatic port assignment for source NAT with port translation, except for deterministic NAT. Automatic port assignment uses the port range 1024 through 65535. Specify either random allocation or round-robin allocation. Random allocation randomly assigns a port from the range 1024 through 65535 for each port translation. Round robin allocation first assigns port 1024, and uses the next higher port for each successive port assignment. Round robin allocation is the default.

Options

random-allocation	Randomly assigns a port from the range 1024 through 65535 for each port translation.
round-robin	First assigns port 1024, and uses the next higher port for each successive port assignment. Round robin allocation is the default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

autonomous

IN THIS SECTION

- Syntax | 2762
- Hierarchy Level | 2762
- Description | 2762
- Default | 2762
- Required Privilege Level | 2762
- Release Information | 2762

Syntax

```
(autonomous | no-autonomous);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name prefix prefix],
[edit protocols router-advertisement interface interface-name prefix prefix]
```

Description

Specify whether prefixes in the router advertisement messages are used for stateless address autoconfiguration:

- **autonomous**—Use prefixes for address autoconfiguration.
- **no-autonomous**—Do not use prefixes for address autoconfiguration.

Default

autonomous

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

autonomous (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 2763
- Hierarchy Level | 2763
- Description | 2763
- Default | 2763
- Required Privilege Level | 2764
- Release Information | 2764

Syntax

```
(autonomous | no-autonomous);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols router-advertisement interface interface-name
prefix prefix]
```

Description

Specify whether prefixes in the router advertisement messages are used for stateless address autoconfiguration:

- `autonomous`—Use prefixes for address autoconfiguration.
- `no-autonomous`—Do not use prefixes for address autoconfiguration.

Default

`autonomous`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA

Dynamic Router Advertisement Configuration Overview

autonomous-system

IN THIS SECTION

- Syntax | 2764
- Hierarchy Level | 2765
- Description | 2765
- Options | 2765
- Required Privilege Level | 2767
- Release Information | 2767

Syntax

```
autonomous-system autonomous-system <asdot-notation> <loops number> {
    independent-domain <no-attrset>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Specify the routing device's AS number.

An autonomous system (AS) is a set of routing devices that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it. ASs are identified by a number that is assigned by the Network Information Center (NIC) in the United States (http://www.isi.edu).

If you are using BGP on the routing device, you must configure an AS number.

The AS path attribute is modified when a route is advertised to an EBGP peer. Each time a route is advertised to an EBGP peer, the local routing device prepends its AS number to the existing path attribute, and a value of 1 is added to the AS number.

In Junos OS Release 9.1 and later, the numeric range is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <*16-bit high-order value in decimal*>.<*16-bit low-order value in decimal*>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options

autonomous-system—AS number. Use a number assigned to you by the NIC.

- **Range:** 1 through 4,294,967,295 (2^{32} – 1) in plain-number format for 4-byte AS numbers

In this example, the 4-byte AS number 65,546 is represented in plain-number format:

```
[edit]
routing-options {
    autonomous-system 65546;
}
```

- **Range:** 0.0 through 65535.65535 in AS-dot notation format for 4-byte numbers

In this example, 1.10 is the AS-dot notation format for 65,546:

```
[edit]
routing-options {
    autonomous-system 1.10;
}
```

- **Range:** 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

In this example, the 2-byte AS number 60,000 is represented in plain-number format:

```
[edit]
routing-options {
    autonomous-system 60000;
}
```

asdot-notation—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

- **Default:** Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

loops ***number***—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

- **Range:** 1 through 10

- **Default:** 1

> **NOTE**: When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the `loops` statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.
>
> Use the **independent-domain** option if the `loops` statement must be enabled only on a subset of routing instances.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

asdot-notation option introduced in Junos OS Release 9.3.

asdot-notation option introduced in Junos OS Release 9.3 for EX Series switches.

no-attrset option introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Examples: Configuring External BGP Peering
- Examples: Configuring Internal BGP Peering

autonomous-system-type (Flow Monitoring)

IN THIS SECTION

- Syntax | 2768

- Hierarchy Level | 2768
- Description | 2768
- Default | 2768
- Options | 2768
- Required Privilege Level | 2769
- Release Information | 2769

Syntax

```
autonomous-system-type (origin | peer);
```

Hierarchy Level

```
[edit forwarding-options accountingoutput cflowd hostname],
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-
server hostname],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server hostname]
```

Description

Specify the type of AS numbers that cflowd exports.

Default

`origin`

Options

`origin`—Export origin AS numbers of the packet source address in the Source Autonomous System cflowd field.

`peer`—Export peer AS numbers through which the packet passed in the Source Autonomous System cflowd field.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Flow Aggregation

Directing Traffic Sampling Output to a Server Running the cflowd Application

autostate-exclude

IN THIS SECTION

- Syntax | 2769
- Hierarchy Level | 2769
- Description | 2770
- Required Privilege Level | 2770
- Release Information | 2770

Syntax

```
autostate-exclude;
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
```

Description

Specify not to include an IRB interface in the state calculation for VLAN members. The default behavior is not to exclude an IRB interface in the state calculation unless all the ports on the interface go down. Because an IRB interface often has multiple ports in a single VLAN, the state calculation for a VLAN member might include a port that is down, possibly resulting in traffic loss. This feature enables you to exclude a trunk or access interface from the state calculation, which results in the IRB interface being marked as down as soon as the port specifically assigned to a VLAN goes down.

IRB interfaces are used to bind specific VLANs to Layer 3 interfaces, enabling a switch to forward packets between those VLANs— without having to configure another device, such as a router, to connect VLANs. In a typical scenario, a port on the interface is assigned to a specific VLAN, while a different port on that interface is assigned to an 802.1Q trunk interface to carry traffic between multiple VLANs, and a third port on that interface is assigned to an access interface used to connect the VLAN to network devices.

To ensure that an interface is marked as down and thereby excluded from the state calculation for VLAN members when the port assigned to the VLAN goes down, configure this statement on the trunk or access interface. The trunk or port interface is automatically excluded from the state calculation of the IRB interface. In this way, when a port assigned to a specified VLAN goes down, the IRB interface assigned to that VLAN is also marked as down.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 14.1x53-D40 and Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Excluding an IRB Interface from State Calculations on a QFX Series Switch

port-mode | **12841**

show ethernet-switching interface

auxiliary

IN THIS SECTION

- Syntax | 2771
- Hierarchy Level | 2771
- Description | 2771
- Default | 2771
- Options | 2772
- Required Privilege Level | 2772
- Release Information | 2772

Syntax

```
auxiliary {
    disable;
    insecure;
    type terminal-type;
    port-type (mini-usb | rj45);
}
```

Hierarchy Level

```
[edit system ports]
```

Description

Configure the characteristics of the auxiliary port.

Default

The auxiliary port is disabled. `disable` is the default option.

Options

disable Disable the port.

insecure Disable super user access or root logins to establish terminal connection.

type *terminal-type* Type of terminal that is connected to the port.

- **Range:** `ansi`, `vt100`, `small-xterm`, `xterm`
- **Default:** The terminal type is unknown, and the user is prompted for the terminal type.

port-type Set the RJ-45 console port or the Mini-USB console port as the active console port.

- mini-usb—Set the Mini USB type-B console port as the active console port.
- rj45—Set the RJ-45 console port as the active console port.
- **Default:** The RJ-45 console port is the active port.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`port-type` introduced in Junos OS Release 11.3 for EX Series switches.

RELATED DOCUMENTATION

- Configuring Junos OS to Set Console and Auxiliary Port Properties
- Configuring Console and Auxiliary Port Properties
- Configuring the Console Port Type (CLI Procedure)

auxiliary-spi

IN THIS SECTION

- Syntax | 2773
- Hierarchy Level | 2773
- Description | 2773
- Options | 2773
- Required Privilege Level | 2773
- Release Information | 2774

Syntax

```
auxiliary-spi spi-value;
```

Hierarchy Level

```
[edit services (IPsec VPN) ipsec-vpn rule rule-name term term-name then manual direction direction]
```

Description

Configure an auxiliary Security Parameter Index (SPI) for a manual SA. Use the auxiliary SPI when you configure the **protocol** statement to use the **bundle** option.

Options

spi-value—An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).

- **Range:** 256 through 16,639

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Security Associations

auxiliary-spi (Security IPsec)

IN THIS SECTION

- Syntax | 2774
- Hierarchy Level | 2774
- Description | 2775
- Options | 2775
- Required Privilege Level | 2775
- Release Information | 2775

Syntax

```
auxiliary-spi auxiliary-spi-value;
```

Hierarchy Level

```
[edit security ipsec security-association sa-name manual direction (inbound | outbound | bi-
directional)]
```

Description

Configure the auxiliary Security Parameter Index (SPI) for a manual SA. Use the auxiliary SPI when you configure the `protocol` statement to use the `bundle` option.

Options

`auxiliary-spi-value`—Arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).

- **Range:** `256 through 16,639`

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Manual IPsec Security Associations for an ES PIC

spi (Junos OS) | **28387**

avg-fd-twoway-threshold

IN THIS SECTION

- Syntax | **2776**
- Hierarchy Level | **2776**
- Description | **2776**
- Options | **2776**
- Required Privilege Level | **2777**

Release Information | 2777

Syntax

```
avg-fd-twoway-threshold avg-fd-twoway-threshold-value{
flap-trap-monitor seconds;}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
profile-name]
```

Description

Configure the threshold value for average frame delay, in microseconds, for two-way Ethernet frame delay measurement (ETH-DM). When the configured threshold for average frame delay is exceeded, an SNMP trap is generated for ETH-DM. SNMP traps are triggered if you configure either the average frame-delay threshold or the average inter-frame delay variance threshold. If you do not configure either the frame-delay threshold or the frame delay variation threshold, no SNMP traps are generated. You can configure these threshold values only with a two-way ETH-DM SLA iterator.

Frame delay refers to the difference, in microseconds, between the time a frame is sent and when it is received. Frame delay variation refers to the difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.” For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet frame delay measurement, only the initiator MEP (on the local system) collects statistics.

Options

avg-fd-twoway-threshold-value—Threshold value for average frame delay, in microseconds, for two-way ETH-DM.

- **Range:** 1 through 4294967295 microseconds

`flap-trap-monitor` *seconds*—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

- **Range:** 1 through 360 seconds

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

flap-trap-monitor | 7215

Configuring an Iterator Profile

Configuring an Iterator Profile on a Switch (CLI Procedure)

avg-flr-backward-threshold

IN THIS SECTION

- Syntax | 2777
- Hierarchy Level | 2778
- Description | 2778
- Options | 2778
- Required Privilege Level | 2778
- Release Information | 2778

Syntax

```
avg-flr-backward-threshold avg-flr-backward-threshold-value;{
flap-trap-monitor seconds; }
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
profile-name]
```

Description

Configure the threshold value for average frame loss ratio, in milli-percent, in the backward or downstream direction for Ethernet loss measurement (ETH-LM) and Ethernet synthetic loss measurement (ETH-SLM). When the configured threshold for average backward frame loss ratio is exceeded, an SNMP trap is generated for ETH-LM and ETH-SLM. SNMP traps are triggered if you configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold. If you do not configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold, no SNMP traps are generated. You can configure these threshold values with an SLA iterator for ETH-SLM and ETH-LM.

ETH-SLM is an application that enables the calculation of frame loss by using synthetic frames instead of data traffic. This mechanism can be considered as a statistical sample to approximate the frame loss ratio of data traffic.

Options

avg-flr-backward-threshold-value—Threshold value for average frame loss ratio in the backward or downstream direction, in milli-percent, for ETH-SLM and ETH-LM.

- **Range:** 1 through 100000 milli-percent

`flap-trap-monitor` *seconds*—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

- **Range:** 1 through 360 seconds

Required Privilege Level

configure—To enter configuration mode.

control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

flap-trap-monitor | 7215

Configuring an Iterator Profile

Configuring an Iterator Profile on a Switch (CLI Procedure)

avg-flr-forward-threshold

IN THIS SECTION

- Syntax | 2779
- Hierarchy Level | 2779
- Description | 2779
- Options | 2780
- Required Privilege Level | 2780
- Release Information | 2780

Syntax

```
avg-flr-forward-threshold avg-flr-forward-threshold-value;{
flap-trap-monitor seconds; }
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
profile-name]
```

Description

Configure the threshold value for average frame loss ratio, in milli-percent, in the upstream or forward direction for Ethernet loss measurement (ETH-LM) and Ethernet synthetic loss measurement (ETH-SLM). When the configured threshold for average forward frame loss ratio is exceeded, an SNMP trap is

generated for ETH-LM and ETH-SLM. SNMP traps are triggered if you configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold. If you do not configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold, no SNMP traps are generated. You can configure these threshold values with an SLA iterator for ETH-SLM and ETH-LM.

ETH-SLM is an application that enables the calculation of frame loss by using synthetic frames instead of data traffic. This mechanism can be considered as a statistical sample to approximate the frame loss ratio of data traffic.

Options

avg-flr-forward-threshold-value—Threshold value for average frame loss ratio in the forward or upstream direction, in milli-percent, for ETH-SLM and ETH-LM.

- **Range:** 1 through 100000 milli-percent

flap-trap-monitor *seconds*—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

- **Range:** 1 through 360 seconds

Required Privilege Level

configure—To enter configuration mode.

control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

flap-trap-monitor | 7215

Configuring an Iterator Profile

Configuring an Iterator Profile on a Switch (CLI Procedure)

avg-ifdv-twoway-threshold

IN THIS SECTION

- Syntax | 2781
- Hierarchy Level | 2781
- Description | 2781
- Options | 2782
- Required Privilege Level | 2782
- Release Information | 2782

Syntax

```
avg-ifdv-twoway-threshold avg-ifdv-twoway-threshold-value; {
flap-trap-monitor seconds; }
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
profile-name]
```

Description

Configure the threshold value for average frame delay variation, in microseconds, for two-way Ethernet frame delay measurement (ETH-DM). When the configured threshold for average frame delay variation is exceeded, an SNMP trap is generated for ETH-DM. SNMP traps are triggered if you configure either the average frame-delay threshold or the average inter-frame delay variance threshold. If you do not configure either the frame-delay threshold or the frame delay variation threshold, no SNMP traps are generated. You can configure these threshold values only with a two-way ETH-DM SLA iterator.

Frame delay refers to the difference, in microseconds, between the time a frame is sent and when it is received. Frame delay variation refers to the difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.” For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet

frame delay measurement, only the initiator MEP (on the local system) collects statistics. In two-way ETH-DM mode, frame delay and frame delay variation values are based on the time difference between when the initiator MEP transmits a request frame and receives a reply frame from the responder MEP, subtracting the time elapsed at the responder MEP.

Options

avg-ifdv-twoway-threshold-value—Threshold value for average frame delay variation, in microseconds, for two-way ETH-DM.

- **Range:** 1 through 4294967295 microseconds

flap-trap-monitor *seconds*—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

- **Range:** 1 through 360 seconds

Required Privilege Level

configure—To enter configuration mode.

control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

flap-trap-monitor | **7215**

Configuring an Iterator Profile

Configuring an Iterator Profile on a Switch (CLI Procedure)

avira-engine

IN THIS SECTION

- Syntax | 2783
- Hierarchy Level | 2783
- Description | 2784
- Required Privilege Level | 2784
- Release Information | 2784

Syntax

```
avira-engine {
    pattern-update {
        email-notify {
            admin-email admin-email;
            custom-message custom-message;
            custom-message-subject custom-message-subject;
        }
        interval interval;
        no-autoupdate;
        proxy-profile proxy-profile;
        routing-instance routing-instance;
        start-time start-time;
        url url;
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration anti-virus]
```

Description

The Avira scan engine is a licensed feature provided as a downloadable module. Download and install the Avira scan engine either through SRX with Internet connectivity to Juniper hosted URL, user hosted URL, or manually.

The Antivirus engine provides a full file-based virus scanning function which is available through a licensed subscription service. When your antivirus license key expires, you can continue to use the locally stored antivirus signatures without any updates. In case, the local database is deleted, antivirus scanning is also disabled.

Required Privilege Level

security— To view this statement in the configuration.

security-control— To add this statement to the configuration.

Release Information

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Statement introduced in Junos OS Release 18.4R1.

avp (L2TP Tunnel Switching)

IN THIS SECTION

- Syntax | 2785
- Hierarchy Level | 2785
- Description | 2785
- Required Privilege Level | 2785
- Release Information | 2785

Syntax

```
avp {
    bearer-type;
    calling-number;
    cisco-nas-port-info;
}
```

Hierarchy Level

```
[edit access tunnel-switch-profile profile-name]
```

Description

Specify the action taken on L2TP AVPs that are negotiated when the first session is created; these AVPs are contained in the L2TP packets that are switched by the tunnel switch profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

B

IN THIS SECTION

- backoff-period | 2791
- backup-bridge-priority | 2793
- backup-destination | 2795
- backup-interface | 2796
- backup-interface (Layer 2 Circuits) | 2797
- backup-liveness-detection | 2799
- backup (MBGP MVPN) | 2800
- backup-neighbor | 2802
- backup (OCS Partition) | 2804
- backup-on-failure (Accounting Options) | 2806
- backup-options | 2807
- backup-pe-group | 2808
- backup-peer-ip | 2810
- backup-remote-gateway | 2811
- backup-router | 2812
- backup-selection (Protocols IS-IS) | 2814
- backup-selection (Protocols OSPF) | 2819
- backup-spf-options (Protocols IS-IS) | 2821
- backup-spf-options (Protocols OSPF) | 2823
- backups | 2826
- backward-frr-enable | 2827
- bad-inner-header | 2829
- bad-option (IDS Screen Next Gen Services) | 2830
- bandwidth (multicast) | 2831
- bandwidth (storm control) | 2833
- bandwidth-based-metrics (Protocols OSPF) | 2835
- bandwidth-based-metrics | 2837
- bandwidth (DDoS) | 2839

- bandwidth-degradation | **2841**
- bandwidth (Fast Reroute, Signaled, and Multiclass LSPs) | **2843**
- bandwidth (Inline Services) | **2845**
- bandwidth (Interfaces) | **2847**
- bandwidth (Interfaces) | **2849**
- bandwidth-kbps | **2850**
- bandwidth-level | **2851**
- bandwidth-limit | **2853**
- bandwidth-limit (Hierarchical Policer) | **2855**
- bandwidth-limit (Policer) | **2857**
- bandwidth-limit (Policer for Gigabit Ethernet Interfaces) | **2860**
- bandwidth-model | **2861**
- bandwidth (MPC Bandwidth) | **2863**
- bandwidth-percent (protocols mpls) | **2864**
- bandwidth-percent | **2866**
- bandwidth-percentage | **2869**
- bandwidth (Protocols Layer 2 Circuit) | **2870**
- bandwidth (Protocols RSVP) | **2872**
- bandwidth (RFC 2544 Benchmarking) | **2874**
- bandwidth-scale (DDoS) | **2875**
- bandwidth (Static LSP) | **2876**
- bandwidth-threshold | **2878**
- bandwidth (Tunnel Services) | **2880**
- banner (Access FTP HTTP Telnet Authentication) | **2882**
- banner (Access Web Authentication) | **2884**
- base-distinguished-name | **2885**
- batch query | **2886**
- bb-sc-n | **2889**
- beacon-period | **2890**
- bearer-type (L2TP Tunnel Switching) | **2892**
- ber-threshold-clear | **2894**
- ber-threshold-signal-degrade | **2897**

- bert-algorithm | **2900**
- bert-error-rate | **2903**
- bert-period | **2905**
- best-site | **2908**
- bfd | **2909**
- bfd (protocols) | **2911**
- bfd-liveliness | **2914**
- bfd-liveliness (Activeness Probe) | **2917**
- bfd-liveness-detection (BGP) | **2918**
- bfd-liveness-detection (EVPN Node Detection) | **2923**
- bfd-liveness-detection (LAG) | **2926**
- bfd-liveness-detection (Layer 2 VPN and VPLS) | **2929**
- bfd-liveness-detection (LSP) | **2931**
- bfd-liveness-detection (Protocols IS-IS) | **2933**
- bfd-liveness-detection (Protocols LDP) | **2937**
- bfd-liveness-detection (Protocols MPLS) | **2939**
- bfd-liveness-detection (Protocols OSPF) | **2941**
- bfd-liveness-detection (Protocols PIM) | **2945**
- bfd-liveness-detection (Protocols RIP) | **2947**
- bfd-liveness-detection (Routing Options Static Route) | **2950**
- bfd-liveness-detection (Source Packet Routing Template) | **2955**
- bgp | **2956**
- bgp (RPM) | **2958**
- bgp-error-tolerance (Protocols BGP) | **2959**
- bgp-identifier | **2962**
- bgp-orf-cisco-mode | **2963**
- bgp-peer | **2965**
- bgp-static | **2966**
- bgpmcast | **2968**
- bidirectional (Interface) | **2969**
- bidirectional (RP) | **2971**
- bier | **2972**

- bier-sub-domain (Protocols IS-IS) | **2974**
- bits (Chassis Synchronization) | **2975**
- black-list | **2978**
- blackhole-action | **2979**
- block | **2980**
- block-allocation (Source NAT Next Gen Services) | **2983**
- block-command | **2985**
- block-content-type | **2986**
- block-extension | **2987**
- block-frag (IDS Screen Next Gen Services) | **2988**
- block-message (Application Firewall) | **2989**
- block-message (Security Content Security) | **2993**
- block-mime | **2994**
- block (Spanning Trees) | **2995**
- blocklists (Security IKE) | **2997**
- bmp | **3000**
- boot-loader-authentication | **3008**
- boot-server (DHCP) | **3010**
- bootp | **3012**
- bootp (System Processes) | **3013**
- bootstrap | **3015**
- bootstrap-export | **3016**
- bootstrap-import | **3018**
- bootstrap-priority | **3019**
- bounded-delay | **3021**
- bpdu-block | **3022**
- bpdu-block-on-edge | **3024**
- bpdu-destination-mac-address | **3026**
- bpdu-destination-mac-address (Spanning Tree) | **3027**
- bridge-domain | **3029**
- bridge-domain (Analyzer) | **3031**
- bridge-domains | **3032**

- bridge-options | 3034
- bridge-priority | 3036
- bridge-template | 3038
- brief | 3039
- broadcast (Interfaces) | 3041
- broadcast (Router Discovery) | 3043
- broadcast (Forwarding class) | 3044
- bucket-size | 3046
- buffer-dynamic-threshold | 3047
- buffer-partition (Egress) | 3049
- buffer-partition (Ingress) | 3052
- buffer-size | 3055
- buffer-size (Dynamic Scheduling) | 3061
- buffer-size (Schedulers) | 3063
- buffer-rate (Schedulers) | 3065
- bulk-call-mode | 3067
- bulk-leasequery (DHCP Relay Agent) | 3068
- bum-hashing | 3070
- bundle | 3072
- bundle (multilink) | 3073
- bundle (Interfaces) | 3074
- burst (DDoS) | 3075
- burst-scale (DDoS) | 3077
- burst-size | 3078
- burst-size (Default Local Policy) | 3081
- burst-size-limit | 3083
- burst-size-limit (Hierarchical Policer) | 3084
- burst-size-limit (Policer) | 3086
- burst-size-limit (Policer for Gigabit Ethernet Interfaces) | 3089
- burst-size (TDF Domain) | 3090
- by-destination (IDS MS-DPC) | 3092
- by-destination (IDS MS-MPC) | 3093

- by-destination (IDS Screen Next Gen Services) | **3095**
- by-pair (IDS MS-DPC) | **3098**
- by-protocol (IDS MS-MPC) | **3099**
- by-protocol (IDS Screen Next Gen Services) | **3102**
- by-source (IDS MS-DPC) | **3105**
- by-source (IDS MS-MPC) | **3106**
- by-source (IDS Screen Next Gen Services) | **3108**
- bypass | **3111**
- bypass-aggregate (DDoS) | **3112**
- bypass-queuing-chip | **3114**
- bypass (Signaled LSP) | **3115**
- bypass (Static LSP) | **3117**
- bypass-traffic-on-exceeding-flow-limits | **3119**
- bypass-traffic-on-pic-failure | **3120**
- bytes (Dynamic Traffic Shaping) | **3121**
- bytes (otn-options) | **3123**

backoff-period

IN THIS SECTION

- Syntax | **2792**
- Hierarchy Level | **2792**
- Description | **2792**
- Options | **2792**
- Required Privilege Level | **2793**
- Release Information | **2793**

Syntax

```
backoff-period milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name
bidirectional df-election],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name bidirectional df-election],
[edit protocols pim interface (Protocols PIM)  interface-name bidirectional df-election],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name bidirectional df-election]
```

Description

Configure the designated forwarder (DF) election backoff period for bidirectional PIM. The **backoff-period** statement configures the period that the acting DF waits between receiving a better DF Offer and sending the Pass message to transfer DF responsibility.

> **NOTE**: Junos OS checks rendezvous point (RP) unicast reachability before accepting incoming DF messages. DF messages for unreachable rendezvous points are ignored. This is needed to prevent the following example scenario. Routers A and B are downstream routers on the same LAN, and both are supposed to send DF election messages with an infinite metric on their upstream interfaces (reverse-path forwarding [RPF] interfaces). Router A has a higher IP address than Router B. When both routers lose the path to the RP, both send an Offer message with the infinite metric onto the LAN. Router A wins the election because it has a higher IP address, and Router B backs off as a result. After three Offer messages, according to RFC 5015, Router A looks up the RP and finds no path to the RP. As a result, Router A transitions to the Lose state and sends nothing. On the other hand, after backing off for an interval of 3 x the Offer period, Router B does not receive any messages, and resumes the DF election by sending a new Offer message. Hence, the pattern repeats indefinitely.

Options

milliseconds—Period that the acting DF waits between receiving a better DF Offer and sending the Pass message to transfer DF responsibility.

- **Range:** 100 through 65,535 milliseconds
- **Default:** 1000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Bidirectional PIM

Example: Configuring Bidirectional PIM

backup-bridge-priority

IN THIS SECTION

- Syntax | 2793
- Hierarchy Level | 2794
- Description | 2794
- Options | 2794
- Required Privilege Level | 2794
- Release Information | 2794

Syntax

```
backup-bridge-priority priority;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp)],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id],
[edit protocols (mstp | rstp)],
[edit protocols vstp vlan vlan-id],
[edit routing-instances routing-instance-name protocols (mstp | rstp)],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id]
```

Description

Determine the priority of the backup bridge in a VPLS multihomed Layer 2 ring with MPLS infrastructure.

Options

priority—The backup bridge priority can be set only in increments of 4096.

- **Range:** 0 through 61,440
- **Default:** 32,768

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Understanding VPLS Multihoming

Configuring VPLS Root Protection Topology Change Actions to Control Individual VLAN Spanning-Tree Behavior

backup-destination

IN THIS SECTION

- Syntax | 2795
- Hierarchy Level | 2795
- Description | 2795
- Options | 2795
- Required Privilege Level | 2796
- Release Information | 2796

Syntax

```
backup-destination destination-address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number tunnel], [edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number tunnel]
```

Description

For tunnel interfaces, specify the remote address of the backup tunnel.

Options

destination-address—Address of the remote side of the connection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

destination (Interfaces) | 5433

destination (Tunnel Remote End) | 5526

Configuring IPsec Tunnel Redundancy

backup-interface

IN THIS SECTION

- Syntax | 2796
- Hierarchy Level | 2797
- Description | 2797
- Options | 2797
- Required Privilege Level | 2797
- Release Information | 2797

Syntax

```
backup-interface interface-name;
```

Hierarchy Level

```
[edit interfaces interface-name es-options]
```

Description

Configure a backup ES Physical Interface Card (PIC). When the primary ES PIC has a servicing failure, the backup becomes active, inherits all the tunnels and security associations (SAs), and acts as the new next hop for IPsec traffic.

Options

interface-name—Name of ES interface to serve as the backup.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring ES PIC Redundancy

backup-interface (Layer 2 Circuits)

IN THIS SECTION

- Syntax | 2798
- Hierarchy Level | 2798
- Description | 2798

Required Privilege Level | 2798
Release Information | 2798

Syntax

```
backup-interface interface-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name end-interface interface interface-name],
[edit protocols l2circuit local-switching interface interface-name end-interface interface
interface-name]
```

Description

Specify the interface to be used by the protection pseduowire in connection protection configurations for Layer 2 circuits.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring Layer 2 Circuit Switching Protection

backup-liveness-detection

IN THIS SECTION

- Syntax | 2799
- Hierarchy Level | 2799
- Description | 2799
- Required Privilege Level | 2800
- Release Information | 2800

Syntax

```
backup-liveness-detection {
    backup-peer-ip ipv4-address;
}
```

Hierarchy Level

```
[edit protocols iccp peer]
```

Description

Determine whether a peer is up or down by exchanging keepalive messages over the management link between the two Inter-Chassis Control Protocol (ICCP) peers.

When an ICCP connection is operationally down, the status of the peers hosting a multichassis link aggregation group (MC-LAG) is detected by sending liveness detection requests to each other. Peers must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, the liveness detection check fails, and a failure action is implemented. Backup liveness detection must be configured on both peers hosting the MC-LAG.

For more information on the ICCP failure scenarios and handling the failures, refer to ICCP Failure Scenarios for EX9200 Switches.

The remaining statement is explained separately. See CLI Explorer.

> **NOTE**: If backup liveness detection is configured, the peer status is always up when either the ICCP TCP Connection is established, or Bidirectional Forwarding Protocol (BFD) is configured and the peer is up. The backup liveness check is only performed when the ICCP connection is down.

> **NOTE**: When an MC-LAG node *(n0)* is rebooted with ICCP disabled on it, the backup liveness of the other node *(n1)* remains active even though the remote link goes down. This is because the status of the last action performed on the Backup Liveness process is shown in the CLI, the Backup Liveness process is not active on the other node *(n1)*, and therefore, the node *(n1)* is unable to read the current status of the first node *(n0)*. However, there is no traffic loss. The Backup Liveness process is not active all the time; it runs only when initiated.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- Configuring Multichassis Link Aggregation on MX Series Routers

backup (MBGP MVPN)

IN THIS SECTION

- Syntax | 2801
- Hierarchy Level | 2801

- Description | 2801
- Options | 2801
- Required Privilege Level | 2801
- Release Information | 2801

Syntax

```
backup address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvpn
static-umh],
[edit routing-instances routing-instance-name protocols mvpn static-umh]
```

Description

Define a backup upstream multicast hop (UMH) for type 7 (S,G) routes.

If the primary UMH is unavailable, the backup is used. If neither UMH is available, no UMH is selected.

Options

address Address of the backup UMH.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

Example: Configuring Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

sender-based-rpf (MBGP MVPN) | **16552**

static-umh (MBGP MVPN) | **28585**

unicast-umh-election | **30353**

backup-neighbor

IN THIS SECTION

- Syntax | **2802**
- Hierarchy Level | **2803**
- Description | **2803**
- Options | **2803**
- Required Privilege Level | **2803**
- Release Information | **2803**

Syntax

```
backup-neighbor address {
    community name;
    mtu number;
    hot-standby;
    psn-tunnel-endpoint address;
    standby;
    static;
    virtual-circuit-id number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor address],
[edit protocols l2circuit local-switching interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name],
[edit routing-instances routing-instance-name protocols vpls neighbor address]
```

Description

Configure pseudowire redundancy for Layer 2 circuits and VPLS. A redundant pseudowire can act as a backup connection and can be configured between a PE router and a CE device or between PE routers, maintaining Layer 2 circuit and VPLS services after certain types of failures. This feature can help improve the reliability of certain types of networks where a single point of failure could interrupt service for customers.

NOTE: The `psn-tunnel-endpoint` statement is not supported at the `[edit protocols l2circuit local-switching interface interface-name end-interface interface interface-name]` hierarchy level.

Options

address—Specifies the address for the backup neighbor.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Pseudowire Redundancy on the PE Router

Example: Configuring Layer 2 Circuit Switching Protection

community (Protocols Layer 2 Circuit) | **4469**

psn-tunnel-endpoint | **13851**

virtual-circuit-id | **30749**

backup (OCS Partition)

IN THIS SECTION

- Syntax | **2804**
- Hierarchy Level | **2804**
- Description | **2805**
- Options | **2805**
- Required Privilege Level | **2805**
- Release Information | **2805**

Syntax

```
backup {
    limit;
    timeout seconds;
    overflow (deny-login | drop-oldest);
}
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Provides file backup for OCS data when both primary and alternative paths to the OCS are down.

Options

limit	The limit on the total number of backup entries. The login fails if the backup entries exceeds the prescribed limit.
timeout *seconds*	The timeout for backup operation.
overflow (deny-login | drop-oldest)	Controls action on the number of backup entries over limit. The values deny-login denies andy logins until the number of backup entries drops under backup-limit), and the value discard-oldest discards the oldest backup entries.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1.

RELATED DOCUMENTATION

Configuring the OCS Partition

Gy File Backup Overview

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

backup-on-failure (Accounting Options)

IN THIS SECTION

- Syntax | 2806
- Hierarchy Level | 2806
- Description | 2806
- Options | 2806
- Required Privilege Level | 2807
- Release Information | 2807

Syntax

```
backup-on-failure (master-and-slave | master-only);
```

Hierarchy Level

[edit accounting-options file *filename*]

Description

Configure the router to save a copy of the accounting file locally, to the **/var/log/pfedBackup** directory of the relevant Routing Engine, in the event that file transfer to the remote archive sites cannot be completed.

Options

master-and-slave Back up accounting files from both the primary Routing Engine and the backup Routing Engine.

master-only Back up accounting files from only the primary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- Configuring Accounting-Data Log Files
- *Configuring Flat-File Accounting for Layer 2 Wholesale*
- *Configuring Flat-File Accounting for Extensible Subscriber Services Management*
- *Flat-File Accounting Overview*

backup-options

IN THIS SECTION

- Syntax | 2807
- Hierarchy Level | 2808
- Description | 2808
- Required Privilege Level | 2808
- Release Information | 2808

Syntax

```
backup-options {
    interface interface-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure an interface to be used as a backup interface if the primary interface goes down. This is used to support ISDN dial backup operation.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

backup-pe-group

IN THIS SECTION

- Syntax | 2809
- Hierarchy Level | 2809
- Description | 2809
- Options | 2809
- Required Privilege Level | 2809
- Release Information | 2809

Syntax

```
backup-pe-group group-name {
    backups [ addresses ];
    local-address address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Description

Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options

group-name—Name of the group for PE backups.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Ingress PE Redundancy

backup-peer-ip

IN THIS SECTION

- Syntax | 2810
- Hierarchy Level | 2810
- Description | 2810
- Required Privilege Level | 2810
- Release Information | 2811

Syntax

```
backup-peer-ip ipv4-address;
```

Hierarchy Level

```
[edit protocols iccp peer backup-liveness-detection]
```

Description

Specify the IP address of the peer being used as a backup peer in the Bidirectional Forwarding Detection (BFD) configuration. You can specify any IP address in the default routing-instance. We recommend using the management IP addresses associated with `fxp0` or `em0` for redundancy in the event of a network disruption.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

backup-remote-gateway

IN THIS SECTION

- Syntax | 2811
- Hierarchy Level | 2811
- Description | 2811
- Options | 2812
- Required Privilege Level | 2812
- Release Information | 2812

Syntax

```
backup-remote-gateway address;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Define the backup remote address to which the IPsec traffic is directed when the primary remote gateway is down. Configuring this statement also enables the dead peer detection (DPD) protocol.

Options

address—Backup remote IPv4 or IPv6 address.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

backup-router

IN THIS SECTION

- Syntax | 2812
- Hierarchy Level | 2813
- Description | 2813
- Options | 2813
- Required Privilege Level | 2813
- Release Information | 2813

Syntax

```
backup-router address <destination destination-address>;
```

Hierarchy Level

```
[edit system]
```

Description

Set a default router (running IP version 4 [IPv4]) to use while the local router (running IPv4) is booting and if the routing protocol processes fail to start. The Junos OS removes the route to this router as soon as the software starts.

Options

address	Address of the default router.
destination ***destination-address***	(Optional) Destination address that is reachable through the backup router. You can include this option to achieve network reachability while loading, configuring, and recovering the router, but without the risk of installing a default route in the forwarding table.

- **Default:** All hosts (default route) are reachable through the backup router.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Backup Router

backup-selection (Protocols IS-IS)

IN THIS SECTION

- Syntax | 2814
- Hierarchy Level | 2815
- Description | 2815
- Options | 2815
- Required Privilege Level | 2818
- Release Information | 2818

Syntax

```
backup-selection {
    destination prefix {
        interface (interface-name| all){
            admin-group {
                exclude  [ group-name ];
                include-all  [ group-name ];
                include-any  [ group-name ];
                preference  [ group-name ];
            }
            bandwidth-greater-equal-primary;
            dest-metric (highest | lowest);
            downstream-paths-only;
            metric-order [ root dest ];
            node {
                exclude [ node-address ];
                preference [ node-address ];
            }
            node-tag {
                exclude  [ route-tag ];
                preference  [ route-tag ];
            }
            protection-type (link | node | node-link);
            root-metric (highest | lowest);
            srlg (loose | strict);
```

```
            evaluation-order [ admin-group srlg bandwidth protection-type node node-tag
metric ] ;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-instances instance-name routing-options],
[edit routing-instances instance-name routing-options],
[edit routing-options]
```

Description

Define backup selection policies, per prefix per primary next-hop interface, to enforce loop-free alternate (LFA) selection based on admin-group, srlg, bandwidth, protection-type, node, node-tag, and metric attributes of the backup path.

The remaining statements are explained separately. See CLI Explorer.

Options

destination *prefix*
Define the backup selection policy for a particular destination prefix or for all the prefixes. The *prefix* is the destination prefix name and prefix length. You can specify 0/0 for the IPv4 least-specific prefix or 0::0/0 for the IPv6 least-specific prefix.

interface
Define the backup selection policy for a specific primary next hop.

interface-name
Name of the primary next-hop interface.

all
All the interfaces.

bandwidth-greater-equal-primary
Allow the selection of the backup next hop only if the bandwidth is greater than or equal to the bandwidth of the primary next hop.

dest-metric (highest | lowest)
Specifiy the metric from the one-hop neighbor or from the remote router such as an RSVP backup label-switched-path (LSP) tail-end router to the final destination.

- **highest** Select the backup path that has the highest destination metric.

lowest Select the backup path that has the lowest destination metric.

downstream-paths-only Select the backup path that is a downstream path to the destination.

evaluation-order [admin-group srlg bandwidth protection-type neighbor neighbor-tag metric] Control the order and the criteria of evaluating the backup path. The default order of evaluation is admin-group, srlg, bandwidth, protection-type, neighbor, neighbor-tag, and metric.

> **NOTE**: For the explicitly configured evaluation order, only the listed attributes influence the selection of the backup path.

metric-order [root dest] Specify the order of preference of the root and the destination metric during the backup path selection. The preference order can be:

- `[root dest]` — Backup path selection or preference is first based on the root-metric criteria. If the criteria of all the root-metric is the same, then the selection or preference is based on the dest-metric.
- `[dest root]` — Backup path selection or preference is first based on the dest-metric criteria. If the criteria of all the dest-metric is the same, then the selection is based on the root-metric.

> **NOTE**: By default, backup paths with lower destination metric criteria are selected or preferred. If the criteria is the same, then the lowest root metric criteria is preferred or selected.

root The metric to a one-hop neighbor or a remote router.

dest The metric from a one-hop neighbor or remote router to the final destination.

protection-type (link | node | node-link) Specify the required protection type of the backup path.

> **NOTE**: If no protection-type is configured, then by default the first best path that matches all the other criteria is executed.

link Select the backup path that provides link protection.

node Select the backup path that provides node protection.

node-link Allow either node or link protection LFA where node-protection LFA is preferred over link-protection LFA.

root-metric (highest | lowest) Specify the metric to the one-hop neighbor or to the remote router such as an RSVP backup label-switched-path (LSP) tail-end router.

highest Select the highest root metric.

lowest Select the lowest root metric.

srlg (loose | strict) Define the backup selection to either allow or reject the common shared risk link groups (SRLGs) between the primary link and any link in the backup path.

loose Allow the backup path that has common srlgs between the primary link and any link in the backup path. A backup path with a fewer number of srlg collisions is preferred.

strict Reject the backup path that has common srlgs between the primary link and each link in the backup path.

node Define a list of loop-back IP addresses of the adjacent nodes to either prefer or exclude in the backup path selection. The node can be a local (adjacent router) node, remote node, or any other router in the backup path.

> **NOTE**: The nodes are identified through the TE-router-ID TLV advertised by a node in the LSP.

exclude [*node-address*] Specify the list of nodes to be excluded. The backup path that has a router from the list is not selected as the loop-free alternative or backup next hop.

node-address Name of one or more nodes to be excluded during backup path selection.

preference [*node-address*] Define an ordered set of nodes to be preferred. The backup path having the leftmost node is selected.

node-address Name of one or more nodes to be preferred in the backup path selection.

node-tag Define per-neighbor policy to either prefer or exclude a backup path.

> **NOTE**: This statement identifies a group of nodes in the network based on criteria such as the same neighbor tag values for all PE nodes. This is implemented using IS-IS admin-tags.

exclude [*node-address*]
: Specify the list of nodes to be excluded. The backup path that has a router from the list is not selected as the loop-free alternative or backup next hop.

  node-address
  : Name of one or more nodes to be excluded during backup path selection.

preference [*node-address*]
: Define an ordered set of nodes to be preferred. The backup path having the leftmost node is selected.

  node-address
  : Name of one or more nodes to be preferred in the backup path selection.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Backup Selection Policy for IS-IS Protocol

Example: Configuring Backup Selection Policy for IS-IS Protocol

backup-selection (Protocols OSPF)

IN THIS SECTION

- Syntax | 2819
- Hierarchy Level | 2820
- Description | 2820
- Options | 2820
- Required Privilege Level | 2821
- Release Information | 2821

Syntax

```
backup-selection {
    destination  prefix {
        interface (interface-name| all){
            admin-group {
                exclude  [ group-name ];
                include-all  [ group-name ];
                include-any  [ group-name ];
                preference  [ group-name ];
            }
            bandwidth-greater-equal-primary;
            dest-metric (highest | lowest);
            downstream-paths-only;
            metric-order [ root dest ];
            node {
                exclude [ node-address ];
                preference [ node-address ];
            }
            protection-type (link | node | node-link);
            root-metric (highest | lowest);
            srlg (loose | strict);
            evaluation-order [ admin-group srlg bandwidth protection-type node metric ] ;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-instances instance-name routing-options],
[edit routing-instances instance-name routing-options],
[edit routing-options]
```

Description

Define backup selection policies, per prefix per primary next-hop interface, to enforce loop-free alternate (LFA) selection based on admin-group, srlg, bandwidth, protection-type, node, and metric attributes of the backup path.

Options

destination *prefix*

Define the backup selection policy for a particular destination prefix or for all the prefixes. The value *prefix* defines the destination prefix name and prefix length. You can specify 0/0 for the IPv4 least-specific prefix or 0::0/0 for the IPv6 least-specific prefix.

node

Define a list of loop-back IP addresses of the adjacent nodes to either prefer or exclude in the backup path selection. The node can be a local (adjacent router) node, remote node, or any other router in the backup path.

> **NOTE**: The nodes are identified through the route-id advertised by a node in the LSP.

exclude [*node-address*]

Specify one or more nodes to be excluded. The backup path that has a router from the list is not selected as the loop-free alternative or backup next hop.

preference [*node-address*]

Define an ordered set of one or more nodes to be preferred. The backup path having the leftmost node is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Example: Configuring Backup Selection Policy for the OSPF or OSPF3 Protocol

Configuring Backup Selection Policy for the OSPF Protocol

Understanding Backup Selection Policy for OSPF Protocol

backup-spf-options (Protocols IS-IS)

IN THIS SECTION

- Syntax | 2821
- Hierarchy Level | 2822
- Description | 2822
- Options | 2822
- Required Privilege Level | 2823
- Release Information | 2823

Syntax

```
backup-spf options {
    node-link-degradation;
    per-prefix-calculation;
    remote-backup-calculation;
```

```
    use-source-packet-routing;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Configure options for running the shortest-path-first (SPF) algorithm for backup next hops for protected interfaces. Use these options to override the default behavior of having Junos OS calculate backup paths for all the topologies in an instance when at least one interface is configured with link protection or node-link protection.

This configuration option is used in conjunction with the `node-link-protection` statement to compute per-prefix loop-free alternate routes (LFAs).

Options

node-link-degradation Degrade an interface from node-link to link protection in case no node protection LFA route is found for a given destination node.

per-prefix-calculation Calculate backup next hops for non-best prefix originators.

remote-backup-calculation Determine the remote LFA backup paths from the point of local repair (PLR) in an IS-IS network. For every protected link on the PLR, Junos OS creates a dynamic LDP label-switched path to reach the remote LFA node. When the primary link fails, the PLR uses these remote LFA backup paths to reach all the destinations reachable through the primary-link.

use-source-packet-routing Enable use of source packet routing node segment labels for computing backup paths for normal IPv4 or IPv6 IS-IS prefixes and primary IS-IS source packet routing node segments.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4R3.

`node-link-degradation` option added in Junos OS Release 14.2.

`use-source-packet-routing` option introduced in Junos OS Release 16.2R1 for MX Series and PTX Series routers.
Support for this option introduced in Junos OS Release 17.2R1 for QFX5100 and QFX10000 switches.

Support for `remote-backup-calculation` option introduced in Junos OS Release 18.2R1 for QFX5100, QFX5110, and QFX5200 switches.

RELATED DOCUMENTATION

Example: Configuring MPLS Egress Protection for Layer 3 VPN Services

Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP

backup-spf-options (Protocols OSPF)

IN THIS SECTION

- Syntax | 2824
- Hierarchy Level | 2824
- Description | 2824
- Options | 2825
- Required Privilege Level | 2825
- Release Information | 2825

Syntax

```
backup-spf-options {
    disable;
    downstream-paths-only;
    no-install;
    node-link-degradation;
    per-prefix-calculation {
        all;
        externals;
        stubs;
        summary;
    }
    remote-backup-calculation;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf topology (default | name)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
topology (default | name)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf topology (default | name)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf topology (default | name)]
```

Description

Configure options for running the shortest-path-first (SPF) algorithm for backup next hops for protected interfaces. Use these options to override the default behavior of having Junos OS calculate backup paths for all the topologies in an instance when at least one interface is configured with link protection or node-link protection. These options also enable you to change the default behavior for a specific topology in an OSPF instance.

Options

disable Do not calculate backup next hops for the specified instance or topology.

downstream-paths-only Calculate and install only downstream paths as defined in RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates* for the specified instance or topology.

no-install Do not install the backup next hops for the specified instance or topology.

node-link-degradation Degrade an interface from node-link to link protection in case no node protection LFA route is found for a given destination node. A link protecting loop-free alternate (LFA) is used when node-link protecting LFA is not available in the topology for any of the protected links.

per-prefix-calculation Calculate backup next hops for non-best prefix originators.

- `all`—Calculate per-prefix loop free alternate (LFA) for all.
- `externals`—Calculate per-prefix LFA for not-so-stubby and externals only.
- `stubs`—Calculate per-prefix LFA for stubs only.
- `summary`—Calculate per-prefix LFA for summary originators only.

remote-backup-calculation Determine the remote LFA backup paths from the point of local repair (PLR) in an OSPF network. For every protected link on the PLR, Junos OS creates a dynamic LDP label-switched path to reach the remote LFA node. When the primary link fails, the PLR uses these remote LFA backup paths to reach all the destinations reachable through the primary-link.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

`node-link-degradation`, `per-prefix-calculation`, and `remote-backup-calculation` options introduced in Junos OS Release 15.1.

Support for `remote-backup-calculation` option introduced in Junos OS Release 18.2R1 for QFX5100, QFX5110, and QFX5200 switches.

RELATED DOCUMENTATION

Configuring Backup SPF Options for Protected OSPF Interfaces

backups

IN THIS SECTION

- Syntax | 2826
- Hierarchy Level | 2826
- Description | 2827
- Options | 2827
- Required Privilege Level | 2827
- Release Information | 2827

Syntax

```
backups [ addresses ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast backup-pe-group group-name],
[edit logical-systems logical-system-name routing-options multicast backup-pe-group group-name],
[edit routing-instances routing-instance-name routing-options multicast backup-pe-group group-
name],
[edit routing-options multicast backup-pe-group group-name]
```

Description

Configure the address of backup PEs for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options

addresses—Addresses of other PEs in the backup group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Ingress PE Redundancy

backward-frr-enable

IN THIS SECTION

- Syntax | 2828
- Hierarchy Level | 2828
- Description | 2828
- Default | 2828
- Options | 2828
- Required Privilege Level | 2829
- Release Information | 2829

Syntax

```
(backward-frr-enable | no-backward-frr-enable);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

Description

Enable or disable backward fast reroute status insertion.

Enable backward fast reroute to insert local pre-forward error correction (FEC) bit error rate (BER) status into transmitted OTN frames, notifying the remote interface. The remote interface can use the information to reroute traffic to a different interface. When you enable backward fast reroute and also enable pre-FEC BER monitoring including the `signal-degrade-monitor-enable` statement, notification of signal degradation and rerouting of traffic occurs in less time than that required through a Layer 3 protocol.

> **NOTE**: When you configure pre-FEC BER signal degrade monitoring, we recommend that you configure both the `signal-degrade-monitor-enable` and `backward-frr-enable` statements.

You can also configure the pre-FEC BER thresholds that raise or clear a signal degrade alarm and the time interval for the thresholds. If the BER thresholds and interval are not configured, the default values are used. Include the `ber-threshold-signal-degrade` *value*, `ber-threshold-clear` *value*, and `interval` *value* statements at the `[edit interfaces` *interface-name* `otn-options signal-degrade]` hierarchy level to configure the BER thresholds and time interval. See *Understanding Pre-FEC BER Monitoring and BER Thresholds* for more information about pre-FEC BER monitoring and determining BER threshold settings.

Default

By default, backward fast reroute insertion is disabled.

Options

`backward-frr-enable`—Enable backward fast reroute status insertion.

`no-backward-frr-enable`—Do not enable backward fast reroute status insertion.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding Pre-FEC BER Monitoring and BER Thresholds

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

bad-inner-header

IN THIS SECTION

- Syntax | 2829
- Hierarchy Level | 2830
- Description | 2830
- Required Privilege Level | 2830
- Release Information | 2830

Syntax

```
bad-inner-header;
```

Hierarchy Level

```
[edit security screen ids-option ids-option-name ip tunnel]
```

Description

Enable IP tunnel bad inner header IDS option.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D10.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

bad-option (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 2831
- Hierarchy Level | 2831
- Description | 2831
- Required Privilege Level | 2831
- Release Information | 2831

Syntax

```
bad-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop any packet with incorrectly formatted IPv4 options or IPv6 extension headers. Incorrectly formatted IPv4 options or IPv6 extension headers can cause unpredictable issues, depending on the IP stack implementation of routers and the target.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

bandwidth (multicast)

IN THIS SECTION

Syntax | 2832

- Hierarchy Level | 2832
- Description | 2832
- Options | 2832
- Required Privilege Level | 2832
- Release Information | 2833

Syntax

```
bandwidth ( bps | adaptive );
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast flow-map],
[edit logical-systems logical-system-name routing-options multicast flow-map],
[edit routing-instances routing-instance-name routing-options multicast flow-map],
[edit routing-options multicast flow-map]
```

Description

Configure the bandwidth property for multicast flow maps.

Options

adaptive—Specify that the bandwidth is measured for the flows that are matched by the flow map.

bps—Bandwidth, in bits per second, for the flow map.

- **Range:** 0 through any amount of bandwidth
- **Default:** 2 Mbps

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Example: Configuring a Multicast Flow Map

bandwidth (storm control)

IN THIS SECTION

- Syntax | 2833
- Hierarchy Level | 2834
- Description | 2834
- Default | 2834
- Options | 2834
- Required Privilege Level | 2835
- Release Information | 2835

Syntax

```
bandwidth bandwidth;
```

Hierarchy Level

```
[edit ethernet-switching-options storm-control interface (all |
interface-name)]
```

Description

Configure the storm control level as the bandwidth in kilobits per second of the applicable traffic streams, as follows:

- On EX2200, EX3200, EX3300, EX4200, and EX6200 switches—Applies to the combined broadcast and unknown unicast streams by default. Storm control does not apply to multicast traffic by default on these switches. If you enable storm control for multicast traffic on a specific interface, the configured bandwidth allocation applies to the combined broadcast, unknown unicast, and multicast traffic on that interface.
- On EX4500 and EX8200 switches—Applies to the combined broadcast, multicast, and unknown unicast streams.

> **NOTE**: When you configure storm control bandwidth on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control bandwidth of 15,000 Kbps on **ae1**, and **ae1** has two members, **ge-0/0/0** and **ge-0/0/1**, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on **ae1** allows a traffic rate of up to 30,000 Kbps of combined broadcast, multicast, and unknown unicast traffic.

Default

If you omit the `bandwidth` statement when you configure storm control on an interface, the storm control level defaults to 80 percent of the available bandwidth used by the combined applicable traffic streams. Depending upon the configuration, applicable traffic could include broadcast, unknown unicast, and multicast traffic.

Options

bandwidth—Traffic rate in kilobits per second of the combined applicable traffic streams.

- **Range:** 100 through 10,000,000
- **Default:** None

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

level (OAM) | 9716

Configuring Autorecovery for Port Security Events

Enabling and Disabling Storm Control (non-ELS)

bandwidth-based-metrics (Protocols OSPF)

IN THIS SECTION

- Syntax | 2835
- Hierarchy Level | 2836
- Description | 2836
- Options | 2836
- Required Privilege Level | 2837
- Release Information | 2837

Syntax

```
bandwidth-based-metrics {
    bandwidth value;
    metric number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface
interface-name],
[edit logical-systems logical-system-name protocols ospf area area-id interface interface-name
topology topology-name],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast) area area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-id interface interface-name],
[edit logical-systems logical-system-name protocols ospf area area-id interface interface-name
topology topology-name],
[edit logical-systems logical-system-name routing-instances routing-instances protocols ospf3
realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit protocols (ospf | ospf3) area area-id interface interface-name],
[edit protocols ospf area area-id interface interface-name topology topology-name],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id
interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface
interface-name],
[edit routing-instances routing-instance-name protocols ospf area area-id interface interface-
name topology topology-name],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast) area area-id interface interface-name]
```

Description

Specify a set of bandwidth threshold values and associated metric values for an OSPF interface or for a topology on an OSPF interface. When the bandwidth of an interface changes, Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value.

Options

bandwidth *value*—Specify the bandwidth threshold in bits per second.

- **Range:** 9600 through 1,000,000,000,000,000

metric *number*—Specify a metric value to associate with a specific bandwidth value.

- **Range:** 1 through 65,535

> **NOTE**: You must also configure a static metric value for the OSPF interface or topology with the `metric` statement. Junos OS uses this value to calculate the cost of a route from the OSPF interface or topology if the bandwidth for the interface is higher than of any bandwidth threshold values configured for bandwidth-based metrics.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth

Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth

bandwidth-based-metrics

IN THIS SECTION

- Syntax | 2838
- Hierarchy Level | 2838
- Description | 2838
- Options | 2838
- Required Privilege Level | 2839
- Release Information | 2839

Syntax

```
bandwidth-based-metrics {
    bandwidth threshold {
        metric metric;
    }
}
```

Hierarchy Level

```
[edit logical-systems name protocols isis interface-group name level name topology],
[edit logical-systems name routing-instances name protocols isis interface-group name level name
topology],
[edit logical-systems name tenants name routing-instances name protocols isis interface-group
name level name topology],
[edit protocols isis interface-group name level name topology],
[edit routing-instances name protocols isis interface-group name level name topology],
[edit tenants name routing-instances name protocols isis interface-group name level name topology]
```

Description

Specify a set of bandwidth threshold values and associated metric values for an IS-IS interface or a group of interfaces configured under IS-IS `interface-group`. Depending on the combined bandwidth of the member interfaces, IS-IS advertises the specified metric value for the interface. When the cumulative bandwidth of an interface group changes, Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value. We recommend a maximum of 3 sets of bandwidth threshold values and associated metric values for easy interface management.

NOTE: You must also configure a default static metric value for the interface or topology with the `bandwidth-based-metric` statement. Junos OS uses this value to calculate the cost of a route from the IS-IS interface or topology if the bandwidth for the interface is higher than any of the specified bandwidth threshold values. In case you have not configured the default static metric value IS-IS calculates the metric based on the reference-bandwidth.

Options

bandwidth *threshold* Specify the bandwidth threshold in bits per second.

- **Range:** 9600 through 1,000,000,000,000,000

metric ***metric metric***

Specify a metric value to associate with a specific bandwidth value.

- **Range:** 0 through 16,777,215

Required Privilege Level

routing—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

link-group-protection (Protocols IS-IS) | 9794

show isis interface-group | 21746

bandwidth (DDoS)

IN THIS SECTION

- Syntax | 2840
- Hierarchy Level | 2840
- Description | 2840
- Options | 2840
- Required Privilege Level | 2840
- Release Information | 2841

Syntax

```
bandwidth packets-per-second;
```

Hierarchy Level

- For MX Series routers, T4000 routers, and EX9200 switches:

```
[edit system ddos-protection protocols protocol-group  (aggregate | packet-type)]
```

- For ACX Series routers, PTX Series routers and QFX Series switches:

```
[edit system ddos-protection protocols protocol-group  (aggregate | packet-type)]
```

Description

Configure the DDoS bandwidth rate limit; that is, the maximum traffic rate (packets per second) allowed by the specified policer. When the value is exceeded, a violation is declared.

Options

packets-per-second—Number of packets per second that are allowed by the aggregate or packet-type policer.

- **Range:** 1 through 100,000 packets per second
- **Default:** The default bandwidth value varies by packet type or protocol. You can view the default values for all packet types or protocols before you begin control plane DDoS protection configuration by entering the `show ddos-protection protocols parameters brief` command from operational mode. For ACX Series routers, PTX Series routers and QFX Series switches, the default bandwidth limits are also provided in "protocols (DDoS) (ACX Series, PTX Series, and QFX Series)" on page 13765.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers

bandwidth-degradation

IN THIS SECTION

- Syntax | 2841
- Hierarchy Level | 2841
- Description | 2842
- Options | 2842
- Required Privilege Level | 2842
- Release Information | 2842

Syntax

```
bandwidth-degradation {
    action (log-only | restart | offline | restart-then-offline);
    no-fabric-switchover;
    percentage (1-99);
}
```

Hierarchy Level

```
[edit chassis fpc slot-numberfabric]
```

Description

Configure a FPC to take a specific action once bandwidth degradation reaches a certain percentage to avoid causing a null route in the chassis.

> **NOTE**: This configuration statement is mutually exclusive with the `offline-on-fabric-bandwidth-reduction` statement. If both statements are configured, the commit check fails and returns an error.

> **NOTE**: MX Series devices will count fabric planes in the `offline` or `fault` state as part of the bandwidth degradation percentage calculation. On PTX Series devices, when a fabric plane is offline it won't be counted in the bandwidth degradation percentage calculation.

Options

log-only	A message gets logged in the chassisd and message files when the fabric degradation threshold is reached. No other actions are taken.
restart	The FPC with a degraded fabric plane is restarted once the threshold is reached.
offline	The FPC with a degraded fabric plane is taken offline once the threshold is reached. The FPC requires manual intervention to be brought back online. This is the default action if no action attribute configured.
restart-then-offline	The FPC with a degraded fabric plane is restarted once the threshold is reached, and if fabric plane degradation is detected again within 10 minutes, the FPC is taken offline. The FPC requires manual intervention to be brought back online.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Managing Bandwidth Degradation

blackhole-action | 2979

offline-on-fabric-bandwidth-reduction | 11786

bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)

IN THIS SECTION

- Syntax | 2843
- Hierarchy Level | 2843
- Description | 2844
- Options | 2844
- Required Privilege Level | 2844
- Release Information | 2845

Syntax

```
bandwidth bps {
    ct0 bps;
    ct1 bps;
    ct2 bps;
    ct3 bps;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-
reroute],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
```

```
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name fast-reroute],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

When configuring an LSP, specify the traffic rate associated with the LSP.

When configuring fast reroute, allocate bandwidth for the reroute path. By default, no bandwidth is reserved for the rerouted path. The fast reroute bandwidth does not need to be identical to that allocated for the LSP itself.

When configuring a multiclass LSP, use the `ct`*number bandwidth* statements to specify the bandwidth to be allocated for each class type.

Options

bps—Bandwidth, in bits per second. You can specify this as an integer value. You can also use the abbreviations `k` (for a thousand), `m` (for a million), or `g` (for a billion).

- **Range:** Any positive integer
- **Default:** 0 (no bandwidth is reserved)

> **NOTE**: On the ACX Series, *bps* is the only supported option.

`ct`*number bps*—Bandwidth for the specified class type, in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations `k` (for a thousand), `m` (for a million), or `g` (for a billion [also called a thousand million]).

- **Range:** Any positive integer
- **Default:** 0 (no bandwidth is reserved)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

RELATED DOCUMENTATION

Configuring Fast Reroute

Configuring the Bandwidth Value for LSPs

Configuring Traffic-Engineered LSPs

Configuring Class-Type Bandwidth Constraints for Multiclass LSPs

bandwidth (Inline Services)

IN THIS SECTION

- Syntax | 2845
- Hierarchy Level | 2845
- Description | 2846
- Options | 2846
- Required Privilege Level | 2846
- Release Information | 2846

Syntax

```
bandwidth bandwidth-value;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number inline-services]
```

Description

Configure the amount of bandwidth in gigabits per second reserved on each Packet Forwarding Engine for tunnel traffic using inline services. Configuring the bandwidth creates a virtual tunnel interface that is represented as `si-<fpc/pic/port>`.

Starting in Junos OS Release 16.2, you are not longer required to explicitly specify a bandwidth for L2TP LNS tunnel traffic using inline services. When you do not specify a bandwidth, the maximum bandwidth supported on the PIC is automatically available for the inline services; inline services can use up to this maximum value. In earlier releases, you must specify a bandwidth when you enable inline services with the `inline-services` statement.

Options

bandwidth-value Amount of bandwidth in Gbps to reserve for tunnel traffic using inline services. You can configure bandwidth values can be as follows:

- `1g`
- `10g` through `100g` in 10 Gbps increments: `10g`, `20g`, `30g`, `40g`, `50g`, `60g`, `70g`, `80g`, `90g`, `100g`
- `100g` through `400g` in 100 Gbps increments: `100g`, `200g`, `300g`, `400g`

> **NOTE**: Values of 100 Gbps and up are available only on MPC7E, MPC8E, and MPC9E line cards.

- **Default:** Maximum bandwidth supported on the PIC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

`20g`, `30g`, and `40g` options added in Junos OS Release 14.1R3.

`100g` option added in Junos OS Release 18.2R1 on MX Series Routers with MPC7E, MPC8E, and MPC9E line cards.

`50g`, `60g`, `70g`, `80g`, `90g`, `200g`, `300g` and `400g` options added in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

bandwidth (Interfaces)

IN THIS SECTION

- Syntax | 2847
- Hierarchy Level | 2847
- Description | 2847
- Options | 2848
- Required Privilege Level | 2848
- Release Information | 2848

Syntax

```
bandwidth rate;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure the bandwidth value for an interface. This statement is valid for all logical interface types except multilink and aggregated interfaces.

NOTE: We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the `bandwidth` statement affects how the interface cost is calculated for a dynamic routing protocol, such as OSPF. By default, the interface cost for a dynamic routing protocol is calculated using the following formula:

```
cost = reference-bandwidth/bandwidth,
```

where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the `bandwidth` statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

Options

rate—Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation `c`; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps.

- **Range:** Not limited.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Interface Bandwidth

bandwidth (Interfaces)

IN THIS SECTION

- Syntax | 2849
- Hierarchy Level | 2849
- Description | 2849
- Required Privilege Level | 2849
- Release Information | 2849

Syntax

```
bandwidth bandwidth;
```

Hierarchy Level

```
[edit interfaces interface-name radio-router]
```

Description

This option controls the weight of the current (vs. maximum) data rate (value 0–100).

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

bandwidth-kbps

IN THIS SECTION

- Syntax | 2850
- Hierarchy Level | 2850
- Description | 2850
- Options | 2851
- Required Privilege Level | 2851
- Release Information | 2851

Syntax

```
bandwidth-kbps kbps;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarking profilestest-profile profile-name]
```

Description

Define the theoretical maximum bandwidth, in kilobits per second, for the test. The theoretical limit of the media for the frame size configured for the test. This value is typically set to the bandwidth of the server being tested. The range is 1,000 Kbps through 1,000,000 Kbps (1 Gbps). The value defined is the highest bandwidth value used for this test.

Options

kbps Bandwidth limit, in kilobits per second (kpbs).

- **Range:** 1,000 kbps through 1,000,000 Kbps.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

bandwidth-level

IN THIS SECTION

- Syntax | **2852**
- Hierarchy Level | **2852**
- Description | **2852**
- Default | **2852**
- Options | **2852**
- Required Privilege Level | **2853**
- Release Information | **2853**

Syntax

```
bandwidth-level kbps;
```

Hierarchy Level

```
[edit forwarding-options storm-control-profiles profile-name all]
```

Description

Configure the storm control level as the bandwidth in kilobits per second of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams.

> **NOTE**: When you configure storm control level on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control level of 15,000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30,000 Kbps of combined broadcast, multicast, and unknown unicast traffic.

Default

On EX4300 switches—If you do not specify the storm control level using either the `bandwidth-level` or the `bandwidth-percentage` statements, the storm control level defaults to 80 percent of the available bandwidth used by the combined broadcast, unknown unicast, and multicast traffic streams.

On EX9200 switches—Storm control is not enabled by default.

On MX Series routers—Storm control is not enabled by default.

Options

`bandwidth-level` *kbps*—Traffic rate in kilobits per second of the combined broadcast, multicast, and unknown unicast traffic streams.

- **Range:** 100 through 10,000,000

- **Range:** 100 through 100,000,000 on QFX10000 Series switches
- **Default:** None

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

bandwidth-percentage | 2869
Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches
Example: Using Storm Control to Prevent Network (MX Routers)
Enabling and Disabling Storm Control (ELS)

bandwidth-limit

IN THIS SECTION

- Syntax | 2854
- Hierarchy Level | 2854
- Description | 2854
- Options | 2854
- Required Privilege Level | 2854
- Release Information | 2854

Syntax

```
bandwidth-limit bps;
```

Hierarchy Level

```
[edit firewall policer policer-name if-exceeding]
[edit logical-systems logical-system-name firewall policer policer-name if-exceeding]
```

Description

Specify the traffic rate in bits per second.

Options

bps —Traffic rate to be specified in bits per second. Specify *bps* as a decimal value or as a decimal number followed by one of the following abbreviations:

- k (thousand)
- m (million)
- g (billion, which is also called a thousand million)

Range:

-
  - 1000 (1k) through 102,300,000,000 (102.3g) bps (EX Series switches)
  - 8000 (8k) through 40,000,000,000 (40g) bps (routers)

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches

Configuring Policers to Control Traffic Rates (CLI Procedure)

Understanding the Use of Policers in Firewall Filters

Basic Single-Rate Two-Color Policers

bandwidth-limit (Hierarchical Policer)

IN THIS SECTION

- Syntax | 2855
- Hierarchy Level | 2855
- Description | 2856
- Options | 2856
- Required Privilege Level | 2856
- Release Information | 2856

Syntax

```
bandwidth-limit bps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer aggregate if-exceeding],
[edit dynamic-profiles profile-name firewall hierarchical-policer premium if-exceeding],
[edit firewall hierarchical-policer aggregate if-exceeding],
[edit firewall hierarchical-policer premium if-exceeding]
```

Description

On M40e, M120, and M320 (with FFPC and SFPC) edge routers; on MPCs hosted on MX Series routers; on T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs; and on T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure the maximum average bandwidth for premium or aggregate traffic in a hierarchical policer.

Options

bps—You can specify the number of bits per second either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

Range:

- - 32,000 through 50,000,000,000 on M Series routers
  - 32,000 through 100,000,000,000 on T Series routers
  - 32,000 through 18,446,744,073,709,551,615 on MX Series routers

> **NOTE**: When you specify a numeric value beyond the supported bandwidth of the PFE, the router caps the bandwidth at the maximum supported bandwidth of the PFE.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles ... if-exceeding]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Hierarchical Policer Configuration Overview

Policer Bandwidth and Burst-Size Limits

Policer Color-Marking and Actions

Single Token Bucket Algorithm

Determining Proper Burst Size for Traffic Policers
aggregate (Hierarchical Policer) | 2064
burst-size-limit (Hierarchical Policer) | 3084
premium (Hierarchical Policer) | 13110

bandwidth-limit (Policer)

IN THIS SECTION

- Syntax | 2857
- Hierarchy Level | 2857
- Description | 2857
- Options | 2858
- Required Privilege Level | 2859
- Release Information | 2859

Syntax

```
bandwidth-limit bps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name if-exceeding],
[edit firewall policer policer-name if-exceeding],
[edit logical-systems logical-system-name policer policer-name if-exceeding]
```

Description

For a single-rate two-color policer, configure the bandwidth limit as a number of bits per second. Single-rate two-color policing uses the single token bucket algorithm to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the bandwidth limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with `low` packet loss priority (PLP) and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

> **NOTE**: This statement specifies the bandwidth limit as an absolute number of bits per second. Alternatively, for single-rate two-color policers only, you can use the `bandwidth-percent` *percentage* statement to specify the bandwidth limit as a percentage of either the physical interface port speed or the configured logical interface shaping rate.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allows bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options

bps—You can specify the number of bits per second either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

Range:

- (M Series and T Series routers) 8000 through 100,000,000,000
- (Mx Series routers) 8000 through 18,446,744,073,709,551,615
- (ACX series routers)
  - *Single-Rate Two-Color Policer*

    22000..400000000000
  - *Single-Rate Three-Color Policer*

    committed-information-rate - 22000..400000000000
  - *Two-Rate Three-Color Policer*

committed-information-rate - 22000..400000000000

peak-information-rate - 22000..400000000000

- *Hierarchical Policer*

  committed-information-rate - 22000..400000000000

  max-committed-information-rate - 22000..400000000000

> **NOTE**: When you specify a numeric value beyond the supported bandwidth of the PFE, the router caps the bandwidth at the maximum supported bandwidth of the PFE.

- **Default:** None.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit dynamic-profiles ... if-exceeding]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Two-Color Policer Configuration Overview
- Policer Bandwidth and Burst-Size Limits
- Policer Color-Marking and Actions
- Single Token Bucket Algorithm
- Determining Proper Burst Size for Traffic Policers
- bandwidth-percent | **2866**
- burst-size-limit (Policer) | **3086**

bandwidth-limit (Policer for Gigabit Ethernet Interfaces)

IN THIS SECTION

- Syntax | 2860
- Hierarchy Level | 2860
- Description | 2860
- Options | 2860
- Required Privilege Level | 2861
- Release Information | 2861

Syntax

```
bandwidth-limit bps;
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile
policer cos-policer-name aggregate],
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile
policer cos-policer-name premium]
```

Description

Define a policer to apply to nonpremium traffic.

Options

bps—Bandwidth limit, in bits per second. Specify either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 32 Kbps through 32 gigabits per second (Gbps). For IQ2 and IQ2-E interfaces 65,536 bps through 1 Gbps. For 10-Gigabit IQ2 and IQ2-E interfaces 65,536 bps through 10 Gbps.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

burst-size-limit (Policer for Gigabit Ethernet Interfaces) | 3089

bandwidth-model

IN THIS SECTION

- Syntax | 2861
- Hierarchy Level | 2862
- Description | 2862
- Options | 2862
- Required Privilege Level | 2862
- Release Information | 2862

Syntax

```
bandwidth-model {
    extended-mam;
    mam;
    rdm;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls diffserv-te],
[edit protocols mpls diffserv-te]
```

Description

Configure the bandwidth model for differentiated services. Note that you cannot configure both bandwidth models at the same time.

Options

`extended-mam`—The extended maximum allocation model (MAM) is a bandwidth model based on MAM.

`mam`—The MAM is defined in RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*.

`rdm`—The Russian dolls bandwidth allocation model (RDM) is defined in RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*. RDM makes efficient use of bandwidth by allowing the class types to share bandwidth.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Bandwidth Model

bandwidth (MPC Bandwidth)

IN THIS SECTION

- Syntax | 2863
- Hierarchy Level | 2863
- Description | 2863
- Options | 2864
- Required Privilege Level | 2864
- Release Information | 2864

Syntax

`bandwidth` *`bandwidth`*;

Hierarchy Level

`[edit chassis fpc` *`slot`*`]`

Description

Configure the MPC bandwidth.

By default, MPC8E provides a maximum bandwidth of 960 Gbps. You can upgrade MPC8E to provide an increased bandwidth of 1600 Gbps (1.6 Tbps) by using an add-on license. After you purchase the license and configure the `bandwidth 1.6T` statement, MPC8E provides an increased bandwidth of 1.6 Tbps. After you upgrade the MPC, you can use the `show chassis fpc` *`slot`* `detail` command to verify the status as shown in the following example:

```
[edit]
user@router# run show chassis fpc 3 detail
Slot 3 information:
  State                                   Online
  Temperature                          29
  Total CPU DRAM                     3200 MB
...
```

```
  Max Power Consumption            1150 Watts
  Configured Bandwidth            1600 G
  Operating Bandwidth             1600 G

[edit]
user@router#
```

Options

bandwidth Configure the MPC to operate at a specified bandwidth. You can specify `1.6T` as the value to configure the MPC to operate at a bandwidth of 1.6 Tbps.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

MPC8E

bandwidth-percent (protocols mpls)

IN THIS SECTION

- Syntax | 2865
- Hierarchy Level | 2865
- Description | 2865
- Options | 2865

Required Privilege Level | 2865
Release Information | 2865

Syntax

```
bandwidth-percent percentage;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-reroute],
[edit protocols mpls label-switched-path lsp-name fast-reroute]
```

Description

Configure the percentage of bandwidth to reserve for the detour path in case the primary path for a traffic engineered LSP or a multiclass LSP fails. The percentage configured indicates the percentage of the protected path's bandwidth that is reserved for the detour path.

Options

percentage—The percentage of the protected path's bandwidth that is reserved for the detour path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Fast Reroute

Configuring Fast Reroute for Traffic-Engineered LSPs

Configuring Fast Reroute for Multiclass LSPs

bandwidth-percent

IN THIS SECTION

- Syntax | 2866
- Hierarchy Level | 2866
- Description | 2867
- Options | 2867
- Required Privilege Level | 2868
- Release Information | 2868

Syntax

```
bandwidth-percent percentage;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name if-exceeding],
[edit firewall policer policer-name if-exceeding],
[edit logical-systems logical-system-name policer policer-name if-exceeding]
```

Description

For a single-rate two-color policer, configure the bandwidth limit as a percentage value. Single-rate two-color policing uses the *single token bucket algorithm* to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the bandwidth limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with `low` packet loss priority and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

> **NOTE**: This statement specifies the bandwidth limit as a percentage of either the physical interface port speed or the configured logical interface shaping rate. Alternatively, you can use the `bandwidth-limit` *`bps`* statement to specify the bandwidth limit as an absolute number of bits per second.

The function of the bandwidth limit is extended by the burst size (configured using the `burst-size-limit` *`bytes`* statement) to allow bursts of traffic up to a limit based on the overall traffic load:

- When a single-rate two-color policer is applied to the input or output traffic at an interface, the initial capacity for traffic bursting is equal to the number of bytes specified by this statement.
- During periods of relatively low traffic (traffic that arrives at or departs from the interface at overall rates below the token arrival rate), unused tokens accumulate in the bucket, but only up to the configured token bucket depth.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allows bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options

`percentage`—Traffic rate as a percentage of either the physical interface media rate or the logical interface configured shaping rate. You can configure a shaping rate on a logical interface by using class-of-service statement.

> **NOTE**: The bandwidth percentage policer cannot be used to rate-limit tunnel or software interfaces, or for forwarding table filters. It is only valid for interface-specific filters. When used for matching bandwidth or burst-size on aggregated Ethernet or SONET bundles, bandwidth percentage policers must be used in conjunction with `shared-bandwidth-policer`.

- **Range:** 0 through 100
- **Default:** None.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit dynamic-profiles ... if-exceeding]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Two-Color Policer Configuration Overview
- Policer Bandwidth and Burst-Size Limits
- Policer Color-Marking and Actions
- Single Token Bucket Algorithm
- Determining Proper Burst Size for Traffic Policers
- Bandwidth Policers
- bandwidth-limit (Policer) | **2857**
- burst-size-limit (Policer) | **3086**

bandwidth-percentage

IN THIS SECTION

- Syntax | 2869
- Hierarchy Level | 2869
- Description | 2869
- Default | 2870
- Required Privilege Level | 2870
- Release Information | 2870

Syntax

```
bandwidth-percentage percentage;
```

Hierarchy Level

```
[edit forwarding-options storm-control-profiles profile-name all]
```

Description

Configure the storm control level as the percentage of total interface bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams on an interface. The storm control level is configured as part of the storm control profile. As an example, a 10 Gbps interface with a 10% storm-control bandwidth configuration provides 1Gbps limit on BUM (broadcast, unknown, and multicast) traffic. So rate of BUM traffic can be up to 1Gbps and if the rate exceeds, the BUM packets are dropped.

NOTE: When you configure storm control level on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control level of 15,000 Kbps on ae1,

and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30,000 Kbps of combined broadcast, multicast, and unknown unicast traffic.

Default

On EX4300 switches—The storm control level is 80 percent of the available bandwidth used by the combined broadcast, unknown unicast, and multicast traffic streams.

On EX9200 switches—Storm control is not enabled by default.

On MX Series routers—Storm control is not enabled by default.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

bandwidth-level | **2851**

Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches

Example: Using Storm Control to Prevent Network (MX Routers)

Enabling and Disabling Storm Control (ELS)

bandwidth (Protocols Layer 2 Circuit)

IN THIS SECTION

- Syntax | **2871**
- Hierarchy Level | **2871**

- Description | 2871
- Options | 2871
- Required Privilege Level | 2871
- Release Information | 2872

Syntax

```
bandwidth (bandwidth | ctnumber bandwidth);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit protocols l2circuit neighbor address interface interface-name]
```

Description

Specify bandwidth allocation for a Layer 2 circuit or for the class types of a Layer 2 circuit.

Options

bandwidth—Configure the bandwidth in bits per second for the Layer 2 circuit. You cannot configure the bandwidth for the Layer 2 circuit and for the class types at the same time.

`ct`*number bandwidth*—Configure the bandwidth in bits per second for a class type on the Layer 2 circuit. You can configure bandwidth for up to four class types (`ct0`, `ct1`, `ct2`, `ct3`) per Layer 2 circuit. If you configure the class types, you must configure them in order, starting with class type `ct0`.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Bandwidth Allocation and Call Admission Control in Layer 2 Circuits

bandwidth (Protocols RSVP)

IN THIS SECTION

- Syntax | 2872
- Hierarchy Level | 2872
- Description | 2873
- Default | 2873
- Options | 2873
- Required Privilege Level | 2873
- Release Information | 2873

Syntax

```
bandwidth bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection bypass bypass-name],
[edit protocols rsvp interface interface-name],
```

```
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Description

For certain logical interfaces (such as Asynchronous Transfer Mode [ATM], Permanent Virtual Circuit [PVC], or Frame Relay), you cannot determine the correct bandwidth from the hardware. This statement enables you to specify the actual available bandwidth.

This statement also enables you to specify the bandwidth for a bypass label switched path (LSP). If you have configured multiple bypasses, this statement is mandatory and is applied to all of the bypass LSPs.

Default

The hardware raw bandwidth is used.

Options

bps—Bandwidth in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations `k` (for a thousand), `m` (for a million), or `g` (for a billion [also called a thousand million]).

- **Range:** Any positive integer
- **Default:** 0 (no bandwidth is reserved)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring the Bandwidth for Bypass LSPs*
>
> *Configuring Link Protection on Interfaces Used by LSPs*
>
> *Configuring Bypass LSPs*

bandwidth (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 2874
- Junos OS Evolved Hierarchy Level | 2874
- Description | 2874
- Default | 2874
- Required Privilege Level | 2875
- Release Information | 2875

Syntax

```
bandwidth bits-per-second (kbps | mbps | gbps)
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 test-profile profile-name]
```

Description

For an RFC 2544 benchmarking test, configure the rate of the traffic generated and sent out at the interface, in kilobits per second (kbps), megabits per second (mbps), or gigabits per second (gbps). The maximum rate is 100 gbps. This rate determines the characteristics of the network under test. This value is typically set to the bandwidth of the server being tested.

Default

***bits-per-second* (kbps | mbps | gbps)**	Bandwidth limit, in kilobits per second (kbps), megabits per second (mbps), or gigabits per second (gbps). • **Range:** 1,000 kbps through 100 gbps.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.4R1.

bandwidth-scale (DDoS)

IN THIS SECTION

- Syntax | 2875
- Hierarchy Level | 2875
- Description | 2876
- Options | 2876
- Required Privilege Level | 2876
- Release Information | 2876

Syntax

```
bandwidth-scale percentage;
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group (aggregate | packet-type) fpc slot-number]
```

Description

(MX Series routers with only MPCs, T4000 routers with only FPC5s, PTX Series routers, EX9200 switches, or QFX Series switches) Configure the percentage by which the DDoS bandwidth rate limit is scaled down for the aggregate or packet-type policer on the card in the specified slot.

Options

percentage—Percentage multiplied by the bandwidth rate limit to reduce the number of packets per second allowed for the packet type or protocol.

- **Range:** 1 through 100 percent
- **Default:** 100

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers

bandwidth (Static LSP)

IN THIS SECTION

- Syntax | 2877
- Hierarchy Level | 2877
- Description | 2877
- Options | 2877

Required Privilege Level | 2877
Release Information | 2878

Syntax

```
bandwidth bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
bypass],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
transit incoming-label],
[edit protocols mpls static-label-switched-path lsp-name bypass],
[edit protocols mpls static-label-switched-path lsp-name ingress],
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Description

When configuring a static LSP, specify the traffic rate associated with the LSP.

Options

bps—Bandwidth, in bits per second. You can specify this as an integer value. You can also use the abbreviations `k` (for a thousand), `m` (for a million), or `g` (for a billion).

- **Range:** Any positive integer
- **Default:** 0 (no bandwidth is reserved)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring Static LSPs

bandwidth-threshold

IN THIS SECTION

- Syntax | 2878
- Hierarchy Level | 2878
- Description | 2879
- Options | 2879
- Required Privilege Level | 2879
- Release Information | 2879

Syntax

```
bandwidth-threshold bits-per-second priority-cost priority;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-
group group-id track interface interface-name],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-
group group-id track interface interface-name],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
```

```
family inet address address vrrp-group group-id track interface interface-name],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number
family inet6 address address vrrp-inet6-group group-id track interface interface-name]
```

Description

Specify the bandwidth threshold for Virtual Router Redundancy Protocol (VRRP) logical interface tracking.

Options

bits-per-second—Bandwidth threshold for the tracked interface. When the bandwidth of the tracked interface drops below the specified value, the VRRP group uses the bandwidth threshold priority cost value. You can include up to five bandwidth threshold statements for each interface you track.

- **Range:** 1 through 10000000000000 bits per second

priority-cost *priority*—The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new primary router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring a Logical Interface to Be Tracked for a VRRP Group

bandwidth (Tunnel Services)

IN THIS SECTION

- Syntax | 2880
- Hierarchy Level | 2880
- Description | 2880
- Options | 2881
- Required Privilege Level | 2881
- Release Information | 2882

Syntax

```
bandwidth bandwidth-value;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number tunnel-services]
```

```
[edit chassis fpc slot-number pfe pfe-id core core-id channel channel-id tunnel-services]
```

```
[edit chassis feb slot slot-number pfe pfe-id core core-id channel channel-id tunnel-services]
```

Description

(ACX Series, MX Series 5G Universal Routing Platforms and T4000 Core Routers only) Configure the amount of bandwidth in gigabits per second reserved on each Packet Forwarding Engine for tunnel traffic using tunnel services. Configuring the bandwidth creates a virtual tunnel interface that is represented as *interface-<fpc/pic/port>*. Refer to Tunnel Services Overview for a list of interfaces that support tunneling services.

Options

bandwidth-value—Amount of bandwidth in Gbps to reserve for tunnel traffic using tunnel services:

- On ACX Series routers, the bandwidth values can be `1g` or `10g`.
- On ACX7K Series routers (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X) the bandwidth values can be as follows:
- - `1g`
  - `10g` through `100g` in 10 Gbps increments: `10g`, `20g`, `30g`, `40g`, `50g`, `60g`, `70g`, `80g`, `90g`, `100g`
  - `100g` through `400g` in 100 Gbps increments: `100g`, `200g`, `300g`, `400g`
- On MX Series routers, the bandwidth values can be as follows:
  - `1g`
  - `10g` through `100g` in 10 Gbps increments: `10g`, `20g`, `30g`, `40g`, `50g`, `60g`, `70g`, `80g`, `90g`, `100g`
  - `100g` through `400g` in 100 Gbps increments: `100g`, `200g`, `300g`, `400g`
- On T4000 routers, the bandwidth values can be `10g` through `100g` in 10 Gbps increments: `10g`, `20g`, `30g`, `40g`, `50g`, `60g`, `70g`, `80g`, `90g`, `100g`.

> **NOTE**: The bandwidth that you specify determines the port number of the tunnel interfaces that are created. When you specify a bandwidth of `1g`, the port number is always 10. When you specify any other bandwidth, the port number is always 0.

> **NOTE**: If you specify a bandwidth that is not compatible with the type of DPCs or MPCs and their respective Packet Forwarding Engine, tunnel services are not activated. For example, you cannot specify 1 gigabit per second bandwidth for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC or 16x10GE 3D MPC.
>
> When the tunnel bandwidth is unspecified in the Routing Engine CLI, the maximum tunnel bandwidth for MPC3E is 60G.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Statement introduced in Junos OS Evolved Release 24.1R1 for ACX7K Series Routers.

RELATED DOCUMENTATION

- Example: Configuring Tunnel Interfaces on a Gigabit Ethernet 40-Port DPC
- Tunnel Interface Configuration on MX Series Routers Overview
- Configuring Tunnel Interfaces on T4000 Routers
- Example: Configuring Tunnel Interfaces on a 10-Gigabit Ethernet 4-Port DPC
- Example: Configuring Tunnel Interfaces on the MPC3E
- tunnel-services (Chassis) | **30234**

banner (Access FTP HTTP Telnet Authentication)

IN THIS SECTION

- Syntax | **2882**
- Hierarchy Level | **2883**
- Description | **2883**
- Options | **2883**
- Required Privilege Level | **2883**
- Release Information | **2883**

Syntax

```
banner {
   fail string;
   login string;
```

```
    success string;
}
```

Hierarchy Level

```
[edit access firewall-authentication pass-through (ftp | http | telnet)]
```

Description

Configure the banners that appear to users during the FTP, HTTP, HTTPS, and Telnet firewall authentication process. The banners appear during login, after successful authentication, and after failed authentication.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

HTTPS for Web authentication is supported on SRX5400, SRX5600, and SRX5800 devices and SRX Series Firewalls from Junos OS Release 12.1X44-D10 and on SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, and SRX1500 Services Gateways from Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

Understanding Firewall Authentication Banner Customization

banner (Access Web Authentication)

IN THIS SECTION

- Syntax | 2884
- Hierarchy Level | 2884
- Description | 2884
- Options | 2884
- Required Privilege Level | 2884
- Release Information | 2885

Syntax

```
banner {
    success string;
}
```

Hierarchy Level

```
[edit access firewall-authentication web-authentication]
```

Description

Configure the banner that appears to users during the Web authentication process. The banner appears during login, after successful authentication, and after failed authentication.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

base-distinguished-name

IN THIS SECTION

- Syntax | 2885
- Hierarchy Level | 2885
- Description | 2885
- Options | 2886
- Required Privilege Level | 2886
- Release Information | 2886

Syntax

```
base-distinguished-name base-distinguished-name;
```

Hierarchy Level

```
[edit access ldap-options],
[edit access profileprofile-name ldap-options]
```

Description

Specify the base distinguished name (DN), which can be used in one of the following ways:

- If you are using the `assemble` statement so that the user's distinguished name is being assembled, the base distinguished name is appended to a username to generate the user's distinguished name. The resulting distinguished name is used in the LDAP bind call.
- If you are using the `search` statement so that the user's distinguished name is found by a search, the search is restricted to the subtree of the base distinguished name.

Options

base-distinguished-name—Series of basic properties that define the user. For example in the base distinguished name `o=juniper, c=us`, where `c` stands for country, and `o` for organization.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

batch query

IN THIS SECTION

- Syntax | 2887
- Hierarchy Level | 2887
- Description | 2887
- Options | 2888
- Required Privilege Level | 2888
- Release Information | 2889

Syntax

```
batch-query {
    items-per-batch items-per-batch;
    query-interval seconds;
}
```

Hierarchy Level

```
[edit services user-identification identity-management]
```

Description

Configure the SRX Series Firewall to communicate with the Juniper Identity Management Service server to obtain an access token to use to query the server for identity information for an individual user (IP query and user query) or a group of users (batch query). The access token allows the SRX Series Firewall to connect to the Juniper Identity Management Service server to query it for this information.

The batch-query statement allows the SRX Series Firewall to periodically query the Juniper Identity Management Service server automatically for user identity information. When you start the SRX Series Firewall, it automatically sends a batch query request to the Juniper Identity Management Service server to obtain all of the user identity information that it expects. After it receives the user identity information, the SRX Series Firewall periodically issues a query to the Juniper Identity Management Service server requesting that a new report be generated to include any newly available user identity items so as to keep its authentication table entries up-to-date.

You can configure an interval for when the batch query request is to be issued and the maximum number of user identity items to be sent in response to the query in one batch. Only remaining available user identity items are sent if their number is fewer than the configured maximum.

NOTE: If you need to refresh the user identities in the authentication table—that is, everything that was received automatically when you started the system and from subsequent batch queries or IP queries—you can clear the authentication table by disabling the user-identification feature configuration. Afterward, you can reconfigure the advanced-query feature to retrieve all available user identities. To accomplish this, you use the following sequence of CLI statements: deactivate services user-identification, commit, activate services user-identification, commit.

WARNING: Before you use this feature, you must disable active-directory-access and authentication-source options under the [edit services user-identification] hierarchy. You cannot commit this configuration if active directory authentication or the ClearPass query and webapi functions are configured and committed.

The advanced query feature queries the Juniper Identity Management Service for user identification information that the SRX Series stores in its authentication table and uses to authenticate users. Use of the Juniper Identity Management Service allows you to provision users locally and have their authentication information made available to other sites in your network for policy enforcement and reporting.

To obtain device information, such as device identity, groups, and the operating system, from the Juniper Identity Management Service server using either the batch-query or ip-query configuration, you must set the device authentication source, as follows.

```
user@host# set services user-identification device-information authentication-source network-
access-controller
```

Options

items-per-batch The maximum number of user identity items that the SRX Series Firewall will accept in one batch in response to the query.

- **Default:** `200`
- **Range:** `100-1000`

query-interval. Interval in seconds after which the SRX Series Firewall will issue a query request for newly generated user identities.

- **Default:** `5`
- **Range:** `1-60`

Required Privilege Level

1. services—To view this statement in the configuration.
2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Understanding Advanced Query Feature for Obtaining User Identity Information from JIMS

filter (Identity Management Advanced Query) | **7138**

primary connection (Identity Management Advanced Query) | **13131**

secondary connection (Identity Management Advanced Query) | **16409**

bb-sc-n

IN THIS SECTION

- Syntax | **2889**
- Hierarchy Level | **2889**
- Description | **2890**
- Options | **2890**
- Required Privilege Level | **2890**
- Release Information | **2890**

Syntax

```
bb-sc-n bb-sc-n;
```

Hierarchy Level

```
[edit interfaces interface-name                fibrechannel-options]
```

Description

Configure the buffer-to-buffer credit state change number to prevent the permanent loss of Fibre Channel credits over time (buffer-to-buffer credit recovery).

Options

bb-sc-n—Number of buffer-to-buffer state change credits.

- **Range:** 0 through 15
- **Default:** 0 (disabled)

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel interfaces | **19980**

Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric

Converting an Ethernet Interface To a Fibre Channel Interface

Configuring a Physical Fibre Channel Interface

Understanding Interfaces on an FCoE-FC Gateway

beacon-period

IN THIS SECTION

- Syntax | **2891**
- Hierarchy Level | **2891**

- Description | 2891
- Options | 2892
- Required Privilege Level | 2892
- Release Information | 2892

Syntax

```
beacon-period milliseconds;
```

Hierarchy Level

Original CLI

```
[edit ethernet-switching options secure-access-port vlan (all | vlan-name) examine-fip examine-
vn2vn]
```

ELS CLI for Platforms that Support FCoE

```
[edit vlans vlan-name forwarding-options fip-security]
```

NOTE: The `beacon-period` configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.

Description

Set the interval between periodic beacons. Beacons perform virtual link maintenance for VN_Ports in a way that is similar to FIP keepalive advertisements.

The ENode sends periodic beacons every 90 seconds on behalf of the VN_Port. Each received beacon resets the session timer for the virtual link connection to the other VN_Port. If the FCF does not receive a beacon before the beacon timer expires, the VN_Port is considered as “down” and the virtual link is terminated. The beacon timer expires in 2.5 times the configured beacon timer value.

Options

milliseconds—Time in milliseconds between beacons.

- **Range:** 250 through 90000 milliseconds
- **Default:** 8000 milliseconds

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)

Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches)

Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch)

bearer-type (L2TP Tunnel Switching)

IN THIS SECTION

- Syntax | 2893
- Hierarchy Level | 2893
- Description | 2893
- Options | 2893
- Required Privilege Level | 2893

Release Information | 2894

Syntax

```
bearer-type action;
```

Hierarchy Level

```
[edit access tunnel-switch-profile profile-name avp]
```

Description

Specify the action taken on the Bearer Type AVP (18) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.

Options

action—One of the following actions:

- `drop`—Drop the AVP.
- `regenerate`—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.
- `relay`—Forward the AVP transparently as is and send it in the switched packet.
- **Default:** `relay`

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

ber-threshold-clear

IN THIS SECTION

- Syntax | 2894
- Hierarchy Level | 2894
- Description | 2894
- Options | 2896
- Required Privilege Level | 2896
- Release Information | 2897

Syntax

```
ber-threshold-clear value;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options signal-degrade]
[edit interfaces interface-name otn-options odu-signal-degrade]
```

Description

Specify bit error rate (BER) threshold to clear the interface alarm for signal degradation.

You can configure the BER clear threshold to customize the BER that will clear an interface alarm when signal degrade monitoring is enabled.

> **NOTE**: Configuring a high BER threshold for signal degradation and a long interval might cause the internal counter register to be saturated. Such a configuration is ignored by the router, and the default values are used instead. A system log message is logged for this error.

If you configure the BER thresholds at the [`edit interfaces` *`interface-name`* `otn-options signal-degrade`] hierarchy level, then the thresholds are calculated using the pre-forward error correction (pre-FEC) BER (the BER before FEC correction). These thresholds are used for pre-FEC BER monitoring. See *Understanding Pre-FEC BER Monitoring and BER Thresholds* for more information about pre-FEC BER monitoring and determining BER threshold settings.

If you configure the BER thresholds at the [`edit interfaces` *`interface-name`* `otn-options odu-signal-degrade`] hierarchy level, then the thresholds are calculated using the post-FEC BER (the BER after FEC correction). This BER is referred to as the optical channel data unit (ODU) BER.

> **NOTE**: You can configure ODU BER thresholds only at the [`edit interfaces` *`interface-name`* `otn-options odu-signal-degrade`] hierarchy level on the P2-100GE-OTN PIC.

Table 5 on page 2895 shows the default values for pre-FEC BER and ODU BER signal degrade threshold values for different PICs. If the BER signal degrade threshold is not configured, the default value is used.

Table 5: Default Clear Threshold Values

PIC or MPC	Default Pre-FEC BER Clear Threshold Value	Default ODU BER Clear Threshold Value
P1-PTX-2-100G-WDM	3.0E–3	Not supported
P2-100GE-OTN	3.0E–3	1.0E–9
P1-PTX-24-10G-W-SFPP	3.0E–3	Not supported
MIC6-100G-CFP2	1.0E-6	1.0E-9
MPC5E	1.0E-6	1.0E-9

To configure the threshold that raises the signal degrade alarm, include the `ber-threshold-signal-degrade` statement at the same hierarchy level. To configure the time interval during which the BER must stay above or below the configured thresholds to raise or clear the alarm, include the `interval` statement at the same hierarchy level.

> **NOTE**: For the P1-PTX-2-100G-WDM PIC, the BER must stay above the signal degradation threshold for ten consecutive intervals for the alarm to be raised and the BER must stay below the clear threshold for ten consecutive intervals for the alarm to be cleared. For example, if the interval is configured as 10 ms, then the BER must stay above the signal degradation threshold for 100 ms (10 ms * 10 intervals) for the alarm to be raised, or below the clear threshold for 100 ms for the alarm to be cleared.

Options

- **Values:** *value*—BER threshold for clearing the signal degradation in scientific notation. Both the mantissa and exponent are configurable. Enter the value in the format *x*E-*n*, where *x* is the mantissa and *n* is the exponent. For example, 4.5E-3.
- **Range:** The mantissa must be a decimal number. There is no limit on the number of digits before or after the decimal point.

  The exponent must be an integer from 0 through 9.
- **Default:** See Table 5 on page 2895 for the default values.

> **BEST PRACTICE**: Always set the `ber-threshold-clear` *value* lower than the `ber-threshold-signal-degrade` *value*. For the FEC limits, see the table describing the signal degrade and clear thresholds after configuration in *Understanding Pre-FEC BER Monitoring and BER Thresholds*.

> **NOTE**: In Junos OS Release 13.2R1, only the exponent is valid input for the BER threshold value, and the mantissa is not configurable. The BER threshold value is 1.0E−*n* where *n* > 0, and the valid range of *n* is from 1 through 10.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding Pre-FEC BER Monitoring and BER Thresholds

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

ber-threshold-signal-degrade

IN THIS SECTION

- Syntax | 2897
- Hierarchy Level | 2897
- Description | 2898
- Options | 2899
- Required Privilege Level | 2900
- Release Information | 2900

Syntax

```
ber-threshold-signal-degrade value;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options signal-degrade]
[edit interfaces interface-name otn-options odu-signal-degrade]
```

Description

Specify the bit error rate (BER) threshold to raise an interface alarm for signal degradation.

You can configure the BER signal degrade threshold to customize the BER that will raise an interface alarm when signal degrade monitoring is enabled.

> **NOTE**: Configuring a high BER threshold for signal degradation and a long interval might cause the internal bit error counter register to get saturated. For example, for the P1-PTX-2-100G-WDM PIC, the internal bit error counter gets saturated when the error count reaches 2E+29. Therefore, the value of ber-threshold-signal-degrade * line rate / interval must be less than 2E+29 to avoid saturation. Assuming a fixed PIC line rate of 1.27E+11 bits per second and an interval of 1000 ms, the `ber-threshold-signal-degrade` value must be less than 4.22E−3.
>
> If the value of the ber-threshold-signal-degrade * line rate / interval exceeds the saturation limit, the configuration is ignored by the router, and the default values are used instead. A system log message is logged for this error.

If you configure the BER thresholds at the [`edit interfaces` *`interface-name`* `otn-options signal-degrade`] hierarchy level, then the thresholds are calculated using the pre-forward error correction (pre-FEC) BER (the BER before FEC correction). These thresholds are used for pre-FEC BER monitoring. See *Understanding Pre-FEC BER Monitoring and BER Thresholds* for more information about pre-FEC BER monitoring and determining BER threshold settings.

If you configure the BER thresholds at the [`edit interfaces` *`interface-name`* `otn-options odu-signal-degrade`] hierarchy level, then the thresholds are calculated using the post-FEC BER (the BER after FEC correction). This BER is referred to as the optical channel data unit (ODU) BER.

> **NOTE**: You can configure ODU BER thresholds only at the [`edit interfaces` *`interface-name`* `otn-options odu-signal-degrade`] hierarchy level on the P2-100GE-OTN PIC.

Table 6 on page 2898 shows the default values for pre-FEC BER and ODU BER signal degrade threshold values for different PICs. If the BER signal degrade threshold is not configured, the default value is used.

Table 6: Default Signal Degrade Threshold Values

PIC or MPC	Default Pre-FEC BER Signal Degrade Threshold Value	Default ODU BER Signal Degrade Threshold Value
P1-PTX-2-100G-WDM	7.5E−3	Not supported

Table 6: Default Signal Degrade Threshold Values *(Continued)*

PIC or MPC	Default Pre-FEC BER Signal Degrade Threshold Value	Default ODU BER Signal Degrade Threshold Value
P2-100GE-OTN	7.5E−3	1.0E−6
P1-PTX-24-10G-W-SFPP	7.5E−3	Not supported
MIC6-100G-CFP2	1.14E-5	1.0E-06
MPC5E	1.14E-5	1.0E-06

To configure the threshold that clears the signal degrade alarm, include the `ber-threshold-clear` statement at the same hierarchy level. To configure the time interval during which the BER must stay above or below the configured thresholds to raise or clear the alarm, include the `interval` statement at the same hierarchy level.

> **NOTE**: For the P1-PTX-2-100G-WDM PIC, the BER must stay above the signal degradation threshold for ten consecutive intervals for the alarm to be raised and the BER must stay below the clear threshold for ten consecutive intervals for the alarm to be cleared. For example, if the interval is configured as 10 ms, then the BER must stay above the signal degradation threshold for 100 ms (10 ms * 10 intervals) for the alarm to be raised, or below the clear threshold for 100 ms for the alarm to be cleared.

Options

value—BER threshold for signal degradation in scientific notation. Both the mantissa and exponent are configurable. Enter the value in the format *x*E−*n*, where *x* is the mantissa and *n* is the exponent. For example, 4.5E−3.

- **Range:** The mantissa must be a decimal number. There is no limit on the number of digits before or after the decimal point.

  The exponent must be an integer from 0 through 9.

- **Default:** See Table 6 on page 2898.

NOTE: In Junos OS Release 13.2R1, only the exponent is valid input for the BER threshold value, the mantissa is not configurable. The BER threshold value is 1.0E–*n* where *n* > 0, and the valid range of *n* is from 1 through 10.

BEST PRACTICE: To enable proactive protection before packet loss occurs, set the `ber-threshold-signal-degrade` *value* below the FEC limit. For the FEC limits, see the table describing the signal degrade and clear thresholds after configuration in *Understanding Pre-FEC BER Monitoring and BER Thresholds*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding Pre-FEC BER Monitoring and BER Thresholds

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

bert-algorithm

IN THIS SECTION

- Syntax | **2901**
- Hierarchy Level | **2901**
- Description | **2901**
- Options | **2901**

Required Privilege Level | 2903
Release Information | 2903

Syntax

```
bert-algorithm algorithm;
```

Hierarchy Level

```
[edit interfaces ce1-fpc/pic/port],
[edit interfaces ct1-fpc/pic/port],
[edit interfaces interface-name ds0-options],
[edit interfaces interface-name e1-options],
[edit interfaces interface-name e3-options],
[edit interfaces interface-name t1-options],
[edit interfaces interface-name t3-options]
```

Description

Configure the pattern to send in the bit stream during a bit error rate test (BERT). Applies to T1, E3, T3, and multichannel DS3 interfaces, the channelized interfaces (DS3, OC12, STM1), and channelized IQ and IQE interfaces (E1, E3 and DS3).

NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the `bert-algorithm` statement must be included at the `[edit interfaces ce1-fpc/pic/port]` or `[edit interfaces ct1-fpc/pic/port]` hierarchy level as appropriate.

Options

algorithm—Pattern to send in the bit stream. There are two categories of test patterns: pseudorandom and repetitive. Both patterns conform to CCITT/ITU O.151, O.152, O.153, and O.161 standards. The algorithm can be one of the following patterns:

- `all-ones-repeating`—Pattern is all ones.

- `all-zeros-repeating`—Pattern is all zeros.
- `alternating-double-ones-zeros`—Pattern is alternating pairs of ones and zeros.
- `alternating-ones-zeros`—Pattern is alternating ones and zeros.
- `pseudo-2e3`—Pattern is $2^3 - 1$.
- `pseudo-2e4`—Pattern is $2^4 - 1$.
- `pseudo-2e5`—Pattern is $2^5 - 1$.
- `pseudo-2e6`—Pattern is $2^6 - 1$.
- `pseudo-2e7`—Pattern is $2^7 - 1$.
- `pseudo-2e9-o153`—Pattern is $2^9 - 1$, as defined in the O153 standard.
- `pseudo-2e10`—Pattern is $2^{10} - 1$.
- `pseudo-2e11-o152`—Pattern is $2^{11} - 1$, as defined in the O152 standard.
- `pseudo-2e15-o151`—Pattern is $2^{15} - 1$, as defined in the O151 standard.
- `pseudo-2e17`—Pattern is $2^{17} - 1$.
- `pseudo-2e18`—Pattern is $2^{18} - 1$.
- `pseudo-2e20-o151`—Pattern is $2^{20} - 1$, as defined in the O151 standard.
- `pseudo-2e20-o153`—Pattern is $2^{20} - 1$, as defined in the O153 standard.
- `pseudo-2e21`—Pattern is $2^{21} - 1$.
- `pseudo-2e22`—Pattern is $2^{22} - 1$.
- `pseudo-2e23-o151`—Pattern is $2^{23} - 1$, as defined in the O151 standard.
- `pseudo-2e25`—Pattern is $2^{25} - 1$.
- `pseudo-2e28`—Pattern is $2^{28} - 1$.
- `pseudo-2e29`—Pattern is $2^{29} - 1$.
- `pseudo-2e31`—Pattern is $2^{31} - 1$.
- `pseudo-2e32`—Pattern is $2^{32} - 1$.
- `repeating-1-in-4`—One bit in four is set to 1; the others are set to 0.

- `repeating-1-in-8`—One bit in eight is set to 1; the others are set to 0.
- `repeating-3-in-24`—Three bits in twenty four are set to 1; the others are set to 0.
- **Default:** `pseudo-2e3`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Interface Diagnostics
Configuring E1 BERT Properties
Configuring E3 BERT Properties
Configuring T1 BERT Properties
Configuring T3 BERT Properties
Examples: Configuring T3 Interfaces
bert-error-rate | 2903
bert-period | 2905

bert-error-rate

IN THIS SECTION

- Syntax | 2904
- Hierarchy Level | 2904
- Description | 2904
- Options | 2904

Required Privilege Level | 2905
Release Information | 2905

Syntax

```
bert-error-rate rate;
```

Hierarchy Level

```
[edit interfaces ce1-fpc/pic/port],
[edit interfaces ct1-fpc/pic/port],
[edit interfaces interface-name ds0-options],
[edit interfaces interface-name e1-options],
[edit interfaces interface-name e3-options],
[edit interfaces interface-name t1-options],
[edit interfaces interface-name t3-options]
```

Description

Configure the bit error rate to use in a BERT procedure. Applies to E1, E3, T1, or T3 interfaces, and to the channelized interfaces (DS3, OC3, OC12, and STM1).

NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the `bert-error-rate` statement must be included at the `[edit interfaces ce1-fpc/pic/port]` or `[edit interfaces ct1-fpc/pic/port]` hierarchy level as appropriate.

Options

rate—Bit error rate.

- **Range:** 0 through 7, which corresponds to 10^{-1} (1 error per bit) to 10^{-7} (1 error per 10 million bits)
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- bert-algorithm | 2900
- bert-period | 2905
- ds0-options
- e1-options
- e3-options
- t1-options
- t3-options
- *Interface Diagnostics*
- Configuring E1 BERT Properties
- Configuring E3 BERT Properties
- Configuring T1 BERT Properties
- Configuring T3 BERT Properties
- Examples: Configuring T3 Interfaces

bert-period

IN THIS SECTION

- Syntax | 2906
- Hierarchy Level | 2906
- Description | 2906

- Options | 2906
- Required Privilege Level | 2907
- Release Information | 2907

Syntax

```
bert-period seconds;
```

Hierarchy Level

```
[edit interfaces ce1-fpc/pic/port],
[edit interfaces ct1-fpc/pic/port],
[edit interfaces interface-name ds0-options],
[edit interfaces interface-name e1-options],
[edit interfaces interface-name e3-options],
[edit interfaces interface-name t1-options],
[edit interfaces interface-name t3-options]
```

Description

Configure the duration of a BERT test. Applies to E1, E3, T1, and T3 interfaces, and to E1, E3, T1, and T3 partitions on the channelized interfaces (CE1, CT1, DS3, OC3, OC12, OC48, STM1, STM4, and STM16).

E1 and T1 IQ, IQE, and standard interfaces support an extended BERT period range, up to 86,400 seconds (24 hours).

NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the `bert-period` statement must be included at the `[edit interfaces ce1-fpc/pic/port]` or `[edit interfaces ct1-fpc/pic/port]` hierarchy level as appropriate.

Options

seconds—Test duration. Range and default values vary by interface type.

Range:

- - PIC-dependent—Normal BERT period: either 1 through 239 seconds or 1 through 240 seconds
  - PIC-dependent—Extended BERT period: from 1 through 86,400 seconds

Default:

- - Normal BERT period: 10 seconds
  - Extended BERT period (on supported E1 interfaces): 10 seconds
  - Extended BERT period (on supported T1 interfaces): 240 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Interface Diagnostics

Configuring E1 BERT Properties

Configuring E3 BERT Properties

Configuring T1 BERT Properties

Configuring T3 BERT Properties

bert-algorithm | **2900**

bert-error-rate | **2903**

best-site

IN THIS SECTION

- Syntax | 2908
- Hierarchy Level | 2908
- Description | 2908
- Required Privilege Level | 2908
- Release Information | 2909

Syntax

```
best-site;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
site site-name],
[edit routing-instances routing-instance-name protocols vpls site site-name]
```

Description

Enables the VPLS multihoming best site functionality, allowing the site on which it has been enabled to be the preferred site for this PE router. This statement must be configured on all PE routers within the optimized VPLS routing instance.

NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: VPLS Multihoming, Improved Convergence Time

bfd

IN THIS SECTION

- Syntax | 2909
- Hierarchy Level | 2910
- Description | 2910
- Required Privilege Level | 2910
- Release Information | 2910

Syntax

```
bfd {
    version (0 | 1 | automatic);
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    detection-time {
        threshold milliseconds;
```

```
    }
    session-mode (automatic | multihop | singlehop);
    holddown-interval milliseconds;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method],
[edit system services dhcp-local-server dhcpv6 liveness-detection method],
[edit forwarding-options dhcp-relay liveness-detection method],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method],
[edit system services dhcp-local-server group group-name liveness-detection method],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method],
[edit forwarding-options dhcp-relay group group-name liveness-detection method],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method]
```

Description

Configure Bidirectional Forwarding Detection (BFD) as the liveness detection method.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients
- Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

bfd (protocols)

IN THIS SECTION

- Syntax | 2911
- Hierarchy Level | 2911
- Description | 2912
- Default | 2912
- Options | 2912
- Required Privilege Level | 2914
- Release Information | 2914

Syntax

```
bfd {
        traceoptions {
            file filename <files number> <match regular-expression> <size size> <world-
readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
        sbfd {
            pool value;
            local-discriminator value {
                minimum-receive-interval value;
                local-ip-address IP address;
            }
            local-discriminator-ip IP address;
        }
    }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
```

```
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Configure trace options for Bidirectional Forwarding Protocol (BFD) traffic. Configure seamless BFD (S-BFD) over non-colored static segment-routing LSPs with first-hop label resolution and use S-BFD as a fast mechanism to detect path failures and to trigger global convergence.

Default

If you do not include this statement, no BFD tracing operations are performed.

Options

disable—(Optional) Disable the BFD tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the **/var/log** directory . We recommend that you place global routing protocol tracing output in the **routing-log** file.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 2 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. These are the BFD protocol tracing options:

- **adjacency**—Trace adjacency messages.
- **all**—Trace all options for BFD.
- **error**—Trace all errors.
- **event**—Trace all events.
- **issu**—Trace in-service software upgrade (ISSU) packet activity.

- **nsr-packet**—Trace non-stop-routing (NSR) packet activity.
- **nsr-synchronization**—Trace NSR synchronization events.
- **packet**—Trace all packets.
- **pipe**—Trace pipe messages.
- **pipe-detail**—Trace pipe messages in detail.
- **ppm-packet**—Trace packet activity by periodic packet management (PPM).
- **state**—Trace state transitions.
- **timer**—Trace timer processing.

match ***regular-expression***—(Optional) Regular expression for lines to be logged.

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the trace file again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

world-readable—(Optional) Allow any user to read the log file.

sbfd—Seamless BFD parameters. In S-BFD, each responder might have multiple discriminators. The discriminators might be advertised by IGP to other routers, or they might be statically configured on these routers.

local-discriminator ***value***—Local discriminator for seamless BFD responder.

Range: 1 through 4294967295.

- **local-ip-address** ***IP address***—IPv4 source address.
- **minimum-receive-interval** ***value***—Minimum receive interval for seamless BFD responder.

  Range: 10 through 255000 milliseconds.

local-discriminator-ip ***IP address***—Local discriminator IP address for sBFD responder.

pool ***value***—List of seamless BFD endpoints.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace–control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

bfd-liveliness

IN THIS SECTION

- Syntax | **2915**
- Hierarchy Level | **2915**
- Description | **2915**
- Options | **2916**
- Required Privilege Level | **2916**
- Release Information | **2916**

Syntax

Syntax (for monitor)

```
bfd-liveliness destination-ip {
 interface interface;
 routing-instance routing-instance;
 session-type(multihop | singlehop);
 src-ip src-ip;
}
```

Syntax (for monitor-object)

```
bfd-liveliness {
 destination-ip destination-ip {
    interface interface;
    routing-instance routing-instance;
    session-type(multihop | singlehop);
    src-ip src-ip;
    weight value;
   }
threshold value;
```

Hierarchy Level

```
[edit chassis high-availability services-redundancy-group name monitor],
[edit chassis high-availability services-redundancy-group name monitor monitor-object name]
```

Description

Configure Bidirectional Forwarding Detection (BFD) monitoring to detect failures in a network. You can configure Multinode High Availability to monitor one or more links using BFD. This configuration

triggers a failover in the event of BFD failure. Configure BFD liveliness by specifying source and destination IP and the interface where the peer device is directly connected to.

Options

<dst-ip> Destination IPV4/IPV6 address for BFD monitoring

interface *interface*— Name of the interface for single-hop sessions

routing-instance *routing-instance*— Routing instance to locate the route

session-type— BFD Monitoring session type. Configure BFD session mode to be single-hop or multihop. By default, BGP uses single-hop BFD sessions if the peer is directly connected to the router's interface.

src-ip *src-ip*— Source IPV4/IPV6 address for BFD monitoring

weight *number* (monitor-object) — BFD monitoring weight value for SRG monitoring object.

- **Range:** 1 through 1000

threshold (monitor-object) —BFD-liveliness monitoring threshold value.

- **Range:** 1 through 1000

Within a monitor-object, whenever the particular object fails to trigger a failover as result of BFD monitoring, the system considers the event as monitoring failure and adds a count to BFD-liveliness.

When the count exceeds the threshold value BFD-liveliness, the system adds the count to parent monitoring-object's threshold value. Monitoring-object's threshold also includes threshold values of interface monitoring and IP monitoring.

When the sum of the weights of all monitor objects bound to the SRG reaches a value equal to or greater than the threshold value configured on the SRG, the system triggers a monitor failure for that SRG. SRG fails over to the other node.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

bfd-liveliness (Activeness Probe)

IN THIS SECTION

- Syntax | 2917
- Hierarchy Level | 2917
- Description | 2917
- Options | 2918
- Required Privilege Level | 2918
- Release Information | 2918

Syntax

```
bfd-liveliness {
    destination-ip destination-ip;
    interface interface;
    minimum-interval milliseconds;
    multiplier multiplier;
    routing-instance routing-instance;
    source-ip source-ip;
    }
```

Hierarchy Level

```
[edit chassis high-availability services-redundancy-group name activeness-probe]
```

Description

Configure BFD monitoring for split-brain detection in a Multinode High Availability setup.

Options

destination-ip ***destination-ip—***	Destination IPV4/IPV6 address for BFD monitoring
interface ***interface-name—***	Name of the interface for single-hop sessions
minimum-interval ***milliseconds—***	Minimum transmit and receive interval (milliseconds) • **Range:** 1 through 255000
multiplier ***multiplier—***	Detection time multiplier • **Range:** 1 through 255
routing-instance ***routing-instance—***	(Optional) Routing instance to locate the route
source-ip ***source-ip—***	Source IPV4/IPV6 address for BFD monitoring

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

bfd-liveness-detection (BGP)

IN THIS SECTION

- Syntax | 2919
- Hierarchy Level | 2919
- Description | 2920
- Options | 2920
- Required Privilege Level | 2922
- Release Information | 2922

Syntax

```
bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    holddown-interval milliseconds;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    session-mode (automatic | multihop | single-hop);
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
```

Hierarchy Level

```
[edit logical-systems name protocols              bgp],
[edit logical-systems name protocols              bgp group group-name],
[edit logical-systems name protocols              bgp group group-name neighbor address],
[edit logical-systems name              routing-instances              name protocols bgp],
[edit logical-systems name              routing-instances              name protocols bgp group
group-name],
[edit logical-systems name              routing-instances              name protocols bgp group
group-name neighbor address],
[edit logical-systems name tenants name routing-instances name protocols              bgp ],
[edit logical-systems name tenants name routing-instances name protocols              bgpgroup
group-name],
[edit logical-systems name tenants name routing-instances name protocols              bgpgroup
group-name neighbor address],
[edit protocols bgp],
```

```
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances name protocols              bgp],
[edit routing-instances name protocols              bgp group group-name],
[edit routing-instances name protocols              bgp group group-name neighbor address],
[edit tenants name routing-instances name protocols bgp]
[edit tenants name routing-instances name protocols bgp group group-name]
[edit tenants name routing-instances name protocols bgp group group-name neighbor address]
```

Description

Configure bidirectional failure detection (BFD) timers and authentication for BGP.

For IBGP and multihop EBGP support, configure the `bfd-liveness-detection` statement at the global `[edit bgp protocols]` hierarchy level. You can also configure IBGP and multihop support for a routing instance or a logical system.

Options

holddown-interval *milliseconds*

(Optional) Configure an interval specifying how long a BFD session must remain up before a state change notification is sent.

When you configure the hold-down interval for the BFD protocol for EBGP, the BFD session is unaware of the BGP session during this time. In this case, if the BGP session goes down during the configured hold-down interval, BFD already assumes the BGP session is down and does not send a state change notification. The `holddown-interval` statement is supported only for EBGP peers at the `[edit protocols bgp group group-name neighbor address]` hierarchy level. If the BFD session goes down and then comes back up during the configured hold-down interval, the timer is restarted. You must configure the hold-down interval on both EBGP peers. If you configure the hold-down interval for a multihop EBGP session, you must also configure a local IP address by including the `local-address` statement at the `[edit protocols bgp group group-name]` hierarchy level.

- **Range:** 0 through 255,000 milliseconds
- **Default:** 0

minimum-interval *milliseconds*

(Required) Configure the minimum interval at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the minimum-interval (specified under the `transmit-interval` statement) and minimum-receive-interval statements.

- **Range:** 1 through 255,000 milliseconds

minimum-receive-interval *milliseconds*

(Optional) Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the `minimum-interval` statement.

- **Range:** 1 through 255,000 milliseconds

multiplier *number*

(Optional) Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

- **Range:** 1 through 255
- **Default:** 3

no-adaptation

(Optional) Configure BFD sessions not to adapt to changing network conditions. We recommend that you *do not* disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. However, include the `no-adaptation` statement in the configuration if you do not want BFD sessions to adapt to changing network conditions.

You can use the `clear bfd adaptation` command to return BFD interval timers to their configured values. The `clear bfd adaptation` command does not affect traffic flow on the routing device.

- **Default:** BFD sessions adapt to changing network conditions.

session-mode (automatic | multihop | single-hop)

(Optional) Configure BFD session mode to be single-hop or multihop. By default, BGP uses single-hop BFD sessions if the peer is directly connected to the router's interface. BGP uses multihop BFD sessions if the peer is not directly connected to the router's interface. If the peer session's `local-address` option is configured, the directly connected check is based partly on the source address that would be used for BGP and BFD.

For backward compatibility, you can override the default behavior by configuring the `single-hop` or `multihop` option. Before Junos OS Release 11.1, the behavior was to assume that IBGP peer sessions were multihop.

- **Values:** Specify one of these values:
  - automatic—Configure BGP to use single-hop BFD sessions if the peer is directly connected to the router's interface, and multihop BFD sessions if the peer is not directly connected to the router's interface.
  - multihop—Configure BGP to use multihop BFD sessions.
  - single-hop—Configure BGP to use single-hop BFD sessions.
- **Default:** automatic

version (1 | automatic)

(Optional) Specify the BFD version.

- **Values:** Specify one of the following:
  - 1—Configure BFD version 1.
  - automatic—Configure the device to automatically detect the BFD version.
- **Default:** The device automatically detects the BFD version.

The remaining statements are explained separately. See CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

`detection-time threshold` and `transmit-interval threshold` options introduced in Junos OS Release 8.2.

Support for logical routers introduced in Junos OS Release 8.3.

Support for IBGP and multihop EBGP sessions introduced in Junos OS Release 8.3.

`holddown-interval` option introduced in Junos OS Release 8.5. You can configure this option only for EBGP peers at the `[edit protocols bgp group` *`group-name`* `neighbor` *`address`*`]` hierarchy level.

`no-adaptation` option introduced in Junos OS Release 9.0.

Support for BFD authentication introduced in Junos OS Release 9.6.

`session-mode` option introduced in Junos OS Release 11.1.

Support for BFD on IPv6 interfaces with BGP introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

Example: Configuring BFD Authentication for Securing Static Routes

Example: Configuring BFD on Internal BGP Peer Sessions

Example: Configuring BFD Authentication for BGP

Understanding BFD for BGP

bfd-liveness-detection (EVPN Node Detection)

IN THIS SECTION

- Syntax | **2923**
- Hierarchy Level | **2924**
- Description | **2924**
- Options | **2924**
- Required Privilege Level | **2925**

Syntax

```
bfd-liveness-detection {
 no-adaptation;
 detection-time {
  threshold milliseconds;
 }
 inline-disable;
```

```
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier multiplier;
    neighbor neighbor;
    pdu-size pdu-size;
    transmit-interval {
     minimum-interval milliseconds;
     threshold milliseconds;
    }
    version(0 | 1 | automatic);
   }
```

Hierarchy Level

```
[edit logical-systems name protocols evpn node-detection],
[edit protocols evpn node-detection]
```

Description

Configure Bidirectional Forwarding Detection (BFD) timers and options for EVPN node detection.

Options

no-adaptation Disable the BFD adaptation. Include this statement to prevent the BFD sessions from adapting to changing network conditions.

detection-time Mandatory

- threshold *milliseconds* — Configure a threshold for the adaptation of the BFD session detection time. When the detection time equals or exceeds the threshold, a single trap and a single system log message are sent.
  - **Range:** 1 through 255000

inline-disable Disable inline mode for this BFD session

minimum-interval *milliseconds* Configure the minimum transmit and receive interval the local device transmits to, and expects a reply from, the BFD neighbor. Optionally, you can specify the minimum receive interval separately using the `minimum-receive-interval` option. Or you can use the `transmit-interval` statement's `minimum-interval` and `minimum-receive-interval` options.

- **Range:** 1 through 255000

minimum-receive-interval *milliseconds*

Configure the minimum interval the local device expects to receive a reply from the BFD neighbor. Optionally, instead of using this statement, you can configure the minimum receive interval using the `minimum-interval` statement.

- **Range:** 1 through 255000

multiplier *number*

Configure the number of BFD packets not received before declaring the BFD peer interface down.

- **Default:** 3
- **Range:** 1 through 255

neighbor

Configure the interface address of the BFD neighbor.

pdu-size

Configure the BFD transport protocol payload size

- **Default:** 24
- **Range:** 24 through 9000

transmit-interval

- minimum-interval *milliseconds* — Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the minimum-interval statement.
  - **Range:** 1 through 255000
- threshold *milliseconds* — Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

version

BFD protocol version number

- 0—BFD version 0 (deprecated)
- 1—BFD version 1
- automatic—Choose BFD version automatically

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Backup Liveness Detection on EVPN Dual Homing

bfd-liveness-detection (LAG)

IN THIS SECTION

- Syntax | 2926
- Hierarchy Level | 2927
- Description | 2927
- Options | 2927
- Required Privilege Level | 2929
- Release Information | 2929

Syntax

```
bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    damping {
        half-life-interval seconds;
        increment increment;
        max-suppress-time seconds;
        reuse reuse;
        suppress suppress;
    }
    detection-time {
        threshold milliseconds;
```

```
    }
    holddown-interval milliseconds;
    local-address bfd-local-address;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    neighbor bfd-neighbor-address;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
```

Hierarchy Level

```
[edit interfaces name aggregated-ether-options],
[edit dynamic-profiles name interfaces name aggregated-ether-options],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-ether-options]
```

Description

Configure Bidirectional Forwarding Detection (BFD) timers and authentication for aggregated Ethernet interfaces.

For an aggregated ethernet interface, you cannot configure all three configuration options, `bfd-liveness-detection`, `minimum-links`, and `sync-reset` at the same time.

Options

holddown-interval *milliseconds*

Specify a time limit, in milliseconds, indicating the time that a BFD session remains up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

- **Range:** 0 through 255,000
- **Default:** 0

local-address *bfd-local-address*

Specify the loopback address or the AE interface address of the source of the BFD session.

> **NOTE**: Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD `local-address` against the interface or loopback IP address before the configuration commit. Junos OS performs this check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails. The configured micro-BFD `local-address` should match with the micro-BFD `neighbour-address` configured on the peer router.

minimum-interval *milliseconds*

Specify a minimum time interval after which the local routing device transmits a BFD packet and then expects to receive a reply from the BFD neighbor. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the `transmit-interval minimum-interval` statement.

- **Range:** 1 through 255,000

minimum-receive-interval *milliseconds*

Specify the minimum time interval after which the routing device expects to receive a reply from the BFD neighbor.

- **Range:** 1 through 255,000

multiplier *number*

Specify the number of BFD packets that were not received by the BFD neighbor before the originating interface is declared down.

- **Range:** 1 through 255

neighbor *bfd-neighbor-address*

Specify the loopback address or the AE interface address of a remote destination to send BFD packets.

no-adaptation

Disable the BFD adaptation. Include this statement if you do not want the BFD sessions to adapt to changing network conditions. We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

version

Configure the BFD version to detect (BFD version 1) or autodetect (the BFD version).

> **NOTE**: The `version` option is not supported on the QFX Series. Starting in Junos OS Release 17.2R1, a warning will appear if you attempt to use this command.

- **Default:** automatic

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

- authentication
- detection-time
- transmit-interval
- *Configuring Micro BFD Sessions for LAG*
- Example: Configuring Independent Micro BFD Sessions for LAG
- Understanding Independent Micro BFD Sessions for LAG

bfd-liveness-detection (Layer 2 VPN and VPLS)

IN THIS SECTION

- Syntax | 2930
- Hierarchy Level | 2930
- Description | 2930
- Required Privilege Level | 2931
- Release Information | 2931

Syntax

```
bfd-liveness-detection {
    detection-time {
        threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances instance_name protocols l2vpn oam],
[edit logical-systems name routing-instances instance_name protocols vpls neighbor neighbor-id
oam],
[edit logical-systems name routing-instances instance_name protocols vpls mesh-group mesh-group-
name neighbor neighbor-id oam],
[edit logical-systems name routing-instances instance_name protocols vpls oam],
[edit routing-instances instance_name protocols l2vpn oam],
[edit routing-instances instance_name protocols vpls neighbor neighbor-id oam],
[edit routing-instances instance_name protocols vpls mesh-group mesh-group-name neighbor
neighbor-id oam],
[edit routing-instances instance_name protocols vpls oam]
```

Description

Configure bidirectional failure detection timers.

The BFD failure detection timers are adaptive and can be adjusted to be more or less aggressive. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission

(Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the `clear bfd adaptation` command to return BFD interval timers to their configured values. The `clear bfd adaptation` command is hitless, meaning that the command does not affect traffic flow on the routing device.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Statement introduced in Junos OS Release 13.2 for Layer 2 VPNs and VPLS.

RELATED DOCUMENTATION

Configuring BFD for Layer 2 VPN and VPLS

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

bfd-liveness-detection (LSP)

IN THIS SECTION

- Syntax | 2932
- Hierarchy Level | 2932
- Description | 2932
- Options | 2932
- Required Privilege Level | 2933
- Release Information | 2933

Syntax

```
bfd-liveness-detection {
    minimum-interval milliseconds;
    multiplier multiplier;
    no-router-alert-option;
    sbfd {
        remote-discriminator remote-discriminator;
    }
}
```

Hierarchy Level

```
[edit protocols source-packet-routing segment-list],
```

```
[edit protocols source-packet-routing source-routing-path  lsp-name primary segment-list-name],
```

```
[edit protocols source-packet-routing source-routing-path  lsp-name secondary segment-list-name]
```

Description

Bidirectional forwarding detection options.

Options

minimum-interval

Minimum transmit and receive interval (milliseconds).

- **Range:** 1 through 255000

multiplier

Detection time multiplier.

- **Default:** 3
- **Range:** 1 through 255

no-router-alert-option

Do not set the Router Alert option in IP header.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.4R1.

Support at the following hierarchy levels introduced in Junos OS Release 19.1R1: [edit protocols source-packet-routing source-routing-path *lsp-name* primary *segment-list-name*], and [edit protocols source-packet-routing source-routing-path *lsp-name* secondary *segment-list-name*].

> **NOTE**: Starting in Junos OS Release 19.1R1, the bfd-liveness-detection statement is not supported at the [edit protocols source-packet-routing segment-list] hierarchy level.

bfd-liveness-detection (Protocols IS-IS)

IN THIS SECTION

- Syntax | 2933
- Hierarchy Level | 2934
- Description | 2934
- Options | 2935
- Required Privilege Level | 2936
- Release Information | 2936

Syntax

```
bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
```

```
    detection-time {
        threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis
interface interface-name],
[edit protocols isis interface interface-name],
[edit protocols isis interface interface-name family inet|inet6],
[edit routing-instances routing-instance-name protocols isis interface interface-name]
```

Description

Configure bidirectional failure detection timers and authentication.

> **NOTE**: Starting with Junos OS Release 16.1, you can configure IS-IS BFD sessions for the IPv6 address family. Therefore, this configuration statement is also available under the `[edit protocols isis interface interface-name family inet|inet6]` hierarchy level.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (RX) interval by two if the local BFD instance is the reason for the session flap. The transmission (TX) interval is increased by two if the remote BFD instance is the reason for the session flap.

You can use the `clear bfd adaptation` command to return BFD interval timers to their configured values. The `clear bfd adaptation` command does not affect traffic flow on the routing device.

Options

`authentication algorithm` *`algorithm-name`* —Configure the algorithm used to authenticate the specified BFD session: `simple-password`, `keyed-md5`, `keyed-sha-1`, `meticulous-keyed-md5`, `meticulous-keyed-sha-1`.

`authentication key-chain` *`key-chain-name`*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the `authentication-key-chains key-chain` statement at the `[edit security]` hierarchy level.

`authentication loose-check`—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.

`detection-time threshold` *`milliseconds`*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

`minimum-interval` *`milliseconds`*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the `transmit-interval minimum-interval` and `minimum-receive-interval` statements.

- **Range:** 1 through 255,000

`minimum-receive-interval` *`milliseconds`*—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the `minimum-interval` statement.

- **Range:** 1 through 255,000

`multiplier` *`number`*—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

- **Range:** 1 through 255
- **Default:** 3

`no-adaptation`—Specify that BFD sessions not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

`transmit-interval threshold` *`milliseconds`*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

- **Range:** 0 through 4,294,967,295 (2^{32} – 1)

`transmit-interval minimum-interval` *`milliseconds`*—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the `minimum-interval` statement.

- **Range:** 1 through 255,000

`version`—Configure the BFD version to detect: `1` (BFD version 1) or `automatic` (autodetect the BFD version)

- **Default:** `automatic`

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`detection-time threshold` and `transmit-interval threshold` options added in Junos OS Release 8.2.

Support for logical systems introduced in Junos OS Release 8.3.

`no-adaptation` statement introduced in Junos OS Release 9.0.

`authentication algorithm`, `authentication key-chain`, and `authentication loose-check` options introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring BFD for IS-IS

Understanding BFD Authentication for IS-IS

bfd-liveness-detection (Protocols LDP)

IN THIS SECTION

- Syntax | 2937
- Hierarchy Level | 2938
- Description | 2938
- Options | 2938
- Required Privilege Level | 2938
- Release Information | 2939

Syntax

```
bfd-liveness-detection {
    detection-time threshold milliseconds;
    ecmp;
    failure-action {
        remove-nexthop;
        remove-route;
    }
    holddown-interval seconds;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (0 | 1 | automatic);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp oam],
[edit logical-systems logical-system-name protocols ldp oam fec address],
[edit protocols ldp oam],
[edit protocols ldp oam fec address]
```

Description

Enable Bidirectional Forwarding Detection (BFD) for all MPLS LSPs or for just a specific LSP.

Options

`minimum-interval`—Minimum transmit and receive interval.

- **Range:** 50 through 255,000 milliseconds
- **Default:** 50

`minimum-receive-interval`—Minimum receive interval.

- **Range:** 50 through 255,000 milliseconds
- **Default:** 50

`minimum-transmit-interval`—Minimum transmit interval.

- **Range:** 50 through 255,000 milliseconds
- **Default:** 50

`multiplier`—Detection time multiplier.

- **Range:** 50 through 255
- **Default:** 3

The other options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Support for the `bfd-liveness-detection` statement at the `[edit protocols ldp oam fec address]` hierarchy level and the `ecmp` option added in Junos OS Release 9.0.

Support for the `failure-action` statement with the `remove-nexthop` and `remove-route` options and the `holddown-interval` statement added in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring BFD for LDP LSPs

bfd-liveness-detection (Protocols MPLS)

IN THIS SECTION

- Syntax | 2939
- Hierarchy Level | 2940
- Description | 2940
- Options | 2940
- Required Privilege Level | 2940
- Release Information | 2941

Syntax

```
bfd-liveness-detection {
    failure-action {
        make-before-break teardown-timeout seconds;
        teardown;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
```

```
    multiplier detection-time-multiplier;
}
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name oam],
[edit protocols mpls oam]
```

Description

Enable Bidirectional Forwarding Detection (BFD) for all of the MPLS LSPs or for just a specific LSP.

Options

`minimum-interval`—Minimum transmit and receive interval.

- **Range:** 50 through 255,000 milliseconds
- **Default:** 50

`minimum-receive-interval`—Minimum receive interval.

- **Range:** 50 through 255,000 milliseconds
- **Default:** 50

`minimum-transmit-interval`—Minimum transmit interval.

- **Range:** 50 through 255,000 milliseconds
- **Default:** 50

`multiplier`—Detection time multiplier.

- **Range:** 1 through 255
- **Default:** 3

The `failure-action` statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

`failure-action` option added in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring BFD for MPLS IPv4 LSPs

Configuring Bidirectional Forwarding Detection for MPLS (CLI Procedure)

bfd-liveness-detection (Protocols OSPF)

IN THIS SECTION

- Syntax | 2941
- Hierarchy Level | 2942
- Description | 2942
- Options | 2943
- Required Privilege Level | 2944
- Release Information | 2944

Syntax

```
bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
```

```
    }
    full-neighbors-only
    holddown-interval holddown-interval;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    pdu-size bytes
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface
interface-name],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast) area area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-
name],
[edit protocols (ospf | ospf3) area area-id interface interface-name],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id
interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface
interface-name],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast) area area-id interface interface-name]
```

Description

Configure bidirectional failure detection timers and authentication for OSPF.

The remaining statements are explained separately. See CLI Explorer.

Options

`authentication algorithm` *`algorithm-name`* —Configure the algorithm used to authenticate the specified BFD session: `simple-password`, `keyed-md5`, `keyed-sha-1`, `meticulous-keyed-md5`, or `meticulous-keyed-sha-1`.

`authentication key-chain` *`key-chain-name`*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the `authentication-key-chains key-chain` statement at the `[edit security]` hierarchy level.

`authentication loose-check`—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

`detection-time threshold` *`milliseconds`*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

`full-neighbors-only`—Establish BFD sessions only for OSPF neighbors in the full state. The default behavior is to establish BFD sessions for all OSPF neighbors.

`holddown-interval` *`holddown-interval`*—Time to hold the session-UP notification to the client.

- **Range:** 0 through 255000 milliseconds

`minimum-interval` *`milliseconds`*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the `transmit-interval minimum-interval` and `minimum-receive-interval` statements.

- **Range:** 1 through 255,000 milliseconds

`minimum-receive-interval` *`milliseconds`*—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the `minimum-interval` statement.

- **Range:** 1 through 255,000 milliseconds

`multiplier` *`number`*—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

- **Range:** 1 through 255
- **Default:** 3

`no-adaptation`—Specify that BFD sessions should not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

`pdu-size`—Adjust the size of BFD protocol data units (PDUs).

- **Range:** 24 through 9000 bytes
- **Default:** 24

`transmit-interval threshold` *`milliseconds`*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

- **Range:** 0 through 4,294,967,295 (2^{32} – 1)

`transmit-interval minimum-interval` *`milliseconds`*—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the `minimum-interval` statement.

- **Range:** 1 through 255,000

`version`—Configure the BFD version to detect: `1` (BFD version 1) or `automatic` (autodetect the BFD version).

- **Default:** `automatic`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`detection-time threshold` and `transmit-interval threshold` options added in Junos OS Release 8.2.

Support for logical systems introduced in Junos OS Release 8.3.

`no-adaptation` option introduced in Junos OS Release 9.0.

`no-adaptation` option introduced in Junos OS Release 9.0 for EX Series switches.

Support for OSPFv3 introduced in Junos OS Release 9.3.

Support for OSPFv3 introduced in Junos OS Release 9.3 for EX Series switches.

`full-neighbors-only` option introduced in Junos OS Release 9.5.

`full-neighbors-only` option introduced in Junos OS Release 9.5 for EX Series switches.

`holddown-interval` option introduced in Junos OS Release 19.4 for MX Series routers.

`authentication algorithm`, `authentication key-chain`, and `authentication loose-check` options introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring BFD for OSPF

Example: Configuring BFD Authentication for OSPF

bfd-liveness-detection (Protocols PIM)

IN THIS SECTION

- Syntax | 2945
- Hierarchy Level | 2946
- Description | 2946
- Required Privilege Level | 2946
- Release Information | 2946

Syntax

```
bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
```

```
        threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (0 | 1 | automatic);
}
```

Hierarchy Level

```
[edit protocols pim interface (Protocols PIM) interface-name family (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM) interface-
name family (inet | inet6)]
```

Description

Configure bidirectional forwarding detection (BFD) timers and authentication for PIM.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

authentication option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD for PIM

bfd-liveness-detection (Protocols RIP)

IN THIS SECTION

- Syntax | 2947
- Hierarchy Level | 2948
- Description | 2948
- Options | 2948
- Required Privilege Level | 2949
- Release Information | 2949

Syntax

```
bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip group group-name],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

Configure bidirectional failure detection timers and authentication.

The remaining statements are explained separately. See CLI Explorer.

Options

`authentication algorithm` *`algorithm-name`* —Configure the algorithm used to authenticate the specified BFD session: `simple-password`, `keyed-md5`, `keyed-sha-1`, `meticulous-keyed-md5`, or `meticulous-keyed-sha-1`.

`authentication key-chain` *`key-chain-name`*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the `authentication-key-chains key-chain` statement at the `[edit security]` hierarchy level.

`authentication loose-check`—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication is not configured at both ends of the BFD session.

`detection-time threshold` *`milliseconds`*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

`minimum-interval` *`milliseconds`*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the `transmit-interval minimum-interval` and `minimum-receive-interval` statements.

- **Range:** 1 through 255,000 milliseconds

`minimum-receive-interval` *`milliseconds`*—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the `minimum-interval` statement.

- **Range:** 1 through 255,000 milliseconds

`multiplier` *`number`*—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

- **Range:** 1 through 255
- **Default:** 3

`no-adaptation`—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

`transmit-interval threshold` *`milliseconds`*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

- **Range:** 0 through 4,294,967,295 (2^{32} – 1)

`transmit-interval minimum-interval` *`milliseconds`*—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the `minimum-interval` statement.

- **Range:** 1 through 255,000

`version`—Configure the BFD version to detect: `1` (BFD version 1) or `automatic` (autodetect the BFD version).

- **Default:** `automatic`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

Options `detection-time threshold` and `transmit-interval threshold` introduced in Junos OS Release 8.2.

Support for logical systems introduced in Junos OS Release 8.3.

Option `no-adaptation` introduced in Junos OS Release 9.0.

Options `authentication algorithm`, `authentication key-chain`, and `authentication loose-check` introduced in Junos OS Release 9.6.

Options `authentication algorithm`, `authentication key-chain`, and `authentication loose-check` introduced in Junos OS Release 9.6 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring BFD for RIP

Example: Configuring BFD Authentication for RIP

bfd-liveness-detection (Routing Options Static Route)

IN THIS SECTION

- Syntax | 2950
- Hierarchy Level | 2951
- Description | 2952
- Options | 2952
- Required Privilege Level | 2954
- Release Information | 2954

Syntax

```
bfd-liveness-detection {
    description Site- xxx;
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    echo
        minimum-interval milliseconds
    echo-lite
```

```
        minimum-interval milliseconds
    holddown-interval milliseconds;
    local-address ip-address;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-receive-ttl number;
    multiplier number;
    neighbor address;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name static route destination-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name static route destination-prefix qualified-next-hop (interface-
name | address)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options static route destination-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options static route destination-prefix qualified-next-hop (interface-name | address)],
[edit logical-systems logical-system-name routing-options rib routing-table-name static route
destination-prefix],
[edit logical-systems logical-system-name routing-options rib routing-table-name static route
destination-prefix qualified-next-hop (interface-name | address)],
[edit logical-systems logical-system-name routing-options static route destination-prefix],
[edit logical-systems logical-system-name routing-options static route destination-prefix
qualified-next-hop (interface-name | address)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
route destination-prefix],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
route destination-prefix qualified-next-hop (interface-name | address)],
[edit routing-instances routing-instance-name routing-options static route destination-prefix],
[edit routing-instances routing-instance-name routing-options static route destination-prefix
qualified-next-hop (interface-name | address)],
```

```
[edit routing-options rib routing-table-name static route destination-prefix],
[edit routing-options rib routing-table-name static route destination-prefix qualified-next-hop
(interface-name | address)],
[edit routing-options static route destination-prefix],
[edit routing-options static route destination-prefix qualified-next-hop (interface-name |
address)]
```

Description

Configure bidirectional failure detection timers and authentication criteria for static routes.

Options

authentication algorithm *algorithm-name* —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the `authentication-key-chains key-chain` statement at the `[edit security]` hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

echo—(Optional) Enable BFD `echo` mode. This feature sends BFD echo packets back and forth from a neighboring device to ensure that a forwarding path is available. BFD `echo` mode only works if the neighboring device supports BFD.

echo-lite—(Optional) Enable BFD `echo-lite` mode. This feature performs the same function as BFD `echo` mode but it works even if the neighboring device does not support BFD.

holddown-interval *milliseconds*—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

- **Range:** 0 through 255,000
- **Default:** 0

local-address ***ip-address***—Enable a multihop BFD session and configure the source address for the BFD session.

minimum-interval ***milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and `minimum-receive-interval` statements.

- **Range:** 1 through 255,000

minimum-receive-interval ***milliseconds***—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the `minimum-interval` statement at the `[edit routing-options static route` *`destination-prefix`* `bfd-liveness-detection]` hierarchy level.

- **Range:** 1 through 255,000

minimum-receive-ttl ***number***—Configure the time to live (TTL) for the multihop BFD session.

- **Range:** 1 through 255
- **Default:** 255

multiplier ***number***—Configure number of hello packets not received by the neighbor that causes the originating interface to be declared down.

- **Range:** 1 through 255
- **Default:** 3

neighbor ***address***—Configure a next-hop address for the BFD session for a next hop specified as an interface name.

no-adaptation—Specify for BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold ***milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

- **Range:** 0 through 4,294,967,295

transmit-interval minimum-interval ***milliseconds***—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the `minimum-`

`interval` statement at the `[edit routing-options static route destination-prefix bfd-liveness-detection]` hierarchy level.

- **Range:** 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

- **Default: automatic**

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

detection-time threshold and **transmit-interval threshold** options introduced in Junos OS Release 8.2.

`local-address` statement introduced in Junos OS Release 8.2.

`minimum-receive-ttl` statement introduced in Junos OS Release 8.2.

Support for logical routers introduced in Junos OS Release 8.3.

`holddown-interval` statement introduced in Junos OS Release 8.5.

`no-adaptation` statement introduced in Junos OS Release 9.0.

Support for IPv6 static routes introduced in Junos OS Release 9.1.

authentication algorithm, **authentication key-chain**, and `authentication loose-check` statements introduced in Junos OS Release 9.6.

`echo` and `echo-lite` statements introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

> *Example: Configuring BFD for Static Routes for Faster Network Failure Detection*
>
> *Example: Configuring BFD Authentication for Securing Static Routes*

bfd-liveness-detection (Source Packet Routing Template)

IN THIS SECTION

- Syntax | 2955
- Hierarchy Level | 2955
- Description | 2955
- Options | 2956
- Required Privilege Level | 2956
- Release Information | 2956

Syntax

```
bfd-liveness-detection {
    minimum-interval milliseconds;
    multiplier multiplier;
    no-router-alert-option;
    sbfd;
}
```

Hierarchy Level

```
[edit protocols source-packet-routing source-routing-path-template]
```

Description

Configure Bidirectional forwarding detection (BFD) options for PCE-initiated segment routing LSP template.

NOTE: The BFD configuration is applied at the top level and not individually per segment list. As a result, each path of the PCE-initiated LSP inherits the same BFD configuration.

Options

minimum-interval Specify in milliseconds the minimum transmit and receive interval.

- **Range:** 1 through 255000 milliseconds

multiplier Specify the detection time multiplier.

- **Default:** 3
- **Range:** 1 through 255

no-router-alert-option Do not set the router alert option in the IP header.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

PCEP Configuration

bgp

IN THIS SECTION

- Syntax | 2957
- Hierarchy Level | 2957
- Description | 2957
- Default | 2957
- Required Privilege Level | 2957

Release Information | 2957

Syntax

```
bgp { ... }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable BGP on the routing device or for a routing instance.

Default

BGP is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

BGP User Guide

bgp (RPM)

IN THIS SECTION

- Syntax | 2958
- Hierarchy Level | 2958
- Description | 2959
- Options | 2959
- Required Privilege Level | 2959
- Release Information | 2959

Syntax

```
bgp {
    data-fill data;
    data-size size;
    destination-port port;
    history-size size;
    logical-system logical-system-name <routing-instances routing-instance-name>;
    moving-average-size size;
    probe-count count;
    probe-interval seconds;
    probe-type type;
    routing-instances instance-name;
    test-interval interval;
    ttl;
}
```

Hierarchy Level

```
[edit services rpm bgp]
```

Description

Configure BGP neighbor discovery through Real-Time Performance Monitoring (RPM).

Options

`bgp`—Define properties for configuring BGP neighbor discovery.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: On MX Series routers, you can configure all the statements. On M Series and T Series routers, you can configure only the `logical-system` and `routing-instances` statements.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring BGP Neighbor Discovery Through RPM

bgp-error-tolerance (Protocols BGP)

IN THIS SECTION

- Syntax | 2960
- Hierarchy Level | 2960
- Description | 2960
- Options | 2960

Required Privilege Level | 2961
Release Information | 2961

Syntax

```
bgp-error-tolerance {
    (malformed-route-limit number | no-malformed-route-limit);
    malformed-update-log-interval seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address]
```

Description

Enable error handling for BGP update messages.

Starting in Junos OS Release 24.2R1, BGP error handling is enabled by default. You can still configure sub-options such as, `malformed-route-limit`, `malformed-update-log-interval`, and `no-malformed-route-limit` under this configuration statement. Note that If you delete the `bgp-error-tolerance` statement, the feature will still remain enabled and the sub-options are reset to their default values.

Options

malformed-route-limit *number*

Configure a limit on the number of malformed hidden routes stored in memory.

> **NOTE**: When the value of `malformed-route-limit` is reduced, only new malformed BGP update messages are affected and the existing malformed routes are retained.

- **Default:** 1000
- **Range:** 0 through 4294967295

malformed-update-log-interval *seconds*

Configure the duration for which the logging of malformed BGP update messages are suppressed.

On configuring the malformed update log interval:

1. The first malformed BGP update message is logged.
2. All subsequent malformed update messages are suppressed until the log interval expires.
3. On log interval expiry, the total number of malformed attributes received during the interval are logged.

This process repeats when the next malformed update message is received.

- **Default:** 300 seconds
- **Range:** 10 through 65535 seconds

no-malformed-route-limit

Disable the limit on the number of malformed hidden routes stored in memory.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding Error Handling for BGP Update Messages

Example: Configuring Error Handling for BGP Update Messages

bgp-identifier

IN THIS SECTION

- Syntax | 2962
- Hierarchy Level | 2962
- Description | 2962
- Required Privilege Level | 2962
- Release Information | 2963

Syntax

```
bgp-identifier identifier;
```

Hierarchy Level

```
[edit protocols bgp],
[edit protocols bgp group groupname],
[edit protocols bgp group groupname neighbor peer address ]
```

Description

Configure BGP using BGP identifier.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control— To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

bgp-orf-cisco-mode

IN THIS SECTION

- Syntax | 2963
- Hierarchy Level | 2963
- Description | 2964
- Default | 2964
- Required Privilege Level | 2964
- Release Information | 2964

Syntax

```
bgp-orf-cisco-mode;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp outbound-route-filter],
[edit logical-systems logical-system-name protocols bgp group group-name outbound-route-filter],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
outbound-route-filter],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
outbound-route-filter],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name outbound-route-filter],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address outbound-route-filter],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options outbound-route-filter],
```

```
[edit logical-systems logical-system-name routing-options outbound-route-filter],
[edit protocols bgp outbound-route-filter],
[edit protocols bgp group group-name outbound-route-filter],
[edit protocols bgp group group-name neighbor address outbound-route-filter],
[edit routing-instances routing-instance-name protocols bgp outbound-route-filter],
[edit routing-instances routing-instance-name protocols bgp group group-name outbound-route-
filter],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
outbound-route-filter],
[edit routing-instances routing-instance-name routing-options outbound-route-filter],
[edit routing-options outbound-route-filter]
```

Description

Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.

NOTE: To enable interoperability for all BGP peers configured on the routing device, include the statement at the `[edit routing-options outbound-route-filter]` hierarchy level.

Default

Disabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.

Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring BGP Prefix-Based Outbound Route Filtering

bgp-peer

IN THIS SECTION

- Syntax | 2965
- Hierarchy Level | 2965
- Description | 2965
- Options | 2965
- Required Privilege Level | 2966
- Release Information | 2966

Syntax

```
bgp-peer ip-address;
```

Hierarchy Level

```
[edit routing-instances name protocols evpn mclag]
```

Description

Configure an aggregation device in a Junos Fusion Enterprise or a multichassis link aggregation group (MC-LAG) topology to interwork with an Ethernet VPN-MPLS (EVPN-MPLS) device.

Options

ip-address IP address of the BGP peer. Typically, a BGP peer is identified by the IP address of the device's loopback interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG

bgp-static

IN THIS SECTION

- Syntax | 2966
- Hierarchy Level | 2967
- Description | 2967
- Options | 2967
- Required Privilege Level | 2967
- Release Information | 2967

Syntax

```
bgp-static {
    route destination-prefix/prefix-length {
        as-path <as-path> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator as-
number in-address>;
        community [ community-ids ];
        (metric | metric2 | metric3 | metric4) value <type type>;
        (preference | preference2 | color | color2) preference <type type>;
```

```
        (tag | tag2) metric type number;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-options rib routing-table-name],
[edit routing-options],
[edit routing-options rib routing-table-name]
```

Description

Specify a BGP-static route. You can specify any number of routes within a single BGP-static statement, and you can specify any number of BGP-static options in the configuration.

Options

`route` ***destination-prefix/prefix-length***	*destination-prefix* is the network portion of the IP address, and *prefix-length* is the destination prefix length.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

advertise-bgp-static | 2009

Configuring BGP-Static Routes for Preventing Route Flaps

Example: Configuring BGP-Static Routes to Prevent Route Flaps

bgpmcast

IN THIS SECTION

- Syntax | 2968
- Hierarchy Level | 2968
- Description | 2968
- Required Privilege Level | 2969
- Release Information | 2969

Syntax

```
bgpmcast {
  traceoptions {
     file <file name>;
     flag all detail;
  }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Enable this protocol to accept and carry BGP multicast network layer reachability information (NLRI) and program forwarding states on routers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1

RELATED DOCUMENTATION

Controller-Based BGP Multicast Signaling

accept-prpd-connection | **1639**

inet-bgpmcast | **8557**

inet6-bgpmcast | **8599**

show route table | **23789**

show multicast route | **22306**

bidirectional (Interface)

IN THIS SECTION

- Syntax | **2969**
- Hierarchy Level | **2970**
- Description | **2970**
- Required Privilege Level | **2970**
- Release Information | **2970**

Syntax

```
bidirectional {
    df-election {
```

```
        backoff-period milliseconds;
        offer-period milliseconds;
        robustness-count number;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name],
[edit protocols pim interface (Protocols PIM)  interface-name],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name]
```

Description

Configure parameters for bidirectional PIM.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

> *Understanding Bidirectional PIM*
>
> *Example: Configuring Bidirectional PIM*

bidirectional (RP)

IN THIS SECTION

- Syntax | 2971
- Hierarchy Level | 2971
- Description | 2971
- Required Privilege Level | 2972
- Release Information | 2972

Syntax

```
bidirectional {
    address address {
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
        hold-time seconds;
        priority number;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Configure the routing device's rendezvous-point (RP) properties for bidirectional PIM.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Bidirectional PIM

Example: Configuring Bidirectional PIM

bier

IN THIS SECTION

- Syntax | 2972
- Hierarchy Level | 2973
- Description | 2973
- Options | 2973
- Default | 2974
- Required Privilege Level | 2974
- Release Information | 2974

Syntax

```
bier {
   sub-domain sub-domain-id {
       bfr-id bfr-id;
```

```
            bfr-prefix bfr-prefix;
            encapsulation {
                          mpls;
                          bitstringlen 256;
                          number-sets number of sets;
            }
        }
    }
```

Hierarchy Level

```
[edit protocols]
```

Description

Enable Bit Index Explicit Replication (BIER) on the routing device.

You must include the `bier` statement in the configuration to enable BIER on the router.

Options

sub-domain Assign a numeric value to the sub-domain. It represents a specific topology within a BIER domain.

- **Range:** 0 through 255.

bfr-id A BFR-ID is a number assigned to an ingress or egress BFR. A BFR-ID must be unique within that BIER sub-domain.

- **Range:** 1 through 1024.

bfr-prefix A BFR prefix is typically the configured loopback address of a BFR and must be a routable IP address in the BIER domain. Mandatory.

encapsulation Specify encapsulation type.

- mpls—Currently Junos only supports MPLS encapsulation.

bitstringlen Configure bit string length.

- Currently Junos only supports a bit string length of 256.

number-sets Specify the number of BIER sets.

- **Default:** 4
- **Range:** 1 through 4

Default

BIER is disabled on the routing device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos Evolved OS Release 23.4R2-S1.

bier-sub-domain (Protocols IS-IS)

IN THIS SECTION

- Syntax | 2974
- Hierarchy Level | 2975
- Description | 2975
- Options | 2975
- Required Privilege Level | 2975
- Release Information | 2975

Syntax

```
bier-sub-domain sub-domain-id;
```

Hierarchy Level

```
[edit protocols isis]
```

Description

To include the BIER sub-domain in the IS-IS sub-TLVs, configure `bier-sub-domain` under `protocols isis`. Key BIER information like BFR-IDs and BFER-Prefixes in each sub-domain are flooded through the IS-IS domain to generate the BIER forwarding table.

Options

name Enable bier subdomain-id

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R2-S1.

bits (Chassis Synchronization)

IN THIS SECTION

- Syntax | 2976
- Syntax (ACX710: Interfaces Hierarchy) | 2976
- Syntax (ACX710: Output Interfaces Hierarchy) | 2976
- Hierarchy Level | 2977
- Description | 2977
- Options | 2977

Required Privilege Level | 2977
Release Information | 2977

Syntax

```
bits {
    signal-type (e1 | t1);
    e1-options {
        framing (g704 | g704-no-cro4);
    }
    t1-options {
        framing (esf | sf);
    }
}
```

Syntax (ACX710: Interfaces Hierarchy)

```
bits {
    signal-type (e1 | t1);
    e1-options {
        framing (g704 | g704-no-cro4);
    }
    t1-options {
        framing (esf | sf);
    }
}
```

Syntax (ACX710: Output Interfaces Hierarchy)

```
bits {
    source-mode (chassis | ptp-clock);
    squelch {
        method (ais | ql-dnu | rai | shut);
    }
```

```
    ssm option-2  (GEN-1 | GEN-2);
}
```

Hierarchy Level

```
[edit chassis synchronization interfaces]
[edit chassis synchronization output interfaces]
```

Description

Configure the parameters of the building-integrated timing supply (BITS) interface.

NOTE: `bits` is not supported on the ACX1000 router.

Options

ssm option-2 (GEN-1 | GEN-2) (ACX710 only) Specify the type of synchronization status message (SSM). Choose message type Generation 1 or Generation 2.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

`source-mode`, `squelch`, and `ssm` options introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

External Clock Synchronization Overview for ACX Series Routers

Configuring External Clock Synchronization for ACX Series Routers

Synchronous Ethernet Overview

show chassis synchronization | **18879**

black-list

IN THIS SECTION

- Syntax | **2978**
- Hierarchy Level | **2978**
- Description | **2978**
- Options | **2979**
- Required Privilege Level | **2979**
- Release Information | **2979**

Syntax

```
black-list {
          IP-address-list IP-address-list-name;
           file-name file-name;
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name security-intelligence-policy]
```

Description

Configure a global blocklist of IP addresses to block packets to the destined countries.

Options

- `IP-address-list` *`ip-address-list-name`* —Contains the list of IP addresses that must be blocked. *ip-address-list-name* is the name of the IP address-list specified at the `[edit services web-filter]` hierarchy.
- `file-name` *`file-name`*—Name of the file which contains the IP addresses that must be blocked. The file is in the **/var/db/url-filterd** location and has the same name as in the configuration.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1 on MX Series routers with Juniper Advanced Threat Prevention Cloud (Juniper ATP Cloud) .

RELATED DOCUMENTATION

URL Filtering

blackhole-action

IN THIS SECTION

- Syntax | 2980
- Hierarchy Level | 2980
- Description | 2980
- Required Privilege Level | 2980
- Release Information | 2980

Syntax

```
blackhole-action (log-only | restart | offline | restart-then-offline);
```

Hierarchy Level

```
[edit chassis fpc slot-number fabric]
```

Description

Configure an FPC to take a specific action when fabric plane degradation reaches 100 percent.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Managing Bandwidth Degradation

bandwidth-degradation | **2841**

block

IN THIS SECTION

- Syntax | **2981**
- Hierarchy Level | **2981**

- Description | 2981
- Required Privilege Level | 2982
- Release Information | 2982

Syntax

```
block block-name {
      address;
      global-micro-sid{
            maximum-sids maximum-sids;
            maximum-static-sids maximum-static-sids
      }
      local-micro-sid{
            maximum-static-sids maximum-static-sids
      }
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options source-packet-routing srv6],
[edit logical-systems name routing-options source-packet-routing srv6],
[edit logical-systems name tenants name routing-instances name routing-options source-packet-
routing srv6],
[edit routing-instances name routing-options source-packet-routing srv6],
[edit routing-options source-packet-routing srv6],
[edit tenants name routing-instances name routing-options source-packet-routing srv6]
```

Description

Specify an SRv6 prefix that defines a micro SID block. Configure micro SID block parameters such as global and local micro SIDs. In segment routing domain, a global micro SID. identifies a shortest-path to a node. A local micro SID identifies a cross connect to a direct neighbor over a specific interface.

name Specify a block name.

address Specify a block prefix IPv6 address for segment routing. Default block-length is 32 bits, block-length of 32 only is allowed on the ACX Series.

global-micro-sid (Optional) Specify the size of global micro SIDs.Global SIDs are static by default. Configure the global micro SID block range parameters such as the micro SID length. The length of a micro SID is 16 by default.

- maximum-sids : Specify the maximum number of SIDs used for global block
  - Range: 0 through 65536
  - Default: 57343
- maximum-static-sids: Specify the maximum number of static SIDs used for a global block (0..65536). Note that currently we do not support dynamic global micro SIDS.
  - Range: 0 through 65536
  - Default: If you do not configure the maximum static SIDs, then maximum-sids value is used.

local-micro-sid (Optional) the local micro SID block range parameters.

- maximum-static-sids: Specify the size of the local micro SID block.
- Default: 0

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.2R1 and Junos OS 23.4R1.

RELATED DOCUMENTATION

micro-sid | **10697**

locator | **9997**

No Link Title

block-allocation (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 2983
- Hierarchy Level | 2983
- Description | 2983
- Options | 2984
- Required Privilege Level | 2984
- Release Information | 2984

Syntax

```
block-allocation {
    active-block-timeout timeout-interval;
    block-size block-size;
    interim-logging-interval timeout-interval;
    maximum-blocks-per-host maximum-block-number
    log disable | enable
}
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Description

Allocate a block of ports for each subscriber to use for source NAT with port translation, except for deterministic NAT. New requests for NAT ports for the subscriber are served from the active block.

With port block allocation, we generate one syslog log per set of ports allocated for a subscriber. This reduces the number of logs, making it easier to track subscribers.

Options

active-block-timeout *timeout-interval*

The interval, in seconds, for which the block is active. After the timeout, a new block is allocated, even if ports are available in the active block. If you set the timeout to 0, port blocks are filled completely before a new port block is allocated, and the last port block remains active indefinitely.

- **Range:** 0 through 86,400
- **Default:** 0

block-size *block-size*

Number of ports in a block.

- **Range:** 1 through 64,512
- **Default:** 128

interim-logging-interval *timeout-interval*

The interval, in seconds, at which to send interim system logs for active port blocks and for inactive port blocks with live sessions. This increases the reliability of system logs, which are UDP-based and can get lost in the network.

- **Range:** 1800 through 86,400
- **Default:** 0 (interim logs are disabled)

maximum-blocks-per-host *maximum-block-number*

The maximum number of blocks that can be allocated to a subscriber address.

- **Range:** 1 through 512
- **Default:** 8

log disable

Disable logs for port block allocation. Logs are enbled by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

block-command

IN THIS SECTION

- Syntax | 2985
- Hierarchy Level | 2985
- Description | 2985
- Required Privilege Level | 2985
- Release Information | 2985

Syntax

```
block-command protocol-command-list;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name]
```

Description

Apply protocol block command custom-objects to the content-filtering profile.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

block-content-type

IN THIS SECTION

- Syntax | 2986
- Hierarchy Level | 2986
- Description | 2986
- Options | 2986
- Required Privilege Level | 2987
- Release Information | 2987

Syntax

```
block-content-type (activex | exe | http-cookie | java-applet | zip);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name]
```

Description

Apply blocks to other available content such as exe, http-cookie, java-applet. This is for HTTP only.

Options

- `activex`—Block ActiveX.
- `exe`—Block EXE files.
- `http-cookie`—Block cookies.
- `java-applet`—Block Java applets.
- `zip`—Block ZIP files.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

block-extension

IN THIS SECTION

- Syntax | 2987
- Hierarchy Level | 2987
- Description | 2988
- Required Privilege Level | 2988
- Release Information | 2988

Syntax

```
block-extension extension-list;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name]
```

Description

Apply block extensions to the content-filtering profile.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

block-frag (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 2988
- Hierarchy Level | 2988
- Description | 2989
- Required Privilege Level | 2989
- Release Information | 2989

Syntax

```
block-frag;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop fragmented IP packets. IP fragments might contain an attacker's attempt to exploit the vulnerabilities in the packet reassembly code of specific IP stack implementations. When the target receives these packets, the results can range from processing the packets incorrectly to crashing the entire system.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

block-message (Application Firewall)

IN THIS SECTION

- Syntax | 2990
- Hierarchy Level | 2990
- Description | 2990
- Options | 2991
- Required Privilege Level | 2992
- Release Information | 2992

Syntax

```
block-message type {
    custom-text content custom-html-text;
    custom-redirect-url content custom-redirect-url;
}
```

Hierarchy Level

```
[edit security application-firewall profile profile-name]
```

Description

Defines the profile of the notification to be sent to clients when HTTP or HTTPS traffic is blocked by a reject or deny action from an application firewall.

NOTE: The block message option is not supported for non-HTTP traffic such as FTP, SSH, Telnet, and so on. In these instances, if the action is drop or reject, the traffic is silently dropped or rejected. The user is not informed of the action and no redirection occurs. The associated system log message identifies the action taken for this traffic.

The reject or deny message actions are logged with the reason field containing one of the following phrases:

- appfw deny
- appfw reject

Following sample shows a system log message for SSH traffic, where the traffic was rejected:

```
RT_FLOW_SESSION_DENY [junos@2636.1.1.1.2.134 source-address="1.2.0.100" source-
port="53540" destination-address="1.1.0.100" destination-port="22" connection-tag="0"
service-name="junos-ssh" protocol-id="6" icmp-type="0" policy-name="p1" source-zone-
name="untrust" destination-zone-name="trust" application="SSH" nested-
application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-interface="reth2.0"
encrypted="No" reason="appfw reject"]
```

> **NOTE**: You need to enable SSL forward proxy for the HTTPS traffic that needs to be blocked by a reject or a deny action from an application firewall.

When the `block-message` option is specified, a splash screen and message inform the client that the traffic has been blocked. The default message text is:

```
"username, Application Firewall has blocked your request to application application-name at dest-ip:dest-port  accessed from src-ip:source-port "
```

The variables in the message are replaced with specific traffic values. For clarity, the prefix **junos:** is truncated from the application name.

> **NOTE**: You need to enable SSL forward proxy for the HTTPS traffic,that needs to be blocked by a reject or a deny action from an application firewall.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Options

Use the following option pairs to customize the default message or to redirect the client to a custom webpage instead of the default splash screen.

> **NOTE**: Both the `type` and `content` fields must be used to add custom text or redirect the client to a URL.

- `type`—(Optional) The message type to be displayed after a reject or deny action.
  - `custom-text`—Text message in HTML to be added to the default text. If `custom-text` is specified, the splash screen displays both the default block message and the custom-defined block message.

    When specified, the user is redirected when a reject or deny action is taken during one of the following HTTP methods: GET, POST, OPTIONS, HEAD, PUT, DELETE, TRACE, CONNECT, PROPFIND, PROPPATCH, LOCK, UNLOCK, COPY, MOVE, MKCOL, BCOPY, BDELETE, BCOPY, BMOVE, BPROPFIND, BPROPPATCH, POLL, SEARCH, SUBSCRIBE, and UNSUBSCRIBE. If the reject or deny action occurs during a different HTTP method, the traffic is silently dropped.

- `custom-redirect-url`—URL redirection.

- `content`—(Optional) Message content for the selected message type.

> **NOTE**: The `content` value must match the `type` option selected: `custom-text` requires text, and `custom-redirect-url` requires a URL value.

  - *custom-text*—Custom text to be added to the splash screen. Custom text is inserted below the default message. Add the characters \n to insert a line break in the displayed text.

  - *custom-redirect-url*—The URL of the webpage to which the client is directed. When traffic is rejected or denied, the client is redirected to the specified webpage for further action. The URL can be hosted on either the SRX Series Firewall or an external server.

    Enter the redirect URL in quotation marks for an HTTP or HTTPS site, as shown in the following examples:

```
"http://custom-redirect-url"
"https://custom-redirect-url"
```

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Example: Configuring AppQoS

block-message (Security Content Security)

IN THIS SECTION

- Syntax | 2993
- Hierarchy Level | 2993
- Description | 2993
- Options | 2993
- Required Privilege Level | 2994
- Release Information | 2994

Syntax

```
block-message {
    type {
        custom-redirect-url;
    }
    url url;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name]
```

Description

Configure Juniper enhanced block message settings.

Options

- `type`—Specify the following type of the block message:
  - `custom-redirect-url`—Specify Custom redirect URL server.

- `url` *`url`*—Specify an URL of the block message.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

block-mime

IN THIS SECTION

- Syntax | 2994
- Hierarchy Level | 2995
- Description | 2995
- Options | 2995
- Required Privilege Level | 2995
- Release Information | 2995

Syntax

```
block-mime {
    exception list-name;
    list list-name;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name]
```

Description

Apply MIME pattern list custom-objects to the content-filtering profile for blocking MIME types.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

block (Spanning Trees)

IN THIS SECTION

- Syntax | 2996
- Hierarchy Level | 2996
- Description | 2996
- Required Privilege Level | 2996
- Release Information | 2996

Syntax

```
block;
```

Hierarchy Level

```
[edit protocols mstp interface (all | interface-name) bpdu-timeout-action],
[edit protocols rstp interface (all | interface-name) bpdu-timeout-action],
[edit protocols stp interface (all | interface-name) bpdu-timeout-action],
[edit protocols vstp vlan vlan-id interface (all | interface-name) bpdu-timeout-action]
```

Description

Configures loop protection on a specific interface.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

RELATED DOCUMENTATION

- Understanding Loop Protection for STP, RSTP, VSTP, and MSTP
- Understanding VSTP
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
- Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree
- *Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on non-ELS EX Series Switches*

show spanning-tree bridge | 26548
show spanning-tree interface | 26555

blocklists (Security IKE)

IN THIS SECTION

- Syntax | 2997
- Hierarchy Level | 2998
- Description | 2998
- Options | 2998
- Required Privilege Level | 2999
- Release Information | 2999

Syntax

```
    blocklists {
        blocklist-name {
            description text-description;
            rule rule-name {
            match {
                    role (initiator | responder);
                    id-type (inet | inet6 | hostname | distinguished-name | user-at-hostname |
key-id);
                    id-pattern value;
             }
                then {
                    (discard | reject);
                    backoff timeout-value;
                }
            }
      }
  }
```

Hierarchy Level

```
[edit security ike]
```

Description

Defines the remote peer IKE identity blocklist(s) configuration.

In the configuration hierarchy, you can configure different blocklists for blocking IKE IDs during IKE SA negotiation authentication phase. Once you configure the blockslists, you can use a blocklist in the corresponding IKE policy at the `[edit security ike policy policy-name blocklist blocklist-name]` hierarchy level.

Options

blocklists *blocklist-name*

Specify the blocklist name.

- Value: upto 32 characters
- Max allowed: Upto 1000 blocklists

rule *rule-name*

Specify the rule name.

- Value: upto 32 characters
- Max allowed: Upto 1000 rules per blocklist

description *text-description*

(Optional) Specify the text description.

match

Match criteria.

id-type (inet | inet6 | hostname | distinguished-name | user-at-hostname | key-id)

Specify the remote peer IKE identity type with one of the following:

- `inet`—Specify to use the IPv4 address.
- `inet6`—Specify to use the IPv6 address.
- `hostname` —Specify to use the hostname.
- `distinguished-name`—Specify to use the distinguished name.
- `user-at-hostname`—Specify to use the email address
- `key-id`—Specify to use the Key ID

id-pattern *<value>*

Specify a valid remote peer IKE identity value based on the `id-type` selection. Use the standard regular expression for the value.

role (initiator| responder)

Specify the remote peer role as the `initiator` or the `responder`.

- Default: `initiator`

then

Action for the match criteria.

discard

Discard the peer connection immediately without sending back a response.

reject

Discard the connection immediately and send back a failure response.

backoff *timeout-value*

(Optional) Discard the connection immediately and set a backoff timer. If you do not configure the option explicitly, the `auth-phase-failure` *value* option at the [`edit security ike session half-open backoff-timeouts`] hierarchy level, enables backoff action.

- Value: Range is 1-180 seconds.
- Default: 60 seconds

For other statements, see CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

`blocklists` statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

IKE Protection from DDoS Attacks

policy (Security IKE) | **12713**

session (Security IKE) | **16848**

bmp

IN THIS SECTION

- Syntax | 3000
- Hierarchy Level | 3002
- Description | 3002
- Options | 3003
- Required Privilege Level | 3008
- Release Information | 3008

Syntax

```
bmp {
    authentication-algorithm algorithm;
    authentication-key key;
    authentication-key-chain authentication-key-chain;
    connection-mode (active | passive);
    hold-down {
        seconds;
        flaps flaps;
        period seconds;
    }
    initiation-message text;
    local-address address;
    local-port port;
    max-loc-rib-buffer-count count;
    monitor (disable | enable);
    priority (high | low | medium);
    route-monitoring {
        none;
        loc-rib;
        post-policy {
            exclude-non-eligible;
        }
        pre-policy {
            exclude-non-feasible;
```

```
        }
        rib-out {
            post-policy;
            pre-policy;
        }
    }
    routing-instance routing-instance-name;
    station station-name {
        authentication-algorithm algorithm;
        authentication-key key;
        authentication-key-chain authentication-key-chain;
        connection-mode (active | passive);
        hold-down {
            seconds;
            flaps flaps;
            period seconds;
        }
        initiation-message text;
        in-memory-profiling;
        keep-advertisement-state;
        local-address address;
        local-port port;
        max-loc-rib-buffer-count count;
        monitor (disable | enable);
        priority (high | low | medium);
        routing-instance routing-instance-name;
        route-monitoring {
            none;
            loc-rib;
            post-policy {
                exclude-non-eligible;
            }
            pre-policy {
                exclude-non-feasible;
            }
            rib-out {
                post-policy;
                pre-policy;
            }
                }
        station-address (ip-address | name);
        station-port port-number;
        statistics-timeout seconds;
```

```
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier>;
        }
    }
    station-address (ip-address | name);
    station-port port-number;
    statistics-timeout seconds;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier>;
    }
}
```

Hierarchy Level

```
[edit routing-options],
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address]
```

NOTE: 1. Complete BMP configuration, as mentioned in the syntax, can be done under the first two hierarchy levels only

2. Under other hierarchy levels, only the following configurations are supported:

- Either we can inherit or not inherit the configuration data
- Enable/disable monitoring
- Control route monitoring settings

Description

Configure the BGP Monitoring Protocol (BMP), which enables the routing device to collect data from the BGP Adjacency-RIB-In routing tables and periodically send that data to a monitoring station. The

`Adjacency-RIB-In` tables are the pre-policy tables, meaning that the routes in these tables have not been filtered or modified by routing policies.

BGP Monitoring Protocol (BMP) allows the Junos OS to send the BGP route information from the router to a monitoring application on a separate device. The monitoring application is called the BMP monitoring station or BMP station. To deploy BMP in your network, you need to configure BMP on each router and you also need to configure at least one BMP station.

> **NOTE**: When BMP is configured at multiple hierarchy levels, the order of preference from highest to lowest is as follows:
>
> 1. `[edit protocols bgp group group-name neighbor address bmp]`
> 2. `[edit protocols bgp group group-name bmp]`
> 3. `[edit protocols bgp bmp]`
> 4. `[edit routing-optionsbmp station station-name]`
> 5. `[edit routing-options bmp]`
>
> For example, if BMP is configured at both `[edit routing-options bmp]` and `[edit protocols bgp bmp]` hierarchy levels, the configuration at the protocols BGP level takes precedence over the routing options configuration.

Options

authentication-algorithm *algorithm*

Configure an authentication algorithm.

- **Values:** Specify one of the following types of authentication algorithms:
  - `aes-128-cmac-96`—Cipher-based message authentication code (AES128, 96 bits).
  - `hmac-sha-1-96`—Hash-based message authentication code (SHA1, 96 bits).
  - `md5`—Message digest 5.
- **Default:** `hmac-sha-1-96`

> **NOTE**: The default is not displayed in the output of the `show bgp bmp` command unless a key or key-chain is also configured.

authentication-key *key*

Configure an MD5 authentication key (password), which can be up to 126 characters. Characters can include any ASCII strings. If you include spaces,

enclose all characters in quotation marks (“ ”). Neighboring routing devices use the same password to verify the authenticity of BMP packets sent from this system.

authentication-key-chain *key-chain*

Apply and enable an authentication key chain to the routing device, which can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

> **NOTE**: You must also configure an authentication algorithm by including the `authentication-algorithm` *`algorithm`* statement.

Note that the referenced key chain must be defined. When configuring the authentication key update feature for BGP, you cannot commit the `0.0.0.0/allow` statement with authentication keys or key chains. The CLI issues a warning and fails to commit the configuration.

connection-mode (active | passive)

Specify whether the BMP station connection is `active` or `passive`. If you configure the connection-mode statement as `active`, do not also configure the `local-port` statement. If you configure the connection-mode statement as `passive`, you must configure the `local-port` statement.

If the connection-mode statement is configured as `active`, a station port number is required. If the `connection-mode` statement is configured as `passive`, you must not configure a station port number.

- **Values:** Specify one of the following:
  - active—BMP initiates the connection to the BMP station.
  - passive—BMP does not initiate a connection to the BMP station. However, it does listen for a connection request from active BMP stations and will connect if a station is available.

initiation-message *text*

(Optional) Specify a character string for a type 0 TLV to send to the BMP station with the initiation message. The message is transmitted when a BMP station establishes a connection to the device. You can provide some information to the BMP station system administrator (for example, a contact phone number). The initiation message includes a type 1 TLV containing the SNMP sysDescr value specified in RFC 1213 *Management Information Base for Network Management of TCP/IP-based internets: MIB-II* and a type 2 TLV containing the SNMP sysName value also from RFC 1213. The string in the initiation-message message is UTF-8.

The normal time for sending an initiation message is when the BMP session is first established. However, an initiation message change also triggers the transmission of an initiation message to current BMP sessions.

Another event that triggers the transmission of an initiation message is when you change in the sysName or sysDescr values in the SNMP configuration. The initiation message is sent to current BMP sessions.

- **Range:** 1 through 255 characters

in-memory-profiling

(Optional) Configure this option to collect time series data. Note that collection of time series data uses additional memory resources. Disable this option to free the memory used to store the time series data.

- **Default:** in-memory-profiling is disabled by default to ensure efficient memory utilization.

keep-advertisement-state

(Optional) Enable this option to store BMP advertisement state for rib-out monitoring. Note that this option can use up valuable resources. Disable this option to free the memory used to store BMP advertisement state. We recommend toggling the BMP session after changing deleting this option to have a cosistent view for all prefixes advertised to the station..

- **Default:** keep-advertisement-state is disabled by default to ensure efficient memory utilization.

local-address *address*

(Optional) Specifies the IPv4 or IPv6 address for the BMP connection on the device. We recommend that you configure a local address. For both active and passive connection modes, configure a loopback local address. This provides a consistent local endpoint, is useful for debugging, and assures greater reliability for the BMP connection since it is not tied to a single router interface.

For passive mode, specifying a local address is required. It also provides some security against a malicious BMP connection. For active mode, we also recommend configuring a local address to help ensure reliability.

If you change the local address, the BMP station connection flaps when you commit the configuration.

local-port *port*

Specify the listening port for the BMP station connection.

If you configure the connection-mode statement as `active`, do not configure the `local-port` statement. If you configure the connection-mode statement as `passive`, you must configure `local-port` statement.

If you change the local port, the BMP station connection flaps when you commit the configuration.

- **Range:** 1 through 65,535

max-loc-rib-buffer-count *count*

Specify the maximum number of local-rib outstanding buffers before blocking.

- **Range:** 1 through 50

monitor (enable | disable)

(Optional) Explicitly enable or disable BMP monitoring of BGP peers. You can also selectively enable or disable BMP monitoring at various hierarchy levels (for example, [`edit protocols bgp group` *`group-name`*] or [`edit protocols bgp group` *`group-name`* `neighbor` *`address`*]). If you disable BMP monitoring, withdrawal messages are sent for any previously advertised routes. These are followed by a down message. If you enable BMP monitoring, an up message is sent first and then the route advertisements follow.

- **Default:** BMP monitoring is enabled by default.
- **Values:** Specify one of the following:
  - disable—Disable BMP monitoring.
  - enable—Enable BMP monitoring.

priority (high | medium | low)

(Optional) Specify the dispatch priority for BMP. The dispatch priority controls the frequency with which the device is able to forward BMP messages to BMP stations.

> **NOTE**: Specifying `high` or `medium` priority may reduce the performance of the routing protocol process in its handling of route convergence or other work.

- **Default:** The default dispatch priority is `low`, to minimize interference with other routing protocol process priorities and to match the behavior of previous versions of BMP.
- **Values:** Specify one of the following:
  - high—The routing protocol process handles BMP requests with high urgency.
  - medium—The routing protocol process handles BMP requests with medium urgency.
  - low—The routing protocol process handles BMP requests with low urgency.

routing-instance *routing-instance-name*

Specify the name of the routing instance you want the BGP Monitoring Protocol (BMP) to use. This can be any routing instance name. If you want to use the reserved non-default management routing instance mgmt_junos, make sure you configure the `management-instance` statement.

When `routing-instance` is configured at both hierarchy levels— `[edit routing-options bmp station station-name]` and `[edit routing-options bmp]`— the configuration at the `[edit routing-options bmp station station-name]` hierarchy level takes precedence.

> **NOTE**: You must also configure the routing instance you want to use under the `[edit routing-instances]` hierarchy level.

- Values:
  - default—Default routing instance
  - routing-instance-name—Name of the routing instance

station *station-name*

(Required) Specify a name for the BMP monitoring station. Each station can use a significant amount of a device's resources. You can configure up to three BMP monitoring stations.

station-address (*address* | *station-name*

Specify the address or name for the BMP monitoring station. You can specify one or the other but not both.

- **Values:** Specify one of the following:
  - *station-address*—Specify the address for the BMP station. The address should be a valid IPv4 or IPv6 address.
  - *station-name*—Specify the name for the BMP station.

station-port *port*

Specify the port number for the BMP monitoring station. If the connection-mode statement is configured as `active`, a station port number is required. If the `connection-mode` statement is configured as `passive`, you must not configure a station port number.

- **Range:** 1 though 65535

statistics-timeout *seconds*

(Optional) Specify how often statistics messages are sent to the BMP monitoring station. If you configure a value of 0, no statistics messages are sent.

- **Default:** 3600 seconds
- **Range:** 15 though 65535 seconds

The remaining statements are explained separately. See CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.

Support for BMP version 3 introduced in Junos OS Release 13.3.

`initiation-message`, `local-address`, `local-port`, `monitor`, `priority`, `route-monitoring`, `station`, `station-address`, `station-port`, and `statistics-timeout` options introduced in Junos OS Release 13.3.

`routing-instance` option introduced in Junos OS Release 18.3R1.

`rib-out` option introduced in Junos OS Release 19.1R1.

`loc-rib` option introduced in Junos OS Release 19.2R1.

`in-memory-profiling` and `keep-advertisement-state` options introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Example: Configuring the BGP Monitoring Protocol

Example: Configuring Router Authentication for BGP

Configuring BGP Monitoring Protocol Version 3

management-instance | **10309**

Management Interface in a Non-Default Instance

Configuring BGP Monitoring Protocol to Run Over a Different Routing Instance

boot-loader-authentication

IN THIS SECTION

Syntax | **3009**

- Hierarchy Level | 3009
- Description | 3009
- Options | 3010
- Required Privilege Level | 3010
- Release Information | 3010

Syntax

```
boot-loader-authentication {
    (encrypted-password password | plain-text-password);
}
```

Hierarchy Level

```
[edit system]
```

Description

Set the boot-loader password for accessing the U-Boot CLI during the boot process. The password can be entered either as a plain-text password or as an encrypted password.

Encrypted passwords must be entered in Message Digest 5 (MD5) format. Plain-text passwords are encrypted by using MD5 by default. The encryption format for plain-text passwords can be changed by using the `set system login password format` command.

Encrypted passwords must be between 1 and 128 characters long. The password must be enclosed in quotation marks and cannot be blank within the quotation marks (“ “).

The default requirements for plain-text passwords are as follows:

- The password must be between 6 and 128 characters long
- You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
- The password must contain at least one change of case or character class.

Options

`encrypted-password` *password*— Enter a password that has already been encrypted. You can specify only one encrypted password.

`plain-text-password`—Enter a plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D20.

RELATED DOCUMENTATION

Using Unattended Mode for U-Boot to Prevent Unauthorized Access

password (Login) | **12235**

unattended-boot | **30326**

boot-server (DHCP)

IN THIS SECTION

- Syntax | **3011**
- Hierarchy Level | **3011**
- Description | **3011**
- Options | **3011**
- Required Privilege Level | **3011**
- Release Information | **3011**

Syntax

```
boot-server (address | hostname);
```

Hierarchy Level

```
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]
```

Description

Configure the name of the boot server advertised to DHCP clients. The client uses a boot file located on the boot server to complete DHCP setup.

Options

- *address*—IP address of a DHCP boot server.
- *hostname*—Hostname of a DHCP boot server.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

boot-file

bootp

IN THIS SECTION

- Syntax | 3012
- Hierarchy Level | 3013
- Description | 3013
- Options | 3013
- Required Privilege Level | 3013
- Release Information | 3013

Syntax

```
bootp {
    client-response-ttl number;
    description text-description;
    interface (interface-name | interface-group) {
        client-response-ttl number;
        description text-description;
        maximum-hop-count number;
        minimum-wait-time seconds;
        no-listen;
        server address {
            logical-system logical-system-name <routing-instance [ <default> routing-instance-
names ]>;
            routing-instance [ <default> routing-instance-names ];
        }
    }
    maximum-hop-count number;
    minimum-wait-time seconds;
    relay-agent-option;
    server address {
        <logical-system logical-system-name> <routing-instance
 [ routing-instance-names ]>;
    }
}
```

Hierarchy Level

```
[edit forwarding-options helpers]
```

Description

Configures a router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent. For MX Series (MX80, MX240, MX480 and MX960) routers connected via IRB, see *Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents* for information on how to prevent BOOTP reply packets from being dropped.

DHCP relaying is disabled.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

bootp (System Processes)

IN THIS SECTION

Syntax | 3014

- Hierarchy Level | 3014
- Description | 3014
- Options | 3014
- Required Privilege Level | 3015
- Release Information | 3015

Syntax

```
bootp {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the booting process.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable` —Disable the booting process.
- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.

- `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

bootstrap

IN THIS SECTION

- Syntax | 3015
- Hierarchy Level | 3016
- Description | 3016
- Required Privilege Level | 3016
- Release Information | 3016

Syntax

```
bootstrap {
    family (inet | inet6) {
        export [ policy-names ];
        import [ policy-names ];
        priority number;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Configure parameters to control bootstrap routers and messages.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring PIM Bootstrap Properties for IPv4

Configuring PIM Bootstrap Properties for IPv4 or IPv6

bootstrap-export

IN THIS SECTION

- Syntax | 3017
- Hierarchy Level | 3017

- Description | 3017
- Options | 3017
- Required Privilege Level | 3017
- Release Information | 3017

Syntax

```
bootstrap-export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Apply one or more export policies to control outgoing PIM bootstrap messages.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring PIM Bootstrap Properties for IPv4

Configuring PIM Bootstrap Properties for IPv4 or IPv6

bootstrap-import | 3018

bootstrap-import

IN THIS SECTION

- Syntax | 3018
- Hierarchy Level | 3018
- Description | 3018
- Options | 3019
- Required Privilege Level | 3019
- Release Information | 3019

Syntax

```
bootstrap-import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Apply one or more import policies to control incoming PIM bootstrap messages.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring PIM Bootstrap Properties for IPv4

Configuring PIM Bootstrap Properties for IPv4 or IPv6

bootstrap-export | 3016

bootstrap-priority

IN THIS SECTION

- Syntax | 3020
- Hierarchy Level | 3020
- Description | 3020
- Options | 3020
- Required Privilege Level | 3020
- Release Information | 3020

Syntax

```
bootstrap-priority number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Configure whether this routing device is eligible to be a bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router.

Options

number—Priority for becoming the bootstrap router. A value of 0 means that the routing device is not eligible to be the bootstrap router.

- **Range:** 0 through 255
- **Default:** 0

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring PIM Bootstrap Properties for IPv4*

bounded-delay

IN THIS SECTION

- Syntax | 3021
- Hierarchy Level | 3021
- Description | 3021
- Default | 3022
- Required Privilege Level | 3022
- Release Information | 3022

Syntax

```
bounded-delay;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name mka (MX Series)]
```

Description

Configure bounded delay to ensure that a Media Access Control Security (MACsec) frame will not be delivered after a delay of two seconds or more. This ensures that a delay of MACsec frames resulting from a man-in-the-middle attack will not go undetected.

When you configure bounded delay, you must also configure replay protection by setting the "replay-window-size" on page 14385. This is the window during which duplicate and replay packets are allowed. Bounded delay takes precedence over replay protection. You can increase the effectiveness of bounded delay protection by configuring a lower value for the window size.

NOTE: Bounded delay impacts CPU utilization which can degrade performance. We recommend only configuring bounded delay on interfaces on which it is absolutely required.

Default

Bounded delay is not enabled by default.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 21.1R1.

bpdu-block

IN THIS SECTION

- Syntax | 3022
- Hierarchy Level | 3023
- Description | 3023
- Required Privilege Level | 3023
- Release Information | 3023

Syntax

```
bpdu-block {
    interface (interface-name disable | all);
    disable-timeout seconds;
}
```

Hierarchy Level

```
[edit protocols layer2-control ]
```

```
[edit ethernet-switching-options]
```

Description

Configures bridge protocol data unit (BPDU) protection on a specified interface or on all interfaces. If the interface receives incompatible BPDUs, it is disabled.

If the `disable-timeout` statement is included in the BPDU configuration, the interface is automatically reenabled after the timer expires. Otherwise, you must use the operational mode command `clear ethernet-switching bpdu-error interface` to unblock and reenable the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

- *Understanding BPDU Protection for Spanning-Tree Instance Interfaces*
- *Configuring BPDU Protection for Individual Spanning-Tree Instance Interfaces*
- *Understanding BPDU Protection for STP, RSTP, and MSTP*
- *Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches*
- *Unblocking an Interface on non-ELS EX Series Switches That Receives BPDUs in Error (CLI Procedure)*
- *Understanding BPDU Protection for STP, RSTP, and MSTP*
- show spanning-tree bridge | **26548**

show spanning-tree interface | 26555

clear ethernet-switching bpdu-error interface | 3570

bpdu-block-on-edge

IN THIS SECTION

- Syntax | 3024
- Hierarchy Level | 3024
- Description | 3024
- Default | 3025
- Required Privilege Level | 3025
- Release Information | 3025

Syntax

```
bpdu-block-on-edge;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp | vstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp | vstp)],
[edit protocols ( mstp | rstp| vstp )],
[edit routing-instances routing-instance-name protocols (mstp | rstp | vstp)]
```

Description

Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch. When the `bpdu-block-on-edge` statement is configured and the interface encounters an incompatible BPDU, the interface shuts down.

If the `disable-timeout` statement is included in the BPDU configuration, the interface is automatically reenabled after the timer expires. Otherwise, you must use the operational mode command `clear ethernet-switching bpdu-error interface` to unblock and reenable the interface.

Default

Not enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support. Statement updated in Junos OS Release 11.1 for EX Series switches to change blocking behavior to port shutdown.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding BPDU Protection for Spanning-Tree Instance Interfaces

Understanding BPDU Protection for STP, RSTP, and MSTP

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

rstp | **16046**

mstp | **11006**

vstp | **30983**

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

clear ethernet-switching bpdu-error interface | **3570**

bpdu-destination-mac-address

IN THIS SECTION

- Syntax | 3026
- MX Series and EX Series | 3026
- SRX Series | 3026
- Description | 3027
- Default | 3027
- Required Privilege Level | 3027
- Release Information | 3027

Syntax

```
bpdu-destination-mac-address provider-bridge-group;
```

MX Series and EX Series

```
[edit logical-systems logical-system-name protocols mvrp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
mvrp] (for virtual switch instance type),
[edit protocols mvrp],
[edit routing-instances routing-instance-name protocols mvrp] (for virtual switch instance type)
```

SRX Series

```
[edit protocols mvrp],
[edit routing-instances routing-instance-name protocols mvrp] (for virtual switch instance type)
```

Description

For Multiple VLAN Registration Protocol (MVRP) configurations, specifies the multicast address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the Junos OS uses the customer MVRP multicast MAC address.

Default

By default, the provider MVRP multicast MAC address is used (if configured). Otherwise, the customer MVRP MAC address is used.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers

Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration

Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration on Security Devices

Verifying That MVRP Is Working Correctly

Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

bpdu-destination-mac-address (Spanning Tree)

IN THIS SECTION

Syntax | 3028

Hierarchy Level | 3028

- Description | 3028
- Default | 3028
- Options | 3028
- Required Privilege Level | 3029
- Release Information | 3029

Syntax

```
bpdu-destination-mac-address provider-bridge-group;
```

Hierarchy Level

```
[edit logical-systems  logical-system-name protocols (mstp | rstp)],
[edit logical-systems  logical-system-name routing-instances routing-instance-name
protocols (mstp | rstp)],
[edit protocols (mstp | rstp)],
[edit routing-instances routing-instance-name protocols (mstp | rstp)]
```

Description

Enable MX Series router to participate in the provider Rapid Spanning Tree Protocol (RSTP) instance or a provider Multiple Spanning Tree Protocol (MSTP) instance.

Default

If the `bpdu-destination-mac-address` statement is not configured, the bridge participates in the customer RSTP instance, transmitting and receiving standard RSTP BPDU packets.

Options

provider-bridge-group—The destination MAC address of the BPDU packets transmitted is the provider bridge group address **01:80:c2:00:00:08**. Received BPDU packets with this destination MAC address are accepted and passed to the Routing Engine.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding BPDUs Used for Exchanging Information Among Bridges

Understanding Provider Bridge Participation in RSTP or MSTP Instances

Configuring Rapid Spanning Tree Protocol

Configuring Multiple Spanning Tree Protocol

bridge-domain

IN THIS SECTION

- Syntax | 3030
- Hierarchy Level | 3030
- Description | 3030
- Options | 3030
- Required Privilege Level | 3030
- Release Information | 3030

Syntax

```
bridge-domain name;
    vlan-id [ vlan-identifiers ];
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain maintenance-domain-name],
[edit protocols oam ethernet connectivity-fault-management maintenance-domain maintenance-domain-name virtual-switch virtual-switch-name]
```

Description

(MX Series routers only) Specify the OAM Ethernet CFM maintenance domain bridge domain.

Options

name—Specify the name of the bridge domain.

vlan-identifiers—Specify one or more VLAN identifiers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Maintenance Intermediate Points (MIPs)

maintenance-domain | **10287**

bridge-domain (Analyzer)

IN THIS SECTION

- Syntax | 3031
- Hierarchy Level | 3031
- Description | 3031
- Options | 3031
- Required Privilege Level | 3032
- Release Information | 3032

Syntax

```
bridge-domain bridge-domain-name;
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input egress],
[edit forwarding-options analyzer analyzer-name input egress routing-instance instance-name],
[edit forwarding-options analyzer analyzer-name input ingress],
[edit forwarding-options analyzer analyzer-name input ingress routing-instance instance-name],
[edit forwarding-options analyzer analyzer-name output],
[edit forwarding-options analyzer analyzer-name output routing-instance instance-name]
```

Description

Configure the bridge domain to monitor outgoing traffic.

Options

bridge-domain-name	Name of the bridge domain that monitors outgoing traffic.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

> *Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use*
>
> *Example: Configuring Mirroring to Multiple Interfaces for Remote Monitoring of Employee Resource Use on EX9200 Switches*

bridge-domains

IN THIS SECTION

- Syntax | 3032
- Hierarchy Level | 3033
- Description | 3033
- Options | 3033
- Required Privilege Level | 3034
- Release Information | 3034

Syntax

```
bridge-domains {
    bridge-domain-name {
        bridge-options {
            ...bridge-options-configuration...
        }
        domain-type bridge;
```

```
        interface interface-name;
        no-irb-layer-2-copy;
        no-local-switching;
        routing-interface routing-interface-name;
        vlan-id (all | none | number);
        vlan-id-list [ vlan-id-numbers ];
        vlan-tags outer number inner number;
        bridge-options {
            interface interface-name {
                mac-pinning
                static-mac mac-address;
            }
            interface-mac-limit limit;
            mac-statistics;
            mac-table-size limit;
            no-mac-learning;
        }
    }
}
```

Hierarchy Level

```
[edit],
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

(MX Series routers only) Configure a domain that includes a set of logical ports that share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Options

bridge-domain-name—Name of the bridge domain.

> **NOTE**: You cannot use the slash (/) character as part of the bridge domain name. If you do, the configuration will not commit.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

Support for the `no-irb-layer-2-copy` statement added in Junos OS Release 10.2.

RELATED DOCUMENTATION

- Configuring a Bridge Domain
- Configuring a Layer 2 Virtual Switch

bridge-options

IN THIS SECTION

- Syntax | 3034
- Hierarchy Level | 3035
- Description | 3035
- Required Privilege Level | 3035
- Release Information | 3035

Syntax

```
bridge-options {
    interface interface-name;
        static-mac static-mac-address;
    }
```

```
    global-mac-ip-limit limit;
    interface-mac-ip-limit limit;
    interface-mac-limit limit;
        packet-action drop;
    }
    mac-pinning
    mac-statistics;
    mac-ip-table-size limit;
    mac-table-size limit;
    mac-table-aging-time time;
    no-mac-learning;
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]
```

Description

(MX Series routers only) Configure Layer 2 learning and forwarding properties for a bridge domain or a virtual switch.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

Statment (mac-pinning) introduced in Junos OS Release 16.2.

`global-mac-ip limit`, `interface-mac-ip-limit`, and `mac-ip-table-size` statements introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains

bridge-priority

IN THIS SECTION

- Syntax | 3036
- Hierarchy Level | 3036
- Description | 3037
- Default | 3037
- Options | 3037
- Required Privilege Level | 3037
- Release Information | 3038

Syntax

```
bridge-priority priority;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp)],
[edit logical-systems logical-system-name protocols mstp msti msti-id],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name
protocols mstp msti msti-id],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id],
```

```
[edit routing-instances routing-instance-name protocols (mstp | rstp)],
[edit routing-instances routing-instance-name protocols mstp msti msti-id],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id]
```

```
[edit protocols mstp],
[edit protocols mstp msti msti-id],
[edit protocols rstp],
[edit protocols stp],
[edit protocols vstp vlan vlan-id]
```

Description

Configures the bridge priority, which determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.

Default

32,768

Options

priority—The bridge priority can be set only in increments of 4096.

- **Range:** 0 through 61,440
- **Default:** 32,768

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding MSTP

Understanding VSTP

Understanding Bridge Priority for Election of Root Bridge and Designated Bridge

Example: Configuring Network Regions for VLANs with MSTP on Switches

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

bridge-template

IN THIS SECTION

- Syntax | **3038**
- Hierarchy Level | **3039**
- Description | **3039**
- Required Privilege Level | **3039**
- Release Information | **3039**

Syntax

```
bridge-template;
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix |version9 template template-name]
```

Description

Specify that the template is used for bridge records or for VPLS records.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

brief

IN THIS SECTION

- Syntax | 3040
- Hierarchy Level | 3040
- Description | 3040
- Default | 3040
- Required Privilege Level | 3041
- Release Information | 3041

Syntax

```
(brief | full);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options (aggregate | generate) (defaults | route)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name (aggregate | generate) (defaults | route)],
[edit logical-systems logical-system-name routing-options (aggregate | generate) (defaults |
route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name (aggregate |
generate) (defaults | route)],
[edit routing-instances routing-instance-name routing-options (aggregate | generate) (defaults |
route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name (aggregate |
generate) (defaults | route)],
[edit routing-options (aggregate | generate) (defaults | route)],
[edit routing-options rib routing-table-name (aggregate | generate) (defaults | route)]
```

Description

Configure all AS numbers from all contributing paths to be included in the aggregate or generated route's path.

- **brief**—Include only the longest common leading sequences from the contributing AS paths. If this results in AS numbers being omitted from the aggregate route, the BGP **ATOMIC_ATTRIBUTE** path attribute is included with the aggregate route.
- **full**—Include all AS numbers from all contributing paths in the aggregate or generated route's path. Include this option when configuring an individual route in the **route** portion of the `generate` statement to override a **retain** option specified in the **defaults** portion of the statement.

Default

full

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation

Example: Configuring a Conditional Default Route Policy

Understanding Conditionally Generated Routes

aggregate (Routing) | **2070**

generate | **7672**

broadcast (Interfaces)

IN THIS SECTION

- Syntax | **3042**
- Hierarchy Level | **3042**
- Description | **3042**
- Default | **3042**
- Options | **3042**
- Required Privilege Level | **3042**
- Release Information | **3043**

Syntax

```
broadcast address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family address address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family address address]
```

Description

Set the broadcast address on the network or subnet. On a subnet you cannot specify a host address of 0 (0.0.0.0), nor can you specify a broadcast address (255.255.255.255). For example, in the statement `set interface ge-0/0/0 unit 0 family inet address 10.1.1.0/24`, the subnet address `10.1.1.0` has the host address of 0. Hence, you cannot configure this address. Similarly, for the subnet, you cannot use the broadcast address `10.1.1.255/24`.

Default

The default broadcast address has a host portion of all ones.

Options

address—Broadcast address. The address must have a host portion of either all ones or all zeros. You cannot specify the addresses `0.0.0.0` or `255.255.255.255`.

NOTE: The `edit logical-systems` hierarchy is not available on QFabric systems.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Interface Address

broadcast (Router Discovery)

IN THIS SECTION

- Syntax | 3043
- Hierarchy Level | 3043
- Description | 3044
- Default | 3044
- Required Privilege Level | 3044
- Release Information | 3044

Syntax

```
(broadcast | multicast);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery address address],
[edit protocols router-discovery address address]
```

Description

Specify when the server should include the IP addresses in router advertisement packets. On the same physical interfaces, some addresses might be included only in multicast packets, while others might be included only in broadcast packets.

If you specify **broadcast**, the server includes the addresses in router advertisement packets only if the packets are broadcast.

Default

multicast if the router supports IP multicast; **broadcast** if the router does not support IP multicast.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

multicast (Protocols Router Discovery) | **11077**

broadcast (Forwarding class)

IN THIS SECTION

- Syntax | **3045**
- Hierarchy Level | **3045**
- Description | **3045**

- Options | 3045
- Required Privilege Level | 3045
- Release Information | 3045

Syntax

```
broadcast forwarding-class-name;
```

Hierarchy Level

```
[edit class-of-service multi-destination family ethernet]
```

Description

Specify the forwarding class for the broadcast traffic belonging to the Ethernet family.

Options

forwarding-class-name —Name of the forwarding class:

- **mcast-af**—Default forwarding class for assured forwarding of multicast traffic.
- **mcast-be**—Default best-effort forwarding class for multicast traffic.
- **mcast-ef**—Default forwarding class for expedited forwarding of multicast traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding CoS Schedulers

Understanding CoS Forwarding Classes

Understanding CoS Classifiers

bucket-size

IN THIS SECTION

- Syntax | 3046
- Hierarchy Level | 3046
- Description | 3046
- Default | 3047
- Options | 3047
- Required Privilege Level | 3047
- Release Information | 3047

Syntax

```
bucket-size number;
```

Hierarchy Level

```
[edit snmp rmon history index]
```

Description

Configure the sampling of Ethernet statistics for network fault diagnosis, planning, and performance tuning.

Default

50

Options

number—Number of discrete samples of Ethernet statistics requested.

Required Privilege Level

snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- *Understanding RMON*
- *RMON MIB Event, Alarm, Log, and History Control Tables*
- *Configuring RMON History Sampling*
- *Monitoring RMON MIB Tables*
- Junos OS Network Management Configuration Guide

buffer-dynamic-threshold

IN THIS SECTION

- Syntax | 3048
- Hierarchy Level | 3048
- Description | 3048
- Options | 3048
- Required Privilege Level | 3049

Release Information | 3049

Syntax

```
buffer-dynamic-threshold value;
```

Hierarchy Level

```
[edit class-of-service schedulers]
```

Description

Adjust the maximum size of the shared-buffer pool at the egress buffer partition for each queue.

NOTE: This option is not supported for lossless queues.

During congestion or micro-burst scenarios, dedicated buffers can be exhausted, and queues start consuming from the shared buffer pool. The device dynamically calculates the limit of buffers that each queue can consume from the shared buffer pool based on the dynamic threshold setting called *alpha*. The purpose of alpha, or `buffer-dynamic-threshold`, is to scale the current free resources to determine an instantaneous threshold.

NOTE: You adjust the maximum size of the *global* shared-buffer pool by configuring the `dynamic-threshold` value at the `[edit class-of-service shared-buffer egress buffer-partition]` hierarchy level.

When you configure the buffer dynamic threshold at both the global and the scheduler (per-queue) level, the per-queue value takes precedence.

Options

`buffer-dynamic-threshold` ***`value`***	Threshold for maximum buffer share for a queue at the egress buffer partition.

Range: 0-10

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 22.3R1.

RELATED DOCUMENTATION

xon (Input Congestion Notification) | 31161

buffer-partition (Egress)

IN THIS SECTION

- Syntax | 3049
- Hierarchy Level | 3050
- Description | 3050
- Default | 3050
- Options | 3051
- Required Privilege Level | 3051
- Release Information | 3051
- Platform-Specific Behavior | 3051

Syntax

```
buffer-partition (lossless | lossy | multicast) {
    percent percent;
```

```
    dynamic-threshold value;
}
```

Hierarchy Level

```
[edit class-of-service shared-buffer egress]
```

Description

The egress shared buffer pool is divided into three partitions. Each partition reserves a percentage of the available shared buffer pool for a type of traffic, so that the switch provides enough resources to support a mix of best-effort, lossless, and multicast traffic (multicast also includes broadcast and destination lookup fail traffic). To better support the mix of traffic on your network, you can optimize the allocation of egress shared buffers to different types of traffic by fine-tuning the shared buffer partitions.

The percentages you configure for the three egress shared buffer partitions must total exactly 100 percent. If the total of the three shared buffer percentages is not 100 percent, the system returns a commit error and does not commit the configuration. You can configure any partition to 0 (zero) percent as long as the allocation to other partitions totals 100 percent.

This is a global allocation that applies to all ports. All ports on the switch receive the same allocation of egress shared buffers.

If you do not configure buffer partitions, the switch uses the default partitioning.

CAUTION: Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Default

The default egress buffer partition shown in Table 7 on page 3051 supports networks with a balanced mix of best-effort, multicast, and lossless traffic. It is the recommended configuration if you are using the default configuration with two lossless forwarding classes.

Table 7: Default Egress Shared Buffer Partitioning

Lossless Partition	Lossy Partition	Multicast Partition
50%	31%	19%

The sum of the default percentages configured for each partition is 100 percent. The sum of the partition percentages must always total 100 percent.

Options

`dynamic-threshold` *value*	Threshold for maximum buffer share for a queue at the egress buffer partition.
`lossless`	Shared buffer space reserved for all lossless egress traffic.
`lossy`	Shared buffer space for best-effort unicast egress traffic.
`multicast`	Shared buffer space reserved for all multicast (including broadcast and destination lookup fail) egress traffic.
`percent` *percent*	The percentage of buffer space to allocate to the specified buffer partition (lossless, lossy, or multicast buffers). The sum of the percentages for the three buffer partitions must total 100 percent.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

`dynamic-threshold` option introduced in Junos OS Release 19.1R1.

`dynamic-threshold` option introduced in Junos OS Evolved Release 19.4R1.

Platform-Specific Behavior

Use Feature Explorer to confirm platform and release support for specific features.

Use the following table to review platform-specific behavior for your platforms:

Platform	Difference
QFX5000 Series	QFX5000 Series switches running Junos OS Evolved support the `dynamic-threshold` option *only* for lossy egress buffer partitions and not for lossless or multicast partitions.

RELATED DOCUMENTATION

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic

Configuring Global Ingress and Egress Shared Buffers

Understanding CoS Buffer Configuration

Aggregated Ethernet Interfaces

buffer-partition (Ingress)

IN THIS SECTION

- Syntax | 3053
- Hierarchy Level | 3053
- Description | 3053
- Default | 3053
- Options | 3054
- Required Privilege Level | 3054
- Release Information | 3054
- Platform-Specific Behavior | 3055

Syntax

```
buffer-partition (lossless | lossless-headroom | lossy) {
    percent percentage;
    dynamic-threshold value;
}
```

Hierarchy Level

```
[edit class-of-service shared-buffer ingress]
```

Description

The ingress shared buffer pool is divided into three partitions. Each partition reserves a percentage of the available shared buffer pool for a type of traffic, so that the switch provides enough resources to support a mix of best effort (best-effort unicast and multicast) and lossless traffic. To better support the mix of traffic on your network, you can optimize the allocation of ingress shared buffers to different types of traffic by fine-tuning the shared buffer partitions.

The percentages you configure for the three ingress shared buffer partitions must total exactly 100 percent. If the total of the three shared buffer percentages is not 100 percent, the system returns a commit error and does not commit the configuration. You can configure any partition to 0 (zero) percent as long as the allocation to other partitions totals 100 percent.

This is a global allocation that applies to all ingress traffic. All ports on the switch receive the same allocation of ingress shared buffers.

If you do not configure buffer partitions, the switch uses the default partitioning.

> ⚠ **CAUTION**: Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Default

The default ingress buffer partition shown in Table 8 on page 3054 supports networks with a balanced mix of best-effort, multicast, and lossless traffic. It is the recommended configuration if you are using the default configuration with two lossless forwarding classes.

Table 8: Default Ingress Shared Buffer Partitioning

Lossless Partition	Lossless-Headroom Partition	Lossy Partition
9%	45%	46%

The sum of the default percentages configured for each partition is 100 percent. The sum of the partition percentages always must total 100 percent.

Options

`dynamic-threshold` *`value`* — Threshold for maximum buffer share for a queue at the ingress buffer partition.

`lossless` — Shared buffer space reserved for all lossless ingress traffic.

`lossless-headroom` — Shared buffer space reserved to store packets received while either an 802.3x Ethernet PAUSE or a priority-based flow control (PFC) pause is asserted. (When an ingress interface pauses traffic, it must have the buffer space to store all of the packets currently in the buffer, and also all of the packets received before the connected peer stops sending traffic and the wire is cleared of packets.)

`lossy` — Shared buffer space for best-effort ingress traffic.

`percent` *`percent`* — The percentage of buffer space to allocate to the specified buffer partition (lossless, lossless-headroom, or lossy buffers). The sum of the percentages for the three buffer partitions must total 100 percent.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

`dynamic-threshold` option introduced in Junos OS Release 19.1R1.

`dynamic-threshold` option introduced in Junos OS Evolved Release 19.4R1.

Platform-Specific Behavior

Use Feature Explorer to confirm platform and release support for specific features.

Use the following table to review platform-specific behavior for your platforms:

Platform	Difference
QFX5000 Series	QFX5000 Series switches use the `dynamic-threshold` option on lossless ingress buffer partitions to calculate the ECN threshold on lossless queues.

RELATED DOCUMENTATION

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic

Configuring Global Ingress and Egress Shared Buffers

Understanding CoS Buffer Configuration

Aggregated Ethernet Interfaces

buffer-size

IN THIS SECTION

- Syntax | 3056
- Hierarchy Level | 3056
- Description | 3056
- Default | 3058
- Options | 3060

Required Privilege Level | 3060
Release Information | 3060

Syntax

```
buffer-size (percent percent | remainder);
```

```
buffer-size (exact | percent percentage | remainder | temporal microseconds);
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

```
[edit class-of-service schedulers scheduler-name]
```

Description

Specify buffer size in a scheduler configuration.

On all switches, you configure the proportion of port buffers allocated to a particular output queue using the following process:

1. Configure a scheduler and set the `buffer-size` option.
2. Use a scheduler map to map the scheduler to the forwarding class that is mapped to the queue to which you want to apply the buffer size.

   For example, suppose that you want to change the dedicated buffer allocation for FCoE traffic. FCoE traffic is mapped to the fcoe forwarding class, and the fcoe forwarding class is mapped to queue 3 (this is the default configuration). To use default FCoE traffic mapping, in the scheduler map configuration, map the scheduler to the `fcoe` forwarding class.
3. If you are using enhanced transmission selection (ETS) hierarchical scheduling, associate the scheduler map with the traffic control profile you want to use on the egress ports that carry FCoE traffic. If you are using direct port scheduling, skip this step.

4. If you are using ETS, associate the traffic control profile that includes the scheduler map with the desired egress ports. For this example, you associate the traffic control profile with the ports that carry FCoE traffic. If you are using port scheduling, associate the scheduler map with the desired egress ports.

   Queue 3, which is mapped to the fcoe forwarding class and therefore to the FCoE traffic, receives the dedicated buffer allocation specified in the `buffer-size` statement.

> **NOTE**: The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

QFX10000 Switches

On QFX10000 switches, the buffer size is the amount of time in milliseconds of port bandwidth that a queue can use to continue to transmit packets during periods of congestion, before the buffer runs out and packets begin to drop.

The switch can use up to 100 ms total (combined) buffer space for all queues on a port. A buffer-size configured as one percent is equal to 1 ms of buffer usage. A buffer-size of 15 percent (the default value for the best effort and network control queues) is equal to 15 ms of buffer usage.

The total buffer size of the switch is 4 GB. A 40-Gigabit port can use up to 500 MB of buffer space, which is equivalent to 100 ms of port bandwidth on a 40-Gigabit port. A 10-Gigabit port can use up to 125 MB of buffer space, which is equivalent to 100 ms of port bandwidth on a 10-Gigabit port. The total buffer sizes of the eight output queues on a port cannot exceed 100 percent, which is equal to the full 100 ms total buffer available to a port. The maximum amount of buffer space any queue can use is also 100 ms (which equates to a 100 percent buffer-size configuration), but if one queue uses all of the buffer, then no other queue receives buffer space.

There is no minimum buffer allocation, so you can set the buffer-size to zero (0) for a queue. However, we recommend that on queues on which you enable PFC to support lossless transport, you allocate a minimum of 5 ms (a minimum buffer-size of 5 percent). The two default lossless queues, fcoe and no-loss, have buffer-size default values of 35 ms (35 percent).

Queue buffer allocation is dynamic, shared among ports as needed. However, a queue cannot use more than its configured amount of buffer space. For example, if you are using the default CoS configuration, the best-effort queue receives a maximum of 15 ms of buffer space because the default transmit rate for the best-effort queue is 15 percent.

If a switch experiences congestion, queues continue to receives their full buffer allocation until 90 percent of the 4 GB buffer space is consumed. When 90 percent of the buffer space is in use, the amount of buffer space per port, per queue, is reduced in proportion to the configured buffer size for each queue. As the percentage of consumed buffer space rises above 90 percent, the amount of buffer space per port, per queue, continues to be reduced.

On 40-Gigabit ports, because the total buffer is 4 GB and the maximum buffer a port can use is 500 MB, up to seven 40-Gigabit ports can consume their full 100 ms allocation of buffer space. However, if an eighth 40-Gigabit port requires the full 500 MB of buffer space, then the buffer allocations are proportionally reduced because the buffer consumption is above 90 percent.

On 10-Gigabit ports, because the total buffer is 4 GB and the maximum buffer a port can use is 125 MB, up to 28 10-Gigabit ports can consume their full 100 ms allocation of buffer space. However, if a 29th 10-Gigabit port requires the full 125 MB of buffer space, then the buffer allocations are proportionally reduced because the buffer consumption is above 90 percent.

QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems

Set the dedicated buffer size of the egress queue that you bind the scheduler to in the scheduler map configuration. The switch allocates space from the global dedicated buffer pool to ports and queues in a hierarchical manner. The switch allocates an equal number of dedicated buffers to each egress port, so each egress port receives the same amount of dedicated buffer space. The amount of dedicated buffer space per port is not configurable.

However, the `buffer-size` statement allows you to control the way each port allocates its share of dedicated buffers to its queues. For example, if a port only uses two queues to forward traffic, you can configure the port to allocate all of its dedicated buffer space to those two queues and avoid wasting buffer space on queues that are not in use. We recommend that the buffer size should be the same size as the minimum guaranteed transmission rate (the `transmit-rate`).

Default

The default behavior differs on different switches.

QFX10000 Switches

If you do not configure buffer-size and you do not explicitly configure a queue scheduler, the default buffer-size is the default transmit rate of the queue. If you explicitly configure a queue scheduler, the default buffer allocations are not used. If you explicitly configure a queue scheduler, configure the buffer-size for each queue in the scheduler, keeping in mind that the total buffer-size of the queues cannot exceed 100 percent (100 ms).

Table 9 on page 3058 shows the default queue buffer sizes on QFX10000 switches. The default buffer size is the same as the default transmit rate for each default queue:

Table 9: Default Output Queue Buffer Sizes (QFX10000 Switches)

Queue Number	Forwarding Class	Transmit Rate	Buffer Size
0	best-effort	15%	15%

Table 9: Default Output Queue Buffer Sizes (QFX10000 Switches) *(Continued)*

Queue Number	Forwarding Class	Transmit Rate	Buffer Size
3	fcoe	35%	35%
4	no-loss	35%	35%
7	network-control	15%	15%

By default, only the queues mapped to the default forwarding classes receive buffer space from the port buffer pool. (Buffers are not wasted on queues that do not carry traffic.)

QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems

The port allocates dedicated buffers to queues that have an explicitly configured scheduler buffer size. If you do not explicitly configure a scheduler buffer size for a queue, the port serves the explicitly configured queues first. Then the port divides the remaining dedicated buffers equally among the queues that have an explicitly attached scheduler *without* an explicitly configured buffer size configuration. (If you configure a scheduler, but you do not configure the buffer size parameter, the default is equivalent to configuring the buffer size with the `remainder` option.)

If you use the default scheduler and scheduler map on a port (no explicit scheduler configuration), then the port allocates its dedicated buffer pool to queues based on the default scheduling. Table 10 on page 3059 shows the default queue buffer sizes. The default buffer size is the same as the default transmit rate for each default queue:

Table 10: Default Output Queue Buffer Sizes (QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)

Queue Number	Forwarding Class	Transmit Rate	Buffer Size
0	best-effort	5%	5%
3	fcoe	35%	35%
4	no-loss	35%	35%
7	network-control	5%	5%

Table 10: Default Output Queue Buffer Sizes (QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems) *(Continued)*

Queue Number	Forwarding Class	Transmit Rate	Buffer Size
8	mcast	20%	20%

By default, only the queues mapped to the default forwarding classes receive buffer space from the port buffer pool. (Buffers are not wasted on queues that do not carry traffic.)

On EX Series switches except EX4300 switches, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent, respectively. On EX4300 switches, the default scheduler transmission rate and buffer size for queues 0 through 11 are 75, 0, 0, 5, 0, 0, 0, 0, 15, 0, 0 and 5 percent, respectively, of the total available buffer.

Options

percent *percent* — Percentage of the port dedicated buffer pool allocated to the queue (or queues) mapped to the scheduler.

remainder — Remaining dedicated buffer pool after the port satisfies the needs of the explicitly configured buffers. The port divides the remaining buffers equally among the queues that are explicitly attached to a scheduler but that do not have an explicit buffer size configuration (or are configured with `remainder` as the buffer size).

exact — Enforce the exact buffer size. When this option is configured, sharing is disabled on the queue, restricting the usage to guaranteed buffers only.

temporal *microseconds* — Buffer size as a temporal value.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches

Defining CoS Schedulers and Scheduler Maps (CLI Procedure)

Defining CoS Schedulers (J-Web Procedure)

Understanding CoS Schedulers

buffer-size (Dynamic Scheduling)

IN THIS SECTION

- Syntax | 3061
- Hierarchy Level | 3061
- Description | 3061
- Default | 3062
- Options | 3062
- Required Privilege Level | 3062
- Release Information | 3062

Syntax

```
buffer-size (percent (percentage | $junos-cos-scheduler-bs) | remainder | temporal (microseconds
| $junos-cos-scheduler-bs));
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
```

Description

Specify buffer size.

Default

If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

Options

`percent` *`percentage`*—Buffer size as a percentage of total buffer.

`remainder`—Remaining buffer available.

`temporal` *`microseconds`*—Buffer size as a temporal value. The queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This maximum is computed by multiplying the logical interface speed by the configured temporal value.

`$junos-scheduler-bs`—Junos predefined variable that is replaced with the buffer size obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

The `$junos-cos-scheduler-bs` predefined variable introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Schedulers in a Dynamic Profile for Subscriber Access

scheduler (Dynamic Scheduler Maps) | **16344**

buffer-size (Schedulers)

IN THIS SECTION

- Syntax | 3063
- Hierarchy Level | 3063
- Description | 3063
- Default | 3064
- Options | 3064
- Required Privilege Level | 3065
- Release Information | 3065

Syntax

```
buffer-size (percent percentage | remainder | shared | temporal microseconds);
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Specify buffer size in a scheduler configuration.

You configure the proportion of port buffers allocated to a particular output queue using the following process:

1. Configure a scheduler and set the `buffer-size` option.
2. Use a scheduler map to map the scheduler to the forwarding class that is mapped to the queue to which you want to apply the buffer size.

   For example, suppose that you want to change the dedicated buffer allocation for best effort traffic. Best effort traffic is mapped to the best effort forwarding class, and the best effort forwarding class is

mapped to queue 0 (by default). To use default best effort traffic mapping, in the scheduler map configuration, map the scheduler to the `best-effort` forwarding class.

3. If you are using hierarchical scheduling, associate the scheduler map with the traffic control profile you want to use on the egress ports that carry best effort traffic. If you are using direct port scheduling, skip this step.

4. If you are using hierarchical scheduling, associate the traffic control profile that includes the scheduler map with the desired egress ports. For this example, you associate the traffic control profile with the ports that carry best effort traffic. If you are using port scheduling, associate the scheduler map with the desired egress ports.

   Queue 0, which is mapped to the best effort forwarding class and therefore to best effort traffic, receives the dedicated buffer allocation specified in the `buffer-size` statement.

> **NOTE**: The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

> **NOTE**: On PTX Series Packet Transport Routers, you cannot configure buffer-size on rate-limited queues.

Default

If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent, respectively.

Options

`percent` *`percentage`*
Buffer size as a percentage of the total buffer.

- **Range:** 0 through 100

`remainder`
Remaining buffer available.

`shared`
On PTX Series routers, set a queue's buffer to be up to 100 percent of the interface's buffer. This option allows the queue's buffer to grow as large as 100 percent of the interface's buffer if and only if it is the only active queue for the interface.

`temporal` *`microseconds`*
Buffer size as a temporal value. The queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This maximum is computed by multiplying the logical interface speed by the configured temporal value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`shared` option introduced in Junos OS Release 18.1.

RELATED DOCUMENTATION

Managing Congestion on the Egress Interface by Configuring the Scheduler Buffer Size

buffer-rate (Schedulers)

IN THIS SECTION

- Syntax | 3065
- Hierarchy Level | 3066
- Description | 3066
- Options | 3066
- Required Privilege Level | 3066
- Release Information | 3066

Syntax

```
buffer-rate (percent percent |  rate);
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Configure the base rate for buffer size calculation. The buffer rate is the target rate of a VOQ, which is the intended egress queue rate during typical congestion. The target rate is an import parameter for VOQs, as the ratio of the drain rate to the target rate determines how much buffer adjustment happens at ingress.

The VOQ target rate is equal to the lesser of the buffer rate and the calculated shaping rate. If you don't define the buffer rate, then the VOQ target rate is equal to the scheduler transmit rate. If you also don't define the scheduler transmit rate, then the VOQ target rate is equal to the calculated shaping rate.

For platforms that support VOQs each ingress Packet Forwarding Engine (PFE) provides up to eight virtual output queues for a single egress port. The egress PFE distributes the bandwidth to each ingress virtual output queue; therefore they receive equal treatment regardless of their presented load. The drain rate of a queue is the rate at which a queue is draining. The egress Packet Forwarding Engine divides its bandwidth for each output queue equally across the ingress PFEs. So, the drain rate of each ingress PFE = drain rate of the output queue/number of ingress PFEs.

Options

percent *percentage*	Buffer rate as a percentage of the total buffer. • **Range:** 0 through 100
rate	Buffer rate as an absolute value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

buffer-size

bulk-call-mode

IN THIS SECTION

- Syntax | 3067
- Hierarchy Level | 3067
- Description | 3067
- Required Privilege Level | 3068
- Release Information | 3068

Syntax

```
bulk-call-mode;
```

Hierarchy Level

```
[edit security alg sip ]
```

Description

When there are multiple users behind SIP proxies at either end of an SRX, with multiple SIP calls ongoing between the users via these SIP proxies and through the device, and all SIP traffic is using the same 5 tuples, this results in a single SIP control session on the device, referred to as SIP bulk call or a SIP "fat" session.

To load balance such SIP calls on the device, the command `set security alg sip bulk-call-mode` can be used. When enabled, SIP call ID is used to identify and distribute the SIP calls among different SPUs instead of the regular 5 tuples, resulting in individual control sessions for each SIP call.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

sip (Security) | 27977

bulk-leasequery (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 3068
- Hierarchy Level | 3069
- Description | 3069
- Options | 3069
- Required Privilege Level | 3070
- Release Information | 3070

Syntax

```
bulk-leasequery {
    attempts number-of-attempts;
    timeout seconds;
    trigger automatic;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name ...],
[edit routing-instances routing-instance-name ...]
```

Description

Enable support for bulk leasequery on a DHCPv4 or DHCPv6 relay agent. You can also configure parameters that the DHCP relay agent uses when sending DHCP bulk leasequery messages to obtain lease information from the DHCP local servers in the logical system/routing instance.

> **NOTE**: You must also configure support on the relevant DHCP local servers with the "allow-leasequery" on page 2191 statement.

Options

attempts *number-of-attempts*

Specify the number of times the DHCP relay agent attempts to send DHCP bulk leasequery messages to the configured DHCP servers in the logical system/routing instance. DHCP relay agent resends the query message if the configured `timeout` value is reached, and either a confirmed reply or a reply from all configured DHCP servers has not been received. DHCP relay agent sends the subsequent messages only to the DHCP servers that have not replied to previous queries.

- Range:
  - (DHCPv4 bulk leasequery) 1 through 10
  - (DHCPv6 bulk leasequery) 1 through 720
- Default:
  - (DHCPv4 bulk leasequery) 6
  - (DHCPv6 bulk leasequery) 360

timeout *seconds*

Specify the number of seconds that DHCP relay agent waits before resending a leasequery message to the DHCP servers when all servers have not responded to a previous message.

- **Range:** 1 through 10
- **Default:** 10

trigger automatic (DHCPv6 only) Specify that the DHCPv6 relay agent always sends a bulk leasequery message whenever the jdhcp daemon is started or restarted and there are no bound subscribers in the session database. The automatic trigger updates the DHCP relay agent's binding information for clients associated with the requesting DHCPv6 relay agent.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring and Using DHCP Bulk Leasequery

bum-hashing

IN THIS SECTION

- Syntax | 3071
- Hierarchy Level | 3071
- Description | 3071
- Required Privilege Level | 3071
- Release Information | 3071

Syntax

```
bum-hashing;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Load balance VPLS BUM (broadcast, unknown, and multicast) traffic across all members of an aggregate interface for the routing instance.

> **NOTE**: On MX Series routers, the bum-hashing configuration is not required when the router is operating in `enhanced-ip` or `enhanced-ethernet` modes. If set, the configuration is ignored.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Load Balancing VPLS Non-Unicast Traffic Across Member Links of an Aggregate Interface

Example: Configuring Multicast Load Balancing over Aggregated Ethernet Links

bundle

IN THIS SECTION

- Syntax | 3072
- Hierarchy Level | 3072
- Description | 3072
- Options | 3072
- Required Privilege Level | 3072
- Release Information | 3073

Syntax

```
bundle (lsq-fpc/pic/port | ... );
```

Hierarchy Level

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number family mlppp]
```

Description

Associate the voice services interface with the logical interface it is joining.

Options

`lsq-`*`fpc/pic/port`*—Name of the voice services interface you are linking.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Voice Services Bundles with MLPPP Encapsulation

bundle (multilink)

IN THIS SECTION

- Syntax | 3073
- Hierarchy Level | 3073
- Description | 3073
- Options | 3074
- Required Privilege Level | 3074
- Release Information | 3074

Syntax

```
bundle (ml-fpc/pic/port | ls-fpc/pic/port);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family mlfr-end-to-end],
[edit interfaces interface-name unit logical-unit-number family mlfr-uni-nni]
```

Description

Associate the multilink interface with the logical interface it is joining.

Options

`ml-`*`fpc/pic/port`*—Name of the multilink interface you are linking.

`ls-`*`fpc/pic/port`*—Name of the link services interface you are linking.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Links in a Multilink or Link Services Bundle

bundle (Interfaces)

IN THIS SECTION

- Syntax | 3074
- Hierarchy Level | 3075
- Description | 3075
- Required Privilege Level | 3075
- Release Information | 3075

Syntax

```
bundle bundle-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family mlppp]
```

Description

Specify the logical interface name the link joins.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Interfaces

burst (DDoS)

IN THIS SECTION

- Syntax | 3076
- Hierarchy Level | 3076
- Description | 3076
- Options | 3076
- Required Privilege Level | 3076
- Release Information | 3077

Syntax

```
burst size;
```

Hierarchy Level

- For MX Series routers, T4000 routers, and EX9200 switches:

```
[edit system ddos-protection protocols protocol-group  (aggregate | packet-type)]
```

- For ACX Series routers, PTX Series routers and QFX Series switches:

```
[edit system ddos-protection protocols protocol-group  (aggregate | packet-type)]
```

Description

Configure the DDoS burst limit; that is, the maximum number of packets that is allowed in a burst of traffic by the specified policer. When this value is exceeded, a violation is declared.

Options

size—Number of packets that are allowed in a burst by the aggregate or packet-type policer.

- **Range:** 1 through 100,000 packets
- **Default:** The default burst value varies by packet type or protocol. You can view the default values for all packet types or protocols on an unconfigured router or switch by entering the `show ddos-protection protocols parameters brief` command from operational mode. For ACX Series routers, PTX Series routers and QFX Series switches, the default bandwidth limits are also provided in the "protocols (DDoS) (ACX Series, PTX Series, and QFX Series)" on page 13765 statement description.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers

burst-scale (DDoS)

IN THIS SECTION

- Syntax | 3077
- Hierarchy Level | 3077
- Description | 3078
- Options | 3078
- Required Privilege Level | 3078
- Release Information | 3078

Syntax

```
burst-scale percentage;
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group (aggregate | packet-type) fpc slot-number]
```

Description

(MX Series routers with only MPCs, T4000 routers with only FPC5s, PTX Series routers, EX9200 switches, or QFX Series switches) Configure the percentage by which the DDoS burst limit is scaled down for the aggregate or packet-type policer on the card in the specified slot.

Options

percentage—Percentage multiplied by the burst limit to reduce the number of packets allowed in a burst for the packet type or protocol.

- **Range:** 1 through 100 percent
- **Default:** 100

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers

burst-size

IN THIS SECTION

- Syntax | 3079
- Hierarchy Level | 3079
- Description | 3079
- Default | 3080

- Options | 3080
- Required Privilege Level | 3080
- Release Information | 3081

Syntax

```
burst-size bytes;
```

Hierarchy Level

```
[edit forwarding-options storm-control-profiles profile-name all]
```

Description

Configure the number of bytes of bursting traffic allowed to pass through a storm control interface. The burst size allows for short periods of back-to-back traffic at average rates that exceed the storm control level. If either the burst size or rate exceeds the limit, traffic will be dropped.

Storm control enables the switch to monitor traffic levels and to drop broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level—is exceeded, thus preventing these packets from proliferating and causing a traffic storm. Storm control can also shut down the interface in the event of a traffic storm.

The storm control level is a limit on bandwidth, which is configured using either the `bandwidth-limit` statement or the `bandwidth-percent` statement. The burst size extends the bandwidth limit so that sudden bursts of traffic do not cause storm control to drop packets or shut down the interface.

Because one burst size is not suitable for every traffic pattern, select the best burst size for an interface by performing experimental configurations. For your first test configuration, select the burst size limit by using one of the following methods.

The preferred method for determining the burst size is by using the following values:

- bandwidth—line rate of the storm control interface (in bps units)

- burst-period—allowable traffic-burst time (in milliseconds)

BEST PRACTICE: We recommend you set the burst size value to the amount of traffic that can be sent over the interface in 5 milliseconds. For example, if your interface is configured with a 100 Mbps bandwidth limit, the recommended value would be 62,500 bytes. This value is derived using the bandwidth and burst-period, as follows:

- Convert the bandwidth (100 Mbps) to bps units.
- Multiply by the burst-period (5 milliseconds or 0.005 seconds).
- Divide by 8 to convert from bits to bytes.

```
                    100,000,000 bps x 0.005 s
100 Mbps x 5 ms  =  ------------------------  =  62,500 bytes
                       8 bits per byte
```

If you do not know the bandwidth of the interface, you can calculate the burst size using the maximum transmission unit (MTU). The burst size limit should not be set lower than 10 times the MTU of the traffic on the interface.

For more information on burst size calculation, see *Determining Proper Burst Size for Traffic Policers*.

Default

If you do not configure the allowed burst size, a traffic burst in excess of the storm control level might cause storm control to drop traffic or to shut down the interface.

Options

bytes—Burst size limit in bytes.

- **Range:** 1,500 through 100,000,000 bytes
- **Default:** 1,500 kilobits (187,500 bytes)

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R3-S7, 17.4R3-S2, and 18.1R1.

RELATED DOCUMENTATION

Determining Proper Burst Size for Traffic Policers

Enabling and Disabling Storm Control (ELS)

burst-size (Default Local Policy)

IN THIS SECTION

- Syntax | 3081
- Hierarchy Level | 3081
- Description | 3082
- Options | 3082
- Required Privilege Level | 3082
- Release Information | 3082

Syntax

```
burst-size uplink uplink-burst-size downlink downlink-burst-size;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber default-local-
policy]
```

Description

Specify the allowed burst size for a subscriber's uplink and downlink traffic during the TDF IP-based subscriber creation process. Uplink traffic originates from the subscriber towards the public data network (PDN), and downlink traffic comes from the PDN and is destined for the subscriber.

Options

uplink-burst-size

Burst size value for the uplink direction.

- **Range:** 1500 through 1,500,000,000 bytes.

downlink-burst-size

Burst size value for the downlink direction.

- **Range:** 1500 through 1,500,000,000 bytes

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

burst-size-limit

IN THIS SECTION

- Syntax | 3083
- Hierarchy Level | 3083
- Description | 3083
- Options | 3083
- Required Privilege Level | 3084
- Release Information | 3084

Syntax

```
burst-size-limit bytes;
```

Hierarchy Level

```
[edit firewall policer policer-name if-exceeding]
[edit logical-systems logical-system-name firewall policer policer-name if-exceeding]
```

Description

Specify the maximum allowed burst size to control the amount of traffic bursting.

Options

bytes —Decimal value or a decimal number followed by k (thousand) or m (million).

Range:

-
  - 1 through 2,147,450,880 bytes (EX Series switches)
  - 1500 through 1,00,000,000,000 bytes (routers)

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches

Configuring Policers to Control Traffic Rates (CLI Procedure)

Understanding the Use of Policers in Firewall Filters

Basic Single-Rate Two-Color Policers

burst-size-limit (Hierarchical Policer)

IN THIS SECTION

- Syntax | 3084
- Hierarchy Level | 3085
- Description | 3085
- Options | 3085
- Required Privilege Level | 3085
- Release Information | 3085

Syntax

```
burst-size-limit bytes;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer aggregate if-exceeding],
[edit dynamic-profiles profile-name firewall hierarchical-policer premium if-exceeding],
[edit firewall hierarchical-policer aggregate if-exceeding],
[edit firewall hierarchical-policer premium if-exceeding]
```

Description

On M40e, M120, and M320 (with FFPC and SFPC) edge routers; on MPCs hosted on MX Series routers; on T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs; and on T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure the burst-size limit for premium or aggregate traffic in a hierarchical policer.

Options

bytes—Burst-size limit in bytes. The minimum recommended value is the maximum transmission unit (MTU) of the IP packets being policed. You can specify the value either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1500 through 2,147,450,880 (1500 through 100,000,000,000 on MPCs hosted on MX Series routers)

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles ... if exceeding]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Hierarchical Policer Configuration Overview
- Policer Bandwidth and Burst-Size Limits

Policer Color-Marking and Actions
Single Token Bucket Algorithm
Determining Proper Burst Size for Traffic Policers
Hierarchical Policers
aggregate (Hierarchical Policer) | 2064
bandwidth-limit (Hierarchical Policer)
premium (Hierarchical Policer) | 13110

burst-size-limit (Policer)

IN THIS SECTION

- Syntax | 3086
- Hierarchy Level | 3086
- Description | 3087
- Options | 3088
- Required Privilege Level | 3088
- Release Information | 3088

Syntax

```
burst-size-limit bytes;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name if-exceeding],
[edit firewall policer policer-name if-exceeding],
[edit logical-systems logical-system-name policer policer-name if-exceeding]
```

Description

For a single-rate two-color policer, configure the burst size as a number of bytes. The burst size allows for short periods of traffic bursting (back-to-back traffic at average rates that exceed the configured bandwidth limit). Single-rate two-color policing uses the *single token bucket algorithm* to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the bandwidth limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with `low` packet loss priority and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

The burst size extends the function of the bandwidth limit (configured using either the `bandwidth-limit` *`bps`* statement or the `bandwidth-percent` *`percentage`* statement) to allow bursts of traffic up to a limit based on the overall traffic load:

- When a single-rate two-color policer is applied to the input or output traffic at an interface, the initial capacity for traffic bursting is equal to the number of bytes specified by this statement.
- During periods of relatively low traffic (traffic that arrives at or departs from the interface at overall rates below the token arrival rate), unused tokens accumulate in the bucket, but only up to the configured token bucket depth.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allows bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

The burst-size limit enforced is based on the burst-size limit you configure. For a rate-limited logical interface, the Packet Forwarding Engine calculates the optimum burst-size-limit values and then applies the value closest to the burst-size-limit value specified in the policer configuration.

On MX Series routers and EX Series switches, the burst-size limit is not as freely configurable as it is on other platforms. Junos OS does not support an unlimited combination of policer bandwidth and burst-size limits on MX Series routers and EX Series switches. For a single-rate two-color policer on an MX Series router and on an EX Series switch, the minimum supported burst-size limit is equivalent to the amount of traffic allowed by the policer bandwidth limit in a time span of 1 millisecond. For example, for a policer configured with a `bandwidth-limit` value of 1 Gbps, the minimum supported value for `burst-size-`

`limit` on an MX Series router is 125 KB. If you configure a value that is smaller than the minimum, Junos OS overrides the configuration and applies the actual minimum.

Options

bytes—Burst-size limit in bytes. The minimum recommended value is the maximum transmission unit (MTU) of the IP packets being policed. You can specify the value either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1500 through 100,000,000,000
- **Default:** None

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit dynamic-profiles ... if-exceeding]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Two-Color Policer Configuration Overview
- Policer Bandwidth and Burst-Size Limits
- Policer Color-Marking and Actions
- Single Token Bucket Algorithm
- Determining Proper Burst Size for Traffic Policers
- bandwidth-limit (Policer) | **2857**
- bandwidth-percent | **2866**

burst-size-limit (Policer for Gigabit Ethernet Interfaces)

IN THIS SECTION

- Syntax | 3089
- Hierarchy Level | 3089
- Description | 3089
- Options | 3089
- Required Privilege Level | 3090
- Release Information | 3090

Syntax

```
burst-size-limit bytes;
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile
policer cos-policer-name aggregate],
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile
policer cos-policer-name premium]
```

Description

Define a policer to apply to nonpremium traffic.

Options

bytes—Burst length.

- **Range:** 1500 through 100,000,000 bytes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

bandwidth-limit (Policer for Gigabit Ethernet Interfaces) | 2860

burst-size (TDF Domain)

IN THIS SECTION

- Syntax | 3090
- Hierarchy Level | 3091
- Description | 3091
- Options | 3091
- Required Privilege Level | 3091
- Release Information | 3091

Syntax

```
burst-size {
    apply-groups [group-names];
    apply-groups-except [group-names];
    downlink downlink-burst-size;
```

```
    uplink uplink-burst-size ;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
```

Description

Configure the TDF domain's default TDF subscriber allowed burst size for uplink and downlink traffic. Uplink traffic originates from the subscriber towards the public data network (PDN), and downlink traffic comes from the PDN and is destined for the subscriber.

Options

downlink-burst-size

Burst size value for the downlink direction.

- **Range:** 1500 through 1,500,000,000 bytes.

uplink-burst-size

Burst size value for the uplink direction.

- **Range:** 1500 through 1,500,000,000 bytes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based and IFL-Based TDF Subscribers Overview

IP-Based Subscriber Setup Overview

by-destination (IDS MS-DPC)

IN THIS SECTION

- Syntax | 3092
- Hierarchy Level | 3092
- Description | 3092
- Options | 3093
- Required Privilege Level | 3093
- Release Information | 3093

Syntax

```
by-destination {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then session-limit]
```

Description

Apply limit to sessions based on numbers generated from the configured destination (IP or subnet) or application when using the MS-DPC.

Options

`hold-time` *seconds*—Length of time for which to stop all new flows once the rate of events exceeds the threshold set by one or more of the `maximum`, `packets`, or `rate` statements.

`maximum` *number*—Maximum number of open sessions per application or IP address.

`packets` *number*—Maximum peak packets per second per application or IP address.

`rate` *number*—Maximum number of sessions per second per application or IP address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in IDS Rules

by-destination (IDS MS-MPC)

IN THIS SECTION

- Syntax | 3094
- Hierarchy Level | 3094
- Description | 3094
- Options | 3095
- Required Privilege Level | 3095
- Release Information | 3095

Syntax

```
by-destination {
    by-protocol {
        icmp {
            maximum number;
            packets number;
            rate number;
        }
        tcp {
            maximum number;
            packets number;
            rate number;
        }
        udp {
            maximum number;
            packets number;
            rate number;
        }
    }
    maximum number;
    packets number;
    rate number;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then session-limit]
```

Description

Configure the IDS rule session limits for an individual destination address or subnet. This protects against network probing attacks and network flooding attacks. This IDS rule can only be assigned to a service set on an MS-MPC.

When a session limit is exceeded for a destination, packets to the destination are dropped until the session limit is no longer exceeded.

To specify limits for destination subnets rather than individual addresses, include the `aggregation` statement at the `[edit services ids rule rule-name term term-name then]` hierarchy level.

Options

maximum *number* Specify the maximum number of concurrent sessions allowed for an individual destination address or subnet.

packets *number* Specify the maximum number of packets per second allowed for an individual destination address or subnet.

rate *number* Specify the maximum number of connections per second allowed for an individual destination address or subnet.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

by-destination (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 3096
- Hierarchy Level | 3096
- Description | 3097
- Options | 3097

Required Privilege Level | 3097
Release Information | 3097

Syntax

```
by-destination {
    by-protocol {
        icmp {
            maximum-sessions number;
            packets-rate number;
            session-rate number;
        }
        tcp {
            maximum-sessions number;
            packet-rate number;
            session-rate number;
        }
        udp {
            maximum-sessions number;
            packet-rate number;
            session-rate number;
        }
    }
    maximum-sessions number;
    packet-rate number;
    session-rate number;
    ;
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name limit-session]
```

Description

Configure session limits for individual destination addresses or for individual destination subnets. This protects against network probing attacks and network flooding attacks. You can specify limits for specific protocols (ICMP, TCP, and UDP), or specify limits independent of a protocol. When a session limit is exceeded for a destination, packets to the destination are dropped until the session limit is no longer exceeded.

To specify limits for destination subnets rather than individual addresses, include the `aggregations` statement at the `[edit services screen ids-option screen-name]` hierarchy level.

Options

maximum-sessions *number*	Specify the maximum number of concurrent sessions allowed for an individual destination address or subnet.
packet-rate *number*	Specify the maximum number of packets per second allowed for an individual destination address or subnet.
session-rate *number*	Specify the maximum number of connections per second allowed for an individual destination address or subnet.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

by-pair (IDS MS-DPC)

IN THIS SECTION

- Syntax | 3098
- Hierarchy Level | 3098
- Description | 3098
- Options | 3098
- Required Privilege Level | 3099
- Release Information | 3099

Syntax

```
by-pair {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then session-limit]
```

Description

Apply limit to paired stateful firewall and NAT flows (forward and reverse) when using the MS-DPC.

Options

`hold-time` *seconds*—Length of time for which to stop all new flows once the rate of events exceeds the threshold set by one or more of the `maximum`, `packets`, or `rate` statements.

`maximum` *number*—Maximum number of open sessions per application or IP address.

`packets` *number*—Maximum peak packets per second per application or IP address.

`rate` *number*—Maximum number of sessions per second per application or IP address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in IDS Rules

by-protocol (IDS MS-MPC)

IN THIS SECTION

- Syntax | 3099
- Hierarchy Level | 3100
- Description | 3100
- Options | 3100
- Required Privilege Level | 3101
- Release Information | 3101

Syntax

```
by-protocol {
    icmp {
        maximum number;
```

```
        packets number;
        rate number;
    }
    tcp {
        maximum number;
        packets number;
        rate number;
    }
    udp {
        maximum number;
        packets number;
        rate number;
    }
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then session-limit by-destination],
[edit services ids rule rule-name term term-name then session-limit by-source]
```

Description

Configure the IDS rule session limits for an individual destination or source address or subnet for the specified protocol. This protects against network probing attacks and network flooding attacks. This IDS rule can only be assigned to a service set on an MS-MPC.

When a session limit is exceeded for a source or destination for the protocol, packets from the source or to the destination are dropped until the session limit is no longer exceeded.

To specify limits for destination or source subnets rather than individual addresses, include the `aggregation` statement at the `[edit services ids rule rule-name term term-name then]` hierarchy level.

Options

icmp Apply session limits to ICMP packets.

- **maximum *number*** Specify the maximum number of concurrent ICMP sessions allowed for an individual destination or source address or subnet.

packets *number* Specify the maximum number of ICMP packets per second allowed for an individual destination or source address or subnet.

rate *number* Specify the maximum number of ICMP connections per second allowed for an individual destination or source address or subnet.

tcp Session limits apply to TCP packets.

maximum *number* Specify the maximum number of concurrent TCP sessions allowed for an individual destination or source address or subnet.

packets *number* Specify the maximum number of TCP packets per second allowed for an individual destination or source address or subnet.

rate *number* Specify the maximum number of TCP connections per second allowed for an individual destination or source address or subnet.

udp Session limits apply to UDP packets.

maximum *number* Specify the maximum number of concurrent UDP sessions allowed for an individual destination or source address or subnet.

packets *number* Specify the maximum number of UDP packets per second allowed for an individual destination or source address or subnet.

rate *number* Specify the maximum number of UDP connections per second allowed for an individual destination or source address or subnet.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

by-protocol (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 3102
- Hierarchy Level | 3103
- Description | 3103
- Options | 3103
- Required Privilege Level | 3104
- Release Information | 3104

Syntax

```
by-protocol {
    icmp {
        maximum-sessions number;
        packet-rate number;
        session-rate number;
    }
    tcp {
        maximum-sessions number;
        packet-rate number;
        session-rate number;
    }
    udp {
        maximum-sessions number;
        packet-rate number;
        session-rate number;
    }
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name limit-session by-destination],
[edit services screen ids-option screen-name limit-session by-source]
```

Description

Configure session limits for individual destination or source addresses, or for individual destination or source subnets, for the specified protocol. This protects against network probing attacks and network flooding attacks. When a session limit is exceeded for a source or destination for the protocol, packets from the source or to the destination are dropped until the session limit is no longer exceeded.

To specify limits for destination or source subnets rather than individual addresses, include the `aggregations` statement at the `[edit services screen ids-option screen-name]` hierarchy level.

Options

icmp Apply session limits to ICMP packets.

- **maximum-sessions *number*** Specify the maximum number of concurrent ICMP sessions allowed for individual destination or source addresses, or for individual destination or source subnets.
- **packet-rate *number*** Specify the maximum number of ICMP packets per second allowed for individual destination or source addresses, or for individual destination or source subnets.
- **session-rate *number*** Specify the maximum number of ICMP connections per second allowed for individual destination or source addresses, or for individual destination or source subnets.

tcp Apply session limits to TCP packets.

- **maximum-sessions *number*** Specify the maximum number of concurrent TCP sessions allowed for individual destination or source addresses, or for individual destination or source subnets.
- **packet-rate *number*** Specify the maximum number of TCP packets per second allowed for individual destination or source addresses, or for individual destination or source subnets.

session-rate *number* — Specify the maximum number of TCP connections per second allowed for individual destination or source addresses, or for individual destination or source subnets.

udp — Apply session limits to UDP packets.

maximum-sessions *number* — Specify the maximum number of concurrent UDP sessions allowed for individual destination or source addresses, or for individual destination or source subnets.

packet-rate *number* — Specify the maximum number of UDP packets per second allowed for individual destination or source addresses, or for individual destination or source subnets.

session-rate *number* — Specify the maximum number of UDP connections per second allowed for individual destination or source addresses, or for individual destination or source subnets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

by-source (IDS MS-DPC)

IN THIS SECTION

- Syntax | 3105
- Hierarchy Level | 3105
- Description | 3105
- Options | 3105
- Required Privilege Level | 3106
- Release Information | 3106

Syntax

```
by-source {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then session-limit (IDS MS-DPC)]
```

Description

Apply limit to sessions based on numbers generated from the configured source (IP or subnet) or application when using the MS-DPC.

Options

`hold-time` *seconds*—Length of time for which to stop all new flows once the rate of events exceeds the threshold set by one or more of the `maximum`, `packets`, or `rate` statements.

`maximum` *number*—Maximum number of open sessions per application or IP address.

`packets` *number*—Maximum peak packets per second per application or IP address.

`rate` *number*—Maximum number of sessions per second per application or IP address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in IDS Rules

by-source (IDS MS-PC)

IN THIS SECTION

- Syntax | 3106
- Hierarchy Level | 3107
- Description | 3107
- Options | 3108
- Required Privilege Level | 3108
- Release Information | 3108

Syntax

```
by-source {
    by-protocol {
```

```
        icmp {
            maximum number;
            packets number;
            rate number;
        }
        tcp {
            maximum number;
            packets number;
            rate number;
        }
        udp {
            maximum number;
            packets number;
            rate number;
        }
    }
    maximum number;
    packets number;
    rate number;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then session-limit]
```

Description

Configure the IDS rule session limits for an individual source address or subnet. This protects against network probing attacks and network flooding attacks. When a session limit is exceeded for a source, packets from the source are dropped until the session limit is no longer exceeded. This IDS rule can only be assigned to a service set on an MS-MPC.

When a session limit is exceeded for a source, packets from the source are dropped until the session limit is no longer exceeded.

To specify limits for source subnets rather than individual addresses, include the `aggregation` statement at the `[edit services ids rule rule-name term term-name then]` hierarchy level.

Options

maximum *number* Specify the maximum number of concurrent sessions allowed for an individual source address or subnet.

packets *number* Specify the maximum number of packets per second allowed for an individual source address or subnet.

rate *number* Specify the maximum number of connections per second allowed for an individual source address or subnet.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

by-source (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 3109
- Hierarchy Level | 3109
- Description | 3110
- Options | 3110

Required Privilege Level | 3110
Release Information | 3110

Syntax

```
by-source {
    by-protocol {
        icmp {
            maximum-sessions number;
            packet-rate number;
            session-rate number;
        }
        tcp {
            maximum-sessions number;
            packet-rate number;
            session-rate number;
        }
        udp {
            maximum-sessions number;
            packet-rate number;
            session-rate number;
        }
    }
    maximum-sessions number;
    packet-rate number;
    session-rate number;
    ;
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name limit-session]
```

Description

Configure session limits for individual source addresses or for individual source subnets. This protects against network probing attacks and network flooding attacks. You can specify limits for specific protocols (ICMP, TCP, and UDP), or specify limits independent of a protocol. When a session limit is exceeded for a source, packets from the source are dropped until the session limit is no longer exceeded.

To specify limits for source subnets rather than individual addresses, include the `aggregations` statement at the `[edit services screen ids-option screen-name]` hierarchy level.

Options

maximum-sessions ***number***	Specify the maximum number of concurrent sessions allowed for an individual source address or subnet.
packet-rate ***number***	Specify the maximum number of packets per second allowed for an individual source address or subnet.
session-rate ***number***	Specify the maximum number of connections per second allowed for an individual source address or subnet.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

bypass

IN THIS SECTION

- Syntax | 3111
- Hierarchy Level | 3111
- Description | 3111
- Default | 3111
- Options | 3111
- Required Privilege Level | 3112
- Release Information | 3112

Syntax

```
(bypass | no-bypass);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

Description

Pass or do not pass the delay measurement (DM) value through a node.

Default

If you omit the bypass statement, the default behavior is to disable ODU delay management options.

By default, do not pass the DM value through a node.

Options

bypass Pass the DM value through a node.

no-bypass Do not pass the DM value through a node.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

bypass-aggregate (DDoS)

IN THIS SECTION

- Syntax | 3112
- Hierarchy Level | 3113
- Description | 3113
- Required Privilege Level | 3113
- Release Information | 3113

Syntax

```
bypass-aggregate;
```

Hierarchy Level

- For MX Series routers, T4000 routers, and EX9200 switches:

```
[edit system ddos-protection protocols protocol-group  packet-type]
```

- For QFX10000 Series and QFX5200 switches:

```
[edit system ddos-protection protocols protocol-group  packet-type]
```

> **NOTE**: The `bypass-aggregate` option is not supported on PTX Series routers and QFX10002-60C switches.

Description

Prevent this packet type from being considered by the DDoS aggregate policer. Traffic for the packet type is still included in traffic statistics.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Control Plane DDoS Protection

bypass-queuing-chip

IN THIS SECTION

- Syntax | 3114
- Hierarchy Level | 3114
- Description | 3114
- Required Privilege Level | 3115
- Release Information | 3115

Syntax

```
bypass-queuing-chip;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

On MPCs that contain a queuing chip and on vMX routers, bypass the queuing chip to increase the available bandwidth when rich queuing features are not required. When applied on a vMX router, this option saves a vCPU.

CAUTION: When enabling this option, do not oversubscribe the port bandwidth. Also, after committing this change on a vMX router, you must reboot the MPC that contains the interface.

NOTE: This option requires "flexi-queuing-mode" on page 7233 to be enabled.

NOTE: Once you apply this option:

- You cannot configure either `per-unit-scheduler` or `hierarchical-scheduler` on the interface.
- You can no longer apply schedulers or shapers to the interface.
- Ingress queuing does not work on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Increasing Available Bandwidth on Rich-Queuing MPCs by Bypassing the Queuing Chip

flexible-queuing-mode | 7233

bypass (Signaled LSP)

IN THIS SECTION

- Syntax | 3116
- Hierarchy Level | 3116
- Description | 3116
- Options | 3116
- Required Privilege Level | 3117
- Release Information | 3117

Syntax

```
bypass bypass-name {
    bandwidth bps;
    description text;
    hop-limit number;
    no-cspf;
    path address <strict | loose>;
    priority setup-priority reservation-priority;
    subscription subscription-percentage;
    to address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit protocols rsvp interface interface-name link-protection]
```

Description

Enables you to configure specific bandwidth and path constraints for a bypass LSP. It is possible to individually configure multiple bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints.

If you specify the `bandwidth`, `hop-limit`, and `path` statements for the bypass LSP, these values take precedence over the values configured at the `[edit protocols rsvp interface interface-name link-protection]` hierarchy level. The other attributes (`subscription`, `no-node-protection`, and `optimize-timer`) are inherited from the general constraints.

Options

bypass-name (Required) Specify a name for the bypass LSP. The name can be up to 64 characters.

description Provides a textual description of the bypass LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the `show mpls lsp bypass detail` command and has no effect on the operation of the bypass LSP. The description text can be no more than 80 characters in length.

subscription ***subscription-percentage***

(Optional) Specify the subscription percentage per manual bypass LSP. The subscription percentage configured under a particular manual bypass LSP overrides the subscription percentage configured commonly for all manual bypass LSPs under an interface.

- **Range:** 0 through 65000

to ***address***

(Required) Specify the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`description` option was added in Junos OS Release 10.4.

`subscription` *`subscription-percentage`* option introduced in Junos OS Release 19.2R1 on all platforms.

RELATED DOCUMENTATION

Configuring Bypass LSPs

bypass (Static LSP)

IN THIS SECTION

Syntax | **3118**

- Hierarchy Level | 3118
- Description | 3118
- Required Privilege Level | 3118
- Release Information | 3119

Syntax

```
bypass bypass-name {
    bandwidth bps;
    description string;
    next-hop (address | interface-name | address/interface-name);
    next-table string
    push out-label;
    to address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name]
```

Description

Configure specific bandwidth and path constraints for a bypass ingress LSP. It is possible to configure multiple bypass LSPs individually. If you do not, they all share the same path and bandwidth constraints.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring Static LSPs

bypass-traffic-on-exceeding-flow-limits

IN THIS SECTION

- Syntax | 3119
- Hierarchy Level | 3119
- Description | 3119
- Required Privilege Level | 3120
- Release Information | 3120

Syntax

```
bypass-traffic-on-exceeding-flow-limits;
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Description

Bypass traffic when exceeding the maximum flow limit.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

Statement introduced in Junos OS Release 19.3R2 on MX240, MX480 and MX960 routers using the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring Service Sets to be Applied to Services Interfaces

bypass-traffic-on-pic-failure

IN THIS SECTION

- Syntax | 3120
- Hierarchy Level | 3121
- Description | 3121
- Required Privilege Level | 3121
- Release Information | 3121

Syntax

```
bypass-traffic-on-pic-failure;
```

Hierarchy Level

```
[edit services service-set service-set-options]
```

Description

When the MultiServices PIC configured for a service set is either administratively taken offline or undergoes a failure, all the traffic entering the configured interface with an IDP service set would be dropped without notification. To avoid this traffic loss, include the `bypass-traffic-on-pic-failure` statement. When this statement is configured, the affected packets are forwarded in the event of a MultiServices PIC failure or offlining, as though interface-style services were not configured.

This issue applies only to Junos Application Aware (previously known as Dynamic Application Awareness) configurations with IDP service sets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring Service Sets to be Applied to Services Interfaces

bytes (Dynamic Traffic Shaping)

IN THIS SECTION

- Syntax | 3122
- Hierarchy Level | 3122

- Description | 3122
- Options | 3122
- Required Privilege Level | 3122
- Release Information | 3123

Syntax

```
bytes bytes | $junos-cos-byte-adjust;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name overhead-accounting],
[edit class-of-service traffic-control-profiles profile-name overhead-accounting]
```

Description

Configure the number of overhead bytes.

Options

bytes—Byte adjustment value for the `cell-mode` or `frame-mode` shaping options. This can be the predefined variable `$junos-cos-byte-adjust`, which is the variable for byte adjustment that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

> **BEST PRACTICE**: We recommend using the `cell-mode cell-mode-bytes cell-mode-bytes` option or the `frame-mode frame-mode-bytes frame-mode-bytes` option rather than the `bytes` option.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

CoS Adjustment Control Profiles Overview

Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates

Bandwidth Management for Downstream Traffic in Edge Networks Overview

egress-shaping-overhead

bytes (otn-options)

IN THIS SECTION

- Syntax | 3123
- Hierarchy Level | 3123
- Description | 3124
- Options | 3124
- Required Privilege Level | 3124
- Release Information | 3124

Syntax

```
bytes transmit-payload-type value;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify the transmit payload type on OTN header bytes.

Options

value—Transmit payload type.

- **Range:** 0 through 255 bytes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

C

IN THIS SECTION

- c-timeout | 3152
- ca-chain (REST API) | 3154
- ca-identity | 3155
- ca-name | 3156

- ca-profile (Security PKI) | **3157**
- ca-profile (Services) | **3160**
- ca-type | **3161**
- cable-length (Congestion Notification) | **3164**
- cac (TDF Gateway) | **3165**
- cache | **3167**
- cache (Security Log) | **3168**
- cache-size | **3170**
- cache-timeout-negative | **3171**
- cacheable (Application Identification) | **3173**
- cak | **3174**
- cak (MX Series) | **3176**
- calculation-weight | **3177**
- call-flood | **3179**
- call-rate-statistics | **3180**
- callback | **3182**
- callback-wait-period | **3183**
- called-station-id | **3185**
- called-station-id (OCS Partition) | **3186**
- caller | **3188**
- calling-number | **3189**
- calling-number (L2TP Tunnel Switching) | **3190**
- calling-station-id | **3192**
- calling-station-id-format (Subscriber Management) | **3193**
- candidate-uplink-port-policy (satellite-policies) | **3196**
- capabilities (Protocols RIFT) | **3197**
- captive-portal | **3199**
- captive-portal-content-delivery-profile (Services) | **3202**
- captive-portal (Services UAC) | **3203**
- captive-portal (Services UAC Policy) | **3205**
- capture-file (Security) | **3206**
- capture-group | **3208**

- casb | **3210**
- casb-policy | **3212**
- cascade-port | **3214**
- cascade-ports | **3216**
- cascade-ports (Junos Fusion Satellite Device Cluster) | **3218**
- categories | **3219**
- category | **3221**
- category (Security Dynamic Attack Group) | **3223**
- category (Security Logging) | **3224**
- category (Security Web Filtering) | **3226**
- category (System Logging) | **3236**
- cbr | **3238**
- cbr rate | **3240**
- ccc | **3241**
- ccc (Routing Policy Condition) | **3242**
- cdvt | **3244**
- cell-mode (Dynamic Traffic Shaping) | **3245**
- cell-overhead | **3248**
- cellular-options | **3249**
- centralized | **3250**
- cert-file (Scripts) | **3252**
- certificate | **3254**
- certificate-authority (REST API) | **3256**
- certificate-id | **3257**
- certificate (Juniper Secure Connect) | **3259**
- certificate-verification | **3260**
- certificates | **3262**
- certification-authority | **3263**
- ces-psn-channel (tdm-options) | **3265**
- cfeb | **3267**
- cflowd (Discard Accounting) | **3268**
- cflowd (Flow Monitoring) | **3270**

- cfp-to-et | **3272**
- cgn-pic | **3273**
- chain | **3274**
- chain-order (Application Identification) | **3276**
- chained-composite-next-hop | **3277**
- challenge-length (Static and Dynamic PPP) | **3281**
- challenge-password | **3283**
- change-configuration (Event Policy) | **3284**
- channel | **3286**
- channel-speed | **3287**
- chap | **3289**
- chap (Dynamic PPP) | **3291**
- chap (L2TP) | **3292**
- charging-id (OCS Partition) | **3293**
- charging-service-list | **3295**
- chassis | **3296**
- chassis-id (LLDP) | **3301**
- chassis-id (Aggregated Ethernet Interface) | **3302**
- chassis (QFabric System) | **3303**
- chassis (Subscriber Limits) | **3306**
- check-bytes (Application Identification) | **3308**
- check-zero | **3309**
- checksum (Scripts) | **3311**
- checksum-validate | **3313**
- child-inactivity-timeout | **3314**
- children | **3315**
- cipher-list (REST API) | **3318**
- cipher-suite (MACsec) | **3320**
- circuit-id | **3322**
- circuit-id (DHCP Relay Agent) | **3324**
- circuit-id (VLAN Authentication Username) | **3328**
- circuit-type | **3329**

- circuit-type (DHCP Local Server) | **3330**
- circuit-type (DHCP Relay Agent) | **3332**
- cisco-interoperability | **3334**
- cisco-nas-port-info (L2TP Tunnel Switching) | **3335**
- civic-based | **3337**
- ckn | **3339**
- ckn (MX Series) | **3341**
- class | **3342**
- class (Unified Edge) | **3345**
- class (CoS-Based Forwarding) | **3346**
- class (Defining Login Classes) | **3347**
- class (Forwarding Class Sets) | **3361**
- class (Forwarding classes) | **3362**
- class (Forwarding Classes) | **3364**
- class-of-service (Dynamic Profiles) | **3366**
- class-of-service (Protocols MPLS) | **3369**
- class-of-service (Protocols RSVP) | **3370**
- class-of-service (saegw system) | **3372**
- class-usage-profile | **3373**
- classification-key (DHCP Local Server) | **3375**
- classification-key (DHCP Relay Agent) | **3376**
- classification-override | **3379**
- classifier | **3380**
- classifiers (VRF) | **3382**
- classifiers (Class of Service) | **3383**
- classifiers (Definition) | **3386**
- classifiers (Definition) | **3388**
- classifiers (Dynamic CoS Application) | **3390**
- classifiers (Interfaces) | **3391**
- classifiers (Physical Interface) | **3393**
- classifiers (Routing Instance) | **3395**
- clat-ipv6-prefix-length | **3396**

- clat-prefix | 3398
- clat-prefix (Source NAT Next Gen Services) | 3399
- cleanup-interval (Accounting Options) | 3400
- clear | 3402
- clear access-security router-advertisement statistics | 3403
- clear access-security slaac-snooping binding | 3404
- clear access-security slaac-snooping statistics | 3406
- clear accounting server statistics archival-transfer | 3407
- clear-action (CFM) | 3409
- clear amt statistics | 3410
- clear amt tunnel | 3412
- clear ancp access-loop | 3414
- clear ancp neighbor | 3416
- clear ancp statistics | 3418
- clear ancp subscriber | 3421
- clear arp | 3424
- clear arp inspection statistics | 3426
- clear auto-configuration interfaces | 3428
- clear auto-configuration interfaces interface-set | 3430
- clear bfd adaptation | 3432
- clear bfd session | 3434
- clear bgp damping | 3436
- clear bgp neighbor | 3437
- clear bgp table | 3441
- clear bridge mac-ip-table | 3443
- clear bridge mac-table | 3444
- clear bridge recovery-timeout | 3446
- clear bridge statistics | 3447
- clear captive-portal | 3450
- clear chassis alarms fabric degraded | 3453
- clear chassis cluster control-plane statistics | 3455
- clear chassis cluster data-plane statistics | 3457

- clear chassis cluster failover-count | 3458
- clear chassis cluster ip-monitoring failure-count | 3460
- clear chassis cluster ip monitoring failure count ip address | 3462
- clear chassis cluster statistics | 3463
- clear chassis display message | 3465
- clear chassis fpc errors | 3468
- clear chassis hardware history | 3470
- clear chassis high-availability data-plane statistics | 3472
- clear chassis high-availability information | 3473
- clear chassis sib errors | 3474
- clear class-of-service application-traffic-control counter | 3476
- clear class-of-service application-traffic-control rate-limiters | 3478
- clear class-of-service application-traffic-control statistics rule | 3479
- clear cli logical-system | 3481
- clear ddos-protection protocols | 3482
- clear dhcp client binding | 3484
- clear dhcp client statistics | 3486
- clear dhcp relay active-leasequery statistics | 3488
- clear dhcp relay binding (Security) | 3491
- clear dhcp relay binding | 3493
- clear dhcp relay lockout-entries | 3496
- clear dhcp relay statistics (Security) | 3499
- clear dhcp relay statistics | 3500
- clear dhcp-security binding | 3503
- clear dhcp-security ipv6 binding | 3505
- clear dhcp server active-leasequery statistics | 3506
- clear dhcp server binding | 3508
- clear dhcp server lockout-entries | 3512
- clear dhcp server statistics | 3515
- clear dhcp snooping binding | 3517
- clear dhcp snooping statistics | 3519
- clear dhcpv6 client binding | 3521

- clear dhcpv6 client statistics | 3522
- clear dhcpv6 relay active-leasequery statistics | 3524
- clear dhcpv6 relay binding | 3527
- clear dhcpv6 relay lockout-entries | 3531
- clear dhcpv6 relay statistics | 3534
- clear dhcpv6 server active-leasequery statistics | 3536
- clear dhcpv6 server binding | 3538
- clear dhcpv6 server binding (Local Server) | 3542
- clear dhcpv6 server lockout-entries | 3543
- clear dhcpv6 server statistics | 3546
- clear dhcpv6 server statistics (Local Server) | 3548
- clear dhcpv6 snooping binding | 3549
- clear dhcpv6 snooping statistics | 3551
- clear diameter function statistics | 3553
- clear diameter peer | 3554
- clear-dont-fragment-bit (fti) | 3556
- clear-dont-fragment-bit (Interfaces GRE Tunnels) | 3557
- clear-dont-fragment-bit (NAT Next Gen Services) | 3559
- clear-dont-fragment-bit (Services IPsec VPN) | 3560
- clear-dont-fragment-bit (Services NAT Options) | 3561
- clear-dont-fragment-bit (Services Service Set) | 3562
- clear dot1x | 3564
- clear edge-virtual-bridging | 3566
- clear error bpdu interface | 3567
- clear error mac-rewrite | 3568
- clear ethernet-switching bpdu-error interface | 3570
- clear ethernet-switching evpn arp-table | 3572
- clear ethernet-switching evpn nd-statistics | 3573
- clear ethernet-switching evpn nd-table | 3574
- clear ethernet-switching layer2-protocol-tunneling error | 3576
- clear ethernet-switching layer2-protocol-tunneling statistics | 3577
- clear ethernet-switching mac-ip-table | 3579

- clear ethernet-switching port-error | 3580
- clear ethernet-switching recovery-timeout | 3582
- clear ethernet-switching recovery-timeout (SRX) | 3583
- clear ethernet-switching statistics mac-learning | 3584
- clear ethernet-switching table persistent-learning | 3586
- clear evpn duplicate-mac-suppression | 3588
- clear evpn mac-ip-table | 3589
- clear evpn nd-table | 3590
- clear evpn statistics | 3592
- clear extensible-subscriber-services counters | 3593
- clear extensible-subscriber-services sessions | 3594
- clear fibre-channel fc2 statistics | 3596
- clear fibre-channel fip enode | 3597
- clear fibre-channel fip statistics | 3599
- clear fibre-channel fip vn-port | 3600
- clear fibre-channel flogi statistics | 3602
- clear fibre-channel proxy statistics | 3603
- clear fip snooping enode | 3604
- clear fip snooping statistics | 3606
- clear fip snooping vlan | 3608
- clear fip vlan-discovery statistics | 3610
- clear firewall (QFX) | 3611
- clear firewall | 3613
- clear forwarding-options port-mirroring instance family | 3616
- clear gtp tunnels | 3618
- clear igmp membership | 3619
- clear igmp snooping membership | 3623
- clear igmp snooping statistics | 3625
- clear igmp statistics | 3627
- clear-ike-sas-on-pic-restart | 3631
- clear ike security-associations | 3632
- clear ilmi statistics | 3634

- clear inband-flow-telemetry stats | **3635**
- clear interfaces aeX forwarding-options load-balance state | **3637**
- clear interfaces aggregate forwarding-options load-balance state | **3638**
- clear interfaces interface-set statistics | **3640**
- clear interfaces interval | **3641**
- clear interfaces lb-stats af | **3643**
- clear interfaces mac-database | **3648**
- clear interfaces mac-database statistics | **3649**
- clear interfaces statistics | **3651**
- clear interfaces statistics swfabx | **3653**
- clear interfaces transport pm | **3654**
- clear-ipsec-sas-on-pic-restart | **3656**
- clear ipsec security-associations | **3657**
- clear ipv6 neighbors | **3659**
- clear ipv6 neighbors (SRX Series) | **3661**
- clear ipv6 router-advertisement | **3663**
- clear isis adjacency | **3664**
- clear isis database | **3667**
- clear isis overload | **3670**
- clear isis spring traffic-statistics | **3673**
- clear isis statistics | **3674**
- clear l2-learning mac-move-buffer | **3677**
- clear lacp statistics interfaces | **3679**
- clear ldp neighbor | **3680**
- clear ldp session | **3682**
- clear ldp statistics | **3683**
- clear lldp neighbors | **3685**
- clear lldp statistics | **3686**
- clear log | **3688**
- clear log (JDM) | **3690**
- clear log (Next Gen Services) | **3692**
- clear loop-detect enhanced interface | **3693**

- clear mac-vrf forwarding mac-ip-table | 3695
- clear mac-vrf forwarding mac-learning-log | 3697
- clear mac-vrf forwarding mac-table | 3698
- clear mac-vrf forwarding recovery-timeout | 3700
- clear mac-vrf forwarding redundancy-group | 3701
- clear mld membership | 3702
- clear mld snooping membership | 3704
- clear mld snooping statistics | 3705
- clear mld statistics | 3707
- clear monitor security flow filter | 3709
- clear mpls container-lsp | 3711
- clear mpls lsp | 3713
- clear mpls tunnel manager statistics | 3716
- clear msdp cache | 3717
- clear msdp statistics | 3719
- clear multicast bandwidth-admission | 3720
- clear multicast forwarding-cache | 3723
- clear multicast scope | 3724
- clear multicast sessions | 3726
- clear multicast snooping statistics | 3728
- clear multicast statistics | 3730
- clear mvrp statistics | 3732
- clear neighbor-discovery-inspection statistics | 3734
- clear network-access aaa statistics | 3736
- clear network-access aaa subscriber | 3741
- clear network-access gx-plus replay | 3743
- clear network-access gx-plus statistics | 3745
- clear network access nas-5g statistics | 3746
- clear network-access ocs statistics | 3747
- clear network-access pcrf | 3749
- clear network-access radsec state | 3751
- clear network-access radsec statistics | 3752

- clear network-access requests pending | 3753
- clear network-access requests statistics | 3755
- clear network-access securid-node-secret-file | 3756
- clear network-access aaa saml assertion-cache | 3757
- clear network-access aaa saml idp-metadata | 3759
- clear node reboot | 3760
- clear oam ethernet connectivity-fault-management continuity-measurement | 3761
- clear oam ethernet connectivity-fault-management delay-statistics | 3763
- clear oam ethernet connectivity-fault-management event | 3765
- clear oam ethernet connectivity-fault-management linktrace path-database | 3767
- clear oam ethernet connectivity-fault-management loss-statistics | 3768
- clear oam ethernet connectivity-fault-management path-database | 3770
- clear oam ethernet connectivity-fault-management policer | 3772
- clear oam ethernet connectivity-fault-management sla-iterator-statistics | 3774
- clear oam ethernet connectivity-fault-management statistics | 3775
- clear oam ethernet connectivity-fault-management synthetic-loss-measurement | 3777
- clear oam ethernet link-fault-management state | 3779
- clear oam ethernet link-fault-management statistics | 3781
- clear-on-abort (DHCP Local Server) | 3782
- clear (ospf | ospf3) database | 3784
- clear (ospf | ospf3) database-protection | 3788
- clear (ospf | ospf3) io-statistics | 3789
- clear (ospf | ospf3) neighbor | 3791
- clear (ospf | ospf3) overload | 3793
- clear (ospf | ospf3) statistics | 3795
- clear ovsdb commit failures | 3798
- clear passive-monitoring statistics | 3800
- clear path-computation-client statistics | 3802
- clear performance-monitoring mpls lsp | 3803
- clear pfe sflow egress-mcast-stats | 3805
- clear pim join | 3806
- clear pim join-distribution | 3809

- clear pim register | 3812
- clear pim snooping join | 3814
- clear pim snooping statistics | 3817
- clear pim statistics | 3819
- clear policy statistics | 3823
- clear pppoe lockout | 3824
- clear pppoe lockout atm-identifier | 3827
- clear pppoe lockout vlan-identifier | 3829
- clear pppoe statistics | 3832
- clear protection-group ethernet-ring statistics | 3834
- clear ptp statistics stream | 3836
- clear remote-device-management statistics | 3837
- clear rip general-statistics | 3838
- clear rip statistics | 3840
- clear ripng general-statistics | 3842
- clear ripng statistics | 3844
- clear rsvp session | 3846
- clear rsvp statistics | 3848
- clear security advance-policy-based-routing sla statistics | 3850
- clear security alarms | 3851
- clear security alg h323 counters | 3855
- clear security alg ike-esp-nat | 3857
- clear security alg sccp calls | 3858
- clear security alg sccp counters | 3860
- clear security alg sip calls | 3862
- clear security alg sip counters | 3863
- clear security application-firewall rule-set statistics | 3865
- clear security application-firewall rule-set statistics logical-system | 3866
- clear security datapath-debug counters | 3868
- clear security dns-cache | 3869
- clear security dynamic-vpn all | 3871
- clear security dynamic-vpn user | 3872

- clear security firewall-authentication history | 3874
- clear security firewall-authentication history address | 3876
- clear security firewall-authentication history identifier | 3878
- clear security firewall-authentication users | 3881
- clear security firewall-authentication users address | 3883
- clear security firewall-authentication users identifier | 3885
- clear security flow ip-action | 3887
- clear security flow session | 3890
- clear security flow session application | 3893
- clear security flow session application-traffic-control | 3896
- clear security flow session conn-tag | 3899
- clear security flow session destination-port | 3901
- clear security flow session destination-prefix | 3903
- clear security flow session family | 3905
- clear security flow session IDP | 3907
- clear security flow session interface | 3910
- clear security flow session nat | 3912
- clear security flow session protocol | 3915
- clear security flow session resource-manager | 3918
- clear security flow session services-offload | 3920
- clear security flow session session-identifier | 3924
- clear security flow session source-port | 3926
- clear security flow session source-prefix | 3928
- clear security flow session tenant | 3931
- clear security flow session tunnel | 3933
- clear security forward-options mirror filter | 3935
- clear security group-vpn member group | 3937
- clear security group-vpn member ike security-associations | 3938
- clear security group-vpn member ipsec security-associations | 3939
- clear security group-vpn member ipsec security-associations statistics | 3941
- clear security group-vpn member ipsec statistics | 3942
- clear security group-vpn member kek security-associations | 3943

- clear security group-vpn server | 3944
- clear security group-vpn server server-cluster statistics | 3946
- clear security group-vpn server statistics | 3947
- clear security gtp counters | 3949
- clear security gtp sctp association | 3952
- clear security gtp sctp counters | 3954
- clear security host-vpn security-associations | 3956
- clear security idp | 3957
- clear security idp attack table | 3959
- clear security idp counters action | 3961
- clear security idp counters application-identification | 3962
- clear security idp counters dfa | 3964
- clear security idp counters flow | 3965
- clear security idp counters http-decoder | 3966
- clear security idp counters ips | 3968
- clear security idp counters log | 3969
- clear security idp counters memory | 3971
- clear security idp counters packet | 3972
- clear security idp counters packet-log | 3974
- clear security idp counters pdf-decoder | 3975
- clear security idp counters policy-manager | 3977
- clear security idp counters ssl-inspection | 3978
- clear security idp counters tcp-reassembler | 3980
- clear security idp ssl-inspection session-id-cache | 3981
- clear security ike active-peer aaa-username | 3983
- clear security ike peers | 3984
- clear security ike peers statistics | 3985
- clear security ike respond-bad-spi-count | 3987
- clear security ike security-associations | 3988
- clear security ike stats | 3990
- clear security ipsec security-associations | 3993
- clear security ipsec statistics | 3995

- clear security ipsec tunnel-events-statistics | 3997
- clear security log | 3998
- clear security log file | 4001
- clear security log stream file | 4002
- clear security mka statistics | 4004
- clear security mka statistics (MX Series) | 4006
- clear security nat incoming-table | 4007
- clear security nat source persistent-nat-table | 4009
- clear security nat statistics destination pool | 4010
- clear security nat statistics destination rule | 4011
- clear security nat statistics source pool | 4013
- clear security nat statistics source rule | 4014
- clear security nat statistics static rule | 4015
- clear security pki ca-certificate | 4017
- clear security pki certificate-request | 4018
- clear security pki crl | 4020
- clear security pki key-pair | 4021
- clear security pki key-pair (Local Certificate) | 4023
- clear security pki local-certificate | 4024
- clear security pki local-certificate | 4026
- clear security pki local-certificate (Device) | 4027
- clear security pki node-local key-pair | 4030
- clear security pki node-local local-certificate | 4031
- clear security pki statistics | 4033
- clear security policies file-serialization counters | 4034
- clear security policies hit-count | 4035
- clear security policies statistics | 4037
- clear security screen statistics | 4038
- clear security screen statistics interface | 4039
- clear security screen statistics zone | 4042
- clear security ssh key-pair-identity | 4044
- clear security tcp-encap statistics | 4046

- clear security user-identification local-authentication-table | 4047
- clear security utm anti-spam statistics | 4048
- clear security utm antivirus statistics | 4051
- clear security utm content-filtering statistics | 4054
- clear security utm session | 4057
- clear security utm web-filtering statistics | 4058
- clear service l2tp session routing-instance | 4062
- clear service user-identification identity-management counter | 4063
- clear services accounting statistics inline-jflow | 4064
- clear services advanced-anti-malware statistics | 4066
- clear services agf nas statistics | 4068
- clear services agf ngap amf | 4069
- clear services alg statistics | 4070
- clear services application-aware-access-list statistics | 4071
- clear services application-identification application-statistics | 4072
- clear services application-identification application-statistics cumulative | 4073
- clear services application-identification application-statistics interval | 4074
- clear services application-identification application-system-cache | 4075
- clear services application-identification application-system-cache (Junos OS) | 4076
- clear services application-identification counter (Values) | 4078
- clear services application-identification packet-capture counters | 4081
- clear services application-identification statistics | 4082
- clear services captive-portal-content-delivery statistics | 4085
- clear services cos statistics | 4086
- clear services crtp statistics | 4088
- clear services dns-filtering cache | 4089
- clear services dynamic-flow-capture | 4093
- clear services flow-collector statistics | 4094
- clear services icap-redirect statistic | 4096
- clear services ids | 4098
- clear services ids destination-table | 4100
- clear services ids pair-table | 4101

- clear services ids source-table | 4103
- clear services inline-monitoring statistics | 4105
- clear services inline nat pool | 4106
- clear services inline nat statistics | 4107
- clear services inline softwire statistics | 4109
- clear services ipsec-vpn certificates | 4110
- clear services ipsec-vpn ike security-associations | 4112
- clear services ipsec-vpn ipsec security-associations | 4113
- clear services ipsec-vpn ipsec statistics | 4115
- clear services l2tp destination | 4117
- clear services l2tp destination lockout | 4119
- clear services l2tp destination statistics | 4121
- clear services l2tp multilink | 4123
- clear services l2tp session | 4125
- clear services l2tp session statistics | 4129
- clear services l2tp tunnel | 4131
- clear services l2tp tunnel statistics | 4134
- clear services lrf collector statistics | 4136
- clear services lrf statistics | 4138
- clear services mobile-edge sessions | 4139
- clear services monitoring rfc2544 | 4140
- clear services monitoring twamp server control-connection | 4141
- clear services nat flows | 4142
- clear services nat mappings | 4144
- clear services nat mappings app | 4147
- clear services nat mappings eim | 4149
- clear services nat mappings pcp | 4151
- clear services nat source mappings | 4154
- clear services redundancy-set last-saved-state id | 4157
- clear services rpm rfc2544-benchmarking | 4158
- clear services rpm twamp server connection | 4159
- clear services security-intelligence dns-statistics | 4160

- clear services security-metadata-streaming | **4161**
- clear services service-set statistics ids drops | **4164**
- clear services service-sets statistics ids session-limits counters | **4165**
- clear services service-sets statistics integrity-drops | **4166**
- clear services service-sets statistics jflow-log | **4168**
- clear services service-sets statistics packet-drops | **4170**
- clear services service-sets statistics syslog | **4171**
- clear services sessions | **4173**
- clear services sessions analysis | **4178**
- clear services ssl proxy statistics | **4179**
- clear services stateful-firewall flows | **4180**
- clear services stateful-firewall sip-call | **4183**
- clear services stateful-firewall sip-register | **4187**
- clear services stateful-firewall statistics | **4191**
- clear services subscriber analysis | **4193**
- clear services user-identification active-directory-access | **4194**
- clear services user-identification authentication-table | **4195**
- clear services user-identification logical-domain-identity-management counters | **4197**
- clear services video-monitoring mdi errors fpc-slot | **4199**
- clear services video-monitoring mdi statistics fpc-slot | **4200**
- clear-services-web-filter-dns-feed-statistics-feed-name | **4201**
- clear services web-filter statistics profile | **4202**
- clear sflow collector statistics (QFX Series) | **4204**
- clear sflow collectors statistics | **4206**
- clear snmp history | **4207**
- clear snmp statistics | **4208**
- clear spanning-tree protocol-migration | **4210**
- clear spanning-tree statistics | **4212**
- clear spanning-tree statistics bridge | **4214**
- clear spanning-tree stp-buffer | **4215**
- clear synchronous-ethernet esmc statistics | **4216**
- clear synchronous-ethernet wait-to-restore | **4217**

- clear system commit | 4219
- clear system commit prepared | 4221
- clear system errors | 4223
- clear system login lockout | 4225
- clear system reboot | 4226
- clear system services dhcp binding | 4230
- clear system services dhcp conflict | 4232
- clear system services dhcp statistics | 4233
- clear system services dns dns-proxy | 4235
- clear system subscriber-management statistics | 4236
- clear tcp-forwarding connections | 4237
- clear tcp-forwarding statistics | 4241
- clear-threshold | 4245
- clear trace | 4246
- clear unified-edge tdf aaa radius client statistics | 4248
- clear unified-edge tdf aaa radius network-element statistics | 4250
- clear unified-edge tdf aaa radius server statistics | 4252
- clear unified-edge tdf aaa radius snoop-segment statistics | 4253
- clear unified-edge tdf aaa statistics | 4255
- clear unified-edge tdf address-assignment pool | 4257
- clear unified-edge tdf address-assignment statistics | 4259
- clear unified-edge tdf call-admission-control statistics | 4261
- clear unified-edge tdf diameter network-element statistics | 4262
- clear unified-edge tdf diameter pcc-gx statistics | 4264
- clear unified-edge tdf diameter peer statistics | 4266
- clear unified-edge tdf statistics | 4268
- clear unified-edge tdf subscribers | 4269
- clear unified-edge tdf subscribers peer | 4271
- clear validation database | 4273
- clear validation session | 4275
- clear validation statistics | 4277
- clear virtual-chassis heartbeat (MX Series Virtual Chassis) | 4278

- clear virtual-chassis protocol statistics | **4281**
- clear virtual-chassis vc-port statistics | **4283**
- clear vpls mac-address | **4285**
- clear vpls mac-move-action | **4287**
- clear vpls mac-table | **4288**
- clear vrrp | **4290**
- client | **4292**
- client (TWAMP) | **4295**
- client (interfaces) | **4299**
- client (Unified Edge) | **4300**
- client-address | **4302**
- client-config (Juniper Secure Connect) | **4303**
- client-delegate-probes | **4307**
- client-discover-match (DHCP Local Server) | **4308**
- client-discover-match (DHCP Relay Agent) | **4310**
- client-group | **4312**
- client-ia-type | **4313**
- client-id (DHCP Local Server) | **4315**
- client-id (DHCP Relay Agent) | **4317**
- client-identifier (DHCP Client) | **4319**
- client-identifier (DHCPv6 Client) | **4321**
- client-identifier (Interfaces) | **4322**
- client-identity (NETCONF TLS) | **4324**
- client-idle-timeout (Access Profile) | **4326**
- client-list | **4328**
- client-list (SNMP) | **4329**
- client-list-name | **4331**
- client-name-filter | **4332**
- client-negotiation-match (DHCPv6 Local Server) | **4333**
- client-negotiation-match (DHCPv6 Relay Agent) | **4335**
- client-response-ttl | **4336**
- client-session-timeout (Access Profile) | **4338**

- client (System Services) | **4339**
- client-type | **4340**
- client-type (Subscriber Limits) | **4342**
- clients (SNMP Community) | **4344**
- clients | **4345**
- clients (Security) | **4347**
- clns-routing | **4349**
- clns-updown-compatibility | **4350**
- clock-class | **4352**
- clock-class-to-quality-level-mapping | **4353**
- clock-ipv6-client | **4355**
- clock-ipv6-source | **4356**
- clock-mode | **4357**
- clock-mode (Clock Synchronization) | **4359**
- clock-rate | **4360**
- clock-source (hybrid) | **4362**
- clocking-mode | **4364**
- close-timeout | **4365**
- cluster | **4367**
- cluster (Chassis) | **4369**
- cluster-id (Junos Fusion Satellite Device Cluster) | **4373**
- cluster (Junos Fusion) | **4375**
- cmcast-joins-limit-inet (MVPN Selective Tunnels) | **4376**
- cmcast-joins-limit-inet6 (MVPN Selective Tunnels) | **4379**
- coa-accounting (AAA Profile) | **4381**
- coa-immediate-update | **4382**
- coa-no-override service-class-attribute | **4384**
- code | **4385**
- code (IDP) | **4386**
- code (AAA Profile) | **4388**
- code (Application Identification) | **4389**
- code-point | **4390**

- code-point-aliases | 4392
- code-point (Congestion Notification) | 4393
- code-point (Fibre Channel Interfaces) | 4395
- code-point (Input Congestion Notification) | 4397
- code-point (Output Congestion Notification) | 4398
- code-point (Rewrite Rules) | 4400
- code-points (Application Maps) | 4402
- code-points (Classifiers) | 4403
- code-points (Classifiers) | 4405
- collapsed-forward | 4406
- collector | 4408
- collector (Flow Collection) | 4410
- collector (Analytics) | 4412
- collector (Flow Monitoring Logs for NAT) | 4413
- collector (Flow Template Profiles for NAT) | 4415
- collector-group (Flow Monitoring Logs for NAT) | 4416
- collector-group (Flow Template Profiles for NAT) | 4418
- collector (Inline Monitoring) | 4420
- collector (LRF Profile) | 4422
- collector (LRF Rule) | 4424
- color | 4425
- color-aware | 4428
- color-blind | 4429
- command (Op Scripts) | 4431
- commands (Event Policy Change Configuration) | 4432
- commands (Event Policy Execute Commands) | 4434
- commit | 4436
- commit activate | 4442
- commit (Commit Scripts) | 4444
- commit-delay | 4446
- commit-interval | 4447
- commit-options | 4449

- commit prepare | 4451
- commit-synchronize-server | 4452
- commit (System) | 4455
- committed-burst-size | 4458
- committed-information-rate | 4460
- community | 4462
- community-name | 4464
- community (Policy Options) | 4465
- community (Protocols Layer 2 Circuit) | 4469
- community (RMON) | 4471
- community (Routing Options) | 4472
- community (SNMP) | 4475
- community-vlan | 4477
- community-vlans (MX Series) | 4478
- compatibility (Application Identification) | 4480
- compatibility-version | 4481
- compliance (Juniper Secure Connect) | 4483
- compress (Accounting Options) | 4488
- compress-configuration-files (System) | 4489
- compression | 4491
- compression-device (Interfaces) | 4492
- compression-device (Interfaces) | 4493
- compression (PPP Properties) | 4494
- compute-options | 4496
- compute-profile | 4497
- concurrent-data-sessions | 4500
- condition | 4502
- conditional-match | 4504
- conditional-metric | 4506
- confederation | 4508
- configuration-database | 4510
- configuration-database (Enhanced Subscriber Management) | 4513

- configuration-file | **4514**
- configuration-name | **4516**
- configuration-servers | **4517**
- configuration-synchronize (Chassis Cluster) | **4519**
- configure | **4520**
- configured-flow-control | **4523**
- congestion-notification-profile | **4525**
- connect-actively | **4529**
- connection-flood | **4531**
- connection (Identity Management Advanced Query) | **4533**
- connection-limit | **4537**
- connection-limit (REST API) | **4538**
- connection-protection | **4540**
- connection-protection-tlv | **4541**
- connections (Circuits) | **4542**
- connections (Host VPN) | **4544**
- connections (MPLS) | **4547**
- connectivity-association | **4548**
- connectivity-association (MACsec Interfaces) | **4551**
- connectivity-association (MACsec Interfaces for MX Series) | **4553**
- connectivity-association (MX Series) | **4554**
- connectivity-fault-management | **4556**
- connectivity-fault-management-switches | **4560**
- connectivity-type | **4563**
- console (System Logging) | **4565**
- console (System Ports) | **4567**
- constant | **4570**
- constituent-list | **4571**
- cont-stats-collection-interval | **4573**
- contact (SNMP) | **4575**
- container-devices | **4576**
- container-label-switched-path | **4577**

- content-destination | **4579**
- content-filtering (Logical System Security Feature Profile) | **4580**
- content-filtering (Security Content Security Policy) | **4583**
- content-filtering (Security Feature Profile) | **4586**
- Content Security | **4589**
- Content Security-policy (Application Services) | **4600**
- content-size | **4601**
- content-size-limit | **4603**
- content-size (Security Antivirus Sophos Engine) | **4604**
- context (Application Identification) | **4606**
- context-identifier (Protocols IS-IS) | **4611**
- context-identifier (Protocols OSPF) | **4613**
- context (Security Custom Attack) | **4614**
- context (SNMPv3) | **4615**
- continuity-check | **4616**
- continuity-check (EX Series Switch Only) | **4618**
- control-channel | **4620**
- control-channel (Protocols Link Management Peer) | **4621**
- control-channel (Protocols OAM) | **4623**
- control-connection (Junos OS) | **4624**
- control-connection (Junos OS Evolved) | **4627**
- control-cores | **4631**
- control-link-recovery | **4633**
- control-plane-bandwidth-percent (Node Slicing) | **4634**
- control-polarity | **4636**
- control-ports | **4637**
- control (REST API) | **4639**
- control-signal | **4640**
- control-source | **4641**
- control-vlan | **4643**
- control-word | **4644**
- control-word (Protocols Layer 2 Circuit Neighbor) | **4647**

- control-word (Protocols Layer 2 VPN) | **4648**
- controller | **4650**
- controller (OVSDB) | **4652**
- controller (Protocols OpenFlow) | **4654**
- convert-clock-class-to-quality-level | **4656**
- convert-sa-to-msdp | **4657**
- convey-loss-threshold | **4658**
- copy | **4660**
- copy-dont-fragment-bit (Services IPsec VPN) | **4661**
- copy-dont-fragment-bit (Services Set) | **4662**
- copy-plp-all | **4664**
- copy-tos-to-outer | **4665**
- copy-tos-to-outer-ip-header | **4666**
- copy-tos-to-outer-ip-header | **4668**
- copy-tos-to-outer-ip-header-transit | **4669**
- core-dump | **4670**
- core-facing | **4672**
- corouted-bidirectional | **4673**
- corouted-bidirectional-passive | **4674**
- corrupt-file | **4676**
- cos-rule-sets (Service Set Next Gen Services) | **4677**
- cos-rules (Service Set) | **4679**
- cos-rules (Service Set Next Gen Services) | **4680**
- cost | **4681**
- count (Protocols PIM) | **4684**
- count | **4685**
- count (HTTP Header Enrichment) | **4686**
- count (Security Custom Attack) | **4688**
- count (Security Policies) | **4689**
- counter | **4690**
- counters | **4691**
- cpu | **4693**

- cpu-factor (Dampen Options) | **4694**
- cpu-load-threshold | **4696**
- cpu (resource-manager) | **4697**
- cpu (TDF Gateway) | **4699**
- cpu-throttle (Next Gen Services) | **4701**
- craft-lockout | **4703**
- crb-proxy-mac | **4704**
- create-new-ucast-tunnel | **4707**
- create-req | **4708**
- credibility | **4710**
- credit (Interfaces) | **4712**
- critical-event (Process) | **4713**
- critical-event (System) | **4715**
- crl | **4717**
- crl (Adaptive Services Interface) | **4718**
- crl (Encryption Interface) | **4720**
- crl (Security) | **4721**
- cross-connect | **4724**
- cross-connect-ccm | **4726**
- cryptographic-address | **4728**
- csds | **4729**
- cspf-backoff-time | **4731**
- cts | **4732**
- cts-polarity | **4734**
- current-hop-limit | **4735**
- current-hop-limit (Dynamic Router Advertisement) | **4737**
- custom-attack | **4738**
- custom-attack-group | **4746**
- custom-attack-groups (Security IDP) | **4747**
- custom-attacks | **4748**
- custom-block-message | **4750**
- custom-ciphers | **4751**

- custom-message (Security Content Filtering) | 4758
- custom-message (Security Email Notify) | 4759
- custom-message (Security Fallback Block) | 4761
- custom-message (Security Fallback Non-Block) | 4762
- custom-message (Security Virus Detection) | 4763
- custom-message (Security Web Filtering) | 4764
- custom-message-subject (Security Email Notify) | 4766
- custom-message-subject (Security Fallback Block) | 4767
- custom-message-subject (Security Fallback Non-Block) | 4769
- custom-message-subject (Security Virus Detection) | 4770
- custom-objects | 4771
- custom-options | 4773
- custom-page | 4777
- custom-page-file | 4778
- custom-tag-string | 4780
- custom-url-category | 4781
- customer-vlans | 4783
- customization (SNMP) | 4784
- cut-through | 4786
- cvss-score | 4787
- cycle-time | 4788

c-timeout

IN THIS SECTION

- Syntax | 3153
- Hierarchy Level | 3153
- Description | 3153
- Options | 3153

Required Privilege Level | 3153
Release Information | 3154

Syntax

```
c-timeout value-in-minutes;
```

Hierarchy Level

```
[edit logical-systems name security alg sip],
[edit logical-systems name tenants name security alg sip],
[edit security alg sip],
[edit services alg sip],
[edit tenants name security alg sip]
```

Description

Configure the c-timeout parameter to Specify the time out interval for Session Initiation Protocol (SIP) transactions in minutes.

Options

value-in-minutes—Timeout interval.

- **Range:** 3 through 10 minutes
- **Default:** 3 minutes

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sip (Security) | 27977

ca-chain (REST API)

IN THIS SECTION

- Syntax | 3154
- Hierarchy Level | 3154
- Description | 3154
- Required Privilege Level | 3155
- Release Information | 3155

Syntax

```
ca-chain ca-chain;
```

Hierarchy Level

```
[edit system services rest https]
```

Description

Configure Certificate authority chain profile to perform https REST API request using mutual and server authentications.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 20.4R1.

ca-identity

IN THIS SECTION

- Syntax | 3155
- Hierarchy Level | 3155
- Description | 3155
- Options | 3156
- Required Privilege Level | 3156
- Release Information | 3156

Syntax

```
ca-identity ca-identity;
```

Hierarchy Level

```
[edit security pki ca-profile ca-profile-name]
```

Description

Specify the certificate authority (CA) identity to use in requesting digital certificates for J Series Services Routers and Adaptive Services (AS) and MultiServices PICs installed in M Series and T Series routers.

Options

ca-identity—The name of the CA identity. This name is typically the domain name of the CA.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Specifying the CA Profile Name

ca-name

IN THIS SECTION

- Syntax | 3156
- Hierarchy Level | 3157
- Description | 3157
- Options | 3157
- Required Privilege Level | 3157
- Release Information | 3157

Syntax

```
ca-name ca-identity;
```

Hierarchy Level

```
[edit security certificates certification-authority]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Specify the certificate authority (CA) identity to use in the certificate request.

Options

ca-identity—CA identity to use in the certificate request.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Specifying the Certificate Authority Name

ca-profile (Security PKI)

IN THIS SECTION

- Syntax | 3158
- Hierarchy Level | 3158
- Description | 3158

- Options | 3158
- Required Privilege Level | 3160
- Release Information | 3160

Syntax

```
ca-profile ca-profile-name {
    administrator {
        e-mail-address e-mail-address;
    }
    ca-identity ca-identity ;
    enrollment {
            retry number;
            retry-interval seconds;
            url url-name;
    }
    proxy-profile;
    revocation-check;
    routing-instance routing-instance-name ;
    source-address ip-address;
}
```

Hierarchy Level

```
[edit security pki]
```

Description

Configure certificate authority (CA) profile. The CA profile contains the name and URL of the CA or RA, as well as retry-timer settings.

Options

ca-profile-name Name of a trusted CA.

administrator *email-address*
Specify an administrator e-mail address to which the certificate request is sent. By default, there is no preset e-mail address.

ca-identity
Specify the certificate authority (CA) identity to use in requesting digital certificates. This name is typically the domain name of the CA.

enrollment
Specify the enrollment parameters for a certificate authority (CA).

retry *number*
Number of automated attempts for online enrollment to be retried in case enrollment response is pending.

- **Range:** 0 through 1080
- **Default:** 10

retry-interval *seconds*
Time interval between the enrollment retries.

- **Range:** 0 through 3600
- **Default:** 900 seconds

url *url-name*
Enrollment URL where the Simple Certificate Enrollment Protocol (SCEP) or CMPv2 request is sent to the certification authority (CA) as configured in this profile. With SCEP, you enroll CA certificates with the `request security pki ca-certificate enroll` command and specify the CA profile. There is no separate command to enroll CA certificates with CMPv2. The IP address in the enrollment URL can be an IPv4 or an IPv6 address.

proxy-profile
Use specified proxy server. If proxy profile is configured in CA profile, the device connects to the proxy host instead of the CA server while certificate enrollment, verification or revocation. The proxy host communicates with the CA server with the requests from the device, and then relay the response to the device.

Public key infrastructure (PKI) uses proxy profile configured at the system-level. The proxy profile being used in the CA profile must be configured at the `[edit services proxy]` hierarchy. There can be more than one proxy profile configured under `[edit services proxy]` hierarchy. Each CA profile is referred to the most one such proxy profile. You can configure host and port of the proxy profile at the `[edit system services proxy]` hierarchy.

revocation-check
Specify the method the device uses to verify the revocation status of digital certificates.

routing-instance
Specify the routing-instance to be used.

source-address Specifies a source IPv4 or IPv6 address to be used instead of the IP address of the egress interface for communications with external servers. External servers are used for certificate enrollment and reenrollment using Simple Certificate Enrollment Protocol (SCEP) or Certificate Management Protocol version 2 (CMPv2), downloading certificate revocation lists (CRLs) using HTTP or LDAP, or checking certificate revocation status with Online Certificate Status Protocol (OCSP). If this option is not specified then the IP address of the egress interface is used as the source address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5. Support for `ca-identity` option is added in Junos OS Release 11.1. Support for `ocsp` and `use-ocsp` options added in Junos OS Release 12.1X46-D20.

Support for `proxy-profile` option is added in Junos OS Release 18.2R1.

Support for `source-address` is introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Basic Elements of PKI in Junos OS

ca-profile (Services)

IN THIS SECTION

- Syntax | 3161
- Hierarchy Level | 3161
- Description | 3161
- Required Privilege Level | 3161
- Release Information | 3161

Syntax

```
ca-profile ca-profile;
```

Hierarchy Level

```
[edit services unified-access-control infranet-controller hostname]
```

Description

Specify the certificate authority (CA) of the certificate that the SRX Series Firewall should use in communications with an Infranet Enforcer. The SRX Series Firewall uses the CA to validate the IC Series UAC Appliance server certificate.

Use this statement if you have loaded certificates from multiple certificate authorities (CAs) onto your SRX Series Firewall and you need to configure the device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series appliance .

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

ca-type

IN THIS SECTION

- Syntax | 3162
- Hierarchy Level | 3162

- Description | 3162
- Default | 3162
- Options | 3162
- Required Privilege Level | 3163
- Release Information | 3163

Syntax

```
ca-type type {
    ca-value value;
}
```

Hierarchy Level

```
[edit protocols lldp-med interface (all | interface-name location civic-based)]
```

Description

For Link Layer Discovery Protocol–Media Endpoint Device (LLDP-MED), configure the location types and values that comprise the location information advertised from the device to the MED. This information is used during emergency calls to identify the location (civic or postal address) of the caller.

For further information about the types and values that can be used to comprise the location, refer to RFC 4776, *Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information*. A subset of those types, with sample values, is provided below.

Default

Disabled.

Options

ca-type *type* Specify category codes that together with values represent information about the civic or postal address of the caller's location. The address is divided into information types, with

each information type represented by a code. Some of the codes, with corresponding sample values for the ca-value option, are:

- 0—A code that specifies the language used to describe the location.
- 16—The leading-street direction, such as “N”.
- 17—A trailing street suffix, such as “SW”.
- 18—A street suffix or type, such as “Ave” or “Road”.
- 19—A house number, such as “6450”.
- 20—A house-number suffix, such as “A” or “1/2”.
- 21—A landmark, such as “Stanford University”.
- 22—Additional location information, such as “South Wing”.
- 23—The name and occupant of a location, such as “Carrillo's Holiday Market”.
- 24—A house-number suffix, such as “95684”.
- 25—A building structure, such as “East Library”.

ca-value *value* Configure location information (civic or postal address) that is indexed by the ca-type code. This value has a maximum size of 248 characters. See the description of the ca-type option for examples.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show lldp | **21942**

Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch

cable-length (Congestion Notification)

IN THIS SECTION

- Syntax | 3164
- Hierarchy Level | 3164
- Description | 3164
- Default | 3165
- Options | 3165
- Required Privilege Level | 3165
- Release Information | 3165

Syntax

```
cable-length cable-length-value;
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name input]
```

Description

Specify the length of the cable between the interface and its peer interface in meters. The system uses the cable length and the maximum receive unit (MRU) to calculate the amount of buffer headroom reserved to support priority-based flow control (PFC). The the shorter the cable length and lower the MRU, the less headroom buffer space is required for PFC.

NOTE: You can also set a maximum transmission unit (MTU) value (the largest packet size the interface sends) for interfaces by including the `mtu` statement at the `[edit interfaces interface-name]` hierarchy level.

Default

The default cable length value is 100 meters (approximately 328 feet).

Options

cable-length-value—Length of the cable in meters. (Generally from 1 to 300 meters, but there is no configuration restriction.)

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring CoS PFC (Congestion Notification Profiles)

Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows

cac (TDF Gateway)

IN THIS SECTION

- Syntax | 3166
- Hierarchy Level | 3166
- Description | 3166
- Required Privilege Level | 3166

Release Information | 3166

Syntax

```
cac {
    cpu cpu-pct;
    maximum-sessions max-sessions;
    maximum-sessions-trap-percentage max-sessions-pct;
    memory memory-pct;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name]
```

Description

Configure the call admissions control (CAC) parameters for the TDF gateway.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a TDF Gateway

cache

IN THIS SECTION

- Syntax | 3167
- Hierarchy Level | 3167
- Description | 3167
- Options | 3167
- Required Privilege Level | 3168
- Release Information | 3168

Syntax

```
cache {
    size value;
    timeout value;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated]
[edit security utm feature-profile web-filtering juniper-enhanced]
```

Description

Set the cache parameters for Surf-Control-Integrated Web filtering and Enhanced Web Filtering.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Surf-Control feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5 for surf-control integrated.

Statement introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Surf-Control feature is not supported from Junos OS Release 15.1X49-D10 onwards.

cache (Security Log)

IN THIS SECTION

- Syntax | 3169
- Hierarchy Level | 3169
- Description | 3169
- Options | 3169
- Required Privilege Level | 3169
- Release Information | 3170

Syntax

```
cache {
    exclude exlude-name {
        destination-address destination-address;
        destination-port destination-port;
        event-id event-id;
        failure;
        interface-name interface-name;
        policy-name policy-name;
        process process-name;
        protocol protocol;
        source-address source-address;
        source-port source-address;
        success;
        user-name user-name;
    }
    limit value;
}
```

Hierarchy Level

```
[edit security log]
```

Description

Cache security log events in the audit log buffer.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

syslog (System) | 28836

cache-size

IN THIS SECTION

- Syntax | 3170
- Hierarchy Level | 3170
- Description | 3170
- Options | 3171
- Required Privilege Level | 3171
- Release Information | 3171

Syntax

```
cache-size bytes;
```

Hierarchy Level

```
[edit security certificates]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the cache size for digital certificates.

Options

bytes—Cache size for digital certificates.

- **Range:** 64 through 4,294,967,295
- **Default:** 2 megabytes (MB)

> **NOTE**: We recommend that you limit your cache size to 4 MB.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Cache Size

cache-timeout-negative

IN THIS SECTION

- Syntax | 3172
- Hierarchy Level | 3172
- Description | 3172
- Options | 3172
- Required Privilege Level | 3172
- Release Information | 3172

Syntax

```
cache-timeout-negative seconds;
```

Hierarchy Level

```
[edit security certificates]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Configure a negative cache for digital certificates.

Options

seconds—Negative time to cache digital certificates, in seconds.

- **Range:** 10 through 4,294,967,295
- **Default:** 20

> **CAUTION**: Configuring a large negative cache value can lead to a denial-of-service attack.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Negative Cache

cacheable (Application Identification)

IN THIS SECTION

- Syntax | 3173
- Hierarchy Level | 3173
- Description | 3173
- Required Privilege Level | 3173
- Release Information | 3173

Syntax

```
cacheable;
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Description

Enable the application system cache (ASC), which saves the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service. The ASC is disabled by default.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

cak

IN THIS SECTION

- Syntax | 3174
- Hierarchy Level | 3174
- Description | 3174
- Default | 3175
- Options | 3175
- Required Privilege Level | 3175
- Release Information | 3175

Syntax

```
ckn hexadecimal-number;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name pre-shared-key],
[edit security macsec connectivity-association connectivity-association-name fallback-key]
```

Description

Specifies the connectivity association key (CAK) for a preshared key.

A preshared key includes a connectivity association key name (CKN) and a CAK. A preshared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security keys. The CKN is a 64-digit hexadecimal number and the CAK is a 32-digit hexadecimal number.

The MACsec Key Agreement (MKA) protocol is enabled once the preshared keys are successfully exchanged. The preshared key—the CKN and CAK—must match on both ends of a link.

To configure MACsec on the supported ports, you need to create the preshared key by configuring the connectivity association key name (CKN) and connectivity association key (CAK).

Default

No CAK exists, by default.

Options

hexadecimal-number

The key name, in hexadecimal format.

For AES-128 the CAK key length is 32 hexadecimal characters, and for AES-256 it is 64. If you enter a key name with fewer characters than the standard (32 for cipher-suite gcm-aes-128, and 64 for cipher-suite gcm-aes-256), Junos will automatically fill in the remaining characters with zeros.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced for SRX devices in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

cak (MX Series)

IN THIS SECTION

- Syntax | 3176
- Hierarchy Level | 3176
- Description | 3176
- Default | 3176
- Options | 3177
- Required Privilege Level | 3177
- Release Information | 3177

Syntax

```
cak hexadecimal-number;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name pre-shared-key]
```

Description

Specifies the connectivity association key (CAK) for a pre-shared key.

A pre-shared key includes a connectivity association key name (CKN) and a CAK. A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security keys. The MACsec Key Agreement (MKA) protocol is enabled once the pre-shared keys are successfully exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link

Default

No CAK exists, by default.

Options

hexadecimal-number The key name, in hexadecimal format.

For AES-128 the CAK key length is 32 hexadecimal characters, and for AES-256 it is 64. If you enter a key name with fewer characters than the standard (32 for cipher-suite gcm-aes-128, and 64 for cipher-suite gcm-aes-256), Junos will automatically fill in the remaining characters with zeros.

On MX10003 routers, to maximize the security, it is recommended to configure CAK of even length. If you configure the length of CAK to an odd value, you'll receive a warning message.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

calculation-weight

IN THIS SECTION

- Syntax | 3178
- Hierarchy Level | 3178
- Description | 3178
- Required Privilege Level | 3178
- Release Information | 3178

Syntax

```
calculation-weight {
    delay delay-value;
    delay-variation delay-variation-value;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
profile-name]
```

Description

Configure the calculation weight for delay and delay variation.

NOTE: This option is applicable only for two-way delay measurement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring an Iterator Profile

Configuring an Iterator Profile on a Switch (CLI Procedure)

delay | 5238

delay-variation | 5265

call-flood

IN THIS SECTION

- Syntax | 3179
- Hierarchy Level | 3179
- Description | 3179
- Options | 3179
- Required Privilege Level | 3180
- Release Information | 3180

Syntax

```
call-flood {
    threshold rate;
}
```

Hierarchy Level

```
[edit logical-systems name security alg sccp application-screen],
[edit logical-systems name tenants name security alg sccp application-screen],
[edit security alg sccp application-screen],
[edit services alg sccp application-screen],
[edit tenants name security alg sccp application-screen]
```

Description

Limit the number of calls per second allowed to Skinny Client Control Protocol (SCCP) client. Calls exceeding the threshold are dropped by the SCCP Application Layer Gateway (ALG).

Options

`threshold` *rate*—Number of calls per second per client.

- **Range:** 2 through 1000
- **Default:** 20

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sccp | 16341

call-rate-statistics

IN THIS SECTION

- Syntax | 3180
- Hierarchy Level | 3181
- Description | 3181
- Options | 3181
- Required Privilege Level | 3181
- Release Information | 3181

Syntax

```
call-rate-statistics {
    history records;
```

```
    interval minutes;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name],
[edit unified-edge gateways tdf gateway-name domains domain-name]
```

Description

Configure call rate statistics for a TDF gateway or a TDF domain.

Options

records Number of call-rate statistics records to save. When the number of call-rate records equals this value and a new record is received, the oldest record is replaced by the new record.

minutes Length of statistics collection interval.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Call-Rate Statistics Collection

callback

IN THIS SECTION

- Syntax | 3182
- Hierarchy Level | 3182
- Description | 3182
- Required Privilege Level | 3183
- Release Information | 3183

Syntax

```
callback;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options incoming-map],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options
incoming-map]
```

Description

On J Series Services Routers with interfaces configured for ISDN, configure the dialer to terminate the incoming call and call back the originator after the callback wait period. The default wait time is 5 seconds. To configure the wait time, include the `callback-wait-period` statement at the `[edit interfaces dl n unit logical-unit-number dialer-options]` hierarchy level.

NOTE: The `incoming-map` statement is mandatory for the router to accept any incoming ISDN calls.

If the `callback` statement is configured, you cannot use the `caller caller-id` statement at the `[edit interfaces dln unit logical-unit-number dialer-options]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

callback-wait-period | 3183

callback-wait-period

IN THIS SECTION

- Syntax | 3183
- Hierarchy Level | 3184
- Description | 3184
- Options | 3184
- Required Privilege Level | 3184
- Release Information | 3184

Syntax

```
callback-wait-period time;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with interfaces configured for ISDN with callback, specify the amount of time the dialer waits before calling back the caller. The default wait time is 5 seconds. The wait time is necessary because, when a call is rejected, the switch waits for up to 4 seconds on point-to-multipoint connections to ensure no other device accepts the call before sending the DISCONNECT message to the originator of the call. However, the default time of 5 seconds may not be sufficient for different switches or may not be needed on point-to-point connections.

To configure callback mode, include the `callback` statement at the [`edit interfaces dl`*n* `unit` *logical-unit-number* `dialer-options`] hierarchy level.

If the `callback` statement is configured, you cannot use the `caller` *caller-id* statement at the [`edit interfaces dl`*n* `unit` *logical-unit-number* `dialer-options`] hierarchy level.

Options

time—Time the dialer waits before calling back the caller.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

called-station-id

IN THIS SECTION

- Syntax | 3185
- Hierarchy Level | 3185
- Description | 3185
- Required Privilege Level | 3185
- Release Information | 3186

Syntax

```
called-station-id {
    equals value;
    matches value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from]
```

Description

Specify the RADIUS AVP called station ID for the incoming RADIUS request from the IP-based subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

called-station-id (OCS Partition)

IN THIS SECTION

- Syntax | 3186
- Hierarchy Level | 3187
- Description | 3187
- Options | 3187
- Required Privilege Level | 3187
- Release Information | 3187

Syntax

```
called-station-id station-name;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Include the Called-Station-Id AVP value in all CCR-GY packets.

Options

station-name Called station name defined for the Called-Station-Id AVP to include in all CCR-GY packets.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

caller

IN THIS SECTION

- Syntax | 3188
- Hierarchy Level | 3188
- Description | 3188
- Options | 3188
- Required Privilege Level | 3189
- Release Information | 3189

Syntax

```
caller (caller-id | accept-all);
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options incoming-map],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options
incoming-map]
```

Description

On J Series Services Routers with interfaces configured for ISDN, specify the dialer to accept a specified caller number or accept all incoming calls.

Options

caller-id—Incoming caller number. You can configure multiple caller IDs on a dialer. The caller ID of the incoming call is matched against all caller IDs configured on all dialers. The dialer matching the caller ID is looked at for further processing. Only a precise match is a valid match, For example, the configured caller ID 1-222-333-4444 or 222-333-4444 will match the incoming caller ID 1-222-333-4444.

If the incoming caller ID has fewer digits than the number configured, it is not a valid match. Duplicate caller IDs are not allowed on different dialers; however, for example, the numbers 1-408-532-1091, 408-532-1091, and 532-1091 can still be configured on different dialers.

Only one B-channel can map to one dialer. If one dialer is already mapped, any other call mapping to the same dialer is rejected (except in the case of a multilink dialer). If no dialer caller is configured on a dialer, that dialer will not accept any calls.

`accept-all`—Any incoming call in an associated interface is accepted.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

calling-number

IN THIS SECTION

- Syntax | 3189
- Hierarchy Level | 3190
- Description | 3190
- Options | 3190
- Required Privilege Level | 3190
- Release Information | 3190

Syntax

```
calling-number number;
```

Hierarchy Level

```
[edit interfaces br-pim/0/port isdn-options]
```

Description

On J Series Services Routers with ISDN interfaces, configure the calling number to include in outgoing calls.

Options

number—Calling number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring ISDN Physical Interface Properties

calling-number (L2TP Tunnel Switching)

IN THIS SECTION

- Syntax | 3191
- Hierarchy Level | 3191
- Description | 3191

- Options | 3191
- Required Privilege Level | 3191
- Release Information | 3192

Syntax

```
calling-number action;
```

Hierarchy Level

```
[edit access tunnel-switch-profile profile-name avp]
```

Description

Specify the action taken on the Calling Number AVP (22) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.

Options

action—One of the following actions:

- `drop`—Drop the AVP.
- `regenerate`—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.
- `relay`—Forward the AVP transparently as is and send it in the switched packet.
- **Default:** `relay`

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

calling-station-id

IN THIS SECTION

- Syntax | 3192
- Hierarchy Level | 3192
- Description | 3193
- Required Privilege Level | 3193
- Release Information | 3193

Syntax

```
calling-station-id {
    equals value;
    matches value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from]
```

Description

Specify the RADIUS AVP calling station ID for the incoming RADIUS request from the IP-based subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers*
- *Understanding Selection of Properties for an IP-Based TDF Subscriber*
- *Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber*
- *IP-Based Subscriber Setup Overview*

calling-station-id-format (Subscriber Management)

IN THIS SECTION

- Syntax | 3194
- Hierarchy Level | 3194
- Description | 3194
- Default | 3194

- Options | 3194
- Required Privilege Level | 3195
- Release Information | 3195

Syntax

```
calling-station-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    mac-address;
    nas-identifier;
    stacked-vlan;
    vlan;
}
```

Hierarchy Level

```
[edit access profile profile-name radius options]
```

Description

Specify the information that the router includes in the Calling-Station-ID (RADIUS IETF attribute 31) that is passed to the RADIUS server during authentication and accounting. You can include one or more optional values in any combination.

Default

The router displays the Calling-Station-ID set by the client.

Options

`agent-circuit-id`—Include the agent circuit identifier (ACI) string, which uniquely identifies the subscriber's access node and the digital subscriber line (DSL) on the access node. The ACI string is stored in either

the DHCP option 82 field of DHCP messages (for DHCP traffic), or in the DSL Forum Agent-Circuit-ID VSA [26-1] of PPPoE PADI and PADR control packets (for PPPoE traffic).

`agent-remote-id`—Include the agent remote identifier (ARI) string, which identifies the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request. The ARI string is stored in either the DHCP option 82 field of DHCP messages (for DHCP traffic), or in the DSL Forum Agent-Remote-ID VSA [26-2] of PPPoE PADI and PADR control packets (for PPPoE traffic).

`interface-description`—Include the interface description value.

`interface-text-description`—Include the interface text description.

`mac-address`—Include the MAC address of the source device for the subscriber.

`nas-identifier`—Include the NAS-identifier (RADIUS IETF attribute 32), which specifies the name of the NAS that orginated the authentication or accounting request.

`stacked-vlan`—Include the stacked VLAN tag value.

`vlan`—Include the VLAN tag value.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

`mac-address`, `interface-text-description`, `stacked-vlan`, and `vlan` options added in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring a Calling-Station-ID with Additional Options

candidate-uplink-port-policy (satellite-policies)

IN THIS SECTION

- Syntax | 3196
- Hierarchy Level | 3196
- Description | 3197
- Default | 3197
- Options | 3197
- Required Privilege Level | 3197
- Release Information | 3197

Syntax

```
candidate-uplink-port-policy policy-name{
    <holddown holddown-time>;
    <minimum-links number-of-links>;
    <uplink-port-group uplink-port-group-name>;
    term term-name {
        from {
            product-model model-name;
            <holddown holddown-time>;
            <minimum-links number-of-links>;
            <uplink-port-group uplink-port-group-name>;
        }
    }
}
```

Hierarchy Level

```
[edit policy-options satellite-policies]
```

Description

Configures a candidate uplink port profile, which contains uplink failure detection feature configuration that can be applied to satellite devices in a Junos Fusion.

Default

There is no configured candidate uplink port profile, by default.

Options

policy-name User-defined name for the policy.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

capabilities (Protocols RIFT)

IN THIS SECTION

- Syntax | 3198
- Hierarchy Level | 3198

- Description | 3198
- Options | 3198
- Required Privilege Level | 3198
- Release Information | 3199

Syntax

```
capabilities {
    (flood-reduction | no-flood-reduction);
}
```

Hierarchy Level

```
[edit protocols rift]
```

Description

Enable the RIFT capability of reducing flooding globally.

If this capability is enabled, the node refloods TIEs only if it is the flood leader on the incoming link.

If you do not specify any of the options under the `capabilities` statement, the default behavior is to enable flood reduction.

Options

flood-reduction (Optional) Enable flood reduction globally.

no-flood-reduction (Optional) Disable flood reduction globally.

- **Default:** Flood reduction is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

captive-portal

IN THIS SECTION

- Syntax | 3199
- Hierarchy Level | 3200
- Description | 3200
- Default | 3200
- Options | 3200
- Required Privilege Level | 3201
- Release Information | 3201

Syntax

```
captive-portal {
    authentication-profile-name authentication-profile-name
    custom-options {
        banner-message string;
        footer-bgcolor color;
        footer-message string;
        footer-text-color color;
        form-header-bgcolor color;
```

```
        form-header-message string;
        form-header-text-color color;
        form-reset-label label name;
        form-submit-label label name;
        header-bgcolor color;
        header-logo filename;
        header-message string;
        header-text-color color;
        post-authentication-url url-string;
    }
    interface (all | [interface-names]) {
        quiet-period seconds;
        retries number-of-retries;
        server-timeout seconds;
        session-expiry seconds;
        supplicant (multiple | single | single-secure);
        user-keepalive minutes;
    }
    secure-authentication (http | https);
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure captive portal to authenticate clients connected to the switch for access to the network.

Default

Captive portal is disabled.

Options

authentication-profile-name *access-profile-name*

Specify the name of the access profile to be used for captive portal authentication. You configure the access profile at the [edit access profile] hierarchy level. The access profile contains the RADIUS server IP address and other information used for authentication.

- **Default:** No access profile is specified.

secure-authentication (http | https)

Enable HTTP or HTTPS access on the captive portal interface.

- **Default:** http
- **Values:** Specify one of the following:
  - http—Enables HTTP access on the captive portal interface.
  - https—Enables HTTPS access on the captive portal interface. HTTPS is recommended.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`user-keepalive` introduced in Junos OS Release 16.1 for EX Series switches.

RELATED DOCUMENTATION

Example: Setting Up Captive Portal Authentication on an EX Series Switch

Designing a Captive Portal Authentication Login Page on Switches

Configuring Captive Portal Authentication (CLI Procedure)

Configuring an EX Series Switch to Use Junos Pulse Access Control Service for Network Access Control (CLI Procedure)

captive-portal-content-delivery-profile (Services)

IN THIS SECTION

- Syntax | 3202
- Hierarchy Level | 3202
- Description | 3202
- Options | 3203
- Required Privilege Level | 3203
- Release Information | 3203

Syntax

```
captive-portal-content-delivery-profile profile-name
    interface-service {
        service-interface name;
    }
    next-hop-service {
        inside-service-interface interface-name.unit-number;
        outside-service-interface interface-name.unit-number;
        outside-service-interface-type local;
    }
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

NOTE: Starting in Junos OS Release 17.2R1, you can configure converged services for MS-MPCs and MS-MICs. Starting in Junos OS Release 19.3R2, you can configure

converged services for SPC3s if you have enabled Next Gen Services on the MX Series router. At the `edit service-set` *`service set name`* `captive-portal-content-delivery-profile` *`profile-name`* `interface-service` hierarchy level, you can configure captive portal content delivery (CPCD) profiles for MS-MICs and MS-MPCs by including the `service-interface ms-`*`fpc/pic/port`* statement, and configure captive portal content delivery (CPCD) profiles for SPC3s by including the `service-interface vms-`*`fpc/pic/port`* statement.

Options

`service-interface` *`name`*—Name of the service device associated with the interface-wide service set.

Required Privilege Level

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 17.2.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

HTTP Redirect Service Overview

dynamic-profile (Dynamic PPPoE) | **6124**

captive-portal (Services UAC)

IN THIS SECTION

- Syntax | **3204**
- Hierarchy Level | **3204**
- Description | **3204**

- Options | 3204
- Required Privilege Level | 3204
- Release Information | 3205

Syntax

```
captive-portal redirect-policy-name{
    redirect-traffic (all | unauthenticated);
    redirect-url redirect-url;
}
```

Hierarchy Level

```
[edit services unified-access-control]
```

Description

Specify the preconfigured security policy for captive portal on the Junos OS Enforcer to enable the captive portal feature. The captive portal policy is configured as part of the UAC policy.

By configuring the captive portal feature, you can redirect traffic destined for protected resources to the IC Series device or to the URL you configure on the Junos OS Enforcer.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

captive-portal (Services UAC Policy)

IN THIS SECTION

- Syntax | 3205
- Hierarchy Level | 3205
- Description | 3205
- Required Privilege Level | 3206
- Release Information | 3206

Syntax

```
captive-portal captive-portal-policy-name;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
application-services uac-policy]
```

Description

Create the captive portal policy in the UAC security policy. You use the captive portal policy to configure the captive portal feature on the Junos OS Enforcer. By configuring the captive portal feature, you can redirect traffic destined for protected resources to the IC Series device or to the URL you configure on the Junos OS Enforcer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Security Policies Overview

capture-file (Security)

IN THIS SECTION

- Syntax | 3206
- Hierarchy Level | 3207
- Description | 3207
- Options | 3207
- Required Privilege Level | 3207
- Release Information | 3208

Syntax

```
capture-file {
    filename;
    files number;
    format pacp-format;
    size maximum-file-size;
```

```
    (world-readable | no-world-readable);
}
```

Hierarchy Level

```
[edit security datapath-debug]
```

Description

Sets packet capture for performing the datapath-debug action.

Options

- *filename*—Name of the file to receive the output of the packet capturing operation.
- `files`—Maximum number of capture files.

  If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

  Range: 1 through 10 files
- `format`—Describes the format of the capture file. The default format file is pcap. You can also set it as private (binary) format.
- `size`—Describes the size limit of the capture file.

  If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

  Range: 10 KB through 100 MB
- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

System Log Messages

capture-group

IN THIS SECTION

- Syntax | 3208
- Hierarchy Level | 3209
- Description | 3209
- Required Privilege Level | 3209
- Release Information | 3209

Syntax

```
capture-group client-name {
    content-destination identifier {
        address address;
        hard-limit bandwidth;
        hard-limit-target bandwidth;
        soft-limit bandwidth;
        soft-limit-clear bandwidth;
        ttl hops;
    }
    control-source identifier {
        allowed-destinations [ destinations ];
        minimum-priority value;
        no-syslog;
        notification-targets address port port-number;
```

```
        service-port port-number;
        shared-key value;
        source-addresses [ addresses ];
    }
    duplicates-dropped-periodicity seconds;
    input-packet-rate-threshold rate;
    interfaces interface-name;
    max-duplicates number;
    pic-memory-threshold percentage percentage;
}
```

Hierarchy Level

```
[edit services dynamic-flow-capture]
```

Description

Define the capture group values.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Capture Group

casb

IN THIS SECTION

- Syntax | 3210
- Hierarchy Level | 3211
- Description | 3211
- Options | 3211
- Required Privilege Level | 3212
- Release Information | 3212

Syntax

```
casb {
    instance instance-name {
     application application-name;
        domain domain-name;
     instance-id instnace-id;
     tag tag-value;
     type type;
    }
    casb-policy policy-name {
        rules rule-name {
            match {
                  application application-name {
                activity activity-name {
                [param-name param-value param-name param-value]
            }
            instance instance-name;
           }
              application-group application-group-name {
                   activity activity-name
               }

            }
            then {
            allow;
```

```
                deny;
                log-action;
            }
        }
        log-activity [ login upload download ];
        default-rule {
            deny;
            log-action;
            }
        }
    default-policy;
    traceoptions {
        file (file-name | files files | match match | no-world-readable | size size | world-
readable);
        flag ( all | flow | memory | pfeman);
        level;
        no-remote-trace;
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure Cloud Access Security Broker (CASB).

CASB provides visibility into the security of your cloud applications. To use CASB on your firewalls, you must configure CASB policies and apply CASB policy rules in a security policy.

Options

"instance " on page 8754*instance-name*	CASB instance name. Configure CASB instance with application, instance ID, domain, and type. Creating instances for SaaS applications ensures controlled access, data security, and differentiation between corporate and non-corporate usage.
"casb-policy" on page 3212*policy-name*	CASB policy name. Configure CASB policies to control specific actions on each cloud application.

default-policy ***policy-name***	The default CASB policy that defines the actions the device takes on a packet that does not match any user-defined policy. You must set up a default CASB policy for the unified policy configuration. This default policy applies to the session until a dynamic application match occurs. Once the final application match available for the security policy, the corresponding CASB policy will be applied. If no CASB policy is explicitly configured in the final firewall policy, the CASB service disengages for the session.
"traceoption" on page 29503	(Optional) Define tracing operations for CASB functionality.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

casb-policy

IN THIS SECTION

- Syntax | 3213
- Hierarchy Level | 3213
- Description | 3213
- Options | 3214
- Required Privilege Level | 3214
- Release Information | 3214

Syntax

```
casb-policy policy-name {
    rules rule-name {
        match {
              application application-name {
            actity activity-name {
            [param-name param-value param-name param-value]
            }
        instance instance-name;
          }
          application-group application-group-name {
               activity activity-name
             }
        }
        then {
        allow;
           deny;
           log-action;
        }
    }
    log-activity [ login upload download ];
    default-rule {
        deny;
        log-action;
        }
    }
```

Hierarchy Level

```
[edit security]
```

Description

Configure Cloud Access Security Broker (CASB) policy. CASB policy includes set of rules to match a section of traffic based on applications or instances and then specify actions to apply on the matched traffic and take active or passive preventive actions.

Options

"rules" on page 16087 *rule-name*	Configure rules for the CASB policy.
log-activity	(optional) Define activity logging for the CASB policy. Available options are login, download, and upload.
default-rule	The default CASB policy that defines the actions the device takes on a packet that does not match any other rule. This is mandatory statement.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

cascade-port

IN THIS SECTION

- Syntax | 3215
- Hierarchy Level | 3215
- Description | 3215
- Default | 3215
- Options | 3215
- Required Privilege Level | 3215
- Release Information | 3215

Syntax

```
cascade-port;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure the specified interface on the aggregation device in a Junos Fusion into a cascade port.

Additional configuration is required to configure cascade ports on a Junos Fusion. See *Configuring Junos Fusion Provider Edge* or Configuring or Expanding a Junos Fusion Enterprise.

Default

No interfaces are cascade ports, by default.

Options

interface-name	Specifies the name of the interface.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

> Configuring or Expanding a Junos Fusion Enterprise
>
> *Configuring Junos Fusion Provider Edge*

cascade-ports

IN THIS SECTION

- Syntax | 3216
- Hierarchy Level | 3216
- Description | 3216
- Default | 3217
- Options | 3217
- Required Privilege Level | 3217
- Release Information | 3217

Syntax

```
cascade-ports interface-name;
```

Hierarchy Level

```
[edit chassis satellite-management fpc slot-id]
```

Description

Associate a cascade port with an FPC slot ID number in a Junos Fusion.

The FPC slot ID of the satellite device is determined by the value entered as the FPC *slot-id*. For instance, if the **set chassis satellite-management fpc 105 cascade-ports xe-0/0/1** statement is used to configure interface xe-0/0/1 into a cascade port, the satellite device that connects to interface xe-0/0/1 has an FPC slot ID of 105 in the Junos Fusion.

A Junos Fusion provides two methods of assigning an FPC identifier: Unique ID-based FPC identification and connectivity-based FPC identification. Unique ID-based FPC identification maps an FPC slot ID to a satellite device's MAC address or serial number, while connectivity-based FPC identification maps an FPC slot ID to a cascade port. This statement is used to assign an FPC ID using connectivity-based FPC identification by mapping an FPC slot ID to a cascade port.

In a Junos Fusion, each satellite device must be mapped to an FPC identifier (FPC ID). The FPC ID is used for Junos Fusion configuration, monitoring, and maintenance. Interface names—which are identified using the *type-fpc* / *pic* / *port* format—use the FPC ID as the *fpc* variable when the satellite device is participating in a Junos Fusion. For instance, built-in port 2—a Gigabit Ethernet interface on a satellite device that is using 101 as it's FPC ID— uses ge-101/0/2 as its interface name.

For additional information on the role of FPC slot IDs in a Junos Fusion, see *Understanding Junos Fusion Provider Edge Components* or Understanding Junos Fusion Enterprise Components.

Default

No FPC slot IDs are associated with satellite devices, by default.

Options

interface-name	Specifies the name of the interface.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

cascade-ports (Junos Fusion Satellite Device Cluster)

IN THIS SECTION

- Syntax | 3218
- Hierarchy Level | 3218
- Description | 3218
- Default | 3218
- Options | 3219
- Required Privilege Level | 3219
- Release Information | 3219

Syntax

```
cascade-ports interface-name;
```

Hierarchy Level

```
[edit chassis satellite-management cluster cluster-name]
```

Description

Associate a cascade port with a satellite device cluster in a Junos Fusion.

This command defines which cascade ports are associated with a satellite device cluster only. An interface still needs to be converted into a cascade port before it performs cascade port functions. There are multiple ways to convert an interface on the aggregation device into a cascade port. See *Configuring or Expanding a Junos Fusion Enterprise*.

Default

Cascade ports are not associated with satellite device clusters, by default.

Options

interface-name Specifies the name of the interface on the aggregation device that is associated with the satellite device cluster.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

categories

IN THIS SECTION

- Syntax | 3220
- Hierarchy Level | 3220
- Description | 3220
- Default | 3220
- Options | 3220
- Required Privilege Level | 3221
- Release Information | 3221

Syntax

```
categories {
    category;
}
```

Hierarchy Level

```
[edit snmp trap-group group-name]
```

Description

Define the types of traps that are sent to the targets of the named trap group.

Default

If you omit the `categories` statement, all trap types are included in trap notifications, except the `timing-events` category. You must explicitly configure generation of timing traps by including the `timing-events` statement at the [`edit snmp trap-group name`] hierarchy level for the timing traps to be included in the trap notifications.

Options

Categories

`authentication`	Authentication failures.
`chassis`	Chassis or environment notifications.
`chassis-cluster`	(For SRX Series) Clustering notifications.
`configuration`	Configuration notifications.
`link`	Link up-down transitions.
`otn-alarms`	OTN alarm trap subcategories.
`remote-operations`	Remote operations.
`rmon-alarm`	RMON rising and falling alarms.

`routing`	Routing protocol notifications.
`services`	Services notifications.
`sonet-alarms`	SONET alarm trap subcategories.
`startup`	System warm and cold starts.
`vrrp-events`	VRRP notifications.
`timing-events`	Timing notifications.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring SNMP Trap Groups

category

IN THIS SECTION

- Syntax | 3222
- Hierarchy Level | 3222
- Description | 3222
- Options | 3222
- Required Privilege Level | 3222
- Release Information | 3223

Syntax

```
category category-name {
   inline-monitoring-instance inline-monitoring-instance;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pfe identifier exception-reporting]
```

Description

(Required) Configure the PFE exception category type and assign a particular inline-monitoring instance to it.

Options

category-name

Name of the category. Include more than one `category` statement in your configuration if you want reporting for more than one exception type (Junos OS only). You can specify one of the following types on each `category` statement:

- all—(Junos OS Evolved only) All exceptions
- firewall—(Junos OS only) Firewall exceptions
- forwarding-state—(Junos OS only) Forwarding-state-related exceptions
- layer2—(Junos OS only) Layer 2 exceptions
- layer3—(Junos OS only) Layer 3 exceptions
- packet-errors—(Junos OS only) Packet-format-error-related exceptions

inline-monitoring-instance *inline-monitoring-instance*

Assign a particular inline-monitoring instance to the exception reporting category.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 for MX Series routers.

Statement introduced in Junos Evolved OS Release 22.2R1 for PTX Series routers.

RELATED DOCUMENTATION

Juniper Resiliency Interface

category (Security Dynamic Attack Group)

IN THIS SECTION

- Syntax | 3223
- Hierarchy Level | 3223
- Description | 3224
- Options | 3224
- Required Privilege Level | 3224
- Release Information | 3224

Syntax

```
category {
    values [category-value];
}
```

Hierarchy Level

```
[edit security idp dynamic-attack-group dynamic-attack-group-name filters]
```

Description

Specify a category filter to add attack objects based on the category.

Options

`values`—Name of the category filter. You can configure multiple filters separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

category (Security Logging)

IN THIS SECTION

- Syntax | 3224
- Hierarchy Level | 3225
- Description | 3225
- Options | 3225
- Required Privilege Level | 3226
- Release Information | 3226

Syntax

```
category (all | content-security | fw-auth | screen | alg | nat | flow | sctp | gtp | ipsec |
idp | rtlog | pst-ds-lite | appqos | secintel | aamw | dnsf)
```

Hierarchy Level

```
[edit security log stream stream-name]
[edit logical-systems name security log stream stream-name]
[edit tenants tenant-name security log stream stream-name]
```

Description

Set the category of logging to `all` or `content-security`. Note that for the WELF format, the category must be set to `content-security`.

Options

- `all`—All events are logged. By default, all the events listed in the `category` parameter are logged.
- `content-security`—Only content security events are logged.
- `fw-auth`—Firewall authentication events are logged.
- `screen`—Screen events are logged.
- `alg`—Application Layer Gateway (ALG) events are logged.
- `nat`—Network Address Translation (NAT) events are logged.
- `flow`—Flow events are logged.
- `sctp`—Stream Control Transmission Protocol (SCTP) events are logged.
- `gtp`—GPRS Tunneling Protocol (GTP) events are logged.
- `ipsec`—IPsec events are logged.
- `idp`—Intrusion Detection and Prevention (IDP) events are logged.
- `rtlog`—RTLOG system log events are logged.
- `pst-ds-lite`—PST dual-stack lite (DS-Lite) events are logged.
- `appqos`—Application quality of service (AppQoS) events are logged.
- `secintel`—Juniper Networks Security Intelligence (SecIntel) events are logged.
- `aamw`—Advanced-anti-malware (AAMW) events are logged.
- `dnsf`—DNS events are logged.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. Statement modified in Junos OS Release 15.1X49-D40.

The [edit `logical-systems` *name* security log stream *stream-name*] hierarchy level introduced in Junos OS Release 18.2R1.

The [edit `tenants` *tenant-name* security log stream *stream-name*] hierarchy levels introduced in Junos OS Release 18.3R1.

The dnsf category is introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

- Application Security User Guide for Security Devices
- Logical Systems and Tenant Systems User Guide for Security Devices

category (Security Web Filtering)

IN THIS SECTION

- Syntax | 3227
- Hierarchy Level | 3227
- Description | 3227
- Options | 3235
- Required Privilege Level | 3236
- Release Information | 3236

Syntax

```
category name{
    action  (block | log-and-permit | permit | quarantine);
    custom-message message-name;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated profile profile-name]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name]
```

Description

Select a custom URL category list you created (custom objects) for filtering against. The `custom-message` configuration option is used to notify the users when the URL is blocked or quarantined for each EWF category. You can customize the message with options such as user message or redirect URL. User messages indicate that website access has been blocked by an organization's access policy. Redirect URLs redirect a blocked or quarantined URL to any user-defined URL. Table 11 on page 3227 shows the list of categories predefined by Websense.

Starting with Junos OS Release 17.4R1, you can download and dynamically load new Enhanced Web Filtering (EWF) categories. The downloading and dynamic loading of the new EWF categories do not require a software upgrade. Websense occasionally releases new EWF categories. EWF classifies websites into categories according to host, URL, or IP address and performs filtering based on the categories.

NOTE: Existing configurations are not affected by the new categories but can be modified to make use of the new categories.

Table 11: List of Categories Predefined by Websense

Category ID	Category Name	Parent ID
1	Adult Material	0

Table 11: List of Categories Predefined by Websense *(Continued)*

Category ID	Category Name	Parent ID
2	Business and Economy	0
3	Education	0
4	Government	0
5	News and Media	0
6	Religion	0
7	Society and Lifestyles	0
8	Special Events	0
9	Information Technology	0
10	Abortion	0
11	Advocacy Groups	0
12	Entertainment	0
13	Gambling	0
14	Games	0
15	Illegal or Questionable	0
16	Job Search	0

Table 11: List of Categories Predefined by Websense *(Continued)*

Category ID	Category Name	Parent ID
17	Shopping	0
18	Sports	0
19	Tasteless	0
20	Travel	0
21	Vehicles	0
22	Violence	0
23	Weapons	0
24	Drugs	0
25	Militancy and Extremist	0
26	Intolerance	0
27	Health	0
28	Website Translation	9
29	Advertisements	110
64	User-Defined	0
65	Nudity	1

Table 11: List of Categories Predefined by Websense *(Continued)*

Category ID	Category Name	Parent ID
66	Adult Content	1
67	Sex	1
68	Financial Data and Services	2
69	Cultural Institutions	3
70	Media File Download	12
72	Military	4
73	Political Organizations	4
74	General Email	91
75	Proxy Avoidance	9
76	Search Engines and Portals	9
78	Web Hosting	9
79	Web Chat	91
80	Hacking	9
81	Alternative Journals	5
82	Non-Traditional Religions	6

Table 11: List of Categories Predefined by Websense *(Continued)*

Category ID	Category Name	Parent ID
83	Traditional Religions	6
84	Restaurants and Dining	7
85	Gay or Lesbian or Bisexual Interest	7
86	Personals and Dating	7
87	Alcohol and Tobacco	7
88	Prescribed Medications	24
89	Nutrition	24
90	Abused Drugs	24
91	Internet Communication	0
92	Pro-Choice	10
93	Pro-Life	10
94	Sex Education	1
95	Lingerie and Swimsuit	1
96	Online Brokerage and Trading	110
97	Educational Institutions	3

Table 11: List of Categories Predefined by Websense *(Continued)*

Category ID	Category Name	Parent ID
98	Instant Messaging	110
99	Application and Software Download	110
100	Pay-to-Surf	110
101	Internet Auctions	17
102	Real Estate	17
103	Hobbies	7
107	Sport Hunting and Gun Clubs	18
108	Internet Telephony	116
109	Streaming Media	116
110	Productivity	0
111	Marijuana	24
112	Message Boards and Forums	110
113	Personal Network Storage and Backup	116
114	Internet Radio and TV	116
115	Peer-to-Peer File Sharing	116

Table 11: List of Categories Predefined by Websense *(Continued)*

Category ID	Category Name	Parent ID
116	Bandwidth	0
117	Social Networking and Personal Sites	7
118	Educational Materials	3
121	Reference Materials	3
122	Social Organizations	0
123	Service and Philanthropic Organizations	122
124	Social and Affiliation Organizations	122
125	Professional and Worker Organizations	122
126	Security	0
128	Malicious Web Sites	126
138	Computer Security	9
146	Miscellaneous	0
147	Web Infrastructure	146
148	Web Images	146
149	Private IP Addresses	146

Table 11: List of Categories Predefined by Websense *(Continued)*

Category ID	Category Name	Parent ID
150	Content Delivery Networks	146
151	Dynamic Content	146
152	Network Errors	146
153	Uncategorized	146
154	Spyware	126
156	File Download Servers	146
164	Phishing and Other Frauds	126
166	Keyloggers	126
167	Potentially Unwanted Software	126
172	Bot Networks	126
191	Extended Protection	0
192	Elevated Exposure	191
193	Emerging Exploits	191
194	Suspicious Content	191
195	Organizational Email	91

Table 11: List of Categories Predefined by Websense *(Continued)*

Category ID	Category Name	Parent ID
196	Text and Media Messaging	91
200	Web and Email Spam	9
220	Compromised Websites	0
221	Newly Registered Websites	0
222	Collaboration Office	0
223	Office Mail	222
224	Office Drive	222
225	Office Documents	222
226	Office Apps	222
227	Web Analytics	9
228	Web and Email Marketing	9
1529	Classifieds Posting	0
1530	Blog Posting	0
1531	Blog Commenting	0

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4 .

Starting with Junos OS Release 15.1X49-D10, the SurfControl integrated feature is no longer supported. For previous releases, statement introduced in Junos OS Release 9.5. The `custom-message` option introduced in Junos OS Release 15.1X49-D110.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Redirect Web Filtering

category (System Logging)

IN THIS SECTION

- Syntax | 3236
- Hierarchy Level | 3237
- Description | 3237
- Options | 3237
- Required Privilege Level | 3238
- Release Information | 3238

Syntax

```
category category, category....category;
```

Hierarchy Level

```
[edit services service-set service-set-name syslog stream]
```

Description

Specify the categories for which you want to collect logs.

Options

all	All events are logged
content-security	Content security events are logged
fw-auth	Fw-auth events are logged
screen	Screen events are logged
alg	ALG events are logged
nat	NAT events are logged
flow	Flow events are logged
sctp	Sctp events are logged
gtp	Gtp events are logged
ipsec	Ipsec events are logged
idp	Idp events are logged
rtlog	Rtlog events are logged
pst-ds-lite	Pst-ds-lite events are logged
appqos	Appqos events are logged
secintel	Secintel events are logged
aamw	AAMW events are logged
sfw	Stateful Firewall events are logged

session	Session open and close events are logged
session-open	Session open events are logged
session-close	Session close events are logged
urlf	DNS request filtering events are logged
ha	Stateful High-Availability open and close events are logged
ha-open	Stateful High-Availability open events are logged
ha-close	Stateful High-Availability close events are logged
pcp	PCP logs

Required Privilege Level

system—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

cbr

IN THIS SECTION

- Syntax | 3239
- Hierarchy Level | 3239
- Description | 3239
- Default | 3239
- Options | 3239
- Required Privilege Level | 3240
- Release Information | 3240

Syntax

```
cbr rate;
```

Hierarchy Level

```
[edit interfaces at-fpc/pic/port atm-options vpi vpi-identifier shaping],
[edit interfaces at-fpc/pic/port unit logical-unit-number address address family family
multipoint-destination address shaping],
[edit interfaces at-fpc/pic/port unit logical-unit-number shaping],
[edit logical-systems logical-system-name interfaces at-fpc/pic/port unit logical-unit-number
address address family family multipoint-destination address shaping],
[edit logical-systems logical-system-name interfaces at-fpc/pic/port unit logical-unit-number
shaping]
```

Description

For ATM encapsulation only, define a constant bit rate bandwidth utilization in the traffic-shaping profile.

Default

Unspecified bit rate (UBR); that is, bandwidth utilization is unlimited.

Options

rate Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation `c`; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps.

For ATM1 and ATM2 OC3 interfaces, the maximum available rate is 100 percent of *line-rate*, or 135,600,000 bps. For ATM1 OC12 interfaces, the maximum available rate is 50 percent of *line-rate*, or 271,263,396 bps. For ATM2 IQ interfaces, the maximum available rate is 542,526,792 bps.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Defining the ATM Traffic-Shaping Profile Overview

rtvbr

shaping | **16925**

vbr

cbr rate

IN THIS SECTION

- Syntax | **3240**
- Hierarchy Level | **3241**
- Description | **3241**
- Options | **3241**
- Required Privilege Level | **3241**
- Release Information | **3241**

Syntax

```
cbr rate;
```

Hierarchy Level

```
[edit interfaces interface-name atm-options vpi vpi-identifier shaping]
```

Description

For ATM encapsulation only, define a constant bit rate bandwidth utilization in the traffic-shaping profile.

Options

- CBR Value–Constant bandwidth utilization (range: 33,000 through 1,199,920)
- CDVT–Cell delay variation tolerance in microseconds (range: 1 through 9999)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Release 9.5 of Junos OS.

ccc

IN THIS SECTION

- Syntax | 3242
- Hierarchy Level | 3242
- Description | 3242
- Default | 3242
- Required Privilege Level | 3242
- Release Information | 3242

Syntax

```
ccc;
```

Hierarchy Level

```
[edit interfaces ge-fpc/slot/port unit logical-unit-number family]
```

Description

Configure the logical interface as a circuit cross-connect (CCC).

> **NOTE**: On QFX10002 switches, circuit cross-connects are not supported on aggregated Ethernet interfaces.

Default

You must configure a logical interface to be able to use the physical device.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

ccc (Routing Policy Condition)

IN THIS SECTION

Syntax | 3243

- Hierarchy Level | 3243
- Description | 3243
- Options | 3243
- Required Privilege Level | 3243
- Release Information | 3244

Syntax

```
ccc {
    interface-name;
    standby;
    peer-unit unit-number;
    table table-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options condition if-route-exists address-
family],
[edit policy-options condition if-route-exists address-family]
```

Description

Specify that the route must correspond to a CCC prefix to be considered a match.

Options

interface-name—Interface used to establish the CCC route.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring Pseudowire Redundancy in a Mobile Backhaul Scenario

cdvt

IN THIS SECTION

- Syntax | 3244
- Hierarchy Level | 3244
- Description | 3245
- Options | 3245
- Required Privilege Level | 3245
- Release Information | 3245

Syntax

```
cdvt rate;
```

Hierarchy Level

```
[edit firewall atm-policer atm-policer-name]
```

Description

(MX Series routers) Define the Cell Delay/Variation Tolerance (CDVT) rate on a clear-channel multirate circuit emulation MIC.

Options

rate—CDVT rate in microseconds.

- **Range:** 1 through 1,800,000,000

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

show class-of-service traffic-control-profile | 19145

cell-mode (Dynamic Traffic Shaping)

IN THIS SECTION

- Syntax | 3246
- Hierarchy Level | 3246
- Description | 3246
- Options | 3246
- Required Privilege Level | 3247
- Release Information | 3247

Syntax

```
cell-mode (bytes bytes | $junos-cos-byte-adjust | cell-mode-bytes cell-mode-bytes |$junos-cos-
byte-adjust-cell);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name overhead-
accounting],
[edit class-of-service traffic-control-profiles profile-name overhead-accounting],
```

Description

Configure the mode to shape downstream ATM traffic as cells.

Options

bytes	Byte adjustment value for the `cell-mode` or `frame-mode` shaping options.
$junos-cos-byte-adjust	Predefined variable for byte adjustment that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.
cell-mode-bytes *cell-mode-bytes*	Shaping is based on the number of bytes in cells, and accounts for the ATM cell encapsulation and padding overhead. The resulting traffic stream conforms to the policing rates configured in downstream ATM switches, reducing the number of packet drops in the Ethernet network.
$junos-cos-byte-adjust-cell	Predefined variable for the cell mode shaping. This variable can not be used when the `overhead-accounting bytes` *bytes* option is configured.

BEST PRACTICE: We recommend using the `cell-mode-bytes` *cell-mode-bytes* option rather than the `bytes` option.

- **Range:** –120 through 124 bytes

> **NOTE**: If you specify a value for the `bytes` *`bytes`* option, you cannot specify a value for either the `cell-mode-bytes` option.

> **NOTE**: Cell mode is supported only on logical interfaces and interface sets; it is not supported on physical interfaces (ifd or ifd-remaining).

- **Default:** The default is `frame-mode`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Variable *`$junos-cos-byte-adjust-cell`* introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

CoS Adjustment Control Profiles Overview

Configuring CoS Adjustment Control Profiles

adjustment-control-profiles | **1961**

Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates

Bandwidth Management for Downstream Traffic in Edge Networks Overview

egress-shaping-overhead

bytes (Dynamic Traffic Shaping) | **3121**

frame-mode (Dynamic Traffic Shaping) | **7552**

cell-overhead

IN THIS SECTION

- Syntax | 3248
- Hierarchy Level | 3248
- Description | 3248
- Required Privilege Level | 3248
- Release Information | 3248

Syntax

```
cell-overhead;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp],
[edit access profile profile-name client client-name ppp]
```

Description

Configure the session to use Asynchronous Transfer Mode (ATM)-aware egress shaping on the IQ2 PIC.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

cellular-options

IN THIS SECTION

- Syntax | 3249
- Hierarchy Level | 3250
- Description | 3250
- Options | 3250
- Required Privilege Level | 3250
- Release Information | 3250

Syntax

```
cellular-options {
    roaming-mode (home only | automatic)
    gsm-options {
        select-profile profile-name;
        profiles {
            profile-name {
                sip-user-id simple-ip-user-id;
                sip-password simple-ip-password;
                access-point-name apn;
                authentication-method (pap | chap | none);
            }
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure options for connecting a 3G wireless modem interface to a cellular network.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

centralized

IN THIS SECTION

- Syntax | 3251
- Hierarchy Level | 3251
- Description | 3251
- Default | 3251
- Required Privilege Level | 3251
- Release Information | 3251

Syntax

```
centralized;
```

Hierarchy Level

```
[edit protocols lacp ppm]
```

Description

Disable distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all PPM processing for LACP packets on the Routing Engine.

This statement disables distributed PPM processing for only LACP packets. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the `no-delegate-processing` statement in the `[edit routing-options ppm]` hierarchy.

> **BEST PRACTICE**: We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

Default

Distributed PPM processing is enabled for all packets that use PPM.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

inline | 8666

Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)

Configuring Aggregated Ethernet LACP (CLI Procedure)

Configuring Distributed Periodic Packet Management

Configuring Link Aggregation

cert-file (Scripts)

IN THIS SECTION

- Syntax | 3252
- Hierarchy Level | 3252
- Description | 3253
- Default | 3253
- Options | 3253
- Required Privilege Level | 3253
- Release Information | 3253

Syntax

```
cert-file path;
```

Hierarchy Level

```
[edit event-options event-script file filename],
[edit system extensions extension-service application file filename],
[edit system scripts commit file filename],
[edit system scripts op file filename],
[edit system scripts snmp file filename]
```

Description

Specify the certificate file (Root CA or self-signed) that the device uses to validate the server's certificate when you refresh a local script from an HTTPS server.

When you issue the `set refresh` or `set refresh-from` configuration mode command to refresh a single script or all scripts of a given type from an HTTPS server, the device uses the certificate configured for each script to validate the server's certificate and ensure the server is authentic. If certificate validation is successful, the device refreshes the script.

> **NOTE**: We do not recommend refreshing scripts from an HTTPS server when the device is unable to authenticate the server because malicious users could compromise the integrity of a script.

Default

If you do not configure the `cert-file` statement for a script and you attempt to refresh the script from an HTTPS server, the device issues an error that `cert-file` is mandatory for HTTPS communication.

Options

path Absolute path to the certificate (Root CA or self-signed) in Privacy-Enhanced Mail (PEM) format that is used to validate the certificate of the server from which you refresh the script.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Configure and Use a Master Source Location for a Script

certificate

IN THIS SECTION

- Syntax | 3254
- Hierarchy Level | 3254
- Description | 3254
- Options | 3255
- Required Privilege Level | 3255
- Release Information | 3256

Syntax

```
certificate {
    local-certificate certificate-id;
    peer-certificate-type (pkcs7 | x509-signature);
    policy-oids oid;
    trusted-ca {
        ca-profile ca-profile-name;
        trusted-ca-group trusted-ca-group-name;
    }
}
```

Hierarchy Level

```
[edit security ike policy policy-name]
```

Description

Specify usage of a digital certificate to authenticate the virtual private network (VPN) initiator and recipient.

Options

`local-certificate` *`certificate-id`* —Specify a particular certificate when the local device has multiple loaded certificates. The device deletes existing IKE and IPsec SAs when you update the `local-certificate` configuration in the IKE policy. Starting in Junos OS Release 19.1R1, a commit check is added to prevent user from adding `.`, `/`, `%`, and space in a certificate identifier while generating a local or remote certificates or a key pair.

`peer-certificate-type`—Specify a preferred type of certificate (PKCS7 or X509).

- `pkcs7`—Public-Key Cryptography Standard #7.
- `x509-signature`—X509 is an ITU-T standard for public key infrastructure. This is the default value.

`policy-oids` *`oid`*—Configure policy object identifiers (OIDs). This configuration is optional. Policy OID contained in a peer's certificate or certificate chain. Up to five policy OIDs can be configured. Each OID can be up to 63 bytes long. You must ensure that at least one of the configured policy OIDs is included in a peer's certificate or certificate chain. Note that the **policy-oids** field in a peer's certificate is optional. If you configure policy OIDs in an IKE policy and the peer's certificate chain does not contain any policy OIDs, certificate validation for the peer fails.

`trusted-ca`—Specify a name for the trusted CA group. A minimum of one CA profile is mandatory to create a trusted CA group and a maximum of 20 CAs are allowed in one trusted CA group. Any CA from a particular group can validate the certificate for that particular entity. Specify the preferred certificate authority (CA) to use when requesting a certificate from the peer. You can associate an IKE policy to a single trusted CA profile or a trusted CA group. During certificate validation the IKE policy will limit itself to the configured group of CAs while establishing a secure connection. Any certificate issued other than the single trusted CA or the trusted CA group are not validated.

- `ca-profile` *`ca-profile-name`*—Specify a name for the CA profiles. A Certificate Authority (CA) is an entity that issues digital certificates which helps to establish secure connection between peers through certificate validation.
- `trusted-ca-group` *`trusted-ca-group-name`*—Specify a name for the trusted CA group. A minimum of one CA profile is mandatory to create a trusted CA group and a maximum of 20 CAs are allowed in one trusted CA group. Any CA from a particular group can validate the certificate for that particular topology.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. `policy-oids` option added in Junos OS Release 12.3X48-D10. Support for `trusted-ca` option added in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IPsec Overview

Understanding Certificate Authority Profiles

certificate-authority (REST API)

IN THIS SECTION

- Syntax | 3256
- Hierarchy Level | 3256
- Description | 3257
- Options | 3257
- Required Privilege Level | 3257
- Release Information | 3257

Syntax

```
certificate-authority certificate-authority-profile-name;
```

Hierarchy Level

```
[edit system services rest https mutual-authentication]
```

Description

Set the server certificate authority profile when configuring mutual authentication.

Options

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

certificate-id

IN THIS SECTION

- Syntax | 3257
- Hierarchy Level | 3258
- Description | 3258
- Options | 3258
- Required Privilege Level | 3258
- Release Information | 3258

Syntax

```
certificate-id {
    ca-profile ca-profile-name;
    challenge-password password;
    re-enroll-trigger-time-percentage percentage;
```

```
    re-generate-keypair;
    validity-period days;
}
```

Hierarchy Level

```
[edit security auto-re-enrollment]
```

Description

Specify a router certificate for auto-reenrollment. The ID is the same as that used to get the end entity's certificate from the issuing certificate authority.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Auto-Reenrollment of a Router Certificate

auto-re-enrollment | 2727

certificate (Juniper Secure Connect)

IN THIS SECTION

- Syntax | 3259
- Hierarchy Level | 3259
- Description | 3259
- Options | 3259
- Required Privilege Level | 3260
- Release Information | 3260

Syntax

```
certificate {
    no-expiry-warning;
    no-pin-request-per-connection;
    warn-before-expiry days;
}
```

Hierarchy Level

```
[edit security remote-access client-config]
```

Description

Define certificate identifier parameters for Juniper Secure Connect.

Options

no-expiry-warning	Disable certificate expiry warning.
no-pin-request-per-connection	Disable certificate pin request per connection.
warn-before-expiry	Enable certificate expiration warning in days before the expiry date.

- **Default:** 60 days
- **Range:** 1 through 90

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Juniper Secure Connect Administrator Guide

certificate-verification

IN THIS SECTION

- Syntax | 3260
- Hierarchy Level | 3261
- Description | 3261
- Options | 3261
- Required Privilege Level | 3261
- Release Information | 3261

Syntax

```
certificate-verification [ optional | required | warning ]
```

Hierarchy Level

```
[edit services unified-access-control]
```

Description

This option determines whether server certificate verification is required when initiating a connection between an SRX Series Firewall and a Junos Pulse Access Control Service in a UAC configuration. If no CA profile contains the certificate authority (CA) that signed the configured server certificate for the Access Control Service, this option determines whether the commit check should fail, a warning should be displayed, or the connection should be made without any warning.

> **NOTE**: For strict security, this option should be reset to `required`, and the proper CA certificate should be specified in the CA profile.

Options

- `optional`—Certificate verification is not required. If the CA certificate is not specified in the `ca-profile` option, the commit check passes and no warning is issued.
- `required`—Certificate verification is required. If the CA certificate is not specified in the `ca-profile` option, an error message is displayed, and the commit check fails. Use this option to ensure strict security.
- **Default:** `warning`—Certificate verification is not required. A warning message is displayed during commit check if the CA certificate is not specified in the `ca-profile` option.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Communications Between the Junos OS Enforcer and the IC Series UAC Appliance

Understanding User Role Firewalls

certificates

IN THIS SECTION

- Syntax | 3262
- Hierarchy Level | 3263
- Description | 3263
- Required Privilege Level | 3263
- Release Information | 3263

Syntax

```
certificates {
    cache-size bytes;
    cache-timeout-negative seconds;
    certification-authority ca-profile-name {
        ca-name ca-identity;
        crl file-name;
        encoding (binary | pem);
        enrollment-url url-name;
        file certificate-filename;
        ldap-url url-name;
    }
    enrollment-retry attempts;
    local certificate-name {
        certificate-key-string;
        load-key-file URL filename;
    }
    maximum-certificates number;
```

```
    path-length certificate-path-length;
}
```

Hierarchy Level

```
[edit security]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the digital certificates for IPsec.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Digital Certificates for an ES PIC

certification-authority

IN THIS SECTION

- Syntax | 3264
- Hierarchy Level | 3264

- Description | 3264
- Options | 3264
- Required Privilege Level | 3265
- Release Information | 3265

Syntax

```
certification-authority ca-profile-name {
    ca-name ca-identity;
    crl file-name;
    encoding (binary | pem);
    enrollment-url url-name;
    file certificate-filename;
    ldap-url url-name;
}
```

Hierarchy Level

```
[edit security certificates]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Configure a certificate authority profile name.

Configure certification authority (CA) for X.509 certificate.

Options

- *profile-name*—Name of this CA configuration.
- `ca-name` *`name`*—Name of the CA.
- `crl` *`filename`*—Certificate revocation list (CRL) filename.
- `encoding`—Certificate encoding, either `binary` or `pem` (privacy-enhanced mail).

- `enrollment-url` *`url`*—Enrollment URL.
- `file` *`filename`*—Certificate filename.
- `ldap-url` *`url`*—Lightweight Directory Access Protocol (LDAP) URL.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Certificate Authority Properties for an ES PIC

Configuring the Certificate Authority Properties for an ES PIC

ces-psn-channel (tdm-options)

IN THIS SECTION

- Syntax | 3266
- Hierarchy Level | 3266
- Description | 3266
- Options | 3266
- Required Privilege Level | 3266
- Release Information | 3267

Syntax

```
ces-psn-channel {
    dmac-address address;
    mode mode;
    vlan-id-1 vlanid;
    vlan-id-2 vlanid;
}
```

Hierarchy Level

```
[edit interfaces interface-name tdm-options]
```

Description

Configure the CES (Circuit Emulation Service) PSN (Packet Switched Networks) Channel parameters. Use the parameters to specify the destination MAC address, encapsulation type, and vlan tagging.

Options

dmac-address *address* Destination MAC address to be paired with the smart SFP.

mode *mode* Encapsulation mode for the TDM traffic for further network processing. Possible values are: MEF8 and MPLS. Default mode is MEF8 for E1, T1, and DS3 smart SFPs. MPLS encapsulation is not supported for STM1, STM4, and STM16 smart SFPs.

vlan-id-1 *vlanid* A valid VLAN identifier for single VLAN tagging. Possible values: 0 through 4094.

vlan-id-2 *vlanid* A valid outer VLAN identifier for dual VLAN tagging. Possible values: 0 through 4094. Dual VLAN tagging is not supported on STM1, STM4, and STM16 smart SFPs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4.

RELATED DOCUMENTATION

show interfaces smart-sfp-statistics | 21372

show interfaces smart-sfp-defects | 21359

cfeb

IN THIS SECTION

- Syntax | 3267
- Hierarchy Level | 3267
- Description | 3268
- Default | 3268
- Options | 3268
- Required Privilege Level | 3268
- Release Information | 3268

Syntax

```
cfeb slot-number (always | preferred);
```

Hierarchy Level

```
[edit chassis redundancy]
```

Description

On M10i routers only, configure which Compact Forwarding Engine Board (CFEB) is the primary and which is the backup.

Default

By default, the CFEB in slot 0 is the primary and the CFEB in slot 1 is the backup.

Options

slot-number—Specify which slot is the primary and which is the backup.

always—Define this CFEB as the sole device.

preferred—Define this CFEB as the preferred device of at least two.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring CFEB Redundancy on the M10i Router

cflowd (Discard Accounting)

IN THIS SECTION

- Syntax | 3269
- Hierarchy Level | 3269

- Description | 3269
- Options | 3270
- Required Privilege Level | 3270
- Release Information | 3270

Syntax

```
cflowd hostname {
    aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
            caida-compliant;
        }
        source-prefix;
    }
    autonomous-system-type (origin | peer);
    port port-number;
    version format;
}
```

Hierarchy Level

```
[edit forwarding-options accounting name output]
```

Description

Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect.

You can configure up to one version 5 and one version 8 flow format at the `[edit forwarding-options accounting name output]` hierarchy level.

Options

hostname—The IP address or identifier of the host system (the workstation running the cflowd utility).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Directing Traffic Sampling Output to a Server Running the cflowd Application

Enabling Flow Aggregation

cflowd (Flow Monitoring)

IN THIS SECTION

- Syntax | 3271
- Hierarchy Level | 3271
- Description | 3271
- Options | 3271
- Required Privilege Level | 3271
- Release Information | 3271

Syntax

```
cflowd hostname {
    port port-number;
}
```

Hierarchy Level

```
[edit forwarding-options monitoring name inet output]
```

Description

Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect.

You can configure up to eight version 5 flow formats at the `[edit forwarding-options monitoring name output]` hierarchy level. Version 8 flow formats are not supported for flow-monitoring applications.

Options

hostname—IP address or identifier of the host system (the workstation running the cflowd utility).

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Flow Aggregation

Directing Traffic Sampling Output to a Server Running the cflowd Application

cfp-to-et

IN THIS SECTION

- Syntax | 3272
- Hierarchy Level | 3272
- Description | 3272
- Required Privilege Level | 3273
- Release Information | 3273

Syntax

```
cfp-to-et;
```

Hierarchy Level

```
[edit chassis fpc slot]
```

Description

Make the interface et-0/1/0 (on the QSFP28 port) available for use. After you configure the `set chassis fpc 0 cfp-to-et` command and commit the configuration, you need to restart the FPC by executing the `restart chassis-control` command. After the FPC comes online, interface et-0/1/0 is created and et-0/2/1(on the CFP2 port) is deleted.

NOTE: Before executing this command, plan to handle disruption of services.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1-S1.

cgn-pic

IN THIS SECTION

- Syntax | 3273
- Hierarchy Level | 3273
- Description | 3273
- Required Privilege Level | 3274
- Release Information | 3274

Syntax

```
cgn-pic;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Description

Restrict usage of the service PIC to carrier-grade NAT (CGN) or associated services (intrusion detection, stateful firewall, and softwire). All memory is available for CGN or related services and can be used for CGN scaling.

The `cgn-pic` statement is supported only on the MS-DPC, MS-100, MS-400, and MS-500 line cards. The `cgn-pic` statement is *not* supported on MS-MPCs and MS-MICs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

> *Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card*

chain

IN THIS SECTION

- Syntax | 3274
- Hierarchy Level | 3275
- Description | 3275
- Options | 3276
- Required Privilege Level | 3276
- Release Information | 3276

Syntax

```
chain {
    expression boolean-expression;
    member (Security IDP) member-name {
        attack-type (Security Anomaly) {
```

```
            (anomaly ...same statements as in [edit security idp custom-attack attack-name
attack-type anomaly] hierarchy level | signature ...same statements as in [edit security idp
custom-attack attack-name attack-type signature] hierarchy level);
        }
    }
    order;
    protocol-binding {
        application (Security Custom Attack) application-name;
        icmp;
        icmpv6;
        ip (Security IDP Custom Attack) {
            protocol-number transport-layer-protocol-number;
        }
        ipv6 {
            protocol-number transport-layer-protocol-number;
        }
        rpc {
            program-number rpc-program-number;
        }
        tcp (Security IDP Protocol Binding) {
            minimum-port port-number <maximum-port port-number>;
        }
        udp (Security IDP Protocol Binding) {
            minimum-port port-number <maximum-port port-number>;
        }
    }
    reset;
    scope (session | transaction);
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type]
```

Description

Chain attack object combines multiple signatures and/or protocol anomalies into a single object. Traffic must match all of the combined signatures and/or protocol anomalies to match the chain attack object.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

chain-order (Application Identification)

IN THIS SECTION

- Syntax | 3276
- Hierarchy Level | 3276
- Description | 3277
- Required Privilege Level | 3277
- Release Information | 3277

Syntax

```
chain-order;
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type
signature l4-l7-signature-name]
```

Description

Read members in order. By default, chain ordering is turned off. If there is only one member, this option is ignored.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

chained-composite-next-hop

IN THIS SECTION

- Syntax | 3278
- Hierarchy Level | 3278
- Description | 3278
- Default | 3280
- Options | 3280
- Required Privilege Level | 3280
- Release Information | 3281

Syntax

```
chained-composite-next-hop {
    ingress;
    transit;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table],
[edit routing-options forwarding-table]
```

> **NOTE**: The `[edit logical-systems]` hierarchy level is not supported on the QFX10000 switches.

Description

Allows you to configure the chained composite next hops for devices handling ingress or transit traffic in the network.

Chained composite next hops help to facilitate the handling of large volumes of transit traffic in the core of large networks by allowing the router to process much larger volumes of routes. A chained composite next hop allows the router to direct sets of routes sharing the same destination to a common forwarding next hop, rather than having each route also include the destination. In the event that a network destination is changed, rather than having to update all of the routes sharing that destination with the new information, just the shared forwarding next hop is updated with the new information. The chained composite next hops continue to point to this forwarding next hop which now contains the new destination.

You can configure the chained composite next hops for devices handling ingress or transit traffic in the network. We have added support only for the following options on the MX204, MX480, MX960, MX10003, and VMX routers:

- BGP-LU for IPV4 in the ingress router—`set routing-options forwarding-table chained-composite-next-hop ingress labeled-bgp inet`
- BGP-LU on the transit router—`set routing-options forwarding-table chained-composite-next-hop transit labeled-bgp`

You can also configure class of service (CoS) and define rewrite rules for ingress and transit chained composite next hop for BGP-LU.

> **NOTE**: This feature is supported only on Modular Port Concentrator 9 (MPC9) and the previous models of line cards.

On platforms containing only MPCs, such as PTX Series Packet Transport Routers, the MX80 router, the MX2020 router, and the QFX10000 switches, chained composite next hops are enabled by default. On MX Series 5G Universal Routing Platforms containing both DPC and MPC FPCs and on T4000 Core Routers containing MPC and FPCs, chained composite next hops are disabled by default and need to be explicitly configured.

To explicitly configure chained composite next hops, include the `enhanced-ip` statement at the `[edit chassis network-services]` hierarchy level. However, take the following into consideration when enabling the `enhanced-ip` mode:

- Non-service DPCs do not work with enhanced network services mode options. Only MPCs, MS-DPCs, and MS-MPCs provide support for the `enhanced-ip` configuration.
- If you configure chained composite next hops on MX Series routers with both MPCs, and DPCs or DPCEs, the network services mode must be changed from `enhanced-ip` to `ethernet` for the DPC or DPCE to come online.

  You can verify the FPC status in such cases, using the `show chassis fpc` command output, where the DPCs and DPCEs are marked as `FPC misconfiguration`.

> **NOTE**:
>
> - When chained composite next hops are enabled on a device, the BGP connections are reset.
> - On MX Series routers with DPCs or DPCEs, only the "l3vpn" on page 9571 option is supported under the "ingress" on page 8628 configuration statement; all other configuration options and functionality are not supported.
> - On MX series routers with Modular Port Concentrator 9 (MPC9) and the previous models of line cards, only the `labeled-bgp` option and the "l3vpn" on page 9571 option are supported under the "ingress" on page 8628 configuration statement; all other configuration options and functionality are not supported.
> - The `transit` statement and the associated functionality is supported only on PTX Packet Transport Routers and QFX10000 switches.

- On MX Series routers, removing the `chained-composite-next-hop` statement from a PE device configuration causes all IBGP sessions to be torn down and triggers the BGP session to flap as well. A similar change on a router configured as a route reflector does not have any effect, however.

  The following is a sample system log message that is generated to record such an event:

  ```
  Nov  6 15:16:21.670 host PE1: rpd[6947]: bgp_peer_mgmt_clear:5995: NOTIFICATION
  sent to 10.0.100.2 (External AS 100): code 6 (Cease) subcode 4 (Administratively
  Reset), Reason: Management session cleared BGP neighbor
  ```

- Starting with Junos OS Release 17.2, you cannot configure `chained-composite-next-hop ingress l3vpn extended-space` on a logical system.

- `chained-composite-next-hop ingress l2ckt` is supported in Junos OS Release 17.3 and versions thereafter. Earlier releases of Junos do not support this statement.

The remaining statements are explained separately. See CLI Explorer.

Default

This statement is disabled by default.

Options

ingress Enable or disable composite chained next hop for ingress traffic.

transit (PTX and QFX10000) Enable or disable composite chained next hop for transit traffic. Starting in Junos OS Release 14.1, the `transit l3vpn` statement is enabled by default on PTX Series Packet Transport Routers only.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs

Chained Composite Next Hops for Transit Devices for VPNs

Example: Configuring Chained Composite Next Hops for Direct PE-PE Connections in VPNs

ingress (Chained Composite Next Hop) | **8628**

transit (Chained Composite Next Hops)

challenge-length (Static and Dynamic PPP)

IN THIS SECTION

- Syntax | **3281**
- Hierarchy Level | **3281**
- Description | **3282**
- Options | **3282**
- Required Privilege Level | **3282**
- Release Information | **3282**

Syntax

```
challenge-length minimum minimum-length maximum maximum-length;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-
interface-unit” ppp-options chap],
```

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" ppp-options
chap],
[edit interfaces pp0 unit unit-number ppp-options chap]
```

Description

Modify the length of the Challenge Handshake Authentication Protocol (CHAP) challenge by specifying the minimum and maximum allowable length, in bytes.

BEST PRACTICE: We recommend that you configure both the minimum length and the maximum length of the CHAP challenge to at least 16 bytes.

Options

minimum-length Minimum length, in bytes, of the CHAP challenge.

- **Range:** 8 through 63
- **Default:** 16

maximum-length Maximum length, in bytes, of the CHAP challenge. The *maximum-length* must be equal to or greater than the *minimum-length*.

- **Range:** 8 through 63
- **Default:** 32

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

PPP Subscriber Access Networks Overview

challenge-password

IN THIS SECTION

- Syntax | 3283
- Hierarchy Level | 3283
- Description | 3283
- Options | 3283
- Required Privilege Level | 3283
- Release Information | 3284

Syntax

```
challenge-password password;
```

Hierarchy Level

```
[edit security pki auto-re-enrollment certificate-id]
```

Description

Specify the challenge password used by the certificate authority (CA) for router certificate enrollment and revocation. This challenge password must be the same used when the router certificate was originally configured.

Options

password—The password required by the CA.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Auto-Reenrollment of a Router Certificate

auto-re-enrollment | 2727

change-configuration (Event Policy)

IN THIS SECTION

- Syntax | 3284
- Hierarchy Level | 3285
- Description | 3285
- Required Privilege Level | 3285
- Release Information | 3285

Syntax

```
change-configuration {
    commands {
        "command";
    }
    commit-options {
        check <synchronize>;
        force;
        log "comment-string";
        synchronize;
    }
    retry count number interval seconds;
```

```
    user-name username;
}
```

Hierarchy Level

```
[edit event-options policy policy-name then]
```

Description

When the associated event policy is triggered, execute the Junos OS configuration mode commands to update the candidate configuration and commit the changes.

> **NOTE**: Do not use the `change-configuration` statement to modify the configuration on dual Routing Engine devices that have nonstop active routing (NSR) enabled, because both Routing Engines might attempt to acquire a lock on the configuration database, which can cause the commit to fail. To modify the configuration through an event policy when NSR is enabled, invoke an event script that executes the commit on only the primary Routing Engine.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configure an Event Policy to Change the Configuration

Example: Change the Configuration Using an Event Policy

Example: Change the Interface Configuration in Response to an Event

commands (Event Policy Change Configuration) | **4432**

channel

IN THIS SECTION

- Syntax | 3286
- Hierarchy Level | 3286
- Description | 3287
- Options | 3287
- Required Privilege Level | 3287
- Release Information | 3287

Syntax

```
channel channel-id {
        tunnel-services {
            bandwidth(100g | 10g | 1g | 2.5g | 200g | 20g | 300g | 30g | 400g | 40g |
480g | 50g | 60g | 70g | 80g | 90g | bandwidth);
        }
    }
```

Hierarchy Level

```
[edit chassis afeb slot name pfe core],
[edit chassis cfeb slot name pfe core],
[edit chassis feb slot name pfe core],
[edit chassis fpc name pfe core],
[edit chassis lcc name fpc name pfe core],
[edit chassis member name fpc name pfe core],
[edit chassis tfeb slot name pfe core]
```

Description

Configure the Channel number.

- The channel (per core) number for ACX7024 and ACX7024X is 2.
- The channel (per core) number for ACX7100-48L and ACX7100-32C is 8.
- The channel (per core) number for ACX7509 is 8.
- The channel (per core) number for ACX7348 and ACX7332 is 4.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 24.1R1 for ACX Series Routers.

channel-speed

IN THIS SECTION

- Syntax | 3288
- Hierarchy Level | 3288
- Description | 3288
- Default | 3288
- Options | 3288
- Required Privilege Level | 3288
- Release Information | 3288

Syntax

```
channel-speed (10g | 25g | 50g | 100g | disable-auto-speed-detection) ;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number (port port-number | port-range port-range-low port-range-high)]
```

Description

Enable the specified port on the Physical Interface Card (PIC) to perform in the specified channel speed. Additionally, you can disable auto-speed detection.

Default

The default speed of the platform is the maximum speed supported on the platform. Refer *Port Settings* for the supported speed details.

Options

10g—Set the channel speed to 10g (10-Gigabit Ethernet).

25g—Set the channel speed to 25g (25-Gigabit Ethernet).

50g—Set the channel speed to 50g (50-Gigabit Ethernet).

100g—Set the channel speed to 100g (100-Gigabit Ethernet).

disable-auto-speed-detection—Disable auto-speed detection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Port Settings

chap

IN THIS SECTION

- Syntax | 3289
- Hierarchy Level | 3289
- Description | 3290
- Required Privilege Level | 3290
- Release Information | 3290

Syntax

```
chap {
    access-profile name;
    challenge-length minimum minimum-length maximum maximum-length;
    default-chap-secret name;
    local-name name;
    passive;
}
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options],
[edit interfaces interface-name unit logical-unit-number ppp-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options]
```

Description

Allow each side of a link to challenge its peer, using a “secret” known only to the authenticator and that peer. The secret is not sent over the link.

By default, PPP CHAP is disabled. If CHAP is not explicitly enabled, the interface makes no CHAP challenges and denies all incoming CHAP challenges.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- `atm-ppp-llc`—PPP over AAL5 LLC encapsulation.
- `atm-ppp-vc-mux`—PPP over AAL5 multiplex encapsulation.

> **BEST PRACTICE**: On inline service (si) interfaces for L2TP, only the `chap` statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the PPP Challenge Handshake Authentication Protocol

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

chap (Dynamic PPP)

IN THIS SECTION

- Syntax | 3291
- Hierarchy Level | 3291
- Description | 3291
- Required Privilege Level | 3291
- Release Information | 3292

Syntax

```
chap {
    challenge-length minimum minimum-length maximum maximum-length;
    local-name name;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-
interface-unit” ppp-options]
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options],
```

Description

Specify CHAP authentication in a PPP dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit” ppp-options]` hierarchy level introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Dynamic Profiles Overview

Configuring Dynamic Authentication for PPP Subscribers

Attaching Dynamic Profiles to Static PPP Subscriber Interfaces

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

chap (L2TP)

IN THIS SECTION

- Syntax | 3292
- Hierarchy Level | 3292
- Description | 3293
- Required Privilege Level | 3293
- Release Information | 3293

Syntax

```
chap;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp ppp-options]
```

Description

(MX Series routers only) Specify CHAP authentication for PPP subscribers in an L2TP LNS user group profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

charging-id (OCS Partition)

IN THIS SECTION

- Syntax | 3294
- Hierarchy Level | 3294
- Description | 3294
- Options | 3294
- Required Privilege Level | 3294
- Release Information | 3294

Syntax

```
charging-id number;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Include the 3GPP-Charging-Id AVP value in all CCR-GY messages.

Options

number 3GPP-Charging-Id AVP value to include in all CCR-GY messages.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

charging-service-list

IN THIS SECTION

- Syntax | 3295
- Hierarchy Level | 3295
- Description | 3295
- Options | 3295
- Required Privilege Level | 3296
- Release Information | 3296

Syntax

```
charging-service-list ocs;
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure the list of charging services with which to communicate for the broadband Policy and Charging Enforcement Function (BPCEF).

NOTE: Currently, if you configure this statement, you must also configure the `provisioning-order` statement to `pcrf`.

Options

ocs Use Online Charging Services (OCS) as the list of charging services with which to communicate.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

chassis

IN THIS SECTION

- Syntax | 3296
- Hierarchy Level | 3299
- Description | 3299
- Options | 3299
- Required Privilege Level | 3300
- Release Information | 3300

Syntax

```
chassis {
    aggregated-devices {...}
    alarm {...}
```

```
alarm-port {...}
ambient-temperature;
auto-image-upgrade;
auto-satellite-conversion;
cb name {...}
config-button (no-clear | no-rescue);
container-devices {...}
copy-tos-to-outer {...}
craft-lockout;
dedicated-ukern-cpu (BFD);
disable-fm;
disable-power-management;
disk-partition {...}
ecmp-alb {...}
effective-shaping-rate;
enhanced-policer;
epacl-firewall-optimization;
error {...}
extended-statistics;
fabric {...}
forwarding-options {...}
fpc name {...}
fpc-empolicy-profile low-power-mode;
fpc-ifl-ae-statistics;
fpc-resync;
fru-local-storage-disable;
fru-poweron-sequence;
host-outbound media-interface {...}
icmp {...}
icmp6 {...}
input-firewall-optimized-profile;
ioc-npc-connectivity {...}
ipvacl-firewall-l2-profile;
iracl-firewall-ipv4-profile;
ivacl-firewall-no-portrange-profile;
jnu-management {...}
lcc name {...}
lcd-menu {...}
limited-ifl-scaling;
local-station-mac local-station-mac;
loopback-firewall-optimization;
loopback-firewallv6-optimization;
maximum-ecmp;
```

```
    memory-enhanced {...}
    mode power-optimized;
    monitor-pem-inputs;
    multicast-loadbalance {...}
    network-services;
    network-slices {...}
    no-power-budget;
    no-psu-redundancy;
    nssu {...}
    onbox-av-load-flavor heavy;
    optics display-3po-pn;
    oss-map {...}
    p4rt node-id node-id;
    pem {...}
    pem {...}
    periodic {...}
    per-logical-interface-firewall;
    policer-drop-probability-low;
    policer-limit;
    power-off-ports-on-no-master-re (disable | enable);
    ppp-subscriber-services;
    preserve-fpc-poweron-sequence;
    provider-instance-devices device-count;
    pseudowire-service {...}
    psm {...}
    psu {...}
    realtime-ukern-thread (BFD);
    redundancy {...}
    redundancy-group {...}
    resource-watch {...}
    route-localization {...}
    routing-engine {...}
    satellite satellite-name;
    serdes test-mode test-mode;
    sfm {...}
    sib;
    slow-pfe-alarm;
    synchronization {...}
    synchronization {...}
    synchronization {...}
    synchronization {...}
    turbotx-disable;
    usb {...}
```

```
    vrf-mtu-check;
}
```

Hierarchy Level

```
[edit]
```

Description

Configure router chassis properties.

> **NOTE**: Use Feature Explorer to confirm platform and release support for specific features.

Options

config-button Config button behavior settings.

- no-clear—Don't reset to factory-default configuration.
- no-rescue—Don't reset to rescue configuration.

disable-fm Disable Fabric Manager.

disable-power-management Disable Power Management in this chassis.

fpc-empolicy-profile low-power-mode (Only for MX10004 and MX10008 devices with LC9600 line card) Enable fpc-empolicy-profile low-power-mode.

fpc-ifl-ae-statistics Enable fpc ifl ae child statistics.

fru-local-storage-disable Disable and erase local storage on FPC. (FPC reboot required)

local-station-mac Local station mac address configuration.

mode power-optimized Set power mode to power-optimized.

monitor-pem-inputs Enable low current usage monitoring.

no-psu-redundancy Disable N+1 Redundancy for PSUs.

onbox-av-load-flavor heavy (Only for SRX devices) Select onbox anti-virus traffic load flavor.

optics display-3po-pn To display third party optics part number.

p4rt node-id Chassis p4rt identifier.

- **Range**: 1 through 18446744073709551615

per-logical-interface-firewall Ingress port firewall can be supported per logical interface, restarts PFE.

policer-limit Limit the policer tick unit to 32 bytes.

power-off-ports-on-no-master-re Powering off ports if no master RE present.

- disable—Disable powering off ports if no master RE present.
- enable—Enable powering off ports if no master RE present.

provider-instance-devices Number of provider instance port devices.

satellite *satellite-name* Display JNU satellite configuration schema. Connected Security Distributed Services (CSDS) Architecture support this option.

serdes test-mode Serdes test mode for link training.

turbotx-disable Disable turbotx processing. Turbotx host path enables higher packets-per-second (pps) for certain junos applications. Only supported for Linux based line cards.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

`satellite` *`satellite-name`* option is added in Junos OS Release 24.4R1.

Statement introduced before Junos OS Release 7.4.

chassis-id (LLDP)

IN THIS SECTION

- Syntax | 3301
- Hierarchy Level | 3301
- Description | 3301
- Options | 3302
- Required Privilege Level | 3302
- Release Information | 3302

Syntax

```
chassis-id {
 chassis-id-type (chassis-component | interface-alias | interface-name | locally-assigned | mac-address | network-address | port-component);
 chassis-id-value chassis-id-value;
}
```

Hierarchy Level

```
[edit logical-systems name protocols lldp],
[edit protocols lldp]
```

Description

Configure the chassis-id to be used to generate the Chassis ID TLV advertised to the LLDP neighbors.

Options

chassis-id-type Chassis-id-type to be used for Chassis ID TLV generation

- chassis-component—User define chassis component name.
- interface-alias—Interface description or name on which LLDPDU is sent.
- interface-name—Interface name on which LLDPDU is sent.
- locally-assigned—Locally assigned value (SNMP index of the interface).
- mac-address—MAC Address of Interface on which LLDPDU is sent.
- network-address—Netwok Address of Interface on which LLDPDU is sent.
- port-component—User define port component name.

chassis-id-value Chassis-id value as per chassis id type.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

chassis-id (Aggregated Ethernet Interface)

IN THIS SECTION

- Syntax | 3303
- Hierarchy Level | 3303
- Description | 3303
- Required Privilege Level | 3303
- Release Information | 3303

Syntax

```
chassis-id chassis-id;
```

Hierarchy Level

```
[edit interfaces aggregated-ether-options mc-ae]
```

Description

Specify the chassis ID of the multichassis aggregated Ethernet interface device. LACP uses the chassis ID to calculate the port number of the multichassis link aggregation group (MC-LAG) physical member links.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

chassis (QFabric System)

IN THIS SECTION

- Syntax | 3304
- Hierarchy Level | 3306
- Description | 3306
- Required Privilege Level | 3306
- Release Information | 3306

Syntax

```
chassis {
    interconnect-device {
        alarm {
            (ethernet (Alarm) | management-ethernet) {
                link-down (red | yellow | ignore);
            }
        }
        container-devices {
            device-count number;
        }
        craft-lockout {
            alarm {
                interface-type {
                    link-down (red | yellow | ignore);
                }
            }
            container-devices {
                device-count number;
            }
            fpc slot {
                power (on | off);
            }
            routing-engine {
                on-disk-failure {
                    disk-failure-action (halt | reboot);
                }
            }
        }
        fpc slot {
            power (on | off);
        }
        routing-engine {
            on-disk-failure {
                disk-failure-action (halt | reboot);
            }
        }
    }
    node-group name {
        aggregated-devices {
            ethernet {
```

```
                device-count number;
            }
        }
        alarm {
            interface-type {
                link-down (ignore | red | yellow);
            }
        }
        container-devices {
            device-count number;
        }
        node-device name {
            fibre-channel {
                port-range {
                    port-range-low port-range-high;
                }
            }
            pic pic-number {
                fte {
                    port port-number;
                    port-range port-range-low port-range-high;
                }
                xe {
                    port port-number;
                    port-range port-range-low port-range-high;
                }
                xle {
                    port port-number;
                    port-range port-range-low port-range-high;
                }
            }
        }
        routing-engine {
            on-disk-failure {
                disk-failure-action (halt | reboot);
            }
        }
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure chassis-specific properties for the switch.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

chassis (Subscriber Limits)

IN THIS SECTION

- Syntax | 3307
- Hierarchy Level | 3307
- Description | 3307
- Options | 3307
- Required Privilege Level | 3307
- Release Information | 3307

Syntax

```
chassis {
    limit limit;
}
```

Hierarchy Level

```
[edit system services resource-monitor subscribers-limit client-type name]
```

Description

Configure the maximum number of subscribers of the specified client type allowed to be logged in on the chassis. When that number is reached, subsequent logins on the chassis are denied until the current number of subscribers drops below the maximum allowed. You can also specify the maximum number of subscribers of a client type allowed per port, per MIC, and per MPC.

Options

limit Maximum number of subscribers.

- **Range:** `1` through `1,000,000`

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

> *Limiting Subscribers by Client Type and Hardware Element with Resource Monitor*
>
> *Resource Monitoring for Subscriber Management and Services*

check-bytes (Application Identification)

IN THIS SECTION

- Syntax | 3308
- Hierarchy Level | 3308
- Description | 3308
- Options | 3308
- Required Privilege Level | 3309
- Release Information | 3309

Syntax

```
check-bytes max-bytes-to-check;
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type
signature l4-l7-signature-name member member-name]
```

Description

Specify the maximum number of bytes to be inspected. This statement applies to TCP and UDP protocols for stream context. It is not considered for other protocols and contexts.

Options

max-bytes-to-check

Number of bytes to be inspected.

Range: 1 through 5000

Default: Not configured

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

check-zero

IN THIS SECTION

- Syntax | 3309
- Hierarchy Level | 3310
- Description | 3310
- Default | 3310
- Required Privilege Level | 3310
- Release Information | 3310

Syntax

```
(check-zero | no-check-zero);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip],
[edit protocols rip group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

Check whether the reserved fields in a RIP packet are zero:

- **check-zero**—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.
- **no-check-zero**—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.

Default

check-zero

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

checksum (Scripts)

IN THIS SECTION

- Syntax | 3311
- Hierarchy Level | 3311
- Description | 3311
- Options | 3312
- Required Privilege Level | 3312
- Release Information | 3312

Syntax

```
checksum (md5 | sha-256 | sha1) hash;
```

Hierarchy Level

```
[edit event-options event-script file filename],
[edit system scripts commit file filename],
[edit system scripts op file filename],
[edit system scripts snmp file filename],
[edit system extensions extension-service application file filename]
```

Description

For Junos OS commit scripts, event scripts, op scripts, SNMP scripts, and scripts developed using the Juniper Extension Toolkit (JET) specify the *MD5*, *SHA-1*, or SHA-256 checksum hash. When Junos OS executes a local commit, event, op, SNMP, or JET script, the system verifies the integrity of the script by using the configured checksum hash.

Options

md5 *hash* MD5 checksum of this script. This option is deprecated starting in Junos OS Releases 18.2R2 and 18.3R1.

sha-256 *hash* SHA-256 checksum of this script.

sha1 *hash* SHA-1 checksum of this script. This option is deprecated starting in Junos OS Releases 18.2R2 and 18.3R1.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Statement introduced for the `[edit system extensions extension-service application file filename]` hierarchy level in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.

`md5` and `sha1` options deprecated in Junos OS Releases 18.2R2 and 18.3R1.

RELATED DOCUMENTATION

Configuring Checksum Hashes for a Commit Script

Configuring Checksum Hashes for an Event Script

Configuring Checksum Hashes for an Op Script

Configuring Checksum Hashes for an SNMP Script

Executing an Op Script from a Remote Site

file checksum md5 | **7037**

file checksum sha-256 | **7039**

file checksum sha1 | **7041**

checksum-validate

IN THIS SECTION

- Syntax | 3313
- Hierarchy Level | 3313
- Description | 3313
- Options | 3313
- Required Privilege Level | 3314
- Release Information | 3314

Syntax

```
checksum-validate {
    match (equal | greater-than | less-than | not-equal);
    value checksum-value;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol ipv4]
[edit security idp custom-attack attack-name attack-type signature protocol tcp]
[edit security idp custom-attack attack-name attack-type signature protocol udp]
[edit security idp custom-attack attack-name attack-type signature protocol icmp]
[edit security idp custom-attack attack-name attack-type signature protocol icmpv6]
```

Description

Allow IDP to validate checksum field against the calculated checksum.

Options

match (equal | greater-than | less-than | not-equal)	Match an operand.

value *checksum-value*

Match a decimal value.

- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

child-inactivity-timeout

IN THIS SECTION

- Syntax | 3314
- Hierarchy Level | 3314
- Description | 3315
- Options | 3315
- Required Privilege Level | 3315
- Release Information | 3315

Syntax

```
child-inactivity-timeout seconds;
```

Hierarchy Level

```
[edit applications application ike-esp-nat]
```

Description

For an IKE ALG application, configure the ESP session (IPsec data traffic) idle timeout. If no IPsec data traffic is passed on the ESP session in this time, the session is deleted.

The IKE ALG enables the passing of IKEv1 and IPsec packets through NAPT-44 and NAT64 rules between IPsec peers that are not NAT-T compliant.

Options

seconds

Number of seconds.

- **Default:** 800 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

ALG Descriptions

Configuring Application Sets

Configuring Application Properties

children

IN THIS SECTION

- Syntax | **3316**
- Hierarchy Level | **3316**

- Description | 3316
- Options | 3317
- Required Privilege Level | 3318
- Release Information | 3318

Syntax

```
children {
    child-name {
        esp-proposal esp-proposal;
        local-traffic-selector {
            (ipv4-prefix ipv4-prefix | ipv6-prefix ipv6-prefix);
            port port;
            protocol protocol;
        }
        mode (transport | tunnel);
        rekey-time rekey-time;
        remote-traffic-selector{
            (ipv4-prefix ipv4-prefix | ipv6-prefix ipv6-prefix);
            port port;
            protocol protocol;
        }
    }
}
```

Hierarchy Level

```
[edit security host-vpn connections connection-name]
```

Description

Configure child details to establish a security association (SA). An SA describes a specific negotiated set of parameters to protect traffic between two host for a certain period of time.

Options

child-name Specify the child SA name.

esp-proposal *esp-proposal* Specify the algorithms to use in negotiating the child SA from among the pre-selected combinations available, which represent the encryption algorithm, integrity algorithm, and Diffie Hellman group. There are the following options:

- **3des-sha1-modp1536** Propose 3des SHA1 and DH group modp1536.
- **aes256gcm128-ecp384** Propose aes256gcm128 and DH group ecp384.
- **aes256gcm128-modp3072** Propose aes256gcm128 and DH group modp3072.
- **aes256-sha384-ecp384** Propose aes256 CBC, sha384 and DH group ecp384.
- **aes256-sha384-modp3072** Propose aes256 CBC, sha384 and DH group modp3072.
- **[ ]** Propose a set composed from the values permitted.

- **Default:** aes256gcm128-ecp384

mode (transport | tunnel) Specify the IPsec usage mode to negotiate: `transport` or `tunnel`.

- **tunnel** In tunnel mode, the entire IP packet is encrypted and authenticated. It is then encapsulated into a new IP packet with a new IP header. Tunnel mode may be used with any kind of IP traffic. Of the two modes, only tunnel mode supports NAT transversal. Tunnel mode is required if you are communicating with a server behind a gateway.
- **transport** In transport mode, only the payload of the IP packet is encrypted or authenticated. The IP header is neither modified nor encrypted. Transport mode does not support NAT transversal. Transport mode or tunnel mode can be used when communications is between two hosts, for example, between a router and a Syslog server.

- **Default:** `tunnel`

rekey-time *rekey-time* Specify how long, in seconds, before the child SA is rekeyed. Actual rekeying occurs slightly sooner than the rekey time specified because of rekey randomization.

- **Default:** 14,400
- **Range:** 60 through 86,400

The remaining statements are explained separately.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 18.3R1.

cipher-list (REST API)

IN THIS SECTION

- Syntax | 3318
- Hierarchy Level | 3318
- Description | 3319
- Options | 3319
- Required Privilege Level | 3320
- Release Information | 3320

Syntax

```
cipher-list [ cipher-1 cipher-2 cipher-3 ...];
```

Hierarchy Level

```
[edit system services rest https]
```

Description

Specify the set of ciphers the server can use to perform encryption and decryption functions. If this option is not configured, the server accepts any supported suite that is available.

Options

- `rsa-with-RC4-128-md5`— RSA, 128- bit RC4, MD5 hash
- `rsa-with-RC4-128-sha`— RSA, 128-bit RC4, SHA hash
- `rsa-with-3DES-ede-cbc-sha`— RSA, 3DES EDE/CBC, SHA hash
- `dhe-rsa-with-3DES-ede-cbc-sha`— DHE/RSA, 3ES/EDE CBC, SHA hash
- `rsa-with-aes-128-cbc-sha`— RSA, 128-bit AES/CBC, SHA hash
- `dhe-rsa-with-aes-128-cbc-sha`— DHE/RSA, 128-bit AES/CBC, SHA hash
- `rsa-with-aes-256-cbc-sha`— RSA, 256 bit AES/CBC, SHA hash
- `dhe-rsa-with-aes-256-cbc-sha`— DHE/RSA, 256 bit AES/CBC, SHA hash
- `ecdhe-rsa-with-RC4-128-sha`— ECDHE/RSA, 128-bit RC4, SHA hash
- `ecdhe-rsa-with-3DES-ede-cbc-sha`— ECDHE/RSA, 128-bit 3DES EDE/CBC SHA hash
- `ecdhe-rsa-with-aes-128-cbc-sha`— ECDHE/RSA, 128-bit AES/CBC, SHA hash
- `ecdhe-rsa-with-aes-256-cbc-sha`— ECDHE/RSA, 256 bit AES/CBC, SHA hash
- `rsa-with-aes-128-cbc-SHA256`— RSA, 128-bit AES/CBC, SHA256 hash
- `rsa-with-aes-256-cbc-SHA256`— RSA, 256 bit AES/CBC, SHA256 hash
- `dhe-rsa-with-aes-128-cbc-SHA256`— DHE/RSA, 128-bit AES/CBC, SHA256 hash
- `dhe-rsa-with-aes-256-cbc-SHA256`— DHE/RSA, 256 bit AES/CBC, SHA256 hash
- `rsa-with-aes-128-gcm-SHA256`— RSA, 128-bit AES/GCM, SHA256 hash
- `rsa-with-aes-256-gcm-SHA384`— RSA, 256 bit AES/GCM, SHA384 hash
- `dhe-rsa-with-aes-128-gcm-SHA256`— DHE/RSA, 128-bit AES/GCM, SHA256 hash
- `dhe-rsa-with-aes-256-gcm-SHA384`— DHE/RSA, 256 bit AES/GCM, SHA384 hash
- `ecdhe-rsa-with-aes-128-cbc-SHA256`— ECDHE/RSA, 128-bit AES/CBC, SHA256 hash

- `ecdhe-rsa-with-aes-256-cbc-SHA384`— ECDHE/RSA, 256 bit AES/CBC, SHA384 hash
- `ecdhe-rsa-with-aes-128-gcm-SHA256`— ECDHE/RSA, 128-bit AES/GCM, SHA256 hash
- `ecdhe-rsa-with-aes-256-gcm-SHA384`— ECDHE/RSA, 256 bit AES/GCM, SHA384 hash

> **NOTE**: For Junos OS in FIPS mode, only the following FIPS-compliant cipher algorithms are supported:
>
> - `rsa-with-aes-256-gcm-SHA384`— RSA, 256 bit AES/GCM, SHA384 hash
> - `dhe-rsa-with-aes-128-gcm-SHA256`— DHE/RSA, 128-bit AES/GCM, SHA256 hash
> - `dhe-rsa-with-aes-256-gcm-SHA384`— DHE/RSA, 256 bit AES/GCM, SHA384 hash
> - `ecdhe-rsa-with-aes-128-gcm-SHA256`— ECDHE/RSA, 128-bit AES/GCM, SHA256 hash
> - `ecdhe-rsa-with-aes-256-gcm-SHA384`— ECDHE/RSA, 256 bit AES/GCM, SHA384 hash

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

cipher-suite (MACsec)

IN THIS SECTION

- Syntax | 3321
- Hierarchy Level | 3321
- Description | 3321
- Default | 3322
- Options | 3322

Required Privilege Level | 3322
Release Information | 3322

Syntax

```
cipher-suite encryption-algorithm-name;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Specify the set of ciphers used to encrypt traffic on an Ethernet link that is secured with Media Access Control Security (MACsec). The encryption used by MACsec ensures that the data in the Ethernet frame cannot be viewed by anybody monitoring traffic on the link. MACsec encryption is optional and user-configurable. The configured cipher suites should be the same between MACsec peers.

MACsec utilizes the Galois/Counter Mode Advanced Encryption Standard (GCM-AES). The default cipher suite used for MACsec is GCM-AES-128, with a maximum key length of 128 bits. MACsec also supports GCM-AES-256, with a maximum key length of 256 bits.

GCM– AES– 128 and GCM– AES– 256 use a 32-bit packet number as part of the initial value that has to be unique for every packet sent with a given secure association key (SAK). When the permutations of the 32-bit packet number are exhausted, the SAK much be refreshed. The frequency of SAK refreshes can be reduced by using a cipher suite with Extended Packet Numbering (XPN), which increases the size of the packet number to 64-bits. Both GCM-AES-128 and GCM-AES-256 are available with XPN.

NOTE: When enabling MACsec on *et* interfaces, use either the GCM-AES-XPN-128 or GCM-AES-XPN-256 cipher suite.

NOTE: EX4300-48MP switches support GCM-AES-256 and GCM-AES-XPN-256 cipher suites on the 2-port QSFP+/1-port QSFP28 uplink module.

Default

If the `cipher-suite` statement is not configured, the default cipher suite used for encryption is GCM-AES-128.

Options

`gcm-aes-128`	GCM-AES-128 has a maximum key size of 128 bits.
`gcm-aes-xpn-128`	GCM-AES-XPN-128 has a maximum key size of 128 bits and extended packet number.
`gcm-aes-256`	GCM-AES-256 has a maximum key size of 256 bits.
`gcm-aes-xpn-256`	GCM-AES-XPN-256 has a maximum key size of 256 bits and extended packet number.

Required Privilege Level

admin— To view this statement in the configuration.

admin-control— To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2R1.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Configuring Media Access Control Security (MACsec) on Routers

circuit-id

IN THIS SECTION

- Syntax | 3323
- Hierarchy Level | 3323

- Description | 3323
- Default | 3324
- Required Privilege Level | 3324
- Release Information | 3324

Syntax

```
circuit-id {
    prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
    }
    use-interface-description (device | logical);
    use-vlan-id;
}
```

Hierarchy Level

- For platforms with Enhanced Layer 2 Software (ELS):

```
[edit vlans vlan-name forwarding-options dhcp-security option-82 ]
```

- For MX Series platforms:

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security option-82]
```

Description

Configure the `circuit-id` suboption (suboption 1) of DHCP option 82 (the DHCP relay agent information option) in DHCP packets destined for a DHCP server. This suboption identifies the circuit (the interface, the VLAN, or both) on which the DHCP request arrived.

The remaining statements are explained separately. See CLI Explorer.

Default

If DHCP option 82 is enabled on the switch, the circuit ID is supplied by default in the format *interface-name*:*vlan-name* or, on a Layer 3 interface, just *interface-name*.

> **NOTE**: When you configure `circuit-id`, `remote-id` is also enabled, even if you do not explicitly configure `remote-id` .

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

> *Example: Setting Up DHCP Option 82*
>
> http://tools.ietf.org/html/rfc3046

circuit-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 3325
- Hierarchy Level | 3325
- Description | 3325
- Required Privilege Level | 3327
- Release Information | 3327

Syntax

```
circuit-id {
    include-irb-and-l2;
    keep-incoming-circuit-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-vlan-id;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82],
[edit forwarding-options dhcp-relay group group-name relay-option-82],
[edit logical-systems logical-system-name forwarding-options dhcp-relay relay-option-82],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name relay-
option-82],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay relay-option-82],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name relay-option-82],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
relay-option-82]
```

Description

Specify the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option (option 82) to include in DHCP packets destined for a DHCP server. Optionally specify that the suboption includes a prefix, textual description, or VLAN tag.

> **NOTE**: For Layer 3 interfaces, when you configure relay-option-82 only, the Agent Circuit ID is the default. If no VLAN tags are configured, then the default is the logical interface device (IFL) name. For integrated routing and bridging (IRB) interfaces, the default is the Layer 2 IFL name and bridge domain name.

> The interface to bridge domain relationship may be implicit (the interface is mapped to the bridge domain by the system based on VLAN tag) or explicit (the interface is mapped to the bridge domain by configuring it in the bridge domain definition). For the explicit case, tagging might not be relevant for the mapping.

The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual LANs (VLANs), stacked VLANs (S-VLANs), or bridge domains is as follows:

```
(fe | ge)-fpc/pic/port.subunit
```

> **NOTE**: For remote systems, the *subunit* is required and is used to differentiate an interface for remote systems.

The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that use VLANs is as follows:

```
(fe | ge)-fpc/pic/port:vlan-id
```

The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs is as follows:

```
(fe | ge)-fpc/pic/port:svlan-id-vlan-id
```

In the case of an IRB interface, the format displays the Layer 2 interface instead of the IRB interface along with the bridge domain name. For IRB interfaces (or other pseudo devices) the default format is as follows:

- IRB interfaces that use bridge domains but do not use VLANs or S-VLANs:

  ```
  (fe | ge)-fpc/pic/port.subunit:bridge-domain-name
  ```

- IRB interfaces that use VLANs:

  ```
  (fe | ge)-fpc/pic/port.subunit:vlan-name
  ```

To include the IRB interface name with the Layer 2 interface name, configure the `include-irb-and-l2` statement. The format is as follows:

- IRB interfaces that use bridge domains but do not use VLANs or S-VLANs:

```
(fe | ge)-fpc/pic/port:bridge-domain-name+irb.subunit
```

- IRB interfaces that use VLANs:

```
(fe | ge)-fpc/pic/port:vlan-name+irb.subunit
```

To include only the IRB interface name without the Layer 2 interface and bridge domain or VLAN, configure the `no-vlan-interface-name` statement. The format is as follows:

```
irb.subunit
```

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

`include-irb-and-l2` , `no-vlan-interface-name`, and `use-vlan-id` options added in Junos OS Release 14.1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
21.2R1	Starting with Junos OS Release 21.2R1, QFX Series switches support the `circuit-id` CLI statement in a stateless DHCP relay configuration. You can configure stateless DHCP relay using the `forward-only` CLI statement at the [`edit forwarding-options dhcp-relay`] hierarchy level.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information

Configuring Option 82 Information

circuit-id (VLAN Authentication Username)

IN THIS SECTION

- Syntax | 3328
- Hierarchy Level | 3328
- Description | 3328
- Required Privilege Level | 3329
- Release Information | 3329

Syntax

```
circuit-id;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include]
```

Description

Include the agent circuit identifier (ACI) in the username sent to RADIUS for authentication of the dynamic VLAN. The ACI is conveyed by the Access-Loop-Circuit-ID TLV in an out-of-band ANCP Port Up message.

NOTE: This statement is not supported for stacked VLANs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs

Configuring VLAN Interface Username Information for AAA Authentication

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

circuit-type

IN THIS SECTION

- Syntax | 3329
- Hierarchy Level | 3330
- Description | 3330
- Required Privilege Level | 3330
- Release Information | 3330

Syntax

```
circuit-type;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication username-include],
```

Description

Specify that the circuit type is concatenated with the username during the subscriber authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

circuit-type (DHCP Local Server)

IN THIS SECTION

- Syntax | 3331
- Hierarchy Level | 3331
- Description | 3332
- Required Privilege Level | 3332
- Release Information | 3332

Syntax

```
circuit-type;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication username-
include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  authentication
username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
```

```
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify that the circuit type is concatenated with the username during the subscriber authentication or client authentication process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

circuit-type (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 3333
- Hierarchy Level | 3333
- Description | 3333
- Required Privilege Level | 3333
- Release Information | 3333

Syntax

```
circuit-type;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-
include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify that the circuit type is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Specifying Authentication Support

Creating Unique Usernames for DHCP Clients

cisco-interoperability

IN THIS SECTION

- Syntax | 3334
- Hierarchy Level | 3334
- Description | 3334
- Options | 3334
- Required Privilege Level | 3335
- Release Information | 3335

Syntax

```
cisco-interoperability send-lip-remove-link-for-link-reject;
```

Hierarchy Level

```
[edit interfaces interface-name mlfr-uni-nni-bundle-options]
```

Description

FRF.16 interoperability settings.

Options

`send-lip-remove-link-for-link-reject`—Send Link Integrity Protocol remove link when an add-link rejection message is received.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring SONET APS Interoperability with Cisco Systems FRF.16

cisco-nas-port-info (L2TP Tunnel Switching)

IN THIS SECTION

- Syntax | 3335
- Hierarchy Level | 3335
- Description | 3336
- Options | 3336
- Required Privilege Level | 3336
- Release Information | 3336

Syntax

```
cisco-nas-port-info action;
```

Hierarchy Level

```
[edit access tunnel-switch-profile profile-name avp]
```

Description

Define a tunnel profile for subscriber access.

Specify the action taken on the Cisco NAS Port Info AVP (100) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.

Options

action—One of the following actions:

- `drop`—Drop the AVP.
- `regenerate`—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.
- `relay`—Forward the AVP transparently as is and send it in the switched packet.
- **Default:** `relay`

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

civic-based

IN THIS SECTION

- Syntax | 3337
- Hierarchy Level | 3337
- Description | 3337
- Default | 3338
- Options | 3338
- Required Privilege Level | 3338
- Release Information | 3338

Syntax

```
civic-based {
    ca-type name {
        ca-value  ca-value;
    }
    country-code country-code;
    what what;
}
```

Hierarchy Level

```
[edit protocols lldp-med interface (all | interface-name) location]
```

Description

For Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED), configure the geographic location to be advertised from the device to the MED. This information is used during emergency calls to identify the location of the MED.

Default

Disabled.

Options

country-code *code*

(Required) Configure the two-letter ISO 3166 country code in capital ASCII letters; for example, US or DE. The code is part of the location information. Location information is advertised from the device to the MED, and is used during emergency calls to identify the location of the MED. The country code is required when configuring LLDP-MED based on location.

- **Default:** Disabled.

what *number*

Configure the location to which the DHCP entry refers. This information is advertised, along with other location information, from the switch to the MED. It is used during emergency calls to identify the location of the MED.

Options 0 and 1 should not be used unless you know that the DHCP client is in close physical proximity to the server or network element.

- **Values:** Location to which the DHCP entry refers:
  - 0—Location of the DHCP server.
  - 1—Location of a network element believed to be closest to the client.
  - 2—Location of the client.
- **Default:** 1

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`what` modified in Junos OS Release 9.2 for EX Series to display new default.

RELATED DOCUMENTATION

show lldp | 21942

Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch

Configuring LLDP-MED (CLI Procedure)

ckn

IN THIS SECTION

- Syntax | 3339
- Hierarchy Level | 3339
- Description | 3339
- Default | 3340
- Options | 3340
- Required Privilege Level | 3340
- Release Information | 3340

Syntax

```
ckn hexadecimal-number;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name pre-shared-key],
[edit security macsec connectivity-association connectivity-association-name fallback-key]
```

Description

Specifies the connectivity association key name (CKN) for a preshared key.

A preshared key includes a CKN and a connectivity association key (CAK). A preshared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security

keys. The CKN is a 64-digit hexadecimal number and the CAK is a 32-digit hexadecimal number. The MACsec Key Agreement (MKA) protocol is enabled once the preshared keys are successfully exchanged. The preshared key—the CKN and CAK—must match on both ends of a link.

To configure MACsec on supported ports, you need to create the preshared key by configuring the connectivity association key name (CKN) and connectivity association key (CAK).

Default

No CKN exists, by default.

Options

hexadecimal-number The key name, in hexadecimal format.

The key name is 64 hexadecimal characters in length. To maximize security, we recommend configuring all 64 digits of a CKN. If you enter a key name that is less then 64 characters long, the remaining characters are set to 0. However, you will receive a warning message when you commit the configuration.

On SRX devices, the key name is 32 hexadecimal characters in length. If you enter a key name that is less than 32 characters long, the remaining characters are set to 0.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced on SRX devices in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

ckn (MX Series)

IN THIS SECTION

- Syntax | 3341
- Hierarchy Level | 3341
- Description | 3341
- Default | 3341
- Options | 3342
- Required Privilege Level | 3342
- Release Information | 3342

Syntax

```
ckn hexadecimal-number;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name pre-shared-key]
```

Description

Specifies the connectivity association key name (CKN) for a pre-shared key.

A pre-shared key includes a CKN and a connectivity association key (CAK). A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security keys. The MACsec Key Agreement (MKA) protocol is enabled once the pre-shared keys are successfully exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link

Default

No CKN exists, by default.

Options

hexadecimal-number

The key name, in hexadecimal format.

The key name is 32 hexadecimal characters in length. If you enter a key name that is less then 32 characters long, the remaining characters are set to 0.

- On MX10003 router, if you configure the length of CKN to the value less than 64 hexadecimal digits, then the following warning message is displayed:

  ```
  warning: To maximize security, recommend configuring all 64 digits of pre-shared-key ckn
  ```

- On MX10003 router, if you configure the length of CKN to an odd value, then the commit will not be successful and the following error message is displayed:

  ```
  error: ckn: 'abcde': Must be an even-length string up to 64 hexadecimal digits (0-9, a-f, A-F)
  ```

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

class

IN THIS SECTION

- Syntax | 3343
- Hierarchy Level | 3343

- Description | 3343
- Options | 3343
- Required Privilege Level | 3344
- Release Information | 3344

Syntax

```
class {
    alg-logs;
    deterministic-nat-configuration-log;
    ids-logs;
    nat-logs;
    packet-logs;
    pcp-logs;
    session-logs <open | close>;
    stateful-firewall-logs ;
    urlf-logs;
}
```

Hierarchy Level

```
[edit services service-set service-set-name syslog host hostname]
```

Description

Set the class of applications to be logged to the system log.

Starting in Junos OS Release 17.4R1, you can configure up to a maximum of four system log servers (combination of local system log hosts and remote system log collectors) for each service set at the `[edit services service-set service-set-name]` hierarchy level.

Options

class-name Enter one of the following values:

- `alg-logs`—Log application-level gateway events.
- `deterministic-nat-configuration-log`—Log deterministic NAT sessions.
- `ids-logs`—Log intrusion detection system events.
- `nat-logs`—Log Network Address Translation events.
- `packet-logs`—Log general packet-related events.
- `pcp-logs`—Log Port Control Protocol events.
- `session-logs`—Log session open and close events.
- `session-logs open`—Log session open events only.
- `session-logs close`—Log session close events.
- `urlf-logs`—Log events for the filtering of DNS requests for blocklisted website domains.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

`deterministic-nat-configuration-log` option introduced in Junos OS Release 17.3R1.

You can configure multiple system log hosts from Junos OS Release 17.4R1 onwards.

`urlf-logs` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Configuring System Logging for Service Sets

class (Unified Edge)

IN THIS SECTION

- Syntax | 3345
- Hierarchy Level | 3345
- Description | 3345
- Required Privilege Level | 3345
- Release Information | 3346

Syntax

```
class {
    equals value;
    has-prefix value;
    has-suffix value;
    matches value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from]
```

Description

Specify the RADIUS AVP class for the incoming RADIUS request from the IP-based subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

class (CoS-Based Forwarding)

IN THIS SECTION

- Syntax | 3346
- Hierarchy Level | 3347
- Description | 3347
- Options | 3347
- Required Privilege Level | 3347
- Release Information | 3347

Syntax

```
class class-name {
    classification-override {
        forwarding-class class-name;
    }
}
```

Hierarchy Level

```
[edit class-of-service forwarding-policy]
```

Description

Configure CoS-based forwarding class.

Options

class-name—Name of the routing policy class.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Overriding the Input Classification

class (Defining Login Classes)

IN THIS SECTION

- Syntax | 3348
- Hierarchy Level | 3349
- Description | 3349

- Options | 3349
- Required Privilege Level | 3359
- Release Information | 3359

Syntax

```
class class-name {
 allow-hidden-commands;
 no-hidden-commands {
  except [“regular expression or command 1” “regular expression or command 2” ...];

 }
 access-end hh:mm;
 access-start hh:mm;
 allow-commands (regular-expression1)|(regular-expression2)...;
 allow-commands-regexps [“regular expression 1” “regular expression 2 ” ... ];
 allow-configuration (regular-expression1)|(regular-expression2)...;
 allow-configuration-exact-match configuration;
 allow-configuration-regexps (regular-expression1)|(regular-expression2)...;
 allowed-days(monday | tuesday | wednesday | thursday | friday | saturday | sunday);
 allow-sources [ source-addresses ... ];
 allow-times [ times ... ];
 cli {
  prompt prompt;
  timestamp {
   format format;
  }
 }
 configuration-breadcrumbs;
 confirm-commands (regular expression or command 1)|(regular expression or command 2) {
   confirmation-message;
 }
 deny-commands (regular-expression1)|(regular-expression2)...;
 deny-commands-regexps (regular-expression1)|(regular-expression2)...;
 deny-configuration (regular-expression1)|(regular-expression2)...;
 deny-configuration-exact-match configuration;
 deny-configuration-regexps [“regular expression 1” “regular expression 2 ” ... ];
```

```
  deny-grpc-rpc-regexps [ deny-grpc-rpc-regexps ... ];
  deny-sources [ source-addresses ... ];
  deny-times [ times ... ];
  idle-timeout minutes;
  logical-system logical-system-name;
  login-alarms;
  login-script login-script;
  login-tip;
  no-scp-server;
  no-sftp-server;
  permissions [ permissions ... ];
  satelliteall;
  security-role(audit-administrator | crypto-administrator | ids-administrator | security-
administrator);
  tenant tenant-system-name;
  web-ui-hidden-menus web-ui-hidden-menus;
  web-ui-read-only-menus web-ui-read-only-menus;
 }
```

Hierarchy Level

```
[edit system login]
```

Description

Define a login class. All users who log in to the router or switch must be in a login class. Therefore, you must define a Junos OS login class for each user or type of user. You can define any number of login classes depending on the types of permissions the users need. You may not need to define any login classes; Junos OS has several predefined login classes, to suit a variety of needs. However, the predefined login classes cannot be modified. If you define a class with the same name as a predefined class, Junos OS appends `-local` to the login class name and creates a new login class. See Predefined System Login Classes for more information.

Options

class-name A name you choose for the login class.

access-end Specify the end time in *HH*:*MM* (24-hour) format, where *HH* represents the hours and *MM* represents the minutes.

> **NOTE**: Access start and end times that span across 12:00 AM starting on a specified day results in the user having access until the next day, even if the access day is not explicitly configured on the `allowed-days` statement.

access-start

Specify the start time in *HH*:*MM* (24-hour) format, where *HH* represents the hours and *MM* represents the minutes.

> **NOTE**: Access start and end times that span across 12:00 AM starting on a specified day results in the user having access until the next day, even if the access day is not explicitly configured on the `allowed-days` statement.

(allow-commands | allow-commands-regexps)

Specify one or more regular expressions to allow users in this class to issue operational mode commands. You use the `allow-commands` or the `allow-commands-regexps` statement to explicitly allow authorization for commands that would otherwise be denied by the access privilege levels for a login class.

For the `allow-commands` statement, each expression separated by a pipe (|) symbol must be a complete standalone expression, and must be enclosed in parentheses ( ). Do not use spaces between regular expressions separated with parentheses and connected with the pipe (|) symbol.

For the `allow-commands-regexps` statement, you configure a set of strings in which each string is a regular expression, enclosed in double quotes and separated with a space operator. Each string is evaluated against the full path of the command, which provides faster matching than the `allow-command` statement. You can also include values for variables in the regular expressions, which is not supported using the `allow-commands` statement.

The `deny-commands` or the `deny-commands-regexps` statement takes precedence if it is used in the same login class definition.

> **NOTE**: The `allow/deny-commands` and `allow/deny-commands-regexps` statements are mutually exclusive and cannot be configured together for a login class. At a given point in time, a login class can include either the `allow/deny-commands` statements, or the `allow/deny-commands-regexps` statements. If you have existing configurations using the `allow/deny-commands` statements, using the same configuration options with

the `allow/deny-commands-regexps` statements might not produce the same results, as the search and match methods differ in the two forms of these statements.

Authorizations can also be configured remotely by specifying Juniper Networks vendor-specific TACACS+ attributes in your authentication server's configuration. For a remote user, when the authorization parameters are configured both remotely and locally, authorization parameters configured remotely and locally are both considered together for authorization. For a local user, only the authorization parameters configured locally for the class are considered.

- **Default:** If you do not configure authorizations for operational mode commands using the `allow/deny-commands` or `allow/deny-commands-regexps` statements, users can edit only those commands for which they have access privileges set with the `permissions` statement.
- **Syntax:** *`regular-expression`*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Enter as many expressions as needed.

(allow-configuration | allow-configuration-regexps)

Specify one or more regular expressions to explicitly allow users in this class to access the specified levels in the configuration hierarchy even if the permissions set with the `permissions` statement do not grant such access.

For the`allow-configuration` statement, each expression separated by a pipe (|) symbol must be a complete standalone expression, and must be enclosed in parentheses ( ). Do not use spaces between regular expressions separated with parentheses and connected with the pipe (|) symbol.

For the `allow-configuration-regexps` statement, you configure a set of strings in which each string is a regular expression, enclosed in double quotes and separated with a space operator. Each string is evaluated against the full path of the command, which provides faster matching than the `allow/deny-configuration` statements. You can also include values for variables in the regular expressions, which is not supported using the `allow/deny-configuration` statements.

The `deny-configuration` or `deny-configuration-regexps` statement takes precedence if it is used in the same login class definition.

NOTE: The `allow/deny-configuration` and `allow/deny-configuration-regexps` statements are mutually exclusive and cannot be configured together

for a login class. At a given point in time, a login class can include either the `allow/deny-configuration` statements, or the `allow/deny-configuration-regexps` statements. If you have existing configurations using the `allow/deny-configuration` statements, using the same configuration options with the `allow/deny-configuration-regexps` statements might not produce the same results, as the search and match methods differ in the two forms of these statements.

- **Default:** If you omit the `allow-configuration/allow-configuration-regexps` statement and the `deny-configuration/deny-configuration-regexps` statement, users can edit only those commands for which they have access privileges through the `permissions` statement.
- **Syntax:** *`regular-expression`*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Enter as many expressions as needed.

allow-configuration-exact-match

This option uses complete hierarchy strings and wildcards in addition to regular expressions. Specify a hierarchy including `set`, `delete`, `active`, or `deactive` to explicitly allow users in this class to access the specified configuration.

The `allow-configuration-exact-match` option offers more granular control over access privileges than the `allow-configuration` option. For example if you use `allow-configuration-exact-match set interfaces <*>` to allow a login class access to the `set interfaces` configuration, it will not grant access to the `delete interfaces` configuration.

If you set the same configuration for both `allow-configuration-exact-match` and `deny-configuration-exact-match`, then the then the configuration access will be denied.

allow-grpc-rpc-regexps

Specify one or more regular expressions to explicitly allow users in this class to execute matching GRPC RPCs

For the `allow-grpc-rpc-regexps` statement, you configure a set of strings in which each string is a regular expression, enclosed in double quotes and separated with a space operator.

The `deny-grpc-rpc-regexps` statement takes precedence if it is used along with `allow-grpc-rpc-regexps` in the same login class definition.

- **Syntax:** *`regular-expression`*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Enter as many expressions as needed.

allow-hidden-commands

Allow all hidden commands to be run. If the no-hidden-commands statement is specified at the `[edit system]` hierarchy level, overrides that restriction for this login class. Hidden commands are Junos OS commands that are not published but could be run on a router. Hidden commands serve a specific purpose, but for most part are not expected to be used, and as such are not actively supported. The no-hidden-commands statement at the `[edit system]` hierarchy level allows you to block all hidden commands to all users except the root users.

- **Default:** Hidden commands are enabled by default.

allowed-days [*days of the week*]

Specify one or more days of the week when users in this class are allowed to log in.

- Values:
  - monday—Monday
  - tuesday—Tuesday
  - wednesday—Wednesday
  - thursday—Thursday
  - friday—Friday
  - saturday—Saturday
  - sunday—Sunday

cli

Set the CLI prompt specified for the login class. If a CLI prompt is also set at the [edit system login user cli] hierarchy level, the prompt set for the login user has precedence over the prompt set for the login class.

prompt *prompt* Specify the prompt string you want to see displayed in the CLI prompt.

configuration-breadcrumbs

Enable the configuration breadcrumbs view in the CLI to display the location in the configuration hierarchy. For an example of how to enable this view, see *Enabling Configuration Breadcrumbs* .

confirm-commands

Specify that confirmation for particular commands is explicitly required and, optionally, specify the wording of the message displayed at confirm time. You can specify the commands using a list of regular expressions or commands.

- **Syntax:** *message*

- **Default:** If you omit this option, then confirmation for commands is not required. If the optional message is not set, then the default "Do you want to continue?" message is displayed.

(deny-commands | deny-commands-regexps)

Specify one or more regular expressions to explicitly deny users in this class permission to issue operational mode commands, even though the permissions set with the `permissions` statement would allow it.

For the `deny-commands` statement, each expression separated by a pipe (|) symbol must be a complete standalone expression, and must be enclosed in parentheses ( ). Do not use spaces between regular expressions separated with parentheses and connected with the pipe (|) symbol.

For the `deny-commands-regexps` statement, you configure a set of strings in which each string is a regular expression, enclosed in double quotes and separated with a space operator. Each string is evaluated against the full path of the command, which provides faster matching than the `allow/deny-command` statements. You can also include values for variables in the regular expressions, which is not supported using the `allow/deny-commands` statements.

Expressions configured with the `deny-commands` or the `deny-commands-regexps` statement take precedence over expressions configured with `allow-commands/allow-commands-regexps` if the two statements are used in the same login class definition.

> **NOTE**: The `allow/deny-commands` and `allow/deny-commands-regexps` statements are mutually exclusive and cannot be configured together for a login class. At a given point in time, a login class can include either the `allow/deny-commands` statements, or the `allow/deny-commands-regexps` statements. If you have existing configurations using the `allow/deny-commands` statements, using the same configuration options with the `allow/deny-commands-regexps` statements might not produce the same results, as the search and match methods differ in the two forms of these statements.

Authorizations can also be configured remotely by specifying Juniper Networks vendor-specific TACACS+ attributes in your authentication server's configuration. For a remote user, when the authorization parameters are configured both remotely and locally, authorization parameters configured remotely and locally are both considered together for authorization. For a local user, only the authorization parameters configured locally for the class are considered.

- **Default:** If you do not configure authorizations for operational mode commands using `allow/deny-commands` or `allow/deny-commands-regexps`, users can edit only those commands for which they have access privileges set with the `permissions` statement.

- **Syntax:** *`regular-expression`*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Enter as many expressions as needed.

(deny-configuration | deny-configuration-regexps)

Specify one or more regular expressions to explicitly deny users in this class access to the specified levels in the configuration hierarchy even if the permissions set with the `permissions` statement grant such access. Note that the user cannot view a particular hierarchy if configuration access is denied for that hierarchy.

For the `deny-configuration` statement, each expression separated by a pipe (|) symbol must be a complete standalone expression, and must be enclosed in parentheses ( ). Do not use spaces between regular expressions separated with parentheses and connected with the pipe (|) symbol.

For the `deny-configuration-regexps` statement, you configure a set of strings in which each string is a regular expression, enclosed in double quotes and separated with a space operator. Each string is evaluated against the full path of the command, which provides faster matching than the `allow/deny-configuration` statements. You can also include values for variables in the regular expressions, which is not supported using the `allow/deny-configuration` statements.

Expressions configured with `deny-configuration/deny-configuration-regexps` take precedence over expressions configured with `allow-configuration/allow-configuration-regexps` if the two statements are used in the same login class definition.

> **NOTE**: The `allow/deny-configuration` and `allow/deny-configuration-regexps` statements are mutually exclusive and cannot be configured together for a login class. At a given point in time, a login class can include either the `allow/deny-configuration` statements, or the `allow/deny-configuration-regexps` statements. If you have existing configurations using the `allow/deny-configuration` statements, using the same configuration options with the `allow/deny-configuration-regexps` statements might not produce the same results, as the search and match methods differ in the two forms of these statements.

- **Default:** If you omit the `deny-configuration/deny-configuration-regexps` statement and the `allow-configuration/allow-configuration-regexps` statement, users can edit those levels in the configuration hierarchy for which they have access privileges through the `permissions` statement.

- **Syntax:** *`regular-expression`*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Enter as many expressions as needed.

deny-configuration-exact-match

This option uses complete hierarchy strings and wildcards in addition to regular expressions. Specify a hierarchy including `set`, `delete`, `active`, or `deactive` to explicitly deny users in this class to access the specified configuration.

The `deny-configuration-exact-match` option offers more granular control over denying privileges than the `deny-configuration` option. For example if you use `deny-configuration-exact-match set interfaces <*>` to deny a login class access to the `set interfaces` configuration, it will not grant access to the `set interfaces` configuration.

If you set the same configuration for both `allow-configuration-exact-match` and `deny-configuration-exact-match`, then the then the configuration access will be denied.

Configurations that are denied with `deny-configuration-exact-match` will return a `Permission denied` error if configured. Denied configurations will still appear as autocomplete entries when using the Tab or Spacebar keys in the Junos CLI.

deny-grpc-rpc-regexps

Specify one or more regular expressions to prevent users in this class from running matching GRPC RPCs

For the `deny-grpc-rpc-regexps` statement, you configure a set of strings in which each string is a regular expression, enclosed in double quotes and separated with a space operator.

The `deny-grpc-rpc-regexps` statement takes precedence over the `allow-grpc-rpc-regexps` if it is used in the same login class definition.

- **Syntax:** *`regular-expression`*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Enter as many expressions as needed.

idle-timeout

For a login class, configure the maximum time in minutes that a session can be idle before the session times out and the user is logged out of the device. The session times out after remaining at the CLI operational mode prompt for the specified time.

> **NOTE**: After the user logs in to a device from a shell prompt such as csh, if the user starts another program to run in the foreground of the CLI, the idle-timer control is stopped from being computed. The calculation of the idle time of the CLI session is restarted only after the foreground process exits and the control is returned to the shell prompt. When the restart of the idle-timer control occurs, if no interaction from the user occurs on the shell, the user is automatically logged out after the time set on this statement.

- **Default:** If you omit this statement, a user is never forced off the system after extended idle times.
- **Syntax:** *minutes*—Maximum time in minutes that a session can be idle before a user is logged out.
- **Range:** Range: 0 through 4294967295 minutes

> **NOTE**: The idle-timeout feature is disabled if the value of *minutes* is set to 0.

login-alarms

Display system alarms when a user with `admin` permissions logs in to the device. For more information about configuring this statement, see Configuring System Alarms to Appear Automatically Upon Login.

login-script

Run the specified op script when a user belonging to the class logs in to the CLI. The script must be enabled in the configuration.

logical-system

Assign the users in this login class to a logical system. If you specify a logical system, you can't include the satellite configuration statement in the configuration for this login class.

login-tip

Display CLI tips when logging in.

- **Default:** If this statement is not configured, CLI tips are not displayed.

no-hidden-commands

Deny all hidden commands, except for those specified, for users in this login class. Each command listed as an exception must be enclosed in quotation marks.

- **Default:** Hidden commands are enabled by default.
- **Syntax:** except *["command 1" "command 2"...]*

no-scp-server Disable incoming SCP connections for this login class.

no-sftp-server Disable incoming SFTP connections for this login class.

permissions Specify login access privileges for the login class.

- **Syntax:** *permissions*—One or more permission flags, which together specify the access privileges for the login class. Permission flags are not cumulative, so for each class, you must list all the permission flags needed, including `view` to display information and `configure` to enter configuration mode. For a list of permission flags, see Login Class Permission Flags.

satellite Specify access to Junos Fusion satellite devices for the login class. All users assigned to the login class are satellite users. If you include this statement, you can't include the logical-system configuration statement in the configuration for this login class.

- Values:
  - all—Specify all Junos Fusion satellite devices.

security-role Specify one or more Common Criteria (ISO/IEC 15408) security roles for the login class.

- Values:

audit-administrator	Specify which users are responsible for the regular review of specific target of evaluation (TOE) audit data and audit trail deletion. Audit administrators can also invoke the non-cryptographic self-test.
crypto-administrator	Specify which users are responsible for the configuration and maintenance of cryptographic elements related to the establishment of secure connections to and from the TOE audit data.
ids-administrator	Specify which users can act as intrusion detection service (IDS) administrators, who are responsible for all of the activities regarding identity and access management of the organization's employees.
security-administrator	Specify which users are responsible for ensuring that the organization's security policy is in place.

tenant Assign the users in this class to a tenant system. Tenant systems are used when you need to separate departments, organizations, or customers and each of them can be limited to one virtual router. The main difference between a logical system and a tenant system is that a logical system supports advanced routing functionality using multiple routing instances. In comparison, a tenant system supports only one routing instance, but supports the deployment of significantly more tenants per system.

web-ui-hidden-menus Enable hidden menus in the J-Web interface.

web-ui-read-only-menus Enable read-only menus in the J-Web interface.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

The `class`, `allow-commands`, `deny-commands`, `allow-configuration`, `deny-configuration`, `idle-timeout`, `login-alarms`, `login-tip`, and `permissions` statements were introduced before Junos OS Release 7.4.

All of the previously mentioned statements were introduced in Junos OS Release 9.0 for the EX Series.

The `login-script` statement was introduced in Junos OS Release 9.5.

The `access-end`, `access-start`, and `allowed-days` statements were introduced in Junos OS Release 10.1.

All of the previously mentioned statements were introduced in Junos OS Release 11.1 for the QFX Series.

All of the previously mentioned statements were introduced in Junos OS Release 11.2 for the SRX Series.

The `allow-configuration-regexps`, `deny-configuration-regexps`, and `security-role` statements were introduced in Junos OS Release 11.2.

The `configuration-breadcrumbs` statement was introduced in Junos OS Release 12.2.

All of the previously mentioned statements were introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

All of the previously mentioned statements were introduced in Junos OS Release 15.1X49-D70 for the vSRX Virtual Firewall, SRX4100, SRX4200 and SRX1500 devices.

All of the previously mentioned statements were introduced in Junos OS Release 16.1 for the MX Series and PTX Series.

The `allow-hidden-commands`, `confirm-commands`, `no-hidden-commands`, and `satellite` statements were introduced in Junos OS Release 16.1.

The `cli` statement was introduced in Junos OS Release 17.3.

The `allow-commands-regexps` and `deny-commands-regexps` statements were introduced in Junos OS Release 18.1.

The `tenant` statement was introduced in Junos OS 18.4.

The `no-scp-server` and `no-sftp-server` statements were introduced in Junos OS Release 19.2.

The `web-ui-hidden-menus` and `web-ui-read-only-menus` statements were introduced in Junos OS 21.3 for the SRX platforms.

RELATED DOCUMENTATION

- Defining Junos OS Login Classes
- Configuring Time-Based User Access
- Understanding Junos OS Access Privilege Levels
- Example: Configuring User Permissions with Access Privileges for Operational Mode Commands
- Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands, Configuration Statements, and Hierarchies
- Example: Configuring User Permissions with Access Privileges for Configuration Statements and Hierarchies
- Example: Configuring User Permissions with Access Privilege Levels
- Configuring System Alarms to Appear Automatically Upon Login
- *Executing an Op Script on the Local Device*
- Understanding Administrative Roles
- Example: Configuring Administrative Roles
- Tenant Systems Overview
- user (Access) | **30547**

class (Forwarding Class Sets)

IN THIS SECTION

- Syntax | 3361
- Hierarchy Level | 3361
- Description | 3361
- Options | 3361
- Required Privilege Level | 3361
- Release Information | 3362

Syntax

```
class class-name;
```

Hierarchy Level

```
[edit class-of-service forwarding-class-sets forwarding-class-set-name]
```

Description

Group forwarding classes into sets of forwarding classes (priority groups). You can group some or all of the configured forwarding classes into up to three unicast forwarding class sets and one multidestination forwarding class set.

Options

class-name —Name of the forwarding class.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Forwarding Class Sets

Understanding CoS Forwarding Class Sets (Priority Groups)

class (Forwarding classes)

IN THIS SECTION

- Syntax | 3362
- Hierarchy Level | 3362
- Description | 3363
- Options | 3363
- Required Privilege Level | 3363
- Release Information | 3363

Syntax

```
class class-name queue-num queue-number priority ( high | low );
```

Hierarchy Level

```
[edit class-of-service forwarding-classes]
```

Description

Configure up to 16 forwarding classes with multiple forwarding classes mapped to single queues. If you want to configure up to eight forwarding classes with one-to-one mapping to output queues, use the `queue` statement instead of the `class` statement at the `[edit class-of-service forwarding-classes]` hierarchy level.

Options

`class-name`	Name of forwarding class.
`priority (high | low)`	(Optional) Fabric priority. • **Default:** `low`
`queue-num` *`queue-number`*	Output queue number. • **Range:** 0 through 7

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches

Defining CoS Forwarding Classes (CLI Procedure)

Defining CoS Forwarding Classes (J-Web Procedure)

class (Forwarding Classes)

IN THIS SECTION

- Syntax for QFX Series Routers | 3364
- Syntax for MX Series routers and EX Series switches | 3364
- Hierarchy Level | 3365
- Description | 3365
- Options | 3365
- Required Privilege Level | 3365
- Release Information | 3365

Syntax for QFX Series Routers

```
class {
    class-name {
        pfc-priority pfc-priority;
        queue-num queue-number <no-loss>;
    }
}
```

Syntax for MX Series routers and EX Series switches

```
class {
    class-name {
        queue-num queue-number ;
        priority (high | low) ;
    }
}
```

Hierarchy Level

```
[edit class-of-service forwarding-classes]
```

Description

Specify the output transmission queue to which to map all input from an associated forwarding class.

This statement enables you to configure up to 16 forwarding classes with multiple forwarding classes mapped to single queues. If you want to configure up to eight forwarding classes with one-to-one mapping to output queues, use the "queue" on page 13944 statement instead of the `class` statement at the `[edit class-of-service forwarding-classes]` hierarchy level.

Map one or more forwarding classes to a single queue. Also, when configuring DSCP-based PFC, map a forwarding class to a PFC priority value to use in pause frames when traffic on a DSCP value becomes congested (see *Configuring DSCP-based PFC for Layer 3 Untagged Traffic* for details).

You can map unicast forwarding classes to a unicast queue (0 through 7) and multidestination forwarding classes to a multicast queue (8 through 11). The queue to which you map a forwarding class determines if the forwarding class is a unicast or multicast forwarding class.

Options

class-name—Name of the forwarding class.

queue-number—Output queue number.

- **Range:** 0 through 7.

The remaining statements are explained separately. See CLI Explorer for details.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

No-loss option introduced in Junos OS Release 12.3 for the QFX Series.

PFC-priority statement introduced in Junos OS Release 17.4R1 for the QFX Series.

RELATED DOCUMENTATION

Example: Configuring Forwarding Classes

Understanding CoS Forwarding Classes

Understanding CoS Forwarding Classes

Configuring a Custom Forwarding Class for Each Queue

queue (Global Queues) | **13944**

class-of-service (Dynamic Profiles)

IN THIS SECTION

- Syntax | **3366**
- Hierarchy Level | **3368**
- Description | **3368**
- Default | **3368**
- Required Privilege Level | **3368**
- Release Information | **3368**

Syntax

```
class-of-service {
    dynamic-class-of-service-options {
        vendor-specific-tags tag;
    }
    interfaces {
        interface-name ;
        }
        unit logical-unit-number {
            classifiers {
                type (classifier-name | default);
            }
            output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
            report-ingress-shaping-rate bps;
```

```
            rewrite-rules {
                dscp (rewrite-name | default);
                dscp-ipv6 (rewrite-name | default);
                ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                inet-precedence (rewrite-name | default);
                }
            }
        }
    }
    scheduler-maps {
        map-name {
            forwarding-class class-name scheduler scheduler-name;
        }
    }
    schedulers {
        (scheduler-name) {
            buffer-size (seconds | percent percentage | remainder | temporal microseconds);
            drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol
(any | non-tcp | tcp) drop-profile profile-name;
            excess-priority (low | high | $junos-cos-scheduler-excess-priority);
            excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
            overhead-accounting (shaping-mode) <bytes (byte-value>;
            priority priority-level;
            shaping-rate (rate | predefined-variable);
            transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
        }
    }
    traffic-control-profiles profile-name {
        adjust-minimum rate;
        delay-buffer-rate (percent percentage | rate);
        excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
        excess-rate-high (percent percentage | proportion value);
        excess-rate-low (percent percentage | proportion value);
        guaranteed-rate (percent percentage | rate) <burst-size  bytes>;
        max-burst-size cells;
        overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
        peak-rate rate;
        scheduler-map map-name;
        shaping-rate (percent percentage | rate | predefined-variable) <burst-size  bytes>;
        shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
        shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
        shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
        shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
```

```
        shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
        shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
        shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
        shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
        shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
        sustained-rate rate;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Configure Junos OS CoS features in a dynamic client profile or a dynamic service profile.

Default

If you do not configure any CoS features, all packets are transmitted from output transmission queue 0.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

> *Guidelines for Configuring Dynamic CoS for Subscriber Access*
>
> Configuring Static Hierarchical Scheduling in a Dynamic Profile

class-of-service (Protocols MPLS)

IN THIS SECTION

- Syntax | 3369
- Hierarchy Level | 3369
- Description | 3369
- Options | 3370
- Required Privilege Level | 3370
- Release Information | 3370

Syntax

```
class-of-service cos-value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress]
```

Description

Class-of-service (CoS) value given to all packets in the LSP.

The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.

Options

cos-value—CoS value. A higher value typically corresponds to a higher level of service.

- **Range:** 0 through 7
- **Default:** If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Class of Service for MPLS LSPs

Configuring the Ingress Router for Static LSPs

Configuring the Intermediate (Transit) and Egress Routers for Static LSPs

class-of-service (Protocols RSVP)

IN THIS SECTION

- Syntax | 3371
- Hierarchy Level | 3371
- Description | 3371
- Options | 3371
- Required Privilege Level | 3371
- Release Information | 3371

Syntax

```
class-of-service cos-value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection bypass bypass-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Description

Class-of-service (CoS) value given to all packets in the bypass LSP. You can specify a single CoS value for all the bypass LSPs traversing an interface. You can also configure CoS values for specific bypass LSPs traversing an interface.

The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.

Options

cos-value—CoS value. A higher value typically corresponds to a higher level of service.

- **Range:** 0 through 7
- **Default:** If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Class of Service for Bypass LSPs

class-of-service (saegw system)

IN THIS SECTION

- Syntax | 3372
- Hierarchy Level | 3372
- Description | 3373
- Required Privilege Level | 3373
- Release Information | 3373

Syntax

```
class-of-service {
      reserved-bandwidth {
            express percentage;
            gbr percentage;
      }
      downlink-dscp-to-forwarding-class {
          classifier classifier-name;
          classifier classifier-name {
              routing-instance instance-name;
          }
      }
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw system]
```

Description

Configure class-of-service features on Junos Multi-Access User Plane.

See CLI Explorer for the remaining statements.

Required Privilege Level

system

Release Information

Juniper Networks introduced this statement in Junos OS Release 22.2R1.

class-usage-profile

IN THIS SECTION

- Syntax | 3373
- Hierarchy Level | 3374
- Description | 3374
- Options | 3374
- Required Privilege Level | 3374
- Release Information | 3374

Syntax

```
class-usage-profile profile-name {
    file filename;
    interval minutes;
    start-time time;
    source-classes {
        source-class-name;
    }
    destination-classes {
```

```
        destination-class-name;
    }
}
```

Hierarchy Level

```
[edit accounting-options]
```

Description

Create a class usage profile, which is used to log class usage statistics to a file in the `/var/log` directory. The class usage profile logs class usage statistics for the configured source classes on every interface that has `destination-class-usage` configured.

For information about configuring source classes, see the *Junos Routing Protocols Configuration Guide*. For information about configuring source class usage, see the *Junos Network Management Configuration Guide*.

Options

profile-name Name of the destination class profile.

start-time ***time*** (Optional) Specify the start time in *yyyy-mm-dd.hh:mm* local time format for the accounting profile. Sometimes, if the time at which the records are written to the accounting file is beyond the transfer interval window of the file, drifting of accounting records occurs. To prevent drifting and to have predictable timestamps of when accounting data is written to an accounting file, we can specify the start time option for the accounting profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Class Usage Profiles

classification-key (DHCP Local Server)

IN THIS SECTION

- Syntax | 3375
- Hierarchy Level | 3375
- Description | 3376
- Options | 3376
- Required Privilege Level | 3376
- Release Information | 3376

Syntax

```
classification-key {
    circuit-id circuit-id;
    mac-address mac-address;
    remote-id remote-id;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server dual-stack-group
dual-stack-group-name ],
[edit logical-systems name system services dhcp-local-server dual-stack-group dual-stack-group-
name ],
[edit routing-instances name system services dhcp-local-server dual-stack-group dual-stack-group-
name ],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name ]
```

Description

Different mechanisms to identify a single household.

Options

circuit-id	Circuit-id as key.
mac-address	MAC address of client.
remote-id	Remote-id as key.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

Single-Session DHCP Dual-Stack Overview

classification-key (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 3377
- Hierarchy Level | 3377
- Description | 3378
- Options | 3378
- Required Privilege Level | 3378

Release Information | 3379

Syntax

```
classification-key {
    circuit-id circuit-id;
    mac-address mac-address;
    remote-id remote-id;
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-
stack-group-name ],
[edit bridge-domains name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name
dual-stack-group dual-stack-group-name ],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dual-stack-group
dual-stack-group-name ],
[edit logical-systems name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-
stack-group-name ],
[edit logical-systems name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dhcpv6 group name
dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dual-stack-group
dual-stack-group-name ],
[edit logical-systems name routing-instances name system services dhcp-local-server dual-stack-
group dual-stack-group-name ],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dhcpv6
group name ddual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dual-
```

```
stack-group dual-stack-group-name ],
[edit logical-systems name system services dhcp-local-server dual-stack-group dual-stack-group-name ],
[edit logical-systems name vlans name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit logical-systems name vlans name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit routing-instances name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group ],
[edit routing-instances name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit routing-instances name system services dhcp-local-server dual-stack-group dual-stack-group-name ],
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit routing-instances name vlans name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit system services dhcp dhcp-local-server dual-stack-group dual-stack-group-name ],
[edit vlans name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name ],
[edit vlans name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ]
```

Description

Different mechanisms to identify a single household.

Options

circuit-id	Circuit-id as key
mac-address	MAC address of client
remote-id	Remote-id as key

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

Single-Session DHCP Dual-Stack Overview

classification-override

IN THIS SECTION

- Syntax | 3379
- Hierarchy Level | 3379
- Description | 3380
- Required Privilege Level | 3380
- Release Information | 3380

Syntax

```
classification-override {
    forwarding-class class-name;
}
```

Hierarchy Level

```
[edit class-of-service forwarding-policy class class-name]
```

Description

Override the incoming packet classification, assigning all packets sent to a destination prefix to the same output transmission queue.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Overriding the Input Classification

Junos OS Routing Protocols Library for Routing Devices

classifier

IN THIS SECTION

- Syntax | 3380
- Hierarchy Level | 3381
- Description | 3381
- Required Privilege Level | 3381
- Release Information | 3381

Syntax

```
classifier {
    per-unit-scheduler {
```

```
        forwarding-class class-name {
            loss-priority (high | low);
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile
output-priority-map]
```

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the classifier for the output priority map to be applied to outgoing frames on this interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

input-priority-map | 8716

classifiers (VRF)

IN THIS SECTION

- Syntax | 3382
- Hierarchy Level | 3382
- Description | 3382
- Default | 3382
- Options | 3383
- Required Privilege Level | 3383
- Release Information | 3383

Syntax

```
classifiers {
    exp (classifier-name | default);
}
```

Hierarchy Level

```
[edit class-of-service routing-instances routing-instance-name]
```

Description

For routing instances with VRF table labels enabled, apply a custom MPLS EXP classifier to the routing instance. You can apply the default
MPLS EXP classifier or one that is previously defined.

Default

If you do not include this statement, the default MPLS EXP classifier is applied to the routing instance.

Options

classifier-name—Name of the behavior aggregate MPLS EXP classifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Custom MPLS EXP Classifiers to Routing Instances in Layer 3 VPNs

Junos OS Network Interfaces Library for Routing Devices

classifiers (Class of Service)

IN THIS SECTION

- Syntax (BA Classifiers) | 3384
- Syntax (Multidestination BA Classifiers) | 3384
- Syntax (Interface Classifier Association: DSCP, DSCP IPv6, IEEE) | 3384
- Syntax (Global EXP Interface Classifier Association with Interfaces) | 3385
- Hierarchy Level (BA Classifiers) | 3385
- Hierarchy Level (Multidestination BA Classifiers) | 3385
- Hierarchy Level (Interface Classifier Association: DSCP, DSCP IPv6, IEEE) | 3385
- Hierarchy Level (Global EXP Classifier) | 3385
- Description | 3385
- Options | 3385
- Required Privilege Level | 3385

Release Information | 3386

Syntax (BA Classifiers)

```
classifiers {
    (dscp | dscp-ipv6 | ieee-802.1) classifier-name {
        import (classifier-name | default);
        forwarding-class class-name {
            loss-priority level {
               code-points [ aliases ] [ bit-patterns ];
            }
        }
    }
}
```

Syntax (Multidestination BA Classifiers)

Multidestination BA Classifiers

```
classifiers {
    (dscp | ieee-802.1) classifier-name;
}
```

Syntax (Interface Classifier Association: DSCP, DSCP IPv6, IEEE)

Interface Classifier Association

```
classifiers {
    (dscp | dscp-ipv6 | ieee-802.1) (default | classifier-name);
}
```

Syntax (Global EXP Interface Classifier Association with Interfaces)

```
classifiers {
    exp classifier-name;
}
```

Hierarchy Level (BA Classifiers)

```
[edit class-of-service],
```

Hierarchy Level (Multidestination BA Classifiers)

```
[edit class-of-service multi-destination],
```

Hierarchy Level (Interface Classifier Association: DSCP, DSCP IPv6, IEEE)

```
[edit class-of-service interfaces interface-name  unit logical-unit-number]
```

Hierarchy Level (Global EXP Classifier)

```
[edit class-of-service system-defaults]
```

Description

Define a unicast or multidestination CoS behavior aggregate (BA) classifier.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

Configuring a Global MPLS EXP Classifier

Example: Configuring Unicast Classifiers

Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers

Understanding CoS Classifiers

Understanding CoS Classifiers

Understanding CoS MPLS EXP Classifiers and Rewrite Rules

classifiers (Definition)

IN THIS SECTION

- Syntax | 3386
- Hierarchy Level | 3387
- Description | 3387
- Options | 3387
- Required Privilege Level | 3387
- Release Information | 3388

Syntax

```
classifiers {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) classifier-name {
        forwarding-class forwarding-class-name {
```

```
        loss-priority (high  | low  | medium-high  | medium-low) {
            code-point alias-or-bit-string ;
        }
        import (default | user-defined;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Configure a user-defined behavior aggregate (BA) classifier.

Options

- *classifier-name*—User-defined name for the classifier.
- import (default | *user-defined*)—Specify the template to use to map any code points not explicitly mapped in this configuration. For example, if the classifier is of type `dscp` and you specify `import default`, code points you do not map in your configuration will use the predefined DSCP default mapping; if you specify `import mymap`, for example, code points not mapped in the forwarding-class configuration would use the mappings in a user-defined classifier named `mymap`.
- forwarding-class *class-name*—Specify the name of the forwarding class. You can use the default forwarding class names or define new ones.
- loss-priority *level*—Specify a loss priority for this forwarding class: `high`, `low`, `medium-high`, `medium-low`.
- code-points (*alias* | *bits*)—Specify a code-point alias or the code points that map to this forwarding class.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2

RELATED DOCUMENTATION

Understanding Interfaces

classifiers (Definition)

IN THIS SECTION

- Syntax | 3388
- Hierarchy Level | 3389
- Description | 3389
- Options | 3389
- Required Privilege Level | 3389
- Release Information | 3389

Syntax

```
classifiers {
    type classifier-name {
        import (classifier-name | default);
        forwarding-class class-name {
            loss-priority level code-points [ aliases ] [ bit-patterns ];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Define a CoS behavior aggregate (BA) classifier for classifying packets. You can associate the classifier with a forwarding class or code-point mapping, and import a default classifier or one that is previously defined.

Options

classifier-name—Name of the aggregate behavior classifier.

type—Traffic type: `dscp`, `dscp-ipv6`, `exp`, `ieee-802.1`, `ieee-802.1ad`, `inet-precedence`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`ieee-802.1ad` option introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic

classifiers (Dynamic CoS Application)

IN THIS SECTION

- Syntax | 3390
- Hierarchy Level | 3390
- Description | 3390
- Options | 3390
- Required Privilege Level | 3391
- Release Information | 3391

Syntax

```
classifiers {
    dscp (classifier-name | default);
    dscp-ipv6 (classifier-name | default);
    ieee-802.1 (classifier-name | default) vlan-tag (inner | outer)
    inet-precedence (classifier-name | default);
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Apply a CoS behavior aggregate classifier to a dynamic interface. You can apply a default classifier or one that is previously defined.

Options

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying a Classifier to a Subscriber Interface in a Dynamic Profile

classifiers (Definition)

classifiers (Interfaces)

IN THIS SECTION

- Syntax | 3391
- Hierarchy Level | 3392
- Description | 3392
- Options | 3392
- Required Privilege Level | 3392
- Release Information | 3393

Syntax

```
classifiers {
    type (classifier-name | default) family (mpls | inet);
    no-default;
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name ]
```

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Apply a CoS aggregate behavior classifier to a logical or physical interface. You can apply a default classifier or one that is previously defined.

> **NOTE**: On MX Series routers, you must apply the classifier to a logical interface. On ACX Series routers, you must apply the classifier to a physical interface.

If you do not explicitly apply a classifier configuration to the interface, the default classifier is applied to the interface. Apply the `no-default` option to disable the application of any default classifier to the routing instance.

> **NOTE**: On supported devices, you can apply the `no-default` option to a flexible tunnel interface (`ftin`) to enable the device to classify traffic exiting the tunnel based on the classifier assigned to the underlay (ingress) interface of the tunnel's egress device.

Options

`classifier-name` Name of the aggregate behavior classifier.

`type` Traffic type.

- **Values:** `dscp`, `dscp-ipv6`, `exp`, `ieee-802.1`, `inet-precedence`

> **NOTE**: You can only specify a family for the `dscp` and `dscp-ipv6` types.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`no-default` option added in Junos OS Release 16.1.

RELATED DOCUMENTATION

Default DSCP and DSCP IPv6 Classifiers

Applying Behavior Aggregate Classifiers to Logical Interfaces

classifiers (Physical Interface)

IN THIS SECTION

- Syntax | 3393
- Hierarchy Level | 3394
- Description | 3394
- Options | 3394
- Required Privilege Level | 3394
- Release Information | 3394

Syntax

```
classifiers {
    type (classifier-name | default) ;
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name ]
```

Description

Apply a CoS aggregate behavior classifier to a physical interface. You can apply a default classifier or one that is previously defined.

> **NOTE**: You cannot apply a classifier to a physical interface on MX Series routers. On MX Series routers, you must apply the classifier to a logical interface.

Options

classifier-name Name of the aggregate behavior classifier.

type Traffic type.

- **Values:** `dscp`, `ieee-802.1`, and `inet-precedence`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

dscp (CoS Interfaces) | 5999

inet-precedence (Classifier on Physical Interface) | 8580

ieee-802.1 (Classifier on Physical Interface) | 8348

classifiers (Interfaces) | 3391

classifiers (Routing Instance)

IN THIS SECTION

- Syntax | 3395
- Hierarchy Level | 3395
- Description | 3395
- Options | 3396
- Required Privilege Level | 3396
- Release Information | 3396

Syntax

```
classifiers {
    exp (classifier-name | default);
    dscp (classifier-name | default);
    dscp-ipv6 (classifier-name | default);
    no-default;
}
```

Hierarchy Level

```
[edit class-of-service routing-instances routing-instance-name]
```

Description

For routing instances with VRF table labels enabled, apply a custom MPLS EXP, DSCP, or IEEE802.1 classifier to the routing instance. You can apply the default classifier or one that is previously defined.

If you do not explicitly apply a classifier configuration to the routing instance, the default classifier is applied to the routing instance. Apply the `no-default` option to disable the application of any default classifier to the routing instance.

Options

classifier-name—Name of the behavior aggregate MPLS EXP, DSCP, or IEEE802.1 classifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`dscp` and `dscp-ipv6` support introduced in Junos OS Release 9.6.

`no-default` option added in Junos OS Release 16.1.

Support for PTX Series routers introduced in Junos OS Evolved Release 22.4R1.

Support for SRX Series firewalls introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Applying MPLS EXP Classifiers to Routing Instances

Applying Behavior Aggregate Classifiers to Logical Interfaces

clat-ipv6-prefix-length

IN THIS SECTION

- Syntax | 3397
- Hierarchy Level | 3397
- Description | 3397
- Options | 3397
- Required Privilege Level | 3397
- Release Information | 3397

Syntax

```
clat-ipv6-prefix-length (32 | 40 | 48 | 56 | 64 | 96);
```

Hierarchy Level

```
[edit services nat rule name term name then translated]
```

Description

Specify the ipv6 prefix length for CLAT source address. Once you configure this command, source-address and source-prefix are no more mandatory configuration. It allows the NAT rules to accept the traffic from different CLAT prefix and apply XLAT464 based on destination-address of the traffic. This configuration supports overlapping ipv4 address inside CLAT from the same client. With the help of this configuration, the administrator can choose to assign same or different ipv4 address inside CLAT prefix.

Options

32	The ipv6 prefix length of 32.
40	The ipv6 prefix length of 40.
48	The ipv6 prefix length of 48.
56	The ipv6 prefix length of 56.
64	The ipv6 prefix length of 64.
96	The ipv6 prefix length of 96.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

464XLAT Overview

Configuring 464XLAT Provider-Side Translater for IPv4 Connectivity Across IPv6-Only Network

clat-prefix

IN THIS SECTION

- Syntax | 3398
- Hierarchy Level | 3398
- Description | 3398
- Required Privilege Level | 3398
- Release Information | 3399

Syntax

```
clat-prefix clat-prefix;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated]
```

Description

Specify the IPv6 prefix that the customer-side translator (CLAT) uses when performing 464XLAT translation to IPv6 (the CLAT is not a Juniper Networks product). The Provider-Side Translater (PLAT) on the MX Series uses the CLAT IPv6 prefix to translate the IPv6 packet back to IPv4.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

464XLAT Overview

Configuring 464XLAT Provider-Side Translater for IPv4 Connectivity Across IPv6-Only Network

clat-prefix (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 3399
- Hierarchy Level | 3399
- Description | 3400
- Required Privilege Level | 3400
- Release Information | 3400

Syntax

```
clat-prefix clat-prefix;
```

Hierarchy Level

```
[edit services nat source rule-set rule-set rule rule-name then source-nat]
```

Description

Specify the customer-side translator (CLAT) IPv6 source prefix, which is used for 464XLAT.

464XLAT lets an IPv4 client with a private IP address connect to an IPv4 host over an IPv6 network. The CLAT translates IPv4 source addresses to IPv6 by embedding the IPv4 source address in this IPv6 source prefix. The CLAT then sends the packets over an IPv6 network to the MX Series router, which acts as a provider-side translator (PLAT). The MX translates the embedded IPv4 private IP address to a public IPv4 address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

cleanup-interval (Accounting Options)

IN THIS SECTION

- Syntax | 3400
- Hierarchy Level | 3401
- Description | 3401
- Options | 3401
- Required Privilege Level | 3401
- Release Information | 3401

Syntax

```
cleanup-interval days;
```

Hierarchy Level

```
[edit accounting-options]
```

Description

Configure the interval to delete files from the local backup directory.

Options

days Number of days after which accounting-options files are to be deleted from the backup directory.

- **Range:** 1 through 31 days
- **Default:** 1 day

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Accounting-Data Log Files

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Flat-File Accounting Overview

clear

IN THIS SECTION

- Syntax | 3402
- Hierarchy Level | 3402
- Description | 3402
- Required Privilege Level | 3402

Syntax

```
request protection-group ethernet-aps clear md <md> ma <ma>
```

Hierarchy Level

```
[edit protocols protection-group ethernet-aps]
```

Description

Clears the lockout, force switch, manual switch, exercise, and wait-to-restore (WTR) states.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Ethernet Automatic Protection Switching Overview

clear access-security router-advertisement statistics

IN THIS SECTION

- Syntax | 3403
- Description | 3403
- Options | 3403
- Required Privilege Level | 3404
- Output Fields | 3404
- Release Information | 3404

Syntax

```
clear access-security router-advertisement statistics (fail | success) (all | interface interface-name |
vlan vlan-name)
```

Description

Clear the IPv6 Router Advertisement (RA) guard entries for received RA messages. IF RA guard is enabled on a switch, the switch examines incoming RA messages and filters them on the basis of a predefined set of criteria. If the switch validates the sender of the RA message as a legitimate IPv6 router, it forwards the RA message to its destination; otherwise, the RA message is dropped.

Options

all	Clear the RA guard entries on all VLANs.
fail	Clear RA guard entries for RA messages that were discarded.
interface *interface-name*	Clear the RA guard entries for the specified interface.
success	Clear the RA guard entries for RA messages that were accepted.
vlan *vlan-name*	Clear the RA guard entries for the specified VLAN.

Required Privilege Level

view

Output Fields

This command generates no output.

Release Information

Command introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

show access-security router-advertisement statistics | 17001

clear access-security slaac-snooping binding

IN THIS SECTION

- Syntax | 3404
- Description | 3405
- Options | 3405
- Required Privilege Level | 3405
- Output Fields | 3405
- Sample Output | 3405
- Release Information | 3405

Syntax

```
clear access-security slaac-snooping binding
<interface (interface-name | all)>
<vlan vlan-name>
```

```
<vlan vlan-name routing-instance routing-instance-name>
<vlan vlan-name logical-system logical-system-name>
```

Description

Clear the SLAAC snooping database information.

Options

interface (*interface-name* | all)	Clear SLAAC snooping information for the specified interface or all interfaces.
vlan (*vlan-name*)	Clear SLAAC snooping information for the specified VLAN.
routing-instance(*routing-instance-name*)	Clear SLAAC snooping information for the specified routing instance.
logical-system(*logical-system-name*)	Clear SLAAC snooping information for the specified logical system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear access-security slaac-snooping binding

```
user@switch> clear access-security slaac-snooping binding
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show access-security slaac-snooping binding | 17003

clear access-security slaac-snooping statistics

IN THIS SECTION

- Syntax | 3406
- Description | 3406
- Options | 3406
- Required Privilege Level | 3406
- Output Fields | 3407
- Sample Output | 3407
- Release Information | 3407

Syntax

```
clear access-security slaac-snooping statistics
<interface (interface-name | all)>
```

Description

Clear the SLAAC snooping statistics.

Options

`interface (interface-name | all)`	Clear SLAAC snooping statistics for the specified interface or all interfaces.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear access-security slaac-snooping statistics

```
user@switch> clear access-security slaac-snooping statistics
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show access-security slaac-snooping statistics | 17007

clear accounting server statistics archival-transfer

IN THIS SECTION

- Syntax | 3408
- Description | 3408
- Options | 3408
- Required Privilege Level | 3408
- Output Fields | 3408
- Sample Output | 3408
- Release Information | 3408

Syntax

```
clear accounting server statistics archival-transfer
```

Description

Clears the statistics of transfer attempted, succeeded, and failed for accounting statistics files and router configuration archives.

Options

This command has no options.

Required Privilege Level

clear

Output Fields

When you enter this command, the transfer statistics are cleared.

Sample Output

clear accounting server statistics archival-transfer

```
user@host> clear accounting server statistics archival-transfer
```

Release Information

Command introduced in Junos OS Release 19.2.

clear-action (CFM)

IN THIS SECTION

- Syntax | 3409
- Hierarchy Level | 3409
- Description | 3409
- Options | 3409
- Required Privilege Level | 3410
- Release Information | 3410

Syntax

```
clear-action {
    interface-down peer-interface;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile profile-name]
```

Description

Clear the action or actions to be taken when the connectivity fault management event occurs. You cannot configure multiple actions at this time. Only one action can be configured. This limitation affects both the **action** and `clear-action` statements.

Options

peer-interface—Name of the peer interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring a CFM Action Profile to Specify CFM Actions for CFM Events

clear amt statistics

IN THIS SECTION

- Syntax | 3410
- Description | 3411
- Options | 3411
- Required Privilege Level | 3411
- Output Fields | 3411
- Sample Output | 3411
- Release Information | 3411

Syntax

```
clear amt statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Clear Automatic Multicast Tunneling (AMT) statistics.

Options

none	Clear the multicast statistics for all AMT tunnel interfaces.
instance ***instance-name***	(Optional) Clear AMT multicast statistics for the specified instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear amt statistics

```
user@host> clear amt statistics
```

Release Information

Command introduced in JUNOS Release 10.2.

RELATED DOCUMENTATION

show amt statistics | 17033

clear amt tunnel

IN THIS SECTION

- Syntax | 3412
- Description | 3412
- Options | 3412
- Required Privilege Level | 3413
- Output Fields | 3413
- Sample Output | 3413
- Release Information | 3413

Syntax

```
clear amt tunnel
<gateway gateway-ip-addr> <port port-number>
<instance instance-name>
<logical-system (all | logical-system-name)>
<statistics>
<tunnel-interface interface-name>
```

Description

Clear the Automatic Multicast Tunneling (AMT) multicast state. Optionally, clear AMT protocol statistics.

Options

none	Clear multicast state for all AMT tunnel interfaces.
gateway *gateway-ip-addr* port *port-number*	(Optional) Clear the AMT multicast state for the specified gateway address. If no port is specified, clear the AMT multicast state for all AMT gateways with the given IP address.
instance *instance-name*	(Optional) Clear the AMT multicast state for the specified instance.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
statistics	(Optional) Clear multicast statistics for all AMT tunnels or for specified tunnels.
tunnel-interface *interface-name*	(Optional) Clear the AMT multicast state for the specified AMT tunnel interface.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear amt tunnel

```
user@host> clear amt tunnel
```

clear amt tunnel statistics gateway-address

```
user@host> clear amt tunnel statistics gateway-address 100.31.1.21 port 4000
```

Release Information

Command introduced in JUNOS Release 10.2.

RELATED DOCUMENTATION

show amt tunnel | 17039

clear ancp access-loop

IN THIS SECTION

- Syntax | 3414
- Description | 3414
- Options | 3414
- Required Privilege Level | 3415
- Output Fields | 3415
- Sample Output | 3415
- Release Information | 3415

Syntax

```
clear ancp access-loop
(neighbor ip-address | subscriber-interface physical-interface-name)
circuit-id aci
remote-id ari
outer-vlan-id vlan-id
```

Description

Clear the connection for the subscriber on the specified access loop for an ANCP-triggered, autosensed dynamic VLAN. The autoconfiguration daemon (autoconfd) deletes any existing cached information about the subscriber. This command simulates a CPE connection reset as seen when the access node sends a Port Down message followed by a Port Up message.

Options

aci ANCP Access-Loop-Circuit-ID TLV that identifies the subscriber-side access loop logical port and partially identifies an access loop to clear.

ari ANCP Access-Loop-Remote-ID TLV that uniquely identifies the subscriber on the access loop and partially identifies an access loop to clear.

ip-address	ANCP neighbor's IP address that specifies an access loop to clear.
physical-interface-name	Subscriber interface that specifies an access loop to clear.
vlan-id	ANCP Access-Aggregation-Circuit-ID-Binary TLV that identifies the logical circuit identifier on the NAS side and partially identifies an access loop to clear.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp neighbor` command before and after clearing the access loop to verify the clear operation.

Sample Output

clear ancp access-loop

```
user@host> clear ancp-access-loop neighbor 192.168.25.31 circuit-id line-aci-1 remote-id line-
ari-1 outer-vlan-id 126
```

Release Information

Command introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Clearing ANCP Access Loops

Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

clear ancp neighbor

IN THIS SECTION

- Syntax | 3416
- Description | 3416
- Options | 3416
- Required Privilege Level | 3417
- Output Fields | 3417
- Sample Output | 3417
- Release Information | 3418

Syntax

```
clear ancp neighbor
<ip-address ip-address>
<system-name mac-address>
```

Description

Clear the ANCP agent connection with all ANCP neighbors or with the specified ANCP neighbor. This command deletes information for subscribers associated with the neighbor, causing the adjusted traffic rates to revert to the configured rate for the subscriber interfaces. The neighbor remains configured (its administrative state is *enabled*) and can reestablish adjacencies.

This command initiates logout of ANCP-triggered dynamic VLAN sessions on the physical interface associated with the specified neighbor; conventionally autosensed dynamic VLAN sessions and their associated logical interfaces are not affected.

Options

none	Clear all ANCP neighbors.
ip-address *ip-address*	(Optional) Clear the ANCP neighbor specified by the IP address.

system-name ***mac-address*** (Optional) Clear the ANCP neighbor specified by the MAC address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp neighbor` command before and after clearing the ANCP neighbors to verify the clear operation.

Sample Output

clear ancp neighbor

```
user@host> clear ancp neighbor
```

show ancp neighbor

The following sample output displays the connections with ANCP neighbors before and after the `clear ancp neighbor` command was issued.

```
user@host> show ancp neighbor

  IP Address        MAC Address         State         Subscriber   Capabilities
                                                        Count
  203.0.113.102     00:00:5e:00:53:10  Established    5              Topo
  203.0.113.122     00:00:5e:00:53:12  Established    5              Topo
  203.0.113.132     00:00:5e:00:53:13  Established    5              Topo
  203.0.113.142     00:00:5e:00:53:14  Established    5              Topo


user@host> clear ancp neighbor ip-address 203.0.113.102

user@host> show ancp neighbor

  IP Address        MAC Address         State         Subscriber   Capabilities
```

```
                                                        Count
  203.0.113.122     00:00:5e:00:53:12  Established    5              Topo
  203.0.113.132     00:00:5e:00:53:13  Established    5              Topo
  203.0.113.142     00:00:5e:00:53:14  Established    5              Topo
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

show ancp neighbor | 17054

clear ancp statistics

IN THIS SECTION

- Syntax | 3418
- Description | 3419
- Options | 3419
- Required Privilege Level | 3419
- Output Fields | 3419
- Sample Output | 3419
- Release Information | 3421

Syntax

```
clear ancp statistics
<ip-address ip-address>
<system-name mac-address>
```

Description

Clear current statistics accumulated by the ANCP agent for all ANCP neighbors or the specified neighbor.

Options

none	Clear all ANCP statistics.
ip-address *ip-address*	(Optional) Clear statistics for the ANCP neighbor specified by the IP address.
system-name *mac-address*	(Optional) Clear statistics for the ANCP neighbor specified by the MAC address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp neighbor` command before and after clearing the ANCP neighbor statistics to verify the clear operation.

Sample Output

clear ancp statistics

```
user@host> clear ancp statistics
```

show ancp neighbor

The following sample output displays statistics for an ANCP neighbor before and after the `clear ancp statistics` command was issued.

```
user@host> show ancp neighbor ip-address 192.168.10.1 detail
Neighbor Information
```

```
      IP Address              : 192.168.10.1
      System Name             : 00:00:5E:00:53:02
          Up Time                              : 38
          TCP Port                             : 64959
          State                                : Established
          Subscriber Count                     : 7
          Capabilities                         : Topology Discovery
          System Instance                      : 11
          Peer Instance                        : 1
          Adjacency Timer (in 100ms)           : 50
          Peer Adjacency Timer (in 100ms)      : 100
          Partition Type                       : 0
          Partition Flag                       : 1
          Partition Identifier                 : 0
          Dead Timer                           : 22
          Received Syn Count                   : 47
          Received Synack Count                : 48
          Received Rstack Count                : 2
          Received Ack Count                   : 12
          Received Port Up Count               : 8
          Received Port Down Count             : 2
          Received Other Count                 : 0
          Sent Syn Count                       : 48
          Sent Synack Count                    : 47
          Sent Rstack Count                    : 1
          Sent Ack Count                       : 12
          Max Discovery Limit Exceed Count     : 0

user@host> clear ancp statistics ip-address 192.168.10.1

user@host> show ancp neighbor ip-address 192.168.10.1 detail
Neighbor Information
      IP Address              : 192.168.10.1
      System Name             : 00:00:5E:00:53:02
          Up Time                              : 38
          TCP Port                             : 64959
          State                                : Established
          Subscriber Count                     : 7
          Capabilities                         : Topology Discovery
          System Instance                      : 11
          Peer Instance                        : 1
          Adjacency Timer (in 100ms)           : 50
          Peer Adjacency Timer (in 100ms)      : 100
```

```
            Partition Type                      : 0
            Partition Flag                      : 1
            Partition Identifier                : 0
            Dead Timer                          : 22
            Received Syn Count                  : 0
            Received Synack Count               : 0
            Received Rstack Count               : 0
            Received Ack Count                  : 0
            Received Port Up Count              : 0
            Received Port Down Count            : 0
            Received Other Count                : 0
            Sent Syn Count                      : 0
            Sent Synack Count                   : 0
            Sent Rstack Count                   : 0
            Sent Ack Count                      : 0
            Max Discovery Limit Exceed Count    : 0
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

show ancp neighbor | 17054

clear ancp subscriber

IN THIS SECTION

- Syntax | 3422
- Description | 3422
- Options | 3422
- Required Privilege Level | 3422
- Output Fields | 3423

Sample Output | 3423
Release Information | 3423

Syntax

```
clear ancp subscriber
<identifier identifier>
<ip-address ip-address>
<system-name mac-address>
```

Description

Clear the ANCP agent connection with all ANCP subscribers or with the specified ANCP subscriber. This command deletes information for the subscribers, causing the adjusted traffic rate to revert to the configured rate for the subscriber interface, but otherwise has no affect on ANCP neighbors.

Options

none	Clear all ANCP subscribers.
identifier *identifier-string*	(Optional) Clear the ANCP subscriber identified by the access loop ID.
ip-address *ip-address*	(Optional) Clear all ANCP subscribers on the neighbor specified by the IP address.
system-name *mac-address*	(Optional) Clear all ANCP subscribers on the neighbor specified by the MAC address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp subscriber` command before and after clearing the ANCP neighbors to verify the clear operation.

Sample Output

show ancp subscriber brief

```
user@host> show ancp subscriber brief
Loop Identifier          Type      Interface        Rate          Neighbor
                                                      Kbps
  port-1-10                VDSL2    set-ge-10410       64       203.0.113.102
  port-1-11                VDSL2    set-ge-10411       64       203.0.113.112
  port-2-10                VDSL2    ge-1/0/4.12        64       203.0.113.122
  port-2-10                VDSL2    ge-1/0/4.12        64       203.0.113.123
  port-2-11                VDSL2    ge-1/0/4.13        64       203.0.113.132
user@host> clear ancp subscriber identifier port-2-10

user@host> show ancp subscriber brief
Loop Identifier          Type      Interface        Rate          Neighbor
                                                      Kbps
  port-1-10                VDSL2    set-ge-10410       64       203.0.113.102
  port-1-11                VDSL2    set-ge-10411       64       203.0.113.112
  port-2-11                VDSL2    ge-1/0/4.13        64       203.0.113.132
```

clear ancp subscriber

```
user@host> clear ancp subscriber
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show ancp subscriber | 17072

clear arp

IN THIS SECTION

- Syntax | 3424
- Description | 3424
- Options | 3425
- Required Privilege Level | 3425
- Output Fields | 3425
- Sample Output | 3425
- Release Information | 3426

Syntax

```
clear arp
<all>
<hostname hostname>
<interface interface-name>
<logical-system logical-system-name>
<tenant name>
<vpn vpn>
```

Description

Remove entries from the Address Resolution Protocol (ARP) table for the current CLI view. To clear entries for a specific logical system, you must first enter the `set cli logical-system` *`logical-system-name`* command, and then issue the `clear arp` command.

Options

all	Clear all entries from the ARP table.
hostname *hostname*	(Optional) Clear only the specified host entry from the ARP table.
interface *interface-name*	(Optional) Clear entries only for the specified interface from the ARP table.
logical-system *logical-system-name*	(Optional) Clear entries for only the specified logical system from the ARP table (only available in main router context).
tenant *name*	(Optional) Clear entries for only the specified tenant from the ARP table (only available in main router context).
vpn *<routing-instance-name>*	(Optional) Clear entries from the ARP table for the specified routing instance. You can also use `clear arp vpn <routing-instance-name> hostname <hostname-learned-from-routing-instance>` to delete a specific entry.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear arp all

```
user@host> clear arp all
192.168.71.254   deleted
192.168.65.46    deleted
192.168.64.10    deleted
10.0.12.14       deleted
10.0.17.14       deleted
```

clear arp logical-system ls1

```
user@host> clear arp logical-system ls1
192.168.71.254   deleted
192.168.65.46    deleted
192.168.64.10    deleted
10.0.12.14       deleted
10.0.17.14       deleted
```

Release Information

Command introduced before Junos OS Release 7.4.

`all` option introduced in Junos OS Release 14.2.

tenant option added in Junos OS Release 18.3.

RELATED DOCUMENTATION

- set cli logical-system
- show arp | 17096
- show dhcp-security arp inspection statistics
- Port Security Features

clear arp inspection statistics

IN THIS SECTION

- Syntax | 3427
- Description | 3427
- Options | 3427
- Required Privilege Level | 3427
- Output Fields | 3427
- Sample Output | 3427

Release Information | 3427

Syntax

```
clear arp inspection statistics
<interface interface>
```

Description

Clear ARP inspection statistics.

Options

none	Clears ARP statistics on all interfaces.
interface *interface-names*	(Optional) Clear ARP statistics on one or more interfaces.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear arp inspection statistics

```
user@switch> clear arp inspection statistics
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show arp inspection statistics | 17103

Example: Configuring Port Security (non-ELS)

Verifying That DAI Is Working Correctly

clear auto-configuration interfaces

IN THIS SECTION

- Syntax | 3428
- Description | 3428
- Options | 3428
- Required Privilege Level | 3429
- Output Fields | 3429
- Sample Output | 3429
- Release Information | 3429

Syntax

```
clear auto-configuration interfaces interface-name
```

Description

Clear dynamically created VLAN interfaces.

NOTE: For the clear command to be successful, no interface bindings (for example, DHCP server bindings) can exist on the dynamic interface.

Options

interface-name	Name of a physical or logical interface.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear auto-configuration interfaces (All Interfaces)

```
user@host> clear auto-configuration interfaces ge-1/0/0
```

command-name

```
10 interfaces removed from device ge-1/0/0
```

clear auto-configuration interfaces (Single Dynamically Created Interface)

```
user@host> clear auto-configuration interfaces ge-1/0/0.1073741824
```

command-name

```
Interface ge-1/0/0.1073741824 deleted
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/subscriber-access/subscriber-mgmt-vlan-interface.html

clear auto-configuration interfaces interface-set

IN THIS SECTION

- Syntax | 3430
- Description | 3430
- Options | 3430
- Required Privilege Level | 3431
- Output Fields | 3431
- Sample Output | 3431
- Release Information | 3431

Syntax

```
clear auto-configuration interfaces interface-set interface-set-name
```

Description

Clear a specified dynamic agent circuit identifier (ACI) interface set or access-line identifier (ALI) interface set on the router. An ACI or ALI interface set is a logical collection of dynamic VLAN subscriber interfaces that originate at the same household or on the same access-loop port.

You can clear only those ACI or ALI interface sets that have no active subscriber interface members. If the ACI or ALI interface set that you want to clear still has valid member interfaces, you must first remove these interfaces before issuing the `clear auto-configuration interfaces interface-set` *`interface-set-name`* command.

Options

interface-set-name Name of the empty ACI or ALI interface set that you want to clear. Use the ACI or ALI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the

actual ACI or ALI string found in the DHCP or PPPoE control packets. To view the names of the ACI or ALI interface sets configured on the router, you can issue the `show subscribers` command.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear auto-configuration interfaces interface-set

```
user@host> clear auto-configuration interfaces interface-set aci-1003-ge-1/0/0.4001
Interface-set aci-1003-ge-1/0/0.4001 deleted
```

clear auto-configuration interfaces interface-set (Error Message for ACI Interface Set with Active Members)

```
user@host> clear auto-configuration interfaces interface-set aci-1005-ge-1/0/0.2800
error: Interface set aci-1005-ge-1/0/0.2800 has references.
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Clearing Agent Circuit Identifier Interface Sets

Clearing Access-Line-Identifier Interface Sets

clear bfd adaptation

IN THIS SECTION

- Syntax | 3432
- Description | 3432
- Options | 3432
- Additional Information | 3433
- Required Privilege Level | 3433
- Output Fields | 3433
- Sample Output | 3433
- Release Information | 3433

Syntax

```
clear bfd adaptation
<all>
<address session-address>
<discriminator discr-number>
```

Description

Clear adaptation for Bidirectional Forwarding Detection (BFD) sessions. BFD is a simple hello mechanism that detects failures in a network. Configured BFD interval timers can change, adapting to network situations. Use this command to return BFD interval timers to their configured values.

The `clear bfd adaptation` command is hitless, meaning that the command does not affect traffic flow on the routing device.

Options

all	Clear adaptation for all BFD sessions.
address *session-address*	(Optional) Clear adaptation for all BFD sessions matching the specified address.

discriminator *discr-number* (Optional) Clear adaptation for the local BFD session matching the specified discriminator.

Additional Information

For more information, see the description of the `bfd-liveness-detection` configuration statement in the *Junos Routing Protocols Configuration Guide*.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bfd adaptation

```
user@host> clear bfd adaptation
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show bfd session | 17134

clear bfd session

IN THIS SECTION

- Syntax | 3434
- Syntax (EX Series Switch and QFX Series) | 3434
- Description | 3434
- Options | 3435
- Required Privilege Level | 3435
- Output Fields | 3435
- Sample Output | 3435
- Release Information | 3435

Syntax

```
clear bfd session
<all>
<address session-address>
<discriminator discr-number>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and QFX Series)

```
clear bfd session
<all>
<address session-address>
<discriminator discr-number>
```

Description

Drop one or more Bidirectional Forwarding Detection (BFD) sessions.

Options

all	Drop all BFD sessions.
address *session-address*	(Optional) Drop all BFD sessions matching the specified address.
discriminator *discr-number*	(Optional) Drop the local BFD session matching the specified discriminator.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bfd session all

```
user@host> clear bfd session all
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show bfd session | 17134

clear bgp damping

IN THIS SECTION

- Syntax | 3436
- Syntax (EX Series Switch and QFX Series) | 3436
- Description | 3436
- Options | 3436
- Required Privilege Level | 3437
- Output Fields | 3437
- Sample Output | 3437
- Release Information | 3437

Syntax

```
clear bgp damping
<logical-system (all | logical-system-name)>
<prefix>
```

Syntax (EX Series Switch and QFX Series)

```
clear bgp damping
<prefix>
```

Description

Clear BGP route flap damping information.

Options

none	Clear all BGP route flap damping information.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
prefix	(Optional) Clear route flap damping information for only the specified destination prefix.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear bgp damping

```
user@host> clear bgp damping
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show policy damping | 23206
show route damping | 23564

clear bgp neighbor

IN THIS SECTION

Syntax | 3438

- Syntax (EX Series Switch and QFX Series) | 3438
- Description | 3439
- Options | 3439
- Required Privilege Level | 3440
- Output Fields | 3440
- Sample Output | 3440
- Release Information | 3440

Syntax

```
clear bgp neighbor
<all>
<as as-number>
<gracefully>
<instance instance-name>
<logical-system (all | logical-system-name)>
<malformed-route>
<neighbor>
<soft | soft-inbound>
<soft-minimum-igp>
<soft-reconfig>
<stale-routes>
```

Syntax (EX Series Switch and QFX Series)

```
clear bgp neighbor
<all>
<as as-number>
<instance instance-name>
<malformed-route>
<neighbor>
<soft | soft-inbound>
<soft-minimum-igp>
```

Description

Perform one of the following tasks:

- Change the state of one or more BGP neighbors to `IDLE`. For neighbors in the `ESTABLISHED` state, this command drops the TCP connection to the neighbors and then reestablishes the connection.
- (`soft` keyword only) Reapply export policies and send refresh updates to one or more BGP neighbors without changing their state.
- (`soft-inbound` keyword only) Send a route-refresh message to one or more BGP neighbors without changing their state, and reapply import policies on the received updates.

Options

all	Change the state of all BGP neighbors to `IDLE`.
as *as-number*	(Optional) Apply this command only to neighbors in the specified autonomous system (AS).
gracefully	(Optional) Enable the BGP peer to start graceful-restart receiving-speaker mode. The receiving speaker also sends its own routes to the restarted speaker, and sends an End-of-RIB marker when it completes the update. The `clear bgp neighbor` *`neighbor-address`* `gracefully` command is the same as `clear bgp neighbor hard` (the default for `clear bgp neighbor`), but it does not use the new Hard Reset subcode on the Notify and Cease messages that are sent. This allows the neighbor to enter GR or LLGR helper mode, if negotiated. The session is still cleared on this router, and this router does not enter GR or LLGR helper mode.
instance *instance-name*	(Optional) Apply this command only to neighbors for the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
malformed-route	(Optional) Remove malformed routes. If a specific neighbor is provided, Junos OS removes malformed routes for that particular neighbor. Otherwise, Junos OS removes malformed routes for all BGP neighbors. To find routes that have malformed attributes, run the `show route hidden` command, and look for routes marked with `MalformedAttr` in the AS path field.
neighbor	(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.

soft	(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.
soft-inbound	(Optional) Send a route-refresh message to BGP neighbors and reapply import policies on the route updates received from the BGP neighbors without clearing the BGP state.
soft-minimum-igp	(Optional) Provide soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.
soft-reconfig	(Optional) Trigger the import re-evaluation of a set of BGP neighbors. soft reset inbound state by import re-evaluation.
stale-routes	(Optional) Any stale route currently being held for the specified neighbor because of BGP graceful restart (GR) or long-lived graceful restart (LLGR) receiver mode operations.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bgp neighbor

```
user@host> clear bgp neighbor
```

Release Information

Command introduced before Junos OS Release 7.4.

`malformed-route` option introduced in Junos OS Release 13.2.

`all` option introduced in Junos OS Release 14.2.

`gracefully` and `stale-routes` options introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

show bgp neighbor | 17175

clear bgp table

IN THIS SECTION

- Syntax | 3441
- Syntax (EX Series Switch and QFX Series) | 3441
- Description | 3441
- Options | 3442
- Additional Information | 3442
- Required Privilege Level | 3442
- Output Fields | 3442
- Sample Output | 3442
- Release Information | 3443

Syntax

```
clear bgp table table-name
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and QFX Series)

```
clear bgp table table-name
```

Description

Request that BGP refresh routes in a specified routing table.

Options

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
table-name	Request that BGP refresh routes in the specified table.

Additional Information

In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the `clear bgp table` command to request that BGP refresh routes in a VPN instance table.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear bgp table private.inet.0

```
user@host> clear bgp table private.inet.0
```

clear bgp table inet.6 logical-system all

```
user@host> clear bgp table inet.6 logical-system all
```

clear bgp table private.inet.6 logical-system ls1

```
user@host> clear bgp table private.inet.6 logical-system ls1
```

clear bgp table logical-system all inet.0

```
user@host> clear bgp table logical-system all inet.0
```

clear bgp table logical-system ls2 private.inet.0

```
user@host> clear bgp table logical-system ls2 private.inet.0
```

Release Information

Command introduced in Junos OS Release 9.0.

clear bridge mac-ip-table

IN THIS SECTION

- Syntax | 3443
- Description | 3444
- Options | 3444
- Required Privilege Level | 3444
- Release Information | 3444

Syntax

```
clear bridge mac-ip-table
<address>
<bridge-domain (all | bridge-domain-name>
<instance instance-name>
<logical-system (all | logical-system-name>
```

Description

Clear the locally learnt ARP entries for EVPN routing instances where the instance-type is virtual-switch. Remotely learnt entries cannot be cleared

Options

address	(Optional) Clear the specific suppressed MAC address from the MAC-IP address table .
bridge-domain (all | *bridge-domain-name*)	(Optional) Clear the MAC-IP address table for all bridge domains or for a specified bridge domain.
instance *instance-name*	(Optional) Clear the MAC-IP address table for a specific routing instance.
logical-system (all | *logical-system-name*)	(Optional) Clear the MAC-IP address for all logical systems or on a specific logical system.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 17.4R2.

clear bridge mac-table

IN THIS SECTION

- Syntax | 3445
- Description | 3445
- Options | 3445
- Required Privilege Level | 3445
- Output Fields | 3446
- Sample Output | 3446

Release Information | 3446

Syntax

```
clear bridge mac-table
<bridge-domain (all | bridge-domain-name)>
<instance instance-name>
<interface interface-name>
<learning-vlan id (all-vlan | learning-vlan-id)>
<mac-address>
```

Description

(MX Series routers only) Clear learned Layer 2 address information from the media access control (MAC) address table.

Options

none	Clear all learned Layer 2 address information from the MAC address table.
bridge-domain (all | *bridge-domain-name*)	(Optional) Clear learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.
instance *instance-name*	(Optional) Clear learned Layer 2 MAC addresses for the specified routing instance.
interface *interface-name*	(Optional) Clear learned Layer 2 MAC addresses for the specified interface.
learning-vlan-id (all-vlan | *learning-vlan-id*)	(Optional) Clears learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.
mac-address	(Optional) Clear the specified learned Layer 2 address from the MAC address table.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bridge mac-table

```
user@host> clear bridge mac-table
```

Release Information

Command introduced in Junos OS Release 8.4.

clear bridge recovery-timeout

IN THIS SECTION

- Syntax | 3446
- Description | 3447
- Options | 3447
- Required Privilege Level | 3447
- Sample Output | 3447
- Release Information | 3447

Syntax

```
clear bridge recovery-timeout
<interface interface-name>
```

Description

Clear all storm control errors from all the Ethernet switching interfaces on the switch, and restore the interfaces to service.

Options

interface *interface-name*	Clear all storm control errors from the Ethernet switching interfaces on the interface specified in the command and restore this interface to service.

Required Privilege Level

clear

Sample Output

clear bridge recovery-timeout (interface interface-name)

```
user@host> clear bridge recovery-timeout interface ae0.0
user@host> clear bridge recovery-timeout interface ae0.0
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring Autorecovery for Port Security Events

clear bridge statistics

IN THIS SECTION

Syntax | 3448

- Description | 3448
- Options | 3448
- Required Privilege Level | 3448
- Output Fields | 3449
- Sample Output | 3449
- Sample Output | 3449
- Sample Output | 3449
- Release Information | 3449

Syntax

```
clear bridge statistics
<bridge-domain bridge-domain>
<instance instance-name bridge-domain bridge-domain>
<logical-system logical-system-name instance instance-name bridge-domain bridge-domain>
```

Description

Clear bridge statistics of all interfaces in a routing instance.

Options

bridge-domain *bridge-domain*	(Optional) Clear bridge statistics of all interfaces for the specified bridge domain in the default routing instance in the global logical system.
instance *instance-name* bridge-domain *bridge-domain*	(Optional) Clear bridge statistics of all interfaces for the specified bridge domain in the specified routing instance in the global logical system.
logical-system *logical-system-name* instance *instance-name* bridge-domain *bridge-domain*	(Optional) Clear bridge statistics of all interfaces for the specified bridge domain, routing instance, and logical system name.

Required Privilege Level

maintenance

Output Fields

These commands produces no output.

Sample Output

clear bridge statistics

```
user@host> clear bridge statistics bridge-domain bridge-domain-name
```

Sample Output

clear bridge statistics

```
user@host> clear bridge statistics instance instance-name bridge-domain bridge-domain-name
```

Sample Output

clear bridge statistics

```
user@host> clear bridge statistics logical-system logical-system-name instance instance-name
bridge-domain bridge-domain-name
```

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

clear evpn statistics | 3592

clear captive-portal

IN THIS SECTION

- Syntax | 3450
- Description | 3450
- Options | 3450
- Required Privilege Level | 3451
- Output Fields | 3451
- Sample Output | 3452
- Release Information | 3452

Syntax

```
clear captive-portal (firewall  [interface-names] | interface (802.1X) (all | [interface-names])
| mac-address [mac-addresses])
```

Description

Reset the authentication state of a captive portal interface or captive portal firewall statistics on one or more interfaces.

Options

firewall [*interface-names*]	Resets captive portal statistics on all interfaces or on the specified interface.
interface (all | *interface-names*)	Resets the authentication state of users connected to all interfaces or the specified interfaces.
mac-address *mac-addresses*	Resets the authentication state for the specified MAC addresses.

Required Privilege Level

view

Output Fields

Table 12 on page 3451 lists the output fields for the `clear captive-portal interface` command. (The **clear captive-portal firewall** and `clear captive-portal mac-address` commands have no output). Output fields are listed in the approximate order in which they appear.

Table 12: clear captive-portal interface Output Fields

Field Name	Field Description
Interface	Interface on which captive portal has been configured.
State	The state of the port: • **Authenticated**—The client has been authenticated through the RADIUS server or has been permitted access through server fail fallback. • **Authenticating**—The client is authenticating through the RADIUS server. • **Connecting**—Switch is attempting to contact the RADIUS server. • **Initialize**—The interface link is down. • **Held**—An action has been triggered through server fail fallback during a RADIUS server timeout. A supplicant is denied access, permitted access through a specified VLAN, or maintains the authenticated state granted to it before the RADIUS server timeout occurred.
MAC address	The MAC address of the connected client on the interface.
User	Users connected to the captive portal interface.

Sample Output

clear captive-portal interface

```
user@switch> clear captive-portal interface
ge-0/0/3.0
```

clear captive-portal interface

```
user@switch> clear captive-portal interface
Captive Portal Information:
Interface      State           MAC address          User
ge-0/0/3.0     Authenticated   00:03:47:e1:ba:b9    aclallow
ge-0/0/5.0     Connecting
ge-0/0/7.0     Connecting
ge-0/0/9.0     Connecting
```

clear captive-portal mac-address

```
user@switch> clear captive-portal mac-address 00:03:47:e1:ba:b9
```

This command has no output.

clear captive-portal firewall

```
user@switch> clear captive-portal firewall
```

This command has no output.

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

show captive-portal authentication-failed-users | 17281

show captive-portal interface | 17286
show captive-portal firewall | 17283
Example: Setting Up Captive Portal Authentication on an EX Series Switch
Configuring Captive Portal Authentication (CLI Procedure)

clear chassis alarms fabric degraded

IN THIS SECTION

- Syntax (TX Matrix Plus Router with 3D SIBs) | 3453
- Description | 3453
- Options | 3453
- Required Privilege Level | 3454
- Sample Output | 3454
- Release Information | 3455

Syntax (TX Matrix Plus Router with 3D SIBs)

```
clear chassis alarms fabric degraded lcc number fpc number
```

Description

Clear the fabric degraded alarm for an FPC.

Options

lcc *number*

Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

fpc ***number***

Flexible PIC Concentrator (FPC) slot number. On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, specify the number of a T1600 or T4000 router by using the `lcc` *`number`* option and replace `fpc` *`number`* with a value from 0 through 7.

Required Privilege Level

clear

Sample Output

show system alarms (TX Matrix Plus router with 3D SIBs)

```
user@host> show system alarms


sfc0-re0:
--------------------------------------------------------------------------
2 alarms currently active
Alarm time               Class  Description
2013-05-08 18:13:58 UTC  Major  LCC 0 Major Errors
2013-05-08 17:48:46 UTC  Major  LCC 7 Major Errors

lcc0-re1:
--------------------------------------------------------------------------
3 alarms currently active
Alarm time               Class  Description
2013-05-08 17:35:34 UTC  Minor  SIB 3 Not Online
2013-05-08 17:35:34 UTC  Minor  SIB 2 Not Online
2013-05-08 18:19:24 UTC  Major  FPC 5 degraded fabric condition detected


user@host> clear chassis alarms fabric degraded
lcc 0 fpc 5
lcc0-re1:
--------------------------------------------------------------------------
```

```
user@host> show system alarms



sfc0-re0:
--------------------------------------------------------------------------
2 alarms currently active
Alarm time               Class  Description
2013-05-08 18:13:58 UTC  Major  LCC 0 Major Errors
2013-05-08 17:48:46 UTC  Major  LCC 7 Major Errors

lcc0-re1:
--------------------------------------------------------------------------
2 alarm currently active
Alarm time               Class  Description
2013-05-08 17:36:34 UTC  Minor  SIB 3 Not Online
2013-05-08 17:36:34 UTC  Minor  SIB 2 Not Online
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

show system alarms | 26683

clear chassis cluster control-plane statistics

IN THIS SECTION

- Syntax | 3456
- Description | 3456
- Required Privilege Level | 3456
- Output Fields | 3456
- Sample Output | 3456

- Release Information | 3456

Syntax

```
clear chassis cluster control-plane statistics
```

Description

Clear the control plane statistics of a chassis cluster.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear chassis cluster control-plane statistics

```
user@host> clear chassis cluster control-plane statistics
Cleared control-plane statistics
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show chassis cluster control-plane statistics | 17351

clear chassis cluster data-plane statistics

IN THIS SECTION

- Syntax | 3457
- Description | 3457
- Required Privilege Level | 3457
- Output Fields | 3457
- Sample Output | 3457
- Release Information | 3458

Syntax

```
clear chassis cluster data-plane statistics
```

Description

Clear the data plane statistics of a chassis cluster.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear chassis cluster data-plane statistics

```
user@host> clear chassis cluster data-plane statistics
Cleared data-plane statistics
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show chassis cluster data-plane statistics | 17356

clear chassis cluster failover-count

IN THIS SECTION

- Syntax | 3458
- Description | 3458
- Required Privilege Level | 3458
- Output Fields | 3459
- Sample Output | 3459
- Release Information | 3460

Syntax

```
clear chassis cluster failover-count
```

Description

Clear the failover count of all redundancy groups.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

The following example displays the redundancy groups before and after the failover-counts are cleared.

show chassis cluster status

```
user@host> show chassis cluster status

Cluster ID: 3
Node name         Priority     Status     Preempt  Manual failover

Redundancy group: 0 , Failover count: 1
   node0            200        secondary   no        no
   node1            100        primary     no        no

Redundancy group: 1 , Failover count: 2
   node0            100        primary     no        no
   node1            10         secondary   no        no
```

clear chassis cluster failover-count

```
user@host> clear chassis cluster failover-count
Cleared failover-count for all redundancy-groups
```

show chassis cluster status

```
user@host> show chassis cluster status

Cluster ID: 3
Node name        Priority   Status     Preempt  Manual failover

Redundancy group: 0 , Failover count: 0
  node0           200       secondary    no       no
  node1           100       primary      no       no
```

```
Redundancy group: 1 , Failover count: 0
  node0           100       primary      no       no
  node1           10        secondary    no       no
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

- request chassis cluster failover node | 14454
- request chassis cluster failover reset | 14457
- show chassis cluster status | 17405

clear chassis cluster ip-monitoring failure-count

IN THIS SECTION

- Syntax | 3460
- Description | 3461
- Required Privilege Level | 3461
- Output Fields | 3461
- Sample Output | 3461
- Release Information | 3461

Syntax

```
clear chassis cluster ip-monitoring failure-count
```

Description

Clear the failure count for all IP addresses.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
user@host> clear chassis cluster ip-monitoring failure-count
```

command-name

```
node0:
--------------------------------------------------------------------------
Cleared failure count for all IPs

node1:
--------------------------------------------------------------------------
Cleared failure count for all IPs
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

clear chassis cluster ip monitoring failure count ip address | 3462

clear chassis cluster ip monitoring failure count ip address

IN THIS SECTION

- Syntax | 3462
- Description | 3462
- Required Privilege Level | 3462
- Output Fields | 3462
- Sample Output | 3463
- Release Information | 3463

Syntax

```
clear chassis cluster ip-monitoring failure-count ip-address IP-address
```

Description

Clear the failure count for a specified IP address.

NOTE: Entering an IP address at the end of this command is optional. If you do not specify an IP address, the failure count for all monitored IP addresses is cleared.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
user@host> clear chassis cluster ip-monitoring failure-count ip-address 10.1.1.1
```

clear chassis cluster ip-monitoring failure-count ip-address 10.1.1.1

```
node0:
--------------------------------------------------------------------------
Cleared failure count for IP: 10.1.1.1

node1:
--------------------------------------------------------------------------
Cleared failure count for IP: 10.1.1.1
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

clear chassis cluster failover-count | 3458

clear chassis cluster ip-monitoring failure-count | 3460

clear chassis cluster statistics

IN THIS SECTION

- Syntax | 3464
- Description | 3464
- Required Privilege Level | 3464
- Output Fields | 3464
- Sample Output | 3464

- Release Information | 3464

Syntax

```
clear chassis cluster statistics
```

Description

Clear the control plane and data plane statistics of a chassis cluster.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear chassis cluster statistics

```
user@host> clear chassis cluster statistics
Cleared control-plane statistics
Cleared data-plane statistics
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show chassis cluster statistics | 17399

clear chassis display message

IN THIS SECTION

- Syntax | 3465
- Syntax (TX Matrix Router) | 3465
- Syntax (TX Matrix Plus Router) | 3465
- Syntax (QFabric Systems) | 3466
- Description | 3466
- Options | 3466
- Required Privilege Level | 3467
- Output Fields | 3467
- Sample Output | 3467
- Release Information | 3468

Syntax

```
clear chassis display message
```

Syntax (TX Matrix Router)

```
clear chassis display message
<lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
clear chassis display message
<lcc number | sfc number>
```

Syntax (QFabric Systems)

```
clear chassis display message
<node-device name | interconnect-device name>
```

Description

(M40e, M160, M320, T Series routers, EX Series, and QFabric systems only) Clear or stop a text message on the craft interface display, which is on the front of the router or switch or on the LCD panel display on the router or switch. The craft interface alternates the display of text messages with standard craft interface messages, switching between messages every 2 seconds. By default, on both the router and the switch, the text message is displayed for 5 minutes. The craft interface display has four 20-character lines. The LCD panel display has two 16-character lines, and text messages appear only on the second line.

Options

none

Clear or stop a text message on the craft interface display.

interconnect-device *name*

(QFabric systems only) (Optional) On a QFabric system, clear or stop a text message on the LCD panel display on the specified Interconnect device.

lcc *number*

(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

node-device *name*

(QFabric systems only) (Optional) On a QFabric system, clear or stop a text message on the LCD panel display on the specified Node device in a Node group.

scc (TX Matrix routers only) (Optional) Clear or stop a text message on the craft interface on the TX Matrix router (switch-card chassis).

sfc *number* (TX Matrix Plus routers only) (Optional) Clear or stop a text message on the craft interface on the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with 0.

Required Privilege Level

clear

Output Fields

See "show chassis craft-interface" on page 17410 for an explanation of output fields.

Sample Output

clear chassis display message

The following example displays and then clears the text message on the craft interface display:

```
user@host> show chassis craft-interface
Red alarm:     LED off, relay off
Yellow alarm:  LED off, relay off
Host OK LED:   On
Host fail LED: Off
FPCs     0  1  2  3  4  5  6  7
-------------------------------
Green  ..  *..  *  *.
Red    ........
LCD screen:
     +--------------------+
     |NOC contact Dusty   |
     |(888) 526-1234      |
     +--------------------+

user@host> clear chassis display message

user@host> show chassis craft-interface
Red alarm:     LED off, relay off
```

```
Yellow alarm:  LED off, relay off
Host OK LED:   On
Host fail LED: Off
FPCs     0  1  2  3  4  5  6  7
-------------------------------
Green  ..  *..  *  *.
Red    ........
LCD screen:
     +--------------------+
     |host                |
     |Up: 0+17:05:47      |
     |                    |
     |Temperature OK      |
     +--------------------+
```

Release Information

Command introduced in Junos OS Release 7.5.

sfc option for the TX Matrix Plus routers introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring the LCD Panel on EX Series Switches (CLI Procedure)

set chassis display message | **16865**

show chassis craft-interface | **17410**

clear chassis fpc errors

IN THIS SECTION

- Syntax | **3469**
- Description | **3469**
- Options | **3469**
- Required Privilege Level | **3469**

- Output Fields | 3469
- Sample Output | 3470
- Release Information | 3470

Syntax

```
clear chassis fpc errors fpc-slot fpc-slot (all | error-id error-id)
```

Description

Clear the FPC errors on the device. You can choose to clear a particular error or all errors on the FPC.

> **NOTE:**
> `show chassis errors active detail`

Options

fpc-slot *fpc-slot*	The slot number of the FPC in which you want to run this command.
all	Clear all the errors on the FPC.
error-id *error-id*	Clear a particular error identified by an error-id. An *error-id*, a unique error identifier, is represented as a Uniform Resource Identifier (URI). For example, "`/cpu/0/memory/0/memory-uncorrected-error`" is an error-id that indicates an uncorrectable error under CPU memory module instance `0`.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear chassis fpc errors

```
user@host> clear chassis fpc errors fpc-slot 1 all
Clearing error(s) on fpc 1, option all
```

clear chassis fpc errors error-id *error-id*

```
user@host> clear chassis fpc errors error-id /fpc/0/pfe/0/cm/0/Eth_Port_Error/0/
ETH_CMERROR_MAJOR_0 fpc-slot 0

Clearing error(s) on fpc 0, option error-id error-id=/fpc/0/pfe/0/cm/0/Eth_Port_Error/0/
ETH_CMERROR_MAJOR_0.
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show chassis errors active | 17895

clear chassis hardware history

IN THIS SECTION

- Syntax | 3471
- Description | 3471
- Required Privilege Level | 3471
- Output Fields | 3471
- Sample Output | 3471

Release Information | 3471

Syntax

```
clear chassis hardware history
```

Description

Clears chassis hardware history database on Junos OS Evolved platforms.

Required Privilege Level

clear

Output Fields

When you enter this command, you are prompted to provide a confirmation of the request.

Sample Output

clear chassis hardware history

```
user@host> clear chassis hardware history
Are you sure you want to delete all history from both routing engines? [yes,no] (no) yes
```

Release Information

Command introduced in Junos OS Evolved 23.4R1.

RELATED DOCUMENTATION

No Link Title

clear chassis high-availability data-plane statistics

IN THIS SECTION

- Syntax | 3472
- Description | 3472
- Required Privilege Level | 3472
- Output Fields | 3472
- Sample Output | 3472
- Release Information | 3473

Syntax

```
clear chassis high-availability data-plane statistics
```

Description

Clear the data plane statistics of Multinode High Availability.

Required Privilege Level

clear

Output Fields

Sample Output

clear chassis high-availability data-plane statistics

```
user@host> clear chassis high-availability data-plane statistics
Cleared data-plane statistics
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show chassis high-availability data-plane statistics | 18476

clear chassis high-availability information

IN THIS SECTION

- Syntax | 3473
- Description | 3473
- Required Privilege Level | 3473
- Output Fields | 3474
- Sample Output | 3474
- Release Information | 3474

Syntax

```
clear chassis high-availability information
```

Description

Clear Multinode High Availability information.

Required Privilege Level

clear

Output Fields

Sample Output

clear chassis high-availability information

```
user@host> clear chassis high-availability information

Cleared chassis l3-ha information
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show chassis high-availability information | 18482

clear chassis sib errors

IN THIS SECTION

- Syntax | 3475
- Description | 3475
- Options | 3475
- Required Privilege Level | 3475
- Output Fields | 3475
- Sample Output | 3475
- Release Information | 3476

Syntax

```
clear chassis sib errors sib-slot sib-slot (all | error-id error-id)
```

Description

Clear the chassis SIB errors. You can choose to clear a particular error or all errors on the SIB.

Options

fpc-slot ***fpc-slot***	The slot number of the SIB in which you want to run this command.
all	Clear all the errors on the SIB.
error-id ***error-id***	Clear a particular error identified by an error-id. An *error-id*, a unique error identifier, is represented as a Uniform Resource Identifier (URI). For example, "`/cpu/0/memory/0/memory-uncorrected-error`" is an error-id that indicates an uncorrectable error under CPU memory module instance `0`.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear chassis sib errors

```
user@host> clear chassis sib errors sib-slot
1 all

Clearing error(s) on sib 1, option all
```

Release Information

Command introduced in Junos OS Evolved Release 19.4R1.

RELATED DOCUMENTATION

error | 6441

clear class-of-service application-traffic-control counter

IN THIS SECTION

- Syntax | 3476
- Description | 3476
- Options | 3477
- Required Privilege Level | 3477
- Output Fields | 3477
- Sample Output | 3477
- Release Information | 3477

Syntax

```
clear class-of-service application-traffic-control counter
<all-logical-systems-tenants>
<logical-system logical-system-name>
<node>
<root-logical-system>
<tenant tenant-name>
```

Description

Clears the application traffic control counters.

Options

all-logical-systems-tenants	(Optional) Clears the application traffic control counters of logical systems and tenant systems.
logical-system *logical-system-name*	(Optional) Clears the application traffic control counters of a specified logical system.
node	(Optional) Clears the application traffic control counters of a specified node.
root-logical-system	(Optional) Clears the application traffic control counters of the root logical-system (default).
tenant *tenant-name*	(Optional) Clears the application traffic control counters of a specified tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear class-of-service application-traffic-control counter

```
user@host>clear class-of-service application-traffic-control counter
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

show class-of-service application-traffic-control counter | 18968

clear class-of-service application-traffic-control rate-limiters

IN THIS SECTION

- Syntax | 3478
- Description | 3478
- Required Privilege Level | 3478
- Output Fields | 3478
- Sample Output | 3478
- Release Information | 3479

Syntax

```
clear class-of-service application-traffic-control rate-limiters
```

Description

Clears the application traffic control rate limiters information.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear class-of-service application-traffic-control rate-limiters

```
user@host>clear class-of-service application-traffic-control rate-limiters
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show class-of-service application-traffic-control rate-limiters | 18974

clear class-of-service application-traffic-control statistics rule

IN THIS SECTION

- Syntax | 3479
- Description | 3479
- Options | 3480
- Required Privilege Level | 3480
- Output Fields | 3480
- Sample Output | 3480
- Release Information | 3480

Syntax

```
clear class-of-service application-traffic-control statistics rule
<all-logical-systems-tenants>
<logical-system logical-system-name>
<node>
<root-logical-system>
<tenant tenant-name>
```

Description

Clears the application traffic control rule hit statistics. AppQoS counters identifying rule hits.

Options

all-logical-systems-tenants	Clears the application traffic control rule hit statistics of logical systems and tenant systems.
logical-system *logical-system-name*	(Optional) Clears the application traffic control rule hit statistics of a specified logical-system.
node	(Optional) Clears the application traffic control rule hit statistics of a specified node.
root-logical-system	(Optional) Clears the application traffic control rule hit statistics of the root logical system (default).
tenant *tenant-name*	(Optional) Clears the application traffic control rule hit statistics of a specified tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear class-of-service application-traffic-control statistics rule

```
user@host> clear class-of-service application-traffic-control statistics rule
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

show class-of-service application-traffic-control statistics rule

clear cli logical-system

IN THIS SECTION

- Syntax | 3481
- Description | 3481
- Options | 3481
- Required Privilege Level | 3481
- Output Fields | 3481
- Sample Output | 3482
- Release Information | 3482

Syntax

```
clear cli logical-system
```

Description

Clear the logical system view and return to a full router view. In a logical system view, the output of the command displays information related to the logical system only.

Options

This command has no options.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear cli logical-system

```
user@host:lr1> clear cli logical-system

Cleared default logical system


user@host>
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

set cli logical-system | 16875

Logical Systems User Guide for Routers and Switches

clear ddos-protection protocols

IN THIS SECTION

- Syntax | 3483
- Description | 3483
- Options | 3483
- Required Privilege Level | 3483
- Output Fields | 3483
- Sample Output | 3484
- Release Information | 3484

Syntax

```
clear ddos-protection protocols
<protocol-group <packet-type>> (culprit-flows | states  | statistics)
```

Description

Clear current DDoS protection statistics, violation states, or culprit flows for all packet types in all protocol groups, for all packet types in a particular protocol group, or for a particular packet type in a particular protocol group.

Options

protocol-group (Optional) Protocol group that is cleared. See "show ddos-protection protocols" on page 19222 for a list of available groups.

packet-type (Optional) Packet type in a particular protocol group that is cleared. See "show ddos-protection protocols" on page 19222 for a list of available packet types.

culprit-flows Clear culprit flows for a packet type, for a protocol group, or for all protocol groups. This option is not supported on QFX Series switches.

states Clear DDoS protection violation states for a packet type, for a protocol group, or for all protocol groups.

statistics Clear DDoS protection statistics such as packet counts and rates for a packet type, for a protocol group, or for all protocol groups.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ddos-protection protocols (Clear Statistics for All Protocols)

```
user@host> clear ddos-protection protocols statistics
```

clear ddos-protection protocols (Clear Violation States for Packet Type)

```
user@host> clear ddos-protection protocols radius server states
```

Release Information

Command introduced in Junos OS Release 11.2.

Option `culprit-flows` introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

show ddos-protection protocols | 19222
show ddos-protection statistics | 19295
show ddos-protection version | 19298

clear dhcp client binding

IN THIS SECTION

- Syntax | 3485
- Description | 3485
- Options | 3485
- Required Privilege Level | 3485
- Output Fields | 3485

Release Information | 3485

Syntax

```
clear dhcp client binding
[all|interface <interface-name>]
[routing-instance <routing-instance-name>]
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the DHCP client table.

Options

all	(Optional) Clear the binding state for all DHCP clients.
interface <interface-name>	(Optional) Clear the binding state for DHCP clients on the specified interface.
routing-instance <routing-instance-name>	(Optional) Clear the binding state for DHCP clients on the specified routing instance. If you do not specify a routing instance, binding state is cleared for DHCP clients on the default routing instance.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

NOTE: This command is not supported in Junos OS Evolved because it is not needed. Stateless DHCPv4 and DHCPv6 relay are supported and enabled by default, so Junos OS Evolved does not maintain states or records of statistics or bindings for DHCP clients.

RELATED DOCUMENTATION

show dhcp client binding | 19300

clear dhcp client statistics

IN THIS SECTION

- Syntax | 3486
- Description | 3486
- Options | 3487
- Required Privilege Level | 3487
- Output Fields | 3487
- Release Information | 3487

Syntax

```
clear dhcp client statistics
<all>
<interface>
<routing-instance>
```

Description

Clear all Dynamic Host Configuration Protocol (DHCP) client statistics.

Options

all	(Optional) Clear all the DHCP client statistics.
interface	(Optional) Clear the statistics for DHCP clients on the specified interface.
routing-instance	(Optional) Clear the statistics for DHCP clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

> **NOTE**: This command is not supported in Junos OS Evolved because it is not needed. Stateless DHCPv4 and DHCPv6 relay are supported and enabled by default, so Junos OS Evolved does not maintain states or records of statistics or bindings for DHCP clients.

RELATED DOCUMENTATION

show dhcp client statistics | **19305**

clear dhcp relay active-leasequery statistics

IN THIS SECTION

- Syntax | 3488
- Description | 3488
- Options | 3488
- Required Privilege Level | 3488
- Output Fields | 3489
- Sample Output | 3489
- Sample Output | 3490
- Release Information | 3491

Syntax

```
clear dhcp relay active-leasequery statistics
(interface interface-name | peer ip-address)
```

Description

Clear active leasequery statistics for the specified access interface or for the specified DHCPv4 active leasequery peer relay agent. The statistics for bindings sent, received, installed and failed are cleared.

Options

interface *interface-name* (Optional) Clear all active leasequery statistics for a specific access interface.

peer *ip-address* (Optional) Clear all active leasequery statistics for a specific peer relay agent.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

show dhcp relay active-leasequery (Interface Statistics)

```
user@host> show dhcp relay active-leasequery statistics interface ge-2/1/1.1020

     Interface : ge-2/1/1.1020
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 400
     Bindings Received                      : 320
     Bindings Installed Successfully        : 310
     Bindings Failed to install             : 10
     Last Synchronization Time              : 2019-02-15 16:20:05 IST
     ALQ Transmit Buffer count              : ffff
     Max Leasequery Transmit Rate           : 60
     Local Interface count                  : 1
     Remote Interface count                 : 1
```

clear dhcp relay active-leasequery statistics (Interface)

```
user@host> clear dhcp relay active-leasequery statistics interface ge-2/1/1.1020
```

show dhcp relay active-leasequery (Interface Statistics)

```
user@host> show dhcp relay active-leasequery statistics interface ge-2/1/1.1020

     Interface : ge-2/1/1.1020
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 0
     Bindings Received                      : 0
     Bindings Installed Successfully        : 0
     Bindings Failed to install             : 0
```

```
    Last Synchronization Time              : 2019-02-15 16:21:58 IST
    ALQ Transmit Buffer count              : ffff
    Max Leasequery Transmit Rate           : 60
    Local Interface count                  : 1
    Remote Interface count                 : 1
```

Sample Output

show dhcp relay active-leasequery (Peer Statistics)

```
user@host> show dhcp relay active-leasequery statistics peer 192.0.2.2

    peer : 192.0.2.2
    Topology-Discover Configured           : Yes
    State                                  : Done
    Bindings Sent                          : 225
    Bindings Received                      : 210
    Bindings Installed Successfully        : 210
    Bindings Failed to install             : 0
    Last Synchronization Time              : 2019-02-15 16:28:36 IST
    ALQ Transmit Buffer count              : ffff
    Max Leasequery Transmit Rate           : 60
    Local Interface count                  : 4
    Remote Interface count                 : 4
```

clear dhcp relay active-leasequery statistics (Peer)

```
user@host> clear dhcp relay active-leasequery statistics
```

show dhcp relay active-leasequery (Peer Statistics)

```
user@host> show dhcp relay active-leasequery statistics peer 192.0.2.2

    peer : 192.0.2.2
    Topology-Discover Configured           : Yes
    State                                  : Done
    Bindings Sent                          : 0
    Bindings Received                      : 0
```

```
    Bindings Installed Successfully        : 0
    Bindings Failed to install             : 0
    Last Synchronization Time              : 2019-02-15 16:30:01 IST
    ALQ Transmit Buffer count              : ffff
    Max Leasequery Transmit Rate           : 60
    Local Interface count                  : 4
    Remote Interface count                 : 4
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show dhcp relay active-leasequery | 19309

clear dhcpv6 relay active-leasequery statistics | 3524

clear dhcp relay binding (Security)

IN THIS SECTION

- Syntax | 3491
- Description | 3492
- Options | 3492
- Required Privilege Level | 3492
- Output Fields | 3492
- Release Information | 3492

Syntax

```
clear dhcp relay binding
<all | ip-address | mac-address>
```

```
<interface interface-name>
<routing-instance routing-instance-name>
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table.

Options

all	(Optional) Clear the binding state for all DHCP clients.
ip-address	(Optional) Clear the binding state for the DHCP client, using the specified IP address.
mac-address	(Optional) Clear the binding state for the DHCP client, using the specified MAC address.
interface interface-name	(Optional) Clear the binding state for DHCP clients on the specified interface
routing-instance routing-instance-name	(Optional) Clear the binding state for DHCP clients on the specified routing instance. If you do not specify a routing instance, the binding state is cleared for the default routing instance.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings

show dhcp relay binding | **19319**

clear dhcp relay binding

IN THIS SECTION

- Syntax | 3493
- Description | 3493
- Options | 3493
- Required Privilege Level | 3494
- Output Fields | 3494
- Sample Output | 3494
- Release Information | 3496

Syntax

```
clear dhcp relay binding
<address>
<all>
<dual-stack>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table.

Options

address (Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.

- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

all (Optional) Clear the binding state for all DHCP clients.

dual-stack (Optional) Clear the binding state for DHCPv4 clients and the associated DHCPv6 bindings in the single-session DHCP dual stack. DHCPv6 clients created in a DHCPv6-only stack are not affected.

interface ***interface-name*** (Optional) Clear the binding state for DHCP clients on the specified interface.

interfaces-vlan (Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard (Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).

logical-system ***logical-system-name*** (Optional) Clear the binding state for DHCP clients on the specified logical system.

routing-instance ***routing-instance-name*** (Optional) Clear the binding state for DHCP clients on the specified routing instance.

Required Privilege Level

view

Output Fields

See "show dhcp relay binding" on page 19319 for an explanation of output fields.

Sample Output

clear dhcp relay binding

The following sample output displays the address bindings in the DHCP client table before and after the `clear dhcp relay binding` command is issued.

```
user@host> show dhcp relay binding
IP address       Hardware address   Type     Lease expires at
198.51.100.32     00:00:5e:00:53:01  active   2007-02-08 16:41:17 EST
192.168.14.8      00:00:5e:00:53:02  active   2007-02-10 10:01:06 EST
```

```
user@host> clear dhcp relay binding 198.51.100.32

user@host> show dhcp relay binding
IP address       Hardware address   Type     Lease expires at
192.168.14.8      00:00:5e:00:53:02  active   2007-02-10 10:01:06 EST
```

clear dhcp relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcp relay binding all
```

clear dhcp relay binding dual-stack all

The following command clears all DHCP relay agent bindings for all DHCPv4 clients and the associated DHCPv6 bindings in the single-session DHCP dual stack. DHCPv6 clients created in a DHCPv6-only stack are not affected.

```
user@host> clear dhcp relay binding dual-stack all
```

clear dhcp relay binding interface

The following command clears DHCP relay agent bindings on a specific interface:

```
user@host> clear dhcp relay binding interface fe-0/0/3
```

clear dhcp relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP relay agent bindings on top of the underlying interface `ae0`, which clears DHCP bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcp relay binding interface ae0
```

clear dhcp relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP relay agent bindings over a specific interface:

```
user@host> clear dhcp relay binding ge-1/0/0.*
```

Release Information

Command introduced in Junos OS Release 8.3.

Options `all` and `interface` added in Junos OS Release 8.4.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Option `dual-stack` added in Junos OS Release 15.1.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings

show dhcp relay binding | **19319**

clear dhcp relay lockout-entries

IN THIS SECTION

- Syntax | **3497**
- Description | **3497**

- Options | 3497
- Required Privilege Level | 3497
- Output Fields | 3497
- Sample Output | 3498
- Release Information | 3498

Syntax

```
clear dhcp relay lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv4 relay agent lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared. The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all Clear all client entries from the lockout database.

index *index* Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcp-relay lockout-entries` command.

Required Privilege Level

view

Output Fields

See "show dhcp relay lockout-entries" on page 19328 for an explanation of the output fields.

Sample Output

clear dhcp relay lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcp relay lockout-entries` command is issued for a specific entry.

```
user@host> show dhcp relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcp relay lockout-entries index 2

user@host> show dhcp relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

clear dhcp relay lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcp relay lockout-entries` command is issued for all entries.

```
user@host> show dhcp relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcp relay lockout-entries all

user@host> show dhcp relay lockout-entries all
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show dhcp relay lockout-entries | 19328

clear dhcp relay statistics (Security)

IN THIS SECTION

- Syntax | 3499
- Description | 3499
- Options | 3499
- Required Privilege Level | 3499
- Output Fields | 3500
- Release Information | 3500

Syntax

```
clear dhcp relay statistics
<routing-instance routing-instance-name>
```

Description

Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

routing-instance routing-instance-name	(Optional) Clear the DHCP relay statistics on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

show dhcp relay statistics (Security) | 19331

clear dhcp relay statistics

IN THIS SECTION

- Syntax | 3500
- Syntax | 3501
- Description | 3501
- Options | 3501
- Required Privilege Level | 3501
- Output Fields | 3501
- Sample Output | 3502
- Release Information | 3503

Syntax

```
clear dhcp relay statistics
<bulk-leasequery-connections>
<leasequery>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Syntax

Syntax for EX Series switches:

```
show dhcp relay statistics
<routing-instance routing-instance-name>
```

Description

Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

bulk-leasequery-connections	(Optional) Clear DHCP relay bulk leasequery statistics.
leasequery	(Optional) Clear DHCP relay individual leasequery statistics.
logical-system *logical-system-name*	(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.
routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.

Required Privilege Level

view

Output Fields

See `show dhcp relay statistics` for an explanation of output fields.

Sample Output

clear dhcp relay statistics

The following sample output displays the DHCP relay statistics before and after the `clear dhcp relay statistics` command is issued.

```
user@host> show dhcp relay statistics
Packets dropped:
    Total                      1
    Lease Time Violated        1

Messages received:
    BOOTREQUEST                116
    DHCPDECLINE                0
    DHCPDISCOVER               11
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                105

Messages sent:
    BOOTREPLY                  44
    DHCPOFFER                  11
    DHCPACK                    11
    DHCPNAK                    11


user@host> clear dhcp relay statistics

user@host> show dhcp relay statistics
Packets dropped:
    Total                      0

Messages received:
    BOOTREQUEST                0
    DHCPDECLINE                0
    DHCPDISCOVER               0
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                0

Messages sent:
```

```
    BOOTREPLY                  0
    DHCPOFFER                  0
    DHCPACK                    0
    DHCPNAK                    0
```

Release Information

Command introduced in Junos OS Release 8.3.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

show dhcp relay statistics | 19334

clear dhcp-security binding

IN THIS SECTION

- Syntax | 3503
- Description | 3504
- Options | 3504
- Required Privilege Level | 3504
- Release Information | 3504

Syntax

```
clear dhcp-security binding
<interface interface-name>
<ip-address ip-address>
<statistics>
<vlan vlan-name>
```

Description

Clear the DHCP snooping database information.

Options

interface ***interface-name***	(Optional) Clear DHCP snooping database information for the specified interface.
ip-address ***ip-address***	(Optional) Clear DHCP snooping database information for the specified IP address.
statistics	(Optional) Clear all DHCP snooping database statistics.
vlan ***vlan-name***	(Optional) Clear DHCP snooping database information for the specified VLAN.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

show dhcp-security binding | **19344**

Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing

Example: Configuring IP Source Guard and Dynamic ARP Inspection on a Specified Bridge Domain to Protect the Devices Against Attacks

Port Security Features

clear dhcp-security ipv6 binding

IN THIS SECTION

- Syntax | 3505
- Description | 3505
- Options | 3505
- Required Privilege Level | 3506
- Output Fields | 3506
- Sample Output | 3506
- Release Information | 3506

Syntax

```
clear dhcp-security ipv6 binding
<all>
<interface interface-name>
<ipv6-address ipv6-address>
<vlan vlan-name>
```

Description

Clear the DHCPv6 snooping database information.

Options

all	(Optional) Clear all DHCPv6 snooping database statistics.
interface *interface-name*	(Optional) Clear DHCPv6 snooping database information for the specified interface.
ipv6-address *ipv6-address*	(Optional) Clear DHCPv6 snooping database information for the specified IPv6 address.

vlan *vlan-name*	(Optional) Clear DHCPv6 snooping database information for the specified VLAN.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear dhcp-security ipv6 binding

```
user@switch> clear dhcp-security ipv6 binding
```

Release Information

Command introduced in Junos OS Release 13.2X51-D20.

RELATED DOCUMENTATION

show dhcp-security ipv6 binding | 19351

Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing

clear dhcp server active-leasequery statistics

IN THIS SECTION

- Syntax | 3507
- Description | 3507

- Required Privilege Level | 3507
- Sample Output | 3507
- Release Information | 3508

Syntax

```
clear dhcp server active-leasequery statistics
```

Description

Clear the active leasequery statistics of the DHCP local server.

Required Privilege Level

clear

Sample Output

clear dhcp server active-leasequery statistics

```
user@host> clear dhcp server active-leasequery statistics peer 192.0.2.0
```

```
     peer : 192.0.2.0
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 0
     Bindings Received                      : 0
     Bindings Installed Successfully        : 0
     Bindings Failed to install             : 0
     Last Synchronization Time              : 1970-01-01 05:30:00 IST
     ALQ Transmit Buffer count              : ffff
     Max Leasequery Transmit Rate           : 60
     Local Interface count                  : 1
     Remote Interface count                 : 1
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

show dhcp server active-leasequery statistics | 19364

clear dhcp server binding

IN THIS SECTION

- Syntax | 3508
- Description | 3509
- Options | 3509
- Required Privilege Level | 3510
- Output Fields | 3510
- Sample Output | 3510
- Release Information | 3512

Syntax

```
clear dhcp server binding
<address>
<all>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<dual-stack>
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table on the extended DHCP local server.

> **NOTE**: If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

Options

address
(Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

all
(Optional) Clear the binding state for all DHCP clients.

interface ***interface-name***
(Optional) Clear the binding state for DHCP clients on the specified interface.

> **NOTE**: This option clears all bindings whose initial login requests were received over the specified interface. Dynamic demux login requests are not received over the dynamic demux interface, but rather the underlying interface of the dynamic demux interface. To clear a specific dynamic demux interface, use the *ip-address* or *mac-address* options.

interfaces-vlan
(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard
(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).

logical-system ***logical-system-name***
(Optional) Clear the binding state for DHCP clients on the specified logical system.

routing-instance ***routing-instance-name***
(Optional) Clear the binding state for DHCP clients on the specified routing instance.

dual-stack (Optional) Remove either both arms or single arm of dual-stack.

> **NOTE:**
>
> - The `dual-stack` command is added in the syntax removes both arms of the dual-stack with a single command entry.
> - When the `dual-stack` command is not added in the syntax, the `clear dhcpv6 server binding` command clears only the family specific arm of the dual-stack.

Required Privilege Level

view

Output Fields

See "show dhcp server binding" on page 19367 for an explanation of output fields.

Sample Output

clear dhcp server binding <ip-address>

The following sample output displays the address bindings in the DHCP client table on the extended DHCP local server before and after the `clear dhcp server binding` command is issued.

```
user@host> show dhcp server binding

2 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)

IP address       Hardware address   Type     Lease expires at
198.51.100.1      00:00:5e:00:53:01  active   2007-01-17 11:38:47 PST
198.51.100.3      00:00:5e:00:53:02  active   2007-01-17 11:38:41 PST


user@host> clear dhcp server binding 198.51.100.1

user@host> show dhcp server binding

1 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)
```

```
IP address       Hardware address   Type     Lease expires at
198.51.100.3      00:00:5e:00:53:02  active   2007-01-17 11:38:41 PST
```

clear dhcp server binding all

The following command clears all DHCP local server bindings:

```
user@host> clear dhcp server binding all
```

clear dhcp server binding interface

The following command clears DHCP local server bindings on a specific interface:

```
user@host> clear dhcp server binding interface fe-0/0/2
```

clear dhcp server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP local server bindings on top of the underlying interface `ae0`, which clears DHCP bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcp server binding ae0
```

clear dhcp server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP local server bindings over a specific interface:

```
user@host> clear dhcp server binding ge-1/0/0.*
```

clear dhcp server binding dual-stack all

The following command clears all the dual-stack local server bindings.

```
user@host> clear dhcp server binding dual-stack all
```

Release Information

Command introduced in Junos OS Release 9.0.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Command updated with `dual-stack` statement in Junos OS Release 17.3.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings

show dhcp server binding | 19367

clear dhcp server lockout-entries

IN THIS SECTION

- Syntax | 3513
- Description | 3513
- Options | 3513
- Required Privilege Level | 3513
- Output Fields | 3513
- Sample Output | 3513
- Release Information | 3514

Syntax

```
clear dhcp server lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv4 local server lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared. The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all Clear all client entries from the lockout database.

index *index* Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcp-server lockout-entries` command.

Required Privilege Level

view

Output Fields

See "show dhcp server lockout-entries" on page 19377 for an explanation of the output fields.

Sample Output

clear dhcp server lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcp server lockout-entries` command is issued for a specific entry.

```
user@host> show dhcp server lockout-entries all
Index     Key               State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
```

```
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcp server lockout-entries index 2

user@host> show dhcp server lockout-entries all
Index     Key               State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

clear dhcp server lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcp server lockout-entries` command is issued for all entries.

```
user@host> show dhcp server lockout-entries all
Index     Key               State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcp server lockout-entries all

user@host> show dhcp server lockout-entries all
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show dhcp server lockout-entries | 19377

clear dhcp server statistics

IN THIS SECTION

- Syntax | 3515
- Description | 3515
- Options | 3515
- Required Privilege Level | 3516
- Output Fields | 3516
- Sample Output | 3516
- Release Information | 3517

Syntax

```
clear dhcp server statistics
<bulk-leasequery-connections>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options

bulk-leasequery-connections	(Optional) Clear DHCPv4 local server bulk leasequery statistics.
logical-system *logical-system-name*	(Optional) Clear the statistics for DHCP clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.
routing-instance *routing-instance-name*	(Optional) Clear the statistics for DHCP clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.

Required Privilege Level

view

Output Fields

See "show dhcp server statistics" on page 19381 for an explanation of output fields.

Sample Output

clear dhcp server statistics

The following sample output displays the extended DHCP local server statistics before and after the `clear dhcp server statistics` command is issued.

```
user@host> show dhcp server statistics
Packets dropped:
    Total                       1
    Lease Time Violation        1

Messages received:
    BOOTREQUEST                 89163
    DHCPDECLINE                 0
    DHCPDISCOVER                8110
    DHCPINFORM                  0
    DHCPRELEASE                 0
    DHCPREQUEST                 81053

Messages sent:
    BOOTREPLY                   32420
    DHCPOFFER                   8110
    DHCPACK                     8110
    DHCPNAK                     8100


user@host> clear dhcp server statistics

user@host> show dhcp server statistics
Packets dropped:
    Total                       0
```

```
Messages received:
    BOOTREQUEST                0
    DHCPDECLINE                0
    DHCPDISCOVER               0
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                0

Messages sent:
    BOOTREPLY                  0
    DHCPOFFER                  0
    DHCPACK                    0
    DHCPNAK                    0
```

Release Information

Command introduced in Junos OS Release 9.0.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

clear dhcp snooping binding

IN THIS SECTION

- Syntax | 3518
- Description | 3518
- Options | 3518
- Required Privilege Level | 3518
- Output Fields | 3518
- Sample Output | 3518
- Release Information | 3518

Syntax

```
clear dhcp snooping binding
<mac (all | mac-address)>
<vlan (all | vlan-name)>
<vlan (all | vlan-name) mac (all | mac-address)>
```

Description

Clear the DHCP snooping database information.

Options

mac (all | *mac-address*) (Optional) Clear DHCP snooping information for the specified MAC address or all MAC addresses.

vlan (all | *vlan-name*) (Optional) Clear DHCP snooping information for the specified VLAN or all VLANs.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear dhcp snooping binding

```
user@switch> clear dhcp snooping binding
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show dhcp snooping binding | 19387

Example: Configuring Port Security (non-ELS)

Enabling DHCP Snooping (non-ELS)

clear dhcp snooping statistics

IN THIS SECTION

- Syntax | 3519
- Description | 3519
- Required Privilege Level | 3519
- Output Fields | 3519
- Sample Output | 3520
- Release Information | 3520

Syntax

```
clear dhcp snooping statistics
```

Description

Clear all Dynamic Host Configuration Protocol (DHCP) snooping statistics.

Required Privilege Level

view

Output Fields

See "show dhcp snooping statistics" on page 19390 for an explanation of the output fields.

Sample Output

clear dhcp snooping statistics

The following sample output displays the DHCP snooping statistics before and after the `clear dhcp snooping statistics` command is issued.

```
user@switch>                  show dhcp snooping statistics

Successful Transfers :        0   Failed Transfers :       21
Successful Reads     :        0   Failed Reads     :        0
Successful Writes    :        0   Failed Writes    :       21


user@switch> clear dhcp snooping statistics

user@switch> show dhcp snooping statistics

Successful Transfers :        0   Failed Transfers :        0
Successful Reads     :        0   Failed Reads     :        0
Successful Writes    :        0   Failed Writes    :        0
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

show dhcp snooping statistics | **19390**

Understanding DHCP Snooping (non-ELS)

clear dhcpv6 client binding

IN THIS SECTION

- Syntax | 3521
- Description | 3521
- Options | 3521
- Required Privilege Level | 3522
- Output Fields | 3522
- Release Information | 3522

Syntax

```
clear dhcpv6 client binding
[all | interface interface-name]
[routing-instance routing-instance-name]
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCPv6) client from the DHCPv6 client table.

Options

all	(Optional) Clear the binding state for all DHCPv6 clients.
interface *interface-name*	(Optional) Clear the binding state for DHCPv6 clients on the specified interface.
routing-instance *routing-instance-name*	(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance. If you do not specify a routing instance, the binding state is cleared for DHCPv6 clients on the default routing instance.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

> **NOTE**: This command is not supported in Junos OS Evolved because it is not needed. Stateless DHCPv4 and DHCPv6 relay are supported and enabled by default, so Junos OS Evolved does not maintain states or records of statistics or bindings for DHCP clients.

RELATED DOCUMENTATION

show dhcpv6 client binding | 19392

clear dhcpv6 client statistics

IN THIS SECTION

- Syntax | 3523
- Description | 3523
- Options | 3523
- Required Privilege Level | 3523
- Output Fields | 3523
- Release Information | 3523

Syntax

```
clear dhcpv6 client statistics
routing-instance routing-instance-name
```

Description

Clear all DHCPv6 client statistics.

Options

routing-instance *routing-instance-name*	(Optional) Clear the statistics for DHCPv6 clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

NOTE: This command is not supported in Junos OS Evolved because it is not needed. Stateless DHCPv4 and DHCPv6 relay are supported and enabled by default, so Junos OS Evolved does not maintain states or records of statistics or bindings for DHCP clients.

RELATED DOCUMENTATION

show dhcpv6 client statistics | **19395**

clear dhcpv6 relay active-leasequery statistics

IN THIS SECTION

- Syntax | 3524
- Description | 3524
- Options | 3524
- Required Privilege Level | 3524
- Output Fields | 3525
- Sample Output | 3525
- Sample Output | 3526
- Release Information | 3527

Syntax

```
clear dhcpv6 relay active-leasequery statistics
(interface interface-name | peer ipv6-address)
```

Description

Clear active leasequery statistics for the specified access interface or for the specified DHCPv6 active leasequery peer relay agent. The statistics for bindings sent, received, installed and failed are cleared.

Options

interface *interface-name* (Optional) Clear all active leasequery statistics for a specific access interface.

peer *ip-address* (Optional) Clear all active leasequery statistics for a specific peer relay agent.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

show dhcp relay active-leasequery (Interface Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics interface ge-0/0/0.1

     Interface : ge-0/0/0.1
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 325
     Bindings Received                      : 150
     Bindings Installed Successfully        : 150
     Bindings Failed to install             : 0
     Last Synchronization Time              : 2019-02-28 14:11:32 IST
     ALQ Transmit Buffer count              : ffff
     Max Leasequery Transmit Rate           : 60
     Local Interface count                  : 1
     Remote Interface count                 : 1
```

clear dhcp relay active-leasequery statistics (Interface)

```
user@host> clear dhcp relay active-leasequery statistics interface ge-0/0/0.1
```

show dhcpv6 relay active-leasequery (Interface Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics interface ge-0/0/0.1

     Interface : ge-0/0/0.1
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 0
     Bindings Received                      : 0
     Bindings Installed Successfully        : 0
     Bindings Failed to install             : 0
```

```
     Last Synchronization Time              : 2019-02-28 14:12:10 IST
     ALQ Transmit Buffer count              : ffff
     Max Leasequery Transmit Rate           : 60
     Local Interface count                  : 1
     Remote Interface count                 : 1
```

Sample Output

show dhcpv6 relay active-leasequery (Peer Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics peer 2001:db8::fa:2

     peer : 2001:db8::fa:2
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 420
     Bindings Received                      : 220
     Bindings Installed Successfully        : 220
     Bindings Failed to install             : 0
     Last Synchronization Time              : 2019-03-22 09:15:15 IST
     ALQ Transmit Buffer count              : ffff
     Max Leasequery Transmit Rate           : 60
     Local Interface count                  : 4
     Remote Interface count                 : 4
```

clear dhcp relay active-leasequery statistics (Peer)

```
user@host> clear dhcp relay active-leasequery statistics
```

show dhcpv6 relay active-leasequery (Peer Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics peer 2001:db8::fa:2

     peer : 2001:db8::fa:2
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 0
     Bindings Received                      : 0
```

```
    Bindings Installed Successfully        : 0
    Bindings Failed to install             : 0
    Last Synchronization Time              : 2019-03-22 09:17:05 IST
    ALQ Transmit Buffer count              : ffff
    Max Leasequery Transmit Rate           : 60
    Local Interface count                  : 4
    Remote Interface count                 : 4
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show dhcpv6 relay active-leasequery | 19398

clear dhcp relay active-leasequery statistics | 3488

clear dhcpv6 relay binding

IN THIS SECTION

- Syntax | 3527
- Description | 3528
- Options | 3528
- Required Privilege Level | 3529
- Output Fields | 3529
- Sample Output | 3529
- Release Information | 3531

Syntax

```
clear dhcpv6 relay binding
<address>
```

```
<all>
<dual-stack>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear the binding state of Dynamic Host Configuration Protocol for IPv6 (DHCPv6) clients from the client table.

Options

address (Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:

- *CID*—The specified Client ID (CID).
- *ipv6-prefix*—The specified IPv6 prefix.
- *session-id*—The specified session ID.

all (Optional) Clear the binding state for all DHCPv6 clients.

dual-stack (Optional) Clear the binding state for DHCPv6 clients and the associated DHCPv4 bindings in the single-session DHCP dual stack. DHCPv4 clients created in a DHCPv4-only stack are not affected.

interfaces-vlan (Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard (Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).

interface *interface-name* (Optional) Clear the binding state for DHCPv6 clients on the specified interface.

logical-system *logical-system-name* (Optional) Clear the binding state for DHCPv6 clients on the specified logical system.

routing-instance *routing-instance-name* (Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.

Required Privilege Level

view

Output Fields

See `show dhcpv6 relay binding` for an explanation of output fields.

Sample Output

clear dhcpv6 relay binding

The following sample output displays the DHCPv6 bindings before and after the `clear dhcpv6 relay binding` command is issued.

```
user@host> show dhcpv6 relay binding

Prefix                    Session Id  Expires  State    Interface    Client DUID
2001:db8:3c4d:15::/64    1           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
2001:db8:3c4d:16::/64    2           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:db8:3c4d:17::/64    3           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:db8:3c4d:18::/64    4           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:db8:3c4d:19::/64    5           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:db8:3c4d:20::/64    6           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06
```

clear dhcpv6 relay binding <prefix>

```
user@host> clear dhcpv6 relay binding 2001:db8:3c4d:15::/64
user@host> show dhcpv6 relay binding

Prefix                    Session Id  Expires  State    Interface    Client DUID
2001:db8:3c4d:16::/64    2           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:db8:3c4d:17::/64    3           83720    BOUND    ge-1/0/0.0
```

```
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:db8:3c4d:18::/64    4           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:db8:3c4d:19::/64    5           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:db8:3c4d:20::/64    6           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06
```

clear dhcpv6 relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcpv6 relay binding all
```

clear dhcpv6 relay binding dual-stack all

The following command clears all DHCPv6 relay agent bindings for all DHCPv6 clients and the associated DHCPv4 bindings in the single-session DHCP dual stack. DHCPv4 clients created in a DHCPv4-only stack are not affected.

```
user@host> clear dhcpv6 relay binding dual-stack all
```

clear dhcv6p relay binding interface

The following command clears DHCPv6 relay agent bindings on a specific interface:

```
user@host> clear dhcpv6 relay binding interface fe-0/0/2
```

clear dhcpv6 relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 relay agent bindings on top of the underlying interface `ae0`, which clears DHCPv6 bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcpv6 relay binding interface ae0
```

clear dhcpv6 relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 relay agent bindings over a specific interface:

```
user@host> clear dhcpv6 relay binding ge-1/0/0.*
```

Release Information

Command introduced in Junos OS Release 11.4.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Option `dual-stack` added in Junos OS Release 15.1.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings

show dhcpv6 relay binding | **19404**

clear dhcpv6 relay lockout-entries

IN THIS SECTION

Syntax | **3532**

- Description | 3532
- Options | 3532
- Required Privilege Level | 3532
- Output Fields | 3532
- Sample Output | 3533
- Release Information | 3533

Syntax

```
clear dhcpv6 relay lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv6 relay agent lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared. The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all Clear all client entries from the lockout database.

index *index* Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcpv6-relay lockout-entries` command.

Required Privilege Level

view

Output Fields

See "show dhcpv6 relay lockout-entries" on page 19417 for an explanation of the output fields.

Sample Output

clear dhcpv6 relay lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 relay lockout-entries` command is issued for a specific entry.

```
user@host> show dhcpv6 relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcpv6 relay lockout-entries index 2

user@host> show dhcpv6 relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

clear dhcpv6 relay lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 relay lockout-entries` command is issued for all entries.

```
user@host> show dhcpv6 relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcpv6 relay lockout-entries all

user@host> show dhcpv6 relay lockout-entries all
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show dhcpv6 relay lockout-entries | 19417

clear dhcpv6 relay statistics

IN THIS SECTION

- Syntax | 3534
- Description | 3534
- Options | 3534
- Required Privilege Level | 3535
- Output Fields | 3535
- Sample Output | 3535
- Release Information | 3536

Syntax

```
clear dhcpv6 relay statistics
<bulk-leasequery-connections>
<leasequery>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear all Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.

Options

bulk-leasequery-connections	(Optional) Clear DHCPv6 relay bulk leasequery statistics.
leasequery	(Optional) Clear DHCPv6 relay individual leasequery statistics.

logical-system *logical-system-name*	(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.
routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.

Required Privilege Level

view

Output Fields

See `show dhcpv6 relay statistics` for an explanation of output fields.

Sample Output

clear dhcpv6 relay statistics

The following sample output displays the DHCPv6 relay statistics before and after the `clear dhcpv6 relay statistics` command is issued.

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                       0
    Lease Time Violated         1

Messages received:
    DHCPV6_DECLINE              0
    DHCPV6_SOLICIT              10
    DHCPV6_INFORMATION_REQUEST  0
    DHCPV6_RELEASE              0
    DHCPV6_REQUEST              10
    DHCPV6_CONFIRM              0
    DHCPV6_RENEW                0
    DHCPV6_REBIND               0
    DHCPV6_RELAY_REPL           0

Messages sent:
```

```
    DHCPV6_ADVERTISE            0
    DHCPV6_REPLY                0
    DHCPV6_RECONFIGURE          0
    DHCPV6_RELAY_FORW           0
user@host> clear dhcpv6 relay statistics
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                       0

Messages received:
    DHCPV6_DECLINE              0
    DHCPV6_SOLICIT              0
    DHCPV6_INFORMATION_REQUEST  0
    DHCPV6_RELEASE              0
    DHCPV6_REQUEST              0
    DHCPV6_CONFIRM              0
    DHCPV6_RENEW                0
    DHCPV6_REBIND               0
    DHCPV6_RELAY_REPL           0

Messages sent:
    DHCPV6_ADVERTISE            0
    DHCPV6_REPLY                0
    DHCPV6_RECONFIGURE          0
    DHCPV6_RELAY_FORW           0
```

Release Information

Command introduced in Junos OS Release 11.4.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

clear dhcpv6 server active-leasequery statistics

IN THIS SECTION

Syntax | 3537

- Description | 3537
- Options | 3537
- Required Privilege Level | 3537
- Sample Output | 3538
- Release Information | 3538

Syntax

```
clear dhcpv6 server active-leasequery statistics
<peer peer-address>
<interface interface-name>
```

Description

Clear the active leasequery statistics of the DHCPv6 local server.

Options

peer *peer-address*	IP address of the peer DHCP server on which you want to view the active leasequery statistics.
interface *interface-name*	Interface name of the peer DHCP server on which you want to view the active leasequery statistics.

Required Privilege Level

clear

Sample Output

clear dhcpv6 server active-leasequery statistics peer ip-address

```
user@host> clear dhcpv6 server active-leasequery statistics peer 2001:db8:3::2
```

```
    peer : 2001:db8:3::2
    Topology-Discover Configured          : Yes
    State                                 : Done
    Bindings Sent                         : 0
    Bindings Received                     : 0
    Bindings Installed Successfully       : 0
    Bindings Failed to install            : 0
    Last Synchronization Time             : 1970-01-01 05:30:00 IST
    ALQ Transmit Buffer count             : ffff
    Max Leasequery Transmit Rate          : 60
    Local Interface count                 : 1
    Remote Interface count                : 1
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

show dhcpv6 server active-leasequery statistics | 19428

clear dhcpv6 server binding

IN THIS SECTION

- Syntax | 3539
- Description | 3539

- Options | 3539
- Required Privilege Level | 3540
- Output Fields | 3540
- Sample Output | 3540
- Release Information | 3542

Syntax

```
clear dhcpv6 server binding
<address>
<all>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<dual-stack>
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) client from the client table on the extended DHCPv6 local server.

Options

address (Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:

- *CID*—The specified Client ID (CID).
- *ipv6-prefix*—The specified IPv6 prefix.
- *session-id*—The specified session ID.

all (Optional) Clear the binding state for all DHCPv6 clients.

interface *interface-name* (Optional) Clear the binding state for DHCPv6 clients on the specified interface.

interfaces-vlan	(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.
interfaces-wildcard	(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).
logical-system *logical-system-name*	(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.
routing-instance *routing-instance-name*	(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.
dual-stack	(Optional) Remove either both arms or single arm of dual-stack.

NOTE:

- The `dual-stack` command is added in the syntax removes both arms of the dual-stack with a single command entry.
- When the `dual-stack` command is not added in the syntax, the `clear dhcpv6 server binding` command clears only the family specific arm of the dual-stack.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear dhcpv6 server binding all

The following command clears all DHCPv6 local server bindings:

```
user@host> clear dhcpv6 server binding all
```

clear dhcpv6 server binding <ipv6-prefix>

The following command clears DHCPv6 local server bindings for a specific IPv6 prefix:

```
user@host> clear dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
```

clear dhcpv6 server binding interface

The following command clears DHCPv6 local server bindings on a specific interface:

```
user@host> clear dhcpv6 server binding interface fe-0/0/2
```

clear dhcpv6 server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 local server bindings on top of the underlying interface `ae0`, which clears DHCPv6 bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcpv6 server binding interface ae0
```

clear dhcpv6 server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 local server bindings over a specific interface:

```
user@host> clear dhcpv6 server binding ge-1/0/0.*
```

clear dhcpv6 server binding dual-stack all

The following command clears all the dual-stack local server bindings.

```
user@host> clear dhcpv6 server binding dual-stack all
```

Release Information

Command introduced in Junos OS Release 9.6.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Command updated with `dual-stack` statement in Junos OS Release 17.3.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings

show dhcpv6 server binding | **19431**

clear dhcpv6 server binding (Local Server)

IN THIS SECTION

- Syntax | **3542**
- Description | **3542**
- Options | **3543**
- Required Privilege Level | **3543**
- Release Information | **3543**

Syntax

```
clear dhcpv6 server binding
<all | client-id | ip-address | session-id>
<interface interface-name>
<routing-instance routing-instance-name>
```

Description

Clear the binding state of a DHCPv6 client from the client table on the DHCPv6 local server.

Options

- all—(Optional) Clear the binding state for all DHCPv6 clients.
- *client-id*—(Optional) Clear the binding state for the DHCPv6 client with the specified client ID (option 1).
- *ip-address*—(Optional) Clear the binding state for the DHCPv6 client with the specified address.
- *session-id*—(Optional) Clear the binding state for the DHCPv6 client with the specified session ID.
- interface *interface-name*—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.
- routing-instance *routing-instance-name*—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

show dhcpv6 server binding (View) | 19440

clear dhcpv6 server lockout-entries

IN THIS SECTION

Syntax | 3544
Description | 3544
Options | 3544
Required Privilege Level | 3544

- Output Fields | 3544
- Sample Output | 3545
- Release Information | 3545

Syntax

```
clear dhcpv6 server lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv6 local server lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared. The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all Clear all client entries from the lockout database.

index *index* Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcpv6-server lockout-entries` command.

Required Privilege Level

view

Output Fields

See "show dhcpv6 server lockout-entries" on page 19446 for an explanation of the output fields.

Sample Output

clear dhcpv6 server lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 server lockout-entries` command is issued for a specific entry.

```
user@host> show dhcpv6 server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     LT          30           5200            2
2     00:00:5E:00:53:11     GT         120           780             2
3     00:00:5E:00:53:22     LT         180           2300            1

user@host> clear dhcpv6 server lockout-entries index 2

user@host> show dhcpv6 server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     LT          30           5200            2
3     00:00:5E:00:53:22     LT         180           2300            1
```

clear dhcpv6 server lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 server lockout-entries` command is issued for all entries.

```
user@host> show dhcpv6 server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     LT          30           5200            2
2     00:00:5E:00:53:11     GT         120           780             2
3     00:00:5E:00:53:22     LT         180           2300            1

user@host> clear dhcpv6 server lockout-entries all

user@host> show dhcpv6 server lockout-entries all
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show dhcpv6 server lockout-entries | 19446

clear dhcpv6 server statistics

IN THIS SECTION

- Syntax | 3546
- Description | 3546
- Options | 3546
- Required Privilege Level | 3547
- Output Fields | 3547
- Sample Output | 3547
- Release Information | 3547

Syntax

```
clear dhcpv6 server statistics
<bulk-leasequery-connections>
<interface interface-name>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear all extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.

Options

bulk-leasequery-connections (Optional) Clear DHCPv6 local server bulk leasequery statistics.

logical-system ***logical-system-name***	(Optional) Clear the statistics for DHCPv6 clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.
routing-instance ***routing-instance-name***	(Optional) Clear the statistics for DHCPv6 clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear dhcpv6 server statistics

```
user@host> clear dhcpv6 server statistics
```

Release Information

Command introduced in Junos OS Release 9.6.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

show dhcpv6 server statistics | 19450

clear dhcpv6 server statistics (Local Server)

IN THIS SECTION

- Syntax | 3548
- Description | 3548
- Options | 3548
- Required Privilege Level | 3548
- Release Information | 3549

Syntax

```
clear dhcpv6 server statistics
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear all DHCPv6 local server statistics.

Options

logical-system *logical-system-name*	(Optional) Clear the statistics for DHCPv6 clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.
routing-instance *routing-instance-name*	(Optional) Clear the statistics for DHCPv6 clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

show dhcpv6 server statistics (View) | 19456

clear dhcpv6 snooping binding

IN THIS SECTION

- Syntax | 3549
- Description | 3549
- Options | 3550
- Required Privilege Level | 3550
- Output Fields | 3550
- Sample Output | 3550
- Release Information | 3550

Syntax

```
clear dhcpv6 snooping binding
<mac (all | mac-address)>
<vlan (all | vlan-name)>
<vlan (all | vlan-name) mac (all | mac-address)>
```

Description

Clear the DHCPv6 snooping database information.

Options

mac (all | *mac-address*) (Optional) Clear DHCPv6 snooping information for the specified MAC address or all MAC addresses.

vlan (all | *vlan-name*) (Optional) Clear DHCPv6 snooping information for the specified VLAN or all VLANs.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear dhcpv6 snooping binding

```
user@switch> clear dhcpv6 snooping binding
```

Release Information

Command introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

show dhcpv6 snooping binding | **19460**

Example: Configuring Port Security (non-ELS)

Enabling DHCP Snooping (non-ELS)

clear dhcpv6 snooping statistics

IN THIS SECTION

- Syntax | 3551
- Description | 3551
- Required Privilege Level | 3551
- Output Fields | 3551
- Sample Output | 3552
- Release Information | 3552

Syntax

```
clear dhcpv6 snooping statistics
```

Description

Clear all Dynamic Host Configuration Protocol for IPv6 (DHCPv6) snooping statistics.

Required Privilege Level

view

Output Fields

See "show dhcpv6 snooping statistics" on page 19463 for an explanation of the output fields.

Sample Output

clear dhcpv6 snooping statistics

The following sample output displays the DHCPv6 snooping statistics before and after the `clear dhcpv6 snooping statistics` command is issued.

```
user@switch>                  show dhcpv6 snooping statistics

Successful Transfers :        0   Failed Transfers :       21
Successful Reads     :        0   Failed Reads     :        0
Successful Writes    :        0   Failed Writes    :       21


user@switch> clear dhcpv6 snooping statistics

user@switch> show dhcpv6 snooping statistics

Successful Transfers :        0   Failed Transfers :        0
Successful Reads     :        0   Failed Reads     :        0
Successful Writes    :        0   Failed Writes    :        0
```

Release Information

Command introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

> show dhcpv6 snooping statistics | **19463**
>
> *Understanding DHCP Snooping (non-ELS)*

clear diameter function statistics

IN THIS SECTION

- Syntax | 3553
- Description | 3553
- Options | 3553
- Required Privilege Level | 3553
- Output Fields | 3553
- Sample Output | 3554
- Release Information | 3554

Syntax

```
clear diameter function <function-name> statistics
```

Description

Clear current statistics accumulated for a specified function (application) or for all functions associated with the Diameter instance.

Options

function-name (Optional) Clear statistics for the specified function. Gx-Plus, JSRC, and packet-triggered-subscribers are supported functions.

Required Privilege Level

clear

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear diameter function statistics

```
user@host> clear diameter function jsrc statistics
```

Release Information

Command introduced in Junos OS Release 9.6.

Support for PTSP introduced in Junos OS Release 10.2.

Support for Gx-Plus introduced in Junos OS Release 11.2.

Support for PTSP discontinued in Junos OS Release 13.1.

RELATED DOCUMENTATION

Gx-Plus for Provisioning Subscribers

JSRC for Subscriber Provisioning and Accounting

show diameter | 19472

show diameter function | 19482

show diameter function statistics | 19488

clear diameter peer

IN THIS SECTION

- Syntax | 3555
- Description | 3555
- Options | 3555
- Required Privilege Level | 3555
- Output Fields | 3555
- Sample Output | 3555

Release Information | 3556

Syntax

```
clear diameter peer peer-name
<connection | statistics>
```

Description

Delete the specified Diameter peer and clear all statistics or only current statistics for the specified peer.

Options

peer-name Delete the Diameter peer.

connection (Optional) Clear all peer statistics and restart the peer state machine for the specified Diameter peer. This is the default action.

statistics (Optional) Clear current statistics for the specified Diameter peer.

Required Privilege Level

clear

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear diameter peer

```
user@host> clear diameter peer peer5 connection
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- show diameter | 19472
- show diameter peer | 19504
- show diameter peer map | 19510
- show diameter peer statistics | 19514

clear-dont-fragment-bit (fti)

IN THIS SECTION

- Syntax | 3556
- Hierarchy Level | 3556
- Description | 3557
- Required Privilege Level | 3557
- Release Information | 3557

Syntax

```
clear-dont-fragment-bit;
```

Hierarchy Level

```
[edit interfaces fti number unit logical-unit-number family inet address]
```

Description

Clear the Don't Fragment (DF) bit on all IP version 4 (IPv4) packets entering the (UDP/GRE/IP-in-IP) tunnel on FTI interfaces. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation.

When the option is not configured and the DF bit is set, the packets will be dropped if the packet size is larger than the MTU of the FTI.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

encapsulation

clear-dont-fragment-bit (Interfaces GRE Tunnels)

IN THIS SECTION

- Syntax | 3558
- Hierarchy Level | 3558
- Description | 3558
- Required Privilege Level | 3558
- Release Information | 3558

Syntax

```
clear-dont-fragment-bit;
```

Hierarchy Level

```
[edit interfaces gr-fpc/pic/port unit logical-unit-number],
[edit logical-systems logical-system-name interfaces gr-fpc/pic/port unit logical-unit-number]
```

Description

Clear the Don't Fragment (DF) bit on all IP version 4 (IPv4) packets entering the generic routing encapsulation (GRE) tunnel on Adaptive Services (AS) or Multiservices interfaces. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. The statement is supported only on MX Series routers and all M Series routers except the M320 router.

When you configure the `clear-dont-fragment-bit` statement on an interface with the MPLS protocol family enabled, you must specify an MTU value. This MTU value must not be greater than maximum supported value, which is 9192.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 10.2 for MPC1 and MPC1Q.

> **NOTE**: On SRX platforms the clearing of the DF bit on a GRE tunnel is supported only when the device is in packet or selective packet mode; This feature is not supported in flow mode. As a result, when in flow mode, a packet that exceeds the MTU of the GRE interface with the DF bit set is dropped, despite having the `clear-dont-fragment-bit` configured on the GRE interface.

RELATED DOCUMENTATION

Enabling Fragmentation on GRE Tunnels

clear-dont-fragment-bit (NAT Next Gen Services)

IN THIS SECTION

- Syntax | 3559
- Hierarchy Level | 3559
- Description | 3559
- Required Privilege Level | 3559
- Release Information | 3560

Syntax

```
set clear-dont-fragment-bit;
```

Hierarchy Level

```
[edit services nat natv6v4]
```

Description

Specify that the don't fragment (DF) bit for IPv4 packet headers is cleared when the packet length is less than 1280 bytes. Use this statement when configuring stateful NAT64, deterministic NAPT64, and 464XLAT. This prevents unnecessary creation of an IPv6 fragmentation header when translating IPv4 packets that are less than 1280 bytes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

clear-dont-fragment-bit (Services IPsec VPN)

IN THIS SECTION

- Syntax | 3560
- Hierarchy Level | 3560
- Description | 3560
- Required Privilege Level | 3561
- Release Information | 3561

Syntax

```
clear-dont-fragment-bit;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Clear the do not fragment (DF) bit on all IP version 4 (IPv4) packets entering the IPsec tunnel. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation.

By default, this statement is disabled (the DF bit value is not cleared on the inner header and outer header by default).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

clear-dont-fragment-bit (Services NAT Options)

IN THIS SECTION

- Syntax | 3561
- Hierarchy Level | 3561
- Description | 3562
- Required Privilege Level | 3562
- Release Information | 3562

Syntax

```
clear-dont-fragment-bit;
```

Hierarchy Level

```
[edit services service-set service-set-name nat-options stateful-nat64]
```

Description

Clear the DF (don't fragment) bit in a translated IPv4 packet if its packet size is less than 1280 bytes. If the packet is greater than or equal to 1280 bytes, the DF bit is not cleared.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced with Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Service Rules

clear-dont-fragment-bit (Services Service Set)

IN THIS SECTION

- Syntax | 3562
- Hierarchy Level | 3563
- Description | 3563
- Required Privilege Level | 3563
- Release Information | 3563

Syntax

```
clear-dont-fragment-bit;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Clear the Don't Fragment (DF) bit on all IP version 4 (IPv4) packets entering the IPsec tunnel. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. This statement is useful for dynamic endpoint tunnels, for which you cannot configure the `clear-dont-fragment-bit` statement at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level.

For static IPsec tunnels, setting this statement clears the DF bit on packets entering all the static tunnels within this service set. If you want to clear the DF bit on packets entering a specific tunnel, set the `clear-dont-fragment-bit` statement at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level.

By default, this statement is disabled (the DF bit value is not cleared on the inner header and outer header by default).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring IPsec Service Sets

Configuring IPsec Rules

clear dot1x

IN THIS SECTION

- Syntax | 3564
- Description | 3564
- Options | 3564
- Required Privilege Level | 3565
- Sample Output | 3565
- Release Information | 3565

Syntax

```
clear dot1x
interface  <interface-name>
mac-address <static-mac-address>
statistics <interface interface-name>
```

Description

Reset the authentication state of an interface or delete 802.1X statistics from the device. When you reset an interface using the `interface` or `mac-address` options, reauthentication on the interface is also triggered. The device sends out a multicast message on the interface to restart the authentication of all connected supplicants. If a MAC address is reset, then the device sends out a unicast message to that specific MAC address to restart authentication.

Options

interface <[*interface-name*]>	Reset the authentication state of all the supplicants (also, clear all the authentication bypassed clients) connected to the specified interface (when the interface is an authenticator) or reset the authentication state for the interface itself (when the interface is a supplicant).
mac-address [*mac-addresses*]	Reset the authentication state of the specified MAC addresses.

statistics <interface *interface-name*> Clear 802.1X statistics on all 802.1X-enabled interfaces. If the `interface` option is specified, clear 802.1X firewall statistics for that interface or interfaces.

Required Privilege Level

view

Sample Output

clear dot1x interface

```
user@host> clear dot1x interface ge-0/0/1
```

clear dot1x mac-address

```
user@host> clear dot1x mac-address 00:04:ae:cd:23:5f
```

clear dot1x statistics interface

```
user@host> clear dot1x statistics interface ge-0/0/1
```

Release Information

Command introduced in Junos OS Release 15.1X49-D80.

RELATED DOCUMENTATION

dot1x | 5917

clear edge-virtual-bridging

IN THIS SECTION

- Syntax | 3566
- Description | 3566
- Options | 3566
- Required Privilege Level | 3566
- Release Information | 3567

Syntax

```
clear edge-virtual-bridging
    <edge-control-protocol-statistics>
    <firewall <interface interface-name>
<vsi-profiles <interface interface-name>
```

Description

Clear edge-virtual-bridging (EVB).

Options

none	Clear EVB.
edge-control-protocol-statistics	(Optional) Clear Edge Control Protocol (ECP) statistics.
firewall <interface *interface-name*>	(Optional) Clear EVB implicit filter counters on all interfaces or on a specific interface.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on an EX Series Switch

Configuring Edge Virtual Bridging on an EX Series Switch

clear error bpdu interface

IN THIS SECTION

- MX Series | 3567
- QFX Series, EX Series, NFX Series | 3567
- Description | 3568
- Required Privilege Level | 3568
- Output Fields | 3568
- Sample Output | 3568
- Release Information | 3568

MX Series

```
clear error bpdu interface interface-name
```

QFX Series, EX Series, NFX Series

```
clear error bpdu interface (all | interface-name)
```

Description

Clear a bridge protocol data unit (BPDU) error condition caused by the detection of a possible bridging loop from Spanning Tree Protocol (STP) operation.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear error bpdu interface ge-1/1/1

```
user@host> clear error bpdu interface ge-1/1/1
```

Release Information

Command introduced in Junos OS Release 9.4.

Command supports `all` option in Junos OS Release 15.1 for EX Series switches.

RELATED DOCUMENTATION

Configuring BPDU Protection on ACX Router, EX Switch and MX Router Edge Ports

Unblocking a Switch Interface That Receives BPDUs in Error (CLI Procedure)

clear error mac-rewrite

IN THIS SECTION

Syntax | 3569

- Description | 3569
- Options | 3569
- Required Privilege Level | 3569
- Output Fields | 3569
- Sample Output | 3570
- Release Information | 3570

Syntax

```
clear error mac-rewrite
<interface interface-name>
```

Description

Clear a MAC rewrite error condition on an interface receiving tunneled Layer 2 protocol packets.

On interfaces with L2PT configured, customer-facing ports should not receive packets with the L2PT MAC address as the destination address unless there is a network tolopogy or configuration error. The device sets the status of such interfaces to “Disabled”. Use this command to clear the error and re-enable the interface.

Options

interface *interface-name* (Optional) Clear the MAC rewrite error condition for the specified interface.

Required Privilege Level

clear

Output Fields

When you enter this command, the device returns feedback on the status of the request.

Sample Output

clear error mac-rewrite interface

```
user@host> clear error mac-rewrite interface ge-1/0/1
```

Release Information

Command introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Configuring Layer 2 Protocol Tunneling

Clearing a MAC Rewrite Error on an Interface with Layer 2 Protocol Tunneling

show mac-rewrite interface | 21993

clear ethernet-switching bpdu-error interface

IN THIS SECTION

- Syntax | 3570
- Description | 3571
- Options | 3571
- Required Privilege Level | 3571
- Sample Output | 3571
- Release Information | 3571

Syntax

```
clear ethernet-switching bpdu-error interface interface-name
```

Description

Clear bridge protocol data unit (BPDU) errors from an interface and bring up the interface.

> **NOTE**: This command is only available on switches running Junos OS without support for the Enhanced Layer 2 Software (ELS) configuration style. To clear BPDU errors from switches that support ELS, see "clear error bpdu interface" on page 3567.

Options

interface-name	Clear BPDU errors on the specified interface.

Required Privilege Level

clear

Sample Output

clear ethernet-switching bpdu-error interface

```
user@switch> clear ethernet-switching bpdu-error interface xe-0/0/1.0
```

Release Information

Command introduced in Junos OS Release 9.1.

Command updated in Junos OS Release 11.1 for EX Series switches—a BPDU error shuts down the interface and this command brings the interface back up.

RELATED DOCUMENTATION

show spanning-tree interface | 26555

clear ethernet-switching evpn arp-table

IN THIS SECTION

- Syntax | 3572
- Description | 3572
- Options | 3572
- Required Privilege Level | 3572
- Output Fields | 3572
- Sample Output | 3573
- Release Information | 3573

Syntax

```
clear ethernet-switching evpn arp-table
```

Description

Clear the ARP proxy table from the Ethernet VPN (EVPN) for integrated routing and bridging (IRB) interfaces. This command applies to EVPN instances of type `virtual-switch`.

Options

none Clear information about the ARP proxy tables for an EVPN.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ethernet-switching evpn arp-table

```
user@host> clear ethernet-switching evpn arp-table
```

Release Information

Command introduced in Junos OS Release 17.3R1.

clear ethernet-switching evpn nd-statistics

IN THIS SECTION

- Syntax | 3573
- Description | 3573
- Options | 3574
- Required Privilege Level | 3574
- Output Fields | 3574
- Sample Output | 3574
- Release Information | 3574

Syntax

```
clear ethernet-switching evpn nd-statistics
```

Description

Clear the Neighbor Discovery (ND) proxy table statistics for integrated routing and bridging (IRB) interfaces participating in the Ethernet VPN (EVPN). ND proxy is a kernel module that implements IPv6 Neighbor Discovery proxying over Ethernet-like access networks.

Options

none Clear the ND proxy table statistics for IRB interfaces participating in the EVPN.

Required Privilege Level

clear

Output Fields

When you enter this command, the ND proxy table statistics for IRB interfaces participating in the EVPN are cleared.

Sample Output

clear ethernet-switching evpn nd-statistics

```
user@host> clear ethernet-switching evpn nd-statistics
```

Release Information

Command introduced in Junos OS Release 17.3R1.

clear ethernet-switching evpn nd-table

IN THIS SECTION

- Syntax | 3575
- Description | 3575
- Options | 3575
- Required Privilege Level | 3575
- Output Fields | 3575
- Sample Output | 3575

Release Information | 3575

Syntax

```
clear ethernet-switching evpn nd-table
```

Description

Clear the Neighbor Discovery (ND) proxy table from the Ethernet VPN (EVPN) for integrated routing and bridging (IRB) interfaces. This command applies to EVPN instances of type `virtual-switch`.

Options

none Clear information about the ND proxy tables for an EVPN.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ethernet-switching evpn nd-table

```
user@host> clear ethernet-switching evpn nd-table
```

Release Information

Command introduced in Junos OS Release 17.3R1.

clear ethernet-switching layer2-protocol-tunneling error

IN THIS SECTION

- Syntax | 3576
- Description | 3576
- Options | 3576
- Required Privilege Level | 3576
- Sample Output | 3577
- Release Information | 3577

Syntax

```
clear ethernet-switching layer2-protocol-tunneling error
<interface interface-name>
```

Description

Clear Layer 2 protocol tunneling (L2PT) errors on one or more interfaces. If an interface has been disabled because the amount of Layer 2 protocol traffic exceeded the shutdown threshold or because the switch has detected an error in the network topology or configuration, use this command to reenable the interface.

Options

none	Clears L2PT errors on all interfaces.
interface *interface-name*	(Optional) Clear L2PT errors on the specified interface.

Required Privilege Level

view

Sample Output

clear ethernet-switching layer2-protocol-tunneling error

```
user@switch> clear ethernet-switching layer2-protocol-tunneling error
```

clear ethernet-switching layer2-protocol-tunneling error interface xe-0/0/1.0

```
user@switch> clear ethernet-switching layer2-protocol-tunneling error interface xe-0/0/1.0
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

clear ethernet-switching layer2-protocol-tunneling statistics

IN THIS SECTION

- Syntax | 3578
- Description | 3578
- Options | 3578
- Required Privilege Level | 3578
- Sample Output | 3578
- Release Information | 3579

Syntax

```
clear ethernet-switching layer2-protocol-tunneling statistics
<interface interface-name>
<vlan vlan-name>
```

Description

Clear Layer 2 protocol tunneling (L2PT) statistics on one or more interfaces or VLANs.

Options

none	Clear L2PT statistics on all interfaces and VLANs.
interface *interface-name*	(Optional) Clear L2PT statistics on the specified interface.
vlan *vlan-name*	(Optional) Clear L2PT statistics on the specified VLAN.

Required Privilege Level

view

Sample Output

clear ethernet-switching layer2-protocol-tunneling statistics

```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics
```

clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0

```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics interface xe-0/1/1.0
```

clear ethernet-switching layer2-protocol-tunneling error vlan v2

```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics vlan v2
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show ethernet-switching layer2-protocol-tunneling statistics | 19654

Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

clear ethernet-switching mac-ip-table

IN THIS SECTION

- Syntax | 3579
- Description | 3579
- Options | 3580
- Required Privilege Level | 3580
- Release Information | 3580

Syntax

```
clear ethernet-switching mac-ip-table
<address>
<vlan-name (all | vlan-domain-name)>
<instance instance-name>
```

Description

Clear the locally learnt ARP entries for VLANs in routing instances where the instance-type is ethernet switching. Remotely learnt entries cannot be cleared.

Options

address	(Optional) Clear the specific MAC address from the MAC-IP address table .
vlan-name (all | *vlan-domain-name*)	(Optional) Clear the MAC-IP address table for all VLANs or for the specified VLAN.
instance *instance-name*	(Optional) Clear the MAC-IP address table for a specific routing instance.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 17.4R2.

clear ethernet-switching port-error

IN THIS SECTION

- Syntax | 3580
- Description | 3581
- Options | 3581
- Required Privilege Level | 3581
- Output Fields | 3581
- Sample Output | 3581
- Release Information | 3581

Syntax

```
clear ethernet-switching port-error
<interface interface-name>
```

Description

Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch or from the specified interface, and restore all interfaces or the specified interface to service.

Options

none	Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch and restore these interfaces to service.
interface *interface-name*	(Optional) Clear all MAC limiting, MAC move limiting, and storm control errors from the specified interface and restore the interface to service.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear ethernet-switching port-error

```
user@switch> clear ethernet-switching port-error
```

Release Information

Command introduced in JUNOS Release 9.6.

RELATED DOCUMENTATION

Configuring Autorecovery for Port Security Events

Configuring Port Security (non-ELS)

Configuring Autorecovery for Port Security Events

clear ethernet-switching recovery-timeout

IN THIS SECTION

- Syntax | 3582
- Description | 3582
- Options | 3582
- Required Privilege Level | 3582
- Output Fields | 3582
- Release Information | 3583

Syntax

```
clear ethernet-switching recovery-timeout
```

Description

Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch, and restore the interfaces to service.

Options

`interface` *`interface-name`* `vlan` *`vlan-name`*—(EX9200 switches) Unblock an interface on the basis of its membership in the specified VLAN. This option can be used to restore an interface that is blocked because of a `vlan-member-shutdown` action.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Configuring Autorecovery for Port Security Events

clear ethernet-switching recovery-timeout (SRX)

IN THIS SECTION

- Syntax | 3583
- Description | 3583
- Options | 3584
- Required Privilege Level | 3584
- Release Information | 3584

Syntax

```
clear ethernet-switching recovery-timeout
<interface interface-name>
```

Description

Clear all MAC limiting errors from all the Ethernet switching interfaces on the device or from the specified interface, and restore the interfaces or the specified interface to service. You can also use this command to manually recover interfaces that duplicate MAC detection loop protection has shut down.

To manually recover the interface that was shut down,

Options

interface *interface-name*	(Optional) Clear all MAC limiting errors from the specified interface and restore the interface to service.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 15.1X49-D70

RELATED DOCUMENTATION

- Understanding MAC Limiting
- Example: Configuring MAC Limiting on a Security Device
- Configuring Autorecovery From the Disabled State on Secure Interfaces of a Security Device (CLI Procedure)

clear ethernet-switching statistics mac-learning

IN THIS SECTION

- Syntax | 3585
- Description | 3585
- Options | 3585
- Required Privilege Level | 3585
- Sample Output | 3585
- Release Information | 3585

Syntax

clear ethernet-switching statistics mac-learning

Description

Clear the media access control (MAC) learning statistics.

Options

- `none`—Clear MAC learning statistics on all interfaces.
- `interface` *`interface-name`*—(Optional) Clear MAC learning statistics on the specified interface.

Required Privilege Level

view

Sample Output

clear ethernet-switching statistics mac-learning

```
user@host> clear ethernet-switching statistics mac-learning
```

clear ethernet-switching statistics mac-learning interface interface-name

```
user@host> clear ethernet-switching statistics mac-learning interface interface-name
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

show ethernet-switching table | 19700

clear ethernet-switching table persistent-learning

IN THIS SECTION

- Syntax | 3586
- Description | 3586
- Options | 3586
- Required Privilege Level | 3587
- Output Fields | 3587
- Sample Output | 3587
- Release Information | 3587

Syntax

```
clear ethernet-switching table
mac <mac-address> | interface <interface-name> | vlan-name <vlan-name> | vlan-id <vlan-id> |
instance <instance>
<persistent-learning (interface <interface-name> | mac <mac-address>)>
```

Description

NOTE: On a QFabric system, using this command on an FCoE-enabled VLAN when FCoE sessions are active can cause traffic flooding and FCoE traffic drop. The FCoE sessions are not terminated and the traffic reconverges after a short period of time.

Clear learned entries, which are media access control (MAC) addresses, in the Ethernet switching table (also called the forwarding database table).

Options

none Clear learned entries in the Ethernet switching table, except for persistent MAC addresses.

interface *interface-name*	(Optional) Clear all learned MAC addresses for the specified interface from the Ethernet switching table.
mac *mac-address*	(Optional) Clear the specified learned MAC address from the Ethernet switching table.
persistent-learning *interface* | *mac-address*	(Optional) Clear all persistent MAC addresses. Use the **interface** option to clear all persistent MAC addresses on the interface, or use the **mac-address** option to clear a specific persistent MAC address. Use this command whenever you move a device in your network that has a persistent MAC address on the switch. If you move the device to another port on the switch and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address and the device will not be able to connect.
vlan *vlan-name*	(Optional) Clear all MAC addresses learned for the specified VLAN from the Ethernet switching table.

Required Privilege Level

view

Output Fields

This command produces no output.

Sample Output

clear ethernet-switching table

```
user@switch> clear ethernet-switching table
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show ethernet-switching table | 19700

clear evpn duplicate-mac-suppression

IN THIS SECTION

- Syntax | 3588
- Description | 3588
- Options | 3588
- Required Privilege Level | 3589
- Release Information | 3589

Syntax

```
clear evpn duplicate-mac-suppression
<instance instance |l2-domain-id l2-domain-id | logical-system (all |logical-system-name) | mac-address mac-address>
```

Description

Clear suppressed duplicate MAC addresses in an EVPN network. You can also use this command to re-enable any blocked duplicate MAC addresses. After you have enabled these duplicate MAC addresses, these MAC addresses will be available for relearning and forwarding.

Options

instance *instance*	(Optional) Clear suppressed MAC address for a specific routing instance. On MX Series routers , the routing instance can be an EVPN instance or virtual-switch instance. On QFX Series switches, the routing instance can be an EVPN instance, virtual-switch instance or an implicit default-switch instance.
l2-domain-id *l2-domain-id*	(Optional) Clear suppressed MAC address for a specific L2 domain.

logical-system (all | *logical-system-name*)	(Optional) Clear suppressed MAC address on all logical systems or on a specific logical system.
mac-address *mac-address*	(Optional) Clear a specific suppressed MAC address.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

Overview of MAC Mobility

Changing Duplicate MAC Address Detection Settings

clear evpn mac-ip-table

IN THIS SECTION

- Syntax | 3589
- Description | 3590
- Options | 3590
- Required Privilege Level | 3590
- Release Information | 3590

Syntax

```
clear evpn mac-ip-table
<address>
```

```
<instance instance-name>
<logical-system (all | logical-system-name>
```

Description

Clear the locally learned ARP entries for routing instances where the instance-type is evpn. Remotely learned entries cannot be cleared.

Options

address	(Optional) Clear the specific MAC address from the MAC-IP address table.
instance *instance-name*	(Optional) Clear the MAC-IP address table for a specific routing instance.
logical-system (all | *logical-system-name***)**	(Optional) Clear the MAC-IP address for all logical systems or on a specific logical system.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 17.4R2.

clear evpn nd-table

IN THIS SECTION

- Syntax | 3591
- Description | 3591
- Options | 3591
- Required Privilege Level | 3591
- Output Fields | 3591
- Sample Output | 3591

Release Information | 3591

Syntax

```
clear evpn nd-table
```

Description

Clear the Neighbor Discovery (ND) proxy table from the Ethernet VPN (EVPN). ND proxy is a kernel module that implements IPv6 Neighbor Discovery proxying over Ethernet-like access networks.

Options

none Clear learned ND proxy tables from the EVPN.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear evpn nd-table

```
user@host> clear evpn nd-table
```

Release Information

Command introduced in Junos OS Release 16.2.

clear evpn statistics

IN THIS SECTION

- Syntax | 3592
- Description | 3592
- Options | 3592
- Required Privilege Level | 3592
- Output Fields | 3593
- Sample Output | 3593
- Release Information | 3593

Syntax

```
clear evpn statistics instance evpn-instance
<logical-system (all | logical-system-name)>
```

Description

Clear EVPN statistics of all interfaces in a routing instance of type EVPN.

Options

evpn-instance	Name of the EVPN routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

maintenance

Output Fields

This command produces no output.

Sample Output

clear evpn statistics

```
user@host> clear evpn statistics instance evpn-instance
```

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

clear bridge statistics | 3447

clear extensible-subscriber-services counters

IN THIS SECTION

- Syntax | 3593
- Description | 3594
- Required Privilege Level | 3594
- Sample Output | 3594
- Release Information | 3594

Syntax

```
clear extensible-subscriber-services counters
```

Description

Clear values of all event counters to zero. Active sessions and services counters are not reset.

Required Privilege Level

view

Sample Output

command-name

```
#clear extensible-subscriber-services counters
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

show extensible-subscriber-services counters | 19849

clear extensible-subscriber-services sessions

IN THIS SECTION

- Syntax | 3595
- Description | 3595
- Options | 3595
- Required Privilege Level | 3595
- Sample Output | 3595
- Release Information | 3595

Syntax

```
clear extensible-subscriber-services sessions <accounting-session-id>
```

Description

Clear extensible subscriber service sessions in any state by executing the operational script or the application to remove all services created. Specify an accounting session ID to clear a specific session. If you do not specify an accounting session ID in the command, the command clears all sessions.

Options

accounting-session-id	(Optional) Identifier of the ESSM session you want cleared.

Required Privilege Level

view

Sample Output

command-name

```
#clear extensible-subscriber-services sessions "jnpr demux0.1073762028:46422"
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

show extensible-subscriber-services sessions | 19869

clear fibre-channel fc2 statistics

IN THIS SECTION

- Syntax | 3596
- Description | 3596
- Options | 3596
- Required Privilege Level | 3596
- Sample Output | 3597
- Release Information | 3597

Syntax

```
clear fibre-channel fc2 statistics
<fabric fabric-name>
```

Description

Clear FC-2 (network layer) Fibre Channel statistics globally or on a specified Fibre Channel fabric.

Options

fabric *fabric-name*	(Optional) Clear FC-2 statistics only on the specified fabric.

Required Privilege Level

view

Sample Output

clear fibre-channel fc2 statistics

```
user@switch> clear fibre-channel fc2 statistics
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fc2 statistics | 19936

show fibre-channel fc2 sessions | 19932

clear fibre-channel fip enode

IN THIS SECTION

- Syntax | 3597
- Description | 3598
- Options | 3598
- Required Privilege Level | 3598
- Sample Output | 3598
- Release Information | 3598

Syntax

```
clear fibre-channel fip enode enode-mac
```

Description

Clear Fibre Channel over Ethernet (FCoE) node (ENode) information for a specified ENode. This operation deletes the ENode state from the switch database and from the FIP snooping firewall filters, which causes the ENode to lose the connection to the Fibre Channel (FC) fabric and to log in to the fabric again. If you clear an ENode, all VN_Ports associated with that ENode are also cleared and lose their connection to the FC fabric and must log in to the fabric again.

Options

enode-mac	MAC address of the ENode.

Required Privilege Level

view

Sample Output

clear fibre-channel fip enode

```
user@switch> clear fibre-channel fip enode 00:10:94:00:00:02
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- show fibre-channel fip enode | 19946
- clear fibre-channel fip statistics | 3599
- clear fibre-channel fip vn-port | 3600

clear fibre-channel fip statistics

IN THIS SECTION

- Syntax | 3599
- Description | 3599
- Options | 3599
- Required Privilege Level | 3599
- Sample Output | 3600
- Release Information | 3600

Syntax

```
clear fibre-channel fip statistics
<fabric fabric-name>
```

Description

Clear Fibre Channel over Ethernet (FCoE) initialization protocol (FIP) statistics.

Options

fabric *fabric-name*	(Optional) Clear FIP statistics only on the specified fabric.

Required Privilege Level

view

Sample Output

clear fibre-channel fip statistics

```
user@switch> clear fibre-channel fip statistics
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fip statistics | 19966

show fibre-channel fip | 19938

clear fibre-channel fip vn-port

IN THIS SECTION

- Syntax | 3600
- Description | 3601
- Options | 3601
- Required Privilege Level | 3601
- Sample Output | 3601
- Release Information | 3601

Syntax

```
clear fibre-channel fip vn-port vn-port--mac
```

Description

Clear virtual N_Port (VN_Port) information for a specified VN_Port. This operation deletes the VN_Port state from the switch database and from the FIP snooping firewall filters, which causes the VN_Port to lose its connection to the Fibre Channel fabric and to log in to the fabric again. When you clear a VN_Port, other VN_Ports associated with the same Fibre Channel over Ethernet (FCoE) Node (ENode) are not affected and are not cleared.

Options

vn-port-mac	MAC address of the VN_Port.

Required Privilege Level

view

Sample Output

clear fibre-channel fip vn-port

```
user@switch> clear fibre-channel fip vn-port 00:10:94:00:00:08
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fip enode | 19946

clear fibre-channel fip enode | 3597

clear fibre-channel fip statistics | 3599

clear fibre-channel flogi statistics

IN THIS SECTION

- Syntax | 3602
- Description | 3602
- Options | 3602
- Required Privilege Level | 3602
- Sample Output | 3603
- Release Information | 3603

Syntax

```
clear fibre-channel flogi statistics
<fabric fabric-name>
```

Description

Clear fabric login (FLOGI) statistics globally or on a specified Fibre Channel fabric.

Options

fabric *fabric-name*	(Optional) Clear FLOGI statistics only on the specified fabric.

Required Privilege Level

view

Sample Output

clear fibre-channel flogi statistics

```
user@switch> clear fibre-channel flogi statistics
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- show fibre-channel flogi statistics | 19977
- show fibre-channel flogi fport | 19971
- show fibre-channel flogi nport | 19973

clear fibre-channel proxy statistics

IN THIS SECTION

- Syntax | 3603
- Description | 3604
- Options | 3604
- Required Privilege Level | 3604
- Sample Output | 3604
- Release Information | 3604

Syntax

```
clear fibre-channel proxy statistics
<fabric fabric-name>
```

Description

Clear Fibre Channel gateway statistics globally or on a specified Fibre Channel fabric.

Options

`fabric` *`fabric-name`* (Optional) Clear proxy statistics only on the specified fabric.

Required Privilege Level

view

Sample Output

clear fibre-channel proxy statistics

```
user@switch> clear fibre-channel proxy statistics
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- show fibre-channel proxy statistics | 20003
- show fibre-channel proxy login-table | 19993
- show fibre-channel proxy np-port | 19997

clear fip snooping enode

IN THIS SECTION

- Syntax | 3605

- Syntax (Junos Fusion) | 3605
- Description | 3605
- Options | 3606
- Required Privilege Level | 3606
- Sample Output | 3606
- Release Information | 3606

Syntax

```
                  clear fip snooping enode enode-mac
<vlan vlan-name>
```

Syntax (Junos Fusion)

```
                  clear fip snooping satellite enode enode-mac
<vlan vlan-name>
```

Description

Clear FIP snooping information for the specified FCoE Node (ENode) or (optionally) only on a specified FCoE VLAN.

This operation deletes the ENode state from the switch database and from the FIP snooping firewall filters, which causes the ENode to lose its connection to the FCoE forwarder (FCF) and to log in to the FCF again.

The command syntax in a Junos Fusion environment includes the `satellite` keyword and clears FIP snooping Enode information on satellite device FCoE VLANs, which have FCoE and FIP functions and status information distributed between the aggregation devices and satellite devices. The command validates that a specified VLAN is a satellite FCoE VLAN, and displays an error message if the satellite syntax is not used for a satellite FCoE VLAN, or if the satellite syntax is used with a VLAN that is not a satellite FCoE VLAN.

Options

enode-mac	MAC address of the ENode.
vlan *vlan-name*	(Optional) Name of the VLAN.

Required Privilege Level

view

Sample Output

clear fip snooping enode enode-mac

```
user@switch> clear fip snooping enode 00:10:94:00:00:02
```

Release Information

Command introduced in Junos OS Release 10.4.

Command introduced with satellite keyword in Junos OS Release 17.4R1 for Junos Fusion Data Center.

RELATED DOCUMENTATION

show fip snooping enode | 20015

clear fip snooping statistics

IN THIS SECTION

- Syntax | 3607
- Syntax (Junos Fusion) | 3607
- Description | 3607
- Required Privilege Level | 3607

Sample Output | 3607
Release Information | 3608

Syntax

```
clear fip snooping statistics
<vlan vlan-name>
```

Syntax (Junos Fusion)

```
clear fip snooping satellite statistics
<vlan vlan-name>
```

Description

Clear FIP snooping statistics globally or on a specified VLAN.

The command syntax in a Junos Fusion environment includes the `satellite` keyword and clears FIP snooping information for satellite device FCoE VLANs, which have FCoE and FIP functions and status information distributed between the aggregation devices and satellite devices. The command validates that a specified VLAN is a satellite FCoE VLAN, and displays an error message if the satellite syntax is not used for a satellite FCoE VLAN, or if the satellite syntax is used with a VLAN that is not a satellite FCoE VLAN.

Required Privilege Level

view

Sample Output

clear fip snooping statistics

```
user@switch> clear fip snooping statistics
```

Release Information

Command introduced in Junos OS Release 10.4.

Command introduced with satellite keyword in Junos OS Release 17.4R1 for Junos Fusion Data Center.

RELATED DOCUMENTATION

show fip snooping statistics | 20029

clear fip snooping vlan

IN THIS SECTION

- Syntax | 3608
- Syntax (Junos Fusion) | 3608
- Description | 3609
- Options | 3609
- Required Privilege Level | 3609
- Sample Output | 3609
- Release Information | 3609

Syntax

```
clear fip snooping vlan vlan-name
```

Syntax (Junos Fusion)

```
clear fip snooping satellite vlan vlan-name
```

Description

Clear FIP snooping information for the specified FCoE VLAN.

This operation deletes all ENode and FCF information for the specified VLAN from the switch database and causes the ENodes to lose their connections to the FCFs. After clearing a VLAN, the switch relearns all of the FCFs and ENodes on the VLAN, and the ENodes must log in to the FCF again.

The command syntax in a Junos Fusion environment includes the `satellite` keyword to clear FIP snooping information for satellite device FCoE VLANs, which have FCoE and FIP functions distributed between the aggregation devices and satellite devices. The command validates that a specified VLAN is a satellite FCoE VLAN, and displays an error message if the satellite syntax is not used for a satellite FCoE VLAN, or if the satellite syntax is used with a VLAN that is not a satellite FCoE VLAN.

Options

`vlan-name`	Name of the VLAN.

Required Privilege Level

view

Sample Output

clear fip snooping vlan vlan-name

```
user@switch> clear fip snooping vlan fcoevlan1
```

Release Information

Command introduced in Junos OS Release 10.4.

Command introduced with satellite keyword in Junos OS Release 17.4R1 for Junos Fusion Data Center.

RELATED DOCUMENTATION

show fip snooping vlan | 20034

clear fip vlan-discovery statistics

IN THIS SECTION

- Syntax | 3610
- Description | 3610
- Required Privilege Level | 3610
- Sample Output | 3610
- Release Information | 3610

Syntax

```
clear fip vlan-discovery statistics
```

Description

Clear FIP VLAN discovery statistics.

Required Privilege Level

view

Sample Output

clear fip vlan-discovery statistics

```
user@switch> clear fip vlan-discovery statistics
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

show fip vlan-discovery | 20040

clear firewall (QFX)

IN THIS SECTION

- Syntax | 3611
- Description | 3611
- Options | 3612
- Required Privilege Level | 3612
- Output Fields | 3612
- Sample Output | 3612
- Release Information | 3613

Syntax

```
clear firewall (all | counter counter-name | filter filter-name)
```

Description

Clear statistics provided by firewall filters.

> **NOTE**: The clear firewall command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for GRES.

When you clear the counters of a filter, this not only impacts the counters shown by the CLI, but also the ones tracked by SNMP 2.

If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the prefix-action action on matched packets, wait at least 5 seconds before you enter the show firewall prefix-action-stats command. A 5-second pause between issuing the clear firewall and show firewall prefix-action-stats commands avoids a possible timeout of the show firewall prefix-action-stats command.

Options

all	Clear the packet and byte counts for all firewall filter counters and clear the packet counts for all policer counters.
counter *counter-name*	Clear the packet and byte counts for the specified firewall filter counter.
filter *filter-name*	Clear the packet and byte counts for the specified firewall filter.
policer counter (all | counter-id *counter-index***)**	Clear all policer counters using the `policer counter all` command, or clear a specific policer counter using the `policer counter counter-id` *`counter-index`* command. The value of *counter-index* can be 0, 1, or 2.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear firewall all

```
user@switch> clear firewall all
```

clear firewall counter

```
user@switch> clear firewall counter port-filter-counter
```

clear firewall filter

```
user@switch> clear firewall filter ingress-port-filter
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Verifying That Two-Color Policers Are Operational

Overview of Firewall Filters (QFX Series)

Overview of Policers

show firewall (View) | **20073**

clear firewall

IN THIS SECTION

- Syntax | **3613**
- Syntax (EX Series Switches) | **3614**
- Syntax (QFX Series) | **3614**
- Description | **3614**
- Options | **3614**
- Required Privilege Level | **3615**
- Sample Output | **3615**
- Release Information | **3616**

Syntax

```
clear firewall (all | counter counter-name | filter filter-name | log (all | logical-system-
name ) | logical-system logical-system-name)
```

Syntax (EX Series Switches)

```
clear firewall (all | counter counter-name | filter filter-name | log (all | logical-system-name) | policer counter (all | counter-id counter-index))
```

Syntax (QFX Series)

```
clear firewall (all | counter counter-name | filter filter-name
```

Description

Clear statistics about configured firewall filters.

When you clear the counters of a filter, this impacts not only the counters shown by the CLI, but also the ones tracked by SNMP2.

Subscriber management uses firewall filters to capture and report the volume-based service accounting counters that are used for subscriber billing. The `clear firewall` command also clears the service accounting counters that are reported to the RADIUS accounting server. For this reason, you must be cautious in specifying which firewall statistics you want to clear.

> **NOTE**: The `clear firewall` command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover (GRES).

If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the `prefix-action` action on matched packets, wait at least 5 seconds before you enter the `show firewall prefix-action-stats` command. A 5-second pause between issuing the `clear firewall` and `show firewall prefix-action-stats` commands avoids a possible timeout of the `show firewall prefix-action-stats` command.

Options

all	Clear the packet and byte counts for all filters. On EX Series switches, this option also clears the packet counts for all policer counters.
counter *counter-name*	Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.
filter *filter-name*	Clear the packet and byte counts for the specified firewall filter.

log (all | *logical-system-name*)	Clear log entries for IPv4 firewall filters that have `then log` as an action. Use `log all` to clear all log entries or `log logical-system-name` to clear log entries for the specified logical system.
logical-system *logical-system-name*	Clear the packet and byte counts for the specified logical system.
policer counter (all | counter-id *counter-index*)	(EX8200 switches only) Clear all policer counters using the `policer counter all` command, or clear a specific policer counter using the `policer counter counter-id counter-index` command. The value of *counter-index* can be 0, 1, or 2.

Required Privilege Level

clear

Sample Output

clear firewall all

```
user@host> clear firewall all
```

clear firewall (counter counter-name)

```
user@host> clear firewall counter port-filter-counter
```

clear firewall (filter filter-name)

```
user@host> clear firewall filter ingress-port-filter
```

clear firewall (policer counter all) (EX8200 Switch)

```
user@switch> clear firewall policer counter all
```

clear firewall (policer counter counter-id counter-index) (EX8200 Switch)

```
user@switch> clear firewall policer counter counter-id 0
```

Release Information

Command introduced before Junos OS Release 7.4.

`logical-system` option introduced in Junos OS Release 9.3.

Command introduced in Junos OS Release 11.1.

`log` option introduced before Junos OS Release 11.4.

RELATED DOCUMENTATION

- show firewall | 20048
- *Verifying That Two-Color Policers Are Operational*
- *Overview of Firewall Filters (QFX Series)*
- *Overview of Policers*
- show firewall (View) | 20073

clear forwarding-options port-mirroring instance family

IN THIS SECTION

- Syntax | 3617
- Description | 3617
- Options | 3617
- Required Privilege Level | 3617
- Output Fields | 3617
- Sample Output | 3617
- Release Information | 3617

Syntax

```
clear forwarding-options port-mirroring instance instance-name family family-name
```

Description

Clear the self-mirroring statistics and delete the associated PCAP files.

Options

family *family-name* (Optional) Protocol family to which this command is applied.

instance *instance-name* (Optional) Port-mirroring instance to which this command is applied.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show forwarding-options port-mirroring self-mirror statistics` command to verify the operation.

Sample Output

clear forwarding-options port-mirroring instance I1 family inet

```
user@host> clear forwarding-options port-mirroring instance I1 family inet
```

Release Information

Command introduced in Junos OS Evolved Release 24.2R1.

clear gtp tunnels

IN THIS SECTION

- Syntax | 3618
- Description | 3618
- Options | 3618
- Required Privilege Level | 3618
- Release Information | 3618

Syntax

```
clear security gtp tunnel <all | identifier>
```

Description

Use this command to clear all or specified GTP tunnels on the device.

Options

- all—Clears all existing tunnels.
- *identifier*—Clears a single tunnel by entering the tunnel ID. To view current tunnel IDs, type `show security gtp tunnels`.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 10.0.

Support for GTPv2 added in Junos OS Release 11.4.

RELATED DOCUMENTATION

show gtp tunnels | 20122

clear igmp membership

IN THIS SECTION

- Syntax | 3619
- Syntax (EX Series Switch and the QFX Series) | 3619
- Description | 3620
- Options | 3620
- Required Privilege Level | 3620
- Output Fields | 3620
- Sample Output | 3620
- Release Information | 3623

Syntax

```
clear igmp membership
<all>
<group address-range>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
clear igmp membership
<group address-range>
<interface interface-name>
```

Description

Clear Internet Group Management Protocol (IGMP) group members.

Options

all	Clear IGMP members for groups and interfaces in the master instance.
group *address-range*	(Optional) Clear all IGMP members that are in a particular address range. An example of a range is `233.252/16`. If you omit the destination prefix length, the default is `/32`.
interface *interface-name*	(Optional) Clear all IGMP group members on an interface.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

See "show igmp group" on page 20135 for an explanation of output fields.

Sample Output

clear igmp membership all

The following sample output displays IGMP group information before and after the `clear igmp membership` command is entered:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         198.51.100.253   203.0.113.1          186
so-0/0/0         198.51.100.254   203.0.113.1          186
so-0/0/0         198.51.100.255   203.0.113.1          187
so-0/0/0         198.51.100.240   203.0.113.1          188
local            198.51.100.6     (null)                0
local            198.51.100.5     (null)                0
```

```
local            198.51.100.25    (null)                0
local            198.51.100.22    (null)                0
local            198.51.100.2     (null)                0
local            198.51.100.13    (null)                0


user@host> clear igmp membership all
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0


user@host> show igmp group
Interface        Group           Last Reported   Timeout
local            198.51.100.6            (null)                0
local            198.51.100.5            (null)                0
local            198.51.100.254          (null)                0
local            198.51.100.255          (null)                0
local            198.51.100.2            (null)                0
local            198.51.100.13           (null)                0
```

clear igmp membership interface

The following sample output displays IGMP group information before and after the `clear igmp membership interface` command is issued:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         198.51.100.253   203.0.113.1          210
so-0/0/0         198.51.100.200   203.0.113.1          210
so-0/0/0         198.51.100.255   203.0.113.1          215
so-0/0/0         198.51.100.254   203.0.113.1          216
local            198.51.100.6     (null)                0
local            198.51.100.5     (null)                0
local            198.51.100.254   (null)                0
local            198.51.100.255   (null)                0
local            198.51.100.2     (null)                0
local            198.51.100.13    (null)                0


user@host> clear igmp membership interface so-0/0/0
```

```
Clearing Group Membership Info for so-0/0/0



user@host> show igmp group
Interface        Group           Last Reported   Timeout
local            198.51.100.6       (null)               0
local            198.51.100.5       (null)               0
local            198.51.100.254     (null)               0
local            198.51.100.255     (null)               0
local            198.51.100.2       (null)               0
local            198.51.100.13      (null)               0
```

clear igmp membership group

The following sample output displays IGMP group information before and after the `clear igmp membership group` command is entered:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         198.51.100.253   203.0.113.1          210
so-0/0/0         198.51.100.25    203.0.113.1          210
so-0/0/0         198.51.100.255   203.0.113.1          215
so-0/0/0         198.51.100.254   203.0.113.1          216
local            198.51.100.6     (null)                0
local            198.51.100.5     (null)                0
local            198.51.100.254   (null)                0
local            198.51.100.25    (null)                0
local            198.51.100.2     (null)                0
local            198.51.100.13    (null)                0


user@host> clear igmp membership group 233.252/16
Clearing Group Membership Range 198.51.100.0/16 on so-0/0/0
Clearing Group Membership Range 198.51.100.0/16 on so-1/0/0
Clearing Group Membership Range 198.51.100.0/16 on so-2/0/0


user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         198.51.100.255   203.0.113.1          231
so-0/0/0         198.51.100.254   203.0.113.1          233
```

```
so-0/0/0          198.51.100.253   203.0.113.1          236
local             198.51.100.6     (null)                0
local             198.51.100.5     (null)                0
local             198.51.100.254   (null)                0
local             198.51.100.255   (null)                0
local             198.51.100.2     (null)                0
local             198.51.100.13    (null)                0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show igmp group | 20135

show igmp interface | 20142

clear igmp snooping membership

IN THIS SECTION

- Syntax | 3623
- Description | 3624
- Options | 3624
- Required Privilege Level | 3624
- Output Fields | 3625
- Sample Output | 3625
- Release Information | 3625

Syntax

```
clear igmp snooping membership
<vlan                 vlan-name>
```

```
<group | source address>
<instance instance-name>
<interface interface-name>
<learning-domain learning-domain-name>
<logical-system logical-system-name>
<vlan-id vlan-identifier>
```

Description

Clear IGMP snooping dynamic membership information from the multicast forwarding table.

Options

none	Clear IGMP snooping membership for all supported address families on all interfaces.
vlan *vlan-name*	(Optional) Clear dynamic membership information for the specified VLAN.
group | source *address*	(Optional) Clear IGMP snooping membership for the specified multicast group or source address.
instance *instance-name*	(Optional) Clear IGMP snooping membership for the specified instance.
interface *interface-name*	(Optional) Clear IGMP snooping membership on a specific interface.
learning-domain *learning-domain-name*	(Optional) Perform this operation on all learning domains or on a particular learning domain.
logical-system *logical-system-name*	(Optional) Display information about a particular logical system, or for all logical systems.
vlan-id *vlan-identifier*	(Optional) Perform this operation on a particular VLAN.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear igmp snooping membership

```
user@host> clear igmp snooping membership
```

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

show igmp snooping membership | 20165

clear igmp snooping statistics

IN THIS SECTION

- Syntax | 3626
- Description | 3626
- Options | 3626
- Required Privilege Level | 3626
- Output Fields | 3626
- Sample Output | 3627
- Release Information | 3627

Syntax

```
clear igmp snooping statistics
<instance instance-name>
<interface interface-name>
<learning-domain (all | learning-domain-name)>
<logical-system logical-system-name>
```

Description

Clear IP IGMP snooping statistics.

Options

none	Clear IGMP snooping statistics for all supported address families on all interfaces.
instance *instance-name*	(Optional) Clear IGMP snooping statistics for the specified instance.
interface *interface-name*	(Optional) Clear IGMP snooping statistics on a specific interface.
learning-domain (all | *learning-domain-name*)	(Optional) Perform this operation on all learning domains or on a particular learning domain.
logical-system *logical-system-name*	(Optional) Delete the IGMP snooping statistics for a given logical system or for all logical systems.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear igmp snooping statistics

```
user@host> clear igmp snooping statistics
```

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

show igmp snooping statistics | 20187

clear igmp statistics

IN THIS SECTION

- Syntax | 3628
- Syntax (EX Series) | 3628
- Syntax (MX Series) | 3628
- Description | 3628
- Options | 3629
- Required Privilege Level | 3629
- Output Fields | 3629
- Sample Output | 3629
- Release Information | 3631

Syntax

```
clear igmp statistics
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series)

```
clear igmp statistics
<interface interface-name>
```

Syntax (MX Series)

```
clear igmp statistics
(<continuous> | <interface interface-name>)
<logical-system (all | logical-system-name)>
```

Description

Clear Internet Group Management Protocol (IGMP) statistics. Clearing IGMP statistics zeros the statistics counters as if you rebooted the device.

By default, Junos OS multicast devices collect statistics of received and transmitted IGMP control messages that reflect currently active multicast group subscribers. Some devices also automatically maintain *continuous* IGMP statistics globally on the device in addition to the default active subscriber statistics—these are persistent, continuous statistics of received and transmitted IGMP control packets that account for both past and current multicast group subscriptions processed on the device. The device maintains continuous statistics across events or operations such as routing daemon restarts, graceful Routing Engine switchovers (GRES), in-service software upgrades (ISSU), or line card reboots. The default active subscriber-only statistics are not preserved in these cases.

Run this command to clear the currently active subscriber statistics. On devices that support continuous statistics, run this command with the `continuous` option to clear the continuous statistics. You must run these commands separately to clear both types of statistics because the device maintains and clears the two types of statistics separately.

Options

none	Clear IGMP statistics on all interfaces. This form of the command clears statistics for currently active subscribers only.
continuous	Clear only the continuous IGMP statistics that account for both past and current multicast group subscribers instead of the default statistics that only reflect currently active subscribers. This option is not available with the `interface` option for interface-specific statistics.
interface ***interface-name***	(Optional) Clear IGMP statistics for the specified interface only. This option is not available with the `continuous` option.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

See "show igmp statistics" on page 20203 for an explanation of output fields.

Sample Output

clear igmp statistics

The following sample output displays IGMP statistics information before and after the `clear igmp statistics` command is entered:

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query           8883        459          0
V1 Membership Report          0          0          0
DVMRP                     19784      35476          0
PIM V1                    18310          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
```

```
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                  0
IGMP v3 unsupported type                             0
IGMP v3 source required for SSM                      0
IGMP v3 mode not applicable for SSM                  0

IGMP Global Statistics
Bad Length                    0
Bad Checksum                  0
Bad Receive If                0
Rx non-local               1227


user@host> clear igmp statistics
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          0          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                  0
IGMP v3 unsupported type                             0
IGMP v3 source required for SSM                      0
IGMP v3 mode not applicable for SSM                  0
IGMP Global Statistics
Bad Length                    0
Bad Checksum                  0
Bad Receive If                0
Rx non-local                  0
```

Release Information

Command introduced before Junos OS Release 7.4.

`continuous` option added in Junos OS Release 19.4R1 for MX Series routers.

RELATED DOCUMENTATION

show igmp statistics | 20203

clear-ike-sas-on-pic-restart

IN THIS SECTION

- Syntax | 3631
- Hierarchy Level | 3631
- Description | 3631
- Required Privilege Level | 3632
- Release Information | 3632

Syntax

```
clear-ike-sas-on-pic-restart;
```

Hierarchy Level

```
[edit services ipsec-vpn]
```

Description

Clear IKE security associations (SAs) when the corresponding PIC restarts or is taken offline.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Security Associations

clear ike security-associations

IN THIS SECTION

- Syntax | 3632
- Description | 3633
- Options | 3633
- Required Privilege Level | 3633
- Output Fields | 3633
- Sample Output | 3633
- Release Information | 3633

Syntax

```
clear ike security-associations
<destination-ip-address>
```

Description

(Encryption interface on M Series and T Series routers only) Clear information about the current Internet Key Exchange (IKE) security association. This command is valid for dynamic security associations only. For IKEv2, this command creates new security associations for IKE SA and IPSEC SAs.

Options

none	Clear all IKE security associations.
destination-ip-address	(Optional) Clear the IKE security association at the specified destination address.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ike security-associations

```
user@host> clear ike security-associations
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ike security-associations | 20210

clear ilmi statistics

IN THIS SECTION

- Syntax | 3634
- Description | 3634
- Options | 3634
- Required Privilege Level | 3634
- Output Fields | 3634
- Sample Output | 3635
- Release Information | 3635

Syntax

```
clear ilmi statistics
```

Description

Set Integrated Local Management Interface (ILMI) statistics to zero.

Options

This command has no options.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ilmi statistics

```
user@host> clear ilmi statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ilmi statistics | 20217

clear inband-flow-telemetry stats

IN THIS SECTION

- Syntax | 3635
- Description | 3636
- Options | 3636
- Required Privilege Level | 3636
- Output Fields | 3636
- Sample Output | 3636
- Release Information | 3636

Syntax

```
clear inband-flow-telemetry stats
```

Description

Clear Inband Flow Analyzer 2.0 (IFA 2.0) statistics.

You retrieve IFA statistics directly from the Packet Forwarding Engine. The Routing Engine does not maintain these statistics. Therefore, restarting the Packet Forwarding Engine process clears the statistics, whereas a Routing Engine process restart does not have any impact on the statistics.

Options

None

Required Privilege Level

network

Output Fields

This command produces no output.

Sample Output

clear inband-flow-telemetry stats (QFX5120-48Y and QFX5120-32C)

```
user@host> clear inband-flow-telemetry stats
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

show services inband-flow-telemetry | 25574

inband-flow-telemetry | 8512

Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Performance Monitoring

clear interfaces aeX forwarding-options load-balance state

IN THIS SECTION

- Syntax | 3637
- Description | 3637
- Options | 3637
- Required Privilege Level | 3638
- Sample Output | 3638
- Release Information | 3638

Syntax

```
clear interfaces aeX unit logical-unit-number aggregate forwarding-options load-balance state
```

Description

Clear the specified aggregate Ethernet interface load balancing state and re-create it newly. If the traffic flows become aged frequently, then the device needs to remove or refresh the load balancing states. As a result, you must configure rebalancing or run the clear command at periodic intervals for proper load-balancing. Otherwise, traffic skewing can occur.

If you observe load distribution to be not very effective, you can clear the load-balancing states or use rebalancing functionality to cause an automatic clearance of the hardware states. When you configure the rebalancing facility, traffic flows can get redirected to different links, which can cause packet reordering.

Options

ae*X*	Name of a particular aggregated Ethernet interface.
logical-unit-number	Number of the logical unit of the interface.
forwarding-options load-balance state	Cause the load-balancing state to be cleared for the specific interface.

Required Privilege Level

clear

Sample Output

clear interfaces aeX aggregate forwarding-options

```
user@host> clear interfaces  ae1 aggregate forwarding-options load-balance state
```

Release Information

Command introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

show interfaces interval | 20934

clear interfaces aggregate forwarding-options load-balance state

IN THIS SECTION

- Syntax | 3639
- Description | 3639
- Options | 3639
- Required Privilege Level | 3639
- Sample Output | 3639
- Release Information | 3639

Syntax

```
clear interfaces aggregate forwarding-options load-balance state
```

Description

Clear all the aggregate Ethernet interface load balancing states and re-create them newly. If the traffic flows become aged frequently, then the device needs to remove or refresh the load balancing states. As a result, you must configure rebalancing or run the clear command at periodic intervals for proper load-balancing. Otherwise, traffic skewing can occur.

Options

interface-name	Name of a particular interface.

Required Privilege Level

clear

Sample Output

clear interfaces aggregate forwarding-options

```
user@host> clear interfaces  aggregate forwarding-options load-balance state
```

Release Information

Command introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

show interfaces interval | 20934

clear interfaces interface-set statistics

IN THIS SECTION

- Syntax | 3640
- Description | 3640
- Options | 3640
- Required Privilege Level | 3640
- Output Fields | 3640
- Sample Output | 3641
- Release Information | 3641

Syntax

```
clear interfaces interface-set statistics interface-set-name
```

Description

Set interface set statistics to zero.

Options

interface-set-name	Set statistics on a specified interface set to zero. Wildcard values can be used in the interface set name. This command will not clear the statistics of the member logical interfaces.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear interfaces interface-set statistics

```
user@host> clear interfaces interface-set statistics
```

Release Information

Command introduced in Junos OS Release 8.5.

clear interfaces interval

IN THIS SECTION

- Syntax | 3641
- Description | 3641
- Options | 3642
- Required Privilege Level | 3642
- Output Fields | 3642
- Sample Output | 3642
- Release Information | 3643

Syntax

```
clear interfaces interval interface-name
```

Description

Clear the channel service unit (CSU) alarm and defect counters so that only the current time interval is displayed. This operation affects the `show interface interval` command, but not an SNMP query.

Options

interface-name	Name of a particular interface.

Required Privilege Level

clear

Output Fields

See "show interfaces interval" on page 20934 for an explanation of output fields.

Sample Output

clear interfaces interval

The following example displays the output for a T3 interface before and after the `clear interfaces` command is entered:

```
user@host> show interfaces interval t3-0/3/0:4
Physical interface: t3-0/3/0:4, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  16:28-16:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
CES: 195, CSES: 195, SEFS: 195, UAS: 206
  14:58-15:13:
```

```
    LCV: 35, PCV: 163394, CCV: 54485, LES: 0, PES: 35, PSES: 35, CES:
35, CSES: 35, SEFS: 35, UAS: 32
  Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

```
user@host> clear interfaces interval t3-0/3/0:4
```

```
user@host> show interfaces interval t3-0/3/0:4
Physical interface: t3-0/3/0:4, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  Interval Total:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0, SEFS: 0, UAS: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show interfaces interval | 20934

clear interfaces lb-stats af

IN THIS SECTION

- Syntax | 3644
- Description | 3644
- Required Privilege Level | 3644
- Sample Output | 3644

Release Information | 3647

Syntax

```
clear interfaces lb-stats af-interface-name
```

Description

Clears the information about the load balancing of transmit traffic on each peer packet forwarding engine of the abstracted fabric (af) interface. This command restarts the load balancing transmit values from zero.

Required Privilege Level

View

Sample Output

show interfaces lb-stats af2

```
user@host> show interfaces lb-stats af2
```

```
Physical interface: af2, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 501
  Type: Ethernet, Link-level type: Ethernet, MTU: 1518, Speed: 3000000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Current address: 00:90:69:15:20:47, Hardware address: 00:90:69:15:20:47
  Last flapped   : 2021-09-01 00:40:40 PDT (2w6d 01:56 ago)

  Peer GNF id    : 2
  Collapsed Forwarding : Disabled
  Peer GNF Forwarding element(FE) view :
  FPC slot:FE num  FE Bandwidth(Gbps) Status      Transmit Packets
```

```
Trans                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                mit
Bytes
       2:2                   500         Up
0
             0
       2:3                   500         Up
0
             0
       2:4                   500         Up
0
             0
       2:5                   500         Up
0
             0
       2:6                   500         Up
203674
           19873016
       2:7                   500         Up
834
           38364

  Residual Transmit Statistics :
  Packets :                    0 Bytes :                    0

  Fabric Queue Statistics :
  FPC slot:FE num    High priority(pkts)        Low priority(pkts)
       2:2                            0                          0
       2:3                            0                          0
       2:4                            0                          0
       2:5                            0                          0
       2:6                            0                     203674
       2:7                            0                        834
  FPC slot:FE num    High priority(bytes)      Low priority(bytes)
       2:2                              0                        0
       2:3                              0                        0
       2:4                              0                        0
       2:5                              0                        0
       2:6                              0                 19873016
       2:7                              0                    38364
  Residual Queue Statistics :
      High priority(pkts)       Low priority(pkts)
                       0                         0
```

```
      High priority(bytes)      Low priority(bytes)
                         0                        0
```

clear interfaces lb-stats af2

```
user@host> clear interfaces lb-stats af2
```

show interfaces lb-stats af2

```
user@host> show interfaces lb-stats af2
```

```
Physical interface: af2, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 501
  Type: Ethernet, Link-level type: Ethernet, MTU: 1518, Speed: 3000000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Current address: 00:90:69:15:20:47, Hardware address: 00:90:69:15:20:47
  Last flapped   : 2021-09-01 00:40:40 PDT (2w6d 22:40 ago)

  Peer GNF id    : 2
  Collapsed Forwarding : Disabled
  Peer GNF Forwarding element(FE) view :
  FPC slot:FE num  FE Bandwidth(Gbps) Status      Transmit Packets
Trans                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                
mit
Bytes
       2:2                  500         Up
0
             0
       2:3                  500         Up
0
             0
       2:4                  500         Up
0
             0
       2:5                  500         Up
0
             0
```

```
        2:6                   500         Up
8
            784
        2:7                   500         Up
0
              0

  Residual Transmit Statistics :
  Packets :                    0 Bytes :                    0

  Fabric Queue Statistics :
  FPC slot:FE num    High priority(pkts)        Low priority(pkts)
       2:2                            0                          0
       2:3                            0                          0
       2:4                            0                          0
       2:5                            0                          0
       2:6                            0                          8
       2:7                            0                          0
  FPC slot:FE num    High priority(bytes)      Low priority(bytes)
       2:2                              0                        0
       2:3                              0                        0
       2:4                              0                        0
       2:5                              0                        0
       2:6                              0                      784
       2:7                              0                        0
  Residual Queue Statistics :
      High priority(pkts)       Low priority(pkts)
                       0                        0
      High priority(bytes)      Low priority(bytes)
                        0                        0
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

show interfaces lb-stats af | **20959**

clear interfaces mac-database

IN THIS SECTION

- Syntax | 3648
- Description | 3648
- Options | 3648
- Required Privilege Level | 3649
- Output Fields | 3649
- Sample Output | 3649
- Release Information | 3649

Syntax

```
clear interfaces mac-database (interface-name | aex) <mac-address mac-address>
```

Description

Clear learned media access control (MAC) addresses from the hardware and MAC database for Gigabit Ethernet IQ2 interfaces or aggregated Ethernet interfaces. Static MAC addresses configured by the operator are not cleared.

Options

interface-name	Name of a physical or logical interface. When you clear a physical interface, all learned MAC addresses on all the logical interfaces under the physical interface are cleared.
ae*x*	Name of aggregated Ethernet interface.
mac-address *mac-address*	(Optional) Clear only the specified MAC address.

Required Privilege Level

view

Output Fields

This command produces no output.

Sample Output

clear interfaces mac-database

```
user@host> clear interfaces mac-database ge-0/0/0.0
```

Release Information

Command introduced in Junos OS Release 8.3.

Support for statement with the `ae*x*` option introduced in Junos OS Release 15.1.

clear interfaces mac-database statistics

IN THIS SECTION

- Syntax | 3650
- Description | 3650
- Options | 3650
- Required Privilege Level | 3650
- Output Fields | 3650
- Sample Output | 3650
- Release Information | 3650

Syntax

```
clear interfaces mac-database statistics (interface-name /all)
```

Description

Clear statistics that are collected for every MAC address, including policer statistics, on a physical or logical interface or all interfaces.

Options

(*interface-name* | all) Clear MAC database statistics for the specified physical or logical gigabit or 10-Gigabit Ethernet interface. Specify **all** to clear the MAC database statistics for all interfaces.

Required Privilege Level

view

Output Fields

This command produces no output.

Sample Output

clear interfaces mac-database statistics (Gigabit Ethernet)

```
user@host> clear interfaces mac-database statistics ge-0/1/0
```

Release Information

Command introduced in Junos OS Release 8.3.

clear interfaces statistics

IN THIS SECTION

- Syntax | 3651
- Description | 3651
- Options | 3652
- Required Privilege Level | 3652
- Output Fields | 3652
- Sample Output | 3652
- Release Information | 3652

Syntax

```
clear interfaces statistics (all | interface-name)
```

Description

Set interface statistics to zero. If you issue the `clear interfaces statistics` *`interface-name`* command and then perform a graceful Routing Engine switchover, the interface statistics are not cleared on the new primary node. Reissue the command to clear the interface statistics again.

Starting in Junos OS Release 17.3R1, this command supports the clearing of Packet Forwarding Engine accounting statistics on logical interfaces configured with accounting options. On these interfaces, the current statistics values are stored as the new current baseline values and then the counters are reset to zero. If the `allow-clear` statement is included in the interface profile, then the cleared statistics values are reported to the accounting options flat file associated with the interface. Reporting is disabled by default; if `allow-clear` is not configured, then the CLI displays cleared statistics counters, but they are not reported to the flat file.

Starting in Junos OS Release 19.1R1, this command supports the clearing of unicast Reverse Path Forwarding (RPF) statistics.

Starting in Junos OS Release 21.3R1, on MX Series Routers, this command supports the clearing of Reverse Path Forwarding (RPF) statistics of the Ethernet interfaces.

Options

all	Set statistics on all interfaces to zero.
interface-name	Set statistics on a particular interface to zero.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear interfaces statistics

```
user@host> clear interfaces statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, this command supports the clearing of Packet Forwarding Engine accounting statistics on logical interfaces configured with accounting options.

clear interfaces statistics swfabx

IN THIS SECTION

- Syntax | 3653
- Description | 3653
- Required Privilege Level | 3653
- Output Fields | 3653
- Sample Output | 3653
- Release Information | 3654

Syntax

```
clear interfaces statistics <swfab0 | swfab1>
```

Description

Clear interface statistics for the specified swfab interface.

Required Privilege Level

clear

Output Fields

When you enter this command, interface statistics for swfab0 and swfab1 are cleared.

Sample Output

clear interfaces statistics <swfab0 | swfab1>

```
user@host> clear interfaces statistics <swfab0 | swfab1>
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show interfaces swfabx | 21492

clear interfaces transport pm

IN THIS SECTION

- Syntax | 3654
- Description | 3654
- Options | 3654
- Required Privilege Level | 3655
- Output Fields | 3655
- Sample Output | 3655
- Release Information | 3655

Syntax

```
clear interfaces transport pm (all | optics | otn) (all | current | current-day) (all |
interface-name)
```

Description

Clear optics and OTN information from the transport performance monitoring data.

Options

(all | optics | otn)—Clear both optics and OTN information or either only optics or only OTN information.

(all | current | current-day)—Clear information for the current 15-minute interval, the ninety-six 15-minute intervals, the current day, and the previous day; information only for the current 15-minute interval; or information only for the current 24 hours.

(all | *interface-name*)—Clear information for all interfaces or only for the specified interface (for example, *et-fpc/pic/port*).

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear interfaces transport pm

```
user@host> clear interfaces transport pm transport otn current all
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

show interfaces transport pm | **21509**

100-Gigabit Ethernet OTN Options Configuration Overview

tca | **28973**

clear-ipsec-sas-on-pic-restart

IN THIS SECTION

- Syntax | 3656
- Hierarchy Level | 3656
- Description | 3656
- Required Privilege Level | 3656
- Release Information | 3656

Syntax

```
clear-ipsec-sas-on-pic-restart;
```

Hierarchy Level

```
[edit services ipsec-vpn]
```

Description

Clear IPsec security associations (SAs) when the corresponding PIC restarts or is taken offline.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Security Associations

clear ipsec security-associations

IN THIS SECTION

- Syntax | 3657
- Description | 3657
- Options | 3657
- Required Privilege Level | 3658
- Output Fields | 3658
- Sample Output | 3658
- Release Information | 3659

Syntax

```
clear ipsec security-associations
<sa-name>
```

Description

(Encryption interface on M Series and T Series routers only) Clear information about the current IP Security (IPsec) security association. This command is valid for dynamic security associations only. For IKEv1, this command creates new security associations for IKE SA and IPSEC SAs.

Options

none	Clear all IPsec security associations.
sa-name	(Optional) Clear the specified security association.

Required Privilege Level

view

Output Fields

See the "show ipsec security-associations" on page 21675 for an explanation of output fields.

Sample Output

clear ipsec security-associations

The following output from the `show ipsec security-associations detail` command is displayed before and after the `clear ipsec security-associations` command is issued:

```
user@host> show ipsec security-associations detail
Security association: sa-dynamic, Interface family: Up

  Direction: inbound, SPI: 242379418, State: Installed
  Mode: tunnel, Type: dynamic
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
  Soft lifetime: Expires in 22979 seconds
  Hard lifetime: Expires in 28739 seconds

  Direction: outbound, SPI: 368592771, State: Installed
  Mode: tunnel, Type: dynamic
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
  Soft lifetime: Expires in 22979 seconds
  Hard lifetime: Expires in 28739 seconds


user@host> clear ipsec security-associations


user@host> show ipsec security-associations detail
Security association: sa-dynamic, Interface family: Up

  Direction: inbound, SPI: 1031597683, State: Installed
  Mode: tunnel, Type: dynamic
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
  Soft lifetime: Expires in 23037 seconds
```

```
  Hard lifetime: Expires in 28797 seconds

  Direction: outbound, SPI: 1618419878, State: Installed
  Mode: tunnel, Type: dynamic
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
  Soft lifetime: Expires in 23037 seconds
  Hard lifetime: Expires in 28797 seconds
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ipsec security-associations | 21675

clear ipv6 neighbors

IN THIS SECTION

- Syntax | 3659
- Description | 3660
- Options | 3660
- Required Privilege Level | 3660
- Output Fields | 3661
- Sample Output | 3661
- Release Information | 3661

Syntax

```
clear ipv6 neighbors
<all>
< host hostname>
```

```
< interface interface-name>
< logical-system logical-system-name>
< tenant name>
< vpn vpn-name>
```

Description

Clear IPv6 neighbor cache information.

On Junos OS, issuing the `clear ipv6 neighbors` command does not clear the cache for IPv6 neighbors in a reachable state.

On Junos OS Evolved, issuing the `clear ipv6 neighbors` command clears the cache for IPv6 neighbors in a reachable state.

Options

none	Clear all IPv6 neighbor cache information.
all	(Optional) Clear all IPv6 neighbor cache information.
host ***hostname***	(Optional) Clear the information for the specified IPv6 neighbors.
interface ***interface-name***	(Optional) Clear the information about IPv6 for the specified logical interface.
logical-system ***logical-system-name***	(Optional) Clear the IPv6 entries for the specified logical system; only available in the main router context.
tenant ***name***	(Optional) Clear the IPv6 entries for the specified tenant system; only available in the main router context.
vpn ***vpn-name***	(Optional) Clear entries in the IPv6 table for the specified virtual private network's (VPN) routing table.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ipv6 neighbors | 21680

clear ipv6 neighbors (SRX Series)

IN THIS SECTION

- Syntax | 3662
- Description | 3662
- Options | 3662
- Required Privilege Level | 3662
- Output Fields | 3662
- Sample Output | 3662
- Release Information | 3662

Syntax

```
clear ipv6 neighbors
<all | host hostname>
```

Description

Clear IPv6 neighbor cache information.

Options

none	Clear all IPv6 neighbor cache information.
all	(Optional) Clear all IPv6 neighbor cache information.
host *hostname*	(Optional) Clear the information for the specified IPv6 neighbors.

Required Privilege Level

clear

Output Fields

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
11:11::2            00:19:e2:4b:61:83  deleted
            12:12::2            00:19:e2:4b:61:83  deleted
            10:1::2             00:00:0a:00:00:00  deleted
```

Release Information

Command introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

show ipv6 neighbors (SRX Series) | 21683

clear ipv6 router-advertisement

IN THIS SECTION

- Syntax | 3663
- Description | 3663
- Options | 3663
- Required Privilege Level | 3664
- Output Fields | 3664
- Sample Output | 3664
- Release Information | 3664

Syntax

```
clear ipv6 router-advertisement
<interface interface>
<logical-system (all | logical-system-name)>
```

Description

Clear IPv6 router advertisement counters.

Options

none	Clear IPv6 router advertisement counters for all interfaces.
interface *interface*	(Optional) Clear IPv6 router advertisement counters for the specified interface.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 router-advertisement

```
user@host> clear ipv6 router-advertisement
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ipv6 router-advertisement | 21685

clear isis adjacency

IN THIS SECTION

- Syntax | 3665
- Syntax (EX Series Switches and QFX Series) | 3665
- Description | 3665
- Options | 3665

- Required Privilege Level | 3666
- Output Fields | 3666
- Sample Output | 3666
- Release Information | 3667

Syntax

```
clear isis adjacency
<all>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
<neighbor>
```

Syntax (EX Series Switches and QFX Series)

```
clear isis adjacency
<all>
<instance instance-name>
<interface interface-name>
<neighbor>
```

Description

Remove entries from the IS-IS adjacency database.

Options

all	Remove all entries from the adjacency database.
instance *instance-name*	(Optional) Clear all adjacencies for the specified routing instance only.
interface *interface-name*	(Optional) Clear all adjacencies for the specified interface only.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
neighbor	(Optional) Clear adjacencies for the specified neighbor only.

Required Privilege Level

clear

Output Fields

See "show isis adjacency" on page 21694 for an explanation of output fields.

Sample Output

clear isis adjacency

The following sample output displays IS-IS adjacency database information before and after the `clear isis adjacency` command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface     System         L State        Hold (secs) SNPA
so-1/0/0.0    karakul        3 Up                    26
so-1/1/3.0    1921.6800.5080 3 Up                    23
so-5/0/0.0    1921.6800.5080 3 Up                    19

user@host> clear isis adjacency karakul



user@host> show isis adjacency
IS-IS adjacency database:
Interface     System         L State        Hold (secs) SNPA
so-1/0/0.0    karakul        3 Initializing          26
so-1/1/3.0    1921.6800.5080 3 Up                    24
so-5/0/0.0    1921.6800.5080 3 Up                    21
```

clear isis adjacency all

```
user@host> clear isis adjacency all
IS-IS adjacency database:
Interface     System         L State        Hold (secs) SNPA
so-1/0/0.0    karakul        3 Initializing          26
so-1/1/3.0    1921.6800.5080 3 Initializing          24
so-5/0/0.0    1921.6800.5080 3 Initializing          21
```

Release Information

Command introduced before Junos OS Release 7.4.

`all` option introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

show isis adjacency | 21694

clear isis database

IN THIS SECTION

- Syntax | 3668
- Syntax (EX Series Switches and QFX Series) | 3668
- Description | 3668
- Options | 3668
- Required Privilege Level | 3668
- Output Fields | 3669
- Sample Output | 3669
- Release Information | 3669

Syntax

```
clear isis database
<all>
<entries>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
clear isis database
<all>
<entries>
<instance instance-name>
```

Description

Remove the entries from the IS-IS link-state database, which contains prefixes and topology information.

Options

all	Remove all entries from the IS-IS link-state database for all routing instances.
entries	(Optional) Name of the database entry.
instance *instance-name*	(Optional) Clear all entries for the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

See "show isis database" on page 21722 for an explanation of output fields.

Sample Output

clear isis database

The following sample output displays IS-IS link-state database information before and after the `clear isis database all` command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)
crater.00-00                0x12   0x84dd            1139
  1 LSPs
IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)
crater.00-00                0x19   0xe92c            1134
badlands.00-00              0x16   0x1454             985
carlsbad.00-00              0x33   0x220b            1015
ranier.00-00                0x2e   0xfc31            1007
1921.6800.5066.00-00        0x11   0x7313             566
1921.6800.5067.00-00        0x14   0xd9d4             939
  6 LSPs


user@host> clear isis database all


user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)

IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show isis database | 21722

clear isis overload

IN THIS SECTION

- Syntax | 3670
- Syntax (EX Series Switches and QFX Series) | 3670
- Description | 3671
- Options | 3671
- Required Privilege Level | 3671
- Output Fields | 3671
- Sample Output | 3671
- Release Information | 3672

Syntax

```
clear isis overload
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
clear isis overload
<instance instance-name>
```

Description

Reset the IS-IS dynamic overload bit. This command can appear to not work, continuing to display `overload` after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.

When other routers detect that the overload bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes.

Options

none	Reset the IS-IS dynamic overload bit.
instance ***instance-name***	(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

See "show isis database" on page 21722 for an explanation of output fields.

Sample Output

clear isis overload

The following sample output displays IS-IS database information before and after the `clear isis overload` command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
pro3-c.00-00                     0x4   0x10db     1185 L1 L2 Overload

  1 LSPs
```

```
IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
pro3-c.00-00                     0x5   0x429f     1185 L1 L2 Overload


pro2-a.00-00                   0x91e   0x2589      874 L1 L2
pro2-a.02-00                     0x1    0xcbc      874 L1 L2
  3 LSPs


user@host> clear isis overload


user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
pro3-c.00-00                     0xa   0x429e     1183 L1 L2
  1 LSPs

IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
pro3-c.00-00                     0xc   0x9c39     1183 L1 L2
pro2-a.00-00                   0x91e   0x2589      783 L1 L2
pro2-a.02-00                     0x1    0xcbc      783 L1 L2
  3 LSPs
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show isis database | 21722

clear isis spring traffic-statistics

IN THIS SECTION

- Syntax | 3673
- Description | 3673
- Options | 3673
- Required Privilege Level | 3674
- Release Information | 3674

Syntax

```
clear isis spring traffic-statistics
<all>
<interface interface>
<logical-system (all | logical-system-name>
```

Description

Clear all IS-IS SPRING traffic statistics. The Packet Forwarding Engine (PFE) does not support clearing the counters. The IS-IS routing protocol must maintain the traffic counters corresponding to the time when clear was issued and subtract it from the next sample.

Use this command when the routing protocol process (rpd) restarts and the IS-IS time stamps are lost. The existing `clear statistics` command might not clear all SPRING traffic statistics.

Options

none	Clear IS-IS SPRING traffic statistics.
all	(Optional) Clear IS-IS SPRING traffic statistics for all interfaces.
interface *interface*	(Optional) Clear all IS-IS SPRING traffic statistics for an interface.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

auto-bandwidth | 2682
traffic-statistics (Protocols IS-IS) | 30048
update-threshold-max-reservable | 30458
show auto-bandwidth | 17116
show isis spring interface traffic-statistics | 21779

clear isis statistics

IN THIS SECTION

- Syntax | 3675
- Syntax (EX Series Switches and QFX Series) | 3675
- Description | 3675
- Options | 3675
- Required Privilege Level | 3675
- Output Fields | 3675
- Sample Output | 3676
- Release Information | 3677

Syntax

```
clear isis statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
clear isis statistics
<instance instance-name>
```

Description

Set statistics about IS-IS traffic to zero.

Options

none	Set IS-IS traffic statistics to zero for all routing instances.
instance *instance-name*	(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

See "show isis statistics" on page 21783 for an explanation of output fields.

Sample Output

clear isis statistics

The following sample output displays IS-IS statistics before and after the `clear isis statistics` command is entered:

```
user@host> show isis statistics
IS-IS statistics for merino:

PDU type     Received  Processed      Drops       Sent     Rexmit
LSP             12793      12793          0       8666        719
IIH            116751     116751          0     118834          0
CSNP           203956     203956          0     204080          0
PSNP             7356       7350          6       8635          0
Unknown             0          0          0          0          0
Totals         340856     340850          6     340215        719

Total packets received: 340856 Sent: 340934

SNP queue length:           0 Drops:          0
LSP queue length:           0 Drops:          0

SPF runs:                1064
Fragments rebuilt:       1087
LSP regenerations:        436
Purges initiated:           0


user@host> clear isis statistics


user@host> show isis statistics
IS-IS statistics for merino:

PDU type     Received  Processed      Drops       Sent     Rexmit
LSP                 0          0          0          0          0
IIH                 3          3          0          3          0
CSNP                2          2          0          4          0
PSNP                0          0          0          0          0
Unknown             0          0          0          0          0
Totals              5          5          0          7          0
```

```
Total packets received: 5 Sent: 7

SNP queue length:           0 Drops:          0
LSP queue length:           0 Drops:          0

SPF runs:                   0
Fragments rebuilt:          0
LSP regenerations:          0
Purges initiated:           0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show isis statistics | 21783

clear l2-learning mac-move-buffer

IN THIS SECTION

- Syntax | 3678
- Description | 3678
- Options | 3678
- Required Privilege Level | 3678
- Output Fields | 3678
- Sample Output | 3678
- Release Information | 3679

Syntax

```
clear l2-learning mac-move-buffer
<active>
```

Description

Clear the MAC move buffer entries.

Options

none Clear the MAC move buffer entries.

active (Optional) Unblock the interfaces that were blocked by the MAC move action feature. This allows the user to keep the `reopen-time` configured to a large value, but when the looping error is fixed, the user can manually release the blocking.

Required Privilege Level

clear

Output Fields

When you enter this command, the MAC move buffer entries are deleted.

Sample Output

clear l2-learning mac-move-buffer

```
user@host> clear l2-learning mac-move-buffer
```

clear l2-learning mac-move-buffer active

```
user@host> clear l2-learning mac-move-buffer active
```

Release Information

Command introduced in Junos OS Release 13.2.

clear lacp statistics interfaces

IN THIS SECTION

- Syntax | 3679
- Description | 3679
- Options | 3679
- Required Privilege Level | 3679
- Output Fields | 3680
- Release Information | 3680

Syntax

```
clear lacp statistics interfaces <interface-name>
```

Description

Clear the LACP statistics. If you do not specify an interface name, LACP statistics for all interfaces are cleared.

Options

interface-name	(Optional) Name of an interface.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command modified in Junos OS Release 10.2.

RELATED DOCUMENTATION

show lacp statistics interfaces (View) | 21843

Verifying LACP on Redundant Ethernet Interfaces

clear ldp neighbor

IN THIS SECTION

- Syntax | 3680
- Description | 3681
- Options | 3681
- Required Privilege Level | 3681
- Output Fields | 3681
- Sample Output | 3681

Syntax

```
clear ldp neighbor
<instance instance-name>
<logical-system (all | logical-system-name)>
<neighbor>
```

Description

Tear down Label Distribution Protocol (LDP) neighbor connections.

Options

none	Tear down connections with all LDP neighbors for all routing instances.
instance ***instance name***	(Optional) Clear the LDP session for the specified routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
neighbor	(Optional) Clear an LDP session for the specified neighbor (IP address) only.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ldp neighbor

```
user@host> clear ldp neighbor
```

RELATED DOCUMENTATION

show ldp neighbor | 21868

clear ldp session

IN THIS SECTION

- Syntax | 3682
- Description | 3682
- Options | 3682
- Required Privilege Level | 3683
- Output Fields | 3683
- Sample Output | 3683
- Release Information | 3683

Syntax

```
clear ldp session
<all>
<destination>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Clear Label Distribution Protocol (LDP) sessions.

Options

all	Clear LDP sessions for all destinations for all routing instances.
destination	(Optional) Clear an LDP session for the specified destination (IP address).
instance *instance-name*	(Optional) Clear the LDP session for the specified routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ldp session

```
user@host> clear ldp session all
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ldp session | 21901

clear ldp statistics

IN THIS SECTION

- Syntax | 3684
- Description | 3684
- Options | 3684
- Required Privilege Level | 3684
- Output Fields | 3684
- Sample Output | 3684
- Release Information | 3684

Syntax

```
clear ldp statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Set all Label Distribution Protocol (LDP) statistics to zero.

Options

none	Set all LDP statistics to zero for all routing instances.
instance *instance-name*	(Optional) Clear the LDP session for the specified routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ldp statistics

```
user@host> clear ldp statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ldp statistics | 21912

show ldp traffic-statistics | 21918

clear lldp neighbors

IN THIS SECTION

- Syntax | 3685
- Description | 3685
- Options | 3685
- Required Privilege Level | 3686
- Sample Output | 3686
- Release Information | 3686

Syntax

```
clear lldp neighbors
<interface interface>
```

Description

Clear the learned remote neighbor information on all or selected interfaces.

Options

none	Clear the remote neighbor information on all interfaces.
interface *interface*	(Optional) Clear the remote neighbor information from one or more selected interfaces.

Required Privilege Level

clear

Sample Output

clear lldp neighbors

```
user@switch> clear lldp neighbors
```

clear lldp neighbors interface ge-0/1/1.0

```
user@switch> clear lldp neighbors interface ge-0/1/1.0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Understanding LLDP

Configuring LLDP (CLI Procedure)

Understanding LLDP and LLDP-MED on EX Series Switches

clear lldp statistics

IN THIS SECTION

- Syntax | 3687
- Description | 3687
- Options | 3687
- Required Privilege Level | 3687

- Output Fields | 3687
- Sample Output | 3688
- Release Information | 3688

Syntax

```
clear lldp statistics
<interface interface-name>
```

Description

Clear all Link Layer Discovery Protocols (LLDP) statistics or LLDP statistics associated with the specified interface.

For information about interface names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .

Options

none	Clear LLDP statistics on all interfaces.
interface *interface-name*	(Optional) Clear LLDP statistics on the specified interface.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show lldp statistics` command before and after clearing the LLDP statistics to verify the clear operation.

Sample Output

clear lldp statistics

```
user@switch> clear lldp statistics
```

clear lldp statistics interface ge-0/1/1.0

```
user@switch> clear lldp statistics interface ge-0/1/1.0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

clear lldp neighbors | 3685

Device Discovery Using LLDP and LLDP-MED on Switches

Device Discovery Using LLDP and LLDP-MED on Switches

Device Discovery Using LLDP and LLDP-MED on Switches

clear log

IN THIS SECTION

- Syntax | 3689
- Description | 3689
- Options | 3689
- Required Privilege Level | 3689
- Output Fields | 3689
- Sample Output | 3689

Release Information | 3690

Syntax

```
clear log filename
<all>
```

Description

Remove contents of a log file.

Options

filename Name of the specific log file to delete. Note that the file name cannot contain any special characters, including: ![=;|(){}]

all (Optional) Delete the specified log file and all archived versions of it.

Required Privilege Level

clear

Output Fields

See file list for an explanation of output fields.

Sample Output

clear log

The following sample commands list log file information, clear the contents of a log file, and then display the updated log file information:

```
user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
```

```
--------------------------------------------------------------------------
-rw-r-----  1 root  wheel      26450 Jun 23 18:47 /var/log/sampled
total 1

user@host> clear log lcc0-re0:sampled
lcc0-re0:
--------------------------------------------------------------------------

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
--------------------------------------------------------------------------
-rw-r-----  1 root  wheel         57 Sep 15 03:44 /var/log/sampled
total 1
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show log

clear log (JDM)

IN THIS SECTION

- Syntax | 3691
- Description | 3691
- Required Privilege Level | 3691
- Output Fields | 3691
- Sample Output | 3691
- Release Information | 3691

Syntax

```
clear log file-name
```

Description

Clear the system log or trace files.

Required Privilege Level

View

Output Fields

This command produces no output.

Sample Output

clear log

```
user@jdm> clear log syslog
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

monitor start (JDM) | 10930

clear log (Next Gen Services)

IN THIS SECTION

- Syntax | 3692
- Description | 3692
- Options | 3692
- Required Privilege Level | 3692
- Output Fields | 3692
- Sample Output | 3693
- Release Information | 3693

Syntax

clear log service-set | interface | file-name

Description

Clear log for service-set, interface, or file.

Options

service-set	Specify the name of the service-set for which you want to clear the log.
interface-name	Specify the name of the interface for which you want to clear the log.
file-name	Specify the file-name for which you want to clear the log.

Required Privilege Level

View

Output Fields

This command produces no output.

Sample Output

clear log

```
user@host> clear log vms 1/0/0
```

Release Information

Command introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

monitor start (JDM) | 10930

clear loop-detect enhanced interface

IN THIS SECTION

- Syntax | 3693
- Description | 3694
- Options | 3694
- Required Privilege Level | 3694
- Sample Output | 3694
- Release Information | 3695

Syntax

```
clear loop-detect enhanced interface
<interface-name>
```

Description

Restore a specified logical interface or all logical interfaces to their prior state after the device detects a loop and applies a configured action to break the loop.

You can configure loop detection on leaf devices for server-facing Layer 2 logical interfaces using the `loop-detect` statement in the `[edit protocols]` hierarchy. With that statement, you specify the action that the leaf device performs on the interface when it detects a loop, such as bringing the interface down. By default, the interface remains in the state imposed by the loop detection action until you run this command to manually restore the prior state. For example, if you configured the loop detection feature to bring the interface down upon detecting a loop, then this command brings the interface up again. You can also configure the `revert-interval` option to set a timer for after the loop condition has been repaired; when the timer expires, the device reverts the interface to its state prior to the loop detection action.

Use this command after you've repaired the loop condition to return the interface or interfaces to their prior state in either of these cases:

- You didn't configure `revert-interval`. In this case, you must run this command to restore the interface state.
- You want to restore the interface to its prior state before the configured `revert-interval` expires.

Options

interface-name

Logical interface to revert to the state prior to the loop detection action.

If you don't specify an interface name, this command defaults to restoring the state of all logical interfaces affected by the loop detection action.

Required Privilege Level

clear

Sample Output

clear loop-detect enhanced interface

```
user@host> clear loop-detect enhanced interface

user@host>
```

clear loop-detect enhanced interface *interface-name*

```
user@host> clear loop-detect enhanced interface xe-0/0/10.0

user@host>
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

EVPN-VXLAN Lightweight Leaf to Server Loop Detection

loop-detect | **10113**

show loop-detect enhanced interface | **21987**

clear mac-vrf forwarding mac-ip-table

IN THIS SECTION

- Syntax | **3695**
- Description | **3696**
- Options | **3696**
- Additional Information | **3696**
- Required Privilege Level | **3696**
- Release Information | **3696**

Syntax

```
clear mac-vrf forwarding mac-ip-table
<address>
<instance routing instance>
```

```
<ipv4>
<ipv6>
<vlan-name vlan name>
```

Description

Clear all learned media access control (MAC) address to IP address (MAC+IP) mappings from MAC-VRF type routing instances.

Options

none Clear all learned MAC+IP addresses learned through ARP from the device.

ipv4 *ipv4-address* (Optional) Clear the specified (learned) IPv4 address from the device.

ipv6 *ipv6-address* (Optional) Clear the specified (learned) IPv6 address from the device.

instance *mac-vrf-routing -nstance-name* (Optional) Clear MAC addresses from the specified MAC-VRF routing instance.

vlan-name *VLAN-Name | all* (Optional) Clear MAC address learned on the specified VLAN or all VLANs.

Additional Information

On MX Series routers and the EX9200 line of switches, this command is an alias for the `clear bridge mac-ip-table` command.

On QFX Series switches, this command is an alias for the `clear ethernet switching mac-ip-table` command.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

clear mac-vrf forwarding mac-table | 3698

clear mac-vrf forwarding mac-learning-log

IN THIS SECTION

- Syntax | 3697
- Description | 3697
- Options | 3697
- Sample Output | 3697
- Required Privilege Level | 3697
- Release Information | 3697

Syntax

```
clear mac-vrf forwarding mac-learning-log
```

Description

(QFX Series switches only) Clear the mac learning log. This action applies to all entries in the mac learning log, regardless of source.

Options

There are no options for this command.

Sample Output

The device displays no output when you issue this command.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

clear mac-vrf forwarding redundancy-group | 3701

clear mac-vrf forwarding mac-table

IN THIS SECTION

- Syntax | 3698
- Description | 3698
- Options | 3698
- Additional Information | 3699
- Required Privilege Level | 3699
- Release Information | 3699

Syntax

```
clear mac-vrf forwarding mac-table
<address>
<instance routing instance>
<interface interface name>
<isid isis id>
<persistent-learning>
<vlan-id vlan id>
<vlan-name vlan name>
```

Description

Clear all learned media access control (MAC) addresses from the MAC-VRF type routing instances.

Options

none Clear all learned MAC address entries from the device.

address (Optional) Clear the specified MAC address from the device.

instance *routing instance* (Optional) Clear MAC addresses from the specified MAC-VRF routing instance.

interface *interface name* (Optional) Clear the MAC table for the specified interface.

isid *isis id* (Optional) Clear MAC addresses learned on a specified IS-IS ID. Range: (256..16777215).

persistent-learning (Optional) Clear persistent MAC entries.

vlan-id *VLAN ID* (Optional) Clear MAC addresses learned on the specified VLAN id. Range: (0..4095).

vlan-name (all | *VLAN Name***)** (Optional) Clear MAC addresses learned on the specified VLAN or all VLAN names.

Additional Information

On MX Series routers and the EX9200 line of switches, this command is an alias for the `clear bridge evpn` command.

On QFX Series switches, this command is an alias for the `clear ethernet switching evpn` command.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

clear mac-vrf forwarding mac-ip-table | **3695**

clear mac-vrf forwarding redundancy-group | **3701**

clear mac-vrf forwarding recovery-timeout

IN THIS SECTION

- Syntax | 3700
- Description | 3700
- Options | 3700
- Required Privilege Level | 3700
- Release Information | 3700

Syntax

```
clear mac-vrf forwarding recovery-timeout
<interface interface name>
```

Description

Clear all port errors that the device registered due to MAC address limits. If you configure MAC limiting with the `shutdown` option, the device shuts down the affected interfaces until you clear the error.

Options

No options exist for this command.

interface ***interface name***	(Optional) Clear only those port errors that the device has registered on a specific interface due to MAC address limits.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

clear mac-vrf forwarding mac-ip-table | 3695

clear mac-vrf forwarding redundancy-group

IN THIS SECTION

- Syntax | 3701
- Description | 3701
- Options | 3701
- Required Privilege Level | 3702
- Release Information | 3702

Syntax

```
clear mac-vrf forwarding redundancy-group
<redundancy-group-id>
<arp-statistics>
<nd-statistics>
```

Description

Clear redundancy group statistics counters within all MAC-VRF routing instances.

Options

none Clear all redundancy group statistics counters.

redundancy-group-id (Optional) Clear the counters for the specified redundancy group.

arp-statistics (Optional) Clear multichassis aggregated-Ethernet ARP synchronization statistics.

nd-statistics (Optional) Clear multichassis aggregated-Ethernet neighbor discovery protocol (ND) synchronization statistics.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

clear mac-vrf forwarding mac-table | 3698

clear mld membership

IN THIS SECTION

- Syntax | 3702
- Description | 3703
- Options | 3703
- Required Privilege Level | 3703
- Output Fields | 3703
- Sample Output | 3703
- Release Information | 3703

Syntax

```
clear mld membership
<all>
<group group-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Description

Clear Multicast Listener Discovery (MLD) group membership.

Options

all	Clear MLD memberships for groups and interfaces in the master instance.
group *group-name*	(Optional) Clear MLD membership for the specified group.
interface *interface-name*	(Optional) Clear MLD group membership for the specified interface.
logical-system (all | *logical-system-name)*	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mld membership all

```
user@host> clear mld membership all
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show mld group | 22066

clear mld snooping membership

IN THIS SECTION

- Syntax | 3704
- Description | 3704
- Options | 3704
- Required Privilege Level | 3704
- Sample Output | 3705
- Release Information | 3705

Syntax

```
clear mld snooping membership
<vlan vlan-name>
```

Description

Clear MLD snooping dynamic membership information from the multicast forwarding table.

Options

none	Clear dynamic membership information for all VLANs.
vlan *vlan-name*	(Optional) Clear dynamic membership information for the specified VLAN.

Required Privilege Level

view

Sample Output

clear mld snooping membership vlan employee-vlan

```
user@host> clear mld snooping membership vlan employee-vlan
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding MLD Snooping

Example: Configuring MLD Snooping on SRX Series Devices

mld-snooping | **10783**

show mld snooping membership | **22090**

clear mld snooping statistics | **3705**

clear mld snooping statistics

IN THIS SECTION

- Syntax | **3706**
- Description | **3706**
- Required Privilege Level | **3706**
- Sample Output | **3706**
- Release Information | **3706**

Syntax

```
clear mld snooping statistics
```

Description

Clear MLD snooping statistics.

Required Privilege Level

view

Sample Output

clear mld snooping statistics

```
user@host> clear mld snooping statistics
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding MLD Snooping

Example: Configuring MLD Snooping on SRX Series Devices

mld-snooping | **10783**

show mld snooping statistics | **22098**

clear mld snooping membership | **3704**

clear mld statistics

IN THIS SECTION

- Syntax | 3707
- Syntax (MX Series) | 3707
- Description | 3707
- Options | 3708
- Required Privilege Level | 3708
- Output Fields | 3708
- Sample Output | 3708
- Release Information | 3709

Syntax

```
clear mld statistics
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (MX Series)

```
clear mld statistics
(<continuous> | <interface interface-name>)
<logical-system (all | logical-system-name)>
```

Description

Clear Multicast Listener Discovery (MLD) statistics. Clearing MLD statistics zeros the statistics counters as if you rebooted the device.

By default, Junos OS multicast devices collect statistics of received and transmitted MLD control messages that reflect currently active multicast group subscribers. Some devices also automatically maintain *continuous* MLD statistics globally on the device in addition to the default active subscriber statistics—these are persistent, continuous statistics of received and transmitted MLD control packets

that account for both past and current multicast group subscriptions processed on the device. The device maintains continuous statistics across events or operations such as routing daemon restarts, graceful Routing Engine switchovers (GRES), in-service software upgrades (ISSU), or line card reboots. The default active subscriber-only statistics are not preserved in these cases.

Run this command to clear the currently active subscriber statistics. On devices that support continuous statistics, run this command with the `continuous` option to clear the continuous statistics. You must run these commands separately to clear both types of statistics because the device maintains and clears the two types of statistics separately.

Options

none	(Same as **logical-system all**) Clear MLD statistics for all interfaces. This form of the command clears statistics for currently active subscribers only.
continuous	Clear only the continuous MLD statistics that account for both past and current multicast group subscribers instead of the default statistics that only reflect currently active subscribers. This option is not available with the `interface` option for interface-specific statistics.
interface ***interface-name***	(Optional) Clear MLD statistics for the specified interface. This option is not available with the `continuous` option.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mld statistics

```
user@host> clear mld statistics
```

clear mld statistics continuous

```
user@host> clear mld statistics continuous
```

Release Information

Command introduced before Junos OS Release 7.4.

`continuous` option added in Junos OS Release 19.4R1 for MX Series routers.

RELATED DOCUMENTATION

show mld statistics | 22104

clear monitor security flow filter

IN THIS SECTION

- Syntax | 3709
- Description | 3710
- Options | 3710
- Required Privilege Level | 3710
- Output Fields | 3710
- Sample Output | 3710
- Release Information | 3710

Syntax

```
clear monitor security flow filter <filter-name>
```

Description

Specify the security flow filters to be deleted. Once deleted, the filters are removed from the Packet Forwarding Engine and the Routing Engine. .

NOTE: Specifying the filter name is optional. If no filter is specified, all filters are deleted.

Options

This command has no options.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

Release Information

Command introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Monitoring Security Flow Sessions Overview

monitor security flow start | **10918**

monitor security flow filter | **10916**

clear mpls container-lsp

IN THIS SECTION

- Syntax | 3711
- Description | 3711
- Options | 3711
- Required Privilege Level | 3712
- Output Fields | 3712
- Sample Output | 3712
- Release Information | 3713

Syntax

```
clear mpls  container-lsp
<autobandwidth>
<history>
<logical-system (all | logical-system-name)>
<member>
<name name>
<optimize | optimize-aggressive>
<statistics>
```

Description

Release the routes and states associated with MPLS container label-switched paths (LSPs), and start new LSPs.

Options

none Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.

autobandwidth	(Optional) Clear LSP autobandwidth counters.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
name *name*	(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the *Junos Network Interfaces Configuration Guide*.
optimize | optimize-aggressive	(Optional) Run nonpreemptive optimization or aggressive optimization computation now.
statistics	(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (`name` and `path` options) on transit routers.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mpls container-lsp

```
user@host> clear mpls  container-lsp
```

clear mpls container-lsp name

```
user@host> clear mpls  container-lsp name name
```

clear mpls container-lsp statistics

```
user@host> clear mpls container-lsp statistics
```

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.

RELATED DOCUMENTATION

show mpls container-lsp | 22136

request mpls container-lsp | 14689

clear mpls lsp

IN THIS SECTION

- Syntax | 3713
- Syntax (EX and QFX Series Switches) | 3714
- Description | 3714
- Options | 3714
- Required Privilege Level | 3715
- Output Fields | 3715
- Sample Output | 3715
- Release Information | 3715

Syntax

```
clear mpls lsp
<all>
<autobandwidth>
<counters>
<logical-system (all | logical-system-name)>
<name name>
<optimize | optimize-aggressive>
```

```
<path regular-expression>
<statistics>
```

Syntax (EX and QFX Series Switches)

```
clear mpls lsp
<all>
<autobandwidth>
<name name>
<optimize | optimize-aggressive>
<path regular-expression>
<statistics>
```

Description

Release the routes and states associated with MPLS label-switched paths (LSPs), and start new LSPs.

> **CAUTION**: This command disconnects existing Resource Reservation Protocol (RSVP) sessions on the ingress routing device. If there is a time lag between the old path being torn down and the new path being set up, this command might impact traffic traveling along the LSPs.

Options

all	Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.
autobandwidth	(Optional) Clear LSP autobandwidth counters.
counters	(Optional) Reset the flap and the MBB counters to zero.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
name *name*	(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the *Junos Network Interfaces Configuration Guide*.

optimize | optimize-aggressive	(Optional) Run nonpreemptive optimization or aggressive optimization computation now.
path ***regular-expression***	(Optional) Clear the specific LSP path matching the specified regular expression.
statistics	(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (`name` and `path` options) on transit routers.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mpls lsp all

```
user@host> clear mpls lsp all
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show mpls lsp | 22172

show rsvp session | 23860

clear mpls tunnel manager statistics

IN THIS SECTION

- Syntax | 3716
- Description | 3716
- Options | 3716
- Required Privilege Level | 3716
- Sample Output | 3717
- Release Information | 3717

Syntax

```
clear mpls tunnel manager statistics
<logical-system (all | logical-system-name)>
```

Description

Clear the counters that track the number of Multiprotocol Label Switching (MPLS) label-switched paths (LSPs) that have been updated in-place without any change to LSP identifier (ID).

Options

none	Clear in-place updated LSP counter statistics.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Sample Output

clear mpls tunnel manager statistics

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show mpls tunnel manager statistics | 22219

clear msdp cache

IN THIS SECTION

- Syntax | 3717
- Description | 3718
- Options | 3718
- Required Privilege Level | 3718
- Output Fields | 3718
- Sample Output | 3718
- Release Information | 3718

Syntax

```
clear msdp cache
<all>
<instance instance-name>
<logical-system (all | logical-system-name)>
<peer peer-address>
```

Description

Clear the entries in the Multicast Source Discovery Protocol (MSDP) source-active cache.

Options

all	Clear all MSDP source-active cache entries in the master instance.
instance *instance-name*	(Optional) Clear entries for a specific MSDP instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
peer *peer-address*	(Optional) Clear the MSDP source-active cache entries learned from a specific peer.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear msdp cache all

```
user@host> clear msdp cache all
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show msdp source-active | 22228

clear msdp statistics

IN THIS SECTION

- Syntax | 3719
- Description | 3719
- Options | 3719
- Required Privilege Level | 3720
- Output Fields | 3720
- Sample Output | 3720
- Release Information | 3720

Syntax

```
clear msdp statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
<peer peer-address>
```

Description

Clear Multicast Source Discovery Protocol (MSDP) peer statistics.

Options

none	Clear MSDP statistics for all peers.
instance *instance-name*	(Optional) Clear statistics for the specified instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
peer *peer-address*	(Optional) Clear the statistics for the specified peer.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear msdp statistics

```
user@host> clear msdp statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show msdp statistics | 22232

clear multicast bandwidth-admission

IN THIS SECTION

- Syntax | 3721
- Description | 3721
- Options | 3721
- Required Privilege Level | 3722
- Output Fields | 3722
- Sample Output | 3722
- Release Information | 3722

Syntax

```
clear multicast bandwidth-admission
<group group-address>
<inet | inet6>
<instance instance-name>
<interface interface-name>
<source source-address>
```

Description

Reapply IP multicast bandwidth admissions.

Options

none	Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.
group *group-address*	(Optional) Reapply multicast bandwidth admissions for the specified group.
inet	(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.
inet6	(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.
instance *instance-name*	(Optional) Reapply multicast bandwidth admission settings for the specified instance. If you do not specify an instance, the command applies to the master routing instance.
interface *interface-name*	(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows: • If the interface is congested, and it was admitted previously, it is removed. • If the interface was rejected previously, the `clear multicast bandwidth-admission` command enables the interface to be admitted as long as enough bandwidth exists on the interface. • If you do not specify an interface, issuing the `clear multicast bandwidth-admission` command readmits any previously rejected interface for the relevant entries as long as enough bandwidth exists on the interface.

To manually reject previously admitted outbound interfaces, you must specify the interface.

source *source-address* (Optional) Use with the **group** option to reapply multicast bandwidth admission settings for the specified (source, group) entry.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast bandwidth-admission

```
user@host> clear multicast bandwidth-admission
```

Release Information

Command introduced in Junos OS Release 8.3.

inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.

RELATED DOCUMENTATION

show multicast interface | 22286

clear multicast forwarding-cache

IN THIS SECTION

- Syntax | 3723
- Description | 3723
- Options | 3723
- Required Privilege Level | 3724
- Output Fields | 3724
- Sample Output | 3724
- Release Information | 3724

Syntax

```
clear multicast forwarding-cache
<all>
<inet | inet6>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Clear IP multicast forwarding cache entries.

This command is not supported for next-generation multiprotocol BGP multicast VPNs (MVPNs).

Options

all	Clear all multicast forwarding cache entries in the master instance.
inet	(Optional) Clear multicast forwarding cache entries for IPv4 family addresses.
inet6	(Optional) Clear multicast forwarding cache entries for IPv6 family addresses.

instance *instance-name*	(Optional) Clear multicast forwarding cache entries on a specific routing instance.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast forwarding-cache all

```
user@host> clear multicast forwarding-cache all
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

show multicast forwarding-cache statistics | 22283

clear multicast scope

IN THIS SECTION

- Syntax | 3725
- Syntax (EX Series Switch and the QFX Series) | 3725

- Description | 3725
- Options | 3725
- Required Privilege Level | 3726
- Output Fields | 3726
- Sample Output | 3726
- Release Information | 3726

Syntax

```
clear multicast scope
<inet | inet6>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
clear multicast scope
<inet | inet6>
<interface interface-name>
```

Description

Clear IP multicast scope statistics.

Options

none	(Same as **logical-system all**) Clear multicast scope statistics.
inet	(Optional) Clear multicast scope statistics for IPv4 family addresses.
inet6	(Optional) Clear multicast scope statistics for IPv6 family addresses.
interface *interface-name*	(Optional) Clear multicast scope statistics on a specific interface.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast scope

```
user@host> clear multicast scope
```

Release Information

Command introduced in Junos OS Release 7.6.

inet6 option introduced in Junos OS Release 10.0 for EX Series switches.

RELATED DOCUMENTATION

show multicast scope | 22328

clear multicast sessions

IN THIS SECTION

- Syntax | 3727
- Syntax (EX Series Switch and the QFX Series) | 3727
- Description | 3727

- Options | 3727
- Required Privilege Level | 3727
- Output Fields | 3728
- Sample Output | 3728
- Release Information | 3728

Syntax

```
clear multicast sessions
<logical-system (all | logical-system-name)>
<regular-expression>
```

Syntax (EX Series Switch and the QFX Series)

```
clear multicast sessions
<regular-expression>
```

Description

Clear IP multicast sessions.

Options

none	(Same as **logical-system all**) Clear multicast sessions.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
regular-expression	(Optional) Clear only multicast sessions that contain the specified regular expression.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast sessions

```
user@host> clear multicast sessions
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show multicast sessions | 22331

clear multicast snooping statistics

IN THIS SECTION

- Syntax | 3729
- Description | 3729
- Options | 3729
- Required Privilege Level | 3729
- Output Fields | 3729
- Sample Output | 3729
- Release Information | 3730

Syntax

```
clear multicast snooping statistics
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Description

Clear IP multicast snooping statistics.

Options

none	Clear multicast snooping statistics for all supported address families on all interfaces.
instance *instance-name*	(Optional) Clear multicast snooping statistics for the specified instance.
interface *interface-name*	(Optional) Clear multicast snooping statistics on a specific interface.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast snooping statistics

```
user@host> clear multicast snooping statistics
```

Release Information

Command introduced in Junos OS Release 8.5.

clear multicast statistics

IN THIS SECTION

- Syntax | 3730
- Syntax (EX Series Switch and the QFX Series) | 3730
- Syntax (EX4300 Switch) | 3731
- Description | 3731
- Options | 3731
- Required Privilege Level | 3731
- Output Fields | 3731
- Sample Output | 3732
- Release Information | 3732

Syntax

```
clear multicast statistics
<inet | inet6>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
clear multicast statistics
<inet | inet6>
<instance instance-name>
<interface interface-name>
```

Syntax (EX4300 Switch)

```
clear system-packet-forwarding-options multicast-statistics
```

There are no available options for the EX4300.

Description

Clear IP multicast statistics.

Options

none	Clear multicast statistics for all supported address families on all interfaces.
inet	(Optional) Clear multicast statistics for IPv4 family addresses.
inet6	(Optional) Clear multicast statistics for IPv6 family addresses.
instance *instance-name*	(Optional) Clear multicast statistics for the specified instance.
interface *interface-name*	(Optional) Clear multicast statistics on a specific interface.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you get feedback on the status of your request.

Sample Output

clear multicast statistics

```
user@host> clear multicast statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.

Syntax added in Junos OS Release 19.2R1 for clearing multicast route statistics (EX4300 switches).

RELATED DOCUMENTATION

show multicast statistics | 22356

clear mvrp statistics

IN THIS SECTION

- Syntax (EX Series) | 3733
- Syntax (SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320) | 3733
- Description | 3733
- Options | 3733
- Required Privilege Level | 3733
- Output Fields | 3733
- Sample Output | 3734
- Release Information | 3734

Syntax (EX Series)

```
clear mvrp statistics <interface interface-name>
```

Syntax (SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320)

```
clear mvrp statistics
<interface interface-name>
<routing-instance routing-instance-name>
```

Description

Clear all Multiple VLAN Registration Protocol (MVRP) interface and, for SRX Series Firewalls, routing instances statistics.

Options

none	Clear all MVRP statistics.
interface *interface-name*	Clear the MVRP statistics on the specified interface.
routing-instance name	Clear the MVRP statistics on the specified SRX Series Firewall's named routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mvrp statistics

```
user@switch> clear mvrp statistics
```

clear mvrp statistics interface ge-0/0/1.0

```
user@switch> clear mvrp statistics interface ge-0/0/1.0
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show mvrp statistics | 22406

Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches

show mvrp | 22390

clear neighbor-discovery-inspection statistics

IN THIS SECTION

- Syntax | 3735
- Description | 3735
- Options | 3735
- Required Privilege Level | 3735
- Output Fields | 3735
- Sample Output | 3735
- Release Information | 3735

Syntax

```
clear neighbor-discovery-inspection statistics
<interface interface-name>
```

Description

Clear IPv6 neighbor discovery inspection statistics.

Options

none	Clear neighbor discovery inspection statistics on all interfaces.
interface *interface-name*	(Optional) Clear neighbor discovery inspection statistics on one or more interfaces.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear neighbor-discovery-inspection statistics

```
user@switch> clear neighbor-discovery-inspection statistics
```

Release Information

Command introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

show neighbor-discovery-inspection statistics | 22413

Example: Configuring Port Security (non-ELS)

clear network-access aaa statistics

IN THIS SECTION

- Syntax | 3736
- Description | 3736
- Options | 3737
- Required Privilege Level | 3737
- Output Fields | 3737
- Sample Output | 3738
- Release Information | 3740

Syntax

```
clear network-access aaa statistics
<accounting>
<address-assignment (client | pool pool-name)>
<authentication>
<dynamic-requests>
<radius>
<re-authentication>
<saml>
<session-limit-per-username username username access-profile profile-name>
<terminate-code>
```

Description

Clear AAA statistics.

Options

accounting (Optional) Clear AAA accounting statistics.

address-assignment client (Optional) Clear AAA address-assignment statistics for the client.

address-assignment pool *pool-name* (Optional) Clear AAA address-assignment pool statistics.

authentication (Optional) Clear AAA authentication statistics.

dynamic-requests (Optional) Clear AAA dynamic-request statistics.

radius (Optional) Clears the values in the `Peak` and `Exceeded` columns only.

re-authentication (Optional) Clear AAA reauthentication statistics.

saml (Optional) Clear Security Assertion Markup Language (SAML) AAA statistics.

session-limit-per-username (MX Series routers only) (Optional) Clear all blocked request statistics for all access profiles from the username session-limit table. You can also specify additional options:

- `username` *`username`*—Clear the blocked request statistics for the specified username across all access profiles. A given username can be used in more than one access profile.
- `access-profile` *`profile-name`*—Clear the blocked request statistics for all usernames in the specified access profile.

> **NOTE**: This command does not clear (delete) the entry in the session-limit table. Entries in the table are added or deleted during session login or logout processing.

terminate-code (Optional) Clear AAA termination code statistics.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear network-access aaa statistics accounting

```
user@host> clear network-access aaa statistics accounting
```

clear network-access aaa statistics address-assignment pool

```
user@host> clear network-access aaa statistics address-assignment pool isp_1
```

clear network-access aaa statistics radius

```
user@host> clear network-access aaa statistics radius
```

clear network-access aaa statistics saml

```
user@host> clear network-access aaa statistics saml
```

clear network-access aaa statistics session-limit-per-username (All Usernames Across All Access Profiles)

```
user@host> show network-access aaa statistics on-limit-per-username detail
Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               1                    4
xyz@example.net    BNG1               3                    5
abc@example.net    BNG2               2                    5
pqr@example.net    BNG2               3                    4

user@host> clear network-access aaa statistics session-limit-per-username
user@host> show network-access aaa statistics on-limit-per-username detail
Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               0                    4
xyz@example.net    BNG1               0                    5
abc@example.net    BNG2               0                    5
pqr@example.net    BNG2               0                    4
```

clear network-access aaa statistics session-limit-per-username (Specific Username Across All Access Profiles)

```
user@host> show network-access aaa statistics on-limit-per-username detail

Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               1                    4
xyz@example.net    BNG1               3                    5
rkv@example.net    BNG2               2                    5
pqr@example.net    BNG2               3                    4
user@host> clear network-access aaa statistics session-limit-per-username username
rkv@example.net
user@host> show network-access aaa statistics on-limit-per-username detail

Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               0                    4
xyz@example.net    BNG1               3                    5
rkv@example.net    BNG2               0                    5
pqr@example.net    BNG2               3                    4
```

clear network-access aaa statistics session-limit-per-username (All Usernames for Specific Access Profiles)

```
user@host> show network-access aaa statistics on-limit-per-username detail

Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               1                    4
xyz@example.net    BNG1               3                    5
rkv@example.net    BNG2               2                    5
pqr@example.net    BNG2               3                    4
user@host> clear network-access aaa statistics session-limit-per-username access-profile BNG2
user@host> show network-access aaa statistics on-limit-per-username detail

Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               1                    4
xyz@example.net    BNG1               3                    5
123@example.net    BNG2               0                    5
pqr@example.net    BNG2               0                    4
```

clear network-access aaa statistics session-limit-per-username (Specific Username in Specific Access Profile)

```
user@host> show network-access aaa statistics on-limit-per-username detail

Username            Access-profile     Blocked requests     Session count
rkv@example.net     BNG1               1                    4
xyz@example.net     BNG1               3                    5
rkv@example.net     BNG2               2                    5
pqr@example.net     BNG2               3                    4
user@host> clear network-access aaa statistics session-limit-per-username username
rkv@example.net access-profile BNG2
user@host> show network-access aaa statistics on-limit-per-username detail

Username            Access-profile     Blocked requests     Session count
rkv@example.net     BNG1               3                    4
xyz@example.net     BNG1               3                    5
rkv@example.net     BNG2               0                    5
pqr@example.net     BNG2               3                    4
```

Release Information

Command introduced in Junos OS Release 10.0.

`radius` option introduced in Junos OS Release 11.4

`terminate-code` option introduced in Junos OS Release 11.4.

`session-limit-per-username` option introduced in Junos OS Release 18.4R1 on MX Series routers.

Support for `saml` added in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

- Verifying and Managing Subscriber AAA Information
- Understanding Session Options for Subscriber Access
- Limiting the Number of Active Sessions per Username and Access Profile
- show network-access aaa statistics

clear network-access aaa subscriber

IN THIS SECTION

- Syntax | 3741
- Description | 3741
- Options | 3741
- Required Privilege Level | 3742
- Output Fields | 3742
- Sample Output | 3742
- Release Information | 3743

Syntax

```
clear network-access aaa subscriber
<session-id identifier <reconnect>>
<statistics username username>
<username username <reconnect>>
```

Description

Clear AAA subscriber statistics and log out subscribers. You can log out subscribers based on the username or on the subscriber session identifier. Use the session identifier when more than one session has the same username string.

Options

reconnect (Optional) Reconnect as a Layer 2 wholesale session when the subscriber session has been fully logged out. This option is equivalent to issuing a RADIUS-initiated disconnect with reconnect semantics; that is, when the message includes Acct-Terminate-Cause (RADIUS attribute 49) with a value of callback (16). You can apply this option to either a Layer 2 wholesale session or a conventionally auto-sensed dynamic VLAN supporting a PPPoE session.

In the latter case, this option triggers a PPPoE session logout and removal of the dynamic VLAN logical interface. This is followed by authorization of the access-line to attempt creation of a dynamic VLAN IFL supporting Layer 2 wholesale session in its place.

session-id *identifier* — (Optional) Log out the subscriber based on the subscriber session identifier.

statistics username *username* — (Optional) Clear AAA subscriber statistics and log out the subscriber.

username *username* — (Optional) Log out the AAA subscriber.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear network-access aaa subscriber statistics username

```
user@host> clear network-access aaa subscriber statistics username user22@example.com
```

clear network-access aaa subscriber statistics username (Tenant systems)

```
user@host:TSYS1> clear network-access aaa subscriber statistics username user22@example.com
```

clear network-access aaa subscriber username

```
user@host> clear network-access aaa subscriber username user22@example.com
```

clear network-access aaa subscriber username (Tenant systems)

```
user@host:TSYS1> clear network-access aaa subscriber username user22@example.com
```

clear network-access aaa subscriber session-id

```
user@host> clear network-access aaa subscriber session-id 18367425
```

clear network-access aaa subscriber session-id (Tenant systems)

```
user@host:TSYS1> clear network-access aaa subscriber session-id 1
```

Release Information

Command introduced in Junos OS Release 9.1.

`reconnect` and `session-id` options added in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information

clear network-access gx-plus replay

IN THIS SECTION

- Syntax | 3744
- Description | 3744
- Options | 3744
- Required Privilege Level | 3744
- Output Fields | 3744
- Sample Output | 3744

Release Information | 3744

Syntax

```
clear network-access gx-plus replay
```

Description

Clear pending Gx-Plus login and logout requests (replays). Sends JSER message to PCRF that includes the Juniper-Event-Type AVP (AVP code 2103) with a value of 3 indicating a discovery request. The PCRF returns a JDER message to initiate discovery of all subscribers. When this discovery completes, all pending subscriber requests are cleared.

Options

This command has no options.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear network-access gx-plus replay

```
user@host> clear network-access gx-plus replay
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

clear network-access gx-plus statistics | 3745

show network-access gx-plus | 22492

clear network-access gx-plus statistics

IN THIS SECTION

- Syntax | 3745
- Description | 3745
- Options | 3745
- Required Privilege Level | 3745
- Output Fields | 3746
- Sample Output | 3746
- Release Information | 3746

Syntax

```
clear network-access gx-plus statistics
```

Description

Clear Gx-Plus statistics.

Options

This command has no options.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear network-access gx-plus statistics

```
user@host> clear network-access gx-plus statistics
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

show network-access gx-plus | 22492

clear network access nas-5g statistics

IN THIS SECTION

- Syntax | 3746
- Description | 3747
- Options | 3747
- Required Privilege Level | 3747
- Release Information | 3747

Syntax

```
clear network-access nas-5g statistics
```

Description

Clear the counter tracking network NAS 5G messages.

Options

none

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 22.3R1.

clear network-access ocs statistics

IN THIS SECTION

- Syntax | 3747
- Description | 3748
- Required Privilege Level | 3748
- Output Fields | 3748
- Sample Output | 3748
- Release Information | 3748

Syntax

```
clear network-access ocs statistics
```

Description

Clear Online Charging System (OCS) provisioning statistics information.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear network-access ocs statistics

```
user@host> clear network-access ocs statistics
```

Release Information

Command introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

ocs (Diameter Applications) | **11771**

show network-access ocs | **22520**

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

clear network-access pcrf

IN THIS SECTION

- Syntax | 3749
- Description | 3749
- Options | 3749
- Required Privilege Level | 3749
- Output Fields | 3750
- Sample Output | 3750
- Sample Output | 3750
- Release Information | 3750

Syntax

```
clear network-access pcrf (statistics | subscribers)
```

Description

Clear Policy and Charging Rules Function (PCRF) provisioning statistics and subscribers information.

Options

statistics (Optional) Clear PCRF provisioning statistics.

subscribers (Optional) Force logout of all PCRF subscribers. Used with draining mode to flush all subscribers before making substantial configuration changes.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear network-access pcrf statistics

```
user@host> clear network-access pcrf statistics
```

Sample Output

clear network-access pcrf subscribers

```
user@host> clear network-access pcrf subscribers
```

Release Information

Command introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

pcrf (Diameter Applications) | **12325**

show network-access pcrf | **22523**

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

clear network-access radsec state

IN THIS SECTION

- Syntax | 3751
- Description | 3751
- Options | 3751
- Required Privilege Level | 3751
- Output Fields | 3751
- Release Information | 3752

Syntax

```
clear network-access radsec state
<destination destination-id>
```

Description

Clear the connection state information for RADSEC destinations.

Options

destination ***destination-id*** (Optional) Clear connection state information for the specified RADSEC destination.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

RADIUS over TLS (RADSEC)

clear network-access radsec statistics

IN THIS SECTION

- Syntax | 3752
- Description | 3752
- Options | 3752
- Required Privilege Level | 3753
- Output Fields | 3753
- Release Information | 3753

Syntax

```
clear network-access radsec statistics
<destination destination-id>
```

Description

Clear the connection statistics for RADSEC destinations.

Options

destination ***destination-id*** (Optional) Clear connection statistics for the specified RADSEC destination.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

RADIUS over TLS (RADSEC)

clear network-access requests pending

IN THIS SECTION

- Syntax | 3753
- Description | 3754
- Options | 3754
- Required Privilege Level | 3754
- Output Fields | 3754
- Sample Output | 3754
- Release Information | 3755

Syntax

```
clear network-access requests pending
<index            index-number        >
```

Description

Clear or cancel all pending authentication requests.

Options

- none—Clear all network access requests pending.
- `index` *`index-number`* —Clear the specified authentication request. To display index numbers, use the `show network-access requests pending` command.

Required Privilege Level

clear

Output Fields

Sample Output

The following example displays the network access requests that are pending, clears the requests, and displays the results of the clear operation:

clear network-access requests pending

```
user@host> show network-access requests pending
Information about pending authentication entries
  Total pending authentication requests: 2
Index User                  Status
1     Sun                   Processing
2     Sam                   Processed

user@host> clear network-access requests pending
user@host> show network-access requests pending
Information about pending authentication entries
  Total pending authentication requests: 2
Index User                  Status
1     Sun                   Cancelled by Admin
2     Sam                   Cancelled by Admin
```

Release Information

Command introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

show network-access requests pending | 22538

clear network-access requests statistics

IN THIS SECTION

- Syntax | 3755
- Description | 3755
- Required Privilege Level | 3755
- Output Fields | 3755
- Release Information | 3756

Syntax

```
clear network-access requests statistics
```

Description

Clear general authentication statistics for the configured authentication type.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

authentication-order (Access Profile) | 2635

show network-access requests statistics | 22542

clear network-access securid-node-secret-file

IN THIS SECTION

- Syntax | 3756
- Description | 3756
- Required Privilege Level | 3756
- Output Fields | 3757
- Sample Output | 3757
- Release Information | 3757

Syntax

```
clear network-access securid-node-secret-file
```

Description

Delete the node secret file for the SecurID authentication type.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear network-access securid-node-secret-file

```
user@host> clear network-access securid-node-secret-file
```

Release Information

Command introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Firewall User Authentication Overview

configuration-file | 4514

securid-server | 16447

show network-access securid-node-secret-file | 22547

clear network-access aaa saml assertion-cache

IN THIS SECTION

- Syntax | 3758
- Description | 3758
- Options | 3758
- Required Privilege Level | 3758
- Release Information | 3758

Syntax

```
clear network-access aaa saml assertion-cache
<domain-name domain-name>
<username username>
```

Description

Clear Security Assertion Markup Language (SAML) authentication assertion cache information. If you don't specify an option, the firewalls clears all the entries by default.

Options

domain-name *domain-name*
Specify the domain name such as example.net or any. This domain should match the domain configured in the profile at `set access profile demo saml identity-provider`.

username *username*
Specify the SAML authenticated username. This username is the same as the user login ID.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

saml | 16281

show network-access aaa saml assertion-cache | 22560

clear network-access aaa saml idp-metadata

IN THIS SECTION

- Syntax | 3759
- Description | 3759
- Options | 3759
- Required Privilege Level | 3759
- Release Information | 3759

Syntax

```
clear network-access aaa saml idp-metadata
<all>
<id id>
```

Description

Clear Security Assertion Markup Language (SAML) authentication identity provider (IdP) metadata.

Options

all Clear all IdP metadata.

id *id* Specify the IdP metadata ID.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

saml | 16281

request network-access aaa saml load-idp-metadata | 14707

clear node reboot

IN THIS SECTION

- Syntax | 3760
- Description | 3760
- Sample Output | 3760
- Release Information | 3760

Syntax

```
clear node reboot node-name
```

Description

When you use this command, Junos OS Evolved cancels any pending reboots or shutdowns. If there are no such commands, a fail message gets returned.

Sample Output

clear node reboot

```
user@host> clear node reboot re0
Cancel reboot of node at Tue Feb 22 13:31:09 2022
```

Release Information

Command introduced in Junos OS Evolved 19.2R1.

clear oam ethernet connectivity-fault-management continuity-measurement

IN THIS SECTION

- Syntax | 3761
- Description | 3761
- Options | 3761
- Required Privilege Level | 3762
- Output Fields | 3762
- Sample Output | 3762
- Release Information | 3762

Syntax

```
clear oam ethernet connectivity-fault-management continuity-measurement
maintenance-domain md-name
maintenance-association ma-name
<local-mep local-mep-id>
<remote-mep remote-mep-id>
```

Description

For all routers that support IEEE 802.1ag OAM connectivity fault management (CFM), clear the existing continuity measurement and restart counting the operational uptime (that is, the total time during which CCM adjacency is active for a particular remote MEP.).

Options

maintenance-domain *md-name* Name of an existing CFM maintenance domain.

maintenance-association *ma-name* Name of an existing CFM maintenance association.

local-mep *local-mep-id*	(Optional) Display connectivity fault management information for the specified local MEP only.
remote-mep *remote-mep-id*	(Optional) Display connectivity fault management information for the specified remote MEP only.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault- management continuity-measurement

```
user@host> clear oam ethernet connectivity-fault-management continuity-measurement maintenance-
domain md5 maintenance-association ma5 local-mep 100 remote-mep 102
Continuity measurement restarted.
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Managing Continuity Measurement Statistics

Ethernet Interfaces User Guide for Routing Devices

clear oam ethernet connectivity-fault-management delay-statistics

IN THIS SECTION

- Syntax | 3763
- Description | 3763
- Options | 3763
- Required Privilege Level | 3764
- Output Fields | 3764
- Sample Output | 3764
- Release Information | 3764

Syntax

```
clear oam ethernet connectivity-fault-management delay-statistics
maintenance-association maintenance-association-name
maintenance-domain maintenance-domain-name
<logical-system logical-system-name>
<one-way>
<two-way>
```

Description

On MX Series routers and EX Series switches, clear ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM) delay statistics and ETH-DM frame counts.

Options

maintenance-association *maintenance-association-name*	Clear ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance association.
maintenance-domain *maintenance-domain-name*	Clear ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance domain.

logical-system *logical-system-name*	(MX Series routers only) (Optional) Clear ETH-DM delay statistics and ETH-DM frame counts for the specified logical system.
one-way	(Optional) Clear one-way ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance association, maintenance domain, or (on the routers only) logical system.
two-way	(Optional) Clear two-way ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance association, maintenance domain, or (on the routers only) logical system.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault- management delay statistics

```
user@switch> clear oam ethernet connectivity-fault-management delay-statistics maintenance-
domain md1 maintenance-association ma1
Delay statistics entries cleared
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

clear oam ethernet connectivity-fault-management statistics | 3775

show oam ethernet connectivity-fault-management delay-statistics | 22585

show oam ethernet connectivity-fault-management interfaces | 22596

clear oam ethernet connectivity-fault-management event

IN THIS SECTION

- Syntax | 3765
- Description | 3765
- Options | 3765
- Required Privilege Level | 3766
- Output Fields | 3766
- Sample Output | 3766
- Release Information | 3767

Syntax

```
clear oam ethernet connectivity-fault-management event
(adjacency-loss| all-defects | cross-connect-ccm | erroneous-ccm | receive-ais)
maintenance-domain md-name
maintenance-association ma-name
local-mep local-mep-id
logical-system
remote-mep remote-mep-id
```

Description

Clear the event for the configured action profile. If you have configured the action profile at the local MEP level, when you clear the event for the configured action profile, you can specify only the local MEP numeric identifier to clear the event. You need not specify the remote MEP numeric identifier. If you have configured the action profile at the remote MEP level, when you clear the event for the configured action profile, you must specify both the local MEP numeric identifier and remote MEP numeric identifier to clear the event.

Options

adjacency loss Clear the adjacency loss event.

all-defects	Clear the all-defects event.
cross-connect-ccm	Clear the cross-connect-ccm event.
erroneous-ccm	Clear the erroneous-ccm event.
local-mep *local-mep-id*	Specify the local MEP numeric identifier. If you have configured the action profile at the local MEP, you can clear the event for the configured action profile at the local MEP level. You need not specify the remote MEP numeric identifier.
logical-system	Name of the logical system
maintenance-association *ma-name*	Name of an existing CFM maintenance association.
maintenance-domain *md-name*	Name of an existing CFM maintenance domain.
receive-ais	Clear the receive-ais event.
remote-mep *remote-mep-id*	Specify the remote MEP numeric identifier. If you have configured the action profile at the remote MEP, you can clear the event for the configured action profile at the remote MEP level. You must specify both the local MEP numeric identifier and the remote MEP numeric identifier.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault-management event

```
user@host> clear oam ethernet connectivity-fault-management event maintenance-domain md5
maintenance-association ma5 local-mep 2 remote-mep 1
Action Profile Event Cleared.
```

Release Information

Command introduced in Junos OS Release 20.2.

clear oam ethernet connectivity-fault-management linktrace path-database

IN THIS SECTION

- Syntax | 3767
- Description | 3767
- Options | 3767
- Required Privilege Level | 3768
- Sample Output | 3768
- Release Information | 3768

Syntax

```
clear oam ethernet connectivity-fault-management linktrace path-database mac-address maintenance-
association ma-name maintenance-domain md-name
```

Description

Clear all the linktrace entries and the relevant path information from the database for a particular remote host on M320, MX Series, T320, and T640 routers.

Options

mac-address	Clear connectivity fault management path database information for the specified MAC address of the remote host.
maintenance-association *ma-name*	Clear connectivity fault management path database information for the specified maintenance association.

maintenance-domain ***md-name***	Clear connectivity fault management path database information for the specified maintenance domain.

Required Privilege Level

view

Sample Output

clear oam ethernet connectivity-fault- management linktrace path-database

```
user@host> clear oam ethernet connectivity-fault-management linktrace path-database maintenance-
domain md1 maintenance-association ma3 00058573e483
This command produces no output.
```

Release Information

Command introduced in Junos OS Release 9.0.

clear oam ethernet connectivity-fault-management loss-statistics

IN THIS SECTION

- Syntax | 3769
- Description | 3769
- Options | 3769
- Required Privilege Level | 3769
- Output Fields | 3769
- Sample Output | 3770
- Release Information | 3770

Syntax

```
clear oam ethernet connectivity-fault-management loss-statistics
<interface ethernet-interface-name>
<level md-level>
```

Description

For all routers that support IEEE 802.1ag OAM connectivity fault management (CFM), clear all loss statistics maintained by CFM for a given maintenance domain and maintenance association.

In addition, for Ethernet interfaces on MX Series routers , clear any ITU-T Y.1731 Ethernet frame loss measurement (ETH-LM) statistics.

By default, the command clears ETH-LM statistics for CFM maintenance association end points (MEPs) attached to any interface on the router.

Options

interface *ethernet-interface-name*	(Optional) Clear ETH-LM statistics and ETH-LM frame counts only for MEPs attached to the specified Ethernet physical interface.
level *md-level*	(Optional) Clear ETH-LM statistics and ETH-LM frame counts only for MEPs within CFM maintenance domains (MDs) of the specified level.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault- management loss-statistics

```
user@host> clear oam ethernet connectivity-fault-management loss-statistics
Cleared loss measurements statistics of all CFM sessions
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Ethernet Frame Loss Measurement

clear oam ethernet connectivity-fault-management path-database

IN THIS SECTION

- Syntax | 3770
- Description | 3771
- Options | 3771
- Required Privilege Level | 3771
- Sample Output | 3771
- Release Information | 3771

Syntax

```
clear oam ethernet connectivity-fault-management path-database maintenance-domain md-name
maintenance-association ma-name host <mac-addr>
```

Description

Clear the relevant path information from the database for the specified remote host.

Options

host	MAC address of remote host in xx:xx:xx:xx:xx:xx format.
maintenance-association	Name of the maintenance association.
maintenance-domain	Name of the maintenance domain.

Required Privilege Level

clear

Sample Output

clear oam ethernet connectivity-fault- management path-database

```
user@host> clear oam ethernet connectivity-fault-management path-database maintenance-domain
private maintenance-association private-ma 00:00:5E:00:53:AA
Path database entries cleared for the remote-host
```

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

show oam ethernet connectivity-fault-management path-database

clear oam ethernet connectivity-fault-management policer

IN THIS SECTION

- Syntax | 3772
- Description | 3772
- Options | 3772
- Required Privilege Level | 3773
- Output Fields | 3773
- clear oam ethernet connectivity- fault-management policer | 3773
- clear oam ethernet connectivity- fault-management policer maintenance-domain md-name maintenance-association ma-name | 3773
- Release Information | 3773

Syntax

```
clear oam ethernet connectivity-fault-management policer maintenance-domain md-name maintenance-
association ma-name
```

Description

On M7i and M10i with the Enhanced CFEB (CFEB-E), M320, M120, MX Series, T320, and T640 routers, clear connectivity-fault-management policer statistics.

Options

The following options are supported:

maintenance-domain md-name Name of an existing CFM maintenance domain. If this option is not specified, policer statistics are cleared for all maintenance associations for all maintenance domains.

maintenance-association ma-name Name of an existing CFM maintenance association. If this option is not specified, policer statistics are cleared for all maintenance associations for given

maintenance domain. This option cannot be specified without specifying maintenance-domain name.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

clear oam ethernet connectivity- fault-management policer

command-name

```
user@host> clear oam ethernet connectivity-fault-management
policer
Policer statistics cleared
```

clear oam ethernet connectivity- fault-management policer maintenance-domain md-name maintenance-association ma-name

command-name

```
user@host> clear oam ethernet connectivity-fault-management
policer maintenance-domain md5 maintenance-association ma5-1
Policer statistics cleared
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show oam ethernet connectivity-fault-management policer | 22660

clear oam ethernet connectivity-fault-management sla-iterator-statistics

IN THIS SECTION

- Syntax | 3774
- Description | 3774
- Options | 3774
- Required Privilege Level | 3775
- Output Fields | 3775
- Sample Output | 3775
- Release Information | 3775

Syntax

```
clear oam ethernet connectivity-fault-management sla-iterator-statistics
maintenance-association maintenance-association-name
maintenance-domain maintenance-domain-name
<local-mep local-mep-id>
<remote-mep remote-mep-id>
sla-iterator sla-iterator
```

Description

Clear Ethernet Operation, Administration, and Maintenance (OAM) service-level agreement (SLA) iterator statistics. For MX Series routers, clear the SLA iterator statistics and proactive Ethernet synthetic loss measurement (ETH-SLM) statistics.

Options

maintenance-association ***maintenance-association-name***	Name of the maintenance association.
maintenance-domain ***maintenance-domain-name***	Name of the maintenance domain.
local-mep ***local-mep-id***	(Optional) Identifier of the local MEP.

remote-mep ***remote-mep-id***	(Optional) Identifier of the remote MEP.
sla-iterator ***sla-iterator***	Name of the SLA iterator profile.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault -management sla-iterator- statistics

```
user@switch> clear oam ethernet connectivity-fault-management sla-iterator-statistics
maintenance-domain md1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator i1
Iterator statistics entries cleared
```

Release Information

Command introduced in Junos OS Release 11.4.

clear oam ethernet connectivity-fault-management statistics

IN THIS SECTION

- Syntax | 3776
- Description | 3776
- Options | 3776
- Required Privilege Level | 3776
- Output Fields | 3776
- Sample Output | 3777

Release Information | 3777

Syntax

```
clear oam ethernet connectivity-fault-management statistics
interface ethernet-interface-name
level md-level
```

Description

For all routers that support IEEE 802.1ag OAM connectivity-fault management (CFM), clear all statistics maintained by CFM.

In addition, for Ethernet interfaces on Dense Port Concentrators (DPCs) in MX Series routers only, also clear any ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM) statistics and ETH-DM frame counts.

By default, the command clears CFM statistics and ETH-DM statistics and frame counts for CFM maintenance association end points (MEPs) attached to any interface on the router.

Options

interface *ethernet-interface-name*	(Optional) Clear CFM statistics, ETH-DM statistics, and ETH-DM frame counts only for MEPs attached to the specified Ethernet physical interface.
level *md-level*	(Optional) Clear CFM statistics, ETH-DM statistics, and ETH-DM frame counts only for MEPs within CFM maintenance domains (MDs) of the specified level.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault- management statistics

```
user@host> clear oam ethernet connectivity-fault-management statistics
Cleared statistics of all CFM sessions
```

Release Information

Command introduced in Junos OS Release 8.4.

Support for ETH-DM statistics and frame counts added in Junos OS Release 9.5.

RELATED DOCUMENTATION

- show oam ethernet connectivity-fault-management delay-statistics | 22585
- show oam ethernet connectivity-fault-management interfaces | 22596
- show oam ethernet connectivity-fault-management mep-database | 22620
- show oam ethernet connectivity-fault-management mep-statistics | 22637

clear oam ethernet connectivity-fault-management synthetic-loss-measurement

IN THIS SECTION

- Syntax | 3778
- Description | 3778
- Options | 3778
- Required Privilege Level | 3778
- Output Fields | 3778
- Sample Output | 3779
- Release Information | 3779

Syntax

```
clear oam ethernet connectivity-fault-management synthetic-loss-measurement
maintenance-domain md-name
maintenance-association ma-name
<local-mep local-mep-id>
<remote-mep remote-mep-id>
```

Description

(MX Series routers)—For all routers that support IEEE 802.1ag OAM connectivity fault management (CFM), clear the existing on-demand Ethernet synthetic loss measurement (ETH-SLM) statistics and restart counting the ETH-SLM frame counts and statistics.

Options

maintenance-domain *md-name*	Name of an existing CFM maintenance domain.
maintenance-association *ma-name*	Name of an existing CFM maintenance association.
local-mep *local-mep-id*	(Optional) Clear connectivity fault management information for the specified local MEP only.
remote-mep *remote-mep-id*	(Optional) Clear connectivity fault management information for the specified remote MEP only.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault- management synthetic-loss- measurement

```
user@host> clear oam ethernet connectivity-fault-management synthetic-loss-measurement
maintenance-domain md5 maintenance-association ma5 local-mep 100 remote-mep 102
Synthetic loss measurement restarted.
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

monitor ethernet synthetic-loss-measurement | 10850

show oam ethernet connectivity-fault-management synthetic-loss-statistics | 22677

clear oam ethernet link-fault-management state

IN THIS SECTION

- Syntax | 3780
- Description | 3780
- Options | 3780
- Required Privilege Level | 3780
- Output Fields | 3780
- Sample Output | 3780
- Release Information | 3780

Syntax

```
clear oam ethernet link-fault-management state <interface-name>
```

Description

On all M Series, MX Series, ACX series, PTX Series, T320, and T640 routers, clear link fault management state information, restart the link discovery process, and reset OAM loopback state (if set previously) on Ethernet interfaces.

Options

none	Clear OAM link fault management state information, restart the link discovery process, and reset OAM loopback state (if set previously) on all Ethernet interfaces.
interface-name	(Optional) Clear OAM link fault management state information, restart the link discovery process, and reset OAM loopback state (if set previously) on the specified Ethernet interface only.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet link-fault-management state

```
user@host> clear oam ethernet link-fault-management state ge-0/3/3
Cleared link-fault-management state for interface ge-0/3/3
```

Release Information

Command introduced in Junos OS Release 8.4.

clear oam ethernet link-fault-management statistics

IN THIS SECTION

- Syntax | 3781
- Description | 3781
- Options | 3781
- Required Privilege Level | 3781
- Output Fields | 3782
- Sample Output | 3782
- Release Information | 3782

Syntax

```
clear oam ethernet link-fault-management <interface-name>
```

Description

On M320, M120, MX Series, PTX Series, T320, and T640 routers, clear Operation, Administration, and Management (OAM) link fault management statistics or state information from Ethernet interfaces.

Options

none	Clear OAM link fault management statistics from all Ethernet interfaces.
interface-name	(Optional) Clear OAM link fault management statistics from the specified Ethernet interface only.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet link-fault-management statistics

```
user@host> clear oam ethernet link-fault-management statistics
Cleared link-fault-management statistics for all interfaces
```

Release Information

Command introduced in Junos OS Release 8.2.

clear-on-abort (DHCP Local Server)

IN THIS SECTION

- Syntax | 3782
- Hierarchy Level | 3783
- Description | 3783
- Default | 3783
- Required Privilege Level | 3784
- Release Information | 3784

Syntax

```
clear-on-abort;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure],
[edit system services dhcp-local-server reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server group group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Delete all DHCP clients or only the DHCP clients serviced by the specified group of interfaces when reconfiguration fails; that is, when the maximum number of retry attempts have been made without success. A group configuration takes precedence over a DHCP local server configuration.

Default

Restores the original client configuration when reconfiguration fails.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview

Configuring Deletion of the Client When Dynamic Reconfiguration Fails

clear (ospf | ospf3) database

IN THIS SECTION

- Syntax | 3784
- Syntax (EX Series Switch and QFX Series) | 3785
- Description | 3786
- Options | 3786
- Required Privilege Level | 3787
- Output Fields | 3787
- Sample Output | 3787
- Release Information | 3787

Syntax

```
clear (ospf | ospf3) database
<all>
```

```
<advertising-router (router-id | self)>
 <area area-id>
<asbrsummary>
<external>
<instance instance-name>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<logical-system (all | logical-system-name)>
<lsa-id lsa-id>
<netsummary>
<network>
<nssa>
<opaque-area>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>
```

Syntax (EX Series Switch and QFX Series)

```
clear (ospf | ospf3) database
<all>
<advertising-router (router-id | self)>
 <area area-id>
<asbrsummary>
<external>
<instance instance-name>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<lsa-id lsa-id>
<netsummary>
<network>
<nssa>
<opaque-area>
<router>
```

Description

With the primary Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and sync the new database with the primary Routing Engine.

Options

all	Delete all LSAs other than the system's own LSAs, which are regenerated. To resynchronize the database, the system destroys all adjacent neighbors that are in the state `EXSTART` or higher. The neighbors are then reacquired and the databases are synchronized.
advertising-router (*router-id* | self)	(Optional) Discard entries for the LSA entries advertised by the specified routing device or by this routing device.
area *area-id*	(Optional) Discard entries for the LSAs in the specified area.
asbrsummary	(Optional) Discard summary AS boundary router LSA entries.
external	(Optional) Discard external LSAs.
instance *instance-name*	(Optional) Delete or discard entries for the specified routing instance only.
inter-area-prefix	(OSPFv3 only) (Optional) Discard interarea prefix LSAs.
inter-area-router	(OSPFv3 only) (Optional) Discard interarea router LSAs.
intra-area-prefix	(OSPFv3 only) (Optional) Discard intra-area prefix LSAs.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
link-local	(Optional) Delete link-local LSAs.
lsa-id *lsa-id*	(Optional) Discard the LSA entries with the specified LSA identifier.
netsummary	(Optional) Discard summary network LSAs.
network	(Optional) Discard network LSAs.
nssa	(Optional) Discard not-so-stubby area (NSSA) LSAs.
opaque-area	(Optional) Discard opaque area-scope LSAs.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)	(OSPFv3 only) (Optional) Delete the entries for the specified OSPFv3 realm, or address family. Use the `realm` option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
router	(Optional) Discard router LSAs.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf database all

```
user@host> clear ospf database all
```

Release Information

Command introduced before Junos OS Release 7.4.

`advertising-router` *`router-id`*, `netsummary`, `network`, `nssa`, `opaque-area`, and `router` options added in Junos OS Release 8.3.

`area` *`area-id`* option added in Junos OS Release 8.3.

`realm` option added in Junos OS Release 9.2.

`advertising-router` (*`router-id`* | `self`) option added in Junos OS Release 9.5.

`advertising-router` (*`router-id`* | `self`) option introduced in Junos OS Release 9.5 for EX Series switches.

RELATED DOCUMENTATION

show ospf database | 22780

show ospf3 database | 22869

clear (ospf | ospf3) database-protection

IN THIS SECTION

- Syntax | 3788
- Description | 3788
- Options | 3788
- Required Privilege Level | 3788
- Output Fields | 3788
- Sample Output | 3789
- Release Information | 3789

Syntax

```
clear (ospf | ospf3) database-protection
<instance instance-name>
```

Description

Clear the Open Shortest Path First (OSPF) link-state database from its isolated state. Reset the ignore count, ignore timer, and reset timer, and resume normal operations.

Options

instance *instance-name*	(Optional) Clear the OSPF link-state database for the specified routing instance only.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear ospf database-protection

```
user@host> clear ospf database-protection
```

Release Information

Command introduced in Junos OS Release 10.2.

clear (ospf | ospf3) io-statistics

IN THIS SECTION

- Syntax | 3789
- Syntax (EX Series Switch and QFX Series) | 3790
- Description | 3790
- Options | 3790
- Required Privilege Level | 3790
- Output Fields | 3790
- Sample Output | 3790
- Release Information | 3790

Syntax

```
clear (ospf | osfp3) io-statistics
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and QFX Series)

```
clear (ospf | osfp3) io-statistics
```

Description

Clear Open Shortest Path First (OSPF) input and output statistics.

Options

none	Clear OSPF input and output statistics.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf io-statistics

```
user@host> clear ospf io-statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

clear (ospf | ospf3) neighbor

IN THIS SECTION

- Syntax | 3791
- Syntax (EX Series Switch and QFX Series) | 3791
- Description | 3792
- Options | 3792
- Required Privilege Level | 3792
- Output Fields | 3792
- Sample Output | 3792
- Release Information | 3793

Syntax

```
clear (ospf | ospf3) neighbor
<all>
<area area-id>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
<neighbor>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
```

Syntax (EX Series Switch and QFX Series)

```
clear (ospf | ospf3) neighbor
<all>
<area area-id>
<instance instance-name>
<interface interface-name>
<neighbor>
```

Description

Tear down Open Shortest Path First (OSPF) neighbor connections.

Options

all	Tear down OSPF connections with all neighbors for all routing instances.
area *area-id*	(Optional) Tear down neighbor connections for the specified area only.
instance *instance-name*	(Optional) Tear down neighbor connections for the specified routing instance only.
interface *interface-name*	(Optional) Tear down neighbor connections for the specified interface only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
neighbor	(Optional) Clear the state of the specified neighbor only.
realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)	(Optional) (OSPFv3 only) Clear the state of the specified OSPFv3 realm, or address family. Use the `realm` option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf neighbor all

```
user@host> clear ospf neighbor all
```

Release Information

Command introduced before Junos OS Release 7.4.

`realm` option introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

show (ospf | ospf3) neighbor | 22835

clear (ospf | ospf3) overload

IN THIS SECTION

- Syntax | 3793
- Syntax (EX Series Switches) | 3794
- Description | 3794
- Options | 3794
- Required Privilege Level | 3794
- Output Fields | 3794
- Sample Output | 3794
- Release Information | 3794

Syntax

```
clear (ospf | ospf3) overload
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
clear (ospf | ospf3) overload
<instance instance-name>
```

Description

Clear the Open Shortest Path First (OSPF) overload bit and rebuild link-state advertisements (LSAs).

Options

none	Clear the overload bit and rebuild LSAs for all routing instances.
instance *instance-name*	(Optional) Clear the overload bit and rebuild LSAs for the specified routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf overload

```
user@host> clear ospf overload
```

Release Information

Command introduced before Junos OS Release 7.4.

clear (ospf | ospf3) statistics

IN THIS SECTION

- Syntax | 3795
- Syntax (EX Series Switch and QFX Series) | 3795
- Description | 3795
- Options | 3795
- Required Privilege Level | 3796
- Output Fields | 3796
- Sample Output | 3796
- Release Information | 3797

Syntax

```
clear (ospf | osfp3) statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
```

Syntax (EX Series Switch and QFX Series)

```
clear (ospf | osfp3) statistics
<instance instance-name>
```

Description

Clear Open Shortest Path First (OSPF) statistics.

Options

none	Clear OSPF statistics.

instance *instance-name*	(Optional) Clear statistics for the specified routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)	(Optional) (OSPFv3 only) Clear statistics for the specified OSPFv3 realm, or address family. Use the `realm` option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

clear

Output Fields

See "show (ospf | ospf3) statistics" on page 22859 for an explanation of output fields.

Sample Output

clear ospf statistics

The following sample output displays OSPF statistics before and after the `clear ospf statistics` command is entered:

```
user@host> show ospf statistics

Packet type             Total                  Last 5 seconds
                    Sent      Received        Sent      Received
   Hello            3254          2268           3             1
     DbD              41            46           0             0
   LSReq               8             7           0             0
LSUpdate             212           154           0             0
   LSAck              65            98           0             0

DBDs retransmitted     :                    3, last 5 seconds :          0
LSAs flooded           :                   12, last 5 seconds :          0
LSAs flooded high-prio :                    0, last 5 seconds :          0
LSAs retransmitted     :                    0, last 5 seconds :          0
LSAs transmitted to nbr:                    3, last 5 seconds :          0
LSAs requested         :                    5, last 5 seconds :          0
LSAs acknowledged      :                   19, last 5 seconds :          0
```

```
Flood queue depth      :               0
Total rexmit entries   :               0
db summaries           :               0
lsreq entries          :               0

Receive errors:
  626 subnet mismatches


user@host> clear ospf statistics


user@host> show ospf statistics
Packet type             Total                  Last 5 seconds
                   Sent      Received        Sent      Received
   Hello              3             1           3             1
     DbD              0             0           0             0
   LSReq              0             0           0             0
LSUpdate              0             0           0             0
   LSAck              0             0           0             0

DBDs retransmitted     :                    0, last 5 seconds :          0
LSAs flooded           :                    0, last 5 seconds :          0
LSAs flooded high-prio :                    0, last 5 seconds :          0
LSAs retransmitted     :                    0, last 5 seconds :          0
LSAs transmitted to nbr:                    0, last 5 seconds :          0
LSAs requested         :                    0, last 5 seconds :          0
LSAs acknowledged      :                    0, last 5 seconds :          0

Flood queue depth      :               0
Total rexmit entries   :               0
db summaries           :               0
lsreq entries          :               0

Receive errors:
  None
```

Release Information

Command introduced before Junos OS Release 7.4.

`realm` option introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

show (ospf | ospf3) statistics | 22859

clear ovsdb commit failures

IN THIS SECTION

- Syntax | 3798
- Description | 3798
- Options | 3799
- Required Privilege Level | 3799
- Sample Output | 3800
- Release Information | 3800

Syntax

```
clear ovsdb commit failures
<transaction-id>
```

Description

Remove a transaction from a queue maintained by a Juniper Networks switch that supports the *Open vSwitch Database (OVSDB)* management protocol and Virtual Extensible LANs (VXLANs). The transaction includes OVSDB-managed VXLANs and associated logical interfaces that the Juniper Networks switch dynamically configured and tried to commit but was unable to because of an issue with one or more of the configurations. In addition to removing the transaction, entering the `clear ovsdb commit failures` command causes the Juniper Networks switch to automatically retry committing all configurations in the transaction.

If there is an issue with one or more of the configurations in a transaction, this causes all configurations in the transaction, even the ones that are correctly configured, to remain uncommitted and in the queue until you troubleshoot and resolve the configuration issue(s).

You can display an erroneous transaction by entering the `show ovsdb commit failures` command. In the output that appears, you must determine which configuration(s) are erroneous and therefore prevent the Juniper Networks switch from committing the configurations in the transaction.

Issues that can cause commitment errors include but are not limited to the detection of the same VXLAN name or VXLAN network identifier (VNI) in a dynamically configured VXLAN and in a VXLAN that was previously configured using the Junos OS CLI.

To monitor for issues with dynamically configured OVSDB-managed VXLANs and their associated interfaces, we recommend checking for system log messages and traceoptions files for OVSDB.

After resolving the error(s), enter the `clear ovsdb commit failures` command to remove the transaction from the queue and retry committing all configurations in the transaction.

> **NOTE**: While an erroneous transaction exists in the queue, the Juniper Networks switch cannot commit the dynamic configurations of additional VXLANs and their associated logical interfaces. The commitment of these VXLANs and logical interfaces remain in a pending state until all VXLAN and logical interface configurations in the erroneous transaction are resolved and successfully committed.

Options

none Remove the transaction that currently appears in the `show ovsdb commit failures` command output, and retry committing all configurations in the transaction.

transaction-id Remove the transaction with the specified numerical ID, and retry committing the configurations in the transaction.

Required Privilege Level

clear

Sample Output

clear ovsdb commit failures

```
user@host> clear ovsdb commit failures
```

clear ovsdb commit failures (Specific Transaction)

```
user@host> clear ovsdb commit failures 1
```

Release Information

Command introduced in Junos OS Release 14.1X53-D26.

RELATED DOCUMENTATION

show ovsdb commit failures | 22881

clear passive-monitoring statistics

IN THIS SECTION

- Syntax | 3801
- Description | 3801
- Options | 3801
- Required Privilege Level | 3801
- Output Fields | 3801
- Sample Output | 3801
- Release Information | 3801

Syntax

```
clear passive-monitoring statistics (all | interface interface-name)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Clear statistics for one passive monitoring interface or for all passive monitoring interfaces.

Options

all	Clear statistics for all configured passive monitoring interfaces.
interface *interface-name*	Clear statistics for the specified passive monitoring interface (**mo-***fpc*/*pic*/*port*).

Required Privilege Level

network

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear passive-monitoring statistics

```
user@host> clear passive-monitoring statistics interface mo-5/0/0
```

Release Information

Command introduced in Junos OS Release 7.6.

clear path-computation-client statistics

IN THIS SECTION

- Syntax | 3802
- Description | 3802
- Options | 3802
- Required Privilege Level | 3802
- Output Fields | 3802
- Sample Output | 3803
- Release Information | 3803

Syntax

```
clear path-computation-client statistics
<pce-id | all>
```

Description

Clear Path Computation Element (PCE) statistics.

Options

pce-id (Optional) Clear statistics of the specified PCE.

all (Optional) Clear statistics of all available PCEs configured on the path computation client (PCC).

Required Privilege Level

clear

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear path-computation-client statistics pce-id

```
user@host> clear path-computation-client statistics pce1
```

clear path-computation-client statistics all

```
user@host> clear path-computation-client statistics all
```

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

show path-computation-client statistics | 22936

clear performance-monitoring mpls lsp

IN THIS SECTION

- Syntax | 3804
- Description | 3804
- Options | 3804
- Required Privilege Level | 3804
- Output Fields | 3804
- Sample Output | 3804
- Release Information | 3804

Syntax

```
clear performance-monitoring mpls lsp
<name lsp-name>
```

Description

Restart the performance monitoring statistics.

Options

none	Reset and restart all performance monitoring for all LSPs.
name *lsp-name*	(Optional) Reset and restart performance monitoring for the specified LSP.

Required Privilege Level

clear

Output Fields

When you enter this command, performance monitoring is restarted.

Sample Output

clear performance-monitoring mpls lsp

```
user@host> clear performance-monitoring mpls lsp
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

performance-monitoring (Protocols MPLS) | 12434

show performance-monitoring mpls lsp | 22947

clear pfe sflow egress-mcast-stats

IN THIS SECTION

- Syntax | 3805
- Description | 3805
- Required Privilege Level | 3805
- Sample Output | 3806
- Release Information | 3806

Syntax

```
clear pfe sflow egress-mcast-stats
```

Description

Clears sFlow egress multicast related statistics like dropped samples count and replicated samples count from the Packet Forwarding Engine.

Required Privilege Level

view

Sample Output

clear pfe sflow egress-mcast-stats

```
user@host> clear pfe sflow egress-mcast-stats
```

```
Slot 0




Sflow VxLAN Stats cleared
```

Release Information

Command introduced in Junos OS Release 21.2R1.

clear pim join

IN THIS SECTION

- Syntax | 3807
- Syntax (EX Series Switch and the QFX Series) | 3807
- Description | 3807
- Options | 3807
- Additional Information | 3808
- Required Privilege Level | 3808
- Output Fields | 3808
- Sample Output | 3808
- Release Information | 3809

Syntax

```
clear pim join
<all>
<group-address>
<bidirectional | dense | sparse>
<exact>
<inet | inet6>
<instance instance-name>
<logical-system (all | logical-system-name)>
<rp ip-address/prefix | source ip-address/prefix>
<sg | star-g>
```

Syntax (EX Series Switch and the QFX Series)

```
clear pim join
<all>
<group-address>
<dense | sparse>
<exact>
<inet | inet6>
<instance instance-name>
<rp ip-address/prefix | source ip-address/prefix>
<sg | star-g>
```

Description

Clear the Protocol Independent Multicast (PIM) join and prune states.

Options

all	To clear PIM join and prune states for all groups and family addresses in the master instance, you must specify “all”..
group-address	(Optional) Clear the PIM join and prune states for a group address.
bidirectional | dense | sparse	(Optional) Clear PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.

exact	(Optional) Clear only the group that exactly matches the specified group address.
inet | inet6	(Optional) Clear the PIM entries for IPv4 or IPv6 family addresses, respectively.
instance *instance-name*	(Optional) Clear the entries for a specific PIM-enabled routing instance.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system.
rp *ip-address/prefix* **| source** *ip-address/prefix*	(Optional) Clear the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.
sg | star-g	(Optional) Clear PIM (S,G) or (*,G) entries.

Additional Information

The `clear pim join` command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pim join all

```
user@host> clear pim join all
Cleared 8 Join/Prune states
```

clear pim join inet6 all

```
user@host> clear pim join inet6 all
Cleared 4 Join/Prune states
```

clear pim join inet6 star-g all

```
user@host> clear pim join inet6 star-g all
Cleared 1 Join/Prune states
```

Release Information

Command introduced before Junos OS Release 7.4.

inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.

Multiple new filter options introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

show pim join | 23070

clear pim join-distribution

IN THIS SECTION

- Syntax | 3810
- Description | 3810
- Options | 3810
- Additional Information | 3811
- Required Privilege Level | 3811
- Output Fields | 3811
- Sample Output | 3811

Release Information | 3811

Syntax

```
clear pim join-distribution
<all>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Clear the PIM `join-redistribute` states.

Use the `show pim source` command to find out if there are multiple paths available for a source (for example, an RP).

When you include the `join-load-balance` statement in the configuration, the PIM join states are distributed evenly on available equal-cost multipath links. When an upstream neighbor link fails, Junos OS redistributes the PIM join states to the remaining links. However, when new links are added or the failed link is restored, the existing PIM joins are not redistributed to the new link. New flows will be distributed to the new links. However, in a network without new joins and prunes, the new link is not used for multicast traffic. The `clear pim join-distribution` command redistributes the existing flows to the new upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you run the `clear pim join-distribution` command during a maintenance window.

Options

all	(Optional) Clear the PIM `join-redistribute` states for all groups and family addresses in the master instance.
none	Automatically clear all PIM join/prune states.
instance *instance-name*	(Optional) Redistribute the join states for a specific PIM-enabled routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The `clear pim join-distribution` command cannot be used to redistribute the PIM join states on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show pim join` command before and after distributing the join state to verify the operation.

Sample Output

clear pim join-distribution all

```
user@host> clear pim join-distribution all
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- show pim neighbors | 23106
- show pim join | 23070
- join-load-balance | 9431

clear pim register

IN THIS SECTION

- Syntax | 3812
- Syntax (EX Series Switch and the QFX Series) | 3812
- Syntax (PTX Series) | 3813
- Description | 3813
- Options | 3813
- Additional Information | 3813
- Required Privilege Level | 3813
- Output Fields | 3813
- Sample Output | 3814
- Release Information | 3814

Syntax

```
clear pim register
<all>
<inet | inet6>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
clear pim register
<inet | inet6>
<instance instance-name>
<interface interface-name>
```

Syntax (PTX Series)

```
clear pim register
<inet | inet6>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Clear Protocol Independent Multicast (PIM) register message counters.

Options

all	Required to clear the PIM register message counters for all groups and family addresses in the master instance.
inet | inet6	(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.
instance *instance-name*	(Optional) Clear register message counters for a specific PIM-enabled routing instance.
interface *interface-name*	(Optional) Clear PIM register message counters for a specific interface.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The `clear pim register` command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pim register all

```
user@host> clear pim register all
```

Release Information

Command introduced in Junos OS Release 7.6.

inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.

RELATED DOCUMENTATION

show pim statistics | 23152

clear pim snooping join

IN THIS SECTION

- Syntax | 3814
- Description | 3815
- Options | 3815
- Required Privilege Level | 3815
- Output Fields | 3815
- Sample Output | 3815
- Release Information | 3816

Syntax

```
clear pim snooping join
<instance instance-name>
```

```
<logical-system logical-system-name>
<vlan-id vlan-id>
```

Description

Clear information about Protocol Independent Multicast (PIM) snooping joins.

Options

none	Display detailed information.
instance ***instance-name***	(Optional) Clear PIM snooping join information for the specified routing instance.
logical-system ***logical-system-name***	(Optional) Delete the IGMP snooping statistics for a given logical system or for all logical systems.
vlan-id ***vlan-identifier***	(Optional) Clear PIM snooping join information for the specified VLAN.

Required Privilege Level

view

Output Fields

See "show pim snooping join" on page 23128 for an explanation of the output fields.

Sample Output

clear pim snooping join

The following sample output displays information about PIM snooping joins before and after the `clear pim snooping join` command is entered:

```
user@host> show pim snooping join extensive
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20

Group: 198.51.100.2
```

```
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.5, port: ge-1/3/7.20
Downstream port: ge-1/3/1.20
Downstream neighbors:
192.0.2.2 State: Join Flags: SRW Timeout: 185

Group: 198.51.100.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.4, port: ge-1/3/5.20
Downstream port: ge-1/3/3.20
Downstream neighbors:
192.0.2.3 State: Join Flags: SRW Timeout: 175
user@host> clear pim snooping join
Clearing the Join/Prune state for 203.0.113.0/24
Clearing the Join/Prune state for 203.0.113.0/24
user@host> show pim snooping join extensive
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

PIM Snooping for VPLS

clear pim snooping statistics

IN THIS SECTION

- Syntax | 3817
- Description | 3817
- Options | 3817
- Required Privilege Level | 3818
- Output Fields | 3818
- Sample Output | 3818
- Release Information | 3819

Syntax

```
clear pim snooping statistics
<instance instance-name>
<interface interface-name>
<logical-system logical-system-name>
<vlan-id vlan-id>
```

Description

Clear Protocol Independent Multicast (PIM) snooping statistics.

Options

none	Clear PIM snooping statistics for all family addresses, instances, and interfaces.
instance *instance-name*	(Optional) Clear statistics for a specific PIM-snooping-enabled routing instance.
interface *interface-name*	(Optional) Clear PIM snooping statistics for a specific interface.

logical-system *logical-system-name*	(Optional) Delete the IGMP snooping statistics for a given logical system or for all logical systems.
vlan-id *vlan-identifier*	(Optional) Clear PIM snooping statistics information for the specified VLAN.

Required Privilege Level

clear

Output Fields

See "show pim snooping statistics" on page 23141 for an explanation of the output fields.

Sample Output

clear pim snooping statistics

The following sample output displays PIM snooping statistics before and after the `clear pim snooping statistics` command is entered:

```
user@host> show pim snooping statistics
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 660
Rx J/P messages -- seen 0
Rx J/P messages -- received 660
Rx Hello messages 1396
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0

Learning-Domain: vlan-id 20
user@host> clear pim snooping statistics
```

```
user@host> show pim snooping statistics
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 0
Rx J/P messages -- seen 0
Rx J/P messages -- received 0
Rx Hello messages 0
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0

Learning-Domain: vlan-id 20
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

- PIM Snooping for VPLS

clear pim statistics

IN THIS SECTION

- Syntax | 3820
- Syntax (EX Series Switch and the QFX Series) | 3820
- Description | 3820
- Options | 3820

- Additional Information | 3821
- Required Privilege Level | 3821
- Output Fields | 3821
- Sample Output | 3821
- Release Information | 3823

Syntax

```
clear pim statistics
<inet | inet6>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
clear pim statistics
<inet | inet6>
<instance instance-name>
<interface interface-name>
```

Description

Clear Protocol Independent Multicast (PIM) statistics.

Options

none	Clear PIM statistics for all family addresses, instances, and interfaces.
inet | inet6	(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.

instance ***instance-name***	(Optional) Clear statistics for a specific PIM-enabled routing instance.
interface ***interface-name***	(Optional) Clear PIM statistics for a specific interface.
logical-system (all | *logical-system-name)*	(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The `clear pim statistics` command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Output Fields

See "show pim statistics" on page 23152 for an explanation of output fields.

Sample Output

clear pim statistics

The following sample output displays PIM statistics before and after the `clear pim statistics` command is entered:

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
Assert                        0          0          0
Graft                         0          0          0
Graft Ack                     0          0          0
Candidate RP                  0          0          0
```

```
V1 Query                    2111       4222          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune              14200      13115          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
PIM statistics summary for all interfaces:
Unknown type                              0
V1 Unknown type                           0
Unknown Version                           0
Neighbor unknown                          0
Bad Length                                0
Bad Checksum                              0
Bad Receive If                            0
Rx Intf disabled                       2007
Rx V1 Require V2                          0
Rx Register not RP                        0
RP Filtered Source                        0
Unknown Reg Stop                          0
Rx Join/Prune no state                 1040
Rx Graft/Graft Ack no state               0
...


user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                          0          0          0
Register                       0          0          0
Register Stop                  0          0          0
Join Prune                     0          0          0
Bootstrap                      0          0          0
Assert                         0          0          0
Graft                          0          0          0
Graft Ack                      0          0          0
Candidate RP                   0          0          0
V1 Query                       1          0          0
V1 Register                    0          0          0
...
```

Release Information

Command introduced before Junos OS Release 7.4.

inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.

RELATED DOCUMENTATION

show pim statistics | 23152

clear policy statistics

IN THIS SECTION

- Syntax | 3823
- Description | 3823
- Options | 3824
- Required Privilege Level | 3824
- Output Fields | 3824
- Sample Output | 3824
- Release Information | 3824

Syntax

```
clear policy statistics
<policy -name>
```

Description

Clear policy statistics.

Options

none	Clear statistics for all policies.
policy-name	(Optional) Clear statistics for the specified policy only.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear policy statistics

```
user@host> clear policy statistics
```

Release Information

Command introduced in Junos OS Release 15.1F6.

RELATED DOCUMENTATION

show policy | 23199

test policy | 29209

clear pppoe lockout

IN THIS SECTION

Syntax | 3825

- Description | 3825
- Options | 3825
- Required Privilege Level | 3826
- Sample Output | 3826
- Release Information | 3826

Syntax

```
clear pppoe lockout
<aci circuit-id | mac-address mac-address >
<underlying-interfaces underlying-interface-name>
```

Description

Clear the lockout condition for the PPPoE client associated with the specified media access control (MAC) source address or agent circuit identifier (ACI) value.

Options

none
Clear the lockout condition for the PPPoE clients associated with all MAC source addresses on all PPPoE underlying interfaces.

aci *circuit-id*
(Optional) Clear the lockout condition for the PPPoE client associated with the specified ACI value. To clear the lockout condition by a specified ACI value, you must specify the `filter aci` option in the `short-cycle-protection` statement when you configure PPPoE subscriber session lockout. If the `filter aci` option is missing from the `short-cycle-protection` statement , no PPPoE client sessions are cleared using the ACI filter. The `aci` option and the `mac-address` option are mutually exclusive.

mac-address *mac-address*
(Optional) Clear the lockout condition for the PPPoE client associated with the specified MAC source address. The `mac-address` option and the `aci` option are mutually exclusive.

underlying-interfaces *underlying-interface-name*
(Optional) Clear the lockout condition for all PPPoE clients associated with the specified PPPoE underlying interface.

Required Privilege Level

clear

Sample Output

clear pppoe lockout (All MAC Source Addresses on All Underlying Interfaces)

```
user@host> clear pppoe lockout
```

clear pppoe lockout mac-address (Specified MAC Source Address)

```
user@host> clear pppoe lockout mac-address 00:00:5e:00:53:30
```

clear pppoe lockout mac-address underlying-interfaces (Specified MAC Source Address on Specified Underlying Interface)

```
user@host> clear pppoe lockout mac-address 00:00:5e:00:53:30 underlying-interfaces ge-1/0/0.101
```

clear pppoe lockout underlying-interfaces (All MAC Source Addresses on Specified Underlying Interface)

```
user@host> clear pppoe lockout underlying-interfaces ge-1/0/0.101
```

clear pppoe lockout underlying-interfaces aci (ACI on Specified Underlying Interface)

```
user@host> clear pppoe lockout underlying-interfaces demux0.214 aci “Relay-identifier atm
3/0:100\.*”
```

Release Information

Command introduced in Junos OS Release 11.4.

`aci` option introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Clearing Lockout of PPPoE Subscriber Sessions

Configuring Lockout of PPPoE Subscriber Sessions

clear pppoe lockout atm-identifier

IN THIS SECTION

- Syntax | 3827
- Description | 3827
- Options | 3828
- Required Privilege Level | 3828
- Sample Output | 3828
- Release Information | 3829

Syntax

```
clear pppoe lockout atm-identifier device-name device-name vpi vpi-identifier vci vci-identifier
<aci circuit-id | mac-address mac-address >
```

Description

Clear the lockout condition for the PPPoE client associated with the specified ATM encapsulation type and, optionally, media access control (MAC) source address or agent circuit identifier (ACI) value. Because the lockout condition persists even in the absence of an underlying interface or after automatic removal of the VLAN or VLAN demux interface, using the `clear pppoe lockout atm-identifier` command enables you to clear the lockout condition for PPPoE clients by specifying ATM identifying characteristics instead of the ATM interface name.

The following characteristics comprise the ATM encapsulation type identifier:

- Device name (physical interface or aggregated Ethernet bundle)
- Virtual path identifier (VPI)

- Virtual circuit identifier (VCI)

Options

circuit-id (Optional) ACI value associated with the PPPoE client for which you want to clear lockout. To clear the lockout condition by a specified ACI value, you must specify the `filter aci` option in the `short-cycle-protection` statement when you configure PPPoE subscriber session lockout. If the `filter aci` option is missing from the `short-cycle-protection` statement, no PPPoE client sessions are cleared using the ACI filter. The `aci` option and the `mac-address` option are mutually exclusive.

device-name Name of the ATM physical interface or aggregated Ethernet bundle associated with the PPPoE client for which you want to clear lockout.

mac-address (Optional) MAC address value associated with the PPPoE client for which you want to clear lockout. The `mac-address` option and the `aci` option are mutually exclusive.

vci-identifier ATM VCI value associated with the PPPoE client for which you want to clear lockout.

- **Range:** 0 through 65535

vpi-identifier ATM VPI value associated with the PPPoE client for which you want to clear lockout.

- **Range:** 0 through 255

Required Privilege Level

clear

Sample Output

clear pppoe lockout atm-identifier device-name vpi vci (PPPoE Client with Specified VPI and VCI on ATM Physical Interface)

```
user@host> clear pppoe lockout atm-identifier device-name at-1/0/0 vpi 10 vci 40
```

clear pppoe lockout atm-identifier device-name vpi vci aci (PPPoE Client with Specified VPI and VCI on Aggregated Ethernet Bundle Where ACI Matches Regular Expression)

```
user@host> clear pppoe lockout atm-identifier device-name ae1 vpi 1 vci 30 aci ""Relay-
identifier atm 1/0:100\.*"
```

clear pppoe lockout atm-identifier device-name vpi vci mac-address (PPPoE Client with Specified VPI, VCI, and MAC Address on ATM Logical Interface)

```
user@host> clear pppoe lockout atm-identifier device-name at-1/1/0.20 vpi 1 vci 20 mac-address
00:00:5e:00:53:30
```

Release Information

Command introduced in Junos OS Release 15.2.

RELATED DOCUMENTATION

Clearing Lockout of PPPoE Subscriber Sessions

Configuring Lockout of PPPoE Subscriber Sessions

clear pppoe lockout vlan-identifier

IN THIS SECTION

- Syntax | 3830
- Description | 3830
- Options | 3830
- Required Privilege Level | 3831
- Sample Output | 3831
- Release Information | 3832

Syntax

```
clear pppoe lockout vlan-identifier device-name device-name
<aci circuit-id | mac-address mac-address >
<svlan-id svlan-identifier>
<vlan-id vlan-identifier>
```

Description

Clear the lockout condition for the PPPoE client associated with the specified VLAN encapsulation type and, optionally, media access control (MAC) source address and agent circuit identifier (ACI) value. Because the lockout condition persists even in the absence of an underlying interface or after automatic removal of the VLAN or VLAN demux interface, using the `clear pppoe lockout vlan-identifier` command enables you to clear the lockout condition for PPPoE clients by specifying VLAN identifying characteristics rather than by specifying the underlying interface name.

The following characteristics comprise the VLAN encapsulation type identifier:

- Device name (physical interface or aggregated Ethernet bundle)
- Stacked VLAN (S-VLAN) ID (also known as the *outer tag*)
- VLAN ID (also known as the *inner tag*)

You can configure PPPoE subscriber session lockout, also known as PPPoE short-cycle protection, for VLAN, VLAN demux, and PPPoE-over-ATM dynamic subscriber interfaces.

Options

circuit-id (Optional) ACI value associated with the PPPoE client for which you want to clear lockout. To clear the lockout condition by a specified ACI value, you must specify the `filter aci` option in the `short-cycle-protection` statement when you configure PPPoE subscriber session lockout. If the `filter aci` option is missing from the `short-cycle-protection` statement, no PPPoE client sessions are cleared using the ACI filter. The `aci` option and the `mac-address` option are mutually exclusive.

device-name Name of the Ethernet physical interface or aggregated Ethernet bundle associated with the PPPoE client for which you want to clear lockout.

mac-address (Optional) MAC address value associated with the PPPoE client for which you want to clear lockout. The `mac-address` option and the `aci` option are mutually exclusive.

svlan-identifier (Optional) A valid S-VLAN identifier associated with the PPPoE client for which you want to clear lockout.

- **Range:** 1 through 4094

vlan-identifier (Optional) A valid VLAN identifier associated with the PPPoE client for which you want to clear lockout.

- **Range:** 1 through 4094

Required Privilege Level

clear

Sample Output

clear pppoe lockout vlan-identifier device-name (Untagged VLAN on Aggregated Ethernet Bundle)

```
user@host> clear pppoe lockout vlan-identifier device-name ae3
```

clear pppoe lockout vlan-identifier device-name vlan-id (Single-Tagged VLAN on Gigabit Ethernet Interface)

```
user@host> clear pppoe lockout vlan-identifier device-name ge-2/0/0 vlan-id 2000
```

clear pppoe lockout vlan-identifier device-name svlan-id vlan-id aci (Dual-Tagged VLAN on 10-Gigabit Ethernet Interface Where ACI Matches Regular Expression)

```
user@host> clear pppoe lockout vlan-identifier device-name xe-1/0/0 svlan-id 10 vlan-id 20 aci
““Relay-identifier atm 1/0:100\.*”
```

clear pppoe lockout vlan-identifier device-name svlan-id vlan-id mac-address (Dual-Tagged VLAN on Aggregated Ethernet Bundle with Specified MAC Address)

```
user@host> clear pppoe lockout vlan-identifier device-name ae0 svlan-id 1 vlan-id 100 mac-
address 00:00:5e:00:53:30
```

Release Information

Command introduced in Junos OS Release 15.2.

RELATED DOCUMENTATION

Clearing Lockout of PPPoE Subscriber Sessions

Configuring Lockout of PPPoE Subscriber Sessions

clear pppoe statistics

IN THIS SECTION

- Syntax | 3833
- Description | 3833
- Options | 3833
- Required Privilege Level | 3833
- Output Fields | 3833
- Sample Output | 3833
- Release Information | 3834

Syntax

```
clear pppoe statistics
<interface interface-name>
<underlying-interface-name>
```

Description

Reset PPPoE session statistics information.

Options

none	Reset PPPoE statistics for all interfaces.
underlying-interface-name	(Optional) Reset PPPoE statistics for the specified underlying PPPoE interface.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pppoe statistics

```
user@host> clear pppoe statistics
```

clear pppoe statistics

```
user@host> clear pppoe statistics ge-4/0/3.2
```

Release Information

Command introduced before Junos OS Release 7.4.

underlying-interface-name option introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

show pppoe statistics | 23271

clear protection-group ethernet-ring statistics

IN THIS SECTION

- Syntax | 3834
- Description | 3834
- Options | 3835
- Required Privilege Level | 3835
- Output Fields | 3835
- Sample Output | 3835
- Sample Output | 3835
- Release Information | 3835

Syntax

```
clear protection-group ethernet-ring statistics
 <group-name  group-name>
```

Description

On MX Series routers, clear the statistics for all Ethernet ring protection groups or a specific Ethernet ring protection group.

Options

group-name *group-name* (Optional) Clear the Ethernet ring protection statistics for the specified group.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear protection-group ethernet-ring statistics

To clear all Ethernet ring protection group statistics for all protection groups, use the following command:

```
user@host> clear protection-group ethernet-ring statistics
```

Sample Output

clear protection-group ethernet-ring statistics

To clear Ethernet ring protection group statistics for the group *my_prot_group*, use the following command:

```
user@host> clear protection-group ethernet-ring statistics group-name my_prot_group
```

Release Information

Command introduced in Junos OS Release 9.4.

clear ptp statistics stream

IN THIS SECTION

- Syntax | 3836
- Description | 3836
- Options | 3836
- Required Privilege Level | 3836
- Release Information | 3836

Syntax

```
clear ptp statistics stream [all | <stream id>]
```

Description

Clear the stream statistics for primary, client and stateful interfaces based on their stream ID or for all the streams.

Options

all Clears all stream statistics.

stream id Clears the statistics for the specified stream ID.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

show ptp statistics | 23408

clear remote-device-management statistics

IN THIS SECTION

- Syntax | 3837
- Description | 3837
- Options | 3837
- Required Privilege Level | 3838
- Output Fields | 3838
- Sample Output | 3838
- Release Information | 3838

Syntax

```
clear remote-device-management statistics
(summary | service-devices device-name)
```

Description

Clear service statistics for all remote devices globally or statistics for a specific remote service device.

Options

service-devices *device-name*	(Optional) Clear statistics for the specified service device.
summary	(Optional) Clear service statistics for all remote devices.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the "show remote-device-management statistics" on page 23445 command before and after clearing the service statistics to verify the clear operation.

Sample Output

clear remote-device-management statistics (Service Device)

```
user@host> clear remote-device-management statistics service-device olt-xyz
```

clear remote-device-management statistics (Global)

```
user@host> clear remote-device-management statistics summary
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

show remote-device-management statistics | 23445

clear rip general-statistics

IN THIS SECTION

Syntax | 3839

- Syntax (EX Series Switches and QFX Series) | 3839
- Description | 3839
- Options | 3839
- Required Privilege Level | 3839
- Output Fields | 3840
- Sample Output | 3840
- Release Information | 3840

Syntax

```
clear rip general-statistics
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
clear rip general-statistics
```

Description

Clear RIP general statistics.

Options

none	Clear RIP general statistics.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rip general-statistics

```
user@host> clear rip general-statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show rip general-statistics | 23493

clear rip statistics

IN THIS SECTION

- Syntax | 3841
- Syntax (EX Series Switches and QFX Series) | 3841
- Description | 3841
- Options | 3841
- Required Privilege Level | 3841
- Output Fields | 3842
- Sample Output | 3842
- Release Information | 3842

Syntax

```
clear rip statistics
<instance (all | instance-name)>
<logical-system (all | logical-system-name)>
<neighbor>
<peer (all | address)>
```

Syntax (EX Series Switches and QFX Series)

```
clear rip statistics
<instance (all | instance-name)>
<neighbor>
```

Description

Clear RIP statistics.

Options

none	Reset RIP counters for all neighbors for all routing instances.
instance (all | *instance-name*)	(Optional) Clear RIP statistics for all instances or for the specified routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
neighbor	(Optional) Clear RIP statistics for the specified neighbor only.
peer (all | *address*)	(Optional) Clear RIP statistics for a single peer or all peers.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rip statistics

```
user@host> clear rip statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show rip statistics | 23499

clear ripng general-statistics

IN THIS SECTION

- Syntax | 3843
- Syntax (EX Series Switches) | 3843
- Description | 3843
- Options | 3843
- Required Privilege Level | 3843
- Output Fields | 3843
- Sample Output | 3843
- Release Information | 3844

Syntax

```
clear ripng general-statistics
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
clear ripng general-statistics
```

Description

Clear RIP next generation (RIPng) general statistics.

Options

none	Clear RIPng general statistics.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ripng general-statistics

```
user@host> clear ripng general-statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ripng general-statistics | 23504

clear ripng statistics

IN THIS SECTION

- Syntax | 3844
- Syntax (EX Series Switch) | 3844
- Description | 3845
- Options | 3845
- Required Privilege Level | 3845
- Output Fields | 3845
- Sample Output | 3845
- Release Information | 3845

Syntax

```
clear ripng statistics
<instance   | name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch)

```
clear ripng statistics
<instance   | name>
```

Description

Clear RIP next-generation (RIPng) statistics.

Options

none	Reset RIPng counters for all neighbors for all routing instances.
instance	(Optional) Reset RIPng counters for the specified instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
name	(Optional) Reset RIPng counters for the specified neighbor.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ripng statistics

```
user@host> clear ripng statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ripng statistics | 23508

clear rsvp session

IN THIS SECTION

- Syntax | 3846
- Syntax (EX and QFX Series Switches) | 3846
- Description | 3847
- Options | 3847
- Required Privilege Level | 3847
- Output Fields | 3847
- Sample Output | 3848
- Release Information | 3848

Syntax

```
clear rsvp session
<all>
<connection-destination address>
<connection-source address>
<gracefully>
<logical-system (all | logical-system-name)>
<lsp-id identifier>
<name name>
<optimize-fast-reroute>
<tunnel-id identifier>
```

Syntax (EX and QFX Series Switches)

```
clear rsvp session
<connection-destination address>
<connection-source address>
<gracefully>
<lsp-id identifier>
<name name>
```

```
<optimize-fast-reroute>
<tunnel-id identifier>
```

Description

Reset and restart Resource Reservation Protocol (RSVP) sessions.

Options

all	Clear all RSVP sessions for which this routing device is the ingress, transit, or egress routing device.
connection-source *address*	(Optional) Source address for GMPLS and MPLS LSPs from the RSVP sender template.
connection-destination *address*	(Optional) Destination address for GMPLS and MPLS LSPs from the RSVP sender template.
gracefully	(Optional) Gracefully reset an RSVP session for a nonpacket LSP in two passes. In the first pass, the Admin-Status object is signaled along the path to the other endpoint of the RSVP session. In the second pass, the path used by the RSVP session is torn down. This option can only be used on the ingress or egress routing device of the RSVP session and is only valid for nonpacket LSPs.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
lsp-id *identifier*	(Optional) LSP identifier (source port) for the RSVP sender template.
name *name*	(Optional) Reset and restart the specified RSVP session.
optimize-fast-reroute	(Optional) Begin fast reroute optimization.
tunnel-id *identifier*	(Optional) Tunnel identifier (destination port) for the RSVP session.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rsvp session all

```
user@host> clear rsvp session all
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear mpls lsp | 3713

show rsvp session | 23860

clear rsvp statistics

IN THIS SECTION

- Syntax | 3849
- Syntax (EX Series Switches) | 3849
- Description | 3849
- Options | 3849
- Required Privilege Level | 3849
- Output Fields | 3849
- Sample Output | 3849
- Release Information | 3850

Syntax

```
clear rsvp statistics
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
clear rsvp statistics
```

Description

Clear Resource Reservation Protocol (RSVP) packet and error statistics.

Options

none	Clear RSVP packet and error statistics.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rsvp statistics

```
user@host> clear rsvp statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show rsvp statistics | 23878

clear security advance-policy-based-routing sla statistics

IN THIS SECTION

- Syntax | 3850
- Description | 3850
- Required Privilege Level | 3850
- Output Fields | 3851
- Sample Output | 3851
- Release Information | 3851

Syntax

```
clear security advance-policy-based-routing sla statistics
```

Description

Clears SLA rule-specific statistics and counters.

Required Privilege Level

view

Output Fields

This command produces no output.

Sample Output

Release Information

Command introduced in Junos OS Release 15.1X49-D130.

RELATED DOCUMENTATION

Application Quality of Experience

Advanced Policy-Based Routing

clear security alarms

IN THIS SECTION

- Syntax | 3851
- Description | 3852
- Options | 3852
- Required Privilege Level | 3853
- Output Fields | 3853
- Sample Output | 3853
- Release Information | 3855

Syntax

```
clear security alarms
<all>
<alarm-id id-number>
```

```
<alarm-type [ types ]>
<newer-than YYYY-MM-DD.HH:MM:SS>
<older-than YYYY-MM-DD.HH:MM:SS>
<process process>
<severity severity>
```

Description

Clear (acknowledge) the alarms that are active on the device.

Options

all (Optional) Clear all active alarms.

alarm-id *id-number* (Optional) Clear the specified alarm.

alarm-type [*types*] (Optional) Clear the specified alarm type or a set of types.

You can specify one or more of the following alarm types:

- `authentication`
- `cryptographic-self-test`
- `decryption-failures`
- `encryption-failures`
- `ike-phase1-failures`
- `ike-phase2-failures`
- `key-generation-self-test`
- `non-cryptographic-self-test`
- `policy`
- `replay-attacks`

newer-than *YYYY-MM-DD.HH:MM:SS* (Optional) Clear active alarms that were raised after the specified date and time.

older-than *YYYY-MM-DD.HH:MM:SS* (Optional) Clear active alarms that were raised before the specified date and time.

process *process*

(Optional) Clear active alarms that were raised by the specified system process.

severity *severity*

(Optional) Clear active alarms of the specified severity.

You can specify the following severity levels:

- `alert`
- `crit`
- `debug`
- `emerg`
- `err`
- `info`
- `notice`
- `warning`

Required Privilege Level

security—To view this statement in the configuration.

Output Fields

This command produces no output, except a statement about the number of security alarms cleared.

Sample Output

clear security alarms all

```
[3 SECURITY ALARMS] user@router> clear security alarms all

3 security alarms cleared
```

clear security alarms alarm-id <alarm-id>

```
[3 SECURITY ALARMS] user@router> clear security alarms alarm-id 1

1 security alarm cleared
```

clear security alarms alarm-type authentication

```
[3 SECURITY ALARMS] user@router> clear security alarms alarm-type authentication

3 security alarms cleared
```

clear security alarms newer-than <time>

```
[3 SECURITY ALARMS] user@router> clear security alarms newer-than 2010-01-19.13:41:59

1 security alarm cleared
```

show security alarms older-than <time>

```
[3 SECURITY ALARMS] user@router> clear security alarms older-than 2010-01-19.13:41:59

2 security alarms cleared
```

show security alarms process <process>

```
[3 SECURITY ALARMS] user@router> clear security alarms process sshd

3 security alarms cleared
```

show security alarms severity <severity>

```
[3 SECURITY ALARMS] user@router> clear security alarms severity notice

3 security alarms cleared
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

show security alarms | 23952

Troubleshooting Security Policies

clear security alg h323 counters

IN THIS SECTION

- Syntax | 3855
- Description | 3856
- Options | 3856
- Required Privilege Level | 3856
- Output Fields | 3856
- Sample Output | 3856
- Release Information | 3856

Syntax

```
clear security alg h323 counters
<logical-system | root-logical-system | tenant>
```

Description

use this command to clear information about H.323 Application Layer Gateway (ALG) counters.

Options

- `logical-system`—Name of the logical system.
- `root-logical-system`—Displays root logical system as default.
- `tenant`—Name of the tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear security alg h323 counters

```
user@host> clear security alg h323 counters
```

clear security alg h323 counters tenant TN1

```
user@host> clear security alg h323 counters tenant TN1
```

clear security alg h323 counters tenant all

```
user@host> clear security alg h323 counters tenant all
```

Release Information

Statement introduced in Junos OS Release 8.5.

The `logical-system` option is introduced in Junos OS Release 18.2R1.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

h323 | 7961

show security alg h323 counters | 23958

clear security alg ike-esp-nat

IN THIS SECTION

- Syntax | 3857
- Description | 3857
- Required Privilege Level | 3857
- Output Fields | 3858
- Sample Output | 3858
- Release Information | 3858

Syntax

```
clear security alg ike-esp-nat
```

Description

Use this command to clear state information about Application Layer Gateway (ALG) for IKE and ESP.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear security alg ike-esp-nat

```
user@host> clear security alg ike-esp-nat
10 active IKE-ESP alg state cleared
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security alg ike-esp-nat summary | 23963

clear security alg ike-esp-nat | 3857

clear security alg sccp calls

IN THIS SECTION

- Syntax | 3859
- Description | 3859
- Options | 3859
- Required Privilege Level | 3859
- Output Fields | 3859
- Release Information | 3859

Syntax

```
clear security alg sccp calls
<node (          node-id         | all | local | primary)>
```

Description

use this command to clear Skinny Client Protocol (SCCP) Application Layer Gateway (ALG) call information.

Options

- none—Clear all SCCP ALG calls.
- `node`—(Optional) For chassis cluster configurations, clear SCCP calls on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.
  - `local` —Clear the local node.
  - `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 8.5; `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

sccp | **16341**

show security alg sccp calls | 23966

clear security alg sccp counters

IN THIS SECTION

- Syntax | 3860
- Description | 3860
- Options | 3860
- Required Privilege Level | 3861
- Output Fields | 3861
- Sample Output | 3861
- Release Information | 3861

Syntax

```
clear security alg sccp counters
logical-system
root-logical-system
tenant
```

Description

Use this command to clear Skinny Client Control Protocol (SCCP) Application Layer Gateway (ALG) counters.

Options

- `logical-system`—Name of the logical system.
- `root-logical-system`—Displays root logical system as default.
- `tenant`—Name of the tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear security alg sccp counters

```
user@host> clear security alg sccp counters
```

clear security alg sccp counters tenant TN1

```
user@host> clear security alg sccp counters tenant TN1
```

clear security alg sccp counters tenant all

```
user@host> clear security alg sccp counters tenant all
```

Release Information

Command introduced in Junos OS Release 8.5.

The `logical-system` option is introduced in Junos OS Release 18.2R1.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

sccp | 16341

show security alg sccp counters | 23969

clear security alg sip calls

IN THIS SECTION

- Syntax | 3862
- Description | 3862
- Options | 3862
- Required Privilege Level | 3863
- Output Fields | 3863
- Sample Output | 3863
- Release Information | 3863

Syntax

```
clear security alg sip calls
<node (          node-id         | all | local | primary)>
```

Description

Use this command to clear Session Initiation Protocol (SIP) Application Layer Gateway (ALG) call information.

Options

- none—Clear all SIP ALG calls.
- `node`—(Optional) For chassis cluster configuration, clear SIP calls on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.
  - `local` —Clear the local node.
  - `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

Sample Output

clear security alg sip calls

```
user@host> clear security alg sip calls
```

Release Information

Command modified in Junos OS Release 9.2; `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

sip (Security) | 27977

show security alg sip calls | 23973

clear security alg sip counters

IN THIS SECTION

- Syntax | 3864
- Description | 3864
- Options | 3864
- Required Privilege Level | 3864
- Output Fields | 3864
- Sample Output | 3864
- Release Information | 3865

Syntax

```
clear security alg sip counters
<logical-system | root-logical-system | tenant>
```

Description

Use this command to clear Session Initiation Protocol (SIP) Application Layer Gateway (ALG) counters.

Options

- `logical-system`—Name of the logical system.
- `root-logical-system`—Displays root logical system as default.
- `tenant`—Name of the tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear security alg sip counters

```
user@host> clear security alg sip counters
```

clear security alg sip counters tenant TN1

```
user@host> clear security alg sip counters tenant TN1
```

clear security alg sip counters tenant all

```
user@host> clear security alg sip counters tenant all
```

Release Information

Statement modified in Junos OS Release 9.2.

The `logical-system` option is introduced in Junos OS Release 18.2R1.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

sip (Security) | 27977

show security alg sip counters | 23975

clear security application-firewall rule-set statistics

IN THIS SECTION

- Syntax | 3865
- Description | 3866
- Required Privilege Level | 3866
- Output Fields | 3866
- Release Information | 3866

Syntax

```
clear security application-firewall rule-set statistics
```

Description

Clear all the security application firewall rule set statistics information.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show security application-firewall rule-set | 24012

clear security application-firewall rule-set statistics logical-system

IN THIS SECTION

- Syntax | 3867
- Description | 3867
- Options | 3867
- Required Privilege Level | 3868
- Output Fields | 3868

Release Information | 3868

Syntax

The primary, or root, administrator can issue the following statements:

```
clear security application-firewall rule-set statistics [logical-system logical-system-name |
all | root-logical-system]
```

The user logical system administrator can issue the following statement:

```
clear security application-firewall rule-set statistics all
```

Description

Clear all security application firewall rule set statistics.

> **NOTE**: User logical system administrators can clear statistics only for the logical systems they can access. For information about primary and user administrator roles in logical systems, see Understanding the Primary Logical Systems and the Primary Administrator Role.

Starting in Junos OS Release 18.2R1 application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration.

Options

`logical-system-name`—Name of a specific logical system.

`all`—(default) Clear all rule set statistics for a specific logical system or all logical systems.

`root-logical-system`—Clear application firewall rule set statistics on the root logical system (primary administrator only).

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show security application-firewall rule-set logical-system | 24017

clear security datapath-debug counters

IN THIS SECTION

- Syntax | 3868
- Description | 3868
- Required Privilege Level | 3869
- Output Fields | 3869
- Release Information | 3869

Syntax

```
clear security datapath-debug counters
```

Description

Clear all data path-debugging counters.

NOTE: Data path debugging is supported on SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show security datapath-debug capture | 24025
show security datapath-debug counter | 24027

clear security dns-cache

IN THIS SECTION

- Syntax | 3870
- Description | 3870
- Options | 3870
- Required Privilege Level | 3870
- Release Information | 3870

Syntax

```
clear security dns-cache <dns-name dns-name>
```

Description

Reset DNS cache information.

> **NOTE**: This command is only available to the primary administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

Options

- `dns-name`—Clear DNS cache information for the specified name.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

show security dns-cache | 24028

Understanding the Primary Logical Systems and the Primary Administrator Role

clear security dynamic-vpn all

IN THIS SECTION

- Syntax | 3871
- Description | 3871
- Required Privilege Level | 3871
- Output Fields | 3871
- Sample Output | 3872
- Release Information | 3872

Syntax

```
clear security dynamic-vpn all
```

Description

Clear all dynamic VPN user connections. This feature is supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

In Junos OS Release 21.4R1, we’ve deprecated the `clear security dynamic-vpn all` operational mode command and we might remove it completely in a future release.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security dynamic-vpn all

```
user@host> clear security dynamic-vpn all
2 user connection entries cleared
```

Release Information

Command introduced in Junos Release 10.4.

RELATED DOCUMENTATION

show security dynamic-vpn users | 24053

show security dynamic-vpn users terse | 24055

clear security dynamic-vpn user

IN THIS SECTION

- Syntax | 3872
- Description | 3873
- Required Privilege Level | 3873
- Output Fields | 3873
- Sample Output | 3873
- Release Information | 3873

Syntax

```
clear security dynamic-vpn user username ike-id id
```

Description

Clear the dynamic VPN user connection for the specified username. This feature is supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

In Junos OS Release 21.4R1, we've deprecated the `clear security dynamic-vpn user` operational mode command and we might remove it completely in a future release.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security dynamic-vpn user

```
user@host> clear security dynamic-vpn user user ike-id bob.example.net
Connection entry for user user has been cleared
```

Release Information

Command introduced in Junos Release 10.4.

RELATED DOCUMENTATION

show security dynamic-vpn users | 24053

show security dynamic-vpn users terse | 24055

clear security firewall-authentication history

IN THIS SECTION

- Syntax | 3874
- Description | 3874
- Options | 3874
- Required Privilege Level | 3875
- Output Fields | 3875
- Sample Output | 3875
- Sample Output | 3876
- Release Information | 3876

Syntax

```
clear security firewall-authentication history
<node (node-id | all | local | primary)>
<address>
<identifier>
<logical-system (logical-system-name | all)>
<root-logical-system (address | auth-type | from-zone | identifier | tenant | to-zone>
<tenant (tenant-name |all)>
```

Description

Clear all firewall authentication history information.

Options

- `node`—(Optional) For chassis cluster configurations, clear all firewall authentication history on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.

  - `local` —Clear the local node.
  - `primary`—Clear the primary node.
- address—Display authentication entries based on ip address.
- identifier—Display authentication entries by id.
- logical-system—Display firewall authentication tables based on logical system name.
- `node`—(Optional) For chassis cluster configurations, display firewall authentication details for all users on a specific node.
  - *node-id*—Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.
- root-logical-system—Display firewall authentication tables for root logical system.
- tenant—Display firewall authentication tables based on tenant name.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

This command produces no output.

Sample Output

clear security firewall-authentication history

```
user@host> clear security firewall-authentication history
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication history node 1

```
user@host> clear security firewall-authentication history node 1
node1:
--------------------------------------------------------------------------
```

clear security firewall-authentication history tenant all

```
user@host> clear security firewall-authentication history tenant all
```

Release Information

Command introduced in Junos OSRelease 8.5. The `node` options added in Junos OS Release 9.0.

The `tenant` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show security firewall-authentication history | 24058

clear security firewall-authentication history address

IN THIS SECTION

- Syntax | 3877
- Description | 3877
- Options | 3877
- Required Privilege Level | 3877
- Output Fields | 3877

Sample Output | 3878
Sample Output | 3878
Release Information | 3878

Syntax

```
clear security firewall-authentication history address          address
<node (          node-id         | all | local | primary)>
```

Description

Clear firewall authentication history for this source IP address.

Options

- `address` *address* —Source IP address for which to clear firewall authentication history.
- none—Clear all firewall authentication history for this address.
- `node`—(Optional) For chassis cluster configurations, clear firewall authentication history for this address on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.
  - `local` —Clear the local node.
  - `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security firewall-authentication history address 10.0.0.1

```
user@host> clear security firewall-authentication history address 10.0.0.1
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication history address 192.0.2.2 node 1

```
user@host> clear security firewall-authentication history address 192.0.2.2  node 1
node1:
--------------------------------------------------------------------------
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Firewall User Authentication Overview

clear security firewall-authentication history identifier

IN THIS SECTION

- Syntax | 3879
- Description | 3879
- Options | 3879

- Required Privilege Level | 3879
- Output Fields | 3880
- Sample Output | 3880
- Sample Output | 3880
- Release Information | 3880

Syntax

```
clear security firewall-authentication history identifier          identifier
<node (          node-id         | all | local | primary)>
```

Description

Clear firewall authentication history information for the authentication with this identifier.

Options

- `identifier` *`identifier`* —Identification number of the authentication for which to clear authentication history.
- none—Clear all firewall authentication history information for the authentication with this identifier.
- `node`—(Optional) For chassis cluster configurations, clear firewall authentication history on a specific node for the authentication with this identifier.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.
  - `local` —Clear the local node.
  - `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security firewall-authentication history identifier 2

```
user@host> clear security firewall-authentication history identifier 2
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication history identifier 2 node 1

```
user@host> clear security firewall-authentication history identifier 2 node 1
node1:
--------------------------------------------------------------------------
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Firewall User Authentication Overview

clear security firewall-authentication users

IN THIS SECTION

- Syntax | 3881
- Description | 3881
- Options | 3881
- Required Privilege Level | 3882
- Output Fields | 3882
- Sample Output | 3882
- Sample Output | 3883
- Release Information | 3883

Syntax

```
clear security firewall-authentication users
<node (node-id | all | local | primary)>
<address>
<identifier>
<logical-system (logical-system-name | all)>
<root-logical-system (address | auth-type | from-zone | identifier | tenant | to-zone>
<tenant (tenant-name |all)>
```

Description

Clear firewall authentication tables for all users.

Options

- `node`—(Optional) For chassis cluster configurations, clear firewall authentication details for all users on a specific node.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.

  - `local` —Clear the local node.
  - `primary`—Clear the primary node.
- address—Display authentication entries based on ip address.
- identifier—Display authentication entries by id.
- logical-system—Display firewall authentication tables based on logical system name.
- `node`—(Optional) For chassis cluster configurations, display firewall authentication details for all users on a specific node.
  - *node-id*—Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.
- root-logical-system—Display firewall authentication tables for root logical system.
- tenant—Display firewall authentication tables based on tenant name.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

This command produces no output.

Sample Output

clear security firewall-authentication users

```
user@host> clear security firewall-authentication users node 1
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication users node 1

```
user@host> clear security firewall-authentication users node 1
node1:
--------------------------------------------------------------------------
```

clear security firewall-authentication users tenant all

```
user@host> clear security firewall-authentication users tenant all
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

The `tenant` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show security firewall-authentication users | 24071

clear security firewall-authentication users address

IN THIS SECTION

- Syntax | 3884
- Description | 3884
- Options | 3884
- Required Privilege Level | 3884
- Output Fields | 3884

Sample Output | 3885
Sample Output | 3885
Release Information | 3885

Syntax

```
clear security firewall-authentication users address          address
<node (          node-id         | all | local | primary)>
```

Description

Clear information about the users at the specified IP address that are currently authenticated.

Options

- `address` *address* —IP address for which to clear user firewall authentication information.
- none—Clear all the firewall authentication information for users at this IP address.
- `node`—(Optional) For chassis cluster configurations, clear user firewall authentication entries on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.
  - `local` —Clear the local node.
  - `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security firewall-authentication users address 198.51.100.11

```
user@host> clear security firewall-authentication users address 198.51.100.11
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication users address 198.51.100.11 node 1

```
user@host> clear security firewall-authentication users address 198.51.100.11 node 1
node1:
--------------------------------------------------------------------------
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Firewall User Authentication Overview

clear security firewall-authentication users identifier

IN THIS SECTION

- Syntax | 3886
- Description | 3886
- Options | 3886

- Required Privilege Level | 3886
- Output Fields | 3886
- Sample Output | 3887
- Sample Output | 3887
- Release Information | 3887

Syntax

```
clear security firewall-authentication users identifier          identifier
<node (          node-id         | all | local | primary)>
```

Description

Clear firewall authentication details about the user with this identification number.

Options

- none—Identification number of the user for which to clear authentication details.
- `node`—(Optional) For chassis cluster configurations, clear the firewall authentication details on a specific node (device) in the cluster for the user with this identification number.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.
  - `local` —Clear the local node.
  - `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security firewall-authentication users identifier 2

```
user@host> clear security firewall-authentication users identifier 2
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication users identifier 2 node 1

```
user@host> clear security firewall-authentication users identifier 2 node 1
node1:
--------------------------------------------------------------------------
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Firewall User Authentication Overview

clear security flow ip-action

IN THIS SECTION

- Syntax | 3888
- Description | 3888
- Options | 3888

- Required Privilege Level | 3889
- Output Fields | 3889
- Sample Output | 3889
- Release Information | 3890

Syntax

```
clear security flow ip-action   [filter]
```

Description

Clear IP-action entries, based on filtered options, for IP sessions running on the device.

Options

filter—Filter the display based on the specified criteria.

The following filters display those sessions that match the criteria specified by the filter. Refer to the sample output for filtered output examples.

`all` | [*filter*]—All active sessions on the device.

`destination-port` *`destination-port`*—Destination port number of the traffic. Range is 1 through 65,535.

`destination-prefix` *`destination-prefix`*—Destination IP prefix or address.

`family (inet | inet6)` [*filter*]—IPv4 traffic or IPv6-NATPT traffic and filtered options.

`logical-system` *`logical-system-name`* | `all` [*filter*]—Specified logical system or all logical systems.

`protocol` *`protocol-name`* / *`protocol-number`* [*filter*]—Protocol name or number and filtered options.

- `ah` or `51`
- `egp` or `8`
- `esp` or `50`
- `gre` or `47`
- `icmp` or `1`

- `icmp6` or `58`
- `ipip` or `4`
- `ospf` or `89`
- `pim` or `103`
- `rsvp` or `46`
- `sctp` or `132`
- `tcp` or `6`
- `udp` or `17`

`root-logical-system` [*filter*]—Default logical system information and filtered options.

`source-port` *`source-port`*—Source port number of the traffic. Range is 1 through 65,535.

`source-prefix` *`source-prefix`*—Source IP prefix or address of the traffic.

Required Privilege Level

clear

Output Fields

When you enter this command, the system responds with the status of your request.

Sample Output

clear security flow ip-action all

```
user@host>clear security flow ip-action all
1008 ip-action entries cleared
```

clear security flow ip-action destination-prefix

```
user@host>clear security flow ip-action destination-prefix 192.0.2.5/24
87 ip-action entries cleared
```

clear security flow ip-action family inet

```
user@host>clear security flow ip-action family inet
2479 ip-action entries cleared
```

clear security flow ip-action protocol udp

```
user@host>clear security flow ip-action protocol udp
270 ip-action entries cleared
```

Release Information

Command introduced in Junos OS Release 10.4. Logical systems option introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

show security flow ip-action | 24147

clear security flow session

IN THIS SECTION

- Syntax | 3891
- Description | 3891
- Options | 3891
- Required Privilege Level | 3891
- Output Fields | 3891
- Sample Output | 3892
- Sample Output | 3892
- Release Information | 3892

Syntax

```
clear security flow session
all <node (          node-id         | all | local | primary)>
<nat number>
<nat-port-overload-index number>
<source-nat-pool source-nat-pool-name>
```

Description

Clear all currently active security sessions on the device.

Options

- all—Clear information about all active sessions.
- `node`—(Optional) For chassis cluster configurations, clear all security sessions on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.
  - `local` —Clear the local node.
  - `primary`—Clear the primary node.
- `nat`—Clear sessions with network address translation.
- `nat-port-overload-index`—Clear NAT port overload index.
- `source-nat-pool`—Clear the source NAT pool name.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session all

```
user@host> clear security flow session all
node0:
--------------------------------------------------------------------------
1 active sessions cleared
node1:
-------- ------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session all node 0

```
user@host> clear security flow session all node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 8.5 ; `node` options added in Junos OS Release 9.0.

The `nat`, `nat-port-overload-index`, and `source-nat-pool` are introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

show security flow session

clear security flow session application

IN THIS SECTION

- Syntax | 3893
- Description | 3893
- Options | 3893
- Required Privilege Level | 3894
- Output Fields | 3894
- Sample Output | 3895
- Sample Output | 3895
- Release Information | 3895

Syntax

```
clear security flow session application
                    application-name
<node (          node-id         | all | local | primary)>
```

Description

Use this command to clear currently active sessions for application types or application sets.

Options

- *application-name* —Name of the specified application type or application set.
  - dns—Domain Name System
  - ftp—File Transfer Protocol
  - ignore—Ignore application type
  - mgcp-ca—Media Gateway Control Protocol with Call Agent
  - mgcp-ua—MGCP with User Agent

  - `ms-rpc`—Microsoft RPC
  - `pptp`—Point-to-Point Tunneling Protocol
  - `q931`—ISDN connection control protocol
  - `ras`—RAS
  - `realaudio`—RealAudio
  - `rsh`—UNIX remote shell services
  - `rtsp`—Real-Time Streaming Protocol
  - `sccp`—Skinny Client Control Protocol
  - `sip`—Session Initiation Protocol
  - `sqlnet-v2`—Oracle SQLNET
  - `sun-rpc`—Sun Microsystems RPC
  - `talk`—TALK program
  - `tftp`—Trivial File Transfer Protocol
- `node`—(Optional) For chassis cluster configurations, clear sessions for applications on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.
  - `local` —Clear the local node.
  - `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session application dns

```
user@host> clear security flow session application dns
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session application dns node 0

```
user@host> clear security flow session application dns node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 8.5.

The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

show security flow session application | 24184

clear security flow session application-traffic-control

IN THIS SECTION

- Syntax | 3896
- Description | 3896
- Options | 3896
- Required Privilege Level | 3898
- Output Fields | 3898
- Sample Output | 3898
- Release Information | 3898

Syntax

```
clear security flow session application-traffic-control [Enter] | application-traffic-control-
session-options
```

Description

Clear all application traffic control sessions or the session associated with the specified option.

Application traffic control enables application-aware quality of service, as specified in the rules of the rules set defined for it.

Options

application Name of the application or application set.

- `dns`—Domain Name System
- `ftp`—File Transfer Protocol
- `ignore`—Ignore application type
- `mgcp-ca`—Media Gateway Control Protocol with Call Agent

- `mgcp-ua`—MGCP with User Agent
- `ms-rpc`—Microsoft RPC
- `pptp`—Point-to-Point Tunneling Protocol
- `q931`—ISDN connection control protocol
- `ras`—RAS
- `realaudio`—RealAudio
- `rsh`—UNIX remote shell services
- `rtsp`—Real-Time Streaming Protocol
- `sccp`—Skinny Client Control Protocol
- `sip`—Session Initiation Protocol
- `sqlnet-v2`—Oracle SQLNET
- `sun-rpc`—Sun Microsystems RPC
- `talk`—TALK program
- `tftp`—Trivial File Transfer Protocol

To display the supported applications list on an SRX Series device, enter the following command from configuration mode:

```
user@host# show groups junos-default applications
```

conn-tag

A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U), and the Stream Control Transmission Protocol (STCP) sessions.

The connection tag for GTP-U is the tunnel endpoint identifier (TEID). For SCTP, it is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.

- **Range:** 0 through 4294967295

destination-port

Destination port.

- **Range:** 1 through 65535.

destination-prefix Destination IP prefix or address.

family Protocol family:

- inet—Clear IPv4 sessions.
- inet6–Clear IPv6 sessions.

interface Name of incoming or outgoing interface.

protocol IP protocol number.

source-port Source port.

- **Range:** 1 through 65535.

source-prefix Source IP prefix or address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

The same output is displayed when you enter an option for this command.

Sample Output

command-name

```
user@host> clear security flow session application-traffic-control
number active sessions cleared
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show security flow session

clear security flow session conn-tag

IN THIS SECTION

- Syntax | 3899
- Description | 3899
- Options | 3899
- Required Privilege Level | 3900
- Output Fields | 3900
- Sample Output | 3900
- Release Information | 3900

Syntax

```
clear security flow session conn-tag conn-tag-id-number
```

Description

Clear the session identified by the session connection (conn-tag) identification tag.

The flow session connection tag allows you to add a filter to further distinguish GRSP tunneling protocol, user plane (GTP-U) flow sessions, and Stream Control Transmission Protocol (SCTP) flow sessions.

Options

conn-tag A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U) and the Stream Control Transmission Protocol (STCP) sessions. The connection tag for GTP-U is the tunnel endpoint identifier (TEID) and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.

You can configure the system to include the session connection tag tuple to identify GTP-U session and SCTP sessions by adding the session connection tag to the standard six tuples that identify a session. The system determines the DCP for GTP-U/SCTP by hashing the session connection tag.

The central point architecture distributes GTP-U traffic handled by a gateway GPRS support node (GGSN) and SGSN pair on all SPUs by switching to tunnel endpoint identifier (TEID)-based hash distribution. To handle load-balancing issues, tag-based hash distribution is used to ensure even distribution of SCTP traffic from different associations among all SPUs. (The connection tag for GTP-U is the TEID and for SCTP is the vTag.)

- **Range:** 0 through 4294967295.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

command-name

```
user@host> clear-security-flow-session conn-tag conn-tag-id-number
number active sessions cleared
```

Release Information

Command introduced in Junos OS Release 15.1X49-D40.

clear security flow session destination-port

IN THIS SECTION

- Syntax | 3901
- Description | 3901
- Options | 3901
- Required Privilege Level | 3902
- Output Fields | 3902
- Sample Output | 3902
- Sample Output | 3902
- Release Information | 3902

Syntax

```
clear security flow session destination-port
                    destination-port-number
<node (          node-id         | all | local | primary)>
```

Description

Clear each session that uses the specified destination port

Options

- *destination-port-number* —Number of the destination port.
- node—(Optional) For chassis cluster configurations, clear security sessions on the port on a specific node (device) in the cluster.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - all —Clear all nodes.
  - local —Clear the local node.

- `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session destination-port 1

```
user@host> clear security flow session destination-port 1
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session destination-port 1 node 0

```
user@host> clear security flow session destination-port 1 node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 8.5 ; `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

show security flow session destination-port | 24205

clear security flow session destination-prefix

IN THIS SECTION

- Syntax | 3903
- Description | 3903
- Options | 3903
- Required Privilege Level | 3904
- Output Fields | 3904
- Sample Output | 3904
- Sample Output | 3905
- Release Information | 3905

Syntax

```
clear security flow session destination-prefix
                   destination-IP-prefix
<node (          node-id         | all | local | primary)>
```

Description

Clear sessions that match this destination IP prefix or address.

Options

- *destination-IP-prefix*—Destination IP prefix or address.
- node—(Optional) For chassis cluster configurations, clear sessions that match the IP prefix or address on a specific node (device) in the cluster.

- *node-id*—Identification number of the node. It can be 0 or 1.
- `all`—Clear all nodes.
- `local`—Clear the local node.
- `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session destination-prefix 100.0.0.1

```
clear security flow session destination-prefix 100.0.0.1
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

clear security flow session destination-prefix 10::10

```
user@host> clear security flow session destination-prefix 10::10
1 active sessions cleared
```

Sample Output

Clear security flow session destination-prefix 100.0.0.1 node 0

```
user@host> clear security flow session destination-prefix 100.0.0.1 node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```

Release Information

Command introduced in Release 8.5 of Junos OS.

The `node` options added in Release 9.0 of Junos OS.

Support for IPv6 addresses added in Release 10.2 of Junos OS.

RELATED DOCUMENTATION

show security flow session destination-prefix | 24211

clear security flow session family

IN THIS SECTION

- Syntax | 3906
- Description | 3906
- Options | 3906
- Required Privilege Level | 3906
- Output Fields | 3906
- Sample Output | 3906
- Release Information | 3906

Syntax

```
clear security flow session family (inet | inet6)
```

Description

Clear sessions that match the specified protocol family.

Options

- `inet`—Clear IPv4 sessions.
- `inet6`—Clear IPv6 sessions.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session family inet

```
user@host> clear security flow session family inet
1 active sessions cleared
```

clear security flow session family inet6

```
user@host> clear security flow session family inet6
1 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow session family | 24227

clear security flow session IDP

IN THIS SECTION

- Syntax | 3907
- Description | 3907
- Options | 3907
- Required Privilege Level | 3909
- Output Fields | 3909
- Sample Output | 3909
- Release Information | 3910

Syntax

```
clear security flow session IDP [Enter] | idp-session-parameters
```

Description

Clear all active Intrusion Detection and Prevention (IDP) sessions or an IDP session based on the specified session parameter.

IDP allows you to enforce various attack detection and prevention techniques on network traffic passing through the SRX Series Firewall. The SRX Series offers the same set of IDP signatures that are available on Juniper Networks IDP Series Intrusion Detection and Prevention Appliances to secure networks against attacks.

Options

application Name of the application or application set.

- `dns`—Domain Name System
- `ftp`—File Transfer Protocol
- `ignore`—Ignore application type
- `mgcp-ca`—Media Gateway Control Protocol with Call Agent
- `mgcp-ua`—MGCP with User Agent
- `ms-rpc`—Microsoft RPC
- `pptp`—Point-to-Point Tunneling Protocol
- `q931`—ISDN connection control protocol
- `ras`—RAS
- `realaudio`—RealAudio
- `rsh`—UNIX remote shell services
- `rtsp`—Real-Time Streaming Protocol
- `sccp`—Skinny Client Control Protocol
- `sip`—Session Initiation Protocol
- `sqlnet-v2`—Oracle SQLNET
- `sun-rpc`—Sun Microsystems RPC
- `talk`—TALK program
- `tftp`—Trivial File Transfer Protocol

To display the supported applications list on an SRX Series Firewall, enter the following command from configuration mode:

```
user@host# show groups junos-default applications
```

conn-tag

A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U), and the Stream Control Transmission Protocol (STCP) sessions.

The connection tag for GTP-U is the tunnel endpoint identifier (TEID). For SCTP, it is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.

- **Range:** 0 through 4294967295

destination-port Destination port.

- **Range:** 1 through 65535.

destination-prefix Destination IP prefix or IP address.

family Protocol family:

- inet—Clear IPv4 sessions.
- inet6–Clear IPv6 sessions.

interface Name of incoming or outgoing interface.

protocol IP protocol number.

source-port Source port.

- **Range:** 1 through 65535.

source-prefix Source IP prefix or address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

The same output is displayed when you enter an option for this command.

Sample Output

command-name

```
user@host> clear security flow session idp
number active sessions cleared
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow session

clear security flow session interface

IN THIS SECTION

- Syntax | 3910
- Description | 3910
- Options | 3911
- Required Privilege Level | 3911
- Output Fields | 3911
- Sample Output | 3911
- Sample Output | 3911
- Release Information | 3912

Syntax

```
clear security flow session interface
                    interface-name
<node (          node-id         | all | local | primary)>
```

Description

Clear sessions that use the specified interface.

Options

- *interface-name* —Name of a specific incoming or outgoing interface.
- node—(Optional) For chassis cluster configurations, clear security sessions on the interface on a specific node (device) in the cluster.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - all —Clear all nodes.
  - local —Clear the local node.
  - primary—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session interface ge-0/0/0.0

```
user@host> clear security flow session interface ge-0/0/0.0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session interface ge/0/0.0 node 0

```
user@host> clear security flow session interface ge-0/0/0.0 node 0
node0:
```

```
---------------------------------------------------------------------------
0 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 8.5; `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

show security flow session interface | 24240

clear security flow session nat

IN THIS SECTION

- Syntax | 3912
- Description | 3912
- Options | 3913
- Required Privilege Level | 3914
- Output Fields | 3915
- Sample Output | 3915
- Release Information | 3915

Syntax

```
clear security flow session nat <Enter> | nat-session-parameter
```

Description

Clear all active sessions with NAT configurations or the active NAT session identified by a session parameter.

Network Address Translation (NAT) is now used primarily to improve traffic security. But it also offers solutions to IP address constraints prior to the advent and implementation of IPv6. NAT allows you to remap one IP address space to another. Network address information in the IP datagram packet headers are modified to achieve the translation.

Options

application Name of the application protocol

- `dns`—Domain Name System
- `ftp`—File Transfer Protocol
- `ignore`—Ignore application type
- `mgcp-ca`—Media Gateway Control Protocol with Call Agent
- `mgcp-ua`—MGCP with User Agent
- `ms-rpc`—Microsoft RPC
- `pptp`—Point-to-Point Tunneling Protocol
- `q931`—ISDN connection control protocol
- `ras`—RAS
- `realaudio`—RealAudio
- `rsh`—UNIX remote shell services
- `rtsp`—Real-Time Streaming Protocol
- `sccp`—Skinny Client Control Protocol
- `sip`—Session Initiation Protocol
- `sqlnet-v2`—Oracle SQLNET
- `sun-rpc`—Sun Microsystems RPC
- `talk`—TALK program
- `tftp`—Trivial File Transfer Protocol

To display a list of the supported applications on an SRX Series device, enter the following command from configuration mode:

```
user@host# show groups junos-default applications
```

conn-tag Session identified by the specified conn-tag.

A conn-tag is a 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U), and the Stream Control Transmission Protocol (STCP) sessions.

The connection tag for GTP-U is the tunnel endpoint identifier (TEID). For SCTP, it is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.

Session connection identifiers are in the following range:

- **Range:** 0 through 4294967295

destination-port Destination port.

- **Range:** 1 through 65535

destination-prefix Destination IP prefix or address.

family Protocol family:

- inet—Clear IPv4 sessions
- inet6–Clear IPv6 sessions

interface Name of incoming or outgoing interface.

protocol IP protocol number.

source-port Source port.

- **Range:** 1 through 65535.

source-prefix Source IP prefix or IP address.

Required Privilege Level

clear

Output Fields

Displays a message reporting the number of active sessions cleared. The same message is displayed when any specific option is entered.

Sample Output

command-name

```
user@host> clear security flow session nat
number active sessions cleared
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow session

clear security flow session protocol

IN THIS SECTION

- Syntax | 3916
- Description | 3916
- Options | 3916
- Required Privilege Level | 3917
- Output Fields | 3917
- Sample Output | 3917
- Sample Output | 3917
- Release Information | 3918

Syntax

```
clear security flow session protocol          protocol-name        |          protocol-
number
<node (          node-id         | all | local | primary)>
```

Description

Clear each session that uses the specified IP protocol.

Options

- *protocol-name* — (Optional) Networking protocol name. The following text values are supported.
  - ah—IP Security Authentication Header
  - egp—Exterior gateway protocol
  - esp—IPsec Encapsulating Security Payload
  - gre—Generic routing encapsulation
  - icmp—Internet Control Message Protocol
  - igmp—Internet Group Management Protocol
  - ipip—IP over IP
  - ospf—Open Shortest Path First
  - pim—Protocol Independent Multicast
  - rsvp—Resource Reservation Protocol
  - sctp—Stream Control Transmission Protocol
  - tcp—Transmission Control Protocol
  - udp—User Datagram Protocol
- *protocol-number* —(Optional) Numeric protocol value. For a complete list of possible numeric values, see RFC 1700, *Assigned Numbers (for the Internet Protocol Suite)*.

  Range: 0 through 255

- `node`—(Optional) For chassis cluster configurations, clear security on a specific node (device) in the cluster for the user with this identification number.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.
  - `local` —Clear the local node.
  - `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session protocol pim

```
user@host> clear security flow session protocol pim
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session protocol 0

```
user@host> clear security flow session protocol 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
```

```
---------------------------------------------------------------------------
0 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 8.5 ; `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

show security flow session protocol | 24263

clear security flow session resource-manager

IN THIS SECTION

- Syntax | 3918
- Description | 3918
- Options | 3919
- Required Privilege Level | 3919
- Output Fields | 3919
- Sample Output | 3919
- Sample Output | 3919
- Release Information | 3920

Syntax

```
clear security flow session resource-manager
<node (          node-id         | all | local | primary)>
```

Description

Clear resource-manager sessions.

Options

- `node`—(Optional) For chassis cluster configurations, clear the resource manager sessions on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.
  - `local` —Clear the local node.
  - `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session resource-manager

```
user@host> clear security flow session resource-manager
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session resource-manager node 0

```
user@host> clear security flow session resource-manager node 0
node0:
```

```
---------------------------------------------------------------------------
0 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 8.5; `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

show security flow session resource-manager | 24270

clear security flow session services-offload

IN THIS SECTION

- Syntax | 3920
- Description | 3920
- Options | 3921
- Required Privilege Level | 3922
- Output Fields | 3922
- Sample Output | 3922
- Release Information | 3924

Syntax

```
clear security flow session services-offload  [filter]
```

Description

Clear services-offload security sessions, based on filtered options, on the device. This command also clears a services-offload security session from both the network processor and the Services Processing Unit (SPU) on which the specified session was installed.

Options

filter—Filter the display based on the specified criteria.

The following filters clear those sessions that match the criteria specified by the filter. Refer to the sample output for filtered output examples.

`application` *`application-name`*—Application name.

`destination-port` *`destination-port`*—Destination port number. Range is from 1 through 65,535.

`destination-prefix` *`destination-prefix`*—Destination IP prefix or address.

`family (inet | inet6)`—IPv4 traffic or IPv6-NAT-PT traffic.

`interface` *`interface-name`*—Incoming or outgoing interface name.

`logical-system` *`logical-system-name`* `| all`—Specified logical system name or all logical systems.

`protocol` *`protocol-name`* / *`protocol-number`* —Protocol name or number.

- `ah` or `51`
- `egp` or `8`
- `esp` or `50`
- `gre` or `47`
- `icmp` or `1`
- `icmp6` or `58`
- `igmp` or `2`
- `ipip` or `4`
- `ospf` or `89`
- `pim` or `103`
- `rsvp` or `46`
- `sctp` or `132`
- `tcp` or `6`
- `udp` or `17`

`root-logical-system` [*filter*]—Root logical system information and filtered options.

`source-port` *`source-port`*—Source port number of the traffic. Range is from 1 through 65,535.

`source-prefix` *`source-prefix`*—Source IP prefix or address of the traffic.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session services-offload

```
user@host>clear security flow session services-offload
0 active sessions cleared
```

clear security flow session services-offload application

```
user@host>clear security flow session services-offload dns
0 active sessions cleared
```

clear security flow session services-offload destination-port

```
user@host>clear security flow session services-offload destination-port 1
0 active sessions cleared
```

clear security flow session services-offload destination-prefix

```
user@host>clear security flow session services-offload destination-prefix 10.0.0.1
0 active sessions cleared
```

clear security flow session services-offload family

```
user@host>clear security flow session services-offload family inet
1 active sessions cleared
```

clear security flow session services-offload interface

```
user@host>clear security flow session services-offload interface ge-0/0/0.0
0 active sessions cleared
```

clear security flow session services-offload logical-system

```
user@host>clear security flow session services-offload logical-system all
0 active sessions cleared
```

clear security flow session services-offload protocol

```
user@host>clear security flow session services-offload protocol pim
0 active sessions cleared
```

clear security flow session services-offload root-logical-system

```
user@host>clear security flow session services-offload root-logical-system application dns
0 active sessions cleared
```

clear security flow session services-offload source-port

```
user@host>clear security flow session services-offload source-port 1
0 active sessions cleared
```

clear security flow session services-offload source-prefix

```
user@host>clear security flow session services-offload source-prefix 10.0.0.1
0 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 11.4.

Starting with Junos OS Release 15.1X49-D10 and and Junos OS Release 17.3R1, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) with Express Path (formerly known as *services offloading*) support are introduced for SRX5400, SRX5600, and SRX5800 devices.

RELATED DOCUMENTATION

show security flow session services-offload | 24276

clear security flow session session-identifier

IN THIS SECTION

- Syntax | 3925
- Description | 3925
- Options | 3925
- Required Privilege Level | 3925
- Output Fields | 3925
- Sample Output | 3925
- Sample Output | 3926
- Release Information | 3926

Syntax

```
clear security flow session session-identifier
                   session-identifier
<node (          node-id         | all | local | primary)>
```

Description

Clear the session with the specific identifier.

Options

- *session-identifier* —Number from 1 through 4,294,967,295 that identifies the security session.
- node—(Optional) For chassis cluster configurations, clear the specified session on a specific node (device) in the cluster.
  - *node-id*—Identification number of the node. It can be 0 or 1.
  - all—Clear all nodes.
  - local—Clear the local node.
  - primary—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session session-identifier 1

```
user@host> clear security flow session session-identifier 1
0 active sessions cleared
```

Sample Output

clear security flow session session-identifier 1 node 0

```
user@host> clear security flow session session-identifier 1 node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 8.5; `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

show security flow session session-identifier | 24283

clear security flow session source-port

IN THIS SECTION

- Syntax | 3927
- Description | 3927
- Options | 3927
- Required Privilege Level | 3927
- Output Fields | 3927
- Sample Output | 3927
- Sample Output | 3928
- Release Information | 3928

Syntax

```
clear security flow session source-port
                   source-port-number
<node (          node-id         | all | local | primary)>
```

Description

Clear each session that uses the specified source port.

Options

- *source-port-number* —Number that identifies the source port.

  Range: 1 through 65,535

- node—(Optional) For chassis cluster configurations, clear sessions on the specified source port on a specific node (device) in the cluster.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - all —Clear all nodes.
  - local —Clear the local node.
  - primary—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session source-port 1

```
user@host> clear security flow session source-port 1
node0:
```

```
-----------------------------------------------------------------------
0 active sessions cleared
node1:
-----------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session source-port 1 node 0

```
user@host> clear security flow session source-port 1 node 0
node0:
-----------------------------------------------------------------------
0 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 8.5; `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

show security flow session source-port | 24295

clear security flow session source-prefix

IN THIS SECTION

- Syntax | 3929
- Description | 3929
- Options | 3929
- Required Privilege Level | 3929
- Output Fields | 3929
- Sample Output | 3930

Sample Output | 3930
Release Information | 3930

Syntax

```
clear security flow session source-prefix
                   source-prefix-number
<node (          node-id         | all | local | primary)>
```

Description

Clear sessions that match the source prefix.

Options

- *source-prefix-number*—Source IP prefix or address.
- node—(Optional) For chassis cluster configurations, clear security sessions matching the source prefix on a specific node (device) in the cluster.
  - *node-id*—Identification number of the node. It can be 0 or 1.
  - all—Clear all nodes.
  - local—Clear the local node.
  - primary—Clear the primary node.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session source-prefix 10.0.0.1

```
user@host> clear security flow session source-prefix 10.0.0.1
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

clear security flow session source-prefix 10::10

```
user@host> clear security flow session source-prefix 10::10
1 active sessions cleared
```

Sample Output

clear security flow session source-prefix 10.0.0.1 node 0

```
user@host> clear security flow session source-prefix 10.0.0.1 node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```

Release Information

Command introduced in Release 8.5 of Junos OS.

The `node` options added in Release 9.0 of Junos OS.

Support for IPv6 addresses added in Release 10.2 of Junos OS.

RELATED DOCUMENTATION

show security flow session source-prefix | 24300

clear security flow session tenant

IN THIS SECTION

- Syntax | 3931
- Description | 3931
- Options | 3931
- Required Privilege Level | 3931
- Output Fields | 3932
- Sample Output | 3932
- Release Information | 3932

Syntax

```
clear security flow session tenant (tenant-name| all)
```

Description

Clears the information about the currently active security flow sessions of the tenant systems on the device. You can either clear the currently active security flow sessions for a specific tenant system or for all the tenant systems.

Options

tenant-name Name of the tenant system.

all Clears the security flow session information for all the tenant systems.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session tenant T1

```
user@host> clear security flow session tenant T1
0 active sessions cleared
1 active sessions cleared
0 active sessions cleared
0 active sessions cleared
```

clear security flow session tenant all

```
user@host> clear security flow session tenant all
0 active sessions cleared
2 active sessions cleared
0 active sessions cleared
0 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

show security flow session tenant | 24327

clear security flow session tunnel

IN THIS SECTION

- Syntax | 3933
- Description | 3933
- Options | 3933
- Required Privilege Level | 3935
- Output Fields | 3935
- Sample Output | 3935
- Release Information | 3935

Syntax

```
clear security flow session tunnel [Enter] | tunnel-session-parameters
```

Description

Clear all active tunnel sessions by entering the command without parameters, or clear the tunnel session whose session parameters are specified.

Options

application Name of the application or application set.

- `dns`—Domain Name System
- `ftp`—File Transfer Protocol
- `ignore`—Ignore application type
- `mgcp-ca`—Media Gateway Control Protocol with Call Agent
- `mgcp-ua`—MGCP with User Agent

- `ms-rpc`—Microsoft RPC
- `pptp`—Point-to-Point Tunneling Protocol
- `q931`—ISDN connection control protocol
- `ras`—RAS
- `realaudio`—RealAudio
- `rsh`—UNIX remote shell services
- `rtsp`—Real-Time Streaming Protocol
- `sccp`—Skinny Client Control Protocol
- `sip`—Session Initiation Protocol
- `sqlnet-v2`—Oracle SQLNET
- `sun-rpc`—Sun Microsystems RPC
- `talk`—TALK program
- `tftp`—Trivial File Transfer Protocol

To display the supported applications list on an SRX Series device, enter the following command from configuration mode:

```
user@host# show groups junos-default applications
```

conn-tag

session connection identifiers are in the following range:

- **Range:** 0 through 4294967295

destination-port

Destination port.

- **Range:** 1 through 65535

destination-prefix

Destination IP prefix or address.

family

Protocol family:

- inet—Clear IPv4 sessions
- inet6–Clear IPv6 sessions

interface Name of incoming or outgoing interface.

protocol IP protocol number.

source-port Source port.

- **Range:** 1 through 65535.

source-prefix Source IP prefix or address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

The same output is displayed when you enter an option for this command.

Sample Output

command-name

```
user@host> clear security flow session tunnel
number active sessions cleared
```

Release Information

Command introduced in Junos OS Release 8.5.

clear security forward-options mirror filter

IN THIS SECTION

Syntax | 3936

- Description | 3936
- Options | 3936
- Required Privilege Level | 3936
- Release Information | 3936

Syntax

```
clear security forward-options mirror-filter (all | filter-name)
```

Description

Clear statistics about configured mirror filters.

Options

all	Clear statistics for all configured mirror filters.
filter-name	Clear statistics for the specified mirror filter.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

mirror-filter (Security Forwarding Options) | 10768

show security forwarding-options mirror-filter | 24363

clear security group-vpn member group

IN THIS SECTION

- Syntax | 3937
- Description | 3937
- Options | 3937
- Required Privilege Level | 3937
- Output Fields | 3938
- Release Information | 3938

Syntax

```
clear security group-vpn member group <vpn vpn-name> <group-id group-id>
```

Description

Clear all current information for IKE, TEK, and KEK SAs. Group VPNv2 is supported on MX Series routers, SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewalls.

Options

none	Clear SA information for all groups.
vpn *vpn-name*	(Optional) Clear SA information for the specified VPN name.
group-id *group-id*	(Optional) Clear SA information for the specified group identifier.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Group VPNv2 Overview

clear security group-vpn member ike security-associations

IN THIS SECTION

- Syntax | 3938
- Description | 3938
- Options | 3939
- Required Privilege Level | 3939
- Output Fields | 3939
- Release Information | 3939

Syntax

```
clear security group-vpn member ike security-associations [index SA-index] [peer-ipaddress]
```

Description

Clear IKE security association (SA) for a group member. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewalls.

Options

- none—Clear all IKE SAs for the group member.
- `index`—(Optional) Clear the IKE SA with this index number.
- *peer-ipaddress*—(Optional) Clear the IKE SA with this peer.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security group-vpn member ike security-associations | 24366
Group VPNv2 Overview

clear security group-vpn member ipsec security-associations

IN THIS SECTION

Syntax | 3940
Description | 3940
Options | 3940
Required Privilege Level | 3940
Output Fields | 3940
Release Information | 3940

Syntax

```
clear security group-vpn member ipsec security-associations [index SA-index]
```

Description

Clear group VPN SA for a group member. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewalls.

Options

- none—Clear all group VPN SAs for the group member.
- `index`—(Optional) Clear the group VPN SA with this index number.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security group-vpn member ipsec security-associations | 24375

Group VPNv2 Overview

clear security group-vpn member ipsec security-associations statistics

IN THIS SECTION

- Syntax | 3941
- Description | 3941
- Options | 3941
- Required Privilege Level | 3941
- Output Fields | 3941
- Release Information | 3942

Syntax

```
clear security group-vpn member ipsec security-associations statistics <group-id group-id>
```

Description

Clear IPsec SA statistics. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewalls.

Options

none	Clear IPsec SA statistics for all groups.
group-id *group-id*	(Optional) Clear IPsec SA statistics for the specified group identifier.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Group VPNv2 Overview

clear security group-vpn member ipsec statistics

IN THIS SECTION

- Syntax | 3942
- Description | 3942
- Options | 3942
- Required Privilege Level | 3943
- Output Fields | 3943
- Release Information | 3943

Syntax

```
clear security group-vpn member ipsec statistics <index index>
```

Description

Clear IPsec statistics. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewalls.

Options

none Clear IPsec statistics for all groups.

index *index* (Optional) Clear the IPsec statistics for the SA with this index number.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Group VPNv2 Overview

clear security group-vpn member kek security-associations

IN THIS SECTION

- Syntax | 3943
- Description | 3944
- Options | 3944
- Required Privilege Level | 3944
- Output Fields | 3944
- Release Information | 3944

Syntax

```
clear security group-vpn member kek security-assocations  [index SA-index]
```

Description

Clear key encryption key (KEK) SAs for a group member.

Options

- none—Clear all KEK SAs for the group member.
- `index`—(Optional) Clear the KEK SA with this index number.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

- show security group-vpn member kek security-associations
- Group VPNv2 Overview

clear security group-vpn server

IN THIS SECTION

- Syntax | 3945
- Description | 3945
- Options | 3945
- Required Privilege Level | 3945

Output Fields | 3945

Syntax

```
clear security group-vpn server [group group-name | group-id group-id] [now]
```

Description

Clear active members for a specified group. If no options are specified, members are cleared from all groups. After this command is issued, members will need to reregister. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewalls.

An IKE SA can be used by a group member to register to multiple groups. When you clear members for a specified group, all existing IKE SAs that could be used to register to the group are also cleared.

Options

- none—All members are cleared from all groups.
- `group`—(Optional) Clear members and SAs for the specified group name.
- `group-id`—(Optional) Clear members and SAs for the specified group identifier.
- `now`—(Optional) Immediately clear all group-related information.

Required Privilege Level

clear

Output Fields

If there is a problem with the command, one of the following messages appears:

- Group does not exist
- Group is in the process of deletion
- Error in clear members

- Warning Message; Fail to push delete to members as server-member-communication is not configured.

RELATED DOCUMENTATION

show security group-vpn server registered-members | 24409

Group VPNv2 Overview

clear security group-vpn server server-cluster statistics

IN THIS SECTION

- Syntax | 3946
- Description | 3946
- Options | 3946
- Required Privilege Level | 3947
- Output Fields | 3947
- Release Information | 3947

Syntax

```
clear security group-vpn server server-cluster statistics <group group-name> <group-id group-id>
```

Description

Clear Group VPNv2 server cluster statistics. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewalls.

Options

none Clear Group VPNv2 server cluster statistics for all groups.

group *group-name* (Optional) Clear Group VPNv2 server cluster statistics for the specified group name.

group-id *group-id* (Optional) Clear Group VPNv2 server cluster statistics for the specified group identifier.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Group VPNv2 Overview

Understanding Group VPNv2 Server Clusters

clear security group-vpn server statistics

IN THIS SECTION

- Syntax | 3948
- Description | 3948
- Options | 3948
- Required Privilege Level | 3948
- Output Fields | 3948
- Release Information | 3948

Syntax

```
clear security group-vpn server statistics <group group-name> <group-id group-id>
```

Description

Clear group statistics. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewalls.

Options

none	Clear statistics for all groups.
group *group-name*	(Optional) Clear statistics for the specified group name.
group-id *group-id*	(Optional) Clear statistics for the specified group identifier.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

show security group-vpn server statistics | **24416**

Group VPNv2 Overview

clear security gtp counters

IN THIS SECTION

- Syntax | 3949
- Description | 3949
- Options | 3949
- Required Privilege Level | 3950
- Output Fields | 3950
- Sample Output | 3950
- Release Information | 3951

Syntax

```
clear security gtp counters <all | data-packet | error | ha | message <message-name> | packet |
request | tunnel | path-rate-limit>
```

Description

Use this command to clear all or specified GTP counters on the device.

From Junos OS Release 20.4R1 onwards, `clear security gprs gtp counters` is replaced by `clear security gtp counters`.

Options

- **all**—Clears all GTP counters.
- **data-packet**—Clears GTP-U data packet counters.
- **error**—Clears GTP error counters.
- **ha**—Clears GTP HA counters.
- **message *message-name***—Clears GTP message counters.
- **packet**—Clears GTP packet counters.

- **request**—Clears GTP request counters.
- **tunnel**—Clears GTP tunnel counters.
- **path-rate-limit**—Clears path-rate-limit counters.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security gtp counters all

```
user@host> clear security gtp counters all
All GTP counters have been cleared
```

clear security gtp counters data-packet

```
user@host> clear security gtp counters data-packet
All GTP-U data packet counters have been cleared
```

clear security gtp counters error

```
user@host> clear security gtp counters error
GTP error counter has been cleared
```

clear security gtp counters ha

```
user@host> clear security gtp counters ha
GTP HA counter has been cleared
```

clear security gtp counters message v0-create-aa-pdp-req

```
user@host> clear security gtp counters message v0-create-aa-pdp-req
GTPv0 create AA PDP request message counter has been cleared
```

clear security gtp counters packet

```
user@host> clear security gtp counters packet
GTP packet counter has been cleared
```

clear security gtp counters request

```
user@host> clear security gtp counters request
GTP request counter has been cleared
```

clear security gtp counters tunnel

```
user@host> clear security gtp counters tunnel
GTP tunnel counter has been cleared
```

clear security gtp counters path-rate-limit

```
user@host> clear security gtp counters path-rate-limit
GTP path-rate-limit counter has been cleared
```

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

show security gtp sctp counters | 24465

clear security gtp sctp association

IN THIS SECTION

- Syntax | 3952
- Description | 3952
- Options | 3952
- Required Privilege Level | 3953
- Sample Output | 3953
- Release Information | 3954

Syntax

```
clear security gtp sctp association
<all>
<destination-ip>
<desitnation-port>
<guid>
<init>
<source-ip>
<source-port>
```

Description

Use this command to clear the Stream Control Transmission Protocol (SCTP) association.

Options

none Clears the live SCTP associations.

all Clears all the SCTP associations, both initiated and live. All SCTP traffic is blocked while the associations are being cleared, which can take up to one minute.

destination-ip Clears the destination IP SCTP association.

destination-port Clears the destination port SCTP association.

guid Clears the globally unique identifier SCTP association.

init Clears the initiated SCTP associations.

source-ip Clears the source IP address SCTP association.

source-port Clears the source port SCTP association.

Required Privilege Level

clear

Sample Output

clear security gtp sctp association

```
user@host> clear security gtp sctp association


Clear Association Information for FPC: 2    PIC: 0
Cleared matched SCTP association information:
Has cleared matched association: 0


Clear Association Information for FPC: 2    PIC: 1
Cleared matched SCTP association information:
Has cleared matched association: 9


Clear Association Information for FPC: 2    PIC: 2
Cleared matched SCTP association information:
Has cleared matched association: 8


Clear Association Information for FPC: 2    PIC: 3
Cleared matched SCTP association information:
Has cleared matched association: 10
```

```
Clear Association Information for FPC: 5    PIC: 0
Cleared matched SCTP association information:
Has cleared matched association: 7


Clear Association Information for FPC: 5    PIC: 1
Cleared matched SCTP association information:
Has cleared matched association: 6
```

Release Information

Command introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

show security gtp sctp association | 24461

clear security gtp sctp counters

IN THIS SECTION

- Syntax | 3954
- Description | 3955
- Options | 3955
- Required Privilege Level | 3955
- Output Fields | 3955
- Sample Output | 3955
- Release Information | 3955

Syntax

```
clear security gtp sctp counters
```

Description

Use this command to clear the statistics of the dropped Stream Control Transmission Protocol (SCTP) counters.

Options

none—Clears all dropped SCTP counters.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security gprs sctp counters

```
user@host> clear security gtp sctp counters
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security gtp sctp counters | 24465

clear security host-vpn security-associations

IN THIS SECTION

- Syntax | 3956
- Description | 3956
- Options | 3956
- Required Privilege Level | 3956
- Sample Output | 3957
- Release Information | 3957

Syntax

```
clear security host-vpn security-associations connection-name
```

Description

Clear IPsec security association information for the security association specified.

Options

connection-name Specify the protection information to be cleared by connection name. If no connection name is specified, information for all security associations is cleared.

Required Privilege Level

view

Sample Output

clear security host-vpn security-associations connection-name

```
user@host> clear security host-vpn security-associations leftT1
terminate of sa:leftT1 complete
```

clear security host-vpn security-associations

```
user@host> clear security host-vpn security-associations
terminate without IKE-SA name unsupported
```

Release Information

Statement introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- show security host-vpn security-associations | 24473
- host-vpn | 8174

clear security idp

IN THIS SECTION

- Syntax | 3958
- Description | 3958
- Required Privilege Level | 3958
- Output Fields | 3958
- Sample Output | 3958
- Release Information | 3959

Syntax

```
clear security idp
(application-identification | application-statistics | attack | counters | status)
```

Description

Clear the following IDP information:

- `application-identification`—Clear IDP application identification data.
- `application-statistics`—Clear IDP application statistics.
- `attack`—Clear IDP attack data
- `counters`—Clear IDP counters
- `status`—Clear IDP Status

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security idp status

```
user@host> clear security idp status
State of IDP: 2-default, Up since: 2010-02-04 13:37:16 UTC (17:13:45 ago)

Packets/second: 0 Peak: 0 @ 2010-02-05 06:49:51 UTC
KBits/second: 0 Peak: 0 @ 2010-02-05 06:49:51 UTC
Latency (microseconds): [min: 0] [max: 0] [avg: 0]

Packet Statistics:
 [ICMP: 0] [TCP: 0] [UDP: 0] [Other: 0]
```

```
Flow Statistics:
  ICMP: [Current: 0] [Max: 0 @ 2010-02-05 06:49:51 UTC]
  TCP: [Current: 0] [Max: 0 @ 2010-02-05 06:49:51 UTC]
  UDP: [Current: 0] [Max: 0 @ 2010-02-05 06:49:51 UTC]
  Other: [Current: 0] [Max: 0 @ 2010-02-05 06:49:51 UTC]

Session Statistics:
 [ICMP: 0] [TCP: 0] [UDP: 0] [Other: 0]
  Policy Name: sample
  Running Detector Version: 10.4.160091104
```

Release Information

Command introduced in Junos OS Release 10.1.

clear security idp attack table

IN THIS SECTION

- Syntax | 3959
- Description | 3960
- Options | 3960
- Required Privilege Level | 3960
- Output Fields | 3960
- Release Information | 3960

Syntax

```
clear security idp attack table
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Clears the details of the IDP attack table.

Options

none	Clears the details of the IDP attack table.
logical-system ***logical-system-name***	(Optional) Clears the details of the IDP attack table for a specific logical system.
logical-system all	(Optional) Clears the details of the IDP attack table for all logical systems.
tenant ***tenant-name***	(Optional) Clears the details of the IDP attack table for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp attack table | **24500**

clear security idp counters action

IN THIS SECTION

- Syntax | 3961
- Description | 3961
- Options | 3961
- Required Privilege Level | 3961
- Output Fields | 3962
- Release Information | 3962

Syntax

```
clear security idp counters action
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the action counter values.

Options

none	Resets all the action counter values.
logical-system ***logical-system-name***	(Optional) Resets all the action counter values for a specific logical system.
logical-system all	(Optional) Resets all the action counter values for all logical systems.
tenant ***tenant-name***	(Optional) Resets all the action counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters action | 24504

clear security idp counters application-identification

IN THIS SECTION

- Syntax | 3962
- Description | 3963
- Options | 3963
- Required Privilege Level | 3963
- Output Fields | 3963
- Release Information | 3963

Syntax

```
clear security idp counters application-identification
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the application identification counter values.

Options

none	Resets all the application identification counter values.
logical-system ***logical-system-name***	(Optional) Resets all the application identification counter values for a specific logical system.
logical-system all	(Optional) Resets all the application identification counter values for all logical systems.
tenant ***tenant-name***	(Optional) Resets all the application identification counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

application-identification

show security idp counters application-identification | 24507

clear security idp counters dfa

IN THIS SECTION

- Syntax | 3964
- Description | 3964
- Options | 3964
- Required Privilege Level | 3964
- Output Fields | 3965
- Release Information | 3965

Syntax

```
clear security idp counters dfa
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the DFA counter values.

Options

none	Resets all the DFA counter values.
logical-system *logical-system-name*	(Optional) Resets all the DFA counter values for a specific logical system.
logical-system all	(Optional) Resets all the DFA counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the DFA counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters dfa | 24514

clear security idp counters flow

IN THIS SECTION

- Syntax | 3965
- Description | 3965
- Required Privilege Level | 3966
- Output Fields | 3966
- Release Information | 3966

Syntax

```
clear security idp counters flow
clear security idp counters flow logical-system logical-system
```

Description

Reset all the IDP flow-related counter values.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

Command introduced for user logical systems in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

flow (Security IDP)

show security idp counters flow | 24517

clear security idp counters http-decoder

IN THIS SECTION

- Syntax | 3967
- Description | 3967
- Options | 3967
- Required Privilege Level | 3967
- Output Fields | 3967
- Release Information | 3967

Syntax

```
clear security idp counters http-decoder
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the HTTP decoder counter values.

Options

none	Resets all the HTTP decoder counter values.
logical-system *logical-system-name*	(Optional) Resets all the HTTP decoder counter values for a specific logical system.
logical-system all	(Optional) Resets all the HTTP decoder counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the HTTP decoder counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 11.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters http-decoder | 24534

clear security idp counters ips

IN THIS SECTION

- Syntax | 3968
- Description | 3968
- Options | 3968
- Required Privilege Level | 3969
- Output Fields | 3969
- Release Information | 3969

Syntax

```
clear security idp counters ips
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the IPS counter values.

Options

none	Resets all the IPS counter values.
logical-system *logical-system-name*	(Optional) Resets all the IPS counter values for a specific logical system.
logical-system all	(Optional) Resets all the IPS counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the IPS counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

ips

show security idp counters ips | 24537

clear security idp counters log

IN THIS SECTION

- Syntax | 3970
- Description | 3970
- Options | 3970
- Required Privilege Level | 3970
- Output Fields | 3970
- Release Information | 3970

Syntax

```
clear security idp counters log
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the IDP log counter values.

Options

none	Resets all the IDP log counter values.
logical-system *logical-system-name*	(Optional) Resets all the IDP log counter values for a specific logical system.
logical-system all	(Optional) Resets all the IDP log counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the IDP log counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

event-rate

show security idp counters log | 24545

clear security idp counters memory

IN THIS SECTION

- Syntax | 3971
- Description | 3971
- Options | 3971
- Required Privilege Level | 3972
- Output Fields | 3972
- Release Information | 3972

Syntax

```
clear security idp counters memory
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the memory allocation and reallocation counter values.

Options

none	Resets all the memory allocation and reallocation counter values.
logical-system ***logical-system-name***	(Optional) Resets all the memory allocation and reallocation counter values for a specific logical system.
logical-system all	(Optional) Resets all the memory allocation and reallocation counter values for all logical systems.

tenant ***tenant-name*** (Optional) Resets all the memory allocation and reallocation counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters memory | 24552

clear security idp counters packet

IN THIS SECTION

- Syntax | 3973
- Description | 3973
- Options | 3973
- Required Privilege Level | 3973
- Output Fields | 3973
- Release Information | 3973

Syntax

```
clear security idp counters packet
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the IDP packet counter values.

Options

none	Resets all the IDP packet counter values.
logical-system *logical-system-name*	(Optional) Resets all the IDP packet counter values for a specific logical system.
logical-system all	(Optional) Resets all the IDP packet counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the IDP packet counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters packet | 24556

clear security idp counters packet-log

IN THIS SECTION

- Syntax | 3974
- Description | 3974
- Options | 3974
- Required Privilege Level | 3975
- Output Fields | 3975
- Release Information | 3975

Syntax

```
clear security idp counters packet-log
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the IDP counters value for packet log.

Options

none	Resets all the IDP counters value for packet log.
logical-system *logical-system-name*	(Optional) Resets all the IDP counters value for packet log for a specific logical system.
logical-system all	(Optional) Resets all the IDP counters value for packet log for all logical systems.

tenant ***tenant-name*** (Optional) Resets all the IDP counters value for packet log for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters packet-log | 24567

clear security idp counters pdf-decoder

IN THIS SECTION

- Syntax | 3976
- Description | 3976
- Options | 3976
- Required Privilege Level | 3976
- Output Fields | 3976
- Release Information | 3976

Syntax

```
clear security idp counters pdf-decoder
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the PDF-Decoder enabled sessions.

Options

none	Resets all the PDF-Decoder enabled sessions.
logical-system *logical-system-name*	(Optional) Resets all the PDF-Decoder enabled sessions for a specific logical system.
logical-system all	(Optional) Resets all the PDF-Decoder enabled sessions for all logical systems.
tenant *tenant-name*	(Optional) Resets all the PDF-Decoder enabled sessions for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters pdf-decoder | **24581**

clear security idp counters policy-manager

IN THIS SECTION

- Syntax | 3977
- Description | 3977
- Options | 3977
- Required Privilege Level | 3978
- Output Fields | 3978
- Release Information | 3978

Syntax

```
clear security idp counters policy-manager
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the IDP policies counter values.

Options

none	Resets all the IDP policies counter values.
logical-system ***logical-system-name***	(Optional) Resets all the IDP policies counter values for a specific logical system.
logical-system all	(Optional) Resets all the IDP policies counter values for all logical systems.
tenant ***tenant-name***	(Optional) Resets all the IDP policies counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters policy-manager | 24585

clear security idp counters ssl-inspection

IN THIS SECTION

- Syntax | 3979
- Description | 3979
- Options | 3979
- Required Privilege Level | 3979
- Output Fields | 3979
- Release Information | 3979

Syntax

```
clear security idp counters ssl-inspection
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the session encryption and decryption values for IDP counters.

Options

none	Resets all the session encryption and decryption values for IDP counters.
logical-system ***logical-system-name***	(Optional) Resets all the session encryption and decryption values for IDP counters for a specific logical system.
logical-system all	(Optional) Resets all the session encryption and decryption values for IDP counters for all logical systems.
tenant ***tenant-name***	(Optional) Resets all the session encryption and decryption values for IDP counters for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters ssl-inspection | 24588

clear security idp counters tcp-reassembler

IN THIS SECTION

- Syntax | 3980
- Description | 3980
- Options | 3980
- Required Privilege Level | 3981
- Output Fields | 3981
- Release Information | 3981

Syntax

```
clear security idp counters tcp-reassembler
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the TCP reassembler counter values.

Options

none	Resets all the TCP reassembler counter values.
logical-system *logical-system-name*	(Optional) Resets all the TCP reassembler counter values for a specific logical system.
logical-system all	(Optional) Resets all the TCP reassembler counter values for all logical systems.

tenant ***tenant-name*** (Optional) Resets all the TCP reassembler counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

re-assembler

show security idp counters tcp-reassembler | 24592

clear security idp ssl-inspection session-id-cache

IN THIS SECTION

- Syntax | 3982
- Description | 3982
- Required Privilege Level | 3982
- Output Fields | 3982
- Sample Output | 3982
- Release Information | 3982

Syntax

```
clear security idp ssl-inspection session-id-cache
```

Description

Clear all the entries stored in the SSL session ID cache.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security idp ssl-inspection session-id-cache

```
user@host> clear security idp ssl-inspection session-id-cache
Total SSL session cache entries cleared : 2
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show security idp ssl-inspection session-id-cache | 24621

clear security ike active-peer aaa-username

IN THIS SECTION

- Syntax | 3983
- Description | 3983
- Options | 3983
- Required Privilege Level | 3983
- Output Fields | 3983
- Release Information | 3984

Syntax

```
clear security ike active-peer aaa-username username
```

Description

Clears all IKE security association (SA) of a user.

Options

- none—Clear all IKE SAs for the user.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 22.3R1.

clear security ike peers

IN THIS SECTION

- Syntax | 3984
- Description | 3984
- Options | 3984
- Required Privilege Level | 3985
- Output Fields | 3985
- Release Information | 3985

Syntax

```
clear security ike peers <backoff | blocked | failed | in-progress> role <initiator | responder>
index <value> remote-ip-address <value> local-ip-address <value> gateway-name <value>
```

Description

Clear IKE peers that are backed off, blocked, failed or in progress.

Options

backoff Clear the remote peer entries that are temporarily denied from initiating a new IKE negotiation.

blocked Clear the remote peer entries for which the device blocked the IKE negotiations.

failed Clear the remote peer entries for which the IKE negotiation failed.

in-progress	Clear the remote peer entries for which the IKE negotiation is in progress.
gateway-name *value*	(Optional) Enter the IKE gateway name.
index	(Optional) Enter the index number that needs to be cleared.
local-ip-address	(Optional) Enter the local IP address.
remote-ip-address	(Optional) Enter the remote IP address.
role	(Optional) Enter `initiator` or `responder`.

Required Privilege Level

clear

Output Fields

The command produces no output.

Release Information

Command introduced in Junos OS Release 23.4R1

RELATED DOCUMENTATION

IKE Protection from DDoS Attacks

show security ike peers | 24640

clear security ike peers statistics

IN THIS SECTION

- Syntax | 3986
- Description | 3986
- Options | 3986

- Required Privilege Level | 3986
- Output Fields | 3986
- Release Information | 3986

Syntax

```
clear security ike peers statistics <gateway-name value>
```

Description

Clear the global or per IKE gateway level statistics. It clears *Failed* and *Blocked* field statistics.

Options

gateway-name *value* (Optional) Enter the peer IKE gateway name as the value.

Required Privilege Level

clear

Output Fields

The command produces no output.

Release Information

Command introduced in Junos OS Release 23.4R1

RELATED DOCUMENTATION

IKE Protection from DDoS Attacks

show security ike peers | 24640

clear security ike respond-bad-spi-count

IN THIS SECTION

- Syntax | 3987
- Description | 3987
- Options | 3987
- Required Privilege Level | 3987
- Output Fields | 3987
- Release Information | 3988

Syntax

```
clear security ike respond-bad-spi-count <gateway-name>
```

Description

Clear information about invalid Internet Key Exchange (IKE) security parameter index (SPI) counters.

Options

- none—Clear all invalid SPI counters.
- *gateway-name* —(Optional) Clear the invalid SPI counters for the given gateway.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

ike (Security) | 8432

clear security ike security-associations

IN THIS SECTION

- Syntax | 3988
- Description | 3989
- Options | 3989
- Required Privilege Level | 3990
- Output Fields | 3990
- Release Information | 3990

Syntax

```
clear security ike security-associations
<peer-address>
<family (inet  | inet6)>
<fpc slot-number>
<index SA-index-number>
<kmd-instance (all | kmd-instance-name)>
<pic slot-number>
<node-local>
<port port-number>
<sa-type shortcut>
<srg-id id-number>
<ha-link-encryption>
```

Description

Clear information about the current Internet Key Exchange security associations (IKE SAs). For IKEv2, the device clears the information about the IKE SAs and the associated IPSec SA.

Options

- none—Clear all IKE SAs.
- *peer-address* —(Optional) Clear IKE SAs for the destination peer at this IP address.
- family—(Optional) Clear IKE SAs by family.
  - inet—IPv4 address family.
  - inet6—IPv6 address family.
- fpc *slot-number* —Specific to SRX Series Firewalls. Clear information about existing IKE SAs in this Flexible PIC Concentrator (FPC) slot.
- index *SA-index-number* —(Optional) Clear the IKE SA with this index number.
- kmd-instance—Clear information about existing IKE SAs in the key management process (the daemon, which in this case is KMD) identified by FPC *slot-number* and PIC *slot-number*. Specific to SRX Series Firewalls.
  - all—All KMD instances running on the Services Processing Unit (SPU).
  - *kmd-instance-name*—Name of the KMD instance running on the SPU.
- node-local—(Optional) Clear information about IKE SAs for node-local tunnels in a Multinode High Availability setup.
- pic *slot-number* —Clear information about existing IKE SAs in this PIC slot. Specific to SRX Series Firewalls.
- port *port-number*—(Optional) Port number of SA (1 through 65,535).
- sa-type shortcut—(Optional) It's applicable for ADVPN. Clear information about IKE SAs by type shortcut.
- ha-link-encryption—(Optional) Clear information about the current IKE SAs for high availability (HA) link tunnel only. When you enable High Availability feature, you cannot delete customer tunnels on the backup node.
- srg-id—(Optional) Clear information related to a specific services redundancy group (SRG) in a Multinode High Availability setup.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 8.5. The `fpc`, `pic`, and `kmd-instance` options added in Junos OS Release 9.3. The `port` option added in Junos OS Release 10.0. The `family` option added in Junos OS Release 11.1.

Support for the `ha-link-encryption` option added in Junos OS Release 20.4R1.

Support for the `srg-id` option added in Junos OS Release 22.4R1.

Support for the `node-local` option added in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

> https://www.juniper.net/documentation/us/en/software/junos/vpn-ipsec/topics/ref/command/show-security-ike-security-associations.html

clear security ike stats

IN THIS SECTION

- Syntax | 3991
- Description | 3991
- Required Privilege Level | 3991
- Sample Output | 3991
- Release Information | 3993

Syntax

```
clear security ike stats
```

Description

Clears the global IKE statistics.

In releases before Junos OS Release 22.4R3, in a Chassis Cluster mode, when you run the command on the secondary node, you'll see error message `error: IKE-Config-Management not responding to management requests`.

Required Privilege Level

clear

Sample Output

clear security ike stats

```
user@host> clear security ike stats
```

command-name

The `clear security ike stats` command does not display any output. To view the IKE statistics, run the `show security ike stats detail` command.

show security ike stats detail

```
user@host> show security ike stats detail
Total IKE SA and Tunnel Count Statistics:
  Number of IKE SAs: 2             Number of IPsec Tunnels: 2

IKE_SA_INIT exchange stats:
 Initiator stats:                                  Responder stats:
  Request Out             : 0                       Request In               : 0
  Response In             : 0                       Response Out             : 0
  Invalid KE Payload In   : 0                       Invalid KE Payload Out   : 0
```

```
  No Proposal Chosen In   : 0                        No Proposal Chosen Out   : 0
  Cookie Request In       : 0                        Cookie Request Out       : 0
  Cookie Response Out     : 0                        Cookie Response In       : 0
  Res Invalid IKE SPI     : 0                        Res DH Gen Key Fail      : 0
  Res Verify SA Fail      : 0                        Res Invalid DH Group Conf: 0
  Res IKE SA Fill Fail    : 0                        Res Get CAs Fail         : 0
  Res Verify DH Group Fail: 0                        Res Get VID Fail         : 0
  Res DH Compute Key Fail : 0                        Res DH Compute Key Fail  : 0

IKE_AUTH exchange stats:
 Initiator stats:                                  Responder stats:
  Request Out             : 0                        Request In               : 0
  Response In             : 0                        Response Out             : 0
  No Proposal Chosen In   : 0                        No Proposal Chosen Out   : 0
  TS Unacceptable In      : 0                        TS Unacceptable Out      : 0
  Authentication Failed In: 0                        Authentication Failed Out: 0

IKE SA Rekey CREATE_CHILD_SA exchange stats:
 Initiator stats:                                  Responder stats:
  Request Out             : 0                        Request In             : 0
  Response In             : 0                        Response Out           : 0
  No Proposal Chosen In   : 0                        No Proposal Chosen Out : 0
  Invalid KE In           : 0                        Invalid KE Out         : 0
  Res DH Compute Key Fail : 0                        Res DH Compute Key Fail: 0
  Res Verify SA Fail      : 0
  Res Fill IKE SA Fail    : 0
  Res Verify DH Group Fail: 0

IPsec SA Rekey CREATE_CHILD_SA exchange stats:
 Initiator stats:                                  Responder stats:
  Request Out             : 0                        Request In             : 0
  Response In             : 0                        Response Out           : 0
  No Proposal Chosen In   : 0                        No Proposal Chosen Out : 0
  Invalid KE In           : 0                        Invalid KE Out         : 0
  TS Unacceptable In      : 0                        TS Unacceptable Out    : 0
  Res DH Compute Key Fail : 0                        Res DH Compute Key Fail: 0
  Res Verify SA Fail      : 0
  Res Verify DH Group Fail: 0
  Res Verify TS Fail      : 0

Total IKE message failure stats:
  Discarded            : 0                           ID error      : 0
  Integrity fail       : 0                           Invalid SPI   : 0
```

```
  Invalid exchange type: 0                        Invalid length: 0
  Disorder             : 0
```

Release Information

Command is introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

Configure the Certificate Expiration Trap

Enable Peer Down and IPsec Tunnel Down Traps

show security ipsec statistics | 24744

clear security ipsec security-associations

IN THIS SECTION

- Syntax | 3993
- Description | 3994
- Options | 3994
- Required Privilege Level | 3995
- Output Fields | 3995
- Release Information | 3995

Syntax

```
clear security ipsec security-associations
<family (inet  | inet6)>
<fpc slot-number>
<index SA-index-number>
<kmd-instance (all | kmd-instance-name)>
<node-local>
```

```
<pic slot-number>
<sa-type shortcut>
<srg-id id-number>
<ha-link-encryption>
```

Description

Clear information about IPsec security associations (SAs).

Options

- none—Clear all IPsec SAs.
- `family`—(Optional) Clear SAs by family.
  - `inet`—IPv4 address family.
  - `inet6`—IPv6 address family.
- `fpc` *`slot-number`* —Clear information about existing IPsec SAs in this Flexible PIC Concentrator (FPC) slot. Specific to SRX Series Firewalls.
- `index` *`SA-index-number`* —(Optional) Clear the IPsec SA with this index number.
- `kmd-instance`—Clear information about existing IPsec SAs in the key management process (the daemon, which in this case is KMD) identified by FPC *`slot-number`* and PIC *`slot-number`* Specific to SRX Series Firewalls.
  - `all`—All KMD instances running on the Services Processing Unit (SPU).
  - *`kmd-instance-name`*—Name of the KMD instance running on the SPU.
- `node-local`—(Optional) Clear information about IPsec SAs for node-local tunnels in a Multinode High Availability setup.
- `pic` *`slot-number`* —Clear information about existing IPsec SAs in this PIC slot. Specific to SRX Series Firewalls.
- `ha-link-encryption`—(Optional) Clear information about IPsec SAs for interchassis link tunnel only. See "ipsec (High Availability)" on page 9301. When you enable High Availability feature, you cannot delete customer tunnels on the backup node.
- `sa-type shortcut`—(Optional) It's applicable for ADVPN. Clear information about IPsec SAs by type `shortcut`.

- `srg-id`—(Optional) Clear statistics related to a specific services redundancy group (SRG) in a Multinode High Availability setup..

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 8.5. The `fpc`, `pic`, and `kmd-instance` options added in Junos OS Release 9.3. The `family` option added in Junos OS Release 11.1.

Support for the `ha-link-encryption` option added in Junos OS Release 20.4R1.

Support for the `srg-id` option added in Junos OS Release 22.4R1.

Support for the `node-local` option added in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

show security ipsec security-associations | 24701

clear security ipsec statistics

IN THIS SECTION

- Syntax | 3996
- Description | 3996
- Options | 3996
- Required Privilege Level | 3996
- Output Fields | 3997
- Release Information | 3997

Syntax

```
clear security ike statistics
<fpc slot-number>
<index SA-index-number>
<kmd-instance (all |           kmd-instance-name        )>
<pic slot-number>
<srg-id id-number>
```

Description

Clear IPsec statistics on the device.

Options

- none—Clear all IPsec statistics.
- `fpc` *`slot-number`* —Specific to SRX Series Firewalls. Clear statistics about existing IPsec security associations (SAs) in this Flexible PIC Concentrator (FPC) slot.
- `index` *`SA-index-number`* —(Optional) Clear the IPsec statistics for the SA with this index number.
- `kmd-instance`—Specific to SRX Series Firewalls. Clear information about existing IKE SAs in the key management process (the daemon, which in this case is KMD) identified by FPC *`slot-number`* and PIC *`slot-number`* .
  - `all`—All KMD instances running on the Services Processing Unit (SPU).
  - *`kmd-instance-name`*—Name of the KMD instance running on the SPU.
- `pic` *`slot-number`* —Specific to SRX Series Firewalls. Clear statistics about existing IPsec SAs in this PIC slot.
- `srg-id` *`id-number`* —Clear statistics related to a specific services redundancy group (SRG) in a Multinode High Availability setup.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 8.5. `fpc and pic` options added in Junos OS Release 9.3. `kmd-instance` option added in Junos OS Release 10.4.

`srg-id` option added in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

show security ipsec statistics | 24744

clear security ipsec tunnel-events-statistics

IN THIS SECTION

- Syntax | 3997
- Description | 3997
- Required Privilege Level | 3998
- Output Fields | 3998
- Release Information | 3998

Syntax

```
clear security ipsec tunnel-events-statistics
```

Description

Clear IPsec tunnel event statistics.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 12.3X48-D10.

RELATED DOCUMENTATION

show security ipsec tunnel-events-statistics | 24761

clear security log

IN THIS SECTION

- Syntax | 3998
- Description | 3999
- Options | 3999
- Required Privilege Level | 4000
- Sample Output | 4000
- Release Information | 4000

Syntax

```
clear security log
<all>
<all-logical-systems-tenants>
<destination-address>
```

```
<destination-port>
<event-id>
<failure>
<file>
<interface-name>
<newer-than>
<older--than>
<policy-name>
<process>
<protocol>
<query>
<report>
<severity>
<source-address>
<source-port>
<stream>
<success>
<username>
```

Description

Delete the event log.

Options

all	Clear all audit event logs stored in the device memory.
all-logical-systems-tenants	Clear all audit event logs for all the logical systems and tenant systems.
destination-address	Clear audit event logs with the specified destination address.
destination-port	Clear audit event logs with the specified destination port.
event-id	Clear audit event logs with the specified event identification number.
failure	Clear failed audit event logs.
file	Clear security logs in binary format.
interface-name	Clear audit event logs with the specified interface.
newer-than	Clear audit event logs newer than the specified date and time.

older-than	Clear audit event logs older than the specified date and time
policy-name	Policy name filter.
process	Clear audit event logs with the specified process that generated the event.
protocol	Clear audit event logs generated through the specified protocol.
query	Clear logs in database.
report	Clear on-box reports for system traffic logs.
severity	Clear audit event logs generated with the specified severity.
source-address	Clear audit event logs with the specified source address.
source-port	Clear audit event logs with the specified source port.
stream	Clear security logs stream info.
success	Clear successful audit event logs.
username	Clear audit event logs generated for the specified user.

Required Privilege Level

clear

Sample Output

clear security log all

```
user@host> clear security log all
7905 security log events cleared
```

Release Information

Command introduced in Junos OS Release 11.2.

`all-logical-systems-tenants` option introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

exclude (Security Log) | 6627

show security log | 24769

clear security log file

IN THIS SECTION

- Syntax | 4001
- Description | 4001
- Required Privilege Level | 4001
- Sample Output | 4002
- Release Information | 4002

Syntax

```
clear security log file
```

Description

Deletes the content of an event mode security log file stored on the device in binary format.

Required Privilege Level

clear

Sample Output

clear security log file

```
user@host> clear security log file
7905 security log events cleared
```

clear security log file logical-system LSYS1

```
user@host> clear security log file logical-system LSYS1
7905 security log events cleared
```

clear security log file tenant TSYS1

```
user@host> clear security log file tenant TSYS1
7905 security log events cleared
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

show security log file | 24773

clear security log stream file

IN THIS SECTION

- Syntax | 4003
- Description | 4003
- Required Privilege Level | 4003

Output Fields | 4003
Sample Output | 4004
Release Information | 4004

Syntax

```
clear security log query
clear security log stream
file <file-name>
```

Description

- `clear security log query`—Clear the content of the database.
- `clear security log stream file`—Clear the content of the current log file.
- `clear security log stream file logical-system LSYS1` —Clear the content of the current log file for logical system.
- `clear security log stream file tenant TSYS1` —Clear the content of the current log file for tenant system.

Required Privilege Level

clear

Output Fields

The following outputs are occurred in two conditions:

- Clear log stream file successfully, when there is log file.
- Clear log stream file error or does not exits, when there is no log file.

Sample Output

clear security log stream file LSYS_f1.bin logical-system LSYS1

```
user@host> clear security log stream file LSYS_f1.bin logical-system LSYS1
Request to clear log stream file succeeds
```

clear security log stream file TSYS_f1.bin tenant-system TSYS1

```
user@host> clear security log stream file TSYS_f1.bin tenant-system TSYS1
Request to clear log stream file succeeds
```

Release Information

Command introduced in Junos OS Release 15.1X49-D70.

The `logical-system` and `tenant` options are introduced in Junos OS Release 19.2R1.

clear security mka statistics

IN THIS SECTION

- Syntax | 4005
- Description | 4005
- Options | 4005
- Required Privilege Level | 4005
- Sample Output | 4005
- Release Information | 4005

Syntax

```
clear security mka statistics
<interface interface-name>
```

Description

Clear—reset to zero (0)—all MACsec Key Agreement (MKA) protocol statistics.

You are clearing the statistics that are viewed using the `show security mka statistics` when you enter this command.

Options

none	Clear all MKA counters for all interfaces on the switch.
interface *interface-name*	(Optional) Clear MKA traffic counters for the specified interface only.

Required Privilege Level

clear

Sample Output

clear security mka statistics

```
user@switch> clear security mka statistics
```

Release Information

Command introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

show security mka statistics | 24854

show security mka sessions | 24840

clear security mka statistics (MX Series)

IN THIS SECTION

- Syntax | 4006
- Description | 4006
- Options | 4006
- Required Privilege Level | 4007
- Sample Output | 4007
- Release Information | 4007

Syntax

```
clear security mka statistics
<interface interface-name>
```

Description

Clear—reset to zero (0)—all MACsec Key Agreement (MKA) protocol statistics.

You are clearing the statistics that are viewed using the `show security mka statistics` when you enter this command.

Options

none	Clear all MKA counters for all interfaces on the switch.
interface *interface-name*	(Optional) Clear MKA traffic counters for the specified interface only.

Required Privilege Level

clear

Sample Output

clear security mka statistics

```
user@switch> clear security mka statistics
```

Release Information

Command introduced in Junos OS Release 15.1.

clear security nat incoming-table

IN THIS SECTION

- Syntax | 4007
- Description | 4008
- Options | 4008
- Required Privilege Level | 4008
- Output Fields | 4008
- Release Information | 4008

Syntax

```
clear security nat incoming-table
<node (          node-id         | all | local | primary)>
```

Description

Clear Network Address Translation (NAT) incoming table information.

Options

- none—Clear all information NAT incoming table.
- `node`—(Optional) For chassis cluster configurations, clear incoming table information on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Clear all nodes.
  - `local` —Clear the local node.
  - `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

show security nat incoming-table | 24893

clear security nat source persistent-nat-table

IN THIS SECTION

- Syntax | 4009
- Description | 4009
- Options | 4009
- Required Privilege Level | 4010
- Output Fields | 4010
- Release Information | 4010

Syntax

```
clear security nat source persistent-nat-table
( all | interface | internal-ip          ip-address       <internal-port port> | pool poolname )
```

Description

Clear Network Address Translation (NAT) persistent NAT bindings that are in query mode, where all sessions of the binding are gone.

Options

- `all`—Clear all persistent NAT bindings that are in query mode.
- `interface`—Clear persistent NAT bindings that are in query mode for the specified interface.
- `internal-ip` *`ip-address`*—Clear persistent NAT bindings for the specified internal IP address.
- `internal-ip` *`ip-address`* `internal-port` *`port`*—Clear persistent NAT bindings that are in query mode for the specified internal IP address and port.
- `pool`—Clear persistent NAT bindings that are in query mode for the specified source NAT pool.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show security nat source persistent-nat-table | 24910

clear security nat statistics destination pool

IN THIS SECTION

- Syntax | 4010
- Description | 4011
- Options | 4011
- Required Privilege Level | 4011
- Output Fields | 4011
- Release Information | 4011

Syntax

```
clear security nat statistics destination pool
<pool-name>
all
```

Description

Clear the destination NAT pool information.

Options

pool-name—Clear specified destination nat pool information.

all—Clear all destination nat pool information.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show security nat destination pool | 24877

show security nat destination summary | 24889

clear security nat statistics destination rule

IN THIS SECTION

- Syntax | 4012
- Description | 4012
- Options | 4012
- Required Privilege Level | 4012

Output Fields | 4012
Release Information | 4012

Syntax

```
clear security nat statistics destination rule
<rule-name>
alll
```

Description

Clear the destination NAT rule information.

Options

rule-name—Clear specified destination nat rule-set information.

all—Clear all destination nat rule-set information.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show security nat destination rule | 24881
show security nat destination summary | 24889

clear security nat statistics source pool

IN THIS SECTION

- Syntax | 4013
- Description | 4013
- Options | 4013
- Required Privilege Level | 4013
- Output Fields | 4013
- Release Information | 4014

Syntax

```
clear security nat statistics source pool
<pool-name>
all
```

Description

Clear the source NAT statistic pool information.

Options

pool-name—Clear the specified source nat pool information.

all—Clear all source pool information.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show security nat source pool | 24914

show security nat source summary | 24940

clear security nat statistics source rule

IN THIS SECTION

- Syntax | 4014
- Description | 4014
- Options | 4015
- Required Privilege Level | 4015
- Output Fields | 4015
- Release Information | 4015

Syntax

```
clear security nat statistics source rule
<rule-name>
all
```

Description

Clear the source NAT statistic rule-set information.

Options

rule-name—Clear the specified source rule-set information.

all—Clear all source nat rule-set information.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show security nat source summary | 24940

show security nat source rule | 24931

clear security nat statistics static rule

IN THIS SECTION

- Syntax | 4016
- Description | 4016
- Options | 4016
- Required Privilege Level | 4016
- Output Fields | 4016
- Release Information | 4016

Syntax

```
clear security nat statistics static rule
<rule-name>
all
```

Description

Clear the static NAT rule-set information.

Options

rule-name—Clear specified static nat rule-set information.

all—Clear all static nat rule-set information.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show security nat static rule | 24948

clear security pki ca-certificate

IN THIS SECTION

- Syntax | 4017
- Description | 4017
- Options | 4017
- Required Privilege Level | 4017
- Output Fields | 4017
- Sample Output | 4018
- Release Information | 4018

Syntax

```
clear security pki ca-certificate (all | ca-profile ca-profile-name)
```

Description

Delete certificate authority (CA) digital certificates from the router.

Options

all	Delete all CA digital certificates from the router.
ca-profile *ca-profile-name*	Delete the specified CA profile.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security pki ca-certificate all

```
user@host> clear security pki ca-certificate all
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

- request security pki ca-certificate enroll | 14784
- request security pki ca-certificate load | 14787
- show security pki ca-certificate | 24957

clear security pki certificate-request

IN THIS SECTION

- Syntax | 4019
- Description | 4019
- Options | 4019
- Required Privilege Level | 4019
- Output Fields | 4019
- Sample Output | 4019
- Release Information | 4019

Syntax

```
clear security pki certificate-request (all | certificate-id certificate-id-name)
```

Description

Delete manually generated local digital certificate requests from the router.

Options

all	Delete all local digital certificate requests from the router.
certificate-id *certificate-id-name*	Delete the specified local digital certificate and corresponding public/private key pair.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security pki certificate-request all

```
user@host> clear security pki certificate-request all
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

show security pki certificate-request | 24968

clear security pki crl

IN THIS SECTION

- Syntax | 4020
- Description | 4020
- Options | 4020
- Required Privilege Level | 4020
- Output Fields | 4020
- Sample Output | 4021
- Release Information | 4021

Syntax

```
clear security pki crl (all | ca-profile ca-profile-name)
```

Description

Delete certificate revocation lists (CRLs) from the router.

Options

all	Delete all CRLs from the router.
ca-profile *ca-profile-name*	Delete CRLs associated with the specified CA profile.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security pki crl ca-profile all

```
user@host> clear security pki crl ca-profile all
```

Release Information

Command introduced in Junos 8.1

clear security pki key-pair

IN THIS SECTION

- Syntax | 4021
- Description | 4021
- Options | 4022
- Required Privilege Level | 4022
- Output Fields | 4022
- Sample Output | 4022
- Release Information | 4022

Syntax

```
clear security pki key-pair (all | certificate-id certificate-id-name)
```

Description

Clear public key infrastructure (PKI) key pair information for local digital certificates from the router.

Options

all	Delete all local digital certificates, certificate requests, and the corresponding public and private key pairs from the router.
certificate-id *certificate-id-name*	Delete the specified local digital certificate and corresponding public/private key pair.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

command-name

```
user@host> clear security pki key pair
```

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

request security pki local-certificate enroll | 14812

show security pki local-certificate | 24982

clear security pki key-pair (Local Certificate)

IN THIS SECTION

- Syntax | 4023
- Description | 4023
- Options | 4023
- Required Privilege Level | 4023
- Output Fields | 4023
- Release Information | 4024

Syntax

```
clear security pki key-pair (all | certificate-id          certificate-id       )
```

Description

Clear public key infrastructure (PKI) key pair information for local digital certificates on the device.

Options

- `all`—Clear key pair information for all local certificates.
- `certificate-id` *certificate-id* —Clear key pair information for the local certificate with this certificate ID.

Required Privilege Level

clear and security

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

show security pki certificate-request (View) | 24972

clear security pki local-certificate

IN THIS SECTION

- Syntax | 4024
- Description | 4024
- Options | 4025
- Required Privilege Level | 4025
- Output Fields | 4025
- Sample Output | 4025
- Release Information | 4025

Syntax

```
clear security pki local-certificate
<all | certificate-id certificate-id-name | system-generated>
```

Description

Delete local digital certificates, certificate requests, and the corresponding public/private key pairs from the router.

Options

all	(Optional) Delete all local digital certificates, certificate requests, and the corresponding public and private key pairs from the router.
certificate-id *certificate-id-name*	(Optional) Delete the specified local digital certificate and corresponding public and private key pair.
system-generated	(Optional) Auto-generated self-signed certificate.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear security pki local-certificate all

```
user@host> clear security pki local-certificate all
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

request security pki local-certificate enroll | 14812

show security pki local-certificate | 24982

clear security pki local-certificate

IN THIS SECTION

- Syntax | 4026
- Description | 4026
- Options | 4026
- Required Privilege Level | 4027
- Output Fields | 4027
- Sample Output | 4027
- Release Information | 4027

Syntax

```
clear security pki local-certificate
<all | certificate-id certificate-id-name | system-generated>
```

Description

Delete local digital certificates, certificate requests, and the corresponding public/private key pairs from the switch.

Options

all

(Optional) Delete all local digital certificates, certificate requests, and the corresponding public and private key pairs from the router.

> **NOTE**: This option does not delete the automatically generated self-signed certificate or its public/private key pair.

certificate-id ***certificate-id-name***

(Optional) Delete the specified local digital certificate and corresponding public and private key pair.

system-generated (Optional) Delete the automatically generated self-signed certificate.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear security pki local-certificate all

```
user@switch> clear security pki local-certificate all
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Deleting Self-Signed Certificates (CLI Procedure)

clear security pki local-certificate (Device)

IN THIS SECTION

- Syntax | 4028
- Description | 4028
- Options | 4028
- Required Privilege Level | 4028
- Output Fields | 4028

- Sample Output | 4029
- Sample Output | 4029
- Release Information | 4029

Syntax

```
clear security pki local-certificate (all | certificate-id certificate-id | system-generated)
```

Description

Clear public key infrastructure (PKI) information for local digital certificates on the device.

Options

- `all`—Clear information for all the local digital certificates on the device.

  You cannot clear the automatically generated self-signed certificate using `clear security pki local-certificate all` command. To clear the self-signed certificate you need to use `system-generated` as an option.

- `certificate-id` *`certificate-id`* —Clear the specified local digital certificate with this certificate ID.

- `system-generated`—Clear the existing automatically generated self-signed certificate and generate a new self-signed certificate.

Required Privilege Level

clear and security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security pki local-certificate all

```
user@host> clear security pki local-certificate all
```

Sample Output

clear security pki local-certificate system-generated

```
user@host> clear security pki local-certificate system-generated
```

Release Information

Command modified in Junos OS Release 9.1.

Starting in Junos OS Release 20.1R1 on vSRX Virtual Firewall 3.0, you can safeguard the private keys used by PKID and IKED using Microsoft Azure Key Vault hardware security module (HSM) service. You can establish a PKI based VPN tunnel using the keypairs generated at the HSM. The hub `certificate-id` option under certificate-id is not available for configuration after generating HSM key-pair.

Starting in Junos OS Release 20.4R1 on vSRX Virtual Firewall 3.0, you can safeguard the private keys used by PKID and IKED using AWS Key Management Service (KMS). You can establish a PKI based VPN tunnel using the keypairs generated by the KMS. The hub `certificate-id` option under certificate-id is not available for configuration after generating PKI key-pair.

NOTE: You cannot manually re-enroll the local certificates when you re-generate key-pairs, if you are not generating key-pairs during re-enrollment. A warning **HSM does not support auto re-enrollment with new keypair error: configuration check-out failed** is displayed in the output of the `show security pki auto-re-enrollment` command.

Also, when you clear the local certificates using the `run clear security pki local-certificate all` and `run clear security pki key-pair all` commands you will receive a warning **Key pair deleted successfully but still present at HSM. Please purge the keypair from keyvault before re-using the name**.

RELATED DOCUMENTATION

show security pki local-certificate (View) | 24992

request security pki local-certificate generate-self-signed (Security) | 14827

clear security pki node-local key-pair

IN THIS SECTION

- Syntax | 4030
- Description | 4030
- Options | 4030
- Required Privilege Level | 4031
- Output Fields | 4031
- Release Information | 4031

Syntax

```
clear security pki node-local key-pair (all | certificate-id certificate-id)                )
```

Description

Clear public key infrastructure (PKI) key pair information for local digital certificates on the local device in a Multinode High Availability system.

Options

- `all`—Clear key pair information for all local certificates.
- `certificate-id` *certificate-id* —Clear key pair information for the local certificate with this certificate ID.

Required Privilege Level

clear and security

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 22.3R1

RELATED DOCUMENTATION

Multinode High Availability

request security pki node-local generate-key-pair

request security pki node-local key-pair export

clear security pki node-local local-certificate

IN THIS SECTION

- Syntax | 4032
- Description | 4032
- Options | 4032
- Required Privilege Level | 4032
- Output Fields | 4032
- Sample Output | 4032
- Release Information | 4032

Syntax

```
clear security pki node-local local-certificate (all | certificate-id certificate-id | system-generated)
```

Description

Clear public key infrastructure (PKI) information for local digital certificates on the local device in a Multinode High Availability system.

Options

- `all`—Clear information for all the local digital certificates on the device.

  You cannot clear the automatically generated self-signed certificate using `clear security pki local-certificate all` command. To clear the self-signed certificate you need to use `system-generated` as an option.

- `certificate-id` *`certificate-id`* —Clear the specified local digital certificate with this certificate ID.

- `system-generated`—Clear the existing automatically generated self-signed certificate and generate a new self-signed certificate.

Required Privilege Level

clear and security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

This command produces no output.

Release Information

Command modified in Junos OS Release 22.3R1

RELATED DOCUMENTATION

Multinode High Availability

show security pki node-local local-certificate

request security pki node-local local-certificate load

request security pki node-local local-certificate export

request security pki node-local local-certificate enroll

request security pki node-local local-certificate re-enroll

clear security pki statistics

IN THIS SECTION

- Syntax | 4033
- Description | 4033
- Options | 4033
- Required Privilege Level | 4034
- Output Fields | 4034
- Release Information | 4034

Syntax

```
clear security pki statistics
```

Description

Clear PKI statistics on the device.

Options

None

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

show security pki statistics | 25009

clear security policies file-serialization counters

IN THIS SECTION

- Syntax | 4034
- Description | 4035
- Options | 4035
- Required Privilege Level | 4035
- Sample Output | 4035
- Output Fields | 4035
- Release Information | 4035

Syntax

```
clear security policies file-serialization counters
```

Description

Clear the file serialization counters.

Options

This command has no options.

Required Privilege Level

view

Sample Output

clear security policies file-serialization counters

```
user@host> clear security policies file-serialization counters
```

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Release Information

Command introduced in Junos OS Release 24.4R1

clear security policies hit-count

IN THIS SECTION

- Syntax | 4036
- Description | 4036
- Options | 4036
- Additional Information | 4036

- Required Privilege Level | 4036
- Output Fields | 4036
- Release Information | 4036

Syntax

```
clear security policies hit-count
<from-zone zone-name>
<to-zone zone-name>
```

Description

Clear the hit-count values for security policies.

Options

- `from-zone` *`zone-name`*—(Optional) Clear the number of hits for security policies associated with the named source zone.
- `to-zone` *`zone-name`*—(Optional) Clear the number of hits for security policies associated with the named destination zone.

Additional Information

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos Release 12.1.

RELATED DOCUMENTATION

show security policies hit-count | 25061

Monitoring Security Policy Statistics

clear security policies statistics

IN THIS SECTION

- Syntax | 4037
- Description | 4037
- Required Privilege Level | 4037
- Output Fields | 4037
- Release Information | 4037

Syntax

```
clear security policies statistics
```

Description

Clear systemwide policies statistics and security policies statistics configured on the device.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 8.5. Support.

RELATED DOCUMENTATION

show security policies | 25028

Monitoring Security Policy Statistics

clear security screen statistics

IN THIS SECTION

- Syntax | 4038
- Description | 4038
- Options | 4038
- Required Privilege Level | 4039
- Output Fields | 4039
- Sample Output | 4039
- Release Information | 4039

Syntax

```
clear security screen statistics
<node (          node-id         | all | local | primary)>
```

Description

Clear intrusion detection service (IDS) security screen statistics on the device.

Options

`node`—(Optional) For chassis cluster configurations, clear security screen statistics on a specific node.

- *`node-id`* —Identification number of the node. It can be 0 or 1.
- `all` —Clear all nodes.
- `local` —Clear the local node.

- `primary`—Clear the primary node.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear security screen statistics node 0

```
user@host> clear security screen statistics node 0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show security screen statistics | 25116

Example: Configuring Multiple Screening Options

clear security screen statistics interface

IN THIS SECTION

- Syntax | 4040
- Description | 4040
- Options | 4040
- Required Privilege Level | 4040

- Output Fields | 4040
- Sample Output | 4041
- Sample Output | 4041
- Release Information | 4041

Syntax

```
clear security screen statistics interface interface-name
logical-system
root-logical-system
tenant
```

Description

Clear intrusion detection service (IDS) security screen statistics for an interface under one specific logical system or a tenant system.

Options

interface-name	Name of the interface on which to clear security screen statistics.
logical-system	Name of the logical system.
root-logical-system	(Default) Displays root logical system as default.
tenant	The name of the tenant system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security screen statistics interface fab0

```
user@host> clear security screen statistics interface fab0
node0:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
node1:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
```

Sample Output

clear security screen statistics interface fab0 node 0

```
user@host> clear security screen statistics interface fab0 node 0
node0:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
```

clear security screen statistics tenant all interface ge-10/0/1

```
user@host> clear security screen statistics tenant all interface ge-10/0/1
IDS statistics has been cleared.
```

Release Information

Command introduced in Junos OS Release 8.5.

The `node` option is added in Junos OS Release 9.0.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

show security screen statistics | 25116

Example: Configuring Multiple Screening Options

clear security screen statistics zone

IN THIS SECTION

- Syntax | 4042
- Description | 4042
- Options | 4042
- Required Privilege Level | 4043
- Output Fields | 4043
- Sample Output | 4043
- Sample Output | 4043
- Release Information | 4044

Syntax

```
clear security screen statistics zone <zone-name>
logical-system
root-logical-system
tenant
```

Description

Clear IDS security screen statistics for a security zone under all or a specific logical systems or tenant systems.

Options

zone <*zone-name*>	Name of the security zone.
logical-system	Name of the logical system.

root-logical-system	(Default) Displays root logical system as default.
tenant	The name of the tenant system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security screen statistics zone abc node all

```
user@host> clear security screen statistics zone abc node all
node0:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
node1:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
```

Sample Output

clear security screen statistics node 0 zone my-zone

```
user@host> clear security screen statistics node 0 zone my-zone
node0:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
```

clear security screen statistics tenant TN1 zone trust

```
user@host> clear security screen statistics tenant TN1 zone trust
IDS statistics has been cleared.
```

clear security screen statistics tenant all zone trust

```
user@host> clear security screen statistics tenant all zone trust
IDS statistics has been cleared.
```

Release Information

Command introduced in Junos OS Release 8.5.

The `node` option is added in Junos OS Release 9.0.

The `tenant` option is added in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

show security screen statistics | 25116

Example: Configuring Multiple Screening Options

clear security ssh key-pair-identity

IN THIS SECTION

- Syntax | 4045
- Description | 4045
- Options | 4045
- Required Privilege Level | 4045
- Output Fields | 4045
- Sample Output | 4045

Release Information | 4045

Syntax

```
clear security ssh key-pair-identity <all> | <identity-name>
```

Description

Clear private and public SSH key pair for the specified files.

Options

- `all`— Clear all the key-pair files.
- *`identity-name`*—Clear identity name.

Required Privilege Level

clear

Output Fields

Sample Output

clear security ssh key-pair-identity sample

```
user@host> clear security ssh key-pair-identity sample
SSH key sample was removed
```

Release Information

Command introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

request security ssh key-pair-identity generate | 14865

show security ssh key-pair-identity | 25144

clear security tcp-encap statistics

IN THIS SECTION

- Syntax | 4046
- Description | 4046
- Required Privilege Level | 4046
- Output Fields | 4046
- Release Information | 4046

Syntax

```
clear security tcp-encap statistics
```

Description

Clear TCP encapsulation statistics.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X49-D80.

RELATED DOCUMENTATION

show security tcp-encap statistics | 25149

clear security user-identification local-authentication-table

IN THIS SECTION

- Syntax | 4047
- Description | 4047
- Required Privilege Level | 4047
- Output Fields | 4047
- Sample Output | 4047
- Release Information | 4048

Syntax

```
clear security user-identification local-authentication-table
```

Description

This command removes all entries from the local authentication table.

Required Privilege Level

clear

Output Fields

When you enter this command, all entries are cleared from the local authentication table.

Sample Output

clear security user-identification local-authentication-table

```
user@host> clear security user-identification local-authentication-table
user@host> show security user-identification local-authentication-table all
Total entries: 0
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Application Security

Firewall User Authentication Overview

clear security utm anti-spam statistics

IN THIS SECTION

- Syntax | 4048
- Description | 4049
- Options | 4049
- Required Privilege Level | 4049
- Sample Output | 4050
- Release Information | 4051

Syntax

```
clear security utm anti-spam statistics
<root-logical-system>
<logical-system (logical-system-name | all)>
```

```
<all-logical-systems-tenants>
<tenant (tenant-name | all)>
```

Description

Clears antispam statistics information. With chassis cluster support for Content Security, statistics from both the nodes is cleared.

Starting in Junos OS Release 18.3R1, you can clear the antispam statistics information for the primary logical system or for a specific user logical system or for all the user logical systems.

Starting in Junos OS Release 19.2R1, you can clear the antispam statistics information for a specific tenant system or for all the tenant systems.

Options

none Clears the antispam statistics information for the primary logical system.

root-logical-system (Optional) Clears the antispam statistics information for the primary logical system.

logical-system *logical-system-name* (Optional) Clears the antispam statistics information for a specific user logical system.

all (Optional) Clears the antispam statistics information for all the user logical systems.

all-logical-systems-tenants (Optional) Clears the antispam statistics information for all the logical systems and tenant systems.

tenant *tenant-name* (Optional) Clears the antispam statistics information for a specific tenant system.

all (Optional) Clears the antispam statistics information for all the tenant systems.

Required Privilege Level

clear

Sample Output

clear security utm anti-spam statistics

```
user@host> clear security utm anti-spam statistics
Anti-spam clear statistics result: clear done
```

clear security utm anti-spam statistics root-logical-system

```
user@host> clear security utm anti-spam statistics root-logical-system
Anti-spam clear statistics result: clear done
```

clear security utm anti-spam statistics logical-system LSYS1

```
user@host> clear security utm anti-spam statistics logical-system LSYS1
Anti-spam clear statistics result: clear done
```

clear security utm anti-spam statistics logical-system all

```
user@host> clear security utm anti-spam statistics logical-system all
Anti-spam clear statistics result: clear done
```

clear security utm anti-spam statistics tenant TSYS1

```
user@host> clear security utm anti-spam statistics tenant TSYS1
Anti-spam clear statistics result: clear done
```

clear security utm anti-spam statistics tenant all

```
user@host> clear security utm anti-spam statistics tenant all
Anti-spam clear statistics result: clear done
```

clear security utm anti-spam statistics all-logical-systems-tenants

```
user@host> clear security utm anti-spam statistics all-logical-systems-tenants
Anti-spam clear statistics result: clear done
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for Content Security in chassis cluster added in Junos OS Release 11.4.

Support for Content Security in logical system added in Junos OS Release 18.3R1.

Support for Content Security in tenant system added in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security utm anti-spam statistics | 25176

show security utm anti-spam status | 25182

clear security utm antivirus statistics

IN THIS SECTION

- Syntax | 4052
- Description | 4052
- Options | 4052
- Required Privilege Level | 4052
- Sample Output | 4053
- Release Information | 4054

Syntax

```
clear security utm anti-virus statistics
<root-logical-system>
<logical-system (logical-system-name | all)>
<all-logical-systems-tenants>
<tenant (tenant-name | all)>
```

Description

Clears antivirus statistics information. With chassis cluster support for Content Security, statistics from both the nodes are cleared.

Starting in Junos OS Release 18.3R1, you can clear the antivirus statistics information for the primary logical system or for a specific user logical system or for all the user logical systems.

Starting in Junos OS Release 19.2R1, you can clear the antivirus statistics information for a specific tenant system or for all the tenant systems.

Options

none Clears the antivirus statistics information for the primary logical system.

root-logical-system (Optional) Clears the antivirus statistics information for the primary logical system.

logical-system *logical-system-name* (Optional) Clears the antivirus statistics information for a specific user logical system.

all (Optional) Clears the antivirus statistics information for all the user logical systems.

all-logical-systems-tenants (Optional) Clears the antivirus statistics information for all the logical systems and tenant systems.

tenant *tenant-name* (Optional) Clears the antivirus statistics information for a specific tenant system.

all (Optional) Clears the antivirus statistics information for all the tenant systems.

Required Privilege Level

clear

Sample Output

clear security utm anti-virus statistics

```
user@host> clear security utm anti-virus statistics
Anti-virus clear statistics result: clear done
```

clear security utm anti-virus statistics root-logical-system

```
user@host> clear security utm anti-virus statistics root-logical-system
Anti-virus clear statistics result: clear done
```

clear security utm anti-virus statistics logical-system LSYS1

```
user@host> clear security utm anti-virus statistics logical-system LSYS1
Anti-virus clear statistics result: clear done
```

clear security utm anti-virus statistics logical-system all

```
user@host> clear security utm anti-virus statistics logical-system all
Anti-virus clear statistics result: clear done
```

clear security utm anti-virus statistics tenant TSYS1

```
user@host> clear security utm anti-virus statistics tenant TSYS1
Anti-virus clear statistics result: clear done
```

clear security utm anti-virus statistics tenant all

```
user@host> clear security utm anti-virus statistics tenant all
Anti-virus clear statistics result: clear done
```

clear security utm anti-virus statistics all-logical-systems-tenants

```
user@host> clear security utm anti-virus statistics all-logical-systems-tenants
Anti-virus clear statistics result: clear done
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for Sophos Antivirus added in Junos OS Release 11.1.

Support for Content Security in chassis cluster added in Junos OS Release 11.4.

Support for Content Security in logical system added in Junos OS Release 18.3R1.

Support for Content Security in tenant system added in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

- show security utm anti-virus statistics | 25183
- show security utm anti-virus status | 25191
- request security utm anti-virus juniper-express-engine | 14881
- request security utm anti-virus kaspersky-lab-engine | 14883

clear security utm content-filtering statistics

IN THIS SECTION

- Syntax | 4055
- Description | 4055
- Options | 4055
- Required Privilege Level | 4056
- Sample Output | 4056
- Release Information | 4057

Syntax

```
clear security utm content-filtering statistics
<root-logical-system>
<logical-system (logical-system-name | all)>
<all-logical-systems-tenants>
<tenant (tenant-name | all)>
```

Description

Clears content-filtering statistics information. With chassis cluster support for Content Security, statistics from both the nodes are cleared.

Starting in Junos OS Release 18.3R1, you can clear the content filtering statistics information for the primary logical system or for a specific user logical system or for all the user logical systems.

Starting in Junos OS Release 19.2R1, you can clear the content filtering statistics information for a specific tenant system or for all the tenant systems.

Options

none Clears the content filtering statistics information for the primary logical system.

root-logical-system (Optional) Clears the content filtering statistics information for the primary logical system.

logical-system *logical-system-name* (Optional) Clears the content filtering statistics information for a specific user logical system.

all (Optional) Clears the content filtering statistics information for all the user logical systems.

all-logical-systems-tenants (Optional) Clears the content filtering statistics information for all the logical systems and tenant systems.

tenant *tenant-name* (Optional) Clears the content filtering statistics information for a specific tenant system.

all (Optional) Clears the content filtering statistics information for all the tenant systems.

Required Privilege Level

clear

Sample Output

clear security utm content-filtering statistics

```
user@host> clear security utm content-filtering statistics
Content-filtering clear statistics result: clear done
```

clear security utm content-filtering statistics root-logical-system

```
user@host> clear security utm content-filtering statistics root-logical-system
Content-filtering clear statistics result: clear done
```

clear security utm content-filtering statistics logical-system LSYS1

```
user@host> clear security utm content-filtering statistics logical-system LSYS1
Content-filtering clear statistics result: clear done
```

clear security utm content-filtering statistics logical-system all

```
user@host> clear security utm content-filtering statistics logical-system all
Content-filtering clear statistics result: clear done
```

clear security utm content-filtering statistics tenant TSYS1

```
user@host> clear security utm content-filtering statistics tenant TSYS1
Content-filtering clear statistics result: clear done
```

clear security utm content-filtering statistics tenant all

```
user@host> clear security utm content-filtering statistics tenant all
Content-filtering clear statistics result: clear done
```

clear security utm content-filtering statistics all-logical-systems-tenants

```
user@host> clear security utm content-filtering statistics all-logical-systems-tenants
Content-filtering clear statistics result: clear done
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for Content Security in chassis cluster added in Junos OS Release 11.4.

Support for Content Security in logical system added in Junos OS Release 18.3R1.

Support for Content Security in tenant system added in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security utm content-filtering statistics | 25195

clear security utm session

IN THIS SECTION

- Syntax | 4058
- Description | 4058
- Required Privilege Level | 4058
- Output Fields | 4058
- Release Information | 4058

Syntax

```
clear security utm session
```

Description

Clear Content Security session information. With chassis cluster support for Content Security, sessions on both the nodes are cleared.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.5.

Support for Content Security in chassis cluster added in Junos OS Release 11.4.

RELATED DOCUMENTATION

show security utm session | 25212

show security utm status | 25213

clear security utm web-filtering statistics

IN THIS SECTION

Syntax | 4059

Description | 4059

- Options | 4059
- Required Privilege Level | 4060
- Sample Output | 4060
- Release Information | 4061

Syntax

```
clear security utm web-filtering statistics
<root-logical-system>
<logical-system (logical-system-name | all)>
<all-logical-systems-tenants>
<tenant (tenant-name | all)>
```

Description

Clear web filtering statistics information. With chassis cluster support for Content Security, statistics from both the nodes is cleared.

Starting in Junos OS Release 18.3R1, you can clear the Web filtering statistics information for the primary logical system or for a specific user logical system or for all the user logical systems.

Starting in Junos OS Release 19.2R1, you can clear the Web filtering statistics information for a specific tenant system or for all the tenant systems.

Options

none Clears the Web filtering statistics information for the primary logical system.

root-logical-system (Optional) Clears the Web filtering statistics information for the primary logical system.

logical-system *logical-system-name* (Optional) Clears the Web filtering statistics information for a specific user logical system.

all (Optional) Clears the Web filtering statistics information for all the user logical systems.

all-logical-systems-tenants (Optional) Clears the Web filtering statistics information for all the logical systems and tenant systems.

tenant *tenant-name* (Optional) Clears the Web filtering statistics information for a specific tenant system.

all (Optional) Clears the Web filtering statistics information for all the tenant systems.

Required Privilege Level

clear

Sample Output

clear security utm web-filtering statistics

```
user@host> clear security utm web-filtering statistics
Web-filtering clear statistics result: clear done
```

clear security utm web-filtering statistics root-logical-system

```
user@host> clear security utm web-filtering statistics root-logical-system
Web-filtering clear statistics result: clear done
```

clear security utm web-filtering statistics logical-system LSYS1

```
user@host> clear security utm web-filtering statistics logical-system LSYS1
Web-filtering clear statistics result: clear done
```

clear security utm web-filtering statistics logical-system all

```
user@host> clear security utm web-filtering statistics logical-system all
Web-filtering clear statistics result: clear done
```

clear security utm web-filtering statistics tenant TSYS1

```
user@host> clear security utm web-filtering statistics tenant TSYS1
Web-filtering clear statistics result: clear done
```

clear security utm web-filtering statistics tenant all

```
user@host> clear security utm web-filtering statistics tenant all
Web-filtering clear statistics result: clear done
```

clear security utm web-filtering statistics all-logical-systems-tenants

```
user@host> clear security utm web-filtering statistics all-logical-systems-tenants
Web-filtering clear statistics result: clear done
```

Release Information

Command introduced in Junos OS Release 9.5 .

Support for Content Security in chassis cluster added in Junos OS Release 11.4 .

Support for Content Security in logical system added in Junos OS Release 18.3R1.

Support for Content Security in tenant system added in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security utm web-filtering statistics | 25220

show security utm web-filtering status | 25228

clear service l2tp session routing-instance

IN THIS SECTION

- Syntax | 4062
- Description | 4062
- Options | 4062
- Required Privilege Level | 4062
- Sample Output | 4063
- Release Information | 4063

Syntax

```
clear service l2tp session routing-instance name
```

Description

Enables you to log out all the L2TP subscriber sessions associated with a routing instance at once. This feature supports IPv4, IPv6, and dual stack subscribers.

Options

routing-instance *name*	Name of the routing instance from which you need to log out all the L2TP subscriber sessions.

Required Privilege Level

clear

Sample Output

clear service l2tp session routing-instance *name*

```
user@host> clear service l2tp session routing-instance cust-vrf
Session 50924 closed
```

Release Information

Command introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

show service l2tp session routing-instance | 25246

clear service user-identification identity-management counter

IN THIS SECTION

- Syntax | 4063
- Description | 4064
- Options | 4064
- Required Privilege Level | 4064
- Release Information | 4064

Syntax

```
clear service user-identification identity-management counter
```

Description

Clear the counters associated with the batch queries and IP queries for the advanced user query feature.

Options

This command has no options.

Required Privilege Level

clear

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

clear services accounting statistics inline-jflow

IN THIS SECTION

- Syntax | 4064
- Description | 4065
- Options | 4065
- Required Privilege Level | 4065
- Output Fields | 4065
- Sample Output | 4065
- Release Information | 4065

Syntax

```
clear services accounting statistics inline-jflow
<inline-jflow (fpc-slot slot-number)>
```

Description

Clear inline flow statistics for a specified FPC.

Options

fpc-slot *slot-number* Clear inline flow statistics for the specified FPC.

- MX80 Series routers only—Replace *slot-number* with a value from 0 through 1.
- MX104 Series routers only—Replace *slot-number* with a value from 0 through 2.
- MX240 Series routers only—Replace *slot-number* with a value from 0 through 2.
- MX480 Series routers only—Replace *slot-number* with a value from 0 through 5.
- MX960 Series routers only—Replace *slot-number* with a value from 0 through 11.
- MX2010 Series routers only—Replace *slot-number* with a value from 0 through 9.
- MX2020 Series routers only—Replace *slot-number* with a value from 0 through 19.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services accounting statistics inline-jflow

```
user@host> run clear services accounting statistics inline-jflow fpc-slot 5
Statistics Cleared
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

show services accounting flow | 25277

clear services advanced-anti-malware statistics

IN THIS SECTION

- Syntax | 4066
- Description | 4066
- Required Privilege Level | 4067
- Output Fields | 4067
- Sample Output | 4067
- Release Information | 4067

Syntax

```
clear services advanced-anti-malware statistics
```

Description

Sets the following advanced anti-malware statistics to 0. See "show services advanced-anti-malware statistics" on page 25321 for information on these statistics.

Use this command when you change the Juniper Advanced Threat Prevention Cloud profile or policy settings and want to reset the statistics in order to view the effects of the changes.

- Number of session checked
- Number of session ignored
- Number of session permitted
- Number of session blocked
- Number of log notification

- Number of file sent success
- Number of file sent failure
- Number of file sent ignore
- Blocklist-hit-counter
- Whitelist-hit-counter
- Number of http session actives
- Number of https session actives
- Number of http session processed
- Number of https session processed

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear services advanced-anti-malware statistics

```
user@host> clear services advanced-anti-malware statistics
```

Release Information

Command introduced in Junos OS Release 15.1X49-D33.

RELATED DOCUMENTATION

show services advanced-anti-malware statistics | 25321

clear services agf nas statistics

IN THIS SECTION

- Syntax | 4068
- Description | 4068
- Options | 4068
- Required Privilege Level | 4068
- Release Information | 4068

Syntax

```
clear services agf nas statistics
```

Description

Clear Non-Access Stratum (NAS) messaging information that the Access Gateway Function (AGF) has collected.

Options

none

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 22.3R1.

clear services agf ngap amf

IN THIS SECTION

- Syntax | 4069
- Description | 4069
- Options | 4069
- Required Privilege Level | 4069
- Release Information | 4069

Syntax

```
clear services agf ngap amf
<amf-name>
statistics
```

Description

Clear N1 messages sent to the Access Gateway Function (AGF) from the Access and Mobility Management Functions (AMFs).

Options

amf-name	(Optional) Clear the statistics for a particular AMF.
statistics	Clear the statistics for all the AMFs.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 22.3R1.

clear services alg statistics

IN THIS SECTION

- Syntax | 4070
- Description | 4070
- Options | 4070
- Required Privilege Level | 4070
- Release Information | 4070

Syntax

```
clear services alg statistics
```

Description

Clear ALG statistics for Junos OS extension-provider packages.

Options

application-profile	Clear all sessions for the application profile.
interface	Clear all sessions for the interface.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 10.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

clear services application-aware-access-list statistics

IN THIS SECTION

- Syntax | 4071
- Description | 4071
- Options | 4071
- Required Privilege Level | 4071
- Release Information | 4071

Syntax

```
clear services application-aware-access-list statistics
```

Description

Clear application-aware access list (AACL) statistics.

Options

This command has no options.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

show services application-aware-access-list statistics | 25386

clear services application-identification application-statistics

IN THIS SECTION

- Syntax | 4072
- Description | 4072
- Required Privilege Level | 4072
- Output Fields | 4072
- Release Information | 4072

Syntax

```
clear services application-identification application-statistics
```

Description

Clears all Junos OS application statistics such as cumulative, interval, applications, and application groups.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show services application-identification statistics applications | 25451

show services application-identification statistics application-groups | 25446
clear services application-identification application-statistics interval | 4074
clear services application-identification application-statistics cumulative | 4073

clear services application-identification application-statistics cumulative

IN THIS SECTION

- Syntax | 4073
- Description | 4073
- Required Privilege Level | 4073
- Output Fields | 4073
- Release Information | 4073

Syntax

```
clear services application-identification application-statistics cumulative
```

Description

Clear all Junos OS application cumulative statistics.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show services application-identification statistics applications | 25451
show services application-identification statistics application-groups | 25446
clear services application-identification application-statistics | 4072
clear services application-identification application-statistics interval | 4074

clear services application-identification application-statistics interval

IN THIS SECTION

- Syntax | 4074
- Description | 4074
- Required Privilege Level | 4074
- Output Fields | 4074
- Release Information | 4075

Syntax

```
clear services application-identification application-statistics interval
```

Description

Clear all Junos OS application interval statistics.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- show services application-identification statistics applications | 25451
- show services application-identification statistics application-groups | 25446
- clear services application-identification application-statistics | 4072
- clear services application-identification application-statistics cumulative | 4073

clear services application-identification application-system-cache

IN THIS SECTION

- Syntax | 4075
- Description | 4075
- Options | 4076
- Required Privilege Level | 4076
- Output Fields | 4076
- Sample Output | 4076
- Release Information | 4076

Syntax

```
clear services application-identification application-system-cache
```

Description

Clear entries from the application system cache.

Options

This command has no options.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear services application-identification application-system-cache

```
user@host> clear services application-identification application-system-cache
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show services application-identification application-system-cache | 25408

clear services application-identification application-system-cache (Junos OS)

IN THIS SECTION

- Syntax | 4077
- Description | 4077
- Options | 4077

- Required Privilege Level | 4078
- Output Fields | 4078
- Release Information | 4078

Syntax

```
clear services application-identification application-system-cache
<node ( node-id | all | local | primary ) >
<logical-system (logical-system-name | all | root-logical-system)>
<tenant (tenant-name | all)>
```

Description

Clear Junos OS application identification application system cache.

Options

none Clears the application system cache on the device.

node (Optional) For chassis cluster configurations, clear application system cache on the specified nodes.

- `node-id`—Specific node number
- `all`—All nodes
- `local`—Local node
- `primary`—Primary node

logical-system *logical-system-name* (Optional) Clears the application system cache of the specified logical system.

logical-system *all* (Optional) Clears the application system cache of all the logical systems.

root-logical-system (Optional) Clears the application system cache of the root logical system.

tenant *tenant-name* (Optional) Clears the application system cache of the specified tenant system.

tenant *all* (Optional) Clears the application system cache of all the tenant systems.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 10.2.

Command syntax updated in Junos OS Release 12.1.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

show services application-identification application-system-cache (View) | 25413

clear services application-identification counter (Values)

IN THIS SECTION

- Syntax | 4079
- Description | 4079
- Options | 4079
- Required Privilege Level | 4079
- Output Fields | 4079
- Sample Output | 4080
- Release Information | 4080

Syntax

```
clear services application-identification counter
<ssl-encrypted-sessions>
<logical-system (logical-system-name | all | root-logical-system)>
<tenant (tenant-name | all)>
```

Description

Resets all the Junos OS application identification counter values.

Options

ssl-encrypted-sessions	(Optional) Resets application identification counter values for SSL encrypted sessions.
logical-system *logical-system-name*	(Optional) Resets application identification counter values of the specified logical system.
logical-system *all*	(Optional) Resets application identification counter values of all the logical systems.
root-logical-system	(Optional) Resets application identification counter values of the root logical system.
tenant *tenant-name*	(Optional) Resets application identification counter values of the specified tenant system.
tenant *all*	(Optional) Resets application identification counter values of all the tenant systems.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services application-identification counter

```
user@host> clear services application-identification counter
clear_counter_class: counters cleared, status = 0
```

clear services application-identification counter logical-system all

```
user@host> clear services application-identification counter logical-system all
appid counter cleared
```

clear services application-identification counter

```
user@host:TSYS1> clear services application-identification counter
appid counter cleared
```

clear services application-identification counter tenant all

```
user@host> clear services application-identification counter tenant all
appid counter cleared
```

Release Information

Command introduced in Junos OS Release 10.2. Command updated in Junos OS Release 12.1-X47-D15.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

show services application-identification counter (AppSecure) | **25426**

clear services application-identification packet-capture counters

IN THIS SECTION

- Syntax | 4081
- Description | 4081
- Required Privilege Level | 4081
- Output Fields | 4081
- Sample Output | 4081
- Release Information | 4082

Syntax

```
clear services application-identification packet-capture counters
```

Description

Clear all packet capture counters generated for unknown application traffic.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

command-name

```
user@host> clear services application-identification packet-capture counters
Packet-capture counters cleared
```

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Configure Packet Capture For Unknown Application Traffic

packet-capture | **12137**

show services application-identification packet-capture counters | **25443**

request services application-identification clear packet-capture all | **14918**

clear services application-identification statistics

IN THIS SECTION

- Syntax | **4082**
- Description | **4083**
- Options | **4083**
- Required Privilege Level | **4083**
- Output Fields | **4083**
- Sample Output | **4083**
- Release Information | **4084**

Syntax

```
clear services application-identification statistics
<cumulative>
<interval>
<logical-system (logical-system-name | all | root-logical-system)>
<tenant (tenant-name | all)>
```

Description

Clears all Junos OS application statistics such as cumulative, interval, applications, and application groups.

Options

cumulative	(Optional) Clears the cumulative application statistics.
interval	(Optional) Clears the application interval statistics. Interval statistics are displayed in Top-N format, such that the first application group displayed has the largest byte count. If this parameter is not specified, then the default is 1, which is the current interval.
logical-system *logical-system-name*	(Optional) Clears application identification statistics of the specified logical system.
logical-system *all*	(Optional) Clears application identification statistics of all the logical systems.
root-logical-system	(Optional) Clears application identification statistics of the root logical system.
tenant *tenant-name*	(Optional) Clears application identification statistics of the specified tenant system.
tenant *all*	(Optional) Clears application identification statistics of all the tenant systems.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear services application-identification statistics

```
user@host> clear services application-identification statistics
appid statistics cleared
```

clear services application-identification statistics logical-system all

```
user@host> clear services application-identification statistics logical-system all
appid statistics cleared
```

clear services application-identification statistics cumulative tenant TSYS1

```
user@host> clear services application-identification statistics cumulative tenant TSYS1
appid statistics cleared
```

clear services application-identification statistics cumulative tenant all

```
user@host> clear services application-identification statistics cumulative tenant all
appid statistics cleared
```

clear services application-identification statistics cumulative

```
user@host:TSYS1> clear services application-identification statistics cumulative
appid statistics cleared
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

`logical-system` option introduced in Junos OS Release 18.3R1 on SRX Series.

`tenant` option introduced in Junos OS Release 19.4R1 on SRX Series.

RELATED DOCUMENTATION

show services application-identification statistics applications | 25451

show services application-identification statistics application-groups | 25446

clear services captive-portal-content-delivery statistics

IN THIS SECTION

- Syntax | 4085
- Description | 4085
- Options | 4085
- Required Privilege Level | 4085
- Output Fields | 4085
- clear services captive-portal-content-delivery statistics | 4086
- Release Information | 4086

Syntax

```
clear services captive-portal-content-delivery statistics
<interface pic-name>
```

Description

Clear captive portal content delivery statistics.

Options

interface	Clear statistics by PIC name.

Required Privilege Level

clear

Output Fields

When you enter this command, you receive feedback on the status of your request.

clear services captive-portal-content-delivery statistics

command-name

```
user@host> clear services captive-portal-content-delivery statistics interface ms-5/0/0

user@host> show services captive-portal-content-delivery statistics interface ms-5/0/0

service-set interface: ms-5/0/0

Packets received   Packets altered
0                     0

Note that the stats are cleared.
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

show services captive-portal-content-delivery | 25476

clear services cos statistics

IN THIS SECTION

- Syntax | 4087
- Description | 4087
- Options | 4087
- Required Privilege Level | 4087
- Output Fields | 4087
- Sample Output | 4087

Release Information | 4088

Syntax

```
clear services cos statistics
<interface interface-name>
 <service-set service-set-name>
```

Description

Clear statistics for class-of-service (CoS) code point bit patterns and forwarding classes as configured in CoS services for the AS PIC.

Options

none	Clear all services CoS statistics.
interface *interface-name*	(Optional) Clear statistics for the specified interface only.
service-set *service-set-name*	(Optional) Clear statistics for the specified service set only.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services cos statistics

```
user@host> clear services cos statistics
```

Release Information

Command introduced in Junos OS Release 8.1.

clear services crtp statistics

IN THIS SECTION

- Syntax | 4088
- Description | 4088
- Options | 4088
- Required Privilege Level | 4089
- Output Fields | 4089
- Sample Output | 4089
- Release Information | 4089

Syntax

```
clear services crtp statistics
<interface interface-name>
```

Description

Clear Compressed Real-Time Transport Protocol (CRTP) flow statistics.

Options

none	Clear CRTP flow statistics on all interfaces.
interface *interface-name*	(Optional) Clear CRTP flow statistics for the specified interface. On M Series and T Series routers, a link services IQ (**lsq-***fpc/pic/port*) or redundant link services IQ (**rlsq-***fpc/pic/port*) interface.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services crtp statistics

```
user@host> clear services crtp statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

clear services dns-filtering cache

IN THIS SECTION

- Syntax | 4090
- Description | 4090
- Options | 4090
- Required Privilege Level | 4090
- Output Fields | 4090
- Sample Output | 4091
- Release Information | 4092

Syntax

```
clear services dns-filtering cache
clear services dns-filtering cache logical-system all
clear services dns-filtering cache tenant all
clear services dns-filtering cache logical-system <logical-system-name>
clear services dns-filtering cache tenant <tenant-name>
clear services dns-filtering cache all-logical-systems-tenants
clear services dns-filtering cache <fqdn>
clear services dns-filtering cache logical-system all <fqdn>
clear services dns-filtering cache tenant all <fqdn>
clear services dns-filtering cache logical-system <logical-system-name> <fqdn>
clear services dns-filtering cache tenant <tenant-name> <fqdn>
clear services dns-filtering cache all-logical-systems-tenants <fqdn>
```

Description

Clear all entries in the DNS cache. See "show services dns-filtering cache" on page 25498 for information on the DNS cache entries.

Options

fqdn Clear a specific fqdn from the DNS cache.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear services dns-filtering cache

```
user@host> clear services dns-filtering cache
```

clear services dns-filtering cache logical-system all

```
user@host> clear services dns-filtering cache logical-system all
```

clear services dns-filtering cache tenant all

```
user@host> clear services dns-filtering cache tenant all
```

clear services dns-filtering cache logical-system <logical-system-name>

```
user@host> clear services dns-filtering cache logical-system <logical-system-name>
```

clear services dns-filtering cache tenant <tenant-name>

```
user@host> clear services dns-filtering cache tenant <tenant-name>
```

clear services dns-filtering cache all-logical-systems-tenants

```
user@host> clear services dns-filtering cache all-logical-systems-tenants
```

clear services dns-filtering cache <fqdn>

```
user@host> clear services dns-filtering cache <fqdn>
```

clear services dns-filtering cache logical-system all <fqdn>

```
user@host> clear services dns-filtering cache logical-system all <fqdn>
```

clear services dns-filtering cache tenant all <fqdn>

```
user@host> clear services dns-filtering cache tenant all <fqdn>
```

clear services dns-filtering cache logical-system <logical-system-name> <fqdn>

```
user@host> clear services dns-filtering cache logical-system <logical-system-name> <fqdn>
```

clear services dns-filtering cache tenant <tenant-name> <fqdn>

```
user@host> clear services dns-filtering cache tenant <tenant-name> <fqdn>
```

clear services dns-filtering cache all-logical-systems-tenants <fqdn>

```
user@host> clear services dns-filtering cache all-logical-systems-tenants <fqdn>
```

Release Information

dns-filtering cache command introduced in Junos OS Release 21.2R1

RELATED DOCUMENTATION

show services dns-filtering cache | 25498

clear services dynamic-flow-capture

IN THIS SECTION

- Syntax | 4093
- Description | 4093
- Options | 4093
- Required Privilege Level | 4094
- Output Fields | 4094
- Sample Output | 4094
- Release Information | 4094

Syntax

```
clear services dynamic-flow-capture capture-group group-name
 <criteria-identifier identifier>
<destination-identifier identifier>
<force>
<static>
```

Description

(M320 Series routers and T Series routers only) Clear dynamic flow capture information for specified capture group.

Options

capture-group *group-name*	Use the specified capture-group identifier.
criteria-identifier *identifier*	(Optional) Use the specified criteria identifier.
destination-identifier *identifier*	(Optional) Use the specified content destination identifier.
force	(Optional) Force clearing of criteria.

static	(Optional) Clear static criteria.

Required Privilege Level

network

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear services dynamic-flow-capture

```
user@host> clear services dynamic-flow-capture capture-group flow-a
```

Release Information

Command introduced in Junos OS Release 7.4.

clear services flow-collector statistics

IN THIS SECTION

- Syntax | 4095
- Description | 4095
- Options | 4095
- Required Privilege Level | 4095
- Output Fields | 4095
- Sample Output | 4095
- Release Information | 4095

Syntax

```
clear services flow-collector statistics (all | interface interface-name)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Clear statistics for one flow collector interface or for all flow collector interfaces.

Options

all	Clear statistics for all configured flow collector interfaces.
interface *interface-name*	Clear statistics for the specified flow collector interface (**cp-*fpc/pic/port***).

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services flow-collector statistics

```
user@host> clear services flow-collector statistics interface cp-5/0/0
Flow collector interface: cp-5/0/0
Interface state: Collecting flows
Statistics cleared successfully
```

Release Information

Command introduced before Junos OS Release 7.4.

clear services icap-redirect statistic

IN THIS SECTION

- Syntax | 4096
- Description | 4096
- Options | 4096
- Required Privilege Level | 4097
- Output Fields | 4097
- Sample Output | 4097
- Release Information | 4098

Syntax

```
clear services icap-redirect statistic
<all-logical-systems-tenants>
<root-logical-system>
<logical-system (logical-system-name | all)>
<tenant (tenant-name | all)>
```

Description

Clears the ICAP services redirects statistic. ICAP services redirect redirects the HTTP or HTTPS traffic to any third-party server. The security device acts as an SSL proxy server and decrypts the pass-through traffic with the proper SSL profile under a security policy.

Options

all-logical-systems-tenants	(Optional) Clears the ICAP services redirects statistic for the root logical system, all logical systems, and all tenant systems.
logical-system *logical-system-name*	(Optional) Clears the ICAP services redirects statistic for the specified logical system.
logical-system *all*	(Optional) Clears the ICAP services redirects statistic for all the logical systems.

root-logical-system	(Optional) Clears the ICAP services redirects statistic for the root logical system.
tenant *tenant-name*	(Optional) Clears the ICAP services redirects statistic for the specified tenant system.
tenant *all*	(Optional) Clears the ICAP services redirects statistic for all the tenant systems.

Required Privilege Level

clear

Output Fields

Sample Output

clear services icap-redirect statistic root-logical-system

```
user@host> clear services icap-redirect statistic root-logical-system
STATISTICS CLEARED
```

clear services icap-redirect statistic all-logical-systems-tenants

```
user@host> clear services icap-redirect statistic all-logical-systems-tenants
STATISTICS CLEARED
```

clear services icap-redirect statistic logical-system LSYS1

```
user@host> clear services icap-redirect statistic logical-system LSYS1
STATISTICS CLEARED
```

clear services icap-redirect statistic tenant TSYS1

```
user@host> clear services icap-redirect statistic tenant TSYS1
STATISTICS CLEARED
```

clear services icap-redirect statistic

```
user@host:TSYS1> clear services icap-redirect statistic
STATISTICS CLEARED
```

Release Information

Command introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

ICAP Service Redirect

clear services ids

IN THIS SECTION

- Syntax | 4099
- Description | 4099
- Options | 4099
- Required Privilege Level | 4099
- Output Fields | 4099
- Sample Output | 4099
- Release Information | 4099

Syntax

```
clear services ids
<interface interface-name>
<service-set service-set-name>
```

Description

Clear intrusion detection service (IDS) events.

Options

none	Clear all IDS events for all adaptive services interfaces for all service sets, and clear and reset IDS.
interface ***interface-name***	(Optional) On M Series and T Series routers, the ***interface-name*** can be **sp-***fpc/pic/port* or **rsp***number*.
service-set ***service-set-name***	(Optional) Clear all IDS events for a particular service set.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services ids

```
user@host> clear services ids
```

Release Information

Command introduced before Junos OS Release 7.4.

clear services ids destination-table

IN THIS SECTION

- Syntax | 4100
- Description | 4100
- Options | 4100
- Required Privilege Level | 4101
- Output Fields | 4101
- Sample Output | 4101
- Release Information | 4101

Syntax

```
clear services ids destination-table
<destination-prefix destination-prefix-name>
<interface interface-name>
<service-set service-set-name>
```

Description

Clear the intrusion detection service (IDS) events for a particular address that might be under attack.

Options

none	Clear the attack destination address table.
destination-prefix *destination-prefix-name*	(Optional) Clear the attack destination table for a particular destination prefix.
interface *interface-name*	(Optional) Clear the attack destination table for a particular interface. On M Series and T Series routers, the *interface-name* can be **sp-***fpc/pic/port* or **rsp***number*.
service-set *service-set-name*	(Optional) Clear the attack destination table for a particular service set.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services ids destination-table

```
user@host> clear services ids destination-table
```

Release Information

Command introduced before Junos OS Release 7.4.

clear services ids pair-table

IN THIS SECTION

- Syntax | 4102
- Description | 4102
- Options | 4102
- Required Privilege Level | 4102
- Output Fields | 4102
- Sample Output | 4103
- Release Information | 4103

Syntax

```
clear services ids pair-table
<destination-prefix destination-prefix-name>
<interface interface-name>
<service-set service-set-name>
<source-prefix source-prefix-name>
```

Description

Clear the intrusion detection service (IDS) attack source and destination address pair table.

Options

none	Clear the attack source and destination address pair table.
destination-prefix *destination-prefix-name*	(Optional) Clear the attack source and destination address pair table for a particular destination prefix.
interface *interface-name*	(Optional) Clear the attack destination table for a particular interface. On M Series and T Series routers, the *interface-name* can be **sp-***fpc/pic/port* or **rsp***number*.
service-set *service-set-name*	Optional) Clear the attack source and destination address pair table for a particular service set.
source-prefix *source-prefix-name*	(Optional) Clear the attack source and destination address pair table for a particular source prefix.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services ids pair-table

```
user@host> clear services ids pair-table
```

Release Information

Command introduced before Junos OS Release 7.4.

clear services ids source-table

IN THIS SECTION

- Syntax | 4103
- Description | 4103
- Options | 4104
- Required Privilege Level | 4104
- Output Fields | 4104
- Sample Output | 4104
- Release Information | 4104

Syntax

```
clear services ids source-table
<interface interface-name>
<service-set service-set-name>
<source-prefix source-prefix-name>
```

Description

Clear all intrusion detection service (IDS) events for addresses that are suspected attackers.

Options

none	Clear the attack source address table.
interface *interface-name*	(Optional) On M Series and T Series routers, the *interface-name* can be **sp-***fpc/pic/port* or **rsp***number*.
service-set *service-set-name*	(Optional) Clear the attack source address table for a particular service set.
source-prefix source-prefix-name	(Optional) Clear the attack source address table for a particular source prefix.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services ids source-table

```
user@host> clear services ids source-table
```

Release Information

Command introduced before Junos OS Release 7.4.

clear services inline-monitoring statistics

IN THIS SECTION

- Syntax | 4105
- Description | 4105
- Options | 4105
- Required Privilege Level | 4105
- Release Information | 4106

Syntax

```
clear services inline-monitoring statistics fpc-slot fpc-slot
collector-name collector-name
instance-name instance-name
```

Description

Clear statistics for inline monitoring services.

Options

collector-name ***collector-name***	Clear collector level statistics.
fpc-slot ***fpc-slot***	Clear statistics for the specified FPC slot. • **Range:** 0 through 11
instance-name ***instance-name***	Clear instance level statistics.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 19.4R1.

Command introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

show services inline-monitoring statistics fpc-slot | 25588

Understanding Inline Monitoring Services

clear services inline nat pool

IN THIS SECTION

- Syntax | 4106
- Description | 4106
- Options | 4107
- Required Privilege Level | 4107
- Output Fields | 4107
- Sample Output | 4107
- Release Information | 4107

Syntax

```
clear services inline nat pool pool-name
```

Description

Clear global inline NAT statistics.

Options

pool-name Name of the NAT pool for which statistic are cleared.

Required Privilege Level

clear

Output Fields

When you enter this command, the NAT pool statistics are cleared. There is no specific output.

Sample Output

clear services inline nat pool

```
user@host> clear services inline nat pool p1
```

Release Information

Command introduced in Junos OS Release 11.4.

clear services inline nat statistics

IN THIS SECTION

- Syntax | 4108
- Description | 4108
- Options | 4108
- Required Privilege Level | 4108
- Output Fields | 4108
- Sample Output | 4108
- Release Information | 4108

Syntax

```
clear services inline nat statistics
<interface interface-name>
```

Description

Clear global inline NAT statistics.

Options

interface *interface-name* (Optional) Clear inline NAT statistics for the specified interface only.

Required Privilege Level

clear

Output Fields

When you enter this command, the global inline NAT statistics are cleared. There is no specific output.

Sample Output

clear services inline nat statistics

```
user@host> clear services inline nat statistics
```

Release Information

Command introduced in Junos OS Release 11.4.

clear services inline softwire statistics

IN THIS SECTION

- Syntax | 4109
- Description | 4109
- Options | 4110
- Required Privilege Level | 4110
- Output Fields | 4110
- Sample Output | 4110
- Release Information | 4110

Syntax

```
clear services inline softwire statistics
<interface interface-name>
```

Description

Clear global inline softwire statistics.

NOTE: The following two limitations apply to the clearing of data plane statistics using the `clear services inline softwire statistics` command:

- When traffic is continuously flowing and the counters are being updated in the data plane, none of the statistical values except the counter for 6rd decapsulation errors is reset.
- When you delete the softwire concentrator or the service set associated with an inline services (`si-`) interface, the counter for 6rd decapsulation errors might display all the previously accumulated values.

Options

interface *interface-name* (Optional) Clear inline softwire statistics for the specified interface only.

Required Privilege Level

clear

Output Fields

When you enter this command, the global inline softwire statistics are cleared. There is no specific output.

Sample Output

clear services inline softwire statistics

```
user@host> clear services inline softwire statistics
```

Release Information

Command introduced in Junos OS Release 13.3R3.

clear services ipsec-vpn certificates

IN THIS SECTION

- Syntax | 4111
- Description | 4111
- Options | 4111
- Required Privilege Level | 4111
- Output Fields | 4111
- Sample Output | 4111

Release Information | 4112

Syntax

```
clear services ipsec-vpn certificates (all | service-set service-set)
<certificate-cache-entry number>
```

Description

(Adaptive services interfaces only) Delete digital certificates from the IPsec configuration memory cache. Issuing this command also clears the certificate revocation list (CRL) from the cache along with the certificates.

Options

all	Delete digital certificates for all service sets.
service-set *service-set*	Delete digital certificates for the specified service set.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services ipsec-vpn certificates all

```
user@host> clear services ipsec-vpn certificates all
```

Release Information

Command introduced in Junos OS Release 7.5.

clear services ipsec-vpn ike security-associations

IN THIS SECTION

- Syntax | 4112
- Description | 4112
- Options | 4112
- Required Privilege Level | 4113
- Output Fields | 4113
- Sample Output | 4113
- Release Information | 4113

Syntax

```
clear services ipsec-vpn ike security-associations
<peer-address-name>
<service-set service-set-name>
```

Description

(Adaptive services interfaces only) Clear Internet Key Exchange (IKE) security associations.

Options

peer-address-name	(Optional) Clear only the security association specified by the peer address.
service-set *service-set-name*	(Optional) Clear only the security association specified by the service-set name.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services ipsec-vpn ike security-associations

```
user@host> clear services ipsec-vpn ike security-associations
```

Release Information

Command introduced before Junos OS Release 7.4.

`service-set` option added in Junos OS Release 8.5.

RELATED DOCUMENTATION

show services ipsec-vpn ike security-associations | 25615

clear services ipsec-vpn ipsec security-associations

IN THIS SECTION

- Syntax | 4114
- Description | 4114
- Options | 4114
- Required Privilege Level | 4114
- Output Fields | 4114
- Sample Output | 4115

Release Information | 4115

Syntax

```
clear services ipsec-vpn security-associations
<peer-address-name>
<remote-gateway remote-gateway-address>
<service-set-name>
<tunnel-index tunnel-index-number>
```

Description

(Adaptive services interfaces only) Clear IP Security (IPsec) security associations. You can combine the options for greater specificity.

Options

peer-address-name	(Optional) Clear only the security association specified by the peer address.
remote-gateway *remote-gateway-address*	(Optional) Clear only the security association specified by the remote gateway address.
service-set-name	(Optional) Clear only the security association specified by the service-set name.
tunnel-index *tunnel-index-number*	(Optional) Clear only the security association specified by the tunnel index number.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services ipsec-vpn ipsec security-associations

```
user@host> clear services ipsec-vpn ipsec security-associations
```

Release Information

Command introduced before Junos OS Release 7.4.

`remote-gateway`, *`service-set-name`*, and `tunnel-index` options added in Junos OS Release 8.4.

RELATED DOCUMENTATION

show services ipsec-vpn ipsec security-associations | 25622

clear services ipsec-vpn ipsec statistics

IN THIS SECTION

- Syntax | 4116
- Description | 4116
- Options | 4116
- Required Privilege Level | 4116
- Output Fields | 4116
- Sample Output | 4116
- Release Information | 4116

Syntax

```
clear services ipsec-vpn ipsec statistics
<remote-gateway address>
<service-set service-set-name>
```

Description

(Adaptive services interface only) Clear IP Security (IPsec) statistics.

Options

remote-gateway *address*	(Optional) Clear statistics for the specified remote system.
service-set *service-set-name*	(Optional) Clear statistics for the specified service set.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services ipsec-vpn ipsec statistics

```
user@host> clear services ipsec-vpn ipsec statistics
```

Release Information

Command introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

show services ipsec-vpn ipsec statistics | 25632

clear services l2tp destination

IN THIS SECTION

- Syntax | 4117
- Description | 4117
- Options | 4118
- Required Privilege Level | 4118
- Output Fields | 4118
- Sample Output | 4119
- Release Information | 4119

Syntax

```
clear services l2tp destination
<all | local-gateway gateway-address | peer-gateway gateway-address>
```

Description

Clear all Layer 2 Tunneling Protocol (L2TP) destinations and all tunnels and sessions that belong to the destinations. This command is available only for LAC on MX Series routers.

NOTE: You cannot issue the `clear services l2tp destination` command in parallel with statistics-related `show services l2tp` commands from separate terminals. If this `clear` command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the `show` commands listed in the following table:

`show services l2tp destination extensive`	`show services l2tp summary statistics`

`show services l2tp destination statistics`	`show services l2tp tunnel extensive`
`show services l2tp session extensive`	`show services l2tp tunnel statistics`
`show services l2tp session statistics`	

Options

all

Close all L2TP destinations.

> **BEST PRACTICE**: The `all` option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the `all` option in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

local-gateway *gateway-address*

Clear only the L2TP destinations and all tunnels and sessions associated with the specified local gateway address.

peer-gateway *gateway-address*

Clear only the L2TP destinations and all tunnels and sessions associated with the peer gateway with the specified address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services l2tp destination all

```
user@host> clear services l2tp destination all

  Destination 2 closed
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

show services l2tp destination | 25641

clear services l2tp destination lockout

IN THIS SECTION

- Syntax | 4120
- Description | 4120
- Options | 4120
- Required Privilege Level | 4121
- Output Fields | 4121
- Sample Output | 4121
- Release Information | 4121

Syntax

```
clear services l2tp destination lockout
<all | local-gateway gateway-address | peer-gateway gateway-address>
```

Description

Clear the lockout timer for all or only the specified Layer 2 Tunneling Protocol (L2TP) destinations and all tunnels and sessions that belong to the destinations. Clearing the lockout timer removes the destination from the lockout list. This command is available only for LAC on MX Series routers.

NOTE: You cannot issue the `clear services l2tp destination` command in parallel with statistics-related `show services l2tp` commands from separate terminals. If this `clear` command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the `show` commands listed in the following table:

`show services l2tp destination extensive`	`show services l2tp summary statistics`
`show services l2tp destination statistics`	`show services l2tp tunnel extensive`
`show services l2tp session extensive`	`show services l2tp tunnel statistics`
`show services l2tp session statistics`	

Options

all (Optional) Unlock all L2TP destinations.

local-gateway *gateway-address* (Optional) Unlock only the L2TP destination with the specified local gateway address.

peer-gateway *gateway-address* (Optional) Unlock only the L2TP destination with the specified address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear services l2tp destination lockout all

```
user@host> clear services l2tp destination lockout all
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

clear services l2tp destination | 4117
show services l2tp destination | 25641

clear services l2tp destination statistics

IN THIS SECTION

- Syntax | 4122
- Description | 4122
- Options | 4122
- Required Privilege Level | 4122
- Output Fields | 4122
- Sample Output | 4123

Release Information | 4123

Syntax

```
clear services l2tp destination statistics
<all | local-gateway gateway-address | peer-gateway gateway-address >
```

Description

Clear all statistics associated with the Layer 2 Tunneling Protocol (L2TP) destinations and all tunnels and sessions that belong to the destinations. This command is available only for LAC on MX Series routers.

Options

all	Clear all statistics associated with the L2TP destinations.
local-gateway *gateway-address*	Clear statistics related to L2TP destination and all tunnels and sessions associated with the specified local gateway address.
peer-gateway *gateway-address*	Clear statistics related to L2TP destination and all tunnels and sessions associated with the specified peer gateway address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services l2tp destination statistics

```
user@host>clear services l2tp destination statistics all
Destination 1 statistics cleared
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

show services l2tp destination | 25641

clear services l2tp multilink

IN THIS SECTION

- Syntax | 4123
- Description | 4124
- Options | 4124
- Required Privilege Level | 4124
- Output Fields | 4124
- Sample Output | 4124
- Release Information | 4124

Syntax

```
clear services l2tp multilink (all <statistics> | bundle-id number <statistics> |
statistics (all | bundle-id number))
```

Description

(M10i and M7i routers only) Close Layer 2 Tunneling Protocol (L2TP) multilink sessions or clear session statistics.

Options

all <statistics>	Close all L2TP multilink sessions or clear statistics for all L2TP multilink sessions.
bundle-id *number* <statistics>	L2TP multilink bundle ID. The value is an internally generated number from 1 to 65535. Close the specified L2TP multilink session, or using the **statistics** keyword with this option, clear statistics for the specified session.
statistics (all | bundle-id *number*)	Clear all session statistics or clear statistics for the specified multilink bundle ID.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services l2tp multilink statistics all

```
user@host> clear services l2tp multilink statistics all
  Multilink 1 statistics cleared
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

L2TP Services Configuration Overview

L2TP Minimum Configuration

show services l2tp multilink | 25649

clear services l2tp session

IN THIS SECTION

- Syntax | 4125
- Description | 4125
- Options | 4126
- Required Privilege Level | 4127
- Output Fields | 4127
- Sample Output | 4127
- Sample Output | 4128
- Release Information | 4128

Syntax

```
clear services l2tp session (all | interface interface-name | local-gateway gateway-address |
local-gateway-name gateway-name | local-session-id session-id  | local-tunnel-id tunnel-id |
peer-gateway gateway-address | peer-gateway-name gateway-name | routing-instance routing-
instance-name | tunnel-group group-name | user username)
```

Description

(M10i and M7i routers only) Clear Layer 2 Tunneling Protocol (L2TP) sessions on LNS.

(MX Series routers only) Clear L2TP sessions on LAC and LNS.

NOTE: On MX Series routers, you cannot issue the `clear services l2tp session` command in parallel with statistics-related `show services l2tp` commands from separate terminals. If this `clear` command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the `show` commands listed in the following table:

Table 13: Show Commands

`show services l2tp destination extensive`	`show services l2tp summary statistics`
`show services l2tp destination statistics`	`show services l2tp tunnel extensive`
`show services l2tp session extensive`	`show services l2tp tunnel statistics`
`show services l2tp session statistics`	

Options

all

Close all L2TP sessions.

BEST PRACTICE: The `all` option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the `all` option in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

interface *interface-name*

Clear only the L2TP sessions using the specified adaptive services or inline services interface. The interface type depends on the line card as follows:

- `si-`*`fpc/pic/port`*—MPCs on MX Series routers only. This option is not available for L2TP on M Series routers.
- `sp-`*`fpc/pic/port`*—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers.

local-gateway *gateway-address*

Clear only the L2TP sessions associated with the specified local gateway address.

local-gateway-name *gateway-name*	Clear only the L2TP sessions associated with the specified local gateway name.
local-session-id *session-id*	Clear only the L2TP sessions with this identifier for the local endpoint of the L2TP session.
local-tunnel-id *tunnel-id*	Clear only the L2TP sessions associated with the specified local tunnel identifier.
peer-gateway *gateway-address*	Clear only the L2TP sessions associated with the peer gateway with the specified address.
peer-gateway-name *gateway-name*	Clear only the L2TP sessions associated with the peer gateway with the specified name.
routing-instance *routing-instance-name*	Clear only the L2TP sessions associated with the specified routing instance.
tunnel-group *group-name*	Clear only the L2TP sessions associated with the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.
user *username*	(M Series routers only) Clear only the L2TP sessions for the specified username.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services l2tp session

```
user@host> clear services l2tp session 31694
  Session 31694 closed
```

Sample Output

clear services l2tp session interface

```
user@host> show services l2tp session Tunnel local ID: 17185
Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
  5117   1        Established          1073741828        si-2/0/0
  34915  2        Established          1073741829        si-2/1/0
  6454   3        Established          1073741830        si-2/0/0
  46142  4        Established          1073741831        si-2/1/0
```

command-name

```
user@host> clear services l2tp session interface si-2/0/0
  Session  5117 closed
  Session  6454 closed
```

command-name

```
user@host> show services l2tp session Tunnel local ID: 17185
  Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
  34915  2        Established          1073741829        si-2/1/0
  46142  4        Established          1073741831        si-2/1/0
```

Release Information

Command introduced before Junos OS Release 7.4.

routing-instance *routing-instance-name* option introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

- L2TP Services Configuration Overview
- L2TP Minimum Configuration

clear services l2tp session statistics
show services l2tp session | 25661

clear services l2tp session statistics

IN THIS SECTION

- Syntax | 4129
- Description | 4129
- Options | 4129
- Required Privilege Level | 4130
- Output Fields | 4130
- Sample Output | 4131
- Release Information | 4131

Syntax

```
clear services l2tp session statistics (all | interface interface-name | local-gateway gateway-address | local-gateway-name gateway-name | local-session-id session-id | local-tunnel-id tunnel-id | peer-gateway gateway-address | peer-gateway-name gateway-name  | tunnel-group group-name | user username)
```

Description

(M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Clear statistics for Layer 2 Tunneling Protocol (L2TP) sessions.

Options

all — Clear statistics for all L2TP sessions.

interface *interface-name* — Clear only the L2TP sessions using the specified adaptive services or inline services interface. The interface type depends on the line card as follows:

- `si-`*`fpc/pic/port`*—MPCs on MX Series routers only. This option is not available for L2TP on M Series routers.
- `sp-`*`fpc/pic/port`*—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers.

local-gateway ***gateway-address***	Clear statistics for only the L2TP sessions associated with the local gateway with the specified address.
local-gateway-name ***gateway-name***	Clear statistics for only the L2TP sessions associated with the local gateway with the specified name.
local-session-id ***session-id***	Clear statistics for only the L2TP sessions with this identifier for the local endpoint of the L2TP session.
local-tunnel-id ***tunnel-id***	Clear statistics for only the L2TP sessions associated with the specified local tunnel identifier.
peer-gateway ***gateway-address***	Clear statistics for only the L2TP sessions associated with the peer gateway with the specified address.
peer-gateway-name ***gateway-name***	Clear statistics for only the L2TP sessions associated with the peer gateway with the specified name.
tunnel-group ***group-name***	Clear statistics for only the L2TP sessions associated with the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.
user ***username***	Clear statistics for only the L2TP sessions for the specified username. This option is not available for L2TP LAC on MX Series routers.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services l2tp session statistics all

```
user@host> clear services l2tp session statistics all
  Session 26497 statistics cleared
```

Release Information

Command introduced before Junos OS Release 7.4.

Support for MX Series routers added in Junos OS Release 10.4.

RELATED DOCUMENTATION

- L2TP Services Configuration Overview
- L2TP Minimum Configuration
- clear services l2tp session | 4125
- show services l2tp session | 25661

clear services l2tp tunnel

IN THIS SECTION

- Syntax | 4132
- Description | 4132
- Options | 4132
- Required Privilege Level | 4133
- Output Fields | 4133
- Sample Output | 4133
- Release Information | 4133

Syntax

```
clear services l2tp tunnel (all | interface sp-fpc/pic/port | local-gateway gateway-address |
local-gateway-name gateway-name | local-tunnel-id tunnel-id | peer-gateway gateway-address |
peer-gateway-name gateway-name | tunnel-group group-name)
```

Description

(M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Clear Layer 2 Tunneling Protocol (L2TP) tunnels.

NOTE: On MX Series routers, you cannot issue the `clear services l2tp tunnel` command in parallel with statistics-related `show services l2tp` commands from separate terminals. If this `clear` command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the `show` commands listed in the following table:

`show services l2tp destination extensive`	`show services l2tp summary statistics`
`show services l2tp destination statistics`	`show services l2tp tunnel extensive`
`show services l2tp session extensive`	`show services l2tp tunnel statistics`
`show services l2tp session statistics`	

Options

all Clear all L2TP tunnels.

BEST PRACTICE: The `all` option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the `all` option in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

sp-*fpc/pic/port*	(Optional) Clear only the L2TP tunnels using the specified adaptive services interface. This option is not available for L2TP on MX Series routers.
local-gateway *gateway-address*	Clear only the L2TP tunnels associated with the local gateway with the specified address.
local-gateway-name *gateway-name*	Clear only the L2TP tunnels associated with the local gateway with the specified name.
local-tunnel-id *tunnel-id*	Clear only the L2TP tunnels that have the specified local tunnel identifier.
peer-gateway *gateway-address*	Clear only the L2TP tunnels associated with the peer gateway with the specified address.
peer-gateway-name *gateway-name*	Clear only the L2TP tunnels associated with the peer gateway with the specified name.
tunnel-group *group-name*	Clear only the L2TP tunnels in the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services l2tp tunnel

```
user@host> clear services l2tp tunnel 17185

  Tunnel 17185 closed
```

Release Information

Command introduced before Junos OS Release 7.4.

Support for LAC on MX Series routers introduced in Junos OS Release 10.4.

Support for LNS on MX Series routers introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- L2TP Services Configuration Overview
- L2TP Minimum Configuration
- clear services l2tp tunnel statistics | 4134
- show services l2tp tunnel | 25685

clear services l2tp tunnel statistics

IN THIS SECTION

- Syntax | 4134
- Description | 4134
- Options | 4135
- Required Privilege Level | 4135
- Output Fields | 4135
- Sample Output | 4135
- Release Information | 4136

Syntax

```
clear services l2tp tunnel statistics (all | interface sp-fpc/pic/port | local-gateway gateway-address | local-gateway-name gateway-name | local-tunnel-id tunnel-id | peer-gateway gateway-address | peer-gateway-name gateway-name | tunnel-group group-name)
```

Description

(M10i and M7i routers: LNS only. MX Series routers: LAC only.) Clear statistics for Layer 2 Tunneling Protocol (L2TP) tunnels.

Options

all	Clear statistics for all L2TP tunnels.
interface sp-*fpc/pic/port*	Clear statistics for only the L2TP tunnels using the specified adaptive services interface. This option is not available for L2TP LAC on MX Series routers.
local-gateway *gateway-address*	Clear statistics for only the L2TP tunnels associated with the local gateway with the specified address.
local-gateway-name *gateway-name*	Clear statistics for only the L2TP tunnels associated with the local gateway with the specified name.
local-tunnel-id *tunnel-id*	Clear statistics for only the L2TP tunnels that have the specified local tunnel identifier.
peer-gateway *gateway-address*	Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified address.
peer-gateway-name *gateway-name*	Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified name.
tunnel-group *group-name*	Clear statistics for only the L2TP tunnels in the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services l2tp tunnel statistics all

```
user@host> clear services l2tp tunnel statistics all
  Tunnel  9933 statistics cleared
```

Release Information

Command introduced before Junos OS Release 7.4.

Support for MX Series routers added in Junos OS Release 10.4.

RELATED DOCUMENTATION

- L2TP Services Configuration Overview
- L2TP Minimum Configuration
- clear services l2tp tunnel
- show services l2tp tunnel

clear services lrf collector statistics

IN THIS SECTION

- Syntax | 4136
- Description | 4137
- Options | 4137
- Required Privilege Level | 4137
- Output Fields | 4137
- Sample Output | 4137
- Release Information | 4137

Syntax

```
clear services lrf collector statistics
<collector-name>
```

Description

Clear all the LRF statistics for the specified collector. If a collector is not specified, statistics are cleared for all collectors.

Options

none	Clear LRF statistics for all collectors.
collector-name	(Optional) Clear LRF statistics for the specified collector.

Required Privilege Level

clear

Output Fields

A message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

clear services lrf collector statistics

```
user@host> clear services lrf collector statistics coll1
Interface: ms-0/1/0, Status: LRF collector statistics successfully cleared
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show services lrf collector statistics | 25730

clear services lrf statistics

IN THIS SECTION

- Syntax | 4138
- Description | 4138
- Required Privilege Level | 4138
- Output Fields | 4138
- Sample Output | 4139
- Release Information | 4139

Syntax

```
clear services lrf statistics
```

Description

Clear all the LRF statistics.

Required Privilege Level

clear

Output Fields

A message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

clear services lrf statistics

```
user@host> clear services lrf statistics
Interface: ms-0/1/0, Status: LRF statistics successfully cleared
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show services lrf statistics | 25735

clear services mobile-edge sessions

IN THIS SECTION

- Syntax | 4139
- Description | 4139
- Required Privilege Level | 4140
- Release Information | 4140

Syntax

```
clear services mobile-edge sessions
<local-seid>
```

Description

Clear mobile-edge sessions with the session ID.

Required Privilege Level

system

Release Information

Command introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

No Link Title

clear services monitoring rfc2544

IN THIS SECTION

- Syntax | 4140
- Description | 4140
- Options | 4141
- Required Privilege Level | 4141
- Output Fields | 4141
- Release Information | 4141

Syntax

```
clear services monitoring rfc2544
<active-tests | all-tests | completed-tests | terminated-tests | test-id test-id>
```

Description

Clear the indicated type of RFC 2544 benchmarking test sessions.

Options

active-tests	Clears the active RFC 2544 benchmarking test sessions.
all-tests	Clears all RFC 2544 benchmarking test sessions.
completed-tests	Clears the completed RFC 2544 benchmarking test sessions.
terminated-tests	Clears the terminated RFC 2544 benchmarking test sessions.
test-id ***test-id***	Clears the particular RFC 2544 benchmarking test by specifying its test ID.

Required Privilege Level

clear

Output Fields

When you enter this command, if it is successful, there is no response. If it is not successful, the system displays an error message if the FPC for the interface specified on the `test-interface` statement reports an error.

Release Information

Command introduced in Junos OS Evolved Release 21.1R1.

`all-tests` and `test-id` options introduced in Junos OS Evolved 22.4R1.

clear services monitoring twamp server control-connection

IN THIS SECTION

- Syntax | 4142
- Description | 4142
- Options | 4142
- Required Privilege Level | 4142
- Release Information | 4142

Syntax

```
clear services monitoring twamp server control-connection
<connection-name>
```

Description

Clear connections established between the Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, all established connections are cleared, along with the sessions on those connections. To clear only a particular connection, specify the connection name when you issue the command.

Options

connection-name	(Optional) Specify which connection to clear.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

clear services nat flows

IN THIS SECTION

- Syntax | 4143
- Description | 4143

- Options | 4143
- Required Privilege Level | 4143
- Output Fields | 4143
- Sample Output | 4144
- Release Information | 4144

Syntax

```
clear services nat flows
<b4address b4address>
<service-set service-set>
<subscriber subscriber-address>
```

Description

Clear NAT flows.

Options

none	Clear all NAT flows.
b4address *b4address*	(Optional) Clear NAT flows for a particular B4 address.
service-set *service-set*	(Optional) Clear NAT flows for a particular service set.
subscriber *ip*	(Optional) Clear NAT flows for a particular subscriber, identified by IPv4 address.

Required Privilege Level

view

Output Fields

Table 14 on page 4144 lists the output fields for the `clear services nat flows` command. Output fields are listed in the approximate order in which they appear.

Table 14: clear services nat flows Output Fields

Field Name	Field Description
Interface	Name of a services interface.
Service set	Name of the service set from which flows are being cleared.
Flows removed	Number of flows removed.

Sample Output

clear services nat flows subscriber (IPv4 address)

```
user@host> clear services nat flows subscriber ip 198.51.100.3
Interface   Service set  Flows removed

sp-2/0/0    ss1               0
```

Release Information

Command introduced in Junos OS Release 14.1.

clear services nat mappings

IN THIS SECTION

- Syntax | 4145
- Description | 4145
- Options | 4145
- Required Privilege Level | 4145
- Output Fields | 4145

Sample Output | 4146
Release Information | 4146

Syntax

```
clear services nat mappings
<app>
<eim>
<pcp>
<service-set service-set>
```

Description

Clear NAT mappings.

Options

none	Clear all NAT mappings.
app	(Optional) Clear address-pooling paired NAT mappings.
eim	(Optional) Clear endpoint-independent NAT mappings.
pcp	(Optional) Clear Port Control Protocol NAT mappings.
service-set *service-set*	(Optional) Clear NAT mappings for a specified service set..

Required Privilege Level

clear

Output Fields

Table 15 on page 4146 lists the output fields for the `clear services nat mappings` command. Output fields are listed in the approximate order in which they appear.

Table 15: clear services nat mappings Output Fields

Field Name	Field Description
Interface	Name of a services interface.
Service set	Name of the service set from which flows are being cleared.
Mappings removed	Number of mappings removed.
Flows removed	Number of flows removed.

Sample Output

clear services nat mappings

```
user@host> clear services nat mappings
Interface   Service set                        Mappings removed   Flows removed
sp-2/0/0    ss1                                               0               0
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

show services nat source mappings address-pooling-paired | 25864

clear services nat mappings app | 4147

clear services nat mappings eim | 4149

clear services nat mappings pcp | 4151

clear services nat mappings app

IN THIS SECTION

- Syntax | 4147
- Description | 4147
- Options | 4147
- Required Privilege Level | 4148
- Output Fields | 4148
- Sample Output | 4148
- Release Information | 4148

Syntax

```
clear services nat mappings app
<b4address b4address/prefix>
<service-set service-set>
<subscriber subscriber-ipv4-address>
```

Description

Clear NAT mappings for address pooling paired (app).

Options

none	Clear all NAT app mappings.
b4address *b4address/prefix*	(Optional) Clear NAT APP mappings for a particular subscriber b4address/prefix
service-set *service-set*	(Optional) Clear NAT APP mappings for a specified service set..
subscriber *subscriber-ipv4-address/prefix*	(Optional) Clear NAT APP mappings for a particular subscriber ipv4-address/prefix

Required Privilege Level

clear

Output Fields

Table 16 on page 4148 lists the output fields for the `clear services nat mappings app` command. Output fields are listed in the approximate order in which they appear.

Table 16: clear services nat mappings app Output Fields

Field Name	Field Description
Interface	Name of a services interface.
Service set	Name of the service set from which flows are being cleared.
Mappings removed	Number of mappings removed.
Flows removed	Number of flows removed.

Sample Output

clear services nat mappings app

```
user@host> clear services nat mappings app
Interface   Service set                       Mappings removed   Flows removed
sp-2/0/0    ss1                                             0               0
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

show services nat source mappings address-pooling-paired | 25864

clear services nat mappings eim

IN THIS SECTION

- Syntax | 4149
- Description | 4149
- Options | 4149
- Required Privilege Level | 4150
- Output Fields | 4150
- Sample Output | 4151
- Release Information | 4151

Syntax

```
clear services nat mappings eim
<b4address b4address/prefix>
<subscriber subscriber-ipv4-address>
```

Description

Clear endpoint independent (EIM) and port control protocol (PCP) mappings .

Options

none Clear all EIM and PCP mappings.

b4address *b4address/prefix* (Optional) Clear EIM and PCP mappings for a particular subscriber b4address/prefix

- **internal-host** *ipv4address/prefix* (Optional) Clear EIM and PCP mappings matching the specified b4address and internal-host..
  - **port** *port* (Optional) Clear EIM and PCP mappings matching the specified b4address, internal host, and port.

service-set ***service-set*** (Optional) Clear EIM and PCP mappings for the specified service set.

subscriber ***subscriber-ipv4-address/prefix*** (Optional) Clear EIM and PCP mappings for a particular subscriber ipv4-address/prefix

- **port** ***port*** (Optional) Clear EIM and PCP mappings matching the specified ipv4-address/prefix and port.
- **service-set** ***service-set*** (Optional) Clear EIM and PCP mappings for the specified service set.

Required Privilege Level

clear

Output Fields

Table 17 on page 4150 lists the output fields for the `clear services nat mappings eim` command. Output fields are listed in the approximate order in which they appear.

Table 17: clear services nat mappings eim Output Fields

Field Name	Field Description
Interface	Name of a services interface.
Service set	Name of the service set from which flows are being cleared.
Mappings removed	Number of mappings removed.
Flows removed	Number of flows removed.

Sample Output

clear services nat mappings eim

```
user@host> clear services nat mappings eim
Interface    Service set                        Mappings removed   Flows removed
sp-2/0/0     ss1                                               0               0
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

show services nat source mappings address-pooling-paired | 25864

clear services nat mappings pcp

IN THIS SECTION

- Syntax | 4152
- Description | 4152
- Options | 4152
- Required Privilege Level | 4152
- Output Fields | 4152
- Sample Output | 4153
- Release Information | 4153

Syntax

```
clear services nat mappings pcp
<b4address b4address/prefix>
<subscriber subscriber-ipv4-address>
```

Description

Clear NAT mappings for Port Control Protocol (PCP).

Options

none	Clear all NAT PCP mappings.	
b4address *b4address/prefix*	(Optional) Clear NAT PCP mappings for a particular subscriber b4address/prefix	
	port *port*	(Optional) Clear NAT PCP mappings matching the specified b4address internal host, and port.
	service-set *service-set*	(Optional) Clear NAT PCP mappings for the specified service set.
subscriber *ipv4-address/prefix*	(Optional) Clear NAT PCP mappings for a particular subscriber ipv4-address/prefix	
	port *port*	(Optional) Clear NAT PCP mappings matching the specified ipv4-address/prefix, and port.

Required Privilege Level

clear

Output Fields

Table 18 on page 4153 lists the output fields for the `clear services nat mappings pcp` command. Output fields are listed in the approximate order in which they appear.

Table 18: clear services nat mappings pcp Output Fields

Field Name	Field Description
Interface	Name of a services interface.
Service set	Name of the service set from which flows are being cleared.
Mappings removed	Number of mappings removed.
Flows removed	Number of flows removed.

Sample Output

clear services nat mappings pcp

```
user@host> clear services nat mappings pcp
Interface   Service set                        Mappings removed   Flows removed
sp-2/0/0    ss1                                               0               0
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

show services nat source mappings address-pooling-paired | 25864

clear services nat source mappings

IN THIS SECTION

- Syntax | 4154
- Description | 4154
- Options | 4154
- Required Privilege Level | 4155
- Output Fields | 4155
- Sample Output | 4155
- Release Information | 4156

Syntax

```
clear services nat source mappings
<app | eim | pcp>
subscriber private-ip [port port-num]  [service-set service-set]
```

Description

Clear services NAT source mappings. After one mapping is cleared, all the port block alloation blocks referring to that mapping are released.

Options

app	Clear all APP mappings.
app subscriber *private-ip* [port *port-num*] [service-set *service-set*]	Clear one APP mapping by matching conditions
eim	Clear all EIM mappings.
eim subscriber *private-ip* [port *port-num*] [service-set *service-set*]	Clear one EIM mapping by matching conditions
pcp	Clear all PCP mappings.

Required Privilege Level

view

Output Fields

Table 19 on page 4155 lists the output fields for the `clear services nat source mappings` command. Output fields are listed in the approximate order in which they appear.

Table 19: clear services nat source mappings Output Fields

Field Name	Field Description
NAT pool	Name of the NAT pool.
Mappings removed	Number of mappings removed.
Sessions removed	Number of sessions removed.

Sample Output

clear services nat source mappings eim

```
user@host> clear services nat source mappings eim
NAT pool                        Mappings removed   Sessions removed
Test-pool                                      1                  0
```

clear service nat source mappings eim subscriber 2.1.1.1

```
user@host> clear service nat source mappings eim subscriber 2.1.1.1
NAT pool                        Mappings removed   Sessions removed
Test-pool                                      1                  0
```

clear services nat source mappings subscriber 2.1.1.1 port 1026 service-set ss1

```
user@host> clear services nat source mappings subscriber 2.1.1.1 port 1026 service-set
ss1
NAT pool                        Mappings removed   Sessions removed
Test-pool                                                1                  0
```

clear services nat source mappings app

```
user@host> clear services nat source mappings app
  NAT pool                        Mappings removed   Sessions removed
Test-pool                                                1                  0
```

clear services nat source mappings app subscriber 2.1.1.1

```
user@host> clear services nat source mappings app subscriber 2.1.1.1
  NAT pool                        Mappings removed   Sessions removed
Test-pool                                                1                  0
```

clear services nat source mappings app subscriber 2.1.1.1 port 1026 service-set ss1

```
user@host> clear services nat source mappings app subscriber 2.1.1.1 port 1026 service-set
ss1
  NAT pool                        Mappings removed   Sessions removed
Test-pool                                                1                  0
```

Release Information

Command introduced in Junos OS Release 19.3R2.

clear services redundancy-set last-saved-state id

IN THIS SECTION

- Syntax | 4157
- Description | 4157
- Required Privilege Level | 4157
- Release Information | 4157

Syntax

```
clear services redundancy-set last-saved-state id
 (redundancy-set{x | all})
```

Description

Clear the last saved state of a redundancy set or of all redundancy sets.

You can clear the warnings on the standby after inspection by issuing the `clear services redundancy-set {x|all} standby-warning` command on it. The command will also send an update notification regarding this peer's clean state to all the other peers in that redundancy set.

Required Privilege Level

clear

Release Information

Statement introduced in Junos OS Release 16.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-MX-SPC3 services card.

clear services rpm rfc2544-benchmarking

IN THIS SECTION

- Syntax | 4158
- Description | 4158
- Options | 4158
- Additional Information | 4158
- Required Privilege Level | 4159
- Output Fields | 4159
- Release Information | 4159

Syntax

```
clear services rpm rfc2544-benchmarking
<aborted-tests | active-tests | completed-tests>
```

Description

Clear the indicated type of RFC 2544 benchmarking test sessions.

Options

aborted-tests Clears the terminated RFC 2544 benchmarking test sessions.

active-tests Clears the active RFC 2544 benchmarking test sessions.

completed-tests Clears the completed RFC 2544 benchmarking test sessions.

Additional Information

Required Privilege Level

clear

Output Fields

When you enter this command, if it is successful, there is no response. If it is not successful, the system displays an error message if the FPC for the interface specified on the `test-interface` statement reports an error.

Release Information

Command introduced in Junos OS Release 12.3X52.

clear services rpm twamp server connection

IN THIS SECTION

- Syntax | 4159
- Description | 4159
- Options | 4160
- Required Privilege Level | 4160
- Release Information | 4160

Syntax

```
clear services rpm twamp server connection <connection-id>
```

Description

Clear connections established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default all established connections are cleared (along with the sessions on those connections). To clear only a specific connection, specify the connection ID when you issue the command.

Options

connection-id	(Optional) Specific connection to clear.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 9.3.

clear services security-intelligence dns-statistics

IN THIS SECTION

- Syntax | 4160
- Description | 4160
- Required Privilege Level | 4161
- Output Fields | 4161
- Sample Output | 4161
- Release Information | 4161

Syntax

```
clear services security-intelligence dns-statistics
clear services security-intelligence dns-statistics logical-system all
clear services security-intelligence dns-statistics logical-system logical-system-name
```

Description

Sets DNS statistics to zero. See "show services security-intelligence dns-statistics" on page 26037 for information on the DNS statistics.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear services security-intelligence dns-statistics

```
user@host> clear services security-intelligence dns-statistics
```

clear services security-intelligence dns-statistics logical-system

```
user@host> clear services security-intelligence dns-statistics logical-system all
user@host> clear services security-intelligence dns-statistics logical-system logical-system-name
```

Release Information

dns-statistics command introduced in Junos OS Release 20.4R1

RELATED DOCUMENTATION

security-intelligence(services) | 16469

clear services security-metadata-streaming

IN THIS SECTION

- Syntax | 4162
- Description | 4162

- Required Privilege Level | 4162
- Output Fields | 4162
- Sample Output | 4163
- Release Information | 4163

Syntax

```
clear services security-metadata-streaming dns statistics
clear services security-metadata-streaming dns statistics logical-system <logical-system-name>
clear services security-metadata-streaming dns statistics logical-system all
clear services security-metadata-streaming dns statistics tenant <tenant-name>
clear services security-metadata-streaming dns statistics tenant all
clear services security-metadata-streaming http statistics
```

Description

Sets the dns and http security-metadata-streaming statistics to 0. See "show services security-metadata-streaming" on page 26042 for information on these statistics.

Use this command when you change the Juniper Advanced Threat Prevention Cloud profile or policy settings and want to reset the statistics in order to view the effects of the changes.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear services security-metadata-streaming dns statistics

```
user@host> clear services security-metadata-streaming dns statistics
```

clear services security-metadata-streaming dns statistics logical-system <logical-system-name>

```
user@host> clear services security-metadata-streaming dns statistics logical-system <logical-
system-name>
```

clear services security-metadata-streaming dns statistics logical-system all

```
user@host> clear services security-metadata-streaming dns statistics logical-system all
```

clear services security-metadata-streaming dns statistics tenant <tenant-name>

```
user@host> clear services security-metadata-streaming dns statistics tenant <tenant-name>
```

clear services security-metadata-streaming dns statistics tenant all

```
user@host> clear services security-metadata-streaming dns statistics tenant all
```

clear services security-metadata-streaming http statistics

```
user@host> clear services security-metadata-streaming http statistics
```

Release Information

dns-statistics command introduced in Junos OS Release 20.4R1

clear services service-set statistics ids drops

IN THIS SECTION

- Syntax | 4164
- Description | 4164
- Options | 4164
- Required Privilege Level | 4164
- Sample Output | 4165
- Release Information | 4165

Syntax

```
clear services service-set statistics ids drops
<interface interface-name>
<service-set service-set-name>
```

Description

Clear statistics for packet drops resulting from header-integrity, suspicious packet pattern, and session-limit checks performed by an MS-MPC or MS-MIC.

Options

none	Clear statistics for all configured services interfaces and service sets.
interface ***interface-name***	(Optional) Clear statistics for the specified services interface.
service-set ***service-set-name***	(Optional) Clear statistics for the specified service set.

Required Privilege Level

network

Sample Output

clear services service-set statistics ids drops

```
user@host> clear services service-set statistics ids drops
```

Release Information

Command introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show services service-set statistics ids drops | 26049

clear services service-sets statistics ids session-limits counters

IN THIS SECTION

- Syntax | 4165
- Description | 4166
- Options | 4166
- Required Privilege Level | 4166
- Sample Output | 4166
- Release Information | 4166

Syntax

```
clear services service-sets statistics ids session-limits counters
<interface interface-name>
```

Description

Clear counters for session drops and packet drops resulting from session-limit checks performed by an IDS rule on an MS-MPC or MS-MIC.

Options

none	Clear counters for all configured services interfaces.
interface *interface-name*	(Optional) Clear counters for the specified services interface.

Required Privilege Level

network

Sample Output

clear services service-sets statistics ids session-limits counters

```
user@host> clear services service-sets statistics ids session-limits counters
```

Release Information

Command introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show services service-sets statistics ids session-limits counters | 26084

clear services service-sets statistics integrity-drops

IN THIS SECTION

Syntax | 4167

- Description | 4167
- Options | 4167
- Required Privilege Level | 4167
- Release Information | 4167

Syntax

```
clear services service-sets statistics integrity-drops
<interface interface-name>
```

Description

Clear integrity-drops statistics for one adaptive services interface, for all adaptive services interfaces, or for one service-set.

Options

none	Clear integrity-drops statistics for all configured adaptive service interfaces/service-set.
Service-set *service-set-name*	(Optional) Clear integrity-drops statistics for the specified service-set
interface *interface-name*	(Optional) Clear integrity-drops statistics for the specified adaptive services interface.

Required Privilege Level

network

Release Information

Command introduced in Junos OS Release 13.3

RELATED DOCUMENTATION

show services service-sets statistics packet-drops | 26111

clear services service-sets statistics jflow-log

IN THIS SECTION

- Syntax | 4168
- Description | 4168
- Options | 4168
- Required Privilege Level | 4169
- Output Fields | 4169
- Sample Output | 4169
- Release Information | 4169

Syntax

```
clear services service-sets statistics jflow-log
<service-set service-set-name>
<interface interface-name>
```

Description

Clear flow monitoring log statistics for the logs generated in IPFIX or version 9 format for one services interface or for all services interfaces, and for one named service set or all service sets on the interface or interfaces.

Options

none	Clear flow monitoring log for all configured services interfaces and their service sets.

interface ***interface-name***	(Optional) Clear flow monitoring log statistics for the specified services interface. On M Series, MX Series, and T Series routers, the ***interface-name*** can be **ms-*fpc/pic/port***. It is supported only on MS-MICs and MS-MPCS.
service-set ***service-set-name***	(Optional) Clear flow monitoring log statistics for the specified services interface.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services service-sets statistics jflow-log interface

```
user@host> clear services service-sets statistics jflow-log interface ms-5/0/0
Interface: ms-5/0/0
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

show services service-sets statistics syslog | **26114**

clear services service-sets statistics packet-drops

IN THIS SECTION

- Syntax | 4170
- Description | 4170
- Options | 4170
- Required Privilege Level | 4170
- Output Fields | 4171
- Sample Output | 4171
- Release Information | 4171

Syntax

```
clear services service-sets statistics packet-drops
<interface interface-name>
```

Description

Clear dropped-packet statistics for one adaptive services interface or for all adaptive services interfaces.

Options

none	Clear dropped-packet statistics for all configured adaptive services interfaces.
interface *interface-name*	(Optional) Clear dropped-packet statistics for the specified adaptive services interface. On M Series and T Series routers, the *interface-name* can be ms-*fpc*/*pic*/*port*, sp-*fpc*/*pic*/*port* or rsp*number*.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services service-sets statistics packet-drops

```
user@host> clear services service-sets statistics packet-drops interface sp-5/0/0
Flow collector interface: cp-5/0/0
Interface state: Collecting flows
Statistics cleared successfully
```

Release Information

Command introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

show services service-sets statistics packet-drops | 26111

clear services service-sets statistics syslog

IN THIS SECTION

- Syntax | 4172
- Description | 4172
- Options | 4172
- Required Privilege Level | 4172
- Output Fields | 4172
- Sample Output | 4172
- Release Information | 4173

Syntax

```
clear services service-sets statistics syslog
<service-set service-set-name>
<interface interface-name>
```

Description

Clear system log statistics for one services interface or for all services interfaces, and for one named service set or all service sets on the interface or interfaces.

Options

none	Clear system log for all configured services interfaces and their service sets.
interface ***interface-name***	(Optional) Clear system log statistics for the specified services interface. On M Series, MX Series, and T Series routers, the ***interface-name*** can be **ms-*fpc/pic/port***, **sp-*fpc/pic/port***, or **rsp*number***.
service-set ***service-set-name***	(Optional) Clear system log statistics for the specified services interface.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services service-sets statistics syslog

```
user@host> clear services service-sets statistics syslog interface sp-5/0/0
Flow collector interface: cp-5/0/0
Interface state: Collecting flows
Statistics cleared successfully
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show services service-sets statistics syslog | 26114

clear services sessions

IN THIS SECTION

- Syntax | 4173
- Description | 4174
- Options | 4174
- Required Privilege Level | 4176
- Output Fields | 4176
- Sample Output | 4177
- Release Information | 4177

Syntax

```
clear services sessions
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<ip-action>
<nat number>
<nat-port-overload-index number>
<protocol protocol>
<service-set service-set>
<source-nat-pool source-nat-pool-name>
```

```
<source-port source-port>
<source-prefix source-prefix>
```

Description

Clear services sessions currently active on the embedded PIC or MIC. When you enter this command, the sessions are marked for deletion and are cleared thereafter. The time that is taken to clear the currently active sessions varies, depending on the scaled nature of the environment.

Options

none Clear all sessions.

application-protocol *protocol* (Optional) Clear sessions for one of the following application protocols:

- `bootp`—Bootstrap protocol
- `dce-rpc`—Distributed Computing Environment-Remote Procedure Call protocols
- `dce-rpc-portmap`—Distributed Computing Environment-Remote Procedure Call protocols portmap service
- `dns`—Domain Name System protocol
- `exec`—Exec
- `ftp`—File Transfer Protocol
- `h323`—H.323 standards
- `icmp`—Internet Control Message Protocol
- `iiop`—Internet Inter-ORB Protocol
- `ip`—IP
- `login`—Login
- `netbios`—NetBIOS
- `netshow`—NetShow
- `pptp`—Point-to-Point Tunneling Protocol
- `realaudio`—RealAudio

- `rpc`—Remote Procedure Call protocol
- `rpc-portmap`—Remote Procedure Call protocol portmap service
- `rtsp`—Real-Time Streaming Protocol
- `shell`—Shell
- `sip`—Session Initiation Protocol
- `snmp`—Simple Network Management Protocol
- `sqlnet`—SQLNet
- `talk`—Talk Program
- `tftp`—Trivial File Transfer Protocol
- `traceroute`—Traceroute
- `winframe`—WinFrame

destination-port *destination-port*
(Optional) Clear sessions for the specified destination port. The range of values is from 0 to 65535.

destination-prefix *destination-prefix*
(Optional) Clear sessions for the specified destination prefix.

interface *interface-name*
(Optional) Clear sessions for the specified interface. On M Series and T Series routers, the *`interface-name`* can be `ms-`*`fpc/ pic/ port`* or `rsp`*`number`*.

ip-action
(Optional) Clear `ip-action` entries generated by the router to log, drop, or block traffic based on previous matches. The IP action options and targets are configured at the `{edit security idp idp-policy` *`policy-name`* `rulebase-ips rule` *`rule-name`* `then]` hierarchy level.

nat
Clear sessions with network address translation.

nat-port-overload-index
Clear NAT port overload index.

protocol *protocol*
(Optional) Clear sessions for one of the following IP types:

- *`number`*—Numeric protocol value from `0` to `255`
- `ah`—IPsec Authentication Header protocol
- `egp`—An exterior gateway protocol

- `esp`—IPsec Encapsulating Security Payload protocol
- `gre`—A generic routing encapsulation protocol
- `icmp`—Internet Control Message Protocol
- `icmp6`—Internet Control Message Protocol version 6
- `igmp`—Internet Group Management Protocol
- `ipip`—IP-over-IP Encapsulation Protocol
- `ospf`—Open Shortest Path First protocol
- `pim`—Protocol Independent Multicast protocol
- `rsvp`—Resource Reservation Protocol
- `sctp`—Stream Control Transmission Protocol
- `tcp`—Transmission Control Protocol
- `udp`—User Datagram Protocol

service-set *service-set*	(Optional) Clear sessions for the specified service set.
source-nat-pool	Clear the source NAT pool name.
source-port *source-port*	(Optional) Clear sessions for the specified source port. The range of values is from 0 through 65535.
source-prefix *source-prefix*	(Optional) Clear sessions for the specified source prefix.

Required Privilege Level

clear

Output Fields

Table 20 on page 4177 lists the output fields for the `clear services sessions` command. Output fields are listed in the approximate order in which they appear.

Table 20: clear services sessions Output Fields

Field Name	Field Description
`Interface`	Name of an interface.
`Service set`	Name of the service set from which sessions are being cleared.
`Sessions marked for deletion`	Number of sessions that are marked for deletion and are subsequently cleared.

Sample Output

clear services sessions

```
user@host>clear services sessions
Interface   Service set           Sessions marked for deletion
ms-0/0/0    sset                        10
```

Release Information

Command introduced in Junos OS Release 13.1.

The `nat`, `nat-port-overload-index`, and `source-nat-pool` are introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

show services sessions | 26129

clear services sessions analysis

IN THIS SECTION

- Syntax | 4178
- Description | 4178
- Options | 4178
- Required Privilege Level | 4178
- Release Information | 4178

Syntax

```
clear services sessions analysis
```

Description

Clear session statistics.

Options

interface *interface-name*	(Optional) Clear sessions statistics for the specified interface. The *interface-name* can be vms-*fpc*/ *pic*/ *port*.

Required Privilege Level

view

Release Information

Statement introduced in Junos OS Release 17.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

clear services ssl proxy statistics

IN THIS SECTION

- Syntax | 4179
- Description | 4179
- Options | 4179
- Required Privilege Level | 4179
- Output Fields | 4179
- Release Information | 4179

Syntax

```
clear services ssl proxy statistics
```

Description

Clear services SSL proxy statistics. An SSL proxy profile defines SSL behavior for the SRX Series Firewall.

Options

logical-system	Clear the ssl proxy statistics.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 12.1.

The `logical system` option is introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

show services ssl proxy statistics | 26252

clear services stateful-firewall flows

IN THIS SECTION

- Syntax | 4180
- Description | 4181
- Options | 4181
- Required Privilege Level | 4182
- Output Fields | 4182
- Sample Output | 4183
- Release Information | 4183

Syntax

```
clear services stateful-firewall flows
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Description

Clear stateful firewall flows. Issue this command to clear the stateful firewall flows for the specified option. The default option is "none", that is, to close all stateful firewall flows unless another option is specified.

Starting in Junos Release 14.1, the method for closing flows has changed. With the change, even for peak flows, the command prompt now returns to an active state after 30 seconds and the clear command completes in 90 to 120 seconds. In previous releases, closing peak flows could take as long as 4 minutes, after which the command prompt would return. Note too that during the first 30 seconds of issuing the command, the flows to be deleted remain visible in the `show services stateful-firewall flows` command output.

Options

none
Clear all stateful firewall flows.

destination-port *destination-port*
(Optional) Clear stateful firewall flows for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*
(Optional) Clear stateful firewall flows for a particular destination prefix.

interface *interface-name*
(Optional) Clear stateful firewall flows for a particular interface. On M Series and T Series routers, the *interface-name* can be **ms-***fpc/pic/port* or **rsp***number*.

protocol
(Optional) Clear stateful firewall flows for one of the following IP types:

- *number*—Numeric protocol value from **0** to **255**.
- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-over-IP Encapsulation Protocol
- **ospf**—Open Shortest Path First protocol

- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*	(Optional) Clear stateful firewall flows for a particular service set.
source-port *source-port*	(Optional) Clear stateful firewall flows for a particular source port. The range of values is from 0 through 65535.
source-prefix *source-prefix*	(Optional) Clear stateful firewall flows for a particular source prefix.

Required Privilege Level

view

Output Fields

Table 21 on page 4182 lists the output fields for the `clear services stateful-firewall flows` command. Output fields are listed in the approximate order in which they appear.

Table 21: clear services stateful-firewall flows Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of the service set from which flows are being cleared.
Conv removed	Number of conversations removed.

Sample Output

clear services stateful-firewall flows

```
user@host> clear services stateful-firewall flows
Interface   Service set                                          Conv removed
ms-0/3/0    svc_set_trust                                                 0
ms-0/3/0    svc_set_untrust                                               0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show services stateful-firewall flows | 26286

clear services stateful-firewall sip-call

IN THIS SECTION

- Syntax | 4184
- Description | 4184
- Options | 4184
- Required Privilege Level | 4186
- Output Fields | 4186
- Sample Output | 4187
- Release Information | 4187

Syntax

```
clear services stateful-firewall sip-call
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Description

Clear Session Initiation Protocol (SIP) call information in stateful firewall flows.

Options

none Clear stateful firewall statistics for all interfaces and all service sets.

application-protocol (Optional) Clear information about one of the following application protocols:

- **bootp**—(SIP only) Bootstrap protocol
- **dce-rpc**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—(SIP only) Domain Name System protocol
- **exec**—(SIP only) Exec
- **ftp**—(SIP only) File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **login**—Login

- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol
- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*

(Optional) Clear information for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*

(Optional) Clear information for a particular destination prefix.

interface *interface-name*

(Optional) Clear information for a particular adaptive services interface. On M Series and T Series routers, the *interface-name* can be **sp-***fpc/pic/port* or **rsp***number*.

protocol

(Optional) Clear information about one of the following IP types:

- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol

- **ipip**—IP-within-IP Encapsulation Protocol
- **ipv6**—IPv6 within IP
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set* (Optional) Clear information for a particular service set.

source-port *source-port* (Optional) Clear information for a particular source port. The range of values is 0 to 65535.

source-prefix *source-prefix* (Optional) Clear information for a particular source prefix.

Required Privilege Level

view

Output Fields

Table 22 on page 4186 lists the output fields for the `clear services stateful-firewall sip-call` command. Output fields are listed in the approximate order in which they appear.

Table 22: clear services stateful-firewall sip-call Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of the service set from which flows are being cleared.
SIP calls removed	Number of SIP calls removed.

Sample Output

clear services stateful-firewall sip-call

```
user@host> clear services stateful-firewall sip-call
Interface    Service set                                        SIP calls removed
sp-0/3/0     test_sip_777                                                     1
```

Release Information

Command introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

show services stateful-firewall sip-call | 26294

clear services stateful-firewall sip-register

IN THIS SECTION

- Syntax | 4187
- Description | 4188
- Options | 4188
- Required Privilege Level | 4190
- Output Fields | 4190
- Sample Output | 4191
- Release Information | 4191

Syntax

```
clear services stateful-firewall sip-register
<application-protocol protocol>
```

```
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Description

Clear Session Initiation Protocol (SIP) register information in stateful firewall flows.

Options

application-protocol

(Optional) Clear information about one of the following application protocols:

- **bootp**—(SIP only) Bootstrap protocol
- **dce-rpc**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—(SIP only) Domain Name System protocol
- **exec**—(SIP only) Exec
- **ftp**—(SIP only) File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol

- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*
(Optional) Clear information for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*
(Optional) Clear information for a particular destination prefix.

interface *interface*
(Optional) Clear information about a particular interface. On M Series and T Series routers, the *interface-name* can be **sp-***fpc/pic/port* or **rsp***number*.

protocol
(Optional) Clear information about one of the following IP types:

- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ipv6**—IPv6 within IP
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol

- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set* — (Optional) Clear information for a particular service set.

source-port *source-port* — (Optional) Clear information for a particular source port. The range of values is 0 through 65535.

source-prefix *source-prefix* — (Optional) Clear information for a particular source prefix.

Required Privilege Level

view

Output Fields

Table 23 on page 4190 lists the output fields for the `clear services stateful-firewall sip-register` command. Output fields are listed in the approximate order in which they appear.

Table 23: clear services stateful-firewall sip-register Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of the service set from which flows are being cleared.
SIP registration removed	Number of SIP registers removed.

Sample Output

clear services stateful-firewall sip-register

```
user@host> clear services stateful-firewall sip-register
Interface    Service set                                SIP registration removed
sp-0/3/0     test_sip_777                                                    1
```

Release Information

Command introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

show services stateful-firewall sip-register | 26300

clear services stateful-firewall statistics

IN THIS SECTION

- Syntax | 4192
- Description | 4192
- Options | 4192
- Required Privilege Level | 4192
- Output Fields | 4192
- Sample Output | 4192
- Release Information | 4192

Syntax

```
clear services stateful-firewall statistics
<interface interface-name>
<service-set service-set>
```

Description

Clear stateful firewall statistics.

Options

none	Clear stateful firewall statistics for all interfaces and all service sets.
interface ***interface-name***	(Optional) Clear stateful firewall statistics for the specified interface. On M Series and T Series routers, the ***interface-name*** can be **ms-*fpc/pic/port*** or **rsp*number***.
service-set ***service-set***	(Optional) Clear stateful firewall statistics for the specified service set.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services stateful-firewall statistics

```
user@host> clear services stateful-firewall statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show services stateful-firewall statistics | 26305

clear services subscriber analysis

IN THIS SECTION

- Syntax | 4193
- Description | 4193
- Options | 4193
- Required Privilege Level | 4193
- Release Information | 4193

Syntax

```
clear services subscriber analysis
```

Description

Clear information about the number of active subscribers on the services PIC.

Options

interface *interface-name* (Optional) Display information about a particular interface.

Required Privilege Level

view

Release Information

Statement introduced in Junos OS Release 17.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

clear services user-identification active-directory-access

IN THIS SECTION

- Syntax | 4194
- Description | 4194
- Options | 4194
- Required Privilege Level | 4195
- Output Fields | 4195
- Release Information | 4195

Syntax

```
clear services user-identification active-directory-access (active-directory-authentication-table |
statistics (ip-user-mapping | ip-user-probe | user-group-mapping))
```

Description

Delete entries from the Active Directory authentication table or statistics related to integrated user firewall mappings.

Options

- `active-directory-authentication-table`—Remove all entries from the Active Directory authentication table.
- `statistics`—Remove the specified type of statistics:
  - `ip-user-mapping`—IP address-to-user mappings
  - `ip-user-probe`—PC probe statistics

- `user-group-mapping`—User-to-group mappings

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

ip-user-mapping | 9283

request services user-identification active-directory-access ip-user-probe | 14985

show services user-identification active-directory-access statistics | 26389

show services user-identification active-directory-access user-group-mapping | 26393

user-group-mapping | 30555

user-identification (Services) | 30560

clear services user-identification authentication-table

IN THIS SECTION

- Syntax | 4196
- Description | 4196
- Options | 4196
- Additional Information | 4196
- Required Privilege Level | 4196
- Output Fields | 4196
- Sample Output | 4197

Release Information | 4197

Syntax

```
clear services user-identification authentication-table authentication-source authentication-source (all | active-directory | aruba-clearpass | identity-management)
```

Description

Clears the user identity and authentication entries content of the specified authentication source's authentication table.

Options

authentication-source	Active Directoy, Aruba ClearPass, or the identity management server, which could be the Juniper Identity Management Service (JIMS) or any third-party authentication source.

Additional Information

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

clear

Output Fields

For Aruba ClearPass, if there are no entries in the ClearPass authentication table, the following warning message is displayed after you enter the clear command.

```
There is no authentication-table entry.
```

If there are entries in the ClearPass authentication table, no messages are displayed after you enter the `clear` command.

Sample Output

clear services user-identification authentication-table authentication-source

```
user@host> clear services user-identification authentication-table authentication-source aruba-clearpass
warning: "There is no authentication-table entry."
```

clear services user-identification authentication-table authentication-source identity-management

```
user@host> clear services user-identification authentication-table authentication-source identity-management
warning: "There is no authentication-table entry."
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

Support added for Aruba ClearPass as an authentication source in Junos OS release 12.3X48-D30.

Support added for identity-management as an authentication source in Junos OS Release 15.1X49-D100.

clear services user-identification logical-domain-identity-management counters

IN THIS SECTION

- Syntax | 4198
- Description | 4198
- Options | 4198
- Required Privilege Level | 4198
- Sample Output | 4198

Release Information | 4198

Syntax

```
clear services user-identification logical-domain-identity-management counters
```

Description

Clears logical domain identity management information.

Options

counters Clears logical domain identity management query counters.

Required Privilege Level

clear

Sample Output

clear services user-identification logical-domain-identity-management counters

```
user@host> clear services user-identification logical-domain-identity-management counters
node0:
--------------------------------------------------------------------------
warning: "There is no logical-domain-identity-management."
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

show services user-identification logical-domain-identity-management | 26426

clear services video-monitoring mdi errors fpc-slot

IN THIS SECTION

- Syntax | 4199
- Description | 4199
- Options | 4199
- Required Privilege Level | 4199
- Output Fields | 4200
- Sample Output | 4200
- Release Information | 4200

Syntax

```
clear services video-monitoring mdi errors <fpc-slot fpc-slot>
```

Description

Clear all media delivery index error counters for the specified FPC slot or for all FPC slots.

Options

none	Clear error counters for all FPC slots.
fpc-slot	(Optional) Clear error counters for the specified FPC slot.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services video-monitoring mdi errors

```
user@host> clear services video-monitoring mdi errors
Errors counters cleared
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

show services video-monitoring mdi stats fpc-slot | 26441

clear services video-monitoring mdi statistics fpc-slot

IN THIS SECTION

- Syntax | 4201
- Description | 4201
- Options | 4201
- Required Privilege Level | 4201
- Release Information | 4201

Syntax

```
clear services video-monitoring mdi statistics fpc-slot fpc-slot
```

Description

Clear all media delivery index statistics counters except for active flows.

Options

fpc-slot	Number of the FPC slot.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

show services video-monitoring mdi stats fpc-slot | 26441

clear-services-web-filter-dns-feed-statistics-feed-name

IN THIS SECTION

- Syntax | 4202
- Description | 4202
- Required Privilege Level | 4202
- Release Information | 4202

Syntax

```
clear services web-filter dns-feed-statistics feed-name feed-name
```

Description

Clear statistics for DNS request filtering for the specified feed-name.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

Configuring URL Filtering

clear services web-filter statistics profile

IN THIS SECTION

- Syntax | 4203
- Description | 4203
- Options | 4203
- Required Privilege Level | 4203
- Output Fields | 4203
- Sample Output | 4203
- Release Information | 4204

Syntax

```
clear services web-filter statistics profile profile-name
<dns-filter-template template-name>
<fpc-slot fpc-slot pic-slot pic-slot>
<url-filter-template template-name>
```

Description

Clear statistics for DNS request filtering or URL filtering for the specified filter profile.

Options

dns-filter-template ***template-name***	(Optional) Name of the DNS filter template for which statistics are cleared.
fpc-slot ***fpc-slot*** **pic-slot** ***pic-slot***	(Optional) Location of the services PIC for which statistics are cleared.
profile ***profile-name***	Name of the filter profile for which statistics are cleared.
url-filter-template ***template-name***	(Optional) Name of the URL filter template for which statistics are cleared.

Required Privilege Level

clear

Output Fields

When you enter this command, the statistics for DNS request filtering are cleared. There is no specific output.

Sample Output

clear services web-filter statistics profile

```
user@host> clear services web-filter statistics profile profile1
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

> *DNS Request Filtering for Disallowed Website Domains*
>
> *Configuring URL Filtering*

clear sflow collector statistics (QFX Series)

IN THIS SECTION

- Syntax | 4204
- Description | 4204
- Required Privilege Level | 4204
- Sample Output | 4205
- Release Information | 4205

Syntax

```
clear sflow collector statistics
```

Description

Clear the sample counters for all sFlow collectors.

Required Privilege Level

view

Sample Output

clear sflow collector statistics

The following example shows two output examples for the `show sflow collector` command, one before and one after the `clear sflow collector statistics` command was issued.

```
user@host> show sflow collector
Collector        Udp-port    No. of samples
   address
   10.1.1.1            6343        3174
   10.1.2.1            6343        3562
```

command-name

```
user@host> clear sflow collector statistics
```

command-name

```
user@host> show sflow collector
Collector        Udp-port    No. of samples
   address
   10.1.1.1             6343        0
   10.1.2.1             6343        0
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

show sflow collector

clear sflow collectors statistics

IN THIS SECTION

- Syntax | 4206
- Description | 4206
- Required Privilege Level | 4206
- Sample Output | 4206
- Release Information | 4206

Syntax

```
clear sflow collectors statistics
```

Description

Clear the sFlow collector's statistics.

Required Privilege Level

clear

Sample Output

```
clear sflow collectors statistics
```

Release Information

Command introduced in JUNOS Release 9.5 .

clear snmp history

IN THIS SECTION

- Syntax | 4207
- Description | 4207
- Options | 4207
- Required Privilege Level | 4207
- Release Information | 4207

Syntax

```
clear snmp history (index | all)
```

Description

Delete the history record of collected samples of SNMP Ethernet statistics.

Options

all Clear all the entries in the history index.

index Clear the contents of the specified entry in the history index.

Required Privilege Level

clear

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear snmp statistics | 4208

clear snmp statistics

IN THIS SECTION

- Syntax | 4208
- Description | 4208
- Options | 4208
- Required Privilege Level | 4208
- Output Fields | 4209
- Sample Output | 4209
- Release Information | 4210

Syntax

```
clear snmp statistics
```

Description

Clear Simple Network Management Protocol (SNMP) statistics.

Options

This command has no options.

Required Privilege Level

clear

Output Fields

See "show snmp statistics" on page 26529 for an explanation of output fields.

Sample Output

clear snmp statistics

In the following example, SNMP statistics are displayed before and after the `clear snmp statistics` command is issued:

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 8, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 8, Total set varbinds: 0,
    Get requests: 0, Get nexts: 8, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 2298, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 8, Traps: 2290


user@host> clear snmp statistics


user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 0, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 0, Total set varbinds: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
```

```
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 0, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 0, Traps: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show snmp statistics | 26529

clear spanning-tree protocol-migration

IN THIS SECTION

- Syntax | 4211
- Description | 4211
- Options | 4211
- Additional Information | 4211
- Required Privilege Level | 4211
- Sample Output | 4211
- Release Information | 4211

Syntax

```
clear spanning-tree protocol-migration
<interface interface-name>
<routing-instance routing-instance-name>
```

Description

Revert from the original IEEE 802.1D Spanning Tree Protocol (STP) back to the Rapid Spanning Tree Protocol after the `force-version` statement has been removed from the configuration.

Options

none	Reset the STP protocol for all interfaces and all routing instances.
interface *interface-name*	(Optional) Reset the STP protocol for the specified interface only.
routing-instance *routing-instance-name*	(Optional) Reset the STP protocol for a particular routing instance.

Additional Information

For information about the `force-version` statement, see the *Junos Routing Protocols Configuration Guide*.

Required Privilege Level

clear

Sample Output

clear spanning-tree protocol-migration

```
user@host> clear spanning-tree protocol-migration
```

Release Information

Command introduced in Junos OS Release 9.0.

clear spanning-tree statistics

IN THIS SECTION

- Syntax | 4212
- Syntax (EX Series Switches and the QFX Series) | 4212
- Description | 4212
- Options | 4212
- Required Privilege Level | 4213
- Output Fields | 4213
- Sample Output | 4213
- Release Information | 4213

Syntax

```
clear spanning-tree statistics
<interface interface-name>
<logical-system logical-system-name>
```

Syntax (EX Series Switches and the QFX Series)

```
clear spanning-tree statistics
<interface interface-name>
```

Description

Clear Spanning Tree Protocol statistics for all interfaces or a specified interface.

Options

none	Reset STP counters for all interfaces for all routing instances.
interface *interface-name*	(Optional) Clear STP statistics for the specified interface only.

logical-system ***logical-system-name*** (Optional) Clear STP statistics on a particular logical system.

> **NOTE**: The **logical-system** option is not available on QFabric systems.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear spanning-tree statistics

```
user@switch> clear spanning-tree statistics
```

Release Information

Command introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Understanding STP

show spanning-tree statistics | 26572

show spanning-tree bridge | 26548

show spanning-tree interface | 26555

clear spanning-tree statistics bridge

IN THIS SECTION

- Syntax | 4214
- Description | 4214
- Required Privilege Level | 4214
- Sample Output | 4214
- Release Information | 4215

Syntax

```
clear spanning-tree statistics bridge
```

Description

Clear the statistics of the bridge.

Required Privilege Level

clear

Sample Output

clear spanning-tree statistics bridge (MX Series)

command-name

```
user@host> clear spanning-tree statistics bridge
```

Release Information

Command introduced in Junos OS Release 14.1.

clear spanning-tree stp-buffer

IN THIS SECTION

- Syntax | 4215
- Description | 4215
- Required Privilege Level | 4215
- Sample Output | 4215
- Release Information | 4216

Syntax

```
clear spanning-tree stp-buffer
```

Description

Clear the stp-buffer.

Required Privilege Level

clear

Sample Output

clear spanning-tree stp-buffer (MX Series)

command-name

```
user@host> clear spanning-tree stp-buffer
```

Release Information

Command introduced in Junos OS Release 14.1.

clear synchronous-ethernet esmc statistics

IN THIS SECTION

- Syntax | 4216
- Description | 4216
- Options | 4216
- Required Privilege Level | 4217
- Sample Output | 4217
- Release Information | 4217

Syntax

```
clear synchronous-ethernet esmc statistics <interface-name>
```

Description

Clear the ESMC statistics for all the interfaces.

Options

interface-name	(Optional) Clear ESMC statistics for the specified interface.

Required Privilege Level

clear

Sample Output

clear synchronous-ethernet esmc statistics

The following example displays the message after the `clear synchronous-ethernet esmc statistics` command is entered:

```
user@host> clear synchronous-ethernet esmc statistics
Cleared ESMC statistics for all interfaces
```

Release Information

Command introduced before Junos OS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring Synchronous Ethernet on MX Series Routers

Synchronous Ethernet Overview

show synchronous-ethernet esmc statistics | **26674**

clear synchronous-ethernet wait-to-restore

IN THIS SECTION

- Syntax | **4218**
- Description | **4218**
- Options | **4218**
- Required Privilege Level | **4218**
- Sample Output | **4219**

Release Information | 4219

Syntax

```
clear synchronous-ethernet wait-to-restore
<interface interface-name>
```

Description

Clear the wait-to-restore timer for all the interfaces in an MX Series router.

NOTE: When you perform GRES on MX Series routers, you must execute the `clear synchronous-ethernet wait-to-restore` operational mode command on the new primary Routing Engine to clear the wait-to-restore timer on it. This is because the `clear synchronous-ethernet wait-to-restore` operational mode command clears the wait-to-restore timer only on the local Routing Engine.

Options

interface *interface-name* (Optional) Clear wait to restore timer for the specified interface.

Required Privilege Level

clear

Sample Output

clear synchronous-ethernet wait-to-restore interface <interface-name>

The following example displays the message after the `clear synchronous-ethernet wait-to-restore interface ge-2/1/6` command is entered in operational mode command:

```
user@host> clear synchronous-ethernet wait-to-restore interface ge-2/1/6
Cleared wait-to-restore timer for interface ge-2/1/6
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Example: Configuring Synchronous Ethernet on MX Series Routers

Synchronous Ethernet Overview

clear system commit

IN THIS SECTION

- Syntax | 4220
- Description | 4220
- Options | 4220
- Required Privilege Level | 4220
- Output Fields | 4220
- Sample Output | 4220
- Release Information | 4221

Syntax

```
clear system commit
<synchronize-server pending-jobs>
```

Description

Clear pending commit operations.

Options

none Clear pending commit operations on the static configuration database.

synchronize-server pending-jobs (Optional) Clear pending commit synchronize operations for all instances of the ephemeral configuration database on an MX Series Virtual Chassis or a device with dual Routing Engines. This option can only be executed on the primary Routing Engine of the Virtual Chassis primary router or the dual Routing Engine system.

> **NOTE**: The `clear system commit synchronize-server pending-jobs` command should be used only under the supervision of the Juniper Networks Technical Assistance Center (JTAC) during a troubleshooting session.

Required Privilege Level

maintenance (or the actual user who scheduled the commit)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system commit

```
user@host> clear system commit
Pending commit cleared.
```

clear system commit (None Pending)

```
user@host> clear system commit
No commit scheduled.
```

clear system commit (User Does Not Have Required Privilege Level)

```
user@host> clear system commit
error: Permission denied
```

clear system commit synchronize-server pending-jobs

```
user@host> clear system commit synchronize-server pending-jobs
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `synchronize-server` introduced in Junos OS Release 17.2R1 and Junos OS Evolved Release 22.1R1.

RELATED DOCUMENTATION

show system commit | 26754

clear system commit prepared

IN THIS SECTION

- Syntax | 4222
- Description | 4222
- Options | 4222
- Required Privilege Level | 4222

- Output Fields | 4222
- Sample Output | 4222
- Release Information | 4223

Syntax

```
clear system commit prepared
```

Description

Clear the prepared commit. This initiates cleanup of the saved database data structures and the necessary files that are generated as a result of the commit preparation stage and unlinks the pending activation file. A log message is generated upon successful clearing of the pending commit.

Options

This command has no options.

Required Privilege Level

Maintenance (or the actual user who scheduled the commit)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system commit prepared

```
user@host> clear system commit prepared
Prepared commit cleared.
```

clear system commit prepared (None Prepared)

```
user@host> clear system commit prepared
No commit prepared.
```

Release Information

Command introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

clear system commit | 4219

clear system errors

IN THIS SECTION

- Syntax | 4223
- Description | 4224
- Options | 4224
- Required Privilege Level | 4224
- Output Fields | 4224
- Release Information | 4224

Syntax

```
clear system errors fpc fpc-slot fpc-slot
<all>
<error-id error-id-uri>
```

Description

Clear system errors associated with J-Insight fault monitoring.

Options

all	(Optional) Clear all systems errors.
error-id *error-id-uri*	(Optional) Clear system errors for a specified error ID URI.
fpc-slot *fpc-slot*	Clear system errors for a specified FPC.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

J-Insight Device Monitor Basic Configuration

show system errors active | **26822**

show system errors count | **26829**

show system errors fru | **26834**

clear system login lockout

IN THIS SECTION

- Syntax | 4225
- Description | 4225
- Options | 4225
- Required Privilege Level | 4225
- Output Fields | 4225
- Release Information | 4226

Syntax

```
clear system login lockout
<all>
<user username>
```

Description

Use this command to unlock the locked user account.

Options

all	Clear all locked user accounts.
user *username*	Clear the specified locked user account.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

retry-options

show system login lockout | 26902

clear system reboot

IN THIS SECTION

- Syntax | 4226
- Syntax (EX Series Switches) | 4227
- Syntax (TX Matrix Router) | 4227
- Syntax (TX Matrix Plus Router) | 4227
- Syntax (QFX Series) | 4227
- Description | 4227
- Options | 4228
- Required Privilege Level | 4229
- Output Fields | 4229
- Sample Output | 4229
- Release Information | 4230

Syntax

```
clear system reboot
<both-routing-engines>
```

Syntax (EX Series Switches)

```
clear system reboot
<all-members>
<both-routing-engines>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
clear system reboot
<both-routing-engines>
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
clear system reboot
<both-routing-engines>
<all-chassis | all-lcc | lcc number | sfc number>
```

Syntax (QFX Series)

```
clear system reboot
<infrastructure name>
<interconnect-device name>
<node-group name>
```

Description

Clear any pending system software reboots or halts. When issued on a TX Matrix router without any options, the default behavior clears all pending system software reboots or halts on all T640 routers connected to the TX Matrix router. When issued on a TX Matrix Plus router without any options, the default behavior clears all pending system software reboots or halts on all T1600 or T4000 routers connected to the TX Matrix Plus router.

Options

none
Clear all pending system software reboots or halts.

all-chassis
(TX Matrix routers and TX Matrix Plus routers only) (Optional) Clear all halt or reboot requests for all the Routing Engines in the chassis.

all-lcc
(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, clear all halt or reboot requests for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, clear all halt or reboot requests on the l connected T1600 or T4000 LCCs.

all-members
(EX4200 switches only) (Optional) Clear all halt or reboot requests on all members of the Virtual Chassis configuration.

both-routing-engines
(Systems with multiple Routing Engines) (Optional) Clear all halt or reboot requests on both Routing Engines. On a TX Matrix router, clear both Routing Engines on all chassis connected to the TX Matrix router. Likewise, on a TX Matrix Plus router, clear both Routing Engines on all chassis connected to the TX Matrix Plus router.

infrastructure *name*
(QFabric systems) (Optional) Clear all halt or reboot requests on the fabric control Routing Engines or fabric manager Routing Engines.

interconnect-device *name*
(QFabric systems) (Optional) Clear all halt or reboot requests on the Interconnect device.

lcc *number*
(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, clear all halt or reboot requests for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, clear all halt or reboot requests for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local	(EX4200 switches only) (Optional) Clear all halt or reboot requests on the local Virtual Chassis member.
member *member-id*	(EX4200 switches only) (Optional) Clear all halt or reboot requests on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.
node-group *name*	(QFabric systems) (Optional) Clear all halt or reboot requests on the Node group.
scc	(TX Matrix routers only) (Optional) Clear all halt or reboot requests for the TX Matrix router (or switch-card chassis).
sfc *number*	(TX Matrix Plus routers only) (Optional) Clear all halt or reboot requests for the TX Matrix Plus router. Replace *number* with 0.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system reboot

```
user@host> clear system reboot
reboot requested by root at Sat Dec 12 19:37:34 1998
[process id 17855]
Terminating...
```

clear system reboot (TX Matrix Router)

```
user@host> clear system reboot
scc-re0:
--------------------------------------------------------------------------
No shutdown/reboot scheduled.
lcc0-re0:
```

```
---------------------------------------------------------------------------
No shutdown/reboot scheduled.
lcc2-re0:
---------------------------------------------------------------------------
No shutdown/reboot scheduled.
```

clear system reboot (QFX Series)

```
user@switch> clear system reboot node-group node1
No shutdown/reboot scheduled.
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

- request system reboot (Junos OS) | **15140**
- Routing Matrix with a TX Matrix Plus Router Solutions Page

clear system services dhcp binding

IN THIS SECTION

- Syntax | **4231**
- Description | **4231**
- Options | **4231**
- Required Privilege Level | **4231**
- Output Fields | **4231**
- Sample Output | **4231**
- Release Information | **4231**

Syntax

```
clear system services dhcp binding
<address>
```

Description

(EX Series switches only) Remove obsolete IP address bindings on a Dynamic Host Configuration Protocol (DHCP) server and return them to the IP address pool.

Options

address (Optional) Remove a specific IP address binding and return it to the address pool.

Required Privilege Level

view and system

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system services dhcp binding

```
user@host> clear system services dhcp binding
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show system services dhcp binding | 27034

clear system services dhcp conflict

IN THIS SECTION

- Syntax | 4232
- Description | 4232
- Options | 4232
- Required Privilege Level | 4232
- Output Fields | 4232
- Sample Output | 4233
- Release Information | 4233

Syntax

```
clear system services dhcp conflict
<address>
```

Description

(J Series routers and EX Series switches only) Remove IP addresses from the Dynamic Host Configuration Protocol (DHCP) server conflict list and return them to the IP address pool.

Options

address (Optional) Remove a specific IP address from the conflict list and return it to the address pool.

Required Privilege Level

view and system

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system services dhcp conflict

```
user@host> clear system services dhcp conflict
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show system services dhcp conflict | 27043

clear system services dhcp statistics

IN THIS SECTION

- Syntax | 4233
- Description | 4234
- Options | 4234
- Required Privilege Level | 4234
- Output Fields | 4234
- Sample Output | 4234
- Release Information | 4234

Syntax

```
clear system services dhcp statistics
```

Description

(J Series routers and EX Series switches only) Clear Dynamic Host Configuration Protocol (DHCP) server statistics.

Options

This command has no options.

Required Privilege Level

view and system

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear system services dhcp statistics

```
user@host> clear system services dhcp statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show system services dhcp statistics | 27053

clear system services dns dns-proxy

IN THIS SECTION

- Syntax | 4235
- Description | 4235
- Options | 4235
- Required Privilege Level | 4235
- Output Fields | 4236
- Release Information | 4236

Syntax

```
clear system services dns dns-proxy
```

Description

Clear DNS proxy cache information. This option is supported on the SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550M devices.

Options

- `cache`—Clear DNS proxy cache information.
- `statistics`—Clear DNS proxy statistics.
- `none`—Clear all DNS proxy cache information.
- `hostname`—(Optional) Clear DNS proxy cache information from the specified host.

Required Privilege Level

clear

Output Fields

When you enter this command no output is produced.

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

show system services dns dns-proxy | 27057
show system services dynamic-dns | 27061

clear system subscriber-management statistics

IN THIS SECTION

- Syntax | 4236
- Description | 4236
- Options | 4237
- Required Privilege Level | 4237
- Output Fields | 4237
- Sample Output | 4237
- Release Information | 4237

Syntax

```
clear system subscriber-management statistics
```

Description

Clear subscriber-management statistics.

Options

This command has no options.

Required Privilege Level

view and system

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear subscriber-management statistics

```
user@host> clear subscriber-management statistics
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

- show system subscriber-management statistics | 27229
- clear system subscriber-management arp
- clear system subscriber-management ipv6 neighbors

clear tcp-forwarding connections

IN THIS SECTION

- Syntax | 4238
- Description | 4238

- Options | 4238
- Required Privilege Level | 4239
- Output Fields | 4239
- Sample Output | 4239
- Release Information | 4241

Syntax

```
clear tcp-forwarding connections
listening-port listening-port-number listening-address ipv4-listening-address
source-address source-ipv4-address source-port source-port-number
routing-instance routing-instance-name
```

Description

Clear all TCP connections, all connections associated with a specific listening port/listening address combination, or a single connection pair represented by a specific source address/source port combination. For either combination, you can optionally specify a routing instance. If you do not specify a routing instance, the default routing instance is assumed. This command enables you to disconnect TCP port forwarding connections that are not behaving properly.

Options

listening-address *ipv4-listening-address*

IPv4 address that is part of a listening port/listening address combination. The listening address is one on the BNG that external management systems or remote devices must use when attempting to trigger connections on the listening port. You must also specify a listening port.

listening-port *port-number*

Port number that is part of a listening port/listening address combination. The listening port is one that the BNG monitors for connections to be triggered by external management systems or remote devices.

- **Range:** 8000 through 8031

source-address *source-ipv4-address*

Source address of the triggering entity—the remote device or external management system—that appear in the TCP header.

source-port *source-port-number*

Source port of the triggering entity—the remote device or external management system—that appear in the TCP header.

- **Range:** 1 through 65,535

routing-instance *routing-instance-name*

Name of the routing instance for the TCP mapping.

Required Privilege Level

view

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear tcp-forwarding connections

```
user@host> clear tcp-forwarding connections
```

clear tcp-forwarding connections (Listening Port and Address)

The following sample output displays the TCP connection status for a specific listening port/address combination before and after the connection is cleared.

```
user@host> show tcp-forwarding status listening-port 203.0.113.50 listening-address 8002
Listening on: [default:]203.0.113.50:8002
    Status: listening
    Total Bytes Rx: 1230 Tx: 482
  Forwarding to: [default:]192.0.0.4:830
    Total Bytes Rx: 482 Tx: 1230
  Allowed Source Prefixes:
    198.51.100.3/32
  Connections Max: 1 Active: 1
    Source: 198.51.100.3:55002
      Listening: connected Bytes Rx: 0 Tx: 0
      Forwarding: connected Bytes Rx: 0 Tx: 0
```

```
user@host> clear tcp-forwarding connections listening-port 8002 listening-address 203.0.113.50

user@host> show tcp-forwarding status listening-port 8002 listening-address 203.0.113.50
```

clear tcp-forwarding connections (Source Address and Port)

The following sample output displays the TCP connection status before and after the connection is cleared.

```
user@host> show tcp-forwarding status
Listening on: [default:]203.0.113.50:8020
    Status: listening
    Total Bytes Rx: 292 Tx: 112
  Forwarding to: [default:]198.51.100.1:49
    Total Bytes Rx: 112 Tx: 292
  Allowed Source Prefixes:
    192.0.0.1/24
  Connections Max: 4 Active: 3
    Source: 192.0.0.2:55000
      Listening: connected Bytes Rx: 380 Tx: 223
      Forwarding: connected Bytes Rx: 223 Tx: 380
    Source: 192.0.0.3:55000
      Listening: connected Bytes Rx: 855 Tx: 411
      Forwarding: connected Bytes Rx: 411 Tx: 855
    Source: 192.0.0.4:56022
      Listening: connected Bytes 642 Tx: 350
      Forwarding: connected Bytes Rx: 350 Tx: 642

user@host> clear tcp-forwarding connections source-address 192.0.0.2 source-port 55000
user@host> show tcp-forwarding status
Listening on: [default:]203.0.113.50:8020
    Status: listening
    Total Bytes Rx: 292 Tx: 112
  Forwarding to: [default:]198.51.100.1:49
    Total Bytes Rx: 112 Tx: 292
  Allowed Source Prefixes:
    192.0.0.1/24
  Connections Max: 4 Active: 2
    Source: 192.0.0.3:55000
      Listening: connected Bytes Rx: 855 Tx: 411
```

```
    Forwarding: connected Bytes Rx: 411 Tx: 855
  Source: 192.0.0.4:56022
    Listening: connected Bytes 642 Tx: 350
    Forwarding: connected Bytes Rx: 350 Tx: 642
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

TCP Port Forwarding for Remote Device Management

clear tcp-forwarding statistics

IN THIS SECTION

- Syntax | 4241
- Description | 4242
- Options | 4242
- Required Privilege Level | 4242
- Output Fields | 4242
- Sample Output | 4243
- Release Information | 4245

Syntax

```
clear tcp-forwarding statistics
listening-port listening-port-number listening-address ipv4-listening-address
source-address source-ipv4-address source-port source-port-number
routing-instance routing-instance-name
```

Description

Clear the statistics displayed by the `show tcp-forwarding status` command. You can clear statistic for all TCP mappings, for all connections associated with a specific listening port/listening address combination, or for only a single connection pair represented by a specific source address/source port combination. For either combination, you can optionally specify a routing instance. If you do not specify a routing instance, the default routing instance is assumed.

Options

listening-address *ipv4-listening-address*	IPv4 address that is part of a listening port/listening address combination. The listening address is one on the BNG that external management systems or remote devices must use when attempting to trigger connections on the listening port. You must also specify a listening port.
listening-port *port-number*	Port number that is part of a listening port/listening address combination. The listening port is one that the BNG monitors for connections to be triggered by external management systems or remote devices. • **Range:** 8000 through 8031
source-address *source-ipv4-address*	Source address of the triggering entity—the remote device or external management system—that appear in the TCP header.
source-port *source-port-number*	Source port of the triggering entity—the remote device or external management system—that appear in the TCP header. • **Range:** 1 through 65,535
routing-instance *routing-instance-name*	Name of the routing instance for the TCP mapping.

Required Privilege Level

view

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear tcp-forwarding statistics (Listening Port and Address)

The following sample output displays the TCP connection status for a specific listening port/address combination before and after the connection is cleared.

```
user@host> show tcp-forwarding status listening-port 8002 listening-address 203.0.113.50
Listening on: [default:]203.0.113.50:8002
    Status: listening
    Total Bytes Rx: 1230 Tx: 482
  Forwarding to: [default:]192.0.0.4:830
    Total Bytes Rx: 482 Tx: 1230
  Allowed Source Prefixes:
    198.51.100.3/32
  Connections Max: 1 Active: 1
    Source: 198.51.100.3:55002
      Listening: connected Bytes Rx: 587 Tx: 621
      Forwarding: connected Bytes Rx: 621 Tx: 587

user@host> clear tcp-forwarding statistics listening-port 8002 listening-address 203.0.113.50

user@host> show tcp-forwarding status listening-port 8002 listening-address 203.0.113.50
Listening on: [default:]203.0.113.50:8002
    Status: listening
    Total Bytes Rx: 0 Tx: 0
  Forwarding to: [default:]192.0.0.4:830
    Total Bytes Rx: 0 Tx: 0
  Allowed Source Prefixes:
    198.51.100.3/32
  Connections Max: 1 Active: 1
    Source: 198.51.100.3:55002
      Listening: connected Bytes Rx: 0 Tx: 0
      Forwarding: connected Bytes Rx: 0 Tx: 0
```

clear tcp-forwarding statistics (Source Address and Port)

The following sample output displays the TCP connection status before and after the connection is cleared.

```
user@host> show tcp-forwarding status
Listening on: [default:]203.0.113.50:8020
    Status: listening
    Total Bytes Rx: 292 Tx: 112
  Forwarding to: [default:]198.51.100.1:49
    Total Bytes Rx: 112 Tx: 292
  Allowed Source Prefixes:
    192.0.0.1/24
  Connections Max: 4 Active: 3
    Source: 192.0.0.2:55000
      Listening: connected Bytes Rx: 380 Tx: 223
      Forwarding: connected Bytes Rx: 223 Tx: 380
    Source: 192.0.0.3:55000
      Listening: connected Bytes Rx: 855 Tx: 411
      Forwarding: connected Bytes Rx: 411 Tx: 855
    Source: 192.0.0.4:56022
      Listening: connected Bytes 642 Tx: 350
      Forwarding: connected Bytes Rx: 350 Tx: 642

user@host> clear tcp-forwarding statistics source-address 192.0.0.4 source-port 56022
user@host> show tcp-forwarding status
Listening on: [default:]203.0.113.50:8020
    Status: listening
    Total Bytes Rx: 292 Tx: 112
  Forwarding to: [default:]198.51.100.1:49
    Total Bytes Rx: 112 Tx: 292
  Allowed Source Prefixes:
    192.0.0.1/24
  Connections Max: 4 Active: 3
    Source: 192.0.0.2:55000
      Listening: connected Bytes Rx: 380 Tx: 223
      Forwarding: connected Bytes Rx: 223 Tx: 380
    Source: 192.0.0.3:55000
      Listening: connected Bytes Rx: 855 Tx: 411
      Forwarding: connected Bytes Rx: 411 Tx: 855
    Source: 192.0.0.4:56022
```

```
      Listening: connected Bytes 0 Tx: 0
      Forwarding: connected Bytes Rx: 0 Tx: 0
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

TCP Port Forwarding for Remote Device Management

clear-threshold

IN THIS SECTION

- Syntax | 4245
- Hierarchy Level | 4245
- Description | 4246
- Options | 4246
- Required Privilege Level | 4246
- Release Information | 4246

Syntax

```
clear-threshold value;
```

Hierarchy Level

```
[edit logical-systems name security nat source pool name pool-utilization-alarm],
    [edit logical-systems name security nat source pool-utilization-alarm],
    [edit logical-systems name tenants name security nat source pool name pool-utilization-
```

```
alarm],
    [edit logical-systems name tenants name security nat source pool-utilization-alarm],
    [edit security nat source pool name pool-utilization-alarm],
    [edit security nat source pool-utilization-alarm],
    [edit services nat source pool name pool-utilization-alarm],
    [edit services nat source pool-utilization-alarm],
    [edit tenants name security nat source pool name pool-utilization-alarm],
    [edit tenants name security nat source pool-utilization-alarm]
```

Description

Configure the lower threshold at which an SNMP trap is triggered when pool utilization for a source pool without Port Address Translation (PAT) falls below the threshold.

Options

clear-threshold *value* Threshold at which an SNMP trap is triggered.

- **Range:** 40 through 100

Required Privilege Level

security–To view this statement in the configuration.

security-control–To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

clear trace

IN THIS SECTION

- Syntax | 4247
- Description | 4247

- Options | 4247
- Required Privilege Level | 4247
- Sample Output | 4248
- Release Information | 4248

Syntax

```
clear trace
<all-traces | application application-name | node node-name>
```

Description

Clear traces on the system. Trace data from all nodes is collected in a file on the Routing Engine. By default, applications are traced at the info level, which is informational messages.

Options

```
<all-traces | application application-name | node node-name>
```

all-traces	(Optional) Remove all traces.
application *application-name*	(Optional) Remove all traces for the specified application.
node *node-name*	(Optional) Remove all traces for the specified node.

Required Privilege Level

view

Sample Output

clear trace

```
user@host> clear trace
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

Options `all-traces`, `application`, and `node` introduced in Junos OS Evolved Release 19.1R1.

RELATED DOCUMENTATION

show trace | 27392

clear unified-edge tdf aaa radius client statistics

IN THIS SECTION

- Syntax | 4248
- Description | 4249
- Options | 4249
- Required Privilege Level | 4249
- Output Fields | 4249
- Sample Output | 4249
- Release Information | 4250

Syntax

```
clear unified-edge tdf aaa radius client statistics
<all>
```

```
<client name>
<fpc-slot fpc-slot>
<gateway gateway>
<pic-slot pic-slot>
```

Description

Clear statistics for the accounting requests and responses transmitted and received by the RADIUS client for one or more TDF gateways. If a TDF gateway is not specified, then information for all TDF gateways is cleared.

Options

none	(Same as all) Clear statistics for all clients on all TDF gateways.
all	(Optional) Clear statistics for all the clients.
client *name*	(Optional) Clear statistics for the specified client.
fpc-slot *fpc-slot*	(Optional) Clear statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway*	(Optional) Clear statistics for the specified TDF gateway.
pic-slot *pic-slot*	(Optional) Clear statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

clear, unified-edge

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear unified-edge tdf aaa radius client statistics all

```
user@host> clear unified-edge tdf aaa radius client statistics all
```

```
Cleared all RADIUS statistics
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf aaa radius client statistics | 27396

clear unified-edge tdf aaa radius network-element statistics

IN THIS SECTION

- Syntax | 4250
- Description | 4251
- Options | 4251
- Required Privilege Level | 4251
- Output Fields | 4251
- Sample Output | 4251
- Release Information | 4251

Syntax

```
clear unified-edge tdf aaa radius network-element statistics
<fpc-slot fpc-slot>
<gateway gateway>
<pic-slot pic-slot>
<name name>
```

Description

Clear all the statistics for the specified network element.

Options

none	Clear statistics for all network elements for all TDF gateways.
fpc-slot ***fpc-slot***	(Optional) Clear statistics for the specified Flexible PIC Concentrator (FPC).
gateway ***gateway***	(Optional) Clear statistics for the specified TDF gateway.
pic-slot ***pic-slot***	(Optional) Clear statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.
name ***name***	(Optional) Clear statistics for the specified network element.

Required Privilege Level

clear, unified-edge

Output Fields

No message is displayed on successful execution of this command; otherwise, an error message is displayed.

Sample Output

clear unified-edge tdf aaa radius network-element statistics

```
user@host> clear unified-edge tdf aaa radius network-element statistics
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf aaa radius network-element statistics | 27406

clear unified-edge tdf aaa radius server statistics

IN THIS SECTION

- Syntax | 4252
- Description | 4252
- Options | 4252
- Required Privilege Level | 4253
- Output Fields | 4253
- Sample Output | 4253
- Release Information | 4253

Syntax

```
clear unified-edge tdf aaa radius server statistics
<fpc-slot fpc-slot>
<gateway gateway>
<pic-slot pic-slot>
<name name>
```

Description

Clear all the statistics for the specified RADIUS server.

Options

none	Clear statistics for all RADIUS servers for all TDF gateways.
fpc-slot *fpc-slot*	(Optional) Clear statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway*	(Optional) Clear statistics for the specified TDF gateway.

pic-slot *pic-slot* (Optional) Clear statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

name *name* (Optional) Clear statistics for the specified RADIUS server.

Required Privilege Level

clear, unified-edge

Output Fields

No message is displayed on successful execution of this command; otherwise, an error message is displayed.

Sample Output

clear unified-edge tdf aaa radius server statistics

```
user@host> clear unified-edge tdf aaa radius server statistics
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf aaa radius server statistics | 27410

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

clear unified-edge tdf aaa radius snoop-segment statistics

IN THIS SECTION

Syntax | 4254

- Description | 4254
- Options | 4254
- Required Privilege Level | 4254
- Output Fields | 4255
- Sample Output | 4255
- Release Information | 4255

Syntax

```
clear unified-edge tdf aaa radius snoop-segment  statistics
<fpc-slot fpc-slot>
<gateway gateway>
<pic-slot pic-slot>
<segment snoop-segment-name>
```

Description

Clear all snoop segment statistics for FPCs, PICs, TDF gateways, or snoop segments that you specify.

Options

none	Clear all snoop-segment statistics for all TDF gateways.
fpc-slot ***fpc-slot***	(Optional) Clear statistics for the specified Flexible PIC Concentrator (FPC).
gateway ***gateway***	(Optional) Clear statistics for the specified TDF gateway.
pic-slot ***pic-slot***	(Optional) Clear statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.
segment ***snoop-segment-name***	(Optional) Clear statistics for the specified snoop segment.

Required Privilege Level

clear, unified-edge

Output Fields

A message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

clear unified-edge tdf aaa radius snoop-segment statistics

```
user@host> clear unified-edge tdf aaa radius snoop-segment statistics
Cleared Radius snoop-segment Statistics
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf aaa radius snoop-segment statistics | 27420

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview

clear unified-edge tdf aaa statistics

IN THIS SECTION

- Syntax | 4256
- Description | 4256
- Options | 4256
- Required Privilege Level | 4256
- Output Fields | 4256
- Sample Output | 4257
- Release Information | 4257

Syntax

```
clear unified-edge tdf aaa statistics
<all>
<fpc-slot fpc-slot>
<gateway gateway>
<pic-slot pic-slot>
```

Description

Clear global authentication, authorization, and accounting (AAA) statistics for one or more TDF gateways. If a TDF gateway is not specified, then information for all TDF gateways is cleared.

Options

none	(Same as all) Clear AAA statistics for all TDF gateways.
all	(Optional) Clear AAA statistics for all the TDF gateways.
fpc-slot *fpc-slot*	(Optional) Clear AAA statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway*	(Optional) Clear AAA statistics for the specified TDF gateway.
pic-slot *pic-slot*	(Optional) Clear AAA statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

clear, unified-edge

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear unified-edge tdf aaa statistics all

```
user@host> clear unified-edge tdf aaa statistics all

Cleared all AAA statistics
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf aaa statistics | 27425

clear unified-edge tdf address-assignment pool

IN THIS SECTION

- Syntax | 4258
- Description | 4258
- Options | 4258
- Required Privilege Level | 4258
- Output Fields | 4258
- Sample Output | 4258
- Release Information | 4259

Syntax

```
clear unified-edge tdf address-assignment pool name pool-name
<gateway gateway>
<routing-instance routing-instance>
```

Description

Clear the sessions that have been assigned addresses from the specified mobile pool for one or more TDF gateways. If a TDF gateway is not specified, then the sessions for all TDF gateways are cleared.

Options

none	Clear the sessions for all TDF gateways associated with the specified mobile pool.
name ***pool-name***	Clear the sessions for the specified mobile pool.
gateway ***gateway***	(Optional) Clear the sessions on the specified TDF gateway.
routing-instance ***routing-instance***	(Optional) Clear the sessions on the specified routing instance.

Required Privilege Level

clear, unified-edge

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear unified-edge tdf address-assignment pool name

```
user@host> clear unified-edge tdf address-assignment pool name pool-1

Initiated clearing of sessions in the pool
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf address-assignment pool | 27437

clear unified-edge tdf address-assignment statistics

IN THIS SECTION

- Syntax | 4259
- Description | 4259
- Options | 4260
- Required Privilege Level | 4260
- Output Fields | 4260
- Sample Output | 4260
- Release Information | 4260

Syntax

```
clear unified-edge tdf address-assignment statistics
<fpc-slot fpc-slot>
<gateway gateway>
<pic-slot pic-slot>
```

Description

Clear the global address assignment statistics for one or more TDF gateways. If a TDF gateway is not specified, then the statistics for all TDF gateways are cleared.

Options

none	Clear statistics for all TDF gateways.
fpc-slot ***fpc-slot***	(Optional) Clear the statistics for the session PIC in the specified FPC slot.
gateway ***gateway***	(Optional) Clear the statistics for the specified TDF gateway.
pic-slot ***pic-slot***	(Optional) Clear information about the session PIC in this particular PIC slot. For routers, replace *pic-slot* with a value from 0 through 3.

Required Privilege Level

clear, unified-edge

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear unified-edge tdf address-assignment statistics

```
user@host> clear unified-edge tdf address-assignment statistics

Cleared address-assignment statistics
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf address-assignment statistics | 27445

clear unified-edge tdf call-admission-control statistics

IN THIS SECTION

- Syntax | 4261
- Description | 4261
- Options | 4261
- Required Privilege Level | 4261
- Output Fields | 4262
- Sample Output | 4262
- Release Information | 4262

Syntax

```
clear unified-edge tdf call-admission-control statistics gateway gateway-name
<fpc-slot fpc-slot>
<pic-slot pic-slot>
```

Description

Clear call admission control (CAC) statistics for the specified TDF gateway.

Options

fpc-slot *fpc-slot*	(Optional) Clear statistics for the session PIC in the specified FPC slot.
gateway *gateway-name*	Clear CAC statistics for the specified TDF gateway.
pic-slot *pic-slot*	(Optional) Clear statistics for the session PIC in the specified PIC slot.

Required Privilege Level

clear, unified-edge

Output Fields

No message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

clear unified-edge tdf call-admission-control statistics gateway

```
user@host> clear unified-edge tdf call-admission-control statistics gateway TDF
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf call-admission-control statistics | 27448

clear unified-edge tdf diameter network-element statistics

IN THIS SECTION

- Syntax | 4263
- Description | 4263
- Options | 4263
- Required Privilege Level | 4263
- Output Fields | 4263
- Sample Output | 4264
- Release Information | 4264

Syntax

```
clear unified-edge tdf diameter network-element statistics
<fpc-slot fpc-slot>
<gateway gateway-name>
<network-element-name network-element-name>
<pic-slot pic-slot>
```

Description

Clear the statistics for network elements for one or more TDF gateways. If a network element is not specified, then statistics for all network elements are cleared. If a TDF gateway is not specified, then statistics for all TDF gateways are cleared.

Options

none	Clear statistics for all network elements and TDF gateways.
fpc-slot *fpc-slot*	(Optional) Clear statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway-name*	(Optional) Clear statistics for the specified TDF gateway.
network-element-name *network-element-name*	(Optional) Clear statistics for the specified network element.
pic-slot *pic-slot*	(Optional) Clear statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

clear, unified-edge

Output Fields

No message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

clear unified-edge tdf diameter network-element statistics

```
user@host> clear unified-edge tdf diameter network-element statistics
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf diameter network-element statistics | 27456

clear unified-edge tdf diameter pcc-gx statistics

IN THIS SECTION

- Syntax | 4264
- Description | 4265
- Options | 4265
- Required Privilege Level | 4265
- Output Fields | 4265
- Sample Output | 4265
- Release Information | 4265

Syntax

```
clear unified-edge tdf diameter pcc-gx statistics
<fpc-slot fpc-slot>
```

```
<gateway gateway-name>
<pic-slot pic-slot>
```

Description

Clear all statistics for the Gx application for one or more TDF gateways. If a TDF gateway is not specified, then statistics for all TDF gateways are cleared.

Options

none	Clear Gx application statistics for all TDF gateways.
fpc-slot *fpc-slot*	(Optional) Clear statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway-name*	(Optional) Clear statistics for the specified TDF gateway.
pic-slot *pic-slot*	(Optional) Clear statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

clear, unified-edge

Output Fields

No message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

clear unified-edge tdf diameter pcc-gx statistics

```
user@host> clear unified-edge tdf diameter pcc-gx statistics
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf diameter pcc-gx statistics | 4264

clear unified-edge tdf diameter peer statistics

IN THIS SECTION

- Syntax | 4266
- Description | 4266
- Options | 4266
- Required Privilege Level | 4267
- Output Fields | 4267
- Sample Output | 4267
- Release Information | 4267

Syntax

```
clear unified-edge tdf diameter peer statistics
<fpc-slot fpc-slot>
<gateway gateway-name>
<peer-name peer-name>
<pic-slot pic-slot>
```

Description

Clear statistics for Diameter peers for one or more TDF gateways. If a peer is not specified, then statistics for all peers are cleared. If a TDF gateway is not specified, then statistics for all TDF gateways are cleared.

Options

none Clear Diameter peer statistics for all TDF gateways.

fpc-slot *fpc-slot*	(Optional) Clear statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway-name*	(Optional) Clear statistics for the specified TDF gateway.
peer-name *peer-name*	(Optional) Clear statistics for the specified peer.
pic-slot *pic-slot*	(Optional) Clear statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

clear, unified-edge

Output Fields

No message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

clear unified-edge tdf diameter peer statistics

```
user@host> clear unified-edge tdf diameter peer statistics
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf diameter peer statistics | 27469

clear unified-edge tdf statistics

IN THIS SECTION

- Syntax | 4268
- Description | 4268
- Options | 4268
- Required Privilege Level | 4269
- Output Fields | 4269
- Sample Output | 4269
- Release Information | 4269

Syntax

```
clear unified-edge tdf statistics
<data-plane>
<gateway gateway-name>
<domain domain-name
```

Description

Clear all the statistics for the specified TDF gateway, domain, or control plane.

Options

none	Clear statistics for all TDF control planes, domains, and gateways.
data-plane	(Optional) Clear statistics for the data plane.
domain *domain-name*	(Optional) Clear statistics for the specified TDF domain.
gateway *gateway-name*	(Optional) Clear statistics for the specified TDF gateway.

Required Privilege Level

clear, unified-edge

Output Fields

No message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

clear unified-edge tdf statistics gateway

```
user@host> clear unified-edge tdf statistics gateway TDF
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf statistics | 27496

clear unified-edge tdf subscribers

IN THIS SECTION

- Syntax | 4270
- Description | 4270
- Options | 4270
- Required Privilege Level | 4271
- Output Fields | 4271

Sample Output | 4271
Release Information | 4271

Syntax

```
clear unified-edge tdf subscribers [option]
```

Description

Clear the subscribers identified by the option values. You must include at least one option. For IFL-based subscribers, use the `revert` option to re-create the cleared subscribers identified by the option values.

Options

option One or more of the following options:

- `domain` *domain-name*—Clear the subscribers for the specified TDF domain.
- `gateway` *gateway-name*—Clear the subscribers for the specified TDF gateway.
- `interface` *interface-name*—Clear the subscribers on the specified multiservices interface, aggregated multiservices interface, Packet Forwarding Engine interface, or aggregated Packet Forwarding Engine interface names.
- `peer` *peer-name*—Clear the subscriber matching GPRS tunneling protocol (GTP) peer on the specified TDF gateway.
- `revert`—For an IFL-based subscriber, recreate an IFL-subscriber that was cleared.
- `routing-instance` *routing-instance*—Clear the subscriber information for the specified routing instance.
- `subscriber-name` *subscriber-name*—Clear the specified IFL-based subscriber.
- `v4-addr` *v4-addr*—Clear the subscriber information for the specified IPv4 address of the IP-based subscriber's user equipment (UE).
- `v6-addr` *v6-addr*—Clear the subscriber information for the specified IPv6 address of the IP-based subscriber's user equipment.

Required Privilege Level

clear, unified-edge

Output Fields

No message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

clear unified-edge tdf subscribers gateway tdf

```
user@host> clear unified-edge tdf subscribers gateway tdf
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf statistics | 4268
clear unified-edge tdf subscribers peer | 4271
show unified-edge tdf subscribers | 27513
IP-Based and IFL-Based TDF Subscribers Overview
Understanding IFL-Based Subscriber Setup

clear unified-edge tdf subscribers peer

IN THIS SECTION

Syntax | 4272

- Description | 4272
- Options | 4272
- Required Privilege Level | 4273
- Output Fields | 4273
- Sample Output | 4273
- Release Information | 4273

Syntax

```
clear unified-edge tdf subscribers peer
<gateway gateway>
<remote-addr remote-addr>
<nas-id nas-id
<routing-instance routing-instance>
```

Description

Clear the information for IP-based subscribers anchored on the specified RADIUS client, TDF gateway, or both, or for IP-based subscribers matching the specified routing instance.

Options

none	Clear information for all IP-based subscribers.
gateway *gateway*	(Optional) Clear IP-based subscriber information for the TDF gateway.
nas-id *nas-id*	(Optional) Clear IP-based subscriber information for the specified NAS identifier of the RADIUS client.
remote-addr *remote-addr*	(Optional) Clear IP-based subscriber information for the specified IPv4 address of the RADIUS client.
routing-instance *routing-instance*	(Optional) Clear IP-based subscriber information for the specified routing instance.

Required Privilege Level

clear, unified-edge

Output Fields

No message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

clear unified-edge tdf subscribers peer gateway remote-addr

```
user@host> clear unified-edge tdf subscribers peer gateway TDF remote-addr 198.0.2.2
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf subscribers | 4269

show unified-edge tdf subscribers | 27513

clear validation database

IN THIS SECTION

- Syntax | 4274
- Description | 4274
- Options | 4274
- Required Privilege Level | 4274

Sample Output | 4275
Release Information | 4275

Syntax

```
clear validation database
<name database-name>
<instance instance-name>
<logical-system logical-system-name>
```

Description

Clear the route validation database.

Options

none	Clear the route validation database for all routing instances.
name database-name	(Optional) Display the records of the specified database in that routing-instance or logical-system. If unspecified, the records of all databases are displayed.
instance *instance-name*	(Optional) Clear the route validation database for the specified instance.
logical-system *logical-system-name*	(Optional) Perform this operation on a particular logical system.

Required Privilege Level

clear

Sample Output

clear validation database

```
user@host> clear validation database
Clearing database
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

clear validation session

IN THIS SECTION

- Syntax | 4276
- Description | 4276
- Options | 4276
- Required Privilege Level | 4276
- Sample Output | 4276
- Release Information | 4277

Syntax

```
clear validation session
<destination session-ip-address>
<instance instance-name>
<logical-system logical-system-name>
<soft-inbound>
```

Description

Clear the route validation session to the cache server.

Options

none	Clear all route validation sessions for all routing instances.
destination ***session-ip-address***	(Optional) Clear the specified route validation session.
instance ***instance-name***	(Optional) Clear the route validation session for the specified instance.
logical-system ***logical-system-name***	(Optional) Perform this operation on a particular logical system.
soft-inbound	(Optional) Rather than flapping the session to the cache server and removing its contents from the database, refresh the session information without removing the database entries.

Required Privilege Level

clear

Sample Output

clear validation session

```
user@host> clear validation session
Cleared 3 sessions
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

clear validation statistics

IN THIS SECTION

- Syntax | 4277
- Description | 4277
- Options | 4278
- Required Privilege Level | 4278
- Sample Output | 4278
- Release Information | 4278

Syntax

```
clear validation statistics
<instance instance-name>
<logical-system logical-system-name>
```

Description

Clear the route validation statistics.

Options

none	Clear the route validation statistics for all routing instances.
instance *instance-name*	(Optional) Clear the route validation statistics for the specified instance.
logical-system *logical-system-name*	(Optional) Perform this operation on a particular logical system.

Required Privilege Level

clear

Sample Output

clear validation statistics

```
user@host> clear validation statistics
Statistics cleared
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

clear virtual-chassis heartbeat (MX Series Virtual Chassis)

IN THIS SECTION

Syntax | 4279

- Description | 4279
- Options | 4279
- Required Privilege Level | 4280
- Output Fields | 4280
- Sample Output | 4280
- Release Information | 4280

Syntax

```
clear virtual-chassis heartbeat
<(all-members | local | member member-id)>
```

Description

Clear statistics counter and timestamp fields associated with a heartbeat packet connection on both member routers in an MX Series Virtual Chassis. You can issue the `clear virtual-chassis heartbeat` command from the console of either member router in the Virtual Chassis.

> **NOTE**: When you clear heartbeat statistics, the connection time is preserved to retain a record of previous heartbeat connectivity.

Options

none	Clear heartbeat statistics for both member routers in an MX Series Virtual Chassis.
all-members	(Optional) Clear heartbeat statistics for both member routers in an MX Series Virtual Chassis. This is the default behavior if no options are specified.
local	(Optional) Clear heartbeat statistics for the member router from which you are issuing the command.
member ***member-id***	(Optional) Clear heartbeat statistics for the specified member router. Replace *member-id* with the value 0 or 1.

Required Privilege Level

clear

Output Fields

Table 24 on page 4280 lists the output fields for the `clear virtual-chassis heartbeat` command. Output fields are listed in the approximate order in which they appear.

Table 24: clear virtual-chassis heartbeat Output Fields

Field Name	Field Description
`member`*n*	Member ID of the Virtual Chassis member router for which output is displayed.
`heartbeat statistics cleared`	Confirmation that heartbeat statistics are cleared on the specified member router.

Sample Output

clear virtual-chassis heartbeat

```
{master:member0-re0}
user@host> clear virtual-chassis heartbeat
member0:
--------------------------------------------------------------------------
heartbeat statistics cleared

member1:
--------------------------------------------------------------------------
heartbeat statistics cleared
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Verifying and Managing the Virtual Chassis Heartbeat Connection

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in the Same Subnet

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in Different Subnets

clear virtual-chassis protocol statistics

IN THIS SECTION

- Syntax | 4281
- Description | 4281
- Options | 4282
- Required Privilege Level | 4282
- Sample Output | 4282
- Release Information | 4283

Syntax

```
clear virtual-chassis protocol statistics
<all-members>
<local>
<member member-id>
```

Description

Clear—reset to zero (0)—all Virtual Chassis Control Protocol (VCCP) traffic counters for the Virtual Chassis.

Options

none	Clear VCCP counters for all members of the Virtual Chassis.
all-members	(Optional) Clear VCCP traffic counters for all members of the Virtual Chassis.
local	(Optional) Clear VCCP traffic counters on the switch or external Routing Engine on which this command is entered.
member ***member-id***	(Optional) Clear VCCP traffic counters for the specified member of the Virtual Chassis.

Required Privilege Level

clear

Sample Output

clear virtual-chassis protocol statistics

```
user@switch> clear virtual-chassis protocol statistics
member0:
--------------------------------------------------------------------------
Statistics cleared

member1:
--------------------------------------------------------------------------
Statistics cleared

member8:
--------------------------------------------------------------------------
Statistics cleared

member9:
--------------------------------------------------------------------------
Statistics cleared
```

This output shows that the VCCP counters have been cleared for both member switches (member 0 and member 1) and both XRE200 External Routing Engines (member 8 and member 9) in the EX8200 Virtual Chassis.

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

show virtual-chassis protocol statistics | 27663

Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis

clear virtual-chassis vc-port statistics

IN THIS SECTION

- Syntax | 4283
- Description | 4283
- Options | 4284
- Required Privilege Level | 4284
- Sample Output | 4284
- Release Information | 4285

Syntax

```
clear virtual-chassis vc-port statistics
<all-members>
<interface-name>
<local>
<member member-id>
```

Description

Clear—reset to zero (0)—the traffic statistics counters on Virtual Chassis ports (VCPs).

Options

none	Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.
all-members	(Optional) Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.
interface-name	(Optional) Clear traffic statistics for the specified VCP.
local	(Optional) Clear traffic statistics for VCPs from the switch or external Routing Engine on which this command is entered.
member *member-id*	(Optional) Clear traffic statistics for VCPs from the specified member of a Virtual Chassis or VCF.

Required Privilege Level

clear

Sample Output

clear virtual-chassis vc-port statistics (EX4200 Virtual Chassis)

```
user@switch> clear virtual-chassis vc-port statistics
fpc0:
--------------------------------------------------------------------------
Statistics cleared
```

clear virtual-chassis vc-port statistics (EX8200 Virtual Chassis)

```
user@external-routing-engine> clear virtual-chassis vc-port statistics
member0:
--------------------------------------------------------------------------
Statistics cleared

member1:
--------------------------------------------------------------------------
Statistics cleared
```

```
member8:
--------------------------------------------------------------------------
Statistics cleared

member9:
--------------------------------------------------------------------------
Statistics cleared
```

clear virtual-chassis vc-port statistics member 3

```
user@switch> clear virtual-chassis vc-port statistics member 3
Cleared statistics on member 3
```

Release Information

Command introduced in Junos OS Release 9.0.

The options **all-members** and **local** were added in Junos OS Release 9.3 for EX Series switches.

RELATED DOCUMENTATION

- show virtual-chassis vc-port statistics | 27713
- show virtual-chassis vc-port | 27681
- Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

clear vpls mac-address

IN THIS SECTION

- Syntax | 4286
- Description | 4286
- Options | 4286
- Required Privilege Level | 4286
- Output Fields | 4286

Sample Output | 4287
Release Information | 4287

Syntax

```
clear vpls mac-address
<instance instance-name>
<logical-system (all | logical-system-name)>
<mac-address>
```

Description

(T Series and M Series routers, except for the M160 router) Clear media access control (MAC) address entries from the virtual private LAN service (VPLS) table.

Options

none	Clear all MAC address entries from the VPLS table for all routing instances.
instance *instance-name*	(Optional) Clear all MAC address entries for a VPLS instance from the VPLS table.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
mac-address	(Optional) Clear a specific MAC address in a VPLS instance from the VPLS table.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear vpls mac-address

```
user@host> clear vpls mac-address
```

Release Information

Command introduced before Junos OS Release 7.4.

clear vpls mac-move-action

IN THIS SECTION

- Syntax | 4287
- Description | 4287
- Options | 4288
- Required Privilege Level | 4288
- Output Fields | 4288
- Sample Output | 4288
- Release Information | 4288

Syntax

```
clear vpls mac-move-action
<interface ifl-name.unit>
```

Description

Clear the learning interfaces (IFLs) disabled due to a MAC move.

Options

none	Clear all the IFLs disabled due to a MAC move.
`interface` *ifl-name.unit*	(Optional) Clear the VPLS interface disabled due to a MAC move.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear vpls mac-move-action (MX Series)

```
user@host> clear vpls mac-move-action
```

clear vpls mac-move-action interface ifl-name.unit (MX Series)

```
user@host> clear vpls mac-move-action interface xe-1/0/0.1
```

Release Information

Command introduced in Junos OS Release 14.2.

clear vpls mac-table

IN THIS SECTION

Syntax | 4289

- Description | 4289
- Options | 4289
- Required Privilege Level | 4290
- Output Fields | 4290
- Sample Output | 4290
- Release Information | 4290

Syntax

```
clear vpls mac-table
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
<mac-address>
<vlan-id>
```

Description

(MX Series routers) Clear media access control (MAC) addresses from the virtual private LAN service (VPLS) MAC table.

Options

none	Clear all MAC addresses from the VPLS table for all routing instances.
instance *instance-name*	(Optional) Clear all MAC addresses for a VPLS instance from the VPLS table.
interface *interface-name*	(Optional) Clear all MAC addresses for a VPLS interface from the VPLS table.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
mac-address	(Optional) Clear a specific MAC address in a VPLS instance from the VPLS table.

vlan-id (Optional) Clear MAC addresses on a specified VLAN (0 through 4095).

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear vpls mac-table

```
user@host> clear vpls mac-table
```

Release Information

Command introduced before Junos OS Release 9.5.

clear vrrp

IN THIS SECTION

- Syntax | 4291
- Description | 4291
- Options | 4291
- Required Privilege Level | 4291
- Output Fields | 4291
- Sample Output | 4291
- Release Information | 4291

Syntax

```
clear vrrp (all | interface interface-name)
```

Description

Set Virtual Router Redundancy Protocol (VRRP) interface statistics to zero.

Options

all	Clear statistics on all interfaces.
interface *interface-name*	Clear statistics on the specified interface only.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear vrrp all

```
user@host> clear vrrp all
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show vrrp | 27888

client

IN THIS SECTION

- Syntax | 4292
- Hierarchy Level | 4293
- Description | 4293
- Options | 4293
- Required Privilege Level | 4294
- Release Information | 4294

Syntax

```
client client-name {
    chap-secret chap-secret;
    group-profile profile-name;
    ike {
        allowed-proxy-pair {
            remote remote-proxy-address local local-proxy-address;
        }
        pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
        ike-policy policy-name;
        interface-id string-value;
    }
    l2tp {
        aaa-access-profile profile-name;
        interface-id interface-id;
        lcp-renegotiation;
        local-chap;
        maximum-sessions number;
        maximum-sessions-per-tunnel number;
        multilink {
            drop-timeout milliseconds;
            fragment-threshold bytes;
        }
        override-result-code session-out-of-resource;
        ppp-authentication (chap | pap);
```

```
        ppp-profile profile-name;
        sessions-limit-group;
        service-profile profile-name(parameter)&profile-name;
        shared-secret shared-secret;
    }
    pap-password pap-password;
    ppp {
        cell-overhead;
        encapsulation-overhead bytes;
        framed-ip-address ip-address;
        framed-pool framed-pool;
        idle-timeout seconds;
        interface-id interface-id;
        keepalive seconds;
        primary-dns primary-dns;
        primary-wins primary-wins;
        secondary-dns secondary-dns;
        secondary-wins secondary-wins;
    }
    user-group-profile profile-name;
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure the peer identity.

> **NOTE**: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

client-name A peer identity. For L2TP clients, you can use a special name to configure a default client. This client enables the LNS to accept any LAC to establish the session. On M Series routers, use `*` for the default client configuration. On MX Series routers, use `default`.

chap-secret For interfaces with PPP encapsulation on which the PPP Challenge Handshake Authentication Protocol (CHAP) is configured, configure the shared secret (the CHAP secret key associated with a peer), as defined in RFC 1994. This statement is not supported for L2TP LNS on MX Series routers.

- Values:
  - *chap-secret*—The secret key associated with a peer.

group-profile Associate a group profile with a client. This statement is not supported for L2TP LNS on MX Series routers.

- Values:
  - *profile-name*—Name assigned to the group profile.

pap-password Configure the Password Authentication Protocol (PAP) password. This statement is not supported for L2TP LNS on MX Series routers.

- Values:
  - *password*—PAP password.

user-group-profile Apply a configured PPP group profile to PPP users. If `user-group-profile` is modified or deleted, the existing LNS subscribers, which were using this Layer 2 Tunneling Protocol client configuration, go down.

- Values:
  - *profile-name*—Name of a PPP group profile configured at the `[edit access group-profile profile-name]` hierarchy level.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the PPP Challenge Handshake Authentication Protocol

Point-to-Point Protocol (PPP)

Layer 2 Tunneling Protocol (L2TP)

L2TP LNS Inline Service Interfaces

client (TWAMP)

IN THIS SECTION

- Syntax (Junos OS) | **4295**
- Syntax (Junos OS Evolved) | **4297**
- Junos OS Hierarchy Level | **4298**
- Junos OS Evolved Hierarchy Level | **4298**
- Description | **4298**
- Required Privilege Level | **4298**
- Release Information | **4298**

Syntax (Junos OS)

```
client {
    control-connection control-client- name {
        authentication-mode;
        control-type (managed | light);
        destination-interface interface-name;
        destination-port port;
        history-size size;
        moving-average-size number;
        persistent-results;
        routing-instance instance-name;
        target (url url | address address);
        tcp-keepcnt;
        tcp-keepidle;
```

```
        tcp-keepintvl;
        test-interval interval;
        traps {
                control-connection-closed;
                test-iteration-done;
        }
        data-fill-with zeros;
        data-size size;
        dscp-code-points (RPM) dscp-bits;
        probe-count count;
        probe-interval seconds;
        thresholds thresholds;
        test-session session-name {
            data-fill-with zeros data;
            data-size size;
            dscp-code-points (RPM) dscp-bits;
            probe-count count;
            probe-interval seconds;
            source-address source-address;
            target-address target-address local-link IPv6-link-local-interface-name;
            traps {
                    egress-jitter-exceeded;
                    egress-std-dev-exceeded;
                    egress-time-exceeded;
                    ingress-jitter-exceeded;
                    ingress-std-dev-exceeded;
                    ingress-time-exceeded;
                    jitter-exceeded;
                    max-rtt-exceeded;
                    probe-failure;
                    rtt-exceeded;
                    std-dev-exceeded;
                    test-completion;
                    test-failure;
             }
        }
    }
}
```

Syntax (Junos OS Evolved)

```
client {
    control-connection  control-client-name {
        control-type (managed | light);
        destination-port destination-port;
        routing-instance routing-instance-name;
        source-address source-address;
        source-routing {
            forward (srmpls-nodes <label hop1...n> | srv6-nodes <IPv6 SID hop1...n>);
            return (srmpls-nodes <label hop1...n> | srv6-nodes <IPv6 SID hop1...n>);
        }
        target target-address;
        test-start (auto | manual);
         test-interval   seconds;
         traps {
                control-connection-closed;
                test-iteration-done;
         }
         test-session name {
            data-size data-size;
            destination-port destination-port;
            dscp-code-points dscp-code-points;
            history-size history-size;
            moving-average-size moving-average-size;
            offload-type (none | inline-timestamping | pfe-timestamp);
            probe-count  probe-count;
            probe-interval  seconds;
            source-address source-address;
            target target-address local-link IPv6-link-local-interface-name;
            thresholds {
                control-failure (on | off);
                successive-loss number;
                total-loss number;
                threshold-type (microseconds | average);
            }
            traps {
                    egress-jitter-exceeded;
                    egress-time-exceeded;
                    ingress-jitter-exceeded;
                    ingress-time-exceeded;
                    jitter-exceeded;
```

```
                    probe-failure;
                    rtt-exceeded;
                    test-completion;
                    test-failure;
            }
            ttl hop-count;
        }
    }
}
```

Junos OS Hierarchy Level

```
[edit services rpm twamp]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring twamp]
```

Description

Defines the Two-Way Active Measurement Protocol (TWAMP) client configuration. A client opens a TCP connection to the server on well-known port 862. The host that initiates the TCP connection takes on the roles of control client and session sender. The host that acknowledges the TCP connection accepts the roles of server and session reflector.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

Statement introduced in Junos OS Evolved Release 20.3R1.

`control-type` option introduced in Junos OS Release 21.1R1 and Junos OS Evolved Release 20.3R1.

`traps` option introduced in Junos OS Evolved Release 21.3R1.

`source-address` option and the `local-link` sub-option of the `target-address` option for TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

The `local-link` sub-option of the `target` option for TWAMP Light test sessions introduced in Junos OS Evolved Release 22.3R1.

The `inline-timestamping` suboption of the `offload-type` option introduced in Junos OS Evolved Release 22.4R1.

`source-routing` option introduced in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

client (interfaces)

IN THIS SECTION

- Syntax | 4299
- Hierarchy Level | 4300
- Description | 4300
- Required Privilege Level | 4300
- Release Information | 4300

Syntax

```
client;
```

Hierarchy Level

```
[edit interfaces pp0 unit logical-unit-number pppoe-options],
[edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number pppoe-options]
```

Description

Configure the router to operate in the PPPoE client mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring the PPPoE Client Mode

client (Unified Edge)

IN THIS SECTION

- Syntax | 4301
- Hierarchy Level | 4301
- Description | 4301
- Options | 4301
- Required Privilege Level | 4301
- Release Information | 4301

Syntax

```
client client-name;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from]
```

Description

Specify the RADIUS client for the incoming RADIUS request from an IP-based subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

Options

client-name	Name of the RADIUS client.

> **NOTE**: The RADIUS client must have been previously configured at the `[edit access radius]` hierarchy level, and specified as the `aaa-client` at the `[edit unified-edge gateways tdf gateway-name]` hierarchy level.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers*

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers

client-address

IN THIS SECTION

- Syntax | 4302
- Hierarchy Level | 4302
- Description | 4302
- Options | 4303
- Required Privilege Level | 4303
- Release Information | 4303

Syntax

```
client-address ipv4-address;
```

Hierarchy Level

```
[edit services hosted-services server-profile server-profile-name]
```

Description

Configure the source address to include in the header of each sampled packet. You must specify an IPv4 address. You can also specify the loopback address or the management interface address as the client address.

Options

ipv4-address—IPv4 address of the client.

- **Default:** 0.0.0.0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

client-config (Juniper Secure Connect)

IN THIS SECTION

- Syntax | 4304
- Hierarchy Level | 4304
- Description | 4305
- Options | 4305
- Required Privilege Level | 4306
- Release Information | 4306

Syntax

```
client-config name {
    application-bypass {
        term name {
            description description;
            protocol protocol;
            domain-name domain-name;
     }
}
biometric-authentication;
    domain-name domain-name;
    certificate {        no-expiry-warning;
        no-pin-request-per-connection;
        warn-before-expiry days;
    }
    connection-mode (always | manual);
    credentials (username | password);
    dead-peer-detection {
        interval seconds;
        threshold threshold;
    }
    no-dead-peer-detection;
    no-eap-tls;
    no-tcp-encap;
    windows-logon {
        auto-dialog-open;
        disconnect-at-logoff;
        domain domain;
        eap-auth;
        flush-credential-at-logoff;
        lead-time-duration seconds;
        mode (automatic | manual);
    }
}
```

Hierarchy Level

```
[edit security remote-access]
```

Description

Define Juniper Secure Connect remote client configuration parameters. The parameters define how Juniper Secure Connect client establishes VPN tunnel with your security device.

Options

name Name of configuration object name.

application-bypass Define application-bypass configuration.

biometric-authentication Enable biometric authentication.

domain-name Define the set of search domain name. As a system administrator, you can configure the set of search domain name that the Juniper Secure Connect application will use to handle DNS lookups. This is applicable to both full tunnels and split tunnel configurations.

You can provide more than one search domain names by executing the `set security remote-access client-config name domain-name domain-name` multiple times. When you enter more than one domain name, it automatically adds a separator (comma) to that value. The number of domain names are limited to the total number of characters and must not exceed 1023 characters. For example, the two domain names `juniper.net,lab.juniper.net` consumes 27 characters while `juniper.net` consumes 11 characters.

Range: 0-1023 characters including comma.

connection-mode Set one of the following connection mode for clients:

- Values:
  - always—Connect to the VPN automatically when user logs in to remote client device. In always mode, the first VPN connection established when the user clicks the "Connect" button. After that, whenever VPN connection gets disconnected without manual intervention, the client device always attempts to re-establish the connection automatically.
  - manual—Connect to the VPN manually.
- **Default:** manual

credentials Set one of the following to save the user credentials in Juniper Secure Connect application:

Values:

- username—To save the username in Juniper Secure Connect application. When you enable this option, user will not be required to provide username every time they connect to Juniper Secure Connect application.
- password—To save both the username and password in Juniper Secure Connect application. When you enable this option, user will not be required to provide username and password every time they connect to Juniper Secure Connect application.

Note that you cannot configure both the options at the same time. If you have not configured the `credentials` configuration options, then the Juniper Secure Connect application does not remember any user credentials.

dead-peer-detection— Enable dead-peer-detection on the client.

Interval The time between DPD probe messages in seconds.

- **Default:** 60 seconds

threshold Maximum number of DPD retransmissions.

- **Default:** 5

no-dead-peer-detection Disable dead-peer-detection on client

no-eap-tls Disable EAP-TLS IKEV2 method.

no-tcp-encap Disable tcp encapsulation.

windows-logon Specify windows logon options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.3R1.

Support for `domain-name` option at the `[edit security remote-access client-config name]` hierarchy level added in Junos OS Release 22.1R1.

Support for `application-bypass` option added in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Juniper Secure Connect Administrator Guide

client-delegate-probes

IN THIS SECTION

- Syntax | 4307
- Hierarchy Level | 4307
- Description | 4307
- Required Privilege Level | 4308
- Release Information | 4308

Syntax

```
rpm client-delegate-probes;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Generate real-time performance monitoring (RPM) probes on an MS-MPC or MS-MIC services interface, which increases the number of RPM probes that can run at the same time.

The `destination-interface` statement must be configured at the `[edit services rpm probe owner test test-name]` hierarchy level to point to the interface and logical unit number and for which you configure `client-delegate-probes`. Configure the `delegate-probes` statement at the `[edit services rpm probe owner]` hierarchy level to complete the configuration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

client-discover-match (DHCP Local Server)

IN THIS SECTION

- Syntax | 4308
- Hierarchy Level | 4309
- Description | 4309
- Default | 4309
- Options | 4309
- Required Privilege Level | 4309
- Release Information | 4310

Syntax

```
client-discover-match <option60-and-option82 | incoming-interface>;
```

Hierarchy Level

```
[edit system services dhcp-local-server overrides],
[edit system services dhcp-local-server group group-name overrides],
[edit system services dhcp-local-server group group-name interface interface-name overrides]
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ... overrides],
[edit logical-systems logical-system-name system services dhcp-local-server ...overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server ...overrides]
```

Description

Configure the match criteria DHCP local server uses to uniquely identify DHCP subscribers or clients when primary identification fails. The options are mutually exclusive.

Default

By default, DHCP uses the `option60-and-option82` option.

Options

`incoming-interface` (Optional) Allow only one client device to connect on the interface. If the client device changes, the router deletes the existing client binding and creates a binding for the newly connected device.

> **NOTE**: The `overrides client-discover-match incoming-interface` configuration deletes and replaces the existing binding when a new device connects. This action differs from the `overrides interface-client-limit 1` statement, which retains the existing binding and rejects the newly connected client.

`option60-and-option82` (Optional) Use option 60 and option 82 information to identify subscribers.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

`incoming-interface` option added in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Differences Between Legacy DHCP and Extended DHCP

Overriding the Default DHCP Local Server Configuration Settings

DHCP Auto Logout Overview

Allowing Only One DHCP Client Per Interface

client-discover-match (DHCP Relay Agent)

IN THIS SECTION

- Syntax | **4310**
- Hierarchy Level | **4311**
- Description | **4311**
- Default | **4311**
- Options | **4311**
- Required Privilege Level | **4311**
- Release Information | **4312**

Syntax

```
client-discover-match <option60-and-option82 | incoming-interface>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ... overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ... overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group ... overrides]
```

Description

Configure the match criteria DHCP relay uses to uniquely identify DHCP subscribers or clients when primary identification fails. The options are mutually exclusive.

Default

By default, DHCP uses the `option60-and-option82` option.

Options

`incoming-interface` (Optional) Allow only one client device to connect on the interface. If the client device changes, the router deletes the existing client binding and creates a binding for the newly connected device.

> **NOTE**: The `overrides client-discover-match incoming-interface` configuration deletes and replaces the existing binding when a new device connects. This action differs from the `overrides interface-client-limit 1` statement, which retains the existing binding rejects the newly connected client.

`option60-and-option82` (Optional) Use option 60 and option 82 information to identify subscribers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

`incoming-interface` option added in Junos OS Release 13.3.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Overriding the Default DHCP Relay Configuration Settings

DHCP Auto Logout Overview

Allowing Only One DHCP Client Per Interface

client-group

IN THIS SECTION

- Syntax | **4312**
- Hierarchy Level | **4313**
- Description | **4313**
- Options | **4313**
- Required Privilege Level | **4313**
- Release Information | **4313**

Syntax

```
client-group [          group-names        ];
```

Hierarchy Level

```
[edit access profile profile-name client client-name]
[edit access profile profile-name session-options]
```

Description

Specify a list of client groups that the client belongs to. If the group list is not defined as part of the client profile, the client group configured at the `[edit access profile session-options]` hierarchy level is used.

Options

group-names —Names of one or more groups the client belongs to, separated by spaces—for example `g1`, `g2`, `g3`. The total length of the group name string cannot exceed 256 characters.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

client-ia-type

IN THIS SECTION

- Syntax | 4314
- Hierarchy Level | 4314
- Description | 4314
- Options | 4314
- Required Privilege Level | 4314

Release Information | 4315

Syntax

```
client-ia-type {
    ia-na;
    ia-pd;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6 dhcpv6-client]
[edit logical-systems logical -system-name interfaces interface-name unit logical-unit-number
family inet6 dhcpv6-client]
[edit tenants tenant-name interfaces interface-name unit logical-unit-number family inet6 dhcpv6-
client]
```

Description

Configure the DHCPv6 client identity association type.

Options

ia-na Identity association for nontemporary address

ia-pd Identity association for prefix delegation

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

The `logical-systems` and `tenants` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

DHCPv6 Client Overview

client-id (DHCP Local Server)

IN THIS SECTION

- Syntax | 4315
- Hierarchy Level | 4315
- Description | 4316
- Options | 4316
- Required Privilege Level | 4316
- Release Information | 4317

Syntax

```
client-id <exclude-headers> <use-automatic-ascii-hex-encoding>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6
authentication username-include],
```

```
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit system services dhcp-local-server dhcp authentication username-include],
[edit system services dhcp-local-server dhcp group group-name authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server dual-stack-group <group-name> authentication username-include]
```

Description

Specify that the client identifier (DHCP option 61) is concatenated with the username during the subscriber authentication or client authentication process.

Options

exclude-headers

By default, all headers that are part of the client identifier format in option 61 are included in the username string used for RADIUS authentication. Configure the `exclude-headers` option to exclude the use of headers in the username string.

use-automatic-ascii-hex-encoding

By default, all components of the client identifier are converted to ASCII hex to encode the username. Configure the `use-automatic-ascii-hex-encoding` option to use ASCII hex encoding only if there are non-ASCII characters in the client identifier.

Use this option instead of the `interface-name` option when the outer VLAN tag is unique across the system and you do not need the underlying physical interface name to be part of the format.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients

client-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 4317
- Hierarchy Level | 4317
- Description | 4318
- Options | 4318
- Required Privilege Level | 4318
- Release Information | 4318

Syntax

```
client-id <exclude-headers> <use-automatic-ascii-hex-encoding>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
```

```
options dhcp-relay dhcpv6 ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 ...]
```

Description

Specify that the client identifier (DHCP option 61) is concatenated with the username during the subscriber authentication or client authentication process.

Options

exclude-headers

By default, all headers that are part of the client identifier format in option 61 are included in the username string used for RADIUS authentication. Configure the `exclude-headers` option to exclude the use of headers in the username string.

use-automatic-ascii-hex-encoding

By default, all components of the client identifier are converted to ASCII hex to encode the username. Configure the `use-automatic-ascii-hex-encoding` option to use ASCII hex encoding only if there are non-ASCII characters in the client identifier.

Use this option instead of the `interface-name` option when the outer VLAN tag is unique across the system and you do not need the underlying physical interface name to be part of the format.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Specifying Authentication Support

Creating Unique Usernames for DHCP Clients

client-identifier (DHCP Client)

IN THIS SECTION

- Syntax | 4319
- Hierarchy Level | 4319
- Description | 4319
- Options | 4319
- Required Privilege Level | 4320
- Release Information | 4320

Syntax

```
client-identifier {
    hardware-type type-number
    user-id {ascii ascii hexadecimal hexadecimal;
    use-interface-description {logical |device};
    prefix [host-name routing-instance-name];
}
```

Hierarchy Level

```
[edit interfaces name unit name family  inet dhcp],
[edit interfaces interface-range name unit name family inet dhcp]
```

Description

The DHCP server identifies a client by a client-identifier value.

Options

hardware-type Specify the type of hardware used for the interface on which the client access. Ethernet is the common hardware type and the value for Ethernet is 1.

- **Default:** Zero
- **Range:** Zero through 255

ascii Client identifier as an ASCII string

hexadecimal Client identifier as a hexadecimal string

prefix Add prefix to client-id option

- Values:
  - host-name—Add router host name to client-id option
  - logical-system-name—Add logical system name to client-id option
  - routing-instance-name—Add routing instance name to client-id option

use-interface-description Use the interface description

- Values:
  - device—Use the device interface description
  - logical—Use the logical interface description

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

Starting with Junos OS Release 17.3R1, on all SRX Series Firewalls and vSRX Virtual Firewall instances, the CLI option `dhcp-client` at `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy is changed to `dhcp` to keep consistent with other Junos platforms. There is no change in the functionality with the changed CLI option `dhcp`.

The option `hardware-type` is added in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

DHCPv6 Client Overview

client-identifier (DHCPv6 Client)

IN THIS SECTION

- Syntax | 4321
- Hierarchy Level | 4321
- Description | 4321
- Options | 4322
- Required Privilege Level | 4322
- Release Information | 4322

Syntax

```
client-identifier duid-type (duid-ll | duid-llt | vendor);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcpv6-client]
[edit logical-systems logical -system-name interfaces interface-name unit logical-unit-number
family inet6 dhcpv6-client]
[edit tenants tenant-name interfaces interface-name unit logical-unit-number family inet6 dhcpv6-
client]
```

Description

The DHCPv6 server identifies a client by a client-identifier value.

Options

duid-type The DHCPv6 client is identified by a DHCP unique identifier (DUID).

duid-ll Link Layer address.

duid-llt Link Layer address plus time.

vendor Vendor-assigned unique ID based on the enterprise number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

The `logical-systems` and `tenants` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

DHCPv6 Client Overview

client-identifier (Interfaces)

IN THIS SECTION

- Syntax | 4323
- Hierarchy Level | 4323
- Description | 4323
- Options | 4323
- Required Privilege Level | 4323
- Release Information | 4323

Syntax

```
client-identifier {
    (ascii string | hexadecimal string);
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family-name dhcp]
```

Description

Specify an ASCII or hexadecimal identifier for the Dynamic Host Configuration Protocol (DHCP) client. The DHCP server identifies a client by a client-identifier value.

Options

- `ascii` *`ascii`* —Identifier consisting of ASCII characters.
- `hexadecimal` *`hexadecimal`* —Identifier consisting of hexadecimal characters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Interfaces

client-identity (NETCONF TLS)

IN THIS SECTION

- Syntax | 4324
- Hierarchy Level | 4324
- Description | 4324
- Default | 4325
- Options | 4325
- Required Privilege Level | 4326
- Release Information | 4326

Syntax

```
client-identity client-id {
    fingerprint fingerprint;
    map-type (san-dirname-cn | specified);
    username username;
}
```

Hierarchy Level

```
[edit system services netconf tls]
```

Description

For NETCONF sessions over Transport Layer Security (TLS), configure the method to derive the NETCONF username for a given client certificate.

Each configured client must include a client's certificate fingerprint and a map type. If the fingerprint of a client's presented certificate matches the fingerprint for a configured client, then the Junos device uses the corresponding map type to derive the NETCONF username for that certificate. If the certificate fingerprint does not match that of any configured client, then the Junos device uses the default map type defined at the `[edit system services netconf tls default-client-identity]` hierarchy level to derive the

NETCONF username. If the certificate fingerprint does not match a configured client, and there is no default client identity configured, the device does not establish the NETCONF session.

Junos devices support local users and Lightweight Directory Access Protocol (LDAP) remote users for NETCONF sessions over TLS. The username must either have a user account defined locally on the device, or it must be authenticated by an LDAP server, which then maps it to a user template account that is defined locally on the Junos device.

Default

If you do not include the `client-identity` statement, then you must define a default client at the `[edit system services netconf tls default-client-identity]` hierarchy level, or the device does not establish the NETCONF session.

Options

client-id User-defined name that uniquely identifies the client.

fingerprint ***fingerprint*** Client's certificate fingerprint, which is a cryptographic hash of an X.509 certificate in x509c2n:tls-fingerprint format.

The fingerprint's first octet value is the hashing algorithm identifier as defined in RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2*. The remaining octets are the result of the hashing algorithm.

Acceptable hash algorithms and their identifiers are:

- md5: 1
- sha1: 2
- sha224: 3
- sha256: 4
- sha384: 5
- sha512: 6

map-type ***type*** Map type that defines how to derive the NETCONF username.

- Values:

- `san-dirname-cn`—Use the common name (CN) defined for the SubjectAltName's (SAN) DirName field (`DirName:/CN`) in the client certificate as the NETCONF username.

  If you specify `san-dirname-cn` as the map type, but the client certificate does not have a username in this field, the connection fails.

- `specified`—Use the NETCONF username defined in the `username` statement at the same hierarchy level.

`username` *`username`*

Username under whose access privileges the NETCONF operations are executed when `map-type specified` is configured.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

Statement introduced in Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

default-client-identity (NETCONF TLS) | **5185**

client-idle-timeout (Access Profile)

IN THIS SECTION

- Syntax | **4327**
- Hierarchy Level | **4327**
- Description | **4327**
- Options | **4327**

Required Privilege Level | 4327
Release Information | 4327

Syntax

```
client-idle-timeout minutes;
```

Hierarchy Level

```
[edit access profile profile-name session-options]
```

Description

Specify the grace period that begins after an authenticated user terminates all sessions and connections. Authentication is not required if a new connection is initiated during the grace period by the same user.

Options

minutes —Number of minutes of idle time that elapse before the session is terminated.

- **Range:** 10 through 255 minutes
- **Default:** 10 minutes

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

client-list

IN THIS SECTION

- Syntax | 4328
- Junos OS Hierarchy Level | 4328
- Junos OS Evolved Hierarchy Level | 4328
- Description | 4328
- Options | 4329
- Required Privilege Level | 4329
- Release Information | 4329

Syntax

```
client-list list-name {
    address address;
}
```

Junos OS Hierarchy Level

```
[edit services rpm twamp server]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring twamp server]
```

Description

Specify the list of allowed control client hosts that can connect to this server. Each entry is a Classless Interdomain Routing (CIDR) address (IP address plus mask) that represents a network of allowed hosts. You can configure more than one list, but you must configure at least one client address to enable TWAMP. Each list can contain up to 64 entries.

Options

list-name—Name of client address list.

address—Address and mask for an allowed client.

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Statement introduced in Junos OS Evolved 20.3R1.

IPv6 address support introduced in Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

client-list (SNMP)

IN THIS SECTION

- Syntax | 4330
- Hierarchy Level | 4330
- Description | 4330
- Options | 4330
- Required Privilege Level | 4330
- Release Information | 4330

Syntax

```
client-list client-list-name {
    ip-addresses;
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Define a list of SNMP clients.

Options

client-list-name—Name of the client list.

ip-addresses—IP addresses of the SNMP clients to be added to the client list,

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Adding a Group of Clients to an SNMP Community*

client-list-name

IN THIS SECTION

- Syntax | 4331
- Hierarchy Level | 4331
- Description | 4331
- Options | 4331
- Required Privilege Level | 4331
- Release Information | 4332

Syntax

```
client-list-name client-list-name;
```

Hierarchy Level

```
[edit snmp community community-name]
```

Description

Add a client list or prefix list to an SNMP community.

Options

client-list-name—Name of the client list or prefix list.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Adding a Group of Clients to an SNMP Community

client-name-filter

IN THIS SECTION

- Syntax | 4332
- Hierarchy Level | 4332
- Description | 4333
- Options | 4333
- Required Privilege Level | 4333
- Release Information | 4333

Syntax

```
client-name-filter client-name {
    count number;
    domain-name domain-name;
    separator special-character;
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Define client-name-related restrictions. Clients whose names follow these restrictions are authenticated on the server.

Options

client-name—Name of the client.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Obtaining Username and Role Information Through Firewall Authentication

client-negotiation-match (DHCPv6 Local Server)

IN THIS SECTION

- Syntax | 4334
- Hierarchy Level | 4334
- Description | 4334
- Options | 4334
- Required Privilege Level | 4334
- Release Information | 4334

Syntax

```
client-negotiation-match incoming-interface;
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name
overrides]
```

Description

Configure the match criteria the DHCPv6 local server uses to uniquely identify IPv6 subscribers or clients.

Options

incoming-interface Allow only one client device to connect on the interface. If the client device changes, the router deletes the existing client binding and creates a binding for the newly connected device.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Automatically Logging Out DHCPv6 Clients

Understanding Differences Between Legacy DHCP and Extended DHCP

Allowing Only One DHCP Client Per Interface

client-negotiation-match (DHCPv6 Relay Agent)

IN THIS SECTION

- Syntax | 4335
- Hierarchy Level | 4335
- Description | 4335
- Options | 4335
- Required Privilege Level | 4336
- Release Information | 4336

Syntax

```
client-negotiation-match incoming-interface;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides]
```

Description

Configure the match criteria the DHCPv6 relay agent uses to uniquely identify IPv6 subscribers or clients.

Options

incoming-interface Allow only one client device to connect on the interface. If the client device changes, the router deletes the existing client binding and creates a binding for the newly connected device.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Automatically Logging Out DHCPv6 Clients

Extended DHCP Relay Agent Overview

Allowing Only One DHCP Client Per Interface

client-response-ttl

IN THIS SECTION

- Syntax | 4336
- Hierarchy Level | 4337
- Description | 4337
- Options | 4337
- Required Privilege Level | 4337
- Release Information | 4337

Syntax

```
client-response-ttl number;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface (interface-name | interface-group)]
```

Description

Set the IP time-to-live (TTL) value in BOOTP response messages sent to a BOOTP client. If you do not include the `client-response-ttl` statement, the default is to leave the TTL field unchanged.

Options

number—IP time-to-live (TTL) value.

- **Range:** 1 through 255
- **Default:** Leave the TTL field in the BOOTP response message unchanged.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

client-session-timeout (Access Profile)

IN THIS SECTION

- Syntax | 4338
- Hierarchy Level | 4338
- Description | 4338
- Options | 4338
- Required Privilege Level | 4339
- Release Information | 4339

Syntax

```
client-session-timeout minutes;
```

Hierarchy Level

```
[edit access profile profile-name session-options]
```

Description

Specify the amount of time after which user sessions are terminated, regardless of user activity (also known as a forced or hard authentication timeout).

Options

minutes —Number of minutes after which user sessions are terminated.

- **Range:** 1 through 10,000 minutes
- **Default:** Off

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Obtaining Username and Role Information Through Firewall Authentication

client (System Services)

IN THIS SECTION

- Syntax | 4339
- Hierarchy Level | 4339
- Description | 4340
- Options | 4340
- Required Privilege Level | 4340
- Release Information | 4340

Syntax

```
client name;
```

Hierarchy Level

```
[edit system services webapi (System Services)]
```

Description

Configures the IP address of the client. For the Integrated ClearPass Authentication and Enforcement feature Web API process configuration, the client is the ClearPass Policy Manager (CPPM).

The SRX Series Web API process acts as an HTTP(S) server. The CPPM client sends POST request messages containing user authentication and identity information to the Web API process. The SRX Series Firewall accepts information only from the configured address of the client.

Options

name	IP Address of the client.

Required Privilege Level

system	To view this statement in the configuration.
system-control	To add this statement to the configuration.

Release Information

Statement introduced in Junos OS release 12.3X48-D30.

client-type

IN THIS SECTION

- Syntax | 4341
- Hierarchy Level | 4341
- Description | 4341
- Options | 4341
- Required Privilege Level | 4341
- Release Information | 4341

Syntax

```
client-type (autoconfig | stateful);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6 dhcpv6-client]
[edit logical-systems logical -system-name interfaces interface-name unit logical-unit-number
family inet6 dhcpv6-client]
[edit tenants tenant-name interfaces interface-name unit logical-unit-number family inet6 dhcpv6-
client]
```

Description

The type of DHCPv6 client.

Options

- autoconfig—Autoconfig client type for router advertisement
- stateful— Stateful client type for address assignment

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

The `logical-systems` and `tenants` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

DHCPv6 Client Overview

client-type (Subscriber Limits)

IN THIS SECTION

- Syntax | 4342
- Hierarchy Level | 4342
- Description | 4343
- Options | 4343
- Required Privilege Level | 4343
- Release Information | 4343

Syntax

```
client-type (any | dhcp | l2tp | pppoe) {
    chassis {
        limit limit;
    }
    fpc slot-number {
        limit limit;
        pic number {
            limit limit;
            port number {
                limit limit;
            }
        }
    }
}
```

Hierarchy Level

```
[edit system services resource-monitor subscribers-limit]
```

Description

Configure the maximum number of subscribers of the client type that are allowed to be logged in. You can configure limits for subscribers per chassis, per MPC, per MIC, and per port. When the configured maximum number of subscribers is logged in for any level, subsequent logins at that level are denied until the current number of subscribers drops below the maximum allowed.

Options

name Type of client for which subscriber limits are configured.

- Values:
  - `any`—Apply the limit to the sum of all DHCP, L2TP, and PPPoE clients.
  - `dhcp`—Apply the limit to DHCP clients.
  - `l2tp`—Apply the limit to L2TP clients.
  - `pppoe`—Apply the limit to PPPoE clients.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Limiting Subscribers by Client Type and Hardware Element with Resource Monitor

Resource Monitoring for Subscriber Management and Services

clients (SNMP Community)

IN THIS SECTION

- Syntax | 4344
- Hierarchy Level | 4344
- Description | 4344
- Default | 4344
- Options | 4345
- Required Privilege Level | 4345
- Release Information | 4345

Syntax

```
clients {
    address <restrict>;
}
```

Hierarchy Level

```
[edit snmp communitycommunity community-name]
```

Description

Specify the IPv4 or IPv6 addresses of the SNMP client hosts that are authorized to use this community.

Default

If you omit the `clients` statement, all SNMP clients using this community string are authorized to access the router.

Options

address—Address of an SNMP client that is authorized to access this router. You must specify an address, not a hostname. To specify more than one client, include multiple ***address*** options.

restrict—(Optional) Do not allow the specified SNMP client to access the router.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring SNMP Communities

clients

IN THIS SECTION

- Syntax | 4346
- Hierarchy Level | 4346
- Description | 4346
- Options | 4346
- Required Privilege Level | 4346
- Release Information | 4346

Syntax

```
clients client-name {
    accounting {
        secret password;
        response-cache-timeout seconds;
    }
    address client-address;
    <prefer-framed-ip-address>
    <prefer-framed-ipv6-prefix>
    source-interface interface ipv4-address address;
}
```

Hierarchy Level

```
[edit access radius]
```

Description

Configure a RADIUS client for each GGSN, PGW, or BNG that sends subscriber session requests to the MX Series router and identifies it as a RADIUS server.

Options

client-name Name for the client.

- **Range:** 1 through 50 alphanumeric characters. Allowed characters are a-z, A-Z, 0-9.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers

clients (Security)

IN THIS SECTION

- Syntax | 4347
- Hierarchy Level | 4347
- Description | 4347
- Options | 4348
- Required Privilege Level | 4348
- Release Information | 4348

Syntax

```
clients configuration-name {
    ipsec-vpn vpn-name;
    remote-exceptions ip-address/mask;
    remote-protected-resources ip-address/mask;
    user username;
    user-groups user-group-name;
}
```

Hierarchy Level

```
[edit security dynamic-vpn]
```

Description

Create a client configuration for the dynamic VPN feature. Within the configuration, specify a name for the configuration, reference a standard VPN configuration to use for IPsec negotiations, specify which

resources to protect, define any exceptions, and list the users to which the dynamic VPN configuration applies.

In Junos OS Release 21.4R1, we've deprecated the `clients` configuration statement and we might remove it completely in a future release.

Options

configuration-name Name of the client configuration.

ipsec-vpn Use this statement to specify which IPsec VPN configuration the dynamic VPN feature should use to secure traffic.

remote-exceptions Use this statement to specify exceptions to the remote protected resources list for the specified dynamic VPN configuration. Traffic to the specified IP address will not go through the dynamic VPN tunnel and therefore will not be protected by the firewall's security policies.

remote-protected-resources Use this statement to specify which resources to protect using the dynamic VPN feature. Traffic to the protected resource will go through the specified dynamic VPN tunnel and will therefore be protected by the firewall's security policies.

user Specify which users can access the selected dynamic VPN configuration.

user-group Specify which users can access the selected dynamic VPN configuration.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Statement `user-groups` introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Dynamic VPN Overview

clns-routing

IN THIS SECTION

- Syntax | 4349
- Hierarchy Level | 4349
- Description | 4349
- Required Privilege Level | 4349
- Release Information | 4349

Syntax

```
clns-routing;
```

Hierarchy Level

```
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Enable IS-IS to exchange Connectionless Network Service (CLNS) routes. CLNS is a Layer 3 protocol, similar to IPv4. CLNS uses network service access points (NSAPs) to address end systems and intermediate systems.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IS-IS for CLNS

Example: Configuring Static Routes for CLNS When No IGP is Present

Example: Configuring BGP for CLNS VPNs

Understanding BGP for CLNS VPNs

Understanding Static Routes for CLNS

clns-updown-compatibility

IN THIS SECTION

- Syntax | 4350
- Hierarchy Level | 4350
- Description | 4351
- Required Privilege Level | 4351
- Release Information | 4351

Syntax

```
clns-updown-compatibility;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

When you enable IS-IS to exchange Connectionless Network Service (CLNS) routes, Junos OS sets the reserved (R) bit in the default metric field inside type, length, and value (TLV) type-3 (ES-Neighbor) as a marker for routing loop prevention. Junos OS uses the up/down bit for marking prefixes on the Level 2-to-Level 1 boundary as being propagated Down, such that any router in that area never propagates it Up on a Level 1-to-Level 2 boundary. For detailed information about how this works in IP routing environments, see RFC 2966, *Domain-wide Prefix Distribution with Two-Level IS-IS*.

Some other vendors' platforms might not support up/down bit setting in CLNS route TLVs. If one of these vendors' platforms receives this TLV with the R bit set, the platform discards the information.

When you use the `clns-updown-compatibility` statement in the IS-IS configuration, the R bit is set to 0, and the issue is resolved. The `clns-updown-compatibility` statement causes Junos OS to use the Internal/External metric-type bit in the TLV header instead of using the R bit as the up/down bit marker. This has the advantage that older end system (ES) equipment does not receive TLV headers with the R bit set.

> **CAUTION**: Not using the R bit can lead to potential routing loops. You can use the site-of-origin (SoO) extended community to prevent a looped BGP update from being injected back to IS-IS when received from a remote provider edge (PE) device. The receiving PE device can check against the SoO community, and if the value matches its own, the NLRI is not accepted.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4R7.

RELATED DOCUMENTATION

- *Example: Configuring IS-IS for CLNS*
- *Example: Configuring Static Routes for CLNS When No IGP is Present*
- *Example: Configuring BGP for CLNS VPNs*
- *Understanding BGP for CLNS VPNs*
- *Understanding Static Routes for CLNS*

clock-class

IN THIS SECTION

- Syntax | 4352
- Hierarchy Level | 4352
- Description | 4352
- Default | 4352
- Required Privilege Level | 4353
- Release Information | 4353

Syntax

```
clock-class clock-class-value;
```

Hierarchy Level

```
[edit protocols ptp client clock-class-to-quality-level-mapping ]
```

Description

Configure the clock class to the set ESMC quality level.

Default

Following are the default clock class values for various clocks:

- Boundary clock–248
- Ordinary clock (timeTransmitter)–52
- Ordinary clock (timeReceiver)–255

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2R2.

RELATED DOCUMENTATION

Configuring Hybrid Mode and ESMC Quality Level Mapping

Example: Configuring Hybrid Mode and ESMC Quality Level Mapping

Understanding Hybrid Mode

Precision Time Protocol Overview

Synchronous Ethernet Overview

clock-class-to-quality-level-mapping

IN THIS SECTION

- Syntax | 4353
- Hierarchy Level | 4354
- Description | 4354
- Required Privilege Level | 4354
- Release Information | 4354

Syntax

```
clock-class-to-quality-level-mapping {
    clock-class clock-class-value;
    {
```

```
        quality-level ql-value;
    }
}
```

Hierarchy Level

```
[edit protocols ptp slave]
```

Description

Configure the timeReceiver to override the default Precision Time Protocol (PTP) clock class to Ethernet Synchronization Message Channel (ESMC) mapping.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2R2.

RELATED DOCUMENTATION

Configuring Hybrid Mode and ESMC Quality Level Mapping

Example: Configuring Hybrid Mode and ESMC Quality Level Mapping

Understanding Hybrid Mode

Precision Time Protocol Overview

Synchronous Ethernet Overview

clock-ipv6-client

IN THIS SECTION

- Syntax | 4355
- Hierarchy Level | 4355
- Description | 4355
- Options | 4355
- Required Privilege Level | 4355
- Release Information | 4356

Syntax

```
clock-ipv6-client  slave-ip-address;
```

Hierarchy Level

```
[edit protocols ptp master interface interface-name unicast-mode transport ipv4]
```

Description

Configure the IPv6 address of the timeReceiver.

Options

slave-ip-address—The IPv6 address for the timeReceiver.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

clock-ipv6-source

IN THIS SECTION

- Syntax | 4356
- Hierarchy Level | 4356
- Description | 4357
- Options | 4357
- Required Privilege Level | 4357
- Release Information | 4357

Syntax

```
clock-ipv6-source master-ipv6-address {
    local-ip-address local-ip-address;
}
```

Hierarchy Level

```
[edit protocols ptp client interface interface-name unicast-mode transport]
```

Description

Configure the IPv6 address of the timeTransmitter.

Options

master-ipv6-address—IPv6 address for the timeTransmitter.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

clock-mode

IN THIS SECTION

- Syntax | 4358
- Hierarchy Level | 4358
- Description | 4358
- Options | 4358
- Required Privilege Level | 4358

Release Information | 4358

Syntax

```
clock-mode (boundary | ordinary);
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

Configure the clock mode as either boundary clock or ordinary clock. The clock mode determines whether the node is going to act as a timeReceiver, or timeTransmitter, or both. This attribute is mandatory and has no default value.

Options

`boundary`—The clock mode of the node is a boundary clock where the clock acts as both timeTransmitter and timeReceiver.

`ordinary`—The clock mode of the node is a system clock where the clock acts either as a timeTransmitter or as a timeReceiver. The `ordinary` option is not supported on ACX710 devices.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

IEEE 1588v2 Precision Timing Protocol (PTP)

Configuring the Precision Time Protocol G.8275.2 Enhanced Profile (Telecom Profile)

clock-mode (Clock Synchronization)

IN THIS SECTION

- Syntax | 4359
- Hierarchy Level | 4359
- Description | 4359
- Default | 4360
- Options | 4360
- Required Privilege Level | 4360
- Release Information | 4360

Syntax

```
clock-mode (auto-select | free-run);
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Configure the mode of operation to select the clock source from a free-run local oscillator or from an external qualified clock. On MX240, MX480, and MX960 routers with enhanced MPCs, the free-run

clock is provided by a local oscillator. On other MX Series routers, the free-run clock is provided by the SCB.

Default

auto-select

Options

auto-select Select the best external clock source as a clock source.

free-run Select the free-run local oscillator as a clock source.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

clock-rate

IN THIS SECTION

- Syntax | 4361
- Hierarchy Level | 4361
- Description | 4361
- Options | 4361

Required Privilege Level | 4362
Release Information | 4362

Syntax

```
clock-rate rate;
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

For EIA-530 and V.35 interfaces, configure the interface speed, in megahertz (MHz).

Options

rate—You can specify one of the following rates:

- 2.048 MHz
- 2.341 MHz
- 2.731 MHz
- 3.277 MHz
- 4.096 MHz
- 5.461 MHz
- 8.192 MHz
- 16.384 MHz
- **Default:** 16.384 MHz

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Clocking Mode

clock-source (hybrid)

IN THIS SECTION

- Syntax | 4362
- Hierarchy Level | 4363
- Description | 4363
- Options | 4363
- Required Privilege Level | 4363
- Release Information | 4363

Syntax

```
clock-source ip-address {
    interface interface1-name;
    interface interface2-name;
}
```

Hierarchy Level

```
[edit protocols ptp slave hybrid synchronous-ethernet-mapping]
```

Description

Configure the IP address of the PTP timeTransmitter and its possible Synchronous Ethernet source interfaces.

Options

`interface` *`interface1-name`*—Synchronous Ethernet interface traceable to the same PTP timeTransmitter clock.

`interface` *`interface2-name`*—Synchronous Ethernet interface traceable to the same PTP timeTransmitter clock.

> **NOTE**: You must first configure these interfaces at the `[edit chassis synchronization]` hierarchy level as Synchronous Ethernet sources. For information about configuring the interfaces, see "synchronization (MX Series)" on page 28795.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2R2.

RELATED DOCUMENTATION

- *Configuring Hybrid Mode and ESMC Quality Level Mapping*
- *Example: Configuring Hybrid Mode and ESMC Quality Level Mapping*
- *Understanding Hybrid Mode*
- *Precision Time Protocol Overview*
- *Synchronous Ethernet Overview*

clocking-mode

IN THIS SECTION

- Syntax | 4364
- Hierarchy Level | 4364
- Description | 4364
- Options | 4364
- Required Privilege Level | 4365
- Release Information | 4365

Syntax

```
clocking-mode (dce | internal | loop);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

For EIA-530 and V.35 interfaces, configure the clock mode. You cannot configure `clocking-mode dce` on a DTE router using an X.21 serial line protocol (detected automatically when an X.21 cable is plugged into the serial interface).

Options

`dce`—DCE timing (DTE mode only, not valid for X.21).

`internal`—Internal baud timing.

`loop`—Loop timing.

- **Default:** `loop`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Clocking Mode

close-timeout

IN THIS SECTION

- Syntax | 4365
- Hierarchy Level | 4366
- Description | 4366
- Options | 4366
- Required Privilege Level | 4366
- Release Information | 4366

Syntax

```
close-timeout seconds;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
[edit services service-set service-set-name service-set-options tcp-session
```

Description

Configure the timeout period for Transmission Control Protocol (TCP) session tear-down.

Options

seconds

Timeout period.

- **Default:** 1 second
- **Range:** 2 through 300 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Next Gen Services added in Junos OS Release 19.3R2 on MX Series MX240, MX480 and MX960 using MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring Default Timeout Settings for Services Interfaces

cluster

IN THIS SECTION

- Syntax | 4367
- Hierarchy Level | 4367
- Description | 4368
- Options | 4368
- Required Privilege Level | 4368
- Release Information | 4368

Syntax

```
cluster cluster-identifier;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.

CAUTION: If you configure both route reflection and VPNs on the same routing device, the following modifications to the route reflection configuration cause current BGP sessions to be reset:

- Adding a cluster ID—If a BGP session shares the same AS number with the group where you add the cluster ID, all BGP sessions are reset regardless of whether the BGP sessions are contained in the same group.
- Creating a new route reflector—If you have an IBGP group with an AS number and create a new route reflector group with the same AS number, all BGP sessions in the IBGP group and the new route reflector group are reset.

NOTE: If you change the address family specified in the `[edit protocols bgp family]` hierarchy level, all current BGP sessions on the routing device are dropped and then reestablished.

Options

cluster-identifier—4-byte number (such as an IPv4 address).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring BGP Route Reflectors

Understanding External BGP Peering Sessions

no-client-reflect | **11424**

cluster (Chassis)

IN THIS SECTION

- Syntax | 4369
- Hierarchy Level | 4371
- Description | 4371
- Options | 4371
- Required Privilege Level | 4372
- Release Information | 4373

Syntax

```
cluster {
    configuration-synchronize (Chassis Cluster) {
        no-secondary-bootup-auto;
    }
    control-interface {
        node name {
            disable;
        }
    }
    control-link-recovery;
    control-ports fpc {
        port;
    }
    health-monitoring;
    heartbeat-interval milliseconds;
    heartbeat-threshold heartbeat-threshold;
    network-management {
        cluster-master;
    }
    redundancy-group (Chassis Cluster) name {
        gratuitous-arp-count gratuitous-arp-count;
        hold-down-interval seconds;
        interface-monitor name {
            weight weight;
```

```
        }
        ip-monitoring {
            family {
                inet {
                    address name {
                        interface logical-interface-name {
                            secondary-ip-address;
                        }
                        weight weight;
                    }
                }
            }
            global-threshold global-threshold;
            global-weight global-weight;
            retry-count retry-count;
            retry-interval (Chassis Cluster) retry-interval;
        }
        node (Chassis Cluster Redundancy Group) (0 | 1) {
            priority priority;
        }
        preempt (Chassis Cluster) {
            delay seconds;
            limit limit;
            period seconds;
        }
    }
    redundancy-mode(active-active | active-backup);
    redundant-interface name {
        mapping-interface mapping-interface;
    }
    reth-count (Chassis Cluster) reth-count;
    traceoptions (Chassis Cluster) {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        level (alert | all | critical | debug | emergency | error | info | notice | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure a chassis cluster. Perform the configuration under the `[edit chassis cluster]` configuration stanza to define chassis cluster configuration, operations, and monitoring. This configuration must specify configuration synchronization, control link recovery, heartbeat interval and threshold, network management, redundancy group, and traceoptions.

Options

configuration-synchronize

Disable automatic chassis cluster synchronization. See "configuration-synchronize (Chassis Cluster)" on page 4519.

control-interface

Disable and Enable chassis cluster control-link interface on SRX1500, SRX4100, SRX4200, and SRX4600 devices. If you want to reboot any of the nodes, then you must disable the control interfaces on both nodes using the CLI configuration to ensure control link remains disabled after reboot. When control link is down, configuration change on local node will not be synchronized to the remote node.

From configuration mode, run the `set chassis cluster control-interface <node0/node1> disable` on node 0 or node 1 to disable the control link and to enable the control link run the `delete chassis cluster control-interface <node0/node1> disable` on both the nodes. If you have disabled the links using the configuration command, then the links will remain disabled even after system reboot.

From the operational mode, run the `request chassis cluster control-interface <node0/node1> disable` or the `request chassis cluster control-interface <node0/node1> enable` commands to enable or disable the control link from the local node. If you have disabled control link using the operational mode CLI commands, the links will be enabled after system reboot.

control-link-recovery

Enable automatic control link recovery option.

control-ports

Enable specific chassis cluster control ports.

- Values:

- `fpc`—FPC slot number
- `port`—Port number

health-monitoring

Enable to monitor the health status of the SRX Series devices operating in chassis cluster mode. The health status between the two nodes is monitored and shared over control links and fabric links. Failover between the nodes occurs based on the heart beat status and health status of the control links and fabric links. By default, the option is disabled.

heartbeat-interval

Interval between successive heartbeats (milliseconds)

- **Default:** `1000`
- **Range:** `1000-2000`

heartbeat-threshold

Number of consecutive missed heartbeats to indicate device failure

- **Default:** `3`
- **Range:** `3-8`

network-management

Define parameters for network management. See "network-management" on page 11317.

redundancy-group *name*

Define a redundancy group. See "redundancy-group (Chassis Cluster)" on page 14173.

redundancy-mode

Redundancy mode for the cluster (Reboot all nodes simultaneously after change)

- active-active—Primary and secondary nodes in active mode.
- active-backup—Primary in active, secondary in backup mode.

reth-count

Number of redundant ethernet interfaces

- **Range:** `1-128`

traceoptions

Define chassis cluster redundancy process tracing operations. See "traceoptions (Chassis Cluster)" on page 29506.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

The `health-monitoring` option is introduced in Junos OS Release 18.4R1.

The `control-interface` option is introduced in Junos OS Release 20.4R1 for SRX1500 devices.

The `Redundancy mode` option is introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

ip-monitoring | 9244

show chassis cluster information | 17364

cluster-id (Junos Fusion Satellite Device Cluster)

IN THIS SECTION

- Syntax | 4373
- Hierarchy Level | 4374
- Description | 4374
- Default | 4374
- Options | 4374
- Required Privilege Level | 4374
- Release Information | 4374

Syntax

```
cluster-id cluster-id-number;
```

Hierarchy Level

```
[edit chassis satellite-management cluster cluster-name]
```

Description

Assign a cluster identification number to a satellite device cluster in a Junos Fusion.

The cluster identification number is used to identify a satellite device cluster in a Junos Fusion.

Default

Cluster identification numbers are not assigned in a Junos Fusion, by default.

Options

cluster-id-number Specifies the cluster identification number of the satellite device cluster in the Junos Fusion.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

cluster (Junos Fusion)

IN THIS SECTION

- Syntax | 4375
- Hierarchy Level | 4375
- Description | 4375
- Default | 4376
- Options | 4376
- Required Privilege Level | 4376
- Release Information | 4376

Syntax

```
cluster cluster-name{
    cascade-ports interface-name;
    cluster-id cluster-id-number;
    fpc slot-id{
        alias alias;
        description description;
        member-id member-id-number;
        system-id mac-address;
    }
}
```

Hierarchy Level

```
[edit chassis satellite-management ]
```

Description

Create and name a satellite device cluster.

Default

Satellite device clusters are not present, by default.

Options

cluster-name Specifies the name of the satellite device cluster.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

cmcast-joins-limit-inet (MVPN Selective Tunnels)

IN THIS SECTION

- Syntax | 4377
- Hierarchy Level | 4377
- Description | 4377
- Default | 4378
- Options | 4378
- Required Privilege Level | 4378
- Release Information | 4378

Syntax

```
cmcast-joins-limit-inet number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name provider-tunnel
selective],
[edit routing-instances instance-name provider-tunnel selective]
```

Description

Configure the maximum number of IPv4 customer multicast entries

The purpose of the `cmcast-joins-limit-inet` statement is to supplement the multicast forwarding-cache limit when the MVPN `rpt-spt` mode is configured and when traffic is flowing through selective service provider multicast service inteface (S-PMSI) tunnels and is forwarded by way of the (*,G) entry, even though the forwarding cache limit has already blocked the forwarding entries from being created.

The `cmcast-joins-limit-inet` statement limits the number of Type-6 and Type-7 routes. These routes contain customer-route control information.

You can configure the `cmcast-joins-limit-inet` statement only when the MVPN mode is `rpt-spt`.

This statement is independent of the `leaf-tunnel-limit-inet` statement and of the `forwarding-cache threshold` statement.

The `cmcast-joins-limit-inet` statement is applicable on the egress PE router. It limits the customer multiccast entries created in response to PIM (*,G) and (S,G) join messages. This statement is applicable to both type-6 and type-7 routes because the intention is to limit the egress forwarding entries, and in `rpt-spt` mode, an MVPN creates forwarding entries for both of these route types (in other words, for both (*,G) and (S,G) entries). However, this statement does not block BGP-created customer multicast entries because the purpose of this statement is to prevent the creation of forwarding entries on the egress PE router only and only for non-remote receivers. If remote-side customer multicast entries or forwarding entries need to be limited, you can use forwarding-cache threshold on the ingress routers, in which case this statement is not required.

By placing a limit on the customer multicast entries, you can ensure that when the limit is reached or the maximum forwarding state is created, all further local join messages will be blocked by the egress PE router. This ensures that traffic is flowing for only those multicast entries that are permitted.

If another PE router is interested in the traffic, it might pull the traffic from the ingress PE router by sending type-6 and type-7 routes. To prevent forwarding in this case, you can configure the leaf tunnel limit (`leaf-tunnel-limit-inet`). By preventing type-4 routes from being sent in response to type-3 routes, the formation of selective tunnels is blocked when the tunnel limit is reached. This ensures that traffic flows only for the routes within the tunnel limit . For all other routes, traffic flows only to the PE routers that have not reached the configured limit.

Setting the `cmcast-joins-limit-inet` statement or reducing the value of the limit does not alter or delete the already existing and installed routes. If needed, you can run the `clear pim join` command to force the limit to take effect. Those routes that cannot be processed because of the limit are added to a queue, and this queue is processed when the limit is removed or increased and when existing routes are deleted.

Default

Unlimited

Options

number Maximum number of customer multicast entries for IPv4.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Examples: Configuring the Multicast Forwarding Cache

Example: Configuring MBGP Multicast VPN Topology Variations

cmcast-joins-limit-inet6 (MVPN Selective Tunnels)

IN THIS SECTION

- Syntax | 4379
- Hierarchy Level | 4379
- Description | 4379
- Default | 4380
- Options | 4380
- Required Privilege Level | 4380
- Release Information | 4381

Syntax

```
cmcast-joins-limit-inet6 number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name provider-tunnel
selective],
[edit routing-instances instance-name provider-tunnel selective]
```

Description

Configure the maximum number of IPv4 customer multicast entries

The purpose of the `cmcast-joins-limit-inet6` statement is to supplement the multicast forwarding-cache limit when the MVPN `rpt-spt` mode is configured and when traffic is flowing through selective service provider multicast service inteface (S-PMSI) tunnels and is forwarded by way of the (*,G) entry, even though the forwarding cache limit has already blocked the forwarding entries from being created.

The `cmcast-joins-limit-inet6` statement limits the number of Type-6 and Type-7 routes. These routes contain customer-route control information.

You can configure the `cmcast-joins-limit-inet6` statement only when the MVPN mode is `rpt-spt`.

This statement is independent of the `leaf-tunnel-limit-inet6` statement and of the `forwarding-cache threshold` statement.

The `cmcast-joins-limit-inet6` statement is applicable on the egress PE router. It limits the customer multiccast entries created in response to PIM (*,G) and (S,G) join messages. This statement is applicable to both type-6 and type-7 routes because the intention is to limit the egress forwarding entries, and in `rpt-spt` mode, an MVPN creates forwarding entries for both of these route types (in other words, for both (*,G) and (S,G) entries). However, this statement does not block BGP-created customer multicast entries because the purpose of this statement is to prevent the creation of forwarding entries on the egress PE router only and only for non-remote receivers. If remote-side customer multicast entries or forwarding entries need to be limited, you can use forwarding-cache threshold on the ingress routers, in which case this statement is not required.

By placing a limit on the customer multicast entries, you can ensure that when the limit is reached or the maximum forwarding state is created, all further local join messages will be blocked by the egress PE router. This ensures that traffic is flowing for only those multicast entries that are permitted.

If another PE router is interested in the traffic, it might pull the traffic from the ingress PE router by sending type-6 and type-7 routes. To prevent forwarding in this case, you can configure the leaf tunnel limit (`leaf-tunnel-limit-inet6`). By preventing type-4 routes from being sent in response to type-3 routes, the formation of selective tunnels is blocked when the tunnel limit is reached. This ensures that traffic flows only for the routes within the tunnel limit . For all other routes, traffic flows only to the PE routers that have not reached the configured limit.

Setting the `cmcast-joins-limit-inet6` statement or reducing the value of the limit does not alter or delete the already existing and installed routes. If needed, you can run the `clear pim join` command to force the limit to take effect. Those routes that cannot be processed because of the limit are added to a queue, and this queue is processed when the limit is removed or increased and when existing routes are deleted.

Default

Unlimited

Options

number Maximum number of customer multicast entries for IPv4.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Examples: Configuring the Multicast Forwarding Cache

Example: Configuring MBGP Multicast VPN Topology Variations

coa-accounting (AAA Profile)

IN THIS SECTION

- Syntax | 4381
- Hierarchy Level | 4381
- Description | 4382
- Options | 4382
- Required Privilege Level | 4382
- Release Information | 4382

Syntax

```
coa-accounting (enable | disable);
```

Hierarchy Level

```
[edit unified-edge aaa profiles aaa-profile-name radius policy]
```

Description

Enable or disable the initiation of a RADIUS accounting start from the MX Series router to the RADIUS server. Enabling this feature is required if the RADIUS server cannot initiate a change of authorization request without an accounting record. Specifying **enable** does not cause the MX Series router to report any billing information.

Options

enable Initiate a RADIUS accounting start from the MX Series Router to the RADIUS server.

disable Do not initiate a RADIUS accounting start from the MX Series Router to the RADIUS server.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring an AAA Profile

Understanding AAA Profiles

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

coa-immediate-update

IN THIS SECTION

Syntax | 4383

Hierarchy Level | 4383

Description | 4383
Required Privilege Level | 4383
Release Information | 4383

Syntax

```
coa-immediate-update;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router to send an Acct-Update message to the RADIUS accounting server immediately following a CoA operation.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

Configuring Per-Subscriber Session Accounting

coa-no-override service-class-attribute

IN THIS SECTION

- Syntax | 4384
- Hierarchy Level | 4384
- Description | 4384
- Required Privilege Level | 4384
- Release Information | 4384

Syntax

```
coa-no-override service-class-attribute;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Specify that, after a CoA action that changes the RADIUS Class attribute, accounting reports for the subscriber's service sessions continue to use the original Class attribute that was assigned when the service sessions were created. The new Class attribute value is used in accounting reports for the subscriber session only.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

Configuring Per-Subscriber Session Accounting

code

IN THIS SECTION

- Syntax | 4385
- Hierarchy Level | 4385
- Description | 4385
- Options | 4386
- Required Privilege Level | 4386
- Release Information | 4386

Syntax

```
code numeric-code;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name]
```

Description

Specify the custom attribute's AVP code for the incoming RADIUS request from the subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

Options

numeric-code Numeric value for the code.

- **Range:** 0 through 255.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

code (IDP)

IN THIS SECTION

- Syntax | 4387
- Hierarchy Level | 4387
- Description | 4387
- Options | 4387
- Required Privilege Level | 4387
- Release Information | 4387

Syntax

```
code {
    match (equal | greater-than | less-than | not-equal);
    value code-value;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol icmp],

[edit security idp custom-attack attack-name attack-type signature protocol icmpv6]
```

Description

Specify the secondary code that identifies the function of the request/reply within a given type.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`code-value`*—Match a decimal value.
- **Range:** 0 through 255

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3. Statement modified in Junos OS Release 12.3X48-D25 to add ICMPv6 protocol support.

code (AAA Profile)

IN THIS SECTION

- Syntax | 4388
- Hierarchy Level | 4388
- Description | 4388
- Options | 4388
- Required Privilege Level | 4389
- Release Information | 4389

Syntax

```
code numeric-code;
```

Hierarchy Level

```
[edit unified-edge aaa profiles aaa-profile-name radius policy activation-attribute],
[edit unified-edge aaa profiles aaa-profile-name radius policy deactivation-attribute]
```

Description

Configure the RADIUS attributes that you want to carry the PCC rulebase name for rulebase activations and deactivations from the RADIUS policy server to the MX Series router. By default, the rulebase name is carried in the ERX-Service-Activate Juniper vendor-specific attribute (VSA) for activations and in the ERX-Service-Deactivate Juniper VSA for deactivations.

Options

numeric-code	Numeric value for the RADIUS AVP.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring an AAA Profile

Understanding AAA Profiles

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

code (Application Identification)

IN THIS SECTION

- Syntax | 4389
- Hierarchy Level | 4390
- Description | 4390
- Options | 4390
- Required Privilege Level | 4390
- Release Information | 4390

Syntax

```
code icmp-code;
```

Hierarchy Level

```
[edit services application-identification application application-name icmp-mapping]
```

Description

Match the specified ICMP code to create a custom application signature.

Options

value Numeric value for the ICMP code.

- **Range:** 0 through 254

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

code-point

IN THIS SECTION

Syntax | 4391

- Hierarchy Level | 4391
- Description | 4391
- Options | 4391
- Required Privilege Level | 4391
- Release Information | 4391

Syntax

```
code-point [ aliases ] [ bit-patterns ];
```

Hierarchy Level

```
[edit class-of-service rewrite-rules type rewrite-name forwarding-class class-name]
```

Description

Specify one or more code-point aliases or bit sets for association with a forwarding class.

Options

aliases—Name of each alias.

bit-patterns—Value of the code-point bits, in decimal form.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Rewrite Rules

code-point-aliases

IN THIS SECTION

- Syntax | 4392
- Hierarchy Level | 4392
- Description | 4392
- Options | 4393
- Usage Guidelines | 4393
- Required Privilege Level | 4393
- Release Information | 4393

Syntax

```
code-point-aliases {
    type {
        alias-name bits;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Define an alias for a CoS marker. You can use the alias instead of the bit pattern when you specify the code point during configuration.

Options

alias-name—Name of the code-point alias.

bits—6-bit value of the code-point bits, in decimal form.

type—CoS marker type.

- **Values:** `dscp`, `dscp-ipv6`, `exp`, `ieee-802.1`, `ieee-802.1ad`, `inet-precedence`

Usage Guidelines

See *Defining Aliases for CoS Value Bit Patterns*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Defining Aliases for CoS Value Bit Patterns

code-point (Congestion Notification)

IN THIS SECTION

- Syntax | 4394
- Hierarchy Level | 4394
- Description | 4394
- Options | 4394
- Required Privilege Level | 4394

Release Information | 4394

Syntax

```
code-point up-bits pfc;
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name input ieee-802.1],
[edit class-of-service interfaces interface-name congestion-notification-profile profile-name
input ieee-802.1]
```

Description

Configure the IEEE 802.1p (User Priority) code point bits as input for creating the priority-based flow control (PFC) congestion notification profile, which you will associate with a particular traffic class.

Options

- pfc—PFC flow control method
- *up-bits*—Three-bit pattern of the User Priority field in an IEEE 802.1Q tag

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring an FCoE Transit Switch

Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)

code-point (Fibre Channel Interfaces)

IN THIS SECTION

- Syntax | 4395
- Hierarchy Level | 4395
- Description | 4395
- Options | 4396
- Required Privilege Level | 4396

Syntax

```
code-point code-point-bits;
```

Hierarchy Level

```
[edit class-of-service interfaces fibre-channel-interface-name rewrite-value input ieee-802.1]
```

Description

Configure the IEEE 802.1p code point value assigned to all traffic received from the Fibre Channel (FC) network on the specified FC interface (NP_Port). When native FC traffic from the FC SAN arrives at the NP_Port interface, the NP_Port interface encapsulates it in Ethernet to create FCoE packets before forwarding the traffic onto the FCoE network. Instead of using the default value of priority 3 (code point 011) for the FCoE traffic, the interface rewrites the IEEE 802.1p code point to the value specified in the rewrite value code points.

After the code point value is rewritten, the interface forwards the traffic to the Ethernet (FCoE) network. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.

FCoE traffic typically uses priority 3 (IEEE code point 011). The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011).

The system supports only one IEEE 802.1p code point value per FC interface. You cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.

Options

code-point-bits—Value of the code-point bits, in decimal form.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

forwarding-class (Fibre Channel Interfaces) | **7421**

Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

Understanding CoS Classifiers

code-point (Input Congestion Notification)

IN THIS SECTION

- Syntax | 4397
- Hierarchy Level | 4397
- Description | 4397
- Options | 4398
- Required Privilege Level | 4398
- Release Information | 4398

Syntax

```
code-point [code-point-bits] {
    pfc {
        mru mru-value;
    }
    xon
}
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name input dscp]
[edit class-of-service congestion-notification-profile profile-name input ieee-802.1]
```

Description

Enable priority-based flow control (PFC) on an IEEE 802.1p code point (priority) or a Differentiated Services code point (DSCP) value (to implement DSCP-based PFC at Layer 3).

Use this statement to configure PFC to operate based on an IEEE 802.1p code point in the VLAN-tagged packet header at Layer 2, or on a DSCP value from the Layer 3 IP header to support PFC with untagged traffic.

The `pfc` statement is mandatory when setting this statement.

Options

code-point-bits—Code point bit pattern corresponding to an IEEE 802.1p 3-bit priority value when used in the `ieee-802.1` hierarchy, or a 6-bit DSCP value when used in the `dscp` hierarchy.

The remaining statements are described separately.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Support for DSCP values introduced in Junos OS Release 17.4R1 for the QFX Series.

RELATED DOCUMENTATION

- *Example: Configuring CoS PFC for FCoE Traffic*
- *Configuring CoS PFC (Congestion Notification Profiles)*
- *Understanding CoS Flow Control (Ethernet PAUSE and PFC)*
- *Understanding PFC Using DSCP at Layer 3 for Untagged Traffic*
- *Configuring DSCP-based PFC for Layer 3 Untagged Traffic*

code-point (Output Congestion Notification)

IN THIS SECTION

- Syntax | 4399
- Hierarchy Level | 4399
- Description | 4399
- Default | 4399
- Options | 4400

Required Privilege Level | 4400
Release Information | 4400

Syntax

```
code-point [ code-point-bits ] {
    flow-control-queue [ queue | list-of-queues ];
}
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name output ieee-802.1]
```

Description

Specify the IEEE 802.1p code point bits that identify the traffic you want to enable for priority-based flow control (PFC) pause.

Default

By default, IEEE 802.1p priorities 3 and 4 (code points 011 and 100, respectively) are enabled for PFC pause on all Ethernet interfaces. If you explicitly configure priorities to pause and the output queues on which to enable pause, the explicit configuration overrides the default configuration. When you apply an explicit output congestion notification profile to an interface, only the priorities and queues specified in the output congestion notification profile are enabled for pause on that interface.

For example, if you configure an output congestion notification profile that specifies priority 2 (code point 010), then traffic with IEEE 802.1p priority 2 is paused on the configured output queue during periods of congestion. However, traffic with priority 3 and priority 4 is not programmed to pause, because the explicit configuration overwrites the default configuration, and the explicit configuration does not pause priority 3 and priority 4. If you configure an explicit output congestion notification profile, all of the priorities you want to enable for PFC and all of the output queues you want to pause must be explicitly configured.

Options

code-point-bits—3-bit value in decimal form.

The remaining statements are described separately.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring CoS PFC (Congestion Notification Profiles)

Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces

Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface

Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)

Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)

Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows

code-point (Rewrite Rules)

IN THIS SECTION

- Syntax | 4401
- Hierarchy Level | 4401
- Description | 4401

- Options | 4401
- Required Privilege Level | 4401
- Release Information | 4401

Syntax

```
code-point [ alias ] [ bit-pattern ];
```

Hierarchy Level

```
[edit class-of-service rewrite-rules (dscp | ieee-802.1) forwarding-class class-name loss-priority
level]
```

Description

Configure a code-point alias or bit set to apply to a forwarding class for a rewrite rule.

Options

alias—Name of the alias.

bit-pattern—Value of the code-point bits, in decimal form.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Defining CoS Rewrite Rules

Understanding CoS Classifiers

Understanding CoS Classifiers

code-points (Application Maps)

IN THIS SECTION

- Syntax | 4402
- Hierarchy Level | 4402
- Description | 4402
- Options | 4402
- Required Privilege Level | 4403
- Release Information | 4403

Syntax

```
code-points [ aliases ] [ bit-patterns ];
```

Hierarchy Level

```
[edit policy-options application-maps application-map-name application application-name]
```

Description

Define one or more code-point aliases or bit sets for an application.

Options

aliases—Name of the alias or aliases.

bit-patterns—Value of the code-point bits, in decimal form.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- *Configuring an Application Map for DCBX Application Protocol TLV Exchange*
- *Example: Configuring DCBX Application Protocol TLV Exchange*
- Example: Configuring DCBX to Support an iSCSI Application
- *Understanding DCBX Application Protocol TLV Exchange*
- Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

code-points (Classifiers)

IN THIS SECTION

- Syntax | 4404
- Hierarchy Level | 4404
- Description | 4404
- Options | 4404
- Required Privilege Level | 4404
- Release Information | 4404

Syntax

```
code-points [ aliases ] [ 6 bit-patterns ];
```

Hierarchy Level

```
[edit class-of-service classifiers (dscp | ieee-802.1 | inet-precedence) forwarding-class class-name
loss-priority level]
```

Description

Specify one or more DSCP code-point aliases or bit sets for association with a forwarding class.

Options

aliases —Name of the DSCP alias.

6 bit-patterns —Value of the code-point bits, in decimal form.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches

Defining CoS Classifiers (CLI Procedure)

Defining CoS Classifiers (J-Web Procedure)

Understanding CoS Classifiers

code-points (Classifiers)

IN THIS SECTION

- Syntax | 4405
- Hierarchy Level | 4405
- Description | 4405
- Options | 4405
- Required Privilege Level | 4405
- Release Information | 4406

Syntax

```
code-points ([ aliases ] | [ bit-patterns ]);
```

Hierarchy Level

```
[edit class-of-service classifiers type classifier-name forwarding-class class-name loss-
priority level]
```

Description

Specify one or more DSCP code-point aliases or bit sets to apply to a forwarding class.

Options

aliases—Name of the DSCP alias.

bit-patterns—Value of the code-point bits, in six-bit binary form.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Understanding Interfaces
- *Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic*
- *Example: Configuring Behavior Aggregate Classifiers*
- Example: Configuring BA Classifiers on Transparent Mode Security Devices

collapsed-forward

IN THIS SECTION

- Syntax | 4406
- Hierarchy Level | 4406
- Description | 4407
- Options | 4407
- Required Privilege Level | 4407
- Release Information | 4407

Syntax

```
collapsed-forward (monitor | optimize);
```

Hierarchy Level

```
[edit chassis network-slices guest-network-functions gnf id af-name]
```

Description

Configure a fabric path optimization mode to optimize the fabric path of the traffic flowing over abstracted fabric (af) interfaces between two guest network functions (GNFs). After configuring fabric path optimization, you can use the command `show interfaces` *`af-interface-name`* in GNF to view the number of packets that are currently flowing on the optimal / non-optimal path.

NOTE:

- You must configure the `collapsed-forward` statement on both the peer guest network functions (GNFs) on the abstracted fabric interface for the commit to be successful.
- The `collapsed-forward` statement is configured at the abstracted fabric interface level. All logical interfaces on the abstracted fabric interface inherit this configuration.
- We strongly recommend that you do not configure the `collapsed-forward` statement on GNFs that have a mix of supported and unsupported MPCs (MPCs that support fabric path optimization and MPCs that do not) to avoid any erratic behavior.

Options

monitor If you configure this mode, the peer GNF monitors the traffic flow and sends information to the source GNF about the Packet Forwarding Engine to which the traffic is being forwarded currently and the desired Packet Forwarding Engine that could provide an optimized path. However, the source GNF does not forward the traffic towards the desired Packet Forwarding Engine.

optimize If you configure this mode, the peer GNF monitors the traffic flow and sends information to the source GNF about the Packet Forwarding Engine to which the traffic is being forwarded currently and the desired Packet Forwarding Engine that could provide an optimized path. The source GNF then forwards the traffic towards the desired Packet Forwarding Engine.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

Understanding Junos Node Slicing

network-slices | 11322

guest-network-functions | 7949

gnf | 7746

control-plane-bandwidth-percent (Node Slicing) | 4634

description (GNF) | 5332

collector

IN THIS SECTION

- Syntax (EX Series) | 4408
- Syntax (ACX Series, MX Series, PTX Series, and QFX Series) | 4408
- Syntax (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)) | 4409
- Hierarchy Level | 4409
- Description | 4409
- Options | 4409
- Required Privilege Level | 4410
- Release Information | 4410

Syntax (EX Series)

```
collector ip-v4/v6-address {
    udp-port port-number;
}
```

Syntax (ACX Series, MX Series, PTX Series, and QFX Series)

```
collector ip-v4/v6-address {
    dscp dscp;
```

```
    forwarding-class forwarding-class;
    udp-port port-number;
}
```

Syntax (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD))

```
collector ip-v4/v6-address {
    dscp dscp;
    forwarding-class forwarding-class;
    routing-instance routing-instance name;
    udp-port port-number;
}
```

Hierarchy Level

```
[edit protocols sflow]
```

Description

Configure a remote collector for sFlow network traffic monitoring. The device sends sFlow UDP datagrams to this collector for analysis. You can configure up to four collectors on the device. You configure a collector by specifying its IP address and a UDP port.

Options

dscp *dscp*	Configure Differentiated Services Code Point (DSCP) values per collector. These values are applied to all the packets destined for that collector. A change in the DSCP configuration resets all the counters for the collector and the new configuration comes into effect.
forwarding-class *forwarding-class*	Configure forwarding class values per collector. These values are applied to all the packets destined for that collector. A change in the forwarding-class configuration does not reset the counters. The collector maintains its state and the new configuration comes into effect.
ip-address (*ipv4-address* | *ipv6-address*)	Configure the IP address.
(*ipv4-address* | *ipv6-address*)	(EX Series only) Set the IPv4 or IPv6 address.

udp-port *port-number*

Configure the collector UDP port number.

- **Default:** 6343

routing-instance *routing-instance name*

Configure the routing instance on which collector is reachable. Options:

- mgmt_default
- mgmt_junos
- any user-defined VPN routing and forwarding (VRF) routing instance

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

`dscp` and `forwarding-class` options introduced in Junos OS Release 19.1R1.

`routing-instance` options introduced in Junos OS Evoloved Release 24.4R1.

RELATED DOCUMENTATION

Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches

Configuring sFlow Technology for Network Monitoring (CLI Procedure)

Example: Configuring sFlow Technology to Monitor Network Traffic on MX Series Routers

collector (Flow Collection)

IN THIS SECTION

Syntax | 4411

Hierarchy Level | 4411

- Description | 4411
- Options | 4411
- Required Privilege Level | 4411
- Release Information | 4411

Syntax

```
collector interface-name;
```

Hierarchy Level

```
[edit services flow-collector interface-map]
```

Description

Configure the default flow collector interface for interface mapping.

Options

interface-name—Default flow collector interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Interface Mappings

collector (Analytics)

IN THIS SECTION

- Syntax | 4412
- Hierarchy Level | 4412
- Description | 4413
- Required Privilege Level | 4413
- Release Information | 4413

Syntax

```
collector {
    local {
        file filename {
            size size;
            files number;
        }
    }
    address ip-address {
        port number {
            transport protocol {
                export-profile profile-name;
            }
        }
    }
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Configure a local file for storing network analytics statistics and/or a remote server for receiving streamed statistics data.

> **NOTE**: The `collector` statement is not available in Junos OS Releases prior to 13.2X51-D15.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

collector (Flow Monitoring Logs for NAT)

IN THIS SECTION

- Syntax | 4414
- Hierarchy Level | 4414
- Description | 4414
- Options | 4414
- Required Privilege Level | 4414
- Release Information | 4414

Syntax

```
collector collector-name {
    source-ip address;
    destination-address address;
    destination-port port-number;
}
```

Hierarchy Level

```
[edit services jflow-log]
```

Description

Specify the name of the collector to which flow monitoring log messages in IPFIX or version 9 flow template format for NAT events must be sent. The generated flow monitoring logs for NAT events in flow template format are sent to the specified host or external device that functions as the NetFlow collector. You must associate a collector with a template profile for the template characteristics, such as refresh rate of messages and the template format, to be used for generated flow monitoring logs.

Options

collector-name—Name of the collector to which flow monitoring log messages for NAT events in flow monitoring format (IPFIX or version 9 flow template format) must be sent. The name can be up to 32 alphanumeric characters in length. Allowed characters are [a-zA-Z0-9_]

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

collector (Flow Template Profiles for NAT)

IN THIS SECTION

- Syntax | 4415
- Hierarchy Level | 4415
- Description | 4415
- Options | 4416
- Required Privilege Level | 4416
- Release Information | 4416

Syntax

```
collector collector-name;
```

Hierarchy Level

```
[edit services jflow-log template-profile template-profile-name]
```

Description

Specify the name of the collector to be associated with a template profile. The generated flow monitoring logs for NAT events in flow template format are sent to the specified collector. You must

have previously configured the collector by using the `collector collector-name` statement at the `[edit services jflow-log]` hierarchy level before you associate a collector with a template profile.

Options

`collector-name`—Name of the collector to which flow monitoring log messages for NAT events in flow monitoring format (IPFIX or version 9 flow template format) must be sent. The name can be up to 32 alphanumeric characters in length. Allowed characters are [a-zA-Z0-9_]

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

collector-group (Flow Monitoring Logs for NAT)

IN THIS SECTION

- Syntax | **4417**
- Hierarchy Level | **4417**
- Description | **4417**
- Options | **4417**

Required Privilege Level | 4418
Release Information | 4418

Syntax

```
collector-group collector-group-name {
    [collector-name1 collector-name2];
}
```

Hierarchy Level

```
[edit services jflow-log]
```

Description

Specify the name of the collector group that contains a set of NetFlow collectors to which flow monitoring log messages in IPFIX or version 9 flow template format for NAT events must be sent. You must define at least one collector in the group. A maximum of up to eight collectors can be aggregated into a collector group.

The generated flow monitoring logs for NAT events in flow template format are sent to the specified collector group. By using a collector group, you can effectively and optimally transmit flow monitoring logs to a cluster of collectors in a single, one-step operation.

Options

collector-group-name—Name of the collector group to which flow monitoring log messages for NAT events in flow monitoring format (IPFIX or version 9 flow template format) must be sent. The name can be up to 32 alphanumeric characters in length. Allowed characters are [a-zA-Z0-9_]

collector-name—Name of the collector to be assigned to the group of collectors. You must have previously defined the collector by including the `collector` *collector-name* statement at the `[edit services jflow-log]` hierarchy level. You can specify a list of valid collector names. Specify the names individually by using a space to separate each collector name. The name can be up to 32 alphanumeric characters in length. Allowed characters are [a-zA-Z0-9_]

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

collector-group (Flow Template Profiles for NAT)

IN THIS SECTION

- Syntax | 4418
- Hierarchy Level | 4419
- Description | 4419
- Options | 4419
- Required Privilege Level | 4419
- Release Information | 4419

Syntax

```
collector-group collector-group-name;
```

Hierarchy Level

```
[edit services jflow-log template-profile template-profile-name]
```

Description

Specify the name of the collector group to be associated with a template profile. The generated flow monitoring logs for NAT events in flow template format are sent to the specified collector group. By using a collector group, you can effectively and optimally transmit flow monitoring logs to a cluster of collectors in a single, one-step operation. A maximum of up to eight collectors can be aggregated into a collector group. You must have previously configured the collector group by using the `collector-group collector-group-name` statement at the `[edit services jflow-log]` hierarchy level before you associate a collector-group with a template profile.

Options

`collector-group-name`—Name of the collector group to which log messages for NAT events in flow monitoring format (IPFIX or version 9 flow template format) must be sent. The name can be up to 32 alphanumeric characters in length. Allowed characters are [a-zA-Z0-9_]

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

collector (Inline Monitoring)

IN THIS SECTION

- Syntax (Junos OS) | 4420
- Syntax (Junos OS Evolved) | 4420
- Hierarchy Level | 4421
- Description | 4421
- Options | 4421
- Required Privilege Level | 4422
- Release Information | 4422

Syntax (Junos OS)

```
collector collector-name {
    destination-address destination-IP-address;
    destination-port destination-port-number;
    dscp dscp;
    forwarding-class forwarding-class;
    routing-instance routing-instance-name;
    sampling-rate sampling-rate;
    source-address source-IP-address;
}
```

Syntax (Junos OS Evolved)

```
collector collector-name {
    destination-address destination-IP-address;
    destination-port destination-port;
    source-address source-IP-address;
}
```

Hierarchy Level

```
[edit services inline-monitoring instance instance-name]
```

Description

Configure an collector for inline monitoring. The monitored packets are exported to the collector in an IPFIX format. The actual packet is exported in an IPFIX format up to the configured clip length. By default, Junos OS supports a maximum packet length of 126 bytes, starting with the Ethernet header. The IPFIX format exports information on the original packet size, and the incoming or outgoing interface for further processing on the collector.

Options

collector-name	Name of the collector.
destination-address *destination-IP-address*	IPv4 destination IP address.
destination-port *destination-port-number*	Destination port value. • **Range:** 1 through 65535
dscp *dscp*	(Junos OS only) DSCP value. • **Default:** 0 • **Range:** 0 through 63
forwarding-class *forwarding-class*	(Junos OS only) Forwarding class for exported frames. • **Default:** best-effort
routing-instance *routing-instance-name*	(Junos OS only) Name of the routing instance. • **Default:** default.inet
sampling-rate *sampling-rate*	(Junos OS only) Rate at which the packets are sampled. In Junos OS Evolved, you do not specify the sampling rate here; you specify the sampling rate at the `[edit services inline-monitoring instance instance-name]` hierarchy level. For example, if you specify 1000, then 1 packet out of every 1000 packets is sampled.

- **Range:** 1 through 16000000
- **Default:** 1

source-address ***source-IP-address*** IPv4 source IP address.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 19.4R1.

Statement introduced in Junos Evolved OS Release 22.2R1.

RELATED DOCUMENTATION

Understanding Inline Monitoring Services

collector (LRF Profile)

IN THIS SECTION

- Syntax | 4423
- Hierarchy Level | 4423
- Description | 4423
- Options | 4423
- Required Privilege Level | 4423
- Release Information | 4423

Syntax

```
collector collector-name {
    destination {
        address collector-address;
        port collector-port-number;
    }
    source-address source-address;
}
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Description

Configure a collector that receives logging and reporting data. This collector can be specified in LRF rules.

Options

collector-name

Name for the collector.

- **Range:** Up to 32 characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

collector (LRF Rule)

IN THIS SECTION

- Syntax | 4424
- Hierarchy Level | 4424
- Description | 4424
- Options | 4425
- Required Privilege Level | 4425
- Release Information | 4425

Syntax

```
collector collector-name;
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name then report]
```

Description

Specify the collector that receives the data if the LRF rule is matched.

Options

collector-name Name of the collector that receives the data. The referenced collector must already be defined at the `[edit services lrf profile profile-name]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

color

IN THIS SECTION

- Syntax | 4426
- Hierarchy Level | 4426
- Description | 4426
- Options | 4427
- Required Privilege Level | 4427
- Release Information | 4427

Syntax

```
color {
    metric-value;
    <type metric_type>
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options (aggregate | generate | static) (defaults | route)],

[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name (aggregate | generate | static) (defaults | route)],

[edit logical-systems logical-system-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-instances routing-instance-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib routing-table-name (aggregate | generate | static) (defaults | route)]
```

Description

Fine-grained preference value for a static, aggregate, or generated route.

You can also specify a primary route preference (by including the **color** statement in the configuration), and a secondary preference that is used as a tiebreaker (by including the **color2** statement). You can also mark route preferences with additional route tiebreaker information by specifying a primary route preference and a tiebreaker route preference (by including the **preference** and **preference2** statements in the configuration).

If the Junos OS routing table contains a dynamic route to a destination that has a better (lower) preference value than the static, aggregate, or generated route, the dynamic route is chosen as the active route and is installed in the forwarding table.

Options

metric_value—The metric value for an aggregate, a generated, or a static route.

- **Range:** 0 through 4,294,967,295 (2^{32} – 1)

type *metric_type*—(Optional) External metric type for routes exported by OSPF. When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric.

- **Range:** 1 through 16

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation

aggregate (Routing) | **2070**

generate | **7672**

static (Routing Options) | **28570**

preference (Routing Options) | **13046**

color-aware

IN THIS SECTION

- Syntax | 4428
- Hierarchy Level | 4428
- Description | 4428
- Default | 4428
- Required Privilege Level | 4429
- Release Information | 4429

Syntax

```
color-aware;
```

Hierarchy Level

```
[edit firewall three-color-policer policer-name single-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

Configure the way preclassified packets are metered. In color-aware mode, the switch can assign a higher packet-loss priority, but cannot assign a lower packet loss priority (PLP). For example, suppose an upstream device assigns medium-high PLP to a packet because the packet exceeded its committed information rate (CIR). The switch cannot change the PLP to low even if the packet conforms to the configured CIR of the appropriate interface. On the other hand, if an upstream device assigns low PLP to a packet but the packet exceeds the CIR and committed burst size (CBS) of the switch interface, the switch can increase the PLP to medium-high.

Default

If you omit the `color-aware` statement, the default behavior is color-aware mode.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Overview of Policers

Understanding Color-Aware Mode for Single-Rate Tricolor Marking

Understanding Color-Aware Mode for Two-Rate Tricolor Marking

color-blind | 4429

color-blind

IN THIS SECTION

- Syntax | 4429
- Hierarchy Level | 4430
- Description | 4430
- Default | 4430
- Required Privilege Level | 4430
- Release Information | 4430

Syntax

```
color-blind;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name single-rate],
[edit dynamic-profiles profile-name firewall three-color-policer name two-rate],
[edit firewall three-color-policer policer-name single-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

For a three-color policer, configure the way preclassified packets are metered. In color-blind mode, the local router ignores the preclassification of packets and can assign a higher or lower packet loss priority.

For example, suppose an upstream router assigned medium-high packet loss priority to a packet because the packet exceeded the committed information rate on the upstream router interface.

- If the local router applies color-aware policing to the packet, the router *cannot* change the packet loss priority to low, even if the packet conforms to the configured committed information route on the local router interface.

  NOTE: A color-aware policer cannot be applied to Layer 2 traffic.

- If the local router applies color-blind policing to the packet, the router *can* change the packet loss priority to low if the packet conforms to the configured committed information route on the local router interface.

Default

If you omit the `color-blind` statement, the default behavior is color-aware mode.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

Support at the `[edit dynamic-profiles ... single-rate]` and `[edit dynamic-profiles ... two-rate]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Three-Color Policer Configuration Overview

Color Modes for Three-Color Policers

color-aware | 4428

command (Op Scripts)

IN THIS SECTION

- Syntax | 4431
- Hierarchy Level | 4431
- Description | 4431
- Options | 4432
- Required Privilege Level | 4432
- Release Information | 4432

Syntax

```
command filename-alias;
```

Hierarchy Level

```
[edit system scripts op file filename]
```

Description

Configure a filename alias for an op script file. This allows you to run the op script by referencing either the script filename or the filename alias.

Options

filename-alias	Alias for the op script file.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Enabling an Op Script and Defining a Script Alias

commands (Event Policy Change Configuration)

IN THIS SECTION

- Syntax | 4432
- Hierarchy Level | 4433
- Description | 4433
- Options | 4433
- Required Privilege Level | 4433
- Release Information | 4433

Syntax

```
commands {
    "command";
}
```

Hierarchy Level

```
[edit event-options policy policy-name then change-configuration]
```

Description

Specify the configuration mode commands to execute when the corresponding event policy is triggered.

When the event policy is triggered upon receipt of the specified event or events, the event process (eventd) invokes the configured commands to update the candidate configuration, which is then committed, provided that no commit errors occur. The eventd process executes the configuration commands in the order in which they appear in the event policy configuration.

Options

command Configuration mode command to be executed. Enclose each command in quotation marks (" "), and specify the complete statement path to the element, identifier, or value as you do in configuration mode when issuing commands at the `[edit]` hierarchy level.

The following configuration mode commands are supported:

- `activate`
- `deactivate`
- `delete`
- `set`

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configure an Event Policy to Change the Configuration

Example: Change the Configuration Using an Event Policy

Example: Change the Interface Configuration in Response to an Event

change-configuration (Event Policy) | **3284**

commands (Event Policy Execute Commands)

IN THIS SECTION

- Syntax | **4434**
- Hierarchy Level | **4434**
- Description | **4434**
- Options | **4435**
- Required Privilege Level | **4435**
- Release Information | **4435**

Syntax

```
commands {
    "command";
}
```

Hierarchy Level

```
[edit event-options policy policy-name then execute-commands]
```

Description

Specify the operational mode commands to be issued when the corresponding event policy is triggered. When the event policy is triggered, the event process (eventd) executes the commands in the order in which they appear in the event policy configuration.

> **NOTE**: If you include the `op url` command to execute a remote script as an event policy action, Python scripts are always executed using the access privileges of the generic, unprivileged user and group `nobody`. If you do not configure the `user-name` statement, SLAX and XSLT scripts are executed with root privileges.

Options

command Operational mode command to be executed. Enclose each command in quotation marks (" ").

You can include variables in commands. The eventd process replaces each variable with values contained in the events that trigger the policy. You can use command variables of the following forms:

- {$$.*attribute-name*}—The double dollar sign ($$) notation represents the event that triggers the policy. When combined with an attribute name, the variable resolves to the value of the attribute associated with the triggering event.
- {$*event*.*attribute-name*}—The dollar sign with the event name ($*event*) notation represents the most recent event that matches *event*. When combined with an attribute name, the variable resolves to the value of the attribute associated with that event.
- {$*.*attribute-name*}—The dollar sign with the asterisk ($*) notation represents the most recent event that matches any of the correlating events. The variable resolves to the value of the attribute associated with most recent event that matches any of the correlated events specified in the policy configuration.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

How to Represent the Correlating Event in an Event Policy

Configure an Event Policy to Execute Operational Mode Commands

commit

IN THIS SECTION

- Syntax | 4436
- Description | 4436
- Options | 4437
- Additional Information | 4441
- Required Privilege Level | 4441
- Release Information | 4442

Syntax

```
commit
<activate>
<and-quit>
<at "string">
<check>
<comment <comment-string>
<confirmed>
<peers-synchronize>
<prepare>
<scripts>
<synchronize | no-synchronize>
< | >
```

Description

Commit the set of changes to the database and cause the changes to take operational effect.

NOTE: The `fast-synchronize` option is not supported in the QFX Series Virtual Chassis. The `peers-synchronize` option is not supported in SRX Series Firewalls.

> **NOTE**: Beginning in Junos OS 12.3, it is possible that FPCs brought offline using the `request chassis fpc slot` *`fpc-slot`* `offline` operational-mode CLI command can come online during a configuration commit or power-supply replacement procedure. As an alternative, use the `set fpc` *`fpc-slot`* `power off` configuration-mode command at the `[edit chassis]` hierarchy level to ensure that the FPCs remain offline.

In Junos OS Evolved, if an FPC or PIC is brought offline, neither will be started when you enter a `commit` command that configures an element of the offline FPC or PIC.

Options

none

Execute the commit command without any options to commit the configuration changes to the configuration database.

activate (Optional)

Complete commit in two steps of preparing the configuration for commit and later activating the configuration. This enables you configure a number of devices and simultaneously activate the configurations on multiple devices.

and-quit (Optional)

Commit the configuration and, if the configuration contains no errors and the commit succeeds, exit from configuration mode.

at *string*

(Optional) Save software configuration changes and activate the configuration at a future time, or upon reboot. The variable *string* is `reboot` or the future time to activate the configuration changes. Enclose the *string* value (including `reboot`) in quotation marks (“ ”). You can specify time in two formats:

- A time value in the form *hh*:*mm*[:*ss*] (hours, minutes, and optionally seconds)— Commit the configuration at the specified time, which must be in the future by at least one minute but before 11:59:59 PM on the day the `commit at` configuration command is issued. Use 24-hour time for the *hh* value; for example, `04:30:00` is 4:30:00 AM, and `20:00` is 8:00 PM. The time is interpreted with respect to the clock and time zone settings on the device.

- A date and time value in the form *yyyy-mm-dd hh*:*mm*[:*ss*] (year, month, date, hours, minutes, and, optionally, seconds)—Commit the configuration at the specified day and time, which must be after the `commit at` command is issued. Use 24-hour time for the *hh* value. For example, `2003-08-21 12:30:00` is 12:30 PM on August 21, 2003. The time is interpreted with respect to the clock and time zone settings on the router.

For example, `commit at "18:00:00"`. For date and time, include both values in the same set of quotation marks. For example, `commit at "2018-03-10 14:00:00"`.

- A commit check is performed when you issue the `commit at` configuration mode command. If the result of the check is successful, then the current user is logged out of configuration mode, and the configuration data is left in a read-only state. No other commit can be performed until the scheduled commit is completed.

> **NOTE**: If Junos OS fails before the configuration changes become active, all configuration changes are lost.
>
> You cannot enter the `commit at` configuration mode command when there is a pending reboot.
>
> You cannot enter the `request system reboot` command once you schedule a commit operation for a specific time in the future.
>
> You cannot commit a configuration when a scheduled commit is pending. For information about how to use the `clear system commit` command to cancel a scheduled commit configuration, see *clear system commit*.

check

(Optional) Verify the syntax of the configuration, but do not activate it.

comment *comment-string*

(Optional) Add a comment that describes the committed configuration. The comment can be as long as 512 bytes and must be typed on a single line. You cannot include a comment with the `commit check` command. Enclose *`comment-string`* in quotation marks (" "). For example, `commit comment "Includes changes recommended by user"`.

confirmed *in minutes*

(Optional) Require that the commit be confirmed within the specified amount of time.

- To confirm a commit, enter either a `commit` or `commit check` command.
- If the commit is not confirmed within the time limit, the configuration rolls back automatically to the precommit configuration and a broadcast message is sent to all logged-in users. To show when a rollback is scheduled, enter the `show system commit` command. The allowed range is `1` through `65,535` minutes, and the default is `10` minutes.
- The timeout for the `commit confirmed` command is calculated based on the system time, when the `commit confirmed` command is issued. In case the system time is modified while a `commit confirmed` is pending, the remaining time until commit

execution might get shortened (in case the old system time is behind) or prolonged (in case the old system time is ahead) from the intended interval.

- In Junos OS Release 11.4 and later, you can also use the `commit confirmed` command in the `[edit private]` configuration mode.

no-synchronize

(Optional) Configure the `commit` command to run without synchronization. This can be useful in situations, for example, where a Routine Engine configuration is corrupted such that a commit synchronization is not possible or will block the commit.

- This option allows you to commit only on the current Routing Engine even if `set system commit synchronize` is configured.
- This option overrides the `commit peer-synchronize` configuration as well. If you have configured the commit synchronize using `set system commit synchronize` and then use the command `commit no-synchronize`, the commit will happen only on the device issuing the command.
- When using `commit synchronize`, the commit is first done in the other Routing Engine and then in the current one. If the other Routine Engine is corrupted, the commit will fail. In such cases, you can use `commit no-synchronize`. This command cannot be configured using `set`. It can only be run.

peers-synchronize

(Optional) Configure the commit command to automatically perform a peers-synchronize action between peers. The local peer (or requesting peer) on which you enable the peers-synchronize statement copies and loads its configuration to the remote (or responding) peer. Each peer then performs a syntax check on the configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both peers.

synchronize

(Optional) If your router has two Routing Engines, you can manually direct one Routing Engine to synchronize its configuration with the other by issuing the `commit synchronize` command. The Routing Engine on which you execute this command (the request Routing Engine) copies and loads its candidate configuration to the other Routing Engine (the responding Routing Engine). Both Routing Engines then perform a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.

The synchronize option has the following two additional options:

- force—(Optional) Enforce commit synchronization on the Routing Engines by using the `force` option.

The `commit synchronize` command does not work if the responding Routing Engine has uncommitted configuration changes. You can enforce commit synchronization on the Routing Engines by using the `force` option. When you issue the `commit synchronize` command with the `force` option from one Routing Engine, the configuration sessions on the other Routing Engine are terminated and the configuration is synchronized with that on the Routing Engine from which you issued the command.

- scripts—(Optional) Synchronize all commit, event, lib, op, and SNMP scripts from the requesting Routing Engine to the responding Routing Engine and commit and synchronize the configuration.

  If the `commit check` operation fails for the requesting Routing Engine, the process stops, and the scripts are not copied to the responding Routing Engine. If the `commit check` or `commit` operation fails for the responding Routing Engine, the scripts are still synchronized, since the synchronization occurs prior to the `commit check` operation on the responding Routing Engine.

- If the `load-scripts-from-flash` statement is configured at the `[edit system scripts]` hierarchy level for the requesting Routing Engine, the device synchronizes the scripts from flash memory on the requesting Routing Engine to flash memory on the responding Routing Engine. Otherwise, the device synchronizes the scripts from the hard disk on the requesting Routing Engine to the hard disk on the responding Routing Engine. The device synchronizes all scripts regardless of whether they are enabled in the configuration or have been updated since the last synchronization.

> **NOTE**: It can happen that the `commit synchronize` command is initiated at the same time from both Routing Engines, which causes the process to hang. As of Junos OS Release 15.1, this is a temporary (20 seconds) anomaly, after which the user can try the `commit sychronize` command again.

> **NOTE**: When you issue the `commit synchronize` command, you must use the `apply-groups re0` and `re1` commands. For information about how to use groups, see *Disabling Inheritance of a Configuration Group*.
>
> The responding Routing Engine must use Junos OS Release 5.0 or later.

prepare (Optional) Prepare the configuration to activate at a later stage. During the preparation stage, all the required files and databases are generated and the configuration is validated. A file is created that indicates if the commit is pending for activation. In the event of failure during the preparation stage, the log message commit preparation failed is generated.

scripts (Optional) Commit newly enabled scripts during the commit operation and push scripts to the other Routing Engine.

| (pipe) (Optional) Use the | (pipe)) options to filter the output of the `commit` command.

Additional Information

> **NOTE**: Beginning in Junos OS 12.3, it is possible that FPCs brought offline using the `request chassis fpc slot` *`fpc-slot`* `offline` operational-mode CLI command can come online during a configuration commit or power-supply replacement procedure. As an alternative, use the `set fpc` *`fpc-slot`* `power off` configuration-mode command at the `[edit chassis]` hierarchy level to ensure that the FPCs remain offline.

`| display detail`—(Optional) Monitors the commit process.

> **NOTE**: In Junos OS Release 10.4 and later, if the number of commit details or messages exceeds a page when used with the `| display detail` pipe option, the `more` pagination option on the screen is no longer available. Instead, the messages roll up on the screen by default, just like using the `commit` command with the `| no more` pipe option.

Required Privilege Level

configure—To enter configuration mode.

> **NOTE**: If you are using Junos OS in a Common Criteria environment, system log messages are created whenever a `secret` attribute is changed (for example, password changes or changes to the RADIUS shared secret). These changes are logged during the following configuration load operations:
>
> ```
> load merge
> load replace
> load override
> load update
> ```
>
> For more information, see the *Secure Configuration Guide for Common Criteria and Junos-FIPS*

Release Information

Command introduced before Junos OS Release 7.4.

Option `synchronize scripts` introduced in Junos OS Release 13.2.

Option `no-synchronize` introduced in Junos OS Release 17.2R1

RELATED DOCUMENTATION

- *Adding a Comment to Describe the Committed Configuration*
- Committing Configurations on a Routing Matrix with a TX Matrix Plus Router
- Commit Script Overview
- *Deactivating and Reactivating Statements and Identifiers in a Device Configuration*
- *Monitoring the Commit Process*
- *Scheduling a Commit Operation*
- *Verifying the Configuration*
- *Committing a Configuration*

commit activate

IN THIS SECTION

- Syntax | 4443
- Hierarchy Level | 4443
- Description | 4443
- Options | 4443
- Required Privilege Level | 4443
- Release Information | 4444

Syntax

```
commit activate{
    comment;
    and-quit;
    peers-synchronize;
    synchronize;
}
```

Hierarchy Level

```
[edit system]
```

Description

Activate a previously prepared commit. Upon successful validation, during the activation stage, previously prepared commits are activated. Also, pending activation files are checked during this stage. If there are pending activation files, the existence of required files and daemon map present in the database data structures are checked. If there is any failure, a log message is generated that informs you that the commit has failed.

Options

and-quit (Optional) Commit the configuration and, if the configuration contains no errors and the commit succeeds, exit from configuration mode.

no-synchronize (Optional) Do not synchronize the commit. Configure the `commit prepare` statement to run without synchronization.

peers-synchronize (Optional) Synchronize the commit on remote peers.

synchronize (Optional) Synchronize the commit on both Routing Engines.

Required Privilege Level

configure—To enter configuration mode.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Committing a Configuration

Committing a Configuration

commit (Commit Scripts)

IN THIS SECTION

- Syntax | 4444
- Hierarchy Level | 4445
- Description | 4445
- Options | 4445
- Required Privilege Level | 4445
- Release Information | 4446

Syntax

```
commit {
    allow-transients;
    dampen {
        dampen-options {
            cpu-factor cpu-factor;
            line-interval line-interval;
            time-interval time-interval;
        }
    }
```

```
        direct-access;
        file filename {
            allow-transients;
            cert-file path;
            checksum hash-algorithm hash-value;
            optional;
            refresh;
            refresh-from url;
            routing-instance routing-instance-name;
            source url;
        }
        max-datasize size;
        refresh;
        refresh-from url;
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
```

Hierarchy Level

```
[edit system scripts]
```

Description

For Junos OS commit scripts, configure the commit-time scripting mechanism.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Commit Script Overview

How Commit Scripts Work

Storing and Enabling Scripts

commit-delay

IN THIS SECTION

- Syntax | 4446
- Hierarchy Level | 4446
- Description | 4447
- Options | 4447
- Required Privilege Level | 4447
- Release Information | 4447

Syntax

```
commit-delay seconds;
```

Hierarchy Level

```
[edit snmp nonvolatile]
```

Description

Configure the timer for the SNMP **Set** reply and start of the commit.

Options

seconds—Delay between an affirmative SNMP **Set** reply and start of the commit.

- **Default:** 5 seconds

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

commit-interval

IN THIS SECTION

- Syntax | 4447
- Hierarchy Level | 4448
- Description | 4448
- Options | 4448
- Required Privilege Level | 4448
- Release Information | 4448

Syntax

```
commit-interval interval;
```

Hierarchy Level

```
[edit system services extensible-subscriber-services]
```

Description

Specify the interval at which Extensible Subscriber Services Manager issues requests for committing op script configurations. Requests that are received and processed within the interval are committed in a batch at the end of the interval. Requests that are processed after the interval are committed at the end of the next commit interval. When no operation script is to be committed, no request for committing operation script configurations is issued.

Options

interval Length of the interval.

- **Range:** 10 through 3600 seconds
- **Default:** 20 seconds

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Op Script Overview

commit-options

IN THIS SECTION

- Syntax | 4449
- Hierarchy Level | 4449
- Description | 4449
- Options | 4449
- Required Privilege Level | 4450
- Release Information | 4450

Syntax

```
commit-options {
    check <synchronize>;
    force;
    log "comment-string";
    synchronize;
}
```

Hierarchy Level

```
[edit event-options policy policy-name then change-configuration]
```

Description

Customize the commit options for configuration updates made through an event policy. The `check` statement and the other `commit-options` statements are mutually exclusive.

Options

check <synchronize>

Verify that the candidate configuration is syntactically correct, but do not commit the changes. On dual control plane systems, when the `check synchronize` statement is

configured, the candidate configuration on one control plane is copied to the other control plane, and the system verifies that both candidate configurations are syntactically correct. The `check` statement and the other `commit-options` statements are mutually exclusive.

force Force the commit on the other Routing Engine, ignoring any warnings. By default, the `synchronize` command does not work if the responding Routing Engine has uncommitted configuration changes. However, you can enforce commit synchronization on the Routing Engines by using the `force` option.

log "*comment-string*" Include a comment describing changes to the committed configuration. Enclose the comment in quotation marks and include it on a single line. To view commit comments, issue the `show system commit` operational mode command.

synchronize Synchronize the commit on both Routing Engines. The Routing Engine on which you execute this command copies and loads its candidate configuration to the other Routing Engine. Both Routing Engines perform a syntax check on the candidate configuration file. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configure an Event Policy to Change the Configuration

Example: Change the Configuration Using an Event Policy

Example: Change the Interface Configuration in Response to an Event

change-configuration (Event Policy) | **3284**

commands (Event Policy Change Configuration) | **4432**

retry (Event Policy) | **15704**

user-name (Event Policy) | **30569**

commit prepare

IN THIS SECTION

- Syntax | 4451
- Hierarchy Level | 4451
- Description | 4451
- Options | 4452
- Required Privilege Level | 4452
- Release Information | 4452

Syntax

```
commit prepare{
    and-quit;
    no-synchronize;
    peers-synchronzie;
synchronize;
    }
```

Hierarchy Level

```
[edit system]
```

Description

Prepare for an upcoming commit activation. Prepare the configurations that can be activated at a later stage. During the preparation stage, all the required files and databases are generated and the configuration is validated. A file is created that indicates if the commit is pending for activation. In the event of failure during the preparation stage, the log message `commit preparation failed` is generated.

Options

and-quit (Optional) Commit the configuration and, if the configuration contains no errors and the commit succeeds, exit from configuration mode.

no-synchronize (Optional) Do not synchronize the commit. Configure the `commit prepare` statement to run without synchronization.

peers-synchronize (Optional) Synchronize the commit on remote peers.

synchronize (Optional) Synchronize the commit on both Routing Engines.

Required Privilege Level

configure—To enter configuration mode.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

Committing a Configuration

Committing a Configuration

commit-synchronize-server

IN THIS SECTION

- Syntax | 4453
- Hierarchy Level | 4453
- Description | 4453

- Options | 4453
- Required Privilege Level | 4454
- Release Information | 4454

Syntax

```
commit-synchronize-server {
    traceoptions {
        file {
            filename;
            files number;
            microsecond-stamp;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag (all | debug | ephemeral-commit | operational-command);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system commit]
```

Description

For commit synchronize server actions, configure tracing operations.

Options

filename	Name of the file to receive the output of the tracing operation.

> **NOTE**: If you configure traceoptions and do not explicitly specify a filename for logging the events, the events are logged in the file `/var/log/commitd` by default.

files *number*
Maximum number of trace files.

flag *flag*
Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—All tracing operations flags.
- `debug`—Trace operations for debug events.
- `ephemeral-commit`—Trace operations for ephemeral database commit synchronize events.
- `operational-command`—Trace operations for operational command events.

microsecond-stamp
Include microsecond in the timestamp.

no-remote-trace
Disable remote tracing.

size
Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

world-readable | no-world-readable
Grant all users permission to read archived log files, or restrict the permission only to the root user and users who have the Junos OS maintenance permission.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R3 and Junos OS Evolved Release 22.1R1.

RELATED DOCUMENTATION

commit (System) | 4455

commit (System)

IN THIS SECTION

- Syntax | 4455
- Hierarchy Level | 4455
- Description | 4455
- Options | 4456
- Required Privilege Level | 4457
- Release Information | 4457

Syntax

```
commit {
    commit-synchronize-server;
    delta-export;
    fast-synchronize;
    force-commit-log;
    notification;
    peers;
    peers-synchronize;
    persist-groups-inheritance | no-persist-groups-inheritance;
    server;
    synchronize;
}
```

The parameters for fast-synchronize and synchronize do not apply for the SRX Series.

Hierarchy Level

```
[edit system]
```

Description

Configure options for Junos OS commit.

Options

`commit-synchronize-server`—(Optional) Specify traceoptions for commit synchronize server actions.

`delta-export`—(Optional) Configure system commit to export only the changes made in the candidate configuration instead of exporting the entire candidate configuration to the configuration database. This helps to reduce the time taken to commit the configuration changes.

`fast-synchronize`—(Optional) Configure commits to run in parallel (simultaneously) on both the primary and backup Routing Engines to reduce the time required for commit synchronization. The fast-synchronize configuration is valid only on systems with two Routing Engines. (Option not available for SRX Series.)

`force-commit-log`—(Optional) Enforce log or message for each commit. Once this is configured each commit triggered by an user mandates a comment. Only users with admin-control privilege can enable or disable the `force-commit-log` option. The following commands mandate a comment once you enable `force-commit-log` option:

- `commit at`
- `commit confirmed`
- `commit and-quit`
- `batch mode commits`
- `commit synchronize, no-synchronize, and peers-synchronize`
- `commit scripts`

`notification`—(Optional) Notify applications upon commit completion.

`peers`—(Optional) Specify the host names or IP addresses of the MC-LAG peers and the user authentication details for the users administering the MC-LAG peers that are participating in commit synchronization.

> **NOTE**: Starting in Junos OS Release 17.1R1, the `peers` option at the [`edit system commit`] hierarchy level is not supported in batch configuration mode.

`peers-synchronize`—(Optional) Configure a commit synchronization on MC-LAG peers.

`persist-group-inheritance`—(Optional) Configure this option to improve commit performance for systems that use many configuration groups that use wildcards. This option causes the full inheritance paths of the configuration groups to be built in the database instead of in the process memory. To disable this option, use `no-persist-groups-inheritance`. Starting in Junos OS Evolved Release 19.2R1 and Junos OS Release 19.4R1, this option is enabled by default. Junos OS Evolved requires this feature to be enabled.

The `no-persist-groups-inheritance` option to disable `persist-group-inheritance` is not supported in Junos OS Evolved.

`server`—(Optional) Configure a default batch commit.

`synchronize`—(Optional) For devices with multiple Routing Engines only. Configure the commit command to automatically perform a commit synchronize action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the commit command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engine. Each Routing Engine then performs a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines. (Option not available for SRX Series.)

The remaining statements are explained separately. See CLI Explorer.

NOTE: Junos OS Evolved does not support the following `commit` configuration options:

- `fast-synchronize`
- `peers`
- `peer-synchronize`
- `commit-synchronize-server`

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

Option `persist-groups-inheritance` added in Junos OS Release 13.2.

Option `delta-export` added in Junos OS Release 14.2.

Option `peers` added in Junos OS Release 14.2R6.

Option `peers-synchronize` added in Junos OS Release 14.2R6.

Option `no-persist-groups-inheritance` added in Junos OS Evolved Release 19.2R1 and Junos OS Release 19.4R1.

RELATED DOCUMENTATION

Improving Commit Time When Using Configuration Groups

committed-burst-size

IN THIS SECTION

- Syntax | 4458
- Hierarchy Level | 4458
- Description | 4458
- Options | 4459
- Required Privilege Level | 4459
- Release Information | 4459

Syntax

```
committed-burst-size bytes;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name single-rate],
[edit dynamic-profiles profile-name firewall three-color-policer name two-rate],
[edit firewall three-color-policer policer-name single-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

For a three-color policer, configure the committed burst size (CBS) as a number of bytes.

> **NOTE**: When you include the `committed-burst-size` statement in the configuration, you must also include the `committed-information-rate` statement at the same hierarchy level.

In three-color policing, a committed information rate (CIR) defines the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions. A flow of traffic at an average rate that conforms to the CIR is categorized green.

During periods of average traffic rates below the CIR, any unused bandwidth capacity accumulates up to a maximum amount defined by the CBS. Short periods of bursting traffic (back-to-back traffic at averages rates that exceed the CIR) are also categorized as green provided that unused bandwidth capacity is available.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policers use a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red, based on the `excess-burst-size` statement included in the policer configuration.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits. Nonconforming traffic is categorized as yellow or red based on the `peak-information-rate` and `peak-burst-rate` statements included in the policer configuration.

Options

bytes—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1500 through 100,000,000,000 bytes

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

Support at the `[edit dynamic-profiles ... single-rate]` and `[edit dynamic-profiles ... two-rate]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Three-Color Policer Configuration Overview
- Policer Bandwidth and Burst-Size Limits
- Policer Color-Marking and Actions
- Dual Token Bucket Algorithms
- Determining Proper Burst Size for Traffic Policers
- committed-information-rate | 4460
- excess-burst-size | 6580
- peak-burst-size | 12332
- peak-information-rate | 12336

committed-information-rate

IN THIS SECTION

- Syntax | 4460
- Hierarchy Level | 4460
- Description | 4461
- Options | 4461
- Required Privilege Level | 4462
- Release Information | 4462

Syntax

```
committed-information-rate bps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name single-rate],
[edit dynamic-profiles profile-name firewall three-color-policer name two-rate],
```

```
[edit firewall three-color-policer policer-name single-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

For a three-color policer, configure the committed information rate as a number of bits per second. The committed information rate (CIR) is the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions.

> **NOTE**: When you include the `committed-information-rate` statement in the configuration, you must also include the `committed-burst-size` statement at the same hierarchy level.

In three-color policing, a CIR defines the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions. A flow of traffic at an average rate that conforms to the CIR is categorized green.

During periods of average traffic rates below the CIR, any unused bandwidth capacity accumulates up to a maximum amount defined by the committed burst size (CBS). Short periods of bursting traffic (back-to-back traffic at averages rates that exceed the CIR) are also categorized as green provided that unused bandwidth capacity is available.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policers use a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red, based on the `excess-burst-size` statement included in the policer configuration.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits. Nonconforming traffic is categorized as yellow or red based on the `peak-information-rate` and `peak-burst-rate` statements included in the policer configuration.

Options

bps—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

Range:

-
  - 1500 through 100,000,000,000 bps on EX, M, and T Series routers
  - 1500 through 18,446,744,073,709,551,615 bps on Mx Series routers

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

Support at the `[edit dynamic-profiles ... single-rate]` and `[edit dynamic-profiles ... two-rate]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Three-Color Policer Configuration Overview
- Policer Bandwidth and Burst-Size Limits
- Policer Color-Marking and Actions
- Dual Token Bucket Algorithms
- Determining Proper Burst Size for Traffic Policers
- committed-burst-size | **4458**
- excess-burst-size | **6580**
- peak-burst-size | **12332**
- peak-information-rate | **12336**

community

IN THIS SECTION

- Syntax | **4463**
- Hierarchy Level | **4463**
- Description | **4463**
- Options | **4463**
- Required Privilege Level | **4463**

Release Information | 4463

Syntax

```
community community-name;
```

Hierarchy Level

```
[edit snmp rmon event index]
```

Description

The trap group that is used when generating a trap (if `eventType` is configured to send traps). If that trap group has the `rmon-alarm` trap category configured, a trap is sent to all the targets configured for that trap group. The community string in the trap matches the name of the trap group (and hence, the value of `eventCommunity`). If nothing is configured, traps are sent to each group with the `rmon-alarm` category set.

Options

community-name—Identifies the trap group that is used when generating a trap if the event is configured to send traps.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an RMON Event Entry and Its Attributes

community-name

IN THIS SECTION

- Syntax | 4464
- Hierarchy Level | 4464
- Description | 4464
- Options | 4464
- Required Privilege Level | 4465
- Release Information | 4465

Syntax

```
community-name community-name;
```

Hierarchy Level

```
[edit snmp v3 snmp-community community-index]
```

Description

The community name defines an SNMP community. The SNMP community authorizes SNMPv1 or SNMPv2 clients. The access privileges associated with the configured security name define which MIB objects are available and the operations (notify, read, or write) allowed on those objects.

Options

community-name—Community string for an SNMPv1 or SNMPv2c community. If unconfigured, it is the same as the community index. If the name includes spaces, enclose it in quotation marks (" ").

> **NOTE**: Community names must be unique. You cannot configure the same community name at the `[edit snmp community]` and `[edit snmp v3 snmp-community community-index]` hierarchy levels.
>
> The community name at the `[edit snmp v3 snmp-community community-index]` hierarchy level is encrypted and not displayed in the command-line interface (CLI).

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the SNMPv3 Community

community (Policy Options)

IN THIS SECTION

- Syntax | 4466
- Hierarchy Level | 4466
- Description | 4466
- Options | 4466
- Required Privilege Level | 4468
- Release Information | 4468

Syntax

```
community name {
    invert-match;
    members [ community-ids ];
}
```

Hierarchy Level

```
[edit dynamic policy-options],
[edit logical-systems logical-system-name policy-options],
[edit policy-options]
```

Description

Define a community, extended community or large community for use in a routing policy match condition.

Options

name—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters. To include spaces in the name, enclose it in quotation marks (“ ”).

`invert-match`—Invert the results of the community expression matching. The `community` match condition defines a regular expression and if it matches the community attribute of the received prefix, Junos OS returns a TRUE result. If not, Junos OS returns a FALSE result. The `invert-match` statement makes Junos OS behave to the contrary. If there is a match, Junos OS returns a FALSE result. If there is no match, Junos OS returns a TRUE result.

`members` *community-ids*—One or more community members. If you specify more than one member, you must enclose all members in brackets.

The format for *community-ids* is:

```
as-number:community-value
```

Starting in Junos OS Release 15.1, you can apply a wildcard member **segmented-nh:.*:0** to apply the BGP policy to all the S-PMSI A-D routes carrying extended community information.

as-number is the AS number and can be a value in the range from 0 through 65,535. *community-value* is the community identifier and can be a number in the range from 0 through 65,535.

You also can specify *community-ids* for communities as one of the following well-known community names, which are defined in RFC 1997, *BGP Communities Attribute*:

- `no-export`—Routes containing this community name are not advertised outside a BGP confederation boundary.
- `no-advertise`—Routes containing this community name are not advertised to other BGP peers.
- `no-export-subconfed`—Routes containing this community name are not advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation.

You can explicitly exclude BGP community information with a static route using the `none` option. Include `none` when configuring an individual route in the `route` portion of the `static` statement to override a `community` option specified in the `defaults` portion of the statement.

The format for extended *community-ids* is the following:

```
type:administrator:assigned-number
```

type is the type of extended community and can be either a `bandwidth`, `target`, `origin`, `domain-id`, `src-as`, `rt-import`, or a 16-bit number that identifies a specific BGP extended community. The `target` community identifies the destination to which the route is going. The `origin` community identifies where the route originated. The `domain-id` community identifies the OSPF domain from which the route originated. The `src-as` community identifies the autonomous system from which the route originated. The `rt-import` community identifies the route to install in the routing table.

> **NOTE**: For `src-as`, you can specify only an AS number and not an IP address. For `rt-import`, you can specify only an IP address and not an AS number.

administrator is the administrator. It is either an AS number or an IPv4 address prefix, depending on the type of extended community.

assigned-number identifies the local provider.

The format for linking a bandwidth with an AS number is:

`bandwidth:`*as-number*`:`*bandwidth*

- *as-number* specifies the AS number and *bandwidth* specifies the bandwidth in bytes per second.

`bandwidth-non-transitive:value`—Routes containing this extended community name are advertised to external BGP peers. All non-transitive link bandwidth community are advertised including the originated, received, and readvertised link bandwidth community.

> **NOTE**: In Junos OS Release 9.1 and later, you can specify 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*, as well as the 2-byte AS numbers that are supported in earlier releases of the Junos OS. In plain-number format, you can configure a value in the range from 1 through 4,294,967,295. To configure a `target` or `origin` extended community that includes a 4-byte AS number in the plain-number format, append the letter "L" to the end of number. For example, a target community with the 4-byte AS number 334,324 and an assigned number of 132 is represented as `target:334324L:132`.
>
> In Junos OS Release 9.2 and later, you can also use AS-dot notation when defining a 4-byte AS number for the `target` and `origin` extended communities. Specify two integers joined by a period: *16-bit high-order value in decimal.16-bit low-order value in decimal*. For example, the 4-byte AS number represented in plain-number format as 65546 is represented in AS-dot notation as 1.10.

As defined in RFC 8092, BGP large community uses 12-byte encoding and the format for BGP large *community-ids* is:

```
large: global-administrator:assigned-number:assigned-number
```

`large` indicates BGP large community.

`global-administrator` is the administrator. It is a 4-byte AS number.

`assigned-number` is a 4-byte value used to identify the local provider. BGP large community uses two 4-byte assigned number to identify the local provider.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.

Support for BGP large community introduced in Junos OS Release 17.3 for MX Series, PTX Series, and QFX Series.

RELATED DOCUMENTATION

Understanding BGP Communities, Extended Communities, and Large Communities as Routing Policy Match Conditions

Understanding How to Define BGP Communities and Extended Communities

dynamic-db | **6105**

community (Protocols Layer 2 Circuit)

IN THIS SECTION

- Syntax | **4469**
- Hierarchy Level | **4470**
- Description | **4470**
- Options | **4470**
- Required Privilege Level | **4470**
- Release Information | **4470**

Syntax

```
community community-name {
    invert-match;
    members community-members;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options],
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name backup-neighbor address],
[edit logical-systems logical-system-name protocols l2circuit neighbor address
interface interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name backup-neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor address backup-neighbor address],
[edit policy-options],
[edit protocols l2circuit local-switching interface interface-name backup-neighbor address],
[edit protocols l2circuit neighbor address interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name backup-neighbor address],
[edit routing-instances routing-instance-name protocols vpls neighbor address backup-neighbor
address]
```

Description

Specify the community for the Layer 2 circuit.

Options

community-name—Name of the Layer 2 circuit community.

`invert-match`—Invert the results of the community expression match.

`members` *community-members*—Specify the members of the community.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.
Hierarchy levels associated with the `backup-neighbor` statement (pseudowire redundancy) added in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring the Layer 2 Circuit Community

Configuring Pseudowire Redundancy on the PE Router

Example: Configuring Layer 2 Circuit Switching Protection

community (RMON)

IN THIS SECTION

- Syntax | 4471
- Hierarchy Level | 4471
- Description | 4471
- Options | 4472
- Required Privilege Level | 4472
- Release Information | 4472

Syntax

```
community community-name;
```

Hierarchy Level

```
[edit snmp rmon event index]
```

Description

Configure the SNMP trap group that is used when generating a trap (if the eventType object is configured to send traps). If that trap group has the rmon-alarm trap category configured, a trap is sent to all the targets configured for that trap group. The community string in the trap matches the name of the trap group (and hence, the value of eventCommunity). If nothing is configured, traps are sent to each group that has the rmon-alarm category configured.

The event community is not the same as an SNMP community.

Options

community-name—Name of the trap group that is used when generating a trap if the event is configured to send traps.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- *RMON MIB Event, Alarm, Log, and History Control Tables*
- *Configuring RMON Alarms and Events*
- *Monitoring RMON MIB Tables*
- *Understanding RMON*
- Junos OS Network Management Configuration Guide

community (Routing Options)

IN THIS SECTION

- Syntax | 4473
- Hierarchy Level | 4473
- Description | 4473
- Default | 4473
- Options | 4474
- Required Privilege Level | 4475

Release Information | 4475

Syntax

```
community ([ community-ids ] | no-advertise | no-export | no-export-subconfed | none | llgr-
stale | no-llgr);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options (aggregate | generate | static) (defaults | route)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name (aggregate | generate | static) (defaults | route)],
[edit logical-systems logical-system-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-instances routing-instance-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib routing-table-name (aggregate | generate | static) (defaults | route)]
```

Description

Associate BGP community information with a static, aggregate, or generated route.

NOTE: BGP large community is available only for static routes.

Default

No BGP community information is associated with static routes.

Options

community-ids—One or more community identifiers. The ***community-ids*** format varies according to the type of attribute that you use.

The BGP community attribute format is ***as-number:community-value***:

- ***as-number***—AS number of the community member. It can be a value from 1 through 65,535. The AS number can be a decimal or hexadecimal value.
- ***community-value***—Identifier of the community member. It can be a number from 0 through 65,535.

For more information about BGP community attributes, see the "Configuring the Extended Communities Attribute" section in the Routing Policies, Firewall Filters, and Traffic Policers User Guide.

For specifying the BGP community attribute only, you also can specify ***community-ids*** as one of the following well-known community names defined in RFC 1997:

- **no-advertise**—Routes containing this community name are not advertised to other BGP peers.
- **no-export**—Routes containing this community name are not advertised outside a BGP confederation boundary.
- **no-export-subconfed**—Routes containing this community are advertised to IBGP peers with the same AS number, but not to members of other confederations.
- `llgr-stale`—Adds a community to a long-lived stale route when it is readvertised.
- `no-llgr`—Marks routes which a BGP speaker does not want to be retained by LLGR. The Notification message feature does not have any associated configuration parameters.

> **NOTE**: Extended community attributes are not supported at the `[edit routing-options]` hierarchy level. You must configure extended communities at the `[edit policy-options]` hierarchy level. For information about configuring extended communities, see the Routing Policies, Firewall Filters, and Traffic Policers User Guide.

As defined in RFC 8092, BGP large community uses 12-byte encoding and the format for BGP large *`community-ids`* is:

```
large: global-administrator:assigned-number:assigned-number
```

`large` indicates BGP large community.

`global-administrator` is the administrator. It is a 4-byte AS number.

`assigned-number` is a 4-byte value used to identify the local provider. BGP large community uses two 4-byte assigned number to identify the local provider.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`llgr-stale` and `no-llgr` options added in Junos OS Release 15.1.

Support for BGP large community introduced in Junos OS Release 17.3 for MX Series, PTX Series, and QFX Series.

RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation

aggregate (Routing) | **2070**

generate | **7672**

static (Routing Options) | **28570**

community (SNMP)

IN THIS SECTION

- Syntax | **4476**
- Hierarchy Level | **4476**
- Description | **4476**
- Default | **4476**
- Options | **4476**
- Required Privilege Level | **4477**
- Release Information | **4477**

Syntax

```
community community-name {
    authorization authorization;
    client-list-name client-list-name;
    clients {
        address restrict;
    }
    view view-name;
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Define an SNMP community. An SNMP community authorizes SNMP clients based on the source IP address of incoming SNMP request packets. A community also defines which MIB objects are available and the operations (read-only or read-write) allowed on those objects.

> **NOTE**: The `authorization read-write` option is not supported on the QFX3000 QFabric system.

The SNMP client application specifies an SNMP community name in `Get`, `GetBulk`, `GetNext`, and `Set` SNMP requests.

Default

If you omit the `community` statement, all SNMP requests are denied.

Options

community-name—Community string. If the name includes spaces, enclose it in quotation marks (" ").

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring SNMP Communities

Configuring the SNMP Community String

community-vlan

IN THIS SECTION

- Syntax | 4477
- Hierarchy Level | 4477
- Description | 4478
- Required Privilege Level | 4478
- Release Information | 4478

Syntax

```
community-vlan vlan community-vlan-name;
```

Hierarchy Level

```
[edit vlans primary-vlan-name vlan-id primary-vlan-vlan-id]
```

Description

Configure the specified community VLAN to be a secondary VLAN of the specified primary VLAN. A *community* VLAN is used to transport frames among members of a community (a subset of users within the VLAN), and to forward frames upstream to the primary VLAN.

> **NOTE**: Before you specify this configuration statement, you must have already configured the specified community VLAN and assigned a VLAN ID to it. See "private-vlan" on page 13220.

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Creating a Private VLAN on a Single Switch with ELS Support (CLI Procedure)

Creating a Private VLAN Spanning Multiple EX Series Switches with ELS Support (CLI Procedure)

community-vlans (MX Series)

IN THIS SECTION

- Syntax | 4479
- Hierarchy Level | 4479
- Description | 4479
- Options | 4479
- Required Privilege Level | 4479

Release Information | 4480

Syntax

```
community-vlans [ number number-number ];
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name,
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]
```

Description

Configure the specified community VLAN to be a secondary VLAN of the specified primary VLAN. A *community* VLAN is used to transport frames among members of a community, which is a subset of users within the VLAN, and to forward frames upstream to the primary VLAN.

NOTE: When you specify this configuration statement, the VLAN ID of a logical interface that you associate with a bridge domain that matches with the VLAN ID or list of IDs that you specify using the `community-vlans` state is treated as a community port.

Options

number Individual VLAN IDs separated by a space.

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

compatibility (Application Identification)

IN THIS SECTION

- Syntax | 4480
- Hierarchy Level | 4480
- Description | 4480
- Options | 4480
- Required Privilege Level | 4481
- Release Information | 4481

Syntax

```
compatibility junos-compatibility-version;
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Description

Specify the Junos OS release for compatibility.

Options

junos-compatibility-version Name of the Junos OS software release compatibility version, such as 17.1.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

compatibility-version

IN THIS SECTION

- Syntax | 4481
- Hierarchy Level | 4482
- Description | 4482
- Options | 4482
- Required Privilege Level | 4482
- Release Information | 4482

Syntax

```
compatibility-version;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Specify the compatible version mode to be used. When compatibility-version is set to value 1, the node operates in ITU-T Recommendation G.8032/Y.1344 version 1 compatible mode. In this mode all the supported external commands are blocked, ring-id is forced to be 1 and mode of operation is set to revertive mode.

Options

- 1—Use ITU-T Recommendation G.8032/Y.1344 compatibile mode version 1.
- 2—Use ITU-T Recommendation G.8032/Y.1344 compatibile mode version 2.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

compliance (Juniper Secure Connect)

IN THIS SECTION

- Syntax | 4483
- Hierarchy Level | 4484
- Description | 4484
- Options | 4484
- Required Privilege Level | 4487
- Release Information | 4487

Syntax

```
compliance {
    pre-logon name {
        term term-name {
             match {
                 platform {
                    (android | ios | macos | windows) {
                        (app-version | os-version) {
                            (equal | greater-than | greater-than-or-equal | less-than | less-
than-or-equal) version;
                        }
                    }
                }
                  hostname value;
                  ms-domain value;
                  ms-workgroup value;
                  deviceid value;
                }
             action (accept | reject);
      }
   }
}
```

Hierarchy Level

```
[edit security remote-access]
```

Description

The statement defines Juniper Secure Connect remote-access prelogon compliance policies. You associate a single compliance rule object per remote-access connection profile. This means, a remote-access connection profile can have one associated compliance policy. The Juniper Secure Connect application sends details to the SRX Series Firewall. The device performs prelogon compliance checks. Based on the prelogon compliance rule match, action is taken to accept or reject a connection.

You can create multiple prelogon compliance policies and each policy can contain multiple term rules. The term rules are a set of individual rules containing match conditions and their actions based on the compliance parameters listed in the options below. You can associate a single compliance rule object per remote-access connection profile.

Evaluation Criteria

For every connection request, SRX Series Firewall processes each rule as follows –

1. SRX Series Firewall evaluates the term rules in the order they appear in the configuration.
2. If there is no match in the current term rule, it evaluates the next term rule.
3. Based on the match, it takes an action.
4. When there is no action specified, the default action for a match rule is `reject`.
5. When no further term rule is specified for an unmatched rule, the default action is `reject`
6. When no compliance rule is attached to the profile, the default action is `accept`.

Based on this evaluation criteria, the administrator defines rules.

Options

pre-logon Define pre-login compliance rule; specify the compliance rule *name*.

- Values:

- Format: String beginning with a number or letter and consisting of letters, numbers, dashes and underscores.
- Range: Supports upto 255 prelogon compliance rules.

term Define compliance rule term; specify the term rule *name*.

- Values:
  - Format: Must be a string beginning with a number or letter and consisting of letters, numbers, dashes and underscores.
  - Range—Supports upto 10 term rules per compliance rule.

action Specify the action based on the rule match.

- Values:
  - `accept`—To approve the request.
  - `reject`—To reject the request. This is the default action for an unmatched rule when no further term rule is specified.
  - Default: When no action is specified, the default action is `reject`.

match Specify rules to match.

platform Specify rule to match OS and Client information for the specified OS.

- Values:
  - `android`—Specify android version and app-version.
  - `ios`—Specify iOS version and app-version.
  - `macos`—Specify macOS version and app-version.
  - `windows`—Specify Windows OS version and app-version.

app-version Match remote access client version with the specified operational values.

- Values:
  - `equal`—Perform operation 'equal'.

- `greater-than`—Perform operation 'greater-than'.
- `greater-than-or-equal`—Perform operation 'greater-than-or-equal'.
- `less-than`—Perform operation 'less-than'.
- `less-than-or-equal`—Perform operation 'less-than-or-equal.

os-version Match operating system version with the specified operational values.

- Values:
  - `equal`—Perform operation 'equal'.
  - `greater-than`—Perform operation 'greater-than'.
  - `greater-than-or-equal`—Perform operation 'greater-than-or-equal'.
  - `less-than`—Perform operation 'less-than'.
  - `less-than-or-equal`—Perform operation 'less-than-or-equal.

version Specify version.

- Values:
  - `app-version` supports numeric and . (period) characters.
  - `os-version` supports alphanumeric and . (period) characters.
- Range:
  - `app-version` supports upto 16 entries per term rule.
  - `os-version` supports upto 16 entries per term rule.

deviceid Specify set of device IDs.

- Values:
  - Supports alphanumeric, +, /, and = characters.

    This is a list of values.
- Range:
  - Supports upto 1024 entries per term.

hostnames

Specify set of host names.

- Values:
  - Supports alphanumeric, - and _ characters.

    This is a list of values.
- Range:
  - Supports upto 1024 entries per term rule.

ms-domain

Specify set of domain names.

- Values:
  - Supports alphanumeric, - and _ characters.

    This is a list of values.
- Range:
  - Supports upto 16 entries per term rule.

ms-workgroup

Specify set of work groups.

- Values:
  - Supports alphanumeric, - and _ characters.

    This is a list of values.
- Range:
  - Supports upto 16 entries per term rule.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Juniper Secure Connect Administration Guide

compress (Accounting Options)

IN THIS SECTION

- Syntax | 4488
- Hierarchy Level | 4488
- Description | 4488
- Required Privilege Level | 4488
- Release Information | 4489

Syntax

```
compress;
```

Hierarchy Level

```
[edit accounting-options file filename]
```

Description

Compress the accounting file during file transfer to the backup site.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- Configuring Accounting-Data Log Files
- *Configuring Flat-File Accounting for Layer 2 Wholesale*
- *Configuring Flat-File Accounting for Extensible Subscriber Services Management*
- *Flat-File Accounting Overview*

compress-configuration-files (System)

IN THIS SECTION

- Syntax | 4489
- Hierarchy Level | 4489
- Description | 4490
- Default | 4490
- Options | 4490
- Required Privilege Level | 4490
- Release Information | 4490

Syntax

```
(compress-configuration-files | no-compress-configuration-files);
```

Hierarchy Level

```
[edit system]
```

Description

Compress the current operational configuration file. By default, the current operational configuration file is compressed, and is stored in the file **juniper.conf**, in the **/config** file system, along with the last three committed versions of the configuration. However, with large networks, the current configuration file might exceed the available space in the **/config** file system. Compressing the current configuration file allows the file to fit in the file system, typically reducing the size of the file by 90 percent. The current configuration file is compressed on the second commit of the configuration after the first commit is made to include the `compress-configuration-files` statement.

> **NOTE**: We recommend that you enable compression of the router configuration files to minimize the amount of disk space that they require.

Default

The current operational configuration file is compressed.

Options

This statement has no options.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Compressing the Current Configuration File

compression

IN THIS SECTION

- Syntax | 4491
- Hierarchy Level | 4491
- Description | 4491
- Required Privilege Level | 4492
- Release Information | 4492

Syntax

```
compression {
    rtp {
        f-max-period number;
        maximum-contexts number <force>;
        port {
            minimum port-number;
            maximum port-number;
        }
        queues [ queue-numbers ];
    }
}
```

Hierarchy Level

```
[edit interfaces interface-nameunit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure the compression properties for voice services traffic.

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Compression of Voice Traffic

compression-device (Interfaces)

IN THIS SECTION

- Syntax | 4492
- Hierarchy Level | 4492
- Description | 4493
- Options | 4493
- Required Privilege Level | 4493
- Release Information | 4493

Syntax

```
compression-device name;
```

Hierarchy Level

```
[edit interfaces interface-name unit (Interfaces) logical-unit-number]
```

Description

Specify the compression interface for voice services traffic.

Options

name—Name of the AC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Interfaces

compression-device (Interfaces)

IN THIS SECTION

- Syntax | 4494
- Hierarchy Level | 4494
- Description | 4494
- Required Privilege Level | 4494
- Release Information | 4494

Syntax

```
compression-device interface-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Specify the compression interface for voice services traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configuring the Compression Interface with PPP Encapsulation

compression (PPP Properties)

IN THIS SECTION

- Syntax | 4495
- Hierarchy Level | 4495

- Description | 4495
- Required Privilege Level | 4495
- Release Information | 4495

Syntax

```
compression {
    acfc;
    pfc;
}
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options],
[edit interfaces interface-name unit logical-unit-number ppp-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options]
```

Description

For interfaces with PPP encapsulation, set Link Control Protocol (LCP) compression options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

compute-options

IN THIS SECTION

- Syntax | 4496
- Hierarchy Level | 4496
- Description | 4496
- Required Privilege Level | 4496
- Release Information | 4496

Syntax

```
compute-options;
```

Hierarchy Level

```
[edit protocols source-packet-routing]
```

Description

Configure compute options applicable to all the computed paths.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.2R1-S1.

RELATED DOCUMENTATION

Enabling Distributed CSPF for Segment Routing LSPs

compute-profile | 4497

compute-profile

IN THIS SECTION

- Syntax | 4497
- Hierarchy Level | 4498
- Description | 4498
- Options | 4498
- Required Privilege Level | 4500
- Release Information | 4500

Syntax

```
compute-profile name {
    algorithm
    protected {
        mandatory;
    }
    unprotected {
        mandatory;
    }
    bound-metric {
        (delay | igp | te);
    }
    admin-group include-any [ include-any ... ]include-all [ include-all ... ]exclude
[ exclude ... ];
    compute-segment-list compute-segment-list;
    maximum-computed-segment-lists maximum-computed-segment-lists;
    maximum-segment-list-depth maximum-segment-list-depth;
    metric-type {
```

```
        (igp | te | delay);
    }
    no-label-stack-compression;
}
```

Hierarchy Level

```
[edit protocols source-packet-routing]
```

Description

Configure the compute profile for dynamically computed paths. You can use a compute profile to logically group the computation constraints. These compute profiles are referenced by the segment routing paths for computing the primary and secondary segment routing LSPs.

Options

name Name of the computation-profile.

algorithm Algorithm that you would prefer to participate (1 to 255).

- Values:
  - 1—Strict SPF Algorithm
  - 128 to 255—Flex Algorithm
  - 2 to 127—Invalid range and proper commit check error will be given to user

> **NOTE**: Only one algorithm can be configured in the compute profile. The configured algorithm is used as a constraint and SIDs corresponding to the given algorithm will be used for compression.
>
> By default, the value is 0 and hence would result in just using algorithm 0 SIDs for compression.

protected Choose protected labels if available.

- Values:
  - mandatory—Mandatorily choose protected labels.

unprotected

Choose unprotected labels if available.

- Values:
  - mandatory—Mandatorily choose unprotected labels.

bound-metric

Bound metric used for computation.

bound-metric

Bound metric used for computation.

delay

Delay bound metric value.

- **Range:** 1 through 4294967295

igp

IGP bound metric value

- **Range:** 1 through 4294967295

te

TE bound metric value

- **Range:** 1 through 4294967295

compute-segment-list

Name of the compute type segment list.

maximum-computed-segment-lists

Maximum number of segment-lists (ECMP paths) to be computed.

- **Range:** 1 through 128

maximum-segment-list-depth

Maximum depth of computed path.

- **Range:** 1 through 16

metric-type

Specify the metric type used for computation.

- Values:
  - igp—Interior gateway protocol metric.
  - te—Traffic-engineering metric.
  - delay—Compute delay metrics.

    For metric-type delay, you can assign the following metric types:

- minimum—Minimum delay metric value from TED for cumulative lowest latency path calculation. The minimum delay metric is enabled by default.
- average—Average delay metric value from TED for cumulative lowest latency path calculation.
- maximum—Maximum delay metric value from TED for cumulative lowest latency path calculation.
- delay-variation-threshold <value>—Link delay variation threshold (1..16777215). Any link exceeding the delay variation threshold would be excluded from path calculation. The delay variation threshold is independent of whether you are doing optimization for minimum, maximum, or average.

no-label-stack-compression Provide fully expanded path, using adjacency segment identifiers.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.2R1-S1.

RELATED DOCUMENTATION

Enabling Distributed CSPF for Segment Routing LSPs

compute-options | 4496

concurrent-data-sessions

IN THIS SECTION

Syntax | 4501

- Hierarchy Level | 4501
- Description | 4501
- Options | 4501
- Required Privilege Level | 4501
- Release Information | 4501

Syntax

```
concurrent-data-sessions max-session-number;
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name enable service-name]
```

Description

Specify the maximum number of sessions that are concurrently enabled for the named service. The system randomly selects the number of sessions and enables the named service, whereas other sessions are not allotted the named service. This facilitates to increase the limit on the number of resources a service can use.

Options

max-session-number—Maximum number of sessions concurrently enabled for the named service.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

condition

IN THIS SECTION

- Syntax | 4502
- Hierarchy Level | 4503
- Description | 4503
- Options | 4503
- Required Privilege Level | 4503
- Release Information | 4503

Syntax

```
condition name {
    if-route-exists {
        address;
        address-family {
            ccc {
                address;
                peer-unit peer-unit;
                standby;
                table table;
                }
            inet {
                address;
                table table;
                 }
        }
        table table;
        }
    route-active-on (node0 | node1);
    dynamic-db;
 }
```

Hierarchy Level

```
[edit dynamic policy-options],
[edit logical-systems logical-system-name policy-options],
[edit policy-options]
```

Description

Define a policy condition based on the existence of routes in specific tables for use in BGP export policies.

Options

name	Name to identify Condition
route-active-on	Route is active on a specific node • node0—Route active on node 0 • node1—Route active on node 1

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.

Support for the address families introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Conditional Advertisement and Import Policy (Routing Table) with certain match conditions

Example: Configuring Pseudowire Redundancy in a Mobile Backhaul Scenario

dynamic-db | 6105

conditional-match

IN THIS SECTION

- Syntax | 4504
- Hierarchy Level | 4505
- Description | 4505
- Options | 4505
- Required Privilege Level | 4505
- Release Information | 4505

Syntax

```
conditional-match name {
    offset1 {
        base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader |
start-of-L4-Header);
        matchdata1 matchdata1;
        matchdata1-mask matchdata1-mask;
        offset1-value offset1-value;
    }
    offset2 {
        base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader |
start-of-L4-Header);
        matchdata2 matchdata2;
        matchdata2-mask matchdata2-mask;
        offset2-value offset2-value;
    }
    offset3 {
```

```
        base-offset3 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader |
start-of-L4-Header);
        matchdata3 matchdata3;
        matchdata3-mask matchdata3-mask;
        offset3-value offset3-value;
    }
    offset4 {
        base-offset4 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader |
start-of-L4-Header);
        matchdata4 matchdata4;
        matchdata4-mask matchdata4-mask;
        offset4-value offset4-value;
    }
}
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key]
```

Description

Configure selected bytes to be directly used for user-defined flex hashing or to be used only when the data pattern in these bytes matches with specific values (conditional match).

Options

name Name identifier

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

flex-hashing | 7227

no-inner-payload | 11474

conditional-metric

IN THIS SECTION

- Syntax | 4506
- Hierarchy Level | 4506
- Description | 4507
- Options | 4507
- Required Privilege Level | 4507
- Release Information | 4507

Syntax

```
conditional-metric {
 igp-metric-threshold threshold conditional-metric-value;
}
```

Hierarchy Level

```
[edit logical-systems name protocols mpls label-switched-path name],
[edit logical-systems name routing-instances name protocols mpls label-switched-path name],
[edit logical-systems name tenants name routing-instances name protocols mpls label-switched-
path name],
[edit protocols mpls label-switched-path name],
[edit routing-instances name protocols mpls label-switched-path name],
```

```
[edit tenants name routing-instances name protocols mpls label-switched-path name]
```

Description

Set a conditional metric value for the LSP.

Options

igp-metric-threshold

Set conditional LSP metric based on the IGP metric.

- **Range:** Maximum 4

threshold

Set conditional IGP metric threshold.

- **Range:** 1 through 16777215

conditional-metric-value

Set conditional LSP metric value.

- **Range:** 1 through 16777215

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1R1.

RELATED DOCUMENTATION

Configuring LSP Metrics

confederation

IN THIS SECTION

- Syntax | 4508
- Hierarchy Level | 4508
- Description | 4508
- Options | 4509
- Required Privilege Level | 4509
- Release Information | 4509

Syntax

```
confederation confederation-autonomous-system members [ autonomous-systems ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit routing-options]
```

Description

Specify the routing device's confederation AS number.

If you administer multiple ASs that contain a very large number of BGP systems, you can group them into one or more *confederations*. Each confederation is identified by its own AS number, which is called a *confederation AS number*. To external ASs, a confederation appears to be a single AS. Thus, the internal topology of the ASs making up the confederation is hidden.

The BGP path attributes **NEXT_HOP**, **LOCAL_PREF**, and **MULTI_EXIT_DISC**, which normally are restricted to a single AS, are allowed to be propagated throughout the ASs that are members of the same confederation.

Because each confederation is treated as if it were a single AS, you can apply the same routing policy to all the ASs that make up the confederation.

Grouping ASs into confederations reduces the number of BGP connections required to interconnect ASs.

If you are using BGP, you can enable the local routing device to participate as a member of an AS confederation. To do this, include the `confederation` statement.

Specify the AS confederation identifier, along with the peer AS numbers that are members of the confederation.

Note that peer adjacencies do not form if two BGP neighbors disagree about whether an adjacency falls within a particular confederation.

Options

autonomous-systems—AS numbers of the confederation members.

- **Range:** 1 through 65,535

confederation-autonomous-system—Confederation AS number. Use one of the numbers assigned to you by the NIC.

- **Range:** 1 through 65,535

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Example: Configuring BGP Confederations*
>
> *Understanding BGP Confederations*

configuration-database

IN THIS SECTION

- Syntax | 4510
- Hierarchy Level | 4511
- Description | 4511
- Options | 4511
- Required Privilege Level | 4512
- Release Information | 4513

Syntax

```
configuration-database {
    ephemeral {
        allow-commit-synchronize-with-gres;
        commit-synchronize-model (asynchronous | synchronous);
        cyclic-version-for-ephemeral-default version;
        delete-ephemeral-default;
        ignore-ephemeral-default;
        instance instance-name {
            cyclic-version version;
            resize {
                database-size-diff size MB;
                database-size-on-disk size MB;
            }
        }
        resize-ephemeral-default {
           database-size-diff size MB;
           database-size-on-disk size MB;
        }
    }
    extend-size;
    max-db-size max-db-size;
    resize {
        database-size-diff size MB;
        database-size-on-disk size MB;
```

```
    }
    virtual-memory-mapping {
        process process-name {
            fixed-size size KB;
            page-pooling-size size KB;
        }
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Define parameters for the configuration databases.

Configure the `virtual-memory-mapping` statement hierarchy to increase the memory space available for the configuration database size per process, as needed, by using virtual memory mapping and page-pooling.

Configure the `ephemeral` statement hierarchy to specify options for the ephemeral configuration database.

Configure the `resize` statement hierarchy when it is necessary to resize the configuration database.

Options

ephemeral Configure settings for the ephemeral configuration database. `commit-synchronize-model` specifies the commit model used to synchronize ephemeral data to the backup Routing Engine during a commit synchronize operation; `instance` sets the instance name; `allow-commit-synchronize-with-gres` enables a device to synchronize ephemeral configuration data to the other Routing Engine when GRES is enabled on the device and a commit synchronize operation is requested; and `ignore-ephemeral-default` disables the default instance of the ephemeral configuration database.

extend-size Increase the memory space available for the configuration database, per process, up to 1.5 GB. The extended default size is based on the original default database size. A device with a default database size of 409.99 MB will extend to a maximum of 1049.99 MB, and a device with a default database size of 665.99 MB will extend to a maximum of 1305.99.

Any operation on the `extend-size` configuration statement such as, `deactivate`, `delete`, or `set` requires that you reboot the device to make the change effective. The device generates the following warning message:

```
Change in 'system configuration-database extend-size' will be effective at next reboot only.
```

The `extend-size` and `max-db-size` statements are mutually exclusive.

> **NOTE**: The `extend-size` configuration statement is supported for Junos OS Evolved starting in release 22.1.
>
> You should remove the `extend-size` statement before performing a rollback from Junos OS Evolved release 22.1 or later to any release prior to 22.1.

max-db-size Configures the maximum database size. The `max-db-size` and `extend-size` statements are mutually exclusive.

resize Resizes the configuration database. The database will be resized automatically during the commit based on the values specified in the `database-size-diff` and `database-size-on-disk` attributes. The system displays `Database resize completed` when the configuration database is successfully resized.

virtual-memory-mapping (Junos OS only) Define parameters for using virtual memory mapping for the configuration database on a per-process basis. The `process` statement sets the process name. For each process, you can define the following options:

- `fixed-size`—Size in kilobytes to directly map for the initial portion of the configuration database for the specified process.
  - Range: 512 KB through 696,320 KB
- `page-pooling-size`—Size in kilobytes to use for page-pooling the remaining data in the database.
  - Range: 512 KB through 696,320 KB

> **NOTE**: Junos OS Evolved does not support the virtual-memory-mapping option.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Using Virtual Memory for Process Configuration Data

configuration-database (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 4513
- Hierarchy Level | 4513
- Description | 4514
- Required Privilege Level | 4514
- Release Information | 4514

Syntax

```
configuration-database {
    max-db-size size;
}
```

Hierarchy Level

```
[edit system]
```

Description

Enhanced subscriber management leverages system shared memory to improve performance and scaling. Since this memory is used by both the Junos configuration and enhanced subscriber management, you need to set an upper limit on the memory available to the Junos configuration, which in turn determines the allocation available for enhanced subscriber management. Starting in Junos OS Release 20.1R1, the upper limit applies to the Junos configuration database together with the schema database.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

Junos OS Enhanced Subscriber Management Overview

configuration-file

IN THIS SECTION

- Syntax | 4515
- Hierarchy Level | 4515
- Description | 4515
- Options | 4515
- Required Privilege Level | 4515
- Release Information | 4515

Syntax

```
server-name configuration-file filepath;
```

Hierarchy Level

```
[edit access securid-server]
```

Description

Specify the path of the SecurID server configuration file. The file is copied on the devices in some directory location—for example, `/var/db/securid/sdconf.rec`.

Options

- *`server-name`*—Name of the SecurID authentication server.
- *`filepath`*—Path of the SecurID server configuration file.

Required Privilege Level

secret—To view this statement in the configuration.

secret-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 9.1 of Junos OS.

RELATED DOCUMENTATION

> *Example: Configuring SecurID User Authentication*

configuration-name

IN THIS SECTION

- Syntax | 4516
- Hierarchy Level | 4516
- Description | 4516
- Required Privilege Level | 4516
- Release Information | 4517

Syntax

```
configuration-name configuration-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mstp],
[edit logical-systems logical-system-name routing-instances routing-instance-name
protocols mstp],
```

```
[edit protocols mstp],
```

```
[edit routing-instances routing-instance-name protocols mstp]
```

Description

Specify the configuration name, which is the MSTP region name carried in the MSTP BPDUs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding MSTP

Understanding BPDUs Used for Exchanging Information Among Bridges

Configuring Multiple Spanning Tree Protocol

Configuring MSTP on Switches

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP

Understanding MSTP

configuration-servers

IN THIS SECTION

- Syntax | **4518**
- Hierarchy Level | **4518**
- Description | **4518**
- Options | **4518**
- Required Privilege Level | **4518**
- Release Information | **4519**

Syntax

```
configuration-servers {
    url {
        password password;
    }
}
```

Hierarchy Level

```
[edit system autoinstallation]
```

Description

Configure the URL address of a server from which the configuration files must be obtained.

You can download a configuration file automatically from an FTP, Hypertext Transfer Protocol (HTTP), or Trivial FTP (TFTP) servers. Examples of URLs:

- tftp://hostname/path/filename
- ftp://username:password@ftp.hostname.net
- http://hostname/path/filename
- http://usenamer:password@httpconfig.sp.com

Options

- `url`—Specify the URL address of the server containing the configuration files.
- *password*—Specify the password for authentication with the configuration server. Specifying a password in URLs and in the *password* option might result in commit failure. We recommend you to use the *password* option for specifying the password.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Autoinstallation of Configuration Files

configuration-synchronize (Chassis Cluster)

IN THIS SECTION

- Syntax | 4519
- Hierarchy Level | 4519
- Description | 4520
- Options | 4520
- Required Privilege Level | 4520
- Release Information | 4520

Syntax

```
configuration-synchronize {
    no-secondary-bootup-auto;
}
```

Hierarchy Level

```
{primary:node0}
[edit chassis cluster]
```

Description

Disables the automatic chassis cluster synchronization between the primary and secondary nodes.

The chassis cluster synchronization feature automatically synchronizes the configuration from the primary node to the secondary node when the secondary joins the primary as a cluster. If you want to disable automatic chassis cluster synchronization between the primary and secondary nodes, you can do so by entering the `set chassis cluster configuration-synchronize no-secondary-bootup-auto` command in configuration mode.

Options

no-secondary-bootup-auto Disable the automatic chassis cluster synchronization between the primary and secondary nodes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Understanding Automatic Chassis Cluster Synchronization Between Primary and Secondary Nodes

request chassis cluster configuration-synchronize | **14452**

show chassis cluster information configuration-synchronization | **17378**

configure

IN THIS SECTION

Syntax | **4521**

- Syntax (Junos OS Evolved) | 4521
- Description | 4521
- Options | 4521
- Additional Information | 4522
- Required Privilege Level | 4522
- Output Fields | 4522
- Sample Output | 4522
- Release Information | 4523

Syntax

```
configure
<batch>
<dynamic>
<exclusive>
<private>
```

Syntax (Junos OS Evolved)

```
configure
<batch>
<exclusive>
<private>
```

Description

Enter configuration mode. When this command is entered without any optional keywords, everyone can make configuration changes and commit all changes made to the configuration.

Options

none Enter configuration mode.

batch (Optional) Work in the batch commit mode where commit operations are executed in batches.

dynamic (Optional) (Not available for Junos OS Evolved) Configure routing policies and certain routing policy objects in a dynamic database that is not subject to the same verification required in the standard configuration database. As a result, the time it takes to commit changes to the dynamic database is much shorter than for the standard configuration database. You can then reference these policies and policy objects in routing policies you configure in the standard database.

exclusive (Optional) Lock the candidate configuration for as long as you remain in configuration mode, allowing you to make changes without interference from other users. Other users can enter and exit configuration mode, but they cannot change the configuration.

private (Optional) Allow multiple users to edit different parts of the configuration at the same time and to commit only their own changes, or to roll back without interfering with one another's changes. You cannot commit changes in configure private mode when another user is in configure exclusive mode. This mode does not support configuring statements corresponding to third-party YANG data models, for example, OpenConfig or custom YANG data models.

Additional Information

For more information about the different methods of entering configuration mode and the restrictions that apply, see the Junos OS Administration Library for Routing Devices.

Required Privilege Level

configure

Output Fields

When you enter this command, you are placed in configuration mode and the system prompt changes from *hostname>* to *hostname#*.

Sample Output

configure

```
user@host> configure
Entering configuration mode
```

```
[edit]
user@host#
```

Release Information

Command introduced before Junos OS Release 7.4.

The `dynamic` option of the `configure` command is deprecated for Junos OS Evolved.

RELATED DOCUMENTATION

show configuration | 19167

configured-flow-control

IN THIS SECTION

- Syntax | 4523
- Hierarchy Level | 4524
- Description | 4524
- Default | 4524
- Options | 4524
- Required Privilege Level | 4524
- Release Information | 4525

Syntax

```
configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
}
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
```

Description

Configure Ethernet PAUSE asymmetric flow control on an interface. You can set an interface to generate and send PAUSE messages, and you can set an interface to respond to PAUSE messages sent by the connected peer. You must set both the `rx-buffers` and the `tx-buffers` values when you configure asymmetric flow control.

Use the `flow-control` and `no-flow-control` statements to enable and disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.

> **NOTE**: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC) by applying a congestion notification profile to the interface.
>
> Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

Default

Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- congestion-notification-profile | 4525
- flow-control | 7267
- *Configuring CoS Asymmetric Ethernet PAUSE Flow Control*
- *Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control*
- *Understanding CoS Flow Control (Ethernet PAUSE and PFC)*

congestion-notification-profile

IN THIS SECTION

- Syntax | 4525
- EX4550 Switches | 4526
- Hierarchy Level | 4526
- Description | 4527
- Options | 4527
- Required Privilege Level | 4528
- Release Information | 4528

Syntax

```
congestion-notification-profile profile-name {
    input {
        (dscp | ieee-802.1) {
            code-point [code-point-bits] {
                pfc {
                    mru mru-value;
                }
```

```
                xon value;
            }
        }
        cable-length cable-length-value;
    }
    output {
        ieee-802.1 {
            code-point [code-point-bits] {
                flow-control-queue [queue | list-of-queues];
            }
        }
    }
    pfc-watchdog {
        detection number of polling intervals;
        poll-interval time;
        recovery time;
        watchdog-action {
            drop;
        }
    }
}
```

EX4550 Switches

```
congestion-notification-profile profile-name {
    input {
        ieee-802.1 {
            code-point up-bits pfc;
```

Hierarchy Level

```
[edit class-of-service],
[edit class-of-service interfaces interface-name]
```

Description

Configure a congestion notification profile (CNP) to enable priority-based flow control (PFC) on traffic and apply the profile to an interface. You can apply a CNP to most interfaces, including aggregated ethernet (AE) interfaces and their individual members.

A congestion notification profile can be configured to enable PFC on incoming traffic (`input` stanza) that matches the following:

- A Differentiated Services code point (DSCP) value in the Layer 3 IP header (for traffic that is not VLAN-tagged).
- An IEEE 802.1 code point at Layer 2 in the VLAN header (for VLAN-tagged traffic).

A congestion notification profile can be configured to enable PFC on outgoing traffic (`output` stanza) specified only by an IEEE 802.1 code point at Layer 2 in the VLAN header.

> **NOTE**: You must configure PFC for FCoE traffic. Each interface that carries FCoE traffic should be configured for PFC on the FCoE code point (usually `011`).

There is no limit to the total number of congestion notification profiles you can create. However:

- You can attach a maximum of one congestion notification profile to an interface.
- DSCP-based PFC and IEEE 802.1p PFC cannot be configured under the same congestion notification profile.

> **NOTE**: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

Options

profile-name Name of the congestion notification profile.

pfc-watchdog Enable the Priority Flow Control (PFC) watchdog. If you do not configure any options, the default values are used.

- detection *number of polling intervals*—How many polling intervals the PFC watchdog waits before it determines that a PFC queue has stalled.

    - Default: 2
    - Range: 2-10
  - poll-interval *time*—How often the PFC watchdog checks the status of PFC queues. Configure the polling interval in milliseconds.
    - Default: 100
    - Range: 100-1000
  - recovery *time*—Configure in milliseconds how long the PFC watchdog disables the affected queues before it re-enables PFC.
    - Default: 200
    - Range: 200-10,000
  - watchdog-action drop—When the PFC watchdog detects that a PFC queue has stalled, it drops all queued packets and all newly arriving packets for the stalled PFC queue. This option is the default.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support for DSCP values introduced in Junos OS Release 17.4R1 for the QFX Series.

`pfc-watchdog` option introduced in Junos OS Evolved Release 20.4R1.

RELATED DOCUMENTATION

- Configuring CoS PFC (Congestion Notification Profiles)
- *Understanding CoS Flow Control (Ethernet PAUSE and PFC)*
- Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows
- Understanding PFC Using DSCP at Layer 3 for Untagged Traffic

Configuring DSCP-based PFC for Layer 3 Untagged Traffic

PFC Watchdog

connect-actively

IN THIS SECTION

- Syntax | 4529
- Hierarchy Level | 4529
- Description | 4530
- Options | 4530
- Required Privilege Level | 4531
- Release Information | 4531

Syntax

```
connect-actively  {
    <capabilities-exchange-timeout seconds>;
    <port port-number>;
    <repeat-timeout seconds>;
    <retry-timeout seconds>;
    <timeout seconds>;
    transport transport-name;
}
```

Hierarchy Level

```
[edit access diameter peer peer-name]
```

Description

Define the destination port and transport connection used to establish active connections to the Diameter peer.

Options

`capabilities-exchange-timeout` *`seconds`*

(Optional) Use the specified amount of time to wait for a Capabilities-Exchange-Answer message.

- **Range:** 1 through 65,535 seconds
- **Default:** 10 seconds

`port` *`port-number`*

(Optional) Use the specified destination TCP port.

- **Default:** 3868

`repeat-timeout` *`seconds`*

(Optional) Use the specified amount of time to wait before attempting to reconnect to this peer after receiving the DO_NOT_WANT_TO_TALK_TO_YOU value for the Disconnect-Cause AVP in the Disconnect-Peer-Request message. A value of zero means that there is no attempt to reconnect to the peer.

- **Range:** 0 through 65,535 seconds
- **Default:** 0

`retry-timeout` *`seconds`*

(Optional) Use the specified amount of time to wait between connection attempts for this peer.

- **Range:** 1 through 65,535 seconds
- **Default:** 30 seconds

`timeout` *`seconds`*

(Optional) Use the specified amount of time to wait for connection acknowledgement for this peer.

- **Range:** 1 through 65,535 seconds
- **Default:** 10 seconds

`transport` *`transport-name`*

Use the specified name of the transport layer connection.

> **NOTE**: The specified transport must already be configured at the `[edit access diameter transport]` hierarchy level.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

connection-flood

IN THIS SECTION

- Syntax | 4532
- Hierarchy Level | 4532
- Description | 4532
- Options | 4532
- Required Privilege Level | 4532
- Release Information | 4532

Syntax

```
connection-flood {
    threshold rate;
}
```

Hierarchy Level

```
[edit logical-systems name security alg mgcp application-screen],
[edit logical-systems name tenants name security alg mgcp application-screen],
[edit security alg mgcp application-screen],
[edit services alg mgcp application-screen],
[edit tenants name security alg mgcp application-screen]
```

Description

Configure the connection-flood parameter to limit the number of new connection requests allowed per Media Gateway (MG) per second. Media Gateway Control Protocol (MGCP) Application Layer Gateway (ALG) drops the messages exceeding the threshold.

Options

`threshold` *rate*—Number of connection requests per second allowed per MG.

- **Range:** 2 through 10000
- **Default:** 200

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

connection (Identity Management Advanced Query)

IN THIS SECTION

- Syntax | 4533
- Hierarchy Level | 4534
- Description | 4534
- Options | 4535
- Required Privilege Level | 4536
- Release Information | 4536

Syntax

```
connection {
    connect-method (http | https);
    port port;
    primary {
        address address;
        ca-certificate ca-certificate;
        client-id client-id;
        client-secret client-secret;
        interface interface-name;
        routing-instance routing-instance -name;
        source source-address;
    }
    query-api query-api;
    secondary{
        address address;
        ca-certificate ca-certificate;
        client-id client-id;
        client-secret client-secret;
        interface interface-name;
        routing-instance routing-instance -name;
        source source-address;
    }
```

```
    token-api token-api;
}
```

Hierarchy Level

```
[edit services user-identification identity-management]
```

Description

Configure parameters for connecting the SRX Series to the Juniper Identity Management Service (JIMS) server to obtain user identity and device information.

For the SRX Series Firewall to obtain user identity information, you must first establish a connection to the JIMSserver. The parameters to specify for the connection include the protocol, the IP address of the JIMS server, and the information to authenticate the SRX Series Firewall to the JIMS server.

If you are using more than one JIMS server, you must configure each server separately. The SRX Series Firewall always attempts to connect to the primary server first. If the primary server fails, the SRX Series Firewall falls back to the secondary server. The SRX Series Firewall periodically probes the failed primary server and reverts to it when it is available.

Only configuration of the primary server is mandatory. You are not required to use a secondary server.

The SRX Series advanced user identity query feature queries the JIMS for user identity information that the SRX Series stores in its authentication table and uses to authenticate users. Use of the JIMS allows you to provision users locally and have their authentication information made available to other sites in your network for policy enforcement and reporting.

> **WARNING**: Before you use this feature, you must disable any other actively used options under the [edit services user-identification] hierarchy. You cannot commit this configuration if active directory authentication and the ClearPass query and webapi functions are configured and committed.

To obtain device information, such as device identity, groups, and the operating system, from the JIMS server using either the batch-query or ip-query configuration, you must set the device authentication source, as follows.

```
user@host# set services user-identification device-information authentication-source network-
access-controller
```

connect-method- Configure the protocol to be used for the SRX Series Firewall connection to JIMS. The SRX Series Firewall connects to the JIMS to obtain user identity information.

port- Configure the port on the JIMS server that the SRX Series Firewall uses to connect to the server.

query-api- Configure the prefix of the URL path for querying user identities. This value is used to construct the prefix of the path for queries for individual users, as well as for ip-query and batch-query requests, each of which has a unique suffix:

- For IP query, *query-api*/ip/
- For batch query, *query-api*/users/
- For user-query *query-api*/user

The default value for query-api is user-query/v2.

For example, for a batch query, assume that the query API is configured as `user-query/v2`. To generate the complete URL, the prefix is combined with the connection method, which is `HTTPS`, the IP address of the JIMS server, expressed as a variable in this example (*JIMS*), the beginning timestamp, `begintime={timestamp}`, and the number of user identity information items to be provided in the record that the JIMS server returns, `entry_count={count}`.

```
'https://JIMS/user_query/v2/users/endpoints?begintime={timestamp}&entry_count={count}'
```

token-api- The path of the URL for acquiring the access token for OAuth2 authentication (RFC 6749). The JIMS server requires that the SRX Series Firewall authenticate to it using OAuth2. The SRX Series Firewall uses the Client Credentials grant type for this purpose.

The following example shows the default tokenAPI, `oauth_token/oauth`, combined with the connection method, `https`, and the JIMS server IP address placeholder to create the complete URL:

```
https://JIMS/oauth_token/oauth.
```

The advanced user identity query feature, to which this statement belongs, allows you to obtain user identity information from the JIMS through queries. It allows you to provision users locally and have their authentication information made available to other sites in your network for policy enforcement and reporting.

Options

connect-method

Method of connection

- Values:

- http—HTTP connection
- https—HTTPS connection

port Server port

- **Default:** 443
- **Range:** 1 through 65535

query-api Query API

token-api API of acquiring token for OAuth2 authentication

The remaining statements are described separately.

Required Privilege Level

1. services—To view this statement in the configuration.
2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

Source, interface, and routing-instance options are introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Configure Juniper Identity Management Service to Obtain User Identity Information

authentication-entry-timeout (Identity Management Advanced Query) | **2597**

batch query | **2886**

filter (Identity Management Advanced Query) | **7138**

ip-query (Identity Management Advanced Query) | **9259**

primary connection (Identity Management Advanced Query) | **13131**

secondary connection (Identity Management Advanced Query) | **16409**

connection-limit

IN THIS SECTION

- Syntax | 4537
- Hierarchy Level | 4537
- Description | 4537
- Options | 4537
- Required Privilege Level | 4538
- Release Information | 4538

Syntax

```
connection-limit limit;
```

Hierarchy Level

```
[edit system services finger],
[edit system services ftp],
[edit system services netconf ssh],
[edit system services ssh],
[edit system services telnet],
[edit system services tftp-server],
[edit system services xnm-clear-text],
[edit system services xnm-ssl]
```

Description

Configure the maximum number of connections sessions for each type of system service (finger, ftp, ssh, telnet, xnm-clear-text, or xnm-ssl) per protocol (either IPv6 or IPv4).

Options

limit—(Optional) Maximum number of established connections per protocol (either IPv6 or IPv4).

- **Range:** 1 through 250
- **Default:** 75
- **Default (EX Series and QFX Series):** 10

> **NOTE**: The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring DTCP-over-SSH Service for the Flow-Tap Application
- Configuring SSH Service for Remote Access to the Router or Switch

connection-limit (REST API)

IN THIS SECTION

- Syntax | 4539
- Hierarchy Level | 4539
- Description | 4539
- Options | 4539
- Required Privilege Level | 4539
- Release Information | 4539

Syntax

```
connection-limit limit;
```

Hierarchy Level

```
[edit system services rest control]
```

Description

Specify the maximum number of simultaneous connections for the REST API process.

Options

limit Maximum number of simultaneous connections (IPv4 only).

- **Range:** 1 through 1024
- **Default:** 64

> **NOTE**: The maximum number of simultaneous connections that the SRX series supports:
> SRX300 and SRX320 - 8 sessions
>
> SRX34X and SRX550 - 16 sessions

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

connection-protection

IN THIS SECTION

- Syntax | 4540
- Hierarchy Level | 4540
- Description | 4540
- Required Privilege Level | 4540
- Release Information | 4541

Syntax

```
connection-protection;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface],
[edit protocols l2circuit local-switching interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name]
```

Description

Enable connection protection on the Layer 2 circuit. Connection protection enables you to configure a redundant pseudowire to act as a backup connection and can be configured between PE routers, maintaining Layer 2 circuit and VPLS services after certain types of failures. This feature helps to improve the reliability of networks where a single point of failure could interrupt service for customers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring Layer 2 Circuit Switching Protection

connection-protection-tlv

IN THIS SECTION

- Syntax | 4541
- Hierarchy Level | 4541
- Description | 4541
- Required Privilege Level | 4542

Syntax

```
connection-protection-tlv;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name continuity-check]
```

Description

Includes connection protection OUI TLV in continuity check messages (CCM) .The TLV is responsible for carrying the flag information within CCM PDUs. Though this OUI TLV will be included in the CCM frames by provider edge devices, the value is updated by the provider routers in case the traffic to the other end of the network is forwarded by the facility protection tunnel .

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

connectivity-fault-management | 4556

Configuring MAC Flush Message Processing in CET Mode

Example: Configuring an Action Profile Based on Connection Protection TLVs

connections (Circuits)

IN THIS SECTION

- Syntax | 4542
- Hierarchy Level | 4543
- Description | 4543
- Options | 4543
- Required Privilege Level | 4543
- Release Information | 4543

Syntax

```
connections {
    interface-switch connection-name {
        interface interface-name.unit-number;
    }
    lsp-switch connection-name {
        transmit-lsp label-switched-path;
        receive-lsp label-switched-path;
    }
    p2mp-receive-switch {
```

```
        output-interface [ interface-name.unit-number ];
        receive-p2mp-lsp receiving-point-to-multipoint-lsp;
    }
    p2mp-transmit-switch {
        input-interface interface-name.unit-number;
        transmit-p2mp-lsp transmitting-point-to-multipoint-lsp;
    }
    remote-interface-switch connection-name {
        interface interface-name.unit-number;
        receive-lsp label-switched-path;
        transmit-lsp label-switched-path;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit protocols]
```

Description

Define the connection between two circuits in a CCC connection.

Options

The remaining statements are explained separately. See CLI Explorer.

NOTE: The **edit logical-systems** hierarchy is not available on QFabric systems.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Layer 2 Switching Cross-Connects Using CCC

Configuring MPLS LSP Tunnel Cross-Connects Using CCC

Configuring TCC

Configuring CCC Switching for Point-to-Multipoint LSPs

connections (Host VPN)

IN THIS SECTION

- Syntax | 4544
- Hierarchy Level | 4545
- Description | 4545
- Options | 4545
- Required Privilege Level | 4546
- Release Information | 4547

Syntax

```
connections {
    connection-name {
        children {
            child-name {
                esp-proposal esp-proposal;
                local-traffic-selector {
                    (ipv4-prefix ipv4-prefix | ipv6-prefix ipv6-prefix);
                    port port;
                    protocol protocol;
                }
                mode (transport | tunnel);
                rekey-time rekey-time;
                remote-traffic-selector {
                    (ipv4-prefix ipv4-prefix | ipv6-prefix ipv6-prefix);
                    port port;
```

```
                    protocol protocol;
                }
            }
        }
        dpd-delay dpd-delay;
        ike-proposal ike-proposal;
        local {
            id local-id;
        }
        local-address {
            (ipv4 ipv4-address | ipv6 ipv6-address);
        }
        rekey-time rekey-time;
        remote {
            id remote-id;
        }
    }
    remote {
        id remote-id;
    }
    remote-address {
        (ipv4 ipv4-address | ipv6 ipv6-address);
    }
}
```

Hierarchy Level

```
[edit security host-vpn]
```

Description

Configure IPsec connection details. The Internet Key Exchange (IKE) protocol is a component of IPsec used for performing mutual authentication and establishing and maintaining security associations (SAs). Each SA describes a specific negotiated set of parameters to protect traffic for a certain time period for an IPsec VPN.

Options

connection-name Specify the name of the IKE SA connection.

dpd-delay Specify the Dead Peer Detection delay used on the connection. This is the Interval between sending liveness messages.

- **Default:** 0, which is disabled Dead Peer Dectection.
- **Range:** 0 through 3600

ike-proposal *ike-proposal* Specify the algorithms to use in negotiating the IKE SA from among the pre-selected combinations available, which represent the encryption algorithm, integrity algorithm, and Diffie Hellman group.

3des-sha1-modp1536	Propose 3des SHA1 and DH group modp1536.
aes256gcm128-ecp384	Propose aes256gcm128 and DH group ecp384.
aes256gcm128-modp3072	Propose aes256gcm128 and DH group modp3072.
aes256-sha384-ecp384	Propose aes256 CBC, sha384 and DH group ecp384.
aes256-sha384-modp3072	Propose aes256 CBC, sha384 and DH group modp3072.
[ ]	Propose a set composed from the values permitted.

- **Default:** aes256-sha384-ecp384

local-address Specify the local endpoint's IPv4 or IPv6 address.

rekey-time *rekey-time* Specify how long in seconds before the IKE SA is rekeyed. Actual rekeying occurs slightly sooner than that specified because of rekey randomization.

- **Default:** 14,400
- **Range:** 60 through 86,400

remote-address Specify the remote endpoint's IPv4 or IPv6 address.

The remaining statements are explained separately.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 18.3R1.

connections (MPLS)

IN THIS SECTION

- Syntax | 4547
- Hierarchy Level | 4547
- Description | 4547
- Required Privilege Level | 4548
- Release Information | 4548

Syntax

```
connections {
                              remote-interface-switch                        connection-name {
        interface interface-name.unit-number;
        transmit-lsp label-switched-path;
        receive-lsp label-switched-path;
    }
```

Hierarchy Level

```
[edit protocols]
```

Description

Define the connection between two circuits in a CCC connection.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Example: Configuring MPLS on EX8200 and EX4500 Switches

connectivity-association

IN THIS SECTION

- Syntax | 4548
- Syntax (SRX Series Firewalls) | 4549
- Hierarchy Level | 4550
- Description | 4550
- Default | 4550
- Options | 4550
- Required Privilege Level | 4550
- Release Information | 4551

Syntax

```
connectivity-association connectivity-association-name {
    cipher-suite (MACsec) encryption-algorithm-name;
    exclude-protocol protocol-name;
    fallback-key {
        cak hexadecimal-number;
```

```
        ckn hexadecimal-number;
    }
    include-sci;
    mka {
        must-secure;
        key-server-priority priority-number;
        transmit-interval interval;
    }
    no-encryption;
    offset (0|30|50);
    pre-shared-key {
        cak hexadecimal-number;
        ckn hexadecimal-number;
    }
    replay-protect{
        replay-window-size number-of-packets;
    }
    secure-channel secure-channel-name {
        direction (inbound | outbound);
        encryption (MACsec);
        id {
            mac-address mac-address;
            port-id port-id-number;
        }
        offset (0|30|50);
        security-association security-association-number {
            key key-string;
        }
    }
    security-mode security-mode;
}
```

Syntax (SRX Series Firewalls)

```
connectivity-association connectivity-association-name;
    exclude-protocol protocol-name;
    include-sci;
    mka {
        must-secure;
        key-server-priority priority-number;
        transmit-interval interval;
```

```
    }
    no-encryption;
    offset (0|30|50);
    pre-shared-key {
        cak hexadecimal-number;
        ckn hexadecimal-number;
    }
    replay-protect{
        replay-window-size number-of-packets;
    }
    security-mode security-mode;
}
```

Hierarchy Level

```
[edit security macsec]
```

Description

Create or configure a MACsec connectivity association.

A connectivity association is not applying MACsec to traffic until it is associated with an interface. MACsec connectivity associations are associated with interfaces using the "interfaces" on page 9147 statement in the [edit security macsec] hierarchy.

Default

No connectivity associations are present, by default.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced on SRX Series Firewalls in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

connectivity-association (MACsec Interfaces)

IN THIS SECTION

- Syntax | **4551**
- Hierarchy Level | **4551**
- Description | **4552**
- Default | **4552**
- Required Privilege Level | **4552**
- Release Information | **4552**

Syntax

```
connectivity-association connectivity-association-name;
```

Hierarchy Level

```
[edit security macsec interfaces interface-name]
```

Description

Applies a connectivity association to an interface, which enables Media Access Control Security (MACsec) on that interface.

The point-to-point links are secured after security keys are matched at the endpoints of the links. If you enable MACsec by using the static connectivity association key (CAK) security mode, user-configured, preshared keys are matched. If you enable MACsec by using the static secure association key (SAK) security mode, user-configured static security association keys are matched.

When you enable MACsec using static CAK, you have to create and configure a connectivity association.

Default

No connectivity associations are associated with any interfaces.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced on SRX devices in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

connectivity-association (MACsec Interfaces for MX Series)

IN THIS SECTION

- Syntax | 4553
- Hierarchy Level | 4553
- Description | 4553
- Default | 4553
- Options | 4553
- Required Privilege Level | 4554
- Release Information | 4554

Syntax

```
connectivity-association connectivity-association-name;
```

Hierarchy Level

```
[edit security macsec interfaces interface-name]
```

Description

Applies a connectivity association to an interface, which enables Media Access Control Security (MACsec) on that interface.

Default

No connectivity associations are associated with any interfaces.

Options

connectivity-association-name	Name of the MACsec connectivity association.

- **Range:** 1 through 32 alphanumeric characters. Allowed characters are [a-z, A-Z, 0-9]

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

connectivity-association (MX Series)

IN THIS SECTION

- Syntax | 4554
- Hierarchy Level | 4555
- Description | 4555
- Default | 4556
- Options | 4556
- Required Privilege Level | 4556
- Release Information | 4556

Syntax

```
connectivity-association connectivity-association-name {
    exclude-protocol protocol-name;
```

```
    include-sci;
    mka {
        bounded-delay;
        key-server-priority priority-number;
        must-secure;
        should-secure;
        transmit-interval interval;
        eapol-address (pae | provider-bridge | lldp-multicast | destination unicast-address);
    }
    no-encryption;
    offset (0|30|50);
    pre-shared-key {
        cak hexadecimal-number;
        ckn hexadecimal-number;
    }
    replay-protect{
        replay-window-size number-of-packets;
    }
    secure-channel secure-channel-name {
        direction (inbound | outbound);
        encryption ;
        id {
            mac-address mac-address;
            port-id port-id-number;
        }
        offset (0|30|50);
        security-association security-association-number {
            key key-string;
        }
    }
    security-mode security-mode;
}
```

Hierarchy Level

```
[edit security macsec]
```

Description

Create or configure a MACsec connectivity association.

A connectivity association is not applying MACsec to traffic until it is associated with an interface. MACsec connectivity associations are associated with interfaces using the `interfaces` statement in the `[edit security macsec]` hierarchy.

Default

No connectivity associations are present, by default.

Options

connectivity-association-name

Name of the MACsec connectivity association.

- **Range:** 1 through 32 alphanumeric characters. Allowed characters are [a-z, A-Z, 0-9]

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

connectivity-fault-management

IN THIS SECTION

Syntax | 4557

Hierarchy Level | 4559

- Description | 4559
- Options | 4560
- Required Privilege Level | 4560
- Release Information | 4560

Syntax

```
connectivity-fault-management {
    action-profile profile-name {
        action {synchronous-notification
            interface-down;
            log-and-generate-ais {
                interval(1m | 1s);
                levelvalue;
                priority value;
            }
        }
        default-actions {
            interface-down;
        }
        event {
            ais-trigger-condition {
                adjacency-loss;
                all-defects;
                cross-connect-ccm;
                erroneous-ccm;
                receive-ais;
            }
            adjacency-loss;
            interface-status-tlv (down | lower-layer-down);
            port-status-tlv blocked;
            rdi;
        }
    }
    linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
    }
```

```
    expected-defect {
        rx-enable;
        rx-max-duration seconds;
        tx-enable;
        tx-duration seconds;
    }
    maintenance-domain domain-name {
        bridge-domain <vlan-id [ vlan-ids ]>;
        instance routing-instance-name;
        interface interface-name;
        level number;
        name-format (character-string | none | dns | mac+2oct);
        maintenance-association ma-name {
            protect-maintenance-association protect-ma-name;
            remote-maintenance-association remote-ma-name;
            short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
            continuity-check {
                convey-loss-threshold;
                hold-interval minutes;
                interface-status-tlv;
                interval (10m | 10s | 1m | 1s| 100ms);
                loss-threshold number;
                port-status-tlv;
            }
            mep mep-id {
                auto-discovery;
                direction (up | down);
                interface interface-name (protect | working);
                lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
rem-err-xcon | xcon );
                priority number;
                remote-mep mep-id {
                    action-profile profile-name;
                    sla-iterator-profile profile-name {
                        data-tlv-size size;
                        iteration-count count-value;
                        priority priority-value;
                        detect-loc;
                    }
                }
            }
        }
        virtual-switch routing-instance-name {
```

```
            bridge-domain name <vlan-ids [ vlan-ids ]>;
        }
    }
    no-aggregate-delegate-processing;
    performance-monitoring {
        delegate-server-processing;
        hardware-assisted-timestamping;
        hardware-assisted-keepalives;
        sla-iterator-profiles {
            profile-name {
                avg-fd-twoway-threshold;
                avg-ifdv-twoway-threshold;
                avg-flr-forward-threshold;
                avg-flr-backward-threshold;
                disable;
                calculation-weight {
                    delay delay-weight;
                    delay-variation delay-variation-weight;
                }
                cycle-time milliseconds;
                iteration-period connections;
                measurement-type (loss | statistical-frame-loss | two-way-delay);
            }
        }
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet]
```

Description

For Ethernet interfaces on M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M120, M320, MX Series, and T Series routers, specify connectivity fault management for IEEE 802.1ag Operation, Administration, and Management (OAM) support.
In Junos OS Release 9.3 and later, this statement is also supported on aggregated Ethernet interfaces.

The remaining statements are explained separately. See CLI Explorer.

Options

Asynchronous-notification—Generate asynchronous-notification (Laser off for optical physical) option is introduced in Junos OS release 20.4R1 for ACX and MX.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

IEEE 802.1ag OAM Connectivity Fault Management Overview

connectivity-fault-management-switches

IN THIS SECTION

- Syntax | 4561
- Hierarchy Level | 4562
- Description | 4562
- Options | 4562
- Required Privilege Level | 4562
- Release Information | 4563

Syntax

```
connectivity-fault-management {
    action-profile profile-name {
        action {
            asynchronous-notification
interface-down;
        }
        default-actions {
            interface-down;
        }
        event {
            adjacency-loss;
        }
    }
    linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
    }
    maintenance-domain domain-name {
        level number;
        mip-half-function (none | default |explicit);
        name-format (character-string | none | dns | mac+2oct);
        maintenance-association ma-name {
            continuity-check {
                hold-interval minutes;
                interface-status-tlv;
                interval (10m | 10s | 1m | 1s| 100ms);
                loss-threshold number;
                port-status-tlv;
            }
            mep mep-id {
                auto-discovery;
                direction down;
                interface interface-name;
                remote-mep mep-id {
                    action-profile profile-name;
                }
            }
        }
    }
    performance-monitoring {
```

```
        sla-iterator-profiles {
            profile-name {
                calculation-weight {
                    delay delay-value;
                    delay-variation delay-variation-value;
                }
send-zero-interval-ccm;
                cycle-time cycle-time-value;
                iteration-period iteration-period-value;
                measurement-type two-way-delay;
                passive;
            }
        }
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet]
```

Description

Configure connectivity fault management for IEEE 802.1ag Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately. See CLI Explorer.

Options

Asynchronous-notification—Generate asynchronous-notification (Laser off for optical physical) option is

introduced in Junos OS release 20.4R1 for ACX and MX.

send-zero-interval-ccm—Programs end point hardware for TX with 0s CCI. This option is introduced in Junos OS release 21.2R1 for ACX5448, ACX5448-M, and ACX5448-D only.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

performance-monitoring introduced in Junos OS Release 11.4 for EX Series switches.

RELATED DOCUMENTATION

Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)

Junos OS Network Interfaces Configuration Guide

connectivity-type

IN THIS SECTION

- Syntax | 4563
- Hierarchy Level | 4564
- Description | 4564
- Default | 4564
- Options | 4564
- Required Privilege Level | 4564
- Release Information | 4565

Syntax

```
connectivity-type (ce | irb | permanent);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Specify when a VPLS connection is taken down depending on whether or not the interface for the VPLS routing instance is customer-facing or integrated routing and bridging (IRB).

> **NOTE**: The `connectivity-type` statement is not supported for FEC 129 VPLS (also known as LDP VPLS with BGP-based autodiscovery).

Default

`ce`

Options

`ce`—Require that for the VPLS connection to be up, the customer-facing interface for the VPLS routing instance must also be up. If the customer-facing interface fails, the VPLS connection is taken down.

`irb`—Allow a VPLS connection to remain up so long as an IRB interface is configured for the VPLS routing instance.

`permanent`—Allow a VPLS connection to remain up until specifically taken down. This option is reserved for use in configuring Layer 2 Wholesale subscriber networks. See the *Broadband Subscriber Management Solutions Guide* for details about configuring a Layer 2 Wholesale network.

> **NOTE**: To specifically take down a VPLS routing instance that is using the `permanent` option, all associated static logical interfaces must also be down.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`irb` option introduced in Junos OS Release 9.3.

`permanent` option introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring VPLS Routing Instance and VPLS Interface Connectivity

Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers

console (System Logging)

IN THIS SECTION

- Syntax | 4565
- Hierarchy Level | 4566
- Description | 4566
- Options | 4566
- Required Privilege Level | 4566
- Release Information | 4567

Syntax

```
console {
    facility severity;
    application-status (any | error);
}
```

Hierarchy Level

```
[edit system syslog]
```

Description

Configure the logging of system messages to the system console. Log messages include priority information, which is information about log messages' facility and severity levels.

Options

facility Class (type) of messages to log. To specify multiple classes, include multiple *facility severity* statements.

severity Severity of the messages that belong to the facility specified by the paired *facility* name. Messages with severities of the specified level and higher are logged. You can specify the minimum severity level of a message.

> **NOTE**: For a list of the facilities and message severities, see *Junos OS System Logging Facilities*.

application-status (any | error) (Junos OS Evolved) Enable logging of application start or stop status. During system reboot, by default, the systemd process logs messages for applications to the journal log. By default, application stop and start status are written to the journal log. Only application stop and start errors are logged to the console.

This configuration statement allows each application's stop and start application status logs to be logged to the console. You configure the `any` option to allow all systemd application stop and start messages to be written to the console, not just error messages. You configure the `error` option to return to the default setting.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`application-status` option introduced in Junos OS Evolved 23.1R1.

RELATED DOCUMENTATION

> *Directing System Log Messages to the Console*
>
> System Log Explorer

console (System Ports)

IN THIS SECTION

- Syntax | 4567
- Hierarchy Level | 4568
- Description | 4568
- Default | 4568
- Options | 4568
- Required Privilege Level | 4569
- Release Information | 4569

Syntax

```
console {
    authentication-order [authentication-methods ];
    disable;
    insecure;
    log-out-on-disconnect;
    type terminal-type;
}
```

Hierarchy Level

```
[edit system ports]
```

Description

Configure the characteristics of the console port.

Default

The console port is enabled and its speed is 9600 baud.

Options

The following section lists all possible options for the `console` statement. The options that appear may vary depending on your platform and software release.

`authentication-order`—Specify the order in which the authentication methods such as **password** (for local password authentication), **radius** (for RADIUS server authentication), or **tacplus** (for TACACS+ server authentication) should be attempted.

`authentication-methods`—One or more authentication methods, listed in the order in which they should be tried. The method can be one or more of the following:

- `password`—Use the password configured for the user with the `authentication` statement at the `[edit system login user]` hierarchy level.
- `radius`—Use RADIUS authentication services.
- `tacplus`—Use TACACS+ authentication services.

`disable`—Disable console login connections.

`insecure`—Disable root login connections to the console and auxiliary ports. Configuring the console port as insecure also prevents superusers and anyone with a user identifier (UID) of 0 from establishing terminal connections in multiuser mode. This option can be used to prevent a user from attempting password recovery by booting into single-user mode, if the user does not know the root password.

`log-out-on-disconnect`—Log out the session when the data carrier on the console port is lost.

NOTE:

- The `log-out-on-disconnect` option is not operational on MX80 routers. On MX80 routers you must manually log out from the console with the `request system logout u0` command.
- The `log-out-on-disconnect` option is not operational on guest network functions (GNFs), which are managed using Juniper Device Manager (JDM). You must use the `exit` command to log out from the GNF console. For more information, see Junos Node Slicing user guide.

`type` *`terminal-type`*—Type of terminal that is connected to the port.

- **Range:** `ansi`, `vt100`, `small-xterm`, `xterm`
- **Default:** The terminal type is unknown, and the user is prompted for the terminal type.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`disable` option added in Junos OS Release 7.6.

`authentication-order` option added in Junos OS Release 12.2R3.

RELATED DOCUMENTATION

- Configuring Junos OS to Set Console and Auxiliary Port Properties
- Configuring Console and Auxiliary Port Properties
- Determine the Authentication Order for LDAPS, RADIUS, TACACS+, and Password Authentication

constant

IN THIS SECTION

- Syntax | 4570
- Hierarchy Level | 4570
- Description | 4570
- Options | 4570
- Required Privilege Level | 4570
- Release Information | 4571

Syntax

```
constant value;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber subscription-id]
```

Description

Specify a constant string for the Subscription-Id-Data value for IP-based subscribers. This constant value is used if none of the `subscription-id-options` methods can be used. In such a case, the Subscription-Id-Type is END_USER_PRIVATE.

Options

value String that is used for the Subscription-Id-Data value.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

constituent-list

IN THIS SECTION

- Syntax | 4571
- Hierarchy Level | 4572
- Description | 4572
- Options | 4572
- Required Privilege Level | 4572
- Release Information | 4572

Syntax

```
constituent-list constituent-list-name {
    (administrative-group [ group-names ] | administrative-group-extended [ extended-
administrative-group-names ] | srlg [ srlg-names ]);
}
```

Hierarchy Level

```
[edit logical-systems logical-systems-name protocols mpls],
[edit protocols mpls]
```

Description

Create a list of traffic engineering attributes called constituent attributes, which are the link and node attributes whose logical combination makes up an abstract hop. The constituent attributes are listed under administrative groups, extended administrative groups, and Shared Risk Link Groups (SRLGs).

Options

constituent-list-name	Name of the constituent list that includes constituent traffic engineering attributes for use in the abstract hop definition.
administrative-group *[group-names]*	Names of administrative groups to include in the constituent list.
administrative-group-extended *[extended-administrative-group-names]*	Names of extended administrative groups to include in the constituent list.
srlg *[srlg-names]*	Names of SRLGs to include in the constituent list.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Example: Configuring Abstract Hops for MPLS LSPs

abstract-hop | **1625**

show mpls abstract-hop-membership | **22125**

show mpls lsp abstract-computation | 22203

cont-stats-collection-interval

IN THIS SECTION

- Syntax | 4573
- Hierarchy Level | 4573
- Description | 4573
- Default | 4574
- Options | 4574
- Required Privilege Level | 4574
- Release Information | 4574

Syntax

```
cont-stats-collection-interval interval;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options multicast],
[edit logical-systems name routing-options multicast],
[edit logical-systems name tenants name routing-instances name routing-options multicast],
[edit routing-instances name routing-options multicast],
[edit routing-options multicast],
[edit tenants name routing-instances name routing-options multicast]
```

Description

Change the default interval (in seconds) at which continuous, persistent IGMP and MLD statistics are stored on devices that support continuous statistics collection.

Junos OS multicast devices collect statistics of received and transmitted IGMP and MLD control packets for active subscribers. Devices that support continuous IGMP and MLD statistics collection also maintain persistent, continuous statistics of IGMP and MLD messages for past and currently active subscribers. The device preserves these continuous statistics across routing daemon restarts, graceful Routing Engine switchovers, ISSU, or line card reboot operations. Junos OS stores continuous statistics in a shared database and copies it to the backup Routing Engine at this configured interval to avoid too much processing overhead on the Routing Engine.

The `show igmp statistics` and `show mld statistics` CLI commands display currently active subscriber IGMP or MLD statistics by default, or you can include the `continuous` option with either of those commands to display the continuous statistics instead.

Default

300 seconds (5 minutes)

Options

interval Interval in seconds at which you want the device to store collected continuous IGMP and MLD statistics.

- **Range:** 60 seconds to 3600 seconds (5 minutes to 1 hour).

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

show igmp statistics | 20203

show mld statistics | 22104

clear igmp statistics | 3627

clear mld statistics | 3707

contact (SNMP)

IN THIS SECTION

- Syntax | 4575
- Hierarchy Level | 4575
- Description | 4575
- Options | 4575
- Required Privilege Level | 4575
- Release Information | 4576

Syntax

```
contact contact;
```

Hierarchy Level

```
[edit snmp]
```

Description

Defines the value of the SNMP MIB II **sysContact** object, which is the contact person for the managed system.

Options

contact—Name of the contact person. If the name includes spaces, enclose it in quotation marks (" ").

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the System Contact on a Device Running Junos OS

container-devices

IN THIS SECTION

- Syntax | 4576
- Hierarchy Level | 4576
- Description | 4577
- Options | 4577
- Required Privilege Level | 4577
- Release Information | 4577

Syntax

```
container-devices {
    device-count number;
}
```

Hierarchy Level

```
[edit chassis]
[edit chassis interconnect-device name]
[edit chassis node-group name]
```

Description

Specify the container devices configuration. The **number** option specifies the number of sequentially numbered container interfaces, from **ci0** to **ci127** maximum.

Options

number—Number of container devices.

- **Range:** 1 through 128

Required Privilege Level

chassis—To view this statement in the configuration.

chassis-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

container-label-switched-path

IN THIS SECTION

- Syntax | 4578
- Hierarchy Level | 4578
- Description | 4578
- Options | 4578
- Required Privilege Level | 4578
- Release Information | 4579

Syntax

```
container-label-switched-path lsp-name {
    disable;
    description description;
    label-switched-path-template;
    splitting-merging;
    suffix string;
    to ip-address;
}
```

Hierarchy Level

```
[edit protocols mpls]
```

Description

Configure a multi-label-switched path (LSP) tunnel between the ingress and the egress routers. The container LSP consists of several member LSPs to the same destination.

Options

disable	Disable MPLS container-label-switched path.
description *description*	Text describing the container LSP.
suffix *string*	Suffix to generate names of member LSPs of the container LSP.
to *ip-address*	IP address of the egress router.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.

content-destination

IN THIS SECTION

- Syntax | 4579
- Hierarchy Level | 4579
- Description | 4580
- Options | 4580
- Required Privilege Level | 4580
- Release Information | 4580

Syntax

```
content-destination identifier {
    address address;
    hard-limit bandwidth;
    hard-limit-target bandwidth;
    soft-limit bandwidth;
    soft-limit-clear bandwidth;
    ttl hops;
}
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Identify the destination for captured packets.

Options

identifier—Name of the destination.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Content Destination

content-filtering (Logical System Security Feature Profile)

IN THIS SECTION

- Syntax | 4581
- Hierarchy Level | 4581
- Description | 4581
- Options | 4582
- Required Privilege Level | 4582
- Release Information | 4582

Syntax

```
content-filtering {
    profile name {
        block-command block-command;
        block-content-type {
            activex;
            exe;
            http-cookie;
            java-applet;
            zip;
        }
        block-extension block-extension;
        block-mime {
            exception (Security Content Filtering) exception;
            list (Security Content Filtering Block Mime) list;
        }
        notification-options (Security Content Filtering) {
            custom-message (Security Content Filtering) custom-message;
            (notify-mail-sender (Security Content Filtering Notification Options) | no-notify-mail-sender
(Security Content Filtering Notification Options));
            type (Security Content Filtering Notification Options) (message | protocol-only);
        }
        permit-command permit-command;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-systems-name security utm feature-profile]
[edit tenants tenant-name security utm feature-profile]
```

Description

Configures the Content Security content-filtering feature for logical systems. The content filtering feature controls file transfers across the gateway by checking traffic against configured filter lists. It evaluates the traffic before all other Content Security features, except Web filtering. You can also configure the default Content Security configuration for content filtering feature profile. If you do not

configure any option in the content filtering feature profile, the values configured in the default Content Security configuration are applied.

> **NOTE**: A license check for the content filtering configuration is performed at the time of commit and provides a warning if a valid license is not installed on the device. Once a valid license is installed on the device then the custom content filtering profile or the default content filtering profile is able to process the traffic. If a license is expired or license is not installed, the content filtering service does not process the traffic.

Options

`block-command`	Protocol block command custom-objects to the content-filtering profile.
`block-content-type`	Blocks other available content such as exe, http-cookie, java-applet. This is for HTTP only.
`block-extension`	Block extensions to the content-filtering profile.
`block-mime`	MIME pattern list custom-objects to the content-filtering profile for blocking MIME types.
`notification-options`	A message notification to trigger when a content filter is matched.
`permit-command`	Protocol permit command custom-objects to the content-filtering profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

utm default-configuration | 30614

content-filtering (Security Content Security Policy)

IN THIS SECTION

- Syntax | 4583
- Hierarchy Level | 4584
- Description | 4584
- Options | 4584
- Required Privilege Level | 4585
- Release Information | 4585

Syntax

```
content-filtering {
 rule-set name {
  rule name {
   match {
    applications [applications];
    direction (any | download | upload);
    file-types [file-types];
   }
   then {
    action (block | close-client | close-client-and-server | close-server | no-action);
    notification {
     endpoint {
      custom-message custom-message;
      (notify-mail-sender | no-notify-mail-sender);
      type (message | protocol-only);
     }
     log;
    }
   }
```

```
    }
   }
  }
```

Hierarchy Level

```
[edit security utm utm-policy policy-name]
[edit logical-systems logical-systems-name security utm utm-policy policy-name]
[edit tenants tenant-name security utm utm-policy policy-name]
```

Description

Configures a Content Security policy for the content filtering protocols and attach this policy to a security profile to implement it. Each supported protocol may implement available content filters differently. Not all filtering capabilities are supported for each protocol. The HTTP protocol supports all content filtering features. The FTP protocol supports only lock Extension List and Protocol Command Block List. The e-mail protocols (SMTP, IMAP, POP3) supports limited to Block Extension List, Protocol Command Block List, and MIME Pattern Filtering.

Starting in Junos OS Release 21.4R1, content filtering is performed by detecting the file content and not the file extensions. We have introduced the `rule-set` and rules configurations under the `[edit security utm utm-policy <utm-policy-name> content-filtering]` hierarchy level. These rules and rule-set allows you to configure direction specific content filters and connection reset.

So, content filtering options based on mime-type, content-type, and protocol command is not supported. After you upgrade to Junos OS Release 21.4R1, previously existing file extension based content filtering options under the `[edit security utm utm-policy <utm-policy-name> content-filtering]` hierarchy are no more available for configuration.

Junos OS Release 21.4R1 allows you to use legacy functionality if you don't want to migrate to this modern functionality. You will be allowed to use the legacy configurations but all the legacy configuration knobs are deprecated and are hidden. Also, you will receive system logs and error message warnings when you use all the legacy deprecated knobs.

Options

rule-set *name* Name of the content filtering rule set.

rule *name* Name of the content filtering rule.

application List of applications to be inspected.

direction Direction of the content to be inspected (upload or download).

file-type List of file-types in match critera.

action Action to perform for the content filtering match condition. The possible options are:

- `no-action`: No action is required.
- `block`: Block and drop the connection
- `close-client`: Close the client connection.
- `close-server`: Close the server connection.
- `close-client-and-server`: Close the client and the server connection.

notification Notification for the content filtering action taken based on the rules configured.

endpoint Endpoint notification options for the content filtering action taken based on the rules configured.

type Endpoint notification type.

notify-mail-sender Notify mail sender.

no-notify-mail-sender Don't notify mail sender.

custom-message Custom notification message.

log Generate security event if content is blocked by rule.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support for configuration in logical systems introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

`rule-set` option introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Content Filtering

content-filtering (Security Feature Profile)

IN THIS SECTION

- Syntax | 4586
- Hierarchy Level | 4587
- Description | 4587
- Options | 4587
- Required Privilege Level | 4588
- Release Information | 4588

Syntax

```
content-filtering {
    block-command;
    block-content-type {
        activex;
        exe;
        http-cookie;
        java-applet;
        zip;
    }
    block-extension;
    block-mime {
        exception;
        list;
    }
    notification-options {
```

```
        custom-message;
        (notify-mail-sender | no-notify-mail-sender);
        log;
    type (message | protocol-only);
    } permit-command;
    rule-set rule-set-name {
          rule rule-name {
          }
        }
    traceoptions {
        flag name;
    }
    type (content-filtering-none | local);
}
```

Hierarchy Level

```
[edit security utm feature-profile]
[edit security utm default-configuration]
```

Description

Configure Content Security content-filtering features. You can also configure the default Content Security configuration for content filtering feature profile. If you do not configure any option in the content filtering feature profile, the values configured in the default Content Security configuration are applied. The content filtering feature controls file transfers across the gateway by checking traffic against configured filter lists. It evaluates the traffic before all other Content Security features, except Web filtering.

> **NOTE**: A license check for the content filtering configuration is performed at the time of a commit and will provide a warning if a valid license is not installed on the device. Once a valid license is installed on the device then a custom content filtering profile or the default profile will be able to process traffic. If a license is expired or is not installed, the content filtering service will not process traffic.

Options

block-command Protocol block command custom-objects to the content-filtering profile.

block-content-type	Blocks to other available content such as exe, http-cookie, java-applet. This is for HTTP only.
block-extension	Block extensions to the content-filtering profile.
block-mime	MIME pattern list custom-objects to the content-filtering profile for blocking MIME types.
notification-options	A message notification to trigger when a content filter is matched.
permit-command	Protocol permit command custom-objects to the content-filtering profile.
traceoptions	Defines tracing operations for default Content Security configuration for content filtering feature.
type	Type of content filtering solution or URL filtering solution used by the device.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Starting in Junos OS Release 21.4R1, content filtering is performed by detecting the file content and not the file extensions. So, content filtering options based on mime-type, content-type, and protocol command is not supported. After you upgrade to Junos OS Release 21.4R1, content filtering option under the `feature-profile` hierarchy are no more available for configuration. The rule-set and rules configurations are introduced under the `[edit security utm utm-policy <utm-policy-name> content-filtering]` hierarchy level. These rules and rule-set allows you to configure direction specific content filters and connection reset.

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Content Security

IN THIS SECTION

- Syntax | 4589
- Hierarchy Level | 4599
- Description | 4599
- Options | 4599
- Required Privilege Level | 4599
- Release Information | 4600

Syntax

```
utm {
    application-proxy {
        traceoptions {
            flag flag;
        }
    }
    custom-objects {
        custom-url-category object-name {
            value [value];
        }
        filename-extension object-name {
            value [value];
        }
        mime-pattern object-name {
            value [value];
        }
        protocol-command object-name {
            value [value];
        }
        url-pattern object-name {
            value [value];
        }
    }
    feature-profile {
```

```
anti-spam {
    address-blacklist list-name;
    address-whitelist list-name;
    sbl {
        profile profile-name {
            custom-tag-string [string];
            (sbl-default-server | no-sbl-default-server);
            spam-action (block | tag-header | tag-subject);
        }
    }
    traceoptions {
        flag flag;
    }
}
anti-virus {
    juniper-express-engine {
        pattern-update {
            email-notify {
                admin-email email-address;
                custom-message message;
                custom-message-subject message-subject;
            }
            interval value;
            no-autoupdate;
            proxy {
                password password-string;
                port port-number;
                server address-or-url;
                username name;
            }
            url url;
        }
        profile profile-name {
             fallback-options {
                content-size (block | log-and-permit);
                default (block | log-and-permit);
                engine-not-ready (block | log-and-permit);
                out-of-resources (block | (log-and-permit);
                timeout (block | log-and-permit);
                too-many-requests (block | log-and-permit);
            }
            notification-options {
                fallback-block {
```

```
                    administrator-email email-address;
                    allow-email;
                    custom-message message;
                    custom-message-subject message-subject;
                    display-host;
                    (notify-mail-sender | no-notify-mail-sender);
                     type (message | protocol-only);
                }
                fallback-non-block {
                    custom-message message;
                    custom-message-subject message-subject;
                    (notify-mail-recipient | no-notify-mail-recipient);
                }
                virus-detection {
                    custom-message message;
                    custom-message-subject message-subject;
                    (notify-mail-sender | no-notify-mail-sender);
                     type (message | protocol-only);
                }
            }
            scan-options {
                content-size-limit value;
                (intelligent-prescreening | no-intelligent-prescreening);
                timeout value;
            }
            trickling {
                timeout value;
            }
        }
    }
    kaspersky-lab-engine {
        pattern-update {
            email-notify {
                admin-email email-address;
                custom-message message;
                custom-message-subject message-subject;
            }
            interval value;
            no-autoupdate;
            proxy {
                password password-string;
                port port-number;
                server address-or-url;
```

```
                    username name;
                }
                url url;
            }
            profile profile-name {
                 fallback-options {
                    content-size (block | log-and-permit);
                    corrupt-file (block | log-and-permit);
                    decompress-layer (block | log-and-permit);
                    default (block | log-and-permit);
                    engine-not-ready (block | log-and-permit);
                    out-of-resources (block | (log-and-permit);
                    password-file (block | (log-and-permit);
                    timeout (block | log-and-permit);
                    too-many-requests (block | log-and-permit);
                }
                notification-options {
                    fallback-block {
                        administrator-email email-address;
                        allow-email;
                        custom-message message;
                        custom-message-subject message-subject;
                        display-host;
                        (notify-mail-sender | no-notify-mail-sender);
                         type (message | protocol-only);
                    }
                    fallback-non-block {
                        custom-message message;
                        custom-message-subject message-subject;
                        (notify-mail-recipient | no-notify-mail-recipient);
                    }
                    virus-detection {
                        custom-message message;
                        custom-message-subject message-subject;
                        (notify-mail-sender | no-notify-mail-sender);
                         type (message | protocol-only);
                    }
                }
                scan-options {
                    content-size-limit value;
                    decompress-layer-limit value;
                    (intelligent-prescreening | no-intelligent-prescreening);
                    scan-extension filename;
```

```
                scan-mode  (all | by-extension);
                timeout value;
            }
            trickling {
                timeout value;
            }
        }
    }
    mime-whitelist {
        exception listname;
        list listname {
            exception listname;
        }
    }
    sophos-engine {
        pattern-update {
            email-notify {
                admin-email email-address;
                custom-message message;
                custom-message-subject message-subject;
            }
            interval value;
            no-autoupdate;
            proxy {
                password password-string;
                port port-number;
                server address-or-url;
                username name;
            }
            url url;
        }
        profile <name> {
            fallback-options {
                content-size (block | log-and-permit | permit);
                default (block | log-and-permit | permit);
                engine-not-ready (block | log-and-permit | permit);
                out-of-resources (block | log-and-permit | permit);
                timeout (block | log-and-permit | permit);
                too-many-requests (block | log-and-permit | permit);
            }
            notification-options {
                fallback-block {
                    administrator-email email-address;
```

```
                        allow-email;
                        custom-message message;
                        custom-message-subject message-subject;
                        display-host;
                        (notify-mail-sender | no-notify-mail-sender);
                         type (message | protocol-only);
                    }
                    fallback-non-block {
                        custom-message message;
                        custom-message-subject message-subject;
                        (notify-mail-recipient | no-notify-mail-recipient);
                    }
                    virus-detection {
                        custom-message message;
                        custom-message-subject message-subject;
                        (notify-mail-sender | no-notify-mail-sender);
                         type (message | protocol-only);
                    }
                }
                scan-options {
                    content-size-limit value;
                    (no-uri-check | uri-check);
                    timeout value;
                }
                trickling {
                    timeout value;
                }
            }
            sxl-retry value;
            sxl-timeout seconds;
        }
        traceoptions {
            flag flag;
        }
         type (juniper-express-engine | kaspersky-lab-engine | sophos-engine);
        url-whitelist listname;
    }
    content-filtering {
        profile profile-name {
            block-command protocol-command-list;
            block-content-type (activex | exe | http-cookie | java-applet | zip);
            block-extension extension-list;
            block-mime {
```

```
                exception list-name;
                list list-name;
            }
            notification-options {
                custom-message message;
                (notify-mail-sender | no-notify-mail-sender);
                 type (message | protocol-only);
            }
            permit-command protocol-command-list;
        }
        traceoptions {
            flag flag;
        }
    }
    web-filtering {
        juniper-enhanced {
            cache {
                size value;
                timeout value;
            }
            profile profile-name {
                block-message {
                    type {
                        custom-redirect-url;
                    }
                    url url;
                    }
                quarantine-message {
                    type {
                        custom-redirect-url;
                    }
                        url url;
                        }
                category customurl-list name {
                    action  (block | log-and-permit | permit | quarantine);
                }
                custom-block-message value;
                custom-quarantine-message value;
                default  (block | log-and-permit | permit | quarantine);
                fallback-settings {
                    default  (block | log-and-permit);
                    server-connectivity   (block | log-and-permit);
                    timeout    (block | log-and-permit);
```

```
                too-many-requests  (block | log-and-permit);
            }
            no-safe-search;
            site-reputation-action {
                fairly-safe   (block | log-and-permit | permit | quarantine);
                harmful   (block | log-and-permit | permit | quarantine);
                moderately-safe  (block | log-and-permit | permit | quarantine);
                suspicious  (block | log-and-permit | permit | quarantine);
                very-safe   (block | log-and-permit | permit |quarantine);
            }
            timeout value;
        }
        server {
            host host-name;
            port number;
        }
    }
    juniper-local {
        profile profile-name {
            custom-block-message value;
            default  (block | log-and-permit | permit);
            fallback-settings {
                default (block | log-and-permit);
                server-connectivity (block | log-and-permit);
                timeout  (block | log-and-permit);
                too-many-requests  (block | log-and-permit);
            }
            timeout value;
        }
    }
    surf-control-integrated {
        cache {
            size value;
            timeout value;
        }
        profile profile-name {
            category customurl-list name {
                action  (block | log-and-permit | permit);
            }
            custom-block-message value;
            default (block | log-and-permit | permit);
            fallback-settings {
                default  (block | log-and-permit);
```

```
                        server-connectivity   (block | log-and-permit);
                        timeout    (block | log-and-permit);
                        too-many-requests  (block | log-and-permit);
                    }
                    timeout value;
                }
                server {
                    host host-name;
                    port number;
                }
            }
            traceoptions {
                flag flag;
            }
            type (juniper-enhanced | juniper-local | surf-control-integrated | websense-
redirect);
            url-blacklist listname;
            url-whitelist listname;
            websense-redirect {
                 profile profile-name {
                    account value;
                    custom-block-message value;
                    fallback-settings {
                        default  (block | log-and-permit);
                        server-connectivity  (block | log-and-permit);
                        timeout   (block | log-and-permit);
                        too-many-requests  (block | log-and-permit);
                    }
                    server {
                        host host-name;
                        port number;
                    }
                    sockets value;
                    timeout value;
                }
            }
        }
    }
    ipc {
        traceoptions flag flag;
    }
    traceoptions {
        flag flag;
```

```
    }
    utm-policy policy-name {
        anti-spam {
            smtp-profile profile-name;
        }
        anti-virus {
            ftp {
                download-profile profile-name;
                upload-profile profile-name;
            }
            http-profile profile-name;
            imap-profile profile-name;
            pop3-profile profile-name;
            smtp-profile profile-name;
        }
        content-filtering {
            ftp {
                download-profile profile-name;
                upload-profile profile-name;
            }
            http-profile profile-name;
            imap-profile profile-name;
            pop3-profile profile-name;
            rule-set rule-set-name {
            rule rule-name {
            match {
              application any; /*http, pop3, impa, smtp, ftp */
              direction  any; /* upload or download */
              file-type exe; /*predetected file types*/
            }
            then {
              action {
                no-action;            /* No action */
                /* block                Block and drop connection */
                /* close-client         Close client */
                /* close-server         Close server */
                /* close-client-and-server Close client and server */
              }
              notification {
                seclog;  /* event logging */
        endpoint { /* endpoint notification options */
             type protocol-only;
             notify-mail-sender;
```

```
            custom-message "CF Blocks content";
               }
             }
           }
         }
       }
     }
   }
           smtp-profile profile-name;
       }
       traffic-options {
           sessions-per-client {
               limit value;
               over-limit (block | log-and-permit);
           }
       }
       web-filtering {
           http-profile profile-name;
       }
   }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security]
```

Description

Configure Content Security features.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Starting in Junos OS Release 21.4R1, content filtering is performed by detecting the file content and not the file extensions. So, content filtering options based on mime-type, content-type, and protocol command is not supported. After you upgrade to Junos OS Release 21.4R1 above listed options are no more available for configuration. The rule-set and rules configurations are introduced under the `[edit security utm utm-policy <utm-policy-name> content-filtering]` hierarchy level. These rules and rule-set allows you to configure direction specific content filters and connection reset.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

The Kaspersky, surf-control-integrated, and express antivirus features are not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5 .

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Kaspersky, surf-control-integrated, and express antivirus features are not supported from Junos OS Release 15.1X49-D10 onwards.

Content Security-policy (Application Services)

IN THIS SECTION

- Syntax | 4601
- Hierarchy Level | 4601
- Description | 4601
- Options | 4601
- Required Privilege Level | 4601
- Release Information | 4601

Syntax

```
utm-policy policy-name;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
application-services]
```

Description

Configure a Content Security policy for application services and attach this policy to a security profile to implement it.

Options

policy-name—Specify the name of the Content Security policy.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

content-size

IN THIS SECTION

- Syntax | 4602
- Hierarchy Level | 4602

- Description | 4602
- Options | 4602
- Required Privilege Level | 4603
- Release Information | 4603

Syntax

```
content-size (block | log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name fallback-
options]
```

Description

If the content size exceeds a set limit, the content is either passed or blocked. The default action is log-and-permit.

> **NOTE**: When you configure the content-size value, keep in mind that in certain cases, content size is available in the protocol headers, so the max-content-size fallback is applied before a scan request is sent. However, in many cases, content size is not provided in the protocol headers. In these cases, the TCP payload is sent to the antivirus scanner and accumulates until the end of the payload. If the accumulated payload exceeds the maximum content size value, then max-content-size fallback is applied. The default fallback action is log and permit, so you may want to change this option to block, in which case such a packet is dropped and a block message is sent to the client.

Options

- `block`—Log the error and deny the traffic

- `log-and-permit`—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

content-size-limit

IN THIS SECTION

- Syntax | 4603
- Hierarchy Level | 4603
- Description | 4604
- Required Privilege Level | 4604
- Release Information | 4604

Syntax

```
content-size-limit value;
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name scan-
options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name scan-
```

```
options]
[edit security utm default-configuration anti-virus scan-options]
```

Description

The content size check occurs before the scan request is sent. The content size refers to accumulated TCP payload size. The maximum configurable content size varies with different platforms. For example, the content size ranges from 20 through 40,000 for SRX4100.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for Avira engine added in Junos OS Release 18.4R1.

content-size (Security Antivirus Sophos Engine)

IN THIS SECTION

- Syntax | 4605
- Hierarchy Level | 4605
- Description | 4605
- Options | 4605
- Required Privilege Level | 4605
- Release Information | 4606

Syntax

```
content-size (block | log-and-permit | permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name fallback-
options]
```

Description

If the content size exceeds a set limit, the content is either passed or blocked.

> **NOTE**: When you configure the content-size value, keep in mind that in certain cases, content size is available in the protocol headers, so the max-content-size fallback is applied before a scan request is sent. However, in many cases, content size is not provided in the protocol headers. In these cases, the TCP payload is sent to the antivirus scanner and accumulates until the end of the payload. If the accumulated payload exceeds the maximum content size value, then max-content-size fallback is applied. You might want to set the fallback action to block, in which case such a packet is dropped and a block message is sent to the client.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic
- `permit`—Permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Sophos Antivirus Configuration Overview

context (Application Identification)

IN THIS SECTION

- Syntax | 4606
- Hierarchy Level | 4606
- Description | 4607
- Options | 4607
- Required Privilege Level | 4611
- Release Information | 4611

Syntax

```
context {
    context;
}
```

Hierarchy Level

[edit services application-identification application *application-name* over *protocol-type* signature *name* member *name*]

Description

Specify context for matching application running over TCP, UDP, or Layer 7.

Application identification supports custom application signatures to detect applications as they pass through the device. You can create custom application signatures for applications based on ICMP, IP protocol, IP address, and Layer 7. While configuring custom application signatures, you must specify context values that the device can use to match patterns in the application traffic.

Options

context	Specify the context type. For example, Following options are available in application signature package version 3284.
ftp-content-type	Content type of the transferred file.
ftp-file-name	Filename being transferred.
ftp-greeting-message	First line of the server banner.
ftp-load-way	File transfer way—upload or download.
ftp-method	FTP command sent.
ftp-return-content	Message of server's response.
http-filename	The name of the file being fetched or posted. Extracted if content-disposition field has a filename.
http-get-url-parsed-param-parsed	The decoded, normalized GET URL in an HTTP request along with the decoded CGI parameters (if any).
http-header-content-type	Content-type header in an HTTP transaction.
http-header-cookie	Cookie header in an HTTP transaction.
http-header-host	Host header in an HTTP transaction.
http-header-user-agent	User-agent header in an HTTP transaction.
http-post-url-parsed-param-parsed	Decoded, normalized POST URL in an HTTP request along with the decoded CGI parameters (if any).
http-post-variable-parsed	Decoded POST URL or form data variables.

http-url-parsed	Decoded, normalized URL in an HTTP request.
http-url-parsed-param-parsed	Decoded, normalized URL in an HTTP request along with the decoded CGI parameters (if any).
imap-attach-filename	Name of the file attached.
imap-attach-transfer-encoding	Encoding of the attached content.
imap-attach-type	Content type of the sent attached file
imap-auth-type	Used authentication type.
imap-content-language	Language of the message content.
imap-content-transfer-encoding	The encoding of the content
imap-content-type	Content type of the transferred file.
imap-greeting-message	Greeting message of the server
imap-method	Command sent by the client.
imap-mime-version	Version of the message body format standard used in the mail protocol.
imap-received-by-name	Receiving host name.
imap-received-from-name	Sending host name.
smtp-attach-filename	Attachment file name.
smtp-attach-transfer-encoding	Encoding of the attached content.
smtp-attach-type	Content type of the sent attached file.
smtp-content-language	Language of the message content.
smtp-content-transfer-encoding	Encoding of the content
smtp-content-type	Content type of transferred file
smtp-greeting-message	Greeting message of the server
smtp-method	Command sent by the client.

smtp-mime-version	Version of the message body format standard.
smtp-received-by-name	Name of the receiving host.
smtp-received-from-name	Name of the sending host.
smtp-server	The SMTP server name
ssl-common-name	Domain name in the certificate.
ssl-issuer	Certificate Authority.
ssl-organization-name	Organisation name in the certificate.
ssl-protocol-version	SSL/TLS protocol version chosen by the server.
ssl-server-name	Server name in TLS server name extension or SSL server certificate.
ssl-subject-alt-name	Subject Alternative Name (SAN). An SAN certificate allows users to specify multiple domain names and IP addresses in a single certificate. Custom application signatures can detect applications based on the application's host names available in the SAN field of the SSL certificate.
ssl-version	SSL major version in the handshake.
ssl-server-name	Server name in the TLS server name extension or the SSL server certificate. This is also known as Server Name Indication (SNI).
stream	TCP or UDP stream data.

Examples of context types with direction. When configuring custom application signatures, the context-direction combinations as mentioned in Table 25 on page 4609 is supported. Any other combination other than this is not supported.

Table 25: Supported Context-Direction Combination for Custom Application Signatures

Context	Direction
http-get-url-parsed-param-parsed	client-to-server
http-header-host	client-to-server

Table 25: Supported Context-Direction Combination for Custom Application Signatures *(Continued)*

Context	Direction
http-header-user-agent	client-to-server
http-post-url-parsed-param-parsed	client-to-server
http-post-variable-parsed	client-to-server
http-url-parsed	client-to-server
http-url-parsed-param-parsed	client-to-server
ssl-server-name	client-to-server
stream	any/client-to-server/server-to-client
http-header-content-type	any/client-to-server/server-to-client
http-header-cookie	any/client-to-server/server-to-client

NOTE: If you are planning to upgrade the device to Junos OS release 15.1X49-D60 from the previous versions of the Junos OS, you must change the configuration to the valid combination of context-direction as mentioned in Table 25 on page 4609 to avoid any commit failure and possible disabling of the secondary node.

NOTE: If the MX Series router is running Next Gen Services, then the following restrictions apply:

- Only the http-header context types are available at the `[edit services application-identification application application-name over http signature l4-l7-signature-name member member-name]` hierarchy level.

- Only the `ssl-server` context type is available at the `[edit services application-identification application` *`application-name`* `over ssl signature` *`l4-l7-signature-name`* `member` *`member-name`*`]` hierarchy level.
- Only the `stream` context type is available at the `[edit services application-identification application` *`application-name`* `over (tcp | udp) signature` *`l4-l7-signature-name`* `member` *`member-name`*`]` hierarchy level.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D40.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Understanding Junos OS Application Identification Custom Application Signatures

context-identifier (Protocols IS-IS)

IN THIS SECTION

- Syntax | 4612
- Hierarchy Level | 4612
- Description | 4612
- Options | 4612
- Required Privilege Level | 4612
- Release Information | 4612

Syntax

```
context-identifier identifier;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit protocols isis]
```

Description

Configure IS-IS context identifier information.

Options

identifier IPv4 address that defines a protection pair. The context identifier is manually configured on both primary and protector PEs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

show isis context-identifier | 21720

context-identifier (Protocols OSPF)

IN THIS SECTION

- Syntax | 4613
- Hierarchy Level | 4613
- Description | 4613
- Options | 4613
- Required Privilege Level | 4613
- Release Information | 4614

Syntax

```
context-identifer identifier
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf |ospf3) area area-id],
[edit protocols (ospf | ospf3) area area-id ]
```

Description

Configure OSPF context-identifier information.

Options

identifer IPv4 address that defines a protection pair. The context identifier is manually configured on both the primary and protector provider edge (PE) devices.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

show ospf context-identifier | 22776

context (Security Custom Attack)

IN THIS SECTION

- Syntax | 4614
- Hierarchy Level | 4614
- Description | 4615
- Options | 4615
- Required Privilege Level | 4615
- Release Information | 4615

Syntax

```
context context-name;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature]
```

Description

Define the location of the signature where IDP should look for the attack in a specific Application Layer protocol.

Options

context-name—Name of the context under which the attack has to be matched.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

context (SNMPv3)

IN THIS SECTION

- Syntax | 4615
- Hierarchy Level | 4616
- Description | 4616
- Required Privilege Level | 4616
- Release Information | 4616

Syntax

```
context context-name;
```

Hierarchy Level

```
[edit snmp v3 snmp-community community-index]
```

Description

Specify the SNMPv3 context for access control. A context identifies a collection of information accessible for an SNMP entity.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the SNMPv3 Community

continuity-check

IN THIS SECTION

- Syntax | 4617
- Hierarchy Level | 4617
- Description | 4617
- Options | 4617
- Required Privilege Level | 4617
- Release Information | 4618

Syntax

```
continuity-check {
    convey-loss-threshold;
    hold-interval minutes;
    interface-status-tlv;
    interval (10m | 10s | 1m | 1s| 100ms | 10ms| 3.3ms);
    loss-threshold number;
    port-status-tlv;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name]
```

Description

Specify continuity check protocol options.

> **NOTE**: When the continuity-check interval is set to 3.3ms, the actual interval is 4ms (250 pps).

Options

`convey-loss-threshold`—Enable loss-threshold-tlv transmission.

`hold-interval` *minutes*—Specify the continuity check hold-interval, in minutes.

`interface-status-tlv`—Enable interface-status-tlv transmission.

`interval` *(10m | 10s | 1m | 1s| 100ms | 10ms | 3.3ms)*—Specify the continuity check interval.

`loss-threshold` *minutes*—Specify the loss-threshold, in minutes.

`port-status-tlv`—Enable port-status-tlv transmission.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring Continuity Check Protocol Parameters for Fault Detection

continuity-check (EX Series Switch Only)

IN THIS SECTION

- Syntax | 4618
- Hierarchy Level | 4619
- Description | 4619
- Options | 4619
- Required Privilege Level | 4619
- Release Information | 4619

Syntax

```
continuity-check {
    hold-interval minutes;
    interface-status-tlv;
    interval (10m | 10s | 1m | 1s| 100ms);
    loss-threshold number;
    port-status-tlv;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name]
```

Description

Specify continuity check protocol options.

The remaining statements are explained separately. See CLI Explorer.

Options

interface-status-tlv Includes interface status TLV in CCM.

port-status-tlv Includes port status TLV in CCM.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)

Junos OS Network Interfaces Configuration Guide

control-channel

IN THIS SECTION

- Syntax | 4620
- Hierarchy Level | 4620
- Description | 4620
- Options | 4620
- Required Privilege Level | 4621
- Release Information | 4621

Syntax

```
control-channel channel-name {
    vlan vlan-id;
    interface name interface-name
}
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring name (east-interface | west-interface)]
```

Description

Configure the Ethernet RPS control channel logical interface to carry the RAPS PDU. The related physical interface is the physical ring port.

Options

`vlan` *`vlan-id`*—If the control channel logical interface is a trunk port, then a dedicated `vlan` *`vlan-id`* defines the dedicated VLAN channel to carry the RAPS traffic. Only configure the *`vlan-id`* when the control channel logical interface is the trunk port.

`interface name` *`interface-name`*—Interface name of the control channel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Example: Configuring Ethernet Ring Protection Switching on EX Series Switches

Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS

Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

control-channel (Protocols Link Management Peer)

IN THIS SECTION

- Syntax | 4622
- Hierarchy Level | 4622
- Description | 4622
- Options | 4622
- Required Privilege Level | 4622
- Release Information | 4622

Syntax

```
control-channel control-channel-interface;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name],
[edit protocols link-management peer peer-name]
```

Description

Specify the control channel interface for the peer.

Options

control-channel-interface—Name of the control channel interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring LMP Peers

control-channel (Protocols OAM)

IN THIS SECTION

- Syntax | 4623
- Hierarchy Level | 4623
- Description | 4623
- Options | 4624
- Required Privilege Level | 4624
- Release Information | 4624

Syntax

```
control-channel {
    pwe3-control-word;
    pw-label-ttl-1;
    router-alert-label;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn oam],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
oam],
[edit routing-instances routing-instance-name protocols l2vpn oam],
[edit routing-instances routing-instance-name protocols vpls oam]
```

Description

Configure the Virtual Circuit Connection Verification (VCCV) BFD control channel. VCCV provides a control channel associated with a pseudowire. You can configure a number of different CV types for this control channel, based on the configuration of the pseudowire.

Options

`pwe3-control-word`—For BGP-based pseudowires that send OAM packets with a control word that has 0001b as the first nibble.

`pw-label-ttl-1`—For BGP-based pseudowires that send OAM packets with the MPLS pseudowire label and time-to-live (TTL) set to 1.

`router-alert-label`—For BGP-based pseudowires that send OAM packets with router alert label.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring BFD for VCCV for Layer 2 VPNs, Layer 2 Circuits, and VPLS

control-connection (Junos OS)

IN THIS SECTION

- Syntax | 4625
- Hierarchy Level | 4625
- Description | 4626
- Options | 4626
- Required Privilege Level | 4626
- Release Information | 4626

Syntax

```
control-connection control-client-name {
    authentication-mode
    control-type (managed | light);
    destination-interface interface-name;
    destination-port port;
    history-size size;
    moving-average-size number;
    persistent-results
    routing-instance instance-name;
    target-address (url url | address);
    tcp-keepcnt
    tcp-keepidle
    tcp-keepintvl
    test-interval interval;
    traps traps;
    data-fill-with zeros data;
    data-size size;
    dscp-code-points dscp-bits;
    probe-count count;
    probe-interval seconds;
    thresholds thresholds;
    test-session session-name{
        data-fill-with zeros data;
        data-size size;
        dscp-code-points dscp-bits;
        destination-port   destination-port;
        probe-count count;
        probe-interval seconds;
        source-address source-address;
        target-address (url url |  address) local-link IPv6-link-local-interface-name;
    }
}
```

Hierarchy Level

```
[edit services rpm twamp client]
```

Description

List all the TWAMP control clients that can connect to this server. You must configure at least one client to enable TWAMP.

Options

control-client-name	Name of the control client.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

Statement introduced in Junos OS Evolved Release 20.3R1.

`control-type` option introduced in Junos OS Release 21.1R1 and Junos OS Evolved Release 20.3R1.

`source-address` option and the `local-link` sub-option of the `target-address` option for TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

control-connection (Junos OS Evolved) | **4627**

control-connection (Junos OS Evolved)

IN THIS SECTION

- Syntax | 4627
- Hierarchy Level | 4628
- Description | 4628
- Options | 4629
- Required Privilege Level | 4630
- Release Information | 4631

Syntax

```
control-connection  control-client-name {
    control-type (managed | light);
    destination-port destination-port;
    routing-instance routing-instance-name;
    source-address source-address;
    target target-address;
    test-start (auto | manual);
    test-interval seconds;
    traps {
          control-connection-closed;
          test-iteration-done;
    }
    test-session name {
        data-size data-size;
        destination-port destination-port;
        dscp-code-points dscp-code-points;
        history-size    history-size;
         moving-average-size   moving-average-size;
        offload-type (none | inline-timestamping | pfe-timestamp);
        probe-count  probe-count;
        probe-interval  seconds;
        source-address source-address;
        source-routing {
            forward (srmpls-nodes <label hop1...n> | srv6-nodes <IPv6 SID hop1...n>);
```

```
            return (srmpls-nodes <label hop1...n> | srv6-nodes <IPv6 SID hop1...n>);
        }
        target target-address local-link IPv6-link-local-interface-name;
        thresholds {
            control-failure (on | off);
            successive-loss number;
            total-loss number;
            threshold-type (microseconds | average);
        }
        traps {
                    egress-jitter-exceeded
                    egress-time-exceeded;
                    ingress-jitter-exceeded;
                    ingress-time-exceeded;
                    jitter-exceeded;
                    probe-failure;
                    rtt-exceeded;
                    test-completion;
                    test-failure;
        }
        ttl hop-count;
    }
}
```

Hierarchy Level

```
[edit services monitoring twamp client]
```

Description

For Junos OS Evolved, configure all the Two-Way Active Measurement Protocol (TWAMP) control clients that can connect to the server. A TWAMP control connection is responsible for initiating, starting, and ending the test sessions between a TWAMP client and a TWAMP server for performance measurement. However, when the control-type option is set to `light`, the test session parameters are predefined and not negotiated, and no control session is created. Therefore, with TWAMP light, UDP probe generation and reception, as well as reflection, is not part of a control session.

Options

control-client-name

Name of the control client.

control-type (managed | light)

Specify how you want to manage the control connection.

- **Values:** Configure one of the following:
  - managed—Specify that you want to have a stateful version of TWAMP where test parameters are negotiated. You can specify particular source or target addresses, or non-default destination ports for the control connections. The ports are negotiated and agreed upon between the client and server. All test packets received by the server on a destination port are reflected back and forgotten right away.
  - light—Specify that you want a stateless version of TWAMP where test parameters are predefined instead of negotiated. All test packets received by the server on a destination port are reflected back and forgotten right away.
- **Default:** managed

destination-port *port*

Specify the Transmission Control Protocol (TCP) port number for the endpoint of the control connection. You must set the control-type option to `managed` to set the destination port for the entire control connection. Otherwise, if the control-type option is set to `light`, you will have to configure the destination-port option on each test session.

- **Range:** You can specify port 862, or any port from 49152 through 65535.
- **Default:** 862 (IANA port for TWAMP)

routing-instance *instance-name*

Specify the routing instance used by the probes. This routing instance is configured at the [edit routing-instance] hierarchy level.

> **NOTE**: The media interface from where the TWAMP control and test or data packets arrive and exit must be part of the same routing instance.

- **Default:** No routing instance is used. The Internet (IPv4) routing table inet.0 is used instead.

source-address *address*

Specify the IPv4 or IPv6 source address used for control connections. If the source address is not one of the device's assigned addresses, the control connection uses the outgoing interface's address as its source.

You cannot use IPv6 link-local addresses or the following addresses as the source address for control connections:

- 0.0.0.0
- 127.0.0.0/8 (loopback)
- 224.0.0.0/4 (multicast)
- 255.255.255.255 (broadcast)

target *target-address*

Specify the IPv4 or IPv6 address for the target destination for the control connection. You cannot use IPv6 link-local addresses as the target destination for the control connection.

> **NOTE**: You can configure this option only when the control-type option is set or defaulted to managed. The target option is required for a managed control connection.

test-start (auto | manual)

Specify whether or not you want to manually start TWAMP tests.

- **Values:** Configure one of the following:
  - auto—Start test sessions as soon as you commit the configuration.
  - manual—Wait to start test sessions until you issue the `request services monitoring twamp client start` operational mode command.
- **Default:** auto

test-interval *seconds*

Specify the time to wait between tests.

- **Range:** 1 through 86400 seconds
- **Default:** 1

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.3R1.

`traps` option introduced in Junos OS Evolved Release 21.3R1.

IPv6 address support introduced in Junos OS Evolved Release 21.4R1.

The `local-link` sub-option of the `target` option for TWAMP Light test sessions introduced in Junos OS Evolved Release 22.3R1.

The `inline-timestamping` suboption of the `offload-type` option introduced in Junos OS Evolved Release 22.4R1.

`source-routing` option introduced in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

control-connection (Junos OS) | **4624**

request services monitoring twamp client | **14953**

control-cores

IN THIS SECTION

- Syntax | **4632**
- Hierarchy Level | **4632**
- Description | **4632**
- Options | **4632**
- Required Privilege Level | **4632**
- Release Information | **4632**

Syntax

```
control-cores control-number;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-
provider]
```

Description

Configure control cores. Any cores not configured as either control or data cores are treated as user cores. When the number of control cores is changed, the PIC reboots.

Options

control-number—Number of control cores. At least one core must be a control core.

- **Range:** 1 through 8

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

data-cores | 4827

control-link-recovery

IN THIS SECTION

- Syntax | 4633
- Hierarchy Level | 4633
- Description | 4633
- Required Privilege Level | 4634
- Release Information | 4634

Syntax

```
control-link-recovery;
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Enable control link recovery to be done automatically by the system.

If the control link fails, the secondary node enters to ineligible and then to disabled state. If the control link recovers with in three consecutive responses, the device auto reboots irrespective of the state (ineligible or disable) of the node. You can reboot automatically using the `control-link-recovery` configuration option.

After the control link recovers, the system checks whether it receives at least three consecutive heartbeats on the control link. This is to ensure that the control link is not flapping and is perfectly healthy. Once this criterion is met, the system issues an automatic reboot on the node that was disabled when the control link failed. When the disabled node reboots, the node rejoins the cluster. There is no need for any manual intervention.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

interface (Chassis Cluster) | 8850

Configuring Cluster Failover Parameters

control-plane-bandwidth-percent (Node Slicing)

IN THIS SECTION

- Syntax | 4634
- Hierarchy Level | 4635
- Description | 4635
- Options | 4635
- Required Privilege Level | 4635
- Release Information | 4635

Syntax

```
control-plane-bandwidth-percent percent;
```

Hierarchy Level

```
[edit chassis network-slices guest-network-functions gnf]
```

Description

Allocate a percentage of the bandwidth that exists on the control plane on the router to the specified guest network function (GNF). Allocating bandwidth prevents potential overutilization by one GNF over another.

Options

percent	Percentage of control plane bandwidth.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

network-slices | **11322**

guest-network-functions | **7949**

gnf | **7746**

description (GNF) | **5332**

fpcs (Node Slicing) | **7536**

control-polarity

IN THIS SECTION

- Syntax | 4636
- Hierarchy Level | 4636
- Description | 4636
- Options | 4636
- Required Privilege Level | 4636
- Release Information | 4637

Syntax

```
control-polarity (negative | positive);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

For X.21 interfaces only, configure the control signal polarity.

Options

`positive`—Positive signal polarity.

`negative`—Negative signal polarity.

- **Default:** `positive`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Serial Signal Polarities

control-ports

IN THIS SECTION

- Syntax | 4637
- Hierarchy Level | 4637
- Description | 4638
- Options | 4638
- Required Privilege Level | 4638
- Release Information | 4638

Syntax

```
control-ports   fpc {
    port;
}
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Enable the specific control port to use as a control link for the chassis cluster. By default, all control ports are disabled.

After connecting the control ports, you need to configure the control ports to setup control links as a step to establish the chassis cluster.

You need to configure a minimum of one control port per chassis of the cluster. If you configure port 0 only, the Juniper Services Redundancy Protocol process (jsrpd) does not send control heartbeats on control link 1 and the counters it sends will show zeroes.

Options

- fpc *slot-number* —Flexible PIC Concentrator (FPC) slot number.

> **NOTE**: FPC slot range depends on platform. The maximum range of 0 through 23 applies to SRX5800 devices; for SRX5600 devices, the only applicable range is 0 through 11; for SRX5400 devices, the applicable slot range is 0 through 5. See *Chassis Cluster Control Plane Interfaces* for details.

- port *port-number* —Port number on which to configure the control port.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. Support.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

control (REST API)

IN THIS SECTION

- Syntax | 4639
- Hierarchy Level | 4639
- Description | 4639
- Required Privilege Level | 4639
- Release Information | 4640

Syntax

```
control {
    allowed-sources [ value-list ];
    connection-limit limit;
}
```

Hierarchy Level

```
[edit system services rest]
```

Description

Specify the allowed source IP addresses and maximum number of simultaneous connections for the REST API process.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

control-signal

IN THIS SECTION

- Syntax | 4640
- Hierarchy Level | 4640
- Description | 4640
- Options | 4640
- Required Privilege Level | 4641
- Release Information | 4641

Syntax

```
control-signal (assert | de-assert | normal);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options dce-options],
[edit interfaces interface-name serial-options dte-options]
```

Description

For X.21 interfaces only, configure the to-DCE signal.

Options

`assert`—The to-DCE signal must be asserted.

`de-assert`—The to-DCE signal must be deasserted.

`normal`—Normal request-to-send (RTS) signal handling, as defined by ITU-T Recommendation X.21.

- **Default:** normal

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring the Serial Signal Handling*

control-source

IN THIS SECTION

- Syntax | 4641
- Hierarchy Level | 4642
- Description | 4642
- Options | 4642
- Required Privilege Level | 4642
- Release Information | 4642

Syntax

```
control-source identifier {
    allowed-destinations [ destinations ];
```

```
    minimum-priority value;
    no-syslog;
    notification-targets address port port-number;
    service-port port-number;
    shared-key value;
    source-addresses [ addresses ];
}
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Identify the source of the dynamic flow capture request.

Options

identifier—Name of control source.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Control Source

control-vlan

IN THIS SECTION

- Syntax | 4643
- Hierarchy Level | 4643
- Description | 4643
- Required Privilege Level | 4643
- Release Information | 4644

Syntax

```
control-vlan  (vlan-id | vlan-name)
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring]
```

```
[edit protocols protection-group ethernet-ring name (east-interface |west-interface)]
```

Description

Specify the VLAN that carries the protocol data units (PDUs) between the nodes in the protected Ethernet ring. This is a control VLAN, meaning that it carries data for one instance of an Ethernet ring protection switching (ERPS) in the control channel. Use a control VLAN on trunk port interfaces. One control channel can contain multiple control VLANs.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

RELATED DOCUMENTATION

Example: Configuring Ethernet Ring Protection Switching on EX Series Switches

Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS

Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

control-word

IN THIS SECTION

- Syntax | 4644
- Hierarchy Level | 4644
- Description | 4645
- Required Privilege Level | 4646
- Release Information | 4646

Syntax

```
control-word;
```

Hierarchy Level

```
edit logical-systems logical-system-name routing-instances routing-instance-name protocols evpn],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name neighbor address static],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor address static],
[edit routing-instances routing-instance-name protocols evpn],
[edit routing-instances routing-instance-name protocols vpls],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name neighbor
address static],
[edit routing-instances routing-instance-name protocols vpls neighbor address static]
```

Description

Set `control-word` to request that other routers insert a control word between the label stack and the MPLS payload.

Enabling control-word prevents provider edge routers from incorrectly identifying a VPLS payload as an IPv4 or IPv6 payload. This prevents out-of-order packet delivery in a VPLS network that is configured to load-balance VPLS traffic across multiple paths.

The `control-word` status for static pseudowires in an VPLS instance depends on the VPLS instance and mesh-group configurations. Table 26 on page 4645 displays the `control-word` status based on the `control-word` configuration in the VPLS instance and static pseudowire mesh-group.

Table 26: Status of control-word for static pseudowire with respect to control-word at VPLS instance

Control-word Configuration at VPLS Instance	Control-word Configuration at Static Pseudowire in Mesh-group	Control-word Enabled or Disabled for Static Pseudowire
default	default	disabled
default	control-word	enabled
default	no-control-word	disabled
control-word	default	enabled
control-word	control-word	enabled

Table 26: Status of control-word for static pseudowire with respect to control-word at VPLS instance *(Continued)*

Control-word Configuration at VPLS Instance	Control-word Configuration at Static Pseudowire in Mesh-group	Control-word Enabled or Disabled for Static Pseudowire
control-word	no-control-word	disabled
no-control-word	default	disabled
no-control-word	control-word	enabled
no-control-word	no-control-word	disabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced for static pseudowires in mesh-group and in VPLS instance in Junos OS Release 20.3R1.

Statement introduced for the MPC-10E line card on the MX series router in Junos OS Release 21.3R1 and Junos OS Release 21.2R3-S3.

RELATED DOCUMENTATION

Control Word for BGP VPLS Overview

Configuring a Control Word for BGP VPLS

no-control-word | **11426**

control-word (Protocols Layer 2 Circuit Neighbor)

IN THIS SECTION

- Syntax | 4647
- Hierarchy Level | 4647
- Description | 4647
- Options | 4647
- Required Privilege Level | 4648
- Release Information | 4648

Syntax

```
(control-word | no-control-word);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address
interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name]
```

Description

Specify the control word. The control word is four bytes long and is inserted between the Layer 2 protocol data unit (PDU) being transported and the virtual circuit (VC) label that is used for demultiplexing.

Options

`control-word`—Enable the use of the control word.

- **Default:** A null control word is enabled by default. You can also configure the control word explicitly using the `control-word` statement.

`no-control-word`—Disable the use of the control word.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Control Word for Frame Relay Interfaces

control-word (Protocols Layer 2 VPN)

IN THIS SECTION

- Syntax | 4648
- Hierarchy Level | 4649
- Description | 4649
- Default | 4649
- Required Privilege Level | 4649
- Release Information | 4649

Syntax

```
control-word;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn],
[edit routing-instances routing-instance-name protocols l2vpn]
```

Description

Specify the control word. The control word is 4 bytes long and is inserted between the Layer 2 protocol data unit (PDU) being transported and the virtual connection (VC) label that is used for demultiplexing.

NOTE: The following configuration statements are ignored for time-division multiplexing pseudowires at the `[edit protocols l2vpn]` hierarchy level:

- control-word
- no-control-word

Default

The control word is enabled by default. You can also configure the control word explicitly using the `control-word` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Disabling the Control Word for Layer 2 VPNs

no-control-word (Protocols Layer 2 VPN) | **11427**

controller

IN THIS SECTION

- Syntax | 4650
- Hierarchy Level | 4650
- Description | 4650
- Options | 4652
- Required Privilege Level | 4652
- Release Information | 4652

Syntax

```
controller(p4 | re);
```

Hierarchy Level

```
[edit services inline-monitoring instance]
```

Description

PTX10004, PTX10008, and PTX10016 devices running Junos OS Evolved Release 22.1R1 can provide an SDN-based backbone (data center interconnect) for target networks. It works by using a firewall filter to redirect matching packets, via a P4Runtime agent running in Junos OS Evolved on the Routing Engine, to a P4 controller that is also running on the PTX router. (The P4Runtime agent registers with the Juniper Extension Toolkit (JET) services daemon (JSD) to open the gRPC connections and listen for P4 requests from clients.)

You can match:

- IPv4, IPv6, UDP, and TCP protocol packets according to the destination IP address

- Google Discovery protocol packets (matched by specifying VLAN ID: 4000, EtherType: 0x6007)
- traceroute redirect packets (matched by specifying TTL=0 for IPv4 and TTL=1 for IPv6)

The P4Runtime agent supports packet I/O from network devices to the SDN controller, as well as OpenConfig for switch configuration, and gRIBI for route programming.

To configure inline-monitoring services for packet redirects to the P4 controller, you need to configure inline-monitoring, create an instance, and set the instance type to controller P4, as shown here:

```
{master}
[edit services inline-monitoring]
instance {
    Instance-1 {
        controller p4;
    }
}
```

And configure a firewall filter **action** to redirect matching packets to your *instance*, as shown here:

```
{master}
[edit firewall family (any | inet | inet6) filter f1]
term t1 {
    then redirect Instance-1;
}
```

Optionally, you may also want to modify the DDoS protocol parameters:

```
{master}
[edit system ddos-protection protocols custom]
aggregate {
    bandwidth value;
    burst value;
}
```

or disable DDoS, as the case may be:

```
{master}
[edit system ddos-protection protocols custom]
aggregate {
```

```
    disable-fpc;
}
```

Controller type:

Options

p4 Select this option to have the P4Runtime agent send packets to the P4 controller (instead of the default Routing Engine).

re Select this option to send packets to the default Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 22.1R1.

RELATED DOCUMENTATION

Firewall Filter Match Conditions and Actions (PTX Series Routers)

controller (OVSDB)

IN THIS SECTION

- Syntax | 4653
- Hierarchy Level | 4653
- Description | 4653
- Options | 4654
- Required Privilege Level | 4654

Release Information | 4654

Syntax

```
controller ip-address {
    inactivity-probe-duration milliseconds;
    maximum-backoff-duration milliseconds;
    protocol  protocol {
        port number;
    }
}
```

Hierarchy Level

```
[edit protocols ovsdb]
```

Description

Configure a connection between a Juniper Networks device running the *Open vSwitch Database (OVSDB)* management protocol and a software-defined networking (SDN) controller. You can connect a Juniper Networks device to more than one SDN controller for redundancy.

In a VMware NSX environment, one cluster of NSX controllers typically includes three or five controllers. To implement the OVSDB management protocol on a Juniper Networks device, you must explicitly configure a connection to one NSX controller, using the Junos OS CLI. If the NSX controller to which you explicitly configure a connection is in a cluster, the controller pushes information about other controllers in the same cluster to the device, and the device establishes connections with the other controllers. However, you can also explicitly configure connections with the other controllers in the cluster, using the Junos OS CLI.

To implement the OVSDB management protocol on a Juniper Networks device in a Contrail environment, you must configure a connection to a Contrail controller, using the Junos OS CLI.

Connections to all SDN controllers are made on the management interface of the Juniper Networks device.

Options

ip-address—IPv4 address of the SDN controller.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Setting Up OVSDB on Juniper Networks Devices That Support the Dynamic Configuration of VXLANs

Setting Up the OVSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs

Understanding How to Set Up OVSDB Connections on a Juniper Networks Device

controller (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 4655
- Hierarchy Level | 4655
- Description | 4655
- Options | 4655
- Required Privilege Level | 4655
- Release Information | 4655

Syntax

```
controller {
    address address;
    id id;
    protocol tcp {
        port port;
    }
    role equal;
}
```

Hierarchy Level

```
[edit protocols openflow switch switch-name]
```

Description

Configure the OpenFlow controller connection information for a virtual switch on an OpenFlow-enabled device running Junos OS. If you configure a virtual switch with a single controller, by default, the controller is in active mode, and the switch automatically initiates a connection to the controller.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Support for OpenFlow on Devices Running Junos OS

Understanding the Virtual Switch Connection to the OpenFlow Controller on Devices Running Junos OS

address (Protocols OpenFlow) | **1912**

protocol (Protocols OpenFlow) | **13455**

role (Protocols OpenFlow) | **15829**

switch (Protocols OpenFlow) | **28755**

convert-clock-class-to-quality-level

IN THIS SECTION

- Syntax | **4656**
- Hierarchy Level | **4656**
- Description | **4656**
- Required Privilege Level | **4657**
- Release Information | **4657**

Syntax

```
convert-clock-class-to-quality-level;
```

Hierarchy Level

```
[edit protocols ptp slave]
```

Description

Configure the timeReceiver to enable it to retrieve Ethernet Synchronization Message Channel (ESMC) information from the Precision Time Protocol (PTP) clock class.

When this option is set, the outgoing quality level depends on the PTP clock class mapping, irrespective of the clock being configured in hybrid mode or pure PTP mode. This is the default mapping mode of the ESMC quality level value to the clock class.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Hybrid Mode and ESMC Quality Level Mapping

Example: Configuring Hybrid Mode and ESMC Quality Level Mapping

Precision Time Protocol Overview

Synchronous Ethernet Overview

Understanding ESMC Quality Level Mapping

Understanding Hybrid Mode

convert-sa-to-msdp

IN THIS SECTION

- Syntax | 4658
- Hierarchy Level | 4658
- Description | 4658
- Default | 4658
- Required Privilege Level | 4658
- Release Information | 4658

Syntax

```
convert-sa-to-msdp;
```

Hierarchy Level

```
[edit routing-instances instance-name protocols mvpn mvpn-mode spt-only]
```

Description

Enable conversion of MVPN Type 5 routes to MSDP Source Active (SA) routes.

Default

Conversion of MVPN Type 5 routes to MSDP SA routes is disabled.

MVPN Type 5 route to MSDP SA conversion works in `spt-only` mode. `rpt-spt` mode is not supported.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1.

convey-loss-threshold

IN THIS SECTION

- Syntax | 4659
- Hierarchy Level | 4659

Description | 4659
Required Privilege Level | 4659

Syntax

```
convey-loss-threshold;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name continuity-check]
```

Description

Enable loss-threshold-tlv transmission.

During a unified ISSU, the control plane may go down for several seconds and cause CFM continuity check packets to get dropped. This may cause the remote maintenance endpoint (MEP) to detect a connectivity loss and mark the MEP as down. To keep the MEP active during a unified ISSU, the loss threshold TLV communicates the minimum threshold value the receiving MEP requires to keep the MEP active. The receiving MEP parses the TLV and updates the loss threshold value, but only if the new threshold value is greater than the locally configured threshold value. You can control the transmission of the loss threshold TLV in continuity check messages PDUs. The `convey-loss-threshold` statement specifies that the loss threshold TLV must be transmitted as part of the continuity check messages. If the statement is not specified, continuity check messages transmit this TLV only when a unified ISSU is in progress.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

copy

IN THIS SECTION

- Syntax | 4660
- Description | 4660
- Options | 4660
- Required Privilege Level | 4660
- Release Information | 4660

Syntax

```
copy existing-statement to new-statement
```

Description

Make a copy of an existing statement in the configuration.

Options

existing-statement	Statement to copy.
new-statement	Copy of the statement.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Copying a Statement in the Configuration

copy-dont-fragment-bit (Services IPsec VPN)

IN THIS SECTION

- Syntax | 4661
- Hierarchy Level | 4661
- Description | 4661
- Required Privilege Level | 4662
- Release Information | 4662

Syntax

```
copy-dont-fragment-bit;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Copy the do not fragment (DF) bit value to only the outer header and not modify the inner header of the IPsec packet. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs. These settings apply for static endpoint tunnels and not for dynamic tunnels, for which you need to include the `copy-dont-fragment-bit` statement at the `[edit services service-set service-set-name ipsec-vpn-options]` hierarchy level to copy the DF bit value to only the outer header of the packet in a static IPsec tunnel. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs.

By default, this statement is disabled on MS-MICs and MS-MPCs (the DF bit value is not copied to the outer header by default).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring IPsec Rules

copy-dont-fragment-bit (Services Set)

IN THIS SECTION

- Syntax | 4662
- Hierarchy Level | 4663
- Description | 4663
- Required Privilege Level | 4663
- Release Information | 4663

Syntax

```
copy-dont-fragment-bit;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Copy the do not fragment (DF) bit value to only the outer header and not modify the inner header of the IPsec packet in dynamic endpoint tunnels. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs. These settings apply for dynamic endpoint tunnels and not for static tunnels, for which you need to include the `copy-dont-fragment-bit` statement at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level to copy the DF bit value to only the outer header of the packet in a static IPsec tunnel. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs.

By default, this statement is disabled on MS-MICs and MS-MPCs (the DF bit value is not copied to the outer header by default).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring IPsec Service Sets

Configuring IPsec Rules

copy-plp-all

IN THIS SECTION

- Syntax | 4664
- Hierarchy Level | 4664
- Description | 4664
- Required Privilege Level | 4664
- Release Information | 4664

Syntax

```
copy-plp-all;
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Enable PLP bit copying for ingress and egress for unicast and multicast traffic when traffic is ingressing one FPC and egressing the other (from E3-FPC to non-E3 FPC on M320 routers, or from ES-FPC to non-ES FPC on T Series routers).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.3.

RELATED DOCUMENTATION

Managing Congestion by Setting Packet Loss Priority for Different Traffic Flows

copy-tos-to-outer

IN THIS SECTION

- Syntax | 4665
- Hierarchy Level | 4665
- Description | 4665
- Required Privilege Level | 4666
- Release Information | 4666

Syntax

```
copy-tos-to-outer service-type (gre | mt);
```

Hierarchy Level

```
[edit chassis]
```

Description

For static and dynamic GRE tunnel interfaces and MT tunnel interfaces on MPCs, copy the inner IP header's ToS bits to the outer IP packet header for traffic transiting the router. This statement affects all MPC interfaces and takes precedence over the "copy-tos-to-outer-ip-header-transit" on page 4669, which applies to individual interfaces.

Once committed, this configuration statement only affects new gr- and mt- interfaces. To affected an existing gr- or mt- interface, you must delete and re-add the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

copy-tos-to-outer-ip-header-transit | 4669

copy-tos-to-outer-ip-header | 4668

copy-tos-to-outer-ip-header

IN THIS SECTION

- Syntax | 4666
- Hierarchy Level | 4667
- Description | 4667
- Default | 4667
- Required Privilege Level | 4667
- Release Information | 4667

Syntax

```
copy-tos-to-outer-ip-header;
```

Hierarchy Level

```
[edit interfaces gr-fpc/pic/port unit logical-unit-number],
[edit interfaces gre unit logical-unit-number],
[edit logical-systems logical-system-name interfaces gr-fpc/pic/port unit logical-unit-number],
[edit logical-systems logical-system-name interfaces gre unit logical-unit-number]
```

Description

For GRE tunnel interfaces and GRE interfaces for GMPLS control channels only, enable the inner IP header's ToS bits to be copied to the outer IP packet header for traffic originating in the Routing Engine. To verify that this option is enabled at the interface level, use the `show interfaces interface-name detail` command.

Default

If you omit this statement, the ToS bits in the outer IP header are set to 0.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Support for GRE interfaces for Generalized MPLS (GMPLS) introduced in Junos OS Release 12.3R7.

RELATED DOCUMENTATION

Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header

copy-tos-to-outer-ip-header

IN THIS SECTION

- Syntax | 4668
- Hierarchy Level | 4668
- Description | 4668
- Default | 4668
- Required Privilege Level | 4669
- Release Information | 4669

Syntax

```
copy-tos-to-outer-ip-header;
```

Hierarchy Level

```
[edit interfaces gr-fpc/pic/port unit logical-unit-number],
[edit interfaces gre unit logical-unit-number],
[edit logical-systems logical-system-name interfaces gr-fpc/pic/port unit logical-unit-number],
[edit logical-systems logical-system-name interfaces gre unit logical-unit-number]
```

Description

For GRE tunnel interfaces and GRE interfaces for GMPLS control channels only, enable the inner IP header's ToS bits to be copied to the outer IP packet header for traffic originating in the Routing Engine.

To verify that this option is enabled at the interface level, use the `show interfaces interface-name detail` command.

Default

If you omit this statement, the ToS bits in the outer IP header are set to 0.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Support for GRE interfaces for Generalized MPLS (GMPLS) introduced in Junos OS Release 12.3R7.

RELATED DOCUMENTATION

Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header

copy-tos-to-outer-ip-header-transit | 4669

force-control-packets-on-transit-path | 7367

copy-tos-to-outer-ip-header-transit

IN THIS SECTION

- Syntax | 4669
- Hierarchy Level | 4670
- Description | 4670
- Default | 4670
- Required Privilege Level | 4670
- Release Information | 4670

Syntax

```
copy-tos-to-outer-ip-header-transit;
```

Hierarchy Level

```
[edit interfaces gr-fpc/pic/port unit logical-unit-number]
```

Description

For static GRE tunnel interfaces on MPCs, copy the inner IP header's ToS bits to the outer IP packet header for traffic transiting the router.

Default

If you omit this statement, the ToS bits in the outer IP header are set to 0.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

copy-tos-to-outer-ip-header | 4668

force-control-packets-on-transit-path | 7367

core-dump

IN THIS SECTION

- Syntax | 4671
- Hierarchy Level | 4671

- Description | 4671
- Required Privilege Level | 4671
- Release Information | 4671

Syntax

```
(core-dump | no-core-dump);
```

Hierarchy Level

```
[edit interfaces mo-fpc/pic/port multiservice-options]
```

Description

A useful tool for isolating the cause of a problem. Core dumping is enabled by default. The directory `/var/tmp` contains core files. Junos OS saves the current core file (0) and the four previous core files, which are numbered from 1 through 4 (from newest to oldest):

> **NOTE**: By default, all members of a configured user group (with read-only permissions) can access the core dump files and attach them to cases associated with JTAC.

- `core-dump`—Enable the core dumping operation.
- `no-core-dump`—Disable the core dumping operation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Flow Monitoring

core-facing

IN THIS SECTION

- Syntax | 4672
- Hierarchy Level | 4672
- Description | 4672
- Required Privilege Level | 4672
- Release Information | 4673

Syntax

```
core-facing;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family]
```

Description

Specifies that the VLAN is physically connected to a core-facing ISP router and ensures that the network does not improperly treat the interface as a client interface. When specified, the interface is inserted into the core-facing default mesh group where traffic from pseudowires that belong to the default mesh group is not forwarded on the core-facing link.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

Support for EVPN-MPLS on PTX Series routers added in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

- Junos OS Subscriber Management and Services Library
- *Tunnel Services Overview*
- *Tunnel Interface Configuration on MX Series Routers Overview*

corouted-bidirectional

IN THIS SECTION

- Syntax | 4673
- Hierarchy Level | 4674
- Description | 4674
- Default | 4674
- Required Privilege Level | 4674
- Release Information | 4674

Syntax

```
corouted-bidirectional;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Specify that the label-switched path be established as a corouted bidirectional packet LSP. You cannot configure this statement at the same time as the `corouted-bidirectional-passive` statement.

Default

This statement is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Corouted Bidirectional LSPs

corouted-bidirectional-passive | 4674

corouted-bidirectional-passive

IN THIS SECTION

- Syntax | 4675
- Hierarchy Level | 4675

- Description | 4675
- Default | 4675
- Required Privilege Level | 4675
- Release Information | 4675

Syntax

```
corouted-bidirectional-passive;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Specify that the label-switched path be a passive LSP associated with a bidirectional LSP when it is signaled at the ingress router. This passive LSP enables the MPLS application to utilize the reverse LSP. You cannot configure this statement at the same time as the `corouted-bidirectional` statement.

Default

This statement is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Corouted Bidirectional LSPs

corouted-bidirectional | 4673

corrupt-file

IN THIS SECTION

- Syntax | 4676
- Hierarchy Level | 4676
- Description | 4676
- Options | 4677
- Required Privilege Level | 4677
- Release Information | 4677

Syntax

```
corrupt-file (block | log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name fallback-
options]
```

Description

Corrupt file is the error returned by the scan engine when engine detects a corrupted file. The default action is log-and-permit.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Full Antivirus Configuration Overview

cos-rule-sets (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 4678
- Hierarchy Level | 4678
- Description | 4678
- Options | 4678
- Required Privilege Level | 4678
- Release Information | 4678

Syntax

```
cos-rule-sets [cos-rule-set-name];
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the services CoS rule set to apply to the service set. The service set processes the rules in the order they appear in the rule set.

The service set that the CoS rule set is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

cos-rule-set-name	Name of the services CoS rule set.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Class of Service for Services PICs (Next Gen Services)

cos-rules (Service Set)

IN THIS SECTION

- Syntax | 4679
- Hierarchy Level | 4679
- Description | 4679
- Options | 4679
- Required Privilege Level | 4680
- Release Information | 4680

Syntax

```
cos-rules [cos-rule-name];
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the CoS rules to apply to the service set. You can configure multiple rules.

The service set that the CoS rule is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

cos-rule-name	CoS rule name.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring CoS Rules

cos-rules (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 4680
- Hierarchy Level | 4680
- Description | 4681
- Options | 4681
- Required Privilege Level | 4681
- Release Information | 4681

Syntax

```
cos-rules [cos-rule-name];
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the CoS rules to apply to the service set. You can configure multiple rules.

The service set that the CoS rule is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

cos-rule-name	CoS rule name.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Class of Service for Services PICs (Next Gen Services)

cost

IN THIS SECTION

- Syntax | 4682
- Hierarchy Level | 4682
- Description | 4683
- Default | 4683
- Options | 4683

- Required Privilege Level | 4683
- Release Information | 4683

Syntax

```
cost cost;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp | vstp) interface interface-name],
[edit logical-systems logical-system-name protocols mstp msti msti-id interface interface-name],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp | vstp) interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name
protocols mstp msti msti-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id interface interface-name],
```

```
[edit protocols (mstp | rstp | vstp) interface interface-name],
[edit protocols mstp msti msti-id interface interface-name],
[edit protocols vstp vlan vlan-id interface interface-name],
```

```
[edit routing-instances routing-instance-name protocols (mstp | rstp | vstp) interface interface-
name],
[edit routing-instances routing-instance-name protocols mstp msti msti-id interface interface-
name],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id interface interface-name]
```

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure the link cost to control which bridge is the designated bridge and which interface is the designated interface.

Default

The link cost is determined by the link speed.

Options

cost—(Optional) Link cost associated with the port.

- **Range:** 1 through 200,000,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Spanning-Tree Instance Interfaces

Understanding STP

Understanding RSTP

Understanding MSTP

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

count (Protocols PIM)

IN THIS SECTION

- Syntax | 4684
- Hierarchy Level | 4684
- Description | 4684
- Required Privilege Level | 4684
- Release Information | 4685

Syntax

```
count number;
```

Hierarchy Level

```
[edit protocols pim interface (Protocols PIM) interface-name multiple-triggered-joins]
```

Description

Specify the count for the number of triggered joins to be sent between PIM neighbors through the PIM interface. Optionally, you can configure the count number using the `count` statement at the `[edit protocols pim interface interface-name multiple-triggered-joins]` hierarchy level.

- **Range:** 5 through 15
- **Default:** 5

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

interface (Protocols PIM) | 8991
multiple-triggered-joins | 11119

count

IN THIS SECTION

- Syntax | 4685
- Hierarchy Level | 4685
- Description | 4686
- Options | 4686
- Required Privilege Level | 4686
- Release Information | 4686

Syntax

```
count number;
```

Hierarchy Level

```
[edit access profile profile-name client-name-filter client-name]
```

Description

Specify the number of characters to be stripped from a client name, from right to left, until the specified number of characters are deleted. The resulting name is sent to the authentication server.

Options

number—Number of characters to be stripped in a client name.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Obtaining Username and Role Information Through Firewall Authentication

count (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 4687
- Hierarchy Level | 4687
- Description | 4687
- Required Privilege Level | 4687
- Release Information | 4687

Syntax

```
count;
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number then]
```

Description

Enable the collection of statistics for HTTP header enrichment for the tag rule term. The collection of statistics for a term is disabled by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

Configuring HCM Profiles and Assigning Tag Rules

Junos Web Aware HTTP Header Enrichment Overview

show services hcm statistics | **25550**

count (Security Custom Attack)

IN THIS SECTION

- Syntax | 4688
- Hierarchy Level | 4688
- Description | 4688
- Options | 4688
- Required Privilege Level | 4688
- Release Information | 4689

Syntax

```
count count-value;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name time-binding]
```

Description

Specify the number of times that IDP detects the attack within the specified scope before triggering an event.

Options

count-value—Number of times IDP detects the attack.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

count (Security Policies)

IN THIS SECTION

- Syntax | 4689
- Hierarchy Level | 4689
- Description | 4689
- Required Privilege Level | 4689
- Release Information | 4690

Syntax

```
count;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then]
```

Description

Enable a count, in bytes or kilobytes, of all network traffic the policy allows to pass through the device in both directions: the originating traffic from the client to the server (from the from-zone to the to-zone), and the return traffic from the server to the originating client.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Statement updated in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

show security policies | 25028

Security Policies Overview

counter

IN THIS SECTION

- Syntax | 4690
- Hierarchy Level | 4690
- Description | 4691
- Options | 4691
- Required Privilege Level | 4691
- Release Information | 4691

Syntax

```
counter {
    counter-id counter-index;
}
```

Hierarchy Level

```
[edit firewall policer policer-name]
```

Description

(On EX8200 switches only) Configure a policer counter.

Options

counter-id *counter-index*—Global management counter ID.

- **Range:** 0 through 2

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Policers to Control Traffic Rates (CLI Procedure)

Configuring Firewall Filters (CLI Procedure)

Understanding the Use of Policers in Firewall Filters

counters

IN THIS SECTION

- Syntax | 4692
- Hierarchy Level | 4692
- Description | 4692
- Options | 4692
- Required Privilege Level | 4692
- Release Information | 4692

Syntax

```
counters {
    counter-name;
}
```

Hierarchy Level

```
[edit accounting-options filter-profile profile-name]
```

Description

Names of counters for which filter profile statistics are collected. The packet and byte counts for the counters are logged to a file in the `/var/log` directory.

Options

counter-name—Name of the counter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Counters

cpu

IN THIS SECTION

- Syntax | 4693
- Hierarchy Level | 4693
- Description | 4693
- Options | 4694
- Required Privilege Level | 4694
- Release Information | 4694

Syntax

```
cpu {
    reserved percent;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the percentage of CPU utilization that is always available for a logical system. The available CPU space percentage value is configured in a security profile that is bound to a logical system.

Only the primary administrator can create security profiles and bind them to logical systems.

NOTE: The `cpu-control` option at the [`edit system security-profile resources`] hierarchy level must be specified for the `reserved` value to take effect.

Options

`reserved` *`percent`*—A reserved quota that guarantees that the percentage of CPU specified is always available to the logical system.

- **Range:** 0 through 100 percent (decimal point allowed).
- **Default:** 1 percent for the primary logical system and 0 percent for user logical systems.

> **CAUTION**: The primary logical system must not be bound to a security profile that is configured with a 0 percent reserved CPU quota as traffic loss could occur.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only)

cpu-factor (Dampen Options)

IN THIS SECTION

- Syntax | 4695
- Hierarchy Level | 4695
- Description | 4695
- Options | 4695
- Required Privilege Level | 4695

Release Information | 4696

Syntax

```
cpu-factor cpu-factor;
```

Hierarchy Level

```
[edit event-options event-script dampen dampen-options],
[edit event-options event-script file filename dampen dampen-options],
[edit system scripts commit dampen dampen-options],
[edit system scripts op dampen dampen-options],
[edit system scripts op file filename dampen dampen-options]
```

Description

Slow down the operation of scripts executed on the device based on how heavily the CPU is being used.

Script dampening is initiated when CPU use exceeds the value specified using the `cpu-factor` statement. Junos OS then executes the number of lines of script specified using the `line-interval` statement and then pauses script execution for the time specified using the `time-interval` statement. If you do not configure values for these statements, the defaults are used.

Options

cpu-factor Specify the CPU use factor at which script execution is slowed down.

- Range: 0 through 100
- Default: 20

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Dampen Script Execution

cpu-load-threshold

IN THIS SECTION

- Syntax | 4696
- Hierarchy Level | 4696
- Description | 4696
- Options | 4697
- Required Privilege Level | 4697
- Release Information | 4697

Syntax

```
cpu-load-threshold percentage;
```

Hierarchy Level

```
[edit interfaces interface-name services-options session-limit]
```

Description

Regulate the usage of CPU resources on services cards. When the CPU usage exceeds the configured value (percentage of the total available CPU resources), the system reduces the rate of new sessions so

that the existing sessions are not affected by low CPU availability. The CPU utilization is constantly monitored, and if the CPU usage remains above the configured `cpu-load-threshold` value for a continuous period of 5 seconds, Junos OS reduces the session rate value configured at `edit interfaces` *`interface-name`* `services-options session-limit rate (Interface Services)` by 10%. This is repeated until the CPU utilization comes down to the configured limit.

Options

percentage Percentage of total available CPU resources.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 13.2.

Support added in Junos OS Release 19.3R2for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

cpu (resource-manager)

IN THIS SECTION

- Syntax | 4698
- Hierarchy Level | 4698
- Description | 4698
- Options | 4698
- Required Privilege Level | 4699
- Release Information | 4699

Syntax

```
cpu {
    disable-management;
    match {
        threshold threshold;
    }
    then {
        alarm;
        drop-cps;
        drop-flows;
        log;
        max-session-setup-rate max-session-setup-rate;
    }
}
```

Hierarchy Level

```
[edit security resource-manager]
```

Description

Configure resource manager cpu options to control cpu load. Resource manager takes protective action when the traffic load reaches the threshold.

Options

- `disable-management`-Disables the cpu management. By default, cpu management is on.
- `match` *`threshold`*-Configure the threshold limit to let `resource-manager` take actions. The range is 1 to 100 percent.
- `drop-cps`-This is default action when the CPU usage crosses the default or configured threshold. This configuration drops the packets which do not have an existing associated session.
- `max-session-setup-rate`-This throttles new session creation to the specified rate when the cpu threshold is exceeded.

- `drop-flows`-When you mark a session as a drop-flow, it bypasses the first path for each new packet, which matches the original drop session for next 4 seconds only. This configuration drops the packets that match the `drop-flows`.
- `log`-This configuration logs messages when the resource threshold is exceeded and falls back under the threshold.
- `alarm`-This configuration sends an SNMP Trap when the CPU load exceeds the threshold.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Traffic Processing on SRX Series Devices Overview

cpu (TDF Gateway)

IN THIS SECTION

- Syntax | 4700
- Hierarchy Level | 4700
- Description | 4700
- Options | 4700
- Required Privilege Level | 4700
- Release Information | 4700

Syntax

```
cpu cpu-pct;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name cac]
```

Description

Configure the threshold for the maximum amount of CPU that the TDF gateway can use. If the amount of CPU that the TDF gateway uses reaches the threshold, the SNMP trap **jnxScgSMCPUThreshHigh** is generated.

Options

cpu-pct Maximum percentage of CPU.

- **Range:** 1 through 90.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a TDF Gateway

cpu-throttle (Next Gen Services)

IN THIS SECTION

- Syntax | 4701
- Hierarchy Level | 4701
- Description | 4701
- Options | 4702
- Required Privilege Level | 4702
- Release Information | 4702

Syntax

```
cpu-throttle {
    percentage percent;
}
```

Hierarchy Level

```
[edit services screen]
```

Description

Specify the services card CPU utilization percentage that triggers the installation of a dynamic filter on the PFEs of the line cards for suspicious activity. The dynamic filter drops the suspicious traffic.

In addition to this threshold, at least one of the following conditions is required to trigger the installation of a dynamic filter:

- The packet rate from an individual source address or to an individual destination address must exceed four times the configured `packet-rate` at the `[edit services screen ids-option screen-name limit-session by-source]` or `[edit services screen ids-option screen-name limit-session by-destination]` hierarchy level.

- The connection rate from an individual source address or to an individual destination address must exceed four times the configured `session-rate` at the `[edit services screen ids-option screen-name limit-session by-source]` or `[edit services screen ids-option screen-name limit-session by-destination]` hierarchy level.

Dynamic filters are not created from IDS screens that use subnet aggregation.

The dynamic filter drops the suspicious traffic at the PFE, without the traffic being processed by the IDS screen. When the packet or connection rate no longer exceeds four times the limit in the IDS screen, the dynamic filter is removed.

Options

percentage *percent*

The CPU utilization percentage.

- **Range:** 1 through 100
- **Default:** 90

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

craft-lockout

IN THIS SECTION

- Syntax | 4703
- Hierarchy Level | 4704
- Description | 4704
- Required Privilege Level | 4704
- Release Information | 4704

Syntax

```
craft-lockout {
    alarm {
        interface-type {
            link-down (red | yellow | ignore);
        }
    }
    container-devices {
        device-count number;
            }
        }
    fpc slot {
        pic pic-number {
            fibre-channel {
                port-range {
                    port-range-low port-range-high;
                }
            }
        }
    }
    routing-engine
        on-disk-failure {
            disk-failure-action (halt | reboot);
            }
        }
```

```
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Disable the physical operation of the craft interface front panel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos Release 11.3.

RELATED DOCUMENTATION

- Configuring the Junos OS to Disable the Physical Operation of the Craft Interface

crb-proxy-mac

IN THIS SECTION

- Syntax | 4705
- Hierarchy Level | 4705
- Description | 4705
- Options | 4707
- Required Privilege Level | 4707

Release Information | 4707

Syntax

```
crb-proxy-mac {
    family {
        inet IPv4-proxy-MAC-address;
        inet6 IPv6-proxy-MAC-address;
    }
}
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

Enable this statement on Ethernet VPN (EVPN) leaf devices that act as Layer 2 (L2) gateways to manage learning and aging for IPv4 Address Resolution Protocol (ARP) and IPv6 Neighbor Discovery Protocol (NDP) entries.

After you enable the `crb-proxy-mac` option on an EVPN provider edge (PE) leaf device, the leaf device starts learning MAC-IP bindings based on transit traffic from the attached customer edge (CE) devices. The leaf device then uses EVPN Type 2 routes to advertise the MAC-IP information it learns to other PE devices. In an EVPN centrally-routed bridging (CRB) overlay fabric, as part of the L2 gateway role, leaf devices with `crb-proxy-mac` set are also responsible to refresh MAC-IP information toward the CE devices. To do that refresh operation, the leaf devices use the IPv4 or IPv6 proxy addresses you configure with the `crb-proxy-mac` option to send ARP request packets or NDP solicitation packets to the CE devices.

L2 gateway devices typically don’t use integrated routing and bridging (IRB) interfaces. L2 gateway devices also don’t re-send ARP requests or do ARP refresh operations. In EVPN bridged overlay (BO) topologies, all of the devices act as EVPN L2 gateways. In EVPN CRB overlay topologies, the leaf devices act as L2 gateways, so you don’t usually configure IRB interfaces for EVPN traffic on the leaf devices either. The L2 gateway devices advertise only MAC (EVPN type 2) routes for attached hosts and can’t resolve MAC-IP bindings for the hosts. As a result, the L3 gateway spine devices in the network need to use ARP and NDP to learn those bindings.

Because Junos OS and Junos OS Evolved spine devices in a CRB overlay network don't do ARP refresh operations by default, their MAC-IP binding information expires after a timeout interval. Upon receiving the next IPv4 or IPv6 packet after the MAC-IP entry expires, a spine device uses ARP or NDP to try to resolve the MAC-IP binding again. The spine device might drop traffic until it records the new MAC-IP binding.

Use the following guidelines to manage ARP and NDP learning and entry refresh operations in EVPN BO and CRB overlay topologies:

- CRB overlay environments:
  - To enable ARP refresh operations on Junos OS CRB spine devices to remote hosts, set the "proxy-macip-advertisement" on page 13831 option at the `[edit interfaces irb unit logical-unit-number]` hierarchy level on the IRB interfaces. With this setting, the L3 gateway devices can use ARP and NDP for MAC-IP learning and aging.

    NOTE: The `proxy-mac-ip-advertisement` statement is available only on Junos OS devices.

  - If your CRB overlay network has only Junos OS Evolved spine devices or a mix of spine devices that are Junos OS and Junos OS Evolved devices:
    - Set the `proxy-mac-ip-advertisement` option on the Junos OS spine devices as described above to enable ARP refresh operations to remote hosts.
    - Also set the `crb-proxy-mac` option on all the leaf devices in the CRB overlay network, including any Junos OS and Junos OS Evolved leaf devices.
- BO environments:
  - In BO EVPN networks, the EVPN leaf devices (as L2 gateways) don't participate in ARP and NDP operations by default. The ARP request packets or NDP solicitation packets are sent as broadcast traffic by the IP fabric.

    To avoid this ARP and NDP broadcast traffic, configure the `crb-proxy-mac` option on the leaf devices. With `crb-proxy-mac` set, the leaf devices learn MAC-IP information from the transit traffic, and use EVPN Type 2 routes to update other EVPN leaf devices. Then the leaf devices can reply to ARP request packets or NDP solicitation packets from their connected hosts. The ARP request packets and NDP packets don't need to be broadcast in the IP fabric.

NOTE: Multihomed peers must share the same MAC address.

Options

inet *IPv4-proxy-MAC-address*	Configure the IPv4 proxy MAC address. The device uses this address to send ARP requests toward attached CE devices. This address must be the same on all leaf nodes.
inet6 *IPv6-proxy-MAC-address*	Configure the IPv6 proxy MAC address. The device uses this address to send NDP solicitation packets toward attached CE devices. This address must be the same on all leaf nodes.

Required Privilege Level

interface—To view this statement in the configuration

interface-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release and Junos OS Evolved Release 21.2R1.

create-new-ucast-tunnel

IN THIS SECTION

- Syntax | 4707
- Hierarchy Level | 4708
- Description | 4708
- Required Privilege Level | 4708
- Release Information | 4708

Syntax

```
create-new-ucast-tunnel;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name provider-tunnel ingress-replication],
[edit routing-instances routing-instance-name provider-tunnel selective group address source
source-address ingress-replication]
```

Description

One of two modes for building unicast tunnels when ingress replication is configured for the provider tunnel. When this statement is configured, each time a new destination is added to the multicast distribution tree, a new unicast tunnel to the destination is created in the ingress replication tunnel. The new tunnel is deleted if the destination is no longer needed. Use this mode for RSVP LSPs using ingress replication.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs

mpls-internet-multicast | **10978**

ingress-replication | **8647**

create-req

IN THIS SECTION

Syntax | **4709**

- Hierarchy Level | 4709
- Description | 4709
- Options | 4709
- Required Privilege Level | 4710
- Release Information | 4710

Syntax

```
create-req {
    alarm-threshold (Security GPRS) {
        forward number;
        reverse number;
    }
    drop-threshold (Security GPRS) {
        forward number;
        reverse number;
    }
}
```

Hierarchy Level

```
[edit security gtp profile profile-name path-rate-limit message-type]
```

Description

Configure the create-req parameter to limit or restrict the number of packets per second for the GPRS tunneling protocol (GTP) create request.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name path-rate-limit message-type]` hierarchy level is replaced by `[edit security gtp profile profile-name path-rate-limit message-type]`.

Options

`alarm-threshold`—To set alarm threshold for path rate limiting.

`drop-threshold`—To set drop threshold for path rate limiting.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding Path Rate Limiting for GTP Control Messages

Example: Limiting the Message Rate and Path Rate for GTP Control Messages

credibility

IN THIS SECTION

- Syntax | 4710
- Hierarchy Level | 4711
- Description | 4711
- Options | 4711
- Required Privilege Level | 4712
- Release Information | 4712

Syntax

```
credibility {
    direct;
    isis-level-1;
    isis-level-2;
```

```
    ospf;
    static;
    unknown;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls traffic-engineering database export],
[edit protocols mpls traffic-engineering database export]
```

Description

Configure preference values for entries from BGP-TE to the traffic engineering database. A protocol with a higher credibility value is preferred over a protocol with a lower credibility value.

The credibility order for the BGP-TE protocols is as follows:

- `Unknown`—80
- `OSPF`—81
- `ISIS Level 1`—82
- `ISIS Level 2`—83
- `Static`—84
- `Direct`—85

Options

direct	Entries sourced from directly connected links.
isis-level-1	Entries sourced from IS-IS Level 1.
isis-level-2	Entries sourced from IS-IS Level 2.
ospf	Entries sourced from OSPF.
static	Entries sourced from static configuration.
unknown	Entries sourced from unknown entities.

- **Range:** 0 through 512

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

traffic-engineering (Protocols MPLS) | 30031

credit (Interfaces)

IN THIS SECTION

- Syntax | 4712
- Hierarchy Level | 4713
- Description | 4713
- Required Privilege Level | 4713
- Release Information | 4713

Syntax

```
credit {
    interval number;
}
```

Hierarchy Level

```
[edit interfaces interface-name radio-router]
```

Description

This parameter controls credit-based scheduling parameters and includes an interval option to set the grant rate interval to a value between 1–60 seconds.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Understanding Interfaces

critical-event (Process)

IN THIS SECTION

- Syntax | 4714
- Hierarchy Level | 4714
- Description | 4714
- Options | 4714
- Required Privilege Level | 4714
- Release Information | 4715

Syntax

```
critical-event {
    threshold process name memory-limit value ;
    action [snmp-trap | live-core | (terminate-process| terminate-process-force)];
}
```

Hierarchy Level

```
[edit system monitor memory process]
```

Description

Indicates an event that, if left unchecked, may result in serious loss of service or interruption of traffic. You can specify the process threshold value for specific processes and their memory limit. When the threshold is crossed, you can configure the action to be performed. For instance, you can specify the process threshold for the `rpd` process and the memory limit for that process which you would classify as a critical event.

Options

threshold Specifies the process threshold level for a minor event. You must specify the process name and the memory limit for the process.

memory Specifies the permissible memory limit for the process after which the critical event is triggered.

action Specifies the action to be performed when the threshold is crossed. For critical events, you can force terminate the process. You can also display a list of system core files created when the system failed which you can use to debug the issue. You can also raise an SNMP trap.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1.

RELATED DOCUMENTATION

System Log Overview

System Log Explorer

critical-event (System)

IN THIS SECTION

- Syntax | 4715
- Hierarchy Level | 4715
- Description | 4716
- Options | 4716
- Required Privilege Level | 4716
- Release Information | 4716

Syntax

```
critical-event {
    threshold swap-usage value ;
    action [snmp-trap | live-vmcore | (RE-restart| RE-switchover)];
}
```

Hierarchy Level

```
[edit system monitor memory system]
```

Description

Indicates an event that, if left unchecked, may result in serious loss of service or interruption of traffic. At the system level you can monitor the overall free pages (cache and free pages), kernel memory allocations and compare them with threshold values. You can also monitor the swap usage and specify the threshold for swap space usage. Most devices have a swap partition that you can configure to provide additional transient memory. Ideally, the swap partition is used only under exceptional circumstances. In current dynamic systems, you can sporadically use the swap partition as long as the usage is not high or persistent. For instance, you could categorise a critical event as one where the swap usage is greater than 75%.

You can only configure swap usage threshold for memory monitoring. If you do not configure the swap usage threshold, memory monitoring for swap usage is based on the default threshold. Memory monitoring for other features is based on the default threshold.

Options

threshold Specifies the process threshold level for a minor event. You must specify the process name and the memory limit for the process.

swap-usage Specifies the swap usage in percentage.

action Specifies the action to be performed when the threshold is crossed.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1.

RELATED DOCUMENTATION

System Log Overview

System Log Explorer

crl

IN THIS SECTION

- Syntax | 4717
- Hierarchy Level | 4717
- Description | 4717
- Options | 4718
- Required Privilege Level | 4718
- Release Information | 4718

Syntax

```
crl {
    disable disable;
    if-not-present (allow | drop);
    ignore-hold-instruction-code ignore-hold-instruction-code;
}
```

Hierarchy Level

```
[edit services ssl initiation profile profile-name  actions]
[edit services ssl proxy profile profile-name actions]
```

Description

Specify certificate revocation actions.

CRL validation on SRX Series Firewall involves checking for revoked certificates from servers. You can enable or disable the CRL validation to meet your specific security requirements. You can allow or drop the sessions when a CRL information is not available.

To enhance security, the certificate revocation checking feature has been enabled by default on SRX Series Firewalls on any SSL proxy profile.

Options

disable Disable CRL validation.

if-not-present Specify an action if CRL information is not present.

- Values:
  - `allow`—Allow session if CRL information is not present.
  - `drop`—Drop session if CRL information is not present.

ignore-hold-instruction-code Allow the sessions when a certificate is revoked and the revocation reason is on hold.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 15.1X49-D30. This statement is supported in the SRX340, SRX345, SRX550M, SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall instances.

RELATED DOCUMENTATION

Working with the Certificate Revocation Lists for SSL Proxy

crl (Adaptive Services Interface)

IN THIS SECTION

- Syntax | 4719
- Hierarchy Level | 4719
- Description | 4719

- Options | 4719
- Required Privilege Level | 4720
- Release Information | 4720

Syntax

```
crl {
    disable on-download-failure;
    refresh-interval number-of-hours;
    url {
        url-name;
        password;
    }
}
```

Hierarchy Level

```
[edit security pki ca-profile ca-profile-name revocation-check]
```

Description

Configure the certificate revocation list (CRL). A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis.

Options

`disable on-download-failure`—Permit the authentication of the IPsec peer when the CRL is not downloaded.

`password`—Password to access the URLs.

`refresh-interval` *`number-of-hours`*—Time interval, in hours, between CRL updates.

- **Range:** 0 through 8784
- **Default:** 24

`url` *`url-name`*—Location from which to retrieve the CRL through the Lightweight Directory Access Protocol (LDAP). You can configure as many as three URLs for each configured CA profile.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring the Certificate Revocation List

crl (Encryption Interface)

IN THIS SECTION

- Syntax | 4720
- Hierarchy Level | 4721
- Description | 4721
- Options | 4721
- Required Privilege Level | 4721
- Release Information | 4721

Syntax

```
crl file-name;
```

Hierarchy Level

```
[edit security certificates]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the certificate revocation list (CRL). A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis.

Options

file-name—Specify the file from which to read the CRL.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Certificate Authority Properties for an ES PIC

crl (Security)

IN THIS SECTION

- Syntax | 4722
- Hierarchy Level | 4722
- Description | 4722

- Options | 4722
- Required Privilege Level | 4723
- Release Information | 4723

Syntax

```
crl {
    disable {
        on-download-failure;
    }
    refresh-interval hours;
    url url-name;
}
```

Hierarchy Level

```
[edit security pki ca-profile ca-profile-name revocation-check]
```

Description

Configure the certificate revocation list (CRL). A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis.

Options

disable on-download-failure (Optional) Override the default behavior and permit certificate verification even if the CRL fails to download.

refresh-interval *hours* Specify the amount of time interval in hours between certificate revocation list (CRL) updates.

- **Range:** 0 through 8784 hours.

Configuring `refresh-interval` value as 0 or not configuring `refresh-interval` is considered as same in Junos. In both the cases, CRL is updated based on the value specified for the next-update time in the received CRL.

- **Default:** The CRL is updated based on the value specified for the next-update time in the received CRL. This update occurs in the following cases:
  - if the `refresh-interval` is not configured.
  - if the `refresh-interval` value is configured as 0.

url ***url-name*** Name of the location from which to retrieve the CRL through HTTP or Lightweight Directory Access Protocol (LADP). You can specify one URL for each configured CA profile. By default, no location is specified. Use a fully qualified domain name (FQDN) or an IP address and, optionally, a port number. If no port number is specified, port 80 is used for HTTP and port 443 is used for LDAP.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

`disable` option is introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Basic Elements of PKI in Junos OS

revocation-check (Security PKI) | **15739**

cross-connect

IN THIS SECTION

- Syntax | 4724
- Hierarchy Level | 4725
- Description | 4725
- Options | 4726
- Required Privilege Level | 4726
- Release Information | 4726

Syntax

```
cross-connect {
    cross-connect-name {
        physical-interface {
            hsxe0 {
                vlan-id vlan-id;
            }
            hsxe1 {
                vlan-id vlan-id;
            }
        }
        virtual-network-function vnf-name {
            interfaceinterface-name;
            vlan-id vlan-id;
        }
    }
    push-pop-cross-connect-name {
        virtual-network-function vnf-name {
            interfaceinterface-name;
            vlan-id vlan-id;
        }
        virtual-network-function vnf-name {
            interfaceinterface-name;
        }
    }
```

```
    swap-cross-connect-name {
        virtual-network-function vnf-name {
            interfaceinterface-name;
            vlan-id vlan-id;
        }
        virtual-network-function vnf-name {
            interfaceinterface-name;
            vlan-id vlan-id;
        }
    }
    unconditional-cross-connect-name {
        virtual-network-function vnf-name {
            interfaceinterface-name;
        }
        virtual-network-function vnf-name {
            interfaceinterface-name;
        }
    }
    vlan-based-cross-connect-name{
        virtual-network-function vnf-name {
            interfaceinterface-name;
            vlan-id vlan-id;
        }
        virtual-network-function vnf-name {
            interfaceinterface-name;
            vlan-id vlan-id;
        }
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Connects any two VNF interfaces, VLANs on physical interfaces such as hsxe0 and hsxe1, and also provides features such as unconditional cross-connect, VLAN-based cross-connect, cross-connect of native VLAN traffic, and cross-connect with operations such as PUSH, POP, and SWAP.

Options

virtual-network-functions ***vnf-name***	Name of the VNF instance.
cross-connect-name	Name of the cross-connect.
vlan-id	Virtual LAN identifier.

Required Privilege Level

set—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D47.

RELATED DOCUMENTATION

Configuring Cross-connect

cross-connect-ccm

IN THIS SECTION

- Syntax | 4727
- Hierarchy Level | 4727
- Description | 4727
- Required Privilege Level | 4727
- Release Information | 4727

Syntax

```
cross-connect-ccm;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
event ais-trigger-condition]
```

Description

Configure the defect condition that raises an alarm indication signal when any cross-connect continuity check messages (CCMs) are received by the MEP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Alarm Indication Signal (ETH-AIS) Function Overview

Configuring ETH-AIS on a CFM MEP

cryptographic-address

IN THIS SECTION

- Syntax | 4728
- Hierarchy Level | 4728
- Description | 4728
- Required Privilege Level | 4729
- Release Information | 4729

Syntax

```
cryptographic-address {
    key-length number;
    key-pair pathname;
}
```

Hierarchy Level

```
[edit protocols neighbor-discovery secure]
```

Description

Configure parameters for cryptographically generated addresses for Secure Neighbor Discovery.

The Secure Neighbor Discovery (SEND) Protocol uses cryptographically generated addresses (CGAs), as defined in RFC 3972, *Cryptographically Generated Addresses*, to ensure that the sender of a Neighbor Discovery Protocol (NDP) message is the “owner” of the claimed address. Each node must generate a public-private key pair before it can claim an address. The CGA is included in all outgoing neighbor solicitation and neighbor advertisement messages.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing level—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Example: Configuring Secure IPv6 Neighbor Discovery

Understanding Secure IPv6 Neighbor Discovery

csds

IN THIS SECTION

- Syntax | 4729
- Hierarchy Level | 4730
- Description | 4730
- Options | 4730
- Required Privilege Level | 4730
- Release Information | 4730

Syntax

```
csds {
    node-instance node-instance-id {
        gateway-ip gateway-ip;
        host-ip host-ip;
        ip-prefix-range ip-prefix-range;
```

```
    }
    }
```

Hierarchy Level

```
[edit system]
```

Description

Configures the Ubuntu host servers with the IP prefix range to provision the JDM in Connected Security Distributed Services (CSDS) architecture. JDM installation assigns the first IP in the range to the JDM container and the subsequent IPs to the vSRX Virtual Firewalls. Run this command on MX Series Junos Node Unifier (JNU) controller.

Options

node-instance Add new node instance. Set the host instance ID for the baremetal server host.

- Range: 0-100

gateway-ip Provide gateway IP for JDM and vSRX.

host-ip Provide host IP address.

ip-prefix-range Provide IP prefix-range. Prefix range must be contiguous.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

cspf-backoff-time

IN THIS SECTION

- Syntax | 4731
- Hierarchy Level | 4731
- Release Information | 4732
- Description | 4732
- Options | 4732
- Required Privilege Level | 4732

Syntax

```
cspf-backoff-time seconds;
```

Hierarchy Level

```
[edit logical-systems  name protocols mpls],
[edit logical-systems name routing-instances name protocols mpls],
[edit logical-systems name tenants name routing-instances name protocols mpls],
[edit logical-systems name protocols mpls label-switched-path],
[edit logical-systems name routing-instances name protocols mpls label-switched-path],
[edit logical-systems name tenants name routing-instances name protocols mpls label-switched-
path]
[edit protocols  mpls],
[edit routing-instances name protocols mpls],
[edit tenants name routing-instances name protocols mpls],
[edit protocols mpls label-switched-path],
```

```
[edit routing-instances name protocols mpls label-switched-path],
[edit tenants name routing-instances name protocols mpls label-switched-path]
```

Release Information

Statement introduced in Junos OS Release 20.3R1 and Junos OS Evolved Release 20.3R1 on all platforms.

Description

The statement enables you to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

Options

seconds—CSPF delay time in seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

adjust-interval | 1946

rsvp-error-hold-time | 16056

cts

IN THIS SECTION

Syntax | 4733

- Hierarchy Level | 4733
- Description | 4733
- Options | 4733
- Required Privilege Level | 4733
- Release Information | 4734

Syntax

```
cts (ignore | normal | require);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options dce-options],
[edit interfaces interface-name serial-options dte-options]
```

Description

For EIA-530 and V.35 interfaces only, configure the from-DCE signal, clear-to-send (CTS).

Options

`ignore`—The from-DCE signal is ignored.

`normal`—Normal CTS signal handling as defined by the TIA/EIA Standard 530.

`require`—The from-DCE signal must be asserted.

- **Default:** `normal`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Signal Handling

cts-polarity

IN THIS SECTION

- Syntax | 4734
- Hierarchy Level | 4734
- Description | 4734
- Options | 4735
- Required Privilege Level | 4735
- Release Information | 4735

Syntax

```
cts-polarity (negative | positive);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

Configure CTS signal polarity.

Options

`positive`—Positive signal polarity.

`negative`—Negative signal polarity.

- **Default:** `positive`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Serial Signal Polarities

current-hop-limit

IN THIS SECTION

- Syntax | 4736
- Hierarchy Level | 4736
- Description | 4736
- Options | 4736
- Required Privilege Level | 4736
- Release Information | 4736

Syntax

```
current-hop-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Description

Set the default value placed in the hop count field of the IP header for outgoing packets.

Options

number—Hop limit. A value of 0 means the limit is unspecified by this router.

- **Range:** 0 through 255
- **Default:** 64

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

current-hop-limit (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 4737
- Hierarchy Level | 4737
- Description | 4737
- Options | 4737
- Required Privilege Level | 4737
- Release Information | 4738

Syntax

```
current-hop-limit number;
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Default value placed in the hop count field of the IP header for outgoing packets.

Options

number—Hop limit. A value of 0 means the limit is unspecified by this router.

- **Range:** 0 through 255
- **Default:** 64

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA

Dynamic Router Advertisement Configuration Overview

custom-attack

IN THIS SECTION

- Syntax | 4738
- Hierarchy Level | 4745
- Description | 4745
- Options | 4745
- Required Privilege Level | 4746
- Release Information | 4746

Syntax

```
custom-attack attack-name {
    attack-type (Security Anomaly) {
        anomaly {
            direction (any | client-to-server | server-to-client);
            service service-name;
            shellcode (all | intel | no-shellcode | sparc);
            test test-condition;
        }
        chain {
```

```
            expression boolean-expression;
            member member-name {
                attack-type (Security Anomaly) {
                    (anomaly ...same statements as in [edit security idp custom-attack attack-
name attack-type anomaly] hierarchy level | signature ...same statements as in [edit security
idp custom-attack attack-name attack-type signature] hierarchy level);
                }
            }
            order;
            protocol-binding {
                application application-name;
                icmp;
                icmpv6;
                ip {
                    protocol-number transport-layer-protocol-number;
                }
                ipv6 {
                    protocol-number transport-layer-protocol-number;
                }
                rpc {
                    program-number rpc-program-number;
                }
                tcp {
                    minimum-port port-number <maximum-port port-number>;
                }
                udp {
                    minimum-port port-number <maximum-port port-number>;
                }
            }
            reset;
            scope (session | transaction);
        }
        signature (Security IDP) {
            context context-name;
            direction (any | client-to-server | server-to-client);
            negate;
            pattern signature-pattern;
            pattern-pcre signature-pattern-pcre;
            protocol (Security IDP Signature Attack) {
                icmp {
                    checksum-validate {
                        match (equal | greater-than | less-than | not-equal);
                        value checksum-value;
```

```
            }
            code {
                match (equal | greater-than | less-than | not-equal);
                value code-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
            }
            identification {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            sequence-number {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
            type {
                match (equal | greater-than | less-than | not-equal);
                value type-value;
            }
        }
        icmpv6 (Security IDP Custom Attack) {
            checksum-validate {
                match (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
            code {
                match (equal | greater-than | less-than | not-equal);
                value code-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
            }
            identification {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            sequence-number {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
```

```
            type {
                match (equal | greater-than | less-than | not-equal);
                value type-value;
            }
        }
        ipv4 (Security IDP Signature Attack) {
            checksum-validate {
                match (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
            destination {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            identification {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            ihl {
                match (equal | greater-than | less-than | not-equal);
                value ihl-value;
            }
            ip-flags {
                (df | no-df);
                (mf | no-mf);
                (rb | no-rb);
            }
            protocol (Security IDP Signature Attack) {
                match (equal | greater-than | less-than | not-equal);
                value transport-layer-protocol-id;
            }
            source {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            tos {
                match (equal | greater-than | less-than | not-equal);
                value type-of-service-in-decimal;
            }
            total-length {
                match (equal | greater-than | less-than | not-equal);
                value total-length-of-ip-datagram;
            }
```

```
                ttl {
                    match (equal | greater-than | less-than | not-equal);
                    value time-to-live;
                }
            }
            ipv6 {
                destination {
                    match (equal | greater-than | less-than | not-equal);
                    value ip-address-or-hostname;
                }
                flow-label {
                    match (equal | greater-than | less-than | not-equal);
                    value flow-label-value;
                }
                hop-limit {
                    match (equal | greater-than | less-than | not-equal);
                    value hop-limit-value;
                }
                next-header {
                    match (equal | greater-than | less-than | not-equal);
                    value next-header-value;
                }
                payload-length {
                    match (equal | greater-than | less-than | not-equal);
                    value payload-length-value;
                }
                source {
                    match (equal | greater-than | less-than | not-equal);
                    value ip-address-or-hostname;
                }
                traffic-class {
                    match (equal | greater-than | less-than | not-equal);
                    value traffic-class-value;
                }
            tcp (Security IDP Signature Attack){
                ack-number {
                    match (equal | greater-than | less-than | not-equal);
                    value acknowledgement-number;
                }
                checksum-validate {
                    match (equal | greater-than | less-than | not-equal);
                    value checksum-value;
                }
```

```
data-length {
    match (equal | greater-than | less-than | not-equal);
    value tcp-data-length;
}
destination-port (Security Signature Attack) {
    match (equal | greater-than | less-than | not-equal);
    value destination-port;
}
header-length {
    match (equal | greater-than | less-than | not-equal);
    value header-length;
}
mss (Security IDP) {
    match (equal | greater-than | less-than | not-equal);
    value maximum-segment-size;
}
option (Security IDP) {
    match (equal | greater-than | less-than | not-equal);
    value tcp-option;
}
reserved (Security IDP Custom Attack) {
    match (equal | greater-than | less-than | not-equal);
    value reserved-value;
}
sequence-number (Security IDP ICMP Headers) {
    match (equal | greater-than | less-than | not-equal);
    value sequence-number;
}
source-port {
    match (equal | greater-than | less-than | not-equal);
    value source-port;
}
tcp-flags {
    (ack | no-ack);
    (fin | no-fin);
    (psh | no-psh);
    (r1 | no-r1);
    (r2 | no-r2);
    (rst | no-rst);
    (syn | no-syn);
    (urg | no-urg);
}
urgent-pointer {
```

```
                match (equal | greater-than | less-than | not-equal);
                value urgent-pointer;
            }
            window-scale {
                match (equal | greater-than | less-than | not-equal);
                value window-scale-factor;
            }
            window-size {
                match (equal | greater-than | less-than | not-equal);
                value window-size;
            }
        }
        udp (Security IDP Signature Attack) {
            checksum-validate {
                match (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
            }
            destination-port {
                match (equal | greater-than | less-than | not-equal);
                value destination-port;
            }
            source-port {
                match (equal | greater-than | less-than | not-equal);
                value source-port;
            }
        }
    }
    protocol-binding {
        application application-name;
        icmp;
        icmpv6;
        ip {
            protocol-number transport-layer-protocol-number;
        }
        ipv6 {
            protocol-number transport-layer-protocol-number;
        }
        rpc {
            program-number rpc-program-number;
```

```
                }
                tcp {
                    minimum-port port-number <maximum-port port-number>;
                }
                udp {
                    minimum-port port-number <maximum-port port-number>;
                }
            }
            regexp regular-expression;
            shellcode (all | intel | no-shellcode | sparc);
        }
    }
    recommended-action (close | close-client | close-server | drop | drop-packet | ignore |
none);
    severity (critical | info | major | minor | warning);
    time-binding {
        count count-value;
        scope (destination | peer | source);
    }
}
```

Hierarchy Level

```
[edit security idp]
[edit tenants tenant-name security idp]
```

Description

Configure custom attack objects to detect a known or unknown attack that can be used to compromise your network.

Options

attack-name—Name of the custom attack object. The maximum number of characters allowed for a custom attack object name is 60.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.3.

custom-attack-group

IN THIS SECTION

- Syntax | 4746
- Hierarchy Level | 4746
- Description | 4747
- Options | 4747
- Required Privilege Level | 4747
- Release Information | 4747

Syntax

```
custom-attack-group custom-attack-group-name {
    group-members [attack-or-attack-group-name];
}
```

Hierarchy Level

```
[edit security idp]
[edit tenants tenant-name security idp]
```

Description

Configure custom attack group. A custom attack group is a list of attacks that would be matched on the traffic if the group is selected in a policy.

Options

custom-attack-group-name—Name of the custom attack group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

custom-attack-groups (Security IDP)

IN THIS SECTION

- Syntax | 4748
- Hierarchy Level | 4748
- Description | 4748
- Options | 4748
- Required Privilege Level | 4748
- Release Information | 4748

Syntax

```
custom-attack-groups attack-group-name;
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match attacks]
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match attacks]
```

Description

Specify a name for the custom attack group.

Options

attack-group-name—Name of the custom attack group.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

custom-attacks

IN THIS SECTION

- Syntax | 4749
- Hierarchy Level | 4749
- Description | 4749

- Options | 4749
- Required Privilege Level | 4749
- Release Information | 4749

Syntax

```
custom-attacks [attack-name];
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match attacks],
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match attacks]
```

Description

Select custom attacks defined under [`edit security idp custom-attack`] by specifying their names.

Options

`attack-name`—Name of the new custom attack object.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

custom-block-message

IN THIS SECTION

- Syntax | 4750
- Hierarchy Level | 4750
- Description | 4750
- Required Privilege Level | 4750
- Release Information | 4750

Syntax

```
custom-block-message value;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated profile profile-name]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name]
```

Description

Enter a custom message to be sent when HTTP requests are blocked.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Surf-Control feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

custom-ciphers

IN THIS SECTION

- Syntax (Prior to Junos OS Release 21.2R1) | 4751
- Syntax (Starting in Junos OS Release 21.2R1) | 4752
- Syntax (Starting in Junos OS Release 24.2R1) | 4752
- Hierarchy Level | 4754
- Description | 4754
- Options | 4754
- Required Privilege Level | 4757
- Release Information | 4758

Syntax (Prior to Junos OS Release 21.2R1)

```
custom-ciphers [ecdhe-rsa-with-3des-ede-cbc-sha | ecdhe-rsa-with-aes-128-cbc-sha | ecdhe-rsa-
with-aes-128-cbc-sha256 |
ecdhe-rsa-with-aes-128-gcm-sha256 | ecdhe-rsa-with-aes-256-cbc-sha | ecdhe-rsa-with-aes-256-cbc-
sha384 |
ecdhe-rsa-with-aes-256-gcm-sha384 | rsa-with-aes-128-cbc-sha256 RSA | rsa-with-aes-128-gcm-
sha256 RSA |
rsa-with-aes-256-cbc-sha256 RSA | rsa-with-aes-256-gcm-sha384 RSA | rsa-with-rc4-128-md5  RSA |
128bit rc4 |
md5 hash  rsa-with-rc4-128-sha  RSA | 128bit rc4 |sha hash  rsa-with-des-cbc-sha  RSA | des cbc
|
sha hash  rsa-with-3des-ede-cbc-sha  RSA | 3des ede/cbc | sha hash  rsa-with-aes-128-cbc-sha
RSA |
128 bit aes/cbc | sha hash  rsa-with-aes-256-cbc-sha  RSA | 256 bit aes/cbc |sha hash  rsa-
export-with-rc4-40-md5  RSA-export |
40 bit rc4 | md5 hash  rsa-export-with-des40-cbc-sha  RSA-export | 40 bit des/cbc | sha hash
rsa-with-null-md5
```

```
RSA | no symmetric cipher | md5 hash  rsa-with-null-sha    RSA | no symmetric cipher | sha hash
|
ecdhe-ecdsa-with-aes-256-gcm-sha384 | ecdhe-ecdsa-with-aes-256-cbc-sha384 | ecdhe-ecdsa-with-
aes-256-cbc-sha |
ecdhe-ecdsa-with-aes-128-gcm-sha256 |ecdhe-ecdsa-with-aes-128-cbc-sha256 |ecdhe-ecdsa-with-
aes-128-cbc-sha
|ecdhe-ecdsa-with-3des-ede-cbc-sha]
```

Syntax (Starting in Junos OS Release 21.2R1)

```
custom-ciphers [tls12-rsa-aes-128-cbc-sha | tls12-rsa-aes-256-cbc-sha | tls12-rsa-aes-256-gcm-
sha384 | tls12-rsa-aes-256-cbc-sha256 |
tls12-rsa-aes-128-gcm-sha256 | tls12-rsa-aes-128-cbc-sha256 | tls12-ecdhe-rsa-aes-256-gcm-sha384
|
tls12-ecdhe-rsa-aes-256-cbc-sha | tls12-ecdhe-rsa-aes-256-cbc-sha384 | tls12-ecdhe-rsa-3des-ede-
cbc-sha |
tls12-ecdhe-rsa-aes-128-gcm-sha256 | tls12-ecdhe-rsa-aes-128-cbc-sha | tls12-ecdhe-rsa-aes-128-
cbc-sha256 |
tls12-rsa-aes-128-ccm-sha256 | tls12-rsa-aes-128-ccm8-sha256 | tls12-rsa-aes-256-ccm-sha256 |
tls12-rsa-aes-256-ccm8-sha256 |
tls12-ecdhe-ecdsa-aes-128-ccm-sha256 | tls12-ecdhe-ecdsa-aes-256-ccm-sha256 | tls12-ecdhe-ecdsa-
aes-128-ccm8-sha256 |
tls12-ecdhe-ecdsa-aes-256-ccm8-sha256 | tls12-ecdhe-rsa-chacha20-poly1305-sha256 | tls12-ecdhe-
ecdsa-aes-256-gcm-sha384 |
tls12-ecdhe-ecdsa-aes-256-cbc-sha | tls12-ecdhe-ecdsa-aes-256-cbc-sha384 | tls12-ecdhe-ecdsa-
aes-128-gcm-sha256 |
tls12-ecdhe-ecdsa-aes-128-cbc-sha | tls12-ecdhe-ecdsa-aes-128-cbc-sha256 | tls12-ecdhe-
ecdsa-3des-ede-cbc-sha |
dhe-rsa-with-aes-128-ccm8-sha256 | dhe-rsa-with-aes-256-ccm8-sha256 | dhe-rsa-with-aes-256-ccm-
sha256 | dhe-rsa-with-aes-128-ccm-sha256 |
tls12-ecdhe-ecdsa-chacha20-poly1305-sha256 | tls13-with-aes-256-gcm-sha384 | tls13-with-aes-128-
gcm-sha256 |
tls13-with-chacha20-poly1305-sha256 | tls13-with-aes-128-ccm-sha256 | tls13-with-aes-128-ccm8-
sha256]
```

Syntax (Starting in Junos OS Release 24.2R1)

```
custom-ciphers [dhe-rsa-with-3des-ede-cbc-sha | dhe-rsa-with-aes-128-cbc-sha | dhe-rsa-with-
aes-128-cbc-sha256 |
```

```
dhe-rsa-with-aes-128-ccm8-sha256 | dhe-rsa-with-aes-128-ccm-sha256 | dhe-rsa-with-aes-128-gcm-
sha256 |
dhe-rsa-with-aes-256-cbc-sha | dhe-rsa-with-aes-256-cbc-sha256 | dhe-rsa-with-aes-256-ccm8-
sha256 |
dhe-rsa-with-aes-256-ccm-sha256 | dhe-rsa-with-aes-256-gcm-sha384 | ecdhe-ecdsa-with-3des-ede-
cbc-sha |
ecdhe-ecdsa-with-aes-128-cbc-sha | ecdhe-ecdsa-with-aes-128-cbc-sha256 | ecdhe-ecdsa-with-
aes-128-gcm-sha256 |
ecdhe-ecdsa-with-aes-256-cbc-sha | ecdhe-ecdsa-with-aes-256-cbc-sha384 | ecdhe-ecdsa-with-
aes-256-gcm-sha384 |
ecdhe-rsa-with-3des-ede-cbc-sha | ecdhe-rsa-with-aes-128-cbc-sha | ecdhe-rsa-with-aes-128-cbc-
sha256 |
ecdhe-rsa-with-aes-128-gcm-sha256 | ecdhe-rsa-with-aes-256-cbc-sha | ecdhe-rsa-with-aes-256-cbc-
sha384 |
ecdhe-rsa-with-aes-256-gcm-sha384 | rsa-export1024-with-des-cbc-sha | rsa-export1024-with-rc4-56-
md5 |
rsa-export1024-with-rc4-56-sha | rsa-export-with-des40-cbc-sha | rsa-export-with-rc4-40-md5 |
rsa-with-3des-ede-cbc-sha |
rsa-with-aes-128-cbc-sha | rsa-with-aes-128-cbc-sha256 | rsa-with-aes-128-gcm-sha256 | rsa-with-
aes-256-cbc-sha |
rsa-with-aes-256-cbc-sha256 | rsa-with-aes-256-gcm-sha384 | rsa-with-des-cbc-sha | rsa-with-
rc4-128-md5 | rsa-with-rc4-128-sha |
tls12-ecdhe-ecdsa-3des-ede-cbc-sha | tls12-ecdhe-ecdsa-aes-128-cbc-sha | tls12-ecdhe-ecdsa-
aes-128-cbc-sha256 |
tls12-ecdhe-ecdsa-aes-128-ccm8-sha256 | tls12-ecdhe-ecdsa-aes-128-ccm-sha256 | tls12-ecdhe-ecdsa-
aes-128-gcm-sha256 |
tls12-ecdhe-ecdsa-aes-256-cbc-sha | tls12-ecdhe-ecdsa-aes-256-cbc-sha384 | tls12-ecdhe-ecdsa-
aes-256-ccm8-sha256 |
tls12-ecdhe-ecdsa-aes-256-ccm-sha256 | tls12-ecdhe-ecdsa-aes-256-gcm-sha384 | tls12-ecdhe-ecdsa-
chacha20-poly1305-sha256 |
tls12-ecdhe-rsa-3des-ede-cbc-sha | tls12-ecdhe-rsa-aes-128-cbc-sha | tls12-ecdhe-rsa-aes-128-cbc-
sha256 |
tls12-ecdhe-rsa-aes-128-gcm-sha256 | tls12-ecdhe-rsa-aes-256-cbc-sha | tls12-ecdhe-rsa-aes-256-
cbc-sha384 |
tls12-ecdhe-rsa-aes-256-gcm-sha384 | tls12-ecdhe-rsa-chacha20-poly1305-sha256 | tls12-rsa-
aes-128-cbc-sha |
tls12-rsa-aes-128-cbc-sha256 | tls12-rsa-aes-128-ccm8-sha256 | tls12-rsa-aes-128-ccm-sha256 |
tls12-rsa-aes-128-gcm-sha256 |
tls12-rsa-aes-256-cbc-sha | tls12-rsa-aes-256-cbc-sha256 | tls12-rsa-aes-256-ccm8-sha256 | tls12-
rsa-aes-256-ccm-sha256 |
tls12-rsa-aes-256-gcm-sha384 | tls13-with-aes-128-ccm8-sha256 | tls13-with-aes-128-ccm-sha256 |
```

```
tls13-with-aes-128-gcm-sha256 |
tls13-with-aes-256-gcm-sha384 | tls13-with-chacha20-poly1305-sha256]
```

Hierarchy Level

```
[edit services ssl proxy profile profile-name]
[edit services ssl termination profile profile-name]
[edit services ssl initiation profile profile-name]
```

Description

Configure custom cipher for an SSL profile.

Custom ciphers allow you to define your own cipher list. If you do not want to use one of the three categories (strong, medium, or week) of preferred ciphers, you can select ciphers from each of the categories to form a custom cipher set.

To configure custom ciphers, you must set preferred-ciphers to custom. See "preferred-ciphers" on page 13055 for more details.

Options

dhe-rsa-with-3des-ede-cbc-sha	DHE/rsa, 3des ede/cbc, sha hash
dhe-rsa-with-aes-128-cbc-sha	DHE/rsa, 128 bit aes/cbc, sha hash
dhe-rsa-with-aes-128-cbc-sha256	DHE/rsa, 128 bit aes/cbc, sha256 hash
dhe-rsa-with-aes-128-ccm8-sha256	DHE/rsa, 128 bit aes/ccm8, sha256 hash
dhe-rsa-with-aes-128-ccm-sha256	DHE/rsa, 128 bit aes/ccm, sha256 hash
dhe-rsa-with-aes-128-gcm-sha256	DHE/rsa, 128 bit aes/gcm, sha256 hash
dhe-rsa-with-aes-256-cbc-sha	DHE/rsa, 256 bit aes/cbc, sha hash
dhe-rsa-with-aes-256-cbc-sha256	DHE/rsa, 256 bit aes/cbc, sha256 hash
dhe-rsa-with-aes-256-ccm8-sha256	DHE/rsa, 256 bit aes/ccm8, sha256 hash
dhe-rsa-with-aes-256-ccm-sha256	DHE/rsa, 256 bit aes/ccm, sha256 hash

dhe-rsa-with-aes-256-gcm-sha384	DHE/rsa, 256 bit aes/gcm, sha384 hash
ecdhe-ecdsa-with-3des-ede-cbc-sha	ECDHE,ECDSA, 3des ede/cbc, sha hash
ecdhe-ecdsa-with-aes-128-cbc-sha	ECDHE,ECDSA, 128 bit aes/cbc, sha hash
ecdhe-ecdsa-with-aes-128-cbc-sha256	ECDHE,ECDSA, 128 bit aes/cbc, sha256 hash
ecdhe-ecdsa-with-aes-128-gcm-sha256	ECDHE,ECDSA, 128 bit aes/gcm, sha256 hash
ecdhe-ecdsa-with-aes-256-cbc-sha	ECDHE,ECDSA, 256 bit aes/cbc, sha hash
ecdhe-ecdsa-with-aes-256-cbc-sha384	ECDHE,ECDSA, 256 bit aes/cbc, sha384 hash
ecdhe-ecdsa-with-aes-256-gcm-sha384	ECDHE,ECDSA, 256 bit aes/gcm, sha384 hash
ecdhe-rsa-with-3des-ede-cbc-sha	ECDHE/rsa, 3des ede/cbc, sha hash
ecdhe-rsa-with-aes-128-cbc-sha	ECDHE/rsa, 128 bit aes/cbc, sha hash
ecdhe-rsa-with-aes-128-cbc-sha256	ECDHE/rsa, 128 bit aes/cbc, sha256 hash
ecdhe-rsa-with-aes-128-gcm-sha256	ECDHE/rsa, 128 bit aes/gcm, sha256 hash
ecdhe-rsa-with-aes-256-cbc-sha	ECDHE/rsa, 256 bit aes/cbc, sha hash
ecdhe-rsa-with-aes-256-cbc-sha384	ECDHE/rsa, 256 bit aes/cbc, sha384 hash
ecdhe-rsa-with-aes-256-gcm-sha384	ECDHE/rsa, 256 bit aes/gcm, sha384 hash
rsa-export1024-with-des-cbc-sha	RSA 1024 bit export, des/cbc, sha hash
rsa-export1024-with-rc4-56-md5	RSA 1024 bit export, 56 bit rc4, md5 hash
rsa-export1024-with-rc4-56-sha	RSA 1024 bit export, 56 bit rc4, sha hash
rsa-export-with-des40-cbc-sha	RSA-export, 40 bit des/cbc, sha hash
rsa-export-with-rc4-40-md5	RSA-export, 40 bit rc4, md5 hash
rsa-with-3des-ede-cbc-sha	RSA, 3des ede/cbc, sha hash
rsa-with-aes-128-cbc-sha	RSA, 128 bit aes/cbc, sha hash
rsa-with-aes-128-cbc-sha256	RSA, 128 bit aes/cbc, sha256 hash

rsa-with-aes-128-gcm-sha256	RSA, 128 bit aes/gcm, sha256 hash
rsa-with-aes-256-cbc-sha	RSA, 256 bit aes/cbc, sha hash
rsa-with-aes-256-cbc-sha256	RSA, 256 bit aes/cbc, sha256 hash
rsa-with-aes-256-gcm-sha384	RSA, 256 bit aes/gcm, sha384 hash
rsa-with-des-cbc-sha	RSA, des cbc, sha hash
rsa-with-rc4-128-md5	RSA, 128bit rc4, md5 hash
rsa-with-rc4-128-sha	RSA, 128bit rc4, sha hash
tls12-ecdhe-ecdsa-3des-ede-cbc-sha	ECDHE,ECDSA, 3des ede/cbc, sha hash
tls12-ecdhe-ecdsa-aes-128-cbc-sha	ECDHE,ECDSA, 128 bit aes/cbc, sha hash
tls12-ecdhe-ecdsa-aes-128-cbc-sha256	ECDHE,ECDSA, 128 bit aes/cbc, sha256 hash
tls12-ecdhe-ecdsa-aes-128-ccm8-sha256	ECDHE,ECDSA, 128 bit aes/ccm8, sha256 hash
tls12-ecdhe-ecdsa-aes-128-ccm-sha256	ECDHE,ECDSA, 128 bit aes/ccm, sha256 hash
tls12-ecdhe-ecdsa-aes-128-gcm-sha256	ECDHE,ECDSA, 128 bit aes/gcm, sha256 hash
tls12-ecdhe-ecdsa-aes-256-cbc-sha	ECDHE,ECDSA, 256 bit aes/cbc, sha hash
tls12-ecdhe-ecdsa-aes-256-cbc-sha384	ECDHE,ECDSA, 256 bit aes/cbc, sha384 hash
tls12-ecdhe-ecdsa-aes-256-ccm8-sha256	ECDHE,ECDSA, 256 bit aes/ccm8, sha256 hash
tls12-ecdhe-ecdsa-aes-256-ccm-sha256	ECDHE,ECDSA, 256 bit aes/ccm, sha256 hash
tls12-ecdhe-ecdsa-aes-256-gcm-sha384	ECDHE,ECDSA, 256 bit aes/gcm, sha384 hash
tls12-ecdhe-ecdsa-chacha20-poly1305-sha256	ECDHE,ECDSA, chacha_poly, sha256 hash
tls12-ecdhe-rsa-3des-ede-cbc-sha	ECDHE/rsa, 3des ede/cbc, sha hash
tls12-ecdhe-rsa-aes-128-cbc-sha	ECDHE/rsa, 128 bit aes/cbc, sha hash
tls12-ecdhe-rsa-aes-128-cbc-sha256	ECDHE/rsa, 128 bit aes/cbc, sha256 hash
tls12-ecdhe-rsa-aes-128-gcm-sha256	ECDHE/rsa, 128 bit aes/gcm, sha256 hash

tls12-ecdhe-rsa-aes-256-cbc-sha	ECDHE/rsa, 256 bit aes/cbc, sha hash
tls12-ecdhe-rsa-aes-256-cbc-sha384	ECDHE/rsa, 256 bit aes/cbc, sha384 hash
tls12-ecdhe-rsa-aes-256-gcm-sha384	ECDHE/rsa, 256 bit aes/gcm, sha384 hash
tls12-ecdhe-rsa-chacha20-poly1305-sha256	ECDHE/rsa, chacha_poly , sha256 hash
tls12-rsa-aes-128-cbc-sha	RSA, 128 bit aes/cbc, sha hash
tls12-rsa-aes-128-cbc-sha256	RSA, 128 bit aes/cbc, sha256 hash
tls12-rsa-aes-128-ccm8-sha256	RSA, 128 bit aes/ccm8, sha256 hash
tls12-rsa-aes-128-ccm-sha256	RSA, 128 bit aes/ccm, sha256 hash
tls12-rsa-aes-128-gcm-sha256	RSA, 128 bit aes/gcm, sha256 hash
tls12-rsa-aes-256-cbc-sha	RSA, 256 bit aes/cbc, sha hash
tls12-rsa-aes-256-cbc-sha256	RSA, 256 bit aes/cbc, sha256 hash
tls12-rsa-aes-256-ccm8-sha256	RSA, 128 bit aes/ccm8, sha256 hash
tls12-rsa-aes-256-ccm-sha256	RSA, 256 bit aes/ccm, sha256 hash
tls12-rsa-aes-256-gcm-sha384	RSA, 256 bit aes/gcm, sha384 hash
tls13-with-aes-128-ccm8-sha256	Any key-exchange, Any authentication,128bit aes, sha256 hash
tls13-with-aes-128-ccm-sha256	Any key-exchange, Any authentication,128bit aes, sha256 hash
tls13-with-aes-128-gcm-sha256	Any key-exchange, Any authentication,128bit aes, sha256 hash
tls13-with-aes-256-gcm-sha384	Any key-exchange, Any authentication,256bit aes, sha384 hash
tls13-with-chacha20-poly1305-sha256	Any key-exchange, Any authentication,chacha, sha256 hash

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

This statement is supported in the SRX340, SRX345, SRX380, SRX550M, SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall instances. Options to support Elliptic Curve Digital Signature Algorithm (ECDSA) added in Junos OS Release 18.3R1.

Transport Layer Security (TLS) protocol version 1.3 support for SSL proxy introduced in Junos OS 21.2R1.

CCM, CCM8, and TLS12 Chacha-poly ciphers introduced in Junos OS 22.1R1.

ECDSA ciphers support for SSL initiation and SSL termination profiles is available in Junos OS 24.2R1

RELATED DOCUMENTATION

SSL Proxy Overview

Configuring SSL Forward Proxy

Enabling Debugging and Tracing for SSL Proxy

custom-message (Security Content Filtering)

IN THIS SECTION

- Syntax | 4758
- Hierarchy Level | 4759
- Description | 4759
- Required Privilege Level | 4759
- Release Information | 4759

Syntax

```
custom-message message;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name notification-options]
```

Description

Custom message notifications are generally used when content is blocked by the content filter.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

custom-message (Security Email Notify)

IN THIS SECTION

- Syntax | 4760
- Hierarchy Level | 4760
- Description | 4760
- Required Privilege Level | 4760
- Release Information | 4760

Syntax

```
custom-message message;
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus sophos-engine pattern-update email-notify]
[edit security utm default-configuration anti-virus avira-engine pattern-update email-notify]
```

Description

You can configure the device to notify a specified administrator when patterns are updated. This is an e-mail notification with a custom message.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for Avira engine added in Junos OS Release 18.4R1.

custom-message (Security Fallback Block)

IN THIS SECTION

- Syntax | 4761
- Hierarchy Level | 4761
- Description | 4761
- Required Privilege Level | 4761
- Release Information | 4762

Syntax

```
custom-message message;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-block]
```

Description

Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

custom-message (Security Fallback Non-Block)

IN THIS SECTION

- Syntax | 4762
- Hierarchy Level | 4762
- Description | 4763
- Required Privilege Level | 4763
- Release Information | 4763

Syntax

```
custom-message message;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-non-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-non-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-non-block]
```

Description

Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

custom-message (Security Virus Detection)

IN THIS SECTION

- Syntax | 4763
- Hierarchy Level | 4764
- Description | 4764
- Required Privilege Level | 4764
- Release Information | 4764

Syntax

```
custom-message message;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options virus-detection]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options virus-detection]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options virus-detection]
```

Description

Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

custom-message (Security Web Filtering)

IN THIS SECTION

- Syntax | 4765
- Hierarchy Level | 4765
- Description | 4765

- Options | 4766
- Required Privilege Level | 4766
- Release Information | 4766

Syntax

```
custom-message custom-message;
```

Hierarchy Level

```
[edit security utm feature-profile web-filtering juniper-local profile profile-name]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name]
[edit logical-systems name security utm feature-profile web-filtering juniper-local profile
profile-name]
[edit logical-systems name security utm feature-profile web-filtering juniper-enhanced profile
profile-name]
[edit logical-systems name security utm feature-profile web-filtering websense-redirect profile
profile-name]
[edit tenants name security utm feature-profile web-filtering juniper-enhanced profile profile-
name]
[edit tenants name security utm feature-profile web-filtering juniper-local profile profile-name]
[edit tenants name security utm feature-profile web-filtering websense-redirect profile profile-
name]
[edit security utm default-configuration web-filtering juniper-enhanced]
[edit security utm default-configuration web-filtering juniper-local]
[edit security utm default-configuration web-filtering websense-redirect]
```

Description

You can reference the custom-page type custom-message object at the Web filtering profile level. If the category action is block or quarantine and no `custom-message` is configured in the category, the `custom-message` under the profile will take effect.

Options

custom-message	Specify the custom message object.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

custom-page | 4777
custom-page-file | 4778

custom-message-subject (Security Email Notify)

IN THIS SECTION

- Syntax | 4766
- Hierarchy Level | 4767
- Description | 4767
- Required Privilege Level | 4767
- Release Information | 4767

Syntax

```
custom-message-subject message-subject;
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus sophos-engine pattern-update email-notify]
[edit security utm default-configuration avira-engine pattern-update email-notify]
```

Description

You can configure the device to notify a specified administrator when patterns are updated. This is an e-mail notification with a custom message and a custom subject line.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for Avira engine added in Junos OS Release 18.4R1.

custom-message-subject (Security Fallback Block)

IN THIS SECTION

- Syntax | 4768
- Hierarchy Level | 4768
- Description | 4768
- Required Privilege Level | 4768

Release Information | 4768

Syntax

```
custom-message-subject message-subject;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-block]
```

Description

Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file. As part of a custom message, you can customize the message subject line.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

custom-message-subject (Security Fallback Non-Block)

IN THIS SECTION

- Syntax | 4769
- Hierarchy Level | 4769
- Description | 4769
- Required Privilege Level | 4769
- Release Information | 4770

Syntax

```
custom-message-subject message-subject;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-non-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-non-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-non-block]
```

Description

Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file. As part of a custom message, you can customize the message subject line.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

custom-message-subject (Security Virus Detection)

IN THIS SECTION

- Syntax | 4770
- Hierarchy Level | 4770
- Description | 4771
- Required Privilege Level | 4771
- Release Information | 4771

Syntax

```
custom-message-subject message-subject;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options virus-detection]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options virus-detection]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options virus-detection]
```

Description

Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file. As part of a custom message, you can customize the message subject line.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

custom-objects

IN THIS SECTION

- Syntax | 4771
- Hierarchy Level | 4772
- Description | 4772
- Options | 4773
- Required Privilege Level | 4773
- Release Information | 4773

Syntax

```
custom-objects {
    custom-url-category object-name {
```

```
        value [value];
    }
    custom-message {
        name message-name;
        type redirect-url | user-message;
        content redirect-url by user| user-message by user;
    }
    filename-extension object-name {
        value [value];
    }
    mime-pattern object-name {
        value [value];
    }
    protocol-command object-name {
        value [value];
    }
    url-pattern object-name {
        value [value];
    }
}
```

Hierarchy Level

```
[edit security utm]
[edit security utm default-configuration]
[edit logical-systems logical-system-name security utm]
[edit tenants tenant-name security utm]
```

Description

Configure custom objects before configuring Content Security feature-profile features.

WARNING: Custom category does not take precedence over predefined categories when it has the same name as one of the predefined categories. We do not recommend having a custom category name be the same as the predefined category name.

NOTE: Starting from Junos OS Release 17.4R1, support for custom category configuration is available for EWF, local, and Websense redirect profiles.

Options

The statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level introduced in Junos OS Release 18.2R1.

Support for configuration in logical systems introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Unified Threat Management Overview

custom-options

IN THIS SECTION

- Syntax | 4774
- Hierarchy Level | 4774
- Description | 4774
- Options | 4774

Required Privilege Level | 4776
Release Information | 4776

Syntax

```
custom-options {
    banner-message string;
    footer-bgcolor color;
    footer-message string;
    footer-text-color color;
    form-header-bgcolor color;
    form-header-message string;
    form-header-text-color color;
    form-reset-label label name;
    form-submit-label label name;
    header-bgcolor color;
    header-logo filename;
    header-message string;
    header-text-color color;
    post-authentication-url url-string;
}
```

Hierarchy Level

```
[edit services captive-portal]
```

Description

Specify the design elements of a captive portal login page.

Options

banner-message—The first screen displayed before the captive portal login page is displayed—for example, a disclaimer message or a terms and conditions of use page.

- **Range:** 1–2047 characters

footer-bgcolor —The hexadecimal color code for the color of the footer bar across the bottom of the captive portal login page—for example, #2E8B57 (sea green).

- **Values:** # symbol followed by six characters.

footer-message—Text message displayed in the footer bar across the bottom of the captive portal login page.

- **Range:** 1–2047 characters
- **Default:** Copyright @2010, Juniper Networks Inc.

footer-text-color — Color of the text in the footer.

- **Default:** The default color is white.

form-header-bgcolor —The hexadecimal color code for the background color of the header bar across the top of the form area of the captive portal login page.

- **Values:** # symbol followed by six characters.

form-header-message—Text message displayed in the header bar across the top of the form area of the captive portal login page.

- **Range:** 1–255 characters
- **Default:** Captive Portal User Authentication

form-header-text-color—Color of the text in the form header.

- **Default:** The default color is black.

form-reset-label—Label displayed in the button that the user can select to clear the username and password fields on the form.

- **Range:** 1–255 characters
- **Default:** Reset

form-submit-label —Label displayed in the button that the user selects to submit their login information —for example, Log In.

- **Range:** 1–255 characters
- **Default:** Log In

header-bgcolor—The hexadecimal color code for the color of the header bar across the top of the captive portal login page.

- **Values:** # symbol followed by six characters.

header-logo—Filename of the file containing the image of the logo displayed at the top of the captive portal login page. The image file can be in GIF, JPEG, or PNG format.

- **Default:** The Juniper Networks logo

header-message—Text displayed in the header bar across the bottom of the captive portal login page.

- **Range:** 1–2047 characters
- **Default:** User Authentication

header-text-color—Color of the text in the header.

- **Default:** The default color is white.

post-authentication-url—URL to which the users are directed upon successful authentication—for example www.mycafe.com.

- **Range:** 1–255 characters
- **Default:** The page originally requested by the user.

Required Privilege Level

routing—To view this statement in the configuration.

routing–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *Designing a Captive Portal Authentication Login Page on Switches*
>
> *Configuring Captive Portal Authentication (CLI Procedure)*

custom-page

IN THIS SECTION

- Syntax | 4777
- Hierarchy Level | 4777
- Description | 4777
- Options | 4777
- Required Privilege Level | 4778
- Release Information | 4778

Syntax

```
custom-page {
custom-page-file custom-page-file;
}
```

Hierarchy Level

```
[edit security utm custom-objects custom-message name type]
```

Description

You can create your own custom response pages for a URL that is configured with block or quarantine actions in the Content Security Web filtering profile. The custom response pages can include predefined page variables, your corporate branding, acceptable use policies, and links to your internal resources.

Options

- *custom-page-file*—Name of custom page file.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

custom-page-file | 4778

custom-page-file

IN THIS SECTION

- Syntax | 4778
- Hierarchy Level | 4779
- Description | 4779
- Options | 4780
- Required Privilege Level | 4780
- Release Information | 4780

Syntax

```
custom-page-file custom-page-file;
```

Hierarchy Level

```
[edit security utm custom-objects custom-message]
[edit logical-systems name security utm custom-objects custom-message],
[edit tenants name security utm custom-objects custom-message]
```

Description

You can create your own custom response pages based on a predefined template HTML file for a URL that is configured with block or quarantine actions in the Content Security Web filtering profile. By default, Your security device provides a predefined HTML template file, which is saved at the local folder `/var/db/utm/ custom_page/wf_block_page_tmpl.html`. You can update the following parameters in the HTML file based on your requirements:

- Change the title and text content.
- Remove one or all lines of variables.
- Add links to your internal resources to communicate your specific acceptable use policies and corporate branding.

The maximum size of the page file is limited to 2K bytes. We recommend you not add JavaScript in the contents of the HTML template page, because the JavaScript will have full access to the cookies of the blocked website.

If you accidentally change or delete the HTML file does not affect the response page until you reboot the system or execute the `request security utm web-filtering custom-page reload` command.

If you configure a non-existent HTML file, the system displays a warning message when you commit the configuration. The configuration will be successfully submitted, but this configuration will not be available. The system displays the warning message every time you commit the configuration.

You cannot upload an HTML file for a user logical system. Therefore, you can only use the existing HTML files in the device until Jweb supports HTML file upload for a user logical systems.

In the chassis cluster scenario, you need to ensure that there is the same HTML file in the specified path of each node.

If the custom-message configuration is invalid or you do not configure an HTML file, the system displays the default block page for a blocked or quarantined webpage.

Options

custom-page-file	Name of custom page file.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

custom-page | 4777

custom-tag-string

IN THIS SECTION

- Syntax | 4780
- Hierarchy Level | 4781
- Description | 4781
- Required Privilege Level | 4781
- Release Information | 4781

Syntax

```
custom-tag-string [string];
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-spam sbl profile profile-name]
```

Description

Configure a custom string for identifying a message as spam.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

custom-url-category

IN THIS SECTION

- Syntax | 4782
- Hierarchy Level | 4782
- Description | 4782
- Options | 4782
- Required Privilege Level | 4782
- Release Information | 4782

Syntax

```
custom-url-category object-name {
    value [value];
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm custom-objects]
```

Description

Use URL pattern lists to create Custom URL category lists. These are lists of patterns that bypass scanning.

WARNING: Custom category does not take precedence over predefined categories when it has the same name as one of the predefined categories. We do not recommend having a custom category name be the same as the predefined category name.

Options

- *object-name*—Name of the URL category-list object.
- `value` *value*—Value of the URL category-list object. You can configure multiple values separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Unified Threat Management Overview

customer-vlans

IN THIS SECTION

- Syntax | 4783
- Hierarchy Level | 4783
- Description | 4783
- Options | 4783
- Required Privilege Level | 4784
- Release Information | 4784

Syntax

```
customer-vlans (id | native | range);
```

Hierarchy Level

```
[edit vlans vlan-name dot1q-tunneling]
```

Description

Limit the set of accepted customer VLAN tags to a range or to discrete values when mapping customer VLANs to service VLANs.

Options

id—Numeric identifier for a VLAN.

native—Accepts untagged and priority-tagged packets from access interfaces and assigns the configured S-VLAN to the packet.

range—Range of numeric identifiers for VLANs. On the QFX series, you can include as many as eight separate customer VLAN ranges for a given service VLAN. Do not configure more than this number of ranges.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Option **native** introduced in Junos OS Release 9.6 for EX Series switches.

RELATED DOCUMENTATION

- dot1q-tunneling | **5915**
- ether-type (VLAN) | **6474**
- *Example: Setting Up Q-in-Q Tunneling on EX Series Switches*
- *Configuring Q-in-Q Tunneling on EX Series Switches*
- *Understanding Q-in-Q Tunneling and VLAN Translation*

customization (SNMP)

IN THIS SECTION

- Syntax | **4785**
- Hierarchy Level | **4785**
- Description | **4785**
- Options | **4785**
- Required Privilege Level | **4785**

Release Information | 4785

Syntax

```
customization {
    ether-stats-ifd-only;
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Customize SNMP behavior.

Options

ether-stats-ifd-only Stop exposing logical interfaces (IFLs) in the etherStatsTable. When this statement is configured, the output for the `show snmp mib walk etherStatsTable` command displays data for physical interfaces (IFDs) only.

Required Privilege Level

snmp

Release Information

Statement introduced in Junos OS Release 18.4R1.

cut-through

IN THIS SECTION

- Syntax | 4786
- Hierarchy Level | 4786
- Description | 4786
- Required Privilege Level | 4786

Syntax

```
cut-through;
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configures all the interfaces in the QFX series switch or QFabric to use cut-through forwarding mode instead of store-and-forward mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Layer 2 Forwarding Tables

cvss-score

IN THIS SECTION

- Syntax | 4787
- Hierarchy Level | 4787
- Description | 4787
- Options | 4788
- Required Privilege Level | 4788
- Release Information | 4788

Syntax

```
cvss-score
    {
    greater-than value;
    less-than value;
}
```

Hierarchy Level

```
[edit security idp  dynamic-attack-group              name              filters]
```

Description

The Common Vulnerability Scoring System (CVSS) score of attack is a free and open industry standard for assessing the severity of computer system security vulnerabilities. CVSS attempts to assign severity scores to vulnerabilities, allowing responders to prioritize responses and resources according to threats.

Scores range from 0 to 10, with 10 being the most severe. While mostly CVSS base score is used for determining severity, temporal and environmental scores, to factor in availability of mitigations and how widespread vulnerable systems are within an organization.

The CVSS assessment measures three areas of concern:

- Base Metrics for qualities intrinsic to a vulnerability.
- Temporal Metrics for characteristics that evolve over the lifetime of vulnerability.
- Environmental Metrics for vulnerabilities that depend on a particular implementation or environment.

A numerical score is generated for each of these metric groups.

Options

greater-than *value*

Match when CVSS score is greater than the value specified. The value is a real number and can include decimal values. For example, the value 5.5 is a valid CVSS score.

- **Range:** 0 to 10

less-than *value*

Match when CVSS score is less than the value specified.

- **Range:** 0 to 10

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 18.2R1.

cycle-time

IN THIS SECTION

- Syntax | 4789
- Hierarchy Level | 4789
- Description | 4789
- Options | 4789
- Required Privilege Level | 4789

Release Information | 4789

Syntax

```
cycle-time cycle-time-value;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
profile-name]
```

Description

Configure the time (in milliseconds) taken between back-to-back transmissions of SLA frames for a single connection.

Options

cycle-time-value—Cycle time value in milliseconds.

- **Range:** 10 through 3,600,000
- **Default:** 1000

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring an Iterator Profile

Configuring an Iterator Profile on a Switch (CLI Procedure)

D

IN THIS SECTION

- dad-proxy | **4809**
- dampen | **4811**
- dampen (Protocols MVPN) | **4813**
- dampen-options | **4814**
- damping (BFD Liveness Detection) | **4815**
- damping (Interfaces) | **4818**
- damping (Policy Options) | **4820**
- damping (Protocols BGP) | **4822**
- data-channel | **4825**
- data-cores | **4827**
- data-encapsulation | **4828**
- data-fill | **4830**
- data-fill-with zeros | **4831**
- data-flow-affinity | **4832**
- data-format | **4834**
- data-forwarding | **4835**
- data (FTP) | **4837**
- data-length | **4838**
- data-mdt-reuse | **4840**
- data-rate | **4841**
- data-record-fields | **4842**
- data-size | **4847**
- data-tlv-size | **4849**

- database | **4850**
- database-protection | **4852**
- database-replication (Subscriber Session Database) | **4854**
- datapath-debug | **4855**
- dcbx | **4858**
- dcbx-version | **4860**
- dcd | **4861**
- dcd-polarity | **4863**
- dce-options | **4864**
- ddos-protection | **4865**
- ddos-protection (DDoS) | **5137**
- deactivate | **5142**
- deactivation-attribute (AAA Profile) | **5143**
- deactivation-delay | **5144**
- dead-criteria-retries (RADIUS Server) | **5146**
- dead-peer-detection | **5147**
- dead-peer-detection (Security Group VPN Server) | **5149**
- dead-peer-detection (Services IPsec VPN) | **5151**
- deaggregate | **5153**
- debounce-interval | **5154**
- debug-level (System Services) | **5155**
- debug-log (System Services) | **5157**
- decapsulate-accept-inner-vlan | **5158**
- decapsulate (Firewall Filter) | **5160**
- decapsulate-inner-vlan | **5163**
- decompress-layer | **5165**
- decompress-layer-limit | **5166**
- decryption-failures | **5168**
- dedicated-buffer | **5170**
- dedicated-buffer-profile | **5172**
- dedicated-ukern-cpu (BFD) | **5174**
- deep-vlan-qualified-learning | **5176**

- default-action (DHCP Relay Agent Option) | **5177**
- default-action (Protocols OpenFlow) | **5179**
- default-actions | **5180**
- default-address-selection | **5182**
- default-chap-secret | **5184**
- default-client-identity (NETCONF TLS) | **5185**
- default (CoS) | **5187**
- default (CoS Host Outbound Traffic) | **5189**
- default-gateway | **5190**
- default-isid | **5191**
- default-lifetime | **5193**
- default-lifetime (Dynamic Router Advertisement) | **5194**
- default-local-policy | **5195**
- default-lsa | **5197**
- default-pap-password | **5199**
- default-peer | **5200**
- default-policy | **5202**
- default-policy (IDP) | **5203**
- default-pool (Address Pools) | **5205**
- default-prefixes (Protocols RIFT) | **5206**
- default-profile | **5208**
- default-profile (Juniper Secure Connect) | **5209**
- default-rule | **5211**
- default (Security Antivirus) | **5213**
- default (Security Antivirus Sophos Engine) | **5214**
- default (Security Content Security) | **5216**
- default (Security Web Filtering) | **5217**
- default-trusted-ca-certs (Security) | **5218**
- default-value | **5221**
- default-vpn-source | **5222**
- defaults | **5224**
- defaults (ACX series, JRR200, MX204, PTX5000, and VRR) | **5225**

- defaults (early-aging) | **5226**
- defaults (Policy Options) | **5228**
- defer-initial-multipath-build | **5229**
- definition | **5231**
- definition (Protocols OSPF) | **5234**
- degraded | **5236**
- delay | **5238**
- delay-advertise (DHCPv6) | **5239**
- delay-authentication (DHCP Relay Agent) | **5242**
- delay-buffer-rate | **5244**
- delay-buffer-rate (Dynamic Traffic Shaping) | **5246**
- delay-factor | **5248**
- delay-measurement (Protocols IS-IS) | **5249**
- delay-measurement (Protocols OSPF) | **5251**
- delay-metric | **5252**
- delay-metric (Protocols OSPF) | **5253**
- delay-offer (DHCPv4) | **5254**
- delay (PPPoE Service Name Tables) | **5257**
- delay (querier) | **5259**
- delay (responder) | **5260**
- delay-route-advertisements | **5262**
- delay-variation | **5265**
- delegate-probes | **5266**
- delegate-processing (VRRP) | **5268**
- delegate-server-processing | **5270**
- delegated-pool (DHCP Local Server) | **5271**
- delegation-cleanup-timeout | **5272**
- delegation-priority | **5274**
- delete | **5276**
- delete-after-commit (JNU Satellites) | **5278**
- delete-binding-on-renegotiation (DHCP Local Server and Relay Agent) | **5279**
- delete chassis fpc | **5281**

- delete-req | **5282**
- delete services jinsightd subscribe health-monitor | **5284**
- delimiter | **5285**
- delimiter (Access Profile) | **5286**
- delimiter (DHCP Local Server) | **5288**
- delimiter (DHCP Relay Agent) | **5290**
- delimiter (Domain Map) | **5293**
- demand-circuit (Protocols RIP) | **5294**
- demux-destination (Demux Interface) | **5296**
- demux-destination (Underlying Interface) | **5297**
- demux (Interfaces) | **5299**
- demux-options (All Demux Interfaces) | **5301**
- demux-options (Dynamic Interface) | **5302**
- demux-options (Static Interface) | **5304**
- demux-source (Demux Interface) | **5305**
- demux-source (Dynamic IP Demux Interface) | **5306**
- demux-source (Dynamic Underlying Interface) | **5308**
- demux-source (Underlying Interface) | **5309**
- demux0 (Dynamic Interface) | **5311**
- demux0 (Static Interface) | **5313**
- dense-groups | **5315**
- deny (Security Policies) | **5316**
- deny (Security SIP) | **5317**
- depth-threshold | **5319**
- description (RMON) | **5321**
- description (SNMP) | **5322**
- description (Abstracted Fabric) | **5324**
- description (Application Identification) | **5325**
- description (Applications) | **5326**
- description (Chassis Redundancy) | **5328**
- description (Fibre Channel Fabrics) | **5329**
- description (Forwarding Options) | **5330**

- description (GNF) | 5332
- description (IKE policy) | 5333
- description (Interfaces) | 5334
- description (Junos Fusion) | 5336
- description (Junos Fusion Satellite Device Cluster) | 5337
- description (Op Scripts) | 5339
- description (Protocols BGP) | 5340
- description (Protocols Layer 2 Circuit Neighbor) | 5342
- description (Protocols Layer 2 VPN) | 5343
- description (Protocols MPLS) | 5345
- description (Routing Instances) | 5346
- description (Security Address Book) | 5347
- description (Security IDP Policy) | 5349
- description (Security NAT Pool) | 5350
- description (Security NAT Rule) | 5351
- description (Security NAT Rule Set) | 5353
- description (Security Policies Next Gen Services) | 5354
- description (Security Screen) | 5355
- description (Security Zone) | 5357
- description (Services IPsec VPN) | 5358
- description (VLAN) | 5360
- deselect-on-bandwidth-failure | 5361
- designated-forwarder-election-hold-time | 5363
- destination | 5364
- destination (Accounting) | 5366
- destination-address | 5368
- destination-address (IPsec-VPN) | 5369
- destination-address (Firewall) | 5371
- destination-address-excluded | 5372
- destination-address (Flow Monitoring Logs for NAT) | 5373
- destination-address (HTTP Header Enrichment) | 5375
- destination-address (IDS MS-DPC) | 5377

- destination-address-name (NAT Next Gen Services) | **5378**
- destination-address-name (Security Destination NAT) | **5379**
- destination-address-name (Security Source NAT) | **5380**
- destination-address-name (Security Static NAT) | **5382**
- destination-address (NAT Next Gen Services) | **5383**
- destination-address (PCP) | **5384**
- destination-address-range | **5386**
- destination-address-range (HTTP Header Enrichment) | **5387**
- destination-address-range (IDS MS-DPC)) | **5389**
- destination-address-range (PCP) | **5390**
- destination-address-range (Services NAT) | **5391**
- destination-address-range (Services Stateful Firewall) | **5393**
- destination-address (Security Destination NAT) | **5394**
- destination-address (Security IDP Policy) | **5396**
- destination-address (Security Policies) | **5397**
- destination-address (Security Policies Flag) | **5399**
- destination-address (Security Source NAT) | **5400**
- destination-address (Security Static NAT) | **5401**
- destination-address (Services CoS) | **5403**
- destination-address (Services NAT) | **5404**
- destination-address (Services Stateful Firewall) | **5405**
- destination-address (Subscriber Secure Policy) | **5407**
- destination (Application Identification) | **5408**
- destination (Chassis) | **5410**
- destination-class-usage | **5411**
- destination-classes | **5412**
- destination (Diameter Network Element) | **5414**
- destination-equal-load-balancing (L2TP LAC) | **5415**
- destination (Event Policy) | **5416**
- destination-except | **5418**
- destination-header | **5419**
- destination-host | **5421**

- destination-host (Gx-Plus) | **5422**
- destination-host (OCS Partition) | **5423**
- destination-host (PCRF Partition) | **5424**
- destination-hostname | **5426**
- destination-identity-context | **5427**
- destination-identity-context-profile | **5429**
- destination-interface | **5430**
- destination (Interfaces) | **5433**
- destination-ip-address (RADIUS Snoop Segment) | **5434**
- destination-ip-based | **5435**
- destination-ip (Security Alarms) | **5437**
- destination (IPCP) | **5439**
- destination-ipv4-address | **5440**
- destination-ipv4-address (RFC 2544 Benchmarking Tests) | **5442**
- destination-ipv6-address | **5443**
- destination (L2TP) | **5445**
- destination (LRF Profile) | **5446**
- destination-mac-address | **5448**
- destination-nat | **5449**
- destination-networks | **5451**
- destination-option | **5452**
- destination-override | **5454**
- destination-path-group | **5455**
- destination-pool | **5457**
- destination-port (firewall) | **5458**
- destination-port (protocols pcep) | **5460**
- destination-port (RPM and TWAMP) | **5461**
- destination-port (applications) | **5464**
- destination-port (Applications) | **5465**
- destination-port (Flow Monitoring Logs for NAT) | **5471**
- destination-port (PCP) | **5472**
- destination-port (RADIUS Snoop Segment) | **5474**

- destination-port range | **5475**
- destination-port-range (HTTP Header Enrichment) | **5476**
- destination-port (Security Destination NAT) | **5478**
- destination-port (Security Forwarding Options) | **5479**
- destination-port (Security Signature Attack) | **5485**
- destination-port (Security Source NAT) | **5486**
- destination-port (Security Static NAT) | **5487**
- destination-port (SNMP) | **5488**
- destination-port (Subscriber Secure Policy) | **5490**
- destination-ports (HTTP Header Enrichment) | **5491**
- destination-prefix (Destination NAT Next Gen Services) | **5493**
- destination-prefix (IDS) | **5494**
- destination-prefix-ipv6 (IDS) | **5495**
- destination-prefix-list (HTTP Header Enrichment) | **5497**
- destination-prefix-list (PCP) | **5498**
- destination-prefix-list (Services AACL) | **5500**
- destination-prefix-list (Services CoS) | **5501**
- destination-prefix-list (Services IDS) | **5502**
- destination-prefix-list (Services NAT) | **5504**
- destination-prefix-list (Services Stateful Firewall) | **5505**
- destination-prefix (Security Forwarding Options) | **5507**
- destination-prefix (Services NAT) | **5508**
- destination-profile | **5509**
- destination (Protocols OSPF) | **5511**
- destination-realm (Gx-Plus) | **5512**
- destination-realm (JSRC) | **5514**
- destination-realm (OCS Partition) | **5515**
- destination-realm (PCRF Partition) | **5517**
- destination (Routing Instance) | **5518**
- destination (saegw system) | **5520**
- destination (Security Destination NAT) | **5521**
- destination (Security IP Headers Attack) | **5524**

- destination-threshold | **5525**
- destination (Tunnel Remote End) | **5526**
- destination (Tunnels) | **5528**
- destination-udp-port | **5529**
- destination-udp-port (FTI) | **5531**
- destinations | **5532**
- destinations (Event Policy) | **5533**
- destined-port | **5535**
- destruct-timeout (L2TP) | **5536**
- detect-loc | **5538**
- detection (Network Isolation Groups) | **5540**
- detection-time | **5542**
- detection-time (BFD for PIM) | **5543**
- detection-time (BFD Liveness Detection) | **5545**
- detection-time (LAG) | **5549**
- detection-time (Liveness Detection) | **5550**
- detector | **5551**
- deterministic-ar-policy | **5553**
- deterministic-nat-configuration-log-interval (Source NAT Next Gen Services) | **5555**
- deterministic-port-block-allocation | **5557**
- deterministic (Source NAT Next Gen Services) | **5559**
- device-attribute (Easy EVPN LAG Configuration) | **5560**
- device-authentication | **5563**
- device-count | **5565**
- device-count (MX Series) | **5567**
- device-count (Pseudowire Subscriber Interfaces) | **5569**
- devices | **5571**
- df-election | **5572**
- df-election-granularity | **5574**
- df-election-type | **5576**
- dh-group (services) | **5578**
- dh-group | **5579**

- dh-group (Security IKE) | **5581**
- dhcp | **5583**
- dhcp-attributes (Access IPv4 Address Pools) | **5585**
- dhcp-attributes (Access IPv6 Address Pools) | **5587**
- dhcp-attributes (Address-Assignment Pools) | **5589**
- dhcp-client | **5596**
- dhcp (DHCP Client) | **5598**
- dhcp-local-server | **5601**
- dhcp-local-server (System Services) | **5613**
- dhcp-option82 | **5619**
- dhcp-relay | **5620**
- dhcp-relay (DHCP Spoofing Prevention) | **5637**
- dhcp-relay (Easy EVPN LAG Configuration) | **5638**
- dhcp-security | **5640**
- dhcp-security (MX Series) | **5644**
- dhcp-service | **5646**
- dhcp-snooping-file (EX) | **5649**
- dhcp-snooping-file | **5650**
- dhcp-tags (Adjustment Control Profiles) | **5652**
- dhcp-trusted | **5654**
- dhcpv6-client | **5655**
- dhcpv6 (DHCP Local Server) | **5658**
- dhcpv6 (DHCP Relay Agent) | **5665**
- dhcpv6-options | **5673**
- dhcpv6-snooping-file | **5675**
- dhcpv6 (System Services) | **5677**
- diag-port-authentication | **5682**
- dial-options | **5684**
- dial-options (Dynamic Profiles) | **5685**
- dial-string | **5687**
- dialer | **5688**
- dialer-options | **5689**

- dialin | **5691**
- diameter | **5693**
- diameter-instance (Diameter Applications) | **5695**
- diameter-instance (JSRC) | **5696**
- diameter-profile (PCEF Profile) | **5698**
- diameter (Subscriber Aware Policy Control) | **5699**
- diameter (TDF Gateway) | **5702**
- dictionary | **5703**
- diffserv-te | **5705**
- direct-access (Commit Scripts) | **5706**
- direct-connect | **5708**
- direction (RFC 2544 Benchmarking Tests) | **5709**
- direction (services) | **5711**
- direction (Secure-channel) | **5713**
- direction (Application Identification) | **5715**
- direction (forwarding-class-accounting) | **5716**
- direction (Junos-FIPS Software) | **5718**
- direction (Junos OS) | **5719**
- direction (MEP) | **5721**
- direction (MX Series) | **5722**
- direction (Security Custom Attack) | **5724**
- direction (Security Dynamic Attack Group) | **5726**
- direction (Service Data Flow Filters) | **5727**
- director-device (Aliases) | **5729**
- disable (Protocols) | **5731**
- disable | **5736**
- disable (Routing Graceful Restart) | **5738**
- disable (STP) | **5739**
- disable-4byte-as | **5742**
- disable-action | **5743**
- disable-all-instances | **5744**
- disable (Application Tracking) | **5745**

- disable (BGP Graceful Restart) | **5747**
- disable-calling-number-avp (L2TP LAC) | **5748**
- disable (DCBX) | **5749**
- disable (Dynamic IGMP) | **5751**
- disable (Dynamic MLD) | **5752**
- disable (Extensible Subscriber Services Manager) | **5753**
- disable-failover-protocol (L2TP) | **5755**
- disable (Forwarding Options) | **5756**
- disable-fpc (DDoS) | **5758**
- disable-global-timeout-override | **5760**
- disable (GMPLS) | **5761**
- disable-grant-bypass | **5762**
- disable (IGMP Snooping) | **5764**
- disable (Interface) | **5765**
- disable (Link Protection) | **5767**
- disable-logging (DDoS) | **5769**
- disable (Long-Lived Graceful Restart for BGP Restarter) | **5770**
- disable-mlppp-inner-ppp-pfc | **5772**
- disable (Multicast Load Balancing) | **5773**
- disable (Multicast Snooping) | **5775**
- disable (MVRP) | **5776**
- disable-natt (Services IPsec VPN) | **5777**
- disable-notification-extensions (BGP Graceful Restart) | **5779**
- disable-notification-flag (BGP Graceful Restart) | **5780**
- disable (PIM) | **5782**
- disable (PoE) | **5784**
- disable (Power over Ethernet) | **5785**
- disable-preceding-key | **5787**
- disable (Protocols BGP) | **5788**
- disable (Protocols LDP) | **5789**
- disable (Protocols MLD) | **5791**
- disable (Protocols MLD Snooping) | **5792**

- disable (Protocols MPLS) | **5793**
- disable (Protocols MSDP) | **5795**
- disable (Protocols Router Discovery) | **5797**
- disable (Protocols RSVP) | **5798**
- disable (Protocols SAP) | **5800**
- disable-relay | **5801**
- disable-routing-engine (DDoS) | **5802**
- disable (Routing Options) | **5804**
- disable-signature-check | **5805**
- disable-timeout | **5806**
- disable-timeout (BPDU) | **5808**
- disable-timeout (Port Error Disable) | **5810**
- disable-https-filtering | **5812**
- disable-packet-register | **5813**
- discard (forwarding-options) | **5814**
- discard (routing-options) | **5816**
- discard-action-for-unresolved-redir-addr | **5818**
- discard (Forwarding Class) | **5819**
- disconnect-peer-timeout | **5820**
- disk-failure-action | **5821**
- disk-failure-action | **5823**
- disk-partition | **5824**
- display-host (Security Fallback Block) | **5826**
- display-host (Security Virus Detection) | **5827**
- distinguished-name | **5828**
- distinguished-name (Access) | **5830**
- distinguished-name (Security) | **5831**
- distributed-dr | **5833**
- distributed (IGMP) | **5834**
- distribution-ip (Security SIP) | **5836**
- distribution-profile | **5838**
- dlb | **5841**

- dlci | **5843**
- dlv | **5844**
- dne-origin (Diameter Network Element) | **5846**
- dns | **5847**
- dns-alg-pool | **5849**
- dns-alg-prefix | **5850**
- dns-cache | **5852**
- dns-filter | **5853**
- dns-filter-template | **5856**
- dns-filtering | **5859**
- dns-proxy | **5862**
- dns-server-address | **5864**
- dns-server-address (Dynamic Profiles) | **5865**
- dns (System Services) | **5867**
- dnssec | **5869**
- do-not-fragment | **5871**
- dod-request-policy | **5873**
- domain (DNS Forwarding Options) | **5874**
- domain (Precision Time Protocol) | **5875**
- domain (Domain Map) | **5877**
- domain-name | **5879**
- domain-name (System) | **5880**
- domain-name (Access Profile) | **5881**
- domain-name (DHCP) | **5882**
- domain-name (DHCP Local Server) | **5884**
- domain-name (DHCP Relay Agent) | **5886**
- domain-name-server-inet (Routing Instances and Access Profiles) | **5888**
- domain-name-server-inet6 (Routing Instances and Access Profiles) | **5890**
- domain-name-server (Routing Instances and Access Profiles) | **5891**
- domain-name (Static Subscribers) | **5892**
- domain-path-id | **5894**
- domain-path-id (Protocols BGP) | **5898**

- domain-search | **5899**
- domain-selection | **5901**
- domain (TDF Domain Selection) | **5905**
- domain-type (Bridge Domains) | **5906**
- domains | **5908**
- dont-help-shared-fate-bfd-down | **5912**
- dot1p-priority | **5913**
- dot1q-tunneling | **5915**
- dot1x | **5917**
- dot1x (MX Series in Enhanced LAN Mode) | **5920**
- downlink-dscp-to-egress-forwarding-class | **5922**
- downlink-dscp-to-forwarding-class | **5923**
- ignore-server-validation (Application Identification) | **5924**
- download-profile (Security Antivirus FTP) | **5926**
- download-profile (Security Content Filtering FTP) | **5927**
- download (Services) | **5928**
- download-timeout | **5930**
- downstream | **5932**
- downstream-mode | **5933**
- downstream-on-demand | **5934**
- downstream-rate (Traffic Shaping) | **5935**
- dr-election-on-p2p | **5937**
- dr-register-policy | **5938**
- drain | **5940**
- draining (Diameter Applications) | **5941**
- draining-response-timeout (Diameter Applications) | **5943**
- drop (BPDU Block) | **5944**
- drop (DHCP Relay Agent Option) | **5946**
- drop-flow | **5947**
- drop-member-traffic (Aggregated Multiservices) | **5949**
- drop-policy (Subscriber Secure Policy) | **5950**
- drop (PPPoE Service Name Tables) | **5952**

- drop-probability | **5953**
- drop-probability (Fill Level) | **5955**
- drop-probability (Interpolated Value) | **5956**
- drop-probability (Percentage) | **5958**
- drop-profile (Schedulers) | **5959**
- drop-profile (Dynamic Schedulers) | **5960**
- drop-profile-map (Dynamic Schedulers) | **5962**
- drop-profile-map (Schedulers) | **5964**
- drop-profile-map (Schedulers) | **5965**
- drop-profile (Schedulers) | **5966**
- drop-profiles (Definition) | **5968**
- drop (Security GTP) | **5969**
- drop (Security SCTP) | **5974**
- drop-threshold | **5978**
- drop-threshold (Security GPRS) | **5980**
- drop-timeout | **5981**
- drop-timeout (Forwarding Class) | **5983**
- drop-unknown-mac-ip | **5984**
- ds-lite | **5986**
- dscp (Class of Service) | **5988**
- dscp-code-point (CoS Host Outbound Traffic) | **5991**
- dscp-code-points (RFC 2544 Benchmarking) | **5993**
- dscp-code-points (RPM) | **5995**
- dscp (CoS Classifiers) | **5998**
- dscp (CoS Interfaces) | **5999**
- dscp (Dynamic Classifiers) | **6000**
- dscp (Dynamic Rewrite Rules) | **6002**
- dscp (flow-server) | **6003**
- dscp (Input Congestion Notification) | **6005**
- dscp-ipv6 (Class of Service) | **6007**
- dscp-ipv6 (Class of Service) | **6009**
- dscp-ipv6 (Dynamic Classifiers) | **6013**

- dscp-ipv6 (Dynamic Rewrite Rules) | **6014**
- dscp-mapping-classifier | **6016**
- dscp (Multifield Classifier) | **6017**
- dscp-rewrite | **6019**
- dscp (Rewrite Rules on Interfaces) | **6020**
- dscp (Services CoS) | **6022**
- dscp (Subscriber Secure Policy) | **6023**
- dsl (Access-Line Rate Adjustment) | **6025**
- dsl-sfp-options | **6030**
- dslite-softwire-initiator | **6033**
- dsr | **6035**
- dsr-polarity | **6037**
- dtcp-only (System Services) | **6038**
- dte-options | **6039**
- dtr | **6041**
- dtr-circuit | **6042**
- dtr-polarity | **6044**
- dual-dr | **6045**
- dual-phase-bootup | **6047**
- dual-stack (DHCP Local Server Overrides) | **6049**
- dual-stack (DHCP Relay Agent Overrides) | **6050**
- dual-stack-group (DHCP Local Server) | **6052**
- dual-stack-group (DHCP Relay Agent) | **6055**
- dual-stack-interface-client-limit (DHCP Local Server and Relay Agent) | **6058**
- dual-transport | **6059**
- dualstack-support (JSRC) | **6061**
- dump-on-flow-control | **6063**
- duplicate-clients (DHCPv6 Local Server and Relay Agent) | **6064**
- duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent) | **6066**
- duplicate-mac-detection | **6068**
- duplicate-protection (Dynamic PPPoE) | **6071**
- duplicates-dropped-periodicity | **6072**

- duplication (Access Profile) | **6074**
- duplication-filter (Access Profile) | **6075**
- duplication-vrf (Duplicate Accounting) | **6077**
- duration | **6078**
- duration (PoE) | **6080**
- dvmrp | **6081**
- dynamic (services) | **6083**
- dynamic | **6084**
- dynamic-address | **6086**
- dynamic-address | **6088**
- dynamic-application | **6090**
- dynamic-application-group | **6092**
- dynamic-application (Security) | **6093**
- dynamic-application (Security Policies) | **6095**
- dynamic-attack-group | **6097**
- dynamic-attack-groups (Security IDP) | **6100**
- dynamic-bidirectional-transport | **6101**
- dynamic-class-of-service-options (Dynamic Traffic Shaping) | **6102**
- dynamic-db | **6105**
- dynamic-dns | **6106**
- dynamic-filter | **6108**
- dynamic-flow-capture | **6109**
- dynamic-peers | **6111**
- dynamic-policy-control | **6112**
- dynamic-pool | **6114**
- dynamic-profile (Demux) | **6115**
- dynamic-profile (DHCP Local Server) | **6117**
- dynamic-profile (DHCP Relay Agent) | **6119**
- dynamic-profile (Dynamic Access-Line-Identifier VLANs) | **6121**
- dynamic-profile (Dynamic ACI VLANs) | **6122**
- dynamic-profile (Dynamic PPPoE) | **6124**
- dynamic-profile (L2TP) | **6126**

- dynamic-profile-options | 6127
- dynamic-profile (PPP) | 6128
- dynamic-profile (PPPoE Service Name Tables) | 6129
- dynamic-profile (Stacked VLAN) | 6131
- dynamic-profile (VLAN) | 6133
- dynamic-profiles | 6134
- dynamic-requests | 6148
- dynamic-requests-secret (RADIUS Server) | 6149
- dynamic (Security) | 6151
- dynamic-server | 6154
- dynamic (Source Packet Routing) | 6156
- dynamic-tunnel-gre-key | 6158
- dynamic-tunnel-reassembly | 6159
- dynamic-tunnels | 6161
- dynamic-vpn | 6163

dad-proxy

IN THIS SECTION

- Syntax | 4810
- Hierarchy Level | 4810
- Description | 4810
- Options | 4810
- Required Privilege Level | 4810
- Release Information | 4810

Syntax

```
dad-proxy {
 interface-restricted;
 interface-unrestricted;
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit name family inet6],
[edit dynamic-profiles name logical-systems name interfaces name unit name family inet6],
[edit interfaces name unit name family inet6]
```

Description

DAD proxy on interface

Options

interface-restricted Enable DAD interface proxy restricted to interface

interface-unrestricted Enable DAD interface proxy unrestricted to interface

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1

Option `interface-unrestricted` introduced in Junos OS Release 22.1R1 (on SRX Series and vSRX Virtual Firewall 3.0 platform).

NOTE: You cannot configure both interface restricted and interface unrestricted modes on the same interface simultaneously.

dampen

IN THIS SECTION

- Syntax | 4811
- Hierarchy Level | 4811
- Description | 4812
- Options | 4812
- Required Privilege Level | 4812
- Release Information | 4812

Syntax

```
dampen {
    dampen-options {
        cpu-factor cpu-factor;
        line-interval line-interval;
        time-interval time-interval;
    }
}
```

Hierarchy Level

```
[edit event-options event-script],
[edit event-options event-script file filename],
[edit system scripts commit],
[edit system scripts op],
[edit system scripts op file filename]
```

Description

Slow down the operation of SLAX and XSLT scripts executed on the device.

To dampen all enabled scripts of a given script type, configure the statement at the hierarchy level for that script type. To dampen a single op or event script, configure the statement at the hierarchy level where the script is configured. Script dampening is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the `op url` command, Junos OS does not dampen the script. You cannot dampen individual commit scripts.

> **NOTE**: Junos OS does not support script dampening for Python automation scripts. If you configure script dampening for all enabled scripts of a given script type, Junos OS logs the following message in the system log file if it executes a Python op or event script: `dampen is not supported for python scripts`. No warning is logged for Python commit scripts.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Dampen Script Execution

dampen-options | **4814**

event-script (Event Options) | **6533**

dampen (Protocols MVPN)

IN THIS SECTION

- Syntax | 4813
- Hierarchy Level | 4813
- Description | 4813
- Required Privilege Level | 4814
- Release Information | 4814

Syntax

```
dampen minutes
```

Hierarchy Level

```
[edit logical-systems logical-system--name protocols mvpn mvpn-mode spt-only source-active-advertisement],
[edit logical-systems logical-system--name routing-instances instance-name protocols mvpn mvpn-mode spt-only source-active-advertisement],
[edit routing-instances protocols mvpn mvpn-mode spt-only source-active-advertisement],
[edit routing-instances instance-name protocols mvpn mvpn-mode spt-only source-active-advertisement]
```

Description

Time to wait before re-advertising the source-active route (1 to 30 minutes). After traffic on the ingress PE falls below the threshold set for "min-rate" on page 10718, this is length of time that resuming traffic must continue to exceed the `min-rate` before the ingress PE can start re-advertising Source-Active A-D routes.

The default is 1 minute.

To verify that the value is set as expected, you can check whether the Type 5 (Source-Active route) has been advertised using the `show route table vrf.mvpn.0` command. It may take several minutes before you can see the changes in the Source-Active A-D route advertisement after making changes to the `min-rate`.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs

dampen-options

IN THIS SECTION

- Syntax | 4814
- Hierarchy Level | 4815
- Description | 4815
- Options | 4815
- Required Privilege Level | 4815
- Release Information | 4815

Syntax

```
dampen-options {
    cpu-factor cpu-factor;
    line-interval line-interval;
    time-interval time-interval;
}
```

Hierarchy Level

```
[edit event-options event-script dampen],
[edit event-options event-script file filename dampen],
[edit system scripts commit dampen],
[edit system scripts op dampen],
[edit system scripts op file filename dampen]
```

Description

Slow down the operation of automation scripts executed on the device.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Dampen Script Execution

damping (BFD Liveness Detection)

IN THIS SECTION

- Syntax | 4816
- Hierarchy Level | 4816

- Description | 4816
- Options | 4817
- Required Privilege Level | 4817
- Release Information | 4817

Syntax

```
damping {
 half-life-interval seconds;
 increment increment;
 max-suppress-time seconds;
 reuse reuse;
 suppress suppress;
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options bfd-liveness-detection],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-ether-options bfd-
liveness-detection],
[edit interfaces name aggregated-ether-options bfd-liveness-detection]
```

Description

Enable BFD session damping to dampen excess flap notifications from a BFD session that is going up and down repeatedly. BFD session damping using a merit system where each flap adds to a merit value, and once a BFD session reaches the configured merit value, notifications will be dampened. The merit value decreases by half after a configured time period, and when it falls below a configured threshold, BFD state change notifications will resume.

BFD session damping happens independently on each router locally, so BFD session configuration values should match on both ends of a BFD session to ensure consistent behavior.

Options

half-life-interval

Time duration before the accumulated merit value for a BFD session is reduced by half.

- **Default:** 30 seconds
- **Range:** 1 through 600 seconds

increment

Increments applied to merit value for every flap

- **Default:** 1000
- **Range:** 1000 through 20000

max-suppress-time

Maximum amound of time time a BFD session can be suppressed no matter how unstable it has been prior to this period of stability.

- **Default:** 600 seconds
- **Range:** 1 through 20000 seconds

reuse

Maximum merit value that will allow a suppressed BFD session to start notifications again.

- **Default:** 1000
- **Range:** 1000 through 200000

suppress

Minimum merit value that will start suppressing BFD notifications.

- **Default:** 5000
- **Range:** 1000 through 200000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

damping (Interfaces)

IN THIS SECTION

- Syntax | 4818
- Hierarchy Level | 4818
- Description | 4818
- Options | 4819
- Required Privilege Level | 4820
- Release Information | 4820

Syntax

```
damping {
    enable;
    half-life seconds;
    max-suppress seconds;
    reuse number;
    suppress number;
}
```

Hierarchy Level

```
[edit interfaces interface--name],
[edit interfaces interface--range]
```

Description

Limit the number of advertisements of the up and down transitions (flapping) on an interface. Each time a transition occurs, the interface state is changed, which generates an advertisement to the upper-level routing protocols. Damping helps reduce the number of these advertisements. Every time an interface goes down, a penalty is added to the interface penalty counter. Penalty added on every interface flap is 1000.

If at some point the accumulated penalty exceeds the suppress level `max-suppress`, the interface is placed in the suppress state, and further interface state up and down transitions are not reported to the upper-level protocols.

Options

`enable`—Enable damping on a per-interface basis. If damping is enabled on an interface, it is suppressed during interface flaps that match the configuration settings.

- **Default:** Disabled

`half-life` *`seconds`*—Decay half-life. *`seconds`* is the interval after which the accumulated interface penalty counter is reduced by half if the interface remains stable.

> **NOTE**: For the half-life, configure a value that is less than the max-suppress value. If you do not, the configuration is rejected.

- **Range:** 1 through 30
- **Default:** 5

`max-suppress` *`seconds`*—Maximum hold-down time. *`seconds`* is the maximum time that an interface can be suppressed no matter how unstable the interface has been.

> **NOTE**: For max-suppress, configure a value that is greater than the half-life. If you do not, the configuration is rejected.

- **Range:** 1 through 20,000
- **Default:** 20

`reuse` *`number`*—Reuse threshold. When the accumulated interface penalty counter falls below *number*, the interface is no longer suppressed.

- **Range:** 1 through 20,000
- **Default:** 1000

`suppress` *`number`*—Cutoff (suppression) threshold. When the accumulated interface penalty counter exceeds *number*, the interface is suppressed.

- **Range:** 1 through 20,000
- **Default:** 2000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Damping Shorter Physical Interface Transitions

show interfaces extensive | **20676**

hold-time (ae interface) | **8098**

damping (Policy Options)

IN THIS SECTION

- Syntax | **4820**
- Hierarchy Level | **4821**
- Description | **4821**
- Options | **4821**
- Required Privilege Level | **4822**
- Release Information | **4822**

Syntax

```
damping name {
    disable;
    half-life minutes;
    max-suppress minutes;
```

```
    reuse number;
    suppress number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options],
[edit policy-options]
```

Description

Define route flap damping properties to set on BGP routes.

Options

`disable`—Disable damping on a per-prefix basis. Any damping state that is present in the routing table for a prefix is deleted if damping is disabled.

`half-life` *`minutes`*—Decay half-life. *`minutes`* is the interval after which the accumulated figure-of-merit value is reduced by half if the route remains stable.

- **Range:** 1 through 45
- **Default:** 15 minutes

> **NOTE**: For the half-life, configure a value that is less than the max-suppress. If you do not, the configuration is rejected.

`max-suppress` *`minutes`*—Maximum hold-down time. *`minutes`* is the maximum time that a route can be suppressed no matter how unstable it has been.

- **Range:** 1 through 720
- **Default:** 60 minutes

> **NOTE**: For the max-suppress, configure a value that is greater than the half-life. If you do not, the configuration is rejected.

name—Name that identifies the set of damping parameters. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

`reuse` *number*—Reuse threshold. *number* is the figure-of-merit value below which a suppressed route can be used again.

- **Range:** 1 through 20,000
- **Default:** 750 (unitless)

`suppress` *number*—Cutoff (suppression) threshold. *number* is the figure-of-merit value above which a route is suppressed for use or inclusion in advertisements.

- **Range:** 1 through 20,000
- **Default:** 3000 (unitless)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring BGP Flap Damping Parameters

Example: Configuring BGP Route Flap Damping Parameters

Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family

damping (Protocols BGP)

IN THIS SECTION

Syntax | 4823

- Hierarchy Level | 4823
- Description | 4824
- Default | 4824
- Required Privilege Level | 4824
- Release Information | 4824

Syntax

```
damping;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp family family],
[edit logical-systems logical-system-name protocols bgp family family],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name family family],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
family],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
```

```
group group-name neighbor address family family],
[edit protocols bgp],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name family family],
[edit protocols bgp group group-name neighbor address],
[edit protocols bgp group group-name neighbor address family family],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp family family],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name family family],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family family]
```

Description

Enable route flap damping. BGP route flapping describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability information. Flap damping reduces the number of update messages sent between BGP peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.

You typically apply flap damping to external BGP (EBGP) routes (that is, to routes in different ASs). You can also apply it within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.) The exception to this rule is when flap damping is applied at the address family level. When you apply flap damping at the address family level, it works for both IBGP and EBGP.

Default

Flap damping is disabled on the routing device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for flap damping at the address family level introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Examples: Configuring BGP Flap Damping

Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family

data-channel

IN THIS SECTION

- Syntax | 4825
- Hierarchy Level | 4825
- Description | 4826
- Options | 4826
- Required Privilege Level | 4826
- Release Information | 4826

Syntax

```
data-channel {
    vlan number;
}
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

For Ethernet ring protection, configure a data channel to define a set of VLAN IDs that belong to a ring instance.

VLANs specified in the data channel use the same topology used by the ERPS PDU in the control channel. Therefore, if a ring interface is blocked in the control channel, all traffic in the data channel is also blocked on that interface.

Options

`vlan` *`number`*—Specify (by VLAN ID) one or more VLANs that belong to a ring instance.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

RELATED DOCUMENTATION

- Ethernet Ring Protection Using Ring Instances for Load Balancing
- Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

data-cores

IN THIS SECTION

- Syntax | 4827
- Hierarchy Level | 4827
- Description | 4827
- Options | 4827
- Required Privilege Level | 4828
- Release Information | 4828

Syntax

```
data-cores data-number;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-
provider]
```

Description

Configure data cores. Any cores not configured as either data or control cores are treated as user cores. When the number of data cores is changed, the PIC reboots.

Options

data-number—Number of data cores. Although it is not mandatory to dedicate any cores as data cores, it is advisable, depending on the nature of the application, to dedicate a minimum of five as data cores to achieve good performance.

- **Range:** 0 through 7

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

control-cores | 4631

data-encapsulation

IN THIS SECTION

- Syntax | 4828
- Hierarchy Level | 4829
- Description | 4829
- Default | 4829
- Options | 4829
- Required Privilege Level | 4829
- Release Information | 4829

Syntax

```
data-encapsulation (disable | enable);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
msdp],
[edit protocols msdp],
[edit routing-instances routing-instance-name protocols msdp]
```

Description

Configure a rendezvous point (RP) using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages.

Default

If you do not include this statement, the RP encapsulates multicast data.

Options

disable—(Optional) Do not use MSDP data encapsulation.

enable—Use MSDP data encapsulation.

- **Default: enable**

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Example: Configuring MSDP with Active Source Limits and Mesh Groups*

data-fill

IN THIS SECTION

- Syntax | 4830
- Junos OS Hierarchy Levels | 4830
- Junos OS Evolved Hierarchy Level | 4830
- Description | 4830
- Options | 4831
- Required Privilege Level | 4831
- Release Information | 4831

Syntax

```
data-fill data;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],
[edit servicesrpm probe owner test test-name],
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes. The `data-fill` statement is not valid with the `http-get` or `http-metadata-get` probe types.

Options

data—A hexadecimal value; for example, `0-9`, `A-F`.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Configuring BGP Neighbor Discovery Through RPM

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

data-fill-with zeros

IN THIS SECTION

- Syntax | 4832
- Hierarchy Level | 4832
- Description | 4832
- Required Privilege Level | 4832
- Release Information | 4832

Syntax

```
data-fill-with-zeros;
```

Hierarchy Level

```
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Description

If this statement is configured, then the contents of the test packet are zeros, if the statement is not configured, then the data content is a pseudo-random number.

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

data-flow-affinity

IN THIS SECTION

- Syntax | 4833
- Hierarchy Level | 4833

Description | 4833
Required Privilege Level | 4833
Release Information | 4833

Syntax

```
data-flow-affinity {
    hash-key (layer-3 | layer-4);
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-
provider]
```

Description

Enable flow affinity distribution for packets over data CPUs on the PIC. Once enabled, the default behavior distributing data packets changes from a round-robin distribution to a flow affinity distribution based on a hash distribution. Adding or deleting this statement causes the PIC to reboot.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

data-format

IN THIS SECTION

- Syntax | 4834
- Hierarchy Level | 4834
- Description | 4834
- Options | 4834
- Required Privilege Level | 4834
- Release Information | 4835

Syntax

```
data-format format;
```

Hierarchy Level

```
[edit services flow-collector file-specification variant variant-number]
```

Description

Specify the data format for a specific file format variant.

Options

format—Data format. Specify **flow-compressed** as the data format.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring File Formats

data-forwarding

IN THIS SECTION

- Syntax | 4835
- Hierarchy Level | 4836
- Description | 4836
- Default | 4836
- Required Privilege Level | 4836
- Release Information | 4836

Syntax

```
data-forwarding {
    receiver {
        install;
        mode (proxy | transparent);
        (source-list | source-vlans) vlan-list;
        translate;
    }
    source {
        groups group-prefix;
    }
}
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan vlan-name]
[edit logical-systems name protocols igmp-snooping vlan vlan-name]
```

Description

Configure a data-forwarding VLAN as a multicast source VLAN (MVLAN) or a receiver VLAN using the multicast VLAN registration (MVR) feature.

You can configure a data-forwarding VLAN as either a multicast source VLAN (an MVLAN) or a multicast receiver VLAN (an MVR receiver VLAN), but not both.

- When you configure an MVR receiver VLAN, you must also configure the MVLANs you list as source VLANs for that MVR receiver VLAN.
- When you configure a source MVLAN, you aren't required to set up MVR receiver VLANs at the same time; you can configure those later.

MVR is only supported with IGMP version 2 (IGMPv2).

> **NOTE**: The `mode`, `source-list`, and `translate` statements are only applicable to MVR configuration on EX Series switches that support the Enhanced Layer 2 Software (ELS) configuration style. The `source-vlans` statement is applicable only to EX Series switches that do not support ELS, and is equivalent to the ELS `source-list` statement.

The `receiver`, `source`, and `mode` statements and options are explained separately. See CLI Explorer.

Default

Disabled

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Multicast VLAN Registration

Example: Configuring Multicast VLAN Registration on EX Series Switches Without ELS

Configuring Multicast VLAN Registration on EX Series Switches

data (FTP)

IN THIS SECTION

- Syntax | 4837
- Hierarchy Level | 4837
- Description | 4837
- Default | 4838
- Required Privilege Level | 4838
- Release Information | 4838

Syntax

```
data {
    dscp (alias | bits);
    forwarding-class class-name;
}
```

Hierarchy Level

```
[edit services cos application-profile profile-name ftp]
```

Description

Set the appropriate `dscp` and `forwarding-class` value for FTP data.

Default

By default, the system will not alter the DSCP or forwarding class for FTP data traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

video (Application Profile) | **30728**

voice (Application Profile) | **30917**

data-length

IN THIS SECTION

- Syntax | **4839**
- Hierarchy Level | **4839**
- Description | **4839**
- Options | **4839**
- Required Privilege Level | **4839**
- Release Information | **4839**

Syntax

```
data-length {
    match (equal | greater-than | less-than | not-equal);
    value tcp-data-length;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol udp],

[edit security idp custom-attack attack-name attack-type signature protocol icmp],

[edit security idp custom-attack attack-name attack-type signature protocol icmpv6],

[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify the number of bytes in the data payload. In the TCP header, for SYN, ACK, and FIN packets, this field should be empty.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`data-length`*—Match the number of bytes in the data payload.
- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

data-mdt-reuse

IN THIS SECTION

- Syntax | 4840
- Hierarchy Level | 4840
- Description | 4840
- Required Privilege Level | 4840
- Release Information | 4841

Syntax

```
data-mdt-reuse;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel  pim mdt],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel family inet / inet6 mdt],
[edit routing-instances routing-instance-name provider-tunnel family inet / inet6 mdt]
```

Description

Enable dynamic reuse of data MDT group addresses.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. In Junos OS Release 17.3R1, the `mdt` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default MDT in Rosen 7, and data MDT for Rosen 6 and Rosen 7. The `provider-tunnel mdt` hierarchy is now hidden for backward compatibility with existing scripts.

RELATED DOCUMENTATION

Example: Enabling Dynamic Reuse of Data MDT Group Addresses

data-rate

IN THIS SECTION

- Syntax | 4841
- Hierarchy Level | 4841
- Description | 4842
- Options | 4842
- Required Privilege Level | 4842
- Release Information | 4842

Syntax

```
data-rate weight;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number radio-router]
```

Description

Configure the weight of the resource factor when calculating an effective data rate.

Options

weight—Factor used to calculate data rate.

- **Range:** 0 through 100
- **Default:** 100

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.2 of Junos OS .

RELATED DOCUMENTATION

Configuring PPPoE-Based Radio-to-Router Protocols

data-record-fields

IN THIS SECTION

- Syntax | 4843
- Hierarchy Level | 4843
- Description | 4843
- Default | 4844
- Options | 4844
- Additional Information | 4845

Required Privilege Level | 4846
Release Information | 4846

Syntax

```
data-record-fields {
        include-policy-map-name;
        source-prefix-as-path count;
        destination-prefix-as-path count;
        bgp-source-standard-community count;
        bgp-destination-standard-community count;
        bgp-source-extended-community count;
        bgp-destination-extended-community count;
        bgp-source-large-community count;
        bgp-destination-large-community count;
}
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Use this configuration statement to add Information Elements to a data record template. You can add Information Elements concerning class of service (CoS) information or concerning BGP community and AS path information. Any new or changed configuration of this statement is effective immediately and the exported packets contain the new information. Each Information Element is independent of the others and can be exported separately in the collector record.

For CoS, for IPFIX and version 9, you can export the policy map name in the data record in the ingress direction. You configure policy maps at the `[edit class-of-service policy-map]` hierarchy level. See Assigning Rewrite Rules on a Per-Customer Basis Using Policy Maps to learn more about policy maps.

You can also export the BGP community and AS path information in the data record using IPFIX, per RFC 8549. The IPFIX collector should be RFC 6313 compliant to decode data records carrying fields with the basicList format.

count defines the size of the list that will be exported for that Information Element in the data record. If the total number of items exceeds the configured list size, the exported list contains the first and subsequent items up to and including the configured list size. For example, if the AS path list for the source prefix contains 8 AS numbers, but *count* is set to 4, only the first four AS numbers are exported in the data record.

Default

When this statement is not configured, the routing protocol process (rpd) only sends the peer AS number, origin AS number, and the BGP next-hop information. When this statement is configured with at least one Information Element, rpd sends all of the community and AS-path-related information. You do not have to configure all of the Information Elements.

Options

include-policy-map-name

Export the Juniper-specific enterprise Information Element ID 32765 in the data record templates `ip4-template` and `ipv6-template`. This new IE ID is 4 bytes long and contains the first 4 characters of the policy-map name. Therefore, the first 4 letters of your policy-map names should be unique. For example, if you have a policy map that sets DSCP code point 001011, you could begin your policy map name with the characters `DS11`.

source-prefix-as-path *count*

Export the source prefix AS path (Information Element 16, per RFC 6313).

- Default: 8

destination-prefix-as-path *count*

Export the destination prefix AS path (Information Element 17, per RFC 6313).

- Default: 8

bgp-source-standard-community *count*

Export the standard BGP community list for the source prefix, corresponding to a specific flow's source IP address (Information Element 484, per Appendix A of RFC 8549).

- Default: 8

bgp-destination-standard-community *count*

Export the standard BGP community list for the destination prefix, corresponding to a specific flow's destination IP address (Information Element 485, per Appendix A of RFC 8549).

- Default: 8

bgp-source-extended-community *count*

Export the extended BGP community list corresponding to a specific flow's source IP address (Information Element 487, per Appendix A of RFC 8549).

- Default: 8

bgp-destination-extended-community *count*

Export the extended BGP community list corresponding to a specific flow's destination IP address (Information Element 488, per Appendix A of RFC 8549).

- Default: 8

bgp-source-large-community *count*

Export the large BGP community list corresponding to a specific flow's source IP address (Information Element 490, per Appendix A of RFC 8549).

- Default: 8

bgp-destination-large-community *count*

Export the large BGP community list corresponding to a specific flow's destination IP address (Information Element 491, per Appendix A of RFC 8549).

- Default: 8

Additional Information

AS path and BGP community information are exported using the following Information Elements, as shown in Table 27 on page 4845:

Table 27: AS Path and BGP Community Information Elements

IE Name	IE ID	Export Data Type Format
BGP Community	483	Unsigned 32-bit value
BGP Source Community List, which is a list of BGP community numbers (IE 483)	484	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.
BGP Destination Community List, which is a list of BGP community numbers (IE 483)	485	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.
BGP Extended Community	486	8-byte octetArray
BGP Source Extended Community List, which is a list of BGP extended community numbers (IE 486)	487	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.

Table 27: AS Path and BGP Community Information Elements *(Continued)*

IE Name	IE ID	Export Data Type Format
BGP Destination Extended Community List, which is a list of BGP extended community numbers (IE 486)	488	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.
BGP Large Community	489	12-byte octetArray
BGP Source Large Community List, which is a list of BGP large community numbers (IE 489)	490	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.
BGP Destination Large Community List, which is a list of BGP large community numbers (IE 489)	491	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.
Source AS Path List	16	basicList: data is exported in the format defined in RFC 6313.
Destination AS Path List	17	basicList: data is exported in the format defined in RFC 6313.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 21.4R1.

Option `include-policy-map-name` introduced in Junos OS Evolved 24.4R1.

RELATED DOCUMENTATION

Assigning Rewrite Rules on a Per-Customer Basis Using Policy Maps

data-size

IN THIS SECTION

- Syntax | 4847
- Junos OS Hierarchy Levels | 4847
- Junos OS Evolved Hierarchy Level | 4847
- Description | 4848
- Options | 4848
- Required Privilege Level | 4848
- Release Information | 4848

Syntax

```
data-size size;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the size of the data portion of ICMP probes. The `data-size` statement is not valid with the `http-get` or `http-metadata-get` probe type.

Options

size—0 through 65400 for RPM, for TWAMP the value is from 60 through 1400.

- **Default:** 0 for RPM and 60 for TWAMP.

> **NOTE**: If you configure the hardware timestamp feature (see *Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches*):
>
> - The default value of `data-size` is 32 bytes and `32` is the minimum value for explicit configuration. The UDP timestamp probe type is an exception; it requires a minimum data size of 52 bytes.
> - The data size must be at least 100 bytes smaller than the default MTU of the interface of the RPM client interface.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.

Support at the [`edit services rpm twamp client control-connection` *`control-client-name`*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Configuring BGP Neighbor Discovery Through RPM

data-tlv-size

IN THIS SECTION

- Syntax | 4849
- Hierarchy Level | 4849
- Description | 4849
- Options | 4849
- Required Privilege Level | 4850
- Release Information | 4850

Syntax

```
data-tlv-size size;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-
association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile profile-name]
```

Description

Configure the size of the data TLV portion of the Y.1731 data frame.

Options

size—Size of the data TLV portion of the Y.1731 data frame.

NOTE: This option is applicable only for two-way delay measurement.

- **Range:** 1 through 1400 bytes
- **Default:** 1

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

sla-iterator-profile | 28012

Configuring a Remote MEP with an Iterator Profile

database

IN THIS SECTION

- Syntax | 4851
- Hierarchy Level | 4851
- Description | 4851
- Required Privilege Level | 4851
- Release Information | 4851

Syntax

```
database {
    export;
    import;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls traffic-engineering],
[edit protocols mpls traffic-engineering]
```

Description

Include link and node entries from the traffic engineering database into the `lsdist.0` routing information base (RIB), so it gets picked up by the BGP export policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

traffic-engineering (Protocols MPLS) | **30031**

database-protection

IN THIS SECTION

- Syntax | 4852
- Hierarchy Level | 4852
- Description | 4853
- Default | 4853
- Options | 4853
- Required Privilege Level | 4854
- Release Information | 4854

Syntax

```
database-protection {
    ignore-count number;
    ignore-time seconds;
    maximum-lsa number;
    reset-time seconds;
    warning-only;
    warning-threshold percent;
}
```

Hierarchy Level

```
[edit protocols (ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-unicast | ipv6-multicast)]
```

Description

Configure the maximum number of link-state advertisements (LSAs) that are not generated by the router or switch in a given OSPF instance.

Default

By default, OSPF database protection is not enabled.

Options

ignore-count ***number***—Configure the number of times the database can enter the ignore state. When the ignore count is exceeded, the database enters the isolate state.

- **Range:** 1 through 32
- **Default:** 5

ignore-time ***seconds***—Configure the time the database must remain in the ignore state before it resumes regular operations (enters retry state).

- **Range:** 30 through 3,600 seconds
- **Default:** 300 seconds

maximum-lsa ***number***—Configure the maximum number of LSAs whose advertising router ID is different from the local router ID in a given OSPF instance. This includes external LSAs as well as LSAs with any scope, such as the link, area, and autonomous system (AS). This value is mandatory.

- **Range:** 1 through 1,000,000
- **Default:** None

reset-time ***seconds***—Configure the time period during which the database must operate without being in the ignore or isolate state before it is reset to a normal operating state.

- **Range:** 60 through 86,400 seconds
- **Default:** 600 seconds

warning-only—Specify that only a warning should be issued when the maximum LSA number is exceeded. If configured, no other action is taken against the database.

warning-threshold ***percent***—Configure the percentage of the maximum number of LSAs to be exceeded before a warning message is logged.

- **Range:** 30 through 100 percent

- **Default:** 75 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

> *OSPF Database Protection Overview*
>
> *Configuring OSPF Database Protection*

database-replication (Subscriber Session Database)

IN THIS SECTION

- Syntax | 4854
- Hierarchy Level | 4855
- Description | 4855
- Required Privilege Level | 4855
- Release Information | 4855

Syntax

```
database-replication {
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
```

```
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Define operations for subscriber management session database replication processes.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Tracing Subscriber Management Session Database Replication Events for Troubleshooting

datapath-debug

IN THIS SECTION

- Syntax | 4856
- Hierarchy Level | 4857

- Description | 4857
- Options | 4857
- Required Privilege Level | 4857
- Release Information | 4857

Syntax

```
datapath-debug {
    action-profile name {
        event name {
            count;
            packet-dump;
            packet-summary;
            trace;
        }
        module name {
            flag name;
        }
        preserve-trace-order;
        record-pic-history;
    }
    capture-file (Security) filename <files files> <format pcap> <size size> <(world-readable | no-
world-readable)>;
    maximum-capture-size bytes;
    packet-filter name {
        action-profile (default | profile);
        destination-port (afs | bgp | biff | bootpc | bootps | cmd | cvspserver | dhcp | domain
| eklogin | ekshell | exec | finger | ftp | ftp-data | http | https | ident | imap | kerberos-
sec | klogin | kpasswd | krb-prop | krbupdate | kshell | ldap | ldp | login | mobileip-agent |
mobilip-mn | msdp | netbios-dgm | netbios-ns | netbios-ssn | nfsd | nntp | ntalk | ntp | pop3 |
pptp | printer | radacct | radius | range | rip | rkinit | smtp | snmp | snmptrap | snpp | socks
| ssh | sunrpc | syslog | tacacs | tacacs-ds | talk | telnet | tftp | timed | who | xdmcp |
zephyr-clt | zephyr-hm | zephyr-srv);
        destination-prefix destination-prefix;
        interface interface;
        protocol (ah | egp | esp | gre | icmp | icmp6 | igmp | ipip | number | ospf | pim | rsvp
| sctp | tcp | udp);
        source-port (afs | bgp | biff | bootpc | bootps | cmd | cvspserver | dhcp | domain |
```

```
eklogin | ekshell | exec | finger | ftp | ftp-data | http | https | ident | imap | kerberos-sec
| klogin | kpasswd | krb-prop | krbupdate | kshell | ldap | ldp | login | mobileip-agent |
mobilip-mn | msdp | netbios-dgm | netbios-ns | netbios-ssn | nfsd | nntp | ntalk | ntp | pop3 |
pptp | printer | radacct | radius | range | rip | rkinit | smtp | snmp | snmptrap | snpp | socks
| ssh | sunrpc | syslog | tacacs | tacacs-ds | talk | telnet | tftp | timed | who | xdmcp |
zephyr-clt | zephyr-hm | zephyr-srv);
        source-prefix source-prefix;
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure the data path debugging options.

> **NOTE**: Data path debugging is supported on SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Understanding Data Path Debugging for Logical Systems

dcbx

IN THIS SECTION

- Syntax | 4858
- Hierarchy Level | 4859
- Description | 4859
- Options | 4859
- Required Privilege Level | 4859
- Release Information | 4859

Syntax

```
dcbx {
    disable;
    interface (interface-name | all) {
        disable;
        application-map application-map-name;
        applications {
            no-auto-negotiation;
        }
        enhanced-transmission-selection {
            no-auto-negotiation;
            no-recommendation-tlv;
            recommendation-tlv {
                no-auto-negotiation;
            }
        }
        dcbx-version (auto-negotiate | ieee-dcbx | dcbx-version-1.01);
        priority-flow-control {
            no-auto-negotiation;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure DCBX properties. DCBX is an extension of Link Layer Discovery Protocol (LLDP), and LLDP must remain enabled on every interface for which you want to use DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit fails.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

`mode` and `recommendation-tlv` statements introduced in Junos OS Release 12.2 for the QFX Series.

RELATED DOCUMENTATION

- show dcbx neighbors | **19191**
- *Understanding DCB Features and Requirements*
- *Configuring DCBX Autonegotiation*

dcbx-version

IN THIS SECTION

- Syntax | 4860
- Hierarchy Level | 4860
- Description | 4860
- Default | 4860
- Options | 4861
- Required Privilege Level | 4861
- Release Information | 4861

Syntax

```
dcbx-version (auto-negotiate | ieee-dcbx | dcbx-version-1.01);
```

Hierarchy Level

```
[edit protocols dcbx interface (all | interface-name)]
```

Description

Set the DCBX version for the specified interface or interfaces.

QFX3500 switches come up in IEEE DCBX mode and then autonegotiate with the connected peer to set the DCBX version.

QFabric system Node devices come up using DCBX version 1.01, and then autonegotiate with the connected peer to set the DCBX mode.

Default

The default DCBX mode is autonegotation.

Options

`auto-negotiate`—Automatically negotiate the DCBX version with the connected peer.

`ieee-dcbx`—Force the interface to use IEEE DCBX mode, regardless of the peer configuration.

`dcbx-version-1.01`—Force the interface to use version 1.01 DCBX mode, regardless of the peer configuration.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

show dcbx neighbors | **19191**

Configuring DCBX Autonegotiation

Understanding DCBX

dcd

IN THIS SECTION

- Syntax | **4862**
- Hierarchy Level | **4862**
- Description | **4862**
- Options | **4862**
- Required Privilege Level | **4862**
- Release Information | **4862**

Syntax

```
dcd (ignore | normal | require);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options dce-options],
[edit interfaces interface-name serial-options dte-options]
```

Description

For EIA-530 and V.35 interfaces only, configure the from-DCE signal, data-carrier-detect (DCD).

Options

`ignore`—The from-DCE signal is ignored.

`normal`—Normal DCD signal handling as defined by the TIA/EIA Standard 530.

`require`—The from-DCE signal must be asserted.

- **Default:** `normal`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Signal Handling

dcd-polarity

IN THIS SECTION

- Syntax | 4863
- Hierarchy Level | 4863
- Description | 4863
- Options | 4863
- Required Privilege Level | 4863
- Release Information | 4864

Syntax

```
dcd-polarity (negative | positive);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

Configure DCD signal polarity.

Options

`positive`—Positive signal polarity.

`negative`—Negative signal polarity.

- **Default:** `positive`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Serial Signal Polarities

dce-options

IN THIS SECTION

- Syntax | 4864
- Hierarchy Level | 4865
- Description | 4865
- Required Privilege Level | 4865
- Release Information | 4865

Syntax

```
dce-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
}
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

For J Series Services Routers, configure the serial interface signal characteristics.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Statement previously known as `control-leads`.

RELATED DOCUMENTATION

Configuring the Serial Signal Handling

ddos-protection

IN THIS SECTION

- Syntax | 4866
- Hierarchy Level | 5136
- Description | 5136
- Required Privilege Level | 5136
- Release Information | 5137

Syntax

```
ddos-protection {
global {
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection;
 flow-detection-mode(automatic | off | on);
 flow-level-control(drop | keep | police);


 flow-report-rate reports per second;
 violation-report-rate reports per second;
}
protocols {
 all-fiber-channel-enode {
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
   flow-detect-time seconds;
   flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
```

```
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  amtv4 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
amtv6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
ancp {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ancpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
```

```
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  arp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }


    recover-time seconds;
```

```
   }
   arp {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   bcast {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ucast {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  atm {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;


  }
 }


 bfd {
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
   flow-detect-time seconds;
   flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
```

```
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    bundle-bfd {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
```

```
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    multihop-bfd {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
```

```
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  bfdv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  bgp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
```

```
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  bgpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
 }
 custom {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
demux-autosense {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  dhcpv4 {
   ack {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
```

```
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
aggregate {
 bandwidth packets per second;
 burst packets;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;

 recover-time seconds;

}
bad-packets {
 bandwidth packets per second;
```

```
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   bootp {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   decline {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    discover {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
```

```
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    force-renew {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
```

```
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    inform {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    lease-active {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
```

```
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
lease-query {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
    }
    lease-unassigned {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;



    }
    lease-unknown {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   nak {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   no-message-type {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   offer {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   rebind {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   release {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
```

```
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    renew {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
```

```
 timeout-active-flows;


}
request {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
unclassified {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  dhcpv4v6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
```

```
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  dhcpv6 {
   advertise {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;





}
confirm {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   decline {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
```

```
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    information-request {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
```

```
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   leasequery {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
```

```
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    leasequery-data {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   leasequery-done {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
```

```
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
leasequery-reply {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
}
rebind {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
```

```
reconfigure {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
relay-forward {
 bandwidth packets per second;
 burst packets;
```

```
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   relay-reply {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   release {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   renew {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    reply {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
```

```
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    request {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   solicit {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    unclassified {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  diameter {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  dns {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
```

```
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   dtcp {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
```

```
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  dynamic-vlan {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
```

```
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  egpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
```

```
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  eoam {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
```

```
 timeout-active-flows;


}
oam-cfm {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
    }
    unclassified {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   esmc {
```

```
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;



 }
}
exceptions {
 aggregate {
  bandwidth packets per second;
```

```
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    recover-time seconds;
   }
   mcast-rpf-err {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    mtu-exceeded {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
```

```
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    unclassified {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  fab-probe {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
```

```
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



   }
  }
  filter-action {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;



    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }


     recover-time seconds;





    }
    filter-v4 {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   filter-v6 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
```

```
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
other {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
 }
}
fip-snooping {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;
```

```
 }
}
firewall-host {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;



 }
}
```

```
frame-relay {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
  }
  recover-time seconds;
 }
 frf15 {
  bandwidth packets per second;
  burst packets;
  bypass-aggregate;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   frf16 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ftp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ftpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  garp-reply {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  gre {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;





}
hag {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    hbc {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
```

```
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   punt {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
```

```
    recover-time seconds;
    timeout-active-flows;


  }
 }
 gtp-path-mgmt {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }


    recover-time seconds;
```

```
  }
 }
 icmp {
  aggregate {
```

```
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  icmpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  igmp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   query {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   report {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  igmpv4v6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  igmpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    mld {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   mld-query {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
```

```
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    mld-report {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
```

```
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
```

```
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   inline-ka {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
```

```
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;


  }
 }
 inline-svcs {
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
   flow-detect-time seconds;
   flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
```

```
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ip-fragments {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }


    recover-time seconds;
```

```
   }
```

```
first-fragment {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
trail-fragment {
 bandwidth packets per second;
 burst packets;
```

```
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ipmcast-miss {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
```

```
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ipmc-reserved {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ip-options {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;




}
non-v4v6 {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   router-alert {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    unclassified {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



   }
  }
  ipsec {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
```

```
 recover-time seconds;




}
ike {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
```

```
 recover-time seconds;
 timeout-active-flows;


}
mtu-error {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
}
on-trigger {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
sn-alarm {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   spi-inval {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   tx-alarm {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;


  }
 }




 isis {
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
```

```
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   isis-data {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
```

```
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    isis-hello {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
```

```
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  jfm {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  l2pt {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   l2tp {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  l3dest-miss {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
```

```
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  l3mc-sgv-hit-icl {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   l3mtu-fail {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
 }
 l3nhop {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  lacp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
```

```
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ldp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
```

```
    recover-time seconds;
    timeout-active-flows;


  }
  ldp-hello {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;
```

```
 }
 unclassified {
  bandwidth packets per second;
  burst packets;
  bypass-aggregate;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
}
```

```
ldpv6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
}
lldp {
 aggregate {
```

```
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  lmp {
   aggregate {
    bandwidth packets per second;
    burst packets;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  lmpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  localnh {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
```

```
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  mac-host {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  martian-address {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  mcast-snoop {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;




}
igmp {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   mld {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   pim {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
 }
 mld {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  mlp {
   add {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
```

```
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
aggregate {
 bandwidth packets per second;
 burst packets;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;
```

```
}
delete {
 bandwidth packets per second;
```

```
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   lookup {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
```

```
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   macpin-exception {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
```

```
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
```

```
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  msdp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
```

```
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  msdpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
```

```
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }


  mvrp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
```

```
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ndpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;




}
invalid-hop-limit {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   neighbor-advertisement {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    neighbor-solicitation {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
```

```
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   redirect {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
```

```
 recover-time seconds;
 timeout-active-flows;


}
router-advertisement {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
  }
  router-solicitation {
   bandwidth packets per second;
   burst packets;
   bypass-aggregate;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
   flow-detect-time seconds;
   flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;


  }
 }
```

```
nonucast-switch {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;



 }
}
ntp {
 aggregate {
```

```
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  oam-lfm {
   aggregate {
    bandwidth packets per second;
    burst packets;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ospf {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ospf-hello {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ospf-uncls {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ospf-hello {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ospfv3v6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    ospfv3v6-hello {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ospfv3v6-uncls {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
```

```
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  overlay {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   arp {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    ndpv6 {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
```

```
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;


  }
 }
 pfcp {
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
```

```
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
   }


   recover-time seconds;
```

```
 }
}
pfe-alive {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
```

```
}
pim {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
}
pim-ctrl {
```

```
aggregate {
 bandwidth packets per second;
 burst packets;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;



}
hello {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   join {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pim-data {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pimv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
```

```
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pkt-inject {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
```

```
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pmvrp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
```

```
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pos {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
```

```
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ppp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;




}
authentication {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
```

```
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    echo-rep {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
  echo-req {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
```

```
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
ipcp {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
}
ipv6cp {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
```

```
isis {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
lcp {
 bandwidth packets per second;
 burst packets;
```

```
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   mlppp-lcp {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   mplscp {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pppoe {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;




}
padi {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   padm {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   padn {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    pado {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
```

```
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   padr {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
```

```
    recover-time seconds;
    timeout-active-flows;


  }
  pads {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;
```

```
}
padse {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
padt {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



   }
  }
  proto-802-1x {
   aggregate {
    bandwidth packets per second;
```

```
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ptp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
```

```
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ptpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pvstp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  radius {
   accounting {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;


  }
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;




}
authorization {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   server {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  redirect {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
```

```
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   reject {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
```

```
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  rejectv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
```

```
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  re-services {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
```

```
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
captive-portal {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
 }
}
re-services-v6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;
```

```
    }
    captive-portal {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   resolve {
```

```
aggregate {
 bandwidth packets per second;
 burst packets;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }
 recover-time seconds;
}
mcast-v4 {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
```

```
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    mcast-v6 {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
```

```
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   other {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    ucast-v4 {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ucast-v6 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
```

```
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   rip {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ripv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  rsvp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
```

```
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    rsvp-hello {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   rsvp-uncls {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
```

```
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  rsvpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
```

```
    timeout-active-flows;



   }
  }
  sample {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }


    recover-time seconds;
```

```
   }
   host {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   pfe {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   sflow {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    syslog {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
```

```
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   tap {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  sample-dest {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  sample-source {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
```

```
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
sctp {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
```

```
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



   }
  }
  services {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;



    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
```

```
    disable-fpc;
   }


   recover-time seconds;




  }
  bsdt {
   bandwidth packets per second;
   burst packets;
   bypass-aggregate;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
   flow-detect-time seconds;
   flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
```

```
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
packet {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
```

```
    timeout-active-flows;


  }
 }
 snmp {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;
```

```
 }
}
snmpv6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
```

```
}
ssh {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
}
sshv6 {
```

```
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;



 }
}
stp {
 aggregate {
  bandwidth packets per second;
```

```
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  syslog {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
```

```
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  tacacs {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  tcc {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ethernet-tcc {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    iso-tcc {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
```

```
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    unclassified {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  tcp-flags {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;




}
established {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    initial {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
```

```
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
```

```
    recover-time seconds;
    timeout-active-flows;


   }
  }
  telnet {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;
```

```
 }
}
telnetv6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;
```

```
 }
}
ttl {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
}
```

```
tunnel-fragment {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;



 }
}
tunnel-ka {
 aggregate {
```

```
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  unclassified {
   aggregate {
    bandwidth packets per second;
    burst packets;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
```

```
    recover-time seconds;
```

```
   }
   control-layer2 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
```

```
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   control-v4 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    control-v6 {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }




   fw-host {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   host-route-tcp-v4 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    host-route-udp-v4 {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   host-route-v4 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    host-route-v6 {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
```

```
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    mcast-copy {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
```

```
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
other {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
```

```
    timeout-active-flows;



   }





  }
  unknown-l2mc {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
```

```
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  urpf-fail {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
```

```
    timeout-active-flows;


  }
 }
 vcipc-udp {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;
```

```
 }
}
virtual-chassis {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
  }
  recover-time seconds;
 }
 control-high {
  bandwidth packets per second;
  burst packets;
  bypass-aggregate;
  disable-fpc;
  disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   control-low {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    vc-packets {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
```

```
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    vc-ttl-errors {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  vrrp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
```

```
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
vrrpv6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
  }
  flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
  }
  flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
```

```
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  vxlan {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   bfd {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
  }
  traceoptions {
   file< filename><files files><match match><size size><(world-readable | no-world-readable)>;
   flag name;

   level(all | error | info | notice | verbose | warning);
   no-remote-trace;


  }
  }
```

Hierarchy Level

```
[edit system]
```

Description

Configure DDOS process

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

ddos-protection (DDoS)

IN THIS SECTION

- Syntax (ACX Series Routers) | 5137
- Syntax (PTX Series Routers and QFX Series, EX3400, EX4100, EX4300-MP, and EX4400 Switches) | 5138
- Syntax (Other Routers and EX9200 Switches) | 5139
- Hierarchy Level | 5140
- Description | 5140
- Required Privilege Level | 5141
- Release Information | 5141

Syntax (ACX Series Routers)

```
ddos-protection
    global {
        disable-routing-engine;
        disable-logging;
    }
    protocols protocol-group aggregate {
        bandwidth packets-per-second;
        burst size;
        disable-logging;
        disable-routing-engine;
        priority level;
    }
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
```

```
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Syntax (PTX Series Routers and QFX Series, EX3400, EX4100, EX4300-MP, and EX4400 Switches)

```
ddos-protection
    global {
        disable-fpc;
        disable-logging;
    }
    protocols protocol-group (aggregate | packet-type) {
        bandwidth packets-per-second;
        burst size;
        bypass-aggregate;
        disable-fpc;
        disable-logging;
        fpc slot-number {
            bandwidth-scale percentage;
            burst-scale percentage;
            disable-fpc;
        }
        priority level;
    }
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Syntax (Other Routers and EX9200 Switches)

```
ddos-protection
    global {
        disable-fpc;
        disable-logging;
        disable-routing-engine;
        flow-detection;
        flow-level-control;
        flow-detection-mode;
        flow-report-rate;
        violation-report-rate;
    }
    protocols protocol-group (aggregate | packet-type) {
        bandwidth packets-per-second;
        burst size;
        bypass-aggregate;
        disable-fpc;
        disable-logging;
        disable-routing-engine;
        flow-detection-mode (automatic | off | on);
        flow-detect-time seconds;
        flow-level-bandwidth {
            logical-interface flow-bandwidth;
            physical-interface flow-bandwidth;
            subscriber flow-bandwidth;
        }
        flow-level-control {
            logical-interface flow-control-mode;
            physical-interface flow-control-mode;
            subscriber flow-control-mode;
        }
        flow-level-detection {
            logical-interface flow-detection-mode;
            physical-interface flow-detection-mode;
            subscriber flow-detection-mode;
        }
        flow-recover-time seconds;
        flow-timeout-time seconds;
        fpc slot-number {
            bandwidth-scale percentage;
            burst-scale percentage;
```

```
            disable-fpc;
        }
        ipsec {
               ike;
               on-trigger;
               mtu-error;
               sn-alarm;
               spi-inval;
               tx-alarm;
               unclassified;
             }
        no-flow-logging
        priority level;
        recover-time seconds;
        timeout-active-flows;
    }
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure DDoS protection policers for control plane DDoS protection.

DDoS attacks typically use network control packets to trigger large numbers of exceptions to a device's control plane that disrupts normal network operations. DDoS protection polices traffic to enable the device to continue functioning under a DDoS attack.

DDoS protection is enabled by default on supporting devices for the protocol groups and packet types available on the device. You can disable particular policers or change default policer parameters, including:

- Set the maximum allowed traffic rate, maximum burst size, and traffic priority.
- (On some devices) Define how much time must pass since the last violation before the traffic flow is considered to have recovered from the attack.
- (On some devices) Scale bandwidth and burst values for individual line cards so that the policers at this level trigger at lower thresholds than the overall protocol or packet thresholds.

> **NOTE**: Some EX Series switches might have control plane DDoS protection but don't support CLI options to show or change the default policer parameters.

DDoS protection supports policers for many protocol groups. On some devices, you can change policer parameters for specific packet types within some protocol groups. Protocol group and packet type support varies across platforms and Junos OS releases. See the `protocols` statement for details on the main differences, as follows:

- For ACX Series routers, PTX Series routers and QFX Series switches, see protocols (DDoS) (ACX Series, PTX Series, and QFX Series).
- For all other routing devices and EX9200 switches, see "protocols (DDoS)" on page 13751.

The remaining statements in this configuration statement hierarchy are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Starting in Junos OS Release 24.2R1, you can configure the DDOS protocol using CLI on EX4100 and EX4400 devices.

Starting in Junos OS Release 24.4R1, you can configure the DDOS protocol using CLI on EX3400 and EX4300-MP devices.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

Statement introduced in Junos OS Release 24.2R1 for MX304 devices.

RELATED DOCUMENTATION

- Control Plane Distributed Denial-of-Service (DDoS) Protection Overview
- Configuring Control Plane DDoS Protection

deactivate

IN THIS SECTION

- Syntax | 5142
- Description | 5142
- Options | 5142
- Required Privilege Level | 5143
- Release Information | 5143

Syntax

```
deactivate (statement | identifier)
```

Description

Add the `inactive:` tag to a statement, effectively commenting out the statement or identifier from the configuration. Statements or identifiers marked as inactive do not take effect when you issue the `commit` command.

Options

identifier Identifier to which you are adding the `inactive:` tag. It must be an identifier at the current hierarchy level.

statement Statement to which you are adding the `inactive:` tag. It must be a statement at the current hierarchy level.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

activate | 1795

delete | 5276

Deactivating and Reactivating Statements and Identifiers in a Device Configuration

deactivation-attribute (AAA Profile)

IN THIS SECTION

- Syntax | 5143
- Hierarchy Level | 5144
- Description | 5144
- Required Privilege Level | 5144
- Release Information | 5144

Syntax

```
deactivation-attribute {
    <code numeric-code;>
    <vendor-id vendor-id;>
}
```

Hierarchy Level

```
[edit unified-edge aaa profiles aaa-profile-name radius policy]
```

Description

Configure the RADIUS attribute that you want to carry the PCC rulebase name for rulebase deactivations from the RADIUS policy server to the MX Series router. By default, the rulebase name is carried in the ERX-Service-Deactivate Juniper vendor-specific attribute.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring an AAA Profile

Understanding AAA Profiles

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

deactivation-delay

IN THIS SECTION

- Syntax | 5145
- Hierarchy Level | 5145
- Description | 5145

- Options | 5145
- Required Privilege Level | 5145
- Release Information | 5145

Syntax

```
deactivation-delay seconds;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN interfaces, configure the ISDN deactivation delay. Used only for dialer backup and dialer watch cases.

Options

seconds—Interval before the backup interface is deactivated after the primary interface has comes up.

- **Range:** 1 through 4,294,967,295 seconds
- **Default:** 0 (zero)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

dead-criteria-retries (RADIUS Server)

IN THIS SECTION

- Syntax | 5146
- Hierarchy Level | 5146
- Description | 5146
- Default | 5146
- Options | 5147
- Required Privilege Level | 5147
- Release Information | 5147

Syntax

```
dead-criteria-retries retry-number interval seconds;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure a limit to the number of times the MX Series router can resend a request to the RADIUS server when no response from the RADIUS server is received. If the number of retries reaches this limit, the RADIUS server is marked as dead, and the MX Series router begins to send requests to other RADIUS servers in the network element.

Default

The dead server detection function is disabled.

Options

retry-number Number of retries.

- **Range:** 10 through 65535

seconds Time interval in seconds.

- **Range:** 5 through 300

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

dead-peer-detection

IN THIS SECTION

- Syntax | 5148
- Hierarchy Level | 5148
- Description | 5148
- Options | 5148
- Required Privilege Level | 5149

Release Information | 5149

Syntax

```
dead-peer-detection {
    (always-send | optimized | probe-idle-tunnel);
    interval seconds;
    threshold number;
}
```

Hierarchy Level

```
[edit security ike gateway gateway-name]
```

Description

Enable the device to use dead peer detection (DPD). DPD is a method used by devices to verify the current existence and availability of IPsec peers. A device performs this verification by sending encrypted IKE Phase 1 notification payloads (R-U-THERE messages) to a peer and waiting for DPD acknowledgements (R-U-THERE-ACK messages) from the peer.

Options

interval Specify the amount of time that the peer waits for traffic from its destination peer before sending a dead-peer-detection (DPD) request packet.

- **Default:** 10 seconds
- **Range:** 2 through 60 seconds

always-send Instructs the device to send dead peer detection (DPD) requests regardless of whether there is outgoing IPsec traffic to the peer.

optimized Send dead peer detection (DPD) messages if there is no incoming IKE or IPsec traffic within the configured interval after outgoing packets are sent to the peer. This is the default DPD mode.

probe-idle-tunnel Send dead peer detection (DPD) messages during idle traffic time between peers.

threshold Specify the maximum number of unsuccessful dead peer detection (DPD) requests to be sent before the peer is considered unavailable.

- **Default:** 5
- **Range:** 1 through 5

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for the `optimized` and `probe-idle-tunnel` options added in Junos OS Release 12.1X46-D10.

Support for multiple peer addresses in DPD configuration with IPsec VPN running iked process is introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Understanding AutoVPN

IPsec VPN Overview

dead-peer-detection (Security Group VPN Server)

IN THIS SECTION

- Syntax | 5150
- Hierarchy Level | 5150
- Description | 5150
- Options | 5150

Required Privilege Level | 5151
Release Information | 5151

Syntax

```
dead-peer-detection {
    always-send;
    interval seconds;
    threshold number;
}
```

Hierarchy Level

```
[edit security group-vpn server ike gateway gateway-name]
```

Description

Enable the device to use dead peer detection (DPD). DPD is a method used by devices to verify the current existence and availability of IPsec peers. A device performs this verification by sending encrypted IKE Phase 1 notification payloads (R-U-THERE messages) to a peer and waiting for DPD acknowledgements (R-U-THERE-ACK messages) from the peer.

Options

`always-send`—Send probes periodically regardless of incoming and outgoing data traffic. Starting in Junos OS Release 24.2R1, MX304 supports dead-peer-detection in `always-send mode` by default or when explicitly configured.

`interval` *`seconds`*—Specify the interval time in seconds between DPD probe messages.

- **Range:** 10 through 60 seconds
- **Default:** 10 seconds

`threshold` *`number`*—Specify the maximum number of DPD retransmissions.

- **Range:** 1 through 5

- **Default:** 5

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Support for the Group VPN server added in Junos OS Release 15.1X49-D30 for vSRX Virtual Firewall.

RELATED DOCUMENTATION

Group VPNv2 Overview

gateway (Security Group VPN Server IKE) | **7660**

dead-peer-detection (Services IPsec VPN)

IN THIS SECTION

- Syntax | **5151**
- Hierarchy Level | **5152**
- Description | **5152**
- Options | **5152**
- Required Privilege Level | **5152**
- Release Information | **5152**

Syntax

```
dead-peer-detection {
    interval seconds;
```

```
    threshold number;
}
```

Hierarchy Level

```
[edit services ipsec-vpn rule-name term term-name then]
```

Description

Sets dead peer detection options when dead peer detection has been enabled with the `initiate-dead-peer-detection` command. The `dead-peer-detection` options are used for IKEv1 security associations (SAs). Starting in Junos OS Release 17.2R1, the `dead-peer-detection` options are also applicable to IKEv2 SAs. In Junos OS Release 17.1 and earlier, the `dead-peer-detection` options are not applicable to IKEv2 SAs, which use the default values.

Options

- interval—Specify the amount of time that the peer waits for dead-peer-detection (DPD) response from its destination peer before sending next DPD request packet. Range is 1 through 180 seconds. The default value is 10 seconds.
- threshold—Specify the maximum number of unsuccessful dead peer detection (DPD) requests to be sent before the peer is considered unavailable. Range is 1 through 10. The default value is 3.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

IKEv2 support introduced in Junos OS Release 17.2.

deaggregate

IN THIS SECTION

- Syntax | 5153
- Hierarchy Level | 5153
- Description | 5153
- Default | 5153
- Options | 5154
- Required Privilege Level | 5154
- Release Information | 5154

Syntax

```
deaggregate | no-deaggregate;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Control forwarding equivalence class (FEC) deaggregation on the router. The use of the `deaggregate` statement in LDP is a standard practice that we recommend for LDP deployments.

Default

Deaggregation is disabled on the router.

Options

`deaggregate`—Deaggregate FECs.

`no-deaggregate`—Aggregate FECs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring FEC Deaggregation

debounce-interval

IN THIS SECTION

- Syntax | 5155
- Hierarchy Level | 5155
- Description | 5155
- Options | 5155
- Required Privilege Level | 5155
- Release Information | 5155

Syntax

```
debounce-interval seconds;
```

Hierarchy Level

```
[edit logical-systems name protocols uplink-failure-detection group],
[edit protocols uplink-failure-detection group]
```

Description

The amount of time, in seconds, that elapses before the downlink interfaces are brought up after a state change of the uplink interfaces.

Options

seconds—Delay, in seconds, before bringing up the `link-to-disable` port.

- **Range:** 5 to 300 seconds

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

debug-level (System Services)

IN THIS SECTION

Syntax | 5156

- Hierarchy Level | 5156
- Description | 5156
- Options | 5156
- Required Privilege Level | 5157
- Release Information | 5157

Syntax

```
debug-level (alert | crit | emerg | error | info | notice | warn);
```

Hierarchy Level

```
[edit system services webapi (System Services)]
```

Description

Specify the trace level for the integrated ClearPass authentication and enforcement Web API process (webapi). Level describes a flag that specifies the type of logs to be write into the log file for the Web API process (webapi).

Options

debug-level

Debug level for webapi process.

- Values:
  - alert— Matches alert messages.
  - crit— Matches critical messages.
  - emerg— Matches emergency messages.
  - error— Matches error messages.

- notice— Matches notification messages.
- warn— Matches warning messages.

Required Privilege Level

1. system—To view this statement in the configuration.
2. system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

debug-log (System Services)

IN THIS SECTION

- Syntax | 5157
- Hierarchy Level | 5157
- Description | 5158
- Options | 5158
- Required Privilege Level | 5158
- Release Information | 5158

Syntax

```
debug-log file;
```

Hierarchy Level

```
[edit system services webapi (System Services)]
```

Description

Specify the name of the log file to which trace messages for the integrated ClearPass authentication and enforcement Web API process (webapi) are written.

The debug level flag determines the kind of logs that are written to this file.

Options

file Debug file for Web API process.

Required Privilege Level

1. system—To view this statement in the configuration.
2. system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

decapsulate-accept-inner-vlan

IN THIS SECTION

- Syntax | 5158
- Hierarchy Level | 5159
- Description | 5159
- Required Privilege Level | 5159
- Release Information | 5159

Syntax

```
decapsulate-accept-inner-vlan
```

Hierarchy Level

```
[edit protocols l2-learning]
[edit routing-instances routing-instance-name vlans vlan-name vxlan]
```

Description

Configure the switch to de-encapsulate and to accept original VLAN tags in Virtual Extensible LAN (VXLAN) packets.

We require this option on some platforms to process tagged VXLAN payloads so the device doesn't drop the tagged packet. You don't need to configure this option in some use cases, such as:

- On the QFX10000 line of switches, you only need to include the `encapsulate-inner-vlan configuration` statement at the `[edit vlans vlan-name vxlan]` hierarchy level. You don't need to also configure the `decapsulate-accept-inner-vlan` option.
- On ACX7100 routers, the device doesn't drop the tagged packets, and can process the packets whether you configure the `decapsulate-accept-inner-vlan` option or not.

NOTE: Starting in Junos OS release 22.3R1, IS-IS logical link control (LLC) packets that are encapsulated in a VXLAN tunnel now have their VLAN tags removed by default. You can configure this option to accept the original VLAN tags on IS-IS LLC packets.

NOTE: ARP suppression is automatically disabled when you encapsulate VLAN traffic for EVPN-VXLAN by configuring `encapsulate-inner-vlan` and its equivalent converse statement `decapsulate-accept-inner-vlan`. As a result, you might observe ARP packets being sent by the device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 14.1X53 for the QFX Series.

RELATED DOCUMENTATION

Understanding VXLANs

Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network

encapsulate-inner-vlan | 6282

decapsulate (Firewall Filter)

IN THIS SECTION

- Syntax | 5160
- Hierarchy Level | 5161
- Description | 5161
- Caveats | 5162
- Options | 5162
- Required Privilege Level | 5162
- Release Information | 5162

Syntax

```
decapsulate {
    gre {
        apply-groups;
        apply-groups-except;
        forwarding-class;
        interface-group(0 -255)
        no-decrement-ttl;
        routing-instance;
        sample;
    }
    inet-in-udp {
        apply-groups;
        apply-groups-except;
        forwarding-class;
        no-decrement-ttl;
```

```
        routing-instance;
    }
    inet6-in-udp {
        apply-groups;
        apply-groups-except;
        forwarding-class;
        no-decrement-ttl;
        routing-instance;
    }
    mpls-in-udp {
        apply-groups;
        apply-groups-except;
        forwarding-class;
        no-decrement-ttl;
        routing-instance;
    }
    gre-in-udp{
    l2tp {
        apply-groups;
        apply-groups-except;
        cookie;
        forwarding-class;
        no-decrement-ttl;
        output-interface;
        sample;
    }
}
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name then]
```

Description

Define the termination action for GRE, UDP, and L2TP tunnels.

NOTE: `inet-in-udp`, `inet6-in-udp`, and `mpls-in-udp` are supported only on PTX10003, PTX10004, PTX10008, PTX10016 and PTX10001-36MR device models.

Caveats

The following are the caveats for using the decapsulate firewall filter action.

- For GRE tunnel decapsulate action if "payload type" in GRE header is a value other than IPv4, IPv6, or MPLS, the packet will be dropped.
- For `UDP` tunnel decapsulation, it is expected that the user will configure term match conditions for UDP destination port correctly to term action for decapsulate of IPv4, IPv6 or MPLS. In case of mismatch traffic can be dropped.
- No decrement TTL attribute is not supported for MPLS payload.
- Output filter cannot match TTL value of outgoing packet for ingress filter decapsulation traffic with no-decrement-ttl configured.

Options

`gre`—(Optional) Terminate a GRE tunnel for the filter conditions that are matched.

`inet-in-udp`—(Optional) Terminate a UDP tunnel with IPv4 payload.

`inet6-in-udp`—(Optional) Terminate a UDP tunnel with IPv6 payload.

`mpls-in-udp`—(Optional) Terminate a UDP tunnel with MPLS payload.

`l2tp`—(Optional) Terminate an L2TP tunnel for the filter conditions that are matched.

`output-interface` *`interface-name`*—(Optional) For L2TP tunnels, enable the packet to be duplicated and sent towards the customer or the network (based on the MAC address in the Ethernet payload),

`cookie` *`l2tpv3-cookie`*—(Optional) For L2TP tunnels, specify the L2TP cookie for the duplicated packets. If the tunnel does not contain the receive-cookie configured, packet injection does not happen. In such a case, any received tunnel packet is counted and dropped in the same manner in which packets that arrive with a wrong cookie are counted and dropped.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

`output-interface` and `cookie` options introduced in Junos OS Release 15.1.

`decapsulate gre` introduced in Junos OS Release 15.1F3 and 16.1R2 for PTX5000 routers with third generation FPCs and Junos OS Release 15.1F6 and 16.1R2 for PTX3000 routers with third-generation FPCs.

`no-decrement-ttl` attribute for the `decapsulate gre` filter action introduced in Junos OS Release 15.1F6 and 16.2R1 for PTX5000 routers with third-generation FPCs.

`inet-in-udp`, `inet6-in-udp`, and `mpls-in-udp` introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Guidelines for Configuring Firewall Filters

Configuring Multifield Classifiers

Guidelines for Configuring and Applying Firewall Filters in Logical Systems

decapsulate-inner-vlan

IN THIS SECTION

- Syntax | **5163**
- Hierarchy Level | **5164**
- Description | **5164**
- Default | **5164**
- Options | **5164**
- Required Privilege Level | **5164**
- Release Information | **5164**

Syntax

```
decapsulate-inner-vlan
```

Hierarchy Level

```
[edit routing-instances routing-instance-name vlans vlan-name vxlan],
[edit vlans vlan-name vxlan]
```

Description

Configure the switch to de-encapsulate a preserved original VLAN tag (in the inner Ethernet packet) from a VXLAN encapsulated packet.

Default

A preserved VLAN tag is dropped when the packet is de-encapsulated.

Options

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Understanding VXLANs

Manually Configuring VXLANs on QFX Series and EX4600 Switches

Examples: Manually Configuring VXLANs on QFX Series and EX4600 Switches

encapsulate-inner-vlan | **6282**

decompress-layer

IN THIS SECTION

- Syntax | 5165
- Hierarchy Level | 5165
- Description | 5165
- Options | 5165
- Required Privilege Level | 5166

Syntax

```
decompress-layer (block | log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name fallback-options]
```

Description

The Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, decompress layer error is the error returned by the scan engine when the scanned file has too many compression layers. The default action is block.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Full Antivirus Configuration Overview

decompress-layer-limit

IN THIS SECTION

- Syntax | 5166
- Hierarchy Level | 5166
- Description | 5167
- Options | 5167
- Required Privilege Level | 5168
- Release Information | 5168

Syntax

```
decompress-layer-limit decompress-layer-limit;
```

Hierarchy Level

```
[edit security utm default-configuration anti-virus scan-options]
```

Description

When an antivirus scan engine scans a file for viruses, the scan engine decompresses the layers of nested compressed files and files with embedded extractable objects. Embedded extractable objects include files such as archive files (tar), MS Word, and PowerPoint files. For example, if a message contains a compressed .zip file that contains another compressed .zip file, then there are two compression layers. Decompressing both files requires a decompress layer setting of 2. You can set the decompression layer limit for the scan engine.

During the transfer of data, some protocols use content encoding. Before an antivirus scan engine scans viruses, the scan engine decodes this layer, which is considered as a decompression level.

The decompression layer limit is applicable to the following compressed files:

- zip, rar, and gzip
- Encoded data such as Multipurpose Internet Mail Extension (MIME)
- Packaged data such as OLE, CAP, MSI, TAR, EML
- Files with internal extractable objects, such as archive files (tar), MS Word, and PowerPoint files.

If a file exceeds the compression layer limit, the scan engine drops or forwards the file based on the fallback options.

The scan engine scans each layer before unpacking the next layer. The scan engine continues to scan until any of the following conditions are met, whichever happens first:

- Reaches the decompression limit
- Exceeds the system resource allocated for decompression
- Finds a virus or other malware
- Decompresses the data completely.

When the virus signature database becomes larger and the scan algorithms are more sophisticated, the scan engine can look deeper into the data for embedded malware. As a result, the scan engine uncovers more layers of compressed data.

The Juniper Networks device's level of security is limited by the decompress layer limit. You define the decompress layer limit based on the memory allocated to the security service.

Options

decompress-layer-limit	Specify the number of decompression layer limit.

- **Default:** 3
- **Range:** 0 through 10

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

Support at the `[edit security utm default-configuration]` hierarchy level introduced in Junos OS Release 18.2R1.

The `[edit anti-virus scan-options]` hierarchy level is deprecated in Junos OS Release 18.2R1.

Support for Avira Antivirus scan engine added in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

On-Device Avira Antivirus

Full Antivirus Configuration Overview

decryption-failures

IN THIS SECTION

- Syntax | 5169
- Hierarchy Level | 5169
- Description | 5169
- Default | 5169
- Options | 5169

Required Privilege Level | 5169
Release Information | 5170

Syntax

```
decryption-failures {
    threshold value;
}
```

Hierarchy Level

```
[edit security alarms potential-violation]
```

Description

Raise a security alarm after exceeding a specified number of decryption failures.

Default

Multiple decryption failures do not cause an alarm to be raised.

Options

failures—Number of decryption failures up to which an alarm is not raised. When the configured number is exceeded, an alarm is raised.

- **Range:** 1 through 1,000,000,000.
- **Default:** 1000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

IPsec Overview

Monitoring VPN Traffic

dedicated-buffer

IN THIS SECTION

- Syntax | 5170
- Hierarchy Level | 5171
- Description | 5171
- Options | 5171
- Required Privilege Level | 5171
- Release Information | 5171

Syntax

```
dedicated-buffer {
    egress {
        percent percent;
    }
    ingress {
        percent percent;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Set the global dedicated buffer pool of the device as a percent of its default. This, in effect, increases the shared buffer pool.

> **NOTE**: You can descrease the `dedicated-buffer` (thereby increasing the shared buffer pool), or decrease the `shared-buffer` (thereby increasing the dedicated buffer pool), but not both.

Options

`egress percent` ***`percent`***

Set the global egress dedicated buffer size as a percentage of its default.

Range: 15-100 percent

Default: 100 percent

`ingress percent` ***`percent`***

Set the global ingress dedicated buffer size as a percentage of its default.

Range: 15-100 percent

Default: 100 percent

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Configuring Ingress and Egress Dedicated Buffers

dedicated-buffer-profile | 5172

shared-buffer | 16964

dedicated-buffer-profile

IN THIS SECTION

- Syntax | 5172
- Hierarchy Level | 5173
- Description | 5173
- Options | 5173
- Required Privilege Level | 5174
- Release Information | 5174

Syntax

```
dedicated-buffer-profile name {
    egress {
        buffer-size {
            ( cells | none);
        }
    }
    ingress {
        buffer-size {
            ( cells | none);
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Dedicated buffer profile configuration.

By default, the operating system calculates port level dedicated buffers internally. Therefore, even ports that are down or unused also get an equal amount of dedicated buffer space that is then unavailable to any traffic burst. `dedicated-buffer-profile` provides the ability to increase or decrease the default dedicated buffer at a physical interface level.

Once you define a `dedicated-buffer-profile`, you can attach it directly to a physical interface. For example:

```
[edit class-of-service interfaces interface-name]
user@host# set dedicated-buffer-profile name
```

Remaining dedicated buffers that you do not allocate to any ports by a `dedicated-buffer-profile` are equally shared among ports (based on speed) that don't have a `dedicated-buffer-profile` assigned to them.

NOTE: You cannot assign a `dedicated-buffer-profile` to aggregated Ethernet (`ae-`) interfaces. You can only assign a `dedicated-buffer-profile` to a physical interface.

CAUTION: If the `buffer-size` of all dedicated buffer profiles combined exceeds the total available dedicated buffer pool, the system logs a syslog error and does not implement the new configuration even though the commit succeeds.

Options

name Name the buffer profile. You apply a dedicated buffer profile to a physical interface by this name.

egress buffer-size *cells* Set the egress dedicated buffer size for a port as a number of cells or `none` for no dedicated buffer.

Range: 20-100,000 (Each cell is 254 bytes.)

`ingress buffer-size cells`

Set the ingress dedicated buffer size for a port as a number of cells or `none` for no dedicated buffer.

Range: 20-50,000 (Each cell is 254 bytes.)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Configuring Ingress and Egress Dedicated Buffers

dedicated-buffer | **5170**

shared-buffer | **16964**

show class-of-service dedicated-buffer-profile | **19007**

dedicated-ukern-cpu (BFD)

IN THIS SECTION

- Syntax | **5175**
- Hierarchy Level | **5175**
- Description | **5175**
- Required Privilege Level | **5175**
- Release Information | **5175**

Syntax

```
dedicated-ukern-cpu;
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable the dedicated Bidirectional Forwarding Detection (BFD) protocol. One dedicated CPU core is allocated for the flowd ukernel thread to handle the dedicated BFD. This ensures that the BFD packet processing does not compete with the Routing Engine daemons.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Enabling Dedicated and Real-Time BFD

show chassis dedicated-ukern-cpu | **17418**

bfd-liveness-detection (LAG) | **2926**

authentication (LAG) | **2622**

detection-time (LAG) | **5549**

transmit-interval (LAG) | **30090**

deep-vlan-qualified-learning

IN THIS SECTION

- Syntax | 5176
- Hierarchy Level | 5176
- Description | 5176
- Required Privilege Level | 5176
- Release Information | 5177

Syntax

```
deep-vlan-qualified-learning vlan_tag_number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical_unit_number]
```

Description

Enable qualified MAC learning on the third VLAN tag (innermost) of an ingress 3-tagged packet, without any kind of implicit VLAN manipulation. If the packet has 2 tags, MAC learning happens on the second VLAN. If the ingress packet has more than 3 tags, all tags beyond the third tag are treated as part of data and not considered for learning. For bidirectional traffic flow `input-vlan-map pop` must be configured.

NOTE: In Junos OS Release 17.4R1 the *vlan_tag_number* is 3 only.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

vlan-id inner-all | 30804

Understanding Qualified MAC Learning

default-action (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 5177
- Hierarchy Level | 5177
- Description | 5178
- Required Privilege Level | 5178
- Release Information | 5178

Syntax

```
default-action {
    drop;
    forward-only;
    local-server-group local-server-group;
    relay-server-group relay-server-group;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay dhcpv6 relay-option],
```

```
[edit forwarding-options dhcp-relay group group-name relay-option],
[edit forwarding-options dhcp-relay dhcpv6 group group-name relay-option],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify the action DHCP relay agent takes when the option string in client traffic does not satisfy any match criteria or when no match criteria are configured.

The `default-action` statement is optional. If the match criteria are not satisfied or not configured and no `default-action` is specified, DHCP relay processes the traffic in the normal manner.

The `local-server-group` option is not supported for DHCPv6 relay agent.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address

default-action (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 5179
- Hierarchy Level | 5179
- Description | 5179
- Default | 5179
- Options | 5180
- Required Privilege Level | 5180
- Release Information | 5180

Syntax

```
default-action (drop | packet-in);
```

Hierarchy Level

```
[edit protocols openflow switch switch-name]
```

Description

Specify the default action that is executed when an OpenFlow packet does not match an existing flow entry. The default action is specific to the OpenFlow virtual switch and is the same across all filters associated with that virtual switch.

Default

If you do not include the `default-action` statement, the default action is `packet-in`.

Options

drop Drop packets that do not match an existing flow entry.

packet-in Accept packets that do not match an existing flow entry, and forward the packet to the controller.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Support for OpenFlow on Devices Running Junos OS

Understanding OpenFlow Flows and Filters on Devices Running Junos OS

openflow (Protocols OpenFlow) | **11842**

switch (Protocols OpenFlow) | **28755**

default-actions

IN THIS SECTION

- Syntax | **5181**
- Hierarchy Level | **5181**
- Description | **5181**
- Default | **5181**
- Options | **5181**
- Required Privilege Level | **5181**

Release Information | 5181

Syntax

```
default-actions {
    interface-down;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile profile-name]
```

Description

Define the action to be taken when connectivity to the remote MEP is lost.

Default

If no action is configured, no action is taken.

Options

`interface-down`—When a remote MEP connectivity failure is detected, bring the interface down.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring a CFM Action Profile to Specify CFM Actions for CFM Events

default-address-selection

IN THIS SECTION

- Syntax | 5182
- Hierarchy Level | 5182
- Description | 5182
- Required Privilege Level | 5183
- Release Information | 5183

Syntax

```
default-address-selection;
```

Hierarchy Level

```
[edit system]
```

Description

Use the loopback interface, `lo0`, as the source address for all locally generated IP packets when the packet is sent through a routed interface, and also when the packet is sent through a local interface such as `fxp0`. The `lo0` interface is the interface to the router's or switch's Routing Engine.

The default address is used as the source address for all locally generated IP packets on outgoing interfaces that are unnumbered. If an outgoing interface is numbered, the default address is chosen using the following sequence:

The default address of the router is chosen using the following sequence:

1. The primary address on the loopback interface `lo0` that is not `127.0.0.1` is used.
2. The primary address on the primary interface is used.
3. When there are multiple interfaces with "primary" and "preferred" addresses, the interface with the lowest interface index is selected, and the primary address is used. In the case that none of the interface's IP addresses are explicitly marked with the `primary` statement, the numerically lowest address on that interface is used as the system default address.
4. Any remaining interface with an IP address may be selected. This includes the router's management or internal interfaces. For this reason, it's recommended that you assign a loopback address, or explicitly configure a primary interface, to control default address selection.

An interface's *primary address* is used by default as the local address for broadcast and multicast packets sourced locally and sent out through the interface. By default, the primary address on an interface is selected as the numerically lowest local address configured on the interface.

An interface's *preferred address* is the default local address used for packets sourced by the local router or switch to destinations on the subnet. By default, the numerically lowest local address configured for the interface is chosen as the preferred address on the subnet.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configure TCP Options

Configure Default, Primary, and Preferred Addresses and Interfaces

default-chap-secret

IN THIS SECTION

- Syntax | 5184
- Hierarchy Level | 5184
- Description | 5184
- Default | 5185
- Required Privilege Level | 5185
- Release Information | 5185

Syntax

```
default-chap-secret name;
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options chap],
[edit interfaces interface-name unit logical-unit-number ppp-options chap],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options chap]
```

Description

Define the default CHAP secret to be used when no matching CHAP access profile exists.

For ATM2 IQ interfaces only, you can configure a default CHAP secret on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- `atm-ppp-llc`—PPP over AAL5 LLC encapsulation.
- `atm-ppp-vc-mux`—PPP over AAL5 multiplex encapsulation.

Default

If you do not include the `default-chap-secret` statement in the configuration, and an interface receives a CHAP challenge or response from a peer that is not in the applied access profile, the link is immediately dropped.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring a Default CHAP Secret

access-profile

default-client-identity (NETCONF TLS)

IN THIS SECTION

- Syntax | 5186
- Hierarchy Level | 5186
- Description | 5186
- Default | 5186
- Options | 5186
- Required Privilege Level | 5187
- Release Information | 5187

Syntax

```
default-client-identity {
    map-type (san-dirname-cn | specified);
    username username;
}
```

Hierarchy Level

```
[edit system services netconf tls]
```

Description

For NETCONF sessions over Transport Layer Security (TLS), configure the default method to derive the NETCONF username for clients that do not match any configured clients.

If the fingerprint of a client's presented certificate does not match the fingerprint for any client configured at the `[edit system services netconf tls client-identity]` hierarchy level, then the Junos device uses the `default-client-identity` map type to derive the NETCONF username for the client.

Junos devices support local users and LDAP remote users for NETCONF sessions over TLS. The username must either have a user account defined locally on the device, or it must be authenticated by an LDAP server, which then maps it to a user template account that is defined locally on the Junos device.

Default

If you do not include the `default-client-identity` statement, and a NETCONF-over-TLS client does match any clients configured at the `[edit system services netconf tls client-identity]` hierarchy level, then the device does not establish the NETCONF session.

Options

`map-type` *`type`*

Map type that defines how to derive the NETCONF username.

- Values:

- `san-dirname-cn`—Use the common name (CN) defined for the SubjectAltName's (SAN) DirName field (`DirName:/CN`) in the client certificate as the NETCONF username.

  If you specify `san-dirname-cn` as the map type, but the client certificate does not have a username in this field, the connection fails.

- `specified`—Use the NETCONF username defined in the `username` statement at the same hierarchy level.

`username` ***`username`***

Username under whose access privileges the NETCONF operations are executed when `map-type specified` is configured.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

Statement introduced in Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

client-identity (NETCONF TLS) | 4324

default (CoS)

IN THIS SECTION

- Syntax | 5188
- Hierarchy Level | 5188
- Description | 5188
- Required Privilege Level | 5188

- Release Information | 5188

Syntax

```
default;
```

Hierarchy Level

```
[edit class-of-service virtual-channel-groups group-name virtual-channel-name]
```

Description

Specify the default channel. You must configure one of the virtual channels in the group to be the default. Any traffic not explicitly directed to a virtual channel is transmitted by way of this default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

scheduler-map (CoS Virtual Channels) | 16350

shaping-rate (CoS Virtual Channels) | 16936

virtual-channel-group (CoS Interfaces) | 30740

virtual-channel-groups | 30742

virtual-channels | 30743

default (CoS Host Outbound Traffic)

IN THIS SECTION

- Syntax | 5189
- Hierarchy Level | 5189
- Description | 5189
- Options | 5189
- Required Privilege Level | 5189
- Release Information | 5190

Syntax

```
default value;
```

Hierarchy Level

```
[edit class-of-service host-outbound-traffic ieee-802.1]
```

Description

Apply a global default value to the IEEE 802.1p—priority code point (PCP)—field in the Ethernet frame header for all host outbound traffic.

Options

value Three-bit binary number.

- **Range:** 000 through 111

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic

Rewriting Packet Headers to Ensure Forwarding Behavior

default-gateway

IN THIS SECTION

- Hierarchy Level | 5190
- Description | 5191
- Options | 5191


```
default-gateway {
    (advertise | do-not-advertise | no-gateway-community);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
evpn],
[edit protocols evpn],
[edit routing-instances routing-instance-name protocols evpn]
```

Description

Specify the behavior for IRB interfaces defined as default gateways in an EVPN fabric.

In an EVPN-VXLAN environment with multihoming, you might use multiple EVPN routing instances on peer provider edge (PE) devices that share an Ethernet segment (ES). When you configure anycast gateways with this statement, we don't support mixing the default behavior (same as the `advertise` option) with the `no-gateway-community` option on the links that participate in the same ES.

As a result, if you configure the `default-gateway` statement with the `no-gateway-community` option in any EVPN routing instances on any peer PE device that share an ES, you must configure this statement:

- In all the routing instances that share the ES on a PE device,
- On all the peer PE devices that share the ES,
- Only with either the "no-gateway-community" option or the "do-not-advertise" option.

Options

advertise	Advertise the default gateway MAC address. This is the default behavior.
do-not-advertise	Don't advertise the default gateway MAC address.
no-gateway-community	Advertise virtual gateway addresses and IRB MAC addresses to EVPN peer devices so that Ethernet-only PE devices learn those MAC addresses.

RELATED DOCUMENTATION

Anycast Gateways

default-isid

IN THIS SECTION

- Syntax | 5192
- Hierarchy Level | 5192
- Description | 5192
- Options | 5192

Required Privilege Level | 5192
Release Information | 5192

Syntax

```
default-isid isid-number;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name service-groups service-group-name pbb-service-options]
```

Description

For IEEE 802.1ah provider backbone bridge (PBB) configurations, configure the default service identifier (I-SID) for all unmapped service VLANs (S-VLANs) for the customer routing instance (I-component) service group.

Options

default-isid— Default service identifier. Enter an I-SID in the range from **256** through **16777214**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.0.

RELATED DOCUMENTATION

Example: Configuring E-LINE and E-LAN Services for a PBB Network on MX Series Routers

default-lifetime

IN THIS SECTION

- Syntax | 5193
- Hierarchy Level | 5193
- Description | 5193
- Options | 5193
- Required Privilege Level | 5194
- Release Information | 5194

Syntax

```
default-lifetime seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Description

Configure the lifetime associated with a default router.

Options

seconds—Default lifetime. A value of 0 means this router is not the default router.

- **Range:** Maximum advertisement interval value through 9000 seconds
- **Default:** Three times the maximum advertisement interval value

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

max-advertisement-interval (Protocols IPv6 Neighbor Discovery) | 10415

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

default-lifetime (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 5194
- Hierarchy Level | 5195
- Description | 5195
- Options | 5195
- Required Privilege Level | 5195
- Release Information | 5195

Syntax

```
default-lifetime seconds;
```

Hierarchy Level

```
[edit protocols router-advertisement interface interface-name]
```

Description

Lifetime associated with a default router.

Options

seconds—Default lifetime. A value of 0 means this router is not the default router.

- **Range:** Maximum advertisement interval value through 9000 seconds
- **Default:** Three times the maximum advertisement interval value

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA

Dynamic Router Advertisement Configuration Overview

default-local-policy

IN THIS SECTION

Syntax | 5196

- Hierarchy Level | 5196
- Description | 5196
- Required Privilege Level | 5196
- Release Information | 5196

Syntax

```
default-local-policy {
    flow-action (drop | forward);
    maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value;
    burst-size uplink uplink-burst-size downlink downlink-burst-size;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
```

Description

Specify the default local policy, which is applied to the IP-based subscriber's data packets entering the access interface of the TDF domain when a TDF subscriber session does not exist.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

default-lsa

IN THIS SECTION

- Syntax | 5197
- Hierarchy Level | 5197
- Description | 5198
- Options | 5198
- Required Privilege Level | 5198
- Release Information | 5199

Syntax

```
default-lsa {
    default-metric metric;
    metric-type type;
    type-7;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id nssa],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast) area area-id nssa],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-id nssa],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
```

```
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id nssa],
[edit protocols (ospf | ospf3) area area-id nssa],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id nssa],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id nssa],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast) area area-id nssa]
```

Description

On area border routers only, for a not-so-stubby area (NSSA), inject a default link-state advertisement (LSA) with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.

The remaining statements are explained separately. See CLI Explorer.

Options

default-metric *metric*
Metric value, ranging from 1 through 16,777,215.

metric-type *type*
Metric type 1 or 2. The configured metric determines the method used to compute the cost to a destination:

- The Type 1 external metric is equivalent to the link-state metric. The path cost uses the advertised external path cost and the path cost to the AS boundary router (the route is equal to the sum of all internal costs and the external cost).
- The Type 2 external metric uses the cost assigned by the AS boundary router (the route is equal to the external cost alone). By default, OSPF uses the Type 2 external metric.

type-7
Flood Type 7 default link-state advertisements (LSAs) if the `no-summaries` statement is configured. By default, when the `no-summaries` statement is configured, a Type 3 LSA is injected into not-so-stubby areas (NSSAs) for Junos OS Release 5.0 and later. To support backward compatibility with earlier Junos OS releases, include the `type-7` statement. This statement enables NSSA ABRs to advertise a Type 7 default LSA into the NSSA if you have also included the `no-summaries` statement in the configuration.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Understanding OSPF Areas

Example: Configuring OSPF Not-So-Stubby Areas

nssa | **11717**

stub | **28665**

default-pap-password

IN THIS SECTION

- Syntax | **5199**
- Hierarchy Level | **5200**
- Description | **5200**
- Required Privilege Level | **5200**
- Release Information | **5200**

Syntax

```
default-pap-password password;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number ppp-options pap],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options pap]
```

Description

For PAP authentication, the default PAP password.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

> *Configuring the PPP Password Authentication Protocol*
>
> access-profile

default-peer

IN THIS SECTION

- Syntax | 5201
- Hierarchy Level | 5201
- Description | 5201
- Required Privilege Level | 5201
- Release Information | 5202

Syntax

```
default-peer;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name],
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
peer address],
[edit protocols msdp],
[edit protocols msdp group group-name],
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit routing-instances routing-instance-name protocols msdp],
[edit routing-instances routing-instance-name protocols msdp group group-name],
[edit routing-instances routing-instance-name protocols msdp group group-name peer address],
[edit routing-instances routing-instance-name protocols msdp peer address]
```

Description

Establish this peer as the default MSDP peer and accept source-active messages from the peer without the usual peer-reverse-path-forwarding (peer-RPF) check.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP with Active Source Limits and Mesh Groups

default-policy

IN THIS SECTION

- Syntax | 5202
- Hierarchy Level | 5202
- Description | 5202
- Options | 5203
- Required Privilege Level | 5203
- Release Information | 5203

Syntax

```
default-policy (deny-all | permit-all);
```

Hierarchy Level

```
[edit security policies]
```

Description

Configure the default security policy that defines the actions the device takes on a packet that does not match any user-defined policy.

Options

`deny-all`—Deny all traffic. Packets are dropped. This is the default.

`permit-all`—Permit all traffic that does not match a policy.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

default-policy (IDP)

IN THIS SECTION

- Syntax | 5204
- Hierarchy Level | 5204
- Description | 5204
- Options | 5204
- Required Privilege Level | 5204
- Release Information | 5204

Syntax

```
default-policy default-policy;
```

Hierarchy Level

```
[edit security idp]
```

Description

Specify which policy among the configured policies to be configured as the default IDP policy.

When you have multiple IDP policies configured and when policy conflict occurs, then the policy configured as default the IDP policy will be applied for a given session.

Options

default-policy—Name of the default policy.

NOTE: The default policy must be enforced in the data plane.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

An IDP policy is directly assigned in the security policy rule. This is to simplify IDP policy usage. As a part of session interest check, IDP is enabled if an IDP policy is present in any of the matched rules. An IDP policy is activated in security policies by permitting the IDP policy within the application services using the `set security policies from-zone zone-name to-zone zone-name policy policy-name then permit application-services idp-policy idp-policy-name` command. Because the IDP policy name is directly used in the security policy rule, the `[edit security idp active-policy policy-name]` statement is deprecated.

When the device is configured with unified policies, you can configure multiple IDP policies to provide the flexibility to have multiple policies active at the same time and to configure one of the IDP policies as the default IDP policy.

> **NOTE**: If you have configured two or more IDP policies in a unified security policy, then you must configure the default IDP policy.

default-pool (Address Pools)

IN THIS SECTION

- Syntax | 5205
- Hierarchy Level | 5205
- Description | 5205
- Required Privilege Level | 5206
- Release Information | 5206

Syntax

```
default-pool;
```

Hierarchy Level

```
[edit access address-assignment address-pools name]
```

Description

Configure the address pool as a default pool. A TDF domain uses the default address pool to specify the source IP addresses of packets that undergo TDF processing when an address pool is not specified for the TDF domain.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Understanding Source IP Filtering with Address Pools in TDF Domains for IP-Based Subscribers

Configuring Address Pools for Source-IP Filtering of IP-Based Subscribers

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

default-prefixes (Protocols RIFT)

IN THIS SECTION

- Syntax | 5206
- Hierarchy Level | 5207
- Description | 5207
- Options | 5207
- Required Privilege Level | 5207
- Release Information | 5207

Syntax

```
default-prefixes {
    family {
    (inet sets-of-ipv4-address | inet6 sets-of-ipv6-address);
```

```
}
}
```

Hierarchy Level

```
[edit protocols rift]
```

Description

Explicitly specify default prefixes generated in southbound direction.

This is helpful in cases where you have an internet default route and you want to inject the route from a leaf node without overlays. If the IP fabric originates default routes, you cannot differentiate between them and this might cause a null route. However, if you have all the internal fabric default routes on well known address prefixes, you can originate them and continue to use internet defaults in the IP fabric.

Options

family	Specify the family type for the default prefixes to be generated in the southbound direction.
inet *ipv4-address*	Specify the IPv4 IP addresses for the default prefixes to be generated in the southbound direction.
inet6 *ipv6-address*	Specify the IPv6 IP addresses for the default prefixes to be generated in the southbound direction.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

default-profile

IN THIS SECTION

- Syntax | 5208
- Hierarchy Level | 5208
- Description | 5208
- Options | 5208
- Required Privilege Level | 5209
- Release Information | 5209

Syntax

```
default-profile profile-name;
```

Hierarchy Level

```
[edit access firewall-authentication pass-through]
```

Description

Specify the authentication profile to use if no profile is specified in a policy.

Options

profile-name—Name of the profile.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Understanding Captive Portal Configuration on the Junos OS Enforcer

default-profile (Juniper Secure Connect)

IN THIS SECTION

- Syntax | 5209
- Hierarchy Level | 5209
- Description | 5210
- Required Privilege Level | 5210
- Release Information | 5211

Syntax

```
default-profile default-profile;
```

Hierarchy Level

```
[edit security remote-access]
```

Description

Configure default profile. On your security device, you must specify one of the remote-access profiles as the default profile.

> **NOTE**: Starting in Junos OS Release 23.1R1, we've hidden the `default-profile` option at the [`edit security remote-access`] hierarchy level. In releases before Junos OS Release 23.1R1, you use this option to specify one of the remote-access profiles as the default profile in Juniper Secure Connect. But with changes to the format of remote-access profile names, we no longer require the `default-profile` option.
>
> We've deprecated `default-profile` option—rather than immediately removing it—to provide backward compatibility and a chance to make your existing configuration conform to the changed configuration. You'll receive a warning message if you continue to use the `default-profile` option in your configuration. However existing deployments are not affected if you modify the current configuration. See profile (Juniper Secure Connect).

In existing deployments, to ensure a smooth transition with this change, we recommend that you modify the current configuration with profile *hr* to *ra.example.com/hr* or *192.168.1.10/hr* at the [`edit`] hierarchy using below commands -

- ```
  user@host# rename security remote-access profile hr to profile ra.example.net/hr
  ```
- ```
  user@host# rename edit security remote-access profile hr to profile 192.168.1.10/hr
  ```

For new configurations, consider the following scenarios to create a new remote-access profile based on how your end users connect using the Juniper Secure Connect application -

- If your end users connect using an *IP address*, specify the *IP address* in the profile name.
- If your end users connect using *FQDN*, specify the *FQDN* in the profile name.
- If you need to separate users with different realm values such as *hr*, append */hr* to the *IP address* or *FQDN* as follows -
  - [`edit security remote-access profile ra.example.net/hr`]
  - [`edit security remote-access profile 192.168.1.10/hr`]

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.3R1

The *default-profile* option is not available in SRX Series Firewalls starting Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Juniper Secure Connect Administrator Guide

default-rule

IN THIS SECTION

- Syntax | 5211
- Hierarchy Level | 5211
- Description | 5212
- Options | 5212
- Required Privilege Level | 5212
- Release Information | 5213

Syntax

```
default-rule {
    (deny [block-message] | permit | reject [block-message]);
}
```

Hierarchy Level

```
[edit security application-firewall rule-sets rule-set-name]
```

Description

Configure the default rule that defines the actions to be performed on a packet that does not match any defined rule.

An application firewall permits, rejects, or denies traffic based on the application of the traffic. The firewall consists of one or more rule sets with rules that specify match criteria, including dynamic applications, and the action to be taken for matching traffic. The application firewall rule set must contain a single default rule. The default rule defines the action to be taken for any traffic that does not match one of the rules.

Starting in Junos OS Release 18.2R1 application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration.

Options

- `deny`—Block the traffic at the firewall. The device drops the packet. No message is returned to the sender.
  - `block-message`—(Optional) In application firewall rules, provide information to the user regarding blocked traffic. Depending on the content of the `profile` option for this rule set, including the `block-message` option displays a default message or customized message, or redirects the user for denied HTTP or HTTPS traffic. All other traffic is dropped silently.
- `permit`—Permit traffic at the firewall.
- `reject`—Block the traffic at the firewall. For TCP traffic, by default the device drops the packet and returns a TCP reset (RST) message to the source host and to the server in some cases. For UDP and other protocol traffic, by default the device drops the packet and returns an ICMP “destination unreachable, port unreachable” message to both the client and the server.
  - `block-message`—(Optional) In application firewall rules, provide information to the user regarding blocked traffic. Depending on the content of the `profile` option for this rule set, including the `block-message` option displays a default message or customized message, or redirects the user for rejected HTTP or HTTPS traffic. All other traffic is dropped as specified in the default action for the `reject` option.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1. Statement updated in Junos OS Release 12.1X44-D10 with the `reject` option. The `block-message` option added in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Example: Configuring Application Firewall Rule Sets Within a Security Policy

default (Security Antivirus)

IN THIS SECTION

- Syntax | 5213
- Hierarchy Level | 5213
- Description | 5214
- Options | 5214
- Required Privilege Level | 5214
- Release Information | 5214

Syntax

```
default (block | log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name fallback-
options]
```

Description

All errors other than those specifically listed fall into this category. This could include either unhandled system exceptions (internal errors) or other unknown errors.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

default (Security Antivirus Sophos Engine)

IN THIS SECTION

- Syntax | 5215
- Hierarchy Level | 5215
- Description | 5215
- Options | 5215
- Required Privilege Level | 5215
- Release Information | 5215

Syntax

```
default (block | log-and-permit | permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name fallback-
options]
```

Description

All errors other than those specifically listed fall into this category. This could include either unhandled system exceptions (internal errors) or other unknown errors.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic
- `permit`—Permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

default (Security Content Security)

IN THIS SECTION

- Syntax | 5216
- Hierarchy Level | 5216
- Description | 5216
- Options | 5216
- Required Privilege Level | 5217
- Release Information | 5217

Syntax

```
default (block |log-and-permit | permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name]
```

Description

Specify the default action to take for a URL.

Options

- block—Log the error and deny the traffic.
- log-and-permit—Log the error and permit the traffic.
- permit—Permit the traffic.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

default (Security Web Filtering)

IN THIS SECTION

- Syntax | 5217
- Hierarchy Level | 5217
- Description | 5218
- Options | 5218
- Required Privilege Level | 5218
- Release Information | 5218

Syntax

```
default (block | log-and-permit | permit | quarantine);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated profile profile-name
fallback-settings]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name fallback-
settings]
```

```
[edit security utm feature-profile web-filtering juniper-local profile profile-name fallback-
settings]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name fallback-
settings]
```

Description

Specify an action for the profile, for requests that experience internal errors in the Web-filtering module.

Options

- block—Log the error and deny the traffic.
- log-and-permit—Log the error and permit the traffic.
- permit —Permit the traffic.
- quarantine—Show the warning message and permit/block the traffic based on user input.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Surf-Control feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 11.4.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

default-trusted-ca-certs (Security)

IN THIS SECTION

Syntax | 5219

- Hierarchy Level | 5219
- Description | 5219
- Options | 5220
- Required Privilege Level | 5220
- Release Information | 5221

Syntax

```
default-trusted-ca-certs {
    automatic-download {
        deactivate;
        interval {
             hours value;
        }
        url value;
        routing-instance value;
     }
  }
```

Hierarchy Level

```
[edit security pki]
```

Description

Dynamic update of trusted CA bundle requires -

- Downloading of trusted CA bundle from Juniper Networks security website, https://signatures.juniper.net/cacert or a custom URL.
- Uploading trusted CA bundle to PKI.
- Periodic polling of trusted CA bundle.

This functionality is configured using the statement `default-trusted-ca-certs`.

Options

automatic-download Sets automatic download of CA certs configuration.

deactivate Disables automatic download of default CA certs.

Use this option when automatic download is configured and you plan to disable it.

interval Specify default trusted CA certs automatic download interval.

- Value:
  - `hours`—Specify a value between 1-336 hours.
- Default:
  - If nothing is specified, it considers 24 hours as the interval.

routing-instance Specify a routing instance for trusted CA cert download. Use this option to configure non-default routing instance.

- Value:
  - *routing-instance-name*. Example: R1

url Specify HTTP URL for OCSP (Online Certificate Status Protocol) access location. This option sets the base URL for downloading trusted CA certs.

- Value:
  - *url*.
- Default:
  - If nothing is specified, default Juniper CDN server URL http://signatures.juniper.net/cacert is considered.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

- pki
- Dynamic Updated of CA Bundle

default-value

IN THIS SECTION

- Syntax | 5221
- Hierarchy Level | 5221
- Description | 5222
- Options | 5222
- Required Privilege Level | 5222
- Release Information | 5222

Syntax

```
default-value default-value;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name variables variable-name]
```

Description

Configure a default value for a user-defined variable in a dynamic profile. The values that the system uses for these variables are applied when the subscriber authenticates.

Options

default-value—Default value for the variable.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

User-Defined Variables

Using Variable Expressions in User-Defined Variables

Configuring User-Defined Dynamic Variables in Dynamic Profiles

default-vpn-source

IN THIS SECTION

- Syntax | 5223
- Hierarchy Level | 5223
- Description | 5223
- Default | 5223
- Required Privilege Level | 5223
- Release Information | 5223

Syntax

```
default-vpn-source {
    interface-name interface-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit protocols pim]
```

Description

Enable the router to use the primary loopback address configured in the default routing instance as the source address when PIM hello messages, join messages, and prune messages are sent over multicast tunnel interfaces for interoperability with other vendors' routers.

The remaining statements are explained separately. See CLI Explorer.

Default

By default, the router uses the loopback address configured in the VRF routing instance as the source address when sending PIM hello messages, join messages, and prune messages over multicast tunnel interfaces.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

interface-name | 8934

defaults

IN THIS SECTION

- Syntax | 5224
- Hierarchy Level | 5224
- Description | 5224
- Required Privilege Level | 5225
- Release Information | 5225

Syntax

```
defaults {
    (accounting | no-accounting);
    group-policy [ policy-names ];
    query-interval seconds;
    query-response-interval seconds;
    robust-count number;
    ssm-map ssm-map-name;
    version version;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name statement-name protocols igmp amt relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name statement-name
protocols igmp amt relay],
[edit protocols igmp amt relay],
[edit routing-instances routing-instance-name statement-name protocols igmp amt relay]
```

Description

Configure default IGMP attributes for all Automatic Multicast Tunneling (AMT) interfaces.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

defaults (ACX series, JRR200, MX204, PTX5000, and VRR)

IN THIS SECTION

- Syntax | 5225
- Hierarchy Level | 5226
- Description | 5226
- Required Privilege Level | 5226
- Release Information | 5226

Syntax

```
defaults {
    ebgp {
        no-policy {
            advertise (accept | reject | reject-always);
            receive (accept | reject | reject-always);
        }
    }
}
```

Hierarchy Level

```
[edit fabric protocols bgp],
[edit logical-systems name protocols bgp],
[edit logical-systems name routing-instances name protocols bgp],
[edit logical-systems name tenants name routing-instances name protocols bgp],
[edit protocols bgp],
[edit routing-instances name protocols bgp],
[edit tenants name routing-instances name protocols bgp]
```

Description

You can specify the route propagation of EBGP speakers when no explicit policy is configured.

Required Privilege Level

routing

Release Information

Hierarchy `defaults ebgp no-policy` introduced at `protocols bgp` in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Implicit filter for Default EBGP Route Propagation Behavior without Policies

defaults (early-aging)

IN THIS SECTION

- Syntax | 5227
- Hierarchy Level | 5227
- Description | 5227
- Required Privilege Level | 5227

Release Information | 5227

Syntax

```
defaults {
 cutoff-time {
  (percent percent | seconds seconds);
 }
}
```

Hierarchy Level

```
[edit security flow aging early-aging]
```

Description

It is the default cutoff time of early-ageout, it is not configured using the options like protocol or application.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

This statement is introduced in Junos OS Release 23.4R1.

defaults (Policy Options)

IN THIS SECTION

- Syntax | 5228
- Hierarchy Level | 5228
- Description | 5228
- Required Privilege Level | 5228
- Release Information | 5229

Syntax

```
defaults {
    route-filter (no-walkup | walkup);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options],
[edit logical-systems logical-system-name policy-options policy-statement policy-statement-name],
[edit policy-options],
[edit policy-options policy-statement policy-statement-name]
```

Description

Establish defaults for a particular policy statement or globally. Defaults include the walkup feature, which examines more than the longest match route filters in a policy statement term with more than one route filter, allowing consolidation of terms and a potential performance enhancement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

- no-walkup | 11623
- walkup | 31019
- route-filter | 15861
- *Walkup for Route Filters Overview*
- *Configuring Walkup for Route Filters to Improve Operational Efficiency*
- *Example: Configuring Walkup for Route Filters Globally to Improve Operational Efficiency*
- *Example: Configuring Walkup for Route Filters Locally to Improve Operational Efficiency*

defer-initial-multipath-build

IN THIS SECTION

- Syntax | 5229
- Hierarchy Level | 5230
- Description | 5230
- Options | 5231
- Required Privilege Level | 5231
- Release Information | 5231

Syntax

```
defer-initial-multipath-build {
    maximum-delay maximum-delay;
}
```

Hierarchy Level

```
[edit logical-systems name protocols bgp family inet unicast],
[edit logical-systems name protocols bgp family inet inet6 unicast],
[edit logical-systems name protocols bgp group name family inet unicast],
[edit logical-systems name protocols bgp group name family inet inet6 unicast],
[edit logical-systems name protocols bgp group name neighbor name family inet unicast],
[edit logical-systems name protocols bgp group name neighbor name family inet inet6 unicast],
[edit logical-systems name routing-instances name protocols bgp family inet unicast],
[edit logical-systems name routing-instances name protocols bgp family inet inet6 unicast],
[edit logical-systems name routing-instances name protocols bgp group name family inet unicast],
[edit logical-systems name routing-instances name protocols bgp group name family inet inet6
unicast],
[edit logical-systems name routing-instances name protocols bgp group name neighbor name family inet
unicast],
[edit logical-systems name routing-instances name protocols bgp group name neighbor name family inet
inet6 unicast],
[edit protocols bgp family inet unicast],
[edit protocols bgp family inet inet6 unicast],
[edit protocols bgp group name family inet unicast],
[edit protocols bgp group name family inet inet6 unicast],
[edit protocols bgp group name neighbor name family inet unicast],
[edit protocols bgp group name neighbor name family inet inet6 unicast],
[edit routing-instances name protocols bgp family inet unicast],
[edit routing-instances name protocols bgp family inet inet6 unicast],
[edit routing-instances name protocols bgp group name family inet unicast],
[edit routing-instances name protocols bgp group name family inet inet6 unicast],
[edit routing-instances name protocols bgp group name neighbor name family inet unicast],
[edit routing-instances name protocols bgp group name neighbor name family inet inet6 unicast]
```

Description

Defer initial multipath calculation until all BGP routes are received. When multipath is enabled, BGP inserts the route into the multipath queue each time a new route is added or whenever an existing route changes. When multiple paths are received through BGP add-path feature, BGP might calculate one multipath route multiple times. This slows down the BGP RIB (also known as the routing table) learning rate. With this feature enabled the router does not start BGP multipath calculation until end-of-RIB marker is received. Configure this option to delay multipath calculation.

Alternatively to delay the multipath calculation the BGP multipath job priority can be modified using "multipath-build-priority" on page 11107 configuration statement at `[edit protocols bgp]` hierarchy level.

Options

maximum-delay Specify a value in seconds to indicate the maximum time that a device must delay the multipath calculation after a peer is up.

- **Range:** `1` through `3600` seconds

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

multipath-build-priority | **11107**

Understanding BGP Multipath

definition

IN THIS SECTION

- Syntax | **5232**
- Hierarchy Level | **5232**
- Description | **5232**
- Options | **5232**
- Required Privilege Level | **5233**
- Release Information | **5233**

Syntax

```
definition {
    admin-group (exclude | include-all | include-any)
    (spf | strict-spf);
    exclude-srlg (test1 test2)
    metric-type (delay-metric | igp-metric | te-metric);
    priority priority;
    use-flex-algorithm-prefix-metric;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options flex-algorithm],
[edit logical-systems name routing-options flex-algorithm],
[edit logical-systems name tenants name routing-instances name routing-options flex-algorithm],
[edit routing-instances name routing-options flex-algorithm],
[edit routing-options flex-algorithm],
[edit tenants name routing-instances name routing-options flex-algorithm]
```

Description

Configure the flex-algorithm definition (FAD) and specify the parameters of the definition. IS-IS calculates the path based on these specified parameters of the FAD. We recommend configuring flexible algorithm on only a couple of routers to provide redundancy and to avoid conflicts.

Options

admin-group Specify the administrative groups to include or exclude along a specified path.

- Values:
  - exclude— Specify this option to exclude an admin-group, all of which must be absent.
  - include-all— Specify this option to include an admin-group, all of which must be present.
  - include-any— Specify this option to include an admin-group, one or more of which must be present.

spf Normal SPF of type 0

strict-spf Strict SPF

exclude-srlg *value* Specify SRLG constraints in an FAD,

metric-type Specify the metric type that you would like to use in your network.

- Values:
  - delay-metric— Specify this option to enable minimum unidirectional link delay metric (minimum latency).
  - igp-metric— Specify this option to use the IGP route metric instead of the traffic engineering metric.
  - te-metric— Specify this option to use the configured traffic engineering metric instead of the IGP metric if you have enabled traffic engineering on the device.

priority Specify a priority to the flexible algorithm advertisement. IS-IS prioritizes a particular FAD advertisement over another FAD based on your assigned priority.

- **Range:** 0 through 255

use-flex-algorithm-prefix-metric Specify this option to advertise FAD Flags sub-TLV with M flag set. Flex-Algo path computation will use FAPM only if the winner FAD has this flag set to true.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

flex-algorithm | **7221**

Configuring Flexible Algorithm for Segment Routing Traffic Engineering

How to Configure Flexible Algorithms in IS-IS for Segment Routing Traffic Engineering

show isis flex-algorithm | 21732

definition (Protocols OSPF)

IN THIS SECTION

- Syntax | 5234
- Hierarchy Level | 5234
- Description | 5235
- Options | 5235
- Required Privilege Level | 5235
- Release Information | 5235

Syntax

```
definition {
    (spf | strict-spf);
    metric-type (igp-metric | te-metric);
    priority priority;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options flex-algorithm],
[edit logical-systems name routing-options flex-algorithm],
[edit logical-systems name tenants name routing-instances name routing-options flex-algorithm],
[edit routing-instances name routing-options flex-algorithm],
[edit routing-options flex-algorithm],
[edit tenants name routing-instances name routing-options flex-algorithm]
```

Description

Configure the flex-algorithm definition (FAD) and specify the parameters of the definition. OSPFv2 calculates the path based on these specified parameters of the FAD. We recommend configuring flexible algorithm on only a couple of routers to provide redundancy and to avoid conflicts.

Options

spf Normal SPF of type 0

strict-spf Strict SPF

metric-type Specify the metric type that you would like to use in your network.

- Values:
  - igp-metric— Specify this option to use the IGP route metric instead of the traffic engineering metric.
  - te-metric— Specify this option to use the configured traffic engineering metric instead of the IGP metric if you have enabled traffic engineering on the device.

priority Specify a priority to the flexible algorithm advertisement. OSPFv2 prioritizes a particular FAD advertisement over another FAD based on your assigned priority.

- **Range:** 0 through 255

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Understanding OSPF Flexible Algorithm for Segment Routing

Example: OSPF Flexible Algorithm

flex-algorithm (Protocols OSPF)

definition (Protocols OSPF)

degraded

IN THIS SECTION

- Syntax | 5236
- Hierarchy Level | 5236
- Description | 5236
- Options | 5237
- Required Privilege Level | 5237
- Release Information | 5237

Syntax

```
degraded {
    action-fpc-restart-disable;
    degraded-fabric-detection-enable;
    degraded-fpc-bad-plane-threshold number-of-bad-planes;
}
```

Hierarchy Level

```
[edit chassis fabric]
```

Description

Configure options that apply to degraded chassis fabric conditions.

In Junos OS Evolved, for the fabric planes to be available for the Packet Forwarding Engines, the following conditions must be met:

- The associated SIBs are in online state.

- The fabric ASICs are initialized.
- The links between the Packet Forwarding Engine ASICs and the fabric ASICs are brought up.

However, events such as link faults, fabric ASIC faults, SIB faults, and boot time Fabric OAM errors could make the fabric planes unreachable for Packet Forwarding Engines. If you have configured the degraded fabric detection along with a threshold for the number of bad planes, when the number of unreachable planes becomes equal to or more than the threshold, the software reports a degraded fabric condition.

When a Packet Forwarding Engine is affected by a degraded fabric condition, the software disables the interfaces associated with that Packet Forwarding Engine.

Options

- `action-fpc-restart-disable`—Allow the user to disable restarting of the FPCs during healing from a degraded fabric condition. The device can automatically recover from degraded fabric conditions by restarting both the fabric planes and the FPCs. If the action-fpc-restart-disable statement is configured, the healing attempt is limited to restarting the fabric planes only.

  The system will detect a null-route filtering condition and try to heal the system.

  > **NOTE**: The `action-fpc-restart-disable` option is not applicable to Junos OS Evolved.

- `degraded-fabric-detection-enable`—Enable detection of an FPC with degraded fabric.
- `degraded-fpc-bad-plane-threshold` **`number-of-bad-planes`**—Configure the number of bad planes that indicate an FPC is degraded.

  Options: *number-of-bad-planes*—Number of bad planes.

  Range: The range could vary depending on the platform. For example, the PTX10001-36MR supports a range from 4 to 32.

  Default: 4

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement “action-fpc-restart-disable” introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Fabric Resiliency and Degradation

Disabling Line Card Restart to Limit Recovery Actions from Degraded Fabric Conditions

delay

IN THIS SECTION

- Syntax | 5238
- Hierarchy Level | 5238
- Description | 5238
- Options | 5239
- Required Privilege Level | 5239
- Release Information | 5239

Syntax

```
delay delay-value;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
profile-name calculation-weight]
```

Description

Configure the calculation weight for delay.

Options

delay-value—Calculation weight for delay.

> **NOTE**: This option is applicable only for two-way delay measurement.

- **Range:** 1 through 65,535
- **Default:** 1

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring an Iterator Profile

Configuring an Iterator Profile on a Switch (CLI Procedure)

calculation-weight | 3177

delay-advertise (DHCPv6)

IN THIS SECTION

- Syntax | 5240
- Hierarchy Level | 5240
- Description | 5241
- Options | 5241
- Required Privilege Level | 5242

Release Information | 5242

Syntax

```
delay-advertise {
    based-on (option-15 | option-16 | option-18 | option-37) {
        equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
        not-equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
        starts-with {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
    }
    delay-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server dhcpv6 group
group-name interface interface-name overrides]
[edit logical-systems name routing-instances name system services dhcp-local-server dhcpv6 group
group-name overrides]
[edit logical-systems name routing-instances name system services dhcp-local-server dhcpv6
overrides]

[edit logical-systems name system services dhcp-local-server dhcpv6 group group-name interface
interface-name overrides]
[edit logical-systems name system services dhcp-local-server dhcpv6 group group-name overrides]
[edit logical-systems name system services dhcp-local-server dhcpv6 overrides]

[edit routing-instances name system services dhcp-local-server dhcpv6 group group-name interface
```

```
interface-name overrides]
[edit routing-instances name system services dhcp-local-server dhcpv6 group group-name overrides]
[edit routing-instances name system services dhcp-local-server dhcpv6 overrides]

[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name overrides]
[edit system services dhcp-local-server dhcpv6 group group-name overrides]
[edit system services dhcp-local-server dhcpv6 overrides]
```

Description

Delay sending a DHCPv6-advertise message to specified clients. The configured server sends the advertise message when the delay timer expires. If the client is already bound to another server when it receives this advertise message, the DHCPv6 server that delayed the message releases the client. You can change the delay value at any time, but the change applies only to clients from which a solicit is received after the change; the new value does not apply to clients for which the response is already delayed.

Configuring the delay on DHCPv6 local servers enables load-balancing among multiple local servers on the network. When a client sends a solicit message, the delay prevents more than one server from replying at the same time. The delay applies on a per-client basis. You can configure the clients affected based on DHCPv6 option 15 (user class identifier), option 16 (vendor class identifier), option 18 (interface identifier; ACI). or option 37 (ARI).

If the configured server receives a second solicit message from the client, that means that no other server has responded. In this case, the configured server immediately replies to the client. This behavior enables the server to act as a back-up for other servers on the network.

Options

ascii-string Value of the option expressed as an ASCII string.

based-on Specify the DHCPv6 option received in the solicit message that is compared with the configured value on the local server.

- `option-15`—DHCPv6 option 15, user class identifier.
- `option-16`—DHCPv6 option 16, vendor class identifier.
- `option-18`—DHCPv6 option 18, interface identifier, equivalent to the Agent Circuit ID.
- `option-37`—DHCPv6 option 37, the Agent Remote ID.

seconds	Time delay between receiving DHCPv6 solicit message and responding to the client with an advertise message. • **Range:** 1 through 30 • **Default:** 3
equals	Specify that the received option value from the client must match the configured string.
hexadecimal-string	Value of the option expressed as a hexadecimal.
not-equals	Specify that the received option value from the client must not match the configured string.
starts-with	Specify that the received option value from the client must start with the configured string.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Delaying DHCP Offer and Advertise Responses to Load Balance DHCP Servers

delay-authentication (DHCP Relay Agent)

IN THIS SECTION

Syntax | 5243
Hierarchy Level | 5243

- Description | 5243
- Required Privilege Level | 5243
- Release Information | 5243

Syntax

```
delay-authentication;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay ... overrides],
[edit forwarding-options dhcp-relay dhcpv6 ... overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ... overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ... overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ... overrides]
```

Description

Delay authentication of subscribers until the DHCP client sends a Request packet. This conserves managed resources by delaying the authorization process and the creation of an entry in the subscriber database until the DHCP request processing phase.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Overriding the Default DHCP Relay Configuration Settings

delay-buffer-rate

IN THIS SECTION

- Syntax | 5244
- Hierarchy Level | 5244
- Description | 5244
- Default | 5245
- Options | 5245
- Required Privilege Level | 5245
- Release Information | 5246

Syntax

```
delay-buffer-rate (percent percentage | rate);
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name]
```

Description

For Gigabit Ethernet IQ, Channelized IQ PICs, and FRF.15 and FRF.16 LSQ interfaces only, base the delay-buffer calculation on a delay-buffer rate.

Default

If you do not include this statement, the delay-buffer calculation is based on the guaranteed rate if one is configured, or the shaping rate if no guaranteed rate is configured. For more information, see *Bandwidth and Delay Buffer Allocations by Configuration Scenario*.

Options

percent *percentage*

For LSQ interfaces, delay-buffer rate as a percentage of the available interface bandwidth.

- **Range:** 1 through 100 percent

rate

Delay-buffer rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000 bps

> **NOTE**: Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Starting with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.

> **NOTE**: Starting with Junos 17.2R1, the default minimum delay buffer threshold for MPCs is changed from 4096 bytes to 1568 bytes. This change improves the delay of time that packets can spend in the buffer. However, this lower minimum delay buffer threshold is easier to exceed than before and can cause tail drops. Juniper recommends you configure sufficient delay buffer rate to reduce this impact. In CIR mode, for nodes that do not have a guaranteed rate configured, Juniper also recommends that you configure a delay buffer rate equal to the shaping rate.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Oversubscribing Interface Bandwidth

Providing a Guaranteed Minimum Rate

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

output-traffic-control-profile | **12016**

delay-buffer-rate (Dynamic Traffic Shaping)

IN THIS SECTION

- Syntax | **5246**
- Hierarchy Level | **5246**
- Description | **5247**
- Default | **5247**
- Options | **5247**
- Required Privilege Level | **5247**
- Release Information | **5247**

Syntax

```
delay-buffer-rate (percent percentage | rate | $junos-cos-delay-buffer-rate);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

Base the delay-buffer calculation on a delay-buffer rate.

Default

If you do not include this statement, the delay-buffer calculation is based on the guaranteed rate if one is configured, or the shaping rate if no guaranteed rate is configured.

Options

rate—Delay-buffer rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000 bps

`$junos-cos-delay-buffer-rate`—Junos predefined variable that is replaced with the delay-buffer rate obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

The `$junos-cos-delay-buffer-rate` variable introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Traffic Scheduling and Shaping for Subscriber Access

output-traffic-control-profile (Dynamic CoS Definition) | **12018**

delay-factor

IN THIS SECTION

- Syntax | 5248
- Hierarchy Level | 5248
- Description | 5248
- Required Privilege Level | 5249
- Release Information | 5249

Syntax

```
delay-factor {
    no-syslog-generation;
    generate-snmp-traps;
    storm-control {
        count number;
        interval number;
    }
    alarm-mode {
        mdi-records-count number;
        average;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the maximum observed time difference between the arrival of media data and the drain of media data. The delay factor suggests the minimum size of the buffer required at the next downstream node. As a stream progresses, the variation of the delay factor indicates packet bunching or packet gaps

(jitter). Greater delay factor values also indicate that more network latency is needed to deliver a stream because of the need to pre-fill a receive buffer before beginning the drain to guarantee no underflow.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Inline Video Monitoring on MX Series Routers

alarms (Video Monitoring) | 2125

delay-measurement (Protocols IS-IS)

IN THIS SECTION

- Syntax | 5249
- Hierarchy Level | 5250
- Description | 5250
- Options | 5250
- Required Privilege Level | 5250

Syntax

```
delay-measurement {
    advertisement (Delay Measurement);normalize {
```

```
        interval value
        offset value
    }
    probe-count seconds;
    probe-interval seconds;
}
```

Hierarchy Level

```
[edit protocols isis interface name]
```

Description

Enable delay measurement.

Options

normalize Normalize Measured Link Delay

- interval *value*The value of the normalized interval in microseconds.

  Range: 1 through 5000000 microseconds)

- offset *value*The value of the normalized offset in microseconds (Must be < normalized interval).

  Range: 1 through 5000000 microseconds)

"probe-count" on page 13231 Probe count in seconds.

"probe-interval" on page 13233 Probe interval in seconds.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

RELATED DOCUMENTATION

Understanding Link Delay Measurement and Advertising in IS-IS

delay-measurement (Protocols OSPF)

IN THIS SECTION

- Syntax | 5251
- Hierarchy Level | 5252
- Description | 5252
- Options | 5252
- Required Privilege Level | 5252
- Release Information | 5252

Syntax

```
delay-measurement {
    advertisement {
       accelerated {
          threshold percentage;
}
       periodic {
          interval seconds;
          threshold percentage;
    }
  }
     probe-count seconds;
     probe-interval seconds;
}
```

Hierarchy Level

```
[edit protocols ospf interface name]
```

Description

Enable delay measurement.

Options

advertisement delay advertisement

probe-count Probe count (seconds)

- Range: 1 through 15
- Values:

probe-interval Probe interval (seconds)

- Range: 1 through 255

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 21.4R1

delay-metric

IN THIS SECTION

- Syntax | 5253
- Hierarchy Level | 5253
- Description | 5253

- Required Privilege Level | 5253

Syntax

```
delay-metric microseconds;
```

Hierarchy Level

```
[edit protocols isis interface name]
```

Description

Delay metric (microseconds)

Required Privilege Level

routing

delay-metric (Protocols OSPF)

IN THIS SECTION

- Syntax | 5254
- Hierarchy Level | 5254
- Description | 5254
- Required Privilege Level | 5254
- Release Information | 5254

Syntax

```
delay-metric microseconds;
```

Hierarchy Level

```
[edit protocols ospf interface name]
```

Description

Delay metric (microseconds)

Range:

- 0 through 16777215

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 21.4R1

delay-offer (DHCPv4)

IN THIS SECTION

- Syntax | 5255
- Hierarchy Level | 5255
- Description | 5256
- Options | 5256
- Required Privilege Level | 5257
- Release Information | 5257

Syntax

```
delay-offer {
    based-on (option-60 | option-77 | option-82) {
        equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
        not-equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
        starts-with {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
    }
    delay-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server group group-name interface interface-name overrides]
[edit logical-systems name routing-instances name system services dhcp-local-server group group-name overrides]
[edit logical-systems name routing-instances name system services dhcp-local-server overrides]

[edit logical-systems name system services dhcp-local-server group group-name interface interface-name overrides]
[edit logical-systems name system services dhcp-local-server group group-name overrides]
[edit logical-systems name system services dhcp-local-server overrides]

[edit routing-instances name system services dhcp-local-server group group-name interface interface-name overrides]
[edit routing-instances name system services dhcp-local-server group group-name overrides]
[edit routing-instances name system services dhcp-local-server overrides]

[edit system services dhcp-local-server group group-name interface interface-name overrides]
```

```
[edit system services dhcp-local-server group group-name overrides]
[edit system services dhcp-local-server overrides]
```

Description

Delay sending a DHCP-offer message to specified clients. The configured server sends the offer message when the delay timer expires. If the client is already bound to another server when it receives this offer message, the DHCPv4 server that delayed the message releases the client. You can change the delay value at any time, but the change applies only to clients from which an offer is received after the change; the new value does not apply to clients for which the response is already delayed.

Configuring the delay on DHCPv4 local servers enables load-balancing among multiple local servers on the network. When a client sends a discover message, the delay prevents more than one server from replying at the same time. The delay applies on a per-client basis. You can configure the clients affected based on DHCPv4 option 60 (vendor class identifier), option 77 (user class identifier), or option 82 (ACI and ARI).

If the configured server receives a second discover message from the client, that means that no other server has responded. In this case, the configured server immediately replies to the client. This behavior enables the server to act as a back-up for other servers on the network.

Options

ascii *ascii-string*

Value of the option expressed as an ASCII string.

based-on

Specify the DHCPv4 option received in the discover message that is compared with the configured value on the local server.

- `option-60`—DHCPv4 option 60, vendor class identifier.
- `option-77`—DHCPv4 option 77, the user class identifier.
- `option-82`—DHCPv4 option 82, the ACI/ARI.

delay-time *seconds*

Time delay between receiving DHCPv4 discover message and responding to the client with an offer message.

- **Range:** 1 through 30
- **Default:** 3

equals

Specify that the received option value from the client must match the configured string.

hexadecimal *hexadecimal-string*	Value of the option expressed as a hexadecimal.
not-equals	Specify that the received option value from the client must not match the configured string.
starts-with	Specify that the received option value from the client must start with the configured string.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Delaying DHCP Offer and Advertise Responses to Load Balance DHCP Servers

delay (PPPoE Service Name Tables)

IN THIS SECTION

- Syntax | 5258
- Hierarchy Level | 5258
- Description | 5258
- Options | 5258
- Required Privilege Level | 5258
- Release Information | 5258

Syntax

```
delay seconds;
```

Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name],
[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci
circuit-id-string ari remote-id-string]
```

Description

Configure the PPPoE underlying interface on the router to wait a specified number of seconds after receiving a PPPoE Active Discovery Initiation (PADI) control packet from a PPPoE client before sending a PPPoE Active Discovery Offer (PADO) packet to indicate that it can service the client request

The router (PPPoE server) does not check whether another server has already sent a PADO packet during the delay period in response to the PPPoE client's PADI packet. It is up to the PPPoE client to determine whether another PPPoE server has responded to its PADI request, or if it must respond to the delayed PADO packet to establish a PPPoE session.

Options

seconds—Number of seconds that the PPPoE underlying interface waits after receiving a PADI packet from a PPPoE client before sending a PADO packet in response.

- **Range:** 1 through 120 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at `[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci circuit-id-string ari remote-id-string]` hierarchy level introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

delay (querier)

IN THIS SECTION

- Syntax | 5259
- Hierarchy Level | 5259
- Description | 5260
- Required Privilege Level | 5260
- Release Information | 5260

Syntax

```
delay {
    traffic-class tc-value {
        average-sample-size sample size;
        padding-size size;
        query-interval milliseconds;
        rtt-delay-threshold rtt threshold value;
        twcd-delay-threshold twcd threshold value;
    }
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring
querier],
```

```
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring
querier]
```

Description

Configure delay measurement options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Pro-Active Loss and Delay Measurements

On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview

performance-monitoring (Protocols MPLS) | **12434**

delay (responder)

IN THIS SECTION

- Syntax | **5261**
- Hierarchy Level | **5261**
- Description | **5261**
- Options | **5261**
- Required Privilege Level | **5261**
- Release Information | **5261**

Syntax

```
delay {
    min-query-interval milliseconds;
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring responder],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring responder],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring
responder],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring
responder]
```

Description

Configure delay measurement options.

Options

min-query-interval *milliseconds*

(Optional) Specify the minimum query interval that the responder supports. If the minimum query interval of the responder is greater than the query interval configured at querier, the effective message query rate will be the minimum query interval configured for the responder.

- **Default:** 10 seconds
- **Range:** 1000 through 4294967295 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Pro-Active Loss and Delay Measurements

On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview

performance-monitoring (Protocols MPLS) | **12434**

delay-route-advertisements

IN THIS SECTION

- Syntax | **5262**
- Hierarchy Level | **5263**
- Description | **5263**
- Options | **5263**
- Required Privilege Level | **5264**
- Release Information | **5264**

Syntax

```
delay-route-advertisements {
    minimum-delay {
        routing-uptime routing-uptime;
        inbound-convergence inbound-convergence;
     }
    maximum-delay {
        route-age routing-age;
        routing-uptime routing-uptime;
     }
    always-wait-for-krt-drain;
}
```

Hierarchy Level

```
[edit logical-systems name protocols bgp family name],
[edit logical-systems name protocols bgp group name family name],
[edit logical-systems name protocols bgp group name neighbor ip-address family name],
[edit logical-systems name routing-instances name protocols bgp family name],
[edit logical-systems name routing-instances name protocols bgp group name family name],
[edit logical-systems name routing-instances name protocols bgp group name neighbor ip-address
family name],
[edit protocols bgp familyname],
[edit protocols bgp group name family name],
[edit protocols bgp group name neighbor ip-address familyname],
[edit routing-instances name protocols bgp family name],
[edit routing-instances name protocols bgp group name family name],
[edit routing-instances name protocols bgp neighbor ip-address family name]
```

Description

Configure this option to delay route updates for a specified family until the forwarding table is synchronized. When a device starts up, BGP establishes peering sessions with its neighbors and receives route updates. These route updates are then readvertised as more specific BGP routes or less specific aggregates. Advertising routes prematurely, that is, before all the available routes are installed in the forwarding table, might result in traffic loss.

In multihomed networks, this behavior might cause unnecessary loss of service when a BGP session at the primary provider edge comes up. This problem is more pronounced when the primary provider edge device advertises route aggregates, because few aggregate prefixes can be announced more quickly to the network peers than a full routing table with thousands of more specific prefixes to the forwarding table. In order to avoid this problem, the device must delay a BGP route advertisement until the associated forwarding state is installed into the forwarding table. This feature allows a Junos OS device to do so, and allows you to configure the minimum and maximum delay periods.

Options

minimum-delay (Optional) Specify a minimum delay, in seconds, in advertising the routes.

inbound-convergence *inbound-convergence* (Optional) Specify a minimum delay in route advertisement after the source peer has sent all route updates. The device waits at least for the configured duration after inbound convergence has completed at the source of the route. For BGP routes, the source peer sends the initial route updates, for example after end-of-rib is received.

- **Default:** 120 seconds
- **Range:** 1 through 36000 seconds

routing-uptime *routing-uptime*

(Optional) Specify the minimum delay, in seconds, before sending a route advertisement after the routing protocol process (rpd) starts. The device waits for at least the configured duration before sending out route advertisements to its peers.

- **Default:** 0 seconds
- **Range:** 1 through 36000 seconds

maximum-delay

(Optional) Specify a maximum delay, in seconds, before advertising routes to peers.

route-age *routing-age*

(Optional) Specify a maximum delay in sending a route advertisement after route aggregates have been created, that is, the route age. The device suspends waiting for the routes to be downloaded to the forwarding table at the configured route age and starts sending route advertisements to its peers.

- **Default:** 0 seconds
- **Range:** 1 through 36000 seconds

routing-uptime *routing-uptime*

(Optional) Specify the maximum delay in seconds before sending a route advertisement after the routing protocol process (rpd) starts. The device does not wait more than the configured duration before sending out route advertisements to its peers.

- **Default:** 0 seconds
- **Range:** 1 through 36000 seconds

always-wait-for-krt-drain

Delay route advertisement until KRT queue is drained.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F6.

RELATED DOCUMENTATION

routing-instances (Multiple Routing Entities) | 15976

delay-variation

IN THIS SECTION

- Syntax | 5265
- Hierarchy Level | 5265
- Description | 5265
- Options | 5265
- Required Privilege Level | 5266
- Release Information | 5266

Syntax

```
delay-variation delay-variation-value;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
profile-name calculation-weight]
```

Description

Configure the calculation weight for delay variation.

Options

delay-variation-value—Calculation weight for delay variation.

NOTE: This option is applicable only for two-way delay measurement.

- **Range:** 1 through 65,535
- **Default:** 1

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring an Iterator Profile

Configuring an Iterator Profile on a Switch (CLI Procedure)

calculation-weight | **3177**

delegate-probes

IN THIS SECTION

- Syntax | **5267**
- Hierarchy Level | **5267**
- Description | **5267**
- Required Privilege Level | **5267**
- Release Information | **5268**

Syntax

```
delegate-probes;
```

Hierarchy Level

```
[edit services rpm probe owner]
```

Description

Generate real-time performance monitoring (RPM) probes on an MS-MPC or MS-MIC card, which increases the number of RPM probes that can run at the same time.

To use the `delegate-probes` statement, you must first configure the `destination-interface` statement at the `[edit services rpm probe owner test test-name]` hierarchy level to point to a valid logical unit number of a multiservices interface. Then configure the same unit and multiservice interface with the `rpm client-delegate-probes` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level.

The `probe-type type` at the `[edit services rpm probe owner test test-name]` hierarchy level can be `icmp-ping` or `icmp-ping-timestamp` starting in Junos OS Release 17.3R1, and `icmp6-ping` starting in Junos OS Release 18.1R1.

To avoid packet bursts in the network due to RPM, probes will be distributed in a better way.

The chances of multiple tests starting and ending at the same time are smaller. This way RPM syslog bursts and a potential performance bottleneck in `event-processing` are avoided.. This does not exclude potential syslog drops on the RE if more than 12000 RPM tests are running simultaneously. For scaled configurations (with more than 12000 RPM tests) we recommend you to configure syslogs to sent to an external hosts for offloaded processing.

> **NOTE**: You cannot configure the `routing-instance` statement at the `[edit services rpm probe owner test test-name]` hierarchy level for RMP probes that are generated on an MS-MPC or MS-MIC card.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

client-delegate-probes | 4307

delegate-processing (VRRP)

IN THIS SECTION

- Syntax | 5268
- Hierarchy Level | 5268
- Description | 5269
- Options | 5269
- Required Privilege Level | 5269
- Release Information | 5269

Syntax

```
delegate-processing {
    ae-irb;
}
```

Hierarchy Level

```
[edit protocols vrrp]
```

Description

Configure the distributed periodic packet management process (ppmd) to send Virtual Router Redundancy Protocol (VRRP) advertisements .

Using a hash logic based on ifIIndex, the vrrp group ID, and the IP version, select one of the Flexible OIC Concentrators (FPCs) for distribution. The selected FPC is called the *anchor FPC*. All transmit instances and receive instances are from and to the anchor FPC. The anchor FPC is static, and VRRP is not guaranteed to get distributed to all available FPCs uniformly for all VRRP sessions.

Options

`ae-irb` Enable distributed ppmd for VRRP over aggregated Ethernet and integrated routing and bridging (IRB) interfaces.

Using the `ae-irb` option is only for MPC line cards. `ae-irb` is not supported on small MX Series routing devices with built-in MPCs such as the MX104 and below. Using the `ae-irb` option requires use of the `enhanced-ip` mode.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

`ae-irb` option introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Enabling the Distributed Periodic Packet Management Process for VRRP

delegate-server-processing

IN THIS SECTION

- Syntax | 5270
- Hierarchy Level | 5270
- Description | 5270
- Required Privilege Level | 5270
- Release Information | 5270

Syntax

```
delegate-server-processing;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
```

Description

For Ethernet interfaces on MX Series routers , enable server-side processing for two-way delay measurement and loss measurement.

By default, the processing is done by the Routing Engine.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Ethernet Frame Delay Measurements Overview

delegated-pool (DHCP Local Server)

IN THIS SECTION

- Syntax | 5271
- Hierarchy Level | 5271
- Description | 5272
- Options | 5272
- Required Privilege Level | 5272
- Release Information | 5272

Syntax

```
delegated-pool pool-name;
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 ...],
[edit logical-systems logical-system-name system services system services dhcp-local-server
dhcpv6 ...],
[edit routing-instances routing-instance-name system services system services dhcp-local-server
dhcpv6 ...]
```

Description

Specify the address pool that assigns the IA_PD address. A pool specified by RADIUS VSA 26-161 takes precedence over the pool specified by this `delegated-pool` statement.

Options

pool-name	Name of the address-assignment pool.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Specifying the Delegated Address-Assignment Pool to Be Used for DHCPv6 Prefix Delegation

Overriding the Default DHCP Local Server Configuration Settings

delegation-cleanup-timeout

IN THIS SECTION

- Syntax | 5273
- Hierarchy Level | 5273
- Description | 5273
- Options | 5273
- Required Privilege Level | 5273
- Release Information | 5274

Syntax

```
delegation-cleanup-timeout seconds;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Description

Specify the amount of time (in seconds) that a Path Computation Client (PCC) must wait before returning control of all LSPs to the routing protocol process after a PCEP session with the main active stateful Path Computation Element (PCE) is disconnected.

> **NOTE**: In compliance with *draft-ietf-pce-stateful-pce-09*, revoking of PCE-initiated LSP delegations by a PCC happens in a make-before-break fashion before the LSPs are redelegated to an alternate PCE. Starting in Junos OS Release 18.1R1, the `lsp-cleanup-timer` must be greater than or equal to the `delegation-cleanup-timeout` for the PCC to revoke the LSP delegations. If not, the redelegation timeout interval for the PCC can be set to infinity, where the LSP delegations to that PCE remain intact until specific action is taken by the PCC to change the parameters set by the PCE.

Options

seconds Time (in seconds) that a PCC must wait before returning control of all LSPs to the routing protocol process after a PCEP session with the main active stateful PCE is disconnected.

A value of 0 indicates immediate delegation cleanup.

- **Range:** 0 through 2,147,483,647 seconds

  Prior to Junos OS Release 18.4R1, the maximum range value is 600 seconds.

- **Default:** 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

pce | 12303

delegation-priority

IN THIS SECTION

- Syntax | 5274
- Hierarchy Level | 5275
- Description | 5275
- Options | 5275
- Required Privilege Level | 5275
- Release Information | 5275

Syntax

```
delegation-priority priority-number;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Description

Specify the priority number of the active stateful Path Computation Element (PCE). This value is used by the Path Computation Client (PCC) to elect a PCE to delegate LSPs. No two PCEs can have the same delegation-priority value. The PCC elects the PCE with a lower priority as the main active stateful PCE to delegate LSPs.

Options

priority-number Priority number of the active stateful PCE.

- **Range:** 1 through 65535
- **Default:** 0 (no priority is set)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

pce | 12303

delete

IN THIS SECTION

- Syntax | 5276
- Description | 5276
- Options | 5276
- Required Privilege Level | 5277
- Sample Output | 5277
- Release Information | 5277

Syntax

```
delete <statement-path> <identifier>
```

Description

Delete a statement or identifier. All subordinate statements and identifiers contained within the specified statement path are deleted with it.

Deleting a statement or an identifier effectively “unconfigures” or disables the functionality associated with that statement or identifier.

If you do not specify *statement-path* or *identifier*, the entire hierarchy, starting at the current hierarchy level, is removed.

NOTE: For Junos OS Evolved, if you use the `delete` configuration command at the top level of the configuration, you cannot commit the resulting empty configuration. At a minimum, the root authentication password is required.

Options

statement-path (Optional) Path to an existing statement or identifier. Include this if the statement or identifier to be deleted is not at the current hierarchy level.

identifier (Optional) Name of the statement or identifier to delete.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Sample Output

delete (Junos OS Evolved)

```
[edit]
user@host# delete
This will delete the entire configuration
Delete everything under this level? [yes,no] (no) yes
```

If you then try to commit the change, you get this:

```
user@host# commit
error: cannot commit an empty configuration
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

deactivate | 5142

Deleting a Statement from a Device Configuration

delete-after-commit (JNU Satellites)

IN THIS SECTION

- Syntax | 5278
- Hierarchy Level | 5278
- Description | 5278
- Required Privilege Level | 5278
- Release Information | 5279

Syntax

```
delete-after-commit;
```

Hierarchy Level

```
[edit system autoinstallation]
```

Description

Specify that during the subsequent commit operation of configuration settings (after the autoinstallation process successfully retrieves, installs, and commits the configuration), the autoinstallation configuration parameters be removed from the router. Removal of the autoinstallation parameters and statements from the committed configuration on the router ensures that the router does not attempt to perform an autoinstallation process when it is powered on the next time. Although you can optionally specify the interfaces to perform autoinstallation or configuration servers from which the files are to be downloaded, you must include the `delete-after-commit` statement to prevent the router from entering a recursive loop and repeatedly performing an autoinstallation every time it is powered on.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

- Autoinstallation of Satellite Devices in a Junos Node Unifier Group
- Autoinstallation Process on Satellite Devices in a Junos Node Unifier Group
- *Configuring Autoinstallation of Configuration Files*
- *Configuring Autoinstallation of Configuration Files*
- autoinstallation (JNU Satellite Devices) | **2755**

delete-binding-on-renegotiation (DHCP Local Server and Relay Agent)

IN THIS SECTION

- Syntax | **5279**
- Hierarchy Level | **5279**
- Description | **5280**
- Required Privilege Level | **5280**
- Release Information | **5281**

Syntax

```
delete-binding-on-renegotiation;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
```

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server overrides],
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides]
```

Description

Configure DHCP to override the default behavior when the local server or relay agent receives a DHCPv4 Discover or DHCPv6 Solicit message while in a bound state. In this case, DHCP drops the message and it is not processed. On a DHCP relay agent, the agent sends a Release message to the local server. DHCP cleans up the existing session and deletes the existing client entry, removing the binding. When a second Discover or Solicit message is received from the client, the message is processed and DHCP negotiation proceeds.

A consequence of the override behavior is that the time to complete a DHCP negotiation is prolonged if the client begins negotiation before the existing client entry is expired. The delay can be up to several seconds.

The default behavior (this statement is not configured) is that DHCP maintains the client entry if it receives a Discover or Solicit message that has a client ID that matches the existing client. DHCP then processes the new message using the existing client entry and responds to the client with an Offer or Advertise message.

NOTE: In releases earlier than Junos OS Release 15.1, the default behavior for DHCPv6 local server and relay agent is the same as the override behavior in Junos OS Release 15.1 and later. For any release, the default behavior for DHCPv4 local server and relay agent is to maintain the existing client entry and respond without waiting for a second Discover or Solicit message.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

DHCP Behavior When Renegotiating While in Bound State

delete chassis fpc

IN THIS SECTION

- Syntax | 5281
- Description | 5281
- Options | 5282
- Usage | 5282
- Required Privilege Level | 5282
- Output Fields | 5282
- delete chassis fpc (PTX) | 5282
- Release Information | 5282

Syntax

```
delete chassis fpc fpc_id no-vlan-tagging
```

Description

(PTX Series routers) Disable no VLAN tagging for a specific FPC.

Options

fpc *fpc_id* For PTX devices, *fpc_id* value is 0 to 15.

Usage

Use `set chassis fpc <fpc_id> no-vlan-tagging` command to re-enable **no-vlan-tagging** on a specific FPC.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

delete chassis fpc (PTX)

```
user@host> delete chassis fpc 1 no-vlan-tagging
```

Release Information

Command introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

fpc (Chassis) | 7501

delete-req

IN THIS SECTION

Syntax | 5283

- Hierarchy Level | 5283
- Description | 5283
- Options | 5283
- Required Privilege Level | 5284
- Release Information | 5284

Syntax

```
delete-req {
    alarm-threshold (Security GPRS) {
        forward number;
        reverse number;
    }
    drop-threshold (Security GPRS) {
        forward number;
        reverse number;
    }
}
```

Hierarchy Level

```
[edit security gtp profile profile-name path-rate-limit message-type]
```

Description

Configure the delete-req parameter to limit the number of packets per second for the GPRS tunneling protocol (GTP) delete request.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name path-rate-limit message-type]` hierarchy level is replaced by `[edit security gtp profile profile-name path-rate-limit message-type]`.

Options

`alarm-threshold`—To set alarm threshold for path rate limiting.

`drop-threshold`—To set drop threshold for path rate limiting.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding Path Rate Limiting for GTP Control Messages

Example: Limiting the Message Rate and Path Rate for GTP Control Messages

delete services jinsightd subscribe health-monitor

IN THIS SECTION

- Syntax | 5284
- Description | 5285
- Options | 5285
- Required Privilege Level | 5285
- Release Information | 5285

Syntax

```
delete services jinsightd subscribe health-monitor
```

Description

Disables the J-Insight health monitor. Starting in Junos OS Release 18.2R1, J-Insight provides health monitoring capabilities for FPC FRUs on the MX series routers. The health monitor is disabled by default.

Options

This command has no options.

Required Privilege Level

system

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

J-Insight Device Monitor Basic Configuration

set services jinsightd subscribe health-monitor | **16892**

delimiter

IN THIS SECTION

- Syntax | **5286**
- Hierarchy Level | **5286**
- Description | **5286**
- Required Privilege Level | **5286**
- Release Information | **5286**

Syntax

```
delimiter delimiter-character;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication username-include]
```

Description

Specify the character used as the delimiter between the concatenated components of the username. You cannot use the semicolon (;) as a delimiter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

delimiter (Access Profile)

IN THIS SECTION

Syntax | 5287

- Hierarchy Level | 5287
- Description | 5287
- Default | 5287
- Options | 5287
- Required Privilege Level | 5288
- Release Information | 5288

Syntax

```
delimiter delimiter;
```

Hierarchy Level

```
[edit access profile profile-name session-options strip-user-name]
```

Description

Specify up to eight characters that the router uses to determine which part of the subscriber login string to discard—leaving the remainder for use as a modified username—when subscriber username stripping is configured in a subscriber access profile. The characters to the right of the delimiter are discarded along with the delimiter. Use the "parse-direction" on page 12176 statement when more than one delimiter appears in a username to determine the characters that are stripped by identifying the desired delimiter. A given subscriber login string can result in multiple different modified usernames depending on the number and placement of delimiters and the direction of stripping.

Default

None. You must always configure a delimiter.

Options

delimiter Character that specifies the boundary between the part of the original username that is kept and the part that is discarded.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access

Configuring Username Modification for Subscriber Sessions

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

delimiter (DHCP Local Server)

IN THIS SECTION

- Syntax | 5288
- Hierarchy Level | 5289
- Description | 5290
- Options | 5290
- Required Privilege Level | 5290
- Release Information | 5290

Syntax

```
delimiter delimiter-character;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication username-
include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  authentication
username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group group-
name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify the character used as the delimiter between the concatenated components of the username.

Options

delimiter-character—Character that separates components that make up the concatenated username. You cannot use the semicolon (;) as a delimiter.

- **Default:** . (period)

> **NOTE**: When you include the "interface-description " on page 8854 in the username, the delimiter must not be a character that is part of the interface description. For example, if the text description is configured as “Backbone connection/PHL01”, you cannot use the forward slash (/) as the delimiter.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

delimiter (DHCP Relay Agent)

IN THIS SECTION

- Syntax | **5291**
- Hierarchy Level | **5291**

- Description | 5292
- Options | 5292
- Required Privilege Level | 5292
- Release Information | 5292

Syntax

```
delimiter delimiter-character;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-
include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-
include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication
username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
```

```
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
authentication username-include]
```

Description

Specify the character used as the delimiter between the concatenated components of the username. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Options

`delimiter-character`—Character that separates components that make up the concatenated username. You cannot use the semicolon (;) as a delimiter.

- **Default:** . (period)

> **NOTE**: When you include the "interface-description " on page 8854 in the username, the delimiter must not be a character that is part of the interface description. For example, if the text description is configured as “Backbone connection/PHL01”, you cannot use the forward slash (/) as the delimiter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Specifying Authentication Support

Creating Unique Usernames for DHCP Clients

delimiter (Domain Map)

IN THIS SECTION

- Syntax | 5293
- Hierarchy Level | 5293
- Description | 5293
- Default | 5293
- Options | 5294
- Required Privilege Level | 5294
- Release Information | 5294

Syntax

```
delimiter [delimiter-character];
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the characters that the router uses to separate usernames from domain names.

Default

The `@` character.

Options

delimiter-character—One or more characters used as delimiters. You can specify a maximum of eight delimiters. You cannot use the semicolon (;) as a delimiter. Do not include spaces between characters.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Specifying Domain and Realm Name Delimiters

Configuring Domain and Realm Name Usage for Domain Maps

demand-circuit (Protocols RIP)

IN THIS SECTION

- Syntax | 5295
- Hierarchy Level | 5295
- Description | 5295
- Default | 5295
- Required Privilege Level | 5295
- Release Information | 5295

Syntax

```
demand-circuit;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip group group-name],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip group group-name],
[edit protocols rip group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip group group-name],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

Configure a neighboring interface to act as a RIP demand circuit. To complete the demand circuit, you must configure both ends of the pair as demand circuits. When configured, the device sends RIP information only when changes occur in the routing database.

Default

Disabled. You must explicitly configure two neighboring interfaces to act as a RIP demand circuit.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 11.1 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring RIP Demand Circuits

RIP Demand Circuits Overview

max-retrans-time | 10464

demux-destination (Demux Interface)

IN THIS SECTION

- Syntax | 5296
- Hierarchy Level | 5296
- Description | 5297
- Required Privilege Level | 5297
- Release Information | 5297

Syntax

```
demux-destination {
    destination-prefix;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family],
[edit logical-systems logical-system-name routing-instances routing-instance-name interfaces
interface-name unit logical-unit-number family family]
```

Description

Configure one or more logical demultiplexing (demux) destination prefixes. The prefixes are matched against the destination address of packets that the underlying interface receives. When a match occurs, the packet is processed as if it was received on the demux interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for aggregated Ethernet added in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring an IP Demultiplexing Interface

Configuring a VLAN Demultiplexing Interface

demux-destination (Underlying Interface)

IN THIS SECTION

- Syntax | 5298
- Hierarchy Level | 5298
- Description | 5298
- Required Privilege Level | 5298
- Release Information | 5298

Syntax

```
demux-destination family;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name routing-instances routing-instance-name interfaces
interface-name unit logical-unit-number]
```

Description

Configure the logical demultiplexing (demux) destination family type on the IP demux underlying interface.

NOTE: The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for aggregated Ethernet added in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring an IP Demultiplexing Interface

Configuring a VLAN Demultiplexing Interface

demux (Interfaces)

IN THIS SECTION

- Syntax | 5299
- Hierarchy Level | 5300
- Description | 5300
- Required Privilege Level | 5301
- Release Information | 5301

Syntax

```
demux {
    inet {
        address source;
        auto-configure {
            address-ranges {
                authentication {
                    password password-string;
                    username-include {
                        auth-server-realm realm-string;
                        delimiter delimiter-character;
                        domain-name domain-name;
                        interface-name;
                        source-address;
                        user-prefix user-prefix-string;
                    }
                }
                dynamic-profile profile-name {
                    network ip-address {
                        range name {
                            low lower-limit;
                            high upper-limit;
                        }
                    }
                }
            }
```

```
        }
    }
    inet6 {
        address source;
        auto-configure {
            address-ranges {
                authentication {
                    password password-string;
                    username-include {
                        auth-server-realm realm-string;
                        delimiter delimiter-character;
                        domain-name domain-name;
                        interface-name;
                        source-address;
                        user-prefix user-prefix-string;
                    }
                }
                dynamic-profile profile-name {
                    network ip-address {
                        range name {
                            low lower-limit;
                            high upper-limit;
                        }
                    }
                }
            }

        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number]
```

Description

Configure demultiplexing (demux) interface options. The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles

demux-options (All Demux Interfaces)

IN THIS SECTION

- Syntax | 5301
- Hierarchy Level | 5302
- Description | 5302
- Required Privilege Level | 5302
- Release Information | 5302

Syntax

```
demux-options {
    use-underlying-interface-mac {
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure demultiplexing (demux) interface options for all demux interfaces.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

> *Configuring the IPv6 Link-Local Address for Dynamic Demux Interfaces over Static Demux VLAN Interfaces*

demux-options (Dynamic Interface)

IN THIS SECTION

- Syntax | 5303
- Hierarchy Level | 5303
- Description | 5303
- Required Privilege Level | 5303
- Release Information | 5303

Syntax

```
demux-options {
    underlying-interface interface-name
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 interface-name unit logical-unit-number]
```

Description

Configure logical demultiplexing (demux) interface options in a dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

> *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
>
> Demultiplexing Interface Overview

demux-options (Static Interface)

IN THIS SECTION

- Syntax | 5304
- Hierarchy Level | 5304
- Description | 5304
- Required Privilege Level | 5304
- Release Information | 5305

Syntax

```
demux-options {
    underlying-interface interface-name
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure logical demultiplexing (demux) interface options.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring an IP Demultiplexing Interface

Configuring a VLAN Demultiplexing Interface

demux-source (Demux Interface)

IN THIS SECTION

- Syntax | 5305
- Hierarchy Level | 5305
- Description | 5306
- Required Privilege Level | 5306
- Release Information | 5306

Syntax

```
demux-source {
    source-prefix;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family],
[edit logical-systems logical-system-name routing-instances routing-instance-name interfaces
interface-name unit logical-unit-number family family]
```

Description

Configure one or more logical demultiplexing (demux) source prefixes. The prefixes are matched against the source address of packets that the underlying interface receives. When a match occurs, the packet is processed as if it was received on the demux interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for aggregated Ethernet added in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring an IP Demultiplexing Interface

Configuring a VLAN Demultiplexing Interface

demux-source (Dynamic IP Demux Interface)

IN THIS SECTION

- Syntax | 5307
- Hierarchy Level | 5307
- Description | 5307
- Options | 5307
- Required Privilege Level | 5307
- Release Information | 5307

Syntax

```
demux-source {
    source-address;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family]
```

Description

Configure a logical demultiplexing (demux) source address for a subscriber in a dynamic profile.

Options

source-address—Either the specific source address you want to assign to the subscriber interface or the source address variable. For IPv4, specify `$junos-subscriber-ip-address`; for IPv6, specify `$junos-subscriber-ipv6-address`. The source address for the interface is dynamically supplied by DHCP when the subscriber accesses the router.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles

Demultiplexing Interface Overview

demux-source (Dynamic Underlying Interface)

IN THIS SECTION

- Syntax | 5308
- Hierarchy Level | 5308
- Description | 5308
- Options | 5308
- Required Privilege Level | 5309
- Release Information | 5309

Syntax

```
demux-source family;
```

Hierarchy Level

```
[edit dynamic-profiles interfaces interface-name unit logical-unit-number]
```

Description

Configure the logical demultiplexing (demux) source family type on the IP demux underlying interface within a dynamic profile.

NOTE: The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

Options

family—Protocol family:

- `inet`—Internet Protocol version 4 suite

- `inet6`—Internet Protocol version 6 suite

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

demux-source (Underlying Interface)

IN THIS SECTION

- Syntax | 5309
- Hierarchy Level | 5309
- Description | 5310
- Options | 5310
- Required Privilege Level | 5310
- Release Information | 5310

Syntax

```
demux-source family;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name interfaces
interface-name unit logical-unit-number]
```

Description

Configure the logical demultiplexing (demux) source family type on the IP demux underlying interface.

> **NOTE**: The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

Options

family—Protocol family:

- `inet`—Internet Protocol version 4 suite
- `inet6`—Internet Protocol version 6 suite

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for aggregated Ethernet added in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring an IP Demultiplexing Interface

Configuring a VLAN Demultiplexing Interface

demux0 (Dynamic Interface)

IN THIS SECTION

- Syntax | 5311
- Hierarchy Level | 5312
- Description | 5312
- Required Privilege Level | 5312
- Release Information | 5312

Syntax

```
demux0 {
    unit logical-unit-number {
        demux-options {
            underlying-interface interface-name
        }
        family family {
            access-concentrator name;
            address address;
            demux-source {
                source-prefix;
            }
            direct-connect;
            duplicate-protection;
            dynamic-profile profile-name;
            filter {
                input filter-name;
                output filter-name;
            }
            mac-validate (loose | strict):
            max-sessions number;
            max-sessions-vsa-ignore;
            rpf-check {
                fail-filter filter-name;
                mode loose;
            }
```

```
            service-name-table table-name
            short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-
seconds>;
            unnumbered-address interface-name <preferred-source-address address>;
        }
        filter {
            input filter-name;
            output filter-name;
        }
        vlan-id number;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces]
```

Description

Configure the logical demultiplexing (demux) interface in a dynamic profile.

Logical IP demux interfaces do not support IPv4 and IPv6 dual stack.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

> *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
>
> Demultiplexing Interface Overview

demux0 (Static Interface)

IN THIS SECTION

- Syntax | 5313
- Hierarchy Level | 5314
- Description | 5314
- Required Privilege Level | 5314
- Release Information | 5314

Syntax

```
demux0 {
    unit logical-unit-number {
        demux-options {
            underlying-interface interface-name
        }
        family family {
            access-concentrator name;
             {
                destination-prefix;
            }
            direct-connect;
            duplicate-protection;
            dynamic-profile profile-name;
             {
                source-prefix;
            }
            max-sessions number;
            service-name-table table-name
            targeted-distribution;
            unnumbered-address interface-name <preferred-source-address address>;
        }
        vlan-id number;
        vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
```

```
    }
}
```

Hierarchy Level

```
[edit interfaces],
[edit logical-systems logical-system-name interfaces]
```

Description

Configure the logical demultiplexing (demux) interface.

Logical IP demux interfaces do not support IPv4 and IPv6 dual stack.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring an IP Demultiplexing Interface

Configuring a VLAN Demultiplexing Interface

dense-groups

IN THIS SECTION

- Syntax | 5315
- Hierarchy Level | 5315
- Description | 5315
- Options | 5315
- Required Privilege Level | 5316
- Release Information | 5316

Syntax

```
dense-groups {
    addresses;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Configure which groups are operating in dense mode.

Options

addresses—Address of groups operating in dense mode.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring PIM Sparse-Dense Mode Properties

deny (Security Policies)

IN THIS SECTION

- Syntax | 5316
- Hierarchy Level | 5316
- Description | 5317
- Required Privilege Level | 5317
- Release Information | 5317

Syntax

```
deny;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then]
```

Description

Block the service at the firewall. The device drops the packets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

deny (Security SIP)

IN THIS SECTION

- Syntax | 5317
- Hierarchy Level | 5318
- Description | 5318
- Options | 5318
- Required Privilege Level | 5318
- Release Information | 5319

Syntax

```
deny {
    all {
        timeout value-in-seconds;
```

```
    }
    destination-ip address;
    timeout value-in-seconds;
}
```

Hierarchy Level

```
[edit logical-systems name security alg sip application-screen protect],
[edit logical-systems name tenants name security alg sip application-screen protect],
[edit security alg sip application-screen protect],
[edit services alg sip application-screen protect],
[edit tenants name security alg sip application-screen protect]
```

Description

Protect servers against INVITE attacks.

Options

- `all`—Configure the Session Initiation Protocol (SIP) application screen to protect servers at all destination IP addresses against INVITE attacks.
- `destination-ip` *`address`*—Configure the SIP application screen to protect the server at this destination IP address against INVITE attacks. You can include up to 16 destination IP addresses of servers to be protected. Enabling this option disables the all option.
- `timeout` *`value-in-seconds`*—Amount of time (in *`seconds`*) to make an attack table entry for each INVITE, which is listed in the application screen.
- **Range:** 1 through 3600 seconds
- **Default:** 5 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sip (Security) | 27977

depth-threshold

IN THIS SECTION

- Syntax | 5319
- Hierarchy Level | 5319
- Description | 5320
- Options | 5320
- Required Privilege Level | 5320
- Release Information | 5320

Syntax

```
depth-threshold {
    high number;
    low number;
}
```

Hierarchy Level

```
[edit services analytics interfaces]
[edit services analytics resource-profiles]
```

Description

If network analytics queue statistics monitoring is enabled, specify the high and low values (in bytes) of the queue depth (buffer) threshold. If you configure a depth threshold, you cannot configure the latency threshold. You can configure the depth threshold for one interface or all interfaces. Specify the high and low queue depth threshold numbers:

> **NOTE**: The configuration for a specific interface supersedes the global configuration for all interfaces.

Options

high *number*

Specify the maximum value for the depth threshold.

- **Range:** 1 to 1,250,000,000 bytes
- Default:
  - Junos OS Release 13.2X51-D10 or later—0 bytes
  - Junos OS Release 13.2X50-D15—14,680,064 bytes (14 MB)

low *number*

Specify the minimum value for the depth threshold.

- **Range:** 1 to 1,250,000,000 bytes
- Default:
  - Junos OS Release 13.2X51-D10 or later—0 bytes
  - Junos OS Release 13.2X50-D15—1024 bytes (1 KB)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Statement in the `[edit services analytics resource-profiles]` hierarchy level introduced in Junos OS Release 13.2X51-D15.

RELATED DOCUMENTATION

Network Analytics Overview

latency-threshold | 9639

resource-profiles (Analytics) | 15620

description (RMON)

IN THIS SECTION

- Syntax | 5321
- Hierarchy Level | 5321
- Description | 5321
- Options | 5322
- Required Privilege Level | 5322
- Release Information | 5322

Syntax

```
description description;
```

Hierarchy Level

```
[edit snmp rmon alarm index],
[edit snmp rmon event index]
```

Description

Text description of alarm or event.

Options

description—Text description of an alarm or event entry. If the description includes spaces, enclose it in quotation marks (" ").

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Description

Configuring an RMON Event Entry and Its Attributes

description (SNMP)

IN THIS SECTION

- Syntax | 5323
- Hierarchy Level | 5323
- Description | 5323
- Options | 5323
- Required Privilege Level | 5323
- Release Information | 5323

Syntax

```
description description;
```

Hierarchy Level

```
[edit snmp]
```

Description

Define the value of the MIB II `sysDescription` object, which is the description of the system being managed.

Options

`description`—System description. If the name includes spaces, enclose it in quotation marks (" ").

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the System Description on a Device Running Junos OS

description (Abstracted Fabric)

IN THIS SECTION

- Syntax | 5324
- Hierarchy Level | 5324
- Description | 5324
- Options | 5324
- Required Privilege Level | 5324
- Release Information | 5325

Syntax

```
description description;
```

Hierarchy Level

```
[edit chassis network-slices guest-network-functions gnf af-name]
```

Description

Provide a description string for the specified abstracted fabric (af) interface.

Options

description A description string for the specified abstracted fabric (af) interface.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Setting Up Junos Node Slicing

gnf | 7746

control-plane-bandwidth-percent (Node Slicing) | 4634

fpcs (Node Slicing) | 7536

description (Application Identification)

IN THIS SECTION

- Syntax | 5325
- Hierarchy Level | 5325
- Description | 5326
- Options | 5326
- Required Privilege Level | 5326
- Release Information | 5326

Syntax

```
description description
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Description

Provide a description of the application.

Options

description Textual description of the application.

- **Range:** 1 through 255 characters

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

description (Applications)

IN THIS SECTION

- Syntax | 5327
- Hierarchy Level | 5327
- Description | 5327
- Options | 5327
- Required Privilege Level | 5327
- Release Information | 5327

Syntax

```
description text;
```

Hierarchy Level

```
[edit applications application application-name],
[edit applications application-set application-set-name]
```

Description

Specify descriptive text for an application or an application set.

> **NOTE**: The descriptive text should not include characters “<”, “>”, “&”, or “\n”.

Options

text—Descriptive text about an application or an application set.

- **Range:** 1 through 900 characters

> **NOTE**: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 900 characters, the configuration might fail to take effect.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Security Policy Application Sets Overview

description (Chassis Redundancy)

IN THIS SECTION

- Syntax | 5328
- Hierarchy Level | 5328
- Description | 5328
- Options | 5328
- Required Privilege Level | 5329
- Release Information | 5329

Syntax

```
description description;
```

Hierarchy Level

```
[edit chassis redundancy feb redundancy-group group-name]
```

Description

Provide a description of the FEB redundancy group.

Options

description—Provide a description for the FEB redundancy group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring FEB Redundancy on the M120 Router

description (Fibre Channel Fabrics)

IN THIS SECTION

- Syntax | 5329
- Hierarchy Level | 5329
- Description | 5330
- Options | 5330
- Required Privilege Level | 5330

Syntax

```
description description
```

Hierarchy Level

```
[edit fc-fabrics fabric-name]
```

Description

Text string that describes the Fibre Channel fabric. The text string has no effect on the operation of the fabric.

Options

description Text that describes the fabric. Text can include letters, numbers, and hyphens (-) and can be up to 255 characters in length. If the text includes spaces, enclose the entire text string in quotation marks.

Required Privilege Level

storage—To view this statement in the configuration.

storage-control—To add this statement to the configuration.

RELATED DOCUMENTATION

show fibre-channel fabric | 19929

description (Forwarding Options)

IN THIS SECTION

- Syntax | 5331
- Hierarchy Level | 5331
- Description | 5331
- Required Privilege Level | 5331
- Release Information | 5331

Syntax

```
description text-description;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface (interface-name | interface-group)],
[edit forwarding-options helpers domain],
[edit forwarding-options helpers domain interface interface-name],
[edit forwarding-options helpers port port-number],
[edit forwarding-options helpers port port-number interface interface-name],
[edit forwarding-options helpers tftp],
[edit forwarding-options helpers tftp interface interface-name]
```

Description

Describe a BOOTP, DHCP, Domain Name System (DNS), Trivial File Transfer Protocol (TFTP), or port-based LAN broadcast packet forwarding service, or an interface that is configured for the service.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding

Configuring Port-based LAN Broadcast Packet Forwarding

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

description (GNF)

IN THIS SECTION

- Syntax | 5332
- Hierarchy Level | 5332
- Description | 5332
- Options | 5332
- Required Privilege Level | 5332
- Release Information | 5333

Syntax

```
description description;
```

Hierarchy Level

```
[edit chassis network-slices guest-network-functions gnf]
```

Description

Provide a description string for the specified guest network function (GNF).

Options

description A description string for the specified guest network function (GNF).

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

- network-slices | 11322
- guest-network-functions | 7949
- gnf | 7746
- control-plane-bandwidth-percent (Node Slicing) | 4634
- fpcs (Node Slicing) | 7536

description (IKE policy)

IN THIS SECTION

- Syntax | 5333
- Hierarchy Level | 5333
- Description | 5334
- Required Privilege Level | 5334

Syntax

```
description description;
```

Hierarchy Level

```
[edit security ike policy ike-peer-address],
[edit security ike proposal ike-proposal-name],
[edit security ipsec policy ipsec-policy-name],
```

```
[edit security ipsec proposal ipsec-proposal-name],
[edit security ipsec security-association sa-name]
```

Description

Specify a text description for an IKE proposal or policy, or an IPsec proposal, policy, or SA.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Configuring Security Associations for IPsec on an ES PIC*
- Configuring the Description for an IKE Proposal
- *Configuring an IKE Policy for Preshared Keys*
- *Configuring an IPsec Proposal for an ES PIC*
- *Configuring the IPsec Policy for an ES PIC*

description (Interfaces)

IN THIS SECTION

- Syntax | 5335
- Hierarchy Level | 5335
- Description | 5335
- Options | 5335
- Required Privilege Level | 5335
- Release Information | 5335

Syntax

```
description text;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the `show interfaces` commands, and is also exposed in the `ifAlias` Management Information Base (MIB) object. It has no effect on the operation of the interface on the router or switch.

The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.

Options

text—Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

Configuring Interface Description

Adding a Logical Unit Description to the Configuration

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches

Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

Using DHCP Relay Agent Option 82 Information

Junos OS Network Interfaces Library for Routing Devices

Example: Connecting Access Switches with ELS Support to a Distribution Switch with ELS Support

description (Junos Fusion)

IN THIS SECTION

- Syntax | 5336
- Hierarchy Level | 5336
- Description | 5337
- Default | 5337
- Options | 5337
- Required Privilege Level | 5337
- Release Information | 5337

Syntax

```
description description;
```

Hierarchy Level

```
[edit chassis satellite-management fpc slot-id]
```

Description

Configure a description for the satellite device.

The description is optional and used for information purposes only.

Default

Satellite devices do not have descriptions, by default.

Options

description	A text description of the satellite device.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

description (Junos Fusion Satellite Device Cluster)

IN THIS SECTION

- Syntax | 5338
- Hierarchy Level | 5338

- Description | 5338
- Default | 5338
- Options | 5338
- Required Privilege Level | 5338
- Release Information | 5339

Syntax

```
description description;
```

Hierarchy Level

```
[edit chassis satellite-management cluster cluster-name fpc slot-id ]
```

Description

Configure a description for the satellite device in the satellite device cluster.

The description is optional and used for information purposes only.

Default

Satellite devices in satellite device clusters do not have descriptions, by default.

Options

description A text description of the satellite device in the satellite device cluster.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

description (Op Scripts)

IN THIS SECTION

- Syntax | 5339
- Hierarchy Level | 5339
- Description | 5340
- Options | 5340
- Required Privilege Level | 5340
- Release Information | 5340

Syntax

```
description description;
```

Hierarchy Level

```
[edit system scripts op file filename]
[edit system scripts op file filename arguments argument-name]
```

Description

Provide help text to describe an op script or its command-line arguments. The help text is included in the context-sensitive help in the *command-line interface* (*CLI*).

Options

description At the script hierarchy level, specify the purpose of the script. At the arguments hierarchy level, provide a description for the command-line argument.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Help Text for Op Scripts

Declaring and Using Command-Line Arguments in Op Scripts

file (Op Scripts) | **7067**

description (Protocols BGP)

IN THIS SECTION

- Syntax | **5341**
- Hierarchy Level | **5341**
- Description | **5341**
- Options | **5341**
- Required Privilege Level | **5341**
- Release Information | **5342**

Syntax

```
description text-description;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],

[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Provide a description of the global, group, or neighbor configuration. If the text includes one or more spaces, enclose it in quotation marks (“ “). The test is displayed in the output of the `show` command and has no effect on the configuration.

Options

text-description—Text description of the configuration. It is limited to 255 characters.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

BGP User Guide

description (Protocols Layer 2 Circuit Neighbor)

IN THIS SECTION

- Syntax | 5342
- Hierarchy Level | 5342
- Description | 5342
- Required Privilege Level | 5343
- Release Information | 5343

Syntax

```
description text;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit protocols l2circuit neighbor address interface interface-name]
```

Description

Provide a text description for the Layer 2 circuit. If the text includes one or more spaces, enclose the entire text string in quotation marks (" ").

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routing Instances on PE Routers in VPNs

description (Protocols Layer 2 VPN)

IN THIS SECTION

- Syntax | 5343
- Hierarchy Level | 5344
- Description | 5344
- Options | 5344
- Required Privilege Level | 5344
- Release Information | 5344

Syntax

```
description text;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn site site-name interface interface-name],
[edit routing-instances routing-instance-name protocols l2vpn site site-name interface interface-
name]
```

Description

Describe the VPN or virtual private LAN service (VPLS) routing instance.

Options

text—Provide a text description. If the text includes one or more spaces, enclose it in quotation marks (" "). Any descriptive text you include is displayed in the output of the `show route instance detail` command and has no effect on operation.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Site

Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)

description (Protocols MPLS)

IN THIS SECTION

- Syntax | 5345
- Hierarchy Level | 5345
- Description | 5345
- Options | 5346
- Required Privilege Level | 5346
- Release Information | 5346

Syntax

```
description text;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
bypass],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
transit incoming-label],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name bypass],
[edit protocols mpls static-label-switched-path lsp-name ingress],
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Description

Provides a textual description of the LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the `show mpls lsp detail` command and has no effect on the operation of the LSP.

Options

text—Provide a textual description of the LSP. The description text can be no more than 80 characters in length.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Text Description for LSPs

description (Routing Instances)

IN THIS SECTION

- Syntax | 5346
- Hierarchy Level | 5347
- Description | 5347
- Required Privilege Level | 5347
- Release Information | 5347

Syntax

```
description text;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Provide a text description for the routing instance. If the text includes one or more spaces, enclose it in quotation marks (" "). Any descriptive text you include is displayed in the output of the `show route instance detail` command and has no effect on the operation of the routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

description (Security Address Book)

IN THIS SECTION

- Syntax | 5348
- Hierarchy Level | 5348
- Description | 5348
- Options | 5348
- Required Privilege Level | 5348
- Release Information | 5348

Syntax

```
description text;
```

Hierarchy Level

```
[edit security address-book book-name]
```

Description

Specify descriptive text for an address book.

NOTE: The descriptive text should not include characters “<”, “>”, “&”, or “\n”.

Options

text—Descriptive text about an address book.

- **Range:** 1 through 900 characters

NOTE: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 900 characters, the configuration might fail to take effect.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Address Books

Understanding Address Sets

description (Security IDP Policy)

IN THIS SECTION

- Syntax | 5349
- Hierarchy Level | 5349
- Description | 5349
- Options | 5349
- Required Privilege Level | 5350
- Release Information | 5350

Syntax

```
description text;
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name]
[edit security idp idp-policy policy-name rulebase-ips rule rule-name]
```

Description

Specify descriptive text for an exempt rule, or IPS rule.

Options

text—Descriptive text about an exempt rule, or IPS rule.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

description (Security NAT Pool)

IN THIS SECTION

- Syntax | 5350
- Hierarchy Level | 5350
- Description | 5351
- Options | 5351
- Required Privilege Level | 5351
- Release Information | 5351

Syntax

```
description text;
```

Hierarchy Level

```
[edit security nat destination pool pool-name]
[edit security nat source pool pool-name]
```

Description

Specify descriptive text for a source or destination NAT pool.

> **NOTE**: The descriptive text should not include characters, such as “<”, “>”, “&”, or “\n”.

Options

text—Descriptive text about a source or destination NAT pool.

- **Range:** 1 through 300 characters

> **NOTE**: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 300 characters, the configuration might fail to take effect.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

description (Security NAT Rule)

IN THIS SECTION

- Syntax | 5352
- Hierarchy Level | 5352
- Description | 5352
- Options | 5352
- Required Privilege Level | 5352

Release Information | 5353

Syntax

```
description text;
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name rule rule-name]
[edit security nat source rule-set rule-set-name rule rule-name]
[edit security nat static rule-set rule-set-name rule rule-name]
```

Description

Specify descriptive text for a source, destination, or static NAT rule.

NOTE: The descriptive text should not include characters, such as “<”, “>”, “&”, or “\n”.

Options

text—Descriptive text about a source, destination, or static NAT rule.

- **Range:** 1 through 300 characters

NOTE: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 300 characters, the configuration might fail to take effect.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

description (Security NAT Rule Set)

IN THIS SECTION

- Syntax | 5353
- Hierarchy Level | 5353
- Description | 5353
- Options | 5354
- Required Privilege Level | 5354
- Release Information | 5354

Syntax

```
description text;
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name]
[edit security nat source rule-set rule-set-name]
[edit security nat static rule-set rule-set-name]
```

Description

Specify descriptive text for a source, destination, or static NAT rule set.

NOTE: The descriptive text should not include characters, such as “<”, “>”, “&”, or “\n”.

Options

text—Descriptive text about a source, destination, or static NAT rule set.

- **Range:** 1 through 300 characters

> **NOTE**: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 300 characters, the configuration might fail to take effect.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

description (Security Policies Next Gen Services)

IN THIS SECTION

- Syntax | 5355
- Hierarchy Level | 5355
- Description | 5355
- Options | 5355
- Required Privilege Level | 5355
- Release Information | 5355

Syntax

```
description description;
```

Hierarchy Level

```
[edit security ike policy policy-name],
[edit security ike proposal proposal-name],
[edit security ipsec policy policy-name],
[edit security ipsec proposal proposal-name]
```

Description

Enter descriptive text for an IKE policy, an IPsec policy, an IKE proposal, or an IPsec proposal.

Options

description	Descriptive text.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

description (Security Screen)

IN THIS SECTION

Syntax | 5356

- Hierarchy Level | 5356
- Description | 5356
- Options | 5356
- Required Privilege Level | 5356
- Release Information | 5357

Syntax

```
description text;
```

Hierarchy Level

```
[edit security screen ids-option screen-name]
```

Description

Specify descriptive text for a screen.

NOTE: The descriptive text should not include characters, such as “<”, “>”, “&”, or “\n”.

Options

text—Descriptive text about a screen.

- **Range:** 1 through 300 characters

NOTE: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 300 characters, the configuration might fail to take effect.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

description (Security Zone)

IN THIS SECTION

- Syntax | 5357
- Hierarchy Level | 5357
- Description | 5358
- Options | 5358
- Required Privilege Level | 5358
- Release Information | 5358

Syntax

```
description text;
```

Hierarchy Level

```
[edit security zones functional-zone management]
[edit security zones security-zone zone-name]
```

Description

Specify descriptive text for a security zone.

> **NOTE**: The descriptive text should not include characters as “<”, “>”, “&”, or “\n”.

Options

text—Descriptive text about a security zone.

- **Range:** 1 through 900 characters

> **NOTE**: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 900 characters, the configuration might fail to take effect.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

description (Services IPsec VPN)

IN THIS SECTION

- Syntax | 5359
- Hierarchy Level | 5359
- Description | 5359
- Required Privilege Level | 5359
- Release Information | 5359

Syntax

```
description description;
```

Hierarchy Level

```
[edit services ipsec-vpn ike policy policy-name],
[edit services ipsec-vpn ike proposal proposal-name],
[edit services ipsec-vpn ipsec policy policy-name],
[edit services ipsec-vpn ipsec proposal proposal-name]
```

Description

Specify the text description for an IKE or IPsec policy or proposal.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

description (Services IPsec VPN) | 5358

Configuring IPsec Proposals

Configuring IPsec Policies

description (VLAN)

IN THIS SECTION

- Syntax | 5360
- Hierarchy Level | 5360
- Description | 5360
- Options | 5360
- Required Privilege Level | 5360
- Release Information | 5361

Syntax

```
description text-description;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Description

Provide a textual description for the VLAN. The text has no effect on the operation of the VLAN or switch.

Options

text-description—Text to describe the interface. It can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. If the text includes spaces, enclose the entire text in quotation marks.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Option **text-description** enhanced from supporting up to 128 characters to supporting up to 256 characters in Junos OS Release 10.2 for EX Series switches.

RELATED DOCUMENTATION

Understanding Bridging and VLANs on Switches

Understanding Bridging and VLANs on Switches

Example: Setting Up Basic Bridging and a VLAN on Switches

Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch

Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch with ELS Support

show vlans | 27738

deselect-on-bandwidth-failure

IN THIS SECTION

- Syntax | 5362
- Hierarchy Level | 5362
- Description | 5362
- Options | 5362
- Required Privilege Level | 5362
- Release Information | 5362

Syntax

```
deselect-on-bandwidth-failure {
    tear-lsp;
}
```

Hierarchy Level

```
[edit protocols mpls],
[edit protocols mpls label-switched-path path-name]
```

Description

Deselect an active path if it does not meet the auto-bandwidth criteria required for path selection. The `deselect-on-bandwidth-failure` statement does not apply to static bandwidth.

Options

tear-lsp Bring down an active path if none of the paths are able to reserve the required bandwidth.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

designated-forwarder-election-hold-time

IN THIS SECTION

- Syntax | 5363
- Hierarchy Level | 5363
- Description | 5363
- Options | 5364
- Required Privilege Level | 5364
- Release Information | 5364

Syntax

```
designated-forwarder-election-hold-time seconds;
```

Hierarchy Level

```
[edit protocols evpn],
[edit routing-instances protocols evpn]
```

Description

Set the time that the device waits before electing a designated forwarder (DF).

Customer edge (CE) devices require a DF when they are multihomed to more than one provider edge (PE) device. Without a designated forwarder, multihomed hosts would receive duplicate packets. The connections between a multihomed CE device and its multihoming peer PE devices form an Ethernet segment identified by an Ethernet segment identifier (ESI). The multihomed peer devices choose DFs for an ESI based on EVPN Type 4 route advertisements. EVPN Type 4 routes are called Ethernet segment routes.

Options

seconds Time in seconds that a multihomed CE device waits before electing a DF based on the EVPN Type 4 routes it receives.

- **Range:** 1 through 1800

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D15.

RELATED DOCUMENTATION

EVPN Multihoming Overview

destination

IN THIS SECTION

- Syntax | 5365
- Hierarchy Level | 5365
- Description | 5365
- Options | 5365
- Required Privilege Level | 5365
- Release Information | 5366

Syntax

```
destination {
    address address;
}
```

Hierarchy Level

```
[edit interfaces fti number unit logical-unit-number tunnel encapsulation ( gre | ipip | udp )]
```

Description

Interface address prefix of the destination where the tunnel terminates. Configuring this interface is mandatory.

> **NOTE**: Starting in Junos OS Evolved Release 20.3R1, for GRE tunnels, configure the address for the tunnel encapsulation destination address.
>
> When you do not configure "tunnel-termination" on page 30245, you can use `destination address address` to configure the destination IP address in the outer IP header while encapsulating the packet.

Options

address You can configure the destination address with an optional destination mask. The mask range is from 0 to 32 for IPv4, and 0 to 128 for IPv6 addresses. /32 is the default mask for IPv4, and /128 is the default mask for IPv6.

> **NOTE**: Starting in Junos OS Evolved Release 20.3R1, for GRE tunnels, if you do not configure "tunnel-termination" on page 30245 for tunnel decapsulation, you cannot configure the mask for the destination address.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Evolved Release 20.2R1.

Statement introduced for IPIP tunnels in Junos OS Evolved Release 20.4R1.

Statement introduced for UDP tunnels in Junos OS Evolved Release 21.1R1.

destination (Accounting)

IN THIS SECTION

- Syntax | 5366
- Hierarchy Level | 5367
- Description | 5367
- Options | 5367
- Required Privilege Level | 5367
- Release Information | 5367

Syntax

```
destination {
    radius {
        server {
            server-address {
                accounting-port port-number;
                accounting-retry number;
                accounting-timeout seconds;
                dynamic-request-port number;
                max-outstanding-requests value;
                port number;
                preauthentication-port number;
                preauthentication-secret secret;
                retry number;
                routing-instance routing-instance-name;
                secret password;
```

```
                source-address source-address;
                timeout seconds;
            }
    }
    tacplus {
        server {
            server-address {
                port port-number;
                routing-instance routing-instance;
                secret password;
                single-connection;
                source-address address
                timeout seconds;
            }
        }
    }
}
```

Hierarchy Level

```
[edit system accounting]
```

Description

Configure the authentication server.

Options

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`radius` statement added in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring RADIUS System Accounting

Configuring TACACS+ System Accounting

destination-address

IN THIS SECTION

- Syntax | 5368
- Hierarchy Level | 5368
- Description | 5368
- Options | 5369
- Required Privilege Level | 5369
- Release Information | 5369

Syntax

```
destination-address address;
```

Hierarchy Level

```
[edit services aacl rule rule-name term term-name from]
```

Description

Specify the destination address for rule matching.

Options

address—Destination IPv4 or IPv6 address or prefix value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

IPv6 support introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Match Conditions in AACL Rules

destination-address (IPsec-VPN)

IN THIS SECTION

- Syntax | 5370
- Hierarchy Level | 5370
- Description | 5370
- Options | 5370
- Required Privilege Level | 5370
- Release Information | 5370

Syntax

```
destination-address address;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name from]
```

Description

Specify the destination address for rule matching.

Options

address—Destination IP address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

destination-address (Firewall)

IN THIS SECTION

- Syntax | 5371
- Hierarchy Level | 5371
- Description | 5371
- Required Privilege Level | 5371
- Release Information | 5372

Syntax

```
destination-address address;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name from],
[edit firewall family family-name filter filter-name term term-name from ip-version ip-version]
```

Description

IP destination address field, which is the address of the destination node for the packet. You cannot specify the `destination-address` and `address` match conditions in the same term.

For IPv4, the `destination-address` field is 32 bits in length. The filter description syntax supports either a mask value that can be noncontiguous, such as 10.0.0.10/255.0.0.255, or prefix notation such as 10.0.0.0/8. Simple filters do not support noncontiguous mask values.

For IPv6, the `destination-address` field is 128 bits in length. The filter description syntax supports the text representations for IPv6 addresses that are described in RFC 2373, *IP Version 6 Addressing Architecture*.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1R1.

RELATED DOCUMENTATION

family (Firewall) | 6867

family (Firewall Filter) | 6869

Firewall Filter Match Conditions Based on Address Fields

Firewall Filter Match Conditions for IPv4 Traffic

Firewall Filter Match Conditions for IPv6 Traffic

Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic

Guidelines for Configuring Simple Filters

destination-address-excluded

IN THIS SECTION

- Syntax | 5372
- Hierarchy Level | 5373
- Description | 5373
- Required Privilege Level | 5373
- Release Information | 5373

Syntax

```
destination-address-excluded;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy  policy-name match]
```

Description

Exclude the destination address(es) from the policy.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Security Policies Overview

destination-address (Flow Monitoring Logs for NAT)

IN THIS SECTION

- Syntax | 5374
- Hierarchy Level | 5374
- Description | 5374
- Options | 5374
- Required Privilege Level | 5374
- Release Information | 5374

Syntax

```
destination-address address;
```

Hierarchy Level

```
[edit services jflow-log collector collector-name]
```

Description

Specify the destination IP address or identifier of the host or external device that functions as the collector for receiving the generated flow monitoring logs that are sent from the exporter. You can configure an IPv4 address, or an identifier of the host system (the workstation either running the Jflow utility or collecting traffic flows using version 9 or IPFIX format). For external NetFlow collectors or servers, the hostname must be reachable from the same routing instance to which the initial data packet (that triggered session establishment) is delivered. You can specify a maximum of eight collectors per profile.

Options

address—Destination hostname, or IPv4 or IPv6 address of the collector.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

destination-address (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 5375
- Hierarchy Level | 5375
- Description | 5375
- Options | 5376
- Required Privilege Level | 5376
- Release Information | 5376

Syntax

```
destination-address {
    (any-ipv4 | any-ipv4 except);
    (any-ipv6 | any-ipv6 except);
    (any-unicast | any-unicast except);
    (prefix | prefix except);
}
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number from]
```

Description

Specify the prefix or address type that the HTTP request destination IP address must match. You can specify multiple prefixes or address types by including the `destination-address` statement multiple times.

After this criterion and the other match criteria specified for the `term` are matched, the HTTP header enrichment actions specified for the `term` are applied to the HTTP traffic.

Options

any-ipv4 Match any IPv4 address.

any-ipv4 except Exclude IPv4 addresses from addresses that are in a `destination-address`, `destination-address-range`, or `destination-prefix-list` statement configured at the `[edit services hcm tag-rule` *`rule-name`* `term` *`term-number`* `from]` hierarchy level. You cannot use `except` without also configuring addresses that do match.

any-ipv6 Match any IPv6 address.

any-ipv6 except Exclude IPv6 addresses from addresses that are in a `destination-address`, `destination-address-range`, or `destination-prefix-list` statement configured at the `[edit services hcm tag-rule` *`rule-name`* `term` *`term-number`* `from]` hierarchy level. You cannot use `except` without also configuring addresses that do match.

any-unicast Match any IPv4 unicast address. This option does not match any IPv6 addresses.

any-unicast except Exclude IPv4 unicast addresses from addresses that are in a `destination-address`, `destination-address-range`, or `destination-prefix-list` statement configured at the `[edit services hcm tag-rule` *`rule-name`* `term` *`term-number`* `from]` hierarchy level. You cannot use `except` without also configuring IPv4 addresses that do match.

prefix IP prefix for the addresses that are matched.

***prefix* except** Exclude the specified IP prefixes from addresses that are in a `destination-address`, `destination-address-range`, or `destination-prefix-list` statement configured at the `[edit services hcm tag-rule` *`rule-name`* `term` *`term-number`* `from]` hierarchy level. You cannot use `except` without also configuring addresses that do match.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm (HTTP Header Enrichment) | 8005

destination-address (IDS MS-DPC)

IN THIS SECTION

- Syntax | 5377
- Hierarchy Level | 5377
- Description | 5377
- Options | 5377
- Required Privilege Level | 5378
- Release Information | 5378

Syntax

```
destination-address (address | any-unicast) <except>;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name from]
```

Description

Specify the destination address for rule matching when using the MS-DPC.

Options

address—Destination IPv4 or IPv6 address or prefix value.

`any-unicast`—Any unicast packet.

`except`—(Optional) Exempt the specified address, prefix, or unicast packets from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Match Conditions in IDS Rules

destination-address-name (NAT Next Gen Services)

IN THIS SECTION

- Syntax | 5378
- Hierarchy Level | 5379
- Description | 5379
- Required Privilege Level | 5379
- Release Information | 5379

Syntax

```
destination-address-name address-name;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name match],
[edit services nat source rule-set rule-set rule rule-name match]
```

Description

Specify the name of the range of destination addresses that the packet must match for the NAT rule to take effect. The range of addresses is configured with the `address` statement at the `[edit services address-book global]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

destination-address-name (Security Destination NAT)

IN THIS SECTION

- Syntax | 5380
- Hierarchy Level | 5380
- Description | 5380
- Options | 5380
- Required Privilege Level | 5380
- Release Information | 5380

Syntax

```
destination-address-name <address-name>;
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name rule rule-name match]
```

Description

Specify a destination address name to match the rule. You can configure multiple address names.

Options

address-name—Destination address name.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

destination-address-name (Security Source NAT)

IN THIS SECTION

- Syntax | 5381
- Hierarchy Level | 5381
- Description | 5381
- Options | 5381

Required Privilege Level | 5381
Release Information | 5381

Syntax

```
destination-address-name <address-name>;
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name match]
```

Description

Specify a destination address name to match the rule. You can configure multiple address names.

Options

address-name—Destination address name.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

destination-address-name (Security Static NAT)

IN THIS SECTION

- Syntax | 5382
- Hierarchy Level | 5382
- Description | 5382
- Options | 5382
- Required Privilege Level | 5382
- Release Information | 5383

Syntax

```
destination-address-name <address-name>;
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name match]
```

Description

Specify a destination address name to match the rule.

Options

destination-address-name—Name of the destination address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

destination-address (NAT Next Gen Services)

IN THIS SECTION

- Syntax | 5383
- Hierarchy Level | 5383
- Description | 5383
- Options | 5383
- Required Privilege Level | 5384
- Release Information | 5384

Syntax

```
destination-address (address | any | any-ipv4 | any-ipv6);
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name match],
[edit services nat source rule-set rule-set rule rule-name match]
```

Description

Specify the destination address that the packet must match for the NAT rule to take effect.

Options

address A specific address that must be matched.

any	Any unicast destination address results in a match.
any-ipv4	Any IPv4 destination address results in a match.
any-ipv6	Any IPv6 destination address results in a match.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

destination-address (PCP)

IN THIS SECTION

- Syntax | 5384
- Hierarchy Level | 5385
- Description | 5385
- Options | 5385
- Required Privilege Level | 5385
- Release Information | 5385

Syntax

```
destination-address address <except>;
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name from]
```

Description

Specify the destination address that must be matched for the PCP rule. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

address	Destination address or prefix value.
minimum-value	Lower boundary for the address range.
except	(Optional) Prevent the specified address range from matching the PCP rule.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

destination-address-range

IN THIS SECTION

- Syntax | 5386
- Hierarchy Level | 5386
- Description | 5386
- Options | 5386
- Required Privilege Level | 5386
- Release Information | 5387

Syntax

```
destination-address-range low minimum-value high maximum-value;
```

Hierarchy Level

```
[edit services aacl rule rule-name term term-name from]
```

Description

Specify the destination address range for rule matching.

Options

minimum-value—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

IPv6 support introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Match Conditions in AACL Rules

destination-address-range (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 5387
- Hierarchy Level | 5387
- Description | 5388
- Options | 5388
- Required Privilege Level | 5388
- Release Information | 5388

Syntax

```
destination-address-range {
    high address low address <except>;
}
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number from]
```

Description

Specify an IP address range that the HTTP request destination IP address must match. You can specify multiple address ranges by including the `destination-address-range` statement multiple times.

After this criterion and the other match criteria specified for the `term` are matched, the HTTP header enrichment actions specified for the `term` are applied to the HTTP traffic.

Options

except (Optional) Exclude addresses in the specified address range from addresses that are in a `destination-address`, `destination-address-range`, or `destination-prefix-list` statement configured at the [`edit services hcm tag-rule` *`rule-name`* `term` *`term-number`* `from`] hierarchy level. You cannot use `except` without also configuring addresses that do match.

high *address* Upper limit of the address range.

low *address* Lower limit of the address range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm (HTTP Header Enrichment) | **8005**

destination-address-range (IDS MS-DPC))

IN THIS SECTION

- Syntax | 5389
- Hierarchy Level | 5389
- Description | 5389
- Options | 5389
- Required Privilege Level | 5389
- Release Information | 5390

Syntax

```
destination-address-range low minimum-value high maximum-value <except>;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name from]
```

Description

Specify the destination address range for rule matching when using the MS-DPC.

Options

minimum-value—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

`except`—(Optional) Exempt the specified address range from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

minimum-value and *maximum-value* options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Match Conditions in IDS Rules

destination-address-range (PCP)

IN THIS SECTION

- Syntax | 5390
- Hierarchy Level | 5390
- Description | 5391
- Options | 5391
- Required Privilege Level | 5391
- Release Information | 5391

Syntax

```
destination-address-range high maximum-value low minimum-value <except>;
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name from]
```

Description

Specify the destination address range that must be matched for the PCP rule. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

maximum-value	Upper boundary for the address range.
minimum-value	Lower boundary for the address range.
except	(Optional) Prevent the specified address range from matching the PCP rule.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

destination-address-range (Services NAT)

IN THIS SECTION

- Syntax | 5392
- Hierarchy Level | 5392

- Description | 5392
- Options | 5392
- Required Privilege Level | 5392
- Release Information | 5393

Syntax

```
destination-address-range low minimum-value high maximum-value <except>;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name from]
```

Description

Specify the destination address range for rule matching.

If the `translation-type` statement in the `then` statement of the nat rule is set to **stateful-nat-64**, the destination address range for rule matching must be within the range specified by the `destination-prefix` statement in the `then` statement.

Options

minimum-value—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

except—(Optional) Prevent the specified address range from being translated.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

minimum-value and ***maximum-value*** options enhanced to support IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Network Address Translation Rules Overview*

destination-address-range (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 5393
- Hierarchy Level | 5393
- Description | 5394
- Options | 5394
- Required Privilege Level | 5394
- Release Information | 5394

Syntax

```
destination-address-range low minimum-value high maximum-value <except>;
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name term term-name from]
```

Description

Specify the destination address range for rule matching.

Options

minimum-value—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

except—(Optional) Exclude the specified address range from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

minimum-value and *maximum-value* options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Match Conditions in Stateful Firewall Rules

destination-address (Security Destination NAT)

IN THIS SECTION

- Syntax | 5395
- Hierarchy Level | 5395
- Description | 5395
- Options | 5395

Required Privilege Level | 5395
Release Information | 5396

Syntax

```
destination-address <ip-address>;
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name rule rule-name match]
```

Description

Specify a destination address to match the rule. You can configure one address or a subnet.

NOTE:

- If the destination address is IPv4 and the pool is an IPv6 prefix, the length of the IPv6 prefix must be 96.
- If the destination address is an IPv6 prefix and the pool is an IPv6 prefix, their length must be the same.

Options

ip-address— Destination address or a subnet.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

destination-address (Security IDP Policy)

IN THIS SECTION

- Syntax | 5396
- Hierarchy Level | 5396
- Description | 5396
- Options | 5396
- Required Privilege Level | 5397
- Release Information | 5397

Syntax

```
destination-address ([address-name] | any | any-ipv4 | any-ipv6);
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match]
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match]
```

Description

Specify a destination IP address or IP address set object to be used as the match destination address object. The default value is any.

Options

- *address-name*—IP address or IP address set object.

- *any*—Specify any IPv4 or IPv6 address.
- *any-ipv4*—Specify any IPv4 address.
- *any-ipv6*—Specify any IPv6 address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

destination-address (Security Policies)

IN THIS SECTION

- Syntax | 5397
- Hierarchy Level | 5398
- Description | 5398
- Options | 5398
- Required Privilege Level | 5398
- Release Information | 5398

Syntax

```
destination-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name match]
```

```
[edit security policies global policy policy-name match]
```

Description

Define the matching criteria. You can specify one or more IP addresses, address sets, or wildcard addresses. You can specify wildcards `any`, `any-ipv4`, or `any-ipv6`.

Options

address—IP address (`any`, `any-ipv4`, `any-ipv6`), IP address set, or address book entry, or wildcard address (represented as A.B.C.D/wildcard-mask). You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support.

RELATED DOCUMENTATION

Security Policies Overview

destination-address (Security Policies Flag)

IN THIS SECTION

- Syntax | 5399
- Hierarchy Level | 5399
- Description | 5399
- Options | 5400
- Required Privilege Level | 5400
- Release Information | 5400

Syntax

```
destination-address {
    drop-translated;
    drop-untranslated;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit]
```

Description

Specify whether the traffic permitted by the security policy is limited to packets where the destination IP address has been translated by means of a destination NAT rule or to packets where the destination IP address has not been translated.

On Juniper Networks security devices, destination NAT rules are processed before security policy lookup. Therefore, it is possible for a security policy to permit traffic from a source `S` to a destination `D` (where no destination NAT is performed) and also to permit traffic from the source `S` to the destination `d` (where `d` has been translated to `D`).

Options

- `drop-translated`—Drop packets with translated destination IP addresses. Traffic permitted by the security policy is limited to packets where the destination IP address has not been translated.
- `drop-untranslated`—Drop packets without translated destination IP addresses. Traffic permitted by the security policy is limited to packets where the destination IP address has been translated by means of a destination NAT rule.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Security Policies Overview

destination-address (Security Source NAT)

IN THIS SECTION

- Syntax | 5401
- Hierarchy Level | 5401
- Description | 5401
- Options | 5401
- Required Privilege Level | 5401
- Release Information | 5401

Syntax

```
destination-address <ip-address>;
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name match]
```

Description

Specify a destination address to match the rule. You can configure multiple addresses or subnets.

Options

ip-address—Destination address or a subnet.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

destination-address (Security Static NAT)

IN THIS SECTION

- Syntax | 5402
- Hierarchy Level | 5402
- Description | 5402
- Options | 5402

Required Privilege Level | 5402
Release Information | 5402

Syntax

```
destination-address <ip-address>;
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name match]
```

Description

Specify a destination address to match the rule. You can configure one address or a subnet.

Options

ip-address—Destination address or a subnet.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

destination-address (Services CoS)

IN THIS SECTION

- Syntax | 5403
- Hierarchy Level | 5403
- Description | 5403
- Options | 5403
- Required Privilege Level | 5403
- Release Information | 5404

Syntax

```
destination-address (address | any-unicast) <except>;
```

Hierarchy Level

```
[edit services cos rule rule-name term term-name from]
```

Description

Specify the destination address for rule matching.

Options

address—Destination IPv4 or IPv6 address or prefix value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

Configuring Match Conditions In CoS Rules

destination-address (Services NAT)

IN THIS SECTION

- Syntax | 5404
- Hierarchy Level | 5404
- Description | 5405
- Options | 5405
- Required Privilege Level | 5405
- Release Information | 5405

Syntax

```
destination-address (address | any-unicast) <except>;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name from]
```

Description

Specify the destination address for rule matching.

Options

address—Destination IPv4 or IPv6 address or prefix value.

any-unicast—Any unicast packet.

except—(Optional) Prevent the specified address, prefix, or unicast packets from being translated.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

any-unicast and **except** options introduced in Junos OS Release 7.6.

address option enhanced to support IPv6 and addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

destination-address (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 5406
- Hierarchy Level | 5406
- Description | 5406
- Options | 5406

Required Privilege Level | 5406
Release Information | 5407

Syntax

```
destination-address (address | any-ipv4 | any-ipv6 | any-unicast) <except>;
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name term term-name from]
```

Description

Specify the destination address for rule matching.

Options

address—Destination IPv4 or IPv6 address or prefix value. Using a value of 0::0/0 with IPv6 is not allowed for M-Series and MX-Series routers.

`any-ipv4`—Any IPv4 packet.

`any-ipv6`—Any IPv6 packet.

`any-unicast`—Match all unicast packets.

`except`—(Optional) Exclude the specified address, prefix, IPv4, IPv6, or unicast packets from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`any-unicast` and `except` options introduced in Junos OS Release 7.6.

address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Match Conditions in Stateful Firewall Rules

destination-address (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 5407
- Hierarchy Level | 5407
- Description | 5408
- Options | 5408
- Required Privilege Level | 5408
- Release Information | 5408

Syntax

```
destination-address address;
```

Hierarchy Level

```
[edit services radius-flow-tap policy policy-name inet drop-policy rule-name from],
[edit services radius-flow-tap policy policy-name inet6 drop-policy rule-name from]
```

Description

Specify destination IP address or prefix value for radius-flow-tap policy rule mapping.

Options

address— IPv4 or IPv6 address for the radius-flow-tap policy.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Subscriber Secure Policy Overview

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

destination (Application Identification)

IN THIS SECTION

- Syntax | 5409
- Hierarchy Level | 5409
- Description | 5409
- Options | 5409
- Required Privilege Level | 5409
- Release Information | 5409

Syntax

```
destination ip ip-address-prefix;
```

Hierarchy Level

```
[edit services application-identification application application-name address-mapping]
```

Description

Specify the destination IP address for address mapping-based application identification.

Options

ip-address-prefix	IP address and prefix for matching.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> *Application Identification Overview*
>
> *Configuring Custom Application Signatures*

destination (Chassis)

IN THIS SECTION

- Syntax | 5410
- Hierarchy Level | 5410
- Description | 5410
- Options | 5410
- Required Privilege Level | 5411
- Release Information | 5411

Syntax

```
destination destination;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-
provider syslog facility]
```

Description

Configure where log messages go. By default, all messages go to the **/var/log** directory on the Routing Engine. Enhancements to the existing infrastructure make debugging on the Multiservices PIC easier by giving the user the option of redirecting log messages. When the `syslog destination` statement is configured to redirect the log messages, you can use the `set system syslog` command, a command available in the native Junos OS CLI, to override the syslog settings made on the Multiservices PIC.

Options

destination—Choose one of the following options:

- `routing-engine`—Forward log messages to the Routing Engine.

- `pic-console`—Forward log messages to the console of the PIC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

extension-provider | 6716

destination-class-usage

IN THIS SECTION

- Syntax | 5411
- Hierarchy Level | 5412
- Description | 5412
- Required Privilege Level | 5412
- Release Information | 5412

Syntax

```
destination-class-usage;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet accounting],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet accounting]
```

Description

Enable packet counters on an interface that count packets that arrive from specific customers and are destined for specific prefixes on the provider core router.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Source Class and Destination Class Usage

accounting (interfaces) | **1709**

source-class-usage | **28191**

destination-classes

IN THIS SECTION

- Syntax | **5413**
- Hierarchy Level | **5413**
- Description | **5413**

- Options | 5413
- Required Privilege Level | 5413
- Release Information | 5413

Syntax

```
destination-classes {
    destination-class-name;
}
```

Hierarchy Level

```
[edit accounting-options class-usage-profile profile-name]
```

Description

Specify the destination classes for which statistics are collected.

Options

destination-class-name—Name of the destination class to include in the source class usage profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Class Usage Profile

destination (Diameter Network Element)

IN THIS SECTION

- Syntax | 5414
- Hierarchy Level | 5414
- Description | 5414
- Options | 5414
- Required Privilege Level | 5415
- Release Information | 5415

Syntax

```
destination realm realm-name <host hostname>;
```

Hierarchy Level

```
[edit diameter network-element element-name forwarding route dne-route-name]
```

Description

Associate the route with all hosts of the specified realm or with a specific host of the specified realm. Together with the function and metric, defines a route reachable through a Diameter network element.

Options

`host` *`hostname`*—(Optional) Name of the destination host associated with the route.

`realm` *`realm-name`*—Name of the destination realm associated with the route.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Network Elements

destination-equal-load-balancing (L2TP LAC)

IN THIS SECTION

- Syntax | 5415
- Hierarchy Level | 5415
- Description | 5416
- Required Privilege Level | 5416
- Release Information | 5416

Syntax

```
destination-equal-load-balancing;
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Enable the LAC to balance the L2TP session load equally across multiple LNSs by selecting tunnels according to how many sessions currently exist for the destination and tunnel.

Disabled by default. By default, tunnel selection within a preference level is strictly random. The **weighted-load-balancing** statement must be disabled to successfully enable this statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

L2TP LAC Tunneling for Subscribers

L2TP LAC Tunneling for Subscribers

destination (Event Policy)

IN THIS SECTION

- Syntax | 5417
- Hierarchy Level | 5417
- Description | 5417
- Options | 5417
- Required Privilege Level | 5417
- Release Information | 5417

Syntax

```
destination destination-name {
    retry-count count retry-interval seconds;
    transfer-delay seconds;
}
```

Hierarchy Level

```
[edit event-options policy policy-name then event-script filename],
[edit event-options policy policy-name then execute-commands]
```

Description

Assign a location to which to upload command or script output for the specified event policy.

Options

destination-name Name of a destination defined in the `destinations` statement at the `[edit event-options]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

Support at the `[edit event-options policy policy-name then event-script filename]` hierarchy level introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configure an Event Policy to Execute Operational Mode Commands

Execute Event Scripts in an Event Policy

destinations (Event Policy) | **5533**

destination-except

IN THIS SECTION

- Syntax | **5418**
- Hierarchy Level | **5418**
- Description | **5418**
- Options | **5418**
- Required Privilege Level | **5419**
- Release Information | **5419**

Syntax

```
destination-except [address-name];
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match]
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match]
```

Description

Specify a destination IP address or IP address set object to specify all destination address objects except the specified address objects. The default value is any.

Options

address-name—IP address or IP address set object.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

destination-header

IN THIS SECTION

- Syntax | 5419
- Hierarchy Level | 5420
- Description | 5420
- Options | 5420
- Required Privilege Level | 5420
- Release Information | 5420

Syntax

```
destination-header {
    ILNP-nonce-option;
    home-address-option;
    line-identification-option;
    tunnel-encapsulation-limit-option;
    user-defined-option-type low | <to high>;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name ip ipv6-extension-header]
```

Description

Define the IPv6 destination header screen option.

Options

ILNP-nonce-option	Enable the Identifier-Locator Network Protocol nonce screen option.
home-address-option	Enable the home address screen option.
line-identification-option	Enable the line identification screen option.
tunnel-encapsulation-limit-option	Enable the tunnel encapsulation limit screen option.
user-defined-header-type *low* **| <to** *high***>**	Define the type of header range. • **Range:** 1 through 255.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Understanding IPv6 Support for Screens

ipv6-extension-header | **9350**

hop-by-hop-header | **8132**

destination-host

IN THIS SECTION

- Syntax | 5421
- Hierarchy Level | 5421
- Description | 5421
- Options | 5421
- Required Privilege Level | 5421
- Release Information | 5422

Syntax

```
destination-host hostname
```

Hierarchy Level

```
[edit jsrc partition partition-name]
```

Description

Configure the host on which the SAE application resides.

Options

hostname—Host on which the SAE is installed.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

JSRC Configuration Overview

Configuring the JSRC Partition

destination-host (Gx-Plus)

IN THIS SECTION

- Syntax | 5422
- Hierarchy Level | 5422
- Description | 5422
- Options | 5423
- Required Privilege Level | 5423
- Release Information | 5423

Syntax

```
destination-host hostname;
```

Hierarchy Level

```
[edit access gx-plus partition partition-name]
```

Description

Configure the host on which the PCRF application resides.

Options

hostname—Host on which the PCRF is installed.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Gx-Plus

Configuring the Gx-Plus Partition

destination-host (OCS Partition)

IN THIS SECTION

- Syntax | 5423
- Hierarchy Level | 5424
- Description | 5424
- Options | 5424
- Required Privilege Level | 5424
- Release Information | 5424

Syntax

```
destination-host ocs-hostname;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Configure the Destination-Host AVP value used in the CCR-GY-I message.

Options

ocs-hostname—Value of the Destination-Host AVP to be used in the CCR-GY-I message.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- *Configuring the OCS Partition*
- *Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers*
- *Understanding Gy Interactions Between the Router and the OCS*
- *Understanding Interactions Between the PCRF, PCEF, and OCS*

destination-host (PCRF Partition)

IN THIS SECTION

Syntax | 5425

- Hierarchy Level | 5425
- Description | 5425
- Options | 5425
- Required Privilege Level | 5425
- Release Information | 5425

Syntax

```
destination-host pcrf-hostname;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the Destination-Host AVP value used in the CCR-GX-I message.

Options

pcrf-hostname—(Optional) Value of the Destination-Host AVP to be used in the CCR-GX-I message.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

Understanding Upstream and Downstream Messages for the PCRF

destination-hostname

IN THIS SECTION

- Syntax | 5426
- Hierarchy Level | 5426
- Description | 5427
- Required Privilege Level | 5427
- Release Information | 5427

Syntax

```
destination-hostname destination-hostname;
```

Hierarchy Level

```
[edit interfaces name unit name tunnel]
```

Description

Define the hostname of the tunnel destination for the IP tunnel interface.

After the DNS resolves the IP address for the hostname, the resolved address is used as tunnel destination address. The ipip interface (IFL) is not created unless the hostname is resolved.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 for SRX Series Firewalls (SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, SRX5800, and vSRX3.0)

Statement introduced in Junos OS Release 23.4R1 for SRX Series Firewalls (SRX38, SRX300, SRX320, SRX340, and SRX345)

RELATED DOCUMENTATION

Tunnel Services Overview

tunnel | 30190

destination-identity-context

IN THIS SECTION

- Syntax | 5428
- Hierarchy Level | 5428
- Description | 5428
- Options | 5428
- Required Privilege Level | 5429
- Release Information | 5429

Syntax

```
destination-identity-context {
      [user-or-role-name];
      any;
      identity-context-authenticated;
      identity-context-unauthenticated;
      identity-context-unavailable;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name]
```

Description

Identifies users and roles to be used as match criteria for a policy. If a value other than `any` is specified as match criteria for a policy within a zone pair, the traffic is matched to table entries to retrieve associated user and roles before policy lookup occurs. Users and roles are retrieved from the local authentication table or from a UIT pushed to the SRX Series firewall from an access control service when a user is authenticated.

We recommend that you don't configure captive portal, when you use `destination-identity-context` with `identity-context-unauthenticated` option.

Options

The following entries specify the destination identities that match a policy:

user-or-role-name A list of specific users and roles.

- Range: 0 through 39 characters.

  SRX Series firewall truncate imported roles to 39 characters. You need to ensure that all of your roles are 39 characters or less.

any Any user or role, as well as the keywords identity-context-authenticated, identity-context-unauthenticated, and identity-context-unavailable.

identity-context-authenticated All users and roles that have been authenticated.

identity-context-unauthenticated Any user or role that does not have an IP-address mapped to authentication destination and the authentication destination is up and running.

identity-context-unavailable Any user or role that does not have an IP address mapped to authentication destination, because the authentication destination is disconnected from the SRX Series firewall. In this case, users are unable to be authenticated due to an authentication server disconnection, such as a power outage.

Unavailable user must be configured for non-domain users to be able to authenticate and log in.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

destination-identity-context-profile

IN THIS SECTION

- Syntax | 5430
- Hierarchy Level | 5430
- Description | 5430
- Options | 5430
- Required Privilege Level | 5430
- Release Information | 5430

Syntax

```
destination-identity-context-profile device-identity-profile-name;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name]
```

Description

The `destination-identity-context-profile` field in a security policy enables you to specify a device identity profile. The device identity feature allows you to use the identity of a device and its attributes to control access to network resources instead of the identity of the user of that device.

Options

device-identity-profile-name	Device identity profile that specifies characteristics that can apply to one or more devices.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

destination-interface

IN THIS SECTION

Syntax | 5431

- Hierarchy Level | 5431
- Description | 5431
- Options | 5432
- Required Privilege Level | 5432
- Release Information | 5432

Syntax

```
destination-interface interface-name;
```

Hierarchy Level

```
[edit services rpm probe owner test test-name],
[edit services rpm probe-server (tcp | udp)],
[edit services rpm twamp client control-connection control-client-name]
```

Description

On M Series and T Series routers, specify a services (sp-) interface that adds a timestamp to RPM probe messages. This feature is supported only with icmp-ping, icmp-ping-timestamp, udp-ping, and udp-ping-timestamp probe types. You must also configure the rpm statement on the sp- interface and include the unit 0 family inet statement with a /32 address.

On M Series, MX Series, and T Series routers, specify a multiservices (ms-) interface that adds a timestamp to RPM probe messages. This feature is supported only with icmp-ping, icmp-ping-timestamp, udp-ping, and udp-ping-timestamp probe types. You must also configure the rpm statement on the ms- interface and include the unit 0 family inet statement with a /32 address.

The inline service interface (si- interface) is a virtual physical service interface that resides on the Packet Forwarding Engine to provide L2TP services without a special services PIC. The inline service interface is supported only by MPCs on MX Series routers. Four inline service interfaces are configurable per MPC-occupied chassis slot. Specify a multiservices (si-) interface that adds a timestamp to TWAMP probe messages. You must also configure the rpm twamp-client or twamp-server statement on the si- interface and include the unit 0 family inet statement with a /32 address.

To enable RPM for the extension-provider packages on the adaptive services interface, configure the `object-cache-size`, `policy-db-size`, and `package` statements at the [`edit chassis fpc` *`slot-number`* `pic` *`pic-number`* `adaptive-services service-package extension-provider`] hierarchy level. For the extension-provider package, *`package-name`* in the `package` *`package-name`* statement is `jservices-rpm`.

Starting in Junos OS Release 17.3R1, you can use `destination-interface` *`interface-name.logical-unit-number`* at the [`edit services rpm probe` *`owner`* `test` *`test-name`*] hierarchy level to configure the generation of probes on an MS-MPC or MS-MIC. You must also inlcude the `delegate-probes` statement at the [`edit services rpm probe` *`owner`*] hierarchy level and the `rpm client-delegate-probes` and the `family (inet | inet6) address` *`address`* statements at the [`edit interfaces` *`interface-name`* `unit` *`logical-unit-number`*] hierarchy level.

Options

`interface-name`—Name of the adaptive services interface.

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

Support at the [`edit services rpm twamp client control-connection` *`control-client-name`*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

RELATED DOCUMENTATION

Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches

Configuring RPM Receiver Servers

Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches

hardware-timestamp | **7975**

rpm (Interfaces) | **16019**

Enabling RPM on MX, M and T Series Routers and SRX Firewalls for the Services SDK

destination (Interfaces)

IN THIS SECTION

- Syntax | 5433
- Hierarchy Level | 5433
- Description | 5433
- Options | 5433
- Required Privilege Level | 5434
- Release Information | 5434

Syntax

```
destination address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number tunnel],
[edit interfaces interface-name unit logical-unit-number family inet address address],
[edit interfaces interface-name unit logical-unit-number tunnel],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet address address]
```

Description

For point-to-point interfaces only, specify the address of the interface at the remote end of the connection.

For tunnel and encryption interfaces, specify the remote address of the tunnel.

Options

address—Address of the remote side of the connection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Link and Multilink Services Logical Interfaces

Configuring Encryption Interfaces

Configuring Traffic Sampling on MX, M and T Series Routers

Configuring Flow Monitoring

Configuring Unicast Tunnels

destination-ip-address (RADIUS Snoop Segment)

IN THIS SECTION

- Syntax | 5434
- Hierarchy Level | 5435
- Description | 5435
- Options | 5435
- Required Privilege Level | 5435
- Release Information | 5435

Syntax

```
destination-ip-address destination-address;
```

Hierarchy Level

```
[edit access radius snoop-segments segment-name]
```

Description

Specify the destination IP address for accounting messages to snoop.

Options

destination-address	Destination IPv4 address of accounting messages.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview

IP-Based Subscriber Setup Overview

destination-ip-based

IN THIS SECTION

- Syntax | 5436
- Hierarchy Level | 5436

- Description | 5436
- Options | 5436
- Required Privilege Level | 5436
- Release Information | 5436

Syntax

```
destination-ip-based number;
```

Hierarchy Level

```
[edit security screen ids-option screen-name limit-session]
```

Description

Limit the number of concurrent sessions the device can direct to a single destination IP address.

Options

number —Maximum number of concurrent sessions that can be directed to a destination IP address.

- **Range:** 1 through 1,000,000

For SRX Series Firewalls, the applicable range is 1 through 2,000,000.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

destination-ip (Security Alarms)

IN THIS SECTION

- Syntax | 5437
- Hierarchy Level | 5437
- Description | 5437
- Options | 5438
- Required Privilege Level | 5438
- Release Information | 5438

Syntax

```
destination-ip {
    duration interval;
    size count;
    threshold value;
}
```

Hierarchy Level

```
[edit security alarms potential-violation policy]
```

Description

Configure alarms for a number of policy violations to a destination network identifier within a specified time period.

Options

- `duration` *`interval`*—Indicate the duration of counters.
- **Range:** 1 through 3600 seconds.
- **Default:** 1 second.
- `size` *`count`*—Indicate the number of destination IP addresses for which policy violation checks can be done concurrently.
- **Range:** 1 through 10240.
- **Default:** 1024.
- `threshold` *`value`*—Indicate the maximum number of destination IP address matches required to raise an alarm.
- **Range:** 1 through 4294967295.
- **Default:** 1000.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Security Policies Overview

destination (IPCP)

IN THIS SECTION

- Syntax | 5439
- Hierarchy Level | 5439
- Description | 5439
- Options | 5439
- Required Privilege Level | 5440
- Release Information | 5440

Syntax

```
destination address destination-profile profile-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet unnumbered-address
interface-name],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet unnumbered-address interface-name]
```

Description

For unnumbered interfaces with PPP encapsulation, specify the IP address of the remote interface.

Options

address—IP address of the remote interface.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPCP Options for Interfaces with PPP Encapsulation

address | 1846

negotiate-address | 11261

Junos OS Administration Library for Routing Devices

destination-ipv4-address

IN THIS SECTION

- Syntax | 5440
- Hierarchy Level | 5441
- Description | 5441
- Options | 5441
- Required Privilege Level | 5441
- Release Information | 5441

Syntax

```
destination-ipv4-address ipv4-address;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Description

Specify the IPv4 address of the Path Computation Element (PCE) to which the Path Computation Client (PCC) should connect.

Options

ipv4-address	IPv4 address of the PCE.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

pce | 12303

destination-ipv4-address (RFC 2544 Benchmarking Tests)

IN THIS SECTION

- Syntax | 5442
- Junos OS Hierarchy Level | 5442
- Junos OS Evolved Hierarchy Level | 5442
- Description | 5442
- Options | 5443
- Required Privilege Level | 5443
- Release Information | 5443

Syntax

```
destination-ipv4-address address;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the destination IPv4 address to be used in generated test frames. You must configure this option if you specify inet as the family. This option is not required if you specify ccc as the family.

Options

address (Required if you specify `inet` as the family.) Valid IPv4 address.

- **Default:** When you specify `ccc` as the family, if you do not configure the destination IPv4 address, the default value of 192.168.1.20 is used.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

rfc2544-benchmarking | **15759**

destination-ipv6-address

IN THIS SECTION

- Syntax | **5444**
- Hierarchy Level | **5444**
- Release Information | **5444**
- Description | **5444**
- Required Privilege Level | **5444**

Syntax

```
destination-ipv6-address destination-ipv6-address;
```

Hierarchy Level

```
[edit logical-systems name protocols pcep pce],
[edit protocols pcep pce]
```

Release Information

Statement introduced in Junos OS Release 21.3R1.

Description

Specify the IPv6 address of the Path Computation Element (PCE) to which the Path Computation Client (PCC) should connect.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

local-ipv6-address | 9956

destination (L2TP)

IN THIS SECTION

- Syntax | 5445
- Hierarchy Level | 5445
- Description | 5445
- Required Privilege Level | 5446
- Release Information | 5446

Syntax

```
destination {
    address ip-address {
        access-line-information <connection-speed-update>;
        drain;
        routing-instance routing-instance-name {
            drain;
        }
    }
    lockout-timeout seconds;
    name destination-name {
        drain;
    }
}
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Configure attributes for all L2TP destinations or a specified L2TP destination.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

L2TP LNS Inline Service Interfaces

L2TP Subscriber Access Lines and Connection Speeds

destination (LRF Profile)

IN THIS SECTION

- Syntax | 5446
- Hierarchy Level | 5447
- Description | 5447
- Required Privilege Level | 5447
- Release Information | 5447

Syntax

```
destination {
    address collector-address;
```

```
    port collector-port-number;
}
```

Hierarchy Level

```
[edit services lrf profile profile-name collector collector-name]
```

Description

Specify the destination IP address and port number of the collector.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

destination-mac-address

IN THIS SECTION

- Syntax | 5448
- Junos OS Hierarchy Level | 5448
- Junos OS Evolved Hierarchy Level | 5448
- Description | 5448
- Options | 5449
- Required Privilege Level | 5449
- Release Information | 5449

Syntax

```
destination-mac-address mac-address;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the destination MAC address used in the generated test frames. This is a mandatory parameter for family `bridge` and `vpls` and optional for family `ccc`.

Options

mac-address Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn*.*nnnn*.*nnnn* or *nn*:*nn*:*nn*:*nn*:*nn*:*nn*—for example, `0000:5e00:5355` or `00:00:5e:00:53:55`.

Default: 0x00:0x11:0xAE:0x92:0x2F:0x28

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

rfc2544-benchmarking | **15759**

rfc2544 | **15757**

RFC 2544-Based Benchmarking Tests Overview

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

Configuring an RFC 2544-Based Benchmarking Test

destination-nat

IN THIS SECTION

- Syntax | **5450**
- Hierarchy Level | **5450**
- Description | **5450**
- Options | **5450**

Required Privilege Level | 5450
Release Information | 5451

Syntax

```
destination-nat (off | pool pool-name | rule-session-count-alarm (clear-threshold value | raise-threshold value));
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name rule rule-name then]
```

Description

Specify the action of the destination NAT rule.

Options

off Do not perform destination NAT operation.

pool Use user-defined destination NAT pool to perform destination NAT.

rule-session-count-alarm Define session count alarm thresholds for a specific destination NAT rule. When the session count exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered.

NOTE: If you enter a value for `raise-threshold` but not for `clear-threshold`, `clear-threshold` is automatically set to 80 percent of `raise-threshold`.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. The `rule-session-count-alarm` option added in Junos OS Release 12.1X45-D10.

destination-networks

IN THIS SECTION

- Syntax | 5451
- Hierarchy Level | 5451
- Description | 5452
- Options | 5452
- Required Privilege Level | 5452
- Release Information | 5452

Syntax

```
destination-networks prefix;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options
dynamic-tunnels tunnel-name],
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name rsvp-te
entry],
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name],
[edit routing-instances routing-instance-name routing-options dynamic-tunnels tunnel-name],
[edit routing-instances routing-instance-name routing-options dynamic-tunnels tunnel-name  rsvp-
te entry],
[edit routing-options dynamic-tunnels tunnel-name],
[edit routing-options dynamic-tunnels tunnel-name rsvp-te entry]
```

Description

Specify the IPv4 prefix range for the destination network. Only tunnels within the specified IPv4 prefix range can be created.

Options

prefix—Destination prefix of the network.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Configuring GRE Tunnels for Layer 3 VPNs*
- *Dynamic Tunnels Overview*
- Configuring RSVP Automatic Mesh

destination-option

IN THIS SECTION

- Syntax | 5453
- Hierarchy Level | 5453
- Description | 5453
- Required Privilege Level | 5453
- Release Information | 5453

Syntax

```
destination-option {
    home-address {
        match (equal | greater-than | less-than | not-equal);
        value header-value;
    }
    option-type {
        match (equal | greater-than | less-than | not-equal);
        value header-value;
    }
}
```

Hierarchy Level

```
[edit  set security idp custom-attack attack-name attack-type signature protocol ipv6 extension-
header]
```

Description

Specify the IPv6 destination option for the extension header. The `destination-option` option inspects the header option type of `home-address` field in the `extension header` and reports a custom attack if a match is found. The `destination-option` supports the `home-address` field type of inspection.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

destination-override

IN THIS SECTION

- Syntax | 5454
- Hierarchy Level | 5454
- Description | 5454
- Options | 5454
- Required Privilege Level | 5455
- Release Information | 5455

Syntax

```
destination-override {
    syslog host ip-address;
}
```

Hierarchy Level

```
[edit system tracing]
```

Description

This option overrides the system-wide configuration under `[edit system tracing]` and has no effect if system tracing is not configured.

Options

These options specify the system logs and the host to which remote tracing output is sent:

- `syslog`—Specify the system process log files to send to the remote tracing host.
- `host` *`ip-address`*—Specify the IP address to which to send tracing information.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Junos OS Tracing and Logging Operations

tracing | 29987

destination-path-group

IN THIS SECTION

- Syntax | 5455
- Hierarchy Level | 5456
- Description | 5456
- Options | 5456
- Required Privilege Level | 5457
- Release Information | 5457

Syntax

```
destination-path-group group-name {
    active-probe-properties {
         active-probe-only ;
    }
    inline-gre-encap
    overlay-path {
```

```
        overlay-path-name;
    }
    probe-routing-instance {
        routing-instance-name;
    }
}
```

Hierarchy Level

[edit security advance-policy-based-routing]

Description

Define a group containing multiple overlay paths terminating at a same destination.

In releases prior to Junos OS release 20.2R1, AppQoE determines the applicable destination path group and binds the application sessions to that particular destination-path-group. Based on the application's SLA requirements and link preferences, AppQoE determines the best link among all the links in that destination-path-group. All the instances of the applications use the same best path in the chosen DPG. If there is SLA violation on the current link, then AppQoE determines a new best-link only among all the links in the chosen destination-path-group. In the case of primary hub failover or routing change, AppQoE selects another DPG and then binds the session to the new destination-path-group.

Staring in Junos OS release 20.2R1, with support for the active-active deployment, you can select the links across the destination path groups if an end-point is reachable through them. The best link selected among all the links in the selected destination path group sends the application traffic.

Options

group-name	Name that identifies the destination path group.	
active-probe-properties	Specify additional active probe property.	
	active-probe-only	Enable active probing only to the destination path group
inline-gre-encap	Enable inline GRE encapsulation.	
overlay-path ***overlay-path-name***	Overlay path name.	
probe-routing-instance ***routing-instance-name***	Routing instance for the probe path.	

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Application Quality of Experience

Advanced Policy-Based Routing

destination-pool

IN THIS SECTION

- Syntax | 5457
- Hierarchy Level | 5458
- Description | 5458
- Options | 5458
- Required Privilege Level | 5458
- Release Information | 5458

Syntax

```
destination-pool nat-pool-name;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated]
```

Description

Specify the destination address pool for translated traffic.

Options

nat-pool-name—Destination pool name.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

destination-port (firewall)

IN THIS SECTION

- Syntax | 5459
- Hierarchy Level | 5459
- Description | 5459
- Options | 5459

Required Privilege Level | 5459
Release Information | 5459

Syntax

```
destination-port <destination-port>;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name from ip-version ip-version
protocol (tcp|udp)]
```

Description

Configure the destination port of the Layer 4 header.

Options

destination-port The destination port of the Layer 4 header.

- **Range:** 0 through 65,535

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1R1.

RELATED DOCUMENTATION

Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic

destination-port (protocols pcep)

IN THIS SECTION

- Syntax | 5460
- Hierarchy Level | 5460
- Description | 5460
- Options | 5460
- Required Privilege Level | 5461
- Release Information | 5461

Syntax

```
destination-port port-number;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Description

Specify the TCP port number of the Path Computation Element (PCE) to which the Path Computation Client (PCC) should connect.

Options

port-number	Destination TCP port number.

- **Range:** 1 through 65535
- **Default:** 4189

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFXswitches added in Junos OS Release 16.1R3.

RELATED DOCUMENTATION

pce | 12303

destination-port (RPM and TWAMP)

IN THIS SECTION

- Syntax | 5462
- Junos OS Hierarchy Levels | 5462
- Junos OS Evolved Hierarchy Level | 5462
- Description | 5462
- Options | 5463
- Required Privilege Level | 5463
- Release Information | 5463

Syntax

```
destination-port port;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name]
[edit services rpm twamp client control-connection control-client-name test-session test-session-
name
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) port to which a probe is sent. This statement is used only for TCP or UDP probe types.

The value for the `destination-port` can be only 7 when you configure the destination port along with hardware timestamping. A constraint check prevents you for configuring any other value for the destination port in this case.

This constraint does not apply when you are using one-way hardware timestamping along with `destination-port` and either `probe-type udp-ping` or `probe-type udp-ping-timestamp`.

For a managed TWAMP control connection (the default or configured with the `control-type managed` statement), the `destination-port` statement is required at the `[edit services rpm twamp control-connection control-connection-name]` hierarchy level (Junos OS) or the `[edit services monitoring twamp control-connection control-connection-name]` hierarchy level (Junos OS Evolved) and is not allowed at the individual test-session level, at `[edit services rpm twamp control-connection control-connection-name test-session test-session-name]` (Junos OS) or `[edit services monitoring twamp control-connection control-connection-name test-session test-session-name]` (Junos OS Evolved).

For a TWAMP control connection configuration that includes the `control-type light` statement, the `destination-port` statement is not allowed at the `[edit services rpm twamp control-connection control-connection-`

name] hierarchy level (Junos OS) or the [`edit services monitoring twamp control-connection` *control-connection-name*] hierarchy level (Junos OS Evolved), but is required for each test session, at [`edit services rpm twamp control-connection` *control-connection-name* `test-session` *test-session-name*] (Junos OS) or [`edit services monitoring twamp control-connection` *control-connection-name* `test-session` *test-session-name*] (Junos OS Evolved).

Options

port—Port number `7` or from `862` through `65,535`.

- **Default:** The default value for the port is 862 to which the TWAMP client establishes the control connection.

> **NOTE**: The specified port numbers are recommended for RPM only.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.

Support at the [`edit services rpm twamp client control-connection` *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

Support at the [`edit services rpm twamp client control-connection` *control-client-name* `test-session` *test-session-name*] hierarchy level for Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Configuring BGP Neighbor Discovery Through RPM

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

destination-port (applications)

IN THIS SECTION

- Syntax | 5464
- Hierarchy Level | 5464
- Description | 5464
- Options | 5464
- Required Privilege Level | 5465
- Release Information | 5465

Syntax

```
destination-port port-value;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number.

Options

port-value—Identifier for the port or range of ports. For a complete list of supported application destination port requirements, see .

- **Range:** 1 through 65,535

> **NOTE**: If you specify a value of 0 as a destination port or beginning of a destination report range, you will receive the following error:

```
'application application-name'
    TCP Destination Port 0 Invalid
 error: configuration check-out failed
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

ALG Descriptions

Configuring Application Sets

Configuring Application Properties

Examples: Configuring Application Protocols

Verifying the Output of ALG Sessions

Understanding Two-Way Active Measurement Protocol on Routers

destination-port (Applications)

IN THIS SECTION

- Syntax | 5466
- Hierarchy Level | 5466
- Description | 5466
- Options | 5466
- Required Privilege Level | 5470
- Release Information | 5470

Syntax

```
destination-port port-identifier;
```

Hierarchy Level

```
[edit applications application application-name],
[edit applications application  application-name term term-name]
```

Description

Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number, which combines with `protocol` to identify an application type. The Internet Assigned Numbers Authority (IANA) assigns port numbers. See the IANA *Service Name and Transport Protocol Port Number Registry* at http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml for a list of assigned port numbers.

Options

port-identifier —Range of ports. You can use a numeric value or one of the text synonyms listed in Table 28 on page 5466.

Table 28: Port Supported by Services Interfaces

Port Name	Corresponding Port Number
`afs`	1483
`bgp`	179
`biff`	512
`bootpc`	68
`bootps`	67

Table 28: Port Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
cmd	514
cvspserver	2401
dhcp	67
domain	53
eklogin	2105
ekshell	2106
excc	512
finger	79
ftp	21
ftp-data	20
http	80
https	443
ident	113
imap	143
kerberos-sec	88

Table 28: Port Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
klogin	543
kpasswd	761
krb-prop	754
krbupdate	760
kshell	544
ldap	389
ldp	646
login	513
mobileip-agent	434
mobilip-mn	435
msdp	639
netbios-dgm	138
netbios-ns	137
netbios-ssn	139
nfsd	2049

Table 28: Port Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
nntp	119
ntalk	518
ntp	123
pop3	110
pptp	1723
printer	515
radacct	1813
radius	1812
rip	520
rkinit	2108
smtp	25
snmp	161
snmp-trap	162
snpp	444
socks	1080

Table 28: Port Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
ssh	22
sunrpc	111
syslog	514
tacacs	49
tacacs-ds	65
talk	517
telnet	23
tftp	69
timed	525
who	513
xdmcp	177

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policy Application Sets Overview

destination-port (Flow Monitoring Logs for NAT)

IN THIS SECTION

- Syntax | 5471
- Hierarchy Level | 5471
- Description | 5471
- Options | 5471
- Required Privilege Level | 5472
- Release Information | 5472

Syntax

```
destination-port port-number;
```

Hierarchy Level

```
[edit services jflow-log collector collector-name]
```

Description

Specify the UDP port of the destination to be used in the UDP header for the generated flow monitoring logs. This is a required setting.

Options

port-number—UDP port number for the test frames.

- **Default:** 4041

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

destination-port (PCP)

IN THIS SECTION

- Syntax | 5472
- Hierarchy Level | 5473
- Description | 5473
- Options | 5473
- Required Privilege Level | 5473
- Release Information | 5473

Syntax

```
destination-port high maximum-value low minimum-value;
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name from]
```

Description

Specify the destination port range that must be matched for the PCP rule. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

maximum-value	Upper boundary for the port range.
minimum-value	Lower boundary for the port range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

destination-port (RADIUS Snoop Segment)

IN THIS SECTION

- Syntax | 5474
- Hierarchy Level | 5474
- Description | 5474
- Options | 5474
- Required Privilege Level | 5475
- Release Information | 5475

Syntax

```
destination-port destination-port;
```

Hierarchy Level

```
[edit access radius snoop-segments segment-name]
```

Description

Specify the destination port for accounting messages to snoop.

Options

destination-port

Destination port of accounting messages.

- **Default:** 1813
- **Range:** 1 through 65,535

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview

IP-Based Subscriber Setup Overview

destination-port range

IN THIS SECTION

- Syntax | 5475
- Hierarchy Level | 5476
- Description | 5476
- Options | 5476
- Required Privilege Level | 5476
- Release Information | 5476

Syntax

```
destination-port range high maximum-value low minimum-value;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name from]
```

Description

Specify the destination port range for rule matching.

Options

maximum-value—Upper limit of port range for matching.

minimum-value—Lower limit of port range for matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Port Forwarding for Static Destination Address Translation

destination-port-range (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 5477
- Hierarchy Level | 5477
- Description | 5477

- Options | 5477
- Required Privilege Level | 5477
- Release Information | 5478

Syntax

```
destination-port-range {
    high port-number low port-number;
}
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number from]
```

Description

Specify a port range that the HTTP request destination port number must match. You can specify multiple port ranges by including the `destination-port-range` statement multiple times.

After this criterion and the other match criteria specified for the `term` are matched, the HTTP header enrichment actions specified for the `term` are applied to the HTTP traffic.

Options

high *port-number*	Upper limit of the port range.
low *port-number*	Lower limit of the port range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm (HTTP Header Enrichment) | 8005

destination-port (Security Destination NAT)

IN THIS SECTION

- Syntax | 5478
- Hierarchy Level | 5478
- Description | 5478
- Options | 5479
- Required Privilege Level | 5479
- Release Information | 5479

Syntax

```
destination-port (port-or-low <to high>);
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name rule rule-name match]
```

Description

Specify a destination port or port range to match the rule. Up to eight port or port ranges are supported.

Options

port —Specify a destination port number.

low—Specify the lower limit of the destination port range.

<to *high*>—Specify the upper limit of the destination port range.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X47-D10.

destination-port (Security Forwarding Options)

IN THIS SECTION

- Syntax | 5479
- Hierarchy Level | 5480
- Description | 5480
- Required Privilege Level | 5484
- Release Information | 5484

Syntax

```
destination-port port-number;
```

Hierarchy Level

```
[edit security forwarding-options mirror-filter filter-name]
```

Description

Specify a Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number to be matched for mirroring. You can specify a numeric value or one of the text synonyms listed in Table 29 on page 5480.

Table 29: Ports Supported by Services Interfaces

Port Name	Corresponding Port Number
afs	1483
bgp	179
biff	512
bootpc	68
bootps	67
cmd	514
cvspserver	2401
dhcp	67
domain	53
eklogin	2105

Table 29: Ports Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
ekshell	2106
excc	512
finger	79
ftp	21
ftp-data	20
http	80
https	443
ident	113
imap	143
kerberos-sec	88
klogin	543
kpasswd	761
krb-prop	754
krbupdate	760
kshell	544

Table 29: Ports Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
ldap	389
ldp	646
login	513
mobileip-agent	434
mobilip-mn	435
msdp	639
netbios-dgm	138
netbios-ns	137
netbios-ssn	139
nfsd	2049
nntp	119
ntalk	518
ntp	123
pop3	110
pptp	1723

Table 29: Ports Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
printer	515
radacct	1813
radius	1812
rip	520
rkinit	2108
smtp	25
snmp	161
snmp-trap	162
snpp	444
socks	1080
ssh	22
sunrpc	111
syslog	514
tacacs	49
tacacs-ds	65

Table 29: Ports Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
`talk`	517
`telnet`	23
`tftp`	69
`timed`	525
`who`	513
`xdmcp`	177

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

mirror-filter (Security Forwarding Options) | **10768**

show security forwarding-options mirror-filter | **24363**

destination-port (Security Signature Attack)

IN THIS SECTION

- Syntax | 5485
- Hierarchy Level | 5485
- Description | 5485
- Options | 5485
- Required Privilege Level | 5486
- Release Information | 5486

Syntax

```
destination-port {
    match (equal | greater-than | less-than | not-equal);
    value destination-port;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol udp],

[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify the port number of the attack target.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`destination-port`*—Match the port number of the attack target.

- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

destination-port (Security Source NAT)

IN THIS SECTION

- Syntax | 5486
- Hierarchy Level | 5486
- Description | 5487
- Options | 5487
- Required Privilege Level | 5487
- Release Information | 5487

Syntax

```
destination-port (port-or-low <to high>);
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name match]
```

Description

Specify a destination port or port range to match the rule. Up to eight port or port ranges are supported.

Options

port —Specify a destination port number.

low—Specify the lower limit of the destination port range.

<to *high*>—Specify the upper limit of the destination port range.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X47-D10.

destination-port (Security Static NAT)

IN THIS SECTION

- Syntax | 5488
- Hierarchy Level | 5488
- Description | 5488
- Options | 5488
- Required Privilege Level | 5488
- Release Information | 5488

Syntax

```
destination-port (port-or-low | <to high>);
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name match]
```

Description

Specify a destination port or port range to allow static NAT to map ports.

Options

port-or-low Specify the port name or the lower limit of the port range.

to *high* Specify the upper limit of the port range.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

destination-port (SNMP)

IN THIS SECTION

- Syntax | **5489**
- Hierarchy Level | **5489**

- Description | 5489
- Default | 5489
- Options | 5489
- Required Privilege Level | 5489
- Release Information | 5490

Syntax

```
destination-port port-number;
```

Hierarchy Level

```
[edit snmp trap-group]
```

Description

Assign a trap port number other than the default.

Default

If you omit this statement, the default port is 162.

Options

port-number—SNMP trap port number.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring SNMP Trap Groups

destination-port (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 5490
- Hierarchy Level | 5490
- Description | 5490
- Options | 5491
- Required Privilege Level | 5491
- Release Information | 5491

Syntax

```
destination-port port-number;
```

Hierarchy Level

```
[edit services radius-flow-tap policy policy-name inet drop-policy rule-name from],
[edit services radius-flow-tap policy policy-name inet6 drop-policy rule-name from]
```

Description

Specify the destination IP address for the radius-flow-tap policy.

Options

port-number— Number of the IPv4 or IPv6 destination port for the radius-flow-tap policy.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Subscriber Secure Policy Overview

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

destination-ports (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 5491
- Hierarchy Level | 5492
- Description | 5492
- Options | 5492
- Required Privilege Level | 5492
- Release Information | 5492

Syntax

```
destination-ports value;
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number from]
```

Description

Specify the HTTP request destination port number that must be matched. You can specify multiple ports by including the `destination-ports` statement multiple times.

After this criterion and the other match criteria specified for the `term` are matched, the HTTP header enrichment actions specified for the `term` are applied to the HTTP traffic.

Options

value—Port number.

- **Range:** 0 through 65,535

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm (HTTP Header Enrichment) | **8005**

destination-prefix (Destination NAT Next Gen Services)

IN THIS SECTION

- Syntax | 5493
- Hierarchy Level | 5493
- Description | 5493
- Required Privilege Level | 5493
- Release Information | 5493

Syntax

```
destination-prefix destination-prefix;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name then destination-nat]
```

Description

Specify the IPv6 prefix that is used to embed an IPv4 destination address in an IPv6 address. The `destination-prefix` statement is used in Stateful NAT64 and 464XLAT translations.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

destination-prefix (IDS)

IN THIS SECTION

- Syntax | 5494
- Hierarchy Level | 5494
- Description | 5494
- Options | 5494
- Required Privilege Level | 5495
- Release Information | 5495

Syntax

```
destination-prefix prefix-value;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then aggregation]
```

Description

Specify a prefix length for destination IPv4 address aggregation for the IDS rule. This applies session limits to an aggregation of all attacks to a subnet of the specified length.

For example, if you configure a value of 24 for `destination-prefix`, then attacks to 10.1.1.2 and 10.1.1.3 are counted as attacks to the 10.1.1/24 subnet.

Options

prefix-value—Integer value.

- **Range:** 1 through 32

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in IDS Rules

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

destination-prefix-ipv6 (IDS)

IN THIS SECTION

- Syntax | 5495
- Hierarchy Level | 5496
- Description | 5496
- Options | 5496
- Required Privilege Level | 5496
- Release Information | 5496

Syntax

```
destination-prefix-ipv6 prefix;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then aggregation]
```

Description

Specify a prefix length for destination IPv6 address aggregation for the IDS rule. This applies session limits to an aggregation of all attacks to a subnet of the specified length.

For example, if you configure a value of 64 for `destination-prefix-ipv6`, then attacks to 2001:db8:1234:72a2::2 and 2001:db8:1234:72a2::3 are counted as attacks to the 2001:db8:1234:72a2::/64 subnet.

Options

`prefix-value`—Integer value.

- **Range:** 1 through 128

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Configuring Actions in IDS Rules*
>
> *Configuring Protection Against Network Attacks on an MS-MPC*
>
> *Understanding IDS on an MS-MPC*

destination-prefix-list (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 5497
- Hierarchy Level | 5497
- Description | 5497
- Options | 5497
- Required Privilege Level | 5498
- Release Information | 5498

Syntax

```
destination-prefix-list {
    (prefix-name | prefix-name except);
}
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number from]
```

Description

Specify the destination prefix list that the HTTP request destination IP address must match. You can specify multiple prefix lists by including the `destination-prefix-list` statement multiple times.

After this criterion and the other match criteria specified for the `term` are matched, the HTTP header enrichment actions specified for the `term` are applied to the HTTP traffic.

Options

prefix-name	Name of the prefix list.

> **NOTE**: The prefix list must already be defined at the `[edit policy-options prefix-list]` hierarchy level.

***prefix-name* except**

Exclude addresses that are in the specified prefix list from addresses that are in the `destination-address` or `destination-address-range` statement configured at the `[edit services hcm tag-rule` *`rule-name`* `term` *`term-number`* `from]` hierarchy level. You cannot use `except` without also configuring addresses that do match.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm (HTTP Header Enrichment) | **8005**

destination-prefix-list (PCP)

IN THIS SECTION

- Syntax | **5499**
- Hierarchy Level | **5499**
- Description | **5499**
- Options | **5499**
- Required Privilege Level | **5499**
- Release Information | **5499**

Syntax

```
destination-prefix-list list-name;
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name from]
```

Description

Specify the destination prefix list that must be matched for the PCP rule. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

list-name Destination prefix list.

except (Optional) Prevent the specified prefix list from matching the PCP rule.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

destination-prefix-list (Services AACL)

IN THIS SECTION

- Syntax | 5500
- Hierarchy Level | 5500
- Description | 5500
- Options | 5500
- Required Privilege Level | 5500
- Release Information | 5501

Syntax

```
destination-prefix-list list-name;
```

Hierarchy Level

```
[edit services aacl rule rule-name term term-name from]
```

Description

Specify the destination prefix list for rule matching. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level.

Options

`list-name`—Destination prefix list.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Match Conditions in AACL Rules

destination-prefix-list (Services CoS)

IN THIS SECTION

- Syntax | 5501
- Hierarchy Level | 5501
- Description | 5501
- Options | 5502
- Required Privilege Level | 5502
- Release Information | 5502

Syntax

```
destination-prefix-list list-name <except>;
```

Hierarchy Level

```
[edit services cos rule rule-name term term-name from]
```

Description

Specify the destination prefix list for rule matching. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level.

Options

list-name—Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

> *Configuring CoS Rules*
>
> https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-policy/config-guide-policy.html

destination-prefix-list (Services IDS)

IN THIS SECTION

- Syntax | 5503
- Hierarchy Level | 5503
- Description | 5503
- Options | 5503
- Required Privilege Level | 5503
- Release Information | 5503

Syntax

```
destination-prefix-list list-name <except>;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name from]
```

Description

Specify the destination prefix list for rule matching. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level.

Options

`list-name`—Destination prefix list.

`except`—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

> https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-policy/config-guide-policy.html
>
> *Configuring Match Conditions in IDS Rules*

destination-prefix-list (Services NAT)

IN THIS SECTION

- Syntax | 5504
- Hierarchy Level | 5504
- Description | 5504
- Options | 5504
- Required Privilege Level | 5505
- Release Information | 5505

Syntax

```
destination-prefix-list list-name <except>;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name from]
```

Description

Specify the destination prefix list for rule matching. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level.

If the `translation-type` statement in the `then` statement of the nat rule is set to **stateful-nat-64**, the destination prefix list for rule matching must be within the range specified by the `destination-prefix` statement in the `then` statement.

Options

list-name—Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

> *Network Address Translation Rules Overview*
>
> https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-policy/config-guide-policy.html

destination-prefix-list (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 5505
- Hierarchy Level | 5506
- Description | 5506
- Options | 5506
- Required Privilege Level | 5506
- Release Information | 5506

Syntax

```
destination-prefix-list list-name <except>;
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name term term-name from]
```

Description

Specify the destination prefix list for rule matching. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level.

Options

`list-name`—Destination prefix list.

`except`—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

> *Configuring Match Conditions in Stateful Firewall Rules*
>
> https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-policy/config-guide-policy.html

destination-prefix (Security Forwarding Options)

IN THIS SECTION

- Syntax | 5507
- Hierarchy Level | 5507
- Description | 5507
- Required Privilege Level | 5507
- Release Information | 5507

Syntax

```
destination-prefix destination-prefix;
```

Hierarchy Level

```
[edit security forwarding-options mirror-filter filter-name]
```

Description

Specify the destination IP prefix or address to be matched for mirroring.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

mirror-filter (Security Forwarding Options) | 10768

show security forwarding-options mirror-filter | 24363

destination-prefix (Services NAT)

IN THIS SECTION

- Syntax | 5508
- Hierarchy Level | 5508
- Description | 5508
- Options | 5508
- Required Privilege Level | 5509
- Release Information | 5509

Syntax

```
destination-prefix destination-prefix;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated]
```

Description

Specify the destination prefix for translated traffic.

Options

destination-prefix—IPv4 or IPv6 destination prefix value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

destination-prefix option enhanced to support IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

destination-profile

IN THIS SECTION

- Syntax | 5509
- Hierarchy Level | 5510
- Description | 5510
- Options | 5510
- Required Privilege Level | 5510
- Release Information | 5510

Syntax

```
destination-profile name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address],
[edit interfaces interface-name unit logical-unit-number family inet unnumbered-address
interface-name destination address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet address address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet unnumbered-address interface-name destination address]
```

Description

For interfaces with PPP encapsulation, assign PPP properties to the remote destination end. You define the profile at the `[edit access group-profile name ppp]` hierarchy level.

Options

name—Profile name defined at the `[edit access group-profile name ppp]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPCP Options for Interfaces with PPP Encapsulation

destination (IPCP) | **5439**

Junos OS Administration Library for Routing Devices

destination (Protocols OSPF)

IN THIS SECTION

- Syntax | 5511
- Hierarchy Level | 5512
- Description | 5512
- Options | 5512
- Required Privilege Level | 5512
- Release Information | 5512

Syntax

```
destination prefix {
    interface (interface-name |all) {
        admin-group {
            exclude  [ group-name ];
            include-all  [ group-name ];
            include-any  [ group-name ];
            preference  [ group-name ];
        }
        bandwidth-greater-equal-primary;
        dest-metric  (highest | lowest);
        downstream-paths-only ;
        evaluation-order [ admin-group srlg bandwidth protection-type neighbor neighbor-tag
metric ];
        metric-order  [ root dest ];
        node {
            exclude [ node-address ];
            preference [ node-address ];
        }
        protection-type  (link |node | node-link);
        root-metric  (highest | lowest) ;
        srlg (loose |strict);
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options backup-selection],
[edit logical-systems logical-system-name routing-instances instance-name routing-options backup-
selection],
[edit routing-instances instance-name routing-options backup-selection],
[edit routing-options backup-selection]
```

Description

Define the backup selection policy for a particular destination prefix or for all the prefixes.

Options

The remaining statements are explained separately. See CLI Explorer.

prefix Destination prefix name and prefix length. You can specify 0/0 for the IPv4 least-specific prefix or 0::0/0 for the IPv6 least-specific prefix.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

destination-realm (Gx-Plus)

IN THIS SECTION

Syntax | 5513

- Hierarchy Level | 5513
- Description | 5513
- Options | 5513
- Required Privilege Level | 5513
- Release Information | 5513

Syntax

```
destination-realm realm;
```

Hierarchy Level

```
[edit access gx-plus partition partition-name]
```

Description

Configure the realm in which the PCRF host resides.

Options

realm—Realm in which the PCRF host resides.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Gx-Plus

Configuring the Gx-Plus Partition

destination-realm (JSRC)

IN THIS SECTION

- Syntax | 5514
- Hierarchy Level | 5514
- Description | 5514
- Options | 5514
- Required Privilege Level | 5515
- Release Information | 5515

Syntax

```
destination-realm realm
```

Hierarchy Level

```
[edit jsrc partition partition-name]
```

Description

Configure the realm in which the SAE host resides.

Options

realm—Realm in which the SAE host resides.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

JSRC Configuration Overview

Configuring the JSRC Partition

destination-realm (OCS Partition)

IN THIS SECTION

- Syntax | 5515
- Hierarchy Level | 5516
- Description | 5516
- Options | 5516
- Required Privilege Level | 5516
- Release Information | 5516

Syntax

```
destination-realm ocs-realm-name;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Configure the Destination-Realm AVP value in all CCR-GY messages.

Options

ocs-realm-name—Name of the Destination-Real AVP value to be used in all CCR-GY messages.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

destination-realm (PCRF Partition)

IN THIS SECTION

- Syntax | 5517
- Hierarchy Level | 5517
- Description | 5517
- Options | 5517
- Required Privilege Level | 5517
- Release Information | 5518

Syntax

```
destination-realm pcrf-realm-name;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the Destination-Realm AVP value in all CCR-GX messages.

Options

pcrf-realmname—Value of the Destination-Real AVP to be used in all CCR-GX messages.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

Understanding Upstream and Downstream Messages for the PCRF

destination (Routing Instance)

IN THIS SECTION

- Syntax | 5518
- Hierarchy Level | 5519
- Description | 5519
- Default | 5519
- Options | 5519
- Required Privilege Level | 5519
- Release Information | 5519

Syntax

```
destination routing-instance-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number tunnel routing-instance]
```

Description

Specify the destination routing instance that points to the routing table containing the tunnel destination address.

Default

The default Internet routing table `inet.0`.

Options

routing-instance-name—Name of the destination routing instance.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Tunnel Interfaces for Routing Table Lookup

destination (saegw system)

IN THIS SECTION

- Syntax | 5520
- Hierarchy Level | 5520
- Description | 5520
- Options | 5521
- Required Privilege Level | 5521
- Release Information | 5521

Syntax

```
destination {
    peer-address address;
    port port-number;
    routing-instance instance-name;
    source-address address;
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw system]
```

Description

Configure the destination address of the IPsec peer or content of communication point of aggregation (CC-PAG) peer on the user plane for lawful intercept data. The control plane provides the address of the Mediation and Delivery Function (MDF), but it might not provide the address of the IPsec peer, which you need to configure. This configuration is applicable to both 4G and 5G.

NOTE: Currently, this configuration supports only IPv4 addresses.

If you configure the peer address, you must also configure the source address because you must define both endpoints of the IPsec tunnel. You can configure the routing instance independent of the source and peer address.

A change in the configuration takes effect only when the control plane deactivates and reactivates the lawful intercept task. Existing tasks maintain their original CLI-configured values.

Options

peer-address *address*	Address of the IPsec peer.
port *port-number*	TCP/UDP port number of the IPsec peer.
routing-instance *instance-name*	Name of the routing instance.
source-address *address*	Address of the user plane (SAEGW-U or UPF).

Required Privilege Level

system

Release Information

Juniper Networks introduced this statement in Junos OS Release 22.2R1

RELATED DOCUMENTATION

Lawful Intercept on Junos Multi-Access User Plane

destination (Security Destination NAT)

IN THIS SECTION

- Syntax | 5522
- Hierarchy Level | 5523
- Description | 5523
- Options | 5523

- Required Privilege Level | 5523
- Release Information | 5523

Syntax

```
destination {
    pool pool-name {
        address <ip-address> {
            (port port-number | to ip-address);
        }
        description text;
        routing-instance (routing-instance-name | default);
    }
    rule-set rule-set-name {
        description text;
        from {
            interface [interface-name];
            routing-instance [routing-instance-name];
            zone [zone-name];
        }
        rule rule-name {
            description text;
            match {
                application {
                    [application];
                    any;
                }
                (destination-address ip-address| destination-address-name address-name);
                destination-port (port-or-low <to high>);
                protocol [protocol-name-or-number];
                source-address [ip-address];
                source-address-name [address-name];
            }
            then {
                destination-nat (off | pool pool-name |rule-session-count-alarm (clear-threshold
value | raise-threshold value));
            }
        }
```

```
    }
}
```

Hierarchy Level

```
[edit security nat]
```

Description

Configure destination NAT, which allows you to configure the following:

- Translate destination IP address or addresses to a specific IP address.
- Translate destination IP address or addresses and port number(s) to a specific IP address and one port number.
- Translate a range of destination IP addresses to another range of IP addresses. This mapping is one-to-one, static, and without PAT.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. The `description` option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

destination (Security IP Headers Attack)

IN THIS SECTION

- Syntax | 5524
- Hierarchy Level | 5524
- Description | 5524
- Options | 5524
- Required Privilege Level | 5525
- Release Information | 5525

Syntax

```
destination {
    match (equal | greater-than | less-than | not-equal);
    value ip-address-or-hostname;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol ipv4],

[edit security idp custom-attack attack-name attack-type signature protocol ipv6]
```

Description

Specify the IP address of the attack target.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`ip-address-or-hostname`*—Match an IP address or a hostname.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

destination-threshold

IN THIS SECTION

- Syntax | 5525
- Hierarchy Level | 5525
- Description | 5526
- Options | 5526
- Required Privilege Level | 5526
- Release Information | 5526

Syntax

```
destination-threshold number;
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp syn-flood]
```

Description

Specify the number of SYN segments received per second for a single destination IP address before the device begins dropping connection requests to that destination. If a protected host runs multiple services, you might want to set a threshold based only on the destination IP address, regardless of the destination port number.

Options

number —Number of SYN segments received per second before the device begins dropping connection requests.

- **Range:** 4 through 1,000,000 per second
- **Default:** 4000 per second

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

attack-threshold | **2528**

destination (Tunnel Remote End)

IN THIS SECTION

Syntax | **5527**

- Hierarchy Level | 5527
- Description | 5527
- Options | 5527
- Required Privilege Level | 5527
- Release Information | 5527

Syntax

```
destination destination-address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number tunnel],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
tunnel]
```

Description

For tunnel interfaces, specify the remote address of the tunnel.

Options

destination-address—Address of the remote side of the connection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Unicast Tunnels

Configuring Traffic Sampling on MX, M and T Series Routers

Configuring Flow Monitoring

destination (Tunnels)

IN THIS SECTION

- Syntax | 5528
- Hierarchy Level | 5528
- Description | 5529
- Options | 5529
- Required Privilege Level | 5529
- Release Information | 5529

Syntax

```
destination address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address],
[edit interfaces interface-name unit logical-unit-number family inet unnumbered-address interface-
name],
[edit interfaces interface-name unit logical-unit-number tunnel],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet address address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet unnumbered-address interface-name],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
tunnel]
```

Description

For encrypted, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection.

Options

address—Address of the remote side of the connection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring the Interface Address
- Configuring Generic Routing Encapsulation Tunneling
- Junos OS Services Interfaces Library for Routing Devices

destination-udp-port

IN THIS SECTION

- Syntax | 5530
- Junos OS Hierarchy Level | 5530
- Junos OS Evolved Hierarchy Level | 5530
- Description | 5530
- Options | 5530
- Required Privilege Level | 5530
- Release Information | 5530

Syntax

```
destination-udp-port  port-number;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

(Required if you specify `inet` as the family.) Specify the UDP port of the destination to be used in the UDP header for the generated frames. For other families, if you do not specify the UDP port, the default value of 4041 is used.

Options

port-number UDP port number for the test frames

- **Default:** 4041

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

rfc2544-benchmarking | **15759**

destination-udp-port (FTI)

IN THIS SECTION

- Syntax | **5531**
- Hierarchy Level | **5531**
- Description | **5531**
- Required Privilege Level | **5531**
- Release Information | **5532**

Syntax

```
destination-udp-port destination-udp-port;
```

Hierarchy Level

```
[edit interfaces name unit name tunnel encapsulation vxlan-gpe]
```

Description

Assign a numeric value to write to the destination-udp-port-field to identify a Virtual Extensible LAN (VXLAN). The range is 1 through 65,535.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Flexible Tunnel Interfaces Overview

Configuring Flexible Tunnel Interfaces

show interfaces fti | **20866**

vni(Interfaces) | **30911**

Example: Configuring Flexible Tunnel Interfaces on MX Series Routers

destinations

IN THIS SECTION

- Syntax | **5532**
- Hierarchy Level | **5533**
- Description | **5533**
- Required Privilege Level | **5533**
- Release Information | **5533**

Syntax

```
destinations {
    ftp:url {
        password "password";
    }
}
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Specify the primary and secondary destination FTP servers.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Destination FTP Servers for Flow Records

destinations (Event Policy)

IN THIS SECTION

- Syntax | 5534
- Hierarchy Level | 5534
- Description | 5534
- Options | 5534
- Required Privilege Level | 5534
- Release Information | 5534

Syntax

```
destinations {
    destination-name {
        archive-sites {
            url <password password>;
        }
        transfer-delay seconds;
    }
}
```

Hierarchy Level

```
[edit event-options]
```

Description

Define one or more destinations, each with a unique name. You can reference a destination within an event policy to indicate where to upload command output and various files, such as system log files and core files.

Options

destination-name	User-defined name of a destination.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Define Destinations for Event Policy File Archiving

destined-port

IN THIS SECTION

- Syntax | 5535
- Hierarchy Level | 5535
- Description | 5535
- Options | 5536
- Required Privilege Level | 5536
- Release Information | 5536

Syntax

```
destined-port port id;
```

Hierarchy Level

```
[edit services nat port-forwarding map-name]
```

Description

Specify the destination port number that needs to be translated to another port.

The `destined-port` statement is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, the `destined-port` statement is also supported on the MS-MPC and MS-MIC.

Options

port id The destination port number from where traffic will be forwarded.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Port Forwarding for Static Destination Address Translation

Configuring Port Forwarding Without Destination Address Translation

destruct-timeout (L2TP)

IN THIS SECTION

- Syntax | 5537
- Hierarchy Level | 5537
- Description | 5537
- Options | 5537
- Required Privilege Level | 5537
- Release Information | 5537

Syntax

```
destruct-timeout seconds;
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Set how long the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed.

BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing `no services l2tp destruct-timeout`.

Options

seconds—Length of the destruct timeout.

- **Range:** 10 through 3600
- **Default:** 300

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

L2TP LAC Subscriber Configuration

L2TP LNS Inline Service Interfaces

detect-loc

IN THIS SECTION

- Syntax | 5538
- Hierarchy Level | 5538
- Description | 5538
- Default | 5539
- Required Privilege Level | 5539
- Release Information | 5539

Syntax

```
detect-loc;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-
association ma-name mep mep-id remote-mep mep-id ]
```

```
[edit protocols oam ethernet link-fault-management interface interface-name ]
```

Description

Specify whether Ethernet OAM continuity checks are performed for an individual remote maintenance end point (MEP).

When you configure the `detect-loc` statement at [`edit protocols oam ethernet link-fault-management interface` *`interface-name`*] hierarchy level, a loss-of-continuity (LOC) defect is raised when the peer is not found within a period that is equal to 3 times the current keepalive pdu interval. When an LOC defect is raised, a syslog error message is generated .

> **NOTE**: When you configure the `detect-loc` statement at the [`edit protocols oam ethernet link-fault-management interface` *`interface-name`*] hierarchy level, any `action-profile` configured to bring down the interface is executed when an LOC defect is detected. However, the `action-profile` is not executed if you have not configured `detect-loc` statement the `detect-loc` statement at the [`edit protocols oam ethernet link-fault-management interface` *`interface-name`*] hierarchy level.

To view the current LOC status of an interface, execute the `show oam ethernet link-fault-management` command.

Default

The MEP does not generate LOC defect messages by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring a MEP to Generate and Respond to CFM Protocol Messages

remote-mep | **14352**

detection (Network Isolation Groups)

IN THIS SECTION

- Syntax | 5540
- Hierarchy Level | 5540
- Description | 5541
- Options | 5541
- Required Privilege Level | 5541
- Release Information | 5541

Syntax

```
detection {
    hold-time (<down milliseconds> <up milliseconds>);
    link-tracking {
      failure-type link;
      interface name;
      minimum-links minimum-links;
    }
    service-tracking {
      core-isolation;
    }
}
```

Hierarchy Level

```
[edit logical-systems name protocols network-isolation group],
[edit protocols network-isolation group]
```

Description

Configure detection parameters in a network isolation group for an EVPN network. The device uses these parameters to detect network isolation condition changes. You also specify how long to delay before the device records and acts on a network isolation status change.

A network isolation condition means a bridge domain, VLAN, or EVPN instance (EVI) becomes isolated from the EVPN core network, usually for one or more of the following reasons:

- The associated interconnecting physical or logical links are down.
- The device doesn't have routes to the remote endpoints in the bridge domain, VLAN, or EVI.

You can configure either the `link-tracking` option or the `service-tracking` option in a particular network isolation group, and assign that group as a network isolation profile. You can't define link tracking and service tracking parameters in the same network isolation group.

Options

hold-time down ***milliseconds***

When evaluating the device has experienced a network isolation condition, the time the device waits in milliseconds before declaring the network is not reachable. Specify a value of 1 or more milliseconds.

This setting is supported only with the `link-tracking` option.

hold-time up ***milliseconds***

When evaluating if the device has recovered from a network isolation condition, the time the device waits in milliseconds before declaring the network is reachable and acting to bring up the associated interfaces. Specify a value of 1 or more milliseconds.

You usually want to set a hold time greater than 0 so the device isn't excessively processing rapid status changes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 23.1R1.

`service-tracking` option added in Junos OS and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Determine IRB Interface State Changes from Local L2 Interface or Remote Connectivity Status in EVPN Fabrics

Layer 2 Interface Status Tracking and Shutdown Actions for EVPN Core Isolation Conditions

interface-state | **9064**

network-isolation-profile | **11314**

detection-time

IN THIS SECTION

- Syntax | **5542**
- Hierarchy Level | **5542**
- Description | **5543**
- Required Privilege Level | **5543**
- Release Information | **5543**

Syntax

```
detection-time {
    threshold milliseconds;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-
relay dhcpv6 liveness-detection method bfd],
```

```
[edit system services dhcp-local-server group group-name liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd]
```

Description

Enable failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients
- Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

detection-time (BFD for PIM)

IN THIS SECTION

- Syntax | 5544
- Hierarchy Level | 5544
- Description | 5544

Required Privilege Level | 5544
Release Information | 5545

Syntax

```
detection-time {
    threshold milliseconds;
}
```

Hierarchy Level

```
[edit protocols pim interface interface-name bfd-liveness-detection],
[edit routing-instances routing-instance-name protocols pim interface interface-name bfd-
liveness-detection]
```

Description

Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the `clear bfd adaptation` command to return BFD interval timers to their configured values. The `clear bfd adaptation` command is hitless, meaning that the command does not affect traffic flow on the routing device.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Support for BFD authentication introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD for PIM

bfd-liveness-detection (Protocols PIM) | 2945

threshold (PIM BFD Detection Time) | 29331

detection-time (BFD Liveness Detection)

IN THIS SECTION

- Syntax | 5545
- BGP | 5546
- EVPN, L2VPN, VPLS | 5546
- Description | 5547
- Options | 5548
- Required Privilege Level | 5548
- Release Information | 5548

Syntax

```
detection-time {
    threshold milliseconds;
}
```

BGP

```
[edit logical-systems name protocols bgp bfd-liveness-detection],
[edit logical-systems name protocols bgpgroup bfd-liveness-detection],
[edit logical-systems name protocols bgp group name neighbor address bfd-liveness-
detection],
[edit logical-systems name routing-instances name
protocols bgp bfd-liveness-detection],
[edit logical-systems name routing-instances name
protocols bgpgroup bfd-liveness-detection],
[edit logical-systems name routing-instances name
protocols bgpgroup neighbor address bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols
bgp bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols
bgpgroup bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols bgpgroup
neighbor address bfd-liveness-detection],
[edit protocols bgp bfd-liveness-detection],
[edit protocols bgp group bfd-liveness-detection],
[edit protocols bgp group neighbor address bfd-liveness-detection],
[edit routing-instances name protocols bgp bfd-liveness-detection],
[edit routing-instances name protocols bgp group bfd-liveness-detection],
[edit routing-instances name protocols bgp group neighbor address bfd-liveness-
detection],
[edit tenants name routing-instances name protocols bgp bfd-liveness-detection]
[edit tenants name routing-instances name protocols bgp group bfd-liveness-detection]
[edit tenants name routing-instances name protocols bgp groupneighbor address bfd-liveness-detection]
```

EVPN, L2VPN, VPLS

```
[edit logical-systems name routing-instances name protocols (evpn |
l2vpn | vpls) oam bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn |
l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn |
l2vpn | vpls)mesh-group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn |
vpls) oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn |
```

```
vpls)neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-
group mesh-group-name neighbor    neighbor-id oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls)              oam              bfd-
liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-
detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-name neighbor
neighbor-id oam bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls)
oam              bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id
oam              bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-
name neighbor    neighbor-id oam bfd-liveness-detection],
```

Description

Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the `clear bfd adaptation` command to return BFD interval timers to their configured values. The `clear bfd adaptation` command is hitless, meaning that the command does not affect traffic flow on the routing device.

Starting in Junos OS Release 20.3R1, we support distributed mode for BFD failure detection on the SRX5000 line of devices with SPC3 card. The distributed mode provides faster BFD failure detection of 300 (3 x 100) ms. You can enable distributed mode when you configure the BFD failure detection timer value less than 500 ms.

Starting in Junos OS Release 21.1R1, we support distributed mode for BFD on SRX1500, SRX4100, SRX4200, and SRX4600. This mode provides a faster BFD failure detection time of 3 x 300 ms.

For optimization and performance enhancement, you must configure the BFD failure detection timer value in multiples of 50 ms.

Starting in Junos OS Release 24.3R1, we've introduced distributed mode for BFD failure detection on vSRX 3.0. To enable distributed mode, configure the BFD failure detection time value as 500 ms. This feature is supported in stand-alone and L3-HA mode on vSRX 3.0.

With this support, we've also added the dedicated offload CPU feature on vSRX 3.0. The dedicated offload CPU feature reschedules a flow thread and uses the DPDK flow filters on the NIC to move high priority packets (BGP, RIPv2, OSPF, PIM, Multicast, IGMP, Single-Hop BFD, and Multihop BFD) onto the dedicated flow thread. With this feature, packets are processed by their own dedicated flow thread or queue, even in cases where the forwarding plane is oversubscribed and dropping packets.

To enable the dedicated offload CPU feature on vSRX 3.0, run the `set security forwarding-options dedicated-offload-cpu` command.

To view the Packet Forwarding Engine's current dedicated offload CPU status use the `show security forward-options dedicated-offload-cpu` command. This command displays whether the Packet Forwarding Engine has dedicated offload CPU feature enabled or disabled.

Options

threshold ***milliseconds***

Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

> **NOTE**: The threshold value must be equal to or greater than the transmit interval.
>
> The threshold time must be equal to or greater than the value specified in the `minimum-interval` or the `minimum-receive-interval` statement.

- **Range:** 1 through 255,000 ms

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 8.2.

Support for BFD authentication introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding How BFD Detects Network Failures

Configuring BFD for Layer 2 VPN and VPLS

Example: Configuring BFD for BGP

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

bfd-liveness-detection

detection-time (LAG)

IN THIS SECTION

- Syntax | 5549
- Hierarchy Level | 5549
- Description | 5549
- Options | 5550
- Required Privilege Level | 5550
- Release Information | 5550

Syntax

```
detection-time {
    threshold milliseconds;
}
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options bfd-liveness-detection]
```

Description

Configure BFD timers for aggregated Ethernet interfaces.

Options

threshold *milliseconds* — Specify the maximum time interval for detecting a BFD neighbor. If the transmit interval is greater than this value, the device triggers a trap.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

- authentication (LAG) | **2622**
- bfd-liveness-detection (LAG) | **2926**
- transmit-interval (LAG) | **30090**
- *Configuring Micro BFD Sessions for LAG*
- Example: Configuring Independent Micro BFD Sessions for LAG
- Understanding Independent Micro BFD Sessions for LAG

detection-time (Liveness Detection)

IN THIS SECTION

- Syntax | **5551**
- Hierarchy Level | **5551**
- Description | **5551**
- Required Privilege Level | **5551**
- Release Information | **5551**

Syntax

```
detection-time {
    threshold milliseconds;
}
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection]
```

Description

The Bidirectional Forwarding Detection (BFD) timers are adaptive and can be adjusted to be faster or slower.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

detector

IN THIS SECTION

- Syntax | 5552
- Hierarchy Level | 5552
- Description | 5552

- Options | 5552
- Required Privilege Level | 5552
- Release Information | 5552

Syntax

```
detector {
    protocol-name protocol-name {
        tunable-name  tunable-name {
            tunable-value protocol-value;
        }
    }
}
```

Hierarchy Level

```
[edit security idp sensor-configuration]
```

Description

Configure protocol detector engine for a specific service.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

deterministic-ar-policy

IN THIS SECTION

- Syntax | 5553
- Hierarchy Level | 5553
- Description | 5553
- Options | 5554
- Required Privilege Level | 5555
- Release Information | 5555

Syntax

```
deterministic-ar-policy [ deterministic-ar-policy-name ... ];
```

Hierarchy Level

```
[edit routing-instances name protocols evpn assisted-replication leaf]
```

Description

Assign routing policies to an EVPN instance on an assisted replication (AR) leaf device to deterministically steer multicast flows to specific preferred AR replicator devices. With this feature, you can control how the leaf device load-balances multicast flows among the available AR replicator devices, rather than using the default load-balancing method on the device.

> **NOTE**: AR is a multicast traffic flow optimization feature where an ingress leaf device sends one copy of a multicast stream to another device called an AR replicator in the EVPN network. AR replicators are devices that can efficiently handle replicating and forwarding multicast streams to the other multicast receivers in the network. See *Assisted Replication Multicast Optimization in EVPN Networks* for details on how AR works and how to configure AR leaf and AR replicator devices.

We support setting deterministic AR replicator policies only in EVPN-VXLAN networks running AR with optimized intersubnet multicast (OISM).

You can define and assign AR replicator routing policies as follows:

- You can include match conditions for IPv4 or IPv6 multicast traffic with IGMP or MLD multicast groups.

  With IGMPv2 (IPv4) or MLDv1 (IPv6) multicast traffic, you can match a multicast flow for a multicast group and/or a bridge domain (a VLAN-mapped virtual network identifier [VNI]).

  With IGMPv3 (IPv4) or MLDv2 (IPv6 traffic), you can match a multicast flow for any combination of the following:

  - A multicast group
  - A bridge domain (a VLAN-mapped VNI)
  - A multicast source address, because these are source-specific multicast (SSM) protocols.

- You can configure the policy in either of the following modes—these options are mutually exclusive:

  - Strict mode—Include the `strict` option to strictly direct matching flows only to the preferred AR replicator. If that AR replicator goes down, the AR leaf device drops any matching flows even if the device detects other available AR replicators in the network. When the preferred AR replicator comes back up again, the AR leaf resumes sending matching traffic to that AR replicator.

  - Loose mode—The default mode. You can optionally include a fallback AR replicator to use if the preferred AR replicator goes down. If the fallback AR replicator is also down or you don't specify a fallback, but the AR leaf device detects other AR replicators in the network, the device uses its default AR replicator load balancing method to distribute flows among the available other AR replicators. (See *Default AR Leaf Load-Balancing When Detecting Multiple AR Replicators*.)

    If the AR leaf device doesn't detect any available AR replicators in the network, the device defaults to ingress replication (replicates and forwards the multicast flows itself to the other devices with receivers interested in those flows).

Options

deterministic-ar-policy-name

Routing policy name that deterministically directs matching multicast flows to specified AR replicators.

You set deterministic AR routing policy conditions to match multicast flows in the `from` clause of the policy statement at the `[edit policy-options policy-statement deterministic-ar-policy-name]` hierarchy.

You set the policy action in the `then` clause of the policy using the `assisted-replication replicator-ip` *`replicator-ip-addr`* statement at the `[edit policy-options policy-statement` *`deterministic-ar-policy-name`* `then]` hierarchy.

The policy action must include the `replicator-ip` statement to specify a preferred AR replicator for this feature. Include the `strict` option if you don't want the matching flows to go to any AR replicator other than the preferred one. Or, you can optionally include a fallback AR replicator to which to send the matching flows if the specified AR replicator is down (use the `fallback-replicator-ip` option).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.4R2.

RELATED DOCUMENTATION

Deterministic Load Balancing and Traffic Steering to AR Replicators

Routing Policy Match Conditions

Actions in Routing Policy Terms

assisted-replication (Deterministic AR Replicator Policy Actions) | **2507**

deterministic-nat-configuration-log-interval (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax (MX Series Devices) | **5556**
- Hierarchy Level | **5556**
- Description | **5556**

- Options | 5556
- Required Privilege Level | 5556
- Release Information | 5556

Syntax (MX Series Devices)

```
deterministic-nat-configuration-log-interval seconds;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Description

Configure the interval at which the syslog is generated for the deterministic NAT configuration. (Deterministic NAPT for Next Gen Services is available only for MX series devices.)

Options

deterministic-nat-configuration-log-interval *seconds*

Number of seconds in the interval.

- **Range:** 1800 through 86400
- **Default:** 1800

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Deterministic NAPT for Next Gen Services

deterministic-port-block-allocation

IN THIS SECTION

- Syntax | 5557
- Hierarchy Level | 5557
- Description | 5557
- Options | 5558
- Required Privilege Level | 5558
- Release Information | 5558

Syntax

```
deterministic-port-block-allocation {
    block-size block-size;
    include-boundary-addresses;
}
```

Hierarchy Level

```
[edit services nat pool pool-name port]
```

Description

Configure algorithm-based allocation of blocks of destination ports. By specifying this method, you ensure that an incoming (source) IP address and port always map to the same destination IP address and port block, thus eliminating the need for logging address translations.

Options

block-size—Maximum number of ports that can be allocated to a user.

If `block-size` is configured as zero, the method for computing the block size is as follows:

block-size = int(65412/ceil[(Number of subscribers/Number of IP addresses in the NAT pool)]).

where

64512 is derived from (65535 - 1023) because the regular port assignments start from 1024.

Number of subscribers is derived from the from clause of the applicable NAT rule.

- **Default:** 256
- **Range:** 0 through 32,000

`include-boundary-addresses`—(Optional) Specifies that the lowest and highest addresses (the network and broadcast addresses) in the source address range of a NAT rule should be translated when the NAT pool is used. If the source address has a prefix of /32, the lowest and highest address are automatically translated.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Deterministic NAPT

deterministic (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 5559
- Hierarchy Level | 5559
- Description | 5559
- Options | 5560
- Required Privilege Level | 5560
- Release Information | 5560

Syntax

```
deterministic {
    block-size block-size;
    host {
        address address;
    }
    include-boundary-addresses;
}
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Description

Configure deterministic NAT to ensure that the original internal source IPv4 or IPv6 address and port always map to the same post-NAT IPv4 address and block of ports. In addition, the reverse mapping of a given translated external IPv4 address and port are always mapped to the same internal IP address.

This eliminates the need for address translation logging.

Options

block-size *block-size*

The number of ports in the port block.

- **Range:** 1 to 64,512
- **Default:** 256

host address *address*

The first usable pre-NAT subscriber address, which is used to perform the deterministic NAT mapping.

include-boundary-addresses

Include the translation of the lowest and highest IPv4 addresses (the network and broadcast addresses) in the source address range of a NAT rule. This does not apply to IPv6 source addresses.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Deterministic NAPT for Next Gen Services

device-attribute (Easy EVPN LAG Configuration)

IN THIS SECTION

- Syntax | 5561
- Hierarchy Level | 5561
- Description | 5561
- Options | 5562

Required Privilege Level | 5563
Release Information | 5563

Syntax

```
device-attribute {
  loopback {
    peer1-subnet peer1-subnet;
    peer2-subnet peer2-subnet;
    }
  peer-id peer-id;
  peer-to-peer {
    overlay-connectivity {
      ibgp;
    }
    peer-subnet {
     inet inet;
     inet6 inet6;
     interface-name [ interface-name ... ];
    }
    underlay-connectivity {
      ospf;
    }
  }
  system-id system-id;
}
```

Hierarchy Level

```
[edit services evpn]
```

Description

Specify device-specific parameters as part of a generated Ethernet VPN (EVPN) fabric configuration.

You use the statements in the `[edit services evpn]` hierarchy level to simplify configuring an EVPN fabric with Ethernet segment identifier (ESI) link aggregation groups (LAGs) in the topology. The feature you enable with statements in this hierarchy is an easy EVPN LAG configuration feature, also called EZ-LAG.

You only need to configure a few statements to provide the parameters required for an EVPN core fabric with multihomed or single-homed customer edge (CE) devices (or end server devices). When you commit the `[edit services evpn]` configuration statements, the device uses a commit script to generate the corresponding standard EVPN fabric configuration statements.

The statements at the `device-attribute` hierarchy level provide parameters for configuration elements that are specific to:

- The peer provider edge (PE) devices in the EVPN fabric.
- Setting up LACP on the ESI LAG interfaces to multihomed customer edge (CE) devices (or end servers).

Options

loopback (peer1-subnet *peer1-subnet* peer2-subnet *peer2-subnet*)

Specify the device loopback IPv4 subnet address for each of the peer PE devices. The `[edit services evpn]` commit script uses the *peer-id* of the device on which it runs to derive the statements in the generated configuration for:

- The router ID
- The device overlay local address and peer neighbor address

> **NOTE**: The commit process doesn't do a commit check to enforce setting the same loopback subnet addresses across both devices. As a result, you should make sure to set the easy EVPN LAG configuration loopback *peer1-subnet* address to the same value on both peer PE devices, and likewise for the *peer2-subnet* address.

peer-id *peer-id*

Assign a unique identifier to the PE device among its peer PE devices in the EVPN fabric. The `[edit services evpn]` commit script uses this value to derive the statements to set values in the generated configuration that must be unique across the peer PE devices.

- **Values for *peer-id*:** 1 or 2

system-id *system-id*

Specify the PE device's system ID, which is in the format of a MAC address. The `[edit services evpn]` commit script uses this value to derive statements to set up Link Aggregation Control Protocol (LACP) on the ESI LAG interfaces.

To configure LACP on an aggregated Ethernet interface in an ESI LAG, the commit script adds this *system-id* to the interface ESI parameter. The commit script uses that combined value as the LACP system-id in the generated LACP configuration statement. For example, if *system-id* is 10:11:12:13:14:10 and the interface ESI parameter is 1, the commit script generates the following statement in the configuration:

```
set interfaces ae3 aggregated-ether-options lacp system-id 10:11:12:13:14:11
```

> **NOTE**: You set an aggregated Ethernet interface's ESI parameter using the `esi-lag-id esi-lag-id` statement at the `[edit services evpn evpn-vxlan server server-name]` hierarchy level.

The remaining statements are explained in detail on separate pages.

You can also consult CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Easy EVPN LAG (EZ-LAG) Configuration

device-authentication

IN THIS SECTION

Syntax | 5564

- Hierarchy Level | 5564
- Description | 5564
- Options | 5564
- Required Privilege Level | 5565
- Release Information | 5565

Syntax

```
device-authentication {
    (encrypted-password "password" | plain-text-password);
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the authentication password used when accessing individual QFabric system components with the `request component login` command.

> **NOTE**: Configuring this statement overrides the password for individual QFabric system components that was set during the initial QFabric system setup procedure.

Options

encrypted-password "*password*"— Specify the MD5 or other encrypted authentication password. You can specify only one encrypted password.

You cannot configure a blank password for the **encrypted-password** option using blank quotation marks (" "). You must configure a password of 1 through 128 characters and enclose the password in quotation marks.

plain-text-password—Plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring QFabric System Login Classes

Performing the QFabric System Initial Setup on a QFX3100 Director Group

device-count

IN THIS SECTION

- Syntax (EX, NFX, QFX Series, EX4600, OCX1100, QFabric System) | **5566**
- Hierarchy Level (EX Series) | **5566**
- Hierarchy Level (EX, NFX, QFX Series, EX4600, OCX1100, QFabric System) | **5566**
- Description | **5566**
- Options | **5566**
- Required Privilege Level | **5566**
- Release Information | **5567**

Syntax (EX, NFX, QFX Series, EX4600, OCX1100, QFabric System)

```
device-count number;
```

Hierarchy Level (EX Series)

```
[edit chassis aggregated-devices ethernet (Aggregated Devices)]
```

Hierarchy Level (EX, NFX, QFX Series, EX4600, OCX1100, QFabric System)

```
[edit chassis aggregated-devices ethernet],
[edit chassis node-group name aggregated-devices ethernet]
```

Description

Configure the number of aggregated Ethernet logical devices available to the switch.

Options

number—Maximum number of aggregated Ethernet logical interfaces on the switch.

- **Range:** 1 through 32 for EX2200, EX3200, and standalone EX3300 switches and for EX3300 Virtual Chassis
- **Range:** 1 through 64 for standalone EX4200, standalone EX4500, and EX6200 switches and for EX4200 and EX4500 Virtual Chassis
- **Range:** 1 through 239 for EX8200 Virtual Chassis
- **Range:** 1 through 255 for standalone EX8200 switches
- **Range:** 1 through 480 for standalone EX9200 switches

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Range updated in Junos OS Release 9.5 for EX Series switches.

RELATED DOCUMENTATION

- *Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion*
- *Configuring Link Aggregation*
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- Configuring Aggregated Ethernet Links (CLI Procedure)
- Junos OS Network Interfaces Configuration Guide

device-count (MX Series)

IN THIS SECTION

- Syntax | 5567
- Hierarchy Level | 5568
- Description | 5568
- Options | 5568
- Required Privilege Level | 5568
- Release Information | 5568

Syntax

```
device-count number;
```

Hierarchy Level

```
[edit chassis aggregated-devices ethernet]
[edit chassis aggregated-devices sonet]
```

Description

Configure the number of aggregated logical devices available to the router.

Starting in Junos release 14.2, for MX series routers, aggregated Ethernet interfaces created under a logical system can be individually named. Prior to 14.2, ae interfaces were named automatically (AE1, AE2) etc. upon setting the device count. This change allows administrators to use custom naming schemes. System resources are only allocated for named ae interfaces, regardless of how many were declared in the device count. (In Junos 14.2 and earlier, ae naming occurred automatically up to the number specified for device count, and system resources were allocated whether a given ae interface was used or not.)

Options

number—Set the number of aggregated logical devices that will be available for configuration.

> **NOTE**: Starting with Junos OS Release 13.2, a maximum of 64 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces. In releases before Junos OS Release 13.2, a maximum of 16 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces.
>
> For Junos OS Evolved, you can specify up to 512 aggregated Ethernet devices.

- **Range:** 1 - 1000. The upper limit for this value is system specific.
- **Range:** 1 - 1052 for Junos OS Evolved.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement functionality updated in Junos OS Release 14.2, as described below.

RELATED DOCUMENTATION

- Configuring Junos OS for Supporting Aggregated Devices
- Configuring Aggregated SONET/SDH Interfaces

device-count (Pseudowire Subscriber Interfaces)

IN THIS SECTION

- Syntax | 5569
- Hierarchy Level | 5569
- Description | 5569
- Options | 5570
- Required Privilege Level | 5570
- Release Information | 5570

Syntax

```
device-count number;
```

Hierarchy Level

```
[edit chassis pseudowire-service]
```

Description

Configure the number of pseudowire logical devices available to the router. The statement also defines the available interface names for the pseudowire interfaces.

> **NOTE**: When you subsequently configure the pseudowire interfaces, you must specify the interface names in the range from ps0 up to ps*(device-count - 1)*. For example, if you set the maximum number of devices to 5, then you can only configure interfaces ps0, ps1, ps2, ps3, and ps4. If you specify an interface name outside that range, the pseudowire interface is not created.

Options

number Number of devices.

- **Range:** 1 through 7000, 1 through 18000 for MX2010 and MX2020 routers with the MX2K-MPC9E or MX2K-MPC11E line card

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

Increased upper limit in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

devices

IN THIS SECTION

- Syntax | 5571
- Hierarchy Level | 5571
- Description | 5571
- Default | 5571
- Options | 5572
- Required Privilege Level | 5572
- Release Information | 5572

Syntax

```
devices device-names;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Specifies one of the virtual tunnel (VT) interfaces to de-encapsulate the egress traffic for ultimate-hop popping on point-to-multipoint LSPs. If no device is specified, the selection process is performed automatically.

Default

The device selection process is performed automatically if no device is configured. Junos OS selects one of the available VT interfaces to de-encapsulate the egress traffic.

Options

device-names—Specify which VT interfaces are used to handle the RSVP traffic.

- **Range:** 0 to 8 devices

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs

Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs

df-election

IN THIS SECTION

- Syntax | 5573
- Hierarchy Level | 5573
- Description | 5573
- Required Privilege Level | 5573
- Release Information | 5573

Syntax

```
df-election {
    backoff-period milliseconds;
    offer-period milliseconds;
    robustness-count number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name
bidirectional],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name bidirectional],
[edit protocols pim interface (Protocols PIM)  interface-name bidirectional],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name bidirectional]
```

Description

Optionally, configure the designated forwarder (DF) election parameters for bidirectional PIM.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Bidirectional PIM

Example: Configuring Bidirectional PIM

df-election-granularity

IN THIS SECTION

- Syntax | 5574
- Hierarchy Level | 5574
- Description | 5574
- Options | 5575
- Required Privilege Level | 5575
- Release Information | 5575

Syntax

```
df-election-granularity {
    per-esi {
        lacp-oos-on-ndf;
    }
    per-esi-vlan;
    per-evi;
}
```

Hierarchy Level

```
[edit interfaces name esi],
[edit interfaces name unit number esi]
```

Description

Configure the designated forwarder (DF) election granularity per member link for an aggregated Ethernet interface in an EVPN multihoming configuration.

The `per-esi lacp-oos-on-ndf` setting enables EVPN multihoming when:

- The peer provider edge (PE) devices in the multihomed Ethernet segment (ES) operate in active-standby (single-active) mode.

- The connection to the multihomed CE device is an aggregated Ethernet (AE) link with member links that connect to more than one of the peer PE devices.

When an aggregated Ethernet interface connects to the DF and you enable this option with single-active multihoming, the aggregated Ethernet interface emulates the behavior of an MC-LAG configuration (without requiring you to configure an ICCP or ICL interface). This behavior is different from in a standard EVPN multihoming configuration, where the logical interfaces you configure on an aggregated Ethernet interface can take different designated forwarder election roles.

In an EVPN single-active configuration, the provider edge (PE) device that is not the DF (the non-DF) sends LACP out-of-sync packets to the CE device. As a result, LACP goes down on the CE device, and the CE device doesn't send traffic on the links connected to the non-DF.

If the connection between a multihomed CE device and the DF PE device fails, the DF election process selects a new DF. If the connection between the CE device and a non-DF PE device fails, the current DF assignment remains the same.

NOTE: You can't configure both the `per-esi` option and the `per-esi-vlan` option with this statement at the same time. If you configure one of these statements when the other is already set, the CLI replaces the existing setting with the new one you've configured.

Options

The options are explained separately. You can click the linked statements in the Syntax section or search for a statement in CLI Explorer.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

MC-LAG Concepts

df-election-type

IN THIS SECTION

- Syntax | 5576
- Hierarchy Level | 5576
- Description | 5577
- Options | 5577
- Required Privilege Level | 5577
- Release Information | 5577

Syntax

```
df-election-type {
    mod;
    preference {
      least;
      non-revertive;
      value value;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name esi],
[edit dynamic-profiles name logical-systems name interfaces name esi],
[edit interfaces name esi],
[edit interfaces name unit number esi],
[edit protocols evpn interconnect esi],
[edit routing-instances name protocols evpn interconnect esi]
```

Description

Choose how the device elects the designated forwarder (DF) for an Ethernet segment among the multihoming peer provider edge (PE) devices in an Ethernet VPN (EVPN) network.

> **NOTE**: If you configure the "df-election-granularity" on page 5574 "per-evi" on page 12409 option, you must also configure the `df-election-type mod` option without configuring a `df-election-type preference` option to prevent an error when committing the configuration.

Options

mod Use MOD-based DF election.

The device uses MOD-based DF election by default if you don't configure the `df-election-type` statement at all.

See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3 for MX Series Routers with MPC and MIC interfaces.

Support added in Junos OS Release 20.3R1 at the `[edit protocols evpn interconnect esi]` and `[edit routing-instances *routing-instance-name* protocols evpn interconnect esi]` hierarchy levels.

RELATED DOCUMENTATION

evpn | 6541

interconnect | 8797

dh-group (services)

IN THIS SECTION

- Syntax | 5578
- Hierarchy Level | 5578
- Description | 5578
- Options | 5578
- Required Privilege Level | 5579
- Release Information | 5579

Syntax

```
dh-group (group1 | group2 | group5 |group14 | group15 | group16 | group19 | group20 | group24);
```

Hierarchy Level

```
[edit services ipsec-vpn ike proposal proposal-name]
```

Description

Configure the IKE Diffie-Hellman prime modulus group to use for performing the new Diffie-Hellman exchange.

Options

`group1`—768-bit.

`group2`—1024-bit.

`group5`—1536-bit.

`group14`—2048-bit.

`group15`—3072-bit.

`group16`—4096-bit.

`group19`—256-bit random Elliptic Curve Group.

`group20`—384-bit random Elliptic Curve Group.

`group24`—2048-bit with 256-bit Prime Order Subgroup.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`group15`, `group16`, and `group24` options added in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Configuring IKE Proposals

dh-group

IN THIS SECTION

- Syntax | 5580
- Hierarchy Level | 5580
- Description | 5580
- Options | 5580
- Required Privilege Level | 5580
- Release Information | 5580

Syntax

```
dh-group (group1 | group2 | group5 | group14);
```

Hierarchy Level

```
[edit security ike proposal ike-proposal-name]
```

Description

Configure the IKE Diffie-Hellman group.

Options

`dh-group`—Type of Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange. It can be one of the following:

- `group1`—768-bit.
- `group2`—1024-bit.
- `group5`—1536-bit.
- `group14`—2048-bit.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Diffie-Hellman Group for an IKE Proposal

dh-group (Security IKE)

IN THIS SECTION

- Syntax | 5581
- Hierarchy Level | 5581
- Description | 5581
- Options | 5581
- Required Privilege Level | 5582
- Release Information | 5582

Syntax

```
dh-group (group1 | group2 | group5 | group14 | group15 | group16 | group19 | group20 | group21 |
group24);
```

Hierarchy Level

```
[edit security ike proposal proposal-name]
```

Description

Specify the IKE Diffie-Hellman group. The device does not delete existing IPsec SAs when you update the `dh-group` configuration in the IKE proposal.

Options

`dh-group`—Diffie-Hellman group for key establishment.

- `group1`—768-bit Modular Exponential (MODP) algorithm.
- `group2`—1024-bit MODP algorithm.
- `group5`—1536-bit MODP algorithm.

- `group14`—2048-bit MODP group.
- `group15`—3072-bit MODP algorithm.
- `group16`—4096-bit MODP algorithm.
- `group19`—256-bit random Elliptic Curve Groups modulo a Prime (ECP groups) algorithm.
- `group20`—384-bit random ECP groups algorithm.
- `group21`—521-bit random ECP groups algorithm.
- `group24`—2048-bit MODP Group with 256-bit prime order subgroup.

We recommend that you use `group14`, `group15`, `group16`, `group19`, `group20`, or `group21` instead of `group1`, `group2`, or `group5`.

We support `group15`, `group16`, and `group21` options only with iked process when `junos-ike` package is installed.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Support for the `group14` option added in Junos OS Release 11.1.

Support for `group19`, `group20`, and `group24` options added in Junos OS Release 12.1X45-D10.

Support for `group19` and `group20` options added in Junos OS Release 15.1X49-D70 for vSRX Virtual Firewall.

Support for `group15`, `group16`, and `group21` options added in Junos OS Release 19.1R1 on SRX5000 line of devices with `junos-ike` package installed.

Starting in Junos OS Release 20.2R1, we've changed the help text description as `NOT RECOMMENDED` for the CLI options `group1`, `group2`, and `group5` for devices running IKED with `junos-ike` package installed.

Support for `group15`, `group16`, and `group21` options added in Junos OS Release 20.3R1 on vSRX Virtual Firewall instances with `junos-ike` package installed.

Support for `group15`, `group16`, and `group21` options added in Junos OS Release 21.1R1 on vSRX Virtual Firewall 3.0 instances with `junos-ike` package installed.

RELATED DOCUMENTATION

IPsec Overview

proposal (Security IKE) | 13389

Installing Junos IKE package

dhcp

IN THIS SECTION

- Syntax | 5583
- Hierarchy Level | 5584
- Description | 5584
- Required Privilege Level | 5584
- Release Information | 5584

Syntax

```
dhcp {
    boot-file filename;
    boot-server (address | hostname);
    default-lease-time seconds;
    domain-name domain-name;
    domain-search [domain-list];
    maximum-lease-time seconds;
    name-server {
        address;
    }
    next-servernext-server
    option option-identifier-code ;
    pool address/prefix-length {
        address-range {
            low address;
            high address;
        }
```

```
        exclude-address {
            address;
        }
    }
    router {
        address;
    }
    static-binding mac-address {
        fixed-address {
            address;
        }
        host-name hostname;
        client-identifier (ascii client-id | hexadecimal client-id);
    }
    wins-server {
        address;
    }
}
```

Hierarchy Level

```
[edit system services]
[edit logical-systems logical -system-name system services]
```

Description

Configure a router, or a switch as a DHCP server. A DHCP server can allocate network addresses and deliver configuration information to client hosts on a TCP/IP network.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The `logical-systems` option is introduced in Junos OS Release 18.4R1.

dhcp-attributes (Access IPv4 Address Pools)

IN THIS SECTION

- Syntax | 5585
- Hierarchy Level | 5586
- Description | 5586
- Required Privilege Level | 5587
- Release Information | 5587

Syntax

```
dhcp-attributes {
    boot-file boot-file-name;
    boot-server boot-server-name;
    domain-name domain-name;
    grace-period seconds;
    maximum-lease-time (seconds | infinite);
    name-server ipv4-address;
    netbios-node-type (b-node | h-node | m-node | p-node);
    next-server next-server-name;
    option dhcp-option-identifier-code {
        array {
            byte [8-bit-value];
            flag [ false| off |on |true];
            integer [32-bit-numeric-values];
            ip-address [ip-address];
            short [signed-16-bit-numeric-value];
            string [character string value];
            unsigned-integer [unsigned-32-bit-numeric-value];
            unsigned-short [16-bit-numeric-value];
        }
        byte 8-bit-value;
```

```
        flag (false | off | on | true);
        integer 32-bit-numeric-values;
        ip-address ip-address;
        short signed-16-bit-numeric-value;
        string character string value;
        unsigned-integer unsigned-32-bit-numeric-value;
        unsigned-short 16-bit-numeric-value;
    }
    option-match {
        option-82 {
            circuit-id match-value {
                range range-name;
            }
            remote-id match-value;
                range range-name;
            }
        }
    }
    propagate-ppp-settings [interface-name];
    propagate-settings interface-name;
    router ipv4-address;
    server-identifier ip-address;
    sip-server {
        ip-address ipv4-address;
        name sip-server-name;
    }
    tftp-server server-name;
    wins-server ipv4-address;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet]
```

Description

Configure attributes for IPv4 address pools that can be used by different clients. The DHCP attributes for this statement uses standard IPv4 DHCP options.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

DHCP Overview

dhcp-attributes (Access IPv6 Address Pools)

IN THIS SECTION

- Syntax | 5587
- Hierarchy Level | 5588
- Description | 5588
- Options | 5588
- Required Privilege Level | 5589
- Release Information | 5589

Syntax

```
dhcp-attributes {
    dns-server ipv6-address;
    grace-period seconds;
    maximum-lease-time (seconds | infinite);
    option dhcp-option-identifier-code {
        array {
            byte [8-bit-value];
```

```
            flag [ false| off |on |true];
            integer [32-bit-numeric-values];
            ip-address [ip-address];
            short [signed-16-bit-numeric-value];
            string [character string value];
            unsigned-integer [unsigned-32-bit-numeric-value];
            unsigned-short [16-bit-numeric-value];
        }
        byte 8-bit-value;
        flag (false | off | on | true);
        integer 32-bit-numeric-values;
        ip-address ip-address;
        short signed-16-bit-numeric-value;
        string character string value;
        unsigned-integer unsigned-32-bit-numeric-value;
        unsigned-short 16-bit-numeric-value;
    }
    propagate-ppp-settings [interface-name];
    sip-server-address ipv6-address;
    sip-server-domain-name domain-name;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet6]
```

Description

Configure attributes for address pools that can be used by different clients.

Options

- `dns-server` *IPv6-address*—Specify a DNS server to which clients can send DNS queries.
- `grace-period` *seconds* —Specify the grace period offered with the lease.
- **Range:** 0 through 4,294,967,295 seconds
- **Default:** 0 (no grace period)

- `maximum-lease-time` *`seconds`*—Specify the maximum length of time in seconds for which a client can request and hold a lease on a DHCP server.
- **Range:** 30 through 4,294,967,295 seconds
- **Default:** 86,400 seconds (24 hours)
- `option` *`dhcp-option-identifier-code`*—Specify the DHCP option identifier code.
- `propagate-ppp-settings [`*`interface-name`*—Specify PPP interface name for propagating DNS or WINS settings.
- `sip-server-address` *`IPv6-address`*—Specify the IPv6 address of the SIP outbound proxy server.
- `sip-server-domain-name` *`domain-name`*—Specify the domain name of the SIP outbound proxy server.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

DHCP Overview

dhcp-attributes (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 5590
- Hierarchy Level | 5591
- Description | 5591
- Options | 5591

Required Privilege Level | 5595
Release Information | 5596

Syntax

```
dhcp-attributes {
    boot-file filename;
    boot-server (address | hostname);
    dns-server [ ipv6-address ];
    domain-name domain-name;
    exclude-prefix-len exclude-prefix-length;
    grace-period seconds;
    maximum-lease-time seconds;
    name-server [ server-list ];
    netbios-node-type node-type;
    option {
        [ (id-number option-type option-value)
        (id-number array option-type option-value) ];
    }
    option-match {
        option-82 {
            circuit-id value range named-range;
            remote-id value range named-range;
        }
    }
    preferred-lifetime seconds;
    router [ router-address ];
    server-identifier ip4-address;
    sip-server-address [ ipv6-address ];
    sip-server-domain-name domain-name;
    t1-percentage percentage;
    t1-renewal-time;
    t2-percentage percentage;
    t2-rebinding-time;
    tftp-server address;
    valid-lifetime seconds;
    wins-server [ servers ];
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family family]
```

Description

Configure DHCP attributes for the protocol family in a specific address pool. The attributes determine options and behaviors for the DHCP clients.

Options

boot-file

Set the boot file advertised to DHCP clients. After the client receives an IP address and the boot file location from the DHCP server, the client uses the boot image stored in the boot file to complete DHCP setup. This configuration is equivalent to DHCP option 67.

- **Values:** *filename*—Location of the boot file on the boot server. The filename can include a pathname.

boot-server

Configure the name of the boot server advertised to DHCP clients. The client uses a boot file located on the boot server to complete DHCP setup. This configuration is equivalent to DHCP option 66.

- Values:
  - *address*—IPv4 address of a boot server.
  - *hostname*—Fully qualified hostname of a boot server.

dns-server

Specify a DNS server to which clients can send DNS queries. This is equivalent to DHCPv6 option 23. To specify multiple DNS servers, add multiple `dns-server` statements in order of preference.

- **Values:** *ipv6-address*—IPv6 address of a DNS server.

domain-name

Configure the name of the domain in which clients search for a DHCP server host. This is the default domain name that is appended to hostnames that are not fully qualified. This is equivalent to DHCP option 15.

- **Values:** *domain-name*—Name of the domain.

exclude-prefix-len *exclude-prefix-length*

Specify the length of the IPv6 prefix to be excluded from the delegated prefix. Range: 1 through 128.

grace-period

Configure the amount of time that the client retains the address lease after the lease expires. The address cannot be reassigned to another client during the grace period.

- **Values:** *seconds*—Number of seconds the lease is retained.
- **Range:** 0 through 4,294,967,295 seconds.
- **Default:** 0 (no grace period).

maximum-lease-time

Specify the maximum length of time, in seconds, that the lease is held for a client if the client does not renew the lease. This is equivalent to DHCP option 51. The `maximum-lease-time` is mutually exclusive with both the `preferred-lifetime` and the `valid-lifetime`, and cannot be configured with either timer.

- **Values:** *seconds*—Maximum number of seconds the lease can be held.
- **Range:** 30 through 4,294,967,295 seconds.
- **Default:** 86,400 (24 hours).

name-server

Configure one or more Domain Name System (DNS) name servers available to the client to resolve hostname-to-client mappings. This is equivalent to DHCP option 6.

- **Values:** *server-names*—IP addresses of the domain name servers, listed in order of preference.

netbios-node-type

Specify the NetBIOS node type. This is equivalent to DHCP option 46.

- **Values:** *node-type*—One of the following node types:
  - b-node—Broadcast node.
  - h-node—Hybrid node.
  - m-node—Mixed node.
  - p-node—Peer-to-peer node.

option

Specify user-defined options that are added to client packets. Starting in Junos OS Release 13.3, the `hex-string` option type was introduced.

- Values:
  - *array*—An option can include an array of option types.

- *id-number*—Any whole number. The ID number is used to index the option and must be unique across a DHCP server.
- *option-type*—Any of the following types: byte, byte-stream, flag, hex-string, integer, ip-address, short, string, unsigned-integer, or unsigned-short.
- *option-value*—Value associated with an option. The option value must be compatible with the option type (for example, an On or Off value for a flag type).

preferred-lifetime

Specify the length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix active. When the lifetime expires, the address is deprecated. If the `valid-lifetime` is also configured, the `preferred-lifetime` must be less than the `valid-lifetime`. The `preferred-lifetime` and the `maximum-lease-time` are mutually exclusive and cannot both be configured.

- **Values:** *seconds*—Number of seconds that the IPv6 prefix is active.
- **Range:** 30 through 4,294,967,295 seconds.
- **Default:** 86,400 (24 hours).

router

Specify one or more routers located on the client's subnet. This statement is the equivalent of DHCP option 3.

- **Values:** *router-address*—IP address of one or more routers.

server-identifier

Specify the IP address that is used as the source address the DHCP server includes in IP packets when communicating with clients. The address is included in the DHCP packet in option 54.

- **Values:** *ipv4-address*—IP address.

sip-server-address

Specify a SIP outbound proxy server that DHCPv6 local server clients can use. This is equivalent to DHCPv6 option 22. To specify multiple servers, add multiple `sip-server-address` statements in order of preference.

- **Values:** *`ipv6-address`*—IPv6 address of a SIP outbound proxy server.

sip-server-domain-name

Configure the domain name of the SIP outbound proxy server that DHCPv6 local server clients can use. This is equivalent to DHCPv6 option 21.

- **Values:** *domain-name*—Name of the domain.

t1-percentage

Specify a percentage of the `preferred-lifetime` value. After this percentage of the `preferred-lifetime` value elapses, the DHCPv4 or DHCPv6 client requests an extension on its lease from the originating DHCP local server. The `t1-percentage` is also referred to as

the renewal time. The `t1-percentage` value must be less than the `t2-percentage` value. DHCPv4 server support was added in Junos OS Release 17.2.

- **Values:** *percentage*—Percentage of the `preferred-lifetime` value.
- **Range:** 0 through 100.
- **Default:** If the `t1-percentage` value is not configured, the default is based on the `preferred-lifetime` value:
  - If the `preferred-lifetime` value is finite, the default is 50 percent of the `preferred-lifetime` value.
  - If the `preferred-lifetime` value is infinite, the default is also infinite.

t1-renewal-time

Specify the time (T1) at which the DHCPv4 or DHCPv6 client requests an extension (renewal) of the existing lease. This time is expressed as the number of seconds since the beginning of the lease. Using this statement to configure a duration in seconds is an alternative to using the `t1-percentage` statement.

- **Values:** *seconds*—Number of seconds.
- **Range:** 30 through 4,294,967,295 seconds.
- **Default:** 50 percent of the lease duration (`preferred-lifetime`).

t2-percentage

Specify a percentage of the `preferred-lifetime` value. After this percentage of the `preferred-lifetime` value elapses, the DHCPv4 or DHCPv6 client requests an extension on its lease from any available DHCPv4 or DHCPv6 server. The `t2-percentage` is also referred to as the rebinding time. The `t2-percentage` value must be greater than the `t1-percentage` value. DHCPv4 server support was added in Junos OS Release 17.2.

- **Values:** *percentage*—Percentage of the `preferred-lifetime` value.
- **Range:** 0 through 100.
- **Default:** Default: If the `t2-percentage` value is not configured, the default is based on the `preferred-lifetime` value:
  - If the `preferred-lifetime` value is finite, the default is 80 percent of the `preferred-lifetime` value.
  - When the `preferred-lifetime` value is infinite, the default is also infinite.

t2-rebinding-time

Specify the time (T2) at which the DHCPv4 or DHCPv6 client attempts to contact any DHCP server to request an extension (rebinding) of the existing lease. This time is

expressed as the number of seconds since the beginning of the lease. Using this statement to configure a duration in seconds is an alternative to using the `t2-percentage` statement.

- **Values:** *seconds*—Number of seconds.
- **Range:** 30 through 4,294,967,295 seconds.
- **Default:** The default value depends on the client:
  - (DHCPv4 clients) 87.5 percent of the lease duration (`preferred-lifetime`).
  - (DHCPv6 clients) 80 percent of the lease duration (`preferred-lifetime`).

tftp-server Specify the Trivial File Transfer Protocol (TFTP) server that the client uses to obtain the client configuration file. This is equivalent to DHCP option 150.

- **Values:** *ip-address*—IP address of the TFTP server.

valid-lifetime Specify the length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix valid. When the lifetime expires, the address becomes invalid. If the `preferred-lifetime` is also configured, the `valid-lifetime` must be greater than the `preferred-lifetime`. The `valid-lifetime` and the `maximum-lease-time` are mutually exclusive and cannot both be configured.

- **Values:** *seconds*—Number of seconds that the IPv6 prefix is valid.
- **Range:** 30 through 4,294,967,295 seconds.
- **Default:** 86,400 (24 hours).

wins-server Specify one or more NetBIOS name servers (NBNS) that the client uses to resolve NetBIOS names. This is equivalent to DHCP option 44.

- **Values:** *ipv4-address*—IP address of each NetBIOS name server. Add them to the configuration in order of preference.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`exclude-prefix-len` statement introduced in Junos OS Release 17.3 for MX Series.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management

DHCP Client Attribute and Address Assignment

DHCP Lease Times for IP Addresses

dhcp-client

IN THIS SECTION

- Syntax | 5596
- Hierarchy Level | 5597
- Description | 5597
- Options | 5597
- Required Privilege Level | 5598
- Release Information | 5598

Syntax

```
dhcp-client {
    client-identifier {
        prefix {
            host-name;
            logical-system-name;
            routing-instance-name;
        }
        use-interface-description (device | logical);
        user-id (ascii string| hexadecimal string);
```

```
    }
    force-discover;
    lease-time (length | infinite);
    no-dns-install;
    options no-hostname;
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family]
```

Description

Configure the Dynamic Host Configuration Protocol (DHCP) client.

Options

- **force-discover**—Forces the DHCP client to send a DHCP discover packet after one to three failed `dhcp-request` attempts. The `force-discover` option ensures that the DHCP server will assign the same or a new IP address to the client.
- **lease-time**—Specify the time to negotiate and exchange DHCP information.
  - **infinite**—Lease never expires.
  - **length**—Number of seconds.
- **no-dns-install**—Do not add DNS information to the DHCP client (resolve.conf) even after learn from DHCP server
- **retransmission-attempt**—Specify the number of times the device attempts to retransmit a DHCP packet fallback. Range is 0-6.
- **server-address**—Specify the preferred DHCP server address that is sent to DHCP clients.
- **update-server**—Propagate TCP/IP settings to a local DHCP server.
- **vendor-id**—Vendor class ID for the DHCP client.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

Starting with Junos OS Release 17.4R1 and Junos OS Release 15.1X49-D60, on all SRX Series Firewalls and vSRX Virtual Firewall instances, the CLI option `dhcp-client` at `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy is changed to `dhcp` to align with other Junos platforms. There is no change in the functionality.

RELATED DOCUMENTATION

DHCP Overview

dhcp (DHCP Client)

IN THIS SECTION

- Syntax (EX Series) | **5599**
- Syntax (SRX Series) | **5599**
- Hierarchy level (EX Series) | **5599**
- Hierarchy level (SRX Series) | **5600**
- Description | **5600**
- Options | **5600**
- Required Privilege Level | **5600**
- Release Information | **5601**

Syntax (EX Series)

```
dhcp {
    client-identifier duid-type (duid-ll | duid-llt | vendor);
    no-dns-install;
    rapid-commit;
    options name;
}
```

Syntax (SRX Series)

```
dhcp {
    client-identifier {
        (ascii string | hexadecimal string);
    }
    force-discover;
    lease-time (length | infinite);
    metric;
    no-dns-install;
    options;
    requested-options;
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
}
```

Hierarchy level (EX Series)

```
[edit interfaces interface-name unit logical-unit-number family inet]
[edit logical-systems name interfaces interface-name unit logical-unit-number family inet]
[edit tenants tenant-name interfaces interface-name unit logical-unit-number family inet]
```

Hierarchy level (SRX Series)

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

Configure a Dynamic Host Configuration Protocol (DHCP) client for an IPv4 interface for logical systems and tenant systems.

The remaining statements are described separately.

> **NOTE**: Starting in Junos OS Release 18.1R1, DHCPv4 and DHCPv6 clients are supported on management interfaces (fxp0 and em0) configured in the non-default management routing instance, `mgmt_junos`.

Options

client-identifier duid-type	Identify a client by a client-identifier value. This statement is mandatory.
no-dns-install	Do not add DNS information to the DHCP client even after it is learned from the DHCP server.
options	Specify options requested by the DHCPv4 client.
force-discover	Send DHCPDISCOVER after DHCPREQUEST retransmission failure
lease-time	Specify lease time in seconds requested in DHCP client protocol packet (60 through 2147,,483,647 seconds for SRX devices)
metric	client initiated default-route metric (0..255 for SRX Series devices)
requested-options	Specify the DHCP options.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

The `logical-systems` and `tenants` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Understanding Interfaces

Configuring a DHCP Client

dhcp-local-server

IN THIS SECTION

- Syntax | **5601**
- Hierarchy Level | **5612**
- Description | **5612**
- Required Privilege Level | **5613**
- Release Information | **5613**

Syntax

```
dhcp-local-server {
    access-profile profile-name;
    allow-active-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
    }
    allow-bulk-leasequery {
        max-connections number-of-connections;
        max-empty-replies number-of-replies;
        restricted-requestor;
```

```
        timeout seconds;
    }
    allow-leasequery {
        restricted-requestor;
    }
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name ;
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dhcpv6 {
        access-profile profile-name;
        allow-active-leasequery {
            idle-timeout seconds;
            peer-address address;
            timeout seconds;
        }
        allow-bulk-leasequery {
            max-connections number-of-connections;
            max-empty-replies number-of-replies;
            restricted-requestor;
            timeout seconds;
        }
        allow-leasequery {
            restricted-requestor;
        }
        authentication {
            ...
        }
        duplicate-clients incoming-interface;
        group group-name {
```

```
access-profile profile-name;
authentication {
    ...
}
interface interface-name {
    access-profile profile-name;
    exclude;
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
        rapid-commit;
    }
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    trace;
    upto upto-interface-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
```

```
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                        session-mode (automatic | multihop | singlehop);
                        holddown-interval milliseconds;
                    }
                    layer2-liveness-detection {
                        max-consecutive-retries number;
                        transmit-interval interval;
                    }
                }
            }
            overrides {
                asymmetric-lease-time seconds;
                asymmetric-prefix-lease-time seconds;
                delay-advertise {
                    based-on (option-15 | option-16 | option-18 | option-37) {
                        equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        not-equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        starts-with {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                    }
                    delay-time seconds;
                }
```

```
                delegated-pool;
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                multi-address-embedded-option-response;
                process-inform {
                    pool pool-name;
                }
                protocol-attributes attribute-set-name;
                rapid-commit;
            }
            route-suppression;
            server-duid-type type;
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                    session-mode (automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        overrides {
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
```

```
            delay-advertise {
                based-on (option-15 | option-16 | option-18 | option-37) {
                    equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    not-equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    starts-with {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                }
                delay-time seconds;
            }
            delegated-pool;
            dual-stack dual-stack-group-name;
            include-option-82 {
                forcerenew;
                nak;
            }
            interface-client-limit number;
            multi-address-embedded-option-response;
            process-inform {
                pool pool-name;
            }
            protocol-attributes attribute-set-name;
            rapid-commit;
        }
        reconfigure {
            attempts attempt-count;
            clear-on-terminate;
            strict;
            support-option-pd-exclude;
            timeout timeout-value;
            token token-value;
             trigger {
                radius-disconnect;
            }
        }
        reauthenticate (<lease-renewal> <remote-id-mismatch >);
```

```
        requested-ip-network-match subnet-mask;
        route-suppression;
        server-duid-type type;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    dual-stack-group name {
        access-profile access-profile;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name ;
                logical-system-name;
                mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                routing-instance-name;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        classification-key {
            circuit-id circuit-id;
            mac-address mac-address;
            remote-id remote-id;
        }
        dual-stack-interface-client-limit number;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
```

```
        }
        on-demand-address-allocation;
        protocol-primary (inet | inet6);
        reauthenticate (<lease-renewal> <remote-id-mismatch >);
        service-profile service-profile;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    duplicate-clients-in-subnet  (incoming-interface | option-82);
    dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary primary-profile-
name>;
    forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
    group group-name {
        authentication {
            ...
        }
        dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary primary-
profile-name>;
        interface interface-name {
            exclude;
            overrides {
                asymmetric-lease-time seconds;
                client-discover-match (option60-and-option82 | incoming-interface);
                delay-offer {
                    based-on (option-60 | option-77 | option-82) {
                        equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        not-equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        starts-with {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                    }
                    delay-time seconds;
                }
                include-option-82 {
                    forcerenew;
                    nak;
                }
```

```
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                process-inform {
                    pool pool-name;
                }
                protocol-attributes attribute-set-name;
            }
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
            trace;
            upto upto-interface-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                    session-mode(automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        overrides {
            asymmetric-lease-time seconds;
            client-discover-match  (option60-and-option82 | incoming-interface);
            delay-offer {
                based-on (option-60 | option-77 | option-82) {
                    equals {
```

```
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        include-option-82 {
            forcerenew;
            nak;
        }
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
    }
    requested-ip-network-match subnet-mask
    route-suppression;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
```

```
                detection-time {
                    threshold milliseconds;
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    on-demand-address-allocation;
    overrides {
        asymmetric-lease-time seconds;
        client-discover-match  <option60-and-option82 | incoming-interface>;
        delay-offer {
            based-on (option-60 | option-77 | option-82) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
    }
    pool-match-order {
        external-authority;
        ip-address-first;
```

```
        option-82;
    }
    protocol-primary;
    reauthenticate (<lease-renewal> <remote-id-mismatch >);
    reconfigure {
        attempts attempt-count;
        clear-on-terminate;
        strict;
        timeout timeout-value;
        token token-value;
         trigger {
            radius-disconnect;
        }
    }
    requested-ip-network-match subnet-mask;
    route-suppression;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services],
[edit logical-systems logical-system-name system services],
[edit routing-instances routing-instance-name system services],
[edit system services]
```

Description

Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch to enable the router or switch to function as an extended DHCP local server. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client.

The extended DHCP local server is incompatible with the DHCP server on J Series routers and, therefore, is not supported on J Series routers. Also, the DHCP local server and the DHCP/BOOTP relay server, which are configured under the `[edit forwarding-options helpers]` hierarchy level, cannot both be enabled on the router or switch at the same time. The extended DHCP local server is fully compatible with the extended DHCP relay feature.

The `dhcpv6` stanza configures the router or switch to support Dynamic Host Configuration Protocol for IPv6 (DHCPv6). The DHCPv6 local server is fully compatible with the extended DHCP local server and the extended DHCP relay feature.

> **NOTE**: When you configure the `dhcp-local-server` statement at the routing instance hierarchy level, you must use a routing instance type of `virtual-router`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Understanding Differences Between Legacy DHCP and Extended DHCP

DHCPv6 Local Server Overview

dhcp-local-server (System Services)

IN THIS SECTION

- Syntax | 5614
- Hierarchy Level | 5618
- Description | 5618
- Options | 5618
- Required Privilege Level | 5618
- Release Information | 5618

Syntax

```
dhcp-local-server {
    dhcpv6 {
        authentication {
            password password;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name;
                interface-name;
                logical-system-name;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix;
            }
        }
        dynamic-profile {
            profile-name;
            aggregate-clients {
                merge;
                replace;
            }
            junos-default-profile;
            use-primary dynamic-profile-name;
        }
        group group-name {
            authentication {
                password password;
                username-include {
                    circuit-type;
                    client-id;
                    delimiter delimiter-character;
                    domain-name domain-name;
                    interface-name;
                    logical-system-name;
                    relay-agent-interface-id;
                    relay-agent-remote-id;
                    relay-agent-subscriber-id;
```

```
                routing-instance-name;
                user-prefix user-prefix;
            }
        }
        dynamic-profile {
            profile-name;
            aggregate-clients {
                merge;
                replace;
            }
            junos-default-profile;
            use-primary dynamic-profile;
        }
        interface interface-name {
            dynamic-profile {
                profile-name;
                aggregate-clients {
                    merge;
                    replace;
                }
                junos-default-profile;
                use-primary dynamic-profile-name;
            }
            exclude;
            overrides {
                delegated-pool pool-name;
                interface-client-limit number;
                process-inform {
                    pool pool-name;
                }
                rapid-commit ;
            }
            service-profile service-profile-name
            trace ;
            upto interface-name;
        }
        liveness-detection {
            failure-action {
                 clear-binding;
                clear-binding-if-interface-up;
                log-only;
            }
            method {
```

```
                bfd {
                    detection-time {
                        threshold milliseconds;
                    }
                    holddown-interval interval;
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    session-mode (automatic | multihop | single-hop);
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                version (0 | 1 | automatic);
            }
        }
        overrides {
            delegated-pool pool-name;
            interface-client-limit number;
            process-inform {
                pool pool-name;
            }
            rapid-commit ;
        }
        reconfigure {
            attempts number;
            clear-on-abort;
            strict;
            timeout number;
            token token-name;
            trigger {
                radius-disconnect;
            }
        }
        service-profile service-profile-name;
    }
    liveness-detection {
        failure-action {
             clear-binding;
            clear-binding-if-interface-up;
            log-only;
        }
```

```
            method {
                bfd {
                    detection-time {
                        threshold milliseconds;
                    }
                    holddown-interval interval;
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    session-mode (automatic | multihop | single-hop);
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                version (0 | 1 | automatic);
            }
        }
        overrides {
            delegated-pool pool-name;
            interface-client-limit number;
            process-inform {
                pool pool-name;
            }
            rapid-commit ;
        }
        reconfigure {
            attempts number;
            clear-on-abort;
            strict;
            timeout number;
            token token-name;
            trigger {
                radius-disconnect;
            }
        }
        service-profile service-profile-name;
    }
    group group-name {
        interface interface-name {
            exclude;
            upto upto-interface-name;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure DHCP Local Server for DHCPv6, forwarding snoop (unicast) packets, and setting traceoptions.

NOTE: SRX Series Firewalls do not support client authentication.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

DHCP Overview

dhcp-option82

IN THIS SECTION

- Syntax | 5619
- Hierarchy Level | 5619
- Description | 5620
- Default | 5620
- Required Privilege Level | 5620
- Release Information | 5620

Syntax

```
dhcp-option82 {
    circuit-id {
        prefix hostname;
        use-interface-description;
        use-vlan-id;
    }
    remote-id {
        prefix hostname | mac | none;
        use-interface-description;
        use-string                          string;
    }
    vendor-id <string>;
}
```

Hierarchy Level

```
[edit ethernet-switching-options                secure-access-port vlan (all | vlan-name)]
[edit forwarding-options helpers bootp]
[edit forwarding-options helpers bootp interface interface-name]
```

Description

When the switch receives a DHCP request from a DHCP client connected on one of the switch's interfaces, have the switch insert DHCP option 82 (also known as the DHCP relay agent information option) information in the DHCP request packet header before it forwards or relays the request to a DHCP server. The server uses the option 82 information, which provides details about the circuit and host the request came from, in formulating the reply; the server does not, however, make any changes to the option 82 information in the packet header. The switch receives the reply and then removes the DHCP option 82 information before forwarding the reply to the client.

The remaining statements are explained separately. See CLI Explorer.

Default

Insertion of DHCP option 82 information is not enabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Example: Setting Up DHCP Option 82 Using the Same VLAN

Example: Setting Up DHCP Option 82

http://tools.ietf.org/html/rfc3046.

dhcp-relay

IN THIS SECTION

Syntax | 5621

- Hierarchy Level | 5635
- Description | 5636
- Required Privilege Level | 5636
- Release Information | 5636

Syntax

```
dhcp-relay {
    access-profile profile-name;
    active-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
        topology-discovery;
    }
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name;
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            routing-instance-name;
            user-prefix user-prefix-string;
stale-timer            vlan-tags;
        }
    }
    bulk-leasequery {
        attempts number-of-attempts;
        timeout seconds;
    }
    dhcpv6 {
```

```
        access-profile profile-name;
        active-leasequery {
            idle-timeout seconds;
            peer-address address;
            timeout seconds;
            topology-discovery;
        }
        active-server-group server-group-name;
        }
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name interface-name;
                logical-system-name;
                mac-address mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        bulk-leasequery {
            attempts number-of-attempts;
            timeout seconds;
            trigger automatic;
        }
        duplicate-clients incoming-interface;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        forward-only {
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        forward-only-replies;
```

```
        }
        forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-
interfaces);
        group group-name {
            access-profile profile-name;
            active-server-group server-group-name;
            authentication {
                password password-string;
                username-include {
                    circuit-type;
                    client-id;
                    delimiter delimiter-character;
                    domain-name domain-name-string;
                    interface-description  (device-interface | logical-interface);
                    interface-name interface-name;
                    logical-system-name;
                    mac-address mac-address;
                    relay-agent-interface-id;
                    relay-agent-remote-id;
                    relay-agent-subscriber-id;
                    routing-instance-name;
                    user-prefix user-prefix-string;
                    vlan-tags;
                }
            }
            dynamic-profile profile-name {
                aggregate-clients (merge | replace);
                use-primary primary-profile-name;
            }
            forward-only {
                logical-system <current | default | logical-system-name>;
                routing-instance <current | default | routing-instance-name>;
            }
            interface interface-name {
                access-profile profile-name;
                dynamic-profile profile-name {
                    aggregate-clients (merge | replace);
                    use-primary primary-profile-name;
                }
                exclude;
                overrides {
                    allow-snooped-clients;
                    asymmetric-lease-time seconds;
```

```
                asymmetric-prefix-lease-time seconds;
                client-negotiation-match incoming-interface;
                delay-authentication;
                delete-binding-on-renegotiation;
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                no-allow-snooped-clients;
                no-bind-on-request;
                relay-source interface-name;
                send-release-on-delete;
            }
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
            trace;
            upto upto-interface-name;
        }
        }
        lease-time-validation {
            lease-time-threshold seconds;
            violation-action action;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                    session-mode(automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
```

```
                }
            }
        }
        overrides {
            allow-snooped-clients;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-negotiation-match incoming-interface;
            delay-authentication;
            delete-binding-on-renegotiation;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            no-allow-snooped-clients;
            no-bind-on-request;
            relay-source interface-name;
            send-release-on-delete;
        }
        relay-agent-interface-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
        relay-agent-remote-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
        relay-option {
            option-number option-number;
            default-action {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            equals (ascii ascii-string | hexadecimal hexadecimal-string) {
```

```
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
        }
        remote-id-mismatch disconnect;
        route-suppression;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    leasequery {
        attempts number-of-attempts;
        timeout seconds;
    }
    lease-time-validation {
        lease-time-threshold seconds;
        violation-action action;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
```

```
                max-consecutive-retries number;
                transmit-interval interval;
            }
            route-suppression;
            service-profile dynamic-profile-name;
        }
    }
    no-snoop;
    overrides {
        allow-snooped-clients;
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-negotiation-match incoming-interface;
        delay-authentication;
        delete-binding-on-renegotiation;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        no-allow-snooped-clients;
        no-bind-on-request;
        relay-source interface-name;
        send-release-on-delete;
    }
    relay-agent-interface-id {
        include-irb-and-l2;
        keep-incoming-interface-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    relay-agent-remote-id {
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    relay-option {
        option-number option-number;
        default-action {
```

```
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            equals (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
        }
        relay-option-vendor-specific{
            host-name;
            location;
        remote-id-mismatch disconnect;
        route-suppression;
        server-group {
            server-group-name {
                server-ip-address;
            }
        }
        server-response-time seconds;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    dual-stack-group dual-stack-group-name {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name;
                logical-system-name;
                mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
```

```
                routing-instance-name;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        classification-key {
            circuit-id circuit-id;
            mac-address mac-address;
            remote-id remote-id;
        }
        dual-stack-interface-client-limit number;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        protocol-primary (inet | inet6);
        relay-agent-interface-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
        relay-agent-remote-id {
            include-irb-and-l2;
            keep-incoming-remote-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
```

```
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    duplicate-clients-in-subnet (incoming-interface | option-82):
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    forward-only {
        logical-system <current | default | logical-system-name>;
        routing-instance <current | default | routing-instance-name>;
    }
    forward-only-replies;
    forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
    group group-name {
        access-profile profile-name;
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 [circuit-id] [remote-id];
                routing-instance-name;
                user-prefix user-prefix-string;
                }
            vlan-tags;
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        forward-only {
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        forward-only {
```

```
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        interface interface-name {
            access-profile profile-name;
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                        session-mode (automatic | multihop | singlehop);
                        holddown-interval milliseconds;
                    }
                }
            }
            overrides {
                allow-no-end-option;
                allow-snooped-clients;
                always-write-giaddr;
                always-write-option-82;
                asymmetric-lease-time seconds;
                client-discover-match <option60-and-option82 | incoming-interface>;
                delay-authentication;
                delete-binding-on-renegotiation;
                disable-relay;
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                layer2-unicast-replies;
                no-allow-snooped-clients;
                no-bind-on-request;
                proxy-mode;
```

```
                relay-source
                replace-ip-source-with;
                send-release-on-delete;
                trust-option-82;
            }
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
            trace;
            upto upto-interface-name;
        }
        overrides {
            allow-no-end-option
            allow-snooped-clients;
            always-write-giaddr;
            always-write-option-82;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-discover-match (option60-and-option82 | incoming-interface);
            delay-authentication;
            delete-binding-on-renegotiation;
            disable-relay;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            layer2-unicast-replies;
            no-allow-snooped-clients;
            no-bind-on-request;
            proxy-mode;
            relay-source
            replace-ip-source-with;
            send-release-on-delete;
            trust-option-82;
        }
        relay-option {
            option-number option-number;
            default-action {
                drop;
                forward-only;
                relay-server-group group-name;
            }
            equals (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
```

```
            }
            starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                local-server-group local-server-group;
                relay-server-group relay-server-group;
            }
        }
        relay-option-82 {
            circuit-id {
                prefix prefix;
                use-interface-description (logical | device);
            }
            remote-id {
                prefix prefix;
                use-interface-description (logical | device);
            }
            server-id-override
        }
        remote-id-mismatch disconnect;
        route-suppression:
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    leasequery {
        attempts number-of-attempts;
        timeout seconds;
    }
    lease-time-validation {
        lease-time-threshold seconds;
        violation-action action;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
```

```
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    no-snoop;
    overrides {
        allow-no-end-option
        allow-snooped-clients;
        always-write-giaddr;
        always-write-option-82;
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-discover-match (option60-and-option82 | incoming-interface);
        delay-authentication;
        delete-binding-on-renegotiation;
        disable-relay;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        layer2-unicast-replies;
        no-allow-snooped-clients;
        no-bind-on-request;
        proxy-mode;
        relay-source
        replace-ip-source-with;
        send-release-on-delete;
        trust-option-82;
    }
    relay-option {
        option-number option-number;
        default-action {
            drop;
            forward-only;
            relay-server-group group-name;
```

```
        }
        equals (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            relay-server-group relay-server-group;
        }
        starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
    }
    relay-option-82 {
        circuit-id {
            prefix prefix;
            use-interface-description (logical | device);
        }
        remote-id {
            prefix prefix;
            use-interface-description (logical | device);
        }
        server-id-override
    }
    }
    remote-id-mismatch disconnect;
    route-suppression:
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
    server-response-time seconds;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit forwarding-options],
[edit logical-systems logical-system-name forwarding-options],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options],
[edit routing-instances routing-instance-name forwarding-options]
```

Description

Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch to enable the router (or switch) to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the `dhcp-relay` and `dhcpv6` statements are incompatible with the DHCP/BOOTP relay agent options configured with the `bootp` statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router (or switch) at the same time.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

DHCPv6 Relay Agent Overview

DHCP Relay Proxy Overview

Specifying Authentication Support

dhcp-relay (DHCP Spoofing Prevention)

IN THIS SECTION

- Syntax | 5637
- Hierarchy Level | 5637
- Description | 5637
- Required Privilege Level | 5638
- Release Information | 5638

Syntax

```
dhcp-relay {
    group group-name {
        interface interface-name;
    }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name forwarding-
options],
[edit routing-instances routing-instance-name forwarding-options]
```

Description

Configure Dynamic Host Configuration Protocol (DHCP) snooping on the router. When acting as a snooping agent, the MX Series router typically is located between the client and the DHCP relay agent. It creates filters by “snooping” DHCP messages and binding DHCP-issued IP addresses to the MAC address of the client. These filters help prevent DHCP spoofing.

Configure DHCP snooping by including the appropriate interfaces in the DHCP relay configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4 (MX Series routers only).

RELATED DOCUMENTATION

Preventing DHCP Spoofing on MX Series 5G Universal Routing Platforms

dhcp-relay (Easy EVPN LAG Configuration)

IN THIS SECTION

- Syntax | 5638
- Hierarchy Level | 5639
- Description | 5639
- Options | 5639
- Required Privilege Level | 5640
- Release Information | 5640

Syntax

```
dhcp-relay name {
    dhcp-server-address dhcp-server-address;
    relay-source DHCP-relay-source-interface;
    vrf-instance vrf-instance-id;
};
```

Hierarchy Level

```
[edit services evpn evpn-vxlan]
```

Description

Specify Dynamic Host Configuration Protocol (DHCP) relay parameters for a generated configuration for an Ethernet VPN (EVPN) fabric with Virtual Extensible LAN (VXLAN) encapsulation.

You use the statements in the `[edit services evpn]` hierarchy level to simplify configuring an EVPN fabric with Ethernet segment identifier (ESI) link aggregation groups (LAGs) in the topology. The feature you enable with statements in this hierarchy is an easy EVPN LAG configuration feature, also called EZ-LAG.

You only need to configure a few statements to provide the parameters required for an EVPN core fabric with multihomed or single-homed customer edge (CE) devices (or end server devices). When you commit the `[edit services evpn]` configuration statements, the device uses a commit script to generate the corresponding standard EVPN fabric configuration statements.

The statements at the `[edit services evpn evpn-vxlan dhcp-relay]` hierarchy level provide the DHCP relay parameters the commit script needs to generate a default configuration for a DHCP relay group in an EVPN-VXLAN fabric.

Options

name	The group name to identify a DHCP server and relay group. The commit script generates configuration statements to create the DHCP relay group.
dhcp-server-address *dhcp-server-address*	(Required) The IP address of the DHCP server. The commit script uses this address in the generated configuration to enable the DHCP relay server group.
relay-source *DHCP-relay-source-interface*	The DHCP relay source loopback interface IP address. The commit script generates DHCP `overrides relay-source` statements that use this source address.
vrf-instance *vrf-instance-id*	(Required) The virtual routing and forwarding (VRF) instance identifier for which you want to configure DHCP relay. You specify VRF instance IDs by number in this easy EVPN LAG configuration feature. The commit script appends the VRF instance ID to a default VRF instance name to generate the configuration statements for different tenant VRF instances. • **Range:** 1 through 255

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Easy EVPN LAG (EZ-LAG) Configuration

DHCP Relay Agent over EVPN-VXLAN

https://www.juniper.net/documentation/us/en/software/junos/dhcp/index.html

dhcp-security

IN THIS SECTION

- Syntax | 5640
- Hierarchy Level | 5642
- Description | 5643
- Required Privilege Level | 5643
- Release Information | 5643

Syntax

```
dhcp-security {
    arp-inspection;
    dhcpv6-options {
        light-weight-dhcpv6-relay;
        option-16 {
            use-string string;
```

```
        }
        option-18 {
            prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
                vlan-id;
                vlan-name;
            }
            use-interface-mac;
            use-interface-index (device | logical);
            use-interface-description (device | logical);
            use-interface-name (device | logical);
            use-string string;
        }
        option-37{
            prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
                vlan-id;
                vlan-name;
            }
            use-interface-mac;
            use-interface-index (device | logical);
            use-interface-description (device | logical);
            use-interface-name (device | logical);
            use-string string;
        }
        option-79;
    }
    group group-name {
        interface interface-name {
            static-ip ip-address {
                mac mac-address;
            }
            static-ipv6 ip-address {
                mac mac-address;
            }
        }
        overrides {
            no-dhcpv6-options;
            no-option16;
```

```
            no-option18;
            no-option37;
            no-option82;
            trusted;
            untrusted;
        }
    }
    ip-source-guard;
    ipv6-source-guard;
    mine-dhcp-client-options;
    mine-dhcpv6-client-options;
    neighbor-discovery-inspection;
    no-dhcp-snooping;
    no-dhcpv6-snooping;
    option-82 {
        circuit-id {
            prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
            }
            use-interface-description (device | logical);
            use-vlan-id;
        }
        remote-id {
            host-name hostname;
            use-interface-description (device | logical);
            mac;
            use-string string;
        }
        vendor-id {
            use-string string;
        }
    }
}
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options]
```

Description

Configure DHCP or DHCPv6 snooping on the switch. DHCP snooping is also enabled automatically if you configure any of the following port security features within this hierarchy:

- Dynamic ARP inspection (DAI)
- IP source guard
- DHCP option 82
- Static IP

For switches that support DHCPv6, both DHCP snooping and DHCPv6 snooping are enabled automatically if you configure any of the afore-mentioned features or any of the following IPv6 features:

- IPv6 neighbor discovery inspection
- IPv6 source guard
- Static IPv6

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10 .

Support for `static-ipv6`, `neighbor-discovery-inspection`, `ipv6-source-guard`, `no-dhcpv6-snooping`, and `no-option37` introduced in Junos OS Release 13.2X51-D20 for EX Series switches.

Support for `dhcpv6-options`, `option-16`, `option-18`, `option-37`, `no-dhcpv6-options`, `no-option16`, `no-option18`, and `no-option37` introduced in Junos OS Release 14.2 for EX Series switches.

Support for `mine-dhcp-client-options` and `mine-dhcpv6-client-options` introduced in Junos OS Release 23.1.

RELATED DOCUMENTATION

Enabling Dynamic ARP Inspection (ELS)

Configuring IP Source Guard (ELS)

Setting Up DHCP Option 82 on the Switch with No Relay (ELS)

Configuring Static DHCP IP Addresses for DHCP snooping (ELS)

dhcp-security (MX Series)

IN THIS SECTION

- Syntax | 5644
- Hierarchy Level | 5645
- Description | 5645
- Options | 5645
- Required Privilege Level | 5646
- Release Information | 5646

Syntax

```
dhcp-security {
    arp-inspection;
    group group-name {
        interface interface-name {
            static-ip ip-address {
                mac mac-address;
            }
        }
        overrides {
            no-option82;
            trusted;
            untrusted;
        }
    }
    ip-source-guard;
    no-dhcp-snooping;
    option-82 {
```

```
        circuit-id {
            prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
            }
            use-interface-description (device | logical);
            use-vlan-id;
        }
        remote-id {
            host-name;
            use-interface-description (device | logical);
            use-string string;
        }
        vendor-id {
            use-string string;
        }
    }
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]
```

Description

Configure port security features on the switching device. DHCP snooping is enabled automatically if you configure any of the following port security features within this hierarchy:

- Dynamic ARP inspection (DAI)
- IP source guard
- DHCP option 82
- Static IP

Options

mac-address—Value (in hexadecimal format) of the address assigned to this device.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

- *Applying CoS Forwarding Classes to Prioritize Inspected Packets*
- *Configuring IP Source Guard (non-ELS)*
- *Configuring DHCP Option 82 on a Router with Bridge Domain*
- *Configuring Static DHCP IP Addresses for DHCP snooping (MX routers)*

dhcp-service

IN THIS SECTION

- Syntax | 5646
- Hierarchy Level | 5648
- Description | 5648
- Required Privilege Level | 5648
- Release Information | 5648

Syntax

```
dhcp-service {
    accept-max-tcp-connections max-tcp-connections;
    dhcp-snooping-file(local_pathname | remote_URL) {
        write-interval interval;
    }
```

```
    dhcpv6-snooping-file {
        location;
        write-interval seconds;
    }
    (disable | enable);
    interface-traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        level  (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
    log {
        session {
            client;
            all;
            dhcpv6 {
                client;
                server;
                relay;
                dynamic-server;
                all;
            }
            server;
            relay;
        }
    }
    ltv-syslog-interval seconds;
    persistent-storage {
        backup-interval backup-interval;
        file-name;
    }
    request-max-tcp-connections max-tcp-connections;
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        level  (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Enable DHCP services on the device. DHCP services automate network-parameter assignment to network devices. The DHCP service process is enabled by default. However, by default, IP-MAC bindings in the DHCP snooping database do not persist through device reboots. You can inprove performance after rebooting by configuring the IP-MAC bindings to persist, by configuring a storage location for the DHCP database file. When specifying the location for the DHCP database, you must also specify how frequently the switch writes the database entries into the DHCP snooping database file.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

Support for log option introduced in Junos OS Release 19.1R1 for SRX Series Firewalls.

RELATED DOCUMENTATION

- Configuring Persistent Bindings in the DHCP or DHCPv6 (ELS)

dhcp-snooping-file (EX)

IN THIS SECTION

- Syntax | 5649
- Hierarchy Level | 5649
- Description | 5649
- Default | 5649
- Required Privilege Level | 5650
- Release Information | 5650

Syntax

```
dhcp-snooping-file {
                        location                    ( local_pathname | remote_URL);
                        timeout                     seconds;
                        write-interval                    seconds;
}
```

Hierarchy Level

```
[edit ethernet-switching-options                secure-access-port]
```

Description

Ensure that IP-MAC bindings persist through switch reboots by specifying a local pathname or a remote URL for the storage location of the DHCP snooping database file.

The remaining statements are explained separately. See CLI Explorer.

Default

The IP-MAC bindings in the DHCP snooping database file are not persistent. If the switch is rebooted, the bindings are lost.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Persistent Bindings in the DHCP or DHCPv6 (non-ELS)

Understanding DHCP Snooping (non-ELS)

dhcp-snooping-file

IN THIS SECTION

- Syntax | 5650
- Hierarchy Level | 5651
- Description | 5651
- Default | 5651
- Required Privilege Level | 5651
- Release Information | 5651

Syntax

```
dhcp-snooping-file  (local_pathname | remote_URL);
    write-interval seconds;
}
```

Hierarchy Level

```
[edit system processes  dhcp-service]
```

Description

Ensure that IP-MAC bindings persist through the device reboots by specifying a local pathname or a remote URL for the storage location of the DHCP snooping database file. You *must* specify how frequently the device writes the database entries into the DHCP snooping database file.

The remaining statement is explained separately. See CLI Explorer.

Default

The IP-MAC bindings in the DHCP snooping database file are not persistent by default. If the device is rebooted, the bindings are lost, and the table must be rebuilt on reboot.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Configuring Persistent Bindings in the DHCP or DHCPv6 (ELS)

Configuring Static DHCP IP Addresses for DHCP snooping (MX routers)

Understanding DHCP Snooping (non-ELS)

dhcp-tags (Adjustment Control Profiles)

IN THIS SECTION

- Syntax | 5652
- Hierarchy Level | 5652
- Description | 5652
- Options | 5653
- Required Privilege Level | 5653
- Release Information | 5653

Syntax

```
dhcp-tags {
    algorithm algorithm;
    priority priority;
}
```

Hierarchy Level

```
[edit class-of-service adjustment-control-profiles profile-name application]
```

Description

Configure the shaping rate adjustment controls for the DHCP tags application. DHCP tags are supported for session negotiation for DHCPv4 and DHCPv6 over IP demux and VLAN demux interfaces. This means that shaping rates received from DHCP option 82 in the DISCOVER message or DHCPv6 option 17 in the SOLICIT can be used to apply a shaping rate to the logical interface.

NOTE: Single-session dual-stack DHCP is not fully supported; for example, when rates vary between the individual DHCP sessions during negotiation.

Options

algorithm Rate adjustment algorithm used by the DHCP Tags application.

- Values:
  - adjust-always—Adjust the shaping rate unconditionally.
  - adjust-greater—Adjust the shaping rate if it is greater than the configured value.
  - adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.
  - adjust-less—Adjust the shaping rate if it is less than the configured value.
  - adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
  - adjust-never—Do not perform rate adjustments.
- **Default:** adjust-less

priority Priority of the DHCP tags application in the adjustment control profile.

- **Range:** 1 through 10; 1 is the highest priority.
- **Default:** 2

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R1.

RELATED DOCUMENTATION

Configuring CoS Adjustment Control Profiles

CoS Adjustment Control Profiles Overview

Verifying the CoS Adjustment Control Profile Configuration

dhcp-trusted

IN THIS SECTION

- Syntax | 5654
- Hierarchy Level | 5654
- Description | 5654
- Default | 5654
- Required Privilege Level | 5655
- Release Information | 5655

Syntax

```
(dhcp-trusted | no-dhcp-trusted);
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port interface (all |
interface-name)]
```

Description

Allow DHCP responses from the specified interfaces (ports) or all interfaces.

- **dhcp-trusted**—Allow DHCP responses.
- **no-dhcp-trusted**—Deny DHCP responses.

Default

Trusted for trunk ports, untrusted for access ports.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Port Security (non-ELS)

Example: Protecting against Rogue DHCP Server Attacks

Enabling a Trusted DHCP Server (non-ELS)

dhcpv6-client

IN THIS SECTION

- Syntax | 5655
- Hierarchy Level | 5656
- Description | 5656
- Options | 5656
- Required Privilege Level | 5657
- Release Information | 5657

Syntax

```
dhcpv6-client {
    client-ia-type {
        ia-na;
        ia-pd;
    }
```

```
    client-identifier duid-type (duid-ll | duid-llt | vendor);
    client-type (autoconfig | stateful);
    interface-identifier interface_id;
rapid-commit;
    req-option (dns-server | domain | fqdn | nis-domain | nis-server | ntp-server | sip-domain |
sip-server |time-zone | vendor-spec);
    retransmission-attempt number;
    update-router-advertisement {
        interface interface-name {
            prefix-identifier subnet-id;
        interface-identifier interface_id;
            }
    }
    update-server;
    use-ra-prefix;
    vendor-id <model-number>:<serial-number>;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6]
[edit logical-systems name interfaces name unit name family inet6]
[edit tenants tenant-name interfaces name unit name family inet6]
```

Description

Configure the Dynamic Host Configuration Protocol version 6 (DHCPv6) client.

> **NOTE**: NOTE: Starting in Junos OS Release 17.3R1 (on MX Series) and Junos OS Release 18.3R1 (on SRX Series), DHCPv4 and DHCPv6 clients are supported on management interfaces (fxp0 and em0) configured in the non-default management routing instance, `mgmt_junos`.

Options

client-ia-type	Identity association type for DHCPv6 client. This statement is mandatory.
client-identifier duid-type	Identity a client by a client-identifier value. This statement is mandatory.

client-type	Identify the type of DHCPv6 client. This statement is mandatory.
interface-identifier	(Optional) Specify the interface ID during assignment of WAN interface IP through Stateless Address Auto Configuration (SLAAC).
prefix-identifier	(Optional) Specify last bits of the prefix for the respective dynamic pool. The subnet-id must be of at least network-prefix-length bits. In absence of dynamic pool, the received prefix is used to configure LAN interfaces for router advertisement
rapid-commit	The use of the two-message exchange for address assignment.
req-option	Specify options requested by the DHCPv6 client.
retransmission-attempt number	Specify the number of times the device retransmits a DHCPv6 client packet if a DHCPv6 server fails to respond. After the specified number of attempts, no further attempts at reaching a server are made.
update-router-advertisement	Specify the interface used to delegate prefixes.
update-server	Propagate TCP/IP settings to the DHCPv6 server.
use-ra-prefix	(Optional) Enable WAN IP assignment from router advertisement prefix in stateful mode.

For detailed information about these commands, see CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

Statement introduced in Junos OS Release 24.2R1 for NFX Series Routers.

The `logical-systems` and `tenants` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

DHCP Overview

Minimum DHCPv6 Client Configuration on SRX Series Devices

dhcpv6 (DHCP Local Server)

IN THIS SECTION

- Syntax | 5658
- Hierarchy Level | 5664
- Description | 5664
- Required Privilege Level | 5664
- Release Information | 5664

Syntax

```
dhcpv6 {
    access-profile profile-name;
    allow-active-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
    }
    allow-bulk-leasequery {
        max-connections number-of-connections;
        max-empty-replies number-of-replies;
        restricted-requestor;
        timeout seconds;
    }
    allow-leasequery {
        restricted-requestor;
    }
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
```

```
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            logical-system-name;
            mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    duplicate-clients incoming-interface;
    group group-name {
        access-profile profile-name;
        authentication {
        ...
        interface interface-name {
            access-profile profile-name;
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                        session-mode(automatic | multihop | singlehop);
                        holddown-interval milliseconds;
                    }
                }
            }
            overrides {
```

```
                asymmetric-lease-time seconds;
                asymmetric-prefix-lease-time seconds;
                client-negotiation-match incoming-interface;
                delay-advertise {
                    based-on (option-15 | option-16 | option-18 | option-37) {
                        equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        not-equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        starts-with {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                    }
                    delay-time seconds;
                }
                delete-binding-on-renegotiation;
                interface-client-limit number;
                multi-address-embedded-option-response;
                process-inform {
                    pool pool-name;
                }
                protocol-attributes attribute-set-name;
                rapid-commit;
            }
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
            trace;
            upto upto-interface-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
```

```
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-negotiation-match incoming-interface;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delegated-pool;
        delete-binding-on-renegotiation;
        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
```

```
            }
            protocol-attributes attribute-set-name;
            rapid-commit;
        }
        route-suppression;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-negotiation-match incoming-interface;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
```

```
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delegated-pool;
        delete-binding-on-renegotiation;
        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
        rapid-commit;
        reconfigure {
            attempts attempt-count;
            clear-on-terminate;
            strict;
            timeout timeout-value;
            token token-value;
             trigger {
                radius-disconnect;
            }
        }
    }
    reauthenticate (<lease-renewal> <remote-id-mismatch >);
    reconfigure {
        attempts attempt-count;
        clear-on-terminate;
        strict;
        support-option-pd-exclude;
        timeout timeout-value;
        token token-value;
         trigger {
            radius-disconnect;
        }
    }
```

```
    requested-ip-network-match subnet-mask;
    route-suppression;
    server-duid-type type;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server],
[edit logical-systems logical-system-name system services dhcp-local-server],
[edit routing-instances routing-instance-name system services dhcp-local-server],
[edit system services dhcp-local-server]
```

Description

Configure DHCPv6 local server options on the router or switch to enable the router or switch to function as a server for the DHCP protocol for IPv6. The DHCPv6 local server sends and receives packets using the IPv6 protocol and informs IPv6 of the routing requirements of router clients. The local server works together with the AAA service framework to control subscriber access (or DHCP client access) and accounting.

The DHCPv6 local server is fully compatible with the extended DHCP local server and DHCP relay agent.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

DHCPv6 Local Server Overview

dhcpv6 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 5665
- Hierarchy Level | 5672
- Description | 5672
- Required Privilege Level | 5672
- Release Information | 5672

Syntax

```
dhcpv6 {
    access-profile profile-name;
    active-leasequery {
        idle-timeout seconds;
        peer-address address;
        timeout seconds;
        topology-discovery;
    }
    active-server-group server-group-name;
    }
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name interface-name;
            logical-system-name;
```

```
            mac-address mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;

        }
    }
    bulk-leasequery {
        attempts number-of-attempts;
        timeout seconds;
        trigger automatic;
    }
    duplicate-clients incoming-interface;
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    forward-only {
        logical-system <current | default | logical-system-name>;
        routing-instance <current | default | routing-instance-name>;
    }
    forward-only-replies;
    }
    forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
    group group-name {
        access-profile profile-name;
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name;
                logical-system-name;
                mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
```

```
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        forward-only {
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        interface interface-name {
            access-profile profile-name;
            dynamic-profile profile-name {
                aggregate-clients (merge | replace);
                use-primary primary-profile-name;
            }
            exclude;
            overrides {
                allow-snooped-clients;
                asymmetric-lease-time seconds;
                asymmetric-prefix-lease-time seconds;
                client-negotiation-match incoming-interface;
                delay-authentication;
                delete-binding-on-renegotiation;
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                no-allow-snooped-clients;
                no-bind-on-request;
                relay-source interface-name;
                send-release-on-delete;
            }
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
            trace;
            upto upto-interface-name;
        }
        }
        lease-time-validation {
            lease-time-threshold seconds;
```

```
            violation-action action;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                    session-mode(automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
                layer2-liveness-detection {
                    max-consecutive-retries number;
                    transmit-interval interval;
                }
            }
        }
        overrides {
            allow-snooped-clients;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-negotiation-match incoming-interface;
            delay-authentication;
            delete-binding-on-renegotiation;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            no-allow-snooped-clients;
            no-bind-on-request;
            relay-source interface-name;
            send-release-on-delete;
        }
        relay-agent-interface-id {
            include-irb-and-l2;
```

```
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82;
        }
        relay-agent-remote-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
        }
        relay-option {
            option-number option-number;
            default-action {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            equals (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
        }
        remote-id-mismatch disconnect;
        route-suppression;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    leasequery {
        attempts number-of-attempts;
        timeout seconds;
    }
    lease-time-validation {
        lease-time-threshold seconds;
```

```
        violation-action action;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
            route-suppression;
            service-profile dynamic-profile-name;
        }
    }
    no-snoop;
    overrides {
        allow-snooped-clients;
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-negotiation-match incoming-interface;
        delay-authentication;
        delete-binding-on-renegotiation;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        no-allow-snooped-clients;
        no-bind-on-request;
        relay-source interface-name;
        send-release-on-delete;
```

```
    }
    relay-agent-interface-id {
        include-irb-and-l2;
        keep-incoming-interface-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82;
    }
    relay-agent-remote-id {
        include-irb-and-l2;
        keep-incoming-interface-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
    }
    relay-option {
        option-number option-number;
        default-action {
            drop;
            forward-only;
            relay-server-group relay-server-group;
        }
        equals (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            relay-server-group relay-server-group;
        }
        starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            relay-server-group relay-server-group;
        }
    }
    relay-option-vendor-specific{
        host-name;
        location;
    remote-id-mismatch disconnect;
    route-suppression;
    server-group {
        server-group-name {
            server-ip-address;
```

```
        }
    }
    server-response-time seconds;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay]
```

Description

Configure DHCPv6 relay options on the router or switch and enable the router or switch to function as a DHCPv6 relay agent. A DHCPv6 relay agent forwards DHCPv6 request and reply packets between a DHCPv6 client and a DHCPv6 server.

The DHCPv6 relay agent server is fully compatible with the extended DHCP local server and DHCP relay agent. However, the options configured with the `dhcpv6` statement are incompatible with the DHCP/BOOTP relay agent options configured with the `bootp` statement. As a result, the DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot be enabled on the router or switch at the same time.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support for `forward-snooped-clients` introduced in Junos OS Release 15.1X53-D56 for EX Series switches and Junos OS Release 17.1R1.

RELATED DOCUMENTATION

dhcp-relay | 5620

DHCPv6 Relay Agent Overview

Specifying Authentication Support

dhcpv6-options

IN THIS SECTION

- Syntax | 5673
- Hierarchy Level | 5674
- Description | 5674
- Required Privilege Level | 5674
- Release Information | 5674

Syntax

```
dhcpv6-options {
    option-16 {
        use-string string;
    }
    option-18 {
        prefix {
            host-name;
            logical-system-name;
            routing-instance-name;
            vlan-id;
            vlan-name;
        }
        use-interface-mac;
        use-interface-index (device | logical);
        use-interface-description (device | logical);
        use-interface-name (device | logical);
        use-string string;
```

```
    }
    option-37 {
        prefix {
            host-name;
            logical-system-name;
            routing-instance-name;
            vlan-id;
            vlan-name;
        }
        use-interface-mac;
        use-interface-index (device | logical);
        use-interface-description (device | logical);
        use-interface-name (device | logical);
        use-string string;
    }
}
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security]
```

Description

Configure optional information to be included in DHCPv6 packets during the DHCPv6 snooping process.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10 .

RELATED DOCUMENTATION

no-dhcpv6-options | 11440

Setting Up DHCP Option 82 on the Switch with No Relay (ELS)

Configuring Static DHCP IP Addresses for DHCP snooping (ELS)

dhcpv6-snooping-file

IN THIS SECTION

- Syntax | 5675
- Hierarchy Level | 5675
- Description | 5676
- Default | 5676
- Required Privilege Level | 5676
- Release Information | 5676

Syntax

```
dhcpv6-snooping-file (local_pathname | remote_URL);
    location local_pathname | remote_URL;
    timeout seconds;
    write-interval seconds;
}
```

Hierarchy Level

- For platforms with Enhanced Layer 2 Software (ELS):

```
[edit system processes  dhcp-service];
```

- For platforms without ELS:

```
[edit ethernet-switching-options secure-access-port]
```

Description

Ensure that IP-MAC bindings persist through switch reboots by specifying a local pathname or a remote URL for the storage location of the DHCPv6 snooping database file.

The remaining statements are explained separately. See CLI Explorer.

Default

The IP-MAC bindings in the DHCPv6 snooping database are not persistent. If the switch is rebooted, the bindings are lost.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D20.

Support at the `[edit ethernet-switching-options secure-access-port]` hierarchy level introduced in Junos OS 14.1X53-D10 for EX Series switches.

RELATED DOCUMENTATION

> *Configuring Persistent Bindings in the DHCP or DHCPv6 (ELS)*
>
> *Understanding DHCP Snooping (non-ELS)*

dhcpv6 (System Services)

IN THIS SECTION

- Syntax | 5677
- Hierarchy Level | 5681
- Description | 5681
- Options | 5681
- Required Privilege Level | 5681
- Release Information | 5681

Syntax

```
dhcpv6 {
    authentication {
        password password;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name;
            interface-name;
            logical-system-name;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix;
        }
    }
    dynamic-profile {
        profile-name;
        aggregate-clients {
            merge;
            replace;
        }
        junos-default-profile;
```

```
        use-primary dynamic-profile-name;
    }
    group group-name {
        authentication {
            password password;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name;
                interface-name;
                logical-system-name;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix;
            }
        }
        dynamic-profile {
            profile-name;
            aggregate-clients {
                merge;
                replace;
            }
            junos-default-profile;
            use-primary dynamic-profile;
        }
        interface interface-name {
            dynamic-profile {
                profile-name;
                aggregate-clients {
                    merge;
                    replace;
                }
                junos-default-profile;
                use-primary dynamic-profile-name;
            }
            exclude;
            overrides {
                delegated-pool pool-name;
                interface-client-limit number;
                process-inform {
```

```
                pool pool-name;
            }
            rapid-commit ;
        }
        service-profile service-profile-name
        trace ;
        upto interface-name;
    }
    liveness-detection {
        failure-action {
             clear-binding;
            clear-binding-if-interface-up;
            log-only;
        }
        method {
            bfd {
                detection-time {
                    threshold milliseconds;
                }
                holddown-interval interval;
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                session-mode (automatic | multihop | single-hop);
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
            version (0 | 1 | automatic);
        }
    }
    overrides {
        delegated-pool pool-name;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        rapid-commit ;
    }
    reconfigure {
        attempts number;
        clear-on-abort;
```

```
            strict;
            timeout number;
            token token-name;
            trigger {
                radius-disconnect;
            }
        }
        service-profile service-profile-name;
    }
    liveness-detection {
        failure-action {
             clear-binding;
            clear-binding-if-interface-up;
            log-only;
        }
        method {
            bfd {
                detection-time {
                    threshold milliseconds;
                }
                holddown-interval interval;
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                session-mode (automatic | multihop | single-hop);
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
            version (0 | 1 | automatic);
        }
    }
    overrides {
        delegated-pool pool-name;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        rapid-commit ;
    }
    reconfigure {
        attempts number;
```

```
        clear-on-abort;
        strict;
        timeout number;
        token token-name;
        trigger {
            radius-disconnect;
        }
    }
    service-profile service-profile-name;
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure DHCPv6 server to provide IPv6 addresses to clients.

NOTE: SRX Series Firewalls do not support client authentication.

Options

- `duplicate-clients-on-interface`—Allow duplicate clients on different interfaces in a subnet.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

DHCP Overview

diag-port-authentication

IN THIS SECTION

- Syntax | 5682
- Hierarchy Level | 5682
- Description | 5682
- Default | 5683
- Options | 5683
- Required Privilege Level | 5683
- Release Information | 5683

Syntax

```
diag-port-authentication (encrypted-password "password" | plain-text-password);
```

Hierarchy Level

```
[edit system]
```

Description

Configure a password for performing diagnostics on the router's System Control Board (SCB), System and Switch Board (SSB), Switching and Forwarding Module (SFM), or Forwarding Engine Board (FEB) port.

For routers that have more than one SSB, the same password is used for both SSBs.

> **NOTE**: Do not run diagnostics on the SCB, SSB, SFM, or FEB unless you have been instructed to do so by Customer Support personnel.

Default

No password is configured on the diagnostics port.

Options

`encrypted-password` *password*—Use MD5 or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.

You cannot configure a blank password for `encrypted-password` using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

`plain-text-password`—Use a plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password for each user.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Diagnostic Tools Overview
- show interfaces diagnostics optics

dial-options

IN THIS SECTION

- Syntax | 5684
- Hierarchy Level | 5684
- Description | 5684
- Options | 5685
- Required Privilege Level | 5685
- Release Information | 5685

Syntax

```
dial-options {
    ipsec-interface-id name;
    l2tp-interface-id name;
    (shared | dedicated);
}
```

Hierarchy Level

```
[edit interfaces sp-fpc/pic/port unit logical-unit-number],
[edit interfaces si-fpc/pic/port unit logical-unit-number],
[edit logical-systems logical-system-name interfaces sp-fpc/pic/port unit logical-unit-number],
[edit logical-systems logical-system-name interfaces si-fpc/pic/port unit logical-unit-number]
```

Description

Specify the options for configuring logical interfaces for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.

Options

`dedicated`—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.

`ipsec-interface-id` *name*—(M Series routers only) Interface identifier for group of dynamic peers. This identifier must be replicated at the `[edit access profile` *name* `client * ike]` hierarchy level.

`l2tp-interface-id` *name*—Interface identifier that must be replicated at the `[edit access profile` *name*`]` hierarchy level.

`shared`—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The `[edit ...si-...]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Configuring the Identifier for Logical Interfaces that Provide L2TP Services
- Configuring Dynamic Endpoints for IPsec Tunnels
- *Configuring Options for the LNS Inline Services Logical Interface*

dial-options (Dynamic Profiles)

IN THIS SECTION

- Syntax | 5686
- Hierarchy Level | 5686

- Description | 5686
- Options | 5686
- Required Privilege Level | 5687
- Release Information | 5687

Syntax

```
dial-options {
    ipsec-interface-id name;
    l2tp-interface-id name;
    (shared | dedicated);
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

Description

Specify the options for configuring logical interfaces in dynamic profiles for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.

Options

dedicated (LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.

ipsec-interface-id *name* Interface identifier for group of dynamic peers. This identifier must be replicated at the `[edit access profile name client * ike]` hierarchy level. This option is not currently supported for dynamic profiles.

l2tp-interface-id *name* (MX Series routers only) L2TP interface identifier that must be replicated at the `[edit access profile name]` hierarchy level.

shared (LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

dial-string

IN THIS SECTION

- Syntax | 5687
- Hierarchy Level | 5688
- Description | 5688
- Options | 5688
- Required Privilege Level | 5688
- Release Information | 5688

Syntax

```
dial-string [ dial-string-numbers ];
```

Hierarchy Level

```
[edit interfaces br-pim/0/port unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces br-pim/0/port unit logical-unit-number
dialer-options]
```

Description

On J Series Services Routers with ISDN interfaces, specify one or more ISDN dial strings used to reach a destination subnetwork.

Options

dial-string-numbers—One or more strings of numbers to call.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

dialer

IN THIS SECTION

- Syntax | 5689
- Hierarchy Level | 5689
- Description | 5689
- Options | 5689
- Required Privilege Level | 5689
- Release Information | 5689

Syntax

```
dialer filter-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family]
```

Description

Apply a dialer filter to an interface. To create the dialer filter, include the `dialer-filter` statement at the `[edit firewall filter family family]` hierarchy level.

Options

filter-name—Dialer filter name.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

dialer-options

IN THIS SECTION

Syntax | 5690

- Hierarchy Level | 5690
- Description | 5691
- Required Privilege Level | 5691
- Release Information | 5691

Syntax

```
dialer-options {
    activation-delay seconds;
    callback;
    callback-wait-period time;
    deactivation-delay seconds;
    dial-string [ dial-string-numbers ];
    idle-timeout seconds;
    incoming-map {
        caller caller-number | accept-all;
        initial-route-check seconds;
        load-interval seconds;
        load-threshold percent;
        pool pool-name;
        redial-delay time;
        self-recover-time
        watch-list {
            [ routes ];
        }
    }
}
```

Hierarchy Level

```
[edit interfaces umd0],
[edit interfaces dln unit logical-unit-number],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number]
```

Description

Specify the dialer options for configuring logical interfaces for group and user sessions.

The remaining statements are explained separately. See CLI Explorer.

You can use the new CLI option `self-recover-time` to configure the amount of time the `cl` interface waits to reconnect to a network after a disconnect occurs. In certain ISP networks, the modem disconnects and then reconnects after several seconds. The `self-recover-time` option provides enough time for the `cl` interface to reconnect instead of failing over to another `cl` interface immediately. If the reconnection attempt times out, then the connection fails over to another `cl` interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Services Interfaces Library for Routing Devices

dialin

IN THIS SECTION

- Syntax | 5692
- Hierarchy Level | 5692
- Description | 5692
- Options | 5692
- Required Privilege Level | 5692
- Release Information | 5692

Syntax

```
dialin (console | routable);
```

Hierarchy Level

```
[edit interfaces umd0 modem-options]
```

Description

For J Series Services Routers, configure a USB modem port to act as a dial-in console or WAN backup port.

Options

console Configure the USB modem port to operate as a dial-in console for management.

routable Configure the USB modem port to operate as a dial-in WAN backup interface.

- **Default:** console

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

diameter

IN THIS SECTION

- Syntax | 5693
- Hierarchy Level | 5694
- Description | 5694
- Options | 5694
- Required Privilege Level | 5694
- Release Information | 5694

Syntax

```
diameter {
    network-element element-name {
        dne-origin realm realm-name <host hostname>;
        forwarding {
            route dne-route-name {
                destination realm realm-name <host hostname>;
                function function-name <partition partition-name>;
                metric route-metric;
            }
        }
        function function-name;
        peer peer-name {
            priority priority-number;
        }
    }
    origin realm realm-name host hostname;
    peer peer-name {
        address ip-address;
        connect-actively {
            port port-number;
            transport transport-name;
        }
        logical-system logical-system-name  <routing-instance routing-instance-name> ;
        peer-origin realm realm-name host hostname;
```

```
        routing-instance routing-instance-name;
    }
    product-name name;
    transport transport-name {
        address;
        logical-system logical-system-name <routing-instance routing-instance-name>;
        routing-instance routing-instance-name;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure the Diameter base protocol for subscriber management.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

product-name *name*

Name of a product to advertise in the Capability Exchange Request and Capability-Exchange Answer messages to a Diameter peer. The name is included in the Diameter AVP Product-Name (269). If you change the product-name, the MX series router acting as a BNG disconnects all existing connections to Diameter peers and reconnects to those peers using the new product name.

- **Default:** JUNOS

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

`product-name` option added in Junos OS Release 18.3R1 on MX Series routers.

RELATED DOCUMENTATION

Configuring Diameter

diameter-instance (Diameter Applications)

IN THIS SECTION

- Syntax | 5695
- Hierarchy Level | 5695
- Description | 5695
- Options | 5696
- Required Privilege Level | 5696
- Release Information | 5696

Syntax

```
diameter-instance instance-name;
```

Hierarchy Level

```
[edit access gx-plus partition partition-name]
[edit access ocs partition partition-name],
[edit access pcrf partition partition-name]
```

Description

Specify the Diameter instance associated with the Gx-Plus, OCS, or PCRF partition.

Options

instance-name—Name of the Diameter instance. Currently, only `master` is supported.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support at the `[edit access ocs partition partition-name]` and `[edit access pcrf partition partition-name]` hierarchy levels introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring Gx-Plus

Configuring the Gx-Plus Partition

Configuring the OCS Partition

Understanding Gx Interactions Between the Router and the PCRF

Understanding Gy Interactions Between the Router and the OCS

diameter-instance (JSRC)

IN THIS SECTION

- Syntax | 5697
- Hierarchy Level | 5697
- Description | 5697
- Options | 5697
- Required Privilege Level | 5697
- Release Information | 5697

Syntax

```
diameter-instance instance-name
```

Hierarchy Level

```
[edit jsrc partition partition-name]
```

Description

Specify the Diameter instance associated with the JSRC partition.

Options

instance-name—Name of the Diameter instance. Currently, only `master` is supported.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

JSRC Configuration Overview

Configuring the JSRC Partition

diameter-profile (PCEF Profile)

IN THIS SECTION

- Syntax | 5698
- Hierarchy Level | 5698
- Description | 5698
- Options | 5698
- Required Privilege Level | 5699
- Release Information | 5699

Syntax

```
diameter-profile gx-profile-name;
```

Hierarchy Level

```
[edit unified-edge pcef profiles profile-name dynamic-policy-control],
[edit services pcef profiles profile-name dynamic-policy-control]
```

Description

Specify the Diameter Gx profile to use for the PCEF dynamic policy control profile. A PCEF profile with dynamic policy control must reference a defined Diameter Gx profile.

If you are using Junos OS Broadband Subscriber Management, specify the Diameter Gx profile at the `[edit services pcef profiles profile-name dynamic-policy-control]` hierarchy level.

If you are using Junos OS Subscriber Aware, specify the Diameter Gx profile at the `[edit unified-edge pcef profiles profile-name dynamic-policy-control]` hierarchy level.

Options

gx-profile-name Name of the Diameter Gx profile to use with this dynamic policy control profile.

Required Privilege Level

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef profiles profile-name dynamic-policy-control]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 18.2R1 on MX Series.

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Dynamic Policies

diameter (Subscriber Aware Policy Control)

IN THIS SECTION

- Syntax | 5700
- Hierarchy Level | 5701
- Description | 5701
- Required Privilege Level | 5701
- Release Information | 5701

Syntax

```
diameter {
    applications {
        pcc-gx {
            <maximum-pending-requests requests>;
        }
    }
    <firmware-revision version>;
    network-element element-name {
        function function-name;
        peer peer-name {
            priority priority-value;
            <timeout seconds>;
        }
    }
    origin {
        host hostname;
        realm realm-name;
    }
    peer peer-name {
        address ip-address;
        connect-actively {
            <capabilities-exchange-timeout seconds>;
            <port port-number>;
            <repeat-timeout seconds>;
            <retry-timeout seconds>;
            <timeout seconds>;
            transport transport-name;
        }
        <disconnect-peer-timeout seconds>;
        <incoming-queue> {
            size size;
        }
        <outgoing-queue> {
            <high-watermark high-watermark>;
            <low-watermark low-watermark>;
            size size;
        }
        <watchdog-timeout seconds>;
    }
<product-name product-name>;
```

```
    traceoptions {
        file diameter;
        flag flag;
        level all;
        peer {
            peer-name;
        }
    }
```

Hierarchy Level

```
[edit access]
```

Description

Configure the Diameter base protocol parameters for subscriber-aware dynamic policy control, so that Diameter applications can connect to remote peers. The Diameter base protocol configuration includes configuration of the endpoint origin, the transport layer connection, the remote peers, and the network elements.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Diameter Profiles

diameter (TDF Gateway)

IN THIS SECTION

- Syntax | 5702
- Hierarchy Level | 5702
- Description | 5702
- Required Privilege Level | 5703
- Release Information | 5703

Syntax

```
diameter {
    network-element {
        element-name {
            session-pics {
                group {
                    group-name {
                        [session-pic interface-name];
                    }
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit unified-edge tdf gateway-name]
```

Description

Configure the Diameter protocol parameters associated with Diameter bindings for this TDF gateway.

> **NOTE**: If you want to set up Diameter bindings for session PICs on the TDF gateway, contact Juniper Networks Professional Services for assistance.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Diameter Profiles

dictionary

IN THIS SECTION

- Syntax | 5704
- Hierarchy Level | 5704
- Description | 5704
- Options | 5704
- Required Privilege Level | 5704
- Release Information | 5704

Syntax

```
dictionary dictionary-path;
```

Hierarchy Level

[edit system services extensible-subscriber-services]

Description

Configure an XML-based dictionary file. The dictionary path is the complete path to the dictionary file and includes the dictionary filename. Extensible Subscriber Services Manager acts on the extensible-subscriber-service request on the basis of the services configured in the dictionary file. This configuration is mandatory.

Options

dictionary-path Path to the dictionary file. The complete path including the filename must not be more than 127 characters.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

show extensible-subscriber-services dictionary | **19853**

show extensible-subscriber-services dictionary attributes | **19859**

request services extensible-subscriber-services reload-dictionary | **14942**

show extensible-subscriber-services dictionary services | **19863**

Understanding the Dictionary File

diffserv-te

IN THIS SECTION

- Syntax | 5705
- Hierarchy Level | 5705
- Description | 5706
- Options | 5706
- Required Privilege Level | 5706
- Release Information | 5706

Syntax

```
diffserv-te {
    bandwidth-model {
        extended-mam;
        mam;
        rdm;
    }
    te-class-matrix {
        tenumber {
            priority priority;
            traffic-class {
                ctnumber priority priority;
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Specify properties for differentiated services in traffic engineering.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routers for DiffServ-Aware Traffic Engineering

direct-access (Commit Scripts)

IN THIS SECTION

- Syntax | 5707
- Hierarchy Level | 5707
- Description | 5707
- Required Privilege Level | 5707
- Release Information | 5707

Syntax

```
direct-access;
```

Hierarchy Level

```
[edit system scripts commit]
```

Description

Enable SLAX and XSLT commit scripts to read the input configuration directly from the database.

During a commit operation, the management process exports the post-inheritance candidate configuration in XML format and passes it as input to the commit script. If the configuration is large, the script driver might have trouble reading it into memory for commit script processing. When this occurs, you can configure the `direct-access` statement to enable the script driver to retrieve the *candidate configuration* directly from the configuration database. We recommend configuring the `direct-access` statement only if the configuration is large, because directly accessing the configuration data is processor-intensive and can affect system performance

> **NOTE**: Junos OS supports configuring the `direct-access` statement only when SLAX and XSLT commit scripts are configured. It does not support configuring the `direct-access` statement when Python commit scripts are also configured.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Processing Large Configurations Against Commit Scripts

direct-connect

IN THIS SECTION

- Syntax | 5708
- Hierarchy Level | 5708
- Description | 5708
- Required Privilege Level | 5709
- Release Information | 5709

Syntax

```
direct-connect;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number pppoe-
underlying-options],
[edit interfaces interface-name unit logical-unit-numberfamily pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
pppoe-underlying-options]
```

Description

Configure the router to ignore any DSL Forum VSAs that it receives in PPPoE control packets when the router is directly connected to CPE devices.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 13.3.

RELATED DOCUMENTATION

Ignoring DSL Forum VSAs from Directly Connected Devices

Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces

Configuring the PPPoE Family for an Underlying Interface

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

direction (RFC 2544 Benchmarking Tests)

IN THIS SECTION

- Syntax | 5709
- Junos OS Hierarchy Level | 5710
- Junos OS Evolved Hierarchy Level | 5710
- Description | 5710
- Options | 5710
- Required Privilege Level | 5710
- Release Information | 5710

Syntax

```
direction (egress | ingress);
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the direction of the interface on which the test must be run. This parameter is valid only for a `ccc` , a `ethernet-switching` or a `bridge` family. RFC2544 tests are supported only in the egress direction or the user-to-network interface (UNI) direction of an E-line or E-LAN service parameters in a bridge domain between two routers for unicast traffic. You cannot compute the NNI direction of Ethernet services between two routers for multicast or broadcast traffic.

Options

egress Run the test in the egress direction of the interface (network-to-network interface (NNI)). This option is applicable for a `ccc`, `bridge`, or `ethernet-switching` family.

ingress Run the test in the ingress direction of the interface (user-to-network interface (UNI)). You cannot configure this option for a `bridge` or `ethernet-switching` family.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

rfc2544-benchmarking | 15759

rfc2544 | 15757

RFC 2544-Based Benchmarking Tests Overview

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

Configuring an RFC 2544-Based Benchmarking Test

direction (services)

IN THIS SECTION

- Syntax | 5711
- Hierarchy Level | 5712
- Description | 5712
- Options | 5712
- Required Privilege Level | 5712
- Release Information | 5712

Syntax

```
direction (inbound | outbound | bidirectional) {
    protocol (ah | bundle | esp);
    spi spi-value;
    auxiliary-spi spi-value;
    authentication {
        algorithm (hmac-md5-96 | hmac-sha1-96);
        key (ascii-text key | hexadecimal key);
    }
    encryption {
        algorithm algorithm;
        key (ascii-text key | hexadecimal key);
    }
}
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then manual]
```

Description

Specify the direction in which manual SAs are applied.

Options

`bidirectional`—Apply the SA in both directions.

`inbound`—Apply the SA on inbound traffic.

`outbound`—Apply the SA on outbound traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

direction (Secure-channel)

IN THIS SECTION

- Syntax | 5713
- Hierarchy Level | 5713
- Description | 5713
- Default | 5714
- Options | 5714
- Required Privilege Level | 5714
- Release Information | 5714

Syntax

```
direction (inbound | outbound);
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel secure-channel-name]
```

Description

Configure whether the secure channel applies MACsec security to traffic entering or leaving an interface.

If you need to apply MACsec on traffic entering and leaving an interface, you need to create one secure channel to apply MACsec on incoming traffic and another secure channel to apply MACsec on outgoing traffic within the same connectivity association. When you associate the connectivity association with an interface, MACsec is applied on traffic entering and leaving that interface.

You only use this configuration option when you are configuring MACsec using static secure association keys (SAK) security mode. When you are configuring MACsec using static connectivity association keys

(CAK) security mode, two secure channels that are not user-configurable—one inbound secure channel and one outbound secure channel—are automatically created within the connectivity association.

Default

This statement does not have a default value.

If you have configured a secure channel to enable MACsec using static SAK security mode, you must specify whether the secure channel applies MACsec to traffic entering or leaving an interface. A candidate configuration that contains a secure channel that has not configured a direction cannot be committed.

Options

inbound Enable MACsec security on traffic entering the interface that has applied the secure channel.

outbound Enable MACsec security on traffic leaving the interface that has applied the secure channel.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

direction (Application Identification)

IN THIS SECTION

- Syntax | 5715
- Hierarchy Level | 5715
- Description | 5715
- Options | 5716
- Required Privilege Level | 5716
- Release Information | 5716

Syntax

```
direction {
    any;
    client-to-server;
    server-to-client;
}
```

Hierarchy Level

[edit services application-identification application *application-name* over *protocol-type* signature *name* member *name*]

Description

The connection direction of the packets to apply pattern matching. You can specify match patterns on both client to server and server to client while configuring custom application signatures.

Options

any The directions of packets are either from client-side to server-side or from server-side to client-side.

client-to-server The direction of packets is from client-side to server-side.

server-to-client The direction of packets is from server-side to client-side.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D40.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Understanding Junos OS Application Identification Custom Application Signatures

direction (forwarding-class-accounting)

IN THIS SECTION

- Syntax | 5717
- Hierarchy Level | 5717
- Description | 5717
- Options | 5717
- Required Privilege Level | 5717
- Release Information | 5717

Syntax

```
direction (ingress | egress | both)
```

Hierarchy Level

```
[edit interfaces interface-name forwarding-class-accounting],
[edit interfaces interface-name unit logical-unit-number forwarding-class-accounting]
```

Description

Specify the direction of traffic for which you want to apply counters. A single aggregate counter per forwarding class is used for flows. Forwarding class accounting applies to IPv4, IPv6, MPLS, Layer 2 and Other traffic.

Options

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3R3 in MX Series.

RELATED DOCUMENTATION

show class-of-service interface | 19056

clear interfaces statistics | 3651

direction (Junos-FIPS Software)

IN THIS SECTION

- Syntax | 5718
- Hierarchy Level | 5718
- Description | 5718
- Options | 5719
- Required Privilege Level | 5719
- Release Information | 5719

Syntax

```
direction (bidirectional | inbound | outbound) {
    protocol esp;
    spi spi-value;
    encryption {
        algorithm 3des-cbc;
        key ascii-text ascii-text-string;
    }
}
```

Hierarchy Level

```
[edit security ipsec internal security-association manual],
[edit security trusted-channel ipsec security-association manual]
```

Description

Establish a manual security association (SA) for internal Routing-Engine–to–Routing-Engine communication.

Options

`bidirectional`—Apply the same SA values in both directions between Routing Engines.

`inbound`—Apply these SA properties only to the inbound IPsec tunnel.

`outbound`—Apply these SA properties only to the outbound IPsec tunnel.

The remaining statements are explained separately.

Required Privilege Level

Crypto Officer—To view and add this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

direction (Junos OS)

IN THIS SECTION

- Syntax | 5719
- Hierarchy Level | 5720
- Description | 5720
- Options | 5720
- Required Privilege Level | 5720
- Release Information | 5720

Syntax

```
direction (inbound | outbound | bidirectional) {
    authentication {
        algorithm (hmac-md5-96 | hmac-sha1-96);
        key (ascii-text key | hexadecimal key);
```

```
    }
    auxiliary-spi auxiliary-spi-value;
    encryption {
        algorithm (des-cbc | 3des-cbc);
        key (ascii-text key | hexadecimal key);
    }
    protocol (ah | esp | bundle);
    spi spi-value;
}
```

Hierarchy Level

```
[edit security ipsec security-association sa-name manual]
```

Description

Define the direction of IPsec processing.

Options

inbound—Inbound SA—Define algorithms, keys, or security parameter index (SPI) values to decrypt and authenticate incoming traffic coming from the peer.

outbound—Outbound SA—Define algorithms, keys, or SPI values to decrypt and authenticate outbound traffic to the peer.

bidirectional—Bidirectional SA—Decrypt and authenticate the incoming and outgoing traffic using the same algorithm, keys, or SPI in both directions, unlike inbound and outbound SAs that use different attributes in both directions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Manual IPsec Security Associations for an ES PIC

Example: Using IPsec to Protect BGP Traffic

direction (MEP)

IN THIS SECTION

- Syntax | 5721
- Hierarchy Level | 5721
- Description | 5721
- Options | 5722
- Required Privilege Level | 5722
- Release Information | 5722

Syntax

```
direction (up | down);
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id]
```

Description

Configure the direction of the MEP.

Options

`up`—An UP MEP CCM is transmitted out of every logical interface which is part of the same bridging or vpls instance except for the interface configured on this MEP.

> **NOTE**: The `up` direction for MEP is not supported on T Series routers.

`down`—Down MEP CCMs are transmitted only out the interface configured on this MEP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring a MEP to Generate and Respond to CFM Protocol Messages

IEEE 802.1ag OAM Connectivity Fault Management Overview

direction (MX Series)

IN THIS SECTION

- Syntax | 5723
- Hierarchy Level | 5723
- Description | 5723
- Default | 5723
- Options | 5724
- Required Privilege Level | 5724

Release Information | 5724

Syntax

```
direction (inbound | outbound);
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel
secure-channel-name]
```

Description

Configure whether the secure channel applies MACsec security to traffic entering or leaving an interface.

If you need to apply MACsec on traffic entering and leaving an interface, you need to create one secure channel to apply MACsec on incoming traffic and another secure channel to apply MACsec on outgoing traffic within the same connectivity association. When you associate the connectivity association with an interface, MACsec is applied on traffic entering and leaving that interface.

You only use this configuration option when you are configuring MACsec using static secure association keys (SAK) security mode. When you are configuring MACsec using static connectivity association keys (CAK) security mode, two secure channels that are not user-configurable—one inbound secure channel and one outbound secure channel—are automatically created within the connectivity association.

Default

This statement does not have a default value.

If you have configured a secure channel to enable MACsec using static SAK security mode, you must specify whether the secure channel applies MACsec to traffic entering or leaving an interface. A candidate configuration that contains a secure channel that has not configured a direction cannot be committed.

Options

inbound Enable MACsec security on traffic entering the interface that has applied the secure channel.

outbound Enable MACsec security on traffic leaving the interface that has applied the secure channel.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

direction (Security Custom Attack)

IN THIS SECTION

- Syntax | 5725
- Hierarchy Level | 5725
- Description | 5725
- Options | 5725
- Required Privilege Level | 5725
- Release Information | 5725

Syntax

```
direction (any | client-to-server | server-to-client);
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type anomaly],

[edit security idp custom-attack attack-name attack-type signature]
```

Description

Define the connection direction of the attack.

Options

- `any`—Detect the attack in either direction.
- `client-to-server`—Detect the attack only in client-to-server traffic.
- `server-to-client`—Detect the attack only in server-to-client traffic.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

direction (Security Dynamic Attack Group)

IN THIS SECTION

- Syntax | 5726
- Hierarchy Level | 5726
- Description | 5726
- Options | 5726
- Required Privilege Level | 5727
- Release Information | 5727

Syntax

```
direction {
    expression (and | or);
    values [any client-to-server exclude-any exclude-client-to-server exclude-server-to-client
server-to-client];
}
```

Hierarchy Level

```
[edit security idp dynamic-attack-group  dynamic-attack-group-name filters]
```

Description

Specify a direction filter to add predefined attacks to the dynamic group based on the direction specified in the attacks.

Options

`expression`—Boolean operators:

- `and`— If both the member name patterns match, the expression matches.

- `or`— If either of the member name patterns match, the expression matches.

`values`—Name of the direction filter. You can select from the following directions:

- `any`—Monitors traffic from client to server and server to client.
- `client-to-server`—Monitors traffic from client to server (most attacks occur over `client-to-server` connections) only.
- `exclude-any`—Allows traffic from client to server and server to client.
- `exclude-client-to-server`—Allows traffic from client to server only.
- `exclude-server-to-client`—Allows traffic from server to client only.
- `server-to-client`—Monitors traffic from server to client only.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3. The `expression` option added in Junos OS Release 11.4.

direction (Service Data Flow Filters)

IN THIS SECTION

- Syntax | 5728
- Hierarchy Level | 5728
- Description | 5728
- Default | 5728
- Options | 5728
- Required Privilege Level | 5728
- Release Information | 5729

Syntax

```
direction (uplink | downlink | both);
```

Hierarchy Level

```
[edit unified-edge pcef flow-descriptions flow-identifier],
[edit services pcef flow-descriptions flow-identifier]
```

Description

Specify the direction in which service data flow (SDF) filters will detect service flow IP packets.

If you are using Junos OS Subscriber Aware, specify the direction at the `[edit unified-edge pcef flow-descriptions flow-identifier]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the direction at the `[edit services pcef flow-descriptions flow-identifier]` hierarchy level.

Default

If you do not configure the `direction` statement, the default direction is `both`.

Options

uplink SDF filters are applied in the uplink direction.

downlink SDF filters are applied in the downlink direction.

both SDF filters are applied in both the uplink and downlink directions.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef flow-descriptions flow-identifier]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

Configuring Service Data Flow Filters

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Understanding Application-Aware Policy Control for Subscriber Management

director-device (Aliases)

IN THIS SECTION

- Syntax | 5729
- Hierarchy Level | 5730
- Description | 5730
- Options | 5730
- Required Privilege Level | 5730
- Release Information | 5730

Syntax

```
director-device director-device-name {
    assigned-director-device-name;
}
```

Hierarchy Level

```
[edit fabric aliases]
```

Description

(QFabric systems only) Specify the mapping of user-defined names to QFabric system Director devices.

NOTE: The following rules apply to QFabric component alias naming:

- Alias names must use alphabetic (**A** through **Z** and **a** through **z**), numeric (**0** through **9**), or dash (**-**) characters.
- The maximum length of an alias name is 30 characters.
- Alias names are case sensitive. For example, **MY-NG-1** and **my-ng-1** refer to different components.
- You cannot use the reserved names `all`, `fabric`, or `director-group` as an alias name.

Options

director-device-name—Specify a user-defined name for a QFabric system Director device.

assigned-director-device-name—Specify the Director device identifier or name that has been provided. Identifiers are usually auto-generated by the Director software, and names are usually provided by the administrator of the default partition.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Aliases for the QFabric System

Understanding the Director Group

disable (Protocols)

IN THIS SECTION

- Syntax | 5731
- Hierarchy: disable (Protocols IGMP) | 5732
- Hierarchy: disable (Protocols SAP) | 5732
- Hierarchy: disable (Protocols MSDP) | 5732
- Hierarchy: disable (Protocols MLD) | 5733
- disable (PIM Graceful Restart) | 5733
- Hierarchy: disable (Protocols DVMRP) | 5733
- Hierarchy: disable (PIM) | 5733
- disable (Multicast Snooping) | 5734
- Hierarchy: disable (Protocols MLD Snooping) | 5734
- disable (IGMP Snooping) | 5734
- disable (MLD Snooping) | 5734
- Hierarchy: disable (IGMP Snooping) | 5734
- Description | 5734
- Default | 5735
- Required Privilege Level | 5735
- Release Information | 5735

Syntax

```
disable;
```

Hierarchy: disable (Protocols IGMP)

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Hierarchy: disable (Protocols SAP)

```
[edit logical-systems logical-system-name protocols sap],
[edit protocols sap]
```

Hierarchy: disable (Protocols MSDP)

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name],
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
peer address],
[edit protocols msdp],
[edit protocols msdp group group-name],
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit routing-instances routing-instance-name protocols msdp],
[edit routing-instances routing-instance-name protocols msdp group group-name],
[edit routing-instances routing-instance-name protocols msdp group group-name peer address],
[edit routing-instances routing-instance-name protocols msdp peer address]
```

Hierarchy: disable (Protocols MLD)

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

disable (PIM Graceful Restart)

```
[edit logical-systems logical-system-name protocols pim graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
graceful-restart],
[edit protocols pim graceful-restart],
[edit routing-instances routing-instance-name protocols pim graceful-restart]
```

Hierarchy: disable (Protocols DVMRP)

```
[edit logical-systems logical-system-name protocols dvmrp],
[edit logical-systems logical-system-name protocols dvmrp interface interface-name],
[edit protocols dvmrp],
[edit protocols dvmrp interface interface-name]
```

Hierarchy: disable (PIM)

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name protocols pim family  (inet | inet6)],
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name],
[edit logical-systems logical-system-name protocols pim rp local family  (inet | inet6)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp local family  (inet | inet6)],
[edit protocols pim],
[edit protocols pim family  (inet | inet6)],
[edit protocols pim interface (Protocols PIM)  interface-name],
[edit protocols pim rp local family  (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim],
[edit routing-instances routing-instance-name protocols pim family  (inet | inet6)],
```

```
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name],
[edit routing-instances routing-instance-name protocols pim mvpn family  (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim rp local family  (inet | inet6)]
```

disable (Multicast Snooping)

```
[edit multicast-snooping-options graceful-restart]
```

Hierarchy: disable (Protocols MLD Snooping)

```
[edit protocols mld-snooping vlan (all | vlan-name)]
```

disable (IGMP Snooping)

```
[edit protocols igmp-snooping vlan (all | vlan-name)]
```

disable (MLD Snooping)

```
[edit protocols mld-snooping vlan (all | vlan-name)]
```

Hierarchy: disable (IGMP Snooping)

```
[edit protocols igmp-snooping vlan vlan-name]
```

Description

`disable (Protocols IGMP)`disables IGMP on the system.

`disable (Protocols SAP)`explicitly disables SAP.

`disable (Protocols MSDP)` explicitly disables MSDP.

`disable (Protocols MLD)`disables MLD on the system.

`disable (PIM Graceful Restart)`explicitly disables PIM sparse mode graceful restart.

`disable (Protocols DVMRP)`explicitly disables DVMRP on the system or on an interface.

`disable (PIM)`explicitly disable PIM at the protocol, interface or family hierarchy levels.

`disable (Multicast Snooping)`explicitly disables graceful restart for multicast snooping.

`disable (Protocols MLD Snooping)`disables MLD snooping on the VLAN. Multicast traffic will be flooded to all interfaces in the VLAN except the source interface.

`disable (IGMP Snooping)`disables IGMP snooping on the VLAN. Multicast traffic will be flooded to all interfaces on the VLAN except the source interface.

`disable (IGMP Snooping)`disables IGMP snooping on all interfaces in a VLAN.

Default

If you do not include this statement, MLD snooping is enabled on all interfaces in the VLAN.

If you do not include this statement in the configuration for a VLAN, IGMP snooping is enabled on the VLAN.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

`address (Local RPs)` and `disable (Protocols IGMP)` and `disable (Protocols SAP)` and `disable (PIM)` and `disable (Protocols MLD)` and `disable (Protocols MSDP)` introduced before Junos OS Release 7.4.

`address (Local RPs)` and `disable (Protocols IGMP)`introduced in Junos OS Release 9.0 for EX Series switches.

`disable (IGMP Snooping)`introduced in Junos OS Release 9.2 for EX Series switches.

disable statement extended to the **[family]** hierarchy level of disable (PIM) in Junos OS Release 9.6.

`disable (IGMP Snooping)` introduced in Junos OS Release 11.1 for the QFX Series.

`disable (MLD Snooping)` introduced in Junos OS Release 18.1R1 for the SRX1500 devices.

`address (Local RPs)` introduced in Junos OS Release 11.3 for the QFX Series.

`disable (Protocols IGMP)` and `disable (Protocols MLD Snooping)` and `disable (Protocols MSDP)`introduced in Junos OS Release 12.1 for the QFX Series.

`disable (Protocols MLD Snooping)`introduced in Junos OS Release 12.1 for EX Series switches.

`disable (Multicast Snooping)` introduced in Junos OS Release 12.3.

`address (Local RPs)` and `disable (Protocols MSDP)` introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

> **NOTE**: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

RELATED DOCUMENTATION

mld-snooping | **10783**

Disabling IGMP

Disabling MLD

Disabling PIM

family (Protocols PIM) | **6907**

Configuring the Session Announcement Protocol

Configuring PIM Sparse Mode Graceful Restart

Example: Configuring Multicast Snooping

Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)

show mld-snooping vlans | **22100**

disable

IN THIS SECTION

- Syntax | **5737**
- Hierarchy Level | **5737**
- Description | **5737**
- Options | **5737**
- Required Privilege Level | **5737**

Release Information | 5737

Syntax

```
disable service-name;
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name]
```

Description

Disable the service name of the subscriber profile.

Options

service-name—Name of the disabled service.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

disable (Routing Graceful Restart)

IN THIS SECTION

- Syntax | 5738
- Hierarchy Level | 5738
- Description | 5738
- Required Privilege Level | 5739
- Release Information | 5739

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (bgp | isis | ldp | ospf | ospf3 | pim | rip
| ripng | rsvp) graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (bgp
| ldp | ospf | ospf3 | pim) graceful-restart],
[edit protocols (bgp | esis | isis | ospf | ospf3 | ldp | pim | rip | ripng | rsvp) graceful-
restart],
[edit protocols bgp group group-name graceful-restart],
[edit protocols bgp group group-name neighbor ip-address graceful-restart],
[edit routing-instances routing-instance-name protocols (bgp | ldp | ospf | ospf3 | pim) graceful-
restart],
[edit routing-instances routing-instance-name routing-options graceful-restart],
[edit routing-options graceful-restart]
```

Description

Disable graceful restart.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Enabling Graceful Restart*
- *Configuring Routing Protocols Graceful Restart*
- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- Graceful Restart Configuration Statements
- *Configuring Graceful Restart for QFabric Systems*

disable (STP)

IN THIS SECTION

- Syntax | 5740
- Hierarchy Level | 5740
- Description | 5741
- Default | 5741
- Required Privilege Level | 5741
- Release Information | 5741

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mstp],
[edit logical-systems logical-system-name routing-instances routing-instance-name
protocols mstp],
```

```
[edit protocols mpls],
[edit protocols mpls interface (all | interface-name)],
```

```
[edit protocols mstp],
[edit protocols mstp (all | interface-name)],
[edit protocols mstp (all | interface-name) arp-on-stp],
[edit protocols mstp msti msti-id vlan (vlan-id | vlan-name) interface (all | interface-name)],
[edit protocols mstp msti msti-id vlan (vlan-id | vlan-name) interface interface-name arp-on-stp],
[edit protocols rstp],
[edit protocols rstp interface (all | interface-name)],
[edit protocols rstp interface (all | interface-name) arp-on-stp],
[edit protocols stp],
[edit protocols stp interface (all | interface-name)],
[edit protocols stp interface (all | interface-name) arp-on-stp],
[edit protocols vstp],
[edit protocols vstp vlan vlan-id interface (all | interface-name)]
[edit protocols vstp vlan vlan-id interface (all | interface-name) arp-on-stp]
[edit protocols mstp interface interface-name,
[edit protocols mstp msti msti-id vlan (all vlan-id | vlan-name) interface interface-name],
```

```
[edit protocols rstp interface interface-name],
[edit protocols vstp vlan vlan-id interface interface-name]
```

```
[edit routing-instances routing-instance-name protocols mstp]
```

```
[edit ethernet-switching-options bpdu-block interface]
```

Description

Disable the entire MSTP, RSTP, or VSTP instance on a specific interface.

NOTE: You cannot disable spanning tree parameters for all interfaces.

Default

Not enabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

Statement updated in Junos OS Release 15.1 for EX Series switches.

RELATED DOCUMENTATION

Understanding RSTP

Configuring Multiple Spanning Tree Protocol

Configuring RSTP on EX Series Switches (CLI Procedure)

Configuring MSTP on Switches

Disabling MSTP

Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP

interface (BPDU Blocking) | **8842**

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

disable-4byte-as

IN THIS SECTION

- Syntax | **5742**
- Hierarchy Level | **5742**
- Description | **5743**
- Required Privilege Level | **5743**
- Release Information | **5743**

Syntax

```
disable-4byte-as;
```

Hierarchy Level

```
[edit protocols bgp],
[edit protocols bgp group group-name ],
[edit protocols bgp group group-name neighbor ip-address  ]
```

Description

Use to enable a BGP peer that uses a 4-Byte to interact with another BGP peer old speaker that uses 2-Byte.

> **NOTE**: We recommend that you configure `disable-4byte-as` only if the BGP peer does not support or ignores the capability advertisement of 4byte-as, and brings up the session as a 2byte AS.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 19.3.

disable-action

IN THIS SECTION

- Syntax | 5743
- Hierarchy Level | 5744
- Description | 5744
- Required Privilege Level | 5744
- Release Information | 5744

Syntax

```
disable-action;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

(MX Series routers only) Disable the MAC move action feature globally. MAC move detection configuration does exist, but the action is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring MAC Move Parameters

disable-all-instances

IN THIS SECTION

- Syntax | 5745
- Hierarchy Level | 5745
- Description | 5745
- Usage Guidelines | 5745
- Required Privilege Level | 5745
- Release Information | 5745

Syntax

```
disable-all-instances;
```

Hierarchy Level

```
[edit forwarding-options port-mirror]
```

Description

Disable all port mirroring instances globally.

Usage Guidelines

See Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

disable (Application Tracking)

IN THIS SECTION

- Syntax | 5746
- Hierarchy Level | 5746
- Description | 5746
- Required Privilege Level | 5746

Release Information | 5746

Syntax

```
disable;
```

Hierarchy Level

```
[edit security application-tracking]
```

Description

Disable application tracking on a device without deleting the zone configuration.

Application tracking is enabled by default. If application tracking has been previously disabled and you want to reenable it, delete the configuration statement that specifies disabling of application tracking as shown in the following statement:

```
[edit]
user@host# delete security application-tracking disable
```

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring AppTrack

disable (BGP Graceful Restart)

IN THIS SECTION

- Syntax | 5747
- Hierarchy Level | 5747
- Description | 5747
- Required Privilege Level | 5748
- Release Information | 5748

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
graceful-restart],
[edit protocols bgp graceful-restart],
[edit protocols bgp group group-name graceful-restart],
[edit protocols bgp group group-name neighbor address graceful-restart]
```

Description

Disable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition.

NOTE: When you disable graceful restart at one level in the configuration statement hierarchy, it is also disabled at lower levels in the same hierarchy. For example, if you disable graceful restart at the `[edit protocols bgp group group-name]` hierarchy level, it is

disabled for all the peers in the group. Therefore, if you want to enable graceful restart for some peers in a group and disable it for others, enable graceful restart at the `[edit protocols bgp group group-name]` hierarchy level and disable graceful restart for each peer at the `[edit protocols bgp group group-name neighbor address]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

graceful-restart (Protocols BGP) | **7769**

restart-time (BGP Graceful Restart) | **15663**

stale-routes-time | **28478**

disable-calling-number-avp (L2TP LAC)

IN THIS SECTION

- Syntax | **5749**
- Hierarchy Level | **5749**
- Description | **5749**
- Required Privilege Level | **5749**
- Release Information | **5749**

Syntax

```
disable-calling-number-avp;
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Prevent the LAC from sending L2TP Calling Number AVP 22 in incoming-call request (ICRQ) packets to the LNS. By default, the LAC in an L2TP network generates this AVP from the Calling-Station-Id and sends it to the LNS.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP Subscriber Access Lines and Connection Speeds

disable (DCBX)

IN THIS SECTION

Syntax | 5750

- Hierarchy Level | 5750
- Description | 5750
- Default | 5750
- Required Privilege Level | 5750
- Release Information | 5750

Syntax

```
disable
```

Hierarchy Level

```
[edit protocols dcbx]

[edit protocols dcbx interface interface-name]
```

Description

Disable Data Center Bridging Capability Exchange protocol (DCBX) on one or more 10-Gigabit Ethernet interfaces.

Default

DCBX is enabled by default on all 10-Gigabit or higher Ethernet interfaces.

DCBX is enabled by default on all 10-Gigabit Ethernet interfaces on EX4500 CEE-enabled switches.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring DCBX Autonegotiation

Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)

Understanding DCB Features and Requirements

Understanding DCB Features and Requirements on EX Series Switches

disable (Dynamic IGMP)

IN THIS SECTION

- Syntax | 5751
- Hierarchy Level | 5751
- Description | 5751
- Required Privilege Level | 5752
- Release Information | 5752

Syntax

```
"disable:$junos-igmp-enable";
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name],
```

Description

Disable IGMP on the interface.

NOTE: Though the purpose of this statement is to disable IGMP on interfaces, under the `dynamic-profiles` hierarchy you can use this statement and an enable variable

(`disable:$junos-igmp-enable`) to ensure that IGMP is not disabled by a AAA-based authentication and management method (RADIUS).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Disabling IGMP

disable (Dynamic MLD)

IN THIS SECTION

- Syntax | 5752
- Hierarchy Level | 5753
- Description | 5753
- Required Privilege Level | 5753
- Release Information | 5753

Syntax

```
disable;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Description

Disable MLD on the dynamic interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Disabling MLD

disable (Extensible Subscriber Services Manager)

IN THIS SECTION

- Syntax | 5754
- Hierarchy Level | 5754
- Description | 5754
- Default | 5754
- Required Privilege Level | 5754
- Release Information | 5754

Syntax

```
disable
```

Hierarchy Level

`[edit system processes extensible-subscriber-services]`

Description

Disable the Extensible Subscriber Services Manager process when there is an extensible-subscriber-services configuration and the user wants to stop the process. The process is disabled by default when there is no configuration under the `[edit system services extensible-subscriber-services]` hierarchy level.

Default

The process is disabled by default.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

restart extensible-subscriber-services | **15656**

disable-failover-protocol (L2TP)

IN THIS SECTION

- Syntax | 5755
- Hierarchy Level | 5755
- Description | 5755
- Required Privilege Level | 5756
- Release Information | 5756

Syntax

```
disable-failover-protocol;
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Configure the LAC or LNS to use only the silent failover method when resynchronizing with its peer in the event of a control plane failover. This statement prevents the default behavior, where the LAC first attempts to negotiate the failover protocol when it establishes control connections with the peer. If the remote peer does not support the failover protocol, then the LAC falls back on the silent failover method. Including this configuration is useful when the peers configured for silent failover or incorrectly negotiate use of the failover protocol even though they do not support it.

BEST PRACTICE: We recommend that you include this statement on both the LAC and LNS to prevent the use of failover protocol. When failover protocol is used, the nonfailed peer (LAC or LNS) keeps the tunnel open with the failed peer, in case the failed peer is able to recover from the failure and resynchronize with the nonfailed peer. This behavior keeps the tunnel up and the subscribers logged in while traffic is not flowing, preventing service level agreements from being met.

Starting in Junos OS Releases 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, the `disable-failover-protocol` statement is deprecated and no longer needs to be used. The default failover resynchronization method is changed to silent failover, rather than the previous default method of failover-protocol-fall-back-to-silent-failover. The new default method conforms to our recommendation to use silent failover. Consequently, there is no need to disable the failover protocol. Configurations that include this statement are still supported when you upgrade to a release in which it is deprecated; The CLI informs you of the deprecation if the statement is included.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Statement deprecated in Junos OS Release 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1R6	Starting in Junos OS Releases 15.1R6, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, the `disable-failover-protocol` statement is deprecated and no longer needs to be used.

RELATED DOCUMENTATION

Configuring the L2TP Peer Resynchronization Method

disable (Forwarding Options)

IN THIS SECTION

Syntax | 5757

- Hierarchy Level | 5757
- Description | 5757
- Required Privilege Level | 5757
- Release Information | 5758

Syntax

```
disable;
```

Hierarchy Level

```
[edit forwarding-options port-mirror],
[edit forwarding-options port-mirror instance instance-name],
[edit forwarding-options sampling],
[edit forwarding-options sampling instance instance-name],
[edit forwarding-options sampling family  (inet | inet6 | mpls | vpls) ],
[edit forwarding-options sampling family (inet | inet6 | mpls | vpls) output file]
```

Description

Disable traffic accounting, port mirroring, or sampling.

NOTE: The `disable` statement at the `[edit forwarding-options sampling]` hierarchy level disables only Routing Engine-based sampling. To disable PIC-based sampling and inline sampling, include the `disable` statement at the `[edit forwarding-options sampling instance instance-name]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement added to `port-mirror` hierarchy in Junos OS Release 9.6.

> **NOTE**: Beginning in Junos OS Release 15.1F5 and later 15.1 releases and Junos OS Release 16.1 and later, the `disable` option has been deprecated at the `forwarding-options sampling instance` *`instance-name`* `family (inet | inet6 | mpls)` hierarchy level on PTX3000 Series routers. When configured, the option does not take effect, so packets continue to be sampled. Instead of the `disable` option, use the `deactivate forwarding-options sampling instance` *`instance-name`* `family (inet | inet6 | mpls)` command to prevent sampling.

RELATED DOCUMENTATION

Disabling Traffic Sampling

Configuring Traffic Sampling on MX, M and T Series Routers

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

disable-fpc (DDoS)

IN THIS SECTION

- Syntax | 5758
- Hierarchy Level | 5759
- Description | 5759
- Required Privilege Level | 5759
- Release Information | 5759

Syntax

```
disable-fpc;
```

Hierarchy Level

```
[edit system ddos-protection global],
[edit system ddos-protection protocols protocol-group (aggregate | packet-type)],
[edit system ddos-protection protocols protocol-group (aggregate | packet-type) fpc slot-number]
```

Description

(MX Series routers with only MPCs, T4000 routers with only FPC5s, PTX Series routers, EX9200 switches, or QFX Series switches) Disable DDoS policers for debugging purposes on the card in the specified slot for a particular packet type within a protocol group, on all cards for a particular packet type within a protocol group, or globally on all cards and for all packet types in all protocols. This statement does not affect the state of the Routing Engine policers.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support at the `[edit system ddos-protection protocols protocol-group (aggregate | packet-type)]` hierarchy level introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Disabling Control Plane DDoS Protection Policers and Logging Globally

Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers

disable-global-timeout-override

IN THIS SECTION

- Syntax | 5760
- Hierarchy Level | 5760
- Description | 5760
- Required Privilege Level | 5760
- Release Information | 5760

Syntax

```
disable-global-timeout-override;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
[edit services service-set service-set-name service-set-options
```

Description

Disallow overriding a global inactivity or session timeout.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support added in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480, and MX960 routers.

RELATED DOCUMENTATION

Defining an Application Identification

disable (GMPLS)

IN THIS SECTION

- Syntax | 5761
- Hierarchy Level | 5761
- Description | 5761
- Default | 5762
- Required Privilege Level | 5762
- Release Information | 5762

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],
[edit protocols link-management te-link te-link-name]
```

Description

Disable a traffic engineering link.

Default

The configured object is enabled (operational) unless explicitly disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Disabling the Traffic Engineering Link for LMP Peers

disable-grant-bypass

IN THIS SECTION

- Syntax | 5762
- Hierarchy Level | 5763
- Description | 5763
- Default | 5763
- Required Privilege Level | 5763
- Release Information | 5763

Syntax

```
disable-grant-bypass;
```

Hierarchy Level

```
[edit chassis fabric]
```

Description

Disable the fabric grant bypass feature. This feature is used for communication between Packet Forwarding Engines. For instance, when a Packet Forwarding Engine wants to send a packet to another Packet Forwarding Engine (on the same MPC or on a different MPC), a request is sent to the Packet Forwarding Engine across the fabric plane. Only after the request is granted, can the source Packet Forwarding Engine send the packet to the destination Packet Forwarding Engine. Disabling the default behavior controls congestion and thus improves system behavior and performance on MX2010 and MX2020 routers.

> **NOTE**: After disabling fabric grant bypass feature on the MX2010 and MX2020, you must reboot the router for the changes to take effect. MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFP) do not power on after you disable the fabric grant bypass feature and reboot the router.

Default

Enabled for all MPCs on MX2010 and MX2020 routers with Switch Fabric Boards (SFBs). On MX2010 and MX2020 routers with the Switch Fabric Board SFB2, this feature is enabled only on the MPCs, MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFP), and disabled on all other MPCs. The feature is disabled on routers with Switch Control Boards (SCBs).

Required Privilege Level

view

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding Fabric Grant Bypass

Disabling Fabric Grant Bypass to Control Congestion and Improve Performance

Re-Enabling Fabric Grant Bypass

disable (IGMP Snooping)

IN THIS SECTION

- Syntax | 5764
- Hierarchy Level | 5764
- Description | 5764
- Required Privilege Level | 5765
- Release Information | 5765

Syntax

```
disable;
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan vlan-name]
```

Description

Disable IGMP snooping on the VLAN. Without IGMP snooping, multicast traffic will be flooded to all interfaces on the VLAN except the source interface.

This option is available only on legacy switches that do not support the Enhanced Layer 2 Software (ELS) configuration style. On these switches, IGMP snooping is enabled by default on all VLANs, and this statement includes the `disable` option if you want to disable IGMP snooping selectively on some VLANs or disable it on all VLANs.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping on Switches

Configuring IGMP Snooping on Switches

disable (Interface)

IN THIS SECTION

- Syntax | 5765
- Hierarchy Level | 5766
- Description | 5766
- Required Privilege Level | 5767
- Release Information | 5767

Syntax

```
disable;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Disable a physical or a logical interface, effectively unconfiguring it.

CAUTION:

- Dynamic subscribers and logical interfaces use physical interfaces for connection to the network. The Junos OS allows you to set the interface to disable and commit the change while dynamic subscribers and logical interfaces are still active. This action results in the loss of all subscriber connections on the interface. Use care when disabling interfaces.
- If aggregated SONET links are configured between a T1600 router and a T4000 router, interface traffic is disrupted when you disable the physical interface configured on the T1600 router. If you want to remove the interface, we recommend that you deactivate the interface instead of disabling it.

NOTE:

- When you use the `disable` statement at the `[edit interfaces]` hierarchy level, depending on the PIC type, the interface might or might not turn off the laser. Older PIC transceivers do not support turning off the laser, but newer Gigabit Ethernet (GE) PICs with SFP and XFP transceivers and ATM MIC with SFP do support it and the laser will be turned off when the interface is disabled. If the ATM MIC with SFP is part of an APS group, then the laser will not be turned off when you use the `disable` statement at the `[edit interfaces]` hierarchy level..
- When you disable or deactivate an interface, then all the references made to the deactivated interface must be removed from the routing instance.
- For abstracted fabric interfaces, the `disable` command disables AF interface on the local GNF only.

LASER WARNING: Do not stare into the laser beam or view it directly with optical instruments even if the interface has been disabled.

Starting in Junos OS Release 22.1R1, you can use the `disable` statement under `[edit interfaces interface-name]` and `[edit interfaces interface-name unit logical-unit-number]` hierarchy level to disable an IFD and IFL in PS interfaces on MX Series routers.

If you disable a PS IFD, then the IFD will be admin down and IFLs will do link down. If you disable a PS IFL, then the IFL will be admin down and link up.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

Disabling a Physical Interface

Disabling a Logical Interface

disable (Link Protection)

IN THIS SECTION

- Syntax | 5768
- Hierarchy Level | 5768
- Description | 5768
- Required Privilege Level | 5768

Release Information | 5768

Syntax

```
disable;
```

Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
```

Description

Disable LACP link protection on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring LACP for Aggregated Ethernet Interfaces

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

disable-logging (DDoS)

IN THIS SECTION

- Syntax | 5769
- Hierarchy Level | 5769
- Description | 5769
- Required Privilege Level | 5769
- Release Information | 5769

Syntax

```
disable-logging;
```

Hierarchy Level

```
[edit system ddos-protection global],
[edit system ddos-protection protocols protocol-group (aggregate | packet-type)]
```

Description

Disable device-wide logging of all DDoS violation and flow detection events globally. Disable only logging of events other than flow detection culprit flow events for a particular packet type or for a protocol group. Typically used for debugging purposes.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support at the `[edit system ddos-protection protocols protocol-group (aggregate | packet-type)]` hierarchy level introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Disabling Control Plane DDoS Protection Policers and Logging Globally

Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers

Disabling Automatic Logging of Culprit Flow Events for a Packet Type

disable (Long-Lived Graceful Restart for BGP Restarter)

IN THIS SECTION

- Syntax | 5770
- Hierarchy Level | 5770
- Description | 5771
- Required Privilege Level | 5772
- Release Information | 5772

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family (l2vpn | route-target | inet)
(labeled-unicast | flow) graceful-restart long-lived restarter],
[edit logical-systems logical-system-name protocols bgp group group-name family (l2vpn | route-
target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
(l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
```

```
family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived
restarter],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart
long-lived restarter],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family (l2vpn | route-target | inet) (labeled-unicast | flow)
graceful-restart long-lived restarter],
[edit routing-instances routing-instance-name protocols bgp family (l2vpn | route-target | inet)
(labeled-unicast | flow) graceful-restart long-lived restarter],
[edit routing-instances routing-instance-name protocols bgp group group-name family (l2vpn |
route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived
restarter],
[edit routing-instances routing-instance-name protocols bgp family (l2vpn | route-target | inet)
(labeled-unicast | flow) graceful-restart long-lived restarter],
[edit routing-instances routing-instance-name protocols bgp group group-name family (l2vpn |
route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived
restarter]
```

NOTE: Each routing table is identified by the protocol family or address family indicator (AFI) and a subsequent address family identifier (SAFI). The AFI parameter can be one of the `(l2vpn | inet | route-target)` protocols and the SAFI parameter can be either of the `(flow | labeled-unicast)` protocols for inet family and one of the `(auto-discovery-mspw | auto-discovery-only | signaling)` protcols for L2VPN family..

Configuring LLGR does not require that BGP graceful restart also be configured. The long-lived-graceful-restart section is visible only for families l2vpn, inet labeled-unicast, inet flow and route-target. It is prohibited for inet-mvpn, inet6-mvpn and inet-mdt. It is hidden for other families.

Description

Disable the long-lived graceful restart capability for BGP sessions on the restarting router. A hidden `enable` attribute can be used to override an inherited disable attribute. Configuring graceful-restart long-lived restarter at the neighbor level (when it is not configured at the containing group level or globally) causes an internal group to be split.

When LLGR restarter is enabled or disabled for a family or the stale- time is changed, the session is reset so that the new capability can be sent to the neighbor.

The stanzas in the per-family graceful-restart long-lived restarter configuration section enables LLGR restarter mode negotiation for BGP globally, or for a group or neighbor. The values are inherited by groups from the global configuration, and by neighbors from the group configuration. The disable attribute is used to override configuration inherited from a higher level. It does not disable LLGR receiver mode; you must disable LLGR receiver mode explicitly for all families as necessary.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

disable-mlppp-inner-ppp-pfc

IN THIS SECTION

- Syntax | 5773
- Hierarchy Level | 5773
- Description | 5773
- Required Privilege Level | 5773
- Release Information | 5773

Syntax

```
disable-mlppp-inner-ppp-pfc;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For MLPPP interfaces only, disable compression of the inner PPP header in the MLPPP payload. By default, compression is enabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring Link and Multilink Services Logical Interfaces

disable (Multicast Load Balancing)

IN THIS SECTION

Syntax | 5774

- Hierarchy Level | 5774
- Description | 5774
- Required Privilege Level | 5774
- Release Information | 5774

Syntax

```
disable;
```

Hierarchy Level

```
[edit chassis multicast-loadbalance]
```

Description

(EX8200 switches only) Disable multicast load balancing. After you configure this statement and commit it, multicast load balancing no longer balances traffic on aggregated 10-gigabit Ethernet interfaces that are currently configured as well as those that are configured later.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2 for EX Series switches.

RELATED DOCUMENTATION

> Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure)

disable (Multicast Snooping)

IN THIS SECTION

- Syntax | 5775
- Hierarchy Level | 5775
- Description | 5775
- Required Privilege Level | 5775
- Release Information | 5775

Syntax

```
disable;
```

Hierarchy Level

```
[edit multicast-snooping-options graceful-restart]
```

Description

Explicitly disable graceful restart for multicast snooping.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring Multicast Snooping

disable (MVRP)

IN THIS SECTION

- Syntax | 5776
- Hierarchy Level | 5776
- Description | 5776
- Default | 5776
- Required Privilege Level | 5777
- Release Information | 5777

Syntax

```
disable;
```

Hierarchy Level

```
[edit protocols mvrp],

[edit protocols mvrp interface(all | interface-name)]
```

Description

Disable the MVRP configuration on the interface.

Default

MVRP is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring Multiple VLAN Registration Protocol (MVRP) on Switches

disable-natt (Services IPsec VPN)

IN THIS SECTION

- Syntax | 5777
- Hierarchy Level | 5777
- Description | 5778
- Required Privilege Level | 5778
- Release Information | 5778

Syntax

```
disable-natt;
```

Hierarchy Level

```
[edit services ipsec-vpn]
```

Description

Before Junos OS Release 17.4R1, Network Address Translation-Traversal (NAT-T) is not supported for the Junos VPN Site Secure suite of IPsec features on the MX Series routers. In Junos OS releases before 17.4R1, disable NAT-traversal (NAT-T) when a NAT device is present between two IPsec gateways to cause the Encapsulating Security Payload (ESP) protocol to be used for encapsulation.

In traditional network deployments, IPsec does not work when packets traverse across a device that is configured for network address translation (NAT) or network address port translation (NAPT) for translating packets, IPsec does not work when either one of the device or both the devices that terminate the IPsec tunnel is behind a NAT device. This behavior occurs because NAT checks the port information, which is not present for IPsec-protected traffic.

When NAT-T is configured, IPsec traffic is encapsulated using the UDP header and port information is provided for the NAT devices. By default, Junos OS detects whether either one of the IPsec tunnel is behind a NAT device and automatically switches to using NAT-T for the protected traffic. However, in certain cases, NAT-T support on MX Series routers running a Junos OS Release before 17.4R1 might not work as desired. Also, you might require NAT-traversal to be disabled if you are aware that the network uses IPsec-aware NAT.

To avoid problems with NAT-T on MX series routers, you can disable NAT-T. When you disable NAT-T, the NAT-T functionality is globally switched off. Also, even when a NAT device is present between the two IPsec gateways, only ESP encapsulation is used when you disable NAT-T.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Disabling NAT-T on MX Series Routers for Handling NAT with IPsec-Protected Packets

disable-notification-extensions (BGP Graceful Restart)

IN THIS SECTION

- Syntax | 5779
- Hierarchy Level | 5779
- Description | 5780
- Required Privilege Level | 5780
- Release Information | 5780

Syntax

```
disable-notification-extensions {
    omit-no-export;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
graceful-restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp graceful-
restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name graceful-restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name neighbor address graceful-restart],
[edit protocols bgp graceful-restart],
[edit protocols bgp group group-name graceful-restart],
[edit protocols bgp group group-name neighbor address graceful-restart]
```

Description

Disables the transmission of the N flag in the graceful restart capability negotiation, but in addition, it disables the new rules for invoking graceful restart receiver mode as specified in the IETF bgp-gr-notification draft, and disables the transmission of the Hard Reset subcode. The Hard Reset subcode is continued to be observed when received in a Notify or a Cease message. The BGP protocol sends a notification NOTIFICATION message and reset the peering session to handle the error condition. BGP graceful restart that permits the operational procedures to be performed when the BGP speaker receives a notification message.

You can define this setting at the `[edit protocols bgp graceful-restart]`, `[edit protocols bgp group group-name graceful-restart]`, or `[edit protocols bgp group group-name neighbor neighbor-address graceful-restart]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

disable-notification-flag (BGP Graceful Restart)

IN THIS SECTION

- Syntax | 5781
- Hierarchy Level | 5781
- Description | 5781

Required Privilege Level | 5782
Release Information | 5782

Syntax

```
disable-notification-flag {
    omit-no-export;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
graceful-restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp graceful-
restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name graceful-restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name neighbor address graceful-restart],
[edit protocols bgp graceful-restart],
[edit protocols bgp group group-name graceful-restart],
[edit protocols bgp group group-name neighbor address graceful-restart]
```

Description

Disables the transmission of the notification (N) flag in the graceful restart capability negotiation. The BGP protocol sends a notification NOTIFICATION message and reset the peering session to handle the error condition. BGP graceful restart that permits the operational procedures to be performed when the BGP speaker receives a notification message. This behavior permits the BGP speaker to avoid flapping reachability and continue forwarding while the BGP speaker restarts the session to handle errors detected in the BGP protocol. You can define this setting at the `[edit protocols bgp graceful-restart]`, `[edit protocols bgp group group-name graceful-restart]`, or `[edit protocols bgp group group-name neighbor neighbor-address graceful-restart]` hierarchy level.

Disables the transmission of the N flag in the graceful restart capability negotiation, but in addition, it disables the new rules for invoking graceful restart receiver mode as specified in the IETF bgp-gr-notification draft, and disables the transmission of the Hard Reset subcode. The Hard Reset subcode is continued to be observed when received in a Notify or a Cease message.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

disable (PIM)

IN THIS SECTION

- Syntax | 5782
- Hierarchy Level | 5783
- Description | 5783
- Required Privilege Level | 5783
- Release Information | 5783

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name protocols pim family (inet | inet6)],
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM) interface-name],
[edit logical-systems logical-system-name protocols pim rp local family (inet | inet6)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM) interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp local family (inet | inet6)],
[edit protocols pim],
[edit protocols pim family (inet | inet6)],
[edit protocols pim interface (Protocols PIM) interface-name],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim],
[edit routing-instances routing-instance-name protocols pim family (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM) interface-
name],
[edit routing-instances routing-instance-name protocols pim mvpn family (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim rp local family (inet | inet6)]
```

Description

Explicitly disable PIM at the protocol, interface or family hierarchy levels.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

disable statement extended to the **[family]** hierarchy level in Junos OS Release 9.6.

RELATED DOCUMENTATION

Disabling PIM

family (Protocols PIM) | 6907

disable (PoE)

IN THIS SECTION

- Syntax | 5784
- Hierarchy Level | 5784
- Description | 5784
- Default | 5785
- Required Privilege Level | 5785
- Release Information | 5785

Syntax

```
disable;
```

Hierarchy Level

```
[edit poe interface  (all | interface-name) ]
 [edit poe interface (all | interface-name) telemetries]
```

Description

Disables the PoE capabilities of the port. If PoE capabilities are disabled for a port, the port operates as a standard network access port. If the disable statement is specified after the telemetries statement, logging of PoE power consumption for the port is disabled. To disable monitoring and retain the stored interval and duration values for possible future use, you can specify the disable sub statement in the sub

stanza for telemetries. Similarly for retaining the port configuration but disabling the PoE feature on the port, disable can be used in sub stanza for interface.

Default

The PoE capabilities are automatically enabled when a PoE interface is set. Specifying the telemetries statement enables monitoring of PoE per-port power consumption.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

disable (Power over Ethernet)

IN THIS SECTION

- Syntax | 5785
- Hierarchy Level | 5786
- Description | 5786
- Required Privilege Level | 5786
- Release Information | 5786

Syntax

```
disable;
```

Hierarchy Level

```
[edit poe interface (all | all-extended | interface-name)],
[edit poe interface (all | all-extended | interface-name) telemetries],
[edit poe notification-control fpc slot-number]
```

Description

Disable a PoE interface, disable the collection of power consumption data for a PoE interface, or disable the generation of the PoE SNMP traps. The action of the `disable` statement depends on which statement it is used with:

- When used with **interface**—Disable the PoE capability of this interface. The interface operates as a standard network access interface, and power is no longer allocated to it from the PoE power budget. Although the PoE capability is disabled, the PoE configuration for the interface is retained. To reenable the PoE capability of this interface, delete the `disable` statement from the **interface** entry in the configuration.
- When used with **telemetries**—Disable the collection of PoE power consumption records for this interface. Any previously collected records are deleted. However, the **telemetries** configuration is retained, including the values for **interval** and **duration**. To reenable record collection, delete the `disable` statement from the **telemetries** entry in the configuration.
- When used with **notification-control**—Disable the generation of PoE SNMP traps. To reenable PoE traps, delete the `disable` statement from the **notification-control** entry in the configuration.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

all-extended option introduced in Junos OS Release 16.1R1.

Statement introduced in Junos OS Release 17.2R1 for a Junos Fusion Provider Edge.

RELATED DOCUMENTATION

- Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch
- Configuring PoE Interfaces on EX Series Switches
- *Configuring Power over Ethernet in a Junos Fusion*

disable-preceding-key

IN THIS SECTION

- Syntax | 5787
- Hierarchy Level | 5787
- Description | 5787
- Default | 5788
- Required Privilege Level | 5788
- Release Information | 5788

Syntax

```
disable-preceding-key;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Disable the preceding pre-shared key (PSK) so that the fallback PSK is used to establish a MACsec connection during a pre-shared-key change event.

When you enable MACsec using static CAK security mode, a preshared PSK is exchanged between the devices on each end of the point-to-point Ethernet link. The PSK must match across devices for a MACsec session to be established. If the primary PSK is changed on one device but not the other, the

mismatch is resolved by using the older primary PSK. This is a temporary key known as the preceding PSK, and is not configurable.

If a fallback PSK is configured, it will not take effect if the MACsec session is live with the preceding PSK. You can configure the `disable-preceding-key` statement so that the session immediately switches to using the fallback PSK if there is a change to the primary PSK.

Default

By default, the preceding PSK takes priority over the fallback PSK. Configure the `disable-preceding-key` to override this behavior.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Configuring MACsec with Fallback PSK

disable (Protocols BGP)

IN THIS SECTION

- Syntax | 5789
- Hierarchy Level | 5789
- Description | 5789
- Required Privilege Level | 5789
- Release Information | 5789

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit protocols bgp],
[edit routing-instances routing-instance-name protocols bgp]
```

Description

Disable BGP on the system.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

disable (Protocols LDP)

IN THIS SECTION

- Syntax | 5790
- Hierarchy Level | 5790
- Description | 5790
- Default | 5790
- Required Privilege Level | 5790

Release Information | 5791

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp graceful-restart],
[edit logical-systems logical-system-name protocols ldp interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp
interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options graceful-restart],
[edit protocols ldp graceful-restart],
[edit protocols ldp interface interface-name],
[edit routing-instances routing-instance-name protocols ldp interface interface-name],
[edit routing-instances routing-instance-name routing-options graceful-restart]
```

Description

Explicitly disable LDP on an interface, or explicitly disable LDP graceful restart.

Default

LDP is enabled on interfaces configured with the LDP `interface` statement. LDP graceful restart is automatically enabled when graceful restart is enabled under the `[edit routing-options]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling and Disabling LDP

Configuring LDP Graceful Restart

disable (Protocols MLD)

IN THIS SECTION

- Syntax | 5791
- Hierarchy Level | 5791
- Description | 5791
- Required Privilege Level | 5792
- Release Information | 5792

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

Disable MLD on the system.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Disabling MLD

disable (Protocols MLD Snooping)

IN THIS SECTION

- Syntax | 5792
- Hierarchy Level | 5793
- Description | 5793
- Default | 5793
- Required Privilege Level | 5793
- Release Information | 5793

Syntax

```
disable;
```

Hierarchy Level

```
[edit protocols mld-snooping vlan (all | vlan-name)]
```

Description

Disable MLD snooping on the VLAN. Multicast traffic will be flooded to all interfaces in the VLAN except the source interface.

Default

If you do not include this statement, MLD snooping is enabled on all interfaces in the VLAN.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)

show mld-snooping vlans | 22100

disable (Protocols MPLS)

IN THIS SECTION

- Syntax | 5794
- Hierarchy Level | 5794
- Description | 5794

- Default | 5794
- Required Privilege Level | 5794
- Release Information | 5795

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls interface interface-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls],
[edit protocols mpls interface interface-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Disable the functionality of the configured object.

Default

The configured object is enabled (operational) unless explicitly disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

label-switched-path (Protocols MPLS) | 9585

interface (Protocols MPLS) | 8983

disable (Protocols MSDP)

IN THIS SECTION

- Syntax | 5795
- Hierarchy Level | 5795
- Description | 5796
- Required Privilege Level | 5796
- Release Information | 5796

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name],
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
```

```
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
peer address],
[edit protocols msdp],
[edit protocols msdp group group-name],
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit routing-instances routing-instance-name protocols msdp],
[edit routing-instances routing-instance-name protocols msdp group group-name],
[edit routing-instances routing-instance-name protocols msdp group group-name peer address],
[edit routing-instances routing-instance-name protocols msdp peer address]
```

Description

Explicitly disable MSDP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Disabling MSDP

disable (Protocols Router Discovery)

IN THIS SECTION

- Syntax | 5797
- Hierarchy Level | 5797
- Description | 5797
- Default | 5797
- Required Privilege Level | 5797
- Release Information | 5798

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery],
[edit protocols router-discovery]
```

Description

Disable router discovery.

Default

The configured object is enabled (operational) unless explicitly disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

disable (Protocols RSVP)

IN THIS SECTION

- Syntax | 5798
- Hierarchy Level | 5798
- Description | 5799
- Default | 5799
- Required Privilege Level | 5799
- Release Information | 5799

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit logical-systems logical-system-name protocols rsvp graceful-restart],
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit logical-systems logical-system-name protocols rsvp peer-interface peer-interface-name],
```

```
[edit protocols rsvp],
[edit protocols rsvp graceful-restart],
[edit protocols rsvp interface interface-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp peer-interface peer-interface-name]
```

Description

Explicitly disable RSVP or RSVP graceful restart. Explicitly disable link protection on the specified interface.

Default

RSVP is enabled on interfaces and peer interfaces configured with the RSVP `interface` statement. RSVP graceful restart is enabled on the router. Link protection is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Minimum RSVP Configuration

Configuring RSVP Graceful Restart

Configuring Link Protection on Interfaces Used by LSPs

disable (Protocols SAP)

IN THIS SECTION

- Syntax | 5800
- Hierarchy Level | 5800
- Description | 5800
- Required Privilege Level | 5800
- Release Information | 5800

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols sap],
[edit protocols sap]
```

Description

Explicitly disable SAP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Session Announcement Protocol

disable-relay

IN THIS SECTION

- Syntax | 5801
- Hierarchy Level | 5801
- Description | 5802
- Required Privilege Level | 5802
- Release Information | 5802

Syntax

```
disable-relay;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
```

```
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Disable DHCP relay on specific interfaces in a group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

disable-routing-engine (DDoS)

IN THIS SECTION

- Syntax | 5803
- Hierarchy Level | 5803
- Description | 5803
- Required Privilege Level | 5803
- Release Information | 5803

Syntax

```
disable-routing-engine;
```

Hierarchy Level

```
[edit system ddos-protection global],
[edit system ddos-protection protocols protocol-group (aggregate | packet-type)],
[edit system ddos-protection protocols protocol-group aggregate]
```

Description

(ACX Series routers, MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Disable DDoS Routing Engine policers for debugging purposes for a particular packet type as follows:

- Within a protocol group—Use this option in the `protocols` statement hierarchy.
- Globally for all packet types in all protocol groups—Use this option in the `global` statement hierarchy.

This statement does not affect the state of the line card policers.

> **NOTE**: ACX Series routers only support aggregate policers and don't support policing at the line card level. Changes to default settings at the Routing Engine level propagate down to the Packet Forwarding Engine (PFE) chipset level. As a result, disabling Routing Engine policers globally on ACX Series routers effectively disables control plane DDoS protection on the device.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Disabling Control Plane DDoS Protection Policers and Logging Globally

disable (Routing Options)

IN THIS SECTION

- Syntax | 5804
- Hierarchy Level | 5804
- Description | 5804
- Required Privilege Level | 5804
- Release Information | 5805

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options
graceful-restart],
[edit logical-systems logical-system-name routing-options graceful-restart],
[edit routing-instances routing-instance-name routing-options graceful-restart],
[edit routing-options graceful-restart]
```

Description

Disable graceful restart.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS High Availability User Guide

disable-signature-check

IN THIS SECTION

- Syntax | 5805
- Junos OS Hierarchy Level | 5805
- Junos OS Evolved Hierarchy Level | 5806
- Description | 5806
- Required Privilege Level | 5806
- Release Information | 5806

Syntax

```
disable-signature-check;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Disable signature verification on the received test frames. This statement is valid only if you configure the test mode to be a reflector. The configuration is useful when the test traffic is generated using a third-party vendor tool, instead of an ACX Series router.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

Supported RFC 2544-Based Benchmarking Statements on MX Series Routers

Configuring an RFC 2544-Based Benchmarking Test

disable-timeout

IN THIS SECTION

Syntax | 5807

Hierarchy Level | 5807

- Description | 5807
- Default | 5807
- Options | 5808
- Required Privilege Level | 5808
- Release Information | 5808

Syntax

```
disable-timeout timeout;
```

Hierarchy Level

```
[edit ethernet-switching-options port-error-disable]
```

Description

Specify how long the Ethernet switching interfaces remain in a disabled state because of MAC limiting, MAC move limiting, or storm control errors.

> **NOTE**: If you modify the timeout value of an existing disable timeout setting, the new timeout value does not impact the timing of restoration to service of currently disabled interfaces that have been configured for automatic recovery. The new timeout value is applied only during the next occurrence of a port error.
>
> You can bring up the currently disabled interfaces by running the operational command `clear ethernet-switching port-error`.

Default

The disable timeout is not enabled.

Options

timeout—Time, in seconds, that the disabled state remains in effect. The disabled interface is automatically restored to service when the specified timeout value is reached.

- **Range:** 10 through 3600 seconds

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Autorecovery for Port Security Events

Configuring Autorecovery for Port Security Events

disable-timeout (BPDU)

IN THIS SECTION

- Syntax | 5809
- Hierarchy Level | 5809
- Description | 5809
- Default | 5809
- Options | 5809
- Required Privilege Level | 5809
- Release Information | 5810

Syntax

```
disable-timeout seconds;
```

Hierarchy Level

```
[edit protocols layer2-control bpdu-block]
```

```
[edit ethernet-switching-options                 bpdu-block]
```

Description

For interfaces configured for BPDU protection, specify the amount of time an interface receiving BPDUs is disabled.

Configure the timeout value to periodically check to see if an interface is still disabled with BPDU blocking. If this option is not configured, the interface is not periodically checked and remains disabled.

Default

The disable timeout is not enabled.

Options

seconds——Amount of time, in seconds, the interface receiving BPDUs is disabled. Once the timeout expires, the interface is brought back into service.

- **Range:** 10 through 3600
- **Default:** If this option is not configured, the interface is not periodically checked and remains disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Understanding BPDU Protection for Spanning-Tree Instance Interfaces

Understanding BPDU Protection for STP, RSTP, and MSTP

Configuring BPDU Protection for Individual Spanning-Tree Instance Interfaces

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches

Example: Configuring Network Regions for VLANs with MSTP on Switches

show spanning-tree bridge | 26548

show spanning-tree interface | 26555

disable-timeout (Port Error Disable)

IN THIS SECTION

- Syntax | 5811
- Hierarchy Level | 5811
- Description | 5811
- Default | 5811
- Options | 5811
- Required Privilege Level | 5811
- Release Information | 5811

Syntax

```
disable-timeout timeout;
```

Hierarchy Level

```
[edit ethernet-switching-options  port-error-disable]
```

Description

Specify how long Ethernet switching interfaces remain in a disabled state due to MAC limiting, MAC move limiting, or storm control errors.

> **NOTE**: If you modify an existing timeout value, the new timeout value does not affect currently disabled interfaces are configured for automatic recovery. The new timeout value applies only to subsequent port errors. Run the `clear ethernet-switching port-error` command to restore currently disabled interfaces.

Default

The disable timeout statement is not enabled.

Options

timeout—Time, in seconds, that an interface remains disabled. The disabled interface automatically returns to service when the specified time expires.

- **Range:** 10 through 3600 seconds

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- Understanding MAC Limiting and MAC Move Limiting for Port Security
- *Understanding Storm Control*
- *Example: Using Storm Control to Prevent Network Outages (non-ELS)*
- *Example: Using Storm Control to Prevent Network Outages*
- action-shutdown

disable-https-filtering

IN THIS SECTION

- Syntax | 5812
- Hierarchy Level | 5812
- Description | 5813
- Required Privilege Level | 5813
- Release Information | 5813

Syntax

```
disable-https-filtering;
```

Hierarchy Level

```
[edit services web-filter profile name url-filter-template]
```

Description

When you configure `disable-https-filtering`, the system permits or bypasses all HTTPS traffic. The TERM action takes priority over the `disable-https-filtering` command.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.3R1.

disable-packet-register

IN THIS SECTION

- Syntax | 5813
- Hierarchy Level | 5813
- Description | 5814
- Release Information | 5814

Syntax

```
disable-packet-register;
```

Hierarchy Level

```
[edit dynamic-profiles name protocols pim],
[edit logical-systems name protocols pim],
[edit logical-systems name routing-instances name protocols pim],
[edit logical-systems name tenants name routing-instances name protocols pim],
```

```
[edit protocols pim],
[edit routing-instances name protocols pim],
[edit tenants name routing-instances name protocols pim]
```

Description

Disable PIM packet registers on the first-hop router.

With PIM, the first-hop router (FHR) is the device connected to the PIM rendezvous point (RP). With the usual PIM behavior, the FHR encapsulates multicast packets in PIM register messages (which registers the source with the PIM RP) and sends the register messages to the PIM RP. The PIM RP registers the multicast (S,G) route, and decapsulates the multicast packets to forward them on to the receivers. After the RP has calculated and added the path to the source to its multicast routing tree, the RP sends a PIM register stop message toward the FHR.

Devices that might not be able to encapsulate (or decapsulate) multicast packets in PIM register messages will not generate the PIM register stop message. Without a PIM register stop message, the source registration process doesn't terminate properly. As a result, you can set this option on devices that act as FHRs to avoid keeping the PIM registration process in an unterminated state.

Release Information

Statement introduced in Junos OS Release 23.3R1.

discard (forwarding-options)

IN THIS SECTION

- Syntax | 5815
- Hierarchy Level | 5815
- Description | 5815
- Required Privilege Level | 5815
- Release Information | 5816

Syntax

```
discard {
    prefix-list-name prefix-list-name;
    source-ip-address-list address-list-name;
    source-mac-address-list address-list-name;
}
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard policy policy-name]
```

Description

Configure a discard policy for an IPv6 Router Advertisement (RA) guard policy. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.

When RA guard is enabled, the switch compares the information contained in the attributes of RA messages to the criteria configured in the policy. When RA guard is enabled by using a discard policy, any RA messages that match the conditions defined in the policy are dropped, and RA messages that do not match the conditions are forwarded.

The criteria are configured as one or more lists of IPv6 addresses, MAC addresses, or IPv6 address prefixes associated with the policy. RA guard compares the source address or address prefix of incoming RA messages with the configured lists. You configure the lists at the [`edit policy-options`] hierarchy level, by using the `prefix-list` option for an IPv6 address or address prefix list, and the `mac-list` option for a MAC address list.

If more than one list is associated with a discard policy, then an incoming RA message that meets the criteria in any of the lists is discarded.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

discard (routing-options)

IN THIS SECTION

- Syntax | 5816
- Hierarchy Level | 5816
- Description | 5817
- Default | 5817
- Required Privilege Level | 5817
- Release Information | 5817

Syntax

```
discard;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options (aggregate | generate) (defaults | route)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name (aggregate | generate) (defaults | route)],
[edit logical-systems logical-system-name routing-options (aggregate | generate) (defaults |
route)],
```

```
[edit logical-systems logical-system-name routing-options rib routing-table-name (aggregate |
generate) (defaults | route)],
[edit routing-instances routing-instance-name routing-options (aggregate | generate) (defaults |
route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name (aggregate |
generate) (defaults | route)],
[edit routing-options (aggregate | generate) (defaults | route)],
[edit routing-options rib routing-table-name (aggregate | generate) (defaults | route)]
```

Description

Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

To propagate static routes into the routing protocols, include the `discard` statement when you define the route, along with a routing policy.

> **NOTE**: In other vendors' software, a common way to propagate static routes into routing protocols is to configure the routes so that the next-hop routing device is the loopback address (commonly, **127.0.0.1**). However, configuring static routes in this way (by including a statement such as **route *address/ mask-length* next-hop 127.0.0.1**) does not propagate the static routes, because the forwarding table ignores static routes whose next-hop routing device is the loopback address.

Default

When an aggregate route becomes active, it is installed in the routing table with a reject next hop, which means that ICMP unreachable messages are sent.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation

aggregate (Routing) | 2070

generate | 7672

discard-action-for-unresolved-redir-addr

IN THIS SECTION

- Syntax | 5818
- Hierarchy Level | 5818
- Description | 5818
- Required Privilege Level | 5819
- Release Information | 5819

Syntax

```
discard-action-for-unresolved-redir-addr;
```

Hierarchy Level

```
[edit routing-instances name routing-options flow],
[edit routing-options flow]
```

Description

Configure the discard action for BGP flow specification routes that were not resolved using the redirect to IP action. If you do not configure this option, then Junos OS by default accepts the unresolved redirect to IP addresses.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

legacy-redirect-ip-action | 9711

Configuring BGP Flow Specification Action Redirect to IP to Filter DDoS Traffic

Understanding BGP Flow Routes for Traffic Filtering

discard (Forwarding Class)

IN THIS SECTION

- Syntax | 5819
- Hierarchy Level | 5819
- Description | 5820
- Required Privilege Level | 5820
- Release Information | 5820

Syntax

```
discard;
```

Hierarchy Level

```
[edit class-of-service forwarding-policy next-hop-map map-name forwarding-class class-name]
```

Description

Discard traffic sent to this forwarding class for the next-hop map referenced by this forwarding policy.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Configuring CoS-Based Forwarding

non-lsp-next-hop | **11659**

disconnect-peer-timeout

IN THIS SECTION

- Syntax | **5820**
- Hierarchy Level | **5821**
- Description | **5821**
- Options | **5821**
- Required Privilege Level | **5821**
- Release Information | **5821**

Syntax

```
disconnect-peer-timeout seconds;
```

Hierarchy Level

```
[edit access diameter peer peer-name]
```

Description

Configure the amount of time to wait in the Closing state while disconnecting this peer.

Options

seconds Amount of time to wait in the Closing state.

- **Range:** 1 through 65,535 seconds
- **Default:** 10 seconds

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

disk-failure-action

IN THIS SECTION

Syntax | 5822

- Hierarchy Level | 5822
- Description | 5822
- Options | 5822
- Required Privilege Level | 5822
- Release Information | 5822

Syntax

```
disk-failure-action (halt | reboot);
```

Hierarchy Level

```
[edit chassis routing-engine on-disk-failure]
```

Description

Configure the Routing Engine to halt or reboot when the Routing Engine hard disk fails.

Options

halt—Specify the Routing Engine to halt.

reboot—Specify the Routing Engine to reboot.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Enabling a Routing Engine to Reboot on Hard Disk Errors

disk-failure-action

IN THIS SECTION

- Syntax | 5823
- Hierarchy Level | 5823
- Description | 5823
- Options | 5823
- Required Privilege Level | 5824
- Release Information | 5824

Syntax

```
disk-failure-action (halt | reboot);
```

Hierarchy Level

```
[edit chassis redundancy on-disk-failure],
[edit chassis routing-engine on-disk-failure]
```

Description

Configure the Routing Engine to halt or reboot when the Routing Engine hard disk fails.

Options

halt—Specify the Routing Engine to halt.

reboot—Specify the Routing Engine to reboot.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

graceful-switchover (Chassis) | 7786

Enabling a Routing Engine to Reboot on Hard Disk Errors

Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)

High Availability Features for EX Series Switches Overview

disk-partition

IN THIS SECTION

- Syntax | 5824
- Hierarchy Level | 5825
- Description | 5825
- Options | 5825
- Required Privilege Level | 5825
- Release Information | 5825

Syntax

```
disk-partition partition-name (/config |/var) {
    level state (full | high) {
        free-space threshold-value (mb | percent);
```

```
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the disk usage monitoring level. When the specified disk usage monitoring level is reached, a system alarm is activated.

To view the partition size and the current unused space, use the `show system storage partitions` command.

To clear disk space, use the `request system storage cleanup` command.

The remaining statements are explained separately.

NOTE: The `disk-partition` configuration option is not available for EX9200, QFX10008, QFX10016, PTX and MX series devices.

Options

partition-name Select the partition to be monitored for disk usage.

- `/config`—Monitor the /config partition for disk usage.
- `/var`— Monitor the /var partition for disk usage.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

display-host (Security Fallback Block)

IN THIS SECTION

- Syntax | 5826
- Hierarchy Level | 5826
- Description | 5826
- Required Privilege Level | 5826
- Release Information | 5827

Syntax

```
display-host;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-block]
```

Description

Display the computer host name in the notification e-mail sent to the administrator when a fallback-block notification occurs.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

display-host (Security Virus Detection)

IN THIS SECTION

- Syntax | 5827
- Hierarchy Level | 5828
- Description | 5828
- Required Privilege Level | 5828
- Release Information | 5828

Syntax

```
display-host;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus profile profile name notification-options virus-
detection]
```

Description

Display the computer host name in the notification e-mail sent to the administrator when a virus is detected by Sophos antivirus.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

distinguished-name

IN THIS SECTION

- Syntax | 5829
- Hierarchy Level | 5829
- Description | 5829
- Options | 5829
- Required Privilege Level | 5829
- Release Information | 5829

Syntax

```
distinguished-name container container-string-values |wildcard wildcard-string-values
```

Hierarchy Level

```
[edit services ipsec-vpn ikepolicy policy-name local-id|remote-id
```

Description

Specify one or more distinguished name values.

A distinguished name is a name used with digital certificates to uniquely identify a user. For example a distinguished name can be:

- CN=user
- DC=example
- DC=com

Optionally, you can use the container keyword to specify that the order of the fields in a DN and their values exactly match the configured DN, or use the wildcard keyword to specify that the values of fields in a DN must match but the order of the fields does not matter.

Options

`container` *`container-string-values`* —One or more distinguished name container string.

`wildcard` *`wildcard-string-values`* —One or more distinguished name wildcard string.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

`distinguished-name` option added in Junos OS Release 19.1.

RELATED DOCUMENTATION

Configuring IKE Policies

distinguished-name (Access)

IN THIS SECTION

- Syntax | 5830
- Hierarchy Level | 5830
- Description | 5830
- Options | 5830
- Required Privilege Level | 5831
- Release Information | 5831

Syntax

```
distinguished-name distinguished-name;
```

Hierarchy Level

```
[edit access ldap-options search admin-search],
[edit access profile profile-name ldap-options search admin-search]
```

Description

Specify the distinguished name of an administrative user. The distinguished name is used in the bind for performing the LDAP search.

Options

distinguished-name—Set of properties that define the user. For example, `cn`=admin, `ou`=eng, `o`=juniper, `dc`=net.

Required Privilege Level

secret—To view this statement in the configuration.

secret-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring RADIUS and LDAP User Authentication

distinguished-name (Security)

IN THIS SECTION

- Syntax | 5831
- Hierarchy Level | 5831
- Description | 5832
- Options | 5832
- Required Privilege Level | 5832
- Release Information | 5832

Syntax

```
distinguished-name <container container-string> <wildcard wildcard-string>
```

Hierarchy Level

```
[edit security ike gateway gateway-name dynamic]
```

Description

Specify a distinguished name as the identifier for the remote gateway with a dynamic IP address.

Options

container-string

DN field and value to be matched. For example, `cn`=admin, `ou`=eng, `o`=example, `dc`=net. Specify one or more distinguished name (DN) field and value pairs that must match the DN in the VPN peer's digital certificate. The order of the fields and their values must exactly match the DN in the peer's digital certificate.

Add a space between each field and value pair. For example, `edit security ike gateway jsr_gateway dynamic distinguished-name container o=example, dc=net`.

wildcard-string

DN field and value pairs to be matched. For example, `cn`=admin, `ou`=eng, `o`=example, `dc`=net. Specify one or more distinguished name (DN) field and value pairs that must match the DN in the VPN peer's digital certificate. The configured field and value must match the DN in the peer's digital certificate but the order of the fields in the DN does not matter.

Add a space between each field and value pair. For example, `edit security ike gateway jsr_gateway dynamic distinguished-name wildcard o=example, dc=net`.

Starting in Junos OS Release 19.4R1, you can now configure only one dynamic DN attribute among `container-string` and `wildcard-string` at `[edit security ike gateway` *`gateway_name`* `dynamic distinguished-name]` hierarchy. If you try configuring the second attribute after you configure the first attribute, the first attribute is replaced with the second attribute. Before your upgrade your device, you must remove one of the attributes if you have configured both the attributes.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

IPsec Overview

distributed-dr

IN THIS SECTION

- Syntax | 5833
- Hierarchy Level | 5833
- Description | 5833
- Required Privilege Level | 5834
- Release Information | 5834

Syntax

```
distributed-dr;
```

Hierarchy Level

```
[edit dynamic-profiles name protocols pim interface (Protocols PIM) interface-name],
[edit logical-systems name protocols pim interface (Protocols PIM) interface-name],
[edit logical-systems name routing-instances name protocols pim interface (Protocols PIM) interface-
name],
[edit protocols pim interface (Protocols PIM) interface-name],
[edit routing-instances name protocols pim interface (Protocols PIM) interface-name]
```

Description

Enable PIM distributed designated router (DR) functionality on IRB interfaces associated with EVPN virtual LANs (VLANs) that have been configured with IGMP snooping or MLD snooping. By effectively disabling certain PIM features that are not required in this scenario, this statement supports using PIM to perform intersubnet, that is, inter-VLAN, multicast routing more efficiently.

When you configure this statement on an interface on a device, PIM ignores the DR status of the interface when processing IGMP reports received on the interface. When the interface receives the IGMPor MLD report, the device sends PIM upstream join messages to pull the multicast stream and forward it to the interface regardless of the DR status of the interface. This setting also disables the PIM assert mechanism on the interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

- Configure Multicast Forwarding with IGMP Snooping in an EVPN-MPLS Environment
- *Configure Multicast Forwarding with MLD Snooping in an EVPN-MPLS Environment*
- *Example: Preserving Bandwidth with IGMP Snooping in an EVPN-VXLAN Environment*

distributed (IGMP)

IN THIS SECTION

- Syntax | 5835
- Hierarchy Level | 5835
- Description | 5835
- Required Privilege Level | 5835
- Release Information | 5835

Syntax

```
distributed;
```

Hierarchy Level

```
[edit protocols igmp interface interface-name],
[edit dynamic-profiles protocols igmp interface $junos-interface-name]
```

Description

Enable distributed IGMP by moving IGMP processing from the Routing Engine to the Packet Forwarding Engine. Distributed IGMP reduces the join and leave latency of IGMP memberships.

Distributed IGMP is only available when `chassis network-services enhanced-ip` is configured.

> **NOTE**: When you enable distributed IGMP, the following interface options are not supported on the Packet Forwarding Engine: `oif-map`, `group-limit`, `ssm-map`, and `static`. However, the `ssm-map-policy` option is supported on distributed IGMP interfaces. The `traceoptions` and `accounting` statements can only be enabled for IGMP operations still performed on the Routing Engine; they are not supported on the Packet Forwarding Engine. The `clear igmp membership` command is not supported when distributed IGMP is enabled.

When the `distributed` command is enabled in conjunction with mldp-inband-signalling, (so PIM act as a multipoint LDP inband edge router), it supports interconnecting separate PIM domains via a MPLS-based core.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X50.

Support added in Junos OS Release 18.2R1 for using `distributed` IGMP in conjunction with Multipoint LDP (mLDP) in-band signalling.

RELATED DOCUMENTATION

Enabling Distributed IGMP

Configuring Dynamic DHCP Client Access to a Multicast Network

Junos OS Multicast Protocols User Guide

distribution-ip (Security SIP)

IN THIS SECTION

- Syntax | 5836
- Hierarchy Level | 5837
- Description | 5837
- Options | 5837
- Required Privilege Level | 5837
- Release Information | 5838

Syntax

```
distribution-ip ip-address {
    distribute-to {
        fpc fpc-slot {
            pic pic-slot;
            }
        }
}
```

Hierarchy Level

```
[edit security alg sip]
```

Description

The `distribution-ip` configuration statement allows you to configure maximum of up to 20 IPv4 or IPv6 distribution IP addresses that enable SIP ALG to distribute the traffic based on the available IP addresses. Use the optional `fpc` and `pic` commands to add Flexible PIC Concentrator (FPC) card numbers and PIC card numbers to distribute the SIP traffic with IP address to the corresponding port. If the `fpc` and `pic` options are not configured, the SIP traffic with IP address distributes to the port decided by hash value of IP address.

Before configuring the `set security alg sip distribution-ip` *`ip-address`* `fpc` *`fpc-slot`* `pic` *`pic-slot`* command, perform the following steps:

- Use the `delete security alg sip distribution-ip` *`ip-address`* command to remove a specific IP address or use the `delete security alg sip distribution-ip` command to remove all the existing IP addresses.
- Use the `clear security nat incoming-table` command to remove the dynamically added IP addresses.

Options

fpc Flexible PIC Concentrator (FPC) card number.

fpc-slot FPC slot number.

- **Range:** 0 through 127

pic PIC card number.

pic-slot PIC slot number.

- **Range:** 0 through 3

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Understanding the SIP ALG

clear security nat incoming-table | 4007

distribution-profile

IN THIS SECTION

- Syntax | 5838
- Hierarchy Level | 5839
- Description | 5839
- Options | 5840
- Required Privilege Level | 5840
- Release Information | 5840

Syntax

```
distribution-profile (fat-core | name) {
    description description;
    fpc fpc {
        pic pic {
            thread-id thread-id;
        }
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

The `distribution-profile` option is introduced to give the administrator an option to define a profile to handle tunnels associated with a certain VPN object. If the default profiles such as `default-spc3-profile` or `default-spc2-profile` are not selected, a new user-defined profile can be created. In a profile, you should mention the Flexible PIC Concentrator (FPC) slot and the PIC slot. When this profile is associated with a VPN object, all matching tunnels will be distributed across these PICs. The `thread-id` is an optional value. If you specify a thread ID, then the tunnel is distributed in the specified thread.

Starting in Junos OS Release 20.2R1, when you add, change, or delete the thread ID from distribution profile, all tunnels part of modified distribution profile anchored on modified SPU member of distribution profile are teared down and re-negotiated. See Table 30 on page 5839 for catastrophic changes when you change the distribution profile configuration.

Table 30: Distribution Profile Catastrophic Change

Distribution Profile Change	Catastrophic Change
Add new distribution profile added to the VPN.	All tunnels part of this new distribution profile are brought down and re-negotiated.
Add new SPU information to a profile already part of the VPN.	No impact on any tunnel.
Delete SPU information from a distribution profile of the VPN.	Only those tunnels part of the distribution profile that is modified and anchored on a deleted SPU are brought down and re-negotiated.
Add first thread ID to an SPU part of the distribution profile.	All tunnels part of this distribution profile are brought down and re-negotiated.
Add next set of thread IDs to SPU part of the distribution profile.	No impact for any tunnel.

Table 30: Distribution Profile Catastrophic Change *(Continued)*

Distribution Profile Change	Catastrophic Change
Delete a thread ID from an SPU part of the distribution profile.	Only those tunnels part of the modified distribution profile and anchored on the deleted SPU are brought down and re-negotiated.
Delete last thread ID from SPU part of the distribution profile.	No impact on any tunnel.
Delete distribution profile from the VPN	No impact on any tunnel.
Change distribution profile name from profileA to profileB in VPN.	All tunnels part of this profile are brought down and re-negotiated.

Options

description Text description of the distribution profile.

fpc FPC slot number.

pic PIC slot number.

thread-id (Optional) Thread ID number. Only valid for SPC3.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

dlb

IN THIS SECTION

- Syntax | 5841
- Hierarchy Level | 5841
- Description | 5841
- Options | 5842
- Required Privilege Level | 5842
- Release Information | 5842

Syntax

```
dlb {
    assigned-flow;
    per-packet;
    flowlet inactivity-interval;
    ether-type (Dynamic
       Load Balancing) (ipv4|ipv6|mpls);
}
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

Description

Configure dynamic load balancing (DLB) on LAG with specific mode type—assigned flow, per-packet, or flowlet) and ether type—ipv4, ipv6, or mpls.

NOTE: Since ether-type is a global configuration and applicable to all LAGs in the system with DLB enabled, the CLI for specifying ether-type for LAG is added in the global

format under forwarding-options. See "ether-type (Dynamic Load Balancing)" on page 6473 for details.

NOTE: QFX switches in a virtual chassis configuration do not support dynamic load balancing.

Options

assigned-flow	Fixed link assignment.
per-packet	Per-packet link assignment.
flowlet ***inactivity-interval***	Minimum inactivity interval for link re-assignment. • **Range:** 16 through 65535 (in micro seconds).
ether-type ***(ipv4 | ipv6 | mpls)***	EtherType for DLB: • ipv4—Set dynamic load balancing for IPv4. • ipv6—Set dynamic load balancing for IPv6. • mpls—Set dynamic load balancing for MPLS.

Required Privilege Level

interface - To view this statement in the configuration.

interface-control - To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

Load Balancing for Aggregated Ethernet Interfaces

Load Balancing for Aggregated Ethernet Interfaces

Load Balancing for Aggregated Ethernet Interfaces

dlci

IN THIS SECTION

- Syntax | 5843
- Hierarchy Level | 5843
- Description | 5843
- Options | 5843
- Required Privilege Level | 5844
- Release Information | 5844

Syntax

```
dlci dlci-identifier;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For Frame Relay and Multilink Frame Relay user-to-network interface (UNI) network-to-network interface (NNI) encapsulation only, and for link services and point-to-point interfaces only, configure the data-link connection identifier (DLCI) for a permanent virtual circuit (PVC) or a switched virtual circuit (SVC).

To configure a DLCI for a point-to-multipoint interface, use the `multipoint-destination` statement to specify the DLCI.

Options

dlci-identifier—Data-link connection identifier.

- **Range:** 16 through 1022

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring DLCIs on Link Services Logical Interfaces

Junos OS Network Interfaces Library for Routing Devices

Data-Link Connection Identifiers on Channelized Interfaces

Configuring Frame Relay DLCIs

Junos OS Services Interfaces Library for Routing Devices

encapsulation (Logical Interface) | **6300**

multicast-dlci

multipoint-destination

dlv

IN THIS SECTION

- Syntax | **5845**
- Hierarchy Level | **5845**
- Description | **5845**
- Options | **5845**
- Required Privilege Level | **5845**
- Release Information | **5845**

Syntax

```
dlv {
    domain-name domain-name trusted-anchor  trusted-anchor;
    }
```

Hierarchy Level

```
[edit system services dns dnssec]
```

Description

Configure DNSSEC Lookaside Validation (DLV) (RFC 5074).

Options

domain-name *domain-name*	Specify the secure domain server name.
trusted-anchor *trusted-anchor*	Specify the trusted DLV anchor.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

dne-origin (Diameter Network Element)

IN THIS SECTION

- Syntax | 5846
- Hierarchy Level | 5846
- Description | 5846
- Options | 5846
- Required Privilege Level | 5847
- Release Information | 5847

Syntax

```
dne-origin realm realm-name <host hostname>;
```

Hierarchy Level

```
[edit diameter network-element element-name]
```

Description

Specify values of Origin-Realm-AVP and Origin-Host-AVP used in messages sent for the specified network element by the Diameter instance.

NOTE: Only the realm is mandatory for the DNE origin.

Options

host *hostname* (Optional) Name of the message origin host that is supplied as the value of the Origin-Host AVP for Diameter messages associated with the network element.

realm ***realm-name***

Name of the message origin realm, that is supplied as the value of the Origin-Realm AVP for Diameter messages associated with the network element.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring Diameter Network Elements

Configuring Diameter

dns

IN THIS SECTION

- Syntax | 5847
- Hierarchy Level | 5848
- Description | 5848
- Options | 5848
- Required Privilege Level | 5849
- Release Information | 5849

Syntax

```
dns {
    disable;
```

```
    doctoring (none | sanity-check);
    maximum-message-length bytes;
    oversize-message-drop;
    traceoptions {
        flag {
            all {
                extensive
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Specify the Domain Name System (DNS) Application Layer Gateway (ALG) on the device.

Options

- `disable`—Disable the DNS ALG. By default, the DNS ALG is enabled. You can enable or disable the DNS ALG for both IPv4 and IPv6 mode.
- `doctoring`—Configure DNS ALG doctoring.
  - `none`— Disable all DNS ALG doctoring.
  - `sanity-check`—Perform only DNS ALG sanity checks.
- `maximum-message-length`—Set maximum message length (bytes).

  Range: 512 through 819.

  Default: 512.
- `oversize-message-drop`—Configure to drop the oversized DNS packets.

- `traceoptions`— Configure DNS ALG tracing options.
  - `flag`—Trace operation to perform.
    - `all`—Trace all events.
    - *`extensive`*—Display extensive amount of data.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

DNS Overview

dns-alg-pool

IN THIS SECTION

- Syntax | 5850
- Hierarchy Level | 5850
- Description | 5850
- Required Privilege Level | 5850
- Release Information | 5850

Syntax

```
dns-alg-pool dns-alg-pool;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated]
```

Description

Specify the Network Address Translation (NAT) pool for destination translation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

dns-alg-prefix

IN THIS SECTION

- Syntax | 5851
- Hierarchy Level | 5851
- Description | 5851

Required Privilege Level | 5851
Release Information | 5851

Syntax

```
dns-alg-prefix dns-alg-prefix;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated]
```

Description

Set the Domain Name System (DNS) application-level gateway (ALG) 96-bit prefix for mapping IPv4 addresses to IPv6 addresses.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

dns-cache

IN THIS SECTION

- Syntax | 5852
- Hierarchy Level | 5852
- Description | 5852
- Options | 5853
- Required Privilege Level | 5853
- Release Information | 5853

Syntax

```
dns-cache {
    error-response-delete-ip {
        retry-interval seconds;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name security policies],
[edit logical-systems logical-system-name tenants tenant-system-name security policies],
[edit security policies],
[edit tenants tenant-system-name security policies]
```

Description

Define security policy DNS cache behaviors. The following behaviors can occur:

- Flushing of the DNS cache entry on error response.
- Keeping the DNS cache entry unchanged on all DNS servers time out.
- Updating the DNS cache entry when there is a positive response.

Options

error-response-delete-ip Clear DNS cache entry IP on error DNS response

- Values:

  retry-interval Interval between deleting IP and resending DNS requests

  - **Range:** 5 through 86400 seconds.
  - **Default:** 300 seconds.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

- show security dns-cache
- clear security dns-cache

dns-filter

IN THIS SECTION

- Syntax | 5854
- Hierarchy Level | 5854
- Description | 5854
- Options | 5854
- Required Privilege Level | 5855
- Release Information | 5855

Syntax

```
dns-filter {
    database-file filename;
    dns-resp-ttl seconds;
    dns-server [ ip-address ];
    hash-key key-string;
    hash-method hash-method-name;
    statistics-log-timer minutes;
    wildcarding-level level;
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name],
[edit services web-filter profile profile-name dns-filter-template template-name]
```

Description

Configure the settings for filtering DNS requests for disallowed website domains. Filtering can result in either:

- Blocking access to the site by sending the client a DNS response that includes an IP address or domain name of a sinkhole server instead of the disallowed domain.
- Logging the DNS request and allowing access.

Settings at the [edit services web-filter profile *profile-name* dns-filter-template *template-name*] hierarchy level override the corresponding settings at the [edit services web-filter profile *profile-name*] hierarchy level.

Options

database-file *filename*
: Name of the domain filter database file to use when filtering DNS requests.

dns-resp-ttl *seconds*
: Number of seconds to live while sending the DNS response after taking the DNS sinkhole action.

  - **Default:** 1800
  - **Range:** 0 through 86,400

dns-server [*ip-address*]
(Optional) IP addresses (IPv4 or IPv6) for up to three specific DNS servers. DNS filtering examines only DNS requests that are destined for those DNS servers.

hash-key *key-string*
Hash key that you used to create the hashed domain name in the domain filter database file.

hash-method *hash-method-name*
Hash method that you used to create the hashed domain name in the domain filter database file. The only supported hash method is `hmac-sha2-256`.

statistics-log-timer *minutes*
Number of minutes in the interval for logging statistics for DNS requests and for sinkhole actions performed for each customer IP address.

- **Default:** 5
- **Range:** 0 through 60

wildcarding-level *level*
Level of subdomains that are searched for a match. A value of 0 indicates that subdomains are not searched.

For example, if you set the `wildcarding-level` to 4 and the database file includes an entry for **example.com**, the following comparisons are made for a DNS request that arrives with the domain **198.51.100.0.example.com**:

- **198.51.100.0.example.com**: no match
- **51.100.0.example.com**: no match for one level down
- **100.0.example.com**: no match for two levels down
- **0.example.com**: no match for three levels down
- **example.com**: match for four levels down
- **Range:** 0 through 10

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series.

Support added for Next Gen Services on MX Series routers MX240, MX480 and MX960 with MX-SPC3 services cards in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

dns-filter-template

IN THIS SECTION

- Syntax | 5856
- Hierarchy Level | 5857
- Description | 5857
- Options | 5857
- Required Privilege Level | 5858
- Release Information | 5858

Syntax

```
dns-filter-template template-name {
    client-interfaces [ client-interface-name ];
    client-routing-instance client-routing-instance-name;
    dns-filter {
        database-file filename;
        dns-resp-ttl seconds;
        dns-server [ ip-address ];
        hash-key key-string;
        hash-method hash-method-name;
        statistics-log-timer minutes;
        wildcarding-level level;
    }
    server-interfaces  [ server-interface-name ];
    server-routing-instance server-routing-instance-name;
```

```
    term term-name {
        from {
            src-ip-prefix [ source-prefix ];
        }
        then {
            accept;
            dns-sinkhole;
        }
    }
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name]
```

Description

Configure filtering of DNS requests for disallowed website domains for requests on specific uplink and downlink logical interfaces or routing instances, or for requests from specific source IP address prefixes. The DNS filter template overrides the corresponding settings at the DNS profile level. You can configure up to 32 DNS filter templates in a profile.

Filtering can result in either:

- Blocking access to the site by sending the client a DNS response that includes an IP address or domain name of a sinkhole server instead of the disallowed domain.
- Logging the DNS request and allowing access.

Options

accept	Accept DNS requests for DNS filtering.
client-interfaces [*client-interface-name*]	(Optional) Client-facing (uplink) logical interfaces on which the DNS filter template settings are applied.
client-routing-instance *client-routing-instance-name*	(Optional) Client-facing (uplink) routing instance on which the DNS filter template settings are applied.
dns-filter-template *template-name*	Name of the DNS filter template.

dns-sinkhole Perform the sinkhole action identified in the domain filter database for disallowed DNS requests.

server-interfaces [*server-interface-name*] (Optional) Server-facing logical interfaces (downlink) on which the DNS filter template settings are applied.

server-routing-instance *server-routing-instance-name* (Optional) Server-facing (downlink) routing instance on which the DNS filter template settings are applied.

> **NOTE**: If you configure the client and server interfaces or the client and server routing instances, implicit filters are installed on the interfaces or routing instances to direct DNS traffic to the MS-MPC for DNS filtering. If you configure neither the client and server interfaces nor the routing instances, you must provide a way to direct DNS traffic to the MS-MPC (for example, via routes).

src-ip-prefix [*source-prefix*] (Optional) Source IP address prefixes of DNS requests you want to filter. You can configure a maximum of 64 prefixes in a term. If you do not specify any source prefixes, then all DNS requests are filtered.

term *term-name* Name for a term. You can configure a maximum of 64 terms in a template.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

dns-filtering

IN THIS SECTION

- Syntax | 5859
- Hierarchy Level | 5860
- Description | 5860
- Options | 5860
- Required Privilege Level | 5862
- Release Information | 5862

Syntax

```
dns-filtering {
    sinkhole {
        dns-resp-ttl seconds;
        fqdn domain-name;
        ipv4-address ipv4-address;
        ipv6-address ipv4-address;
        srv-resp-err-code (Noerror | Refused);
        txt-resp-err-code (Noerror | Refused);
        wildcarding-level value;
    }
    traceoptions {
        file {
            file-name;
            files number-of-trace-files;
            match match-expression;
            no-world-readable
            size maximum-trace-file-size;
            world-readable
        }
        flag (all | config | connect | dns | filter | gencfg | normal | operational commands |
parse | routing | snmp | statistics | system | timer);
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
```

```
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure DNS filtering to identify DNS requests for disallowed domains.

NOTE: The DNS sinkhole must be configured only for DNS profile category.

Options

dns-resp-ttl Configure the time-to-live (TTL) value in seconds to send_IN_the DNS response after taking the DNS sinkhole action. The maximum value for DNS response TTL is 86400 (24 hour). Default value is 1800 (30 min).

fqdn Configure the fully qualified domain name (FQDN) that must be sent in the DNS response for the sinkhole domains.

ipv4-address Configure the sinkhole IPv4 address.

ipv6-address Configure the sinkhole IPv6 address.

srv-resp-err-code Configure the DNS response error code that must be sent for bad domains for server query type.

txt-resp-err-code Configure the DNS response error code that must be sent for bad domains for text query type.

wildcarding-level Configure the number of levels that will be iteratively examined for a domain match. The range is 0 to 10. The default value is 10.

file Configure traffic sampling tracing operations. You can configure the following:

- file-name—Specify the file name in which to write trace information.

- files—Specify the maximum number of trace files to be saved with samples or trace data. The range is 2 to 1000.
- match—Specify the regular expression for lines to be logged for tracing.
- size—Specify the maximum trace file size, in kilobytes (KB), megabytes (MB), or gigabytes (GB). Syntax: xk to specify KB, xm to specify MB, or xg to specify GB. The range is 10240 to 1073741824.
- no-world-readable—Disable unrestricted file access.
- world-readable—Enable unrestricted file access. Allows any user to read the log file.

flag

Tracing flag parameters. The available options are:

- all—Trace everything
- config—Trace DNS filtering configuration events
- connect—Trace DNS filtering IPC events
- dns—Trace DNS filtering crawler events
- filter—Trace DNS filtering filter programming events
- gencfg—Trace DNS filtering gencfg events
- normal— Trace normal events
- operation-commands—Trace DNS filtering show events
- parse—Trace DNS filtering parse events
- routing—Trace DNS filtering route programming events
- snmp— Trace DNS filtering snmp events
- statistics—Trace DNS filtering statistics events
- system—Trace DNS filtering system events
- timer—Trace DNS filtering timer events

level

Level of debugging output. The available options are:

- all—Match all levels

- error—Match error conditions
- info—Match informational messages
- notice—Match conditions that should be handled specially
- verbose—Match verbose messages
- warning—Match warning messages

no-remote-trace Disable remote tracing.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 20.4R1.

dns-proxy

IN THIS SECTION

- Syntax | 5863
- Hierarchy Level | 5863
- Description | 5863
- Options | 5863
- Required Privilege Level | 5863
- Release Information | 5864

Syntax

```
dns-proxy {

    cache hostname inet ip-address;
    default-domain domain-name {
        forwarders ip-address;
    }
    interface interface-name;
    propogate-setting (enable | disable);
    view view-name {
        domain domain-name {
            forward-only;
            forwarders ip-address;
        }
        match-clients subnet-address;
    }
}
```

Hierarchy Level

```
[edit system services dns dns-proxy]
```

Description

Configure the device as a DNS proxy server by enabling DNS proxy on a logical interface. This option is supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550M devices.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

- DNS Proxy Overview
- Configuring the Device as a DNS Proxy

dns-server-address

IN THIS SECTION

- Syntax | 5864
- Hierarchy Level | 5864
- Description | 5865
- Required Privilege Level | 5865
- Release Information | 5865

Syntax

```
dns-server-address address {
    lifetime seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Description

Specify the recursive DNS server address that the device must use to resolve DNS names. The recursive DNS server address is the 128-bit IPv6 address of the recursive DNS server. You can configure a maximum of three recursive DNS server addresses at the interface level.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Example: Configuring Recursive DNS Server Addresses for IPv6 Hosts

Understanding Recursive DNS Servers for IPv6

lifetime (Router Advertisement) | **9745**

dns-server-address (Dynamic Profiles)

IN THIS SECTION

- Syntax | **5866**
- Hierarchy Level | **5866**
- Description | **5866**
- Options | **5866**
- Required Privilege Level | **5867**
- Release Information | **5867**

Syntax

```
dns-server-address [dns-server-address | $junos-ipv6-dns-server-address] {
    lifetime seconds;
}
```

Hierarchy Level

```
[edit dynamic-profiles dynamic-profile-name protocols router-advertisement interface interface-name]
```

Description

Specify the address of the DNS server that is used to resolve IPv6 DNS names. You can use RADIUS to provide the address dynamically in the $junos-ipv6-dns-server variable within Access-Accept messages, or you can statically configure up to three IPv6 addresses.

You can also specify the maximum time in seconds for which the DNS server address remains valid. The device can use the recursive DNS server address for DNS name resolution until the time specified expires.

Options

$junos-ipv6-dns-server-address
Dynamically receive the address of the DNS server from RADIUS in Access-Accept messages.

dns-server-address
IPv6 address of the DNS server. You can configure up to three DNS server addresses.

lifetime *seconds*
Maximum time for which the recursive DNS server address remains valid.

- **Range:** 0 through 4,294,967,295 seconds
- **Default:** 1800 seconds
- **Values:** 0 indicates that the advertised recursive DNS server address is no longer valid and that this recursive DNS server address entry can be deleted.

  A value of 4,294,967,295 seconds indicates an infinite lifetime and a persistent entry in the device for this recursive DNS server address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

DNS Resolver for IPv6 DNS Overview

Configuring a DNS Server Address for IPv6 Hosts

dns (System Services)

IN THIS SECTION

- Syntax | 5867
- Hierarchy Level | 5869
- Description | 5869
- Options | 5869
- Required Privilege Level | 5869
- Release Information | 5869

Syntax

```
dns {
        dns-proxy {

            cache hostname inet ip-address;
            default-domain domain-name {
                forwarders ip-address;
```

```
            }
            interface interface-name;
            propogate-setting (enable | disable);
            view view-name {
                domain domain-name {
                    forward-only;
                    forwarders ip-address;
                }
                match-clients subnet-address;
            }
        }
    }
    dnssec {
        disable;
        dlv {
            domain-name domain-name trusted-anchor  trusted-anchor;
            }
    secure-domains domain-name;
        trusted-keys (key dns-key | load-key-file url);
    forwarders {
        ip-address;
    }
    max-cache-ttl seconds;
    max-ncache-ttl seconds;
    traceoptions {
        category {
            category-type;
        }
        debug-level level;
        file {
            filename;
            files number;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure the DNS server.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

DNS Overview

dnssec

IN THIS SECTION

- Syntax | 5870
- Hierarchy Level | 5870
- Description | 5870
- Options | 5870

Required Privilege Level | 5871
Release Information | 5871

Syntax

```
dnssec {
    disable;
    dlv {
        domain-name domain-name trusted-anchor  trusted-anchor;
        }
secure-domains domain-name;
    trusted-keys {
        (key dns-key | load-key-file url);
    }
```

Hierarchy Level

```
[edit system services dns]
```

Description

Configure domain name service security extensions (DNSSEC) in the DNS server. DNSSEC is an extension of DNS that provides authentication and integrity verification of data by using public-key-based signatures.

Options

disable

Disable DNSSEC.

- **Default:** DNSSEC is enabled.

secure-domains [*domain-name*]

Configure one or more secure domains in the DNS server. The server accepts only signed responses for this domain. For unsigned responses, the server returns SERVFAIL error to the client.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.2 of Junos OS.

RELATED DOCUMENTATION

DNSSEC Overview

Example: Configuring Secure Domains and Trusted Keys for DNSSEC

do-not-fragment

IN THIS SECTION

- Syntax | 5871
- Hierarchy Level | 5872
- Description | 5872
- Default | 5872
- Required Privilege Level | 5872
- Release Information | 5872

Syntax

```
do-not-fragment;
```

Hierarchy Level

```
[edit interfaces gr-fpc/pic/port unit logical-unit-number tunnel],
[edit logical-systems logical-system-name interfaces gr-fpc/pic/port unit logical-unit-number
tunnel]
```

Description

For a generic routing encapsulation (GRE) tunnel, disable fragmentation of GRE-encapsulated packets. This sets the do-not-fragment (DF) bit in the outer IP header of the GRE-encapsulated packets so that they do not get fragmented anywhere in the path. When the size of a GRE-encapsulated packet is greater than the MTU of a link that the packet passes through, the GRE-encapsulated packet is dropped.

Default

By default, fragmentation of GRE-encapsulated packets is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Statement introduced in Junos OS Release 19.3 for MPC10E line card.

RELATED DOCUMENTATION

- allow-fragmentation | **2179**
- reassemble-packets | **14100**
- *Enabling Fragmentation and Reassembly on Packets After GRE-Encapsulation*
- Junos OS Services Interfaces Library for Routing Devices

dod-request-policy

IN THIS SECTION

- Syntax | 5873
- Hierarchy Level | 5873
- Description | 5873
- Options | 5873
- Required Privilege Level | 5873
- Release Information | 5874

Syntax

```
dod-request-policy dod-request-policy-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit protocols ldp]
```

Description

Specify the name of the LDP downstream on demand request policy. The `dod-request-policy` statement performs exact match, as a result, LDP sends label request messages only for those FECs matching in the downstream on demand request policy.

Options

dod-request-polcy-name Specify the name of the downstream on demand request policy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring LDP Downstream on Demand

domain (DNS Forwarding Options)

IN THIS SECTION

- Syntax | 5874
- Hierarchy Level | 5875
- Description | 5875
- Required Privilege Level | 5875
- Release Information | 5875

Syntax

```
domain {
    description text-description;
    interface interface-name {
        broadcast;
        description text-description;
        no-listen;
        server address <logical-system logical-system-name> <routing-instance routing-instance-
name>;
    }
    server address <logical-system logical-system-name> <routing-instance routing-instance-name>;
}
```

Hierarchy Level

```
[edit forwarding-options helpers]
```

Description

Enable DNS request packet forwarding.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding

domain (Precision Time Protocol)

IN THIS SECTION

- Syntax | 5876
- Hierarchy Level | 5876
- Description | 5876
- Options | 5876
- Required Privilege Level | 5876
- Release Information | 5876

Syntax

```
domain domain-value;
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

Configure multiple independent Precision Time Protocol (PTP) domains.

NOTE: Only one PTP domain is supported at any given point in time.

Options

domain-value—The PTP domain value. The range depends on the selected profile type.

- **Range:** 0 through 127
- **Default:** 0
- **Range:** 44 to 63 for profile type g.8275.2
- **Default:** 44 for profile type g.8275.2

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

domain (Domain Map)

IN THIS SECTION

- Syntax | 5877
- Hierarchy Level | 5878
- Description | 5878
- Required Privilege Level | 5878
- Release Information | 5878

Syntax

```
domain {
    delimiter [delimiter-character];
    map domain-map-name {
        aaa-logical-system logical-system-name {
            aaa-routing-instance routing-instance-name;
        }
        aaa-routing-instance routing-instance-name;
        access-profile profile-name;
        address-pool pool-name;
        dynamic-profile profile-name;

        strip-domain;
        target-logical-system logical-system-name {
            target-routing-instance routing-instance-name;
        }
        target-routing-instance routing-instance-name;
        tunnel-profile profile-name;
```

```
    }
    parse-direction (left-to-right | right-to-left);
    parse-order (domain-first | realm-first);
    realm-delimiter [delimiter-character];
    realm-parse-direction  (left-to-right | right-to-left);
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure a domain map, which is used to map access options and session parameters for subscriber sessions.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring a Domain Map

domain-name

IN THIS SECTION

- Syntax | 5879
- Hierarchy Level | 5879
- Description | 5879
- Options | 5879
- Required Privilege Level | 5879
- Release Information | 5880

Syntax

```
domain-name domain-name-string;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication username-include]
```

Description

Specify the domain name that is concatenated with the username during the subscriber authentication process.

Options

domain-name-string—The domain name formatted string.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

domain-name (System)

IN THIS SECTION

- Syntax | 5880
- Hierarchy Level | 5880
- Description | 5880
- Options | 5881
- Required Privilege Level | 5881
- Release Information | 5881

Syntax

```
domain-name domain-name;
```

Hierarchy Level

```
[edit system]
```

Description

Configure the name of the domain in which the router or switch is located. This is the default domain name that is appended to hostnames that are not fully qualified.

Options

domain-name Name of the domain.

> **NOTE**: The length of the domain name cannot exceed 255 characters.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

domain-name (Access Profile)

IN THIS SECTION

- Syntax | 5881
- Hierarchy Level | 5882
- Description | 5882
- Options | 5882
- Required Privilege Level | 5882
- Release Information | 5882

Syntax

```
domain-name domain-name;
```

Hierarchy Level

```
[edit access profile profile-name client-name-filter client-name]
```

Description

Specify a domain name that must be in a client's name during the authentication process.

Options

domain-name—Domain name that must be in a client name. The name must not exceed 128 characters.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Obtaining Username and Role Information Through Firewall Authentication

domain-name (DHCP)

IN THIS SECTION

- Syntax | 5883
- Hierarchy Level | 5883
- Description | 5883
- Options | 5883

Required Privilege Level | 5883
Release Information | 5883

Syntax

```
domain-name domain-name;
```

Hierarchy Level

```
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]
```

Description

Configure the name of the domain in which clients search for a DHCP server host. This is the default domain name that is appended to hostnames that are not fully qualified.

Options

domain-name—Name of the domain.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a DHCP Server on Switches

domain-name (DHCP Local Server)

IN THIS SECTION

- Syntax | 5884
- Hierarchy Level | 5884
- Description | 5885
- Options | 5885
- Required Privilege Level | 5885
- Release Information | 5886

Syntax

```
domain-name domain-name-string;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication username-
include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  authentication
username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
```

```
name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify the domain name that is concatenated with the username during the subscriber authentication or DHCP client authentication process.

Options

domain-name-string—Domain name formatted string.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

domain-name (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 5886
- Hierarchy Level | 5886
- Description | 5887
- Options | 5887
- Required Privilege Level | 5887
- Release Information | 5888

Syntax

```
domain-name domain-name-string;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-
include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-
```

```
include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication
username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include]
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
authentication username-include]
```

Description

Specify the domain name that is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Options

domain-name-string—Domain name formatted string.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Specifying Authentication Support

Creating Unique Usernames for DHCP Clients

domain-name-server-inet (Routing Instances and Access Profiles)

IN THIS SECTION

- Syntax | 5888
- Hierarchy Level | 5889
- Description | 5889
- Options | 5889
- Required Privilege Level | 5889
- Release Information | 5889

Syntax

```
domain-name-server-inet dns-address;
```

Hierarchy Level

```
[edit access],
[edit access profile]
```

Description

Configure an IPv4 address for a DNS name server. You can configure an address globally for a routing instance at the `[edit access]` hierarchy level or for an access profile at the `[edit access profile profile-name]` hierarchy level. You can configure more than one address by including the statement multiple times.

> **NOTE**: A DNS name server address configured with this statement is higher in preference than one configured with the "domain-name-server" on page 5891 statement. That is, the server with the address configured with the "domain-name-server-inet" on page 5888 takes precedence over a server configured with the "domain-name-server" on page 5891 statement.

Options

dns-address	IPv4 address of the DNS name server.

Required Privilege Level

admin—To view this statement in the configuration

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring DNS Name Server Addresses for Subscriber Management

DNS Name Server Address Overview

domain-name-server-inet6 (Routing Instances and Access Profiles)

IN THIS SECTION

- Syntax | 5890
- Hierarchy Level | 5890
- Description | 5890
- Options | 5890
- Required Privilege Level | 5890
- Release Information | 5891

Syntax

```
domain-name-server-inet6 dns-address;
```

Hierarchy Level

```
[edit access],
[edit access profile]
```

Description

Configure an IPv6 address for a DNS name server. You can configure an address globally for a routing instance at the `[edit access]` hierarchy level or for an access profile at the `[edit access profile profile-name]` hierarchy level. You can configure more than one address by including the statement multiple times.

Options

dns-address	IPv6 address of the DNS name server.

Required Privilege Level

admin—To view this statement in the configuration

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring DNS Name Server Addresses for Subscriber Management

DNS Name Server Address Overview

domain-name-server (Routing Instances and Access Profiles)

IN THIS SECTION

- Syntax | 5891
- Hierarchy Level | 5891
- Description | 5892
- Options | 5892
- Required Privilege Level | 5892
- Release Information | 5892

Syntax

```
domain-name-server dns-address;
```

Hierarchy Level

```
[edit access],
[edit access profile]
```

Description

Configure an IPv4 address for a DNS name server. You can configure an address globally for a routing instance at the `[edit access]` hierarchy level or for an access profile at the `[edit access profile profile-name]` hierarchy level. You can configure more than one address by including the statement multiple times.

> **NOTE**: A DNS name server address configured with this statement is less preferred than one configured with the "domain-name-server-inet" on page 5888 statement. That is, the server with the address configured with the "domain-name-server-inet" on page 5888 takes precedence over a server configured with this statement.

Options

dns-address	IPv4 address of the DNS name server.

Required Privilege Level

admin—To view this statement in the configuration

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring DNS Name Server Addresses for Subscriber Management

DNS Name Server Address Overview

domain-name (Static Subscribers)

IN THIS SECTION

Syntax | 5893

- Hierarchy Level | 5893
- Description | 5893
- Options | 5894
- Required Privilege Level | 5894
- Release Information | 5894

Syntax

```
domain-name domain-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication username-include],
[edit logical-systems logical-system-name system services static-subscribers authentication
username-include],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],
[edit system services static-subscribers authentication username-include],
[edit system services static-subscribers group group-name authentication username-include]
```

Description

Specify the domain name that is included at the end of the username created for all static subscribers or for the static subscribers in a specified group. The group version of the statement takes precedence over the global version.

Options

domain-name—Domain name that ends the username created for all static subscribers. The username is also sent to RADIUS in the Access-Request message. The string can include the following characters: a through z, A through Z, 0 through 9, "-", or ".".

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Configuring the Static Subscriber Global Username

Configuring the Static Subscriber Group Username

domain-path-id

IN THIS SECTION

- Syntax | 5895
- Hierarchy Level | 5895
- Description | 5895
- Required Privilege Level | 5896
- Options | 5896
- Release Information | 5897

Syntax

```
[edit routing-options]
 domain-path-id {
     <domain-id-string>;
 }
```

Hierarchy Level

```
[edit routing-options]
```

Description

NOTE: In Junos OS Release 22.2R1 and 22.3R1, and Junos OS Evolved Releases and 22.2R1 and 22.3R1, the `dpath-domain-id` statement was called `uniform-propagation-mode` with a `domain-id` option. The `uniform-propagation-mode` statement with the `domain-id` option has the same functionality as the `dpath-domain-id` statement.

Starting in Junos OS Release 22.3R2 and Junos OS Evolved Release 22.3R2, the `dpath-domain-id` statement is now called `domain-path-id`.

Starting in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1, the domain-path-id is neither exported nor imported by default. Refer to "domain-path-id (Protocols BGP)" on page 5898 for more information.

When you configure the `domain-path-id` statement on a gateway provider edge (PE) device, you're enabling the BGP D-PATH attribute for all of the gateway PE's virtual routing and forwarding (VRF) instances. You can use the BGP D-PATH attribute for intersubnet forwarding (ISF) and subsequent address family identifiers (SAFIs) that advertise IP prefix routes. The BGP D-PATH attribute supports EVPN Type 5 and IP-VPN routes.

The BGP D-PATH attribute enables a gateway provider edge (PE) device to add a domain ID to a BGP route. In many situations, a tenant system is located in different data centers, where an IP Virtual Private Network (IP-VPN) is configured in the WAN. In this situation, gateway PEs aren't aware of the domains they traverse and can install suboptimal routes or routing loops in their IP virtual routing and forwarding (IP-VRF) tables. Gateway PEs can then send BGP updates that advertise those routes to other gateway PEs.

The BGP D-PATH attribute extends its path selection algorithm to prevent routing loops. This attribute also ensures that the gateway PEs install the best routes in their IP-VRF tables. When a gateway PE receives an ISF route, it imports the route into its IP-VRF table and adds the BGP D-PATH attribute to its existing path attributes. The gateway PE uses the information in the BGP D-PATH attribute and the ISF routes in its IP-VRF table to identify the domains through which EVPN IP prefix routes and IP-VPN routes have traversed. With this information, the gateway PE determines if a route should be installed in its IP-VRF table. If the route forms a loop, the gateway PE doesn't install the route in its IP-VRF table.

The BGP D-PATH attribute contains a sequence of domain segments. Each domain segment includes the following fields:

- Domain segment length

  This segment is a 1-octet field that contains the number of domains in a domain segment.

- Domain ID

  This segment is a 6-octet field that represents a particular domain. This segment also contains a 4-octet global administrator sub-field and a 2-octet local administrator sub-field. The global administrator sub-field can contain an AS number, an IPv4 address, or any other value that guarantees the uniqueness of the domain ID.

- ISF_SAFI_TYPE

  This segment is a 1-octet field that identifies the ISF route and the SAFI.

  Table 31: ISF_SAFI Types and Values

ISF_SAFI Types	Values
IPVPN (SAFI 128)	128
EVPN (SAFI 128)	70

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Options

domain-id-string—(Required) Specify the domain ID type and value.

NOTE: You can configure only two domain IDs.

The domain ID contains the following elements:

- Autonomous system (AS) number, IPv4 address, or 4-octet value.

  NOTE: If you use a 4-octet value instead of the AS number or IPv4 value, the value needs to be unique.

- Local administration subfield.

  Configure a local 2-octet value between 0 to 65535 (0xFFFF) for your local data center.

- Inter-subnet forwarding identifier (ISF).
- Subsequent address family identifier (SAFI).

Use this format to create the domain ID type: `<ASN | IPv4 address | 4-octet value>:local administration subfield:ISF_SAFI_TYPE`.

Here's an example that configures two EVPN ISF_SAFI_TYPE domain IDs with AS number 6500, and local administration subfields 1 and 2 for ISF_SAFI types EVPN and IPVPN, respectively:

```
domain-path-id {
    6500:1:evpn;
    6500:2:ipvpn;
}
```

Release Information

Statement introduced in Junos OS Release 22.2R1.

domain-path-id (Protocols BGP)

IN THIS SECTION

- Syntax | 5898
- Hierarchy Level | 5898
- Description | 5898
- Options | 5899
- Required Privilege Level | 5899
- Release Information | 5899

Syntax

```
domain-path-id {
 receive;
 send;
}
```

Hierarchy Level

```
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor neighbor-address],
[edit routing-instances instance-name protocols bgp],
[edit routing-instances instance-name protocols bgp group group-name],
[edit routing-instances instance-name protocols bgp group group-name neighbor neighbor-address]
```

Description

You can use the `receive` statement to control importing the BGP domain path attribute. You can use the `send` statement to control exporting the BGP domain path attribute. The BGP domain path attribute is not exported or imported by default.

Options

receive Allow receiving BGP domain path attribute

send Allow sending BGP domain path attribute

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

`receive` statement introduced in Junos OS and Junos OS Evolved Release 23.2R1.

`send` statement introduced in Junos OS and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

domain-path-id | 5894

domain-search

IN THIS SECTION

- Syntax | 5900
- Hierarchy Level | 5900
- Description | 5900
- Options | 5900
- Required Privilege Level | 5900
- Release Information | 5901

Syntax

```
domain-search [domain-list ];
```

Hierarchy Level

```
[edit system],
[edit system services dhcp],
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]
```

Description

Configure a list of domains to search (in the case where you want to configure access to multiple DNS servers for redundancy, and/or to resolve hosts that the previous server could not).

Options

domain-list List of domain servers to search. The list can contain up to six domain names, separated by a space, with a total of up to 256 characters.

For example to search `domain1.net`, and if it fails to resolve the host, `domain2.net`, and if fails to resolve the host, `domain3.net`, you would configure the following domain list at the `domain-search` hierarchy level:

```
[edit system]
set domain-search [domain1.net domain2.net domain3.net]
```

NOTE: In gNMI, the user cannot delete individual domain lists. The entire domain search command, including all the configured domain lists, must be deleted.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

domain-selection

IN THIS SECTION

- Syntax | 5901
- Hierarchy Level | 5903
- Description | 5904
- Required Privilege Level | 5904
- Release Information | 5904

Syntax

```
domain-selection {
    term term-name {
        from {
            3gpp-imsi {
                equals value;
                has-prefix value;
                has-suffix value;
                matches value;
            }
            attribute name {
                code numeric-code;
                vendor-id vendor-id;
                format {
                    integer {
                        equals {
                            value;
                        }
                        greater-than value;
                        less-than value;
```

```
                }
                string {
                    equals {
                        value;
                    }
                    has-prefix {
                        value;
                    }
                    has-suffix {
                        value;
                    }
                    matches {
                        value;
                    }
                }
                time {
                    equals {
                        value;
                    }
                    greater-than value;
                    less-than value;
                }
                v4address {
                    equals {
                        value;
                    }
                }
                v6address {
                    equals {
                        value;
                    }
                }
                v6prefix {
                    equals {
                        value;
                    }
                }
            }
        }
        called-station-id {
            equals value;
            matches value;
        }
```

```
            calling-station-id {
                equals value;
                matches value;
            }
            class {
                equals value;
                has-prefix value;
                has-suffix value;
                matches value;
            }
            client client-name;
            framed-ip-address {
                equals value;
            }
            framed-ipv6-prefix {
                equals value;
            }
            nas-ip-address {
                equals value;
            }
            snoop-segment snoop-segment-name;
            user-name {
                equals value;
                has-prefix value;
                has-suffix value;
                matches value;
            }
        }
        then {
            domain tdf-domain-name;
            pcef-profile pcef-profile-name;
        }
    }
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name]
```

Description

Specify the TDF domain to be used for an IP-based subscriber. You can configure multiple terms under `domain-selection`, and each term is applied in the order in which it is configured. You can specify multiple match conditions within a term and all of the conditions have to match. If the incoming RADIUS request from the subscriber matches the criteria in a term, then the TDF domain specified in the `then` statement of the term is used to create the TDF subscriber.

You can also specify a PCEF profile for an IP-based subscriber. This is required if the TDF domain selected for a subscriber does not specify a PCEF profile or you want to allow different members of the same TDF domain to have different PCEF profiles.

After a term matches and a TDF domain is selected, further terms are not evaluated if the PCEF profile is specified in either the `then` statement or in the selected TDF domain. If a PCEF profile is not specified in either the `then` statement or in the selected TDF domain, further terms are evaluated to find a PCEF profile for the subscriber.

If no TDF domain is selected for a subscriber, then a TDF subscriber session is not created.

> **NOTE**: The TDF domain must have been previously configured at the `[edit unified-edge gateways tdf` *`gateway-name`* `domains]` hierarchy level.
>
> The PCEF profile must have been previously configured at the `[edit unified-edge pcef]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

domain (TDF Domain Selection)

IN THIS SECTION

- Syntax | 5905
- Hierarchy Level | 5905
- Description | 5905
- Options | 5906
- Required Privilege Level | 5906
- Release Information | 5906

Syntax

```
domain tdf-domain-name;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name then]
```

Description

Specify the TDF domain to be selected when the criteria specified in the domain selection statement are matched.

NOTE: This statement is required even if you have not specified any match criteria.

Options

tdf-domain-name	Name of the TDF domain to use.

> **NOTE**: The TDF domain must have been previously configured at the `[edit unified-edge gateways tdf gateway-name domains]` hierarchy level.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers*
- *Understanding Selection of Properties for an IP-Based TDF Subscriber*
- *Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber*
- *IP-Based Subscriber Setup Overview*

domain-type (Bridge Domains)

IN THIS SECTION

- Syntax | 5907
- ACX Series and MX Series | 5907
- SRX Series | 5907
- Description | 5907
- Required Privilege Level | 5908

Release Information | 5908

Syntax

```
domain-type bridge;
```

ACX Series and MX Series

```
[edit bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]
```

SRX Series

```
[edit bridge-domains bridge-domain-name]
```

Description

Define the domain type `bridge` for a Layer 2 bridge domain.

NOTE: There is only one domain type `bridge`, that can be configured on SRX Series Firewalls. Domain type `bridge` is not enabled by default. An SRX Series Firewall operates in the Layer 2 transparent mode when all physical bridge domains on the device are partitioned into logical bridge domains.

NOTE: Starting with Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the CLI `domain-type` is not available.

> **NOTE**: Starting in Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the hierarchy `[edit bridge-domains bridge-domain-name]` is renamed to `[edit vlans vlan-name]`. For detailed information about the modified hierarchies, see *Enhanced Layer 2 CLI Configuration Statement and Command Changes for Security Devices*.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement modified in Junos OS Release 9.5.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Ethernet Switching and Layer 2 Transparent Mode Overview

Configuring a Bridge Domain

Configuring a Layer 2 Virtual Switch

domains

IN THIS SECTION

- Syntax | 5909
- Hierarchy Level | 5911
- Description | 5911
- Options | 5911
- Required Privilege Level | 5911

Release Information | 5912

Syntax

```
domains domain-name {
    apply-groups [group-names];
    apply-groups-except [group-names];
    burst-size {
        apply-groups [group-names];
        apply-groups-except [group-names];
        downlink downlink-burst-size;
        uplink uplink-burst-size ;
    }
    ifl-subscriber [subscriber-name] {
        access-interfaces [interface-name];
        apply-groups [group-names];
        apply-groups-except [group-names];
    }
    ip-subscriber {
        access-interfaces [interface-name];
        apply-groups [group-names];
        apply-groups-except [group-names];
        default-local-policy {
            flow-action (drop | forward);
            maximum-bit-rate {
                uplink mbr-uplink-value ;
                downlink mbr-downlink-value;
            }
            burst-size {
                uplink uplink-burst-size;
                downlink downlink-burst-size;
            }
        }
        idle-timeout idle-timout;
        immediate-accounting-response (enabled | disabled);
        maximum-subscribers number;
        subscriber-address {
            apply-groups [group-names];
            apply-groups-except [group-names];
```

```
            inet {
                apply-groups [group-names];
                apply-groups-except [group-names];
                pool pool-name;
            }
            inet6 {
                apply-groups [group-names];
                apply-groups-except [group-names];
                pool pool-name;
            }
        }
        subscription-id {
            apply-groups [group-names];
            apply-groups-except [group-names];
            constant ;
            subscription-id-options {
                entry-name {
                    id-components {
                        use-imsi;
                        use-msisdn;
                        use-nai;
                        use-username;
                        use-realm;
                        use-nas-port;
                        use-nas-port-id;
                    }
                }
            }
        }
    }
    maximum-bit-rate {
        apply-groups [group-names];
        apply-groups-except [group-names];
        downlink mbr-downlink-value;
        uplink mbr-uplink-value;
    }
    pcef-profile name;
    service-mode service-mode-options;
    subscriber-exclude-prefix {
        apply-groups [group-names];
        apply-groups-except [group-names];
        family {
            inet {
```

```
                apply-groups [group-names];
                apply-groups-except [group-names];
                network address net-mask;
            }
            inet6 {
                apply-groups [group-names];
                apply-groups-except [group-names];
                network address net-mask;
            }
        }
    }
    subscriber-type (ip | ifl);
    tdf-interface  mif.number;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name]
```

Description

Configure a TDF domain, which specifies a set of properties for creating TDF subscriber sessions and for handling subscriber traffic.

Options

domain-name Name of the TDF domain.

- **Range:** 1 through 50 alphanumeric characters.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IFL-Based Subscriber Properties with a TDF Domain

dont-help-shared-fate-bfd-down

IN THIS SECTION

- Syntax | 5912
- Hierarchy Level | 5912
- Description | 5913
- Default | 5913
- Required Privilege Level | 5913
- Release Information | 5913

Syntax

```
dont-help-shared-fate-bfd-down
```

Hierarchy Level

```
[edit protocols bgp graceful-restart]
```

Description

When BFD is control plane dependent and the device detects a BFD down event and is not already entering the graceful restart helper mode, this is treated as a regular BFD down event and the device enters the graceful restart helper mode. This behavior makes the control plane dependent BFD unusable in conjunction with graceful restart.

Include the `dont-help-shared-fate-bfd-down` statement at the `[edit protocols bgp graceful-restart]` hierarchy to ensures that the device does not enter the graceful restart helper mode and data traffic continues to be fowarded to an alternate path even if there is an interface failure (without a control plane restart on the BGP neighbor).

Default

By default, this option is not enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Configuring Routing Protocols Graceful Restart

Understanding External BGP Peering Sessions

dot1p-priority

IN THIS SECTION

- Syntax | 5914
- Hierarchy Level | 5914

- Description | 5914
- Options | 5914
- Required Privilege Level | 5914
- Release Information | 5914

Syntax

```
dot1p-priority number;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Specify the IEEE 802.1p priority to be used in the transmitted RAPS protocol data units.

Options

number—802.1p priority number.

- **Range:** 0 through 7
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

dot1q-tunneling

IN THIS SECTION

- Syntax (Ethernet Switching) | 5915
- Syntax (VLANs) | 5915
- Ethernet Switching | 5916
- VLANs | 5916
- Description | 5916
- Required Privilege Level | 5916
- Release Information | 5916

Syntax (Ethernet Switching)

```
dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
}
```

Syntax (VLANs)

```
dot1q-tunneling {
    customer-vlans (id | native | range);
    layer2-protocol-tunneling all | protocol-name {
        drop-threshold number;
        shutdown-threshold number;
    }
}
```

Ethernet Switching

```
[edit ethernet-switching-options]
```

VLANs

```
[edit vlans vlan-name]
```

Description

For Ethernet switching, sets a global value for the EtherType for Q-in-Q tunneling.

For VLANs, enables Q-in-Q tunneling on the specified VLAN.

NOTE:

- The VLAN on which you enable Q-in-Q tunneling must be a tagged VLAN.
- You cannot configure 802.1X user authentication on interfaces that have been enabled for Q-in-Q tunneling.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Option **native** introduced in Junos OS Release 9.6 for EX Series switches.

Options **layer2-protocol-tunneling**, **drop-threshold**, and **shutdown-threshold** introduced in Junos OS Release 10.0 for EX Series switches.

RELATED DOCUMENTATION

Example: Setting Up Q-in-Q Tunneling on EX Series Switches

Configuring Q-in-Q Tunneling on EX Series Switches

Configuring Q-in-Q Tunneling on QFX Series Switches

Example: Setting Up Q-in-Q Tunneling on QFX Series Switches

dot1x

IN THIS SECTION

- Syntax | 5917
- Hierarchy Level | 5919
- Description | 5919
- Default | 5919
- Options | 5919
- Required Privilege Level | 5920
- Release Information | 5920

Syntax

```
dot1x {
    authenticator {
        authentication-profile-name  access-profile-name;
          interface (all | [ interface-names ]) {
            authentication-order (captive-portal | dot1x | mac-radius);
            disable;
            guest-bridge-domain guest-bridge-domain;
            guest-vlan guest-vlan;
            ignore-port-bounce;
            mac-radius {
                authentication-protocol {
                    eap-md5;
                    eap-peap {
```

```
                        resume;
                    }
                    pap;
                }
                flap-on-disconnect;
                restrict;
            }
            maximum-requests number;
            multi-domain {
                max-data-session max-data-session;
                packet-action (drop-and-log | shutdown);
                recovery-timeout seconds;
            }
            (no-reauthentication | reauthentication interval );
            no-tagged-mac-authentication;
            quiet-period seconds;
            redirect-url redirect-url;
            retain-mac-aged-session;
            retries number;
            server-fail (bridge-domain bridge-domain | deny | permit | use-cache | vlan-name
vlan-name);
            server-fail-voip (deny | permit | use-cache | vlan-name vlan-name);
            server-reject-bridge-domain bridge-domain {
                block-interval seconds;
                eapol-block;
            }
            server-reject-vlan (vlan-id | vlan-name)  {
                block-interval block-interval;
                eapol-block;
            }
            server-timeout seconds;
            supplicant (single | single-secure | multiple);
            supplicant-timeout seconds;
            transmit-period seconds;
        }
        ip-mac-session-binding;
        no-mac-table-binding;
        radius-options {
            add-interface-text-description;
            use-vlan-id;
            use-vlan-name;
        }
        static mac-address {
```

```
            bridge-domain-assignment bridge-domain-assignment;
            interface interface;
            vlan-assignment vlan-identifier;
        }
    }
    }
    ssl-certificate-path path-name;
    traceoptions {
        file filename <files files>  <size size> <(world-readable | no-world-readable)>;
        flag (all | config-internal | dot1x-debug | dot1x-event | dot1x-ipc | eapol | esw-if |
general | iccp | normal | parse | state | task | timer | vlan) {
            disable;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name protocols],
[edit protocols]
```

Description

Configure IEEE 802.1X authentication for Port-Based Network Access Control. 802.1X authentication is supported on interfaces that are members of private VLANs (PVLANs).

Default

802.1X is disabled.

Options

ssl-certificate-path *path-name*

Specify the file path for SSL certificates if you are not using the default path. The default path for SSL certificates is `/var/tmp`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`ssl-certificate-path` introduced in Junos OS Release 19.4.

`ip-mac-session-binding` introduced in Junos OS Release 20.2

RELATED DOCUMENTATION

show dot1x | 19548

Example: Setting Up 802.1X for Single-Supplicant or Multiple-Supplicant Configurations on an EX Series Switch

Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an EX Series Switch

Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch

Example: Configuring Static MAC Bypass of 802.1X and MAC RADIUS Authentication on an EX Series Switch

Example: Configuring MAC RADIUS Authentication on an EX Series Switch

Configuring RADIUS Server Fail Fallback (CLI Procedure)

dot1x (MX Series in Enhanced LAN Mode)

IN THIS SECTION

- Syntax | 5921
- Hierarchy Level | 5921
- Description | 5921
- Default | 5921
- Required Privilege Level | 5922

Release Information | 5922

Syntax

```
dot1x {
    disable;
    guest-vlan (vlan-id  / vlan-name);
    mac-radius {
        flap-on-disconnect;
        restrict;
    }
    maximum-requests number;
    no-reauthentication;
    server-fail (deny | permit | use-cache | vlan-id | vlan-name);
    server-reject-vlan (vlan-id | vlan-name)  {
        eapol-block;
        block-interval block-interval;
    }
    supplicant-timeout seconds;
    transmit-period seconds;
}
```

Hierarchy Level

```
[edit protocols authentication-access-control interface (all | interface-names)]
```

Description

Configure 802.1X authentication for Port-Based Network Access Control. 802.1X authentication is supported on interfaces that are members of private VLANs (PVLANs).

The remaining statements are explained separately. See CLI Explorer.

Default

802.1X is disabled.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

downlink-dscp-to-egress-forwarding-class

IN THIS SECTION

- Syntax | 5922
- Hierarchy Level | 5922
- Description | 5923
- Required Privilege Level | 5923
- Release Information | 5923

Syntax

```
downlink-dscp-to-egress-forwarding-class {
    classifier classifier-name;
    classifier classifier-name; {
        routing-instance instance-name;
    }
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw system class-of-service]
```

Description

You will use these classifiers to map DiffServ code point (DSCP) values provided in Transport Level Marking (TLM) to specific forwarding/egress classes on the WAN interface.

The Egress Forwarding Class will be assigned to the Service Data Flow (SDF). All packets matching that SDF inherit the forwarding class. If the Egress Forwarding Class is not assigned to the SDF, the forwarding class assigned by the ingress WAN classifier will be used.

To assign Egress Forwarding Classes to an SDF, Transport Level Marking (TLM) must be provided. If the Fowarding Action (FAR) does not include TLM, the Egress Forwarding Class will not be assigned.

Required Privilege Level

system

Release Information

This statement was introduced in Junos OS Release 23.1R1.

downlink-dscp-to-forwarding-class

IN THIS SECTION

- Syntax | 5923
- Hierarchy Level | 5924
- Description | 5924
- Required Privilege Level | 5924
- Release Information | 5924

Syntax

```
downlink-dscp-to-forwarding-class {
    classifier classifier-name;
    classifier classifier-name {
```

```
        routing-instance instance-name;
    }
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw system class-of-service]
```

Description

Assign the classifiers to use for mapping the DiffServ code point (DSCP) value provided in the Transport Level Marking to the forwarding queues.

The DSCP value is relevant in the context of the routing instance associated with the outgoing interface. Therefore, you need to assign classifiers to peer routing instances. If you do not assign any routing instance to a classifier, the system uses the classifier for all routing instances without an assigned classifier.

See CLI Explorer for the remaining statements.

Required Privilege Level

system

Release Information

Juniper Networks introduced this statement in Junos OS Release 22.2R1.

ignore-server-validation (Application Identification)

IN THIS SECTION

- Syntax | 5925
- Hierarchy Level | 5925
- Description | 5925

Required Privilege Level | 5925
Release Information | 5925

Syntax

```
ignore-server-validation;
```

Hierarchy Level

```
[edit services application-identification download]
```

Description

Disable server authentication for application signature package download.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1 on SRX Series Firewalls.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Application Identification Overview

Downloading and Installing Predefined Junos OS Application Signature Packages

download-profile (Security Antivirus FTP)

IN THIS SECTION

- Syntax | 5926
- Hierarchy Level | 5926
- Description | 5926
- Required Privilege Level | 5926
- Release Information | 5926

Syntax

```
download-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name anti-virus ftp]
```

Description

Configure a Content Security policy for the antivirus FTP (download) protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The [edit security utm default-configuration] hierarchy level is introduced in Junos OS Release 18.2R1.

download-profile (Security Content Filtering FTP)

IN THIS SECTION

- Syntax | 5927
- Hierarchy Level | 5927
- Description | 5927
- Required Privilege Level | 5927
- Release Information | 5928

Syntax

```
download-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name content-filtering ftp]
```

Description

Configure a Content Security policy for the content-filtering FTP (download) protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

download (Services)

IN THIS SECTION

- Syntax | 5928
- Hierarchy Level | 5929
- Description | 5929
- Options | 5929
- Required Privilege Level | 5929
- Release Information | 5930

Syntax

```
download {
    automatic {
        interval hours;
        minor-only;
        start-time start-time;
    }
    ignore-server-validation;
    proxy-profile proxy-profile;
    url url;
}
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Configure automatic download for the application identification services application package.

The application package contains definitions for known applications, such as: DNS, Facebook, FTP, Skype, and SNMP. The application package is located at **https://signatures.juniper.net**. If you do not have access to the default download site from your device, you can use the URL option to download from a different location.

> **NOTE**: You need to download the application package before configuring application identification services.

Options

automatic Download the application package automatically at a certain time of day or at intervals.

- *interval*—Download the application package at intervals.

  Range: 6 through 720 hours

- minor-only—Specify signature package version type as minor.
- *start-time*—Start time in which the application package will be download. Format is MM-DD.hh:mm. Example: 04-15.09:00 will start the download on April 15 at 9 AM.

ignore-server-validation Disable server authentication for Applicaton Signature download

proxy-profile Configure web proxy for Application signature download

url URL for application package download. Use this option to change the default download location of the application package.

Required Privilege Level

For Routing Devices

- interface—To view this statement in the configuration.

- interface-control—To add this statement to the configuration.

For Security Devices

- security—To view this statement in the configuration.
- security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5 for Routing Devices and Junos OS Release 10.2 for Security Devices

RELATED DOCUMENTATION

Example: Scheduling the Application Signature Package Updates

download-timeout

IN THIS SECTION

- Syntax | 5930
- Hierarchy Level | 5931
- Description | 5931
- Options | 5931
- Required Privilege Level | 5931
- Release Information | 5931

Syntax

```
download-timeout minutes;
```

Hierarchy Level

```
[edit security idp security-package automatic]
```

Description

Specify the time that the device automatically times out and stops downloading the updated signature database from the specified URL.

> **NOTE**: The default value for download-timeout is one minute. If download is completed before the download times out, the signature is automatically updated after the download. If the download takes longer than the configured period, the automatic signature update is terminated.

Options

minutes—Time in minutes.

- **Range:** 1 through 60 minutes
- **Default:** 1 minute

> **NOTE**: For SRX Series Firewalls the applicable range is 1 through 4000000 per second.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 9.6R3 of Junos OS.

downstream

IN THIS SECTION

- Syntax | 5932
- Hierarchy Level | 5932
- Description | 5932
- Options | 5932
- Required Privilege Level | 5933
- Release Information | 5933

Syntax

```
downstream name;
```

Hierarchy Level

```
[edit logical-systems name protocols router-advertisement interface],
[edit protocols router-advertisement interface]
```

Description

Configure the list of downstream interfaces for RA proxy

Options

name Configure the downstream interface for RA proxy

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1R1

downstream-mode

IN THIS SECTION

- Syntax | 5933
- Hierarchy Level | 5933
- Description | 5934
- Required Privilege Level | 5934
- Release Information | 5934

Syntax

```
downstream-mode;
```

Hierarchy Level

```
[edit logical-systems name protocols router-advertisement interface],
[edit protocols router-advertisement interface]
```

Description

Configure the interface as downstream interface for RA proxy

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1R1

downstream-on-demand

IN THIS SECTION

- Syntax | 5934
- Hierarchy Level | 5934
- Description | 5935
- Required Privilege Level | 5935
- Release Information | 5935

Syntax

```
downstream-on-demand;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp session session-address],
[edit protocols ldp session session-address]
```

Description

Enable LDP downstream on demand on the LDP session. LDP is widely deployed in downstream unsolicited advertisement mode. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between access and aggregation networks to minimize the workload for the access node (AN) control plane and to avoid the storage of tens of thousands of label bindings from upstream aggregation nodes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring LDP Downstream on Demand

downstream-rate (Traffic Shaping)

IN THIS SECTION

- Syntax | 5936
- Hierarchy Level | 5936
- Description | 5936
- Options | 5936
- Required Privilege Level | 5936
- Release Information | 5936

Syntax

```
downstream-rate rate;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-interface-
unit advisory-options],
[edit dynamic-profiles profile-name interfaces interface-set $junos-interface-set-name interface
$junos-interface-ifd-name advisory-options],
[edit interfaces demux0 unit logical-unit-number advisory-options],
[edit interfaces interface-name unit logical-unit-number advisory-options]
```

Description

Specify a recommended shaping rate to be applied to downstream traffic on an interface.

For ANCP interfaces, this configured rate is used as the default value for the Juniper VSA Downstream-Calculated-Qos-Rate (26–141) when the router has not received and processed the attributes from the access node.

For L2TP, the rate is configured on an underlying PPPoE logical interface for a subscriber on an MX Series router acting as a LAC. When the subscriber is tunneled, this rate, referred to as speed for L2TP, is sent to the LNS in the ICCN message as AVP 24.

Options

rate—Traffic rate in bits per second.

- **Range:** 1000 through 4,294,967,295 bits per second

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the `[edit interfaces demux0 ...]` hierarchy level introduced in Junos OS Release 12.2.

Support at the `[edit dynamic-profiles ...]` hierarchy level introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces

Configuring the ANCP Agent

Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS

dr-election-on-p2p

IN THIS SECTION

- Syntax | 5937
- Hierarchy Level | 5937
- Description | 5938
- Default | 5938
- Required Privilege Level | 5938
- Release Information | 5938

Syntax

```
dr-election-on-p2p;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Enable PIM designated router (DR) election on point-to-point (P2P) links.

Default

No PIM DR election is performed on point-to-point links.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Configuring PIM Designated Router Election on Point-to-Point Links

dr-register-policy

IN THIS SECTION

- Syntax | 5939
- Hierarchy Level | 5939
- Description | 5939
- Options | 5939
- Required Privilege Level | 5939
- Release Information | 5939

Syntax

```
dr-register-policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Apply one or more policies to control outgoing PIM register messages.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

> *Configuring Register Message Filters on a PIM RP and DR*
>
> rp-register-policy | **16000**

drain

IN THIS SECTION

- Syntax | 5940
- Hierarchy Level | 5940
- Description | 5940
- Required Privilege Level | 5940
- Release Information | 5941

Syntax

```
drain;
```

Hierarchy Level

```
[edit services l2tp],
[edit services l2tp destination address ip-address],
[edit services l2tp destination address ip-address routing-instance routing-instance-name],
[edit services l2tp destination name destination-name],
[edit services l2tp tunnel name name],
[edit services l2tp tunnel name name address ip-address],
[edit services l2tp tunnel name name address ip-address routing-instance routing-instance-name]
```

Description

Prevent the creation of new sessions, destinations, and tunnels globally at an L2TP access concentrator (LAC) or an L2TP network server (LNS). Prevent the creation of new tunnels and sessions for a specific destination. Prevent the creation of new sessions for a specific tunnel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

L2TP LAC Subscriber Configuration

L2TP LNS Inline Service Interfaces

draining (Diameter Applications)

IN THIS SECTION

- Syntax | 5941
- Hierarchy Level | 5941
- Description | 5942
- Required Privilege Level | 5942
- Release Information | 5942

Syntax

```
draining;
```

Hierarchy Level

```
[edit access pcrf partition partition-name],
[edit access ocs partition partition-name]
```

Description

Configure the Policy and Charging Rule Function (PCRF) or Online Charging System (OCS) partition to the draining state to make substantial configuration changes quickly. To log out all subscribers quickly and immediately, issue the `clear network-access pcrf subscribers` command.

After you set draining for either of the PCRF or OCS partitions, any new subscriber logins are denied, and the time limit you set in the `draining-response-timeout` statement is used instead of the `logout-response-timeout` time limit.

After the partition has drained, you must issue the `clear network-access pcrf` or `clear network-access ocs statistics` command to clear out the respective partition's subscribers and wait for all of the subscribers to log out before you can make configuration changes and resume normal operations.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

draining-response-timeout (Diameter Applications)

IN THIS SECTION

- Syntax | 5943
- Hierarchy Level | 5943
- Description | 5943
- Options | 5943
- Required Privilege Level | 5944
- Release Information | 5944

Syntax

```
draining-response-timeout seconds;
```

Hierarchy Level

```
[edit access pcrf partition partition-name],
[edit access ocs partition partition-name]
```

Description

Configure the amount of time in seconds before a PCRF or an Online Charging System OCS partition responds and begins to drain. Configuring this statement is optional if you set `draining` statement for either of the Policy and Charging Rule Function (PCRF) or Online Charging System (OCS) partitions.

After you set draining for either of the PCRF or OCS partitions, any new subscriber logins are denied, and the time limit you set in this statement is used instead of the `logout-response-timeout` time limit.

Options

seconds Number of seconds to wait before the partition begins to drain.

- **Default:** 400
- **Range:** 0 through 86,400 seconds (24 hours)

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

drop (BPDU Block)

IN THIS SECTION

- Syntax | 5945
- Hierarchy Level | 5945
- Description | 5945
- Default | 5945
- Required Privilege Level | 5945
- Release Information | 5945

Syntax

```
drop;
```

Hierarchy Level

```
[edit ethernet-switching-options bpdu-block bpdu-block  (all | [interface-name])]
```

Description

Drop bridge protocol data units (BPDUs) that enter any interface or a specified interface, but do not disable the interface. Configure the `drop` statement *only* on interfaces on which no spanning-tree protocol (STP, MSTP, or RSTP) is configured.

> **NOTE**: Do not configure `drop` on any interface on which a spanning-tree protocol has been configured. Doing so could cause STP misconfiguration and consequent network outages.

Default

Not enabled.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Understanding BPDU Protection for STP, RSTP, and MSTP

Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on non-ELS EX Series Switches

Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches

bpdu-block-on-edge | 3024

drop (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 5946
- Hierarchy Level | 5946
- Description | 5947
- Required Privilege Level | 5947
- Release Information | 5947

Syntax

```
drop;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option (default-action | equals | starts-with)],
[edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action | equals | starts-with)],
[edit forwarding-options dhcp-relay group group-name relay-option (default-action | equals | starts-
with)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name relay-option (default-action | equals
| starts-with)],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Drop (discard) specified DHCP client packets when you use DHCP relay agent selective processing. You can configure the drop operation globally or for a group of interfaces, and for either DHCP or DHCPv6 relay agent.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic

drop-flow

IN THIS SECTION

- Syntax | 5948
- Hierarchy Level | 5948
- Description | 5948
- Options | 5948
- Required Privilege Level | 5948
- Release Information | 5949

Syntax

```
drop-flow {
 max-sessions percent;
 timeout seconds;
}
```

Hierarchy Level

```
[edit security flow]
```

Description

To prevent the probable security attack, we introduce `drop-flow` configuration. Sessions in the drop-flow is valid for 4 seconds by default.

Options

max-sessions

Maximum Drop-flow Sessions (default 10%) (percent)

- **Default:** 10
- **Range:** 0 through 30

timeout

Timeout (default 4s) (seconds)

- **Default:** 4
- **Range:** 2 through 300

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

This statement is introduced in Junos OS Release 23.3R1.

drop-member-traffic (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 5949
- Hierarchy Level | 5949
- Description | 5949
- Default | 5950
- Required Privilege Level | 5950
- Release Information | 5950

Syntax

```
drop-member-traffic {
    rejoin-timeout rejoin-timeout;
}
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options member-failure-options]
```

Description

Specify whether the broadband gateway should drop traffic to a services PIC when it fails.

For many-to-one (N:1) high availability (HA) for service applications like Network Address Translation (NAT), this configuration is valid only when two or more services PICs have failed.

The remaining statement is explained separately. See CLI Explorer.

Default

If this statement is not configured, then the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

member-failure-options (Aggregated Multiservices) | **10608**

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

drop-policy (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | **5951**
- Hierarchy Level | **5951**
- Description | **5951**
- Options | **5951**
- Required Privilege Level | **5951**
- Release Information | **5951**

Syntax

```
drop-policy rule-name {
    from {
        apply-groups group-name;
        apply-groups-except group-name;
        destination-address address;
        destination-port port-number;
        dscp dscp-value;
        protocol protocol;
        source-address address;
        source-port port-number;
    }
}
```

Hierarchy Level

```
[edit services radius-flow-tap policy policy-name inet| inet6]
```

Description

Specify the drop-policy that is applied to mirrored packets sent to a mediation device.

Options

rule-name–Define the term name.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Subscriber Secure Policy Overview

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

drop (PPPoE Service Name Tables)

IN THIS SECTION

- Syntax | 5952
- Hierarchy Level | 5952
- Description | 5952
- Required Privilege Level | 5953
- Release Information | 5953

Syntax

```
drop;
```

Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name],
[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci
circuit-id-string ari remote-id-string]
```

Description

Direct the router to drop (ignore) a PPPoE Active Discovery Initiation (PADI) control packet received from a PPPoE client that contains the specified service name tag or agent circuit identifier/agent remote identifier (ACI/ARI) information. This action effectively denies the client's request to provide the specified service, or to accept requests from the subscriber or subscribers represented by the ACI/ARI information.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at `[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci circuit-id-string ari remote-id-string]` hierarchy level introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

drop-probability

IN THIS SECTION

- QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | **5953**
- QFX10000 Switches | **5954**
- Hierarchy Level | **5954**
- Description | **5954**
- Options | **5954**
- Required Privilege Level | **5955**
- Release Information | **5955**

QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems

```
drop-probability 0 drop-probability high-value;
```

QFX10000 Switches

```
drop-probability percentage1 percentage2 ... percentage32;
```

Hierarchy Level

```
[edit class-of-service drop-profiles profile-name interpolate]
```

Description

When configuring WRED, map the packet `drop-probability` to the fullness of a queue ("fill-level" on page 7110). You configure the `fill-level` and `drop-probability` statements in related pairs. The pairs of fill level and drop probability values set a probability of dropping packets at a specified queue fullness value.

On switches that support only two fill level/drop probability pairs, the first drop probability is always zero. The first fill level/drop probability pair sets the drop start point, and the second fill level/drop probability pair sets the drop end point.

On switches that support 32 fill level/drop probability pairs, the first fill level/drop probability pair sets the drop start point, and the last fill level/drop probability pair sets the drop end point.

As the queue fills from the drop start point to the drop end point, the rate of packet drop increases in a curve pattern. The higher the queue fill level, the higher the probability of dropping packets.

Options

0 (switches that support only two fill level/drop probability pairs)—Probability that packets will drop at the lowest `fill-level` value. This is always zero, because until the queue reaches the specified low `fill-level` value, no packets are scheduled to drop.

- **Range:** 0

high-value (switches that support only two fill level/drop probability pairs)—The maximum probability that packets will drop before queue fullness exceeds the high value of the queue `fill-level`, expressed as a percentage. If the queue fills beyond the high `fill-level` value, all packets drop.

- **Range:** 0 through 100 percent

percentage1 percentage2 ... percentage32 (switches that support 32 fill level/drop probability pairs)—The probability that packets will drop before the queue fullness exceeds the `fill-level` value, expressed

as a percentage. Each drop probability pairs with a queue fill level to define the probability of a packet dropping at a specified queue fullness.

- **Range:** 0 through 100 percent

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

drop-probability (Fill Level)

IN THIS SECTION

- Syntax | 5955
- Hierarchy Level | 5955
- Description | 5956
- Options | 5956
- Required Privilege Level | 5956
- Release Information | 5956

Syntax

```
drop-probability percentage;
```

Hierarchy Level

```
[edit class-of-service drop-profiles profile-name fill-level percentage]
```

Description

Drop packets at the rate of the drop-probability value when the queue fills to the percentage configured with the fill-level value. This way you can manage network congestion.

Options

percentage—The probability (expressed in percentage) for a packet to be dropped from the queue.

- **Range:** 0 through 100

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS 11.4.

RELATED DOCUMENTATION

Configuring CoS Tail Drop Profiles (CLI Procedure)

Understanding Junos OS CoS Components for EX Series Switches

drop-probability (Interpolated Value)

IN THIS SECTION

- Syntax | 5957
- Hierarchy Level | 5957
- Description | 5957
- Options | 5957
- Required Privilege Level | 5957
- Release Information | 5957

Syntax

```
drop-probability [values];
```

Hierarchy Level

```
[edit class-of-service drop-profiles profile-name interpolate]
```

Description

Define values for interpolated drop probabilities. The maximum number of drop probability values supported per drop profile is based on the line card.

On EX Series switches, this statement is supported only on the EX9200 switch, EX8200 standalone switches, and EX8200 Virtual Chassis.

Options

percentage—The probability (expressed in percentage) for a packet to be dropped from the queue.

- **Range:** 0 through 100

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Managing Congestion Using RED Drop Profiles and Packet Loss Priorities

Defining Packet Drop Behavior by Configuring RED Drop Profiles

drop-probability (Percentage)

IN THIS SECTION

- Syntax | 5958
- Hierarchy Level | 5958
- Description | 5958
- Options | 5958
- Required Privilege Level | 5958
- Release Information | 5959

Syntax

```
drop-probability percentage;
```

Hierarchy Level

```
[edit class-of-service drop-profiles profile-name]
```

Description

Define drop probability percentages. The maximum number of drop probability values supported per drop profile is based on the line card.

Options

percentage—Probability that a packet is dropped, expressed as a percentage. A value of 0 means that a packet is never dropped, and a value of 100 means that all packets are dropped.

- **Range:** 0 through 100 percent

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Managing Congestion Using RED Drop Profiles and Packet Loss Priorities

Defining Packet Drop Behavior by Configuring RED Drop Profiles

drop-profile (Schedulers)

IN THIS SECTION

- Syntax | 5959
- Hierarchy Level | 5959
- Description | 5960
- Options | 5960
- Required Privilege Level | 5960
- Release Information | 5960

Syntax

```
drop-profile profile-name;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name drop-profile-map loss-priority (low | medium-high |
high) protocol protocol]
```

Description

Define drop profiles for weighted random early detection (WRED). When a packet arrives, WRED checks the queue fill level specified in the drop profile. If the fill level corresponds to a nonzero drop probability, the WRED algorithm determines whether to drop the arriving packet.

Options

profile-name—Name of the drop profile.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Drop Profile Maps

drop-profile (Dynamic Schedulers)

IN THIS SECTION

- Syntax | 5961
- Hierarchy Level | 5961
- Description | 5961
- Options | 5961
- Required Privilege Level | 5962
- Release Information | 5962

Syntax

```
drop-profile (profile-name | predefined-variable);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name drop-profile-map loss-
priority (any | low | medium-low | medium-high | high) protocol (any | non-tcp | tcp)]
```

Description

Within the drop-profile map, specify the name of the drop profile to use for random early detection (RED) for a specific packet-loss priority (PLP) level and protocol type. A drop profile maps a fill level (fullness of a queue) to a drop probability (probability that a packet will be dropped). When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.

You enable RED by applying a drop profile to a scheduler.

You configure drop profiles statically (at the `[edit class-of-service drop-profiles]` hierarchy level).

Options

profile-name—Name of the drop profile.

predefined-variable—One of the following Junos predefined variable that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached:

- `$junos-cos-scheduler-dropfile-low`—Name of the drop profile for PLP level `low` and protocol `any`, specified for a scheduler configured in a dynamic profile for subscriber access.
- `$junos-cos-scheduler-dropfile-medium-low`—Name of the drop profile for PLP level `medium-low` and protocol `any`, specified for a scheduler configured in a dynamic profile for subscriber access.
- `$junos-cos-scheduler-dropfile-medium-high`—Name of the drop profile for PLP level `medium-high` and protocol `any`, specified for a scheduler configured in a dynamic profile for subscriber access.
- `$junos-cos-scheduler-dropfile-high`—Name of the drop profile for PLP level `high` and protocol `any`, specified for a scheduler configured in a dynamic profile for subscriber access.

- `$junos-cos-scheduler-dropfile-lny`—Name of the drop profile for PLP level `any` and protocol `any`, specified for a scheduler configured in a dynamic profile for subscriber access.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

The `$junos-cos-scheduler-dropfile-low`, `$junos-cos-scheduler-dropfile-medium-low`, `$junos-cos-scheduler-dropfile-medium-high`, `$junos-cos-scheduler-dropfile-high`, and `$junos-cos-scheduler-dropfile-any` predefined variable introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Schedulers in a Dynamic Profile for Subscriber Access

scheduler (Dynamic Scheduler Maps) | **16344**

Determining Packet Drop Behavior by Configuring Drop Profile Maps for Schedulers

drop-profile-map (Dynamic Schedulers)

IN THIS SECTION

- Syntax | **5963**
- Hierarchy Level | **5963**
- Description | **5963**
- Required Privilege Level | **5963**
- Release Information | **5963**

Syntax

```
drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any | non-
tcp | tcp) drop-profile (profile-name | predefined-variable);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
```

Description

Define loss priority value for drop profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Schedulers in a Dynamic Profile for Subscriber Access

scheduler (Dynamic Scheduler Maps) | **16344**

drop-profile-map (Schedulers)

IN THIS SECTION

- Syntax | 5964
- Hierarchy Level | 5964
- Description | 5964
- Options | 5964
- Required Privilege Level | 5964
- Release Information | 5965

Syntax

```
drop-profile-map loss-priority (low | medium-high | high) protocol protocol drop-profile drop-profile-name;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Map a drop profile to a loss priority and protocol for weighted random early detection (WRED). When a packet arrives, WRED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the WRED algorithm determines whether to drop the arriving packet.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Drop Profile Maps

drop-profile-map (Schedulers)

IN THIS SECTION

- Syntax | 5965
- Hierarchy Level | 5965
- Description | 5966
- Required Privilege Level | 5966
- Release Information | 5966

Syntax

```
drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any | non-
tcp | tcp) drop-profile profile-name;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Map a drop profile to a loss priority and protocol for weighted random early detection (WRED). When a packet arrives, WRED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the WRED algorithm determines whether to drop the arriving packet.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Default Schedulers Overview*
- *Determining Packet Drop Behavior by Configuring Drop Profile Maps for Schedulers*
- *Example: Configuring Drop Profile Maps*
- *Example: Configuring CoS on EX Series Switches*
- *Defining CoS Schedulers and Scheduler Maps (CLI Procedure)*
- *Defining CoS Schedulers (J-Web Procedure)*
- *Understanding CoS Schedulers*

drop-profile (Schedulers)

IN THIS SECTION

- Syntax | 5967
- Hierarchy Level | 5967
- Description | 5967

- Options | 5967
- Required Privilege Level | 5967
- Release Information | 5967

Syntax

```
drop-profile profile-name;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name drop-profile-map loss-priority (any | low | medium-
low | medium-high | high) protocol (any | non-tcp | tcp)]
```

Description

Define drop profiles for RED. When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.

Options

profile-name—Name of the drop profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Determining Packet Drop Behavior by Configuring Drop Profile Maps for Schedulers

Managing Congestion Using RED Drop Profiles and Packet Loss Priorities

drop-profiles (Definition)

IN THIS SECTION

- Syntax | 5968
- Hierarchy Level | 5968
- Description | 5969
- Options | 5969
- Required Privilege Level | 5969
- Release Information | 5969

Syntax

```
drop-profiles {
    profile-name {
        fill-level percentage drop-probability percentage;
        interpolate {
            drop-probability [values];
            fill-level [values]
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Define drop profiles for random early detection (RED).

For a packet to be dropped, it must match the drop profile. When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the packet.

Options

profile-name—Name of the drop profile.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Defining Packet Drop Behavior by Configuring RED Drop Profiles

drop (Security GTP)

IN THIS SECTION

- Syntax | 5970
- Hierarchy Level | 5971
- Description | 5971
- Options | 5971

Required Privilege Level | 5974
Release Information | 5974

Syntax

```
drop {
    aa-create-pdp 0;
    aa-delete-pdp 0;
    bearer-resource 2;
    change-notification 2;
    config-transfer 2;
    context 2;
    create-bearer 2;
    create-data-forwarding 2;
    create-pdp (0 | 1 | all);
    create-session 2;
    create-tnl-forwarding 2;
    cs-paging 2;
    data-record (0 | 1 | all);
    delete-bearer 2;
    delete-command 2;
    delete-data-forwarding 2;
    delete-pdn 2;
    delete-pdp (0 | 1 | all);
    delete-session 2;
    detach 2;
    downlink-notification 2;
    echo (0 | 1 | 2 | all);
    error-indication (0 | 1 | all);
    failure-report (0 | 1 | all);
    fwd-access 2;
    fwd-relocation (1 | 2 | all);
    fwd-srns-context 1;
    g-pdu (0 | 1 | all);
    identification (0 | 1 | 2 | all);
    mbms-session-start (1 | 2 | all);
    mbms-session-stop (1 | 2 | all);
    mbms-session-update (1 | 2 | all);
```

```
    modify-bearer 2;
    modify-command 2;
    node-alive (0 | 1 | all);
    note-ms-present (0 | 1 | all);
    pdu-notification (0 | 1 | all);
    ran-info (1 | 2 | all);
    redirection (0 | 1 | all);
    release-access 2;
    relocation-cancel (1 | 2 | all);
    resume 2;
    send-route (0 | 1 | all);
    sgsn-context (0 | 1 | all);
    stop-paging 2;
    supported-extension 1;
    suspend 2;
    trace-session 2;
    update-bearer 2;
    update-pdn 2;
    update-pdp (0 | 1 | all);
    ver-not-supported (0 | 1 | 2 | all);
}
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

Use the parameters *(0 | 1 | 2 | all)* for the GTP release version number for a specific message type. The possible versions are 0, 1, 2, or all. You can configure the device to drop packets that do not meet your specified minimum or maximum message lengths. A GTP message type includes one or many messages. When you drop a message type, you automatically drop all messages of the specified type.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name]` hierarchy level is replaced by `[edit security gtp profile profile-name]`.

Options

- **aa-create-pdp** —Represents Create AA PDP Context Request and Create AA PDP Context Response messages.

- **aa-delete-pdp** —Represents Delete AA PDP Context Request and Delete AA PDP Context Response messages.
- **bearer-resource**—Represents Bearer Resource Command and Bearer Resource Failure messages.
- **change-notification**—Represents Change Notification Request and Change Notification Response messages.
- **context**—Represents Context Request and Context Response messages.
- **config-transfer**—Represents Configuration Transfer Tunnel messages.
- **create-bearer**—Represents Create Bearer Request and Create Bearer Response messages.
- **create-data-forwarding**—Represents Create Indirect Data Forwarding Request and Create Indirect Data Forwarding Response messages.
- **create-tnl-forwarding**—Represents Create Forwarding Tunnel Request and Create Forwarding Tunnel Response messages.
- **create-pdp**—Represents Create PDP Context Request and Create PDP Context Response messages.
- **create-session**—Represents Create Session Request and Create Session Response messages.
- **cs-paging**—Represents CS Paging Indication messages.
- **data-record**—Represents Data Record Request and Data Record Response messages.
- **delete-bearer**—Represents Delete Bearer Request and Delete Bearer Response messages.
- **delete-command**—Represents Delete Bearer Command and Delete Bearer Failure messages.
- **delete-data-forwarding**—Represents Delete Indirect Data Forwarding Request and Delete Indirect Data Forwarding Response messages.
- **delete-pdn**—Represents Delete PDN Connection Set Request and Delete PDN Connection Set Response messages.
- **delete-pdp**—Represents Delete PDP Context Request and Delete PDP Context Response messages.
- **delete-session**—Represents Delete Session Request and Delete Session Response messages.
- **detach**—Represents Detach Notification and Detach Acknowledgement messages.
- **downlink-notification**—Represents Downlink Data Notification, Downlink Data Acknowledgement, and Downlink Data Notification Failure Indication messages.
- **echo**—Represents Echo Request and Echo Response messages.
- **error-indication**—Represents Error Indication messages.

- **failure-report**—Represents Failure Report Request and Failure Report Response messages.
- **fwd-access**—Represents Forward Access Context Notification and Forward Access Context Acknowledgment messages.
- **fwd-relocation**—Represents Forward Relocation Request, Forward Relocation Response, Forward Relocation Complete, and Forward Relocation Complete Acknowledge messages.
- **fwd-srns-context**—Represents Forward SRNS Context Request and Forward SRNS Context Response messages.
- **g-pdu**—Represents G-PDU and T-PDU messages.
- **identification**—Represents Identification Request and Identification Response messages.
- **mbms-sess-start**—Represents MBMS Session Start Request and MBMS Session Start Response messages.
- **mbms-sess-stop**—Represents MBMS Session Stop Request and MBMS Session Stop Response messages.
- **mbms-sess-update**—Represents MBMS Session Update Request and MBMS Session Update Response messages.
- **modify-bearer**—Represents Modify Bearer Request and Modify Bearer Response messages.
- **modify-command**—Represents Modify Bearer Command and Modify Bearer Failure messages.
- **node-alive**—Represents Node Alive Request and Node Alive Response messages.
- **note-ms-present**—Represents Note MS GPRS Present Request and Note MS GPRS Present Response messages.
- **pdu-notification**—Represents PDU Notification request and PDU Notification response messages.
- **ran-info**—Represents Ran Info Relay messages.
- **redirection**—Represents Redirection Request and Redirection Response messages.
- **relocation-cancel**—Represents Relocations Cancel Request and Relocation Cancel Response messages.
- **resume**—Represents Resume Notification and Resume Acknowledgement messages.
- **send-route**—Represents Send Route Info Request and Send Route Info Response messages.
- **sgsn-context**—Represents SGSN Context Request and SGSN Context Response messages.
- **stop-paging**—Represents Stop Paging Indication messages.

- **supported-extension**—Represents Supported Extension Headers Notification messages.
- **suspend**—Represents Suspend Notification and Suspend Acknowledgement messages.
- **trace-session**—Represents Trace Session Activation and Trace Session Deactivation messages.
- **update-bearer**—Represents Update Bearer Request and Update Bearer Response messages.
- **update-pdn**—Represents Update PDN Set Connection Request and PDN Set Connection Response messages.
- **update-pdp**—Represents Update PDP Request and Update PDP Response messages.
- **ver-not-supported**—Represents Version Not Supported messages.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. New GTP message types added in Junos OS Release 11.4. Support .

RELATED DOCUMENTATION

Understanding GTP Message Filtering

Example: Filtering GTP Message Types

Understanding GTP APN Filtering

drop (Security SCTP)

IN THIS SECTION

Syntax | 5975

Hierarchy Level | 5976

- Description | 5976
- Options | 5976
- Required Privilege Level | 5977
- Release Information | 5977

Syntax

```
drop {
    m3ua-service {
        isup;
       sccp;
        tup;
   }
   payload-protocol {
        id;
        all;
        asap;
        bicc;
       ddp-segment;
       ddp-stream;
       diameter-dtls;
       diameter-sctp;
        dua;
        enrp;
        h248;
        h323;
         iua;
         m2pa;
          m2ua;
         m3ua;
        qipc;
       reserved;
       s1ap;
         simco;
         sua;
         tali;
        v5ua;
       x2ap;
```

```
    }
}
```

Hierarchy Level

```
[edit security sctp profile profile-name]
```

Description

Configure the drop SCTP message types to display information about the configuration of the current Stream Control Transmission Protocol (SCTP) inspection.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs sctp profile profile-name]` hierarchy level is replaced by `[edit security sctp profile profile-name]`.

Options

- **m3ua-services**—M3UA data service indicator. The following values are supported:
  - **isup**—ISDN Upper Part.
  - **sccp**—Signaling Connection Control Part.
  - **tup**—Telephone User Part.
- **payload-protocol**—SCTP payload protocol identifier. The following values are supported:
  - **id**—Specify payload protocol ID.
  - **all**—All SCTP payload protocol identifiers (id:0~63).
  - **asap**—Aggregate Server Access Protocol.
  - **bicc**—Bearer Independent Call Control.
  - **ddp-segement**—Direct Data Placement Segment chunk.
  - **ddp-stream**—Direct Data Placement Stream session control.
  - **diameter-dtls**—Diameter in a DTLS/SCTP DATA chunk.
  - **diameter-sctp**—Diameter in a SCTP DATA chunk.
  - **dua**—DPNSS/DASS 2 extensions to IUA Protocol.

- **enrp**—Endpoint Handlespace Redundancy Protocol.
- **h248**—H.248 Protocol.
- **h323**—H.323 Protocol.
- **iua**—ISDN User Adaptation Layer.
- **m2pa**—MTP2 User Peer-to-Peer Adaption Layer.
- **m2ua**—MTP2 User Adaption Layer.
- **m3ua**—MTP3 User Adaption Layer.
- **qipc**—Q.IPC.
- **reserved**—Reserved by SCTP.
- **s1ap**—S1 Application Protocol (S1AP).
- **simco**—Simple Middlebox Configuration.
- **sua**—SCCP User Adaption Layer.
- **tali**—Transport Adapter Layer Interface.
- **v5ua**—v5.2 User Adaption Layer.
- **x2ap**—X2 Application Protocol (X2AP).

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2. Support for the `payload-protocol` statement was modified n Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

> *Understanding Stream Control Transmission Protocol*
>
> *Understanding SCTP Multichunk Inspection*

drop-threshold

IN THIS SECTION

- Syntax | 5978
- Hierarchy Level | 5978
- Description | 5978
- Default | 5979
- Options | 5979
- Required Privilege Level | 5979
- Release Information | 5979

Syntax

```
drop-threshold number;
```

Hierarchy Level

```
[edit vlans vlan-name dot1q-tunneling layer2-protocol-tunneling (all | protocol-name vlan-name
dot1q-tunneling layer2-protocol-tunneling (all | protocol-name)]
```

```
[edit protocols layer2-control bpdu-block interface interface-name
```

]

Description

Specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the switch begins dropping the Layer 2 PDUs. The drop threshold value must be less than or equal to the shutdown threshold value.

L2PT processing is done by the CPU, and L2PT traffic to the CPU is rate-limited to a maximum of 1000 pps. If traffic is received at a rate faster than this limit, the rate limit causes the traffic to be dropped before it hits the threshold and the dropped packets are not reported in L2PT statistics. This can also

occur if you configure a drop threshold that is less than 1000 pps but traffic is received at a faster rate. For example, if you configure a drop threshold of 900 pps and the VLAN receives traffic at rate of 1100 pps, L2PT statistics will show that 100 packets were dropped. The 100 packets dropped because of the rate limit will not be reported. Similarly, if you do not configure a drop threshold and the VLAN receives traffic at rate of 1100 pps, the 100 packets dropped because of the rate limit are not reported.

> **NOTE**: If the drop threshold value is greater than the shutdown threshold value and you try to commit the configuration, the commit operation fails.

You can specify a drop threshold value without specifying a shutdown threshold value.

Default

No drop threshold is specified.

Options

number—Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the switch begins dropping the Layer 2 PDUs.

- **Range:** 1 through 1000

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

shutdown-threshold | **27937**

drop-threshold (Security GPRS)

IN THIS SECTION

- Syntax | 5980
- Hierarchy Level | 5980
- Description | 5980
- Options | 5981
- Required Privilege Level | 5981
- Release Information | 5981

Syntax

```
drop-threshold {
    forward number;
    reverse number;
}
```

Hierarchy Level

```
[edit security gtp profile profile-name path-rate-limit message-type (create-req | delete-req |
echo-req | other)]
```

Description

Configure the drop-threshold parameter to drop traffic when the number of packets per second or per minute exceeds the configured limit.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name path-rate-limit message-type (create-req | delete-req | echo-req | other)]` hierarchy-level is replaced by `[edit security gtp profile profile-name path-rate-limit message-type (create-req | delete-req | echo-req | other)]`.

Options

number—Limit messages in forward or reverse direction.

- **Range:** 1 through 10,000.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Example: Limiting the Message Rate and Path Rate for GTP Control Messages

drop-timeout

IN THIS SECTION

- Syntax | 5982
- Hierarchy Level | 5982
- Description | 5982
- Options | 5982
- Required Privilege Level | 5982
- Release Information | 5982

Syntax

```
drop-timeout milliseconds;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options],
[edit interfaces (ls-fpc/pic/port | ml-fpc/pic/port) unit logical-unit-number],
[edit logical-systems logical-system-name interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-
options],
[edit logical-systems logical-system-name interfaces (ls-fpc/pic/port| ml-fpc/pic/port) unit
logical-unit-number]
```

Description

For multilink and link services interfaces only, configure the drop timeout period, in milliseconds.

Options

milliseconds—Drop timeout period.

- **Range:** 1 through 2000 milliseconds
- **Default:** 500 ms for bundles greater than or equal to the T1 bandwidth value, and 1500 ms for other bundles. Any CLI-configured value overrides these defaults. Setting a value of 0 reverts to the default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Link and Multilink Services Logical Interfaces

drop-timeout (Forwarding Class)

IN THIS SECTION

- Syntax | 5983
- Hierarchy Level | 5983
- Description | 5983
- Default | 5983
- Options | 5984
- Required Privilege Level | 5984
- Release Information | 5984

Syntax

```
drop-timeout milliseconds;
```

Hierarchy Level

```
[edit class-of-service fragmentation-map map-name forwarding-class class-name]
```

Description

Disable or set the resequencing timeout interval for each forwarding class of a multiclass MLPPP.

Default

If you do not include this statement, the default sequencing timeouts for T1 speeds (500 ms) or lower (1500 ms) apply.

Options

milliseconds—Time to wait for fragments. A value of 0 disables the resequencing logic for that forwarding class.

- **Range:** 0 through 500 milliseconds for bundles with bandwidths or T1 speeds or higher or 1500 ms for bundles with bandwidths of less than T1 speeds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring Drop Timeout Interval for Fragmentation by Forwarding Class

drop-unknown-mac-ip

IN THIS SECTION

- Syntax | 5985
- Hierarchy Level | 5985
- Description | 5985
- Required Privilege Level | 5985
- Release Information | 5985

Syntax

```
drop-unknown-mac-ip;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols evpn],
[edit routing-instances name protocols evpn],
```

```
[edit logical-systems name routing-instances name bridge-domains name bridge-options],
[edit routing-instances name bridge-domains name bridge-options],
```

```
[edit routing-instances name vlans name switch-options],
[edit vlans name switch-options]
```

Description

Disable dynamic learning of MAC-IP entries. The ingress PE router drops all ARP requests and NS messages directed to unknown IP addresses.

Required Privilege Level

[none specified]

Release Information

Statement introduced in Junos OS Release 23.4R1.

ds-lite

IN THIS SECTION

- Syntax | 5986
- Hierarchy Level | 5986
- Description | 5986
- Options | 5987
- Required Privilege Level | 5987
- Release Information | 5987

Syntax

```
    ds-lite ds-lite-softwire-concentrator {
        auto-update-mtu;
        flow-limit flow-limit | session-limit-per-prefix session-limit-per-prefix;
        mtu-v6 bytes;
        softwire-address softwire-address;
    }
}
```

Hierarchy Level

```
[edit services softwire softwire-concentrator]
[edit services softwires softwire-types]
```

Description

Configure settings for a DS-Lite concentrator used to process IPv4 packets encapsulated in IPv6.

The `ds-lite` statement is supported on MX Series routers with MS-DPCs and on M Series routers with MS-100, MS-400, and MS-500 line Multiservices PICs. Starting in Junos OS release 17.4R1, DS-Lite is supported on MX Series routers with MS-MPCs and MS-MICs.

Options

bytes—Maximum transmission unit (MTU), in bytes, for encapsulating IPv4 packets into IPv6. If the final length is greater than the configured value, the IPv6 packet is fragmented. This option is supported on MX Series routers equipped with MS-DPCs. Starting in Junos OS release 18.1R1, this option is also supported on MX Series routers with MS-MPCs or MS-MICs.

ds-lite-softwire-concentrator—Name applied to a DS-Lite softwire concentrator.

auto-update-mtu—This option is not currently supported.

copy-dscp—Copy DSCP information to IPv4 headers during decapsulation.

flow-limit—Maximum number of IPv4 flows per softwire.

- **Range:** 0 through 16384 flows
- **Range:** 0 through 9192 bytes

session-limit-per-prefix—Maximum number of sessions per B4 subnet prefix. This option is supported on MX Series routers equipped with MS-DPCs. Starting in Junos OS Release 18.2R1, this option is also supported on MS-MPCs and MS-MICs.

- **Range:** 0 through 16384 sessions

softwire-address—Address of the DS-Lite softwire concentrator.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

auto-update-mtu option introduced in Junos OS Release 10.4.

copy-dscp option introduced in Junos OS Release 11.2.

mtu-v6 option introduced in Junos OS Release 10.4.

softwire-address option introduced in Junos OS Release 10.4.

Support for DS-Lite at the `[edit services softwires softwire-types]` added in Junos OS release 20.2R1 for Next Gen Services on MX240, MX480 and MX960 routers.

RELATED DOCUMENTATION

Configuring a DS-Lite Softwire Concentrator

dscp (Class of Service)

IN THIS SECTION

- Syntax (Classifiers) | 5988
- Syntax (Code-Point Alias) | 5989
- Syntax (Multidestination Classifier) | 5989
- Syntax (Interface Classifier Association) | 5989
- Syntax (Rewrite Rule) | 5989
- Hierarchy Level (Classifiers) | 5989
- Hierarchy Level (Code-Point Aliases) | 5989
- Hierarchy Level (Multidestination Classifier) | 5990
- Hierarchy Level (Interface Classifier Association) | 5990
- Hierarchy Level (Rewrite Rule) | 5990
- Description | 5990
- Options | 5990
- Required Privilege Level | 5990
- Release Information | 5990

Syntax (Classifiers)

```
dscp classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [ aliases ] [ bit-patterns ];
        }
    }
}
```

Syntax (Code-Point Alias)

```
dscp alias-name bit-pattern;
```

Syntax (Multidestination Classifier)

```
dscp classifier-name;
```

Syntax (Interface Classifier Association)

```
dscp (classifier-name | default);
```

Syntax (Rewrite Rule)

```
dscp rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-point [ aliases ] [ bit-patterns ];
        }
    }
}
```

Hierarchy Level (Classifiers)

```
[edit class-of-service classifiers],
```

Hierarchy Level (Code-Point Aliases)

```
[edit class-of-service code-point-aliases],
```

Hierarchy Level (Multidestination Classifier)

```
[edit class-of-service multi-destination classifiers],
```

Hierarchy Level (Interface Classifier Association)

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers],
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules],
```

Hierarchy Level (Rewrite Rule)

```
[edit class-of-service rewrite-rules]
```

Description

Define the Differentiated Services code point (DSCP) mapping that is applied to the packets.

Options

classifier-name—Name of the classifier.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Unicast Classifiers

Example: Configuring Classifiers

Defining CoS Code-Point Aliases

Defining CoS Rewrite Rules

Assigning CoS Components to Interfaces

Understanding CoS Classifiers

Understanding CoS Classifiers

Understanding CoS Rewrite Rules

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

dscp-code-point (CoS Host Outbound Traffic)

IN THIS SECTION

- Syntax | 5991
- Hierarchy Level | 5991
- Description | 5992
- Options | 5992
- Required Privilege Level | 5992
- Release Information | 5992

Syntax

```
dscp-code-point value;
```

Hierarchy Level

```
[edit class-of-service host-outbound-traffic]
```

Description

Specify the value of the DSCP bits in the type of service (ToS) field of host outbound traffic (packets generated by the local Routing Engine) as they are placed in the default or specified output queue on all egress interfaces. This statement does not affect transit traffic or incoming traffic.

If you use the `ping` operational mode command with the `tos` *`type-of-servce`* option, the value specified in this configuration statement overrides the DSCP value you specify in the `ping` command.

> **NOTE**: With this configuration, QFX Series switches rewrite DSCP code points for *all* packets that are host-generated. EX Series switches and MX Series routers rewrite DSCP code points *only if* the packet L3 TOS value is not already set or is marked with the default code point.
>
> Any DSCP rewrite rules configured on a 10-Gigabit Ethernet LAN/WAN PIC with SFP+ overwrite this DSCP value.

For egress interfaces hosted on MX Series routers, M120 routers, or Enhanced III FPCs in M320 routers, both Routing Engine sourced traffic and distributed protocol handler traffic are affected. For all other egress interfaces, only Routing Engine sourced traffic is affected.

Options

code-point	Six-bit DSCP code point value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for distributed protocol handler traffic introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic

Default DSCP and DSCP IPv6 Classifiers

dscp-code-points (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 5993
- Junos OS Hierarchy Level | 5993
- Junos OS Evolved Hierarchy Level | 5993
- Description | 5994
- Options | 5994
- Required Privilege Level | 5995
- Release Information | 5995

Syntax

```
dscp-code-points dscp-bits;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the value of the Differentiated Services (DiffServ) field within the IP header of the test frames. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern. If you do not specify this value, 0 is used in the DSCP fields in the IP header.

Options

dscp-bits—A valid 6-bit pattern; for example, `001111`, or one of the following configured DSCP aliases:

- `af11`—Default: 001010
- `af12`—Default: 001100
- `af13`—Default: 001110
- `af21`—Default: 010010
- `af22`—Default: 010100
- `af23` —Default: 010110
- `af31` —Default: 011010
- `af32` —Default: 011100
- `af33` —Default: 011110
- `af41` —Default: 100010
- `af42` —Default:100100
- `af43` —Default:100110
- `be`—Default: 000000
- `cs1`—Default: 001000
- `cs2`—Default: 010000
- `cs3`—Default: 011000
- `cs4`—Default: 100000
- `cs5`—Default: 101000
- `cs6`—Default: 110000
- `cs7`—Default: 111000

- `ef`—Default: 101110
- `nc1`—Default: 110000
- `nc2`—Default: 111000

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

dscp-code-points (RPM)

IN THIS SECTION

- Syntax | **5996**
- Junos OS Hierarchy Levels | **5996**
- Junos OS Evolved Hierarchy Level | **5996**
- Description | **5996**
- Options | **5996**
- Required Privilege Level | **5997**
- Release Information | **5997**

Syntax

```
dscp-code-points dscp-bits;
```

Junos OS Hierarchy Levels

```
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the value of the Differentiated Services (DiffServ) field within the IP header. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern.

Options

dscp-bits—A valid 6-bit pattern; for example, `001111`, or one of the following configured DSCP aliases:

- `af11`—Default: 001010
- `af12`—Default: 001100
- `af13`—Default: 001110
- `af21`—Default: 010010
- `af22`—Default: 010100
- `af23` —Default: 010110
- `af31` —Default: 011010
- `af32` —Default: 011100
- `af33` —Default: 011110

- `af41` —Default: 100010
- `af42` —Default:100100
- `af43` —Default:100110
- `be`—Default: 000000
- `cs1`—Default: 001000
- `cs2`—Default: 010000
- `cs3`—Default: 011000
- `cs4`—Default: 100000
- `cs5`—Default: 101000
- `cs6`—Default: 110000
- `cs7`—Default: 111000
- `ef`—Default: 101110
- `nc1`—Default: 110000
- `nc2`—Default: 111000

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release for PTX Series Packet Transport routers.

Support at the [`edit services rpm twamp client control-connection` *`control-client-name`*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

Understanding Two-Way Active Measurement Protocol on Routers

dscp (CoS Classifiers)

IN THIS SECTION

- Syntax | 5998
- Hierarchy Level | 5998
- Description | 5999
- Options | 5999
- Required Privilege Level | 5999
- Release Information | 5999

Syntax

```
dscp classifier-name {
    import (classifier-name | default);
    forwarding-class  class-name {
        loss-priority level ] {
            code-points [ aliases ] [ bit-patterns;
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service classifiers]
```

Description

Define the diffserv code point (DSCP) mapping to apply to incoming packets.

Options

classifier-name—Name of the classifier.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic

Applying DSCP Classifiers to MPLS Traffic

dscp (CoS Interfaces)

IN THIS SECTION

- Syntax | 6000
- Hierarchy Level | 6000
- Description | 6000
- Options | 6000
- Required Privilege Level | 6000
- Release Information | 6000

Syntax

```
dscp (classifier-name | default);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name classifiers]
```

Description

For ACX Series Universal Metro routers, map the DSCP field of the incoming packet to the forwarding class and packet loss priority based on the specified DSCP classifier.

Options

classifier-name—Name of the previously defined DSCP behavior aggregate classifier.

default—Default DSCP behavior aggregate classifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

dscp (Dynamic Classifiers)

IN THIS SECTION

- Syntax | 6001
- Hierarchy Level | 6001

- Description | 6001
- Options | 6001
- Required Privilege Level | 6001
- Release Information | 6001

Syntax

```
dscp (classifier-name | default);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-number classifiers]
```

Description

For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) classifier to a subscriber interface in a dynamic profile.

Options

classifier-name—Name of a `classifier` mapping configured at the `[edit class-of-service classifiers dscp]` hierarchy level.

`default`—The default mapping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying a Classifier to a Subscriber Interface in a Dynamic Profile

classifiers (Definition)

dscp (Dynamic Rewrite Rules)

IN THIS SECTION

- Syntax | 6002
- Hierarchy Level | 6002
- Description | 6002
- Options | 6003
- Required Privilege Level | 6003
- Release Information | 6003

Syntax

```
dscp (rewrite-name | default);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) rewrite rule to a subscriber interface in a dynamic profile.

Options

rewrite-name—Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules dscp]` hierarchy level.

`default`—The default mapping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile

rewrite-rules

dscp (flow-server)

IN THIS SECTION

- Syntax | 6004
- Hierarchy Level | 6004
- Description | 6004
- Options | 6004
- Required Privilege Level | 6004
- Release Information | 6004

Syntax

```
dscp dscp-value;
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet | inet6) output flow-
server hostname]
```

Description

Specify the Differentiated Services Code Point (DSCP) mapping that is applied to exported packets for inline active flow monitoring. This allows different levels of service to be assigned to sampled traffic.

Options

dscp *dscp-value*

Can be a value between 0 and 63 (the default is 0). When the same `flow-server` is configured under both the `inet` and `inet6` families in a sampling instance, use the same `dscp` value for both `flow-server` appearances.

The *dscp-value* is overwritten by the CoS DSCP value if you configure `dscp` under the `[edit class-of-service]` hierarchy.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F4.

Statement introduced in Junos OS Release 16.1 for the MX Series.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring on PTX Series Routers

Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers

Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches

dscp (Input Congestion Notification)

IN THIS SECTION

- Syntax | 6005
- Hierarchy Level | 6005
- Description | 6005
- Required Privilege Level | 6006
- Release Information | 6006

Syntax

```
dscp {
    code-point [code-point-bits] {
        pfc {
            mru mru-value;
        }
        xon (Input Congestion Notification)
    }
}
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile name input]
```

Description

Configure a Differentiated Services code point (DSCP) value and apply priority-based flow control (PFC) to packets with that code point.

When this statement is configured, DSCP-based PFC can be invoked for untagged traffic by matching specified 6-bit DSCP values in the Layer 3 IP header of incoming packets instead of an IEEE 802.1p priority in the Layer 2 VLAN header. Additional configuration parameters associate configured DSCP values with a PFC priority to use in the Layer 2 pause frames sent to peers when the link becomes congested.

DSCP-based PFC can be used to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2).

To enable DSCP-based PFC:

- Use this statement to define a congestion notification profile to enable PFC on traffic specified by a DSCP value.
- Use the `[edit class-of-service forwarding-classes class class-name]` `pfc-priority` statement to map a lossless forwarding class to a PFC priority value to use in the PFC pause frames.
- Use the `[edit class-of-service classifiers]` `dscp` statement to set up a DSCP classifier for the desired DSCP value and forwarding class mapped to a PFC priority above.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Understanding PFC Using DSCP at Layer 3 for Untagged Traffic

Configuring DSCP-based PFC for Layer 3 Untagged Traffic

Configuring CoS PFC (Congestion Notification Profiles)

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

dscp-ipv6 (Class of Service)

IN THIS SECTION

- Syntax (Classifier) | 6007
- Syntax (Code-Point Alias) | 6007
- Syntax (Interface Classifier Association) | 6008
- Syntax (Rewrite Rule) | 6008
- Hierarchy (Classifiers) | 6008
- Hierarchy (Code-Point Aliases) | 6008
- Hierarchy (Interface Classifier Association) | 6008
- Hierarchy (Rewrite Rules) | 6008
- Description | 6009
- Options | 6009
- Required Privilege Level | 6009
- Release Information | 6009

Syntax (Classifier)

```
dscp-ipv6 classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [ aliases ] [ bit-patterns ];
        }
    }
}
```

Syntax (Code-Point Alias)

```
dscp-ipv6 alias-name bit-pattern;
```

Syntax (Interface Classifier Association)

```
dscp-ipv6 (classifier-name | default);
```

Syntax (Rewrite Rule)

```
dscp-ipv6 rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-point [ aliases ] [ bit-patterns ];
        }
    }
}
```

Hierarchy (Classifiers)

```
[edit class-of-service classifiers],
```

Hierarchy (Code-Point Aliases)

```
[edit class-of-service code-point-aliases],
```

Hierarchy (Interface Classifier Association)

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers],
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules],
```

Hierarchy (Rewrite Rules)

```
[edit class-of-service rewrite-rules]
```

Description

Define the Differentiated Services code point (DSCP) IPv6 mapping that is applied to the packets.

> **NOTE**: On switches that use different classifiers for unicast and multidestination (multicast, broadcast, and destination lookup fail) traffic, there is no DSCP IPv6 classifier for multidestination (multicast, broadcast, and destination lookup fail) traffic. Multidestination IPv6 traffic uses the multidestination DSCP classifier.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Defining CoS Code-Point Aliases

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

dscp-ipv6 (Class of Service)

IN THIS SECTION

- Syntax (Classifier) | **6010**
- Syntax (Code-Point Alias) | **6010**
- Syntax (Interface Classifier Association) | **6010**

- Syntax (Rewrite Rule) | 6011
- Hierarchy (Classifiers) | 6011
- Hierarchy (Code-Point Aliases) | 6011
- Hierarchy (Interface Classifier Association) | 6011
- Hierarchy (Rewrite Rules) | 6011
- Description | 6011
- Options | 6012
- Required Privilege Level | 6012
- Release Information | 6012

Syntax (Classifier)

```
dscp-ipv6 classifier-name {
    import (classifier-name | default) protocol mpls;
    forwarding-class class-name {
        loss-priority level {
            code-points [ aliases ] [ bit-patterns ];
        }
    }
}
```

Syntax (Code-Point Alias)

```
dscp-ipv6 alias-name bit-pattern;
```

Syntax (Interface Classifier Association)

```
dscp-ipv6 (classifier-name | default);
```

Syntax (Rewrite Rule)

```
dscp-ipv6 rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-point [ aliases ] [ bit-patterns ];
        }
    }
}
```

Hierarchy (Classifiers)

```
[edit class-of-service classifiers],
```

Hierarchy (Code-Point Aliases)

```
[edit class-of-service code-point-aliases],
```

Hierarchy (Interface Classifier Association)

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers],
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules],
```

Hierarchy (Rewrite Rules)

```
[edit class-of-service rewrite-rules]
```

Description

Define the Differentiated Services code point (DSCP) IPv6 mapping that is applied to the packets.

> **NOTE**: On devices that use different classifiers for unicast and multidestination (multicast, broadcast, and destination lookup fail) traffic, there is no DSCP IPv6 classifier for multidestination (multicast, broadcast, and destination lookup fail) traffic. Multidestination IPv6 traffic uses the multidestination DSCP classifier.

Logical interfaces do not support multiple `dscp-ipv6` rewrite rules for the same protocol.

Options

The remaining statements are explained separately. See CLI Explorer.

`protocol mpls` (Optional for ingress MPLS tunnel nodes) On devices that support this option, rewrite the MPLS EXP bits in the MPLS header independently of the IPv6 DSCP value for IPv6 packets entering an MPLS tunnel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for `protocol mpls` option introduced in Junos OS Release 10.4R2.

RELATED DOCUMENTATION

Configuring Rewrite Rules

protocol (Rewrite Rules) | **13457**

Setting IPv6 DSCP and MPLS EXP Values Independently

Configuring DSCP Values for IPv6 Packets Entering the MPLS Tunnel

Applying Rewrite Rules to Output Logical Interfaces

rewrite-rules (Definition) | **15746**

Defining CoS Code-Point Aliases

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

Example: Configuring CoS on EX Series Switches
Defining CoS Code-Point Aliases (CLI Procedure)
Defining CoS Code-Point Aliases (J-Web Procedure)
Defining CoS Classifiers (CLI Procedure)
Defining CoS Classifiers (J-Web Procedure)
Defining CoS Rewrite Rules (CLI Procedure)
Defining CoS Rewrite Rules (J-Web Procedure)
Assigning CoS Components to Interfaces (CLI Procedure)
Assigning CoS Components to Interfaces (J-Web Procedure)
Understanding CoS Classifiers

dscp-ipv6 (Dynamic Classifiers)

IN THIS SECTION

- Syntax | 6013
- Hierarchy Level | 6013
- Description | 6014
- Options | 6014
- Required Privilege Level | 6014
- Release Information | 6014

Syntax

```
dscp-ipv6 (classifier-name | default);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-
number classifiers]
```

Description

For IPv6 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) classifier to a subscriber interface in a dynamic profile.

Options

classifier-name—Name of a `classifier` mapping configured at the `[edit class-of-service classifiers dscp-ipv6]` hierarchy level.

`default`—The default mapping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 10.1.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying a Classifier to a Subscriber Interface in a Dynamic Profile

classifiers (Definition)

dscp-ipv6 (Dynamic Rewrite Rules)

IN THIS SECTION

- Syntax | 6015
- Hierarchy Level | 6015
- Description | 6015
- Options | 6015

Required Privilege Level | 6015
Release Information | 6015

Syntax

```
dscp-ipv6 (rewrite-name | default);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

For IPv6 traffic, apply a DSCP rewrite rule to a subscriber interface in a dynamic profile.

Options

rewrite-name—Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules dscp-ipv6]` hierarchy level.

`default`—The default mapping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release10.1.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

rewrite-rules

dscp-mapping-classifier

IN THIS SECTION

- Syntax | 6016
- Hierarchy Level | 6016
- Description | 6016
- Required Privilege Level | 6017
- Release Information | 6017

Syntax

```
dscp-mapping-classifier ipv4-classifier-name
```

```
dscp-ipv6-mapping-classifier ipv6-classifier-name
```

Hierarchy Level

```
[edit forwarding-options family (inet | inet6)]
```

Description

Enable the BGP FlowSpec filter on a device. Configure BGP Flow Specification (FlowSpec) DSCP action to forward packets using the forwarding class and loss priority information across the network effectively.

> **NOTE**: On Junos OS Evolved version 22.4 and earlier, we don't support the DSCP marking action as part of BGP Flow Specification (FlowSpec) route installation for PTX platforms. DSCP marking action for BGP FlowSpec filter configuration is blocked from the Control Plane. However, on Junos OS Evolved versions later than 22.4, DSCP marking action as part of BGP FlowSpec route installation is supported.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1 for PTX series and MX series devices.

dscp (Multifield Classifier)

IN THIS SECTION

- Syntax | 6017
- Hierarchy Level | 6018
- Description | 6018
- Options | 6018
- Required Privilege Level | 6018
- Release Information | 6018

Syntax

```
dscp [0 | value];
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name then]
```

Description

Set the IPv4 Differentiated Services code point (DSCP) bit. You can specify a numerical value from `0` through `63`. To specify the value in hexadecimal form, include `0x` as a prefix. To specify the value in binary form, include `b` as a prefix.

The default DSCP value is `be` (best effort), or `0`.

For M320 and T Series routers, set the DSCP field of incoming or outgoing packets to `000000`. On the same packets, you can use a behavior aggregate (BA) classifier and a rewrite rule to rewrite the MPLS EXP field.

For MX Series routers with MPCs and EX Series switches, the DSCP field can be set from a numeric range.

For MX Series routers and EX Series switches, if you configure a firewall filter with a `dscp` action or `traffic-class` action on a DPC, the commit fails.

If you apply a firewall filter with a `dscp` action on an AE member interface, the commit fails. However, you can apply a firewall filter with a `dscp` action on a full AE interface.

Options

value—For MX Series routers with MPCs, specify the field of incoming or outgoing packets in the range from `0` through `63`.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring and Applying Tricolor Marking Policers

Firewall Filter Nonterminating Actions

dscp-rewrite

IN THIS SECTION

- Syntax | 6019
- Hierarchy Level | 6019
- Description | 6019
- Options | 6020
- Required Privilege Level | 6020
- Release Information | 6020

Syntax

```
dscp-rewrite  {
    code-point  6 bit patterns;
}
```

Hierarchy Level

```
[edit security alg sip]
[edit security alg h323]
[edit security alg mgcp]
[edit security alg sccp]
```

Description

Specify a rewrite-rule for the traffic that passes through a voice over IP Application Layer Gateway (VoIP ALG).

Options

`6-bit patterns`—Value of the code point in binary form.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding VoIP DSCP Rewrite Rules

dscp (Rewrite Rules on Interfaces)

IN THIS SECTION

- Syntax | 6020
- Hierarchy Level | 6021
- Description | 6021
- Options | 6021
- Required Privilege Level | 6021
- Release Information | 6021

Syntax

```
dscp (rewrite-name | default) protocol (inet-both | inet-outer | mpls);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name rewrite-rules],
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) rewrite rule to an interface. Supported interfaces for rewrite rules depend on the platform and software release.

Logical interfaces do not support multiple `dscp` rewrite rules for the same protocol.

Options

rewrite-name	Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules dscp]` hierarchy level.
default	The default mapping.
protocol inet-both	Rewrite the DSCP CoS value to both the inner and outer header for Unicast/Multicast IPv4 traffic. The first six bits of the CoS value are rewritten and the final two bits are taken from the incoming CoS value.
protocol inet-outer	Rewrite the DSCP CoS value to the outer header for Unicast/Multicast IPv4 traffic. The first six bits of the CoS value are rewritten and the final two bits are taken from the incoming CoS value.
protocol mpls	(Optional for ingress MPLS tunnel nodes) Rewrite the MPLS EXP bits in the MPLS header independently of the IPv4 DSCP value for IPv4 packets entering an MPLS tunnel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Rewrite Rules

Applying Rewrite Rules to Output Logical Interfaces

protocol (Rewrite Rules) | **13457**

Rewriting MPLS and IPv4 Packet Headers

rewrite-rules (Definition) | **15746**

dscp (Services CoS)

IN THIS SECTION

- Syntax | **6022**
- Hierarchy Level | **6022**
- Description | **6023**
- Options | **6023**
- Required Privilege Level | **6023**
- Release Information | **6023**

Syntax

```
dscp (alias | bits);
```

Hierarchy Level

```
[edit services cos application-profile profile-name (ftp | sip) (data | video | voice)],
[edit services cos rule rule-name term term-name then],
[edit services cos rule rule-name term term-name then reverse]
```

Description

Define the Differentiated Services code point (DSCP) mapping that is applied to the packets. Change the DSCP (or TOS) on the packet to the specified value. Any conformant bit string can be specified, but only the default alias can be used.

Options

alias—Name assigned to a set of CoS markers.

bits—Mapping value in the packet header.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring Actions in CoS Rules

Configuring CoS Rules on Services PICs

dscp (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 6024
- Hierarchy Level | 6024
- Description | 6024
- Options | 6024
- Required Privilege Level | 6024

Release Information | 6024

Syntax

```
dscp value;
```

Hierarchy Level

```
[edit services radius-flow-tap policy policy-name inet drop-policy rule-name from],
[edit services radius-flow-tap policy policy-name inet6 drop-policy rule-name from]
```

Description

Specify the DSCP value for the radius-flow-tap policy.

Options

dscp-value— IPv4 or IPv6 dscp value for the radius-flow-tap policy.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Subscriber Secure Policy Overview

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

dsl (Access-Line Rate Adjustment)

IN THIS SECTION

- Syntax | 6025
- Hierarchy Level | 6027
- Description | 6027
- Options | 6028
- Required Privilege Level | 6030
- Release Information | 6030

Syntax

```
dsl {
    adsl {
        overhead-bytes bytes;
        total-adjust percentage;
    }
    adsl2 {
        overhead-bytes bytes;
        total-adjust percentage;
    }
    adsl2-plus {
        overhead-bytes bytes;
        total-adjust percentage;
    }
    gfast {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    gfast-bonded {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    other {
```

```
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    sdsl {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    sdsl-bonded {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    type tlv-value {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    vdsl {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    vdsl2 {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    vdsl2-annex-q {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    vdsl2-annex-q-bonded {
        overhead-adjust percentage;
        overhead-bytes bytes;
        total-adjust percentage;
    }
    vdsl2-bonded {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
```

```
    }
}
```

Hierarchy Level

```
[edit system access-line]
```

Description

Configure adjustments to actual DSL data rates as follows:

- Multiply the actual data rate by a percentage for only downstream rates or for both upstream an downstream rates
- Adjust the encapsulation overhead by adding to or subtracting from the total cell or frame bytes a specified number of bytes

The actual (unadjusted) downstream and upstream data rates, DSL line type, and encapsulation mode are received from the access node by the ANCP agent in ANCP port messages, or by the PPPoE daemon from the PPPoE intermediate agent (PPPoE-IA) in PADI or PADR messages. The ANCP agent or PPPoE daemon subsequently adjusts rates and bytes based on the configuration.

If the DSL-Type TLV (0x91) is not received in either the ANCP Port Status message or PPPoE-IA tags, the default adjustment leaves the rates and bytes unchanged.

Adjustments are applied to all subscribers using access lines of the specific DSL line type. Depending on the value, it may be reported to AAA, CoS, or both:

- Adjusted and unadjusted downstream and upstream rates are always reported to AAA in response to an AAA request.
- Adjusted and unadjusted downstream rates and overhead byte adjustments are reported to CoS, but only when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.
- Overhead byte adjustments are not reported to AAA.

AAA reports the adjusted values to the RADIUS server in the Access-Request and Accounting-Request messages through Juniper Networks VSAs 26-141, Downstream-Calculated-Qos-Rate Rate, and 26-142, Upstream-Calculated-Qos-Rate. A change in value triggers an immediate Interim-Accounting message to the RADIUS server if you have configured the "ancp-speed-change-immediate-update" on page 2251 statement.

The ANCP agent reports these values to the LAC in an L2TP network. The LAC passes the rates to the LNS in the following AVPs and messages:

- AVP 24, Tx Connect Speed (ICCN message)—For the initial total rate adjustment to Actual-Data-Rate-Downstream TLV (0x82).
- AVP 38, Rx Connect Speed (ICCN message)—For the initial total rate adjustments to Actual-Data-Rate-Upstream TLV (0x81).
- AVP 97, Connect Speed Update (CSUN message)—For subsequent changes to the initial rates reported in AVP 24 and AVP 38.

Options

adsl Sets attributes for ADSL access lines.

adsl2 Sets attributes for ADSL2 access lines.

adsl2-plus Sets attributes for ADSL2+ access lines.

gfast Sets attributes for G.fast high speed access lines connected to a PON tree infrastructure.

gfast-bonded Sets attributes for G.fast high speed bonded access lines connected to a PON tree infrastructure.

other Sets attributes for access lines of type OTHER. For example, when an OLT sends PON rates in DSL TLVs, the DSL type is set to OTHER.

overhead-adjust ***percentage*** Adjusts the actual downstream rates for all subscribers on the specified access line by multiplying the rate by the specified percentage. This adjustment accounts for the Layer 1 overhead for the DSL type. For the subscriber access line, the adjustment is made to the downstream rate reported in the Actual-Data-Rate-Downstream TLV (0x82). The overhead is adjusted as follows:

- When Agent-Remote-Id TLV (0x02) is different than Access-Aggregation-Circuit-ID-ASCII (0x03), the adjustment shapes the logical interface (residential subscribers) or the interface set (business subscribers), as determined by the CoS adjustment control profile.
- When Agent-Remote-Id TLV (0x02) equals Access-Aggregation-Circuit-ID-ASCII (0x03), the adjustment shapes the PON tree or bonded copper line (parent) interface set, as determined by the CoS adjustment control profile.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

overhead-bytes *bytes*

Adjusts the actual downstream cell overhead for all subscribers on the specified access line by adding or subtracting the specified number of bytes. The adjustment accounts for the traffic encapsulation overhead. The adjustment is made to the overhead reported in the Actual-Data-Rate-Downstream TLV (0x82). The overhead is adjusted as follows:

- Agent-Remote-Id TLV (0x02) is different than Access-Aggregation-Circuit-ID-ASCII (0x03)—The adjustment shapes the logical interface (residential subscribers) or the interface set (business subscribers), as determined by the CoS adjustment control profile.
- Agent-Remote-Id TLV (0x02) equals Access-Aggregation-Circuit-ID-ASCII (0x03)—The adjustment shapes the (parent) interface set for a PON tree (FTTB) or bonded copper line (CuTTB) as determined by the CoS adjustment control profile.
- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

sdsl

Sets attributes for SDSL access lines.

sdsl-bonded

Sets attributes for bonded SDSL access lines.

total-adjust *percentage*

Adjusts the downstream and upstream data rates for all subscribers on an access line of the specified types by multiplying the rate by the specified percentage. This adjustment accounts for the total Layer 1 and encapsulation overhead for the DSL type. The adjustment is made to the Actual-Data-Rate-Downstream TLV (0x82) and Actual-Data-Rate-Upstream TLV (0x81).

- **Range:** 1 through 100 percent
- **Default:** 100 percent

type *tlv-value*

Sets attributes for access lines by specifying the unsigned integer value of the DSL-Type TLV (0x91) to reference the DSL type. This option enables the `dsl` statement to be used for DSL access line types that might be introduced in the future.

- **Range:** 14 through 4294967295

vdsl

Sets attributes for VDSL access lines.

vdsl2

Sets attributes for VDSL2 access lines.

vdsl2-annex-q	Sets attributes for VDSL2 Annex Q access lines.
vdsl2-annex-q-bonded	Sets attributes for bonded VDSL2 Annex Q access lines.
vdsl2-bonded	Sets attributes for bonded VDSL access lines.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Neighbors and Operations

dsl-sfp-options

IN THIS SECTION

- Syntax (SRX300, SRX320, SRX340, SRX345, SRX380) | **6031**
- Hierarchy Level | **6031**
- Description | **6031**
- Options | **6032**
- Required Privilege Level | **6033**
- Release Information | **6033**

Syntax (SRX300, SRX320, SRX340, SRX345, SRX380)

```
dsl-sfp-options {
    adsl-options {
        encap encapsulation;
        vci identifier;
        vpi identifier;
        annex  (auto | annexj-off);
    }
    gfast-options{
        carrier carrier setting;
    }
    vdsl-options {
        carrier carrier setting;
        profile profile;
    }
}
```

Hierarchy Level

```
[edit interfaces interface name]
```

Description

Configure ADSL properties on SRX Series Firewalls.

Follow these example pre-configuration steps with the SFP inserted on slot 8 of the device, required for Annex J support. You must configure one logical interface for *x*DSL SFP control path to work.

1. ADSL and Annex J on `inet` interface:

```
user@host# set interfaces ge-0/0/8 unit 0 family inet address 10.1.1.10/24
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 encap llcsnap-
bridged-802.1q
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 annex auto encap
llcsnap-bridged-802.1q
```

2. ADSL and Annex J without vlan on Ethernet-switching interface:

```
user@host# set interfaces ge-0/0/8 unit 0 family ethernet-switching vlan members xdsl-test
user@host# set interfaces irb unit 50 family inet address 10.1.1.10/24
user@host# set vlans xdsl-test vlan-id 50
user@host# set vlans xdsl-test l3-interface irb.50
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 encap llcsnap-
bridged-802.1q
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 annex auto encap
llcsnap-bridged-802.1q
```

3. ADSL and Annex J with vlan on Ethernet-switching interface:

```
user@host# set interfaces ge-0/0/8 native-vlan-id 50
user@host# set interfaces ge-0/0/8 unit 0 family ethernet-switching interface-mode trunk
user@host# set interfaces ge-0/0/8 unit 0 family ethernet-switching vlan members xdsl-test
user@host# set interfaces irb unit 50 family inet address 10.1.1.10/24
user@host# set vlans xdsl-test vlan-id 50
user@host# set vlans xdsl-test l3-interface irb.50
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 encap llcsnap-
bridged-802.1q
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 annex auto encap
llcsnap-bridged-802.1q
```

Similarly, to set up G.fast option under `dsl-sfp-options`, use:

```
user@host# set interfaces ge-0/0/8 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/8 dsl-sfp-options gfast-options carrier a43
```

Options

adsl-options Options of ADSL interface.

annex (auto | annexj-off) The Annex type.

carrier Carrier setting on VDSL. Supported VDSL carriers are A43, B43, and auto. Supported G.fast carriers are A43, A43c, B43, and B43c.

encap The encapsulation type.

gfast-options	Options of G.fast (Gigabit broadband access technology).
profile	The profile type on VDSL.
vci	Virtual circuit identifier.
vdsl-options	Options.
vpi	Virtual path identifier.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Configuring ADSL Interfaces

dslite-softwire-initiator

IN THIS SECTION

- Syntax | 6034
- Hierarchy Level | 6034
- Description | 6034
- Options | 6034
- Required Privilege Level | 6035

Release Information | 6035

Syntax

```
dslite-softwire-initiator {
    maximum amount;
     reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specifies the number of IPv6 dual-stack lite (DS-Lite) softwire initiators that connects to the softwire concentrator configured in either a user logical system or the primary logical system. The option is configured in the security profile that is bound to the logical system.

Only the primary administrator can create security profiles and bind them to logical systems. The primary administrator:

- Uses security profiles to provision logical systems with resources
- Binds security profiles to user logical systems and the primary logical system
- Configures more than one security profile, specifying different amounts of resource allocations in various profiles

Options

- `maximum` *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources. The default is the system maximum.

- `reserved` *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system. The default is 0.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding IPv6 Dual-Stack Lite in Logical Systems

dsr

IN THIS SECTION

- Syntax | 6035
- Hierarchy Level | 6036
- Description | 6036
- Options | 6036
- Required Privilege Level | 6036
- Release Information | 6036

Syntax

```
dsr (ignore | normal | require);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options dce-options],
[edit interfaces interface-name serial-options dte-options]
```

Description

For EIA-530 and V.35 interfaces only, configure the from-DCE signal, data-set-ready (DSR).

Options

`ignore`—The from-DCE signal is ignored.

`normal`—Normal DSR signal handling as defined by the TIA/EIA Standard 530.

`require`—The from-DCE signal must be asserted.

- **Default:** `normal`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Signal Handling

dsr-polarity

IN THIS SECTION

- Syntax | 6037
- Hierarchy Level | 6037
- Description | 6037
- Options | 6037
- Required Privilege Level | 6037
- Release Information | 6038

Syntax

```
dsr-polarity (negative | positive);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

Configure DSR signal polarity.

Options

`positive`—Positive signal polarity.

`negative`—Negative signal polarity.

- **Default:** `positive`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Serial Signal Polarities

dtcp-only (System Services)

IN THIS SECTION

- Syntax | 6038
- Hierarchy Level | 6038
- Description | 6039
- Required Privilege Level | 6039
- Release Information | 6039

Syntax

```
dtcp-only;
```

Hierarchy Level

```
[edit system services]
```

Description

Prevent RADIUS-initiated subscriber secure policy mirroring from being enabled, while allowing both DTCP-initiated mirroring and DTCP-based flow-tap services (FlowTapLite) to be enabled. This statement has no effect on existing RADIUS-initiated mirroring services. You must issue the statement before such services are activated for a subscriber. Subscriber login and session establishment are not affected.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

- *Disabling RADIUS-Initiated Subscriber Secure Policy Mirroring*
- *Subscriber Secure Policy Overview*
- *Configuring Support for Subscriber Secure Policy Mirroring*
- Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs

dte-options

IN THIS SECTION

- Syntax | 6040
- Hierarchy Level | 6040
- Description | 6040
- Required Privilege Level | 6040
- Release Information | 6040

Syntax

```
dte-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
}
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

For M Series and T Series routers, configure the serial interface signal characteristics.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Statement previously known as `control-leads`.

RELATED DOCUMENTATION

Configuring the Serial Signal Handling

dtr

IN THIS SECTION

- Syntax | 6041
- Hierarchy Level | 6041
- Description | 6041
- Options | 6041
- Required Privilege Level | 6042
- Release Information | 6042

Syntax

```
dtr signal-handling-option;
```

Hierarchy Level

```
[edit interfaces interface-name serial-options dce-options],
[edit interfaces interface-name serial-options dte-options]
```

Description

For EIA-530 and V.35 interfaces only, configure the to-DCE signal, data-transmit-ready (DTR).

Options

signal-handling-option—Signal handling for the DTR signal. The signal handling can be one of the following:

`assert`—The to-DCE signal must be asserted.

`auto-synchronize`—Normal DTR signal with automatic synchronization. This statement has two substatements:

`duration` *milliseconds*—Pulse duration of resynchronization.

- **Range:** 1 through 1000 milliseconds
- **Default:** 1000 milliseconds

`interval` *seconds*—Offset interval for resynchronization.

- **Range:** 1 through 31 seconds
- **Default:** 15 seconds

`de-assert`—The to-DCE signal must be deasserted.

`normal`—Normal DTR signal handling as defined by the TIA/EIA Standard 530.

- **Default:** `normal`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Signal Handling

dtr-circuit

IN THIS SECTION

- Syntax | 6043
- Hierarchy Level | 6043
- Description | 6043
- Options | 6043
- Required Privilege Level | 6043

Release Information | 6043

Syntax

```
dtr-circuit (balanced | unbalanced);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

For EIA-530 and V.35 interfaces only, configure a DTR circuit.

Options

`balanced`—Balanced DTR signal.

`unbalanced`—Unbalanced DTR signal.

- **Default:** `balanced`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial DTR Circuit

dtr-polarity

IN THIS SECTION

- Syntax | 6044
- Hierarchy Level | 6044
- Description | 6044
- Options | 6044
- Required Privilege Level | 6044
- Release Information | 6045

Syntax

```
dtr-polarity (negative | positive);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

Configure DTR signal polarity.

Options

`positive`—Positive signal polarity.

`negative`—Negative signal polarity.

- **Default:** `positive`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Serial Signal Polarities

dual-dr

IN THIS SECTION

- Syntax | 6045
- Hierarchy Level | 6045
- Description | 6046
- Options | 6046
- Required Privilege Level | 6047
- Release Information | 6047

Syntax

```
dual-dr {
    enhanced;
}
```

Hierarchy Level

```
[edit dynamic-profiles name protocols pim interface],
```

```
[edit logical-systems name protocols pim interface],
[edit logical-systems name routing-instances name protocols pim interface],
[edit logical-systems name tenants name routing-instances name protocols pim interface],
[edit protocols pim interface],
[edit routing-instances name protocols pim interface],
[edit tenants name routing-instances name protocols pim interface]
```

Description

Enable PIM dual designated router mode for a pair of Multichassis Link Aggregation Group (MC-LAG) peers managing VLAN multicast traffic and Layer 3 multicast traffic over IRB interfaces.

PIM dual designated router mode sets up one device in a pair of MC-LAG peers as a primary designated router (DR) and the other device as a standby or backup DR for redundancy in managing multicast packet forwarding. Both devices join the multicast forwarding tree and receive multicast traffic. If the primary device fails, the standby quickly takes over forwarding multicast packets with minimal traffic disruption.

In the default or backwards-compatible dual designated router mode (no `enhanced` option specified), both DRs forward packets for a short time while negotiating which device will forward the multicast traffic. After PIM forwarder negotiation, only one device continues to forward traffic. If the primary DR fails, the standby can take over forwarding multicast packets upon detecting the failure. When the primary DR comes back online and is receiving multicast traffic again, it resumes forwarding and invokes PIM forwarder negotiation again to reduce duplicate packet load.

In enhanced dual designated router mode (`enhanced` option specified), both the primary and standby DRs receive multicast traffic, but only the primary DR forwards the packets, skipping PIM forwarder negotiation. Upon primary DR failure, the standby takes over forwarding immediately, and also switches to primary role. When the prior primary DR comes back online, it assumes the standby role, receiving but not forwarding multicast traffic. The new primary DR continues forwarding multicast packets without any traffic interruption, and again no duplicate packet traffic is generated.

If you want to enable dual designated router mode, we recommend using the `enhanced` option, if available, to achieve optimal convergence time in resuming multicast traffic flow both when a DR fails and when it comes back online again.

Options

none Enable default dual designated router mode (available for backwards compatibility).

enhanced Enable enhanced PIM dual designated router mode (recommended if available).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51 for the QFX Series.

Statement introduced in Junos OS Release 16.1 for the MX Series.

Statement with enhanced option introduced in Junos OS Release 17.2R1 for Junos Fusion Data Center.

RELATED DOCUMENTATION

Understanding Multichassis Link Aggregation Groups

dual-phase-bootup

IN THIS SECTION

- Syntax | 6047
- Hierarchy Level | 6048
- Description | 6048
- Default | 6048
- Required Privilege Level | 6048
- Release Information | 6049

Syntax

```
dual-phase-bootup;
```

Hierarchy Level

```
[edit system]
```

Description

You can prevent the device from reaching an amnesiac state post-reboot by configuring the dual-phase-bootup feature before the reboot. When a device has a scaled configuration or has a lot of constraints to be validated, upon reboot it may take more than 45 minutes to finish. This lengthy reboot time exceeds the limit set for the watchdog timer. The watchdog timer going off can cause the device to reach an amnesiac state. To avoid reaching an amnesiac state during a future reboot, configure the `dual-phase-bootup` statement.

If you have configured the `dual-phase-bootup` statement before the reboot, the device picks up the rescue configuration from the next reboot. Post-reboot, the device's operational state is active and the device automatically loads the last-configured user configuration (**juniper.conf** file), thus preventing the device from reaching an amnesiac state.

To be able to commit the configuration for the `dual-phase-bootup` statement, you must already have created a rescue configuration (**rescue.conf** file). We recommend that you have a minimal rescue configuration.

During device upgrades or downgrades, this feature enables a two-step restart process. If the `dual-phase-bootup` statement is already configured and committed, the device picks up the **rescue.conf** file during any upgrade or downgrade. The image is then validated against both the **rescue.conf** file and the **juniper.conf** file. An internal commit is performed post-reboot, with the latest last-run user configuration being committed. While this configuration is being committed, you cannot make any changes to the configuration.

To check the status of an ongoing dual-phase bootup or to check if it is enabled, use the `show dual-phase-bootup-status` operational mode command.

Default

Dual-phase bootup is turned off by default.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Rescue and Recovery of Configuration File (Junos OS)

dual-stack (DHCP Local Server Overrides)

IN THIS SECTION

- Syntax | 6049
- Hierarchy Level | 6049
- Description | 6050
- Options | 6050
- Required Privilege Level | 6050
- Release Information | 6050

Syntax

```
dual-stack dual-stack-group-name;
```

Hierarchy Level

```
[edit logical-systems name routing-instances routing-instance-name system services dhcp-local-
server ...],
[edit logical-systems name system services dhcp-local-server ...],
[edit routing-instances name system services dhcp-local-server ...],
[edit system services dhcp-local-server group group-name interface interface-name overrides],
[edit system services dhcp-local-server group group-name overrides],
[edit system services dhcp-local-server overrides]
```

Description

Assigns the specified dual-stack group to both legs (DHCP and DHCPv6) of the DHCP dual stack. The dual-stack group defines the common configuration settings for DHCP and DHCPv6 subscribers on both legs. These settings take precedence over all other configurations, such as those specified in global, group, or interface settings.

Options

dual-stack-group-name	Name of the globally configured dual-stack group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Single-Session DHCP Dual-Stack Overview

Configuring Single-Session DHCP Dual-Stack Support

dual-stack (DHCP Relay Agent Overrides)

IN THIS SECTION

- Syntax | 6051
- Hierarchy Level | 6051
- Description | 6051
- Options | 6051
- Required Privilege Level | 6051

Release Information | 6052

Syntax

```
dual-stack dual-stack-group-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Assigns the specified dual-stack group to both legs (DHCP and DHCPv6) of the DHCP dual stack. The dual-stack group defines the common configuration settings for DHCP and DHCPv6 subscribers on both legs. These settings take precedence over all other configurations, such as those specified in global, group, or interface settings.

Options

dual-stack-group-name	Name of the globally configured dual-stack group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Single-Session DHCP Dual-Stack Overview

Configuring Single-Session DHCP Dual-Stack Support

dual-stack-group (DHCP Local Server)

IN THIS SECTION

- Syntax | 6052
- Hierarchy Level | 6053
- Description | 6054
- Options | 6054
- Required Privilege Level | 6054
- Release Information | 6054

Syntax

```
dual-stack-group name {
    access-profile access-profile;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name ;
            logical-system-name;
            mac-address;
```

```
            relay-agent-interface-id;
            relay-agent-remote-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    classification-key {
        circuit-id circuit-id;
        mac-address mac-address;
        remote-id remote-id;
    }
    dual-stack-interface-client-limit number;
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    on-demand-address-allocation;
    protocol-primary (inet | inet6);
    reauthenticate (<lease-renewal> <remote-id-mismatch >);
    service-profile service-profile;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server],
[edit logical-systems name system services dhcp-local-server],
[edit routing-instances name system services dhcp-local-server],
[edit system services dhcp-local-server]
```

Description

Specifies common configuration settings that are used for both legs (DHCP and DHCPv6) of the DHCP local server dual-stack, and names the dual-stack group.

When applied, the dual-stack configuration takes precedence over all other configurations, such as those specified in global, group, or interface settings.

Options

name	Name of the dual-stack group.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

- *Single-Session DHCP Dual-Stack Overview*
- DHCP Liveness Detection Using ARP and Neighbor Discovery Packets
- *Configuring RADIUS Reauthentication for DHCP Subscribers*
- *RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers*

dual-stack-group (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 6055
- Hierarchy Level | 6056
- Description | 6057
- Options | 6057
- Required Privilege Level | 6057
- Release Information | 6057

Syntax

```
dual-stack-group name {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name;
            logical-system-name;
            mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    classification-key {
        circuit-id circuit-id;
        mac-address mac-address;
        remote-id remote-id;
    }
```

```
    dual-stack-interface-client-limit number;
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    protocol-primary (inet | inet6);
    relay-agent-interface-id {
        include-irb-and-l2;
        keep-incoming-interface-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    relay-agent-remote-id {
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    service-profile dynamic-proflle-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay]
```

Description

Specifies common configuration settings that are used for both legs (DHCP and DHCPv6) of the DHCP dual stack, and names the dual stack group.

The group is assigned to each leg of the DHCP dual-stack with the "dual-stack" on page 6050 statement in the `overrides` stanza. When applied, the dual-stack configuration takes precedence over all other configurations, such as those specified in global, group, or interface settings.

Options

name	Name of the dual-stack group.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

- *Single-Session DHCP Dual-Stack Overview*
- *Configuring Single-Session DHCP Dual-Stack Support*
- DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

dual-stack-interface-client-limit (DHCP Local Server and Relay Agent)

IN THIS SECTION

- Syntax | 6058
- Hierarchy Level | 6058
- Description | 6058
- Options | 6059
- Required Privilege Level | 6059
- Release Information | 6059

Syntax

```
dual-stack-interface-client-limit number;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dual-stack-group
dual-stack-group-name ],
[edit logical-systems name routing-instances name system services dhcp-local-server dual-stack-
group dual-stack-group-name ],
[edit logical-systems name system services dhcp-local-server dual-stack-group dual-stack-group-
name ],
[edit routing-instances name forwarding-options dhcp-relay dual-stack-group dual-stack-group-
name ],
[edit routing-instances name system services dhcp-local-server dual-stack-group dual-stack-group-
name ],
[edit system services dhcp dhcp-local-server dual-stack-group dual-stack-group-name ],
```

Description

Limit the number of clients allowed on an interface.

> **NOTE**: For dual-stack subscribers, always use this statement instead of the "interface-client-limit" on page 9108 (DHCP Relay Agent) or "interface-client-limit" on page 9106 (DHCP Local Server) statements.

Options

number Maximum number of dual-stack subscribers that can log in per interface.

- **Range:** 1 through 500,000

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Single-Session DHCP Dual-Stack Overview

dual-transport

IN THIS SECTION

- Syntax | 6060
- Hierarchy Level | 6060
- Description | 6060
- Options | 6060
- Required Privilege Level | 6060

Release Information | 6060

Syntax

```
dual-transport {
    inet-lsr-id inet-lsr-id;
    inet6-lsr-id inet6-lsr-id;
}
```

Hierarchy Level

```
[edit protocols ldp]
```

Description

Configure to allow Junos LDP to establish the TCP connection over IPv4 with IPv4 neighbors, and over IPv6 with IPv6 neighbors as a single-stack LSR. `inet-lsr-id` and `inet6-lsr-id` are the two LSR IDs that have to be configured to establish an LDP session over IPv4 and IPv6 TCP transport. These two IDs should be non-zero and must be configured with different values.

Options

inet-lsr-id *inet-lsr-id* Configure the LSR ID to establish an LDP session over IPv4 TCP transport.

inet6-lsr-id *inet6-lsr-id* Configure the LSR ID to establish an LDP session over IPv6 TCP transport.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

LDP Native IPv6 Support Overview

Example: Configuring LDP Native IPv6 Support

Configuring LDP Native IPv6 Support

family (protocols ldp) | 6831

dualstack-support (JSRC)

IN THIS SECTION

- Syntax | 6061
- Hierarchy Level | 6061
- Description | 6061
- Default | 6062
- Required Privilege Level | 6062
- Release Information | 6062

Syntax

```
dualstack-support;
```

Hierarchy Level

```
[edit jsrc]
```

Description

Configure JSRC provisioning for dual-stack subscribers so that it reports information about the separate stacks for a given subscriber, using a single JSRC session. By default (and in all cases for releases earlier than Junos OS Release 18.1R1), the DHCPv4 and DHCPv6 stacks are treated as a single subscriber; the

remote SRC peer (SAE) is not informed about whether only one family or both families are active. The statistics are reported as an aggregate of both families rather than separated by family.

When you configure dual-stack support for JSRC, Diameter AA-Request (AAR) provisioning messages sent to the SAE include Diameter AVPs (IANA enterprise number 2636) to convey the IPv4 and IPv6 addressing information that is available in the session database. For IPv4, that includes Framed-IP-Address (AVP 8) and Framed-IP-Netmask (AVP 9). For IPv6, that includes Framed-IPv6-Address (AVP 168), Framed-IPv6-Prefix (AVP 97), Delegated-IPv6-Prefix (AVP 123), Juniper-IPv6-Ndra-Prefix (AVP 2200), and Juniper-Framed-IPv6-Netmask (AVP 2201). JSRC also includes information about the access line if it is available in the session database, by means of Juniper-Agent-Circuit-Id (AVP 2202) and Juniper-Remote-Circuit-Id (AVP 2203).

When the first network family is activated, JSRC sends the addresses for only that family in the initial request to the provisioning server. When the second network family is activated, the AAR message includes the Juniper-Request-Type AVP (2050) with a value of 4 to signify family activation. When the next-to-last family is deactivated, the same AVP is sent with a value of 5 to signify the deactivation.

Default

Disabled.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

JSRC Configuration Overview

dump-on-flow-control

IN THIS SECTION

- Syntax | 6063
- Hierarchy Level | 6063
- Description | 6063
- Required Privilege Level | 6063
- Release Information | 6063

Syntax

```
dump-on-flow-control;
```

Hierarchy Level

```
[edit interfaces interface-name multiservice-options]
```

Description

This option supports high availability functionality and can be used with various service interfaces, including `rsp`, `rms`, `lsq`, and `rlsq`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Multiservice Physical Interface Properties

Junos OS Services Interfaces Library for Routing Devices

passive-monitor-mode | 12212

duplicate-clients (DHCPv6 Local Server and Relay Agent)

IN THIS SECTION

- Syntax | 6064
- Hierarchy Level | 6064
- Description | 6065
- Options | 6065
- Required Privilege Level | 6065
- Release Information | 6066

Syntax

```
duplicate-clients incoming-interface;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server dhcpv6]
```

Description

Specify the criteria that the `jdhcpd` process uses to support duplicate clients. The router uses the additional criteria to distinguish between the duplicate clients.

Duplicate clients have the same DUID (DHCP unique identifier). Typically, the router treats a request from a duplicate client as a renegotiation, and replaces the existing client entry with a new entry. However, in some cases, the duplicate request is from a different client, and replacement is not desired. When you enable duplicate client support, the router uses the additional criteria to distinguish between the two clients, and grants a lease to the new client while retaining the original client entry.

NOTE: The only supported differentiating criterion is `incoming-interface`.

BEST PRACTICE: To allow duplicate clients over the incoming interface for DHCPv6 relay, you must configure the `relay-agent-interface-id` statement to cause the DHCP relay agent to insert the DHCPv6 Interface-ID option (option 18) in DHCPv6 packets destined for the DHCPv6 server.

Do not configure the `use-interface-description` statement, because option 18 must include the interface name rather than an interface description.

CAUTION: We recommend that you do not enable or disable duplicate client support mode when clients are bound, because different client keys are used to store the clients in the database depending on the mode. Changing the mode removes clients from the database and then adds them back with a different key.

Additionally, disabling duplicate client support mode causes all duplicate clients to be deleted.

Options

incoming-interface Allow duplicate clients when the incoming DHCPv6 requests are received over different underlying interfaces.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

DHCPv6 Duplicate Client DUIDs

Configuring the Router to Use Underlying Interfaces to Distinguish Between DHCPv6 Duplicate Client DUIDs

duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)

IN THIS SECTION

- Syntax | 6066
- Hierarchy Level | 6066
- Description | 6067
- Options | 6067
- Required Privilege Level | 6067
- Release Information | 6067

Syntax

```
duplicate-clients-in-subnet (incoming-interface | option-82);
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server],
[edit logical-systems logical-system-name system services dhcp-local-server],
[edit routing-instances routing-instance-name system services dhcp-local-server],
[edit system services dhcp-local-server]
```

Description

Configure how the router distinguishes between duplicate clients in the same subnet. Duplicate clients are defined as clients that have the same hardware address or client ID.

> **NOTE**: You must configure the `duplicate-clients-in-subnet` statement identically for both the DHCP local server (`[edit forwarding-options dhcp-relay]`) and the DHCP relay agent (`[edit system services dhcp-local-server]`).

Options

incoming-interface Use the incoming interface information in packets to differentiate between duplicate clients.

option-82 Use the option 82 information to differentiate between duplicate clients. Starting in Junos OS Release 16.1R5, 16.2R2, 17.1R2, and 17.2R1, only the ACI (suboption 1) and ARI (suboption 2) are used. Other suboptions, such as Vendor-Specific (suboption 9) are ignored.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
16.1R5	Starting in Junos OS Release 16.1R5, 16.2R2, 17.1R2, and 17.2R1, only the ACI (suboption 1) and ARI (suboption 2) are used. Other suboptions, such as Vendor-Specific (suboption 9) are ignored.

RELATED DOCUMENTATION

DHCPv4 Duplicate Client In Subnet Overview

Guidelines for Configuring Support for DHCPv4 Duplicate Clients

Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Their Incoming Interfaces

Configuring the Router to Distinguish Between DHCPv4 Duplicate Clients Based on Option 82 Information

duplicate-mac-detection

IN THIS SECTION

- Syntax | 6068
- Hierarchy Level | 6069
- Description | 6069
- Options | 6069
- Required Privilege Level | 6070
- Release Information | 6070

Syntax

```
duplicate-mac-detection {
    action(block | shutdown);
    auto-recovery-time minutes;
    detection-threshold detection-threshold;
```

```
    detection-window seconds;
    include-local-moves;
}
```

Hierarchy Level

```
[edit logical-systems  logical-systems-name protocols evpn],
[edit logical-systems logical-systems-name routing-instances  routing-instance-name protocols evpn],
[edit protocols evpn]
[edit routing-instances  routing-instance-name protocols evpn],
```

Description

Configure duplicate MAC address detection settings. You can configure a threshold for MAC address mobility events and the window for detecting the number of MAC address mobility events. In addition, you can configure the optional recovery time that the Juniper Networks device waits before the duplicate MAC address is unsuppressed.

Prevent loops at the forwarding plane. You can block a duplicate MAC address or shut down the duplicate MAC address's local interfaces. After you block a duplicate MAC address, you can clear it. After you shut down the interface, you can recover it.

Options

action Block the duplicate MAC address. You can also shut down the duplicate MAC address's corresponding local interface.

- block—When the device detects a duplicate MAC address, the device blocks the duplicate MAC address. If you have enabled the `auto-recovery-time` statement, the auto-recovery timer starts. While the timer's running, any duplicate MAC addresses are blocked. When the auto-recovery timer stops, the system clears the duplicate MAC addresses. Alternatively, you can issue the `clear evpn duplicate-mac-suppression <instance name | l2-domain-id | mac-address>` command to manually clear the duplicate MAC addresses.

- shutdown—When the device detects a duplicate MAC address, the device shuts down the local interface. If you enable the `recovery-timeout`statement--for example, `set interfaces et-0/0/0 unit 0 family ethernet-switching recovery-timeout 10`--the auto-recovery timer starts. When the auto-recovery timer stops, the interface is recovered. Alternatively, you can issue the `clear ethernet-switching recovery-timeout` command to manually recover the interface that was shut down.

auto-recovery-time

The length of time a device suppresses a duplicate MAC address. At the end of this interval, MAC address updates will resume. When an auto recovery time is not specified, duplicate MAC address updates will continue to be suppressed. To manually clear the suppression of duplicate MAC address, use the `clear evpn duplicate-mac-suppression` command.

- **Range:** `1-360` minutes

detection-threshold

Both this statement and the `detection-window` statement help detect if the MAC address is a duplicate address. A certain number of MAC mobility events need to be detected before a MAC address before it is identified as a duplicate MAC address. Once the threshold is reached, updates for this MAC address are suppressed.

- **Default:** 5
- **Range:** `2-20`

detection-window

Both this statement and the `detection-threshold` statement help detect if the MAC address is a duplicate address based on the time interval used in detecting a duplicate MAC address. When the number of MAC mobility events for a MAC address exceeds the `detection-threshold` within the detection window, the MAC address is identified as a duplicate MAC address.

- **Default:** 180 seconds
- **Range:** `5-600` seconds

include-local-moves Track local MAC address mobility movements.

> **NOTE**: To ensure that the mobility advertisements have sufficient time to age out, set an auto-recovery-time greater than the detection window.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Overview of MAC Mobility

Changing Duplicate MAC Address Detection Settings

Configuring Loop Detection for Duplicate MAC Addresses

duplicate-protection (Dynamic PPPoE)

IN THIS SECTION

- Syntax | 6071
- Hierarchy Level | 6071
- Description | 6072
- Required Privilege Level | 6072
- Release Information | 6072

Syntax

```
duplicate-protection;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
pppoe],
[edit interfaces interface-name unit logical-unit-numberfamily pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
pppoe-underlying-options]
```

Description

Prevent the activation of another dynamic PPPoE logical interface on the same underlying interface when a dynamic PPPoE logical interface for a client with the same media access control (MAC) address is already active on that interface. Duplicate protection is disabled by default. Enabling duplicate protection has no effect on dynamic PPPoE logical interfaces that are already active.

> **NOTE**: The `[edit ... family pppoe]` hierarchies are supported only on MX Series routers with MPCs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

Support for the `[edit ... family pppoe]` hierarchies introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

- *Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces*
- *Configuring the PPPoE Family for an Underlying Interface*
- *Configuring Lockout of PPPoE Subscriber Sessions*
- *Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview*

duplicates-dropped-periodicity

IN THIS SECTION

- Syntax | 6073
- Hierarchy Level | 6073

- Description | 6073
- Options | 6073
- Required Privilege Level | 6073
- Release Information | 6073

Syntax

```
duplicates-dropped-periodicity seconds;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Specify the frequency for sending notifications to affected control sources when transmission of duplicate sets of data is restricted because the `max-duplicates` threshold has been reached.

Options

seconds—Period for sending DuplicatesDropped notifications.

- **Default:** 30 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

g-duplicates-dropped-periodicity | 7645

Limiting the Number of Duplicates of a Packet

max-duplicates | 10436

duplication (Access Profile)

IN THIS SECTION

- Syntax | 6074
- Hierarchy Level | 6074
- Description | 6074
- Default | 6075
- Required Privilege Level | 6075
- Release Information | 6075

Syntax

```
duplication;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router to send accounting reports to both the RADIUS accounting server configured in the access profile for the wholesaler and the RADIUS accounting server configured in the access profile for the retailer.

Default

The router sends accounting reports to the accounting servers that are in the context in which the subscriber is authenticated.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Authentication and Accounting Parameters for Subscriber Access

Understanding RADIUS Accounting Duplicate Reporting

duplication-filter (Access Profile)

IN THIS SECTION

- Syntax | 6076
- Hierarchy Level | 6076
- Description | 6076
- Options | 6076
- Required Privilege Level | 6076
- Release Information | 6076

Syntax

```
duplication-filter (interim-duplicated | interim-original) <exclude-attributes>;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure which accounting servers receive accounting messages when RADIUS duplication accounting is configured in the access profile.

Options

exclude-attributes (Optional) Filter RADIUS attributes from duplicated accounting interim messages based on the `exclude` statement configuration in the corresponding duplication access profile. You can configure `exclude-attributes` alone, or with `interim-duplicated` or `interim-original`.

interim-duplicated Do not send accounting interim messages to RADIUS accounting servers that are in a duplication context other than the subscriber's access profile.

interim-original Do not send accounting interim messages to RADIUS accounting servers that are in the subscriber's access profile.

NOTE: The `interim-duplicated` and `interim-original` filters are mutually exclusive.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Authentication and Accounting Parameters for Subscriber Access

Understanding RADIUS Accounting Duplicate Reporting

Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting

Configuring Per-Subscriber Session Accounting

duplication-vrf (Duplicate Accounting)

IN THIS SECTION

- Syntax | 6077
- Hierarchy Level | 6077
- Description | 6077
- Required Privilege Level | 6078
- Release Information | 6078

Syntax

```
duplication-vrf {
    access-profile-name profile-name;
    vrf-name vrf-name;
}
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router to send duplicate accounting information to the RADIUS accounting servers defined in up to five access profiles all in the same nondefault VRF (LS:RI combination).

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.

Statement supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

RELATED DOCUMENTATION

Understanding RADIUS Accounting Duplicate Reporting

Configuring Authentication and Accounting Parameters for Subscriber Access

duration

IN THIS SECTION

- Syntax | 6079
- Hierarchy Level | 6079
- Description | 6079
- Options | 6079
- Required Privilege Level | 6079
- Release Information | 6079

Syntax

```
duration hours;
```

Hierarchy Level

```
[edit poe interface (all | interface-name) telemetries]
```

Description

Modify the duration over which data is collected when you are monitoring the power consumption of a PoE interface.

Options

hours—Number of hours over which the data is to be collected.

- **Range:** 1 through 24
- **Default:** 1

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

> *Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch*
>
> *Configuring PoE Interfaces on EX Series Switches*

duration (PoE)

IN THIS SECTION

- Syntax | 6080
- Hierarchy Level | 6080
- Description | 6080
- Options | 6080
- Required Privilege Level | 6081
- Release Information | 6081

Syntax

```
duration hours;
```

Hierarchy Level

```
[edit poe interface (all | interface-name) telemetries]
```

Description

Modifies the duration for which telemetry records are stored. If telemetry logging continues beyond the specified duration, the older records are discarded one by one as new records are collected.

Options

hours— Hours for which telemetry data should be retained.

- **Range:** 1 through 24 hours
- **Default:** 1 hour

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

dvmrp

IN THIS SECTION

- Syntax | 6081
- Hierarchy Level | 6082
- Description | 6082
- Default | 6082
- Options | 6082
- Required Privilege Level | 6082
- Release Information | 6082

Syntax

```
dvmrp {
    disable;
    export [ policy-names ];
    import [ policy-names ];
    interface interface-name {
        disable;
        hold-time seconds;
        metric metric;
        mode (forwarding | unicast-routing);
    }
    rib-group group-name;
```

```
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit protocols]
```

Description

Enable DVMRP on the router or switch.

Default

DVMRP is disabled on the router or switch.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

NOTE: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring DVMRP

dynamic (services)

IN THIS SECTION

- Syntax | 6083
- Hierarchy Level | 6083
- Description | 6083
- Options | 6084
- Required Privilege Level | 6084
- Release Information | 6084

Syntax

```
dynamic {
    ike-policy policy-name;
    ipsec-policy policy-name;
}
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Define a dynamic IPsec SA.

Options

`ike-policy` *`policy-name`*—Name of the IKE policy. This statement is optional for the non-preshared-key authentication method. For digital signature-based authentication, this statement is optional and the default policy is used if none is supplied.

`ipsec-policy` *`policy-name`*—Name of the IPsec policy. This statement is optional and the default policy is used if none is supplied.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Security Associations

dynamic

IN THIS SECTION

- Syntax | 6085
- Hierarchy Level | 6085
- Description | 6085
- Options | 6085
- Required Privilege Level | 6085
- Release Information | 6085

Syntax

```
dynamic {
    ipsec-policy ipsec-policy-name;
    replay-window-size (32 | 64);
}
```

Hierarchy Level

```
[edit security ipsec security-association name]
```

Description

Define a dynamic IPsec SA.

Options

`ipsec-policy` *`ipsec-policy-name`*—Name of the IPsec policy.

`replay-window-size`—(Optional) Antireplay window size. It can be one of the following values:

- `32`—32-packet window size.
- `64`—64-packet window size.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring Dynamic IPsec Security Associations*
>
> *Associating the Configured Security Association with a Logical Interface*

dynamic-address

IN THIS SECTION

- Syntax | 6086
- Hierarchy Level | 6087
- Description | 6087
- Options | 6087
- Required Privilege Level | 6087
- Release Information | 6088

Syntax

```
dynamic-address {
    address-name name {
        description description;
        profile {
            category name {
                feed feed;
                property name {
                    string name;
                }
            }
            feed-name name;
        }
    }
    feed-server feed-server-name {
         url url);
        description description;
        feed-name name {
            description description;
            hold-interval seconds;
            path path;
            update-interval seconds;
        }
        hold-interval seconds;
        hostname hostname;
```

```
        update-interval seconds;
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security]
[edit logical-systems logical-system-name security]
[edit tenants tenant-name security ]
```

Description

Configure security dynamic address for IPv4 and IPv6 networks within a logical system and tenant system. Each dynamic address belongs to only one instance. Within that instance is a set of categories to which the dynamic address further belongs. A dynamic address entry provides dynamic IP address information to the security policies.

Options

address-name Security dynamic address name.

feed-server Security dynamic address feed-server.

traceoptions Security dynamic address tracing options.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

dynamic-address

IN THIS SECTION

- Syntax | 6088
- Hierarchy Level | 6089
- Description | 6089
- Required Privilege Level | 6089
- Release Information | 6090

Syntax

```
dynamic-address {
address-name name {
 description description;
 profile {
  category name {
   feed feed;
   property name {
    string name;

   }
  }
  feed-name name;

 }
 session-scan;
}
feed-server name {
 (hostname hostname | url url);
```

```
  description description;
  feed-name name {
   description description;
   hold-interval seconds;
   path path;
   update-interval seconds;
  }
  hold-interval seconds;
  tls-profile tls-profile;
  update-interval seconds;
  validate-certificate-attributes {
   subject-or-subject-alternative-names;
  }
 }
 session-scan {
  hold-interval seconds;
 }
 traceoptions {
  file< filename><files files><match match><size size><(world-readable | no-world-readable)>;
  flag name;

  level(all | error | info | notice | verbose | warning);
  no-remote-trace;
 }
 }
```

Hierarchy Level

```
[edit security]
```

Description

Configure security dynamic address

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D25 on SRX Series Firewalls and on cSRX Container Firewall in Junos OS Release 22.2R1.

dynamic-application

IN THIS SECTION

- Syntax | 6090
- Hierarchy Level | 6090
- Description | 6090
- Options | 6091
- Required Privilege Level | 6091
- Release Information | 6091

Syntax

```
dynamic-application [system-application];
```

Hierarchy Level

```
[edit security application-firewall rule-sets rule-set-name rule rule-name match]
```

Description

Specify the dynamic application names for match criteria in application firewall rule set.

An application firewall configuration permits, rejects, or denies traffic based on the application of the traffic. The AppFW consists of one or more rule sets with rules that specify match criteria, including dynamic applications, and the action to be taken for matching traffic.

The junos:UNKNOWN keyword is reserved for unknown dynamic applications. In the following cases, the application ID is set to junos:UNKNOWN:

- The traffic does not match an application signature in the database.
- The system encounters an error when identifying the application.
- The session fails over to another device.

Traffic with an application ID of junos:UNKNOWN matches a rule with a dynamic application of junos:UNKNOWN. If there is no rule defined for junos:UNKNOWN, the default rule is applied.

Starting in Junos OS Release 18.2R1 application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration.

Options

`system-application`—Set of system applications for match criteria.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Application Firewall Overview

dynamic-application-group

IN THIS SECTION

- Syntax | 6092
- Hierarchy Level | 6092
- Description | 6092
- Options | 6093
- Required Privilege Level | 6093
- Release Information | 6093

Syntax

```
dynamic-application-group [system-application-group];
```

Hierarchy Level

```
[edit security application-firewall rule-sets rule-set-name rule rule-name match]
```

Description

Specify the dynamic application group to match. When you define application firewall rules, you can specify dynamic application groups as match criteria.

With application identification, multiple applications can be configured in a dynamic application groups for consistent reuse. AppFW rules permit and deny traffic by specifying application names, dynamic application group names, or both. By using predefined application groups, AppFW rules require no updating when new applications are added to common groups.

Starting in Junos OS Release 18.2R1 application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration.

Options

system-application-group—Set of groups defining one or more system applications for match criteria.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Application Firewall Overview

dynamic-application (Security)

IN THIS SECTION

- Syntax | 6093
- Hierarchy Level | 6094
- Description | 6094
- Options | 6094
- Required Privilege Level | 6094
- Release Information | 6095

Syntax

```
dynamic-application {
    profile profile name {
        redirect-message {
```

```
            type {
                custom-text {
                    content content;
                }
                redirect-url {
                    content content;
                }
            }
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure options for dynamic applications to define a redirect profile and traceoptions.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Unified Security Policies

Unified Security Policies

dynamic-application (Security Policies)

IN THIS SECTION

- Syntax | 6095
- Hierarchy Level | 6095
- Description | 6096
- Options | 6096
- Required Privilege Level | 6097
- Release Information | 6097

Syntax

```
dynamic-application dynamic-application-name | dynamic-application-group-name | any | none;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name match]
[edit security policies global policy policy-name match]
```

Description

Specify the dynamic applications or dynamic application groups used as match criteria within a security policy.

By adding dynamic applications to the match criteria, the data traffic is classified based on the Layer 7 application inspection results. Application Identification (AppID) identifies dynamic or real-time Layer 4 through Layer 7 applications. After a particular application is identified and the matching policy is found, then the actions are applied according to the policy.

Options

dynamic-application-name | dynamic-application-group-name

Specify dynamic applications or dynamic application groups.

Examples for dynamic applications or dynamic application groups are as follows:

- `junos:FTP` (dynamic application).
- Enter `junos:UNKNOWN` when no dynamic application can be determined.
- `junos:web:shopping` (dynamic application group).
- Enter `junos:unassigned` when no dynamic application group can be determined.
- `junos:all-new-apps` (dynamic application group).

  This has newly added application signatures from the latest signature package released in Junos OS Release 21.1R1.

> **NOTE**: When you download the application signature package on your security device, the entire predefined applications and application groups are downloaded. For downloading and installing application signature package, see Predefined Application Signatures for Application Identification

any

Configuring the dynamic application as `any` installs the policy with the application as a wildcard (default). If an application cannot be specified, configure `any` as the default application. Data traffic that match the parameters in a unified policy matches the policy regardless of the application type.

none

Configuring the dynamic application as `none` ignores classification results from AppID and does not use the dynamic application in security policy lookups. Within the list of potential match policies, if there is any policy configured with dynamic application as

none, irrespective of the policy match sequence, this policy is matched as the final policy and is terminal. If any Layer 7 services are configured in this policy, deep packet inspection for the traffic is performed.

When upgrading the Junos OS release from previous releases (where dynamic applications were not supported), all existing traditional policies are considered as policies with the dynamic application configured as none.

> **TIP**: Ensure to match an application if dynamic-application is set to none in the match criteria

If dynamic application is not configured within a security policy, the policy is considered to be a traditional security policy. This policy is similar to a policy with the dynamic application configured as none.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Unified Security Policies

Application Identification Support for Unified Policies

Example: Configure Application Firewall with Unified Policy

dynamic-attack-group

IN THIS SECTION

Syntax | 6098

- Hierarchy Level | 6099
- Description | 6099
- Options | 6099
- Required Privilege Level | 6100
- Release Information | 6100

Syntax

```
dynamic-attack-group name {
     filters {
        age-of-attack
            {
            greater-than value;
            less-than value;
        }
        category (Security Dynamic Attack Group) {
            values [ values ... ];
        }
        cvss-score
            {
            greater-than value;
            less-than value;
        }
        direction (Security Dynamic Attack Group) {
            expression (and | or);
            values (any | client-to-server | exclude-any | exclude-client-to-server | exclude-
server-to-client | server-to-client);
        }
         Excluded {
        }
         no-excluded {
        }
        false-positives {
            values (frequently | occasionally | rarely | unknown);
        }
        file-type {
            values [ values ... ];
        }
```

```
        performance {
            values (fast | normal | slow | unknown);
        }
        (recommended | no-recommended);
        service (Security IDP Dynamic Attack Group) {
            values [ values ... ];
        }
        severity (Security IDP Dynamic Attack Group) {
            values (critical | info | major | minor | warning);
        }
        type (Security IDP Dynamic Attack Group) {
            values (anomaly | signature);
        }
        vendor name {
            product-name product-name;
        }
        vulnerability-type {
            values [ values ... ];
        }
    }
}
```

Hierarchy Level

```
[edit security idp]
[edit tenants tenant-name security idp]
```

Description

Configure a dynamic attack group. A dynamic attack group selects its members based on the filters specified in the group. Therefore, the list of attacks is updated (added or removed) when a new signature database is used.

Options

dynamic-attack-group-name—Name of the dynamic attack group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

The `expression` option added in Junos OS Release 11.4.

Additional tags under filters of dynamic attack groups (CVSS score, age-of-attack, file-type, vulnerability-type) are added in Junos OS Release 18.2R1 for dynamic attacks grouping of IDP signatures. The Product and Vendor tags are already supported under existing filter products. The CLI interface for configuring these tags is now more user friendly with possible completions being available for configuration in 18.2R1.

The `Excluded` and `no-excluded` filters are added in Junos OS Release 19.1R1.

dynamic-attack-groups (Security IDP)

IN THIS SECTION

- Syntax | 6100
- Hierarchy Level | 6101
- Description | 6101
- Options | 6101
- Required Privilege Level | 6101
- Release Information | 6101

Syntax

```
dynamic-attack-groups attack-group-name;
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match attacks]
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match attacks]
```

Description

Specify a name for the dynamic attack group.

Options

attack-group-name—Name of the dynamic attack group.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

dynamic-bidirectional-transport

IN THIS SECTION

- Syntax | 6102
- Hierarchy Level | 6102
- Description | 6102
- Options | 6102
- Required Privilege Level | 6102
- Release Information | 6102

Syntax

```
dynamic-bidirectional-transport {
    template template;
}
```

Hierarchy Level

```
[edit protocols rsvp peer-interface peer-interface-name]
```

Description

Enable the dynamic setup of associated bidirectional packet LSP for transporting non-packet Generalized Multiprotocol Label Switching (GMPLS) label-switched path (LSP).

Options

template *template* Name of the template for the dynamic bidirectional packet LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

dynamic-class-of-service-options (Dynamic Traffic Shaping)

IN THIS SECTION

Syntax | 6103

- Hierarchy Level | 6103
- Description | 6103
- Options | 6103
- Required Privilege Level | 6104
- Release Information | 6104

Syntax

```
dynamic-class-of-service-options {
    vendor-specific-tags tag;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service]
```

Description

Configure CoS attributes applied to dynamic interfaces or interface sets based on access line information received in PPPoE discovery packets on dynamic subscriber interfaces. The packets convey the information in PPPoE-IA tags.

Options

`vendor-specific-tags` *tag*—Specify one or more of the following PPPoE-IA tags to set CoS attributes:

- `access-loop-encapsulation`—Set the overhead-accounting CoS attribute based on the Access-Loop-Encapsulation tag (0x91). The tag enables the router to determine whether the access loop is cell-based or frame-based.
- `actual-data-rate-downstream`—Set the shaping-rate class-of-service attribute based on the Actual-Data-Rate-Downstream tag (0x82).
- `apply-to-interface-set`—Enable CoS to determine the node to which it conditionally applies the adjusted downstream rate. By default, when a CoS traffic control profile is configured for both a logical interface and the underlying (child) interface set, the adjusted rate values (shaping rate,

overhead bytes, and overhead mode) are applied to the logical interface, according to the adjustment control profile. When you specify this option, the following conditions are evaluated:

- If the $junos-interface-set-name has been resolved by AAA authorization, CoS applies the adjusted values to the underlying interface set to which the logical interface belongs. The name for this predefined variable is supplied by the Juniper Networks Qos-Set-Name VSA, 26–130. The application of the adjusted values is governed by the adjustment control profile, regardless of whether a traffic control profile is present for the logical interface.
- If the $junos-interface-set-name has not been resolved by AAA authorization, CoS applies the adjusted values to the logical interface, governed by the adjustment control profile. If a traffic control profile is not present for the logical interface, then the adjusted values are applied to the underlying interface set.

This option is appropriate for both five-level and four-level CoS hierarchies.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

`apply-to-interface-set` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Setting Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags

Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces

Bandwidth Management for Downstream Traffic in Edge Networks Overview

dynamic-db

IN THIS SECTION

- Syntax | 6105
- Hierarchy Level | 6105
- Description | 6105
- Required Privilege Level | 6106
- Release Information | 6106

Syntax

```
dynamic-db;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options as-path path-name],
[edit logical-systems logical-system-name policy-options as-path-group group-name],
[edit logical-systems logical-system-name policy-options community community-name],
[edit logical-systems logical-system-name policy-options condition condition-name],
[edit logical-systems logical-system-name policy-options policy-statement policy-statement-name],
[edit logical-systems logical-system-name policy-options prefix-list prefix-list-name],
[edit policy-options as-path path-name],
[edit policy-options as-path-group group-name],
[edit policy-options community community-name],
[edit policy-options condition condition-name],
[edit policy-options policy-statement policy-statement-name],
[edit policy-options prefix-list prefix-list-name]
```

Description

Define routing policies and policy objects that reference policies configured in the dynamic database at the `[edit dynamic]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Example: Configuring Dynamic Routing Policies

dynamic-dns

IN THIS SECTION

- Syntax | 6106
- Hierarchy Level | 6107
- Description | 6107
- Options | 6107
- Required Privilege Level | 6107
- Release Information | 6107

Syntax

```
dynamic-dns {
    client hostname {
        agent agent-name;
        interface interface-name;
        password server-password;
        server server-name;
        username user-name;
```

```
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure the device as a dynamic DNS server that maintains the list of the changed addresses and their associated domain names registered with it. The device updates these DDNS servers with this information periodically or whenever there is a change in IP addresses.

Options

- `client` —Specifies the hostname of the registered client.
- `agent`—Specifies the name of the dynamic DNS agent.
- `interface` —Specifies the interface whose IP address is mapped to the registered domain name.
- `password`—Specifies the password.
- `server`—Specifies the name of the dynamic DNS server that allows dynamic DNS clients to update the IP address changes associated with the registered hostname.
- `username`—Specifies the dynamic DNS username.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

DNS Overview

dynamic-filter

IN THIS SECTION

- Syntax | 6108
- Hierarchy Level | 6108
- Description | 6108
- Required Privilege Level | 6109
- Release Information | 6109

Syntax

```
dynamic-filter;
```

Hierarchy Level

```
[edit services security-metadata-streaming policy]
```

Description

Configure dynamic filtering options for security metadata streaming policy on SRX Series Firewalls. Security metadata streaming policy sends the metadata and connection patterns of a IoT device traffic to Juniper Networks ATP Cloud. Juniper ATP Cloud uses the streamed data to get details such as brand, device model, class, vendor, IP, MAC address, and other properties of IoT devices.

After configuring the security metadata streaming policy, attach it to the security policy at zone-level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1R1 on SRX Series Firewalls with Juniper Advanced Threat Prevention Cloud (Juniper ATP Cloud).

dynamic-flow-capture

IN THIS SECTION

- Syntax | 6109
- Hierarchy Level | 6110
- Description | 6110
- Required Privilege Level | 6110
- Release Information | 6110

Syntax

```
dynamic-flow-capture {
    capture-group client-name {
        content-destination identifier {
            address address;
            hard-limit bandwidth;
            hard-limit-target bandwidth;
            soft-limit bandwidth;
            soft-limit-clear bandwidth;
            ttl hops;
        }
        control-source identifier {
            allowed-destinations [ destinations ];
```

```
            minimum-priority value;
            no-syslog;
            notification-targets address port port-number;
            service-port port-number;
            shared-key value;
            source-addresses [ addresses ];
        }
        duplicates-dropped-periodicity seconds;
        input-packet-rate-threshold rate;
        interfaces interface-name;
        max-duplicates number;
        pic-memory-threshold percentage percentage;
    }
    g-duplicates-dropped-periodicity seconds;
    g-max-duplicates number;
}
```

Hierarchy Level

```
[edit services]
```

Description

Define the dynamic flow capture properties to be applied to traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding Junos Capture Vision

dynamic-peers

IN THIS SECTION

- Syntax | 6111
- Hierarchy Level | 6111
- Description | 6111
- Required Privilege Level | 6111
- Release Information | 6112

Syntax

```
dynamic-peers;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit protocols rip group group-name neighbor neighbor-name]
```

Description

Configure an interface to have dynamic peers in a point-to-multipoint RIP network.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring Point-to-Multipoint RIP Networks

dynamic-policy-control

IN THIS SECTION

- Syntax | 6112
- Hierarchy Level | 6113
- Description | 6113
- Required Privilege Level | 6113
- Release Information | 6113

Syntax

```
dynamic-policy-control {
    pcc-rules {
        [rule-name  number];
    }
    pcc-rulebases {
        [rulebase-name];
    }
    diameter-profile gx-profile-name;
}
```

Hierarchy Level

```
[edit unified-edge pcef profiles profile-name]
```

Description

Configure the dynamic policy control for the PCC rules, PCC rulebases, or both in a PCEF profile. You can configure a maximum of 32 PCC rules in a PCEF profile. There is no limit to the number of PCC rulebases you can configure in a PCEF profile.

> **NOTE**: If you configure the `dynamic-policy-control` statement for a PCEF profile, you cannot configure the `static-policy-control` statement in the same profile.

The remaining statements are explained separately.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Dynamic Policies

Understanding How Subscriber-Aware Policy and Charging Control Rules Are Provisioned Dynamically by a PCRF

dynamic-pool

IN THIS SECTION

- Syntax | 6114
- Hierarchy Level | 6115
- Description | 6115
- Options | 6115
- Required Privilege Level | 6115
- Release Information | 6115

Syntax

```
dynamic-pool <dynamic-pool>{
    family {
        inet6 {
            from-interface <interface>;
            delegated-prefix-length <network-prefix-length>;
            range <range-name> {
                masked-low <masked-low>;
                masked-high <masked-high>;
                prefix-length <prefix-length>;
            }
            dhcp-attributes {
                dns-server <address>;
                t1-percentage <t1-percentage>;
                t2-percentage <t2-percentage>;
                preferred-lifetime <preferred-lifetime>;
                valid-lifetime <valid-lifetime>;
            }
        }
    }
}
```

Hierarchy Level

```
[edit access address-assignment (Access)]
```

Description

Configure the dynamic pool updated by the client running on the WAN interface.

Options

The remaining statements are explained separately.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

address-assignment (Access) | 1858

IP Address Assignment Pool

dynamic-profile (Demux)

IN THIS SECTION

- Syntax | 6116
- Hierarchy Level | 6116
- Description | 6116

- Options | 6116
- Required Privilege Level | 6117
- Release Information | 6117

Syntax

```
dynamic-profile profile-name {
    network ip-address {
        range name {
            low lower-limit;
            high upper-limit;
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet auto-configure address-ranges]
[edit interfaces interface-name unit unit-number demux inet6 auto-configure address-ranges]
```

Description

Assign a dynamic profile and specify address options for the demultiplexing (demux) interface options.

Options

profile-name	Name of the dynamic profile for the demultiplexing (demux) interface options.
network *ip-address*	Configure an IPv4 or IPv6 address for a dynamic profile for the demultiplexing (demux) interface options.
range <range-name>	Configure an IP name range used within an address-assignment pool for the demultiplexing (demux) interface options.

- `low` *lower-limit*—Lower limit of IPv4 or IPv6 address range.
- `high` *upper-limit*—Upper limit of IPv4 or IPv6 address range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles

dynamic-profile (DHCP Local Server)

IN THIS SECTION

- Syntax | 6118
- Hierarchy Level | 6118
- Description | 6118
- Options | 6118
- Required Privilege Level | 6118
- Release Information | 6119

Syntax

```
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server group group-name interface interface-name],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Specify the dynamic profile that is attached to all interfaces, a named group of interfaces, or a specific interface.

Options

profile-name—Name of the dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Options `aggregate-clients` and `use-primary` introduced in Junos OS Release 9.3.

Support at the `[edit ... interface]` hierarchy levels introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

Configuring a Default Subscriber Service

dynamic-profile (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 6119
- Hierarchy Level | 6120
- Description | 6120
- Options | 6120
- Required Privilege Level | 6120
- Release Information | 6120

Syntax

```
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify the dynamic profile that is attached to all interfaces, to a named group of interfaces, or to a specific interface.

M120 and M320 routers do not support DHCPv6.

Options

profile-name—Name of the dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

dhcp-relay | 5620

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

Grouping Interfaces with Common DHCP Configurations

Configuring a Default Subscriber Service

dynamic-profile (Dynamic Access-Line-Identifier VLANs)

IN THIS SECTION

- Syntax | 6121
- Hierarchy Level | 6121
- Description | 6121
- Options | 6122
- Required Privilege Level | 6122
- Release Information | 6122

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces “$junos-interface-ifd-name” unit “$junos-
interface-unit”  auto-configure line-identity],
[edit interfaces interface-name unit logical-unit-number auto-configure line-identity]
```

Description

Attach a dynamic profile to a static or dynamic underlying VLAN interface to create a dynamic VLAN on the interface, based on receiving a trusted option in the DHCP or PPPoE discovery packet. The trusted

option can be the ACI, ARI, both ACI and ARI, or neither. The VLAN is known as an access-line-identifier VLAN.

Options

profile-name	Name of the dynamic profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 17.1.

RELATED DOCUMENTATION

Configuring Dynamic Underlying VLAN Interfaces to Use Access-Line Identifiers

Configuring Dynamic VLAN Subscriber Interfaces Based on Access-Line Identifiers

dynamic-profile (Dynamic ACI VLANs)

IN THIS SECTION

- Syntax | 6123
- Hierarchy Level | 6123
- Description | 6123
- Options | 6123
- Required Privilege Level | 6123
- Release Information | 6123

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces “$junos-interface-ifd-name” unit “$junos-
interface-unit” auto-configure agent-circuit-identifier],
[edit interfaces interface-name unit logical-unit-number auto-configure agent-circuit-identifier]
```

Description

Attach a dynamic profile for an agent circuit identifier (ACI) interface set to a static or dynamic underlying VLAN interface.

Options

- *profile-name*—Name of the dynamic profile that defines the ACI interface set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

> *Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*
>
> *Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*

dynamic-profile (Dynamic PPPoE)

IN THIS SECTION

- Syntax | 6124
- Hierarchy Level | 6124
- Description | 6124
- Options | 6125
- Required Privilege Level | 6125
- Release Information | 6125

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
pppoe],
[edit interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
pppoe-underlying-options]
```

Description

Attach a PPPoE dynamic profile to an underlying Ethernet interface. This underlying interface is configured with either the `encapsulation ppp-over-ether` statement or the `family pppoe` statement; the two statements are mutually exclusive. When the router creates a dynamic PPPoE logical interface on the underlying interface, it uses the information in the dynamic profile to determine the properties of the dynamic PPPoE logical interface.

NOTE: The `[edit ... family pppoe]` hierarchies are supported only on MX Series routers with MPCs.

Starting in Junos OS Release 17.2R1, you can configure converged services for MS-MPCs and MS-MICs. You can configure captive portal content delivery (CPCD) profiles for MS-MICs and MS-MPCs by including the service interface ms-fpc/pic/port statement at the `edit service-set` *`service set name`* `captive-portal-content-delivery-profile` *`profile name`* `interface-service` heirarchy level.

Options

`profile-name`—Name of a previously configured PPPoE dynamic profile, up to 64 characters in length, defined at the `[edit dynamic-profiles` *`profile-name`* `interfaces pp0]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

Support for the `[edit ... family pppoe]` hierarchies introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces

Configuring the PPPoE Family for an Underlying Interface

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

dynamic-profile (L2TP)

IN THIS SECTION

- Syntax | 6126
- Hierarchy Level | 6126
- Description | 6126
- Options | 6126
- Required Privilege Level | 6126
- Release Information | 6127

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Assign a dynamic profile to the tunnel group for dynamic LNS sessions.

Options

profile-name	Name of the dynamic profile for the tunnel group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

dynamic-profile-options

IN THIS SECTION

- Syntax | 6127
- Hierarchy Level | 6127
- Description | 6127
- Required Privilege Level | 6128
- Release Information | 6128

Syntax

```
dynamic-profile-options {
    versioning;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure global dynamic profile options.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Versioning for Dynamic Profiles

dynamic-profile (PPP)

IN THIS SECTION

- Syntax | 6128
- Hierarchy Level | 6129
- Description | 6129
- Required Privilege Level | 6129
- Release Information | 6129

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number ppp-options]
```

Description

Specify the dynamic profile that is attached to the interface. On the MX Series routers, this statement is supported on PPPoE interfaces only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support for MLPPP on LSQ interfaces introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

- *Dynamic Profiles Overview*
- *Configuring a Basic Dynamic Profile*
- *Attaching Dynamic Profiles to Static PPP Subscriber Interfaces*
- *Attaching Dynamic Profiles to MLPPP Bundles*
- *Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces*

dynamic-profile (PPPoE Service Name Tables)

IN THIS SECTION

Syntax | 6130

- Hierarchy Level | 6130
- Description | 6130
- Options | 6130
- Required Privilege Level | 6131
- Release Information | 6131

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name],
[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci
circuit-id-string ari remote-id-string]
```

Description

Specify a dynamic profile to instantiate a dynamic PPPoE interface. You can associate a dynamic profile with a named service entry, `empty` service entry, or `any` service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The dynamic profile associated with a service entry in a PPPoE service name table overrides the dynamic profile associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the `dynamic-profile` statement at the `[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci circuit-id-string ari remote-id-string]` hierarchy level, you cannot also include the `static-interface` statement at this level. The `dynamic-profile` and `static-interface` statements are mutually exclusive for ACI/ARI pair configurations.

Options

profile-name—Name of the dynamic profile that the router uses to instantiate a dynamic PPPoE interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation

dynamic-profile (Stacked VLAN)

IN THIS SECTION

- Syntax | 6131
- Hierarchy Level | 6132
- Description | 6132
- Options | 6132
- Required Privilege Level | 6132
- Release Information | 6132

Syntax

```
dynamic-profile profile-name {
    accept  (any | dhcp-v4 |dhcp-v6| inet | inet6 | pppoe);
    access-profile vlan-dynamic-profile-name;
```

```
    ranges (any | low-tag-high-tag),(any | low-tag-high-tag);
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure stacked-vlan-ranges]
```

Description

Configure a dynamic profile for use when configuring dynamic stacked VLANs.

Options

profile-name—Name of the dynamic profile that you want to use when configuring dynamic stacked VLANs.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Dynamic Profiles Overview

Configuring a Basic Dynamic Profile

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

dynamic-profile (VLAN)

IN THIS SECTION

- Syntax | 6133
- Hierarchy Level | 6133
- Description | 6133
- Options | 6133
- Required Privilege Level | 6134
- Release Information | 6134

Syntax

```
dynamic-profile profile-name {
    accept  (any | dhcp-v4 |dhcp-v6| inet | inet6 | pppoe);
    accept-out-of-band protocol;
    access-profile vlan-dynamic-profile-name;
    ranges (any | low-tag)-(any | high-tag);
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges]
```

Description

Configure a dynamic profile for use when configuring dynamic VLANs.

Options

profile-name—Name of the dynamic profile that you want to use when configuring dynamic VLANs.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Dynamic Profiles Overview

Configuring a Basic Dynamic Profile

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

dynamic-profiles

IN THIS SECTION

- Syntax | 6134
- Hierarchy Level | 6146
- Description | 6146
- Options | 6146
- Required Privilege Level | 6147
- Release Information | 6147

Syntax

```
dynamic-profiles {
    profile-name {
        class-of-service {
            dynamic-class-of-service-options {
```

```
                vendor-specific-tags tag;
            }
            interfaces {
                interface-name ;
                }
                unit logical-unit-number {
                    classifiers {
                        type (classifier-name | default);
                    }
                    output-traffic-control-profile (profile-name | $junos-cos-traffic-control-
profile);
                    report-ingress-shaping-rate bps;
                    rewrite-rules {
                        dscp (rewrite-name | default);
                        dscp-ipv6 (rewrite-name | default);
                        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                        inet-precedence (rewrite-name | default);
                        }
                    }
                }
            }
            scheduler-maps {
                map-name {
                    forwarding-class class-name scheduler scheduler-name;
                }
            }
            schedulers {
                (scheduler-name) {
                    buffer-size (seconds | percent percentage | remainder | temporal
microseconds);
                    drop-profile-map loss-priority (any | low | medium-low | medium-high | high)
protocol (any | non-tcp | tcp) drop-profile profile-name;
                    excess-priority (low | high | $junos-cos-scheduler-excess-priority);
                    excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
                    overhead-accounting (shaping-mode) <bytes (byte-value>;
                    priority priority-level;
                    shaping-rate (rate | predefined-variable);
                    transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
                }
            }
            traffic-control-profiles profile-name {
                adjust-minimum rate;
                delay-buffer-rate (percent percentage | rate);
```

```
                excess-rate (percent percentage | proportion value | percent $junos-cos-excess-
rate);
                excess-rate-high (percent percentage | proportion value);
                excess-rate-low (percent percentage | proportion value);
                guaranteed-rate (percent percentage | rate) <burst-size  bytes>;
                max-burst-size cells;
                overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
                peak-rate rate;
                scheduler-map map-name;
                shaping-rate (percent percentage | rate | predefined-variable) <burst-size
bytes>;
                shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
                shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
                shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
                shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
                shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
                shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
                shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
                shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
                shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
                sustained-rate rate;
            }
        }
        firewall {
            family family {
                fast-update-filter filter-name {
                    interface-specific;
                    match-order [match-order];
                    term term-name {
                        from {
                            match-conditions;
                        }
                        then {
                            action;
                            action-modifiers;
                        }
                        only-at-create;
                    }
                }
                filter filter-name {
                    enhanced-mode-override;
                    instance-shared;
                    interface-shared;
```

```
            interface-specific;
                term term-name {
                    from {
                        match-conditions;
                    }
                    then {
                        action;
                        action-modifiers;
                    }
                only-at-create;
            filter filter-name {
            interface-specific;
                term term-name {
                    from {
                        match-conditions;
                    }
                    then {
                        action;
                        action-modifiers;
                    }
                }
        hierarchical-policer uid {
            aggregate {
                if-exceeding {
                    bandwidth-limit-limit bps;
                    burst-size-limit bytes;
                }
                then {
                    policer-action;
                }
            }
            premium {
                if-exceeding {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
                then {
                    policer-action;
                }
            }
        }
        policer uid {
            filter-specific;
```

```
                if-exceeding {
                    (bandwidth-limit bps | bandwidth-percent percentage);
                    burst-size-limit bytes;
                }
                logical-bandwidth-policer;
                logical-interface-policer;
                physical-interface-policer;
                then {
                    policer-action;
                }
            }
            three-color-policer uid {
                action {
                    loss-priority high then discard;
                }
                logical-interface-policer;
                single-rate {
                    (color-aware | color-blind);
                    committed-burst-size bytes;
                    committed-information-rate bps;
                    excess-burst-size bytes;
                }
                two-rate {
                    (color-aware | color-blind);
                    committed-burst-size bytes;
                    committed-information-rate bps;
                    peak-burst-size bytes;
                    peak-information-rate bps;
                    }
                }
            }
        }
        interfaces interface-name {
            interface-set interface-set-name {
                interface interface-name {
                    unit logical unit number {
                        advisory-options {
                            downstream-rate rate;
                            upstream-rate rate;
                        }
                    }
                }
            }
```

```
            unit logical-unit-number {
                actual-transit-statistics;
                auto-configure {
                    agent-circuit-identifier {
                        dynamic-profile profile-name;
                    }
                    line-identity {
                        include {
                            accept-no-ids;
                            circuit-id;
                            remote-id;

                        }
                        dynamic-profile profile-name;
                    }
                }
                encapsulation  (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-
vc-mux | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap |
atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-
ppp | ethernet | frame-relay-ccc | frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type |
frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether |
ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls);
                family family {
                    address address;
                    filter {
                        adf {
                            counter;
                            input-precedence precedence;
                            not-mandatory;
                            output-precedence precedence;
                            rule rule-value;
                        }
                        input filter-name (
                            precedence precedence;
                            shared-name filter-shared-name;
                        }
                        output filter-name {
                            precedence precedence;
                            shared-name filter-shared-name;
                        }
                    }
                    rpf-check {
                        fail-filter filter-name;
```

```
                mode loose;
            }
            service {
                 input {
                     service-set service-set-name {
                         service-filter filter-name;
                     }
                     post-service-filter filter-name;
                }
                input-vlan-map {
                     inner-tag-protocol-id tpid;
                     inner-vlan-id number;
                     (push | swap);
                     tag-protocol-id tpid;
                     vlan-id number;
                }
                output {
                     service-set service-set-name {
                         service-filter filter-name;
                     }
                }
                output-vlan-map {
                     inner-tag-protocol-id tpid;
                     inner-vlan-id number;
                     (pop | swap);
                     tag-protocol-id tpid;
                     vlan-id number;
                }
                pcef  pcef-profile-name {
                     activate rule-name | activate-all;
                }
            }
            unnumbered-address interface-name <preferred-source-address address>;
        }
        filter {
            input filter-name (
                shared-name filter-shared-name;
            }
            output filter-name {
                shared-name filter-shared-name;
            }
        }
        host-prefix-only;
```

```
ppp-options {
    aaa-options aaa-options-name;
    authentication [ authentication-protocols ];
    chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
    }
    ignore-magic-number-mismatch;
    initiate-ncp (dual-stack-passive | ipv6 | ip)
    ipcp-suggest-dns-option;
    mru size;
    mtu (size | use-lower-layer);
    on-demand-ip-address;
    pap;
    peer-ip-address-optional;
    local-authentication {
        password password;
        username-include {
            circuit-id;
            delimiter character;
            domain-name name;
            mac-address;
            remote-id;
        }
    }
}
reassemble-packets;
targeted-options {
    backup backup;
    group group;
    primary primary;
    weight ($junos-interface-target-weight | weight-value);
}
telemetry {
    subscriber-statistics;
    queue-statistics {
        interface $junos-interface-name {
            refresh rate;
            queues queue set;
        }
        interface-set $junos-interface-set-name {
            refresh rate;
            queues queue set;
```

```
                    }
                }
            }
            vlan-id number;
            vlan-tags outer  [tpid].vlan-id [inner [tpid].vlan-id];
        }
    }
    interfaces {
        demux0 {...}
    }
    interfaces {
        pp0 {...}
    }
    policy-options {
        prefix-list uid {
            ip-addresses;
            dynamic-db;
        }
    }
    predefined-variable-defaults predefined-variable <variable-option> default-value;
    profile-type remote-device-service;
    protocols {
        igmp {
            interface interface-name {
                accounting;
                disable;
                group-limit limit;
                group-policy;
                group-threshold value;
                immediate-leave
                log-interval seconds;
                no-accounting;
                oif-map;
                passive;
                promiscuous-mode;
                ssm-map ssm-map-name;
                ssm-map-policy ssm-map-policy-name
                static {
                    group group {
                        source source;
                    }
                }
                version version;
```

```
            }
        }
        mld {
             interface interface-name {
                (accounting | no-accounting);
                disable;
                group-limit limit;
                group-policy;
                group-threshold value;
                immediate-leave;
                log-interval seconds;
                oif-map;
                passive;
                ssm-map ssm-map-name;
                ssm-map-policy ssm-map-policy-name;
                static {
                    group multicast-group-address {
                        exclude;
                        group-count number;
                        group-increment increment;
                        source ip-address {
                            source-count number;
                            source-increment increment;
                        }
                    }
                }
                version version;
            }
        }
        router-advertisement {
            interface interface-name {
                 current-hop-limit number;
                default-lifetime seconds;
                dns-server-address
                (managed-configuration | no-managed-configuration);
                max-advertisement-interval seconds;
                min-advertisement-interval seconds;
                (other-stateful-configuration | no-other-stateful-configuration);
                prefixprefix {
                    (autonomous | no-autonomous);
                    (on-link | no-on-link);
                    preferred-lifetime seconds;
                    valid-lifetime seconds;
```

```
                }
                reachable-time milliseconds;
                retransmit-timer milliseconds;
            }
        }
    }
    routing-instances routing-instance-name {
        interface interface-name;
        routing-options {
            access {
                route prefix {
                    next-hop next-hop;
                    metric route-cost;
                    preference route-distance;
                    tag route-tag;
                    tag2 route-tag2;
                }
            }
            access-internal {
                route subscriber-ip-address {
                    qualified-next-hop underlying-interface {
                        mac-address address;
                    }
                }
            }
            multicast {
                interface interface-name {
                    no-qos-adjust;
                }
            }
        }
        rib routing-table-name {
            access {
                route prefix {
                    next-hop next-hop;
                    metric route-cost;
                    preference route-distance;
                    tag route-tag;
                    tag2 route-tag2;
                }
            }
            access-internal {
                route subscriber-ip-address {
```

```
                    qualified-next-hop underlying-interface {
                        mac-address address;
                    }
                }
            }
        }
    }
    routing-options {
        access {
            route prefix {
                next-hop next-hop;
                metric route-cost;
                preference route-distance;
                tag route-tag;
                tag2 route-tag2;
            }
        }
        access-internal {
            route subscriber-ip-address {
                qualified-next-hop underlying-interface {
                    mac-address address;
                }
            }
        }
        multicast {
            interface interface-name {
                no-qos-adjust;
            }
        }
    }
    services {
        captive-portal-content-delivery {
            auto-deactivate value;
            rule name {
                match-direction (input | input-output | output);
                term name {
                    then {
                        accept;
                        redirect url;
                        rewrite destination-address address <destination-port port-number>;
                        syslog;
                    }
                }
```

```
                }
            }
        }
        variables {
            variable-name {
                default-value default-value;
                equals expression;
                mandatory;
                uid;
                uid-reference;
            }
        }
        version-alias profile-alias-string;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Create dynamic profiles for use with DHCP or PPP client access.

Options

profile-name Name of the dynamic profile; string of up to 80 alphanumeric characters.

reassemble-packets (Optional) Enables IPv4 reassembly of fragmented GRE packets conveyed across a soft GRE tunnel from a Wi-Fi access point to a Wi-Fi access gateway on a BNG. Reassembly is supported for fragments that range in size from 256 bytes through 8192 bytes.

> **NOTE**:
>
> - The maximum reassembled packet size is 13,310 bytes; this requires an MTU of 1500 bytes. The router drops reassembled packets that

are larger than 13,310 bytes. The router also drops DHCP discover packets that are smaller than the MTU.

- Ordering is not maintained between fragmented packets and non-fragmented packets.
- The WAG does not support soft GRE packets with keys. Fragmented packets GRE with key are not reassembled.
- Soft GRE packet reassembly is not supported for pseudowires over redundant logical tunnels (RLT).
- The order of the last arriving fragment is not guaranteed when the reassembled packets are forwarded.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the `filter`, `policer`, `hierarchical-policer`, `three-color-policer`, and `policy options` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring a Basic Dynamic Profile

Configuring Dynamic VLANs Based on Agent Circuit Identifier Information

Dynamic Profiles for Subscriber Management

dynamic-requests

IN THIS SECTION

- Syntax | 6148
- Hierarchy Level | 6148
- Description | 6148
- Options | 6149
- Required Privilege Level | 6149
- Release Information | 6149

Syntax

```
dynamic-requests {
    routing-instance routing-instance-name;
    source-address source-address;
    source-port source-port;
}
```

Hierarchy Level

```
[edit access radsec  destination ]
```

Description

Configure RADSEC clients to receive and process dynamic requests. RADSEC servers can be a source of dynamic RADIUS requests such as Change of Authorization (CoA) and RADIUS-initiated disconnect (RID) messages.

You must configure the IP address of the RADSEC server as the source address for the requests. You can also specify the source port that the client monitors for dynamic requests. If the port is not explicitly configured, the default RADSEC port 2083 is used.

Options

routing-instance *routing-instance-name*	Specify the routing instance name.
source-address *ip-address*	(Required) Configure the source IP address, which is the IP address of the RADSEC server. If the source address is not configured, dynamic requests will be rejected.
source-port *port-number*	(Optional) Specify the source port that the client monitors for dynamic requests. If the port is not explicitly configured, the default is used. • **Default:** 2083 • **Range:** 1 through 65535

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

RADIUS over TLS (RADSEC)

dynamic-requests-secret (RADIUS Server)

IN THIS SECTION

- Syntax | 6150
- Hierarchy Level | 6150
- Description | 6150

- Default | 6150
- Options | 6150
- Required Privilege Level | 6150
- Release Information | 6150

Syntax

```
dynamic-requests-secret password;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure the secret password to be used for change of authorization requests from the RADIUS server.

Default

Use the same password that is used for authentication requests.

Options

password	Password for dynamic requests.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

dynamic (Security)

IN THIS SECTION

- Syntax | 6151
- Hierarchy Level | 6152
- Description | 6152
- Options | 6152
- Required Privilege Level | 6153
- Release Information | 6153

Syntax

```
dynamic {
    connections-limit number;
    distinguished-name {
        container container-string;
        wildcard wildcard-string
    }
    general-ikeid;
    hostname domain-name;
    ike-user-type (group-ike-id | shared-ike-id);
    inet ip-address;
    inet6 ipv6-address;
    reject-duplicate-connection;
    user-at-hostname e-mail-address;
}
```

Hierarchy Level

```
[edit security ike gateway gateway-name]
```

Description

Specify the identifier for the remote gateway with a dynamic IPv4 or IPv6 address. Use this statement to set up a VPN with a gateway that has an unspecified IPv4 or IPv6 address.

Options

connections-limit Configure the number of concurrent connections that the group profile supports. When the maximum number of connections is reached, no more dynamic virtual private network (VPN) endpoints dialup users attempting to access an IPsec VPN are allowed to begin Internet Key Exchange (IKE) negotiations. This configuration applies to SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100, SRX4200, and SRX4600 devices and vSRX instances, and to SRX5400, SRX5600, and SRX5800 devices configured for AutoVPN.

distinguished-name Specify a distinguished name as the identifier for the remote gateway with a dynamic IP address.

general-ikeid Disables IKE ID validation. If this option is enabled, the new iked process skips the IKE ID validation. After skipping the IKE ID validation, the new iked process still continues the authentication as per the IKE standard. `general-ikeid` is an optional configuration statement.

hostname Name by which a network-attached device is known on a network. A fully qualified domain name (FQDN), or partial FQDN that can be matched to a peer's X.509 PKI certificate. A partial FQDN is matched to the right-most part of the alternate subject field in the peer device's certificate. For example, the partial FQDN example.net can match devices with host1.example.net or host2.example.net in the alternate subject field of their certificates. Note that the partial FQDN example.net does not match host1.example.network.com or host2.net.com because example.net is not the right-most value in the alternate subject field. For AutoVPN, a partial FQDN combined with ike-user-type group-ike-id can be used to identify a specific remote user or peer when there are multiple peers that share a common domain name.

ike-user-type Configure the type of IKE user for a remote access connection.

- Values:
  - group-ike-id—E-mail address or fully qualified domain name (FQDN) shared by a group of remote access users so that each user does not need to configure a separate IKE profile. When group IKE IDs are configured, the IKE ID of each user is a concatenation of a user-specific part and a part that is common to all group IKE ID users. For example, the user Bob might use ”Bob.example.net“ as his full IKE ID, where ”.example.net“ is common to all users. The full IKE ID is used to uniquely identify each user connection. Group IKE IDs require the generation of a unique preshared key based on the username supplied during VPN connection, which can be viewed with the `show security ike pre-shared-key` command.
  - shared-ike-id—E-mail address shared by a large number of remote access users so that each user does not need to configure a separate IKE profile. When a shared IKE ID is configured, all users share a single IKE ID and a single IKE preshared key. Each user is authenticated through the mandatory XAuth phase, where the credentials of individual users are verified either with an external RADIUS server or with a local access database. XAuth is required for shared IKE IDs.

inet Use an IPV4 address to identify the dynamic peer.

inet6 Use an IPV6 address to identify the dynamic peer.

reject-duplicate-connection Reject new connection from duplicate IKE-id.

user-at-hostname Use an e-mail address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5. Support for the `inet6` option added in Junos OS Release 11.1.

`general-ikeid` option under `[edit security ike gateway gateway-name dynamic]` hierarchy is introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

IPsec Overview

Example: Configuring AutoVPN with Pre-Shared Key

dynamic-server

IN THIS SECTION

- Syntax | 6154
- Hierarchy Level | 6155
- Description | 6155
- Options | 6155
- Required Privilege Level | 6155
- Release Information | 6155

Syntax

```
dynamic-server {
        group group-name {
            neighbor-discovery-router-advertisement <ndra-pool>;
            interface interface-name {
                overrides {
                    delegated-pool <delegated-pool>;
                    ia-na-pool <ia-na-pool>;
                    interface-client-limit interface-client-limit;
                    process-inform {
                    pool pool-name;
                    }
                    rapid-commit;
                }
            }
        }
    }
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server dhcpv6],
[edit logical-systems name system services dhcp-local-server dhcpv6],
[edit logical-systems name tenants name routing-instances name system services dhcp-local-server
dhcpv6],
[edit routing-instances name system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6],
[edit tenants name routing-instances name system services dhcp-local-server dhcpv6]
```

Description

Configure the server running on a LAN interface.

Options

group *group-name*	Name of the group.
neighbor-discovery-router-advertisement *ndra-pool*	Name of the address-assignment pool used to assign the router advertisement prefix.
interface *interface-name*	Name of the interface.
ia-na-pool *ia-na-pool*	Identity association for non-temporary address (IA_NA) pool name for inet6.
interface-client-limit *interface-client-limit*	Limit the number of clients allowed on an interface. **Range**: 1 through 500000

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

dhcp-local-server (System Services) | 5613

dhcp-client | 5596

dynamic (Source Packet Routing)

IN THIS SECTION

- Syntax | 6156
- Hierarchy Level | 6156
- Description | 6157
- Options | 6157
- Required Privilege Level | 6157
- Release Information | 6157

Syntax

```
dynamic {
    protected {
        mandatory;
    }
    unprotected {
        mandatory;
    }
}
```

Hierarchy Level

```
[edit protocols source-packet-routing segment-list]
```

Description

Enable dynamic computation of segment routing label-switched paths (LSPs) based on tunnel destination and translation service to fetch the corresponding segment identifiers (SIDs).

NOTE:

When the dynamic statement is enabled, all the next hops must have an IP address assigned as a minimum configuration. In the case of segment-lists, if a next hop has both IP address and label configured, then the configured label is retained.

NOTE:

- Because translation service use IGP instance of traffic-engineered database (TED), you must include the `igp-topology` statement at the `[edit protocols isis traffic-engineering]` hierarchy level for successful translation.
- The `auto-translation` and `dynamic` statements are mutually exclusive, and only either of them can be configured under a segment-list.

Options

protected (Optional) Ensures a protected adjacency SID is used for the links of the LSP.

unprotected (Optional) Ensures an unprotected adjacency SID is used for the links of the LSP.

mandatory (Optional) Enabled translation to fail if the specified type of adjacency SID cannot be found for a link. This option does not have effect on node SIDs.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Understanding Static Segment Routing LSP in MPLS Networks

dynamic-tunnel-gre-key

IN THIS SECTION

- Syntax | 6158
- Hierarchy Level | 6158
- Description | 6158
- Options | 6159
- Required Privilege Level | 6159
- Release Information | 6159

Syntax

```
dynamic-tunnel-gre-key dynamic-tunnel-gre-key;
```

Hierarchy Level

```
[edit routing-options dynamic-tunnels tunnel-attributes]
```

Description

Enable traffic flow identification by configuring a GRE key that will be used to identify an individual traffic flow in a GRE tunnel. RFC2890 describes the enhancements to the GRE header that includes the key field and its use in identifying individual traffic flow in a tunnel. The encapsulating router uses the GRE key to encapsulate the packets for a particular traffic flow. The decapsulating router uses the GRE key to identify and decapsulate those packets for that traffic flow.

The key must be configured on the routers on both endpoints of a tunnel and a policy must be created and exported to populate the key in the forwarding table.

Options

dynamic-tunnel-gre-key An identifier for a individual traffic flow within a GRE tunnel.

- **Range:** 1 through 4,294,967,295.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.1R1

RELATED DOCUMENTATION

dynamic-tunnels | 6161

tunnel-attributes

show dynamic-tunnels database

dynamic-tunnel-reassembly

IN THIS SECTION

- Syntax | 6160
- Hierarchy Level | 6160
- Description | 6160
- Options | 6160
- Required Privilege Level | 6160
- Release Information | 6160

Syntax

```
dynamic-tunnel-reassembly (off | on);
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options dynamic-tunnels tunnel-attributes],
[edit logical-systems name routing-options dynamic-tunnels tunnel-attributes],
[edit logical-systems name tenants name routing-instances name routing-options dynamic-tunnels
tunnel-attributes],
[edit routing-instances name routing-options dynamic-tunnels tunnel-attributes],
[edit routing-options dynamic-tunnels tunnel-attributes],
[edit tenants name routing-instances name routing-options dynamic-tunnels tunnel-attributes]
```

Description

Use this statement to enable or disable IPV6 fragment reassembly for forwarding Ipv4 traffic. You can configure multiple tunnels and specify different attributes for each tunnel. Service providers with IPv6 infrastructure can configure individual tunnels for each customer to route IPv4 traffic.

Options

off Disable reassembly check

on Enable reassembly check. By default, this is disabled.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

tunnel-attributes | **30194**

dynamic-tunnels

IN THIS SECTION

- Syntax | 6161
- Hierarchy Level | 6162
- Description | 6162
- Options | 6162
- Required Privilege Level | 6163
- Release Information | 6163

Syntax

```
dynamic-tunnels tunnel-name {
    bgp-signal;
    destination-networks prefix;
    gre;
    ipip
    rsvp-te entry-name {
        destination-networks network-prefix;
        label-switched-path-template (Multicast) {
            default-template;
            template-name;
        }
    }
    source-address address;
    spring-te;
    traceoptions;
    tunnel-attributes name {
        dynamic-tunnel-anchor-pfe dynamic-tunnel-anchor-pfe;
        dynamic-tunnel-anti-spoof (off | on);
        dynamic-tunnel-gre-key
        dynamic-tunnel-mtu dynamic-tunnel-mtu;
        dynamic-tunnel-source-prefix dynamic-tunnel-source-prefix;
        dynamic-tunnel-type (BGP-SIGNAL | GRE | UDP | V4oV6);
    }
```

```
  v4ov6 ipv6-anycast-source-duplication;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Configure a dynamic tunnel between two PE routers.

> **NOTE**: ACX Series routers do not support the `gre` statement.

Configure dynamic IPv4-over-IPv6 tunnels and define their attributes to forward IPv4 traffic over an IPv6-only network. IPv4 traffic is tunneled from customer premises equipment to IPv4-over-IPv6 gateways. You must also configure `extended-nexthop` option at `[edit protocols bgp family inet unicast]` hierarchy level to allow BGP to route IPv4 address families over an IPv6 session.

Options

bgp-signal Enable the creation of a tunnel signaled by BGP.

gre Enable dynamic generic routing encapsulation type tunnel mode for IPv4.

- Values:
  - `next-hop-based-tunnel`—Enable next hop base dynamic-tunnel for steering IPv4 traffic with IPv6 next hop address.

ipip Enable dynamic IP in IP encapsulation type tunnel mode for IPv4

- Values:
  - `full-resolved-next-hop-based-tunnel`—Enable fully resolved next hop base dynamic-tunnel for steering IPv4 traffic with IPv6 next hop address.

source-address	Specify the source address of the tunnel.
tunnel-name	Name of the dynamic tunnel.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`bgp-signal` option introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

- extended-nexthop
- tunnel-attributes
- Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks
- Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP

dynamic-vpn

IN THIS SECTION

- Syntax | 6164
- Hierarchy Level | 6164
- Description | 6164
- Options | 6165
- Required Privilege Level | 6165
- Release Information | 6165

Syntax

```
dynamic-vpn {
    access-profile profile-name;
    clients configuration-name {
        ipsec-vpn vpn-name;
        remote-exceptions ip-address/mask;
        remote-protected-resources ip-address/mask;
        user username;
        user-groups user-group-name;
    }
    config-check;
    force-upgrade;
    interface;
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag {
            all;
        }
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure the dynamic VPN feature. The dynamic VPN feature simplifies remote access by enabling users to create IPsec VPN tunnels without having to manually configure settings on their PCs or laptops. This feature is supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

In Junos OS Release 21.4R1, we’ve deprecated the `dynamic-vpn` configuration statement and we might remove it completely in a future release.

Options

access-profile Specify the access profile to use for Extended Authentication for remote users trying to download the Access Manager. This feature is supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

clients Create a client configuration for the dynamic VPN feature.

config-check Enable extra dynamic VPN configuration checking. If you include this statement in your configuration, it is automatically enabled. If the statement is not present in your configuration, the configuration check option is not enabled. This feature is supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

force-upgrade Force upgrade.the dynamic vpn.

interface Specify a list of interfaces to set the interfaces that allow access to dynamic VPN, separated by spaces. This feature is supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

traceoptions Configure dynamic VPN tracing options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release Release 9.5.

`config-check` and `interface` options introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Dynamic VPN Overview

E

IN THIS SECTION

- enhanced-scale (IRB Interface) | **6175**
- e1-options (BITS Interfaces Signal Type) | **6176**
- e1-options (Clock Synchronization) | **6178**
- e2e-transparent | **6179**
- eapol-address (MACSec) | **6181**
- eapol-block | **6183**
- early-ageout | **6185**
- east-interface | **6186**
- echo-req | **6188**
- ecmp | **6190**
- ecmp-alb | **6191**
- ecmp-fast-reroute | **6193**
- ecmp-local-bias | **6195**
- ecmp-resilient-hash | **6196**
- edge | **6197**
- edge-virtual-bridging | **6200**
- edit | **6201**
- effective-aigp | **6203**
- effective-shaping-rate | **6204**
- egress | **6205**
- egress (Analyzer) | **6207**
- egress (Buffer Configuration) | **6208**
- egress-chaining | **6210**
- egress-l2-extended-match | **6212**
- egress-policer-overhead | **6213**
- egress-policy | **6215**
- egress-protection (BGP) | **6216**
- egress-protection (Layer 2 circuit) | **6218**

- egress-protection (MPLS) | **6220**
- egress-protection (Routing Instances) | **6221**
- egress-shaping-overhead | **6223**
- egress-stats (Flat-File Accounting Options) | **6225**
- egress-te | **6227**
- egress-te-adj-segment | **6229**
- egress-te-backup-paths | **6232**
- egress-te-node-segment | **6234**
- egress-te-set-segment | **6236**
- egress-quantization | **6238**
- ei-mapping-timeout | **6241**
- ei-mapping-timeout (Source NAT Next Gen Services) | **6242**
- eif-flow-limit | **6244**
- em-hw-profile | **6245**
- email-notify | **6247**
- embedded-rp | **6248**
- enable | **6249**
- enable | **6251**
- enable-all-ifl | **6252**
- enable-all-qmodules | **6253**
- enable-asymmetric-traffic-processing (Service Set Next Gen Services) | **6254**
- enable-call-distribution | **6255**
- enable-descriptive-session-syslog | **6256**
- enable (Enhanced Subscriber Management) | **6257**
- enable-explorer (REST API) | **6259**
- enable-extended-ql-tlv | **6260**
- enable-flow-tracing (Services) | **6263**
- enable-gtpu-distribution | **6264**
- enable-ipv6-services-for-lac (L2TP) | **6265**
- enable-mac-move-action | **6267**
- enable-packet-pool | **6268**
- enable-performance-mode | **6269**

- enable-performance-mode (Application Identification) | **6271**
- enable-rejoin (Aggregated Multiservices) | **6272**
- enable-reroute-uniform-link-check | **6273**
- enable-reverse-reroute | **6275**
- enable-session-cache | **6276**
- enable-snmp-tunnel-statistics (L2TP) | **6277**
- enable-subscriber-analysis (Services Options VMS Interfaces) | **6279**
- encap-type | **6280**
- encapsulate-inner-vlan | **6282**
- encapsulation (protocols evpn) | **6284**
- encapsulation (logical-link-layer) | **6286**
- encapsulation (interfaces-fti) | **6288**
- encapsulation (interfaces) | **6289**
- encapsulation (Dynamic Interfaces) | **6297**
- encapsulation (Logical Interface) | **6300**
- encapsulation-overhead | **6305**
- encapsulation (Physical Interface) | **6307**
- encapsulation-type (Layer 2 Circuits) | **6308**
- encapsulation-type (Layer 2 VPNs) | **6310**
- encapsulation-udp-fti | **6313**
- encoding (QFX Series) | **6314**
- encoding (security) | **6315**
- encoding | **6317**
- encoding-type | **6318**
- encrypt (HTTP Header Enrichment) | **6320**
- encryption | **6322**
- encryption-algorithm | **6324**
- encryption-algorithm (Security) | **6326**
- encryption-algorithm (Security IKE) | **6327**
- encryption-failures | **6330**
- encryption (Junos-FIPS Software) | **6331**
- encryption (Junos OS) | **6332**

- encryption (MACsec) | **6334**
- encryption (MACsec for MX Series) | **6336**
- end-interface | **6338**
- end-user-profile | **6339**
- endpoint-registration-timeout | **6342**
- enforce-first-as | **6343**
- enforce-strict-scale-limit-license (Subscriber Management) | **6345**
- engine-id | **6346**
- engine-id (Forwarding Options) | **6348**
- engine-not-ready | **6350**
- engine-not-ready (Security Antivirus Sophos Engine) | **6351**
- engine-type | **6352**
- enhanced | **6354**
- enhanced-cfm-mode | **6355**
- enhanced-convergence | **6357**
- enhanced (forwarding-class-accounting) | **6358**
- enhanced-hash-key | **6362**
- enhanced-hash-key (EX) | **6378**
- enhanced-hash-key profile1/profile2 | **6389**
- enhanced-hierarchical-policer | **6394**
- enhanced-mode | **6396**
- enhanced-mode (Network Services) | **6398**
- enhanced-mode-override | **6401**
- enhanced-policer | **6403**
- enhanced-priority-mode | **6404**
- enhanced-transmission-selection | **6406**
- enrollment | **6408**
- enrollment-retry | **6409**
- enrollment-url | **6411**
- enterprise-oid | **6412**
- entropy-label (protocols bgp) | **6413**
- entropy-label (protocols mpls) | **6416**

- environment-monitoring-policy (satellite-management) | **6417**
- environment-monitoring-policy (satellite-policies) | **6419**
- epacl-firewall-optimization | **6421**
- ephemeral | **6423**
- equals | **6426**
- equals (DHCP Relay Agent) | **6428**
- equals (Dynamic Profile) | **6433**
- equals (Event Policy) | **6434**
- eracl-ip6-match (packet-forwarding-options) | **6435**
- eracl-profile (packet-forwarding-options) | **6437**
- erroneous-ccm | **6439**
- error | **6441**
- error (Security) | **6446**
- error-id | **6449**
- es-options | **6451**
- esi | **6453**
- esi-resolution-per-bridge-domain | **6456**
- esmc-transmit (Chassis Synchronization) | **6458**
- establish-tunnels | **6460**
- ether-options | **6461**
- ether-pseudowire | **6471**
- ether-type (Dynamic Load Balancing) | **6473**
- ether-type (VLAN) | **6474**
- ether-type | **6476**
- ether-type (Chassis FPC) | **6477**
- ethernet | **6480**
- ethernet (Aggregated Devices) | **6481**
- ethernet (Alarm) | **6482**
- ethernet (Chassis) | **6484**
- ethernet (CoS for Multidestination Traffic) | **6485**
- ethernet-frame-size | **6486**
- ethernet-policer-profile | **6488**

- ethernet (Protocols OAM) | **6490**
- ethernet-ring | **6497**
- ethernet-switch-profile | **6499**
- ethernet-switching | **6501**
- ethernet-switching | **6503**
- ethernet-switching-options | **6506**
- ethernet-switching (Port Mirroring) | **6513**
- ethernet-vlan (Protocols Link Management) | **6514**
- eui-64 | **6515**
- evc-protocol cfm | **6517**
- evcs | **6518**
- event | **6520**
- event (CFM) | **6521**
- event (Enhanced Subscriber Management) | **6523**
- event (LFM) | **6525**
- event-options | **6527**
- event-rate | **6531**
- event-rate (Next Gen Services Service-Set Local System Logging) | **6532**
- event-script (Event Options) | **6533**
- event-script (Event Policy) | **6536**
- event-thresholds | **6537**
- events (Event Policy Correlated Events) | **6538**
- events (Event Policy Trigger Events) | **6540**
- evpn | **6541**
- evpn-aliasing-optimize | **6547**
- evpn (Easy EVPN LAG Configuration) | **6548**
- evpn-ssm-reports-only | **6553**
- evpn-vxlan | **6555**
- evpn-vxlan (Easy EVPN LAG Configuration) | **6559**
- examine-dhcp | **6561**
- examine-dhcpv6 | **6563**
- examine-fip | **6565**

- examine-vn2vf | **6566**
- examine-vn2vn | **6568**
- exceed-action | **6570**
- exception | **6571**
- exception-reporting | **6572**
- exception (Security Content Filtering) | **6573**
- exceptions | **6575**
- excess-bandwidth-share | **6577**
- excess-burst-size (QFX) | **6579**
- excess-burst-size | **6580**
- excess-priority | **6582**
- excess-priority (Dynamic Schedulers) | **6584**
- excess-rate (Schedulers) | **6585**
- excess-rate (Dynamic Schedulers) | **6588**
- excess-rate (Dynamic Traffic Shaping) | **6589**
- excess-rate-high | **6591**
- excess-rate-high (Dynamic Traffic Shaping) | **6593**
- excess-rate-low | **6594**
- excess-rate-low (Dynamic Traffic Shaping) | **6596**
- excess-rate-medium-high | **6597**
- excess-rate-medium-low | **6599**
- excess-rate (Schedulers) | **6601**
- exclude (Diameter Gx Profiles) | **6602**
- exclude (Dynamic MLD Interface) | **6604**
- exclude (for Administrative Groups) | **6605**
- exclude (for Fast Reroute) | **6606**
- exclude-hostname | **6608**
- exclude (JSRC Attributes) | **6609**
- exclude-protocol | **6610**
- exclude-protocol (MX Series) | **6612**
- exclude (Protocols IGMP) | **6614**
- exclude (Protocols MLD) | **6615**

- exclude-queue-overhead-bytes | **6616**
- exclude (RADIUS Attributes) | **6617**
- exclude (Schedulers) | **6626**
- exclude (Security Log) | **6627**
- exclude-srlg | **6629**
- excluded-address (Address-Assignment Pools) | **6631**
- excluded-range (Address-Assignment Pools) | **6632**
- exclusive-mac | **6634**
- execute-commands (Event Policy) | **6636**
- exercise | **6637**
- exit | **6638**
- exp (Class of Service) | **6639**
- exp-push-push-push | **6642**
- exp (Rewrite rules) | **6643**
- exp-swap-push-push | **6646**
- expand-loose-hop | **6647**
- explicit-congestion-notification | **6649**
- explicit-null-cos | **6650**
- explicit-null (Protocols BGP) | **6652**
- explicit-null (Protocols LDP) | **6654**
- explicit-null (Protocols MPLS) | **6655**
- explicit-priority | **6657**
- export (ISIS) | **6658**
- export (OSPF) | **6660**
- export (RIP) | **6661**
- export (Bootstrap) | **6663**
- export-format (JSON) | **6665**
- export-format (Flow Monitoring) | **6667**
- export (MPLS Traffic engineering database) | **6668**
- export-profile (Junos Telemetry Interface) | **6669**
- export-profiles | **6674**
- export (Protocols BGP) | **6677**

- export (Protocols DVMRP) | **6678**
- export (Protocols LDP) | **6680**
- export (Protocols MSDP) | **6681**
- export (Protocols PIM) | **6683**
- export (Protocols RIFT) | **6685**
- export (Protocols RIPng) | **6686**
- export-rib | **6688**
- export (Routing Options) | **6691**
- export-target | **6693**
- expression | **6694**
- extend-secondary-vlan-id | **6696**
- extended-echo | **6697**
- extended-isid-list | **6698**
- extended-nexthop | **6700**
- extended-nexthop-tunnel | **6701**
- extended-space | **6703**
- extended-statistics | **6704**
- extended-system-id | **6706**
- extended-vlan-list | **6707**
- extended-vni-all | **6709**
- extended-vni-list | **6710**
- Extensible Subscriber Services | **6712**
- extension-header | **6714**
- extension-provider | **6716**
- extension-service (System Extensions) | **6717**
- extension-service (System Services) | **6719**
- extensions | **6722**
- external-assigned (Address Pools) | **6725**
- external-authority | **6726**

enhanced-scale (IRB Interface)

IN THIS SECTION

- Syntax | 6175
- Hierarchy Level | 6175
- Description | 6175
- Options | 6175
- Required Privilege Level | 6176
- Release Information | 6176

Syntax

```
interface irb.<unit> {
    enhanced-scale;
}
```

Hierarchy Level

```
[edit interfaces irb unit <unit>]
```

Description

Enable the IRB interface to allocate Next Hop Identifiers (NHId) from the extended space, thereby supporting up to 1.5 million ARP/ND unicast next hop entries. This feature is essential for environments requiring a high scale of next hop entries.

Options

unit: The logical unit number for the IRB interface.

- Range: `<unit>` specifies the logical unit number.
- Default: None

- Syntax: <unit>
- Values: Any valid logical unit number.

Required Privilege Level

admin: To view this statement in the configuration.

Release Information

admin-control: To add this statement to the configuration.

e1-options (BITS Interfaces Signal Type)

IN THIS SECTION

- Syntax | 6176
- Hierarchy Level | 6177
- Description | 6177
- Options | 6177
- Required Privilege Level | 6177
- Release Information | 6177

Syntax

```
e1-options {
    framing (g704 | g704-no-cro4);
}
```

Hierarchy Level

```
[edit chassis synchronization interfaces bits ]
```

Description

Specify options for the E1 signal type for the building-integrated timing supply (BITS) interface. Valid only if `e1` is the `signal-type` configured at the [`edit chassis synchronization source interfaces` *`interface-name`* `bits signal-type`] hierarchy.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

External Clock Synchronization Overview for ACX Series Routers

Configuring External Clock Synchronization for ACX Series Routers

Synchronous Ethernet Overview

show chassis synchronization | **18879**

e1-options (Clock Synchronization)

IN THIS SECTION

- Syntax | 6178
- Hierarchy Level | 6178
- Description | 6178
- Options | 6178
- Required Privilege Level | 6179
- Release Information | 6179

Syntax

```
e1-options {
     framing (g704 | g704-no-crc4);
    line-encoding (ami | hdb3);
    sabit bit;
}
```

Hierarchy Level

```
[edit chassis synchronization interfaces external]
[edit chassis synchronization interfaces (external-0/0 | external-1/0)]
```

Description

Configure the E1 interface options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

e2e-transparent

IN THIS SECTION

- Syntax | 6179
- Hierarchy Level | 6179
- Description | 6180
- Options | 6180
- Required Privilege Level | 6180
- Release Information | 6180

Syntax

```
e2e-transparent;
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

> **NOTE**: If IGMP snooping is enabled on routers and switches that are configured with a PTP Transparent Clock, the PTP devices will not perform an IGMP join on 224.0.1.129. You will need to configure a static IGMP configuration in order to forward the PTP traffic.
>
> For example:
>
> ```
> set protocols igmp-snooping vlan vlan-name interface interface-name static group 224.0.1.129
> ```

Configure the end-to-end (E2E) transparent clock for Precision Time Protocol (PTP). With an end-to-end transparent clock, only the residence time is included in the timestamp in the packet. Transparent clock functionality is supported on PTP over Ethernet, IPv4, IPv6, unicast, and multicast. With PTP over Ethernet, one or two VLANs are supported. Transparent clock functionality is enabled globally and might be required in scenarios in which the interface on which packets are received and transmitted is unknown.

> **NOTE**: ACX5048 and ACX5096 routers do not support PTP over IPv6 for transparent clocks.

Options

There are no options.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D25.

Statement introduced in Junos OS Release 20.1R1 EVO for QFX5220 devices

Statement introduced in Junos OS Release 24.2R1 EVO for QFX5130-48C devices

RELATED DOCUMENTATION

Configuring Transparent Clock Mode for Precision Time Protocol

eapol-address (MACSec)

IN THIS SECTION

- Syntax | 6181
- Hierarchy Level | 6181
- Description | 6181
- Default | 6182
- Options | 6182
- Required Privilege Level | 6182
- Release Information | 6182

Syntax

```
eapol-address (pae | provider-bridge | lldp-multicast | destination unicast-address);
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name mka]
```

Description

Configure an EAPoL destination MAC address. The `eapol-address` *pae* is the default configuration.

To establish a MACsec session, MACsec Key Agreement PDUs (MKPDUs) are sent or received between nodes. These PDUs are Extensible Authentication Protocol over LAN (EAPoL) packets and, by default, their destination MAC address is the EAPoL multicast address 01:80:C2:00:00:03.

If the nodes are connected through a provider network, the multicast packets might be consumed or dropped, depending on their configuration. To overcome this issue, you can configure the destination MAC address. The configuration must match on both peer nodes to establish the MACsec session.

NOTE:

- The `pae`, `provider-bridge`, and `lldp-multicast` options are multicast addresses. You can configure a unicast address using the `destination` option.
- It is assumed that the adjacency between both the nodes is guaranteed by the provider network. The MKPDUs are not VLAN-tagged and include multicast address as their destination address. It is also also assumed that the provider network has a configuration to transfer the untagged MKPDUs to the destination node.

Default

Port Access Entity (PAE) group address (01:80:C2:00:00:03).

Options

`pae`	The Port Access Entity option is mapped to MAC address 01:80:C2:00:00:03. Do not use if 802.1X authentication is configured on the provider network.
`provider-bridge`	The provider bridge option is mapped to MAC address 01:80:C2:00:00:00. Do not use if STP/RSTP/MSTP protocols are configured on the provider network.
`lldp-multicast`	The Link Level Discovery Protocol multicast option is mapped to MAC address 01:80:C2:00:00:0E. Do not use if LLDP is configured on the provider network.
`destination` ***`unicast-address`***	The unicast address option is a configurable MAC address.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

Option for destination unicast address introduced in Junos OS Release 19.3R1 for MX Series routers.

RELATED DOCUMENTATION

mka (MX Series) | 10777

Configuring Advanced MACsec Features

eapol-block

IN THIS SECTION

- Syntax | 6183
- Hierarchy Level | 6183
- Description | 6184
- Default | 6184
- Options | 6184
- Required Privilege Level | 6185
- Release Information | 6185

Syntax

```
eapol-block {
    captive-portal;
    mac-radius;
    server-fail <seconds>;
}
```

Hierarchy Level

```
[edit logical-systems name protocols dot1x authenticator interface (all | interface-names)],
[edit logical-systems name protocols dot1x authenticator interface (all | interface-names) (server-reject-bridge-domain | server-reject-vlan)],
[edit protocols dot1x authenticator interface (all | interface-names)],
[edit protocols dot1x authenticator interface (all | interface-names) (server-reject-bridge-domain | server-reject-vlan)]
```

Description

Enable the device to ignore Extensible Authentication Protocol over LAN (EAPoL)-Start messages received from a client that has been authenticated so that the device does not trigger re-authentication. The device typically attempts to restart the authentication procedure by contacting the authentication server when it receives an EAPoL-Start message from a client—even for authenticated clients. You can configure the `eapol-block` statement to help prevent unnecessary downtime that can occur when the device waits for a response from the authentication server.

If you configure the device to block EAPoL-Start messages, when the device receives an EAPoL-Start message from an authenticated client, the device ignores the message and does not attempt to contact the authentication server for reauthentication. The existing authentication session that was established for the client remains open.

The EAPoL-Start messages are blocked only if the client is in the authenticated state. EAPoL-Start messages from new clients are accepted.

Default

If the `eapol-block` statement is not configured, the device attempts to contact the authentication server to authenticate the client when it receives an EAPoL-Start message.

Options

`captive-portal` Configure the device to ignore EAPoL-Start messages received from a client that has been authenticated using captive portal authentication.

`mac-radius` Configure the device to ignore EAPoL-Start messages received from a client that has been authenticated using MAC RADIUS authentication. The `mac-radius` option is also valid for clients authenticated using central Web authentication (CWA).

`server-fail` <*seconds*> Configure the device to ignore EAPoL-Start messages received from a client that has been authenticated using server fail fallback or server reject VLAN methods. Optionally, configure the time interval, in seconds, during which the device will not attempt to contact the authentication server to re-authenticate a client that has already been authenticated using server fail fallback.

- **Default:** 120 seconds.
- **Range:** 120 through 65,535 seconds.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support at the [`edit protocols dot1x authenticator interface` *`interface-name`*] hierarchy level introduced in Junos OS Releases 14.1X53-D40 and 15.1X53-D51 for EX Series switches.

`captive-portal` and `mac-radius` introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

> *Example: Configuring Fallback Options on EX Series Switches for EAP-TTLS Authentication and Odyssey Access Clients*

early-ageout

IN THIS SECTION

- Syntax | 6185
- Hierarchy Level | 6186
- Description | 6186
- Options | 6186
- Required Privilege Level | 6186
- Release Information | 6186

Syntax

```
early-ageout seconds;
```

Hierarchy Level

```
[edit security flow aging]
```

Description

Define the value before the device aggressively ages out a session from its session table.

Options

seconds —Amount of time that elapses before the device aggressively ages out a session.

- **Range:** 1 through 65,535 seconds
- **Default:** 20 seconds

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

east-interface

IN THIS SECTION

Syntax | 6187
Hierarchy Level | 6187

- Description | 6187
- Required Privilege Level | 6188
- Release Information | 6188

Syntax

```
east-interface {
    node-id mac-address;
    control-channel channel-name {
        vlan number;
        interface name interface-name
    }
    interface-none
    ring-protection-link-end;
}
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Define one of the two interface ports for Ethernet ring protection, the other being defined by the `west-interface` statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

EX Series switches do not use the node-id statement--the node ID is automatically configured on the switches using the MAC address.

NOTE: Always configure this port first, before configuring the `west-interface` statement.

NOTE: The Node ID is not configurable on EX Series switches. The node ID is automatically configured using the MAC address.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

RELATED DOCUMENTATION

- *Ethernet Ring Protection Switching Overview*
- Ethernet Ring Protection Using Ring Instances for Load Balancing
- west-interface | **31102**
- ethernet-ring | **6497**
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

echo-req

IN THIS SECTION

- Syntax | **6189**
- Hierarchy Level | **6189**
- Description | **6189**
- Options | **6189**
- Required Privilege Level | **6190**

Release Information | 6190

Syntax

```
echo-req {
    alarm-threshold (Security GPRS) {
        forward number;
        reverse number;
    }
    drop-threshold (Security GPRS) {
        forward number;
        reverse number;
    }
}
```

Hierarchy Level

```
[edit security gtp profile profile-name path-rate-limit message-type]
```

Description

Configure the echo-req parameter to limit the number of packets per minute for the GPRS tunneling protocol (GTP) echo request.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name path-rate-limit message-type]` hierarchy level is replaced by `[edit security gtp profile profile-name path-rate-limit message-type]`.

Options

`alarm-threshold`—To set alarm threshold for path rate limiting.

`drop-threshold`—To set drop threshold for path rate limiting.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding Path Rate Limiting for GTP Control Messages

Example: Limiting the Message Rate and Path Rate for GTP Control Messages

ecmp

IN THIS SECTION

- Syntax | 6190
- Hierarchy Level | 6191
- Description | 6191
- Required Privilege Level | 6191
- Release Information | 6191

Syntax

```
ecmp;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp oam bfd-liveness-detection],
[edit logical-systems logical-system-name protocols ldp oam fec address bfd-liveness-detection],
[edit protocols ldp oam bfd-liveness-detection],
[edit protocols ldp oam fec address bfd-liveness-detection]
```

Description

Allows LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the `ecmp` statement, you must also configure the `periodic-traceroute` statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the `periodic-traceroute` statement at the global hierarchy level (`[edit protocols ldp oam]`) while only configuring the `ecmp` statement for a specific FEC (`[edit protocols ldp oam fec address bfd-liveness-detection]`).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring ECMP-Aware BFD for LDP LSPs

ecmp-alb

IN THIS SECTION

- Syntax | 6192
- Hierarchy Level | 6192

- Description | 6192
- Options | 6192
- Required Privilege Level | 6193
- Release Information | 6193

Syntax

```
ecmp-alb {
    apply-groups;
    apply-groups-except;
    tolerance;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable adaptive load balancing for equal-cost multipath (ECMP) next hops.

NOTE: The `ecmp-alb` statement can be enabled only when the `[edit chassis network-services enhanced-ip]` statement is configured.

Starting in Junos OS Release 20.1R1, you can configure adaptive load balancing for ECMP next hops on multiple Packet Forwarding Engines on the same line card for even distribution of the traffic and redundancy.

Options

`apply-groups`—Specify the groups from which to inherit configuration data.

`apply-groups-except`—Specify the groups from which configuration data should not be inherited.

`tolerance`—Specify the adaptive tolerance in percentage.

- **Default:** 20%.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

ecmp-fast-reroute

IN THIS SECTION

- Syntax | 6193
- Hierarchy Level | 6193
- Description | 6194
- Required Privilege Level | 6194
- Release Information | 6194

Syntax

```
ecmp-fast-reroute;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table ],
[edit routing-options forwarding-table]
```

Description

Enables equal-cost multipath (ECMP) fast reroute protection. If a link fails, ECMP uses fast reroute protection to shift packet forwarding to operational links to reduce packet loss. Fast reroute protection updates ECMP information for the interface without waiting for route table updates. When the next route table update occurs, a new set of ECMP interfaces are added with fewer links, or the route points to a single next hop.

Without ECMP fast reroute protection, upon link failure the creation of the new ECMP set is delayed while the routing table information is updated. Once the new ECMP set is created, the hashing algorithm calculates new paths. Enabling the `ecmp-fast-reroute` option eliminates the routing table convergence delay.

ECMP is a global configuration and will impact only the nexthop addresses learnt after the ECMP configuration. However, the existing nexthop addresses are not impacted. Therefore, you must either remove the existing configuration and restore it along with the fast reroute configuration or you can restart routing or the linecard. You can also reboot the device to add the ECMP fast reroute links for the existing nexthop addresses.

> **NOTE**: ECMP works differently with indirect next hops. Please see *ECMP Flow-Based Forwarding* for more information.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Command introduced before Junos OS Release 17.4.

RELATED DOCUMENTATION

ECMP Flow-Based Forwarding

ecmp-local-bias

IN THIS SECTION

- Syntax | 6195
- Hierarchy Level | 6195
- Description | 6195
- Default | 6196
- Required Privilege Level | 6196
- Release Information | 6196

Syntax

```
ecmp-local-bias;
```

Hierarchy Level

```
[edit forwarding-options load-balance]
```

Description

By default, equal cost multi-path (ECMP) traffic flows are distributed more-or-less equally between forwarding next-hops. Enable `ecmp-local-bias` to have ECMP traffic prefer local forwarding next-hops (that is, local to the packet forwarding engine (PFE) that is performing the packet look up) over remote ones, and to distribute the flows among local members. Local bias percentages cannot be assigned. Note that `ecmp-local-bias` is not intended to be used in conjunction with any other load balancing schemes.

This feature applies chassis-wide. It supports BGP prefixes that are directly reachable with IPv4 MPLS ECMP next-hops on Gigabit Ethernet (ge-) and 10-Gigabit Ethernet (xe-) interfaces. Multicast traffic is not affected.

Use `ecmp-local-bias` when you need to direct ECMP traffic towards local links, for example, to ensure that the overall load on the fabric is reduced.

> **NOTE**: For packet forwarding engine (PFE) to achieve the functionality of ECMP local bias, the dependent config knob of ecmp-fast-reroute is required. In MX platforms, `ecmp-fast-reroute` configuration command is available as an explicit router configuration.

Default

Not enabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Configuring Per-Prefix Load Balancing

ecmp-resilient-hash

IN THIS SECTION

- Syntax | 6197
- Hierarchy Level | 6197
- Description | 6197
- Required Privilege Level | 6197
- Release Information | 6197

Syntax

```
ecmp-resilient-hash;
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key]
```

Description

Enable resilient hashing for ECMP groups, to minimize remapping of destination paths.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Configuring Resilient Hashing for LAGs/ECMP Groups

edge

IN THIS SECTION

- Syntax | 6198
- Hierarchy Level | 6198
- Description | 6198
- Default | 6199

Required Privilege Level | 6199
Release Information | 6199

Syntax

```
edge;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp | vstp) interface interface-name],
[edit logical-systems logical-system-name protocols mstp msti msti-id interface interface-name],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp | vstp) interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mstp
msti msti-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id interface interface-name],
```

```
[edit protocols (mstp | rstp | vstp) interface (all | interface-name)],
[edit protocols mstp msti msti-id interface (all | interface-name)],
[edit protocols vstp vlan vlan-id interface (all | interface-name)],
```

```
[edit routing-instances routing-instance-name protocols (mstp | rstp | vstp) interface interface-
name],
[edit routing-instances routing-instance-name protocols mstp msti msti-id interface interface-
name],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id interface interface-name]
```

Description

For Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure interfaces as edge ports or edge interfaces. Edge ports do not expect to

receive BPDUs. If a BPDU is received, the port becomes a non edge port unless `set protocols (vstp|rstp|mstp) bpdu-block-on-edge` is configured.

> **NOTE**: Although the `edge` configuration statement appears in the [`edit protocols stp interface (all | interface-name)`] or [`edit protocols rstp force-version stp interface (all | interface-name)`] hierarchy on the switch, this statement has no effect on the switch operation if you configure it.

Default

Edge interfaces are not enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding STP

Example: Configuring Network Regions for VLANs with MSTP on Switches

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

edge-virtual-bridging

IN THIS SECTION

- Syntax | 6200
- Hierarchy Level | 6200
- Description | 6201
- Default | 6201
- Required Privilege Level | 6201
- Release Information | 6201

Syntax

```
edge-virtual-bridging {
    traceoptions {
        file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
        flag flag ;
    }
    vsi-discovery {
        interface interface-name
        vsi-policy vsi-policy-name
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure edge virtual bridging (EVB). EVB enables a virtualized station (a physical end station, a server, connected to virtual machines (VMs)) to network with an adjacent switch so that applications residing on the virtual machines can interact with each other and external networks through a technology called virtual Ethernet packet aggregator (VEPA).

The remaining statements are explained separately. See CLI Explorer.

Default

EVB is disabled by default.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on an EX Series Switch

Configuring Edge Virtual Bridging on an EX Series Switch

edit

IN THIS SECTION

- Syntax | 6202
- Description | 6202
- Options | 6202
- Required Privilege Level | 6202

Release Information | 6202

Syntax

```
edit statement-path
```

Description

Move inside the specified statement hierarchy. If the statement does not exist, it is created.

You cannot use the `edit` command to change the value of identifiers. You must use the `set` command.

Options

statement-path	Path to the statement.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Displaying the Current Configuration

effective-aigp

IN THIS SECTION

- Syntax | 6203
- Hierarchy Level | 6203
- Description | 6203
- Options | 6203
- Required Privilege Level | 6204
- Release Information | 6204

Syntax

```
effective-aigp <metric-offset>;
```

Hierarchy Level

```
[edit fabric protocols bgp group name metric-out],
[edit logical-systems name protocols bgp group name metric-out],
[edit logical-systems name routing-instances name protocols bgp group name metric-out],
[edit logical-systems name tenants name routing-instances name protocols bgp group name metric-out],
[edit protocols bgp group name metric-out],
[edit routing-instances name protocols bgp group name metric-out],
[edit tenants name routing-instances name protocols bgp group name metric-out]
```

Description

To track the effective AIGP metric. Effective AIGP is the AIGP value advertised with the route plus the IGP cost to reach the nexthop.

Options

metric-offset Metric offset for MED.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.2R1.

effective-shaping-rate

IN THIS SECTION

- Syntax | 6204
- Hierarchy Level | 6204
- Description | 6204
- Required Privilege Level | 6205
- Release Information | 6205

Syntax

```
effective-shaping-rate;
```

Hierarchy Level

```
[edit chassis]
```

Description

Specify that the Cos-Effective-Shaping-Rate VSA [26–177] included in RADIUS Acct-Start, Acct-Stop, and Interim-Acct messages reports the actual rate of the downstream traffic for a subscriber, in kilobits per second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Reporting the Effective Shaping Rate for Subscribers

egress

IN THIS SECTION

- Syntax | 6205
- Hierarchy Level | 6206
- Description | 6206
- Required Privilege Level | 6206
- Release Information | 6206

Syntax

```
egress {
     interface (all | interface-name);
    }
 vlan (vlan-id | vlan-name);
```

Hierarchy Level

For platforms without ELS:

```
[edit ethernet-switching-options analyzer name input]
```

For platforms with ELS:

```
[edit forwarding-options analyzer name input]
```

Description

Specify interface or VLAN for which egressing traffic is mirrored. (The `vlan` statement is not supported on all switches.)

The remaining statement is explained separately. See CLI Explorer.

> **NOTE**: If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs. If you do so, some of the mirrored packets might contain incorrect VLAN IDs.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Port Mirroring and Analyzers

Configuring Port Mirroring

Examples: Configuring Port Mirroring for Local Analysis

egress (Analyzer)

IN THIS SECTION

- Syntax | 6207
- Hierarchy Level | 6207
- Description | 6207
- Required Privilege Level | 6208
- Release Information | 6208

Syntax

```
egress {
    bridge-domain bridge-domain-name;
    interface (all | interface-name);
    routing-instance {
        instance-name {
            bridge-domain bridge-domain-name;
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input]
```

Description

Specify ports where traffic exiting the interface is to be mirrored in a mirroring configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

> *Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX9200 Switches*

egress (Buffer Configuration)

IN THIS SECTION

- Syntax | 6208
- Hierarchy Level | 6209
- Description | 6209
- Default | 6209
- Options | 6209
- Required Privilege Level | 6210
- Release Information | 6210

Syntax

```
egress {
    percent percent;
    buffer-partition (lossless | lossy | multicast) {
        percent percent;
```

```
    }
}
```

Hierarchy Level

```
[edit class-of-service shared-buffer]
```

Description

Configure the global shared buffer pool allocation for egress traffic. The system allocates the shared buffer pool dynamically across its ports as the ports require memory space. Some buffer space is reserved for other buffers such as dedicated buffers (buffers allocated permanently to ports).

The percentage you specify is the percentage of available (user-configurable) buffer space allocated to the global shared egress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports. However, on a port, you can configure the portion of dedicated port buffer space allocated to each queue in the scheduler configuration using the `buffer-size` option.)

> **CAUTION**: Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

You can also partition the shared buffer pool to adjust the egress buffer allocations for different mixes of network traffic using the `buffer-partition` statement.

Default

The default shared buffer percentage is 100 percent. (All available buffer space is allocated to the shared buffer pool.)

Options

`percent` *percent*—Percentage of available egress buffer space allocated to the shared buffer pool. If the percentage is less than 100 percent, the remaining buffer space is allocated to the dedicated buffer pool.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic

Configuring Global Ingress and Egress Shared Buffers

Understanding CoS Buffer Configuration

egress-chaining

IN THIS SECTION

- Syntax | 6211
- Hierarchy Level | 6211
- Description | 6211
- Options | 6211
- Required Privilege Level | 6212
- Release Information | 6212

Syntax

```
egress-chaining {
    spring-te;
}
```

Hierarchy Level

```
[edit routing-options forwarding-table]
```

Description

Enable egress chaining for devices that support a maximum of two labels push at the ingress. You can configure egress chaining to implement more than two labels push using a composite next hop at the egress and pass a pointer from the ingress to the egress about the location of its label stack. Route pointing to this next hop chain can be IPv4, IPV6, or an MPLS label (binding segment ID).

NOTE:

- Label push at the egress should ensure that they are pushed ahead of any other labels programmed at the egress.
- Egress chaining is possible for one chain only. Any more composite next hops after egress chaining are flattened into forwarding next hops.
- Junos OS release 20.4R1 supports only a single egress chain.

Options

spring-te Enable egress chaining for segment routing traffic-engineered (SPRING-TE) routes.

When you configure the `spring-te`, you set the egress chain flag in the SPRING-TE route's next hop. This enables installation of a composite next hop in the egress Packet Forwarding Engine (PFE).

NOTE: Ensure that you do not resolve the SPRING-TE route further over another composite next hop when the `spring-te` option is enabled.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

fib-next-hop-split | 7006

egress-l2-extended-match

IN THIS SECTION

- Syntax | 6212
- Hierarchy Level | 6212
- Description | 6213
- Required Privilege Level | 6213
- Release Information | 6213

Syntax

```
egress-l2-extended-match
```

Hierarchy Level

```
[edit system packet-forwarding-options firewall-profile]
```

Description

Set this configuration to support **TCP-flags, dscp, ip-version, source-port, destination-port, ip-source-address, ip-destination-address, source-prefix-list, destination-prefix-list, and icmp-type** matches for egress port firewall filters and egress VLAN firewall filters.

> **NOTE**: The configuration is applicable for EX4400, EX4100, QFX5120-48Y, EX4650-48Y, QFX5120-32C, QFX5120-48T, and QFX5120-48YM platforms.
>
> This knob will restart PFE for standalone box. For Virtual Chassis, it requires manual reboot to take effect.
>
> If the knob is not configured, then commit error is issued if these matches are present in the firewall filter configuration.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos Release 23.4.

egress-policer-overhead

IN THIS SECTION

- Syntax | 6214
- Hierarchy Level | 6214
- Description | 6214
- Options | 6214
- Required Privilege Level | 6214
- Release Information | 6214

Syntax

```
egress-policer-overhead bytes;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number]
```

Description

Add the specified number of bytes to the actual length of an Ethernet frame when determining the actions of Layer 2 policers, MAC policers, or queue rate limits applied to output traffic on the line card. You can configure egress policer overhead to account for egress *shaping* overhead bytes added to output traffic on the line card.

On M Series and T Series routers, this statement is supported on Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs and Enhanced IQ2 (IQ2E) PICs. On MX Series routers, this statement is supported for interfaces configured on Dense Port Concentrators (DPCs).

> **NOTE**: This statement is not supported on Modular Interface Cards (MICs) or Modular Port Concentrators (MPCs) in MX Series routers.

Options

bytes—Number of bytes added to a packet exiting an interface.

- **Range:** 0–255 bytes
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 11.1.

RELATED DOCUMENTATION

egress-shaping-overhead | 6223

Policer Overhead to Account for Rate Shaping Overview

Example: Configuring Policer Overhead to Account for Rate Shaping

Configuring a Policer Overhead

CoS on Enhanced IQ2 PICs Overview

egress-policy

IN THIS SECTION

- Syntax | 6215
- Hierarchy Level | 6215
- Description | 6216
- Default | 6216
- Options | 6216
- Required Privilege Level | 6216
- Release Information | 6216

Syntax

```
egress-policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Control the prefixes advertised into LDP.

Default

Only the loopback address is advertised.

Options

policy-names—Name of one or more routing policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Prefixes Advertised into LDP from the Routing Table

egress-protection (BGP)

IN THIS SECTION

- Syntax | 6217
- Hierarchy Level | 6217
- Description | 6217
- Options | 6217
- Required Privilege Level | 6217

Release Information | 6218

Syntax

```
egress-protection {
    context-identifier context-id-ip-address;
    keep-import policy-name;
}
```

Hierarchy Level

```
[edit protocols bgp group group-name family inet labeled-unicast],
[edit protocols bgp group group-name family inet-vpn unicast],
[edit protocols bgp group group-name family inet6-vpn unicast],
[edit protocols bgp group group-name family iso-vpn unicast]
```

Description

Enable egress protection for the configured BGP Network Layer Reachability Information (NLRI).

Options

context-identifier *context-id-ip-address*	(Optional) The IPv4 address of the context identifier.
keep-import *policy-name*	The import policy that contains routes with the configured route target community. BGP keeps all these routes in the VPN routing table. The protector PE router scans the policy and builds the remote instance with the configured community name from the policy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4R3.

Statement introduced for labeled unicast in Junos OS Release 14.1.

RELATED DOCUMENTATION

Example: Configuring Egress Protection for Layer 3 VPN Services

Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP

Example: Configuring Egress Protection for BGP Labeled Unicast

egress-protection (Layer 2 circuit)

IN THIS SECTION

- Syntax | 6218
- Hierarchy Level | 6219
- Description | 6219
- Options | 6219
- Required Privilege Level | 6219
- Release Information | 6219

Syntax

```
egress-protection {
    protected-l2circuit {
        egress-pe address;
        ingress-pe address;
        virtual-circuit-id identifier;
    }
    protector-interface interface-name;
    protector-pe address {
        context-identifier identifier;
```

```
        lsp lsp-name;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit protocols l2circuit neighbor address interface interface-name]
```

Description

Configures an egress protection virtual circuit (EPVC).

Options

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS release 10.4.

RELATED DOCUMENTATION

Example: Configuring an Egress Protection LSP for a Layer 2 Circuit

egress-protection (MPLS)

IN THIS SECTION

- Syntax | 6220
- Hierarchy Level | 6220
- Description | 6220
- Options | 6221
- Required Privilege Level | 6221
- Release Information | 6221

Syntax

```
egress-protection {
    context-identifier context-id {
        primary | protector;
        metric igp-metric-value;
        advertise-mode (stub-alias | stub-proxy);
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Enables an Edge Protection Virtual Circuit (EPVC) for the MPLS protocol.

Options

context-identifier *context-id-ip-address*	(Optional) The context identifier IPv4 address.
metric *igp-metric-value*	(Optional) The IGP metric value ranging from `2` through `16777215`.
(primary | protector)	On the primary PE router, configure as type `primary`. On the protector PE router, configure as type `protector`.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Options `primary`, `protector`, and `metric` introduced in Junos OS Release 11.4R3.

Option `advertise-mode` introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Example: Configuring Egress Protection for Layer 3 VPN Services

Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP

egress-protection (Routing Instances)

IN THIS SECTION

- Syntax | 6222
- Hierarchy Level | 6222
- Description | 6222
- Options | 6222

Required Privilege Level | 6222
Release Information | 6222

Syntax

```
egress-protection {
    context-identifier context-id-ip-address;
    }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name]
```

Description

Enable egress instance protection and provides customer edge (CE) VRF-level context-id granularity for each virtual routing and forwarding (VRF) instance.

Options

context-identifier *context-id-ip-address*	(Optional) The IPv4 address for the context identifier of the protected PE router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4R3.

RELATED DOCUMENTATION

Example: Configuring Egress Protection for Layer 3 VPN Services

egress-shaping-overhead

IN THIS SECTION

- Syntax | 6223
- Hierarchy Level | 6223
- Description | 6223
- Options | 6224
- Required Privilege Level | 6224
- Release Information | 6225

Syntax

```
egress-shaping-overhead number;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number traffic-manager],
[edit chassis lcc number fpc slot-number pic pic-number traffic-manager]
```

Description

Number of bytes to add to packet to determine shaped session packet length.

NOTE: On M Series and T Series routers with Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs and Enhanced IQ2 (IQ2E) PICs and on MX Series routers with Dense Port Concentrators (DPCs) only, to account for egress shaping overhead bytes added to

output traffic on the line card, you must use the `egress-policer-overhead` statement to explicitly configure corresponding egress policing overhead for Layer 2 policers, MAC policers, or queue rate limits applied to output traffic on the line card.

NOTE: For MIC and MPC interfaces on MX Series routers, by default the value of `egress-shaping-overhead` is configured to 20, which means that the number of class-of-service (CoS) shaping overhead bytes to be added to the packets is 20. The interfaces on DPCs in MX Series routers, the default value is zero. For interfaces on PICs other than the 10-port 10-Gigabit Oversubscribed Ethernet (OSE) Type 4, you should configure `egress-shaping-overhead` to a minimum of 20 bytes to add a shaping overhead of 20 bytes to the packets.

NOTE: When you change the `egress-shaping-overhead` value, on M Series, T Series, and MX104 routers the PIC on which it is changed is restarted. On MX5 routers, the MIC on which it is changed is restarted. On other MX Series routers, the DPC/MPC on which it is changed is restarted.

Options

number—When traffic management (queuing and scheduling) is configured on the egress side, the number of CoS shaping overhead bytes to add to the packets on the egress interface.

Range:

- –63 through 192.
- –62 through 192 for vSRX Virtual Firewall.

NOTE: The L2 headers (DA/SA + VLAN tags) are automatically a part of the shaping calculation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

- egress-policer-overhead | 6213
- ingress-shaping-overhead | 8649
- mode (Traffic Manager) | 10807
- ingress-shaping-overhead | 8649
- traffic-manager | 30036

egress-stats (Flat-File Accounting Options)

IN THIS SECTION

- Syntax | 6225
- Hierarchy Level | 6226
- Description | 6226
- Options | 6226
- Required Privilege Level | 6227
- Release Information | 6227

Syntax

```
egress-stats {
    all-fields;
    input-bytes;
    input-packets;
    output-bytes;
    output-packets;
    queue-id;
    red-drop-bytes;
```

```
    red-drop-packets;
    tail-drop-packets;
    total-drop-packets;
}
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name fields]
```

Description

Specify egress queue statistics to be collected for the interface.

Options

all-fields Collect all egress queue statistics available for the interface context, logical or physical.

input-bytes Collect the number of octets queued including traffic dropped because of congestion.

input-packets Collect the number of packets queued including traffic dropped because of congestion.

output-bytes Collect the number of octets transmitted by the egress queue.

output-packets Collect the number of packets transmitted by the egress queue.

queue-id Collect the logical identifier for the egress queue; identifies the traffic class.

red-drop-bytes Collect the number of octets dropped on the egress queue because of random early detection.

red-drop-packets Collect the number of packets dropped on the egress queue because of random early detection.

tail-drop-packets Collect the number of packets dropped in the egress queue because of tail drop.

total-drop-packets Collect the total number of packets dropped in the egress queue.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Flat-File Accounting Overview

egress-te

IN THIS SECTION

- Syntax | 6227
- Hierarchy Level | 6228
- Description | 6228
- Options | 6228
- Required Privilege Level | 6229
- Release Information | 6229

Syntax

```
egress-te {
    backup-path backup-path;
    import;
    install-address address;
```

```
    no-install;
    rib (inet.0 | inet6.0);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Enable egress peer engineering to direct core service traffic such as MPLS RSVP to a specific single-hop egress BGP peer. The ingress BGP peer can traffic-engineer the core inet unicast and inet6 unicast service traffic using BGP labeled unicast towards a specific egress BGP peer, which is also the AS boundary router.

You can enable MPLS fast reroute (FRR) on the egress BGP peer, which has traffic engineering enabled. The AS boundary router switches to the backup path when the primary link fails. Specify a predefined template with one or more backup paths. You can define a template using the "egress-te-backup-paths" on page 6232 configuration statement and configure one or more backup path for MPLS FRR. The backup paths specified in the template must belong to the same address family as the BGP peer.

Options

none	Enable traffic engineering on the egress peer.
backup-path ***backup-path***	(Optional) Specify a predefined template that has the configured backup path for MPLS fast reroute.

import	Import policy to set attributes on the egress-te created route.
install-address	Host (/32 or /128) address to install egress-te route in inet[6].3 table.
no-install	Avoid installation to FIB or resolving over egress-te route.
rib (inet.0 | inet6.0)	Install egress-te route in inet[6].0 instead of inet[6].3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R4.

RELATED DOCUMENTATION

egress-te-backup-paths | 6232

Configuring Egress Peer Traffic Engineering by Using BGP Labeled Unicast and Enabling MPLS Fast Reroute

Egress Peer Traffic Engineering Using BGP Labeled Unicast Overview

Example: Configuring Egress Peer Traffic Engineering Using BGP Labeled Unicast

egress-te-adj-segment

IN THIS SECTION

- Syntax | 6230
- Hierarchy Level | 6230
- Description | 6230
- Options | 6231
- Required Privilege Level | 6231

Release Information | 6231

Syntax

```
egress-te-adj-segment name {
    egress-te-backup-ip-forward instance-name;
    egress-te-backup-segment label label-value;
    egress-te-set set-name weight weight;
    label label-value;
    next-hop next-hop-address;
    egress-te-link-attribute {
        delay-metric <value>
        }
}
```

Hierarchy Level

```
[edit logical-systems name protocols bgp group name neighbor],
[edit logical-systems name routing-instances name protocols bgp group
name neighbor],
[edit logical-systems name tenants name routing-instances name protocols
bgp group name neighbor],
[edit protocols bgp group name neighbor],
[edit routing-instances name protocols bgp group name neighbor],
[edit tenants name routing-instances name protocols bgp group name neighbor]
```

Description

Specify a BGP peer adjacency segment for only multihop EBGP peers. You can specify a segment set so that the packet is sent to any member in the set, which is the equal-cost multipath next hop. You can include an adjacent segment or a node segment to the same peer segment set. You must configure "egress-te-node-segment" on page 6234 before configuring the BGP peer adjacency segment.

name Specify a name for the peer adjacency segment.

Options

egress-te-backup-ip-forward *instance-name*
Specify a routing-instance name to use the IP forward backup path for egress TE.

egress-te-backup-segment label *label-value*
Specify the backup label value for a segment for MPLS fast reroute. The label must already be associated with a segment and must be different from the protected segment specified with the "egress-te-node-segment" on page 6234 or "egress-te-adj-segment" on page 6229 statement.

egress-te-set *set-name* **weight** *weight*
Specify the name of the segment set configured with the"egress-te-set-segment" on page 6236 statement. You can include both node segments and adjacency segments in the same segment set. The egress router advertises the peer node SID label for all its peers including the EBGP peers and the controller advertises these SID labels to the ingress router.

- **Range:** weight: 1 through 255

label *label-value*
Specify a label from the static label range for the peer adjacency segment.

next-hop *next-hop-address*
Specify a directly connected next-hop address for the peer adjacency segment.

egress-te-link-attribute
Specify the link attribute for BGP EPE link delay metrics for path computation.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

`egress-te-backup-ip-forward` option introduced in Release 20.1R1.

`delay-metric <value>` option introduced in Junos os Release 21.4R1.

RELATED DOCUMENTATION

egress-te-node-segment | 6234

egress-te-set-segment | 6236

egress-te-backup-paths

IN THIS SECTION

- Syntax | 6232
- Hierarchy Level | 6232
- Description | 6232
- Options | 6233
- Required Privilege Level | 6233
- Release Information | 6234

Syntax

```
egress-te-backup-paths {
    template path-name {
        ip-forward rti-name;
        peer peer-addr;
        remote-nexthop remote-nh-addr;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical system name protocols bgp],
[edit logical-systems logical system name protocols bgp group group-name],
[edit protocols bgp],
[edit protocols bgp group group-name]
```

Description

Specify backup paths for MPLS fast reroute (FRR) on an egress peer, which has traffic engineering enabled. Egress peer engineering directs core service traffic such as MPLS RSVP to a specific egress BGP peer. The ingress BGP peer can traffic-engineer the core inet unicast and inet6 unicast service traffic using BGP labeled unicast towards a specific egress BGP peer, which is also the AS boundary router.

Specify a backup path through another directly connected external BGP peer. The configured backup path provides MPLS fast reroute when the primary link fails, and the AS boundary router redirects the traffic received from the core to the this backup path. You can configure more than one backup path on the egress BGP peer. The specified backup paths are automatically installed into the MPLS forwarding table of the egress BGP peer configured with the egress traffic engineering feature.

Options

template *path-name*

Define a template that can be reused by multiple BGP groups or peers. All addresses listed in one template must belong to the same IP address family as the protected device that is the egress BGP peer.

ip-forward *rti-name*

(Optional) Configure this option if you want the egress peer to perform an IP look up in the `inet6.0` table for backup path that egress BGP peer must use for faster reroute You can optionally specify a routing instance. If you do not specify a routing instance, the device configures the backup path for the master instance.

You cannot use this option with the `remote-nexthop` option.

> **CAUTION**: `ip-forward` option might cause forwarding loops if the IP route chooses an internal path. To avoid forwarding loops configure a virtual routing and forwarding (VRF) instance with leaked external routes only, and use this VRF instance with the `ip-forward` option.

peer *peer-addr*

(Optional) Specify another directly connected external BGP peer that the device must use for faster reroute when the primary link fails. Each template can specify one or more external BGP peers.

remote-nexthop *remote-nh-addr*

(Optional) Specify a remote next-hop address if transit peering is not available locally to tunnel traffic to another AS boundary router in the local AS that has transit connectivity. The specified remote next-hop address must have the ability to forward this redirected traffic to its destination. This option does not support multiple routing instances; therefore, do not use this option with the `ip-forward` option.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2 R4.

RELATED DOCUMENTATION

egress-te | 6227

Configuring Egress Peer Traffic Engineering by Using BGP Labeled Unicast and Enabling MPLS Fast Reroute

Egress Peer Traffic Engineering Using BGP Labeled Unicast Overview

Example: Configuring Egress Peer Traffic Engineering Using BGP Labeled Unicast

egress-te-node-segment

IN THIS SECTION

- Syntax | 6234
- Hierarchy Level | 6235
- Description | 6235
- Options | 6235
- Required Privilege Level | 6236
- Release Information | 6236

Syntax

```
egress-te-node-segment name {
    egress-te-backup-ip-forward instance-name;
    egress-te-backup-segment label label-value;
    egress-te-set set-name weight weight;
    label label-value;
}
```

Hierarchy Level

```
[edit logical-systems name protocols bgp group name neighbor],
[edit logical-systems name routing-instances name protocols bgp group
name neighbor],
[edit logical-systems name tenants name routing-instances name protocols
bgp group name neighbor],
[edit protocols bgp group name neighbor],
[edit routing-instances name protocols bgp group name neighbor],
[edit tenants name routing-instances name protocols bgp group name neighbor]
```

Description

Specify a BGP peer node segment for both single hop and multihop EBGP peers. You can configure a node segment as a member of a segment set so that the packet is sent to any member that is the ECMP next hop in the set. Before configuring a node segment as a member of a segment set, you must create a segment set first. If you fail to create the segment set before assigning a node segment, the commit might fail.

This configuration enables egress peer engineering using BGP link-state distribution in a network configured with segment routing. The egress router advertises the peer node SID label for all its peers and the controller advertises these SID labels to the ingress router.

Options

egress-te-backup-ip-forward *instance-name*

Specify a routing-instance name to use the IP forward backup path for egress TE.

egress-te-backup-segment label *label-value*

Specify the backup label value for a segment for MPLS fast reroute. The label must already be associated with a segment and must be different from the protected segment specified with the "egress-te-node-segment" on page 6234 or "egress-te-adj-segment" on page 6229 statement.

egress-te-set *set-name* weight *weight*

Specify the name of the segment set configured with the"egress-te-set-segment" on page 6236 statement. You can include both node segments and adjacency segments in the same segment set. The egress router advertises the peer node SID label for all its peers including the EBGP peers and the controller advertises these SID labels to the ingress router.

- **Range:** weight: 1 through 255

label *label-value* Specify a label from the static label range for the peer node segment.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

`egress-te-backup-ip-forward` option introduced in Release 20.1R1.

RELATED DOCUMENTATION

egress-te-adj-segment | 6229

egress-te-set-segment | 6236

egress-te-set-segment

IN THIS SECTION

- Syntax | 6237
- Hierarchy Level | 6237
- Description | 6237
- Options | 6237
- Required Privilege Level | 6238
- Release Information | 6238

Syntax

```
egress-te-set-segment name {
    egress-te-backup-ip-forward <instance-name>;
    egress-te-backup-segment label label-value;
    label label-value;
}
```

Hierarchy Level

```
[edit logical-systems name protocols              bgp],
[edit logical-systems name              routing-instances              name protocols bgp],
[edit logical-systems name tenants name routing-instances name protocols              bgp],
[edit protocols bgp],
[edit routing-instances name protocols              bgp],
[edit tenants name routing-instances name protocols bgp]
```

Description

Specify a peer segment set that can include adjacency segments or node segments, or a combination of both as members. You can assign a label, which is represented as equal-cost multipath next hop to send a packet to any member in the set.

name Specify a unique name for the BGP peer segment set.

Options

egress-te-backup-ip-forward *instance-name* Specify a routing-instance name to use the IP forward backup path for egress TE.

egress-te-backup-segment label *label-value* Specify the backup label value for a segment for MPLS fast reroute. The label must already be associated with a segment and must be different from the protected segment specified with the "egress-te-node-segment" on page 6234 or "egress-te-adj-segment" on page 6229 statement.

label *label-value* Specify a label from the static label range for the peer segment set.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

`egress-te-backup-ip-forward` option introduced in Release 20.1R1.

RELATED DOCUMENTATION

egress-te-adj-segment | 6229

egress-te-node-segment | 6234

egress-quantization

IN THIS SECTION

- Syntax | 6238
- Hierarchy Level | 6239
- Description | 6239
- Default | 6240
- Options | 6241
- Required Privilege Level | 6241
- Release Information | 6241

Syntax

```
egress-quantization
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key ecmp-dlb]
```

Description

Selective dynamic load balancing selects an optimal link based on:

- Port load metric - Amount of traffic in bytes transmitted in recent past over each ECMP link.
- Port queue metric - Amount of traffic in number of cells enqueued on each ECMP link for transmission.

Use `egress-quantization` to configure the desired ratio of port load metric to port queue metric based on the traffic pattern. Member link quality is quantized in the range of 0 to 7. A link quality value of 0 indicates the worst quality link whereas a value of 7 indicates a best quality link. Quantization levels can be configured by setting minimum and maximum thresholds along with rate-weightage.

Rate-weightage is the ratio of port load metric to port queue metric. Assume rate-weightage is configured as 60%. Then a link is chosen amongst competing links if the link has a higher port load metric than its port queue metric.

Minimum threshold is the minimum bandwidth threshold level of the total link bandwidth. When bandwidth level drops below this level, the device assigns this link with a best quality band value (value 7). Assume the value set is 20. When the ECMP link transmits traffic lower than 20% of link bandwidth, the device assigns a best quality band value to this link.

Maximum threshold is the maximum bandwidth threshold level of the total link bandwidth. When bandwidth level rises above this level, the device assigns this link with a worst quality band value (value 0). Assume the value is set to 80. When the ECMP link transmits traffic higher than 80% of the link bandwidth, the device assigns worst quality band value to this link.

The following table summarizes egress quantization link quality values.

Table 32: Egress quantization link quality values

Egress quantization when egress rate weightage is configured with value of 100%	
Egress Link Bandwidth utilization	Quality Band
Lesser than 20	7 (Best)

Table 32: Egress quantization link quality values *(Continued)*

Egress quantization when egress rate weightage is configured with value of 100%	
Egress Link Bandwidth utilization	Quality Band
21-30	6
31-40	5
41-50	4
51-60	3
61-70	2
71-80	1
Greater than 80	0 (Worst)

Default

When selective dynamic load balancing is configured, the default ratio of port load metric to port queue rate weightage is 50:50 for per-packet mode and flowlet-mode. The default values of minimum threshold value and maximum threshold value are:

Per-packet mode

- EgressBytes Min Threshold - 80
- EgressBytes Max Threshold - 99

Flowlet mode

- EgressBytes Min Threshold - 20
- EgressBytes Max Threshold - 50

Options

- apply-groups
- max
- min
- rate-weightage

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 24.1

ei-mapping-timeout

IN THIS SECTION

- Syntax | 6241
- Hierarchy Level | 6242
- Description | 6242
- Options | 6242
- Required Privilege Level | 6242
- Release Information | 6242

Syntax

```
mapping-timeout seconds;
```

Hierarchy Level

```
[edit services nat pool nat-pool-name]
```

Description

Specify the duration for endpoint independent translations that use the specified NAT pool. This includes endpoint-independent mapping (EIM) and endpoint-independent filtering (EIF).

Options

seconds—Lifetime of endpoint independent mappings in seconds.

- **Default:** 300
- **Range:** 120 through 864,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

`ei-mapping-timeout` statement introduced in JUNOS Releases 12.3.

RELATED DOCUMENTATION

Network Address Translation Configuration Overview

ei-mapping-timeout (Source NAT Next Gen Services)

IN THIS SECTION

Syntax | 6243

- Hierarchy Level | 6243
- Description | 6243
- Options | 6243
- Required Privilege Level | 6243
- Release Information | 6244

Syntax

```
ei-mapping-timeout ei-mapping-timeout;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name]
```

Description

Specify the timeout period for endpoint independent translations that use the NAT pool. Mappings that are inactive for this amount of time are dropped.

If you do not configure `ei-mapping-timeout` for endpoint independent translations, then the `mapping-timeout` value is used for endpoint independent translations.

Options

ei-mapping-timeout *ei-mapping-timeout*

The timeout period in seconds.

- **Range:** 120 through 86,400
- **Default:** 300 (timeout period for endpoint independent translations is set by `mapping-timeout` value at the `[edit services nat source pool nat-pool-name]` hierarchy level)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

eif-flow-limit

IN THIS SECTION

- Syntax | 6244
- Hierarchy Level | 6244
- Description | 6244
- Required Privilege Level | 6245
- Release Information | 6245

Syntax

```
eif-flow-limit number-of-flows
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated secure-nat-mapping]
```

Description

Specify the maximum number of inbound flows allowed on EIF mapping to the configured value. This limit is per EIF mapping and is per given instance of time. For example, if `eif-flow-limit` is configured as n, then only n inbound connections are allowed at a given instance of time, The n+1 and subsequent connections arriving when n connections are alive are dropped . A new inbound connection is allowed only when one of the n connections times out or is closed. This limit is applied for all type of flows.

Starting in Junos OS Release 15.1R3, `eif-flow-limit` is also supported on the MS-MPC and MS-MIC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3

RELATED DOCUMENTATION

Protecting CGN Devices Against Denial of Service (DOS) Attacks

em-hw-profile

IN THIS SECTION

- Syntax | 6245
- Hierarchy Level | 6245
- Description | 6246
- Default | 6246
- Required Privilege Level | 6246
- Release Information | 6246

Syntax

```
em-hw-profile;
```

Hierarchy Level

```
[edit chassis forwarding-options]
```

Description

(QFX5120 only) Configure a unified forwarding table (UFT) profile to allocate the amount of MAC address, layer 3 host, longest prefix match (LPM), and exact-match (EM) memory available for software-based flow-based telemetry (FBT) for VXLANs. You must reboot the switch after committing this statement to the configuration to create the UFT profile. The software allocates these amounts of memory in the UFT for each of the following types, in bytes:

- MAC address: 48K
- Layer 3 host address: 48K
- LPM memory: 16K
- EM memory: 32K

Default

The UFT uses the default profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Flow-Based Telemetry for VXLANs

email-notify

IN THIS SECTION

- Syntax | 6247
- Hierarchy Level | 6247
- Description | 6247
- Options | 6247
- Required Privilege Level | 6248
- Release Information | 6248

Syntax

```
email-notify {
    admin-email email-address;
    custom-message message;
    custom-message-subject message-subject;
}
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update]
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update]
[edit security utm feature-profile anti-virus sophos-engine pattern-update]
[edit security utm default-configuration anti-virus avira-engine pattern-update]
```

Description

You can configure the device to notify a specified administrator when patterns are updated. This is an e-mail notification with a custom message and a custom subject line.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for Avira engine added in Junos OS Release 18.4R1.

embedded-rp

IN THIS SECTION

- Syntax | 6248
- Hierarchy Level | 6249
- Description | 6249
- Required Privilege Level | 6249
- Release Information | 6249

Syntax

```
embedded-rp {
    group-ranges {
        destination-ip-prefix</prefix-length>;
    }
    maximum-rps limit;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Configure properties for embedded IP version 6 (IPv6) RPs.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring PIM Embedded RP for IPv6

enable

IN THIS SECTION

- Syntax | 6250
- Hierarchy Level | 6250
- Description | 6250

- Options | 6250
- Required Privilege Level | 6250
- Release Information | 6250

Syntax

```
enable service-name {
    concurrent-data-sessions max-session-number;
}
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name]
```

Description

Enable the service name for the subscriber profile.

Options

service-name—Name of the enabled service.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

enable

IN THIS SECTION

- Syntax | 6251
- Hierarchy Level | 6251
- Description | 6251
- Options | 6251
- Required Privilege Level | 6252
- Release Information | 6252

Syntax

```
enable {
    download-timeout minutes;
    interval (Security IDP) hours;
    start-time (Security IDP) start-time;
}
```

Hierarchy Level

```
[edit security idp security-package automatic]
```

Description

Enables the automatic download of the IDP security package.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

enable-all-ifl

IN THIS SECTION

- Syntax | 6252
- Hierarchy Level | 6252
- Description | 6252
- Required Privilege Level | 6253
- Release Information | 6253

Syntax

```
enable-all-ifl;
```

Hierarchy Level

```
[edit protocols layer2-control mac-rewrite interface interface-name]
```

Description

Enable tunneling for STP, VTP, CDP, and other supported protocols on all logical interfaces (VLANs) configured on the interface.

NOTE: Tunneling on all logical interfaces is enabled automatically for PVST/PVST+.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Layer 2 Protocol Tunneling

protocol | 13424

enable-all-qmodules

IN THIS SECTION

- Syntax | 6253
- Hierarchy Level | 6254
- Description | 6254
- Required Privilege Level | 6254
- Release Information | 6254

Syntax

```
(enable-all-qmodules | no-enable-all-qmodules);
```

Hierarchy Level

```
[edit security idp sensor-configuration global]
```

Description

Enable all the qmodules of the global rulebase IDP security policy. By default all the qmodules are enabled.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

enable-asymmetric-traffic-processing (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 6254
- Hierarchy Level | 6255
- Description | 6255
- Required Privilege Level | 6255
- Release Information | 6255

Syntax

```
enable-asymmetric-traffic-processing;
```

Hierarchy Level

```
[edit services service-set service-set-options]
```

Description

Enable the service set to handle unidirectional traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

enable-call-distribution

IN THIS SECTION

- Syntax | 6255
- Hierarchy Level | 6256
- Description | 6256
- Required Privilege Level | 6256
- Release Information | 6256

Syntax

```
enable-call-distribution;
```

Hierarchy Level

```
[edit security alg sip ]
```

Description

When you configure this mode, the call SPU distributes the SIP calls according to the call-id of each SIP call instead of proxy IP.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

sip (Security) | 27977

enable-descriptive-session-syslog

IN THIS SECTION

- Syntax | 6257
- Hierarchy Level | 6257
- Description | 6257
- Required Privilege Level | 6257
- Release Information | 6257

Syntax

```
enable-descriptive-session-syslog;
```

Hierarchy Level

```
edit services service-set service-set-options
```

Description

Display descriptive information of session logs.

You can enable syslog to display information related to inside and outside packets, byte count, and the session id for both open and close sessions.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3.

RELATED DOCUMENTATION

service-set-options | 16746

enable (Enhanced Subscriber Management)

IN THIS SECTION

Syntax | 6258

- Hierarchy Level | 6258
- Description | 6258
- Required Privilege Level | 6258
- Release Information | 6258

Syntax

```
enable (force);
```

Hierarchy Level

```
[edit system services subscriber-management]
```

Description

Enable the Junos OS enhanced subscriber management software architecture to support configuration of dynamic interfaces and services for subscriber management. To use dynamic profiles and authentication to create and manage dynamic subscriber interfaces and services, you must enable enhanced subscriber management and reboot the router.

Use the `force` option to enable packet triggered recovery support for static VLAN subscribers.

Options
force – Enable packet triggered recovery support for static VLAN subscribers.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

Junos OS Enhanced Subscriber Management Overview

enable-explorer (REST API)

IN THIS SECTION

- Syntax | 6259
- Hierarchy Level | 6259
- Description | 6259
- Required Privilege Level | 6259
- Release Information | 6260

Syntax

```
enable-explorer;
```

Hierarchy Level

```
[edit system services rest]
```

Description

Enable the REST API Explorer. This GUI is disabled by default, and can be enabled by specifying `set system services rest enable-explorer`. To disable the REST API Explorer, specify `delete system services rest enable-explorer`. To learn more about the REST API Explorer, see *Example: Using the REST API Explorer*.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Example: Using the REST API Explorer

enable-extended-ql-tlv

IN THIS SECTION

- Syntax | 6260
- Hierarchy Level | 6260
- Description | 6261
- Options | 6261
- Required Privilege Level | 6262
- Release Information | 6262

Syntax

```
enable-extended-ql-tlv;
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Enable enhanced synchronous Ethernet equipment clock (eEEC). eEEC supports new clocks with different Quality Levels (QLs) in the Synchronous Ethernet chain and transmits ESMC with extended QL TLV.

An extended QL TLV for eEEC provides the following three combinations of clocks:

- eEEC with extended QL TLV support
- EEC with no extended QL TLV support
- EEC with extended QL TLV support.

The extended QL TLV allows the count of the number of cascaded eEEC and EEC clocks. This also allows clocks that support enhanced ESMC messages to report the presence of clocks with discarded TLVs.

Options

quality-level

The available quality levels are as follows:

Table 33: Quality Levels

Quality Level	Description
PRTC	Timing quality of a primary reference time clock (option-1/option-2).
eEEC	Timing quality of an enhanced Ethernet equipment clock (option-1/option-2).
ePRC	Timing quality of an enhanced primary reference clock (option-1/option-2).
ePRTC	Timing quality of an enhanced primary reference time clock (option-1/option-2).
prc	Timing quality of a primary reference clock (option-1 only).
prs	Clock traceable to a primary reference source (option-2 only).

Table 33: Quality Levels *(Continued)*

Quality Level	Description
sec	Timing quality of an SDH equipment clock (option-1 only).
smc	Clock traceable to a self-timed SONET (option-2 only).
ssu-a	Timing quality of a type I or IV client clock (option-1 only).
ssu-b	Timing quality of a type VI client clock (option-1 only).
st2	Clock traceable to Stratum 2 (option-2 only).
st3	Clock traceable to Stratum 3 (option-2 only).
st3e	Clock traceable to Stratum 3E (option-2 only).
st4	Clock traceable to Stratum 4 free-run (option-2 only).
stu	Clock traceable to an unknown quality (option-2 only).
tnc	Clock traceable to a transit node clock (option-2 only).

If you do not configure the quality level and the clock source does not receive ESMC messages, the quality level is set to DNU for network option-1 and DUS for network option-2.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 EVO.

RELATED DOCUMENTATION

synchronization (ACX Series)

show chassis synchronization

enable-flow-tracing (Services)

IN THIS SECTION

- Syntax | 6263
- Hierarchy Level | 6263
- Description | 6263
- Required Privilege Level | 6264
- Release Information | 6264

Syntax

```
enable-flow-tracing;
```

Hierarchy Level

```
[edit services ssl proxy profile profile-name]
[edit services ssl termination profile profile-name]
[edit services ssl initiation profile profile-name]
```

Description

Enable flow tracing for the profile.

When you configure `enable-flow-tracing` for SSL profiles, the debug tracing will be enabled on that profile when the flag is set as `selected-profile`.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. This statement is supported.

RELATED DOCUMENTATION

SSL Proxy Overview

Configuring SSL Forward Proxy

Enabling Debugging and Tracing for SSL Proxy

enable-gtpu-distribution

IN THIS SECTION

- Syntax | 6264
- Hierarchy Level | 6265
- Description | 6265
- Required Privilege Level | 6265
- Release Information | 6265

Syntax

```
enable-gtpu-distribution;
```

Hierarchy Level

```
[edit security forwarding-process application-services]
```

Description

Use the enable-gtpu-distribution option to enable GTP control plane (GTP-C) and GTP user plane (GTP-U) session distribution. This helps to distribute the GTP-C and GTP U traffic on all Services Processing Units (SPUs) by using tunnel-based session distribution, where the GTP-C or the GTP-U traffic of different tunnels is spread across different SPUs.

This option is not enabled by default. You need to enable this option to distribute GTP-U and GTP-C traffic on all SPUs.

If this option is not configured, the GTP ALG will not work properly. GTP-C packets will be processed as normal UDP traffic.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

Understanding GTP Support for Central Point Architecture

enable-ipv6-services-for-lac (L2TP)

IN THIS SECTION

Syntax | 6266

- Hierarchy Level | 6266
- Description | 6266
- Default | 6266
- Required Privilege Level | 6266
- Release Information | 6266

Syntax

```
enable-ipv6-services-for-lac;
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Enable the LAC to create the IPv6 address family (inet6) when establishing a tunnel for subscribers, allowing IPv6 filters to be applied. By default, the LAC requires only family inet to enable forwarding into an IP tunnel. It can apply IPv4 firewall filters to the session. Even when family inet6 is included in the dynamic profile, by default it is not created and IPv6 firewall filters cannot be applied.

Default

Disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

L2TP LAC Subscriber Configuration

enable-mac-move-action

IN THIS SECTION

- Syntax | 6267
- Hierarchy Level | 6267
- Description | 6267
- Required Privilege Level | 6267
- Release Information | 6268

Syntax

```
enable-mac-move-action;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name]
```

Description

Enable the MAC move action feature at the bridge domain level. This statement blocks the logical interface for the bridge domain when a MAC move is detected on the interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Statement supported in Junos OS Release 14.2 for MX104 Router.

RELATED DOCUMENTATION

Configuring a Bridge Domain

enable-packet-pool

IN THIS SECTION

- Syntax | 6268
- Hierarchy Level | 6268
- Description | 6268
- Required Privilege Level | 6269
- Release Information | 6269

Syntax

```
(enable-packet-pool | no-enable-packet-pool);
```

Hierarchy Level

```
[edit security idp sensor-configuration global]
```

Description

Enable the packet pool to use when the current pool is exhausted. By default packet pool is enabled.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

enable-performance-mode

IN THIS SECTION

- Syntax | 6269
- Hierarchy Level | 6269
- Description | 6269
- Options | 6270
- Required Privilege Level | 6270
- Release Information | 6270

Syntax

```
enable-performance-mode max-packet-threshold number;
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Set the deep packet inspection (DPI) in performance mode for application identification.

The application traffic throughput can be improved by setting the DPI in performance mode with default packet inspection limit as two packets, including both client-to-server and server-to-client directions. By default, performance mode is disabled on SRX Series Firewalls.

If you want to set DPI to default accuracy mode and disable the performance mode, delete the configuration statement that specifies enabling of the performance mode by using the `delete services application-identification enable-performance-mode` command.

Starting in Junos OS Release 15.1X49-D210R1 and Junos OS Release 19.4R1, the maximum packet threshold for DPI performance mode option is deprecated—rather than immediately removed—to provide backward compatibility and an 6 opportunity to bring your configuration into compliance with the new configuration.

Options

max-packet-threshold *number* Set the maximum packet threshold for DPI performance mode.

- **Range:** 1 through 100
- **Default:** 2

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Improving the Application Traffic Throughput

show services application-identification status | **25456**

enable-performance-mode (Application Identification)

IN THIS SECTION

- Syntax | 6271
- Hierarchy Level | 6271
- Description | 6271
- Required Privilege Level | 6271
- Release Information | 6271

Syntax

```
enable-performance-mode;
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Set the deep packet inspection (DPI) in performance mode for application identification. Performance mode improves application traffic throughput by limiting the maximum DPI processing to four packets per session. This limit includes both client-to-server and server-to-client directions. By default, performance mode is disabled.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Application Identification Overview

enable-rejoin (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 6272
- Hierarchy Level | 6272
- Description | 6272
- Default | 6273
- Required Privilege Level | 6273
- Release Information | 6273

Syntax

```
enable-rejoin;
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options member-failure-options redistribute-all-traffic]
```

Description

Enable the failed member to rejoin the aggregated Multiservices (AMS) interface after the member comes back online.

For many-to-one (N:1) high availability (HA) for service applications like Network Address Translation (NAT), this configuration allows the failed members to rejoin the pool of active members automatically.

Default

If you do not configure this option, then the failed members do not automatically rejoin the `ams` interface even after coming back online.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

redistribute-all-traffic (Aggregated Multiservices) | **14165**

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

enable-reroute-uniform-link-check

IN THIS SECTION

- Syntax | **6274**
- Hierarchy Level | **6274**
- Description | **6274**
- Required Privilege Level | **6274**
- Release Information | **6274**

Syntax

```
enable-reroute-uniform-link-check nat;
```

Hierarchy Level

[set security flow]

Description

Enable retaining an existing session with Network Address Translation (NAT) rule when there is a change in egress interface because of rerouting.

The `enable-reroute-uniform-link-check nat` command is disabled by default.

When the `enable-reroute-uniform-link-check nat` command is enabled:

- If the new egress interface and the previous egress interface are in the same security zone and there is no change in the matched NAT rule or if no rule is applied before and after rerouting, the session is retained with the existing NAT rule.
- If the new egress interface and the previous egress interface are in the same security zone and the matched NAT rule is changed, the session expires.

When the `enable-reroute-uniform-link-check nat` command is disabled:

- If the new egress interface and the previous egress interface are in the same security zone, the traffic is forwarded to the new egress interface.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding NAT Configuration Check on Egress Interfaces after Reroute

enable-reverse-reroute

IN THIS SECTION

- Syntax | 6275
- Hierarchy Level | 6275
- Description | 6275
- Required Privilege Level | 6276
- Release Information | 6276

Syntax

```
enable-reverse-reroute;
```

Hierarchy Level

[edit security zones security-zone *zone-name*]

Description

Reroute the reverse traffic when there is a link switch for the incoming traffic.

When you configure the `enable-reverse-reroute` option for a security zone, then the packets of each session that has been initiated from the zone are checked for the change in the incoming interface. When an incoming packet arrives on an interface that is different from the one cached in session, the route lookup is performed for the reverse path, and the preference is given to the interface on which the packet has arrived when there are ECMP routes available to the source. Ensure that when you configure enable-reverse-reroute option, the new interface on which packets arrive must be part of the same zone as the earlier interface.

You can enable reverse rerouting in hub-and-spoke deployments, where a spoke device uses APBR to re-route the traffic based on the dynamic applications. In such cases reverse re-route can be used on hub device to correctly re-route the reverse traffic.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D123.

RELATED DOCUMENTATION

Understanding Advanced Policy-Based Routing

enable-session-cache

IN THIS SECTION

- Syntax | 6276
- Hierarchy Level | 6276
- Description | 6277
- Required Privilege Level | 6277
- Release Information | 6277

Syntax

```
enable-session-cache;
```

Hierarchy Level

```
[edit services ssl termination profile profile-name]
[edit services ssl initiation profile profile-name]
```

Description

Enable SSL session cache.

You can enable session caching to cache session information, such as the pre-master secret key and agreed-upon ciphers, for both the client and server.

The cached information is identified by a session ID. In subsequent connections both parties agree to use the session ID to retrieve the information rather than create a new pre-master secret key. Session resumption shortens the handshake process and accelerates SSL transactions there by improves the throughput and maintains an appropriate level of security at the same time.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. This statement is supported.

RELATED DOCUMENTATION

SSL Proxy Overview

Configuring SSL Forward Proxy

Enabling Debugging and Tracing for SSL Proxy

enable-snmp-tunnel-statistics (L2TP)

IN THIS SECTION

- Syntax | 6278
- Hierarchy Level | 6278
- Description | 6278
- Default | 6278

Required Privilege Level | 6278
Release Information | 6278

Syntax

```
enable-snmp-tunnel-statistics;
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Enable collection of L2TP tunnel and global counters for SNMP statistics.

> **NOTE**: The system load can increase when you enable these counters and also use RADIUS interim accounting updates. We recommend you enable these counters when you are using only SNMP statistics.

Default

Disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1R4 and supported in later 12.1Rx releases.

Statement supported in Junos OS Release 12.2R2 and later 12.2Rx releases. (Not supported in Junos OS Release 12.2R1.)

Statement supported in Junos OS Release 12.3 and later releases.

RELATED DOCUMENTATION

L2TP for Subscriber Access Overview

enable-subscriber-analysis (Services Options VMS Interfaces)

IN THIS SECTION

- Syntax | 6279
- Hierarchy Level | 6279
- Description | 6279
- Required Privilege Level | 6280
- Release Information | 6280

Syntax

```
enable-subscriber-analysis;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Description

Enable the creation of subscribers if the following are not configured, but you want subscribers to be created:

- NAT

- The `max-sessions-per-subscriber` statement at the `[edit services service-set service-set-name]` hierarchy level

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

How to Configure Services Interfaces for Next Gen Services

encap-type

IN THIS SECTION

- Syntax | 6280
- Hierarchy Level | 6281
- Description | 6281
- Options | 6281
- Required Privilege Level | 6281
- Release Information | 6281

Syntax

```
encap-type(ipv4 | ipv6);
```

Hierarchy Level

```
[edit interfaces name unit name tunnel]
```

Description

Define the type of encapsulation on the IP tunnel interface .

You must configure this option if you configure `source-interface` or `source-hostname` or `destination-hostname`.

Options

ipv4 For IPv4 tunnel encapsulation

ipv6 For IPv6 tunnel encapsulation

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 for SRX Series Firewalls (SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, SRX5800, and vSRX3.0)

Statement introduced in Junos OS Release 23.4R1 for SRX Series Firewalls (SRX38, SRX300, SRX320, SRX340, and SRX345)

RELATED DOCUMENTATION

Tunnel Services Overview

encapsulate-inner-vlan

IN THIS SECTION

- Syntax | 6282
- Hierarchy Level | 6282
- Description | 6282
- Default | 6283
- Required Privilege Level | 6283
- Release Information | 6283

Syntax

```
encapsulate-inner-vlan
```

Hierarchy Level

```
[edit routing-instances routing-instance-name vlans vlan-name vxlan],
[edit vlans vlan-name vxlan]
```

Description

Configure the switch to preserve the original VLAN tag (in the inner Ethernet packet) when performing Virtual Extensible LAN (VXLAN) encapsulation.

Preserving the original VLAN tag is implicit with VLAN bundle services. As a result, in an EVPN-VXLAN environment, you don't need to configure this option with EVPN MAC-VRF instance `vlan-bundle` service type configurations. Explicitly configure this option with other service types if needed.

NOTE: Starting in Junos OS release 23.2R2, you cannot configure the `encapsulate-inner-vlan` statement in a MAC-VRF routing instance if the vlan has an IRB interface associated with it.

NOTE: Starting in Junos OS release 22.3R1, IS-IS logical link control (LLC) packets that are encapsulated in a VXLAN tunnel now have their VLAN tags removed by default. You can configure this option to preserve the VLAN tags on IS-IS LLC packets.

NOTE: ARP suppression is automatically disabled when you encapsulate VLAN traffic for EVPN-VXLAN by configuring `encapsulate-inner-vlan` and its equivalent converse statement `decapsulate-accept-inner-vlan`. As a result, you might observe ARP packets being sent by the device.

Default

When the device encapsulates the packet, it drops the original VLAN tag.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Understanding VXLANs

Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network

Manually Configuring VXLANs on QFX Series and EX4600 Switches

Examples: Manually Configuring VXLANs on QFX Series and EX4600 Switches

decapsulate-accept-inner-vlan | **5158**

decapsulate-accept-inner-vlan | **5158**

encapsulation (protocols evpn)

IN THIS SECTION

- Syntax | 6284
- Hierarchy Level | 6284
- Description | 6284
- Options | 6284
- Required Privilege Level | 6286
- Release Information | 6286

Syntax

```
encapsulation (mpls | srv6 | vxlan);
```

Hierarchy Level

```
[edit protocols evpn],
[edit protocols evpn interconnect],
[edit routing-instances routing-instance-name protocols evpn],
[edit routing-instances routing-instance-name protocols evpn interconnect]
```

Description

Configure an encapsulation type for EVPN features.

Options

encapsulation mpls

Use the `encapsulation mpls` statement with the `interconnected-vlan-list` statement as part of the "interconnect" on page 8797 configuration stanza for seamless stitching of EVPN-VXLAN data centers through EVPN-MPLS.

You configure both `vxlan` and `mpls` encapsulations as follows:

- EVPN-VXLAN encapsulation going to the data center (DC) side
- EVPN-MPLS encapsulation on the WAN Data Center Interconnect (DCI) side under the `interconnect` configuration hierarchy.

> **NOTE**: Configure the `encapsulation mpls` statement with the `interconnected-vlan-list` statement for DCI implementations.

This sample configuration defines which VLANs to interconnect from the EVPN-VXLAN domain to the EVPN-MPLS domain.

```
interconnect {
               vrf-target target:2:2;
               route-distinguisher 100:120;
               esi 00:0a:0b:0c:0d:0a:0b:0c:0d:0a; {
                 all-active;
               }
               interconnected-vlan-list [51 52];
               encapsulation mpls;
            }
```

encapsulation srv6

The `encapsulation srv6` statement is used in EVPN-VPWS and EVPN E-LAN services over an SRv6 underlay network.

> **NOTE**: `encapsulation srv6` supports only pure Type 5 routes for EVPN E-LAN over SRv6.

encapsulation vxlan

The `encapsulation vxlan` statement is required for a VXLAN EVPN instance.

> **NOTE**: If you configure the `encapsulation vxlan` statement, then you must also configure the `extended-vni-list` statement.

> **NOTE**: The `encapsulation vxlan` statement is an exclusive command. You cannot configure the `encapsulation vxlan` statement with the `extended-vlan-list` statement, or other commands associated with MPLS EVPN instances.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

The encapsulation vxlan statement was introduced in Junos OS Release 16.1.

The encapsulation mpls statement was added to the interconnect stanza in Junos OS Release 21.4R1.

The encapsulation srv6 statement was introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

Understanding EVPN with VXLAN Data Plane Encapsulation

EVPN Over VXLAN Encapsulation Configuration Overview for QFX Series and EX4600 Switches

Example: Configuring an EVPN Control Plane and VXLAN Data Plane

Implementing EVPN-VXLAN for Data Centers

Configuring EVPN-VPWS over SRv6

encapsulation (logical-link-layer)

IN THIS SECTION

- Syntax | 6287
- Hierarchy Level | 6287
- Description | 6287
- Options | 6287
- Required Privilege Level | 6287
- Release Information | 6287

Syntax

```
encapsulation type;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Specify the logical link-layer encapsulation type.

Options

`atm-mlppp-llc`—For ATM2 IQ physical interfaces only, use Multilink Point-to-Point Protocol (MLPPP) over AAL5 LLC encapsulation.

`frame-relay-ppp`—For Frame Relay circuits, use Frame Relay PPP encapsulation.

`multilink-ppp`—By default, voice services logical interfaces use MLPPP encapsulation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Encapsulation for Voice Services

Junos OS Network Interfaces Library for Routing Devices

encapsulation (interfaces-fti)

IN THIS SECTION

- Syntax | 6288
- Hierarchy Level | 6289
- Description | 6289
- Required Privilege Level | 6289
- Release Information | 6289

Syntax

```
encapsulation {
    ( gre | ipip | udp )
        destination {
            addressaddress;
        }
        key key;
        source {
            addressaddress;
        }
        tunnel-routing-instance {
            routing-instance routing-instance;
            port-profile profile-name
        }
        tunnel-termination;
    }
     {
        destination {
            address address;
        }
        source {
            address address;
            interface interface;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit interfaces fti number unit logical-unit-number tunnel]
```

Description

Configure encapsulation over tunnel.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Evolved Release 20.2R1.

Statement introduced for IPIP tunnels in Junos OS Evolved Release 20.4R1.

Statement introduced for UDP tunnels in Junos OS Evolved Release 21.1R1.

Statement introduced in Junos OS Evolved 24.1R1 for ACX Series routers (GRE tunnels)

encapsulation (interfaces)

IN THIS SECTION

- Syntax | 6290
- Syntax (SRX Series Physical Interfaces) | 6290
- Syntax (SRX Series Logical Interfaces) | 6290
- Hierarchy Level (Physical Interfaces) | 6291
- Hierarchy Level (Logical Interfaces) | 6291

- Description | 6291
- Default | 6291
- Physical Interface Options and Logical Interface Options | 6291
- Required Privilege Level | 6296
- Release Information | 6296

Syntax

```
encapsulation (atm-ccc-cell-relay | atm-pvc | cisco-hdlc | cisco-hdlc-ccc | cisco-hdlc-tcc |
ethernet-bridge | ethernet-ccc | ethernet-over-atm | ethernet-tcc | ethernet-vpls | ethernet-
vpls-fr | ether-vpls-over-atm-llc | ethernet-vpls-ppp | extended-frame-relay-ccc | extended-
frame-relay-ether-type-tcc | extended-frame-relay-tcc | extended-vlan-bridge | extended-vlan-ccc
| extended-vlan-tcc | extended-vlan-vpls | flexible-ethernet-services | flexible-frame-relay |
frame-relay | frame-relay-ccc | frame-relay-ether-type | frame-relay-ether-type-tcc | frame-
relay-port-ccc | frame-relay-tcc | generic-services | multilink-frame-relay-uni-nni | ppp | ppp-
ccc | ppp-tcc | vlan-ccc | vlan-vci-ccc | vlan-vpls);
```

Syntax (SRX Series Physical Interfaces)

```
encapsulation (ether-vpls-ppp | ethernet-bridge | ethernet-ccc | ethernet-tcc | ethernet-vpls |
extended-frame-relay-ccc | extended-frame-relay-tcc | extended-vlan-bridge | extended-vlan-ccc |
extended-vlan-tcc | extended-vlan-vpls | flexible-ethernet-services  | frame-relay-port-ccc |
vlan-ccc | vlan-vpls);
```

Syntax (SRX Series Logical Interfaces)

```
encapsulation ( dix | ether-vpls-fr | frame-relay-ppp | ppp-over-ether | vlan-bridge | vlan-ccc
| vlan-tcc | vlan-vpls );
```

Hierarchy Level (Physical Interfaces)

```
[edit interfaces interface-name]
[edit interfaces rlsq number:number]
```

Hierarchy Level (Logical Interfaces)

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Specify the physical link layer encapsulation type. For some devices, you may also specify the logical link layer encapsulation type. See Feature Explorer for more information.

> **NOTE**: Not all encapsulation types are supported on all switches. See the switch CLI.

Default

`ppp`—Use serial PPP encapsulation.

Physical Interface Options and Logical Interface Options

For physical interfaces:

> **NOTE**: EX Series switches do not support Frame Relay, ATM, PPP, SONET, or SATSOP encapsulation.
>
> ACX Series routers do not support CISCO HDLC encapsulation.

- `atm-ccc-cell-relay`—Use ATM cell-relay encapsulation.
- `atm-pvc`—Defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. When you configure physical ATM interfaces with ATM PVC encapsulation, an RFC 2684-compliant ATM Adaptation Layer 5 (AAL5) tunnel is set up to route the ATM cells over a Multiprotocol Label Switching (MPLS) path that is typically established between two MPLS-capable routers using the Label Distribution Protocol (LDP).

- `cisco-hdlc`—Use Cisco-compatible High-Level Data Link Control (HDLC) framing. E1, E3, SONET/SDH, T1, and T3 interfaces can use Cisco HDLC encapsulation. Two related versions are supported:
  - CCC version (`cisco-hdlc-ccc`)—The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the `ccc` family only.
  - TCC version (`cisco-hdlc-tcc`)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- `cisco-hdlc-ccc`—Use Cisco-compatible HDLC framing on CCC circuits.
- `cisco-hdlc-tcc`—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.
- `ethernet-bridge`—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.
- `ethernet-over-atm`—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.
- `ethernet-tcc`—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.
- `ethernet-vpls`—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.
- `ethernet-vpls-fr`—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 Frame Relay connection, use the *802.1p* bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.
- `ethernet-vpls-ppp`—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the

packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.

- `ether-vpls-over-atm-llc`—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

- `extended-frame-relay-ccc`—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC. When you use this encapsulation type, you can configure the `ccc` family only.

- `extended-frame-relay-ether-type-tcc`—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.

- `extended-frame-relay-tcc`—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

- `extended-vlan-bridge`—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE *802.1Q* VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

- `extended-vlan-ccc`—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. Extended VLAN CCC encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. When you use this encapsulation type, you can configure the `ccc` family only. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.

- `extended-vlan-tcc`—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

- `extended-vlan-vpls`—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

> **NOTE**: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

- `flexible-ethernet-services`—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and for Gigabit Ethernet interfaces, use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

- `flexible-frame-relay`—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.

- `frame-relay`—Use Frame Relay encapsulation is defined in RFC 1490, *Multiprotocol Interconnect over Frame Relay*. E1, E3, link services, SONET/SDH, T1, T3, and voice services interfaces can use Frame Relay encapsulation.

- `frame-relay-ccc`—Use Frame Relay encapsulation on CCC circuits. This encapsulation is same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. The logical interface must also have `frame-relay-ccc` encapsulation. When you use this encapsulation type, you can configure the `ccc` family only.

- `frame-relay-ether-type`—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay. IETF frame relay encapsulation identifies the payload format using NLPID and SNAP formats. Cisco-compatible Frame Relay encapsulation uses the Ethernet type to identify the type of payload.

  > **NOTE**: When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.

- `frame-relay-ether-type-tcc`—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. This encapsulation is Cisco-compatible Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to TCC.

- `frame-relay-port-ccc`—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. The connection between the two CE routers can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. When you use this encapsulation type, you can configure the `ccc` family only.

- `frame-relay-tcc`—This encapsulation is similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- `generic-services`—Use generic services encapsulation for services with a hierarchical scheduler.
- `multilink-frame-relay-uni-nni`—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.
- `ppp`—Use serial PPP encapsulation. This encapsulation is defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. E1, E3, SONET/SDH, T1, and T3 interfaces can use PPP encapsulation.
- `ppp-ccc`—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the `ccc` family only.
- `ppp-tcc`—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the `tcc` family only.
- `vlan-ccc`—Use Ethernet VLAN encapsulation on CCC circuits. VLAN CCC encapsulation supports TPID 0x8100 only. When you use this encapsulation type, you can configure the `ccc` family only.
- `vlan-vci-ccc`—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the `ccc` family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to `vlan-vci-ccc`.
- `vlan-vpls`—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

> **NOTE**:
>
> - Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.
> - Starting with Junos OS release 13.3, a commit error occurs when you configure `vlan-vpls` encapsulation on a physical interface and configure `family inet` on one of the logical units. Previously, it was possible to commit this invalid configuration.

For logical interfaces:

- `frame-relay`—Configure a Frame Relay encapsulation when the physical interface has multiple logical units, and the units are either point to point or multipoint.
- `multilink-frame-relay-uni-nni`—Link services interfaces functioning as FRF.16 bundles can use Multilink Frame Relay UNI NNI encapsulation.
- `ppp`—For normal mode (when the device is using only one ISDN B-channel per call). Point-to-Point Protocol is for communication between two computers using a serial interface.
- `ppp-over-ether`—This encapsulation is used for underlying interfaces of pp0 interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Physical Encapsulation on an Interface

Configuring Interface Encapsulation on Physical Interfaces

Configuring CCC Encapsulation for Layer 2 VPNs

Configuring Layer 2 Switching Cross-Connects Using CCC

Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits

Configuring ATM Interface Encapsulation

Configuring ATM-to-Ethernet Interworking

Configuring VLAN and Extended VLAN Encapsulation

Configuring VLAN and Extended VLAN Encapsulation

Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces

Configuring Interfaces for Layer 2 Circuits

Configuring Interface Encapsulation on PTX Series Packet Transport Routers

Configuring MPLS LSP Tunnel Cross-Connects Using CCC

Configuring TCC

Configuring VPLS Interface Encapsulation

Configuring Interfaces for VPLS Routing

Defining the Encapsulation for Switching Cross-Connects

Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)

encapsulation (Dynamic Interfaces)

IN THIS SECTION

- Syntax | 6297
- Hierarchy Level | 6297
- Description | 6297
- Options | 6298
- Required Privilege Level | 6300
- Release Information | 6300

Syntax

```
encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux | atm-
mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-vc-mux |
ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp |
ethernet | frame-relay-ccc | frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type | frame-
relay-ether-type-tcc | multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether | ppp-
over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

Description

Dynamic interface configuration of the logical link-layer encapsulation type.

Options

`atm-ccc-cell-relay`—Use ATM cell-relay encapsulation.

`atm-ccc-vc-mux`—Use ATM virtual circuit (VC) multiplex encapsulation on circuit cross-connect (CCC) circuits. When you use this encapsulation type, you can configure the `ccc` family only.

`atm-cisco-nlpid`—Use Cisco ATM network layer protocol ID (NLPID) encapsulation. When you use this encapsulation type, you can configure the `inet` family only.

`atm-mlppp-llc`—For ATM2 IQ interfaces only, use Multilink Point-to-Point Protocol (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a link services or voice services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.

`atm-nlpid`—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the `inet` family only.

`atm-ppp-llc`—For ATM2 IQ interfaces only, use PPP over AAL5 LLC encapsulation.

`atm-ppp-vc-mux`—For ATM2 IQ interfaces only, use PPP over ATM AAL5 multiplex encapsulation.

`atm-snap`—Use ATM subnetwork attachment point (SNAP) encapsulation.

`atm-tcc-snap`—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

`atm-tcc-vc-mux`—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the `tcc` family only.

`atm-vc-mux`—Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the `inet` family only.

`ether-over-atm-llc`—For interfaces that carry IPv4 traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

`ether-vpls-over-atm-llc`—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

`ether-vpls-over-fr`—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, as per *Multiprotocol Interconnect over Frame Relay*(RFC 2427 [1490]).

`ether-vpls-over-ppp`—For E1, T1, E3, T3 and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over PPP encapsulation to support Bridged Ethernet over PPP encapsulated TDM interfaces for VPLS applications.

`ethernet`—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

`ethernet-vpls`—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.

`extended-vlan-vpls`—Use extended virtual LAN (VLAN) VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901.

> **NOTE**: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

`frame-relay-ccc`—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the `ccc` family only.

`frame-relay-ppp`—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the `ppp` family only.

`frame-relay-tcc`—Use Frame Relay encapsulation on TCC circuits for connecting unlike media. When you use this encapsulation type, you can configure the `tcc` family only.

`frame-relay-ether-type`—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with `flexible-frame-relay` encapsulation.

`frame-relay-ether-type-tcc`—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect unlike media. The physical interface must be configured with `flexible-frame-relay` encapsulation.

`multilink-frame-relay-end-to-end`—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services,and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

`multilink-ppp`—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

`ppp-over-ether`—You use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface.

`vlan-bridge`—Use Ethernet VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q tagging, flexible ethernet services, and bridging enabled, and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

`vlan-ccc`—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the `ccc` family only.

`vlan-vci-ccc`—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the `ccc` family only.

`vlan-tcc`—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the `tcc` family only.

`vlan-vpls`—Use Ethernet VLAN encapsulation on VPLS circuits.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution

Configuring PPP over ATM2 Encapsulation Overview

encapsulation (Logical Interface)

IN THIS SECTION

- Syntax | 6301
- Hierarchy Level | 6301
- Description | 6301
- Options | 6301
- Required Privilege Level | 6304
- Release Information | 6304

Syntax

```
encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-mlppp-llc | atm-nlpid
| atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-tcc-vc-mux | atm-vc-mux | ether-
over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet |
ethernet-ccc | ethernet-vpls | ethernet-vpls-fr | frame-relay-ccc | frame-relay-ether-type |
frame-relay-ether-type-tcc | frame-relay-ppp | frame-relay-tcc | gre-fragmentation | multilink-
frame-relay-end-to-end | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-
bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls | vxlan);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number],
[edit interfaces rlsq number unit logical-unit-number]
[edit protocols evpn]
```

Description

Configure a logical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.

Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms. See Configuring VLAN TCC Encapsulation for more details. Non-ethernet media types, SONET and ATM interfaces are only supported. It is expected that the user will have the member links of aggregated ethernet with supported hardware for configuring VLAN TCC encapsulation and no commit check is performed externally for the aggregated ethernet (AE) interfaces.

Options

`atm-ccc-cell-relay`—Use ATM cell-relay encapsulation.

`atm-ccc-vc-mux`—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the `ccc` family only.

`atm-cisco-nlpid`—Use Cisco ATM network layer protocol identifier (NLPID) encapsulation. When you use this encapsulation type, you can configure the `inet` family only.

`atm-mlppp-llc`—For ATM2 IQ interfaces only, use Multilink Point-to-Point (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.

`atm-nlpid`—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the `inet` family only.

`atm-ppp-llc`—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over AAL5 LLC encapsulation.

`atm-ppp-vc-mux`—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over ATM AAL5 multiplex encapsulation.

`atm-snap`—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM subnetwork attachment point (SNAP) encapsulation.

`atm-tcc-snap`—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

`atm-tcc-vc-mux`—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the `tcc` family only.

`atm-vc-mux`—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the `inet` family only.

`ether-over-atm-llc`—(All IP interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) For interfaces that carry IP traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

`ether-vpls-over-atm-llc`—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

`ether-vpls-over-fr`—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, per RFC 2427, *Multiprotocol Interconnect over Frame Relay*.

NOTE: The SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, and the DS3/E3 MIC do not support Ethernet over Frame Relay encapsulation.

`ether-vpls-over-ppp`—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Point-to-Point Protocol (PPP) encapsulation to support Bridged Ethernet over PPP-encapsulated TDM interfaces for VPLS applications.

`ethernet`—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

`ethernet-ccc`—Use Ethernet CCC encapsulation on Ethernet interfaces.

`ethernet-vpls`—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.

> **NOTE**: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

`ethernet-vpls-fr`—Use in a VPLS setup when a CE device is connected to a PE router over a time-division multiplexing (TDM) link. This encapsulation type enables the PE router to terminate the outer layer 2 Frame Relay connection, use the *802.1p* bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

`frame-relay-ccc`—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the `ccc` family only.

`frame-relay-ether-type`—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with flexible-frame-relay encapsulation.

`frame-relay-ether-type-tcc`—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. The physical interface must be configured with flexible-frame-relay encapsulation.

`frame-relay-ppp`—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the `ppp` family only.

`frame-relay-tcc`—Use Frame Relay encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the `tcc` family only.

`gre-fragmentation`—For adaptive services interfaces only, use GRE fragmentation encapsulation to enable fragmentation of IPv4 packets in GRE tunnels. This encapsulation clears the do not fragment (DF) bit in the packet header. If the packet' s size exceeds the tunnel' s maximum transmission unit (MTU) value, the packet is fragmented before encapsulation.

`multilink-frame-relay-end-to-end`—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

`multilink-ppp`—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

`ppp-over-ether`—Use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface on M120 and M320 routers with Intelligent Queuing 2 (IQ2) PICs, and on MX Series routers with MPCs.

`ppp-over-ether-over-atm-llc`—(MX Series routers with MPCs using the ATM MIC with SFP only) For underlying ATM interfaces, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

`vlan-bridge`—Use Ethernet VLAN bridge encapsulation on Ethernet interfaces that have IEEE *802.1Q* tagging, flexible-ethernet-services, and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

`vlan-ccc`—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the `ccc` family only.

`vlan-vci-ccc`—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the `ccc` family only.

`vlan-tcc`—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the `tcc` family only.

`vlan-vpls`—Use Ethernet VLAN encapsulation on VPLS circuits.

`vxlan`—Use VXLAN data plane encapsulation for EVPN.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
20.1R1	Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms.

RELATED DOCUMENTATION

- Configuring Layer 2 Switching Cross-Connects Using CCC
- Configuring the Encapsulation for Layer 2 Switching TCCs
- Configuring Interface Encapsulation on Logical Interfaces
- Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects
- Circuit and Translational Cross-Connects Overview
- Identifying the Access Concentrator
- Configuring ATM Interface Encapsulation
- Configuring VLAN and Extended VLAN Encapsulation
- Configuring ATM-to-Ethernet Interworking
- Configuring Interface Encapsulation on PTX Series Packet Transport Routers
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- Configuring ATM for Subscriber Access
- Understanding CoS on ATM IMA Pseudowire Interfaces Overview
- Configuring Policing on an ATM IMA Pseudowire

encapsulation-overhead

IN THIS SECTION

- Syntax | **6306**
- Hierarchy Level | **6306**
- Description | **6306**

- Options | 6306
- Required Privilege Level | 6306
- Release Information | 6306

Syntax

```
encapsulation-overhead bytes;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp],
[edit access profile profile-name client client-name ppp]
```

Description

Configure the encapsulation overhead for class-of-service calculations.

Options

bytes—The number of bytes used as encapsulation overhead for the session.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

encapsulation (Physical Interface)

IN THIS SECTION

- Syntax | 6307
- Hierarchy Level | 6307
- Description | 6307
- Default | 6307
- Options | 6308
- Required Privilege Level | 6308
- Release Information | 6308

Syntax

```
encapsulation (multilink-frame-relay-uni-nni | ... );
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces rlsqnumber:number]
```

Description

Physical link-layer encapsulation type.

Default

MLFR UNI NNI encapsulation (on link services interfaces).

Options

`multilink-frame-relay-uni-nni`—Use MLFR UNI NNI encapsulation. This encapsulation is used only on link services interfaces functioning as FRF.16 bundles and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Encapsulation for Link Services Physical Interfaces

Junos OS Network Interfaces Library for Routing Devices

encapsulation-type (Layer 2 Circuits)

IN THIS SECTION

- Syntax | 6309
- Hierarchy Level | 6309
- Description | 6309
- Options | 6309
- Required Privilege Level | 6310
- Release Information | 6310

Syntax

```
encapsulation-type (atm-aal5 | atm-cell | atm-cell-port-mode | atm-cell-vc-mode | atm-cell-vp-
mode | cesop | cisco-hdlc | ethernet | ethernet-vlan | frame-relay | frame-relay-port-mode |
interworking | ppp | satop-e1 | satop-e3 | satop-t1 | satop-t3);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit protocols l2circuit local-switching interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name]
```

Description

Specify the type of Layer 2 traffic transiting the Layer 2 circuit.

Options

`atm-aal5`—ATM Adaptation Layer (AAL/5)

`atm-cell`—ATM cell relay

`atm-cell-port-mode`—ATM cell relay port promiscuous mode

`atm-cell-vc-mode`—ATM VC cell relay nonpromiscuous mode

`atm-cell-vp-mode`—ATM virtual path (VP) cell relay promiscuous mode

`cesop`—CESOP-based Layer 2 circuit

`cisco-hdlc`—Cisco Systems–compatible HDLC

`ethernet`—Ethernet

`ethernet-vlan`—Ethernet VLAN

`frame-relay`—Frame Relay

`frame-relay-port-mode`—Frame Relay port mode

`interworking`—Layer 2.5 interworking

`ppp`—PPP

`satsop-e1`—SATSOP-E1-based Layer 2 circuit

`satsop-e3`—SATSOP-E3-based Layer 2 circuit

`satsop-t1`—SATSOP-T1-based Layer 2 circuit

`satsop-t3`—SATSOP-T3-based Layer 2 circuit

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring the Encapsulation Type for the Layer 2 Circuit Neighbor Interface

encapsulation-type (Layer 2 VPNs)

IN THIS SECTION

- Syntax | 6311
- Hierarchy Level | 6311
- Description | 6311
- Options | 6311
- Required Privilege Level | 6312
- Release Information | 6312

Syntax

```
encapsulation-type (atm-aal5 | atm-cell | atm-cell-port-mode | atm-cell-vc-mode | atm-cell-vp-
mode | cesop | cisco-hdlc | ethernet | ethernet-vlan | frame-relay | frame-relay-port-mode |
interworking | ppp | satop-e1 | satop-e3 | satop-t1 | satop-t3);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit protocols l2circuit neighbor address interface interface-name],
[edit routing-instances routing-instance-name protocols l2vpn],
[edit routing-instances routing-instance-name protocols l2vpn neighbor address],
[edit routing-instances routing-instance-name protocols vpls],
[edit routing-instances routing-instance-name protocols vpls neighbor address]
```

Description

Specify the type of Layer 2 traffic originating from the CE device. Only the `ethernet` and `ethernet-vlan` encapsulation types are supported for VPLS. Not all encapsulation types are supported on the switches. See the switch CLI.

Options

`atm-aal5`—ATM Adaptation Layer (AAL/5)

`atm-cell`—ATM cell relay

`atm-cell-port-mode`—ATM cell relay port promiscuous mode

`atm-cell-vc-mode`—ATM VC cell relay nonpromiscuous mode

`atm-cell-vp-mode`—ATM virtual path (VP) cell relay promiscuous mode

`cesop`—CESOP-based Layer 2 VPN

`cisco-hdlc`—Cisco Systems–compatible HDLC

`ethernet`—Ethernet

`ethernet-vlan`—Ethernet VLAN

`frame-relay`—Frame Relay

`frame-relay-port-mode`—Frame Relay port mode

`interworking`—Layer 2.5 interworking VPN

`ppp`—PPP

`satsop-e1`—SATSOP-E1–based Layer 2 VPN

`satsop-e3`—SATSOP-E3–based Layer 2 VPN

`satsop-t1`—SATSOP-T1–based Layer 2 VPN

`satsop-t3`—SATSOP-T3–based Layer 2 VPN

- **Default:** For VPLS networks, the default encapsulation type is `ethernet`.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring the Encapsulation Type

Configuring VPLS Routing Instances

Configuring the Encapsulation Type for the Layer 2 Circuit Neighbor Interface

Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)

encapsulation-udp-fti

IN THIS SECTION

- Syntax | 6313
- Hierarchy Level | 6314
- Description | 6314
- Required Privilege Level | 6314
- Release Information | 6314

Syntax

```
encapsulation {
    udp
        destination {
            address address;
        }
        key key;
        source {
            address address;
        }
        tunnel-routing-instance {
            routing-instance routing-instance;
        }
        tunnel-termination;
port-profile profile-name;

    }

    }
}
```

Hierarchy Level

```
[edit interfaces fti number unit logical-unit-number tunnel]
```

Description

Configure UDP encapsulation over tunnel.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Evolved Release 21.1R1.

encoding (QFX Series)

IN THIS SECTION

- Syntax | 6314
- Hierarchy Level | 6315
- Description | 6315
- Options | 6315
- Required Privilege Level | 6315
- Release Information | 6315

Syntax

```
encoding (differential | non-differential);
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Description

Specify the encoding mode.

Options

`differential`—Differential phase line encoding.

`non-differential`—Non-differential phase line encoding.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Configuring Link Down Notification for Optics Options Alarm or Warning

100-Gigabit Ethernet OTN Options Configuration Overview

encoding (security)

IN THIS SECTION

- Syntax | 6316
- Hierarchy Level | 6316

- Description | 6316
- Options | 6316
- Required Privilege Level | 6316
- Release Information | 6317

Syntax

```
encoding (binary | pem);
```

Hierarchy Level

```
[edit security ike policy ike-peer-address],
[edit security certificates certification-authority ca-profile-name]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Specify the file format used for the `local-certificate` and `local-key-pair` statements.

Options

`binary`—Binary file format.

`pem`—Privacy-enhanced mail (PEM), an ASCII base 64 encoded format.

- **Default:** `binary`

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Type of Encoding Your CA Supports

Configuring an IKE Policy for Digital Certificates for an ES PIC

encoding

IN THIS SECTION

- Syntax | 6317
- Hierarchy Level | 6317
- Description | 6318
- Default | 6318
- Options | 6318
- Required Privilege Level | 6318
- Release Information | 6318

Syntax

```
encoding (nrz | nrzi);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

For serial interfaces, set the line encoding format.

Default

The default line encoding is non-return to zero (NRZ).

Options

`nrz`—Use NRZ line encoding.

`nrzi`—Use non-return to zero inverted (NRZI) line encoding.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Serial Line Encoding

encoding-type

IN THIS SECTION

- Syntax | 6319
- Hierarchy Level | 6319
- Description | 6319
- Default | 6319

Required Privilege Level | 6319
Release Information | 6320

Syntax

```
encoding-type (ethernet | packet | pdh | sonet-sdh);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name lsp-
attributes],
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
```

Description

Specify the encoding type of payload carried by the LSP. It can be any of the following:

- `ethernet`—Ethernet
- `packet`—Packet
- `pdh`—Plesiochronous digital hierarchy (PDH)
- `sonet-sdh`—SONET/SDH

Default

`packet`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Encoding Type

encrypt (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 6320
- Hierarchy Level | 6320
- Description | 6321
- Options | 6321
- Required Privilege Level | 6321
- Release Information | 6321

Syntax

```
encrypt {
    hash algorithm;
    prefix hash-prefix;
}
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
```

Description

Specify the transform to be applied to the header for the HTTP header enrichment so that you can add subscriber attributes in a way that is obscured from the user.

> **NOTE**: If you include this statement, then you also must configure `hash` and `prefix` statements.

Options

hash ***algorithm*** Use the specified hashing algorithm. Currently, only `md5` is supported.

prefix ***hash-prefix*** Use the specified prefix key (up to 63 alphanumeric characters). The prefix key is concatenated with the specified tag attribute and hashed. The resulting hash value is then inserted into the HTTP header.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

Configuring HCM Profiles and Assigning Tag Rules

Junos Web Aware HTTP Header Enrichment Overview

encryption

IN THIS SECTION

- Syntax | 6322
- Hierarchy Level | 6322
- Description | 6322
- Options | 6322
- Required Privilege Level | 6323
- Release Information | 6323

Syntax

```
encryption {
    algorithm algorithm;
    key (ascii-text key | hexadecimal key);
}
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]
```

Description

Configure an encryption algorithm and key for manual SA.

Options

`algorithm`—Type of encryption algorithm. The algorithm can be one of the following:

- `des-cbc`—Has a block size of 8 bytes (64 bits); the key size is 48 bits long.
- `3des-cbc`—Has a block size of 8 bytes (64 bits); the key size is 192 bits long.
- `aes-128-cbc`—Advanced Encryption Standard (AES) 128-bit encryption algorithm.

- `aes-192-cbc`—Advanced Encryption Standard (AES) 192-bit encryption algorithm.
- `aes-256-cbc`—Advanced Encryption Standard (AES) 256-bit encryption algorithm.

> **NOTE**: For `3des-cbc`, the first 8 bytes should differ from the second 8 bytes, and the second 8 bytes should be the same as the third 8 bytes.

`key`—Type of encryption key. The key can be one of the following:

- `ascii-text`—ASCII text key. Following are the key lengths, in ASCII characters, for the different encryption options:
  - `des-cbc` option, 8 ASCII characters
  - `3des-cbc` option, 24 ASCII characters
  - `aes-128-cbc` option, 16 ASCII characters
  - `aes-192-cbc` option, 24 ASCII characters
  - `aes-256-cbc` option, 32 ASCII characters
- `hexadecimal`—Hexadecimal key. Following are the key lengths, in hexadecimal characters, for the different encryption options:
  - `des-cbc` option, 16 hexadecimal characters
  - `3des-cbc` option, 48 hexadecimal characters
  - `aes-128-cbc` option, 32 hexadecimal characters
  - `aes-192-cbc` option, 48 hexadecimal characters
  - `aes-256-cbc` option, 64 hexadecimal characters

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`aes-128-cbc`, `aes-192-cbc`, and `aes-256-cbc` options added in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Security Associations

encryption-algorithm

IN THIS SECTION

- Syntax | 6324
- Hierarchy Level | 6324
- Description | 6324
- Options | 6324
- Required Privilege Level | 6325
- Release Information | 6325

Syntax

```
encryption-algorithm algorithm;
```

Hierarchy Level

```
[edit services ipsec-vpn ike proposal proposal-name],
[edit services ipsec-vpn ipsec proposal proposal-name]
```

Description

Configure an IKE or IPsec encryption algorithm.

Options

3des-cbc Has a block size of 24 bytes; the key size is 192 bits long.

aes-128-cbc Advanced Encryption Standard (AES) 128-bit encryption algorithm.

aes-192-cbc AES 192-bit encryption algorithm.

aes-256-cbc AES 256-bit encryption algorithm.

> **NOTE**: In Junos FIPS mode, AES-GCM is not supported in Junos OS Release 17.3R1. Starting in Junos OS Release 17.4R1, AES-GCM is supported in Junos FIPS mode.

aes-128-gcm (IPsec only) Advanced Encryption Standard in Galois/Counter Mode (AES-GCM) 128-bit encryption algorithm with a 16 octet integrity check value (ICV). This can only be used on an MS-MPC or MS-MIC.

aes-192-gcm (IPsec only) AES-GCM 192-bit encryption algorithm with a 16 octet ICV. This can only be used on an MS-MPC or MS-MIC.

aes-256-gcm (IPsec only) AES-GCM 256-bit encryption algorithm with a 16 octet ICV. This can only be used on an MS-MPC or MS-MIC.

des-cbc Has a block size of 8 bytes; the key size is 48 bits long.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`aes-128-cbc`, `aes-192-cbc`, and `aes-256-cbc` options added in Junos OS Release 7.6.

`aes-128-gcm`, `aes-192-gcm`, and `aes-256-gcm` options added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring IKE Proposals

Configuring IPsec Proposals

encryption-algorithm (Security)

IN THIS SECTION

- Syntax | 6326
- Hierarchy Level | 6326
- Description | 6326
- Options | 6326
- Required Privilege Level | 6327
- Release Information | 6327

Syntax

```
encryption-algorithm (3des-cbc | des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-256-cbc
| aes-256-gcm);
```

Hierarchy Level

```
[edit security ike proposal ike-proposal-name],
[edit security ipsec proposal ipsec-proposal-name]
```

Description

Configure an IKE or IPsec encryption algorithm. Two algorithms are supported: Cipher Block Chaining (CBC) and Galois/Counter Mode (GCM). CBC is older than GCM and is less secure. GCM is considered more secure than CBC since it provides built-in authentication and integrity checks. GCM is also more efficient for parallel processing, which can result in better performance for larger datasets.

Options

- `3des-cbc`—Encryption algorithm with key size of 24 bytes; its key size is 192 bits long.
- `des-cbc`—Encryption algorithm with key size of 8 bytes; its key size is 48 bits long.

- `aes-128-cbc`—Advanced encryption algorithm that has a key size of 16 bytes; its key size is 128 bits long.
- `aes-128-gcm`—AES GCM 128-bit encryption algorithm.
- `aes-192-cbc`—Advanced encryption algorithm that has a key size of 24 bytes; its key size is 192 bits long.
- `aes-256-cbc`—Advanced encryption algorithm that has a key size of 32 bytes; its key size is 256 bits long.
- `aes-256-gcm`—AES GCM 256-bit encryption algorithm. For an IKE proposal, AES 256-bit authenticated encryption algorithm is supported with IKEv2 only. When this option is used, aes-256-gcm should be configured at the [edit security ipsec proposal proposal-name] hierarchy level, and the authentication-algorithm option should not be configured at the [edit security ike proposal proposal-name] hierarchy level.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an IKE Proposal for Dynamic SAs

Configuring an IPsec Proposal for an ES PIC

encryption-algorithm (Security IKE)

IN THIS SECTION

Syntax | 6328

Hierarchy Level | 6328

- Description | 6328
- Options | 6328
- Required Privilege Level | 6329
- Release Information | 6329

Syntax

```
encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-256-cbc | aes-256-
gcm | chacha20-poly1305 | des-cbc);
```

Hierarchy Level

```
[edit security ike proposal proposal-name]
```

Description

Configure an encryption algorithm for an IKE proposal. The device does not delete existing IPsec SAs when you update the `encryption-algorithm` configuration in the IKE proposal.

Options

3des-cbc Has a block size of 24 bytes; the key size is 192 bits long.

aes-128-cbc Advanced Encryption Standard (AES) 128-bit encryption algorithm.

aes-128-gcm AES 128-bit authenticated encryption algorithm supported with IKEv2 only. When this option is used, `aes-128-gcm` should be configured at the [`edit security ipsec proposal` *`proposal-name`*] hierarchy level, and the `authentication-algorithm` option should not be configured at the [`edit security ike proposal` *`proposal-name`*] hierarchy level.

When `aes-128-gcm` or `aes-256-gcm` encryption algorithms are configured in the IPsec proposal, it is not mandatory to configure AES-GCM encryption algorithm in the corresponding IKE proposal.

aes-192-cbc AES 192-bit encryption algorithm.

aes-256-cbc AES 256-bit encryption algorithm.

aes-256-gcm AES 256-bit authenticated encryption algorithm supported with IKEv2 only. When this option is used, `aes-256-gcm` should be configured at the [`edit security ipsec proposal` *`proposal-name`*] hierarchy level, and the `authentication-algorithm` option should not be configured at the [`edit security ike proposal` *`proposal-name`*] hierarchy level.

> **NOTE**: Integrity cannot be set with AES-GCM encryption algorithm.

chacha20-poly1305 Chacha20-Poly1305 authenticated encryption algorithm.

On the supported devices, you can use the algorithm with IKEv2 only for control plane security.

Note that you cannot use the algorithm on SRX1500 Firewalls.

des-cbc Has a block size of 8 bytes; the key size is 48 bits long.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for `aes-128-gcm` and `aes-256-gcm` options added in Junos OS Release 15.1X49-D40.

Starting in Junos OS Release 20.2R1, we've changed the help text description as `NOT RECOMMENDED` for the CLI options `3des-cbc` and `des-cbc` for devices running IKED with `junos-ike` package installed.

Support for the `chacha20-poly1305` option added to SRX1600, SRX2300, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0 in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

IPsec Overview

proposal (Security IKE) | **13389**

Installing Junos IKE package

encryption-failures

IN THIS SECTION

- Syntax | 6330
- Hierarchy Level | 6330
- Description | 6330
- Default | 6330
- Options | 6331
- Required Privilege Level | 6331
- Release Information | 6331

Syntax

```
encryption-failures {
    threshold value;
}
```

Hierarchy Level

```
[edit security alarms potential-violation]
```

Description

Raise a security alarm after exceeding a specified number of encryption failures.

Default

Multiple encryption failures do not cause an alarm to be raised.

Options

failures—Number of encryption failures up to which an alarm is not raised. When the configured number is exceeded, an alarm is raised.

- **Range:** 1 through 1,000,000,000.
- **Default:** 1000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

IPsec Overview

Monitoring VPN Traffic

encryption (Junos-FIPS Software)

IN THIS SECTION

- Syntax | 6332
- Hierarchy Level | 6332
- Description | 6332
- Required Privilege Level | 6332
- Release Information | 6332

Syntax

```
encryption {
    algorithm 3des-cbc;
    key ascii-text ascii-text-string;
}
```

Hierarchy Level

```
[edit security ipsec internal security-association manual direction]
```

Description

Define the encryption parameters for internal Routing-Engine–to–Routing-Engine communication. The remaining statements are explained separately.

> **NOTE**: The hexadecimal format must be used for the encryption key to be FIPS compliant. The hexadecimal keys provide maximum key strength.

Required Privilege Level

Crypto Officer—To view and add this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

encryption (Junos OS)

IN THIS SECTION

- Syntax | 6333
- Hierarchy Level | 6333

- Description | 6333
- Options | 6333
- Required Privilege Level | 6334
- Release Information | 6334

Syntax

```
encryption {
    algorithm (des-cbc | 3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc);
    key (ascii-text key | hexadecimal key);
}
```

Hierarchy Level

```
[edit security ipsec security-association sa-name manual direction (inbound | outbound |
bidirectional)]
```

Description

Configure an encryption algorithm and key for a manual Security Association.

Options

`algorithm`—Type of encryption algorithm. It can be one of the following:

- `des-cbc`—Has a block size of 8 bytes (64 bits); its key size is 48 bits long.
- `3des-cbc`—Has block size of 8 bytes (64 bits); its key size is 192 bits long.

  NOTE: For `3des-cbc`, we recommend that the first 8 bytes be different from the second 8 bytes, and the second 8 bytes be the same as the third 8 bytes.

- `aes-128-cbc`—Has a block size of 128 bits; its key size is 128 bits long.
- `aes-192-cbc`—Has a block size of 128 bits; its key size is 192 bits long.

- aes-256-cbc—Has a block size of 128 bits; its key size is 256 bits long.

> **NOTE**: The aes-*-cbc algorithms support both IKE and IPsec configurations at the [security] hierarchy level.

key—Type of encryption key. It can be one of the following:

- ascii-text—ASCII text key. For the des-cbc option, the key contains 8 ASCII characters; for 3des-cbc, the key contains 24 ASCII characters.
- hexadecimal—Hexadecimal key. For the des-cbc option, the key contains 16 hexadecimal characters; for the 3des-cbc option, the key contains 48 hexadecimal characters.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

aes-128-cbc, aes-192-cbc, and aes-256-cbc algorithm options added in Junos OS Release 15.1.

RELATED DOCUMENTATION

Example: Using IPsec to Protect BGP Traffic

Configuring Manual IPsec Security Associations for an ES PIC

encryption (MACsec)

IN THIS SECTION

Syntax | 6335

Hierarchy Level | 6335

- Description | 6335
- Default | 6336
- Required Privilege Level | 6336
- Release Information | 6336

Syntax

```
encryption;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel secure-channel-name]
```

Description

Enable MACsec encryption within a secure channel.

You can enable MACsec without enabling encryption. If a connectivity association with a secure channel that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic has not been tampered with and does not represent a security threat.

Traffic traversing a MAC-enabled point-to-point Ethernet link traverses the link at the same speed regardless of whether encryption is enabled or disabled. You cannot increase the speed of traffic traversing a MACsec-enabled Ethernet link by disabling encryption.

This command is used to enable encryption when MACsec is configured using secure association key (SAK) security mode only. When MACsec is configuring using static connectivity association key (CAK) security mode, the encryption setting is configured outside of the secure channel using the `no-encryption` configuration statement.

Default

MACsec encryption is disabled when MACsec is configured using static SAK security mode, by default.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

encryption (MACsec for MX Series)

IN THIS SECTION

- Syntax | 6336
- Hierarchy Level | 6337
- Description | 6337
- Default | 6337
- Required Privilege Level | 6337
- Release Information | 6337

Syntax

```
encryption;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel
secure-channel-name]
```

Description

Enable MACsec encryption within a secure channel.

You can enable MACsec without enabling encryption. If a connectivity association with a secure channel that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic has not been tampered with and does not represent a security threat.

Traffic traversing a MAC-enabled point-to-point Ethernet link traverses the link at the same speed regardless of whether encryption is enabled or disabled. You cannot increase the speed of traffic traversing a MACsec-enabled Ethernet link by disabling encryption.

This command is used to enable encryption when MACsec is configured using secure association key (SAK) security mode only. When MACsec is configuring using static connectivity association key (CAK) security mode, the encryption setting is configured outside of the secure channel using the `no-encryption` configuration statement.

Default

MACsec encryption is disabled when MACsec is configured using static SAK security mode, by default.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

end-interface

IN THIS SECTION

- Syntax | 6338
- Hierarchy Level | 6338
- Description | 6338
- Required Privilege Level | 6339
- Release Information | 6339

Syntax

```
end-interface {
    interface interface-name;
    no-revert;
    protect-interface interface-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name],
[edit protocols l2circuit local-switching interface interface-name]
```

Description

Specify the end interface for a local interface switch.

> **NOTE**: The protect interface must be configured prior to configuring the `no-revert` statement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Local Interface Switching in Layer 2 Circuits

end-user-profile

IN THIS SECTION

- Syntax | 6340
- Hierarchy Level | 6340
- Description | 6340
- Options | 6341
- Required Privilege Level | 6341
- Release Information | 6341

Syntax

```
end-user-profile profile-name profile-name
domain-name domain-name;
{
attribute device-category {
    string string-value;
}
attribute device-identity {
    string string-value;
}
attribute device-vendor {
    string string-value;
}
attribute device-type {
    string string-value;
}
attribute device-os {
    string string-value;
}
attribute device-os-version {
    string string-value;
    }
}
```

Hierarchy Level

```
[edit services user-identification device-information]
```

Description

Specify the name of the device identity profile, also referred to as the `end-user-profile`, and either one or more of its attributes or the name of the Active Directory domain to which the device belongs.

The device identity profile is a key component of the SRX Series Firewall identity feature, which enables you to control access to network resources based on the identity of the user's device, not the identity of the user of the device. The device identity profile includes the domain name and a collection of attributes that characterize the device.

NOTE: You cannot configure the device identity profile without specifying either the domain that the device belongs to at least one of its attributes.

Options

- profile-name *profile-name*—Name of the device identity profile; for example, marketing-west-coast. The profile is specified in the `source-end-user-profile` field of a security policy.
- domain *domain-name*—Name of the domain to which the device belongs; for example, domain1.
- attribute device-identity *string*––Name given to the device, for example, my-device1.
- attribute device-category *string*––Category of the device, for example, laptop.
- attribute device-vendor *string*—Name of the manufacturer of the device, for example, Lenovo.
- attribute device-type *string*––Type of device; for example, ThinkPad.
- attribute device-os *string*-—Operating system running on the device; for example, Windows.
- attribute device-os-version *string*-—Version of the operating system that is running on the device; for example, 10.1.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 15.1X49-D70.

RELATED DOCUMENTATION

Understanding Access Control to Network Resources Based on Device Identity Information

Understanding the Device Identity Attributes and Profiles for the Integrated User Firewall Device Identity Authentication Feature

Understanding the Device Identity Authentication Table and Its Entries

endpoint-registration-timeout

IN THIS SECTION

- Syntax | 6342
- Hierarchy Level | 6342
- Description | 6342
- Options | 6342
- Required Privilege Level | 6343
- Release Information | 6343

Syntax

```
endpoint-registration-timeout value-in-seconds;
```

Hierarchy Level

```
[edit logical-systems name security alg h323],
[edit logical-systems name tenants name security alg h323],
[edit security alg h323],
[edit services alg h323],
[edit tenants name security alg h323]
```

Description

Configure the endpoint-registration-timeout to specify the timeout value in seconds for entries in the NAT table.

Options

value-in-seconds—Timeout value.

- **Range:** 10 through 50,000 seconds

- **Default:** 3600 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

h323 | 7961

enforce-first-as

IN THIS SECTION

- Syntax | 6343
- Hierarchy Level | 6344
- Description | 6344
- Required Privilege Level | 6344
- Release Information | 6344

Syntax

```
enforce-first-as;
```

Hierarchy Level

```
[edit fabric protocols bgp],
[edit fabric protocols bgp group group_name],
[edit fabric protocols bgp group group_name neighbor address],
[edit logical-systems name],
[edit logical-systems name protocols bgp],
[edit logical-systems name protocols bgp group group_name],
[edit logical-systems name protocols bgp group group_name neighbor address],
[edit logical-systems name routing-instances instance_name protocols bgp],
[edit logical-systems name routing-instances instance_name protocols bgp group group_name],
[edit logical-systems name routing-instances instance_name protocols bgp group group_name
neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor neighbor-address],
[edit routing-instances instance_name protocols bgp],
[edit routing-instances instance_name protocols bgp group group_name],
[edit routing-instances instance_name protocols bgp group group_name neighbor address]
```

Description

Enforce that the first (left-most) autonomous system number (ASN) in AS-path is the previous neighbor's ASN. When configured, this statement enforces that as the domain is transited, the routes received from an EBGP peer have the peer's ASN in the left-most position of the AS path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Prepending 4-Byte AS Numbers in an AS Path

Example: Enforcing Correct Autonomous System Number in AS-Path in BGP Network

enforce-strict-scale-limit-license (Subscriber Management)

IN THIS SECTION

- Syntax | 6345
- Hierarchy Level | 6345
- Description | 6345
- Required Privilege Level | 6345
- Release Information | 6345

Syntax

```
enforce-strict-scale-limit-license;
```

Hierarchy Level

```
[edit system services subscriber-management]
```

Description

Configure the router to strictly enforce the subscriber scaling license, and to not allow the normal grace period. No additional subscribers are allowed to log in after the number of subscribers reaches the maximum allowed for the license.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring User-Defined Dynamic Variables in Dynamic Profiles

User-Defined Variables

Using Variable Expressions in User-Defined Variables

engine-id

IN THIS SECTION

- Syntax | 6346
- Hierarchy Level | 6346
- Description | 6346
- Default | 6347
- Options | 6347
- Required Privilege Level | 6348
- Release Information | 6348

Syntax

```
engine-id {
    (local engine-id-suffix | use-default-ip-address | use-mac-address);
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Define a unique identifier for an SNMPv3 engine by configuring the suffix of the engine ID. The engine ID is used for identification only and not for addressing. There are two parts of an engine ID: the prefix

and the suffix. The prefix is formatted according to the specifications defined in RFC 3411, *An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks* and cannot be configured. The suffix is configured here.

> **NOTE**: SNMPv3 authentication and encryption keys are generated based on the associated user passwords and the engine ID. If you configure or change the engine ID, you must commit the user passwords and new engine ID before you configure SNMPv3 users, or the authentication will fail.
>
> For the engine ID, we recommend using the MAC address of the management port.
>
> By default the `use-mac-address` option is enabled. This causes the engine ID suffix to be based on a MAC address.
>
> Starting with Junos OS release 14.1X53-D43, QFX Series switches use the chassis MAC address for the engine ID suffix. On OCX Series, the MAC address of the management interface continues to be used.
>
> You can override the default behavior by using the `local` *`engine-id-suffix`* option, or by using the `use-default-ip-address` option.

Default

For QFX series, when no `snmp engine-id` is configured, the default value for the configuration is `use-mac-address`.

For ACX, EX, and MX series, when no `snmp engine-id` is configured, the default value for the configuration is `use-default-ip-address`.

> **NOTE**: For ZTP, if you need `snmp engine-id` to be constant, you must set the `engine-id` configuration to `set snmp engine-id local` *`engine-id-suffix`*.

Options

`local` ***`engine-id-suffix`***	The engine ID suffix is set based on the data entered.
`use-default-ip-address`	The engine ID suffix is generated from the default IP address.
`use-mac-address`	The engine ID suffix is generated from the MAC address of the management interface on all the junos platforms.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring the Local Engine ID

engine-id (Forwarding Options)

IN THIS SECTION

- Syntax | 6348
- Hierarchy Level | 6349
- Description | 6349
- Options | 6349
- Required Privilege Level | 6349
- Release Information | 6349

Syntax

```
engine-id number;
```

Hierarchy Level

```
[edit forwarding-options accounting name output interface interface-name],
[edit forwarding-options monitoring name output interface interface-name],
[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output
interface interface-name],
[edit forwarding-options sampling family (inet | inet6 | mpls) output interface interface-name]
```

Description

Specify the engine ID number for flow monitoring and accounting services.

Options

number—Identity of accounting interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

Configuring Flow Monitoring

Configuring Discard Accounting

engine-not-ready

IN THIS SECTION

- Syntax | 6350
- Hierarchy Level | 6350
- Description | 6350
- Options | 6350
- Required Privilege Level | 6351
- Release Information | 6351

Syntax

```
engine-not-ready (block | log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name fallback-
options]
```

Description

The scan engine is initializing itself, for example, loading the signature database. During this phase, it is not ready to scan a file. A file could either pass or be blocked according to this setting. The default action is block.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

engine-not-ready (Security Antivirus Sophos Engine)

IN THIS SECTION

- Syntax | 6351
- Hierarchy Level | 6351
- Description | 6352
- Options | 6352
- Required Privilege Level | 6352
- Release Information | 6352

Syntax

```
default (block | log-and-permit | permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name fallback-
options]
```

Description

The scan engine is initializing itself, for example, loading the signature database. During this phase, it is not ready to scan a file. A file could either pass or be blocked according to this setting.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic
- `permit`—Permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Sophos Antivirus Configuration Overview

engine-type

IN THIS SECTION

- Syntax | 6353
- Hierarchy Level | 6353
- Description | 6353
- Options | 6353

Required Privilege Level | 6354
Release Information | 6354

Syntax

```
engine-type number;
```

Hierarchy Level

```
[edit forwarding-options accounting name output interface interface-name],
[edit forwarding-options monitoring name output interface interface-name],
[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output
interface interface-name],
[edit forwarding-options sampling family (inet | inet6 | mpls) output interface interface-name]
```

Description

Specify the engine type number for flow monitoring and accounting services. The engine type attribute refers to the type of the flow switching engine, such as the route processor or a line module. The configured engine type is inserted in output `cflowd` packets. The `Source ID`, a 32-bit value to ensure uniqueness for all flows exported from a particular device, is the equivalent of the engine type and the engine ID fields.

NOTE: You must configure a source address in the output interface statements. The interface-level statement of engine-type is added automatically but you can override this value with manually configured statements to track different flows with a single `cflowd` collector.

Options

number—Platform-specific accounting interface type.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

Configuring Flow Monitoring

Configuring Discard Accounting

enhanced

IN THIS SECTION

- Syntax | 6354
- Hierarchy Level | 6355
- Release Information | 6355
- Description | 6355
- Required Privilege Level | 6355

Syntax

```
enhanced;
```

Hierarchy Level

```
[edit forwarding-options storm-control],
[edit logical-systems name routing-instances name forwarding-options storm-control],
[edit logical-systems name tenants name routing-instances name forwarding-options storm-control],
[edit routing-instances name forwarding-options storm-control],
[edit tenants name routing-instances name forwarding-options storm-control]
```

Release Information

Statement introduced in Junos OS Release 17.4R1.

Description

Storm control defines a threshold for broadcast, unknown unicast, or multicast traffic. If traffic exceeds this threshold, storm control drops the packets. When a filter is applied to the storm control interface, the dropped packet count increases even when storm control does not drop packets. You can configure enhanced storm control to collect statistics on packets dropped by storm control when a filter is configured on the same interface.

Required Privilege Level

interface

RELATED DOCUMENTATION

Enabling and Disabling Storm Control (ELS)

enhanced-cfm-mode

IN THIS SECTION

Syntax | 6356

- Hierarchy Level | 6356
- Description | 6356
- Required Privilege Level | 6356
- Release Information | 6356

Syntax

```
enhanced-cfm-mode enhanced-cfm-mode;
```

Hierarchy Level

```
[edit logical-systems name protocols oam ethernet connectivity-fault-management],
[edit protocols oam ethernet connectivity-fault-management]
```

Description

Enables enhanced CFM mode. When you enable enhanced CFM mode, Junos OS supports 32, 000 maintenance association end points (MEPs) and maintenance intermediate points (MIPs) each per chassis for bridge, VPLS, L2VPN, and CCC domains. To support enhanced CFM mode, configure the network services mode on the router as `enhanced-ip`.

> **NOTE**: After enabling CFM mode, restart CFM for the changes to take effect. If you do not restart CFM, CFM automatically restarts after 1 minute.

Required Privilege Level

Routing

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Enabling Enhanced Connectivity Fault Management Mode

enhanced-convergence

IN THIS SECTION

- Syntax | 6357
- Hierarchy Level | 6357
- Description | 6357
- Required Privilege Level | 6358
- Release Information | 6358

Syntax

```
enhanced-convergence;
```

Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options mc-ae]
[edit interfaces irb unit unit-number]
```

Description

NOTE: On EX9200 and QFX10000 switches, enhanced convergence is applicable for unicast traffic only—for example, when a MAC address is learned over an MC-AE interface, or an ARP entry is resolved over an MC-AE interface.

> **NOTE**: Enhanced convergence is not supported on QFX5100, QFX5110, QFX5120, QFX5200-48Y, QFX5200-32C, and QFX5210-64C switches.

Improves Layer 2 and Layer 3 convergence time when a multichassis aggregated Ethernet (MC-AE) link goes down or comes up in a bridge domain or VLAN. Convergence time is improved because the traffic on the MC-AE interface is switched to the interchassis link (ICL) without waiting for a MAC address update.

If you have configured an IRB interface over an MC-AE interface that has enhanced convergences enabled, then you must configure enhanced convergence on the IRB interface as well. Enhanced convergence must be enabled for both Layer 2 and Layer 3 interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R1.

enhanced (forwarding-class-accounting)

IN THIS SECTION

- Syntax | 6359
- Hierarchy Level | 6359
- Description | 6359
- Options | 6359
- Required Privilege Level | 6361
- Release Information | 6361

Syntax

```
enhanced {
    overhead-bytes overhead-value  (0-255 bytes);
    traffic-type (ucast | mcast);
    family (ipv4 | ipv6 | both );
    direction (ingress | egress | both);
}
```

Hierarchy Level

```
[edit interfaces interface-name forwarding-class-accounting],
[edit interfaces interface-name unit logical-unit-number forwarding-class-accounting]
```

Description

Specify the traffic and type for which you want to apply counters. A single aggregate counter per forwarding class can be used for inet and inet6 flows. Forwarding class accounting applies to transit traffic only, not host-generated or host-bound traffic. For ingress, only packets forwarded to the fabric are counted. For egress, only packets forwarded to the WAN are counted. Non-relevant network protocols such as ARP, BFD, and EOAM, as well as dropped packets, are not counted.

NOTE: Forwarding class accounting (enhanced mode) is not supported in "hyper-mode (forwarding-options)" on page 8221.

Options

overhead-bytes *overhead-value*

The number of overhead-bytes to be accounted.

- **Range:** 0 through 255 bytes
- **Default:** If zero, or no value is specified, the number of overhead-bytes will be determined by the encapsulation configured on the interface. For Ethernet, the default values are as follows:

Untagged DIX, includes CRC	18

Single-tagged DIX, includes CRC	22
Double-tagged DIX, includes CRC	26
VLAN-bridge (CRC)	4
VLAN-CCC (CRC)	4
Untagged VLAN-TCC, includes CRC	18
Single-tagged VLAN-TCC, includes CRC	22
Double-tagged VLAN-TCC, includes CRC	26
VLAN-VPLS (CRC)	4

Packet over SONET (POS) interfaces on an MPC are dependent on the frame check sequence (FCS) settings on the interface, which may be two or four bytes. The default values for POS interfaces are as follows:

PPP (Layer 2 overhead plus FCS)	4+ (2 or 4)
Cisco-HDLC (Layer 2 overhead plus FCS)	4+ (2 or 4)
Frame-relay (Layer 2 overhead plus FCS)	4+ (2 or 4)
Cisco-HDLC-CCC (FCS)	2 or 4
Cisco-HDLC-TCC (FCS)	2 or 4
PPP-CCC (FCS)	2 or 4
PPP-TCC (FCS)	2 or 4

Frame-relay-CCC (FCS)	2 or 4
Frame-relay-TCC (FCS)	2 or 4

traffic-type *(ucast | mcast)*

Traffic to be counted can be unicast only or multicast only.

- **Default:** not enabled, i.e., both unicast and multicast are counted

family *(ipv4 | ipv6 | both)*

Traffic to be counted can be IPv4, IPv6, or both.

- **Default:** both

direction *(ingress | egress | both)*

Traffic can be inbound to the interface, outbound, or both.

- **Default:** both

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1 in MX Series.

RELATED DOCUMENTATION

CoS-Based Interface Counters for IPv4 or IPv6 Aggregate on Layer 2

show class-of-service interface | **19056**

clear interfaces statistics | **3651**

enhanced-hash-key

IN THIS SECTION

- Syntax | 6362
- Syntax (EX Series) | 6364
- Syntax (QFX5000 Series Switches) | 6365
- Syntax (QFX10000 Series Switches) | 6368
- Syntax (SRX5000 Series Routers)) | 6369
- Syntax (Junos OS Evolved) | 6370
- Syntax Junos OS Evolved (QFX5220 and QFX5130) | 6371
- Hierarchy Level | 6372
- Description | 6373
- Default | 6375
- Options | 6375
- Required Privilege Level | 6377
- Release Information | 6377

Syntax

```
enhanced-hash-key {
    family any {
        incoming-interface-index;
        no-tunnel-payload;
    }
    family inet {
        gtp-tunnel-endpoint-identifier;
        l2tp-tunnel-session-identifier;
        incoming-interface-index;
        no-destination-port;
        no-source-port;
        type-of-service;
    }
    family inet6 {
        gtp-tunnel-endpoint-identifier;
```

```
        incoming-interface-index;
        no-destination-port;
        no-source-port;
        traffic-class;
    }
    family mpls {
        ether-pseudowire {
            zero-control-word;
        }
        incoming-interface-index;
        label-1-exp;
        no-ether-pseudowire;
        no-labels;
        no-payload;
    }
    family multiservice {
        incoming-interface-index;
        no-mac-addresses;
        no-payload;
        outer-priority;
        source-mac;
        no-destination-mac;
    }
    services-loadbalancing {
        family inet {
            layer-3-services {
                destination-address;
                incoming-interface-index;
                source-address;
            }
        }
        family inet6 {
            layer-3-services {
                destination-address;
                incoming-interface-index;
                source-address;
                src-prefix-len;
            }
        }
    }
    symmetric;
}
```

Syntax (EX Series)

```
enhanced-hash-key {
    ecmp-dlb {
        assigned-flow;
        per-packet;
        flowlet inactivity-interval;
        ether-type (ipv4|ipv6|mpls);
    }
    ecmp-resilient-hash;
    fabric-load-balance {
        flowlet {
            inactivity-interval interval;
        }
        per-packet;
    }
    hash-mode {
        layer2-header;
        layer2-payload;
    }
    family inet {
        no-ipv4-destination-address;
        no-ipv4-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-protocol;
        vlan-id;
    }
    family inet6 {
        no-ipv6-destination-address;
        no-ipv6-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-next-header;
        vlan-id;
    }
    layer2 {
        no-destination-mac-address;
        no-ether-type;
        no-source-mac-address;
        vlan-id;
```

```
    }
}
```

Syntax (QFX5000 Series Switches)

```
enhanced-hash-key {
    conditional-match name {
        offset1 {
            base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader
| start-of-L4-Header);
            matchdata1 matchdata1;
            matchdata1-mask matchdata1-mask;
            offset1-value offset1-value;
        }
        offset2 {
            base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader
| start-of-L4-Header);
            matchdata2 matchdata2;
            matchdata2-mask matchdata2-mask;
            offset2-value offset2-value;
        }
        offset3 {
            base-offset3 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader
| start-of-L4-Header);
            matchdata3 matchdata3;
            matchdata3-mask matchdata3-mask;
            offset3-value offset3-value;
        }
        offset4 {
            base-offset4 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader
| start-of-L4-Header);
            matchdata4 matchdata4;
            matchdata4-mask matchdata4-mask;
            offset4-value offset4-value;
        }
    }
    ecmp-dlb {
        assigned-flow;
        per-packet;
        flowlet inactivity-interval;
        ether-type (ipv4|ipv6|mpls);
```

```
    }

    ecmp-resilient-hash;
    fabric-load-balance {
        flowlet {
            inactivity-interval interval;
        }
        per-packet;
    }
    flex-hashing name {
        ethtype {
            inet {
                conditional-match conditional-match;
                hash-offset {
                    offset1 {
                        base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
OuterHeader | start-of-L4-Header);
                        offset1-mask offset1-mask;
                        offset1-value offset1-value;
                        offset2 {
                            base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-
L3-OuterHeader | start-of-L4-Header);
                            offset2-mask offset2-mask;
                            offset2-value offset2-value;
                        }
                    }
                }
                interface interface;
            }
            inet6 {
                conditional-match conditional-match;
                hash-offset {
                    offset1 {
                        base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
OuterHeader | start-of-L4-Header);
                        offset1-mask offset1-mask;
                        offset1-value offset1-value;
                        offset2 {
                            base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-
L3-OuterHeader | start-of-L4-Header);
                            offset2-mask offset2-mask;
                            offset2-value offset2-value;
                        }
```

```
                    }
                }
                interface interface;
            }
            mpls {
                conditional-match conditional-match;
                hash-offset {
                    offset1 {
                        base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
OuterHeader | start-of-L4-Header);
                        offset1-mask offset1-mask;
                        offset1-value offset1-value;
                        offset2 {
                            base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-
L3-OuterHeader | start-of-L4-Header);
                            offset2-mask offset2-mask;
                            offset2-value offset2-value;
                        }
                    }
                }
                interface interface;
                num-labels num-labels;
            }
        }
    }
    hash-mode {
        layer2-header;
        layer2-payload;
        gtp-header-offset offset-value;
    }
    hash-parameters {
        ecmp {
            function {
                (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
            }
            offset offset;
            preprocess;
        }
        lag {
            function {
                (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
            }
            offset offset;
```

```
            preprocess;
        }
    }
    family inet {
        gtp-tunnel-endpoint-identifier;
        no-incoming-port;
        no-ipv4-destination-address;
        no-ipv4-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-protocol;
        vlan-id;
    }
    family inet6 {
        no-incoming-port;
        no-ipv6-destination-address;
        no-ipv6-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-next-header;
        vlan-id;
    }
    layer2 {
        no-destination-mac-address;
        no-ether-type;
        no-source-mac-address;
        vlan-id;
    }
    symmetric-hash {
        inet;
        inet6;
    }
}
vxlan {
    no-inner-payload;
}
```

Syntax (QFX10000 Series Switches)

```
enhanced-hash-key {
    hash-seed seed-value;
```

```
    family inet {
        gtp-tunnel-endpoint-identifier;
        no-ipv4-destination-address;
        no-ipv4-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-incoming-port;
    }
    family inet6 {
        gtp-tunnel-endpoint-identifier;
        ipv6-flow-label;
        no-ipv6-destination-address;
        no-ipv6-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-incoming-port;
    }
    layer2 {
        destination-mac-address
        inner-vlan-id;
        no-ether-type;
        no-vlan-id;
        source-mac-address;
    }
    no-mpls;
    gre {
        key;
        protocol;
    }
    vxlan-vnid
    }
}
```

Syntax (SRX5000 Series Routers))

```
enhanced-hash-key {
         family inet {
        gtp-tunnel-endpoint-identifier;
        incoming-interface-index;
        l2tp-tunnel-session-identifier;
        session-id;
```

```
        type-of-service;
    }
    family inet6 {
        gtp-tunnel-endpoint-identifier;
        incoming-interface-index;
        no-flow-label;
        session-id;
        traffic-class;
    }
    layer2 {
        destination-mac-address
        inner-vlan-id;
        no-ether-type;
        no-vlan-id;
        source-mac-address;
    }
    no-mpls;
    gre {
        key;
        protocol;
    }
    vxlan-vnid
    }
}
```

Syntax (Junos OS Evolved)

```
enhanced-hash-key {
    family any {
        incoming-interface-index;
        no-tunnel-payload;
    }
    family inet {
        gtp-tunnel-endpoint-identifier;
        no-destination-port;
        no-source-port;
        type-of-service;
    }
    family inet6 {
        gtp-tunnel-endpoint-identifier;
        no-destination-port;
```

```
        no-flow-label;
        no-source-port;
        traffic-class;
    }
    family mpls {
        label-1-exp;
        no-labels;
        no-payload;
    }
    family multiservice {
        no-payload;
    }
    hash-seed;
    }
    resilient-hash-seed;
    }
}
```

Syntax Junos OS Evolved (QFX5220 and QFX5130)

```
enhanced-hash-key {
    ecmp-resilient-hash;
    family inet {
        no-incoming-device;
        no-incoming-port;
        no-ipv6-destination-address;
        no-ipv6-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-protocol;
        vlan-id;
    }
    family inet6 {
        no-incoming-device;
        no-incoming-port;
        no-ipv6-destination-address;
        no-ipv6-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-nxt-hdr;
```

```
        vlan-id;
    }
     hash-mode {
           layer2-header;
           layer2-payload;
    }
     layer2 {
           no-destination-mac-address;
           no-ether-type;
        no-incoming-device;
        no-incoming-device;
           no-smac-address;
           vlan-id;
    }
     hash-parameters {
        ecmp {
            function {
                (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
            }
            offset offset;
            preprocess;
        }
        lag {
            function {
                (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
            }
            offset offset;
            preprocess;
        }
   }
}
```

Hierarchy Level

```
[edit forwarding-options],
[edit logical-systems logical-system-name routing-instances instance-name forwarding-options],
[edit routing-instances instance-name forwarding-options]
```

Junos OS Evolved uses the `set` command instead of `edit`.

Description

Configure the hashing key used to hash link aggregation group (LAG) and equal-cost multipath (ECMP) traffic, or enable adaptive load balancing (ALB) in a Virtual Chassis Fabric (VCF).

> **NOTE**: Starting in Junos OS Release 14.1X53-D46, 15.1R7, 16.1R6, 17.1R3, 17.2R2, 17.3R2, and 17.4R1, the ALB feature is deprecated. If `fabric-load-balance` is enabled in the configuration for a VCF, delete the configuration item upon upgrading Junos OS.

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The computed hash is not only used in selecting an ECMP path, but also is used for load balancing. Starting in Junos OS Release 18.3R1, the `flow-label field` is included by default in the hash computation for IPv6, GRE, and PPPoE packets. This can be beneficial, for example, when you have MX routers operating as designated router (DR) or rendezvous point (RP) and want to load balance traffic on the basis of a single Layer 3 or Layer 4 flow. You can revert to the previous method of hash computation by setting the `no-flow-label` option.

- For GRE packets, if the outer IP packet is non-option packet and the inner packet is IPv4 or IPv6, then the source and destination IP addresses from inner packet will be included in hash computation.

  The Layer 4 ports will also be included in hash computation if the protocol of the inner IP packet is TCP or UDP, and if the inner IP packet is not an options packet.

  If outer IP packet is a non-options packet, and the inner packet is MPLS, then the top inner label is included in hash computation.

- For PPPoE packet, if the inner packet is IPv4 or IPv6, then source and destination IP addresses from inner packet will be included in hash computation.

  The Layer 4 ports are included in the hash computation if the protocol of inner IP packet is TCP or UDP, and the inner IP packet is a non-options packet.

For MX Series routers with MPCs, T4000 routers with Type 5 FPCs, EX9200 switches, and PTX10008 routers, select data used in the hash key for enhanced IP forwarding engines.

By default, MPCs use the following parameters for hashing:

- Source IP address
- Destination IP address
- Layer 3 protocol
- Source port
- Destination port
- Generic routing encapsulation (GRE) for GRE packets only.

You can modify the default hashing mechanism on MPCs and Type 5 FPCs by configuring statements at the `[edit forwarding-options enhanced-hash-key]` hierarchy level.

QFX10000 Series switches

- On QFX10000 Series switches, you can configure the hash seed for load balancing. By default, the QFX10000 Series switches use the system MAC address to generate a hash seed value. You can configure the hash seed value using the **hash-seed** statement at the `[edit forwarding-options enhanced-hash-key]` hierarchy level. Set a value between 0 and 4294967295. If you do not configure a hash seed value, the system generates a hash seed value based on the system MAC address.
- Starting in Junos OS Release 18.4R1, symmetric hashing is supported on the QFX10000 Series switches. You configure the `no-incoming-port` option under the [edit forwarding-options enhanced-hash-key] hierarchy. By default, Destination IP (DIP), SIP, Layer 4 source and destination ports, and the incoming port are used for hashing. You can only configure symmetric hashing at the global level.

QFX5000 Series switches

> **NOTE**: On QFX5000 Series switches, multicast packets gets duplicated or dropped when `vlan-id` is included in the `enhanced-hash-key` configuration. Hence, while configuring `enhanced-hash-key`, you must configure `layer2-payload` and `layer2-header` options without configuring `vlan-id`. This behavior is only seen in multicast traffic.

- Starting in Junos OS Release 19.4R1, the dynamic load balancing on ECMP is supported on QFX5120-32C and QFX5120-48Y switches. You can configure the `ecmp-dlb` option under the [edit forwarding-options enhanced-hash-key] hierarchy. Refer *Dynamic Load Balancing* for more details.
- To enable symmetric hashing on the QFX5000 line of switches, configure the `symmetric-hash` option.

Enhanced hash key is supported on SRX5000 routers by implementing control path for for the configured setting to reach the services processing card (SPC2) and SPC3. You configure the `session-id` option to enable the control path for the configured setting to reach the service processing unit (SPU).

For more details, see CLI Explorer.

Default

In PTX, compared to MX (which is similar to QFX), the source and destination mac address options for hash computation are different. While QFX excludes the default MAC address fields for hash computations, PTX includes the destination MAC while excluding the source-mac.

Options

`services-loadbalancing`—Distributes traffic across PICs based on source IP address when a route pointing to more than one services PICs is installed.

`symmetric`—Enable symmetric load balancing across aggregated Ethernet interfaces. This option is needed on Trio-based MPCs only.

Data selections for `services-loadbalancing`:

- `inet`—IPv4 addressing protocol.
- `inet6`—IPv6 addressing protocol.
- `layer-3-services`—Include layer 3 IP data in the hash key.
- `incoming-interface-index`—Include incoming interface index in the hash key.
- `source-address`—Include source-address in the hash key.
- `destination-address`—Include destination-address in the hash key.
- `src-prefix-len`—Include the source prefix length in the hash key.

Data selections for family `any`:

- `incoming-interface-index`—(PTX10008 only) Include incoming interface index in the hash key.
- `no-tunnel-payload`—(PTX10001-36MR, PTX10004, PTX10008, and PTX10016 only) Omit the tunnel payload data from the hash key.

Data selections for family `inet`:

- **gtp-tunnel-endpoint-identifer**—Include the tunnel endpoint identifier (TEID) field in the hash key for GPRS tunneling protocol (GTP) traffic.

  > **NOTE**: This option is supported only on MX Series routers with MPCs and on the MX80 router.

- `incoming-interface-index`—Include incoming interface index in the hash key.

- `no-destination-port`—Omit IP destination port in the hash key.
- `no-source-port`—Omit IP source port in the hash key.
- `type-of-service`—Include type-of-service (TOS) byte in the hash key.

Data selections for family `inet6`:

- **gtp-tunnel-endpoint-identifer**—Include the tunnel endpoint identifier (TEID) field in the hash key for GPRS tunneling protocol (GTP) traffic.

  > **NOTE**: This option is supported only on MX Series routers with MPCs and on the MX80 router.

- `incoming-interface-index`—Include the incoming interface index in the hash key.
- `no-destination-port`—Omit the IP destination port in the hash key.
- `no-source-port`—Omit the IP source port in the hash key.
- `traffic-class`—Include the traffic class byte in the hash key.

Data selections for family `mpls`:

- `ether-pseudowire`—Load-balance IP over Ethernet pseudowire. Presence of zero-control-word in the payload indicates an Ethernet frame.
- `incoming-interface-index`—Include the incoming interface index in the hash key.
- `label-1-exp`—The EXP bit of the first label is used in the hash calculation.
- `no-ether-pseudowire`—Omit the Ethernet pseudowire payload data from the hash key (MX Series routers with MPCs only).
- `no-labels`—Omit MPLS labels from the hash key (PTX10008 only).
- `no-payload`—Omit the MPLS payload data from the hash key.

Data selections for family `multiservice`:

- `incoming-interface-index`—Include the incoming interface index in the hash key.
- `no-mac-addresses`—Omit source and destination MAC addresses from the hash key.
- `no-payload`—Omit the payload data from the hash key.
- `outer-priority`—Include the outer 802.1 priority bits in the hash key.

- `source-mac`—Includes source MAC address in hash key
- `no-destination-mac`—Excludes destination MAC address in hash key.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`services-loadbalancing` statement introduced in Junos OS Release 11.2.

`gtp-tunnel-endpoint-identifier` statement introduced in Junos OS Release 13.2

The `fabric-load-balance` statement introduced in Junos OS Release 14.1X53-D10.

The `fabric-load-balance` statement deprecated starting in Junos OS Releases 14.1X53-D46, 15.1R7, 16.1R6, 17.1R3, 17.2R2, 17.3R2, and 17.4R1.

The `hash-seed` statement introduced in Junos OS Release 15.1X53-D30.

`ether-pseudowire` statement introduced in Junos OS Release 16.1 for the M Series, MX Series, and PTX Series.

`l2tp-tunnel-session-identifier` statement introduced in Junos OS Release 17.2

Starting in Junos OS Release 18.3R1, the default behavior for IPv6, GRE, and PPPoE packet hash computation is to include the flow-label field for improved load-balancing in certain cases. Use the `no-payload` option to revert to the previous method for hash computation.

The `ecmp-dlb` statement introduced in Junos OS Release 19.4R1 for QFX5120-32C and QFX5120-48Y switches.

Option `symmetric-hash` introduced in Junos OS Release 20.4R1.

Statement introduced in Junos OS Evolved Release 21.3R1.

The `ecmp-dlb` statement is introduced in Junos OS Release 23.2R1 for EX4400-24T, EX4400-24X, and EX-4400-48F switches.

RELATED DOCUMENTATION

Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers

Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic

Understanding Passive Monitoring

Understanding Per-Packet Load Balancing

show forwarding-options enhanced-hash-key

enhanced-hash-key (EX)

IN THIS SECTION

- Syntax (EX Series) | **6378**
- Syntax (QFX5000 Series Switches) | **6379**
- Syntax (QFX10000 Series Switches) | **6383**
- Syntax (SRX5000 Series Routers)) | **6384**
- Syntax (Junos OS Evolved) | **6385**
- Syntax Junos OS Evolved (QFX5220 and QFX5130) | **6386**
- Hierarchy Level | **6387**
- Description | **6387**
- Required Privilege Level | **6388**
- Release Information | **6389**

Syntax (EX Series)

```
enhanced-hash-key {
    ecmp-dlb {
        assigned-flow;
        per-packet;
        flowlet inactivity-interval;
```

```
        ether-type (ipv4|ipv6|mpls);
    }
    ecmp-resilient-hash;
    fabric-load-balance {
        flowlet {
            inactivity-interval interval;
        }
        per-packet;
    }
    hash-mode {
        layer2-header;
        layer2-payload;
    }
    family inet {
        no-ipv4-destination-address;
        no-ipv4-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-protocol;
        vlan-id;
    }
    family inet6 {
        no-ipv6-destination-address;
        no-ipv6-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-next-header;
        vlan-id;
    }
    layer2 {
        no-destination-mac-address;
        no-ether-type;
        no-source-mac-address;
        vlan-id;
    }
}
```

Syntax (QFX5000 Series Switches)

```
enhanced-hash-key {
    conditional-match name {
```

```
        offset1 {
            base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader
| start-of-L4-Header);
            matchdata1 matchdata1;
            matchdata1-mask matchdata1-mask;
            offset1-value offset1-value;
        }
        offset2 {
            base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader
| start-of-L4-Header);
            matchdata2 matchdata2;
            matchdata2-mask matchdata2-mask;
            offset2-value offset2-value;
        }
        offset3 {
            base-offset3 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader
| start-of-L4-Header);
            matchdata3 matchdata3;
            matchdata3-mask matchdata3-mask;
            offset3-value offset3-value;
        }
        offset4 {
            base-offset4 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-OuterHeader
| start-of-L4-Header);
            matchdata4 matchdata4;
            matchdata4-mask matchdata4-mask;
            offset4-value offset4-value;
        }
    }
    ecmp-dlb {
        assigned-flow;
        per-packet;
        flowlet inactivity-interval;
        ether-type (ipv4|ipv6|mpls);
    }

    ecmp-resilient-hash;
    fabric-load-balance {
        flowlet {
            inactivity-interval interval;
        }
        per-packet;
    }
```

```
    flex-hashing name {
        ethtype {
            inet {
                conditional-match conditional-match;
                hash-offset {
                    offset1 {
                        base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
OuterHeader | start-of-L4-Header);
                        offset1-mask offset1-mask;
                        offset1-value offset1-value;
                        offset2 {
                            base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-
L3-OuterHeader | start-of-L4-Header);
                            offset2-mask offset2-mask;
                            offset2-value offset2-value;
                        }
                    }
                }
                interface interface;
            }
            inet6 {
                conditional-match conditional-match;
                hash-offset {
                    offset1 {
                        base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
OuterHeader | start-of-L4-Header);
                        offset1-mask offset1-mask;
                        offset1-value offset1-value;
                        offset2 {
                            base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-
L3-OuterHeader | start-of-L4-Header);
                            offset2-mask offset2-mask;
                            offset2-value offset2-value;
                        }
                    }
                }
                interface interface;
            }
            mpls {
                conditional-match conditional-match;
                hash-offset {
                    offset1 {
                        base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
```

```
OuterHeader | start-of-L4-Header);
                        offset1-mask offset1-mask;
                        offset1-value offset1-value;
                        offset2 {
                            base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-
L3-OuterHeader | start-of-L4-Header);
                            offset2-mask offset2-mask;
                            offset2-value offset2-value;
                        }
                    }
                }
                interface interface;
                num-labels num-labels;
            }
        }
    }
    hash-mode {
        layer2-header;
        layer2-payload;
        gtp-header-offset offset-value;
    }
    hash-parameters {
        ecmp {
            function {
                (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
            }
            offset offset;
            preprocess;
        }
        lag {
            function {
                (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
            }
            offset offset;
            preprocess;
        }
    }
    family inet {
        gtp-tunnel-endpoint-identifier;
        no-incoming-port;
        no-ipv4-destination-address;
        no-ipv4-source-address;
        no-l4-destination-port;
```

```
        no-l4-source-port;
        no-protocol;
        vlan-id;
    }
    family inet6 {
        no-incoming-port;
        no-ipv6-destination-address;
        no-ipv6-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-next-header;
        vlan-id;
    }
    layer2 {
        no-destination-mac-address;
        no-ether-type;
        no-source-mac-address;
        vlan-id;
    }
    symmetric-hash {
        inet;
        inet6;
    }
}
vxlan {
    no-inner-payload;
}
```

Syntax (QFX10000 Series Switches)

```
enhanced-hash-key {
    hash-seed seed-value;
    family inet {
        gtp-tunnel-endpoint-identifier;
        no-ipv4-destination-address;
        no-ipv4-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-incoming-port;
    }
    family inet6 {
```

```
        gtp-tunnel-endpoint-identifier;
        ipv6-flow-label;
        no-ipv6-destination-address;
        no-ipv6-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-incoming-port;
    }
    layer2 {
        destination-mac-address
        inner-vlan-id;
        no-ether-type;
        no-vlan-id;
        source-mac-address;
    }
    no-mpls;
    gre {
        key;
        protocol;
    }
    vxlan-vnid
    }
}
```

Syntax (SRX5000 Series Routers))

```
enhanced-hash-key {
         family inet {
        gtp-tunnel-endpoint-identifier;
        incoming-interface-index;
        l2tp-tunnel-session-identifier;
        session-id;
        type-of-service;
    }
    family inet6 {
        gtp-tunnel-endpoint-identifier;
        incoming-interface-index;
        no-flow-label;
        session-id;
        traffic-class;
    }
```

```
    layer2 {
        destination-mac-address
        inner-vlan-id;
        no-ether-type;
        no-vlan-id;
        source-mac-address;
    }
    no-mpls;
    gre {
        key;
        protocol;
    }
    vxlan-vnid
    }
}
```

Syntax (Junos OS Evolved)

```
enhanced-hash-key {
    family any {
        incoming-interface-index;
        no-tunnel-payload;
    }
    family inet {
        gtp-tunnel-endpoint-identifier;
        no-destination-port;
        no-source-port;
        type-of-service;
    }
    family inet6 {
        gtp-tunnel-endpoint-identifier;
        no-destination-port;
        no-flow-label;
        no-source-port;
        traffic-class;
    }
    family mpls {
        label-1-exp;
        no-labels;
        no-payload;
    }
```

```
    family multiservice {
        no-payload;
    }
    hash-seed;
    }
    resilient-hash-seed;
    }
}
```

Syntax Junos OS Evolved (QFX5220 and QFX5130)

```
enhanced-hash-key {
    ecmp-resilient-hash;
    family inet {
        no-incoming-device;
        no-incoming-port;
        no-ipv6-destination-address;
        no-ipv6-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-protocol;
        vlan-id;
    }
    family inet6 {
        no-incoming-device;
        no-incoming-port;
        no-ipv6-destination-address;
        no-ipv6-source-address;
        no-l4-destination-port;
        no-l4-source-port;
        no-nxt-hdr;
        vlan-id;
    }
     hash-mode {
            layer2-header;
            layer2-payload;
    }
     layer2 {
            no-destination-mac-address;
            no-ether-type;
```

```
        no-incoming-device;
        no-incoming-device;
           no-smac-address;
           vlan-id;
    }
     hash-parameters {
        ecmp {
            function {
                (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
            }
            offset offset;
            preprocess;
        }
        lag {
            function {
                (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
            }
            offset offset;
            preprocess;
        }
   }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Junos OS Evolved uses the `set` command instead of `edit`.

Description

Configure the hashing key used to hash link aggregation group (LAG) and equal-cost multipath (ECMP) traffic, or enable adaptive load balancing (ALB) in a Virtual Chassis Fabric (VCF).

NOTE: Starting in Junos OS Release 14.1X53-D46, 15.1R7, 16.1R6, 17.1R3, 17.2R2, 17.3R2, and 17.4R1, the ALB feature is deprecated. If `fabric-load-balance` is enabled in the configuration for a VCF, delete the configuration item upon upgrading Junos OS.

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

QFX10000 Series switches

- On QFX10000 Series switches, you can configure the hash seed for load balancing. By default, the QFX10000 Series switches use the system MAC address to generate a hash seed value. You can configure the hash seed value using the **hash-seed** statement at the [`edit forwarding-options enhanced-hash-key`] hierarchy level. Set a value between 0 and 4294967295. If you do not configure a hash seed value, the system generates a hash seed value based on the system MAC address.

- Starting in Junos OS Release 18.4R1, symmetric hashing is supported on the QFX10000 Series switches. You configure the `no-incoming-port` option under the [edit forwarding-options enhanced-hash-key] hierarchy. By default, Destination IP (DIP), SIP, Layer 4 source and destination ports, and the incoming port are used for hashing. You can only configure symmetric hashing at the global level.

QFX5000 Series switches

> **NOTE**: On QFX5000 Series switches, multicast packets gets duplicated or dropped when `vlan-id` is included in the `enhanced-hash-key` configuration. Hence, while configuring `enhanced-hash-key`, you must configure `layer2-payload` and `layer2-header` options without configuring `vlan-id`. This behavior is only seen in multicast traffic.

- Starting in Junos OS Release 19.4R1, the dynamic load balancing on ECMP is supported on QFX5120-32C and QFX5120-48Y switches. You can configure the `ecmp-dlb` option under the [edit forwarding-options enhanced-hash-key] hierarchy. Refer *Dynamic Load Balancing* for more details.

- To enable symmetric hashing on the QFX5000 line of switches, configure the `symmetric-hash` option.

Enhanced hash key is supported on SRX5000 routers by implementing control path for for the configured setting to reach the services processing card (SPC2) and SPC3. You configure the `session-id` option to enable the control path for the configured setting to reach the service processing unit (SPU).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D15.

The `fabric-load-balance` statement introduced in Junos OS Release 14.1X53-D10.

The `fabric-load-balance` statement deprecated starting in Junos OS Releases 14.1X53-D46, 15.1R7, 16.1R6, 17.1R3, 17.2R2, 17.3R2, and 17.4R1.

The `hash-seed` statement introduced in Junos OS Release 15.1X53-D30.

The `ecmp-dlb` statement introduced in Junos OS Release 19.4R1 for QFX5120-32C and QFX5120-48Y switches.

Option `symmetric-hash` introduced in Junos OS Release 20.4R1.

Statement introduced in Junos OS Evolved Release 21.3R1.

The `ecmp-dlb` statement is introduced in Junos OS Release 23.2R1 for EX4400-24T, EX4400-24X, and EX-4400-48F switches.

RELATED DOCUMENTATION

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic

Understanding Passive Monitoring

Understanding Per-Packet Load Balancing

show forwarding-options enhanced-hash-key | **20086**

enhanced-hash-key profile1/profile2

IN THIS SECTION

- Syntax | **6390**
- Hierarchy Level | **6390**
- Description | **6390**

- Default | 6392
- Options | 6393
- Required Privilege Level | 6394
- Release Information | 6394

Syntax

```
enhanced-hash-key profile1/profile2 profile-fields
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Enable or disable profile fields for an enhanced-hash-key load balancing profile (profile1 or profile2).

On EVO platforms, users can disable or enable enhanced-hash-key fields for load balancing by configuring "set forwarding-options enhanced-hash-key". However, this enhanced-hash-key config change is chassis-wide, and will be applied to any route. Committing a change on any particular enhanced-hash-key field will affect load balance of all applicable traffic.

Therefore, users can configure two load balance profiles (profile1 and profile2) with different sets of enhanced-hash-key configurations. Users will be able to apply a load balance profile to a specific route, not affecting other routes on the chassis, adding flexibility for load balance management.

The following are the profile fields that can be used for enhanced-hash-key profile1 and profile2:

Table 34: profile fields

Settings	enhanced-hash-key profile field	Description
Interface	no-in-interface	Inner inteface fields
	no-interface	Interface fields

Table 34: profile fields *(Continued)*

Settings	enhanced-hash-key profile field	Description
Flexbytes	no-in-flexbytes	Inner L3/L4 flexbytes fields
	no-flexbytes	L3/L4 flexbytes fields
Tunnel	no-tunnel	Tunnel header fields
	no-in-tunnel	Inner tunnel header fields
Traffic Class	no-in-layer3-pri	Inner L3 dscp/dp fields
	no-layer3-pri	L3 dscp/dp fields
Layer 2	no-in-layer2	Inner layer2 fields
	no-layer2	DMAC, SMAC and VLAN fields
Layer 3	no-in-layer3-proto	Inner L3 protocol fields
	no-layer3-proto	L3 protocol fields
	no-layer3-source-address	L3 source address fields
	no-layer3-destination-address	L3 destination address fields
	no-in-layer3	Inner layer3 fields
Layer 4	no-layer4-source-port	L4 source port fields
	no-layer4-destination-port	L4 destination port fields
	no-in-layer4	Inner layer4 fields

Profile fields are enabled by default and are only disabled if explicitly set. As an example, if load-balancing profile `profile1` has been set for a route, interface fields are enabled by default in profile1, unless explicitly disabled with the `set forwarding-options enhnaced-hash-key profile1 no-interface` configuration command.

Default

Default profile - The default profile or global profile is applied internally when neither profile1 or profile2 is applied to a route. You cannot use "default" as a profile name.

The final hash result is determined from an AND operation between hash fields of default profile and hash fields of profile1/profile2.

Table 35: Final hash results from default profile and profile1/profile2 AND operation

Default profile hash field enabled/ disabled (1/0)	profile1/profile2 hash field enabled/disabled (1/0)	Final hash result
1	1	1
1	0	0
0	1	0
0	0	0

As an example, `no-destination-port` and `no-source-port` enhanced-hash-key fields (disabled by default) are part of the default profile. Their equivalent profile fields in profile1/profile2 are `no-layer4-source-port` and `no-layer4-destination-port` (enabled by default). To enable `no-destination-port` and `no-source-port` enhanced-hash-key fields of the default profile, the following global CLI can be used:

```
set forwarding-options enhanced-hash-key family inet no-destination-port;
set forwarding-options enhanced-hash-key family inet no-source-port;
```

The following table shows the final hash computation when `no-destination-port` and `no-source-port` is enabled globally. Note that you can use `show forwarding-options enhanced-hash-key` *`profile1/profile2`* to view the final hash computation results.

Table 36: no-destination-port and no-source-port Final Hash

Default profile no-destination-port and no-source-port enhanced-hash-key fields enabled/disabled (1/0)	profile1/profile2 no-layer4-source-port and no-layer4-destination-port profile fields enabled/disabled (1/0)	Final hash result
1	1	1

To disable `no-destination-port` and `no-source-port` hashing only for a specific route while enabling `no-destination-port` and `no-source-port` hashing on all other routes, user can configure a profile for the route as follows.

```
set forwarding-options enhanced-hash-key profile1 no-layer4-source-port
set forwarding-options enhanced-hash-key profile1 no-layer4-destination-port
set policy-options policy-statement lb term 1 from route-filter 10.0.0.0/16 orlonger
set policy-options policy-statement lb term 1 then load-balance profile1
set policy-options policy-statement lb term 1 then accept
set policy-options policy-statement lb then accept
set routing-options forwarding-table export lb
```

The final hash result for `no-destination-port` and `no-source-port` hashing for route 10.0.0.0/16 would be "0" or disabled as indicated in the following table. `no-destination-port` and `no-source-port` hashing is still enabled for all other routes on the chassis.

Table 37: no-destination-port and no-source-port Final Hash for Specific Route

Default profile no-destination-port and no-source-port enhanced-hash-key fields enabled/disabled (1/0)	profile1 no-layer4-source-port and no-layer4-destination-port profile fields enabled/disabled (1/0)	Final hash result
1	0	0

Options

- profile1
- profile2

Required Privilege Level

- interface—To view this statement in the configuration.
- interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2

RELATED DOCUMENTATION

hash-seed (profile)

load-balance profile1/profile2

enhanced-hierarchical-policer

IN THIS SECTION

- Syntax | 6394
- Hierarchy Level | 6395
- Description | 6395
- Options | 6395
- Required Privilege Level | 6395
- Release Information | 6395

Syntax

```
enhanced-hierarchical-policer enhanced-hierarchical-policer-name
```

Hierarchy Level

```
[edit firewall]
```

Description

Configure up to four policers in an enhanced hierarchical policer configuration to police traffic from four traffic priority levels, namely, high, medium-high, medium-low, and low.

Options

Each of the four policers must be configured with only one option from the following list.

- high – The policer is assigned to rate-limit traffic from the 'high' traffic priority.
- medium-high - The policer is assigned to rate-limit traffic from the 'medium-high' traffic priority.
- medium-low - The policer is assigned to rate-limit traffic from the 'medium-low' traffic priority.
- low - The policer is assigned to rate-limit traffic from the 'low' traffic priority.

After which, each option needs to be further configured with commited-burst-size, committed-information-rate and max-committed-information-rate. For example, the following is a sample configuration of how medium-low can be configured.

```
set firewall enhanced-hierarchical-policer hpol medium-high committed-information-rate 10m
set firewall enhanced-hierarchical-policer hpol medium-high committed-burst-size 15k
set firewall enhanced-hierarchical-policer hpol medium-high max-committed-information-rate 15m
set firewall enhanced-hierarchical-policer hpol medium-high then discard
```

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.3.

enhanced-mode

IN THIS SECTION

- Syntax | 6396
- Hierarchy Level | 6396
- Description | 6396
- Required Privilege Level | 6398
- Release Information | 6398

Syntax

```
enhanced-mode;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name filter filter-name],
[edit firewall filter filter-name],
[edit firewall family family-name filter filter-name],
[edit logical-systems logical-system-name firewall filter filter-name],
[edit logical-systems logical-system-name firewall family family-name filter filter-name]
```

Description

Limit static service filters or API-client filters to term-based filter format only for inet or inet6 families when enhanced network services mode is configured at the `[edit chassis network-services]` hierarchy level. You cannot attach enhanced mode filters to local loopback, management, or MS-DPC interfaces. These interfaces are processed by the Routing Engine and DPC modules and can accept only compiled firewall filter format. In cases where both filter formats are needed for dynamic service filters, you can use the "enhanced-mode-override" on page 6401 statement on the specific filter definition to override the default filter term-based only format of chassis network-service enhanced IP mode.The `enhanced-mode` and the `enhanced-mode-override` statements are mutually exclusive; you can define the filter with either `enhanced-mode` or `enhanced-mode-override`, but not both.

> **NOTE**:
> For MX Series routers with MPCs, you need to initialize Trio-only match filters (that is, a filter that includes at least one match condition or action that is only supported by the Trio chipset) by walking the corresponding SNMP MIB. For example, for any filter that is configured or changed with respect to their Trio only filters, you need to run a command such as the following: `show snmp mib walk (ascii | decimal) object-id`. This forces Junos to learn the filter counters and ensure that the filter statistics are displayed. This guidance applies to all `enhanced-mode` firewall filters. It also applies to Firewall Filter Match Conditions for IPv4 Traffic with flexible match filter terms for offset-range or offset-mask, `gre-key`, and Firewall Filter Match Conditions for IPv6 Traffic with any of the following match conditions: `payload-protocol`, `extension headers`, `is_fragment`. It also applies to filters with either of the following Firewall Filter Terminating Actions: `encapsulate` or `decapsulate`, or either of the following Firewall Filter Nonterminating Actions: `policy-map`, and `clear-policy-map`.

When used with one of the chassis enhanced network services modes, firewall filters are generated in term-based format for use with MPC modules. Do not use enhanced mode for firewall filters that are intended for control plane traffic. Control plane filtering is handled by the Routing Engine kernel, which cannot use the term-based format of the enhanced mode filters.

If enhanced network services are not configured for the chassis, the `enhanced-mode` statement is ignored and any enhanced mode firewall filters are generated in both term-based and the default, compiled format. Only term-based (enhanced) firewall filters will be generated, regardless of the setting of the `enhanced-mode` statement at the [`edit chassis network-services`] hierarchy level, if any of the following are true:

- Flexible filter match conditions are configured at the [`edit firewall family` *`family-name`* `filter` *`filter-name`* `term` *`term-name`* `from`] or [`edit firewall filter` *`filter-name`* `term` *`term-name`* `from`] hierarchy levels.
- A tunnel header push or pop action, such as GRE encapsulate or decapsulate is configured at the [`edit firewall family` *`family-name`* `filter` *`filter-name`* `term` *`term-name`* `then`] hierarchy level.
- Payload-protocol match conditions are configured at the [`edit firewall family` *`family-name`* `filter` *`filter-name`* `term` *`term-name`* `from`] or [`edit firewall filter` *`filter-name`* `term` *`term-name`* `from`] hierarchy levels.
- An extension-header match is configured at the [`edit firewall family` *`family-name`* `filter` *`filter-name`* `term` *`term-name`* `from`] or [`edit firewall filter` *`filter-name`* `term` *`term-name`* `from`] hierarchy levels.
- A match condition is configured that only works with MPC cards, such as firewall bridge filters for IPv6 traffic.

For packets sourced from the Routing Engine, the Routing Engine processes Layer 3 packets by applying output filters to the packets and forwards Layer 2 packets to the Packet Forwarding Engine for

transmission. By configuring the enhanced mode filter, you explicitly specify that only the term-based filter format is used, which also implies that the Routing Engine cannot use this filter.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement introduced in Junos OS Release 23.2 for SRX4600, SRX5400, SRX5600, and SRX5800.

Statement introduced in Junos OS Release 24.4 for SRX4700.

RELATED DOCUMENTATION

- Network Services Mode Overview
- *Firewall Filters and Enhanced Network Services Mode Overview*
- *Configuring a Filter for Use with Enhanced Network Services Mode*
- Firewall Filter Match Conditions for IPv4 Traffic
- Firewall Filter Match Conditions for IPv6 Traffic
- Firewall Filter Terminating Actions
- Firewall Filter Flexible Match Conditions

enhanced-mode (Network Services)

IN THIS SECTION

- Syntax | 6399
- Hierarchy Level | 6399
- Description | 6399
- Default | 6400

Required Privilege Level | 6401
Release Information | 6401

Syntax

```
enhanced-mode;
```

Hierarchy Level

```
[edit chassis network-services]
```

Description

PTX1000, PTX3000 and PTX5000 Routers—When you configure `enhanced-mode` on PTX series routers, the following features are enabled on the router:

- Filter-based generic routing encapsulation (GRE) for IPV4 and IPV6 tunneling.
- `promote gre-key` statement for configuring gre-key as one of the matches in a filter.
- `promote dscp` under `family inet`, and `promote traffic-class` under `family inet6` statements for configuring these as one of the matches in a filter. (Added in Junos OS 17.1 for all PTX Series routers with third-generation FPCs.)
- `gtp-tunnel-endpoint-identifier` statement for including hash calculation for IPv4 or IPv6 packets in the GPRS tunneling protocol–Tunnel end point ID (GTP-TEID) field hash calculations.
- Wider configuration range for Bidirectional Forwarding Detection (BFD) protocol intervals.
- Support for up to two million routes per chassis.
- Support for Layer 3 VPN. The `vrf-table-label` statement is supported. (Added in Junos OS 15.1F5 for PTX5000 routers.)
- Support for destination class usage (DCU) and source class usage (SCU) accounting. (Added in Junos OS 15.1F5 for PTX5000 routers.)

NOTE:

- When you configure the `enhanced-mode` statement, only third-generation FPCs are allowed to be powered on. All other FPCs are powered off and cannot be brought online.
- When you do not configure the `enhanced-mode` statement, third-generation FPCs do not support the advanced features in the preceding list. Third-generation FPC only provide the same functionality as the first-generation and second-generation FPCs.
- After you configure the `enhanced-mode` statement and commit the configuration, the router must reboot.

T4000 Routers—When you configure `enhanced-mode` on T4000 routers, improved virtual private LAN service (VPLS) MAC address learning by supporting up to 262,143 MAC addresses per VPLS routing instance is enabled.

NOTE:

- The `enhanced-mode` statement supports up to 262,143 MAC addresses per VPLS routing instance. However, the MAC address learning limit for each interface remains the same (that is, 65,535 MAC addresses).
- After you configure the `enhanced-mode` statement and committing the configuration, you receive a warning message that prompts you to reboot the router. You must reboot the router and then modify the size of the VPLS MAC address table; otherwise, the improved VPLS MAC address learning does not take effect.
- When the T4000 router reboots after the `enhanced-mode` statement is configured, only the T4000 Type 5 FPCs are online while the remaining FPCs are offline.

Default

PTX Series Routers, except the PTX10000 line of routers—By default, the `enhanced-mode` statement is disabled.

PTX 10000 line of routers (such as PTX10001, PTX10003, PTX10008, and PTX10016 routers)—By default, the `enhanced-mode` statement is enabled.

T4000 Routers—By default, the improved VPLS MAC address learning feature is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Network Services Mode Overview

show chassis fpc | **18311**

mac-table-size (VPLS) | **10269**

show chassis network-services | **18571**

enhanced-mode-override

IN THIS SECTION

- Syntax | **6401**
- Hierarchy Level | **6402**
- Description | **6402**
- Required Privilege Level | **6402**
- Release Information | **6402**

Syntax

```
enhanced-mode-override;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name filter filter-name],
[edit firewall filter filter-name],
[edit firewall family family-name filter filter-name],
[edit logical-systems logical-system-name firewall filter filter-name],
[edit logical-systems logical-system-name firewall family family-name filter filter-name]
```

Description

Overrides the default filter enhanced-mode of dynamic service filters when the chassis is running in network-services enhanced IP mode. It functions similarly to the "enhanced-mode" on page 6396 statement used to override the default IP mode of static filters when the chassis is running in network-services enhanced IP mode.

When the chassis is running in network-service enhanced IP mode, all dynamic service inet and inet6 firewall filters are automatically generated in term-based filter format only. For any dynamic service filter that must be generated in both term-based and compiled formats, you must specifically configure the `enhanced-mode-override` statement for that filter definition.

Similar to how the filter `enhanced-mode` statement functions, if the chassis is not running in network-services enhanced IP mode, then the `enhanced-mode-override` statement is ignored.

> **NOTE**: The `enhanced-mode` and the `enhanced-mode-override` statements are mutually exclusive; you can define the filter with either `enhanced-mode` or `enhanced-mode-override`, but not both.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

enhanced-mode | 6396

Network Services Mode Overview

Firewall Filters and Enhanced Network Services Mode Overview

Configuring a Filter for Use with Enhanced Network Services Mode

enhanced-policer

IN THIS SECTION

- Syntax | 6403
- Hierarchy Level | 6403
- Description | 6403
- Required Privilege Level | 6404
- Release Information | 6404

Syntax

```
enhanced-policer
```

Hierarchy Level

```
[edit chassis]
```

Description

Collect additional statistics to be displayed using `show` commands. An FPC restart is required after changing this configuration.

A warning log message is generated when you commit a configuration that contains the `enhanced-policer` statement. The log message states that the enhanced policer is enabled on FPCs only after they are restarted. If you do not reboot the FPCs, the FPCs return all 0s (zeros) when you perform a query for the retrieval of detailed statistics—for example, when you issue the `show firewall detail` command.

BEST PRACTICE: We recommend that you do not use enhanced policers in highly scaled environments. Enhanced policer statistics require significantly more memory for statistics counters than do simple policers. The memory requirements for enhanced policers might exceed your resource memory limit, depending on the number of subscribers in your configuration and on the number of policers configured for the subscribers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

- Router Chassis Configuration Statements
- *Enhanced Policer Statistics Overview*
- show policer
- show firewall | 20048

enhanced-priority-mode

IN THIS SECTION

- Syntax | 6405
- Hierarchy Level | 6405
- Description | 6405
- Required Privilege Level | 6405
- Release Information | 6405

Syntax

```
enhanced-priority-mode;
```

Hierarchy Level

```
[edit chassis fpc slot-number traffic-manager]
```

Description

Enable the enhanced priority mode. When you enable the enhanced priority mode, the scheduler supports four additional per-priority shaping rates and two additional excess priorities at the interface and interface set level. The four additional per-priority shaping rates are: Guaranteed Strict-high, Guaranteed Medium-low, Excess medium-high, and Excess medium-low. The two additional excess priorities are: Excess-rate Medium-high and Excess-rate Medium-low.

NOTE: The line card reboots when you enable or delete the enhanced priority mode feature.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Statement introduced in Junos OS Release 18.1 for PTX Series.

RELATED DOCUMENTATION

- excess-rate-medium-high | 6597
- excess-rate-medium-low | 6599
- *Per-Priority Shaping on MIC and MPC Interfaces Overview*
- shaping-rate-excess-medium-high | 16943

shaping-rate-excess-medium-low | 16945
shaping-rate-priority-medium-low | 16956
shaping-rate-priority-strict-high | 16958
traffic-manager | 30036
traffic-control-profiles | 30009

enhanced-transmission-selection

IN THIS SECTION

- Syntax | 6406
- Hierarchy Level | 6406
- Description | 6407
- Options | 6407
- Required Privilege Level | 6407
- Release Information | 6407

Syntax

```
enhanced-transmission-selection {
    no-auto-negotiation;
    no-recommendation-tlv;
    recommendation-tlv {
        no-auto-negotiation;
    }
}
```

Hierarchy Level

```
[edit protocols dcbx interface interface-name]
```

Description

Disable advertising the enhanced transmission selection (ETS) state of the interface to the peer. To disable ETS on the interface, do not enable ETS on the interface in the class-of-service (CoS) configuration.

Disabling ETS autonegotiation stops the QFX Series from advertising the ETS Configuration TLV and the ETS Recommendation TLV.

Disabling the ETS recommendation TLV stops the QFX Series from advertising the ETS Recommendation TLV, but the ETS Configuration TLV is still advertised.

Options

`no-auto-negotiation`—Disable automatic negotiation of ETS (Configuration TLV and Recommendation TLV)

`no-recommendation-tlv`—Disable automatic negotiation of the ETS Recommendation TLV

`recommendation-tlv`—Enable automatic negotiation of ETS Recommendation TLV

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show dcbx neighbors | **19191**

Configuring DCBX Autonegotiation

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Understanding DCB Features and Requirements

enrollment

IN THIS SECTION

- Syntax | 6408
- Hierarchy Level | 6408
- Description | 6408
- Options | 6408
- Required Privilege Level | 6409
- Release Information | 6409

Syntax

```
enrollment {
    url url-name;
    retry number-of-enrollment-attempts;
    retry-interval seconds;
}
```

Hierarchy Level

```
[edit security pki ca-profile ca-profile-name]
```

Description

Specify the URL and enrollment parameters of the certificate authority (CA) for Adaptive Services (AS) and MultiServices PICs installed on MX Series, M Series, and T Series routers.

Options

`url` *`url-name`*—Location of the CA to which the router sends the Simple Certificate Enrollment Protocol-based (SCEP-based) certificate enrollment requests for the configured CA profile. Use the CA host DNS name or IP address.

`retry` *`number-of-enrollment-attempts`*—Number of enrollment retries.

- **Range:** 0 through 100
- **Default:** 0

`retry-interval` *`seconds`*—Length of time, in seconds, that a router should wait between enrollment attempts.

- **Range:** 0 through 3600
- **Default:** 0

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Specifying an Enrollment URL

Specifying the Enrollment Properties

enrollment-retry

IN THIS SECTION

- Syntax | 6410
- Hierarchy Level | 6410
- Description | 6410
- Options | 6410
- Required Privilege Level | 6410

Release Information | 6410

Syntax

```
enrollment-retry attempts;
```

Hierarchy Level

```
[edit security certificates]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Specify how many times a router or switch can resend a digital certificate request.

Options

attempts—Number of enrollment retries.

- **Range:** 0 through 100
- **Default:** 0

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Number of Enrollment Retries

enrollment-url

IN THIS SECTION

- Syntax | 6411
- Hierarchy Level | 6411
- Description | 6411
- Options | 6411
- Required Privilege Level | 6412
- Release Information | 6412

Syntax

```
enrollment-url url-name;
```

Hierarchy Level

```
[edit security certificates certification-authority ca-profile-name]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Specify where your router or switch sends Simple Certificate Enrollment Protocol-based (SCEP-based) certificate enrollment requests (certificate authority URL).

Options

url-name—Certificate authority URL.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Specifying an Enrollment URL

enterprise-oid

IN THIS SECTION

- Syntax | 6412
- Hierarchy Level | 6412
- Description | 6413
- Required Privilege Level | 6413
- Release Information | 6413

Syntax

```
enterprise-oid;
```

Hierarchy Level

```
[edit snmp trap-options]
```

Description

Add the **snmpTrapEnterprise** object, which shows the association between an enterprise-specific trap and the organization that defined the trap, to standard SNMP traps. By default, the **snmpTrapEnterprise** object is added only to the enterprise-specific traps. When the `enterprise-oid` statement is included in the configuration, **snmpTrapEnterprise** is added to all the traps generated from the device.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0

RELATED DOCUMENTATION

Configuring SNMP Trap Options

entropy-label (protocols bgp)

IN THIS SECTION

- Syntax | 6414
- Hierarchy Level | 6414
- Description | 6415
- Options | 6415
- Required Privilege Level | 6415
- Release Information | 6415

Syntax

```
entropy-label {
      elc-v2-compatible;
      import [ import ... ];
      no-next-hop-validation;
     }
```

Hierarchy Level

```
[edit fabric protocols bgp family inet labeled-unicast],
[edit fabric protocols bgp group name family inet labeled-unicast],
[edit fabric protocols bgp group name neighbor name family inet labeled-unicast],
[edit logical-systems name protocols bgp family inet labeled-unicast],
[edit logical-systems name protocols bgp group name family inet labeled-unicast],
[edit logical-systems name protocols bgp group name neighbor name family inet labeled-unicast],
[edit logical-systems name routing-instances name protocols bgp family inet labeled-unicast],
[edit logical-systems name routing-instances name protocols bgp group name family inet labeled-
unicast],
[edit logical-systems name routing-instances name protocols bgp group name neighbor name family
inet labeled-unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp family inet labeled-
unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp group name family
inet labeled-unicast],
[edit logical-systems name tenants name routing-instances name protocols bgp group name neighbor
name family inet labeled-unicast],
[edit protocols bgp family inet labeled-unicast],
[edit protocols bgp group name family inet labeled-unicast],
[edit protocols bgp group name neighbor name family inet labeled-unicast],
[edit routing-instances name protocols bgp family inet labeled-unicast],
[edit routing-instances name protocols bgp group name family inet labeled-unicast],
[edit routing-instances name protocols bgp group name neighbor name family inet labeled-unicast],
[edit tenants name routing-instances name protocols bgp family inet labeled-unicast],
[edit tenants name routing-instances name protocols bgp group name family inet labeled-unicast],
[edit tenants name routing-instances name protocols bgp group name neighbor name family inet
labeled-unicast]
```

Description

Insert the entropy label into the BGP labeled unicast LSP packets, which assists the transit router in load-balancing BGP traffic across equal-cost multipaths or link aggregation groups. The ingress label edge router inspects the flow information of a packet and maps it to an entropy label, which is inserted into the BGP label stack. LSRs in the core use this entropy label as the key to hash the packet and direct it to the correct path.

Options

elc-v2-compatible	Send and accept ELCv2 in addition to ELCv3
import	Policy to select BGP LSPs to use entropy label
no-next-hop-validation	Don't validate next hop field against route next hop

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

`elc-v2-compatible` statement introduced in Junos OS Release 23.1.

RELATED DOCUMENTATION

- labeled-unicast (Protocols BGP) | **9603**
- policy-statement | **12736**
- *Configuring an Entropy Label for a BGP Labeled Unicast LSP*
- *Example: Configuring an Entropy Label for a BGP Labeled Unicast LSP*
- *Understanding Entropy Label for BGP Labeled Unicast LSP*

entropy-label (protocols mpls)

IN THIS SECTION

- Syntax | 6416
- Hierarchy Level | 6416
- Description | 6416
- Options | 6416
- Required Privilege Level | 6417
- Release Information | 6417

Syntax

```
entropy-label {
    ingress-policy ingress-policy-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp]
```

Description

Assists the transit router in load-balancing MPLS traffic across ECMP paths or Link Aggregation groups by introducing the entropy label to the MPLS label stack. The entropy label allows routers to load balance MPLS traffic by using a hash-input without the need to perform deep packet inspection. Deep packet inspection requires more of the router's processing power and is not a capability shared by all routers.

Options

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring the Entropy Label for LSPs

environment-monitoring-policy (satellite-management)

IN THIS SECTION

- Syntax | 6417
- Hierarchy Level | 6418
- Description | 6418
- Default | 6418
- Options | 6418
- Required Privilege Level | 6418
- Release Information | 6419

Syntax

```
environment-monitoring-policy policy-name;
```

Hierarchy Level

```
[edit chassis satellite-management]
[edit chassis satellite-management fpc slot-id]
```

Description

Enable an environment monitoring policy in a Junos Fusion.

You configure environment monitoring policies for a Junos Fusion in the [`edit policy-options environment-monitoring-policy` *`policy-name`*] hierarchy.

You can configure an environment monitoring policy in a Junos Fusion for a single satellite device using the **fpc *slot-id*** option, or for all satellite devices in the Junos Fusion by not specifying the **fpc *slot-id*** option.

You can configure a different environment monitoring policy for a single satellite device using the **fpc *slot-id*** when an environment monitoring policy for all satellite devices is configured. The environment monitoring policy for the FPC is enabled in cases when both an individual and global environment monitoring policy are configured.

Default

No environment monitoring policies for the Junos Fusion are present.

If you enable an environment monitoring policy in a Junos Fusion without specifying the **fpc *slot-id*** option, the environment monitoring policy is applied for all satellite devices in the Junos Fusion.

Options

policy-name

Specifies the name of the environment monitoring policy.

The *policy-name* name is defined as part of the environment monitoring policy configuration procedure, which is handled in the [`edit policy-options environment-monitoring-policy` *`policy-name`*] hierarchy.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Satellite Device Alarm Handling Using an Environment Monitoring Satellite Policy in a Junos Fusion

Understanding Satellite Policies in a Junos Fusion

environment-monitoring-policy (satellite-policies)

IN THIS SECTION

- Syntax | 6419
- Hierarchy Level | 6420
- Description | 6420
- Options | 6420
- Required Privilege Level | 6420
- Release Information | 6420

Syntax

```
environment-monitoring-policy policy-name {
    alarm {
        linkdown [ignore | red | yellow]
    }
    term term-name{
        from {
            product-model model-name;
        }
    }
}
```

Hierarchy Level

```
[edit policy-options satellite-policies]
```

Description

Configure an environment monitoring satellite policy for a device or devices in a Junos Fusion.

An environment monitoring satellite policy is used to configure alarm behavior on satellite devices in a Junos Fusion.

The environment monitoring policy is applied to a Junos Fusion using the `environment-monitoring-policy` statement in the [edit chassis "satellite-management" on page 16310] or [edit chassis "satellite-management" on page 16310 "fpc" on page 7512 *slot-id*] hierarchy levels.

Options

policy-name Specifies the user-defined name of the environment monitoring policy.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

> *Configuring Satellite Device Alarm Handling Using an Environment Monitoring Satellite Policy in a Junos Fusion*
>
> *Understanding Satellite Policies in a Junos Fusion*

epacl-firewall-optimization

IN THIS SECTION

- Syntax | 6421
- Hierarchy Level | 6421
- Description | 6421
- Default | 6422
- Required Privilege Level | 6422
- Release Information | 6422

Syntax

```
epacl-firewall-optimization
```

Hierarchy Level

```
[edit chassis forwarding-options]
```

Description

Enable `epacl-firewall-optimization` to use Layer 2 and Layer 3 match conditions in firewall filters to support micro-segmentation on VXLAN deployments. Filtering in both the ingress and egress directions is supported. (For egress filtering on VLANs, this statement is not needed.)

For example, to create micro-segmentation in a VXLAN, you need to enable the `epacl-firewall-optimization` statement at the `[chassis]` level of the hierarchy, and then create the firewall rules with the match conditions that you want to filter on.

For both VLANs and VXLANs, you can use the following match conditions:

- `ip-source-address`
- `ip-destination-address`
- `destination-port`

- `user-vlan-id`
- `source-mac-address`
- `destination-mac-address`
- `ip-protocol`

Valid actions are `accept`, `count`, and `discard`.

The configuration sample below shows how to configure a QFX5110 Series switch that is part of a VXLAN to provide Layer 2 filtering in the egress direction. First we enable `epacl-firewall-optimization` on the device, and then we create a Layer 2 egress firewall filter named **epacl**, and attach it to the `xe-0/0/10.0` interface. The first term tells the switch to accept and count packets from the specified source MAC address (00:00:5e:00:53:a1/48). The second term tells the interface to count and discard all other packets.

```
set chassis forwarding-options epacl-firewall-optimization
set firewall family ethernet-switching filter epacl term t1 from source-mac-address
00:00:5e:00:53:a1/48
set firewall family ethernet-switching filter epacl term t1 then accept
set firewall family ethernet-switching filter epacl term t1 then count epacl-accept
set firewall family ethernet-switching filter epacl term t2 then discard
set firewall family ethernet-switching filter epacl term t2 then count epacl-discard
set interfaces xe-0/0/10 unit 0 family ethernet-switching filter output epacl
```

Default

Not enabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1 for QFX5110 and QFX5120 Series switches.

ephemeral

IN THIS SECTION

- Syntax | 6423
- Hierarchy Level | 6424
- Description | 6424
- Options | 6425
- Required Privilege Level | 6426
- Release Information | 6426

Syntax

```
ephemeral {
    allow-commit-synchronize-with-gres;
    commit-synchronize-model (asynchronous | synchronous);
    cyclic-version-for-ephemeral-default version;
    delete-ephemeral-default;
    ignore-ephemeral-default;
    instance instance-name {
        cyclic-version version;
        resize {
            database-size-diff size MB;
            database-size-on-disk size MB;
        }
    }
    resize-ephemeral-default {
       database-size-diff size MB;
       database-size-on-disk size MB;
    }
}
```

Hierarchy Level

```
[edit system configuration-database]
```

Description

Configure settings for the ephemeral configuration database.

The ephemeral database is an alternate configuration database that enables Juniper Extension Toolkit (JET) applications and NETCONF and Junos XML protocol client applications to simultaneously load and commit configuration changes on Junos devices and with significantly greater throughput than when committing data to the candidate configuration database. Junos devices provide a default ephemeral database instance as well as the ability to configure multiple user-defined instances of the ephemeral configuration database.

The ephemeral database is not subject to the same verification required in the static configuration database. As a result, the ephemeral configuration database does not support configuration groups or interface ranges, or macros, commit scripts, or translation scripts. Additionally, certain configuration statements cannot be configured through the ephemeral database as described in *Unsupported Configuration Statements in the Ephemeral Configuration Database*. A Junos device validates the syntax but does not validate the semantics of configuration data committed to the ephemeral database. Therefore, all configuration data must be validated before loading it into the ephemeral database and committing it on the device. If you commit invalid configuration data to the database, it can cause Junos processes to restart or even crash and result in disruption to the system or network.

NOTE: When you configure statements at the `[edit system configuration-database ephemeral]` hierarchy level and commit the configuration, all Junos processes must check and evaluate their complete configuration, which might cause a spike in CPU utilization, potentially impacting other critical software processes.

NOTE: When you use the ephemeral configuration database, commit operations on the static configuration database might take longer, because additional operations must be performed to merge the static and ephemeral configuration data.

By default, the ephemeral database performs commit synchronize operations asynchronously. We do *not* recommend using the ephemeral database with the asynchronous commit model on devices that have graceful Routing Engine switchover (GRES) enabled. If you elect to use the ephemeral database when GRES is enabled, you must explicitly configure the `allow-commit-synchronize-with-gres` statement to enable the device to synchronize ephemeral configuration data to the backup Routing Engine when you request a commit synchronize operation on an ephemeral instance. The ephemeral database also

supports a synchronous commit model for commit synchronize operations. The synchronous commit model enables you to use the ephemeral database with greater reliability on devices that have high availability features enabled.

Ephemeral configuration data does not persist across reboots. In addition, when you install a package that requires rebuilding the Junos schema, such as an OpenConfig or YANG package, the device deletes all ephemeral configuration data in the process of rebuilding the schema.

Options

allow-commit-synchronize-with-gres

Enable a device that has GRES enabled and that uses the asynchronous commit synchronize model for the ephemeral database to synchronize an ephemeral instance to the backup Routing Engine when you request a commit synchronize operation on the instance.

commit-synchronize-model (asynchronous | synchronous)

Specify how the device synchronizes ephemeral configuration data in a dual Routing Engine or Virtual Chassis environment.

Synchronous commit operations are slower than asynchronous commit operations but provide better assurance that the ephemeral configuration is synchronized across Routing Engines and Virtual Chassis members.

- **Default:** `asynchronous`

cyclic-version-for-ephemeral-default *version*

Number of versions of the default ephemeral database for which the system stores deleted configuration objects. After each commit, the device reclaims the space occupied by objects deleted in the previous database version relative to the current database version as determined by the cyclic version value. You can disable cyclic versioning by setting the value to 0.

Thus, starting with the 11th commit operation, the device, by default, reclaims the space occupied by objects deleted in version 1 of the database. The process continues with each subsequent commit operation.

- **Default:** 10
- **Range:** 0 through 4,294,967,295

delete-ephemeral-default

Delete the configuration data and files for the default instance of the ephemeral configuration database. When you configure this statement, you must also configure the `ignore-ephemeral-default` statement.

ignore-ephemeral-default

Disable the default instance of the ephemeral configuration database.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.2R2 and Junos OS Evolved Release 22.1R1.

`commit-synchronize-model` option added in Junos OS Release 21.1R1.

`delete-ephemeral-default` statement added in Junos OS Release 22.1R1.

`cyclic-version-for-ephemeral-default` statement added in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Understanding the Ephemeral Configuration Database

Enabling and Configuring Instances of the Ephemeral Configuration Database

Committing and Synchronizing Ephemeral Configuration Data Using the NETCONF or Junos XML Protocol

equals

IN THIS SECTION

- Syntax | 6427
- Hierarchy Level | 6427
- Description | 6427
- Options | 6428
- Required Privilege Level | 6428
- Release Information | 6428

Syntax

```
equals {
    value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from called-station-id],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from calling-station-id],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from class],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from framed-ip-address],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from framed-ipv6-prefix],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from 3gpp-imsi],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from nas-ip-address],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name format integer],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name format string],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name format time],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name format v4address],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name format v6address],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name format v6prefix]
```

Description

Specify the value that the RADIUS attribute must equal.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

Options

value Value that the RADIUS attribute must equal.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

equals (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 6429
- Hierarchy Level | 6429
- Description | 6431
- Options | 6432
- Required Privilege Level | 6432
- Release Information | 6432

Syntax

```
equals {
    ascii name {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
    hexadecimal name {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-77],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-77],
[edit forwarding-options dhcp-relay group name relay-option],
[edit forwarding-options dhcp-relay group name relay-option option-60],
[edit forwarding-options dhcp-relay group name relay-option option-77],
[edit forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay relay-option option-60],
[edit forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
```

```
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-77],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option option-60],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option option-77],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option option-60],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option option-77],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option option-60],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option option-77],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option
option-60],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option
option-77],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option option-60],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option option-77],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option option-60],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option option-77],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option
```

```
option-60],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option
option-77],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option option-77],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-77],
[edit routing-instances name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-77],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-60],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-77],
[edit vlans name forwarding-options dhcp-relay relay-option],
[edit vlans name forwarding-options dhcp-relay relay-option option-60],
[edit vlans name forwarding-options dhcp-relay relay-option option-77]
```

Description

Configure the exact match criteria used with the DHCP relay agent selective processing feature. DHCP relay agent compares the configured match string with the option-specific string received in DHCP

client packets. If there is an exact left-to-right match, DHCP performs the action you define for the match criteria.

You can configure an unlimited number of match strings. Match strings do not support wildcard attributes.

The `local-server-group` option is not supported for DHCPv6 relay agent.

Options

ascii-string ASCII string of 1 through 255 alphanumeric characters.

hexadecimal-string Hexadecimal string of 1 through 255 hexadecimal characters (0 through 9, a through f, A through F).

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Statement updated in Junos OS Release 17.4.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address

equals (Dynamic Profile)

IN THIS SECTION

- Syntax | 6433
- Hierarchy Level | 6433
- Description | 6433
- Options | 6433
- Required Privilege Level | 6433
- Release Information | 6434

Syntax

```
equals expression;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name variables variable-name]
```

Description

Configure an expression—a group of arithmetic operators, string operators, and operands—for a user-defined variable that is evaluated at run time and returned as the variable value.

Options

expression Expression evaluated to return a value for the user-defined variable.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring User-Defined Dynamic Variables in Dynamic Profiles

User-Defined Variables

Using Variable Expressions in User-Defined Variables

equals (Event Policy)

IN THIS SECTION

- Syntax | 6434
- Hierarchy Level | 6434
- Description | 6435
- Default | 6435
- Options | 6435
- Required Privilege Level | 6435
- Release Information | 6435

Syntax

```
event1.attribute-name equals event2.attribute-name;
```

Hierarchy Level

```
[edit event-options policy policy-name attributes-match]
```

Description

Configure an event policy to trigger only if the specified attribute of *event1* equals the specified attribute of *event2*.

Default

Options

event1.attribute-name	Attribute of one event.
event2.attribute-name	Attribute of another event.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Use Correlated Events to Trigger an Event Policy

eracl-ip6-match (packet-forwarding-options)

IN THIS SECTION

- Syntax | 6436
- Hierarchy Level | 6436
- Description | 6436
- Options | 6436

Required Privilege Level | 6437
Release Information | 6437

Syntax

```
eracl-ip6-match {
    (srcip6-and-destip6 | srcip6-only);
    }
}
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

Use the options of this command to allow source and/or destination IPv6 address match conditions for eRACL inet6 filters.

In Junos, firewall filters are classified as *ingress* or *egress* depending on where in the sequence the packet is evaluated and action taken. Filtering IPv6 traffic on an `inet6` egress interface can be useful, for example, for safeguarding a third-party device connected to the Juniper switch.

NOTE: After configuring, modifying, or deleting the `eracl-ip6-match` statement, you must commit the configuration, and the packet forwarding engine (PFE) must be restarted.

Options

eracl-ip6-match Configuring match conditions in a firewall filter for IPv6 source and/or destination IP addresses is only allowed if the `srcip6-and-destip6` or the `srcip6-only` options described below are enabled. The two options cannot both be enabled at the same time. If neither option is configured, the default behavior is to allow match condition to be created for IPv6 destination addresses on egress interfaces only.

- Values:
  - **srcip6-and-destip6**—Choose this option to allow both source and destination IPv6 address match conditions on inet6 interfaces in egress direction. The source and destination port match conditions are also allowed only with this option. Note that when this option is enabled, the scale of eRACLv6 is reduced by half.
  - **srcip6-only**—Choosing this option allows the source IPv6 address match condition in eRACLv6 filters but not a destination address. Both source and destination port match conditions cannot be configured at the same time as this option is enabled (you will get a commit error).

Required Privilege Level

flow-tap

Release Information

Statement introduced in Junos OS Release 19.1 (EX4300 and QFX5100 Series switches only).

RELATED DOCUMENTATION

Example: Configuring an Egress Filter Based on IPv6 Source or Destination IP Addresses

eracl-profile (packet-forwarding-options)

IN THIS SECTION

- Syntax | 6438
- Hierarchy Level | 6438
- Description | 6438
- Options | 6438
- Required Privilege Level | 6439
- Release Information | 6439

Syntax

```
eracl-profile {
    eracl-scale;
}
```

Hierarchy Level

```
[edit system packet-forwarding-options firewall]
[edit system packet-forwarding-options firewall eracl-profile eracl-scale]
```

Description

Use the option of this command to configure egress firewall filters, also known as eRACLs, in scaled mode. This feature is supported only in the egress direction (routed traffic exiting the device).

In Junos, firewall filters are classified as ingress or egress depending on where in the sequence the packet is evaluated and action taken. Filtering traffic on an egress interface can be useful, for example, for safeguarding a third-party device connected to the Juniper switch.

Options

eracl-scale Use this option to increase the number of egress firewall filters to 2000. When you configure an egress filter in scaled mode, the switch uses ingress TCAM space (IFP) to achieve the higher scale.

> **NOTE**: After configuring, modifying, or deleting the `eracl-profile` statement, you must commit the configuration, and the packet forwarding engine (PFE) must be restarted.

When you enable `eracl-scale` mode, the following restrictions apply:

- You can only apply a filter in the egress direction (traffic exiting the VLAN).
- Only `inet` and `inet6` protocol families are supported.
- Generic Routing Encapsulation (GRE) interfaces are not supported.

- You cannot apply filters with the same match condition to different egress VLANs or Layer 3 interfaces. The only supported actions are `accept`, `discard`, and `count`.
- Match conditions are programmed in the ingress firewall filter TCAM. This means that any counters attached to the filter counts traffic on any incoming VLANs.
- The `eracl-scale` option comes configured in global mode. When enabled, any of your existing egress filters will be automatically reinstalled in scaled mode.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 19.4R2 (QFX5220 switches).

RELATED DOCUMENTATION

Planning the Number of Firewall Filters to Create

Firewall Filter Match Conditions and Actions (QFX5220, QFX5700 and the QFX5130-32CD)

Configuring Firewall Filters

erroneous-ccm

IN THIS SECTION

- Syntax | 6440
- Hierarchy Level | 6440
- Description | 6440
- Required Privilege Level | 6440
- Release Information | 6440

Syntax

```
erroneous-ccm;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
event ais-trigger-condition]
```

Description

Configure the defect condition that raises an alarm indication signal when any cross-connect continuity check messages (CCMs) with an unexpected MEP ID or an erroneous maintenance domain level are received by the MEP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Alarm Indication Signal (ETH-AIS) Function Overview

Configuring ETH-AIS on a CFM MEP

error

IN THIS SECTION

- Syntax | 6441
- Junos OS | 6442
- Junos OS Evolved | 6442
- Description | 6442
- Options | 6443
- Required Privilege Level | 6445
- Release Information | 6445

Syntax

```
error {
    error_id {
        severity (fatal | major | minor)
        state {
            disable:
        }
    }
    major {
        threshold threshold-value;
        action
    }
    minor {
        threshold threshold-value;
        action
    }
    fatal {
        threshold threshold-value;
        action
    }
    scope error-scope {
        category category {
            severity-level {
```

```
                threshold threshold-value;
                action
            }
        }
        switch {
            category category {
                switch {
                    severity-level {
                        threshold threshold-value;
                        action
                    }
                }
            }
        }
    }
    reset-pfe {
        pause-period <pause_minutes>;
        pfe-disable-period <pfe_disable_minutes>;
        retry-limit <retry_number>;
    }
}
```

Junos OS

[edit chassis]

[edit chassis fpc *slot-number*]

[edit chassis sib slot *slot-number*]

Junos OS Evolved

[edit chassis fpc *slot-number*]

[edit chassis sib slot *slot-number*]

Description

Configure the threshold at which FPC or SIB errors will take the action you configure to be performed by the device. Starting from Junos OS Release 18.1R3, you can configure error thresholds and actions at the error scope and error category levels on MX Series routers.

Some Juniper devices include an internal framework for detecting and correcting FPC errors that can have the potential to affect services. You can classify the errors according to severity, set an automatic recovery action for each severity, and set a threshold (i.e., the number of times the error must occur before the action is triggered). Using the configuration command for error levels, you can isolate PFE errors, thereby reducing the need for a field replacement of FPCs. For feature details and PTX platforms supported, see FPC self-healing.

NOTE:

- On the MX104 routers, Junos does not initiate restart of the system on encountering a Fatal error. Additionally, though you can configure the action `disable-pfe` for Major errors on the MX104, the router does not disable its only PFE on encountering a Major error.

Options

You can configure the threshold for the following severity levels:

- `fatal`—Fatal error on the FPC. An error that results in blockage of considerable amount of traffic across modules is a fatal error.
- `major`—Major error on the FPC. An error that results in continuing loss of packet traffic but does not affect other modules is a major error.
- `minor`—Minor error on the FPC. An error that results in the loss of a single packet but is fully recoverable is a minor error.
- `threshold` *`threshold-value`*—Configure the threshold value at which to take action. If the severity level of the error is fatal, the action is carried out only once when the total number of errors crosses the threshold value. If the severity level of the error is major, the action is carried out once after the occurrence crosses the threshold. If the severity level is minor, the action is carried out as many times as the value specified by the threshold. For example, when the severity level is minor, and you have configured the threshold value as 10, the action is carried out after the tenth occurrence. On Junos OS Evolved, for the errors belonging to the `internal` category, the default threshold value is 1.

  NOTE: You can set the threshold value to 0 for errors with severity level as minor. This implies that no action is taken for that error. You cannot set the threshold value to 0 for errors with severity level as major or fatal.

  Default: The error count for fatal and major actions is 1. The default error count for minor actions is 10.

  Range: 0—429,496,729

The available detection and recovery actions are as follows:

- `alarm`—Raise an alarm.
- `disable-pfe`—Disable the PFE interfaces on the FPC.
- `get-state`—Get the current state of the FPC.
- `log`—Generate a log for the event.
- `offline`—Take the FPC offline.
- `offline-pic`—Take the PIC (installed in the FPC) offline.
- `reset`—Reset the FPC.
- `reset-pfe`—Reset the PFE. (Supported for PTX series devices including PTX10003)
- `offline-pfe`—Take the PFE offline.
- `trap`—Raise traps for the FPC errors.

> **NOTE**: Starting in Junos OS Evolved Release 19.1R1, the `offline` and `disable-pfe` actions are not available for errors with `minor` severity (under the hierarchy `edit chassis error minor action`).
>
> Starting in Junos OS Release 21.4R3, the additional options for reset PFE are valid only for line cards MPC7, MPC8, and MPC9, for the platforms MX240, MX480, MX960, MX2008, MX2010, and MX2020.

The available detection and recovery actions are as follows for devices running Junos OS Evolved:

- `alarm`—Raise an alarm.
- `fault`—System goes to fault state but stays up (diagnostics can be run on it).
- `get-state`—Get the current state of the FPC.
- `log`—Generate a log for the event.

> **NOTE**: Starting in Junos OS Release 17.2R1, if you configure the `disable-pfe`, `offline`, `offline-pic` or `reset` action on an MX Series or PTX Series router, the `get-state` action is additionally configured on the router. This means, for example, if you configure the `disable-pfe` action on the router, the router gets both `disable-pfe` and `get-state` actions configured.

- `scope` *`error-scope`*—Group the errors of a particular severity into different scopes. Errors belonging to each error scope is further grouped into categories, before thresholds and actions are defined at the group level. The following scopes are available: `board` and `pfe`. Junos OS Evolved also supports the scope `switch`.

- `category` *`category`*—Categorize errors into various subgroups under the scope level. An error category helps you group similar errors belonging to a particular scope and define actions for them at once. This feature eliminates the need for configurations against individual error-ids. Some of the error-categories are `functional`, `io` (input/output errors), `storage` (for example, errors related to HDD, SSD, and flash), `memory` (for example, errors related to static RAM), `processing` (for example, CPU-related errors), and `switch`. Junos OS Evolved also supports the category `internal`. On every occurrence of an error belonging to the `internal` category, the software by default raises an alarm at the individual error level (not at the scope or category level). You cannot configure an action against errors belonging to the `internal` category.

- *severity-level* – Configure the severity levels associated to each `error-id`. The options are `fatal`, `major` and `minor`.

- *`error-id`*—Use the error ID to disable an error or modify the error severity associated with that error. An *error-id*, which is a unique error identifier, is represented as a Uniform Resource Identifier (URI). For example, `/cpu/0/memory/0/memory-uncorrected-error` is an error ID that indicates an uncorrectable error under CPU memory module instance `0`.

- state—State

  - disable—Disable error.

- `reset-pfe`—Configure thresholds associated to reset-pfe action.

  - `pause-period` *`<pause_minutes>`*— Pause period in minutes. Valid range is 0 to 10000000.

  - `pfe-disable-period` *`<pfe_disable_minutes>`*— PFE disable period in minutes before reset PFE. Valid range is 1 to 10000000. The PFE disable period must be greater than the pause period.

  - `retry limit` *`<retry_number>`*— Retry limit for reset PFE. Valid range is 0 to 3.

Required Privilege Level

interface	To view this statement in the configuration.
interface-control	To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Reset-pfe error action support added to PTX10003 in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

Fabric Resiliency

Configuring FPC Error Levels and Actions

show chassis fabric errors | **18068**

show chassis fpc errors | **18387**

action (chassis error) | **1756**

error-id | **6449**

error (Security)

IN THIS SECTION

- Syntax | **6446**
- Hierarchy Level | **6447**
- Description | **6447**
- Options | **6447**
- Required Privilege Level | **6449**
- Release Information | **6449**

Syntax

```
error {
    (fatal | major | minor) {
    threshold threshold number;
    action (alarm | disable-pfe | offline-pic | log | get-state |offline | reset);
    }
}
```

Hierarchy Level

```
[edit chassis]

[edit chassis fpc slot-number]
```

Description

Configure the threshold at which FPC errors will take the action you configure to be performed by the device.

Some devices include an internal framework for detecting and correcting FPC errors that can have the potential to affect services. You can classify FPC errors according to severity, set an automatic recovery action for each severity, and set a threshold (the number of times the error must occur before the action is triggered).

However, the alarm is added to the default fault handling action list for a fatal error. Adding an alarm to the default fault handling list will allow the chassis alarm to trigger the RG1 switchover as soon as the fatal error is detected.

Starting with Junos OS Release 15.1-X49-D50 and Junos OS Release 17.3R1, this feature supports I/O cards (IOCs) and Services Processing Cards (SPCs) on the Junos OS SRX5000 line of devices. The following cards are supported:

- IOC2
- IOC3
- SPC2

CAUTION: Take care when setting the fault handling actions for SPC2 cards on the SRX5000 line of devices. Consider that if you set the fault handling action on an SPC2 card to offline or reset, when the card is either taken offline or the reboot occurs, the chassis daemon (chassisd) will reboot all of its FPC cards, both SPCs and IOCs—that is, the entire chassis will be rebooted.

Options

You can configure the threshold for the following severity levels:

NOTE: You cannot change the severity level of an error.

- `fatal`—Fatal error on the FPC. An error that results in the blockage of a considerable amount of traffic across modules is a fatal error. (default: raise an alarm and reset the FPC)
- `major`—Major error on the FPC. An error that results in continuing loss of packet traffic but does not affect other modules is a major error. (default: get the current state of the FPC and raise an alarm)
- `minor`—Minor error on the FPC. An error that results in the loss of a small number of packets but is fully recoverable is a minor error. (default: write a log for the event.)
- `threshold` *`threshold-value`*—Configure the threshold value at which to take action. If the severity level of the error is fatal, the action is carried out only once when the total number of errors exceeds the threshold value. If the severity level of the error is major, the action is carried out once after the occurrence exceeds the threshold. If the severity level is minor, the action is carried out as many times as the value specified by the threshold. For example, when the severity level is minor, and you have configured the threshold value as 10, the action is carried out after the tenth occurrence.

  > **NOTE**: You can set the threshold value to 0 for errors with a severity level of minor. This implies that no action is taken for that error. You cannot set the threshold value to 0 for errors with a severity level of major or fatal.

  Reset and offline are not listed as default actions for the minor error level for safety purposes.

  The alarm and reset default action is included in the implementation for the SRX5000 line of devices. It is required to trigger the RG1 switchover in a chassis cluster environment when an FPC fatal error occurs and the alarm being raised is a major alarm.

  Default: The error count for fatal and major actions is 1. The default error count for minor actions is 10.

  Table 38 on page 6448 shows the range of values for each error level.

  Table 38: Value Ranges for Error Levels

Error Level	Default Threshold	Value Range
Fatal	1	1 through 1024
Major	1	1 through 1024
Minor	10	0 through 1024

The available detection and recovery actions are as follows:

- `alarm`—Raise an alarm.
- `disable-pfe`—Disable the Packet Forwarding Engine interfaces on the FPC.
- `get-state`—Get the current state of the FPC.
- `log`—Generate a log for the event.
- `offline`—Take the FPC offline.
- `offline-pic`—Take the PIC (installed in the FPC) offline.
- `reset`—Reset the FPC.

Required Privilege Level

interface	To view this in the configuration.
interface-control	To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

fpc error | 7507

error-id

IN THIS SECTION

- Syntax | 6450
- Hierarchy Level | 6450
- Description | 6450
- Options | 6450
- Required Privilege Level | 6451

Release Information | 6451

Syntax

```
error-id {
    pfe;
    severity {
        (fatal | major | minor)
    }
    state {
        disable;
    }
}
```

Hierarchy Level

[edit chassis fpc *slot-number* error]

[edit chassis sib slot *slot-number* error]

Description

Manage an error by using the error identifier. This feature allows you to disable an error or modify its severity. The `disable` option allows you to stop the error from being reported in the system until the error is enabled again.

An *error-id*, a unique error identifier, is represented as a Uniform Resource Identifier (URI), and is composed of a module identifier and an error identifier. For example, the error-id "`/fpc/2/platformd/0/cm/0/precl/3/PRECL_CMERROR_INQ_BUF_OVERFLOW`" indicates a CM error due to buffer overflow.

Options

- `severity`—Apply a new severity to the error ID. You can apply any of the following severities:
  - `fatal`—Fatal error on the FPC. An error that results in blockage of considerable amount of traffic across modules is a fatal error.
  - `major`—Major error on the FPC. An error that results in continuing loss of packet traffic but does not affect other modules is a major error.

- `minor`—Minor error on the FPC. An error that results in the loss of a single packet but is fully recoverable is a minor error.

- `state`—By default, an error is in “enabled” state.
  - `disable`—Disable an error.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show chassis fpc errors | **18387**

Configuring FPC Error Levels and Actions

Managing FPC Errors

error-id | **6449**

es-options

IN THIS SECTION

- Syntax | **6452**
- Hierarchy Level | **6452**
- Description | **6452**
- Required Privilege Level | **6452**
- Release Information | **6452**

Syntax

```
es-options {
    backup-interface interface-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

On ES interfaces, configure ES interface-specific interface properties.

The `backup-interface` statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring ES PIC Redundancy

esi

IN THIS SECTION

- Syntax | 6453
- Hierarchy Level | 6454
- Description | 6454
- Options | 6455
- Required Privilege Level | 6456
- Release Information | 6456

Syntax

```
esi identifier {
    all-active | single-active;
    auto-derive {
        lacp;
    }
    df-election-granularity {
        per-esi {
            lacp-oos-on-ndf;
        }
        per-esi-vlan;
        per-evi;
    }
    df-election-type {
        mod;
        preference {
          least;
          non-revertive;
          value value;
        }
    }
    source-bmac <mac-address>;
}
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit protocols evpn interconnect],
[edit routing-instances routing-instance-name protocols evpn interconnect]
```

Description

Configure an Ethernet Segment Identifier (ESI) in either EVPN multihoming active/standby or active/active mode by using one of the following methods:

- Manually configure an ESI on a physical, aggregated Ethernet, or logical interface using the `esi identifier` configuration statement.
- Configure an aggregated Ethernet interface or aggregated Ethernet logical interface to automatically derive an ESI from the Link Aggregation Control Protocol (LACP) configuration. Use the `auto-derive lacp` configuration statement.

  If you configure automatic ESI on an aggregated Ethernet interface, the device generates an ESI. The device assigns that particular ESI to all logical interfaces on the aggregated Ethernet interface. You can also configure automatic ESI on one or more logical interfaces on an aggregated Ethernet interface. In that case, the device generates an ESI for each logical interface on which you enabled the feature.

> **NOTE**: You can't commit a configuration in which you try to configure the following on the same physical or logical aggregated Ethernet interface:
>
> - Manually configure an ESI on the aggregated Ethernet interface.
> - Enable the automatic ESI feature on the aggregated Ethernet interface.

Starting in Junos OS Releases 15.1F6 and 17.1R1 for MX Series routers and Junos OS Release 17.3R1 for EX9200 switches, you can specify an ESI on a logical interface. In earlier releases, you can configure an ESI only on a physical or aggregated Ethernet interface. For example:

```
set interfaces ae0 esi 00:11:22:33:44:55:66:77:88:99
```

When configuring an ESI, remember that the ESI is a factor in the designated forwarder (DF) election process. For example, let's look at a case in which you configure EVPN multihoming active/standby on aggregated Ethernet interface ae0. In this case, ae0 also has logical interfaces on unit 1 and unit 2. The DF election process might put ae0 into the down state due to the ESI and other determining factors.

The logical interfaces configured on ae0 would also be down, and unable to provide services to their respective customer sites (VLANs).

If you specify an ESI on a logical interface, the DF election process now occurs at the individual logical interface level. You can better utilize logical interfaces. For example, you configure logical interfaces ae0.1 and ae0.2 on aggregated Ethernet interface ae0. You also enable EVPN multihoming active-standby on both logical interfaces. Then the DF election brings ae0.1 down based on the ESI on ae0.1 and other DF determining factors. Despite logical interface ae0.1 being down, logical interface ae0.2 and other ae0 logical interfaces can remain up. The DF election process doesn't disrupt customer services to the respective customer sites (VLANs).

NOTE: A physical interface on which you configure an ESI won't become a DF if any one of its logical units is down. As a result, if you don't want logical interfaces in one unit to forward traffic, then delete the logical unit rather than disabling it. We recommend that you do not explicitly disable any logical units associated with a physical interface on which you configured an ESI.

Options

identifier Ten octet value. Reserved values: ESI value 0 and all fixed filters. You can't use the reserved values for configuring a multihomed Ethernet segment.

NOTE:

- You can't configure two interfaces (physical, logical, or aggregated Ethernet) with the same ESI value. However, starting with Junos 24.2R1 for MX Series routers supporting ESI per EVPN instance (EVI), users can configure multiple logical interfaces (IFLs) with the same ESI value, but they need to be part of the same bridge-domain. See "per-evi" on page 12409 for more information.
- The left most octet must be configured as 00. The other 9 octets are fully configurable.

all-active Configure the EVPN active-active multihoming mode of operation.

single-active Configure the EVPN active-standby multihoming mode of operation.

NOTE: We don't support active-standby multihoming mode on QFX Series switches or in EVPN configurations with VXLAN overlays. As a result, if

you configure the `single-active` option on QFX Series switches or in EVPN-VXLAN configurations, the device ignores that configuration item.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

`source-bmac` option introduced in Junos OS Release 18.1R1 on MX Series routers with MPCs and MICs.

`auto-derive` stanza introduced in Junos OS Release 18.4R1 on MX Series routers and QFX Series switches.

Support at the following hierarchy levels introduced in Junos OS Release 20.3R1 on QFX Series switches: `[edit protocols evpn interconnect]` and `[edit routing-instances` *`routing-instance-name`* `protocols evpn interconnect]`.

RELATED DOCUMENTATION

evpn | 6541

Example: Configuring an ESI on a Logical Interface With EVPN Multihoming

Understanding Automatically Generated ESIs in EVPN Networks

esi-resolution-per-bridge-domain

IN THIS SECTION

Syntax | 6457

- Hierarchy Level | 6457
- Description | 6457
- Required Privilege Level | 6457
- Release Information | 6457

Syntax

```
esi-resolution-per-bridge-domain;
```

Hierarchy Level

```
[edit protocols evpn]
```

Description

Enables a device that does not support per bridge domain routing to interoperate with devices in the EVPN MPLS network that advertises MPLS labels on a per bridge domain by building next hop routing tables based on a per bridge domain basis.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 21.1.

esmc-transmit (Chassis Synchronization)

IN THIS SECTION

- Syntax | 6458
- ACX710 | 6458
- Hierarchy Level | 6459
- Description | 6459
- Options | 6459
- Required Privilege Level | 6459
- Release Information | 6459

Syntax

```
esmc-transmit {
    interfaces (all | interface-name);
}
```

ACX710

```
esmc-transmit {
    interfaces {
        (all | interface-name) {
            squelch {
                method ql-dnu;
                quality-level quality-level;
            }
        }
    }
}
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Enable Ethernet Synchronization Message Channel (ESMC) packet transmission on all interfaces or on a specific interface.

Options

interfaces (all | *interface-name*) Specify the interface. To enable ESMC packet transmission on all interfaces, use `all`. To enable ESMC packet transmission on one interface, use the name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2R4.

`bits` option introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

- Interfaces Fundamentals
- synchronization (ACX Series) | **28787**
- synchronization (MX Series) | **28795**
- synchronization (PTX Series) | **28802**
- show chassis synchronization | **18879**

establish-tunnels

IN THIS SECTION

- Syntax | 6460
- Hierarchy Level | 6460
- Description | 6460
- Options | 6461
- Required Privilege Level | 6461
- Release Information | 6461

Syntax

```
establish-tunnels (immediately | on-traffic | responder-only);
```

Hierarchy Level

```
[edit services ipsec-vpn]
```

Description

Specify when IKE is activated: immediately after VPN information is configured and configuration changes are committed, or only when data traffic flows. In the second case, IKE needs to be negotiated with the peer gateway. Starting in Junos OS Release 18.2R1, you can also specify that the MX Series router only responds to IKE negotiations.

NOTE: The *immediately* option is required to tear down the st0 interface when dead peer detection (DPD) protocol is configured.

Options

immediately IKE is activated immediately after VPN configuration and configuration changes are committed.

on-traffic IKE is activated only when data traffic flows. IKE needs to be negotiated with the peer gateway.

responder-only Responds to IKE negotiations that are initiated by the peer gateway, but does not initiate IKE negotiations. This option is required when another vendor's peer gateway expects the protocol and port values in the traffic selector from the initiating gateway, which the MX Series does not provide.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

`responder-only` option added in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Understanding Junos VPN Site Secure

ether-options

IN THIS SECTION

- Junos OS Syntax | 6462
- Junos OS Evolved Syntax | 6463
- Hierarchy Level | 6464

- Description | 6464
- Default | 6470
- Options | 6470
- Required Privilege Level | 6471
- Release Information | 6471

Junos OS Syntax

```
ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
            (primary |backup);
            port-priority
        }
    }
    asynchronous-notification;
     (auto-negotiation| no-auto-negotiation);

    configured-flow-control {
        rx-buffers (on | off);
        tx-buffers (on | off);
    }
    ethernet-switch-profile {
        ethernet-policer-profile
        (mac-learn-enable | no-mac-learn-enable);
        recovery-timeout time-in-seconds;
        storm-control storm-control-profile;
        tag-protocol-id;
    }
    (flow-control | no-flow-control);
    ieee-802-3az-eee;
    ignore-l3-incompletes;
    link-mode (automatic | full-duplex | half-duplex);
    (loopback | no-loopback);
    mdi-mode (auto | force | mdi | mdix);
```

```
    mpls {
        pop-all-labels <required-depth (1 | 2 | all)>;
    }
    no-auto-mdix;
    redundant-parent (Interfaces) parent;
    source-address-filter name;
    (source-filtering| no-source-filtering);
    speed {
        (auto-negotiation <auto-negotiate-10-100> | ethernet-100m | ethernet-10g | ethernet-10m
| ethernet-1g);
    }
}
```

Junos OS Evolved Syntax

```
ether-options {
    802.3ad  {
        aex;
        (backup | primary);
        lacp {
            force-up;
            (primary |backup);
            port-priority
        }
    }
    asynchronous-notification;
    (auto-negotiation| no-auto-negotiation);

    ethernet-switch-profile {
        ethernet-policer-profile
        (mac-learn-enable | no-mac-learn-enable);
        recovery-timeout time-in-seconds;
        storm-control storm-control-profile;
        tag-protocol-id;
    }
    fec (gigether)
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    (loopback | no-loopback);
    loopback-remote;
    mpls {
```

```
        pop-all-labels <required-depth (1 | 2 | all)>;
    }
    source-address-filter name;
    (source-filtering | no-source-filtering);
}
```

Hierarchy Level

```
[edit interfaces interface-name]
[edit interfaces interface-range range]
```

Description

Configure `ether-options` properties for a Gigabit Ethernet or 10-Gigabit Ethernet interface.

In Junos OS Evolved, when you configure `set interfaces interface ether-options 802.3ad ae name` at the same time as you apply a second configuration to the same interface at the `[edit interfaces interface]` hierarchy, the second configuration will not take effect until the interface joins the aggregated Ethernet interface *ae name*.

> **NOTE**: The `ether-options` statement is not supported for subscriber management on aggregated Ethernet member link interfaces. You must configure `gigether-options` instead.

Table 39 on page 6465 shows the supported and unsupported platforms.

Table 39: Supported Platform Information

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
ACX Series Routers (Junos OS) • ACX500 • ACX710 • ACX1000 • ACX1100 • ACX2100 • ACX2200 • ACX4000 • ACX5400 • ACX5448 • ACX5048 • ACX5096	ACX Series Routers (Junos OS) • ACX5048 • ACX5096	To configure storm control on the ACX Series routers, use `ether-options`.

Table 39: Supported Platform Information *(Continued)*

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
EX Series Switches (Junos OS) • EX4300 Multigigabit • EX4600 • EX4650 • EX9200 • EX9250	EX Series Switches (Junos OS) • EX2300 • EX2300 Multigigabit • EX2300-C • EX3400 • EX4300 • EX4300 Multigigabit • EX4600 • EX4650 • EX9200 • EX9250	EX Series Switches (Junos OS) EX Series switches support both `ether-options` and `gigether-options`. To configure Ethernet configuration options such as loopback, flow-control, auto-negotiation etc., use `ether-options`. To configure Forward Error Correction (FEC), use `gigether-options`. **NOTE**: For all EX Series Switches except EX9200 series, use `ether-options` for all ethernet configuration options except FEC option. You must use `gigether-options` for configuring FEC.

Table 39: Supported Platform Information *(Continued)*

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
MX Series Routers (Junos OS) • MX5 • MX10 • MX40 • MX80 • MX104 • MX150 • MX204 • MX240 • MX480 • MX960 • MX20008 • MX2010 • MX2020 • MX10003 • MX10008 and MX10016	Not Supported	None

Table 39: Supported Platform Information *(Continued)*

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
PTX Series Routers (Junos OS) • PTX1000 • PTX3000 • PTX5000 • PTX10001 • PTX10002 • PTX10008 and PTX10016	PTX Series Routers (Junos OS) • PTX1000 • PTX10001 • PTX10002 • PTX10008 and PTX10016	PTX Series Routers (Junos OS) In Junos OS Release 17.3R3S7, PTX1000 Series routers support both `ether-options` and `gigether-options`. In Junos OS Releases 17.3R1, 17.4R1, and 17.4R2, PTX10000 Series routers support both `ether-options` and `gigether-options`. In Junos OS Release 18.1R1 and later, PTX Series routers support only `gigether-options` for your configuration and do not support `ether-options`. Before Junos OS Release 18.1R1, PTX Series support `ether-options`. See TSB17864 for additional information.
PTX Series routers (Junos OS Evolved) • PTX10003 • PTX10008 and PTX10016	PTX Series routers (Junos OS Evolved) • PTX10003 • PTX10008 and PTX10016	PTX Series routers (Junos OS Evolved) Starting in Junos OS Evolved Release 20.1R1, PTX Series routers support `ether-options` only. Before Junos OS Evolved Release 20.1R1, PTX Series routers supported `gigether-options` only.

Table 39: Supported Platform Information *(Continued)*

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
QFX Series Switches (Junos OS) • QFX5100 (48S) • QFX5100 (48T) • QFX5100 (24Q) • QFX5100 (96S) • QFX5110 (48S) • QFX5110 (32Q) • QFX5120 (48Y) • QFX5120 (48T) • QFX5120 (48ym) • QFX5120 (32C) • QFX5200 (48Y) • QFX5200 (32C) • QFX5210 (64C) • QFX5210 (64C-S) • QFX5220 (32CD) • QFX5220 (128C) • QFX10002 • QFX10008 and QFX10016	QFX Series Switches (Junos OS) • QFX5100 (48S) • QFX5100 (48T) • QFX5100 (24Q) • QFX5100 (96S) • QFX5110 (48S) • QFX5110 (32Q) • QFX5120 (48Y) • QFX5120 (48T) • QFX5120 (48ym) • QFX5120 (32C) • QFX5200 (48Y) • QFX5200 (32C) • QFX5210 (64C) • QFX5210 (64C-S) • QFX5220 (32CD) • QFX5220 (128C) • QFX10002 • QFX10008 and QFX10016	QFX Series Switches (Junos OS) QFX Series switches support both `ether-options` and `gigether-options`. To configure Ethernet configuration options such as loopback, flow-control, auto-negotiation etc., use `ether-options`. To configure Forward Error Correction (FEC), use `gigether-options`.

Table 39: Supported Platform Information *(Continued)*

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
Not Supported	QFX Series Switches (Junos OS Evolved) • QFX5200-32C-L • QFX5220-32CD • QFX5220-128C	To configure auto-negotiation on management interfaces (`re0:mgmt-0`), use `ether-options`.
SRX Series (Junos OS) • SRX300 • SRX550 • SRX1500 • SRX4100 and SRX4200 • SRX4600 • SRX5400 • SRX5600	SRX Series (Junos OS) • SRX300 • SRX550 • SRX1500 • SRX4100 and SRX4200 • SRX4600 • SRX5400 • SRX5600	SRX Series (Junos OS) To configure gigabit-Ethernet interfaces (ge-), use `gigether-options`. To configure ethernet interfaces (et-) and fast ethernet interfaces (fe-), use `ether-options`.

Default

Enabled.

Options

NOTE: The `auto-negotation` and `speed` statements are not supported on the OCX Series.

loopback-remote Starting in Junos OS Evolved Release 20.1R1, enable remote loopback.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`autostate-exclude` option introduced in Junos OS Release 14.1x53-D40 for QFX5100 switches only.

`fec` and `loopback-remote` options introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

Gigabit Ethernet Interface

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support

gigether-options | **7685**

ether-pseudowire

IN THIS SECTION

- Syntax | **6472**
- Hierarchy Level | **6472**
- Description | **6472**
- Required Privilege Level | **6472**
- Release Information | **6472**

Syntax

```
ether-pseudowire {
    zero-control-word;
}
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key family mpls]
```

Description

Load-balance IP over Ethernet pseudowire. Presence of zero-control-word in the payload indicates an Ethernet frame.

zero-control-word Precedes Ethernet packet to indicate the start Ethernet farme in an MPLS ether-pseudowire payload.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1 for the MX Series.

RELATED DOCUMENTATION

enhanced-hash-key | 6362

hash-key (Forwarding Options) | 7986

family mpls | 6886

MPLS Encapsulated Payload Load-balancing Overview

ether-type (Dynamic Load Balancing)

IN THIS SECTION

- Syntax | 6473
- Hierarchy Level | 6473
- Description | 6473
- Options | 6473
- Required Privilege Level | 6474
- Release Information | 6474

Syntax

```
ether-type (ipv4|ipv6|mpls|none);
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key ecmp-dlb]
```

Description

EtherType for dynamic load balancing.

NOTE: On QFX5130 and QFX5700 devices, `ether-type` is not supported on Junos OS Evolved.

Options

ether-type *(ipv4 | ipv6| mpls)*

Ether-type for DLB:

- ipv4—Set dynamic load balancing for IPv4.

- ipv6—Set dynamic load balancing for IPv6.
- mpls—Set dynamic load balancing for MPLS.
- none-Disable dynamic load balancing for all ether types.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

dlb | 5841

Load Balancing for Aggregated Ethernet Interfaces

Load Balancing for Aggregated Ethernet Interfaces

Load Balancing for Aggregated Ethernet Interfaces

ether-type (VLAN)

IN THIS SECTION

- Syntax | 6475
- Hierarchy Level | 6475
- Description | 6475
- Required Privilege Level | 6475
- Release Information | 6475

Syntax

```
ether-type (0x8100 | 0x88a8 | 0x9100)
```

Hierarchy Level

```
[edit ethernet-switching-options dot1q-tunneling]
```

Description

Configure a global value for the tag protocol identifier (EtherType) of the service VLAN tags (outer tags) in Q-in-Q tunneling. Only one Ethertype value is supported at a time. The Ethertype value appears in the Ethernet type field of the packet. It specifies the protocol being transported in the Ethernet frame.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

dot1q-tunneling | **5915**

Example: Setting Up Q-in-Q Tunneling on EX Series Switches

Configuring Q-in-Q Tunneling on EX Series Switches

Configuring Q-in-Q Tunneling on QFX Series Switches

Example: Setting Up Q-in-Q Tunneling on QFX Series Switches

ether-type

IN THIS SECTION

- Syntax | 6476
- Hierarchy Level | 6476
- Description | 6476
- Options | 6476
- Required Privilege Level | 6477
- Release Information | 6477

Syntax

```
ether-type ether-type;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Two-octet field in an Ethernet frame that defines the protocol encapsulated in the frame payload. See http://standards.ieee.org/develop/regauth/ethertype/eth.txt for a list of Institute of Electrical and Electronics Engineers (IEEE) EtherTypes.

NOTE: To create a FIP application, use the EtherType 0x8914.

Options

type—Identifier for the EtherType.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Defining an Application for DCBX Application Protocol TLV Exchange

Example: Configuring DCBX Application Protocol TLV Exchange

Understanding DCBX Application Protocol TLV Exchange

ether-type (Chassis FPC)

IN THIS SECTION

- Syntax | 6477
- Hierarchy Level | 6478
- Description | 6478
- Options | 6479
- Required Privilege Level | 6479
- Release Information | 6479

Syntax

```
ether-type index etype
```

Hierarchy Level

```
[edit chassis fpc slot]
```

Description

Add or remove a custom EtherType value in the ternary content addressable memory (TCAM) table of EtherType values for a specified FPC slot number on the device.

Ethernet frame headers contain an EtherType field (two octets in length) to identify the protocol in the frame's payload so that a receiving device knows how to process the Layer 2 traffic. The device maintains a list of the default EtherType values in TCAM on each FPC on the device for fast access. Use this statement to configure a custom EtherType value in TCAM on a particular FPC in place of the default value at the specified index in the TCAM table.

For example, to replace the default value in TCAM on FCP 0 at index 24 with EtherType value 0x8890:

```
set chassis fpc 0 ether-type 24 0x8890
```

NOTE: This command doesn't check for or prevent you from adding duplicate values.

The device retains the default values stored at each index. If you set a custom EtherType value and then delete that configuration item, the device restores the default value for that index.

To remove a custom setting, delete the configuration item using the index where you configured the custom EtherType. You don't need to specify the EtherType value you set at that index. For example:

```
delete chassis fpc 0 ether-type 24
```

The default EtherType values in TCAM might differ across Junos OS releases and device types. Enter the `show chassis fpc ether-types` command to see the default and active EtherType values on one or all FPCs on the device. If you configured a custom value at a particular index, the ACTIVE column shows that value instead of the default value.

Some EtherType values are reserved, and you can't change or configure those values. For those values, the `show chassis fpc ether-types` command displays NO in the CONFIGURABLE output column for the corresponding index. Reserved values are stored at TCAM indexes 0-5 and 28-31, so you can only configure custom values at indexes 6 through 27.

Options

index (Required) EtherTypes TCAM table index where you want to replace the default EtherType value with a new EtherType value. The TCAM table can hold up to 32 EtherType entries, but ten locations (0-5 and 28-31) are reserved for default EtherTypes that can't be reconfigured.

- Range: 6-27

etype (Required) EtherType value in hexadecimal format (0x*nnnn*) that you want to add at *index* in the EtherTypes TCAM table.

Reserved EtherTypes that you can't change in TCAM include:

- 0x0800—IPv4
- 0x86DD—IPv6
- 0x0806—ARP
- 0x88CC—LLDP
- 0x8809—LACP
- 0x8035—Reverse Address Resolution Protocol (RARP)
- 0x8100—VLAN-tagged frames for customer VLANs
- 0x88a8—Service VLAN tagged frames for provider bridging
- 0x9100—VLAN-tagged frames with double tagging for Q-in-Q tunneling
- 0x9200—VLAN-tagged frames with triple tagging

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

show chassis fpc ether-types | 18392

ethernet

IN THIS SECTION

- Syntax | 6480
- Hierarchy Level | 6480
- Description | 6480
- Required Privilege Level | 6481
- Release Information | 6481

Syntax

```
ethernet {
    device-count number;
}
```

Hierarchy Level

```
[edit chassis aggregated-devices],
[edit chassis node-group aggregated-devices]
```

Description

Configure properties for aggregated Ethernet devices on the switch.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion

Configuring Link Aggregation

Junos OS Network Interfaces Library for Routing Devices

ethernet (Aggregated Devices)

IN THIS SECTION

- Syntax | 6481
- Hierarchy Level | 6482
- Description | 6482
- Required Privilege Level | 6482
- Release Information | 6482

Syntax

```
ethernet {
device-count                    number;
lacp {
    link-protection {
        non-revertive;
    }
```

```
    system-priority;
}
```

Hierarchy Level

```
[edit chassis aggregated-devices]
```

Description

Configure properties for Ethernet aggregated devices on the switch.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Configuring Aggregated Ethernet Links (CLI Procedure)
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*
- Junos OS Ethernet Interfaces Configuration Guide

ethernet (Alarm)

IN THIS SECTION

- Syntax | 6483
- Hierarchy Level | 6483

- Description | 6483
- Options | 6483
- Required Privilege Level | 6483
- Release Information | 6483

Syntax

```
ethernet {
    link-down  (red | yellow | ignore);
}
```

Hierarchy Level

```
[edit chassis alarm],
[edit chassis interconnect-device  name alarm],
[edit chassis node-group name alarm]
```

Description

Configure alarms for an Ethernet interface.

Options

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

ethernet (Chassis)

IN THIS SECTION

- Syntax | 6484
- Hierarchy Level | 6484
- Description | 6484
- Options | 6484
- Required Privilege Level | 6485
- Release Information | 6485

Syntax

```
ethernet {
    device-count number;
    lacp {
        link-protection {
            non-revertive;
        }
        system-priority number;
    }
}
```

Hierarchy Level

```
[edit chassis aggregated-devices]
```

Description

Configure properties for aggregated Ethernet devices.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

ethernet (CoS for Multidestination Traffic)

IN THIS SECTION

- Syntax | 6485
- Hierarchy Level | 6485
- Description | 6485
- Required Privilege Level | 6486
- Release Information | 6486

Syntax

```
ethernet {
    broadcast forwarding-class-name;
}
```

Hierarchy Level

```
[edit class-of-service multi-destination family]
```

Description

Specify the Ethernet broadcast traffic family.

> **NOTE**: On EX4300 switches, including switches in a virtual chassis or virtual chassis fabric configuration, you cannot explicitly configure multi-destination (broadcast, multicast, and destination lookup fail) traffic classification. By default, all multi-destination traffic uses output queue 8. This means all multicast (multi-destination) traffic always uses Q8 and there is no control to modify the queue.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

- *Understanding CoS Schedulers*
- *Understanding CoS Forwarding Classes*
- *Understanding CoS Classifiers*

ethernet-frame-size

IN THIS SECTION

- Syntax | 6487
- Hierarchy Level | 6487
- Description | 6487
- Options | 6487
- Required Privilege Level | 6487
- Release Information | 6487

Syntax

```
ethernet-frame-size ethernet-frame-size;
```

Hierarchy Level

```
[edit protocols oam fabric flow-specs flow-specification-name]
```

Description

Configure the Ethernet frame size in the flow specification parameters for internal fabric monitoring.

Options

ethernet-frame-size Integer defining the size (in bytes) of the Ethernet frame.

- **Range:** 256 to 9116 bytes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | 29924

ethernet-policer-profile

IN THIS SECTION

- Syntax | 6488
- Hierarchy Level | 6489
- Description | 6489
- Required Privilege Level | 6489
- Release Information | 6489

Syntax

```
ethernet-policer-profile {
    input-priority-map {
        ieee802.1p premium [ values ];
    }
    output-priority-map {
        classifier {
            premium {
                forwarding-class class-name {
                    loss-priority (high | low);
                }
            }
        }
    }
    policer cos-policer-name {
        aggregate {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        premium {
            bandwidth-limit bps;
            burst-size-limit bytes;
```

```
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile],
[edit interfaces interface-name aggregated-ether-options ethernet-switch-profile]
```

Description

NOTE: On QFX Series standalone switches, this statement hierarchy is only supported on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet, Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, configure a class of service (CoS)-based policer. Policing applies to the inner VLAN identifiers, not to the outer tag. For Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), the `premium` policer is not supported.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

ethernet (Protocols OAM)

IN THIS SECTION

- Syntax: MX, T, ACX Series Routers, SRX Firewalls, M320 and EX Series Switches | 6490
- Syntax: EX Series Switches and NFX Series Devices | 6493
- Hierarchy Level | 6496
- Description | 6496
- Required Privilege Level | 6496
- Release Information | 6496

Syntax: MX, T, ACX Series Routers, SRX Firewalls, M320 and EX Series Switches

```
ethernet {
    connectivity-fault-management {
        action-profile profile-name {
            default-actions {
                interface-down;
            }
        }
        performance-monitoring {
            delegate-server-processing;
            hardware-assisted-timestamping;
            hardware-assisted-keepalives;
            sla-iterator-profiles {
                profile-name {
                    avg-fd-twoway-threshold;
                    avg-ifdv-twoway-threshold;
                    avg-flr-forward-threshold;
                    avg-flr-backward-threshold;
                    disable;
                    calculation-weight {
                        delay delay-weight;
                        delay-variation delay-variation-weight;
                    }
                    cycle-time milliseconds;
                    iteration-period connections;
```

```
                    measurement-type (loss | statistical-frame-loss | two-way-delay);
                }
            }
        }
        linktrace {
            age (30m | 10m | 1m | 30s | 10s);
            path-database-size path-database-size;
        }
        maintenance-domain domain-name {
            level number;
            name-format (character-string | none | dns | mac+2octet);
            maintenance-association ma-name {
                short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
                protect-maintenance-association protect-ma-name;
                remote-maintenance-association remote-ma-name;
                continuity-check {
                    convey-loss-threshold;
                    hold-interval minutes;
                    interface-status-tlv;
                    interval (10m | 10s | 1m | 1s| 100ms);
                    loss-threshold number;
                    port-status-tlv;
                }
                mep mep-id {
                    auto-discovery;
                    direction (up | down);
                    interface interface-name (protect | working);
                    lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect
| rem-err-xcon | xcon );
                    priority number;
                    remote-mep mep-id {
                        action-profile profile-name;
                        sla-iterator-profile profile-name {
                            data-tlv-size size;
                            iteration-count count-value;
                            priority priority-value;
                        }
                    }
                }
            }
        }
    }
    evcs evc-id {
```

```
        evc-protocol cfm management-domain domain-id (management-association association-id |
vpls (routing-instance instance-id);
        remote-uni-count count;
        multipoint-to-multipoint;
    }
    link-fault-management {
        action-profile profile-name {
            action {
                link-down;
                send-critical-event;
                syslog;
            }
            event {
                link-adjacency-loss;
                link-event-rate {
                    frame-error count;
                    frame-period count;
                    frame-period-summary count;
                    symbol-period count;
                }
                protocol-down;
            }
        }
        interface interface-name {
            apply-action-profile;
            link-discovery (active | passive);
            loopback-tracking;
            pdu-interval interval;
            pdu-threshold threshold-value;
            remote-loopback;
            event-thresholds {
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
            }
            negotiation-options {
                allow-remote-loopback;
                no-allow-link-events;
            }
        }
    }
    lmi {
```

```
        status-counter count;
        polling-verification-timer value;
        interface name {
            uni-id uni-name;
            status-counter number;
            polling-verification-timer value;
            evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
            evc evc-name {
                default-evc;
                vlan-list vlan-id-list;
            }
        }
    }
}
```

Syntax: EX Series Switches and NFX Series Devices

```
ethernet {
    connectivity-fault-management {
        action-profile profile-name {
            action {
                interface-down;
            }
            default-actions {
                interface-down;
            }
            event {
                adjacency-loss;
            }
        }
        esp-traceoptions  {
            file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-
world-readable>;
            flag (all |error | esp | interface | krt | lib |normal |task |timer);
        }
        linktrace {
            age (30m | 10m | 1m | 30s | 10s);
            path-database-size path-database-size;
        }
        maintenance-association domain-name {
            level number;
```

```
            mip-half-function (none | default |explicit);
            name-format (character-string | none | dns | mac+2oct);
            maintenance-association ma-name {
                continuity-check {
                    hold-interval minutes;
                    interface-status-tlv;
                    interval (10m | 10s | 1m | 1s| 100ms);
                    loss-threshold number;
                    port-status-tlv;
                }
                mep mep-id {
                    auto-discovery;
                    direction down;
                    interface interface-name;
                    priority
                    remote-mep mep-id {
                        action-profile profile-name;
                        sla-iterator-profile profile-name {
                            data-tlv-size size;
                            iteration-count count-value;
                            priority priority-value;
                        }
                    }
                }
                short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
            }
        }
        performance-monitoring {
            sla-iterator-profiles {
                profile-name {
                    calculation-weight {
                        delay delay-value;
                        delay-variation delay-variation-value;
                    }
                    cycle-time cycle-time-value;
                    iteration-period iteration-period-value;
                    measurement-type two-way-delay;
                    passive;
                }
            }
        }
        traceoptions {
            file filename <files number> <match regex> <size size> <world-readable | no-world-
```

```
readable>;
            flag flag ;
            no-remote-trace;
        }
    }
    link-fault-management {
        action-profile profile-name;
            action {
                syslog;
                link-down;
            }
            event {
                link-adjacency-loss;
                link-event-rate {
                    frame-error  count;
                    frame-period count;
                    frame-period-summary  count;
                    symbol-period count;
                }
            }
        interface interface-name {
            link-discovery (active | passive);
            pdu-interval interval;
            pdu-threshold threshold-value;
            remote-loopback;
            event-thresholds  {
                frame-error count;
                 frame-period count;
                 frame-period-summary  count;
                 symbol-period count;
            }
            negotiation-options {
                allow-remote-loopback;
                 no-allow-link-events;
            }
        }
        traceoptions {
            file filename <files number> <match regex> <size size> <world-readable | no-world-
readable>;
            flag flag ;
            no-remote-trace;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit protocols oam]
```

Description

Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) support for Ethernet interfaces or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

`connectivity-fault-management` introduced in Junos OS Release 10.2 for EX Series switches.

RELATED DOCUMENTATION

> Enabling IEEE 802.3ah OAM Support
>
> *Example: Configuring Ethernet OAM Link Fault Management*

ethernet-ring

IN THIS SECTION

- Syntax | 6497
- Hierarchy Level | 6498
- Description | 6498
- Options | 6498
- Required Privilege Level | 6498
- Release Information | 6498

Syntax

```
ethernet-ring ring-name {
    control-vlan (vlan-id | vlan-name);
    data-channel {
        vlan number
    }
    east-interface {
        control-channel channel-name {
            vlan number;
            interface name interface-name
        }
    }
    guard-interval number;
    node-id mac-address;
    restore-interval number;
    ring-protection-link-owner;
    west-interface {
        control-channel channel-name {
            vlan number;
        }
    }
}
```

Hierarchy Level

```
[edit protocols protection-group]
```

Description

For Ethernet PICs on MX Series routers or for EX Series switches, , specify the Ethernet ring in an Ethernet ring protection switching configuration.

Options

ring-name—Name of the Ethernet protection ring.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

RELATED DOCUMENTATION

- *Ethernet Ring Protection Switching Overview*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

ethernet-switch-profile

IN THIS SECTION

- Syntax | 6499
- Hierarchy Level | 6500
- Description | 6500
- Default | 6500
- Required Privilege Level | 6501
- Release Information | 6501

Syntax

```
ethernet-switch-profile {
    ethernet-policer-profile {
        input-priority-map {
            ieee802.1p premium [values];
        }
        output-priority-map {
            classifier {
                premium {
                    forwarding-class class-name {
                        loss-priority (high | low);
                    }
                }
            }
        }
        policer cos-policer-name {
            aggregate {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
            premium {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
        }
```

```
        storm-control storm-control-profile;
        tag-protocol-id tpid;
    }
    mac-learn-enable;
}
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options],
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name ether-options]
```

Description

NOTE: On QFX Series standalone switches, the `ethernet-policer-profile` CLI hierarchy and the `mac-learn-enable` statement are supported only on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, aggregated Ethernet with Gigabit Ethernet IQ interfaces, the built-in Gigabit Ethernet port on the M7i router); 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches, configure VLAN tag and MAC address accounting and filtering properties.

The remaining statements are explained separately. See CLI Explorer.

NOTE: When you gather interfaces into a bridge domain, the `no-mac-learn-enable` statement at the `[edit interfaces interface-name gigether-options ethernet-switch-profile]` hierarchy level is not supported. You must use the `no-mac-learning` statement at the `[edit bridge-domains bridge-domain-name bridge-options interface interface-name]` hierarchy level to disable MAC learning on an interface in a bridge domain. For information on disabling MAC learning for a bridge domain, see the *MX Series Layer 2 Configuration Guide*.

Default

If the `ethernet-switch-profile` statement is not configured, Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) behave like Gigabit Ethernet interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

- *Configuring Gigabit Ethernet Policers*
- *Configuring Gigabit Ethernet Policers*
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview
- Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support

ethernet-switching

IN THIS SECTION

- Syntax | 6501
- Hierarchy Level | 6502
- Description | 6502
- Default | 6502
- Required Privilege Level | 6502
- Release Information | 6502

Syntax

```
ethernet-switching {
    filter {
```

```
        group filter-group-number;
        input filter-name;
        input-list [ filter-names ];
        output filter-name;
        output-list [ filter-names ];
    }
    interface-mode (access | trunk);
    recovery-timeout seconds;
    storm-control profile-name;
    vlan {
        members (vlan-name |[-vlan-names] | all);
    }
}
```

Hierarchy Level

```
[edit interfaces ge-chassis/slot/port unit logical-unit-number] family
```

Description

Configure Ethernet switching protocol family information for the logical interface.

The remaining statements are explained separately. See CLI Explorer.

Default

You must configure a logical interface to be able to use the physical device.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

JUNOS Software Network Interfaces Configuration Guide

ethernet-switching

IN THIS SECTION

- Syntax (EX Series and QFX Series) | 6503
- Syntax (SRX Series) | 6504
- Hierarchy Level | 6504
- Description | 6504
- Default | 6504
- Options | 6504
- Required Privilege Level | 6505
- Release Information | 6505

Syntax (EX Series and QFX Series)

```
ethernet-switching {
     filter input filter-name;
     filter output filter-name;
     native-vlan-id vlan-id;
     port-mode mode;
    reflective-relay;
     vlan {
         members [ (all | names | vlan-ids) ];
    }
}
```

Syntax (SRX Series)

```
ethernet-switching {
    block-non-ip-all;
    bpdu-vlan-flooding;
    bypass-non-ip-unicast;
    no-packet-flooding {
        no-trace-route;
    }
}
```

Hierarchy Level

For EX Series and QFX Series switches:

```
[edit interfaces ge-chassis/slot/port unit logical-unit-number] family
```

For SRX Series Firewalls:

```
[edit security flow]
```

Description

Configure Ethernet switching protocol family information for the logical interface. Changes default Layer 2 forwarding behavior.

The remaining statements are explained separately. See CLI Explorer.

Default

You must configure a logical interface to be able to use the physical device.

Options

- `block-non-ip-all`—Block all Layer 2 non-IP and non-ARP traffic, including multicast and broadcast traffic.

- `bpdu-vlan-flooding`—Set 802.1D bridge protocol data unit (BPDU) flooding based on VLAN on which BPDU originate. The default behavior is to receive BPDUs and then flood BPDUs out to all active ports on the SRX Series Firewalls.
- `bypass-non-ip-unicast`—Allow all Layer 2 non-IP traffic to pass through the device.
- `no-packet-flooding`—Stop IP flooding and send ARP or ICMP requests to discover the destination MAC address for a unicast packet.

  NOTE: On all SRX Series Firewalls in transparent mode, packet flooding is enabled by default. If you have manually disabled packet flooding with the set security flow ethernet-switching no-packet-flooding command, then multicast packets such as OSPFv3 hello packets are dropped.

  - `no-trace-route`—Do not send ICMP requests to discover the destination MAC address for a unicast packet. Only ARP requests are sent. This option only allows the device to discover the destination MAC address for a unicast packet if the destination IP address is in the same subnetwork as the ingress IP address.

NOTE: The `block-non-ip-all` and `bypass-non-ip-unicast` options cannot be configured at the same time.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

Understanding Traffic Processing on Security Devices

JUNOS Software Network Interfaces Configuration Guide

ethernet-switching-options

IN THIS SECTION

- EX Series | 6506
- QFX Series, QFabric, EX4600 | 6510
- Hierarchy Level | 6512
- Description | 6512
- Required Privilege Level | 6512
- Release Information | 6512

EX Series

```
ethernet-switching-options {
    analyzer  {
        name {
            loss-priority priority;
            ratio number;
            input {
                ingress {
                    interface (all | interface-name);
                    vlan (vlan-id | vlan-name);
                }
                egress {
                    interface (all | interface-name);
                }
            }
            output {
                interface interface-name;
                vlan (vlan-id | vlan-name) {
                no-tag;
                }
            }
        }
    }
     bpdu-block {
         disable-timeout timeout;
```

```
        interface (all | [interface-name]) {
            (disable | drop | shutdown);
            }
    }
    dot1q-tunneling {
        ether-type (0x8100 | 0x88a8 | 0x9100);
    }
     interfaces interface-name {
        no-mac-learning;
        no-mac-notification;
    }
    mac-lookup-length number-of-entries;
}
    mac-notification {
        notification-interval seconds;
    }
    mac-table-aging-time seconds;
    nonstop-bridging;
    port-error-disable {
        disable-timeout timeout;
    }
    redundant-trunk-group {
        group name {
            interface interface-name <primary>;
            interface interface-name;
        }
    }
    secure-access-port {
        dhcp-snooping-file {
            location local_pathname | remote_URL;
            timeout seconds;
            write-interval seconds;
        }
        dhcpv6-snooping-file {
            location local_pathname | remote_URL;
            timeout seconds;
            write-interval seconds;
        }
        interface (all | interface-name) {
            allowed-mac {
                mac-address-list;
            }
            (dhcp-trusted | no-dhcp-trusted);
```

```
            fcoe-trusted;
            mac-limit limit action (drop | log | none | shutdown);
            no-allowed-mac-log;
            persistent-learning;
            static-ip ip-address {
                vlan  vlan-name;
                mac mac-address;
            }
            static-ipv6 ip-address {
                vlan  vlan-name;
                mac mac-address;
            }
        }
         vlan (all | vlan-name) {
            (arp-inspection | no-arp-inspection) [
                forwarding-class class-name;
            }
             dhcp-option82 {
                 circuit-id {
                     prefix hostname;
                     use-interface-description;
                     use-vlan-id;
                }
                 remote-id {
                     prefix hostname | mac | none;
                     use-interface-description;
                     use-string string;
                }
                 vendor-id [string];
            }
            (examine-dhcp | no-examine-dhcp) {
                forwarding-class class-name;
            }
            (examine-dhcpv6 | no-examine-dhcpv6) {
                forwarding-class class-name;
            }
            examine-fip {
            fc-map fc-map-value;
            }
            (ip-source-guard | no-ip-source-guard);
            (ipv6-source-guard | no-ipv6-source-guard);
             mac-move-limit limit action (drop | log | none | shutdown);
            }
```

```
            (neighbor-discovery-inspection | no-neighbor-discovery-inspection);
        no-option-37;
        static {
            vlan name {
                mac  mac-address                 {
                    next-hop interface-name;
                }
            }
        }
        storm-control {
            action-shutdown;
            interface (all | interface-name) {
                bandwidth bandwidth;
                level level;
                multicast;
                no-broadcast;
                no-multicast;
                no-registered-multicast;
                no-unknown-unicast;
                no-unregistered-multicast;
            }
        }
        traceoptions {
            file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-
world-readable>;
            flag flag <disable>;
        }
        unknown-unicast-forwarding {
            vlan (all | vlan-name) {
                interface interface-name;
            }
        }
        voip {
            interface (all | [interface-name | access-ports]) {
                forwarding-class forwarding-class;
                vlan vlan-name;
            }
        }
    }
```

QFX Series, QFabric, EX4600

```
ethernet-switching-options {
    analyzer  {
        name {
            input {
                egress {
                    interface (all | interface-name);
                }
                ingress {
                    interface (all | interface-name);
                    vlan (vlan-id | vlan-name);
                }
            output {
                interface interface-name;
                 ip-address ip-address;
                vlan (vlan-id | vlan-name);
            }
        }
    }
    bpdu-block {
        interface (all | [interface-name]);
        disable-timeout timeout;
    }
    dot1q-tunneling {
        ether-type (0x8100 | 0x88a8 | 0x9100)
    }
    interfaces interface-name {
        no-mac-learning;
    }
     mac-table-aging-time seconds {
    }
    port-error-disable {
        disable-timeout timeout;
    }
    secure-access-port {
        dhcp-snooping-file {
            location local_pathname | remote_URL;
            timeout seconds;
            write-interval seconds;
        }
        interface (all | interface-name) {
```

```
            allowed-mac {
                mac-address-list;
            }
            (dhcp-trusted | no-dhcp-trusted);
            fcoe-trusted;
            mac-limit limit action action;
            no-allowed-mac-log;
        }
        vlan (all | vlan-name) {
            (arp-inspection | no-arp-inspection) [
                forwarding-class (for DHCP Snooping or DAI Packets) class-name;
            }
            dhcp-option82 {
                circuit-id {
                    prefix (Circuit ID for Option 82) hostname;
                    use-interface-description;
                    use-vlan-id;
                }
                 remote-id {
                    prefix (Remote ID for Option 82) hostname | mac | none;
                    use-interface-description;
                    use-string                                  string;
                }
                vendor-id <string>;
            }
            (examine-dhcp | no-examine-dhcp) {
                forwarding-class (for DHCP Snooping or DAI Packets) class-name;
            }
            examine-fip {
                examine-vn2vn {
                    beacon-period milliseconds;
                }
                fc-map fc-map-value;
            }
            mac-move-limit limit <fabric-limit limit action action;
        }
    }
    static {
        vlan vlan-id {
            mac mac-address next-hop interface-name;
        }
    }
    storm-control {
```

```
        interface (all | interface-name) {
            bandwidth bandwidth;
            no-broadcast;
            no-multicast;
            no-unknown-unicast;
        }
    }
    traceoptions {
        file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
        flag flag <disable>;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure Ethernet switching options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Understanding Port Mirroring and Analyzers
- Understanding How to Protect Access Ports from Common Attacks
- Port Security Features

Understanding BPDU Protection for STP, RSTP, and MSTP
Understanding Redundant Trunk Links (Legacy RTG Configuration)
Understanding Storm Control
Understanding 802.1X and VoIP on EX Series Switches
Understanding Q-in-Q Tunneling and VLAN Translation
Understanding MAC Notification on EX Series Switches
Understanding FIP Snooping
Understanding Nonstop Bridging on EX Series Switches

ethernet-switching (Port Mirroring)

IN THIS SECTION

- Syntax | 6513
- Hierarchy Level | 6514
- Description | 6514
- Required Privilege Level | 6514
- Release Information | 6514

Syntax

```
ethernet-switching;
output {
    interface interface-name {
    }
    no-filter-check;
    }
    vlan vlan-name {
        no-tag;
    }
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring [instance name] family]
```

Description

Specify that the output interface for the port mirror will be configured as an `ethernet-switching` interface.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding Port Mirroring and Analyzers

Configuring Port Mirroring

Examples: Configuring Port Mirroring for Local Analysis

ethernet-vlan (Protocols Link Management)

IN THIS SECTION

- Syntax | 6515
- Hierarchy Level | 6515
- Description | 6515
- Options | 6515
- Required Privilege Level | 6515
- Release Information | 6515

Syntax

```
ethernet-label {
    vlan-id-range vlan-id-range;
}
```

Hierarchy Level

```
[edit protocols link-management te-link te-link-name]
```

Description

Specify the TE-link to be used for Layer 2 VLAN label-switched path (LSP).

Options

vlan-id-range *vlan-id-range*	Pool of VALN IDs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

eui-64

IN THIS SECTION

Syntax | 6516

- Hierarchy Level | 6516
- Description | 6516
- Required Privilege Level | 6516
- Release Information | 6516

Syntax

```
eui-64;
```

Hierarchy Level

```
[edit interfaces interface-name unit number family inet6 address address]
```

Description

For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Interface Address

evc-protocol cfm

IN THIS SECTION

- Syntax | 6517
- Hierarchy Level | 6517
- Description | 6517
- Options | 6517
- Required Privilege Level | 6518
- Release Information | 6518

Syntax

```
evc-protocol cfm {
    maintenance-association association-name | vpls routing-instance routing-id;
    maintenance-domain domain-id;
    mep-id mep-id;
}
```

Hierarchy Level

```
[edit protocols oam ethernet evcs]
```

Description

Specify connectivity fault management (CFM) or virtual private LAN service (VPLS) as the Ethernet virtual connection (EVC) protocol.

Options

management-domain *domain-id*	(Optional) For CFM, specify the CFM management domain.
management-association *association-id*	(Optional) For CFM, specify the CFM management association.

`routing-instance` *`instance-id`* (Optional) For VPLS, specify the VPLS routing instance.

`mep-id` *`mep-id`* (Required for CFM) Identifier for the maintenance association endpoint

> **NOTE**: This option is available on MX Series routers only.

- **Range:** 1 through 8191

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

`mep-id` *`mep-id`* statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Ethernet Local Management Interface

lmi (Ethernet OAM) | **9861**

evcs

IN THIS SECTION

- Syntax | **6519**
- Hierarchy Level | **6519**
- Description | **6519**
- Options | **6519**
- Required Privilege Level | **6519**

Release Information | 6520

Syntax

```
evcs evc-id {
    evc-protocol cfm;
    remote-uni-count count;
    multipoint-to-multipoint;
}
```

Hierarchy Level

```
[edit protocols oam ethernet]
```

Description

On MX Series routers with `ge`, `xe`, or `ae` interfaces, configure an OAM Ethernet virtual connection.

Options

`remote-uni-count` *count*—(Optional) Specify the number of remote UNIs in the EVC configuration, the default is 1.

`multipoint-to-multipoint`—(Optional) Specify multiple points in the EVC configuration, the default is point-to-point if `remote-uni-count` is 1.

Remaining options are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Ethernet Local Management Interface

lmi (Ethernet OAM) | 9861

event

IN THIS SECTION

- Syntax | 6520
- Hierarchy Level | 6520
- Description | 6521
- Options | 6521
- Required Privilege Level | 6521
- Release Information | 6521

Syntax

```
event index {
    community community-name;
    description description;
    type type;
}
```

Hierarchy Level

```
[edit snmp rmon]
```

Description

Configure RMON event entries.

Options

index—Identifier for a specific event entry.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an RMON Event Entry and Its Attributes

alarm (SNMP RMON) | **2115**

event (CFM)

IN THIS SECTION

- Syntax | **6522**
- Hierarchy Level | **6522**
- Description | **6522**
- Options | **6522**
- Required Privilege Level | **6522**
- Release Information | **6523**

Syntax

```
event {
            adjacency-loss;
            interface-status-tlv {
                down;
                lower-layer-down;
            }
            port-status-tlv blocked;
            rdi;
        }
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile profile-name]
```

Description

Configure threshold values for connectivity fault management events in an action profile.

Options

adjacency-loss	Connectivity is lost.
interface-status-tlv [lower-layer-down down]	Values that need to be monitored in interface status TLV.
port-status-tlv	Values that need to be monitored in port status TLV.
rdi	RDI received from some MEP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1

RELATED DOCUMENTATION

Configuring a CFM Action Profile to Specify CFM Actions for CFM Events

interface-status-tlv | 9073

port-status-tlv | 12888

event (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 6523
- Hierarchy Level | 6524
- Description | 6524
- Options | 6524
- Required Privilege Level | 6524
- Release Information | 6524

Syntax

```
event {
    catastrophic-failure {
        reboot (master | standby);
    }
}
```

Hierarchy Level

```
[edit system services subscriber-management overrides]
```

Description

Configure the router to detect a type of event and automatically take action when the event is detected.

Options

catastrophic-failure Specify that the event is a corruption of the shared memory.

reboot (master | standby) Specify which Routing Engine is rebooted when the shared memory corruption is detected. You can configure both Routing Engines to be rebooted, but that requires two separate `set` commands.

> **NOTE**: If you do not configure this automatic reboot, then you must perform a manual reboot to clear out the shared memory. Otherwise, the memory remains corrupted and can cause applications that share the memory to generate core errors repeatedly.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

event (LFM)

IN THIS SECTION

- Syntax: MX, M, T, ACX Series Routers, SRX Firewalls and EX Series Switches | 6525
- Syntax: EX Series Switches and NFX Series Devices | 6525
- Hierarchy Level | 6526
- Description | 6526
- Required Privilege Level | 6526
- Release Information | 6526

Syntax: MX, M, T, ACX Series Routers, SRX Firewalls and EX Series Switches

```
event {
    link-adjacency-loss;
    link-event-rate {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
        }
    protocol-down;
}
```

Syntax: EX Series Switches and NFX Series Devices

```
event {
    link-adjacency-loss;
    link-event-rate {
        frame-error  count;
        frame-period count;
        frame-period-summary  count;
        symbol-period count;
```

```
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile]
```

Description

Configure link events in an action profile for Ethernet OAM link fault management (LFM).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Monitoring Protocol Status

Configuring Ethernet OAM Link Fault Management

event-options

IN THIS SECTION

- Syntax | 6527
- Hierarchy Level | 6530
- Description | 6530
- Required Privilege Level | 6530
- Release Information | 6530

Syntax

```
event-options {
    destinations {
        destination-name {
            archive-sites {
                url <password password>;
            }
            transfer-delay seconds;
        }
    }
    event-script {
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        file filename {
            checksum hash-algorithm hash-value;
            dampen {
                dampen-options {
                    cpu-factor cpu-factor;
                    line-interval line-interval;
                    time-interval time-interval;
                }
```

```
            }
            python-script-user username;
            refresh;
            refresh-from url;
            remote-execution {
                remote-hostname {
                    passphrase user-password;
                    username user-login;
                }
            }
            routing-instance routing-instance-name;
            source url;
            start {
                start-options {
                    mem-factor percentage;
                }
            }
        }
        max-datasize size;
        optional;
        refresh;
        refresh-from url;
        start {
            start-options {
                mem-factor percentage;
            }
        }
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
    generate-event {
        event-name (time-interval seconds <start-time start-time> | time-of-day hh:mm:ss) <no-
drift>;
    }
    max-policies policies;
    policy policy-name {
        attributes-match {
            event1.attribute-name equals event2.attribute-name;
            event.attribute-name matches regular-expression;
            event1.attribute-name starts-with event2.attribute-name;
```

```
        }
        events [events];
        then {
            change-configuration {
                commands {
                    "command";
                }
                commit-options {
                    check <synchronize>;
                    force;
                    log "comment-string";
                    synchronize;
                }
                retry count number interval seconds;
                user-name username;
            }
            event-script filename {
                arguments {
                    argument-name argument-value;
                }
                destination destination-name {
                    retry-count number retry-interval seconds;
                    transfer-delay seconds;
                }
                output-filename filename;
                output-format (text | xml);
                user-name username;
            }
            execute-commands {
                commands {
                    "command";
                }
                destination destination-name {
                    retry-count count retry-interval seconds;
                    transfer-delay seconds;
                }
                output-filename filename;
                output-format (text | xml);
                user-name username;
            }
            ignore;
            priority-override {
                facility facility-type;
```

```
                severity severity-level;
            }
            raise-trap;
            upload filename (filename | committed) destination destination-name {
                retry-count count retry-interval seconds;
                transfer-delay seconds;
                user-name username;
            }
        }
        within seconds {
            events [ events ];
            not events [ events ];
            trigger (after number | on number | until number);
        }
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure event policies and event scripts.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

event-rate

IN THIS SECTION

- Syntax | 6531
- Hierarchy Level | 6531
- Description | 6531
- Options | 6531
- Required Privilege Level | 6532
- Release Information | 6532

Syntax

```
event-rate rate;
```

Hierarchy Level

```
[edit security log]
[edit logical-systems name security log]
[edit tenants tenant-name security log]
```

Description

Limits the rate at which logs will be streamed per second.

Options

rate—Rate at which logs will be streamed per second.

- **Range:** 0 through 1500 logs per second
- **Default:** 1500 logs per second

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

The [edit `logical-systems` *name* security log] and [edit `tenants` *tenant-name* security log] hierarchy levels introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

syslog (System) | 28836

event-rate (Next Gen Services Service-Set Local System Logging)

IN THIS SECTION

- Syntax | 6532
- Hierarchy Level | 6533
- Description | 6533
- Required Privilege Level | 6533
- Release Information | 6533

Syntax

```
event-rate rate-per-second;
```

Hierarchy Level

```
[edit services services-set name syslog]
```

Description

Rate at which log messages are sent per second to the local file.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

event-script (Event Options)

IN THIS SECTION

- Syntax | 6534
- Hierarchy Level | 6535
- Description | 6535
- Required Privilege Level | 6535
- Release Information | 6535

Syntax

```
event-script {
    dampen {
        dampen-options {
            cpu-factor cpu-factor;
            line-interval line-interval;
            time-interval time-interval;
        }
    }
    file filename {
        cert-file path;
        checksum hash-algorithm hash-value;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        python-script-user username;
        refresh;
        refresh-from url;
        remote-execution {
            remote-hostname {
                passphrase user-password;
                username user-login;
            }
        }
        routing-instance routing-instance-name;
        source url;
        start {
            start-options {
                mem-factor percentage;
            }
        }
    }
    max-datasize size;
    optional;
    refresh;
    refresh-from url;
    start {
```

```
        start-options {
            mem-factor percentage;
        }
    }
    traceoptions {
        file <filename> <files number> <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit event-options]
```

Description

For Junos OS event scripts, configure scripting mechanisms.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Store and Enable Junos Automation Scripts

event-script (Event Policy)

IN THIS SECTION

- Syntax | 6536
- Hierarchy Level | 6536
- Description | 6536
- Required Privilege Level | 6537
- Release Information | 6537

Syntax

```
event-script filename {
    arguments {
        argument-name argument-value;
    }
    destination destination-name {
        retry-count count retry-interval seconds;
        transfer-delay seconds;
    }
    output-filename filename;
    output-format (text | xml);
    user-name username;
}
```

Hierarchy Level

```
[edit event-options policy policy-name then]
```

Description

Configure an event policy to execute one or more event scripts when triggered.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Execute Event Scripts in an Event Policy

Configure an Event Policy to Pass Arguments to an Event Script

event-thresholds

IN THIS SECTION

- Syntax | 6537
- Hierarchy Level | 6538
- Description | 6538
- Required Privilege Level | 6538
- Release Information | 6538

Syntax

```
event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
}
```

Hierarchy Level

```
[edit protocols oam link-fault-management interface interface-name]
```

Description

Configure threshold limit values for link events in periodic OAM PDUs.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring Threshold Values for Local Fault Events on an Interface

events (Event Policy Correlated Events)

IN THIS SECTION

- Syntax | 6539
- Hierarchy Level | 6539
- Description | 6539
- Options | 6539
- Required Privilege Level | 6539
- Release Information | 6539

Syntax

```
events [ events ];
```

Hierarchy Level

```
[edit event-options policy policy-name within seconds]
```

Description

Configure events that must occur within a specified time interval before one of the trigger events in order for the event policy to be triggered.

Options

[*events*] List of events. Events can be internally generated, or they can be generated by Junos OS processes.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Use Correlated Events to Trigger an Event Policy

events (Event Policy Trigger Events)

IN THIS SECTION

- Syntax | 6540
- Hierarchy Level | 6540
- Description | 6540
- Options | 6540
- Required Privilege Level | 6540
- Release Information | 6541

Syntax

```
events [ events ];
```

Hierarchy Level

```
[edit event-options policy policy-name]
```

Description

Configure one or more events that will trigger the event policy. If one or more of the listed events occurs, the policy is executed.

Options

[*events*] List of events. Events can be generated by Junos OS processes, or they can be internally generated.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Use Correlated Events to Trigger an Event Policy

evpn

IN THIS SECTION

- Syntax | 6541
- Hierarchy Level | 6544
- Description | 6545
- Options | 6545
- Required Privilege Level | 6546
- Release Information | 6546

Syntax

```
evpn {
    advertise-zero-ethernet-tag;
    assisted-replication {
        leaf {
            replicator-activation-delay seconds;
        }
        replicator {
            inet ip;
            vxlan-encapsulation-source-ip (assisted-replicator-ip | ingress-replication-ip |
retain-incoming-source-ip);
        }
    }
    default-gateway {
```

```
        (advertise | do-not-advertise | no-gateway-community);
    }
    designated-forwarder-election-hold-time seconds;
    (designated-forwarder-preference-highest | designated-forwarder-preference-least);
    duplicate-mac-detection {
        auto-recovery-time minutes;
        detection-threshold detection-threshold;
        detection-window seconds;
    }
    encapsulation (protocols evpn) encapslation-type;
    es-import-oldstyle;
    es-label-oldstyle;
    esi identifier;
    esi-resolution-per-bridge-domain;
    evpn-etree;
    extended-isid-list (all | single-isid | [isid-list] | isid-range);
    extended-vlan-list (vlan-id | [list of vlans]);
    extended-vni-list (all | [list of VNIs]);
    interconnect {
        esi identifier {
            all-active;
            df-election-type {
                mod;
                preference {
                    least;
                    non-revertive;
                    value value (0..65535);
                }
            }
        }
        interconnected-vlan-list (vlan-list);
        interconnected-vni-list (all | [vni-list]);
        route-distinguisher (number:id | ip-address:id);
        vrf-export (policy-name | [policy-names]);
        vrf-import (policy-name | [policy-names]);
        vrf-target {
            community;
            import community;
            export community;
        }
    }
    interconnect-multihoming-peer-gateways (multihoming-peer-gateway-IP-address | [multihoming-peer-
gateway-IP-address ... ]);
```

```
    interconnected-vlan-list (vlan-list);
    interconnected-vni-list (all | [vni-list]);
    interface interface-name {
        ignore-encapsulation-mismatch;
        interface-mac-limit limit {
            packet-action drop;
        }
        mac-pinning ;
        no-mac-learning;
        static-mac mac-address;
    }
    interface-mac-limit limit {
        packet-action drop;
    }
    irb-symmetric-routing {
        vni vni;
    }
    leave-sync-route-oldstyle;
    label-allocation per-instance;
    mac-list name {
        mac-addresses;
    }
    mac-mobility {
        no-sequence-numbers;
    }
    mac-statistics;
    mac-table-size limit {
        packet-action drop;
    }
    multicast-mode client | ingress-replication;
    no-core-isolation;
    no-default-gateway-ext-comm;
    no-mac-learning;
    oism {
        originate-smet-on-revenue-vlan-too;
        pim-evpn-gateway {
            external-irb name;
            peg-df-election {
              delay-time seconds;
              mod;
              preference {
                use-least;
                value value;
```

```
                }
            }
        }
        supplemental-bridge-domain-irb supplemental-bridge-domain-irb;
    }
    p2mp-bud-support;
    pbb-evpn-core;
    peg-df-election-history num;
    pmsi-tunnel-endpoint pmsi-tunnel-endpoint;
    remote-ip-host-routes {
        import [ import ... ];
        no-advertise-community;
    }
    smet-etag-carry-vid;
    source-packet-routing {
        srv6 locator loc-name (end-dt2-sid | micro-dt2-sid);
    }
    static-multihoming-peer (static-multihoming-peer-IP-address | [static-multihoming-peer-IP-
address ... ]);
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier>;
    }
    vni-options vni vxlan-network-identifier {
        designated-forwarder-election-hold-time seconds;
        vrf-target {
            community;
            auto;
            import community;
            export community;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable an Ethernet VPN (EVPN) on the routing instance.

> **NOTE**: Starting in Junos OS Release 24.2R1, we renamed the `interconnect-multihoming-peer-gateways` option to "multihoming-peer-gateways" on page 8806 to support additional multihoming peer feature use cases besides interconnect scenarios. You won't see the `interconnect-multihoming-peer-gateways` statement in the Junos CLI anymore; use the new option name `multihoming-peer-gateways` for the same function.

Options

designated-forwarder-election-hold-time *seconds*

Time in seconds to wait before electing a designated forwarder (DF). Peer provider edge (PE) devices elect a designated forwarder (DF) when customer edge devices (CEs) are multihomed to more than one PE device. Without a designated forwarder, multihomed hosts would receive duplicate packets. The designated forwarder choice for an Ethernet segment identifier (ESI) is based on EVPN Type 4 route advertisements.

- **Range:** 1 through 1800 seconds

designated-forwarder-preference-highest

Use highest preference value for DF election for the EVPN instance (EVI). By default, the preference-based DF election is based on the highest preference value configured.

designated-forwarder-preference-least

Use least preference value for DF election for the EVPN instance (EVI). By default, the preference-based DF election is based on the highest preference value configured.

es-import-oldstyle

Enable non-compliant ES import route-target computation.

es-label-oldstyle

Enable devices to use the lower-order bits for the ESI label field. This non-compliant ESI 24-bit ESI label can be found in previous Junos OS versions.

esi

Ten octet value. Configure an ESI for the EVPN instance (EVI). Use this option when configuring `df-election-granularity` "per-evi" on page 12409 to support the ESI per EVI feature.

evpn-etree

Configure an Ethernet VPN E-TREE service.

leave-sync-route-oldstyle

Enable device to interoperate with devices that do not support the maximum response time attributes on EVPN Type 8 routes.

mac-mobility no-sequence-numbers Disable MAC mobility. This feature is enabled by default. See *Overview of MAC Mobility*.

no-core-isolation Disable the core isolation feature (enabled by default on supporting platforms), which is desirable in certain EVPN fabric use cases. See *Understanding When to Disable EVPN-VXLAN Core Isolation* for details.

peg-df-election-history *num* Number of PIM EVPN gateway (PEG) DF election history entries the device maintains per bridge domain (VLAN) on devices configured as PEG devices with optimized intersubnet multicast (OISM) . The device stores and displays up to this many entries in the output from the `show evpn oism peg-df-status extensive` command.

- **Range:** 0 through 255

pmsi-tunnel-endpoint *pmsi-tunnel-endpoint* Configure the PMSI tunnel endpoint on the ingress router to use the loopback address (router ID) that is the MPLS network PMSI tunnel endpoint for EVPN Type 3 routes.

smet-etag-carry-vid Enable a device to interoperate with other devices in the EVPN network that use VNIs (VLAN ID) as the Ethernet Tag ID.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

`designated-forwarder-election-hold-time` *seconds* statement introduced in Junos OS Release 14.1.

`extended-vlan-list` *vlan-id* | [*vlan-id set*] statement introduced in Junos OS Release 14.1.

> **NOTE**: QFX10000 switches don't support the `extended-vlan-list` statement.

`designated-forwarder-preference-least` and `service-type` statements introduced in Junos OS Release 17.3 for MX Series routers with MPC interfaces and MIC interfaces.

`no-core-isolation` statement added in Junos OS Releases 17.3R3, 17.4R2, 18.1R3, 18.2R2, and 18.3R1.

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

`p2mp-bud-support` statement introduced in Junos OS Release 18.2R1.

`assisted-replication` statement introduced in Junos OS Release 18.4R2 and 19.4R1.

`pmsi-tunnel-endpoint` statement introduced in Junos OS Release 20.3R1.

`oism` statement hierarchy introduced in Junos OS Release 21.2R1.

`irb-symmetric-routing` statement introduced in Junos OS Release 21.4R1.

`peg-df-election-history` statement introduced in Junos OS Release 23.4R1 and Junos OS Evolved Release 23.4R1.

`designated-forwarder-preference-highest` statement introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Implementing EVPN-VXLAN for Data Centers

Configuring EVPN Routing Instances

Tracing EVPN Traffic and Operations

Provider Backbone Bridging (PBB) and EVPN Integration Overview

Optimized Inter-Subnet Multicast in EVPN Networks

evpn-aliasing-optimize

IN THIS SECTION

- Syntax | 6548
- Hierarchy Level | 6548
- Description | 6548
- Required Privilege Level | 6548
- Release Information | 6548

Syntax

```
evpn-aliasing-optimize;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table ],
[edit routing-options forwarding-table]
```

Description

Enables EVPN aliasing optimization. EVPN aliasing optimization allows the device to quickly recover when a link to a multihomed PE device fails. It reduces reduces convergence time, and improves device performance. By default, `evpn-aliasing-optimize` is enabled on the PTX10008 Router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Command introduced in Junos Evolve OS Release 21.1.

evpn (Easy EVPN LAG Configuration)

IN THIS SECTION

- Syntax | 6549
- Hierarchy Level | 6551
- Description | 6551
- Options | 6552
- Required Privilege Level | 6553

Release Information | 6553

Syntax

```
evpn {
  defaults-override {
    no-aggregate-device-count-config;
    no-loop-detect-config;
    no-platform-defaults-config;
    no-policy-and-routing-options-config;
    no-storm-control-config;
    no-underlay-config
  }
  device-attribute {
    loopback-subnet {
      peer1-subnet peer1-subnet;
      peer2-subnet peer2-subnet;
    }
    peer-id peer-id;
    peer-to-peer {
      overlay-connectivity {
        ibgp;
      }
      peer-subnet {
        inet inet;
        inet6 inet6;
        interface-name [ interface-name ... ];
      }
      underlay-connectivity {
        ospf;
      }
    }
    system-id system-id;
  }
  evpn-vxlan {
    dhcp-relay name {
       dhcp-server-address dhcp-server-address;
       relay-source relay-source;
       vrf-instance vrf-instance;
```

```
    }
    irb irb-instance {
      apply-config-groups apply-config-groups;
      instance {
        mac-vrf-instance instance-id;
        vrf-instance instance-id;
      }
      subnet-address {
        inet (inet-addr | [inet-addr1 inet-addr2 ...]);
        inet6 (inet6-addr | [inet6-addr1 inet6-addr2 ...]);
      }
      use-anycast-address;
      vlan-id vlan-id;
      vrf-instance vrf-instance;
    }
    server name {
      apply-config-groups apply-config-groups;
      enable-pxe-boot;
      (esi-lag-id esi-lag-id | single-home-id single-home-id);
      interface [ interface ... ];
      (mac-vrf-instance instance-id { vlan-id-list [ vlan-id-list ... ]} | vlan-id-list [ vlan-
id-list ... ]);
    }
  }
  global-parameters {
    anycast-mac anycast-mac;
    default-vlan-id vlan-id;
    mtu {
      overlay overlay-mtu;
      underlay underlay-mtu;
    }
    start-aggregate-ethernet-index start-aggregate-ethernet-index;
    virtual-gateway {
      v4-mac v4-mac;
      v6-mac v6-mac;
    }
    virtual-gateway-mac virtual-gateway-mac;
  }
}
```

Hierarchy Level

```
[edit services]
```

Description

Specify parameters to generate an Ethernet VPN (EVPN) fabric configuration with one or more customer edge (CE) devices or end servers that are:

- Single-homed to a provider edge (PE) device in the EVPN fabric.
- Multihomed on an Ethernet segment identified by an Ethernet segment identifier (ESI), and connected in a link aggregation group (LAG) to two or more PE devices in the EVPN Fabric.

> **NOTE**: We refer to this type of multihomed Ethernet segment connection as an ESI LAG. The feature you enable with statements in this hierarchy is an easy EVPN LAG configuration feature, also called EZ-LAG.

You use the statements at this hierarchy level to simplify setting up EVPN fabrics with ESI LAG configurations. You only need to configure a few statements to provide the parameters required for an EVPN core fabric with interfaces that connect to multihomed or single-homed end devices.

When you commit the configuration, the device accesses a commit script. The commit script uses the parameters from the [edit services evpn] statements to generate an EVPN fabric configuration that includes:

- Underlay and overlay peering configuration
- EVPN instance configuration (using MAC-VRF instances)
- Interfaces configuration for the ESI LAGs and IRB interfaces
- Host VLAN configuration
- Virtual routing and forwarding (VRF) routing instance, routing options, and policy options configurations
- Additional feature configuration such as:
  - Lightweight leaf to server loop detection
  - DHCP Relay
  - Storm control

The commit script generates a prescribed default EVPN fabric configuration. For mandatory EVPN fabric elements, the commit script uses default values or requires you to specify values for which it doesn't have prescribed defaults. For other elements, the commit script uses default algorithms to derive the values from the configured statements.

For some elements, you can choose to override generating configuration statements for those elements, and omit them or manually configure them instead. You can also add other standard EVPN fabric statements to the configuration manually as needed.

Options

defaults-override Select options to skip statements that the commit script would otherwise include by default when it generates the configuration:

- `no-aggregate-device-count-config`—Don't derive and include the aggregated Ethernet interface count in the generated configuration. You must manually configure the `set chassis aggregated-devices ethernet device-count` *num* statement on each peer PE device.

- `no-loop-detect-config`—Don't derive and add statements for the lightweight loop detection feature in the generated configuration. You can optionally configure loop detection manually.

- `no-platform-defaults-config`—Don't add the default statements for platform-specific options in the generated configuration. You must configure any required platform-specific statements manually.

- `no-policy-and-routing-options-config`—Don't derive and include the default routing options and policy options in the generated configuration.

  If you set this option, you must configure either of the following:

  - A routing policy with the name EXPORT-LO0 using `set policy-options policy-statement` commands. By default, when generating the EBGP peering configuration on a peer PE device, the commit script applies a policy with that name to advertise the device loopback interface to the device's neighbors.

  - The `no-underlay-config` option at the `[edit services evpn defaults-override]` hierarchy level to override the default behavior to generate the underlay peering configuration. In that case, you must manually configure the statements you need to set up the underlay peering.

- `no-storm-control-config`—Don't derive and include the storm control configuration elements. You can optionally configure storm control manually.

- `no-underlay-config`—Don't derive and add the underlay peering statements in the generated configuration. if you specify this override option, you must manually include the underlay peering configuration you need.

The remaining statements are explained in detail on separate pages.

You can also consult CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Easy EVPN LAG (EZ-LAG) Configuration

evpn-ssm-reports-only

IN THIS SECTION

- Syntax | 6554
- Hierarchy Level | 6554
- Description | 6554
- Required Privilege Level | 6555
- Release Information | 6555

Syntax

```
evpn-ssm-reports-only;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping],
[edit bridge-domains bridge-domain-name protocols mld-snooping],
[edit protocols igmp-snooping vlan vlan-name],
[edit protocols mld-snooping vlan vlan-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
snooping],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols mld-
snooping],
[edit routing-instances routing-instance-name protocols igmp-snooping],
[edit routing-instances routing-instance-name protocols mld-snooping]
```

Description

Accept and process only (S,G) reports for source-specific multicast (SSM) groups when:

- You configure Internet Group Management Protocol (IGMP) snooping for IGMP version 3 (IGMPv3) with IPv4 multicast traffic in an EVPN network.
- You configure Multicast Listener Discovery (MLD) snooping for MLD version 2 (MLDv2) with IPv6 multicast traffic in an EVPN network.

Devices use IGMP and MLD to discover and indicate interest in multicast groups. IGMP or MLD snooping restricts the forwarding of multicast traffic to only those interfaces in a bridge domain (VLAN) that have interested listeners.

By default, when IGMP snooping or MLD snooping is enabled, the device accepts and processes (*,G) (any-source multicast [ASM]) reports for:

- IGMP version 2 (IGMPv2)
- IGMPv3
- MLD version 1 (MLDv1)
- MLDv2

IGMPv3 and MLDv2 also support (S,G) source-specific multicast (SSM) reports and queries, but we don't support accepting and processing both (*,G) and (S,G) reports at the same time.

When you enable the `evpn-ssm-reports-only` option with IGMP snooping or MLD snooping, the device honors only IGMPv3 or MLDv2 (S,G) reports. The device drops any (*,G) reports for IGMPv2, IGMPv3, MLDv1, or MLDv2.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.4R1.

Statement added in Junos OS Releases 19.4R2 and 20.3R1 on ACX Series routers with EVPN multicast support.

RELATED DOCUMENTATION

Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment

Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment

Configure Multicast Forwarding with MLD Snooping in an EVPN-MPLS Environment

evpn-vxlan

IN THIS SECTION

- Syntax | 6556
- Hierarchy Level | 6556
- Description | 6556
- Options | 6556
- Default | 6558
- Required Privilege Level | 6558
- Release Information | 6558

Syntax

```
evpn-vxlan {
   gbp {
       ingress-enforcement;
       ingress-src-tag filter <gbp-filter>
       mac-ip-inter-tagging;
       tag-only-policy;
   }
   host-originated-packets {
       gbp-src-tag gbp-src-tag;
   }
   reroute-address {
       inet address;
       inet6 address;
   }
   shared-tunnels;
   vxlan-trans-vni-enable;
}
```

Hierarchy Level

```
[edit forwarding-options],
[edit logical-systems name routing-instances name forwarding-options]
```

Description

Set various global or per-routing-instance options in an EVPN-VXLAN fabric.

Options

gbp ingress-enforcement	Configure Group Based Policy (GBP) policy enforcement at the ingress node.
gbp ingress-src-tag filter *<gbp-filter>*	Associate the specified *<gbp-filter>* with the routing instance.

Use this option to enable EVPN Type 5 tag propagation, which is the ability to pass a source GBP tag within the extended BGP communities string of an EVPN Type 5 advertisement. By propagating the source GBP tag across the network, remote endpoints gain the information necessary to enforce GBP policies at their ingress.

The specified *<gbp-filter>* must be an IP-based (IP address or IP prefix) GBP filter.

This option has no effect if the specified *<gbp-filter>* is not an IP-based filter.

gbp mac-ip-inter-tagging

Allow MAC-based GBP filters and IP-based GBP filters to apply to both routed and switched traffic.

If disabled, then MAC-based GBP filters apply only to switched traffic, and IP-based GBP filters apply only to routed traffic.

This setting has no effect on other GBP filters, such as port and VLAN GBP filters.

> **NOTE**: The Packet Forwarding Engine (PFE) restarts when you set or delete this option.

gbp tag-only-policy

Allow only GBP source and destination tags as matches in GBP policy.

Configuring the IPv4 and IPv6 L4 filters can reduce the supported GBP scale. These matches are supported by default. However, on the EX4650 series, QFX5120-32C, and QFX5120-48Y switches, you can use this option to allow only GBP source and destination tags as matches in GBP policy.

> **NOTE**: The Packet Forwarding Engine (PFE) restarts when you set or delete this option.

host-originated-packets gbp-src-tag *gbp-src-tag*

Insert the specified GBP source tag in all host originated packets.

Range for *gbp-src-tag*: 1 through 65535.

reroute-address (inet *address* | inet6 *address*)

Enable fast reroute egress link protection (ELP) on provider edge (PE) devices for access links to multihomed customer edge (CE) devices in an EVPN-VXLAN fabric. With this statement, you also specify the IPv4 or IPv6 device reroute address, which is a secondary loopback address you configure as the fast reroute ELP tunnel source VTEP address for an Ethernet segment.

This feature helps to minimize load-balanced traffic loss when the link from a PE device to a multihomed CE device goes down.

shared-tunnels Create shared VXLAN tunnel endpoint (VTEP) tunnels to EVPN provider edge (PE) devices to save VTEP resources.

> **NOTE**: When you configure the `shared-tunnels` option, you must reboot the device for the setting to take effect.

vxlan-trans-vni-enable Enable translated VXLAN network identifiers (VNIs) with Data Center Interconnect (DCI).

> **NOTE**: We require this option on QFX5120 switches when you configure EVPN-VXLAN stitching with EVPN Type 2 routes in DCI topologies.
>
> The PFE restarts when you set or delete this option.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Default

Not enabled

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1 for supported EX4100, EX4400, EX4650, and QFX5120 Series switches.

`tag-only-policy` option introduced in 23.2R1 for the QFX5120-48YM, QFX5120-48Y, and EX4650-48Y Series switches.

`reroute-address` option introduced in Junos OS 21.3R1 for the QFX5110 and QFX5120 Series swtiches.

`shared-tunnels` option introduced in Junos OS 20.3R1.

`vxlan-trans-vni-enable` option introduced in Junos OS 23.2R1 for QFX5120 switches.

`gbp ingress-src-tag filter` option introduced in Junos OS 24.2R1 for the EX4400, EX4650, and QFX5120 Series switches.

RELATED DOCUMENTATION

Example: Micro and Macro Segmentation using Group Based Policy in a VXLAN

Fast Reroute for Egress Link Protection with EVPN-VXLAN Multihoming

evpn-vxlan (Easy EVPN LAG Configuration)

IN THIS SECTION

- Syntax | 6559
- Hierarchy Level | 6560
- Description | 6560
- Required Privilege Level | 6560
- Release Information | 6561

Syntax

```
evpn-vxlan {
  dhcp-relay name {
    dhcp-server-address dhcp-server-address;
    relay-source relay-source;
    vrf-instance vrf-instance;
  }
  irb irb-instance {
    apply-config-groups apply-config-groups;
    instance {
      mac-vrf-instance mac-vrf-instance;
      vrf-instance vrf-instance;
    }
    subnet-address {
      inet (inet-addr | [inet-addr1 inet-addr2 ...]);
      inet6 (inet6-addr | [inet6-addr1 inet6-addr2 ...]);
    }
    use-anycast-address;
    vlan-id vlan-id;
  }
```

```
  server name {
    apply-config-groups apply-config-groups;
    enable-pxe-boot;
    (esi-lag-id esi-lag-id | single-home-id single-home-id);
    interface [ interface ... ];
    (mac-vrf-instance name {vlan-id-list [ vlan-id-list ... ]} | vlan-id-list [ vlan-id-
list ... ]);
  }
}
```

Hierarchy Level

```
[edit services evpn]
```

Description

Specify parameters the device uses to generate an Ethernet VPN (EVPN) fabric configuration with Virtual Extensible LAN (VXLAN) encapsulation.

You use the statements in the `[edit services evpn]` hierarchy level to simplify configuring an EVPN fabric with Ethernet segment identifier (ESI) link aggregation groups (LAGs) in the topology. The feature you enable with statements in this hierarchy is an easy EVPN LAG configuration feature, also called EZ-LAG.

You only need to configure a few statements to provide the parameters required for an EVPN core fabric with multihomed or single-homed customer edge (CE) devices (or end server devices). When you commit the `[edit services evpn]` configuration statements, the device uses a commit script to generate the corresponding standard EVPN fabric configuration statements.

The statements at the `evpn-vxlan` hierarchy level provide the parameters the commit script needs to generate the required EVPN-VXLAN configuration.

The statements at this hierarchy level are explained in detail on separate pages.

You can also see CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Easy EVPN LAG (EZ-LAG) Configuration

examine-dhcp

IN THIS SECTION

- Syntax | 6561
- Hierarchy Level | 6561
- Description | 6562
- Default | 6562
- Required Privilege Level | 6562
- Release Information | 6562

Syntax

```
(examine-dhcp | no-examine-dhcp) {
    forwarding-class class-name;
}
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name)]
```

Description

Enable DHCP snooping on all VLANs or on the specified VLAN.

> **NOTE**: If you configure DHCP for all VLANs and you enable a different port security feature on a specific VLAN, you must also explicitly enable DHCP snooping on that VLAN. Otherwise, the default value of no DHCP snooping applies to that VLAN.

- **examine-dhcp**—Enable DHCP snooping.
- **no-examine-dhcp**—Disable DHCP snooping.

When DHCP snooping is enabled, the switch logs DHCP packets (DHCPOFFER, DHCPDECLINE, DHCPACK, and DHCPNAK packets) that it receives on untrusted ports. You can monitor the log for these messages, which can signal the presence of a malicious DHCP server on the network.

> **TIP**: For private VLANs (PVLANs), enable DHCP snooping on the primary VLAN. If you enable DHCP snooping only on a community VLAN, DHCP messages coming from PVLAN trunk ports are not snooped.

The remaining statement is explained separately.

Default

Disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Port Security (non-ELS)

Example: Protecting Against Address Spoofing and Layer 2 DoS Attacks

Example: Protecting Against ARP Spoofing Attacks

Example: Prioritizing Snooped and Inspected Packet

Enabling DHCP Snooping (non-ELS)

Enabling DHCP Snooping (J-Web Procedure)

examine-dhcpv6

IN THIS SECTION

- Syntax | 6563
- Hierarchy Level | 6563
- Description | 6563
- Default | 6564
- Required Privilege Level | 6564
- Release Information | 6564

Syntax

```
examine-dhcpv6 {
    forwarding-class class-name;
}
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name)]
```

Description

Enable DHCPv6 snooping on all VLANs or on the specified VLAN.

> **NOTE**: If you configure DHCP for all VLANs and you enable a different port security feature on a specific VLAN, you must also explicitly enable DHCP snooping on that VLAN. Otherwise, the default value of no DHCP snooping applies to that VLAN.

When DHCP snooping is enabled, the switch logs DHCP packets (DHCPOFFER, DHCPDECLINE, DHCPACK, and DHCPNAK packets) that it receives on untrusted ports. You can monitor the log for these messages, which can signal the presence of a malicious DHCP server on the network.

> **TIP**: For private VLANs (PVLANs), enable DHCP snooping on the primary VLAN. If you enable DHCP snooping only on a community VLAN, DHCP messages coming from PVLAN trunk ports are not snooped.

The remaining statement is explained separately.

Default

Disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

- *Example: Configuring Port Security (non-ELS)*
- *Example: Protecting Against Address Spoofing and Layer 2 DoS Attacks*
- *Example: Protecting Against ARP Spoofing Attacks*
- *Example: Prioritizing Snooped and Inspected Packet*
- *Enabling DHCP Snooping (non-ELS)*
- Enabling DHCP Snooping (J-Web Procedure)

examine-fip

IN THIS SECTION

- Syntax | 6565
- Hierarchy Level | 6565
- Description | 6565
- Required Privilege Level | 6566
- Release Information | 6566

Syntax

```
examine-fip {
    examine-vn2vn {
        beacon-period milliseconds;
    }
    fc-map fc-map-value;
    no-fip-snooping-scaling;
}
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name)]
```

Description

NOTE: This statement supports the original CLI. If your switch runs the Enhanced Layer 2 Software (ELS) CLI, see "examine-vn2vf" on page 6566 for VN_Port to VF_Port (VN2VF_Port) FIP snooping, and see "examine-vn2vn" on page 6568 for VN_Port to VN_Port (VN2VN_Port) FIP snooping. For ELS details, see Using the Enhanced Layer 2 Software CLI.

Enable FIP snooping on a specified VLAN. Ensure that the VLAN is a dedicated FCoE VLAN that transports only FCoE traffic.

(QFX Series only) Enable VN2VN_Port FIP snooping on the specified VLAN. The VLAN must be a dedicated FCoE VLAN that transports only VN2VN_Port traffic. One FCoE VLAN cannot support both VN2VF_Port FIP snooping and VN2VN_Port FIP snooping. Configure separate, dedicated FCoE VLANs for VN2VN_Port FIP snooping and VN2VN_Port FIP snooping.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Statement `examine-vn2vn` introduced in Junos OS Release 12.2 for the QFX Series.

Statement `no-fip-snooping-scaling` introduced in Junos OS Release 13.2X52-D10 for the QFX Series.

RELATED DOCUMENTATION

vlan (Access Port Security) | **30785**

Example: Configuring an FCoE Transit Switch

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

examine-vn2vf

IN THIS SECTION

- Syntax | **6567**
- Hierarchy Level | **6567**
- Description | **6567**
- Required Privilege Level | **6567**

Release Information | 6567

Syntax

```
examine-vn2vf {
}
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options fip-security]
```

Description

NOTE: This statement supports the Enhanced Layer 2 Software (ELS) CLI. If your switch runs the original (non-ELS) software, see examine-fip. For ELS details, see Using the Enhanced Layer 2 Software CLI.

Enable VN_Port to VF_Port (VN2VF_Port) FIP snooping on the specified VLAN. Ensure that the VLAN is a dedicated FCoE VLAN that transports only FCoE traffic.

If the switch also performs VN_Port to VN_Port (VN2VN_Port) FIP snooping, ensure that the VN2VN_Port traffic is on a different VLAN than the VN2VF_Port traffic. You cannot mix VN2VF_Port and VN2VN_Port traffic in the same VLAN, so you must use separate VLANs for VN2VF_Port and VN2VN_Port traffic.

The remaining statement is explained separately. Also, see CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

examine-vn2vn | 6568

Understanding VN_Port to VF_Port FIP Snooping on an FCoE Transit Switch

Understanding FCoE Transit Switch Functionality

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

examine-vn2vn

IN THIS SECTION

- Syntax | 6568
- Hierarchy Level | 6568
- Description | 6569
- Required Privilege Level | 6569
- Release Information | 6569

Syntax

```
examine-vn2vn {
    beacon-period milliseconds;
}
```

Hierarchy Level

Original CLI

```
[edit ethernet-switching options secure-access-port vlan (all | vlan-name) examine-fip]
```

ELS CLI for Platforms that Support FCoE

```
[edit vlans vlan-name forwarding-options fip-security]
```

NOTE: The `examine-vn2vn` configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.

Description

Enable VN_Port to VN_Port (VN2VN) FIP snooping on a specified VLAN. The VLAN must be a dedicated FCoE VLAN that transports only FCoE traffic. A VLAN cannot support VN2VN FIP snooping and VN_Port to VF_Port FIP snooping (VN2VF) simultaneously. Configure separate VLANs for VN2VN FIP snooping and VN2VF FIP snooping.

When you enable VN2VN FIP snooping on a VLAN, the VN2VF session filters are removed and the all existing VN2VF sessions are terminated.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)

Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches)

Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch)

exceed-action

IN THIS SECTION

- Syntax | 6570
- Hierarchy Level | 6570
- Description | 6570
- Required Privilege Level | 6570
- Release Information | 6571

Syntax

```
exceed-action {
    drop;
    syslog;
}
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name max-data-sessions-per-subscriber]
```

Description

Specify the action if the maximum data sessions per subscriber exceed the maximum limit. You must also specify the drop rate of the packets for `drop` and system log details for `syslog`.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

exception

IN THIS SECTION

- Syntax | 6571
- Hierarchy Level | 6571
- Description | 6571
- Required Privilege Level | 6572
- Release Information | 6572

Syntax

```
exception listname;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus mime-whitelist]
[edit security utm feature-profile anti-virus mime-whitelist list listname]
```

Description

Configure the antivirus scanner to use an exception list to the MIME bypass list (custom objects). To use the exception list, you first create a allowlist custom-object list with the `list` statement. The system will first look at any existing allowlist mime pattern. If it matches an item, it will then continue to look for any exceptions to the allowlist and will then scan any item in the exception list.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

exception-reporting

IN THIS SECTION

- Syntax | 6572
- Hierarchy Level | 6573
- Description | 6573
- Options | 6573
- Required Privilege Level | 6573
- Release Information | 6573

Syntax

```
exception-reporting {
   category category-name {
      inline-monitoring-instance inline-monitoring-instance;
   }
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pfe identifier]
```

Description

Enable reporting of exceptions in the forwarding path.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 for MX Series routers.

Statement introduced in Junos Evolved OS Release 22.2R1 for PTX Series routers.

RELATED DOCUMENTATION

Juniper Resiliency Interface

exception (Security Content Filtering)

IN THIS SECTION

- Syntax | 6574
- Hierarchy Level | 6574

- Description | 6574
- Required Privilege Level | 6574
- Release Information | 6574

Syntax

```
exception list-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name block-mime]
```

Description

Configure the content filter to use an exception list to the MIME block list (custom objects).

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

exceptions

IN THIS SECTION

- Syntax (Junos OS) | 6575
- Syntax (Junos OS Evolved) | 6575
- Hierarchy Level | 6576
- Description | 6576
- Default | 6576
- Options | 6576
- Required Privilege Level | 6576
- Release Information | 6576

Syntax (Junos OS)

```
exceptions {
   forwarding;
   os;
   routing;
}
```

Syntax (Junos OS Evolved)

```
exceptions {
   forwarding;
   routing;
}
```

Hierarchy Level

```
[edit system resiliency]
```

Description

Subscribe to forwarding, operating system, and routing exceptions for the on-box collector.

Default

The on-box collector is disabled.

Options

forwarding Subscribe to forwarding exceptions.

os (Junos OS only) Subscribe to operating system exceptions.

routing Subscribe to routing exceptions.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 for MX Series routers.

Statement introduced in Junos Evolved OS Release 22.2R1 for PTX Series routers.

RELATED DOCUMENTATION

Juniper Resiliency Interface

inline-monitoring | **8672**

primary-data-record-fields | **13134**

exception-reporting | **6572**

excess-bandwidth-share

IN THIS SECTION

- Syntax | 6577
- Hierarchy Level | 6577
- Description | 6577
- Options | 6578
- Required Privilege Level | 6578
- Release Information | 6578

Syntax

```
excess-bandwidth-share (proportional value | equal);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-set interface-set-name]
```

Description

Specify the method of sharing excess bandwidth in a hierarchical scheduler environment. With hierarchical schedulers, you can provide shaping and scheduling at the service VLAN level as well as other levels, such as the physical interface. You can also group a set of logical interfaces and then apply scheduling and shaping parameters to the logical interface set.

To configure CoS hierarchical schedulers, you must enable hierarchical scheduling by including the `hierarchical-scheduler` statement at the `[edit interfaces]` hierarchy for the physical interface. If you do not include this statement, the interfaces on the MX Series router cannot use hierarchical interfaces.

The Enhanced Queuing DPC supports the following hierarchical scheduler characteristics:

- Shaping at the physical interface level
- Shaping and scheduling at the service VLAN interface set level

- Shaping and scheduling at the customer VLAN *logical interface* level
- Scheduling at the queue level

Options

equal Share excess bandwidth equally among the configured interfaces.

proportional *value* (Default) Share excess bandwidth proportionally according to the specified value. In this mode, the excess bandwidth is shared at the ratio of the logical interface shaping rates.

- **Default:** 32.64 Mbps

This example sets the excess bandwidth sharing for an Enhanced Queuing DPC interface proportionally at a rate of 100 Mbps and a shaping rate of 80 Mbps applied to the interface through the output-traffic-control profile for scheduling and shaping:

```
[edit class-of-service interfaces interface-set example-interface-set]
user@host# set excess-bandwidth-share proportional 100m
user@host# set output-traffic-control-profile PIR-80Mbps
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Hierarchical Schedulers for CoS

Configuring Interface Sets

Enhanced Queuing DPC CoS Properties

Configuring MDRR on Enhanced Queuing DPCs

excess-burst-size (QFX)

IN THIS SECTION

- Syntax | 6579
- Hierarchy Level | 6579
- Description | 6579
- Options | 6579
- Required Privilege Level | 6580
- Release Information | 6580

Syntax

```
excess-burst-size bytes;
```

Hierarchy Level

```
[edit firewall three-color-policer policer-name single-rate]
```

Description

Configure the maximum number of bytes allowed for incoming traffic to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red).

> **NOTE**: When you include the `excess-burst-size` statement in the configuration, you must also include the **committed-burst-size** and `committed-information-rate` statements at the same hierarchy level.

Options

bytes—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

- **Range:** 512 bytes through 268435456 bytes (268 MB)

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Two-Color and Three-Color Policers to Control Traffic Rates

Overview of Policers

excess-burst-size

IN THIS SECTION

- Syntax | 6580
- Hierarchy Level | 6581
- Description | 6581
- Options | 6581
- Required Privilege Level | 6582
- Release Information | 6582

Syntax

```
excess-burst-size bytes;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name single-rate],
[edit firewall three-color-policer policer-name single-rate]
```

Description

For a single-rate three-color policer, configure the excess burst size (EBS) as a number of bytes. The EBS allows for moderate periods of bursting traffic that exceeds both the committed information rate (CIR) and the committed burst size (CBS).

> **NOTE**: When you include the `excess-burst-size` statement in the configuration, you must also include the `committed-burst-size` and `committed-information-rate` statements at the same hierarchy level.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policing uses a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red based on the `excess-burst-size` statement included in the policer configuration.

During periods of traffic that conforms to the CIR, any unused portion of the guaranteed bandwidth capacity accumulates in the first token bucket, up to the maximum number of bytes defined by the CBS. If any accumulated bandwidth capacity overflows the first bucket, the excess accumulates in a second token bucket, up to the maximum number of bytes defined by the EBS.

A nonconforming traffic flow is categorized yellow if its size conforms to bandwidth capacity accumulated in the first token bucket. Packets in a yellow flow are marked with `medium-high` packet loss priority (PLP) and then passed through the interface.

A nonconforming traffic flow is categorized red if its size exceeds the bandwidth capacity accumulated in the second token bucket. Packets in a red traffic flow are marked with `high` PLP and then either passed through the interface or optionally discarded.

Options

bytes—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1500 through 100,000,000,000 bytes

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

Support at the `[edit dynamic-profiles ... single-rate]` hierarchy level introduced in Junos Release OS 11.4.

RELATED DOCUMENTATION

- Three-Color Policer Configuration Overview
- Policer Bandwidth and Burst-Size Limits
- Policer Color-Marking and Actions
- Dual Token Bucket Algorithms
- Determining Proper Burst Size for Traffic Policers
- committed-burst-size | 4458
- committed-information-rate | 4460

excess-priority

IN THIS SECTION

- Syntax | 6583
- Hierarchy Level | 6583
- Description | 6583
- Options | 6583
- Required Privilege Level | 6583
- Release Information | 6583

Syntax

```
excess-priority [ low | medium-low | medium-high | high | none];
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Determine the priority of excess bandwidth traffic on a scheduler.

> **NOTE**: For Link Services IQ (LSQ) PICs or Multiservices PIC (MS-PICs), the `excess-priority` statement is allowed for consistency, but ignored. If an explicit priority is not configured for these interfaces, a default low priority is used. This default priority is also used in the excess region.

Options

`low`—Excess traffic for this scheduler has low priority.

`medium-low`—Excess traffic for this scheduler has medium-low priority.

`medium-high`—Excess traffic for this scheduler has medium-high priority.

`high`—Excess traffic for this scheduler has high priority.

`none`—System does not demote the priority of guaranteed traffic when the bandwidth exceeds the shaping rate or the guaranteed rate.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Option `none` introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Excess Bandwidth Sharing on IQE PICs

Bandwidth Sharing on Nonqueuing Packet Forwarding Engines Overview

Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs

excess-priority (Dynamic Schedulers)

IN THIS SECTION

- Syntax | 6584
- Hierarchy Level | 6584
- Description | 6584
- Options | 6585
- Required Privilege Level | 6585
- Release Information | 6585

Syntax

```
excess-priority (low | high | $junos-cos-scheduler-excess-priority | none);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
```

Description

Determine the priority of excess bandwidth traffic on a scheduler in a dynamic profile.

Options

`low`—Excess traffic for this scheduler has low priority.

`high`—Excess traffic for this scheduler has high priority.

`$junos-cos-scheduler-excess-priority`—Variable for the excess-priority that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

`none`—System does not demote the priority of guaranteed traffic when the bandwidth exceeds the shaping rate or the guaranteed rate.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

`none` option added in Junos OS Release 11.4.

RELATED DOCUMENTATION

Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces

scheduler (Dynamic Scheduler Maps) | **16344**

excess-rate (Schedulers)

IN THIS SECTION

- Syntax | **6586**
- Hierarchy Level | **6586**
- Description | **6586**

- Options | 6587
- Required Privilege Level | 6587
- Release Information | 6587

Syntax

```
excess-rate (percent percentage | proportion value);
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name],
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
```

Description

Determine the percentage of excess port bandwidth for which a queue (forwarding class) that is not a strict-high priority queue or forwarding class set (priority group) can contend. Excess bandwidth is the extra port bandwidth left after strict-high priority queues and the guaranteed minimum bandwidth requirements of other queues (as determined by each queue's transmit rate) are satisfied. With the exception of strict-high priority queues, the switch allocates extra port bandwidth to queues or to priority groups based on the configured excess rate. If you do not configure an excess rate for a queue, the default excess rate is the same as the transmit rate.

You cannot configure an excess rate on strict-high priority queues. Strict-high priority queues receive extra bandwidth based on an extra bandwidth sharing weight of "1", which is not configurable. However, the switch serves traffic on strict-high priority queues up to the configured transmit rate before it serves any other queues, so by configuring an appropriate transmit rate on a strict-high priority queue, you can guarantee strict-high priority traffic on that queue is treated in the manner you want.

> **NOTE**: For platforms that support the `proportion` option, this option provides a greater range of values over the `percent` option and hence influences the priorities assigned to the queues.

Options

percentage—Percentage of the excess bandwidth to share.

- **Range:** 0 through 100 percent
- **Default:** Excess bandwidth is shared in proportion to the configured transmit rate of each queue.

value—Proportion of the excess bandwidth to share.

- **Range:** 0 through 1000

> **NOTE**: Configure MPCs with increments of 10 from 0 through 1000.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Scheduler Transmission Rate

Configuring Excess Bandwidth Sharing on IQE PICs

Allocating Excess Bandwidth Among Frame Relay DLCIs on Multiservices PICs

Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs

Defining CoS Queue Schedulers for Port Scheduling

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Understanding CoS Port Schedulers on QFX Switches

Understanding CoS Traffic Control Profiles

excess-rate (Dynamic Schedulers)

IN THIS SECTION

- Syntax | 6588
- Hierarchy Level | 6588
- Description | 6588
- Options | 6588
- Required Privilege Level | 6589
- Release Information | 6589

Syntax

```
excess-rate percent (percentage | $junos-cos-scheduler-excess-rate | proportion value);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
```

Description

Determine the percentage of excess bandwidth traffic to share.

Options

percentage—Percentage of the excess bandwidth to share.

- **Range:** 0 through 100 percent

`$junos-cos-scheduler-excess-rate`—Variable for the excess rate that is specified for a scheduler. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

value—Proportion of the excess bandwidth to share.

- **Range:** 0 through 1000

> **NOTE**: The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 0 through 1000.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces

output-traffic-control-profile (Dynamic CoS Definition) | **12018**

excess-rate (Dynamic Traffic Shaping)

IN THIS SECTION

- Syntax | **6590**
- Hierarchy Level | **6590**
- Description | **6590**
- Options | **6590**
- Required Privilege Level | **6590**
- Release Information | **6591**

Syntax

```
excess-rate (percent percentage | $junos-cos-excess-rate) | proportion value);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For an MPC interface, determine the percentage or proportion of excess bandwidth traffic to share for all priorities of traffic.

Options

percentage—Percentage of the excess bandwidth to share.

- **Range:** 0 through 100 percent

value—Proportion of the excess bandwidth to share.

- **Range:** 0 through 1000

> **NOTE**: The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 0 through 1000.

`$junos-cos-excess-rate`—Variable for the excess rate that is specified for the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces

output-traffic-control-profile (Dynamic CoS Definition) | **12018**

excess-rate-high

IN THIS SECTION

- Syntax | **6591**
- Hierarchy Level | **6591**
- Description | **6592**
- Options | **6592**
- Required Privilege Level | **6592**
- Release Information | **6592**

Syntax

```
excess-rate-high (percent percentage | proportion value);
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
```

Description

For a MIC or MPC interface, determine the percentage of excess bandwidth from high-priority traffic to share.

Options

percentage—Percentage of the excess bandwidth to share.

- **Range:** 0 through 100 percent

proportion—Proportion of the excess bandwidth to share.

- **Range:** 0 through 1000

> **NOTE**: The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 0 through 1000.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs

excess-rate-high (Dynamic Traffic Shaping)

IN THIS SECTION

- Syntax | 6593
- Hierarchy Level | 6593
- Description | 6593
- Options | 6593
- Required Privilege Level | 6594
- Release Information | 6594

Syntax

```
excess-rate-high ((percent percentage | $junos-cos-excess-rate-high) | proportion value);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For an MPC/MIC interface, determine the percentage of excess bandwidth for high-priority traffic to share.

Options

percentage—Percentage of the excess bandwidth to share.

- **Range:** 0 through 100 percent

value—Proportion of the excess bandwidth to share.

- **Range:** 0 through 1000

> **NOTE**: The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 0 through 1000.

`$junos-cos-excess-rate-high`—Variable for the excess rate that is specified for high-priority traffic on the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces

output-traffic-control-profile (Dynamic CoS Definition) | **12018**

excess-rate-low

IN THIS SECTION

- Syntax | **6595**
- Hierarchy Level | **6595**
- Description | **6595**
- Options | **6595**
- Required Privilege Level | **6595**
- Release Information | **6595**

Syntax

```
excess-rate-low (percent percentage | proportion value);
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
```

Description

For a MIC or MPC interface, determine the percentage of excess bandwidth from low-priority traffic to share.

Options

percentage—Percentage of the excess bandwidth to share.

- **Range:** 0 through 100 percent

value—Proportion of the excess bandwidth to share.

- **Range:** 0 through 1000

> **NOTE**: The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 0 through 1000.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs

excess-rate-low (Dynamic Traffic Shaping)

IN THIS SECTION

- Syntax | 6596
- Hierarchy Level | 6596
- Description | 6596
- Options | 6596
- Required Privilege Level | 6597
- Release Information | 6597

Syntax

```
excess-rate-low ((percent percentage | $junos-cos-excess-rate-low) | proportion value);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For an MPC/MIC interface, determine the percentage of excess bandwidth for low-priority traffic to share.

Options

percentage—Percentage of the excess bandwidth to share.

- **Range:** 0 through 100 percent

value—Proportion of the excess bandwidth to share.

- **Range:** 0 through 1000

> **NOTE**: The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 0 through 1000.

`$junos-cos-excess-rate-low`—Variable for the excess rate that is specified for low-priority traffic on the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces

output-traffic-control-profile (Dynamic CoS Definition) | **12018**

excess-rate-medium-high

IN THIS SECTION

- Syntax | **6598**
- Hierarchy Level | **6598**
- Description | **6598**
- Options | **6598**

Required Privilege Level | 6598
Release Information | 6598

Syntax

```
excess-rate-medium-high (percent percentage | proportion value);
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
```

Description

For MPC7E-MRATE, MPC7E-10G, MPC8E, and MPC9E interfaces, determine the percentage of excess bandwidth from high-priority traffic to share.

Options

percentage—Percentage of the excess bandwidth to share.

- **Range:** 1 through 100 percent

proportion—Proportion of the excess bandwidth to share.

- **Range:** 0 through 1000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

enhanced-priority-mode | 6404
excess-rate-medium-low | 6599
Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs
Per-Priority Shaping on MIC and MPC Interfaces Overview
shaping-rate-excess-medium-high | 16943
shaping-rate-excess-medium-low | 16945
shaping-rate-priority-medium-low | 16956
shaping-rate-priority-strict-high | 16958

excess-rate-medium-low

IN THIS SECTION

- Syntax | 6599
- Hierarchy Level | 6599
- Description | 6600
- Options | 6600
- Required Privilege Level | 6600
- Release Information | 6600

Syntax

```
excess-rate-medium-low (percent percentage | proportion value);
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
```

Description

For MPC7E-MRATE, MPC7E-10G, MPC8E, and MPC9E interfaces, determine the percentage of excess bandwidth from low-priority traffic to share.

Options

percentage—Percentage of the excess bandwidth to share.

- **Range:** 1 through 100 percent

value—Proportion of the excess bandwidth to share.

- **Range:** 0 through 1000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

enhanced-priority-mode | **6404**

excess-rate-medium-high | **6597**

Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs

Per-Priority Shaping on MIC and MPC Interfaces Overview

shaping-rate-excess-medium-high | **16943**

shaping-rate-excess-medium-low | **16945**

shaping-rate-priority-medium-low | **16956**

shaping-rate-priority-strict-high | **16958**

excess-rate (Schedulers)

IN THIS SECTION

- Syntax | 6601
- Hierarchy Level | 6601
- Description | 6601
- Default | 6601
- Options | 6601
- Required Privilege Level | 6602
- Release Information | 6602

Syntax

```
excess-rate {
    percent percentage;
}
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

(EX4300 switches only) Specify the percentage of excess bandwidth traffic to share.

Default

Excess bandwidth is shared in proportion to the configured transmit rate of each queue.

Options

- `percent`—Percentage of the excess bandwidth to share.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Understanding CoS Schedulers

Defining CoS Schedulers and Scheduler Maps (CLI Procedure)

Defining CoS Schedulers (J-Web Procedure)

Example: Configuring CoS on EX Series Switches

exclude (Diameter Gx Profiles)

IN THIS SECTION

- Syntax | 6602
- Hierarchy Level | 6603
- Description | 6603
- Options | 6603
- Required Privilege Level | 6603
- Release Information | 6603

Syntax

```
exclude {
    an-gw-address;
    default-eps-bearer-qos;
    packet-filter-information;
```

```
    packet-filter-operation;
    rat-type;
}
```

Hierarchy Level

```
[edit unified-edge diameter-profiles gx-profile profile-name attributes]
```

Description

Configure the attribute-value pairs (AVPs) to be excluded from the credit control request (CCR) messages between the MX Series router and the policy and charging enforcement function (PCEF).

Options

`an-gw-address`—Exclude the AN-GW-Address AVP.

`default-eps-bearer-qos`—Exclude the Default-EPS-Bearer-QoS AVP.

`packet-filter-information`—Exclude the Packet-Filter-Information AVP.

`packet-filter-operation`—Exclude the Packet-Filter-Operation AVP.

`rat-type`—Exclude the RAT-Type AVP.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

gx-profile | 7954

exclude (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 6604
- Hierarchy Level | 6604
- Description | 6604
- Required Privilege Level | 6604
- Release Information | 6604

Syntax

```
exclude;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name static group multicast-group-address]
```

Description

Configure the group to operate in exclude mode on the dynamic interface. In exclude mode all sources except the address configured are accepted for the group. By default, the group operates in include mode.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Enabling MLD Static Group Membership

exclude (for Administrative Groups)

IN THIS SECTION

- Syntax | 6605
- Hierarchy Level | 6605
- Description | 6605
- Options | 6606
- Required Privilege Level | 6606
- Release Information | 6606

Syntax

```
exclude [ group-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name admin-group],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name admin-group],
[edit protocols mpls label-switched-path lsp-name admin-group],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name admin-group]
```

Description

Define the administrative groups to exclude for an LSP or for a path's primary and secondary paths.

Options

group-names—Names of one or more groups defined with the "admin-groups" on page 1973 statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Administrative Groups for LSPs

exclude (for Fast Reroute)

IN THIS SECTION

- Syntax | 6606
- Hierarchy Level | 6607
- Description | 6607
- Options | 6607
- Required Privilege Level | 6607
- Release Information | 6607

Syntax

```
(exclude [ group-names ] | no-exclude);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-reroute],
[edit protocols mpls label-switched-path lsp-name fast-reroute]
```

Description

Control exclusion of administrative groups:

- `exclude`—Define the administrative groups to exclude for fast reroute.
- `no-exclude`—Disable administrative group exclusion.

Options

`group-names`—Names of one or more groups defined with the `admin-groups` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Fast Reroute

admin-groups | **1973**

exclude-hostname

IN THIS SECTION

- Syntax | 6608
- Hierarchy Level | 6608
- Description | 6608
- Required Privilege Level | 6608
- Release Information | 6608

Syntax

```
exclude-hostname;
```

Hierarchy Level

```
[edit system syslog host hostname]
```

Description

Disable logging of hostname in the message directed to remote host.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2

exclude (JSRC Attributes)

IN THIS SECTION

- Syntax | 6609
- Hierarchy Level | 6609
- Description | 6609
- Options | 6609
- Required Privilege Level | 6610
- Release Information | 6610

Syntax

```
exclude {
    user-name [ authorization-request | provisioning-request ];
}
```

Hierarchy Level

```
[edit access profile profile-name jsrc attributes]
```

Description

Configure the router to exclude the specified attribute-value pair (AVP) from the specified Diameter message for JSRC.

Options

- `user-name`—User-Name AVP (1). Diameter AVP name and number.
- `authorization-request`—address-authorization request in AAR message sent from JSRC to the SAE.
- `provisioning-request`—provisioning-request in AAR message sent from JSRC to the SAE.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Excluding AVPs from Diameter Messages for JSRC

Understanding JSRC-SAE Interactions

exclude-protocol

IN THIS SECTION

- Syntax | 6610
- Hierarchy Level | 6611
- Description | 6611
- Default | 6611
- Options | 6611
- Required Privilege Level | 6611
- Release Information | 6611

Syntax

```
exclude-protocol protocol-name;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Specifies protocols whose packets are not secured using Media Access Control Security (MACsec) when MACsec is enabled on a link using static connectivity association key (CAK) security mode.

When this option is enabled in a connectivity association that is attached to an interface, MACsec is not enabled for all packets of the specified protocols that are sent and received on the link.

Default

Disabled.

All packets are secured on a link when MACsec is enabled, with the exception of all types of Spanning Tree Protocol (STP) packets.

Options

protocol-name Specifies the name of the protocol that should not be MACsec-secured. Options include:

- **cdp**—Cisco Discovery Protocol.
- **lacp**—Link Aggregation Control Protocol.
- **lldp**—Link Level Discovery Protocol.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced for SRX Series Firewalls in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

exclude-protocol (MX Series)

IN THIS SECTION

- Syntax | **6612**
- Hierarchy Level | **6612**
- Description | **6612**
- Default | **6613**
- Options | **6613**
- Required Privilege Level | **6613**
- Release Information | **6613**

Syntax

```
exclude-protocol protocol-name;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Specifies protocols whose packets are not secured using Media Access Control Security (MACsec) when MACsec is enabled on a link using static connectivity association key (CAK) security mode.

When this option is enabled in a connectivity association that is attached to an interface, MACsec is not enabled for all packets of the specified protocols that are sent and received on the link.

Default

Disabled.

All packets are secured on a link when MACsec is enabled, with the exception of all types of Spanning Tree Protocol (STP) packets.

Options

protocol-name Specifies the name of the protocol that should not be MACsec-secured. Options include:

- **cdp**—Cisco Discovery Protocol.
- **lacp**—Link Aggregation Control Protocol.
- **lldp**—Link Level Discovery Protocol.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

exclude (Protocols IGMP)

IN THIS SECTION

- Syntax | 6614
- Hierarchy Level | 6614
- Description | 6614
- Required Privilege Level | 6614
- Release Information | 6615

Syntax

```
exclude;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name static group
multicast-group-address],
[edit protocols igmp interface interface-name static group multicast-group-address]
```

Description

Configure the static group to operate in exclude mode. In exclude mode all sources except the address configured are accepted for the group. If this statement is not included, the group operates in include mode.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Enabling IGMP Static Group Membership

exclude (Protocols MLD)

IN THIS SECTION

- Syntax | 6615
- Hierarchy Level | 6615
- Description | 6615
- Required Privilege Level | 6616
- Release Information | 6616

Syntax

```
exclude;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name static group
multicast-group-address],
[edit protocols mld interface interface-name static group multicast-group-address]
```

Description

Configure the static group to operate in exclude mode. In exclude mode all sources except the address configured are accepted for the group. By default, the group operates in include mode.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Enabling MLD Static Group Membership

exclude-queue-overhead-bytes

IN THIS SECTION

- Syntax | 6616
- Hierarchy Level | 6617
- Description | 6617
- Options | 6617
- Required Privilege Level | 6617
- Release Information | 6617

Syntax

```
exclude-queue-overhead-bytes {
    include-hierarchy;
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name]
```

Description

By default, the Layer 2 header bytes applied to upper-level logical interfaces are included in CoS per-queue statistics at the physical interface, which can provide inaccurate results. This options enables you to exclude the counting of overhead bytes from aggregate queue statistics.

Options

include-hierarchy Exclude the counting of overhead bytes from aggregate queue statistics of all child interfaces, including logical interfaces and interface sets, as well.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show class-of-service interface | 19056

exclude (RADIUS Attributes)

IN THIS SECTION

- Syntax | 6618
- Hierarchy Level | 6620

- Description | 6620
- Options | 6621
- Required Privilege Level | 6625
- Release Information | 6625

Syntax

```
exclude {
    acc-aggr-cir-id-asc [ access-request | accounting-start | accounting-stop ];
    acc-aggr-cir-id-bin [ access-request | accounting-start | accounting-stop ];
    acc-loop-cir-id [ access-request | accounting-start | accounting-stop ];
    acc-loop-encap [ access-request | accounting-start | accounting-stop ];
    acc-loop-remote-id [ access-request | accounting-start | accounting-stop ];
    accounting-authentic [ accounting-off | accounting-on | accounting-start | accounting-stop ]
    accounting-delay-time [ accounting-off | accounting-on | accounting-start | accounting-
stop ];
    accounting-session-id access-request;
    accounting-terminate-cause accounting-off;
    acct-request-reason [ accounting-start | accounting-stop ];
    acct-tunnel-connection [ access-request | accounting-start | accounting-stop ];
    act-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    act-data-rate-up [ access-request | accounting-start | accounting-stop ];
    act-interlv-delay-dn [ access-request | accounting-start | accounting-stop ];
    act-interlv-delay-up [ access-request | accounting-start | accounting-stop ];
    att-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    att-data-rate-up [ access-request | accounting-start | accounting-stop ];
    called-station-id [ access-request | accounting-start | accounting-stop ];
    calling-station-id [ access-request | accounting-start | accounting-stop ];
    chargeable-user-identity access-request;
    class [ accounting-start | accounting-stop ];
    cos-shaping-rate [ accounting-start | accounting-stop ];
    delegated-ipv6-prefix [ accounting-start | accounting-stop ];
    dhcp-gi-address [ access-request | accounting-start | accounting-stop ];
    dhcp-header access-request;
    dhcp-mac-address [ access-request | accounting-start | accounting-stop ];
    dhcp-options [ access-request | accounting-start | accounting-stop ];
    dhcpv6-header access-request;
    dhcpv6-options [ access-request | accounting-start | accounting-stop ];
```

```
    downstream-calculated-qos-rate [ access-request | accounting-start | accounting-stop ];
    dsl-forum-attributes [ access-request | accounting-start | accounting-stop ];
    dsl-line-state [ access-request | accounting-start | accounting-stop ];
    dsl-type [ access-request | accounting-start | accounting-stop ];
    dynamic-iflset-name [ accounting-start | accounting-stop ];
    event-timestamp [ accounting-off | accounting-on | accounting-start | accounting-stop ];
    filter-id [ accounting-start | accounting-stop ];
    first-relay-ipv4-address [ access-request | accounting-start | accounting-stop ];
    first-relay-ipv6-address [ access-request | accounting-start | accounting-stop ];
    framed-interface-id [ access-request | accounting-start | accounting-stop ];
    framed-ip-address [ access-request | accounting-start | accounting-stop ];
    framed-ip-netmask [ access-request | accounting-start | accounting-stop ];
    framed-ip-route [ accounting-start | accounting-stop ];
    framed-ipv6-address [ access-request | accounting-start | accounting-stop ];
    framed-ipv6-pool [ accounting-start | accounting-stop ];
    framed-ipv6-prefix [ accounting-start | accounting-stop ];
    framed-ipv6-route [ accounting-start | accounting-stop ];
    framed-pool [ accounting-start | accounting-stop ]; input-ipv6-gigawords accounting-stop;
    input-filter [ accounting-start | accounting-stop ];
    input-gigapackets accounting-stop;
    input-gigawords accounting-stop;
    input-ipv6-octets accounting-stop;
    input-ipv6-packets accounting-stop;
    interface-description [ access-request | accounting-start | accounting-stop ];
    l2c-downstream-data [ access-request | accounting-start | accounting-stop ];
    l2c-upstream-data [ access-request | accounting-start | accounting-stop ];
    l2tp-rx-connect-speed [ access-request | accounting-start | accounting-stop ];
    l2tp-tx-connect-speed [ access-request | accounting-start | accounting-stop ];
    max-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    max-data-rate-up [ access-request | accounting-start | accounting-stop ];
    max-interlv-delay-dn [ access-request | accounting-start | accounting-stop ];
    max-interlv-delay-up [ access-request | accounting-start | accounting-stop ];
    min-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    min-data-rate-up [ access-request | accounting-start | accounting-stop ];
    min-lp-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    min-lp-data-rate-up [ access-request | accounting-start | accounting-stop ];
    nas-identifier [ access-request | accounting-off | accounting-on | accounting-start |
accounting-stop ];
    nas-port [ access-request | accounting-start | accounting-stop ];
    nas-port-id [ access-request | accounting-start | accounting-stop ];
    nas-port-type [ access-request | accounting-start | accounting-stop ];
    output-filter [ accounting-start | accounting-stop ];
    output-gigapackets accounting-stop;
```

```
    output-gigawords accounting-stop;
    output-ipv6-gigawords accounting-stop;
    output-ipv6-octets accounting-stop;
    output-ipv6-packets accounting-stop;
    pppoe-description [ access-request | accounting-start | accounting-stop ];
    standard-attribute number {
        packet-type [ access-request | accounting-off | accounting-on | accounting-start |
accounting-stop ];
    }
    tunnel-assignment-id [ access-request | accounting-start | accounting-stop ];
    tunnel-client-auth-id [ access-request | accounting-start | accounting-stop ];
    tunnel-client-endpoint [ access-request | accounting-start | accounting-stop ];
    tunnel-medium-type [ access-request | accounting-start | accounting-stop ];
    tunnel-server-auth-id [ access-request | accounting-start | accounting-stop ];
    tunnel-server-endpoint [ access-request | accounting-start | accounting-stop ];
    tunnel-type [ access-request | accounting-start | accounting-stop ];
    upstream-calculated-qos-rate [ access-request | accounting-start | accounting-stop ];
    vendor-id id-number {
        vendor-attribute vsa-number {
            packet-type [ access-request | accounting-off | accounting-on | accounting-start |
accounting-stop ];
        }
    }
    virtual-router [ access-request | accounting-start | accounting-stop ];
}
```

Hierarchy Level

```
[edit access profile profile-name radius attributes]
```

Description

Configure the router or switch to exclude the specified attributes from being sent in the specified type of RADIUS message. Exclusion can be useful, for example, for attributes that do not change values over the lifetime of a subscriber. By not sending these attributes, you reduce the packet size without losing information. Contrast this behavior with that provided by the `ignore` statement.

You can specify attribute exclusion for multiple RADIUS message types by enclosing the message types, separated by spaces, within brackets ([ ]). You do not need brackets when specifying a single message type.

Starting in Junos OS Release 18.1R1, you can specify standard RADIUS attributes with the attribute number. You can specify VSAs with the IANA-assigned vendor ID and the VSA number. With this flexible configuration method, you can configure any standard attribute and VSA supported by your platform to be excluded. The configuration has no effect if you configure unsupported attributes, vendors, and VSAs.

The legacy method allows you to configure only those attributes and VSAs for which the statement syntax includes a specific option. Consequently, you can use the legacy method to exclude only a subset of all attributes that can be received in Access-Accept messages.

Not all attributes are available in all types of RADIUS messages.

NOTE: If you exclude an attribute from Acct-Off messages, the attributes are then excluded from Interim-Acct messages.

NOTE: VSAs with dedicated option names include Juniper Networks (IANA vendor ID 4874) and DSL Forum (vendor ID 3561) VSAs.

Options

RADIUS attribute—RADIUS standard attribute or VSA:

- `acc-aggr-cir-id-asc`—Exclude Juniper Networks VSA 26-112, Acc-Aggr-Cir-Id-Asc.
- `acc-aggr-cir-id-bin`—Exclude Juniper Networks VSA 26-111, Acc-Aggr-Cir-Id-Bin.
- `acc-loop-cir-id`—Exclude Juniper Networks VSA 26-110, Acc-Loop-Cir-Id.
- `acc-loop-encap`—Exclude Juniper Networks VSA 26-183, Acc-Loop-Encap.
- `acc-loop-remote-id`—Exclude Juniper Networks VSA 26-182, Acc-Loop-Remote-Id.
- `accounting-authentic`—Exclude RADIUS attribute 45, Acct-Authentic.
- `accounting-delay-time`—Exclude RADIUS attribute 41, Acct-Delay-Time.
- `accounting-session-id`—Exclude RADIUS attribute 44, Acct-Session-Id.
- `accounting-terminate-cause`—Exclude RADIUS attribute 49, Acct-Terminate-Cause.
- `acct-request-reason`—Exclude Juniper Networks VSA 26-210, Acct-Request-Reason.
- `acct-tunnel-connection`—Exclude RADIUS attribute 68, Acct-Tunnel-Connection.
- `act-data-rate-dn`—Exclude Juniper Networks VSA 26-114, Act-Data-Rate-Dn.

- `act-data-rate-up`—Exclude Juniper Networks VSA 26-113, Act-Data-Rate-Up.
- `act-interlv-delay-dn`—Exclude Juniper Networks VSA 26-126, Act-Interlv-Delay-Dn.
- `act-interlv-delay-up`—Exclude Juniper Networks VSA 26-124, Act-Interlv-Delay-Up.
- `att-data-rate-dn`—Exclude Juniper Networks VSA 26-118, Att-Data-Rate-Dn.
- `att-data-rate-up`—Exclude Juniper Networks VSA 26-117, Att-Data-Rate-Up.
- `called-station-id`—Exclude RADIUS attribute 30, Called-Station-Id.
- `calling-station-id`—Exclude RADIUS attribute 31, Calling-Station-Id.
- `chargeable-user-identity`—Exclude RADIUS attribute 89, Chargeable-User-Identity.
- `class`—Exclude RADIUS attribute 25, Class.
- `cos-shaping-rate`—Exclude Juniper Networks VSA 26-177, Cos-Shaping-Rate.
- `delegated-ipv6-prefix`—Exclude RADIUS attribute 123, Delegated-IPv6-Prefix.
- `dhcp-gi-address`—Exclude Juniper Networks VSA 26-57, DHCP-GI-Address.
- `dhcp-header`—Exclude Juniper Networks VSA 26-208, DHCP-Header.
- `dhcp-mac-address`—Exclude Juniper Networks VSA 26-56, DHCP-MAC-Address.
- `dhcp-options`—Exclude Juniper Networks VSA 26-55, DHCP-Options.
- `dhcpv6-header`—Exclude Juniper Networks VSA 26-209, DHCPv6-Header.
- `dhcpv6-options`—Exclude Juniper Networks VSA 26-207, DHCPv6-Options.
- `dynamic-iflset-name`—Exclude Juniper Networks VSA 26-130, Qos-Set-Name.
- `downstream-calculated-qos-rate`—Exclude Juniper Networks VSA 26-141.
- `dsl-forum-attributes`—Exclude DSL Forum VSA (vendor ID 3561) as described in RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*.
- `dsl-line-state`—Exclude Juniper Networks VSA 26-127, DSL-Line-State.
- `dsl-type`—Exclude Juniper Networks VSA 26-128, DSL-Type.
- `event-timestamp`—Exclude RADIUS attribute 55, Event-Timestamp.
- `filter-id`—Exclude RADIUS attribute 11, Filter-Id.
- `first-relay-ipv4-address` —Exclude Juniper Networks VSA 26-189, DHCP-First-Relay-IPv4-Address.

- `first-relay-ipv6-address` —Exclude Juniper Networks VSA 26-190, DHCP-First-Relay-IPv6-Address.
- `framed-interface-id`—Exclude RADIUS attribute 96, Framed-Interface-ID.
- `framed-ip-address`—Exclude RADIUS attribute 8, Framed-IP-Address.
- `framed-ip-netmask`—Exclude RADIUS attribute 9, Framed-IP-Netmask.
- `framed-ip-route`—Exclude RADIUS attribute 22, Framed-Route.
- `framed-ipv6-address`—Exclude RADIUS attribute 168, Framed-IPv6-Address.
- `framed-ipv6-pool`—Exclude RADIUS attribute 100, Framed-IPv6-Pool.
- `framed-ipv6-prefix`—Exclude RADIUS attribute 97, Framed-IPv6-Prefix.
- `framed-ipv6-route`—Exclude RADIUS attribute 99, Framed-IPv6-Route.
- `framed-pool`—Exclude RADIUS attribute 88, Framed-Pool.
- `input-filter`—Exclude Juniper Networks VSA 26-10, Ingress-Policy-Name.
- `input-gigapackets`—Exclude Juniper Networks VSA 26-42, Acct-Input-Gigapackets.
- `input-gigawords`—Exclude RADIUS attribute 52, Acct-Input-Gigawords.
- `input-ipv6-gigawords`—Exclude Juniper Networks VSA 26-155, Acct-Input-IPv6-Gigawords.
- `input-ipv6-octets`—Exclude Juniper Networks VSA 26-151, Acct-Input-IPv6-Octets.
- `input-ipv6-packets`—Exclude Juniper Networks VSA 26-153, Acct-Input-IPv6-Packets.
- `interface-description`—Exclude Juniper Networks VSA 26-53, Interface-Desc.
- `l2c-downstream-data`—Exclude Juniper Networks VSA 26-93, L2C-Down-Stream-Data.
- `l2c-upstream-data`—Exclude Juniper Networks VSA 26-92, L2C-Up-Stream-Data.
- `l2tp-rx-connect-speed`—Exclude Juniper Networks VSA 26-163, Rx-Connect-Speed.
- `l2tp-tx-connect-speed`—Exclude Juniper Networks VSA 26-162, Tx-Connect-Speed.
- `max-data-rate-dn`—Exclude Juniper Networks VSA 26-120, Max-Data-Rate-Dn.
- `max-data-rate-up`—Exclude Juniper Networks VSA 26-119, Max-Data-Rate-Up.
- `max-interlv-delay-dn`—Exclude Juniper Networks VSA 26-125, Max-Interlv-Delay-Dn.
- `max-interlv-delay-up`—Exclude Juniper Networks VSA 26-123, Max-Interlv-Delay-Up.
- `min-data-rate-dn`—Exclude Juniper Networks VSA 26-116, Min-Data-Rate-Dn.

- `min-data-rate-up`—Exclude Juniper Networks VSA 26-115, Min-Data-Rate-Up.
- `min-lp-data-rate-dn`—Exclude Juniper Networks VSA 26-122, Min-Lp-Data-Rate-Dn.
- `min-lp-data-rate-up`—Exclude Juniper Networks VSA 26-121, Min-Lp-Data-Rate-Up.
- `nas-identifier`—Exclude RADIUS attribute 32, NAS-Identifier.
- `nas-port`—Exclude RADIUS attribute 5, NAS-Port.
- `nas-port-id`—Exclude RADIUS attribute 87, NAS-Port-Id.
- `nas-port-type`—Exclude RADIUS attribute 61, NAS-Port-Type.
- `output-filter`—Exclude Juniper Networks VSA 26-11, Egress-Policy-Name.
- `output-gigapackets`—Exclude Juniper Networks VSA 26-43, Acct-Output-Gigapackets.
- `output-gigawords`—Exclude RADIUS attribute 53, Acct-Output-Gigawords.
- `output-ipv6-gigawords`—Exclude Juniper Networks VSA 26-156, Acct-Output-IPv6-Gigawords.
- `output-ipv6-octets`—Exclude Juniper Networks VSA 26-152, Acct-Output-IPv6-Octets.
- `output-ipv6-packets`—Exclude Juniper Networks VSA 26-154, Acct-Output-IPv6-Packets.
- `packet-type`—Specify the RADIUS message type to exclude; term required when excluding a standard attribute or VSA by number rather than name. You can enclose multiple values in square brackets to specify a list of message types. Message types include Access-Request, Accounting-Off, Accounting-Off, Accounting-Start, and Accounting-Stop.
- `pppoe-description`—Exclude Juniper Networks VSA 26-24, PPPoE-Description.
- `standard-attribute` *`number`*—RADIUS standard attribute number supported by your platform. If you configure an unsupported attribute, that configuration has no effect. When you use this option, you must use the `packet-type` term to specify the message from which the attribute is excluded.
- `tunnel-assignment-id`—Exclude RADIUS attribute 82, Tunnel-Assignment-ID.
- `tunnel-client-auth-id`—Exclude RADIUS attribute 90. Tunnel-Client-Auth-ID.
- `tunnel-client-endpoint`—Exclude RADIUS attribute 66, Tunnel-Client-Endpoint.
- `tunnel-medium-type`—Exclude RADIUS attribute 65, Tunnel-Medium-Type.
- `tunnel-server-auth-id`—Exclude RADIUS attribute 91, Tunnel-Server-Auth-ID.
- `tunnel-server-endpoint`—Exclude RADIUS attribute 67, Tunnel-Server-Endpoint.
- `tunnel-type`—Exclude RADIUS attribute 64, Tunnel-Type.

- `upstream-calculated-qos-rate`—Exclude Juniper Networks VSA 26-142
- `vendor-attribute` *`vsa-number`*—Number identifying a VSA belonging to the specified vendor; both must be supported by your platform. If you configure an unsupported VSA, that configuration has no effect. When you use this option, you must use the `packet-type` term to specify the message from which the attribute is excluded.
- `vendor-id` *`id-number`*—IANA vendor ID supported by your platform. If you configure an unsupported vendor ID, that configuration has no effect.
- `virtual-router`—Exclude Juniper Networks VSA 26-1.

RADIUS message type:

- `access-request`—RADIUS Access-Request messages.
- `accounting-off`—RADIUS Accounting-Off messages.
- `accounting-on`—RADIUS Accounting-On messages.
- `accounting-start`—RADIUS Accounting-Start messages.
- `accounting-stop`—RADIUS Accounting-Stop messages.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`downstream-calculated-qos-rate`, `dsl-forum-attributes`, and `upstream-calculated-qos-rate` options added in Junos OS Release 11.4.

`cos-shaping-rate` and `filter-id` options added in Junos OS Release 13.2.

`pppoe-description` option added in Junos OS Release 14.2.

`virtual-router` option added in Junos OS Release 15.1.

`first-relay-ipv4-address` and `first-relay-ipv6-address` options added in Junos OS Release 16.1.

`acc-loop-encap` and `acc-loop-remote-id` options added in Junos OS Release 16.1R4.

`access-request` option support for all tunnel attributes added in Junos OS Release 15.1R7, 16.1R5, 16.2R2, 17.1R2, 17.2R2, and 17.3R1 for MX Series.

`packet-type`, `standard-attribute`, `vendor-attribute`, and `vendor-id` options added in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

Standard and Vendor-Specific RADIUS Attributes

exclude (Schedulers)

IN THIS SECTION

- Syntax | 6626
- Hierarchy Level | 6626
- Description | 6627
- Required Privilege Level | 6627
- Release Information | 6627

Syntax

```
exclude;
```

Hierarchy Level

```
[edit schedulers scheduler  scheduler-name  daily],
[edit schedulers scheduler  scheduler-name  friday],
[edit schedulers scheduler  scheduler-name monday],
[edit schedulers scheduler  scheduler-name saturday],
[edit schedulers scheduler  scheduler-name  sunday],
[edit schedulers scheduler  scheduler-name   tuesday],
```

```
[edit schedulers scheduler  scheduler-name  thursday],
[edit schedulers scheduler  scheduler-name  wednesday]
```

Description

Exclude a specified day from the schedule.

Use the `exclude` statement to exclude a day from a daily schedule created with the `daily` statement. You cannot use the `exclude` statement for a particular day unless it is in conjunction with the `daily` statement in a schedule.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policy Schedulers Overview

exclude (Security Log)

IN THIS SECTION

- Syntax | 6628
- Hierarchy Level | 6628
- Description | 6628
- Options | 6628
- Required Privilege Level | 6629
- Release Information | 6629

Syntax

```
exclude exlude-name {
    destination-address destination-address;
    destination-port destination-port;
    event-id event-id;
    failure;
    interface-name interface-name;
    policy-name policy-name;
    process process-name;
    protocol protocol;
    source-address source-address;
    source-port source-port;
    success;
    user-name user-name;
}
```

Hierarchy Level

```
[edit security log cache]
[edit logical-systems name security log cache]
[edit tenants tenant-name security log cache]
```

Description

Configure a list of auditable events that can be excluded from the audit log.

Options

- `destination-ip` *`destination-address`*—Destination IP address.
- `destination-port` *`destination-port`*—Destination port number.
- `event-id` *`event-id`*—Error message identification number.
- `failure`—Failed audit event logs.
- `interface-name` *`interface-name`*—Name of the interface.
- `policy-name` *`policy-name`*—Policy name filter.

- `process` *`process-name`*—Process that generated the event.
- `protocol` *`protocol`*—Protocol that generated the event.
- `source-ip` *`source-address`*—Source IP address.
- `source-port` *`source-port`*—Source port number.
- `success`—Successful audit event logs.
- `username` *`user-name`*—User name filter.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The [edit `logical-systems` *name* security log cache] and [edit `tenants` *tenant-name* security log cache] hierarchy levels introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

show security log | 24769

clear security log | 3998

exclude-srlg

IN THIS SECTION

- Syntax | 6630
- Hierarchy Level | 6630
- Description | 6630
- Required Privilege Level | 6630

Release Information | 6631

Syntax

```
exclude-srlg;
```

Hierarchy Level

```
[edit protocols mpls],
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls label-switched-path path-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path path-name],
[edit protocols rsvp interface interface-name link-protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit protocols rsvp interface interface-name link-protection bypass destination],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection bypass destination]
```

Description

Exclude Shared Risk Link Group (SRLG) links for the secondary path for critical links where it is imperative to keep the secondary and primary label-switched paths completely disjoint from any common SRLG.

When specified, the Constrained Shortest Path First (CSPF) algorithm excludes any link belonging to the set of SRLGs in the primary path. When not specified and if a link belongs to the set of SRLGs in the primary path, CSPF adds the SRLG cost to the metric, but still accepts the link for computing the path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Excluding SRLG Links Completely for the Secondary LSP

excluded-address (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 6631
- Hierarchy Level | 6631
- Description | 6631
- Required Privilege Level | 6632
- Release Information | 6632

Syntax

```
excluded-address ip-address;
```

Hierarchy Level

```
[edit access address assignment-address pool]
```

Description

Allows you to exclude select IPv4 or IPv6 addresses from a DHCP address pool. Within a configured address pool, you can specifically exclude up to 20 addresses. Junos OS will not assign these excluded addresses to any clients. If you configure an excluded address that has already been assigned to a DHCP client, that excluded address will be revoked from the client.

NOTE: Excluded addresses must match the address family of the configured address pool. For example, you cannot exclude an IPv4 address within an IPv6 address pool.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D20.

RELATED DOCUMENTATION

pool (Address-Assignment Pools) | 12766

Preventing Addresses from Being Allocated from an Address Pool

excluded-range (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 6633
- Hierarchy Level | 6633
- Description | 6633
- Options | 6633
- Required Privilege Level | 6633
- Release Information | 6633

Syntax

```
excluded-range name low minimum-value high maximum-value;
```

Hierarchy Level

```
[edit access address-assignment pool name family (inet | inet6)],
[edit logical-systems name access address-assignment pool name family (inet | inet6)],
[edit logical-systems name routing-instances name access address-assignment pool name family (inet
| inet6)],
[edit routing-instances name access address-assignment pool name family (inet | inet6)]
```

Description

Specify a range of consecutive addresses to exclude from consideration when addresses are allocated from the corresponding address pool. For example, you might want to reserve certain addresses or ranges to be used only for static subscribers. If a range you configure for exclusion includes an address that has already been allocated, the subscriber that has that address is logged out. The address is then deallocated and marked for exclusion from future allocation.

Options

high Upper limit of excluded range of addresses.

low Lower limit of range of addresses.

name Name of a range of addresses to exclude from the address pool for the specified family.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Preventing Addresses from Being Allocated from an Address Pool

Address-Assignment Pools Overview

Address-Assignment Pool Configuration Overview

exclusive-mac

IN THIS SECTION

- Syntax | 6634
- Hierarchy Level | 6634
- Description | 6634
- Options | 6635
- Required Privilege Level | 6635
- Release Information | 6635

Syntax

```
exclusive-mac virtual-mac-mac-address/mask;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

Exclude MAC addresses from the MAC move limit algorithm.

The global MAC move feature is used to track MAC addresses when they appear on a different physical interface or within a different unit of the same physical interface. When you configure the `exclusive-mac` *`virtual-mac-mac-address/mask`* parameter at the [`edit protocols l2-learning global-mac-move`] hierarchy level,

specified MAC addresses are excluded and will not be tracked. In addtion, excluded MAC addresses cannot be pinned when `mac-pinning` is set on an interface.

This feature can be useful in OVSDB-managed topologies with VRRP servers deployed in a redundancy configuration (primary/member), and when MAC move limit is configured. Both servers could negotiate primary role, and the same MAC address could be learned under the global MAC move feature while negotiation is occurring. In such cases, excluding the MAC address of the VRRP servers by using the `exclusive-mac` statement prevents this “false” move from being tracked.

The following example excludes VRRP V2 virtual router MAC addresses, as defined in RFC 3768:

```
[edit]
set protocols l2-learning global-mac-move exclusive-mac 00:00:5e:00:01:00/40
```

The following example excludes VRRP V3 virtual router MAC addresses, as defined in RFC 5798:

```
[edit]
set protocols l2-learning global-mac-move exclusive-mac 00:00:5e:00:02:00/40
```

Options

virtual-mac-mac-address/mask	Specify a MAC address and a mask. If the mask is 48, only the exact MAC address is excluded. If the mask is 40, all the MAC addresses that have the same first 5 bytes are excluded.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D45.

RELATED DOCUMENTATION

Configuring MAC Move Parameters

execute-commands (Event Policy)

IN THIS SECTION

- Syntax | 6636
- Hierarchy Level | 6636
- Description | 6636
- Required Privilege Level | 6637
- Release Information | 6637

Syntax

```
execute-commands {
    commands {
        "command";
    }
    destination destination-name {
        retry-count count retry-interval seconds;
        transfer-delay seconds;
    }
    output-filename filename;
    output-format (text | xml);
    user-name username;
}
```

Hierarchy Level

```
[edit event-options policy policy-name then]
```

Description

Specify the operational mode commands to be issued when the corresponding event policy is triggered. You can also specify the format of the command output and the filename identifier and destination for the output file.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configure an Event Policy to Execute Operational Mode Commands

exercise

IN THIS SECTION

- Syntax | 6637
- Hierarchy Level | 6637
- Description | 6638
- Required Privilege Level | 6638

Syntax

```
request protection-group ethernet-aps exercise md <md> ma <ma>
```

Hierarchy Level

```
[edit protocols protection-group ethernet-aps]
```

Description

This configuration statement is used to test if APS is operating correctly, it does not interrupt regular APS operations.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Ethernet Automatic Protection Switching Overview

exit

IN THIS SECTION

- Syntax | 6638
- Description | 6638
- Options | 6639
- Required Privilege Level | 6639
- Release Information | 6639

Syntax

```
exit <configuration-mode>
```

Description

Exit the current level of the statement hierarchy, returning to the level prior to the last `edit` command, or exit from configuration mode. The `quit` and `exit` commands are synonyms.

Options

none	Return to the previous edit level. If you are at the top of the statement hierarchy, exit configuration mode.
configuration-mode	(Optional) Exit from configuration mode.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Displaying the Current Configuration

exp (Class of Service)

IN THIS SECTION

- Syntax | 6640
- Hierarchy Level | 6640
- Description | 6641
- Options | 6641
- Required Privilege Level | 6641
- Release Information | 6641

Syntax

Classifier Configuration

```
exp classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [ aliases ] [ bit-patterns ];
        }
    }
}
```

Rewrite Rule Configuration

```
exp rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-point [ aliases ] [ bit-patterns ];
        }
    }
}
```

Global Classifier Association with Interfaces

```
exp name;
```

Hierarchy Level

```
[edit class-of-service classifiers]
[edit class-of-service code-point-aliases],
[edit class-of-service rewrite-rules]
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules],
[edit class-of-service system-defaults classifiers]
```

Description

Define the EXP code point mapping that is applied to MPLS packets. EXP classifiers are not applied to any traffic except MPLS traffic. EXP classifiers are applied only to interfaces that are configured as `family mpls` (for example, `set interfaces xe-0/0/35 unit 0 family mpls`.)

There are no default EXP classifiers. You can configure up to 64 EXP classifiers.

Starting in Junos OS 20.3R1, on segment routing LSPs, applying an EXP rewrite rule to the egress interface on the ingress (provider edge) router of a segment-routing network imposes the rewrite rule to all transport labels in the stack. As a result, you don't need to configure rewrite rules on every segment in the LSP.

On QFX10000 switches, you can configure and apply EXP classifiers to interfaces in the same way that you configure and apply DSCP, DSCP IPv6, and IEEE classifiers to interfaces. Different interfaces can have different EXP classifiers. QFX10000 switches do not support global EXP classifiers.

QFX5100 and EX4600 switches uses only one EXP classifier as a global MPLS classifier on all interfaces. You specify the global EXP classifier in the `[edit class-of-service system-defaults]` hierarchy.

EX Series switches support only one EXP code mapping on the switch (either default or custom). It is applied globally and implicitly to all the MPLS-enabled interfaces on the switch. You cannot bind it or disable it on individual interfaces.

Options

classifier-name—Name of the EXP classifier.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X50.

RELATED DOCUMENTATION

Configuring Rewrite Rules for MPLS EXP Classifiers

Understanding CoS MPLS EXP Classifiers and Rewrite Rules

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

Understanding Using CoS with MPLS Networks on EX Series Switches

Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect

Configuring CoS on Provider Switches of an MPLS Network

exp-push-push-push

IN THIS SECTION

- Syntax | 6642
- Hierarchy Level | 6642
- Description | 6642
- Options | 6643
- Required Privilege Level | 6643
- Release Information | 6643

Syntax

```
exp-push-push-push default;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

Rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining CoS of an incoming non-MPLS packet.

Options

`default`—Apply the default MPLS EXP rewrite table.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Rewriting the EXP Bits of All Three Labels of an Outgoing Packet

dscp (Rewrite Rules on Interfaces) | **6020**

dscp-ipv6 (Class of Service) | **6009**

exp (Rewrite rules) | **6643**

exp-swap-push-push | **6646**

ieee-802.1 (Rewrite Rules on Logical Interface) | **8360**

ieee-802.1ad | **8363**

inet-precedence (CoS Rewrite Rules) | **8582**

rewrite-rules (Definition) | **15746**

exp (Rewrite rules)

IN THIS SECTION

- Syntax | **6644**
- Hierarchy Level | **6644**
- Description | **6644**
- Options | **6644**

Required Privilege Level | 6645
Release Information | 6645

Syntax

```
exp (rewrite-name | default) protocol protocol-types;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

Apply an MPLS experimental (EXP) rewrite rule.

Starting in Junos OS 20.3R1, on segment routing LSPs, applying an EXP rewrite rule to the egress interface on the ingress (provider edge) router of a segment-routing network imposes the rewrite rule to all transport labels in the stack. As a result, you don't need to configure rewrite rules on every segment in the LSP.

Options

rewrite-name Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules exp]` hierarchy level.

default The default mapping.

By default, IP precedence rewrite rules alter the first three bits on the type-of-service (ToS) byte while leaving the last three bits unchanged. This default behavior applies to rewrite rules you configure for MPLS packets with IPv4 payloads. You configure these types of rewrite rules by including the `mpls-inet-both` or `mpls-inet-both-non-vpn` option at the `[edit class-of-service interfaces interface interface-name unit logical-unit-number rewrite-rules exp rewrite-rule-name protocol]` hierarchy level. The IP precedence rewrite rules explanation does not apply to ACX Series Universal Metro routers.

On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series 5G Universal Routing Platforms and EX Series switches, we highly recommend that you configure the `default` option when you configure a behavior aggregate (BA) classifier that does not include a specific rewrite rule for MPLS packets. Doing so ensures that MPLS exp value is rewritten according to the BA classifier rules configured for forwarding or packet loss priority. This does not apply to ACX Series Universal Metro routers.

protocol-types

Specify one or more protocol matching criteria:

- `mpls-any`—Apply to MPLS packets, write MPLS header only.
- `mpls-inet-both`—Apply to IPv4 MPLS packets, write MPLS and IPv4 header.
- `mpls-inet-both-non-vpn`—Apply to IPv4 MPLS packets, write MPLS and IPv4 header for only non VPN traffic.

> Junos OS Evolved supports only the `mpls-any` option. Devices running Junos OS Evolved do not support different rewrite rules for VPN and non-VPN traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for PTX platforms running Junos OS Evolved added with Junos OS Evolved Release 22.4.

RELATED DOCUMENTATION

Configuring Rewrite Rules

Rewriting the EXP Bits of All Three Labels of an Outgoing Packet

Applying Rewrite Rules to Output Logical Interfaces

protocol (Rewrite Rules) | **13457**

rewrite-rules (Definition) | **15746**

exp-swap-push-push

IN THIS SECTION

- Syntax | 6646
- Hierarchy Level | 6646
- Description | 6646
- Options | 6646
- Required Privilege Level | 6646
- Release Information | 6647

Syntax

```
exp-swap-push-push default;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

Rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining CoS of an incoming MPLS packet.

Options

`default`—Apply the default MPLS EXP rewrite table.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Rewriting the EXP Bits of All Three Labels of an Outgoing Packet

dscp (Rewrite Rules on Interfaces) | **6020**

dscp-ipv6 (Class of Service) | **6009**

exp (Rewrite rules) | **6643**

exp-push-push-push | **6642**

ieee-802.1 (Rewrite Rules on Logical Interface) | **8360**

ieee-802.1ad | **8363**

inet-precedence (CoS Rewrite Rules) | **8582**

rewrite-rules (Definition) | **15746**

expand-loose-hop

IN THIS SECTION

- Syntax | **6647**
- Hierarchy Level | **6648**
- Description | **6648**
- Required Privilege Level | **6648**
- Release Information | **6648**

Syntax

```
expand-loose-hop;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Allow an LSP to traverse multiple OSPF areas within a service provider's network.

Allows a point-to-multipoint LSP to span multiple domains in a network. Effectively, this allows you to configure one or more sub-LSPs (branches) in separate network domains. Examples of such domains include OSPF areas and autonomous systems (ASs). A sub-LSP of an inter-domain point-to-multipoint LSP can be intra-area, inter-area, or inter-AS, depending on the location of the egress node (leaf) with respect to the ingress node (source). Only OSPF areas are supported for inter-domain point-to-multipoint LSPs. IS-IS levels are not supported.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Point-to-multipoint LSP support introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Enabling Interarea Traffic Engineering

Configuring Inter-Domain Point-to-Multipoint LSPs

explicit-congestion-notification

IN THIS SECTION

- Syntax | 6649
- Hierarchy Level | 6649
- Description | 6649
- Required Privilege Level | 6650
- Release Information | 6650

Syntax

```
explicit-congestion-notification;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Enable explicit congestion notification (ECN) on the output queue (forwarding class) or output queues (forwarding classes) mapped to the scheduler. ECN enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality.

A weighted random early detection (WRED) packet drop profile must be applied to the output queues on which ECN is enabled. ECN uses the WRED drop profile thresholds to mark packets when the output queue experiences congestion.

ECN reduces packet loss by forwarding ECN-capable packets during periods of network congestion instead of dropping those packets. (TCP notifies the network about congestion by dropping packets.) During periods of congestion, ECN marks ECN-capable packets that egress from congested queues. When the receiver receives an ECN packet that is marked as experiencing congestion, the receiver

echoes the congestion state back to the sender. The sender then reduces its transmission rate to clear the congestion.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51.

Statement supported on SRX380, SRX300, SRX320, SRX340, SRX345, vSRX3.0, and NFX devices in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Example: Configuring ECN

Understanding CoS Explicit Congestion Notification

explicit-null-cos

IN THIS SECTION

- Syntax | 6651
- Hierarchy Level | 6651
- Description | 6651
- Options | 6651
- Required Privilege Level | 6651
- Release Information | 6651

Syntax

```
explicit-null-cos {
    inet;
    inet6;
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

The default classification for explicit-null packets is based on the payload (IPv4 or IPv6 DSCP bits). Configure this to make the packet classification based on the MPLS EXP value rather than on the payload, thus preserving the MPLS classification of the packet.

Options

NOTE: Include both options below or neither to use the MPLS EXP value for classification regardless of the payload type.

inet Use the MPLS EXP value for classification when the payload is an IPv4 packet.

inet6 Use the MPLS EXP value for classification when the payload is an IPv6 packet.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 18.1.

RELATED DOCUMENTATION

Applying MPLS EXP Classifiers for Explicit-Null Labels

explicit-null (Protocols BGP)

IN THIS SECTION

- Syntax | 6652
- Hierarchy Level | 6652
- Description | 6653
- Default | 6653
- Required Privilege Level | 6654
- Release Information | 6654

Syntax

```
explicit-null;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols bgp family inet labeled-unicast],
[edit logical-systems logical-system-name protocols bgp family inet6 labeled-unicast],
[edit logical-systems logical-system-name protocols bgp group group-name family inet labeled-
unicast],
[edit logical-systems logical-system-name protocols bgp group group-name family inet6 labeled-
unicast],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
inet labeled-unicast],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
inet6 labeled-unicast],
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp family
```

```
inet labeled-unicast],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp family
inet6 labeled-unicast],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name family inet labeled-unicast],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name family inet6 labeled-unicast],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name neighbor address family inet labeled-unicast],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name neighbor address family inet6 labeled-unicast],
[edit logical-systems logical-system-name routing-instances instance-name protocols ldp],
[edit protocols mpls],
[edit protocols bgp family inet labeled-unicast],
[edit protocols bgp family inet6 labeled-unicast],
[edit protocols bgp group group-name family inet labeled-unicast],
[edit protocols bgp group group-name family inet6 labeled-unicast],
[edit protocols bgp group group-name neighbor address family inet labeled-unicast]
 [edit protocols bgp group group-name neighbor address family inet6 labeled-unicast],
[edit protocols ldp],
[edit routing-instances instance-name protocols bgp family inet labeled-unicast],
[edit routing-instances instance-name protocols bgp family inet6 labeled-unicast],
[edit routing-instances instance-name protocols bgp group group-name family inet labeled-
unicast],
[edit routing-instances instance-name protocols bgp group group-name family inet6 labeled-
unicast],
[edit routing-instances instance-name protocols bgp group group-name neighbor address family
inet labeled-unicast],
[edit routing-instances instance-name protocols bgp group group-name neighbor address family
inet6 labeled-unicast],
[edit routing-instances instance-name protocols ldp]
```

Description

Advertise label 0 to the egress routing device of an LSP.

Default

If you do not include the `explicit-null` statement in the configuration, label 3 (implicit null) is advertised.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Advertising Explicit Null Labels to BGP Peers

explicit-null (Protocols LDP)

IN THIS SECTION

- Syntax | 6654
- Hierarchy Level | 6655
- Description | 6655
- Default | 6655
- Required Privilege Level | 6655
- Release Information | 6655

Syntax

```
explicit-null;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Advertise label 0 to the egress router of a label-switched path (LSP).

Default

If you do not include the `explicit-null` statement in the MPLS configuration, label 3 (implicit null) is advertised.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router

explicit-null (Protocols MPLS)

IN THIS SECTION

Syntax | 6656

- Hierarchy Level | 6656
- Description | 6656
- Default | 6656
- Required Privilege Level | 6656
- Release Information | 6657

Syntax

```
explicit-null;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Advertise label 0 to the egress router of an LSP.

Default

If you do not include the `explicit-null` statement in the MPLS configuration, label 3 (implicit null) is advertised.

> **NOTE**: Junos OS does not support explicit null routes with next hops to virtual tunnel (vt-) interfaces.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring RSVP to Pop the Label on the Ultimate-Hop Router

explicit-priority

IN THIS SECTION

- Syntax | 6657
- Hierarchy Level | 6657
- Description | 6658
- Required Privilege Level | 6658
- Release Information | 6658

Syntax

```
explicit-priority;
```

Hierarchy Level

```
[edit logical-systems logical-system-name system syslog  file filename],
[edit logical-systems logical-system-name system syslog host],
[edit system syslog file filename],
[edit system syslog host]
```

Description

Record the priority (facility and severity level) in each standard-format system log message directed to a file or remote destination.

When the `structured-data` statement is also included at the `[edit system syslog file filename]` hierarchy level, this statement is ignored for the file.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Including Priority Information in System Log Messages

System Log Explorer

structured-data | **28664**

export (ISIS)

IN THIS SECTION

- Syntax | **6659**
- Hierarchy Level | **6659**
- Description | **6659**
- Options | **6659**
- Required Privilege Level | **6659**
- Release Information | **6660**

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Apply one or more policies to routes being exported from the routing table into IS-IS.

All routing protocols store the routes that they learn in the routing table. The routing table uses this collected route information to determine the active routes to destinations. The routing table then installs the active routes into its forwarding table and exports them into the routing protocols. It is these exported routes that the protocols advertise.

For each protocol, you control which routes the protocol stores in the routing table and which routes the routing table exports into the protocol from the routing table by defining a *routing policy* for that protocol.

> **NOTE**: For IS-IS, you cannot apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol can easily lead to an inconsistent topology database.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

export (OSPF)

IN THIS SECTION

- Syntax | 6660
- Hierarchy Level | 6660
- Description | 6661
- Options | 6661
- Required Privilege Level | 6661
- Release Information | 6661

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Apply one or more policies to routes being exported from the routing table into OSPF.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Understanding OSPF Routing Policy

Import and Export Policies for Network Summaries Overview

import (Protocols OSPF) | **8463**

import (Protocols OSPF) | **8463**

export (RIP)

IN THIS SECTION

- Syntax | **6662**
- Hierarchy Level | **6662**
- Description | **6662**

- Options | 6662
- Required Privilege Level | 6663
- Release Information | 6663

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name],
[edit protocols rip group group-name],
[edit routing-instances routing-instance-name protocols rip group group-name]
```

Description

Apply a policy to routes being exported to the neighbors.

By default, RIP does not export routes it has learned to its neighbors. To enable RIP to export routes, apply one or more export policies.

If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the `metric-in` and `metric-out` statements.

> **NOTE**: The export policy on RIP does not support manipulating routing information of the next hop.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

import (Protocols RIP) | 8483

export (Bootstrap)

IN THIS SECTION

- Syntax | 6663
- Hierarchy Level | 6664
- Description | 6664
- Options | 6664
- Required Privilege Level | 6664
- Release Information | 6664

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp bootstrap family (inet | inet6)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp bootstrap family (inet | inet6)],
[edit protocols pim rp bootstrap family (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim rp bootstrap family (inet | inet6)]
```

Description

Apply one or more export policies to control outgoing PIM bootstrap messages.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring PIM Bootstrap Properties for IPv4

Configuring PIM Bootstrap Properties for IPv4 or IPv6

import (Protocols PIM Bootstrap) | **8482**

export-format (JSON)

IN THIS SECTION

- Syntax | 6665
- Hierarchy Level | 6665
- Description | 6665
- Options | 6666
- Required Privilege Level | 6666
- Release Information | 6666

Syntax

```
export-format {
    json {
        ietf;
        verbose;
    }
    state-data {
        json {
            compact;
        }
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Specify the default implementation of the serialization to use for exported data in the given format. This statement affects configuration data and command output that is displayed in the requested format.

Options

json Define which implementation of the serialization to use for configuration data emitted in JavaScript Object Notation (JSON) format.

Acceptable values include:

- `ietf`—Emit JSON data according to the encoding rules defined in Internet drafts draft-ietf-netmod-yang-json-09, *JSON Encoding of Data Modeled with YANG*, and draft-ietf-netmod-yang-metadata-06, *Defining and Using Metadata with YANG*.
- `verbose`—Emit JSON data in verbose format, which emits all objects as JSON arrays.

  You can configure the `verbose` statement starting in Junos OS Release 16.1R1, even though the statement is not exposed in the Junos OS CLI until a later release.
- **Default:** `ietf` in Junos OS Release 16.1R1 and later; `verbose` in earlier releases.

state-data Specify the format for emitting state data.

- `json compact`—Emit JSON-formatted state data in compact format where only the objects that have multiple values are emitted as JSON arrays. Junos devices, by default, emit state data in non-compact format, which emits all objects as JSON arrays.

  This statement is deprecated starting in Junos OS Release 24.4R1 and Junos OS Evolved Release 24.4R1

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`state-data json compact` statement added in Junos OS Release 17.3R1.

`state-data json compact` statement deprecated in Junos OS Release 24.4R1 and Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

Map Junos OS Configuration Statements to JSON

Map Junos OS Command Output to JSON in the CLI

export-format (Flow Monitoring)

IN THIS SECTION

- Syntax | 6667
- Hierarchy Level | 6667
- Description | 6667
- Options | 6667
- Required Privilege Level | 6668
- Release Information | 6668

Syntax

```
export-format format;
```

Hierarchy Level

```
[edit forwarding-options monitoring name output]
```

Description

Flow monitoring export format.

Options

format—Format of the flows.

- **Values:** 5 or 8
- **Default:** 5

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

version (Flow Aggregation) | 30692

Exporting Flows

Configuring Passive Flow Monitoring

export (MPLS Traffic engineering database)

IN THIS SECTION

- Syntax | 6668
- Hierarchy Level | 6669
- Description | 6669
- Options | 6669
- Required Privilege Level | 6669
- Release Information | 6669

Syntax

```
export {
    credibility;
    policy policy-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls traffic-engineering database],
[edit protocols mpls traffic-engineering database]
```

Description

Configure the traffic engineering database export-related parameters.

Options

policy *policy-name*	Name of the export policy.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

traffic-engineering (Protocols MPLS) | **30031**

export-profile (Junos Telemetry Interface)

IN THIS SECTION

- Syntax | **6670**
- Hierarchy Level | **6670**

- Description | 6670
- Options | 6671
- Required Privilege Level | 6673
- Release Information | 6673

Syntax

```
export-profile name {
    dscp value;
    format file-format;
    forwarding-class (assured-forwarding |  best-effort |  expedited-forwarding  | network-
control);
    local-address ip-address;
    local-port source-port-number;
    loss-priority (high | low | medium-high | medium-low);
    <payload-size bytes>;
    reporting-rate seconds;
    transport protocol-name;
    routing-instance routing-instance-name
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Configure the parameters of the export process for data generated through Junos Telemetry Interface sensors. You can create one or more export profiles. Each profile can be associated with one or more sensors that define the system resource to monitor and stream data. You can associate only one export profile with a specific sensor configuration.

The IP layer delivers the exported data to the remote server. The export profile configuration allows you to specify a format for exported data, a transport protocol, the rate which the system generates data, and the local source port and IP address that are used to define the transport headers in the exported packets.

To enable Junos Telemetry Interface, you must also configure a sensor that defines the parameters of the system resource to monitor and stream data, and a server to collect the data. To configure a sensor, include the `sensor` *`sensor-name`* statement at the `[edit services analytics]` hierarchy level. To configure the server that functions as a data collector, include `streaming-server` *`server-name`* statement at the `[edit services analytics]` hierarchy level.

If you are configuring the export profile for dialout, the options `local-address`, `local-port`, `dscp`, `forwarding-class`, `loss-priority`, `payload-size`, and `resource-filter` are not supported. Only the `reporting-interval` option is supported with dialout.

> **NOTE**: Junos Telemetry Interface was introduced in Junos OS Release 15.1F3 on MX Series routers with interfaces configured on MPC1 through MPC6E and on PTX Series routers with interfaces configured on FPC3. Starting in Junos OS Release 15.1F5, Junos Telemetry Interface is also supported on MPC7E, MPC8E, and MPC9E on MX Series routers.
>
> Starting with Junos OS Release 16.1R3, FPC1 and FPC2 on PTX Series routers are also supported.
>
> Starting with Junos OS Release 17.2R1, QFX10000 switches and PTX1000 routers are also supported.

Options

name

Name of export profile.

> **NOTE**: To associate this export profile with a configured sensor, include the `name` you configure for the `export-profile` statement at the `[edit services analytics sensor` *`sensor-name`* `export-name]` hierarchy level.

dscp ***value***

Specify the DSCP value for the exported packets.

- **Range:** 0 through 63.
- **Default:** 0

> **NOTE**: Any interface-level DSCP rewrite rules you have configured override the DSCP value you specify for the export profile. You need to specify a DSCP value for the export profile only if you do not configure DSCP rewrite rules on the outgoing interface. For more information, see *Configuring Rewrite Rules*.

format

Specify the format to define the structure of exported data.

gpb Google protocol buffers format.

gpb-gnmi Google protocol buffers encoding support for (gRPC Network Management Interface) gNMI.

json-gnmi Javascript Object Notation (JSON) encoding support for gNMI.

forwarding-class (assured-forwarding | best-effort | expedited-forwarding | network-control)

(Packet Forwarding Engine sensors only) Specify the forwarding class for exported packets.

- **Default:** `best-effort`

loss-priority (high | low | medium-high | medium-low) (MX Series only)

Specify the loss priority for exported packets. Loss priority settings help determine which packets are dropped from the network during periods of congestion.

local-address *ip-address*

Specify the source address of exported packets.

local-port *number*

Specify the source port for the exported packets.

payload-size *bytes* (Optional)

Specify the maximum size of exported packets.

> **NOTE**: The `payload-size` option is supported only on the following sensors:
>
> - /junos/system/linecard/interface/
> - /junos/system/linecard/interface/logical/usage/
> - /junos/system/linecard/firewall/

- **Default:** 5000 bytes.
- **Range:** For the QFX5100 line of switches, 3000 through 9192 bytes. For Packet Forwarding Engine-related statistics (logical and physical interface and firewall statistics) on all platforms, 2000 through 9192. This is a platform limitation. For all other supported platforms, 1400 through 9192 bytes.

> **NOTE**: Junos Telemetry Interface does not export packets larger than 9192 bytes.

reporting-rate *seconds*

Specify the interval at which the Junos Telemetry Interface sensor generates data to export to the collector.

As the configured interval expires, the most recent sample collected by the sensor is gathered and forwarded to the server configured to collect data.

> **NOTE**: For Packet Forwarding Engine sensors, the minimum reporting rate is 2 seconds.

- **Range:** 1 through 3600 (1 hour)

transport *protocol-name*

Specify the transport protocol to use to carry the telemetry data in the IP packets.

udp User Datagram Protocol.

routing-instance *rounting-instance-name*

Specify the routing instance to be used for grpc transport.
The `routing-instance` configuration is for grpc dialout connections only. Configuring `routing-instance` is not mandatory. If it is not configured, the default routing instance is applied. If you specify a `routing-instance` which is not configured, commit fails and an error message is displayed. If `routing-instance` configuration is specified for UDP transport protocol, commit fails and an error message is displayed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F3.

`payload-size` *bytes* option introduced in Junos OS Release 16.1R3.

`loss-priority` option introduced in Junos OS Release 17.3R1 for MX Series routers only.

Statement introduced in Junos OS Release 17.4R1 for virtual MX Series (vMX) routers.

RELATED DOCUMENTATION

sensor (Junos Telemetry Interface) | 16564

export-profiles

IN THIS SECTION

- Syntax | 6674
- Hierarchy Level | 6675
- Description | 6675
- Options | 6675
- Required Privilege Level | 6676
- Release Information | 6676

Syntax

```
export-profiles {
    profile-name {
        interface {
            information;
            statistics {
                queue;
                traffic;
            }
            status {
                link;
                queue;
                traffic;
            }
        }
        stream-format format;
        system {
            information;
            status {
                queue;
```

```
                traffic;
            }
        }
    }
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Configure an profile to specify the network analytics data being streamed to remote servers. Each profile is a template that defines the type of data being streamed.

> **NOTE**: The `export-profile` statement is not available in Junos OS Releases prior to 13.2X51-D15.

Options

profile-name Name of the export profile containing the configuration of the data being streamed.

stream-format ***format*** Format of the streaming data being sent to a server. Only one format can be sent to each port on a server.

- Values:
  - **csv**—Comma-separated Values (CSV). Data sent in this format is newline separated, and each record contains one stream type (queue or traffic data) per interface.Each record contains either a “q” for a queue statistics, or a “t” for a traffic statistics.
  - **gpb**—Google Protocol Buffer (GPB). Data sent in this format has a hierachical format, and is categorized by resource type (system or interfaces), which is specified in the message header. You can generate data formatted in other formats (CSV, TSV, and JSON) from GPB-encoded data.

    Each message includes a 8-byte header containing the following information:

- Bytes 0 to 3—Length of the message.
- Byte 4—Message version.
- Bytes 5 to 7—Reserved for future use.

> **NOTE**: A schema file called **analytics.proto** containing the definitions of the GPB messages is available for downloading from the following location:
> https://www.juniper.net/documentation/en_US/junos13.2/topics/reference/proto-files/analytics-proto.txt

- **json**—JavaScript Object Notation (JSON). Data sent in this format is newline separated, and each record contains one stream type (queue or traffic data) per interface. Each record contains either “queue-statistics” or “traffic-statistics” in the “record type” field.
- **tsv**—Tab-separated Values (TSV). Data sent in this format is newline separated, and each record contains one stream type (queue or traffic data) per interface. Each record contains a “q” for a queue statistics, or a “t” for a traffic statistics.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

export (Protocols BGP)

IN THIS SECTION

- Syntax | 6677
- Hierarchy Level | 6677
- Description | 6678
- Options | 6678
- Required Privilege Level | 6678
- Release Information | 6678

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group name optimal-route-reflection],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Apply one or more policies to routes being exported from the routing table into BGP.

If you specify more than one policy, they are evaluated in the order specified, from left to right, and the first matching filter is applied to the route. If no routes match the filters, the routing table exports into BGP only the routes that it learned from BGP. If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during the evaluation of the remaining policies. For example, if the action specified in the first policy of a chain sets a route's metric to 500, this route matches the criterion of `metric 500` defined in the next policy.

Options

`policy-names`—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routing Policies to Control BGP Route Advertisements

Routing Policies, Firewall Filters, and Traffic Policers User Guide

import (BGP) | **8461**

export (Protocols DVMRP)

IN THIS SECTION

Syntax | **6679**

- Hierarchy Level | 6679
- Description | 6679
- Options | 6679
- Required Privilege Level | 6679
- Release Information | 6680

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols dvmrp],
[edit protocols dvmrp]
```

Description

Apply one or more policies to routes being exported from the routing table into DVMRP. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found, the routing table exports into DVMRP only the routes that it learned from DVMRP and direct routes.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

> **NOTE**: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

import (Protocols DVMRP) | 8474

Example: Configuring DVMRP to Announce Unicast Routes

export (Protocols LDP)

IN THIS SECTION

- Syntax | 6680
- Hierarchy Level | 6681
- Description | 6681
- Options | 6681
- Required Privilege Level | 6681
- Release Information | 6681

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.

Options

policy-names—Name of one or more routing policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Filtering Outbound LDP Label Bindings

export (Protocols MSDP)

IN THIS SECTION

Syntax | 6682

- Hierarchy Level | 6682
- Description | 6683
- Options | 6683
- Required Privilege Level | 6683
- Release Information | 6683

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name],
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
peer address],
[edit protocols msdp],
[edit protocols msdp group group-name],
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit routing-instances routing-instance-name protocols msdp],
[edit routing-instances routing-instance-name protocols msdp group group-name],
[edit routing-instances routing-instance-name protocols msdp group group-name peer address],
[edit routing-instances routing-instance-name protocols msdp peer address]
```

Description

Apply one or more policies to routes being exported from the routing table into MSDP.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP in a Routing Instance

import (Protocols MSDP) | **8478**

export (Protocols PIM)

IN THIS SECTION

- Syntax | **6684**
- Hierarchy Level | **6684**
- Description | **6684**
- Required Privilege Level | **6684**
- Release Information | **6684**

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Apply one or more export policies to control outgoing PIM join and prune messages. PIM join and prune filters can be applied to PIM-SM and PIM-SSM messages. PIM join and prune filters cannot be applied to PIM-DM messages.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Filtering Outgoing PIM Join Messages

export (Protocols RIFT)

IN THIS SECTION

- Syntax | 6685
- Hierarchy Level | 6685
- Description | 6685
- Options | 6686
- Required Privilege Level | 6686
- Release Information | 6686

Syntax

```
export {
    northbound {
        policy;
    }
    southbound {
        policy;
        allow-rift-routes;
    }
}
```

Hierarchy Level

```
[edit protocols rift]
```

Description

Enable exporting of routes from other protocols into the RIFT protocol.

Options

northbound (Optional) Policy to advertise routes from other protocols using north external prefix topology information elements (TIEs).

southbound (Optional) Policy to advertise routes from other protocols using south external prefix TIEs.

policy Name of the policy to be used for northbound or southbound exporting of routes into the RIFT protocol.

allow-rift-routes (Optional) (Southbound Only) Allow calculated RIFT northbound routes in redistribution.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

export (Protocols RIPng)

IN THIS SECTION

- Syntax | 6687
- Hierarchy Level | 6687
- Description | 6687

- Options | 6687
- Required Privilege Level | 6687
- Release Information | 6688

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng group group-name],
[edit protocols ripng group group-name],
[edit routing-instances routing-instance-name protocols ripng group group-name]
```

Description

Apply a policy or list of policies to routes being exported to the neighbors.

By default, RIPng does not export routes it has learned to its neighbors. To have RIPng export routes, apply one or more export policies. To apply export policies and to filter routes being exported from the local routing device to its neighbors, include the `export` statement and list the name of the policy to be evaluated.

You can define one or more export policies. If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the `metric-in` and `metric-out` statements.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

import (Protocols RIPng) | 8485

export-rib

IN THIS SECTION

- Syntax | 6688
- Hierarchy Level | 6688
- Description | 6689
- Options | 6690
- Required Privilege Level | 6691
- Release Information | 6691

Syntax

```
export-rib routing-table-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib-groups group-name],
[edit logical-systems logical-system-name routing-options rib-groups group-name],
```

```
[edit routing-instances routing-instance-name routing-options rib-groups group-name],
[edit routing-options rib-groups group-name]
```

Description

Specify the name of the routing table from which Junos OS should export routing information. For any individual RIB group, only one table can be specified in the `export-rib` statement.

The `export-rib` statement specifies the source table from which routing information is advertised.

One common use of the `export-rib` statement is interdomain routing. The export RIB is the table used when BGP extracts routes to advertise to peers. In multicast interdomain routing, for example, the export RIB is likely to be inet.2.

Another use of `export-rib` is dynamic route leaking between the global routing table (inet.0) and a VRF routing table (*instance*.inet.0). For example, you can use a RIB group to copy routes learned in the VRF into the global routing table, inet.0, or copy routes learned in inet.0 into a VRF. You define the use of this RIB group in the VRF's BGP configuration. In a routing policy you can do dynamic filtering of routes. For instance, you can use an import policy to only copy routes with certain communities into the global routing table.

For example:

```
rib-groups {
    rib-interface-routes-v4 {
        import-rib [ inet.0 VRF.inet.0 ];
    }
    rib-import-VRF-routes-to-inet0-v4 {
        export-rib VRF.inet.0;
        import-rib [ VRF.inet.0 inet.0 ];
        import-policy rib-import-VRF-routes-to-inet0-v4;
    }
    rib-import-inet0-routes-to-VRF-v4 {
        export-rib inet.0;
        import-rib [ inet.0 VRF.inet.0 ];
        import-policy rib-import-inet0-routes-to-VRF-v4;
    }
}
routing-options {
    interface-routes {
        rib-group {
            inet rib-interface-routes-v4;
        }
```

```
    }
}
protocols {
    bgp {
        group iBGP-peers {
            type internal;
            family inet {
                unicast {
                    rib-group rib-import-inet0-routes-to-VRF-v4;
                }
            }
        }
    }
}
routing-instances {
    VRF {
        routing-options {
            interface-routes {
                rib-group {
                    inet rib-interface-routes-v4;
                }
            }
        }
        protocols {
            bgp {
                group peersin-VRF {
                    family inet {
                        unicast {
                            rib-group rib-import-VRF-routes-to-inet0-v4;
                        }
                    }
                }
            }
        }
    }
}
```

Options

routing-table-name—Routing table group name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table

Example: Configuring a PIM RPF Routing Table

Example: Configuring DVMRP to Announce Unicast Routes

Example: Configuring a Dedicated PIM RPF Routing Table

Example: Configuring Any-Source Multicast for Draft-Rosen VPNs

import-rib | **8488**

passive (Routing Options) | **12225**

export (Routing Options)

IN THIS SECTION

- Syntax | **6692**
- Hierarchy Level | **6692**
- Description | **6692**
- Options | **6692**
- Required Privilege Level | **6692**
- Release Information | **6693**

Syntax

```
export [ policy-name ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options forwarding-table],
[edit logical-systems logical-system-name routing-options forwarding-table],
[edit routing-instances routing-instance-name routing-options forwarding-table],
[edit routing-options forwarding-table]
```

Description

Apply one or more policies to routes being exported from the routing table into the forwarding table.

In the `export` statement, list the name of the routing policy to be evaluated when routes are being exported from the routing table into the forwarding table. Only active routes are exported from the routing table.

You can reference the same routing policy one or more times in the same or a different `export` statement.

You can apply export policies to routes being exported from the routing table into the forwarding table for the following features:

- Per-packet load balancing
- Class of service (CoS)

Options

policy-name—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Load Balancing BGP Traffic

export-target

IN THIS SECTION

- Syntax | 6693
- Hierarchy Level | 6693
- Description | 6694
- Options | 6694
- Required Privilege Level | 6694
- Release Information | 6694

Syntax

```
export-target {
    target target-community;
    unicast;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvpn
route-target],
[edit routing-instances routing-instance-name protocols mvpn route-target]
```

Description

Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN network layer reachability information (NLRI).

Options

`target` *`target-community`*—Specify the export target community.

`unicast`—Use the same target community as specified for unicast.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

expression

IN THIS SECTION

- Syntax | 6695
- Hierarchy Level | 6695
- Description | 6695
- Options | 6695
- Required Privilege Level | 6695
- Release Information | 6695

Syntax

```
expression boolean-expression;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain]
```

Description

Configure the Boolean expression. The Boolean expression defines the condition for the individual members of a chain attack that will decide if the chain attack is hit.

For standalone IDP devices, expression overrides order function.

For SRX Series Firewalls, expression and order cannot be configured together. Only one of them can be specified.

Options

boolean-expression—Boolean operators:

- `or`—If either of the member name patterns match, the expression matches.
- `and`—If both of the member name patterns match, the expression matches. It does not matter which order the members appear in.
- `oand`—If both of the member name patterns match, and if they appear in the same order as in the Boolean Expression, the expression matches.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

extend-secondary-vlan-id

IN THIS SECTION

- Syntax | 6696
- Hierarchy Level | 6696
- Description | 6696
- Required Privilege Level | 6696
- Release Information | 6696

Syntax

```
extend-secondary-vlan-id number;
```

Hierarchy Level

```
[edit vlans vlan-name pvlan]
```

Description

Configure traffic that egresses from a secondary VLAN trunk port to retain its secondary VLAN tag instead of getting the tag of the primary VLAN that the secondary port is a member of.

Required Privilege Level

routing—To view this statement in the configuration.

routing–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring PVLANs with Secondary VLAN Trunk Ports and Promiscuous Access Ports on a QFX Series Switch

extended-echo

IN THIS SECTION

- Syntax | 6697
- Hierarchy Level | 6697
- Description | 6697
- Default | 6698
- Required Privilege Level | 6698
- Release Information | 6698

Syntax

```
extended-echo;
```

Hierarchy Level

```
[edit system]
```

Description

Configure ICMP extended echo to enable the `probe` command. On enabling this statement on the proxy node, you can query for the status of the probed interface. The ICMP extended echo and ICMP extended reply messages are used for the communication between the probing interface and the probed interface through the proxy interface.

Default

Extended echo is disabled unless explicitly enabled.

Required Privilege Level

system-To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

How to Use the Probe Command

probe (ICMP Router Discovery) | 13222

extended-isid-list

IN THIS SECTION

- Syntax | 6698
- Hierarchy Level | 6699
- Description | 6699
- Options | 6699
- Required Privilege Level | 6699
- Release Information | 6699

Syntax

```
extended-isid-list (single-isid | isid-list |isid-range | all);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols]
[edit routing-instances routing-instance-name protocol evpn]
```

Description

Specify the service identifiers (ISIDs) that are being extended over the Ethernet VPN (EVPN) core.

> **NOTE**: When you deactivate the `extended-isid-list` statement, with or without configuring the automatic derivation of the BGP route target (auto-RT) for advertised prefixes, both Type 2 and Type 3 routes are retained.

Options

single-isid Specify a single ISID.

isid-list Specify a list of multiple ISIDs.

isid-range Specify a range of ISIDs.

all Specify that all ISIDs are being extended over the EVPN core.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

evpn | 6541

extended-nexthop

IN THIS SECTION

- Syntax | 6700
- Hierarchy Level | 6700
- Description | 6701
- Required Privilege Level | 6701
- Release Information | 6701

Syntax

```
extended-nexthop;
```

Hierarchy Level

```
[edit logical-systems name protocols bgp family inet unicast],
[edit logical-systems name protocols bgp family inet inet6 unicast],
[edit logical-systems name protocols bgp group name family inet unicast],
[edit logical-systems name protocols bgp group name family inet inet6 unicast],
[edit logical-systems name protocols bgp group name neighbor name family inet unicast],
[edit logical-systems name protocols bgp group name neighbor name family inet inet6 unicast],
[edit logical-systems name routing-instances name protocols bgp family inet unicast],
[edit logical-systems name routing-instances name protocols bgp family inet inet6 unicast],
[edit logical-systems name routing-instances name protocols bgp group name family inet unicast],
[edit logical-systems name routing-instances name protocols bgp group name family inet inet6
unicast],
[edit logical-systems name routing-instances name protocols bgp group name neighbor name family inet
unicast],
[edit logical-systems name routing-instances name protocols bgp group name neighbor name family inet
inet6 unicast],
[edit protocols bgp family inet unicast],
[edit protocols bgp family inet inet6 unicast],
[edit protocols bgp group name family inet unicast],
[edit protocols bgp group name family inet inet6 unicast],
```

```
[edit protocols bgp group name neighbor name family inet unicast],
[edit protocols bgp group name neighbor name family inet inet6 unicast],
[edit routing-instances name protocols bgp family inet unicast],
[edit routing-instances name protocols bgp family inet inet6 unicast],
[edit routing-instances name protocols bgp group name family inet unicast],
[edit routing-instances name protocols bgp group name family inet inet6 unicast],
[edit routing-instances name protocols bgp group name neighbor name family inet unicast],
[edit routing-instances name protocols bgp group name neighbor name family inet inet6 unicast]
```

Description

Configure extended next-hop encoding for BGP groups with IPv6 peers to route IPv4 address families over an IPv6 session. Configure dynamic IPv4-over-IPv6 tunnels and define their attributes to forward IPv4 traffic over an IPv6-only network.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

dynamic-tunnels | 6161

tunnel-attributes | 30194

Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP

extended-nexthop-tunnel

IN THIS SECTION

Syntax | 6702

- Hierarchy Level | 6702
- Description | 6702
- Required Privilege Level | 6702
- Release Information | 6702

Syntax

```
extended-nexthop-tunnel;
```

Hierarchy Level

```
[edit protocols bgp group group-name family inet unicast],
[edit protocols bgp group group-name family inet6 unicast]
```

Description

Configure the tunnel attribute to transition the BGP signaling of MPLS-over-UDP dynamic nexthop tunnels to IP-over-IP tunnels.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Overview of Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation

extended-space

IN THIS SECTION

- Syntax | 6703
- Hierarchy Level | 6703
- Description | 6703
- Default | 6704
- Required Privilege Level | 6704
- Release Information | 6704

Syntax

```
extended-space;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-
next-hop ingress l3vpn],
[edit routing-options forwarding-table chained-composite-next-hop ingress l3vpn]
```

Description

Accept more than one million Layer 3 VPN BGP updates with unique inner VPN labels. The neighboring PE routers are typically from other vendors and configured to assign a unique inner label to each Layer 3 VPN BGP route.

NOTE: This statement is supported on Juniper Networks routers containing only MPCs and MICs and with chassis network services configured with the `enhanced-ip` setting.

> **WARNING**: You cannot configure the `extended-space` statement on a MX system with a MS-DPC unless you have also configured FIB localization with the MS-DPC set as a FIB-remote. This is because the MS-DPC cannot install nexthop in `extended-space` so the use of a default nexthop pointing to a FIB-local FPC is necessary to ensure proper data forwarding.

Default

This statement is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs

extended-statistics

IN THIS SECTION

- Syntax | 6705
- Hierarchy Level | 6705
- Description | 6705
- Default | 6705
- Required Privilege Level | 6705
- Release Information | 6705

Syntax

```
extended-statistics;
```

Hierarchy Level

```
[edit chassis]
```

Description

(MX Series routers only) Enable accounting of system statistics for IPv4 and IPv6 traffic.

Default

Accounting of system statistics is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

show system statistics

extended-system-id

IN THIS SECTION

- Syntax | 6706
- Hierarchy Level | 6706
- Description | 6706
- Options | 6706
- Required Privilege Level | 6707
- Release Information | 6707

Syntax

```
extended-system-id identifier;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rstp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
rstp],
[edit protocols rstp],
[edit routing-instances routing-instance-name protocols rstp]
```

Description

The extended system ID is used to specify different bridge identifiers for different RSTP or STP routing instances.

Options

identifier—Specify the system identifier to use for the RSTP or STP instance.

- **Range:** 0 through 4095

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Rapid Spanning Tree Protocol

Configuring RSTP on EX Series Switches (CLI Procedure)

Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP

Understanding RSTP

extended-vlan-list

IN THIS SECTION

- Syntax | 6707
- Hierarchy Level | 6708
- Description | 6708
- Options | 6708
- Required Privilege Level | 6708
- Release Information | 6708

Syntax

```
extended-vlan-list vlan-id | [vlan-id set];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols]
[edit routing-instances routing-instance-name instance-type virtual-switch protocols evpn]
```

Description

Specify the VLAN or range of VLANs that are extended over the WAN, wherein all the single VLAN bridge domains corresponding to these VLANs are stretched.

> **NOTE**: The `extended-vni-list` statement is an exclusive command. You cannot configure the `extended-vni-list` statement with either the `extended-vlan-list` or `extended-vni-all` statements.

Options

vlan-id VLAN ID to be EVPN extended.

vlan-id set List of VLAN IDs to be EVPN extended.

- **Range:** 1 through 4094 VLANs

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

evpn | 6541

extended-vni-all

IN THIS SECTION

- Syntax | 6709
- Hierarchy Level | 6709
- Description | 6709
- Required Privilege Level | 6710
- Release Information | 6710

Syntax

```
extended-vni-all;
```

Hierarchy Level

For QFX5100 and EX4600 switches:

```
{edit protocols evpn]
```

For MX Series routers:

```
[edit routing-instances routing-instance-name protocols evpn]
```

Description

Include all VXLAN Network Identifiers (VNIs) as part of the Virtual Switch (VS) instance. By specifying `all`, you bypass the `commit check` process and all configured bridge domains (BDs) in the Ethernet Virtual Private Network (EVPN) are considered extended.

> **NOTE**: The `extended-vni-list` statement is an exclusive command. You cannot configure the `extended-vni-list` statement with either the `extended-vlan-list` or `extended-vni-all` statements.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D30.

RELATED DOCUMENTATION

Example: Configuring an EVPN Control Plane and VXLAN Data Plane

Understanding EVPN with VXLAN Data Plane Encapsulation

EVPN Over VXLAN Encapsulation Configuration Overview for QFX Series and EX4600 Switches

extended-vni-list | **6710**

extended-vni-list

IN THIS SECTION

- Syntax | **6711**
- Hierarchy Level | **6711**
- Description | **6711**
- Options | **6712**
- Required Privilege Level | **6712**
- Release Information | **6712**

Syntax

```
extended-vni-list [list of VNIs / all];
```

Hierarchy Level

For MX Series routers, EX9200, and QFX Series switches:

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit routing-instances routing-instance-name protocols evpn]
```

Description

Establishes which VXLAN Network Identifiers (VNI) will be part of the Virtual Switch (VS) instance. When you issue the `commit check` command to verify the candidate configuration syntax without committing it, it also checks if the specified VNI(s) is associated with a bridge domain (BD) (`bridge-domain` *`name`* `vxlan vni`).

There are different broadcast, unknown unicast, and multicast (BUM) replication options available in Ethernet Virtual Private Network (EVPN). By using the `extended-vni-list` statement, you forgo a multicast underlay in favor of EVPN and VXLAN ingress-replication.

NOTE: The `extended-vni-list` statement is an exclusive command. You cannot configure the `extended-vni-list` statement with either the `extended-vlan-list` or `extended-vni-all` statements.

NOTE: If you configure the `extended-vni-list` statement, then you must also configure the `encapsulation (protocols evpn)` `vxlan` statement.

NOTE: The `extended-vni-list` statement is optional for routing-instances of type `mac-vrf`. In the mac-vrf type routing instance, all VNI are extended by default.

If you configure an `extended-vni-list` within a specific `mac-vrf` type routing instance, you negate the default behavior for that routing instance.

Options

list of VNIs Specify a single VNI or list of VNIs as part of the VS instance, for example `extended-vni-list [ 10-50 60 70]`.

all Include all VNIs as part of the VS instance. By specifying `all`, you bypass the `commit check` process and all configured BDs in the EVPN are considered extended.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D15.

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

- *Example: Configuring an EVPN Control Plane and VXLAN Data Plane*
- *Understanding EVPN with VXLAN Data Plane Encapsulation*
- *EVPN Over VXLAN Encapsulation Configuration Overview for QFX Series and EX4600 Switches*
- show configuration protocols evpn
- *Implementing EVPN-VXLAN for Data Centers*

Extensible Subscriber Services

IN THIS SECTION

- Syntax | 6713
- Hierarchy Level | 6713
- Description | 6713

- Options | 6713
- Required Privilege Level | 6714
- Release Information | 6714

Syntax

```
extensible-subscribe-services {
     access-profile access-profile-name;
     commit-interval interval;
     dictionary dictionary-path;
     flat-file-profile flat-file-profile;
     logical-interface-init-range {
          high;
          low;
     }
     maximum-subscribers limit;
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure the extensible subscriber services manager. Extensible Subscriber Services Manager (ESSM) is a background process that is part of the Intelligent Customer Extendable authentication, authorization, and accounting (ICE-AAA) framework, which supports customer extensible services for both business and residential subscribers.

Options

access-profile Access profile name for time-accounting specified using this configuration. In most cases access-profile radius server information is same as control session access-profile. It is a mandatory configuration.

commit-interval	The ESSM daemon issues commit request at the interval specified in this configuration.
dictionary	XML dictionary file is specified using this configuration command. It is a complete path with dictionary file name. The daemon takes action for the extensible-service request based on the XML description in the file. It is a mandatory configuration.
flat-file-profile	Refer a flat-file accounting profile that defines the contents of a flat file that records accounting statistics collected from the Packet Forwarding Engine for an interface at regular intervals.
logical-interface-init-range	Logical interface unit range. The daemon assigns the first available unit number in the range specified. Default low is 1 and high is 16385.
maximum-subscribers	Number of control subscriber sessions supported can be limited by using this configuration command. Default maximum is 1000.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R1.

extension-header

IN THIS SECTION

- Syntax | 6715
- Hierarchy Level | 6715
- Description | 6715
- Required Privilege Level | 6715
- Release Information | 6715

Syntax

```
extension-header {
    destination-option {
        home-address {
            match (equal | greater-than | less-than | not-equal);
            value header-value;
        }
            option-type {
                match (equal | greater-than | less-than | not-equal);
                value header-value;
            }
    }
    routing-header {
        header-type {
            match (equal | greater-than | less-than | not-equal);
            value header-value;
            }
        }
}
```

Hierarchy Level

```
[edit  set security idp custom-attack attack-name attack-type signature protocol ipv6]
```

Description

Specify the IPv6 extension header.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

extension-provider

IN THIS SECTION

- Syntax | 6716
- Hierarchy Level | 6717
- Description | 6717
- Required Privilege Level | 6717
- Release Information | 6717

Syntax

```
extension-provider {
    control-cores control-number;
    data-cores data-number;
    data-flow-affinity {
        hash-key (layer-3 | layer-4);
    }
    forwarding-db-size size;
    object-cache-size size;
    package package-name;
    policy-db-size size;
    syslog {
        facility {
            severity;
            destination destination;
        }
    }
    wired-max-processes num-procs;
    wired-process-mem-size mem-size;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package]
```

Description

Configure an application on a PIC.

> **CAUTION**: Do not configure these settings unless it is specified you should do so. The default settings work under most normal scenarios. Unneeded settings can cause negative effects.

> **NOTE**: The PICs of MS-MPC and MS-MIC support only `extension-provider` service-package.

When the `extension-provider` statement is first configured, the PIC reboots.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

extension-service (System Extensions)

IN THIS SECTION

- Syntax | 6718
- Hierarchy Level | 6719

- Description | 6719
- Required Privilege Level | 6719
- Release Information | 6719

Syntax

```
extension-service {
    application {
        file script-name {
            arguments arguments;
            cert-file path;
            checksum (md5 | sha-256 | sha1) hash;
            daemonize;
            interpreter (bash | python | python3);
            optional;
            refresh;
            refresh-from;
            respawn-on-normal-exit;
            routing-instance
            source;
            start {
                start-options {
                    mem-factor percentage;
                }
            }
            traceoptions {
                file <filename> <files number> <size size> <world-readable | no-world-readable>;
                flag flag;
                no-remote-trace;
            }
            username username;
        }
        max-datasize max-datasize;
        start {
            start-options {
                mem-factor percentage;
            }
        }
```

```
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
}
```

Hierarchy Level

```
[edit system extensions]
```

Description

Enable the Junos OS and Junos OS Evolved extension services.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

extensions | 6722

extension-service (System Services)

IN THIS SECTION

Syntax (Junos OS) | 6720

- Syntax (Junos OS Evolved) | 6721
- Hierarchy Level | 6721
- Description | 6722
- Required Privilege Level | 6722
- Release Information | 6722

Syntax (Junos OS)

```
extension-service {
    request-response {
        grpc {
            max-connections max-connections;
            routing-instance routing-instance;
            ssl {
                address ip-address;
                hot-reloading;
                local-certificate local-certificate;
                mutual-authentication {
                    certificate-authority certificate-authority-profile-name;
                    client-certificate-request (no-certificate | request-certificate | request-
certificate-and-verify | require-certificate | require-certificate-and-verify);
                }
                port port;
            }
        }
    }
    notification {
        allow-clients {
            address ip-address;
        }
        broker-socket-send-buffer-size broker-socket-send-buffer-size;
        max-connections max-connections;
        port port;
    }
    traceoptions {
        file <filename> <files number> <match regex> <size size> <world-readable | no-world-
readable>;
        flag flag;
```

```
        level <error>;
        no-remote-trace;
    }
}
```

Syntax (Junos OS Evolved)

```
extension-service {
    request-response {
        grpc {
            max-connections max-connections;
            routing-instance routing-instance;
            ssl {
                address ip-address;
                hot-reloading;
                local-certificate local-certificate;
                mutual-authentication {
                    certificate-authority certificate-authority-profile-name;
                    client-certificate-request (no-certificate | request-certificate | request-
certificate-and-verify | require-certificate | require-certificate-and-verify);
                }
                port port;
                use-pki;
            }
        }
    }
    traceoptions {
        file <filename> <files number> <match regex> <size size> <world-readable | no-world-
readable>;
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Enable Junos OS extension services.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`grpc` option introduced in Junos OS Release 16.2.

extensions

IN THIS SECTION

- Syntax | 6722
- Hierarchy Level | 6724
- Description | 6724
- Required Privilege Level | 6724
- Release Information | 6724

Syntax

```
extensions {
    extension-service {
        application {
            file script-name {
                arguments arguments;
                cert-file path;
```

```
                checksum hash-algorithm hash-value;
                daemonize;
                interpreter (bash | python | python3);
                optional;                 refresh;
                refresh-from;
                routing-instance;
                source;
                start {
                    start-options {
                        mem-factor percentage;
                    }
                }
                traceoptions {
                    file <filename> <files number> <match regex> <size size> <world-readable |
no-world-readable>;
                    flag flag;
                    no-remote-trace;
                }
                username username;
            }
            max-datasize max-datasize;
            start {
                start-options {
                    mem-factor percentage;
                }
            }
            traceoptions {

                file <filename> <files number> <size size> <world-readable | no-world-readable>;
                flag flag;
                no-remote-trace;
            }
        }
    }
    providers {
        provider-id {
            license-type license deployment-scope [ deployments ];
        }
    }
```

```
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure extensions to Junos OS.

You must configure the `providers` statement to enable application packages developed using the Junos SDK to be deployed and run on the router.

You must configure the `extension-service` statement to enable application packages developed using the Juniper Extension Toolkit (JET) to be deployed and run on the device.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: This configuration is optional in Junos OS Evolved. You can run JET applications directly using a Python interpreter instead of configuring and invoking them in the CLI.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`extension-service` option introduced in Junos OS Release 16.1.

external-assigned (Address Pools)

IN THIS SECTION

- Syntax | 6725
- Hierarchy Level | 6725
- Description | 6725
- Required Privilege Level | 6725
- Release Information | 6725

Syntax

```
external-assigned;
```

Hierarchy Level

```
[edit access address-assignment address-pools name family inet network network-prefix],
[edit access address-assignment address-pools name family inet6 network network-prefix]
```

Description

Assign addresses in network prefixes statically.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Understanding Source IP Filtering with Address Pools in TDF Domains for IP-Based Subscribers

Configuring Address Pools for Source-IP Filtering of IP-Based Subscribers

external-authority

IN THIS SECTION

- Syntax | 6726
- Hierarchy Level | 6726
- Description | 6726
- Required Privilege Level | 6727
- Release Information | 6727

Syntax

```
external-authority;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server pool-match-order],
[edit logical-systems logical-system-name system services dhcp-local-server pool-match-order],
[edit routing-instances routing-instance-name system services dhcp-local-server pool-match-order],
[edit system services dhcp-local-server pool-match-order]
```

Description

Specify that an external authority (for example, RADIUS or Diameter) provides the address assignment.

When RADIUS is the external authority, the router uses the Framed-IPv6-Pool attribute (RADIUS attribute 100) to select the pool. When Diameter is the external authority, the router uses the Diameter counterpart of RADIUS Framed-IPv6-Pool attribute.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use

Understanding Differences Between Legacy DHCP and Extended DHCP

Address-Assignment Pools Overview

F

IN THIS SECTION

- f-max-period | 6742
- f-max-period (rtp) | 6744
- fabric | 6745
- fabric (Class-of-Service) | 6747
- fabric-control | 6749
- fabric-id | 6750
- fabric-interfaces | 6751
- fabric-load-balance | 6753
- fabric loopback wan | 6755

- fabric-maintenance-associations | **6756**
- fabric-maintenance-end-points | **6759**
- fabric (OAM) | **6761**
- fabric-options | **6763**
- fabric-tree-root | **6766**
- fabric-type | **6767**
- facility | **6769**
- facility-override | **6770**
- facility-override (Security) | **6771**
- facility-override (Service Sets) | **6772**
- facility-override (System Log Reporting) | **6774**
- fail | **6775**
- fail-filter (Dynamic Profiles) | **6777**
- failover (Chassis) | **6778**
- failover-delay | **6780**
- failover-delay | **6781**
- failover-resync | **6783**
- failover (System Process) | **6784**
- failover-within-preference (L2TP LAC) | **6786**
- failure-action | **6787**
- failure-action (Protocols LDP) | **6789**
- failure-action (Protocols MPLS) | **6790**
- fallback-block (Security Antivirus) | **6792**
- fallback-key | **6794**
- fallback-non-block (Security Antivirus) | **6795**
- fallback-option (ICAP Redirect Service) | **6797**
- fallback-options (Security Antivirus Juniper Express Engine) | **6799**
- fallback-options (Security Antivirus Kaspersky Lab Engine) | **6800**
- fallback-options (Security Antivirus Sophos Engine) | **6802**
- fallback-settings (Security Web Filtering) | **6804**
- fallback-settings (Security Web Filtering Juniper Local) | **6805**
- fallback-settings (Security Web Filtering Websense Redirect) | **6806**

- falling-event-index | **6808**
- falling-event-index (RMON) | **6809**
- falling-threshold (Health Monitor) | **6811**
- falling-threshold (integer) | **6812**
- falling-threshold-interval | **6814**
- falling-threshold (percentage) | **6815**
- falling-threshold (RMON) | **6817**
- false-positives | **6818**
- family (RFC 2544 Benchmarking Tests) | **6820**
- family (logical-interface) | **6822**
- family (ELS) | **6824**
- family (QFX) | **6829**
- family (protocols ldp) | **6831**
- family (interfaces) | **6832**
- family (Link Interfaces) | **6839**
- family (Address-Assignment Pools) | **6841**
- family (Address Pools) | **6843**
- family (Aggregated Multiservices) | **6844**
- family (Bootstrap) | **6846**
- family bridge | **6847**
- family bridge (GRE Interfaces) | **6849**
- family (chassis) | **6851**
- family (Dynamic Demux Interface) | **6853**
- family (Dynamic Firewalls) | **6855**
- family (Dynamic PPPoE) | **6857**
- family (Dynamic Standard Interface) | **6859**
- family (Exclude Prefix) | **6862**
- family fcoe | **6863**
- family (Filtering) | **6865**
- family (Firewall) | **6867**
- family (Firewall Filter) | **6869**
- family inet | **6871**

- family inet (Interfaces) | **6874**
- family inet6 | **6878**
- family (Interfaces) | **6881**
- family (Local RP) | **6883**
- family (Monitoring) | **6885**
- family mpls | **6886**
- family (Multi-destination) | **6890**
- family (Multifield Classifier) | **6891**
- family multiservice | **6893**
- family (Port Mirroring) | **6896**
- family (Protocols AMT Relay) | **6898**
- family (Protocols BGP) | **6900**
- family (Protocols IS-IS) | **6905**
- family (Protocols PIM) | **6907**
- family (Protocols PIM Interface) | **6909**
- family route-target | **6911**
- family (Sampling) | **6913**
- family-state-change-immediate-update | **6916**
- family (TDF Interface) | **6917**
- family (Voice Services) | **6918**
- family vpls (Layer 2 Pseudowires) | **6920**
- family (VRF Advertisement) | **6921**
- fast-aps-switch | **6922**
- fast-failover | **6924**
- fast-interval | **6926**
- fast-lookup-filter | **6928**
- fast-lookup-filter (PTX) | **6930**
- fast-lookup-tuple | **6935**
- fast-lookup-tuple-list | **6937**
- fast-lookup-tuple-list (policy-options) | **6938**
- fast-poe | **6939**
- fast-reroute-priority | **6941**

- fast-reroute (Protocols MPLS) | **6943**
- fast-reroute (Protocols RSVP) | **6944**
- fast-update-filter (Dynamic Firewalls) | **6946**
- fast-lookup-filter (Protocols BGP) | **6947**
- fastether-options | **6949**
- fatal | **6950**
- fate-sharing | **6952**
- fc-fabrics | **6954**
- fc-map | **6957**
- fc-options | **6959**
- fc2 | **6960**
- fcoe | **6962**
- fcoe-lag | **6963**
- fcoe-trusted | **6965**
- feature-profile | **6967**
- features (VNF) | **6978**
- features | **6979**
- feb (Assigning a FEB to a Redundancy Group) | **6981**
- feb (chassis) | **6982**
- feb (Creating a Redundancy Group) | **6984**
- fec | **6985**
- fec (MX Series) | **6987**
- fec (protocols ldp) | **6990**
- fec (ether) | **6992**
- fec (gigether) | **6995**
- feed-name | **6997**
- feed-server | **6999**
- feeds (T640, T1600, and T4000 Routers with Six-Input DC Power Supply) | **7001**
- fib-install-priority | **7003**
- fib-local | **7005**
- fib-next-hop-split | **7006**
- fib-prioritization | **7007**

- fib-remote | **7009**
- fibre-channel (Alarm) | **7010**
- fibre-channel (Family Interfaces) | **7012**
- fibre-channel (Port) | **7013**
- fibrechannel-options | **7015**
- fields (Flat-File Accounting Options) | **7016**
- fields (for Interface Profiles) | **7019**
- fields (for Routing Engine Profiles) | **7021**
- file | **7023**
- file | **7025**
- file (Security Certificates) | **7027**
- file (Analytics) | **7028**
- file (App Engine Virtual Machine Management Service) | **7030**
- file archive | **7032**
- file (Associating with a Profile) | **7036**
- file checksum md5 | **7037**
- file checksum sha-256 | **7039**
- file checksum sha1 | **7041**
- file (Commit Scripts) | **7042**
- file compare | **7044**
- file (Configuring a Log File) | **7047**
- file copy | **7049**
- file delete | **7052**
- file (Event Scripts) | **7054**
- file (Extended DHCP Relay Agent and Helpers Trace Options) | **7056**
- file (Flat-File Accounting Options) | **7058**
- file (JET) | **7059**
- file list | **7062**
- file (Next Gen Services Global System Logging) | **7065**
- file (Op Scripts) | **7067**
- file (Packet Forwarding Engine Traffic Traceoptions) | **7069**
- file (QFabric System) | **7071**

- file rename | **7072**
- file (Sampling) | **7075**
- file (Security Log) | **7076**
- file show | **7078**
- file (SNMP Scripts) | **7080**
- file-specification (File Format) | **7082**
- file-specification (Interface Mapping) | **7083**
- file (System Logging) | **7085**
- file (System Logging) | **7086**
- file (Trace Options) | **7090**
- file (Tracing for Origin AS Validation) | **7091**
- file-type | **7093**
- file-serialization | **7094**
- filename | **7096**
- filename-extension | **7097**
- filename (Next Gen Services Global System Logging) | **7099**
- filename-prefix | **7100**
- filename (Sampling) | **7101**
- files | **7103**
- files (accounting) | **7104**
- files (Next Gen Services Global System Logging) | **7105**
- files (Sampling and Traceoptions) | **7107**
- files (syslog) | **7108**
- fill-level | **7110**
- fill-level (Drop Profiles) | **7112**
- fill-level (Interpolated Value) | **7113**
- filter (Flow Monitoring) | **7115**
- filter (interfaces) | **7116**
- filter (Interfaces) | **7117**
- filter | **7119**
- filter (routing-options) | **7122**
- filter (Applying to a Logical Interface) | **7123**

- filter (Applying to an Interface) | **7125**
- filter (Configuring) | **7127**
- filter-duplicates | **7129**
- filter (Dynamic Interface Unit) | **7130**
- filter (Dynamic Profiles Filter Attachment) | **7132**
- filter (Dynamic Profiles Filter Creation) | **7134**
- filter (export routing policy action) | **7135**
- filter (Firewall Filters) | **7136**
- filter (Identity Management Advanced Query) | **7138**
- filter-interfaces | **7142**
- filter (IPv4, IPv6, and MPLS) | **7144**
- filter (Layer 2 and Layer 3 Interfaces) | **7145**
- filter (Layer 3 Interfaces) | **7146**
- filter-list-template | **7148**
- filter (memory-enhanced) | **7150**
- filter-profile | **7151**
- filter (Security) | **7153**
- filter-specific (QFX) | **7154**
- filter-specific | **7155**
- filter (VLANs) | **7157**
- filter (VPLS) | **7158**
- filtering-type | **7160**
- filtering-type (Source NAT Next Gen Services) | **7161**
- filters | **7162**
- fin-invalidate-session | **7165**
- fin-no-ack | **7167**
- fin-no-ack (IDS Screen Next Gen Services) | **7168**
- final-response-timeout (OCS Partition) | **7170**
- finger | **7171**
- fip | **7173**
- fip-security | **7175**
- firewall | **7176**

- firewall-authentication | **7179**
- firewall-authentication-forced-timeout | **7181**
- firewall-authentication (Security) | **7184**
- firewall-authentication (Security Policies) | **7186**
- firewall-authentication-service | **7188**
- firewall-authentication (tenants) | **7189**
- firewall (Dynamic Firewalls) | **7191**
- firewall-install-disable | **7194**
- firewall-user | **7196**
- firmware-revision | **7197**
- fiveqi-map | **7199**
- fixed-stuff-bytes | **7201**
- fka-adv-period | **7202**
- flag (App Engine Virtual Machine Management Service) | **7204**
- flag (Interfaces) | **7206**
- flag (Next Gen Services Global System Logging) | **7208**
- flag (Services) | **7209**
- flag (SMTP) | **7211**
- flag (Tracing for Origin AS Validation) | **7212**
- flap-trap-monitor | **7215**
- flat-file-profile (Accounting Options) | **7217**
- flat-file-profile (Extensible Subscriber Services) | **7220**
- flex-algorithm | **7221**
- flex-algorithm (Protocols OSPF) | **7223**
- flex-flow-sizing | **7225**
- flex-hashing | **7227**
- flexible-cross-connect-vlan-aware | **7230**
- flexible-cross-connect-vlan-unaware | **7232**
- flexible-queuing-mode | **7233**
- flexible-vlan-tagging | **7235**
- flexible-vlan-tagging (Interfaces) | **7237**
- flood | **7238**

- flood-as-unknown-unicast | **7239**
- flood-groups | **7241**
- flood-reflector (Protocols IS-IS) | **7242**
- flood (Security ICMP) | **7245**
- flood (Security UDP) | **7246**
- flood (VLANs) | **7248**
- flow (protocols bgp) | **7250**
- flow | **7252**
- flow-action | **7259**
- flow-active-timeout (Flow Monitoring) | **7261**
- flow-collector (Flow Collection) | **7263**
- flow-control (ACX Series) | **7265**
- flow-control | **7267**
- flow-control (Interfaces) | **7269**
- flow-control-options | **7270**
- flow-control-queue (Output Congestion Notification) | **7272**
- flow-descriptions | **7274**
- flow-detect-time (DDoS Flow Detection) | **7276**
- flow-detection (DDoS Flow Detection) | **7278**
- flow-detection (DDoS Packet Level) | **7279**
- flow-detection-mode (DDoS Flow Detection) | **7281**
- flow-detection-mode (DDoS Global Flow Detection) | **7283**
- flow-export-destination (Flow Collection) | **7285**
- flow-export-rate | **7286**
- flow-export-timer | **7288**
- flow-gate | **7289**
- flow-groups | **7291**
- flow-inactive-timeout (Flow Monitoring) | **7296**
- flow (IPv6) | **7298**
- flow-key (Flow Monitoring) | **7304**
- flow-label | **7306**
- flow-label-receive | **7307**

- flow-label-receive-static | **7309**
- flow-label-static | **7311**
- flow-label-transmit | **7312**
- flow-label-transmit-static | **7314**
- flow-level-bandwidth (DDoS Flow Detection) | **7316**
- flow-level-control (DDoS Flow Detection) | **7317**
- flow-level-control (DDoS Global Flow Detection) | **7319**
- flow-level-detection (DDoS Flow Detection) | **7320**
- flow-map | **7322**
- flow-monitoring | **7323**
- flow-monitoring (Inline Monitoring Services) | **7326**
- flow-monitoring (Services) | **7328**
- flow-recover-time (DDoS Flow Detection) | **7330**
- flow-report-rate (DDoS Flow Detection) | **7331**
- flow (Security Flow) | **7333**
- flow (Security IDP) | **7339**
- flow-server | **7344**
- flow-session | **7347**
- flow-specs | **7349**
- flow-table-size | **7351**
- flow-table-size (Chassis) | **7354**
- flow-tap | **7355**
- flow-tap-dtcp | **7357**
- flow-timeout-time (DDoS Flow Detection) | **7359**
- flow-sample | **7360**
- flows (PCC Rules) | **7362**
- fnp | **7364**
- force-continue (OCS Partition) | **7365**
- force-control-packets-on-transit-path | **7367**
- force-discover (DHCP Client) | **7369**
- force-entry (IDS MS-DPC) | **7370**
- force-ip-reassembly | **7371**

- force-premium (Firewall Filter Action) | **7373**
- force switch | **7374**
- force-switch | **7375**
- force-up (EX and QFX Series) | **7376**
- force-up | **7378**
- force-version (IEEE 802.1D STP) | **7379**
- force-fbf-terms | **7381**
- format (Flat-File Accounting Options) | **7382**
- format (LRF Profile) | **7383**
- format (Next Gen Services Service-Set Remote System Logging) | **7385**
- format (Security Log) | **7386**
- format (Security Log Stream) | **7388**
- format (Unified Edge Gateways) | **7389**
- forward-and-send-to-re | **7392**
- forward-delay | **7393**
- forward-only | **7395**
- forward-only (DHCP Relay Agent) | **7396**
- forward-only (DHCP Relay Agent Option) | **7398**
- forward-only (DNS) | **7400**
- forward-only-replies (DHCP Relay Agent) | **7401**
- forward-policy-mismatch (Security Group VPN Member) | **7402**
- forward-snooped-clients (DHCP Local Server) | **7404**
- forward-snooped-clients (DHCP Relay Agent) | **7405**
- forwarding-cache (Bridge Domains) | **7407**
- forwarding-cache (Flow Maps) | **7408**
- forwarding-class | **7410**
- forwarding-class-accounting | **7413**
- forwarding-class (BA Classifiers) | **7415**
- forwarding-class (CoS Host Outbound Traffic) | **7417**
- forwarding-class-default (Forwarding Policy) | **7418**
- forwarding-class (Dynamic Scheduler Maps) | **7420**
- forwarding-class (Fibre Channel Interfaces) | **7421**

- forwarding-class (Firewall Filter Action) | **7423**
- forwarding-class (for DHCP Snooping or DAI Packets) | **7425**
- forwarding-class (Forwarding Policy) | **7426**
- forwarding-class (Fragmentation) | **7428**
- forwarding-class (Gigabit Ethernet IQ Classifier) | **7430**
- forwarding-class (Interfaces) | **7432**
- forwarding-class (Multifield Classifiers) | **7433**
- forwarding-class (PCC Action Profiles) | **7434**
- forwarding-class (Restricted Queues) | **7436**
- forwarding-class (RFC 2544 Benchmarking) | **7437**
- forwarding-class (Sampling) | **7439**
- forwarding-class (Services CoS Fragmentation Properties) | **7440**
- forwarding-class (Services PIC Classifiers) | **7442**
- forwarding-class-set | **7443**
- forwarding-class-set (Interconnect Device) | **7444**
- forwarding-class-sets | **7446**
- forwarding-classes | **7447**
- forwarding-context (Protocols BGP) | **7453**
- forwarding-db-size | **7454**
- forwarding (Diameter Network Element) | **7456**
- forwarding-instance identifier | **7457**
- forwarding-mode (100-Gigabit Ethernet) | **7459**
- forwarding-mode (PTX Series Packet Transport Routers) | **7460**
- forwarding-mode (Security Content Security Policy) | **7461**
- forwarding-options (Routing Instance) | **7463**
- forwarding-options (Bridge Domains or VLANs) | **7467**
- forwarding-options (Security) | **7473**
- forwarding-packages | **7478**
- forwarding-policy | **7479**
- forwarding-process | **7481**
- forwarding-rib | **7483**
- forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter) | **7485**

- forwarding-state-bit (Per Family for BGP Graceful Restart) | **7488**
- forwarding-table | **7490**
- forwarding-options (Unified Forwarding Table Profile) | **7493**
- fpc (QFX Series) | **7499**
- fpc (Chassis) | **7501**
- fpc (DDoS) | **7505**
- fpc error | **7507**
- fpc-feb-connectivity | **7509**
- fpc (Interfaces) | **7511**
- fpc (Junos Fusion) | **7512**
- fpc (Junos Fusion Satellite Device Cluster) | **7514**
- fpc (M320, T320 and T640 Devices) | **7516**
- fpc (MX Series 5G Universal Routing Platforms) | **7519**
- fpc-nmi-volt-fail-knob | **7521**
- fpc (Notification Control) | **7523**
- fpc-offline-on-blackholing | **7524**
- fpc (PTX Series Devices) | **7525**
- fpc-restart | **7528**
- fpc-resync | **7529**
- fpc-slice | **7530**
- fpc (Subscriber Limits) | **7533**
- fpc (TX Matrix and TX Matrix Plus Routers) | **7534**
- fpcs (Node Slicing) | **7536**
- fpcs (NSSU Upgrade Groups) | **7538**
- fragment (IDS Screen Next Gen Services) | **7540**
- fragment-limit | **7541**
- fragment-threshold | **7542**
- fragment-threshold (Forwarding Class Maps) | **7544**
- fragment-threshold (Interfaces LSQ) | **7545**
- fragmentation-map | **7547**
- fragmentation-maps | **7548**
- frame-error | **7550**

- frame-mode (Dynamic Traffic Shaping) | **7552**
- frame-period | **7554**
- frame-period-summary | **7555**
- frame-relay-de (CoS Interfaces) | **7557**
- frame-relay-de (CoS Loss Priority) | **7558**
- frame-relay-de (CoS Rewrite Rule) | **7560**
- frame-relay-de (Defining Loss Priority Maps) | **7561**
- frame-relay-de (Defining Loss Priority Rewrites) | **7563**
- framed-ip-address | **7564**
- framed-ipv6-prefix | **7566**
- framed-pool | **7567**
- framing (10-Gigabit Ethernet Interfaces) | **7568**
- framing (E1 Options) | **7571**
- framing (E1 Options for BITS Interfaces) | **7572**
- framing (T1 Options) | **7573**
- framing (T1 Options for BITS Interfaces) | **7575**
- free-fw-memory-watermark (Resource Monitor) | **7576**
- free-heap-memory-watermark (Resource Monitor) | **7578**
- free-nh-memory-watermark (Resource Monitor) | **7579**
- frequency-lock-threshold | **7580**
- from (Application Aware Access List) | **7582**
- from (services-ipsec-vpn) | **7584**
- from (Firewall) | **7585**
- from (Captive Portal Content Delivery Tags) | **7587**
- from (HTTP Header Enrichment) | **7588**
- from (IDS MS-DPC) | **7590**
- from (PCC Rules) | **7591**
- from (PCP) | **7593**
- from (Protocols MPLS) | **7595**
- from (Security NAT) | **7596**
- from (Services CoS) | **7597**
- from (Services NAT) | **7599**

- from (Services Stateful Firewall) | 7600
- from (Subscriber Secure Policy) | 7602
- from (TDF Domain Selection) | 7603
- from-zone (Security IDP Policy) | 7607
- from-zone (Security Policies) | 7608
- from-zone (Security Policies Global) | 7612
- fru-poweron-sequence | 7613
- fte (Port) | 7615
- ftp | 7617
- ftp (Access) | 7619
- ftp (Content Security Policy Anti-Virus) | 7621
- ftp (Content Security Policy Content Filtering) | 7622
- ftp (Flow Collector Files) | 7624
- ftp (Security ALG) | 7626
- ftp (Services CoS) | 7628
- ftp (Services CoS Next Gen Services) | 7629
- ftp (Transfer Log Files) | 7631
- full | 7632
- function (Diameter Network Element for Subscriber Management) | 7633
- function (Diameter Network Element) | 7634
- function (Diameter Route) | 7636
- functional-zone | 7638

f-max-period

IN THIS SECTION

- Syntax | 6743
- Hierarchy Level | 6743
- Description | 6743

- Options | 6743
- Required Privilege Level | 6743
- Release Information | 6743

Syntax

```
f-max-period number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-name unit logical-unit-number rtp]
```

Description

For all adaptive services interfaces and for ISDN interfaces on J Series Services Routers. Specify the maximum number of compressed packets allowed between the transmission of full headers in a compressed Real-Time Transport Protocol (RTP) traffic stream.

Options

number—Maximum number of packets. The value can be from 1 through 65535.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring Bandwidth on Demand
- Junos OS Services Interfaces Library for Routing Devices

f-max-period (rtp)

IN THIS SECTION

- Syntax | 6744
- Hierarchy Level | 6744
- Description | 6744
- Options | 6745
- Required Privilege Level | 6745
- Release Information | 6745

Syntax

```
f-max-period number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number compression rtp],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
compression rtp]
```

Description

Specify the maximum number of compressed packets allowed between the transmission of full headers in a compressed Real-time Transport Protocol (RTP) traffic stream.

Options

number—Maximum number of packets.

- **Default:** 256

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Compression of Voice Traffic

fabric

IN THIS SECTION

- Syntax | 6745
- Hierarchy Level | 6746
- Description | 6747
- Required Privilege Level | 6747
- Release Information | 6747

Syntax

```
fabric
    aliases {
```

```
        director-device director-device-name {
            assigned-director-device-name;
        }
        interconnect-device interconnect-device-name {
            assigned-interconnect-device-name;
        }
        node-device node-device-name {
            assigned-node-device-name;
        }
    }
    flow-groups {
        flow-group-name {
            node-device (node-device-name | serial-ID | alias);
            interconnect-device interconnect-device-name;
        }
        node-device (node-device-name | serial-ID | alias) {
            interconnect-device interconnect-device- name {
                preference (high | normal | never);
            }
        }
    }
    protocols {
        fabric-control {
            graceful-restart {
                restart-timeseconds;
                stale-routes-time seconds;
            }
        }
    }
    resources {
        node-group node-group-name {
            node-device node-device-name;
            network-domain;
        }
    }
}
```

Hierarchy Level

```
[edit]
```

Description

(QFabric systems only) Define resources, routing options, fabric control protocol settings, and the mapping of user-defined names to QFabric system components.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

show fabric administration inventory | 19878

Understanding QFabric System Administration Tasks and Utilities

fabric (Class-of-Service)

IN THIS SECTION

- Syntax | 6748
- Hierarchy Level | 6748
- Description | 6748
- Required Privilege Level | 6748
- Release Information | 6748

Syntax

```
fabric {
    scheduler-map {
        priority (high | low) scheduler scheduler-name;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Define CoS parameters of the switch fabric. For M320 routers, MX Series routers, T Series routers and EX Series switches only, associate a scheduler with a fabric priority.

On EX Series switches, this statement is supported only on EX8200 standalone switches and EX8200 Virtual Chassis.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Associating Schedulers with Fabric Priorities

fabric-control

IN THIS SECTION

- Syntax | 6749
- Hierarchy Level | 6749
- Description | 6749
- Required Privilege Level | 6749
- Release Information | 6750

Syntax

```
fabric-control {
    graceful-restart {
        restart-time seconds;
        stale-routes-time seconds;
    }
}
```

Hierarchy Level

```
[edit fabric protocols]
```

Description

Specify attributes for the fabric control protocol.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X52-D10.

RELATED DOCUMENTATION

Understanding Routing Engines in the QFabric System

fabric-id

IN THIS SECTION

- Syntax | 6750
- Hierarchy Level | 6750
- Description | 6750
- Options | 6751
- Required Privilege Level | 6751
- Release Information | 6751

Syntax

```
fabric-id fc-fabric-id;
```

Hierarchy Level

```
[edit fc-fabrics fc-fabric-name]
```

Description

Configure a unique identifier for the FC fabric.

NOTE: Changing the ID of an FC fabric causes all logins to drop and forces the ENodes to log in again.

Options

fc-fabric-id —Unique identifier of the FC fabric.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fabric | 19929
Configuring an FCoE-FC Gateway Fibre Channel Fabric
Understanding an FCoE-FC Gateway

fabric-interfaces

IN THIS SECTION

- Syntax | 6752
- Hierarchy Level | 6752
- Description | 6752
- Options | 6752
- Required Privilege Level | 6752
- Release Information | 6752

Syntax

```
fabric-interfaces {
    node-group (node-group-name | all) {
        oxid (enable | disable);
    }
}
```

Hierarchy Level

```
[edit forwarding-options hash-key family fcoe]
```

Description

Specify that you are enabling or disabling OxID hash control on the fabric ports of a QFabric system Node group. OxID hash control is enabled or disabled on the fabric ports that face an FCoE forwarder (FCF).

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X52-D10.

RELATED DOCUMENTATION

Enabling and Disabling CoS OxID Hash Control for FCoE Traffic on QFabric Systems

Understanding OxID Hash Control for FCoE Traffic Load Balancing on QFabric Systems

fabric-load-balance

IN THIS SECTION

- Syntax | 6753
- Hierarchy Level | 6753
- Description | 6753
- Default | 6754
- Options | 6754
- Required Privilege Level | 6754
- Release Information | 6755

Syntax

```
fabric-load-balance {
    flowlet {
        inactivity-interval interval;
    }
    per-packet;
}
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key]
```

Description

Enable adaptive load balancing (ALB) for a VCF, and specify how ALB is implemented.

NOTE: Starting in Junos OS Release 14.1X53-D46, 15.1R7, 16.1R6, 17.1R3, 17.2R2, 17.3R2, and 17.4R1, the ALB feature is deprecated. If `fabric-load-balance` is enabled in the configuration, delete the configuration item upon upgrading Junos OS.

When ALB is enabled, the Virtual Chassis ports (VCPs) are reset. Packets are dropped and might potentially arrive out of order for a brief period of time as a result of this VCP reset. Normal operation of the VCF resumes after the VCP reset with no further user action.

Default

ALB is disabled, by default.

If you do not specify a mode when enabled ALB, ALB is enabled using flowlet mode with an inactivity timer of 16 microseconds.

Options

`flowlet` Implement ALB by using flowlets.

When ALB is implemented using flowlets, traffic flows that enter the VCF are spliced into smaller flows—flowlets—and individually forwarded across the VCF to the same destination device over different paths when the inactivity time between packet bursts on the sending interface exceeds the user-configurable inactivity interval.

The inactivity interval is 16 microseconds by default, and can be configured using the `inactivity-interval` statement. You should configure the inactivity interval to ensure in-order packet delivery, so that overall performance is not negatively impacted by the packet reordering process at the receiving device. To ensure in-order packet delivery, the inactivity interval should be larger than the largest latency skew among all the paths in the VCF from any node to any other node.

Implementing ALB using flowlets is especially effective in environments that periodically experience extremely large traffic flows—*elephant flows*—that are substantially larger than the majority of other traffic flowing through the VCF. The VCF is better able to manage elephant flows by splicing them into smaller flowlets using ALB.

`per-packet` Implement ALB using per-packet mode.

When per-packet mode is enabled, the VCF forwarding algorithm dynamically monitors all paths in the VCF and forwards packets to destination devices using the best available path at that moment. Flows are reordered at the destination node when per-packet mode is used to enable ALB, so some performance impact due to packet reordering is experienced.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

Statement deprecated starting in Junos OS Releases 14.1X53-D46, 15.1R7, 16.1R6, 17.1R3, 17.2R2, 17.3R2, and 17.4R1.

RELATED DOCUMENTATION

Understanding Traffic Flow Through a Virtual Chassis Fabric

fabric loopback wan

IN THIS SECTION

- Syntax | 6755
- Hierarchy Level | 6755
- Description | 6756
- Default | 6756
- Options | 6756
- Required Privilege Level | 6756
- Release Information | 6756

Syntax

```
fabric loopback wan (off | on);
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

Allow the fabric bound packets to take new fabric loopback path, freeing up the WAN bandwidth thus improving the sampling and tunneling performance of the router.

You can configure fabric side loopback by using the `fabric loopback wan off` option or switch to WAN side by using the `fabric loopback wan on` option. By default, Junos OS uses fabric loopback for the loopback packets.

Default

`off`

Options

`off`—Configure fabric bound packets to take fabric loopback path.

`on`—Configure fabric bound packets to take the WAN loopback path.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3.

fabric-maintenance-associations

IN THIS SECTION

- Syntax | 6757
- Hierarchy Level | 6757
- Description | 6757
- Options | 6757

Required Privilege Level | 6758
Release Information | 6758

Syntax

```
fabric-maintenance-associations {
    fma-name {
        description string;
        fabric-maintenance-end-points {
            fmep-id {
                description string;
                fmep-interface interface-name;
                fmep-name name;
            }
        }
        vlan-name vlan-name;
    }
}
```

Hierarchy Level

```
[edit protocols oam fabric]
```

Description

Configure the fabric maintenance association (FMA) parameters for internal fabric monitoring. The FMA defines a set of QFabric system interfaces (fabric maintenance endpoints [FMEPs]) that are members of a given VLAN which are used in internal fabric monitoring operations. The FMA associates the set of FMEPs with a VLAN. Configuration of the VLAN is mandatory.

The remaining statements under this hierarchy level are explained separately.

Options

description *string* Description of the FMA.

- **Syntax:** 1 to 128 alphanumeric characters

fma-name Name of the FMA.

- **Syntax:** 1 to 20 alphanumeric characters

vlan-name *vlan-name* Name of the VLAN for which the FMA is defined.

> **NOTE**: Configuring the VLAN option is mandatory.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | **29924**

fabric-maintenance-end-points

IN THIS SECTION

- Syntax | 6759
- Hierarchy Level | 6759
- Description | 6759
- Options | 6760
- Required Privilege Level | 6760
- Release Information | 6760

Syntax

```
fabric-maintenance-end-points {
    fmep-id {
        description string;
        fmep-interface interface-name;
        fmep-name name;
    }
}
```

Hierarchy Level

```
[edit protocols oam fabric fabric-maintenance-associations fma-name]
```

Description

Configure the fabric maintenance association endpoint (FMEP) parameters within a fabric maintenance association (FMA). The FMA associates the FMEP with a VLAN for purposes of internal fabric monitoring.

Options

description *string*

Description of the FMEP.

- **Syntax:** 1 to 128 alphanumeric characters

fmep-id

32-bit integer that is used to identify an FMEP.

- **Range:** 1 to 64,000

fmep-name *name*

Name of the fabric maintenance association endpoint.

- **Syntax:** 1 to 20 alphanumeric characters

fmep-interface *interface-name*

Interface on a QFabric system Node device.

- **Syntax:** *node-device-name:interface-type-fpc/pic/slot.unit*

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | **29924**

fabric (OAM)

IN THIS SECTION

- Syntax | 6761
- Hierarchy Level | 6762
- Description | 6763
- Required Privilege Level | 6763
- Release Information | 6763

Syntax

```
fabric {
    fabric-maintenance-associations {
        fma-name {
            description string;
            fabric-maintenance-end-points {
                fmep-id {
                    description string;
                    fmep-interface interface-name;
                    fmep-name name;
                }
            }
            vlan-name vlan-name;
        }
    }
    flow-specs {
        flow-specification-name {
            ethernet-frame-size size;
            multicast-ipv4 {
                dest-ip-multicast-group ipv4-mcast-address;
                source-ip ipv4-address;
            }
            multicast-vlan-flood;
            unicast-ethernet {
                destination-mac destination-mac-address destination-mac-mask destination-mac-
mask;
```

```
                ethertype ethertype;
                source-mac mac-address source-mac-mask source-mac-mask;
            }
            unicast-ethernet-ipv4 {
                destination-ip destination-ip-address destination-ip-mask destination-ip-mask;
                destination-l4-port destination-l4-port-number;
                destination-mac destination-mac-address;
                ip-proto protocol;
                source-ip source-ip-address source-ip-mask source-ip-mask;
                source-l4-port source-l4-port-number;
            }
        }
    }
    faboam-trace-options {
        file filename{
            files number;
            no-stamp;
            no-world-readable;
            replace;
            size size;
            world-readable;
        }
        flag {
            all;
            debug-all;
            ffping;
            generic;
            netio;
            packets;
            trace-route;
        }
    }
}
```

Hierarchy Level

```
[edit protocols oam]
```

Description

Configure the fabric maintenance association (FMA) and flow specification parameters for internal fabric monitoring.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | **29924**

fabric-options

IN THIS SECTION

- Syntax | **6764**
- Hierarchy Level | **6764**
- Description | **6764**
- Options | **6765**
- Required Privilege Level | **6765**

Release Information | 6765

Syntax

```
fabric-options {
    member-interfaces member-interface-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure fabric interface specific options in chassis clusters.

The fabric is the data link between the nodes and is used to forward traffic between the chassis. Each node in the chassis requires a fabric interface configured. To create a fabric link between the two chassis requires the creation of a special interface called the fab interface. Node 0's fabric interface, called fab0, and node 1's fabric interface, called fab1.

Only the same type of interfaces can be configured as fabric children, and you must configure an equal number of child links for fab0 and fab1.

You can enable and disable the fabric link on SRX1500, SRX4100, SRX4200, and SRX4600 devices.

From configuration mode,

run the `set interfaces fab0 disable` command to disable the fabric link interface fab0 and `set interfaces fab1 disable` command to disable the fabric link interface fab1.

run the `delete interfaces fab0 disable` to enable fabric link interface fab0 and `delete interfaces fab1 disable` to enable fabric link interface fab1.

```
fab0 {
    fabric-options {
        member-interfaces {
            ge-0/0/4;
```

```
        }
    }
}
fab1 {
    fabric-options {
        member-interfaces {
            ge-2/0/4;
        }
    }
}
```

NOTE: When you run the `system autoinstallation` command, the command will configure unit 0 logical interface for all the active state physical interfaces. However, a few commands such as `fabric-options` do not allow the physical interface to be configured with a logical interface. If the `system autoinstallation` and the `fabric-options` commands are configured together, the following message is displayed:

```
incompatible with 'system autoinstallation'
```

Options

member-interface-name—Member interface name.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring the Chassis Cluster Fabric Interfaces

fabric-tree-root

IN THIS SECTION

- Syntax | 6766
- Hierarchy Level | 6766
- Description | 6766
- Required Privilege Level | 6767
- Release Information | 6767

Syntax

```
fabric-tree-root;
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

Optionally specify a member of a VCF to be a root node for creating the shared multicast distribution trees (MDTs) used in forwarding broadcast, unknown unicast, and multicast (BUM) traffic within the VCF. To configure this option, the VCF must be configured as "auto-provisioned" on page 2725 or "preprovisioned" on page 13115.

> **NOTE**: This option only affects MDTs used for directing traffic internally between member nodes of the VCF, and does not affect how BUM traffic is managed on the network outside the VCF.

Setting this option affects BUM traffic routing behavior within the VCF by changing how VCF MDT trees are created. By default, if this option is not configured for any VCF members, every member switch builds an MDT with itself as the root. When this option is configured for one or more VCF members, MDTs are only created with the specified member or members as root nodes. The number of trees created corresponds to the configured number of root nodes.

If you want to use this option in a VCF, we recommend configuring *all* the spine devices in the VCF (regardless of the role—Routing Engine or linecard), and *only* spine devices, as fabric tree roots. Configuring multiple spine devices as fabric tree roots prevents inadvertently reverting to the default behavior if a spine member becomes unavailable. Avoiding leaf members as fabric tree roots helps to prevent traffic interruption during load rebalancing when a leaf node goes offline or is reset.

See *Understanding Traffic Flow Through a Virtual Chassis Fabric* for more information about how this option affects traffic flow.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D35.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric

Adding a Device to a Virtual Chassis Fabric

fabric-type

IN THIS SECTION

- Syntax | 6768
- Hierarchy Level | 6768
- Description | 6768
- Options | 6768
- Required Privilege Level | 6768
- Release Information | 6768

Syntax

```
fabric-type proxy;
```

Hierarchy Level

```
[edit fc-fabrics fc-fabric-name]
```

Description

Specify that the FC fabric be an FCoE-FC gateway fabric.

Options

proxy—Specify that the switch be an FCoE-FC gateway fabric.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fabric | **19929**

Configuring an FCoE-FC Gateway Fibre Channel Fabric

Understanding an FCoE-FC Gateway

facility

IN THIS SECTION

- Syntax | 6769
- Hierarchy Level | 6769
- Description | 6769
- Required Privilege Level | 6769
- Release Information | 6769

Syntax

```
facility facility-type;
```

Hierarchy Level

```
[edit event-options policy policy-name then priority-override]
```

Description

Within an event policy, override the default facility type of the triggering event so that the event is logged based on the configured facility type.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding the Event System Log Priority in an Event Policy

Example: Configure the Event System Log Priority in an Event Policy

Junos OS System Logging Facilities and Message Severity Levels

priority-override (Event Policy) | **13183**

severity (Event Policy) | **16898**

facility-override

IN THIS SECTION

- Syntax | **6770**
- Hierarchy Level | **6770**
- Description | **6770**
- Options | **6771**
- Required Privilege Level | **6771**
- Release Information | **6771**

Syntax

```
facility-override facility-name;
```

Hierarchy Level

```
[edit interfaces interface-name services-options sysloghost hostname]
```

Description

Override default facility for system log reporting.

Options

facility-name—Name of facility that overrides the default assignment.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Services Interfaces Library for Routing Devices

facility-override (Security)

IN THIS SECTION

- Syntax | 6771
- Hierarchy Level | 6772
- Description | 6772
- Required Privilege Level | 6772
- Release Information | 6772

Syntax

```
facility-override facility;
```

Hierarchy Level

```
[edit security log]
[edit logical-systems name security log]
[edit tenants tenant-name security log]
```

Description

Alternate facility for logging to remote host.

System log server is set up to use a facility-override value to filter or write log files received by a system log agent.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D35.

The [edit `logical-systems` *name* security log] and [edit `tenants` *tenant-name* security log] hierarchy levels introduced in Junos OS Release 19.1R1.

facility-override (Service Sets)

IN THIS SECTION

- Syntax | 6773
- Hierarchy Level | 6773
- Description | 6773
- Options | 6773
- Required Privilege Level | 6773
- Release Information | 6773

Syntax

```
facility-override facility-name;
```

Hierarchy Level

```
[edit services service-set service-set-name syslog host hostname]
```

Description

Override the default facility for system log reporting.

Options

facility-name—Name of the facility that overrides the default assignment. Valid entries are:

1. `authorization`
2. `daemon`
3. `ftp`
4. `kernel`
5. `local0` through `local7`
6. `user`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring System Logging for Service Sets

facility-override (System Log Reporting)

IN THIS SECTION

- Syntax | 6774
- Hierarchy Level | 6774
- Description | 6774
- Options | 6774
- Required Privilege Level | 6775
- Release Information | 6775

Syntax

```
facility-override facility-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group group-name syslog host hostname]
```

Description

Override the default facility for system log reporting.

Options

facility-name—Name of the facility that overrides the default assignment. Valid entries include:

1. `authorization`

2. `daemon`
3. `ftp`
4. `kernel`
5. `local0` through `local7`
6. `user`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring System Logging of L2TP Tunnel Activity

fail

IN THIS SECTION

- Syntax | 6776
- Hierarchy Level | 6776
- Description | 6776
- Options | 6776
- Required Privilege Level | 6776
- Release Information | 6776

Syntax

```
fail string;
```

Hierarchy Level

```
[edit access firewall-authentication pass-through default-profile profile-name (ftp | http |
telnet) banner]
```

Description

Specify the banner that a client sees if the authentication process fails.

Options

string—Banner text. Maximum length of the message text is 250 characters. Enclose the banner text within spaces or special characters—for example, quotation marks (” “).

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring RADIUS and LDAP User Authentication

fail-filter (Dynamic Profiles)

IN THIS SECTION

- Syntax | 6777
- Hierarchy Level | 6777
- Description | 6777
- Options | 6777
- Required Privilege Level | 6778
- Release Information | 6778

Syntax

```
fail-filter filter-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family rpf-check],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family family rpf-check]
```

Description

Specify a filter that evaluates packets that fail a unicast RPF check. The filter determines what action to take with the failed packets. If the fail filter is not configured, the failed packets are silently discarded.

Options

filter-name Name of the filter that evaluates packets that fail the RPF check.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Unicast RPF in Dynamic Profiles for Subscriber Interfaces

Unicast RPF in Dynamic Profiles for Subscriber Interfaces

Understanding Unicast RPF (Routers)

failover (Chassis)

IN THIS SECTION

- Syntax | 6778
- Hierarchy Level | 6779
- Description | 6779
- Options | 6779
- Required Privilege Level | 6779
- Release Information | 6779

Syntax

```
failover {
    on-disk-failure;
    on-loss-of-keepalives;
    on-re-to-fpc-stale;
```

```
}
```

Hierarchy Level

```
[edit chassis redundancy]
```

Description

Specify conditions on the primary Routing Engine that cause the backup router to take primary role.

On Junos OS devices this statement enables automatic Routine Engine switchover. By default this statement is disabled on Junos OS Evolved, but it may be enabled if you are upgrading an existing configuration from Junos OS to Junos OS Evolved. To disable automatic failover on Junos OS Evolved devices, delete this statement by entering the **delete chassis redundancy failover** command.

Options

on-disk-failure	Instruct the backup router to take primary role if it detects hard disk errors on the primary Routing Engine.
on-re-to-fpc-stale	Instruct the backup router to take mastership if the `em0` interface fails on the master Routing Engine. This option is not supported on PTX platforms running Junos OS Evolved. These platforms default to a switchover when disk failure is detected.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`on-re-to-fpc-stale` option introduced in Junos OS Release 15.2.

RELATED DOCUMENTATION

On Detection of a Hard Disk Error on the Primary Routing Engine

Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)

failover-delay

IN THIS SECTION

- Syntax | 6780
- Hierarchy Level | 6780
- Description | 6780
- Options | 6780
- Required Privilege Level | 6781
- Release Information | 6781

Syntax

```
failover-delay milliseconds;
```

Hierarchy Level

```
[edit protocols vrrp]
```

Description

Configure the failover delay for VRRP and VRRP for IPv6 operations.

Options

milliseconds—Specify the failover delay time, in milliseconds.

- **Range:** 50 through 2000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring VRRP and VRRP for IPv6

failover-delay

IN THIS SECTION

- Syntax | 6781
- Hierarchy Level | 6781
- Description | 6782
- Options | 6782
- Required Privilege Level | 6782
- Release Information | 6782

Syntax

```
failover-delay milliseconds;
```

Hierarchy Level

```
[edit protocols vrrp]
```

Description

If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new primary must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new primary).

If you configure a failover delay, the new primary delays sending gratuitous ARP replies for the period that you set. This allows the new primary to send the ARP replies for all of the VRRP groups simultaneously.

Options

milliseconds—Specify the failover delay time, in milliseconds.

- **Range:** 50 through 2000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Troubleshooting VRRP

show vrrp

failover-resync

IN THIS SECTION

- Syntax | 6783
- Hierarchy Level | 6783
- Description | 6783
- Options | 6784
- Required Privilege Level | 6784
- Release Information | 6784

Syntax

```
failover-resync (failover-protocol | silent-failover);
```

Hierarchy Level

```
[edit services l2tp tunnel]
```

Description

Configure the method used by the LAC or the LNS to resynchronize with its peer in the event of a control plane failover. Failure can be the result of a Routing Engine switchover, a daemon restart, or some other cause. During tunnel setup, the L2TP endpoints negotiate the resynchronization method; silent failover is the default.

With the silent failover method, only the failed endpoint is involved in recovering the tunnels and sessions; the nonfailed endpoint remains unaware of the failure.

With the failover protocol method, the nonfailed endpoint keeps the tunnel open with the failed peer, in case the failed peer is able to recover from the failure and resynchronize with the nonfailed peer. The detection of tunnel keepalive failures is delayed. This behavior keeps the tunnel up and the subscribers logged in while traffic is not flowing, preventing service level agreements from being met.

BEST PRACTICE: Use the default method, silent failover.

This statement supersedes the deprecated statement, `disable-failover-protocol`.

Options

failover-protocol Specify the L2TP failover protocol as the resynchronization method, but fall back to silent failover if the other endpoint does not support it.

silent-failover Specify silent failover as the resynchronization method.

- **Default:** `silent-failover`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Configuring the L2TP Peer Resynchronization Method

L2TP Failover and Peer Resynchronization

failover (System Process)

IN THIS SECTION

Syntax | 6785

Hierarchy Level | 6785

- Description | 6785
- Options | 6785
- Required Privilege Level | 6785
- Release Information | 6786

Syntax

```
failover (alternate-media | other-routing-engine);
```

Hierarchy Level

```
[edit system processes process-name]
```

Description

Configure the router to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.

Options

process-name—Junos OS process name. Some of the processes that support the `failover` statement are `bootp`, `chassis-control`, `craft-control`, `ethernet-connectivity-fault-management`, `init`, `interface-control`, `neighbor-liveness`, `pfe`, `redundancy-interface-process`, `routing`, `smg-service`, and `vrrp`.

`alternate-media`—Use the Junos OS image on alternate media during the reboot.

`other-routing-engine`—On routers with dual Routing Engines, use the Junos OS image on the other Routing Engine during the reboot. That Routing Engine assumes the primary role; in the usual configuration, the other Routing Engine is the designated backup Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- When a Software Process Fails
- processes

failover-within-preference (L2TP LAC)

IN THIS SECTION

- Syntax | 6786
- Hierarchy Level | 6786
- Description | 6786
- Required Privilege Level | 6787
- Release Information | 6787

Syntax

```
failover-within-preference;
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Enable L2TP LAC tunnel selection within a preference level. When the router is unable to connect to a destination at a given preference level, it attempts to connect to another destination at the same level.

By default, when a connection attempt fails at one preference level, the next attempt is made at the next lower level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

L2TP LAC Tunneling for Subscribers

failure-action

IN THIS SECTION

- Syntax | 6787
- Hierarchy Level | 6788
- Description | 6788
- Options | 6788
- Required Privilege Level | 6788
- Release Information | 6788

Syntax

```
failure-action (clear-binding | clear-binding-if-interface-up | log-only);
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection],
[edit system services dhcp-local-server dhcpv6 liveness-detection],
[edit forwarding-options dhcp-relay liveness-detection],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection],
[edit system services dhcp-local-server group group-name liveness-detection],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection],
[edit forwarding-options dhcp-relay group group-name liveness-detection],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection]
```

Description

Configure the action the router (or switch) takes when a liveness detection failure occurs.

Options

- **Default:** `clear-binding`

`clear-binding`—The DHCP client session is cleared when a liveness detection failure occurs, except when `maintain-subscribers interface-delete` setting is configured and active.

`clear-binding-if-interface-up`—The DHCP client session is cleared only when a liveness detection failure occurs and the local interface is detected as being up. Use this setting to distinguish failures from between a liveness detection failure due to a local network error, and a host disconnecting from the network. If the client binding is in the maintain-binding Finite State Machine (FSM) state when the liveness detection failure detection occurs, then the binding is not deleted. Not supported for Layer 2 ARP/ND liveness detection on MX Series routers.

`log-only`—A message is logged to indicate the event; no action is taken and DHCP is left to manage the failure and maintain the client binding. Not supported for Layer 2 ARP/ND liveness detection on MX Series routers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- DHCP Liveness Detection Overview
- Configuring Detection of DHCP Local Server Client Connectivity with BFD
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD
- Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients
- Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

failure-action (Protocols LDP)

IN THIS SECTION

- Syntax | 6789
- Hierarchy Level | 6789
- Description | 6790
- Options | 6790
- Required Privilege Level | 6790
- Release Information | 6790

Syntax

```
failure-action {
    remove-nexthop;
    remove-route;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp oam bfd-livenesss-detection],
[edit logical-systems logical-system-name protocols ldp oam fec address bfd-livenesss-detection],
[edit protocols ldp oam bfd-livenesss-detection],
[edit protocols ldp oam fec address bfd-livenesss-detection]
```

Description

Configure route and next-hop properties in the event of a BFD session failure event on an LDP LSP. The failure event could be an existing BFD session that has gone down or could be a BFD session that never came up. LDP adds back the route or next hop when the relevant BFD session comes back up.

Options

`remove-nexthop`—Remove a route corresponding to a next hop of the LSP's route at the ingress node when a BFD session failure event is detected.

`remove-route`—Remove the route corresponding to an LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring a Failure Action for the BFD Session on an LDP LSP

failure-action (Protocols MPLS)

IN THIS SECTION

- Syntax | 6791
- Hierarchy Level | 6791
- Description | 6791

- Options | 6791
- Required Privilege Level | 6792
- Release Information | 6792

Syntax

```
failure-action {
    make-before-break teardown-timeout seconds;
    teardown;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls oam bfd-liveness-detection],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name oam bfd-
liveness-detection],
[edit protocols mpls label-switched-path lsp-name oam bfd-liveness-detection],
[edit protocols mpls oam bfd-liveness-detection]
```

Description

Configure route and next-hop properties in the event of a Bidirectional Forwarding Detection (BFD) protocol session failure event on an RSVP label-switched path (LSP). The failure event could be an existing BFD session that has gone down or a BFD session that never came up. RSVP adds back the route or next hop when the relevant BFD session comes back up.

Options

`make-before-break`—When a BFD session fails for an RSVP LSP, an attempt is made to signal a new LSP path before tearing down the old LSP path.

`teardown`—When a BFD session fails for an RSVP LSP, the associated LSP path is taken down and resignaled immediately.

`teardown-timeout` *seconds*—When you configure the `make-before-break` option, you can specify a time in seconds for the `teardown-timeout` option. At the end of the time specified, the associated RSVP LSP is automatically torn down and resignaled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring a Failure Action for the BFD Session on an RSVP LSP

fallback-block (Security Antivirus)

IN THIS SECTION

- Syntax | 6792
- Hierarchy Level | 6793
- Description | 6793
- Options | 6793
- Required Privilege Level | 6793
- Release Information | 6793

Syntax

```
fallback-block {
    administrator-email email-address;
    allow-email;
```

```
    custom-message message;
    custom-message-subject message-subject;
    display-host;
    (notify-mail-sender | no-notify-mail-sender);
     type (message | protocol-only);
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options]
```

Description

Configure notifications for fallback blocking actions. Fallback options tell the system how to handle the errors returned by either the scan engine or the scan manager.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

fallback-key

IN THIS SECTION

- Syntax | 6794
- Hierarchy Level | 6794
- Description | 6794
- Default | 6795
- Options | 6795
- Required Privilege Level | 6795
- Release Information | 6795

Syntax

```
fallback-key {
    cak hexadecimal-number;
    ckn hexadecimal-number;
}
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Specifies the fallback preshared key (PSK) to be used to enable MACsec using static connectivity association key (CAK) security mode. You can configure a fallback PSK to prevent traffic loss in case the primary PSK fails to establish a connection.

When you enable MACsec using static CAK security mode, a preshared PSK is exchanged between the devices on each end of the point-to-point Ethernet link. The PSK is includes a connectivity association name (CKN) and a connectivity association key (CAK). The PSK must match across devices for a MACsec session to be established. If there is a mismatch, the session will not be established and all packets will

be dropped. The fallback PSK is used when the primary PSK does not match for the initial MACsec negotiation.

Default

Fallback PSK is not enabled by default.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Configuring Advanced MACsec Features

fallback-non-block (Security Antivirus)

IN THIS SECTION

- Syntax | 6796
- Hierarchy Level | 6796
- Description | 6796
- Options | 6796
- Required Privilege Level | 6796
- Release Information | 6796

Syntax

```
fallback-non-block {
    custom-message message;
    custom-message-subject message-subject;
    (notify-mail-recipient | no-notify-mail-recipient);
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options]
```

Description

Configure notifications for fallback nonblocking actions.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

fallback-option (ICAP Redirect Service)

IN THIS SECTION

- Syntax | 6797
- Hierarchy Level | 6797
- Description | 6797
- Options | 6798
- Required Privilege Level | 6798
- Release Information | 6799

Syntax

```
fallback-option {
    connectivity (block | log-permit | permit);
    default-action (block | log-permit | permit);
    timeout (block | log-permit | permit);
}
```

Hierarchy Level

```
[edit services icap-redirect profile name],
[edit logical-systems logical-system-name services icap-redirect profile name]
```

Description

Specify fallback options for the device. Fallback settings enable the device to handle errors.

The fallback option is used to define the actions such as permit, log-and-permit, or block. This is the action that occurs when a request fails due to conditions such as too many requests, or a timeout occurred, or connectivity issues.

Options

connectivity Fallback settings when connection-related failures occur.

- Values:
  - block—Log the error and deny the requests.
  - log-permit—Log the error and permit the requests.
  - permit—Permit the requests.

default-action Default failure action.

- Values:
  - block—Log the error and deny the requests.
  - log-permit—Log the error and permit the requests.
  - permit—Permit the requests.

throttle Fallback action when the total number of requests received concurrently exceeds the devices limit.

- Values:
  - block—Log the error and deny the requests.
  - log-permit—Log the error and permit the requests.
  - permit—Permit the requests.

timeout Fallback action when there is a timeout occurrence.

- Values:
  - block—Log the error and deny the requests.
  - log-permit—Log the error and permit the requests.
  - permit—Permit the requests.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.1R1.

The logical system option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Example: Configuring ICAP Redirect Service on SRX Series Firewalls

fallback-options (Security Antivirus Juniper Express Engine)

IN THIS SECTION

- Syntax | 6799
- Hierarchy Level | 6800
- Description | 6800
- Options | 6800
- Required Privilege Level | 6800
- Release Information | 6800

Syntax

```
fallback-options {
   content-size (block | log-and-permit);
   default (block | log-and-permit);
   engine-not-ready (block | log-and-permit);
   out-of-resources (block | (log-and-permit);
   timeout (block | log-and-permit);
   too-many-requests (block | log-and-permit);
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name]
```

Description

Fallback options tell the system how to handle the errors.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Express Antivirus Configuration Overview

fallback-options (Security Antivirus Kaspersky Lab Engine)

IN THIS SECTION

Syntax | 6801

- Hierarchy Level | 6801
- Description | 6801
- Options | 6801
- Required Privilege Level | 6802
- Release Information | 6802

Syntax

```
fallback-options {
   content-size (block | log-and-permit);
   corrupt-file (block | log-and-permit);
   decompress-layer (block | log-and-permit);
   default (block | log-and-permit);
   engine-not-ready (block | log-and-permit);
   out-of-resources (block | (log-and-permit);
   password-file (block | (log-and-permit);
   timeout (block | log-and-permit);
   too-many-requests (block | log-and-permit);
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name]
```

Description

Fallback options tell the system how to handle the errors returned by either the scan engine or the scan manager.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Kaspersky feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Full Antivirus Configuration Overview

fallback-options (Security Antivirus Sophos Engine)

IN THIS SECTION

- Syntax | 6802
- Hierarchy Level | 6803
- Description | 6803
- Options | 6803
- Required Privilege Level | 6803
- Release Information | 6803

Syntax

```
fallback-options {
    content-size (block | log-and-permit | permit);
    default (block | log-and-permit | permit);
    engine-not-ready (block | log-and-permit | permit);
```

```
    out-of-resources (block | log-and-permit | permit);
    timeout (block | log-and-permit | permit);
    too-many-requests (block | log-and-permit | permit);
    server-connection-err (block | log-and-permit | permit);
}
```

Hierarchy Level

```
[edit security utm default-configuration anti-virus]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name]
```

Description

Configure fallback options to instruct the system how to handle errors.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration anti-virus]` hierarchy level is introduced in Junos OS Release 18.2R1.

The `server-connection-err` option is introduced in Junos OS Release 23.1R2.

RELATED DOCUMENTATION

Sophos Antivirus Configuration Overview

fallback-settings (Security Web Filtering)

IN THIS SECTION

- Syntax | 6804
- Hierarchy Level | 6804
- Description | 6804
- Options | 6804
- Required Privilege Level | 6805
- Release Information | 6805

Syntax

```
fallback-settings {
    default  (block | log-and-permit);
    server-connectivity   (block | log-and-permit);
    timeout    (block | log-and-permit);
    too-many-requests  (block | log-and-permit);
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated profile profile-name]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name]
```

Description

Fallback settings tell the system how to handle errors.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Surf-Control feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, Statement introduced in Junos OS Release 9.5 .

Statement introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

fallback-settings (Security Web Filtering Juniper Local)

IN THIS SECTION

- Syntax | 6805
- Hierarchy Level | 6806
- Description | 6806
- Options | 6806
- Required Privilege Level | 6806
- Release Information | 6806

Syntax

```
fallback-settings {
    default (block | log-and-permit);
    server-connectivity (block | log-and-permit);
    timeout  (block | log-and-permit);
    too-many-requests  (block | log-and-permit);
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering juniper-local profile profile-name]
```

Description

Fallback settings tell the system how to handle errors.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Web Filtering Overview

fallback-settings (Security Web Filtering Websense Redirect)

IN THIS SECTION

- Syntax | 6807
- Hierarchy Level | 6807

- Description | 6807
- Options | 6807
- Required Privilege Level | 6807
- Release Information | 6808

Syntax

```
fallback-settings {
    default  (block | log-and-permit);
    server-connectivity  (block | log-and-permit);
    timeout   (block | log-and-permit);
    too-many-requests  (block | log-and-permit);
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name]
```

Description

Fallback settings tell the system how to handle errors.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Understanding Redirect Web Filtering

falling-event-index

IN THIS SECTION

- Syntax | 6808
- Hierarchy Level | 6808
- Description | 6809
- Options | 6809
- Required Privilege Level | 6809
- Release Information | 6809

Syntax

```
falling-event-index index;
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

The index of the event entry that is used when a falling threshold is crossed. If this value is zero, no event is triggered.

Options

index—Index of the event entry that is used when a falling threshold is crossed.

- **Range:** 0 through 65,535
- **Default:** 0

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Falling Event Index or Rising Event Index

rising-event-index | **15805**

falling-event-index (RMON)

IN THIS SECTION

- Syntax | **6810**
- Hierarchy Level | **6810**
- Description | **6810**
- Options | **6810**

Required Privilege Level | 6810
Release Information | 6810

Syntax

```
falling-event-index index;
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Set the index number of the event entry that is used when a falling threshold is crossed. You specify the falling-event index when you configure an SNMP RMON alarm. If this value is zero, no event is triggered.

Options

index—Index of the event entry that is used when a falling threshold is crossed.

- **Range:** 0 through 65,535
- **Default:** 0

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

RMON MIB Event, Alarm, Log, and History Control Tables

Configuring RMON Alarms and Events

Monitoring RMON MIB Tables

Understanding RMON

Junos OS Network Management Configuration Guide

falling-threshold (Health Monitor)

IN THIS SECTION

- Syntax | 6811
- Hierarchy Level | 6811
- Description | 6811
- Options | 6812
- Required Privilege Level | 6812
- Release Information | 6812

Syntax

```
falling-threshold percentage;
```

Hierarchy Level

```
[edit snmp health-monitor]
```

Description

Set the lower threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.

Options

percentage—Lower threshold for the alarm entry.

- **Range:** 1 through 100
- **Default:** 70 percent of the maximum possible value

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D20.

RELATED DOCUMENTATION

rising-threshold

Configuring Health Monitoring

falling-threshold (integer)

IN THIS SECTION

- Syntax | **6813**
- Hierarchy Level | **6813**
- Description | **6813**
- Options | **6813**
- Required Privilege Level | **6813**
- Release Information | **6813**

Syntax

```
falling-threshold integer;
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Set the lower threshold for the sampled variable (monitored object) when you configure an SNMP RMON alarm. By setting a rising and a falling threshold for a variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.

Options

integer—The lower threshold for the alarm entry.

- **Range:** -2,147,483,648 through 2,147,483,647
- **Default:** 20 percent less than `rising-threshold`

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring the Falling Threshold or Rising Threshold*
>
> rising-threshold (RMON) | **15808**

falling-threshold-interval

IN THIS SECTION

- Syntax | 6814
- Hierarchy Level | 6814
- Description | 6814
- Options | 6814
- Required Privilege Level | 6815
- Release Information | 6815

Syntax

```
falling-threshold-interval seconds;
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Set the interval between samples after the rising threshold is exceeded and the value of the sample starts to drop. If the value of the sample drops and exceeds the falling threshold, the regular sampling interval is used.

Options

seconds—Time between samples, in seconds.

- **Range:** 1 through 2,147,483,647 seconds
- **Default:** 60 seconds

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring the Falling Threshold Interval

interval (RMON) | 9206

falling-threshold (percentage)

IN THIS SECTION

- Syntax | 6815
- Hierarchy Level | 6816
- Description | 6816
- Options | 6816
- Required Privilege Level | 6816
- Release Information | 6816

Syntax

```
falling-threshold percentage;
```

Hierarchy Level

```
[edit snmp health-monitor],
[edit snmp health-monitor idp]
```

Description

Set the lower threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.

Options

percentage—The lower threshold for the alarm entry.

- **Range:** `1` through `100`
- **Default:** `70` percent of the maximum possible value

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

> *Configuring the Falling Threshold or Rising Threshold*
>
> rising-threshold (Health Monitor) | **15807**

falling-threshold (RMON)

IN THIS SECTION

- Syntax | 6817
- Hierarchy Level | 6817
- Description | 6817
- Options | 6817
- Required Privilege Level | 6818
- Release Information | 6818

Syntax

```
falling-threshold integer;
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Set the lower threshold for the sampled variable (monitored object) when you configure an SNMP RMON alarm. By setting a rising and a falling threshold for a variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.

Options

integer—Lower threshold for the alarm entry.

- **Range:** –2,147,483,648 through 2,147,483,647
- **Default:** 20 percent less than the `rising-threshold` value

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- *RMON MIB Event, Alarm, Log, and History Control Tables*
- *Configuring RMON Alarms and Events*
- *Monitoring RMON MIB Tables*
- *Understanding RMON*
- Junos OS Network Management Configuration Guide

false-positives

IN THIS SECTION

- Syntax | 6819
- Hierarchy Level | 6819
- Description | 6819
- Options | 6819
- Required Privilege Level | 6819
- Release Information | 6819

Syntax

```
false-positives {
    values [frequently occasionally rarely unknown];
}
```

Hierarchy Level

```
[edit security idp dynamic-attack-group dynamic-attack-group-name filters]
```

Description

Specify a false positives filter to track attack objects based on the frequency that the attack produces a false positive on your network.

Options

`values`—Name of the false positives filter. You can select from the following false positive frequency:

- `frequently`—Frequently track false positive occurrences.
- `occasionally`—Occasionally track false positive occurrences.
- `rarely`—Rarely track false positive occurrences.
- `unknown`—By default, all compound attack objects are set to Unknown. As you fine-tune IDP to your network traffic, you can change this setting to help you track false positives.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

family (RFC 2544 Benchmarking Tests)

IN THIS SECTION

- Syntax | 6820
- Junos OS Hierarchy Level | 6820
- Junos OS Evolved Hierarchy Level | 6820
- Description | 6820
- Options | 6821
- Required Privilege Level | 6821
- Release Information | 6821

Syntax

```
family (bridge | ccc | ethernet-switching | inet | vpls);
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Configure the address type family for the benchmarking test.

Options

bridge	(Junos OS) Run the test on a Layer 2 Ethernet line (E- Line) or an Ethernet LAN (E-LAN) service configured in a bridge domain. You can run the RFC2544-based benchmarking test only in the egress direction or the user-to-network interface (UNI) direction of an Ethernet line.
ccc	Run the test on a layer 2 VPN, circuit cross-connect (CCC), or Ethernet pseudowire service (such as EVPN-VPWS). You can run the RFC2544-based benchmarking test either in the egress or ingress direction.
ethernet-switching	(Junos OS Evolved) Run the test over a Layer 2 bridge service, EVPN-MPLS network, or a Virtual Private LAN Service (VPLS). You can run the RFC2544-based benchmarking test only in the egress direction or the user-to-network interface (UNI) direction of an Ethernet line.
inet	Run the test on an IPv4 service.
vpls	(Junos OS) Run the test over a Virtual Private LAN Service (VPLS).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

`bridge` option introduced in Junos OS Release 12.3X53 for ACX Series routers.

`bridge` option introduced in Junos OS Release 14.2 for MX104 Universal Routing Platforms.

Statement introduced in Junos OS Evolved Release 21.1R1 for the `inet` option only.

`ccc` and `ethernet-switching` options introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

rfc2544-benchmarking | **15759**

rfc2544 | **15757**

RFC 2544-Based Benchmarking Tests Overview

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

family (logical-interface)

IN THIS SECTION

- Syntax | 6822
- Hierarchy Level | 6822
- Description | 6822
- Options | 6822
- Required Privilege Level | 6823
- Release Information | 6823

Syntax

```
family family {
    ipsec-sa sa-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Configure protocol family information for the logical interface.

Options

family—Protocol family:

- `ccc`—Circuit cross-connect protocol suite
- `inet`—IP version 4 suite
- `inet6`—IP version 6 suite
- `iso`—Open Systems Interconnection (OSI) International Organization for Standardization (ISO) protocol suite
- `mlfr-end-to-end`—Multilink Frame Relay FRF.15
- `mlfr-uni-nni`—Multilink Frame Relay FRF.16
- `multilink-ppp`—Multilink Point-to-Point Protocol
- `mpls`—MPLS
- `tcc`—Translational cross-connect protocol suite
- `tnp`—Trivial Network Protocol
- `vpls`—Virtual private LAN service

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring Encryption Interfaces*
>
> Junos OS Network Interfaces Library for Routing Devices

family (ELS)

IN THIS SECTION

- Syntax | 6824
- Hierarchy Level | 6828
- Description | 6828
- Default | 6828
- Options | 6828
- Required Privilege Level | 6828
- Release Information | 6829

Syntax

```
family  {
    ethernet-switching {
        filter {
            group filter-group-number;
            input filter-name;
            input-list [ filter-names ];
            output filter-name;
            output-list [ filter-names ];
        }
        interface-mode (access | trunk);
        recovery-timeout seconds;
        storm-control profile-name;
        vlan {
            members (vlan-name |[-vlan-names] | all);
        }
    }
    fibre-channel {
        port-mode (f-port | np-port);
    }
    inet {
        accounting {
            destination-class-usage;
```

```
            source-class-usage {
                input;
                output;
            }
        }
        address ipv4-address {
            arp ip-address (mac | multicast-mac) mac-address <publish>;
            broadcast address;
            preferred;
            primary;
            vrrp-group group-number {
                (accept-data | no-accept-data);
                advertise-interval seconds;
                advertisements-threshold number;
                authentication-key key;
                authentication-type authentication;
                fast-interval milliseconds;
                (preempt | no-preempt) {
                    hold-time seconds;
                }
                priority number;
                track {
                    interface interface-name {
                        priority-cost number;
                    }
                    priority-hold-time seconds;
                    route ip-address/mask routing-instance instance-name priority-cost cost;
                }
                virtual-address [addresses];
                vrrp-inherit-from {
                    active-group group-number;
                    active-interface interface-name;
                }
            }
        }
        filter {
            group filter-group-number;
            input filter-name;
            input-list [ filter-names ];
            output filter-name;
            output-list [ filter-names ];
        }
        mtu bytes;
```

```
        no-neighbor-learn;
        no-redirects;
        primary;
        rpf-check {
            fail-filter filter-name;
            mode {
                loose;
            }
        }
    }
    inet6 {
        accounting {
            destination-class-usage;
            source-class-usage {
                input;
                output;
            }
        }
        address address {
            eui-64;
            ndp ip-address (mac | multicast-mac) mac-address <publish>;
            preferred;
            primary;
            vrrp-inet6-group group-id {
                accept-data | no-accept-data;
                advertisements-threshold number;
                authentication-key key;
                authentication-type authentication;
                fast-interval milliseconds;
                inet6-advertise-interval milliseconds;
                preempt | no-preempt {
                    hold-time seconds;
                }
                priority number;
                track {
                    interface interface-name {
                        priority-cost number;
                    }
                    priority-hold-time seconds;
                    route ip-address/mask routing-instance instance-name priority-cost cost;
                }
                virtual-inet6-address [addresses];
                virtual-link-local-address ipv6-address;
```

```
                vrrp-inherit-from {
                    active-group group-name;
                    active-interface interface-name;
                }
            }
        }
        (dad-disable | no-dad-disable);
        filter {
            group filter-group-number;
            input filter-name;
            input-list [ filter-names ];
            output filter-name;
            output-list [ filter-names ];
        }
        mtu bytes;
        nd6-stale-time time;
        no-neighbor-learn;
        no-redirects;
        policer {
            input policer-name;
            output policer-name;
        }
        rpf-check {
            fail-filter filter-name;
            mode {
                loose;
            }
        }
        mpls {
            filter {
                group filter-group-number;
                input filter-name;
                input-list [ filter-names ];
                output filter-name;
                output-list [ filter-names ];
            }
        mtu bytes;
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-range interface-name unit logical-unit-number family]
```

Description

Configure protocol family information for the logical interface on QFX Series switches.

Default

Access interfaces on the QFX Series are set to **family ethernet-switching** by default. If you are going to change the family setting for an interface, you might have to delete this default setting or any user-configured family setting first.

You must configure a logical interface to be able to use the physical device.

Options

Interface types on the switch are:

- Aggregated Ethernet (**ae**)
- Gigabit Ethernet (**ge**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Routed VLAN interface (RVI) (**vlan**)
- 10-Gigabit Ethernet (**xe**)

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches*
- *Configuring Link Aggregation*
- Configuring IRB Interfaces on Switches
- Junos OS Network Interfaces Library for Routing Devices

family (QFX)

IN THIS SECTION

- Syntax | 6829
- Hierarchy Level | 6830
- Description | 6830
- Options | 6830
- Required Privilege Level | 6830
- Release Information | 6831

Syntax

```
family family-name {
    filter filter-name {
        interface-specific;
        micro-segmentation;
        term term-name {
            from {
                match-conditions;
                egress-to-ingress;
            }
```

```
            then {
                action;
                action-modifiers;
            }
        }
    }
}
```

Hierarchy Level

```
[edit firewall]
```

Description

Configure the fields a firewall filter can match on.

Options

family-name—Type of addressing protocol:

- `ethernet-switching`—Filter Layer 2 Ethernet packets and Layer 3 (IP) packets (allows some Layer 3 filtering). Not supported on OCX Series switches.
- `egress-to-ingress`—Include this option to increase the number of egress VLAN firewall filter terms from 1024 to 2048.
- `evpn`—Filter Ethernet VPN (EVPN) packets.
- `inet`—Filter Layer 3 IPv4 packets (provides additional Layer 3 filter options).
- `inet6`—Filter Layer 3 IPv6 packets (provides additional Layer 3 filter options).
- `mpls`—Filter multiprotocol label switched packets. Not supported on OCX Series switches.
- `any`—Filter packets based on protocol-independent match conditions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

`evpn` options introduced in Junos OS Release 15.1 for the MX Series.

`egress-to-ingress` option introduced in Junos OS Release 19.1R1 for the QFX5110.

`any` option introduced in Junos OS Release 22.4R1 for supported QFX5120, EX4100, EX4400, and EX4650 Series switches.

micro-segmentation option introduced in Junos OS Release 22.4R1 for supported QFX5120, EX4100, EX4400, and EX4650 Series switches.

RELATED DOCUMENTATION

Firewall Filter Match Conditions and Actions (QFX5100, QFX5110, QFX5120, QFX5200, EX4600, EX4650)

Configuring Firewall Filters

Overview of Firewall Filters (QFX Series)

family (protocols ldp)

IN THIS SECTION

- Syntax | 6831
- Hierarchy Level | 6832
- Description | 6832
- Required Privilege Level | 6832
- Release Information | 6832

Syntax

```
family {
    inet;
```

```
    inet6;
}
```

Hierarchy Level

```
[edit protocols ldp]
```

Description

Configure the address family as inet for IPv4 or inet6 for IPv6, or both. If the address family is not configured, then the default address family is IPv4.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

LDP Native IPv6 Support Overview

Example: Configuring LDP Native IPv6 Support

Configuring LDP Native IPv6 Support

dual-transport | **6059**

family (interfaces)

IN THIS SECTION

Syntax | **6833**

- Hierarchy Level | 6836
- Description | 6836
- Options | 6836
- Required Privilege Level | 6838
- Release Information | 6838

Syntax

```
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            (input | output | input output);
        }
    }
    access-concentrator name;
    address address {
        ... the address subhierarchy appears after the main [edit interfaces interface-name unit
logical-unit-number family family-name] hierarchy ...
    }
    bundle interface-name;
    core-facing;
    demux-destination {
        destination-prefix;
    }
    demux-source {
        source-prefix;
    }
    direct-connect;
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
        group filter-group-number;
        input filter-name;
        input-list [ filter-names ];
        output filter-name;
        output-list [ filter-names ];
    }
```

```
    interface-mode (access | trunk);
    ipsec-sa sa-name;
    keep-address-and-control;
    mac-validate (loose | strict);
    max-sessions number;
    max-sessions-vsa-ignore;
    mtu bytes;
    multicast-only;
    nd6-stale-time seconds;
    negotiate-address;
    no-neighbor-learn;
    no-redirects;
    policer {
        arp policer-template-name;
        input policer-template-name;
        output policer-template-name;
    }
    primary;
    protocols [inet iso mpls];
    proxy inet-address address;
    receive-options-packets;
    receive-ttl-exceeded;
    remote (inet-address address | mac-address address);
    rpf-check {
        fail-filter filter-name
        mode loose;
    }
    sampling {
        input;
        output;
    }
    service {
        input {
            post-service-filter filter-name;
            service-set service-set-name <service-filter filter-name>;
        }
        output {
            service-set service-set-name <service-filter filter-name>;
        }
    }
    service-name-table table-name;
    short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>
<filter [aci]>;
```

```
    (translate-discard-eligible | no-translate-discard-eligible);
    (translate-fecn-and-becn | no-translate-fecn-and-becn);
    translate-plp-control-word-de;
    unnumbered-address interface-name destination address destination-profile profile-name;
    vlan-id number;
    vlan-id-list [number number-number];
    address address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        destination address;
        destination-profile name;
        eui-64;
        primary-only;
        multipoint-destination address dlci dlci-identifier;
        multipoint-destination address {
            epd-threshold cells;
            inverse-arp;
            oam-liveness {
                up-count cells;
                down-count cells;
            }
            oam-period (disable | seconds);
            shaping {
                (cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length
peak rate sustained rate);
                queue-length number;
            }
            vci vpi-identifier.vci-identifier;
        }
        preferred;
        primary;
        vrrp-groupgroup-id {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
```

```
                    bandwidth-threshold bits-per-second priority-cost priority;
                    priority-cost priority;
                }
                priority-hold-time seconds;
                route prefix routing-instance instance-name priority-cost priority;
                }
            }
            virtual-address [ addresses ];
            }
        virtual-link-local-address ipv6-address;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure protocol family information for the logical interface.

Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms. See Configuring VLAN TCC Encapsulation for more details.

NOTE: Not all subordinate statements are available to every protocol family.

Options

family—Protocol family:

- `any`—Protocol-independent family used for Layer 2 packet filtering

  NOTE: This option is not supported on T4000 Type 5 FPCs.

- `bridge`—(M Series and T Series routers only) Configure only when the physical interface is configured with `ethernet-bridge` type encapsulation or when the logical interface is configured with `vlan-bridge`

type encapsulation. You can optionally configure this protocol family for the logical interface on which you configure VPLS.

- `ethernet-switching`—(M Series and T Series routers only) Configure only when the physical interface is configured with `ethernet-bridge` type encapsulation or when the logical interface is configured with `vlan-bridge` type encapsulation
- `ccc`—Circuit cross-connect protocol suite. You can configure this protocol family for the logical interface of CCC physical interfaces. When you use this encapsulation type, you can configure the `ccc` family only.
- `inet`—Internet Protocol version 4 suite. You must configure this protocol family for the logical interface to support IP protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).
- `inet6`—Internet Protocol version 6 suite. You must configure this protocol family for the logical interface to support IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), BGP, and Virtual Router Redundancy Protocol for IPv6 (VRRP).
- `iso`—International Organization for Standardization Open Systems Interconnection (ISO OSI) protocol suite. You must configure this protocol family for the logical interface to support IS-IS traffic.
- `mlfr-end-to-end`—Multilink Frame Relay FRF.15. You must configure this protocol or multilink Point-to-Point Protocol (MLPPP) for the logical interface to support multilink bundling.
- `mlfr-uni-nni`—Multilink Frame Relay FRF.16. You must configure this protocol or `mlfr-end-to-end` for the logical interface to support link services and voice services bundling.
- `multilink-ppp`—Multilink Point-to-Point Protocol. You must configure this protocol (or `mlfr-end-to-end`) for the logical interface to support multilink bundling.
- `mpls`—Multiprotocol Label Switching (MPLS). You must configure this protocol family for the logical interface to participate in an MPLS path.
- `pppoe`—Point-to-Point Protocol over Ethernet
- `tcc`—Translational cross-connect protocol suite. You can configure this protocol family for the logical interface of TCC physical interfaces.
- `tnp`—Trivial Network Protocol. This protocol is used to communicate between the Routing Engine and the router's packet forwarding components. The Junos OS automatically configures this protocol family on the router's internal interfaces only, as discussed in Understanding Internal Ethernet Interfaces.

- `vpls`—(M Series and T Series routers only) Virtual private LAN service. You can optionally configure this protocol family for the logical interface on which you configure VPLS. VPLS provides an Ethernet-based point-to-multipoint Layer 2 VPN to connect customer edge (CE) routers across an MPLS backbone. When you configure a VPLS encapsulation type, the `family vpls` statement is assumed by default.

  MX Series routers support dynamic profiles for VPLS pseudowires, VLAN identifier translation, and automatic bridge domain configuration.

  For more information about VPLS, see the Junos OS VPNs Library for Routing Devices.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Option `max-sessions-vsa-ignore` introduced in Junos OS Release 11.4.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
20.1R1	Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms.

RELATED DOCUMENTATION

Configuring the Protocol Family

family (Link Interfaces)

IN THIS SECTION

- Syntax | 6839
- Hierarchy Level | 6839
- Description | 6839
- Options | 6839
- Required Privilege Level | 6840
- Release Information | 6840

Syntax

```
family family {
    address address {
        destination address;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Configure protocol family information for the logical interface.

Options

family—Protocol family:

- `ccc`—Circuit cross-connect protocol suite
- `inet`—IP version 4 (IPv4)

- `inet6`—IP version 6 (IPv6)
- `iso`—Open Systems Interconnection (OSI) International Organization for Standardization (ISO) protocol suite
- `mlfr-end-to-end`—Multilink Frame Relay FRF.15
- `mlfr-uni-nni`—Multilink Frame Relay FRF.16
- `multilink-ppp`—Multilink Point-to-Point Protocol
- `mpls`—MPLS
- `tcc`—Translational cross-connect protocol suite
- `tnp`—Trivial Network Protocol
- `vpls`—Virtual private LAN service

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> Link and Multilink Services Interfaces User Guide for Routing Devices
>
> Junos OS Network Interfaces Library for Routing Devices

family (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 6841
- Hierarchy Level | 6842
- Description | 6842
- Options | 6842
- Required Privilege Level | 6842
- Release Information | 6842

Syntax

```
family family {
    dhcp-attributes {
        [protocol-specific attributes]
    }
    excluded-address ip-address;
    excluded-range name low minimum-value high maximum-value;
    host hostname {
        hardware-address mac-address;
        ip-address ip-address;
    }
    network ip-prefix/<prefix-length>;
    prefix ipv6-prefix;
    range range-name {
        high upper-limit;
        low lower-limit;
        prefix-length prefix-length;
    }
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name]
```

Description

Configure the protocol family for the address-assignment pool.

> **NOTE**: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

family—Protocol family:

- `inet`—Internet Protocol version 4 suite
- `inet6`—Internet Protocol version 6 suite

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools Overview

Address-Assignment Pool Configuration Overview

family (Address Pools)

IN THIS SECTION

- Syntax | 6843
- Hierarchy Level | 6843
- Description | 6843
- Options | 6844
- Required Privilege Level | 6844
- Release Information | 6844

Syntax

```
family (inet | inet6) {
    network {
        [network-prefix] {
            external-assigned;
        }
    }
}
```

Hierarchy Level

```
[edit access address-assignment address-pools name]
```

Description

Specify the protocol family information for the address pool. Address pools must have either `inet` (IPv4) or `inet6` (IPv6) configured.

NOTE: A address pool can have either `inet` (IPv4) or `inet6` (IPv6) configured, but not both.

Options

inet	IP version 4 (IPv4) suite.
inet6	IP version 6 (IPv6) suite.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Understanding Source IP Filtering with Address Pools in TDF Domains for IP-Based Subscribers

Configuring Address Pools for Source-IP Filtering of IP-Based Subscribers

family (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 6845
- Hierarchy Level | 6845
- Description | 6845
- Options | 6845
- Required Privilege Level | 6845
- Release Information | 6845

Syntax

```
family family;
```

Hierarchy Level

```
[edit interfaces interface-name unit interface-unit-number]
```

Description

Configure protocol family information for the logical interface.

Options

family—Protocol family. Currently, only one option, `inet` (IP version 4 suite), is supported.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

unit (Aggregated Multiservices) | **30380**

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

family (Bootstrap)

IN THIS SECTION

- Syntax | 6846
- Hierarchy Level | 6846
- Description | 6846
- Options | 6846
- Required Privilege Level | 6847
- Release Information | 6847

Syntax

```
family (inet | inet6) {
    export [ policy-names ];
    import [ policy-names ];
    priority number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp bootstrap],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp bootstrap],
[edit protocols pim rp bootstrap],
[edit routing-instances routing-instance-name protocols pim rp bootstrap]
```

Description

Configure which IP protocol type bootstrap properties to apply.

Options

inet—Apply IP version 4 (IPv4) local RP properties.

inet6—Apply IPv6 local RP properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring PIM Bootstrap Properties for IPv4

Configuring PIM Bootstrap Properties for IPv4 or IPv6

family bridge

IN THIS SECTION

- Syntax | 6847
- Hierarchy Level | 6848
- Description | 6848
- Options | 6848
- Required Privilege Level | 6848
- Release Information | 6848

Syntax

```
family bridge {
    apply-groups value;
```

```
    apply-groups-except value;
    core-facing;
    interface-mode access | trunk;
    inner-vlan-id-list inner-vlan-id-range;
    storm-control;
    vlan-id vlan-id ;
    vlan-id-list vlan-id-range;
}
```

Hierarchy Level

```
[edit interfaces]
[edit logical-systems name interfaces]
```

Description

Family bridge is used when you want a port that has more than one logical unit, each with the same or different encapsulations. Bridge domains are associated with GRE interface with the corresponding BD VLAN.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

Starting in Junos OS release 17.4R1 for MX Series routers, support for storm control was added for logical systems.

RELATED DOCUMENTATION

Tunnel Services Overview

Tunnel Interface Configuration on MX Series Routers Overview

family bridge (GRE Interfaces)

IN THIS SECTION

- Syntax | 6849
- Hierarchy Level | 6849
- Description | 6849
- Options | 6850
- Required Privilege Level | 6850
- Release Information | 6850

Syntax

```
family bridge {
    interface-mode (access | trunk);
    core-facing;
    vlan-id number;
    vlan-id-list vlan-list;
}
```

Hierarchy Level

```
[edit interfaces gr-fpc/pic/port.unit]
```

Description

Configure the bridge domain family on GRE interfaces. To enable Layer 2 Ethernet packets to be terminated on GRE tunnels, you must configure the bridge domain protocol family on the `gr-` interfaces

and associate the `gr-` interfaces with the bridge domain. You must configure the GRE interfaces as core-facing interfaces, and they must be access or trunk interfaces. To configure the bridge domain family on `gr-` interfaces, include the `family bridge` statement at the `[edit interfaces gr-`*`fpc/pic/port`* `unit` *`logical-unit-number`*`]` hierarchy level.

Options

`interface-mode`—Specify the type of VLAN tagging on packets that the interface accepts.

`access`—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the `vlan-id` statement.

`trunk`—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the `vlan-id` or `vlan-id-list` statement.

`core-facing`—Specify that the VLAN is physically connected to a core-facing ISP router and ensures that the network does not improperly treat the interface as a client interface. When specified, the interface is inserted into the core-facing default mesh group where traffic from pseudowires that belong to the default mesh group is not forwarded on the core-facing link.

`vlan-id` *`number`*—Individual VLAN IDs separated by a space.

`vlan-id-list` *`vlan-list`*—Starting VLAN ID and ending VLAN ID in an inclusive range. Separate the starting VLAN ID and ending VLAN ID with a hyphen.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement at the `[edit interfaces gr-fpc/pic/port.unit]` hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

RELATED DOCUMENTATION

Layer 2 Services over GRE Tunnel Interfaces on MX Series with MPCs

family (chassis)

IN THIS SECTION

- Syntax | 6851
- Hierarchy Level | 6852
- Description | 6852
- Options | 6852
- Required Privilege Level | 6852
- Release Information | 6852

Syntax

```
family {
    inet {
        layer-3;
        layer-4;
        symmetric-hash {
            complement;
        }
    }
    multiservice {
        source-mac;
        destination-mac;
        payload {
            ip {
                layer-3;
                layer-4;
            }
        }
        symmetric-hash {
            complement;
        }
    }
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number hash-key]
```

Description

(MX Series 5G Universal Routing Platforms only) Configure data used in a hash key for a specific protocol family when configuring PIC-level symmetrical load balancing on an 802.3ad Link Aggregation Group.

Options

inet Configure data used in a hash key for the **inet** protocol family.

multiservice Configure data used in a hash key for the **multiservice** protocol family.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

> *Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers*

family (Dynamic Demux Interface)

IN THIS SECTION

- Syntax | 6853
- Hierarchy Level | 6854
- Description | 6854
- Options | 6854
- Required Privilege Level | 6854
- Release Information | 6854

Syntax

```
family family {
    access-concentrator name;
    address address;
    demux-source {
        source-address;
    }
    direct-connect;
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
        input filter-name;
        output filter-name;
    }
    mac-validate (loose | strict);
    max-sessions number;
    max-sessions-vsa-ignore;
    service-name-table table-name;
    short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>
<filter [aci]>;
    unnumbered-address interface-name <preferred-source-address address>;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number]
```

Description

Configure protocol family information for the logical interface.

NOTE: Not all subordinate stanzas are available to every protocol family.

Options

family—Protocol family:

- inet—Internet Protocol version 4 suite
- inet6—Internet Protocol version 6 suite
- pppoe—(MX Series routers with MPCs only) Point-to-Point Protocol over Ethernet

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

pppoe option added in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles

Subscriber Interfaces and Demultiplexing Overview

family (Dynamic Firewalls)

IN THIS SECTION

- Syntax | 6855
- Hierarchy Level | 6856
- Description | 6856
- Options | 6856
- Required Privilege Level | 6856
- Release Information | 6856

Syntax

```
family family {
    fast-update-filter filter-name {
        interface-specific;
        match-order [match-order];
        term term-name {
            from {
                match-conditions;
            }
            then {
                action;
                action-modifiers;
            }
            only-at-create;
        }
    }
    filter filter-name {
        enhanced-mode-override;
        instance-shared;
        interface-shared;
        interface-specific;
        term term-name {
            from {
                match-conditions;
            }
```

```
            then {
                action;
                action-modifiers;
            }
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall]
```

Description

Configure fast update filters or parameterized filters for a protocol family in a dynamic client profile or a dynamic service profile.

Options

family—Protocol family:

- `any`—Filter packets based on protocol-independent match conditions.
- `inet`—Filter Internet Protocol version 4 suite packets.
- `inet6`—filter Internet Protocol version 6 suite packets.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

`any` option added in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Fast Update Filters

Parameterized Filters Overview

Guidelines for Configuring Firewall Filters

Configuring Unique Identifiers for Parameterized Filters

family (Dynamic PPPoE)

IN THIS SECTION

- Syntax | 6857
- Hierarchy Level | 6858
- Description | 6858
- Options | 6858
- Required Privilege Level | 6858
- Release Information | 6858

Syntax

```
family family {
    unnumbered-address interface-name;
    address address;
    service {
        input {
            service-set service-set-name {
                service-filter filter-name;
            }
            post-service-filter filter-name;
        }
        output {
            service-set service-set-name {
                service-filter filter-name;
            }
        }
```

```
    }
    filter {
        input filter-name {
            precedence precedence;
        }
        output filter-name {
            precedence precedence;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"]
```

Description

Configure protocol family information for the logical interface.

Options

family—Protocol family:

- `inet`—Internet Protocol version 4 suite
- `inet6`—Internet Protocol version 6 suite

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring a PPPoE Dynamic Profile

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

family (Dynamic Standard Interface)

IN THIS SECTION

- Syntax | 6859
- Hierarchy Level | 6860
- Description | 6860
- Options | 6861
- Required Privilege Level | 6861
- Release Information | 6861

Syntax

```
family family {
    access-concentrator name;
    address address;
    direct-connect;
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
        adf {
            counter;
            input-precedence precedence;
            not-mandatory;
            output-precedence precedence;
            rule rule-value;
        }
        input filter-name {
            precedence precedence;
            shared-name filter-shared-name;
```

```
        }
        output filter-name {
            precedence precedence;
            shared-name filter-shared-name;
        }
    }
    mac-validate (loose | strict);
    max-sessions number;
    max-sessions-vsa-ignore;
    rpf-check {
        fail-filter filter-name;
        mode loose;
    }
    service {
         input {
             service-set service-set-name {
                service-filter filter-name;
            }
            post-service-filter filter-name;
        }
        output {
             service-set service-set-name {
                service-filter filter-name;
            }
        }
    }
    service-name-table table-name;
    short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>
<filter [aci]>;
    unnumbered-address interface-name <preferred-source-address address>;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

Description

Configure protocol family information for the logical interface.

NOTE: Not all subordinate stanzas are available to every protocol family.

Options

family—Protocol family:

- `inet`—IP version 4 suite
- `inet6`—IP version 6 suite
- `pppoe`—(MX Series routers with MPCs only) Point-to-Point Protocol over Ethernet
- `vpls`—Virtual private LAN service

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

`pppoe` option added in Junos OS Release 11.2.

RELATED DOCUMENTATION

Static Routing on Logical Systems

Configuring the Protocol Family

family (Exclude Prefix)

IN THIS SECTION

- Syntax | 6862
- Hierarchy Level | 6862
- Description | 6862
- Required Privilege Level | 6863
- Release Information | 6863

Syntax

```
family {
    inet {
        apply-groups [group-names];
        apply-groups-except [group-names];
        network address net-mask;
    }
    inet6 {
        apply-groups [group-names];
        apply-groups-except [group-names];
        network address net-mask;
    }
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name subscriber-exclude-prefix]
```

Description

Specify the IP version for the network prefix of source IP addresses for uplink packets and destination IP addresses for downlink packets that must not undergo TDF processing.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based and IFL-Based TDF Subscribers Overview

IP-Based Subscriber Setup Overview

family fcoe

IN THIS SECTION

- Syntax | 6864
- Hierarchy Level | 6864
- Description | 6864
- Options | 6864
- Required Privilege Level | 6865
- Release Information | 6865

Syntax

QFX Series Standalone Switches

```
family fcoe {
    oxid (enable | disable);
}
```

QFabric Systems

```
family fcoe {
    ethernet-interfaces {
        node-group (node-group-name | all) {
            oxid (enable | disable);
        }
    }
    fabric-interfaces {
        node-group (node-group-name | all) {
            oxid (enable | disable);
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options hash-key]
```

Description

Configure whether or not to use the originator exchange identifier (OxID) field for hash control for FCoE traffic load balancing.

Options

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Ethernet-interfaces and fabric-interfaces statements introduced in Junos OS Release 13.2X52-D10 for the QFX Series.

RELATED DOCUMENTATION

Enabling and Disabling CoS OxID Hash Control for FCoE Traffic on Standalone Switches

Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches

family (Filtering)

IN THIS SECTION

- Syntax | 6865
- Hierarchy Level | 6866
- Description | 6866
- Options | 6866
- Required Privilege Level | 6866
- Release Information | 6866

Syntax

```
family family-name {
    filter {
        input input-filter-name;
```

```
        output output-filter-name;
    }
    flood {
        input filter-name;
    }
    route-accounting;
    source-checking;
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Specify address family for filters.

Options

family-name—Address family. Specify **inet** for IP version 4 (IPv4), **inet6** for IP version 6 (IPv6), **mpls** for MPLS, or **vpls** for virtual private LAN service (VPLS).

> **NOTE**: In Junos OS Release 8.4 and later, the `output` statement is not valid at the `[edit forwarding-options family vpls filter]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

route-accounting option introduced in Junos OS Release 8.3; supported only with IPv6.

`source-checking` option introduced in Junos OS Release 12.3 on MX Series 5G Universal Routing Platforms; supported only with IPv6.

RELATED DOCUMENTATION

Applying Forwarding Table Filters

family (Firewall)

IN THIS SECTION

- Syntax | 6867
- Hierarchy Level | 6868
- Description | 6868
- Options | 6868
- Required Privilege Level | 6869
- Release Information | 6869

Syntax

```
family family-name {
    filter filter-name {
        accounting-profile name;
        enhanced-mode;
        interface-specific;
        physical-interface-filter;
    }
    prefix-action name {
        count;
        destination-prefix-length prefix-length;
        policer policer-name;
        source-prefix-length prefix-length;
        subnet-prefix-length prefix-length;
    }
```

```
    simple-filter filter-name {
        term term-name {
            from {
                match-conditions;
            }
            then {
                action;
                action-modifiers;
            }
        }
    }
}
```

Hierarchy Level

```
[edit firewall],
[edit logical-systems logical-system-name firewall]
```

Description

Configure a firewall filter for IP version 4 (IPv4) or IP version 6 (IPv6) traffic. Only on MX Series routers and EX Series switches, configure a firewall filter for Layer 2 traffic in a bridging environment.

Options

family-name—Version or type of addressing protocol:

- **any**—Protocol-independent match conditions.
- **bridge**—(MX Series routers only) Layer 2 packets that are part of bridging domain.
- **ethernet-switching**—(EX Series switches) Filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets.
- **ccc**—Layer 2 switching cross-connects.
- **inet**—IPv4 addressing protocol.
- **inet6**—IPv6 addressing protocol.
- **mpls**—MPLS.
- **vpls**—Virtual private LAN service (VPLS).

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: The packet lengths that a policer considers depends on the address family of the firewall filter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

`simple-filter` statement introduced in Junos OS Release 7.6.

any family type introduced in Junos OS Release 8.0.

bridge family type introduced in Junos OS Release 8.4 (MX Series routers only).

RELATED DOCUMENTATION

- *Guidelines for Configuring Firewall Filters*
- *Guidelines for Configuring Service Filters*
- *Guidelines for Configuring Simple Filters*

family (Firewall Filter)

IN THIS SECTION

- Syntax | 6870
- Hierarchy Level | 6870
- Description | 6870

- Options | 6870
- Required Privilege Level | 6871
- Release Information | 6871

Syntax

```
family family-name {
    filter filter-name {
        interface-specific;
        micro-segmentation;
        term term-name {
            from {
                match-conditions;
            }
            then {
                action;
                action-modifiers;
            }
        }
    }
}
```

Hierarchy Level

```
[edit firewall]
```

Description

Configure a firewall filter for IP version 4 or IP version 6.

Options

family-name—Version or type of addressing protocol:

- **any**—Filter packets based on protocol-independent match conditions.

- **ethernet-switching**—Filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets.
- **inet**—Filter IPv4 packets.
- **inet6**—Filter IPv6 packets.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Option **interface-specific** introduced in Junos OS Release 9.5 for EX Series switches.

Option **micro-segmentation** introduced in Junos OS Release 22.4R1 for supported EX4100, EX4400, EX4650, and QFX5120 Series switches.

RELATED DOCUMENTATION

Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches

Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches

Configuring Firewall Filters (CLI Procedure)

Firewall Filters for EX Series Switches Overview

family inet

IN THIS SECTION

- Syntax | 6872
- Hierarchy Level | 6872
- Description | 6872
- Options | 6872

Required Privilege Level | 6873
Release Information | 6873

Syntax

```
family inet {
    layer-3;
    layer-4;
    session-id;
    symmetric-hash {
        complement;
    }
}
```

Hierarchy Level

```
[edit forwarding-options hash-key]
```

Description

Configure layer information for the load-balancing specification. Only the IPv4 protocol is supported.

Options

`family inet`—Incorporate port data into the hash key for flow determination. By default, port data is ignored when determining flows.

- `layer-3`—Incorporate Layer 3 (IP) data into the hash key. You must include the `layer-3` statement. If you omit the `layer-3` statement, the management process removes the `hash-key` statement from the configuration and the router behaves as if you specified `layer-3`.

  By default, or if you specify only the `layer-3` statement, the router uses the following Layer 3 information in the packet header for per-flow load balancing:

  - Source IP address
  - Destination IP address

  - Protocol
  - Incoming interface index

- `layer-4`—Incorporate Layer 4 Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) data into the hash key.

  If you include the `layer-4` statement, the router uses the following Layer 4 information to load-balance:

  - Source port number
  - Destination port number
  - IP type of service

- `session-id`—Include the session ID in the hash key.
- `symmetric-hash`—Create the symmetric hash key with source and destination ports.
  - `complement`—Create the complement of the symmetric hash key.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Per-Packet Load Balancing

family inet (Interfaces)

IN THIS SECTION

- Syntax | 6874
- Hierarchy Level | 6877
- Description | 6877
- Options | 6877
- Required Privilege Level | 6877
- Release Information | 6877

Syntax

```
inet {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
    address (source-address/prefix) {
        arp destination-address {
            (mac mac-address | multicast-mac multicast-mac-address);
            publish publish-address;
        }
        broadcast address;
        preferred;
        primary;
        vrrp-group group-id {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            advertisements-threshold number;
            authentication-key key-value;
            authentication-type (md5 | simple);
            fast-interval milliseconds;
            inet6-advertise-interval milliseconds
```

```
            (preempt <hold-time seconds> | no-preempt );
            priority value;
            track {
                interface interface-name {
                    bandwidth-threshold bandwidth;
                     priority-cost value;
                }
                priority-hold-time seconds;
                route route-address{
                    routing-instance routing-instance;
                     priority-cost value;
                }
            }
            virtual-address [address];
            virtual-link-local-address address;
            vrrp-inherit-from {
                active-group value;
                active-interface interface-name;
            }
        }
        web-authentication {
            http;
            https;
            redirect-to-https;
        }
    }
    dhcp {
        client-identifier {
            (ascii string | hexadecimal string);
        }
        lease-time (length | infinite);
        retransmission-attempt value;
        retransmission-interval seconds;
        server-address server-address;
        update-server;
        vendor-id vendor-id ;
    }
    dhcp-client {
        client-identifier {
            prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
```

```
            }
            use-interface-description (device | logical);
            user-id (ascii string| hexadecimal string);
        }
        lease-time (length | infinite);
        retransmission-attempt value;
        retransmission-interval seconds;
        server-address server-address;
        update-server;
        vendor-id vendor-id ;
    }
    filter {
        group number;
        input filter-name;
        input-list [filter-name];
        output filter-name;
        output-list [filter-name];
    }
    mtu value;
    no-neighbor-learn;
    no-redirects;
    policer {
        arp arp-name;
        input input-name;
        output output-name;
    }
    primary;
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
    sampling {
        input;
        output;
        simple-filter;
    }
    targeted-broadcast {
        (forward-and-send-to-re |forward-only);
    }
    unnumbered-address {
        interface-name;
```

```
        preferred-source-address preferred-source-address;
    }
}
```

Hierarchy Level

```
[edit interfaces interface unit unit ]
```

Description

Assign an IP address to a logical interface.

Options

ipaddress—Specify the IP address for the interface. The remaining statements are explained separately.

> **NOTE**: You use family `inet` to assign an IPv4 address. You use family `inet6` to assign an IPv6 address. An interface can be configured with both an IPv4 and IPv6 address.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement supported in Junos 10.2 for SRX Series Firewalls.

RELATED DOCUMENTATION

Understanding Interfaces

family inet6

IN THIS SECTION

- Syntax | 6878
- Hierarchy Level | 6880
- Description | 6880
- Options | 6880
- Required Privilege Level | 6881
- Release Information | 6881

Syntax

```
inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            ouput;
        }
    }
    address source-address/prefix {
        eui-64;
        ndp address {
            (mac mac-address | multicast-mac multicast-mac-address);
            publish;
        }
        preferred;
        primary;
        vrrp-inet6-group group_id {
            (accept-data | no-accept-data);
            advertisements-threshold number;
            authentication-key value;
            authentication-type (md5 | simple);
            fast-interval milliseconds;
            inet6-advertise-interval milliseconds;
            (preempt <hold-time seconds>| no-preempt );
```

```
            priority value;
            track {
                interface interface-name {
                     bandwidth-threshold value;
                    priority-cost value;
                }
                priority-hold-time seconds;
                route route-address{
                    routing-instance routing-instance;
                }
            }
            virtual-inet6-address [address];
            virtual-link-local-address address;
            vrrp-inherit-from {
                active-group value;
                active-interface interface-name;
            }
        }
        web-authentication {
            http;
            https;
            redirect-to-https;
        }
    }
    (dad-disable | no-dad-disable);
    dhcpv6-client {
        client-ia-type (ia-na | ia-pd);
        client-identifier duid-type (duid-ll | duid-llt | vendor);
        client-type (autoconfig | stateful);
        rapid-commit;
        req-option (dns-server | domain | fqdn | nis-domain | nis-server | ntp-server | sip-
domain | sip-server |time-zone | vendor-spec);
        retransmission-attempt number;
        update-router-advertisement {
            interface interface-name;
        }
        update-server;
    }
    filter {
        group number;
        input filter-name;
        input-list [filter-name];
        output filter-name;
```

```
        output-list [filter-name];
    }
    mtu value;
    nd6-stale-time seconds;
    no-neighbor-learn;
    policer {
        input input-name;
        output output-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
    sampling {
        input;
        output;
    }
    unnumbered-address {
        interface-name;
        preferred-source-address preferred-source-address;
    }
    ndp-proxy | dad-proxy {
        interface-restricted
        }
}
```

Hierarchy Level

```
[edit interfaces interface unit unit ]
```

Description

Assign an IPV6 address to a logical interface.

Options

ipaddress—Specify the IP address for the interface. The remaining statements are explained separately.

> **NOTE**: You use family `inet6` to assign an IPv6 address. You use family `inet` to assign an IPv4 address. An interface can be configured with both an IPv4 and IPv6 address.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement supported in Junos 10.2 for SRX Series Firewalls.

RELATED DOCUMENTATION

Understanding Interfaces

family (Interfaces)

IN THIS SECTION

- Syntax | 6881
- Hierarchy Level | 6882
- Description | 6882
- Options | 6882
- Required Privilege Level | 6882
- Release Information | 6882

Syntax

```
family inet {
    address address {
```

```
        ...
    }
    service {
        input {
            [ service-set service-set-name <service-filter filter-name> ];
            post-service-filter filter-name;
        }
        output {
            [ service-set service-set-name <service-filter filter-name> ];
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure protocol family information for the logical interface.

Options

family—Protocol family. Valid settings for service interfaces include `inet` (IPv4) and `mpls`.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Network Interfaces Library for Routing Devices

Configuring the Address and Domain for Services Interfaces

Junos OS Network Interfaces Library for Routing Devices

family (Local RP)

IN THIS SECTION

- Syntax | 6883
- Hierarchy Level | 6884
- Description | 6884
- Options | 6884
- Required Privilege Level | 6884
- Release Information | 6884

Syntax

```
family (inet | inet6) {
    disable;
    address address;
    anycast-pim {
        local-address address;
        rp-set {
            address address <forward-msdp-sa>;
        }
    }
    group-ranges {
        destination-ip-prefix</prefix-length>;
    }
    hold-time seconds;
    override;
```

```
    priority number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp local],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp local],
[edit protocols pim rp local],
[edit routing-instances routing-instance-name protocols pim rp local]
```

Description

Configure which IP protocol type local RP properties to apply.

Options

inet—Apply IP version 4 (IPv4) local RP properties.

inet6—Apply IPv6 local RP properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Local PIM RPs

family (Monitoring)

IN THIS SECTION

- Syntax | 6885
- Hierarchy Level | 6886
- Description | 6886
- Required Privilege Level | 6886
- Release Information | 6886

Syntax

```
family inet {
    output {
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        export-format format;
        cflowd hostname {
            aggregation {
                autonomous-system;
                destination-prefix;
                protocol-port;
                source-destination-prefix {
                    caida-compliant;
                }
                source-prefix;
            }
            port port-number;
        }
        interface interface-name {
            engine-id number;
            engine-type number;
            input-interface-index number;
            output-interface-index number;
            source-address address;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit forwarding-options monitoring name]
```

Description

Specify input and output interfaces and properties for flow monitoring. Only IPv4 (`inet`) is supported.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Flow Monitoring

Configuring Passive Flow Monitoring

family mpls

IN THIS SECTION

- Syntax | 6887
- Hierarchy Level | 6887

- Description | 6888
- Options | 6888
- Required Privilege Level | 6889
- Release Information | 6889

Syntax

```
family mpls {
    all-labels;
    label-1;
    label-2;
    label-3;
    no-labels;
    no-label-1-exp;
    payload {
        ether-pseudowire {
            zero-control-word;
        }
        ip {
            disable;
            layer-3-only;
            port-data {
                source-msb;
                source-lsb;
                destination-msb;
                destination-lsb;
            }
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options hash-key]
```

Description

For aggregated Ethernet and SONET/SDH interfaces only, configure load balancing based on MPLS labels and payload. Only the IPv4 protocol is supported.

Options

`family mpls`—(Aggregated Ethernet interfaces, aggregated SONET/SDH interfaces, and multiple equal-cost MPLS next hops only) Incorporate MPLS label and payload information into the hash key for per-flow load balancing. Only the IPv4 protocol is supported.

- `all-labels`—(PTX Series Packet Transport Routers only) Up to eight MPLS labels are included in the hash key to identify the uniqueness of a flow in the Packet Forwarding Engine. This is the default setting.
- `label-1`—(M120, M320, MX Series, and T Series routers only) Include the first MPLS label into the hash key. This is used for a one-label packet for per-flow load balancing IPv4 VPLS traffic based on IP information and MPLS labels.
- `label-2`—(M120, M320, MX Series, and T Series routers only) Include the second MPLS label into the hash key. This is used for a two-label packet for per-flow load balancing IPv4 VPLS traffic based on IP information and MPLS labels. To use the second MPLS label in the hash key, include both the `label-1` and `label-2` statements at the `[edit forwarding-options hash-key family mpls]` hierarchy level. By default, the router provides hashing on the first and second labels. If both labels are specified, the entire first label and the first 16 bits of the second label are hashed.
- `label-3`—(M120, M320, MX Series, and T Series routers only) Include the third MPLS label into the hash key. To use the third MPLS label, include the `label-1`, `label-2`, and `label-3` statements at the `[edit forwarding-options hash-key family mpls]` hierarchy level.
- `no-labels`—Include no MPLS labels into the hash key.
- `no-label-1-exp`—(M120, M320, MX Series, and T Series routers only) The EXP bit of the first label is not used in the hash calculation to avoid reordering complications.
- `payload`—Incorporate bits from the IP payload into the hash key for per-flow load balancing Layer 2 information based on MPLS labels.
  - `disable`—(PTX Series Packet Transport Routers only) Exclude IP payload from the hash key.
  - `ether-pseudowire`—(M120, M320, MX Series, and T Series routers only) Load-balance IPv4 traffic over Layer 2 Ethernet pseudowires.
  - `zero-control-word`—(M Series, MX Series, and PTX Series) Precedes Ethernet packet to indicate the start of an Ethernet frame in an MPLS ether-pseudowire payload.

- `ip`—Include the IP address of the IPv4 or IPv6 payload into the hash key for per-flow load balancing Layer 2 information based on MPLS labels. For the PTX Series Packet Transport Routers, this is the default setting with both Layer 3 and Layer 4 IP information included in the hash key.
  - `disable`—(PTX Series Packet Transport Routers only) Exclude IP payload from the hash key.
  - `layer-3-only`—Include only Layer 3 IP information from the IP payload data into the hash key for per-flow load balancing Layer 2 information based on MPLS labels.
  - `port-data`—(M120, M320, MX Series, and T Series routers only) Include the source and destination port field information into the hash key. By default, the most significant byte and least significant byte of the source and destination port fields are hashed. To select specific bytes to be hashed, include one or more of the `source-msb`, `source-lsb`, `destination-msb`, and `destination-lsb` options at the `[edit forwarding-options hash-key family mpls payload ip port-data]` hierarchy level. To prevent all four bytes from being hashed, include the `layer-3-only` statement at the `[edit forwarding-options hash-key family mpls payload ip]` hierarchy level.
    - `destination-lsb`—Include the least-significant byte of the destination port.
    - `destination-msb`—Include the most-significant byte of the destination port.
    - `source-lsb`—Include the least-significant byte of the source port.
    - `source-msb`—Include the most-significant byte of the source port.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`no-label-1-exp` option introduced in Junos OS Release 8.0.

`label-3` and `no-labels` options introduced in Junos OS Release 8.1.

`ether-pseudowire` statement introduced in Junos OS Release 9.1 (M320 and T Series routers only); support extended to M120 and MX Series routers in Junos OS Release 9.4.

`all-labels` and `payload ip disable` options introduced in Junos OS Release 12.1X48R2. (PTX Series Packet Transport Routers only).

`zero-control-word` option introduced in Junos OS Release 16.1 for the M Series, MX Series, and PTX Series.

RELATED DOCUMENTATION

> Configuring Load Balancing Based on MPLS Labels
>
> *Configuring Load Balancing for Ethernet Pseudowires*

family (Multi-destination)

IN THIS SECTION

- Syntax | 6890
- Hierarchy Level | 6891
- Description | 6891
- Required Privilege Level | 6891
- Release Information | 6891

Syntax

```
family {
    ethernet {
        broadcast forwarding-class-name;
    }
    inet {
        classifiers{
            (dscp | inet-precedence) classifier-name;
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service multi-destination]
```

Description

Specify the multidestination traffic family.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding CoS Schedulers

Understanding CoS Forwarding Classes

Understanding CoS Classifiers

family (Multifield Classifier)

IN THIS SECTION

- Syntax | 6892
- Hierarchy Level | 6892
- Description | 6892
- Options | 6892

Required Privilege Level | 6893
Release Information | 6893

Syntax

```
family family-name {
    filter filter-name {
        term term-name {
            ... term_configuration ...
        }
    }
}
```

Hierarchy Level

```
[edit firewall]
```

Description

Configure a firewall filter for IP version 4 (IPv4) or IP version 6 (IPv6) traffic.

Options

family-name—Protocol family:

- `ccc`—Circuit cross-connect parameters
- `inet`—IPv4 parameters
- `inet6`—IPv6 protocol parameters
- `iso`—OSI ISO protocol parameters
- `mlppp`—Multilink PPP protocol parameters
- `mpls`—MPLS protocol parameters
- `tcc`—Translational cross-connect parameters

- `vpls`—Virtual private LAN service parameters.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Multifield Classifiers

family multiservice

IN THIS SECTION

- Syntax | 6893
- Hierarchy Level | 6894
- Description | 6894
- Options | 6894
- Required Privilege Level | 6895
- Release Information | 6896

Syntax

```
family multiservice {
    destination-mac;
    label-1;
```

```
    label-2;
    payload {
        ip {
            layer-3 {
                (source-ip-only | destination-ip-only);
            }
            layer-3-only;
            layer-4;
        }
    }
    source-mac;
    symmetric-hash {
        complement;
    }
}
```

Hierarchy Level

```
[edit forwarding-options hash-key]
```

Description

Configure load balancing based on Layer 2 media access control information. On MX Series routers, configure VPLS load balancing. On M120 and M320 routers only, configure VPLS load balancing based on MPLS labels and IP information. For IPv4 traffic, only the IP source and destination addresses are included in the hash key. For MPLS and IPv4 traffic, one or two MPLS labels and IPv4 source and destination addresses are included. For MPLS Ethernet pseudowires, only one or two MPLS labels are included in the hash key.

Options

You can configure one or more options to load-balance using the packet information that you specify.

`destination-mac`—Include the destination-address MAC information in the hash key for Layer 2 load balancing.

`label-1` (M120 and M320 routers only)—Include the first MPLS label in the hash key. Used for including a one-label packet for per-flow load balancing of IPv4 VPLS traffic based on IP information and MPLS labels.

`label-2` (M120 and M320 routers only)—Include the second MPLS label in the hash key. If both `label-1` and `label-2` are specified, the entire first label and the first 16 bits of the second label are hashed.

`payload` (MX Series, M120, and M320 routers only)—Include the packet's IP payload in the hash key.

- `ip` (MX Series, M120, and M320 routers only)—Include the IP address of the IPv4 or IPv6 payload in the hash key.
  - `layer-3` (MX Series routers only)—Use this to include Layer 3 information from the packet's IP payload in the hash key.
    - `destination-ip-only` (MX Series routers only)—Use this to include only the destination IP address in the payload in the hash key.
    - `source-ip-only` (MX Series routers only)—Use this to include only the source IP address in the payload in the hash key.

    NOTE: You can include either the `source-ip-only` or the `destination-ip-only` statement, not both. They are mutually exclusive.

  - `layer-3-only` (M120, and M320 routers only)—Include only the Layer 3 information from the packet's IP payload in the hash key.
  - `layer-4` (MX Series routers only)—Include Layer 4 information from the packet's IP payload in the hash key.

  NOTE: On MX Series routers only, you can configure either Layer 3 or Layer 4 load balancing, or both at the same time.

  NOTE: On I chip platforms, an unknown Layer 4 header is excluded from load-balance hashing to avoid undesired packet reordering.

`source-mac`—Include the source-address MAC information in the hash key.

`symmetric-hash` (MX Series routers only)—Configure the symmetric hash or symmetric hash complement for configuring symmetrical load balancing on an 802.3ad Link Aggregation Group.

- `complement` —Include the complement of the symmetric hash in the hash key.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

`ip`, `label-1`, `label-2`, `layer-3-only`, and `payload` options introduced in Junos OS Release 9.4.

`layer-3`, `layer-`. `source-ip-only`, and `destination-ip-only` options introduced in Junos OS Release 9.5.

`symmetric-hash` and `complement` options introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Configuring Load Balancing Based on MAC Addresses
- *Configuring VPLS Load Balancing Based on IP and MPLS Information*
- *Configuring VPLS Load Balancing on MX Series 5G Universal Routing Platforms*
- *Configuring VPLS Load Balancing*

family (Port Mirroring)

IN THIS SECTION

- MX, M, T Series Routers, EX Series Switches and SRX Series Firewalls | **6896**
- Syntax: QFX Series (but not QFX10000) and EX4600 Switches and NFX Series Devices | **6897**
- Hierarchy Level | **6897**
- Description | **6897**
- Required Privilege Level | **6898**
- Release Information | **6898**

MX, M, T Series Routers, EX Series Switches and SRX Series Firewalls

```
family (inet | inet6) {
    output {
```

```
        interface interface-name {
            next-hop address;
        }
        no-filter-check;
    }
}
```

Syntax: QFX Series (but not QFX10000) and EX4600 Switches and NFX Series Devices

```
family
    ethernet-switching {
        output {
            interface interface-name {
            }
            no-filter-check;
            }
            vlan vlan-name {
                no-tag;
            }
        }
    inet
        output {
            ip-address address {
            }
            routing-instance instance-name {
                ip-address address {
            }
        }
```

Hierarchy Level

```
[edit forwarding-options port-mirroring]
```

Description

Specify the type of interface that will be used to forward port mirrored packet to an analyzer device. Configure the protocol family to be sampled. Only IPv4 (`inet`) and IPv6 (`inet6`) are supported.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4 .

RELATED DOCUMENTATION

- Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers
- *Understanding Port Mirroring and Analyzers*
- *Configuring Port Mirroring*
- *Examples: Configuring Port Mirroring for Local Analysis*

family (Protocols AMT Relay)

IN THIS SECTION

- Syntax | 6899
- Hierarchy Level | 6899
- Description | 6899
- Required Privilege Level | 6899
- Release Information | 6899

Syntax

```
family {
    inet {
        anycast-prefix ip-prefix/<prefix-length>;
        local-address ip-address;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols amt relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols amt
relay],
[edit protocols amt relay],
[edit routing-instances routing-instance-name protocols amt relay]
```

Description

Configure the protocol address family for Automatic Multicast Tunneling (AMT) relay functions. Only the **inet** family for IPv4 protocol addresses is supported.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

family (Protocols BGP)

IN THIS SECTION

- Syntax | 6900
- Hierarchy Level | 6903
- Description | 6903
- Options | 6903
- Required Privilege Level | 6904
- Release Information | 6904

Syntax

```
family {
  (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) {
    (any | flow | labeled-unicast | multicast | unicast | segment-routing-te) {
      accepted-prefix-limit {
        maximum number;
        teardown <percentage-threshold> idle-timeout (forever | minutes);
        drop-excess <percentage>;
        hide-excess <percentage>;
      }
      add-path {
        receive;
        send {
          include-backup-path backup_path_number;
          multipath;
          path-count number;
          path-selection-mode { (all-paths | equal-cost-paths); }
          prefix-policy [ policy-names ];
        }
      }
      aigp [disable];
      loops number;
      prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
```

```
        drop-excess <percentage>;
        hide-excess <percentage>;
      }
      protection;
      rib-group group-name;
      topology name {
        community { target identifier; }
      }
      flow {
        no-install;
        no-validate policy-name;
      }
      labeled-unicast {
        accepted-prefix-limit {
          maximum number;
          teardown <percentage> <idle-timeout (forever | minutes)>;
          drop-excess <percentage>;
          hide-excess <percentage>;
        }
        aggregate-label { community community-name: }
        explicit-null { connected-only; }
        prefix-limit {
          maximum number;
          teardown <percentage> <idle-timeout (forever | minutes)>;
          drop-excess <percentage>;
          hide-excess <percentage>;
        }
        resolve-vpn;
        rib (inet.3 | inet6.3);
        rib-group group-name;
        traffic-statistics {
          file filename <world-readable | no-world-readable>;
          interval seconds;
        }
      }
    }
    route-target {
      accepted-prefix-limit {
        maximum number;
        proxy-generate <route-target-policy route-target-policy-name>;
        teardown <percentage> <idle-timeout (forever | minutes)>;
        drop-excess <percentage>;
        hide-excess <percentage>;
```

```
        }
        advertise-default;
        external-paths number;
        prefix-limit {
          maximum number;
          teardown <percentage> <idle-timeout (forever | minutes)>;
          drop-excess <percentage>;
          hide-excess <percentage>;
        }
      }
      (evpn | inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
        signaling {
          accepted-prefix-limit {
            maximum number;
            teardown <percentage-threshold> idle-timeout (forever | minutes);
            drop-excess <percentage>;
            hide-excess <percentage>;
          }
          add-path {
            receive;
            send {
              include-backup-path backup_path_number;
              multipath;
              path-count number;
              path-selection-mode {  (all-paths | equal-cost-paths); }
              prefix-policy [ policy-names ];
            }
          }
          aigp [disable];
          damping;
          loops number;
          prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
            drop-excess <percentage>;
            hide-excess <percentage>;
          }
          rib-group group-name;
        }
      }
      traffic-engineering;
```

```
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp family prefix-limit (inet | inet6)(any | flow | labeled-unicast | multicast |
unicast)],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Enable multiprotocol BGP (MP-BGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, to specify MP-BGP to carry NLRI for the IPv6 address family, or to carry NLRI for VPNs.

Options

`any`—Configure the family type to be both unicast and multicast.

`evpn`—Configure NLRI parameters for Ethernet VPNs (EVPNs).

`inet`—Configure NLRI parameters for IPv4.

`inet6`—Configure NLRI parameters for IPv6.

`inet-mdt`—Configure NLRI parameters for the multicast distribution tree (MDT) subaddress family identifier (SAFI) for IPv4 traffic in Layer 3 VPNs.

`inet-mvpn`—Configure NLRI parameters for IPv4 for multicast VPNs.

`inet6-mvpn`—Configure NLRI parameters for IPv6 for multicast VPNs.

inet-vpn—Configure NLRI parameters for IPv4 for Layer 3 VPNs.

inet6-vpn—Configure NLRI parameters for IPv6 for Layer 3 VPNs.

`inet6-vpn`—Configure NLRI parameters for IPv6 for Layer 3 VPNs.

`iso-vpn`—Configure NLRI parameters for IS-IS for Layer 3 VPNs.

`l2vpn`—Configure NLRI parameters for IPv4 for MPLS-based Layer 2 VPNs and VPLS.

`labeled-unicast`—Configure the family type to be labeled-unicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by labeled-unicast for resolving the labeled-unicast routes. This statement is supported only with inet and inet6.

`multicast`—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.

`unicast`—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes. The default family type is unicast.

`segment-routing-te`—Configure the family type to be segment routing traffic engineering. This means that BGP peers only carry segment routing policies for traffic steering.

> **NOTE**: If you enable `family inet-vpn unicast` for external BGP peer, Junos considers the node is positioned as an AS boundary router in an inter-provider scenario. Hence, it keeps the VPN-IPv4 routes in the routing information base (bgp.l3vpn.0), and the labels associated with the prefixes in the forwarding information base (FIB).

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`inet-mvpn` and `inet6-mpvn` statements introduced in Junos OS Release 8.4.

`inet-mdt` statement introduced in Junos OS Release 9.4.

Support for the `loops` statement introduced in Junos OS Release 9.6.

`evpn` statement introduced in Junos OS Release 13.2.

`traffic-engineering` statement introduced in Junos OS Release 14.2.

`segment-routing-te` option introduced in Junos OS Release 17.4R1 for QFX Series, MX Series, and PTX Series with FPC-PTX-P1-A.

RELATED DOCUMENTATION

- *Configuring IBGP Sessions Between PE Routers in VPNs*
- Understanding Multiprotocol BGP
- autonomous-system
- local-as

family (Protocols IS-IS)

IN THIS SECTION

- Syntax | 6905
- Hierarchy Level | 6906
- Description | 6906
- Options | 6906
- Required Privilege Level | 6907
- Release Information | 6907

Syntax

```
family inet {
    shortcuts {
        multicast-rpf-routes;
    }
}
```

```
family inet-mpls {
    shortcuts;
}
family inet6 {
    shortcuts;
}
family inet6-mpls {
    shortcuts;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis traffic-engineering],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis
traffic-engineering],
[edit protocols isis traffic-engineering],
[edit routing-instances routing-instance-name protocols isis traffic-engineering]
```

Description

Configure the address family for traffic engineering IS-IS interior gateway protocol (IGP) shortcuts.

Options

inet Configure the IPv4 address family.

inet-mpls Configure the IPv4 MPLS address family.

inet6 Configure the IPv6 address family.

inet6-mpls Configure the IPv6 address family.

shortcuts Configure IS-IS to use MPLS label-switched paths (LSPs) as next hops if possible when installing routing information into the inet.3 or inet6.3 routing table. Internal gateway protocol (IGP) shortcuts allow the IGP to install prefixes in inet.3 or inet6.3. It is only necessary to enable IGP shortcuts on the ingress router because that is the router performing the shortest-path-first (SPF) calculations.

It is important to understand how IGP shortcuts affect the protocol and routing table relationship. The IGP performs SPF calculations to subnets downstream of LSP egress

points, but the results of these calculations are entered into the inet.3 table only. At the same time, the IGP performs its traditional SPF calculations and enters the results of these calculations into the inet.0 table. The result is that although the IGP is making entries into the inet.3 table, BGP is still the only protocol with visibility into that table for the purposes of route resolution. Therefore, forwarding to AS-internal destinations still uses the inet.0 IGP routes, and the LSPs are only used for BGP next-hop resolution. If you want the LSPs to be used for IGP next-hop resolution, you must configure `traffic-engineering bgp-igp`.

multicast-rpf-routes Install unicast IPv4 routes into the multicast routing table (inet.2) for multicast reverse-path-forwarding (RPF) checks. Traffic engineering shortcuts must be enabled. IPv4 multicast topology must not be enabled. Label-switched paths (LSPs) must not be advertised into IS-IS.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Support for IPv6 for IGP shortcuts introduced in Junos OS Release 9.3.

Support for `inet-mpls` and `inet6-mpls` statements introduced in Junos OS Release 17.2 for the MX Series routers, PTX Series routers, QFX5100, and QFX10000 switches.

family (Protocols PIM)

IN THIS SECTION

- Syntax | 6908
- Hierarchy Level | 6908
- Description | 6908
- Options | 6908

Required Privilege Level | 6908
Release Information | 6909

Syntax

```
family (inet | inet6) {
    disable;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name],
[edit protocols pim],
[edit protocols pim interface (Protocols PIM)  interface-name],
[edit routing-instances routing-instance-name protocols pim],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name]
```

Description

Disable the PIM protocol for the specified family.

Options

inet—Disable the PIM protocol for the IP version 4 (IPv4) address family.

inet6—Disable the PIM protocol for the IP version 6 (IPv6) address family.

Required Privilege Level

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Disabling PIM

disable (PIM Graceful Restart)

disable (PIM) | 5782

family (Protocols PIM Interface)

IN THIS SECTION

- Syntax | 6909
- Hierarchy Level | 6910
- Description | 6910
- Options | 6910
- Release Information | 6910

Syntax

```
family  (inet | inet6) {
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
```

```
        multiplier number;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (0 | 1 | automatic);
    }
    disable;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name],
[edit protocols pim interface (Protocols PIM)  interface-name],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name]
```

Description

Configure one of the following PIM protocol settings for the specified family on the specified interface:

- BFD protocol settings
- Disable PIM

Options

inet—Enable the PIM protocol for the IP version 4 (IPv4) address family.

inet6—Enable the PIM protocol for the IP version 6 (IPv6) address family.

The remaining statements are explained separately. See CLI Explorer.

Release Information

Statement introduced in Junos OS Release 9.6.

Support for the Bidirectional Forwarding Detection (BFD) Protocol statements was introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol

Disabling PIM

family route-target

IN THIS SECTION

- Syntax | 6911
- Hierarchy Level | 6911
- Description | 6912
- Options | 6912
- Required Privilege Level | 6912
- Release Information | 6912

Syntax

```
family route-target {
    advertise-default;
    external-paths number;
    prefix-limit number;
    proxy-generate <route-target-policy route-target-policy-name>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address]
```

Description

Enable BGP route target filtering on the VPN.

The `family route-target` statement is useful for filtering VPN routes before they are sent. Provider edge (PE) routers inform the route reflector (RR) which routes to send, using `family route-target` to provide the route-target-interest information. The RR then sends to the PE router only the advertisements containing the specified route target.

Options

`advertise-default`—Cause the router to advertise the default route target route (`0:0:0/0`) and suppress all routes that are more specific. This can be used by a route reflector on BGP groups consisting of neighbors that act as provider edge (PE) routers only. PE routers often need to advertise all routes to the route reflector. Suppressing all route target advertisements other than the default route reduces the amount of information exchanged between the route reflector and the PE routers. The Junos OS further helps to reduce route target advertisement overhead by not maintaining dependency information unless a nondefault route is received.

`external-paths` *number*—Cause the router to advertise the VPN routes that reference a given route target. The number you specify with the `external-paths` statement determines the number of external peer routers (currently advertising that route target) that receive the VPN routes. The default value is `1`.

> **NOTE**: `external-paths` cannot be configured under the `[edit protocols bgp group` *group-name* `neighbor` *address*`]` hierarchy level.

`prefix-limit` *number*—The number of prefixes that can be received from a peer router.

The remaining statement is described separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring BGP Route Target Filtering for VPNs

Understanding Proxy BGP Route Target Filtering for VPNs

Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs

family (Sampling)

IN THIS SECTION

- Syntax | 6913
- Hierarchy Level | 6915
- Description | 6915
- Required Privilege Level | 6915
- Release Information | 6915

Syntax

```
family (inet | inet6 | mpls | vpls | bridge) {
    disable;
    input {
        max-packets-per-second max-packets-per-second;
        maximum-packet-length maximum-packet-length;
        rate rate;
        run-length run-length
        }
    output {
        aggregate-export-interval seconds;
        file {
             disable;
             filename filename;
             files number;
             size bytes;
            (stamp | no-stamp);
            (world-readable | no-world-readable);
```

```
        }
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        flow-server hostname {
            aggregation {
                autonomous-system;
                destination-prefix;
                protocol-port;
                source-destination-prefix {
                    caida-compliant;
                }
                source-prefix;
            }
            autonomous-system-type (origin | peer);
            dscp dscp-value;
            forwarding-class class-name;
            (local-dump | no-local-dump);
            port port-number;
            version format;
            version9 {
                template template-name;
            }
            version-ipfix {
                template template-name;
            }
        }
        interface interface-name {
            engine-id number;
            engine-type number;
            source-address address;
        }
            inline-jflow {
            source-address address;
            flow-export-rate rate;
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options sampling],
[edit forwarding-options sampling instance instance-name]
[edit forwarding-options sampling instance instance-name family (inet | inet6 | bridge)
```

Description

Configure the protocol family to be sampled. IPv4 (`inet`) is supported for most purposes, but you can configure `family mpls` to collect and export MPLS label information, `family inet6` to collect and export IPv6 traffic using flow aggregation version 9, and `vpls` to collect and export VPLS information, and `bridge` to collect and export bridge information.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: The `inline-jflow` statement is valid only under the `[edit forwarding-options sampling instance instance-name family family output]` hierarchy level. The `file` statement is valid only under the `[edit forwarding-options sampling family inet output]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`mpls` option introduced in Release 8.3.

`inet6` option introduced in Release 9.4.

`vpls` option introduced in Junos OS Release 13.2 for MX Series routers.

`bridge` option introduced in Release 18.2R1 for MX Series routers.

Support for IPv6 addresses for the `flow-server` *address* and `source-address` *address* statements within sampling instances introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

Configuring Traffic Sampling

family-state-change-immediate-update

IN THIS SECTION

- Syntax | 6916
- Hierarchy Level | 6916
- Description | 6916
- Required Privilege Level | 6916
- Release Information | 6917

Syntax

```
family-state-change-immediate-update family-state-change-immediate-update;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Send an `Acct-Update` message to notify address family activation state change.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

Exchange of DHCPv4 and DHCPv6 Parameters with the RADIUS Server Overview

family (TDF Interface)

IN THIS SECTION

- Syntax | 6917
- Hierarchy Level | 6917
- Description | 6917
- Options | 6918
- Required Privilege Level | 6918
- Release Information | 6918

Syntax

```
family family-name;
```

Hierarchy Level

```
[edit interfaces mif unit interface-unit-number]
```

Description

Configure the protocol family information for the TDF logical interface.

Options

family-name Protocol family. The following options are supported:

- `inet`—IP version 4 suite.
- `inet6`—IP version 6 suite.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a TDF Logical Interface

family (Voice Services)

IN THIS SECTION

- Syntax | 6919
- Hierarchy Level | 6919
- Description | 6919
- Options | 6919
- Required Privilege Level | 6919
- Release Information | 6919

Syntax

```
family (inet | mlppp | ...) {
    address address {
        ...
    }
    bundle interface-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Configure protocol family information for the logical interface.

Options

family—Protocol family:

- `inet`—IP version 4
- `mlppp`—MLPPP

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Network Interfaces Library for Routing Devices

Configuring Network Interfaces for Voice Services

Junos OS Network Interfaces Library for Routing Devices

family vpls (Layer 2 Pseudowires)

IN THIS SECTION

- Syntax | 6920
- Hierarchy Level | 6920
- Description | 6920
- Required Privilege Level | 6920
- Release Information | 6921

Syntax

```
family vpls;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

Description

Specify that the protocol family for the logical interface is VPLS.

Required Privilege Level

router—To view this statement in the configuration.

router-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Applying the Policers to Dynamic Profile Interfaces

Creating a Dynamic Profile for the Complex Configuration

family (VRF Advertisement)

IN THIS SECTION

- Syntax | 6921
- Hierarchy Level | 6921
- Description | 6922
- Required Privilege Level | 6922
- Release Information | 6922

Syntax

```
family {
    inet-mvpn;
    inet6-mvpn;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name vrf-advertise-selective],
[edit routing-instances routing-instance-name vrf-advertise-selective]
```

Description

Explicitly enable IPv4 or IPv6 MVPN routes to be advertised from the VRF instance while preventing all other route types from being advertised.

The options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring PIM-SSM GRE Selective Provider Tunnels

inet-mvpn (VRF Advertisement) | **8575**

inet6-mvpn (VRF Advertisement) | **8612**

fast-aps-switch

IN THIS SECTION

- Syntax | **6923**
- Hierarchy Level | **6923**
- Description | **6923**
- Required Privilege Level | **6923**
- Release Information | **6924**

Syntax

```
fast-aps-switch;
```

Hierarchy Level

```
[edit interfaces interface-name sonet-options aps]
```

Description

(M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only, EX Series switches, and MX series routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only using container interfaces) Reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits.

NOTE:

- The fast APS switching feature is supported only within a single chassis on a MX series router using a container interface.
- Configuring this statement reduces the APS switchover time only when the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor is SAToP.
- When the `fast-aps-switch` statement is configured in revertive APS mode, you must configure an appropriate value for revert time to achieve reduction in APS switchover time.
- To prevent the logical interfaces in the data path from being shut down, configure appropriate hold-time values on all the interfaces in the data path that support TDM.
- The `fast-aps-switch` statement cannot be configured when the APS `annex-b` option is configured.
- The interfaces that have the `fast-aps-switch` statement configured cannot be used in virtual private LAN service (VPLS) environments.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Reducing APS Switchover Time in Layer 2 Circuits

fast-failover

IN THIS SECTION

- Syntax | 6924
- Hierarchy Level | 6924
- Description | 6925
- Default | 6925
- Options | 6925
- Required Privilege Level | 6925
- Release Information | 6925

Syntax

```
fast-failover (ge | vcp disable | xe);
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Enable the fast failover feature on an uplink Virtual Chassis port (VCP).

Fast failover is not supported on an EX2200 Virtual Chassis or EX8200 Virtual Chassis.

Default

Fast failover is enabled, by default, on dedicated VCPs.

Fast failover is not enabled on uplink VCPs, including the uplink ports that are, by default, configured as VCPs on an EX3300 switch.

Options

- **ge**—Enable fast failover on all Gigabit Ethernet uplink VCPs in the ring topology.
- **vcp disable**—Disable fast failover on all dedicated VCPs in the ring topology.
- **xe**—Enable fast failover on all 10-Gigabit Ethernet uplink VCPs in the ring topology.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

> *Example: Configuring Fast Failover on Uplink Module VCPs to Reroute Traffic When an EX4200 Virtual Chassis Switch or Intermember Link Fails*
>
> *Configuring Fast Failover in an EX Series Virtual Chassis*
>
> *Disabling Fast Failover in an EX Series Virtual Chassis*

fast-interval

IN THIS SECTION

- Syntax | 6926
- Hierarchy Level | 6926
- Description | 6926
- Options | 6927
- Required Privilege Level | 6927
- Release Information | 6927

Syntax

```
fast-interval milliseconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
```

Description

Configure the interval, in milliseconds, between Virtual Router Redundancy Protocol (VRRP) advertisement packets.

All routers in the VRRP group must use the same advertisement interval.

Options

milliseconds—Interval between advertisement packets.

- **Range:**
  - 10 through 40,950 milliseconds for VRRPv3.
  - 100-999 milliseconds for VRRP.

> **NOTE**: When configuring VRRPv2, we recommend using `advertise-interval` instead of `fast-interval`.

- > **NOTE**: When configuring VRRP for IPv4, if you have chosen not to enable VRRPv3, you cannot set a value less than 100 for `fast-interval`. Commit check fails if a value less than 100 is configured.

- **Default:** 1 second

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Advertisement Interval for the VRRP Primary Router

advertise-interval | **2017**

advertise-interval | **2017**

inet6-advertise-interval (VRRP) | **8596**

version-3 | **30691**

fast-lookup-filter

IN THIS SECTION

- Syntax | 6928
- Hierarchy Level | 6928
- Description | 6928
- Required Privilege Level | 6929
- Release Information | 6929

Syntax

```
fast-lookup-filter;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name],
[edit logical-systems logical-system-name firewall family family-name filter filter-name]
```

Description

The `fast-lookup-filter` is available for the inet and inet6 protocol families for both static and dynamic profiles. Junos installs firewall filters created under this hierarchy to the accelerated filter block available in the certain MPCs, which provides enhanced performance.

Juniper recommends that you use the `payload-protocol` term rather than the `next-header` term when configuring a firewall filter with match conditions for IPv6 traffic. Although either can be used, `payload-protocol` provides the more reliable match condition because it uses the actual payload protocol to find a match, whereas `next-header` simply takes whatever appears in the first header following the IPv6 header, which may or may not be the actual protocol. In addition, if `next-header` is used with IPv6, the accelerated filter block lookup process is bypassed and the standard filter used instead.

See *Firewall Filter Match Conditions for IPv6 Traffic* for more information about firewall filters and terms.

Fast lookup filters can boost filtering performance by as much as three to four times for filters under 3000 terms. Firewall instances from the same firewall block can also be attached to multiple interfaces.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3R3 for MX 240, MX 480, MX 960, MX 2010, and MX 2020 routers with MPC5E, MPC5EQ, or MPC6E, and later for MPC7E, MPC8E, and MPC9E MPCs.

Support for the `next-header` firewall match condition was added in Junos OS Release 13.3R6

Support for MPC2E-NG and MPC3E-NG MPCs was added in Junos OS Release 15.1R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1	Support for MPC2E-NG and MPC3E-NG MPCs was added in Junos OS Release 15.1R1.
13.3R6	Support for the `next-header` firewall match condition was added in Junos OS Release 13.3R6
13.3R3	Statement introduced in Junos OS Release 13.3R3 for MX 240, MX 480, MX 960, MX 2010, and MX 2020 routers with MPC5E, MPC5EQ, or MPC6E,

RELATED DOCUMENTATION

- MX Series 5G Universal Routing Platform Interface Module Reference
- Routing Policies, Firewall Filters, and Traffic Policers User Guide
- MX Series Interface Module Reference manual

fast-lookup-filter (PTX)

IN THIS SECTION

- Syntax | 6930
- Hierarchy Level | 6930
- Description | 6930
- Required Privilege Level | 6935
- Release Information | 6935

Syntax

```
fast-lookup-filter
```

Hierarchy Level

```
[edit firewall family family-name filter output filter-name],
[edit logical-systems logical-system-name firewall family family-name filter output filter-name]
```

Description

On PTX platforms, the `fast-lookup-filter` (FFT) is available on Inet, Inet 6, Ethernet-Switching, Any, MPLS (no MPLS `label` match support) and CCC firewall filter families. Using `fast-lookup-filter` on output filters prioritizes the filters for fast processing. Applying the `fast-lookup-filter` configuration on output firewall filters, programs the firewall filters on the fast-lookup-filter block on the device, to enable high-speed line-rate (2 billion packets-per-second) processing of the firewall filters in the egress direction.

NOTE: It is mandatory to set the fast-lookup-filter knob to apply the family mpls output filter.

On PTX platforms, only the following firewall filter match conditions and actions are supported for fast look filters.

Table 40: Supported firewall filter match conditions for fast lookup filters on PTX platforms

CLI	IPv6	IPv4	MPLS	Ethernet switching	Any	CCC	Except match
learn-vlan-id	No	No	No	Yes	Yes	Yes	Yes
vlan-id	No	No	No	Yes	Yes	Yes	Yes
ttl	Yes	Yes	No	No	No	No	Yes
source-port	Yes	Yes	No	Yes	Yes	Yes	Yes
destination-port	Yes	Yes	No	Yes	Yes	Yes	Yes
source-class	Yes	Yes	No	No	No	Yes	Yes
interface-group	No	No	No	No	No	No	Yes
forwarding-class	Yes	Yes	Yes	Yes	Yes	Yes	Yes
loss-priority	Yes	Yes	Yes	Yes	Yes	Yes	Yes
interface	Yes	Yes	No	Yes	Yes	Yes	Yes
destination-mac	No	No	No	Yes	Yes	Yes	Yes
source-mac	No	No	No	Yes	Yes	Yes	Yes

Table 40: Supported firewall filter match conditions for fast lookup filters on PTX platforms *(Continued)*

CLI	IPv6	IPv4	MPLS	Ethernet switching	Any	CCC	Except match
outer-vlan	No	No	No	Yes	Yes	Yes	Yes
inner-vlan	No	No	No	Yes	Yes	Yes	Yes
ether-type	Yes	Yes	No	Yes	Yes	Yes	Yes
source-address	Yes	Yes	No	Yes	Yes	Yes	Yes
source-prefix-list	Yes	Yes	No	Yes	Yes	Yes	Yes
destination-address	Yes	Yes	No	Yes	Yes	Yes	Yes
destination-prefix-list	Yes	Yes	No	Yes	Yes	Yes	Yes
dscp	Yes	Yes	NA	Yes	Yes	Yes	Yes
Is-fragment	Yes	No	NA	Yes	Yes	Yes	Yes
Ip-options	Yes	Yes	NA	No	No	No	Yes
protocol	Yes	Yes	NA	Yes	Yes	Yes	Yes
src-port	Yes	Yes	NA	Yes	Yes	Yes	Yes
dst-port	Yes	Yes	NA	Yes	Yes	Yes	Yes

Table 40: Supported firewall filter match conditions for fast lookup filters on PTX platforms *(Continued)*

CLI	IPv6	IPv4	MPLS	Ethernet switching	Any	CCC	Except match
icmp-type	Yes	Yes	NA	Yes	Yes	Yes	Yes
icmp-code	Yes	Yes	NA	Yes	Yes	Yes	Yes
tcp-flags	Yes	Yes	NA	Yes	Yes	Yes	Yes
tcp-initial	Yes	Yes	NA	Yes	Yes	Yes	Yes
tcp-established	Yes	Yes	NA	Yes	Yes	Yes	Yes
gre-key	No	No	NA	NA	NA	NA	Yes
vxlan-header	No	No	NA	NA	NA	NA	Yes

Table 41: Supported firewall filter actions for fast lookup filters on PTX platforms

CLI	Supported?
accept	Yes
discard	Yes
reject	Yes
count	Yes
forwarding-class	Yes

Table 41: Supported firewall filter actions for fast lookup filters on PTX platforms *(Continued)*

CLI	Supported?
traffic-class	No
loss-priority	Yes
policer	Yes
sample	Yes
syslog	Yes

NOTE:

- Only 8 unique firewall filters per PFE are supported across all supported firewall filter families.
- Fast lookup is applicable only for output filters.
- Unsupported firewall filter matches and actions is processed from the default processing pipeline, and not programmed on the fast-lookup-filter block.
- Fast lookup filters are limited to 4096 rules of 96 bits or 2048 rules of 192 bits or 1024 rules of 384 bits.
- Maximum 256 logical interface (IFL) bindings for interface specific filter per PFE.
- The `fast-lookup-filter` configuration cannot be applied on input filters.
- There is no support for split filter or filter chaining.
- The `fast-lookup-filter` configuration cannot be used to prioritize amongst the 8 unique fast lookup filters.
- Fast lookup filters do not support slow counters. Counter scale is limited to 48 thousand counter per PFE (Packet Forwarding Engine).

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 23.1 R1 for PTX platforms.

RELATED DOCUMENTATION

fast-lookup-filter

fast-lookup-tuple

IN THIS SECTION

- Syntax | 6935
- Hierarchy Level | 6935
- Description | 6936
- Default | 6936
- Required Privilege Level | 6936
- Release Information | 6937

Syntax

```
fast-lookup-tuple
```

Hierarchy Level

```
[edit firewall family inet filter firewall filter name term term name from]
```

Description

Use `fast-lookup-tuple` for better firewall filter processing performance by matching exact values of the five tuple or five combined match conditions. It enables high scale firewall filter performance improvement with only 5-tuple configurations per term without range configurations by avoiding execution of multiple terms in sequence. You can also use `fast-lookup-tuple-list` to match five tuple list template defined under policy options.

`fast-lookup-tuple` is only for INET firewall filter family. No range of values is supported for each of the 5-tuple fields. `fast-lookup-tuple` can be applied at all bind points including chain filters, list filters, and nested filters. Next-term action is not supported for `fast-lookup-tuple` match. `fast-lookup-tuple` cannot be defined for fast lookup filters.

The following is a sample configuration.

```
set firewall family inet filter f1 term t1 from fast-lookup-tuple
10.11.12.13:20.21.22.23:17:168:65535
```

Default

l3 byte value - `<dest-ip-in-regular-format>:<src-ip-in-regular-format>:<proto>:<src-port>:<dest-port>`

- `<dest-ip-in-regular-format>`: Destination IP address as <a.b.c.d>. Where a, b, c and d are in decimal values. Example: 10.11.12.13 as standard IP notation.
- `<src-ip-in-regular-format>`: Source IP address as <a.b.c.d>. Where a, b, c and d are in decimal values. Example: 10.12.13.14 as standard IP notation.
- `<proto>`: Single byte protocol value in decimal format. Note - Standard protocol names are not allowed for protocol value.
- `<src-port>`: Two bytes port values in decimal format.
- `<dest-port>`: Two bytes port values in decimal format.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos 24.2 on MX platforms

fast-lookup-tuple-list

IN THIS SECTION

- Syntax | 6937
- Hierarchy Level | 6937
- Description | 6937
- Default | 6938
- Required Privilege Level | 6938
- Release Information | 6938

Syntax

```
fast-lookup-tuple-list list name
```

Hierarchy Level

```
[edit firewall family inet filter firewall filter name term term name from]
```

Description

Match any of the list of 5-tuple (fast-lookup-tuple) combinations. The list is defined under policy options.

`fast-lookup-tuple-list` is only for INET firewall filter family. No range of values is supported for each of the 5-tuple fields. `fast-lookup-tuple-list` can be applied at all bind points including chain filters, list filters, and nested filters.

The following is an example of a configuration.

```
set policy-options fast-lookup-tuple-list five-tuple-list 10.11.12.13:20.21.22.23:17:168:65535
```

Default

`fast-lookup-tuple-list` *`list name`* - the name of the list of 5-tuple combinations defined under policy options.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos 24.2 on MX platforms

fast-lookup-tuple-list (policy-options)

IN THIS SECTION

- Syntax | 6938
- Hierarchy Level | 6939
- Description | 6939
- Required Privilege Level | 6939
- Release Information | 6939

Syntax

```
fast-lookup-tuple-list list name
```

Hierarchy Level

```
[edit policy-options]
```

Description

Define a list of fast-lookup-tuple or 5-tuple combinations. In the following example, the `five-tuple-list` list has a list of three 5-tuple combinations defined.

```
policy-options {
    fast-lookup-tuple-list five-tuple-list {
        10.11.12.13:20.21.22.23:17:168:65534
        10.11.12.13:20.21.22.23:17:168:65535
        20.11.12.13:20.21.22.23:17:168:65534
    }
}
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos 24.2 on MX platforms

fast-poe

IN THIS SECTION

- Syntax | 6940
- Hierarchy Level | 6940

- Release Information | 6940
- Description | 6940
- Required Privilege Level | 6940

Syntax

```
fast-poe;
```

Hierarchy Level

```
[edit poe]
```

Release Information

Statement introduced in Junos OS Release 21.1.

Description

Fast Power over Ethernet (fast PoE) delivers PoE power to devices connected to the ports even before the switch is fully operational. This is beneficial when the connected devices need only power, and not network connectivity.

> **NOTE**: Fast PoE is disabled if and when you issue commands to zeroize the switch, do a USB install, or load the factory default.

Required Privilege Level

system

fast-reroute-priority

IN THIS SECTION

- Syntax | 6941
- Hierarchy Level | 6941
- Description | 6941
- Default | 6942
- Options | 6942
- Required Privilege Level | 6942
- Release Information | 6942

Syntax

```
fast-reroute-priority (high | low | medium);
```

Hierarchy Level

```
[edit forwarding-options]
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options],
[edit routing-instances routing-instance-name forwarding-options]
```

Description

Specify the fast reroute priority for a VPLS routing instance. You can configure `high`, `medium`, or `low` fast reroute priority to prioritize specific VPLS routing instances for faster convergence and traffic restoration. Because the router repairs next hops for high-priority VPLS routing instances first, the traffic traversing a VPLS routing instance configured with `high` fast reroute priority is restored faster than the traffic for VPLS routing instances configured with `medium` or `low` fast reroute priority.

NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Default

`low`

Options

`high`—Set the fast reroute priority for a VPLS routing instance to high. During a fast reroute event, the router repairs next hops for high-priority VPLS routing instances first.

`low`—Set the fast reroute priority for a VPLS routing instance to low, which is the default. During a fast reroute event, the router repairs next hops for low-priority VPLS routing instances last.

`medium`—Set the fast reroute priority for a VPLS routing instance to medium. During a fast reroute event, the router repairs next hops for medium-priority VPLS instances after high-priority VPLS routing instances but before low-priority VPLS routing instances.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring VPLS Fast Reroute Priority

fast-reroute (Protocols MPLS)

IN THIS SECTION

- Syntax | 6943
- Hierarchy Level | 6943
- Description | 6943
- Options | 6943
- Required Privilege Level | 6944
- Release Information | 6944

Syntax

```
fast-reroute {
    (bandwidth bps | bandwidth-percent percentage);
    (exclude  [ group-names ] | no-exclude );
    hop-limit number;
    (include-all [ group-names ] | no-include-all);
    (include-any [ group-names ] | no-include-any);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Establish detours for the LSP so that if a node or link in the LSP fails, the traffic on the LSP can be rerouted with minimal packet loss.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

RELATED DOCUMENTATION

- *Configuring Fast Reroute*
- *Fast Reroute Overview*
- *MPLS Feature Support on QFX Series and EX4600 Switches*
- *Understanding Interprovider and Carrier-of-Carriers VPNs*

fast-reroute (Protocols RSVP)

IN THIS SECTION

- Syntax | 6944
- Hierarchy Level | 6945
- Description | 6945
- Options | 6945
- Required Privilege Level | 6945
- Release Information | 6945

Syntax

```
fast-reroute optimize-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Configure the optimize timer for fast reroute. The optimize timer triggers a periodic optimization process that recomputes the fast reroute detour LSPs to use network resources more efficiently.

Options

seconds—Specify the number of seconds between fast reroute detour LSP optimizations.

- **Range:** 0 through 65,535 seconds
- **Default:** 0 (disabled)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 7.5.

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

> *Configuring the Optimization Interval for Fast Reroute Paths*

fast-update-filter (Dynamic Firewalls)

IN THIS SECTION

- Syntax | 6946
- Hierarchy Level | 6946
- Description | 6947
- Options | 6947
- Required Privilege Level | 6947
- Release Information | 6947

Syntax

```
fast-update-filter filter-name {
    interface-specific;
    match-order [match-order];
    term term-name {
        from {
            match-conditions;
        }
        then {
            action;
            action-modifiers;
        }
        only-at-create;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family]
```

Description

Configure fast update firewall filters in a dynamic profile.

Options

filter-name—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Fast Update Filters

fast-lookup-filter (Protocols BGP)

IN THIS SECTION

- Syntax | 6948
- Hierarchy Level | 6948
- Description | 6948
- Required Privilege Level | 6948
- Release Information | 6948

Syntax

```
fast-lookup-filter | no-fast-lookup-filter;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options flow],
[edit logical-systems name routing-options flow],
[edit logical-systems name tenants name routing-instances name routing-options flow],
[edit routing-instances name routing-options flow],
[edit routing-options flow]
```

Description

As BGP flowspec filters grow, they can sometimes cause congestion or packet drops. You can enable the fast lookup for the BGP flowspec filters to perform a parallel lookup, which can significantly improve packet throughput and filter performance.

NOTE: Currently we support fast lookup only on MX Series devices with MPCs/MICs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Example: Enabling BGP to Carry Flow-Specification Routes

flow (protocols bgp) | 7250

fastether-options

IN THIS SECTION

- Syntax | 6949
- Hierarchy Level | 6950
- Description | 6950
- Required Privilege Level | 6950
- Release Information | 6950

Syntax

```
fastether-options {
    802.3ad {
        aex (primary | backup);
        lacp {
            port-priority;
        }
    }
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    ingress-rate-limit rate;
    (loopback | no-loopback);
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure Fast Ethernet-specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Ethernet Interfaces Overview

fatal

IN THIS SECTION

- Syntax | 6951
- Hierarchy Level | 6951
- Description | 6951
- Required Privilege Level | 6951
- Release Information | 6951

Syntax

```
fatal {
    threshold threshold-value;
    action (chassis error) {
        alarm;
        disable-pfe;
        get-state;
        log;
        offline;
        reset;
    }
}
```

Hierarchy Level

[edit chassis "fpc" on page 7516 *slot-number* "error " on page 6441]

[edit chassis]

Description

Severity level of the error. An error that results in blockage of considerable amount of traffic across modules is a fatal error. The severity level of an error cannot be configured by a user.

The other statements are explained separately.

Required Privilege Level

interface	To view this statement in the configuration.
interface-control	To add this statement to the configuration.

Release Information

Statement introduced for PTX Series routers in Junos OS Release 13.3.

Statement introduced for MX240, MX480, MX960, and MX2020 routers in Junos OS Release 14.2.

RELATED DOCUMENTATION

Fabric Resiliency

Managing Errors

fate-sharing

IN THIS SECTION

- Syntax | 6952
- Hierarchy Level | 6952
- Description | 6953
- Options | 6953
- Required Privilege Level | 6953
- Release Information | 6953

Syntax

```
fate-sharing {
    group group-name {
        cost value;
        from address <to address>;
        use-for-post-convergence-lfa;    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit routing-options],
[edit routing-instances routing-instance-name routing-options]
```

Description

Specify a backup path in case the primary path becomes unusable.

You specify one or more objects with common characteristics within a group. All objects are treated as /32 host addresses. The objects can be a LAN interface, a router ID, or a point-to-point link. Sequence is insignificant.

Changing the fate-sharing database does not affect existing established LSPs until the next CSPF reoptimization. The fate-sharing database does affect fast-reroute detour path computations.

Options

`cost` *`value`*—Cost assigned to the group.

- **Range:** 1 through 65,535
- **Default:** 1

`from` *`address`*—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate `from` addresses in the group.

`group` *`group-name`*—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.

`to` *`address`*—(Optional) Address of egress router. For point-to-point link objects, you must specify both a `from` and a `to` address.

use-for-post-convergence-lfa—Use this fate-sharing group as a constraint for the backup path computed by topology-independent loop-free alternate (TI-LFA).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Alternate Backup Paths Using Fate Sharing

MPLS Applications User Guide

fc-fabrics

IN THIS SECTION

- Syntax | 6954
- Hierarchy Level | 6956
- Description | 6956
- Options | 6956
- Required Privilege Level | 6956
- Release Information | 6956

Syntax

```
fc-fabrics {
    fc-fabric-name {
        description
        fabric-id fc-fabric-id;
        fabric-type proxy;
        interface {
            interface-name {
                max-login-sessions max-login-sessions;
            }
            interface-name {
                max-login-sessions max-login-sessions;
            }
            <...>;
                max-login-sessions max-login-sessions;
            }
            vlan.interface-name;
        }
```

```
        fc2 {
            traceoptions {
                file filename <replace> <size size> <files number> <no-stamp>;
                <world-readable | no-world-readable>;
                flag flag <flag-modifier>;
            }
        }
        max-login-sessions max-login-sessions;
        protocols {
            fip {
                fcoe-trusted;
                fc-map fc-map-value;
                fka-adv-period milliseconds;
                interface {
                    interface-name {
                        fka-adv-period milliseconds;
                        priority priority;
                    }
                }
                max-sessions-per-enode max-sessions-per-enode;
                priority priority;
                traceoptions {
                    file filename <replace> <size size> <files number> <no-stamp>;
                    <world-readable | no-world-readable>;
                    flag flag <flag-modifier> <disable>;
                }
            }
        }
        proxy {
            auto-load-rebalance
            load-balance-algorithm (simple | enode-based | flogi-based);
            no-fabric-wwn-verify;
            traceoptions {
                file filename <replace> <size size> <files number> <no-stamp>;
                <world-readable | no-world-readable>;
                flag flag <flag-modifier> <disable>;
            }
        }
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure an FC fabric. You can configure a maximum of 12 FC fabrics, one per native FC port.

> **NOTE**: Changing the name of an FC fabric causes all logins to drop and forces the ENodes to log in again.

Options

fc-fabric-name —Unique name of the FC fabric.

The other statements are explained separately.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fabric | **19929**

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

Configuring a Physical Fibre Channel Interface

Converting an Ethernet Interface To a Fibre Channel Interface

Configuring an FCoE VLAN Interface on an FCoE-FC Gateway

Assigning Interfaces to a Fibre Channel Fabric

Configuring an FCoE-FC Gateway Fibre Channel Fabric

Configuring FIP on an FCoE-FC Gateway

Configuring DCBX Autonegotiation

Overview of Fibre Channel

Understanding FCoE-FC Gateway Functions

fc-map

IN THIS SECTION

- Syntax | 6957
- Hierarchy Level | 6957
- Description | 6958
- Options | 6958
- Required Privilege Level | 6959
- Release Information | 6959

Syntax

```
fc-map fc-map-value;
```

Hierarchy Level

Original CLI

```
[edit ethernet-switching options secure-access-port vlan (all | vlan-name) examine-fip]
```

ELS CLI for Platforms that Support FCoE

```
[edit vlans vlan-name forwarding-options fip-security]
```

> **NOTE**: The `fc-map` configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.

QFX Series that Support FCoE-FC Gateway Configuration

```
[edit fc-fabrics fc-fabric-name protocols fip]
```

Description

Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN).

You can configure the FC-MAP value or use the default value. The default FC-MAP value is different for VN_Port to VF_Port (VN2VF_Port) FIP snooping (0x0EFC00) than for VN_Port to VN_Port (VN2VN_Port) FIP snooping.

The FC switch provides the FC-MAP value to FCoE nodes (ENodes) in the FIP discovery advertisement message. If the EX Series switch or the QFX Series FCoE VLAN FC-MAP value does not match the FC switch FC-MAP value, neither device discovers the FC switch on that VLAN, and the ENodes on that VLAN cannot access the FC switch. The FC switch accepts only FCoE traffic that uses the correct FC-MAP value as part of the VN_Port MAC address.

When the QFX Series acts as an FCoE-FC gateway, the FC-MAP value for the gateway and the FCoE devices must match the FC switch FC-MAP value in order to communicate with the FC switch.

> **NOTE**: Changing the FC-MAP value causes all logins to drop and forces the ENodes to log in again.

Options

fc-map-value—FC-MAP value, hexadecimal value preceded by “0x”.

- **Range:** 0x0EFC00 through 0x0EFCFF
- **Default:** 0x0EFC00 for VN2VF_Port FIP snooping

  0x0EFD00 for VN2VN_Port FIP snooping

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

- examine-fip
- show fip snooping | 20008
- Example: Configuring an FCoE Transit Switch
- *Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch*

fc-options

IN THIS SECTION

- Syntax | 6959
- Hierarchy Level | 6960
- Description | 6960
- Options | 6960
- Required Privilege Level | 6960
- Release Information | 6960

Syntax

```
fc-options
    max-login-sessions-per-node max-login-sessions-per-node;
```

```
    no-fip-snooping-scaling;
    traceoptions {
        file filename <replace> <size size> <files number> <no-stamp>;
        <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
```

Hierarchy Level

```
[edit]
```

Description

Set Fibre Channel options.

Options

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Statement `no-fip-snooping-scaling` introduced in Junos OS Release 13.2X52-D10 for the QFabric system.

fc2

IN THIS SECTION

- Syntax | 6961
- Hierarchy Level | 6961

- Description | 6961
- Options | 6961
- Required Privilege Level | 6961

Syntax

```
fc2 {
    traceoptions {
        file filename <replace> <size size> <files number> <no-stamp>;
        <world-readable | no-world-readable>;
        flag flag <flag-modifier>;
    }
}
```

Hierarchy Level

```
[edit fc-fabrics fc-fabric-name]
```

Description

Fibre Channel network layer (FC2) configuration.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

fcoe

IN THIS SECTION

- Syntax | 6962
- Hierarchy Level | 6962
- Description | 6962
- Options | 6962
- Required Privilege Level | 6962
- Release Information | 6963

Syntax

```
fcoe {
    no-auto-negotiation;
}
```

Hierarchy Level

```
[edit protocols dcbx interface interface-name applications]
```

Description

Disable advertising the FCoE state of the interface to the peer. To disable FCoE on the interface, do not configure the FCoE forwarding class on the interface.

Options

no-auto-negotiate—Disable automatic negotiation of FCoE capability.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show dcbx neighbors | 19191

Understanding DCB Features and Requirements

Understanding DCB Features and Requirements on EX Series Switches

fcoe-lag

IN THIS SECTION

- Syntax | 6963
- Hierarchy Level | 6963
- Description | 6964
- Required Privilege Level | 6964
- Release Information | 6964

Syntax

```
fcoe-lag;
```

Hierarchy Level

```
[edit interfaces lag-interface-name aggregated-ether-options]
```

Description

Configure a special link aggregation group (LAG) to transport Fibre Channel over Ethernet (FCoE) traffic and regular Ethernet traffic across the same link aggregation bundle.

An FCoE LAG ensures that FCoE traffic uses the same link within a LAG to transmit and receive information between an FCoE device and a Fibre Channel (FC) SAN switch across a QFabric system Node device. This preserves the point-to-point link emulation that FC requires. A standard LAG uses a hashing algorithm to determine the LAG link used for each communication, so with a standard LAG, you cannot guarantee that communication between an FCoE device and the QFabric system Node device always uses the same link. If communication between the FCoE device and the QFabric system Node device uses different physical links, the SAN terminates the link.

An FCoE LAG treats regular Ethernet traffic (traffic that is not FCoE traffic) in the same way as on a standard LAG, providing link redundancy and load-balancing for the regular Ethernet traffic. An FCoE LAG does not provide link redundancy or load balancing for FCoE traffic.

On FCoE-FC gateways, if the gateway has one or more untrusted FC fabrics, you must also disable FIP snooping scaling on the gateway by including the `no-fip-snooping-scaling` option in the `[edit fc-options]` hierarchy.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X52-D10.

RELATED DOCUMENTATION

no-fcoe-lag | **11457**

no-fip-snooping-scaling | **11461**

Configuring an FCoE LAG

Configuring VLANs for FCoE Traffic on an FCoE Transit Switch

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

Understanding FCoE LAGs

fcoe-trusted

IN THIS SECTION

- Syntax | 6965
- Hierarchy Level | 6965
- Description | 6966
- Required Privilege Level | 6966
- Release Information | 6966

Syntax

```
fcoe-trusted;
```

Hierarchy Level

Original CLI

```
[edit ethernet-switching-options secure-access-port interface interface-name]
```

ELS CLI for Platforms that Support FCoE

```
[edit vlans vlan-name forwarding-options fip-security interface interface-name]
```

NOTE: The `fcoe-trusted` configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.

QFX Series that Support FCoE-FC Gateway Configuration

```
[edit fc-fabrics fc-fabric-name protocols fip]
```

Description

Configure the specified 10-Gigabit Ethernet interface to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.

(QFX Series FCoE-FC gateway) Configure the specified local Fibre Channel fabric to trust FCoE traffic on all ports in the fabric. Changing the fabric ports from untrusted to trusted removes any existing FIP snooping filters from the ports. Changing the fabric ports from trusted to untrusted by removing the `fcoe-trusted` configuration from the fabric forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN_Ports log in again, the switch can build the appropriate FIP snooping filters.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Statement introduced for the FC fabric in Junos OS Release 11.3 for the QFX Series.

Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

show fip snooping | **20008**

Example: Configuring an FCoE Transit Switch

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

feature-profile

IN THIS SECTION

- Syntax | 6967
- Hierarchy Level | 6976
- Description | 6976
- Options | 6976
- Required Privilege Level | 6976
- Release Information | 6976

Syntax

```
feature-profile {
    anti-spam {
        address-blacklist list-name;
        address-whitelist list-name;
        sbl {
            profile profile-name {
                custom-tag-string [string];
                (sbl-default-server | no-sbl-default-server);
                spam-action (block | tag-header | tag-subject);
            }
        }
        traceoptions flag flag;
    }
    anti-virus {
        juniper-express-engine {
            pattern-update {
                email-notify {
                    admin-email email-address;
                    custom-message message;
                    custom-message-subject message-subject;
                }
                interval value;
                no-autoupdate;
                proxy {
```

```
                password password-string;
                port port-number;
                server address-or-url;
                username name;
            }
            url url;
        }
        profile profile-name {
             fallback-options {
                content-size (block | log-and-permit);
                default (block | log-and-permit);
                engine-not-ready (block | log-and-permit);
                out-of-resources (block | (log-and-permit);
                timeout (block | log-and-permit);
                too-many-requests (block | log-and-permit);
            }
            notification-options {
                fallback-block {
                    administrator-email email-address;
                    allow-email;
                    custom-message message;
                    custom-message-subject message-subject;
                    display-host;
                    (notify-mail-sender | no-notify-mail-sender);
                     type (message | protocol-only);
                }
                fallback-non-block {
                    custom-message message;
                    custom-message-subject message-subject;
                    (notify-mail-recipient | no-notify-mail-recipient);
                }
                virus-detection {
                    custom-message message;
                    custom-message-subject message-subject;
                    (notify-mail-sender | no-notify-mail-sender);
                     type (message | protocol-only);
                }
            }
            scan-options {
                content-size-limit value;
                (intelligent-prescreening | no-intelligent-prescreening);
                timeout value;
            }
```

```
                trickling {
                    timeout value;
                }
            }
        }
        kaspersky-lab-engine {
            pattern-update {
                email-notify {
                    admin-email email-address;
                    custom-message message;
                    custom-message-subject message-subject;
                }
                interval value;
                no-autoupdate;
                proxy {
                    password password-string;
                    port port-number;
                    server address-or-url;
                    username name;
                }
                url url;
            }
            profile profile-name {
                 fallback-options {
                    content-size (block | log-and-permit);
                    corrupt-file (block | log-and-permit);
                    decompress-layer (block | log-and-permit);
                    default (block | log-and-permit);
                    engine-not-ready (block | log-and-permit);
                    out-of-resources (block | (log-and-permit);
                    password-file (block | (log-and-permit);
                    timeout (block | log-and-permit);
                    too-many-requests (block | log-and-permit);
                }
                notification-options {
                    fallback-block {
                        administrator-email email-address;
                        allow-email;
                        custom-message message;
                        custom-message-subject message-subject;
                        display-host;
                        (notify-mail-sender | no-notify-mail-sender);
                         type (message | protocol-only);
```

```
                }
                fallback-non-block {
                    custom-message message;
                    custom-message-subject message-subject;
                    (notify-mail-recipient | no-notify-mail-recipient);
                }
                virus-detection {
                    custom-message message;
                    custom-message-subject message-subject;
                    (notify-mail-sender | no-notify-mail-sender);
                     type (message | protocol-only);
                }
            }
            scan-options {
                content-size-limit value;
                decompress-layer-limit value;
                (intelligent-prescreening | no-intelligent-prescreening);
                scan-extension filename;
                scan-mode  (all | by-extension);
                timeout value;
            }
            trickling {
                timeout value;
            }
        }
    }
    mime-whitelist {
        exception listname;
        list listname {
            exception listname;
        }
    }
    sophos-engine {
        pattern-update {
            email-notify {
                admin-email email-address;
                custom-message message;
                custom-message-subject message-subject;
            }
            interval value;
            no-autoupdate;
            proxy {
                password password-string;
```

```
                port port-number;
                server address-or-url;
                username name;
            }
            url url;
        }
        profile <name> {
            fallback-options {
                content-size (block | log-and-permit | permit);
                default (block | log-and-permit | permit);
                engine-not-ready (block | log-and-permit | permit);
                out-of-resources (block | log-and-permit | permit);
                timeout (block | log-and-permit | permit);
                too-many-requests (block | log-and-permit | permit);
            }
            notification-options {
                fallback-block {
                    administrator-email email-address;
                    allow-email;
                    custom-message message;
                    custom-message-subject message-subject;
                    display-host;
                    (notify-mail-sender | no-notify-mail-sender);
                     type (message | protocol-only);
                }
                fallback-non-block {
                    custom-message message;
                    custom-message-subject message-subject;
                    (notify-mail-recipient | no-notify-mail-recipient);
                }
                virus-detection {
                    custom-message message;
                    custom-message-subject message-subject;
                    (notify-mail-sender | no-notify-mail-sender);
                     type (message | protocol-only);
                }
            }
            scan-options {
                content-size-limit value;
                (no-uri-check | uri-check);
                timeout value;
            }
            trickling {
```

```
                    timeout value;
                }
            }
            sxl-retry value;
            sxl-timeout seconds;
        }
        traceoptions flag flag;
         type (juniper-express-engine | kaspersky-lab-engine | sophos-engine);
        url-whitelist listname;
    }
    content-filtering {
        profile profile-name {
            block-command protocol-command-list;
            block-content-type (activex | exe | http-cookie | java-applet | zip);
            block-extension extension-list;
            block-mime {
                exception list-name;
                list list-name;
            }
            notification-options {
                custom-message message;
                (notify-mail-sender | no-notify-mail-sender);
                 type (message | protocol-only);
            }
            permit-command protocol-command-list;
        }
        traceoptions flag flag;
    }
    web-filtering {
        url-whitelist custwhitelist;
        url-blacklist custblacklist;
        http-reassemble;
        type juniper-enhanced;
        juniper-enhanced {
            cache {
                timeout 1800;
                size 500;
            }
            server {
                host rp.cloud.threatseeker.com;
                port 80;
            }
            profile junos-wf-enhanced-default {
```

```
                category {
                    Enhanced_Hacking {
                        action log-and-permit;
                        }
                        Enhanced_Government {
                        action quarantine;
                    }
                }
                site-reputation-action {
                    very-safe permit;
                    moderately-safe log-and-permit;
                    fairly-safe log-and-permit;
                    harmful block;
                    suspicious block;
                }
                default block;
                custom-block-message "***access denied ***";
                fallback-settings {
                    default block;
                    server-connectivity block;
                    timeout block;
                    too-many-requests block;
                }
                timeout 10;
                no-safe-search;
            }
            utm-policy mypolicy {
                web-filtering {
            http-profile my_ewfprofile01;
        }
        }
    }
    web-filtering {
        juniper-enhanced {
            cache {
                size value;
                timeout value;
            }
            profile profile-name {
                category customurl-list name {
                    action  (block | log-and-permit | permit | quarantine);
                }
                custom-block-message value;
```

```
                custom-quarantine-message value;
                default  (block | log-and-permit | permit | quarantine);
                fallback-settings {
                    default  (block | log-and-permit);
                    server-connectivity   (block | log-and-permit);
                    timeout    (block | log-and-permit);
                    too-many-requests  (block | log-and-permit);
                }
                no-safe-search;
                site-reputation-action {
                    fairly-safe   (block | log-and-permit | permit | quarantine);
                    harmful   (block | log-and-permit | permit | quarantine);
                    moderately-safe  (block | log-and-permit | permit | quarantine);
                    suspicious  (block | log-and-permit | permit | quarantine);
                    very-safe   (block | log-and-permit | permit | quarantine);
                }
                timeout value;
            }
            server {
                host host-name;
                port number;
            }
        }
        juniper-local {
            profile profile-name {
                custom-block-message value;
                default  (block | log-and-permit | permit);
                fallback-settings {
                    default (block | log-and-permit);
                    server-connectivity (block | log-and-permit);
                    timeout  (block | log-and-permit);
                    too-many-requests  (block | log-and-permit);
                }
                timeout value;
                no-safe-search;
            }
        }
        surf-control-integrated {
            cache {
                size value;
                timeout value;
            }
            profile profile-name {
```

```
            category customurl-list name {
                action  (block | log-and-permit | permit);
            }
            custom-block-message value;
            default (block | log-and-permit | permit);
            fallback-settings {
                default  (block | log-and-permit);
                server-connectivity   (block | log-and-permit);
                timeout    (block | log-and-permit);
                too-many-requests  (block | log-and-permit);
            }
            timeout value;
        }
        server {
            host host-name;
            port number;
        }
    }
    traceoptions flag flag;
    type (juniper-enhanced | juniper-local | surf-control-integrated | websense-redirect);
    url-blacklist listname;
    url-whitelist listname;
    websense-redirect {
         profile profile-name {
            account value;
            custom-block-message value;
            fallback-settings {
                default  (block | log-and-permit);
                server-connectivity  (block | log-and-permit);
                timeout   (block | log-and-permit);
                too-many-requests  (block | log-and-permit);
            }
            server {
                host host-name;
                port number;
            }
            sockets value;
            timeout value;
            no-safe-search;
        }
    }
```

```
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm]
```

Description

Configure Content Security features, antivirus, antispam, content-filtering, and web-filtering by creating feature profiles.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Starting in Junos OS Release 21.4R1, content evaluation is based of the file content. The file type-based evaluation of content is deprecated and the related configurations are hidden. So, the content filtering options in this hierarchy is deprecated and not supported from Junos OS Release 21.4R1.

You can use the legacy functionality if you do not want to migrate to enhanced content filtering functionality. You will be allowed to use the legacy configurations, but all the legacy configuration knobs are deprecated and hidden. Also, you will receive system logs and error message warnings when you use the legacy configuration options.

The Kaspersky, Express antivirus and Surf-Control features are not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Release 9.5.

Starting with Junos OS Release 18.2R1, the following commands under the `[edit security utm feature-profile]` hierarchy level are deprecated:

- `set web-filtering type`

- `set web-filtering url-blacklist`
- `set web-filtering url-whitelist`
- `set web-filtering http-persist`
- `set web-filtering http-reassemble`
- `set web-filtering traceoptions`
- `set web-filtering juniper-enhanced cache`
- `set web-filtering juniper-enhanced reputation`
- `set web-filtering juniper-enhanced query-type`
- `set anti-virus mime-whitelist`
- `set anti-virus url-whitelist`
- `set anti-virus type`
- `set anti-virus traceoptions`
- `set anti-virus sophos-engine`
- `set anti-spam address-blacklist`
- `set anti-spam address-whitelist`
- `set anti-spam traceoptions`
- `set content-filtering traceoptions`

`no-safe-search` option added for Websense redirect and Juniper local in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

UTM Supported Features

features (VNF)

IN THIS SECTION

- Syntax | 6978
- Hierarchy Level | 6978
- Description | 6978
- Options | 6978
- Required Privilege Level | 6979
- Release Information | 6979

Syntax

```
features {
    hugepages;
    }
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Displays the supported features of a VNF.

Options

features Features of a VNF.

hugepages Option to support memory pages with a size of 2 MB and 1 GB.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D40.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

features

IN THIS SECTION

- Syntax | 6979
- Hierarchy Level | 6980
- Description | 6980
- Options | 6980
- Required Privilege Level | 6981
- Release Information | 6981

Syntax

```
features {
   aggregate-intf-member-id;
   chip-delay;
   egress-drop-reason;
   flow-start-end-time;
```

```
    ingress-drop-reason;
    inter-arrival-time;
    inter-departure-time;
    queue-congestion-level;
    security-enable;
    shared-pool-congestion;
}
```

Hierarchy Level

```
[edit services inline-monitoring feature-profile]
```

Description

When you monitor packets, you need to specify what information you want to collect about them.

Options

aggregate-intf-member-id	ID for a member of a link aggregation group (LAG) or an equal-cost multipath (ECMP) group.
chip-delay	The amount of time the packet takes to transit the ASIC.
egress-drop-reason	The reason the packet is dropped at egress.
flow-start-end-time	The flow start and end time.
ingress-drop-reason	The reason the packet is dropped at ingress.
inter-arrival-time	The time difference between two consecutive packets at ingress.
inter-departure-time	The time difference between two consecutive packets at egress.
queue-congestion-level	Queue congestion level
security-enable	Enable security analytics; specify that Denial-of-Service (DoS) attacks are reported to the collector.
shared-pool-congestion	Shared pool congestion level

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Flow-Based Telemetry

feb (Assigning a FEB to a Redundancy Group)

IN THIS SECTION

- Syntax | 6981
- Hierarchy Level | 6981
- Description | 6982
- Options | 6982
- Required Privilege Level | 6982
- Release Information | 6982

Syntax

```
feb slot-number (backup | primary);
```

Hierarchy Level

```
[edit chassis redundancy feb redundancy-group group-name]
```

Description

On M120 routers only, configure a Forwarding Engine Board (FEB) as part of a FEB redundancy group.

Options

slot-number—Slot number of the FEB. The range of values is from **0** to **5**.

backup—(Optional) For each redundancy group, you must configure exactly one backup FEB.

primary—(Optional) For each redundancy group, you can optionally configure one primary FEB.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring FEB Redundancy on the M120 Router

feb (chassis)

IN THIS SECTION

- Syntax | **6983**
- Hierarchy Level | **6983**
- Description | **6983**
- Required Privilege Level | **6983**
- Release Information | **6983**

Syntax

```
 feb slot slot-number{
        pfe {
            core name {
            channel name {
                tunnel-services {
                    bandwidth(100g | 10g | 1g | 2.5g | 200g | 20g | 300g | 30g | 400g | 40g |
480g | 50g | 60g | 70g | 80g | 90g | bandwidth);
                }
            }
        }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the Forwarding Engine Board (FEB) slot number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 24.1R1 for ACX Series Routers.

feb (Creating a Redundancy Group)

IN THIS SECTION

- Syntax | 6984
- Hierarchy Level | 6984
- Description | 6984
- Options | 6984
- Required Privilege Level | 6985
- Release Information | 6985

Syntax

```
feb {
    redundancy-group group-name {
        description description;
        feb slot-number (backup | primary);
        no-auto-failover;
    }
}
```

Hierarchy Level

```
[edit chassis redundancy]
```

Description

On M120 routers only, configure a Forwarding Engine Board (FEB) redundancy group.

Options

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring FEB Redundancy on the M120 Router

fec

IN THIS SECTION

- Syntax | 6985
- Description | 6985
- Default | 6986
- Options | 6986
- Required Privilege Level | 6986
- Release Information | 6986

Syntax

```
fec (efec | gfec | gfec-sdfec |hgfec | sd-fec | ufec | none);
```

Description

Enable forward error correction (FEC) mode.

Default

The default value is `gfec`.

Options

efec—Enhanced forward error correction (EFEC) is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

gfec—Gneric forward error correction (GFEC) mode is configured to detect and correct bit errors.

gfec-sdfec—GFEC and soft-decision forward error correction (SD-FEC) modes are configured to detect and correct bit errors.

hgfec—High gain forward error correction mode is configured to detect and correct bit errors.

sdfec—Soft-decision forward error correction mode is configured to detect and correct bit errors.

none—FEC mode is not configured.

ufec—Ultra Forward Error Correction (UFEC) mode is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

10-Gigabit Ethernet OTN Options Configuration Overview

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

Understanding Pre-FEC BER Monitoring and BER Thresholds

fec (MX Series)

IN THIS SECTION

- Syntax (M Series, MX Series, PTX Series) | 6987
- Syntax (ACX6360) | 6987
- Syntax (ACX5448-D) | 6987
- Hierarchy Level (M Series, MX Series, PTX Series) | 6987
- Hierarchy Level (ACX6360, ACX5448-D) | 6988
- Description | 6988
- Default | 6988
- Options | 6988
- Required Privilege Level | 6989
- Release Information | 6989

Syntax (M Series, MX Series, PTX Series)

```
fec (efec | gfec | gfec-sdfec |hgfec | sd-fec | ufec | none);
```

Syntax (ACX6360)

```
fec ( sdfec | sdfec15 | none);
```

Syntax (ACX5448-D)

```
fec ( sdfec | hgfec | sdfec15 | none);
```

Hierarchy Level (M Series, MX Series, PTX Series)

```
[edit interfaces interface-name otn-options]
```

Hierarchy Level (ACX6360, ACX5448-D)

```
[edit interfaces interface-name optics-options]
```

Description

Enable forward error correction (FEC) mode.

Default

If you do not specify a mode, the default mode is `gfec`. On PTX Series routers with P1-PTX-2-100G-WDM, the default value is `gfec-sdfec`. On PTX Series routers with PTX-5-100G-WDM and on MX Series routers with MIC3-100G-DWDM, the default value is `sdfec`.

Options

efec—(M Series, MX Series routers and PTX Series routers only) G.975.1 I.4 enhanced forward error correction (EFEC) is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

gfec—(M series, MX Series routers and PTX Series routers only) G.709 generic forward error correction (GFEC) mode is configured to detect and correct bit errors.

gfec-sdfec—(PTX Series routers only) GFEC and soft-decision forward error correction (SD-FEC) modes are configured to detect and correct bit errors.

hgfec—(MX Series routers only) High gain forward error correction mode is configured to detect and correct bit errors.

sdfec—(MX Series routers, PTX Series routers, and ACX6360 routers only) Sky-Compatible Soft-decision forward error correction mode is configured to detect and correct bit errors.

sdfec15—(ACX6360 routers only) Soft Decision Forward Error Correction with 15 percent overhead is configured to detect and correct bit errors.

none—(M Series and MX Series routers only) FEC mode is not configured.

NOTE: On MX Series routers with MIC3-100G-DWDM and PTX Series routers with PTX-5-100G-WDM, none option is not supported. The `fec` mode must be enabled on the MIC3-100G-DWDM MIC and the PTX-5-100G-WDM PIC.

ufec—(MX Series routers and PTX Series routers only) G.975.1 I.7 Ultra Forward Error Correction (UFEC) mode is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Statement and `gfec-sdfec` option introduced in Junos OS Release 13.2 for PTX Series routers. with P1-PTX-2-100G-WDM PIC.

Options `efec`, `gfec`, and `ufec` introduced in Junos OS Release 13.3 for MX Series routers. with MPC5E-100G10G, MPC5E-40G10G, MIC6-10G-OTN, and MIC6-100G-CFP2.

Options `efec`, `gfec`, and `ufec` introduced in Junos OS Release 14.1 for PTX Series routers. with P1-PTX-24-10G-W-SFPP.

Option `hgfec` introduced in Junos OS Release 15.1F5 for MX Series Routers with MIC3-100G-DWDM MIC.

Option `sdfec` introduced in Junos OS Release 15.1F5 for MX Series Routers with MIC3-100G-DWDM MIC.

Option `sdfec` introduced in Junos OS Release 15.1F6 for PTX Series Routers with PTX-5-100G-WDM PIC.

RELATED DOCUMENTATION

10-Gigabit Ethernet OTN Options Configuration Overview

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

Supported Forward Error Correction Modes on MX Series Routers

Supported Forward Error Correction Modes on PTX Series Routers

Understanding Pre-FEC BER Monitoring and BER Thresholds

Supported Forward Error Correction Modes on ACX6360 Router

fec (protocols ldp)

IN THIS SECTION

- Syntax | 6990
- Hierarchy Level | 6991
- Description | 6991
- Options | 6991
- Required Privilege Level | 6991
- Release Information | 6991

Syntax

```
fec fec-address {
    bfd-liveness-detection {
        detection-time threshold milliseconds;
        ecmp;
        failure-action {
            remove-nexthop;
            remove-route;
        }
        holddown-interval milliseconds;
        ingress-policy ingress-policy-name;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-transmit-interval milliseconds;
        multiplier detection-time-multiplier;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (0 | 1 | automatic);
    }
    no-bfd-liveness-detection;
    periodic-traceroute {
        disable;
```

```
        exp exp-value;
        fanout fanout-value;
        frequency minutes;
        paths number-of-paths;
        retries retry-attempts;
        source address;
        ttl ttl-value;
        wait seconds;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-systems-name protocols ldp oam],
[edit protocols ldp oam]
```

Description

Allows you to configure BFD for a specific LDP forwarding equivalence class (*FEC*).

Options

fec-address	Specify the FEC address.

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring BFD for LDP LSPs

fec (ether)

IN THIS SECTION

- Syntax | 6992
- Hierarchy Level | 6992
- Description | 6992
- Options | 6994
- Default | 6994
- Required Privilege Level | 6994
- Release Information | 6995

Syntax

```
fec(fec108 | fec74 | fec91 | none);
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name ether-options],
[edit dynamic-profiles name logical-systems name interfaces name ether-options],
[edit interfaces name ether-options]
```

Description

(PTX10001-36MR, PTX10003-8T, PTX10003-16T, PTX10004, PTX10008, QFX5220-32CD, and QFX5220-128C) Enable or disable RS-FEC (Reed-Solomon Forward Error Correction) for 25, 50, 100, 400-Gigabit Ethernet interface. By default, the Junos OS software enables or disables forward error correction based on the plugged-in optics. For instance, Junos OS software enables RS-FEC for 25G or 50G SR4 optics and disables RS-FEC for 25G or 50G LR4 optics.

If you enable or disable RS-FEC using this statement, this behavior applies to any 25-Gigabit Ethernet or 50-Gigabit Ethernet optical transceiver installed in the port associated with the interface.

You can configure forward error correction (FEC) clauses CL74 on 25-Gigabit and 50-Gigabit interfaces, and CL91 on 25-Gigabit or 50-Gigabit interfaces. Since the FEC clauses are applied by default on these interfaces, you must disable the FEC clauses if you do not want to apply them.

FEC mode is assigned by default. If there is a FEC mismatch, the link between nodes can go down. You must disable FEC mode if you do not want it assigned by default.

To view the FEC mode on an interface, use the `show interfaces interface-name` command.

For instance, consider two peer nodes, Node1 and Node2. Node1 is currently running Junos OS Evolved release 21.1R1 and the peer node (Node2) is currently running any Junos OS Evolved release prior to 21.1R1. The link between the two nodes will go down because of FEC mismatch. To enable the link to come back up, you must manually change the FEC on any one node by using the `fec` option at the `[edit interfaces <interface-name> ether-options]` hierarchy level.

Default FEC for prior to 21.1 Junos OS Evolved 21.1 is FEC74 and from 21.1 Junos OS Evolved is FEC91.

When FEC mode is FEC74 or FEC91 or NONE for all channels, FEC gets locked and interface is UP. For example, when FEC91 is configured together with FEC74 or NONE, on channel-0 and 1 or channel-2 and 3, link goes down for both channels.

Table 42: Link status for FEC combination of channel 2 and 3

Physical channel 0 (2)	Physical channel 1 (3)	Link Status
FEC74	FEC74	UP
FEC91	FEC91	UP
NONE	NONE	UP
FEC74	NONE	UP
NONE	FEC74	UP
FEC74	FEC91	DOWN

FEC91	FEC74	DOWN
FEC91	NONE	DOWN
NONE	FEC91	DOWN

There is no dependency between Channel 0(1) and channel 2(3). FEC74 can be configured with NONE to any interface et-x/y/x:(0/1/2/3) and link comes UP leads to the following:

- When you configure FEC91 on interface-x/y/z:0 then you must configure FEC91 on interface- x/y/z:1.
- When you configure FEC91 on interface-x/y/z:2 then you must configure FEC91 on interface- x/y/z:3.
- When you configure FEC74 or NONE on interface-x/y/z:0/1 then you must configure FEC74 or NONE on interface- x/y/z:2/3.

Options

fec108 IEEE 802.3by Clause 108, Reed-Solomon FEC (RS-FEC) on 25g links

fec74 FEC74 enabled

fec91 IEEE 802.3bj Clause 91, Reed-Solomon FEC (RS-FEC)

none FEC disabled

Default

Junos OS software automatically enables or disables 25-Gigabit Ethernet and 50-Gigabit Ethernet for FEC91 based on the type of pluggable optics used.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1

RELATED DOCUMENTATION

fec (gigether) | 6995

fec (gigether)

IN THIS SECTION

- Syntax | 6995
- Hierarchy Level | 6995
- Description | 6995
- Default | 6997
- Options | 6997
- Required Privilege Level | 6997
- Release Information | 6997

Syntax

```
fec (fec91 | fec74 | none)
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options]
```

Description

Supported on the following:

- ACX6360-OX Transponders
- ACX6360 routers
- MX Series Routers with MPC7E, MPC8E, and MPC9E
- MX10003 Router with MX10003 MPC
- MX204 Router
- PTX1000
- FPC3-PTX-U2 and FPC3-PTX-U3 on PTX5000
- QFX series switches

(ACX6360-OX Transponders, ACX6360 routers, MX Series Routers with MPC7E, MPC8E, and MPC9E, MX10003 Router with MX10003 MPC, MX204 Router, PTX1000, FPC3-PTX-U2 and FPC3-PTX-U3 on PTX5000, and QFX series switches) Enable or disable Reed-Solomon Forward Error Correction (RS-FEC) for a 100-Gigabit Ethernet interface. By default, the Junos OS software enables or disables FEC based on the plugged-in optics. For example, Junos OS enables RS-FEC for 100GbE SR4 optics and disables RS-FEC for 100GbE LR4 optics.

This statement allows you to override the default behavior and explicitly enable or disable RS-FEC. For example, you can extend the reach of 100GbE LR4 optics when you explicitly enable RS-FEC for the optics. RS-FEC is compliant with IEEE 802.3-2015 Clause 91.

Once you enable or disable RS-FEC using this statement, this behavior applies to any 100-Gigabit Ethernet optical transceiver installed in the port associated with the interface.

You can configure forward error correction (FEC) clauses CL74 on 25-Gigabit and 50-Gigabit interfaces, and CL91 on 100-Gigabit interfaces. Since the FEC clauses are applied by default on these interfaces, you must disable the FEC clauses if you do not want to apply them.

> **NOTE**: FPC-PTX-P1-A and FPC2-PTX-P1A on PTX5000 routers do not support RS-FEC. FPC3-SFF-PTX-1H and FP3-SFF-PTX-1T with PE-10-U-QSFP28 PIC and LR4 optics on PTX3000 and PTX5000 routers supports RS-FEC only on port 2. For PE-10-U-QSFP28 with LR4 optics, RS-FEC is the default FEC mode on port 2 and NONE is the default FEC mode on ports 0,1,3 through 9. For PE-10-U-QSFP28 with SR4 optics, RS-FEC is enabled by default on all ports. Do not modify the FEC mode on any port irrespective of the optics installed.

To check the FEC status, use the `show interfaces` *`interface-name`* command.

Default

Junos OS software automatically enables or disables RS-FEC based on the type of pluggable optics used.

Options

fec91 Enables RS-FEC. RS-FEC is compliant with IEEE 802.3-2015 Clause 91.

fec74 Enables RS-FEC. RS-FEC is compliant with IEEE 802.3-2015 Clause 74.

none Disables RS-FEC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1

RELATED DOCUMENTATION

Determining Transceiver Support for the PTX1000

Understanding the features of ACX6360

feed-name

IN THIS SECTION

- Syntax | 6998
- Hierarchy Level | 6998
- Description | 6998
- Required Privilege Level | 6998

Release Information | 6998

Syntax

```
feed-name feed-name
```

Hierarchy Level

```
[edit services web-filter profile profile-name],
[edit services web-filter profile profile-name dns-filter-template template-name],
[edit services web-filter profile profile-name dns-filter-template template-name term term-name]
```

Description

Configure the feed name. The feed name can be specified under profile has lesser precedence compared to the one configured under the template. The feed name configured in term takes higher precedence than the one configured under the template.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

feed-server

IN THIS SECTION

- Syntax | 6999
- Hierarchy Level | 7000
- Description | 7000
- Options | 7000
- Required Privilege Level | 7001
- Release Information | 7001

Syntax

```
feed-server name {
 (hostname hostname | url url);
 description description;
 feed-name name {
  description description;
  hold-interval seconds;
  path path;
  update-interval seconds;
 }
 hold-interval seconds;
 tls-profile tls-profile;
 update-interval seconds;
 validate-certificate-attributes {
  subject-or-subject-alternative-names;
 }
}
```

Hierarchy Level

```
[edit security dynamic-address]
```

Description

Security dynamic address feed-server

Options

name Security dynamic address feed-server name

hostname Hostname or IP address of feed-server

url URL,example:example.url.com;10.1.1.1/<folder-or-file-path>;Browser can verify its validity

description Text description of feed-server

hold-interval Time to keep IP entry when update failed (seconds)

- **Default:** 86400
- **Range:** 0 through 4294967295

tls-profile TLS profile

update-interval Interval to retrieve update (seconds)

- **Default:** 300
- **Range:** 30 through 4294967295

validate-certificate-attributes Validate some certificate attributes

- subject-or-subject-alternative-names—One of them must match feed server host name

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D25 on SRX Series Firewalls and on cSRX Container Firewall in Junos OS Release 22.2R1.

feeds (T640, T1600, and T4000 Routers with Six-Input DC Power Supply)

IN THIS SECTION

- Syntax | 7001
- Hierarchy Level | 7001
- Description | 7002
- Options | 7002
- Required Privilege Level | 7002
- Release Information | 7002

Syntax

```
feeds number-of-input-feeds;
```

Hierarchy Level

```
[edit chassis pem],
[edit chassis lcc lcc-number pem]
```

Description

Configure the number of input feeds connected to the six-input DC power supply on T640, T1600, or T4000 routers. The value assigned to the `feeds` statement must be equal to the number of input feeds provided to the power supply.

When providing four or five input feeds on standalone routers, you must configure the `feeds` statement at the `[edit chassis pem]` hierarchy level. When providing four or five input feeds to an LCC router in a routing matrix, you must configure the `feeds` statement at the `[edit chassis lcc` *`lcc-number`* `pem]` hierarchy level.

NOTE:

- Before configuring input feeds for your router, see the *T640 Core Router Hardware Guide*, *T1600 Core Router Hardware Guide*, or *T4000 Core Router Hardware Guide* for special considerations and for the number of input feeds supported by the router.
- All power supplies in the router must use the same number of inputs feeds.

Options

- **Range:** 4 through 6
- **Default:** 6

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring the Six-Input DC Power Supply on T Series Routers

fib-install-priority

IN THIS SECTION

- Syntax (PTX10001-36MR, PTX10004, PTX10008, PTX10016 Devices) | 7003
- Hierarchy Level | 7003
- Description | 7003
- Options | 7004
- Required Privilege Level | 7004
- Release Information | 7004

Syntax (PTX10001-36MR, PTX10004, PTX10008, PTX10016 Devices)

```
fib-install-priority <high/medium/>
```

Hierarchy Level

```
[edit policy-options policy-statement policy-name term term-name then]
```

Description

Assign a FIB install priority, as high or low, for matching routes. Non-matching routes are by default classified as unprioritized. When the routes are installed in the forwarding table, the number of forwarding table entries that routes of a particular priority get depends on the memory percentage allotted for the priority. Refer "fib-prioritization" on page 7007 to read more.

Issue the `run show route <prefix> extensive` command as shown in the following example to view the FIB install priority of a particular route.

```
user@host# run show route 10.0.1.1 extensive

inet.0: 33 destinations, 35 routes (33 active, 0 holddown, 0 hidden)
10.0.1.1/32 (1 entry, 1 announced)
TSI:
```

```
KRT in-kernel 10.0.1.1/32 -> {10.1.2.1}
   *Static Preference: 5
       FIB Installation Priority: INSTALL PRIORITY MEDIUM
       Next hop type: Router, Next hop index: 11045
       Address: 0x55eaafc82f1c
       Next-hop reference count: 3, key opaque handle: (nil), non-key opaque handle: (nil)
       Next hop: 10.1.2.1 via et-0/0/3.0, selected
       Session Id: 9
       State: <Active Int Ext>
       Local AS:100
       Age: 26
       Validation State: unverified
       Task: RT
       Announcement bits (1): 0-KRT
       AS path: I
       Thread: junos-main
```

Options

- high – Set the FIB install priority as high.
- medium – Set the FIB install priority as medium.

Required Privilege Level

- interface—To view this statement in the configuration.
- interface-control—To add this statement to the configuration.

Release Information

Command support introduced in Junos OS Evolved Release 22.4R1 for the PTX Series.

fib-local

IN THIS SECTION

- Syntax | 7005
- Hierarchy Level | 7005
- Description | 7005
- Required Privilege Level | 7005
- Release Information | 7006

Syntax

```
fib-local;
```

Hierarchy Level

```
[edit chassis fpc fpc-number route-localization]
```

Description

Configure the Packet Forwarding Engine on an FPC as FIB-local.

NOTE: At least, one Packet Forwarding Engine must be configured as **fib-local** for the commit operation to be successful. If you do not configure **fib-local** for the Packet Forwarding Engine, the CLI displays an appropriate error message and the commit fails.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring Packet Forwarding Engine FIB Localization

fib-next-hop-split

IN THIS SECTION

- Syntax | 7006
- Hierarchy Level | 7006
- Description | 7007
- Options | 7007
- Required Privilege Level | 7007
- Release Information | 7007

Syntax

```
fib-next-hop-split {
    labeled-isis;
}
```

Hierarchy Level

```
[edit routing-options forwarding-table]
```

Description

You must configure the `fib-next-hop-split` statement in conjunction with the "egress-chaining" on page 6210 and "transit" on page 30066 statements to enable route resolution over labeled IS-IS routes using flattened router next hops. Configuring the `fib-next-hop-split` statement alone does not have any effect.

When you configure the `no-labeled-isis` statement at the `[edit routing-options forwarding-table chained-composite-next-hop transit]` hierarchy level, the labeled IS-IS routes are flattened to ensure that the segment routing traffic-engineering (SPRING-TE) routes are not resolved any further over another composite next hop.

Options

labeled-isis Create composite-chained next hops for labeled IS-IS routes.

In addition to enabling `fib-next-hop-split`, you can also configure "egress-chaining" on page 6210 to reduce the number of chained next hops present in a route that involves resolution over multiple routes with labels.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

egress-chaining | **6210**

fib-prioritization

IN THIS SECTION

Syntax (PTX10001-36MR, PTX10004, PTX10008, PTX10016 Devices) | **7008**

- Hierarchy Level | 7008
- Description | 7008
- Options | 7009
- Required Privilege Level | 7009
- Release Information | 7009

Syntax (PTX10001-36MR, PTX10004, PTX10008, PTX10016 Devices)

```
chassis {
    fib-prioritization {
        high-priority-memory-percent 0...100
        medium-priority-memory-percent 0...100
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the memory allocation percent to install high and medium priority routes in the Forwarding Information Base (FIB) at device level. Reserve a percentage of memory on chassis for high-priority and medium priority routes, for IPv4, IPv6, and MPLS, BGP, IGP, static, aggregate, generated, and PRPD routes protocol data. The FIB prioritization configuration is a system-wide configuration. The amount of high or medium priority routes supported are different, based on memory available on the FPCs. Please note, that the system must be rebooted every time the FIB prioritization is changed.

```
chassis {
    fib-prioritization {
        high-priority-memory-percent 50
        medium-priority-memory-percent 30
    }
}
```

> **NOTE**: Total value of high priority and medium priority percentage must not exceed 100

Options

fib-prioritization—FIB prefix prioritization to install high and medium priority routes at chassis-level.

- high-priority-memory-percent – Memory allocation for high priority protocol routes. Value range for memory percent is 0 to 100.
- medium-priority-memory-percent – Memory allocation for medium priority protocol routes. Value range for memory percent is 0 to 100.

Required Privilege Level

- interface—To view this statement in the configuration.
- interface-control—To add this statement to the configuration.

Release Information

Command support introduced in Junos OS Evolved Release 22.4R1 for the PTX Series.

fib-remote

IN THIS SECTION

- Syntax | 7010
- Hierarchy Level | 7010
- Description | 7010
- Required Privilege Level | 7010
- Release Information | 7010

Syntax

```
fib-remote;
```

Hierarchy Level

```
[edit chassis fpc fpc-number route-localization]
```

Description

Configure the Packet Forwarding Engine on an FPC as FIB-remote.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring Packet Forwarding Engine FIB Localization

fibre-channel (Alarm)

IN THIS SECTION

- Syntax | 7011
- Hierarchy Level | 7011
- Description | 7011

- Options | 7011
- Required Privilege Level | 7011
- Release Information | 7011

Syntax

```
fibre-channel {
    link-down  (red | yellow | ignore);
}
```

Hierarchy Level

```
[edit chassis alarm],
[edit chassis interconnect-device name alarm],
[edit chassis node-group name alarm]
```

Description

Configure alarms for a Fibre Channel interface.

Options

The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

fibre-channel (Family Interfaces)

IN THIS SECTION

- Syntax | 7012
- Hierarchy Level | 7012
- Description | 7012
- Options | 7012
- Required Privilege Level | 7012
- Release Information | 7013

Syntax

```
fibre-channel {
    port-mode (f-port | np-port);
}
```

Hierarchy Level

```
[edit interfaces vlan unit logical-unit-number family],
[edit interfaces interface-name unit logical-unit-number family]
```

Description

Configure the port mode for FCoE VLAN interfaces and native FC interfaces.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric

Configuring an FCoE VLAN Interface on an FCoE-FC Gateway

Converting an Ethernet Interface To a Fibre Channel Interface

show fibre-channel interfaces | **19980**

show vlans

Understanding Interfaces on an FCoE-FC Gateway

fibre-channel (Port)

IN THIS SECTION

- Syntax | **7013**
- Hierarchy Level | **7014**
- Description | **7014**
- Options | **7014**
- Required Privilege Level | **7014**
- Release Information | **7014**

Syntax

```
fibre-channel {
    port-range {
        port-range-low port-range-high;
```

```
    }
}
```

Hierarchy Level

```
[edit chassis (QFX Series) fpc fpc-id pic pic-id]
```

Description

Specify a range of ports to carry FC traffic when the switch is configured as an FCoE-FC gateway.

Options

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- *Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric*
- *Configuring a Physical Fibre Channel Interface*
- show fibre-channel interfaces | **19980**
- *Understanding Interfaces on an FCoE-FC Gateway*

fibrechannel-options

IN THIS SECTION

- Syntax | 7015
- Hierarchy Level | 7015
- Description | 7015
- Options | 7015
- Required Privilege Level | 7016
- Release Information | 7016

Syntax

```
fibrechannel-options {
    bb-sc-n
    (loopback | no-loopback);
    speed (auto-negotiation | 2g | 4g | 8g);
}
```

Hierarchy Level

```
[edit interfaces  interface-name]
```

Description

Configure FC interface properties such as speed and loopback mode.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel interfaces | 19980

Converting an Ethernet Interface To a Fibre Channel Interface

fields (Flat-File Accounting Options)

IN THIS SECTION

- Syntax | 7016
- Hierarchy Level | 7018
- Description | 7018
- Options | 7018
- Required Privilege Level | 7019
- Release Information | 7019

Syntax

```
fields {
    all-fields;
    egress-stats {
        all-fields;
        input-bytes;
        input-packets;
```

```
        output-bytes;
        output-packets;
        queue-id;
        red-drop-bytes;
        red-drop-packets;
        tail-drop-packets;
        total-drop-packets;
    }
    general-param {
        all-fields;
        accounting-type;
        descr;
        line-id;
        logical-interface;
        nas-port-id;
        physical-interface;
        routing-instance;
        timestamp;
        user-name;
        vlan-id;
    }
    ingress-stats {
        all-fields;
        drop-packets;
        input-bytes;
        input-packets;
        output-bytes;
        output-packets;
        queue-id;
    }
    l2-stats {
        all-fields;
        input-mcast-bytes;
        input-mcast-packets;
    }
    overall-packet {
        all-fields;
        input-bytes;
        input-discards;
        input-errors;
        input-packets;
        inputv6-bytes;
        inputv6-packets;
```

```
        output-bytes;
        output-errors;
        output-packets;
        outputv6-bytes;
        outputv6-packets;
        input-v4-bytes;
        input-v4-packets;
        output-v4-bytes;
        output-v4-packets;
        input-bytes-per-sec;
        input-packets-per-sec;
    }
    service-accounting;
}
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name]
```

Description

Specify the accounting statistics and the nonstatistical information to be collected for an interface and recorded in the accounting flat file created by the profile.

Options

all-fields Include all available statistical fields in the accounting file. Many fields are available for both logical interfaces and physical interfaces, but some fields are available only for one or the other interface type.

service-accounting Include the filter counts in bytes for the inet input filter, inet output filter, inet6 input filter, and inet6 output filter in the service accounting flat file. Statistics reported are the running total values.

> **NOTE**: Starting in Junos OS Release 18.4R1, the `service-accounting` option is no longer supported. If included in a configuration, it is ignored.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Configuring Service Accounting in Local Flat Files

Flat-File Accounting Overview

fields (for Interface Profiles)

IN THIS SECTION

- Syntax | 7019
- Hierarchy Level | 7020
- Description | 7020
- Options | 7020
- Required Privilege Level | 7021
- Release Information | 7021

Syntax

```
fields {
    field-name;
}
```

Hierarchy Level

```
[edit accounting-options interface-profile profile-name]
```

Description

Statistics to collect in an accounting-data log file for an interface.

Options

field-name—Name of the field:

- `input-bytes`—Input bytes
- `input-errors`—Generic input error packets
- `input-multicast`—Input packets arriving by multicast
- `input-packets`—Input packets
- `input-unicast`—Input unicast packets
- `output-bytes`—Output bytes
- `output-errors`—Generic output error packets
- `output-multicast`—Output packets sent by multicast
- `output-packets`—Output packets
- `output-unicast`—Output unicast packets
- `rpf-check-total-packets`—Total packets failing reverse-path-forwarding check
- `rpf-check-total-bytes`—Total bytes failing reverse-path-forwarding check
- `rpf-check-bytes`—Bytes failing reverse-path-forwarding check
- `rpf-check-packets`

  —Packets failing reverse-path-forwarding check
- `rpf-check6-bytes`

—IPv6 bytes failing reverse-path-forwarding check

- `rpf-check6-packets`

  —IPv6 packets failing reverse-path-forwarding check

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Interface Profile

fields (for Routing Engine Profiles)

IN THIS SECTION

- Syntax | 7022
- Hierarchy Level | 7022
- Description | 7022
- Options | 7022
- Required Privilege Level | 7023
- Release Information | 7023

Syntax

```
fields {
    field-name;
}
```

Hierarchy Level

```
[edit accounting-options routing-engine-profile profile-name]
```

Description

Statistics to collect in an accounting-data log file for a Routing Engine.

Options

field-name—Name of the field:

- `cpu-load-1`—Average system load over the last 1 minute
- `cpu-load-5`—Average system load over the last 5 minutes
- `cpu-load-15`—Average system load over the last 15 minutes
- `date`—Date, in YYYYMMDD format
- `host-name`—Hostname for the router
- `time-of-day`—Time of day, in HHMMSS format
- `uptime`—Time since last reboot, in seconds
- `cpu-usage-5sec`—System CPU usage for last 5 seconds
- `cpu-usage-1min`—Average system CPU usage for the last one minute
- `cpu-usage-5min`—Average system CPU usage for the last 5 minutes
- `free-mem`—Free memory available in MB
- `total-mem`—Total memory available in MB

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Routing Engine Profile

file

IN THIS SECTION

- Syntax | 7023
- Hierarchy Level | 7023
- Description | 7024
- Options | 7024
- Required Privilege Level | 7024
- Release Information | 7024

Syntax

```
file <filename> <files files> <match match> <size size> <(world-readable | no-world-readable)>;
```

Hierarchy Level

```
[edit security dynamic-vpn traceoptions (Security Dynamic VPN)],
```

Description

Configure the trace file options. Name of the file to receive the output of the tracing operation.

In Junos OS Release 21.4R1, we've deprecated the `file` configuration statement and we might remove it completely in a future release.

Options

filename Name of file in which to write trace information

files Maximum number of trace files

- **Default:** 3
- **Range:** 2 through 1000

match Regular expression for lines to be logged

no-world-readable Don't allow any user to read the log file

size Maximum trace file size

- **Default:** 128000
- **Range:** 10240 through 1073741824

world-readable Allow any user to read the log file

Required Privilege Level

trace

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

traceoptions (Security Dynamic VPN) | **29786**

Dynamic VPN Overview

file

IN THIS SECTION

- Syntax | 7025
- Description | 7025
- Options | 7025
- Required Privilege Level | 7026
- Release Information | 7026

Syntax

```
file <archive | change-owner | change-permission | checksum | compare | compress | copy | delete |
delete-directory | link | list | make-directory | routing-instance | rename | show | staging-
directory directory-location>
<source-address source-address>
```

Description

Archive files from the device, copy files to and from the router or switch, calculate the file checksum, compare files, delete a file from the device, list files on the device, rename a file, show file contents, or show the local address to initiate a connection.

Options

archive (Optional)	Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.
change-owner (Optional)	Change owner of file.
change-permission (Optional)	Change permission of file.
checksum (Optional)	Calculate the Message Digest 5 (MD5) checksum of a file.

compare (Optional)	Compare two local files and describe the differences between them in default, context, or unified output styles.
Compress (Optional)	file compress.
copy (Optional)	Copy files from one place to another on the local device or between the local device and a remote system.
delete (Optional)	Delete a file on the local device.
delete-directory (Optional)	Delete a directory.
link (Optional)	Create link between files.
list (Optional)	Display a list of files on the local device.
make-directory (Optional)	Create a new directory.
rename (Optional)	Rename a file on the local device.
routing-instance (Optional)	Name of routing instance for host attempt.
show (Optional)	Display the contents of a file.
source-address (Optional)	Specify the source address of the local file.
staging-directory (Optional)	Staging directory on Routing Engine.

Required Privilege Level

maintenance

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Viewing Files and Directories on a Juniper Networks Device

file (Security Certificates)

IN THIS SECTION

- Syntax | 7027
- Hierarchy Level | 7027
- Description | 7027
- Options | 7027
- Required Privilege Level | 7027
- Release Information | 7028

Syntax

```
file certificate-filename;
```

Hierarchy Level

```
[edit security certificates certification-authority ca-profile-name]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Specify the file from which to read the digital certificate.

Options

certificate-filename—File from which to read the digital certificate.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Specifying a File to Read the Digital Certificate

file (Analytics)

IN THIS SECTION

- Syntax | 7028
- Hierarchy Level | 7029
- Description | 7029
- Default | 7029
- Options | 7029
- Required Privilege Level | 7030
- Release Information | 7030

Syntax

```
file filename {
    files number-of-files;
    size size;
}
```

Hierarchy Level

```
[edit services analytics collector local]
[edit services analytics queue-statistics]
[edit services analytics traffic-statistics]
```

Description

Enable the logging of queue or traffic monitoring statistics in a local file. This statement does not enable monitoring.

Default

This feature is disabled by default.

Options

filename Specify a filename for storing queue and traffic monitoring statistics in the Comma-separated Values (CSV) format. The file is stored in the **/var/log/** directory of your device.

If you do not specify a filename, the data is not stored in a file.

> **NOTE**: In Junos OS Release 13.2X51-D15 or later, you configure a single filename to store both queue and traffic monitoring statistics. In Junos OS Release 13.2X51-D10 and earlier, you configure separate files for storing monitoring data, one for queue statistics, and another for traffic statistics.

files ***number-of-files*** Specify the number of files to store locally. After the number of files with the maximum file size is reached, the system starts over and writes the data to the first file.

- **Range:** 2 to 1,000 files.

size ***size*** Configure the file size in megabytes (MB).

- **Syntax:** *xm* to specify MB.
- **Range:** 10 to 4095 MB

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
13.2X51-D15	In Junos OS Release 13.2X51-D15 or later, you configure a single filename to store both queue and traffic monitoring statistics.
13.2X51-D10	In Junos OS Release 13.2X51-D10 and earlier, you configure separate files for storing monitoring data, one for queue statistics, and another for traffic statistics.

RELATED DOCUMENTATION

Network Analytics Overview

file (App Engine Virtual Machine Management Service)

IN THIS SECTION

- Syntax | 7031
- Hierarchy Level | 7031
- Description | 7031
- Options | 7031
- Required Privilege Level | 7032
- Release Information | 7032

Syntax

```
file {filename <files number> | match | no-world-readable | size size | world-readable
}
```

Hierarchy Level

```
[edit system processes app-engine-virtual-machine-management-service traceoptions]
```

Description

Trace file information for the Virtual Machine Management Daemon (vmmd), which communicates with the host OS.

Options

filename Name of the file in which the trace information is stored. By default, the file is created in the **/var/log** directory.

files *number* (Optional) Maximum number of trace files. When a trace file reaches the size specified by the `size` option, the filename is appended with 0 and compressed. For example, when a trace file named trace-file-log reaches size *size*, it is compressed and renamed as trace-file-log.0.gz. When trace-file-log reaches size *size* for the second time, trace-file-log.0.gz is renamed as trace-file-log.1.gz and trace-file-log is compressed and renamed as trace-file-log.0.gz. This renaming sheme ensures that the older logs have a greater index number. When number of trace files reaches *number*, the oldest file is deleted.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000
- **Default:** 10

match Refine the output to include only those lines that match the given regular expression.

no-world-readable Restrict file access to the user who created the trace files.

size*size* Maximum size of each trace file . By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes

(MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the files option.

- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

`world-readable` Enable unrestricted file access.

Required Privilege Level

system-control

Release Information

Statement introduced in Junos OS Release 15.1F3.

> **NOTE**: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

- traceoptions (App Engine Virtual Machine Management Service) | 29497
- level (App Engine Virtual Machine Management Service) | 9719
- flag (App Engine Virtual Machine Management Service) | 7204

file archive

IN THIS SECTION

- Syntax | 7033
- Description | 7033
- Options | 7033
- Required Privilege Level | 7034
- Output Fields | 7034

Sample Output | 7034
Release Information | 7035

Syntax

```
file archive destination destination source source
<compress>
<exclude> filepattern
<routing-instance>
<source-address>
```

Description

Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.

For information on valid filename and URL formats, see *Format for Specifying Filenames and URLs in Junos OS CLI Commands*.

Options

destination *destination*	Destination of the archived file or files. Specify the destination as a URL or filename. The Junos OS adds one of the following suffixes if the destination filename does not already have it: • For archived files—The suffix **.tar** • For archived and compressed files—The suffix **.tgz**
source *source*	Source of the original file or files. Specify the source as a URL or filename.
compress	(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix **.tgz**.
exclude	(Optional) Specify the file pattern to exclude. `exclude` helps to exclude files that delay compression or files that need not be compressed.
routing-instance *routing-instance-name*	(Optional) Name of the routing-instance.

source-address *address* (Optional) Local address to use in originating the connection.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

file archive (Multiple Files)

The following sample command archives all messages files in the local directory `/var/log/` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages* destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.

user@host>
```

file archive (Single File)

The following sample command archives a single messages file in the local directory `/var/log/` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.

user@host>
```

file archive (with Compression)

The following sample command archives and compresses all messages files in the local directory `/var/log/` as the single file `messages-archive.tar`.

```
user@host> file archive compress source /var/log/messages* destination /var/log/messages-
archive.tgz
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

File Archive Using Secure Copy Protocol (scp) with 'source-address' and 'routing-instance' options

To use the scp command to archive a file with the `source-address` and `routing-instance` options, enter the following command:

```
user@host> file archive source source destination scp:// destination source-address address
routing-instance instance-name
```

File Archive Compress (with Exclude Option)

The following sample command archives and compresses all messages files in the local directory `/var/log/` and excludes the single file pattern specified.

```
user@host>file archive compress source /var/log destination /var/tmp/sample exclude trace*
/bin/tar: Removing leading `/' from member names
/bin/tar: /var/log/journal/dbf38b9cae4d11ec862e4fb4fcef8787/system.journal: file changed as we
read it
```

Release Information

Command introduced before Junos OS Release 7.4.

`exclude` option added in Junos OS Release 20.3R1 and Junos OS Evolved Release 20.3R1.

`routing-instance` option added in Junos OS Release 18.4R1.

`source-address` option added in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Format for Specifying Filenames and URLs in Junos OS CLI Commands

file (Associating with a Profile)

IN THIS SECTION

- Syntax | 7036
- Hierarchy Level | 7036
- Description | 7036
- Options | 7037
- Required Privilege Level | 7037
- Release Information | 7037

Syntax

```
file filename;
```

Hierarchy Level

```
[edit accounting-options class-usage-profile profile-name],
[edit accounting-options filter-profile profile-name],
[edit accounting-options interface-profile profile-name],
[edit accounting-options mib-profile profile-name],
[edit accounting-options routing-engine-profile profile-name]
```

Description

Specify the accounting log file associated with the profile.

Options

filename—Name of the log file. You must specify a filename already configured in the `file` statement at the `[edit accounting-options]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The `[edit accounting-options mib-profile profile-name]` hierarchy added in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring the Interface Profile

Configuring the Filter Profile

Configuring the MIB Profile

Configuring the Routing Engine Profile

file checksum md5

IN THIS SECTION

- Syntax | 7038
- Description | 7038
- Options | 7038
- Required Privilege Level | 7038
- Output Fields | 7038
- Sample Output | 7038
- Release Information | 7038

Syntax

```
file checksum md5 <pathname> filename
```

Description

Calculate the Message Digest 5 (MD5) checksum of a file.

Options

pathname (Optional) Path to a filename.

filename Name of a local file for which to calculate the MD5 checksum.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = $ABC123
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Checksum Hashes for a Commit Script

Configuring Checksum Hashes for an Event Script
Configuring Checksum Hashes for an Op Script
Configuring Checksum Hashes for an SNMP Script
Executing an Op Script from a Remote Site
file checksum sha-256 | **7039**
file checksum sha1 | **7041**

file checksum sha-256

IN THIS SECTION

- Syntax | **7039**
- Description | **7039**
- Options | **7039**
- Required Privilege Level | **7040**
- Output Fields | **7040**
- Sample Output | **7040**
- Release Information | **7040**

Syntax

```
file checksum sha-256 <pathname> filename
```

Description

Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.

Options

pathname (Optional) Path to a filename.

filename Name of a local file for which to calculate the SHA-256 checksum.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum sha-256

```
user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =$ABC123
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Checksum Hashes for a Commit Script

Configuring Checksum Hashes for an Event Script

Configuring Checksum Hashes for an Op Script

Configuring Checksum Hashes for an SNMP Script

Executing an Op Script from a Remote Site

file checksum md5 | **7037**

file checksum sha1 | **7041**

file checksum sha1

IN THIS SECTION

- Syntax | 7041
- Description | 7041
- Options | 7041
- Required Privilege Level | 7041
- Output Fields | 7041
- Sample Output | 7042
- Release Information | 7042

Syntax

```
file checksum sha1 <pathname> filename
```

Description

Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.

Options

pathname (Optional) Path to a filename.

filename Name of a local file for which to calculate the SHA-1 checksum.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax

SHA1 (/var/db/scripts/commitscript.slax) = $ABC123
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

- *Configuring Checksum Hashes for a Commit Script*
- *Configuring Checksum Hashes for an Event Script*
- *Configuring Checksum Hashes for an Op Script*
- *Configuring Checksum Hashes for an SNMP Script*
- *Executing an Op Script from a Remote Site*
- file checksum md5 | **7037**
- file checksum sha-256 | **7039**

file (Commit Scripts)

IN THIS SECTION

- Syntax | **7043**
- Hierarchy Level | **7043**
- Description | **7043**
- Options | **7043**
- Required Privilege Level | **7043**
- Release Information | **7043**

Syntax

```
file filename {
    allow-transients;
    cert-file path;
    checksum hash-algorithm hash-value;
    optional;
    refresh;
    refresh-from url;
    routing-instance routing-instance-name;
    source url;
}
```

Hierarchy Level

```
[edit system scripts commit]
```

Description

Enable a Junos OS commit script that is located in the **/var/db/scripts/commit** directory on the hard disk or the **/config/scripts/commit** directory on the flash drive. During a commit operation, enabled commit scripts are executed in the order specified in the configuration.

Options

filename Name of a Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative syntaX (SLAX) file containing a commit script.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Storing and Enabling Scripts

Controlling Execution of Commit Scripts During Commit Operations

file compare

IN THIS SECTION

- Syntax | 7044
- Description | 7044
- Options | 7045
- Required Privilege Level | 7045
- Output Fields | 7045
- Sample Output | 7045
- Release Information | 7047

Syntax

```
file compare (files filename filename)
<context | unified>
<ignore-white-space>
```

Description

Compare two local files and describe the differences between them in default, context, or unified output styles:

- `Default`—In the first line of output, `c` means lines were changed between the two files, `d` means lines were deleted between the two files, and `a` means lines were added between the two files. The numbers preceding this alphabetical marker represent the first file, and the lines after the alphabetical marker represent the second file. A left angle bracket (<) in front of output lines refers to the first file. A right angle bracket (>) in front of output lines refers to the second file.

- `Context`—The display is divided into two parts. The first part is the first file; the second part is the second file. Output lines preceded by an exclamation point (`!`) have changed. Additions are marked with a plus sign (`+`), and deletions are marked with a minus sign (`-`).
- `Unified`—The display is preceded by the line number from the first and the second file (*xx,xxx,x*). Before the line number, additions to the file are marked with a plus sign (`+`), and deletions to the file are marked with a minus sign (`-`). The body of the output contains the affected lines. Changes are viewed as additions plus deletions.

Options

files *filename*	Names of two local files to compare.
context	(Optional) Display output in context format.
ignore-white-space	(Optional) Ignore changes in the amount of white space.
unified	(Optional) Display output in unified format.

Required Privilege Level

none

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

file compare files

```
user@host> file compare files /tmp/one /tmp/two
100c100
<             full-name "File 1";
---
>             full-name "File 2";
102c102
<             class foo; # 'foo' is not defined
---
>             class super-user;
```

file compare files context

```
user@host> file compare files /tmp/one /tmp/two context
*** /tmp/one    Wed Dec  3 17:12:50 2003
--- /tmp/two    Wed Dec  3 09:13:14 2003
***************
*** 97,104 ****
              }
          }
          user bill {
!             full-name "Bill Smith";
!             class foo; # 'foo' is not defined
              authentication {
                  encrypted-password $ABC123;
              }
--- 97,105 ----
              }
          }
          user bill {
!             full-name "Bill Smith";
!             uid 1089;
!             class super-user;
              authentication {
                  encrypted-password $ABC123;
              }
```

file compare files unified

```
user@host> file compare files /tmp/one /tmp/two unified
--- /tmp/one    Wed Dec  3 17:12:50 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -97,8 +97,9 @@
             }
         }
         user bill {
-            full-name "Bill Smith";
-            class foo; # 'foo' is not defined
+            full-name "Bill Smith";
+            uid 1089;
+            class super-user;
```

```
            authentication {
                encrypted-password $ABC123;
            }
```

file compare files unified ignore-white-space

```
user@host> file compare files /tmp/one /tmp/two unified ignore-white-space
--- /tmp/one    Wed Dec  3 09:13:10 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -99,7 +99,7 @@
         user bill {
             full-name "Bill Smith";
             uid 1089;
-            class foo; # 'foo' is not defined
+            class super-user;
             authentication {
                  encrypted-password $ABC123; # SECRET-DATA
             }
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Format for Specifying Filenames and URLs in Junos OS CLI Commands

Viewing Core Files from Junos OS Processes

file (Configuring a Log File)

IN THIS SECTION

- Syntax | 7048
- Hierarchy Level | 7048

- Description | 7048
- Options | 7048
- Required Privilege Level | 7049
- Release Information | 7049

Syntax

```
file filename {
    archive-sites {
        site-name;
    }
    backup-on-failure (master-and-slave | master-only);
    compress;
    files number;
    nonpersistent;
    push-backup-to-master;
    size bytes;
    start-time time;
    transfer-interval minutes;
}
```

Hierarchy Level

```
[edit accounting-options]
```

Description

Specify a log file to be used for accounting data.

Options

filename—Name of the file in which to write accounting data.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Accounting-Data Log Files

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Flat-File Accounting Overview

file copy

IN THIS SECTION

- Syntax | 7049
- Description | 7050
- Options | 7050
- Required Privilege Level | 7051
- Sample Output | 7051
- Release Information | 7052

Syntax

```
file copy source destination
<source-address source-address>
<staging-directory directory-location>
```

Description

Copy files from one location to another location on the local device or to a location on a remote device reachable by the local device.

For information on valid file name and URL formats, see *Format for Specifying Filenames and URLs in Junos OS CLI Commands*.

CAUTION: Starting with Junos OS Release 15.1, the `sslv3-support` option is not available for configuration with the `set system services xnm-ssl` and `file copy` commands. SSLv3 is no longer supported and available.

For all releases prior to and including Junos OS Release 14.2, SSLv3 is disabled by default at runtime. The `sslv3-support` option is hidden and deprecated in Junos OS Release 14.2 and earlier releases. However, you can use the `set system services xnm-ssl sslv3-support` command to enable SSLv3 for a Junos XML protocol client application to use as the protocol to connect to the Junos XML protocol server on a router, and you can use the `file copy source destination sslv3-support` command to enable the copying of files from an SSLv3 URL.

Using SSLv3 presents a potential security vulnerability, and we recommend that you not use SSLv3. For more details about this security vulnerability, see the Juniper Networks Knowledge Base articlehere.

NOTE: If you define an ordered set of ciphers, key exhange methods, or message authentication codes (MACs) at the `[edit system services ssh]` hierarchy level, the newly-defined set is used when copying files using secure copy protocol (scp). For more information, see *Configuring the SSH Service to Support Legacy Cryptography*.

Options

source	Specify the source URL.
destination	Specify the destination URL.
source-address ***source-address***	(Optional) Specify the local address to use in originating the connection.
staging-directory ***directory-location***	(Optional) Specify the staging directory on Routing Engine

> **NOTE**: If you are using a literal IPv6 address in your command entry, the address must be enclosed in [ ] brackets. This conforms to RFC 2732.

Required Privilege Level

maintenance

Sample Output

Following are some `file copy` examples.

Copy a File from the Local Device to a Personal Computer

```
user@host> file copy /var/tmp/rpd.core.4 mypc:/c/junipero/tmp

...transferring.file...... |          0 KB |   0.3 kB/s | ETA: 00:00:00 | 100%
```

Copy a Configuration File between Routing Engines

The following sample command copies a configuration file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy /config/juniper.conf re1:/var/tmp/copied-juniper.conf
```

Copy a File Using File Transfer Protocol which Requires a Password

You can use FTP with the `file copy` command for additional privacy. You will be prompted for a password. Enter the following command to enable FTP:

```
root@host> file copy filename ftp://user@hostname/filename
```

In the following example, **/config/juniper.conf** is the local file and **hostname** is the FTP server:

```
root@host> file copy /config/juniper.conf ftp://user@hostname/juniper.conf
Password:
```

```
Receiving ftp: //user@hostname/juniper.conf (2198 bytes): 100%
2198 bytes transferred in 0.0 seconds (2.69 MBps)
```

NOTE: Starting in Junos OS Evolved release 22.2R2, the password prompt has changed from `Password for user@hostname:` to `Password:`

Copy a File Using a Staging Directory

The following sample command copies a file using a staging directory

```
user@host> file copy re1:/var/tmp/junos-install-x.log /root/ staging-directory /var/tmp/tmp1
```

Release Information

Command introduced before Junos OS Release 7.4.

`source-address` option added in Junos OS Release 7.4.

`staging-directory` option added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Format for Specifying Filenames and URLs in Junos OS CLI Commands

Default Directories for Junos OS File Storage on the Network Device

file delete

IN THIS SECTION

- Syntax | 7053
- Description | 7053
- Options | 7053
- Required Privilege Level | 7053

- Output Fields | 7053
- Sample Output | 7053
- Release Information | 7054

Syntax

```
file delete filename
<purge>
```

Description

Delete a file on the local router or switch.

Options

filename Name of the file to delete. For a routing matrix, include chassis information in the filename if the file to be deleted is not local to the Routing Engine from which the command is issued.

purge (Optional) Overwrite regular files before deleting them.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

file delete

```
user@host> file list /var/tmp
dcd.core
```

```
rpd.core
snmpd.core

user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file delete (Routing Matrix)

```
user@host> file list lcc0-re0:/var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete lcc0-re0:/var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

Release Information

Command introduced before Junos OS Release 7.4.

file (Event Scripts)

IN THIS SECTION

- Syntax | 7055
- Hierarchy Level | 7055
- Description | 7055
- Options | 7056
- Required Privilege Level | 7056
- Release Information | 7056

Syntax

```
file filename {
    cert-file path;
    checksum hash-algorithm hash-value;
    dampen {
        dampen-options {
            cpu-factor cpu-factor;
            line-interval line-interval;
            time-interval time-interval;
        }
    }
    python-script-user username;
    refresh;
    refresh-from url;
    remote-execution {
        remote-hostname {
            passphrase user-password;
            username user-login;
        }
    }
    routing-instance routing-instance-name;
    source url;
    start {
        start-options {
            mem-factor percentage;
        }
    }
}
```

Hierarchy Level

```
[edit event-options event-script]
```

Description

Enable a Junos OS event script that is located in the **/var/db/scripts/event** directory on the hard disk or the **/config/scripts/event** directory on the flash drive. After enabling an event script, an event policy can execute the event script when triggered.

Options

filename Name of a Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative syntaX (SLAX) file containing an event script.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Store and Enable Junos Automation Scripts

Enable and Execute Event Scripts

file (Extended DHCP Relay Agent and Helpers Trace Options)

IN THIS SECTION

- Syntax | 7057
- Hierarchy Level | 7057
- Description | 7057
- Options | 7057
- Required Privilege Level | 7057
- Release Information | 7057

Syntax

```
file filename  <files number> <match regular-expression> <size bytes> <world-readable | no-world-
readable>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay traceoptions],
[edit forwarding-options helpers traceoptions]
```

Description

Configure information about the DNS and TFTP packet-forwarding files that contain trace logging information.

Options

filename—Name of the file containing the trace information.

- **Default: /var/log/sampled**

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing BOOTP, DNS, and TFTP Forwarding Operations

file (Flat-File Accounting Options)

IN THIS SECTION

- Syntax | 7058
- Hierarchy Level | 7058
- Description | 7058
- Options | 7058
- Required Privilege Level | 7058
- Release Information | 7059

Syntax

```
file filename;
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name]
```

Description

Specify the name of the accounting file created by a flat-file profile. By default, the filename becomes the name of the local directory where the accounting file is backed up: `/var/log/pfedBackup/filename`.

Options

filename Name of the accounting file. The complete output filename is in the format *filename.hostname.file-number_timestamp*.gz.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Flat-File Accounting Overview

file (JET)

IN THIS SECTION

- Syntax (Junos OS) | **7059**
- Syntax (Junos OS Evolved) | **7060**
- Hierarchy Level | **7061**
- Description | **7061**
- Options | **7061**
- Required Privilege Level | **7062**
- Release Information | **7062**

Syntax (Junos OS)

```
file filename {
    arguments arguments;
    cert-file path;
    checksum hash-algorithm hash-value;
    daemonize;
    optional;
    refresh;
```

```
    refresh-from;
    respawn-on-normal-exit;
    routing-instance;
    source;
    start {
        start-options {
            mem-factor percentage;
        }
    }
    traceoptions {
        file <filename> <files number> <match regex> <size size> <world-readable | no-world-
readable>;
        flag flag;
        no-remote-trace;
    }
    username username;
}
```

Syntax (Junos OS Evolved)

```
file filename {
    arguments arguments;
    cert-file path;
    checksum hash-algorithm hash-value;
    daemonize;
    interpreter (bash | python | python3);
    optional;
    refresh;
    refresh-from;
    respawn-on-normal-exit;
    routing-instance;
    source;
    start {
        start-options {
            mem-factor percentage;
        }
    }
    traceoptions {
        file <filename> <files number> <match regex> <size size> <world-readable | no-world-
readable>;
        flag flag;
```

```
        no-remote-trace;
    }
    username username;
}
```

Hierarchy Level

```
[edit system extensions extension-service application]
```

Description

For files in the `[edit system extensions extension-service application]` hierarchy level, specify the configuration for each file in the extension-service application.

Options

arguments *arguments*
Specify the command-line arguments called by a JET application. A program can take any number of command-line arguments. Enter the arguments in the way the application expects. Developer must supply this information.

daemonize
Specify the file as daemonized. An application runs as a daemonized process in the background. An application configured to run as a daemonized process is automatically triggered upon commit. A non-daemonized application must be triggered manually from the command-line client.

filename
Local filename of the script file.

interpreter (bash | python | python3)
Specify whether a device running Junos OS Evolved should run a daemonized on-device JET application using Bash, Python 2, or Python 3.

> **NOTE**: Starting in Junos OS Evolved Release 22.3R1, Python 2.7 is no longer supported. The `python` statement is deprecated. Use the `python3` statement instead.

optional
Allow the operating system to commit the configuration even if the file is not available at the **/var/db/scripts/jet** file path.

respawn-on-normal-exit
Automatically restart a daemonized JET application written in Python after you exit the application normally, such as when you upgrade the JET application. If this option is not

configured, and if the application normally exits, it will not restart automatically. This option can only be used with Python JET applications that have been configured to run as a daemonized process using the `daemonize` option.

`username` *`username`*

Specify the name of the user under whose privileges the extension service will execute. This user name is configured at the `[edit system login]` hierarchy level. If you do not associate a user name with an extension-service application, the application is executed as user `nobody`.

- **Default:** nobody

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`refresh`, `refresh-from`, `respawn-on-normal-exit`, `routing-instance`, and `source` options added in Junos OS Release 18.1R1.

`interpreter` option added in Junos OS Evolved Release 20.3R1.

`optional` option added in Junos OS and Junos OS Evolved Release 23.2R1.

file list

IN THIS SECTION

- Syntax | 7063
- Description | 7063
- Options | 7063
- Additional Information | 7063
- Required Privilege Level | 7063
- Sample Output | 7064

Release Information | 7065

Syntax

```
file list <detail | recursive> <path>
```

Description

Display a list of files on the local router or switch.

Options

none Display a list of files in the default directory. The default directory is the home directory of the user logged in to the router or switch.

detail (Optional) Display detailed information about the files. The output is similar to what is displayed by the Linux `ls -l` command.

recursive (Optional) Display detailed information about the files in the directory and all subdirectories below it.

path (Optional) List the files in a specified directory path. The path name cannot contain any special characters (`![ =;|(){}]`).

Additional Information

To view available directories, enter a space and then a slash (/) after the `file list` command. To view files within a specific directory, include a slash followed by the directory and, optionally, subdirectory name after the `file list` command.

Required Privilege Level

maintenance

Sample Output

file list

The following command lists the contents of the **/var/tmp** directory.

```
user@host> file list /var/tmp

/var/tmp:
trace_debug
package.log
pics/
downloads/
```

file list (detailed)

The following command lists detailed information about the contents of the **/var/tmp** directory.

```
user@host> file list /var/tmp detail

/var/tmp/:
total blocks: 4276224
-rw-r--r--  1 user  group       1362 Oct 16 11:11 trace_debug
-rw-r--r--  1 user  group        108 Aug 9   2016 package.log
drwxrwxrwx  2 user  group        512 Jun 30  2016 pics/
drwxr-xr-x  3 user  group        512 Aug 9   2016 downloads/
total files: 2
```

file list (recursive)

The following command lists detailed information about the contents of the **/var/tmp** directory and all subdirectories below it.

```
user@host> file list /var/tmp recursive

/var/tmp/:
total blocks: 4276224
-rw-r--r--  1 user  group       1362 Oct 16 11:11 trace_debug
-rw-r--r--  1 user  group        108 Aug 9   2016 package.log
```

```
drwxrwxrwx  2 user  group        512 Jun 30  2016 pics/
drwxr-xr-x  3 user  group        512 Aug 9   2016 downloads/
total files: 2

/var/tmp/pics:
total blocks: 5120461
-rw-r--r--  1 user  group       1910 Oct 15  2016 image3.png
-rw-r--r--  1 user  group       1852 Oct 15  2016 image2.png
-rw-r--r--  1 user  group       1310 Aug 9   2016 image1.png
total files: 3

/var/tmp/downloads:
total blocks: 24
-rw-r--r--  1 user  group        108 Aug 21  2016 package2.log
-rw-r--r--  1 user  group        108 Aug 9   2016 package1.log
drwxr-xr-x  2 user  group        512 Aug 9   2016 sub-download/
total files: 2

/var/tmp/downloads/sub-download:
total blocks: 16
total files: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

file (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | 7066
- Hierarchy Level | 7066
- Description | 7066
- Options | 7066
- Required Privilege Level | 7066
- Release Information | 7067

Syntax

```
file <filename> <files files> <match match> <size size> <(world-readable | no-world-readable)>;
```

Hierarchy Level

```
[edit services rtlog traceoptions]
```

Description

Trace file information

Options

filename Name of file in which to write trace information

files Maximum number of trace files

- **Default:** 3
- **Range:** 2 through 1000

match Regular expression for lines to be logged

no-world-readable Don't allow any user to read the log file

size Maximum trace file size

- **Default:** 128k
- **Range:** through

world-readable Allow any user to read the log file

All other options are explained separately.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

file (Op Scripts)

IN THIS SECTION

- Syntax | 7067
- Hierarchy Level | 7068
- Description | 7068
- Options | 7068
- Required Privilege Level | 7069
- Release Information | 7069

Syntax

```
file filename {
    allow-commands "regular-expression";
    arguments {
        argument-name {
            description descriptive-text;
        }
    }
    cert-file path;
    checksum hash-algorithm hash-value;
```

```
    command filename-alias;
    dampen {
        dampen-options {
            cpu-factor cpu-factor;
            line-interval line-interval;
            time-interval time-interval;
        }
    }
    description descriptive-text;
    refresh;
    refresh-from url;
    routing-instance routing-instance-name;
    source url;
    start {
        start-options {
            mem-factor percentage;
        }
    }
}
```

Hierarchy Level

```
[edit system scripts op]
```

Description

Enable a Junos OS op script that is located in the **/var/db/scripts/op** directory on the hard disk or the **/config/scripts/op** directory on the flash drive. After enabling an op script, you can execute the script from the CLI using the `op` command.

Options

filename Name of a Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative syntaX (SLAX) file containing an op script.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Storing and Enabling Scripts

Enabling an Op Script and Defining a Script Alias

file (Packet Forwarding Engine Traffic Traceoptions)

IN THIS SECTION

- Syntax | 7069
- Hierarchy Level | 7070
- Description | 7070
- Options | 7070
- Required Privilege Level | 7070
- Release Information | 7070

Syntax

```
file < filename> <files files> <match match> <size size> <(world-readable | no-world-readable)>;
```

Hierarchy Level

```
[edit services pfe traffic traceoptions]
```

Description

Configure the file settings that contain trace logging information.

Options

filename Name of the file containing the trace information.

files Maximum number of trace files.

- **Default:** 3
- **Range:** 2 through 1000

match Regular expression for lines to be logged.

no-world-readable Restrict all users to read the log files.

size Maximum trace file size.

- **Default:** 128k

world-readable Allow all users to read the log files.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

file (QFabric System)

IN THIS SECTION

- Syntax | 7071
- Hierarchy Level | 7071
- Description | 7071
- Options | 7072
- Required Privilege Level | 7072
- Release Information | 7072

Syntax

```
file filename {
    archive {
        size maximum-file-size;
    }
    explicit-priority;
    facility severity;
    match "regular-expression";
    structured-data {
        brief;
    }
}
```

Hierarchy Level

```
[edit system syslog]
```

Description

Configure the logging of system messages to a file.

Options

facility Class of messages to log. To specify multiple classes, include multiple `facility severity` statements.

filename Filename that you specify with the `show log` command.

- **Default:** Filename **messages**

severity Severity of the messages that belong to the facility specified by the paired *facility* name. Messages with severities at the specified level and higher are logged.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

syslog (QFabric System) | 28825

file rename

IN THIS SECTION

- Syntax | 7073
- Description | 7073
- Options | 7073
- Required Privilege Level | 7073
- Output Fields | 7073

Sample Output | 7073
Release Information | 7075

Syntax

```
file rename source destination
```

Description

Rename a file on the local router or switch.

Options

destination New name for the file.

source Original name of the file. For a routing matrix, the filename must include the chassis information.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

file rename

The following example lists the files in **/var/tmp**, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list /var/tmp
dcd.core
```

```
rpd.core
snmpd.core
```

```
user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
```

```
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

file rename (Routing Matrix)

The following example lists the files in **/var/tmp**, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
--------------------------------------------------------------------------

/var/tmp:
.pccardd
sartre.conf
snmpd
syslogd.core-tarball.0.tgz

user@host> file rename lcc0-re0:/var/tmp/snmpd /var/tmp/snmpd.rr
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
--------------------------------------------------------------------------

/var/tmp:
.pccardd
sartre.conf
snmpd.rr
syslogd.core-tarball.0.tgz
```

Release Information

Command introduced before Junos OS Release 7.4.

file (Sampling)

IN THIS SECTION

- Syntax | 7075
- Hierarchy Level | 7075
- Description | 7075
- Required Privilege Level | 7076
- Release Information | 7076

Syntax

```
file {
    disable;
    filename filename;
    files number;
    size bytes;
    (stamp | no-stamp);
    (world-readable | no-world-readable);
}
```

Hierarchy Level

```
[edit forwarding-options sampling family inet output]
```

Description

Collect the traffic samples in a file.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

Collecting Traffic Sampling Output in a File

file (Security Log)

IN THIS SECTION

- Syntax | 7076
- Hierarchy Level | 7077
- Description | 7077
- Options | 7077
- Required Privilege Level | 7077
- Release Information | 7077

Syntax

```
file {
    files max-file-number;
    name file-name;
    path binary-log-file-path;
```

```
    size maximum-file-size;
}
```

Hierarchy Level

```
[edit security log]
[edit logical-systems name security log]
[edit tenants tenant-name security log]
[edit security log mode event]
[edit security log mode stream]
[edit security log mode stream-event]
```

Description

Configure security log file options for logs in binary format.

Options

- `files` *`number`*—Specify the maximum number of binary log files.
- **Range:** 2 through 10 files.
- `name` *`name`* —Name of the file to log messages.
- `path` *`filepath`*—Specify the path of the binary log file.
- `size` *`maximum-file-size`*—Maximum size of each trace file, in megabytes (MB).
- **Range:** 1 KB through 10 MB

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

The [edit `logical-systems` *name* security log] and [edit `tenants` *tenant-name* security log] hierarchy levels introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

syslog (System) | 28836

file show

IN THIS SECTION

- Syntax | 7078
- Description | 7078
- Options | 7078
- Required Privilege Level | 7079
- Output Fields | 7079
- Sample Output | 7079
- Release Information | 7080

Syntax

```
file show filename
<encoding (base64 | raw)>
```

Description

Display the contents of a file.

Options

filename	Name of a file. For a routing matrix, the file name must include the chassis information.
encoding (base64 | raw)	(Optional) Encode file contents with base64 encoding or show raw text.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

file show

```
user@host> file show /var/log/messages
Apr 13 21:00:08 dev1 /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 dev1 /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 dev1 last message repeated 4 times
Apr 13 21:07:04 dev1 last message repeated 8 times
Apr 13 21:07:13 dev1 /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 dev1 /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
...
```

file show (Routing Matrix)

```
user@host> file show lcc0-re0:/var/tmp/.gdbinit
lcc0-re0:
--------------------------------------------------------------------------
####################################################################
# Settings
####################################################################

set print pretty

####################################################################
# Basic stuff
####################################################################

define msgbuf
        printf "%s", msgbufp->msg_ptr
end
# hex dump of a block of memory
```

```
# usage: dump address length
define dump
   p $arg0, $arg1
   set $ch = $arg0
   set $j = 0
   set $n = $arg1
   while ($j < $n)
        #printf "%x %x ",&$ch[$j],$ch[$j]
        printf "%x ",$ch[$j]
        set $j = $j + 1
        if (!($j % 16))
                printf "\n"
        end
   end
end
```

Release Information

Command introduced before Junos OS Release 7.4.

file (SNMP Scripts)

IN THIS SECTION

- Syntax | 7081
- Hierarchy Level | 7081
- Description | 7081
- Options | 7081
- Required Privilege Level | 7082
- Release Information | 7082

Syntax

```
file filename {
    cert-file path;
    checksum hash-algorithm hash-value;
    oid id {
        priority;
    }
    python-script-user username;
    refresh;
    refresh-from;
    routing-instance routing-instance-name;
    source;
    start {
        start-options {
            mem-factor percentage;
        }
    }
}
```

Hierarchy Level

```
[edit system scripts snmp]
```

Description

For Junos OS SNMP scripts, enable an SNMP script that is located in the **/var/db/scripts/snmp** directory. If the `load-scripts-from-flash` option is configured, the system reads SNMP scripts from the **/config/scripts/snmp** directory instead.

Options

filename Name of an Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative syntaX (SLAX) file containing an SNMP script.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

snmp (SNMP Scripts) | 28068

file-specification (File Format)

IN THIS SECTION

- Syntax | 7082
- Hierarchy Level | 7083
- Description | 7083
- Required Privilege Level | 7083
- Release Information | 7083

Syntax

```
file-specification {
    variant variant-number {
        data-format format;
        name-format format;
        transfer {
            record-level number;
            timeout seconds;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure the file format for the flow collection files.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring File Formats

file-specification (Interface Mapping)

IN THIS SECTION

- Syntax | 7084
- Hierarchy Level | 7084
- Description | 7084

- Options | 7084
- Required Privilege Level | 7084
- Release Information | 7084

Syntax

```
file-specification {
    variant variant-number;
}
```

Hierarchy Level

```
[edit services flow-collector interface-map]
```

Description

Configure the default file specification for interface mapping.

Options

variant-number—Default file format variant.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

file (System Logging)

IN THIS SECTION

- Syntax | 7085
- Hierarchy Level | 7085
- Description | 7086
- Options | 7086
- Required Privilege Level | 7086
- Release Information | 7086

Syntax

```
file filename {
    facility severity;
    archive {
        files number;
        size size;
        (no-world-readable | world-readable);
    }
    explicit-priority;
    match "regular-expression";
    match-strings string-name;
    structured-data {
        brief;
    }
}
```

Hierarchy Level

```
[edit system syslog]
```

Description

Configure the logging of system messages to a file.

Options

facility—Class of messages to log. To specify multiple classes, include multiple ***facility severity*** statements. For a list of the facilities, see *Junos OS System Logging Facilities*.

file ***filename***—File in the **/var/log** directory in which to log messages from the specified facility. To log messages to more than one file, include more than one **file** statement.

severity—Severity of the messages that belong to the facility specified by the paired ***facility*** name. Messages with severities of the specified level and higher are logged. For a list of the severities, see *System Log Message Severity Levels*.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Directing System Log Messages to a Log File

file (System Logging)

IN THIS SECTION

- Syntax | **7087**
- Hierarchy Level | **7087**

- Description | 7087
- Options | 7088
- Required Privilege Level | 7089
- Release Information | 7089

Syntax

```
file name {
    allow-duplicates;
    archive name password password routing-instance routing-instance <(binary-data | no-binary-
data)> <files files> <size bytes> <start-time start-time> <transfer-interval minutes> <(world-
readable | no-world-readable)>;
    contents (any | authorization | change-log | conflict-log | daemon | dfc | external |
firewall | ftp | interactive-commands | kernel | local0 | lpr | mail | news | ntp | pfe |
privileged | security | syslog | user | uucp) {
    }
    explicit-priority;
    match match;
    match-strings [match-strings ...];
    structured-data (brief | detail);
}
```

Hierarchy Level

```
[edit logical-systems name system syslog file],
[edit logical-systems name system syslog host],
[edit logical-systems name system syslog user],
[edit system syslog file],
[edit system syslog host],
[edit system syslog user]
```

Description

Specify the file in which to log data. Starting in Junos OS Release 20.3R1, the `change-log` is a default option at [`edit system syslog file name`] hierarchy for SRX Series Firewalls. As the default option, `change-log`

records all the configuration changes. In Junos OS releases earlier than 20.2R1, you need to configure `change-log`.

Options

- *filename*—Specify the name of the file in which to log data.
- *allow-duplicates*—Do not suppress the repeated messages.
- *any*—Specify all facilities information.
  - *alert*—Specify the conditions that should be corrected immediately.
  - *critical*—Specify the critical conditions.
  - *emergency*—Specify the conditions that cause security functions to stop.
  - *error*—Specify the general error conditions.
  - *info*—Specify the information about normal security operations.
  - *none*—Do not specify any messages.
  - *notice*—Specify the conditions that should be handled specifically.
  - *warning*—Specify the general warning conditions.
- *archive*—Specify the archive file information.
  - *archive-sites*—Specify a list of destination URLs for the archived log files.
    - *url*—Specify the primary and failover URLs to receive archive files.
  - *binary-data*—Mark file such that it contains binary data.
  - *no-binary-data*—Do not mark the file such that it contains binary data.
  - *files*—Specify the number of files to be archived. Range: 1 through 1000 files.
  - *size*—Specify the size of files to be archived. Range: 65,536 through 1,073,741,824 bytes.
  - *world-readable*—Allow any user to read the log file.
  - *no-world-readable*—Do not allow any user to read the log file.
  - *start-time*—Specify the start time for file transmission. Enter the start time in the yyyy-mm-dd.hh:mm format.
  - *transfer-interval*—Specify the frequency at which to transfer the files to archive sites.

- *authorization*—Specify the authorization system.
- *change-log*—Specify the configuration change log.
- *conflict-log*—Specify the configuration conflict log.
- *daemon*—Specify the various system processes.
- *dfc*—Specify the dynamic flow capture.
- *explicit-priority*—Include the priority and facility in messages.
- *external*—Specify the local external applications.
- *firewall*—Specify the firewall filtering system.
- *ftp*—Specify the FTP process.
- *interactive-commands*—Specify the commands executed by the UI.
- *kernel*—Specify the kernel information.
- *match*—Specify the regular expression for lines to be logged.
- *ntp*—Specify the NTP process.
- *pfe*—Specify the Packet Forwarding Engine.
- *security*—Specify the security-related information.
- *structured-data*—Log the messages in structured log format.
  - *brief*—Omit English language text from the end of the logged message.
- *user*—Specify the user processes.
  - *info*—Specify the informational messages.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 12.1X47.

file (Trace Options)

IN THIS SECTION

- Syntax | 7090
- Hierarchy Level | 7090
- Description | 7090
- Options | 7090
- Required Privilege Level | 7091
- Release Information | 7091

Syntax

```
file filename {
     files number;
     size bytes;
     world-readable | no-world-readable>;
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions],
[edit forwarding-options sampling traceoptions]
```

Description

Configure information about the files that contain trace logging information.

Options

filename—Name of the file containing the trace information.

- **Default:** `/var/log/sampled`

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing Traffic Sampling Operations

Tracing Traffic-Sampling Operations

file (Tracing for Origin AS Validation)

IN THIS SECTION

- Syntax | 7091
- Hierarchy Level | 7092
- Description | 7092
- Options | 7092
- Required Privilege Level | 7093
- Release Information | 7093

Syntax

```
file filename <files number> <size size> <world-readable | no-world-readable>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation traceoptions],
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation group group-name session address traceoptions],
[edit logical-systems logical-system-name routing-options validation traceoptions],
[edit logical-systems logical-system-name routing-options validation group group-name session
address traceoptions],
[edit routing-instances instance-name routing-options validation traceoptions],
[edit routing-instances instance-name routing-options validation group group-name session
address traceoptions],
[edit routing-options validation traceoptions]
[edit routing-options validation group group-name session address traceoptions]
```

Description

Configure the file settings for tracing resource public key infrastructure (RPKI) BGP route validation.

Options

filename Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

files *number* (Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached (*x*k to specify KB, *x*m to specify MB, or *x*g to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 3 files

no-world-readable (Optional) Restrict file access to the user who created the file.

size *size* (Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `files` option.

- **Syntax:** *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

world-readable (Optional) Enable unrestricted file access.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

file-type

IN THIS SECTION

- Syntax | 7094
- Hierarchy Level | 7094
- Description | 7094
- Options | 7094
- Required Privilege Level | 7094
- Release Information | 7094

Syntax

```
file-type {
    values [ values ];
}
```

Hierarchy Level

```
[edit security idp (Security)              dynamic-attack-group              name
filters]
```

Description

File type the attack is valid for.

Options

values Values for file-type field.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 18.2R1.

file-serialization

IN THIS SECTION

- Syntax | 7095
- Hierarchy Level | 7095

- Description | 7095
- Required Privilege Level | 7096
- Release Information | 7096

Syntax

```
(file-serialization | no-file-serialization);
```

Hierarchy Level

```
[edit security policies],
[edit services policies]
```

Description

Perform policy changes propagation to the dataplane using file serialization. By serializing policy configurations into files, the system ensure that they are read and applied by the PFE in a controlled and reliable manner.

By default, the file-based serialization feature is enabled. You can disable the file-serialization by using the following statement:

```
[edit]
user@host# set security policies no-file-serialization
```

To re-enable the file-serialization feature, use the following statement:

```
[edit]
user@host# delete security policies no-file-serialization
```

Or use the following statement:

```
[edit]
user@host# set security policies file-serialization
```

Required Privilege Level

security—To view this statement in the configuration

security-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 24.4R1

filename

IN THIS SECTION

- Syntax | 7096
- Hierarchy Level | 7097
- Description | 7097
- Options | 7097
- Required Privilege Level | 7097
- Release Information | 7097

Syntax

```
filename filename;
```

Hierarchy Level

```
[edit forwarding-options sampling family (inet |inet6 |mpls) output file]
```

Description

Configure the name of the output file.

Options

filename—Name of the file in which to place the traffic samples. All files are placed in the directory **/var/tmp**.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

filename-extension

IN THIS SECTION

- Syntax | 7098
- Hierarchy Level | 7098
- Description | 7098

- Options | 7098
- Required Privilege Level | 7098
- Release Information | 7099

Syntax

```
filename-extension object-name {
    value [value];
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm custom-objects]
```

Description

When scanning content, you can use a file extension list to define a set of file extensions that are used in file extension scan mode (scan-by-extension).

Options

- *object-name*—Name of the extension-list object.
- `value` *value*—Value of the extension-list object. You can configure multiple values separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

filename (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | 7099
- Hierarchy Level | 7099
- Description | 7099
- Options | 7099
- Required Privilege Level | 7100
- Release Information | 7100

Syntax

```
filename;
```

Hierarchy Level

```
[edit services rtlog traceoptions file]
```

Description

Name of file in which to write trace information

Options

filename Name of file in which to write trace information

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

filename-prefix

IN THIS SECTION

- Syntax | 7100
- Hierarchy Level | 7101
- Description | 7101
- Options | 7101
- Required Privilege Level | 7101
- Release Information | 7101

Syntax

```
filename-prefix prefix;
```

Hierarchy Level

```
[edit services flow-collector transfer-log-archive]
```

Description

Configure the filename prefix for log files.

Options

prefix—Filename identifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Transfer Logs

filename (Sampling)

IN THIS SECTION

- Syntax | 7102
- Hierarchy Level | 7102
- Description | 7102
- Options | 7102

Required Privilege Level | 7102
Release Information | 7102

Syntax

```
filename filename;
```

Hierarchy Level

```
[edit forwarding-options sampling family family-name output file]
```

Description

Configure the name of the output file.

Options

filename—Name of the file in which to place the traffic samples. All files are placed in the directory **/var/tmp**.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in a File

files

IN THIS SECTION

- Syntax | 7103
- Hierarchy Level | 7103
- Description | 7103
- Options | 7103
- Required Privilege Level | 7104
- Release Information | 7104

Syntax

```
files number;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family (inet |inet6 |mpls) output file],
[edit forwarding-options sampling traceoptions file]
```

Description

Configure the total number of files to be saved with samples or trace data.

Options

number—Maximum number of traffic sampling or trace log files. When a file named *sampling-file* reaches its maximum size, it is renamed *sampling-file*.0, then *sampling-file*.1, and so on, until the maximum number of traffic sampling files is reached. Then the oldest sampling file is overwritten.

- **Range:** 1 through 100 files
- **Default:** 5 files for sampling output; 10 files for trace log information

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

Configuring Traffic Sampling on MX, M and T Series Routers

files (accounting)

IN THIS SECTION

- Syntax | 7104
- Hierarchy Level | 7105
- Description | 7105
- Options | 7105
- Required Privilege Level | 7105
- Release Information | 7105

Syntax

```
files number;
```

Hierarchy Level

```
[edit accounting-options file filename]
```

Description

Specify the maximum number of log files to be used for accounting data.

Options

number—The maximum number of files. When a log file (for example, `profilelog`) reaches its maximum size, it is renamed `profilelog.0`, then `profilelog.1`, and so on, until the maximum number of log files is reached. Then the oldest log file is overwritten. The minimum value for *number* is 3 and the default value is 10.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Accounting-Data Log Files

files (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | 7106
- Hierarchy Level | 7106
- Description | 7106

- Options | 7106
- Required Privilege Level | 7106
- Release Information | 7106

Syntax

```
files files;
```

Hierarchy Level

```
[edit services rtlog traceoptions file filename]
```

Description

Maximum number of trace files

Options

files Maximum number of trace files

- **Default:** 3
- **Range:** 2 through 1000

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

files (Sampling and Traceoptions)

IN THIS SECTION

- Syntax | 7107
- Hierarchy Level | 7107
- Description | 7107
- Options | 7108
- Required Privilege Level | 7108
- Release Information | 7108

Syntax

```
files number;
```

Hierarchy Level

```
[edit forwarding-options helperstraceoptions file],
[edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family family-name output file],
[edit forwarding-options sampling traceoptions file]
```

Description

Configure the total number of files to be saved with samples or trace data.

Options

number—Maximum number of traffic sampling or trace log files. When a file named ***sampling-file*** reaches its maximum size, it is renamed ***sampling-file*.0**, then ***sampling-file*.1**, and so on, until the maximum number of traffic sampling files is reached. Then the oldest sampling file is overwritten.

- **Range:** 1 through 100 files
- **Default:** 5 files for sampling output; 10 files for trace log information

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in a File

Tracing Traffic-Sampling Operations

files (syslog)

IN THIS SECTION

- Syntax | 7109
- Hierarchy Level | 7109
- Description | 7109
- Options | 7109
- Required Privilege Level | 7109
- Release Information | 7109

Syntax

```
files number;
```

Hierarchy Level

```
[edit system syslog archive],
[edit system syslog file filename archive]
```

Description

Configure the maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file ***logfile***, it closes the file, compresses it, and renames it to ***logfile*.0.gz** (for information about the maximum file size, see "size" on page 28007). The utility then opens and writes to a new file called ***logfile***. When the new file reaches the maximum size, the ***logfile*.0.gz** file is renamed to ***logfile*.1.gz**, and the new file is closed, compressed, and renamed ***logfile*.0.gz**. By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).

Options

number—Maximum number of archived files.

- **Range:** 1 through 1000
- **Default:** 10 files

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

size (System) | 28007

fill-level

IN THIS SECTION

- QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | 7110
- QFX10000 Switches | 7110
- Hierarchy Level | 7110
- Description | 7111
- Options | 7111
- Required Privilege Level | 7111
- Release Information | 7112

QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems

```
fill-level low-value fill-level high-value;
```

QFX10000 Switches

```
fill-level level1 level2 ... level32;
```

Hierarchy Level

```
[edit class-of-service drop-profiles profile-name interpolate]
```

Description

When configuring weighted random early detection (WRED), map the fullness of a queue to a packet "drop-probability" on page 5953 value. You configure the `fill-level` and `drop-probability` statements in related pairs. The pairs of fill level and drop probability values set a probability of dropping packets at a specified queue fullness value.

The first fill level is the packet drop start point. Packets do not drop until the queue fullness reaches the first fill level. The last fill level is the packet drop end point. After the queue exceeds the fullness set by the drop end point, all non-ECN packets are dropped. As the queue fills from the drop start point to the drop end point, the rate of packet drop increases in a curve pattern. The higher the queue fill level, the higher the probability of dropping packets.

On switches that support only two fill level/drop probability pairs, the two pairs are the drop start point and the drop end point. On switches that support up to 32 fill level/drop probability pairs, you can configure intermediate interpolations between the drop start point and the drop end point, which provides greater flexibility in controlling the packet drop curve.

NOTE: Do not configure the last fill level as 100 percent.

Options

low-value (switches that support only two fill level/drop probability pairs)—Fullness of the queue before packets begin to drop, expressed as a percentage. The low value must be less than the high value.

- **Range:** 0 through 100

high-value (switches that support only two fill level/drop probability pairs)—Fullness of the queue before it reaches the maximum drop probability. If the queue fills beyond the fill level high value, all packets drop. The high value must be greater than the low value.

- **Range:** 0 through 100

level1 level2 ... level32 (switches that support 32 fill level/drop probability pairs)—The queue fullness level, expressed as a percentage. Each fill level pairs with a drop probability to define the probability of a packet dropping at a specified queue fullness.

- **Range:** 0 through 100

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

fill-level (Drop Profiles)

IN THIS SECTION

- Syntax | 7112
- Hierarchy Level | 7112
- Description | 7112
- Options | 7112
- Required Privilege Level | 7113
- Release Information | 7113

Syntax

```
fill-level percentage;
```

Hierarchy Level

```
[edit class-of-service drop-profiles profile-name]
```

Description

When configuring RED, map the fullness of a queue to a drop probability.

Options

percentage How full the queue is, expressed as a percentage. You configure the `fill-level` and `drop-probability` statements in pairs. To specify multiple fill levels, include multiple `fill-level` and

`drop-probability` statements. The values you assign to each statement pair must increase relative to the previous pair's values. This is shown in the discrete graph in *Managing Congestion Using RED Drop Profiles and Packet Loss Priorities*.

- **Range:** 0 through 100 percent

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Managing Congestion Using RED Drop Profiles and Packet Loss Priorities

Defining Packet Drop Behavior by Configuring RED Drop Profiles

fill-level (Interpolated Value)

IN THIS SECTION

- Syntax | 7114
- Hierarchy Level | 7114
- Description | 7114
- Options | 7114
- Required Privilege Level | 7114
- Release Information | 7114

Syntax

```
fill-level [values];
```

Hierarchy Level

```
[edit class-of-service drop-profiles profile-name interpolate]
```

Description

Define up to 64 values for interpolating queue fill level.

Options

values—Data points for mapping queue fill percentage.

- **Range:** 0 through 100
- **Default:** In the default tail drop profile, when the fill level is 0 percent, the drop probability is 0 percent. When the fill level is 100 percent, the drop probability is 100 percent.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Managing Congestion Using RED Drop Profiles and Packet Loss Priorities*
>
> *Defining Packet Drop Behavior by Configuring RED Drop Profiles*

filter (Flow Monitoring)

IN THIS SECTION

- Syntax | 7115
- Hierarchy Level | 7115
- Description | 7115
- Options | 7115
- Required Privilege Level | 7116
- Release Information | 7116

Syntax

```
filter {
    input filter-name;
    output filter-name;
    group filter-group-number;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

Apply a firewall filter to an interface. You can also use filters for encrypted traffic.

Options

`group` *`filter-group-number`*—Use the specified interface to be part of a filter group. The default filter group number is 0.

`input` *`filter-name`*—Use the specified filter to evaluate when packets are received on the interface.

`output` *`filter-name`*—Use the specified filter to evaluate when packets are transmitted on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Flow Monitoring

filter (interfaces)

IN THIS SECTION

- Syntax | 7116
- Hierarchy Level | 7117
- Description | 7117
- Options | 7117
- Required Privilege Level | 7117
- Release Information | 7117

Syntax

```
filter {
    input filter-name;
    output filter-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

Define the filters to be applied on an interface.

Options

`input` *`filter-name`*—Identifier for the input filter.

`output` *`filter-name`*—Identifier for the output filter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Filters for Traffic Transiting the ES PIC

filter (Interfaces)

IN THIS SECTION

- Syntax | 7118
- Hierarchy Level | 7118
- Description | 7118

- Options | 7118
- Required Privilege Level | 7119
- Release Information | 7119

Syntax

```
filter {
    group filter-group-number;
    input filter-name;
    input-list [ filter-names ];
    output filter-name;
    output-list [ filter-names ];
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family]
```

Description

NOTE: On EX Series switches, the `group`, `input-list`, `output-filter` statements are not supported under the [`edit interfaces` *`interface-name`* `unit` *`logical-unit-number`* `family inet`], [`edit interfaces` *`interface-name`* `unit` *`logical-unit-number`* `family inet6`], and [`edit interfaces` *`interface-name`* `unit` *`logical-unit-number`* `family mpls`] hierarchies.

Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure them under the **family ethernet-switching**, **inet**, **inet6**, **mpls**, or **vpls** only.

Options

group *filter-group-number*—Define an interface to be part of a filter group. The default filter group number is 0.

- **Range:** 0 through 255

input ***filter-name***—Name of one filter to evaluate when packets are received on the interface.

output ***filter-name***—Name of one filter to evaluate when packets are transmitted on the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Applying a Filter to an Interface
- Junos OS Administration Library for Routing Devices
- Configuring Gigabit Ethernet Interfaces (CLI Procedure)
- *Configuring Firewall Filters (CLI Procedure)*
- family

filter

IN THIS SECTION

- Syntax | **7120**
- Hierarchy Level | **7120**
- Description | **7120**
- Options | **7120**
- Required Privilege Level | **7121**
- Release Information | **7121**

Syntax

```
filter filter-name {
    interface-specific
    micro-segmentation;
    no-longest-prefix-match;
    term term-name {
        from {
            match-conditions;
        }
        then {
            action;
            action-modifiers;
        }
    }
}
```

Hierarchy Level

```
[edit firewall family family-name]
```

Description

Configure firewall filters.

Options

filter-name—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.

micro-segmentation—Indicates that this firewall filter is a Group Based Policy (GBP) tagging filter. A GBP tagging filter uses firewall terms to assign GBP tags for use with macro and micro segmentation on VXLAN. GBPs make use of existing layer 3 VXLAN network identifiers (VNI), in conjunction with firewall filter policies, to provide micro-segmentation at the level of device or tag, independent of the underlying network topology.

These match conditions are supported for GBP tagging:

- `ip-version ipv4` *<ip address>* | *<prefix-list>*

- `ip-version ipv6` *<ip address>* | *<prefix-list>*
- `mac-address` *<mac address>*
- `interface` *<interface_name>* vlan-id *<vlan id>*
- `vlan-id` *<vlan id>*
- `interface` *<interface_name>*

no-longest-prefix-match—Disables longest prefix match in IP-based GBP tagging filters. This causes IP-based firewall terms in GBP tagging filters to be evaluated in term order instead of by longest prefix match. This parameter is only applicable when **micro-segmentation** is enabled and is ignored otherwise. Set this parameter when you first create the filter. Do not toggle this parameter on an existing filter.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Option **micro-segmentation** introduced in Junos OS Release 22.4R1 to support GBP tagging filters.

Option **no-longest-prefix-match** introduced in Junos OS Release 24.4R1 to disable longest prefix match on IP-based GBP tagging filters.

RELATED DOCUMENTATION

Firewall Filter Match Conditions and Actions (QFX5100, QFX5110, QFX5120, QFX5200, EX4600, EX4650)

Configuring Firewall Filters

Overview of Firewall Filters (QFX Series)

Example: Micro and Macro Segmentation using Group Based Policy in a VXLAN

filter (routing-options)

IN THIS SECTION

- Syntax | 7122
- Hierarchy Level | 7122
- Description | 7122
- Options | 7122
- Required Privilege Level | 7123
- Release Information | 7123

Syntax

```
filter {
    input filter-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name],
[edit logical-systems logical-system-name routing-options rib routing-table-name],
[edit routing-instances routing-instance-name routing-options rib routing-table-name],
[edit routing-options rib routing-table-name]
```

Description

Specify the name of the routing table from which Junos OS should export routing information.

Options

input *filter-name*—Forwarding table filter name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Forwarding Table Filters

Applying Forwarding Table Filters

filter (Applying to a Logical Interface)

IN THIS SECTION

- Syntax | 7123
- Hierarchy Level | 7124
- Description | 7124
- Options | 7124
- Required Privilege Level | 7124
- Release Information | 7125

Syntax

```
filter {
    group filter-group-number;
    input filter-name;
    input-list [ filter-names ];
    output filter-name;
```

```
    output-list [ filter-names ];
}
```

Hierarchy Level

Protocol-independent firewall filter on MX Series router logical interface:

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

All other standard firewall filters on all other devices:

```
[edit interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family]
```

Description

Apply a stateless firewall filter to a logical interface at a specific protocol level.

Options

group *filter-group-number*	(Only Ex, M, MX, and T Series) Number of the group to which the interface belongs. Range: 1 through 255
input *filter-name*	Name of one filter to evaluate when packets are received on the interface.
input-list [*filter-names* **]**	Names of filters to evaluate when packets are received on the interface. Up to 16 filters can be included in a filter input list.
output *filter-name*	Name of one filter to evaluate when packets are transmitted on the interface.
output-list [*filter-names* **]**	Names of filters to evaluate when packets are transmitted on the interface. Up to 16 filters can be included in a filter output list.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Guidelines for Configuring Firewall Filters
- Guidelines for Applying Standard Firewall Filters

filter (Applying to an Interface)

IN THIS SECTION

- Syntax | 7125
- Hierarchy Level | 7126
- Description | 7126
- Options | 7126
- Required Privilege Level | 7126
- Release Information | 7126

Syntax

```
filter {
    input filter-name;
    output filter-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family]
```

Description

Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure the family `inet`, `inet6`, `mpls`, or `vpls` only.

> **NOTE**: MPLS firewall filters applied on output interface are not supported on PTX10003 router due to product limitation.

Options

`input` *`filter-name`*—Name of one filter to evaluate when packets are received on the interface.

`output` *`filter-name`*—Name of one filter to evaluate when packets are transmitted on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

simple-filter (Applying to an Interface) | **27964**

Configuring and Applying Tricolor Marking Policers

Example: Classifying Packets Based on Their Destination Address

Example: Configuring and Verifying a Complex Multifield Filter

Example: Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets

Configuring a Simple Filter

Configuring Policers Based on Logical Interface Bandwidth

Effect of Two-Color Policers on Shaping Rate Changes

filter (Configuring)

IN THIS SECTION

- Syntax | 7127
- Hierarchy Level | 7128
- Description | 7128
- Options | 7128
- Required Privilege Level | 7128
- Release Information | 7128

Syntax

```
filter filter-name {
    accounting-profile name;
    enhanced-mode;
    fast-lookup-filter;
    filter-list-template;
    interface-shared;
    interface-specific;
    physical-interface-filter;
    promote gre-key;
    term term-name {
        ... term configuration ...
    }
}
```

Hierarchy Level

```
[edit firewall family family-name],
[edit logical-systems logical-system-name firewall family family-name]
```

Description

Configure firewall filters.

Options

filter-name—Name that identifies the filter. This must be a non-reserved string of not more than 64 characters. To include spaces in the name, enclose it in quotation marks (“ ”). Firewall filter names are restricted from having the form __.*__ (beginning and ending with underscores) or __.* (beginning with an underscore.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

`physical-interface-filter` statement introduced in Junos OS Release 9.6.

Support for the `interface-shared` statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- Guidelines for Configuring Firewall Filters
- Guidelines for Applying Standard Firewall Filters
- Configuring Multifield Classifiers

Using Multifield Classifiers to Set Packet Loss Priority

simple-filter

filter-duplicates

IN THIS SECTION

- Syntax | 7129
- Hierarchy Level | 7129
- Description | 7129
- Required Privilege Level | 7129
- Release Information | 7130

Syntax

```
filter-duplicates;
```

Hierarchy Level

```
[edit snmp]
```

Description

Filter duplicate **Get**, **GetNext**, or **GetBulk** SNMP requests.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Filtering Duplicate SNMP Requests

Understanding the Implementation of SNMP on the QFabric System

Example: Configuring SNMP

filter (Dynamic Interface Unit)

IN THIS SECTION

- Syntax | 7130
- Hierarchy Level | 7131
- Description | 7131
- Options | 7131
- Required Privilege Level | 7131
- Release Information | 7131

Syntax

```
filter {
    input filter-name (
        shared-name filter-shared-name;
    }
    output filter-name {
        shared-name filter-shared-name;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number],
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number],
```

Description

Apply a dynamic filter to an interface, regardless of its family type.

Options

`input` *`filter-name`*—Name of one filter to evaluate when packets are received on the interface.

`output` *`filter-name`*—Name of one filter to evaluate when packets are transmitted on the interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

- Guidelines for Configuring Firewall Filters
- *Understanding Dynamic Firewall Filters*
- *Classic Filters Overview*
- *Basic Classic Filter Syntax*
- *Dynamically Attaching Statically Created Filters for Any Interface Type*

filter (Dynamic Profiles Filter Attachment)

IN THIS SECTION

- Syntax | 7132
- Hierarchy Level | 7132
- Description | 7133
- Options | 7133
- Required Privilege Level | 7133
- Release Information | 7133

Syntax

```
filter {
    adf {
        counter;
        input-precedence precedence;
        not-mandatory;
        output-precedence precedence;
        rule rule-value;
    }
    input filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
    output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family],
```

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family]
```

Description

Apply a dynamic filter to an interface. You can configure filters for `family any` ,`family inet`, or `family inet6`. The filters can be classic filters, fast update filters, or (for the `adf` statement) Ascend-Data-Filters.

Options

`input` *`filter-name`*—Name of one filter to evaluate when packets are received on the interface.

`output` *`filter-name`*—Name of one filter to evaluate when packets are transmitted on the interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the `[edit dynamic-profiles` *`profile-name`* `interfaces pp0 unit "$junos-interface-unit" family` *`family`*`]` hierarchy level introduced in Junos OS Release 10.1.

`shared-name` statement added in Junos OS Release 12.2.

RELATED DOCUMENTATION

- Junos OS Routing Policies, Firewall Filters and Traffic Policers User Guide for Routing Devices
- Firewall Filters Overview
- *Understanding Dynamic Firewall Filters*
- *Classic Filters Overview*
- *Basic Classic Filter Syntax*
- *Parameterized Filters Overview*

filter (Dynamic Profiles Filter Creation)

IN THIS SECTION

- Syntax | 7134
- Hierarchy Level | 7134
- Description | 7135
- Options | 7135
- Required Privilege Level | 7135
- Release Information | 7135

Syntax

```
filter filter-name {
    enhanced-mode-override;
    instance-shared;
    interface-shared;
    interface-specific;
    term term-name {
        from {
            match-conditions;
        }
        then {
            action;
            action-modifiers;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family]
```

Description

Create firewall filters to be applied by dynamic profile.

Options

filter-name—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (“ ”). The name can also be a predefined variable.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

filter (export routing policy action)

IN THIS SECTION

- Syntax | 7136
- Hierarchy Level | 7136
- Description | 7136
- Required Privilege Level | 7136
- Release Information | 7136

Syntax

```
slice slice-name
```

Hierarchy Level

```
[edit policy-options policy-statement policy-name term term-name then]
```

Description

Enables a firewall filter to be attached to a matched route.

Required Privilege Level

- routing—To view this statement in the configuration
- routing-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS 22.2R1.

filter (Firewall Filters)

IN THIS SECTION

- Syntax | 7137
- Hierarchy Level | 7137
- Description | 7137
- Options | 7137
- Required Privilege Level | 7138
- Release Information | 7138

Syntax

```
filter filter-name {
    interface-specific;
    micro-segmentation;
    term term-name {
        from {
            match-conditions;
        }
        then {
            action;
            action-modifiers;
        }
    }
}
```

Hierarchy Level

```
[edit firewall                family                family-name]
```

Description

Configure firewall filters.

Options

filter-name—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.

micro-segmentation—Enables Group Based Policy (GBP) tagging for use with macro and micro segmentation on VXLAN. GBPs make use of existing layer 3 VXLAN network identifiers (VNI), in conjunction with firewall filter policies, to provide micro-segmentation at the level of device or tag, independent of the underlying network topology.

These match conditions are supported for GBP tagging:

- `ip-version ipv4` *<ip address>* | *<prefix-list>*
- `ip-version ipv6` *<ip address>* | *<prefix-list>*
- `mac-address` *<mac address>*

- `interface` *<interface_name>* vlan-id *<vlan id>*
- `vlan-id` *<vlan id>*
- `interface` *<interface_name>*

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Option **interface-specific** introduced in Junos OS Release 9.5 for EX Series switches.

Option **micro-segmentation** introduced in Junos OS Release 22.4R1 for supported EX4100, EX4400, EX4650, and QFX5120 Series switches.

RELATED DOCUMENTATION

Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches

Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches

Configuring Firewall Filters (CLI Procedure)

Firewall Filters for EX Series Switches Overview

filter (Identity Management Advanced Query)

IN THIS SECTION

- Syntax | 7139
- Hierarchy Level | 7139
- Description | 7139
- Options | 7140

Required Privilege Level | 7141
Release Information | 7141

Syntax

```
filter {
    domain name;
    exclude-ip {
        address-book book-name;
        address-set address-set;
    }
    include-ip {
        address-book book-name;
        address-set address-set;
    }
}
```

Hierarchy Level

```
[edit services user-identification identity-management]
```

Description

The advanced user identity query feature enables the SRX Series Firewall to communicate with the Juniper Identity Management Service (JIMS) server to obtain user identity information for an individual user (ip-query) or a group of users (batch query).

Optionally, you can configure filters to convey to the JIMS server at a more granular level the users for whom you want information, based on their IP addresses. The filter statement gives you the flexibility to specify a range of IP addresses to be excluded from the record that the JIMS server sends in response or a range of IP addresses to be included in it. You can also constrain the query target to users in one or more specific active directory domains. Only IPv4 addresses are supported.

You can configure a filter that includes all three specifications: `include-ip`, `exclude-ip`, and `domain`.

Filters are contextual. That is, you can use a different filter configuration for different requests. If you change the filter configuration, the new filter applies to subsequent user identity requests exclusively. It has no bearing on prior query requests

Use of the JIMS allows you to provision users locally and have their authentication information made available to other sites in your network for policy enforcement and reporting.

> **WARNING**: Before you use this feature, you must disable active-directory-access and authentication-source options under the user-identification hierarchy. You cannot commit this configuration if active directory authentication or the ClearPass query and webapi functions are configured and committed.

Options

include-ip address-book *book-name* address-set *address-set-name*. Optionally, configure a filter that directs the SRX Series Firewall to issue a query to the JIMS server to include in its response record user identity information for users based on IP addresses in certain address-ranges.

The following are the two behaviors when an `include-ip` is configured:

- Batch query—An SRX Series Firewall sends a request to JIMS with the include list of IP addresses.
- IP query—If the IP address to be queried is included, then the SRX Series Firewall queries JIMS only for those IP addresses that need to be included and does not query for other IP addresses; based on the IP query, JIMS does not trigger the PC probe for the IP addresses that are not included in the IP query.

A filter can include up to twenty IP address ranges. Therefore, an address set that contains more than twenty ranges will cause the filter configuration to fail. To specify the ranges, specify the name of a predefined address set which includes them and which is included in an existing address book.

The filter for IP addresses does not support nested address sets in an address book. If an address book contains nested address sets, it is ignored.

Here is an include-ip address configuration:

```
user@host# set security address-book mybook address addr1 range-address 198.51.100.0
to 198.51.120.0
user@host# set security address-book mybook address-set myset address addr1
```

```
user@host# set service user-identification identity-management filter include-ip
address-book mybook address-set myset
```

exclude-ip address-book *book-name* address-set *address-set-name*. Optionally, configure a filter that directs the SRX Series Firewall to issue a query to the JIMS server to exclude from its response record user identity information for users based on the specified address-ranges.

The following are the two behaviors when an `exclude-ip` is configured:

- Batch query—An SRX Series Firewall sends a request to JIMS with the exclude list of IP addresses.
- IP query—If the IP address to be queried is excluded, then no request is sent from an SRX Series Firewall to JIMS.

To specify the ranges, specify the name of a predefined address set which includes them and which is included in an existing address book. The address set must not include more than twenty IP addresses, otherwise the `exclude-ip` filter will fail. Here is an `exclude-ip` address configuration similar to that of the `include-ip` filter:

```
user@host# set security address-book mybook address addr1 range-address
198.51.100.0/24 to 198.51.120.0/24
user@host# set security address-book mybook address-set myset address addr1
user@host# set service user-identification identity-management filter exclude-ip
address-book mybook address-set myset
```

Starting in Junos OS Release 18.3R1, you can include or exclude IPv6 addresses for filtering the IP addresses, in addition to IPv4 addresses.

domain One or more active directory domains of interest to the SRX Series Firewall. You can specify up to twenty domain names for the filter.

Required Privilege Level

1. services—To view this statement in the configuration.
2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Understanding Advanced Query Feature for Obtaining User Identity Information from JIMS

authentication-entry-timeout (Identity Management Advanced Query) | **2597**

batch query | **2886**

invalid-authentication-entry-timeout (Services User Identification Active Directory and ClearPass) | **9227**

primary connection (Identity Management Advanced Query) | **13131**

secondary connection (Identity Management Advanced Query) | **16409**

filter-interfaces

IN THIS SECTION

- Syntax: MX, M, T, ACX, and PTX Series Routers and SRX Firewalls | **7142**
- Syntax: QFX Series switches, QFabric, OCX1100 and EX4600 | **7143**
- Hierarchy Level | **7143**
- Description | **7143**
- Options | **7143**
- Required Privilege Level | **7143**
- Release Information | **7143**

Syntax: MX, M, T, ACX, and PTX Series Routers and SRX Firewalls

```
filter-interfaces {
    interfaces {
        all-internal-interfaces;
        interface 1;
        interface 2;
    }
}
```

Syntax: QFX Series switches, QFabric, OCX1100 and EX4600

```
filter-interfaces {
    all-internal-interfaces;
    interfaces interface
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Configure the options available for the `filter-interfaces` statement to specify the interfaces that you want to exclude from the output of SNMP `Get` and `GetNext` requests performed on interface-related MIBs.

Options

`all-internal-interfaces`—Filters out all the internal interfaces from the output of SNMP `Get` and `GetNext` requests. Internal interfaces include all interfaces belonging to juniper private routing-instances.

`interfaces` *`interfaces`*—Specify the interfaces to filter them out from the output of SNMP `Get` and `GetNext` requests. Use regular expressions to specify the interface.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Filtering Interface Information Out of SNMP Get and GetNext Output

filter (IPv4, IPv6, and MPLS)

IN THIS SECTION

- Syntax | 7144
- Hierarchy Level | 7144
- Description | 7144
- Options | 7144
- Required Privilege Level | 7144
- Release Information | 7145

Syntax

```
filter {
    input input-filter-name;
    output output-filter-name;
}
```

Hierarchy Level

```
[edit forwarding-options family (inet | inet6 | mpls)]
```

Description

Apply a forwarding table filter to a forwarding table.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Table Filters

filter (Layer 2 and Layer 3 Interfaces)

IN THIS SECTION

- Syntax | 7145
- Hierarchy Level | 7145
- Description | 7146
- Default | 7146
- Options | 7146
- Required Privilege Level | 7146
- Release Information | 7146

Syntax

```
filter (input | output) filter-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family-name]
```

Description

Apply a firewall filter to traffic transiting a port or Layer 3 interface.

Default

All incoming traffic is accepted unmodified on the port or Layer 3 interface, and all outgoing traffic is sent unmodified from the port or Layer 3 interface.

Options

filter-name—Name of a firewall filter defined at the `[edit firewall family family-name filter]` hierarchy level.

input—Apply a firewall filter to traffic entering the port or Layer 3 interface.

output—Apply a firewall filter to traffic exiting the port or Layer 3 interface.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Firewall Filters

Overview of Firewall Filters (QFX Series)

filter (Layer 3 Interfaces)

IN THIS SECTION

Syntax | 7147

- Hierarchy Level | 7147
- Description | 7147
- Default | 7147
- Options | 7147
- Required Privilege Level | 7148
- Release Information | 7148

Syntax

```
filter (input | output) filter-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family-name]
```

Description

Apply a firewall filter to traffic transiting a Layer 3 interface.

Default

All incoming traffic is accepted unmodified on the Layer 3 interface, and all outgoing traffic is sent unmodified from the Layer 3 interface.

Options

filter-name—Name of a firewall filter defined at the `[edit firewall family family-name filter]` hierarchy level.

`input`—Apply a firewall filter to traffic entering the Layer 3 interface.

`output`—Apply a firewall filter to traffic exiting the Layer 3 interface.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D20.

RELATED DOCUMENTATION

Configuring Firewall Filters

Overview of Firewall Filters (OCX Series)

filter-list-template

IN THIS SECTION

- Syntax | 7148
- Hierarchy Level | 7149
- Description | 7149
- Required Privilege Level | 7149
- Release Information | 7149

Syntax

```
filter-list-template;
```

Hierarchy Level

```
[edit firewall family (inet | inet6) filter filter-name],
[edit logical-systems logical-system-name firewall family (inet | inet6) filter filter-name]
```

Description

(MX5, MX10, MX40, and MX80 routers, and routers that use MX Series MPC line cards only) Configure all interfaces that use the same filter list to use a common template. This feature can be used to save microkernel memory and DMEM memory.

If the same filter list cannot be used on all interfaces, consider merging the filters and using the `from interface` firewall filter term to group the per-interface terms to produce a new common filter list.

> **NOTE**: If you configure both `fast-lookup-filter` and `interface-specific` statements, filter list templates are also used.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3R9.

RELATED DOCUMENTATION

- input-list
- output-list
- *Firewall Filter Match Conditions for IPv6 Traffic*

filter (memory-enhanced)

IN THIS SECTION

- Syntax | 7150
- Hierarchy Level | 7150
- Description | 7150
- Required Privilege Level | 7150
- Release Information | 7150

Syntax

```
filter;
```

Hierarchy Level

```
[edit chassis memory-enhanced]
```

Description

Enables storing of firewall filters across multiple static RAM (SRAM) segments, resulting in proper utilization of SRAM segments. This feature is useful in routers with small routing tables and large firewall filters. This statement is supported on T Series routers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 11.1.

RELATED DOCUMENTATION

Allocating More Memory for Routing Tables, Firewall Filters, and Layer 3 VPN Labels

filter-profile

IN THIS SECTION

- Syntax | 7151
- Hierarchy Level | 7151
- Description | 7152
- Options | 7152
- Required Privilege Level | 7152
- Release Information | 7152

Syntax

```
filter-profile profile-name {
    counters {
        counter-name;
    }
    file filename;
    interval minutes;
    start-time time;
}
```

Hierarchy Level

```
[edit accounting-options]
```

Description

Create a profile to filter and collect packet and byte count statistics and write them to a file in the `/var/log` directory. To apply the profile to a firewall filter, you include the `accounting-profile` statement at the [`edit firewall filter` *`filter-name`*] hierarchy level. For more information about firewall filters, see *Firewall Filters Overview*.

Options

profile-name Name of the filter profile.

start-time ***time*** (Optional) Specify the start time in *yyyy-mm-dd.hh:mm* local time format for the accounting profile. Sometimes, if the time at which the records are written to the accounting file is beyond the transfer interval window of the file, drifting of accounting records occurs. To prevent drifting and to have predictable timestamps of when accounting data is written to an accounting file, we can specify the start time option for the accounting profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Firewall Filters Overview

Configuring the Filter Profile

filter (Security)

IN THIS SECTION

- Syntax | 7153
- Hierarchy Level | 7153
- Description | 7153
- Required Privilege Level | 7153
- Release Information | 7154

Syntax

```
filter threat-attack
```

Hierarchy Level

```
[edit security log stream stream-name]
```

Description

Configure the log stream filter to transmit only threat and attack logs to the ClearPass Policy Manager (CPPM). The integrated ClearPass authentication and enforcement feature sends to the CPPM threat and attack logs detected by the SRX Series Firewall security modules. You can use these reports to inform your approach to hardening the CPPM security policy. Setting the log stream filter to threat-attack ensures that the SRX Series Firewall and the log server are not overburdened by irrelevant logs.

NOTE: Unlike for other features that use a filter for log streams, threat-attack is the only log stream filter supported for integrated ClearPass. Therefore, it is not shown here as an option.

Required Privilege Level

1. security—To view this statement in the configuration.

2. security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

filter-specific (QFX)

IN THIS SECTION

- Syntax | 7154
- Hierarchy Level | 7154
- Description | 7154
- Required Privilege Level | 7155
- Release Information | 7155

Syntax

```
filter-specific;
```

Hierarchy Level

```
[edit firewall policer policer-name]
```

Description

Configure a policer to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. If you use a filter-specific policer in multiple terms, both of the following are true:

- Traffic is policed at the aggregate rate. For example, if you create a policer that has a bandwidth limit of 100 Mbps and use the policer in two terms, the total allowed bandwidth for both terms is 100 Mbps—not 100 Mbps for each term.

- The implicit counter counts all the packets are that matched by any of the terms. For example, if you reference the same filter-specific policer in term1 and term2, and term1 matches 1000 packets and term2 matches 500 packets, the implicit counter shows 1500 matches for the policer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Two-Color and Three-Color Policers to Control Traffic Rates

Overview of Policers

filter-specific

IN THIS SECTION

- Syntax | 7155
- Hierarchy Level | 7156
- Description | 7156
- Required Privilege Level | 7156
- Release Information | 7156

Syntax

```
filter-specific;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name],
[edit firewall family inet prefix-action name],
[edit firewall policer policer-name],
[edit logical-systems logical-system-name firewall policer policer-name],
[edit logical-systems logical-system-name firewall family inet prefix-action name]
```

Description

By default, a policer operates in *term-specific* mode, which means that for a given *firewall filter*the Junos OS creates a separate policer instance for every filter term that references the policer. You can, however, use a common policer instance for all terms within the same firewall filter by setting the *filter-specific* option in the policer. In addition, for IPv4 firewall filters with multiple terms that reference the same policer, filter-specific mode counts and monitors the activity of the policer at the firewall filter level.

NOTE: Both filter-specific and term-specific apply to prefix-specific policer sets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

Support at the `[edit dynamic-profiles ... policer policer-name]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Filter-Specific Policer Overview
- Prefix-Specific Counting and Policing Overview

Filter-Specific Counter and Policer Set Overview

filter (VLANs)

IN THIS SECTION

- Syntax | 7157
- Hierarchy Level | 7157
- Description | 7157
- Default | 7157
- Options | 7158
- Required Privilege Level | 7158
- Release Information | 7158

Syntax

```
filter (input | output) filter-name;
```

Hierarchy Level

```
[edit vlans vlan-name],
[edit vlans vlan-name forwarding-options]
```

Description

Apply a firewall filter to traffic entering or exiting a VLAN.

Default

All incoming traffic is accepted unmodified to the VLAN, and all outgoing traffic is sent unmodified from the VLAN.

Options

filter-name —Name of a firewall filter defined in a `filter` statement.

- **input**—Apply a firewall filter to VLAN ingress traffic.
- **output**—Apply a firewall filter to VLAN egress traffic.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches

Configuring Firewall Filters

Configuring Firewall Filters (CLI Procedure)

Overview of Firewall Filters (QFX Series)

Firewall Filters for EX Series Switches Overview

Configuring VLANs for EX Series Switches with ELS Support (CLI Procedure)

filter (VPLS)

IN THIS SECTION

Syntax | 7159

Hierarchy Level | 7159

- Description | 7159
- Options | 7159
- Required Privilege Level | 7159
- Release Information | 7159

Syntax

```
filter input filter-name;
```

Hierarchy Level

```
[edit forwarding-options family vpls],
[edit logical-systems logical-system-name forwarding-options family vpls],
[edit logical-systems logical-system-name routing -instances instance-name forwarding-options
family vpls],
[edit routing-instances instance-name forwarding-options family vpls]
```

Description

Apply a forwarding table filter for VPLS.

Options

The other statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Table Filters

filtering-type

IN THIS SECTION

- Syntax | 7160
- Hierarchy Level | 7160
- Description | 7160
- Options | 7160
- Required Privilege Level | 7161
- Release Information | 7161

Syntax

```
filtering-type endpoint-independent;
```

Hierarchy Level

```
[edit services nat rule (Services NAT) rule-name term (Services NAT) term-name then (Services NAT)
translated]
```

Description

Specify the NAT filtering behavior for sessions initiated from outside to inside.

Options

endpoint-independent—Currently, the only valid setting specifies endpoint-independent filtering behavior.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.1.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

filtering-type (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 7161
- Hierarchy Level | 7162
- Description | 7162
- Options | 7162
- Required Privilege Level | 7162
- Release Information | 7162

Syntax

```
filtering-type {
    endpoint-independent {
        prefix-list [allowed-host] except [denied-host ];
    }

}
```

Hierarchy Level

```
[edit services nat source rule-set rule-set rule rule-name then source-nat]
```

Description

Specify prefix lists that contain prefixes of hosts that are allowed to establish inbound connections using endpoint-independent mapping, and prefix lists for hosts that are not allowed to establish inbound connections. (Prefix lists are configured at the `[edit policy-options]` hierarchy level.)

Options

[*allowed-host*]	Names of the prefix lists for hosts that are allowed to establish connections.
except [*denied-host*]	Names of prefix lists for hosts that are not allowed to establish connections.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

filters

IN THIS SECTION

- Syntax | 7163
- Hierarchy Level | 7164
- Description | 7164
- Options | 7164

Required Privilege Level | 7164
Release Information | 7164

Syntax

```
filters {
    age-of-attack
        {
        greater-than value;
        less-than value;
    }
                        category (Security Dynamic Attack Group) {
        values [ values];
    }
    cvss-score
        {
        greater-than value;
        less-than value;
    }
    direction {
        expression (and | or);
        values [any client-to-server exclude-any exclude-client-to-server exclude-server-to-
client server-to-client];
    }
    false-positives {
        values [frequently occasionally rarely unknown];
    }
    file-type {
        values [ values ];
    }
    performance {
        values [fast normal slow unknown];
    }
    recommended;
    service {
        values [service-value];
    }
    severity {
```

```
        values [critical info major minor warning];
    }
    type {
        values [anomaly signature];
    }
    vendor name {
        product-name product-name;
    }
    vulnerability-type {
        values [ values ];
    }
}
```

Hierarchy Level

```
[edit security idp dynamic-attack-group dynamic-attack-group-name]
```

Description

To create a dynamic attack group, set the criteria using different types of filters.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3. The `expression` option added in Junos OS Release 11.4.

Additional tags under filters of dynamic attack groups (CVSS score, age-of-attack, file-type, vulnerability-type) are added in Junos OS Release 18.2R1 for dynamic attacks grouping of IDP signatures. The Product and Vendor tags are already supported under existing filter products. The CLI interface for configuring these tags is more user friendly, with possible completions being available for configuration in 18.2R1.

fin-invalidate-session

IN THIS SECTION

- Syntax | 7165
- Hierarchy Level | 7165
- Description | 7165
- Required Privilege Level | 7167
- Release Information | 7167

Syntax

```
fin-invalidate-session;
```

Hierarchy Level

```
[edit security flow tcp-session]
```

Description

Invalidates a TCP session after the 4-way or 3-way handshake completes, with each session endpoint signalling conclusion of the session independently. New incoming SYN packets will need to establish a new TCP session.

When either session endpoint wants to terminate the session, it sends a FIN(ish) message. When the other session endpoint receives the packet with the FIN flag set, it sends an ACK(nowlege) message. Typically, tearing down a session involves transmission of a pair of FIN-ACK messages from each session endpoint.

After the side that sent the first FIN responds with the final ACK, it waits for a time-out period to expire before closing the connection. During the time-out period, the local port cannot be used for new connections. The time-out period protects against delayed packets from the terminating session being delivered during subsequent connections.

> **NOTE**: On SRX Series Firewalls with `fin-invalidate-session` configured the invalidation of the session occurs immediately whereas without `fin-invalidate-session` configured the session is set to time out 2 seconds after the 4-way or 3-way handshake completes.

Table 43 on page 7166 shows the sequence of packets for a 4-way handshake to terminate a session. In this case, the client signals the server that it is terminating the session. The server responds with an ACK message signaling acknowledgement of the client's FIN message. The ACK is followed immediately by a FIN message that the server sends to the client, signaling that it is terminating the session connection on its end. Finally, the client sends an ACK message to the server signalling that it received the server's FIN message.

Table 43: Terminating a Session with a 4-Way Handshake

Step	Client	Server
1.	FIN	
2.		ACK
3.		FIN • Sets session timer to 150 seconds.
4.	ACK • Sets session timer to 2 seconds.	

A session can be terminated by a 3-way handshake. In this case, the client sends a FIN message to the server. The server responds with message that combines the FIN and ACK messages. The sequence of packet exchange for a three-way handshake session close is as follows:

Step	Client	Server
1.	FIN	

2.		FIN/ ACK • sets session timer to 150 seconds
3.	ACK • sets session timer to 2 seconds	

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4 R13.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

fin-no-ack

IN THIS SECTION

- Syntax | 7168
- Hierarchy Level | 7168
- Description | 7168
- Required Privilege Level | 7168
- Release Information | 7168

Syntax

```
fin-no-ack;
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp]
```

Description

Enable detection of an illegal combination of flags, and reject packets that have this combination.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

fin-no-ack (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 7169
- Hierarchy Level | 7169

- Description | 7169
- Required Privilege Level | 7169
- Release Information | 7169

Syntax

```
fin-no-ack;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Description

Identify and drop any packet with a FIN flag set and without the ACK flag set. The TPC FIN No Ack attack can allow the attacker to identify the operating system of the target or to identify open ports on the target.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

final-response-timeout (OCS Partition)

IN THIS SECTION

- Syntax | 7170
- Hierarchy Level | 7170
- Description | 7170
- Options | 7171
- Required Privilege Level | 7171
- Release Information | 7171

Syntax

```
final-response-timeout seconds;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Configure the amount of time in seconds before an Online Charging System (OCS) partition stops attempting to send the final interrogation during the subscriber logout process. When a subscriber starts to logs out, the OCS sends a final interrogation. If there is no response within 24 hours, then the system continues and starts sending a logout message to the Policy and Charging Rule Function (PCRF). If there is no response within the next 24 hours, then the subscriber logout proceeds.

NOTE: Any configuration changes made to this statement apply to all subscribers currently waiting to log out within a 60 second period.

Options

seconds Number of seconds to wait before an OCS partition stops sending the final interrogation using a CCR-GY-T message.

- **Default:** 7200
- **Range:** 0 through 86,400 seconds (24 hours)

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

finger

IN THIS SECTION

- Syntax | 7172
- Hierarchy Level | 7172
- Description | 7172
- Options | 7172

Required Privilege Level | 7173
Release Information | 7173

Syntax

```
finger {
    connection-limit limit;
    rate-limit limit;
}
```

Hierarchy Level

```
[edit system services]
```

Description

Allow finger requests from remote systems to the local device.

Options

connection-limit *limit*

Configure the maximum number of connection sessions for the finger service per protocol (either IPv4 or IPv6).

> **NOTE**: The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

- **Range:** 1 through 250 connections
- **Default:** 75 connections

rate-limit *limit*

Configure the maximum number of connection attempts per minute, per protocol (either IPv6 or IPv4). For example, a rate limit of 10 allows 10 IPv6 finger session connection attempts per minute and 10 IPv4 finger session connection attempts per minute.

- **Range:** 1 through 250 connections
- **Default:** 150 connections

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Finger Service for Remote Access to the Router

fip

IN THIS SECTION

- Syntax | 7173
- Hierarchy Level | 7174
- Description | 7174
- Options | 7174
- Required Privilege Level | 7174
- Release Information | 7174

Syntax

```
fip {
    fcoe-trusted;
```

```
    fc-map fc-map-value;
    fka-adv-period milliseconds;
    interface {
        interface-name {
            fka-adv-period milliseconds;
            priority priority;
        }
    }
    max-sessions-per-enode max-sessions-per-enode;
    priority priority;
    traceoptions {
        file filename <replace> <size size> <files number> <no-stamp>;
        <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit fc-fabrics fc-fabric-name protocols]
```

Description

Configure global or interface-specific FIP options. Individual interface settings override global settings.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fip | 19938

Configuring FIP on an FCoE-FC Gateway

Overview of FIP

fip-security

IN THIS SECTION

- Syntax (EX4600 Switches and QFX Series Switches) | 7175
- Hierarchy Level | 7175
- Description | 7176
- Required Privilege Level | 7176
- Release Information | 7176

Syntax (EX4600 Switches and QFX Series Switches)

```
fip-security {
    examine-vn2vf;
    examine-vn2vn {
        beacon-period milliseconds;
    }
    fc-map fc-map-value;
    interface interface-name {
        (fcoe-trusted | no-fcoe-trusted;)
    }
}
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options]
```

Description

> **NOTE**: This statement supports the Enhanced Layer 2 Software (ELS) CLI. If your switch runs the original (non-ELS) software, see examine-fip. For ELS details, see Using the Enhanced Layer 2 Software CLI.

Configure FIP snooping and FCoE interface properties.

The remaining statements are explained separately. Also, see CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

- *Understanding VN_Port to VF_Port FIP Snooping on an FCoE Transit Switch*
- *Understanding VN_Port to VN_Port FIP Snooping on an FCoE Transit Switch*
- *Understanding FCoE Transit Switch Functionality*
- *Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch*
- *Enabling VN2VN_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch*

firewall

IN THIS SECTION

- Syntax | 7177
- Hierarchy Level | 7178
- Description | 7178

- Required Privilege Level | 7178
- Release Information | 7178

Syntax

```
firewall {
    family family-name {
        filter filter-name {
            interface-specific;
            micro-segmentation;
            no-longest-prefix-match;
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    action;
                    action-modifiers;
                }
            }
        }
    }
    policer policer-name {
        filter-specific;
        if-exceeding {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        then {
            policer-action;
        }
    }
    three-color-policer policer-name {
        action {
            loss-priority high then discard;
        }
        single-rate {
            (color-aware | color-blind);
```

```
            committed-information-rate bps;
            committed-burst-size bytes;
            excess-burst-size bytes;
        }
        two-rate {
            (color-aware | color-blind);
            committed-information-rate bps;
            committed-burst-size bytes;
            peak-information-rate bps;
            peak-burst-size bytes;
        }
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure firewall filters and policers.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Option **micro-segmentation** introduced in Junos OS Release 22.4R1 for supported switches.

RELATED DOCUMENTATION

> *Firewall Filter Match Conditions and Actions (QFX5100, QFX5110, QFX5120, QFX5200, EX4600, EX4650)*

Configuring Firewall Filters

Configuring Two-Color and Three-Color Policers to Control Traffic Rates

Overview of Firewall Filters (QFX Series)

firewall-authentication

IN THIS SECTION

- Syntax | 7179
- Hierarchy Level | 7180
- Description | 7180
- Options | 7181
- Required Privilege Level | 7181
- Release Information | 7181

Syntax

```
firewall-authentication {
    pass-through {
        default-profile profile-name;
        ftp {
             banner {
                fail string;
                login string;
                success string;
            }
        }
        http {
            logo;
            banner {
                fail string;
                login string;
                logout string;
                success string;
```

```
            }
        telnet {
            banner {
                fail string;
                login string;
                success string;
            }
    }
    traceoptions {
        file {
            filename;
            files number;
            flag flag;
            match regular-expression;
            no-remote-trace;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
    }
     web-authentication {
        banner {
            fail string;
            logout string;
            success string;
        }
        default-profile profile-name;
        logo;
        timeout seconds;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure default firewall authentication settings used by firewall authentication policies that restrict and permit access of firewall users to protected resources behind a firewall.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Dynamic VPN Overview

Firewall User Authentication Overview

firewall-authentication-forced-timeout

IN THIS SECTION

- Syntax | 7181
- Hierarchy Level | 7182
- Description | 7182
- Options | 7183
- Required Privilege Level | 7183
- Release Information | 7183

Syntax

```
firewall-authentication-forced-timeout minutes;
```

Hierarchy Level

```
[edit services user-identification active-directory-access]
```

Description

Configure the firewall authentication forced timeout setting to apply to entries for user who authenticate through captive portal.

When a user authenticates through captive portal, an authentication table entry is generated for that user based on the information that the SRX Series Firewall obtains from the firewall authentication module. At that point, the default traffic-based authentication timeout logic is applied to the entry. This statement gives you control over how long non-domain users who authenticate through captive portal remain authenticated.

When the firewall authentication forced timeout value is configured, it is used in conjunction with the traffic-based timeout logic.

Here is how timeout settings affect active directory authentication entries for users authenticated through captive portal.

- The firewall authentication forced timeout is set for 3 hours.

  Traffic continues to be received and generated by a device associated with an authentication entry for a user. After 3 hours the authentication entry expires, although at that time there are sessions anchored in Packet Forwarding Engine for the authentication entry.

- If set, the firewall authentication forced timeout has no effect.

  An authentication entry does not have sessions anchored to it. It expires after the time set for the authentication entry timeout, for example, 30 minutes.

- The firewall authentication forced timeout configuration is deleted.

  Firewall authentication forced timeout has no effect on new authentication entries. Firewall authentication forced timeout remains enforced for existing authentication entries to which it applied before it was deleted. That is, for those authentication entries, the original forced timeout setting remains in effect.

- The firewall authentication forced timeout configuration setting is changed.

  The new timeout setting is applied to new incoming authentication entries. Existing entries keep the original, former setting.

- The firewall authentication forced timeout is set to 0, disabling it.

If the firewall authentication forced timeout is set to a new value, that value is assigned to all incoming authentication entries. There is no firewall authentication forced timeout setting for existing authentication entries.

- The firewall authentication forced timeout value is not configured.
  - The SRX Series Firewall generates an authentication entry for a user. The default traffic-based timeout logic is applied to the authentication entry.
  - The active directory timeout value is configured for 50 minutes. A traffic-based timeout of 50 minutes is applied to an authentication entry.
  - The active directory timeout is not configured. The default traffic-based timeout of 30 minutes is applied to an authentication entry.

Options

minutes The maximum duration for which the non-domain users who authenticate through captive portal remain authenticated.

- **Default:** 30 minutes
- **Range:** 10 through 1440 minutes

Required Privilege Level

1. services—To view this statement in the configuration.
2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Understanding the Forced Timeout Setting Assigned to Active Directory Authentication Entries for Users Authenticated Through Captive Portal

invalid-authentication-entry-timeout

Understanding the Invalid Authentication Table Entry Timeout Setting

firewall-authentication (Security)

IN THIS SECTION

- Syntax | 7184
- Hierarchy Level | 7184
- Description | 7184
- Options | 7185
- Required Privilege Level | 7185
- Release Information | 7185

Syntax

```
firewall-authentication {
    mtls-profile profile-name;
    mtls-profile-fallback-password;
    traceoptions (Security Firewall Authentication) {
        flag (all | authentication | proxy) {
        }
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Defines the type of firewall authentication available for a logical system. Also specifies the data plane firewall authentication tracing options.

Options

- `mtls-profile`—Enable passwordless user authentication based on client/server certificate validation using a public/private key.
- `mtls-profile-fallback-password`—Enable user authentication with mTLS and fallback to standard login when it fails.
- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple flag statements.
  - `all`—Enable all tracing operations.
  - `authentication`—Trace data-plane firewall authentication events.
  - `proxy`—Trace data-plane firewall authentication proxy events.
- `detail`—Display moderate amount of data.
- `extensive`—Display extensive amount of data.
- `terse`—Display minimum amount of data.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Understanding Logical System Firewall Authentication

firewall-authentication (Security Policies)

IN THIS SECTION

- Syntax | 7186
- Hierarchy Level | 7187
- Description | 7187
- Options | 7187
- Required Privilege Level | 7187
- Release Information | 7187

Syntax

```
firewall-authentication {
    pass-through {
        access-profile profile-name;
        client-match user-or-group-name;
        ssl-termination-profile profile-name;
        web-redirect;
        web-redirect-to-https;
        auth-only-browser
            auth-user-agent
    }
    push-to-identity-management
    user-firewall {
        access-profile profile-name;
        domain domain-name
        ssl-termination-profile profile-name;
        web-redirect;
        web-redirect-to-https;
        auth-only-browser
    }
    web-authentication {
        client-match user-or-group-name;
    }
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit]
```

Description

Configure firewall authentication methods.

On SRX Series Firewall, you must configure security policy with firewall-authentication at initial stage. When ssl-termination-profile is engaged with firewall-authentication in security policy, you can't configure `dynamic-application` statement at `[edit security policies from-zone zone-name to-zone zone-name policy policy-name match]` hierarchy level.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Support added for the `user-firewall` option in Junos OS Release 12.1X45-D10.

Support for the `ssl-termination-profile` and `web-redirect-to-https` options added on SRX5600 and SRX5800 Services Gateways starting from Junos OS Release 12.1X44-D10, on SRX5400 devices starting from 12.1X46-D10, and on vSRX Virtual Firewall, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Starting with Junos OS Release 15.1X49-D70 and Junos OS Release 17.3R1, support for the `web-redirect` and `web-redirect-to-https` options under `user-firewall` added on SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall Services Gateways.

Starting with Junos OS Release 15.1X49-D90 and Junos OS Release 17.3R1, support for the `auth-only-browser` option was added under `pass-through` and `user-firewall` and the `auth-user-agent` option was added under `pass-through auth-only-browser` on SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M,

SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall Services Gateways.

Starting with Junos OS Release 15.1X49-D90 and Junos OS Release 17.3R1, support for the `auth-only-browser` option was added under `pass-through` and `user-firewall` and the `auth-user-agent` option was added under `pass-through auth-only-browser` on SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall Services Gateways. Starting with Junos OS Release 15.1X49-D100 and Junos OS Release 17.3R1, support was added for `push-to-identity-management`.

RELATED DOCUMENTATION

Understanding User Role Firewalls

firewall-authentication-service

IN THIS SECTION

- Syntax | 7188
- Hierarchy Level | 7188
- Description | 7189
- Options | 7189
- Required Privilege Level | 7189
- Release Information | 7189

Syntax

```
firewall-authentication-service (enable | disable);
```

Hierarchy Level

```
[edit system processes]
```

Description

Enable or disable the firewall authentication service process.

Options

- `enable`—Start the firewall authentication service process.
- `disable`—Stop the firewall authentication service process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Firewall User Authentication Overview

firewall-authentication (tenants)

IN THIS SECTION

- Syntax | 7190
- Hierarchy Level | 7191
- Description | 7191
- Required Privilege Level | 7191
- Release Information | 7191

Syntax

```
firewall-authentication {
    pass-through {
        default-profile default-profile;
        ftp {
            banner (Access FTP HTTP Telnet Authentication) {
                fail fail;
                login login;
                success success;
            }
        }
        ftp {
            banner (Access FTP HTTP Telnet Authentication) {
                fail fail;
                login login;
                success success;
            }
        }
        ftp {
            banner (Access FTP HTTP Telnet Authentication) {
                fail fail;
                login login;
                success success;
            }
        }
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace;
    }
    web-authentication {
        banner {
            success success;
        }
        default-profile default-profile;
        timeout timeout;
    }
}
```

Hierarchy Level

```
[edit logical-systems name tenants name access],
[edit tenants name access]
```

Description

Define the type of firewall authentication on tenant system.

Required Privilege Level

access

Release Information

Statement introduced in Junos OS Release 18.3R1

RELATED DOCUMENTATION

Firewall Authentication for Tenant Systems

show security firewall-authentication history | **24058**

show security firewall-authentication users | **24071**

firewall-authentication (tenants) | **7189**

firewall (Dynamic Firewalls)

IN THIS SECTION

- Syntax | **7192**
- Hierarchy Level | **7194**
- Description | **7194**
- Required Privilege Level | **7194**
- Release Information | **7194**

Syntax

```
firewall {
    family family {
        fast-update-filter filter-name {
            interface-specific;
            match-order [match-order];
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    action;
                    action-modifiers;
                }
                only-at-create;
            }
        }
        filter filter-name {
            enhanced-mode-override;
            instance-shared;
            interface-shared;
            interface-specific;
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    action;
                    action-modifiers;
                }
            }
        }
    hierarchical-policer uid {
        aggregate {
            if-exceeding {
                bandwidth-limit-limit bps;
                burst-size-limit bytes;
            }
            then {
                policer-action;
            }
```

```
        }
        premium {
            if-exceeding {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
            then {
                policer-action;
            }
        }
    }
    policer uid {
        filter-specific;
        if-exceeding {
            (bandwidth-limit bps | bandwidth-percent percentage);
            burst-size-limit bytes;
        }
        logical-bandwidth-policer;
        logical-interface-policer;
        physical-interface-policer;
        then {
            policer-action;
        }
    }
    three-color-policer uid {
        action {
            loss-priority high then discard;
        }
        logical-interface-policer;
        single-rate {
            (color-aware | color-blind);
            committed-burst-size bytes;
            committed-information-rate bps;
            excess-burst-size bytes;
        }
        two-rate {
            (color-aware | color-blind);
            committed-burst-size bytes;
            committed-information-rate bps;
            peak-burst-size bytes;
            peak-information-rate bps;
            }
```

```
        }
    }
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Configure firewall filters and policers in a dynamic client profile or a dynamic service profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Methods for Regulating Traffic by Applying Hierarchical Policers

Configuring Fast Update Filters

firewall-install-disable

IN THIS SECTION

Syntax | 7195

- Hierarchy Level | 7195
- Description | 7195
- Default | 7195
- Required Privilege Level | 7195
- Release Information | 7196

Syntax

```
firewall-install-disable;
```

Hierarchy Level

```
[edit routing-options flow],
[edit logical-systems logical-system-name routing-options flow],
[edit routing-instances routing-instance-name routing-options flow],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options flow]
```

Description

Disable installing flow-specification firewall filters in the firewall process (dfwd).

Default

For PTX Series routers, this statement appears in the default configuration, preventing installation of flow-specification firewall filters into dfwd. For other models, this setting is omitted from the default configuration, allowing installation of flow-specification firewall filters into dfwd.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 12.1X48 and 12.3.

RELATED DOCUMENTATION

Example: Enabling BGP to Carry Flow-Specification Routes

Understanding BGP Flow Routes for Traffic Filtering

firewall-user

IN THIS SECTION

- Syntax | 7196
- Hierarchy Level | 7196
- Description | 7197
- Options | 7197
- Required Privilege Level | 7197
- Release Information | 7197

Syntax

```
firewall-user {
    password password;
}
```

Hierarchy Level

```
[edit access profile profile-name client client-name]
```

Description

Specify a client as a firewall user and the associated password (encrypted).

Options

`password` *`password`*—Password used by the firewall user during local authentication.

- **Range:** 1 through 128 characters

Required Privilege Level

secret—To view this statement in the configuration.

secret-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring RADIUS and LDAP User Authentication

firmware-revision

IN THIS SECTION

- Syntax | 7198
- Hierarchy Level | 7198
- Description | 7198
- Options | 7198
- Required Privilege Level | 7198
- Release Information | 7198

Syntax

```
firmware-revision firmware-revision;
```

Hierarchy Level

```
[edit access diameter]
```

Description

Configure the firmware revision that is advertised in the Capabilities-Exchange-Request or Capabilities-Exchange-Answer message.

Options

firmware-revision Number of the firmware revision that is the advertised value of the Firmware-Revision AVP.

- **Default:** 0
- **Range:** 0 through 4294967295

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

fiveqi-map

IN THIS SECTION

- Syntax | 7199
- Hierarchy Level | 7199
- Description | 7199
- Options | 7200
- Required Privilege Level | 7200
- Release Information | 7201

Syntax

```
fiveqi-map identifier {
 upstream-rewrite {
  dscp dscp;
  dscp-n3 dscp-n3;
  forwarding-class forwarding-class;
  loss-priority(high | low);
 }
}
```

Hierarchy Level

```
[edit services agf]
```

Description

Configure the Quality of Service (QoS) attributes for a 5G QoS identifier (5QI) value in the 5QI table.

Customers send data packets across the access Y5 interface to the Access Gateway Function (AGF) in clear IP text. AGF encapsulates the data packet using GPRS Tunnelling Protocol (GTP) and forwards the encapsulated packet to the 5GC. To support GTP encapsulation, AGF sets the QFI and DSCP value in the outer IP header based on the packet priority, forwarding class, and packet loss priority of the data.

AGF looks for 5QI and QoS Flow ID (QFI) values in N2 signaling messages from the Access and Mobility Function (AMF) and matches those values to parameters in the 5QI table. AGF derives the forwarding class and loss priority from the 5QI mapping table. AGF classifies data packet priority using the authorized QoS rules sent in the N1 signaling messages. AGF manages individual subscriber session and applies the QoS rules that corresponds to the 5QI in the 5QI mapping table. AGF can then update differentiated services code point (DSCP) in the payload to match the derived QoS parameters. When you configure DSCP parameters in the 5QI mapping table, AGF overrides the payload DSCP. AGF updates both the inner and outer the IP header when it forwards the GTP packet in the data plane.

This statement allows you to:

- Override the QoS attributes for the non-dynamic 5QI.
- Add a dynamic 5QI and its QoS attributes.

You can configure the following QoS attributes for each defined 5QI. AGF uses the information to update the DSCP value in the IP header as it forwards the data packet from the Y5 interface to the N3 interface.

- Forwarding Class
- Packet Loss Priority
- Upstream DSCP
- Upstream N3 DSCP

Options

identifier 5QI number

Range: 0 through 63

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

CoS for Subscriber Access Overview

upstream-rewrite | **30481**

show services agf fiveqi remarking | **25331**

fixed-stuff-bytes

IN THIS SECTION

- Syntax | **7201**
- Hierarchy Level | **7201**
- Description | **7202**
- Default | **7202**
- Options | **7202**
- Required Privilege Level | **7202**
- Release Information | **7202**

Syntax

```
(fixed-stuff-bytes | no-fixed-stuff-bytes);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

Description

Enable or disable fixed stuff bytes.

Default

By default, no fixed stuff bytes are set.

Options

fixed-stuff-bytes	Fixed stuff bytes 11.0957 Gbps.
no-fixed-stuff-bytes	No fixed stuff bytes 11.0491 Gbps.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

10-Gigabit Ethernet OTN Options Configuration Overview

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

fka-adv-period

IN THIS SECTION

Syntax | 7203

- Hierarchy Level | 7203
- Description | 7203
- Options | 7203
- Required Privilege Level | 7203
- Release Information | 7204

Syntax

```
fka-adv-period milliseconds;
```

Hierarchy Level

```
[edit fc-fabrics fc-fabric-name protocols fip],
[edit fc-fabrics fc-fabric-name protocols fip interface interface-name]
```

Description

Set the global or interface-specific interval between periodic FIP keepalive advertisements. An interval set at the interface level overrides the global setting.

Options

milliseconds—Time in milliseconds between FIP keepalive advertisements.

- **Range:** 250 through 90000 milliseconds
- **Default:** 8000 milliseconds

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fip | 19938

show fibre-channel fip interface | 19961

Overview of FIP

flag (App Engine Virtual Machine Management Service)

IN THIS SECTION

- Syntax | 7204
- Hierarchy Level | 7204
- Description | 7205
- Default | 7205
- Options | 7205
- Required Privilege Level | 7205
- Release Information | 7205

Syntax

```
flag (all |ccif | configuration | heartbeat | init | miscellaneous | platform | pxe | routing-
instances | snmp)
```

Hierarchy Level

```
[edit system processes app-engine-virtual-machine-management-service traceoptions]
```

Description

Perform different tracing operations. To specify more than one tracing operation, include multiple `flag` statements.

Default

Tracing operations are not performed.

Options

`all`	Trace all events.
`ccif`	Trace compute node interface events. This is the default option.
`configuration`	Trace configuration events.
`heartbeat`	Trace compute node heartbeat-related events.
`init`	Trace initialization events.
`miscellaneous`	Trace miscellaneous events.
`platform`	Trace platform-related events.
`pxe`	Trace events related to Preboot Execution Environment (PXE).
`routing-instances`	Trace events related to routing instances.
`snmp`	Trace SNMP events.

Required Privilege Level

system-control

Release Information

Statement introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

traceoptions (App Engine Virtual Machine Management Service) | 29497
file (App Engine Virtual Machine Management Service) | 7030
level (App Engine Virtual Machine Management Service) | 9719

flag (Interfaces)

IN THIS SECTION

- Syntax | 7206
- Hierarchy Level | 7206
- Description | 7206
- Options | 7207
- Required Privilege Level | 7207
- Release Information | 7207

Syntax

```
flag
```

Hierarchy Level

```
[edit interfaces interface-name traceoptions]
```

Description

Define tracing operations for individual interfaces. To specify more than one tracing operation, include multiple flag statements.

Options

- **all**—Enable all interface trace flags.
- **event** —Trace interface events.
- **cache**—Enable interface flags for Web filtering cache maintained on the routing table.
- **enhanced**—Enable interface flags for processing through Enhanced Web Filtering.
- **ipc**—Trace interface IPC messages.
- **media**—Trace interface media changes.
- **critical**—Trace critical events.
- **major**—Trace major events.

> **NOTE**:
>
> - MTU is limited to 1518 on this interface.
> - **Cache** and **enhanced** options are applicable only to Enhanced Web Filtering.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Understanding Interfaces

flag (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | 7208
- Hierarchy Level | 7208
- Description | 7208
- Options | 7208
- Required Privilege Level | 7209
- Release Information | 7209

Syntax

```
flag name;
```

Hierarchy Level

```
[edit services rtlog traceoptions]
```

Description

List of things to include in trace.

Options

name

- Values:
  - all—Enable all interface trace flags. event —Trace interface events.
  - cache—Enable interface flags for Web filtering cache maintained on the routing table.

- enhanced—Enable interface flags for processing through Enhanced Web Filtering.
- ipc—Trace interface IPC messages.
- media—Trace interface media changes.
- critical—Trace critical events.
- major—Trace major events

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

flag (Services)

IN THIS SECTION

- Syntax | 7210
- Hierarchy Level | 7210
- Description | 7210
- Options | 7210
- Required Privilege Level | 7210
- Release Information | 7210

Syntax

```
flag (all | cli-configuration | initiation | proxy | selected-profile | termination);
```

Hierarchy Level

```
[edit services ssl traceoptions]
```

Description

Specify the tracing flag parameters.

Options

- *all*—Trace all the parameters.
- *cli-configuration*—Trace CLI configuration events.
- *initiation*—Trace initiation service events.
- *proxy*—Trace proxy service events.
- *selected-profile*—Trace events for profiles with `enable-flow-tracing` set.
- *termination*—Trace termination service events.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. This statement is supported.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

flag (SMTP)

IN THIS SECTION

- Syntax | 7211
- Hierarchy Level | 7211
- Description | 7211
- Options | 7211
- Required Privilege Level | 7212
- Release Information | 7212

Syntax

```
flag {
    all;
    configuration;
    IPC;
    protocol-exchange;
    send-request;
}
```

Hierarchy Level

```
[edit smtp traceoptions]
```

Description

Set flag for the SMTP traceoptions.

Options

The following flag options are supported:

- `IPC`—Trace interprocesss communication.

- `all`—Trace everything.
- `configuration`—Trace configuration event.
- `protocol-exchange`—Trace SMTP protocol exchanges.
- `send-request`—Trace send mail request event.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 10.0.

RELATED DOCUMENTATION

smtp-profile (Security Content Security Policy Antispam) | 28051

flag (Tracing for Origin AS Validation)

IN THIS SECTION

- Syntax | 7213
- Hierarchy Level | 7213
- Description | 7213
- Options | 7213
- Required Privilege Level | 7214
- Release Information | 7214

Syntax

```
flag flag <flag-modifier> <disable>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation traceoptions],
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation group group-name session address traceoptions],
[edit logical-systems logical-system-name routing-options validation traceoptions],
[edit logical-systems logical-system-name routing-options validation group group-name session
address traceoptions],
[edit routing-instances instance-name routing-options validation traceoptions],
[edit routing-instances instance-name routing-options validation group group-name session
address traceoptions],
[edit routing-options validation traceoptions]
[edit routing-options validation group group-name session address traceoptions]
```

Description

Configure the flags for tracing resource public key infrastructure (RPKI) BGP route validation.

Options

flag Tracing operation to perform. To specify more than one RPKI BGP route validation tracing operation, include multiple `flag` statements.

- Values:
  - all—Trace all events.
  - error—Trace errored packets.
  - keepalive— RPKI-to-router protocol keepalive messages. If you enable the BGP `update` flag only, received keepalive messages do not generate a trace message.

- nsr-synchronization— Nonstop routing synchronization events. (Not valid at [edit routing-options validation group *group-name* session *address*].
- packets—All incoming and outgoing packets.
- policy—Policy processing. (Not valid at [edit routing-options validation group *group-name* session *address*].
- state—State transitions.
- task—Routing protocol task processing.
- timer—Routing protocol timer processing.
- update—Update packets. These packets provide routing updates to BGP systems. If you enable only this flag, received keepalive messages do not generate a trace message. Use the `keepalive` flag to generate a trace message for keepalive messages.

disable (Optional) Disable the specific tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

flag-modifier (Optional) Modifier for the tracing flag. You can specify one or more of these modifiers.

- Values:
  - detail—Provide detailed trace information.
  - receive—Trace the packets being received.
  - send—Trace the packets being transmitted.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring Origin Validation for BGP

Understanding Origin Validation for BGP

Use Case and Benefit of Origin Validation for BGP

flap-trap-monitor

IN THIS SECTION

- Syntax | 7215
- Hierarchy Level | 7215
- Description | 7216
- Options | 7216
- Required Privilege Level | 7216
- Release Information | 7216

Syntax

```
flap-trap-monitor seconds;
```

Hierarchy Level

```
[edit logical-systems name protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
[edit logical-systems name protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles profile-name avg-fd-twoway-threshold],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-
profilesprofile-name avg-fd-twoway-threshold]
[edit logical-systems name protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles profile-name avg-ifdv-twoway-threshold],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-
profilesprofile-name avg-ifdv-twoway-threshold]
[edit logical-systems name protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles profile-name avg-flr-forward-threshold],
```

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profilesprofile-nameavg-flr-forward-threshold ]
[edit logical-systems name protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles profile-name avg-flr-backward-threshold],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profilesprofile-nameavg-flr-backward-threshold]
```

Description

Enables damping of jnxSoamPmThresholdCrossingAlarm traps sent to the network management system (NMS) by summarizing the flap occurrences over a period of time and sends a single jnxSoamPmThresholdFlapAlarm notification to the NMS. You can enable damping at the global level for the iterator or you can enable damping at the individual threshold type of an iterator. You can specify the duration of time for summarizing flap occurrences.

Options

seconds—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the NMS.

- **Range:** 1 through 360 seconds

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

avg-fd-twoway-threshold | **2775**

avg-ifdv-twoway-threshold | **2781**

avg-flr-backward-threshold | **2777**

avg-flr-forward-threshold | **2779**

sla-iterator-profiles | **28013**

Damping CFM performance Monitoring Traps and Notifications to Prevent Congestion of The NMS

flat-file-profile (Accounting Options)

IN THIS SECTION

- Syntax | 7217
- Hierarchy Level | 7219
- Description | 7219
- Options | 7219
- Required Privilege Level | 7219
- Release Information | 7219

Syntax

```
flat-file-profile profile-name{
    fields {
        all-fields;
        egress-stats {
            all-fields;
            input-bytes;
            input-packets;
            output-bytes;
            output-packets;
            queue-id;
            red-drop-bytes;
            red-drop-packets;
            tail-drop-packets;
            total-drop-packets;
        }
        general-param {
            all-fields;
            accounting-type;
            descr;
            line-id;
            logical-interface;
            nas-port-id;
            physical-interface;
            routing-instance;
```

```
            timestamp;
            user-name;
            vlan-id;
        }
        ingress-stats {
            all-fields;
            drop-packets;
            input-bytes;
            input-packets;
            output-bytes;
            output-packets;
            queue-id;
        }
        l2-stats {
            all-fields;
            input-mcast-bytes;
            input-mcast-packets;
        }
        overall-packet {
            all-fields;
            input-bytes;
            input-discards;
            input-errors;
            input-packets;
            inputv6-bytes;
            inputv6-packets;
            output-bytes;
            output-errors;
            output-packets;
            outputv6-bytes;
            outputv6-packets;
            input-v4-bytes;
            input-v4-packets;
            output-v4-bytes;
            output-v4-packets;
            input-bytes-per-sec;
            input-packets-per-sec;
        }
        service-accounting;
    }
    file filename;
    format (csv | ipdr)
    interval minutes;
```

```
    start-time time;
     schema-version schema-name;
     use-fc-ingress-stats;
}
```

Hierarchy Level

[edit accounting-options]

Description

Configure a flat-file accounting profile that defines the contents of a flat file that records accounting statistics collected from the Packet Forwarding Engine for an interface at regular intervals. To be used, the profile is associated with a subscriber interface. The accounting flat file is archived by the accounting-options archiving mechanism.

Options

profile-name	Name of the flat-file profile.
start-time *time*	(Optional) Specify the start time in *yyyy-mm-dd.hh:mm* local time format for the accounting profile. Sometimes, if the time at which the records are written to the accounting file is beyond the transfer interval window of the file, drifting of accounting records occurs. To prevent drifting and to have predictable timestamps of when accounting data is written to an accounting file, we can specify the start time option for the accounting profile.
use-fc-ingress-stats	Enables forwarding class counters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

`service-accounting` option added in Junos OS Release 17.1.

Support for the `service-accounting` option removed in Junos OS Release 18.1R4.

`use-fc-ingress-stats` option added in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

- Configuring Accounting-Data Log Files
- *Configuring Flat-File Accounting for Layer 2 Wholesale*
- *Configuring Flat-File Accounting for Extensible Subscriber Services Management*
- *Configuring Service Accounting in Local Flat Files*
- *Flat-File Accounting Overview*

flat-file-profile (Extensible Subscriber Services)

IN THIS SECTION

- Syntax | **7220**
- Hierarchy Level | **7220**
- Description | **7221**
- Options | **7221**
- Required Privilege Level | **7221**
- Release Information | **7221**

Syntax

```
flat-file-profile profile-name
```

Hierarchy Level

```
[edit system services extensible-subscriber-services]
```

Description

Specify the name of an accounting flat file profile that applies to an ESSM subscriber.

Options

profile name Name of an accounting flat file profile configured at the `[edit accounting-options]` hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Extensible Subscriber Services Management

flex-algorithm

IN THIS SECTION

- Syntax | 7222
- Hierarchy Level | 7222
- Description | 7222
- Options | 7222
- Required Privilege Level | 7223
- Release Information | 7223

Syntax

```
flex-algorithm name {
    color color;
    definition;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options],
[edit logical-systems name routing-options],
[edit logical-systems name tenants name routing-instances name routing-options],
[edit routing-instances name routing-options],
[edit routing-options],
[edit tenants name routing-instances name routing-options],
[edit logical-systems name  routing-instances name protocols isis source-packet-routing],
[edit logical-systems name protocols isissource-packet-routing],
{edit logical-systems name tenants name routing-instances name protocols isis source-packet-
routing],
[edit protocols isis source-packet-routing],
[edit routing-instances name protocols isis source-packet-routing],
[edit tenants name routing-instances nameprotocolsisissource-packet-routing]
```

Description

Define a flexible algorithm for IS-IS to compute a path based on specified parameters to thin slice a network. We recommend configuring flexible algorithms on only a couple of routers to provide redundancy and to avoid conflicts.

Configure participation of routers in a specific flexible algorithm in a network at the `[edit protocols]` hierarchy level.

> **NOTE**: Modifying the flexible algorithm definition could cause traffic disruptions until all the nodes converge on the new paths.

Options

name Specify a flexible algorithm id to identify a flexible algorithm definition.

- **Range:** 128 through 255

color Specify a color to be associated with this flexible algorithm.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

definition | 5231

Configuring Flexible Algorithm for Segment Routing Traffic Engineering

Understanding IS-IS Flexible Algorithms for Segment Routing

flex-algorithm (Protocols OSPF)

IN THIS SECTION

- Syntax | 7224
- Hierarchy Level | 7224
- Description | 7224
- Options | 7225
- Required Privilege Level | 7225
- Release Information | 7225

Syntax

```
flex-algorithm name {
    color color;
   definition (Protocols OSPF) ;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options],
[edit logical-systems name routing-options],
[edit logical-systems name tenants name routing-instances name routing-options],
[edit routing-instances name routing-options],
[edit routing-options],
[edit tenants name routing-instances name routing-options]
[edit logical-systems name routing-instances name
protocols ospf source-packet-routing (Protocols OSPF)],
[edit logical-systems name protocols ospf source-packet-
routing (Protocols OSPF)]
[edit logical-systems name tenants name routing-instances name
protocols ospf source-packet-routing (Protocols OSPF)],
[edit protocols ospf source-packet-routing (Protocols OSPF)],
[edit routing-instances name protocols ospf source-packet-
routing (Protocols OSPF)],
[edit tenants name routing-instances name protocols ospf source-packet-routing (Protocols OSPF)]
```

Description

Define a flexible algorithm for OSPFv2 to compute a path based on specified parameters to thin slice a network. We recommend configuring flexible algorithms on only a couple of routers to provide redundancy and to avoid conflicts.

Configure participation of routers in a specific flexible algorithm in a network at the `[edit protocols]` hierarchy level.

NOTE: Modifying the flexible algorithm definition could cause traffic disruptions until all the nodes converge on the new paths.

Options

name Specify a flexible algorithm id to identify a flexible algorithm definition.

- **Range:** 128 through 255

color Specify a color to be associated with this flexible algorithm.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

definition (Protocols OSPF)

How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering

flex-flow-sizing

IN THIS SECTION

- Syntax | 7226
- Hierarchy Level | 7226
- Description | 7226
- Options | 7226
- Required Privilege Level | 7226
- Release Information | 7226

Syntax

```
flex-flow-sizing;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services]
```

Description

Configure support for the service creation of flows for inline services sampling. This configuration results in a first-come-first-serve creation of flows. Whichever flow comes first, that is allowed to occupy the flow-table if there is space in the table. Otherwise, the flow is dropped and an error count is created.

> **NOTE**: You cannot configure the `explicit flow-table-sizes` because `flex-flow-sizing` and `explicit flow-table-sizes` are mutually exclusive.
>
> You need not perform `fpc` reboot to change from `flex` to `per family` configuration.

Options

- **Default:** 1K flows for IPv6 and VPLS flows each.
- **Range:** 15 through 256K flows for IPv4.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F5.

> **NOTE**: As a workaround for `flex-flow-sizing` on MX204 router for releases earlier than Junos OS Release 17.4R2 or Junos OS Release 18R3, replace the `flex-flow-sizing`

statement with the `flow-table-size` statement configuration and reload the box. Configure the `flow-table-size` statement value within a range of 1 through 15. If the `flex-flow-sizing` statement is configured, deactivate or delete the same. For example:

```
flow-table-size {
    ipv4-flow-table-size 10;
    ipv6-flow-table-size 4;
}
```

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers

flex-hashing

IN THIS SECTION

- Syntax | 7227
- Hierarchy Level | 7229
- Description | 7229
- Options | 7229
- Required Privilege Level | 7229
- Release Information | 7229

Syntax

```
flex-hashing name {
    ethtype {
        inet {
```

```
            conditional-match conditional-match;
            hash-offset {
                offset1 {
                    base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
OuterHeader | start-of-L4-Header);
                    offset1-mask offset1-mask;
                    offset1-value offset1-value;
                    offset2 {
                        base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
OuterHeader | start-of-L4-Header);
                        offset2-mask offset2-mask;
                        offset2-value offset2-value;
                    }
                }
            }
            interface interface;
        }
        inet6 {
            conditional-match conditional-match;
            hash-offset {
                offset1 {
                    base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
OuterHeader | start-of-L4-Header);
                    offset1-mask offset1-mask;
                    offset1-value offset1-value;
                    offset2 {
                        base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
OuterHeader | start-of-L4-Header);
                        offset2-mask offset2-mask;
                        offset2-value offset2-value;
                    }
                }
            }
            interface interface;
        }
        mpls {
            conditional-match conditional-match;
            hash-offset {
                offset1 {
                    base-offset1 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
OuterHeader | start-of-L4-Header);
                    offset1-mask offset1-mask;
                    offset1-value offset1-value;
```

```
                    offset2 {
                        base-offset2 (start-of-L2Header | start-of-L3-InnerHeader | start-of-L3-
OuterHeader | start-of-L4-Header);
                        offset2-mask offset2-mask;
                        offset2-value offset2-value;
                    }
                }
            }
            interface interface;
            num-labels num-labels;
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key]
```

Description

Configure user-defined flex hashing to load balance MPLS traffic based on TCP or UDP source/destination port information.

Options

name Name identifier

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

conditional-match | 4504

no-inner-payload | 11474

flexible-cross-connect-vlan-aware

IN THIS SECTION

- Syntax | 7230
- Hierarchy Level | 7230
- Description | 7230
- Required Privilege Level | 7231
- Release Information | 7231

Syntax

```
flexible-cross-connect-vlan-aware;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols evpn],
[edit routing-instances name protocols evpn]
```

Description

Enable EVPN-VPWS flexible cross-connect (FXC) VLAN aware service. For FXC service, all point-to-point Ethernet services under the same EVPN-VPWS instance share the same incoming MPLS label. FXC allows you to aggregate attachment circuits in the same EVPN-VPWS instance into a single group. For FXC VLAN aware services, you assign a unique service identifier to each attachment circuits in the same group. As a result, each attachment circuit will have its own Ethernet Auto-Discovery per EVPN Instance (EVI) route. When you enable `flexible-cross-connect-vlan-aware` services, you configure the pseudowire service (PS) interface under the routing protocol EVPN hierarchy.

The following sample configuration shows an EVPN-VPWS `flexible-cross-connect-vlan-aware` routing instance:

```
routing-instances evpn-vpws-A {
   instance-type evpn-vpws;
   interface ps0.0;
   route-distinguisher 192.168.2.3:100;
   vrf-target target:100:100;
   protocols {
      evpn {
         flexible-cross-connect-vlan-aware;
         interface ps0.1 {
            vpws-service-id {
               local 101;
               remote 201;
            }
         }
      }
    }
 }
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Overview of Flexible Cross-Connect Support on VPWS with EVPN

auto-service-id | **2739**

flexible-cross-connect-vlan-unaware | **7232**

flexible-cross-connect-vlan-unaware

IN THIS SECTION

- Syntax | 7232
- Hierarchy Level | 7232
- Description | 7232
- Required Privilege Level | 7233
- Release Information | 7233

Syntax

```
flexible-cross-connect-vlan-unaware;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols evpn],
[edit routing-instances name protocols evpn]
```

Description

Enable EVPN-VPWS flexible cross-connect (FXC) VLAN unaware Service. For EVPN-VPWS FXC service, all point-to-point Ethernet services under the same EVPN VPWS instance share the same incoming MPLS label. FXC allows you to aggregate attachment circuits into a single group. When you enable `flexible-cross-connect-vlan-unaware` services, you can bundle a number of attachment circuits in the same group. The group of attachment circuits share the same VPWS service tunnel, local service instance ID, and Ethernet Auto-Discovery per EVPN Instance (EVI) route.

The following sample configuration shows a `flexible-cross-connect-vlan-unaware` routing instance.

```
routing-instances bng-2
```

```
    protocols {
       evpn {
          label-allocation per-instance;
          flexible-cross-connect-vlan-unaware;
          group ac1 {
             interface xe-0/0/32.101;
             interface xe-0/0/30.101;
             interface xe-0/0/32.102;
             service-id {
                 local 101;
                 remote 102;
             }
         }
     }
 }
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Overview of Flexible Cross-Connect Support on VPWS with EVPN

flexible-cross-connect-vlan-aware | 7230

flexible-queuing-mode

IN THIS SECTION

Syntax | 7234

- Hierarchy Level | 7234
- Description | 7234
- Required Privilege Level | 7234
- Release Information | 7235

Syntax

```
flexible-queuing-mode;
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable flexible queuing on a non-HQoS MPCE that is installed in an MPC slot. A maximum of up to 32,000 queues are supported per port and per card, including queues on both ingress and egress interfaces.

When flexible queuing is enabled, the non-HQoS MPC is restarted for the changes to take effect and is brought online only if the power required for the queuing component is available in the power entry module (PEM). The MPC remains offline if the PEM cannot meet the power requirement for the queuing component.

You can configure flexible queuing even when a non-HQoS MPC is not present in the chassis. The configuration takes effect when a non-HQoS MPC is installed.

For more information about the MPCs and the Junos OS release that support this feature, see *Flexible Queuing Mode Overview*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R1.

RELATED DOCUMENTATION

Upgrading non-HQoS MPCs to Support Flexible Queuing

Flexible Queuing Mode Overview

Protocols and Applications Supported by the MPC5E for MX Series Routers

flexible-vlan-tagging

IN THIS SECTION

- Syntax | 7235
- Hierarchy Level | 7235
- Description | 7236
- Required Privilege Level | 7236
- Release Information | 7236

Syntax

```
flexible-vlan-tagging;
```

Hierarchy Level

```
[edit interfaces aex],
[edit interfaces ge-fpc/pic/port],
[edit interfaces et-fpc/pic/port],
[edit interfaces ps0],
[edit interfaces xe-fpc/pic/port]
```

Description

Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.

This statement is supported on M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2, IQ2-E, and IQ PICs or in MX Series DPCs, or on Ethernet interfaces for PTX Series Packet Transport Routers or 100-Gigabit Ethernet Type 5 PIC with CFP.

This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series and QFX Series switches.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

Support for aggregated Ethernet added in Junos OS Release 9.0.

For the list of products and applications that support this feature on Junos OS Evolved, see Flexible VLAN Tagging.

RELATED DOCUMENTATION

- Enabling VLAN Tagging
- Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers
- Configuring Double-Tagged VLANs on Layer 3 Logical Interfaces

flexible-vlan-tagging (Interfaces)

IN THIS SECTION

- Syntax | 7237
- Hierarchy Level | 7237
- Description | 7237
- Options | 7237
- Required Privilege Level | 7238
- Release Information | 7238

Syntax

```
flexible-vlan-tagging;
```

Hierarchy Level

```
[edit interfaces interface ]
```

Description

Simultaneously supports transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port.

NOTE: The `flexible-vlan-tagging` is supported only with either no encapsulation or VPLS VLAN encapsulation.

Options

`native-vlan-id`—Configures a VLAN identifier for single-tag frames, dual-tag frames, or a mixture of single-tag and dual-tag frames.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Configuring VLAN Tagging

flood

IN THIS SECTION

- Syntax | 7238
- Hierarchy Level | 7239
- Description | 7239
- Options | 7239
- Required Privilege Level | 7239
- Release Information | 7239

Syntax

```
flood {
    input filter-name;
}
```

Hierarchy Level

```
[edit forwarding-options family vpls],
[edit routing-instances instance-name forwarding-options family vpls]
```

Description

Apply a forwarding table filter to a flood table.

Options

input *filter-name*—Name of the forwarding table filter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Table Filters

Layer 2 Port Mirroring Firewall Filters

flood-as-unknown-unicast

IN THIS SECTION

- Syntax | 7240
- Hierarchy Level | 7240

- Description | 7240
- Required Privilege Level | 7241
- Release Information | 7241

Syntax

```
flood-as-unknown-unicast;
```

Hierarchy Level

```
[edit interfaces irb unit name family inet address <local-ip-address>/<prefix-length> | arp
<local-ip-address> mac <local-ip-address>]
```

Description

You can configure the `flood-as-unknown-unicast` statement for IRB interfaces located in any Layer 2 domain.

You can enable the `flood-as-unknown-unicast` statement for unicast MAC addresses. When a switch receives traffic from a unicast MAC address that is the destination MAC address, the switch floods the packets on the VLAN where the IRB interface is located.

When a Microsoft load balancer cluster is connected to a device in an EVPN-VXLAN or Layer2 EVPN-MPLS domain, the nodes in the cluster don't export their virtual MAC addresses along with their address resolution protocol (ARP) packets. In this case, the switch doesn't learn the virtual MAC addresses and drops any traffic coming from those addresses. When you enable the `flood-as-unknown-unicast` statement for a virtual MAC address on your switch, the switch can learn the Microsoft load balancer node's virtual MAC address. When the virtual MAC address is the destination MAC, the switch floods the traffic on the Microsoft load balancer nodes.

NOTE: If you've configured assisted replication, flooding occurs on assisted replication virtual tunnel endpoints (VTEP), not regular VTEPs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1.

flood-groups

IN THIS SECTION

- Syntax | 7241
- Hierarchy Level | 7241
- Description | 7242
- Options | 7242
- Required Privilege Level | 7242
- Release Information | 7242

Syntax

```
flood-groups [ ip-addresses ];
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name multicast-snooping-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name multicast-snooping-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name multicast-
snooping-options],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name multicast-snooping-
```

```
options],
[edit routing-instances routing-instance-name multicast-snooping-options]
```

Description

Establish a list of flood group addresses for multicast snooping.

Options

ip-addresses—List of IP addresses subject to flooding.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring Multicast Snooping

flood-reflector (Protocols IS-IS)

IN THIS SECTION

- Syntax | 7243
- Hierarchy Level | 7243
- Description | 7243
- Options | 7243
- Required Privilege Level | 7244
- Release Information | 7244

Syntax

```
flood-reflector {
    client;
        reflector {
        cluster-id;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis interface interface-name level level-number],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis interface interface-name level level-number],
[edit protocols isis interface interface-name level level-number],
[edit protocols isis level level-number],
[edit routing-instances routing-instance-name protocols isis interface interface-name level level-number]
```

Description

IS-IS supports flood-reflector interfaces. Flood reflectors enable the creation of topologies where Level 1 areas provide transit forwarding for Level 2 destinations within a Level 2 topology. This is accomplished by creating Level 2 flood-reflection adjacencies within each Level 1 area.

The Level 2 flood reflection adjacencies are used to flood Level 2 link-state packet data units (PDUs), and are used in the Level 2 SPF computation. Therefore, the flood reflectors limit link-state PDU flooding over IS-IS interfaces. We designate the Flexible Tunnel Interfaces (FTI) as flood-reflector interfaces.

Options

client Configure interface as a flood-reflector client.

reflector Configure interface as a flood reflector.

cluster-id The flood-reflector cluster identifier that allows a flood reflector client to establish flood reflector adjacencies with multiple flood reflectors.

- **Range:** 1 through 42,94,967,295

> **NOTE:**
>
> - `flood-reflector` must be configured on the FTI interface only.
> - `flood-reflector` must be configured at IS-IS Level 2 only.
> - The option `reflector` cannot be configured together with the option `client`. This enforces that the same router cannot be both a flood reflector and a flood-reflector client for the same level.
> - At the `[edit protocols isis level level number flood-reflector reflector]` hierarchy level, the `cluster-id` must be configured.
> - At the `[edit protocols isis level level number flood-reflector client]` hierarchy level, the `cluster-id` must not be configured.
> - To configure the `cluster-id` option at the `[edit protocols isis interface interface name level level number flood-reflector reflector]` hierarchy level, the `client` must be configured at the `[edit protocols isis level level number flood-reflector client]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

How to Configure Flood-Reflector Interfaces in IS-IS Networks

flood (Security ICMP)

IN THIS SECTION

- Syntax | 7245
- Hierarchy Level | 7245
- Description | 7245
- Options | 7245
- Required Privilege Level | 7246
- Release Information | 7246

Syntax

```
flood {
    threshold number;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name icmp]
```

Description

Configure the device to detect and prevent Internet Control Message Protocol (ICMP) floods. An ICMP flood occurs when ICMP echo requests are broadcast with the purpose of flooding a system with so much data that it first slows down, and then times out and is disconnected. The threshold defines the number of ICMP packets per second allowed to ping the same destination address before the device rejects further ICMP packets.

Options

`threshold` *`number`* —Number of ICMP packets per second allowed to ping the same destination address before the device rejects further ICMP packets.

- **Range:** 1 through 1,000,000 per second
- **Default:** 1,000 per second

> **NOTE**: For SRX Series Firewalls the applicable range is 1 through 4,000,000 per second.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

flood (Security UDP) | 7246

flood (Security UDP)

IN THIS SECTION

- Syntax | 7247
- Hierarchy Level | 7247
- Description | 7247
- Options | 7247
- Required Privilege Level | 7248
- Release Information | 7248

Syntax

```
flood {
    threshold number;
    white-list name;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name udp]
```

Description

Configure the device to detect and prevent UDP floods. UDP flooding occurs when an attacker sends UDP packets to slow down the system to the point that it can no longer process valid connection requests.

The threshold defines the number of UDP packets per second allowed to ping the same destination IP address/port pair. When the number of packets exceeds this value within any 1-second period, the device generates an alarm and drops subsequent packets for the remainder of that second.

Options

`threshold` *`number`* —Number of UDP packets per second allowed to ping the same destination address before the device rejects further UDP packets.

- **Range:** 1 through 4,000,000 per second
- **Default:** 1,000 per second

For SRX300, SRX320, SRX340, and SRX345, the applicable range is 1 through 100,000 per second.

For SRX1500, SRX4100, SRX4200, and vSRX Virtual Firewall, the applicable range is 1 to 1,000,000.

For SRX4600, SRX5400, SRX5600, and SRX5800, the applicable range is 1 through 4,000,000 per second.

`white-list` *`name`* —Configure a list of IP addresses that are exempted from UDP flood detection, which occur during the UDP flood screen protection process. This list of exempted addresses is called an allowlist.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

flood (Security ICMP) | **7245**

white-list (Security Screen) | **31107**

flood (VLANs)

IN THIS SECTION

- Syntax | **7249**
- Hierarchy Level | **7249**
- Description | **7249**
- Default | **7249**
- Options | **7249**
- Required Privilege Level | **7249**
- Release Information | **7249**

Syntax

```
flood {
    input filter-name;
}
```

Hierarchy Level

```
[edit vlans vlan-name],
[edit vlans vlan-name forwarding-options (Bridge Domains or VLANs)]
```

Description

Apply a flood filter to traffic ingressing a VLAN. Flood filters are triggered only for broadcast, unknown unicast, and multicast (BUM) traffic.

Flood filters and firewall filters can coexist on the same VLAN. If the actions in the filters are conflicting, then the firewall filter takes priority over the flood filter.

Default

All incoming traffic is accepted unmodified to a VLAN, and all outgoing traffic is sent unmodified from a VLAN.

Options

filter-name—Name of a filter defined at the `[edit firewall family family-name filter]` hierarchy level.

`input`—Apply a flood filter to VLAN ingress traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring Unknown Unicast Forwarding (ELS)

Configuring Firewall Filters

Overview of Firewall Filters (QFX Series)

flow (protocols bgp)

IN THIS SECTION

- Syntax | 7250
- Hierarchy Level | 7250
- Description | 7251
- Options | 7251
- Required Privilege Level | 7251
- Release Information | 7251

Syntax

```
flow {
    no-validate policy-name;
    secondary-independent-resolution;
}
```

Hierarchy Level

```
[edit protocols bgp group group-name family (inet | inet-vpn | inet6 | inet6-vpn)],
[edit protocols bgp group group-name neighbor address family (inet | inet-vpn | inet6 | inet6-
vpn)],
[edit routing-instances routing-instance-name protocols bgp group group-name family (inet | inet-
vpn | inet6 | inet6-vpn)],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (inet | inet-vpn | inet6 | inet6-vpn)]
```

Description

Enables BGP to support flow routes.

> **NOTE**: This statement is supported for the default instance, VRF instance, and virtual-router instance only. It is configured with the `instance-type` statement at the `[edit routing-instance instance-name]` hierarchy level.
>
> The remaining statements are explained separately. See CLI Explorer.

Options

secondary-independent-resolution Configure to resolve flow specification routes in the VRF table independent of VPN flow route. This option is currently supported for inet and inet-vpn families only.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`secondary-independent-resolution` option introduced in Junos OS Release 18.4R1 for the MX Series and PTX Series.

RELATED DOCUMENTATION

Example: Enabling BGP to Carry Flow-Specification Routes

flow

IN THIS SECTION

- Syntax | 7252
- Hierarchy Level | 7254
- Description | 7254
- Default | 7254
- Options | 7254
- Required Privilege Level | 7259
- Release Information | 7259

Syntax

```
    flow {
    discard-action-for-unresolved-redir-addr;
    (fast-lookup-filter | no-fast-lookup-filter);
    firewall-install-disable;
    interface-group (Routing Options) group-name <exclude>;
    (per-route-accounting | no-per-route-accounting);
    route name
        match {
            destination prefix;
            destination-port (afs | bgp | biff | bootpc | bootps | cmd | cvspserver | dhcp |
domain | eklogin | ekshell | exec | expression | finger | ftp | ftp-data | http | https | ident
| imap | kerberos-sec | klogin | kpasswd | krb-prop | krbupdate | kshell | ldap | ldp | login |
mobileip-agent | mobilip-mn | msdp | netbios-dgm | netbios-ns | netbios-ssn | nfsd | nntp |
ntalk | ntp | pop3 | pptp | printer | radacct | radius | rip | rkinit | smtp | snmp | snmptrap |
snpp | socks | ssh | sunrpc | syslog | tacacs | tacacs-ds | talk | telnet | tftp | timed | who |
xdmcp | zephyr-clt | zephyr-hm | zephyr-srv)>;
            dscp [ dscp ... ];
            fragment (dont-fragment | first-fragment | is-fragment | last-fragment | not-a-
fragment);
            icmp-code (communication-prohibited-by-filtering | destination-host-prohibited |
destination-host-unknown | destination-network-prohibited | destination-network-unknown |
```

```
expression | fragmentation-needed | host-precedence-violation | host-unreachable | host-
unreachable-for-tos | ip-header-bad | network-unreachable | network-unreachable-for-tos | port-
unreachable | precedence-cutoff-in-effect | protocol-unreachable | redirect-for-host | redirect-
for-network | redirect-for-tos-and-host | redirect-for-tos-and-net | required-option-missing |
source-host-isolated | source-route-failed | ttl-eq-zero-during-reassembly | ttl-eq-zero-during-
transit);;
            icmp-type (echo-reply | echo-request | expression | info-reply | info-request | mask-
reply | mask-request | parameter-problem | redirect | router-advertisement | router-solicit |
source-quench | time-exceeded | timestamp | timestamp-reply | unreachable);;
            packet-length [ packet-length ... ];
            port (afs | bgp | biff | bootpc | bootps | cmd | cvspserver | dhcp | domain |
eklogin | ekshell | exec | expression | finger | ftp | ftp-data | http | https | ident | imap |
kerberos-sec | klogin | kpasswd | krb-prop | krbupdate | kshell | ldap | ldp | login | mobileip-
agent | mobilip-mn | msdp | netbios-dgm | netbios-ns | netbios-ssn | nfsd | nntp | ntalk | ntp |
pop3 | pptp | printer | radacct | radius | rip | rkinit | smtp | snmp | snmptrap | snpp | socks
| ssh | sunrpc | syslog | tacacs | tacacs-ds | talk | telnet | tftp | timed | who | xdmcp |
zephyr-clt | zephyr-hm | zephyr-srv);;
            protocol (ah | egp | esp | expression | gre | icmp | icmp6 | igmp | ipip | ospf |
pim | rsvp | sctp | tcp | udp);
            source prefix;
            source-port (afs | bgp | biff | bootpc | bootps | cmd | cvspserver | dhcp | domain |
eklogin | ekshell | exec | expression | finger | ftp | ftp-data | http | https | ident | imap |
kerberos-sec | klogin | kpasswd | krb-prop | krbupdate | kshell | ldap | ldp | login | mobileip-
agent | mobilip-mn | msdp | netbios-dgm | netbios-ns | netbios-ssn | nfsd | nntp | ntalk | ntp |
pop3 | pptp | printer | radacct | radius | rip | rkinit | smtp | snmp | snmptrap | snpp | socks
| ssh | sunrpc | syslog | tacacs | tacacs-ds | talk | telnet | tftp | timed | who | xdmcp |
zephyr-clt | zephyr-hm | zephyr-srv);
            tcp-flags (ack | expression | fin | push | rst | syn | urgent);
        }
        no-install;
        then (accept | discard | rate-limit rate-limit)  {
            community name;
            mark mark;
        next-term;
            redirect [ redirect ... ];
            routing-instance routing-instance;
            sample;
        }
    }
    term-order (legacy | standard);
    validation {
        traceoptions {
            file filename <files files> <size size> <(world-readable | no-world-readable)>;
```

```
            flag (all | flash | general | normal | policy | resolution | route | state | task |
timer)<detail> <disable><receive> <send>;
            }
    }
}
```

Hierarchy Level

```
[edit routing-options],
[edit routing-instances routing-instance-name routing-options]
```

Description

Configure a flow route.

Default

`legacy`

Options

`discard-action-for-unresolved-redir-addr` For action redirect IP if redirect address is unresolved install discard action filter.

fast-lookup-filter Enable fast lookup filter.

> **NOTE**: The option fast-lookup-filter is supported on MX Series devices only.

no-fast-lookup-filter Disable fast lookup filter.

> **NOTE**: The option no-fast-lookup-filter is supported on MX Series devices only.

Default: no-fast-lookup-filter

`firewall-install-disable`—(PTX The `firewall-install-disable` statement appears in the default configuration, preventing installation of flow-specification firewall filters into dfwd. For other

Series routers only)

models, this setting is omitted from the default configuration, allowing installation of flow-specification firewall filters into dfwd.

interface-group *group*

Interface-group for applying flowspec filter.

- *group*—Interface group ID to match traffic (1..255).
- exclude—Don't apply flowspec filter to traffic on this group.

per-route-accounting

Enable traffic accounting per flowspec route.

no-per-route-accounting

Disable traffic accounting per flowspec route.

route *name*

Name of the flow route.

match *match-conditions*

Match packets to these conditions.

destination *prefix*

IP destination address field.

destination-port *destination-port-names*

TCP or User Datagram Protocol (UDP) destination port field. You cannot specify both the port and destination-port match conditions in the same term.

In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), xdmcp (177), zephyr-clt (2103), or zephyr-hm (2104).

dscp *value*

Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP.

- **Range:** You can specify DSCP in hexadecimal or decimal form from 0 through 63.

fragment *fragment-value*

IP header Fragment field.

The keywords are grouped by the fragment type with which they are associated:

- first-fragment
- is-fragment

- last-fragment
- not-a-fragment

icmp-code ***icmp-code-value***

ICMP code field. This value or keyword provides more specific information than icmp-type. Because the value's meaning depends on the associated icmp-type value, you must specify icmp-type along with icmp-code.

In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:

- parameter-problem: ip-header-bad (0), required-option-missing (1)
- redirect: redirect-for-host (1), redirect-for-network (0), redirect-for-tos-and-host (3), redirect-for-tos-and-net (2)
- time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0)
- unreachable: communication-prohibited-by-filtering (13), destination-host-prohibited (10), destination-host-unknown (7), destination-network-prohibited (9), destination-network-unknown (6), fragmentation-needed (4), host-precedence-violation (14), host-unreachable (1), host-unreachable-for-TOS (12), network-unreachable (0), network-unreachable-for-TOS (11), port-unreachable (3), precedence-cutoff-in-effect (15), protocol-unreachable (2), source-host-isolated (8), source-route-failed (5)

icmp-type ***icmp-type-value***

ICMP packet type field. Normally, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port.

In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): echo-reply (0), echo-request (8), info-reply (16), info-request (15), mask-request (17), mask-reply (18), parameter-problem (12), redirect (5), router-advertisement (9), router-solicit (10), source-quench (4), time-exceeded (11), timestamp (13), timestamp-reply (14), or unreachable (3).

packet-length ***packet-length***

Total IP packet length value can range from 0 through 65535.

port ***port-names***

TCP or UDP source or destination port field. You cannot specify both the port match condition and either the destination-port or source-port match condition in the same term.

In place of the numeric value, you can specify one of the text synonyms listed under destination-port.

protocol *number*	In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): `ah` (51), `egp` (8), `esp` (50), `gre` (47), `icmp` (1), `igmp` (2), `ipip` (4), (41), `ospf` (89), `pim` (103), `rsvp` (46), `tcp` (6), or `udp` (17).
source *prefix*	IP source address field.
source-port *source-port-names*	TCP or UDP source port field. You cannot specify the `port` and `source-port` match conditions in the same term. In place of the numeric field, you can specify one of the text synonyms listed under `destination-port`.
tcp-flags *tcp-flags*	TCP header format.
no-install	Prohibit installing received routes in the forwarding table.
then	Actions to take on matching packets.
accept	Accept a packet. This is the default value.
community *name*	Replace any communities in the route with the specified communities.
discard	Discard a packet silently, without sending an Internet Control Message Protocol (ICMP) message.
mark *value*	Set a DSCP value for traffic that matches this flow. Specify a value from 0 through 63. This action is supported only on Junos devices with MPCs that are configured for `enhanced-ip` mode.
redirect	Redirect (tunnel) this flow's traffic to given next-hop address.
next-term	Continue to the next match condition for evaluation.
rate-limit *rate-limit*	Limit the bandwidth on the flow route. Express the limit in bits per second (bps).
routing-instance *route-target-extended-community*	Specify a routing instance to which packets are forwarded.
sample	Sample the traffic on the flow route.
term-order (legacy | standard)	• `legacy`—Use version 6 of the flow-specification algorithm.

- `standard`—Use version 7 of the flow-specification algorithm.

Default: legacy

validation

Flow route validation options. Flow specifications received from peers are validated such that it is considered feasible.

traceoptions

Define tracing operations that track all routing protocol functionality in the routing device.

file

Trace file options.

- *`filename`*—Name of file in which to write trace information.
- `files`—Maximum number of trace files (2..1000).
- `no-world-readable`— Don't allow any user to read the log file.
- `size`—Maximum trace file size (10240..4294967295).
- `world-readable`—Allow any user to read the log file.

flag

Tracing parameters.

- `all`—Trace everything.
- `flash`—Trace flash processing.
- `general`—Trace general events.
- `normal`—Trace normal events.
- `policy`—Trace policy processing
- `resolution`—Trace flow to unicast. route resolution
- `route`—Trace routing information.
- `state`—Trace state transitions.
- `task`—Trace routing protocol task processing.
- `timer`—Trace routing protocol timer processing.
  - `detail`—Trace detailed information.

- `disable`—Disable this trace flag.
- `receive`—Trace received packets.
- `send`—Trace transmitted packets.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`term-order` option introduced in Junos OS Release 10.0

`firewall-install-disable` option introduced in Junos OS Releases 12.1X48 and 12.3 for PTX Series routers.

`interface-group` option introduced in Junos OS Release 16.1R1.

`per-route-accounting` option introduced in Junos OS Release 21.1R1

RELATED DOCUMENTATION

Example: Enabling BGP to Carry Flow-Specification Routes

Understanding BGP Flow Routes for Traffic Filtering

flow-action

IN THIS SECTION

- Syntax | 7260
- Hierarchy Level | 7260
- Description | 7260
- Options | 7260

Required Privilege Level | 7260
Release Information | 7260

Syntax

```
flow-action  (drop | forward)
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber default-local-policy]
```

Description

Specify the action to take on a subscriber's data packets entering the access interface of the TDF domain when a TDF IP-based subscriber session does not exist.

Options

drop Drop the subscriber's packets.

forward Forward the subscriber's packets.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

flow-active-timeout (Flow Monitoring)

IN THIS SECTION

- Syntax | 7261
- Hierarchy Level | 7261
- Description | 7262
- Options | 7262
- Required Privilege Level | 7262
- Release Information | 7262

Syntax

```
flow-active-timeout seconds;
```

Hierarchy Level

```
[edit forwarding-options accounting name output],
[edit forwarding-options monitoring name output],
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls | vpls) output],
[edit forwarding-options sampling family (inet |inet6 |mpls |vpls) output],
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name],
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Set the interval after which an active flow is exported.

NOTE:

- The router must include an Adaptive Services, Multiservices, or Monitoring Services PIC for this statement to take effect.
- For Trio-based routers, if you configure the `inline-jflow` statement at the `[edit forwarding-options sampling instance` *`instance-name`* `family inet output]` hierarchy level, timed-out flows are exported even if the router does not contain a Multiservices card.

Options

seconds—Duration of the timeout period.

- **Range:** 60 through 1800 seconds (for `forwarding-options` configurations); 10 through 600 seconds (for `services` configurations)
- **Default:** 1800 seconds (for `forwarding-options` configurations); 60 seconds (for `services` configurations)

NOTE: In active flow monitoring, the cflowd or flow monitoring version 9 records are exported after a time period that is a multiple of 60 seconds and greater than or equal to the configured active timeout value. For example, if the active timeout value is 90 seconds, the cflowd or flow monitoring version 9 records are exported at 120-second intervals. If the active timeout value is 150 seconds, the cflowd or flow monitoring version 9 records are exported at 180-second intervals, and so forth.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit services flow-monitoring version-ipfix template` *`template-name`*`]` hierarchy level added in Junos OS Release 10.2.

Support at the [edit services flow-monitoring version9 template *template-name*] hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.

RELATED DOCUMENTATION

- *Configuring Time Periods when Flow Monitoring is Active and Inactive*
- *Configuring the Version 9 Template Properties*
- *Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices*
- *Configuring Discard Accounting*
- *Configuring Passive Flow Monitoring*
- *Collecting Traffic Sampling Output in a File*

flow-collector (Flow Collection)

IN THIS SECTION

- Syntax | 7263
- Hierarchy Level | 7264
- Description | 7264
- Required Privilege Level | 7265
- Release Information | 7265

Syntax

```
flow-collector {
    analyzer-address address;
    analyzer-id name;
    destinations {
        ftp:url {
            password "password";
        }
    }
```

```
    file-specification {
        variant variant-number {
            data-format format;
            name-format format;
            transfer {
                record-level number;
                timeout seconds;
            }
        }
    }
    interface-map {
        collector interface-name;
        file-specification variant-number;
        interface-name {
            collector interface-name;
            file-specification variant-number;
        }
    }
    retry number;
    retry-delay seconds;
    transfer-log-archive {
        archive-sites {
            ftp:url {
                password "password";
                username username;
            }
        }
        filename-prefix prefix;
        maximum-age minutes;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Define the flow collection.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Flow Collection Overview

flow-control (ACX Series)

IN THIS SECTION

- Syntax | 7265
- Hierarchy Level | 7266
- Description | 7266
- Default | 7266
- Required Privilege Level | 7266
- Release Information | 7266

Syntax

```
(flow-control | no-flow-control);
```

Hierarchy Level

```
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name ether-options],
[edit interfaces interface-name fastether-options],
[edit interfaces interface-name gigether-options],
[edit interfaces interface-name multiservice-options],
[edit interfaces interface-range name aggregated-ether-options],
[edit interfaces interface-range name ether-options]
```

Description

For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router or switch to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.

NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.

Default

Flow control is enabled.

NOTE: Flow control is enabled by default only on physical interfaces and it is disabled by default on aggregated Ethernet interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

Configuring Flow Control

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

flow-control

IN THIS SECTION

- Syntax | 7267
- Hierarchy Level | 7267
- Description | 7267
- Default | 7268
- Required Privilege Level | 7268
- Release Information | 7268

Syntax

```
(flow-control | no-flow-control);
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
```

Description

Explicitly enable or disable symmetric Ethernet PAUSE flow control, which regulates the flow of packets from the switch to the remote side of the connection by pausing all traffic flows on a link during periods of network congestion. Symmetric flow control means that Ethernet PAUSE is enabled in both directions. The interface generates and sends Ethernet PAUSE messages when the receive buffers fill to a certain threshold and the interface responds to PAUSE messages received from the connected peer.

You can configure asymmetric flow control by including the `configured-flow-control` statement at the `[edit interfaces interface-name ether-options` hierarchy level. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.

> **NOTE**: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual "lanes" of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).
>
> Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

- `flow-control`—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.
- `no-flow-control`—Disable flow control.

Default

Whether flow control is enabled or disabled by default depends on the platform. Run `show interfaces interface-name` to see whether flow control is enabled on an interface.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

configured-flow-control | **4523**

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Junos OS Network Interfaces Library for Routing Devices

flow-control (Interfaces)

IN THIS SECTION

- Syntax | 7269
- Hierarchy Level | 7269
- Description | 7269
- Default | 7269
- Required Privilege Level | 7270
- Release Information | 7270

Syntax

```
(flow-control | no-flow-control);
```

Hierarchy Level

```
[edit interfaces interface-name fastether-options]
[edit interfaces interface-name gigether-options]
[edit interfaces interface-name redundant-ether-options]
```

Description

For Fast Ethernet, Gigabit Ethernet, and redundant Ethernet interfaces, flow control regulates the flow of packets from the device to the remote side of the connection.

Default

Flow control is the default behavior for Fast Ethernet and Gigabit Ethernet interfaces. Flow control is disabled by default for redundant Ethernet interfaces

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

flow-control-options

IN THIS SECTION

- Syntax | 7270
- Hierarchy Level | 7271
- Description | 7271
- Usage Guidelines | 7272
- Required Privilege Level | 7272
- Release Information | 7272

Syntax

```
flow-control-options {
    down-on-flow-control;
    dump-on-flow-control;
    reset-on-flow-control;
    up-on-flow-control;
}
```

Hierarchy Level

```
[edit interfaces mo-fpc/pic/port multiservice-options]
```

Description

Configure the flow control options for application recovery in case of a prolonged flow control failure.

- `down-on-flow-control`—Bring interface down during prolonged flow control.
- `dump-on-flow-control`—Cause core dump during prolonged flow control.

> **NOTE**: Starting with Junos OS Release 15.1, on MX Series routers with MS-MICs and MS-MPCs, instead of an eJunos kernel core file, the multiservices PIC management daemon (mspmand) core file is generated when a prolonged flow control failure occurs and when you configure the setting to generate a core dump during prolonged flow control (by using the `dump-on-flow-control` option with the `flow-control-options` statement). The watchdog functionality continues to generate a kernel core file in such scenarios.

- `reset-on-flow-control`—Reset interface during prolonged flow control.

> **NOTE**: Starting in Junos OS Release 16.1R7, the `reset-on-flow-control` option has no effect on the MS-MIC, MS-MPC, MS-DPC, MS-PIC 100, MS-PIC 400, and MS-PIC 500 line cards. This is because starting in Release 16.1R7, Junos OS restarts these line cards to recover them from stuck state due to prolonged flow control.

- `up-on-flow-control`—Cause interface to remain in stuck state until you manually restart the PICs.

> **NOTE**: Starting in Junos OS Release 16.1R7, if interfaces on an MS-PIC or MS-DPC are in stuck state because of prolonged flow control, Junos OS restarts the service PICs to recover them from this state. However, if you want the PICs to remain in stuck state until you manually restart the PICs, configure the `up-on-flow-control` option. In releases before Release 16.1R7, there is no action taken to recover service PICs from this state unless one of the options for the `flow-control-options` statement is configured, or service PIC is manually restarted.

Usage Guidelines

See *Configuring Flow Monitoring*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 8.4.

flow-control-queue (Output Congestion Notification)

IN THIS SECTION

- Syntax | 7272
- Hierarchy Level | 7273
- Description | 7273
- Default | 7273
- Options | 7273
- Required Privilege Level | 7273
- Release Information | 7274

Syntax

```
flow-control-queue [ queue | list-of-queues ];
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name output ieee-802.1 code-point code-point-bits]
```

Description

Specify one or more output queues to pause, to support priority-based flow control (PFC). The specified queues pause when the interface receives a PFC frame with a matching IEEE 802.1p code point.

Default

Queue 3 (mapped to the fcoe forwarding class) and queue 4 (mapped to the no-loss forwarding class) are programmed as flow control queues to pause. No other output queues are programmed to pause by default.

If you configure flow control queues explicitly, only the queues that you specify are programmed to pause. The explicit flow control queue to pause configuration overrides the default setting, so the queues paused in the default configuration are no longer paused by default.

For example, if you configure queue 2 as a flow control queue, then queue 2 pauses when congestion occurs, but queues 3 and 4 do not pause because they were not explicitly specified. To enable pause on output queues 2, 3, and 4, you must explicitly configure all three of the queues as flow control queues.

The same behavior applies to the IEEE 802.1p code points (priorities) on which PFC is enabled. By default, priorities 3 (011) and 4 (100) are enabled for PFC pause. If you explicitly configure flow control queues to pause, you must also explicitly configure pause for each priority (code point) that you want to pause, because the explicit configuration overrides the default configuration.

Options

[*queue* | *list-of-queues*]—The output queue or a list of output queues to pause.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring CoS PFC (Congestion Notification Profiles)

Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces

Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface

Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)

Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)

Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows

flow-descriptions

IN THIS SECTION

- Syntax | 7274
- Hierarchy Level | 7275
- Description | 7275
- Options | 7275
- Required Privilege Level | 7276
- Release Information | 7276

Syntax

```
flow-descriptions flow-identifier {
    direction (uplink | downlink | both);
    local-port-range {
```

```
        low lower-boundary high upper-boundary;
    }
    local-ports number;
    no-send-to-ue;
    protocol protocol-number;
    remote-address (ipv4-address ipv4-address | ipv6-address ipv6-address);
    remote-port-range {
        low lower-boundary high upper-boundary;
    }
    remote-ports number;
}
```

Hierarchy Level

```
[edit unified-edge pcef],
[edit services pcef]
```

Description

Specify a service data flow (SDF) filter (flow identifier) that includes one or more filtering parameters (address, protocol, and port) to identify the subscriber traffic that you want the SDF filter to detect. SDF filters are specified in a PCC rule to identify the Layer 3 or Layer 4 IP packet flows that you want to receive a particular treatment.

> **NOTE**: A PCC rule must include at least one SDF filter and can include a maximum of 15 SDF filters.

If you are using Junos OS Subscriber Aware, specify the name of the SDF filter at the `[edit unified-edge pcef]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the SDF filter at the `[edit services pcef]` hierarchy level.

Options

flow-identifier

Name of the SDF filter.

- **Range:** 1 through 63 characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

Configuring Service Data Flow Filters

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Understanding Application-Aware Policy Control for Subscriber Management

flow-detect-time (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 7277
- Hierarchy Level | 7277
- Description | 7277
- Options | 7277

Required Privilege Level | 7277
Release Information | 7278

Syntax

```
flow-detect-time seconds;
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type flow-detection]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure how much time must pass before a suspicious flow that has exceeded the bandwidth allowed for the packet type is confirmed to be a culprit flow.

BEST PRACTICE: We recommend that you use the default value for the detection period.

Options

seconds—Period of excessive bandwidth required for flow to be a culprit flow.

- **Range:** 1 through 60 seconds
- **Default:** 3 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring the Detection Period for Suspicious Flows

Configuring Flow Detection for Control Plane DDoS Protection

flow-detection (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 7278
- Hierarchy Level | 7278
- Description | 7279
- Required Privilege Level | 7279
- Release Information | 7279

Syntax

```
flow-detection;
```

Hierarchy Level

```
[edit system ddos-protection global]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Enable flow detection globally for all protocol groups and packet types except the following, which do not have typical Ethernet, IP, or IPv6 headers:

- Protocol groups: `fab-probe`, `frame-relay`, `inline-ka`, `isis`, `jfm`, `mlp`, `pfe-alive`, `pos`, and `services`.
- Packet type: `unclassified` in the `ip-options` protocol group.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3. Statement is supported on PTX10003 and PTX10008. Use Feature Explorer to confirm platform support for this feature.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Enabling Flow Detection for All Protocol Groups and Packet Types

Setting Up and Using Flow Detection

flow-detection (DDoS Packet Level)

IN THIS SECTION

- Syntax | 7280
- Hierarchy Level | 7280
- Description | 7280
- Required Privilege Level | 7281
- Release Information | 7281

Syntax

```
flow-detection {
    flow-detect-time detect-period;
    no-flow-logging;
    timeout-active-flows enable-period;
    flow-level-bandwidth {
        logical-interface flow-bandwidth;
        physical-interface flow-bandwidth;
        subscriber flow-bandwidth;
    }
    flow-level-control {
        logical-interface flow-control-mode;
        physical-interface flow-control-mode;
        subscriber flow-control-mode;
    }
    flow-level-detection {
        logical-interface operation-mode;
        physical-interface operation-mode;
        subscriber operation-mode;
    }
    flow-detection-mode  (automatic |off | on);
    flow-recover-time recover-period;
    flow-timeout-time timeout-period;
}
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure control plane DDoS protection suspicious control flow detection for a packet type.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Setting Up and Using Flow Detection

flow-detection-mode (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 7281
- Hierarchy Level | 7282
- Description | 7282
- Default | 7282
- Options | 7282
- Required Privilege Level | 7282
- Release Information | 7282

Syntax

```
flow-detection-mode (automatic | off | on)
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure the mode of operation for flow detection for a protocol group or packet type. Use this statement to override global flow detection settings configured with the `flow-detection-mode` statement at the `[edit system ddos-protection global]` hierarchy level. The operation mode is effective only when flow detection is enabled.

Default

The default mode for all protocol groups and packet types is `automatic`.

Options

automatic Detect flows only when the policer is being violated.

off Disable flow detection.

on Always monitor and detect flows, even when the policer is not being violated.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring How Flow Detection Operates for Individual Protocol Groups or Packets

flow-detection-mode (DDoS Global Flow Detection)

IN THIS SECTION

- Syntax | 7283
- Hierarchy Level | 7283
- Description | 7283
- Default | 7284
- Options | 7284
- Required Privilege Level | 7284
- Release Information | 7284

Syntax

```
flow-detection-mode (automatic | off | on)
```

Hierarchy Level

```
[edit system ddos-protection global]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure the mode of operation for flow detection globally for almost all protocol groups and packet types. The operation mode is effective only when flow detection is enabled.

NOTE: You cannot enable flow detection globally for the following groups and packet type because they do not have typical Ethernet, IP, or IPv6 headers:

- Protocol groups: `fab-probe`, `frame-relay`, `inline-ka`, `isis`, `jfm`, `mlp`, `pfe-alive`, `pos`, and `services`.
- Packet type: `unclassified` in the `ip-options` protocol group.

To override the global configuration for a protocol group or packet type, use the `flow-detection-mode` statement at the `[edit system ddos-protection protocols` *`protocol-group packet-type`*`]` hierarchy level.

Default

The default global mode is `automatic`.

Options

automatic Detect flows only when the policer is being violated.

off Disable flow detection.

on Always monitor and detect flows, even when the policer is not being violated.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring How Flow Detection Operates for Individual Protocol Groups or Packets

Setting Up and Using Flow Detection

flow-export-destination (Flow Collection)

IN THIS SECTION

- Syntax | 7285
- Hierarchy Level | 7285
- Description | 7285
- Options | 7285
- Required Privilege Level | 7285
- Release Information | 7286

Syntax

```
flow-export-destination {
    (cflowd-collector | collector-pic);
}
```

Hierarchy Level

```
[edit forwarding-options monitoring group-name family inet output]
```

Description

Configure flow collection.

Options

`cflowd-collector`—Use the cflowd collector.

`collector-pic`—Use the collector PIC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Exporting Flows

Configuring Passive Flow Monitoring

flow-export-rate

IN THIS SECTION

- Syntax | 7286
- Hierarchy Level | 7286
- Description | 7287
- Options | 7287
- Required Privilege Level | 7287
- Release Information | 7287

Syntax

```
flow-export-rate rate;
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family inet output inline-jflow]
```

Description

Specify the flow export rate of monitored packets in kpps. If you have multiple line cards of different types running on the same router, the `flow-export-rate` will be applied to each card. However, the rate applied to the PFEs on the card will vary in accordance with the number of PFEs that are on the card.

On the MX Series, the actual flow export rate might be less than the configured `flow-export-rate`. In addition, the maximum flow export rate for an FPC or MPC module cannot exceed the configured `flow-export-rate`, regardless of how many PICs or MICs are on the module.

Options

rate Flow export rate of monitored packets in kpps (from 1 through 400).

- **Default:** 1 kpps (applies to all PFEs on the FPC)

> **NOTE**: The maximum rate per PFE is 100 kpps for LU, 800 kpps for XL/EA, so for an FPC with four LU PFEs (such as AS cards) you can set a maximum `flow-export-rate` of 400. For an FPC with two LU PFEs (such as the MPC2), the maximum `flow-export-rate` is 200. For an FPC with one LU PFE (such as the MPC5), the maximum is 100. The Junos CLI accepts as valid any value within the range of 1 to 3200, but when applied the value might trigger an error message such as “The configured flow export rate is higher than supported value/chip” in the Junos message log.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Discard Accounting

Configuring Flow Monitoring

Configuring Traffic Sampling on MX, M and T Series Routers

flow-export-timer

IN THIS SECTION

- Syntax | 7288
- Hierarchy Level | 7288
- Description | 7288
- Default | 7288
- Options | 7288
- Required Privilege Level | 7289

Syntax

```
flow-export-timer seconds
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

Configure the flow export period in seconds. The software exports a IPFIX packet periodically at the configured flow-export timer interval. After you configure the flow export period, you must reboot the device for it to take effect.

Default

10 seconds

Options

seconds	Range: 10 to 600 seconds.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

flow-gate

IN THIS SECTION

- Syntax | 7289
- Hierarchy Level | 7289
- Description | 7289
- Options | 7290
- Required Privilege Level | 7290
- Release Information | 7290

Syntax

```
flow-gate {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specifies the number of flow gates, also known as `pinholes`, that user logical system administrators and primary logical system administrators can configure for their logical systems, if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

`maximum` *`amount`*
: A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

  - **Range:** 0 through 1000000

`reserved` *`amount`*
: A reserved quota that guarantees that the resource amount specified is always available to the logical system.

  - **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only)

flow-groups

IN THIS SECTION

- Syntax | 7291
- Hierarchy Level | 7291
- Description | 7291
- Options | 7295
- Required Privilege Level | 7296
- Release Information | 7296

Syntax

```
flow-groups {
    flow-group-name {
        node-device (node-device-name | serial-ID | alias);
        interconnect-device interconnect-device-name;
    }
    node-device (node-device-name | serial-ID | alias) {
        interconnect-device interconnect-device- name {
            preference (high | normal | never);
        }
    }
}
```

Hierarchy Level

```
[edit fabric]
```

Description

In a QFabric system, a hash function is used to select an Interconnect device to forward traffic between two ingress switches, which are called Node devices. Since this hash function is performed on all Interconnect devices, it is possible for redundant multicast streams to flow through one Interconnect

device, making that Interconnect device a potential single point of failure for the redundant flows. Some operations, such as financial transactions, require redundant data flows to use different Interconnect devices to avoid point-of-failure data loss. To enforce this in a fabric, you can create a flow group for multicast or unicast traffic from a particular ingress switch. In a flow group, switches are forced to use different, specified Interconnect devices to forward redundant traffic.

NOTE: The following rules apply to QFabric system flow control:

- A Node device can belong to only one flow group.
- A fabric can contain up to four flow groups.
- Node devices in flow groups can only use the defined Interconnect devices in their flow groups unless all defined Interconnect devices fail. In case of failure, the flow group checks for Interconnect devices with preference set to `normal`—this includes all Interconnect devices in the default group and any Interconnect devices you specifically set to preference `normal`.

The following example creates a flow group. The name of the flow group, test1, is configured first. Next, the Node device TA3713480203 is configured, and the Interconnect device IC-P6603-C is configured last.

```
[edit fabric]
root@qfabric#  set flow-groups ?
Possible completions:
<flow group name>                Flow Group name
+ apply-groups                   Groups from which to inherit configuration data
+ apply-groups-except            Don't inherit configuration data from these groups
> node-device                    Override the link preference for Node device(s)

root@qfabric#  set flow-groups test1 ?
Possible completions:
+ apply-groups                    Groups from which to inherit configuration data
+ apply-groups-except             Don't inherit configuration data from these groups
> interconnect-device             Interconnect device(s) associated with this flow group
> node-device                     Node device(s) associated with this flow group

root@qfabric#  set flow-groups test1 node-device ?

Possible completions:
<node-device-name> Serial ID or alias of node-device
P3686-C                     Node device
TA3713480203                Node device
```

```
TA3713500165                Node device
TA3713500185                Node device
TA3713500307                Node device

root@qfabric# set flow-groups test1 node-device TA3713480203 ?
Possible completions:

<Enter>                            Execute this command
+ apply-group                   Groups from which to inherit configuration data
+ apply-groups-except           Don't inherit configuration data from these groups
|                               Pipe through a command

root@qfabric# set fabric flow-groups test1 node-device TA3713480203

root@qfabric# set flow-groups test1 interconnect-device ?
Possible completions:
<interconnect-device-name>  Serial ID or alias of interconnect-device
IC-P6601-C                  Interconnect device
IC-P6603-C                  Interconnect device

root@qfabric# set flow-groups test1 interconnect-device IC-P6603-C
root@qfabric# show flow-groups
test1 {
    node-device {
        TA3713480203;
    }
    interconnect-device {
        IC-P6603-C;
    }
}
```

The following example overrides the software settings the flow group test1 Node device TA3713500307 was assigned by the software for Interconnect device IC-P6601-C. Node device TA3713500307 has its preference for Interconnect device IC-P6601-C set to `never` by the software, meaning that Node device TA3713500307 can never use Interconnect device IC-P6601-C. You can change this preference so that Interconnect device IC-P6601-C can be used by Node device TA3713500307 if the specifically-configured Interconnect devices all fail. You use the preference option of the `flow-groups` configuration statement to set Interconnect device IC-P6601-C preference to `normal` for Node device TA3713500307.

To configure the preference setting for a Node device to use a particular Interconnnect device, you first select the Node device, then specify the Interconnect device, and finally specify the preference option:

```
[edit fabric]
root@qfabric# set flow-groups ?

Possible completions:
<flow-group name>          Flow Group name
+ apply-groups             Groups from which to inherit configuration data
+ apply-groups-except      Don't inherit configuration data from these groups >
node-device               Override the link preference for Node device(s)
test1                      Flow Group name

root@qfabric# set flow-groups node-device ?

Possible completions:
<name>    serial ID or alias of node-device

P3686-C                     Node device
TA3713480203                Node device
TA3713500185                Node device
TA3713500307                Node device

root@qfabric# set fabric flow-groups node-device TA3713500307 ?
Possible completions:
+ apply-groups            Groups from which to inherit configuration data
+ apply-groups-except     Don't inherit configuration data from these groups
interconnect device

root@qfabric# set fabric flow-groups node-device TA3713500307 interconnect-device ?
Possible completions:
<name>                      Serial ID or alias of interconnect-device
 IC-P6601-C               Interconnect device
 IC-P6603-C               Interconnect device

root@qfabric# set fabric flow-groups node-device TA3713500307 interconnect-device IC-P6601-C
preference ?
Possible completions:
  high                       High link preference
  never                      Never use this link
  normal                     Normal link reference
```

```
root@qfabric# set fabric flow-groups node-device TA3713500307 interconnect-device IC-P6601-C
preference never

root@qfabric# show fabric flow-groups
test1 {
    node-device {
        TA3713480203;
    }
    interconnect-device {
        IC-P6603-C;
    }
}
node-device TA3713500307 {
    interconnect-device IC-P6601-C {
        preference never;
    }
}
```

TIP: To add a TOR to multiple flow groups:
set fabric resources node-device TOR-B interconnect-device IC-A preference HIGH
To configure a TOR to use an Interconnect device from a different flow group if none in the same flow group are available:
set fabric resources node-device TOR-B interconnect-device IC-A preference NORMAL
To increase the priority of default-flowgroup Interconnect devices for a specific TOR:
set fabric resources node-device TOR-A interconnect-device ALL_ICS preference HIGH
To configure default-flowgroup TORs to never use a particular Interconnect device:
Set fabric resources node-device ALL_TORS interconnect-device IC-A preference NEVER

Options

node-device *node-device-name*	Specify a QFabric system ingress Node switch by device name or alias in the QFabric system.
interconnect-device *interconnect-device-name*	Specify a QFabric system Interconnect device name.

preference *preference* Overrides the algorithm results that assign preference to Node devices to use particular Interconnect devices. You can set preference values for these devices to `normal`, `high`, or `never`. (See *Example: Creating a QFabric Flow Group* for more details on setting `preference` values and their meanings.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D15.

RELATED DOCUMENTATION

Example: Creating a QFabric Flow Group

Understanding QFabric Multicast Flow Groups

Segregating QFabric Traffic Flows With Flow Groups

flow-inactive-timeout (Flow Monitoring)

IN THIS SECTION

- Syntax | 7297
- Hierarchy Level | 7297
- Description | 7297
- Options | 7297
- Required Privilege Level | 7297
- Release Information | 7297

Syntax

```
flow-inactive-timeout seconds;
```

Hierarchy Level

```
[edit forwarding-options accounting name output],
[edit forwarding-options monitoring name output],
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls | vpls) output],
[edit forwarding-options sampling family (inet |inet6 |mpls) output],
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name],
```

Description

Configure the time that elapses before a flow is considered inactive.

> **NOTE**: The router must include an Adaptive Services, Multiservices, or Monitoring Services PIC for this statement to take effect.

Options

seconds—Duration of the timeout period.

- **Range:** 15 through 1800 seconds (for `forwarding-options` configurations); 10 through 600 seconds (for `services` configurations)
- **Default:** 60 seconds (for `forwarding-options` configurations); 60 seconds (for `services` configurations)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [edit services flow-monitoring version-ipfix template *template-name*] hierarchy level added in Junos OS Release 10.2.

Support at the [edit services flow-monitoring version9 template *template-name*] hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.

RELATED DOCUMENTATION

- *Configuring Time Periods when Flow Monitoring is Active and Inactive*
- *Configuring the Version 9 Template Properties*
- *Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices*
- *Configuring Discard Accounting*
- *Configuring Passive Flow Monitoring*
- *Collecting Traffic Sampling Output in a File*

flow (IPv6)

IN THIS SECTION

- Syntax | 7298
- Hierarchy Level | 7300
- Description | 7300
- Options | 7300
- Required Privilege Level | 7304
- Release Information | 7304

Syntax

```
flow {
    discard-action-for-unresolved-redir-addr;
    interface-group (Routing Options) group-name <exclude>;
    (per-route-accounting | no-per-route-accounting);
```

```
    route name {
        match {
            destination {
                ipv6-prefix;
                prefix-offset number;
            }
            destination-port destination-port-names;
            dscp value;
            flow-label numeric-expression;
            fragment fragment-value;
            icmp6-code icmp6-code-value;
            icmp6-type icmp6-type-value;
            packet-length packet-length;
            port port-names;
            protocol number;
            source {
                ipv6-prefix;
                prefix-offset number;
            }
            source-port source-port-names;
            tcp-flags tcp-flags;
            no-install;
     }
        then {
            accept;
            community name;
            discard;
            mark value;
            redirect IP address;
            next-term;
            rate-limit value in bits/sec;
            routing-instance route-target-extended-community;
            sample;
      validation {
         traceoptions
       (routing-options) {
             file filename <files number of files>
             flag(all | flash | general | normal | policy | resolution | route | state | task |
timer) {
                detail;
                   disable;
                receive;
                send;
```

```
                }
              }
            }
          }
      }
  }
```

Hierarchy Level

```
[edit routing-options rib inet6.0],
[edit routing-instances routing-instance-name routing-options rib inet6.0]
```

Description

Configure the BGP flow specification for the IPv6 address family to automate coordination of traffic filtering rules and to allow propagation of traffic flow specification rules for IPv6 and IPv6 VPN in order to mitigate distributed denial-of-service attacks. Flow specification provides protection against denial-of-service attacks and restricts bad traffic that consumes bandwidth and stops it near the source.

> **NOTE**: To propagate IPv6 flow specification routes through BGP, enable family `inet6 flow` or `inet6-vpn flow` at the `[edit protocols bgp family]` hierarchy level on BGP routers in the network.

Options

discard-action-for-unresolved-redir-addr	Configure the discard action for BGP flow specification routes that were not resolved using the redirect to IP action.
interface-group *group*<exclude>	Exclude applying flowspec filter to traffic received on specific interfaces. Use **exclude** to specify the interface group where you do not want to receive the traffic.
per-route-accounting	Enable traffic accounting per flowspec route.
no-per-route-accounting	Disable traffic accounting per flowspec route.
destination *ipv6-prefix*	IP destination address field.

destination-port *destination-port-names*

TCP or User Datagram Protocol (UDP) destination port field. You cannot specify both the `port` and `destination-port` match conditions in the same term.

In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): `afs` (1483), `bgp` (179), `biff` (512), `bootpc` (68), `bootps` (67), `cmd` (514), `cvspserver` (2401), `dhcp` (67), `domain` (53), `eklogin` (2105), `ekshell` (2106), `exec` (512), `finger` (79), `ftp` (21), `ftp-data` (20), `http` (80), `https` (443), `ident` (113), `imap` (143), `kerberos-sec` (88), `klogin` (543), `kpasswd` (761), `krb-prop` (754), `krbupdate` (760), `kshell` (544), `ldap` (389), `login` (513), `mobileip-agent` (434), `mobilip-mn` (435), `msdp` (639), `netbios-dgm` (138), `netbios-ns` (137), `netbios-ssn` (139), `nfsd` (2049), `nntp` (119), `ntalk` (518), `ntp` (123), `pop3` (110), `pptp` (1723), `printer` (515), `radacct` (1813), `radius` (1812), `rip` (520), `rkinit` (2108), `smtp` (25), `snmp` (161), `snmptrap` (162), `snpp` (444), `socks` (1080), `ssh` (22), `sunrpc` (111), `syslog` (514), `tacacs-ds` (65), `talk` (517), `telnet` (23), `tftp` (69), `timed` (525), `who` (513), `xdmcp` (177), `zephyr-clt` (2103), or `zephyr-hm` (2104).

dscp *value*

Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP.

- **Range:** You can specify DSCP in hexadecimal or decimal form from 0 through 63.

flow-label *numeric-expression*

The value of this field ranges from 0 through 1048575.

This match condition is supported only on Junos devices with enhanced MPCs that are configured for `enhanced-ip` mode.

fragment *fragment-value*

The keywords are grouped by the fragment type with which they are associated:

- `first-fragment`
- `is-fragment`
- `last-fragment`
- `not-a-fragment`

This match condition is supported only on Junos devices with enhanced MPCs that are configured for `enhanced-ip` mode.

icmp6-code *icmp6-code-value*

ICMP6 code field. This value or keyword provides more specific information than `icmp6-type`. Because the value's meaning depends on the associated `icmp6-type` value, you must specify `icmp6-type` along with `icmp6-code`.

In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:

- parameter-problem: `ip-header-bad` (0), `required-option-missing` (1)
- redirect: `redirect-for-host` (1), `redirect-for-network` (0), `redirect-for-tos-and-host` (3), `redirect-for-tos-and-net` (2)
- time-exceeded: `ttl-eq-zero-during-reassembly` (1), `ttl-eq-zero-during-transit` (0)
- unreachable: `communication-prohibited-by-filtering` (13), `destination-host-prohibited` (10), `destination-host-unknown` (7), `destination-network-prohibited` (9), `destination-network-unknown` (6), `fragmentation-needed` (4), `host-precedence-violation` (14), `host-unreachable` (1), `host-unreachable-for-TOS` (12), `network-unreachable` (0), `network-unreachable-for-TOS` (11), `port-unreachable` (3), `precedence-cutoff-in-effect` (15), `protocol-unreachable` (2), `source-host-isolated` (8), `source-route-failed` (5)

`icmp6-type` *`icmp6-type-value`*

ICMP6 packet type field. Normally, you specify this match in conjunction with the `protocol` match statement to determine which protocol is being used on the port.

In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): `echo-reply` (0), `echo-request` (8), `info-reply` (16), `info-request` (15), `mask-request` (17), `mask-reply` (18), `parameter-problem` (12), `redirect` (5), `router-advertisement` (9), `router-solicit` (10), `source-quench` (4), `time-exceeded` (11), `timestamp` (13), `timestamp-reply` (14), or `unreachable` (3).

`packet-length` *`packet-length`*

Total IP packet length value can range from 0 through 65535.

`port` *`port-names`*

TCP or UDP source or destination port field. You cannot specify both the `port` match condition and either the `destination-port` or `source-port` match condition in the same term.

In place of the numeric value, you can specify one of the text synonyms listed under `destination-port`.

prefix-offset *number*

(Optional) Specify the number of bits that must be skipped before Junos OS starts matching the prefix.

This match condition is supported only on Junos devices with enhanced MPCs that are configured for `enhanced-ip` mode.

`protocol` *`number`*

For IPv6, the IP protocol field is supported only on Junos devices with MPCs that are configured for `enhanced-ip` mode. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): `ah` (51), `egp` (8),

esp (50), gre (47), icmp (1), igmp (2), ipip (4), ipv6 (41), ospf (89), pim (103), rsvp (46), tcp (6), or udp (17).

source *ipv6-prefix* — IP source address field.

source-port *source-port-names* — TCP or UDP source port field. You cannot specify the port and source-port match conditions in the same term.

In place of the numeric field, you can specify one of the text synonyms listed under destination-port.

tcp-flags *tcp-flags* — TCP header format.

no-install — Prohibit installing received routes in the forwarding table.

accept — Accept a packet. This is the default value.

community *name* — Replace any communities in the route with the specified communities.

discard — Discard a packet silently, without sending an Internet Control Message Protocol (ICMP) message.

mark *value* — Set a DSCP value for traffic that matches this flow. Specify a value from 0 through 63.

This action is supported only on Junos devices with MPCs that are configured for enhanced-ip mode.

> **NOTE**: Junos OS supports traffic marking extended BGP community filtering action. For IPv4 traffic, Junos OS modifies the DiffServ code point (DSCP) bits of a transiting IPv4 packet to the corresponding value of the extended community. For IPv6 packets, Junos OS modifies the first six bits of the traffic class field of the transmitting IPv6 packet to the corresponding value of the extended community.

redirect — Redirect (tunnel) this flow's traffic to given next-hop address.

next-term — Continue to the next match condition for evaluation.

rate-limit *rate-limit* — Limit the bandwidth on the flow route. Express the limit in bits per second (bps).

routing-instance *route-target-extended-community* — Specify a routing instance to which packets are forwarded.

`sample`	Sample the traffic on the flow route.
`traceoptions`	Define tracing operations that track all routing protocol functionality in the routing device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Example: Configuring BGP to Carry IPv6 Flow Specification Routes

Example: Enabling BGP to Carry Flow-Specification Routes

Understanding BGP Flow Routes for Traffic Filtering

flow-key (Flow Monitoring)

IN THIS SECTION

- Syntax | 7305
- Hierarchy Level | 7305
- Description | 7305
- Options | 7305
- Required Privilege Level | 7306
- Release Information | 7306

Syntax

```
flow-key {
    flow-direction;
    vlan-id;
    output-interface;
}
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Include VLAN IDs in both the ingress and egress directions in the flow key, enable flow direction information in a Version 9 or IPFIX flow template, or both, and configure the output-interface for bridge or VPLS family for inline flow monitoring on the MX Series.

Options

flow-direction Enable reporting of the direction of the flow. The field contains 0x00 (ingress) or 0x01 (egress). The flow direction field in the output record contains the invalid value 0xFF if you do not configure `flow-direction`.

vlan-id Include VLAN IDs in both the ingress and egress directions in the flow key.

output-interface Configure the output-interface field as part of flow-key for bridge or VPLS family.

> **NOTE**: If the output-interface (OIF) is configured under flow-key while the flow-monitoring is in progress, all the existing flows (where OIF was not part of flow-key) report OIF field as zero in the next export. Therefore, in progress configuration of output-interface as part of flow-key is not recommended. In order to configure output-interface as part of flow-key, it is recommended to disable the bridge or vpls sampling and wait for the active flows to become zero.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.2.

`output-interface` option added in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates

Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices

flow-label

IN THIS SECTION

- Syntax | 7306
- Hierarchy Level | 7307
- Description | 7307
- Required Privilege Level | 7307
- Release Information | 7307

Syntax

```
flow-label;
```

Hierarchy Level

```
[edit routing-instances name protocols evpn]
```

Description

Advertise that the device is able to support flow-aware transport (FAT) labels for EVPN-MPLS E- LAN services. With FAT flow labels, you can load balance MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs) without the need for deep packet inspection of the payload. RFC 6391 describes FAT flow labels and how they are used.

> **NOTE**: Starting in Junos OS Evolved Release 24.4R1, we support `flow-label` configuration with egress link protection (ELP) on Junos OS Evolved platforms.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 22.2R1.

flow-label-receive

IN THIS SECTION

- Syntax | 7308
- Hierarchy Level | 7308
- Description | 7308
- Required Privilege Level | 7308

Release Information | 7309

Syntax

```
flow-label-receive;
```

Hierarchy Level

```
[edit protocols l2circuit neighbor neighbor-id interface interface-name],
[edit routing-instances routing-instance-name protocols l2vpn site name],
[edit routing-instances routing-instance-name protocols l2vpn site name interface interface-
name],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Configure the router to signal the capability to pop the flow label in the receive direction to the remote provider edge (PE) router.

The flow-aware transport of pseudowires (FAT) flow labels enable load-balancing of MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs) without the need for deep packet inspection of the payload. FAT flow labels can be used for LDP-signaled forwarding equivalence class (FEC 128 and FEC 129) pseudowires for virtual private LAN service (VPLS) and virtual private wire service (VPWS) networks.

NOTE:

- Starting with Junos OS Release 16.1, `flow-label-receive` statement is supported for BGP L2VPN and BGP VPLS.
- The `flow-label-receive` statement is not supported in OAM CFM session over Layer 2 circuit.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring the FAT Flow Label for FEC 128 VPWS Pseudowires for Load-Balancing MPLS Traffic

Configuring the FAT Flow Label for FEC 128 VPLS Pseudowires for Load-Balancing MPLS Traffic

Configuring the FAT Flow Label for FEC 129 VPWS Pseudowires for Load-Balancing MPLS Traffic

Configuring the FAT Flow Label for FEC 129 VPLS Pseudowires for Load-Balancing MPLS Traffic

flow-label-receive-static

IN THIS SECTION

- Syntax | 7309
- Hierarchy Level | 7310
- Description | 7310
- Required Privilege Level | 7310
- Release Information | 7311

Syntax

```
flow-label-receive-static;
```

Hierarchy Level

```
[edit protocols l2circuit neighbor neighbor-id interface interface-name]
[edit routing-instances instance-name protocols evpn],
[edit routing-instances instance-name protocols evpn interface interface-name]
```

Description

Configure the router to pop the flow label on pseudowire packets received from a remote provider edge (PE) router. The ingress PE router inserts the flow label in the pseudowire packet regardless of the information exchanged in the signaling plane. If the egress PE router cannot handle the pseudowire packet with the flow label, it drops the packet.

Flow-aware transport of pseudowires (FAT) flow labels enable load balancing of MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs) without the need for deep packet inspection of the payload. You can use this statement to enable devices to pop FAT flow labels for:

- Forwarding equivalence class (FEC) 128 pseudowires in:
  - Virtual private LAN service (VPLS) networks.
  - Virtual private wire service (VPWS) networks.

  In these cases, configure this statement in the `[edit protocols l2circuit neighbor neighbor-id interface interface-name] hierarchy`.
- Pseudowires in Ethernet VPN (EVPN)-MPLS networks with VPWS (EVPN-VPWS):
  - Globally in an EVPN-VPWS routing instance.

    In this case, configure this statement in the `[edit routing-instances instance-name protocols evpn]` hierarchy.
  - For specific interfaces in an EVPN-VPWS routing instance.

    In this case, configure this statement in the `[edit routing-instances instance-name protocols evpn interface interface-name]` hierarchy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Support in a EVPN-VPWS routing instance added in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Configuring the FAT Flow Label for FEC 128 VPWS Pseudowires for Load-Balancing MPLS Traffic

FAT Flow Labels in EVPN-VPWS Routing Instances

flow-label-static

IN THIS SECTION

- Syntax | 7311
- Hierarchy Level | 7311
- Description | 7312
- Required Privilege Level | 7312
- Release Information | 7312

Syntax

```
flow-label-static;
```

Hierarchy Level

```
[edit routing-instances name protocols evpn]
```

Description

For EVPN-MPLS E-LAN services, configure the router to signal the ability to push the flow-aware transport (FAT) flow label on pseudowire packets received from a provider edge (PE) router or to pop the FAT flow label on pseudowire packets transmitted to a PE router. The ingress PE router inserts the flow label in the pseudowire packet regardless of the information exchanged in the signaling plane. If the egress PE router cannot handle the pseudowire packet with the flow label, it drops the packet. A change in FAT flow label configuration is a catastrophic event, causing traffic loss. RFC 6391 describes FAT flow labels and how they are used.

With FAT flow labels, you can load balance known unicast MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs) without the need for deep packet inspection of the payload. You can use this statement to enable devices to pop or push FAT flow labels for these EVPN service types:

- `vlan-based`
- `vlan-bundle`
- `vlan-aware`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 22.2R1.

flow-label-transmit

IN THIS SECTION

- Syntax | 7313
- Hierarchy Level | 7313
- Description | 7313
- Required Privilege Level | 7313

Release Information | 7314

Syntax

```
flow-label-transmit;
```

Hierarchy Level

```
[edit protocols l2circuit neighbor neighbor-id interface interface-name],
[edit routing-instances routing-instance-name protocols l2vpn site name],
[edit routing-instances routing-instance-name protocols l2vpn site name interface interface-
name],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Configure the router to signal the capability to push the flow label in the transmit direction to the provider edge (PE) router.

The flow-aware transport of pseudowires (FAT) flow labels enable load-balancing of MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs) without the need for deep packet inspection of the payload. FAT flow labels can be used for LDP-signaled forwarding equivalence class (FEC 128 and FEC 129) pseudowires for virtual private LAN service (VPLS) and virtual private wire service (VPWS) networks.

NOTE:

- Starting with Junos OS Release 16.1, `flow-label-transmit` statement is supported for BGP L2VPN and BGP VPLS.
- The `flow-label-transmit` statement is not supported in OAM CFM session over L2Circuit.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring the FAT Flow Label for FEC 128 VPWS Pseudowires for Load-Balancing MPLS Traffic

Configuring the FAT Flow Label for FEC 128 VPLS Pseudowires for Load-Balancing MPLS Traffic

Configuring the FAT Flow Label for FEC 129 VPWS Pseudowires for Load-Balancing MPLS Traffic

Configuring the FAT Flow Label for FEC 129 VPLS Pseudowires for Load-Balancing MPLS Traffic

flow-label-transmit-static

IN THIS SECTION

- Syntax | 7314
- Hierarchy Level | 7315
- Description | 7315
- Required Privilege Level | 7315
- Release Information | 7316

Syntax

```
flow-label-transmit-static;
```

Hierarchy Level

```
[edit protocols l2circuit neighbor neighbor-id interface interface-name],
[edit routing-instances instance-name protocols evpn],
[edit routing-instances instance-name protocols evpn interface interface-name]
```

Description

Configure the router to push the flow label on pseudowire packets it sends to a remote provider edge (PE) router. The ingress PE router inserts the flow label in the pseudowire packet regardless of the information exchanged in the signaling plane. If the egress PE router can't handle a pseudowire packet that contains the flow label, it drops the packet.

Flow-aware transport of pseudowires (FAT) flow labels enable load balancing of MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs) without the need for deep packet inspection of the payload. You can use this statement to enable devices to push FAT flow labels for:

- Forwarding equivalence class (FEC) 128 pseudowires in:
  - Virtual private LAN service (VPLS) networks.
  - Virtual private wire service (VPWS) networks.

  In these cases, configure this statement in the `[edit protocols l2circuit neighbor neighbor-id interface interface-name] hierarchy`.
- Pseudowires in Ethernet VPN (EVPN)-MPLS networks with VPWS (EVPN-VPWS):
  - Globally in an EVPN-VPWS routing instance.

    In this case, configure this statement in the `[edit routing-instances instance-name protocols evpn]` hierarchy.
  - For specific interfaces in an EVPN-VPWS routing instance.

    In this case, configure this statement in the `[edit routing-instances instance-name protocols evpn interface interface-name]` hierarchy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Support in a EVPN-VPWS routing instance added in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Configuring the FAT Flow Label for FEC 128 VPWS Pseudowires for Load-Balancing MPLS Traffic

FAT Flow Labels in EVPN-VPWS Routing Instances

flow-level-bandwidth (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 7316
- Hierarchy Level | 7316
- Description | 7317
- Required Privilege Level | 7317
- Release Information | 7317

Syntax

```
flow-level-bandwidth {
    logical-interface flow-bandwidth;
    physical-interface flow-bandwidth;
    subscriber flow-bandwidth;
}
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure allowed flow bandwidth for the packet type at each flow aggregation level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring the Maximum Flow Bandwidth at Each Flow Aggregation Level

Setting Up and Using Flow Detection

flow-level-control (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 7318
- Hierarchy Level | 7318
- Description | 7318
- Required Privilege Level | 7318
- Release Information | 7318

Syntax

```
flow-level-control {
    logical-interface flow-control-mode;
    physical-interface flow-control-mode;
    subscriber flow-control-mode;
}
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Specify how traffic in the detected flow is handled for the protocol group or packet type at one or more flow aggregation levels. Use this statement to override global flow control mode settings configured with the `flow-level-control` statement at the `[edit system ddos-protection global]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring How Traffic in a Culprit Flow Is Controlled at Each Flow Aggregation Level

Setting Up and Using Flow Detection

flow-level-control (DDoS Global Flow Detection)

IN THIS SECTION

- Syntax | 7319
- Hierarchy Level | 7319
- Description | 7319
- Options | 7319
- Required Privilege Level | 7320
- Release Information | 7320

Syntax

```
flow-level-control flow-control-mode;
```

Hierarchy Level

```
[edit system ddos-protection global]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Specify how traffic in the detected flow is handled globally for all protocol groups and packet types at all flow aggregation levels.

To override the global configuration for a protocol group or packet type, use the `flow-level-control` statement at the `[edit system ddos-protection protocols protocol-group packet-type]` hierarchy level to specify the flow control mode at one or more flow aggregation levels.

Options

flow-control-mode Mode for how traffic in the detected flow is controlled globally.

- `drop`—Drop all traffic in flow.
- `keep`—Keep all traffic in flow.
- `police`—Police the traffic to within its allowed bandwidth.

- **Default:** `drop`

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring How Traffic in a Culprit Flow Is Controlled Globally

Setting Up and Using Flow Detection

Setting Up and Using Flow Detection

flow-level-detection (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 7321
- Hierarchy Level | 7321
- Description | 7321
- Required Privilege Level | 7321
- Release Information | 7321

Syntax

```
flow-level-detection {
    logical-interface flow-detection-mode;
    physical-interface flow-detection-mode;
    subscriber flow-detection-mode;
}
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure the mode of operation for flow detection for the packet type at each flow aggregation level.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: Flow detection operates for individual flow aggregation levels only when the flow detection mode at the packet level is configured to either `automatic` or `on`.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring How Flow Detection Operates at Each Flow Aggregation Level

Setting Up and Using Flow Detection

flow-map

IN THIS SECTION

- Syntax | 7322
- Hierarchy Level | 7322
- Description | 7322
- Options | 7323
- Required Privilege Level | 7323
- Release Information | 7323

Syntax

```
flow-map flow-map-name {
    bandwidth (bps | adaptive);
    forwarding-cache {
        timeout (never non-discard-entry-only | minutes);
    }
    policy [ policy-names ];
    redundant-sources [ addresses ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Description

Configure multicast flow maps.

Options

flow-map-name—Name of the flow-map.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Example: Configuring a Multicast Flow Map

flow-monitoring

IN THIS SECTION

- Syntax | 7323
- Hierarchy Level | 7325
- Description | 7325
- Required Privilege Level | 7325
- Release Information | 7325

Syntax

```
flow-monitoring {
    version9 {
```

```
        template template-name {
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            flow-key {
                flow-direction;
                vlan-id;
                output-interface;
            }
            (ipv4-template | ipv6-template | mpls-template label-position [ positions ] | mpls-ipv4-
template label-position [ positions ] | mpls-ipvx-template);
            multi-bgp-path;
            peer-as-billing-template;
            option-refresh-rate packets packets seconds seconds;
            options-template-id id;
            source-id id;
            template-id id;
            template-refresh-rate packets packets seconds seconds;
            tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
        }
    }
    version-ipfix {
        template template-name {
            data-record-fields {
                source-prefix-as-path count;
                destination-prefix-as-path count;
                bgp-source-standard-community count;
                bgp-destination-standard-community count;
                bgp-source-extended-community count;
                bgp-destination-extended-community count;
                bgp-source-large-community count;
                bgp-destination-large-community count;
            }
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            flow-key {
                flow-direction;
                vlan-id;
            }
            (ipv4-template | ipv6-template | mpls-ipv4-template | mpls-ipvx-template | vpls-template);
            multi-bgp-path;
            nexthop-learning (enable |disable);
            observation-domain-id domain-id;
            option-refresh-rate packets packets seconds seconds;
```

```
            options-template-id id;
            template-id id;
            template-refresh-rate packets packets seconds seconds;
            tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
        }
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Specify the active monitoring properties for flow aggregation version 9 or IPFIX.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

`data-record-fields` option introduced in Junos OS Evolved Release 21.4R1.

`multi-bgp-path` option added in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates

Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices

flow-monitoring (Inline Monitoring Services)

IN THIS SECTION

- Syntax | 7326
- Hierarchy Level | 7326
- Description | 7326
- Default | 7327
- Options | 7327
- Required Privilege Level | 7328
- Release Information | 7328

Syntax

```
flow-monitoring {
    counter-profile profile-identifier;
    flow-rate kbps burst-size bytes;
    flow-limit number;
    sampling-profile profile-name;
    sampling-rate bytes;
    security-enable;
}
```

Hierarchy Level

```
[edit services inline-monitoring template template-name]
```

Description

Configures optional parameters for flow-based telemetry (FBT) for the EX4100, EX4100-F, and EX4400 switches. FBT enables per-flow-level analytics, using inline monitoring services to create flows, collect them, and export them to a collector. A flow is a sequence of packets that have the same source IP, destination IP, source port, destination port, or protocol on an interface. For each flow, various parameters are collected and sent to a collector using the open standard IPFIX template to organize the

flow. Once there is no active traffic for a flow, the flow is aged out after the configured inactive-timeout period (`flow-inactive-timeout` at the [edit services inline-monitoring template *template-name*] hierarchy level).

Default

Disabled

Options

counter-profile *profile-identifier*

Specify which counters should be exported to the collector, by selecting one of the pre-defined profile names.

- `Per_flow_6_counters`: packet range counter (3 counters), time-to-live (TTL) range counter, TCP window range counter, DoS attack (2 counters)
- `Per_flow_4_counters`: flow packet range (packet size 64 to 1500 bytes) counter (1 count), TTL range counter, TCP window range counter, DoS counter

flow-rate *kbps* burst-size *bytes*

Specify the meter rate for each flow, in kbps, and configure the maximum number of bytes allowed for incoming packets to burst above the flow meter rate.

Range: 8 to 10000000 kbps (flow-rate); 512 to 256000000 bytes (burst-size)

flow-limit *number*

Specify the maximum number of flows allowed.

Range: 0 through 32000

sampling-profile *profile-identifier*

Configure one of the following sampling profiles:

- `First_N_Pkt` : Sample contains the first N packets of a flow
- `Deterministic`: Sample contains every Nth packet of a flow
- `Random`:Sample contains randomly chosen packets from a flow
- `Combo1`: Sample contains the first N packets of a flow, followed by a random packet at the configured interval
- `Combo2`: Sample contains the first N packets of a flow, followed by a Deterministic 1 packet at the configured interval

sampling-rate *bytes*

Specify the rate at which packets are sampled to create flows, in bytes

Range: 1 to 65535 bytes

security-enable Enable security analytics; specify that Denial-of-Service (DoS) attacks are reported to the collector.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Flow-Based Telemetry

flow-monitoring (Services)

IN THIS SECTION

- Syntax | 7328
- Hierarchy Level | 7329
- Description | 7329
- Options | 7329
- Required Privilege Level | 7329
- Release Information | 7329

Syntax

```
flow-monitoring {
    version9 {
        template template-name {
```

```
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            ipv4-template;
            ipv6-template;
            option-refresh-rate {
                packets packets;
                seconds seconds;
            }
            template-refresh-rate {
                packets packets;
                seconds seconds;
            }
        }
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure flow monitoring.

Options

`version9`—Version 9 configuration.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding Interfaces

flow-recover-time (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 7330
- Hierarchy Level | 7330
- Description | 7330
- Options | 7330
- Required Privilege Level | 7331
- Release Information | 7331

Syntax

```
flow-recover-time seconds;
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure how much time must pass before a culprit flow for the packet type is considered to have returned to normal. The period starts when the flow drops below the threshold that triggered the last violation.

Options

seconds—Period required for the traffic to recover.

- **Range:** 1 through 3600 seconds
- **Default:** 60 seconds (300 seconds for PFCP and GTP path management)

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring the Recovery Period for a Culprit Flow

Setting Up and Using Flow Detection

flow-report-rate (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 7332
- Hierarchy Level | 7332
- Description | 7332
- Options | 7332
- Required Privilege Level | 7332
- Release Information | 7332

Syntax

```
flow-report-rate report-rate;
```

Hierarchy Level

```
[edit system ddos-protection global]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Set the rate at which culprit flow events are reported by system log messages, for all protocol groups and packet types on all line cards.

Options

report-rate Number of flows per second.

- **Range:** 1 through 50,000
- **Default:** 10

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

> *Configuring the Culprit Flow Reporting Rate for All Protocol Groups and Packet Types*
>
> *Setting Up and Using Flow Detection*

flow (Security Flow)

IN THIS SECTION

- Syntax | 7333
- Hierarchy Level | 7336
- Description | 7336
- Options | 7336
- Required Privilege Level | 7338
- Release Information | 7338

Syntax

```
flow {
    advanced-options {
        drop-matching-link-local-address;
        drop-matching-reserved-ip-address;
        reverse-route-packet-mode-vr;
    }
    aging {
        early-ageout seconds;
        high-watermark percent;
        low-watermark percent;
    }
    allow-dns-reply;
    allow-embedded-icmp;
    Configuring Access Lists for SNMP Access over Routing Instances;
    enable-reroute-uniform-link-check {
        nat;
    }
    enhanced-routing-mode;
    ethernet-switching {
        block-non-ip-all;
        bpdu-vlan-flooding;
        bypass-non-ip-unicast;
        no-packet-flooding {
            no-trace-route;
```

```
        }
    }
    force-ip-reassembly;
     gre-performance-acceleration;
    ipsec-performance-acceleration (Security Flow);
    load-distribution {
        session-affinity {
            ipsec;
        }
    }
    mcast-buffer-enhance;
    multicast-nh-resolve-retry multicast-nh-resolve-retry-value;
    no-local-favor-ecmp;
packet-log (Security Flow) {
        enable;
        packet-filter name {
            conn-tag conn-tag;
            destination-port (afs | bgp | biff | bootpc | bootps | cmd | cvspserver | dhcp |
domain | eklogin | ekshell | exec | finger | ftp | ftp-data | http | https | ident | imap |
kerberos-sec | klogin | kpasswd | krb-prop | krbupdate | kshell | ldap | ldp | login | mobileip-
agent | mobilip-mn | msdp | netbios-dgm | netbios-ns | netbios-ssn | nfsd | nntp | ntalk | ntp |
pop3 | pptp | printer | radacct | radius | range | rip | rkinit | smtp | snmp | snmptrap | snpp
| socks | ssh | sunrpc | syslog | tacacs | tacacs-ds | talk | telnet | tftp | timed | who |
xdmcp | zephyr-clt | zephyr-hm | zephyr-srv);
            destination-prefix destination-prefix;
            interface interface;
            logical-system logical-system;
            protocol (ah | egp | esp | gre | icmp | icmp6 | igmp | ipip | number | ospf | pim |
rsvp | sctp | tcp | udp);
            source-port (afs | bgp | biff | bootpc | bootps | cmd | cvspserver | dhcp | domain |
eklogin | ekshell | exec | finger | ftp | ftp-data | http | https | ident | imap | kerberos-sec
| klogin | kpasswd | krb-prop | krbupdate | kshell | ldap | ldp | login | mobileip-agent |
mobilip-mn | msdp | netbios-dgm | netbios-ns | netbios-ssn | nfsd | nntp | ntalk | ntp | pop3 |
pptp | printer | radacct | radius | range | rip | rkinit | smtp | snmp | snmptrap | snpp | socks
| ssh | sunrpc | syslog | tacacs | tacacs-ds | talk | telnet | tftp | timed | who | xdmcp |
zephyr-clt | zephyr-hm | zephyr-srv);
            source-prefix source-prefix;
        }
        throttle-interval milliseconds;
    }
    pending-sess-queue-length (high | moderate | normal);
    power-mode-ipsec;
    preserve-incoming-fragment-size;
```

```
    route-change-timeout seconds;
    strict-packet-order;
    syn-flood-protection-mode (syn-cookie | syn-proxy);
    sync-icmp-session;
    tcp-mss (Security Flow) {
        all-tcp {
            mss mss;
        }
        gre-in {
            mss mss;
        }
        gre-out {
            mss mss;
        }
        ipsec-vpn (Security Flow) {
            mss mss;
        }
    }
    tcp-session {
        fin-invalidate-session;
        maximum-window (128K | 1M | 256K | 512K | 64K);
        no-sequence-check;
        no-syn-check;
        no-syn-check-in-tunnel;
        rst-invalidate-session;
        rst-sequence-check;
        strict-syn-check;
        tcp-initial-timeout seconds;
        time-wait-state {
            (session-ageout | session-timeout seconds);
            apply-to-half-close-state;
        }
    }
    traceoptions (Security Flow) {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace;
        packet-filter name {
            conn-tag conn-tag;
            destination-port (afs | bgp | biff | bootpc | bootps | cmd | cvspserver | dhcp |
domain | eklogin | ekshell | exec | finger | ftp | ftp-data | http | https | ident | imap |
kerberos-sec | klogin | kpasswd | krb-prop | krbupdate | kshell | ldap | ldp | login | mobileip-
```

```
agent | mobilip-mn | msdp | netbios-dgm | netbios-ns | netbios-ssn | nfsd | nntp | ntalk | ntp |
pop3 | pptp | printer | radacct | radius | range | rip | rkinit | smtp | snmp | snmptrap | snpp
| socks | ssh | sunrpc | syslog | tacacs | tacacs-ds | talk | telnet | tftp | timed | who |
xdmcp | zephyr-clt | zephyr-hm | zephyr-srv);
            destination-prefix destination-prefix;
            interface interface;
            logical-system logical-system;
            protocol (ah | egp | esp | gre | icmp | icmp6 | igmp | ipip | number | ospf | pim |
rsvp | sctp | tcp | udp);
            source-port (afs | bgp | biff | bootpc | bootps | cmd | cvspserver | dhcp | domain |
eklogin | ekshell | exec | finger | ftp | ftp-data | http | https | ident | imap | kerberos-sec
| klogin | kpasswd | krb-prop | krbupdate | kshell | ldap | ldp | login | mobileip-agent |
mobilip-mn | msdp | netbios-dgm | netbios-ns | netbios-ssn | nfsd | nntp | ntalk | ntp | pop3 |
pptp | printer | radacct | radius | range | rip | rkinit | smtp | snmp | snmptrap | snpp | socks
| ssh | sunrpc | syslog | tacacs | tacacs-ds | talk | telnet | tftp | timed | who | xdmcp |
zephyr-clt | zephyr-hm | zephyr-srv);
            source-prefix source-prefix;
        }
        rate-limit rate-limit;
        trace-level {
            (brief | detail | error);
        }
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Determine how the device manages packet flow. The device can regulate packet flow in the following ways:

Options

advanced-options Flow configuration advanced options.

- Values:
  - drop-matching-link-local-address—Drop matching link local address.
  - drop-matching-reserved-ip-address—Drop matching reserved source IP address.
  - reverse-route-packet-mode-vr—Allow reverse route lookup with packet mode vr.

allow-dns-reply Allow unmatched incoming DNS reply packet.

allow-embedded-icmp Allow embedded ICMP packets not matching a session to pass through.

allow-reverse-ecmp Allow reverse ECMP route lookup.

enable-reroute-uniform-link-check Enable reroute check with uniform link.

- Values:
  - nat—Enable NAT check.

enhanced-routing-mode Enable enhanced route scaling.

force-ip-reassembly Force to reassemble IP fragments.

gre-performance-acceleration Accelerate the GRE traffic performance.

ipsec-performance-acceleration Accelerate the IPSec traffic performance.

mcast-buffer-enhance Allow to hold more packets during multicast session creation.

multicast-nh-resolve-retry You can use this command configure the multicast route next-hop resolve attempts. When a multicast route next-hop resolve is unsuccessful, the SRX Series Firewall attempts to resolve the next-hop route based on the specified retry counts.

- **Default:** 0
- **Range:** 0 through 20

no-local-favor-ecmp Does not prefer local node in HA ECMP route lookup.

pending-sess-queue-length Maximum queued length per pending session.

- Values:

- high—Maximum number of queued sessions.
- moderate—Allow more queued sessions than normal.
- normal—Normal number of sessions queued.

power-mode-ipsec Enable power mode ipsec processing.

preserve-incoming-fragment-size Preserve incoming fragment size for egress MTU.

route-change-timeout Timeout value for route change to nonexistent route (seconds).

- **Default:** 6
- **Range:** 6 through 1800

strict-packet-order You can use this command to maintain multicast traffic order and resolve packet drop issue.

syn-flood-protection-mode TCP SYN flood protection mode.

- Values:
  - syn-cookie—Enable SYN cookie protection.
  - syn-proxy—Enable SYN proxy protection.

sync-icmp-session Allow icmp sessions to sync to peer node.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.5. The `power-mode-ipsec` option added in Junos OS Release 18.3R1 for vSRX Virtual Firewall instances, in Junos OS Release 18.4R1 for SRX4100 and SRX4200 devices, and in Junos OS Release 18.2R2 for SRX5400, SRX5600, and SRX5800 devices. The `multicast-nh-resolve-retry` and the `strict-packet-order` options are added in Junos OS Release 20.2R2 for SRX345 and SRX1500 devices.The `gre-performance-acceleration` option added in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

- Understanding Traffic Processing on Security Devices
- Understanding Session Characteristics for SRX Series Services Gateways
- Understanding Packet Flow in Logical Systems for SRX Series Devices

flow (Security IDP)

IN THIS SECTION

- Syntax | 7339
- Hierarchy Level | 7340
- Description | 7340
- Options | 7340
- Required Privilege Level | 7343
- Release Information | 7344

Syntax

```
flow {
    (allow-icmp-without-flow | no-allow-icmp-without-flow);
    drop-if-no-policy-loaded;
    drop-on-failover;
    drop-on-limit;
    fifo-max-size value;
    hash-table-size value;
    idp-bypass-cpu-threshold idp-bypass-cpu-threshold;
    idp-bypass-cpu-tolerance idp-bypass-cpu-tolerance;
    idp-bypass-cpu-usg-overload;
    intel-inspect-cpu-usg-threshold intel-inspect-cpu-usg-threshold;
    intel-inspect-cpu-usg-tolerance intel-inspect-cpu-usg-tolerance;
    intel-inspect-disable-content-decompress;
    intel-inspect-enable;
    intel-inspect-free-mem-threshold intel-inspect-free-mem-threshold;
    intel-inspect-mem-tolerance intel-inspect-mem-tolerance;
```

```
    intel-inspect-protocols [ intel-inspect-protocols ];
    intel-inspect-session-bytes-depth intel-inspect-session-bytes-depth;
    intel-inspect-signature-severity (critical | major | minor);
    (log-errors | no-log-errors);
    max-sessions-offset value;
    max-timers-poll-ticks value;
    min-objcache-limit-lt lower-threshold-value;
    min-objcache-limit-ut upper-threshold-value;
    reject-timeout value;
    (reset-on-policy | no-reset-on-policy);
    udp-anticipated-timeout value;
}
```

Hierarchy Level

```
[edit security idp sensor-configuration]
```

Description

Configure the IDP engine to manage the packet flow.

Options

allow-nonsyn-connection Allow TCP non-syn connection.

drop-if-no-policy-loaded Drop all traffic till IDP policy gets loaded.

drop-on-failover Drop traffic on HA failover sessions.

drop-on-limit Drop connections on exceeding resource limits.

fifo-max-size Maximum fifo size.

Sets the maximum FIFO size (range: 1 through 65535).

- **Range:** 1 through 65535

hash-table-size Flow hash table size. Sets the packet flow hash table size.

- **Range:** 1024 through 1000000

idp-bypass-cpu-threshold

CPU usage in percentage for IDP bypass.

- **Default:** 85
- **Range:** 0 through 99

idp-bypass-cpu-tolerance

CPU usage in percentage for IDP bypass.

- **Default:** 5
- **Range:** 1 through 99

idp-bypass-cpu-usg-overload

Enable IDP bypass of sessions or packets on CPU usage overload.

intel-inspect-cpu-usg-threshold

CPU usage threshold percentage for intelligent inspection.

- **Default:** 80
- **Range:** 0 through 99

intel-inspect-cpu-usg-tolerance

CPU usage tolerance percentage for intelligent inspection.

- **Default:** 5
- **Range:** 1 through 99

intel-inspect-disable-content-decompress

Disable payload content decompression.

intel-inspect-enable

Minimize IDP processing during system overload.

intel-inspect-free-mem-threshold

Free memory threshold percentage for intelligent inspection.

- **Default:** 15
- **Range:** 1 through 100

intel-inspect-mem-tolerance

Memory tolerance percentage for intelligent inspection.

- **Default:** 5
- **Range:** 1 through 100

intel-inspect-protocols

Protocols to be processed in intelligent inspection mode.

intel-inspect-session-bytes-depth

Session bytes scanning depth.

- **Default:** 0

- **Range:** 0 through 1000000

intel-inspect-signature-severity

Signature severities to be considered for IDP processing.

- Values:
  - critical
  - major
  - minor

log-errors

Enable the error log to generate the result of success or failure about the flow. A flow-related error is when IDP receives a packet that does not fit into the expected flow. By default an error log is enabled.

max-sessions-offset

Maximum session offset limit percentage.

Set an offset (percentage) for the maximum IDP session limit. The `max-sessions-offset` option sets an offset for the maximum IDP session limit. When the number of IDP sessions exceeds the maximum session limit, a warning is logged that conditions exist where IDP sessions could be dropped. When the number of IDP sessions drops below the maximum IDP session limit minus the offset value, a message is logged that conditions have returned to normal.

- **Range:** 0 through 99

max-timers-poll-ticks

Specify the time at which timer ticks at regular interval.

- **Syntax:** *value*—Maximum amount of time at which the timer ticks.
- **Range:** 0 through 1000 ticks
- **Default:** 1000 ticks

min-objcache-limit-lt

Memory lower threshold limit percentage.

- **Syntax:** *value*— Memory lower threshold limit percentage.
- **Range:** 1 through 100

min-objcache-limit-ut

Memory upper threshold limit percentage.

- **Syntax:** *value*— Memory upper threshold limit percentage.
- **Range:** 1 through 100

no-log-errors

Do not flow log errors.

reject-timeout

Specify the amount of time in seconds within which a response must be received.

This time-out is applied on flow when drop-connection action is taken by IPS for TCP flow.

- **Syntax:** *value*—Maximum amount of time in seconds.
- **Range:** 1 through 65535
- **Default:** 300 seconds

reset-on-policy

IDP keeps track of connections in a table. If enabled, the security module resets the flow table each time a security policy loads or unloads. If this setting is disabled, then the security module continues to retain a previous security policy until all flows referencing that security policy go away. Juniper Networks recommends that you keep this setting enabled to preserve memory.

When a new IDP policy is loaded, the existing sessions are inspected using the newly loaded policy and the existing sessions not ignored for IDP processing. The `reset-on-policy` command is used to decide whether to continue the IDP inspection with the newly loaded IDP policy or not. This command is disabled by default and all the existing sessions continue to be inspected with newly loaded IDP policy.

> **NOTE**: In Junos OS Release 18.2R1-S1and Junos OS Release 18.3R1, the `no-reset-on-policy` option is not supported on SRX5000 line of devices with SRX5K-SPC3.

session-steering

Session steering for session anticipation.

udp-anticipated-timeout

Sets the maximum UDP anticipated timeout value.

- **Range:** 1 through 65535

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Options intel-inspect-cpu-usg-threshold, intel-inspect-cpu-usg-tolerance, intel-inspect-disable-content-decompress, intel-inspect-enable, intel-inspect-free-mem-threshold, intel-inspect-mem-tolerance, intel-inspect-protocols, intel-inspect-session-bytes-depth, and intel-inspect-signature-severity options added in Junos OS Release 19.2R1.

Starting in Junos OS Release 18.4R1, the `reset-on-policy` command is deprecated—rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration.

flow-server

IN THIS SECTION

- Syntax | 7344
- Hierarchy Level | 7345
- Description | 7345
- Options | 7345
- Required Privilege Level | 7346
- Release Information | 7346

Syntax

```
flow-server hostname {
    aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
            caida-compliant;
        }
        source-prefix;
    }
```

```
    autonomous-system-type (origin | peer);
    dscp dscp-value;
    forwarding-class class-name;
    (local-dump | no-local-dump);
    port port-number;
    source-address address;
    version format;
    version9 {
        template template-name;
    }
    version-ipfix {
        template template-name;
    }
}
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls | vpls |
bridge) output],
[edit forwarding-options sampling family (inet | inet6 | mpls | vpls | bridge) output]
```

Description

Configure a server to either:

- Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect, **or**
- Specify a host system to collect sampled flows using the version 9 or IPFIX format.

You can configure up to one version 5 and one version 8 flow format at the `[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server hostname]` hierarchy level. For the same configuration, you can specify either version 9/IPFIX flow record formats or version 5/8 formats, not both types of formats.

Options

hostname—IP address–IPv4 or IPv6–or identifier of the host system (the workstation either running the cflowd utility or collecting traffic flows using version 9 or IPFIX).

NOTE: Only host systems running IPv4 are supported on QFX10000 switches.

You can configure only one host system for version 9.

NOTE: IPv6 configuration for `flow-server` is supported only in Junos OS Release 12.3 and later.

Note that when you configure an IPv6 address for the `flow-server` statement, you must also configure an IPv6 address for the `inline-jflow source-address` statement at the `[edit forwarding-options sampling instance` *`instance-name`* `family (inet | inet6 | mpls | vpls | bridge) output]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`version9` statement introduced in Junos OS Release 8.3.

Support at the following hierarchy levels introduced in Junos OS Release 18.2R1: `[edit forwarding-options sampling instance instance-name family bridge]`, , `[edit forwarding-options sampling family bridge]`.

Support for IPv6 addresses for the `flow-server` *`address`* and `source-address` *`address`* statements within sampling instances introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

flow-session

IN THIS SECTION

- Syntax | 7347
- Hierarchy Level | 7347
- Description | 7347
- Options | 7348
- Required Privilege Level | 7349
- Release Information | 7349

Syntax

```
flow-session {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specifies the number of flow sessions that user logical system administrators and primary logical system administrators configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum *amount*

A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 1000000

reserved *amount*

A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

NOTE: An IPv6 session consumes twice the memory of an IPv4 session. Therefore the number of sessions available for IPv6 is half the reserved and maximum quotas configured for the flow session resource in a security profile. Use the vty command **show usp flow resource usage cp-session** to check flow session usage.

For performance consideration, flow session rate-limiting is done on distributed central point (DCP) for SRX5400, SRX5600, and SRX5800 devices. DCP session is always created before SPU flow session on these devices. You can count DCP session instead of SPU flow session to check whether the maximum/reserved quotas are reached for a specific logical system or tenant system.

There are no differences between DCP session and SPU flow session count because a session without NAT consumes one DCP session and one SPU flow session.

For sessions with NAT configuration, one connection can consume one SPU flow session and two DCP sessions. If all sessions are NAT sessions and each session consumes two DCP sessions, then the logical system or tenant system can only reach half of the maximum quotas configured for the flow session resources.

Typically, there are sessions with NAT and non-NAT mixed traffic, so the logical system or tenant system can create flow sessions between “maximum”/2 and “maximum”.

This functionality is similar to the reserved quotas that are configured for the flow session resources in a security profile. As a logical system or tenant system can consume more resources than the reserved quotas, you can configure additional flow session resources for reserved quotas with NAT configuration, if there are free resources available in the logical system or tenant system.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only)

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only)

flow-specs

IN THIS SECTION

- Syntax | 7349
- Hierarchy Level | 7350
- Description | 7350
- Options | 7350
- Required Privilege Level | 7351
- Release Information | 7351

Syntax

```
flow-specs {
    flow-specification-name {
        ethernet-frame-size ethernet-frame-size;
        multicast-ipv4 {
            dest-ip-multicast-group ipv4-mcast-address;
            source-ip ipv4-address;
```

```
        }
        multicast-vlan-flood;
        unicast-ethernet {
            destination-mac destination-mac-address destination-mac-mask destination-mac-mask;
            ethertype ethertype;
            source-mac mac-address source-mac-mask source-mac-mask;
        }
        unicast-ethernet-ipv4 {
            destination-ip destination-ip-address destination-ip-mask destination-ip-mask;
            destination-l4-port destination-l4-port-number;
            destination-mac destination-mac-address;
            ip-proto protocol;
            source-ip source-ip-address source-ip-mask source-ip-mask;
            source-l4-port source-l4-port-number;
        }
    }
}
```

Hierarchy Level

```
[edit protocols oam fabric]
```

Description

Configure the fabric flow specification parameters for internal fabric monitoring. The following conditions apply:

- You must specify a protocol type for each flow specification: `unicast-ethernet`, `unicast-ethernet-ipv4`, `multicast-ipv4`, or `multicast-vlan-flood`.
- You can specify only one type for each flow specification.
- All other parameters are optional.

The remaining statements under this hierarchy level are explained separately.

Options

flow-specification-name

Name of the flow specification.

- **Syntax:** 1 to 20 alphanumeric characters

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | **29924**

flow-table-size

IN THIS SECTION

- Syntax | **7352**
- Hierarchy Level | **7352**
- Description | **7352**
- Required Privilege Level | **7353**
- Release Information | **7353**

Syntax

```
flow-table-size {
    ipv4-flow-table-size  units;
    ipv6-extended-attrib;
    ipv6-flow-table-size units;
    mpls-flow-table-size  units;
    vpls-flow-table-size units;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services]
```

Description

Configure the size of hash tables in units of 256K entries for inline services sampling.

Starting with Junos OS Release 15.1F2, by default, the software allocates one 1K IPv4 flow table. To allocate 15 256K IPv4 flow tables, the former default, you can enter this configuration from the `[edit]` hierarchy level:

```
[edit]
user@router# set chassis fpc inline-services flow-table-size ipv4-flow-table-size 15
```

The flow table size is the sum of all the configured flow table types. The sum of all the configured flow table types cannot exceed the maximum supported flow table size. For example, for the MPC2E, MPC3E, and MPC4E line cards, the maximum supported flow table size is 15 units. The different flow table families could be configured as follows for these line cards:

IPv4 table size: 5
IPv6 table size: 5
MPLS table size: 4
VPLS table size: 1

Total table size: 15

Starting with Junos OS Release 17.3R1, the maximum supported flow table size varies by line card type:

Table 44: Maximum Supported Flow Table Size for MX Series Routers

Line Card	Maximum Number of Units (1 unit =256K entries)
MPC2E, MPC3E, and MPC4E	15
MPC5E and MPC6E	220
MPC7E , MPC9E, LC2103, and LC480	48
MPC8E	97
MPC10E and MPC11E	48
JNP10K-LC2301, JNP10K-LC4800, JNP10K-LC9600, and MX304-LMIC16BASE	80

NOTE: We recommend that you configure this statement during a maintenance window:

- Prior to Junos OS 16.1R1 and 15.1F2, the FPC reboots automatically after you commit this configuration change.
- Starting from Junos OS Release 16.1R1 and 15.1F2, any changes in the configured size of the flow table do not require a reboot of the FPC.

The remaining statements are defined separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

`ipv6-extended-attrib` option added in Junos OS Release 14.2 for MX Series routers.

`vpls-flow-table-size` option added in Junos OS Release 13.2 for MX Series routers.

`bridge-flow-table-size` option added in Junos OS Release 18.2R1 for MX Series routers, replacing the `vpls-flow-table-size` option, which has been deprecated.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers

flow-table-size (Chassis)

IN THIS SECTION

- Syntax | 7354
- Hierarchy Level | 7354
- Description | 7355
- Options | 7355
- Required Privilege Level | 7355
- Release Information | 7355

Syntax

```
flow-table-size size;
```

Hierarchy Level

```
[edit chassis fpc slot inline-video-monitoring]
```

Description

Configure the number of video flows that can be measured per Packet Forwarding Engine by an MPC at a given time. This value takes effect the next time the MPC is rebooted.

Options

size Number of video flows per Packet Forwarding Engine.

- **Range:** 16 through 8192

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Inline Video Monitoring on MX Series Routers

flow-tap

IN THIS SECTION

- Syntax | 7356
- Hierarchy Level | 7356
- Description | 7356
- Options | 7356
- Required Privilege Level | 7357
- Release Information | 7357

Syntax

```
flow-tap {
    (interface interface-name | tunnel-interface interface-name);
    family (inet | inet6);
}
```

Hierarchy Level

```
[edit services]
```

Description

Enable the FlowTap service or FlowTapLite service on an interface. FlowTapLite is a lighter version of the flow-tap application that is available only on tunnel interfaces for those MX Series platforms and ACX Series platforms running Junos OS Evolved that support the feature.

Starting in Junos OS Release 17.3R1, the FlowTapLite service can run concurrently with the radius-flow-tap service on the same MX Series router. The radius-flow-tap service (`[edit services radius-flow-tap]`) is required for subscriber secure policy mirroring on MX Series routers.

In earlier releases, the FlowTapLite and radius-flow-tap services cannot run concurrently on an MX Series router, which prevents you from running FlowTapLite monitoring and subscriber secure policy mirroring at the same time. These services cannot run concurrently on an ACX router.

Options

interface *interface-name*
: Use the specified interface for the FlowTap application.

tunnel-interface *interface-name*
: Use the specified tunnel interface for the FlowTapLite application.

family
: (Not applicable for FlowTapLite) Apply flow-tap services to the specified family. If you do not specify an option, the FlowTap service is applied only to IPv4 traffic.

- inet—IPv4 traffic.
- inet6—IPv6 traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

ccc option introduced in Junos OS Release 17.2.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, the FlowTapLite service can run concurrently with the radius-flow-tap service on the same MX Series router.

RELATED DOCUMENTATION

- Configuring Junos Packet Vision on MX, M and T Series Routers

flow-tap-dtcp

IN THIS SECTION

- Syntax | 7358
- Hierarchy Level | 7358
- Description | 7358
- Options | 7358
- Required Privilege Level | 7358
- Release Information | 7359

Syntax

```
flow-tap-dtcp {
    ssh {
        connection-limit limit;
        rate-limit limit;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure Dynamic Tasking Control Protocol (DTCP) sessions to run over SSH in support of the flow-tap, FlowTapLite, or radius-flow-tap services. Note that the flow-tap feature is not supported on outbound, or egress, traffic. Only inbound, or ingress, traffic is supported.

This statement is required for DTCP-initiated subscriber secure policy mirroring (radius-flow-tap service).

Options

`connection-limit` *`limit`*—(Optional) Maximum number of connections allowed.

- **Range:** 1 through 250
- **Default:** 75

`rate-limit` *`limit`*—(Optional) Maximum number of connection attempts allowed per minute.

- **Range:** 1 through 250
- **Default:** 150

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring a DTCP-over-SSH Connection to the Mediation Device

Configuring Flow-Tap Security Properties on MX, M and T Series Routers

Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs

Configuring DTCP-over-SSH Service for the Flow-Tap Application

flow-timeout-time (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 7359
- Hierarchy Level | 7359
- Description | 7360
- Options | 7360
- Required Privilege Level | 7360
- Release Information | 7360

Syntax

```
flow-timeout-time seconds;
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure the period of time that a culprit flow is suppressed for the packet type. The timeout period is effective only when timing out has been enabled with the "timeout-active-flows" on page 29372 statement.

Options

seconds—Period that the traffic is suppressed.

- **Range:** 1 through 7200 seconds
- **Default:** 300 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring the Timeout Period for a Culprit Flow

Setting Up and Using Flow Detection

flow-sample

IN THIS SECTION

Syntax | 7361

- Hierarchy Level | 7361
- Description | 7361
- Options | 7361
- Required Privilege Level | 7362
- Release Information | 7362

Syntax

```
flow-sample {
    ae-parent-index;
}
```

Hierarchy Level

```
[edit forwarding-options sflow]
```

Description

Configure flow sample inputs.

Options

ae-parent-index

Configure to send aggregated Ethernet parent index to an sFlow datagram.

When you configure sFlow, a sampled packet is captured from the traffic stream and the packet details are added to an sFlow datagram. This datagram contains information such as source IP, destination IP, ingress SNMP index, and egress SNMP index of the respective interface.

sFlow is not supported on aggregated Ethernet bundles but you can configure sFlow on any child interface of an aggregated Ethernet interface. When you configure `flow-sample ae-parent-`

`index`, the sFlow datagram includes the SNMP index of the parent aggregated Ethernet interface instead of the child interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

flows (PCC Rules)

IN THIS SECTION

- Syntax | 7362
- Hierarchy Level | 7362
- Description | 7363
- Options | 7363
- Required Privilege Level | 7363
- Release Information | 7363

Syntax

```
flows ([flow-identifier] | any);
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rule-name from],
[edit services pcef pcc-rules rule-name from]
```

Description

Specify the service data flow (SDF) filters (flow identifiers) that define the match criteria for the policy and charging control (PCC) rule. You can configure a maximum of 15 SDF filters. You must include the `flows` statement in a PCC rule. If you do not want to filter subscriber traffic based on SDF filters, use the `any` option.

If you are using Junos OS Subscriber Aware, specify the name of the SDF filter at the `[edit unified-edge pcef pcc-rules` *`rule-name`* `from]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the SDF filter at the `[edit services pcef pcc-rules` *`rule-name`* `from]` hierarchy level.

Options

flow-identifier Name of an SDF filter that is used to detect IP packet flows. You can configure a maximum of 15 SDF filters. The referenced SDF filters must be configured.

- **Range:** 1 through 63 characters.

any All IP packet flows.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-rules` *`rule-name`* `from]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Rules

Configuring Service Data Flow Filters

fnp

IN THIS SECTION

- Syntax | 7364
- Hierarchy Level | 7364
- Description | 7364
- Options | 7365
- Required Privilege Level | 7365
- Release Information | 7365

Syntax

```
fnp {
    interval <100ms | 1s | 10s | 1m | 10m>;;
    loss-threshold number
    interface interface name {
        domain-id domain-id
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet]
```

Description

On routers with ge, xe, or ae interfaces, configure an OAM Ethernet failure notification protocol.

Options

`interval` *`number`*—Specifies the time between the transmission of FNP messages.

`loss-threshold` *`number`*—FNP messages that can be lost before the FNP message is considered aged out and flushed.

`interface` *`interface-name`*—Name of the Ethernet interface.

`domain-id` *`number`*—Domain ID of the access network.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Ethernet Failure Notification Protocol Overview

Configuring the Failure Notification Protocol

force-continue (OCS Partition)

IN THIS SECTION

- Syntax | **7366**
- Hierarchy Level | **7366**
- Description | **7366**
- Required Privilege Level | **7366**
- Release Information | **7366**

Syntax

```
force-continue;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Enable subscriber traffic to flow before the first interrogation with the Online Charging System (OCS) occurs.

Wireline customers often control user services solely through the Policy Control and Charging Rules Function (PCRF) and use the OCS as a convenient real-time usage monitoring mechanism rather than as an enforcement unit. To decrease the number of possible erroneous OCS configurations, include the `force-continue` statement to force the broadband Policy and Charging Enforcement Function (BPCEF) to limit the impact of negative responses from the OCS and quota expirations, and to prevent sending OCS notifications for affected rating groups. Whenever the PCEF receives a negative response to any reported group, it stops reporting this group to the OCS.

NOTE: The `force-continue` state is required; you must configure it in the OCS partition.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

force-control-packets-on-transit-path

IN THIS SECTION

- Syntax | 7367
- Hierarchy Level | 7367
- Description | 7367
- Required Privilege Level | 7368
- Release Information | 7368

Syntax

```
force-control-packets-on-transit-path;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Force host-injected control packets to follow the transit path. Host-injected control traffic reaches the GRE tunnel interface queues at the PFE when the control session is over the GRE tunnel interface. This includes control protocols OSPF, BGP, PIM, RSVP, LDP, OAM, BFD and MSDP.

Injection of control traffic ensures that the kernel includes the interface ID of the GRE logical interface and the unicast next-hop ID of the corresponding GRE physical interface along with the packet that is injected into the PFE. PFE code uses the logical interface ID and next-hop ID information to forward the packet. GRE encapsulation occurs at the PFE level and the packet gets looped back over the GRE tunnel

after being subjected to CoS treatment over the GRE interface queues. The looped back packet is subjected to a second lookup and is then forwarded over the egress physical interface.

Support for control traffic to travel over the GRE tunnel with CoS intact means the `copy-tos-to-outer-ip-header` statement cannot be used. Use the `copy-tos-to-outer-ip-header-transit` configuration statement instead. With this statement, all transit packets on the GRE tunnel logical interface have the TOS copied to the outer header.

To enable this feature, configure the `force-control-packets-on-transit-path` statement on the GRE tunnel logical interface. For example:

```
gr-4/0/10 {
    unit 0 {
        tunnel {
            source 192.168.2.1
            destination 192.16.4.2;
        }
        force-control-packets-on-transit-path;
    }
}
```

This example enables Layer 2.5 injection of control protocols on the GRE tunnel `gr-4/0/10 unit 0` only.

This feature is supported on the MX204 and the MX NG MPCs (MPC2E-NG and MPC3E-NG).

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

tunnel | 30192

force-discover (DHCP Client)

IN THIS SECTION

- Syntax | 7369
- Hierarchy Level | 7369
- Description | 7369
- Required Privilege Level | 7369
- Release Information | 7369

Syntax

```
force-discover;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number dhcp-client]
```

Description

Forces the DHCP client to send a DHCP discover packet after one to three failed `dhcp-request` attempts. The `force-discover` option ensures that the DHCP server will assign the same or a new IP address to the client.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D80.

Starting with Junos OS Release 17.3R1, on all SRX Series Firewalls and vSRX Virtual Firewall instances, the CLI option `dhcp-client` at `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy is changed to `dhcp` to keep consistent with other Junos platforms. There is no change in the functionality with the changed CLI option `dhcp`.

RELATED DOCUMENTATION

DHCP Client

DHCPv6 Client

Minimum DHCP Client Configuration

force-entry (IDS MS-DPC)

IN THIS SECTION

- Syntax | 7370
- Hierarchy Level | 7370
- Description | 7371
- Required Privilege Level | 7371
- Release Information | 7371

Syntax

```
(force-entry | ignore-entry);
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Specify handling of entries in the IDS events cache when using the MS-DPC:

- `force-entry`—Ensure that the entry has a permanent place in the IDS cache after one event is registered.
- `ignore-entry`—Ensure that all IDS events are ignored.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in IDS Rules

force-ip-reassembly

IN THIS SECTION

- Syntax | 7372
- Hierarchy Level | 7372
- Description | 7372
- Required Privilege Level | 7372
- Release Information | 7372

Syntax

```
force-ip-reassembly;
```

Hierarchy Level

```
[edit security flow]
```

Description

Reassemble all IP fragmented packets before forwarding.

This option is disabled by default. You can disable this option by deleting this flag from the CLI.

Force reassembly happens when all the fragment packets are received, and the session has been created (first path) or session has been matched (fast path) in the session SPU when you configure `force-ip-reassembly` command.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

force-premium (Firewall Filter Action)

IN THIS SECTION

- Syntax | 7373
- Hierarchy Level | 7373
- Description | 7373
- Required Privilege Level | 7374
- Release Information | 7374

Syntax

```
force-premium;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name then],
[edit logical-systems logical-system-name firewall family family-name filter filter-name term term-name then]
```

Description

Firewall filter option to force premium treatment for traffic (MX Series routers)— By default, a hierarchical policer processes the traffic it receives according to the traffic's forwarding class. Premium, expedited-forwarding traffic has priority for bandwidth over aggregate, best-effort traffic. Now you can include the `force-premium` option at the `[edit firewall filter filter-name term term-name]` hierarchy level to ensure that traffic matching the term is treated as premium traffic by a subsequent hierarchical policer, regardless of its forwarding class. This traffic is given preference over any aggregate traffic received by that policer. Consider a scenario where a firewall filter is applied to an interface that receives both expedited-forwarding voice traffic and best-effort video traffic. Traffic that matches the first term of the filter is passed to a hierarchical policer in the second term. The hierarchical policer also receives best-effort data traffic from another source. The filtered video traffic is treated the same as this data traffic, as aggregate traffic with a lower priority than the premium voice traffic. Consequently, some of the video traffic might be dropped and some of the data traffic passed on.

To avoid that situation, include the `force-premium` option in the firewall filter term that passes traffic to the hierarchical policer. This term forces the video traffic to be marked as premium traffic. The hierarchical policer gives both the voice traffic and the video traffic priority over the aggregate data traffic.

NOTE: The `force-premium` filter option is supported only on MPCs.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3 for `family inet` and `inet6`.

Support for `family vpls`, `ccc`, and `bridge` added in Junos OS Releases 13.3R8, 13.3R10, 14.1R8, 14.2R7, 15.1R4,16.1R1, and 17.1R1.

RELATED DOCUMENTATION

Firewall Filter Nonterminating Actions

force switch

IN THIS SECTION

- Syntax | 7375
- Hierarchy Level | 7375
- Description | 7375
- Required Privilege Level | 7375

Syntax

```
request protection-group ethernet-aps force-switch md <md> ma <ma>
```

Hierarchy Level

```
[edit protocols protection-group ethernet-aps]
```

Description

Forces traffic to switch from the active path to the alternate path. If the working path is the active path, traffic will be switched to the protection path. If the protection path is the active path, traffic will be switched to the protection path.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Ethernet Automatic Protection Switching Overview

force-switch

IN THIS SECTION

- Syntax | 7376
- Hierarchy Level | 7376
- Description | 7376
- Required Privilege Level | 7376
- Release Information | 7376

Syntax

```
force-switch;
```

Hierarchy Level

```
[edit chassis synchronization source  (external | interface-name)]
```

Description

For MX Series routers operating with Synchronous Ethernet, force a router to use the clock source. provided that the source is enabled and not locked out. Only one configured source may be force-switched.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

force-up (EX and QFX Series)

IN THIS SECTION

Syntax | 7377

- Hierarchy Level (EX Series) | 7377
- Hierarchy Level (EX4600 and QFX Series) | 7377
- Description | 7377
- Required Privilege Level | 7377
- Release Information | 7377

Syntax

```
force-up;
```

Hierarchy Level (EX Series)

```
[edit interfaces interface-name ether-options 802.3ad lacp]
```

Hierarchy Level (EX4600 and QFX Series)

```
[edit interfaces interface-name ether-options 802.3ad lacp]
[edit interfaces interface-name aggregated-ethernet-options lacp]
```

Description

Set the state of the interface as UP when the peer has limited LACP capability.

Required Privilege Level

interface— To view this statement in the configuration.
interface-control— To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Understanding Aggregated Ethernet Interfaces and LACP for Switches

Configuring Aggregated Ethernet LACP (CLI Procedure)

Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch

Junos OS Network Interfaces Library for Routing Devices

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support

Configuring Gigabit Ethernet Interfaces (J-Web Procedure)

Junos OS Ethernet Interfaces Configuration Guide

force-up

IN THIS SECTION

- Syntax | 7378
- Hierarchy Level | 7378
- Description | 7379
- Required Privilege Level | 7379
- Release Information | 7379

Syntax

```
force-up;
```

Hierarchy Level

```
[edit interfaces interface-name aggregated-ethernet-options lacp]
```

Description

Configure the peer interface (in MC-LAG) to remain up even with limited LACP capability.

Required Privilege Level

interface— To view this statement in the configuration.
interface-control— To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

> **NOTE**: For EX9200 switches, you must configure `force-up` on physical interfaces of both MC-LAG peers for this feature to work properly.

RELATED DOCUMENTATION

- Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up

force-version (IEEE 802.1D STP)

IN THIS SECTION

- Syntax | 7380
- Hierarchy Level | 7380
- Description | 7380
- Required Privilege Level | 7380
- Release Information | 7380

Syntax

```
force-version stp;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (rstp | vstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(rstp | vstp)],
```

```
[edit protocols (rstp | vstp)],
[edit routing-instances routing-instance-name protocols (rstp | vstp)]
```

Description

Forces the spanning-tree version to be the original IEEE 803.1D STP protocol.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

Support for logical systems added in Junos OS Release 9.6.

force-fbf-terms

IN THIS SECTION

- Syntax | 7381
- Hierarchy Level | 7381
- Description | 7381
- Required Privilege Level | 7381
- Release Information | 7382

Syntax

```
force-fbf-terms
```

Hierarchy Level

```
[edit system packet-forwarding-options firewall]
```

Description

When a firewall filter is configured with two terms; the first term without filter based forwarding action, and the second term with a filter based forwarding action, the packet is assigned the routing instance specified by the filter based forwarding action in the second term. This happens even if the first term is configured with a terminating action such as *accept*. Ideally, after the first term, the packet should terminate and not be procesed by the second term. But this is the default behavior - a firewall filter term that has a filter based forwarding action is always processed first, regardless of its order of placement.

You use `force-fbf-terms` to change this default behavior. When you apply this configuration, the firewall filter terms in a firewall are always processed in the order of their placement in the configuration, irrespective of whether a firewall filter term has a filter based forwarding action.

Required Privilege Level

firewall – To view this statement in the configuration

firewall-control - To add this statement to the configuration

Release Information

Statement introduced in Junos OS Evolved Release 24.4R1

format (Flat-File Accounting Options)

IN THIS SECTION

- Syntax | 7382
- Hierarchy Level | 7382
- Description | 7382
- Options | 7383
- Required Privilege Level | 7383
- Release Information | 7383

Syntax

```
format (csv | ipdr);
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name]
```

Description

Specify the format for logging the flat-file accounting statistics.

Options

csv Comma-separated values (CSV) format.

ipdr IP Detail Record (IPDR) format.

- **Default:** `ipdr`

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Flat-File Accounting Overview

format (LRF Profile)

IN THIS SECTION

- Syntax | 7384
- Hierarchy Level | 7384
- Description | 7384
- Required Privilege Level | 7384
- Release Information | 7384

Syntax

```
format ipfix;
```

Hierarchy Level

```
[edit services lrf profile profile-name template template-name]
```

Description

Configure a format for the template. Only the IPFIX format is supported for this release.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

format (Next Gen Services Service-Set Remote System Logging)

IN THIS SECTION

- Syntax | 7385
- Hierarchy Level | 7385
- Description | 7385
- Options | 7385
- Required Privilege Level | 7386
- Release Information | 7386

Syntax

```
format format;
```

Hierarchy Level

```
edit services service-set name syslog stream stream-name
```

Description

Specify the file format for the log messages being sent to the remote server.

Options

The file format can be one of the following:

binary Binary syslog defined by Juniper Networks. Requires Juniper Networks decoders on the server side to decode the logs.

sd-syslog Structured syslog (defined by RFC5424)

syslog Traditional syslog (defined by RFC5424)

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

format (Security Log)

IN THIS SECTION

- Syntax | 7386
- Hierarchy Level | 7387
- Description | 7387
- Options | 7387
- Required Privilege Level | 7387
- Release Information | 7387

Syntax

```
format (binary | sd-syslog | syslog)
```

Hierarchy Level

```
[edit security log]
```

Description

Set the default log format for event mode security logging on the device.

Options

- `binary`—Binary encoded text to conserve resources.
- `sd-syslog`—Structured system log file.
- `syslog`—Traditional system log file.
- **Default:** syslog.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced prior to Junos OS Release 10.0. Statement updated in Junos OS Release 12.1.

RELATED DOCUMENTATION

log (Security) | **10026**

format (Security Log Stream)

IN THIS SECTION

- Syntax | 7388
- Hierarchy Level | 7388
- Description | 7388
- Options | 7388
- Required Privilege Level | 7389
- Release Information | 7389

Syntax

```
format (binary | sd-syslog | syslog | welf)
```

Hierarchy Level

```
[edit security log stream stream-name]
[edit logical-systems name security log stream stream-name]
[edit tenants tenant-name security log stream stream-name]
```

Description

Set the format for remote security message logging to `binary`, `syslog` (system log), `sd-syslog` (structured system log), or `welf`. Note that for the WELF format, the category must be set to `content-security` (see "category" on page 3224).

Options

- `binary`—Binary encoded text to conserve resources.
- `sd-syslog`—Structured system log file.
- `syslog`—Traditional system log file.

- `welf`—Web Trends Extended Log Format.
- **Default:** By default `syslog` (system log) is enabled.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0 . Updated in Junos OS Release 12.1 .

The [edit `logical-systems` *name* security log stream *stream-name*] hierarchy level introduced in Junos OS Release 18.2R1.

The [edit `tenants` *tenant-name* security log stream *stream-name*] hierarchy level introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Application Security User Guide for Security Devices

Logical Systems and Tenant Systems User Guide for Security Devices

format (Unified Edge Gateways)

IN THIS SECTION

- Syntax | 7390
- Hierarchy Level | 7391
- Description | 7391
- Required Privilege Level | 7391
- Release Information | 7391

Syntax

```
format {
    integer {
        equals {
            value;
        }
        greater-than value;
        less-than value;
    }
    string {
        equals {
            value;
        }
        has-prefix{
            value;
        }
        has-suffix {
            value;
        }
        matches {
            value;
        }
    }
    time {
        equals {
            value;
        }
        greater-than value;
        less-than value;
    }
    v4address {
        equals {
            value;
        }
    }
    v6address {
        equals {
            value;
        }
    }
    v6prefix {
```

```
        equals {
            value;
        }
    }
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name]
```

Description

Specify the custom AVP attribute's format and value to match for the incoming RADIUS request from the IP-based subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

forward-and-send-to-re

IN THIS SECTION

- Syntax | 7392
- Hierarchy Level | 7392
- Description | 7392
- Required Privilege Level | 7392
- Release Information | 7392

Syntax

```
forward-and-send-to-re;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet targeted-broadcast],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet targeted-broadcast]
```

Description

Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface and the Routing Engine. The packets are broadcast only if the egress interface is a LAN interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Targeted Broadcast

targeted-broadcast | **28956**

Understanding Targeted Broadcast

forward-delay

IN THIS SECTION

- Syntax | **7393**
- Hierarchy Level | **7393**
- Description | **7394**
- Options | **7394**
- Required Privilege Level | **7394**
- Release Information | **7394**

Syntax

```
forward-delay seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp)],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp)],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id],
```

```
[edit protocols (mstp | rstp)],
[edit protocols vstp vlan vlan-id],
[edit routing-instances routing-instance-name protocols (mstp | rstp)],
```

```
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id]
```

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specifies the length of time an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.

Options

seconds—(Optional) Number of seconds the bridge port remains in the listening and learning states.

- **Range:** 4 through 30
- **Default:** 15 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Forward Delay Before Ports Transition to Forwarding State

Understanding STP

Understanding MSTP

Understanding VSTP

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

show spanning-tree bridge | 26548

show spanning-tree interface | 26555

forward-only

IN THIS SECTION

- Syntax | 7395
- Hierarchy Level | 7395
- Description | 7396
- Required Privilege Level | 7396
- Release Information | 7396

Syntax

```
forward-only;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet targeted-broadcast],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet targeted-broadcast]
```

Description

Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Targeted Broadcast

targeted-broadcast | **28956**

Understanding Targeted Broadcast

forward-only (DHCP Relay Agent)

IN THIS SECTION

- Syntax | **7397**
- Hierarchy Level | **7397**
- Description | **7397**
- Default | **7397**
- Options | **7397**
- Required Privilege Level | **7398**
- Release Information | **7398**

Syntax

```
forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify the VRF location of the DHCP server when configuring secure DHCP traffic between the DHCP server and DHCP client when the two reside in different VRFs.

Default

Logical system and routing instance from where the configuration is applied.

Options

logical-system (Optional) Logical system in which the DHCP server resides.

- `current`—Logical system from which the configuration is applied.
- `default`—Root logical system.
- *`logical-system-name`*—A specific logical system.

routing-instance (Optional) Routing instance in which the DHCP server resides.

- `current`—Routing instance from which the configuration is applied.
- `default`—Root routing instance.
- *`logical-system-name`*—A specific routing instance.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 13.3R3, 14.1R2, and 14.2R1.

RELATED DOCUMENTATION

DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs

Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances

forward-only (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | **7399**
- Hierarchy Level | **7399**
- Description | **7399**
- Required Privilege Level | **7399**
- Release Information | **7399**

Syntax

```
forward-only;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action | equals | starts-with)],
[edit forwarding-options dhcp-relay group group-name relay-option (default-action | equals | starts-
with)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name relay-option (default-action | equals
| starts-with)],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Forward specified DHCP client packets, without creating a new subscriber session, when you use DHCP relay selective processing feature. You can configure the forwarding operation globally or for a group of interfaces, and for either DHCP or DHCPv6 relay agent. Use this command to manage subscribers, when large number of DHCP relay requests are received and license limitations are exceeded. This command uses option 82 and applies "scale-subscriber" license.

NOTE: DHCP packets forwarded with the `forward-only` statement do not consider other configurations except for the `trust-option-82` option. The DHCP relay agent ignores all other configured options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic

forward-only (DNS)

IN THIS SECTION

- Syntax | 7400
- Hierarchy Level | 7400
- Description | 7400
- Required Privilege Level | 7400
- Release Information | 7401

Syntax

```
forward-only;
```

Hierarchy Level

```
[edit system services dns dns-proxy view view-name domain domain-name]
```

Description

Specify that the server to forward only DNS queries. This configuration prevents the device from acquiring public IP addresses, in case the IP address specified in `forwarders` option is not reachable, by terminating the DNS query.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

DNS Overview

forward-only-replies (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 7401
- Hierarchy Level | 7401
- Description | 7402
- Required Privilege Level | 7402
- Release Information | 7402

Syntax

```
forward-only-replies;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify that reply packets are for the forward-only support that is configured in option 82 interface ID of the reply packet. You must configure this statement for forward-only support when the client and server are in different logical system/routing instances.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 13.3R3, 14.1R2, and 14.2R1.

RELATED DOCUMENTATION

DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs

Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances

forward-policy-mismatch (Security Group VPN Member)

IN THIS SECTION

- Syntax | 7403
- Hierarchy Level | 7403
- Description | 7403
- Required Privilege Level | 7403
- Release Information | 7403

Syntax

```
vpn vpn-name {
    ike-gateway gateway-number;
    group group-number;
    match-direction (input);
    tunnel mtu mtu-size;
}
```

Hierarchy Level

```
[edit security group-vpn member ipsec]
```

Description

Configure the forward policy mismatch for group VPN member.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

- Group VPNv2 Overview

forward-snooped-clients (DHCP Local Server)

IN THIS SECTION

- Syntax | 7404
- Hierarchy Level | 7404
- Description | 7404
- Options | 7404
- Required Privilege Level | 7405
- Release Information | 7405

Syntax

```
forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server],
[edit logical-systems logical-system-name system services dhcp-local-server],
[edit routing-instances routing-instance-name system services dhcp-local-server],
[edit system services dhcp-local-server]
```

Description

Configure how the DHCP local server filters and handles DHCP snooped packets on the specified interfaces.

Options

`all-interfaces`—Perform the action on all interfaces.

`configured-interfaces`—Perform the action only on interfaces that are configured as part of an interface group.

`non-configured-interfaces`—Perform the action only on interfaces that are not configured as part of a group.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

DHCP Snooping Support

Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server

forward-snooped-clients (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 7405
- Hierarchy Level | 7406
- Description | 7406
- Options | 7406
- Required Privilege Level | 7406
- Release Information | 7406

Syntax

```
forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay]
```

Description

Configure how DHCP relay agent filters and handles DHCP snooped packets on the specified interfaces. The router or switch determines the DHCP snooping action to perform based on a combination of the `forward-snooped-clients` configuration and the configuration of either the `allow-snooped-clients` statement or the `no-allow-snooped-clients` statement.

The router (or switch) also uses this statement to determine how to handle snooped BOOTREPLY packets received on non-configured interfaces.

Options

`all-interfaces`—Perform the action on all interfaces.

- **Default:** On EX Series switches, the action is performed on all interfaces by default.

`configured-interfaces`—Perform the action only on interfaces that are configured as part of an interface group.

`non-configured-interfaces`—Perform the action only on interfaces that are not part of a group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the [`edit forwarding-options dhcp-relay dhcpv6`] hierarchy level introduced in Junos OS Release 15.1X53-D56 for EX Series switches and Junos OS Release 17.1R1.

RELATED DOCUMENTATION

- DHCP Snooping Support
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent
- *DHCP Snooping for Network Security*

forwarding-cache (Bridge Domains)

IN THIS SECTION

- Syntax | 7407
- Hierarchy Level | 7407
- Description | 7408
- Options | 7408
- Required Privilege Level | 7408
- Release Information | 7408

Syntax

```
forwarding-cache {
    threshold suppress value <reuse value>;
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name multicast-snooping-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name multicast-snooping-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name multicast-
snooping-options],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name multicast-snooping-
options],
[edit routing-instances routing-instance-name multicast-snooping-options]
```

Description

Establish multicast snooping forwarding cache parameter values.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring Multicast Snooping

forwarding-cache (Flow Maps)

IN THIS SECTION

- Syntax | 7409
- Hierarchy Level | 7409
- Description | 7409
- Required Privilege Level | 7409
- Release Information | 7409

Syntax

```
forwarding-cache {
    timeout (minutes | never non-discard-entry-only );
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast flow-map flow-map-name],
[edit logical-systems logical-system-name routing-options multicast flow-map flow-map-name],
[edit routing-instances routing-instance-name routing-options multicast flow-map flow-map-name],
[edit routing-options multicast flow-map flow-map-name]
```

Description

Configure multicast forwarding cache properties for the flow map.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Example: Configuring a Multicast Flow Map

forwarding-class

IN THIS SECTION

- Classifier | 7410
- Classifier (EX Series switches) | 7411
- Rewrite Rule | 7411
- Scheduler Map | 7411
- Interface | 7411
- Classifier Hierarchy Level | 7411
- Classifier Hierarchy Level (EX Series Switches) | 7411
- Rewrite Rule Hierarchy Level | 7412
- Rewrite Rule Hierarchy Level (EX Series Switches) | 7412
- Scheduler Map Hierarchy Level | 7412
- Interface Hierarchy Level | 7412
- Host-Outbound Traffic Hierarchy Level (EX Series Switches) | 7412
- Description | 7412
- Options | 7413
- Required Privilege Level | 7413
- Release Information | 7413

Classifier

```
forwarding-class class-name {
    loss-priority level {
        code-points [ aliases ] [ bit-patterns ];
    }
}
```

Classifier (EX Series switches)

```
forwarding-class class-name {
    loss-priority level {
        code-points [aliases] [6-bit-patterns];
    }
}
```

Rewrite Rule

```
forwarding-class class-name {
    loss-priority level {
        code-point [ aliases ] [ bit-patterns ];
    }
}
```

Scheduler Map

```
forwarding-class class-name {
    scheduler scheduler-name;
}
```

Interface

```
forwarding-class class-name;
```

Classifier Hierarchy Level

```
[edit class-of-service classifiers (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name]
```

Classifier Hierarchy Level (EX Series Switches)

```
[edit class-of-service classifiers (dscp | ieee-802.1 | inet-precedence) classifier-name]
```

Rewrite Rule Hierarchy Level

```
[edit class-of-service rewrite-rules] (dscp | dscp-ipv6 | ieee-802.1) rewrite-name | exp]
```

Rewrite Rule Hierarchy Level (EX Series Switches)

```
[edit class-of-service rewrite-rules] (dscp | ieee-802.1 | inet-precedence) rewrite-rule-name]
```

Scheduler Map Hierarchy Level

```
[edit class-of-service scheduler-maps map-name]
```

Interface Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Host-Outbound Traffic Hierarchy Level (EX Series Switches)

```
[edit class-of-service host-outbound-traffic]
```

Description

Define forwarding class name and option values.

- Classifiers—Assign incoming traffic to the specified forwarding class based on the specified code point values and assign that traffic the specified loss priority.
- Rewrite rules—At the egress interface, change (rewrite) the value of the code point bits and the loss priority to specified new values for traffic assigned to the specified forwarding class, before forwarding the traffic to the next hop.
- Scheduler maps—Apply the specified scheduler to the specified forwarding class.
- Interfaces—Assign the specified forwarding class to the interface to use as a fixed classifier (all incoming traffic on the interface is classified into that forwarding class).

Options

class-name—Name of the forwarding class.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- *Example: Configuring CoS on EX Series Switches*
- *Defining CoS Forwarding Classes (CLI Procedure)*
- *Defining CoS Forwarding Classes (J-Web Procedure)*
- *Understanding CoS Forwarding Classes*

forwarding-class-accounting

IN THIS SECTION

- (Forwarding Class Accounting) | 7414
- (Enhanced CoS Accounting) | 7414
- Hierarchy Level | 7414
- Description | 7414
- Required Privilege Level | 7415
- Release Information | 7415

(Forwarding Class Accounting)

```
forwarding-class-accounting {
    direction
}
```

(Enhanced CoS Accounting)

```
forwarding-class-accounting {
    enhanced
}
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number]
```

Description

Forwarding class accounting (non-enhanced mode) is enabled for IPv4, IPv6, MPLS, Layer 2 and other packets simply by enabling the feature.

> **NOTE**: Forwarding class accounting (enhanced mode) is not supported in "hyper-mode (forwarding-options)" on page 8221.

Both bytes and packet total are counted. For Layer 2 traffic, only bytes are counted. Flow rates are measured. Counters can be enabled in ingress, egress or in both directions. Statistics are accounted before queueing drop stage, so will include packets that might get dropped in queueing.

CoS-based interface counters (enhanced mode) are supported for IPv4 and IPv6 aggregate routes on Layer 2.

Both bytes and packet total are counted. Flow rates are not measured. The counters can be configured to exclude overhead bytes (such as protocol encapsulations) so end-customer packets can be differentiated from other traffic.

Disable or delete `forwarding-class-accounting` by removing the statement from the interface configuration–for example, by typing `delete interfaces` *`interface-name`* `forwarding-class-accounting`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Forwarding class accounting statement (non-enhanced mode) introduced in Junos OS Release 13.3R3.

Enhanced CoS Accounting statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

CoS-Based Interface Counters for IPv4 or IPv6 Aggregate on Layer 2

show class-of-service interface | **19056**

clear interfaces statistics | **3651**

show interfaces forwarding-class-counters | **20859**

forwarding-class (BA Classifiers)

IN THIS SECTION

- Syntax | **7416**
- Hierarchy Level | **7416**
- Description | **7416**
- Options | **7416**
- Required Privilege Level | **7416**
- Release Information | **7416**

Syntax

```
forwarding-class class-name {
    loss-priority level code-points [ aliases ] [ bit-patterns ];
}
```

Hierarchy Level

```
[edit class-of-service classifiers type classifier-name]
```

Description

Define forwarding class name and option values.

Options

class-name—Name of the forwarding class.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Behavior Aggregate Classifiers

forwarding-class (CoS Host Outbound Traffic)

IN THIS SECTION

- Syntax | 7417
- Hierarchy Level | 7417
- Description | 7417
- Default | 7418
- Options | 7418
- Required Privilege Level | 7418
- Release Information | 7418

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit class-of-service host-outbound-traffic]
```

Description

Specify the name of the forwarding class to which host outbound traffic is assigned on all egress interfaces. The output queue associated with the forwarding class must be properly configured on all interfaces. In the case of a restricted interface, the traffic flows through a restricted queue.

For egress interfaces hosted on MX Series routers, M120 routers, or Enhanced III FPCs in M320 routers, both Routing Engine sourced traffic and distributed protocol handler traffic are affected. For all other egress interfaces, only Routing Engine sourced traffic is affected.

This statement does not affect transit traffic or incoming traffic.

Default

If you do not configure an output queue for host outbound traffic, the router uses the default queue assignments for host outbound traffic.

Options

class-name—Name of the forwarding class.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for distributed protocol handler traffic introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding How Forwarding Classes Assign Classes to Output Queues

Default Routing Engine Protocol Queue Assignments

Changing the Default Queuing and Marking of Host Outbound Traffic

forwarding-class-default (Forwarding Policy)

IN THIS SECTION

- Syntax | 7419
- Hierarchy Level | 7419
- Description | 7419
- Options | 7419
- Required Privilege Level | 7419

Release Information | 7420

Syntax

```
forwarding-class-default class-name {
    discard;
    lsp-next-hop [ lsp-regular-expression ];
    next-hop [ next-hop-name];
    non-lsp-next-hop;
    match-next-hop-forwarding-class;
}
```

Hierarchy Level

```
[edit class-of-service forwarding-policy next-hop-map map-name]
```

Description

Define the next hop for traffic that does not meet any forwarding class in the next-hop map.

Options

class-name—Name of the forwarding class.

`match-next-hop-forwarding-class`—Match any pre-defined forwarding-class or a forwarding-class that is configured using CLI.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced for QFX10000 Series switches in Junos OS Release 17.1R1.

`match-next-hop-forwarding-class` option introduced in Junos OS Evolved Release 21.3R1 for PTX10008 platform.

RELATED DOCUMENTATION

forwarding-class (Forwarding Policy) | 7426

Load Balancing VPLS Non-Unicast Traffic Across Member Links of an Aggregate Interface

forwarding-class (Dynamic Scheduler Maps)

IN THIS SECTION

- Syntax | 7420
- Hierarchy Level | 7420
- Description | 7421
- Options | 7421
- Required Privilege Level | 7421
- Release Information | 7421

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service scheduler-maps map-name]
```

Description

Associate a scheduler with a scheduler map.

Options

class-name—Name of the forwarding class.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Schedulers in a Dynamic Profile for Subscriber Access

forwarding-class (Fibre Channel Interfaces)

IN THIS SECTION

- Syntax | 7422
- Hierarchy Level | 7422
- Description | 7422
- Options | 7423
- Required Privilege Level | 7423
- Release Information | 7423

Syntax

```
forwarding-class lossless-forwarding-class-name;
```

Hierarchy Level

```
[edit class-of-service interfaces fibre-channel-interface-name]
```

Description

Configure a Layer 3 fixed classifier on a Fibre Channel (FC) interface. The fixed classifier places all traffic received from the FC network into the specified forwarding class. The forwarding class must be lossless. (That is, the forwarding class must be either the default `fcoe` or `no-loss` forwarding class, or the forwarding class must be configured with the `no-loss` drop attribute.) If you attempt to specify a lossy forwarding class, the system returns a commit error.

FCoE networks typically use priority 3 (IEEE code point 011) for FCoE traffic. The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. If the IEEE 802.1p code point value that the Ethernet network uses for FCoE traffic is different than code point 3, you can rewrite the code point to the value used in your Ethernet (FCoE) network. The lossless forwarding class specified in the fixed classifier uses the `rewrite-value` statement as the IEEE 802.1p code point (priority) for FCoE traffic on the FCoE network.

To rewrite the code point value, include the `rewrite-value input ieee code-point` *`code-point-bits`* statement at the [`edit class-of-service interfaces` *`fc-interface-name`*] hierarchy level.

NOTE: If you are not using the default configuration (priority 3 for FCoE traffic), the lossless forwarding class specified in the FC interface fixed classifier must be mapped to the IEEE 802.1p code point specified in the rewrite value statement.

In order to avoid fate sharing (separate flows that affect each other's throughput), the code point (priority) used for the lossless forwarding class (the code point specified in the rewrite value statement) should be the only code point classified to that forwarding class (at the [`edit class-of-service classifiers`] hierarchy level). For example, if the rewrite value uses code point 101 for lossless FCoE forwarding class fcoe_fc1, then in the classifier configuration attached to ingress Ethernet interfaces, code point 101 is the only code point that should be classified to the fcoe_fc1 forwarding class. Now if you also attach a classifier to an interface that maps code point 110 to forwarding class fcoe_fc1, then

congestion on priority 110 unfairly (and unintentionally) affects the FCoE traffic that uses priority 101. Both priorities 101 and 110 are classified into forwarding class fcoe_fc1, so the traffic from both priorities shares the same fate.

Options

lossless-forwarding-class-name—Name of the lossless forwarding class.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

Understanding CoS Classifiers

forwarding-class (Firewall Filter Action)

IN THIS SECTION

- Syntax | 7424
- Hierarchy Level | 7424
- Description | 7424
- Options | 7424
- Required Privilege Level | 7424
- Release Information | 7424

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name then],
[edit logical-systems logical-system-name firewall family family-name filter filter-name term term-
name then]
```

Description

Set the forwarding class of incoming packets.

Options

class-name—Name of the forwarding class.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Firewall Filter Nonterminating Actions*
>
> *Policer Color-Marking and Actions*
>
> *Multifield Classification Overview*

forwarding-class (for DHCP Snooping or DAI Packets)

IN THIS SECTION

- Syntax | 7425
- Hierarchy Level | 7425
- Description | 7425
- Default | 7426
- Options | 7426
- Required Privilege Level | 7426
- Release Information | 7426

Syntax

```
forwarding-class class class-name;
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name) (examine-dhcp | arp-
inspection)]
```

Description

Assign a user-defined or a predefined forwarding class to the packets that have been checked for DHCP snooping or dynamic ARP inspection (DAI).

NOTE: To assign a user-defined class, you must first configure the user-defined class by using the "forwarding-classes" on page 7447 configuration statement at the [edit *class-of-service*] hierarchy level.

Default

Disabled.

Options

class-name—Name of the forwarding class. The forwarding class can be one of the predefined forwarding classes (best-effort, assured-forwarding, expedited-forwarding, network-control) or it can be a user-defined forwarding class.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Example: Prioritizing Snooped and Inspected Packet

Understanding Junos OS CoS Components for EX Series Switches

Understanding DHCP Snooping (non-ELS)

Understanding ARP Spoofing and Inspection

forwarding-class (Forwarding Policy)

IN THIS SECTION

- Syntax | 7427
- Hierarchy Level | 7427
- Description | 7427
- Options | 7427

Required Privilege Level | 7428
Release Information | 7428

Syntax

```
forwarding-class class-name {
    discard;
    lsp-next-hop [ lsp-regular-expression ];
    next-hop [ next-hop-name];
    non-labelled-next-hop;
    non-lsp-next-hop;
    match-next-hop-forwarding-class;
    transport-class {
        color color;
    }
}
```

Hierarchy Level

```
[edit class-of-service forwarding-policy next-hop-map map-name]
[edit class-of-service forwarding-policy class class-name classification-override]
```

Description

Define forwarding class name and associated next hops.

Options

class-name	Name of the forwarding class.
non-labelled-next-hop	Match any non-labelled next hop.
match-next-hop-forwarding-class	Match any pre-defined forwarding-class or a forwarding-class that is configured using CLI.

transport-class color *color* Match transport-class key in FRR indirect keys structure.

> **NOTE**: This option is available only at [edit class-of-service forwarding-policy next-hop-map *next-hop-map-name* forwarding-class *forwarding-class-name*] hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for QFX10000 Series switches in Junos OS Release 17.1R1.

`non-labelled-next-hop` option introduced in Junos OS Release 19.1R1 for all platforms.

`match-next-hop-forwarding-class` option introduced in Junos OS Evolved Release 21.3R1 for PTX10008 platform.

RELATED DOCUMENTATION

Overriding the Input Classification

forwarding-class-default (Forwarding Policy) | **7418**

forwarding-class (Fragmentation)

IN THIS SECTION

- Syntax | **7429**
- Hierarchy Level | **7429**

- Description | 7429
- Default | 7429
- Options | 7429
- Required Privilege Level | 7430
- Release Information | 7430

Syntax

```
forwarding-class class-name {
    drop-timeout milliseconds;
    fragment-threshold bytes;
    multilink-class number;
    no-fragmentation;
}
```

Hierarchy Level

```
[edit class-of-service fragmentation-maps map-name]
```

Description

For Multiservices and Services PICs link services IQ interfaces (`lsq`) only, define a forwarding class name and associated fragmentation properties within a fragmentation map.

The `fragment-threshold` and `no-fragmentation` statements are mutually exclusive.

Default

If you do not include this statement, the traffic in forwarding class *`class-name`* is fragmented.

Options

`class-name`—Name of the forwarding class.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Fragmentation by Forwarding Class

forwarding-class (Gigabit Ethernet IQ Classifier)

IN THIS SECTION

- Syntax | 7430
- Hierarchy Level | 7431
- Description | 7431
- Options | 7431
- Required Privilege Level | 7431
- Release Information | 7431

Syntax

```
forwarding-class class-name {
    loss-priority (high | low);
}
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile
output-priority-map classifier premium]
```

Description

For Gigabit Ethernet IQ interfaces only, define forwarding class name and option values.

Options

class-name—Name of forwarding class.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

input-priority-map | **8716**

Junos OS Class of Service User Guide for Routing Devices

forwarding-class (Interfaces)

IN THIS SECTION

- Syntax | 7432
- Hierarchy Level | 7432
- Description | 7432
- Options | 7432
- Required Privilege Level | 7432
- Release Information | 7433

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Associate a forwarding class configuration or default mapping with a specific interface.

Options

class-name—Name of the forwarding class.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Classes to Interfaces

forwarding-class (Multifield Classifiers)

IN THIS SECTION

- Syntax | 7433
- Hierarchy Level | 7433
- Description | 7433
- Options | 7434
- Required Privilege Level | 7434
- Release Information | 7434

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name then]
```

Description

Set the forwarding class of incoming packets.

Options

class-name—Name of the forwarding class.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Multifield Classifiers

forwarding-class (PCC Action Profiles)

IN THIS SECTION

- Syntax | 7434
- Hierarchy Level | 7435
- Description | 7435
- Options | 7435
- Required Privilege Level | 7435
- Release Information | 7435

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],
[edit services pcef pcc-action-profiles profile-name]
```

Description

Specify the forwarding class to which packets must be assigned.

If you are using Junos OS Subscriber Aware, specify the forwarding class at the `[edit unified-edge pcef pcc-action-profiles profile-name]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the forwarding class at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level.

Options

class-name	Name of the forwarding class.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

- Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware
- Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment
- *Understanding Application-Aware Policy Control for Subscriber Management*
- *Configuring Policy and Charging Control Action Profiles for Subscriber Management*

forwarding-class (Restricted Queues)

IN THIS SECTION

- Syntax | 7436
- Hierarchy Level | 7436
- Description | 7436
- Options | 7437
- Required Privilege Level | 7437
- Release Information | 7437

Syntax

```
forwarding-class class-name queue queue-number;
```

Hierarchy Level

```
[edit class-of-service restricted-queues]
```

Description

For M320 and T Series routers only, map forwarding classes to restricted queues. You can map up to eight forwarding classes to restricted queues.

Options

class-name—Name of the forwarding class.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

forwarding-class (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 7437
- Hierarchy Level | 7438
- Description | 7438
- Options | 7438
- Required Privilege Level | 7438
- Release Information | 7438

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specify the forwarding class to be used for test frames. The forwarding class specifies the manner in which the test frames are processed by the Packet Forwarding Engine of the router. If you do not configure this parameter, test frames are treated as best-effort traffic.

Options

class-name Name of the forwarding class. You must have previously configured this forwarding class by including the `forwarding-class` statement at the `[edit class-of-service interfaces interface-name unit logical-unit-number]` hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

forwarding-class (Sampling)

IN THIS SECTION

- Syntax | 7439
- Hierarchy Level | 7439
- Description | 7439
- Default | 7439
- Options | 7439
- Required Privilege Level | 7440
- Release Information | 7440

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet | inet6) output flow-
server hostname]
```

Description

Specify the forwarding class to which exported packets for inline active flow monitoring are sent.

Default

If you do not include the *forwarding-class* statement, exported packets are sent to the best effort queue.

Options

forwarding-class *class-name*

Name of the forwarding class:

- `assured-forwarding`
- `best-effort`
- `expedited-forwarding`
- `network-control`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches

forwarding-class (Services CoS Fragmentation Properties)

IN THIS SECTION

- Syntax | 7441
- Hierarchy Level | 7441
- Description | 7441
- Default | 7441
- Options | 7441
- Required Privilege Level | 7441
- Release Information | 7441

Syntax

```
forwarding-class class-name {
    (fragment-threshold bytes | no-fragmentation);
    multilink-class number;
}
```

Hierarchy Level

```
[edit class-of-service fragmentation-maps]
```

Description

For link services IQ (`lsq`) interfaces only, define a forwarding class name and associated fragmentation properties within a fragmentation map.

The `fragment-threshold` and `no-fragmentation` statements are mutually exclusive.

Default

If you do not include this statement, the traffic in forwarding class *`class-name`* is fragmented.

Options

`class-name`—Name of the forwarding class.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces

forwarding-class (Services PIC Classifiers)

IN THIS SECTION

- Syntax | 7442
- Hierarchy Level | 7442
- Description | 7442
- Options | 7442
- Required Privilege Level | 7443
- Release Information | 7443

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit services cos application-profile profile-name (ftp | sip) (data | video | voice)],
[edit services cos rule rule-name term term-name then],
[edit services cos rule rule-name term term-name then reflexive | revert | reverse]
```

Description

Define the forwarding class to which packets are assigned.

Options

class-name—Name of the target application.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

forwarding-class-set

IN THIS SECTION

- Syntax | 7443
- Hierarchy Level | 7444
- Description | 7444
- Options | 7444
- Required Privilege Level | 7444
- Release Information | 7444

Syntax

```
forwarding-class-set forwarding-class-set-name {
    output-traffic-control-profile profile-name;
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name]
```

Description

Apply a previously defined forwarding class set to an output traffic control profile.

Options

fowarding-class-set-name—Name of the forwarding class set.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Assigning CoS Components to Interfaces

Understanding CoS Forwarding Class Sets (Priority Groups)

forwarding-class-set (Interconnect Device)

IN THIS SECTION

Syntax | 7445

- Hierarchy Level | 7445
- Description | 7445
- Options | 7445
- Required Privilege Level | 7445
- Release Information | 7445

Syntax

```
forwarding-class-set fabric-forwarding-class-set-name {
    scheduler scheduler-name;
}
```

Hierarchy Level

```
[edit class-of-service scheduler-map-forwarding-class-sets scheduler-map-fcset-name]
```

Description

Specify an existing fabric forwarding class set to map to a scheduler. The scheduler provides bandwidth and drop profile characteristics for all of the forwarding classes assigned to the forwarding class set.

Options

fabric-fowarding-class-set-name—Name of the fabric forwarding class set.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1R2.

RELATED DOCUMENTATION

Example: Configuring CoS Scheduling Across the QFabric System

Understanding CoS Fabric Forwarding Class Sets

Understanding CoS Scheduling Across the QFabric System

forwarding-class-sets

IN THIS SECTION

- Syntax | 7446
- Hierarchy Level | 7446
- Description | 7446
- Options | 7447
- Required Privilege Level | 7447
- Release Information | 7447

Syntax

```
forwarding-class-sets forwarding-class-set-name {
    class class-name;
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Assign forwarding classes to forwarding class sets (priority groups).

Options

fowarding-class-set-name—Name of the forwarding class set.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Forwarding Class Sets

Understanding CoS Forwarding Class Sets (Priority Groups)

forwarding-classes

IN THIS SECTION

- SRX Series | **7448**
- QFX Series | **7448**
- EX Series (Except EX4300) | **7448**
- EX4300 | **7449**
- M320, MX Series, T Series, and PTX Series | **7449**
- Hierarchy Level | **7449**
- Description | **7449**
- Options | **7451**
- Required Privilege Level | **7452**

Release Information | 7452

SRX Series

```
forwarding-classes {
    class class-name {
        priority (high | low);
        queue-num number;
        spu-priority (high | low | medium);
    }
    queue queue-number {
        class class-name {
            priority (high | low);
        }
    }
}
```

QFX Series

```
forwarding-classes {
    class class-name {
        pfc-priority pfc-priority;
        no-loss;
        queue-num queue-number <no-loss>;
    }
}
```

EX Series (Except EX4300)

```
forwarding-classes {
    class class-name {
        queue-num queue-number;
        priority (high | low);
    }
}
```

EX4300

```
forwarding-classes {
    class class-name ;
        queue-num queue-number;
    }
}
```

M320, MX Series, T Series, and PTX Series

```
forwarding-classes {
    class class-name {
        queue queue-number;
        priority (high | low);
    }
    queue queue-number {
        class class-name {
            priority (high | low) [policing-priority (premium | normal)];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Command used to associate forwarding classes with class names and queues with queue numbers.

SRX Series Firewalls

All traffic traversing the SRX Series Firewall is passed to an SPC to have service processing applied. Junos OS provides a configuration option to enable packets with specific Differentiated Services (DiffServ) code points (DSCP) precedence bits to enter a high-priority queue, a medium-priority queue, or a low-priority queue on the SPC. The Services Processing Unit (SPU) draws packets from the highest priority queue first, then from the medium priority queue, and last from the low priority queue. The

processing of the queue is weighted-based not strict-priority-based. This feature can reduce overall latency for real-time traffic, such as voice traffic.

Initially, the spu-priority queue options were "high" and "low". Then, these options (depending on the devices) were expanded to "high", "medium-high", "medium-low", and "low". The two middle options ("medium-high" and "medium-low") have now been deprecated (again, depending on the devices) and replaced with "medium". So, the available options for spu-priority queue are "high", "medium", and "low".

We recommend that the high-priority queue be selected for real-time and high-value traffic. The other options would be selected based on user judgement on the value or sensitivity of the traffic.

M320, MX Series, and T Series Routers and EX Series Switches

For M320, MX Series, and T Series routers, and EX Series switches only, you can configure fabric priority queuing by including the `priority` statement. For Enhanced IQ PICs, you can include the `policing-priority` option.

> **NOTE**: The `priority` and `policing-priority` options are not supported on PTX Series routers.

EX Series Switches

For the EX Series switches, this statement associates the forwarding class with a class name and queue number. It can define the fabric queuing priority as high, medium-high, medium-low, or low.

Map one or more forwarding classes to a single output queue. Also, when configuring DSCP-based priority-based flow control (PFC), map a forwarding class to a PFC priority value to use in pause frames when traffic on a DSCP value becomes congested (see *Configuring DSCP-based PFC for Layer 3 Untagged Traffic* for details).

Switches that use different forwarding classes for unicast and multidestination (multicast, broadcast, and destination lookup fail) traffic support 12 forwarding classes and 12 output queues (0 through 11). You map unicast forwarding classes to a unicast queue (0 through 7) and multidestination forwarding classes to a multidestination queue (8 through 11). The queue to which you map a forwarding class determines if the forwarding class is a unicast or multidestination forwarding class.

Switches that use the same forwarding classes for unicast and multidestination traffic support eight forwarding classes and eight output queues (0 through 7). You map forwarding classes to output queues. All traffic classified into one forwarding class (unicast and multidestination) uses the same output queue.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the `no-loss` packet drop attribute. Do not associate a drop profile with lossless forwarding classes.

NOTE: If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

NOTE: On switches that do not use the Enhanced Layer 2 Software (ELS) CLI, if you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless `fcoe` and `no-loss` forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

NOTE: On switches that do not use the ELS CLI, if you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the `fcoe` and `no-loss` forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the `no-loss` option. If you do not specify the `no-loss` option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the `fcoe` forwarding class and you do not include the `no-loss` option, the `fcoe` forwarding class is lossy, not lossless.

Options

class *class-name*
: Define the forwarding class name.

queue-num *queue-number*
: Output queue number to associate with forwarding class.

- **Range:** 0 through 7.

priority
: Fabric priority value:

  high Forwarding class fabric queuing has high priority.

  low Forwarding class fabric queuing has low priority.

  The `default` priority is `low`.

spu-priority SPU priority queue, `high`, `medium`, or `low`. The default `spu-priority` is `low`.

> **NOTE**: The `spu-priority` option is supported only on the SRX5000 line of firewalls.

The remaining statements are explained separately. See CLI Explorer for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

The `policing-priority` option was introduced in Junos OS Release 9.5.

Statement updated in Junos OS Release 11.4.

The `spu-priority` option was introduced in Junos OS Release 11.4R2.

The `no-loss` option was introduced in Junos OS Release 12.3.

Change from two to four queues made in Junos OS Release 12.3X48-D40 and in Junos OS Release 15.1X49-D70.

The `pfc-priority` statement was introduced in Junos OS Release 17.4R1.

The `medium-high` and `medium-low` priorities for `spu-priority` were deprecated and `medium` priority was added in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Configuring a Custom Forwarding Class for Each Queue

Forwarding Classes and Fabric Priority Queues

Configuring Hierarchical Layer 2 Policers on IQE PICs

Classifying Packets by Egress Interface

forwarding-context (Protocols BGP)

IN THIS SECTION

- Syntax | 7453
- Hierarchy Level | 7453
- Description | 7453
- Required Privilege Level | 7454
- Release Information | 7454

Syntax

```
forwarding-context;
```

Hierarchy Level

```
[edit logical-systems name protocols bgp],
[edit logical-systems name routing-instances name
protocols bgp],
[edit logical-systems name tenants name routing-instances name
protocols bgp],
[edit protocols bgp],
[edit routing-instances name protocols bgp],
[edit tenants name routing-instances name protocols bgp]
```

Description

The MPLS-forwarding type routing-instance can be used for segregating Inter-AS BGP neighbors that require MPLS spoof-protection to ensure the packets remain distinct from other peers.

Setting a forwarding context on a neighbor interface can be useful, for example, when configuring a common AS boundary router so that it only accepts MPLS packets from a peer AS boundary router whose labels were explicitly advertised to the common AS boundary router.

Use this statement in conjunction with `mpls-forwarding` to protect against label spoofing across AS boundary routers in the context of Inter-AS VPN Option B for AS boundary routers. Option B peers are reachable thru local interfaces that are configured as part of the MPLS forwarding type routing instance.

If `forwarding-context` is not set for a VPN BGP peer both the routing instance and forwarding context are provided by the master routing instance. The master instance is the Junos default, global routing-instance, that contains the **protocols bgp** configuration.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- Example: Preventing BGP Session Resets
- Junos OS Administration Library for Routing Devices

forwarding-db-size

IN THIS SECTION

- Syntax | 7455
- Hierarchy Level | 7455
- Description | 7455
- Options | 7455
- Required Privilege Level | 7455
- Release Information | 7455

Syntax

```
forwarding-db-size size;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-
provider]
```

Description

Configure the size of the forwarding database (FDB). When this setting is changed, the PIC reboots.

> **NOTE**: You need to enable the `forwarding-options sampling` statement for the FDB to be created.

Options

size—Size of the FDB, in megabytes (MB). The size of the FDB and the size of the policy database together must be smaller than the size of the object cache.

- **Range:** 0 through 12879 MB

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

policy-db-size | **12674**

wired-process-mem-size | **31133**

object-cache-size | 11765

forwarding (Diameter Network Element)

IN THIS SECTION

- Syntax | 7456
- Hierarchy Level | 7456
- Description | 7456
- Required Privilege Level | 7457
- Release Information | 7457

Syntax

```
forwarding {
    route dne-route-name {
        destination realm realm-name <host hostname>;
        function function-name <partition partition-name>;
        metric route-metric;
    }
}
```

Hierarchy Level

```
[edit diameter network-element element-name]
```

Description

Define the criteria that specify which destinations are reachable through the Diameter network element.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Network Elements

forwarding-instance identifier

IN THIS SECTION

- Syntax | 7457
- Hierarchy Level | 7458
- Description | 7458
- Options | 7458
- Required Privilege Level | 7458
- Release Information | 7458

Syntax

```
forwarding-instance identifier identifier;
```

Hierarchy Level

```
[edit routing-instances mac-vrf-routing-instance-name],
[edit tenants tenant-name routing-instances mac-vrf-routing-instance-name]
```

Description

QFX10000 line of switches, and ACX7100-32C, and ACX7100-48L devices—Configure the forwarding-instance parameters for a MAC-VRF routing instance.

If you configure the forwarding instance with an identifier, then the switch maps the specified MAC-VRF routing instance to the forwarding instance specified by the *identifier* variable.

If you do not configure an identifier, then the switch maps the specified MAC-VRF routing instance to the system default forwarding instance, `default-switch`.

Options

identifier *forwarding-instance-identifier*	Forwarding instance identifier. Specify the forwarding instance identifier as an integer value. QFX10000 line of switches: Identifier number range is 1 - 99. ACX7100-32C and ACX7100-48L routers: Identifier number range is 1 - 6.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

> *MAC-VRF Routing Instance Type Overview*
>
> *Overlapping VLAN Support Using Multiple Forwarding Instances or VLAN Normalization*

forwarding-mode (100-Gigabit Ethernet)

IN THIS SECTION

- Syntax | 7459
- Hierarchy Level | 7459
- Description | 7459
- Required Privilege Level | 7459
- Release Information | 7460

Syntax

```
forwarding-mode {
    (sa-multicast | ...the following vlan-steering statement...);
    vlan-steering {
        vlan-rule (high-low | odd-even);
    }
}
```

Hierarchy Level

```
[edit chassis fpc slot pic slot]
```

Description

Configure the interoperation mode for 100-Gigabit Ethernet PIC or the 100-Gigabit Ethernet MIC.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Statement introduced in Junos OS Release 12.1 for MX Series routers.

RELATED DOCUMENTATION

Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP

Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode

Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP

Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4

sa-multicast (100-Gigabit Ethernet) | **16264**

vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) | **30850**

vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) | **30855**

forwarding-mode (PTX Series Packet Transport Routers)

IN THIS SECTION

- Syntax | **7460**
- Hierarchy Level | **7461**
- Description | **7461**
- Required Privilege Level | **7461**
- Release Information | **7461**

Syntax

```
forwarding-mode {
    sa-multicast
}
```

Hierarchy Level

```
[edit chassis fpc slot pic slot port port-number]
```

Description

Configure interoperability between 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X48R4.

RELATED DOCUMENTATION

Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4

Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP

Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4

forwarding-mode (Security Content Security Policy)

IN THIS SECTION

- Syntax | 7462
- Hierarchy Level | 7462

- Description | 7462
- Options | 7462
- Required Privilege Level | 7463
- Release Information | 7463

Syntax

```
forwarding-mode {
    hold;
    inline-tap;
}
```

Hierarchy Level

```
[edit security utm default-configuration anti-virus]
```

Description

The default configuration for anti-virus is to use the continuous delivery function (CDF). It holds the last packet and sends out all other packets. It saves system memory and makes the packet transmission faster. This mode sends the last packet if the result is “permit” and sends RST message to both the client and the server to reset the connection if the result is “drop”. In CDF mode, you may save an incomplete infected file because it only holds the last packet and sends out others. This file could be executable and harmful, for example, an incomplete script file. CDF mode does not support Mail protocols. Change to `hold` mode to hold all the packets until you get the final result. Configure `inline-tap` mode to permit the traffic even if it is infected. This mode is off by default. You can set the `hold` and `inline-tap` mode separately or simultaneously. When you set both modes simultaneously, `inline-tap` over-rules the `hold` mode and permits the traffic.

To delete `hold` mode use `#delete security utm default-configuration anti-virus forwarding-mode hold`, and to delete `inline-tap` mode use `#delete security utm default-configuration anti-virus forwarding-mode inline-tap`.

Options

`hold` —Hold mode (hold file until analysis is complete, default is CDF mode).

`inline-tap` —Detect-only mode without blocking (default is off).

The statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2.

RELATED DOCUMENTATION

Unified Threat Management Overview

forwarding-options (Routing Instance)

IN THIS SECTION

- Syntax | 7463
- Hierarchy Level | 7464
- Description | 7464
- Options | 7464
- Required Privilege Level | 7466
- Release Information | 7466

Syntax

```
forwarding-options {
    access-security ...
    accounting name ...
```

```
    dhcp-relay ...
    enhanced-hash-key ...
    family family-name ...
    fast-reroute-priority (high | low | medium);
    helpers ...
    inband-telemetry
    ip-options-protocol-queue   protocol-name ...
    link-layer-broadcast-inet-check;
    load-balance ...
    load-balance-label-capability;
    multicast-replication ...
    next-hop-group group-name ...
    no-load-balance-label-capability;
    port-mirroring ...
    rpf-loose-mode-discard ...
    sampling ...
    storm-control enhanced;
    storm-control-profiles profile-name  ...
}
```

Hierarchy Level

```
[edit]
[edit routing-instances routing-instance-name]
```

Description

Configure traffic forwarding options. For more information about each option, click a linked statement in the Syntax section.

Options

access-security Configure IPv6 access security options.

accounting Specify the discard accounting instance name and options.

dhcp-relay Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch to enable the router (or switch) to function as a DHCP relay agent.

enhanced-hash-key	(MX Series routers with MPCs, T4000 routers with Type 5 FPCs, and EX9200 switches) Select data used in the hash key for enhanced IP forwarding engines.
family	Specify address family for filters.
fast-reroute-priority	Specify the fast reroute priority for a VPLS routing instance. You can configure `high`, `medium`, or `low` fast reroute priority to prioritize specific VPLS routing instances for faster convergence and traffic restoration.
helpers	Enable TFTP or DNS request packet forwarding, or configure the router, switch, or interface to act as a DHCP/BOOTP relay agent. Enable forwarding LAN broadcast traffic on custom UDP ports to particular servers as unicast traffic.
inband-telemetry	Enable or disable the Inband Network Telemetry feature at the data path level, without intervention from the control plane.
ip-options-protocol-queue	Configure logical queue-depth in the PFE for ip-options packets for a given protocol such as TCP, UDP, ICMP, and so on, except IGMP.
link-layer-broadcast-inet-check	Enable destination MAC and IP address check to prevent the router from forwarding IPv4 packets that have link layer destination address set to `broadcast` or `multicast`, unless directed to an IPv4 multicast address.
load-balance	Enable per-prefix or per-flow load balancing so that the router or switch elects a next hop independently of the route selected by other routers or switches.
load-balance-label-capability	Enable the router to push and pop the load balancing label and causes LDP and RSVP to advertise the entropy label TLV to neighboring routers.
multicast-replication	Configure the mode of multicast replication that helps to optimize multicast latency.
next-hop-group	Specify the next-hop address for sending copies of packets to an analyzer.
no-load-balance-label-capability	Disable advertisement of entropy label capability in LDP and RSVP.
port-mirroring	Specify the input, output, and traceoptions properties for sending copies of packets to an analyzer.
rpf-loose-mode-discard	Configure unicast reverse path forwarding (unicast RPF) loose mode with the ability to discard packets with the source address pointing to the discard next hop.
sampling	Configure traffic sampling.

storm-control-profiles Configure a storm control profile on a switch or router. Storm control is used to prevent network outages that are caused by broadcast traffic storms.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring Stateless IPv6 Router Advertisement Guard
- Configuring Stateful IPv6 Router Advertisement Guard
- Configuring Discard Accounting
- *Extended DHCP Relay Agent Overview*
- Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers
- Applying Forwarding Table Filters
- *Configuring VPLS Fast Reroute Priority*
- Configuring DNS and TFTP Packet Forwarding
- Configuring Port-based LAN Broadcast Packet Forwarding
- Configuring Per-Flow Load Balancing Based on Hash Values
- Configuring Per-Prefix Load Balancing
- Configuring the Entropy Label for LSPs
- Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers
- Understanding Unicast RPF (Routers)
- Configuring Traffic Sampling on MX, M and T Series Routers
- OBSOLETE: Understanding Storm Control for Managing Traffic Levels

forwarding-options (Bridge Domains or VLANs)

IN THIS SECTION

- Syntax | 7467
- Chassis: EX4600 and QFX Series | 7468
- Chassis: EX4600 and QFX Series | 7469
- Chassis: EX4600 and QFX Series | 7469
- Hierarchy Level | 7469
- Description | 7469
- Options | 7471
- Required Privilege Level | 7472
- Release Information | 7472

Syntax

```
forwarding-options {
    dhcp-security {
        arp-inspection;
        group group-name {
            interface interface-name {
                static-ip ip-address {
                    mac mac-address;
                }
            }
            overrides {
                no-option82;
                (trusted | untrusted);
            }
        }
        ip-source-guard;
        no-dhcp-snooping;
        option-82 {
            circuit-id {
                prefix {
                    host-name;
```

```
                    logical-system-name;
                    routing-instance-name;
                }
                use-interface-description (device | logical);
                use-vlan-id;
            }
            remote-id {
                host-name hostname;
                use-interface-description (device | logical);
                use-string string;
            }
            vendor-id {
                use-string string;
            }
        }
    }
    storm-control {
        enhancedenhanced;
    }
    filter (VLANs) {
        input filter-name;
        output filter-name;
    }
    sflow {
        egress-multicast enable;
        egress-multicast max-replication-rate rate;
    }
    flood {
        input filter-name;
    }
}
```

Chassis: EX4600 and QFX Series

```
forwarding options profile-name {
    num-65-127-prefix number;
}
```

Chassis: EX4600 and QFX Series

```
forwarding-options lpm-profile {
    prefix-65-127-disable;
    unicast-in-lpm;
}
```

Chassis: EX4600 and QFX Series

```
forwarding-options custom-profile {
    l2-entries | l3-entries | lpm-entries {
        num-banks number;
    }
}
```

Hierarchy Level

```
[edit],
[edit bridge-domains bridge-domain-name],
[edit vlans vlan-name]
```

```
[edit chassis (QFX Series)]
```

Description

Configure a unified forwarding table profile to allocate the amount of memory available for the following:

- MAC addresses.
- Layer 3 host entries.
- Longest prefix match table entries.

This feature enables you to select a profile that optimizes the amount of memory available for various types of forwarding-table entries based on the needs of your network. For example, for a switch that handles a great deal of Layer 2 traffic, such as a virtualized network with many servers and virtualized

machines, you would choose the `l2-profile-one`, which allocates the highest amount of memory to MAC addresses.

You configure the memory allocation for LPM table entries differently, depending on whether you using Junos OS Release 13.2X51-D10 or Junos OS Release 13.2X51-D15 and later. For more information about configuring memory allocation for LPM table entries, see Configuring the Unified Forwarding Table on Switches.

The `num-65-127-prefix` *`number`* statement is not supported on the `custom-profile` and the `lpm-profile`. The `prefix-65-127-disable` and `unicast-in-lpm` statements are supported only on the `lpm-profile`.

When you commit a configuration with a forwarding table profile change, in most cases the Packet Forwarding Engine restarts automatically to apply the new parameters, which brings the data interfaces down and then up again.

However, starting with Junos OS Releases 14.1X53-D40, 15.1R5, and 16.1R3, for a Virtual Chassis or Virtual Chassis Fabric (VCF) comprised of EX4600 or QFX5100 switches, the Packet Forwarding Engine in member switches does not automatically restart upon configuring and committing a unified forwarding table profile change. This behavior avoids having Virtual Chassis or VCF instability and a prolonged convergence period if a profile change is propagated to member switches and multiple Packet Forwarding Engines all restart at the same time. In this environment, instead of automatically restarting when you initially commit a profile configuration change, the message `Reboot required for configuration to take effect` is displayed at the primary switch CLI prompt, notifying you that the profile change does not take effect until the next time you restart the Virtual Chassis or VCF. The profile configuration change is propagated to member switches that support this feature, and a reminder that a reboot is required to apply this pending configuration change appears in the system log of the primary switch and applicable member switches. You then enable the profile change subsequently during a planned downtime period using the `request system reboot` command, which quickly establishes a stable Virtual Chassis or VCF with the new configuration.

> **NOTE**: You should plan to make unified forwarding table profile changes only when you are ready to perform a Virtual Chassis or VCF system reboot *immediately* after committing the configuration update. Otherwise, in the intervening period between committing the configuration change and rebooting the Virtual Chassis or VCF, the system can become inconsistent if a member experiences a problem and restarts. In that case, the new configuration takes effect on the member that was restarted, while the change is not yet activated on all the other members.

The remaining statements are explained separately. See CLI Explorer.

Options

profile-name name of the profile to use for memory allocation in the unified forwarding table. Table 45 on page 7471 lists the profiles you can choose that have set values and the associated values for each type of entry.

On QFX5200 Series switches only, you can also select `custom-profile`. This profile enables you to allocate from one to four banks of shared hash memory to a specific type of forwarding-table entry. Each shared hash memory bank can store a maximum of the equivalent of 32,000 IPv4 unicast addresses.

Table 45: Unified Forwarding Table Profiles

Profile Name	MAC Table	Host Table (unicast and multicast addresses)					
	MAC Addresses	IPv4 unicast	IPv6 unicast	IPv4 (*, G)	IPv4 (S, G)	IPv6 (*, G)	IPv6 (S, G)
`l2-profile-one`	288K	16K	8K	8K	8K	4K	4K
`l2-profile-two`	224K	80K	40K	40K	40K	20K	20K
`l2-profile-three` (default)	160K	144K	72K	72K	72K	36K	36K
`l3-profile`	96K	208K	104K	104K	104K	52K	52K
`lpm-profile`*	32K	16K	8K	8K	8K	4K	4K

* This profile supports only IPv4 in Junos OS Release13.2X51-D10. Starting in Junos OS Relase13.2X51-D15, the `lpm-profile` supports IPv4 and IPv6 entries.

> **NOTE**: If the host stores the maximum number of entries for any given type, the entire table is full and is unable to accommodate *any* entries of any other type. For information about valid combinations of table entries see Understanding the Unified Forwarding Table.

l2-entries | l3-entries | lpm-entries

(`custom-profile only`) Select a type of forwarding-table entry—Layer 2, Layer 3, or LPM—to allocate a specific number of shared memory banks. You configure the amount of memory to allocate for each type of entry separately.

num-banks *number*

(`custom-profile only`) Specify the number of shared memory banks to allocate for a specific type of forwarding-table entry. Each shared memory bank stores the equivalent of 32,000 IPv4 unicast addresses.

- **Range:** 0 through 4.

> **NOTE**: There are four shared memory banks, which can be allocated flexibly among the three types of forwarding-table entries. To allocate no shared memory for a particular entry type, specify the number `0`. When you commit the configuration, the system issues a commit check to ensure that you have not configured more than four memory banks. You do not have to configure all four shared memory banks. By default, each entry type is allocated the equivalent of 32,000 IPv4 unicast addresses in shared memory.

egress-multicast enable

Enable egress sFlow sampling of known multicast traffic.

egress-multicast max-replication-rate *rate*

Specify the maximum sFlow sample replication rate per FPC for known multicast traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Hierarchy level [edit vlans *vlan-name*] introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Hierarchy level [edit bridge-domains *bridge-domain-name*] introduced in Junos OS Release 14.1 for MX Series routers.

custom-profile option introduced in Junos OS Release 15.1x53-D30 for QFX5200 Series switches only.

RELATED DOCUMENTATION

- Understanding the Unified Forwarding Table
- Example: Configuring a Unified Forwarding Table Custom Profile
- *Configuring Traffic Forwarding and Monitoring*

forwarding-options (Security)

IN THIS SECTION

- Syntax | 7473
- Hierarchy Level | 7475
- Description | 7475
- Options | 7475
- Required Privilege Level | 7477
- Release Information | 7477

Syntax

```
forwarding-options {
    allow-dataplane-sleep;
    family {
        inet6 {
            mode (drop | flow-based | packet-based);
        }
```

```
        iso {
            mode packet-based;
        }
        mpls {
            mode (flow-based | packet-based);
        }
    }
    mirror-filter filter-name;
    mode {
        tap {
            inspect-pass-through-tunnel {
                gre;
                ipip;
            }
            interface [ interface ... ];
            }
        }

    receive-side-scaling {
        nic-rss {
            mode disable;
        }
        software-rss {
            io-thread-number io-thread-number;
            mode(automatic | disable | enable);
        }
    }
    resource-manager {
        cpu {
            re re;
        }
        enhanced-logging;
        }
    secure-wire name;
    security-service {
        fail-open;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Determine how the `inet6`, `iso`, and `mpls` protocol families manage security forwarding options.

NOTE:

- Packet-based processing is not supported on the following SRX Series Firewalls: SRX5400, SRX5600, and SRX5800.
- On SRX Series Firewalls, the default mode for processing traffic is flow mode. You can configure SRX Series Firewalls to operate in packet mode to process MPLS packets.

  To configure the packet mode on SRX Series Firewall, use the following command:

  ```
  user@host# set security forwarding-options family mpls mode packet-based
  ```

  Selective stateless packet-based services allows you to configure the device to provide only packet-based processing for selected traffic based on input filter terms.

Options

allow-dataplane-sleep Enable sleep on dataplane.

family Specify the protocol family to be used for packet forwarding.

- `inet6`—Specify forwarding options for IPv6 traffic.
  - `drop`—Drop IPv6 packets. This is the default setting.
  - `flow-based`—Perform flow-based packet forwarding.
  - `packet-based`—Perform simple packet forwarding.
- `iso`—Specify forwarding options for IS-IS traffic.
  - `packet-based`—Perform simple packet forwarding.

- `mpls`—Specify forwarding options for MPLS traffic.
  - `flow-based`—Perform flow-based packet forwarding.
  - `packet-based`—Perform simple packet forwarding.

mirror-filter Specify a mirror filter for filtering X2 packets to be mirrored and sent to a packet analyzer.

mode Specify TAP or Sniffer mode.

- `tap`—Specify TAP mode.
  - `inspect-pass-through-tunnel`—Specify TAP mode to inspect pass through IP-IP or GRE tunnel.
  - `interface`—Specify TAP mode interface name. You can configure up to eight TAP interfaces.

receive-side-scaling Receive side scaling (RSS) enables the efficient distribution of network receive processing across multiple CPUs in multiprocessor systems.

- `nic-rss`—The NIC distributes packets by applying a filter to each packet that assigns it to one of a small number of logical flows. Packets for each flow are steered to a separate receive queue, which in turn can be processed by separate CPUs. On reception, a NIC can send different packets to different queues to distribute processing among CPUs. NIC must have same number of receiving (RX) and transmitting (TX) queues as number of vSRX data plane CPU to support multi core vSRX flavors.
  - `mode`—Specify the mode of NIC RSS.
    - `disable`—Disable the NIC RSS mode.
- `software-rss`—Software RSS (SWRSS) removes this limitation of NIC HW queues to run vSRX multi-core flavors by implementing software-based packet spraying across various FLT thread. Software RSS offloads the handling of individual flows to one of the multiple kernel, so the flow thread that takes the packets from the NIC can process more packets. SWRSS implements software model of packet distribution across FLTs after obtaining the packets from NIC receiving queues. By removing NIC HW queue limitation, SWRSS helps to scale vCPUs by supporting various vSRX instance types.

- `io-thread-number`—Specify the software RSS IO thread number.
  - **Range:** 1 through 8
- `mode`—Specify the mode of Software RSS.
  - `automatic`—Auto select the Software RSS mode. This is the default mode.
  - `disable`—Disable the Software RSS mode.
  - `enable`—Enable the Software RSS mode.

resource-manager Display forward option status, the CPU, and memory allocated for the advance services to verify the vCPU allocation between routing engine and flow RT threads.

- `cpu`—Specify CPU resources.
  - **Default:** 1
  - **Range:** 1 through 3
- `enhanced-logging`—Specify a dedicated CPU resource for on-box logging.

secure-wire Specify a name for the secure wire interface mapping.

security-service Security service actions when memory resource is in shortage. The system resource management guarantees the resources are used according to priorities.

- `fail-open`—Ignores Layer 7 services with resource requirements, creates a flow session without Layer 7 services, and forward the packet out.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

`secure-wire` option introduced in Junos OS Release 19.3R1.

`resource-manager` option introduced in Junos OS Release 19.4R1 for vSRX Virtual Firewall.

`mode` option introduced in Junos OS Release 20.1R1.

`enhanced-logging` option introduced in Junos OS Release 23.1R1 to assign a dedicated CPU resource for on-box logging.

RELATED DOCUMENTATION

- MPLS Overview
- *Understanding Packet-Based Processing*
- *Understanding Traffic Processing on Security Devices*
- show security flow status | **24348**

forwarding-packages

IN THIS SECTION

- Syntax | **7478**
- Hierarchy Level | **7479**
- Description | **7479**
- Required Privilege Level | **7479**
- Release Information | **7479**

Syntax

```
forwarding-packages {
     mobility{
        user-plane;
        sgw;
    }
}
```

Hierarchy Level

```
[edit chassis fpc fpc-slot pfe pfe-id]
```

Description

Configure the Packet Forwarding Engine so that it can be used to anchor mobile sessions. If this configuration is changed, then the FPC reboots.

The `forwarding-packages` statement can be configured at the Packet Forwarding Engine level. Therefore, you can configure a subset of Packet Forwarding Engines in an FPC to be mobile anchors.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced for Junos Multi-Acess User Plane in Junos OS Release 19.4R1.

forwarding-policy

IN THIS SECTION

- Syntax | 7480
- Hierarchy Level | 7480
- Description | 7480
- Required Privilege Level | 7480
- Release Information | 7481

Syntax

```
forwarding-policy {
    next-hop-map map-name {
        forwarding-class class-name {
            discard;
            lsp-next-hop [ lsp-regular-expression ];
            next-hop [ next-hop-name ];
            non-lsp-next-hop;
        }
        forwarding-class-default {
            discard;
            lsp-next-hop [ lsp-regular-expression ];
            next-hop [next-hop-name];
            non-lsp-next-hop;
        }
    }
    class class-name {
        classification-override {
            forwarding-class class-name;
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Define CoS-based forwarding policy options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for QFX10000 Series switches in Junos OS Release 17.1R1 to support CoS-based forwarding (CBF). `[set class-of-service forwarding-policy class]` is not supported on QFX10000 Series switches.

RELATED DOCUMENTATION

Configuring CoS-Based Forwarding

forwarding-process

IN THIS SECTION

- Syntax | 7481
- Hierarchy Level | 7482
- Description | 7482
- Options | 7483
- Required Privilege Level | 7483
- Release Information | 7483

Syntax

```
forwarding-process {
                        application-services (Security Forwarding Process) {
        enable-gtpu-distribution;
        inline-fpga-crypto (disabled | enabled);
        maximize-alg-sessions;
                                maximize-idp-sessions {
                                    inline-tap;
                                       weight (Security)  {
                firewall;
```

```
                idp {
                }
            }
        }
        maximize-persistent-nat-capacity;
        packet-ordering-mode (hardware | software);
    }
    enhanced-services-mode;
}
```

Hierarchy Level

```
[edit security]
```

Description

You can configure SRX5400, SRX5600, and SRX5800 devices to switch from an integrated firewall mode to maximize Intrusion Detection and Prevention (IDP) mode to run IDP processing in tap mode and increase the capacity of processing with the `maximize-idp-sessions` option. Inline tap mode can only be configured if the forwarding process mode is set to `maximize-idp-sessions`, which ensures stability and resiliency for firewall services. You also do not need a separate tap or span port to use inline tap mode. When you maximize IDP, you are decoupling IDP processes from firewall processes, allowing the device to support the same number of firewall and IDP sessions, also run the IDP processing in tap mode.

You can configure maximum Application Layer Gateway (ALG) sessions by using the `maximize-alg-sessions` option. By default, the session capacity number for Real-Time Streaming Protocol (RTSP), FTP, and Trivial File Transfer Protocol (TFTP) ALG sessions is 10,000 per flow Services Processing Unit (SPU). You must reboot the device (and its peer in chassis cluster mode) for the configuration to take effect. The `maximize-alg-sessions` option now enables you to increase defaults as follows:

- RTSP, FTP, and TFTP ALG session capacity: 25,000 per flow SPU
- TCP proxy connection capacity: 40,000 per flow SPU

> **NOTE**: Flow session capacity is reduced to half per flow SPU; therefore the aforementioned capacity numbers will not change on central point flow.

Enable GPRS tunneling protocol, user plane(GTP-U) session distribution to distribute GTP-U traffic handled by a Gateway GPRS Support Node (GGSN) and a Serving GPRS Support Node (SGSN) pair on all Services Processing Units (SPUs). You can configure tunnel-base distribution to distribute GTP-U traffic

to multiple SPUs by the `enable-gtpu-distribution` option on SRX5400, SRX5600, and SRX5800 devices , which helps to resolve the GTP-U fat session issue. Also, `enable-gtpu-distribution` command is must for enabling stateful GTP-U inspection.

Options

enhanced-services-mode Enable enhanced application services mode.

When Enhanced Service Mode is enabled, resources such as MBUF, JBUF, SERVICE-MEM, TCP-PROXY TCB, SZ-INFO and user heap are increased in size and Layer 4 session numbers are reduced by half.

The remaining statements are explained separately. See the CLI Explorer.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6. This statement is supported.

Option `enhanced-services-mode` introduced in Junos OS Release 20.3R1 on vSRX Virtual Firewall 3.0.

RELATED DOCUMENTATION

application-services (Security Forwarding Process) | **2375**

Understanding Traffic Processing on Security Devices

forwarding-rib

IN THIS SECTION

- Syntax | **7484**
- Hierarchy Level | **7484**

- Description | 7484
- Options | 7484
- Required Privilege Level | 7484
- Release Information | 7485

Syntax

```
forwarding-rib name {
    inet-import [ inet-import ... ];
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options dynamic-tunnels],
[edit logical-systems name routing-options dynamic-tunnels],
[edit logical-systems name tenants name routing-instances name routing-options dynamic-tunnels],
[edit routing-instances name routing-options dynamic-tunnels],
[edit routing-options dynamic-tunnels],
[edit tenants name routing-instances name routing-options dynamic-tunnels]
```

Description

Configure policy control for forwarding routing table next hops for MPLS-over-UDP dynamic tunnels. With this configuration, you can resolve the dynamic tunnel destination routes over select prefixes.

Options

name Name of the routing table.

inet-import Name of the import policy for IPv4 dynamic-tunnels.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels

forwarding-state-bit (Long-Lived Graceful Restart for BGP Restarter)

IN THIS SECTION

- Syntax | 7485
- Hierarchy Level | 7485
- Description | 7486
- Options | 7487
- Required Privilege Level | 7487
- Release Information | 7487

Syntax

```
forwarding-state-bit (from-fib | set);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family address-family subsequent-address-family graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name family address-family subsequent-address-family graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family address-family subsequent-address-family graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
```

```
family address-family subsequent-address-family graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family address-family subsequent-address-family graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family address-family subsequent-address-family graceful-
restart],
[edit routing-instances routing-instance-name protocols bgp family address-family subsequent-
address-family graceful-restart],
[edit routing-instances routing-instance-name protocols bgp group group-name family address-
family subsequent-address-family graceful-restart],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family address-family subsequent-address-family graceful-restart],
[edit protocols bgp family address-family subsequent-address-family graceful-restart],
[edit protocols bgp group group-name family address-family subsequent-address-family graceful-
restart],
[edit protocols bgp group group-name neighbor address family address-family subsequent-address-
family graceful-restart]
```

NOTE: Each routing table is identified by the protocol family or address family indicator (AFI) and a subsequent address family identifier (SAFI).

Configuring LLGR does not require that BGP graceful restart also be configured. The long-lived-graceful-restart section is visible only for families l2vpn, inet labeled-unicast, inet flow and route-target. It is prohibited for inet-mvpn, inet6-mvpn and inet-mdt. It is hidden for other families.

Description

Configure the forwarding-state bit flag negotiation for BGP for all address families. In addition to the global setting for the Forwarding State bit, the Forwarding State bit behavior can be specified for individual families. Changing the forwarding-state-bit setting has no effect on any existing sessions. Per-family BGP configuration options are added to control the Forwarding State bit in graceful restart and long-lived graceful restart capability advertisements. They can be specified for the default logical system or for a specific logical system, and for the primary routing instance or a specific routing instance. The `per-family forwarding-state-bit` attribute overrides the default rules or the global configuration for setting the Forwarding State bit.

The setting of the F bit (and the "Forwarding State" bit of the accompanying GR capability) depends in part on deployment considerations. The F bit can be interpreted to indicate the helper router needs to flush associated routes (if the bit is left clear). An important scenario in which LLGR is used is for routes that are more similar to configuration than to traditional routing (hop-by-hop forwarding instead of tunnel-based routing). For such routes, it might be useful to always set the F bit, regardless of other

considerations. Similarly, for control-plane-only entities such as dedicated route reflectors, that do not participate in the forwarding plane, it is preferred that the F bit be always set. Overall, the guideline to be adopted is that if loss of state on the restarting router can reasonably be expected to cause a forwarding loop or null route, the F bit must be set judiciously, depending on whether state has been retained. You can determine whether the F bit needs to be set or not, based on your deployment needs and configured settings. It might be necessary to advertise stale routes to a CE in some VPN deployments, even if the CE does not support this specification. In such a scenario, the network operator configuring their PE to advertise such routes must notify the operator of the CE receiving the routes, and the CE must be configured to depreference the routes. Typically, BGP implementations perform this behavior by matching on the LLGR_STALE community, and setting the LOCAL_PREF for matching routes to zero.

Options

`set` Causes the value to be set according to the state of the associated FIB. Changing the per-family forwarding-state-bit setting has no effect on any existing sessions.

`from-fib` Forces the Forwarding State bit to be set to 1.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

forwarding-state-bit (Per Family for BGP Graceful Restart)

IN THIS SECTION

- Syntax | 7488
- Hierarchy Level | 7488
- Description | 7489
- Options | 7489
- Required Privilege Level | 7489
- Release Information | 7490

Syntax

```
forwarding-state-bit (as-rr-client | from-fib);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
graceful-restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp graceful-
restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name graceful-restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name neighbor address graceful-restart],
[edit routing-instances routing-instance-name protocols bgp graceful-restart],
[edit routing-instances routing-instance-name protocols bgp group group-name graceful-restart],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
graceful-restart],
[edit protocols bgp graceful-restart],
[edit protocols bgp group group-name graceful-restart],
[edit protocols bgp group group-name neighbor address graceful-restart]
```

Description

Configure the forwarding-state bit flag negotiation for BGP for individual address families. In addition to the global setting for the Forwarding State bit, the Forwarding State bit behavior can be specified for individual families. Changing the forwarding-state-bit setting has no effect on any existing sessions. Per-family BGP configuration options are added to control the Forwarding State bit in graceful restart and long-lived graceful restart capability advertisements. They can be specified for the default logical system or for a specific logical system, and for the primary routing instance or a specific routing instance. The `per-family forwarding-state-bit` attribute overrides the default rules or the global configuration for setting the Forwarding State bit.

The setting of the F bit (and the "Forwarding State" bit of the accompanying GR capability) depends in part on deployment considerations. The F bit can be interpreted to indicate the helper router needs to flush associated routes (if the bit is left clear). An important scenario in which LLGR is used is for routes that are more similar to configuration than to traditional routing (hop-by-hop forwarding instead of tunnel-based routing). For such routes, it might be useful to always set the F bit, regardless of other considerations. Similarly, for control-plane-only entities such as dedicated route reflectors, that do not participate in the forwarding plane, it is preferred that the F bit be always set. Overall, the guideline to be adopted is that if loss of state on the restarting router can reasonably be expected to cause a forwarding loop or null route, the F bit must be set judiciously, depending on whether state has been retained. You can determine whether the F bit needs to be set or not, based on your deployment needs and configured settings. It might be necessary to advertise stale routes to a CE in some VPN deployments, even if the CE does not support this specification. In such a scenario, the network operator configuring their PE to advertise such routes must notify the operator of the CE receiving the routes, and the CE must be configured to depreference the routes. Typically, BGP implementations perform this behavior by matching on the LLGR_STALE community, and setting the LOCAL_PREF for matching routes to zero.

Options

`set` Causes the value to be set according to the state of the associated FIB. Changing the per-family forwarding-state-bit setting has no effect on any existing sessions.

`from-fib` Forces the Forwarding State bit to be set to 1.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

forwarding-table

IN THIS SECTION

- Syntax | 7490
- Hierarchy Level | 7491
- Description | 7491
- Options | 7491
- Required Privilege Level | 7492
- Release Information | 7492

Syntax

```
forwarding-table {
    chained-composite-next-hop;
    ecmp-fast-reroute;
    evpn-egress-link-protection;
    export [ policy-name ];
    (indirect-next-hop | no-indirect-next-hop);
    (indirect-next-hop-change-acknowledgements | no-indirect-next-hop-change-acknowledgements;)
    krt-nexthop-ack-timeout interval
    multicast-composite-delete-defer-timeout;
    pfe-install-error;
    remnant-holdtime;
    route-origin-proto-tlv;
```

```
    srv6-chain-merge;
    unicast-reverse-path (active-paths | feasible-paths);
    vpls-hotstandby-convergence;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit routing-options]
```

Description

Configure information about the routing device's forwarding table.

Some of the remaining statements are explained separately. See CLI Explorer.

Options

evpn-egress-link-protection

Enables multihomed PE devices to support fast reroutes (FFRs).

multicast-composite-delete-defer-timeout

Use this command to delay the deletion of a composite next-hop from the FIB (in the event of dropped multicast packets following a downstream interface change). You can view the set time by running the **show krt state** command.

Range: 1 to 3600 seconds

pfe-install-error

(Junos OS Evolved only) Enable error checking for Packet Forwarding Engine route install errors. When configured, the forwarding plane can inform the control plane of any errors in next hops. If a route error is received from the Packet Forwarding Engine, the route error is directly added to the route-error list. For next-hop errors, the error is propagated to the parent route and the parent route is added to the route-error list. Two commands are available to check on these route errors after you have configured this feature:

- If the forwarding plane is not forwarding traffic, use the `show route pfe-install-error` command to check if the Packet Forwarding Engine has sent an error for the route that is dropping traffic.

- To check for next-hop errors, use the `show routing nhdb pfe-install-error` command. We support errors for routes pointing to a limited set of next hops: RT_NH_INDIRECT, RT_NH_ROUTER, RT_NH_CHAIN, RT_NH_INDXD, RT_NH_LIST, RT_NH_TUNNEL_COMP, and RT_NH_FRR_INDIRECT.

remnant-holdtime Sets the remnant hold time, which is required for the MXVC-ISSU, where the recommended value is 900.

- **Default:** 180 (For all Junos OS releases and all Junos OS Evolved releases earlier than 20.4R2)
- **Default:** 300 (For all Junos OS Evolved Release 20.4R2 and later)

route-origin-proto-tlv Configure this option to include the origin protocol information in route messages. The origin protocol is the protocol from which the route entry is learned.

srv6-chain-merge Enable merging of SRv6 chain nexthops.

vpls-hotstandby-convergence Enabled on spoke devices to support sending different advertising labels on the active and standby pseudowires.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`srv6-chain-merge` option introduced in Junos OS Release 21.3R1.

`multicast-composite-delete-defer-timeout` option introduced in Junos OS Release 22.1R2 and Junos OS Evolved Release 22.1R2.

`route-origin-proto-tlv` option introduced in Junos OS Release 22.4R1 and Junos OS Evolved Release 22.4R1.

`evpn-egress-link-protection`, `remnant-holdtime`, and `vpls-hotstandby-convergence` options introduced in Junos OS Release 24.2R1.

`pfe-install-error` option introduced in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

- Example: Load Balancing BGP Traffic
- Configuring Per-Packet Load Balancing

forwarding-options (Unified Forwarding Table Profile)

IN THIS SECTION

- Syntax | 7493
- Hierarchy Level | 7494
- Description | 7494
- Options | 7496
- Required Privilege Level | 7498
- Release Information | 7498

Syntax

```
forwarding-options {
    custom-profile {
        exact-match-entries {
            num-banks num-banks;
        }
        l2-entries {
            num-banks num-banks;
        }
        l3-entries {
            num-banks num-banks;
        }
        lpm-entries {
            num-banks(0 | 2 | 4);
            prefix-65-127-disable;
        }
        num-65-127-prefix num-65-127-prefix;
```

```
    }
    l2-profile-one {
        num-65-127-prefix num-65-127-prefix;
    }
    l2-profile-three {
        num-65-127-prefix num-65-127-prefix;
    }
    l2-profile-two {
        num-65-127-prefix num-65-127-prefix;
    }
    l3-profile {
        num-65-127-prefix num-65-127-prefix;
    }
    lpm-profile {
        prefix-65-127-disable;
        unicast-in-lpm;
    }
    vxlan-gbp-l2-profile;
    vxlan-gbp-l3-profile;
    vxlan-gbp-profile;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure a unified forwarding table profile to allocate the amount of memory available for the following:

- MAC addresses.
- Layer 3 host entries.
- Longest prefix match table entries.

This feature enables you to select a profile that optimizes the amount of memory available for various types of forwarding-table entries based on the needs of your network. For example, for a switch that handles a great deal of Layer 2 traffic, such as a virtualized network with many servers and virtualized

machines, you would choose the `l2-profile-one`, which allocates the highest amount of memory to MAC addresses.

You configure the memory allocation for LPM table entries differently, depending on whether you using Junos OS Release 13.2X51-D10 or Junos OS Release 13.2X51-D15 and later. For more information about configuring memory allocation for LPM table entries, see Configuring the Unified Forwarding Table on Switches.

The `num-65-127-prefix` *`number`* statement is not supported on the `custom-profile` and the `lpm-profile`. The `prefix-65-127-disable` and `unicast-in-lpm` statements are supported only on the `lpm-profile`.

When you commit a configuration with a forwarding table profile change, in most cases the Packet Forwarding Engine restarts automatically to apply the new parameters, which brings the data interfaces down and then up again.

However, starting with Junos OS Releases 14.1X53-D40, 15.1R5, and 16.1R3, for a Virtual Chassis or Virtual Chassis Fabric (VCF) comprised of EX4600 or QFX5100 switches, the Packet Forwarding Engine in member switches does not automatically restart upon configuring and committing a unified forwarding table profile change. This behavior avoids having Virtual Chassis or VCF instability and a prolonged convergence period if a profile change is propagated to member switches and multiple Packet Forwarding Engines all restart at the same time. In this environment, instead of automatically restarting when you initially commit a profile configuration change, the message `Reboot required for configuration to take effect` is displayed at the primary switch CLI prompt, notifying you that the profile change does not take effect until the next time you restart the Virtual Chassis or VCF. The profile configuration change is propagated to member switches that support this feature, and a reminder that a reboot is required to apply this pending configuration change appears in the system log of the primary switch and applicable member switches. You then enable the profile change subsequently during a planned downtime period using the `request system reboot` command, which quickly establishes a stable Virtual Chassis or VCF with the new configuration.

> **NOTE**: You should plan to make unified forwarding table profile changes only when you are ready to perform a Virtual Chassis or VCF system reboot *immediately* after committing the configuration update. Otherwise, in the intervening period between committing the configuration change and rebooting the Virtual Chassis or VCF, the system can become inconsistent if a member experiences a problem and restarts. In that case, the new configuration takes effect on the member that was restarted, while the change is not yet activated on all the other members.

The remaining statements are explained separately. See CLI Explorer.

Options

custom-profile

On QFX5200 Series switches only, you can select `custom-profile`. This profile enables you to allocate from one to four banks of shared hash memory to a specific type of forwarding-table entry. Each shared hash memory bank can store a maximum of the equivalent of 32,000 IPv4 unicast addresses.

The the `lpm-profile` supports only IPv4 in Junos OS Release13.2X51-D10. Starting in Junos OS Relase13.2X51-D15, the `lpm-profile` supports IPv4 and IPv6 entries.

> **NOTE**: If the host stores the maximum number of entries for any given type, the entire table is full and is unable to accommodate *any* entries of any other type. For information about valid combinations of table entries see Understanding the Unified Forwarding Table.

l2-entries | l3-entries | lpm-entries

(`custom-profile only`) Select a type of forwarding-table entry—Layer 2, Layer 3, or LPM—to allocate a specific number of shared memory banks. You configure the amount of memory to allocate for each type of entry separately.

num-banks *number*

(`custom-profile only`) Specify the number of shared memory banks to allocate for a specific type of forwarding-table entry. Each shared memory bank stores the equivalent of 32,000 IPv4 unicast addresses.

- **Range:** 0 through 4.

> **NOTE**: There are four shared memory banks, which can be allocated flexibly among the three types of forwarding-table entries. To allocate no shared memory for a particular entry type, specify the number `0`. When you commit the configuration, the system issues a commit check to ensure that you have not configured more than four memory banks. You do not have to configure all four shared memory banks. By default, each entry type is allocated the equivalent of 32,000 IPv4 unicast addresses in shared memory.

profile-name

Configure one of the following profiles to allocate memory:

- l2-profile-one
- l2-profile-three

- l2-profile-two
- l3-profile
- lpm-profile

Review Understanding the Unified Forwarding Table to understand the memory allocation for your platform. You can also use the `set chassis forwarding-options ?` command to view the allocations for your platform. For example:

```
user@host# set chassis forwarding-options ?
Possible completions:
+ apply-groups         Groups from which to inherit configuration data
+ apply-groups-except  Don't inherit configuration data from these groups
> l2-profile-one       MAC: 288K L3-host: 16K LPM: 16K EM:0, needs reboot for VC;
restarts PFE if standalone
> l2-profile-three     (default) MAC: 160K L3-host: 144K LPM: 16K, needs reboot for
VC;restarts PFE if standalone
> l2-profile-two       MAC: 224K L3-host: 80K LPM: 16K EM:0, needs reboot for VC;
restarts PFE if standalone
> l3-profile           MAC: 96K L3-host: 208K LPM: 16K EM:0, needs reboot for VC;
restarts PFE if standalone
> lpm-profile          MAC: 32K L3-host: 16K LPM: 128K EM:0, needs reboot for VC;
restarts PFE if standalone
```

Requires reboot of the Virtual Chassis for profile change to take effect. For standalone deployments, committing the profile configuration restarts the PFE. You can use the `show chassis forwarding-options` command to view the UFT configuration. For example:

```
user@host> show chassis forwarding-options
localre:
--------------------------------------------------------------------------
UFT Configuration:
lpm-profile
prefix-65-127 = enable
--------------Bank details for various types of entries------------------
Entry type Dedicated Bank Size(K) Shared Bank Size(K)
L2(mac) 32 32 * num shared banks
L3(unicast & multicast) 16 32 * num shared banks
Exact Match 0 16 * num shared banks
```

Configure VXLAN GBP profiles and GBP policy filters with L4 matches on EX4400, EX4100, EX4650, and QFX5120 switches:

- vxlan-gbp-l2-profile
- vxlan-gbp-l3-profile
- vxlan-gbp-profile

Review Micro and Macro Segmentation using Group Based Policy in a VXLAN for more details.

You can also use the `set chassis forwarding-options ?` command to view the allocations for your platform. For example:

```
user@host# set chassis forwarding-options ?
Possible completions:
+ apply-groups Groups from which to inherit configuration data
+ apply-groups-except Don't inherit configuration data from these groups
> vxlan-gbp-l2-profile MAC: 160K L3-host: 48K LPM(IPv4): 32K EM: 48K.(restarts PFE
if standalone system and requires reboot for VC systems)
> vxlan-gbp-l3-profile MAC: 96K L3-host: 144K LPM(IPv4): 32K EM: 32K.(restarts PFE
if standalone system and requires reboot for VC systems)
> vxlan-gbp-profile MAC: 160K L3-host: 80K LPM(IPv4): 32K EM: 32K.(restarts PFE if
standalone system and requires reboot for VC systems)
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`custom-profile` option introduced in Junos OS Release 15.1x53-D30 for QFX5200 Series switches only.

VXLAN GBP profile options introduced in Junos OS Release 23.2R1 for EX4100, EX4400, EX4600, QFX5120-32C and QFX5120-48Y switches only.

RELATED DOCUMENTATION

- Understanding the Unified Forwarding Table
- Example: Configuring a Unified Forwarding Table Custom Profile
- *Configuring Traffic Forwarding and Monitoring*

fpc (QFX Series)

IN THIS SECTION

- Syntax | 7499
- Hierarchy Level | 7500
- Description | 7500
- Options | 7500
- Required Privilege Level | 7500
- Release Information | 7500

Syntax

```
fpc slot {
    auto-speed-detection disable;
    pic pic-number{
        tunnel-port port-number tunnel-services;
        port port-number{
            channel-speed (speed|disable-auto-speed-detection) ;
        }
        port-range port-range-low port-range-high {
            channel-speed (speed|disable-auto-speed-detection);
        }
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the FPC slot number. For QFX3500 switches, the slot is a line card slot.

For generic routing encapsulation (GRE) tunneling, use the `tunnel-port` statement to specify the port that you want to convert to a GRE tunnel port.

Options

slot—Number of the FPC slot. For QFX3500, QFX3600, QFX5200, and OCX Series devices, the slot number is always `0`.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show chassis fpc

fpc (Chassis)

IN THIS SECTION

- Syntax | 7501
- Syntax (ACX Series Devices) | 7501
- Syntax (PTX10003 Devices) | 7502
- Syntax (PTX10004, PTX10008, and PTX10016 Devices) | 7502
- Hierarchy Level | 7503
- Description | 7503
- Options | 7503
- Sample Output | 7504
- Required Privilege Level | 7504
- Platform-Specific `fpc (Chassis)` Behavior | 7504
- Release Information | 7505

Syntax

```
fpc fpc-slot {
    error
    max-power <power>
    power (off | on)
    pfe pfe_id power (off | on)
    no-vlan-tagging
    pic pic_slot port port_no no-temperature-monitoring
    mic slot power off
}
```

Syntax (ACX Series Devices)

```
fpc fpc-slot {
    cfp-to-et
    error
```

```
    lossless-flow-control
    pfe pfe_id
    {
       power (off | on)
       core core_id
       {
          channel channel_id
          tunnel-services
          {
             bandwidth bandwidth_speed
          }
       }
    }
    pic
    sampling-instance
    power (off | on)
    max-power <power>
    pfe pfe_id power (off | on)
}
```

Syntax (PTX10003 Devices)

```
fpc fpc-slot {
    error
    power (off | on)
    pfe pfe_id power (off| on)
    no-vlan-tagging
}
```

Syntax (PTX10004, PTX10008, and PTX10016 Devices)

```
fpc fpc-slot {
    error
    pic <slot>
    power (off | on)
    pfe pfe_id power (off| on)
     {
        port <port_no> no-temperature-monitoring
     }
```

```
    no-vlan-tagging
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the FPC slot number. (ACX and PTX Series devices) Configure no VLAN tagging for a specific FPC. Additionally, configure power off configuration (PTX10003 devices only).

Options

Available command options vary based on the device model.

auto-speed detection	Disables automatic speed detection.
cfp-to-et	Enable ET interface and remove CFP client.
lossless-flow-control	Lossless buffer config for Ethernet Mac pause.
pfe *pfe_id* core *core_id* channel *channel_id* tunnel-services bandwidth *bandwidth_speed*	Select a bandwidth speed for your tunnel interface. Speed options start at 1 gigabit.
fpc *fpc_id*	For PTX devices, *fpc_id* value is 0 to 15.
pfe *pfe_id* power (off | on)	Individual PFE Power off configuration option is available for PTX10003 devices only. The reset-pfe error action is available for PFE power off and restart. • Slot range for FPC is 0 to 1 for PTX10003-8T. Slot range for FPC is 0 to 3 for PTX10003-16T. • PFE number range is 0 to 7 for all the FPCs. • If no power configuration is present, then PFE is considered powered ON.
mic *slot* power off	(Only for MX304) Do not provide power to MIC. To power on the MIC, you need to delete this power off configuration. Please note the following information before you commit the

changes: If a commit includes multiple LMICs or separate commits, each LMIC undergoes operations sequentially. When you try to commit back-to-back changes for the same LMIC before the completion of an operation, then the system performs the last commited action and ignores the intermediate actions.

no-vlan-tagging Individual PFE maximum power configuration option is available for PTX10003 devices only. Configure maximum power for the FPC slot.

pic Configure properties for an individual PIC.

no-temperature-monitoring Disable temperature monitoring to explicitly disable optics temperature sensors for EM policy support, on specific WAN ports.

Sample Output

set chassis fpc (MX304 with LMIC)

```
user@host# set chassis fpc 0 mic 0 power off
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Platform-Specific `fpc (Chassis)` Behavior

Use Feature Explorer to confirm platform, and release support for specific features.

Use the following table to review platform-specific behaviors for your platform:

Table 46:

Platform	Difference
MX304 devices	When you try to commit back-to-back changes for the same LMIC before the completion of an operation, then the system performs the last commited action and ignores the intermediate actions. A 120-second delay is enforced between consecutive online or offline requests, irrespective of the request type and LMIC slot. For example, mic 0 offline/online requires 180 seconds to finish. Only after 120 seconds can the mic 1 command begin. If a configuration command is committed within 120 seconds, it will succeed. However, the command's effect will only occur 120 seconds after the last LMIC online/offline completion.

Release Information

Statement introduced in Junos OS Evolved Release 24.1R1 for ACX Series Routers.

`power-off` option introduced for PTX10003 devices in Junos OS Evolved Release 22.4R1.

`tunnel-services` and `bandwidth` options introduced for ACX Series devices in Junos OS Evolved Release 24.2R1.

fpc (DDoS)

IN THIS SECTION

- Syntax | 7506
- Hierarchy Level | 7506
- Description | 7506
- Options | 7506
- Required Privilege Level | 7506
- Release Information | 7507

Syntax

```
fpc slot-number;
    bandwidth-scale percentage;
    burst-scale percentage;
    disable-fpc;
}
```

Hierarchy Level

- For MX Series routers, T4000 routers, and EX9200 switches:

```
[edit system ddos-protection protocols protocol-group  (aggregate | packet-type)]
```

- For PTX Series routers and QFX Series switches:

```
[edit system ddos-protection protocols protocol-group  (aggregate | packet-type)]
```

Description

Modify DDoS protection aggregate or packet-type policer default values on the specified line card.

This configuration statement is supported on MX Series routers with MPCs, T4000 routers with FPC5s, PTX Series routers, EX9200 switches, and QFX Series switches.

Options

slot-number—Slot number of the card. On fixed form-factor devices, specify the line card as FPC 0.

- **Range:** Depends on the router or switch model

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers

fpc error

IN THIS SECTION

- Syntax | 7507
- Hierarchy Level | 7508
- Description | 7508
- Options | 7508
- Required Privilege Level | 7509
- Release Information | 7509

Syntax

```
fpc slot number {
    error {
        (fatal | major | minor) {
        threshold threshold-value;
        action (alarm | disable-pfe | offline-pic | log | get-state |offline | reset);
        }
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the threshold at which FPC errors will take the action you configure to be performed by the device.

Some devices include an internal framework for detecting and correcting FPC errors that can have the potential to affect services. For each FPC on the device, you can classify errors according to severity, set an automatic recovery action for each severity, and set a threshold (the number of times the error must occur before the action is triggered).

Options

- `fatal`—Fatal error on the FPC. An error that results in the blockage of a considerable amount of traffic across modules is a fatal error.
- `major`—Major error on the FPC. An error that results in continuing loss of packet traffic but does not affect other modules is a major error.
- `minor`—Minor error on the FPC. An error that results in the loss of a small number of packets but is fully recoverable is a minor error.
- `threshold` *`threshold-value`*—Configure the threshold value at which to take action. If the severity level of the error is fatal, the action is carried out only once when the total number of errors exceeds the threshold value. If the severity level of the error is major, the action is carried out once after the occurrence exceeds the threshold. If the severity level is minor, the action is carried out as many times as the value specified by the threshold. For example, when the severity level is minor, and you have configured the threshold value as 10, the action is carried out after the tenth occurrence.

  > **NOTE**: You can set the threshold value to 0 for errors with a severity level of minor. This implies that no action is taken for that error. You cannot set the threshold value to 0 for errors with a severity level of major or fatal.

  Default: The error count for fatal and major actions is 1. The default error count for minor actions is 10.

  Range: 0—429,496,729

The available detection and recovery actions are as follows:

- `alarm`—Raise an alarm.
- `disable-pfe`—Disable the Packet Forwarding Engine interfaces on the FPC.
- `get-state`—Get the current state of the FPC.
- `log`—Generate a log for the event.
- `offline`—Take the FPC offline.
- `offline-pic`—Take the PIC (installed in the FPC) offline.
- `reset`—Reset the FPC.

Required Privilege Level

routing—To view this in the configuration.

routing-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D40. This command is supported.

RELATED DOCUMENTATION

error (Security) | 6446

fpc-feb-connectivity

IN THIS SECTION

- Syntax | 7510
- Hierarchy Level | 7510
- Description | 7510
- Options | 7510
- Required Privilege Level | 7510
- Release Information | 7510

Syntax

```
fpc-feb-connectivity {
    fpc number feb (slot-number | none);
}
```

Hierarchy Level

```
[edit chassis]
```

Description

On the M120 router only, configure a connection between any Flexible PIC Concentrator (FPC) and any Forwarding Engine Board (FEB).

Options

fpc ***number***—Specify the FPC slot number.

- **Range:** 0 through 5

feb ***slot-number***—Specify the FEB slot number.

- **Range:** : 0 through 5

none—Disconnect the FPC from the FEB.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring the Junos OS to Support FPC to FEB Connectivity on M120 Routers

fpc (Interfaces)

IN THIS SECTION

- Syntax | 7511
- Hierarchy Level | 7511
- Description | 7511
- Options | 7512
- Required Privilege Level | 7512
- Release Information | 7512

Syntax

```
fpc slot-number ;
```

Hierarchy Level

```
[edit interfaces pic-set pic-set-name]
```

Description

Sets the PIC bundle and the FPC slot.

The `pic-set` bundles all the PICs and corresponding logical interfaces. A PIC can only join only one pic-bundle, and cannot join multiple pic-bundles at same time. When the `pic-set` configuration changes, all the logical interfaces related to the PIC should be synchronized to all member IOC.

Options

- *apply-groups*—Inherit configuration data from these groups.
- *apply-groups-except*—Do not inherit configuration data from these groups.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Interfaces

interface (PIC Bundle) | 8949

fpc (Junos Fusion)

IN THIS SECTION

- Syntax | 7513
- Hierarchy Level | 7513
- Description | 7513
- Options | 7513
- Required Privilege Level | 7514
- Release Information | 7514

Syntax

```
fpc slot-id{
    alias alias;
    cascade-ports interface-name;
    description description;
    environment-monitoring-policy policy;
    serial-number serial-number;
    system-id mac-address;
    uplink-failure-detection {
        candidate-uplink-policy policy;
    }
    local switching;
    selective-vlan-switching{
        routing-instance routing-instance;
    }
}
```

Hierarchy Level

```
[edit chassis satellite-management]
```

Description

Configure an FPC identifier for a satellite device within a Junos Fusion, or modify the configuration of an existing satellite device in a Junos Fusion.

Options

slot-id

Specifies the FPC identifier of the device and functions as the FPC identifier in the interface name when configuring satellite device interfaces.

In a Junos Fusion Enterprise or Junos Fusion Provider Edge, the *slot-id* must have a value of 34 or greater.

NOTE: MPC10E and MPC11E are not supported in Junos Fusion, but can coexist with other MPC line cards in the MX router chassis that conncect

to satellite devices. When using MPC10E/MPC11E line cards in an MX series router that is supporting Junos Fusion, you must use a value range of 160 - 252 as the FPC identifier for the satellite device interfaces.

Junos Fusion does not support the hyper mode feature, so you must also use the `set forwarding-options no-hyper-mode` statement.

local switching Enables local-switching for all the ports on the satellite device.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

`local-switching` and `selective-vlan-switching` introduced in Junos OS Release 17.2R1 for Junos Fusion Provider Edge.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

fpc (Junos Fusion Satellite Device Cluster)

IN THIS SECTION

- Syntax | 7515
- Hierarchy Level | 7515
- Description | 7515

- Options | 7515
- Required Privilege Level | 7515
- Release Information | 7516

Syntax

```
fpc slot-id{
    alias alias;
    description description;
    member-id member-id-number;
    system-id mac-address;
}
```

Hierarchy Level

```
[edit chassis satellite-management cluster cluster-name]
```

Description

Configure an FPC identifier for a satellite device in a satellite device cluster for a Junos Fusion, or modify the configuration of an existing satellite device in a satellite device cluster in a Junos Fusion.

Options

slot-id Specifies the FPC identifier of the device.

In a Junos Fusion, the *slot-id* must be 34 or larger, and functions as the FPC identifier in the interface name when configuring satellite device interfaces.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

fpc (M320, T320 and T640 Devices)

IN THIS SECTION

- Syntax | 7516
- Hierarchy Level | 7518
- Description | 7518
- Options | 7518
- Required Privilege Level | 7518
- Release Information | 7518

Syntax

```
fpc slot-number {
  fpc
            slot slot-number

{  pfe pfe-instance
        { temp-perf-throttle-disable;
          temp-volt-reduction-disable;
          power (off | on);
        }

{  error {
        [fatal | major | minor] {
            threshold threshold-value;
```

```
            action (alarm | disable-pfe | offline-pic | log | get-state |offline | reset | trap
| offline-pfe | reset-pfe);
        }
    }
    optical-options {
        expansion-card {
        fpc fpc-slot;
        }
        express-in {
            fpc fpc-slot;
            }
        tca tca-identifier (enable-tca | no-enable-tca) (threshold number | threshold-24hrs
number) ;
        wavelength nm{
            switch interface-name{
            }
            wss-express-in fpc-slot;
            }
        }
    }
    pic pic-number {
        ce1 {
            e1 port-number {
                channel-group group-number timeslots slot-number;
            }
        }
        ct3 {
            port port-number {
                t1 link-number {
                    channel-group group-number timeslots slot-number;
                }
            }
        }
        framing (sdh | sonet);
        idle-cell-format {
            itu-t;
            payload-pattern payload-pattern-byte;
        }
        max-queues-per-interface (8 | 4);
        no-concatenate;
        q-pic-large-buffer (large-scale | small-scale);
```

```
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure properties for the PICs in individual Flexible PIC Concentrators (FPCs).

Options

slot-number—Slot number in which the FPC is installed.

- **Range:** M320, T640, T1600, T4000, and PTX5000 , PTX10008 routers: 0 through 7. On PTX1000 routers, the FPC number is always 0.
- **Range:** PTX3000 routers: 0, 2, 4, 6, 8, 10, 12, 14

pfe—Slot number in which the PFE is installed.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Error statement introduced for PTX Series routers in Junos OS Release 13.3.

RELATED DOCUMENTATION

Configuring FPC Error Levels and Actions

Automatic Performance Throttling of PFEs

fpc (MX Series 5G Universal Routing Platforms)

IN THIS SECTION

- Syntax | 7519
- Hierarchy Level | 7520
- Description | 7520
- Options | 7520
- Required Privilege Level | 7521
- Release Information | 7521

Syntax

```
fpc slot-number {
    inline-services {
        flow-table-size {
            ipv4-flow-table-size units;
            ipv4-flow-table-size  units;
            ipv6-extended-attrib;
        }
        }
    ir-mode (R | IR);
    pfe identifier {
      exception-reporting {
         category category-name {
            inline-monitoring-instance inline-monitoring-instance;
         }
      }
      forwarding-packages {
         mobility;
      }
      power (off | on);
      tunnel-services;
    }
    pic number {
        inline-services {
            bandwidth  (1g | 10g);
```

```
        }
        port-mirror-instance port-mirroring-instance-name-pic-level;
        tunnel-services {
            bandwidth (1g | 10g)
        }
    }
    port-mirror-instance port-mirroring-instance-name-fpc-level;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure properties for the DPC or MPC and corresponding Packet Forwarding Engines to create tunnel interfaces.

(MX Series Virtual Chassis only) When you configure chassis properties for MPCs installed in a Virtual Chassis member router, statements included at the `[edit chassis member` *`member-id`* `fpc slot` *`slot-number`*`]` hierarchy level apply to the MPC in the specified slot number only on the specified member router in the Virtual Chassis. Statements included at the `[edit chassis fpc slot` *`slot-number`*`]` hierarchy level apply to the MPCs in the specified slot number on *each* member router in the Virtual Chassis.

> **BEST PRACTICE**: To ensure that the statement you use to configure MPC chassis properties in an MX Series Virtual Chassis applies to the intended member router and MPC, we recommend that you always include the `member` *`member-ID`* option before the `fpc` statement, where *`member-id`* is 0 or 1 for a two-member MX Series Virtual Chassis.

Options

`fpc` *`slot-number`*—Specify the slot number of the DPC.

- **Range:** 0 through 11

`pic` *`number`*—Specify the number of the Packet Forwarding Engine. Each DPC includes four Packet Forwarding Engines.

- **Range:** 0 through 4

`port-mirror-instance` *`port-mirroring-instance-name-fpc-level`*—Associate a port-mirroring instance with the DPC and its corresponding PICs. The port-mirroring instance is configured under the `[edit forwarding-options port-mirroring]` hierarchy level.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

`port-mirror-instance` option added in Junos OS Release 9.3.

`ipv6-extended-attrib` option added in Junos OS Release 14.2 for MX Series routers.

`exception-reporting` option added in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

- Configuring Port-Mirroring Instances on MX Series 5G Universal Routing Platforms
- Virtual Chassis Components Overview

fpc-nmi-volt-fail-knob

IN THIS SECTION

- Syntax | 7522
- Hierarchy Level | 7522
- Description | 7522
- Required Privilege Level | 7522
- Release Information | 7522

Syntax

```
fpc-nmi-volt-fail-knob (enable | disable)
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable or disable the non maskable interrupt (NMI) for the voltage failure errors on the flexible pic concentrator (FPC).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 10.4R15, 11.4R8-S2, 11.4R9, 12.1R8, 12.2R6, 12.3R3-S1, 12.3R4, 13.1R3, and 13.2R1

RELATED DOCUMENTATION

show chassis fpc | **18311**

Configuring Voltage Level Monitoring of FPCs

fpc (Notification Control)

IN THIS SECTION

- Syntax | 7523
- Hierarchy Level | 7523
- Description | 7523
- Default | 7523
- Options | 7523
- Required Privilege Level | 7524
- Release Information | 7524

Syntax

```
fpc slot-number {
    disable;
}
```

Hierarchy Level

```
[edit poe notification-control]
```

Description

Enable the specified PoE controller to generate PoE traps.

Default

PoE traps are disabled by default.

Options

slot-number—Indicates the PoE controller by FPC slot number, where ***slot-number*** is:

- **0**—On an EX2200, EX3200, standalone EX3300, standalone EX4200, or standalone EX4300 switch
- Member ID—On an EX3300, EX4200, or EX4300 switch in a Virtual Chassis
- Line card slot number—On an EX6200 or EX8200 switch

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring PoE Interfaces on EX Series Switches

fpc-offline-on-blackholing

IN THIS SECTION

- Syntax | 7524
- Hierarchy Level | 7525
- Description | 7525
- Required Privilege Level | 7525
- Release Information | 7525

Syntax

```
fpc-offline-on-blackholing;
```

Hierarchy Level

```
[edit chassis fabric degraded]
```

Description

Take the FPC offline and raise an alarm if a traffic null route condition is detected in the routing matrix. By default, FPCs remain online when a null route condition is detected.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

show chassis alarms | 17304

fpc (PTX Series Devices)

IN THIS SECTION

- Syntax (PTX Series Devices) | 7526
- Hierarchy Level | 7526
- Description | 7526
- Options | 7526
- Required Privilege Level | 7527
- Release Information | 7527

Syntax (PTX Series Devices)

```
fpc slot-number {
    error
    pfe
    pic
    sampling-instance instance_string;
    traffic-manager {
        (enhanced-priority-mode | no-enhanced-priority-mode)
        }
    optical-options
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure properties for the PICs in individual Flexible PIC Concentrators (FPCs). Configure FPC Restart, offline and online (PTX10003 devices only).

Options

slot-number—Slot number in which the FPC is installed.

- **Range:** From 0 to maximum slot number based on the product type.

  For example:

  - For PTX1000 and PTX10001-36MR devices, the FPC slot number is always 0.
  - For PTX10008 device, the slot number range is 0 to 7.
  - For PTX10016 device, the range is 0 to 15,
  - For PTX10003 device, FPC Restart, offline and online options are available.
  - For PTX3000 devices, alternate slot numbers 0, 2, 4, 6, 8, 10, 12, 14).

`pfe`—Slot number in which the PFE is installed. If FPC power off is set along with PFE power ON, then PFE power ON is not initiated.

`sampling-instance` *`instance_string`*— You can configure active sampling using a sampling instance and associate that sampling instance to a Flexible Port Concentrator (FPC), Modular Port Concentrator (MPC), or Dense Port Concentrator (DPC). In addition, you can define multiple sampling instances associated with multiple destinations and protocol families per sampling instance destination.

Sampling-instances are defined under the `[edit forwarding-options]` hierarchy.

`traffic-manager`— You can configure traffic manager attributes for the FPC. Available options are `enhanced-priority-mode` and `no-enhanced-priority-mode`.

The remaining statements are explained separately, for example ,"error " on page 6441.

> **NOTE**: The `fpc <fpc-slot> pfe <pfe-inst> power-off` configuration is added to power-off specific PFE instances when an FPC is brought online or to power-off a PFE during runtime.

> **NOTE**: Click on the highlighted links in the *Syntax* section to view details of other available configuration options. Using the configuration command for error levels, you can isolate PFE errors, thereby reducing the need for a field replacement of FPCs. For feature details and PTX platforms supported, see FPC self-healing.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Error statement introduced for PTX Series devices in Junos OS Release 13.3.

RELATED DOCUMENTATION

Configuring FPC Error Levels and Actions

Automatic Performance Throttling of PFEs

Managing FPC Errors

fpc-restart

IN THIS SECTION

- Syntax | 7528
- Hierarchy Level | 7528
- Description | 7528
- Default | 7528
- Required Privilege Level | 7529
- Release Information | 7529

Syntax

```
fpc-restart;
```

Hierarchy Level

```
[edit chassis fabric degraded]
```

Description

Allow the user to restart the FPCs when a traffic null route condition is detected in the routing matrix. To enable this feature set the `fpc-restart` statement at the `edit chassis fabric degraded` hierarchy level.

Default

FPCs are not restarted when a null route condition is detected.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 13.2R6.

RELATED DOCUMENTATION

Fabric Resiliency and Degradation

Disabling Line Card Restart to Limit Recovery Actions from Degraded Fabric Conditions

fpc-resync

IN THIS SECTION

- Syntax | 7529
- Hierarchy Level | 7529
- Description | 7530
- Required Privilege Level | 7530
- Release Information | 7530

Syntax

```
fpc-resync;
```

Hierarchy Level

```
[edit chassis]
```

Description

(On M320, T320, T640, T1600, T4000, TX Matrix, and TX Matrix Plus routers only) When a Flexible PIC Concentrator (FPC) is brought online, resynchronize the sequence numbers of the FPC with the other active FPCs.

> **NOTE**: In order to prevent null-route filtering, the `fpc-resync` command will have no effect if a single LMNR based FPC and one or more I-chip FPCs exist in the same chassis.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Resynchronizing FPC Sequence Numbers with Active FPCs when an FPC Comes Online

fpc-slice

IN THIS SECTION

- Syntax | 7531
- Hierarchy Level | 7531
- Description | 7531
- Options | 7531
- Required Privilege Level | 7532
- Release Information | 7532

Syntax

```
fpc-slice fpc fpc-slot {
    slc slc-id {
        pfe-id-list pfe-id-list;
        cores number-of-cores;
        dram dram-value;
        }
    }
```

Hierarchy Level

```
[edit chassis network-slices guest-network-functions gnf gnf-id]
```

Description

Configure smaller logical partitions or sub line cards (SLC) of a line card, allocate resources to the SLCs, and assign each SLC to a guest network function (GNF).

This feature is applicable only to the MPC11E line card. An MPC11E line card supports two sub line cards (SLCs). You can assign each GNF with an SLC, which functions like an independent line card.

NOTE:

- You must configure all the SLCs supported by a line card, assign the required resources such as core, DRAM, and Packet forwarding Engines the SLCs before committing the configuration. You must allocate all the Packet forwarding Engines on the line card to the SLCs.

Options

fpc-slot The slot number of the line card on which you want to configure SLCs.

slc *slc-id* Specify the sub line card (SLC) ID which you want to assign to the selected GNF.

NOTE: Do not assign:

- two or more SLCs of the same line card to the same GNF.
- the same SLC of a line card to more than one GNF.

pfe-id-list *pfe-id-list*

Specify the Packet Forwarding Engine IDs to be allocated to the SLC. You must allocate all the Packet Forwarding Engines on a line card to the SLCs.

cores *number-of-cores*

Specify the number of cores to be allocated to the SLC. Currently, number of cores supported is 4.

dram *dram-value*

Specify the DRAM to be allocated to the SLC. The MPC11E line card supports the following combinations:

SLC1 DRAM (GB)	SLC2 DRAM (GB)	Sub Total (GB)
13	13	26
9	17	26

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

request chassis fpc slc | 14488

fpc (Subscriber Limits)

IN THIS SECTION

- Syntax | 7533
- Hierarchy Level | 7533
- Description | 7533
- Options | 7534
- Required Privilege Level | 7534
- Release Information | 7534

Syntax

```
fpc slot-number {
    limit limit;
    pic number {
        limit limit;
        port number {
            limit limit;
        }
    }
}
```

Hierarchy Level

```
[edit system services resource-monitor subscribers-limit client-type name]
```

Description

Configure the maximum number of subscribers of a client type allowed to be logged in on the MPC in the specified slot. When that number is reached, subsequent logins on the card are denied until the current number of subscribers drops below the maximum allowed. You can also specify the maximum number of subscribers of a client type allowed per port, per MIC, and per chassis.

Options

limit	Maximum number of subscribers. • **Range:** 1 through 256,000
slot-number	Number of the MPC slot in the chassis.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Limiting Subscribers by Client Type and Hardware Element with Resource Monitor

Resource Monitoring for Subscriber Management and Services

fpc (TX Matrix and TX Matrix Plus Routers)

IN THIS SECTION

- Syntax | 7535
- Hierarchy Level | 7535
- Description | 7535
- Options | 7535
- Required Privilege Level | 7535
- Release Information | 7536

Syntax

```
fpc slot-number {
    pic pic-number {
        atm-cell-relay-accumulation;
        atm-l2circuit-mode (cell | aal5 | trunk trunk);
        framing (sdh | sonet);
        idle-cell-format {
            itu-t;
            payload-pattern payload-pattern-byte;
        }
        max-queues-per-interface (8 | 4);
        no-concatenate;
        no-mcast-replication;
        q-pic-large-buffer (large-scale | small-scale);
    }
}
```

Hierarchy Level

```
[edit chassis lcc number]
```

Description

On a TX Matrix or TX Matrix Plus router, configure properties for the PICs in individual FPCs.

Options

slot-number—Slot number in which the FPC is installed.

- **Range:** 0 through 7

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- TX Matrix Router and T640 Router Configuration Overview
- TX Matrix Plus Router Configuration Overview
- Configuring the Junos OS to Enable SONET/SDH Framing for SONET/SDH PICs
- TX Matrix Router Chassis and Interface Names
- TX Matrix Plus Router Chassis and Interface Names

fpcs (Node Slicing)

IN THIS SECTION

- Syntax | 7536
- Hierarchy Level | 7536
- Description | 7537
- Options | 7537
- Required Privilege Level | 7537
- Release Information | 7537

Syntax

```
fpcs fpcs;
```

Hierarchy Level

```
[edit chassis network-slices guest-network-functions gnf]
```

Description

Assign FPCs to a guest network function.

> **NOTE**: A given FPC cannot be assigned to more than one GNF.

Options

fpcs The FPC to be assigned to a guest network function.

- **Range:** `0-63`

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

network-slices | **11322**

guest-network-functions | **7949**

gnf | **7746**

control-plane-bandwidth-percent (Node Slicing) | **4634**

description (GNF) | **5332**

fpcs (NSSU Upgrade Groups)

IN THIS SECTION

- Syntax | 7538
- Hierarchy Level | 7538
- Description | 7538
- Options | 7539
- Required Privilege Level | 7539
- Release Information | 7539

Syntax

```
fpcs (slot-number | [list-of-slot-numbers]);
```

Hierarchy Level

```
[edit chassis nssu upgrade-group group-name],
[edit chassis nssu upgrade-group group-name member member-id]
```

Description

Configure switch line cards, Virtual Chassis member switches, or Virtual Chassis Fabric (VCF) member switches as part of an NSSU upgrade group.

To reduce the time an NSSU takes, you can configure line-card upgrade groups for an EX6200 or EX8200 switch with redundant Routing Engines; an EX8200 Virtual Chassis; an EX4650 Virtual Chassis with more than three member switches; QFX3500, QFX3600, and QFX5100 Virtual Chassis; or a QFX5100 Virtual Chassis Fabric (VCF). NSSU upgrades the devices in the order in which you configure the upgrade groups, so you can also use upgrade groups to control the upgrade sequence.

For switches that have separate line cards, use this statement to assign one or more line cards to an NSSU upgrade group based on their line-card slot numbers.

For Virtual Chassis or VCF member switches that do not have separate line cards, use this statement to assign one or more Virtual Chassis or VCF members to an NSSU upgrade group by specifying their member IDs.

> **NOTE**: For a Virtual Chassis or VCF, you do not use this statement with the `member` option. When to use the `member` statement hierarchy is explained next.

To configure an upgrade group that includes line cards on switches that support multiple line cards and comprise a Virtual Chassis, use this statement with the `member` option to specify the Virtual Chassis member ID and the desired line card slot number or numbers on that member switch to include in the upgrade group. Use multiple statements to add line cards from different Virtual Chassis members to the upgrade group.

Options

list-of-slot-numbers A list of slot numbers of multiple line cards or member IDs of Virtual Chassis or VCF members to be included in the upgrade group. Separate multiple slot numbers or member IDs with spaces and enclose the list in square brackets—for example: [3 4 7].

slot-number The slot number of a single line card or member ID of a Virtual Chassis or VCF member to be included in the upgrade group.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches

Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade

fragment (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 7540
- Hierarchy Level | 7540
- Description | 7540
- Required Privilege Level | 7540
- Release Information | 7540

Syntax

```
fragment;
```

Hierarchy Level

```
[edit services screen ids-option screen-name icmp]
```

Description

Identify and drop ICMP packets that are IP fragments. These are considered suspicious packets because ICMP packets are usually short. When the target receives these packets, the results can range from processing packets incorrectly to crashing the entire system.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

fragment-limit

IN THIS SECTION

- Syntax | 7541
- Hierarchy Level | 7541
- Description | 7541
- Options | 7541
- Required Privilege Level | 7542
- Release Information | 7542

Syntax

```
fragment-limit number-of-fragments;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
[edit security flow]
```

Description

Configure the maximum number of fragments permitted in a packet before the packet is dropped.

Options

number-of-fragments—Maximum number of fragments permitted.

- **Range:** 1 to 250 fragments.

- **Default:** 64 fragments.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

Statement added in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480, and MX960 routers.

RELATED DOCUMENTATION

Configuring Fragmentation Control for MS-DPC and MS-PIC Service Interfaces

fragment-threshold

IN THIS SECTION

- Syntax | 7542
- Hierarchy Level | 7543
- Description | 7543
- Options | 7543
- Required Privilege Level | 7543
- Release Information | 7543

Syntax

```
fragment-threshold bytes;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options],
[edit interfaces (ls-fpc/pic/port| ml-fpc/pic/port) unit logical-unit-number],
[edit logical-systems logical-system-name interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-
options],
[edit logical-systems logical-system-name interfaces (ls-fpc/pic/port| ml-fpc/pic/port) unit
logical-unit-number]
```

Description

For multilink and link services interfaces only, set the fragmentation threshold, in bytes.

Options

bytes—Maximum size, in bytes, for multilink packet fragments. Any nonzero value must be a multiple of 64 bytes.

- **Range:** 128 through 16,320 bytes
- **Default:** 0 bytes (no fragmentation)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Example: Configuring a Multilink Interface with MLPPP*

fragment-threshold (Forwarding Class Maps)

IN THIS SECTION

- Syntax | 7544
- Hierarchy Level | 7544
- Description | 7544
- Default | 7544
- Options | 7545
- Required Privilege Level | 7545
- Release Information | 7545

Syntax

```
fragment-threshold bytes;
```

Hierarchy Level

```
[edit class-of-service fragmentation-maps forwarding-class class-name]
```

Description

For link services IQ (`lsq`) interfaces only, set the fragmentation threshold for an individual forwarding class.

Default

If you do not include this statement, the fragmentation threshold you set at the `[edit interfaces interface-name unit logical-unit-number]` or `[edit interfaces interface-name mlfr-uni-nni-bundle-options]` hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest maximum transmission unit (MTU) of all the links in the bundle.

Options

bytes—Maximum size, in bytes, for multilink packet fragments. Any nonzero value must be a multiple of 64 bytes.

- **Range:** 128 through 16,320 bytes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces

fragment-threshold (Interfaces LSQ)

IN THIS SECTION

- Syntax | 7546
- Hierarchy Level | 7546
- Description | 7546
- Options | 7546
- Required Privilege Level | 7546
- Release Information | 7546

Syntax

```
fragment-threshold bytes;
```

Hierarchy Level

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number],
[edit logical-systems logical-system-name interfaces lsq-fpc/pic/port unit logical-unit-number]
```

Description

For voice services interfaces, set the fragmentation threshold, in bytes.

Options

bytes—Maximum size, in bytes, for multilink packet fragments. The value must be a multiple of 64 bytes, because zero is also a multiple of 64 bytes.

- **Range:** 128 through 16,320 bytes
- **Default:** 0 bytes (no fragmentation)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Delay-Sensitive Packet Interleaving

fragmentation-map

IN THIS SECTION

- Syntax | 7547
- Hierarchy Level | 7547
- Description | 7547
- Default | 7547
- Options | 7547
- Required Privilege Level | 7548
- Release Information | 7548

Syntax

```
fragmentation-map map-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

For AS PIC link services IQ (`lsq`) and virtual LSQ redundancy (`rlsq`) interfaces, associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI.

Default

If you do not include this statement, traffic in all forwarding classes is fragmented.

Options

map-name—Name of the fragmentation map.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Fragmentation by Forwarding Class Overview

Configuring Fragmentation by Forwarding Class

Configuring Fragmentation by Forwarding Class

Example: Configuring Fragmentation by Forwarding Class

Configuring Drop Timeout Interval for Fragmentation by Forwarding Class

fragmentation-maps | 7548

fragmentation-maps

IN THIS SECTION

- Syntax | 7549
- Hierarchy Level | 7549
- Description | 7549
- Default | 7549
- Options | 7549
- Required Privilege Level | 7549
- Release Information | 7550

Syntax

```
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            drop-timeout milliseconds;
            fragment-threshold bytes;
            multilink-class number;
            no-fragmentation;
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

For Multiservices and Services PIC link services IQ (`lsq`) and virtual LSQ redundancy (`rlsq`) interfaces, define fragmentation properties for individual forwarding classes.

Default

If you do not include this statement, traffic in all forwarding classes is fragmented.

Options

map-name—Name of the fragmentation map.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Fragmentation by Forwarding Class Overview
- Configuring Fragmentation by Forwarding Class
- Example: Configuring Fragmentation by Forwarding Class
- Configuring Drop Timeout Interval for Fragmentation by Forwarding Class
- fragmentation-map

frame-error

IN THIS SECTION

- Syntax | 7550
- Hierarchy Level | 7551
- Description | 7551
- Options | 7551
- Required Privilege Level | 7551
- Release Information | 7551

Syntax

```
frame-error count;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile event link-event-rate],
[edit protocols oam link-fault-management interface interface-name event-thresholds]
```

Description

Configure the threshold value for sending frame error events or taking the action specified in the action profile.

Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value.

The window or period during which frame errors are counted is 5 seconds or multiples of it (with a maximum value of 1 minute). This window denotes the duration as intervals of 100 milliseconds, encoded as a 16-bit unsigned integer. This window is not configurable in Junos OS. According to the IEEE 802.3ah standard, the default value of the frame-errors window is 1 second. This window has a lower bound of 1 second and an upper bound of 1 minute.

Options

count—Threshold count for frame error events.

- **Range:** 0 through 100

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

> *Configuring Threshold Values for Local Fault Events on an Interface*
>
> *Configuring Threshold Values for Fault Events in an Action Profile*

frame-mode (Dynamic Traffic Shaping)

IN THIS SECTION

- Syntax | 7552
- Hierarchy Level | 7552
- Description | 7552
- Default | 7552
- Options | 7553
- Required Privilege Level | 7553
- Release Information | 7553

Syntax

```
frame-mode (bytes | $junos-cos-byte-adjust | frame-mode-bytes frame-mode-bytes |$junos-cos-byte-
adjust-frame);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name overhead-
accounting],
[edit class-of-service traffic-control-profiles profile-name overhead-accounting],
```

Description

Configure the mode to shape downstream ATM traffic based as frames.

Default

The default is `frame-mode`.

Options

bytes	Byte adjustment value for the `cell-mode` or `frame-mode` shaping options.
$junos-cos-byte-adjust	Predefined variable for byte adjustment that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.
frame-mode-bytes *frame-mode-bytes*	Overhead bytes when in frame-mode. Traffic shaping is based on the number of bytes in the frame, without regard to cell encapsulation or padding overhead.
$junos-cos-byte-adjust-frame	Predefined variable for frame mode shaping. This variable can not be used when the `overhead-accounting bytes` *bytes* option is configured. • **Range:** –120 through 124 bytes

BEST PRACTICE: We recommend using the `frame-mode-bytes` *frame-mode-bytes* option rather than the `bytes` option.

NOTE: If you specify a value for the `bytes` *bytes* option, you cannot specify a value for either the `frame-mode-bytes` option.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Variable *$junos-cos-byte-adjust-frame* introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

CoS Adjustment Control Profiles Overview

Configuring CoS Adjustment Control Profiles

adjustment-control-profiles | **1961**

Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates

Bandwidth Management for Downstream Traffic in Edge Networks Overview

egress-shaping-overhead

bytes (Dynamic Traffic Shaping) | **3121**

cell-mode (Dynamic Traffic Shaping) | **3245**

frame-period

IN THIS SECTION

- Syntax | **7554**
- Hierarchy Level | **7554**
- Description | **7554**
- Options | **7555**
- Required Privilege Level | **7555**
- Release Information | **7555**

Syntax

```
frame-period count;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile event link-event-rate],
[edit protocols oam link-fault-management interface interface-name event-thresholds]
```

Description

Configure the number of frame errors within the last N frames that has exceeded a threshold. Threshold for sending frame period error events or taking the action specified in the action profile.

A frame error is any frame error on the underlying physical layer. The frame period threshold is reached when the number of frame errors reaches the configured value within the period window. The default period window is the number of minimum-size frames that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.

Currently, Junos does not generate `frame-period` event.

Options

count—Threshold count for frame period error events.

- **Range:** 0 through 100

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring Threshold Values for Local Fault Events on an Interface

Configuring Threshold Values for Fault Events in an Action Profile

frame-period-summary

IN THIS SECTION

- Syntax | 7556
- Hierarchy Level | 7556
- Description | 7556
- Options | 7556

Required Privilege Level | 7556
Release Information | 7556

Syntax

```
frame-period-summary count;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile event link-event-rate],
[edit protocols oam link-fault-management interface interface-name event-thresholds]
```

Description

Configure the threshold value for sending frame period summary error events or taking the action specified in the action profile.

An errored frame second is any 1-second period that has at least one errored frame. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period window. The default window is 60 seconds. The window is not configurable.

Options

count—Threshold count for frame period summary error events.

- **Range:** 0 through 100

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring Threshold Values for Local Fault Events on an Interface

Configuring Threshold Values for Fault Events in an Action Profile

frame-relay-de (CoS Interfaces)

IN THIS SECTION

- Syntax | 7557
- Hierarchy Level | 7557
- Description | 7557
- Options | 7558
- Required Privilege Level | 7558
- Release Information | 7558

Syntax

```
frame-relay-de (name | default);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number loss-priority-maps],
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

Assign the loss priority map or the rewrite rule to a logical interface.

Options

- `default`—Apply default loss priority map or default rewrite rule. The default loss priority map contains the following settings:

```
loss-priority low code-point 0;
loss-priority high code-point 1;
```

- The default rewrite rule contains the following settings:

```
loss-priority low code-point 0;
loss-priority medium-low code-point 0;
loss-priority medium-high code-point 1;
loss-priority high code-point 1;
```

- *name*—Name of loss priority map or rewrite rule to be applied.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

frame-relay-de (CoS Loss Priority)

IN THIS SECTION

- Syntax | 7559
- Hierarchy Level | 7559
- Description | 7559
- Options | 7559

Required Privilege Level | 7559
Release Information | 7559

Syntax

```
frame-relay-de map-name {
    loss-priority level code-point (0 | 1);
}
```

Hierarchy Level

```
[edit class-of-service loss-priority-maps]
```

Description

Define a Frame Relay discard-eligible bit loss-priority map.

Options

- *level*—Level of loss priority to be applied based on the specified CoS values. The level can be **low**, **medium-low**, **medium-high**, or **high**.
- *map-name*—Name of the loss-priority map.

  The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

frame-relay-de (CoS Rewrite Rule)

IN THIS SECTION

- Syntax | 7560
- Hierarchy Level | 7560
- Description | 7560
- Options | 7560
- Required Privilege Level | 7561
- Release Information | 7561

Syntax

```
frame-relay-de rewrite-name {
    forwarding-class class-name {
        loss-priority level code-point (0 | 1);
    }
    import (default | rewrite-name);
}
```

Hierarchy Level

```
[edit class-of-service rewrite-rules]
```

Description

Define a Frame Relay discard-eligible bit rewrite rule.

Options

- *level*—Level of loss priority on which to base the rewrite rule. The loss priority level can be **low**, **medium-low**, **medium-high**, or **high**.

- *rewrite-name*—Name of a rewrite-rules mapping configured at the `[edit class-of-service rewrite-rules]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

frame-relay-de (Defining Loss Priority Maps)

IN THIS SECTION

- Syntax | 7561
- Hierarchy Level | 7562
- Description | 7562
- Options | 7562
- Required Privilege Level | 7562
- Release Information | 7562

Syntax

```
frame-relay-de name {
    loss-priority level code-points [ alias | bits ];
}
```

Hierarchy Level

```
[edit class-of-service loss-priority-maps]
```

Description

Define a Frame Relay discard eligibility (DE) bit loss priority map.

Options

name Name of the loss priority map.

loss-priority *level* Level of the loss priority to be applied based on the specified CoS values. The loss priority level can be one of the following:

- `high`—Packet has high loss priority.
- `low`—Packet has low loss priority.
- `medium-high`—Packet has medium-high loss priority.
- `medium-low`—Packet has medium-low loss priority.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Defining a Custom Frame Relay Loss Priority Map

frame-relay-de (Defining Loss Priority Rewrites)

IN THIS SECTION

- Syntax | 7563
- Hierarchy Level | 7563
- Description | 7563
- Options | 7563
- Required Privilege Level | 7564
- Release Information | 7564

Syntax

```
frame-relay-de name {
    loss-priority level code-point [ alias | bits ];
}
```

Hierarchy Level

```
[edit class-of-service loss-priority-rewrites]
```

Description

Define a Frame Relay discard eligibility (DE) bit loss priority rewrite.

Options

name Name of the loss priority rewrite.

loss-priority *level* Level of the loss priority to be applied based on the specified CoS values. The loss priority level can be one of the following:

- `high`—Packet has high loss priority.

- `low`—Packet has low loss priority.
- `medium-high`—Packet has medium-high loss priority.
- `medium-low`—Packet has medium-low loss priority.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

framed-ip-address

IN THIS SECTION

- Syntax | 7564
- Hierarchy Level | 7565
- Description | 7565
- Required Privilege Level | 7565
- Release Information | 7565

Syntax

```
framed-ip-address {
    equals value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from]
```

Description

Specify the RADIUS AVP Framed-IP-Address (IPv4) for the incoming RADIUS request from the subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

framed-ipv6-prefix

IN THIS SECTION

- Syntax | 7566
- Hierarchy Level | 7566
- Description | 7566
- Required Privilege Level | 7566
- Release Information | 7567

Syntax

```
framed-ipv6-prefix {
    equals value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from]
```

Description

Specify the RADIUS AVP Framed-IPv6-Prefix for the incoming RADIUS request from the subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

framed-pool

IN THIS SECTION

- Syntax | 7567
- Hierarchy Level | 7567
- Description | 7568
- Options | 7568
- Required Privilege Level | 7568
- Release Information | 7568

Syntax

```
framed-pool framed-pool;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp],
[edit access profile profile-name client client-name ppp]
```

Description

Configure the address pool.

Options

framed-pool—References a configured address pool.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

framing (10-Gigabit Ethernet Interfaces)

IN THIS SECTION

- Syntax | 7569
- Hierarchy Level | 7569
- Description | 7569
- Default | 7570
- Options | 7570
- Required Privilege Level | 7570
- Release Information | 7570

Syntax

```
framing (lan-phy | wan-phy);
precise-bandwidth;
```

Hierarchy Level

```
[edit interfaces xe-fpc/pic/port],[edit interfaces et-fpc/pic/port] (PTX Series Packet Transport
Routers and MX Series Routers)
```

Description

For routers supporting the 10-Gigabit Ethernet interface, configure the framing format. WAN PHY mode is supported on MX240, MX480, MX960, T640, T1600,T4000, and PTX Series Packet Transport Routers routers only.

NOTE:

- The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).
- On PTX Series routers, WAN PHY mode is supported only on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ .
- When the PHY mode changes, interface traffic is disrupted because of port reinitialization.
- In Junos OS Releases 17.4R2, 17.4R3, and later, on the following MPCs or routers, you cannot configure `wan-phy` mode at 10-Gbps, 40-Gbps, and 100-Gbps on a per-port basis:
  - MPC7E-10G, MPC7E-MRATE, MX2K-MPC8E, and MX2K-MPC9E

- MPC10003
- MX204 router
- JNP10K-LC2101 MPC

Default

Operates in LAN PHY mode.

Options

`lan-phy`—10GBASE-R interface framing format that bypasses the WIS sublayer to directly stream block-encoded Ethernet frames on a 10-Gigabit Ethernet serial interface.

`wan-phy`—10GBASE-W interface framing format that allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and SONET devices.

`precise-bandwidth`—Enables precise bandwidth for WAN-PHY interface framing format.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

The option `precise-bandwidth` introduced in Junos OS Release 19.3R1 for MX Series Routers.

RELATED DOCUMENTATION

Framing Overview

Configuring SONET Options for 10-Gigabit Ethernet Interfaces

framing (E1 Options)

IN THIS SECTION

- Syntax | 7571
- Hierarchy Level | 7571
- Description | 7571
- Options | 7571
- Required Privilege Level | 7571
- Release Information | 7572

Syntax

```
framing (g704 | g704-no-crc4);
```

Hierarchy Level

```
[edit chassis synchronization interfaces external e1-options]
[edit chassis synchronization interfaces (external-0/0 | external-1/0) e1-options]
```

Description

Configure the framing format for the E1 interface.

Options

g704 Set the G.704 framing format for E1 interfaces.

g704-no-crc4 Set the G.704 framing without CRC4 for E1 interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

framing (E1 Options for BITS Interfaces)

IN THIS SECTION

- Syntax | 7572
- Hierarchy Level | 7572
- Description | 7573
- Options | 7573
- Required Privilege Level | 7573
- Release Information | 7573

Syntax

```
framing (g704 | g704-no-cro4);
```

Hierarchy Level

```
[edit chassis synchronization interfaces bits e1-options]
```

Description

Configure the framing format for the building-integrated timing supply (BITS) interface when the signal type is E1. Valid only when `e1` is the `signal-type` configured at the [`edit chassis synchronization source interfaces` *`interface-name`* `bits signal-type`] hierarchy.

Options

`g704` G.704 framing format for E1 interfaces.

`g704-no-crc4` G.704 framing with no cyclic redundancy check 4 (CRC4) for E1 interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

External Clock Synchronization Overview for ACX Series Routers

Configuring External Clock Synchronization for ACX Series Routers

Synchronous Ethernet Overview

show chassis synchronization | **18879**

framing (T1 Options)

IN THIS SECTION

Syntax | **7574**

- Hierarchy Level | 7574
- Description | 7574
- Options | 7574
- Required Privilege Level | 7574
- Release Information | 7574

Syntax

```
framing (esf | sf);
```

Hierarchy Level

```
[edit chassis synchronization interfaces external t1-options]
[edit chassis synchronization interfaces (external-0/0 | external-1/0) t1-options]
```

Description

Configure the framing format for the T1 interface.

Options

esf Set the framing format as extended super frame.

sf Set the framing format as super frame.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

framing (T1 Options for BITS Interfaces)

IN THIS SECTION

- Syntax | 7575
- Hierarchy Level | 7575
- Description | 7575
- Options | 7575
- Required Privilege Level | 7576
- Release Information | 7576

Syntax

```
framing (esf | sf);
```

Hierarchy Level

```
[edit chassis synchronization interfaces bits t1-options]
```

Description

Configure the framing format for the building-integrated timing supply (BITS) interface when the signal type is T1 . Valid only when `t1` is the `signal-type` configured at the [`edit chassis synchronization source interfaces` *`interface-name`* `bits signal-type`] hierarchy.

Options

esf Extended superframe (ESF) mode for T1 interfaces.

sf Superframe (SF) mode for T1 interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

External Clock Synchronization Overview for ACX Series Routers

Configuring External Clock Synchronization for ACX Series Routers

Synchronous Ethernet Overview

show chassis synchronization | **18879**

free-fw-memory-watermark (Resource Monitor)

IN THIS SECTION

- Syntax | **7577**
- Hierarchy Level | **7577**
- Description | **7577**
- Options | **7577**
- Required Privilege Level | **7577**
- Release Information | **7577**

Syntax

```
free-fw-memory-watermark number;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Configure the percentage of free memory space used for firewall or filters to be monitored with a watermark value. You can configure the resource-monitoring capability on MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.

Options

number Percentage of free memory space used for firewall and filters to be monitored with a watermark value. When the configured watermark is exceeded, error logs are triggered. The default watermark values for the percentage of free ukernel or heap memory, next-hop memory, and firewall filter memory are 20.

- **Range:** 1 through 100

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

free-heap-memory-watermark (Resource Monitor)

IN THIS SECTION

- Syntax | 7578
- Hierarchy Level | 7578
- Description | 7578
- Options | 7578
- Required Privilege Level | 7579
- Release Information | 7579

Syntax

```
free-heap-memory-watermark number;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Configure the percentage of free memory space used for ukernel or heap (ASIC) memory to be monitored with a watermark value. You can configure the resource-monitoring capability on MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.

Options

number Percentage of free memory space used for ukernel or heap to be monitored with a watermark value. When the configured watermark is exceeded, error logs are triggered. The default watermark values for the percentage of free ukernel or heap memory, next-hop memory, and firewall filter memory are 20.

- **Range:** 1 through 100

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

free-nh-memory-watermark (Resource Monitor)

IN THIS SECTION

- Syntax | 7579
- Hierarchy Level | 7579
- Description | 7580
- Options | 7580
- Required Privilege Level | 7580
- Release Information | 7580

Syntax

```
free-nh-memory-watermark number;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Configure the percentage of free memory space used for next-hops to be monitored with a watermark value. The default value and the configured value of the watermark value for the percentage of free next-hop memory also applies to encapsulation memory. You can configure the resource-monitoring capability on MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.

Options

number Percentage of free memory space used for next-hops to be monitored with a watermark value. The NH memory watermark is applicable only for encapsulation memory (output WAN static RAM memory). When the configured watermark is exceeded, error logs are triggered. The default watermark values for the percentage of free ukernel or heap memory, next-hop memory, and firewall filter memory are 20.

- **Range:** 1 through 100

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

frequency-lock-threshold

IN THIS SECTION

- Syntax | 7581
- Hierarchy Level | 7581
- Description | 7581
- Options | 7581

Required Privilege Level | 7581
Release Information | 7582

Syntax

```
frequency-lock-threshold frequency-offset
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

The frequency lock threshold option controls the frequency offset in parts per billion (ppb) that the servo will lock with. Locking with a higher frequency lock threshold allows a quicker frequency lock, although at the expense of a higher phase offset per window in the phase aligned state. This option can be used for PTP applications that do not demand strict requirements that are typically needed with telecom applications.

Options

frequency-offset The frequency offset value in parts per billion (ppb).

- Range: 1 ppb through 15 ppb
- Default: 1 ppb

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1 for MX10008.

Supported on ACX7100-32C, ACX7100-48L, and ACX7024 in Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

Configuring the Precision Time Protocol G.8275.2 Enhanced Profile (Telecom Profile)

from (Application Aware Access List)

IN THIS SECTION

- Syntax | 7582
- Hierarchy Level | 7583
- Description | 7583
- Options | 7583
- Required Privilege Level | 7583
- Release Information | 7583

Syntax

```
from {
    application-group-any;
    application-groups [ application-group-names ];
    application-unknown;
    applications [ application-names ];
    destination-address address <any-unicast>;
    destination-address-range low minimum-value high maximum-value;
```

```
    destination-prefix-list list-name;
    nested-application-unknown;
    source-address address <any-unicast>;
    source-address-range low minimum-value high maximum-value;
    source-prefix-list list-name;
}
```

Hierarchy Level

```
[edit services aacl rule rule-name term term-name]
```

Description

Specify match conditions for the AACL term.

Options

For information on match conditions, see the description of firewall filter match conditions in the Routing Policies, Firewall Filters, and Traffic Policers User Guide.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring AACL Rules

from (services-ipsec-vpn)

IN THIS SECTION

- Syntax | 7584
- Hierarchy Level | 7584
- Description | 7584
- Options | 7584
- Required Privilege Level | 7585
- Release Information | 7585

Syntax

```
from {
    destination-address address;
    ipsec-inside-interface interface-name;
    source-address address;
}
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name]
```

Description

Specify input conditions for the IPsec term.

Options

For information on match conditions, see the description of firewall filter match conditions in the Routing Policies, Firewall Filters, and Traffic Policers User Guide.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

from (Firewall)

IN THIS SECTION

- Syntax | 7585
- Hierarchy Level | 7586
- Description | 7586
- Options | 7586
- Required Privilege Level | 7586
- Release Information | 7586

Syntax

```
from {
    match-conditions;
    egress-to-ingress;
}
```

Hierarchy Level

```
[edit firewall family family-name filter                filter-name term term-name]
```

Description

Match packet fields to values specified in a match condition. If the `from` statement is not included in a firewall filter configuration, all packets are considered to match and the actions and action modifiers in the `then` statement are implemented.

Options

match-conditions—Conditions that define the values or fields that the incoming or outgoing packets must contain for a match. You can specify one or more match conditions. If you specify more than one, they all must match for a match to occur and for the action in the `then` statement to be implemented.

egress-to-ingress—Include this option to increase the number of egress VLAN firewall filter terms from 1024 to 2048.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Option *egress-to-ingress* introduced in Junos OS Release 19.1R1 for the QFX5110.

RELATED DOCUMENTATION

> *Firewall Filter Match Conditions and Actions (QFX5100, QFX5110, QFX5120, QFX5200, EX4600, EX4650)*
>
> *Configuring Firewall Filters*
>
> *Understanding Firewall Filter Match Conditions*

from (Captive Portal Content Delivery Tags)

IN THIS SECTION

- Syntax | 7587
- Hierarchy Level | 7587
- Description | 7587
- Options | 7587
- Required Privilege Level | 7588
- Release Information | 7588

Syntax

```
from {
    destination-address address <except>;
}
```

Hierarchy Level

```
[edit services captive-portal-content-delivery rule rule-name term term-name]
```

Description

Specify one or more addresses in a CPCD rule term to identify traffic that can have header tags inserted into GET messages. Tag insertion is supported only for Routing-Engine-based, static CPCD.

Options

address Destination IPv4 or IPv6 address for traffic to be tagged.

except (Optional) Exclude the specified address from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Configuring Routing Engine-Based, Static HTTP Redirect Services

from (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 7588
- Hierarchy Level | 7589
- Description | 7589
- Required Privilege Level | 7589
- Release Information | 7589

Syntax

```
from {
    destination-address {
        (any-ipv4 | any-ipv4 except);
        (any-ipv6 | any-ipv6 except);
        (any-unicast | any-unicast except);
        (prefix | prefix except);
    }
    destination-address-range {
```

```
        high address low address <except>;
    }
    destination-port-range {
        high port-number low port-number;
    }
    destination-ports value;
    destination-prefix-list {
        (prefix-name | prefix-name except);
    }
}
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number]
```

Description

Specify the match criteria for the term in the tag rule. If all the conditions specified in the match criteria are met, then the HTTP header enrichment actions specified in the `then` statement are applied.

If you want the HTTP header enrichment actions specified in the `then` statement to be applied to all HTTP requests, do not include any matching conditions with the `from` statement.

NOTE: You must include a `from` statement in a tag rule.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm (HTTP Header Enrichment) | 8005

from (IDS MS-DPC)

IN THIS SECTION

- Syntax | 7590
- Hierarchy Level | 7590
- Description | 7591
- Options | 7591
- Required Privilege Level | 7591
- Release Information | 7591

Syntax

```
from {
    application-sets set-name;
    applications [ application-names ];
    destination-address (address | any-unicast) <except>;
    destination-address-range low minimum-value high maximum-value <except>;
    source-address (address | any-unicast) <except>;
    source-address-range low minimum-value high maximum-value <except>;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name]
```

Description

Specify input conditions for the IDS term when using the MS-DPC.

Options

For information on match conditions, see the description of firewall filter match conditions in the Routing Policies, Firewall Filters, and Traffic Policers User Guide.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Match Conditions in IDS Rules

from (PCC Rules)

IN THIS SECTION

- Syntax | 7592
- Hierarchy Level | 7592
- Description | 7592
- Required Privilege Level | 7592
- Release Information | 7593

Syntax

```
from {
    <application-groups [application-group-name]>;
    <applications [application-name]>;
    flows ([flow-identifier] | any);
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rule-name],
[edit services pcef pcc-rules rule-name]
```

Description

Specify the match criteria for the policy and charging control (PCC) rules. Any referenced SDF filter, application, or application group in the `from` statement must be configured.

If you are using Junos OS Subscriber Aware, specify the match criteria at the `[edit unified-edge pcef pcc-rules rule-name]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the match criteria at the `[edit services pcef pcc-rules rule-name]` hierarchy level.

> **NOTE**: You must include the `flows` statement. If you do not want to filter subscriber traffic based on service data flow (SDF) filters, use `flows any`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the [edit services pcef pcc-rules *rule-name*] hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Rules

from (PCP)

IN THIS SECTION

- Syntax | 7593
- Hierarchy Level | 7594
- Description | 7594
- Required Privilege Level | 7594
- Release Information | 7594

Syntax

```
from {
    application-sets set-name;
    applications [ application-name ];
    destination-address address <except>;
    destination-address-range high maximum-value low minimum-value <except>;
    destination-port  high maximum-value low minimum-value;
    destination-prefix-list list-name <except>;
```

```
    source-address address <except>;
    source-address-range high maximum-value low minimum-value <except>;
    source-prefix-list list-name <except>;
}
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name]
```

Description

Specify the match conditions for a PCP rule term. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

from (Protocols MPLS)

IN THIS SECTION

- Syntax | 7595
- Hierarchy Level | 7595
- Description | 7595
- Default | 7595
- Options | 7596
- Required Privilege Level | 7596
- Release Information | 7596

Syntax

```
from address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Specify the source address to use for the LSP.

The address you specify does not affect the outgoing interface used by the LSP.

Default

If you do not include this statement, the software automatically selects the loopback interface as the address.

Options

address—IP address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring the Ingress Router Address for LSPs*

from (Security NAT)

IN THIS SECTION

- Syntax | 7596
- Hierarchy Level | 7597
- Description | 7597
- Options | 7597
- Required Privilege Level | 7597
- Release Information | 7597

Syntax

```
from {
    interface [interface-name];
    routing-instance [routing-instance-name];
```

```
    zone [zone-name];
}
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name]
[edit security nat source rule-set rule-set-name]
[edit security nat static rule-set rule-set-name]
```

Description

Specify the source of the packet among the routing instance, interface, or zone.

Options

- `interface [interface-name]` —Name of the interface.
- `routing-instance [routing-instance-name]` —Name of the routing instance.
- `zone [zone-name]` —Name of the zone.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.3.

from (Services CoS)

IN THIS SECTION

Syntax | 7598

- Hierarchy Level | 7598
- Description | 7598
- Options | 7598
- Required Privilege Level | 7598
- Release Information | 7599

Syntax

```
from {
    applications [ application-name ];
    application-sets [ set-name ];
    destination-address address;
    source-address address;
}
```

Hierarchy Level

```
[edit services cos rule rule-name term term-name]
```

Description

Specify input conditions for a CoS term.

Options

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring CoS Rule Sets on Services PICs

from (Services NAT)

IN THIS SECTION

- Syntax | 7599
- Hierarchy Level | 7600
- Description | 7600
- Options | 7600
- Required Privilege Level | 7600
- Release Information | 7600

Syntax

```
from {
    application-sets set-name;
    applications [ application-names ];
    destination-address (address | any-unicast) <except>;
    destination-address-range low minimum-value high maximum-value <except>;
    destination-port range high maximum-value low minimum-value;
    source-address address (address | any-unicast) <except>;
    source-address-range low minimum-value high maximum-value <except>;
}
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name]
```

Description

Specify input conditions for the NAT term.

Options

For information on match conditions, see the description of firewall filter match conditions in the Routing Policies, Firewall Filters, and Traffic Policers User Guide.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

from (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 7601
- Hierarchy Level | 7601

- Description | 7601
- Options | 7601
- Required Privilege Level | 7602
- Release Information | 7602

Syntax

```
from {
    application-sets set-name;
    applications [ application-names ];
    destination-address (address | any-ipv4 | any-ipv6 | any-unicast) <except>;
    destination-address-range low minimum-value high maximum-value <except>;
    destination-prefix-list list-name <except>;
    source-address (address | any-ipv4 | any-ipv6 | any-unicast) <except>;
    source-address-range low minimum-value high maximum-value <except>;
    source-prefix-list list-name <except>;
}
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name term term-name]
```

Description

Specify input conditions for a stateful firewall term.

Options

For information on match conditions, see the description of firewall filter match conditions in the Routing Policies, Firewall Filters, and Traffic Policers User Guide.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Stateful Firewall Rules

from (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 7602
- Hierarchy Level | 7603
- Description | 7603
- Required Privilege Level | 7603
- Release Information | 7603

Syntax

```
from {
    apply-groups group-name;
    apply-groups-except group-name;
    destination-address address;
    destination-port port-number;
    dscp dscp-value;
    protocol protocol;
    source-address address;
```

```
    source-port port-number;
}
```

Hierarchy Level

```
[edit services radius-flow-tappolicy policy-name inet| inet6]
```

Description

Define the match criteria for the drop-policy rule.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Subscriber Secure Policy Overview

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

from (TDF Domain Selection)

IN THIS SECTION

Syntax | 7604

- Hierarchy Level | 7606
- Description | 7606
- Required Privilege Level | 7606
- Release Information | 7606

Syntax

```
from {
    3gpp-imsi {
        equals value;
        has-prefix value;
        has-suffix value;
        matches value;
    }
    attribute name {
        code numeric-code;
        vendor-id vendor-id;
        format {
            integer {
                equals {
                    value;
                }
                greater-than value;
                less-than value;
            }
            string {
                equals {
                    value;
                }
                has-prefix{
                    value;
                }
                has-suffix {
                    value;
                }
                matches {
                    value;
                }
```

```
            }
            time {
                equals {
                    value;
                }
                greater-than value;
                less-than value;
            }
            v4address {
                equals {
                    value;
                }
            }
            v6address {
                equals {
                    value;
                }
            }
            v6prefix {
                equals {
                    value;
                }
            }
        }
    }
    called-station-id {
        equals value;
        matches value;
    }
    calling-station-id {
        equals value;
        matches value;
    }
    class {
        equals value;
        has-prefix value;
        has-suffix value;
        matches value;
    }
    client client-name;
    framed-ip-address {
        equals value;
    }
```

```
    framed-ipv6-prefix {
        equals value;
    }
    nas-ip-address {
        equals value;
    }
    snoop-segment snoop-segment-name;
    user-name {
        equals value;
        has-prefix value;
        has-suffix value;
        matches value;
    }
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
```

Description

Specify the match criteria for the TDF domain selection or PCEF profile selection term.

> **NOTE**: For any term, the subscriber must match all the match conditions specified in a `from` statement. If you do not configure the `from` statement, then all subscribers are considered a match.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

from-zone (Security IDP Policy)

IN THIS SECTION

- Syntax | 7607
- Hierarchy Level | 7607
- Description | 7607
- Options | 7608
- Required Privilege Level | 7608
- Release Information | 7608

Syntax

```
from-zone (zone-name | any);
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match]
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match]
```

Description

Specify a source zone to be associated with the security policy. The default value is any.

Options

zone-name—Name of the source zone object.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

from-zone (Security Policies)

IN THIS SECTION

- Syntax | 7608
- Hierarchy Level | 7611
- Description | 7611
- Options | 7611
- Required Privilege Level | 7611
- Release Information | 7612

Syntax

```
from-zone zone-name to-zone zone-name {
    policy policy-name {
        description description;
        match {
            application {
                [junos-defaults | application];
                any;
                junos-smtps;
```

```
                junos-imaps;
                junos-pop3s;
                }
            }
            dynamic-application {
                [dynamic-application-name |dynamic-application-group-name];
                any;
                none;
            }
            destination-address {
                [address];
                any;
                any-ipv4;
                any-ipv6;
            }
            source-address {
                [address];
                any;
                any-ipv4;
                any-ipv6;
            }
            source-identity {
                [role-name];
                any;
                authenticated-user;
                unauthenticated-user;
                unknown-user;
            }
            source-end-user-profile {
                profile-name;
            }
        }
        scheduler-name scheduler-name;
        then {
                                count {
            alarm {
                per-minute-threshold number;
                per-second-threshold number;
            }
            }
            deny;
            log {
                session-close;
```

```
                session-init;
            }
            permit {
                application-services {
                    application-firewall {
                        rule-set rule-set-name;
                    }
                    application-traffic-control {
                        rule-set rule-set-name;
                    }
                    gprs-gtp-profile profile-name;
                    gprs-sctp-profile profile-name;
                    idp;
                    redirect-wx | reverse-redirect-wx;
                    ssl-proxy {
                        profile-name profile-name;
                    }
                    uac-policy {
                        captive-portal captive-portal;
                    }
                    utm-policy policy-name;
                }
                destination-address {
                    drop-translated;
                    drop-untranslated;
                }
                firewall-authentication {
                    pass-through {
                        access-profile profile-name;
                        client-match user-or-group-name;
                        ssl-termination-profile profile-name;
                        web-redirect;
                        web-redirect-to-https;
                    }
                    user-firewall {
                        access-profile profile-name;
                        domain domain-name
                        ssl-termination-profile profile-name;
                    }
                    web-authentication {
                        client-match user-or-group-name;
                    }
                }
```

```
                services-offload;
                tcp-options {
                    initial-tcp-mss mss-value;
                    reverse-tcp-mss mss-value;
                    sequence-check-required;
                    sequence-check-required;
                    syn-check-required;
                }
                tunnel {
                    ipsec-group-vpn group-vpn;
                    ipsec-vpn vpn-name;
                    pair-policy pair-policy;
                }
            }
            deny | reject;
            deny | reject [profile name];
        }
    }
}
```

Hierarchy Level

```
[edit security policies]
```

Description

Specify a source zone and destination zone to be associated with the security policy.

Options

- `from-zone` *`zone-name`*—Name of the source zone.
- `to-zone` *`zone-name`*—Name of the destination zone.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for the `services-offload` option added in Junos OS Release 11.4. Support for the `source-identity` option added in Junos OS Release 12.1. Support for the `description` option added in Junos OS Release 12.1. Support for the `ssl-termination-profile` and `web-redirect-to-https` options added in Junos OS Release 12.1X44-D10. Support for the `user-firewall` option added in Junos OS Release 12.1X45-D10. Support for the `initial-tcp-mss` and `reverse-tcp-mss` options added in Junos OS Release 12.3X48-D20. Support for the `dynamic-application` and `deny` options added in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

Unified Policies Configuration Overview

from-zone (Security Policies Global)

IN THIS SECTION

- Syntax | 7612
- Hierarchy Level | 7613
- Description | 7613
- Required Privilege Level | 7613
- Release Information | 7613

Syntax

```
from-zone {
    [zone-name];
    any;
}
```

Hierarchy Level

```
[edit security policies global policy policy-name match]
```

Description

Identify a single source zone or multiple source zones to be used as a match criteria for a policy. You must configure specific zones or default to any zone, but you cannot have both in a global policy.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Security Policies Overview

fru-poweron-sequence

IN THIS SECTION

- Syntax | 7614
- Hierarchy Level | 7614
- Description | 7614
- Options | 7614
- Required Privilege Level | 7615
- Release Information | 7615

Syntax

```
fru-poweron-sequence fru-poweron-sequence;
```

Hierarchy Level

```
[edit chassis]
```

Description

(MX Series 5G Universal Routing Platforms, SRX5xxx Series Platforms) Configure the power-on sequence for the DPCs in the chassis for routers with the enhanced AC Power Entry Module (PEM).

(T640 routers, T1600 routers, T4000 routers, MX2020 routers, and PTX Series packet transport routers) Configure the power-on sequence for Flexible PIC Concentrators (FPCs) installed in the chassis.

Options

(MX Series 5G Universal Routing Platforms only) *fru-poweron-sequence*—Power-on sequence for the DPCs in the chassis. The numbers indicate the slot number of the DPCs.

> **NOTE**: If the power-on sequence is not configured by including the `fru-poweron-sequence` statement, Junos OS uses the `/var/log/fpc_poweron_seq.log` file to determine the power-on sequence for the last power-on operation for the DPCs and the same sequence is used. If the `/var/log/boot_seq.log` file, is not available, Junos OS uses the ascending order of the slot numbers of the DPCs as the sequence to power on the DPCs.

(T640 routers, T1600 routers, T4000 routers, MX2020 routers, and PTX Series packet transport routers) *fru-poweron-sequence*—Power-on sequence for the FPCs in the chassis. The numbers indicate the slot number of the FPCs.

> **NOTE**:
>
> - If the configured sequence contains invalid numbers, Junos OS considers only the valid numbers in the sequence. The invalid numbers are silently discarded.

- If the power-on sequence is not configured by including the `fru-poweron-sequence` statement, Junos OS uses the ascending order of the slot numbers of the FPCs as the sequence to power on the FPCs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Redistributing the Available Power by Configuring Power-On Sequence

fte (Port)

IN THIS SECTION

- Syntax | 7616
- Hierarchy Level | 7616
- Description | 7616
- Options | 7616
- Required Privilege Level | 7616
- Release Information | 7616

Syntax

```
fte {
    (port port-number | port-range port-range-low port-range-high);
}
```

Hierarchy Level

```
[edit chassis fpc slot pic pic-number]
```

Description

Configure a specific port or a range of ports to operate as 40-Gigabit Ethernet data plane uplink (*fte*) type ports.

CAUTION: The Packet Forwarding Engine on the switch is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the switch.

Options

port-number Port number on which you want to configure the port type. Valid values are 2 through 7.

port-range-low Lowest-numbered port in the range of ports. The lowest possible value is 2.

port-range-high Highest-numbered port in the range of ports. The highest possible value is 7.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X52-D10.

RELATED DOCUMENTATION

Configuring the QSFP+ Port Type on QFX3500 Standalone Switches

Configuring the Port Type on QFX3600 Standalone Switches

Configuring the QSFP+ Port Type on QFX5100 Devices

ftp

IN THIS SECTION

- Syntax | 7617
- Hierarchy Level | 7617
- Description | 7618
- Options | 7618
- Required Privilege Level | 7619
- Release Information | 7619

Syntax

```
ftp {
    authentication-order [authentication-methods];
    connection-limit limit;
    rate-limit limit;
    enable-anonymous-login;
}
```

Hierarchy Level

```
[edit system services]
```

Description

Allow FTP requests from remote systems to the local device.

Options

authentication-order [*authentication-methods*]

Configure the order in which the software tries different user authentication methods when attempting to authenticate a user. For each login attempt, the software tries the authentication methods in order, starting with the first one, until the password matches.

- **Values:** Specify one or more of the following authentication methods listed in the order in which they must be tried:
  - `ldaps`—Use LDAP authentication services.
  - `password`—Use the password configured for the user with the `authentication` statement at the `[edit system login user]` hierarchy level.
  - `radius`—Use RADIUS authentication services.
  - `tacplus`—Use TACACS+ authentication services.
- **Default:** If you do not include the `authentication-order` statement, users are verified based on their configured passwords.

connection-limit *limit*

Configure the maximum number of connections sessions for the ftp service per protocol (either IPv6 or IPv4).

> **NOTE**: The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

- **Range:** 1 through 250 connections
- **Default:** 75 connections

rate-limit *limit*

Configure the maximum number of connections attempts per minute, per protocol (either IPv6 or IPv4) on an access service. For example, a rate limit of 10 allows 10 IPv6 ftp session connection attempts per minute and 10 IPv4 ftp session connection attempts per minute.

- **Range:** 1 through 250 connections

- **Default:** 150 connections

enable-anonymous-login

Enable anonymous FTP access. Guests can login and send FTP requests without a password by using the `ftp` username.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Option `ldaps` introduced in Junos OS Release 20.2R1.

Option `enable-anonymous-login` introduced in Junos OS Release 23.2R2.

RELATED DOCUMENTATION

Configuring FTP Service for Remote Access to the Router or Switch

Junos OS User Authentication Methods

ftp (Access)

IN THIS SECTION

- Syntax | 7620
- Hierarchy Level | 7620
- Description | 7620
- Options | 7620
- Required Privilege Level | 7620
- Release Information | 7620

Syntax

```
ftp {
    banner {
        fail string;
        login string;
        success string;
    }
}
```

Hierarchy Level

```
[edit access firewall-authentication pass-through]
```

Description

Configure banners for the FTP login prompt, successful authentication, and failed authentication.

Options

The remaining statements are explained separately.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

ftp (Content Security Policy Anti-Virus)

IN THIS SECTION

- Syntax | 7621
- Hierarchy Level | 7621
- Description | 7621
- Options | 7621
- Required Privilege Level | 7622
- Release Information | 7622

Syntax

```
ftp {
    download-profile profile-name;
    upload-profile profile-name;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name anti-virus]
```

Description

Configure a Content Security policy for the antivirus FTP protocol and attach this policy to a security profile to implement it.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

ftp (Content Security Policy Content Filtering)

IN THIS SECTION

- Syntax | 7622
- Hierarchy Level | 7623
- Description | 7623
- Options | 7623
- Required Privilege Level | 7623
- Release Information | 7623

Syntax

```
ftp {
    download-profile profile-name;
```

```
    upload-profile profile-name;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name content-filtering]
```

Description

Configure a Content Security policy for the content-filtering FTP protocol and attach this policy to a security profile to implement it.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

ftp (Flow Collector Files)

IN THIS SECTION

- Syntax | 7624
- Hierarchy Level | 7624
- Description | 7624
- Options | 7624
- Required Privilege Level | 7625
- Release Information | 7625

Syntax

```
ftp:url;
```

Hierarchy Level

```
[edit services flow-collector destination]
```

Description

Specify the primary and secondary destination FTP server addresses.

Options

url—FTP server address. The URL can include the following macros, typed in braces:

- `{%D}`—Date
- `{%T)`—Time when the file is created
- `{%I}`–Description string for the logical interface configured using the `collector` *`interface-name`* statement at the `[edit services flow-collector interface-map]` hierarchy
- `{%N}`—Unique, sequential number for each new file created

- {am_pm}—AM or PM
- {date}—Current date using the {year} {month} {day} macros
- {day}—From 01 through 31
- {day_abbr}—Sun through Sat
- {day_full}—Sunday through Saturday
- {generation number}—Unique, sequential number for each new file created
- {hour_12}—From 01 through 12
- {hour_24}—From 00 through 23
- {ifalias}—Description string for the logical interface configured using the collector statement at the [edit services flow-collector interface-map] hierarchy
- {minute}—From 00 through 59
- {month}—From 01 through 12
- {month_abbr}—Jan through Dec
- {month_full}—January through December
- {num_zone}—From -2359 to +2359; this macro is not supported
- {second}—From 00 through 60
- {time}—Time the file is created, using the {hour_24} {minute} {second} macros
- {time_zone}—Time zone code name of the locale; for example, gmt (this macro is not supported).
- {year}—In the format YYYY; for example, 1970
- {year_abbr}—From 00 through 99

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Destination FTP Servers for Flow Records

ftp (Security ALG)

IN THIS SECTION

- Syntax | 7626
- Hierarchy Level | 7627
- Description | 7627
- Options | 7627
- Required Privilege Level | 7627
- Release Information | 7628

Syntax

```
ftp {
    allow-mismatch-ip-address;
    disable;
    ftps-extension;
    line-break-extension;
    traceoptions {
        flag {
            all {
                extensive
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Configure the ftp function to Specify the File Transfer Protocol (FTP) Application Layer Gateway (ALG) on the device.

Options

- `allow-mismatch-ip-address`—Control commands in FTP, such as PORT and EPRT, contain IP addresses internally. If the IP address does not match the source IP at Packet Layer 3, which is the client's IP, the packets are discarded. However, when the allow-mismatch-ip-address option is enabled, these packets are accepted.
- `disable`—Disable the FTP ALG. By default, the FTP ALG is enabled. You can enable or disable FTP ALG for both IPV4 and IPV6 mode.
- `ftps-extension`—Enable secure FTP and FTP SSL protocols.
- `line-break-extension`—Enable line-break-extension. This option will enable the FTP ALG to recognize the LF as line break in addition to the standard CR+LF (carriage return, followed by line feed).
- `traceoptions`—Configure FTP ALG tracing options. To specify more than one trace operation, include multiple flag statements.
  - `flag`—Trace operation to perform.
    - `all`—Trace all events.
    - `extensive`—(Optional) Display extensive amount of data.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 11.4.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

FTP ALG Overview

ftp (Services CoS)

IN THIS SECTION

- Syntax | 7628
- Hierarchy Level | 7629
- Description | 7629
- Default | 7629
- Required Privilege Level | 7629
- Release Information | 7629

Syntax

```
ftp {
    data {
        dscp (alias | bits);
        forwarding-class class-name;
    }
}
```

Hierarchy Level

```
[edit services cos application-profile profile-name ftp]
```

Description

Set the appropriate `dscp` and `forwarding-class` value for FTP.

Default

By default, the system does not alter the DSCP or forwarding class for FTP traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

sip (Application Profile) | 27973

ftp (Services CoS Next Gen Services)

IN THIS SECTION

- Syntax | 7630
- Hierarchy Level | 7630
- Description | 7630

- Options | 7630
- Required Privilege Level | 7631
- Release Information | 7631

Syntax

```
ftp {
    data {
        dscp (alias | bits);
        forwarding-class class-name;
    }
}
```

Hierarchy Level

```
[edit services cos application-profile profile-name]
```

Description

Configure CoS actions for FTP traffic in an application profile. The application profile can then be used in CoS rule actions.

Options

dscp (*alias* | *bits*) Either a code point alias or a DSCP bit value to apply to the FTP packets.

forwarding-class *class-name* Forwarding class name to apply to the FTP packets. The choices are:

- assured-forwarding
- best-effort
- expedited-forwarding
- network-control

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Class of Service for Services PICs (Next Gen Services)

ftp (Transfer Log Files)

IN THIS SECTION

- Syntax | 7631
- Hierarchy Level | 7631
- Description | 7632
- Options | 7632
- Required Privilege Level | 7632
- Release Information | 7632

Syntax

```
ftp:url;
```

Hierarchy Level

```
[edit services flow-collector transfer-log-archive archive-sites]
```

Description

Specify the primary and secondary destination FTP server addresses.

Options

url—FTP server address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Transfer Logs

full

IN THIS SECTION

- See | 7632

See

"brief" on page 3039

function (Diameter Network Element for Subscriber Management)

IN THIS SECTION

- Syntax | 7633
- Hierarchy Level | 7633
- Description | 7633
- Options | 7633
- Required Privilege Level | 7634
- Release Information | 7634

Syntax

```
function function-name;
```

Hierarchy Level

```
[edit access diameter network-element element-name]
```

Description

Specify the function associated with a Diameter network element.

Options

function-name—Function associated with the network element.

Functions currently supported:

- Policy charging and control (`pcc-gx`).
- Diameter credit-control application are the functions currently supported.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

function (Diameter Network Element)

IN THIS SECTION

- Syntax | 7634
- Hierarchy Level | 7635
- Description | 7635
- Default | 7635
- Options | 7635
- Required Privilege Level | 7635
- Release Information | 7635

Syntax

```
function [function-name];
```

Hierarchy Level

```
[edit diameter network-element element-name]
```

Description

Specify one or more applications (function) associated with a Diameter network element.

Default

By default, all functions are associated with (supported by) the network element.

Options

function-name—Application (function) associated with the route. You can list multiple functions with the element.

- `gx-plus`—Associate the Gx-Plus application with the network element.
- `jsrc`—Associate the JSRC application with the network element.
- `nasreq`—Associate the NASREQ application with the network element.
- `packet-triggered-subscribers`—Associate the packet-triggered subscribers application with the network element.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

Support for PTSP introduced in Junos OS Release 10.2.

Support for Gx-Plus introduced in Junos OS Release 11.2.

Support for PTSP discontinued in Junos OS Release 13.1.

Support for NASREQ added in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Network Elements

function (Diameter Route)

IN THIS SECTION

- Syntax | 7636
- Hierarchy Level | 7636
- Description | 7636
- Default | 7637
- Options | 7637
- Required Privilege Level | 7637
- Release Information | 7637

Syntax

```
function function-name <partition partition-name>;
```

Hierarchy Level

```
[edit diameter network-element element-name forwarding route dne-route-name]
```

Description

Specify the application (function) associated with a destination and metric. Together, these three elements define a route reachable through a Diameter network element.

Default

All functions are associated with the route.

Options

function-name—Application (function) associated with the route. Gx-Plus, JSRC, NASREQ, and packet-triggered-subscribers are the applications currently supported.

partition *partition-name*—(Optional) Partition associated with the application (function).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support for PTSP introduced in Junos OS Release 10.2.

Support for Gx-Plus introduced in Junos OS Release 11.2.

Support for PTSP discontinued in Junos OS Release 13.1.

Support for NASREQ added in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Network Elements

functional-zone

IN THIS SECTION

- Syntax | 7638
- Hierarchy Level | 7639
- Description | 7639
- Options | 7639
- Required Privilege Level | 7639
- Release Information | 7639

Syntax

```
functional-zone {
    management {
        description text;
        host-inbound-traffic {
            protocols protocol-name {
                                except;
            }
            system-services service-name {
                                except;
            }
        }
        interfaces interface-name  {
            host-inbound-traffic {
                protocols protocol-name {
                                    except;
                }
                system-services service-name {
                                    except;
                }
            }
        }
        screen screen-name;
```

```
    }
}
```

Hierarchy Level

```
[edit security zones]
```

Description

Configure a functional zone.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. The `description` option added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Functional Zones

G

IN THIS SECTION

- g-duplicates-dropped-periodicity | 7645
- g-max-duplicates | 7646
- garp-na-enable | 7648
- gate-status | 7649
- gate-timeout | 7651
- gatekeeper | 7653
- gateway-name (LNS Local Gateway) | 7654
- gateway-name (Tunnel Profile Remote Gateway) | 7655
- gateway-name (Tunnel Profile Source Gateway) | 7657
- gateway (Security Group VPN Member IKE) | 7658
- gateway (Security Group VPN Server IKE) | 7660
- gateway (Security IKE) | 7663
- general-ikeid | 7669
- general-param (Flat-File Accounting Options) | 7670
- generate | 7672
- generate-event | 7674
- generate-snmp-traps | 7676
- geo-ip (global-filter) | 7678
- geo-ip (group-filter) | 7679
- get-route-range | 7681
- ggsn-address (OCS Partition) | 7682
- ggsn-mcc-mnc (OCS Partition) | 7684
- gigether-options | 7685
- gigether-options (Chassis Cluster) | 7694
- global-advertisements-threshold | 7696
- global-arp-prefix-limit (Host Fast Reroute) | 7698
- global-config (Services) | 7700
- global (DDoS) | 7703

- global-dns-stats-log-timer | **7705**
- global (Gx-Plus) | **7706**
- global-limit | **7708**
- global-mac-ip-limit | **7709**
- global-mac-ip-table-aging-time | **7711**
- global-mac-limit | **7712**
- global-mac-limit (Protocols) | **7714**
- global-mac-move | **7716**
- global-mac-statistics | **7717**
- global-mac-table-aging-time | **7719**
- global-mode (Protocols) | **7721**
- global-no-control-mac-aging | **7722**
- global-no-mac-learning | **7723**
- global (OCS) | **7725**
- global-options (Juniper Secure Connect) | **7726**
- global-parameters (Easy EVPN LAG Configuration) | **7728**
- global (PCRF) | **7731**
- global (Security IDP) | **7733**
- global (Security Policies) | **7734**
- global-supplementary-blackout-timer (Host Fast Reroute) | **7737**
- global-threshold | **7740**
- global-wait-to-restore | **7741**
- global-weight | **7743**
- global-load-balancing | **7744**
- gnf | **7746**
- gpid | **7748**
- gprs | **7750**
- gprs-gtp-profile | **7755**
- gprs-sctp-profile | **7757**
- graceful-deletion-timeout | **7758**
- graceful-restart (Enabling Globally) | **7760**
- graceful-restart (Fabric Control) | **7762**

- graceful-restart (Long-Lived for BGP Helper) | **7763**
- graceful-restart (Long-Lived for BGP Restarter) | **7765**
- graceful-restart (Multicast Snooping) | **7768**
- graceful-restart (Protocols BGP) | **7769**
- graceful-restart (Protocols IS-IS) | **7771**
- graceful-restart (Protocols LDP) | **7773**
- graceful-restart (Protocols OSPF) | **7775**
- graceful-restart (Protocols PIM) | **7777**
- graceful-restart (Protocols RIP) | **7779**
- graceful-restart (Protocols RIPng) | **7780**
- graceful-restart (Protocols RSVP) | **7782**
- graceful-shutdown (Protocols BGP) | **7784**
- graceful-switchover (Chassis) | **7786**
- gratuitous-arp-count | **7787**
- gratuitous-arp-reply | **7789**
- gre | **7790**
- gre (Ids-option) | **7791**
- gre-in | **7793**
- gre-out | **7794**
- gre-performance-acceleration (Security Flow) | **7796**
- gre (Routing Options) | **7797**
- greater-than | **7798**
- gres-route-flush-delay (Subscriber Management) | **7800**
- group | **7801**
- group-address (Routing Instances Tunnel Group) | **7802**
- group-address (Routing Instances VPN) | **7804**
- group (Bridge Domains) | **7806**
- group (Configuring Group Name) | **7807**
- group-count (Dynamic MLD Interface) | **7809**
- group-count (Protocols IGMP) | **7810**
- group-count (Protocols MLD) | **7812**
- group (Defining Access Privileges for an SNMPv3 Group) | **7813**

- group (DHCP Local Server) | **7814**
- group (DHCP Relay Agent) | **7820**
- group (DHCP Security) | **7826**
- group (DHCP Security for MX Series) | **7827**
- group (DHCP Spoofing Prevention) | **7829**
- group (Distributed IGMP) | **7830**
- group (Dynamic IGMP Interface) | **7832**
- group (Dynamic MLD Interface) | **7834**
- group (IGMP Snooping) | **7835**
- group-increment (Dynamic MLD Interface) | **7837**
- group-increment (Protocols IGMP) | **7838**
- group-increment (Protocols MLD) | **7840**
- group-limit (Dynamic IGMP Interface) | **7841**
- group-limit (Dynamic MLD Interface) | **7843**
- group-limit (IGMP) | **7844**
- group-limit (IGMP and MLD Snooping) | **7846**
- group-limit (Protocols MLD) | **7847**
- group-members | **7849**
- group (Origin Validation for BGP) | **7850**
- group-policy (Dynamic IGMP Interface) | **7852**
- group-policy (Dynamic MLD Interface) | **7854**
- group-policy (Protocols IGMP) | **7855**
- group-policy (Protocols IGMP AMT Interface) | **7856**
- group-policy (Protocols MLD) | **7858**
- group-profile (Access) | **7859**
- group-profile (Group Profile) | **7861**
- group (Protocols BGP) | **7863**
- group (Protocols IGMP) | **7869**
- group (Protocols MLD) | **7870**
- group (Protocols MLD Snooping) | **7872**
- group (Protocols MSDP) | **7873**
- group (Protocols PIM) | **7876**

- group (Protocols RIP) | **7877**
- group (Protocols RIPng) | **7880**
- group-range (Data MDTs) | **7882**
- group-range (MBGP MVPN Tunnel) | **7883**
- group-ranges | **7885**
- group (Redundant Trunk Groups) | **7887**
- group (Routing Instances) | **7889**
- group-rp-mapping | **7891**
- group (RPF Selection) | **7893**
- group (Security Group VPN) | **7894**
- group (Static Subscribers) | **7897**
- group (System Services DHCP) | **7899**
- group-threshold (Protocols IGMP Interface) | **7903**
- group-threshold (Protocols MLD Interface) | **7905**
- group (Traffic Load Balancer) | **7906**
- group-type (Unknown Unicast Forwarding) | **7908**
- group (Uplink Failure Detection) | **7910**
- group-vpn | **7911**
- grouped-ie-profile | **7915**
- groups | **7916**
- grpc | **7920**
- grpc-replay | **7922**
- grpc-tunnel | **7923**
- gsm-options | **7926**
- gsmp-syn-timeout (ANCP) | **7928**
- gsmp-syn-wait (ANCP) | **7929**
- gtp | **7931**
- gtp-header-offset | **7936**
- gtp-tunnel-endpoint-identifier | **7938**
- guaranteed-rate (Traffic Control Profiles) | **7940**
- guaranteed-rate (Dynamic Traffic Shaping) | **7942**
- guard-band | **7944**

- guard-band (PoE) | 7946
- guard-interval | 7947
- guest-network-functions | 7949
- gw-interface | 7951
- gx-plus (Gx-Plus) | 7952
- gx-profile | 7954

g-duplicates-dropped-periodicity

IN THIS SECTION

- Syntax | 7645
- Hierarchy Level | 7645
- Description | 7646
- Default | 7646
- Options | 7646
- Required Privilege Level | 7646
- Release Information | 7646

Syntax

```
g-duplicates-dropped-periodicity seconds;
```

Hierarchy Level

```
[edit services dynamic-flow-capture]
```

Description

Specify the frequency for sending notifications to affected control sources when transmission of duplicate sets of data is restricted because the `g-max-duplicates` threshold has been reached. This setting is applied globally; the `duplicates-dropped-periodicity` setting applied at the `capture-group` level overrides the global setting.

Default

The default period for sending notifications is 30 seconds.

Options

seconds—Period for sending DuplicatesDropped notifications.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

duplicates-dropped-periodicity | 6072

Limiting the Number of Duplicates of a Packet

g-max-duplicates

IN THIS SECTION

- Syntax | 7647
- Hierarchy Level | 7647

- Description | 7647
- Default | 7647
- Options | 7647
- Required Privilege Level | 7647
- Release Information | 7648

Syntax

```
g-max-duplicates number;
```

Hierarchy Level

```
[edit services dynamic-flow-capture]
```

Description

Specify the maximum number of content destinations to which DFC PICs can send data from a single input set of packets. Limiting the number of duplicates reduces the load on the PIC. This setting is applied globally; the `max-duplicates` setting applied at the `capture-group` level overrides the global setting.

Default

If no value is configured, a default setting of 3 is used.

Options

number—Maximum number of content destinations.

- **Range:** 1 through 64

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

max-duplicates | 10436

Limiting the Number of Duplicates of a Packet

garp-na-enable

IN THIS SECTION

- Syntax | 7648
- Hierarchy Level | 7648
- Description | 7649
- Required Privilege Level | 7649
- Release Information | 7649

Syntax

```
garp-na-enable;
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

Enable sending unsolicited GARP or NA messages.

When learning a new MAC-IP binding on a local PE device, or learning a MAC-IP binding on a remote PE device via EVPN, an unsolicited GARP or NA message is sent to all connected access CE devices. This message enables the connected CE devices to update their ARP or ND caches.

Required Privilege Level

[none specified]

Release Information

Statement introduced in Junos OS Release 23.4R1.

gate-status

IN THIS SECTION

- Syntax | 7649
- Hierarchy Level | 7650
- Description | 7650
- Default | 7650
- Options | 7650
- Required Privilege Level | 7650
- Release Information | 7651

Syntax

```
gate-status (uplink | downlink | uplink-downlink | disable-both);
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],
[edit services pcef pcc-action-profiles profile-name]
```

Description

Configure the gate status in a PCC action profile to enable or disable the forwarding of service flow packets. The gate status determines whether the uplink and downlink gates are opened or closed.

If you are using Junos OS Subscriber Aware, configure the gate status at the `[edit unified-edge pcef pcc-action-profiles profile-name]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the gate status at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level.

Default

By default, if this statement is not configured, forwarding of service data flow packets is enabled in both the uplink and downlink directions.

Options

disable-both Disable forwarding of service data flow packets in the uplink and downlink directions.

downlink Enable forwarding of service data flow packets in the downlink direction.

uplink-downlink Enable forwarding of service data flow packets in the uplink and downlink directions.

uplink Enable forwarding of service data flow packets in the uplink direction.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

- Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware
- Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment
- *Configuring Policy and Charging Control Action Profiles for Subscriber Management*
- *Understanding Application-Aware Policy Control for Subscriber Management*

gate-timeout

IN THIS SECTION

- Syntax | **7651**
- Hierarchy Level | **7652**
- Description | **7652**
- Options | **7652**
- Required Privilege Level | **7652**
- Release Information | **7652**

Syntax

```
gate-timeout seconds;
```

Hierarchy Level

```
[edit applications application ike-esp-nat]
```

Description

For an IKE ALG application, configure the length of time that can pass after IKE establishes the security association between the IPsec client and server and before the ESP traffic starts in both directions. If the ESP traffic has not started before this timeout value, the ESP gates are deleted and the ESP traffic is blocked.

The IKE ALG enables the passing of IKEv1 and IPsec packets through NAPT-44 and NAT64 rules between IPsec peers that are not NAT-T compliant.

Options

seconds

Number of seconds.

- **Default:** 120 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

ALG Descriptions

Configuring Application Sets

Configuring Application Properties

gatekeeper

IN THIS SECTION

- Syntax | 7653
- Hierarchy Level | 7653
- Description | 7653
- Options | 7653
- Required Privilege Level | 7654
- Release Information | 7654

Syntax

```
gatekeeper {
    threshold rate;
}
```

Hierarchy Level

```
[edit logical-systems name security alg h323 application-screen message-flood],
[edit logical-systems name tenants name security alg h323 application-screen message-flood],
[edit security alg h323 application-screen message-flood],
[edit services alg h323 application-screen message-flood],
[edit tenants name security alg h323 application-screen message-flood]
```

Description

Configure the gatekeeper parameter to limit the rate of requests processed from remote access server (RAS) to the gatekeeper. Messages exceeding the threshold are dropped.

Options

`threshold` *rate*—Threshold measured in messages per second.

- **Range:** 1 through 50,000
- **Default:** 1000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

h323 | 7961

gateway-name (LNS Local Gateway)

IN THIS SECTION

- Syntax | 7654
- Hierarchy Level | 7655
- Description | 7655
- Options | 7655
- Required Privilege Level | 7655
- Release Information | 7655

Syntax

```
gateway-name gateway-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group group-name local-gateway]
```

Description

Specify the gateway name for the LNS, which the LNS returns to the LAC in response to the LAC's SCCRQ message. This name must match the remote gateway name configured on the LAC, or the tunnel cannot be established.

Options

gateway-name—Name of the LNS.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

gateway-name (Tunnel Profile Remote Gateway)

IN THIS SECTION

- Syntax | 7656
- Hierarchy Level | 7656
- Description | 7656

- Options | 7656
- Required Privilege Level | 7656
- Release Information | 7656

Syntax

```
gateway-name server-name;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id remote-gateway]
```

Description

Specify the hostname expected by the remote gateway—the LNS—from the source gateway—the LAC—when you set up a tunnel.

Options

server-name—Name of the LNS.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

gateway-name (Tunnel Profile Source Gateway)

IN THIS SECTION

- Syntax | 7657
- Hierarchy Level | 7657
- Description | 7657
- Options | 7657
- Required Privilege Level | 7657
- Release Information | 7658

Syntax

```
gateway-name client-name;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id source-gateway]
```

Description

Specify the hostname provided by the source gateway—the LAC—to the remote gateway—the LNS—when you set up a tunnel.

Options

client-name—Name of the LAC.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

gateway (Security Group VPN Member IKE)

IN THIS SECTION

- Syntax | 7658
- Hierarchy Level | 7659
- Description | 7659
- Options | 7659
- Required Privilege Level | 7660
- Release Information | 7660

Syntax

```
gateway gateway-name {
    ike-policy policy-name;
    local address ip-address;
    local-identity {
        (hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname e-mail-
address);
    }
    remote-identity {
        (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
    }
    routing-instance routing-instance;
```

```
    server-address ip-address;
}
```

Hierarchy Level

```
[edit security group-vpn member ike]
```

Description

Configure IKE gateway for group VPN member. An IKE gateway initiates and terminates network connections between a firewall and a security device.

Options

gateway ***gateway-name***	Name of the gateway.
ike-policy ***policy-name***	Name of the IKE policy.
local address ***ip-address***	Configure the IPv4 address the member uses when accessing the group server.
local-identity ***local-identity***	Specify the local IKE identity to send in the exchange with the destination peer to establish communication.
remote-identity ***remote-identity***	Specify the name of a routing instance. If this is not specified, the default inet.0 routing instance is used.
routing-instance ***routing-instance***	Specify the name of a routing instance. If this is not specified, the default inet.0 routing instance is used.
server-address ***ip-address***	Specify the group server IPv4 address that this member registers through a `groupkey-pull` exchange. Up to four server IP addresses can be configured. The group member attempts to register with the first configured server. If registration with a configured server is not successful, the group member tries to register with the next configured server. We recommend that group members only register with sub-servers in a server cluster and not the root-server.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2. Support for the `routing-instance` option added in Junos OS Release 15.1X49-D30 for vSRX Virtual Firewall.

RELATED DOCUMENTATION

Group VPNv2 Overview

gateway (Security Group VPN Server IKE)

IN THIS SECTION

- Syntax | 7660
- Hierarchy Level | 7661
- Description | 7661
- Options | 7661
- Required Privilege Level | 7662
- Release Information | 7662

Syntax

```
gateway gateway-name {
    address ip-address;
    dead-peer-detection {
        always-send;
        interval seconds;
        threshold number;
```

```
    }
    dynamic {
        (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
    }
    ike-policy policy-name;
    local-address ip-address;
    local-identity {
        (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
    }
    remote-identity {
        (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
    }
    routing-instance routing-instance;
}
```

Hierarchy Level

```
[edit security group-vpn server ike]
```

Description

Configure IKE gateway for group VPN server.

Options

`gateway` *gateway-name* —Name of the gateway.

`address` *ip-address* —Specify the IP address of the peer.

`dead-peer-detection` —Enable DPD between group server cluster servers.

`dynamic`—Specify the identifier for the remote gateway with a dynamic IPv4 address. Use this statement to set up a VPN with a gateway that has an unspecified IPv4 address.

- `hostname` *domain-name* —Specify a fully qualified domain name.
- `inet` *ip-address* —Specify an IPv4 address to identify the dynamic peer.
- `user-at-hostname` *e-mail-address* —Specify an e-mail address.

Configuring `mode main` for group VPN servers or members is not supported when the remote gateway has a dynamic address and the authentication method is `pre-shared-keys`.`ike-policy` *`policy-name`* —Specify the name of the IKE policy.

`local-address` *`ip-address`* —Configure the source IP address the group VPN server uses when communicating with a group member or a root-server. This statement is normally used when there are multiple IP addresses bound to an interface.

`local-identity`—Specify the local IKE identity to send in the exchange with the destination peer to establish communication. If you do not configure a local-identity, the device uses the IPv4 corresponding to the local endpoint by default.

- `hostname` *`hostname`*—Specify identity as a fully qualified domain name (FQDN).
- `inet` *`ip-address`*—Specify identity as an IPv4 address.
- `user-at-hostname` *`e-mail-address`*—Specify identity as an e-mail address.

`remote-identity`—Specify the remote IKE identity of the destination peer. If you do not configure a remote identity, the device uses, by default, the IPv4 address that corresponds to the destination peer.

- `hostname` *`hostname`*—Specify identity as a fully qualified domain name (FQDN).
- `inet` *`ip-address`*—Specify identity as an IPv4 address.
- `user-at-hostname` *`e-mail-address`*—Specify identity as an e-mail address.

`routing-instance` *`routing-instance`*—Configure the routing instance that the group VPN server uses when communicating with a group server. This statement is used when the IKE gateway is not configured in the default routing instance.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Support for the Group VPN server added in Junos OS Release 15.1X49-D30 for vSRX Virtual Firewall.

RELATED DOCUMENTATION

Group VPNv2 Overview

gateway (Security IKE)

IN THIS SECTION

- Syntax | 7663
- Hierarchy Level | 7665
- Description | 7665
- Options | 7665
- Required Privilege Level | 7668
- Release Information | 7668

Syntax

```
gateway gateway-name {
    aaa {
        access-profile access-profile {
            config-payload-password config-payload-password;
        }
        client {
            password;
            username;
        }
        use-routing-instance-address;
        no-push-to-identity-management;
    }
    address [ip-addresses-or-hostnames];
    advpn {
        suggester {
            disable;
        }
        partner {
            connection-limit number;
```

```
            idle-threshold   packets/sec;
            idle-time        seconds;
            disable;
        }
    }
    dead-peer-detection {
        (always-send | optimized | probe-idle-tunnel);
        interval seconds;
        threshold number;
    }
    dynamic {
        connections-limit number;
        distinguished-name {
            container container-string;
            wildcard wildcard-string
        }
        general-ikeid;
        hostname domain-name;
        ike-user-type (group-ike-id | shared-ike-id);
        inet ip-address;
        inet6 ipv6-address;
        reject-duplicate-connection;
        user-at-hostname e-mail-address;
    }
    external-interface external-interface-name;
    fragmentation {
        disable;
        size bytes;
    }
    general-ikeid;
    ike-policy policy-name;
    local-address (ipv4-address | ipv6-address);
    local-identity {
        (distinguished-name | hostname hostname | inet ip-address | inet6 ipv6-address | key-id
| user-at-hostname e-mail-address);
    }
    nat-keepalive seconds;
    no-nat-traversal;
    node-local;
    remote-identity {
        (distinguished-name <container container-string> <wildcard wildcard-string> | hostname
hostname | inet ip-address | inet6 ipv6-address | key-id | user-at-hostname e-mail-address);
    }
```

```
    tcp-encap-profile profile-name;
    version (v1-only | v2-only);
}
```

Hierarchy Level

```
[edit security ike]
```

Description

Configure an IKE gateway.

Options

gateway-name Name of the gateway.

address Specify either the IPv4 or IPv6 addresses or the hostnames of the primary Internet Key Exchange (IKE) gateway (peer) and up to four backup gateways.

Consider the following points before configuring multiple peer addresses with IPsec VPN running iked process:

- The option supports one primary address and upto a maximum of 4 backup addresses, all belonging to the same inet family.
- You must configure DPD before configuring multiple peer addresses.
- Only site-to-site VPN configuration supports multiple peer addresses.
- HA configuration doesn't support multiple peer addresses.

> **CAUTION**: Exercise caution when modifying the peer address:
>
> - If you delete an active peer address, the peer address list is updated and the next available peer address is considered active. This is a catastrophic operation and the VPN tunnel is torn down and a new tunnel is negotiated as per the peer address list.

- If you delete an inactive peer address, the peer address list is updated but does not cause any catastrophic event as there is no impact on the tunnel.
- If you add a new peer address, it will be appended to the end of the list, and there is no impact on the tunnel.

- Values:
  - address—IPv4 or IPv6 addresses or hostnames of an IKE gateway.

    This is a list of values. Priority of the peer addresses is based on the order of addresses in the list.
- Range:
  - Supports upto 5 peer addresses which includes an active peer address and upto 4 backup peer addresses.

aaa Specify that extended authentication is performed in addition to IKE Phase 1 authentication for remote users trying to access a VPN tunnel.

advpn Enable Auto Discovery VPN (ADVPN) protocol on the specified gateway. ADVPN dynamically establishes VPN tunnels between spokes to avoid routing traffic through the Hub.

dead-peer-detection Enable the device to use dead peer detection (DPD).

dynamic Specify the identifier for the remote gateway with a dynamic IPv4 or IPv6 address. Use this statement to set up a VPN with a gateway that has an unspecified IPv4 or IPv6 address.

external-interface Name of the interface to be used to send traffic to the IPsec VPN. Specify the outgoing interface for IKE SAs. This interface is associated with a zone that acts as its carrier, providing firewall security for it.

fragmentation Disable IKEv2 packet fragmentation and, optionally, configure the maximum size of an IKEv2 message before the message is split into fragments that are individually encrypted and authenticated.

disable Disables IKEv2 fragmentation. IKEv2 fragmentation is enabled by default.

size *bytes* Maximum size, in bytes, of an IKEv2 message before it is split into fragments. The size applies to both IPv4 and IPv6 messages.

- **Range:** 500 to 1300 bytes
- **Default:** 576 bytes for IPv4 messages and 1280 bytes for IPv6 messages

general-ikeid Accept peer IKE-ID in general.

ike-policy Specify the IKE policy to be used for the gateway.

local-address Local IP address for IKE negotiations. Specify the local gateway address. Multiple addresses in the same address family can be configured on an external physical interface to a VPN peer. If this is the case, we recommend that `local-address` be configured. If there is only one address configured (IPv4 or IPv6) on an external physical interface, `local-address` configuration is not necessary.

The `local-address` value must be an IP address that is configured on an interface on the SRX Series Firewall. We recommend that `local-address` belong to the external interface of the IKE gateway. If local-address does not belong to the external interface of the IKE gateway, the interface must be in the same zone as the external interface of the IKE gateway and an intra-zone security policy must be configured to permit traffic. The `local-address` value and the remote IKE gateway address must be in the same address family, either IPv4 or IPv6.

local-identity Specify the local IKE identity to send in the exchange with the destination peer to establish communication.

nat-keepalive Specify the interval at which NAT keepalive packets (seconds) can be sent so that NAT translation continues. Default value changed from 5 seconds to 20 seconds in Junos OS Release 12.1X46-D10.

- **Default:** 20
- **Range:** 1 through 300

node-local Mark an IPsec VPN tunnel between Multinode High Availability nodes and a VPN peer device as a node-local tunnel. Node-local tunnels support dynamic routing protocols that facilitate the device to add the routes dynamically. These routes remain local to a node and are not bound to any services redundancy group (SRG). Use this option only for Multinode High Availability.

no-nat-traversal Disable IPSec NAT traversal. Disables UDP encapsulation of IPsec Encapsulating Security Payload (ESP) packets, otherwise known as Network Address Translation Traversal (NAT-T). NAT-T is enabled by default.

tcp-encap-profile Specify the TCP encapsulation profile to be used for TCP connections for remote access clients.

version Specify the IKE version to use to initiate the connection.

- Values:
  - v1-only—The connection must be initiated using IKE version 1. This is the default.
  - v2-only—The connection must be initiated using IKE version 2

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1. The `inet6` option added in Junos OS Release 11.1. Support for the `advpn` option added in Junos OS Release 12.3X48-D10.

Option `fragmentation` is introduced in Junos OS Release 15.1X49-D80.

Option `tcp-encap-profile` is introduced in Junos OS Release 15.1X49-D80.

`general-ikeid` option under `[edit security ike gateway gateway-name dynamic]` hierarchy is introduced in Junos OS Release 21.1R1.

Option `node-local` is introduced in Junos OS Release 23.2R1.

Support for multiple peer addresses in the `address` option for IPsec VPN running iked process is introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

IPsec Overview

Example: Configuring AutoVPN with Pre-Shared Key

tcp-encap | **28980**

general-ikeid

IN THIS SECTION

- Syntax | 7669
- Hierarchy Level | 7669
- Description | 7669
- Required Privilege Level | 7670
- Release Information | 7670

Syntax

```
general-ikeid;
```

Hierarchy Level

```
[edit security ike gateway gateway_name dynamic]
```

Description

During IKE Phase 1 negotiation, when negotiation request is received, there are two identity checks.

1. IKE-ID validation from ID payload.
2. Phase 1 authentication by pre-shared key or RSA/DSA certificate.

Configure `remote-identity` to lookup the certificate of the peer for certificate authentication. This `remote-identity` should match the corresponding field in the `SubjectAltname` extension of the peer certificate for successful detection of peer certificate and authentication.

The identity check with the same IKE-ID is repeated, that is, the IKE-ID validation with remote-identity and the certificate authentication. To avoid this, during authentication of remote peer, use the `general-ikeid` under the`set security ike gateway gateway_name dynamic` hierarchy level to bypass the validation process.

If you enable this option, then during authentication of remote peer, the device accepts all ike-id types like, hostname, user@hostname, and so on.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1

RELATED DOCUMENTATION

Example: Configuring AutoVPN with Pre-Shared Key

general-param (Flat-File Accounting Options)

IN THIS SECTION

- Syntax | 7670
- Hierarchy Level | 7671
- Description | 7671
- Options | 7671
- Required Privilege Level | 7672
- Release Information | 7672

Syntax

```
general-param {
    all-fields;
    accounting-type;
    descr;
    line-id;
    logical-interface;
    nas-port-id;
```

```
    physical-interface;
    routing-instance;
    timestamp;
    user-name;
    vlan-id;
}
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name fields]
```

Description

Specify general, nonstatistical interface parameters that are displayed as part of the header for the accounting file.

Options

all-fields Display all available nonstatistical fields. Many fields are available for both logical interfaces and physical interfaces, but some fields are available for only one interface type.

accounting-type (Logical interfaces only) Display the accounting status type.

descr Display the description of the interface as configured.

line-id (Logical interfaces only) Display the access line identifier.

logical-interface (Logical interfaces only) Display the name of the logical interface.

nas-port-id (Logical interfaces only) Display the NAS port ID.

physical-interface (Physical interfaces only) Display the name of the physical interface.

routing-instance Display the name of the routing instance to which the interface belongs.

timestamp Display the timestamp of the accounting record.

user-name Display the name of the subscriber.

vlan-id (Logical interfaces only) Display the VLAN identifier.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

`user-name` option added in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Configuring Service Accounting in Local Flat Files

Flat-File Accounting Overview

generate

IN THIS SECTION

- Syntax | 7672
- Hierarchy Level | 7673
- Description | 7673
- Options | 7673
- Required Privilege Level | 7674
- Release Information | 7674

Syntax

```
generate {
    defaults {
```

```
        generate-options;
    }
    route destination-prefix {
        policy policy-name;
        generate-options;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name],
[edit routing-options],
[edit routing-options rib routing-table-name]
```

Description

Configure generated routes, which are used as routes of last resort.

Options

defaults—(Optional) Specify global generated route options. These options only set default attributes inherited by all newly created generated routes. These are treated as global defaults and apply to all the generated routes you configure in the `generate` statement.

generate-options—Additional information about generated routes, which is included with the route when it is installed in the routing table. Specify zero or more of the following options in ***generate-options***. Each option is explained separately.

- **(active | passive);**
- **as-path <*as-path*> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator *as-number in-address*>;**
- **(brief | full);**
- **community [*community-ids*];**
- **discard;**

- **(metric | metric2 | metric3 | metric4)** ***value*** **<type** ***type*****>;**
- **(preference | preference2 | color | color2)** ***preference*** **<type** ***type*****>;**
- **tag** ***metric*** **type** ***number*****;**

route ***destination-prefix***—Configure a non-default generated route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/ prefix-length***—/ ***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

The `policy` statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

generate-event

IN THIS SECTION

- Syntax | 7675
- Hierarchy Level | 7675
- Description | 7675
- Options | 7675
- Required Privilege Level | 7676
- Release Information | 7676

Syntax

```
generate-event {
    event-name (time-interval seconds <start-time start-time> | time-of-day time) <no-drift>;
}
```

Hierarchy Level

```
[edit event-options]
```

Description

Generate an internal event based on a time interval or the time of day. You can configure up to 10 internal events.

Configure `time-interval` to specify the frequency at which to generate an event, or configure `time-of-day` to specify the time at which to generate a daily event.

Options

event-name Name of an internally generated event.

no-drift Do not propagate the delay caused in triggering an event to the triggering of the next event.

start-time *start-time* (Optional) Date and time at which to start generating an interval-based event, which is an event configured using the `time-interval` option.

The start time is relative to the local device time and is configured using 24-hour time format *yyyy-mm-dd.hh:mm*. You can optionally include a UTC offset to specify a time relative to UTC (Coordinated Universal Time) by appending `(+|-)`*hhmm* to the time, where the sign is plus (+) for east of UTC and minus (-) for west of UTC, for example, 2020-09-14.11:00-0800. Even if you specify a UTC offset, the `start-time` configuration always displays the time and UTC offset for the device's local time zone.

- **Syntax:** *yyyy-mm-dd.hh:mm*`(+|-)`*hhmm*

time-interval *seconds* Frequency in seconds at which to generate an interval-based internal event.

You can optionally configure the `start-time` option to specify the date and time at which to start generating the events. If you do not specify a start time, Junos OS generates the first event at approximately one time interval after you commit the configuration.

- **Range:** 60 through 2,592,000 seconds

time-of-day *time*

Time of day at which to generate a daily internal event.

The time is relative to the local device time and is configured using 24-hour time format *hh:mm:ss*. You can optionally include a UTC offset to specify a time relative to UTC by appending `(+|-)`*hhmm* to the time, where the sign is plus (+) for east of UTC and minus (-) for west of UTC, for example, 11:00:00-0800. Even if you specify a UTC offset, the `time-of-day` configuration always displays the time and UTC offset for the device's local time zone.

- **Syntax:** *hh:mm:ss*`(+|-)`*hhmm*

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

`no-drift` option added in Junos OS Release 14.1R1.

`start-time` option added in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Generate Internal Events to Trigger Event Policies

generate-snmp-traps

IN THIS SECTION

- Syntax | 7677
- Hierarchy Level | 7677

- Description | 7677
- Required Privilege Level | 7677
- Release Information | 7677

Syntax

```
generate-snmp-traps;
```

Hierarchy Level

```
[edit services]
[edit services video-monitoring]
```

Description

If this statement is configured, the service generates SNMP traps for severity levels such as Info, Warning, Critical, or Cleared. For example, if DF alarm changes from info to warning, or from warning to critical, mdiDFAlarm trap be triggered.

NOTE: SNMP traps are not generated if SNMP trap generation is not enabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Inline Video Monitoring on MX Series Routers

alarms (Video Monitoring) | 2125

geo-ip (global-filter)

IN THIS SECTION

- Syntax | 7678
- Hierarchy Level | 7678
- Description | 7679
- Options | 7679
- Required Privilege Level | 7679
- Release Information | 7679

Syntax

```
geo-ip{
rule {
    match {
    country country-name;
           }
        then {
        action (drop | drop-and-sample);
             }
       }
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name security-intelligence-policy]
```

Description

Configure the global GeoIP policy for the web-filter profile.

Options

- `country`—A list of country codes for which the actions defined under the profile is applied to all templates.
- `drop`—Drop the packets.
- `drop-and-sample`—Drop and sample the packets.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1 on MX Series routers with Juniper Advanced Threat Prevention Cloud .

RELATED DOCUMENTATION

security-intelligence-policy | 16467

URL Filtering

geo-ip (group-filter)

IN THIS SECTION

Syntax | 7680

Hierarchy Level | 7680

Description | 7680

- Options | 7680
- Required Privilege Level | 7681
- Release Information | 7681

Syntax

```
geo-ip{
rule {
    match {
          group {
                  (group-name1, group-name2) ;
                 }
            }
     then {
           action(drop | drop-and-sample);
             }
          }
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name url-filter-template template-name security-
intelligence-policy]
```

Description

Configure the group-based GeoIP policy for the web-filter template.

Options

- `group`—A group of countries to which the defined action is applicable.
- `drop`—Drop the packets.
- `drop-and-sample`—Drop and sample the packets.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1 on MX Series routers with Juniper Advanced Threat Prevention Cloud .

RELATED DOCUMENTATION

url-filter-template | 30496

GeoIP Filtering

get-route-range

IN THIS SECTION

- Syntax | 7681
- Hierarchy Level | 7682
- Description | 7682
- Required Privilege Level | 7682
- Release Information | 7682

Syntax

```
get-route-range;
```

Hierarchy Level

```
[edit logical-systems name policy-options policy-statement name from route-filter],
[edit logical-systems name policy-options policy-statement name term name from route-filter],
[edit policy-options policy-statement name from route-filter],
[edit policy-options policy-statement name term name from route-filter]
```

Description

Get the range of IPv4 prefixes that goes over a particular tunnel.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.2R1.

ggsn-address (OCS Partition)

IN THIS SECTION

- Syntax | 7683
- Hierarchy Level | 7683
- Description | 7683
- Options | 7683
- Required Privilege Level | 7683
- Release Information | 7683

Syntax

```
ggsn-address address;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Include the GGSN-Address AVP value in all CCR-GY messages.

Options

address Address of the GGSN-Address AVP value to include in all CCR-GY messages.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

ggsn-mcc-mnc (OCS Partition)

IN THIS SECTION

- Syntax | 7684
- Hierarchy Level | 7684
- Description | 7684
- Options | 7684
- Required Privilege Level | 7684
- Release Information | 7685

Syntax

```
ggsn-mcc-mnc ggsn-mcc-mnc;
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Configure the value of the 3GPP-GGSN-MCC-MNC AVP value to include in all CCR-GY messages. The value is the mobile country code (MCC) and mobile network code (MNC) of the network that the GGSN belongs to.The combined MCC and MNC uniquely identify the mobile network operator.

Options

ggsn-mcc-mnc Value of the mobile network country and network codes to include in all CCR-GY messages.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

gigether-options

IN THIS SECTION

- Syntax | 7685
- Hierarchy Level | 7687
- Description | 7687
- Required Privilege Level | 7693
- Release Information | 7693

Syntax

```
gigether-options {
    802.3ad {
        aex (primary | backup);
        lacp {
            port-priority;
        }
    }
    (asynchronous-notification | no-asynchronous-notification);
```

```
    (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online | local-interface-
offline>;
    fec (gigether)
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    (loopback | no-loopback);
    loopback-remote
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
    no-auto-mdix
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
    speed
    ethernet-switch-profile {
        (mac-learn-enable | no-mac-learn-enable);
        tag-protocol-id [ tpids ];
        ethernet-policer-profile {
            input-priority-map {
                ieee802.1p premium [ values ];
            }
            output-priority-map {
                classifier {
                    premium {
                        forwarding-class class-name {
                            loss-priority (high | low);
                        }
                    }
                }
            }
            policer cos-policer-name {
                aggregate {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
                premium {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
```

```
            }
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure Gigabit Ethernet specific interface properties.

In Junos OS Evolved, when you configure `set interfaces interface gigether-options 802.3ad ae name` at the same time as you apply a second configuration to the same interface at the `[edit interfaces interface]` hierarchy, the second configuration will not take effect until the interface joins the aggregated Ethernet interface *ae name*.

Table 47 on page 7688 shows the supported and unsupported platforms.

Table 47: Supported Platform Information

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
ACX Series Routers (Junos OS) • ACX500 • ACX710 • ACX1000 • ACX1100 • ACX2100 • ACX2200 • ACX4000 • ACX5448 • ACX5048 • ACX5096	ACX Series Routers (Junos OS) • ACX5048 • ACX5096	To configure storm control on the ACX Series routers, use `ether-options`.

Table 47: Supported Platform Information *(Continued)*

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
EX Series Switches (Junos OS) • EX4300 Multigigabit • EX4400 • EX4600 • EX4650 • EX9200 • EX9250	EX Series Switches (Junos OS) • EX2300 • EX2300 Multigigabit • EX2300-C • EX3400 • EX4300 • EX4300 Multigigabit • EX4400 • EX4600 • EX4650 • EX9200 • EX9250	EX Series Switches (Junos OS) EX Series switches support both `ether-options` and `gigether-options`. To configure Ethernet configuration options such as aggregated Ethernet, loopback, flow-control, auto-negotiation etc., use `ether-options`. To configure Forward Error Correction (FEC), use `gigether-options`.

Table 47: Supported Platform Information *(Continued)*

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
MX Series Routers (Junos OS) • MX5 • MX10 • MX40 • MX80 • MX104 • MX150 • MX204 • MX240 • MX480 • MX960 • MX20008 • MX2010 • MX2020 • MX10003 • MX10008 and MX10016	Not Supported	None

Table 47: Supported Platform Information *(Continued)*

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
PTX Series Routers (Junos OS) • PTX1000 • PTX3000 • PTX5000 • PTX10002 • PTX10008 and PTX10016	PTX Series Routers (Junos OS) • PTX1000 • PTX10002 • PTX10008 and PTX10016	PTX Series Routers (Junos OS) In Junos OS Release 17.3R3S7, PTX1000 Series routers support both `ether-options` and `gigether-options`. In Junos OS Releases 17.3R1, 17.4R1, and 17.4R2, PTX10000 Series routers support both `ether-options` and `gigether-options`. In Junos OS Release 18.1R1 and later, PTX Series routers support only `gigether-options` for your configuration and do not support `ether-options`. Before Junos OS Release 18.1R1, PTX Series support `ether-options`. See TSB17864 for additional information. **NOTE**: In Junos OS Release 19.4R2, PTX10002-60C supports `ether-options`.
PTX Series routers (Junos OS Evolved) • PTX10003	PTX Series routers (Junos OS Evolved) • PTX10001-36MR • PTX10003 • PTX10004 • PTX10008 • PTX10016	PTX Series routers (Junos OS Evolved) Starting in Junos OS Evolved Release 20.1R1, PTX Series routers support `ether-options` only. Before Junos OS Evolved Release 20.1R1, PTX Series routers supported `gigether-options` only.

Table 47: Supported Platform Information *(Continued)*

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
QFX Series Switches (Junos OS) • QFX5100 (48S) • QFX5100 (48T) • QFX5100 (24Q) • QFX5100 (96S) • QFX5110 (48S) • QFX5110 (32Q) • QFX5120 (48Y) • QFX5120 (48T) • QFX5120 (48ym) • QFX5120 (32C) • QFX5200 (48Y) • QFX5200 (32C) • QFX5210 (64C) • QFX5210 (64C-S) • QFX5220 (32CD) • QFX5220 (128C) • QFX10002 • QFX10008 and QFX10016	QFX Series Switches (Junos OS) • QFX5100 (48S) • QFX5100 (48T) • QFX5100 (24Q) • QFX5100 (96S) • QFX5110 (48S) • QFX5110 (32Q) • QFX5120 (48Y) • QFX5120 (32C) • QFX5200 (48Y) • QFX5120 (48T) • QFX5120 (48ym) • QFX5200 (32C) • QFX5210 (64C) • QFX5210 (64C-S) • QFX5220 (32CD) • QFX5220 (128C) • QFX10002 • QFX10008 and QFX10016	QFX Series Switches (Junos OS) QFX Series switches support both `ether-options` and `gigether-options`. To configure Ethernet configuration options such as loopback, flow-control, auto-negotiation etc., use `ether-options`. To configure Forward Error Correction (FEC), use `gigether-options`.

Table 47: Supported Platform Information *(Continued)*

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
Not Supported	QFX Series Switches (Junos OS Evolved) • QFX5200-32C-L • QFX5220-32CD • QFX5220-128C	To configure auto-negotiation on management interfaces (`re0:mgmt-0`), use `ether-options`.
SRX Series (Junos OS) • SRX300 • SRX550 • SRX1500 • SRX4100 and SRX4200 • SRX4600 • SRX5400 • SRX5600	SRX Series (Junos OS) • SRX300 • SRX550 • SRX1500 • SRX4100 and SRX4200 • SRX4600 • SRX5400 • SRX5600	SRX Series (Junos OS) To configure gigabit-Ethernet interfaces (ge-), use `gigether-options`. To configure ethernet interfaces (et-) and fast ethernet interfaces (fe-), use `ether-options`.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Ethernet Interfaces Overview

ether-options | 6461

gigether-options (Chassis Cluster)

IN THIS SECTION

- Syntax | 7694
- Hierarchy Level | 7696
- Description | 7696
- Options | 7696
- Required Privilege Level | 7696
- Release Information | 7696

Syntax

```
gigether-options {
    802.3ad {
        backup | primary | bundle;
        lacp {
            port-priority priority;
        }
    }
    auto-negotiation {
        remote-fault {
            local-interface-offline | local-interface-online;
        }
    }
    no-auto-negotiation;
    ethernet-switch-profile {
        mac-learn-enable;
        tag-protocol-id [tpids];
        ethernet-policer-profile  {
            input-priority-map {
                ieee802.1p {
                    premium [values];
```

```
                }
            }
            output-priority-map {
                classifier {
                    premium {
                        forwarding-class  class-name {
                            loss-priority (high | low);
                        }
                    }
                }
            }
            policer cos-policer-name {
                aggregate {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
                premium  {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
            }
        }
    }
    flow-control | no-flow-control;
    ieee-802-3az-eee;
    ignore-l3-incompletes;
    loopback  | no-loopback;
    mpls {
        pop-all-labels {
            required-depth (1 | 2);
        }
    }
    redundant-parent (Interfaces Gigabit Ethernet) interface-name;
    source-address-filter {
        mac-address;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure Gigabit Ethernet specific interface properties.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Example: Configuring Chassis Cluster Redundant Ethernet Interfaces

global-advertisements-threshold

IN THIS SECTION

- Syntax | 7697
- Hierarchy Level | 7697
- Description | 7697
- Options | 7698

Required Privilege Level | 7698
Release Information | 7698

Syntax

```
global-advertisements-threshold advertisement-value;
```

Hierarchy Level

```
[edit protocols vrrp]
```

Description

Configure the number of fast advertisements that can be missed by a backup router before the primary router is declared as down.

NOTE:

- The advertisement value configured using the `global-advertisements-threshold` statement is applicable to all the Virtual Router Redundancy Protocol (VRRP) groups in the system.
- Setting the advertisement value of the **global-advertisements-threshold** configuration to **1** is not recommended for a scaled configuration with an aggressive advertisement interval. For example, if you have 1000 VRRP groups with an advertisement interval of 100 ms, then do not set the **global-advertisements-threshold** value to 1.
- Changing the advertisement value of the **global-advertisements-threshold** configuration during runtime can result in unpredictable behavior by the VRRP state machine. For example, momentary ownership change from the primary router to the backup router and vice versa. Therefore, avoid changing the advertisement value of the `global-advertisements-threshold` statement during runtime.

Options

advertisement-value—Number of VRRP advertisements missed before the primary router is declared as down.

- **Range:** 1 through 15
- **Default:** 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Improving the Convergence Time for VRRP

Configuring VRRP to Improve Convergence Time

global-arp-prefix-limit (Host Fast Reroute)

IN THIS SECTION

- Syntax | 7699
- Hierarchy Level | 7699
- Description | 7699
- Default | 7700
- Options | 7700
- Required Privilege Level | 7700
- Release Information | 7700

Syntax

```
global-arp-prefix-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options host-fast-reroute],
[edit routing-options host-fast-reroute]
```

Description

Set the ARP prefix limit for all host fast-reroute (HFRR) profiles on the routing device.

When you configure HFRR, an optional ARP prefix limit sets a maximum for the number of ARP routes and, therefore FRR routes created for each HFRR profile in the routing table. This limit prevents ARP attacks from exhausting the virtual memory on the routing devices.

There are two configuration statements (`global-arp-prefix-limit` and `arp-prefix-limit`) that set the ARP prefix limit, one at the global `[edit routing-options host-fast-reroute]` hierarchy level and the other at the `[edit routing-instances instance-name routing-options interface interface-name]` hierarchy level, respectively. The global `global-arp-prefix-limit` statement sets a default ARP prefix limit for all HFRR profiles configured on the routing device. The `arp-prefix-limit` statement overrides the `global-arp-prefix-limit` for that HFRR profile for that protected interface.

Warning system log messages begin when the ARP routes in an HFRR profile reaches 80% of the configured limit. When the number crosses the 100% threshold, the HFRR profile is deactivated. When this happens, all ARP/FRR routes are deleted from the routing table. FRR routes are deleted from forwarding table as well.

After the HFRR profile is deactivated, a blackout timer is started. The timeout value of this timer is the ARP cache timeout (kernel timeout) + the supplementary blackout timer.

There are global and per-HFRR CLI statements (`global-supplementary-blackout-timer` and `supplementary-blackout-timer`) to configure the supplementary blackout timer. The global value is at the `[edit routing-options host-fast-reroute]` hierarchy level and applies to all HFRR profiles on the routing device. The value for the routing-instance interface is at the `[edit routing-instances instance-name routing-options interface interface-name]` hierarchy level, and overrides the global value for that HFRR profile only.

When the blackout timer expires, the HFRR profile is reactivated, and the Junos OS relearns the ARP routes and re-creates the HFRR routes. If the ARP prefix limit is not exceeded again, the HFRR routes will be up.

If an HFRR profile is in the deactivated state, a reevaluation of the ARP state is preformed during every commit operation or whenever the routing process (rpd) is restarted with the `restart routing` command.

Default

If you omit the-`arp-prefix-limit` statement, the `global-arp-prefix-limit` takes effect for all HFRR profiles on the device. If you omit both of these statements, there is no ARP prefix limit for host fast reroute.

Options

number Maximum number of ARP HFRR routes allowed.

- **Range:** 1 through 10,000 HFRR routes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring Link Protection with Host Fast Reroute

global-config (Services)

IN THIS SECTION

- Syntax | **7701**
- Hierarchy Level | **7701**
- Description | **7701**

- Options | 7701
- Required Privilege Level | 7703
- Release Information | 7703

Syntax

```
global-config {
cache-usage-enforcement-threshold cache-usage-enforcement-threshold;
certificate-cache-timeout seconds;
disable-cert-cache;
disable-deferred-profile-selection;
invalidate-cache-on-crl-update;
non-ssl-detection-threshold {
    byte-threshold byte-threshold;
    packet-threshold packet-threshold;
}
session-cache-timeout seconds;
}
```

Hierarchy Level

```
[edit services ssl proxy]
```

Description

Specify the global proxy configuration. When SSL proxy is configured at a global level (within “services ssl proxy”), it is visible across the system configurations on the device.

Options

cache-usage-enforcement-threshold *cache-usage-enforcement-threshold*

Percentage of total cache size after which per logical system limits will be enforced.

- **Range:** 1 through 100

certificate-cache-timeout *seconds*

Regulates the certificate cache timeout.

- **Default:** 600 seconds

disable-cert-cache

Disable the certificate cache. By default certificate cache is enabled.

disable-deferred-profile-selection

Disable the deferred profile selection mechanism. In the defered profile selection mechanism, the SSL proxy module defers SSL profile selection until the dynamic application is detected in a client hello message based on the Server Name Indication (SNI). After detecting dynamic application, SSL proxy module does a firewall rule lookup based on the identified application and selects an appropriate SSL proxy profile.

invalidate-cache-on-crl-update

Invalidate the existing certificate cache. By default, this option is disabled.

non-ssl-detection-threshold

Set limit that allows you to decide how long to wait before ignoring the the session if StartTLS is not received from the client.

- byte-threshold *byte-threshold*—Minimum bytes required to ignore the session. SSL proxy allows the configured number of bytes of plain (unencrypted) SMTP. After reaching the limit, it ignores the session if StartTLS is not received from the client.
  - **Default:** 200
  - **Range:** 1 through 600
- packet-threshold *packet-threshold*—Number of plain (unencrypted) packets in client-to-server direction. SSL proxy allows the configured number of packets of plain (unencrypted) SMTP. After reaching the limit, it ignores the session if StartTLS is not received from the client.
  - **Default:** 5
  - **Range:** 1 through 15

> **NOTE**: Starting in Junos OS Release 23.4R1, we've renamed the `mail-threshold` statement to `non-ssl-detection-threshold`.

> **NOTE**: In case your firewall is processing FTP or SMTP traffic and has unified policy (dynamic-application) and ssl-proxy configured, we recommend the following configuration:

```
[edit]
user@host# set services ssl proxy global-config non-ssl-detection-
threshold
```

session-cache-timeout *seconds*

Specify the session cache timeout.

- **Range:** 300 to 3600 seconds

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. `disable-cert-cache`, `certificate-cache-timeout`, and `Invalidate-cache-on-crl-update` options are introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

SSL Proxy Overview

Configuring SSL Forward Proxy

Enabling Debugging and Tracing for SSL Proxy

global (DDoS)

IN THIS SECTION

- Syntax | 7704
- Hierarchy Level | 7704
- Description | 7704
- Required Privilege Level | 7704

Release Information | 7705

Syntax

```
global {
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection;
    flow-level-control;
    flow-detection-mode;
    flow-report-rate;
    violation-report-rate;
}
```

Hierarchy Level

```
[edit system ddos-protection]
```

Description

(MX Series routers with only MPCs, T4000 routers with only FPC5s, PTX Series routers, EX9200 switches, or QFX Series switches) Modify DDoS policers, event logging, and flow detection globally for all protocols.

> **NOTE**: The following statements are not supported on PTX Series routers and QFX Series switches: `disable-routing-engine`, `flow-detection`, `flow-report-rate`, and `violation-report-rate`.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Disabling Control Plane DDoS Protection Policers and Logging Globally

global-dns-stats-log-timer

IN THIS SECTION

- Syntax | 7705
- Hierarchy Level | 7705
- Description | 7706
- Options | 7706
- Required Privilege Level | 7706
- Release Information | 7706

Syntax

```
global-dns-stats-log-timer minutes;
```

Hierarchy Level

```
[edit services web-filter profile profile-name]
```

Description

Configure the interval for logging per-client statistics for filtering of DNS requests for disallowed website domains.

Options

minutes The number of minutes in the logging interval.

- **Default:** 5
- **Range:** 0 through 60

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

global (Gx-Plus)

IN THIS SECTION

- Syntax | **7707**
- Hierarchy Level | **7707**
- Description | **7707**

Required Privilege Level | 7707
Release Information | 7707

Syntax

```
global {
    include-ipv6;
    max-outstanding-requests number;
}
```

Hierarchy Level

```
[edit access gx-plus]
```

Description

Configure global attributes for the Gx-Plus application.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Gx-Plus

global-limit

IN THIS SECTION

- Syntax | 7708
- Hierarchy Level | 7708
- Description | 7708
- Options | 7708
- Required Privilege Level | 7709
- Release Information | 7709

Syntax

```
global-limit global-limit;
```

Hierarchy Level

```
[edit logical-systems name security utm default-configuration url-feed],
[edit logical-systems name tenants name security utm default-configuration url-feed],
[edit security utm default-configuration url-feed],
[edit tenants name security utm default-configuration url-feed]
```

Description

Configure global URL feed limit. You can define the maximum number of URL feeds on your device.

Options

global-limit	Configure the maximum number of URL feeds.

- **Default:** 20
- **Range:** 1 through 50

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

url-feed | 30489

global-mac-ip-limit

IN THIS SECTION

- Syntax | 7710
- Hierarchy Level | 7710
- Description | 7710
- Options | 7710
- Required Privilege Level | 7710
- Release Information | 7710

Syntax

```
global-mac-ip-limit {
    number;
}
```

Hierarchy Level

```
[edit logical-systems name protocols l2-learning],
[edit protocols l2-learning]
```

Description

Limit the number of entries that can be added systemwide to the MAC-IP bindings database.

Options

number Specify the maximum number of entries that can added systemwide to the MAC-IP bindings database.. When the specified maximum is reached, no new entries are added to the database.

- **Range:** 20 through 1,048,575.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R2.

Statement introduced in Junos OS Release 18.1R1 for QFX Series switches.

RELATED DOCUMENTATION

EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression

global-mac-ip-table-aging-time | 7711
interface-mac-ip-limit | 8908
mac-ip-table-size | 10231

global-mac-ip-table-aging-time

IN THIS SECTION

- Syntax | 7711
- Hierarchy Level | 7711
- Description | 7711
- Options | 7712
- Required Privilege Level | 7712
- Release Information | 7712

Syntax

```
global-mac-ip-table-aging-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2-learning],
[edit protocols l2-learning]
```

Description

Configure the timeout interval systemwide for entries in the MAC-IP address bindings database.

Options

seconds Specify the time that is elapsed before entries in the MAC-IP bindings database are timed out and deleted.

- **Range:** 60 through 1 million seconds
- **Default:** 1200 seconds (20 minutes)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

Statement introduced in Junos OS Release 18.1R1 for QFX Series switches.

RELATED DOCUMENTATION

EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression

global-mac-ip-limit | **7709**

interface-mac-ip-limit | **8908**

mac-ip-table-size | **10231**

global-mac-limit

IN THIS SECTION

- Syntax | **7713**
- Hierarchy Level | **7713**
- Description | **7713**
- Default | **7713**

- Options | 7713
- Required Privilege Level | 7713
- Release Information | 7714

Syntax

```
global-mac-limit limit {
    packet-action drop;
}
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

(MX Series routers and EX Series switches only) Limit the number of media access control (MAC) addresses learned from the logical interfaces on the router or switch.

Default

(MX Series) 393,215 MAC addresses

(EX9200) 524,287 MAC addresses

Options

limit—Number of MAC addresses that can be learned systemwide.

- **Range:** 20 through 1,048,575

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Limiting the Number of MAC Addresses Learned from Each Logical Interface

global-mac-limit (Protocols)

IN THIS SECTION

- Syntax | 7714
- Hierarchy Level | 7715
- Description | 7715
- Default | 7715
- Options | 7715
- Required Privilege Level | 7715
- Release Information | 7715

Syntax

```
global-mac-limit limit {
    packet-action drop;
}
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

Limit the number of media access control (MAC) addresses learned from the logical interfaces on the router.

Default

131,071 MAC addresses

> **NOTE**: SRX300, SRX320, SRX340, and SRX345 devices support 16,383 addresses, and SRX1500 devices support 24,575 addresses.

Options

limit—Number of MAC addresses that can be learned on the device.

- **Range:** 20 through 13,1071 addresses

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.5.

global-mac-move

IN THIS SECTION

- Syntax | 7716
- Hierarchy Level | 7716
- Description | 7716
- Default | 7717
- Required Privilege Level | 7717
- Release Information | 7717

Syntax

```
global-mac-move  {
    cooloff-time seconds;
    disable-action;
    exclusive-mac virtual-mac-mac-address/mask;
    interface-recovery-time seconds;
    notification-time seconds;
    reopen-time seconds;
    statistical-approach-wait-time seconds;
    threshold-count count;
    threshold-time seconds;
    virtual-mac mac-address /mask;
}
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

Set parameters for media access control (MAC) address move reporting.

Default

By default, MAC moves notify every second, with a threshold time of 1 second and a threshold count of 50.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

Support for disable-action and reopen-time added in Junos OS Release 13.2.

Support for exclusive-mac added in Junos OS Release 14.1X53-D45.

Statements `cooloff-time`, `interface-recovery-time`, `statistical-approach-wait-time`, and `virtual-mac` moved from vpls-mac-move to global-mac-move hierarchy level in Junos OS Release 17.4.

RELATED DOCUMENTATION

Configuring MAC Move Parameters

MAC Moves Loop Prevention in VPLS Network Overview

Example: Configuring Loop Prevention in VPLS Network Due to MAC Moves

virtual-mac | **30766**

global-mac-statistics

IN THIS SECTION

- Syntax | **7718**
- Hierarchy Level | **7718**

- Description | 7718
- Default | 7718
- Options | 7718
- Required Privilege Level | 7718
- Release Information | 7718

Syntax

```
global-mac-statistics;
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

(MX Series routers and EX Series switches only) Enable MAC accounting for the entire router or switch.

Default

disabled

Options

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Enabling MAC Accounting

global-mac-table-aging-time

IN THIS SECTION

- Syntax | 7719
- Hierarchy Level | 7719
- Description | 7719
- Default | 7720
- Options | 7720
- Required Privilege Level | 7720
- Release Information | 7720

Syntax

```
global-mac-table-aging-time seconds;
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

Configure the timeout interval for entries in the MAC table.

NOTE: The `global-mac-table-aging-time` statement appears in the Junos OS CLI for devices that support the Enhanced Layer 2 Software (ELS) configuration style. If your device runs software that does *not* support ELS, use the "mac-table-aging-time" on page 10267 statement, which appears in the `[edit ethernet-switching-options]` and the `[edit vlans]` hierarchies. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

Default

300 seconds

Options

seconds—Time elapsed before MAC table entries are timed out and entries are deleted from the table.

- **Range:** For MX Series routers: 10 through 1 million; for EX Series and QFX Series switches: 60 through 1 million; for SRX Series Firewalls: 10 through 64,000 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Statement modified in Junos OS Release 9.5.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring the MAC Table Timeout Interval

Configuring MAC Table Aging on Switches

Example: Configuring VLANs on Security Devices

global-mode (Protocols)

IN THIS SECTION

- Syntax | 7721
- Hierarchy Level | 7721
- Description | 7721
- Default | 7721
- Required Privilege Level | 7722
- Release Information | 7722

Syntax

```
global-mode (switching | transparent-bridge) ;
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

Specify the global mode for the SRX Series Firewall as Layer 2 transparent bridge mode or switching mode. After changing the mode, you must reboot the device for the configuration to take effect.

Default

On SRX1500, the default Layer 2 global mode is transparent-bridge mode.

Starting with Junos OS Release 15.1X49-D100, on SRX300, SRX320, SRX340, SRX345, SRX550, and SRX550M devices, the default Layer 2 global mode configuration is changed from transparent-bridge to switching mode.

NOTE: You must explicitly configure Layer 2 transparent-bridge mode for the SRX300, SRX320, SRX340, SRX345, SRX550, and SRX550M devices that work in transparent mode. Use the command `set protocols l2-learning global-mode transparent-bridge` before rebooting the devices with Junos OS 15.1X49-D100 image.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D40.

global-no-control-mac-aging

IN THIS SECTION

- Syntax | 7722
- Hierarchy Level | 7723
- Description | 7723
- Required Privilege Level | 7723
- Release Information | 7723

Syntax

```
global-no-control-mac-aging;
```

Hierarchy Level

```
[edit logical-systems name protocols l2-learning],
[edit protocols l2-learning]
```

Description

Disable control MAC-address aging from software.

Required Privilege Level

routing—To view this statement in the configuration.

routing–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

global-no-mac-learning

IN THIS SECTION

- Syntax | 7724
- Hierarchy Level | 7724
- Description | 7724
- Default | 7724
- Required Privilege Level | 7724
- Release Information | 7724

Syntax

```
global-no-mac-learning;
```

Hierarchy Level

```
[edit protocols l2-learning],
[edit protocols l2-learning]
```

Description

Disable MAC learning on the entire device.

Default

MAC learning is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Statement modified for SRX Series in Junos OS Release 9.5.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Disabling Layer 2 Learning and Forwarding

Understanding Q-in-Q Tunneling and VLAN Translation

Example: Configuring VLANs on Security Devices

global (OCS)

IN THIS SECTION

- Syntax | 7725
- Hierarchy Level | 7725
- Description | 7725
- Required Privilege Level | 7726
- Release Information | 7726

Syntax

```
global {
    service-context-id service-context ;
}
```

Hierarchy Level

```
[edit access ocs]
```

Description

Configure global attributes and data elements of the 3rd Generation Partnership Project (3GPP) Diameter credit control service charging system for the Online Charging System (OCS), which interacts with the Policy and Charging Enforcement Function (PCEF). The PCEF optionally reports usage and receives additional authorizations from the OCS using the 3GPP Gy protocol. Broadband PCEF (BPCEF) interactions with the OCS use online session charging with centralized unit determination and centralized rating.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring OCS Global Parameters

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

global-options (Juniper Secure Connect)

IN THIS SECTION

- Syntax | 7727
- Hierarchy Level | 7727
- Description | 7727
- Options | 7727
- Required Privilege Level | 7727
- Release Information | 7727

Syntax

```
global-options {
    auth-token-valid-time seconds;
}
```

Hierarchy Level

```
[edit security remote-access]
```

Description

Define global parameters for Juniper Secure Connect remote access configuration.

Options

auth-token-valid-time

Authentication token valid time (seconds).

- **Default:** 60 seconds
- **Range:** 1 through 300

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.3R1

RELATED DOCUMENTATION

Juniper Secure Connect Administrator Guide

global-parameters (Easy EVPN LAG Configuration)

IN THIS SECTION

- Syntax | 7728
- Hierarchy Level | 7728
- Description | 7729
- Options | 7729
- Required Privilege Level | 7731
- Release Information | 7731

Syntax

```
global-parameters {
  anycast-mac anycast-mac;
  default-vlan-id vlan-id;
  mtu {
    overlay overlay-mtu;
    underlay underlay-mtu;
  }
  no-irb-address-auto-derive;
  start-aggregate-ethernet-index num;
  virtual-gateway {
    v4-mac v4-mac;
    v6-mac v6-mac;
  }
  virtual-gateway-mac virtual-gateway-mac;
}
```

Hierarchy Level

```
[edit services evpn]
```

Description

Specify global parameters as part of a generated Ethernet VPN (EVPN) fabric configuration.

You use the statements in the `[edit services evpn]` hierarchy level to simplify configuring an EVPN fabric with Ethernet segment identifier (ESI) link aggregation groups (LAGs) in the topology. The feature you enable with statements in this hierarchy is an easy EVPN LAG configuration feature, also called EZ-LAG.

You only need to configure a few statements to provide the parameters required for an EVPN core fabric with multihomed or single-homed customer edge (CE) devices (or end server devices). When you commit the `[edit services evpn]` configuration statements, the device uses a commit script to generate the corresponding standard EVPN fabric configuration statements.

The statements at the `global-parameters` hierarchy level provide parameters for configuration elements that you want to be common across peer provider edge (PE) devices in the EVPN fabric.

Options

anycast-mac *anycast-mac-address*

Configure the device to globally use the specified anycast virtual gateway MAC address.

You must configure this statement if you include the `[edit services evpn evpn-vxlan irb name] use-anycast-address` option for any integrated routing and bridging (IRB) instances you configure on the PE device.

default-vlan-id *vlan-id*

Specify the VLAN ID you want the commit script to use in the configuration as the default VLAN.

- **Range:** 1 through 4094

mtu (overlay *overlay-mtu* **| underlay** *underlay-mtu***)**

Set the specified maximum transmission unit (MTU) in bytes in the generated configuration for the interfaces used in the overlay peering or underlay peering. The MTU size is the largest amount of data the device can forward without fragmentation. You can specify MTU values for either or both the overlay and underlay interfaces.

- **Range:** The default MTU and configurable range vary by platform and type of interface or protocol for which you configure the value.

  See *Media MTU and Protocol MTU* for more on setting MTU sizes.

no-irb-address-auto-derive

Instruct the commit script to not derive IRB interface subnet addresses in the generated configuration.

For simplicity, with this feature you can set the same IRB interface subnet address on connected PE devices using the `subnet-address` option at the `[edit services evpn evpn-vxlan irb irb-instance]` hierarchy level. By default, the commit script derives different device IRB subnet addresses using the `subnet-address` value configured for the peer PE device with peer ID 1, based on the device peer ID where the script is running.

When you set this option, instead of the default behavior, the commit script uses exactly the IPv4 or IPv6 subnet addresses you specify with the `subnet-address` option for each peer ID. If you set the `no-irb-address-auto-derive` option, make sure you explicitly set the IRB `subnet-address` values you want for each IRB instance on each device.

> **NOTE**: This option takes effect globally for all IRB instances you configure with the `irb irb-instance` statement at the `[edit services evpn evpn-vxlan]` hierarchy level.
>
> You can alternatively set the `no-irb-address-auto-derive` option for individual IRB instances at the `[edit services evpn evpn-vxlan irb irb-instance subnet-address]` hierarchy level.

start-aggregate-ethernet-index *num*

Set a global starting index for numbering the aggregate Ethernet interfaces in multihomed server (ESI LAG) aggregated Ethernet interface configurations.

- **Range:** 1 through 255

virtual-gateway (v4-mac *mac-address* | v6-mac *mac-address*)

Configure the device to globally use the specified virtual gateway MAC address. With this option, you can specify different MAC addresses for IPv4 and IPv6 traffic.

> **NOTE**: Specify either this option *OR* the `virtual-gateway-mac` option (next).

If you don't specify these parameters, the commit script uses:

- ipv4-mac—00:00:5e:00:01:01
- ipv6-mac—00:00:5e:00:02:01

virtual-gateway-mac *mac-address*

Configure the device to globally use the specified virtual gateway MAC address. If you specify this option, the commit script uses the same MAC address for both IPv4 and IPv6 traffic.

> **NOTE**: Specify either this option *OR* the `virtual-gateway (v4-mac mac-address | v6-mac mac-address)` options above.

You can also consult CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Easy EVPN LAG (EZ-LAG) Configuration

global (PCRF)

IN THIS SECTION

- Syntax | 7732
- Hierarchy Level | 7732
- Description | 7732
- Required Privilege Level | 7732
- Release Information | 7732

Syntax

```
global {
    rule-param avp-code;
}
```

Hierarchy Level

```
[edit access pcrf]
```

Description

Configure global attributes and data elements of the 3rd Generation Partnership Project (3GPP) Diameter credit control service charging system for the Online Charging System (OCS), which interacts with the Policy and Charging Enforcement Function (PCEF). The PCEF optionally reports usage and receives additional authorizations from the OCS using the 3GPP Gy protocol. Broadband PCEF (BPCEF) interactions with the OCS use online session charging with centralized unit determination and centralized rating.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

global (Security IDP)

IN THIS SECTION

- Syntax | 7733
- Hierarchy Level | 7733
- Description | 7733
- Options | 7733
- Required Privilege Level | 7734
- Release Information | 7734

Syntax

```
global {
    (enable-all-qmodules | no-enable-all-qmodules);
    (enable-packet-pool | no-enable-packet-pool);
    memory-limit-percent value;
    (policy-lookup-cache | no-policy-lookup-cache);
}
```

Hierarchy Level

```
[edit security idp sensor-configuration]
```

Description

Configure the global rulebase IDP security policy.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

global (Security Policies)

IN THIS SECTION

- Syntax | 7734
- Hierarchy Level | 7737
- Description | 7737
- Options | 7737
- Required Privilege Level | 7737
- Release Information | 7737

Syntax

```
global {
    policy policy-name {
        description description;
        match {
            application {
                [application];
                any;
            }
            destination-address {
                [address];
                any;
                any-ipv4;
```

```
                any-ipv6;
            }
            from-zone {
                [zone-name];
                any;
            }
            source-address {
                [address];
                any;
                any-ipv4;
                any-ipv6;
            }
            source-identity {
                [role-name];
                any;
                authenticated-user;
                unauthenticated-user;
                unknown-user;
            }
                to-zone {
                    [zone-name];
                    any;
                }
        }
        scheduler-name scheduler-name;
        then {
                                count {
                alarm {
                    per-minute-threshold number;
                    per-second-threshold number;
                }
            }
            deny;
            log {
                session-close;
                session-init;
            }
            permit {
                application-services {
                    application-firewall {
                        rule-set rule-set-name;
                    }
                    application-traffic-control {
```

```
                        rule-set rule-set-name;
                    }
                    gprs-gtp-profile profile-name;
                    gprs-sctp-profile profile-name;
                    idp;
                    redirect-wx | reverse-redirect-wx;
                    ssl-proxy {
                        profile-name profile-name;
                    }
                    uac-policy {
                        captive-portal captive-portal;
                    }
                    utm-policy policy-name;
                }
                destination-address {
                    drop-translated;
                    drop-untranslated;
                }
                firewall-authentication {
                    pass-through {
                        access-profile profile-name;
                        client-match user-or-group-name;
                        web-redirect;
                    }
                    web-authentication {
                        client-match user-or-group-name;
                    }
                }
                services-offload;
                tcp-options {
                    sequence-check-required;
                    syn-check-required;
                }
            }
            reject;
        }
    }
}
```

Hierarchy Level

```
[edit security policies]
```

Description

Configure a global policy.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Statement updated with `from-zone` and `to-zone` policy match options in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Global Policy Overview

global-supplementary-blackout-timer (Host Fast Reroute)

IN THIS SECTION

- Syntax | 7738
- Hierarchy Level | 7738
- Description | 7738

- Default | 7739
- Options | 7739
- Required Privilege Level | 7739
- Release Information | 7739

Syntax

```
global-supplementary-blackout-timer minutes;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options host-fast-reroute],
[edit routing-options host-fast-reroute]
```

Description

Set the blackout timer for all host fast-reroute (HFRR) profiles on the routing device.

When you configure HFRR, an optional ARP prefix limit sets a maximum for the number of ARP routes and, therefore FRR routes created for each HFRR profile in the routing table. This limit prevents ARP attacks from exhausting the virtual memory on the routing devices.

There are two configuration statements (`global-arp-prefix-limit` and `arp-prefix-limit`) that set the ARP prefix limit, one at the global `[edit routing-options host-fast-reroute]` hierarchy level and the other at the `[edit routing-instances instance-name routing-options interface interface-name]` hierarchy level, respectively. The global `global-arp-prefix-limit` statement sets a default ARP prefix limit for all HFRR profiles configured on the routing device. The `arp-prefix-limit` statement overrides the `global-arp-prefix-limit` for that HFRR profile for that protected interface.

Warning system log messages begin when the ARP routes in an HFRR profile reaches 80% of the configured limit. When the number crosses the 100% threshold, the HFRR profile is deactivated. When this happens, all ARP/FRR routes are deleted from the routing table. FRR routes are deleted from forwarding table as well.

After the HFRR profile is deactivated, a blackout timer is started. The timeout value of this timer is the ARP cache timeout (kernel timeout) + the supplementary blackout timer.

There are global and per-HFRR CLI statements (`global-supplementary-blackout-timer` and `supplementary-blackout-timer`) to configure the supplementary blackout timer. The global value is at the `[edit routing-options host-fast-reroute]` hierarchy level and applies to all HFRR profiles on the routing device. The value for the routing-instance interface is at the `[edit routing-instances` *`instance-name`* `routing-options interface` *`interface-name`*`]` hierarchy level, and overrides the global value for that HFRR profile only.

When the blackout timer expires, the HFRR profile is reactivated, and the Junos OS relearns the ARP routes and re-creates the HFRR routes. If the ARP prefix limit is not exceeded again, the HFRR routes will be up.

If an HFRR profile is in the deactivated state, a reevaluation of the ARP state is preformed during every commit operation or whenever the routing process (rpd) is restarted with the `restart routing` command.

Default

If you omit the-`supplementary-blackout-timer` statement, the `global-supplementary-blackout-timer` takes effect for all HFRR profiles on the device. If you omit both of these statements, the blackout timeout value is set to the ARP cache timeout value, which by default is 20 minutes and is configurable with the `aging-timer` statement at the `[edit system arp]` hierarchy level.

Options

minutes Number of minutes, in addition to the ARP cache timeout value, before all HFRR profiles on the routing device are reactivated after the ARP prefix limit is exceeded.

- **Range:** 1 through 15 minutes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring Link Protection with Host Fast Reroute

global-threshold

IN THIS SECTION

- Syntax | 7740
- Hierarchy Level | 7740
- Description | 7740
- Options | 7740
- Required Privilege Level | 7741
- Release Information | 7741

Syntax

```
global-threshold number;
```

Hierarchy Level

```
[edit chassis cluster redundancy-group group-number ip-monitoring]
```

Description

Define global threshold for IP monitoring. This is the number that needs to be met or exceeded by all of the cumulative weights of the monitored IP addresses to trigger a failover.

When a monitored address is marked as unreachable, the weight value associated with that address is deducted from the the redundancy group IP address monitoring global threshold. If the accumulated monitored address weight values surpass the global-threshold value, that is, when the global threshold reaches 0, the global weight is deducted from the redundancy group threshold. Every redundancy group has a default threshold of 255. If the threshold reaches 0, a failover is triggered

Options

number—Value at which the IP monitoring weight is applied against the redundancy group failover threshold.

- **Range:** 0 through 255
- **Default:** 255

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

ip-monitoring | 9244

global-wait-to-restore

IN THIS SECTION

- Syntax | 7741
- Hierarchy Level | 7742
- Description | 7742
- Options | 7742
- Required Privilege Level | 7742
- Release Information | 7742

Syntax

```
global-wait-to-restore minutes;
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Use this statement to configure the global time to wait before opening the interface to receive ESMC messages.

When an interface's signal transitions out of the signal fail state, it must be fault-free for the global-wait-to-restore time before it is again considered by the clock selection process.

When the ESMC clock's EEC quality level (QL) mode is enabled, it sends a signal failure to the clock selection process during the global wait-to-restore time. After the global wait-to-restore time ends, a new quality level value is sent to the clock selection process.

To override the global wait-to-restore time on a specific interface, include the "wait-to-restore" on page 31017 statement at the `[edit chassis source interfaces (external-a | external-b | interface interface-name)]` hierarchy level.

Options

minutes Set the time for the port signal to be up before the port is opened to receive and transmit ESMC messages.

- **Range:** 0 through 12 minutes
- **Default:** 5 minutes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

synchronization (ACX Series) | 28787
synchronization (MX Series) | 28795
synchronization (PTX Series) | 28802
wait-to-restore | 31017

global-weight

IN THIS SECTION

- Syntax | 7743
- Hierarchy Level | 7743
- Description | 7743
- Options | 7744
- Required Privilege Level | 7744
- Release Information | 7744

Syntax

```
global-weight number;
```

Hierarchy Level

```
[edit chassis cluster redundancy-groupgroup-numberip-monitoring]
```

Description

Define global weight for IP monitoring. This is the weight that is subtracted from the redundancy group weight for all of the hosts being monitored. This number specifies the relative importance of IP address monitored objects in the operation of the redundancy group.

Every monitored IP address is assigned a weight. If the monitored address becomes unreachable, the weight of the object is deducted from the global-threshold of IP monitoring objects in its redundancy group. When the global-threshold reaches 0, the global-weight is deducted from the redundancy group. Every redundancy group has a default threshold of 255. If the threshold reaches 0, a failover is triggered. Failover is triggered even if the redundancy group is in manual failover mode and preemption is not enabled.

Options

number —Combined weight assigned to all monitored IP addresses. A higher weight value indicates a greater importance.

- **Range:** 0 through 255
- **Default:** 255

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

ip-monitoring | 9244

global-load-balancing

IN THIS SECTION

Syntax | 7745
Hierarchy Level | 7745
Description | 7745

- Options | 7746
- Required Privilege Level | 7746
- Release Information | 7746

Syntax

```
global-load-balancing {
    (helper-only | load-balancer-only);
}
```

Hierarchy Level

```
[edit fabric protocols bgp],
[edit logical-systems name protocols bgp],
[edit logical-systems name routing-instances name protocols bgp],
[edit logical-systems name tenants name routing-instances name protocols bgp],
[edit protocols bgp],
[edit routing-instances name protocols bgp],
[edit tenants name routing-instances name protocols bgp]
```

Description

Enable global load balancing (GLB). Global load balancing is hashing a route with multiple ECMP links onto several links for load balancing. BGP sends next-next-hop node (NNHN) capability for the route it advertises and monitors the link qualities of all local links with EBGP sessions and floods them to all direct neighbors. On receiving link qualities from neighboring nodes, BGP uses the combined link quality of the next-hops and next-next-hops to make load balancing decisions.

NOTE: You can optionally configure `no-global-load-balancing` on some of the BGP groups or peers to disable GLB on a subset of BGP neighbors. If GLB is disabled on a peer, BGP

will neither monitor the link quality, nor send link quality messages over that link. Also, BGP will not send or receive next-next-hop over that peer.

GLB is disabled by default. When you do not configure the options, the node works as both load-balancer and helper.

Options

helper-only (Optional) The node works as a helper and monitors the link quality.

load-balancer-only (Optional) The node works as a load balancer. GLB only receives link qualities from neighboring nodes and uses the combined link quality of next-hops and next-next-hops to make load balancing decisions.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

Global Load Balancing (GLB)

gnf

IN THIS SECTION

- Syntax | 7747
- Hierarchy Level | 7747
- Description | 7747

- Options | 7747
- Required Privilege Level | 7748
- Release Information | 7748

Syntax

```
gnf id {
    control-plane-bandwidth-percent percent;
    description description;
    fpcs fpcs;
    af-name ;
    collapsed-forward
}
```

Hierarchy Level

```
[edit chassis network-slices guest-network-functions]
```

Description

Define a GNF by assigning an ID to it.

Junos node slicing enables a single MX Series router to be partitioned to appear as multiple, independent routers. Each partition has its own Junos control plane, which runs as a virtual machine (VM), and a dedicated set of line cards. Each partition is called a guest network function (GNF).

Options

id GNF ID

- **Range:** 1–10

The remaining statements are explained separately.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Setting Up Junos Node Slicing

network-slices | 11322

guest-network-functions | 7949

gpid

IN THIS SECTION

- Syntax | 7748
- Hierarchy Level | 7749
- Description | 7749
- Default | 7749
- Required Privilege Level | 7749
- Release Information | 7749

Syntax

```
gpid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 | pos-scrambling-
crc-32 | pos-no-scrambling-crc-32 | ppp);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name lsp-attributes],
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
```

Description

Specify the type of payload carried by the LSP. It can be any of the following:

- `ethernet`—Ethernet (GPID value: 33)
- `hdlc`—High-level Data Link Control (HDLC) (GPID value: 44)
- `ipv4`—IP version 4 (GPID value: 0x0800)
- `pos-no-scrambling-crc-16`—for interoperability with other vendors' equipment (GPID value: 29)
- `pos-no-scrambling-crc-32`—for interoperability with other vendors' equipment (GPID value: 30)
- `pos-scrambling-crc-16`—for interoperability with other vendors' equipment (GPID value: 31)
- `pos-scrambling-crc-32`—for interoperability with other vendors' equipment (GPID value: 32)
- `ppp`—Point-to-Point Protocol (PPP) (GPID value: 50)

Default

`ipv4`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`pos-scrambling-crc-16`, `pos-no-scrambling-crc-16`, `pos-scrambling-crc-32`, and `pos-no-scrambling-crc-32` options added in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring the GPID

gprs

IN THIS SECTION

- Syntax | 7750
- Hierarchy Level | 7754
- Description | 7754
- Options | 7755
- Required Privilege Level | 7755
- Release Information | 7755

Syntax

```
gprs {
    gtp {
        handover-default {
            deny;
        }
        ip-group name {
            address-book name {
                address-set set-name;
            }
        }
        profile (Security GTP) profile-name {
            apn pattern-string {
                imsi-prefix imsi-prefix-digits {
                    action (APN GTP) {
                        drop;
                        pass;
                        selection (ms|net|vrf);
                    }
                }
```

```
            }
            mask-uli (1-65535);
            drop (Security GTP) {
                aa-create-pdp 0;
                aa-delete-pdp 0;
                bearer-resource 2;
                change-notification 2;
                config-transfer 2;
                context 2;
                create-bearer 2;
                create-data-forwarding 2;
                create-pdp (0 | 1 | all);
                create-session 2;
                create-tnl-forwarding 2;
                cs-paging 2;
                data-record (0 | 1 | all);
                delete-bearer 2;
                delete-command 2;
                delete-data-forwarding 2;
                delete-pdn 2;
                delete-pdp (0 | 1 | all);
                delete-session 2;
                detach 2;
                downlink-notification 2;
                echo (0 | 1 | 2 | all);
                error-indication (0 | 1 | all);
                failure-report (0 | 1 | all);
                fwd-access 2;
                fwd-relocation (1 | 2 | all);
                fwd-srns-context 1;
                g-pdu (0 | 1 | all);
                identification (0 | 1 | 2 | all);
                mbms-session-start (1 | 2 | all);
                mbms-session-stop (1 | 2 | all);
                mbms-session-update (1 | 2 | all);
                modify-bearer 2;
                modify-command 2;
                node-alive (0 | 1 | all);
                note-ms-present (0 | 1 | all);
                pdu-notification (0 | 1 | all);
                ran-info (1 | 2 | all);
                redirection (0 | 1 | all);
                release-access 2;
```

```
                relocation-cancel (1 | 2 | all);
                resume 2;
                send-route (0 | 1 | all);
                sgsn-context (0 | 1 | all);
                stop-paging 2;
                supported-extension 1;
                suspend 2;
                trace-session 2;
                update-bearer 2;
                update-pdn 2;
                update-pdp (0 | 1 | all);
                ver-not-supported (0 | 1 | 2 | all);
            }
            end-user-address-validated;
            gtp-in-gtp-denied;
            handover-group group-name;
            handover-on-roaming-intf;
            log (Security GTP) {
                forwarded (basic | detail);
                gtp-u name;
                prohibited (basic | detail);
                rate-limited (basic | detail);
            state-invalid (basic | detail);
            max-message-length max-message-length;
            min-message-length min-message-length;
            ne-group group-name;
            path-rate-limit {
                message-type (create-req | delete-req | echo-req | other) {
                    alarm-threshold (Security GPRS) {
                        forward forward;
                        reverse reverse;
                    }
                    drop-threshold (Security GPRS) {
                        forward forward;
                        reverse reverse;
                    }
                }
            }
            rate-limit (Security GTP) limit;
            remove-ie {
                version v1 {
                number ie-number;
                release (R6 | R7 | R8 | R9);
```

```
                }
            }
            req-timeout second;
            restart-path (all | create | echo);
            timeout (Security GTP) hour;
            u-tunnel-validated;
            ue-group group-name;
        }
        traceoptions (Security GTP) {
            file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
            flag name;
            no-remote-trace;
            trace-level {
                (error | info | notice | verbose | warning);
            }
        }
    }
    sctp {
        log (Security SCTP) name;
        multichunk-inspection disable;
        nullpdu {
            protocol (ID-0x0000 | ID-0xFFFF);
        }
        profile (Security SCTP) name {
            association-timeout association-timeout;
            drop (Security SCTP) {
                m3ua-service name;
                payload-protocol name;
            }
            handshake-timeout handshake-timeout;
            limit (Security SCTP) {
                address name {
                    payload-protocol (asap | bicc | ddp-segment | ddp-stream | diameter-dtls |
diameter-sctp | dua | enrp | h248 | h323 | id | iua | m2pa | m2ua | m3ua | others | qipc |
reserved | s1ap | simco | sua | tali | v5ua | x2ap) {
                        rate rate;
                    }
                }
                payload-protocol (asap | bicc | ddp-segment | ddp-stream | diameter-dtls |
diameter-sctp | dua | enrp | h248 | h323 | id | iua | m2pa | m2ua | m3ua | others | qipc |
reserved | s1ap | simco | sua | tali | v5ua | x2ap) {
                    rate rate;
```

```
                }
                rate {
                    address name {
                        sccp sccp;
                        ssp ssp;
                        sst sst;
                    }
                    sccp sccp;
                    ssp ssp;
                    sst sst;
                }
            }
            nat-only;
            permit (Security SCTP) {
                payload-protocol name;
            }
        }
        traceoptions (Security SCTP) {
            file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
            flag name;
            no-remote-trace;
        }
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Use the gprs function to configure all the General Packet Radio Service (GPRS) features.

Use `[edit security gprs]` hierarchy-level only for Junos OS Release 20.3 or previous releases.

From Junos OS Release 20.4R1 onwards, we've removed the term gprs from all the configuration statements and operational commands for GTP and SCTP. As a part of this change:

- We've deprecated the `[edit security gprs]` hierarchy-level and all the configuration options under this hierarchy-level.

- All the configuration statements previously available under the `[edit security gprs gtp]` hierarchy-level are now available under the `[edit security gtp]` hierarchy-level.
- All the configuration statements previously available under the `[edit security gprs sctp]` hierarchy-level are now available under the `[edit security sctp]` hierarchy-level.
- For default applications like junos-gprs-gtp and junos-gprs-sctp, you need not remove the term gprs.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. Statement modified in Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

GPRS Overview

GTPv1 Message Filtering

SCTP Overview

gprs-gtp-profile

IN THIS SECTION

- Syntax | 7756
- Hierarchy Level | 7756
- Description | 7756
- Required Privilege Level | 7756

Release Information | 7756

Syntax

```
gprs-gtp-profile gprs-gtp-profile;
```

Hierarchy Level

```
[edit logical-systems name security policies from-zone],
[edit logical-systems name security policies global policy name then],
[edit logical-systems name security policies stateful-firewall-rule name policy name then],
[edit logical-systems name tenants name security policies from-zone],
[edit logical-systems name tenants name security policies global policy name then],
[edit logical-systems name tenants name security policies stateful-firewall-rule name policy
name then],
[edit security policies from-zone],
[edit security policies global policy name then],
[edit security policies stateful-firewall-rule name policy name then],
[edit tenants name security policies from-zone],
[edit tenants name security policies global policy name then],
[edit tenants name security policies stateful-firewall-rule name policy name then]
```

Description

Use the gprs-gtp-profile option to specify the name of the GPRS tunneling protocol profile.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Enabling GTP Inspection in Policies

Understanding GTP Inspection Objects

Example: Enabling GTPv2 Inspection in Policies

gprs-sctp-profile

IN THIS SECTION

- Syntax | 7757
- Hierarchy Level | 7757
- Description | 7758
- Required Privilege Level | 7758
- Release Information | 7758

Syntax

```
gprs-sctp-profile gprs-sctp-profile;
```

Hierarchy Level

```
[edit logical-systems name security policies from-zone],
[edit logical-systems name security policies global policy name then],
[edit logical-systems name security policies stateful-firewall-rule name policy name then],
[edit logical-systems name tenants name security policies from-zone],
[edit logical-systems name tenants name security policies global policy name then],
[edit logical-systems name tenants name security policies stateful-firewall-rule name policy
name then],
[edit security policies from-zone],
[edit security policies global policy name then],
[edit security policies stateful-firewall-rule name policy name then],
[edit tenants name security policies from-zone],
```

```
[edit tenants name security policies global policy name then],
[edit tenants name security policies stateful-firewall-rule name policy name then]
```

Description

Use the gprs-sctp-profile option to specify the name of the GPRS stream control protocol profile.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

SCTP Packet Structure Overview

graceful-deletion-timeout

IN THIS SECTION

- Syntax | 7759
- Hierarchy Level | 7759
- Description | 7759
- Options | 7759
- Required Privilege Level | 7759
- Release Information | 7759

Syntax

```
graceful-deletion-timeout seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Specify the time, in seconds, before completing graceful deletion of signaling.

Options

seconds—Time before completing graceful deletion of signaling.

- **Range:** 1 through 300 seconds
- **Default:** 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Graceful Deletion Timeout Interval

graceful-restart (Enabling Globally)

IN THIS SECTION

- Syntax | 7760
- Hierarchy Level | 7760
- Description | 7761
- Default | 7761
- Options | 7761
- Required Privilege Level | 7761
- Release Information | 7761

Syntax

```
graceful-restart {
    disable;
    helper-disable;
    maximum-helper-recovery-time seconds;
    maximum-helper-restart-time seconds;
    notify-duration seconds;
    recovery-time seconds;
    restart-duration seconds;
    stale-routes-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit routing-options],
[edit routing-instances routing-instance-name routing-options]
```

Description

You configure the graceful restart routing option globally to enable the feature, but not to enable graceful restart for all routing protocols in a routing instance. To enable graceful restart globally, include the graceful-restart statement under the `[edit routing options]` hierarchy level. This enables graceful restart globally for all routing protocols. You can, optionally, modify the global settings at the individual protocol level.

NOTE:

- For VPNs, the `graceful-restart` statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.
- For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.
- LDP sessions flap when `graceful-restart` configurations change.

Default

Graceful restart is disabled by default.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Enabling Graceful Restart
- Configuring Routing Protocols Graceful Restart

Configuring Graceful Restart for MPLS-Related Protocols
Configuring VPN Graceful Restart
Configuring Logical System Graceful Restart
Configuring Graceful Restart for QFabric Systems

graceful-restart (Fabric Control)

IN THIS SECTION

- Syntax | 7762
- Hierarchy Level | 7762
- Description | 7762
- Required Privilege Level | 7763
- Release Information | 7763

Syntax

```
graceful-restart {
    restart-time seconds;
    stale-routes-time seconds;
}
```

Hierarchy Level

```
[edit fabric protocols fabric-control]
```

Description

Configure graceful restart parameters for the fabric control in a QFabric system.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X52-D10.

RELATED DOCUMENTATION

Understanding Routing Engines in the QFabric System

graceful-restart (Long-Lived for BGP Helper)

IN THIS SECTION

- Syntax | 7763
- Hierarchy Level | 7764
- Description | 7764
- Required Privilege Level | 7765
- Release Information | 7765

Syntax

```
graceful-restart {
    long-lived {
        receiver {
            enable:
            disable;
        }
        advertise-to-non-llgr-neighbor {
            omit-no-export;
```

```
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address]
```

Description

Configure the graceful restart capability for long-lived BGP sessions to enable BGP routing details to be retained for a longer period. It in important to retain BGP data for a longer period when the BGP control plane fails for some reason for slowly-restarting routers for a longer duration. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default.

When nonstop routing (NSR) and long-lived graceful restart (LLGR) are configured together, the router negotiates the LLGR capability in the usual, regular manner, including a long-lived stale time to trigger LLGR receiver mode in its peers. However, full LLGR restarter functionality (delaying the transmission of End of RIB (EoR) markers until EoRs are received from all peers) does not function under NSR. During a full system (both Routing Engines) restart, the routing protocol daemon (rpd) does not wait for EoRs from other peers before sending its own EoR. It transmits the EoR as soon as it has transmitted the current RIB contents. This condition can cause transient outages when the network reconverges. NSR is considered to be adequate to handle all single-Routing Engine restart scenarios. The restarter mode restriction effects only scenarios where both Routing Engines (or both copies of rpd) restart simultaneously. Ordinary restarter mode configuration is not enabled with NSR. Long-lived graceful restart receiver mode is enabled by default, unless ordinary graceful restart receiver mode is disabled.

> **NOTE**: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.

Configure graceful restart globally at the `[edit routing-options]` or `[edit routing-instances instance-name routing-options]` hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

graceful-restart (Long-Lived for BGP Restarter)

IN THIS SECTION

- Syntax | 7766
- Hierarchy Level | 7766
- Description | 7767
- Required Privilege Level | 7767
- Release Information | 7767

Syntax

```
graceful-restart {
    disable-notification-flag;
    disable-notification-extensions {
        omit-no-export;
    }
    forwarding-state-bit (from-fib | set); /* Configurable to be common for all address families
*/
    forwarding-state-bit (as-rr-client | from-fib); /* Configurable for each address family */
    long-lived {
        restarter {
            disable;
            stale-time interval;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family (l2vpn | route-target | inet)
(labeled-unicast | flow)],
[edit logical-systems logical-system-name protocols bgp group group-name family (l2vpn | route-
target | inet) (labeled-unicast | flow)],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
(l2vpn | route-target | inet) (labeled-unicast | flow)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family (l2vpn | route-target | inet) (labeled-unicast | flow)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family (l2vpn | route-target | inet) (labeled-unicast | flow)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family (l2vpn | route-target | inet) (labeled-unicast | flow)],
[edit routing-instances routing-instance-name protocols bgp family (l2vpn | route-target | inet)
(labeled-unicast | flow)],
[edit routing-instances routing-instance-name protocols bgp group group-name family (l2vpn |
route-target | inet) (labeled-unicast | flow)],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (l2vpn | route-target | inet) (labeled-unicast | flow)],
[edit routing-instances routing-instance-name protocols bgp family (l2vpn | route-target | inet)
(labeled-unicast | flow)],
```

```
[edit routing-instances routing-instance-name protocols bgp group group-name family (l2vpn |
route-target | inet) (labeled-unicast | flow)],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (l2vpn | route-target | inet) (labeled-unicast | flow)]
```

NOTE: Each routing table is identified by the protocol family or address family indicator (AFI) and a subsequent address family identifier (SAFI). The AFI parameter can be one of the `(l2vpn | inet | route-target)` protocols and the SAFI parameter can be either of the `(flow | labeled-unicast)` protocols for inet family and one of the `(auto-discovery-mspw | auto-discovery-only | signaling)` protcols for L2VPN family..

Configuring LLGR does not require that BGP graceful restart also be configured. The long-lived-graceful-restart section is visible only for families l2vpn, inet labeled-unicast, inet flow and route-target. It is prohibited for inet-mvpn, inet6-mvpn and inet-mdt. It is hidden for other families.

Description

Configure the graceful restart capability for long-lived BGP sessions on the restarting router to enable BGP routing details to be retained for a longer period. It in important to retain BGP data for a longer period when the BGP control plane fails for some reason for slowly-restarting routers for a longer duration. You can define the time period for which the stale routes need to be maintained.

You can also configure the BGP long-lived graceful restarter mode negotiation mechanism for a particular address family instead of configuring this capability for all address families in a system, logical system, or routing instance.

When LLGR restarter is enabled or disabled for a family or the stale- time is changed, the session is reset so that the new capability can be sent to the neighbor.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

graceful-restart (Multicast Snooping)

IN THIS SECTION

- Syntax | 7768
- Hierarchy Level | 7768
- Description | 7768
- Default | 7769
- Required Privilege Level | 7769
- Release Information | 7769

Syntax

```
graceful-restart {
    disable;
    restart-duration seconds;
}
```

Hierarchy Level

```
[edit multicast-snooping-options]
```

Description

Establish the graceful restart duration for multicast snooping. You can set this value between 0 and 300 seconds. If you set the duration to 0, graceful restart is effectively disabled. Set this value slightly larger than the IGMP query response interval.

Default

180 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Example: Configuring Multicast Snooping

query-response-interval (Bridge Domains) | **13938**

graceful-restart (Protocols BGP)

IN THIS SECTION

- Syntax | **7769**
- Hierarchy Level | **7770**
- Description | **7770**
- Required Privilege Level | **7771**
- Release Information | **7771**

Syntax

```
graceful-restart {
    disable;
    restart-time seconds;
```

```
    stale-routes-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address]
```

Description

Configure graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default. However, helper mode, the ability to assist a neighboring router attempting a graceful restart, is enabled by default.

To configure the duration of the BGP graceful restart period, include the `restart-time` statement at the `[edit protocols bgp graceful-restart]` hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the `stale-routes-time` statement at the `[edit protocols bgp graceful-restart]` hierarchy level.

> **NOTE**: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.

Enable graceful restart mode for BGP (and other protocols) by configuring graceful-restart at the routing-options level. Note that you cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally.

For example, this configuration is required to enable graceful restart:

```
routing-options {
    graceful-restart
}
```

If you want to disable graceful restart for some protocols, you can do this at the protocol's graceful-restart command. The following configuration along with the configuration above will keep graceful restart for all protocols but BGP.

```
protocols{
    bgp{
    graceful-restart;  {
    disable;
}
```

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring Graceful Restart Options for BGP
- *Configuring Graceful Restart for QFabric Systems*
- Junos OS High Availability User Guide

graceful-restart (Protocols IS-IS)

IN THIS SECTION

- Syntax | 7772
- Hierarchy Level | 7772
- Description | 7772

- Options | 7772
- Required Privilege Level | 7773
- Release Information | 7773

Syntax

```
graceful-restart {
    disable;
    helper-disable;
    restart-duration seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit protocols isis]
```

Description

Configure graceful restart parameters for IS-IS.

Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the `[edit routing-options]` hierarchy level. When graceful restart is enabled, the restarting routing device is not removed from the network topology during the restart period. The adjacencies are reestablished after restart is complete.

On LAN interfaces where IS-IS is configured on a transit router that serves as the designated router (DR), a graceful restart causes:

- The ingress router of the label-switched path (LSP), which passes through the DR, to break the LSP.
- The ingress router to re-signal the LSP.

Options

`disable`—Disable graceful restart for IS-IS.

`helper-disable`—Disable graceful restart helper capability. Helper mode is enabled by default.

`restart-duration` *seconds*—Time period for the restart to last, in seconds.

- **Range:** 30 through 300 seconds
- **Default:** 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routing Protocols Graceful Restart

graceful-restart (Protocols LDP)

IN THIS SECTION

- Syntax | 7774
- Hierarchy Level | 7774
- Description | 7774
- Required Privilege Level | 7774
- Release Information | 7774

Syntax

```
graceful-restart {
    disable;
    helper-disable;
    maximum-neighbor-recovery-time value;
    reconnect-time seconds;
    recovery-time value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Configure LDP graceful restart on the LDP primary protocol instance or for a specific routing instance.

> **NOTE**: When you alter the graceful restart configuration at either the `[edit routing-options graceful-restart]` or `[edit protocols ldp graceful-restart]` hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring LDP Graceful Restart

graceful-restart (Protocols OSPF)

IN THIS SECTION

- Syntax | 7775
- Hierarchy Level | 7775
- Description | 7776
- Options | 7776
- Required Privilege Level | 7777
- Release Information | 7777

Syntax

```
graceful-restart {
    disable;
    helper-disable (standard | restart-signaling | both);
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf]
```

Description

Configure graceful restart for OSPF.

Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the `[edit routing-options]` hierarchy level.

Options

disable

Disable graceful restart for OSPF.

helper-disable (standard | restart-signaling| both)

Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The last committed statement takes precedence over the previously configured statement.

- `standard` disables helper mode for standard graceful restart (based on RFC 3623).
- `restart-signaling` disables helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813).
- `both` disables helper mode for both standard and restart signaling-based graceful restart.

Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.

no-strict-lsa-checking

Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.

> **NOTE**: The `helper-disable` statement and the `no-strict-lsa-checking` statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the `show protocols (ospf | ospf3)` command.

notify-duration *seconds*

Estimated time needed to send out purged grace LSAs over all the interfaces. Range is 1 through 3600 seconds, and the default is 30 seconds.

restart-duration *seconds*

Estimated time needed to reacquire a full OSPF neighbor from each area. Range is 1 through 3600 seconds, and the default is 180 seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `no-strict-lsa-checking` statement introduced in Junos OS Release 8.5.

Support for the helper mode `standard`, `restart-signaling`, and `both` options introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring Graceful Restart for OSPF

Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart

Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart

Example: Disabling Strict LSA Checking for OSPF Graceful Restart

graceful-restart (Protocols PIM)

IN THIS SECTION

- Syntax | 7778
- Hierarchy Level | 7778
- Description | 7778
- Required Privilege Level | 7778
- Release Information | 7778

Syntax

```
graceful-restart {
    disable;
    no-bidirectional-mode;
    restart-duration seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Configure PIM sparse mode graceful restart.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring PIM Sparse Mode Graceful Restart

graceful-restart (Protocols RIP)

IN THIS SECTION

- Syntax | 7779
- Hierarchy Level | 7779
- Description | 7779
- Options | 7779
- Required Privilege Level | 7780
- Release Information | 7780

Syntax

```
graceful-restart {
    disable;
    restart-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit protocols rip]
```

Description

Configure graceful restart for RIP.

Options

disable—Disables graceful restart for RIP.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS High Availability User Guide

graceful-restart (Protocols RIPng)

IN THIS SECTION

- Syntax | 7780
- Hierarchy Level | 7781
- Description | 7781
- Options | 7781
- Required Privilege Level | 7781
- Release Information | 7781

Syntax

```
graceful-restart {
    disable;
    restart-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng],
[edit protocols ripng],
[edit routing-instances routing-instance-name protocols ripng]
```

Description

Configure graceful restart for RIPng.

Options

disable—Disables graceful restart for RIPng.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Junos OS High Availability User Guide

graceful-restart (Protocols RSVP)

IN THIS SECTION

- Syntax | 7782
- Hierarchy Level | 7782
- Description | 7782
- Options | 7782
- Required Privilege Level | 7783
- Release Information | 7783

Syntax

```
graceful-restart {
    disable;
    helper-disable;
    maximum-helper-recovery-time seconds;
    maximum-helper-restart-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit protocols rsvp],
[edit routing-options]
```

Description

Configure graceful restart on the router. You must configure the `graceful-restart` statement at the `[edit routing-options]` hierarchy level to enable graceful restart on the router.

Options

`disable`—Disable graceful restart on the router or for RSVP.

`helper-disable`—Disable RSVP graceful restart helper mode (this option is only available at the `[edit protocols rsvp]` hierarchy level).

- **Default:** Helper mode is enabled by default.

`maximum-helper-recovery-time` *seconds*—The maximum length of time the router stores the state of neighboring routers when they undergo a graceful restart. The value applies to all neighboring routers, so it should be based on the time that the slowest RSVP neighbor requires for restart.

- **Default:** 180 seconds
- **Range:** 1 through 3600 seconds

`maximum-helper-restart-time` *seconds*—The maximum length of time the router waits between when it discovers that a neighboring router has gone down and when it declares the neighbor down. This value is applied to all neighboring routers, so it should be based on the time that the slowest RSVP neighbor requires for restart.

- **Default:** 20 seconds
- **Range:** 1 through 1800 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring RSVP Graceful Restart

graceful-shutdown (Protocols BGP)

IN THIS SECTION

- Syntax | 7784
- Hierarchy Level | 7784
- Description | 7785
- Default | 7785
- Required Privilege Level | 7785
- Release Information | 7785

Syntax

```
graceful-shutdown {
    receiver {
        disable;
        local-preference <value>;
    }
    sender {
        local-preference <value>;
    }
}
```

Hierarchy Level

```
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address]
```

Description

Configure graceful shutdown feature for BGP. Graceful shutdown migrates traffic from one BGP next-hop to another without interrupting the traffic flow. This feature can be enabled for both IBGP and EBGP. The router on which the graceful shutdown is performed is known as the sender or initiator. The receiver is the router that has a BGP session with the sender router. A device acting as a sender is by default a receiver. In other words, with the sender feature enabled, the receiver feature cannot be disabled on the same device.

NOTE: The BGP neighbor device must support graceful shutdown feature, without any additional configuration.

NOTE: On completion of the maintenance window, disabling graceful shutdown is recommended.

Default

The graceful shutdown receiver feature is enabled by default.

The graceful shutdown sender feature is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

shutdown (Protocols BGP) | **27936**

graceful-switchover (Chassis)

IN THIS SECTION

- Syntax | 7786
- Hierarchy Level | 7786
- Description | 7786
- Required Privilege Level | 7786
- Release Information | 7787

Syntax

```
graceful-switchover;
```

Hierarchy Level

```
[edit chassis redundancy]
```

Description

For devices with more than one Routing Engine, including those in a Virtual Chassis or a Virtual Chassis Fabric, configure the primary Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.

NOTE: The `graceful-switchover` statement at the `[edit chassis redundancy]` hierarchy level is not supported for Junos OS Evolved. Graceful switchover is enabled on the Junos OS Evolved system by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding Graceful Routing Switchover

Configuring Graceful Routing Engine Switchover

gratuitous-arp-count

IN THIS SECTION

- Syntax | 7787
- Hierarchy Level | 7787
- Description | 7788
- Options | 7788
- Required Privilege Level | 7788
- Release Information | 7788

Syntax

```
gratuitous-arp-count number;
```

Hierarchy Level

```
[edit chassis cluster redundancy-group group-number]
```

Description

Specify the number of gratuitous Address Resolution Protocol (ARP) requests to send on an active interface after failover.

You can configure this option to specify the number of gratuitous ARP requests that an interface can send to notify other network devices of its presence after the redundancy group it belongs to has failed over.

By default, the SRX Series Firewall sends four GARPs per reth on a failover. You can modify the number of GARPs sent per-redundancy-group basis.

Options

number—Number of gratuitous ARP requests that a newly elected primary device in a chassis cluster sends out to announce its presence to the other network devices.

- **Range:** 1 through 16
- **Default:** 4

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

redundancy-group (Chassis Cluster) | **14173**

gratuitous-arp-reply

IN THIS SECTION

- Syntax | 7789
- Hierarchy Level | 7789
- Description | 7789
- Default | 7789
- Options | 7789
- Required Privilege Level | 7790
- Release Information | 7790

Syntax

```
(gratuitous-arp-reply | no-gratuitous-arp-reply);
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-range interface-range-name]
```

Description

Enable processing of ARP updates received via gratuitous ARP reply messages.

Default

Updating of the ARP cache is disabled on all Ethernet interfaces.

Options

gratuitous-arp-reply—Update the ARP cache.

no-gratuitous-arp-reply—Do not update the ARP cache.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

gre

IN THIS SECTION

- Syntax | 7790
- Hierarchy Level | 7791
- Description | 7791
- Required Privilege Level | 7791
- Release Information | 7791

Syntax

```
gre {
    destination {
        address address;
    }
    key key; key;
    source {
        address address;
    }
    tunnel-routing-instance {
        routing-instance routing-instance;
    }
    tunnel-termination;
}
```

Hierarchy Level

```
[edit interfaces fti number unit logical-unit-number tunnel encapsulation]
```

Description

Configure GRE (Generic Routing Encapsulation) encapsulation type.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Evolved Release 20.2R1.

gre (lds-option)

IN THIS SECTION

- Syntax | 7791
- Hierarchy Level | 7792
- Description | 7792
- Options | 7792
- Required Privilege Level | 7792
- Release Information | 7792

Syntax

```
gre {
    gre-4in4;
    gre-4in6;
    gre-6in4;
```

```
        gre-6in6;
    }
```

Hierarchy Level

```
[edit security screen ids-option ids-option-name ip tunnel]
```

Description

Configure IP tunnel GRE IDS option.

Options

gre-4in4 Enable IP tunnel GRE 4in4 IDS option.

gre-4in6 Enable IP tunnel GRE 4in6 IDS option.

gre-6in4 Enable IP tunnel GRE 6in4 IDS option.

gre-6in6 Enable IP tunnel GRE 6in6 IDS option.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D10.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

gre-in

IN THIS SECTION

- Syntax | 7793
- Hierarchy Level | 7793
- Description | 7793
- Options | 7793
- Required Privilege Level | 7794
- Release Information | 7794

Syntax

```
gre-in {
    mss value;
}
```

Hierarchy Level

```
[edit security flow tcp-mss]
```

Description

Enable and specify the TCP maximum segment size (TCP MSS) for Generic Routing Encapsulation (GRE) packets that are coming out from an IPsec VPN tunnel. If the device receives a GRE-encapsulated TCP packet with the SYN bit and TCP MSS option set and the TCP MSS option specified in the packet exceeds the TCP MSS specified by the device, the device modifies the TCP MSS value accordingly. By default, a TCP MSS for GRE packets is not set.

Options

`mss` *`value`* —TCP MSS for GRE packets. Value is optional.

- **Range:** 64 through 63535 bytes

- **Default:** 1320 bytes, if no value is specified

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

gre-out

IN THIS SECTION

- Syntax | 7794
- Hierarchy Level | 7795
- Description | 7795
- Options | 7795
- Required Privilege Level | 7795
- Release Information | 7795

Syntax

```
gre-out {
    mss value;
}
```

Hierarchy Level

```
[edit security flow tcp-mss]
```

Description

Enable and specify the TCP maximum segment size (TCP MSS) for Generic Routing Encapsulation (GRE) packets that are going into an IPsec VPN tunnel. If the device receives a GRE-encapsulated TCP packet with the SYN bit and TCP MSS option set and the TCP MSS option specified in the packet exceeds the TCP MSS specified by the device, the device modifies the TCP MSS value accordingly. By default, a TCP MSS for GRE packets is not set.

Options

`mss` *`value`*—TCP MSS for GRE packets. Value is optional.

- **Range:** 64 through 65,535 bytes
- **Default:** 1320 bytes

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Understanding Traffic Processing on Security Devices*

gre-performance-acceleration (Security Flow)

IN THIS SECTION

- Syntax | 7796
- Hierarchy Level | 7796
- Description | 7796
- Required Privilege Level | 7796
- Release Information | 7796

Syntax

```
gre-performance-acceleration;
```

Hierarchy Level

```
[edit security flow]
```

Description

Enables GRE performance acceleration. This helps to accelerate GRE traffic performance.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1 for SRX Series and NFX Series.

RELATED DOCUMENTATION

show security flow status | 24348

gre (Routing Options)

IN THIS SECTION

- Syntax | 7797
- Hierarchy Level | 7797
- Description | 7797
- Options | 7798
- Required Privilege Level | 7798
- Release Information | 7798

Syntax

```
gre <source-address> <destination-networks>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name],
[edit routing-options dynamic-tunnels tunnel-name]
```

Description

Define a dynamic generic routing encapsulation (GRE) tunnel for IPv4 to automatically establish label-switched paths (LSPs) for any new provider edge (PE) router added to a full mesh of LSPs.

Options

destination-networks A list of one or more destination prefixes that are reachable through the GRE tunnel.

source-address The source address used for the dynamic GRE tunnel.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`next-hop-based-tunnel` option introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring MPLS-Signaled LSPs to Use GRE Tunnels

greater-than

IN THIS SECTION

- Syntax | 7799
- Hierarchy Level | 7799
- Description | 7799
- Options | 7799
- Required Privilege Level | 7799
- Release Information | 7799

Syntax

```
greater-than value;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format integer],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format time]
```

Description

Specify a value for the custom AVP attribute above which the incoming RADIUS request from the subscriber must match.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

Options

value	Value that the attribute must be greater than.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers*

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

gres-route-flush-delay (Subscriber Management)

IN THIS SECTION

- Syntax | 7800
- Hierarchy Level | 7800
- Description | 7800
- Required Privilege Level | 7800
- Release Information | 7801

Syntax

```
gres-route-flush-delay;
```

Hierarchy Level

```
[edit system services subscriber-management]
```

Description

For a subscriber network configured with either nonstop active routing (NSR) or graceful restart, configure the router to wait 180 seconds (3 minutes) before removing (flushing) static or dynamic access routes and access-internal routes from the forwarding table after a graceful Routing Engine switchover (GRES) has taken place.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

High Availability for Subscriber Access Networks

group

IN THIS SECTION

- Syntax | 7801
- Hierarchy Level | 7802
- Description | 7802
- Options | 7802
- Required Privilege Level | 7802
- Release Information | 7802

Syntax

```
group  group-name {
    debounce-interval seconds;
    link-to-monitor interface-name;
    link-to-disable interface-name;
}
```

Hierarchy Level

```
[edit protocols uplink-failure-detection]
```

Description

Configure a group of uplink and downlink interfaces for uplink failure detection.

Options

group-name Name of the uplink failure detection group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Overview of Uplink Failure Detection

Configuring Interfaces for Uplink Failure Detection

Example: Configuring Interfaces for Uplink Failure Detection

group-address (Routing Instances Tunnel Group)

IN THIS SECTION

Syntax | 7803

- Hierarchy Level | 7803
- Description | 7803
- Required Privilege Level | 7803
- Release Information | 7803

Syntax

```
group-address address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel family inet / inet6 pim-ssm],
[edit routing-instances routing-instance-name provider-tunnel family inet / inet6 pim-ssm]
```

Description

Configure the PIM-ASM (Rosen 6) or PIM-SSM (Rosen 7) provider tunnel group address. Each MDT is linked to a group address in the provider space.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

In Junos OS Release 17.3R1, the `pim-ssm` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default multicast distribution tree (MDT) in Rosen 7, and data MDT for Rosen 6 and Rosen 7.

RELATED DOCUMENTATION

Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs

group-address (Routing Instances VPN)

IN THIS SECTION

- Syntax | 7804
- Hierarchy Level | 7804
- Description | 7805
- Options | 7805
- Required Privilege Level | 7805
- Release Information | 7805

Syntax

```
group-address address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel pim-asm],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel pim-asm family inet],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel pim-asm family inet6],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel pim-ssm],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel pim-ssm family inet],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel pim-ssm family inet6],
[edit routing-instances routing-instance-name provider-tunnel pim-asm],
```

```
[edit routing-instances routing-instance-name provider-tunnel pim-asm family inet],
[edit routing-instances routing-instance-name provider-tunnel pim-asm family inet6],
[edit routing-instances routing-instance-name provider-tunnel pim-ssm],
[edit routing-instances routing-instance-name provider-tunnel pim-ssm family inet],
[edit routing-instances routing-instance-name provider-tunnel pim-ssm family inet6]
```

Description

Specify a group address on which to encapsulate multicast traffic from a virtual private network (VPN) instance.

NOTE: IPv6 provider tunnels are not currently supported for draft-rosen MVPNs. They are supported for MBGP MVPNs.

Options

address—For IPv4, IP address whose high-order bits are 1110, giving an address range from 224.0.0.0 through 239.255.255.255, or simply 224.0.0.0/4. For IPv6, IP address whose high-order bits are FF00 (FF00::/8).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Starting with Junos OS Release 11.4, to provide consistency with draft-rosen 7 and next-generation BGP-based multicast VPNs, configure the provider tunnels for draft-rosen 6 anysource multicast VPNs at the `[edit routing-instances routing-instance-name provider-tunnel]` hierarchy level. The `mdt`, `vpn-tunnel-source`, and `vpn-group-address` statements are deprecated at the `[edit routing-instances routing-instance-name protocols pim]` hierarchy level. Use `group-address` in place of `vpn-group-address`.

RELATED DOCUMENTATION

Example: Configuring Any-Source Multicast for Draft-Rosen VPNs

group (Bridge Domains)

IN THIS SECTION

- Syntax | 7806
- Hierarchy Level | 7806
- Description | 7807
- Options | 7807
- Required Privilege Level | 7807
- Release Information | 7807

Syntax

```
group ip-address {
    source ip-address;
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name static],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-name static],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name static],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols vlan vlan-id igmp-snooping interface interface-name static]
```

Description

Configure the IGMP multicast group address that receives data on an interface and (optionally) a source address for certain packets.

Options

ip-address—Group address.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping

group (Configuring Group Name)

IN THIS SECTION

- Syntax | 7808
- Hierarchy Level | 7808
- Description | 7808
- Options | 7808
- Required Privilege Level | 7808
- Release Information | 7809

Syntax

```
group group-name {
    (default-context-prefix | context-prefix context-prefiix){
        security-model (any | usm | v1 | v2c) {
            security-level (authentication | none | privacy) {
                notify-view view-name;
                read-view view-name;
                write-view view-name;
            }
        }
    }
}
```

Hierarchy Level

```
[edit snmp v3 vacm access]
```

Description

Assign the security name to a group, and specify the SNMPv3 context applicable to the group. The `default-context-prefix` statement, when included, adds all the contexts configured on the device to the group, whereas the `context-prefix` *`context-prefix`* statement enables you to specify a context and to add that particular context to the group.

(Not applicable to the QFX Series and OCX Series.) When the context prefix is specified as default (for example, `context-prefix default`), the context associated with the master routing instance is added to the group. To specify a routing instance that is part of a logical system, specify it as `logical system/routing instance`. For example, to specify routing instance ri1 in logical system ls1, include `context-prefix ls1/ri1`.

The remaining statements under this hierarchy are explained separately.

Options

`group-name`—SNMPv3 group name created for the SNMPv3 group.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Group

group-count (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 7809
- Hierarchy Level | 7809
- Description | 7810
- Options | 7810
- Required Privilege Level | 7810
- Release Information | 7810

Syntax

```
group-count number;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name static group multicast-group-address]
```

Description

Configure the number of static groups to be created over the dynamic interface.

Options

number—Number of static groups.

- **Default:** 1
- **Range:** 1 through 512

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Enabling MLD Static Group Membership

group-count (Protocols IGMP)

IN THIS SECTION

- Syntax | **7811**
- Hierarchy Level | **7811**
- Description | **7811**
- Options | **7811**
- Required Privilege Level | **7811**

Release Information | 7811

Syntax

```
group-count number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name static group
multicast-group-address],
[edit protocols igmp interface interface-name static group multicast-group-address]
```

Description

Specify the number of static groups to be created.

Options

number—Number of static groups.

- **Range:** 1 through 512

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Enabling IGMP Static Group Membership

group-count (Protocols MLD)

IN THIS SECTION

- Syntax | 7812
- Hierarchy Level | 7812
- Description | 7812
- Options | 7812
- Required Privilege Level | 7813
- Release Information | 7813

Syntax

```
group-count number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name static group
multicast-group-address],
[edit protocols mld interface interface-name static group multicast-group-address]
```

Description

Configure the number of static groups to be created.

Options

number—Number of static groups.

- **Default:** 1
- **Range:** 1 through 512

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Enabling MLD Static Group Membership

group (Defining Access Privileges for an SNMPv3 Group)

IN THIS SECTION

- Syntax | 7813
- Hierarchy Level | 7813
- Description | 7814
- Options | 7814
- Required Privilege Level | 7814
- Release Information | 7814

Syntax

```
group group-name;
```

Hierarchy Level

```
[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c) security-name security-name]
```

Description

Define access privileges granted to a group.

Options

group-name—Identifies a collection of SNMP security names that belong to the same access policy SNMP.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Group

group (DHCP Local Server)

IN THIS SECTION

- Syntax | 7815
- Hierarchy Level | 7819
- Description | 7819
- Options | 7819
- Required Privilege Level | 7819
- Release Information | 7819

Syntax

```
group group-name {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            relay-agent-interface-id
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary primary-profile-
name>;
    interface interface-name {
        access-profile profile-name;
        exclude;
        overrides {
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-discover-match <option60-and-option82>;
            client-negotiation-match incoming-interface;
            delay-advertise {
                based-on (option-15 | option-16 | option-18 | option-37) {
                    equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    not-equals {
                        ascii ascii-string;
```

```
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delay-offer {
            based-on (option-60 | option-77 | option-82) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        rapid-commit;
    }
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    trace;
    upto upto-interface-name;
}
interface-tag (DHCP Local Server);
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
```

```
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-discover-match <option60-and-option82>;
        client-negotiation-match incoming-interface;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
```

```
        delay-offer {
            based-on (option-60 | option-77 | option-82) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delegated-pool;
        delete-binding-on-renegotiation;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
        rapid-commit;
    }
    reconfigure {
        attempts attempt-count;
        clear-on-terminate;
        strict;
        timeout timeout-value;
        token token-value;
         trigger {
            radius-disconnect;
        }
    }
    route-suppression;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Configure a group of interfaces that have a common configuration, such as authentication parameters. A group must contain at least one interface.

Options

group-name	Name of the group.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`interface-tag` option introduced in Junos OS Release 23.2R1

RELATED DOCUMENTATION

- *Understanding Differences Between Legacy DHCP and Extended DHCP*
- *Grouping Interfaces with Common DHCP Configurations*
- *Specifying Authentication Support*
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*

DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

group (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 7820
- Hierarchy Level | 7824
- Description | 7825
- Options | 7825
- Required Privilege Level | 7825
- Release Information | 7825

Syntax

```
group group-name {
    access-profile profile-name;
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name interface-name;
            logical-system-name;
            mac-address mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
```

```
        }
    }
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    forward-only {
        logical-system <current | default | logical-system-name>;
        routing-instance <current | default | routing-instance-name>;
    }
    interface interface-name {
        access-profile profile-name;
        exclude;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                    session-mode (automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
            }
        }
        overrides {
            allow-no-end-option;
            allow-snooped-clients;
            always-write-giaddr;
            always-write-option-82;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-discover-match <option60-and-option82 | incoming-interface>;
            client-negotiation-match incoming-interface;
```

```
            delay-authentication;
            delete-binding-on-renegotiation;
            disable-relay;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            layer2-unicast-replies;
            no-allow-snooped-clients;
            no-bind-on-request;
            proxy-mode;
            relay-source
            replace-ip-source-with;
            send-release-on-delete;
            trust-option-82;
        }
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
        trace;
        upto upto-interface-name;
    }
    interface-tag (DHCP Relay Server)
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
```

```
        }
    }
    overrides {
        allow-snooped-clients;
        always-write-giaddr;
        always-write-option-82;
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-discover-match <option60-and-option82>;
        client-negotiation-match incoming-interface;
        disable-relay;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        layer2-unicast-replies;
        no-allow-snooped-clients;
        no-bind-on-request;
        proxy-mode;
        relay-source
        replace-ip-source-with;
        send-release-on-delete;
        trust-option-82;
    }
    relay-agent-interface-id {
        include-irb-and-l2;
        keep-incoming-interface-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    relay-agent-remote-id {
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82 <strict>;
        use-vlan-id;
    }
    relay-option {
        option-number option-number;
        default-action {
```

```
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
        equals (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
        starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
    }
    relay-option-82 {
        circuit-id {
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82;
        }
        remote-id {
            prefix prefix;
            use-interface-description (logical | device);
        }
        server-id-override
    }
    remote-id-mismatch disconnect;
    route-suppression;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify the name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration. A group must contain at least one interface. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Options

group-name Name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

`interface-tag` option introduced in Junos OS Release 23.2R1

RELATED DOCUMENTATION

- *Extended DHCP Relay Agent Overview*
- Configuring DHCP Relay Agent
- *Configuring Group-Specific DHCP Relay Options*
- *Grouping Interfaces with Common DHCP Configurations*
- *Specifying Authentication Support*
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*

DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

group (DHCP Security)

IN THIS SECTION

- Syntax | 7826
- Hierarchy Level | 7827
- Description | 7827
- Required Privilege Level | 7827
- Release Information | 7827

Syntax

```
group group-name {
    interface interface-name {
        static-ip ip-address {
            mac mac-address;
        }
        static-ipv6 ip-address {
            mac mac-address;
        }
    }
    overrides {
        no-dhcpv6-options;
        no-option16;
        no-option18;
        no-option37;
        no-option82;
        trusted;
        untrusted;
    }
}
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security]
```

Description

Specify the name of a group of access interfaces that you want to configure for DHCP security attributes that are different from the attributes set for other interfaces in the VLAN. A group must contain at least one interface.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10 .

Support for the `static-ipv6` and `no-option37` statements introduced in Junos OS Release 13.2X51-D20 for EX Series switches.

RELATED DOCUMENTATION

Configuring Static DHCP IP Addresses for DHCP snooping (ELS)

Enabling a Trusted DHCP Server (ELS)

Understanding DHCP Snooping (non-ELS)

group (DHCP Security for MX Series)

IN THIS SECTION

Syntax | 7828

- Hierarchy Level | 7828
- Description | 7828
- Required Privilege Level | 7828
- Release Information | 7829

Syntax

```
group group-name {
    interface interface-name {
        static-ip ip-address {
            mac mac-address;
        }
    }
    overrides {
        no-option82;
        trusted;
        untrusted;
    }
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]
```

Description

Specify the name of a group of access interfaces that you want to configure for DHCP security attributes that are different from the attributes set for other interfaces in the VLAN or bridge domain. A group must contain at least one interface.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10 .

Support for the `static-ipv6` and `no-option37` statements introduced in Junos OS Release 13.2X51-D20 for EX Series switches.

RELATED DOCUMENTATION

Configuring Static DHCP IP Addresses for DHCP snooping (MX routers)

Enabling a Trusted DHCP Server (MX Series Routers)

Understanding DHCP Snooping (non-ELS)

group (DHCP Spoofing Prevention)

IN THIS SECTION

- Syntax | 7829
- Hierarchy Level | 7830
- Description | 7830
- Required Privilege Level | 7830
- Release Information | 7830

Syntax

```
group group-name {
    interface interface-name;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name forwarding-options
dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay]
```

Description

Configure Dynamic Host Configuration Protocol (DHCP) snooping on the router. When acting as a snooping agent, the MX Series router typically is located between the client and the DHCP relay agent. It creates filters by “snooping” DHCP messages and binding DHCP-issued IP addresses with the MAC address of the client. These filters help prevent DHCP spoofing.

Configure DHCP snooping by including the appropriate interfaces under the `group` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4 (MX Series routers only).

RELATED DOCUMENTATION

Preventing DHCP Spoofing on MX Series 5G Universal Routing Platforms

group (Distributed IGMP)

IN THIS SECTION

- Syntax | 7831
- Hierarchy Level | 7831

- Description | 7831
- Options | 7831
- Required Privilege Level | 7831
- Release Information | 7832

Syntax

```
group multicast-group-address {
    <distributed>;
    source  source-address <distributed>;
}
```

Hierarchy Level

```
[edit protocols pim static]
```

Description

Specify the multicast group address for the multicast group that is statically configured on an interface.

Options

distributed (Optional) Preprovision a specific multicast group address (G).

multicast-group-address Specific multicast group address being statically configured on an interface.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X50.

RELATED DOCUMENTATION

Enabling Distributed IGMP

Junos OS Multicast Protocols User Guide

Junos OS Multicast Protocols User Guide

group (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | 7832
- Hierarchy Level | 7833
- Description | 7833
- Options | 7833
- Required Privilege Level | 7833
- Release Information | 7833

Syntax

For group configuration with a source, use the following syntax:

```
group ip-address {
    source ip-address;
}
```

For group configuration without a source, use the following syntax:

```
group group;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name static],
```

Description

When configuring with a source address, configure the IGMP multicast group address that receives data on an interface and a source address for certain packets. For configuration without a source address, configure only the IGMP multicast group address that receives data on an interface.

Options

ip-address—Group IP address.

group—Name of group.

NOTE: You must specify a unique address for each group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Enabling IGMP Static Group Membership

group (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 7834
- Hierarchy Level | 7834
- Description | 7834
- Options | 7835
- Required Privilege Level | 7835
- Release Information | 7835

Syntax

```
group multicast-group-address {
    exclude;
    group-count number;
    group-increment increment;
    source ip-address {
        source-count number;
        source-increment increment;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name static]
```

Description

The MLD multicast group address and (optionally) the source address for the multicast group being dynamically configured on an interface.

Options

multicast-group-address—Address of the group.

> **NOTE**: You must specify a unique address for each group.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Enabling MLD Static Group Membership

group (IGMP Snooping)

IN THIS SECTION

- Syntax | 7836
- Hierarchy Level | 7836
- Description | 7836
- Options | 7836
- Required Privilege Level | 7836
- Release Information | 7836

Syntax

```
group ip-address;
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan  (all | vlan-name) interface (all | interface-
name)                 static]
```

Description

Configure a static multicast group on an interface.

Options

ip-address—IP address of the multicast group receiving data on an interface.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

show igmp-snooping vlans | **20198**

group-increment (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 7837
- Hierarchy Level | 7837
- Description | 7837
- Options | 7837
- Required Privilege Level | 7838
- Release Information | 7838

Syntax

```
group-increment increment;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name static group multicast-group-address source]
```

Description

Configure the number of times the address should be incremented for each static group created on a dynamic interface. The increment is specified in a format similar to an IPv6 address.

Options

increment—Number of times the address should be incremented.

- **Default:** ::1
- **Range:** ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff:

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Enabling MLD Static Group Membership

group-increment (Protocols IGMP)

IN THIS SECTION

- Syntax | 7838
- Hierarchy Level | 7839
- Description | 7839
- Options | 7839
- Required Privilege Level | 7839
- Release Information | 7839

Syntax

```
group-increment increment;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name static group
multicast-group-address],
[edit protocols igmp interface interface-name static group multicast-group-address]
```

Description

Configure the number of times the address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.

Options

increment—Number of times the address should be incremented.

- **Default:** 0.0.0.1
- **Range:** 0.0.0.1 through 255.255.255.255

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Enabling IGMP Static Group Membership

group-increment (Protocols MLD)

IN THIS SECTION

- Syntax | 7840
- Hierarchy Level | 7840
- Description | 7840
- Options | 7840
- Required Privilege Level | 7841
- Release Information | 7841

Syntax

```
group-increment number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name static group
multicast-group-address],
[edit protocols mld interface interface-name static group multicast-group-address]
```

Description

Configure the number of times the address should be incremented for each static group created. The increment is specified in a format similar to an IPv6 address.

Options

increment—Number of times the address should be incremented.

- **Default:** ::1
- **Range:** ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff:

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Enabling MLD Static Group Membership

group-limit (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | 7841
- Hierarchy Level | 7842
- Description | 7842
- Default | 7842
- Options | 7842
- Required Privilege Level | 7842
- Release Information | 7842

Syntax

```
group-limit limit;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name],
```

Description

Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on a dynamic logical interface. After this limit is reached, new reports will be ignored and all related flows are not flooded on the logical interface.

Default

By default, there is no limit to the number of multicast groups that can join the interface.

Options

limit—group limit value for the interface.

- **Range:** 1 through 32767

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

group-limit (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 7843
- Hierarchy Level | 7843
- Description | 7843
- Default | 7843
- Options | 7843
- Required Privilege Level | 7844
- Release Information | 7844

Syntax

```
group-limit limit;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Description

Configure a limit for the number of multicast groups (or [S,G] channels in MLDv2) allowed on a dynamic logical interface. After this limit is reached, new reports will be ignored and all related flows are not flooded on the logical interface.

Default

By default, there is no limit to the number of multicast groups that can join the interface.

Options

limit—group limit value for the interface.

- **Range:** 1 through 32767

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Configuring the Number of MLD Multicast Group Joins on Logical Interfaces

group-limit (IGMP)

IN THIS SECTION

- Syntax | 7845
- Hierarchy Level | 7845
- Description | 7845
- Default | 7845
- Options | 7845
- Required Privilege Level | 7845
- Release Information | 7845

Syntax

```
group-limit limit;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Description

Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.

To confirm the configured group limit on the interface, use the `show igmp interface` command.

Default

By default, there is no limit to the number of multicast groups that can join the interface.

Options

limit—group limit value for the interface.

- **Range:** 1 through 32767

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

group-threshold (Protocols IGMP Interface) | 7903

log-interval (IGMP Interface) | 10009

group-limit (IGMP and MLD Snooping)

IN THIS SECTION

- Syntax | 7846
- Hierarchy Level | 7846
- Description | 7847
- Default | 7847
- Options | 7847
- Required Privilege Level | 7847
- Release Information | 7847

Syntax

```
group-limit limit;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols vlan vlan-id igmp-snooping interface interface-name]
```

Description

Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.

Default

By default, there is no limit to the number of multicast groups joining an interface.

Options

limit—a 32-bit number for the limit on the interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping

group-limit (Protocols MLD)

IN THIS SECTION

- Syntax | 7848
- Hierarchy Level | 7848
- Description | 7848
- Default | 7848

- Options | 7848
- Required Privilege Level | 7848
- Release Information | 7849

Syntax

```
group-limit limit;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

Configure a limit for the number of multicast groups (or [S,G] channels in MLDv2) allowed on a logical interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.

Default

By default, there is no limit to the number of multicast groups that can join the interface.

Options

limit—group value limit for the interface.

- **Range:** 1 through 32767

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring MLD

group-members

IN THIS SECTION

- Syntax | 7849
- Hierarchy Level | 7849
- Description | 7849
- Options | 7850
- Required Privilege Level | 7850
- Release Information | 7850

Syntax

```
group-members [attack-or-attack-group-name];
```

Hierarchy Level

```
[edit security idp custom-attack-group custom-attack-group-name]
```

Description

Specify the group members in a custom group. The members can be predefined attacks, predefined attack groups, custom attacks, or custom dynamic groups.

Use custom groups for the following tasks:

- To define a specific set of attacks to which you know your network is vulnerable.
- To group your custom attack objects.
- To define a specific set of informational attack objects that you use to keep you aware of what is happening on your network.

Options

attack-or-attack-group-name—Name of the attack object or group attack object.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

group (Origin Validation for BGP)

IN THIS SECTION

- Syntax | 7851
- Hierarchy Level | 7851
- Description | 7851
- Options | 7852
- Required Privilege Level | 7852
- Release Information | 7852

Syntax

```
group group-name {
    max-sessions number;
    session address {
        hold-time seconds;
        local-address local-ip-address;
        port port-number;
        preference number;
        record-lifetime seconds;
        refresh-time seconds;
        traceoptions {
            file filename <files number>   <size size> <(world-readable | no-world-readable)>;
            flag flag {
                disable;
                flag-modifier
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation],
[edit logical-systems logical-system-name routing-options validation],
[edit routing-instances instance-name routing-options validation],
[edit routing-options validation]
```

Description

Configure the number of concurrent sessions for each group.

Caches are organized in groups. The Junos OS implementation supports up to 63 sessions per group and both IPv4 and IPv6 address families.

If the number of sessions in a group exceeds the `max-sessions` value, the connections are established in order by the`preference` value. A numerically higher preference results in a higher probability for session establishment. The order of session establishment is random among sessions with equal preferences.

Options

group-name Name of the validation group.

max-sessions *number* Configure the number of concurrent sessions for each group. If the number of sessions in a group exceeds the `max-sessions` value, the connections are established in order by the `preference` value. A numerically higher preference results in a higher probability for session establishment. The order of session establishment is random among sessions with equal preferences.

- **Range:** 1 through 63 sessions
- **Default:** 2 sessions

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

group-policy (Dynamic IGMP Interface)

IN THIS SECTION

Syntax | 7853

- Hierarchy Level | 7853
- Description | 7853
- Options | 7853
- Required Privilege Level | 7853
- Release Information | 7853

Syntax

```
group-policy policy-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name]
```

Description

Compare the IGMPv2 or IGMPv3 group against the specified group policy, after receiving an IGMP report, and perform the action configured in that policy (for example, reject the report).

Options

policy-name—Name of the group policy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Filtering Unwanted IGMP Reports at the IGMP Interface Level

group-policy (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 7854
- Hierarchy Level | 7854
- Description | 7854
- Options | 7855
- Required Privilege Level | 7855
- Release Information | 7855

Syntax

```
group-policy policy-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Description

Compare the MLDv1 or MLDv2 group against the specified group policy, after receiving an MLD report, and perform the action configured in that policy (for example, reject the report).

Options

policy-name	Name of the group policy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Filtering Unwanted MLD Reports at the MLD Interface Level

group-policy (Protocols IGMP)

IN THIS SECTION

- Syntax | 7855
- Hierarchy Level | 7856
- Description | 7856
- Required Privilege Level | 7856
- Release Information | 7856

Syntax

```
group-policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Description

When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Filtering Unwanted IGMP Reports at the IGMP Interface Level

group-policy (Protocols IGMP AMT Interface)

IN THIS SECTION

- Syntax | 7857
- Hierarchy Level | 7857
- Description | 7857
- Options | 7857
- Required Privilege Level | 7857

Release Information | 7857

Syntax

```
group-policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp amt relay defaults],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols igmp
amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances routing-instance-name protocols igmp amt relay defaults]
```

Description

When this statement is enabled on the Automatic Multicast Tunneling (AMT) interfaces running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report).

Options

policy-names—Name of the policy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Default IGMP Parameters for AMT Interfaces

group-policy (Protocols MLD)

IN THIS SECTION

- Syntax | 7858
- Hierarchy Level | 7858
- Description | 7858
- Required Privilege Level | 7858
- Release Information | 7859

Syntax

```
group-policy  [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

When a routing device running MLD version 1 or version 2 (MLDv1 or MLDv2), receives an MLD report, the routing device compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Filtering Unwanted MLD Reports at the MLD Interface Level

group-profile (Access)

IN THIS SECTION

- Syntax | 7859
- Hierarchy Level | 7860
- Description | 7860
- Options | 7860
- Required Privilege Level | 7860
- Release Information | 7861

Syntax

```
group-profile profile-name {
    ppp {
        cell-overhead;
        encapsulated-overhead;
        framed-pool address-pool-name;
        idle-timeout seconds;
        interface-id interface-identifier;
        keepalive seconds;
        primary-dns IP address;
        primary-wins IP address;
        secondary-dns IP address;
```

```
        secondary-dns IP address;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure a group profile to define Point-to-Point Protocol (PPP) attributes. Any client referencing the configured group profile inherits all the group profile attributes.

Options

- `ppp`—Configure Point-to-Point Protocol (PPP) attributes.
- `cell-overhead`—Configure the session to use Asynchronous Transfer Mode (ATM)-aware egress shaping.
- `framed-pool` *`pool-name`*—Configure a framed-pool.
- `idle-timeout`—Configure the idle timeout for a user.
- `interface-id`—Configure the interface identifier.
- `keep-alive`—Configure the keepalive interval for an L2TP tunnel.
- `primary-dns`—Specify the primary-dns IP address.
- `secondary-dns`—Specify the secondary-dns IP address.
- `primary-wins`—Specify the primary-wins IP address.
- `secondary-wins`—Specify the secondary-wins IP address.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.4 of Junos OS.

RELATED DOCUMENTATION

Firewall User Authentication Overview

group-profile (Group Profile)

IN THIS SECTION

- Syntax | 7861
- Hierarchy Level | 7862
- Description | 7862
- Options | 7862
- Required Privilege Level | 7863
- Release Information | 7863

Syntax

```
group-profile profile-name {
    l2tp {
        interface-id interface-id;
        lcp-renegotiation;
        local-chap;
        maximum-sessions-per-tunnel number;
    }
    ppp {
        cell-overhead;
        encapsulation-overhead bytes;
        framed-pool pool-id;
        idle-timeout seconds;
        interface-id interface-id;
```

```
        keepalive seconds;
        ppp-options {
            aaa-options aaa-options-name;
            chap;
            ignore-magic-number-mismatch;
            initiate-ncp (ip | ipv6 | dual-stack-passive)
            ipcp-suggest-dns-option;
            mru;
            mtu;
            pap;
            peer-ip-address-optional;
        }
        primary-dns primary-dns;
        primary-wins primary-wins;
        secondary-dns secondary-dns;
        secondary-wins secondary-wins;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure the group profile.

> **NOTE**: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

profile-name—Name assigned to the group profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

Layer 2 Tunneling Protocol (L2TP)

L2TP LNS Inline Service Interfaces

group (Protocols BGP)

IN THIS SECTION

- Syntax | 7863
- Hierarchy Level | 7867
- Description | 7867
- Options | 7868
- Required Privilege Level | 7868
- Release Information | 7868

Syntax

```
group group-name {
    advertise-bgp-static
    advertise-inactive;
    allow [ network/mask-length ];
```

```
    authentication-key key;
    cluster cluster-identifier;
    damping;
    description text-description;
    enforce-first-as;
    export [ policy-names ];
    family {
        (inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {
            (any | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                    drop-excess <percentage>;
                    hide-excess <percentage>;
                 }
                add-path {
                    receive;
                    send {
                        include-backup-path backup_path_number;
                        multipath;
                        path-count number;
                        path-selection-mode {
                            (all-paths | equal-cost-paths);
                        }
                        prefix-policy [ policy-names ];
                    }
                }
                aigp [disable];
                damping;
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                    drop-excess <percentage>;
                    hide-excess <percentage>;
                  }
                rib-group group-name;
                topology name {
                     community {
                        target identifier;
                    }
                }
            }
            flow {
```

```
                no-validate policy-name;
            }
            labeled-unicast {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                    drop-excess <percentage>;
                    hide-excess <percentage>;
                }
                explicit-null {
                    connected-only;
                }
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                    drop-excess <percentage>;
                    hide-excess <percentage>;
                }
                resolve-vpn;
                rib inet.3;
                rib-group group-name;
            }
        }
        route-target {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
                drop-excess <percentage>;
                hide-excess <percentage>;
             }
            advertise-default;
            external-paths number;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
                drop-excess <percentage>;
                hide-excess <percentage>;
             }
        }
    }
    graceful-restart {
        long-lived {
            receiver {
```

```
                enable:
                disable;
            }
            advertise-to-non-llgr-neighbor {
                omit-no-export;
            }
        }
    }
    graceful-restart {
        long-lived {
            disable-notification-flag;
            disable-notification-extensions {
                omit-no-export;
            }
            forwarding-state-bit (from-fib | set); /* Configurable to be common for all address
families */
            forwarding-state-bit (as-rr-client | from-fib); /* Configurable for each address
family */
            restarter {
                disable;
                stale-time interval;
            }
        }
    }
    hold-time seconds;
    import [ policy-names ];
    ipsec-sa ipsec-sa;
    keep (all | none);
    local-address address;
    local-as autonomous-system <private>;
    local-preference local-preference;
    log-updown;
    link-bandwith auto-sense;
    metric-out metric;
    multihop <ttl-value>;
    multipath {
        multiple-as;
    }
    mvpn-iana-rt-import;
    no-aggregator-id;
    no-client-reflect;
    out-delay seconds;
    passive;
```

```
    peer-as autonomous-system;
    preference preference;
    remove-private;
    rfc6514-compliant-safi129;
    send-non-transitive-link-bandwidth;
    tcp-aggressive-transmission;
    tcp-mss segment-size;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    type type;
    neighbor address {
        ... peer-specific-options ...
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit protocols bgp],
[edit protocols bgp family prefix-limit (inet | inet6)(any | flow | labeled-unicast | multicast |
unicast)],
[edit routing-instances routing-instance-name protocols bgp]
```

Description

Define a BGP peer group. BGP peer groups share a common type, peer autonomous system (AS) number, and cluster ID, if present. To configure multiple BGP groups, include multiple `group` statements.

By default, the group's options are identical to the global BGP options. To override the global options, include group-specific options within the `group` statement.

The `group` statement is one of the statements you must include in the configuration to run BGP on the routing device.

Each group must contain at least one peer and can establish up to 2048 peers. Before establishing a BGP session, RPD checks for established peer count for the group. If the peer count is more than 2048 then the BGP session fails.

Options

group-name	Name of the BGP group
link-bandwith auto-sense	Enable this option to automatically detect the bandwidth towards the peer and attach the link-bandwith value to the interface speed over which the BGP route was received. This prevents a policy churn when there is a change in the link-bandwidth of an interface. Make sure that you also include the `auto-sense` statement in the import policy to enable this feature.
send-non-transitive-link-bandwidth	Enable this option to send non-transitive-link-bandwidth-extended-community towards EBGP neighbors only for link-bandwidth communities. All non-transitive link bandwith community values are advertised including originated or received and readvertised link-bw community.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`link-bandwith auto-sense`, and `send-non-transitive-link-bandwidth` options introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

- BGP User Guide

group (Protocols IGMP)

IN THIS SECTION

- Syntax | 7869
- Hierarchy Level | 7869
- Description | 7869
- Required Privilege Level | 7870
- Release Information | 7870

Syntax

```
group multicast-group-address {
    exclude;
    group-count number;
    group-increment increment;
    source ip-address {
        source-count number;
        source-increment increment;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name static],
[edit protocols igmp interface interface-name static]
```

Description

Specify the IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.

NOTE: You must specify a unique address for each group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling IGMP Static Group Membership

group (Protocols MLD)

IN THIS SECTION

- Syntax | 7871
- Hierarchy Level | 7871
- Description | 7871
- Options | 7871
- Required Privilege Level | 7871
- Release Information | 7872

Syntax

```
group multicast-group-address {
    exclude;
    group-count number;
    group-increment increment;
    source ip-address {
        source-count number;
        source-increment increment;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name static],
[edit protocols mld interface interface-name static]
```

Description

The MLD multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.

Options

multicast-group-address—Address of the group.

> **NOTE**: You must specify a unique address for each group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling MLD Static Group Membership

group (Protocols MLD Snooping)

IN THIS SECTION

- Syntax | 7872
- Hierarchy Level | 7872
- Description | 7873
- Options | 7873
- Required Privilege Level | 7873
- Release Information | 7873

Syntax

```
group multicast-group-address {
    source ip-address;
}
```

Hierarchy Level

```
[edit protocols mld-snooping vlan (all | vlan-name) interface (all | interface-name) static],
[edit routing-instances instance-name protocols mld-snooping vlan vlan-name interface interface-
name static]
```

Description

Configure a static multicast group on an interface and (optionally) the source address for the multicast group.

Options

multicast-group-address—Valid IP multicast address for the multicast group.

source *ip-address*—Valid IP multicast address for the source of the multicast group.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

Support at the [edit routing-instances *instance-name* protocols mld-snooping vlan *vlan-name* interface *interface-name* static] hierarhcy level introduced in Junos OS Release 13.3 for EX Series switches.

Support for the source statement introduced in Junos OS Release 13.3 for EX Series switches.

RELATED DOCUMENTATION

Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)

group (Protocols MSDP)

IN THIS SECTION

- Syntax | 7874
- Hierarchy Level | 7874
- Description | 7875
- Options | 7875

Required Privilege Level | 7875
Release Information | 7875

Syntax

```
group group-name {
    disable;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    mode (mesh-group | standard);
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    peer address; {
        disable;
        active-source-limit {
            maximum number;
            threshold number;
        }
        authentication-key peer-key;
        default-peer;
        export [ policy-names ];
        import [ policy-names ];
        local-address address;
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
```

```
msdp],
[edit protocols msdp],
[edit routing-instances routing-instance-name protocols msdp]
```

Description

Define an MSDP peer group. MSDP peers within groups share common tracing options, if present and not overridden for an individual peer with the "peer" on page 12380 statement. To configure multiple MSDP groups, include multiple `group` statements.

By default, the group's options are identical to the global MSDP options. To override the global options, include group-specific options within the `group` statement.

The group must contain at least one peer.

Options

group-name—Name of the MSDP group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP in a Routing Instance

group (Protocols PIM)

IN THIS SECTION

- Syntax | 7876
- Hierarchy Level | 7876
- Description | 7876
- Options | 7877
- Required Privilege Level | 7877
- Release Information | 7877

Syntax

```
group group-address {
    source source-address {
        rate threshold-rate;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
mdt threshold],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel  family intet | inet6mdt threshold],
[edit routing-instances routing-instance-name protocols pim mdt threshold],
[edit routing-instances routing-instance-name provider-tunnel family inet | inet6mdtthreshold]
```

Description

Specify the explicit or prefix multicast group address to which the threshold limits apply. This is typically a well-known address for a certain type of multicast traffic.

Options

group-address—Explicit group address to limit.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4. In Junos OS Release 17.3R1, the `mdt` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default MDT in Rosen 7, and data MDT for Rosen 6 and Rosen 7. The `provider-tunnel mdt` hierarchy is now hidden for backward compatibility with existing scripts.

RELATED DOCUMENTATION

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

group (Protocols RIP)

IN THIS SECTION

- Syntax | 7878
- Hierarchy Level | 7879
- Description | 7879
- Options | 7879
- Required Privilege Level | 7880
- Release Information | 7880

Syntax

```
group group-name {
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        transmit-interval {
            threshold milliseconds;
            minimum-interval milliseconds;
        }
        multiplier number;
        version (0 | 1 | automatic);
    }
    demand-circuit;
    export policy;
    max-retrans-time  seconds;
    metric-out metric;
    preference number;
    route-timeout seconds;
    update-interval seconds;
    neighbor neighbor-name {
        authentication-key password;
        authentication-type type;
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
```

```
            transmit-interval {
                threshold milliseconds;
                minimum-interval milliseconds;
            }
            multiplier number;
            version (0 | 1 | automatic);
        }
        (check-zero | no-check-zero);
        demand-circuit;
        import policy-name;
        max-retrans-time seconds;
        message-size number;
        metric-in metric;
        metric-out metric;
        receive receive-options;
        route-timeout seconds;
        send send-options;
        update-interval seconds;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit protocols rip],
[edit routing-instances routing-instance-name protocols rip]
```

Description

Configure a set of RIP neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group. Each group must contain at least one neighbor. You should create a group for every export policy.

Options

group-name—Name of a group, up to 16 characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring RIP

group (Protocols RIPng)

IN THIS SECTION

- Syntax | 7880
- Hierarchy Level | 7881
- Description | 7881
- Options | 7881
- Required Privilege Level | 7881
- Release Information | 7881

Syntax

```
group group-name {
    export [ policy-names ];
    metric-out metric;
    neighbor neighbor-name {
        import policy-name;
        metric-in metric;
        receive <none>;
```

```
        route-timeout seconds;
        send <none>;
        update-interval seconds;
    }
    preference number;
    route-timeout seconds;
    update-interval seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng],
[edit protocols ripng],
[edit routing-instances routing-instance-name protocols ripng]
```

Description

Configure a set of RIPng neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group.

Each group must contain at least one neighbor. You should create a group for each export policy that you have.

Options

group-name—Name of a group, up to 16 characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

group-range (Data MDTs)

IN THIS SECTION

- Syntax | 7882
- Hierarchy Level | 7882
- Description | 7882
- Options | 7883
- Required Privilege Level | 7883
- Release Information | 7883

Syntax

```
group-range multicast-prefix;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel family inet / inet6mdt],
[edit routing-instances routing-instance-name  provider-tunnel family inet / inet6mdt]
```

Description

Establish the group range to use for data MDTs created in this VRF instance. Only IPv4 address are valid for group range. This address range cannot overlap the default MDT addresses of any other VPNs on the router, nor can the group range specified under the `inet` and `inet6` hierarchies overlap. If you configure overlapping group ranges, the configuration commit fails. Up to 8000 MDT group ranges are supported for IPv4 and IPv6.

Options

multicast-prefix—Multicast address range to identify data MDTs.

- **Range:** Any valid, nonreserved multicast address range
- **Default:** None (No data MDTs are created for this VRF instance.)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4. In Junos OS Release 17.3R1, the `mdt` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default MDT in Rosen 7, and data MDT for Rosen 6 and Rosen 7. The `provider-tunnel mdt` hierarchy is now hidden for backward compatibility with existing scripts.

RELATED DOCUMENTATION

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

group-range (MBGP MVPN Tunnel)

IN THIS SECTION

- Syntax | 7884
- Hierarchy Level | 7884
- Description | 7884
- Options | 7884
- Required Privilege Level | 7885
- Release Information | 7885

Syntax

```
group-range multicast-prefix;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group group-address source source-address pim-ssm],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group group-address wildcard-source pim-ssm],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet wildcard-source pim-ssm],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet6 wildcard-source pim-ssm],
[edit routing-instances routing-instance-name provider-tunnel selective group group-address
source source-address pim-ssm],
[edit routing-instances routing-instance-name provider-tunnel selective group group-address
wildcard-source pim-ssm],
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet
wildcard-source pim-ssm],
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet6
wildcard-source pim-ssm]
```

Description

Establish the multicast group address range to use for creating MBGP MVPN source-specific multicast selective PMSI tunnels.

Options

multicast-prefix—Multicast group address range to be used to create MBGP MVPN source-specific multicast selective PMSI tunnels.

- **Range:** Any valid, nonreserved IPv4 multicast address range
- **Default:** None

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

group-ranges

IN THIS SECTION

- Syntax | 7885
- Hierarchy Level | 7885
- Description | 7886
- Default | 7886
- Options | 7886
- Required Privilege Level | 7886
- Release Information | 7886

Syntax

```
group-ranges {
    destination-ip-prefix</prefix-length>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp bidirectional address address],
[edit logical-systems logical-system-name protocols pim rp embedded-rp],
[edit logical-systems logical-system-name routing-instances instance-name protocols pim rp
```

```
bidirectional address address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp embedded-rp],
[edit protocols pim rp bidirectional address address],
[edit protocols pim rp embedded-rp],
[edit protocols pim rp local family (inet | inet6)],
[edit protocols pim rp static address address],
[edit routing-instances instance-name protocols pim rp bidirectional address address],
[edit routing-instances routing-instance-name protocols pim rp embedded-rp],
[edit routing-instances routing-instance-name protocols pim rp local family (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim rp static address address]
```

Description

Configure the address ranges of the multicast groups for which this routing device can be a rendezvous point (RP).

Default

The routing device is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12).

Options

destination-ip-prefix</prefix-length>—Addresses or address ranges for which this routing device can be an RP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for bidirectional RP addresses introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Local PIM RPs

Configuring PIM Embedded RP for IPv6

Example: Configuring Bidirectional PIM

group (Redundant Trunk Groups)

IN THIS SECTION

- Syntax | 7887
- Hierarchy Level | 7887
- Description | 7888
- Options | 7888
- Required Privilege Level | 7888
- Release Information | 7888

Syntax

```
group name {
    interface interface-name <primary>;
        interface interface-name;
    preempt-cutover-timer seconds;
}
```

Hierarchy Level

- For platforms with ELS:

```
[edit switch-options redundant-trunk-group]
```

- For platforms without ELS:

```
[edit ethernet-switching-options redundant-trunk-group]
```

Description

Create a redundant trunk group.

Options

name—The name of the redundant trunk group.

- For platforms with ELS:

  The group name must be a string “rtg*n*” where *n* is a number from 0 through 15, such as “rtg2” or “rtg10”.

- For platforms without ELS:

  The group name must start with a letter and can consist of letters, numbers, dashes, and underscores.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Hierarchy level `[edit switch-options]` introduced in Junos OS Release 13.2X50-D10 (ELS). (See *Using the Enhanced Layer 2 Software CLI* for information about ELS.)

RELATED DOCUMENTATION

Example: Configuring Redundant Trunk Links for Faster Recovery on EX Series Switches

Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support

Understanding Redundant Trunk Links (Legacy RTG Configuration)

group (Routing Instances)

IN THIS SECTION

- Syntax | 7889
- Hierarchy Level | 7890
- Description | 7890
- Options | 7890
- Required Privilege Level | 7890
- Release Information | 7890

Syntax

```
group address {
    source source-address {
        inter-region-segmented {
            fan-out fan-out value;
            threshold rate-value;
        }
        ldp-p2mp;
            pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
        threshold-rate number;
    }
    wildcard-source {
        inter-region-segmented {
            fan-out fan-out value;
        }
        ldp-p2mp;
        pim-ssm {
```

```
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective],
[edit routing-instances routing-instance-name provider-tunnel selective]
```

Description

Specify the IP address for the multicast group configured for point-to-multipoint label-switched paths (LSPs) and PIM-SSM GRE selective provider tunnels.

Options

address—Specify the IP address for the multicast group. This address must be a valid multicast group address.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `inter-region-segmented` statement added in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

Configuring PIM-SSM GRE Selective Provider Tunnels

group-rp-mapping

IN THIS SECTION

- Syntax | 7891
- Hierarchy Level | 7891
- Description | 7892
- Options | 7892
- Required Privilege Level | 7892
- Release Information | 7892

Syntax

```
group-rp-mapping {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
    log-interval seconds;
    maximum limit;
    threshold value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
```

```
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Configure a limit for the number of incoming group-to-RP mappings.

> **NOTE**: The maximum limit settings that you configure with the `maximum` and the `family (inet | inet6) maximum` statements are mutually exclusive. For example, if you configure a global maximum group-to-RP mapping limit, you cannot configure a limit at the family level for IPv4 or IPv6. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.

Options

`family (inet | inet6)`—(Optional) Specify either IPv4 or IPv6 messages to be counted towards the configured group-to-RP mapping limit.

- **Default:** Both IPv4 and IPv6 messages are counted towards the configured group-to-RP limit.

The remaining statements are described separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring PIM State Limits

group (RPF Selection)

IN THIS SECTION

- Syntax | 7893
- Hierarchy Level | 7893
- Description | 7893
- Default | 7894
- Options | 7894
- Required Privilege Level | 7894
- Release Information | 7894

Syntax

```
group group-address{
    sourcesource-address{
        next-hop next-hop-address;
    }
    wildcard-source {
        next-hop next-hop-address;
    }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name edit protocols pim rpf-selection]
```

Description

Configure the PIM group address for which you configure RPF selection"group (RPF Selection)" on page 7893.

Default

By default, PIM RPF selection is not configured.

Options

group-address—PIM group address for which you configure RPF selection.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring PIM RPF Selection

group (Security Group VPN)

IN THIS SECTION

- Syntax | 7895
- Hierarchy Level | 7895
- Description | 7896
- Options | 7896
- Required Privilege Level | 7897
- Release Information | 7897

Syntax

```
group name {
    anti-replay-time-window milliseconds;
    description description;
    group-id number;
    ike-gateway gateway-name;
    ipsec-sa name {
        match-policy policy-name {
            destination ip-address/netmask;
            destination-port number;
            protocol number;
            source ip-address/netmask;
            source-port number;
        }
        proposal proposal-name;
    }
    member-threshold number;
    server-cluster {
        ike-gateway gateway-name;
        retransmission-period seconds;
        server-role (root-server | sub-server);
    }
    server-member-communication {
        certificate certificate-id;
        communication-type (unicast);
        encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
        lifetime-seconds seconds;
        number-of-retransmission number;
        retransmission-period seconds;
        sig-hash-algorithm (sha-256 | sha-384);
    }
}
```

Hierarchy Level

```
[edit security group-vpn server]
```

Description

Configure group VPN on the group server. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 devices and vSRX Virtual Firewall instances.

Options

name—Name of the group.

- `anti-replay-time-window` *milliseconds*—Configure antireplay time in milliseconds. Specify a value from 1 to 60,000.

  We recommend that NTP be configured on Group VPNv2 devices to ensure proper antireplay operation.

  Group members that are running on vSRX Virtual Firewall instances on a host machine where the hypervisor is running under a heavy load may experience issues that can be corrected by reconfiguring the `anti-replay-time-window` value. If data that matches the IPsec policy on the group member is not being transferred, check the `show security group-vpn member ipsec statistics` output for D3P errors. Make sure that NTP is operating correctly. If there are errors, adjust the `anti-replay-time-window` value.

- `description` *description*—Description of the group.

- `group-id` *number*—Identifier for this group VPN. Specify a value from 1 to 4,294,967,295.

- `ike-gateway` *gateway-name*—Define the group member for Phase 1 negotiation. There can be multiple instances of this option configured. When a group member sends its registration request to the server, the server checks to see that the member is configured for the group.

- `ipsec-sa` *name*—Configure the group SAs to be downloaded to members. There can be multiple group SAs downloaded to group members.

- `member-threshold` *number*—Specify the maximum number of group VPN members that can be accepted in the group. The same `member-threshold` value must be configured on the root-server and all sub-servers in a group server cluster.

  The maximum number you can configure for a group is dependent upon the group server platform. Also, the sum of the `member-threshold` numbers for all groups configured on the group server must not exceed the capacity of the group server platform.

- `server-cluster`—Configure the Group Domain of Interpretation (GDOI) group controller/key server (GCKS) cluster for the specified group. All servers in a group VPN server cluster must be SRX Series Firewalls.

- `server-member-communication`—Enable and configure server to member communication. When these options are configured, group members receive new keys before current keys expire.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2

`member-threshold` option introduced in Junos OS Release 15.1X49-D30 for vSRX Virtual Firewall.

RELATED DOCUMENTATION

Group VPNv2 Overview

group (Static Subscribers)

IN THIS SECTION

- Syntax | 7897
- Hierarchy Level | 7898
- Description | 7898
- Options | 7899
- Required Privilege Level | 7899
- Release Information | 7899

Syntax

```
group group-name {
    access-profile profile-name;
```

```
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
    }
    authentication {
        password password-string;
        username-include {
            delimiter delimiter-character;
            domain-name domain-name;
            interface;
            logical-system-name;
            routing-instance-name;
             user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    interface  interface-name  <exclude> <upto upto-interface-name>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name system services static-subscribers],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers],
[edit routing-instances routing-instances-name system services static-subscribers],
[edit system services static-subscribers]
```

Description

Configure a static subscriber group with values that override the values configured at the `[edit system services static-subscribers]` hierarchy level for subscribers outside the group. Includes the subscriber access and dynamic profiles, the authentication parameters that trigger the Access-Request message to AAA for static subscribers in the group, and the statically configured interfaces that form the group.

NOTE: The logical system and routing instance in which the group is configured must match the logical system and routing instance where the static interfaces are configured.

Options

group-name—Name of a group that defines authentication parameters for static subscribers to override the global authentication configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Creating a Static Subscriber Group

group (System Services DHCP)

IN THIS SECTION

- Syntax | 7900
- Hierarchy Level | 7902
- Description | 7902
- Options | 7902
- Required Privilege Level | 7902
- Release Information | 7902

Syntax

```
group group-name {
    authentication {
        password password;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name;
            interface-name;
            logical-system-name;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix;
        }
    }
    dynamic-profile {
        profile-name;
        aggregate-clients {
            merge;
            replace;
        }
        junos-default-profile;
        use-primary dynamic-profile;
    }
    interface interface-name {
        dynamic-profile {
            profile-name;
            aggregate-clients {
                merge;
                replace;
            }
            junos-default-profile;
            use-primary dynamic-profile-name;
        }
        exclude;
        overrides {
            delegated-pool pool-name;
            interface-client-limit number;
```

```
            process-inform {
                pool pool-name;
            }
            rapid-commit ;
        }
        service-profile service-profile-name
        trace ;
        upto interface-name;
    }
    liveness-detection {
        failure-action {
             clear-binding;
            clear-binding-if-interface-up;
            log-only;
        }
        method {
            bfd {
                detection-time {
                    threshold milliseconds;
                }
                holddown-interval interval;
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                session-mode (automatic | multihop | single-hop);
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
            version (0 | 1 | automatic);
        }
    }
    overrides {
        delegated-pool pool-name;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        rapid-commit ;
    }
    reconfigure {
        attempts number;
```

```
        clear-on-abort;
        strict;
        timeout number;
        token token-name;
        trigger {
            radius-disconnect;
        }
    }
    service-profile service-profile-name;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6]
```

Description

Configure a group of interfaces that have a common configuration.

The remaining statements are explained separately. See CLI Explorer.

Options

- *group-name*—Name of the group.

> **NOTE**: SRX Series Firewalls do not support DHCP client authentication.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

DHCP Overview

DHCP Server

DHCP Server Configuration

group-threshold (Protocols IGMP Interface)

IN THIS SECTION

- Syntax | 7903
- Hierarchy Level | 7903
- Description | 7903
- Default | 7904
- Options | 7904
- Required Privilege Level | 7904
- Release Information | 7904

Syntax

```
group-threshold value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Description

Specify the threshold at which a warning message is logged for the multicast groups received on a logical interface. The threshold is a percentage of the maximum number of multicast groups allowed on a logical interface.

For example, if you configure a maximum number of 1,000 incoming multicast groups, and you configure a threshold value of 90 percent, warning messages are logged in the system log when the interface receives 900 groups.

To confirm the configured group threshold on the interface, use the `show igmp interface` command.

Default

By default, there is no configured threshold value.

Options

value—Percentage of the maximum number of multicast groups allowed on the interface that starts triggering the warning. You configure a percentage of the **group-limit** value that starts triggering the warnings. You must explicitly configure the **group-limit** to configure a threshold value.

- **Range:** 1 through 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

group-limit (IGMP) | **7844**

log-interval (IGMP Interface) | **10009**

group-threshold (Protocols MLD Interface)

IN THIS SECTION

- Syntax | 7905
- Hierarchy Level | 7905
- Description | 7905
- Default | 7906
- Options | 7906
- Required Privilege Level | 7906
- Release Information | 7906

Syntax

```
group-threshold value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

Specify the threshold at which a warning message is logged for the multicast groups received on a logical interface. The threshold is a percentage of the maximum number of multicast groups allowed on a logical interface.

For example, if you configure a maximum number of 1,000 incoming multicast groups, and you configure a threshold value of 90 percent, warning messages are logged in the system log when the interface receives 900 groups.

To confirm the configured group threshold on the interface, use the `show mld interface` command.

Default

By default, there is no configured threshold value.

Options

value—Percentage of the maximum number of multicast groups allowed on the interface that starts triggering the warning. You configure a percentage of the **group-limit** value that starts triggering the warnings. You must explicitly configure the **group-limit** to configure a threshold value.

- **Range:** 1 through 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring the Number of MLD Multicast Group Joins on Logical Interfaces

group-limit (Protocols MLD) | **7847**

log-interval (MLD Interface) | **10011**

group (Traffic Load Balancer)

IN THIS SECTION

- Syntax | **7907**
- Hierarchy Level | **7907**
- Description | **7907**
- Options | **7907**

Required Privilege Level | 7908
Release Information | 7908

Syntax

```
group group-name {
    health-check-interface-subunit health-check-interface-subunit;
    network-monitoring-profile [profile-name1, <profile-name2>];
    real-service-rejoin-options no-auto-rejoin;
    real-services [server-list];
    <routing-instance routing-instance>;
}
```

Hierarchy Level

```
[edit services traffic-load-balance instance instance-name]
```

Description

Configure a group of servers as a pool for next-hop session distribution.

Options

group-name	Use the specified string identifier for a group of servers to which sessions are distributed using the server distribution table in conjunction with the session distribution API.
group health-check-interface-subunit *health-check-interface-subunit*	Use the specified subunit of the ms- interface used for health checking.
network-monitoring-profile *profile-name1*	Name of the network monitoring profile used to monitor the health of servers in the group.
network-monitoring-profile *profile-name2*	(Optional) Name of a second network monitoring profile used to monitor the health of servers in the group.

real-services ***server-list***	Use the specified list of individual servers to which sessions are distributed using the server distribution table in conjunction with the session distribution API.
real-services-rejoin-options no-auto-rejoin	Disable the default behavior that allows a server to rejoin the group automatically when it comes up.
routing-instance ***routing-instance***	(Optional) Use the specified routing instance if the default `inet.0` is not used.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Traffic Load Balancer Overview

Configuring TLB

group-type (Unknown Unicast Forwarding)

IN THIS SECTION

- Syntax | 7909
- Hierarchy Level | 7909
- Description | 7909
- Options | 7909

Required Privilege Level | 7909
Release Information | 7909

Syntax

```
group-type (none | layer-2)
```

Hierarchy Level

```
[edit forwarding-options next-hop-group]
```

Description

Configure the type of addresses to be used in the next-hop group.

Options

none—Next-hop group uses Layer 2 addresses.

`layer-2`—Specify a next-hop group that uses Layer 2 addresses.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring Unknown Unicast Forwarding (ELS)

Configuring Mirroring on EX9200 Switches to Analyze Traffic (CLI Procedure)

group (Uplink Failure Detection)

IN THIS SECTION

- Syntax | 7910
- Hierarchy Level | 7910
- Description | 7910
- Options | 7911
- Required Privilege Level | 7911
- Release Information | 7911

Syntax

```
group  group-name {
    debounce-interval seconds;
    link-to-monitor interface-name;
    link-to-disable interface-name;
}
```

Hierarchy Level

```
[edit protocols uplink-failure-detection]
```

Description

Configure a group of uplink and downlink interfaces for uplink failure detection. You can also configure the debounce interval, which is the amount of time, in seconds, that elapses before the downlink interfaces are brought up after a state change of the uplink interfaces.

Options

group-name Name of the uplink-failure-detection group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

group-vpn

IN THIS SECTION

- Syntax | 7911
- Hierarchy Level | 7913
- Description | 7913
- Options | 7913
- Required Privilege Level | 7914
- Release Information | 7914

Syntax

```
group-vpn {
    member {
        ike {
            gateway gateway-name;
                policy;
                proposal;
```

```
                traceoptions;
            }
        ipsec {
            vpn vpn-name {
                df-bit (clear | copy | set);
                exclude rule rule-name {
                    source-address ip-address/mask;
                    destination-address ip-address/mask;
                    application application;
                }
                fail-open rule rule-name {
                    source-address ip-address/mask;
                    destination-address ip-address/mask;
                    application application;
                }
                group id;
                group-vpn-external-interface interface;
                ike-gateway gateway-name;
                recovery-probe;
            }
        }
    }
    server {
        group name {
            anti-replay-time-window milliseconds;
            description description;
            group-id number;
            ike-gateway gateway-name;
            ipsec-sa;
            member-threshold number;
            server-cluster;
        }
        ike {
            gateway gateway-name;
                policy;
            proposal;
        }
        ipsec {
            proposal proposal-name;
        }
        traceoptions (Security Group VPN);
```

```
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure Group VPNs in Group VPNv2. Group VPNv2 extends IPsec architecture to support SAs that are shared by a group of security devices. With Group VPNv2, any-to-any connectivity is achieved by preserving the original source and destination IP addresses in the outer header.

Options

member	Configure group VPN member.
ike	Configure IPsec group VPN on the group member.
policy	Configure an IKE policy.
proposal	Define an IKE proposal. You can configure one or more IKE proposals. Each proposal is a list of IKE attributes to protect the IKE connection between the IKE host and its peer.
traceoptions	Configure group VPN tracing options to aid in troubleshooting the IKE or server issues.
ipsec	Configure IPsec for Phase 2 exchange on the group member.
vpn	Configure IPsec VPN for Phase 2 exchange on the group member.
server	Configure group VPN server.
group	Configure group VPN on the group server.
anti-replay-time-window	Configure antireplay time in milliseconds. Specify a value from 1 to 60,000. Each IPsec packet contains a timestamp. The group member checks whether the packet's timestamp falls within the configured `anti-replay-time-window` value. A packet is dropped if the timestamp exceeds the value.

description	Description of the group.
group-id *number*	Identifier for this group VPN. Specify a value from 1 to 4,294,967,295.
ike-gateway *gateway-name*	Define the group member for Phase 1 negotiation. There can be multiple instances of this option configured. When a group member sends its registration request to the server, the server checks to see that the member is configured for the group.
ipsec-sa	Configure the group SAs to be downloaded to members. There can be multiple group SAs downloaded to group members.
member-threshold	Specify the maximum number of group VPN members that can be accepted in the group. There is no default number.
server-cluster	Configure the Group Domain of Interpretation (GDOI) group controller/key server (GCKS) cluster for the specified group. All servers in a group VPN server cluster must be SRX Series Firewalls.
server-member-communication	Enable and configure server to member communication. When these options are configured, group members receive new keys before current keys expire.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

grouped-ie-profile

IN THIS SECTION

- Syntax | 7915
- Hierarchy Level | 7915
- Description | 7915
- Options | 7916
- Required Privilege Level | 7916
- Release Information | 7916

Syntax

```
grouped-ie-profile name {
    ie ie-number {
        instance inst-number {
            grouped-ie-profile grpie-prf-name;
        }
    }
}
```

Hierarchy Level

```
[edit security gtp]
```

Description

Specify a GTP grouped IE filter profile. A grouped IE is a group of IEs or a group of grouped IEs. A grouped-ie-profile is defined and linked to the IE when a Must-IE check is expected to go deep into a grouped IE.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp]` hierarchy level is replaced by `[edit security gtp]`.

Options

name Name of the GTP grouped IE filter profile.

ie-number IE number.

- **Range:** 1 to 255

inst-number instance number.

- **Range:** 0 to 11
- **Default:** 0

grpie-prf-name GTP grouped IE filter profile.

Required Privilege Level

security—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Example: Configure Must-IE check for GTPv1 and GTPv2

groups

IN THIS SECTION

- Syntax | 7917
- Hierarchy Level | 7917
- Description | 7917
- Options | 7918

Required Privilege Level | 7919
Release Information | 7919

Syntax

```
groups {
    group-name {
        configuration-data;
        when {
            chassis chassis-id;
            member member-id;
            model model-id;
            node node-id;
            peers [ names-of-peers ]
            routing-engine routing-engine-id;
            time <start-time> [to <end-time>];
        }
            conditional-data;
    }
    lccn-re0 {
        configuration-data;
    }
    lccn-re1 {
        configuration-data;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Create a configuration group.

> **NOTE**: Junos OS does not support configuring statements corresponding to third-party YANG data models, for example, OpenConfig or custom data models, under the `[edit groups]` hierarchy.

Options

group-name Name of the configuration group. To configure multiple groups, specify more than one group name.

configuration-data The configuration statements that are to be applied elsewhere in the configuration with the `apply-groups` statement, to have the target configuration inherit the statements in the group.

when Define conditions under which the configuration group should be applied. Conditions include the type of chassis, model, or Routing Engine, virtual chassis member, cluster node, and start and optional end time of day. If you specify multiple conditions in a single configuration group, all conditions must be met before the configuration group is applied.

- `chassis` *chassis-id*—Specify the chassis type of the router. Valid types include SCC0, SCC1, LCC0, LCC1 ... LCC3.
- `member` *member-id*—Specify the name of the member of the virtual chassis.
- `model` *model-id*—Specify the model name of the router, such as m7i or tx100.
- `node` *node-id*—Specify the cluster node.
- `peers` *names-of-peers*—Specify the names of the MC-LAG peers participating in commit synchronization.
- `routing-engine` *routing-engine-id*—Specify the type of Routing Engine, re0 or re1.
- `time` *start-time* `[to` *end-time*`]`—Specify the start time or time duration for this configuration group to be applied. If only the start time is specified, the configuration group is applied at the specified time and remains in effect until the time is changed. If the end time is specified, then on each day, the applied configuration group is started and stopped at the specified times. The syntax for specifying the time uses the format yyyy-mm-dd.hh:mm, hh:mm, or hh.

conditional-data Option introduced in Junos 11.3. The conditional statements that are to be applied when this configuration group is applied. On routers that support multiple Routing Engines, you can also specify two special group names:

- `re0`—Configuration statements that are to be applied to the Routing Engine in slot 0.
- `re1`—Configuration statements that are to be applied to the Routing Engine in slot 1.

On routers that support multiple Routing Engines, you can also specify two special group names:

The configuration specified in group `re0` is applied only if the current Routing Engine is in slot 0; likewise, the configuration specified in group `re1` is applied only if the current Routing Engine is in slot 1. Therefore, both Routing Engines can use the same configuration file, each using only the configuration statements that apply to it. Each `re0` or `re1` group contains at a minimum the configuration for the hostname and the management interface (`fxp0`). If each Routing Engine uses a different management interface, the group also should contain the configuration for the backup router and static routes.

(Routing matrix only) The TX Matrix router supports group names for the Routing Engines in each connected T640 router in the following formats:

> **NOTE**: The management Ethernet interface used for the TX Matrix Plus router, T1600 routers in a routing matrix, and PTX Series Packet Transport Routers, is `em0`. Junos OS automatically creates the router's management Ethernet interface, `em0`.

- `lcc`*n*`-re0`—Configuration statements applied to the Routing Engine in slot 0 of the specified T640 router that is connected to a TX Matrix router.
- `lcc`*n*`-re1`—Configuration statements applied to the specified to the Routing Engine in slot 1 of the specified T640 router that is connected to a TX Matrix router.

  n identifies the T640 router and can be from 0 through 3.

Required Privilege Level

configure—To enter configuration mode.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Creating a Configuration Group

apply-groups | 2423

apply-groups-except | 2426

grpc

IN THIS SECTION

- Syntax | 7920
- Hierarchy Level | 7921
- Description | 7921
- Options | 7921
- Required Privilege Level | 7921
- Release Information | 7921

Syntax

```
grpc {
    max-connections max-connections;
    routing-instance routing-instance;
    ssl {
        address ip-address;
        hot-reloading;
        local-certificate (System) local-certificate;
        mutual-authentication {
            certificate-authority certificate-authority-profile-name;
            client-certificate-request (no-certificate | request-certificate | request-
certificate-and-verify | require-certificate | require-certificate-and-verify);
        }
        port port;
        use-pki;
```

```
    }
}
```

Hierarchy Level

```
[edit system services extension-service request-response]
```

Description

Configure the type of connections the gRPC service accepts for API applications.

Options

routing-instance ***routing-instance***	Name of routing instance for grpc.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

request-response | 14741

grpc-replay

IN THIS SECTION

- Syntax | 7922
- Hierarchy Level | 7922
- Description | 7922
- Options | 7922
- Required Privilege Level | 7923
- Release Information | 7923

Syntax

```
grpc-replay <last minute|facility facility|priority priority>
```

Hierarchy Level

```
[edit system syslog]
```

Description

Configure this configuration statement to replay the events.

Options

- `last` *`minute`*–Use this configuration to replay the events. Specify the time in minutes. Once you configure this configuration statement, eventd starts storing the syslog events. Eventd streams upto one million events. Once the replaying of the old data is over, eventd starts streaming the live events. While eventd is replaying the old data as initial sync and any live event takes place, the system immediately streams the live event.
- `facility` *`facility`*–Use this option to filter the events as per facility. You can use this option to filter the live and replay events. This configuration has no relationship with the existing syslog configuration

for writing and sending events to file and remote host. Choose facility from the available options under this configuration.

- `priority` *`priority`*–Use this configuration to filter the events as per priority. You can use this configuration to filter the live and replay events. This configuration has no relationship with the existing syslog configuration for writing and sending events to file and remote host. Choose the priority from the available options under this configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

grpc-tunnel

IN THIS SECTION

- Syntax | 7923
- Hierarchy Level | 7924
- Description | 7924
- Options | 7924
- Required Privilege Level | 7926
- Release Information | 7926

Syntax

```
grpc-tunnel {
```

```
    servers {
        retry-interval interval;
        server server-name {
            address ip-address;
            port port-number;
            credentials {
                tls {
                    certificate-id local-certificate in pki
                    ca-profiles [ specific ca-profiles in pki ];
                }
            }
            targets [ ssh | netconf-ssh | gnmi-gnoi];
            routing-instance routing-instance;
            source-address ip-address;
        }
    }
    target-string-option {
        pattern [ hostname | vendor | model | version | custom ];
        custom-string custom-string;
        delimiter delimiter-string;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure gRPC tunnel settings.

Options

retry-interval

Optional. Specify the time after which the target retries to connect again if the tunnel-server is not reachable. There is no limit on the number of times the target retries to connect to the server.

- Range: 10-300
- Default: 30 seconds

server Specify the tunnel server name. A maximum of 10 tunnel servers can be configured at a time.

address Specify the IPv4 or IPv6 address or hostname of the tunnel server.

port Specify the port number through which the tunnel server listens.

credentials Configure options for TLS-based authentication.

tls Configure TLS certificate related options.

certificate-id Optional. Specify the local tunnel client certificate-id loaded in the PKI hierarchy.

Mandatory in cases where mutual authentication is required and the client must present the certificate to the server.

ca-profiles Optional. Specify the list of certificate authority (CA)s to validate the certificate of the tunnel server. If not specified, the gRPC tunnel utilizes the loaded CAs through PKI hierarchy to validate the server certificate.

targets Specify the list of target applications that are to be made available. It may contain one or more target types.

- Values: `ssh`, `netconf-ssh`, `gnmi-gnoi`

routing-instance Optional. Specify the name of the routing instance through which the target can reach the tunnel server.

source-address Optional. Specify the IPv4 or IPv6 source address that can connect with the tunnel server.

target-string-option Optional. Configure to modify the target string that gets registered with the tunnel server.

pattern Optional. Specify the list of supported options. You can use this statement to select the string to be created using the options.

- Default: hostname
- Values: `hostname`, `vendor`, `model`, `version`, `custom`

custom-string Specify the string to be used when **pattern** contains **custom** as one of the options.

delimiter Specify a character to act as the delimiter.

- Default: `|` (pipe character)

Required Privilege Level

system— To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 22.4R1.

`routing-instance` *`routing-instance`* and `source-address` *`ip-address`* options introduced in Junos OS and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Using gRPC Dial-Out for Secure Telemetry Collection

gsm-options

IN THIS SECTION

- Syntax | 7926
- Hierarchy Level | 7927
- Description | 7927
- Options | 7927
- Required Privilege Level | 7927
- Release Information | 7927

Syntax

```
gsm-options {
    select-profile profile-name;
    profiles {
        profile-name {
```

```
            sip-user-id simple-ip-user-id;
            sip-password simple-ip-password;
            access-point-name apn;
            authentication-method (pap | chap | none);
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options]
```

Description

Configure the 3G wireless modem interface to establish a data call with a Global System for Mobile Communications (GSM) cellular network.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

gsmp-syn-timeout (ANCP)

IN THIS SECTION

- Syntax | 7928
- Hierarchy Level | 7928
- Description | 7928
- Options | 7928
- Required Privilege Level | 7929
- Release Information | 7929

Syntax

```
gsmp-syn-timeout seconds;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure the maximum period that the ANCP agent waits before sending a SYN message to an ANCP neighbor to negotiate the adjacency. If the neighbor sends a SYN message during this period, the ANCP agent uses the partition information in the neighbor's message when generating its own initial SYN message to the neighbor. The agent does not wait for the period to expire if it receives a SYN message from the neighbor.

Options

seconds—Number of seconds the ANCP agent waits.

- **Range:** 1 through 60 seconds
- **Default:** 60 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Configuring the ANCP Agent

Configuring the ANCP Agent to Learn ANCP Partition IDs

gsmp-syn-wait (ANCP)

IN THIS SECTION

- Syntax | 7929
- Hierarchy Level | 7930
- Description | 7930
- Default | 7930
- Required Privilege Level | 7930
- Release Information | 7930

Syntax

```
gsmp-syn-wait;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Enable the ANCP agent to learn partition ID information from neighbors, in support of nonzero ANCP partition IDs. This statement forces the ANCP agent to delay sending a SYN message during adjacency negotiation for a configurable period. When the neighbor sends a SYN message to the ANCP agent during that period, the agent learns the partition ID information from the neighbor and uses that information when it sends its own SYN message. If the agent does not receive the message during the period, then it sends a SYN message to the neighbor when the period times out.

Default

This statement is disabled. The ANCP agent does not wait before sending the initial SYN message and does not support nonzero partition IDs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3

RELATED DOCUMENTATION

Configuring the ANCP Agent

Configuring the ANCP Agent to Learn ANCP Partition IDs

gtp

IN THIS SECTION

- Syntax | 7931
- Hierarchy Level | 7935
- Description | 7935
- Options | 7935
- Required Privilege Level | 7935
- Release Information | 7936

Syntax

```
gtp {
    grouped-ie-profile name {
        ie name {
            instance name {
                grouped-ie-profile grpie-prf-name;
            }
        }
    }
    handover-default {
        deny;
    }
    ie-set name {
        ie name;
    }
    ip-group name {
        address-book name {
            address-set set-name;
        }
    }
    message-ie-profile-v1 name {
        message name {
            ie name;
        }
    }
```

```
message-ie-profile-v2 name {
    message name {
        ie name {
            instance name {
                grouped-ie-profile grpie-prf-name;
            }
        }
    }
}
message-list name {
    message name;
}
profile name {
    apn name {
        imsi-prefix name {
            action {
                drop;
                pass;
                selection {
                    ms;
                    net;
                    vrf;
                }
            }
        }
    }
    drop {
        aa-create-pdp0;
        aa-delete-pdp0;
        bearer-resource2;
        change-notification2;
        config-transfer2;
        context2;
        create-bearer2;
        create-data-forwarding2;
        create-pdp(0 | 1 | all);
        create-session2;
        create-tnl-forwarding2;
        cs-paging2;
        data-record(0 | 1 | all);
        delete-bearer2;
        delete-command2;
        delete-data-forwarding2;
```

```
            delete-pdn2;
            delete-pdp(0 | 1 | all);
            delete-session2;
            detach2;
            downlink-notification2;
            echo(0 | 1 | 2 | all);
            error-indication(0 | 1 | all);
            failure-report(0 | 1 | all);
            fwd-access2;
            fwd-relocation(1 | 2 | all);
            fwd-srns-context1;
            g-pdu(0 | 1 | all);
            identification(0 | 1 | 2 | all);
            mbms-session-start(1 | 2 | all);
            mbms-session-stop(1 | 2 | all);
            mbms-session-update(1 | 2 | all);
            modify-bearer2;
            modify-command2;
            node-alive(0 | 1 | all);
            note-ms-present(0 | 1 | all);
            pdu-notification(0 | 1 | all);
            ran-info(1 | 2 | all);
            redirection(0 | 1 | all);
            release-access2;
            relocation-cancel(1 | 2 | all);
            resume2;
            send-route(0 | 1 | all);
            sgsn-context(0 | 1 | all);
            stop-paging2;
            supported-extension1;
            suspend2;
            trace-session2;
            update-bearer2;
            update-pdn2;
            update-pdp(0 | 1 | all);
            ver-not-supported(0 | 1 | 2 | all);
        }
        end-user-address-validated;
        gtp-in-gtp-denied;
        handover-group group-name;
        handover-on-roaming-intf;
        listening-mode;
        log {
```

```
            forwarded (basic | detail);
            gtp-u name;
            prohibited (basic | detail);
            rate-limited (basic | detail);
            state-invalid (basic | detail);
        }
        max-message-length max-message-length;
        min-message-length min-message-length;
        must-ie-v1 msgie-prf-v1-name;
        must-ie-v2 msgie-prf-v2-name;
        ne-group group-name;
        path-rate-limit {
            message-type (create-req | delete-req | echo-req | other) {
                alarm-threshold {
                    forward forward ;
                    reverse reverse ;
            }
                drop-threshold {
                    forward forward ;
                    reverse reverse ;
                }
            }
        }
        rate-limit {
            alarm-threshold alarm-threshold;
            drop-threshold drop-threshold;
        }
        remove-ie {
            versionv1 {
                release name;
            }
        }
        remove-ie-v1 ieset-name;
        remove-ie-v2 ieset-name;
        req-timeout seconds;
        restart-path (all | create | echo);
        seq-number-validated;
        timeout hours;
        u-tunnel-validated;
        ue-group group-name;
    }
    rate-limit name {
        message {
```

```
            v0 msg-list-name;
            v1 msg-list-name;
            v2 msg-list-name;
        }
    }
    traceoptions {
        file <filename > <files files > <match match > <size size > <(world-readable | no-
worldreadable)>;
        flag name;
        no-remote-trace;
        trace-level {
            (error | info | notice | verbose | warning);
        }
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Use the GTP commands to enable the GTP service, configure GTP objects, set traceoptions, remove GTP inspection object configurations, and obtain configuration information.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs]` hierarchy level is replaced by `[edit security gtp]`.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. The `restart-path` option added in Junos OS Release 11.4. New GPRS tunneling protocol (GTP) message types added in Junos OS Release 11.4. Support for GTPv2 added in Junos OS Release 11.4. Statement modified in Junos OS Release 15.1X49-D40. Support for `handover-default` and `handover-group` options added in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

GPRS Overview

Understanding GTP Support for Central Point Architecture

Enabling GTP Interoperability between 2G and 3G Networks

gtp-header-offset

IN THIS SECTION

- Syntax | 7936
- Hierarchy Level (QFX5000 line of switches) | 7936
- Description | 7937
- Required Privilege Level | 7937
- Release Information | 7938

Syntax

```
gtp-header-offset  offset-value
```

Hierarchy Level (QFX5000 line of switches)

```
[edit forwarding-options enhanced-hash-key hash-mode]
```

Description

You must only provide a decimal value for *gtp-header-offset* command. As per standards, 0x32 is the default header offset value. But, due to some special patterns in the header, offset may vary like, 0x30, 0x28, and so on. In this case, you need to convert it to a proper offset value as follows:

For Example:

- If the offset value is 0x32, the equivalent decimal value is calculated as (16*3+2) which is equal to 50.
- If the offset value is 0x30, the equivalent decimal value is calculated as (16*3+0) which is equal to 48.

After setting *gtp-header-offset* to a proper offset value, run *gtp-tunnel-endpoint-identifier* command to enable GTP hashing. Refer "gtp-tunnel-endpoint-identifier" on page 7938 for more details. For example:

- If the IPv6 traffic has *0x32* as offset value, then use `set forwarding-options enhanced-hash-key family inet` *`gtp-tunnel-endpoint-identifier`* command directly.
- If the IPv6 traffic has *0x30* as offset value, then you must set proper offset value after enabling GTP hashing as follows:

  `set forwarding-options enhanced-hash-key family inet` *`gtp-tunnel-endpoint-identifier`*

  `set forwarding-options enhanced-hash-key hash-mode` *`gtp-header-offset`* `48`

- If the IPv4 traffic has *0x32* as offset value, then use `set forwarding-options enhanced-hash-key family inet gtp-tunnel-endpoint-identifier` command directly.
- If the IPv4 traffic has *0x30* as offset value, then you must set proper offset value after enabling GTP hashing as follows:

  ```
  set forwarding-options enhanced-hash-key family inet gtp-tunnel-endpoint-identifier
  ```

  ```
  set forwarding-options enhanced-hash-key hash-mode gtp-header-offset 48
  ```

> **NOTE**: Both IPv4 and IPv6 share same GTP header offset value. If there are two different traffic streams, that is, if IPv4 has value X as offset value and IPv6 has Y as offset value, then hashing will be successful only for one of the traffic stream.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R3.

RELATED DOCUMENTATION

gtp-tunnel-endpoint-identifier | 7938

gtp-tunnel-endpoint-identifier

IN THIS SECTION

- Syntax | 7938
- Hierarchy Level | 7938
- Hierarchy Level (QFX5000 line of switches) | 7939
- Hierarchy Level (QFX10000 line of switches) | 7939
- Description | 7939
- Required Privilege Level | 7939
- Release Information | 7940

Syntax

```
gtp-tunnel-endpoint-identifier;
```

Hierarchy Level

```
[edit forwarding-options hash-key family inet layer-4],
```

```
[edit forwarding-options hash-key family inet6 layer-4]
```

Hierarchy Level (QFX5000 line of switches)

```
[edit forwarding-options enhanced-hash-key family inet]
```

Hierarchy Level (QFX10000 line of switches)

```
[edit forwarding-options enhanced-hash-key family inet],
```

```
[edit forwarding-options enhanced-hash-key family inet6]
```

Description

When you configure `gtp-tunnel-endpoint-identifier`, the hash calculation of IPv4 or IPv6 packets are included in the GPRS tunneling protocol–tunnel endpoint ID (GTP-TEID) field hash calculations.

> **NOTE**: The `gtp-tunnel-endpoint-identifier` configuration statement is supported on PTX Series routers only when network services is set to `enhanced-mode`. For more information, see `enhanced-mode`.

On the QFX5000 and QFX10000 lines of switches, if the `gtp-tunnel-endpoint-identifier` statement is configured, the default Layer 4 port 2152 (and 2123 as well on QFX5000) is set to use along with the default first byte 0x32.

(On the QFX5000 line of switches only) In most cases, configuring the `gtp-tunnel-endpoint-identifier` statement on QFX5000 switches is sufficient for enabling GTP hashing. After you have enabled GTP hashing, if GTP hashing does not work, we recommend that you capture the packets using relevant tools and identify the offset value. As per standards, 0x32 is the default header offset value. But, due to some special patterns in the header, offset may vary to say 0x30, 0x28, and so on. In these cases, use the `gtp-header-offset` statement to set a proper offset value. Once the header offset value is resolved, run the `gtp-tunnel-endpoint-identifier` command for enabling GTP hashing successfully. Refer to "gtp-header-offset" on page 7936 for more details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F3 and 16.1R2.

RELATED DOCUMENTATION

hash-key (Forwarding Options) | 7986

Understanding Per-Packet Load Balancing

Configuring Per-Packet Load Balancing

gtp-header-offset | 7936

guaranteed-rate (Traffic Control Profiles)

IN THIS SECTION

- Syntax | 7940
- Hierarchy Level | 7940
- Description | 7941
- Default | 7941
- Options | 7941
- Required Privilege Level | 7942
- Release Information | 7942

Syntax

```
guaranteed-rate (percent percentage | rate) <burst-size bytes>;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name]
```

Description

Configure a guaranteed minimum rate of transmission for a traffic control profile. The sum of the guaranteed rates of all of the forwarding class sets (priority groups) on a port should not exceed the total port bandwidth. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group (forwarding class set) can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.

> **NOTE**: You cannot configure a guaranteed rate for a forwarding class set (priority group) that includes strict-high priority queues. If the traffic control profile is for a forwarding class set that contains strict-high priority queues, do not configure a guaranteed rate.

On supported platforms, you can also configure an optional burst size for a logical interface. This can help to ensure that higher priority services do not starve lower priority services.

Default

If you do not include this statement and you do not include the `delay-buffer-rate` statement, the logical interface receives a minimal delay-buffer rate and minimal bandwidth equal to 2 MTU-sized packets.

If you do not specify a guaranteed rate, the guaranteed rate is zero (0) and there is no minimum guaranteed bandwidth.

> **NOTE**: If you do not configure a guaranteed rate for a traffic control profile, the queues that belong to any forwarding class set (priority group) that uses that traffic control profile cannot have a configured transmit rate. The result is that there is no minimum guaranteed bandwidth for those queues and that those queues can be starved during periods of congestion.

Options

`percent` *`percentage`*

Minimum percentage of transmission capacity allocated to the forwarding class set or logical interface.

- **Range:** 1 through 100 percent

`rate`

Minimum transmission rate allocated to the forwarding class set or logical interface, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000 bps

`burst-size` *`bytes`* (Optional) Maximum burst size, in bytes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Option `burst-size` introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Providing a Guaranteed Minimum Rate

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

output-traffic-control-profile | **12016**

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Example: Configuring Minimum Guaranteed Output Bandwidth

Understanding CoS Traffic Control Profiles

output-traffic-control-profile | **12016**

guaranteed-rate (Dynamic Traffic Shaping)

IN THIS SECTION

- Syntax | **7943**
- Hierarchy Level | **7943**
- Description | **7943**

- Default | 7943
- Options | 7943
- Required Privilege Level | 7944
- Release Information | 7944

Syntax

```
guaranteed-rate (rate | $junos-cos-guaranteed-rate) <burst-size [ bytes | $junos-cos-guaranteed-rate-burst]>;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

Configure a guaranteed minimum rate for a logical interface and optionally a burst size for the guaranteed rate.

NOTE: The guaranteed-rate burst size must not exceed the shaping-rate burst size.

Default

If you do not include this statement and you do not include the `delay-buffer-rate` statement, the logical interface receives a minimal delay-buffer rate and minimal bandwidth equal to 2 MTU-sized packets.

Options

rate—Guaranteed rate in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000 bps

`$junos-cos-guaranteed-rate`—Junos predefined variable that is replaced with the guaranteed rate obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

`burst-size` *`bytes`*—(Optional) Maximum burst size, in bytes.

- **Range:** 0 through 1,000,000,000

`$junos-cos-guaranteed-rate-burst`—(Optional) Variable for the burst-size that is specified for the guaranteed rate. Use this variable at the [`edit dynamic-profiles` *`profile-name`* `class-of-service traffic-control-profile`] hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

The `$junos-cos-guaranteed-rate` variable introduced in Junos OS Release 9.4.

Option `burst-size` introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Traffic Scheduling and Shaping for Subscriber Access

output-traffic-control-profile (Dynamic CoS Definition) | **12018**

guard-band

IN THIS SECTION

- Syntax | **7945**
- Hierarchy Level | **7945**

- Description | 7945
- Options | 7945
- Required Privilege Level | 7945
- Release Information | 7946

Syntax

```
guard-band watts;
```

Hierarchy Level

```
[edit poe],
[edit poe (all | fpc slot-number)]
```

Description

Reserve a specified amount of power from the PoE power budget for the switch, line card, or satellite device in case of a spike in PoE consumption.

Options

watts—Amount of power to be reserved in case of a spike in PoE consumption.

- **Range:** 0 through 19 for all switches except EX6200 and EX8200 switches.

  0 through 19 for ACX2000 routers.

  0 through 15 for EX6200 and EX8200 switches.

  0 through 19 for satellite devices in a Junos Fusion.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Metro Routers.

Statement introduced in Junos OS Release 17.2R1 for a Junos Fusion Provider Edge.

RELATED DOCUMENTATION

Configuring PoE Interfaces on EX Series Switches

Configuring Power over Ethernet in a Junos Fusion

guard-band (PoE)

IN THIS SECTION

- Syntax | 7946
- Hierarchy Level | 7946
- Description | 7947
- Options | 7947
- Required Privilege Level | 7947
- Release Information | 7947

Syntax

```
guard-band watts;
```

Hierarchy Level

```
[edit poe]
```

Description

Reserves the specified amount of power for the SRX Series Firewall in case of a spike in PoE consumption.

Options

watts—Amount of power to be reserved for the SRX Series Firewall in case of a spike in PoE consumption.

- **Range:** 0 through 19 W
- **Default:** 0 W

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Power over Ethernet

guard-interval

IN THIS SECTION

- Syntax | 7948
- Hierarchy Level | 7948
- Description | 7948
- Options | 7948

Required Privilege Level | 7948
Release Information | 7949

Syntax

```
guard-interval number;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

When a link goes down, the ring protection link (RPL) activates. When the downed link comes back up, the RPL link receives notification, restores the link, and waits for the restore interval before issuing another block on the same link. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

Options

number—Guard timer interval, in milliseconds.

- **Range:** 10 through 2000 ms
- **Default:** 500 ms

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

RELATED DOCUMENTATION

- *Ethernet Ring Protection Switching Overview*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

guest-network-functions

IN THIS SECTION

- Syntax | 7949
- Hierarchy Level | 7950
- Description | 7950
- Required Privilege Level | 7950
- Release Information | 7950

Syntax

```
guest-network-functions {
    gnf  id {
        control-plane-bandwidth-percent  percent;
        description  description;
        fpcs  fpcs;
        af-name
```

```
    }
}
```

Hierarchy Level

```
[edit chassis network-slices]
```

Description

Configure a guest network function for Junos node slicing.

Junos node slicing enables a single MX Series router to be partitioned to appear as multiple, independent routers. Each partition has its own Junos control plane, which runs as a virtual machine (VM), and a dedicated set of line cards. Each partition is called a guest network function (GNF).

The remaining statements are explained separately.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Setting Up Junos Node Slicing

network-slices | **11322**

gw-interface

IN THIS SECTION

- Syntax | 7951
- Hierarchy Level | 7951
- Description | 7951
- Options | 7952
- Required Privilege Level | 7952
- Release Information | 7952

Syntax

```
gw-interface interface-name.logical-unit-number;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options local-gateway address]
```

Description

Enable the cleanup of IKE triggers and IKE and IPsec SAs when an IPsec tunnel's local gateway IP address goes down or the MS-MIC or MS-MPC being used in the tunnel's service set goes down. If the local gateway IP address for an IPsec tunnel's service set goes down or the MS-MIC or MS-MPC that is being used in the service set goes down, the service set no longer sends IKE triggers. In addition, when the local gateway IP address goes down, the IKE and IPsec SAs are cleared for next-hop service sets, and go to the Not Installed state for interface-style service sets. The SAs that have the Not Installed state are deleted when the local gateway IP address comes back up.

If the local gateway IP address that goes down is for the responder peer, then you need to manually clear the IKE and IPsec SAs on the initiator peer so that the IPsec tunnel comes back up once the local gateway IP address comes back up (see "clear services ipsec-vpn ike security-associations" on page 4112 and "clear services ipsec-vpn ipsec security-associations" on page 4113).

Options

interface-name Name of the interface of the IPsec local gateway.

logical-unit-numer Number of the logical unit of the IPsec local gateway interface. You must include the logical unit number.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Configuring IPsec Service Sets

gx-plus (Gx-Plus)

IN THIS SECTION

- Syntax | 7953
- Hierarchy Level | 7953
- Description | 7953
- Required Privilege Level | 7953
- Release Information | 7953

Syntax

```
gx-plus {
    global {
        include-ipv6;
        max-outstanding-requests number;
    }
    partition partition-name  {
        diameter-instance instance-name;
        destination-host hostname;
        destination-realm realm;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure the Gx-Plus application to interact with a PCRF to authorize and provision subscribers.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

> *Configuring Gx-Plus*

gx-profile

IN THIS SECTION

- Syntax | 7954
- Hierarchy Level | 7955
- Description | 7955
- Options | 7955
- Required Privilege Level | 7955
- Release Information | 7955

Syntax

```
gx-profile profile-name {
    <attributes> {
        exclude {
            an-gw-address;
            default-eps-bearer-qos;
            packet-filter-information;
            packet-filter-operation;
            rat-type;
        }
        include {
            gx-capability-list;
            rule-suggestion;
        }
    }
    <request-timeout seconds>;
    targets {
        target-name {
            <destination-host hostname>;
            destination-realm realm-name;
            network-element element-name;
            priority priority-value;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit unified-edge diameter-profiles]
```

Description

Configure the Diameter profile used for Gx applications.

Options

profile-name	Name of the Diameter profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Diameter Profiles Overview

H

IN THIS SECTION

- h323 | **7961**
- halt-on-prefix-down (RFC 2544 Benchmarking) | **7963**
- handover-default | **7964**
- handover-group | **7966**
- handshake-timeout | **7967**
- hard-limit | **7969**
- hard-limit-target | **7970**
- hardware-assisted-keepalives | **7971**
- hardware-assisted-timestamping | **7973**
- hardware-timestamp | **7975**
- hardware-upgrade | **7976**
- has-prefix (Unified Edge Gateways) | **7977**
- has-suffix | **7979**
- hash-key (Chassis) | **7980**
- hash-key (Chassis LAG) | **7982**
- hash-key (Forwarding Options) | **7984**
- hash-key (Forwarding Options) | **7986**
- hash-keys (Aggregated Multiservices) | **7992**
- hash-keys (Interfaces) | **7995**
- hash-mode | **7997**
- hash-parameters | **8000**
- hash-seed | **8002**
- hash-seed (profile) | **8003**
- hcm (HTTP Header Enrichment) | **8005**
- hcm-profile (HTTP Header Enrichment) | **8007**
- hcm-profile (PCC Action Profiles) | **8008**
- header-integrity-check | **8010**
- header-integrity-check (Next Gen Services) | **8012**

- header-length | **8014**
- header-type | **8015**
- health-monitor | **8016**
- health-monitor (KHMS) | **8018**
- heartbeat-address (MX Series Virtual Chassis) | **8020**
- heartbeat-interval | **8022**
- heartbeat-threshold | **8024**
- heartbeat-timeout (MX Series Virtual Chassis) | **8025**
- heartbeat-tos (MX Series Virtual Chassis) | **8027**
- hello-acknowledgements | **8029**
- hello-dead-interval | **8030**
- hello-interval | **8032**
- hello-interval (L2TP) | **8033**
- hello-interval (LMP) | **8035**
- hello-interval (Protocols LDP) | **8036**
- hello-interval (Protocols RSVP) | **8038**
- hello-message (NETCONF) | **8039**
- hello-time | **8041**
- hello-timer | **8043**
- help | **8045**
- helper-disable (LDP) | **8046**
- helper-disable (Multiple Protocols) | **8048**
- helpers | **8049**
- hide-avps | **8053**
- hide-via-headers | **8054**
- hierarchical-policer | **8055**
- hierarchical-policer (aggregate or premium) | **8057**
- hierarchical-scheduler | **8060**
- hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) | **8062**
- high-availability (Chassis) | **8064**
- high-availability-options (Aggregated Multiservices) | **8068**
- high-availability (security cloud) | **8070**

- high-availability (Security IDP) | **8073**
- high-cos-queue-threshold | **8074**
- high-polarization | **8076**
- high-power | **8077**
- high-threshold (Resource Monitor) | **8078**
- high-watermark | **8080**
- hint | **8081**
- history | **8082**
- history-size | **8084**
- hold-down | **8086**
- hold-down-interval | **8088**
- hold-interval | **8089**
- hold-interval (Clock Synchronization) | **8091**
- hold-interval (OAM) | **8093**
- hold-interval (OAM CFM, for EX Series Switch Only) | **8094**
- hold-interval (Protection Group) | **8095**
- hold off time | **8097**
- hold-time (ae interface) | **8098**
- hold-time | **8101**
- hold-time (OAM) | **8102**
- hold-time (Protocols BGP) | **8104**
- hold-time (Protocols DVMRP) | **8106**
- hold-time (Protocols IS-IS) | **8108**
- hold-time (Protocols LDP) | **8110**
- hold-time (Protocols MSDP) | **8112**
- hold-time (Protocols PIM) | **8114**
- hold-time (Redundant Ethernet Interfaces) | **8115**
- hold-time up | **8117**
- hold-time (VRRP) | **8119**
- holddown (candidate-uplink-port-profile) | **8120**
- holddown-interval | **8122**
- holddown-interval (protocols ldp) | **8124**

- holddown (Protocols RIP) | **8125**
- holddown (Protocols RIPng) | **8127**
- holdover-mode-disable | **8128**
- holdover-time-error-budget | **8129**
- home-address | **8131**
- hop-by-hop-header | **8132**
- hop-limit | **8134**
- host (Address-Assignment Pools) | **8136**
- host-address-base | **8137**
- host-address-base (Source NAT Next Gen Services) | **8139**
- host (Diameter Origin) | **8140**
- host-fast-reroute | **8141**
- host-inbound-traffic | **8142**
- host (Interfaces) | **8144**
- host (L2TP) | **8146**
- host-name | **8147**
- host-name (System) | **8148**
- host-name (DHCP Relay Agent) | **8150**
- host-name (DHCPv6 Relay Agent) | **8151**
- host (Next Gen Services Service-Set Remote System Logging) | **8152**
- host-only-interface | **8154**
- host-os forwarding-options analyzer | **8155**
- host-outbound media-interface | **8157**
- host-outbound-traffic (Class-of-Service) | **8158**
- host-outbound-traffic (Class-of-Service) | **8160**
- host-outbound-traffic (Multicast Snooping) | **8162**
- host-prefix-only | **8163**
- host (Security IDP Sensor Configuration) | **8165**
- host (Security Log) | **8166**
- host (Security Web Filtering) | **8167**
- host (service-set) | **8169**
- host (System) | **8171**

- host-vpn | **8174**
- hosted-service-identifier | **8176**
- hosted-services | **8178**
- hostkey-algorithm | **8179**
- hostkey-algorithm-list | **8182**
- hot-root-standby (MBGP MVPN) | **8184**
- hot-standby (Routing Instance) | **8186**
- hot-standby | **8187**
- hot-standby (Protocols Layer 2 Circuit) | **8189**
- hot-standby-vc-on (Protocols Layer 2 Circuit) | **8190**
- http (Access) | **8192**
- http (icap-redirect profile) | **8193**
- http-log-multiple-transactions (LRF Profile) | **8195**
- http-persist | **8196**
- http-profile (Security Antivirus) | **8198**
- http-profile (Security Content Filtering) | **8199**
- http-profile (Security Web Filtering) | **8200**
- http-reassemble | **8201**
- http (REST API) | **8202**
- http (System Services) | **8204**
- http (Web Management) | **8205**
- https (REST API) | **8207**
- https (System Services) | **8208**
- https (Web Management) | **8211**
- hub-assist | **8213**
- hugepages | **8215**
- hw-db-profile | **8216**
- hybrid | **8220**
- hyper-mode (forwarding-options) | **8221**

h323

IN THIS SECTION

- Syntax | 7961
- Hierarchy Level | 7962
- Description | 7962
- Options | 7962
- Required Privilege Level | 7962
- Release Information | 7962

Syntax

```
h323 {
    application-screen {
        message-flood {
            gatekeeper {
                threshold rate;
            }
        }
        unknown-message {
            permit-nat-applied;
            permit-routed;
        }
    }
    disable;
    dscp-rewrite {
        code-point string;
    }
    endpoint-registration-timeout value-in-seconds;
    media-source-port-any;
    traceoptions {
        flag flag <detail | extensive | terse>;
    }
}
```

Hierarchy Level

```
[edit security alg]
[edit tenants tenant-name security alg]
```

Description

Specify the H.323 Application Layer Gateway (ALG) on the device. H.323 is a control-signaling protocol used to exchange messages between H.323 endpoints.

Options

`disable`—Disable the H.323 ALG. By default, H.323 ALG is enabled.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

application-screen (Security H323) | **2359**

dscp-rewrite | **6019**

halt-on-prefix-down (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 7963
- Junos OS Hierarchy Level | 7963
- Junos OS Evolved Hierarchy Level | 7963
- Description | 7963
- Required Privilege Level | 7964
- Release Information | 7964

Syntax

```
halt-on-prefix-down;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

By default, the RFC 2544-based benchmarking test ignores a prefix-down event (when the prefix associated with the test goes down) and continues to run. If you configure this parameter, a prefix that moves to the down state causes the corresponding tests to be stopped. The `show` command output for the test displays that the test was terminated due to the prefix going down.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

handover-default

IN THIS SECTION

- Syntax | **7965**
- Hierarchy Level | **7965**
- Description | **7965**
- Options | **7965**
- Required Privilege Level | **7965**
- Release Information | **7965**

Syntax

```
handover-default {
    deny;
}
```

Hierarchy Level

```
[edit security gtp]
```

Description

A GPRS tunneling protocol (GTP) handover group is a set of SGSNs or serving gateway (SGW) with a common address-book library. You can configure the handover group using the “handover-group” command. If there is no handover group defined in the GTP profile, and if the traffic reaches the policy configured with this profile, handover between all GTPs matching this policy is permitted by default. Therefore, if you configure handover-default deny command, then handover is denied.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp]` hierarchy level is replaced by `[edit security gtp]`.

Options

deny To deny default handover.

Required Privilege Level

security—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

GTP Handover Group Overview

Understanding GTP Handover Messages

handover-group

IN THIS SECTION

- Syntax | 7966
- Hierarchy Level | 7966
- Description | 7966
- Options | 7967
- Required Privilege Level | 7967
- Release Information | 7967

Syntax

```
handover-group group-name;
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

A GPRS tunneling protocol (GTP) handover group is a set of SGSNs or serving gateway (SGW) with a common address-book library. You can configure a GTP profile and associate an GTP handover group to the GTP profile. When a GTP handover group name is referenced by a GTP profile, the device checks to see if the current SGSN/SGW address and the proposed SGSN/SGW address are both contained within the same GTP handover group. If both SGSN/SGW addresses are contained within the same GTP handover group, then the handover is allowed. If both the current and proposed SGSN/SGW addresses are not within the same GTP handover group, then the profile for the default handover group is used.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name]` hierarchy level is replaced by `[edit security gtp profile profile-name]`.

Options

name	To specify the handover group on the GTP profile.
address-book	To specify the common address-book name.
address-set	To specify the address set for the handover group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

handover-group | 7966

handshake-timeout

IN THIS SECTION

- Syntax | 7968
- Hierarchy Level | 7968
- Description | 7968
- Options | 7968
- Required Privilege Level | 7968
- Release Information | 7968

Syntax

```
handshake-timeout time-in-seconds;
```

Hierarchy Level

```
[edit security sctp profile profile-name]
```

Description

Use the handshake-timeout function to set the handshake time for Stream Control Transmission Protocol (SCTP).

From Junos OS Release 20.4R1 onwards, the `[edit security gprs sctp profile profile-name]` hierarchy level is replaced by `[edit security sctp profile profile-name]`.

Options

time-in-seconds—Number of seconds of handshake time that elapse before the session is terminated.

- **Range:** 10 to 30 seconds.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

> *Understanding Stream Control Transmission Protocol*
>
> *SCTP Packet Structure Overview*

hard-limit

IN THIS SECTION

- Syntax | 7969
- Hierarchy Level | 7969
- Description | 7969
- Options | 7969
- Required Privilege Level | 7969
- Release Information | 7970

Syntax

```
hard-limit bandwidth;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Specify a bandwidth threshold at which the dynamic flow capture application begins deleting criteria, until the bandwidth falls below the `hard-limit-target` value.

Options

bandwidth—Hard limit threshold, in bits per second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

hard-limit-target | 7970

Configuring the Content Destination

hard-limit-target

IN THIS SECTION

- Syntax | 7970
- Hierarchy Level | 7970
- Description | 7970
- Options | 7971
- Required Privilege Level | 7971
- Release Information | 7971

Syntax

```
hard-limit-target bandwidth;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Specify a bandwidth threshold at which the dynamic flow capture application stops deleting criteria.

Options

bandwidth—Target value, in bits per second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

hard-limit | 7969

Configuring the Content Destination

hardware-assisted-keepalives

IN THIS SECTION

- Syntax | 7972
- Hierarchy Level | 7972
- Description | 7972
- Default | 7972
- Options | 7972
- Required Privilege Level | 7973
- Release Information | 7973

Syntax

```
hardware-assisted-keepalives [enable | disable];
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
```

Description

For Ethernet interfaces on Modular Port Concentrators (MPCs) in MX Series routers only, delegate the transmission of the continuity check messages (CCMs) to the forwarding ASIC (that is, to the hardware) by enabling inline transmission of CCMs. Inline transmission of CCMs is also known as Inline-KA.

By default, CCMs are transmitted by the CPU of the MPC and not by the hardware. If the duration between transmissions of CCMs is low or if the CCMs for a specific line card scale, we recommend that you enable delegation of the transmission of CCMs to the hardware. By enabling inline transmission of CCMs, you can achieve maximum scaling of CCMs.

> **NOTE**: Starting in Junos OS Release 16.1R1, the inline support for 1s interval is supported on MPC1 to MPC7 (except MPC3 and MPC4) line cards.

Default

Inline transmission is disabled by default.

Options

enable Enable inline transmission of CCMs.

> **NOTE**: Inline transmission of CCMs is not enabled when there is a CFM session already established. To enable inline transmission, you must first deactivate the CFM session using the `deactivate` command and then reactivate the CFM session using the `activate` command.

disable Disable inline transmission of CCMs.

> **NOTE**: After disabling inline transmission of CCMs, you must reboot the router for the changes to take effect.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R1.

RELATED DOCUMENTATION

Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling

Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades

hardware-assisted-timestamping

IN THIS SECTION

- Syntax | 7974
- Hierarchy Level | 7974
- Description | 7974
- Required Privilege Level | 7974
- Release Information | 7974

Syntax

```
hardware-assisted-timestamping;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
```

Description

For Ethernet interfaces on Enhanced and Enhanced Queuing Dense Port Concentrators (DPCs) in MX Series routers only, enable hardware-assisted timestamping support for Ethernet frame delay measurement.

By default, the ETH-DM feature calculates frame delays using software-based timestamping of the ETH-DM PDU frames sent and received by the MEPs in the session. As an option that can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction, you can enable hardware-assisted timestamping of session frames in the receive direction.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Ethernet Frame Delay Measurements Overview

Guidelines for Configuring Routers to Support an ETH-DM Session

Enabling the Hardware-Assisted Timestamping Option

hardware-timestamp

IN THIS SECTION

- Syntax | 7975
- Hierarchy Level | 7975
- Description | 7975
- Required Privilege Level | 7975
- Release Information | 7975

Syntax

```
hardware-timestamp;
```

Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Description

Enable timestamping of RPM probe messages in the Packet Forwarding Engine host processor. This feature is supported only with `icmp-ping`, `icmp-ping-timestamp`, `udp-ping`, and `udp-ping-timestamp` probe types.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

Statement applied to MX Series routers in Junos OS Release 10.0.

Statement introduced in Junos OS Release 19.1 for PTX Series routers.

hardware-upgrade

IN THIS SECTION

- Syntax | 7976
- Hierarchy Level | 7976
- Description | 7976
- Required Privilege Level | 7976
- Release Information | 7977

Syntax

```
hardware-upgrade
```

Hierarchy Level

```
[edit chassis high-availability]
```

Description

Install an additional SPC3 card in a Multinode High Availability statement. Use this statement when you insert an additional SPC3 on an SRX5000-line device in Multinode High Availability and has an interchassis link (ICL) in encryption mode. You must run the command on both nodes in Multinode High Availability before you start SPC3 installation.

Once you complete the installation, and both nodes are online, delete the statement on both nodes using the command `delete chassis high-availability hardware-upgrade`.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Hardware Upgrade for SRX5000-Line SPC3 in a Multinode High Availability Setup

has-prefix (Unified Edge Gateways)

IN THIS SECTION

- Syntax | 7977
- Hierarchy Level | 7977
- Description | 7978
- Options | 7978
- Required Privilege Level | 7978
- Release Information | 7978

Syntax

```
has-prefix {
    value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from class],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from 3gpp-imsi],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format string]
```

Description

Specify the prefix that the attribute must have.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

Options

value Prefix string.

- **Range:** 1 through 254 alphanumeric characters.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

has-suffix

IN THIS SECTION

- Syntax | 7979
- Hierarchy Level | 7979
- Description | 7979
- Options | 7979
- Required Privilege Level | 7980
- Release Information | 7980

Syntax

```
has-suffix {
    value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from class],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from 3gpp-imsi],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format string]
```

Description

Specify the suffix that the attribute must have.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

Options

value Suffix string.

- **Range:** 1 through 254 alphanumeric characters.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

hash-key (Chassis)

IN THIS SECTION

- Syntax | 7981
- Hierarchy Level | 7981
- Description | 7981
- Default | 7981
- Options | 7981
- Required Privilege Level | 7981
- Release Information | 7981

Syntax

```
hash-key (layer-3 | layer-4);
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-
provider data-flow-affinity]
```

Description

Set the hashing distribution of flow affinity. This is an optional setting. Once the `data-flow-affinity` statement is enabled, you may need to choose the hashing distribution. Modifying this statement causes the PIC to reboot.

Default

If you do not configure the `hash-key` statement, the hashing distribution is 5-tuple hashing, or `layer-4`.

Options

`layer-3`—3-tuple hashing (source IP address, destination IP address, and IP protocol).

`layer-4`—5-tuple hashing (3-tuple plus source and destination TCP or UDP ports).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

extension-provider | 6716

hash-key (Chassis LAG)

IN THIS SECTION

- Syntax | 7982
- Hierarchy Level | 7983
- Description | 7983
- Options | 7983
- Required Privilege Level | 7983
- Release Information | 7983

Syntax

```
hash-key {
    family {
        inet {
            layer-3;
            layer-4;
            symmetric-hash {
                complement;
            }
        }
        multiservice {
            source-mac;
            destination-mac;
            payload {
                ip {
                    layer-3 (source-ip-only | destination-ip-only);
                    layer-4;
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number]
```

Description

(MX Series 5G Universal Routing Platforms only) Configure data used in a hash key for a PIC for symmetrical load balancing on an 802.3ad Link Aggregation Group.

Options

family—Configure data used in a hash key for a protocol family. This statement has the following suboptions:

- **inet**—Configure data used in a hash key for the **inet** protocol family.
- **multiservice**—Configure data used in a hash key for the **multiservice** protocol family.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

> *Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers*

hash-key (Forwarding Options)

IN THIS SECTION

- Syntax | 7984
- Hierarchy Level | 7985
- Description | 7985
- Options | 7985
- Required Privilege Level | 7985
- Release Information | 7986

Syntax

```
hash-key {
    family {
        inet {
            layer-3;
            layer-4;
            inner-vlan-id;
            outer-vlan-id;
        }
    }
    multiservice {
        source-mac;
        destination-mac;
        payload {
            ip {
                layer3-only;
                layer-3 (source-ip-only | destination-ip-only);
                layer-4;
                inner-vlan-id;
                outer-vlan-id;
            }
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
[edit chassis fpc slot-number pic pic-number]
```

Description

(QFX10000 switches only) Select which packet header data to use for per-flow load balancing.

> **NOTE**: You can configure either Layer 3 or Layer 4 load balancing, or both at the same time.

> **NOTE**: On I chip platforms, an unknown Layer 4 header is excluded from load-balance hashing to avoid undesired packet reordering.

Options

- **inet**—IPv4 address family.
- **layer-3**—Incorporate Layer 3 data into the hash key.
- **layer-4**—Incorporate Layer 4 data into the hash key.
- **outer-vlan-id**—Include outer VLAN ID information in the hash key.
- **inner-vlan-id**—Include inner VLAN ID information in the hash key.
- **payload**—Incorporate payload data into the hash key.
- **ip**—Include the IP address of the IPv4 or IPv6 payload into the hash key.
- **layer-3-only**—Include only Layer 3 IP information.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D10.

hash-key (Forwarding Options)

IN THIS SECTION

- Syntax | 7986
- Hierarchy Level | 7988
- Description | 7988
- Options | 7990
- Required Privilege Level | 7991
- Release Information | 7991

Syntax

```
hash-key {
    family inet {
        layer-3;
        layer-4;
        session-id;
        symmetric-hash {
            complement;
        }
    }
    family inet6 {
            layer-3 {
                destination-address;
                incoming-interface-index;
                ipv6-flow-label;
                next-header;
                source-address;
                traffic-class;
            }
```

```
            layer-4 {
                destination-port;
                gtp-tunnel-endpoint-identifier;
                source-port;
            }
            no-incoming-port;
            session-id;
        }
    family mpls {
        all-labels;
        bottom-label-1;
        bottom-label-2;
        bottom-label-3;
        label-1;
        label-2;
        label-3;
        no-labels;
        no-label-1-exp;
        payload {
            ether-pseudowire {
                zero-control-word;
            }
            ip {
                disable;
                layer-3-only;
                port-data {
                    destination-lsb;
                    destination-msb;
                    source-lsb;
                    source-msb;
                }
            }
        }
    }
    family multiservice {
        destination-mac;
        label-1;
        label-2;
        payload {
            ip {
                layer-3-only;
                layer-3 {
                    (source-ip-only | destination-ip-only);
```

```
                }
                layer-4;
            }
        }
        source-mac:
    }
}
```

Hierarchy Level

```
[edit forwarding-options
```

Description

Select which packet header data to use for per-flow load balancing.

The options are explained separately.

NOTE: To modify the default hashing mechanism on Modular Port Concentrators (MPCs) and Type 5 FPCs, you need to configure the statements at the `[edit forwarding-options enhanced-hash-key]` hierarchy level. Statements at the `[edit forwarding-options hash-key]` hierarchy level do not support MPCs and Type 5 FPCs.

NOTE: The following statements are not supported on T Series routers:

- The `symmetric-hash` and the `session-id` statements at the `[edit forwarding-options hash-key family inet]` hierarchy level and all statements at the `[edit forwarding-options hash-key family multiservice]` hierarchy level.
- The `label-1` and `label-2` statements, and the IP configuration at the `[edit forwarding-options hash-key family multiservice]` hierarchy level.

NOTE: The following statements are not supported on QFX Series switches:

- The `symmetric-hash` and the `session-id` statements at the `[edit forwarding-options hash-key family inet]` hierarchy level.

NOTE: QFX 5000 Series switches with an EVPN Type 5 overlay must configure `overlay-ecmp` at the `[edit forwarding-options vxlan-routing` hierarchy.

NOTE: QFX 10000 Series switches in an EVPN-VXLAN environment should configure the `layer2` option with appropriate sub-option at the `[edit forwarding-options enhanced-hash-key]` hierarchy. This applies to the overlay fabric.

Addtionally, you should configure the `family` option with appropriate sub-option at the `[edit forwarding-options enhanced-hash-key` hierarchy. This applies to both EVPN-VXLAN underlay and overlay fabrics.

NOTE: On PTX platforms, *family inet/inet6 layer-3 traffic-class (dscp-v4/dscp-v6)* is disabled by default.

PTX platforms running Junos OS Evolved require the `enhanced-hash-key family` setting for EVPN-VXLAN underlay and overlay fabrics.

NOTE: EVPN-VXLAN configuration is required for underlay ECMP to work in ACX platforms on Junos EVO.

NOTE: On ACX 7000 Series of routers, the following statements are not supported:

- `label-1`, `label-2`, and `payload` statements at the `[edit forwarding-options hash-key family multiservice]` hierarchy level
- `bottom-label-1`, `bottom-label-2`, `bottom-label-3`, `label-1`, `label-2`, `label-3`, `no-labels`, and `no-label-1-exp` at the `[edit forwarding-options hash-key family mpls]` hierarchy level
- `disable` and `port-data` statements at the `[edit forwarding-options hash-key family mpls payload ip]` hierarchy level
- `destination-address`, `incoming-interface-index`, `next-header`, `source-address`, `traffic-class` at the `[edit forwarding-options hash-key family inet6 layer-3]` hierarchy level
- `destination-port` and `source-port` at the `[edit forwarding-options hash-key family inet6 layer-4]` hierarchy level

> **NOTE**: On ACX Series routers, configure the `family (inet|inet6)` option with appropriate sub-option at the `[edit forwarding-options enhanced-hash-key` hierarchy. This applies to both EVPN-VXLAN underlay and overlay fabrics.

Options

- *inet*-IPv4 protocol family.
- *mpls*-MPLS protocol family.
- *layer-3*-Incorporate Layer 3 data into the hash key.
- *layer-4*-Incorporate Layer 4 data into the hash key.
- *session-id*-Incorporate session ID data into the hash key (SRX3000 and SRX5000 lines only). The session ID data has higher precedence than the Layer 3 or 4 information.
- *label-1*-Incorporate the first MPLS label into the hash key.
- *label-2*-Incorporate the second MPLS label into the hash key.
- *label-3*-Incorporate the third MPLS label into the hash key.
- *no-labels*-Include no MPLS labels into the hash key.
- *payload*-Incorporate payload data into the hash key.
- *ip*-Include the IP address of the IPv4 or IPv6 payload into the hash key.
- *layer-3–only*-Include only Layer 3 IP information.
- *port-data*-Include the source and destination port field information.
- *source-msb*-Include the most significant byte of the source port.
- *source-lsb*-Include the least significant byte of the source port.
- *destination-msb*-Include the most significant byte of the destination port.
- *destination-lsb*-Include the least significant byte of the destination port.
- *source-mac*-Include source MAC address in hash key.
- *destination-mac*-Include destination MAC address in hash key.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`family multiservice` and `no-label-1-exp` options introduced in Junos OS Release 8.0.

`label-3` and `no-labels` options introduced in Junos OS Release 8.1.

`ether-pseudowire` statement introduced in Junos OS Release 9.1 (M320 and T Series routers only); support extended to M120 and MX Series routers in Junos OS Release 9.4.

`ip`, `label-1`, `label-2`, `layer-3-only`, and `payload` options for the `family multiservice` statement introduced in Junos OS Release 9.4 (M120 and M320 routers only). For MX Series routers, only the `ip` and `payload` statements apply.

`layer-3`, `source-ip-only`, `destination-ip-only`, and `layer-4` statements introduced for the `family multiservice` statement in Junos OS Release 9.5. (MX Series routers only).

`all-labels` and `payload ip disable` statements introduced for the `family mpls` statement in Junos OS Release 12.1X48R2 PTX Series Packet Transport Routers only.

`bottom-label` statement introduced for the `family mpls` statement in Junos OS Release 14.1 for MX Series routers with DPCs (excluding M7i, M10i, and M120.)

`zero-control-word` option for the `ether-pseudowire` statement introduced in Junos OS Release 16.1 for the M Series, MX Series, and PTX Series.

RELATED DOCUMENTATION

Configuring Per-Packet Load Balancing

Configuring Load Balancing Based on MPLS Labels

Configuring Load Balancing Based on MAC Addresses

hash-keys (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 7992
- Hierarchy Level | 7992
- Description | 7992
- Options | 7994
- Required Privilege Level | 7995
- Release Information | 7995

Syntax

```
hash-keys {
    egress-key (destination-ip | source-ip);
    ingress-key (destination-ip | source-ip);
}
```

Hierarchy Level

```
[edit services service-set service-set-name interface-service load-balancing-options]
```

Description

Configure the hash keys used for load balancing in aggregated multiservices (AMS) for service applications (Network Address Translation [NAT], stateful firewall, application-level gateway [ALG], HTTP header enrichment, and mobility). The hash keys supported in the ingress and egress direction are the source IP address and destination IP address.

Hash keys are used to define the load-balancing behavior among the various members in the AMS group. For example, if **hash-keys** is configured as **source-ip**, then the hashing would be performed based on the source IP address of the packet. Therefore, all packets with the same source IP address land on the same member. Hash keys must be configured with respect to the traffic direction: ingress or egress. For example, if **hash-keys** is configured as **source-ip** in the ingress direction, then it should be configured

as **destination-ip** in the egress direction. This is required to ensure that the packets of the same flow reach the same member of the AMS group.

The configuration of the ingress and egress hash keys is mandatory if you are using AMS for NAT. This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen. Refer to Table 48 on page 7993 for the supported hash keys.

The resource-triggered option enables anchor session PICs to use the load or resource information from the anchor services PICs to select the AMS member will anchor the services for the subscriber for load balancing among AMS members. In addition, for mobile subscriber-aware services (such as HTTP header enrichment), you must configure the `resource-triggered` statement, which means that the load balancing is not done using the ingress and egress keys.

Table 48: Hash Keys Supported for AMS for Service Applications

	Service Set at Ingress Interface		Service Set at Egress Interface	
Hash Keys for NAT				
NAT Type	Ingress hash key	Egress hash key	Ingress hash key	Egress hash key
source static	Destination IP address	Source IP address	Source IP address	Destination IP address
source dynamic	Source IP address	Destination IP address	Destination IP address	Source IP address
Network Address Port Translation (NAPT)	Source IP address	Destination IP address	Destination IP address	Source IP address
destination static	Source IP address	Destination IP address	Destination IP address	Source IP address
Hash Keys for Stateful Firewall				
Stateful Firewall	Destination IP address	Source IP address	Destination IP address	Source IP address

Stateful Firewall	Source IP address	Destination IP address	Source IP address	Destination IP address

NOTE: If NAT is used in the service set (along with stateful firewall and ALG), then the hash keys should be based on the NAT type; otherwise, the hash keys of the stateful firewall should be used.

Options

NOTE: The `egress-keys` option is hidden and is deprecated in Junos OS Release 15.1 and later, and is only maintained for backward compatibility. It might be removed completely in a future software release. Load-balancing or steering of traffic occurs, based on the hash keys in the forward direction. Load-balancing of traffic also occurs, based on the hash keys in the reverse direction except in dynamic NAT scenarios (dynamic NAT, NAT64, and NAPT44). For interface-style services, the ingress hash-key is used for the forward direction and the egress hash-key is used for the reverse direction. These hash-keys are configured within the service-set definition by using the `ingress-key` and `egress-key` statements at the `[edit services service-set` *`service-set-name`* `interface-service load-balancing-options]` hierarchy level. For next-hop style services, the ingress hash-key on the inside-domain next-hop is used in the forward direction and the ingress hash-key (not the egress hash-key) on outside-domain next-hop is used for the reverse direction. These hash-keys are configured at the logical AMS interface level by using the `ingress-key` and `egress-key` statements at the `[edit interfaces ams`*`N`*` unit` *`logical-unit-number`* `load-balancing-options hash-keys]` hierarchy level.

ingress-key destination-ip

Use the destination IP address of the flow to compute the hash used in load balancing in the ingress flow direction.

ingress-key source-ip

Use the source IP address of the flow to compute the hash used in load balancing in the ingress flow direction.

egress-key destination-ip

Use the destination IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the egress flow direction. The configuration is mandatory if you are using AMS for Network Address Translation (NAT). This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen.

egress-key source-ip Use the source IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the egress flow direction.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

load-balancing-options (Aggregated Multiservices) | 9888

hash-keys (Interfaces)

IN THIS SECTION

- Syntax | 7995
- Hierarchy Level | 7996
- Description | 7996
- Options | 7996
- Required Privilege Level | 7997
- Release Information | 7997

Syntax

```
hash-keys {
    egress-key (source-ip | destination-ip);
    ingress-key  (source-ip | destination-ip);
```

```
    ipv6-source-prefix-length ipv6-source-prefix-length;
}
```

Hierarchy Level

```
[edit interfaces unit unit-name load-balancing-options]
```

Description

Configure the hash keys used for load balancing in aggregated multiservices (AMS) for next-hop style services. The hash keys supported in the ingress and egress direction are the source IP address and destination IP address.

Hash keys are used to define the load-balancing behavior among the various members in the AMS. For example, if `hash-keys` is configured as `source-ip`, then the hashing is performed based on the source IP address of the packet, so that all packets with the same source IP address land on the same member. When you use `ingress-key` and `egress-key`, you must configure hash keys to take the traffic direction into consideration. For example, if you configure `hash-keys` as `source-ip` in the ingress direction, then you must configure `hash-keys` as `destination-ip` in the egress direction. This is required to ensure that the packets of the same flow reach the same member of the AMS group.

If you are configuring an AMS interface used in a service set for DS-Lite,

The remaining statements are explained separately. See CLI Explorer.

Options

egress-key destination-ip	Use the destination IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the egress flow direction.
egress-key source-ip	Use the source IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the egress flow direction.
ingress-key destination-ip	Use the destination IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the ingress flow direction.
ingress-key source-ip	Use the source IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the ingress flow direction.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

`ipv6-source-prefix-length` option introduced in Junos OS Release 18.2R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card. The `ipv6-source-prefix-length` option is not supported for Next Gen Services.

RELATED DOCUMENTATION

Configuring Load Balancing on AMS Infrastructure

hash-mode

IN THIS SECTION

- Syntax | 7998
- Hierarchy Level | 7998
- Description | 7998
- Default | 7999
- Options | 7999
- Required Privilege Level | 7999
- Release Information | 7999

Syntax

```
hash-mode {
    layer2-header;
    layer2-payload;
    gtp-header-offset offset-value;
}
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key]
```

Description

Select the mode for the hashing algorithm.

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The hash mode that is set using this statement determines which fields are inspected by the hashing algorithm. You must set the hash mode to `layer2-payload` if you want the hashing algorithm to inspect fields in the Layer 2 payload when making hashing decisions. You must set the hash mode to `layer2-header` if you want the hashing algorithm to inspect fields in the Layer 2 header when making hashing decisions.

If the hash mode is set to `layer2-payload`, you can set the fields used by the hashing algorithm to hash IPv4 traffic using the `set forwarding-options enhanced-hash-key inet` statement. You can set the fields used by the hashing algorithm to hash IPv6 traffic using the `set forwarding-options enhanced-hash-key inet6` statement.

If the hash mode is set to `layer2-header`, you can set the fields that the hashing algorithm inspects in the Layer 2 header using the `set forwarding-options enhanced-hash-key layer2` statement.

Default

layer2-payload

Options

`layer-2-payload`	Set the hashing algorithm to use fields in the Layer 2 payload to make hashing decisions.
`layer-2-header`	Set the hashing algorithm to use fields in the Layer 2 header to make hashing decisions.
`gtp-header-offset`	Set proper offset value for the header based on the different patterns.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D15.

Statement is not supported on QFX10002 and QFX 10008 switches.

Hash mode option *gtp-header-offset* introduced in Junos OS Release 17.3R3 for QFX5000 line of switches.

RELATED DOCUMENTATION

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic

enhanced-hash-key (EX) | **6378**

inet (enhanced-hash-key) | **8561**

inet6 (enhanced-hash-key) | **8602**

layer2 (enhanced-hash-key) | **9644**

gtp-header-offset | **7936**

hash-parameters

IN THIS SECTION

- Syntax | 8000
- Hierarchy Level | 8000
- Description | 8001
- Options | 8001
- Required Privilege Level | 8001
- Release Information | 8001

Syntax

```
hash-parameters {
    ecmp {
        function {
            (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
        }
        offset offset;
        preprocess;
    }
    lag {
        function {
            (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
        }
        offset offset;
        preprocess;
    }
}
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key],
[edit routing-instances name forwarding-options enhanced-hash-key]
```

Description

Set traffic hashing parameters for ECMP or LAG traffic.

Options

ecmp
Set hashing parameters for ECMP traffic.

function (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo)
Set the hash functions for ECMP traffic or LAG traffic.

- Values:
  - crc16-bisync—Use CRC16-BISYNC function for ECMP or LAG traffic received on xe.
  - crc16-ccitt—Use the CRC16-CCITT function for ECMP or LAG traffic received on xe.
  - crc32-hi—Use the CRC32-HI function for ECMP or LAG traffic received on xe.
  - crc32-lo—Use the CRC32-LO function for ECMP or LAG traffic received on xe.

lag
Set hashing parameters for LAG traffic.

offset *offset*
Set the hashing offset for ECMP or LAG traffic.

- **Range:** 0 through 199

preprocess
Enable or disable the preprocess parameter for ECMP or LAG traffic.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Aggregated Ethernet Interfaces

hash-seed

IN THIS SECTION

- Syntax | 8002
- Hierarchy Level | 8002
- Description | 8002
- Options | 8003
- Required Privilege Level | 8003
- Release Information | 8003

Syntax

```
hash-seed seed-value;
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key ]
```

Description

Configure a hash seed for load-balancing functions.

By default, the QFX10002, QFX10008, QFX5220, QFX5130, and QFX5700 switches use the system MAC address to generate a hash seed value. You can configure the hash seed value using the **hash-seed** statement at the [`edit forwarding-options enhanced-hash-key`] hierarchy. Set a value between 0 and 4294967295. If you do not configure a hash seed value, the system will generate a hash seed value based on the system MAC address.

NOTE: The `fabric-load-balance` and `user-defined-fields` statements are not supported at the [`edit forwarding-options enhanced-hash-key`] hierarchy level.

Options

`hash-seed` *`seed-value`*	A hash seed value, in the range from 0 to 4294967295.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D30.

RELATED DOCUMENTATION

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

show forwarding-options enhanced-hash-key | **20086**

hash-seed (profile)

IN THIS SECTION

- Syntax | **8004**
- Hierarchy Level | **8004**
- Description | **8004**
- Options | **8004**
- Required Privilege Level | **8004**
- Release Information | **8004**

Syntax

```
hash-seed seed-value
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key profile1/profile2]
```

Description

Configure a hash seed for load-balancing functions for the given load balancing profile (profile1 or profile2).

Options

hash-seed *seed-value* - A hash seed value, in the range from 0 to 4294967295

Required Privilege Level

- interface—To view this statement in the configuration.
- interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2

RELATED DOCUMENTATION

enhanced-hash-key profile1/profile2

load-balance profile1/profile2

hcm (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 8005
- Hierarchy Level | 8006
- Description | 8006
- Required Privilege Level | 8006
- Release Information | 8007

Syntax

```
hcm {
    tag-attribute [tag-attr-name];
    tag-rule rule-name {
        term term-number {
            from {
                destination-address {
                    (any-ipv4 | any-ipv4 except);
                    (any-ipv6 | any-ipv6 except);
                    (any-unicast | any-unicast except);
                    (prefix | prefix except);
                }
                destination-address-range {
                    high address low address <except>;
                }
                destination-port-range {
                    high port-number low port-number;
                }
                destination-ports value;
            }
            then {
                count;
                tag tag-name {
                    encrypt {
                        hash algorithm;
                        prefix hash-prefix;
```

```
                    }
                    ipv4-mask ipv4-mask;
                    ipv6-mask ipv6-mask;
                    ipv4-or-value ipv4-or-value;
                    ipv6-or-value ipv6-or-value;
                    tag-attribute tag-attr-name;
                    tag-header header;
                    tag-separator separator;
                }
            }
        }
    }
    tag-rule-set rule-set-name {
        [rule rule-name];
    }
    profile profile-name {
        tag rule rule-name;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the parameters required to support subscriber-aware HTTP header enrichment.

You can add content to the HTTP headers sent back and forth as part of the client-server exchange for subscribers accessing Web-based services. You configure HTTP header enrichment as a service for a subscriber.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support for Next Gen Services introduced in Junos OS Release 20.2R1 on MX Series routers MX240, MX480 and MX960.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm-profile (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 8007
- Hierarchy Level | 8007
- Description | 8008
- Options | 8008
- Required Privilege Level | 8008
- Release Information | 8008

Syntax

```
hcm-profile hcm-profile-name;
```

Hierarchy Level

```
[edit services service-set]
```

Description

Specify the HTTP header enrichment profile that was configured at the `[edit services hcm]` hierarchy level. This placeholder profile has no configuration options, but it must be specified to enable HTTP header enrichment functionality on the services plane.

Options

hcm-profile-name	Name of the HCM profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 20.2R1 on MX Series routers MX240, MX480 and MX960.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm (HTTP Header Enrichment) | 8005

hcm-profile (PCC Action Profiles)

IN THIS SECTION

- Syntax | 8009
- Hierarchy Level | 8009
- Description | 8009

- Options | 8009
- Required Privilege Level | 8009
- Release Information | 8010

Syntax

```
hcm-profile hcm-profile-name;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name]
```

Description

Specify the HCM profile that you want a PCC action profile to use for determining which HTTP header enrichment rules to apply.

NOTE: This PCC action profile can be used in a PCC rule that only includes `applications` or `application-groups` statements in the `from` statement, and these statements must identify HTTP-based applications. The HCM profile must have been previously configured at the `[edit services hcm]` hierarchy level.

Options

hcm-profile-name	Name of the HCM profile.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support added in Junos 20.2R1 for Next Gen Services on MX240, MX480, and MX960 routers.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

header-integrity-check

IN THIS SECTION

- Syntax | 8010
- Hierarchy Level | 8010
- Description | 8011
- Required Privilege Level | 8011
- Release Information | 8011

Syntax

```
header-integrity-check {
    enable-all;
}
```

Hierarchy Level

```
[edit services service-set service-set service-set-options]
```

Description

Configure Junos OS to verify the packet header for anomalies in IP, TCP, UDP, and ICMP information and to flag such anomalies and errors.

Starting in Junos OS release 17.1R1, the header integrity check on the MS-MPC or MS-MIC drops any packets with header anomalies and includes the following checks:

- ICMP ping of death
- IP unknown protocol
- TCP no flag
- TCP SYN FIN
- TCP FIN no ACK

> **NOTE**: The `header-integrity-check` option that is supported on MS-MICs and MS-MPCs to verify the packet header for anomalies in IP, TCP, UDP, and ICMP information and flag such anomalies and errors has a functionality that is opposite to the functionality caused by passive mode tunneling. If you configure both the `header-integrity-check` statement and the `passive-mode tunneling` statement on MS-MICs and MS-MPCs, and attempt to commit such a configuration, an error is displayed during commit.
>
> The passive mode tunneling functionality (by including the `passive-mode-tunnelin` statement at the `[edit services service-set` *`service-set-name`* `ipsec-vpn-options]` hierarchy level) is a superset of the capability to disable IPsec tunnel endpoint in the traceroute output (by including `no-ipsec-tunnel-in-traceroute` statement at the `[edit services ipsec-vpn]` hierarchy level). Passive mode tunneling also bypasses the active IP checks and tunnel MTU check in addition to not treating an IPsec tunnel as a next-hop as configured by the `no-ipsec-tunnel-in-traceroute` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 13.2.

RELATED DOCUMENTATION

service-set-options | 16746

Configuring Protection Against Network Attacks on an MS-MPC

header-integrity-check (Next Gen Services)

IN THIS SECTION

- Syntax | 8012
- Hierarchy Level | 8012
- Description | 8012
- Required Privilege Level | 8013
- Release Information | 8013

Syntax

```
header-integrity-check {
    enable-all;
}
```

Hierarchy Level

```
[edit services service-set service-set options]
```

Description

Drop packets that have packet header anomalies. These anomalies include:

- Not an IP packet
- Not an IPv4 packet or an IPv6 packet
- TTL error (TTL is 0)

- Bad source/destination IP
- IP checksum error
- Protocol error
- TCP port zero
- TCP header length error (less than 20 bytes)
- TCP SEQNUM is zero and no flags are set
- TCP SEQNUM is zero and flags are set
- No TCP flags are set
- TCP FIN with no Ack
- TCP FIN & Reset
- TCP SYN & (FIN or URG or RESET)
- UDP port zero
- UDP header length error
- ICMP header length error (not within 48-576 bytes)
- ICMP packet error length
- ICMP large packet (1024)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

header-length

IN THIS SECTION

- Syntax | 8014
- Hierarchy Level | 8014
- Description | 8014
- Options | 8014
- Required Privilege Level | 8015
- Release Information | 8015

Syntax

```
header-length {
    match (equal | greater-than | less-than | not-equal);
    value header-length;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify the number of bytes in the TCP header.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`header-length`*—Match the number of bytes in the TCP header.
- **Range:** 0 through 15 bytes

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

header-type

IN THIS SECTION

- Syntax | 8015
- Hierarchy Level | 8015
- Description | 8016
- Options | 8016
- Required Privilege Level | 8016
- Release Information | 8016

Syntax

```
header-type {
    match (equal | greater-than | less-than | not-equal);
    value header-value;
}
```

Hierarchy Level

```
[edit  set security idp custom-attack attack-name attack-type signature protocol ipv6 extension-
header routing-header]
```

Description

Specify the IPv6 routing header type.

Options

match (equal | greater-than | less-than | not-equal)	Match an operand.
value	Match a decimal value. • **Range:** 0 through 255

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

health-monitor

IN THIS SECTION

- Syntax | 8017
- Hierarchy Level | 8017
- Description | 8017
- Required Privilege Level | 8017
- Release Information | 8017

Syntax

```
health-monitor {
    falling-threshold percentage;
    interval seconds;
    rising-threshold percentage;
    idp {
        falling-threshold percentage;
        interval seconds;
        rising-threshold percentage;
    }
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Configure health monitoring.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

Option `idp` introduced in Junos OS Release 9.3 for all supported platforms.

RELATED DOCUMENTATION

Configuring Health Monitoring on Devices Running Junos OS

health-monitor (KHMS)

IN THIS SECTION

- Syntax | 8018
- Hierarchy Level | 8018
- Description | 8018
- Options | 8019
- Required Privilege Level | 8020
- Release Information | 8020

Syntax

```
health-monitor {
    ifstate-clients {
        (peer-stuck | non-peer-stuck | all-clients-stuck);
    }
    threshold-level {
        (low | medium | high);
    }
    action {
        (alarm | alarm-with-cores | restart);
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the kernel health monitoring system (KHMS). The KHMS is used to detect and take action on stuck ifstate clients. Stuck ifstate clients can affect system performance. With this configuration

statement you can configure a time interval in which the system can detect a stuck ifstate client. You also can configure the action the system takes when it finds a stuck ifstate client.

The ifstate clients receive states from the kernel. There are two kinds of ifstate clients:

- non-peer clients (for example, some daemons, or processes, on the Routing Engine)—ifstate non-peer clients open connections between programs and read states from or write states to the kernel.
- peer clients (for example, FPCs)—ifstate peer clients read peer messages and send updates to the peers.

An ifstate client is stuck if the kernel sends a message and the ifstate client does not send back an ACK. A rt_pfe_veto condition is a log message that indicates that states are sent but no ACK comes back. However, the system won't take the configured action until the configured time interval times outs, in case the ACK comes late.

Options

ifstate-clients

Configure which ifstate clients you want to monitor and manage. There are three options:

- `peer-stuck`—Monitor and manage stuck peers.
- `non-peer-stuck`—Monitor and manage stuck processes.
- `all-clients-stuck`—Monitor and manage both stuck peers and stuck processes.

threshold-level

Configure the time interval in which to detect if a given ifstate client is stuck:

- `high`—540 seconds
- `medium`—360 seconds; this is the default.
- `low`—180 seconds

action

Configure the action to be taken on the stuck ifstate client once the configured time interval times out.

- `alarm`—Only an alarm will be raised about the stuck ifstate client; this is the default.
- `alarm-with-cores`—An alarm will be raised about the stuck ifstate client after collecting live cores from the primary Routing Engine kernel and the stuck peer.

> **CAUTION**: In the case of a stuck peer, collecting lives cores might result in the component being restarted or rebooted.

- `restart`—The stuck ifstate client will be disconnected after collecting live cores from the primary Routing Engine kernel and the stuck peer (depending on supportability).

> **CAUTION**: When choosing this action, be aware of the implications of restarting an ifstate client. For example, some FPCs don't simply restart; they reboot.

Required Privilege Level

admin

Release Information

Statement introduced in Junos OS Release 16.1R1.

heartbeat-address (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 8021
- Hierarchy Level | 8021
- Description | 8021
- Options | 8021
- Required Privilege Level | 8021
- Release Information | 8021

Syntax

```
heartbeat-address (ip-address | ipv6-address);
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Configure an IPv4 address or IPv6 address used to create an IP-based packet connection, known as a *heartbeat connection*, between the primary router and backup router in an MX Series Virtual Chassis. To ensure consistent access to the primary Routing Engine in the Virtual Chassis primary router (VC-Pp) regardless of which Routing Engine is active, you must configure the address for the management Ethernet interface (`fxp0`) with the `master-only` statement at the `[edit groups]` hierarchy level.

NOTE: The `heartbeat-address` statement and the `no-split-detection` statement at the `[edit virtual-chassis]` hierarchy level are mutually exclusive. As a result, the software prevents you from configuring both `heartbeat-address` and `no-split-detection` at the same time. If you attempt to do so, the software displays an error message and causes the `commit` operation to fail.

Options

ip-address IPv4 `master-only` address for the `fxp0` management interface.

ipv6-address IPv6 `master-only` address for the `fxp0` management interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in the Same Subnet

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in Different Subnets

heartbeat-interval

IN THIS SECTION

- Syntax | 8022
- Hierarchy Level | 8022
- Description | 8022
- Options | 8023
- Required Privilege Level | 8023
- Release Information | 8023

Syntax

```
heartbeat-interval milliseconds;
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Set the interval between the periodic signals broadcast to the devices in a chassis cluster to indicate that the active node is operational.

The `heartbeat-interval` option works in combination with the `heartbeat-threshold` option to define the wait time before failover is triggered in a chassis cluster. The default values of these options produce a wait

time of 3 seconds. In a large configuration approaching full capacity on an SRX5400 or SRX5600 or SRX5800 device, however, we recommend that you increase the failover wait time to 5 seconds.

For example, a `heartbeat-threshold` of 3 and a `heartbeat-interval` of 1000 milliseconds result in a total wait of 3 seconds before failover is triggered. To increase this wait to 5 seconds, you could increase the `heartbeat-threshold`, the `heartbeat-interval`, or both. A `heartbeat-threshold` of 5 and a `heartbeat-interval` of 1000 milliseconds would yield a wait time of 5 seconds. Setting the `heartbeat-threshold` to 4 and the `heartbeat-interval` to 1250 milliseconds would also yield a wait time of 5 seconds.

> **NOTE**: In a chassis cluster scaling environment, the `heartbeat-threshold` must always be set to 8.

Options

`milliseconds`—Time interval between any two heartbeat messages.

- **Range:** 1000 through 2000 milliseconds
- **Default:** 1000 milliseconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9. Statement updated in Junos OS Release 10.4.

RELATED DOCUMENTATION

cluster (Chassis) | **4369**

heartbeat-threshold

IN THIS SECTION

- Syntax | 8024
- Hierarchy Level | 8024
- Description | 8024
- Options | 8025
- Required Privilege Level | 8025
- Release Information | 8025

Syntax

```
heartbeat-threshold number;
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Set the number of consecutive missed heartbeat signals that a device in a chassis cluster must exceed to trigger failover of the active node.

The `heartbeat-threshold` option works in combination with the `heartbeat-interval` option to define the wait time before failover is triggered in a chassis cluster. The default values of these options produce a wait time of 3 seconds. In a large configuration approaching full capacity on an SRX5400 or SRX5600 or SRX5800 device, however, we recommend that you increase the failover wait time to 5 seconds.

For example, a `heartbeat-threshold` of 3 and a `heartbeat-interval` of 1000 milliseconds result in a total wait of 3 seconds before failover is triggered. To increase this wait to 5 seconds, you could increase the `heartbeat-threshold`, the `heartbeat-interval`, or both. A `heartbeat-threshold` of 5 and a `heartbeat-interval` of 1000 milliseconds would yield a wait time of 5 seconds. Setting the `heartbeat-threshold` to 4 and the `heartbeat-interval` to 1250 milliseconds would also yield a wait time of 5 seconds.

Options

number—Number of consecutive missed heartbeats.

- **Range:** 3 through 8
- **Default:** 3

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0. Statement updated in Junos OS Release 10.4.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

heartbeat-timeout (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 8026
- Hierarchy Level | 8026
- Description | 8026
- Options | 8026
- Required Privilege Level | 8026
- Release Information | 8027

Syntax

```
heartbeat-timeout seconds;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Configure the maximum time period within which a Virtual Chassis member router must respond to a heartbeat packet sent by the other member router. When an adjacency disruption or split is detected in the Virtual Chassis configuration, each member router sends a single heartbeat packet to the other member router to determine whether that router is still operating and able to respond. If the other member router responds to the heartbeat packet within the configured or default timeout period, the heartbeat connection helps prevent the member routers from changing primary role roles, which can cause undesirable results.

> **NOTE**: Both member routers forming the heartbeat connection must be assigned the `routing-engine` role in the preprovisioned Virtual Chassis configuration. The `routing-engine` role makes the router eligible to function as either the primary router or backup router of the Virtual Chassis.

Options

seconds—Number of seconds within which a Virtual Chassis member router must respond to a heartbeat packet.

- **Range:** 1 through 60 seconds
- **Default:** 2 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in the Same Subnet

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in Different Subnets

heartbeat-tos (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 8027
- Hierarchy Level | 8027
- Description | 8028
- Options | 8028
- Required Privilege Level | 8029
- Release Information | 8029

Syntax

```
heartbeat-tos value;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Configure a network-specific value for the header field that defines either the type of service (ToS), traffic class (TC), or differential services (DS) octet that is transmitted in the header of a heartbeat packet.. This field is a single octet that ensures preferential delivery of heartbeat packets between Virtual Chassis member routers in a congested network. You can set the IP precedence portion of the ToS header field or the differentiated services code point (DSCP) in the DS header field in accordance with your network. When you configure a `heartbeat-tos` value, the software uses this octet value in the IP header for all generated Virtual Chassis heartbeat packets.

Configuring the `heartbeat-tos` Virtual Chassis heartbeat packets is particularly useful in congested networks or in networks with high bandwidth demands. For example, setting the appropriate `heartbeat-tos` value for your network design enables you to prioritize the traffic flows that transmit Virtual Chassis heartbeat packets at a higher service level than the traffic flows for traditional IP applications such as e-mail or Web browsing. During periods of network congestion, configuring the appropriate value can ensure reliable delivery of the heartbeat packets and decrease the likelihood of packet loss.

BEST PRACTICE: To properly prioritize the traffic flows for heartbeat packets based on the network load, make sure you are using a value that is consistent with the protocol standard that was implemented in your network.

Options

`value`—Network-specific value of the ToS, TC, or DS octet transmitted in Virtual Chassis heartbeat packets. Depending on your network design and protocol family implementation, the value you specify must conform to the following RFC standards:

- IPv4 ToS as defined in RFC 791, *INTERNET PROTOCOL - DARPA INTERNET PROGRAM PROTOCOL SPECIFICATION*
- IPv6 traffic class (TC) as defined in RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification*
- IPv4 and IPv6 Differentiated Services as defined in RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
- Assured forwarding per-hop behavior (PHB) as defined in RFC 2597, *Assured Forwarding PHB Group*
- RFC 3260, *New Terminology and Clarifications for Diffserv*
- **Range:** 0 through 255
- **Default:** 0 (`heartbeat-tos` not set)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in the Same Subnet

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in Different Subnets

Configuring a Virtual Chassis Heartbeat Connection

hello-acknowledgements

IN THIS SECTION

- Syntax | 8029
- Hierarchy Level | 8030
- Description | 8030
- Default | 8030
- Required Privilege Level | 8030
- Release Information | 8030

Syntax

```
hello-acknowledgements;
```

Hierarchy Level

```
[edit logical-systems logical-systems-name protocols rsvp],
[edit protocols rsvp]
```

Description

Enable hello messages from nonsession neighbors to be acknowledged with a hello acknowledgment message. Once hello acknowledgments are enabled, the router continues to acknowledge hello messages from any nonsession RSVP neighbors unless the interface itself goes down or the configuration is changed by an administrator.

Default

Hello acknowledgments are disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Hello Acknowledgments for Nonsession RSVP Neighbors

hello-dead-interval

IN THIS SECTION

Syntax | 8031

- Hierarchy Level | 8031
- Description | 8031
- Options | 8031
- Required Privilege Level | 8031
- Release Information | 8032

Syntax

```
hello-dead-interval milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-protocol],
[edit protocols link-management peer peer-name lmp-protocol]
```

Description

Specify how long the Link Management Protocol (LMP) waits before declaring the control channel to be dead. This is an interval during which the router receives no LMP hello packets from the neighbor on a control that is active or up.

Options

milliseconds—Interval to wait before declaring the control channel to be dead.

- **Range:** 500 through 300,000
- **Default:** 500 milliseconds (three times the hello interval)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring Hello Message Intervals for LMP Control Channels

hello-interval (LMP) | 8035

hello-interval

IN THIS SECTION

- Syntax | 8032
- Hierarchy Level | 8032
- Description | 8033
- Options | 8033
- Required Privilege Level | 8033
- Release Information | 8033

Syntax

```
hello-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name],
[edit protocols pim interface (Protocols PIM)  interface-name],
```

```
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM) interface-name]
```

Description

Specify how often the routing device sends PIM hello packets out of an interface.

Options

seconds—Length of time between PIM hello packets.

- **Range:** 0 through 255
- **Default:** 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

hold-time (Protocols PIM) | **8114**

Modifying the PIM Hello Interval

hello-interval (L2TP)

IN THIS SECTION

Syntax | **8034**

Hierarchy Level | **8034**

- Description | 8034
- Options | 8034
- Required Privilege Level | 8034
- Release Information | 8034

Syntax

```
hello-interval seconds;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Specify the keepalive timer for L2TP tunnels.

Options

seconds—Interval, in seconds, after which the server sends a hello message if no messages are received. A value of `0` means that no hello messages are sent.

- **Range:** 0 through 3600
- **Default:** 60 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Timers for L2TP Tunnels

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

hello-interval (LMP)

IN THIS SECTION

- Syntax | 8035
- Hierarchy Level | 8035
- Description | 8035
- Options | 8035
- Required Privilege Level | 8036
- Release Information | 8036

Syntax

```
hello-interval milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-protocol],
[edit protocols link-management peer peer-name lmp-protocol]
```

Description

Specify how often the router sends Link Management Protocol (LMP) hello packets.

Options

milliseconds—Length of time between hello packets.

- **Range:** 150 through 300,000
- **Default:** 150 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring Hello Message Intervals for LMP Control Channels

hello-dead-interval | 8030

hello-interval (Protocols LDP)

IN THIS SECTION

- Syntax | 8036
- Hierarchy Level | 8037
- Description | 8037
- Options | 8037
- Required Privilege Level | 8037
- Release Information | 8037

Syntax

```
hello-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp interface interface-name],
[edit logical-systems logical-system-name protocols ldp targeted-hello],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp
interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp
targeted-hello],
[edit protocols ldp interface interface-name],
[edit protocols ldp targeted-hello],
[edit routing-instances routing-instance-name protocols ldp interface interface-name],
[edit routing-instances routing-instance-name protocols ldp targeted-hello]
```

Description

Control the LDP timer that regulates how often hello messages are sent. You can control the rate both link hello messages and targeted hello messages are sent depending on the hierarchy level at which you configure the `hello-interval` statement.

Options

seconds—Length of time between transmission of hello packets.

- **Range:** 1 through 65,535 seconds
- **Default:** 5 seconds for link hello messages, 15 seconds for targeted hello messages

> **NOTE**: The default hello intervals on EX 9200 switches are not sufficient to avoid protocol flapping during an ISSU Linux upgrade. To avoid protocol flapping during the upgrade, configure the hello intervals for 35 seconds or more.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for LDP targeted hellos added in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring the LDP Timer for Hello Messages

hello-interval (Protocols RSVP)

IN THIS SECTION

- Syntax | 8038
- Hierarchy Level | 8038
- Description | 8038
- Options | 8039
- Required Privilege Level | 8039
- Release Information | 8039

Syntax

```
hello-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit logical-systems logical-system-name protocols rsvp peer-interface peer-interface-name],
[edit protocols rsvp interface interface-name],
[edit protocols rsvp peer-interface peer-interface-name]
```

Description

Enable the sending of hello packets on the interface.

Options

seconds—Length of time between hello packets. A value of 0 disables the sending of hello packets on the interface.

- **Range:** 1 through 60 seconds
- **Default:** 9 seconds

> **NOTE**: The default hello intervals on EX 9200 switches are not sufficient to avoid protocol flapping during an ISSU Linux upgrade. To avoid protocol flapping during the upgrade, configure the hello intervals for 35 seconds or more.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the RSVP Hello Interval

hello-message (NETCONF)

IN THIS SECTION

- Syntax | 8040
- Hierarchy Level | 8040
- Description | 8040
- Default | 8040
- Options | 8040

Required Privilege Level | 8041
Release Information | 8041

Syntax

```
hello-message {
    yang-module-capabilities {
        advertise-native-yang-modules;
        advertise-custom-yang-modules;
        advertise-standard-yang-modules;
    }
}
```

Hierarchy Level

```
[edit system services netconf]
```

Description

Configure the NETCONF server to advertise additional capabilities in the NETCONF capabilities exchange. The NETCONF server emits the `<hello>` element and the `<capabilities>` child element with the list of capabilities at the start of the NETCONF session.

Default

If you do not include the `hello-message` statement, the NETCONF server does not advertise any additional capabilities beyond the default capabilities in the NETCONF capabilities exchange.

Options

`yang-module-capabilities` Configure the YANG modules that the NETCONF server advertises in the NETCONF capabilities exchange. The following statements are supported:

- `advertise-custom-yang-modules`—Advertise third-party YANG modules installed on the device.

- `advertise-native-yang-modules`—Advertise Junos OS native YANG modules.
- `advertise-standard-yang-modules`—Advertise standard YANG modules supported by the device, for example, OpenConfig modules.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1 and Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

Starting the NETCONF Session

hello-time

IN THIS SECTION

- Syntax | 8042
- Hierarchy Level | 8042
- Description | 8042
- Default | 8042
- Options | 8042
- Required Privilege Level | 8043
- Release Information | 8043

Syntax

```
hello-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp)],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id],
```

```
[edit protocols (mstp | rstp)],
[edit protocols vstp vlan vlan-id],
```

```
[edit routing-instances routing-instance-name protocols (mstp | rstp)],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id]
```

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specifies the number of seconds between transmissions of configuration BPDUs by the root bridge.

Default

2 seconds

Options

seconds—(Optional) Number of seconds between transmissions of configuration BPDUs.

- **Range:** 1 through 10
- **Default:** 2 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding MSTP

Understanding VSTP

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

hello-timer

IN THIS SECTION

- Syntax | **8044**
- Hierarchy Level | **8044**
- Description | **8044**
- Options | **8044**
- Required Privilege Level | **8044**
- Release Information | **8044**

Syntax

```
hello-timer milliseconds;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

For link services interfaces only, configure the rate at which hello messages are sent. A hello message is transmitted after a period defined in milliseconds has elapsed.

Options

milliseconds—The rate at which hello messages are sent.

- **Range:** 1 through 180 milliseconds
- **Default:** 10 milliseconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Acknowledgment Timers on Link Services Physical Interfaces

address (Interfaces) | **1878**

acknowledge-timer | **1752**

help

IN THIS SECTION

- Syntax | 8045
- Description | 8045
- Options | 8045
- Required Privilege Level | 8046
- Release Information | 8046

Syntax

```
help < (apropos string | reference statement-name| syslog syslog-tag| tip cli number | topic
word)>
```

Description

Display help about available operational commands, configuration statements, or general information about getting help. Entering the `help` command without an option provides introductory information about how to use the `help` and `?` commands.

Options

`apropos` *string*—(Optional) Display command names and help text that matches the string specified. If the string contains spaces, enclose it in quotation marks (" "). You can also specify a regular expression for the string, using standard UNIX-style regular expression syntax.

`reference` *statement-name*—(Optional) Display summary information for a configuration statement. This information is based on summary descriptions that appear in the Junos configuration guides.

`syslog` *syslog-tag*—(Optional) Display information about system log messages.

`tip cli` *number*—(Optional) Display a tip about using the CLI. Specify the number of the tip you want to view.

`topic` *word*—(Optional) Display usage guidelines for a topic or configuration statement. This information is based on subjects that appear in the Junos configuration guides.

Required Privilege Level

None

Release Information

Command introduced before Junos OS Release 7.4.

The `apropos` option was added in Junos OS Release 8.0.

RELATED DOCUMENTATION

Getting Online Help from the Junos OS Command-Line Interface

helper-disable (LDP)

IN THIS SECTION

- Syntax | 8046
- Hierarchy Level | 8047
- Description | 8047
- Default | 8047
- Required Privilege Level | 8047
- Release Information | 8047

Syntax

```
helper-disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp
graceful-restart],
[edit protocols ldp graceful-restart],
[edit routing-instances routing-instance-name protocols ldp graceful-restart]
```

Description

Disable helper mode for LDP graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart LDP.

Default

Helper mode is enabled by default on all routing protocols (including LDP) that support graceful restart.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring LDP Graceful Restart

helper-disable (Multiple Protocols)

IN THIS SECTION

- Syntax | 8048
- Hierarchy Level | 8048
- Description | 8048
- Default | 8048
- Required Privilege Level | 8049
- Release Information | 8049

Syntax

```
helper-disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (isis | ldp | ospf | ospf3 | rsvp) graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ldp | ospf | ospf3) graceful-restart],
[edit protocols (isis | ldp | ospf | ospf3 | rsvp) graceful-restart],
[edit routing-instances routing-instance-name protocols (ldp | ospf | ospf3) graceful-restart]
```

Description

Disable helper mode for graceful restart. When helper mode is disabled, a router or switch cannot help a neighboring router that is attempting to restart.

Default

Helper mode is enabled by default for these supported protocols: IS-IS, LDP, OSPF/OSPFv3, and RSVP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routing Protocols Graceful Restart

Configuring Graceful Restart for MPLS-Related Protocols

helpers

IN THIS SECTION

- Syntax | 8049
- Hierarchy Level | 8051
- Description | 8051
- Required Privilege Level | 8052
- Release Information | 8052

Syntax

```
helpers {
    bootp {
        client-response-ttl number;
        description text-description;
        interface interface-group {
            client-response-ttl number;
            description text-description;
```

```
            maximum-hop-count number;
            minimum-wait-time seconds;
            no-listen;
            server address {
                logical-system logical-system-name <routing-instance [ <default> routing-
instance-names ]>;
                routing-instance [ <default> routing-instance-names ];
            }
        }
        maximum-hop-count number;
        minimum-wait-time seconds;
        relay-agent-option;
        server address {
            logical-system logical-system-name <routing-instance [ <default> routing-instance-
names ]>;
            routing-instance [ <default> routing-instance-names ];
        }
    }
    domain {
        description text-description;
        interface interface-name {
            broadcast;
            description text-description;
            no-listen;
            server address <logical-system logical-system-name> <routing-instance routing-
instance-name>;
        }
        server address <logical-system logical-system-name> <routing-instance routing-instance-
name>;
    }
    port (Packet Forwarding) port-number {
        description text-description;
        interface interface-name {
            broadcast;
            description text-description;
            no-listen;
            server address <logical-system logical-system-name> <routing-instance routing-
instance-name>;
        }
        server address <logical-system logical-system-name> <routing-instance routing-instance-
name>;
    }
    tftp {
```

```
        description text-description;
        interface interface-name {
            broadcast;
            description text-description;
            no-listen;
            server address <logical-system logical-system-name> <routing-instance routing-
instance-name>;
        }
        server address <logical-system logical-system-name> <routing-instance routing-instance-
name>;
    }
    traceoptions {
        file filename  <files number> <match regular-expression> <size bytes> <world-readable | no-
world-readable>;
        flag flag;
        level level;
        no-remote-trace level;
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Enable TFTP or DNS request packet forwarding, or configure the router, switch, or interface to act as a DHCP/BOOTP relay agent. Use only one server address per interface or global configuration.

You can also use the helpers `port` statement to enable forwarding LAN broadcast traffic on custom UDP ports to particular servers as unicast traffic. Configure the UDP port number and optionally an interface on which to listen for broadcast traffic, and the destination server address to receive that traffic, as shown in either of the following sample configurations:

```
user@ host# show forwarding-options
helpers {
    port 1200 {
        server 10.20.30.40;
```

```
        }
    }
```

```
user@ host# show forwarding-options
helpers {
    port 3000 {
        interface {
            fe-0/0/1.0 {
                server 192.0.2.2;
            }
        }
    }
    port 3001 {
        interface {
            fe-0/0/0.0 {
                server 192.0.2.2;
            }
        }
    }
}
```

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding

Configuring Port-based LAN Broadcast Packet Forwarding

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

hide-avps

IN THIS SECTION

- Syntax | 8053
- Hierarchy Level | 8053
- Description | 8053
- Default | 8053
- Required Privilege Level | 8053
- Release Information | 8054

Syntax

```
hide-avps;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Hide L2TP attribute-value pairs if the secret shared between the two ends of the tunnel is known.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Default

Attribute-value pairs that can be hidden are exposed, even if the secret information is known.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Hiding Attribute-Value Pairs for L2TP Tunnels

hide-via-headers

IN THIS SECTION

- Syntax | 8054
- Hierarchy Level | 8054
- Description | 8055
- Required Privilege Level | 8055
- Release Information | 8055

Syntax

```
hide-via-headers disable;
```

Hierarchy Level

```
[edit security alg sip ]
```

Description

The configuration `set security alg sip hide-via-header disable` controls the stripping of via headers in SIP request messages. By default, the via headers are stripped. If the configuration is set to enable, only the first via header will receive NAT translation from the SIP ALG.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

hierarchical-policer

IN THIS SECTION

- Syntax | 8055
- Hierarchy Level | 8056
- Description | 8056
- Options | 8056
- Required Privilege Level | 8056
- Release Information | 8056

Syntax

```
hierarchical-policer name {

aggregate {
    if-exceeding {
        bandwidth-limit bandwidth;
        burst-size-limit burst;
```

```
    }
    then {
        discard;
    }
}
premium {
    if-exceeding {
        bandwidth-limit bandwidth;
        burst-size-limit burst;
    }
    then {
        discard;
    }
}
```

Hierarchy Level

```
[edit firewall]
```

Description

For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, specify a hierarchical policer.

Options

Options are described separately.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Applying Policers

Junos OS Class of Service User Guide for Routing Devices

hierarchical-policer (aggregate or premium)

IN THIS SECTION

- Syntax (M Series, MX Series, T Series - Bandwidth-Based) | 8057
- Syntax (MX Series - Packets-Per-Second (pps)-Based) | 8058
- Hierarchy Level | 8058
- Description | 8058
- Options | 8059
- Required Privilege Level | 8059
- Release Information | 8059

Syntax (M Series, MX Series, T Series - Bandwidth-Based)

```
hierarchical-policer hierarchical-policer-name | uid {
    aggregate {
        if-exceeding {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        then {
            discard;
        }
    }
    premium {
        if-exceeding {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        then {
```

```
            discard;
        }
    }
}
```

Syntax (MX Series - Packets-Per-Second (pps)-Based)

```
hierarchical-policer hierarchical-policer-name | uid {
    aggregate {
        if-exceeding-pps {
            pps-limit pps;
            packet-burst packets;
        }
        then {
            discard;
        }
    }
    premium {
        if-exceeding-pps (Hierarchical Policer) {
            pps-limit (Hierarchical Policer) pps;
            packet-burst (Hierarchical Policer) packets;
        }
        then {
            discard;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall],
[edit firewall]
```

Description

Use a hierarchical policer to rate-limit ingress Layer 2 traffic at a physical or logical interface and apply different policing actions based on whether the packets are classified as `premium` for expedited forwarding

(EF) or `aggregate` for a lower priority. The two policers defined within the hierarchical policer are `aggregate` and `premium`.

Hierarchical policers are supported on Enhanced Intelligent Queuing (IQE) PICs and SONET interfaces hosted on the M120 and M320 with incoming Flexible PIC Concentrators (FPCs) as SFPC and outgoing FPCs as FFPC; on MPCs hosted on MX Series routers; on the T320, T640, and T1600 with Enhanced Intelligent Queuing (IQE) PICs; and on the T4000 with Type 5 FPC and Enhanced Scaling Type 4 FPC.

NOTE:

- The `if-exceeding-pps` statement is only supported on MX Series routers with MPCs.
- The `if-exceeding` and `if-exceeding-pps` statements are mutually exclusive and, therefore, cannot be applied at the same time.

You can configure the policer in static firewall filters or dynamic firewall filters in a dynamic client profile or a dynamic service profile.

Options

`hierarchical-policer-name`—Name that identifies the policer. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose the name in quotation marks (“ ”).

`uid`—When you configure a hierarchical policer at the `[edit dynamic-profiles profile name firewall]` hierarchy level, you must assign a variable UID as the policer name.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles profile-name firewall]` hierarchy level introduced in Junos OS Release 11.4.

Support for `if-exceeding-pps` statement on MX Series routers with MPCs introduced in Junos OS Release 15.2.

RELATED DOCUMENTATION

- Hierarchical Policer Configuration Overview
- Hierarchical Policers
- aggregate (Hierarchical Policer) | 2064
- bandwidth-limit (Hierarchical Policer)
- burst-size-limit (Hierarchical Policer) | 3084
- pps-limit (Hierarchical Policer)
- packet-burst (Hierarchical Policer)
- if-exceeding (Hierarchical Policer) | 8371
- if-exceeding-pps (Hierarchical Policer)
- premium (Hierarchical Policer) | 13110

hierarchical-scheduler

IN THIS SECTION

- Syntax | 8060
- Hierarchy Level | 8061
- Description | 8061
- Options | 8061
- Required Privilege Level | 8062
- Release Information | 8062

Syntax

```
hierarchical-scheduler {
    implicit-hierarchy;
```

```
    maximum-hierarchy-levels number;
}
```

Hierarchy Level

```
[edit interfaces]
[edit dynamic-profiles name interfaces],
[edit dynamic-profiles name logical-systems name interfaces]
```

Description

Enable the use of hierarchical schedulers. If you do not include this statement, the interfaces on the router cannot use hierarchical interfaces.

NOTE: When configuring hierarchical scheduling for network slices on an interface, you do not need to enable vlan-tagging on that interface.

NOTE: To enable hierarchical scheduling on MX80 and MX104 routers, configure the `hierarchical-scheduler` statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.

Options

implicit-hierarchy For MPC/MIC subscriber interfaces and interface sets running over aggregated Ethernet on MX Series routers, form a hierarchical relationship between the CoS scheduler nodes at level 1, level 2, and level 3.

maximum-hierarchy-levels *number* Set the maximum number of hierarchy levels.

- **Range:** 2 through 4
- **Default:** 3

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Hierarchical Schedulers for CoS

Understanding Hierarchical CoS for Subscriber Interfaces

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) | **8062**

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers)

IN THIS SECTION

- Syntax | **8062**
- Hierarchy Level | **8063**
- Description | **8063**
- Options | **8063**
- Required Privilege Level | **8064**
- Release Information | **8064**

Syntax

```
hierarchical-scheduler {
    implicit-hierarchy;
    maximum-hierarchy-levels number;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure hierarchical scheduling options on the interface.

The statement is supported on the following interfaces:

- MIC and MPC interfaces in MX Series routers
- GRE tunnel interfaces configured on physical interfaces hosted on MIC or MPC line cards in MX Series routers

To enable hierarchical scheduling on MX Series routers, configure the `hierarchical-scheduler` statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.

Options

`implicit-hierarchy`

Configure four-level hierarchical scheduling. When you include the `implicit-hierarchy` option, a hierarchical relationship is formed between the CoS scheduler nodes at level 1, level 2, level 3, and level 4. The `implicit-hierarchy` option is supported only on MPC/MIC subscriber interfaces and interface sets on MX Series routers.

`maximum-hierarchy-levels` *`number`*

Specify the maximum number of hierarchical scheduling levels allowed for node scaling, from 2 through 4 levels. The default number of levels is 3. The `maximum-hierarchy-levels` option is supported on MPC/MIC or EQ DPC subscriber interfaces and interface sets on MX Series routers.

- If you set `maximum-hierarchy-levels` to `2`, interface sets are not allowed. In this case, if you configure a level 2 interface set, you generate Packet Forwarding Engine errors.
- If you do not include the `maximum-hierarchy-levels` option, keeping the default number of hierarchy levels at 3, interface sets can be at either level 2 or level 3, depending on whether the member logical interfaces within the interface set have a traffic control profile. If any member logical interface has a traffic control profile, then the interface set is a level 2 CoS scheduler node. If no member logical interface has a traffic control profile, the interface set is at level 3.

> **CAUTION**: MPC3E, 32x10GE MPC4E, and 2x100GE + 8x10GE MPC4E MPCs support only two levels of scheduling hierarchy. When enabling hierarchical scheduling on these cards, you must explicitly set `maximum-hierarchy-levels` to `2`.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`implicit-hierarchy` option added in Junos OS Release 13.1.

Support on GRE tunnel interfaces configured on physical interfaces on MICs or MPCs in MX Series routers added in Junos OS Release 13.3.

Support for up to four hierarchy levels added in Junos OS Release 16.1.

RELATED DOCUMENTATION

- Understanding Hierarchical CoS for Subscriber Interfaces
- Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links
- Configuring Hierarchical Schedulers for CoS
- Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface
- Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview

high-availability (Chassis)

IN THIS SECTION

Syntax | **8065**

- Hierarchy Level | 8067
- Description | 8067
- Options | 8068
- Required Privilege Level | 8068
- Release Information | 8068

Syntax

```
high-availability {
    local-id id local-ip local-ip;
    peer-id name {
        desc desc;
        interface interface;
        liveness-detection {
            no-adaptation;
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier multiplier;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (0 | 1 | automatic);
        }
        peer-forwarding-ip loopback-address {
            interface interface-name;
            liveness-detection {
                no-adaptation;
                detection-time {
                    threshold milliseconds;
                }
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier multiplier;
                transmit-interval {
```

```
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
            version (0 | 1 | automatic);
        peer-ip peer-ip;
        routing-instance routing-instance;
        vpn-profile vpn-profile;
    }
    services-redundancy-group name {
        active-signal-route {
            ip-address;
            routing-instance routing-instance;
        }
        activeness-priority activeness-priority;
        activeness-probe {
            dest-ip {
                ip-address;
                routing-instance routing-instance;
            }
        }
        backup-signal-route {
            ip-address;
            routing-instance routing-instance;
        }
        monitor {
            bfd-liveliness name {
                interface interface;
                routing-instance routing-instance;
                session-type (multihop | singlehop);
                src-ip src-ip;
            }
            interface name;
            ip name {
                routing-instance routing-instance;
            }
        }
        monitor-object name {
    bfd-liveliness {
         destination-ip name {
      interface interface;
      routing-instance routing-instance;
      session-type(multihop | singlehop);
      src-ip src-ip;
```

```
      weight weight;
     }
     threshold threshold;
    }
    interface {
     interface-name name {
      weight weight;
     }
     threshold threshold;
    }
    ip {
     destination-ip name {
      routing-instance routing-instance;
      weight weight;
     }
     threshold threshold;
    }
    object-threshold object-threshold;
   }
        peer-id id;
        process-packet-on-backup;
        shutdown-on-failure name;
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        level (alert | all | critical | debug | emergency | error | info | notice | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure Multinode High Availability options on SRX Series Firewalls.

Options

"local-id" on page 9948	Define local node identifier.
"peer-id" on page 12367	Define peer node related information.
"services-redundancy-group" on page 16793	Define the services redundancy group details.
"traceoptions" on page 29471	Define high availability traceoptions.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Multinode High Availability

Example: Configure Multinode High Availability in a Layer 3 Network

Example: Configure Multinode High Availability in a Default Gateway Deployment

Example: Configure Multinode High Availability in a Default Gateway Deployment

high-availability-options (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 8069
- Hierarchy Level | 8069
- Description | 8069
- Required Privilege Level | 8069
- Release Information | 8070

Syntax

```
high-availability-options {
    (many-to-one | one-to-one) {
        preferred-backup preferred-backup;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options]
```

Description

Configure the high availability options for the aggregated multiservices (AMS) interface. For service applications, if only the load-balancing feature is being used, then this configuration is optional.

For many-to-one (N:1) high availability support for service applications like Network Address Translation (NAT), the preferred backup services PIC, in hot standby mode, backs up one or more (N) active services PICs.

> **NOTE**: In both cases, if one of the active services PICs goes down, then the backup replaces it as the active PIC. When the failed PIC comes back up, it becomes the new backup. This is called *floating backup*.

One-to-one (1:1) high availability support associates a single backup interface with a single active interface. 1:1 configuration is supported only on the MS-MPC and MX-SPC3. In 1:1 (stateful) configurations, synchronization causes the active and back up PICs to synchronize traffic states and data structures, preventing data loss during a failover event. Stateful synchronization is required for IPsec high availability support. For IPsec connections, AMS supports 1:1 configuration only.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

load-balancing-options (Aggregated Multiservices) | 9888

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

high-availability (security cloud)

IN THIS SECTION

- Syntax | 8070
- Hierarchy Level | 8071
- Description | 8071
- Options | 8071
- Required Privilege Level | 8072
- Release Information | 8072

Syntax

```
high-availability {
          aws {
               eip-based;
               peer-liveliness {
                   probe-ip {
                     destination-ip-address;
                     source-ip source-ip-address;
```

```
                }
            }
            azure {
                hvtimesync-disable;
                secondary-ip-based;
                peer-liveliness {
                    probe-ip {
                      destination-ip-address;
                      source-ip source-ip-address;
                }
            }
            gcp {
                  secondary-ip-based;
                  peer-liveliness {
                    probe-ip {
                      destination-ip-address;
                      source-ip source-ip-address;
                }
            }
        }
    }
```

Hierarchy Level

```
[edit security cloud]
```

Description

Configure vSRX Virtual Firewall instances to work in Multinode High Availability mode with Amazon Web Services (AWS) or Microsoft Azure Cloud or Google Cloud Platform (GCP).

Options

aws—Define deployment type as AWS.

- EIP-based—Deployment type based on EIP.
- peer-liveliness—Configure activeness probe options.

- probe-ip *destination-ip*—IP address for activeness probe configuration. Specify destination IP address (of upstream router) for probing.
- source-ip—Source IP address to initiate probes.

azure—Define deployment type as Azure cloud.

- hvtimesync-disable—Disable the timesync to the host clock.
- secondary-ip-based—Deployment type based on secondary IP address.
- peer-liveliness—Configure activeness probe options.
  - probe-ip *destination-ip*—IP address for activeness probe configuration. Specify destination IP address (of upstream router) for probing.
  - source-ip—Source IP address to initiate probes.

gcp—Define deployment type as GCP.

- secondary-ip-based—Deployment type based on secondary IP address.
- peer-liveliness—Configure activeness probe options.
  - probe-ip *destination-ip*—IP address for activeness probe configuration. Specify destination IP address (of upstream router) for probing.
  - source-ip—Source IP address to initiate probes.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

Multinode High Availability Support for vSRX Instances

high-availability (Security IDP)

IN THIS SECTION

- Syntax | 8073
- Hierarchy Level | 8073
- Description | 8073
- Options | 8073
- Required Privilege Level | 8073
- Release Information | 8074

Syntax

```
high-availability {
    no-policy-cold-synchronization;
}
```

Hierarchy Level

```
[edit security idp sensor-configuration]
```

Description

Configures high availability (chassis cluster) for IDP.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

high-cos-queue-threshold

IN THIS SECTION

- Syntax | 8074
- Hierarchy Level | 8074
- Description | 8074
- Options | 8075
- Required Privilege Level | 8075
- Release Information | 8075

Syntax

```
high-cos-queue-threshold number;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Configure the class-of-service (CoS) queue threshold by percentage.

NOTE: This feature is only available when you enable subscriber management. For more information on enabling subscriber management, see *Configuring Junos OS Enhanced Subscriber Management*.

This CoS resource monitoring feature bases admission decisions only on queues. Other CoS resources are not part of this criteria. This feature does not support throttling for subscribers arriving on pseudo-wire, logical tunnel, or redundant logical tunnel devices. The feature is supported on the following hardware:

- MX240, MX480, and MX960 routers
- MPC2E legacy, MPC2E-NG, MPC3E-NG, MPC5E, and MPC7E line cards

Options

`number` High-threshold percentage for CoS queue utilization per scheduler. When CoS queue utilization on a given FPC reaches that FPC's configured threshold level, further subscriber logins on that FPC are not allowed. This resource monitoring mechanism provides adjustable safety margins to proactively avoid completely exhausting each FPC's available CoS queue resources. Starting in Junos OS Release 19.4R1, you can specify a value of 0, which means that no subscribers are throttled based on CoS queue throttling.

- **Range:** 0 through 90
- **Default:** 100

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Resource Monitoring for Subscriber Management and Services Overview

high-polarization

IN THIS SECTION

- Syntax | 8076
- Hierarchy Level | 8076
- Description | 8076
- Default | 8076
- Required Privilege Level | 8076
- Release Information | 8077

Syntax

```
high-polarization;
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Description

Enable the physical port to rapidly track the state of polarization changes. Enabling this statement reduces the optical signal to noise ratio (OSNR) by few tenths of dB.

Default

By default, the `high-polarization` statement is disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1 for ACX6360 routers.

RELATED DOCUMENTATION

Supported Optics Options on ACX6360 and ACX5448-D Routers

high-power

IN THIS SECTION

- Syntax | 8077
- Hierarchy Level | 8077
- Description | 8077
- Default | 8078
- Required Privilege Level | 8078
- Release Information | 8078

Syntax

```
high-power;
```

Hierarchy Level

```
[edit poe interface (all | interface-name)]
```

Description

Configure a four-pair PoE port to deliver up to 60 W of power. The four-pair PoE standard allows for power to be delivered over all four pairs of wires in a standard CAT5/6 RJ-45 Ethernet cable.

When you configure the four-pair PoE port with the `high-power` option, the port provides up to 60 W of power, even if the powered device does not consume that amount of power.

Default

By default the PoE port operates in two-pair mode and provides power up to 30 W in accordance with the IEEE 802.3at PoE+ standard.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Configuring PoE Interfaces on EX Series Switches

high-threshold (Resource Monitor)

IN THIS SECTION

- Syntax | 8079
- Hierarchy Level | 8079
- Description | 8079
- Options | 8079
- Required Privilege Level | 8079
- Release Information | 8079

Syntax

```
high-threshold number;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Configure the high threshold value, exceeding which warnings or error logs are generated, for all the regions of memory, such as heap or ukernel, next-hop and encapsulation, and firewall filter memory. You can configure the resource-monitoring capability on MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.

Options

number High threshold percentage for memory resource utilization

- **Range:** 1 through 100

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

high-watermark

IN THIS SECTION

- Syntax | 8080
- Hierarchy Level | 8080
- Description | 8080
- Options | 8080
- Required Privilege Level | 8080
- Release Information | 8081

Syntax

```
high-watermark percent;
```

Hierarchy Level

```
[edit security flow aging]
```

Description

Sets the point at which the aggressive aging-out process begins.

Options

percent —Percentage of session-table capacity at which aggressive aging-out starts.

- **Range:** 1 through 100 percent
- **Default:** 100 percent

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

hint

IN THIS SECTION

- Syntax | 8081
- Hierarchy Level | 8081
- Description | 8082
- Default | 8082
- Options | 8082
- Required Privilege Level | 8082
- Release Information | 8082

Syntax

```
hint [ hint-strings ];
```

Hierarchy Level

```
[edit services nat pool nat-pool-name pgcp]
```

Description

Configure a hint that enables the border gateway function (BGF) to choose a NAT pool by direction rather than by virtual interface. The BGF matches the configured hint with a termination hint located in the Direction field of a nonstandard termination ID.

Default

When no hint is configured, the BGF can choose any NAT pool associated with the virtual interface.

Options

hint-string—Alphanumeric string of up to three characters that the BGF uses to match with a termination hint located in the Direction field of a nonstandard termination ID. You can also include underscores (_) and hyphens (-) within the string. To specify a list of hints, use the format: **[hint *xx* hint *yy*]**.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

history

IN THIS SECTION

- Syntax | 8083
- Hierarchy Level | 8083
- Description | 8083
- Default | 8083
- Options | 8083

Required Privilege Level | 8084
Release Information | 8084

Syntax

```
history history-index {
    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
}
```

Hierarchy Level

```
[edit snmp rmon]
```

Description

Configure RMON history group entries. This RMON feature can be used with the Simple Network Management Protocol (SNMP) agent in the switch to monitor all the traffic flowing among switches on all connected LAN segments. The RMON history feature collects statistics in accordance with user-configurable parameters.

The history group controls the periodic statistical sampling of data from various types of networks. This group contains configuration entries that specify an interface, polling period, and other parameters. If you use the `history` statement, you must also configure the "interface" on page 9056 *`interface-name`* statement.

Default

Not configured.

Options

history-index—Tag the history entries with a number.

- **Range:** 1 through 655535 (integer)

Required Privilege Level

snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Understanding RMON

RMON MIB Event, Alarm, Log, and History Control Tables

Configuring RMON History Sampling

Monitoring RMON MIB Tables

Junos OS Network Management Configuration Guide

history-size

IN THIS SECTION

- Syntax | 8085
- Junos OS Hierarchy Levels | 8085
- Junos OS Evolved Hierarchy Level | 8085
- Description | 8085
- Options | 8085
- Required Privilege Level | 8085
- Release Information | 8085

Syntax

```
history-size size;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the number of stored history entries.

Options

size—Value from 0 to 255.

- **Default:** 50

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement at the [`edit services rpm twamp client control-connection` *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

Configuring BGP Neighbor Discovery Through RPM

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

hold-down

IN THIS SECTION

- Syntax | 8086
- Hierarchy Level | 8086
- Description | 8087
- Options | 8087
- Required Privilege Level | 8087
- Release Information | 8088

Syntax

```
hold-down {
    seconds;
    flaps number;
    period seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options bmp],
[edit logical-systems logical-system-name routing-options bmp station station-name],
[edit routing-options bmp],
[edit routing-options bmp station station-name]
```

Description

If the connection to a BMP station flaps and the `hold-down` statement is configured, the station is prevented from reconnecting to the device for the specified period of time. A flap is when the TCP session unexpectedly switches from established to non-established. If you alter the configuration of the `hold-down` statement, the hold down timer and flap counter are reset.

You can effectively disable the `hold-down` statement by setting the `flaps` option to 10 and the `period` option to 30 seconds.

Options

seconds Specify the time in seconds to wait before allowing the BMP station to reconnect to the device.

- **Default:** 600 seconds
- **Range:** 30 through 65,535 seconds

`flaps` ***number*** Specify the number of BMP station flaps allowed before terminating the connection to the BMP station and triggering the hold down timer.

- **Default:** 3 flaps
- **Range:** 2 to 10 flaps

`period` ***seconds*** Specify the time in seconds for the BGP station flaps (specified using the `flaps` option) to occur before triggering the hold down timer. Every time a flap occurs, the number of flaps in the last time period is checked to see if the criteria is met.

For example, if you defined the `period` as 60 seconds and the `flaps` as 4 and the BGP station flaps just 2 times in a 60 second period, the hold down timer would not be triggered. However, if the BGP station flaps 4 times in a 60 second period, the hold down timer would be triggered.

- **Default:** 300 seconds
- **Range:** 30 through 65,535 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D15.

RELATED DOCUMENTATION

Configuring BGP Monitoring Protocol Version 3

hold-down-interval

IN THIS SECTION

- Syntax | 8088
- Hierarchy Level | 8088
- Description | 8088
- Options | 8089
- Required Privilege Level | 8089
- Release Information | 8089

Syntax

```
hold-down-interval number;
```

Hierarchy Level

```
[edit chassis cluster redundancy-group group-number]
```

Description

Set the minimum interval to be allowed between back-to-back failovers for the specified redundancy group (affects manual failovers, as well as automatic failovers associated with monitoring failures).

For redundancy group 0, this setting prevents back-to-back failovers from occurring less than 5 minutes (300 seconds) apart. Note that a redundancy group 0 failover implies a Routing Engine failure.

For some configurations, such as ones with a large number of routes or logical interfaces, the default or specified interval for redundancy group 0 might not be sufficient. In such cases, the system automatically extends the dampening time in increments of 60 seconds until the system is ready for failover.

Options

number—Number of seconds specified for the interval.

- **Range:** For redundancy group 0, 300 through 1800 seconds; for redundancy group 1 through 128, 0 through 1800 seconds.
- **Default:** For redundancy group 0, 300 seconds; for redundancy group 1 through 128, 1 second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

hold-interval

IN THIS SECTION

Syntax | 8090
Hierarchy Level | 8090

- Description | 8090
- Options | 8090
- Required Privilege Level | 8090
- Release Information | 8091

Syntax

```
hold-interval seconds;
```

Hierarchy Level

```
[edit security dynamic-address session-scan]
```

Description

If you enable the `session-scan` option at the `[edit security dynamic-address address-name name]` hierarchy level, the security device initiates session scan when:

- You add a new address into the dynamic-address entry.
- The system reaches the configured `hold-interval` time.

Options

seconds Time interval before the session scan requests in seconds.

- **Range:** 1 thru 3600 seconds
- **Default:** 10 seconds

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

session-scan | 16847

hold-interval (Clock Synchronization)

IN THIS SECTION

- Syntax | 8091
- Hierarchy Level | 8091
- Description | 8092
- Options | 8092
- Required Privilege Level | 8092
- Release Information | 8092

Syntax

```
hold-interval {
    configuration-change seconds;
    restart seconds;
    switchover seconds;
}
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Configure the wait time for clock selection after a change in configuration and after a reboot of the router, and configure the switchover wait time after clock recovery.

Options

`configuration-change` Set the wait time for clock selection after a change in configuration.

- **Range:** 15 seconds through 60 seconds

`restart` Set the wait time for clock selection after reboot of the router.

- **Range:** 60 seconds through 180 seconds
- **Default:** 120 seconds

`switchover` Set the switchover wait time after clock recovery.

- **Range:** 30 seconds through 60 seconds.
- **Default:** 30 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

synchronization (MX Series) | **28795**

hold-interval (OAM)

IN THIS SECTION

- Syntax | 8093
- Hierarchy Level | 8093
- Description | 8093
- Options | 8093
- Required Privilege Level | 8094
- Release Information | 8094

Syntax

```
hold-interval minutes;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name continuity-check]
```

Description

The time to wait in minutes before flushing the maintenance association end point (MEP) database, if no updates occur. The configurable range is 1 minute through 30240 minutes.The default value is 10 minutes.

NOTE: Hold timer based flushing is applicable only for auto discovered remote MEPs and not for statically configured remote MEPs.

Options

minutes—Time to wait, in minutes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Continuity Check Protocol Parameters Overview

Configuring Continuity Check Protocol Parameters for Fault Detection

hold-interval (OAM CFM, for EX Series Switch Only)

IN THIS SECTION

- Syntax | 8094
- Hierarchy Level | 8095
- Description | 8095
- Options | 8095
- Required Privilege Level | 8095
- Release Information | 8095

Syntax

```
hold-interval minutes;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name continuity-check]
```

Description

Configure the time to wait before flushing the maintenance association end point (MEP) database, if no updates occur.

Options

minutes—Time to wait, in minutes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)

Junos OS Network Interfaces Configuration Guide

hold-interval (Protection Group)

IN THIS SECTION

- Syntax | 8096
- Hierarchy Level | 8096

- Description | 8096
- Options | 8096
- Required Privilege Level | 8096
- Release Information | 8096

Syntax

```
hold-interval number;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring name]
```

Description

Specify the hold-off timer interval *for all rings* in 100 millisecond (ms) increments.

Options

number—Hold-timer interval, in milliseconds.

- **Range:** 0 through 10,000 ms
- **Default:** 100 ms

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS

hold off time

IN THIS SECTION

- Syntax | 8097
- Hierarchy Level | 8097
- Description | 8097
- Options | 8097
- Required Privilege Level | 8098
- Release Information | 8098

Syntax

hold-off-time *time*;

Hierarchy Level

[`set chassis synchronization source interfaces`]

Description

Enable hold-off time for Synchronous Ethernet interfaces and external clock source interfaces. If an interface goes down, hold-off time delays short signal failures from being sent to the clock selection process to prevent rapid successive switching.

Options

time Amount of time in milliseconds that a signal is held before being passed to the clock selection process.

Range: 300–1800 milliseconds

Default: 1000 milliseconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring Clock Synchronization Interface on MX Series Routers

Ethernet Synchronization Message Channel Overview

hold-time (ae interface)

IN THIS SECTION

- Syntax | 8099
- Hierarchy Level | 8099
- Hierarchy Level (ae- interfaces) (PTX5000, PTX1000, PTX10002, PTX10008, PTX10016, and MX Series routers) | 8099
- Description | 8099
- Default | 8100
- Options | 8100
- Required Privilege Level | 8100
- Release Information | 8101

Syntax

```
hold-time up milliseconds down milliseconds;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-range interface-range-name]
```

Hierarchy Level (ae- interfaces) (PTX5000, PTX1000, PTX10002, PTX10008, PTX10016, and MX Series routers)

```
[edit interfaces ae{x}]
```

Description

Specify the `hold-time` value to use to damp shorter interface transitions milliseconds. The hold timer enables interface damping by not advertising interface transitions until the hold timer duration has passed. When a hold-down timer is configured and the interface goes from up to down, the down hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still down, then the router begins to advertise the interface as being down. Similarly, when a hold-up timer is configured and an interface goes from down to up, the up hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still up, then the router begins to advertise the interface as being up.

The attribute “up” and “down” is mandatory when we configure the knob “hold-time”. Failing to configure either one, will result in a commit error.

We recommend you to configure sub-second down hold timer value as 500 milliseconds on:

- MX204 router
- MX routers with MPC7E, MPC8E, MPC9E, MPC10E, and MX10003-LC2103 MPCs

The sub-second down hold timer is not supported on the interfaces created on MX routers with MPC3E, MPC4E, MPC5E, MPC5EQ, MPC6E, and MPC11E MPCs. The sub-second hold timer is a configuration of hold timer less than a second.

Starting from Junos OS release 21.4R1, you can specify the `hold-time` value on aggregated ethernet interfaces. When you configure hold-timer for `ae-` interfaces, we recommend not to configure the hold-time for member links.

NOTE:

- We recommend that you configure the `hold-time` value after determining an appropriate value by performing repeated tests in the actual hardware environment. This is because the appropriate value for `hold-time` depends on the hardware (XFP, SFP, SR, ER, or LR) used in the networking environment.
- The `hold-time` option is not available for controller interfaces.

NOTE: On MX Series routers with MPC3E and MPC4E, we recommend that you do not configure the hold-down timer to be less than 1 second. On MX Series routers with MPC5EQ-100G10G (MPC5EQ) or MPC6E (MX2K-MPC6E) with 100-Gigabit Ethernet MIC with CFP2 OTN interfaces, we recommend that you do not configure the hold-down timer to be less than 3 seconds.

Default

Interface transitions are not damped.

Options

`down` *`milliseconds`*—Hold time to use when an interface transitions from up to down. Junos OS advertises the transition within 100 milliseconds of the time value you specify.

- **Range:** 0 through 4,294,967,295
- **Default:** 0 (interface transitions are not damped)

`up` *`milliseconds`*—Hold time to use when an interface transitions from down to up. Junos OS advertises the transition within 100 milliseconds of the time value you specify.

- **Range:** 0 through 4,294,967,295
- **Default:** 0 (interface transitions are not damped)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

- advertise-interval
- interfaces (EX Series switches) | **9135**
- Physical Interface Damping Overview
- *Damping Shorter Physical Interface Transitions*
- Damping Longer Physical Interface Transitions

hold-time

IN THIS SECTION

- Syntax | **8101**
- Hierarchy Level | **8102**
- Description | **8102**
- Options | **8102**
- Required Privilege Level | **8102**
- Release Information | **8102**

Syntax

```
hold-time seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id preempt]
```

Description

Configure the time in seconds after which a backup router with the highest priority preempts the primary router.

Options

seconds—Hold-time period.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring VRRP for IPv6 (CLI Procedure)

hold-time (OAM)

IN THIS SECTION

- Syntax | 8103
- Hierarchy Level | 8103
- Description | 8103

- Options | 8103
- Required Privilege Level | 8103
- Release Information | 8103

Syntax

```
hold-time seconds;
```

Hierarchy Level

```
[edit protocols oam],
[edit protocols oam gre-tunnel interface interface-name]
```

Description

Length of time the originating end of a GRE tunnel waits for keepalive packets from the other end of the tunnel before marking the tunnel as operationally down.

Options

seconds—Hold-time value.

- **Default:** 5 seconds
- **Range:** 5 through 250 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Tunnel Interfaces on an MX Series Router with a 16x10GE 3D MPC

Configuring GRE Keepalive Time

keepalive-time (Protocols OAM) | **9472**

hold-time (Protocols BGP)

IN THIS SECTION

- Syntax | **8104**
- Hierarchy Level | **8104**
- Description | **8105**
- Options | **8105**
- Required Privilege Level | **8106**
- Release Information | **8106**

Syntax

```
hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
```

```
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Specify the hold-time value to use when negotiating a connection with the peer. The hold-time value is advertised in open packets and indicates to the peer the length of time that it should consider the sender valid. If the peer does not receive a keepalive, update, or notification message within the specified hold time, the BGP connection to the peer is closed and routing devices through that peer become unavailable.

The hold time is three times the interval at which keepalive messages are sent.

BGP on the local routing device uses the smaller of either the local hold-time value or the peer's hold-time value received in the open message as the hold time for the BGP connection between the two peers.

Starting in Junos OS Release 12.3, the BGP hold-time value can be zero (0). This implies that the speaker does not expect keepalive messages from its peer to maintain the BGP session. When negotiating between two peers, if one side requests a nonzero hold time and the other requests a zero hold time, the negotiation settles on the nonzero value and keepalive intervals are determined accordingly. Both sides must be set to zero for keepalive messages to stop being sent.

Starting in Junos OS Evolved Release 24.2R1, you can configure up to 400 BGP sessions with highly sensitive hold-times of 10 seconds or less, and up to 8000 total BGP sessions on PTX Series devices that support this feature. When configuring BGP sessions at scale, we recommend using the default hold-time value of 90 seconds for the majority of BGP sessions.

Options

seconds—Hold time.

- **Range:** 3 through 65,535 seconds (or 0 for infinite hold time)
- **Range:** 10 through 65,535 seconds for EX Series switches.
- **Default:** 90 seconds

TIP: When you set a hold-time value of 3 through 19 seconds, we recommend that you also configure the BGP `precision-timers` statement. The `precision-timers` statement ensures

that if scheduler slip messages occur, the routing device continues to send keepalive messages. When the `precision-timers` statement is included, keepalive message generation is performed in a dedicated kernel thread, which helps to prevent BGP session flaps.

Starting in Junos OS Release 17.3R1, the `precision-timers` statement is supported on QFX Series switches.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

BGP Messages Overview

precision-timers | 13010

hold-time (Protocols DVMRP)

IN THIS SECTION

- Syntax | 8107
- Hierarchy Level | 8107
- Description | 8107
- Options | 8107
- Required Privilege Level | 8107
- Release Information | 8107

Syntax

```
hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols dvmrp interface interface-name],
[edit protocols dvmrp interface interface-name]
```

Description

Specify the time period for which a neighbor is to consider the sending router (this router) to be operative (up).

Options

seconds—Hold time.

- **Range:** 1 through 255
- **Default:** 35 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

> **NOTE**: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring DVMRP

hold-time (Protocols IS-IS)

IN THIS SECTION

- Syntax | 8108
- Hierarchy Level | 8108
- Description | 8109
- Options | 8109
- Required Privilege Level | 8110
- Release Information | 8110

Syntax

```
hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis interface interface-name level level-number],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis interface interface-name level level-number],
[edit protocols isis interface interface-name level level-number],
[edit routing-instances routing-instance-name protocols isis interface interface-name level level-number]
```

Description

Set the length of time a neighbor considers this router to be operative (up) after receiving a hello packet. If the neighbor does not receive another hello packet within the specified time, it marks this routing device as inoperative (down). The hold time itself is advertised in the hello packets.

The hold time specifies how long a neighbor should consider this routing device to be operative without receiving another hello packet. If the neighbor does not receive a hello packet from this routing device within the hold time, it marks the routing device as being unavailable.

For systems configured with graceful routing switchover (GRES) with Graceful Restart, the hold time for Primary and Backup Routing Engines should be set to a value higher than 40 seconds. This ensures that adjacencies between the Routing Engine and the neighboring peer 'helper' routers do not time out, stopping graceful restart, and all traffic.

Options

seconds—Hold-time value, in seconds.

- **Range:** 3 through 65,535 seconds, or 1 to send out hello packets every 333 milliseconds
- **Default:** 9 seconds (for designated intermediate system [DIS] routers), 27 seconds (for non-DIS routers; three times the default hello interval)

NOTE: When elected as a designated intermediate system [DIS] router on any LAN adjacency, the hello and hold-timer intervals are scaled down by a factor of 3. This means the default values of 9 and 27 seconds of ISIS hello and hold time intervals are scaled down to 3 and 9 seconds for LAN adjacencies. During switchovers, this hold time is too short to form LAN adjacencies. Therefore, you can configure one of the following solutions:

- Set the hello and hold time interval for LAN adjacencies to 30 seconds and 90 seconds respectively on both the DIS router and a neighboring router.
- Convert the LAN interfaces to point-to-point IS-IS interfaces.

BEST PRACTICE: On QFX10000 switches, we strongly recommend that you configure all IS-IS interfaces, including peer interfaces, as point-to-point interfaces. If you do not, you might experience session flaps, that is, IS-IS sessions that go down and then come back up, when IS-IS is configured in virtual routing and forwarding (VRF) instances. When you scale IS-IS in any scenario, you might also experience scaling issues if you do not configure IS-IS interfaces as point-to-point interfaces.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Configuring Graceful Routing Engine Switchover*
- *Example: Configuring IS-IS*
- *Example: Configuring IS-IS for GRES with Graceful Restart*
- *Understanding IS-IS Configuration*
- hello-interval

hold-time (Protocols LDP)

IN THIS SECTION

- Syntax | 8110
- Hierarchy Level | 8111
- Description | 8111
- Options | 8111
- Required Privilege Level | 8111
- Release Information | 8111

Syntax

```
hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp interface interface-name],
[edit logical-systems logical-system-name protocols ldp targeted-hello],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp
interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp
targeted-hello],
[edit protocols ldp interface interface-name],
[edit protocols ldp targeted-hello],
[edit routing-instances routing-instance-name protocols ldp interface interface-name],
[edit routing-instances routing-instance-name protocols ldp targeted-hello]
```

Description

Specify how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. You can specify times for both link hello messages and targeted hello messages depending on the hierarchy level at which you configure the `hold-time` statement.

Options

seconds—Hold-time value.

- **Range:** 1 through 65,535 seconds
- **Default:** 15 seconds for link hello messages, 45 seconds for targeted hello messages

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for LDP targeted hellos added in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring the Delay Before LDP Neighbors Are Considered Down

hold-time (Protocols MSDP)

IN THIS SECTION

- Syntax | 8112
- Hierarchy Level | 8112
- Description | 8113
- Default | 8113
- Options | 8113
- Required Privilege Level | 8113
- Release Information | 8113

Syntax

```
hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name routing-instances instance-name protocols msdp],
[edit logical-systems logical-system-name routing-instances instance-name protocols msdp group
group-name peer address],
[edit logical-systems logical-system-name routing-instances instance-name protocols msdp peer
address],
[edit protocols msdp],
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit routing-instances instance-name protocols msdp],
```

```
[edit routing-instances instance-name protocols msdp group group-name peer address]
[edit routing-instances instance-name protocols msdp peer address]
```

Description

Specify the hold-time period to use when maintaining a connection with the MSDP peer. If a keepalive message is not received for the hold-time period, the MSDP peer connection is terminated. According to the RFC 3618, *Multicast Source Discovery Protocol (MSDP)*, the recommended value for the hold-time period is 75 seconds.

The hold-time period must be longer than the keepalive interval.

You might want to change the hold-time period and keepalive timer for consistency in a multi-vendor environment.

Default

In Junos OS, the default hold-time period is 75 seconds, and the default keepalive interval is 60 seconds.

Options

seconds—Hold time.

- **Range:** 15 through 150 seconds
- **Default:** 75 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Examples: Configuring MSDP

keep-alive (Protocols MSDP) | **9457**

sa-hold-time (Protocols MSDP) | 16262

hold-time (Protocols PIM)

IN THIS SECTION

- Syntax | 8114
- Hierarchy Level | 8114
- Description | 8114
- Options | 8115
- Required Privilege Level | 8115
- Release Information | 8115

Syntax

```
hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp bidirectional address address],
[edit logical-systems logical-system-name routing-instances instance-name protocols pim rp
bidirectional address address],
[edit protocols pim rp bidirectional address address],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances instance-name protocols pim rp bidirectional address address],
[edit routing-instances routing-instance-name protocols pim rp local family (inet | inet6)]
```

Description

Specify the time period for which a neighbor is to consider the sending routing device (this routing device) to be operative (up).

Options

seconds—Hold time.

- **Range:** 1 through 65535
- **Default:** 150 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for bidirectional RP addresses introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Local PIM RPs

Example: Configuring Bidirectional PIM

hold-time (Redundant Ethernet Interfaces)

IN THIS SECTION

- Syntax | 8116
- Hierarchy Level | 8116
- Description | 8116
- Options | 8116
- Required Privilege Level | 8116
- Release Information | 8117

Syntax

```
hold-time (up | down) timer
```

Hierarchy Level

```
[edit interfaces interface-name ]
```

Description

The hold timer enables interface damping by not advertising interface transitions until the hold timer duration has passed. When a hold-down timer is configured for a parent RETH interface and the primary child interface goes from up to down, the down hold-time timer is triggered. Every interface transition that occurs during the down hold-time is ignored. When the timer expires and the primary child interface state is still down, then the router begins to advertise the parent RETH interface as being down. Similarly, when a hold-up timer is configured for a parent RETH interface and the primary child interface goes from down to up, the up hold-time timer is triggered. Every interface transition that occurs during the up hold-time is ignored. When the timer expires and the primary child interface state is still up, then the router begins to advertise the parent RETH interface as being up.

The hold timer (both up and down) improves the flexibility and resilience of SRX Series Firewalls. Specify the *timer* value in seconds to reduce unnecessary loss of traffic and downtime.

> **NOTE**: Starting in Junos OS release 18.4R1, all SRX Series Firewalls have default delay timer of 11 seconds for both up hold-time and down hold-time.

Options

`down` *seconds*—Hold time to use when an interface transitions from up to down.

`up` *seconds*—Hold time to use when an interface transitions from down to up.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Physical Interface Damping Overview

hold-time (ae interface) | 8098

hold-time up

IN THIS SECTION

- Syntax | 8117
- Hierarchy Level | 8117
- Description | 8118
- Options | 8118
- Required Privilege Level | 8118
- Release Information | 8118

Syntax

```
hold-time up timer-value;
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options lacp]
```

Description

Specifies the time period for which the Link Aggregation Control Protocol (LACP) maintains the state of a child (member) link as expired or default.

When a child link goes from the current state to the expired state, the LACP monitors the reception of protocol data units (PDUs) on the child link for the configured hold-up time interval and does not allow the child link to transition back to the current state. This configuration thus prevents excessive flapping of a child link on an aggregated Ethernet interface.

The configured hold-up timer value is applicable to all the child links within a link aggregated (LAG) interface. By default, this feature is disabled.

Options

timer-value—Hold-up interval in seconds.

- **Range:** 1 to 6000 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring LACP for Aggregated Ethernet Interfaces

Configuring Aggregated Ethernet LACP (CLI Procedure)

hold-time (VRRP)

IN THIS SECTION

- Syntax | 8119
- Hierarchy Level | 8119
- Description | 8119
- Default | 8120
- Options | 8120
- Required Privilege Level | 8120
- Release Information | 8120

Syntax

```
hold-time seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id preempt],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id preempt],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id preempt],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id preempt]
```

Description

In a Virtual Router Redundancy Protocol (VRRP) configuration, set the hold time before a higher-priority backup router preempts the primary router.

Default

VRRP preemption is not timed.

Options

seconds—Hold-time period.

- **Range:** 0 through 3600 seconds
- **Default:** 0 seconds (VRRP preemption is not timed.)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Backup Router to Preempt the VRRP Primary Router

Configuring VRRP Preemption and Hold Time

holddown (candidate-uplink-port-profile)

IN THIS SECTION

- Syntax | 8121
- Hierarchy Level | 8121
- Description | 8121
- Default | 8121
- Options | 8121

Required Privilege Level | 8121
Release Information | 8122

Syntax

```
holddown interval;
```

Hierarchy Level

```
[edit policy-options satellite-policies candidate-uplink-port-profileprofile-name]
[edit policy-options satellite-policies candidate-uplink-port-profile profile-name term term-
name from]
```

Description

Configures the interval of time uplink failure detection waits before trying to try re-enable a satellite device's extended ports after shutting them down due to an uplink port failure. It is intended to avoid port flapping on the extended ports when uplink port connectivity is unstable.

Default

The default holddown interval is 6 seconds.

Options

interval—The holddown interval, in seconds. Valid values are 1-600 seconds.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

holddown-interval

IN THIS SECTION

- Syntax | 8122
- Hierarchy Level | 8122
- Description | 8123
- Options | 8123
- Required Privilege Level | 8123
- Release Information | 8123

Syntax

```
holddown-interval milliseconds;
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options
dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group group-name liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd],
```

```
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd]
```

Description

Configure the time (in milliseconds) for which Bidirectional Forwarding Detection (BFD) holds a session up notification.

Options

milliseconds—Interval specifying how long a BFD session must remain up before a state change notification is sent.

- **Range:** 0 through 255,000
- **Default:** 0

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients

Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

holddown-interval (protocols ldp)

IN THIS SECTION

- Syntax | 8124
- Hierarchy Level | 8124
- Description | 8124
- Options | 8124
- Required Privilege Level | 8125
- Release Information | 8125

Syntax

```
holddown-interval holddown-interval;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp oam bfd-livenesss-detection],
[edit logical-systems logical-system-name protocols ldp oam fec address bfd-livenesss-detection],
[edit protocols ldp oam bfd-livenesss-detection],
[edit protocols ldp oam fec address bfd-livenesss-detection]
```

Description

Specify how long the BFD session should be up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.

Options

holddown-interval—Number of seconds the BFD session should remain up before adding the route or next hop.

- **Default:** 0 seconds

- **Range:** 0 through 65,535 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the Holddown Interval for the BFD Session

holddown (Protocols RIP)

IN THIS SECTION

- Syntax | 8125
- Hierarchy Level | 8126
- Description | 8126
- Options | 8126
- Required Privilege Level | 8126
- Release Information | 8126

Syntax

```
holddown seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit protocols rip],
[edit routing-instances routing-instance-name protocols rip]
```

Description

Configure how long the expired route is retained in the routing table before being removed.

When the hold-down timer runs on RIP demand circuits, routes are advertised as unreachable on other interfaces. When the hold-down timer expires, the route is removed from the routing table if all destinations detect that the route is unreachable or the remaining destinations are down.

Options

seconds—Estimated time to wait before making updates to the routing table.

- **Range:** 10 through 180 seconds
- **Default:** 180 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring RIP Timers

RIP Demand Circuits Overview

holddown (Protocols RIPng)

IN THIS SECTION

- Syntax | 8127
- Hierarchy Level | 8127
- Description | 8127
- Options | 8127
- Required Privilege Level | 8128
- Release Information | 8128

Syntax

```
holddown seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng],
[edit protocols ripng],
[edit routing-instances routing-instance-name protocols ripng]
```

Description

Configure how long the expired route is retained in the routing table before being removed.

Options

seconds—Estimated time to wait before removing expired routes from the routing table.

- **Default:** 180 seconds
- **Range:** 10 through 180 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring RIPng Update Interval

holdover-mode-disable

IN THIS SECTION

- Syntax | 8128
- Hierarchy Level | 8129
- Description | 8129
- Required Privilege Level | 8129
- Release Information | 8129

Syntax

```
holdover-mode-disable;
```

Hierarchy Level

```
[edit chassis synchronization output interfaces external]
[edit chassis synchronization output interfaces (external-0/0 | external-1/0)]
```

Description

Disable the holdover mode on the external output interface—external—on SCBE or on the external output interfaces—external-0/0 and external-1/0—on SCBE2.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

holdover-time-error-budget

IN THIS SECTION

- Syntax | 8130
- Hierarchy Level | 8130
- Description | 8130
- Default | 8130
- Required Privilege Level | 8130
- Release Information | 8130

Syntax

```
holdover-time-error-budget holdover-time-error-budget;
```

Hierarchy Level

```
[edit logical-systems name   protocols ptp]
[edit protocols ptp]
```

Description

Configure the PTP clock hold-over-in-specification time error budget. The range supported is 0 to 10000 nano seconds.

Default

400 nanoseconds

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

> *IEEE 1588v2 Precision Timing Protocol (PTP)*
>
> *Precision Time Protocol Overview*

home-address

IN THIS SECTION

- Syntax | 8131
- Hierarchy Level | 8131
- Description | 8131
- Options | 8131
- Required Privilege Level | 8132
- Release Information | 8132

Syntax

```
home-address {
    match (equal | greater-than | less-than | not-equal);
    value value;
}
```

Hierarchy Level

```
[edit  set security idp custom-attack attack-name attack-type signature protocol ipv6 extension-
header destination-option]
```

Description

Specify the IPv6 home address of the mobile node.

Options

match (equal | greater-than | less-than | not-equal)	Match an operand.
value	Match a decimal value.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

hop-by-hop-header

IN THIS SECTION

- Syntax | 8132
- Hierarchy Level | 8133
- Description | 8133
- Options | 8133
- Required Privilege Level | 8133
- Release Information | 8133

Syntax

```
hop-by-hop-header {
    CALIPSO-option;
    RPL-option;
    SFM-DPD-option;
    jumbo-payload-option;
    quick-start-option;
    router-alert-option;
    user-defined-option-type low | <to high>;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name ip ipv6-extension-header]
```

Description

Define the IPv6 hop-by-hop screen option.

Options

CALIPSO-option	Enable the Common Architecture Label IPv6 Security Option.
RPL-option	Enable the Routing Protocol for Low-Power and Lossy Networks screen option.
SMF-DPD-option	Enable the Simplified Muliticast Forwarding IPv6 Duplicate Packet Detection screen option.
jumbo-payload-option	Enable the IPv6 jumbo payload screen option.
quick-start-option	Enable the IPv6 quick start screen option.
router-alert-option	Enable the IPv6 router alert screen option.
user-defined-header-type ***low*** **| <to** ***high*****>**	Define the type of header range. • **Range:** 1 through 255.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Understanding IPv6 Support for Screens

ipv6-extension-header | 9350

destination-header | 5419

hop-limit

IN THIS SECTION

- Syntax | 8134
- Hierarchy Level | 8134
- Description | 8135
- Options | 8135
- Required Privilege Level | 8135
- Release Information | 8135

Syntax

```
hop-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-
reroute],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection bypass bypass-name],
```

```
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name fast-reroute],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Description

Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:

- LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths.
- Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers.
- Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.

Options

number—Maximum number of hops.

- **Range:** 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)
- **Default:** 255 (for an LSP or for a link protection bypass); 6 (for fast reroute)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Fast Reroute

Limiting the Number of Hops in LSPs

Configuring the Hop Limit for Bypass LSPs

host (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 8136
- Hierarchy Level | 8136
- Description | 8136
- Options | 8137
- Required Privilege Level | 8137
- Release Information | 8137

Syntax

```
host hostname {
    hardware-address mac-address;
    ip-address ip-address;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family (inet | inet6)]
```

Description

Configure a static binding for the specified client.

Options

hostname	Name of the client.
hardware-address *mac-address*	Specify the MAC address of the client. This is the hardware address that identifies the client on the network. • *mac-address*—MAC address of the client.
ip-address *ip-address*	Specify the reserved IP address assigned to the client. • *ip-address*—IP version 4 (IPv4) address.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management

host-address-base

IN THIS SECTION

- Syntax | 8138
- Hierarchy Level | 8138
- Description | 8138
- Options | 8138
- Required Privilege Level | 8138

Release Information | 8138

Syntax

```
host-address-base ip-address;
```

Hierarchy Level

```
[edit security nat source pool pool-name]
```

Description

Specify the base address of the original source IP address range. This is used for IP shifting.

Options

ip-address —IP address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

host-address-base (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 8139
- Hierarchy Level | 8139
- Description | 8139
- Options | 8139
- Required Privilege Level | 8140
- Release Information | 8140

Syntax

```
host-address-base ip-address;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name]
```

Description

Configure static mapping of the source address.

For static NAT that is performed on the services card, configure a one-to-one static shifting of a range of original source addresses to the range of addresses in the source pool by specifying the base address of the original source address range.

For example, if the host address base is 198.51.100.30 and the NAT pool uses the range 203.0.113.10 to 203.0.113.20, then 198.51.100.30 translates to 203.0.113.10, 198.51.100.31 translates to 203.0.113.11, and so on.

Options

host-address-base *ip-address*	The IP address used as the host address base.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

host (Diameter Origin)

IN THIS SECTION

- Syntax | 8140
- Hierarchy Level | 8140
- Description | 8140
- Options | 8141
- Required Privilege Level | 8141
- Release Information | 8141

Syntax

```
host hostname;
```

Hierarchy Level

```
[edit access diameter origin]
```

Description

Specify the name of the host that originates the Diameter message.

Options

hostname Name of the message origin host. Supplied as the value of the Origin-Host AVP for all messages sent by the Diameter instance.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

host-fast-reroute

IN THIS SECTION

- Syntax | 8141
- Hierarchy Level | 8142
- Description | 8142
- Required Privilege Level | 8142
- Release Information | 8142

Syntax

```
host-fast-reroute {
    global-arp-prefix-limit number;
```

```
    global-supplementary-blackout-timer minutes;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit routing-options]
```

Description

Configure global settings for all host fast reroute (HFRR) profiles configured on a routing device.

The remaining statements are described separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring Link Protection with Host Fast Reroute

host-inbound-traffic

IN THIS SECTION

- Syntax | 8143
- Hierarchy Level | 8143

- Description | 8143
- Options | 8143
- Required Privilege Level | 8143
- Release Information | 8144

Syntax

```
host-inbound-traffic {
    protocols protocol-name {
                    except;
    }
    system-services service-name {
                    except;
    }
}
```

Hierarchy Level

```
[edit security zones functional-zone management],
[edit security zones functional-zone management interfaces interface-name],
[edit security zones security-zone zone-name],
[edit security zones security-zone zone-name interfaces  interface-name]
```

Description

Control the type of traffic that can reach the device from interfaces bound to the zone.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding How to Control Inbound Traffic Based on Traffic Types

Understanding How to Control Inbound Traffic Based on Protocols

host (Interfaces)

IN THIS SECTION

- Syntax | 8144
- Hierarchy Level | 8145
- Description | 8145
- Options | 8145
- Required Privilege Level | 8145
- Release Information | 8145

Syntax

```
host hostname {
    services severity-level;
    facility-override facility-name;
    log-prefix prefix-value;
    port port-number;
}
```

Hierarchy Level

```
[edit interfaces interface-name services-options syslog]
```

Description

Specify the hostname for the system logging utility.

Starting with Junos OS release 17.4R1, you can configure up to a maximum of four system log servers (combination of local system log hosts and remote system log collectors) for each service set for ms interface under `[edit interfaces interface-name services-options]` hierarchy.

Options

hostname—Name of the system logging utility host machine. This can be the local Routing Engine or an external server address.

From Junos OS Release 17.4R1, you can configure up to four system log hosts.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

You can configure multiple system log hosts from Junos OS Release 17.4R1 onwards.

RELATED DOCUMENTATION

Applying Filters and Services to Interfaces

host (L2TP)

IN THIS SECTION

- Syntax | 8146
- Hierarchy Level | 8146
- Description | 8146
- Options | 8146
- Required Privilege Level | 8147
- Release Information | 8147

Syntax

```
host hostname {
    services severity-level;
    facility-override facility-name;
    log-prefix prefix-value;
}
```

Hierarchy Level

```
[edit services l2tp tunnel-group group-name syslog]
```

Description

Specify the hostname for the system logging utility.

Options

hostname—Name of the system logging utility host machine. This can be the local Routing Engine or an external server address.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring System Logging of L2TP Tunnel Activity

host-name

IN THIS SECTION

- Syntax | 8147
- Hierarchy Level (EX Series) | 8147
- Hierarchy Level (MX Series) | 8148
- Description | 8148
- Required Privilege Level | 8148
- Release Information | 8148

Syntax

```
host-name host-name;
```

Hierarchy Level (EX Series)

```
[edit vlans vlan-name forwarding-options dhcp-security option-82 remote-id]
```

Hierarchy Level (MX Series)

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security remote-id option-82]
```

Description

Use the hostname of the switching device as the `remote-id` suboption of DHCP option 82 (also known as the DHCP relay agent information option) in DHCP request packet headers before forwarding or relaying requests to a DHCP server. This suboption provides a trusted identifier for the host system that has forwarded or relayed requests to the server.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Setting Up DHCP Option 82 on the Switch with No Relay (ELS)

Configuring DHCP Option 82 on a Router with Bridge Domain

http://tools.ietf.org/html/rfc3046

host-name (System)

IN THIS SECTION

- Syntax | 8149
- Hierarchy Level | 8149
- Description | 8149

- Options | 8149
- Required Privilege Level | 8149
- Release Information | 8149

Syntax

```
host-name hostname;
```

Hierarchy Level

```
[edit system]
```

Description

Set the hostname of the router or switch.

On Junos OS Evolved, when you specify the `host-name` statement, the Routing Engine name where the configuration is committed is appended to the hostname you specify. For example, on Routing Engine 0, `set system host-name my-host` sets the hostname to `my-host-re0`. You can also use the `%s` character to designate where to substitute the Routing Engine name. For example, on Routing Engine 1, `set system host-name %s_my_host` sets the hostname to `re1_my-host`.

Options

hostname	Name of the router or switch.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Hostname of a Router or Switch by Using a Configuration Group

host-name (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 8150
- Hierarchy Level | 8150
- Description | 8150
- Required Privilege Level | 8151
- Release Information | 8151

Syntax

```
host-name name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82],
[edit forwarding-options dhcp-relay group group-name relay-option-82]
```

Description

Supports the addition of vendor-specific hostname in the option 82, suboption 9 field of DHCPv4 control messages on server-facing interfaces. The hostname can be a string of characters such as **Juniper-AB-1**.

Junos OS automatically adds the remaining vendor-specific information as per RFC 4243, *Vendor-Specific Information Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option*. The Juniper Networks enterprise ID is 2636. The hostname is option-data 1 (the location is option-data 2). The DHCPv4 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query the hostname to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 16.2.

host-name (DHCPv6 Relay Agent)

IN THIS SECTION

- Syntax | 8151
- Hierarchy Level | 8151
- Description | 8152
- Required Privilege Level | 8152
- Release Information | 8152

Syntax

```
host-name name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 relay-option-vendor-specific],
[edit forwarding-options dhcp-relaygroup group-name dhcpv6 relay-option-vendor-specific]
```

Description

Supports the addition of vendor-specific hostname in the vendor-specific option (17) of DHCPv6 control messages on server-facing interfaces. The hostname can be a string of characters such as **Juniper-AB-1**.

Junos automatically adds the remaining vendor-specific information as per RFC 4243, *Vendor-Specific Information Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option*. The Juniper Networks enterprise ID is 2636. The hostname is option-data 1 (the location is option-data 2). The DHCPv6 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query the hostname to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 16.2.

host (Next Gen Services Service-Set Remote System Logging)

IN THIS SECTION

- Syntax | **8153**
- Hierarchy Level | **8153**
- Description | **8153**
- Options | **8153**
- Required Privilege Level | **8153**
- Release Information | **8153**

Syntax

```
host host-ip-address;
```

Hierarchy Level

```
edit services service-set name syslog stream stream-name
```

Description

Specify the IP address of syslog server to receive log messages.

Options

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

host-only-interface

IN THIS SECTION

- Syntax | 8154
- Hierarchy Level | 8154
- Description | 8154
- Default | 8154
- Required Privilege Level | 8155
- Release Information | 8155

Syntax

```
host-only-interface;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols vlan vlan-id igmp-snooping interface interface-name]
```

Description

Configure an interface as a host-facing interface. IGMP queries received on these interfaces are dropped.

Default

The interface can either be a host-side or multicast-router interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping

host-os forwarding-options analyzer

IN THIS SECTION

- Syntax | 8155
- Hierarchy Level | 8156
- Description | 8156
- Options | 8156
- Required Privilege Level | 8156
- Release Information | 8157

Syntax

```
forwarding-options {
    analyzer analyzer-instance-name {
        input {
            ingress {
                virtual-network-function vnf-name {
                    interface interface-name;
                }
```

```
            }
            egress {
                virtual-network-function vnf-name {
                    interface interface-name;
                }
            }
        }
        output {
            virtual-network-function analyzer-vnf-name;
                interface interface-name;
        }
    }
```

Hierarchy Level

```
[edit]
```

Description

Configures an analyzer for port mirroring and configures port mirroring for either ingress or egress traffic of a VNF interface to an analyzer VNF.

Options

virtual-network-functions ***vnf-name***	Name of the VNF instance.
analyzer-instance-name	Name of the analyzer.
ingress	Traffic that enters the port.
egress	Traffic that leaves the port.

Required Privilege Level

set host-os—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D40.

host-outbound media-interface

IN THIS SECTION

- Syntax | 8157
- Hierarchy Level | 8157
- Description | 8157
- Required Privilege Level | 8157
- Release Information | 8158

Syntax

```
host-outbound media-interface;
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable Layer 2 port mirroring of host-generated outbound packets only on MPCs on MX Series 5G Universal Routing Platforms.

This statement enables all Routing Engine-generated Layer 2 injections to execute egress logical interface filters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Examples: Layer 2 Port Mirroring at Multiple Levels of the Chassis

Configuring Port Mirroring

Understanding Layer 2 Port Mirroring

host-outbound-traffic (Class-of-Service)

IN THIS SECTION

- Syntax | 8158
- Hierarchy Level | 8159
- Description | 8159
- Options | 8159
- Required Privilege Level | 8159
- Release Information | 8159

Syntax

```
host-outbound-traffic {
    forwarding-class class-name;
    dscp-code-point code-point;
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Allow queue selection for traffic generated by the Routing Engine (host). The selected queue must be configured properly. You can also configure specific DSCP code point bits for the type of service (ToS) field of the generated packets. This configuration does not affect transit packets or incoming packets. This is a global configuration that only affects packets originating on the Routing Engine. If you do not configure an output queue for host outbound traffic, the switch uses the default queue mapping.

Options

`forwarding-class` *`class-name`*—Set the forwarding class name for outbound host traffic (traffic generated by the Routing Engine).

`dscp-code-point` *`code-point`*—Set the six-bit DSCP code point value in the type of service (ToS) field of the packet generated by the Routing Engine (host).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Changing the Host Outbound Traffic Default Queue Mapping

Understanding Host Routing Engine Outbound Traffic Queues and Defaults

host-outbound-traffic (Class-of-Service)

IN THIS SECTION

- Syntax | 8160
- Hierarchy Level | 8160
- Description | 8161
- Default | 8161
- Options | 8161
- Required Privilege Level | 8161
- Release Information | 8161

Syntax

```
host-outbound-traffic {
    dscp-code-point value;
    forwarding-class class-name;
    ieee-802.1 {
        default value;
        rewrite-rules;
    override-firewall;
    protocol {
        isis-over-gre {
            dscp-code-point dscp-code-point;
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Classify and mark host outbound traffic. This statement does not affect transit traffic or incoming traffic.

Default

If you do not specify a forwarding class or DSCP value, the router uses the default queue and DSCP bit assignments for host outbound traffic.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for `ieee-802.1` statement introduced in Junos OS Release 12.3.

Support for distributed protocol handler traffic introduced in Junos OS Release 13.2.

Support for `protocol` statement introduced in Junos OS Release 17.3.

Support for `override-firewall` statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Default Routing Engine Protocol Queue Assignments

Default DSCP and DSCP IPv6 Classifiers

Changing the Default Queuing and Marking of Host Outbound Traffic

Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic

Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface

host-outbound-traffic (Multicast Snooping)

IN THIS SECTION

- Syntax | 8162
- Hierarchy Level | 8162
- Description | 8162
- Options | 8162
- Required Privilege Level | 8163
- Release Information | 8163

Syntax

```
host-outbound-traffic {
    forwarding-class class-name;
    dot1p number;
}
```

Hierarchy Level

```
[edit multicast-snooping-options],
[edit bridge-domains bridge-domain-name multicast-snooping-options],
[edit routing-instances routing-instance-name multicast-snooping-options],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]
```

Description

On an MX Series router in a network enabled for CET service and IGMP snooping, configure multicast forwarding class and IEEE 802.1p value to rewrite of IGMP self generated packets.

Options

- *class-name*—Name of the forwarding class.

- *number*—802.1p priority number.
- **Range:** 0 through 7
- **Default:** 0

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring Multicast Snooping

Configuring IGMP Snooping

host-prefix-only

IN THIS SECTION

- Syntax | 8163
- Hierarchy Level | 8164
- Description | 8164
- Required Privilege Level | 8164
- Release Information | 8164

Syntax

```
host-prefix-only;
```

Hierarchy Level

```
[edit dynamic-profiles interfaces interface-name unit logical-unit-number],
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name routing-instances routing-instance-name interfaces
interface-name unit logical-unit-number]
```

Description

(MPC5 and MPC6 cards) Improve datapath performance by allowing only DHCPv4 subscribers that negotiate a 32-bit prefix to come up on the underlying VLAN interface. All DHCP subscribers on the underlying interface must negotiate a 32-bit prefix. Subscribers that negotiate a subnet prefix are not brought up. You can configure this statement for static or dynamic subscribers.

NOTE: You must add or remove this statement before subscribers become active. The configuration fails if you attempt to configure the statement while subscribers are active.

NOTE: You must also configure `demux-source inet` for the logical interface. Only `inet` is supported. A commit error occurs if you specify `demux-source inet6` or `demux-source [inet inet6]`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

- Configuring an IP Demultiplexing Interface
- Configuring a VLAN Demultiplexing Interface

host (Security IDP Sensor Configuration)

IN THIS SECTION

- Syntax | 8165
- Hierarchy Level | 8165
- Description | 8165
- Options | 8165
- Required Privilege Level | 8165
- Release Information | 8166

Syntax

```
host ip-address <port number>;
```

Hierarchy Level

```
[edit security idp sensor-configuration packet-log]
```

Description

Configure the IP address and port number of the server where the packet capture object will be sent.

Options

- `host` *`ip-address`*—The IP address of the server where the packet capture object will be sent.
- `port` *`number`*—The port number of the server where the packet capture object will be sent.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

host (Security Log)

IN THIS SECTION

- Syntax | 8166
- Hierarchy Level | 8166
- Description | 8166
- Options | 8167
- Required Privilege Level | 8167
- Release Information | 8167

Syntax

```
host {
    ip-address;
    port port-number;
}
```

Hierarchy Level

```
[edit security log stream stream-name]
[edit logical-systems name security log stream stream-name]
[edit tenants tenant-name security log stream stream-name]
```

Description

You can specify the IP address of the server to which the security logs are streamed.

The [edit `logical-systems` *name* security log stream *stream-name*] hierarchy level introduced in Junos OS Release 18.2R1.

The [edit `tenants` *tenant-name* security log stream *stream-name*] hierarchy levels introduced in Junos OS Release 18.3R1.

Options

- *ip-address*—Specify IP address of the host.
- `port` *port-number*—Specify host port number.
- **Default:** The default destination port is the system log port. For UDP or TCP, the default port is 514. For TLS, the default port is 6514.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

syslog (System) | 28836

host (Security Web Filtering)

IN THIS SECTION

Syntax | 8168
Hierarchy Level | 8168
Description | 8168
Required Privilege Level | 8168

Release Information | 8168

Syntax

```
host host-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated server]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name server]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name server]
```

Description

Set server host parameters by entering the server name or IP address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Surf-Control feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

host (service-set)

IN THIS SECTION

- Syntax | 8169
- Hierarchy Level | 8170
- Description | 8170
- Options | 8170
- Required Privilege Level | 8170
- Release Information | 8170

Syntax

```
host hostname {
    class {
        alg-logs;
        deterministic-nat-configuration-log;
        ids-logs;
        nat-logs;
        packet-logs;
        pcp-logurlf-logs;
        session-logs <open | close>;
        stateful-firewall-logs;
        urlf-logs;
    }
    facility-override facility-name;
    interface-service prefix-value;
    log-prefix prefix-value;
    port port-number;
    services severity-level;
    source-address source-address;
}
```

Hierarchy Level

```
[edit services service-set service-set-name syslog]
```

Description

Specify the hostname for the system logging utility.

Starting in Junos OS Release 17.4R1, you can configure up to a maximum of four system log servers (combination of local system log hosts and remote system log collectors) for each service set at the `[edit services service-set service-set-name]` hierarchy level.

> **NOTE**: Starting with Junos OS release 14.1X55, 14.2R5, 15.1R3, and 16.1R1, for multiservices (ms-) interfaces, you cannot configure system logging for PCP and ALGs by including the pcp-logs and alg-logs statements at the `[edit services service-set service-set-name syslog host hostname class]` hierarchy level. An error message is displayed if you attempt to commit a configuration that contains the pcp-logs and alg-logs options to define system logging for PCP and ALGs for ms- interfaces.

Options

hostname—Name of the system logging utility host machine.

From Junos OS Release 17.4R1, you can configure up to four system log hosts.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`class` option introduced in Junos OS Release 13.2.

You can configure multiple system log hosts from Junos OS Release 17.4R1 onwards.

RELATED DOCUMENTATION

Configuring System Logging for Service Sets

host (System)

IN THIS SECTION

- Syntax | 8171
- Hierarchy Level | 8173
- Description | 8173
- Options | 8173
- Required Privilege Level | 8173
- Release Information | 8174

Syntax

```
host (hostname | other-routing-engine) {
    facility severity;
    exclude-hostname
    explicit-priority;
    facility-override facility;
    log-prefix string;
    match "regular-expression";
    match-strings string-name;
    source-address source-address
    structured-data {
        brief;
    }
    transport (System Syslog)
}
```

QFX Series and OCX Series

```
host (hostname {
    facility severity;
    explicit-priority;
    facility-override facility;
    log-prefix string;
    match "regular-expression";
    match-strings string-name;
    port;
    source-address source-address
}
```

TX Matrix Router and EX Series Switches

```
host (hostname | other-routing-engine | scc-master) {
    facility severity;
    explicit-priority;
    facility-override facility;
    log-prefix string;
    match "regular-expression";
    match-strings string-name;
    port;
    source-address source-address
}
```

TX Matrix Plus Router

```
host (hostname | other-routing-engine | sfc0-master) {
    facility severity;
    allow-duplicates;
    explicit-priority;
    facility-override facility;
    log-prefix (System) string;
    match "regular-expression";
    match-strings string-name;
    port;
    source-address source-address
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name system syslog],
[edit system syslog]
```

Description

Configure the logging of system messages to a remote destination.

Options

facility—Class of messages to log. To specify multiple classes, include multiple *facility severity* statements. For a list of the facilities, see *Junos OS System Logging Facilities*.

hostname—IPv4 address, IPv6 address, or fully qualified hostname of the remote machine to which to direct messages. To direct messages to multiple remote machines, include a `host` statement for each one.

`other-routing-engine`—Direct messages to the other Routing Engine on a router or switch with two Routing Engines installed and operational.

> **NOTE**: The `other-routing-engine` option is not applicable to the QFX Series and OCX Series.

`port`—Port number of the remote syslog server that can be modified.

`scc-master`—(TX Matrix routers only) On a T640 router that is part of a routing matrix, direct messages to the TX Matrix router.

severity—Severity of the messages that belong to the facility specified by the paired *facility* name. Messages with severities of the specified level and higher are logged. For a list of the severities, see *System Log Message Severity Levels*.

`sfc0-master`—(TX Matrix Plus routers only) On a T1600 or T4000 router that is part of a routing matrix, direct messages to the TX Matrix Plus router.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Directing System Log Messages to a Remote Machine or the Other Routing Engine*
>
> *Directing Messages to a Remote Destination from the Routing Matrix Based on the TX Matrix Router*
>
> *Directing Messages to a Remote Destination from the Routing Matrix Based on a TX Matrix Plus Router*

host-vpn

IN THIS SECTION

- Syntax | **8174**
- Hierarchy Level | **8176**
- Description | **8176**
- Required Privilege Level | **8176**
- Release Information | **8176**

Syntax

```
host-vpn {
    connections {
        connection-name {
            children {
                child-name {
                    esp-proposal esp-proposal;
                    local-traffic-selector {
                        (ipv4-prefix ipv4-prefix | ipv6-prefix ipv6-prefix);
                        port port;
                        protocol protocol;
                    }
```

```
                    mode (transport | tunnel);
                    rekey-time rekey-time;
                    remote-traffic-selector {
                        (ipv4-prefix ipv4-prefix | ipv6-prefix ipv6-prefix);
                        port port;
                        protocol protocol;
                    }
                }
            }
            dpd-delay dpd-delay;
            ike-proposal ike-proposal;
            local {
                id local-id;
            }
            local-address {
                (ipv4 ipv4-address | ipv6 ipv6-address);
            }
            rekey-time rekey-time;
            remote {
                id remote-id;
            }
        }
        remote {
            id remote-id;
        }
        remote-address {
            (ipv4 ipv4-address | ipv6 ipv6-address);
        }
    }
    ike-log {
        filename filename;
        level level;
    }
    ike-secrets {
        ike-secret {
            id id;
            secret (ascii-text ascii-text | hexadecimal hexadecimal);
        }
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure a host-to-host VPN type of IPsec connection.

The remaining statements are explained separately.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1 Evolved.

RELATED DOCUMENTATION

show security host-vpn security-associations | 24473
clear security host-vpn security-associations | 3956
show security host-vpn version

hosted-service-identifier

IN THIS SECTION

Syntax | 8177
Hierarchy Level | 8177
Description | 8177
Options | 8177

Required Privilege Level | 8177
Release Information | 8177

Syntax

```
hosted-service-identifier identifier;
```

Hierarchy Level

```
[edit services hosted-services server-profile server-profile-name]
```

Description

Configure the identifier for the service performed on the remote server

Options

identifier—Identifer for the service performed on the remote server.

- **Range:** 1 through 63

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

hosted-services

IN THIS SECTION

- Syntax | 8178
- Hierarchy Level | 8178
- Description | 8178
- Options | 8178
- Required Privilege Level | 8179
- Release Information | 8179

Syntax

```
hosted-services {
    server-profile server-profile-name {
        client-address ipv4-address;
        server-address ipv4-address;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure services performed on the remote server.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

hostkey-algorithm

IN THIS SECTION

- Syntax (Prior to Junos OS Release 22.3R1) | **8179**
- Syntax (Starting in Junos OS Release 22.3R1) | **8180**
- Hierarchy Level | **8180**
- Description | **8180**
- Options | **8181**
- Required Privilege Level | **8181**
- Release Information | **8182**

Syntax (Prior to Junos OS Release 22.3R1)

```
hostkey-algorithm {
    (no-ssh-dss | ssh-dss);
    (no-ssh-rsa | ssh-rsa);
    (no-ssh-ecdsa | ssh-ecdsa);
```

```
    (no-ssh-ed25519 | ssh-ed25519);
}
```

Syntax (Starting in Junos OS Release 22.3R1)

```
hostkey-algorithm-list (ecdsa-sha2-nistp256 | ecdsa-sha2-nistp384 | ecdsa-sha2-nistp521 |
ed25519 | rsa)
```

Hierarchy Level

```
[edit system services ssh]
```

Description

Allow or disallow a host-key algorithm to authenticate another host through the SSH protocol. The host-key uses RSA, ECDSA, ED25519, and DSS algorithms.

The following are the behaviors when the `hostkey-algorithm` option is configured with SSH client and SSH server:

- On the SSH client, the host-key algorithms that are supported when talking to a server are:
  1. RSA: Equal or greater-than to 1024 bit
  2. ECDSA: 256, 384, or 521 bit
  3. ED25519: 256 bit
  4. DSS: 1024 bit
- On the SSH server, the host-key algorithms that are generated and stored are:
  1. RSA: 2048 bit
  2. ECDSA: 256 bit (Prior to Junos OS Release 22.3R1).

     ECDSA: 256, 384, or 521 bit (Starting in Junos OS Release 22.3R1).
  3. ED25519: 256 bit
  4. DSS: 1024 bit

Starting in Junos OS Release 22.3R1, we’ve introduced the `hostkey-algorithm-list` statement at the [`edit system services ssh`] hierarchy level. This enhancement enables you to configure only the specified SSH hostkey algorithms. The system automatically disables the remaining unspecified hostkey algorithms. In earlier releases, you need to disable the hostkey algorithms explicitly. All the hostkey algorithms at this hierarchy enabled by default. The DSS algorithm is no longer available at this new hierarchy. In addition, we've deprecated the `hostkey-algorithm` statement at the [`edit system services ssh`] hierarchy level.

Starting in Junos OS Release 24.2R1, the `hostkey-algorithm` statement at the [`edit system services ssh`] hierarchy level is removed and can no longer be used.

Options

ecdsa-sha2-nistp256 Allow generation of ECDSA host-key with NIST P-256 curve.

ecdsa-sha2-nistp384 Allow generation of ECDSA host-key with NIST P-384 curve.

ecdsa-sha2-nistp521 Allow generation of ECDSA host-key with NIST P-521 curve

ed25519 Allow generation of EdDSA host-key with curve25519.

rsa Allow generation of 2048-bit RSA host-key

ssh-ecdsa Allow generation of an ECDSA host-key. Key pair sizes of 256, 384, or 521 bits are compatible with ECDSA.

ssh-dss Allow generation of a 1024-bit DSA host-key.

> **NOTE**: DSA keys are not supported in FIPS, so the `ssh-dss` option is not available on systems operating in FIPS mode.

ssh-rsa Allow generation of RSA host-key. Key pair sizes greater than or equal to 1024 are compatible with RSA.

no-ssh-dss Do not allow generation of a 1024-bit Digital Signature Algorithm (DSA) host-key.

no-ssh-ecdsa Do not allow generation of an Elliptic Curve Digital Signature Algorithm (ECDSA) host-key.

no-ssh-rsa Do not allow generation of an RSA host-key.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

`hostkey-algorithm-list` option added in Junos OS Release 22.3R1.

`hostkey-algorithm` option removed in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Generating SSL Certificates for Secure Web Access (SRX Series Devices)

Generating a Self-Signed SSL Certificate Automatically

hostkey-algorithm-list

IN THIS SECTION

- Syntax | 8182
- Hierarchy Level | 8183
- Description | 8183
- Default | 8183
- Options | 8183
- Required Privilege Level | 8183
- Release Information | 8183

Syntax

```
hostkey-algorithm-list (ecdsa-sha2-nistp256 | ecdsa-sha2-nistp384 | ecdsa-sha2-nistp521 |
ed25519 | rsa)
```

Hierarchy Level

```
[edit system services ssh]
```

Description

Configure the specified SSH hostkey algorithms. The system automatically disables the remaining unspecified hostkey algorithms.

> **NOTE**: This configuration statement replaces the `system services ssh hostkey-algorithm` statement starting in Junos OS Release 22.3R1.

Default

All hostkey algorithms are enabled by default.

Options

ecdsa-sha2-nistp256	Allow generation of ECDSA host-key with NIST P-256 curve.
ecdsa-sha2-nistp384	Allow generation of ECDSA host-key with NIST P-384 curve.
ecdsa-sha2-nistp521	Allow generation of ECDSA host-key with NIST P-521 curve
ed25519	Allow generation of EdDSA host-key with curve25519.
rsa	Allow generation of 2048-bit RSA host-key

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

> *Generating SSL Certificates for Secure Web Access (SRX Series Devices)*

hot-root-standby (MBGP MVPN)

IN THIS SECTION

- Syntax | 8184
- Hierarchy Level | 8184
- Description | 8184
- Required Privilege Level | 8186
- Release Information | 8186

Syntax

```
hot-root-standby {
    min-rate <rate>;
    source-tree;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
mvpn],
[edit routing-instances routing-instance-name protocols mvpn]
```

Description

In a BGP multicast VPN (MVPN) with either RSVP-TE point-to-multipoint or MLDP point-to-multipoint provider tunnels, configure hot-root standby, as defined in *Multicast VPN fast upstream failover*, draft-morin-l3vpn-mvpn-fast-failover-05.

Starting in Junos OS Release 21.1R1, you can configure MLDP point-to-multipoint provider tunnel on MX Series router.

Hot-root standby enables an egress PE router to select two upstream PE routers for an (S,G) and send C-multicast joins to both the PE routers. Multiple ingress PE routers then receive traffic from the source and forward into the core. The egress PE router uses sender-based RPF to forward the one stream received by the primary upstream PE router.

When `hot-root-standby` is configured, based on local policy, as soon as the PE router receives this standby BGP customer multicast route, the PE can install the VRF PIM state corresponding to this BGP source-tree join route. The result is that join messages are sent to the CE device toward the customer source (C-S0, and the PE router receives (C-S,C-G) traffic. Also, based on local policy, as soon as the PE router receives this standby BGP customer multicast route, the PE router can forward (C-S, C-G) traffic to other PE routers through a P-tunnel independently of the reachability of the C-S through some other PE router.

The receivers must join the source tree (SPT) to establish a hot-root standby. Customer multicast join messages continue to be sent to a single upstream provider edge (PE) router for shared-tree state, and duplicate data does not flow through the core in this case.

In case there are more than two PE routers multihomed to the source, backup UMH selection is based on the same algorithm used to select the primary UMH.

There is no limitation in Junos OS on which upstream multicast hop (UMH) selection method is used. For example, you can use `static-umh (MBGP MVPN)` or `unicast-umh-election` .

PIM dense mode as the customer multicast protocol is not supported.

Hot-root standby is supported for RSVP point-to-multipoint and mLDP point-to-multipoint provider tunnels. Other provider tunnels are not supported. A commit error results if `hot-root-standby` is configured and the provider-tunnel is not either RSVP point-to-multipoint or mLDP point-to-multipoint.

Fast failover (sub 50ms) is supported for C-multicast streams within NG-MVPNs in a hot-standby mode. The threshold to trigger fast failover must be set. See "min-rate" on page 10714 for information on fast failover.

When you configure `hot-root-standby` on MPC10 or MPC11 linecards, the failover process takes up to 150 milliseconds.

Cold-root standby and warm-root standby, as specified in draft Morin, are not supported.

The backup attribute is not sent in the customer multicast routes, as this is only needed for warm and cold-root standby.

Internet multicast is not supported with hot-root standby.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support for MLDP point-to-multipoint provider tunnel is introduced in Junos OS Release 21.1R1 for MX Series router.

RELATED DOCUMENTATION

Understanding Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

Example: Configuring Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

sender-based-rpf (MBGP MVPN) | **16552**

unicast-umh-election | **30353**

Example: Configuring Sender-Based RPF in a BGP MVPN with MLDP Point-to-Multipoint Provider Tunnels

hot-standby (Routing Instance)

IN THIS SECTION

- Syntax | **8187**
- Hierarchy Level | **8187**
- Description | **8187**
- Required Privilege Level | **8187**
- Release Information | **8187**

Syntax

```
hot-standby;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols l2vpn site site-name]
```

Description

Turn on the protector behavior for the site. This keeps backup pseudowire in continuous standby mode and ready for traffic forwarding.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Example: Configuring MPLS Egress Protection Service Mirroring for BGP Signaled Layer 2 Services

hot-standby

IN THIS SECTION

- Syntax | 8188
- Hierarchy Level | 8188

- Description | 8188
- Required Privilege Level | 8188
- Release Information | 8188

Syntax

```
hot-standby;
```

Hierarchy Level

```
[edit interfaces rlsqnumber redundancy-options],
[edit interfaces rlsqnumber:number redundancy-options]
[edit interfaces rspnumber redundancy-options]
[edit interfaces rmsnumber redundancy-options]
```

Description

For one-to-one AS, rsp, or rms redundancy configurations, specify that the failure detection and recovery must take place in less than 5 seconds. For FRF.15 (MLFR) and FRF.16 (MFR) configuration, specify the switch over time of 5 seconds and less for FRF.15 and a maximum of 10 seconds for FRF.16.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces

hot-standby (Protocols Layer 2 Circuit)

IN THIS SECTION

- Syntax | 8189
- Hierarchy Level | 8189
- Description | 8189
- Required Privilege Level | 8190
- Release Information | 8190

Syntax

```
hot-standby;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name end-interface interface interface-name backup-neighbor address],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name backup-neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor address backup-neighbor address],
[edit protocols l2circuit neighbor address interface interface-name backup-neighbor address],
[edit routing-instances routing-instance-name protocols vpls neighbor address backup-neighbor
address]
```

Description

Configure the pseudowire to the specified backup neighbor as the hot-standby. When you configure this statement, traffic flows over both the active and hot-standby pseudowires to the backup device (either a CE device or PE router). The backup device drops the traffic from the hot-standby pseudowire, unless

the active pseudowire fails. If the active pseudowire fails, the backup device automatically switches to the hot-standby pseudowire.

Configure the `hot-standby` statement on routers that have both active and standby virtual circuits. Generally, these are access routers. On provider edge routers, configure the `hot-standby-vc-on'` statement to indicate that a hot-standby pseudowire is desired upon arrival of a PW_FWD_STDBY status-TLV.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Hierarchy levels associated with the `backup-neighbor` statement added in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Layer 2 Circuits over Both RSVP and LDP LSPs

Configuring Pseudowire Redundancy on the PE Router

Example: Configuring Layer 2 Circuit Switching Protection

hot-standby-vc-on (Protocols Layer 2 Circuit)

IN THIS SECTION

- Syntax | **8191**
- Hierarchy Level | **8191**
- Description | **8191**
- Required Privilege Level | **8191**
- Release Information | **8192**

Syntax

```
hot-standby-vc-on;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name pseudowire-status-tlv],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name neighbor address pseudowire-status-tlv],
[edit protocols l2circuit neighbor address interface interface-name pseudowire-status-tlv],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name neighbor
address pseudowire-status-tlv]
```

Description

On provider edge (PE) aggregation routers, configure the `hot-standby-vc-on` statement to indicate that a hot-standby pseudowire is desired upon arrival of a PW_FWD_STDBY status-tlv. This flag indicates the standby state. Configure in conjunction with the `hot-standby` statement on routers that have both active and standby virtual circuits. Generally, these are access routers.

The goal of the `hot-standby` statement is to reduce the amount of traffic being discarded during primary-to-backup transition periods. This statement enables the possibility of keeping both the active and standby paths open within the Layer 2 domain. By having both the active and standby VCs able to send and receive traffic, traffic loops could potentially occur within the Layer 2 domain. In consequence, Layer 2 VCs are kept open only in the PE-to-access direction. In other words, aggregation PE devices can send traffic toward access devices, but access devices send traffic exclusively through the active VC.

In this regard, the `hot-standby` statement is quite similar to the `standby` statement. The `hot-standby` statement allows for a faster forwarding-path switchover during transition periods, as compared to what is allowed by the `standby` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Support for VPLS routing instances added in Junos OS Release 15.1R2.

RELATED DOCUMENTATION

Example: Configuring Pseudowire Redundancy in a Mobile Backhaul Scenario

hot-standby (Protocols Layer 2 Circuit) | **8189**

http (Access)

IN THIS SECTION

- Syntax | **8192**
- Hierarchy Level | **8193**
- Description | **8193**
- Options | **8193**
- Required Privilege Level | **8193**
- Release Information | **8193**

Syntax

```
http {
    banner {
        fail string;
        login string;
        success string;
    }
}
```

Hierarchy Level

```
[edit access firewall-authentication pass-through]
```

Description

Configure banners for the HTTP login prompt, successful authentication, and failed authentication.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Obtaining Username and Role Information Through Firewall Authentication

http (icap-redirect profile)

IN THIS SECTION

- Syntax | 8194
- Hierarchy Level | 8194
- Description | 8194

- Options | 8194
- Required Privilege Level | 8195
- Release Information | 8195

Syntax

```
http {
    redirect-request;
    redirect-response;
}
```

Hierarchy Level

```
[edit services icap-redirect profile name],
[edit logical-systems logical-system-name services icap-redirect profile name]
```

Description

Enable the redirect request and the redirect response for the HTTP traffic.

You can forward HTTP requests and HTTP responses to a Internet Content Adaptation Protocol (ICAP) server before sending a request to a Web server or returning a response to the client system.

The SRX Series Firewall decrypts the HTTPS traffic and redirects the HTTP message to a third-party, on-premise, DLP server using the ICAP channel. After DLP processing, the traffic is reflected back to the SRX Series Firewall.

Options

redirect-request	Enable the redirect service on HTTP request
redirect-response	Enable the redirect service on HTTP response

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.1R1.

The logical system option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Example: Configuring ICAP Redirect Service on SRX Series Firewalls

http-log-multiple-transactions (LRF Profile)

IN THIS SECTION

- Syntax | 8195
- Hierarchy Level | 8195
- Description | 8196
- Required Privilege Level | 8196
- Release Information | 8196

Syntax

```
http-log-multiple-transactions;
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Description

Configure HTTP transaction logging to generate and send HTTP metadata for each transaction of a data session. This option is only relevant if the template specified in an LRF rule includes `http` in the `template-type`.

By default, HTTP transaction logging is disabled, and the HTTP transaction records for a TCP session are sent together as one group of records.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

http-persist

IN THIS SECTION

- Syntax | 8197
- Hierarchy Level | 8197
- Description | 8197
- Required Privilege Level | 8197
- Release Information | 8197

Syntax

```
http-persist;
```

Hierarchy Level

```
[edit security utm default-configuration web-filtering]
[edit security utm feature-profile web-filtering]
```

Description

Checks all HTTP requests in a connection. By default, Web filtering first checks the HTTP request method (for example, GET or PUT) in the same session. If there are multiple HTTP request methods in the subsequent HTTP request of the same session, then Web filtering checks are not performed on these methods. If `http-persist` command is enabled for clear text HTTP traffic, then Web filtering checks every HTTP request packet in the same session.

Required Privilege Level

view

Release Information

Statement introduced in Junos OS Release 12.3X48-D25.

The `[edit security utm default-configuration web-filtering]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Example: Configuring Enhanced Web Filtering

http-profile (Security Antivirus)

IN THIS SECTION

- Syntax | 8198
- Hierarchy Level | 8198
- Description | 8198
- Required Privilege Level | 8198
- Release Information | 8198

Syntax

```
http-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name anti-virus]
```

Description

Configure a Content Security policy for the antivirus HTTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The [edit security utm default-configuration] hierarchy level is introduced in Junos OS Release 18.2R1.

http-profile (Security Content Filtering)

IN THIS SECTION

- Syntax | 8199
- Hierarchy Level | 8199
- Description | 8199
- Required Privilege Level | 8199
- Release Information | 8200

Syntax

```
http-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name content-filtering]
```

Description

Configure a Content Security policy for the content-filtering HTTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

http-profile (Security Web Filtering)

IN THIS SECTION

- Syntax | 8200
- Hierarchy Level | 8200
- Description | 8201
- Required Privilege Level | 8201
- Release Information | 8201

Syntax

```
http-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name web-filtering]
```

Description

Configure a Content Security policy for the Web-filtering HTTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Web Filtering Overview

http-reassemble

IN THIS SECTION

- Syntax | 8201
- Hierarchy Level | 8202
- Description | 8202
- Required Privilege Level | 8202
- Release Information | 8202

Syntax

```
http-reassemble;
```

Hierarchy Level

```
[edit security utm default-configuration web-filtering]
[edit security utm feature-profile web-filtering]
```

Description

Reassembles HTTP requests segments. When the `http-reassemble` option is enabled the requested fragment is reassembled. By default, Web filtering checks only HTTP requests in the first HTTP request packet. If HTTP request methods and URLs are fragmented in different packets, then these URLs are not checked. If `http-reassemble` option is enabled for clear text HTTP traffic, then Enhanced Web Filtering (EWF) reassembles the fragmented HTTP request to avoid evasion instead of packet-based inspection.

When a new URL is matched against the active Web Filtering profile and the profile dictates that the URL should be dropped, the entire HTTP session will be blocked by the device.

Required Privilege Level

view

Release Information

Statement introduced in Junos OS Release 12.3X48-D25.

The `[edit security utm default-configuration web-filtering]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Example: Configuring Enhanced Web Filtering

http (REST API)

IN THIS SECTION

Syntax | 8203

- Hierarchy Level | 8203
- Description | 8203
- Required Privilege Level | 8203
- Release Information | 8203

Syntax

```
http {
    addresses [ address-list ];
    port port-number;
}
```

Hierarchy Level

```
[edit system services rest]
```

Description

Specify unencrypted HTTP connection settings, including addresses for incoming connections and the port number.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

http (System Services)

IN THIS SECTION

- Syntax | 8204
- Hierarchy Level | 8204
- Description | 8204
- Options | 8205
- Required Privilege Level | 8205
- Release Information | 8205

Syntax

```
http {
port port;
}
```

Hierarchy Level

```
[edit system services webapi (System Services)]
```

Description

Specify HTTP as the communication protocol for the Web API function of the SRX Series integrated ClearPass authentication and enforcement feature.

The SRX Series Firewall exposes to the ClearPass Policy Manager (CPPM) the Web API for it to use to initiate a connection and then use that connection to send to the SRX Series Firewall user authentication and identity information.

This statement also specifies the port number to use for the HTTP connection. The port number is optional.

If you deploy HTTP along with a Web management application, you must ensure that they run on different service ports.

Options

port Port for HTTP to use for the Web API function.

- **Default:** 8080
- **Range:** 1 through 65535

Required Privilege Level

1. system—To view this statement in the configuration.
2. system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

http (Web Management)

IN THIS SECTION

- Syntax | 8205
- Hierarchy Level | 8206
- Description | 8206
- Options | 8206
- Required Privilege Level | 8206
- Release Information | 8206

Syntax

```
http {
    interface [ interface-names ];
```

```
    port port;
}
```

Hierarchy Level

```
[edit system services web-management]
```

Description

Configure the port and interfaces for the HTTP service, which is unencrypted.

Options

interface [*interface-names*] Specify the name of one or more interfaces on which to accept access through the HTTP service. By default, HTTP access is allowed through built-in Fast Ethernet or Gigabit Ethernet interfaces only.

port *port-number* Configure the TCP port number on which to connect the HTTP service.

- **Range:** 1 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 8.5 for SRX Series.

vSRX Virtual Firewall 3.0 on Hyper-V does not support the web management https configuration.

RELATED DOCUMENTATION

- Configuring Management Access for the EX Series Switch (J-Web Procedure)
- Secure Management Access Configuration Summary

Firewall User Authentication Overview

https (Web Management) | **8211**

https (REST API)

IN THIS SECTION

- Syntax | **8207**
- Hierarchy Level | **8207**
- Description | **8208**
- Required Privilege Level | **8208**
- Release Information | **8208**

Syntax

```
https {
    addresses [ address-list ];
    cipher-list [cipher-1 cipher-2 cipher-3 ... ];
    mutual-authentication {
        certificate-authority certificate-authority-profile-name;
    }
    port port-number;
    server-certificate local-certificate-identifier;
}
```

Hierarchy Level

```
[edit system services rest]
```

Description

Specify encrypted HTTPS connection settings, including addresses for incoming connections, the port number, preferred cipher suites, and server certificate.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

https (System Services)

IN THIS SECTION

- Syntax | 8208
- Hierarchy Level | 8209
- Description | 8209
- Options | 8210
- Required Privilege Level | 8210
- Release Information | 8210

Syntax

```
https {
    certificate certificate;
     certificate-key certificate-key;
    default-certificate;
    pki-local-certificate pki-local-certificate;
```

```
    port port;
}
```

Hierarchy Level

```
[edit system services webapi (System Services)]
```

Description

Specify HTTPS as the communication protocol for the Web API function of the SRX Series integrated ClearPass authentication and enforcement feature. When you configure HTTPS, you specify the service certificate and certificate key. You can also specify the port to be used.

The Web API process, acting as an HTTPS server, allows the ClearPass Policy Manager (CPPM), acting as the client, to send POST request messages to HTTPS server. The CPPM, which is the authentication source for this feature, sends to the SRX Series Firewall user authentication and identity information.

If you deploy HTTPS with a Web management application, ensure that they run on different service ports.

certificate- Configures a custom certificate to be used for the Integrated ClearPass Authentication and Enforcement feature Web API (webapi) configuration when the HTTPS protocol is configured.

When you configure the Web API (webapi) function to use HTTPS, you can use the default certificate, a custom one, or a certificate generated by the PKI local store.

If you configure a custom certificate, you must configure a certificate key with it. Here is an example of how to configure a certificate and certificate key:

```
set system services webapi https certificate /var/tmp/certificate.crt
set system services webapi https certificate-key /var/tmp/certificate.key
```

The Web API supports only the Privacy-Enhanced Mail (PEM) format for the custom certificate and certificate key.

certificate-key- Configures the filename of the certificate key to use with the specified custom certificate for the Web API (webapi) HTTPS configuration. A certificate key is required if a custom certificate file is used.

The Integrated ClearPass Authentication and Enforcement feature Web API supports only the PEM format for the custom certificate and certificate key.

default-certificate- Specify that the default certificate is to be used for the integrated ClearPass authentication and enforcement Web API process (webapi) HTTPS configuration. To ensure security, the Junos OS default certificate key size is 2084 bits.

pki-local-certificate- Configure the Web API process to use the local X.509 PKI certificate for HTTPS when HTTPS is specified as the communication protocol. The SRX Series integrated ClearPass authentication and enforcement feature exposes the Web API to the ClearPass Policy Mananger (CPPM) to allow the CPPM to initiate a connection to the SRX Series Firewall. For this feature, ClearPass acts as the authentication source. The CPPM uses the HTTPS connection to send user authentication and identity information to the SRX Series Firewall.

port- Specify the SRX Series Firewall TCP port to use for incoming HTTP or HTTPS connection requests initiated by the ClearPass Policy Manager (CPPM). The SRX Series Firewall integrated ClearPass authentication and enforcement feature exposes its Web API (webapi) to the CPPM. The CPPM uses the Web API to establish a connection to the SRX Series Firewall and send user authentication and identity information to it.

Options

`certificate`	Configures the Web API process to use the specified, custom certificate file.
`certificate-key`	Configures the Web API process service certificate key. This parameter is required if a custom service certificate file is configured.
`default-certificate`	Configures the Web API process (webapi) to use the default HTTPS certificate.
`pki-local-certificate`	Configures the Web API process to use the local X.509 PKI certificate.
`port`	Configures the HTTPS service port. • **Range:** For port number, 1 through 65,535. • **Default:** For port, 8443.

Required Privilege Level

1. system—To view this statement in the configuration.
2. system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

https (Web Management)

IN THIS SECTION

- Syntax | 8211
- Hierarchy Level | 8211
- Description | 8211
- Options | 8212
- Required Privilege Level | 8212
- Release Information | 8213

Syntax

```
https {
    interface [ interface-names ];
    ( local-certificate name | pki-local-certificate name | system-generated-certificate );
    port port;
    virtual-domain virtual-domain-name {
        pki-local-certificate pki-local-certificate-name;;
}
}
```

Hierarchy Level

```
[edit system services web-management]
```

Description

Configure the secure version of the HTTP service, HTTPS, which is encrypted.

Options

interface [*interface-names*]

Specify the name of one or more interfaces on which to accept access through the HTTPS service. By default, HTTPS access is allowed through any ingress interface, but HTTP access is allowed through built-in Fast Ethernet or Gigabit Ethernet interfaces only.

(local-certificate *name* | pki-local-certificate *name* | system-generated-certificate)

Specify the X.509 certificate type for a Secure Sockets Layer (SSL) connection.

- **Values:** Specify one of the following:
  - local-certificate *name*—Specify the name of the X.509 certificate. You configure the local certificate at the `[edit security certificates local]` hierarchy level.
  - pki-local-certificate *name*—(EX, QFX, and SRX Series only) Specify the name of the X.509 certificate that is generated by the public key infrastructure (PKI) and authenticated by a certificate authority (CA).
  - system-generated-certificate—(EX, QFX, and SRX Series only) Automatically generate a self-signed X.509 certificate for enabling the HTTPS service.

port *port-number*

Configure the TCP port number on which to connect the HTTPS service.

- **Range:** 1 through 65,535

virtual-domain [*virtual-domain-name*]

Specify the virtual-domain name for certificate mapping. Name of the virtual-domain in FQDN format (*host.example.com*).

- Alphabets, numerical, dashes, dot, and underscores are allowed but not exceed 249 characters.
- Only DNS or Server hostnames are supported as virtual domain names but not IP addresses.
- Duplicate entries are not allowed.
- Maximum number of allowed virtual domains is 50.

pki-local-certificate [*pki-local-certificate-name*]

Specify the name of the X.509 certificate that is generated to use from the PKI local store for certificate mapping.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`pki-local-certificate` introduced in Junos OS Release 9.1 for SRX Series.

`system-generated-certificate` introduced in Junos OS Release 11.1 for EX Series.

Statement introduced on the SRX5400, SRX5600, and SRX5800 devices starting from Junos OS Release 12.1X44-D10 and on vSRX Virtual Firewall, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

`virtual-domain` statement is introduced in Junos OS Release 23.1R1

RELATED DOCUMENTATION

Configuring Management Access for the EX Series Switch (J-Web Procedure)

Basic Elements of PKI in Junos OS

Enabling HTTPS and XNM-SSL Services on Switches Using Self-Signed Certificates (CLI Procedure)

http (Web Management) | **8205**

hub-assist

IN THIS SECTION

- Syntax | **8214**
- Hierarchy Level | **8214**
- Description | **8214**
- Options | **8214**
- Required Privilege Level | **8214**
- Release Information | **8214**

Syntax

```
hub-assist weight;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number radio-router]
```

Description

Configure the weight of the resource factor when calculating an effective interface bandwidth.

Options

weight—Factor used to calculate interface bandwidth.

- **Range:** 0 through 100
- **Default:** 100

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE-Based Radio-to-Router Protocols

hugepages

IN THIS SECTION

- Syntax | 8215
- Hierarchy Level | 8215
- Description | 8215
- Required Privilege Level | 8215
- Release Information | 8215

Syntax

```
hugepages;
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

An option to support of 2 MB and 1 GB size memory pages.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D40.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

hw-db-profile

IN THIS SECTION

- Syntax | 8216
- Syntax (ACX7024) | 8216
- Hierarchy Level | 8217
- Description | 8217
- Options | 8217
- Usage scenarios for modular database (MDB) profile | 8218
- Upgrade and Downgrade Scenarios | 8219
- Release Information | 8219

Syntax

```
hw-db-profile
     {
          (lean-edge | bng | carrier-ethernet | cloud-metro);
      }
```

Syntax (ACX7024)

```
hw-db-profile
     {
```

```
            (lean-edge | cloud-metro);
        }
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

Use this configuration to configure various database profiles to allocate tables with different sizes in the hardware. PFE restarts automatically once you change the configuration and commit. Profiles under this configuration support the below databases:

- Large Exact Match (LEM)—This table stores Mac table entries, LSR entries, and IPv4 host entries
- Longest Prefix Match (LPM)—This table stores IPv4 and IPv6 route entries

Options

lean-edge Configures lean-edge profile for Border Network Gateways (BNGs).

bng Selects BNG DB profile and restarts PFE

carrier-ethernet Selects carrier Ethernet high scale MAC DB profile and restarts PFE. Allocation for LPM is less and LEM is more in carrier-ethernet profile.

cloud-metro Selects Cloud-Metro DB profile, restarts PFE

Default profile lean-edge.

Please find the change in profile naming and default as follows:

New MDB Profile	Corresponding old profile
lean-edge (default)	l3-xl
cloud-metro	balanced (default)
carrier-ethernet	l2-xl

bng	balanced-exem

NOTE: For features such as EVPN-VXLAN, if there is a requirement for a higher Layer 2 MAC scale, then it is recommended to migrate to “cloud-metro” profile.

ACX7024 devices support lean-edge and cloud-metro options only.

NOTE: ACX7100-32C, ACX7100-48L, and ACX7509 support FIB compression by default. ACX7024 platform does not support FIB compression.

Usage scenarios for modular database (MDB) profile

Below are the recommended steps to retain an MDB profile which corresponds to the older MDB profile.

Example 1:

To retain the existing default old profile: “balanced” after upgrading to the new software image:

- After the image upgrade to the new software the new default MDB profile is changed to “lean-edge” which corresponds to the old profile l3-xl.
  - Explicitly configure the “cloud-metro” MDB profile by executing the below command:
  - ```
    user@root# set system packet-forwarding-options hw-db-profile cloud-metro
    ```

Example 2: To retain the corresponding MDB profile after upgrading to the newer software:

- The older configured MDB profile, if configured on any network device, is considered and after the upgrade, the corresponding new MDB profile gets activated.
  - If you have configured old profile "balanced-exem", after the upgrade, the corresponding new MDB profile, “bng” is activated.
  - It is recommended that the user explicitly configures the new MDB profile and deletes the older configuration. Use the following commands:

```
user@root# set system packet-forwarding-options hw-db-profile bng
```

```
user@root# delete system packet-forwarding-options hw-db-profile balanced-exem
```
```

Upgrade and Downgrade Scenarios

The below table summarises the expected behaviour during the upgrade and downgrade scenarios with Junos Evolved release 23.4R1.

Scenario	Behaviour
Upgrade to new software when there are no explicit MDB profile configurations present then the default profile is used.	The upgrade will be successful. Default profile will be configured to "lean-edge".
Upgrade to new software when there are existing profile configurations.	MDB profiles are mapped internally with the corresponding new profile.
Downgrade/ rollback to older image from new image if new CLI is present	Downgrade to the older image displays error if new commands are present and upgrade is not completed. In order to successfully downgrade to older version, the new MDB configurations must be removed. If there are no explicit new MDB profile configurations, downgrade to older software is successful. The system runs with the old default profile: "balanced" .
LPM distribution configuration	If lpm-distribution option was configured in the older version software using `set system packet-forwarding-options hw-db-profile lpm-distribution (1|2|3|4|200)` command, upgrade to new software is successful, but LPM distribution configuration is not applied.

Release Information

Statement introduced in Junos OS Evolved Release 23.4R1 for ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L,, ACX7348, ACX7509, ACX753 and ACX7908 routers.

hybrid

IN THIS SECTION

- Syntax | 8220
- Hierarchy Level | 8220
- Description | 8220
- Required Privilege Level | 8220
- Release Information | 8221

Syntax

```
hybrid {
    synchronous-ethernet-mapping {
        clock-source ip-address {
            interface interface1-name;
            interface interface2-name;
        }
    }
}
```

Hierarchy Level

```
[edit protocols ptp slave]
```

Description

Configure hybrid mode.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2R2.

RELATED DOCUMENTATION

Configuring Hybrid Mode and ESMC Quality Level Mapping

Example: Configuring Hybrid Mode and ESMC Quality Level Mapping

Understanding Hybrid Mode

Precision Time Protocol Overview

Synchronous Ethernet Overview

hyper-mode (forwarding-options)

IN THIS SECTION

- Syntax | 8221
- Hierarchy Level | 8221
- Description | 8222
- Required Privilege Level | 8222
- Release Information | 8222

Syntax

```
(hyper-mode | no-hyper-mode);
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure the hyper mode feature to increase the rate at which a data packet is processed. This configuration results in the optimization of the lifetime of a data packet, which further enables the network device (a router or a switch) to provide better performance and throughput. This feature is supported on enhanced MPCs such as MPC3E, MPC4E, MPC5E, and MPC6E.

Configure the `no-hyper-mode` statement when you want to turn off hyper mode and have the device operate in normal mode.

> **NOTE**: The hyper mode feature is configured at the global level and requires a system reboot to turn it on or off. You can enable the feature only if the network-service mode on the device is configured as either `enhanced-ip` or `enhanced-ethernet`.

Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches contains a list of features that are not supported in hyper mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3R4.

RELATED DOCUMENTATION

- *Configuring Hyper Mode on Enhanced MPCs to Speed Up Packet Processing*
- *Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches*
- show forwarding-options hyper-mode | **20100**

I

IN THIS SECTION

- irb-enhanced-scale (Bridge Domain) | **8243**
- icap-redirect | **8245**
- iccp (Traceoptions) | **8247**
- iccp | **8250**
- icl-down-delay | **8252**
- icmp | **8253**
- icmp-code | **8255**
- icmp-code (Applications) | **8256**
- icmp (Error Message Rate Limit) | **8257**
- icmp-fragment-check (IDS MS-MPC) | **8259**
- icmp (IDS Screen Next Gen Services) | **8260**
- icmp-large-packet-check (IDS MS-MPC) | **8262**
- icmp-mapping (Application Identification) | **8263**
- icmp (Security IDP Custom Attack) | **8264**
- icmp (Security IDP Signature Attack) | **8265**
- icmp (Security Screen) | **8267**
- icmp (System) | **8269**
- icmp-type | **8270**
- icmp-type (Applications) | **8271**
- icmp6 (Error Message Rate Limit) | **8273**
- icmpv6-malformed | **8274**
- icmpv6-malformed (IDS Screen Next Gen Services) | **8276**
- icmpv6 (Security IDP) | **8277**
- icmpv6 (Security IDP Custom Attack) | **8278**
- id | **8279**
- id-components | **8281**
- id (MACsec for MX Series) | **8283**
- id (Protocols OpenFlow) | **8284**

- identification (Security ICMP Headers) | **8286**
- identification (Tunnel Profile) | **8287**
- identifier (VPLS Multihoming for FEC 129) | **8288**
- identity | **8290**
- identity-context-profile | **8291**
- identity-management | **8293**
- identity-provider (SAML) | **8298**
- idle-after-switch-over | **8300**
- idle-standby-path-switchover-delay | **8301**
- idle-timeout | **8303**
- idle-timeout (Unified Edge) | **8304**
- idle-timeout (Access) | **8306**
- idle-timeout (L2TP) | **8307**
- idle-timeout (System) | **8309**
- idle-timeout (System Services) | **8311**
- idp (Application Services) | **8312**
- idp (logical-systems) | **8313**
- idp-policy | **8315**
- idp-policy (Application Services) | **8316**
- idp-policy (Security) | **8317**
- idp (Security) | **8321**
- idp (Security Alarms) | **8333**
- idp (SNMP) | **8334**
- ids-option | **8335**
- ids-rule-sets | **8340**
- ids-rules | **8341**
- ie-set | **8343**
- ieee-802.1 (Class of Service) | **8344**
- ieee-802.1 (Classifier on Physical Interface) | **8348**
- ieee-802.1 (Congestion Notification) | **8350**
- ieee-802.1 (Dynamic Classifiers) | **8351**
- ieee-802.1 (Dynamic Rewrite Rules) | **8352**

- ieee-802.1 (Fibre Channel Interfaces) | **8354**
- ieee-802.1 (Host Outbound Traffic) | **8356**
- ieee-802.1 (Input Congestion Notification) | **8357**
- ieee-802.1 (Output Congestion Notification) | **8358**
- ieee-802.1 (Rewrite Rules on Logical Interface) | **8360**
- ieee-802.1 (Rewrite Rules on Physical Interface) | **8362**
- ieee-802.1ad | **8363**
- ieee-802-3az-eee | **8364**
- ieee802.1p | **8366**
- ietf-mode | **8367**
- if-exceeding (QFX) | **8368**
- if-exceeding (EX) | **8370**
- if-exceeding (Hierarchical Policer) | **8371**
- if-exceeding (Policer) | **8373**
- if-exceeding-pps (Hierarchical Policer) | **8374**
- if-exceeding-pps (Policer) | **8376**
- if-route-exists | **8377**
- ifl-subscriber | **8379**
- igmp | **8381**
- igmp (Dynamic Profiles) | **8383**
- igmp-querier (QFabric Systems only) | **8385**
- igmp-snooping | **8387**
- igmp-snooping-options | **8393**
- ignore | **8394**
- ignore-all | **8394**
- ignore-attached-bit | **8396**
- ignore-encapsulation-mismatch | **8397**
- ignore-entry | **8399**
- ignore (Event Policy) | **8399**
- ignore-l3-incompletes | **8400**
- ignore-lsp-metrics | **8402**
- ignore-magic-number-mismatch (Access Group Profile) | **8403**

- ignore-magic-number-mismatch (Dynamic Profiles) | **8405**
- ignore-memory-overflow | **8407**
- ignore-mtu-mismatch | **8408**
- ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-overflow | **8409**
- ignore-reassembly-overflow | **8411**
- ignore-stp-topology-change | **8412**
- ignore-server-validation (Application Identification) | **8413**
- igp-synchronization | **8414**
- ike | **8415**
- ike-access-profile | **8418**
- ike (Access Profile) | **8419**
- ike-esp-nat | **8421**
- ike (High Availability) | **8424**
- ike-log | **8426**
- ike-phase1-failures | **8428**
- ike-phase2-failures | **8429**
- ike-secrets | **8431**
- ike (Security) | **8432**
- ike (Security Group VPN Member) | **8438**
- ike (Security Group VPN Server) | **8440**
- ike (Security IPsec VPN) | **8443**
- ike (Security Services) | **8446**
- image | **8447**
- imap-profile (Security Content Security Policy Antivirus) | **8450**
- imap-profile (Security Content Security Policy Content Filtering) | **8451**
- immediate-accounting-response | **8452**
- immediate-leave | **8454**
- immediate-leave (Dynamic IGMP Interface) | **8456**
- immediate-leave (Dynamic MLD Interface) | **8458**
- immediate-update | **8459**
- import (BGP) | **8461**
- import (Protocols OSPF) | **8463**

- import (Routing Options) | **8465**
- import (Classifiers) | **8466**
- import (CoS) | **8467**
- import-labeled-routes (Routing Instances VPLS) | **8469**
- import (MPLS Traffic Engineering Database) | **8471**
- import-policy | **8473**
- import (Protocols DVMRP) | **8474**
- import (Protocols IS-IS) | **8476**
- import (Protocols LDP) | **8477**
- import (Protocols MSDP) | **8478**
- import (Protocols PIM) | **8480**
- import (Protocols PIM Bootstrap) | **8482**
- import (Protocols RIP) | **8483**
- import (Protocols RIPng) | **8485**
- import (Rewrite Rules) | **8486**
- import-rib | **8488**
- import-target | **8489**
- imsi-prefix | **8491**
- in-place-lsp-bandwidth-update (Protocols MPLS) | **8493**
- in-service | **8494**
- inactive-media-timeout (Security MGCP) | **8496**
- inactive-media-timeout (Security SCCP) | **8497**
- inactive-media-timeout (Security SIP) | **8499**
- inactivity-asymm-tcp-timeout (Service Set Next Gen Services) | **8500**
- inactivity-interval (Fabric Load Balance) | **8501**
- inactivity-probe-duration | **8503**
- inactivity-timeout | **8505**
- inactivity-timeout (services) | **8506**
- inactivity-timeout (Applications) | **8508**
- inactivity-timeout (Security Persistent NAT) | **8509**
- inactivity-timeout (Services RPM) | **8510**
- inband-flow-telemetry | **8512**

- inband-ztp | **8516**
- include-all (for Administrative Groups) | **8518**
- include-all (for Fast Reroute) | **8519**
- include-any (for Administrative Groups) | **8521**
- include-any (for Fast Reroute) | **8522**
- include (Diameter Gx Profiles) | **8524**
- include (Dynamic Access-Line-Identifier VLANs) | **8525**
- include-igp-metric | **8527**
- include-ipv6 (Gx-Plus) | **8528**
- include-irb-and-l2 | **8529**
- include-mp-next-hop | **8532**
- include-option-82 (DHCP Local Server) | **8534**
- include-sci | **8536**
- include-sci (MACsec for MX Series) | **8537**
- inclusive | **8539**
- incoming-map | **8540**
- incoming-queue | **8541**
- independent-domain | **8543**
- indexed-load-balance | **8545**
- indication | **8546**
- indication-polarity | **8548**
- indirect-next-hop | **8549**
- indirect-next-hop-change-acknowledgements | **8551**
- ineligible | **8553**
- inet | **8554**
- inet (AMT Protocol) | **8555**
- inet-bgpmcast | **8557**
- inet (chassis) | **8558**
- inet (CoS) | **8560**
- inet (enhanced-hash-key) | **8561**
- inet (forwarding options) | **8564**
- inet (interfaces (QFX Series)) | **8565**

- inet (Interfaces) | **8568**
- inet-mdt | **8570**
- inet-mdt (Signaling) | **8571**
- inet-mvpn (BGP) | **8574**
- inet-mvpn (VRF Advertisement) | **8575**
- inet (Port Mirroring) | **8576**
- inet-precedence | **8578**
- inet-precedence (Classifier on Physical Interface) | **8580**
- inet-precedence (CoS Classifiers) | **8581**
- inet-precedence (CoS Rewrite Rules) | **8582**
- inet-precedence (Dynamic Classifiers) | **8584**
- inet-precedence (Dynamic Rewrite Rules) | **8585**
- inet-precedence (Rewrite Rules on Physical Interface) | **8587**
- inet (Routing Policy Condition) | **8588**
- inet (Security Static NAT) | **8590**
- inet (Subscriber Secure Policy) | **8591**
- inet (TDF Subscriber Address) | **8593**
- inet (TDF Subscriber Exclude Prefix) | **8594**
- inet6 | **8595**
- inet6-advertise-interval (VRRP) | **8596**
- inet6-backup-router | **8598**
- inet6-bgpmcast | **8599**
- inet6 (CoS Multidestination) | **8601**
- inet6 (enhanced-hash-key) | **8602**
- inet6 (forwarding optons) | **8605**
- inet6 (interfaces (QFX Series)) | **8607**
- inet6 (Interfaces) | **8608**
- inet6-mvpn (BGP) | **8610**
- inet6-mvpn (VRF Advertisement) | **8612**
- inet6-options (Services) | **8613**
- inet6-precedence (CoS Rewrite Rules) | **8614**
- inet6 (Subscriber Secure Policy) | **8616**

- inet6 (TDF Subscriber Address) | **8617**
- inet6 (TDF Subscriber Exclude Prefix) | **8619**
- inet6-vpn | **8620**
- infinity | **8622**
- infranet-controller | **8623**
- ingress (Analyzer) | **8625**
- ingress (Buffer Configuration) | **8626**
- ingress (Chained Composite Next Hop) | **8628**
- ingress-egress-recycle | **8631**
- ingress (LSP) | **8632**
- ingress-node-replication | **8634**
- ingress-node-replication (EVPN) | **8635**
- ingress-only | **8637**
- ingress-policer-overhead | **8638**
- ingress-policy | **8641**
- ingress (Port Mirroring) | **8643**
- ingress-queuing-filter | **8644**
- ingress-rate-limit | **8646**
- ingress-replication | **8647**
- ingress-shaping-overhead | **8649**
- ingress-stats (Flat-File Accounting Options) | **8651**
- ingress (vlans) | **8653**
- inherit-advertisement-interval | **8654**
- init-command-string | **8655**
- init-descriptor | **8657**
- initial-route-check | **8658**
- initial-tcp-mss | **8659**
- initiate-dead-peer-detection | **8661**
- initiate-ncp (Dynamic and Static PPP) | **8662**
- initiation (Services) | **8664**
- inline | **8666**
- inline-fpga-crypto | **8668**

- inline-jflow | **8669**
- inline-jflow (Forwarding Options) | **8671**
- inline-monitoring | **8672**
- inline-services | **8674**
- inline-services (PIC level) | **8676**
- inline-tap | **8677**
- inner-tag-protocol-id | **8679**
- inner-tag-protocol-id (Dynamic VLANs) | **8680**
- inner-vlan-id | **8682**
- inner-vlan-id (Dynamic VLANs) | **8683**
- inner-vlan-id-swap-ranges | **8685**
- input | **8687**
- input (Analyzer) | **8688**
- input-chain | **8690**
- input (Congestion Notification) | **8692**
- input (Congestion Notification) | **8693**
- input-current (T4000 Routers) | **8695**
- input (Dynamic Service Sets) | **8697**
- input-excess-bandwidth-share | **8698**
- input (Fibre Channel Interfaces) | **8700**
- input-firewall-optimized-profile | **8701**
- input (Forwarding Table) | **8703**
- input-hierarchical-policer | **8704**
- input-interface-index | **8705**
- input (Interfaces) | **8707**
- input-list | **8708**
- input-native-vlan-push | **8710**
- input-packet-rate-threshold | **8711**
- input-policer | **8713**
- input (Port Mirroring) | **8714**
- input-priority-map | **8716**
- input (Routing Options RIB) | **8717**

- input (Sampling) | **8718**
- input-scheduler-map | **8720**
- input-shaping-rate (Logical Interface) | **8721**
- input-shaping-rate (Physical Interface) | **8723**
- input-three-color | **8724**
- input-traffic-control-profile | **8726**
- input-traffic-control-profile-remaining | **8728**
- input-vlan-map | **8729**
- input-vlan-map (Aggregated Ethernet) | **8731**
- input-vlan-map (Dynamic Interfaces) | **8732**
- inputs (Analytics) | **8734**
- insert | **8738**
- inspection-profile | **8739**
- install | **8741**
- install-nexthop | **8742**
- install (Protocols MPLS) | **8744**
- install (Routing Options) | **8747**
- instance | **8748**
- instance (Inline Monitoring) | **8750**
- instance (OAM) | **8753**
- instance (CASB) | **8754**
- instance (Ephemeral Database) | **8756**
- instance-export | **8759**
- instance-import | **8760**
- instance (Port Mirroring) | **8762**
- instance-role | **8764**
- instance (Sampling) | **8766**
- instance-shared | **8769**
- instance (Traffic Load Balancer) | **8770**
- instance-type | **8773**
- integer | **8777**
- intelligent-prescreening | **8779**

- intelligent-offload-tunable (IDP) | **8780**
- inter-area-prefix-export | **8782**
- inter-area-prefix-import | **8783**
- inter-as (Routing Instances) | **8785**
- inter-domain | **8788**
- inter-region | **8789**
- inter-region-segmented | **8791**
- inter-region-template | **8793**
- inter-switch-link | **8795**
- interconnect | **8797**
- interconnect-device (Aliases) | **8800**
- interconnect-device (Chassis) | **8802**
- interconnect-device (Flow Groups) | **8804**
- interconnect-multihoming-peer-gateways | **8806**
- interconnected-vlan-list | **8808**
- interconnected-vni-list | **8809**
- interface (services) | **8811**
- interface (VRRP) | **8812**
- interface (Logical Interfaces) | **8814**
- interface (Precision Time Protocol) | **8815**
- interface (LLDP) | **8817**
- interface (802.1X) | **8819**
- interface (Access Port Security) | **8829**
- interface (Accounting or Sampling) | **8831**
- interface (advance-policy-based-routing) | **8833**
- interface (Analyzer) | **8835**
- interface (Backup Selection OSPF) | **8837**
- interface (BOOTP) | **8841**
- interface (BPDU Blocking) | **8842**
- interface (Bridge Domains) | **8845**
- interface (Captive Portal) | **8846**
- interface (Chassis Cluster) | **8850**

- interface (DCBX) | **8851**
- interface-delete (Subscriber Management or DHCP Client Management) | **8853**
- interface-description (DHCP Local Server) | **8854**
- interface-description (DHCP Relay Agent) | **8856**
- interface-description-format | **8858**
- interface (DHCP Local Server) | **8860**
- interface (DHCP Relay Agent) | **8863**
- interface (DHCP Security for MX Series) | **8866**
- interface (DHCP Spoofing Prevention) | **8867**
- interface (DNS, Port, and TFTP Packet Forwarding or Relay Agent) | **8868**
- interface-down | **8870**
- interface (Dynamic IGMP) | **8871**
- interface (Dynamic Interface Sets) | **8873**
- interface (Dynamic MLD) | **8875**
- interface (Dynamic Router Advertisement) | **8877**
- interface (Dynamic Routing Instances) | **8879**
- interface (Dynamic Routing Options) | **8880**
- interface (EVPN Routing Instances) | **8882**
- interface (Export Profiles) | **8884**
- interface (Fibre Channel Fabric) | **8886**
- interface (FIP) | **8887**
- interface (FIP Snooping) | **8889**
- interface-group | **8890**
- interface-group (Decapsulate GRE) | **8892**
- interface-group-down | **8893**
- interface-group (Routing Options) | **8894**
- interface (Hierarchical CoS Schedulers) | **8896**
- interface (Host Fast Reroute) | **8897**
- interface-id | **8898**
- interface (IEEE 802.1ag OAM Connectivity-Fault Management) | **8900**
- interface (IGMP Snooping) | **8902**
- interface-in (Security Forwarding Options) | **8903**

- interface (Layer 2 Protocol Tunneling) | **8904**
- interface (LLDP-MED) | **8906**
- interface-mac-ip-limit | **8908**
- interface-mac-limit | **8910**
- interface-mac-limit (VPLS) | **8913**
- interface-map | **8915**
- interface-mib (Enhanced Subscriber Management) | **8916**
- interface (MLD Snooping) | **8918**
- interface-mode | **8920**
- interface-monitor | **8922**
- interface (Monitoring) | **8923**
- interface (MPLS) | **8925**
- interface (Multicast Scoping) | **8927**
- interface (Multicast Static Routes) | **8928**
- interface (Multichassis Protection) | **8930**
- interface (MVRP) | **8931**
- interface-name | **8934**
- interface-name | **8935**
- interface-name (DHCP Local Server) | **8937**
- interface-name (DHCP Relay Agent) | **8938**
- interface (Next-Hop Group) | **8939**
- interface-none | **8941**
- interface-notification (programmable-rpd) | **8942**
- interface-num | **8943**
- interface (OAM CFM, for EX Series Switch Only) | **8945**
- interface (OAM Link-Fault Management) | **8947**
- interface-out (Security Forwarding Options) | **8948**
- interface (PIC Bundle) | **8949**
- interface (PoE) | **8951**
- interface (Port Mirroring) | **8952**
- interface (Port Mirroring) | **8954**
- interface (Port Mirroring) | **8956**

- interface (Power over Ethernet) | **8957**
- interface-profile | **8959**
- interface (Protocols DVMRP) | **8960**
- interface (Protocols IGMP) | **8962**
- interface (Protocols IPv6 Neighbor Discovery) | **8964**
- interface (Protocols IS-IS) | **8966**
- interface (Protocols Layer 2 Circuit) | **8974**
- interface (Protocols Layer 2 VPN) | **8977**
- interface (Protocols LDP) | **8978**
- interface (Protocols Link Management) | **8980**
- interface (Protocols MLD) | **8981**
- interface (Protocols MPLS) | **8983**
- interface (Protocols OSPF) | **8984**
- interface (Protocols PIM) | **8991**
- interface (Protocols RIFT) | **8994**
- interface (Protocols Router Discovery) | **8998**
- interface (Protocols RSVP) | **8999**
- interface (RA Guard) | **9002**
- interface-range (QFX Series) | **9003**
- interface-range | **9006**
- interface (Redundant Trunk Groups) | **9008**
- interface-routes | **9010**
- interface (Routing Instances) | **9012**
- interface (Routing Options) | **9014**
- interface (Scoping) | **9015**
- interface (Secure Access Port) | **9017**
- interface (Security NAT ARP Proxy) | **9018**
- interface (Security NAT NDP Proxy) | **9020**
- interface (Security Source NAT) | **9021**
- interface (Security Source NAT Rule Set) | **9022**
- interface (Service Interfaces) | **9023**
- interface (Service Redundancy Group Monitor) | **9025**

- interface-service (Services Interfaces) | **9027**
- interface (Services) | **9028**
- interface (Services Flow Tap) | **9029**
- interface (Services PIC) | **9031**
- interface (Session PICs) | **9033**
- interface-set | **9034**
- interface-set (ANCP) | **9036**
- interface-set (Dynamic Profiles) | **9037**
- interface-set (Dynamic VLAN Interface Sets Association) | **9040**
- interface-set (Dynamic VLAN Interface Sets Definition) | **9042**
- interface-set (Ethernet Interfaces) | **9044**
- interface-set (Hierarchical Schedulers) | **9045**
- interface-set (IP Demux Interfaces) | **9046**
- interface-shared | **9048**
- interface-shared-with | **9049**
- interface-shutdown-action | **9050**
- interface (SLAAC Snooping) | **9052**
- interface (SNMP) | **9054**
- interface (SNMP RMON History) | **9056**
- interface (Spanning Tree) | **9057**
- interface-specific | **9060**
- interface-specific (Dynamic Firewalls) | **9061**
- interface-specific (Firewall Filters) | **9063**
- interface-state | **9064**
- interface (Static MAC Bypass) | **9067**
- interface (Static Subscriber Group) | **9068**
- interface (Static Subscriber Username) | **9070**
- interface-status-send-rdi | **9071**
- interface-status-tlv | **9073**
- interface (Storm Control) | **9074**
- interface-switch | **9076**
- interface (Switching Options) | **9078**

- interface (System Services DHCP) | **9079**
- interface-tag (DHCP Local Server) | **9081**
- interface-tag (DHCP Relay Server) | **9083**
- interface-traceoptions (DHCP) | **9085**
- interface-traceoptions (System Services DHCP) | **9087**
- interface-transmit-statistics | **9090**
- interface-type (Protocols OSPF) | **9091**
- interface-type (Protocols RIP) | **9095**
- interface (Unknown Unicast Forwarding) | **9096**
- interface (Virtual Tunnel in Routing Instances) | **9097**
- interface (VoIP) | **9099**
- interface (VPLS Mesh-Group) | **9102**
- interface (VPLS Multihoming for FEC 129) | **9103**
- interface (VPLS Routing Instances) | **9104**
- interface-client-limit (DHCP Local Server) | **9106**
- interface-client-limit (DHCP Relay Agent) | **9108**
- interfaces (VNF) | **9111**
- interfaces | **9113**
- interfaces (Aggregated Multiservices) | **9115**
- interfaces (Analytics) | **9117**
- interfaces (Analytics Resource) | **9119**
- interfaces (ANCP) | **9121**
- interfaces (Autoinstallation) | **9122**
- interfaces (Chassis Synchronization Source) | **9124**
- interfaces (Class of Service) | **9126**
- interfaces (Class of Service) | **9128**
- interfaces (Class of Service) | **9130**
- interfaces (Dynamic CoS Definition) | **9133**
- interfaces (EX Series switches) | **9135**
- interfaces external | **9144**
- interfaces (Junos Node Slicing) | **9145**
- interfaces (MACsec) | **9147**

- interfaces (MACsec for MX Series) | **9149**
- interfaces (OVSDB) | **9151**
- interfaces (Protocols OpenFlow) | **9153**
- interfaces (Q-in-Q Tunneling) | **9154**
- interfaces (QFX Series, ACX Series) | **9156**
- interfaces (Security Zones) | **9165**
- interfaces (Services Dynamic Flow Capture) | **9167**
- interfaces (sFlow) | **9168**
- interfaces (Static and Dynamic Subscribers) | **9170**
- interfaces (Video Monitoring) | **9176**
- interfaces (Voice Services) | **9181**
- interim-logging-interval | **9182**
- interim-rate (Access) | **9183**
- interleave-fragments | **9185**
- internal | **9186**
- internal-forwarding-bandwidth | **9187**
- internal-node | **9189**
- internal (Security IPsec) | **9190**
- internet-options | **9192**
- interpolate (Drop Profiles) | **9197**
- interval | **9198**
- interval (services) | **9201**
- interval (Protocols PIM) | **9202**
- interval (OAM) | **9203**
- interval (RMON) | **9206**
- interval (Accounting Options) | **9207**
- interval (CFM MEP) | **9209**
- interval (EX Series Switch Only) | **9210**
- interval (Flat-File Accounting Options) | **9211**
- interval (Health Monitor) | **9213**
- interval (Interfaces) | **9214**
- interval (PoE) | **9216**

- interval (Power over Ethernet) | **9217**
- interval (Security Antivirus) | **9219**
- interval (Security IDP) | **9220**
- interval (Services) | **9221**
- interval (SNMP RMON) | **9223**
- intra-area-prefix | **9224**
- intra-as | **9226**
- invalid-authentication-entry-timeout (Services User Identification Active Directory and ClearPass) | **9227**
- ioc-npc-connectivity | **9230**
- ip-action (Security IDP Rulebase IPS) | **9231**
- ip-address (Access Profile) | **9234**
- ip-address-first | **9235**
- ip-address-owner | **9236**
- ip-address (Port Mirroring) | **9238**
- ip-can-type (PCRF Partition) | **9239**
- ip (IDS Screen Next Gen Services) | **9241**
- ip-in-udp | **9243**
- ip-monitoring | **9244**
- ip-monitoring (Services) | **9246**
- ip-options-protocol-queue | **9249**
- ip-prefix-routes | **9250**
- ip-prefix-support | **9255**
- ip-protocol-mapping (Application Identification) | **9258**
- ip-query (Identity Management Advanced Query) | **9259**
- ip-reassembly | **9262**
- ip-reassembly (L2TP) | **9263**
- ip-reassembly-rules (Service Set) | **9265**
- ip (Security IDP Custom Attack) | **9266**
- ip (Security Screen) | **9268**
- ip-source-guard | **9272**
- ip-source-guard (MX Series) | **9274**
- ip-subscriber | **9275**

- ip-swap | **9277**
- ip-sweep | **9279**
- ip-tunnel-rpf-check | **9281**
- ip-user-mapping | **9283**
- ip-version | **9285**
- ip-version (Family MPLS) | **9286**
- ipc | **9288**
- ipcp-suggest-dns-option | **9289**
- ipfix-sw-mode | **9291**
- ipip | **9292**
- ipip (lds-option) | **9293**
- ips | **9295**
- ipsec | **9298**
- ipsec-group-vpn (Security Policies) | **9300**
- ipsec (High Availability) | **9301**
- ipsec-inside-interface | **9304**
- ipsec-policy | **9306**
- ipsec-sa | **9307**
- ipsec-sa (Protocols BGP) | **9308**
- ipsec-sa (Security Group VPN) | **9310**
- ipsec (Security) | **9312**
- ipsec (Security Group VPN Member) | **9314**
- ipsec (Security Group VPN Server) | **9317**
- ipsec (Security Interfaces) | **9319**
- ipsec-traffic-selector | **9321**
- ipsec-vpn-options | **9323**
- ipsec-vpn-rules | **9325**
- ipsec-vpn (Security Policies) | **9326**
- ipv4-address (Steering Path) | **9327**
- ipv4-dscp | **9329**
- ipv4 (Family MPLS) | **9330**
- ipv4-flow-table-size | **9332**

- ipv4-mask (HTTP Header Enrichment) | **9335**
- ipv4-or-value (HTTP Header Enrichment) | **9337**
- ipv4-prefix | **9338**
- ipv4 (Security IDP Signature Attack) | **9340**
- ipv4-template | **9344**
- ipv4-template (Services) | **9345**
- ipv6-address (Steering Path) | **9347**
- ipv6-extended | **9348**
- ipv6-extended-attrib | **9349**
- ipv6-extension-header | **9350**
- ipv6-extension-header (IDS Screen Next Gen Services) | **9353**
- ipv6-extension-header-limit | **9355**
- ipv6 (Family MPLS) | **9357**
- ipv6-flow-label | **9359**
- ipv6-flow-table-size | **9360**
- ipv6-malformed-header | **9362**
- ipv6-mask (HTTP Header Enrichment) | **9363**
- ipv6-multicast-interfaces | **9364**
- ipv6-or-value (HTTP Header Enrichment) | **9366**
- ipv6-origin | **9367**
- ipv6-prefix | **9368**
- ipv6-source-guard | **9370**
- ipv6-source-guard-sessions | **9372**
- ipv6-target | **9374**
- ipv6-template | **9375**
- ipv6-template (Services) | **9376**
- ipv6-tunneling | **9377**
- ipvacl-firewall-l2-profile | **9378**
- iq-policing-filter | **9380**
- ir-mode | **9381**
- iracl-firewall-ipv4-profile | **9383**
- irb | **9384**

- irb (Easy EVPN LAG Configuration) | **9386**
- irb (Interfaces) | **9389**
- irb-symmetric-routing | **9394**
- is-ma | **9396**
- isdn-options | **9397**
- isid | **9399**
- isid-list | **9400**
- isis | **9401**
- isis-instance | **9403**
- iso-vpn | **9404**
- isolated | **9406**
- isolated-vlan | **9407**
- isolated-vlan (MX Series) | **9408**
- isolation-id | **9410**
- isolation-vlan-id | **9411**
- iteration-count | **9412**
- iteration-period | **9414**
- ivacl-firewall-no-portrange-profile | **9415**
- ivlan-cfi (RFC 2544 Benchmarking) | **9416**
- ivlan-id (RFC 2544 Benchmarking) | **9417**
- ivlan-priority (RFC 2544 Benchmarking) | **9419**
- iwf-params (tdm-options) | **9420**

irb-enhanced-scale (Bridge Domain)

IN THIS SECTION

- Syntax | **8244**
- Hierarchy Level | **8244**
- Description | **8244**

- Default | 8244
- Options | 8244
- Required Privilege Level | 8245

Syntax

```
bridge-domains <domain-name> {
    irb-enhanced-scale;
}
```

Hierarchy Level

```
[edit bridge-domains <domain-name>]
```

Description

Enable the bridge domain to support enhanced scaling of ARP/ND unicast next hop entries. This configuration allows the bridge domain to handle up to 1.5 million next hop entries for interfaces configured within it.

Default

No default configuration.

Options

domain-name: The name of the bridge domain.

- Range: `<domain-name>` specifies the name of the bridge domain.
- Default: None
- Syntax: <domain-name>
- Values: Any valid bridge domain name.

Required Privilege Level

admin: To view this statement in the configuration.

admin-control: To add this statement to the configuration.

icap-redirect

IN THIS SECTION

- Syntax | 8245
- Hierarchy Level | 8246
- Description | 8246
- Options | 8247
- Required Privilege Level | 8247
- Release Information | 8247

Syntax

```
icap-redirect {
    profile name {
        fallback-option {
            connectivity (block | log-permit | permit);
            default-action (block | log-permit | permit);
            timeout (block | log-permit | permit);
        }
        http {
            redirect-request redirect-request;
            redirect-response redirect-response;
        }
        server name {
            authorization {
                authorization-type authorization-type;
                credentials (ascii ascii | base64 base64);
            }
```

```
            host host;
            port port;
            reqmod-uri reqmod-uri;
            respmod-uri respmod-uri;
            routing-instance ri-name;
            sockets sockets;
            tls-profile tls-profile;
        }
        timeout timeout;
    }
    traceoptions {
        file <filename> <files files>< match match><size size> (world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit services],
[edit logical-systems logical-system-name services],
[edit tenants tenants_name services]
```

Description

Configure the ICAP redirection service.

The SRX Series Firewall acts as an SSL proxy, decrypts HTTP or HTTPS traffic, and redirects the HTTP message to a third-party, on-premise DLP server through the Internet Content Adaptation Protocol (ICAP) channel. To enable ICAP redirection service, you must configure an ICAP redirect profile.

The ICAP server profile allows the ICAP server to process request messages, response messages, fallback options, and so on, to the permitted traffic. This profile is applied as an application service in the security policy.

Starting in Junos OS Release 20.1R1, you can enable ICAP redirect service at the tenant system level, and you can view/clear the ICAP redirect services status and statistics at the tenant systems level. The ICAP service redirect configuration for tenant system is implemented under profile and the ICAP redirect profile capacity is 64 globally. All tenant systems need to share this profile capacity. If 64 tenant systems used the maximum tenants profile capacity, the remaining tenant systems will not be not able to

configure the ICAP redirect profile. Tenant systems can reserve the required or the maximum ICAP redirect profile capacity in their security-profiles using the following CLI commands respectively:

- `edit system security-profile` *`security-profile-name`* `icap-redirect-profile reserved` *`quota`*
- `edit system security-profile` *`security-profile-name`* `icap-redirect-profile maximum` *`quota`*

In addition, we've introduced the `X-Client-IP`, `X-Server-IP`, `X-Authenticated-User`, and `X-Authenticated-Groups` header extensions in an ICAP message to provide information about the source of the encapsulated HTTP message.

Options

The statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1 R1.

Support at the `[edit logical-system` *`logical-system-name`* `services]` hierarchy level introduced in Junos OS Release 18.3R1.

Support at the `[edit tenants` *`tenants_name`* `services]` hierarchy level introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

ICAP Service Redirect

iccp (Traceoptions)

IN THIS SECTION

Syntax | 8248

- Hierarchy Level | 8248
- Description | 8249
- Default | 8249
- Options | 8249
- Required Privilege Level | 8249
- Release Information | 8249

Syntax

```
iccp {
    traceoptions; {
        file <filename> <files number> <match regular-expression> <microsecond-stamp>
<size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
    local-ip-address ip address;
    session-establishment-hold-time value;
    authentication-key string;
    peer ip-address {
        local-ip-address ip address;
        session-establishment-hold-time value;
        authentication-key string;
        redundancy-group-id-list redundancy-group-id-list;
        liveness-detection;
    }
}
```

Hierarchy Level

```
[edit protocols iccp]
[edit logical-systems logical-system-name protocols iccp]
```

Description

Configure Interchassis Control Protocol (ICCP) between the multichassis link aggregation group (MC-LAG) peers. ICCP replicates forwarding information, validates configurations, and propagates the operational state of the MC-LAG members.

Default

If you do not include this statement, no ICCP protocol tracing operations are performed.

Options

`traceoptions`—Set Interchassis Control Protocol (ICCP) tracing options.

`local-ip-address`—Specify the source address where the ICCP packet is routed.

`session-establishment-hold-time`—Specify if the chassis takes over as the primary at the ICCP session.

`authentication-key` —Specify TCP Message Digest 5 (MD5) option for an ICCP TCP session.

`peer ip-address`—Specify the IP address of the peer that hosts an MC-LAG. You must configure ICCP for both peers that host the MC-LAG.

`redundancy-group-id-list`—Specify the redundancy groups between two ICCP peers. The value you assign to this option must match the value you assign to the `redundancy-group-id` option that you configure on the MC-AE interface. If the value differs, when you commit the changes, the system reports a commit check error.

`liveness-detection`—Specify Bidirectional Forwarding Detection (BFD) protocol options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support for logical systems introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring ICCP for MC-LAG

iccp

IN THIS SECTION

- Syntax | 8250
- Hierarchy Level | 8251
- Description | 8251
- Required Privilege Level | 8251
- Release Information | 8251

Syntax

```
iccp {
    authentication-key string;
    local-ip-addr local-ip-addr;
    peer ip-address{
        authentication-key string;
        backup-liveness-detection {
            backup-peer-ip ip-address;
        }
        liveness-detection {
            detection-time {
                threshold milliseconds;
            }
             minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
```

```
            version (1 | automatic);
        }
        local-ip-addr ipv4-address;
        session-establishment-hold-time seconds;
    }
    session-establishment-hold-time seconds;
    traceoptions {
        file <filename> <files number> <match regular-expression> <microsecond-stamp>
<size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure Inter-Chassis Control Protocol (ICCP) between the multichassis link aggregation group (MC-LAG) peers. ICCP replicates forwarding information, validates configurations, and propagates the operational state of the MC-LAG members.

NOTE: Backup liveness detection is not supported on MX Series routers.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

icl-down-delay

IN THIS SECTION

- Syntax | 8252
- Hierarchy Level | 8252
- Release Information | 8252
- Description | 8252
- Options | 8253
- Required Privilege Level | 8253

Syntax

```
icl-down-delay seconds;
```

Hierarchy Level

```
[edit interfaces name multi-chassis-protection IP address of MC-LAG peer]
```

```
[edit multi-chassis multi-chassis-protection IP address of MC-LAG peer]
```

Release Information

Statement introduced in Junos OS Release 19.4R3.

Description

The number of seconds that elapse between when the interchassis link (ICL) goes down and the multichassis aggregated Ethernet interfaces (MCAEs) change to standby mode. Configuring the icl-

down-delay time results in better convergence during a MC-LAG peer reboot when the status control of the peer is active.

Options

`icl-down-delay` *seconds*—The number of seconds that elapse between when the interchassis link (ICL) goes down and the multichassis aggregated Ethernet interfaces (MCAEs) change to standby mode.

- **Range:** 1 through 6000

Required Privilege Level

interface

icmp

IN THIS SECTION

- Junos OS | 8253
- Junos OS Evolved | 8254
- Junos OS Hierarchy Level | 8254
- Junos OS Evolved Hierarchy Level | 8254
- Description | 8254
- Required Privilege Level | 8254
- Release Information | 8254

Junos OS

```
icmp {
    destination-interface interface-name;
}
```

Junos OS Evolved

```
icmp;
```

Junos OS Hierarchy Level

```
[edit services rpm probe-server]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm probe-server]
```

Description

(Required for Junos OS Evolved and for J Series routers running Junos OS) Enable ICMP requests for the RPM probe server. ICMP requests are enabled by default on Junos OS (except for J Series routers); you do not need to explicitly configure them.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

> **NOTE**: The `destination-interface` statement is not supported on PTX Series routers or for Junos OS Evolved.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

> *Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers*

icmp-code

IN THIS SECTION

- Syntax | 8255
- Hierarchy Level | 8255
- Description | 8255
- Options | 8255
- Required Privilege Level | 8256
- Release Information | 8256

Syntax

```
icmp-code value;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Internet Control Message Protocol (ICMP) code value.

Options

value—The ICMP code value. For a complete list, see *Configuring the ICMP Code and Type*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *ALG Descriptions*
- *Configuring Application Sets*
- *Configuring the ICMP Code and Type*
- *Examples: Configuring Application Protocols*
- *Verifying the Output of ALG Sessions*

icmp-code (Applications)

IN THIS SECTION

- Syntax | 8256
- Hierarchy Level | 8257
- Description | 8257
- Options | 8257
- Required Privilege Level | 8257
- Release Information | 8257

Syntax

```
icmp-code value;
```

Hierarchy Level

```
[edit applications application  application-name  ]
[edit applications application  application-name  term  term-name]
```

Description

Specify the Internet Control Message Protocol (ICMP) code value.

Options

value — Specify the Internet Control Message Protocol (ICMP) code value such as `host-unreachable` or `host-unreachable-for-tos`.

Range: 0 through 255.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policy Application Sets Overview

icmp (Error Message Rate Limit)

IN THIS SECTION

Syntax | 8258

- Hierarchy Level | 8258
- Description | 8258
- Options | 8259
- Required Privilege Level | 8259
- Release Information | 8259

Syntax

```
icmp {
    rate-limit rate-limit;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the rate at which ICMP messages are generated for IPv4 packet errors for non-ttl-expired packets. Although you configure the rate limit at the `[edit chassis]` hierarchy level, it is not a chassis-wide limit. Instead, the rate limit applies per interface family. This means, for example, that multiple physical interfaces configured with `family inet` can simultaneously generate the ICMP error messages at the configured rate.

NOTE: This rate limit takes effect only for traffic that lasts 10 seconds or longer. The rate limit is not applied to traffic with a shorter duration, such as 5 seconds or 9 seconds.

NOTE: This statement does not apply to ICMP messages for ttl-expired packets; for these errors, the rate is fixed at 500 pps.

Options

`rate-limit` *`rate-limit`*—Rate in packets per second (pps). Starting in Junos OS Release 18.4R1, the maximum rate is 1000 pps. In earlier releases, the maximum rate is 50 pps.

- **Range:** 1 through 1000
- **Default:** 1

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Rate Limit ICMPv4 and ICMPv6 Error Messages

icmp-fragment-check (IDS MS-MPC)

IN THIS SECTION

- Syntax | 8260
- Hierarchy Level | 8260
- Description | 8260
- Required Privilege Level | 8260
- Release Information | 8260

Syntax

```
icmp-fragment-check;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Identify and drop ICMP packets that are IP fragments. This statement can only be used in IDS rules assigned to a service set on an MS-MPC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

icmp (IDS Screen Next Gen Services)

IN THIS SECTION

Syntax | 8261

- Hierarchy Level | 8261
- Description | 8261
- Required Privilege Level | 8261
- Release Information | 8261

Syntax

```
icmp {
    fragment;
    icmpv6-malformed;
    large;
    ping-death;
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name]
```

Description

Configure ICMP intrusion detection service options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

icmp-large-packet-check (IDS MS-MPC)

IN THIS SECTION

- Syntax | 8262
- Hierarchy Level | 8262
- Description | 8262
- Required Privilege Level | 8262
- Release Information | 8263

Syntax

```
icmp-large-packet-check;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Identify and drop ICMP packets that are larger than 1024. This statement can only be used in IDS rules assigned to a service set on an MS-MPC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

icmp-mapping (Application Identification)

IN THIS SECTION

- Syntax | 8263
- Hierarchy Level | 8263
- Description | 8264
- Required Privilege Level | 8264
- Release Information | 8264

Syntax

```
icmp-mapping {
    code icmp-code;
    order order;
    order-priority (high | low);
    type icmp-type;
}
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Description

Match Internet Control Message Protocol (ICMP) messages identified by unique code and type. This classification is intended to identify and differentiate various types of ICMP messages.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

icmp (Security IDP Custom Attack)

IN THIS SECTION

- Syntax | 8265
- Hierarchy Level | 8265
- Description | 8265
- Required Privilege Level | 8265
- Release Information | 8265

Syntax

```
icmp;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain protocol-binding],

[edit security idp custom-attack attack-name attack-type signature protocol-binding]
```

Description

Allow IDP to match the attack for the specified ICMP.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

icmp (Security IDP Signature Attack)

IN THIS SECTION

- Syntax | 8266
- Hierarchy Level | 8266
- Description | 8266
- Options | 8267
- Required Privilege Level | 8267

Release Information | 8267

Syntax

```
icmp {
    code {
        match (equal | greater-than | less-than | not-equal);
        value code-value;
    }
    data-length {
        match (equal | greater-than | less-than | not-equal);
        value data-length;
    }
    identification (Security ICMP Headers) {
        match (equal | greater-than | less-than | not-equal);
        value identification-value;
    }
    sequence-number (Security IDP ICMP Headers) {
        match (equal | greater-than | less-than | not-equal);
        value sequence-number;
    }
    type (Security IDP ICMP Headers) {
        match (equal | greater-than | less-than | not-equal);
        value type-value;
    }
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol]
```

Description

Allow IDP to match the ICMP header information for the signature attack.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

icmp (Security Screen)

IN THIS SECTION

- Syntax | 8267
- Hierarchy Level | 8268
- Description | 8268
- Options | 8268
- Required Privilege Level | 8268
- Release Information | 8268

Syntax

```
icmp {
    flood {
        threshold number;
    }
    fragment;
    icmpv6-malformed;
    ip-sweep {
        threshold number;
```

```
    }
    large;
    ping-death;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name]
```

Description

Configure ICMP intrusion detection service (IDS) options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Attack Detection and Prevention Overview*
>
> *Example: Configuring Multiple Screening Options*

icmp (System)

IN THIS SECTION

- Syntax | 8269
- Hierarchy Level | 8269
- Description | 8269
- Options | 8269
- Required Privilege Level | 8270
- Release Information | 8270

Syntax

```
icmp {
ttl-expired-source-address ttl-expired-source-address;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure Internet Control Message Protocol (ICMP) features at the system level.

Options

ttl-expired-source-address Use an IPv4 address as the source address for ICMP time-to-live (TTL) expiry error messages.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

Configure ICMP Features

icmp-type

IN THIS SECTION

- Syntax | 8270
- Hierarchy Level | 8270
- Description | 8271
- Options | 8271
- Required Privilege Level | 8271
- Release Information | 8271

Syntax

```
icmp-type value;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

ICMP packet type value.

Options

value—The ICMP type value, such as `echo` or `echo-reply`. For a complete list, see *Configuring the ICMP Code and Type*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *ALG Descriptions*
- *Configuring Application Sets*
- *Configuring the ICMP Code and Type*
- *Examples: Configuring Application Protocols*
- *Verifying the Output of ALG Sessions*

icmp-type (Applications)

IN THIS SECTION

- Syntax | 8272
- Hierarchy Level | 8272
- Description | 8272
- Options | 8272

Required Privilege Level | 8272
Release Information | 8272

Syntax

```
icmp-type value;
```

Hierarchy Level

```
[edit applications application  application-name  ]
[edit applications application  application-name  term  term-name]
```

Description

Specify the ICMP packet type value.

Options

value—ICMP type value, such as `echo` or `echo-reply`.

Range: 0 through 255.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policy Application Sets Overview

icmp6 (Error Message Rate Limit)

IN THIS SECTION

- Syntax | 8273
- Hierarchy Level | 8273
- Description | 8273
- Options | 8274
- Required Privilege Level | 8274
- Release Information | 8274

Syntax

```
icmp6 {
    rate-limit rate-limit;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the rate at which ICMP messages are generated for IPv6 packet errors for non-ttl-expired packets. Although you configure the rate limit at the `[edit chassis]` hierarchy level, it is not a chassis-wide limit. Instead, the rate limit applies per interface family. This means, for example, that multiple physical interfaces configured with `family inet` can simultaneously generate the ICMP error messages at the configured rate.

NOTE: This rate limit takes effect only for traffic that lasts 10 seconds or longer. The rate limit is not applied to traffic with a shorter duration, such as 5 seconds or 9 seconds.

> **NOTE**: This statement does not apply to ICMP messages for ttl-expired packets; for these errors, the rate is fixed at 500 pps.

Options

`rate-limit` *`rate-limit`*—Rate in packets per second (pps). Starting in Junos OS Release 18.4R1, the maximum rate is 1000 pps. In earlier releases, the maximum rate is 50 pps.

- **Range:** 1 through 1000
- **Default:** 1

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Rate Limit ICMPv4 and ICMPv6 Error Messages

icmpv6-malformed

IN THIS SECTION

- Syntax | 8275
- Hierarchy Level | 8275
- Description | 8275
- Options | 8275
- Required Privilege Level | 8275

- Release Information | 8275

Syntax

```
icmpv6-malformed;
```

Hierarchy Level

```
[edit security screen ids-option screen-name icmp]
```

Description

Enable the ICMPv6 malformed intrusion detection service (IDS) option.

Options

This has no options.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Understanding IPv6 Support for Screens

ipv6-extension-header | 9350

icmpv6-malformed (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 8276
- Hierarchy Level | 8276
- Description | 8276
- Required Privilege Level | 8276
- Release Information | 8276

Syntax

```
icmpv6-malformed;
```

Hierarchy Level

```
[edit services screen ids-option screen-name icmp]
```

Description

Identify and drop malformed ICMPv6 packets, which might cause damage to the device and network. Examples of malformed IPv6 packets are packets that are too big (message type 2), that have the next header set to routing (43), or that have a routing header set to hop-by hop.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

icmpv6 (Security IDP)

IN THIS SECTION

- Syntax | 8277
- Hierarchy Level | 8277
- Description | 8277
- Required Privilege Level | 8277
- Release Information | 8278

Syntax

```
icmpv6;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain protocol-binding]
[edit security idp custom-attack attack-name attack-type signature protocol-binding]
```

Description

Specify that the attack is for ICMPv6 packets only.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

icmpv6 (Security IDP Custom Attack)

IN THIS SECTION

- Syntax | 8278
- Hierarchy Level | 8279
- Description | 8279
- Required Privilege Level | 8279
- Release Information | 8279

Syntax

```
icmpv6 {
    checksum-validate {
        match (equal | greater-than | less-than | not-equal);
        value checksum-value;
    }
    code {
        match (equal | greater-than | less-than | not-equal);
        value code-value;
    }
    data-length {
        match (equal | greater-than | less-than | not-equal);
        value data-length;
    }
    identification (Security ICMP Headers) {
        match (equal | greater-than | less-than | not-equal);
        value identification-value;
    }
    sequence-number (Security IDP ICMP Headers) {
        match (equal | greater-than | less-than | not-equal);
```

```
        value sequence-number;
    }
    type (Security IDP ICMP Headers) {
        match (equal | greater-than | less-than | not-equal);
        value type-value;
    }
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol]
```

Description

Allow IDP to match the attack for the specified ICMPv6.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

id

IN THIS SECTION

- Syntax | 8280
- Hierarchy Level | 8280
- Description | 8280
- Options | 8280
- Required Privilege Level | 8280

Release Information | 8280

Syntax

```
id id;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Configure the alphanumeric string that identifies a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.

Options

id—Virtual Chassis ID (VCID), which uses the ISO family address format—for example, **9622.6ac8.5345**.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge

Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge

Configuring an EX4650 or a QFX Series Virtual Chassis

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric

Configuring an EX8200 Virtual Chassis (CLI Procedure)

Understanding Virtual Chassis Member ID Numbering in an EX8200 Virtual Chassis

id-components

IN THIS SECTION

- Syntax | **8281**
- Hierarchy Level | **8282**
- Description | **8282**
- Options | **8282**
- Required Privilege Level | **8283**
- Release Information | **8283**

Syntax

```
id-components {
    use-class;
    use-imsi;
    use-msisdn;
    use-nai;
    use-nas-port;
    use-nas-port-id;
    use-realm;
    use-username;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber subscription-id
subscription-id-options entry-name]
```

Description

Specify a method for constructing the Subscription-Id for the Diameter credit control request (CCR) message that is sent from the TDF to the PCRF for IP-based subscribers belonging to the TDF domain. You may specify more than one option, and the order of preference matches the order in which the options appear.

Options

use-class Subscription-Id-Type is configurable and the Subscription-Id-Data is the entire Class attribute value by default. You can configure a regular expression to parse the Class attribute contents, specify characters to insert between the resulting regular expression groups, and specify the subscription ID type with the `use-class` options under the `[edit unified-edge gateways tdf gateway-name domains domain-name subscription-id]` hierarchy.

use-imsi Subscription-Id-Type is END_USER_IMSI and the Subscription-Id-Data is the 3GPP-IMSI.

use-msisdn Subscription-Id-Type is END_USER_E164 and the Subscription-Id-Data is the Calling-Station-Id.

use-nai Subscription-Id-Type is END_USER_NAI and the Subscription-Id-Data is the entire User-Name.

use-nas-port Subscription-Id-Type is END_USER_PRIVATE and the Subscription-Id-Data is the NAS-Port.

use-nas-port-id Subscription-Id-Type is END_USER_PRIVATE and the Subscription-Id-Data is the NAS-Port-Id.

use-realm Subscription-Id-Type is END_USER_PRIVATE and the Subscription-Id-Data is the realm portion of User-Name in NAI format.

use-username Subscription-Id-Type is END_USER_PRIVATE and the Subscription-Id-Data is the user name portion of User-Name in NAI format.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

id (MACsec for MX Series)

IN THIS SECTION

- Syntax | 8283
- Hierarchy Level | 8284
- Description | 8284
- Options | 8284
- Required Privilege Level | 8284
- Release Information | 8284

Syntax

```
id {
    mac-address mac-address;
    port-id port-id-number;
}
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel
secure-channel-name]
```

Description

Specify a MAC address and a port that traffic on the link must be from to be accepted by the interface when MACsec is enabled using static secure association key (SAK) security mode.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

id (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 8285
- Hierarchy Level | 8285
- Description | 8285

- Options | 8285
- Required Privilege Level | 8285
- Release Information | 8285

Syntax

```
id id;
```

Hierarchy Level

```
[edit protocols openflow switch switch-name controller]
```

Description

Specify an optional numeric identifier for the OpenFlow controller.

Options

id Numeric identifier for the OpenFlow controller.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Support for OpenFlow on Devices Running Junos OS

controller (Protocols OpenFlow) | 4654

identification (Security ICMP Headers)

IN THIS SECTION

- Syntax | 8286
- Hierarchy Level | 8286
- Description | 8286
- Options | 8287
- Required Privilege Level | 8287
- Release Information | 8287

Syntax

```
identification {
    match (equal | greater-than | less-than | not-equal);
    value identification-value;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol icmp],

[edit security idp custom-attack attack-name attack-type signature protocol icmpv6]
```

Description

Specify a unique value used by the destination system to associate requests and replies.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`identification-value`*—Match a decimal value.
- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3. Statement modified in Junos OS Release 12.3X48-D25 to add ICMPv6 protocol support.

identification (Tunnel Profile)

IN THIS SECTION

- Syntax | 8287
- Hierarchy Level | 8288
- Description | 8288
- Options | 8288
- Required Privilege Level | 8288
- Release Information | 8288

Syntax

```
identification name;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Description

Specify the assignment ID of an L2TP tunnel. L2TP sessions with the same tunnel assignment identification and destination are grouped into the same tunnel.

Options

name—Tunnel assignment ID; string of up to 32 alphanumeric characters.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

identifier (VPLS Multihoming for FEC 129)

IN THIS SECTION

- Syntax | 8289
- Hierarchy Level | 8289
- Description | 8289

- Options | 8290
- Required Privilege Level | 8290
- Release Information | 8290

Syntax

```
identifier identifier;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-
homing site site-name],
[edit routing-instances instance-name protocols vpls multi-homing site site-name]
```

Description

Configure a Layer 2 VPN or VPLS multihoming identifier (MHID). An identifier needs to be configured for each multihomed site. Multihoming site identifiers are specific to a VPLS domain. They need not be unique on a provider edge (PE) router when multiple VPLS instances are present. The network layer reachability information (NLRI) advertisements sent to a CE device are identified as candidates for designated forwarder selection because the advertisements have the same multihoming identifier. Thus, you should assign the same identifier on all VPLS PE routers that are multihomed to the same customer site.

NOTE: The route distinguisher must be unique among PE routers participating in a multihomed site, so that the RD:MHID combination is unique across multiple VPLS domains. For example, one PE router might have a route distinguisher of 192.0.2.4:1, and another PE router in the same site might have a route distinguisher of 192.0.2.:1. The first number can be, for example, the loopback interface address that identifies the PE router. The second number is the multihoming identifier.

NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Options

identifier Number that identifies the multihomed site.

- **Range:** 1 through 65535

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring VPLS Multihoming (FEC 129)

identity

IN THIS SECTION

- Syntax | 8290
- Hierarchy Level | 8291
- Description | 8291
- Required Privilege Level | 8291
- Release Information | 8291

Syntax

```
identity identity-name;
```

Hierarchy Level

```
[edit security ike]
```

Description

Define the identity of the remote certificate name if the identity cannot be learned through IKE (ID payload or IP address).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an IKE Policy for Digital Certificates for an ES PIC

identity-context-profile

IN THIS SECTION

- Syntax | 8292
- Hierarchy Level | 8292
- Description | 8292
- Options | 8293
- Required Privilege Level | 8293
- Release Information | 8293

Syntax

```
identity-context-profile profile-name profile-name {
      domain-name domain-name {
            attribute device-category {
                  string string-value;
            }
            attribute device-identity {
                  string string-value;
            }
            attribute device-os {
                  string string-value;
            }
            attribute device-os-version {
                  string string-value;
            }
            attribute device-type {
                  string string-value;
            }
            attribute device-vendor {
                  string string-value;
            }
      }
}
```

Hierarchy Level

```
[edit services user-identification device-information]
```

Description

Specify the name of the device identity profile, also referred to as the `identity-context-profile`, and either one or more of its attributes or the name of the Active Directory domain to which the device belongs.

The device identity profile is a key component of the SRX Series firewall identity feature, which enables you to control access to network resources based on the identity of the user's device, not the identity of the user of the device. The device identity profile includes the domain name and a collection of attributes that characterize the device.

Options

profile-name ***profile-name***	Name of the device identity profile; for example, marketing-west-coast. The profile is specified in the `destination-identity-context-profile` field of a security policy.
domain ***domain-name***	Name of the domain to which the device belongs; for example, domain1.
attribute device-category ***string***	Category of the device, for example, laptop.
attribute device-identity ***string***	Name given to the device, for example, my-device1.
attribute device-os ***string***	Operating system running on the device, for example, Windows.
attribute device-os-version ***string***	Version of the operating system that is running on the device, for example, 10.1.
attribute device-type ***string***	Type of device, for example, ThinkPad.
attribute device-vendor ***string***	Name of the manufacturer of the device, for example, Lenovo.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

identity-management

IN THIS SECTION

- Syntax | 8294
- Hierarchy Level | 8295

- Description | 8296
- Options | 8296
- Required Privilege Level | 8297
- Release Information | 8297

Syntax

```
identity-management {
    authentication-entry-timeout minutes;
    batch-query {
        items-per-batch items-per-batch;
        query-interval seconds;
    }
    connection {
        connect-method (http | https);
        port port;
        primary {
            address address;
            ca-profile ca-profile;
            client-id client-id;
            client-secret client-secret;
            interface interface-name;
            routing-instance routing-instance -name;
            source source-address;
        }
        query-api query-api;
        secondary {
            address address;
            ca-profile ca-profile;
            client-id client-id;
            client-secret client-secret;
            interface interface-name;
            routing-instance routing-instance -name;
            source source-address;
        }
        server-validity-check-interval;
        token-api token-api;
    }
```

```
    filter {
        domain name;
        exclude-ip {
            address-book book-name;
            address-set address-set;
        }
        include-ip {
            address-book book-name;
            address-set address-set;
        }
    }
    invalid-authentication-entry-timeout minutes;
    ip-query {
        no-ip-query;
        query-delay-time seconds;
    }
    jims-validator {
        port port;
    }
    session-batch-query {
        items-per-batch;
        query-interval seconds;
    }
    session-query {
        no-session-query;
        query-delay-time seconds;
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit services user-identification]
```

Description

Configure identity management to collect identity information. The SRX Series or NFX Series device relies on JIMS to obtain user identity information much in the same way that it does LDAP. By configuring the `identity-management`, the device can query JIMS for identity information, populate identity management authentication table with the information that is obtained from JIMS and use the populated identity management authentication table to authenticate a user or a device requesting access to a protected resource.

Options

authentication-entry-timeout

Authentication entry timeout number (0, 10-1440) (minutes)

- **Default:** 60

invalid-authentication-entry-timeout

Invalid authentication entry timeout number (0, 10-1440) (minutes)

- **Default:** 30

file

Configure the trace file options.

- *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.

- `files` *number*—Maximum number of trace files. When a trace file named *`trace-file`* its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

  If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

  Range: 2 through 1000 files

  Default: 10 files

- `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.

- `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme

continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

Syntax: *x* `K` to specify KB, *x* `m` to specify MB, or *x* `g` to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

flag —Trace operation to perform.

level -Level of debugging output.

- `all` —Match all levels.
- `error` —Match error conditions.
- `info` —Match informational messages.
- `notice` —Match conditions that should be handled specially.
- `verbose` —Match verbose messages.
- `warning` —Match warning messages.

no-remote-trace —Set remote tracing as disabled.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Configure Juniper Identity Management Service to Obtain User Identity Information

identity-provider (SAML)

IN THIS SECTION

- Syntax | 8298
- Hierarchy Level | 8299
- Description | 8299
- Options | 8299
- Required Privilege Level | 8299
- Release Information | 8299

Syntax

```
identity-provider identity-provider {
    attribute-mapping {
        address-pool assertion-attribute <mandatory | optional>;
        firstname assertion-attribute <mandatory | optional>;
        group assertion-attribute <mandatory | optional>;
        lastname assertion-attribute <mandatory | optional>;
        attribute-name assertion-attribute <mandatory | optional>;
        username assertion-attribute <mandatory | optional>;
    }
    metadata-id metadata-id;
    settings {
        entity-id entity-id;
        idp-certificate idp-certificate;
        single-logout-url single-logout-url;
        single-signon-url single-signon-url;
    }
}
```

Hierarchy Level

```
[edit access saml]
```

Description

Describe Security Assertion Markup Language (SAML) identity provider (IdP) configuration details. You can describe more than one identity providers.

Options

identity-provider	Specify an IdP name.
attribute-mapping	Set the attribute mapping details for identity provider (IdP).
metadata-id *metadata-id*	Set the IdP configuration from the loaded metadata. Do not use this option if you have configured `settings` option.
settings	Specify IdP settings. Do not use this option if you have configured `metadata-id` option.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

saml | 16281

idle-after-switch-over

IN THIS SECTION

- Syntax | 8300
- Hierarchy Level | 8300
- Description | 8300
- Options | 8301
- Required Privilege Level | 8301
- Release Information | 8301

Syntax

```
idle-after-switch-over (forever | seconds);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address]
```

Description

Configure the routing device so that it does not automatically reestablish BGP peer sessions after a nonstop active routing (NSR) switchover. This feature is particularly useful if you are using dynamic routing policies because the dynamic database is not synchronized with the backup Routing Engine when NSR is enabled.

Options

forever—Do not reestablish a BGP peer session after an non-stop routing switchover until the `clear bgp neighbor` command is issued.

seconds—Do not reestablish a BGP peer session after an non-stop routing switchover until after the specified period.

- **Range:** 1 through 4,294,967,295 (2^{32} – 1)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Preventing Automatic Reestablishment of BGP Peer Sessions After NSR Switchovers

Routing Policies, Firewall Filters, and Traffic Policers User Guide

Junos OS High Availability User Guide

idle-standby-path-switchover-delay

IN THIS SECTION

- Syntax | 8302
- Hierarchy Level | 8302
- Description | 8302
- Options | 8302
- Required Privilege Level | 8302
- Release Information | 8302

Syntax

```
idle-standby-path-switchover-delay <seconds>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Configure the time interval after which an ECMP join is moved to the standby path in the absence of traffic on the path.

In the absence of this statement, ECMP joins are not moved to the standby path until traffic is detected on the path.

Options

<seconds> Time interval after which an ECMP join is moved to the standby RPF path in the absence of traffic on the path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring PIM Make-Before-Break Join Load Balancing

Configuring PIM Join Load Balancing

clear pim join-distribution | 3809

join-load-balance | 9431

standby-path-creation-delay | 28488

idle-timeout

IN THIS SECTION

- Syntax | 8303
- Hierarchy Level | 8303
- Description | 8303
- Options | 8304
- Required Privilege Level | 8304
- Release Information | 8304

Syntax

```
idle-timeout seconds;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN interfaces, configure the number of seconds the link is idle before losing connectivity.

Options

seconds—Time for which the connection can remain idle. For interfaces configured to use a filter for traffic, the idle timeout is based on traffic.

- **Range:** 1 through 429497295
- **Default:** 120 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

idle-timeout (Unified Edge)

IN THIS SECTION

- Syntax | 8304
- Hierarchy Level | 8305
- Description | 8305
- Options | 8305
- Required Privilege Level | 8305
- Release Information | 8305

Syntax

```
idle-timeout idle-timeout;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
```

Description

Configure the idle timeout for the TDF IP-based subscriber session. The idle timeout is the duration that the subscriber session waits to receive a data packet before timing out. After the idle timeout expires, the TDF takes down the session. Setting the idle timeout ensures that if no data is being sent for the duration specified, then the session can be taken down, and the TDF's resources can be freed.

Options

idle-timeout Number of minutes after which the TDF subscriber session times out.

- **Range:** 0 through 300.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

idle-timeout (Access)

IN THIS SECTION

- Syntax | 8306
- Hierarchy Level | 8306
- Description | 8306
- Options | 8307
- Required Privilege Level | 8307
- Release Information | 8307

Syntax

```
idle-timeout seconds;
```

Hierarchy Level

```
[edit access group-profile profile-nameppp ppp],
[edit access profile profile-name client client-nameppp ]
```

Description

Configure the idle timeout for a user. The router might consider a PPP session to be idle because of the following reasons:

- There is no ingress traffic on the PPP session.
- There is no egress traffic.
- There is neither ingress or egress traffic on the PPP session.
- There is no ingress or egress PPP control traffic. This is applicable only if keepalives are enabled.

Options

seconds—Number of seconds a user can remain idle before the session is terminated.

- **Range:** 0 through 4,294,967,295 seconds
- **Default:** 0

> **NOTE**: The `[edit access]` hierarchy is not available on QFabric systems.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

L2TP LNS Inline Service Interfaces

idle-timeout (L2TP)

IN THIS SECTION

- Syntax | 8308
- Hierarchy Level | 8308
- Description | 8308
- Options | 8308
- Required Privilege Level | 8308
- Release Information | 8308

Syntax

```
idle-timeout seconds;
```

Hierarchy Level

```
[edit services l2tp tunnel]
```

Description

Specify how long a tunnel is active after its last session is terminated. The timer starts when the session is terminated and the tunnel is disconnected when the timer expires.

BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing `no services l2tp tunnel idle-timeout`.

Options

seconds—Length of the idle timeout. A value of `0` creates a persistent tunnel; that is, the tunnel remains active indefinitely until the remote peer disconnects it or you issue the `clear services l2tp tunnel` command.

- **Range:** 0 through 86,400
- **Default:** 60

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

L2TP LAC Subscriber Configuration

L2TP LNS Inline Service Interfaces

idle-timeout (System)

IN THIS SECTION

- Syntax | 8309
- Hierarchy Level | 8309
- Description | 8310
- Options | 8310
- Required Privilege Level | 8310
- Release Information | 8310

Syntax

```
idle-timeout idle-timeout;
```

Hierarchy Level

For root users:

```
[edit system login]
```

For non-root users:

```
[edit system login class]
```

Description

Configure the maximum time for which the C shell or CLI session can be idle. The root user is logged out of the CLI session after the expiry of `idle-timeout` and is logged out of the C shell after the expiry of another `idle-timeout`. For instance, if you configure `idle-timeout` as 5 minutes, the root user is logged out of the system after 10 minutes.

> **NOTE**: This config knob is applicable only for root user. To enable `idle-timeout` for non-root user, refer to Configure the Timeout Value for Idle Login Sessions.

The `set system login idle-timeout` is a global setting that applies to the default class and the custom class. If an ideal-timeout is configured for any custom class, it would override the global setting of its user.

Options

idle-timeout

Maximum idle time for root users before logout.

- **Range:** 1 through 60 minutes

> **NOTE**: Although you can set the idle-timeout to 1 minute, it is recommended to set it to 2 minutes or more for proper operation.

Maximum idle time out for non-root users before logout.

- **Range:** Any configured minutes

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

idle-timeout (System Services)

IN THIS SECTION

- Syntax | 8311
- Hierarchy Level | 8311
- Description | 8311
- Options | 8311
- Required Privilege Level | 8312
- Release Information | 8312

Syntax

```
idle-timeout;
```

Hierarchy Level

```
[edit system services web-management session]
```

Description

Configure in minutes the idle-timeout parameter for web-management sessions. The idle-timeout parameter, which applies to all sessions, specifies the length of time a session can be idle before it is terminated. The web-management feature allows you to configure the device using the J-Web interface.

You can also configure the maximum allowed number of concurrent web management sessions using the session-limit parameter.

Options

idle-timeout *minutes*

- **Default:** 30

- **Range:** 1 to 1440

Required Privilege Level

system—To view this in the configuration.

system-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Dynamic VPN Overview

idp (Application Services)

IN THIS SECTION

- Syntax | 8312
- Hierarchy Level | 8313
- Description | 8313
- Required Privilege Level | 8313
- Release Information | 8313

Syntax

```
idp;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
application-services]
```

Description

Configure Intrusion Detection and Prevention (IDP) for application services.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

idp (logical-systems)

IN THIS SECTION

- Syntax | 8314
- Hierarchy Level | 8314
- Description | 8314
- Options | 8314
- Required Privilege Level | 8314
- Release Information | 8314

Syntax

```
idp (default | off | on);
```

Hierarchy Level

```
[edit logical-systems name security]
```

Description

Configure IDP on primary and user logical systems.

Options

default-policy Set the active policy.

max-sessions Maximum number of IDP sessions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 11.4

RELATED DOCUMENTATION

- show security idp logical system | **24600**
- *IDP for Logical Systems*
- *IDP for Logical Systems*

idp-policy

IN THIS SECTION

- Syntax | 8315
- Hierarchy Level | 8315
- Description | 8315
- Options | 8315
- Required Privilege Level | 8315
- Release Information | 8316

Syntax

```
idp-policy idp-policy;
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the IDP policy for the security profile.

Options

idp-policy-name– Name of the IDP policy.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Intrusion Detection and Prevention Overview

idp-policy (Application Services)

IN THIS SECTION

- Syntax | 8316
- Hierarchy Level | 8316
- Description | 8316
- Required Privilege Level | 8317
- Release Information | 8317

Syntax

```
idp-policy idp-policy;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
application-services]
```

Description

Specify IDP policy name.

When you configure a unified policy with a dynamic application as one of the matching condition, the configuration eliminates the additional steps involved in IDP policy configuration. IDP policy configurations are simplified within a unified policy. Configuring source or destination address, source and destination-except, from and to zone, or application is not required with unified policy, as the match happens in the security policy itself.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1

Starting in Junos OS Release 18.2R1, IDP policy is available within unified security policy.

Unified policies are supported on SRX Series Firewalls, allowing granular control and enforcement of Dynamic Layer Applications within the traditional Security Policy. Layer 7 dynamic applications are integrated with security policy match criteria and IDP policy supports Layer 7 application based security policies.

idp-policy (Security)

IN THIS SECTION

- Syntax | **8318**
- Hierarchy Level | **8319**
- Description | **8320**
- Options | **8320**
- Required Privilege Level | **8320**
- Release Information | **8320**

Syntax

```
idp-policy policy-name {
    rulebase-exempt {
        rule rule-name {
            description text;
            match {
                attacks {
                    custom-attack-groups [attack-group-name];
                    custom-attacks [attack-name];
                    dynamic-attack-groups [attack-group-name];
                    predefined-attack-groups [attack-group-name];
                    predefined-attacks [attack-name];
                }
                destination-address ([address-name] | any | any-ipv4 | any-ipv6);
                destination-except [address-name];
                from-zone (zone-name | any );
                source-address ([address-name] | any | any-ipv4 | any-ipv6);
                source-except [address-name];
                to-zone (zone-name | any);
            }
        }
    }
        rulebase-ips {
            rule rule-name {
            description text;
            match {
                application (application-name | any | default);
                attacks {
                    custom-attack-groups [attack-group-name];
                    custom-attacks [attack-name];
                    dynamic-attack-groups [attack-group-name];
                    predefined-attack-groups [attack-group-name];
                    predefined-attacks [attack-name];
                }
                destination-address ([address-name] | any | any-ipv4 | any-ipv6);
                destination-except [address-name];
                from-zone (zone-name | any );
                source-address ([address-name] | any | any-ipv4 | any-ipv6);
                source-except [address-name];
                to-zone (zone-name | any);
            }
```

```
            terminal;
            then {
                action {
                    class-of-service {
                        dscp-code-point number;
                         forwarding-class forwarding-class;
                    }
                    (close-client | close-client-and-server | close-server |drop-connection |
drop-packet | ignore-connection | mark-diffserv value | no-action | recommended);
                }
                ip-action {
                    (ip-block | ip-close | ip-notify);
                    log;
                    log-create;
                    refresh-timeout;
                    target (destination-address | service | source-address | source-zone |
source-zone-address | zone-service);
                    timeout seconds;
                }
                notification {
                    log-attacks {
                        alert;
                    }
                    packet-log {
                        post-attack number;
                        post-attack-timeout seconds;
                        pre-attack number;
                    }
                }
                severity (critical | info | major | minor | warning);
            }
        }
    }
}
```

Hierarchy Level

```
[edit security idp]
[edit tenants tenant-name security idp]
```

Description

Configure a security IDP policy.

Options

policy-name—Name of the IDP policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Starting with Junos OS Release 18.2R1, IDP policy is directly assigned in the security policy rule. This is to simplify IDP policy usage and to provide flexibility to have multiple policies active at the same time. As a part of the session interest check, IDP is enabled if an IDP policy is present in any of the matched rules. An IDP policy is activated in security policies by permitting the IDP policy within the application services using the `set security policies from-zone zone-name to-zone zone-name policy policy-name then permit application-services idp-policy idp-policy-name` command. Because the IDP policy name is directly used in the security policy rule, the `[edit security idp active-policy policy-name]` statement is deprecated.

Additional tags under filters of dynamic attack groups (CVSS score, age-of-attack, file-type, vulnerability-type) are added in Junos OS Release 18.2R1 for dynamic attacks grouping of IDP signatures. The Product and Vendor tags are already supported under existing filter products. The CLI interface for configuring these tags is now more user friendly, with more options available for configuration in Junos OS Release 18.2R1.

Starting in Junos OS Release 18.3R1, with unified policies configured on an SRX Series Firewall, you can configure multiple IDP policies and set one of those policies as the default IDP policy.

> **NOTE**: If you have configured two or more IDP policies in a unified security policy, then you must configure the default IDP policy.

idp (Security)

IN THIS SECTION

- Syntax | 8321
- Hierarchy Level | 8332
- Description | 8332
- Options | 8332
- Required Privilege Level | 8332
- Release Information | 8332

Syntax

```
idp {
    active-policy policy-name;
    custom-attack attack-name {
        attack-type (Security IDP) {
            anomaly {
                direction (any | client-to-server | server-to-client);
                service service-name;
                shellcode (all | intel | no-shellcode | sparc);
                test test-condition;
            }
            chain {
                expression boolean-expression;
                member member-name {
                    attack-type (Security IDP) {
                        (anomaly ...same statements as in [edit security idp custom-attack
attack-name attack-type anomaly] hierarchy level | signature ...same statements as in [edit
security idp custom-attack attack-name attack-type signature] hierarchy level);
                    }
                }
                order;
                protocol-binding {
                    application application-name;
                    icmp;
                    icmpv6;
```

```
                ip (Security IDP Custom Attack) {
                    protocol-number transport-layer-protocol-number;
                }
                ipv6 {
                    protocol-number transport-layer-protocol-number;
                }
                rpc {
                    program-number rpc-program-number;
                }
                tcp (Security IDP Signature Attack) {
                    minimum-port port-number <maximum-port port-number>;
                }
                udp (Security IDP Signature Attack) {
                    minimum-port port-number <maximum-port port-number>;
                }
            }
            reset;
            scope (session | transaction);
        }
        signature (Security IDP) {
            context context-name;
            direction (any | client-to-server | server-to-client);
            negate;
            pattern signature-pattern;
            protocol (Security IDP Signature Attack) {
                icmp {
                    code {
                        match (equal | greater-than | less-than | not-equal);
                        value code-value;
                    }
                    data-length {
                        match (equal | greater-than | less-than | not-equal);
                        value data-length;
                    }
                    identification {
                        match (equal | greater-than | less-than | not-equal);
                        value identification-value;
                    }
                    sequence-number {
                        match (equal | greater-than | less-than | not-equal);
                        value sequence-number;
                    }
                    type {
```

```
                match (equal | greater-than | less-than | not-equal);
                value type-value;
            }
        }
        ipv4 {
            destination {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            identification {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            ip-flags {
                (df | no-df);
                (mf | no-mf);
                (rb | no-rb);
            }
            protocol {
                match (equal | greater-than | less-than | not-equal);
                value transport-layer-protocol-id;
            }
            source {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            tos {
                match (equal | greater-than | less-than | not-equal);
                value type-of-service-in-decimal;
            }
            total-length {
                match (equal | greater-than | less-than | not-equal);
                value total-length-of-ip-datagram;
            }
            ttl {
                match (equal | greater-than | less-than | not-equal);
                value time-to-live;
            }
        }
        ipv6 {
            destination {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
```

```
                }
                flow-label {
                    match (equal | greater-than | less-than | not-equal);
                    value flow-label-value;
                }
                hop-limit {
                    match (equal | greater-than | less-than | not-equal);
                    value hop-limit-value;
                }
                next-header {
                    match (equal | greater-than | less-than | not-equal);
                    value next-header-value;
                }
                payload-length {
                    match (equal | greater-than | less-than | not-equal);
                    value payload-length-value;
                }
                source {
                    match (equal | greater-than | less-than | not-equal);
                    value ip-address-or-hostname;
                }
                traffic-class {
                    match (equal | greater-than | less-than | not-equal);
                    value traffic-class-value;
                }
            tcp {
                ack-number {
                    match (equal | greater-than | less-than | not-equal);
                    value acknowledgement-number;
                }
                data-length {
                    match (equal | greater-than | less-than | not-equal);
                    value tcp-data-length;
                }
                destination-port {
                    match (equal | greater-than | less-than | not-equal);
                    value destination-port;
                }
                header-length {
                    match (equal | greater-than | less-than | not-equal);
                    value header-length;
                }
                mss {
```

```
                match (equal | greater-than | less-than | not-equal);
                value maximum-segment-size;
            }
            option {
                match (equal | greater-than | less-than | not-equal);
                value tcp-option;
            }
            sequence-number {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
            source-port {
                match (equal | greater-than | less-than | not-equal);
                value source-port;
            }
            tcp-flags {
                (ack | no-ack);
                (fin | no-fin);
                (psh | no-psh);
                (r1 | no-r1);
                (r2 | no-r2);
                (rst | no-rst);
                (syn | no-syn);
                (urg | no-urg);
            }
            urgent-pointer {
                match (equal | greater-than | less-than | not-equal);
                value urgent-pointer;
            }
            window-scale {
                match (equal | greater-than | less-than | not-equal);
                value window-scale-factor;
            }
            window-size {
                match (equal | greater-than | less-than | not-equal);
                value window-size;
            }
        }
        udp {
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
            }
```

```
                        destination-port {
                            match (equal | greater-than | less-than | not-equal);
                            value destination-port;
                        }
                        source-port {
                            match (equal | greater-than | less-than | not-equal);
                            value source-port;
                        }
                    }
                }
                protocol-binding {
                    application application-name;
                    icmp;
                    icmpv6;
                    ip {
                        protocol-number transport-layer-protocol-number;
                    }
                    ipv6 {
                        protocol-number transport-layer-protocol-number;
                    }
                    rpc {
                        program-number rpc-program-number;
                    }
                    tcp {
                        minimum-port port-number <maximum-port port-number>;
                    }
                    udp {
                        minimum-port port-number <maximum-port port-number>;
                    }
                }
                regexp regular-expression;
                shellcode (all | intel | no-shellcode | sparc);
            }
        }
        recommended-action (close | close-client | close-server | drop | drop-packet | ignore |
none);
        severity (critical | info | major | minor | warning);
        time-binding {
            count count-value;
            scope (destination | peer | source);
        }
    }
    custom-attack-group custom-attack-group-name {
```

```
        group-members [attack-or-attack-group-name];
    }
    default-policy  default-policy;
    dynamic-attack-group dynamic-attack-group-name {
        filters {
            category {
                values [category-value];
            }
            direction {
                expression (and | or);
                values [any client-to-server exclude-any exclude-client-to-server exclude-server-
to-client server-to-client];
            }
            false-positives {
                values [frequently occasionally rarely unknown];
            }
            performance {
                values [fast normal slow unknown];
            }
            products {
                values [product-value];
            }
            recommended;
            service {
                values [service-value];
            }
            severity {
                values [critical info major minor warning];
            }
            type {
                values [anomaly signature];
            }
        }
    }
    idp-policy (Security)  policy-name {
        rulebase-exempt {
            rule rule-name {
                description text;
                match {
                    attacks {
                        custom-attack-groups [attack-group-name];
                        custom-attacks [attack-name];
                        dynamic-attack-groups [attack-group-name];
```

```
                        predefined-attack-groups [attack-group-name];
                        predefined-attacks [attack-name];
                    }
                    destination-address ([address-name] | any | any-ipv4 | any-ipv6);
                    destination-except [address-name];
                    from-zone (zone-name | any );
                    source-address ([address-name] | any | any-ipv4 | any-ipv6);
                    source-except [address-name];
                    to-zone (zone-name | any);
                }
            }
        }
        rulebase-ips {
            rule rule-name {
                description text;
                match {
                    application (application-name | any | default);
                    attacks {
                        custom-attack-groups [attack-group-name];
                        custom-attacks [attack-name];
                        dynamic-attack-groups [attack-group-name];
                        predefined-attack-groups [attack-group-name];
                        predefined-attacks [attack-name];
                    }
                    destination-address ([address-name] | any | any-ipv4 | any-ipv6);
                    destination-except [address-name];
                    from-zone (zone-name | any );
                    source-address ([address-name] | any | any-ipv4 | any-ipv6);
                    source-except [address-name];
                    to-zone (zone-name | any);
                }
                terminal;
                then {
                    action {
                        class-of-service {
                            dscp-code-point number;
                             forwarding-class forwarding-class;
                        }
                        (close-client | close-client-and-server | close-server |drop-connection
| drop-packet | ignore-connection | mark-diffserv value | no-action | recommended);
                    }
                    ip-action {
                        (ip-block | ip-close | ip-notify);
```

```
                        log;
                        log-create;
                        refresh-timeout;
                        target (destination-address | service | source-address | source-zone |
source-zone-address | zone-service);
                        timeout seconds;
                    }
                    notification {
                        log-attacks {
                            alert;
                        }
                        packet-log {
                            post-attack number;
                            post-attack-timeout seconds;
                            pre-attack number;
                        }
                    }
                    severity (critical | info | major | minor | warning);
                }
            }
        }
    }
    security-package {
        automatic {
            download-timeout minutes;
            enable;
            interval hours;
            start-time start-time;
        }
        install {
            ignore-version-check;
            ignore-appid-failure;
        }
        proxy-profile proxy-profile;
        source-address address;
        url url-name;
    }
    sensor-configuration {
        application-identification {
            max-packet-memory value;
            max-tcp-session-packet-memory value;
            max-udp-session-packet-memory value;
        }
```

```
        detector {
            protocol-name protocol-name {
                tunable-name  tunable-name {
                    tunable-value protocol-value;
                }
            }
        }
        flow (Security IDP) {
            (allow-icmp-without-flow | no-allow-icmp-without-flow);
            fifo-max-size value;
            hash-table-size value;
            (log-errors | no-log-errors);
            max-session-offset value;
            max-timers-poll-ticks value;
            reject-timeout value;
            (reset-on-policy | no-reset-on-policy);
            udp-anticipated-timeout value;
        }
        global (Security IDP) {
            (enable-all-qmodules | no-enable-all-qmodules);
            (enable-packet-pool | no-enable-packet-pool);
            gtp (decapsulation | no-decapsulation);
            memory-limit-percent value;
            (policy-lookup-cache | no-policy-lookup-cache);
        }
        high-availability (Security IDP) {
            no-policy-cold-synchronization;
        }
        ips {
            content-decompression-max-memory-kb value;
            content-decompression-max-ratio value;
            (detect-shellcode | no-detect-shellcode);
            fifo-max-size value;
            (ignore-regular-expression | no-ignore-regular-expression);
            log-supercede-min minimum-value;
            pre-filter-shellcode;
            (process-ignore-s2c | no-process-ignore-s2c);
            (process-override | no-process-override);
            process-port port-number;
        }
        log (Security IDP Sensor Configuration) {
            cache-size size;
            suppression {
```

```
                disable;
                (include-destination-address | no-include-destination-address);
                max-logs-operate value;
                max-time-report value;
                start-log value;
            }
        }
        packet-log (Security IDP Sensor Configuration) {
            host ip-address <port number>;
            max-sessions percentage;
            source-address ip-address;
            total-memory percentage;
        }
        re-assembler {
            action-on-reassembly-failure (drop | drop-session | ignore);
            (force-tcp-window-checks | no-force-tcp-window-checks);
            (ignore-memory-overflow | no-ignore-memory-overflow);
            (ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-overflow);
            ignore-reassembly-overflow;
            max-flow-mem value;
            max-packet-mem value;
            (tcp-error-logging | no-tcp-error-logging);
        }
        ssl-inspection {
            cache-prune-chunk-size number;
            key-protection;
            maximum-cache-size number;
            session-id-cache-timeout seconds;
            sessions number;
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag all;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
```

```
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure Intrusion Detection and Prevention (IDP) to selectively enforce various IDP attack detection and prevention techniques on the network.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.3. The `expression` option added in Junos OS Release 11.4.

Starting in Junos OS Release 18.3R1, when an SRX Series Firewall is configured with unified policies, you can configure multiple IDP policies and set one of those policies as the default IDP policy. If multiple IDP policies are configured for a session and when policy conflict occurs, the device applies the default IDP policy for that session and thus resolves any policy conflicts.

NOTE: If you have configured two or more IDP policies in a unified security policy, then you must configure the default IDP policy.

RELATED DOCUMENTATION

Intrusion Detection and Prevention Overview

idp (Security Alarms)

IN THIS SECTION

- Syntax | 8333
- Hierarchy Level | 8333
- Description | 8333
- Required Privilege Level | 8333
- Release Information | 8333

Syntax

```
idp;
```

Hierarchy Level

```
[edit security alarms potential-violation]
```

Description

Configure alarms for IDP (intrusion, detection, and prevention) attack.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

idp (SNMP)

IN THIS SECTION

- Syntax | 8334
- Hierarchy Level | 8334
- Description | 8334
- Required Privilege Level | 8334
- Release Information | 8335

Syntax

```
idp {
    falling-threshold percentage;
    interval seconds;
    rising-threshold percentage;
}
```

Hierarchy Level

```
[edit snmp health-monitor]
```

Description

Configure health monitoring for Intrusion Detection and Prevention (IDP).

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configure Health Monitoring on Devices Running Junos OS

ids-option

IN THIS SECTION

- Syntax | 8335
- Hierarchy Level | 8338
- Description | 8338
- Options | 8339
- Required Privilege Level | 8339
- Release Information | 8340

Syntax

```
ids-option screen-name {
        alarm-without-drop;
        description text;
        icmp {
            flood {
                threshold number;
            }
            fragment;
            icmpv6-malformed;
            ip-sweep {
                threshold number;
            }
            large;
```

```
            ping-death;
        }
        ip {
            bad-option;
            block-frag {
                white-list name;
            }
            ipv6-extension-header {
                AH-header;
                ESP-header;
                HIP-header;
                destination-header {
                    ILNP-nonce-option;
                    home-address-option;
                    line-identification-option;
                    tunnel-encapsulation-limit-option;
                    user-defined-option-type <type-low> to <type-high>;
                }
                fragment-header;
                hop-by-hop-header {
                    CALIPSO-option;
                    RPL-option;
                    SFM-DPD-option;
                    jumbo-payload-option;
                    quick-start-option;
                    router-alert-option;
                    user-defined-option-type <type-low> to <type-high>;
                }
                mobility-header;
                no-next-header;
                routing-header;
                shim6-header
                user-defined-option-type <type-low> to <type-high>;
            }
            ipv6-extension-header-limit limit;
            ipv6-malformed-header;
            loose-source-route-option;
            record-route-option;
            security-option;
            source-route-option;
            spoofing;
            stream-option;
            strict-source-route-option;
```

```
            tear-drop;
            timestamp-option;
            unknown-protocol;
            tunnel {
                gre {
                    gre-4in4;
                    gre-4in6;
                    gre-6in4;
                    gre-6in6;
                }
                ip-in-udp {
                    teredo;
                }
                ipip {
                    ipip-4in4;
                    ipip-4in6;
                    ipip-6in4;
                    ipip-6in6;
                    ipip-6over4;
                    ipip-6to4relay;
                    isatap;
                    dslite;
                }
                bad-inner-header;
            }
        }
        limit-session {
            destination-ip-based number;
            source-ip-based number;
        }
        tcp {
            fin-no-ack;
            land;
            port-scan {
                threshold number;
            }
            syn-ack-ack-proxy {
                threshold number;
            }
            syn-fin;
            syn-flood {
                alarm-threshold number;
                attack-threshold number;
```

```
                destination-threshold number;
                source-threshold number;
                timeout seconds;
                white-list name {
                    destination-address destination-address;
                    source-address source-address;
                }
            }
            syn-frag;
            tcp-no-flag;
            tcp-sweep {
                threshold threshold number;
            }
            winnuke;
        }
        udp {
            flood {
                threshold number;
            }
            port-scan {
                threshold number;
            }
            udp-sweep {
                threshold threshold number;
            }
        }
    }
}
```

Hierarchy Level

```
[edit security screen],
[edit tenants tenant-name security screen]
```

Description

Define screens for the intrusion detection service (IDS). An `ids-option` can be used for enabling the screen protection on the SRX Series Firewalls. One `ids-option` can be associated with several zones. However each zone can be linked with only one `ids-option`.

Options

`description` *text*—Descriptive text about a screen.

`alarm-without-drop`—Direct the device to generate an alarm when detecting an attack but not block the attack.

`icmp`—Configure the ICMP ids options.

`ip`—Configure the IP layer ids options.

`limit-session`—Limit the number of concurrent sessions the device can initiate from a single source IP address or the number of sessions it can direct to a single destination IP address.

`tcp`—Configure the TCP Layer ids options.

`udp`—Configure the UDP Layer ids options.

`loose-source-route-option`—The device detects packets where the IP option is 3 (Loose Source Routing) and records the event in the screen counters list for the ingress interface. This option specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified.

`source-route-option`—Enable this option to block all IP traffic that employs the loose or strict source route option. Source route options can allow an attacker to enter a network with a false IP address.

`strict-source-route-option`—The device detects packets where the IP option is 9 (Strict Source Routing) and records the event in the screen counters list for the ingress interface. This option specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field. Currently, this screen option is applicable to IPv4 only.

> **NOTE**: Loose source route option and strict source route option will only alarm and will not be dropped when there is overflow of traffic. When only IP source option is configured, the attacked packets are dropped.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Support for the `description` option added in Junos OS Release 12.1.

UDP supports `port-scan` option starting from Junos OS Release 12.1X47-D10.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

ids-rule-sets

IN THIS SECTION

- Syntax | **8340**
- Hierarchy Level | **8341**
- Description | **8341**
- Options | **8341**
- Required Privilege Level | **8341**
- Release Information | **8341**

Syntax

```
(ids-rule-sets rule-set-name);
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the intrusion detection service (IDS) rule set included in this service set. You can configure only one rule set for each service. If the service set is on an MS-MPC, only the first IDS input rule and the first IDS output rule are used.

Options

rule-set-name—Identifier for the set of rules to be included.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Service Rules

Configuring Protection Against Network Attacks on an MS-MPC

ids-rules

IN THIS SECTION

- Syntax | 8342
- Hierarchy Level | 8342

- Description | 8342
- Options | 8342
- Required Privilege Level | 8342
- Release Information | 8342

Syntax

```
ids-rules [rule-name];
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the intrusion detection service (IDS) rules included in this service set. You can configure multiple rules. If the service set is on an MS-MPC, only the first IDS input rule and the first IDS output rule are used.

Options

rule-name—Identifier for the rule to be included.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Service Rules

Configuring Protection Against Network Attacks on an MS-MPC

ie-set

IN THIS SECTION

- Syntax | 8343
- Hierarchy Level | 8343
- Description | 8343
- Options | 8344
- Required Privilege Level | 8344
- Release Information | 8344

Syntax

```
ie-set name {
    ie ie-number;
}
```

Hierarchy Level

```
[edit security gtp]
```

Description

Specify a GTP IE set profile. This helps you to remove IEs of specific types from all messages for GTPv1 and GTPv2.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp]` hierarchy level is replaced by `[edit security gtp]`.

Options

name Name of the GTP IE set profile.

ie-number IE number.

- **Range:** 1 to 255

Required Privilege Level

security—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Example: Configure IE removal for GTPv1 and GTPv2

ieee-802.1 (Class of Service)

IN THIS SECTION

- Syntax (Classifier) | **8345**
- Syntax (Code-Point Alias) | **8345**
- Syntax (Multidestination Classifier) | **8345**
- Syntax (Interface Classifier Association) | **8346**
- Syntax (Rewrite Rule) | **8346**
- Hierarchy Level (Classifiers) | **8346**
- Hierarchy Level (Code-Point Aliases) | **8346**
- Hierarchy Level (Multidestination Classifier) | **8346**
- Hierarchy Level (Interface Classifier Association) | **8347**
- Hierarchy Level (Rewrite Rules) | **8347**

- Description | 8347
- Options | 8347
- Required Privilege Level | 8347
- Release Information | 8347

Syntax (Classifier)

Classifier Configuration

```
ieee-802.1 classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [ aliases ] [ bit-patterns ];
        }
    }
}
```

Syntax (Code-Point Alias)

Code-Point Alias Configuration

```
ieee-802.1 alias-name bit-pattern;
```

Syntax (Multidestination Classifier)

Multidestination Classifier Configuration

```
ieee-802.1 classifier-name;
```

Syntax (Interface Classifier Association)

Interface Classifier Association

```
ieee-802.1 (classifier-name | default);
```

Syntax (Rewrite Rule)

Rewrite Rule Configuration

```
ieee-802.1 rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-point [ aliases ] [ bit-patterns ];
        }
    }
}
```

Hierarchy Level (Classifiers)

```
[edit class-of-service classifiers],
```

Hierarchy Level (Code-Point Aliases)

```
[edit class-of-service code-point-aliases],
```

Hierarchy Level (Multidestination Classifier)

```
[edit class-of-service multi-destination classifiers ],
```

Hierarchy Level (Interface Classifier Association)

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers],
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules],
```

Hierarchy Level (Rewrite Rules)

```
[edit class-of-service rewrite-rules]
```

Description

Configure an IEEE 802.1 classifier, configure an IEEE 802.1 code-point alias, apply a fixed IEEE 802.1 classifier to an interface, or apply an IEEE-802.1 rewrite rule.

Options

classifier-name—Name of the classifier.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Unicast Classifiers

Defining CoS Code-Point Aliases

Defining CoS Rewrite Rules

Assigning CoS Components to Interfaces

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces
Understanding CoS Classifiers
Understanding CoS Classifiers
Understanding CoS Rewrite Rules
Example: Configuring CoS on EX Series Switches
Defining CoS Classifiers (CLI Procedure)
Defining CoS Classifiers (J-Web Procedure)
Defining CoS Code-Point Aliases (CLI Procedure)
Defining CoS Code-Point Aliases (J-Web Procedure)
Defining CoS Rewrite Rules (CLI Procedure)
Defining CoS Rewrite Rules (J-Web Procedure)
Understanding CoS Classifiers
Understanding CoS Rewrite Rules

ieee-802.1 (Classifier on Physical Interface)

IN THIS SECTION

- Syntax | 8348
- Hierarchy Level | 8349
- Description | 8349
- Options | 8349
- Required Privilege Level | 8349
- Release Information | 8349

Syntax

```
ieee-802.1 (classifier-name | default) vlan-tag (inner | outer );
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name classifiers]
```

Description

For ACX Series Universal Metro routers and EX Series switches, map the ieee-802.1p field of the incoming packet to the forwarding class and packet loss priority based on the specified 802.1p classifier. In the case of double tagged packets, you can configure whether to use the 802.1p of the outer or inner VLAN tag.

Options

vlan-tag inner In the case of double tagged packets, classify based on the 802.1p of the inner VLAN tag.

vlan-tag outer Classify based on the 802.1p of the outermost VLAN tag.

classifier-name Name of the previously defined ieee-802.1p behavior aggregate classifier.

default Default ieee-802.1p behavior aggregate classifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

ieee-802.1 (Congestion Notification)

IN THIS SECTION

- Syntax | 8350
- Hierarchy Level | 8350
- Description | 8350
- Required Privilege Level | 8350
- Release Information | 8351

Syntax

```
ieee-802.1 {
    code-point up-bits pfc ;
}
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name],
[edit class-of-service interfaces interface-name congestion-notification-profile profile-name]
```

Description

Set an association between the traffic class and the congestion notification profile.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring an FCoE Transit Switch

Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)

ieee-802.1 (Dynamic Classifiers)

IN THIS SECTION

- Syntax | 8351
- Hierarchy Level | 8351
- Description | 8352
- Options | 8352
- Required Privilege Level | 8352
- Release Information | 8352

Syntax

```
ieee-802.1 (classifier-name | default) vlan-tag (inner | outer);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-
number classifiers]
```

Description

Apply an IEEE-802.1 classifier to a subscriber interface in a dynamic profile.

Options

classifier-name—Name of a `classifier` mapping configured at the `[edit class-of-service classifiers ieee-802.1]` hierarchy level.

`default`—The default mapping.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying a Classifier to a Subscriber Interface in a Dynamic Profile

classifiers (Definition)

ieee-802.1 (Dynamic Rewrite Rules)

IN THIS SECTION

- Syntax | 8353
- Hierarchy Level | 8353

- Description | 8353
- Options | 8353
- Required Privilege Level | 8353
- Release Information | 8354

Syntax

```
ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-
number rewrite-rules]
```

Description

Apply an IEEE-802.1 rewrite rule to a subscriber interface in a dynamic profile.

Options

rewrite-name—Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules ieee-802.1]` hierarchy level.

`default`—The default mapping.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile

rewrite-rules

ieee-802.1 (Fibre Channel Interfaces)

IN THIS SECTION

- Syntax | 8354
- Hierarchy Level | 8354
- Description | 8355
- Required Privilege Level | 8355

Syntax

```
ieee-802.1 {
    code-point code-point-bits;
}
```

Hierarchy Level

```
[edit class-of-service interfaces fibre-channel-interface-name rewrite-value input]
```

Description

Configure the IEEE 802.1p code point value to which all traffic received from the Fibre Channel (FC) network on the specified FC interface is rewritten. After the code point value is rewritten, the interface forwards the traffic to the Ethernet (FCoE) network. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.

FCoE networks typically use priority 3 (IEEE code point 011) for FCoE traffic. The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011).

The system supports only one IEEE 802.1p code point value per FC interface. You cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.

The statement is described separately.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

forwarding-class (Fibre Channel Interfaces) | **7421**

Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

Understanding CoS Classifiers

ieee-802.1 (Host Outbound Traffic)

IN THIS SECTION

- Syntax | 8356
- Hierarchy Level | 8356
- Description | 8356
- Required Privilege Level | 8356
- Release Information | 8357

Syntax

```
ieee-802.1 {
    default value;
    rewrite-rules;
}
```

Hierarchy Level

```
[edit class-of-service host-outbound-traffic]
```

Description

Apply the IEEE 802.1p rewrite rules associated with the egress logical interface to the IEEE 802.1p PCP field for all host outbound traffic on that interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic

Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface

Rewriting Packet Headers to Ensure Forwarding Behavior

Configuring Rewrite Rules

ieee-802.1 (Input Congestion Notification)

IN THIS SECTION

- Syntax | 8357
- Hierarchy Level | 8358
- Description | 8358
- Options | 8358
- Required Privilege Level | 8358
- Release Information | 8358

Syntax

```
ieee-802.1 {
    code-point [code-point-bits] {
        pfc {
            mru mru-value;
        }
        xon
    }
}
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name input]
```

Description

Configure an IEEE 802.1 code point and apply priority-based flow control (PFC) to packets with that code point.

Options

The statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring CoS PFC for FCoE Traffic

Configuring CoS PFC (Congestion Notification Profiles)

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

ieee-802.1 (Output Congestion Notification)

IN THIS SECTION

Syntax | 8359

- Hierarchy Level | 8359
- Description | 8359
- Options | 8359
- Required Privilege Level | 8359
- Release Information | 8360

Syntax

```
ieee-802.1 {
    code-point [ code-point-bits ] {
        flow-control-queue [ queue | list-of-queues ];
    }
}
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name output]
```

Description

Configure an IEEE 802.1 code point and apply priority-based flow control (PFC) to packets with that code point on output queues.

Options

The statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring CoS PFC (Congestion Notification Profiles)

Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces

Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface

Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)

Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)

Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows

ieee-802.1 (Rewrite Rules on Logical Interface)

IN THIS SECTION

- Syntax | **8360**
- Hierarchy Level | **8361**
- Description | **8361**
- Options | **8361**
- Required Privilege Level | **8361**
- Release Information | **8361**

Syntax

```
ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

Apply an IEEE-802.1 rewrite rule. For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.

Options

rewrite-name—Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules ieee-802.1]` hierarchy level.

`default`—The default mapping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`vlan-tag` statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring Rewrite Rules

dscp (Rewrite Rules on Interfaces) | **6020**

dscp-ipv6 (Class of Service) | **6009**

exp (Rewrite rules) | **6643**

exp-push-push-push | **6642**

exp-swap-push-push | **6646**

ieee-802.1ad | **8363**

inet-precedence (CoS Rewrite Rules) | **8582**

rewrite-rules (Definition) | 15746

ieee-802.1 (Rewrite Rules on Physical Interface)

IN THIS SECTION

- Syntax | 8362
- Hierarchy Level | 8362
- Description | 8362
- Options | 8362
- Required Privilege Level | 8363
- Release Information | 8363

Syntax

```
ieee-802.1 (rewrite-name | default) ;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name rewrite-rules]
```

Description

Apply an IEEE-802.1 rewrite rule.

Options

rewrite-name—Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules ieee-802.1]` hierarchy level.

`default`—The default mapping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

ieee-802.1ad

IN THIS SECTION

- Syntax | 8363
- Hierarchy Level | 8363
- Description | 8363
- Options | 8364
- Required Privilege Level | 8364
- Release Information | 8364

Syntax

```
ieee-802.1ad (rewrite-name | default) vlan-tag (outer | outer-and-inner);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

Apply a IEEE-802.1ad rewrite rule.

Options

rewrite-name—Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules ieee-802.1ad]` hierarchy level.

`default`—The default rewrite bit mapping.

`vlan-tag`—The rewrite rule is applied to the `outer` or `outer-and-inner` VLAN tag.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Rewrite Rules

dscp (Rewrite Rules on Interfaces) | **6020**

dscp-ipv6 (Class of Service) | **6009**

exp (Rewrite rules) | **6643**

exp-push-push-push | **6642**

exp-swap-push-push | **6646**

ieee-802.1 (Rewrite Rules on Logical Interface) | **8360**

inet-precedence (CoS Rewrite Rules) | **8582**

rewrite-rules (Definition) | **15746**

ieee-802-3az-eee

IN THIS SECTION

Syntax | **8365**

- Hierarchy Level | 8365
- Description | 8365
- Default | 8365
- Required Privilege Level | 8365
- Release Information | 8365

Syntax

```
ieee-802-3az-eee;
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
```

Description

Configure Energy Efficient Ethernet (EEE) on an EEE-capable Base-T copper interface.

Default

EEE is disabled on EEE-capable interfaces.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Energy Efficient Ethernet on Interfaces

ieee802.1p

IN THIS SECTION

- Syntax | 8366
- Hierarchy Level | 8366
- Description | 8366
- Options | 8367
- Required Privilege Level | 8367
- Release Information | 8367

Syntax

```
ieee802.1p premium [ values ];
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile
input-priority-map]
[edit interfaces interface-name ether-options ethernet-switch-profile ethernet-policer-profile input-
priority-map]
```

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, configure premium priority values for IEEE 802.1p input traffic.

Options

values—Define IEEE 802.1p priority values to be treated as premium.

- **Range:** 0 through 7

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos Release 7.4.

ietf-mode

IN THIS SECTION

- Syntax | 8367
- Hierarchy Level | 8368
- Description | 8368
- Default | 8368
- Required Privilege Level | 8368
- Release Information | 8368

Syntax

```
ietf-mode
```

Hierarchy Level

```
[edit protocols ancp neighbor ip-address]
```

Description

Configure the ANCP agent to run in a mode that is not backward compatible with Internet draft-wadhwa-gsmp-l2control-configuration-00.txt, *GSMP extensions for layer2 control (L2C)*. Include this statement when pre-ietf mode has been configured globally for the ANCP agent, but you want one or more neighbors to run in the default mode.

Default

ANCP does not run in a backward-compatible mode.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring the ANCP Agent

Configuring ANCP Neighbors

if-exceeding (QFX)

IN THIS SECTION

Syntax | 8369

- Hierarchy Level | 8369
- Description | 8369
- Required Privilege Level | 8369
- Release Information | 8369

Syntax

```
if-exceeding {
    bandwidth-limit bps;
    burst-size-limit bytes;
}
```

Hierarchy Level

```
[edit firewall policer policer-name]
```

Description

Configure policer rate limits.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

> *Configuring Two-Color and Three-Color Policers to Control Traffic Rates*

if-exceeding (EX)

IN THIS SECTION

- Syntax | 8370
- Hierarchy Level | 8370
- Description | 8370
- Required Privilege Level | 8371
- Release Information | 8371

Syntax

```
if-exceeding {
                        bandwidth-limit                    bps;
    bandwidth-percent percent
                        burst-size-limit                    bytes;
}
```

Hierarchy Level

```
[edit firewall                policer                policer-name]
[edit logical-systems logical-system-name firewall policer policer-name]
```

Description

Configure policer rate limits.

The `bandwidth-percent` statement is supported on routers only.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches

Configuring Policers to Control Traffic Rates (CLI Procedure)

Understanding the Use of Policers in Firewall Filters

Basic Single-Rate Two-Color Policers

if-exceeding (Hierarchical Policer)

IN THIS SECTION

- Syntax | 8371
- Hierarchy Level | 8372
- Description | 8372
- Required Privilege Level | 8372
- Release Information | 8372

Syntax

```
if-exceeding {
    bandwidth-limit bps;
```

```
    burst-size-limit bytes;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer aggregate],
[edit dynamic-profiles profile-name firewall hierarchical-policer premium],
[edit firewall hierarchical-policer aggregate],
[edit firewall hierarchical-policer premium]
```

Description

For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, specify bandwidth and burst limits for a premium or aggregate component of a hierarchical policer.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles ... aggregate]` and `[edit dynamic-profiles ... premium]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Hierarchical Policer Configuration Overview
- Hierarchical Policers
- aggregate (Hierarchical Policer) | **2064**
- bandwidth-limit (Hierarchical Policer)

burst-size-limit (Hierarchical Policer) | 3084
hierarchical-policer (aggregate or premium) | 8057
premium (Hierarchical Policer) | 13110

if-exceeding (Policer)

IN THIS SECTION

- Syntax | 8373
- Hierarchy Level | 8373
- Description | 8373
- Required Privilege Level | 8374
- Release Information | 8374

Syntax

```
if-exceeding {
    (bandwidth-limit bps | bandwidth-percent number);
    burst-size-limit bytes;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer  policer-name],
[edit firewall policer policer-name],
[edit logical-systems logical-system-name firewall policer policer-name]
```

Description

Configure rate limits for a single-rate two-color policer.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

Support at the `[edit dynamic-profiles ... policer policer-name]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Two-Color Policer Configuration Overview
- Hierarchical Policer Configuration Overview
- Basic Single-Rate Two-Color Policers
- Bandwidth Policers
- Prefix-Specific Counting and Policing Actions
- Multifield Classification
- Policer Overhead to Account for Rate Shaping in the Traffic Manager
- Hierarchical Policers

if-exceeding-pps (Hierarchical Policer)

IN THIS SECTION

- Syntax | 8375
- Hierarchy Level | 8375
- Description | 8375

Required Privilege Level | 8375
Release Information | 8375

Syntax

```
if-exceeding-pps {
    pps-limit pps;
    packet-burst packets;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer hierarchical-policer-name
aggregate],
[edit dynamic-profiles profile-name firewall hierarchical-policer hierarchical-policer-name premium],
[edit firewall hierarchical-policer hierarchical-policer-name aggregate],
[edit firewall hierarchical-policer hierarchical-policer-name premium]
```

Description

For MX Series routers , `if-exceeding-pps` allows you to configure a packets-per-second (pps)-based trigger for a premium or aggregate component of a hierarchical policer. When applied to the loopback interface (lo0), this kind of trigger can help protect the Routing Engine from DDoS attacks. When applied in other areas, to either transit or control traffic, it is a more fine-grained monitor.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.2.

RELATED DOCUMENTATION

Hierarchical Policer Configuration Overview

Hierarchical Policers

aggregate (Hierarchical Policer) | 2064

bandwidth-limit (Hierarchical Policer) | 2855

burst-size-limit (Hierarchical Policer) | 3084

hierarchical-policer (aggregate or premium) | 8057

premium (Hierarchical Policer) | 13110

if-exceeding-pps (Policer)

IN THIS SECTION

- Syntax | 8376
- Hierarchy Level | 8376
- Description | 8377
- Required Privilege Level | 8377
- Release Information | 8377

Syntax

```
if-exceeding-pps {
    pps-limit pps;
    packet-burst packets;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer  policer-name],
[edit firewall policer policer-name],
[edit logical-systems logical-system-name firewall policer policer-name]
```

Description

Configure rate limits in packets per second for a single-rate two-color policer when policing is desired on a packets-per-second basis rather than a bits-per-second basis. This can help protect against distributed denial of service (DDoS) attacks.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Two-Color Policer Configuration Overview*
- *Hierarchical Policer Configuration Overview*
- *Basic Single-Rate Two-Color Policers*
- *Bandwidth Policers*
- *Prefix-Specific Counting and Policing Actions*
- *Multifield Classification*
- *Policer Overhead to Account for Rate Shaping in the Traffic Manager*
- *Hierarchical Policers*

if-route-exists

IN THIS SECTION

- Syntax | 8378
- Hierarchy Level | 8378

- Description | 8378
- Options | 8379
- Required Privilege Level | 8379
- Release Information | 8379

Syntax

```
if-route-exists {
    address;
    address-family {
        inet {
            address;
            table table-name;
        }
        ccc {
            interface-name;
            standby;
            peer-unit unit-number;
            table table-name;
        }
    }
    table table-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options condition],
[edit policy-options condition]
```

Description

Specify the route match conditions.

Options

(Optional) *address*—Specify the IP address that the route must have to be considered a match.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring Pseudowire Redundancy in a Mobile Backhaul Scenario

Example: Configuring a Routing Policy for Conditional Advertisement Enabling Conditional Installation of Prefixes in a Routing Table

ifl-subscriber

IN THIS SECTION

- Syntax | 8380
- Hierarchy Level | 8380
- Description | 8380
- Options | 8380
- Required Privilege Level | 8380
- Release Information | 8380

Syntax

```
ifl-subscriber [subscriber-name] {
    access-interfaces [interface-name];
    apply-groups [group-names];
    apply-groups-except [group-names];
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
```

Description

Specify the name of the IFL-based subscriber for traffic that is carried on a particular interface or interfaces. You can configure up to 32 IFL-based subscribers in a TDF domain.

To configure a subscriber name, you must have set the `subscriber-type` to **ifl** at the `[edit unified-edge gateway tdf gateway-name domains domain-name]` hierarchy.

Options

subscriber-name Name of the subscriber. You can configure up to 32 IFL-based subscribers in a TDF domain.

- **Range:** Up to 63 bytes.

The remaining statements are described separately.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain

Understanding the Definition of a Set of IFL-Based Subscriber Properties with a TDF Domain

igmp

IN THIS SECTION

- Syntax | 8381
- Hierarchy Level | 8382
- Description | 8382
- Default | 8382
- Required Privilege Level | 8383
- Release Information | 8383

Syntax

```
igmp {
    accounting;
    interface interface-name {
        (accounting | no-accounting);
        disable;
        distributed;
        group-limit limit;
        group-policy [ policy-names ];
        group-threshold
        immediate-leave;
        log-interval
        oif-map map-name;
        passive;
        promiscuous-mode;
        ssm-map ssm-map-name;
        ssm-map-policy ssm-map-policy-name;
        static {
```

```
            group multicast-group-address {
                exclude;
                group-count number;
                group-increment increment;
                source ip-address {
                    source-count number;
                    source-increment increment;
                }
            }
        }
        version version;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit protocols]
```

Description

Enable IGMP on the router or switch. IGMP must be enabled for the router or switch to receive multicast packets.

The remaining statements are explained separately. See CLI Explorer.

Default

IGMP is disabled on the router or switch. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling IGMP

igmp (Dynamic Profiles)

IN THIS SECTION

- Syntax | 8383
- Hierarchy Level | 8384
- Description | 8384
- Default | 8384
- Options | 8384
- Required Privilege Level | 8385
- Release Information | 8385

Syntax

```
igmp {
    interface interface-name {
        accounting;
        disable;
        group-limit policy-name;
        group-policy;
```

```
        group-threshold value;
        immediate-leave;
        log-interval seconds;
        no-accounting;
        oif-map map-name;
        passive <allow-receive> <send-general-query> <send-group-query>;
        promiscuous-mode;
        ssm-map ssm-map-name;
        ssm-map-policy ssm-map-policy-name
        static {
            group group {
                source source;
            }
        }
        version version;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols]
```

Description

Enable IGMP on the router. IGMP must be enabled for the router to receive multicast packets.

You can configure IGMP in a dynamic client profile or a dynamic service profile.

Default

IGMP is disabled on the router. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Options

The statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

- *Dynamic IGMP Configuration Overview*
- *Configuring Dynamic DHCP Client Access to a Multicast Network*
- Understanding IGMP
- Enabling IGMP

igmp-querier (QFabric Systems only)

IN THIS SECTION

- Syntax | 8385
- Hierarchy Level | 8386
- Description | 8386
- Options | 8386
- Required Privilege Level | 8386
- Release Information | 8386

Syntax

```
igmp-querier {
    source-address source address;
}
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan vlan-name]
```

Description

Configure a QFabric Node device to be an IGMP querier. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is lower (a smaller number) than the IP addresses for those routers on the network. This ensures that Node is always the IGMP querier on the network.

Options

source-address *source address*
The address that the switch uses as the source address in the IGMP queries that it sends.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D15.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping on Switches

Configuring IGMP Snooping on Switches

show igmp-snooping vlans | **20198**

igmp-snooping

IN THIS SECTION

- Syntax (EX Series and NFX Series) | 8387
- Syntax (ACX Series, EX9200, and MX Series) | 8388
- Syntax (QFX Series) | 8389
- Syntax (SRX Series) | 8390
- Hierarchy Level | 8391
- Description | 8391
- Default | 8392
- Options | 8392
- Required Privilege Level | 8392
- Release Information | 8392

Syntax (EX Series and NFX Series)

```
igmp-snooping   {
    traceoptions  {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
regex>;
        flag flag (detail | disable | receive | send);
    }
    vlan (vlan-name | all) {
        data-forwarding {
            receiver {
                install;
                mode (proxy | transparent);
                (source-list | source-vlans) vlan-list;
                translate;
            }
            source {
                groups group-prefix;
            }
        }
```

```
        disable;
        immediate-leave;
        interface interface-name {
            group-limit limit;
            host-only-interface;
            immediate-leave;
            multicast-router-interface;
            static {
                group multicast-ip-address;
            }
        }
        (l2-querier | igmp-querier (QFabric Systems only)) {
            source-address ip-address;
        }
        proxy {
            source-address ip-address;
        }
        query-interval seconds;
        query-last-member-interval seconds;
        query-response-interval seconds;
        robust-count number;
        version number;
    }
}
```

Syntax (ACX Series, EX9200, and MX Series)

```
igmp-snooping {
    evpn-ssm-reports-only;
    immediate-leave;
    interface interface-name {
        group-limit limit;
        host-only-interface;
        immediate-leave;
        multicast-router-interface;
        static {
            group ip-address {
                source ip-address;
            }
        }
    }
```

```
    proxy {
        source-address ip-address;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
    vlan vlan-id {
        immediate-leave;
        interface interface-name {
            group-limit limit;
            host-only-interface;
            immediate-leave;
            multicast-router-interface;
            static {
                group ip-address {
                    source ip-address;
                }
            }
        }
        proxy {
            source-address ip-address;
        }
        query-interval seconds;
        query-last-member-interval seconds;
        query-response-interval seconds;
        robust-count number;
    }
}
```

Syntax (QFX Series)

```
igmp-snooping   {
    traceoptions  {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
regex>;
        flag flag (detail | disable | receive | send);
    }
    vlan (vlan-name | all) {
        evpn-ssm-reports-only;
        immediate-leave;
```

```
        interface interface-name {
            group-limit limit;
            host-only-interface;
            immediate-leave;
            multicast-router-interface;
            static {
                group multicast-ip-address;
            }
        }
        (l2-querier | igmp-querier (QFabric Systems only)) {
            source-address ip-address;
        }
        proxy {
            source-address ip-address;
        }
        query-interval seconds;
        query-last-member-interval seconds;
        query-response-interval seconds;
        robust-count number;
        version number;
    }
}
```

Syntax (SRX Series)

```
igmp-snooping   {
    vlan (all | vlan-name) {
        immediate-leave;
        interface interface-name {
            group-limit range;
            host-only-interface;
            multicast-router-interface;
            immediate-leave;
            static {
                group multicast-ip-address {
                source ip-address;
                }
            }
        }
        l2-querier {
            source-address ip-address;
```

```
            }
        proxy {
            source-address ip-address;
        }
        qualified-vlan vlan-id;
        query-interval number;
        query-last-member-interval number;
        query-response-interval number;
    robust-count number;
    traceoptions  {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier>;
        }
    }
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols]
[edit routing-instances routing-instance-name protocols]
[edit protocols]
```

Description

Configure IGMP snooping, which constrains multicast traffic to only the ports that have receivers attached.

IGMP snooping enables the device to selectively send out multicast packets on only the ports that need them. Without IGMP snooping, the device floods the packets on every port. The device listens for the exchange of IGMP messages by the device and the end hosts. In this way, the device builds an IGMP snooping table that has a list of all the ports that have requested a particular multicast group.

You can also configure IGMP proxy, IGMP querier, and multicast VLAN registration (MVR) functions on VLANs at this hierarchy level.

NOTE: IGMP snooping must be disabled on the device before running an ISSU operation.

Default

For most devices, IGMP snooping is disabled on the device by default, and you must configure IGMP snooping parameters in this statement hierarchy to enable it on one or more VLANs.

On legacy switches that do not support the Enhanced Layer 2 Software (ELS) configuration style, IGMP snooping is enabled by default on all VLANs, and the `vlan` statement includes a `disable` option if you want to disable IGMP snooping selectively on some VLANs or disable it on all VLANs.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

- *IGMP Snooping Overview*
- *Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment*
- *Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment*
- IGMP Snooping in MC-LAG Active-Active Mode
- *Configuring IGMP Snooping on Switches*
- *Example: Configuring IGMP Snooping on SRX Series Devices*
- *Example: Preserving Bandwidth with IGMP Snooping in an EVPN-VXLAN Environment*

igmp-snooping-options

IN THIS SECTION

- Syntax | 8393
- Hierarchy Level | 8393
- Description | 8393
- Options | 8393
- Required Privilege Level | 8394
- Release Information | 8394

Syntax

```
igmp-snooping-options {
    snoop-pseudowires
    use-p2mp-lsp
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name]
[edit routing-instances routing-instance-name ]
```

Description

Supports the `use-p2mp-lsp` or `snoop-pseudowires` options for independent routing instances and those in a logical system.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

instance-type | 8773

Example: Configuring IGMP Snooping

ignore

IN THIS SECTION

- See | 8394

See

"advertise" on page 2007

ignore-all

IN THIS SECTION

- Syntax | 8395
- Hierarchy Level | 8395

- Description | 8395
- Required Privilege Level | 8395
- Release Information | 8395

Syntax

```
ignore-all;
```

Hierarchy Level

```
[edit interfaces interface-name serial-options dce-options],
[edit interfaces interface-name serial-options dte-options]
```

Description

Ignore all control leads. You can include the `ignore-all` statement in the configuration only if you do not explicitly enable other signal handling options at the `dte-options` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Signal Handling

ignore-attached-bit

IN THIS SECTION

- Syntax | 8396
- Hierarchy Level | 8396
- Description | 8396
- Default | 8397
- Required Privilege Level | 8397
- Release Information | 8397

Syntax

```
ignore-attached-bit;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Ignore the attached bit on IS-IS Level 1 routers. Configuring this statement enables the routing device to ignore the attached bit on incoming Level 1 link-state PDUs. If the attached bit is ignored, no default route, which points to the routing device which has set the attached bit, is installed.

There might be times, such as during a denial-of-service (DoS) attack, that you do not want a Level 1 router to be able to forward traffic based on a default route.

To prevent a routing device from being able to reach interarea destinations, you can prevent the routing device from installing the default route without affecting the status of its IS-IS adjacencies. The `ignore-`

`attached-bit` statement is used to tell the routing device to ignore the presence of the attached bit in Level 1 link-state PDUs, which blocks the installation of the IS-IS default route.

Default

The `ignore-attached-bit` statement is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

ignore-encapsulation-mismatch

IN THIS SECTION

- Syntax | 8397
- Hierarchy Level | 8398
- Description | 8398
- Required Privilege Level | 8398
- Release Information | 8398

Syntax

```
ignore-encapsulation-mismatch;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name neighbor neighbor-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor neighbor-id],
[edit protocols l2circuit local-switching interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name],
[edit routing-instances routing-instance-name protocols evpn interface interface-name],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name neighbor
neighbor-id],
[edit routing-instances routing-instance-name protocols vpls neighbor neighbor-id]
```

Description

Allow a Layer 2 circuit, VPLS, or EVPN to be established even though the encapsulation configured on the CE device interface does not match the encapsulation configured on the Layer 2 circuit, VPLS, or EVPN interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 9.2.

Statement extended to support local switching in Junos OS Release 10.4.

Statement introduced for EVPNs in Junos OS Release 13.2 for MX 3D Series.

RELATED DOCUMENTATION

> *Configuring EVPN Routing Instances*

Enabling the Layer 2 Circuit When the Encapsulation Does Not Match

ignore-entry

IN THIS SECTION

- See | 8399

See

"force-entry" on page 7370

ignore (Event Policy)

IN THIS SECTION

- Syntax | 8399
- Hierarchy Level | 8400
- Description | 8400
- Required Privilege Level | 8400
- Release Information | 8400

Syntax

```
ignore;
```

Hierarchy Level

```
[edit event-options policy policy-name then]
```

Description

Define a policy that ignores particular events. If one or more of the listed events occur, a system log message for the event is not generated, and no further policies associated with this event are processed. If you include the `ignore` statement in a policy configuration, you cannot configure any other actions in the policy.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configure Event Policies to Ignore an Event

ignore-l3-incompletes

IN THIS SECTION

- Syntax | 8401
- Hierarchy Level | 8401
- Description | 8401
- Required Privilege Level | 8401
- Release Information | 8401

Syntax

```
ignore-l3-incompletes;
```

Hierarchy Level

```
[edit interfaces interface-name fastether-options],
[edit interfaces interface-name gigether-options]
```

Description

Ignore the counting of Layer 3 incomplete errors on Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

- Ignoring Layer 3 Incomplete Errors

ignore-lsp-metrics

IN THIS SECTION

- Syntax | 8402
- Hierarchy Level | 8402
- Description | 8402
- Required Privilege Level | 8402
- Release Information | 8403

Syntax

```
ignore-lsp-metrics;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf traffic-engineering shortcuts],
[edit protocols ospf traffic-engineering shortcuts]
```

Description

Cause OSPF to ignore the RSVP LSP metric.

Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This can cause interoperability problems when you configure LDP tunneling over RSVP LSPs in heterogeneous networks.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks

ignore-magic-number-mismatch (Access Group Profile)

IN THIS SECTION

- Syntax | 8403
- Hierarchy Level | 8403
- Description | 8403
- Required Privilege Level | 8404
- Release Information | 8404

Syntax

```
ignore-magic-number-mismatch;
```

Hierarchy Level

```
[edit access group-profile name ppp ppp-options]
```

Description

Prevent the Packet Forwarding Engine from performing a validation check for magic numbers received in LCP keepalive (Echo-Request/Echo-Reply) exchanges for a group of tunneled PPP subscribers at the LNS.

A mismatch occurs when the PPP magic number received from a remote peer in the keepalive exchange does not match the value agreed upon during LCP negotiation. Disabling the validation check prevents PPP from terminating the session when an unexpected number is received. Configuring this statement has no effect on LCP magic number negotiation or on the exchange of keepalives when the remote peer magic number is the expected negotiated number.

NOTE: Because magic number validation is not performed, the Packet Forwarding Engine does not detect whether the remote peer sends the local peer's magic number, which would indicate a loopback or other network issue. This is considered to be an unlikely situation, because LCP negotiation completed successfully, meaning no loopback was present at that time.

NOTE: You can also configure this behavior in a dynamic PPP profile. When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

This means that `ignore-magic-number-mismatch` configured in a group profile is not applied when the dynamic profile includes any PPP option, even when the dynamic profile does not include `ignore-magic-number-mismatch` statement.

NOTE: This statement is not supported on static interfaces.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

ignore-magic-number-mismatch (Dynamic Profiles)

IN THIS SECTION

- Syntax | 8405
- Hierarchy Level | 8405
- Description | 8405
- Required Privilege Level | 8406
- Release Information | 8406

Syntax

```
ignore-magic-number-mismatch;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces  pp0 unit “$junos-interface-unit” ppp-options],
[edit dynamic-profiles profile-name interfaces  "$junos-interface-ifd-name" unit “$junos-
interface-unit” ppp-options]
```

Description

Prevent the Packet Forwarding Engine from performing a validation check for magic numbers received in LCP keepalive (Echo-Request/Echo-Reply) exchanges for dynamic PPP subscriber connections terminated at the router or for dynamic tunneled PPP subscribers on LNS inline service interfaces.

A mismatch occurs when the PPP magic number received from a remote peer in the keepalive exchange does not match the value agreed upon during LCP negotiation. Disabling the validation check prevents PPP from terminating the session when an unexpected number is received. Configuring this statement has no effect on LCP magic number negotiation or on the exchange of keepalives when the remote peer magic number is the expected negotiated number.

For dynamic PPP subscriber connections terminated at the router, configure the statement at the [edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit” ppp-options] hierarchy level.

For dynamic tunneled PPP subscribers on LNS inline service interfaces, configure the statement at the [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit” ppp-options] hierarchy level.

NOTE: Because magic number validation is not performed, the Packet Forwarding Engine does not detect whether the remote peer sends the local peer’s magic number, which would indicate a loopback or other network issue. This is considered to be an unlikely situation, because LCP negotiation completed successfully, meaning no loopback was present at that time.

NOTE: You can also configure this behavior in an L2TP group profile that applies to tunneled PPP subscribers at the LNS. When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

This means that "ignore-magic-number-mismatch" on page 8403 configured in a group profile is not applied when the dynamic profile includes any PPP option, even when the dynamic profile does not include "ignore-magic-number-mismatch" on page 8405.

NOTE: This statement is not supported on static interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

PPP Subscriber Access Networks Overview

PPP Subscriber Access Networks Overview

L2TP LNS Inline Service Interfaces

ignore-memory-overflow

IN THIS SECTION

- Syntax | 8407
- Hierarchy Level | 8407
- Description | 8407
- Required Privilege Level | 8407
- Release Information | 8408

Syntax

```
(ignore-memory-overflow | no-ignore-memory-overflow);
```

Hierarchy Level

```
[edit security idp sensor-configuration re-assembler]
```

Description

Enable the TCP reassembler to ignore the memory overflow to prevent the dropping of IDP custom applications. By default this feature is enabled.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

ignore-mtu-mismatch

IN THIS SECTION

- Syntax | 8408
- Hierarchy Level | 8408
- Description | 8409
- Required Privilege Level | 8409
- Release Information | 8409

Syntax

```
ignore-mtu-mismatch;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit protocols l2circuit local-switching interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name],
```

```
[edit routing-instances routing-instance-name protocols l2vpn interface interface-name],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Ignore the MTU configuration set for the physical interface associated with the local switching interface or with the remote PE router. This allows a pseudowire to be brought up between two logical interfaces that are defined on physical interfaces with different MTU values.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 8.5.

Support for remote PE routers added in Junos OS Release 9.2.

Support for Layer 2 VPNs and VPLS added in Junos OS Release 10.4.

RELATED DOCUMENTATION

Enabling Local Interface Switching When the MTU Does Not Match

Configuring the MTU for Layer 2 Interfaces

ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-overflow

IN THIS SECTION

- Syntax | 8410
- Hierarchy Level | 8410
- Description | 8410

Required Privilege Level | 8410
Release Information | 8410

Syntax

```
(ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-overflow);
```

Hierarchy Level

```
[edit security idp sensor-configuration re-assembler]
```

Description

Reassembly memory overflow occurs when the memory allocated for the reassembly of TCP fragments is exceeded. When the reassembly of TCP fragments exceeds the memory limit, defined with `max-packet-mem-ratio`, you can define the system behavior to ignore or drop the offending packets. If the `ignore-reassembly-memory-overflow` command is enabled on the SRX device, IDP will ignore and permit packets from sessions which trigger a reassembly memory overflow. If you enable the `no-ignore-reassembly-memory-overflow` command when reassembly memory overflow occurs, packets of that session are dropped by the device. By default, the `ignore-reassembly-memory-overflow` command is enabled.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

max-packet-mem-ratio

ignore-reassembly-overflow

IN THIS SECTION

- Syntax | 8411
- Hierarchy Level | 8411
- Description | 8411
- Required Privilege Level | 8411
- Release Information | 8411

Syntax

```
ignore-reassembly-overflow
```

Hierarchy Level

```
[edit security idp sensor-configuration re-assembler]
```

Description

Enable the TCP reassembler to ignore the global reassembly overflow to prevent the dropping of application traffic. This feature is enabled by default.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

ignore-stp-topology-change

IN THIS SECTION

- Syntax | 8412
- Hierarchy Level | 8412
- Description | 8412
- Required Privilege Level | 8412
- Release Information | 8412

Syntax

```
ignore-stp-topology-change;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name multicast-snooping-options],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name multicast-snooping-
options]
```

Description

Ignore messages about spanning tree topology changes. This statement is supported for the **virtual-switch** routing instance type only.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Example: Configuring Multicast Snooping

ignore-server-validation (Application Identification)

IN THIS SECTION

- Syntax | 8413
- Hierarchy Level | 8413
- Description | 8413
- Required Privilege Level | 8413
- Release Information | 8414

Syntax

```
ignore-server-validation;
```

Hierarchy Level

```
[edit services application-identification download]
```

Description

Disable server authentication for application signature package download.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1 on SRX Series Firewalls.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Application Identification Overview

Downloading and Installing Predefined Junos OS Application Signature Packages

igp-synchronization

IN THIS SECTION

- Syntax | 8414
- Hierarchy Level | 8414
- Description | 8415
- Options | 8415
- Required Privilege Level | 8415
- Release Information | 8415

Syntax

```
igp-synchronization holddown-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
```

```
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.

Options

`holddown-interval` *seconds*—Time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational.

- **Default:** 10 seconds
- **Range:** 10 through 60 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring LDP Synchronization with the IGP on the Router

ike

IN THIS SECTION

Syntax | 8416

- Hierarchy Level | 8417
- Description | 8417
- Required Privilege Level | 8417
- Release Information | 8417

Syntax

```
ike {
    proposal proposal-name {
        authentication-algorithm (sha1 | sha-256 | sha-384);
        authentication-method (ecdsa-signatures-256 | ecdsa-signatures-384 | pre-shared-keys | rsa-
signatures);
        description description;
        dh-group  (group1 | group2 | group5 |group14 | group15 | group16 | group19 | group20 |
group24);
        encryption-algorithm algorithm;
        lifetime-seconds seconds;
    }
    policy policy-name {
        description description;
        local-certificate identifier;
        local-id (ipv4_addr ipv4-address | ipv6-addr ipv6-address | key-id identifier);
        version (1 | 2);
        mode (aggressive | main);
        pre-shared-key (ascii-text key | hexadecimal key);
        proposals [ proposal-names ];
        remote-id {
            any-remote-id;
            ipv4_addr [ values ];
            ipv6_addr [ values ];
            key_id [ values ];
        }
    }
}
```

Hierarchy Level

```
[edit services ipsec-vpn]
```

Description

Configure IKE.

The remaining statements are explained separately.

> **NOTE**: In Junos FIPS mode, the `aggressive` option of the "mode" on page 10797 statement is not supported.

> **NOTE**: In Junos FIPS mode, ECDSA options of the `authentication-method` statement are not supported in Junos OS Release 17.3R1. Starting in Junos OS Release 17.4R1, ECDSA is supported in Junos FIPS mode.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IKE Proposals

Configuring IKE Policies

ike-access-profile

IN THIS SECTION

- Syntax | 8418
- Hierarchy Level | 8418
- Description | 8418
- Options | 8418
- Required Privilege Level | 8418
- Release Information | 8419

Syntax

```
ike-access-profile profile-name;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Define the access profile for the IPsec traffic on dynamic tunnels.

Options

profile-name—Identifier for access profile, which must match the name configured at the `[edit access profile name client * ike]` hierarchy level.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Dynamic Endpoints for IPsec Tunnels

Configuring IPsec Service Sets

ike (Access Profile)

IN THIS SECTION

- Syntax | 8419
- Hierarchy Level | 8420
- Description | 8420
- Options | 8420
- Required Privilege Level | 8421
- Release Information | 8421

Syntax

```
ike {
    allowed-proxy-pair {
        remote remote-proxy-address local local-proxy-address;
    }
    dead-peer-detection{
        interval seconds
        threshold number
    }
    ike-policy policy-name;
    initiate-dead-peer-detection;
    interface-id string-value;
    ipsec-policy ipsec-policy;
```

```
    pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
    reverse-route
}
```

Hierarchy Level

```
[edit access profile profile-name client client-name]
```

Description

Configure an IKE access profile.

The remaining statements are explained below.

NOTE: This statement is not supported on MX Series routers.

Options

allowed-proxy-pair

Specify the network address of the local and remote peer associated with an IKE access profile.

- Values:
  - local *local-proxy-address*—Network address of the local peer. Default: 0.0.0.0
  - remote *remote-proxy-address*—Network address of the remote peer. Default: 0.0.0.0

ike-policy *policy-name*

Specify the IKE policy used to authenticate dynamic peers during IKE negotiation.

- Values:
  - *policy-name*—The name of an IKE policy configured at the `[edit services ipsec-vpn ike policy policy-name]` hierarchy level. The IKE policy defines either the local digital certificate or the pre-shared key used for IKE authentication with dynamic peers.

initiate-dead-peer-detection

Detect inactive peers on dynamic IPsec tunnels.

pre-shared-key (ascii-text *character-string* | hexadecimal *hexadecimal-digits*);

Configure the key used to authenticate a dynamic peer during IKE phase 1 negotiation. Specify the key in either ASCII or hexadecimal format.

- Values:
  - ascii-text *character-string*—Authentication key in ASCII format.
  - hexadecimal *hexadecimal-digits*—Authentication key in hexadecimal format.

reverse-route

Configure a reverse route for dynamic endpoint IPsec tunnels.

- **Values:** preference *metric-value*

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

`ike-policy` statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring an IKE Access Profile

Configuring Services Interfaces

ike-esp-nat

IN THIS SECTION

Syntax | 8422

- Hierarchy Level | 8422
- Description | 8422
- Options | 8423
- Required Privilege Level | 8423
- Release Information | 8423

Syntax

```
ike-esp-nat {
    enable;
    esp-gate-timeout seconds;
    esp-session-timeout seconds;
    state-timeout seconds;
    traceoptions {
        flag {
            all {
                extensive
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Configure Application Layer Gateway (ALG) for Internet Key Exchange (IKE) and Encapsulating Security Payload (ESP) traffic with Network Address Translation (NAT).

Options

- `Enable` —Enable the IKE-ESP ALG.
- `esp-gate-timeout` *seconds*—Set the timeout for the ESP gates created after an IKE Phase 2 exchange has completed.

  Range: 2 through 30 seconds.

  Default: 5 seconds.
- `esp-session-timeout` *seconds*—Set the idle timeout for the ESP sessions created from the IPsec gates.

  Range: 60 through 2400 seconds.

  Default: 1800 seconds.
- `state-timeout` *seconds*—Set the timeout for the ALG state information.

  Range: 180 through 86,400 seconds.

  Default: 14,400 seconds.
- `traceoptions`—Set the IKE-ESP ALG trace options.
  - `flag` —Specify which tracing operation to perform.
    - `all`—Trace all operations.
      - `extensive`—Set trace verbosity level to extensive.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Statement supported at hierarchy [`edit logical-systems name tenants name security alg tenant`] in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

> *Understanding the IKE and ESP ALG*

ike (High Availability)

IN THIS SECTION

- Syntax | 8424
- Hierarchy Level | 8425
- Description | 8425
- Options | 8425
- Required Privilege Level | 8425
- Release Information | 8425

Syntax

```
ike {
    gateway name {
        ike-policy policy-name;
        version (v1-only | v2-only);
    }
    policy name {
        description description;
        pre-shared-key (ascii-text ascii-text | hexadecimal hexadecimal);
        proposals [ proposals ... ];
    }
    proposal name {
        authentication-algorithm (md5 | sha-256 | sha-384 | sha-512 | sha1);
        authentication-method (dsa-signatures | ecdsa-signatures-256 | ecdsa-signatures-384 |
ecdsa-signatures-521 | pre-shared-keys | rsa-signatures);
        description description;
        dh-group (group1 | group14 | group15 | group16 | group19 | group2 | group20 | group21 |
group24 | group5);
        encryption-algorithm (aes-256-gcm);
        lifetime-seconds seconds;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Define Internet Key Exchange (IKE) configuration for high availability feature. IKE is a key management protocol that creates dynamic SAs; it negotiates SAs for IPsec. An IKE configuration defines the algorithms and keys used to establish a secure connection with a peer security gateway.

Options

gateway-name Name of the gateway.

ike-policy Specify the IKE policy to be used for the gateway.

version Specify the IKE version to use to initiate the connection.

- Values:
  - v2-only—The connection must be initiated using IKE version 2. For Multinode High availability feature, you must configure the IKE version as `v2-only`.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

ipsec (High Availability) | 9301

High-Availability (Chassis)

Multinode High Availability

services-redundancy-group
local-id
peer-id
traceoptions
show security ike security-associations | 24648
show security ike active-peer | 24630
show security ipsec security-associations | 24701
show security ipsec statistics | 24744
clear security ike security-associations | 3988
clear security ipsec security-associations | 3993

ike-log

IN THIS SECTION

- Syntax | 8426
- Hierarchy Level | 8427
- Description | 8427
- Options | 8427
- Required Privilege Level | 8427
- Release Information | 8427

Syntax

```
ike-log {
    filename filename;
    level level;
}
```

Hierarchy Level

```
[edit security host-vpn]
```

Description

Configure IKE logging details.

Options

filename *filename*

Specify the filename for logging—for example, **/var/log/*filename*.log**.

- **Default: /var/log/charon.log**

level *level*

Specify the logging level (how much detail is included) of the IKE daemon's log messages. Following are the levels of logging available:

- 0: Very basic auditing logs (for example SA up/SA down).
- 1: Generic control flow with errors. This is the default.
- 2: More detailed debugging control flow.
- 3: Including RAW (unprocessed) data dumps in hex.
- **Range:** 0 (minimal logging) through 3
- **Default:** 1

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 18.3R1.

ike-phase1-failures

IN THIS SECTION

- Syntax | 8428
- Hierarchy Level | 8428
- Description | 8428
- Default | 8428
- Options | 8429
- Required Privilege Level | 8429
- Release Information | 8429

Syntax

```
ike-phase1-failures {
    threshold value;
}
```

Hierarchy Level

```
[edit security alarms potential-violation]
```

Description

Raise a security alarm after exceeding a specified number of Internet Key Exchange (IKE) Phase 1 failures.

Default

Multiple IKE phase 1 failures do not cause an alarm to be raised.

Options

failures—Number of IKE phase 1 failures up to which an alarm is not raised. When the configured number is exceeded, an alarm is raised.

- **Range:** 1 through 1,000,000,000.
- **Default:** 20

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

IPsec Overview

Monitoring VPN Traffic

ike-phase2-failures

IN THIS SECTION

- Syntax | 8430
- Hierarchy Level | 8430
- Description | 8430
- Default | 8430
- Options | 8430
- Required Privilege Level | 8430
- Release Information | 8430

Syntax

```
ike-phase2-failures {
    threshold value;
}
```

Hierarchy Level

```
[edit security alarms potential-violation]
```

Description

Raise a security alarm after exceeding a specified number of Internet Key Exchange (IKE) phase 2 failures.

Default

Multiple IKE phase 2 failures do not cause an alarm to be raised.

Options

failures—Number of IKE phase 2 failures up to which an alarm is not raised. When the configured number is exceeded, an alarm is raised.

- **Range:** 1 through 1,000,000,000.
- **Default:** 20

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

IPsec Overview

Monitoring VPN Traffic

ike-secrets

IN THIS SECTION

- Syntax | 8431
- Hierarchy Level | 8431
- Description | 8431
- Options | 8432
- Required Privilege Level | 8432
- Release Information | 8432

Syntax

```
ike-secrets {
    ike-secret {
        id id;
        secret (ascii-text ascii-text | hexadecimal hexadecimal);
    }
}
```

Hierarchy Level

```
[edit security host-vpn]
```

Description

Configure IKE shared secret details. If the shared secret passed between devices does not match, a session is not set up.

NOTE: To configure a host-to-host VPN connection, you must configure the `ike-secrets` statement.

Options

ike-secret	Specify the name of the IKE secret.
id *id*	Specify the identity that the secret belongs to—for example, an IP address, a domain name, or an e-mail address. This identity matches local and remote identities as exchanged in the IKE security association (SA) negotiation at the `[edit security host-vpn connections` hierarchy level.
secret (ascii-text *ascii-text* **| hexadecimal** *hexadecimal***)**	Define the preshared private key associated with the identity in either ASCII or hexadecimal format.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

connections (Host VPN) | 4544

ike (Security)

IN THIS SECTION

Syntax | 8433

- Hierarchy Level | 8436
- Description | 8436
- Options | 8436
- Required Privilege Level | 8436
- Release Information | 8437

Syntax

```
ike {
    gateway (Security IKE) name {
        ( address | dynamic (Security) distinguished-name (Security) < container> < wildcard>
hostname inet inet6  user-at-hostname <connections-limit connections-limit> <ike-user-type
(group-ike-id | shared-ike-id)> <reject-duplicate-connection>);
        aaa {
             access-profile;
            client password password username username;
        }
        advpn {
            partner {
                connection-limit connection-limit;
                disable;
                idle-threshold idle-threshold;
                idle-time seconds;
            }
            suggester {
                disable;
            }
        }
        dead-peer-detection (always-send | optimized | probe-idle-tunnel);
        external-interface external-interface;
        fragmentation {
            disable;
            size size;
        }
        general-ikeid;
         ike-policy;
         local-address;
        local-identity (distinguished-name | hostname identity-hostname | inet identity-ipv4 |
```

```
inet6 identity-ipv6 | key-id string-key-id | user-at-hostname identity-user);
        remote-identity distinguished-name <container container> <wildcard wildcard>hostname
identity-hostnameinet identity-ipv4inet6 identity-ipv6 key-id string-key-id user-at-hostname
identity-user;
         tcp-encap-profile profile-name;
        version (v1-only | v2-only);
    }
    policy name {
        certificate {
            local-certificate (Security) local-certificate;
            peer-certificate-type (pkcs7 | x509-signature);
            policy-oids policy-oids;
            trusted-ca (ca-profile ca-profile | trusted-ca-group trusted-ca-group  );
        }
        description description;
        mode (aggressive | main);
        pre-shared-key (ascii-text ascii-text | hexadecimal hexadecimal);
        seeded-pre-shared-key (ascii-text key | hexadecimal key);
        proposal-set (Security IKE) (basic | compatible | prime-128 | prime-256 | standard | suiteb-
gcm-128 | suiteb-gcm-256);
        proposals [ proposals ... ];
        reauth-frequency reauth-frequency;
    }
    proposal proposal-name {
        authentication-algorithm (md5 | sha-256 | sha-384 | sha-512 | sha1);
        authentication-method (certificates | dsa-signatures | ecdsa-signatures-256 | ecdsa-
signatures-384 | ecdsa-signatures-521 | pre-shared-keys | rsa-signatures);
        description description;
        dh-group dh-group (group1 | group14 | group15 | group16 | group19 | group2 | group20 |
group21 | group24 | group5);
        encryption-algorithm  (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-256-cbc |
aes-256-gcm | des-cbc);
        lifetime-seconds seconds;
    }
    respond-bad-spi <max-responses>;
    session {
        full-open {
            incoming-exchange-max-rates {
                ike-rekey value;
                ipsec-rekey value;
                keepalive value;
            }
        }
```

```
        half-open {
            timeout seconds;
            backoff-timeouts {
               init-phase-failure value;
               auth-phase-failure value;
            }
            discard-duplicate;
            max-count value;
            thresholds {
                send-cookie count;
                reduce-timeout count timeout seconds;
            }
         }
    }
    blocklists {
        blocklist-name {
            description text-description;
            rule rule-name {
                match {
                    role (initiator | responder);
                    id-type (inet | inet6 | hostname | distinguished-name | user-at-hostname |
key-id);
                    id-pattern value;
                }
                then {
                    (discard | reject);
                     backoff timeout-value;
                }
            }
       }
   }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }

        level (critical | error | terse | warning | detail);
        flag flag (all | certificates | config | database | general | high-availability | ike |
next-hop-tunnels | parse | policy-manager | routing-socket | thread | timer)reference/
```

```
configuration-statement/security-edit-ike-security;
        no-remote-trace;
        rate-limit messages-per-second;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Define Internet Key Exchange (IKE) configuration. IKE is a key management protocol that creates dynamic SAs; it negotiates SAs for IPsec. An IKE configuration defines the algorithms and keys used to establish a secure connection with a peer security gateway.

Options

`respond-bad-spi` *`max-responses`*—(Optional) Number of times to respond to invalid SPI values per gateway. Enable response to invalid IPsec Security Parameter Index (SPI) values. If the security associations (SAs) between two peers of an IPsec VPN become unsynchronized, the device resets the state of a peer so that the two peers are synchronized.

- **Range:** 1 through 30
- **Default:** 5

`traceoptions`—Configure IKE tracing options to aid in troubleshooting the IKE issues. This helps troubleshoot one or multiple tunnels negotiation by standard tracefile configuration. IKE tracing allows the user to view the detailed packet exchange and the negotiation information in Phase 1 and Phase 2. IKE tracing is not enabled by default. By default , all IKE or IPsec negotiations are logged into /var/log/kmd. But user can also specify customized file name while configuring the IKE traceoptions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

Support for IPv6 addresses added in Junos OS Release 11.1.

Support for `inet6` option added in Junos OS Release 11.1.

Support for `group15`, `group16`, `group21`, `ecdsa-signatures-521`, and `sha-512` options added in Junos OS Release 19.1R1 on SRX5000 line with `junos-ike` package installed.

Starting in Junos OS Release 20.2R1, we've changed the help text description as `NOT RECOMMENDED` for the CLI options `md5` and `sha1` for devices running IKED with `junos-ike` package installed.

Support for `group15`, `group16`, and `group21` options added in Junos OS Release 20.3R1 on vSRX Virtual Firewall instances with `junos-ike` package installed.

Support for `group15`, `group16`, and `group21` options added in Junos OS Release 21.1R1 on vSRX Virtual Firewall 3.0 instances with `junos-ike` package installed.

`level` option introduced in Junos OS Release 21.1R1.

Support for `seeded-pre-shared-key` option added in Junos OS Release 21.1R1.

Support for `session` and `blocklists` option added in Junos OS Release 23.4R1

RELATED DOCUMENTATION

IPsec Overview

ALG Overview

Understanding Logical Systems for SRX Series Services Gateways

Example: Configuring AutoVPN with Pre-Shared Key

Installing Junos IKE package

IKE Protection from DDoS Attacks

session (Security IKE)

blocklists (Security IKE)

ike (Security Group VPN Member)

IN THIS SECTION

- Syntax | 8438
- Hierarchy Level | 8439
- Description | 8439
- Options | 8439
- Required Privilege Level | 8440
- Release Information | 8440

Syntax

```
ike {
    gateway gateway-name {
        ike-policy policy-name;
        local address ip-address;
        local-identity {
            (hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname e-mail-
address);
        }
        remote-identity {
            (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
        }
        routing-instance routing-instance;
        server-address ip-address;
    }
    policy policy-name {
        description description;
        mode (aggressive | main);
        pre-shared-key (ascii-text key | hexadecimal key);
        proposals proposal-name;
    }
    proposal proposal-name {
        authentication-algorithm (sha-256 | sha-384);
        authentication-method pre-shared-keys;
        description description;
```

```
        dh-group (group14 | group24);
        encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
        lifetime-seconds seconds;
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag (all | certificates | config | database | general | high-availability | ike |
next-hop-tunnels | parse | policy-manager | routing-socket | thread | timer);
        gateway-filter {
            local-address ip-address;
            remote-address ip-address;
        }
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security group-vpn member]
```

Description

Configure IKE group VPN on the group member. A group member encrypts the traffic and is responsible for the actual encryption and decryption of data traffic. A group member is configured with IKE Phase 1 parameters and GC/KS information.

Options

gateway *gateway-name* Configure IKE gateway for group VPN member.

policy *policy-name* Configure an IKE policy.

proposal*proposal-name*	Define an IKE proposal.
traceoptions	Configure group VPN tracing options to aid in troubleshooting the IKE issues.
ipsec	Configure IPsec for Phase 2 exchange on the group member.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

ike (Security Group VPN Server)

IN THIS SECTION

- Syntax | 8441
- Hierarchy Level | 8442
- Description | 8442
- Options | 8442
- Required Privilege Level | 8442
- Release Information | 8442

Syntax

```
ike {
    gateway gateway-name {
        address ip-address;
        dead-peer-detection {
            always-send;
            interval seconds;
            threshold number;
        }
        dynamic {
            (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
        }
        ike-policy policy-name;
        local-address ip-address;
        local-identity {
            (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
        }
        remote-identity {
            (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
        }
        routing-instance routing-instance;
    }
    policy policy-name {
        description description;
        mode (aggressive | main);
        pre-shared-key (ascii-text key | hexadecimal key);
        proposals proposal-name;
    }
    proposal proposal-name {
        authentication-algorithm (sha-256 | sha-384);
        authentication-method pre-shared-keys;
        description description;
        dh-group (group14 | group24);
        encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
    }
}
```

Hierarchy Level

```
[edit security group-vpn server]
```

Description

Configure Phase 1 security association (SA) with a member on the group server. The gateway is the group member. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 devices and vSRX Virtual Firewall instances.

Options

gateway ***gateway-name***	Configure IKE gateway for group VPN server.
policy ***policy-name***	Configure an IKE policy.
proposal ***proposal-name***	Define an IKE proposal.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

ike (Security IPsec VPN)

IN THIS SECTION

- Syntax | 8443
- Hierarchy Level | 8443
- Description | 8443
- Options | 8444
- Required Privilege Level | 8445
- Release Information | 8445

Syntax

```
ike {
    anti-replay-window-size anti-replay-window-size;
    gateway gateway-name;
    idle-time seconds;
    install-interval seconds;
    ipsec-policy ipsec-policy-name;
    no-anti-replay;
    proxy-identity {
        local ip-prefix;
        remote ip-prefix;
        service (any | service-name);
    }
}
```

Hierarchy Level

```
[edit security ipsec vpn vpn-name]
```

Description

Define an IKE-keyed IPsec VPN.

Options

anti-replay-window-size

To enable the `anti-replay-window-size` option, you first need to configure the option for each VPN object or at the global level. You can configure the anti-replay window size in the range of 64 to 8192 (power of 2). If the anti-replay window size is not configured, the window size is 64 by default. If `anti-replay-window-size` command is configured at both the global and VPN object levels, the configuration on VPN object takes precedence over global configuration.

`anti-replay-window-size` is supported only on SRX5000 line with SRX5K-SPC3 card installed.

gateway-name

Name of the remote IKE gateway.

idle-time

Specify the maximum amount of idle time to delete a security association (SA) when there is no traffic flow.

- **Default:** Disabled
- **Range:** 60 through 999,999 seconds

install-interval

Specify the maximum number of seconds to allow the installation of a rekeyed outbound security association (SA) on the device.

- **Default:**
  - 1 second, prior to Junos OS Release 23.4R1 (without iked process)
  - Starting Junos OS Release 23.4R1 with iked process:
    - 3 seconds, for IKEv1 initiator and IKEv2 responder.
    - 0 seconds, for rest of the scenarios
- **Range:** 0 through 10 seconds.

  You can configure 0-10 seconds from CLI, and it takes effect in data plane only for IKEv1 initiator or IKEv2 responder. CLI configured value takes precedence over the default value.

ipsec-policy

Specify the IPsec policy name.

no-anti-replay

Disable the antireplay checking feature of IPsec. Antireplay is an IPsec feature that can detect when a packet is intercepted and then replayed by attackers. By default, antireplay checking is enabled.

proxy-identity Optionally specify the IPsec proxy ID to use in negotiations. The default is the identity based on the IKE gateway. If the IKE gateway is an IPv6 site-to-site gateway, the default proxy ID is ::/0. If the IKE gateway is an IPv4 gateway or a dynamic endpoint or dialup gateway, the default proxy ID is 0.0.0.0/0.

- `local`—Specify the local IPv4 or IPv6 address and subnet mask for the proxy identity.
- `remote`—Specify the remote IPv4 or IPv6 address and subnet mask for the proxy identity.
- `service`—Specify the service (port and protocol combination) to protect. Name of the service is as defined with `system-services` (`Interface Host-Inbound Traffic`) and `system-services` (`Zone Host-Inbound Traffic`).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support.

Statement `anti-replay-window-size` is introduced in Junos OS Release 19.2R1.

Support for `idle-time` and `install-interval` options with IPsec VPN running iked process is added in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

IPsec Overview

ike (Security Services)

IN THIS SECTION

- Syntax | 8446
- Hierarchy Level | 8447
- Description | 8447
- Options | 8447
- Required Privilege Level | 8447
- Release Information | 8447

Syntax

```
ike {
    policy ike-peer-address {
        description policy-description;
        encoding (binary | pem);
        identity identity-name;
        local-certificate certificate-filename;
        local-key-pair private-public-key-file;
        mode (aggressive | main);
        pre-shared-key (ascii-text key | hexadecimal key);
        proposals [ proposal-names ];
    }
    proposal ike-proposal-name {
        authentication-algorithm (md5 | sha1);
        authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
        dh-group (group1 | group2);
        encryption-algorithm (3des-cbc | des-cbc);
        lifetime-seconds seconds;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

(Encryption interface on M Series and T Series routers only) Configure IKE.

Options

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring an IKE Proposal for Dynamic SAs
- *Configuring an IKE Policy for Preshared Keys*

image

IN THIS SECTION

- Syntax | 8448
- Hierarchy Level | 8448
- Description | 8448

- Options | 8448
- Required Privilege Level | 8449
- Release Information | 8449

Syntax

```
image {
    file-path;
    bus-type [ide | virtio];
    driver-cache {
                   default-cache;
                   directsync;
                   none;
                   unsafe;
                   writeback;
                   writethrough;
                   }
    image-type [qcow2 | raw];
}
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Specify the VNF image source file. VNF image is virtual hard disk, which contains the bootable file-system for the VNF.

Options

file-path Path of the image source file

image-type Format of the image. Default value is qcow2.

bus-type Type of the system bus. Default value is virtio.

driver-cache Type of the driver cache mode. Default mode is default-cache

- default-cache—Default mode when you do not specify any driver cache mode.
- directsync—Enables the hypervisor to interact with the disk image file or block device with both `O_DSYNC` and `O_DIRECT` semantics. The hypervisor reports that the writes are completed only when the data is committed to the storage device, and when it is also desirable to bypass the host page cache.
- none—Enables the hypervisor to interact with the disk image file or block device with `O_DIRECT` semantics. The hypervisor bypasses the host page cache and I/O happens directly between the hypervisor user space buffers and the storage device. y.
- unsafe—Ignores all flush commands from the guests. Using this mode implies that the user has accepted the trade-off of performance over risk of data loss in case of a host failure.
- writeback—Enables the hypervisor to interact with the disk image file or block device with neither `O_DSYNC` nor `O_DIRECT` semantics. The hypervisor uses the host page cache and reports the writes to the guest as completed when the writes are in the host page cache.
- writethrough—Enables the hypervisor to interact with the disk image file or block device with `O_DSYNC` semantics. The hypervisor reports the writes as completed only when the data is committed to the storage device.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D45.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

imap-profile (Security Content Security Policy Antivirus)

IN THIS SECTION

- Syntax | 8450
- Hierarchy Level | 8450
- Description | 8450
- Required Privilege Level | 8450
- Release Information | 8450

Syntax

```
imap-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name anti-virus]
```

Description

Configure a Content Security policy for the antivirus IMAP protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The [edit security utm default-configuration] hierarchy level is introduced in Junos OS Release 18.2R1.

imap-profile (Security Content Security Policy Content Filtering)

IN THIS SECTION

- Syntax | 8451
- Hierarchy Level | 8451
- Description | 8451
- Required Privilege Level | 8451
- Release Information | 8452

Syntax

```
imap-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name content-filtering]
```

Description

Configure a Content Security policy for the content-filtering IMAP protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

immediate-accounting-response

IN THIS SECTION

- Syntax | 8452
- Hierarchy Level | 8452
- Description | 8453
- Default | 8453
- Options | 8453
- Required Privilege Level | 8453
- Release Information | 8453

Syntax

```
immediate-accounting-response (enabled | disabled);
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
```

Description

Enable or disable the sending of an immediate RADIUS response message to the accounting start message received from a GGSN, PGW, or BNG RADIUS client.

Default

If you do not specify an option, **disabled** is the default.

Options

enabled Enable immediate response.

disabled Disable immediate response. The response is sent after TDF subscriber creation is complete.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

immediate-leave

IN THIS SECTION

- Syntax | 8454
- Hierarchy Level | 8454
- Description | 8455
- Default | 8455
- Required Privilege Level | 8456
- Release Information | 8456

Syntax

```
immediate-leave;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name],
[edit bridge-domains bridge-domain-name protocols igmp-snooping],
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-
name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
snooping],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
snooping interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols vlan
vlan-id igmp-snooping interface interface-name],
[edit protocols igmp-snooping vlan (all | vlan-name)],
[edit protocols igmp-snooping vlan (all | vlan-name) interface interface-name],
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name],
[edit protocols mld-snooping vlan (all | vlan-name)],
[edit protocols mld-snooping vlan vlan-name interface interface-name],
```

```
[edit routing-instances instance-name protocols mld-snooping vlan vlan-name interface interface-
name]
```

Description

Enable host tracking to allow the device to track the hosts that send membership reports, determine when the last host sends a leave message for the multicast group, and immediately stop forwarding traffic for the multicast group after the last host leaves the group. This setting helps to minimize IGMP or MLD membership leave latency—it reduces the amount of time it takes for the switch to stop sending multicast traffic to an interface when the last host leaves the group.

> **NOTE**: EVPN-VXLAN multicast uses special IGMP group leave processing to handle multihomed sources and receivers, so we don't support the `immediate-leave` option in EVPN-VXLAN networks.

IGMPv2, IGMPv3, MLDv1, and MLDv2 all have immediate leave disabled by default. In this state, the device does not track host memberships. When the device receives a leave report from a host, it sends out a group-specific query to all hosts. If no receiver responds with a membership report within a set interval, the device removes all hosts on the interface from the multicast group and stops forwarding multicast traffic to the interface.

With immediate leave enabled, the device removes an interface from the forwarding-table entry immediately without first sending IGMP group-specific queries out of the interface and waiting for a response. The device prunes the interface from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are active simultaneously.

Immediate leave is supported for IGMPv2, IGMPv3, MLDv1 and MLDv2 on devices that support these protocols.

> **NOTE**: We recommend that you configure immediate leave with IGMPv2 and MLDv1 only when there is only one host on an interface. With IGMPv2 and MLDv1, only one host on a interface sends a membership report in response to a general query—any other interested hosts suppress their reports. Report suppression avoids a flood of reports for the same group, but it also interferes with host tracking because the device knows only about one interested host on the interface at any given time.

Default

Immediate leave is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

- *Specifying Immediate-Leave Host Removal for IGMP*
- *Configuring IGMP Snooping on Switches*
- *Example: Configuring IGMP Snooping on Switches*
- show igmp-snooping vlans | **20198**
- *Example: Configuring IGMP Snooping on SRX Series Devices*
- *IGMP Snooping Overview*
- *Specifying Immediate-Leave Host Removal for MLD*
- *Understanding MLD Snooping*
- *Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)*

immediate-leave (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | **8457**
- Hierarchy Level | **8457**
- Description | **8457**
- Required Privilege Level | **8457**
- Release Information | **8457**

Syntax

```
immediate-leave;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name],
```

Description

Enable the routing device to leave the multicast group immediately after the last host leaves the multicast group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Specifying Immediate-Leave Host Removal for IGMP

immediate-leave (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 8458
- Hierarchy Level | 8458
- Description | 8458
- Required Privilege Level | 8459
- Release Information | 8459

Syntax

```
immediate-leave;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Description

The immediate leave setting is useful for minimizing the leave latency of MLD memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows MLD to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending MLD group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the MLD leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both MLD version 1 and MLD version 2.

> **NOTE**: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Specifying Immediate-Leave Host Removal for MLD

immediate-update

IN THIS SECTION

- Syntax | 8460
- Hierarchy Level | 8460

- Description | 8460
- Required Privilege Level | 8460
- Release Information | 8460

Syntax

```
immediate-update;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router or switch to send an Acct-Update message to the RADIUS accounting server on receipt of a response (for example, an ACK or timeout) to the Acct-Start message.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

RADIUS Accounting for Subscriber Access

import (BGP)

IN THIS SECTION

- Syntax | 8461
- Hierarchy Level | 8461
- Description | 8462
- Options | 8462
- Required Privilege Level | 8462
- Release Information | 8463

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Apply one or more routing policies to routes being imported into the Junos OS routing table from BGP.

If you specify more than one policy, they are evaluated in the order specified, from left to right, and the first matching filter is applied to the route. If no match is found, BGP places into the routing table only those routes that were learned from BGP routing devices. The policy framework software evaluates the routing policies in a chain sequentially. If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during the evaluation of the remaining policies. For example, if the action specified in the first policy of a chain sets a route's metric to 500, this route matches the criterion of `metric 500` defined in the next policy.

It is also important to understand that in Junos OS, although an import policy (inbound route filter) might reject a route, not use it for traffic forwarding, and not include it in an advertisement to other peers, the router retains these routes as hidden routes. These hidden routes are not available for policy or routing purposes. However, they do occupy memory space on the router. A service provider filtering routes to control the amount of information being kept in memory and processed by a router might want the router to entirely drop the routes being rejected by the import policy.

Hidden routes can be viewed by using the `show route receive-protocol bgp` *`neighbor-address`* `hidden` command. The hidden routes can then be retained or dropped from the routing table by configuring the `keep all | none` statement at the `[edit protocols bgp]` or `[edit protocols bgp group` *`group-name`*`]` hierarchy level.

The rules of BGP route retention are as follows:

- By default, all routes learned from BGP are retained, except those where the AS path is looped. (The AS path includes the local AS.)
- By configuring the `keep all` statement, all routes learned from BGP are retained, even those with the local AS in the AS path.
- By configuring the `keep none` statement, all routes received are discarded. When this statement is configured and the inbound policy changes, Junos OS re-advertises all the routes advertised by the peer.

Options

`policy-names`—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Example: Configuring BGP Interactions with IGPs
- Configuring Routing Policies to Control BGP Route Advertisements
- Understanding Routing Policies
- export (Protocols BGP) | 6677

import (Protocols OSPF)

IN THIS SECTION

- Syntax | 8463
- Hierarchy Level | 8463
- Description | 8464
- Options | 8464
- Required Privilege Level | 8464
- Release Information | 8464

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Filter OSPF routes from being added to the routing table.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Understanding OSPF Routing Policy

Import and Export Policies for Network Summaries Overview

export (OSPF) | **6660**

export (OSPF) | **6660**

import (Routing Options)

IN THIS SECTION

- Syntax | 8465
- Hierarchy Level | 8465
- Description | 8465
- Options | 8465
- Required Privilege Level | 8466
- Release Information | 8466

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options resolution rib],
[edit logical-systems logical-system-name routing-options resolution rib],
[edit routing-instances routing-instance-name routing-options resolution rib],
[edit routing-options resolution rib]
```

Description

Specify one or more import policies to use for route resolution.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring Route Resolution on PE Routers

import (Classifiers)

IN THIS SECTION

- Syntax | 8466
- Hierarchy Level | 8466
- Description | 8467
- Options | 8467
- Required Privilege Level | 8467
- Release Information | 8467

Syntax

```
import (classifier-name | default);
```

Hierarchy Level

```
[edit class-of-service classifiers type classifier-name]
```

Description

Specify a default or previously defined classifier.

Options

classifier-name—Name of the classifier mapping configured at the `[edit class-of-service classifiers]` hierarchy level.

`default`—The default classifier mapping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic

import (CoS)

IN THIS SECTION

- Syntax | 8468
- Hierarchy Level | 8468
- Description | 8468
- Options | 8468
- Required Privilege Level | 8468
- Release Information | 8468

Syntax

```
import (import | default);
```

Hierarchy Level

```
[edit class-of-service classifiers (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name],
[edit class-of-service rewrite-rules (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name]
```

Description

Specify a default or previously defined classifier.

Options

import —Name of the classifier mapping configured at the `[edit class-of-service classifiers]` hierarchy level.

`default`—Default classifier mapping.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Unicast Classifiers

Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

Defining CoS Rewrite Rules

Understanding CoS Classifiers

Understanding CoS Classifiers

Understanding CoS Classifiers

Understanding CoS Rewrite Rules

Example: Configuring CoS on EX Series Switches

Defining CoS Classifiers (CLI Procedure)

Defining CoS Classifiers (J-Web Procedure)

Defining CoS Rewrite Rules (CLI Procedure)

Defining CoS Rewrite Rules (J-Web Procedure)

Understanding CoS Classifiers

Understanding CoS Rewrite Rules

import-labeled-routes (Routing Instances VPLS)

IN THIS SECTION

- Syntax | 8469
- Hierarchy Level | 8470
- Description | 8470
- Options | 8470
- Required Privilege Level | 8470
- Release Information | 8470

Syntax

```
import-labeled-routes [ routing-instance-name ]
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit routing-instances routing-instance-name protocols vpls ]
```

Description

(MX Series routers with MPCs and MIC interfaces only) Specify one or more routing instances where MPLS paths (pseudowire labeled routes) are leaked from the mpls.0 routing table in the primary routing instance to the local *routing-instance-name*.mpls.0 routing table.

This capability is useful in an L2VPN/VPLS configuration when the remote PE router is learned from the IGP in a nondefault routing instance, because L2VPN/VPLS installs ingress-labeled routes only in the primary mpls.0 table.

When L2VPN/VPLS traffic is received on the core-facing interface in a nondefault routing instance, the router performs a lookup in the table that corresponds to that interface, *routing-instance-name*.mpls.0. Because the routes are not leaked by default, no routes are found in the *routing-instance-name*.mpls.0 routing table and all the incoming traffic is dropped.

Options

routing-instance-name	Name of a nondefault routing instance where you want routes leaked from the primary MPLS routing table.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring VPLS Routing Instances

import (MPLS Traffic Engineering Database)

IN THIS SECTION

- Syntax | 8471
- Hierarchy Level | 8471
- Description | 8471
- Options | 8472
- Required Privilege Level | 8472
- Release Information | 8472

Syntax

```
import {
    bgp-ls-identifier domain-identifier;
    identifier identifier;
    ipv6;
    l3-unicast-topology {
    bgp-link-state;
    }
    policy policy-name;
    igp-topology{
        bgp-link-state;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls traffic-engineering database],
[edit protocols mpls traffic-engineering database]
```

Description

Configure the traffic engineering database import parameters.

Options

`bgp-ls-identifier` *`domain-identifier`*	BGP-TE domain identifier.
identifier *identifier*	BGP-TE identifier. • **Range:** 2 through 18446744073709551615
`ipv6`	Enable import of IPv6 addresses into the traffic engineering database (TED).
`l3-unicast-topology`	Download L3-Unicast topology into RIB.
`policy` *`policy-name`*	Name of the import policy.
`igp-topology`	Download IGP topology information into the traffic engineering database (TED). In Junos OS, the IGPs install topology information into a database called the traffic engineering database. The traffic engineering database contains the aggregated topology information. The IGP routes are installed by the traffic engineering database on behalf of the corresponding IGP into a user-visible routing table called lsdist.0, subject to route policies.
`bgp-link-state`	Export IGP topology information into BGP-Link State (BGP-LS) from the lsdist.0 routing table. The lsdist.0 routing table stores the network topology information from the traffic engineering database. The BGP-LS reads IGP entries from lsdist.0 and advertises the information to BGP peers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

`ipv6` option introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

traffic-engineering (Protocols MPLS) | **30031**

import-policy

IN THIS SECTION

- Syntax | 8473
- Hierarchy Level | 8473
- Description | 8473
- Options | 8474
- Required Privilege Level | 8474
- Release Information | 8474

Syntax

```
import-policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib-groups group-name],
[edit logical-systems logical-system-name routing-options rib-groups group-name],
[edit routing-instances routing-instance-name routing-options rib-groups group-name],
[edit routing-options rib-groups group-name]
```

Description

Apply one or more policies to routes imported into the routing table group. The `import-policy` statement complements the `import-rib` statement and cannot be used unless you first specify the routing tables to which routes are being imported.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table

export-rib | 6688

passive (Routing Options) | 12225

import (Protocols DVMRP)

IN THIS SECTION

- Syntax | 8475
- Hierarchy Level | 8475
- Description | 8475
- Options | 8475
- Required Privilege Level | 8475
- Release Information | 8475

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols dvmrp],
[edit protocols dvmrp]
```

Description

Apply one or more policies to routes being imported into the routing table from DVMRP. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found, DVMRP shares with the routing table only those routes that were learned from DVMRP routers.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

> **NOTE**: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

export (Protocols DVMRP) | 6678

Example: Configuring DVMRP to Announce Unicast Routes

import (Protocols IS-IS)

IN THIS SECTION

- Syntax | 8476
- Hierarchy Level | 8476
- Description | 8476
- Options | 8477
- Required Privilege Level | 8477
- Release Information | 8477

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Apply one or more routing policies to routes being imported into the Junos OS routing table from IS-IS. You can import a routing policy to filter or restrict routes.

Options

policy-names—Name of one or more routing policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Example: Configuring a Routing Policy to Prioritize IS-IS Routes

Understanding Routing Policies

import (Protocols LDP)

IN THIS SECTION

- Syntax | 8477
- Hierarchy Level | 8478
- Description | 8478
- Options | 8478
- Required Privilege Level | 8478
- Release Information | 8478

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.

Options

policy-names—Name of one or more routing policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Filtering Inbound LDP Label Bindings

import (Protocols MSDP)

IN THIS SECTION

Syntax | 8479

- Hierarchy Level | 8479
- Description | 8480
- Options | 8480
- Required Privilege Level | 8480
- Release Information | 8480

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name],
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
peer address],
[edit protocols msdp],
[edit protocols msdp group group-name],
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit routing-instances routing-instance-name protocols msdp],
[edit routing-instances routing-instance-name protocols msdp group group-name],
[edit routing-instances routing-instance-name protocols msdp group group-name peer address],
[edit routing-instances routing-instance-name protocols msdp peer address]
```

Description

Apply one or more policies to routes being imported into the routing table from MSDP.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP in a Routing Instance

export (Protocols MSDP) | **6681**

import (Protocols PIM)

IN THIS SECTION

- Syntax | **8481**
- Hierarchy Level | **8481**
- Description | **8481**
- Options | **8481**
- Required Privilege Level | **8481**
- Release Information | **8481**

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Apply one or more policies to routes being imported into the routing table from PIM. Use the `import` statement to filter PIM join messages and prevent them from entering the network.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Filtering Incoming PIM Join Messages

import (Protocols PIM Bootstrap)

IN THIS SECTION

- Syntax | 8482
- Hierarchy Level | 8482
- Description | 8482
- Options | 8482
- Required Privilege Level | 8483
- Release Information | 8483

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp bootstrap (inet | inet6)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp bootstrap (inet | inet6)],
[edit protocols pim rp bootstrap (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim rp bootstrap (inet | inet6)]
```

Description

Apply one or more import policies to control incoming PIM bootstrap messages.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring PIM Bootstrap Properties for IPv4

Configuring PIM Bootstrap Properties for IPv4 or IPv6

export (Bootstrap) | 6663

import (Protocols RIP)

IN THIS SECTION

- Syntax | 8483
- Hierarchy Level | 8484
- Description | 8484
- Options | 8484
- Required Privilege Level | 8484
- Release Information | 8484

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip],
[edit protocols rip group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

Apply one or more policies to routes being imported by the local routing device from neighbors.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Applying Policies to RIP Routes Imported from Neighbors

Junos OS Routing Policies, Firewall Filters, and Traffic Policers User Guide for Routing Devices

export (RIP) | **6661**

import (Protocols RIPng)

IN THIS SECTION

- Syntax | 8485
- Hierarchy Level | 8485
- Description | 8485
- Options | 8486
- Required Privilege Level | 8486
- Release Information | 8486

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng],
[edit logical-systems logical-system-name protocols ripng group group-name neighbor neighbor-
name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng group group-name neighbor neighbor-name],
[edit protocols ripng],
[edit protocols ripng group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols ripng],
[edit routing-instances routing-instance-name protocols ripng group group-name neighbor neighbor-
name]
```

Description

Apply one or more policies to routes being imported into the local routing device from its neighbors.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Applying Policies to RIPng Routes Imported from Neighbors

export (Protocols RIPng) | **6686**

import (Rewrite Rules)

IN THIS SECTION

- Syntax | **8487**
- Hierarchy Level | **8487**
- Description | **8487**
- Options | **8487**
- Required Privilege Level | **8487**
- Release Information | **8487**

Syntax

```
import (rewrite-name | default);
```

Hierarchy Level

```
[edit class-of-service rewrite-rules type rewrite-name]
```

Description

Specify a default or previously defined `rewrite-rules` mapping to import.

Options

`rewrite-name`—Name of a rewrite-rules mapping configured at the `[edit class-of-service rewrite-rules]` hierarchy level.

`default`—The default `rewrite-rules` mapping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Rewrite Rules

import-rib

IN THIS SECTION

- Syntax | 8488
- Hierarchy Level | 8488
- Description | 8488
- Options | 8489
- Required Privilege Level | 8489
- Release Information | 8489

Syntax

```
import-rib [ routing-table-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib-groups group-name],

[edit logical-systems logical-system-name routing-options rib-groups group-name],
[edit routing-instances routing-instance-name routing-options rib-groups group-name],
[edit routing-options rib-groups group-name]
```

Description

Specify the name of the routing table into which Junos OS should import routing information. The first routing table name you enter is the primary routing table. Any additional names you enter identify secondary routing tables. When a protocol imports routes, it imports them into the primary and any secondary routing tables. If the primary route is deleted, the secondary route also is deleted. For IPv4 import routing tables, the primary routing table must be **inet.0** or **routing-instance-name.inet.0**. For IPv6 import routing tables, the primary routing table must be **inet6.0**.

In Junos OS Release 9.5 and later, you can configure an IPv4 import routing table that includes both IPv4 and IPv6 routing tables. Including both types of routing tables permits you, for example, to populate an IPv6 routing table with IPv6 addresses that are compatible with IPv4. In releases prior to Junos OS Release 9.5, you could configure an import routing table with only either IPv4 or IPv6 routing tables.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Options

routing-table-names—Name of one or more routing tables.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table

export-rib | **6688**

passive (Routing Options) | **12225**

import-target

IN THIS SECTION

- Syntax | **8490**
- Hierarchy Level | **8490**

- Description | 8490
- Options | 8490
- Required Privilege Level | 8491
- Release Information | 8491

Syntax

```
import-target {
    target {
        target-value;
        receiver target-value;
        sender target-value;
    }
    unicast {
        receiver;
        sender;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvpn
route-target],
[edit routing-instances routing-instance-name protocols mvpn route-target]
```

Description

Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN NLRI.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

imsi-prefix

IN THIS SECTION

- Syntax | 8491
- Hierarchy Level | 8492
- Description | 8492
- Options | 8492
- Required Privilege Level | 8492
- Release Information | 8492

Syntax

```
imsi-prefix imsi-prefix-digits {
    action (APN GTP) {
        drop;
        pass;
        selection (ms|net|vrf);
    }
}
```

Hierarchy Level

```
[edit security gtp profile profile-name apn pattern-string]
```

Description

Use the imsi-prefix option to specify an International Mobile Station Identity (IMSI) prefix for filtering GTP packets. You can also filter GTP packets based on the combination of an IMSI prefix and an access point name (APN).

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name apn pattern-string]` hierarchy level is replaced by `[edit security gtp profile profile-name apn pattern-string]`.

Options

imsi-prefix-digits—To specify an IMSI prefix.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. Support.

RELATED DOCUMENTATION

Understanding GTP APN Filtering

Understanding GTPv2 IMSI Prefix and APN Filtering

Example: Setting a GTP APN and a Selection Mode

in-place-lsp-bandwidth-update (Protocols MPLS)

IN THIS SECTION

- Syntax | 8493
- Hierarchy Level | 8493
- Description | 8493
- Default | 8494
- Required Privilege Level | 8494
- Release Information | 8494

Syntax

```
in-place-lsp-bandwidth-update;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Enable the ingress label edge routers (LER) to re-use the LSP-ID when performing bandwidth change on an intra-domain LSP. the LSP-ID re-use mechanism described in this document will not be applicable for inter-domain LSPs

NOTE: `commit` fails if you have configured either of the following:

- "no-cspf" on page 11429 at the `[edit protocols mpls]` hierarchy.
- "no-cspf" on page 11429 at the `[edit protocols mpls label-switched-path (Protocols MPLS) lsp-name]` hierarchy.

- p2mp (Protocols MPLS) is configured under the same [edit protocols mpls label-switched-path (Protocols MPLS) lsp-name] hierarchy.
- "lsp-external-controller" on page 10187 at the [edit protocols mpls] hierarchy.

Default

in-place-lsp-bandwidth-update is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

in-service

IN THIS SECTION

- Syntax | 8495
- Junos OS Hierarchy Level | 8495
- Junos OS Evolved Hierarchy Level | 8495
- Description | 8495
- Default | 8495
- Required Privilege Level | 8495
- Release Information | 8495

Syntax

```
in-service;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Runs the test in the `in-service` mode. In this mode, while the test is running, the rest of the data traffic sent to and from the UNI port under test on the service are not interrupted. Control protocol packets and control protocol peering are not interrupted.

If this mode is not configured, the test runs in the default `out-of-service` mode. In the `out-of-service` mode, while the test is running, all the data traffic sent to and from the UNI port under test on the service is interrupted. Control protocol peering is not interrupted whereas control protocol packets such as CFM sessions are interrupted.

Default

The default service mode for the reflecting egress interface for an E-LAN service is `out-of-service` mode.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

rfc2544-benchmarking | 15759

rfc2544 | 15757

RFC 2544-Based Benchmarking Tests Overview

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

Configuring an RFC 2544-Based Benchmarking Test

inactive-media-timeout (Security MGCP)

IN THIS SECTION

- Syntax | 8496
- Hierarchy Level | 8496
- Description | 8497
- Options | 8497
- Required Privilege Level | 8497
- Release Information | 8497

Syntax

```
inactive-media-timeout value-in-seconds;
```

Hierarchy Level

```
[edit logical-systems name security alg mgcp],
[edit logical-systems name tenants name security alg mgcp],
[edit security alg mgcp],
[edit services alg mgcp],
[edit tenants name security alg mgcp]
```

Description

Configure inactive-media-timeout parameter to specify the maximum amount of time that the temporary openings in the firewall (pinholes) remain open for media if no activity is detected.

Options

value-in-seconds—Maximum amount of time that the pinholes remain open.

- **Range:** 10 through 2550 seconds
- **Default:** 120 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

mgcp | 10693

inactive-media-timeout (Security SCCP)

IN THIS SECTION

- Syntax | 8498
- Hierarchy Level | 8498
- Description | 8498
- Options | 8498
- Required Privilege Level | 8498

Release Information | 8498

Syntax

```
inactive-media-timeout value-in-seconds;
```

Hierarchy Level

```
[edit logical-systems name security alg sccp],
[edit logical-systems name tenants name security alg sccp],
[edit security alg sccp],
[edit services alg sccp],
[edit tenants name security alg sccp]
```

Description

Configure inactive-media-timeout parameter to specify the maximum amount of time that the temporary openings in the firewall (pinholes) remain open for media if no activity is detected.

Options

value-in-seconds—Maximum amount of time that the pinholes remain open.

- **Range:** 10 through 600 seconds
- **Default:** 120 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sccp | 16341

inactive-media-timeout (Security SIP)

IN THIS SECTION

- Syntax | 8499
- Hierarchy Level | 8499
- Description | 8499
- Options | 8500
- Required Privilege Level | 8500
- Release Information | 8500

Syntax

```
inactive-media-timeout value-in-seconds;
```

Hierarchy Level

```
[edit logical-systems name security alg sip (Security)],
[edit logical-systems name tenants name security alg sip (Security)],
[edit security alg sip (Security)],
[edit services alg sip (Security)],
[edit tenants name security alg sip (Security)]
```

Description

Configure inactive-media-timeout parameter to specify the maximum amount of time that the temporary openings in the firewall (pinholes) remain open for media if no activity is detected.

Options

value-in-seconds—Maximum amount of time that the pinholes remain open.

- **Range:** 0 through 2550 seconds
- **Default:** 120 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sip (Security) | 27977

inactivity-asymm-tcp-timeout (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 8501
- Hierarchy Level | 8501
- Description | 8501
- Required Privilege Level | 8501
- Release Information | 8501

Syntax

```
inactivity-asymm-tcp-timeout seconds;
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options tcp-session]
```

Description

Configure the number of seconds that a unidirectional TCP session can be inactive before it is closed. Valid settings: 4 through 86400 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

inactivity-interval (Fabric Load Balance)

IN THIS SECTION

- Syntax | 8502
- Hierarchy Level | 8502
- Description | 8502
- Default | 8502
- Options | 8503
- Required Privilege Level | 8503

Release Information | 8503

Syntax

```
inactivity-interval interval;
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key fabric-load-balance flowlet]
```

Description

Set the inactivity interval for adaptive load balancing (ALB) using flowlets within a VCF.

NOTE: Starting in Junos OS Release 14.1X53-D46, 15.1R7, 16.1R6, 17.1R3, 17.2R2, 17.3R2, and 17.4R1, the ALB feature is deprecated. If `fabric-load-balance` is enabled in the configuration and `inactivity-interval` is set, delete these configuration items upon upgrading Junos OS.

The inactivity interval is the amount of time that occurs between packet bursts on a sending interface before a traffic flow is spliced into smaller traffic flows—flowlets—when ALB is implemented using flowlets. The flowlets are then individually forwarded across the VCF to the same destination device over different paths.

You should configure the inactivity interval to ensure in-order packet delivery, so that overall performance is not negatively impacted by the packet re-ordering process at the receiving device. To ensure in-order packet delivery, the inactivity interval should be larger than the largest latency skew among all the paths in the VCF from any node to any other node.

Default

ALB is disabled, by default.

If ALB is enabled without specifying a mode, ALB is enabled using flowlet mode with an inactivity interval of 16 microseconds.

If ALB is enabled using flowlet mode without specifying an inactivity interval, the inactivity interval is set to 16 microseconds.

Options

interval The amount of time that occurs between packet bursts on a sending interface before a traffic flow is spliced into flowets.

- **Range:** 16 microseconds (`16us`) to 32 milliseconds(`32ms`).

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

Statement deprecated starting in Junos OS Releases 14.1X53-D46, 15.1R7, 16.1R6, 17.1R3, 17.2R2, 17.3R2, and 17.4R1.

RELATED DOCUMENTATION

Understanding Traffic Flow Through a Virtual Chassis Fabric

inactivity-probe-duration

IN THIS SECTION

- Syntax | 8504
- Hierarchy Level | 8504
- Description | 8504
- Options | 8504
- Required Privilege Level | 8504

Release Information | 8504

Syntax

```
inactivity-probe-duration milliseconds;
```

Hierarchy Level

```
[edit protocols ovsdb controller]
```

Description

Configure the maximum amount of time, in milliseconds, that the connection between a Juniper Networks device that supports the *Open vSwitch Database (OVSDB)* management protocol and a software-defined networking (SDN) controller can be inactive before an inactivity probe is sent.

Options

milliseconds—Number of milliseconds that the connection can be inactive before an inactivity probe is sent.

- **Range:** 0 through 4,294,967,295
- **Default:** 0. This value indicates that an inactivity probe is never sent.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Understanding How to Set Up OVSDB Connections on a Juniper Networks Device

inactivity-timeout

IN THIS SECTION

- Syntax | 8505
- Hierarchy Level | 8505
- Description | 8505
- Options | 8505
- Required Privilege Level | 8506
- Release Information | 8506

Syntax

```
inactivity-timeout ( seconds | never );
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Configure the inactivity timeout period, in seconds.

Options

seconds—Length of time the application is inactive before it times out.

- **Default:** 30 seconds

never—Disables the inactivity-timeout period. Use this option if you do not want an application to time out.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

ALG Descriptions

Configuring Application Sets

Configuring the Inactivity Timeout Period

Examples: Configuring Application Protocols

Verifying the Output of ALG Sessions

inactivity-timeout (services)

IN THIS SECTION

- Syntax | 8507
- Hierarchy Level | 8507
- Description | 8507
- Options | 8507
- Required Privilege Level | 8507
- Release Information | 8507

Syntax

```
inactivity-timeout seconds;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
[edit services service-set-name service-set-options]
```

Description

Configure the inactivity timeout period for established flows. The timeout value configured in the application protocol definition overrides this value.

Options

seconds—Timeout period.

- **Default:** 30 seconds
- **Range:** 4 through 86,400 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for MX-SPC3 services card on MX240, MX480 and MX960 routers.

RELATED DOCUMENTATION

> *Configuring Default Timeout Settings for Services Interfaces*

inactivity-timeout (Applications)

IN THIS SECTION

- Syntax | 8508
- Hierarchy Level | 8508
- Description | 8508
- Options | 8508
- Required Privilege Level | 8509
- Release Information | 8509

Syntax

```
inactivity-timeout (seconds | never) ;
```

Hierarchy Level

```
[edit applications application  application-name  ]
[edit applications application  application-name  term  term-name]
```

Description

Application-specific inactivity timeout period, in seconds.

Options

seconds—Specify the amount of time the application is inactive before it times out in seconds.

Range: 4 through 86400 seconds.

Default: For TCP, 1800 seconds; for UDP, 60 seconds.

never—Disables the inactivity-timeout period. If you do not want an application to time out, use the `never` option to disable the inactivity-timeout period.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policy Application Sets Overview

inactivity-timeout (Security Persistent NAT)

IN THIS SECTION

- Syntax | 8509
- Hierarchy Level | 8510
- Description | 8510
- Options | 8510
- Required Privilege Level | 8510
- Release Information | 8510

Syntax

```
inactivity-timeout seconds;
```

Hierarchy Level

```
[edit security nat source rule-set ruleset rule rule then source-nat interface persistent-nat]
[edit security nat source rule-set ruleset rule rule then source-nat pool persistent-nat]
```

Description

The amount of time, in seconds, that the persistent NAT binding remains in the Juniper Networks device's memory when all the sessions of the binding entry are gone. When the configured timeout is reached, the binding is removed from memory.

Options

seconds—Number of seconds.

- **Range:** 60 through 7200 seconds
- **Default:** 300 seconds (5 minutes)

Required Privilege Level

`security`—To view this statement in the configuration.

`security-control`—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 9.6.

inactivity-timeout (Services RPM)

IN THIS SECTION

- Syntax | 8511
- Hierarchy Level | 8511
- Description | 8511

- Options | 8511
- Required Privilege Level | 8511
- Release Information | 8511

Syntax

```
inactivity-timeout seconds;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Inactivity timeout period, in seconds.

Options

seconds—Length of time the session is inactive before it times out.

- **Default:** 1800 seconds

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

inband-flow-telemetry

IN THIS SECTION

- Syntax | 8512
- Hierarchy Level | 8513
- Description | 8513
- Options | 8513
- Required Privilege Level | 8515
- Release Information | 8516

Syntax

```
inband-flow-telemetry {
   clock-source (ntp|ptp);
   device-id {
      (id-number|auto);
   }
   flow-type (l3|vxlan);
   hop-limit value;
   meta-data-stack-length value;
   no-ipv6-address-match;
   profile {
      ifa-profile-name {
         sample-rate value;
         collector {
            source-address ipv4-address;
            destination-address ipv4-address;
            destination-port port-number;
            maximum-clip-length length;
            mtu size;
```

```
        }
      }
    }
  }
```

Hierarchy Level

```
[edit services]
```

Description

Configure Inband Flow Analyzer 2.0 (IFA 2.0). IFA 2.0 provides insights about complex networks by collecting per-hop flow data on the data plane. IFA uses probe packets to collect network-wide flow data. IFA generates probe packets by sampling the traffic flow of interest. IFA probe packets are representative packets of the original flow and possess the exact same characteristics as the original flow. This means that IFA packets traverse the same path in the network and the same queues in the networking element as the original packet would. Because the IFA probe packets traverse the same network path as the original flow, the packets experience similar latency and congestion.

IFA uses the following processing nodes to monitor and analyze flows:

- IFA initiator node (also known as ingress node)
- IFA transit node
- IFA terminating node (also known as egress node)

Use the `inband-flow-telemetry` configuration options to configure the IFA nodes.

Options

clock-source (ntp|ptp)

Configure the clock source protocol to enable more accurate timestamping. The QFX5120-48YM switch supports PTP as well as NTP; all of the other QFX5120 switch models support NTP only.

Default: ntp

`device-id (`*`id-number`*`| auto)`

(Mandatory for all IFA nodes) Specify a unique device identifier for each hop within an IFA zone. You must configure this value for all three IFA node types: IFA initiator, IFA transit, and IFA terminating. If you configure `auto`, then the device ID is internally generated from the router ID or the management IP address.

Range: 1-1,048,575

flow-type (l3| vxlan)

(Mandatory for IFA initiator node and terminating node and optional for transit node) Specify the IFA flow type—`l3` or `vxlan`. If you configure the flow type as `vxlan`, and the incoming traffic is L3, or vice versa, then the IFA nodes do not behave as expected. This is because you cannot initiate IFA probe packets with invalid fields.

You cannot configure both L3 or VXLAN flows on the same device. This restriction is applicable for the IFA initiator and terminating nodes (generally leaf nodes).

You don't need to configure `flow-type` for a transit node (generally spine nodes).

Default: `l3`

hop-limit *value*

(Optional) Configure the maximum allowed hops in an IFA zone. The initiator node initializes this field. The hop limit is decremented at each hop. If the hop limit of the incoming packets is 0, the current node does not insert the metadata.

You can avoid the metadata insertion at the transit node by using the `hop-limit` configuration.

The IFA terminating node does not perform a hop-limit check. Even if the incoming IFA packet has hop-limit set to 0, the IFA terminating node inserts the metadata and reduces the hop limit by 1. In this case, the `hop-limit` value resets to 255. The `hop-limit` option cannot have a negative value.

- **Range**: 1-250
- **Default**: 250

meta-data-stack-length *value*

(Optional) Configure the maximum allowed length of the metadata stack in multiples of four octets. The initiator node initializes this field. Each node in the path compares the current length with the maximum allowed length. If the current length equals or exceeds the maximum length, the transit node must stop inserting the metadata.

- **Range**: 8-255
- **Default**: 240 (for 30 hops)

no-ipv6-address-match

(Optional) Optimize IFA filter group processing by removing the IPv6 source and destination address match qualifiers from the IFA filter group. IFA cannot be initiated or terminated with these two qualifiers, but you can initiate or terminate IFA with the remaining qualifiers.

Default: off (this statement is not included in the IFA configuration unless you specifically configure it)

profile *ifa-*

IFA profile
name.

profile-name

sample-rate *value*

Configure the average number of samples obtained in one second. For example, if you configure the sample rate as 1000, then out of 1000 packets one packet is sampled per second. You cannot have different sample rates for different flows on an IFA initiator node enabled on a port. All flows within a port must have the same sample rate.

Range: 1-16,777,215

collector

Configure a collector for IFA 2.0 probe packets. The monitored packets are exported to the collector in IPFIX format. By default, Junos OS supports a maximum packet length of 256 bytes starting with the Ethernet header. An IFA IPFIX packet contains IFA headers (8 bytes), IFA metadata (variable length), and the originally monitored packet (256 bytes).

Configure the following collector-related options:

- `source-address` *`ipv4-address`*—IPv4 source address.
- `destination-address` *`ipv4-address`*—IPv4 destination address.
- `destination-port` *`port-number`*—Destination port value.
  Range: 1-65,535.
- `maximum-clip-length` *`length`*—Number of bytes of the original flow packet that should be exported in the IPFIX packet. Because the maximum MTU is 9000 bytes at the IFA termination node, the maximum clip length for the IPFIX packet is equal to or less than: 9000 bytes - (IFA header length + IFA metadata header length + IFA metadata stack length).

  Range: 64 to 9000 bytes

  Default: 256 bytes
- `mtu` *`size`*—Size in bytes of the maximum transmission unit for IPFIX packets leaving the IFA termination node. Any packet exceeding 9000 bytes in length is dropped.

  Range: 256 to 9000 bytes

  Default: 9000 bytes

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

`clock-source`, `maximum-clip-length`, `mtu`, and `no-ipv6-address-match` options introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

- inband-flow-telemetry | 8512
- clear inband-flow-telemetry stats | 3635
- show services inband-flow-telemetry | 25574
- *Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Performance Monitoring*

inband-ztp

IN THIS SECTION

- Syntax | 8516
- Hierarchy Level | 8517
- Description | 8517
- In-band ZTP | 8517
- Options | 8517
- Required Privilege Level | 8517
- Release Information | 8518

Syntax

```
inband-ztp {
    management-vlan-id management-vlan-id;
    name-server [ name-server ... ];
```

```
        }
    }
```

Hierarchy Level

```
[edit system services]
```

Description

Enable in-band ZTP.

In-band ZTP enables your already-provisioned Day 1 upstream devices in your IP fabric to monitor for newly connected Day 0 downstream devices. In-band ZTP invokes an autonomous Python daemon, which uses Link-Level Discovery Protocol (LLDP) telemetry data, to monitor for any new downstream interfaces. When a Day 1 upstream device detects a downstream interface, the Day 1 upstream device can provide Layer 3 connectivity to the Day 0 downstream device. With Layer 3 connectivity, the Day 0 device can then proceed with ZTP.

In-band ZTP

Options

management-vlan-id Provide the management VLAN ID. The management VLAN enables in-band ZTP for a particular VLAN. An upstream device can belong to multiple VLANs, and the management VLAN enables in-band ZTP to monitor for interfaces in those VLANs.

name-server Provide a list of DNS server addresses.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

None.

Release Information

Statement introduced in Junos OS Release 24.2R1

RELATED DOCUMENTATION

In-band Zero Touch Provisioning

include-all (for Administrative Groups)

IN THIS SECTION

- Syntax | 8518
- Hierarchy Level | 8518
- Description | 8519
- Options | 8519
- Required Privilege Level | 8519
- Release Information | 8519

Syntax

```
include-all [ group-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name admin-
group],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name admin-group],
[edit protocols mpls label-switched-path lsp-name admin-group],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name admin-group]
```

Description

Require the LSP to traverse links that include all of the defined administrative groups.

Options

group-names—One or more names of groups defined with the `admin-groups` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Administrative Groups for LSPs

admin-groups | 1973

include-all (for Fast Reroute)

IN THIS SECTION

- Syntax | 8520
- Hierarchy Level | 8520
- Description | 8520
- Options | 8520
- Required Privilege Level | 8520
- Release Information | 8520

Syntax

```
(include-all [ group-names ] | no-include-all);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-
reroute],
[edit protocols mpls label-switched-path lsp-name fast-reroute]
```

Description

Control inclusion of administrative groups:

- `include-all`—Define the administrative groups that must all be included for fast reroute.
- `no-include-all`—Disable administrative group inclusion.

Options

group-names—One or more names of groups defined with the `admin-groups` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Fast Reroute

include-any (for Administrative Groups)

IN THIS SECTION

- Syntax | 8521
- Hierarchy Level | 8521
- Description | 8521
- Options | 8521
- Required Privilege Level | 8522
- Release Information | 8522

Syntax

```
include-any [ group-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name admin-
group],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
 secondary) path-name admin-group],
[edit protocols mpls label-switched-path lsp-name admin-group],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name admin-group]
```

Description

Define the administrative groups to include for an LSP or for a path's primary and secondary paths.

Options

group-names—One or more names of groups defined with the `admin-groups` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Administrative Groups for LSPs

include-any (for Fast Reroute)

IN THIS SECTION

- Syntax | 8522
- Hierarchy Level | 8523
- Description | 8523
- Options | 8523
- Required Privilege Level | 8523
- Release Information | 8523

Syntax

```
(include-any [ group-names ] | no-include-any);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-
reroute],
[edit protocols mpls label-switched-path lsp-name fast-reroute]
```

Description

Control inclusion of administrative groups:

- `include-any`—Define the administrative groups to include for fast reroute.
- `no-include-any`—Disable administrative group inclusion.

Options

`group-names`—One or more names of groups defined with the `admin-groups` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Fast Reroute

include (Diameter Gx Profiles)

IN THIS SECTION

- Syntax | 8524
- Hierarchy Level | 8524
- Description | 8524
- Options | 8524
- Required Privilege Level | 8525
- Release Information | 8525

Syntax

```
include {
    gx-capability-list;
    rule-suggestion;
}
```

Hierarchy Level

```
[edit unified-edge diameter-profiles gx-profile profile-name attributes]
```

Description

Configure the attribute-value pairs (AVPs) to be included in the credit control request (CCR) messages between the MX Series router and the policy and charging enforcement function (PCEF).

Options

`gx-capability-list`—Include the Gx-Capability list AVP.

`rule-suggestion`—Include the Rule-suggestion AVP.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

gx-profile | 7954

include (Dynamic Access-Line-Identifier VLANs)

IN THIS SECTION

- Syntax | 8525
- Hierarchy Level | 8526
- Description | 8526
- Options | 8526
- Required Privilege Level | 8526
- Release Information | 8526

Syntax

```
include {
    accept-no-ids;
    circuit-id;
    remote-id;

}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces “$junos-interface-ifd-name” unit “$junos-
interface-unit”  auto-configure line-identity],
[edit interfaces interface-name unit logical-unit-number auto-configure line-identity]
```

Description

Configure the access-line identifier received in DHCP or PPPoE discovery packets that is a trusted option. Trusted options are accepted for dynamically creating a VLAN on a static or dynamic underlying VLAN interface. These VLANs are known as access-line-identifier-based VLANs.

Options

accept-no-ids (Optional) Enables creation of a VLAN in the absence of the ACI and the ARI string in the received DHCP or PPPoE packet. This VLAN serves an a default VLAN to collect all subscribers for which no sub-option is received.

circuit-id (Optional) Enables creation of the VLAN when the ACI string is received in a DHCP or PPPoE packet.

remote-id (Optional) Enables creation of the VLAN when the ARI string is received in a DHCP or PPPoE packet.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 17.1.

RELATED DOCUMENTATION

Configuring Dynamic Underlying VLAN Interfaces to Use Access-Line Identifiers

Configuring Dynamic VLAN Subscriber Interfaces Based on Access-Line Identifiers

include-igp-metric

IN THIS SECTION

- Syntax | 8527
- Hierarchy Level | 8527
- Description | 8527
- Default | 8527
- Required Privilege Level | 8527
- Release Information | 8528

Syntax

```
include-igp-metric;
```

Hierarchy Level

```
[edit protocols rsvp route]
```

Description

RSVP adds the actual IGP metric of the installed routes to the existing metrics in the RSVP LSP route created with `conditional-metric`.

Default

`include-igp-metric` is disabled by default. If you do not include the `include-igp-metric` statement, RSVP continues to use the currently configured metrics in which the RSVP route will be installed with metrics it is currently using; static, conditional or by default IGP tracking metric.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.1R1

RELATED DOCUMENTATION

Configuring LSP Metrics

include-ipv6 (Gx-Plus)

IN THIS SECTION

- Syntax | 8528
- Hierarchy Level | 8528
- Description | 8529
- Default | 8529
- Required Privilege Level | 8529
- Release Information | 8529

Syntax

```
include-ipv6;
```

Hierarchy Level

```
[edit access gx-plus global]
```

Description

Include IPv6 subscribers in Gx-Plus provisioning requests.

Default

By default, IPv6 subscribers are not included.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring Gx-Plus Global Attributes

Configuring Gx-Plus

include-irb-and-l2

IN THIS SECTION

- Syntax | 8530
- Hierarchy Level | 8530
- Description | 8530
- Required Privilege Level | 8532
- Release Information | 8532

Syntax

```
include-irb-and-l2;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name (relay-agent-interface-id | relay-agent-
remote-id)],
[edit forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)],
[edit forwarding-options dhcp-relay group group-name relay-option-82 (circuit-id | remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay dhcpv6 (relay-agent-
interface-id | relay-agent-remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay ... relay-option-82
(circuit-id | remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6  (relay-agent-
interface-id | relay-agent-remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ... relay-option-82
(circuit-id | remote-id)],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-18],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-37]
```

Description

Include both the integrated routing and bridging (IRB) interface name and Layer 2 interface name in the `circuit-id` or `remote-id` value in the DHCP option 82 information. VLAN tags are global.

For leasequery and bulk leasequery operations that involve integrated routing and bridging (IRB) interfaces, you must configure DHCP relay agent to include the layer 2 interface name along with IRB name in the circuit ID of option 82. DHCP relay agent uses the layer 2 interface name when using leasequery or bulk leasequery to restore the lease database.

When you configure the `include-irb-and-l2` statement without including the `no-vlan-interface` statement, the format is as follows:

- Bridge domain:

```
(fe | ge)-fpc/pic/port.subunit:bridge-domain-name+irb.subunit
```

- VLAN:

```
(fe | ge)-fpc/pic/port.subunit:vlan-name+irb.subunit
```

NOTE: For remote systems, the *subunit* is required and is used to differentiate an interface.

When you configure both the `include-irb-and-l2` statement and the `use-vlan-id` statement, the format is as follows:

```
(fe | ge)-fpc/pic/port.subunit:svlan-id-vlan-id+irb.subunit
```

NOTE: The *svlan-id-vlan-id* represents the VLANs associated with the bridge domain.

When you configure both the `include-irb-and-l2` and `no-vlan-interface-name` statements, the format is as follows:

```
(fe | ge)-fpc/pic/port.subunit+irb.subunit
```

When you configure both the `include-irb-and-l2` and `use-interface-description` statements, the format displays the description for the Layer 2 interface:

```
l2_descr:vlan-name+irb.subunit
```

If you configure both the `include-irb-and-l2` and `use-interface-description` statements, and no description for the Layer 2 interface is found, the format displays the Layer 2 logical interface name:

```
(fe | ge)-fpc/pic/port.subunit:vlan-name+irb.subunit
```

When you configure the `include-irb-and-l2` statement with both the `no-vlan-interface-name` and `use-interface-description` statements, the format displays as follows:

```
l2_descr+irb.subunit
```

If you configure the `include-irb-and-l2` statement with both the `no-vlan-interface-name` and `use-interface-description` statements, and no description is found for the Layer 2 interface, the format displays as follows:

```
(fe | ge)-fpc/pic/port.subunit+irb.subunit
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

> **NOTE**: The EX Series switches that support the `include-irb-and-l2` statement are the EX4300, EX4600, and EX9200 switches.

RELATED DOCUMENTATION

Including a Textual Description in DHCP Options

Using DHCP Relay Agent Option 82 Information

Configuring DHCPv6 Relay Agent Options

include-mp-next-hop

IN THIS SECTION

- Syntax | 8533
- Hierarchy Level | 8533
- Description | 8533
- Required Privilege Level | 8533

Release Information | 8533

Syntax

```
include-mp-next-hop;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Enable multiprotocol updates to contain next-hop reachability information.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IPv6 BGP Routes over IPv4 Transport

Enabling Layer 2 VPN and VPLS Signaling

Understanding Multiprotocol BGP

include-option-82 (DHCP Local Server)

IN THIS SECTION

- Syntax | 8534
- Hierarchy Level | 8534
- Description | 8535
- Options | 8535
- Required Privilege Level | 8535
- Release Information | 8535

Syntax

```
include-option-82 {
    forcerenew;
    nak;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server overrides],
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server (dhcpv6) group group-name overrides],
[edit system services dhcp-local-server (dhcpv6) group group-name interface interface-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server (dhcpv6) ...overrides],
```

```
[edit logical-systems logical-system-name system services dhcp-local-server
(dhcpv6) ...overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server
(dhcpv6) ...overrides]
```

Description

Specify that the DHCP server include option 82 information in NAK and forcerenew messages when you configure secure communications between the DHCP server and DHCP clients that are in different VRFs. You can configure support globally, for a group of interfaces, or for a specific interface.

Options

forcerenew Include option 82 in DHCP forcerenew messages.

nak Include option 82 in DHCP NAK messages.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 13.3R3, 14.1R2, and 14.2R1.

RELATED DOCUMENTATION

DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs

Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances

include-sci

IN THIS SECTION

- Syntax | 8536
- Hierarchy Level | 8536
- Description | 8536
- Default | 8537
- Required Privilege Level | 8537
- Release Information | 8537

Syntax

```
include-sci;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Specifies that the SCI tag should be appended to each packet on a link that has enabled MACsec.

You must enable SCI tagging on a switch that is enabling MACsec on an Ethernet link connecting to an EX4300 or EX4600 switch.

SCI tags are automatically appended to packets leaving a MACsec-enabled interface on EX4300 or EX4600 switches, so this option is not available on these switches.

You should only use this option when connecting a switch to an EX4300 or EX4600 switch, or to a host device that requires SCI tagging. SCI tags are eight octets long, so appending an SCI tag to all traffic on the link adds a significant amount of unneeded overhead.

Default

SCI tagging is enabled by default on SRX Series Firewalls, EX4300 and EX4600 switches that have enabled MACsec using static connectivity association key (CAK) security mode.

SCI tagging is disabled on all other interfaces, by default.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced for SRX Series Firewalls in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

include-sci (MACsec for MX Series)

IN THIS SECTION

- Syntax | **8538**
- Hierarchy Level | **8538**
- Description | **8538**
- Default | **8538**
- Required Privilege Level | **8538**
- Release Information | **8538**

Syntax

```
include-sci;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Specifies that the SCI tag should be appended to each packet on a link that has enabled MACsec.

This option is used only when connecting a router to a host device that requires SCI tagging. SCI tags are eight octets long, so appending an SCI tag to all traffic on the link adds a significant amount of unneeded overhead.

Default

SCI tagging is not enabled by default.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

inclusive

IN THIS SECTION

- Syntax | 8539
- Hierarchy Level | 8539
- Description | 8539
- Required Privilege Level | 8539
- Release Information | 8540

Syntax

```
inclusive;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvpn
family (inet | inet6) autodiscovery-only intra-as],
[edit routing-instances routing-instance-name protocols mvpn family (inet | inet6) autodiscovery-
only intra-as]
```

Description

For Rosen 7, enable the MVPN control plane for autodiscovery only, using intra-AS autodiscovery routes over an inclusive provider multicast service interface (PMSI).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Statement moved to `[..protocols mvpn family inet]` from `[.. protocols mvpn]` in Junos OS Release 13.3.

Support for IPv6 added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs

incoming-map

IN THIS SECTION

- Syntax | 8540
- Hierarchy Level | 8540
- Description | 8541
- Required Privilege Level | 8541
- Release Information | 8541

Syntax

```
incoming-map {
    caller caller-number | accept-all;
}
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with interfaces configured for ISDN, specify the dialer to accept incoming calls.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: The `incoming-map` statement is mandatory for the router to accept any incoming ISDN calls.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

incoming-queue

IN THIS SECTION

- Syntax | 8542
- Hierarchy Level | 8542
- Description | 8542
- Options | 8542
- Required Privilege Level | 8542
- Release Information | 8542

Syntax

```
incoming-queue {
    size size;
}
```

Hierarchy Level

```
[edit access diameter peer peer-name]
```

Description

Configure the incoming queue properties of this peer.

Options

size *size* Size of the queue. The default is 6000.

- **Range:** 1 through 65,535 packets

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

independent-domain

IN THIS SECTION

- Syntax | 8543
- Hierarchy Level | 8543
- Description | 8543
- Options | 8544
- Required Privilege Level | 8544
- Release Information | 8544

Syntax

```
independent-domain <no-attrset>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options autonomous-system autonomous-system],
[edit routing-instances routing-instance-name routing-options autonomous-system autonomous-
system]
```

Description

Configure an independent AS domain.

The independent domain uses transitive path attribute 128 (attribute set) messages to tunnel the independent domain's BGP attributes through the internal BGP (IBGP) core.

This improves the transparency of Layer 3 VPN services for customer networks by preventing the IBGP routes that originate within an autonomous system (AS) in the customer network from being sent to a service provider's AS. Similarly, IBGP routes that originate within an AS in the service provider's network are prevented from being sent to a customer AS.

NOTE: In Junos OS Release 10.3 and later, if BGP receives attribute 128 and you have not configured an independent domain in any routing instance, BGP treats the received attribute 128 as an unknown attribute.

NOTE: The `[edit logical-systems]` hierarchy level is not applicable in ACX Series routers.

Options

`no-attrset` (Optional) Disables attribute set messages on the independent AS domain.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`no-attrset` option introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Example: Tunneling Layer 3 VPN IPv6 Islands over an IPv4 Core Using IBGP and Independent Domains

Configuring Layer 3 VPNs to Carry IBGP Traffic

autonomous-system

indexed-load-balance

IN THIS SECTION

- Syntax | 8545
- Hierarchy Level | 8545
- Description | 8545
- Default | 8546
- Required Privilege Level | 8546
- Release Information | 8546

Syntax

```
indexed-load-balance;
```

Hierarchy Level

```
[edit forwarding-options load-balance],
[edit logical-systems logical-system-name forwarding-options load-balance],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options load-balance],
[edit routing-instances routing-instance-name forwarding-options load-balance]
```

Description

The `indexed-load-balance` statement causes the creation of a nexthop structure that is both a function of the hash, and a function of the low-order bits of the IP address. Include this statement if you notice issues with load-balance distribution for IPv4 or IPv6 traffic to improve load-balance distribution for unicast and aggregated next hops.

For MPC line cards in MX routers, `indexed-load-balance` has been superseded by an internal hash-rotation mechanism to reduce polarization.

CAUTION: Including the `indexed-load-balance` statement causes an increase in memory usage on the device.

Default

Disabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Starting with Junos OS Release 12.1, the indexed-next-hop statement has been renamed as the indexed-load-balance statement.

RELATED DOCUMENTATION

Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers

Configuring Per-Prefix Load Balancing

enhanced-hash-key | **6362**

indication

IN THIS SECTION

- Syntax | **8547**
- Hierarchy Level | **8547**
- Description | **8547**
- Options | **8547**

Required Privilege Level | 8547
Release Information | 8547

Syntax

```
indication (ignore | normal | require);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options dce-options],
[edit interfaces interface-name serial-options dte-options]
```

Description

For X.21 interfaces only, configure the from-DCE signal indication.

Options

`ignore`—The from-DCE signal is ignored.

`normal`—Normal indication signal handling as defined by ITU-T Recommendation X.21.

`require`—The from-DCE signal must be asserted.

- **Default:** `normal`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Signal Handling

indication-polarity

IN THIS SECTION

- Syntax | 8548
- Hierarchy Level | 8548
- Description | 8548
- Options | 8548
- Required Privilege Level | 8549
- Release Information | 8549

Syntax

```
indication-polarity (negative | positive);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

For X.21 interfaces only, configure the indication signal polarity.

Options

`positive`—Positive signal polarity.

`negative`—Negative signal polarity.

- **Default:** `positive`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Serial Signal Polarities

indirect-next-hop

IN THIS SECTION

- Syntax | 8549
- Hierarchy Level | 8550
- Description | 8550
- Default | 8550
- Options | 8550
- Required Privilege Level | 8550
- Release Information | 8551

Syntax

```
(indirect-next-hop | no-indirect-next-hop);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table],
[edit routing-options forwarding-table]
```

Description

Enable indirectly connected next hops for route convergence. This statement is implemented on the Packet Forward Engine to speed up forwarding information base (FIB) updates. Configuring this statement significantly speeds convergence times. The only downside of configuring this statement is that some additional FIB memory overhead is required. Unless routes have an extremely high number of next hops, this increased memory usage should not be noticeable.

NOTE:

- When virtual private LAN service (VPLS) is configured on the routing device, the `indirect-next-hop` statement is configurable at the `[edit routing-options forwarding-table]` hierarchy level. However, this configuration is not applicable to indirect nexthops specific to VPLS routing instances.
- By default, the Junos Trio Modular Port Concentrator (MPC) chipset on MX Series routers is enabled with indirectly connected next hops, and this cannot be disabled using the `no-indirect-next-hop` statement.
- By default, indirectly connected next hops are enabled on PTX Series routers.

Default

Disabled.

Options

indirect-next-hop—Enable indirectly connected next hops.

no-indirect-next-hop—Explicitly disable indirect next hops.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

> *Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine*

indirect-next-hop-change-acknowledgements

IN THIS SECTION

- Syntax | 8551
- Hierarchy Level | 8551
- Description | 8552
- Default | 8552
- Options | 8552
- Required Privilege Level | 8552
- Release Information | 8553

Syntax

```
(indirect-next-hop-change-acknowledgements | no-indirect-next-hop-change-acknowledgements);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table],
[edit routing-options forwarding-table]
```

Description

Configure the routing protocol process (rpd) to request an acknowledgement when creating a new forwarding next hop.

During an indirect next-hop change sequence, the routing device might create a new forwarding next hop that is referenced by the indirect next hop. If the `indirect-next-hop-change-acknowledgements` statement is configured, the routing protocol process requests an acknowledgement when creating the new forwarding next hop. When the routing protocol process receives the acknowledgement, this indicates that all PFEs have received the new forwarding next hop and it is then safe to change the indirect next hop to reference the new forwarding next hop. This prevents packet loss when changing the indirect next hop by ensuring that all PFEs have consistent state information for the new forwarding next hop.

The routing protocol process is not requesting an acknowledgement for the indirect next hop itself. Rather, the routing protocol process is requesting an acknowledgement for the new forwarding next hop that the indirect next hop is going to reference. In the case when the forwarding next hop is an existing one (meaning that it is already installed in the forwarding table), the routing protocol process does not request an acknowledgement, even if the `indirect-next-hop-change-acknowledgements` statement is configured.

We recommend that the `indirect-next-hop-change-acknowledgements` statement be configured when protection mechanisms are being used. This includes MPLS RSVP protection such as fast reroute (FRR) as well as interior gateway protocol (IGP) loop-free alternate (LFA) link or node protection. If there is no protection mechanism being used in the network, the `indirect-next-hop-change-acknowledgements` statement does not provide any benefit and might increase packet loss.

Default

Disabled by default in all platforms except PTX Series, where it is enabled by default.

Options

indirect-next-hop-change-acknowledgements—Enable acknowledgements.

no-indirect-next-hop-change-acknowledgements—Explicitly disable acknowledgements.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine

krt-nexthop-ack-timeout | 9506

ineligible

IN THIS SECTION

- Syntax | 8553
- Hierarchy Level | 8553
- Description | 8554
- Required Privilege Level | 8554
- Release Information | 8554

Syntax

```
ineligible;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery address address],
[edit protocols router-discovery address address]
```

Description

Specify that the address can never become the default router.

Required Privilege Level

routing—To view this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

priority (Protocols Router Discovery) | 13196

inet

IN THIS SECTION

- Syntax | 8554
- Hierarchy Level | 8555
- Description | 8555
- Required Privilege Level | 8555
- Release Information | 8555

Syntax

```
inet;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-next-hop ingress labeled-bgp],
[edit routing-options forwarding-table chained-composite-next-hop ingress labeled-bgp]
```

Description

Allows you to configure chained composite next hops for IPv4 labeled-unicast routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

This statement was introduced in Junos OS Release 22.2.

RELATED DOCUMENTATION

Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs

inet (AMT Protocol)

IN THIS SECTION

- Syntax | 8556
- Hierarchy Level | 8556
- Description | 8556
- Required Privilege Level | 8556
- Release Information | 8556

Syntax

```
inet {
    anycast-prefix ip-prefix/<prefix-length>;
    local-address ip-address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols amt relay family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols amt
relay family],
[edit protocols amt relay family],
[edit routing-instances routing-instance-name protocols amt relay family]
```

Description

Specify the IPv4 local address and anycast prefix for Automatic Multicast Tunneling (AMT) relay functions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

inet-bgpmcast

IN THIS SECTION

- Syntax | 8557
- Hierarchy Level | 8557
- Description | 8557
- Required Privilege Level | 8557
- Release Information | 8558

Syntax

```
inet-bgpmcast {
    signaling;
}
```

Hierarchy Level

```
[edit protocols bgp family],
[edit protocols bgp group group-name family],
[edit protocols bgp neighbor neighbor-address family]
```

Description

Enable the `inet-bgpmcast` address family to advertise BGP multicast NLRI.

IPv4 and IPv6 can be enabled or disabled separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1

RELATED DOCUMENTATION

Controller-Based BGP Multicast Signaling

bgpmcast | 2968

accept-prpd-connection | 1639

inet6-bgpmcast | 8599

show route table | 23789

show multicast route | 22306

inet (chassis)

IN THIS SECTION

- Syntax | 8558
- Hierarchy Level | 8559
- Description | 8559
- Options | 8559
- Required Privilege Level | 8559
- Release Information | 8559

Syntax

```
inet {
    layer-3;
    layer-4;
    symmetric-hash {
        complement;
```

```
    }
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number hash-key family]
```

Description

(MX Series 5G Universal Routing Platforms only) Configure data used in a hash key for the **inet** protocol family when configuring PIC-level symmetrical load balancing on an 802.3ad Link Aggregation Group.

Options

layer-3	Include Layer 3 IP data in the hash key.
layer-4	Include Layer 4 IP data in the hash key.
symmetric-hash	Configure symmetric hash key with source and destination ports.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

> *Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers*

inet (CoS)

IN THIS SECTION

- Syntax | 8560
- Hierarchy Level | 8560
- Description | 8560
- Required Privilege Level | 8560
- Release Information | 8561

Syntax

```
inet {
    classifiers {
        (dscp | inet-precedence) classifier-name ;
    }
}
```

Hierarchy Level

```
[edit class-of-service multi-destination family]
```

Description

Specify the IP multicast family.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Option **inet** introduced in Junos OS Release 9.5 for EX Series switches.

The remaining statements are explained separately. See CLI Explorer.

RELATED DOCUMENTATION

Understanding CoS Schedulers

Understanding CoS Forwarding Classes

Understanding CoS Classifiers

inet (enhanced-hash-key)

IN THIS SECTION

- Syntax (EX Series and QFX5100 Switch) | **8561**
- Syntax (QFX10000 Series Switches) | **8562**
- Hierarchy Level | **8562**
- Description | **8562**
- Default | **8562**
- Options | **8563**
- Required Privilege Level | **8563**
- Release Information | **8563**

Syntax (EX Series and QFX5100 Switch)

```
inet {
    gtp-tunnel-endpoint-identifier;
    no-ipv4-destination-address;
    no-ipv4-source-address;
    no-l4-destination-port;
    no-l4-source-port;
```

```
    no-protocol;
    vlan-id;
}
```

Syntax (QFX10000 Series Switches)

```
inet {
    gtp-tunnel-endpoint-identifier;
    no-ipv4-destination-address;
    no-ipv4-source-address;
    no-l4-destination-port;
    no-l4-source-port;
    no-incoming-port;
}
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key family]
```

Description

Select the payload fields in IPv4 traffic used by the hashing algorithm to make hashing decisions.

When IPv4 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using the `inet` statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or how to forward traffic to the next hop device when ECMP is enabled.

The hashing algorithm, when used to hash LAG bundle traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

The hashing algorithm only inspects the IPv4 fields in the payload to make hashing decisions when the hash mode is set to `layer2-payload`. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the `set forwarding-options enhanced-hash-key hash-mode layer2-payload` statement.

Default

The following fields are used by the hashing algorithm to make hashing decisions for IPv4 traffic:

- IP destination address
- IP source address
- Layer 4 destination port
- Layer 4 source port
- Protocol

Options

`no-ipv4-destination-address`	Exclude the IPv4 destination address field from the hashing algorithm.
`no-ipv4-source-address`	Exclude the IPv4 source address field from the hashing algorithm.
`no-l4-destination-port`	Exclude the Layer 4 destination port field from the hashing algorithm.
`no-l4-source-port`	Exclude the Layer 4 source port field from the hashing algorithm.
`no-protocol`	Exclude the protocol field from the hashing algorithm.
`no-incoming-port`	Exclude the incoming port number from the hashing algorithm.
`vlan-id`	Include the VLAN ID field in the hashing algorithm.

> **NOTE**: The `vlan-id` option is not supported and should not be configured on a Virtual Chassis or Virtual Chassis Fabric (VCF) that contains any of the following switches as members: EX4300, EX4600, QFX3500, QFX3600, QFX5100, or QFX5110 switches.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D15.

RELATED DOCUMENTATION

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)

Understanding Per-Packet Load Balancing

hash-seed | **8002**

enhanced-hash-key (EX) | **6378**

hash-mode | **7997**

inet (enhanced-hash-key) | **8561**

inet (forwarding options)

IN THIS SECTION

- Syntax | **8564**
- Hierarchy Level | **8565**
- Description | **8565**
- Options | **8565**
- Required Privilege Level | **8565**
- Release Information | **8565**

Syntax

```
inet {
    mode(drop | flow-based | packet-based);
}
```

Hierarchy Level

```
[edit security forwarding-options family]
```

Description

The default mode for inet is flow mode.

Family IPv4

Options

mode Forwarding mode

- drop—Disable forwarding
- flow-based—Enable flow-based forwarding
- packet-based—Enable packet-based forwarding

Required Privilege Level

[none specified]

Release Information

Statement supported in Junos 24.2R1 for SRX Series Firewalls

inet (interfaces (QFX Series))

IN THIS SECTION

- Syntax | 8566
- Hierarchy Level | 8566

- Description | 8566
- Default | 8567
- Options | 8568
- Required Privilege Level | 8568
- Release Information | 8568

Syntax

```
inet {
    address address {
        primary;
    filter input filter-name;
    filter output filter-name;
    targeted-broadcast;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family],
[edit interfaces interface-range interface-name unit logical-unit-number family]
```

Description

Configure the primary IP address for the logical interface.

Junos will no longer accept duplicate IPs between different logical interfaces in the same routing instance. It is applicable to both inet and inet6. The below table lists the duplicate addresses found between different interfaces.

Table 49: Allowed Duplicate IP Address Between Different Interfaces

	Point to Point	Management Interfaces	Non-point to point	With same APS	With no or diff APS	LSQ or PSEUDO RLSQ interfaces (fail options)	LSQ o PSEU SQ interf (fail o
Point to Point	Commit accepted	Error displayed	Error displayed	Error displayed	Error displayed	Commit accepted	Comm accep
Management Interfaces	Error displayed	Error displayed	Error displayed	Error displayed	Error displayed	Error displayed	Error displa
Non-point to point	Error displayed	Error displayed	Error displayed	Error displayed	Error displayed	Error displayed	Error displa
With same APS	Error displayed	Error displayed	Error displayed	Commit accepted	Error displayed	Error displayed	Error displa
With no or diff APS	Error displayed	Error displayed	Error displayed	Error displayed	Error displayed	Error displayed	Error displa
LSQ or PSEUDO_RLSQ interfaces (fail options)	Commit accepted	Error displayed	Error displayed	Error displayed	Error displayed	Commit accepted	Comm accep
LSQ or PSEUDO_RLSQ interfaces (no fail options)	Commit accepted	Error displayed	Error displayed	Error displayed	Error displayed	Commit accepted	Error displa

APS configuration is applicable to Sonet (SO) and derivatives of Atm (AT) interfaces only.

Default

You must configure a logical interface to be able to use the physical device.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

inet (Interfaces)

IN THIS SECTION

- Syntax | 8568
- Hierarchy Level | 8569
- Description | 8569
- Options | 8569
- Required Privilege Level | 8570
- Release Information | 8570

Syntax

```
inet {
    address source;
    auto-configure {
```

```
        address-ranges {
            authentication {
                password password-string;
                username-include {
                    auth-server-realm realm-string;
                    delimiter delimiter-character;
                    domain-name domain-name;
                    interface-name;
                    source-address;
                    user-prefix user-prefix-string;
                }
            }
            dynamic-profile profile-name {
                network ip-address {
                    range name {
                        low lower-limit;
                        high upper-limit;
                    }
                }
            }
        }

    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux]
```

Description

Specify the inet family for the demultiplexing (demux) interface options.

Options

address-source *address*

Specify the source IPv4 or IPv6 address or prefix value from which to inherit configuration data for the demultiplexing (demux) interface options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles

inet-mdt

IN THIS SECTION

- Syntax | 8570
- Hierarchy Level | 8571
- Description | 8571
- Required Privilege Level | 8571
- Release Information | 8571

Syntax

```
inet-mdt;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
mvpn family (inet | inet6) autodiscovery],
[edit routing-instances routing-instance-name protocols pim mvpn family (inet | inet6)
autodiscovery]
```

Description

For Rosen 7, configure the PE router in a VPN to use an SSM multicast distribution tree (MDT) subsequent address family identifier (SAFI) NLRI .

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Statement moved to `[..protocols pim mvpn family inet]` from `[.. protocols mvpn]` in Junos OS Release 13.3.

Support for IPv6 added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

> *Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs*

inet-mdt (Signaling)

IN THIS SECTION

- Syntax | 8572
- Hierarchy Level | 8573

- Description | 8573
- Required Privilege Level | 8573
- Release Information | 8573

Syntax

```
signaling {
    accepted-prefix-limit {
        maximum number;
        teardown <percentage-threshold> idle-timeout (forever | minutes);
        drop-excess <percentage>;
        hide-excess <percentage>;
     }
    add-path {
        receive;
        send {
            include-backup-path backup_path_number;
            multipath;
            path-count number;
            path-selection-mode {
                (all-paths | equal-cost-paths);
            }
            prefix-policy [ policy-names ];
        }
    }
    aigp [disable];
    loops number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
        drop-excess <percentage>;
        hide-excess <percentage>;
    }
    rib-group group-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family],
[edit logical-systems logical-system-name protocols bgp group group-name family],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
family],
[edit protocols bgp family],
[edit protocols bgp group group-name family],
[edit protocols bgp group group-name neighbor address family],
[edit protocols bgp family prefix-limit (inet | inet6)(any | flow | labeled-unicast | multicast |
unicast)],
[edit routing-instances instance-name protocols bgp family],
[edit routing-instances instance-name protocols bgp group group-name family],
[edit routing-instances instance-name protocols bgp group group-name neighbor address family]
```

Description

For draft-rosen 7, on the provider edge router enable BGP intra-AS auto-discovery using MDT-SAFI.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs

inet-mvpn (BGP)

IN THIS SECTION

- Syntax | 8574
- Hierarchy Level | 8574
- Description | 8575
- Required Privilege Level | 8575
- Release Information | 8575

Syntax

```
inet-mvpn {
    signaling {
        accepted-prefix-limit {
            maximum number;
            teardown percentage {
                idle-timeout (forever | minutes);
            }
        }
        damping;
        loops number;
        prefix-limit {
            maximum number;
            teardown percentage {
                idle-timeout (forever | minutes);
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family],
[edit protocols bgp family],
```

```
[edit logical-systems logical-system-name protocols bgp group group-name family],
[edit protocols bgp group group-name family]
```

Description

Enable the `inet-mvpn` address family in BGP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

inet-mvpn (VRF Advertisement)

IN THIS SECTION

- Syntax | 8575
- Hierarchy Level | 8576
- Description | 8576
- Required Privilege Level | 8576
- Release Information | 8576

Syntax

```
inet-mvpn;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name vrf-advertise-selective family],
[edit routing-instances routing-instance-name vrf-advertise-selective family]
```

Description

Enable IPv4 MVPN routes to be advertised from the VRF instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Limiting Routes to Be Advertised by an MVPN VRF Instance

inet (Port Mirroring)

IN THIS SECTION

- Syntax | 8577
- Hierarchy Level | 8577
- Description | 8577
- Required Privilege Level | 8577
- Release Information | 8578

Syntax

```
inet {
output {
    ip-addressaddress {
    }
    routing-instance instance-name {
        ip-address address {
    }
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring [instance name] family]
```

Description

Specify that the output interface will be of type `inet`. Use this statement so that you can send the mirrored packets to the IP address of the device running an analyzer application. The device can be on a remote network. When you use this feature, the mirrored packets are GRE-encapsulated. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)

NOTE: An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.

NOTE: If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53.

inet-precedence

IN THIS SECTION

- Syntax | 8578
- Hierarchy Level | 8578
- Description | 8579
- Options | 8579
- Required Privilege Level | 8579
- Release Information | 8579

Syntax

```
inet-precedence classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [ aliases ] [ 6-bit-patterns ];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service classifiers],
[edit class-of-service code-point-aliases],
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers],
[edit class-of-service rewrite-rules]
```

Description

Apply an IPv4 precedence rewrite rule.

Options

classifier-name—Name of the classifier.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- *Example: Configuring CoS on EX Series Switches*
- *Defining CoS Classifiers (CLI Procedure)*
- *Defining CoS Classifiers (J-Web Procedure)*
- *Defining CoS Code-Point Aliases (CLI Procedure)*
- *Defining CoS Code-Point Aliases (J-Web Procedure)*
- *Defining CoS Rewrite Rules (CLI Procedure)*
- *Defining CoS Rewrite Rules (J-Web Procedure)*
- *Understanding CoS Classifiers*
- *Understanding CoS Rewrite Rules*

inet-precedence (Classifier on Physical Interface)

IN THIS SECTION

- Syntax | 8580
- Hierarchy Level | 8580
- Description | 8580
- Options | 8580
- Required Privilege Level | 8581
- Release Information | 8581

Syntax

```
inet-precedence (classifier-name | default);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name  classifiers]
```

Description

On ACX Series Universal Metro routers and EX Series switches, map the inet-precedence field of the incoming packet to the forwarding class and packet loss priority, based on the specified inet-precedence classifier. When no classifier is configured on the physical interface, the default ipprec-compatibility inet-precedence classifier is applied on the physical interface.

Options

classifier-name	Name of the previously defined inet-precedence behavior aggregate classifier.
default	Default inet-precedence behavior aggregate classifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

inet-precedence (CoS Classifiers)

IN THIS SECTION

- Syntax | 8581
- Hierarchy Level | 8581
- Description | 8582
- Options | 8582
- Required Privilege Level | 8582
- Release Information | 8582

Syntax

```
inet-precedence (classifier-name | default) {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level code-points [ aliases ] [ bit-patterns ]
        }
    }
```

Hierarchy Level

```
[edit class-of-service classifiers]
```

Description

Define an IPv4 classifier.

Options

`default`—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic

inet-precedence (CoS Rewrite Rules)

IN THIS SECTION

- Syntax | 8583
- Hierarchy Level | 8583
- Description | 8583
- Options | 8583
- Required Privilege Level | 8583
- Release Information | 8584

Syntax

```
inet-precedence (rewrite-name | default) protocol (inet-both | inet-outer | mpls);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

Apply an IPv4 precedence rewrite rule.

Options

rewrite-name Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules inet-precedence]` hierarchy level.

default The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.

protocol inet-both Rewrite the inet-precedence CoS value to both the inner and outer header for Unicast/Multicast IPv4 traffic. The first three bits of the CoS value are rewritten and the final six bits are taken from the incoming CoS value.

protocol inet-outer Rewrite the inet-precedence CoS value to the outer header for Unicast/Multicast IPv4 traffic. The first three bits of the CoS value are rewritten and the final six bits are taken from the incoming CoS value.

protocol mpls (Optional for ingress MPLS tunnel nodes) Rewrite the MPLS EXP bits in the MPLS header independently of the IPv4 DSCP value for IPv4 packets entering an MPLS tunnel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Rewrite Rules

Applying Rewrite Rules to Output Logical Interfaces

protocol (Rewrite Rules) | **13457**

rewrite-rules (Definition) | **15746**

inet-precedence (Dynamic Classifiers)

IN THIS SECTION

- Syntax | **8584**
- Hierarchy Level | **8584**
- Description | **8585**
- Options | **8585**
- Required Privilege Level | **8585**
- Release Information | **8585**

Syntax

```
inet-precedence (classifier-name | default);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-number classifiers]
```

Description

Apply a IPv4 precedence classifier to a subscriber interface in a dynamic profile.

Options

classifier-name—Name of a classifier mapping configured at the [edit class-of-service classifiers inet-precedence] hierarchy level.

default—The default mapping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying a Classifier to a Subscriber Interface in a Dynamic Profile

classifiers (Definition)

inet-precedence (Dynamic Rewrite Rules)

IN THIS SECTION

- Syntax | 8586
- Hierarchy Level | 8586
- Description | 8586
- Options | 8586
- Required Privilege Level | 8586

Release Information | 8586

Syntax

```
inet-precedence (rewrite-name | default);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

Apply a IPv4 precedence rewrite rule.

Options

rewrite-name—Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules inet-precedence]` hierarchy level.

`default`—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile

rewrite-rules

inet-precedence (Rewrite Rules on Physical Interface)

IN THIS SECTION

- Syntax | 8587
- Hierarchy Level | 8587
- Description | 8587
- Options | 8588
- Required Privilege Level | 8588
- Release Information | 8588

Syntax

```
inet-precedence (rewrite-name | default);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name  rewrite-rules]
```

Description

Apply a IPv4 precedence rewrite rule.

Options

rewrite-name—Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules inet-precedence]` hierarchy level.

`default`—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

inet (Routing Policy Condition)

IN THIS SECTION

- Syntax | 8588
- Hierarchy Level | 8589
- Description | 8589
- Options | 8589
- Required Privilege Level | 8589
- Release Information | 8589

Syntax

```
inet {
    address;
```

```
    table table-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options condition if-route-exists address-
family],
[edit policy-options condition if-route-exists address-family]
```

Description

Specify that the route must correspond to a IPv4 prefix to be considered a match.

Options

(Optional) *address*—Specify the IP address that the route must have to be considered a match.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring Pseudowire Redundancy in a Mobile Backhaul Scenario

inet (Security Static NAT)

IN THIS SECTION

- Syntax | 8590
- Hierarchy Level | 8590
- Description | 8590
- Options | 8590
- Required Privilege Level | 8591
- Release Information | 8591

Syntax

```
inet {
    routing-instance (routing-instance-name| default);
}
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name then static-nat]
```

Description

Specify the automatic translation of IPv6 addresses to IPv4 addresses (and vice versa).

NOTE: If you use this option, you do not need to use the *prefix* option because with this option, the first 96 most significant bits are automatically stripped from the 128-bit IPv6 address.

Options

- `routing-instance` *routing-instance-name* —Use the user-defined static NAT routing-instance to perform static NAT.

- `default`—Use the default routing-instance to perform static NAT. When a `routing-instance-name` is not provided, the default routing-instance primary is used, which refers to the main inet.0 (for IPv4 unicast routes) routing table or inet.6 (for IPv6 unicast routes) routing table.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

inet (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 8591
- Hierarchy Level | 8592
- Description | 8592
- Required Privilege Level | 8592
- Release Information | 8592

Syntax

```
inet {
    drop-policy rule-name {
        from {
            apply-groups group-name;
            apply-groups-except group-name;
            destination-address address;
            destination-port port-number;
            dscp dscp-value;
            protocol protocol;
```

```
            source-address address;
            source-port port-number;
        }
    }
}
```

Hierarchy Level

```
[edit services radius-flow-tap policy policy-name]
```

Description

Specify the inet family for the policy that is applied to mirrored packets sent to a mediation device.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Subscriber Secure Policy Overview

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

inet (TDF Subscriber Address)

IN THIS SECTION

- Syntax | 8593
- Hierarchy Level | 8593
- Description | 8593
- Required Privilege Level | 8593
- Release Information | 8594

Syntax

```
inet {
    apply-groups [group-names];
    apply-groups-except [group-names];
    pool pool-name;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber subscriber-
address]
```

Description

Specify IP version 4 (IPv4) for the address pool that contains the source IP addresses for IP-based subscriber packets that undergo TDF processing.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

inet (TDF Subscriber Exclude Prefix)

IN THIS SECTION

- Syntax | 8594
- Hierarchy Level | 8594
- Description | 8595
- Required Privilege Level | 8595
- Release Information | 8595

Syntax

```
inet {
    network address mask;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name subscriber-exclude-prefix
family]
```

Description

Specify IP version 4 (IPv4) for the network prefix of source IP addresses for uplink packets and destination IP addresses for downlink packets that do not undergo TDF processing.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based and IFL-Based TDF Subscribers Overview

IP-Based Subscriber Setup Overview

inet6

IN THIS SECTION

- Syntax | 8596
- Hierarchy Level | 8596
- Description | 8596
- Required Privilege Level | 8596
- Release Information | 8596

Syntax

```
inet6;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-
next-hop ingress labeled-bgp],
[edit routing-options forwarding-table chained-composite-next-hop ingress labeled-bgp]
```

Description

Allows you to configure chained composite next hops for IPv6 labeled-unicast routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

This statement was introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs

inet6-advertise-interval (VRRP)

IN THIS SECTION

Syntax | 8597

- Hierarchy Level | 8597
- Description | 8597
- Options | 8597
- Required Privilege Level | 8598
- Release Information | 8598

Syntax

```
inet6-advertise-interval milliseconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
```

Description

Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv6 advertisement packets.

All routers in the VRRP group must use the same advertisement interval.

> **NOTE**: When VRRPv3 is enabled, the `inet6-advertise-interval` statement cannot be used to configure advertisement intervals. Instead, use the `fast-interval` statement to configure advertisement intervals.

Options

milliseconds—Interval, in milliseconds, between advertisement packets.

- **Range:** 100 to 40,000 milliseconds (ms)
- **Default:** 1 second

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4R2.

RELATED DOCUMENTATION

No title for referred element: topic-map

Configuring VRRP for IPv6 (CLI Procedure)

version-3 | **30691**

inet6-backup-router

IN THIS SECTION

- Syntax | **8598**
- Hierarchy Level | **8599**
- Description | **8599**
- Options | **8599**
- Required Privilege Level | **8599**
- Release Information | **8599**

Syntax

```
inet6-backup-router address <destination destination-address>;
```

Hierarchy Level

```
[edit system]
```

Description

Set a default router (running IP version 6 [IPv6]) to use while the local router or switch (running IPv6) is booting and if the routing protocol processes fail to start. The Junos OS removes the route to this router or switch as soon as the software starts.

Options

address Address of the default router.

destination *destination-address* (Optional) Destination address that is reachable through the backup router. You can include this option to achieve network reachability while loading, configuring, and recovering the router or switch, but without the risk of installing a default route in the forwarding table.

- **Default:** All hosts (default route) are reachable through the backup router.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

inet6-bgpmcast

IN THIS SECTION

Syntax | 8600

- Hierarchy Level | 8600
- Description | 8600
- Required Privilege Level | 8600
- Release Information | 8600

Syntax

```
inet6-bgpmcast {
    signaling;
}
```

Hierarchy Level

```
[edit protocols bgp family],
[edit protocols bgp group group-name family],
[edit protocols bgp neighbor neighbor-address family]
```

Description

Enable the `inet6-bgpmcast` address family to advertise BGP multicast NLRI.

IPv4 and IPv6 can be enabled or disabled separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1

RELATED DOCUMENTATION

Controller-Based BGP Multicast Signaling

bgpmcast | 2968

accept-prpd-connection | 1639

inet-bgpmcast | 8557

show route table | 23789

show multicast route | 22306

inet6 (CoS Multidestination)

IN THIS SECTION

- Syntax | 8601
- Hierarchy Level | 8601
- Description | 8602
- Required Privilege Level | 8602
- Release Information | 8602

Syntax

```
inet6 {
   classifiers {
       dscp-ipv6 classifier-name;
   }
}
```

Hierarchy Level

```
[edit class-of-service multi-destination family]
```

Description

Specify the IPv6 multicast family.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Option **inet6** introduced in Junos OS Release before Junos OS 11.4 for EX Series switches.

RELATED DOCUMENTATION

Understanding CoS Schedulers

Understanding CoS Forwarding Classes

Understanding CoS Classifiers

inet6 (enhanced-hash-key)

IN THIS SECTION

- Syntax (EX Series and QFX5100 Switch) | **8603**
- Syntax (QFX10000 Series Switches) | **8603**
- Hierarchy Level | **8603**
- Description | **8603**
- Default | **8604**
- Options | **8604**
- Required Privilege Level | **8604**
- Release Information | **8605**

Syntax (EX Series and QFX5100 Switch)

```
inet6 {
    no-ipv6-destination-address;
    no-ipv6-source-address;
    no-l4-destination-port;
    no-l4-source-port;
    no-next-header;
    vlan-id;
}
```

Syntax (QFX10000 Series Switches)

```
inet6 {
    gtp-tunnel-endpoint-identifier;
    ipv6-flow-label;
    no-ipv6-destination-address;
    no-ipv6-source-address;
    no-l4-destination-port;
    no-l4-source-port;
    no-incoming-port;
}
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key family]
```

Description

Select the payload fields in an IPv6 packet used by the hashing algorithm to make hashing decisions.

When IPv6 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or to forward traffic to the next hop device when ECMP is enabled.

The hashing algorithm, when used to hash LAG traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

The hashing algorithm only inspects the IPv6 fields in the payload to make hashing decisions when the hash mode is set to Layer 2 payload. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the `set forwarding-options enhanced-hash-key hash-mode layer2-payload` statement.

Default

The data in the following fields are used by the hashing algorithm to make hashing decisions for IPv6 traffic:

- IP destination address
- IP source address
- Layer 4 destination port
- Layer 4 source port
- Next header

Options

no-ipv6-destination-address	Exclude the IPv6 destination address field from the hashing algorithm.
no-ipv6-source-address	Exclude the IPv6 source address field from the hashing algorithm.
no-l4-destination-port	Exclude the Layer 4 destination port field from the hashing algorithm.
no-l4-source-port	Exclude the Layer 4 source port field from the hashing algorithm.
no-incoming-port	Exclude the incoming port number from the hashing algorithm.
no-next-header	Exclude the Next Header field from the hashing algorithm.
vlan-id	Include the VLAN ID field in the hashing algorithm.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D15.

RELATED DOCUMENTATION

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)

Understanding Per-Packet Load Balancing

hash-seed | **8002**

enhanced-hash-key (EX) | **6378**

hash-mode | **7997**

inet (enhanced-hash-key) | **8561**

inet6 (forwarding optons)

IN THIS SECTION

- Syntax | **8605**
- Hierarchy Level | **8606**
- Description | **8606**
- Options | **8606**
- Required Privilege Level | **8606**
- Release Information | **8606**

Syntax

```
inet6 {
```

```
            mode(drop | flow-based | packet-based);
        }
```

Hierarchy Level

```
[edit security forwarding-options family]
```

Description

The default mode for inet6 is flow mode.

Family IPv6

Options

mode Forwarding mode

- drop—Disable forwarding
- flow-based—Enable flow-based forwarding
- packet-based—Enable packet-based forwarding

Required Privilege Level

[none specified]

Release Information

Statement supported in Junos 24.2R1 for SRX Series Firewalls

inet6 (interfaces (QFX Series))

IN THIS SECTION

- Syntax | 8607
- Hierarchy Level | 8607
- Description | 8607
- Default | 8608
- Options | 8608
- Required Privilege Level | 8608
- Release Information | 8608

Syntax

```
inet6 {
    address address {
        eui-64
            preferred
        primary;
    filter input filter-name;
    filter output filter-name;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family],
[edit interfaces interface-range interface-name unit logical-unit-number family]
```

Description

Configure the primary IP address for the logical interface.

Default

You must configure a logical interface to be able to use the physical device.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches

inet6 (Interfaces)

IN THIS SECTION

- Syntax | 8609
- Hierarchy Level | 8609
- Description | 8609
- Options | 8610
- Required Privilege Level | 8610
- Release Information | 8610

Syntax

```
inet6 {
    address source;
    auto-configure {
        address-ranges {
            authentication {
                password password-string;
                username-include {
                    auth-server-realm realm-string;
                    delimiter delimiter-character;
                    domain-name domain-name;
                    interface-name;
                    source-address;
                    user-prefix user-prefix-string;
                }
            }
            dynamic-profile profile-name {
                network ip-address {
                    range name {
                        low lower-limit;
                        high upper-limit;
                    }
                }
            }
        }

    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux]
```

Description

Specify the inet6 family for the demultiplexing (demux) interface options.

Options

address-source *address* — Specify the source IPv4 or IPv6 address or prefix value from which to inherit configuration data for the demultiplexing (demux) interface options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscribers Services Overview

Configuring Packet Triggered Subscribers Using IP Demux Interfaces in Dynamic Profiles

inet6-mvpn (BGP)

IN THIS SECTION

- Syntax | 8611
- Hierarchy Level | 8611
- Description | 8611
- Required Privilege Level | 8611
- Release Information | 8612

Syntax

```
inet6-mvpn {
    signaling {
        accepted-prefix-limit {
            maximum number;
            teardown percentage {
                idle-timeout (forever | minutes);
            }
        }
        loops number
        prefix-limit {
            maximum number;
            teardown percentage {
                idle-timeout (forever | minutes);
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family],
[edit protocols bgp family],
[edit logical-systems logical-system-name protocols bgp group group-name family],
[edit protocols bgp group group-name family]
```

Description

Enable the `inet6-mvpn` address family in BGP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

BGP Configuration Overview

inet6-mvpn (VRF Advertisement)

IN THIS SECTION

- Syntax | 8612
- Hierarchy Level | 8612
- Description | 8612
- Required Privilege Level | 8613
- Release Information | 8613

Syntax

```
inet6-mvpn;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name vrf-advertise-selective family],
[edit routing-instances routing-instance-name vrf-advertise-selective family]
```

Description

Enable IPv6 MVPN routes to be advertised from the VRF instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

inet6-options (Services)

IN THIS SECTION

- Syntax | 8613
- Hierarchy Level | 8613
- Description | 8614
- Options | 8614
- Required Privilege Level | 8614
- Release Information | 8614

Syntax

```
inet6-options {
    source-address address;
}
```

Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Description

Specify the source IPv6 address used for probes. If the source IPv6 address is not one of the devices' assigned addresses, the packet uses the outgoing interface's address as its source.

Options

`inet6-options`—Use the specified base IPv6 protocol-related settings to be used for RPM probes

`source-address` *`ipv6-address`*—Specify the base IPv6 address for sending the RPM probes from the client to the server (for example, 2001:db8::a:b:c:d).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R4.

RELATED DOCUMENTATION

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

inet6-precedence (CoS Rewrite Rules)

IN THIS SECTION

- Syntax | 8615
- Hierarchy Level | 8615
- Description | 8615
- Options | 8615
- Required Privilege Level | 8615
- Release Information | 8615

Syntax

```
inet6-precedence rewrite-name {
    protocol mpls;
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
```

Description

Apply an IPv6 precedence rule to rewrite the first three bits of the IPv6 DSCP value.

Options

rewrite-name Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules inet6-precedence]` hierarchy level.

`protocol mpls` (Optional for ingress MPLS tunnel nodes) Rewrite the DSCP value for IPv6 packets entering the MPLS tunnel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Configuring Rewrite Rules

Applying Rewrite Rules to Output Logical Interfaces

protocol (Rewrite Rules) | **13457**

rewrite-rules (Definition) | 15746

inet6 (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 8616
- Hierarchy Level | 8616
- Description | 8617
- Required Privilege Level | 8617
- Release Information | 8617

Syntax

```
inet6 {
    drop-policy rule-name {
        from {
            apply-groups group-name;
            apply-groups-except group-name;
            destination-address address;
            destination-port port-number;
            dscp dscp-value;
            protocol protocol;
            source-address address;
            source-port port-number;
        }
    }
}
```

Hierarchy Level

```
[edit services radius-flow-tap policy policy-name]
```

Description

Specify the inet6 family for the policy that is applied to mirrored packets sent to a mediation device.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Subscriber Secure Policy Overview

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

inet6 (TDF Subscriber Address)

IN THIS SECTION

- Syntax | 8618
- Hierarchy Level | 8618
- Description | 8618
- Required Privilege Level | 8618
- Release Information | 8618

Syntax

```
inet6 {
    pool pool-name;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber subscriber-address]
```

Description

Specify IP version 6 (IPv6) for the address pool that contains the source IP addresses for IP-based subscriber packets that undergo TDF processing.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

inet6 (TDF Subscriber Exclude Prefix)

IN THIS SECTION

- Syntax | 8619
- Hierarchy Level | 8619
- Description | 8619
- Required Privilege Level | 8619
- Release Information | 8620

Syntax

```
inet6 {
    network address mask;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name subscriber-exclude-prefix
family]
```

Description

Specify IP version 6 (IPv6) for the network prefix of source IP addresses for uplink packets and destination IP addresses for downlink packets that do not undergo TDF processing.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based and IFL-Based TDF Subscribers Overview

IP-Based Subscriber Setup Overview

inet6-vpn

IN THIS SECTION

- Syntax | **8620**
- Hierarchy Level | **8621**
- Description | **8621**
- Options | **8621**
- Required Privilege Level | **8621**
- Release Information | **8621**

Syntax

```
inet6-vpn (any | multicast | unicast) {
    aggregate-label;
    prefix-limit maximum;
    rib-group rib-group-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family],
[edit logical-systems logical-system-name protocols bgp group group-name family],
[edit protocols bgp family],
[edit protocols bgp group group-name family]
```

Description

Enable IP version 6 (IPv6) on the provider edge (PE) router for the Layer 3 VPN.

NOTE: The `[edit logical-systems]` hierarchy level is not applicable in ACX Series routers.

Options

`any`—Configure the family type to be both multicast and unicast.

`multicast`—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.

`prefix-limit` *`maximum`*—Maximum prefix limit.

- **Range:** 1 through 4,294,967,295
- **Default:** 1

`rib-group` *`rib-group-name`*—The name of the routing table group.

`unicast`—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Layer 3 VPNs to Carry IPv6 Traffic

Junos OS Routing Protocols Library

infinity

IN THIS SECTION

- Syntax | 8622
- Hierarchy Level | 8622
- Description | 8622
- Options | 8623
- Required Privilege Level | 8623
- Release Information | 8623

Syntax

```
infinity [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim spt-threshold],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
spt-threshold],
[edit protocols pim spt-threshold],
[edit routing-instances routing-instance-name protocols pim spt-threshold]
```

Description

Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the **infinity** statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Example: Configuring the PIM SPT Threshold Policy

infranet-controller

IN THIS SECTION

- Syntax | 8623
- Hierarchy Level | 8624
- Description | 8624
- Options | 8624
- Required Privilege Level | 8624
- Release Information | 8624

Syntax

```
infranet-contoller host-name {
    address ip-address;
    ca-profile [ca-profile];
```

```
    interface interface-name;
    password password;
    port port-number;
    server-certificate-subject subject;
}
```

Hierarchy Level

```
[edit services unified-access-control ]
```

Description

To configure an Infranet Controller, specify the hostname of the IC Series device with which the SRX Series Firewall should communicate. Possible values for this statement range from 1 to 31 characters.

This statement is required when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series device.

One or more IC Series devices can be configured as Infranet Controllers on the SRX Series Firewall. There is no maximum number of IC Series devices that can be configured. However, only one IC Series device can be active at any time. The others are failover devices. A round-robin algorithm determines which of the configured IC Series devices is the active Infranet Controller. If the active Infranet Controller becomes inoperative, the algorithm is reapplied to the remaining IC Series devices that are configured to establish the new active Infranet Controller.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Firewall User Authentication Overview

ingress (Analyzer)

IN THIS SECTION

- Syntax | 8625
- Hierarchy Level | 8625
- Description | 8626
- Required Privilege Level | 8626
- Release Information | 8626

Syntax

```
ingress {
    bridge-domain bridge-domain-name;
    interface  (all | interface-name);
    routing-instance {
        instance-name {
            bridge-domain bridge-domain-name;
        }
        vlan (vlan-id | vlan-name);
    }
    vlan (vlan-id | vlan-name);
}
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input]
```

Description

Configure ports, routing instances, VLANs, or bridge domains for which the entering traffic is mirrored as part of a mirroring configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

> *Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX9200 Switches*

ingress (Buffer Configuration)

IN THIS SECTION

- Syntax | 8627
- Hierarchy Level | 8627
- Description | 8627
- Default | 8627
- Options | 8628
- Required Privilege Level | 8628
- Release Information | 8628

Syntax

```
ingress {
    buffer-partition (lossless | lossless-headroom | lossy) {
        percent percent;
    }
    percent percent;
}
```

Hierarchy Level

```
[edit class-of-service shared-buffer]
```

Description

Configure the global shared buffer pool allocation for ingress traffic. The system allocates the shared buffer pool dynamically across its ports as the ports require memory space. Some buffer space is reserved for buffers such as dedicated buffers (buffers allocated permanently to ports) and headroom buffers (buffers that help prevent packet loss on lossless flows).

The percentage you specify is the percentage of available (user-configurable) buffer space allocated to the global shared ingress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports.)

> **CAUTION**: Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

You can also partition the shared buffer pool to adjust the ingress buffer allocations for different mixes of network traffic using the `buffer-partition` statement.

Default

The default shared buffer percentage is 100 percent. (All available buffer space is allocated to the shared buffer pool.)

Options

`percent` *percent*—Percentage of available ingress buffer space allocated to the shared buffer pool. If the percentage is less than 100 percent, the remaining buffer space is allocated to the dedicated buffer pool.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic

Configuring Global Ingress and Egress Shared Buffers

Understanding CoS Buffer Configuration

ingress (Chained Composite Next Hop)

IN THIS SECTION

- Syntax | 8629
- Hierarchy Level | 8629
- Description | 8629
- Default | 8630

- Options | 8630
- Required Privilege Level | 8630
- Release Information | 8630

Syntax

```
ingress {
    (evpn | no-evpn);
    (fec129-vpws | no-fec129-vpws);
    (l2ckt | no-l2ckt);
    (l2vpn | no-l2vpn);
    l3vpn ;
    labeled-bgp;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-next-hop],
[edit routing-options forwarding-table chained-composite-next-hop]
```

Description

Allows you to configure chained composite next hops for devices handling ingress traffic in the network.

NOTE:

- On MX Series routers containing both DPC and MPC FPCs, the chassis must be configured with `enhanced-ip` option for enabling chained composite next hops. Only the "l3vpn" on page 9571 option under the `ingress` statement is supported on DPC and DPCE-based MX Series routers. All other statements and functionality under the `ingress` statement are not supported on MX routers with DPC linecards.

- The `extended-space` statement is not supported in ACX Series routers.
- The `[edit logical-systems]` hierarchy level is not applicable in ACX Series routers.
- On MX series routers with Modular Port Concentrator 9 (MPC9) and the previous models of line cards, only the `labeled-bgp` option and the "l3vpn (ingress)" on page 9571 option are supported under the "ingress (Chained Composite Next Hop)" on page 8628 configuration statement; all other configuration options and functionality are not supported.

Default

This statement is disabled by default.

Options

`evpn | no-evpn`	Enable or disable composite chained next hop for ingress EVPN label-switched paths (LSPs).
`fec129-vpws | no-fec129-vpws`	Enable or disable composite chained next hop for ingress FEC 129 virtual private wire service (VPWS) LSPs.
`l2ckt | no-l2ckt`	Enable or disable composite chained next hop for ingress Layer 2 circuit LSPs.
`l2vpn | no-l2vpn`	Enable or disable composite chained next hop for ingress Layer 2 virtual private network (VPN) LSPs.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

This statement was introduced in Junos OS Release 12.1.

`labeled-bgp` option was introduced in Junos OS Release 12.3.

`l2ckt`, `l2vpn`, `no-l2ckt`, and `no-l2vpn` options were introduced in Junos OS Release 13.3.

`evpn`, `fec129-vpws`, `no-evpn`, and `no-fec129-vpws` options were introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs

chained-composite-next-hop | 3277

ingress-egress-recycle

IN THIS SECTION

- Syntax | 8631
- Hierarchy Level | 8631
- Description | 8632
- Required Privilege Level | 8632
- Release Information | 8632

Syntax

```
ingress-egress-recycle;
```

Hierarchy Level

```
[edit system packet-forwarding-options mcast-repl-type]
```

Description

Use this statement to enable ingress and egress recycle replication that helps restarts Packet Forwarding Engine. If you don't select the replication type, ingress replication (ingress-only) is considered as the default replication type.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

ingress (LSP)

IN THIS SECTION

- Syntax | 8632
- Hierarchy Level | 8633
- Description | 8633
- Required Privilege Level | 8633
- Release Information | 8633

Syntax

```
ingress {
    bandwidth bps;
    class-of-service cos-value;
    description string;
    entropy-label;
    install {
        destination-prefix <active>;
```

```
    }
    link-protection bypass-name name;
    metric metric;
    next-hop (address | interface-name | address/interface-name);
    node-protection bypass-name name next-next-label label;
    no-install-to-address;
    policing {
        filter filter-name;
        no-auto-policing;
    }
    preference preference;
    push out-label;
    to address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name]
```

Description

Configure an ingress LSR for a static LSP.

The remaining statements are explained separately

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`entropy-label` option introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring Static LSPs

ingress-node-replication

IN THIS SECTION

- Syntax | 8634
- Hierarchy Level | 8634
- Description | 8634
- Default | 8635
- Required Privilege Level | 8635
- Release Information | 8635

Syntax

```
ingress-node-replication;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name vxlan],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name vxlan],
[edit routing-instances routing-instance-name vlans vlan-name vxlan],
[edit routing-instances routing-instance-name vxlan]
[edit vlans vlan-name vxlan]
```

Description

Enable ingress node replication for a specified Virtual Extensible LAN (VXLAN) that is managed by the Open vSwitch Database (OVSDB) management protocol.

With this feature enabled, instead of service nodes, Juniper Networks devices with OVSDB implemented handle incoming broadcast, unknown unicast, or multicast (BUM) traffic. For more information about the scenarios in which you can use ingress node replication and how it works, see *Understanding How Layer 2 BUM and Layer 3 Routed Multicast Traffic Are Handled with OVSDB*.

> **NOTE**: When Juniper Networks devices replicate Layer 2 BUM packets to a large number of remote software VTEPs, the performance of the Juniper Networks devices might be impacted.

Default

If you do not include the `ingress-node-replication` statement, one or more service nodes handle BUM traffic.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Configuring OVSDB-Managed VXLANs

ingress-node-replication (EVPN)

IN THIS SECTION

- Syntax | 8636
- Hierarchy Level | 8636
- Description | 8636

Required Privilege Level | 8637
Release Information | 8637

Syntax

```
ingress-node-replication;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name vlans vlan-name vxlan],
[edit vlans vlan-name vxlan]
```

Description

EVPN all-active mode with VXLAN encapsulation is based on the local bias for traffic coming from the access layer (redundant Layer 2 gateway function through a pair of top-of-rack switches). Because the traffic has no MPLS label, the split-horizon filtering rule for multihomed Ethernet segments is modified to be based on the IP address of the EVPN provider edge (PE) device instead of the MPLS ES-label. This is called *local bias* for EVPN-VXLAN.

Each EVPN PE device tracks the IP address of its peer multihomed EVPN PE devices that share the same Ethernet segment. This is the source VTEP IP address (outer SIP) for each VXLAN packet received from other EVPN PE devices. The PE device enforces the local bias filtering rule on both ingress and egress PE devices for the multi-destination traffic. For egress traffic, there is no forwarding of any multi-destination packets to the same multihomed Ethernet segment that an egress PE device shares with its ingress PE, regardless of the egress PE device's DF election status for that Ethernet segment. Ingress traffic is responsible for forwarding multi-destination packets coming from any directly attached access interfaces to the rest of the multi-homed Ethernet segments associated with it, regardless of the ingress PE device's designated forwarder election status on the connected physical device segment.

NOTE: With this statement configured, the device adds all active VTEP interfaces to every bridge domain (VLAN). However, the default routing behavior without this statement results in an optimal flood list with VTEPs and VLANs associated only based

on advertised EVPN Type 3 routes. As a result, in general you don't need to configure this statement on EVPN-VXLAN devices.

If you add or remove this option in an existing VLAN configuration statement and commit the change, the device removes the corresponding VLAN from the forwarding plane and adds it back into the forwarding plane. This causes traffic disruption. As a result, you should plan to make either of those types of changes during a maintenance window.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D30 .

RELATED DOCUMENTATION

Understanding EVPN with VXLAN Data Plane Encapsulation

ingress-only

IN THIS SECTION

- Syntax | 8638
- Hierarchy Level | 8638
- Description | 8638
- Required Privilege Level | 8638
- Release Information | 8638

Syntax

```
ingress-only;
```

Hierarchy Level

```
[edit system packet-forwarding-options mcast-repl-type]
```

Description

Use this statement to to enable full ingress replication model that helps restarts Packet Forwarding Engine. If you don't select the replication type, ingress replication (ingress-only) is considered as the default replication type.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

ingress-policer-overhead

IN THIS SECTION

- Syntax | 8639
- Hierarchy Level | 8639
- Description | 8639

- Options | 8640
- Required Privilege Level | 8641
- Release Information | 8641

Syntax

```
ingress-policer-overhead bytes;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number]
```

Description

Add the configured number of bytes to the length of a packet entering the interface.

Configure a policer overhead to control the rate of traffic received on an interface. Use this feature to help prevent denial-of-service (DoS) attacks or to enforce traffic rates to conform to the service-level agreement (SLA). When you configure a policer overhead, the configured policer overhead value (bytes) is added to the length of the final Ethernet frame. This calculated length of frame is used to determine the policer or the rate-limiting action.

Traffic policing combines the configured policy bandwidth limits and the burst size to determine how to meter the incoming traffic. If you configure a policer overhead on an interface, Junos OS adds those bytes to the length of incoming Ethernet frames. This added overhead fills each frame closer to the burst size, allowing you to control the rate of traffic received on an interface.

You can configure the policer overhead to rate-limit queues and Layer 2 and Layer 3 policers, for standalone (SA) and high-avalability (HA) deployments. The policer overhead and the shaping overhead can be configured simultaneously on an interface.

NOTE: vSRX Virtual Firewall supports policer overhead on Layer 3 policers only.

The policer overhead applies to all interfaces on the PIC. In the following example, Junos OS adds 10 bytes of overhead to all incoming Ethernet frames on ports ge-0/0/0 through ge-0/0/4.

```
set chassis fpc 0 pic 0 ingress-policer-overhead 10
```

NOTE: vSRX Virtual Firewall only supports `fpc 0 pic 0`. When you commit the `ingress-policer-overhead` statement, the vSRX Virtual Firewall takes the PIC offline and then back online.

You need to craft the policer overhead size to match your network traffic. A value that is too low will have minimal impact on traffic bursts. A value that is too high will rate-limit too much of your incoming traffic.

In this example, the policer overhead of 255 bytes is configured for ge-0/0/0 through ge-0/0/4. The firewall policer is configured to discard traffic when the burst size is over 1500 bytes. This policer is applied to ge-0/0/0 and ge 0/0/1. Junos OS adds 255 bytes to every Ethernet frame that comes into the configured ports. If, during a burst of traffic, the combined length of incoming frames and the overhead bytes exceeds 1500 bytes, the policer starts to discard further incoming traffic.

```
set chassis fpc 0 pic 0 ingress-policer-overhead 255
set interfaces ge-0/0/0 unit 0 family inet policer input overhead_policer
set interfaces ge-0/0/0 unit 0 family inet address 10.9.1.2/24
set interfaces ge-0/0/1 unit 0 family inet policer input overhead_policer
set interfaces ge-0/0/1 unit 0 family inet address 10.9.2.2/24
set firewall policer overhead_policer if-exceeding bandwidth-limit 32k
set firewall policer overhead_policer if-exceeding burst-size-limit 1500
set firewall policer overhead_policer then discard
```

Options

bytes—Number of bytes added to a frame entering an interface.

- **Range:** 0–255 bytes
- **Default:** 0

```
[edit chassis fpc 0 pic 0]
user@host# set ingress-policer-overhead 10;
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 11.1.

RELATED DOCUMENTATION

ingress-shaping-overhead | 8649

Policer Overhead to Account for Rate Shaping Overview

Example: Configuring Policer Overhead to Account for Rate Shaping

Configuring a Policer Overhead

CoS on Enhanced IQ2 PICs Overview

ingress-policy

IN THIS SECTION

- Syntax | 8641
- Hierarchy Level | 8642
- Description | 8642
- Options | 8642
- Required Privilege Level | 8642
- Release Information | 8642

Syntax

```
ingress-policy [ingress-policy-names];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp entropy-label],
[edit logical-systems logical-system-name protocols ldp oam],
[edit protocols ldp entropy-label],
[edit protocols ldp oam]
```

Description

Configure an LDP ingress policy for either the entropy label or Operation, Administration, and Management (*OAM*).

For OAM, configure the ingress policy to choose which forwarding equivalence classes (*FEC*s) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under `[edit protocols ldp oam bfd-liveness-detection]` are applied.

Options

ingress-policy-names—Specify the names of the ingress policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Statement introduced at the `[edit protocols ldp entropy-label]` hierarchy level in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring OAM Ingress Policies for LDP

Configuring the Entropy Label for LSPs

ingress (Port Mirroring)

IN THIS SECTION

- Syntax | 8643
- Hierarchy Level | 8643
- Description | 8643
- Required Privilege Level | 8644
- Release Information | 8644

Syntax

```
ingress {
     interface (all | interface-name);
     vlan (vlan-id | vlan-name);
}
```

Hierarchy Level

For platforms without ELS:

```
[edit ethernet-switching-options analyzer name input]
```

For platforms with ELS:

```
[edit forwarding-options analyzer name input]
```

Description

Specify the interfaces or VLANs for which incoming traffic is mirrored as part of a port mirroring configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Understanding Port Mirroring and Analyzers

Configuring Port Mirroring

Examples: Configuring Port Mirroring for Local Analysis

ingress-queuing-filter

IN THIS SECTION

- Syntax | 8644
- Hierarchy Level | 8645
- Description | 8645
- Required Privilege Level | 8645
- Release Information | 8645

Syntax

```
ingress-queuing-filter filter-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number family family-name],
[edit logical-systems logical-system-name interfaces interface-name unit unit-number family
family-name]
```

Description

Use the `ingress-queuing-filter` statement to set the packet loss priority and forwarding class for the packet, or drop the packet prior to input queue selection. This assists in traffic shaping.

The `ingress-queuing-filter` statement is available only for the following protocol families: `bridge`, `ccc`, `inet`, `inet6`, `mpls`, and `vpls`.

`ingress-queuing-filter` takes *`filter-name`* as an argument. The named filter is a normal firewall filter that must be configured with at least one of the following actions: `accept`, `discard`, `forwarding-class`, and `loss-priority`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Example: Configuring a Filter for Use as an Ingress Queuing Filter*

ingress-rate-limit

IN THIS SECTION

- Syntax | 8646
- Hierarchy Level | 8646
- Description | 8646
- Options | 8646
- Required Privilege Level | 8646
- Release Information | 8647

Syntax

```
ingress-rate-limit rate;
```

Hierarchy Level

```
[edit interfaces interface-name fastether-options]
```

Description

Perform port-based rate limiting on ingress traffic arriving on Fast Ethernet 8-port, 12-port, and 48-port PICs.

Options

rate—Traffic rate, in megabits per second (Mbps).

- **Range:** 1 through 100 Mbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Ingress Rate Limit

ingress-replication

IN THIS SECTION

- Syntax | 8647
- Hierarchy Level | 8648
- Description | 8648
- Options | 8648
- Required Privilege Level | 8648
- Release Information | 8648

Syntax

```
ingress-replication {
    create-new-ucast-tunnel;
    label-switched-path {
        label-switched-path-template {
            (template-name | default-template);
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel],
[edit protocols mvpn inter-region-template template template-name all-regions],
[edit protocols mvpn inter-region-template template template-name region region-name],
[edit routing-instances routing-instance-name provider-tunnel],
[edit routing-instances routing-instance-name provider-tunnel selective group address source
source-address]
```

Description

A provider tunnel type used for passing multicast traffic between routers through the MPLS cloud, or between PE routers when using MVPN. The ingress replication provider tunnel uses MPLS point-to-point LSPs to create the multicast distribution tree.

Optionally, you can specify a label-switched path template. If you configure `ingress-replication label-switched-path` and do not include `label-switched-path-template`, ingress replication works with existing LDP or RSVP tunnels. If you include `label-switched-path-template`, the tunnels must be RSVP.

Options

existing-unicast-tunnel—An existing tunnel to the destination is used for ingress replication. If an existing tunnel is not available, the destination is not added. Default mode if no option is specified.

create-new-ucast-tunnel—When specified, a new unicast tunnel to the destination is created and used for ingress replication. The unicast tunnel is deleted later if the destination is no longer included in the multicast distribution tree.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs

create-new-ucast-tunnel | 4707

mpls-internet-multicast | 10978

ingress-shaping-overhead

IN THIS SECTION

- Syntax | 8649
- Hierarchy Level | 8649
- Description | 8649
- Options | 8650
- Required Privilege Level | 8650
- Release Information | 8650

Syntax

```
ingress-shaping-overhead number;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number traffic-manager],
[edit chassis lcc number fpc slot-number pic pic-number traffic-manager]
```

Description

Number of bytes to add to packet to determine shaped session packet length.

> **NOTE**: Although the ingress queuing filter can be used with EX9200 switches and MX-Series routers, it is used only on those MX Series routers that have MPCs. An error is generated at commit if the ingress queuing filter is applied to an interface on any other type of port concentrator.

Options

number—When L2TP session shaping is configured, the number of CoS shaping overhead bytes to add to the packets on the ingress side of the L2TP tunnel to determine the shaped session packet length.

When session shaping is not configured and traffic management (queuing and scheduling) is configured on the ingress side, the number of CoS shaping overhead bytes to add to the packets on the ingress interface.

- **Range:** –63 through 192

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

- egress-shaping-overhead | **6223**
- mode (Traffic Manager) | **10807**
- traffic-manager | **30036**

ingress-stats (Flat-File Accounting Options)

IN THIS SECTION

- Syntax | 8651
- Hierarchy Level | 8651
- Description | 8651
- Options | 8652
- Required Privilege Level | 8652
- Release Information | 8652

Syntax

```
ingress-stats {
    all-fields;
    drop-packets;
    input-bytes;
    input-packets;
    output-bytes;
    output-packets;
    queue-id;
}
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name fields]
```

Description

Specify ingress queue statistics to be collected for the interface.

Options

all-fields Collect all ingress queue statistics available for the interface context, logical or physical.

drop-packets Collect the number of received packets dropped on the Ingress queue.

input-bytes Collect the number of octets received on the queue for the traffic class indicated by the queue identifier.

input-packets Collect the number of packets received on the queue for the traffic class indicated by the queue identifier.

output-bytes Collect the number of octets forwarded for the traffic class indicated by the queue identifier. Same value as `input-bytes` unless oversubscription is present at the ingress.

output-packets Collect the number of packets forwarded for the traffic class indicated by the queue identifier. Same value as `input-packets` unless oversubscription is present at the ingress.

queue-id Collect the logical identifier for the ingress queue; identifies the traffic class.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Flat-File Accounting Overview

ingress (vlans)

IN THIS SECTION

- Syntax | 8653
- Hierarchy Level | 8653
- Description | 8653
- Required Privilege Level | 8653
- Release Information | 8653

Syntax

```
ingress;
```

Hierarchy Level

```
[edit vlans  vlan-name vlan-id number interface interface-name]
```

Description

Specify that the member interface of the VLAN allows only ingress traffic.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

inherit-advertisement-interval

IN THIS SECTION

- Syntax | 8654
- Hierarchy Level | 8654
- Description | 8654
- Options | 8654
- Required Privilege Level | 8655
- Release Information | 8655

Syntax

```
inherit-advertisement-interval seconds;
```

Hierarchy Level

```
[edit protocols vrrp]
```

Description

Set the time interval for advertisement for inherit sessions.

Options

inherit-advertisement-interval *seconds*

Time interval for inherit sessions advertisements in seconds. The default value is the recommended value.

- **Default:** 120
- **Range:** 5 to 120

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

init-command-string

IN THIS SECTION

- Syntax | 8655
- Hierarchy Level | 8655
- Description | 8655
- Options | 8656
- Required Privilege Level | 8656
- Release Information | 8656

Syntax

```
init-command-string initialization-command-string;
```

Hierarchy Level

```
[edit interfaces umd0 modem-options]
```

Description

For J Series Services Routers, configure the command string used to initialize the USB modem.

When you connect the USB modem to the USB port on a Services Router, the router applies the modem AT commands configured in the `init-command-string` command to the initialization commands on the modem.

For example, the initialization command string **ATS0 = 2\n** configures the USB modem to pick up a call after 2 rings.

If you do not include the `init-command-string` statement, the router applies the default initialization string to the modem.

Options

initialization-command-string—Specify an initialization command string using the following AT command values:

- **%C0**—Disables data compression.
- **&C1**—Disables reset of the modem when it loses the carrier signal.
- **&Q8**—Enables Microcom Networking Protocol (MNP) error control mode.
- **AT**—Attention. Informs the modem that a command follows.
- **E0**—Disables the display on the local terminal of commands issued to the modem from the local terminal.
- **Q0**—Enables the display of result codes.
- **S0=0**—Disables the auto-answer feature, whereby the modem automatically answers calls.
- **S7=45**—Instructs the modem to wait 45 seconds for a telecommunications service provider (carrier) signal before terminating the call.
- **V1**—Displays result codes as words.
- **Default: AT S7=45 S0=0 V1 X4 &C1 E0 Q0 &Q8 %C0**

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

init-descriptor

IN THIS SECTION

- Syntax | 8657
- Hierarchy Level | 8657
- Description | 8657
- Options | 8657
- Required Privilege Level | 8657
- Release Information | 8658

Syntax

```
init-descriptor file-path;
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Create an XML descriptor file to launch a VNF. You can launch a VNF by configuring the VNF name, and specifying either the path to the XML descriptor file or to an image.

Options

file-path	Path of the init-descriptor XML file.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D45.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

initial-route-check

IN THIS SECTION

- Syntax | 8658
- Hierarchy Level | 8658
- Description | 8659
- Options | 8659
- Required Privilege Level | 8659
- Release Information | 8659

Syntax

```
initial-route-check seconds;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN interfaces, allows the router to check whether the primary route is up after the initial startup of the router is complete and the timer expires.

Options

seconds—How long to wait to check if the primary interface is up after the router comes up.

- **Range:** 1 through 300 seconds
- **Default:** 120 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

ISDN Interfaces Overview

initial-tcp-mss

IN THIS SECTION

- Syntax | 8660
- Hierarchy Level | 8660
- Description | 8660
- Required Privilege Level | 8660
- Release Information | 8660

Syntax

```
initial-tcp-mss mss-value;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name
policy policy-name then permit tcp-options]
```

Description

Configure the TCP maximum segment size (MSS) for packets that arrive at the ingress interface (initial direction), match a specific policy, and for which a session is created. The value you configure overrides the TCP MSS value in the incoming packet when the value in the packet is higher than the one you specify.

The `initial-tcp-mss` value per policy takes precedence over a global `tcp-mss` value (`all-tcp`, `ipsec-vpn`, `gre-in`, `gre-out`), if one is configured. However, when the `syn-flood-protection-mode syn-proxy` statement at the `[edit security flow]` hierarchy level is used to enable SYN proxy defenses against SYN attacks, the TCP MSS value is not overriden.

Because each policy has two directions, you can configure a value for both directions or for just one direction. To configure a TCP MSS value for the reverse session, use the `reverse-tcp-mss` option.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D20.

RELATED DOCUMENTATION

reverse-tcp-mss | 15722

tcp-mss (Security Flow)

syn-flood-protection-mode

initiate-dead-peer-detection

IN THIS SECTION

- Syntax | 8661
- Hierarchy Level | 8661
- Description | 8661
- Required Privilege Level | 8661
- Release Information | 8662

Syntax

```
initiate-dead-peer-detection;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Enable triggering of dead peer detection (DPD) hello messages to the remote peer for the specified tunnel.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2 .

Support for IKEv2 introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

dead-peer-detection (Services IPsec VPN) | **5151**

backup-remote-gateway | **2811**

initiate-ncp (Dynamic and Static PPP)

IN THIS SECTION

- Syntax | **8662**
- Hierarchy Level | **8662**
- Description | **8663**
- Options | **8663**
- Required Privilege Level | **8663**
- Release Information | **8663**

Syntax

```
initiate-ncp (ip | ipv6 | dual-stack-passive);
```

Hierarchy Level

```
[edit access group-profile profile-name ppp ppp-options],
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options],
[edit dynamic-profiles profile-name interfaces “$junos-interface-ifd-name” unit “$junos-
```

```
interface-unit” ppp-options],
[edit interfaces pp0 unit logical-unit-number ppp-options],
[edit interfaces si-fpc/pic/port unit logical-unit-number ppp-options]
```

Description

Configure PPP Network Control Protocol (NCP) negotiation mode (active or passive) for dynamic and static IPv4 and IPv6 PPP subscriber interfaces. You can also configure PPP NCP negotiation mode for the PPP server in an IPv4/IPv6 dual-stack configuration.

Options

dual-stack-passive Enable passive PPP NCP negotiation for the PPP server in an IPv4/IPv6 dual-stack configuration. The `initiate-ncp dual-stack-passive` statement overrides the `initiate-ncp ip` and `initiate-ncp ipv6` statements if they are configured in an IPv4/IPv6 dual-stack configuration.

ip Enable active PPP NCP negotiation for dynamic and static PPP subscriber interfaces configured with the IPv4 (`inet`) protocol address family, and for which IPv4 address attributes are assigned during authorization. By default, dynamic and static IPv4 subscriber interfaces use passive PPP NCP negotiation. In an IPv4/IPv6 dual-stack configuration, use the `initiate-ncp ip` statement to enable active PPP NCP negotiation for the IPv4 subscriber interface.

ipv6 Enable active PPP NCP negotiation for dynamic and static PPP subscriber interfaces configured with the IPv6 (`inet6`) protocol address family, and for which IPv6 address attributes are assigned during authorization. By default, dynamic and static IPv6 subscriber interfaces use passive PPP NCP negotiation. In an IPv4/IPv6 dual-stack configuration, use the `initiate-ncp ipv6` statement to enable active PPP NCP negotiation for the IPv6 subscriber interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

PPP Network Control Protocol Negotiation

PPP Network Control Protocol Negotiation

initiation (Services)

IN THIS SECTION

- Syntax | 8664
- Hierarchy Level | 8665
- Description | 8665
- Options | 8665
- Required Privilege Level | 8665
- Release Information | 8665

Syntax

```
initiation{
    profile name {
        actions {
            crl {
                disable disable;
                if-not-present (allow | drop);
                ignore-hold-instruction-code ignore-hold-instruction-code;
            }
            ignore-server-auth-failure ignore-server-auth-failure;
        }
        client-certificate client-certificate;
        custom-ciphers  (ecdhe-rsa-with-3des-ede-cbc-sha | ecdhe-rsa-with-aes-128-cbc-sha |
ecdhe-rsa-with-aes-128-cbc-sha256 | ecdhe-rsa-with-aes-128-gcm-sha256 | ecdhe-rsa-with-aes-256-
cbc-sha | ecdhe-rsa-with-aes-256-cbc-sha384 | ecdhe-rsa-with-aes-256-gcm-sha384 | rsa-export-
with-des40-cbc-sha | rsa-export-with-rc4-40-md5 | rsa-export1024-with-des-cbc-sha | rsa-
export1024-with-rc4-56-md5 | rsa-export1024-with-rc4-56-sha | rsa-with-3des-ede-cbc-sha | rsa-
with-aes-128-cbc-sha | rsa-with-aes-128-cbc-sha256 | rsa-with-aes-128-gcm-sha256 | rsa-with-
```

```
aes-256-cbc-sha | rsa-with-aes-256-cbc-sha256 | rsa-with-aes-256-gcm-sha384 | rsa-with-des-cbc-
sha | rsa-with-null-md5 | rsa-with-null-sha | rsa-with-rc4-128-md5 | rsa-with-rc4-128-sha);
        enable-flow-tracing enable-flow-tracing;
        enable-session-cache enable-session-cache;
        preferred-ciphers (custom | medium | strong | weak);
        protocol-version (all | ssl3 | tls1 | tls11 | tls12);
        trusted-ca ;
    }
}
```

Hierarchy Level

```
[edit services ssl]
```

Description

Specify the configuration for Secure Socket Layer (SSL) initiation support service. The SRX Series Firewall, acting as an SSL proxy client, initiates and maintains SSL sessions between itself and an SSL server. SRX Series Firewall receives un-encrypted data from an HTTP client, and encrypts and transmits the data as ciphertext to the SSL server.

Options

- **client-certificate**—Local certificate.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. The `protocol-version` statement is updated to include `tls11` and `tls12` from Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

Firewall User Authentication Overview

inline

IN THIS SECTION

- Syntax | 8666
- Hierarchy Level | 8666
- Description | 8666
- Default | 8667
- Required Privilege Level | 8667
- Release Information | 8667

Syntax

```
inline;
```

Hierarchy Level

```
[edit protocols lacp ppm]
```

Description

(MX Series routers with MPC line cards only) To enable the inline Link Aggregation Control Protocol (LACP) PDU transmission processing.

This statement disables the default distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all Link Aggregation Control Protocol (LACP) PDU transmission processing inline. The inline option can be used in scenarios where the line card CPU is under heavy load and cannot schedule the PPM processing for LACP packets. PPM, by default,

delegates the transmission of PDUs to the PPMAN process on the PFE/line card. But when the inline option is configured, it delegates the transmission of LCAP PDUs even further away from the line card CPU and into the forwarding chipset.

For example, in a system with both MPCs and DPCs, upon configuration of `[protocols lacp ppm inline]`, the PDUs are sent inline on the MPCs and performed by periodic packet management (PPM) on DPCs.

> **BEST PRACTICE**: We recommend to retain the default and disable distributed PPM or enable inline processing only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM or enable inline processing only if you have a compelling reason to disable it.

Refer Disabling or Enabling Inline Periodic Packet Management for LACP Packets for more details.

Default

Distributed PPM processing is enabled for all packets that use PPM.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

centralized | 3250

Disabling or Enabling Inline Periodic Packet Management for LACP Packets

Configuring Link Aggregation

Configuring Aggregated Ethernet LACP (CLI Procedure)

inline-fpga-crypto

IN THIS SECTION

- Syntax | 8668
- Hierarchy Level | 8668
- Description | 8668
- Options | 8669
- Required Privilege Level | 8669
- Release Information | 8669

Syntax

```
inline-fpga-crypto (disabled | enabled);
```

Hierarchy Level

```
[edit security forwarding-process application-services (Security Forwarding Process)]
```

Description

Enable or disable Inline FPGA Decryption (IFD) for cryptographic operation.

By default, the inline-fpga-crypto is disabled. To use the feature, enable the option.

Ensure you enable PMI processing when using `inline-fpga-crypto` option to avoid the risk of traffic loss.

PowerMode IPsec (PMI) uses AES-NI for encryption and FPGA for decryption of cryptographic operation to enhance IPSec VPN performance.

PMI with FPGA is supported on SRX5400, SRX5600, and SRX5800 with SPC3 cards.

Options

disabled	Disable inline FPGA crypto
enabled	Enable inline FPGA crypto

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Improving IPsec Performance with PowerMode IPsec

application-services (Security Forwarding Process)

inline-jflow

IN THIS SECTION

- Syntax | 8669
- Hierarchy Level | 8670
- Description | 8670
- Required Privilege Level | 8670
- Release Information | 8670

Syntax

```
inline-jflow {
    source-address address;
```

```
    flow-export-rate rate;
}
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family inet output]
```

Description

Specify inline active flow monitoring for traffic from the designated address.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: If you configure inline flow monitoring with `inline-jflow` then you have to disable it before performing ISSU. For more information, see Before You Begin a Unified ISSU.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.

Support for IPv6 addresses for the `source-address` *address* statement introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

inline-jflow (Forwarding Options)

IN THIS SECTION

- Syntax | 8671
- Hierarchy Level | 8671
- Description | 8671
- Options | 8671
- Required Privilege Level | 8672
- Release Information | 8672

Syntax

```
inline-jflow {
    flow-export-rate number;
    source-address ip-address;
}
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance- name family inet output]
[edit forwarding-options sampling instance instance- name family inet6 output]
```

Description

Specify Inline processing of sampled packets.

Options

- `flow-export-rate` *value*—Flow export rate of monitored packets in kpps. The range is from 1 through 400.
- `source-address` *address*—Address to use for generating monitored packets.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4. Support.

RELATED DOCUMENTATION

Understanding Interfaces

inline-monitoring

IN THIS SECTION

- Syntax | 8672
- Hierarchy Level | 8673
- Description | 8673
- Options | 8673
- Required Privilege Level | 8673
- Release Information | 8674

Syntax

```
inline-monitoring {
    instance instance-name;
    observation-cloud-id observation-cloud-identifier;
    template template-name;
    traceoptions;
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure inline monitoring services. When you enable inline monitoring, you can monitor actual IPv4 and IPv6 packets at different sampling rates, and export the actual packet up to the configured clip length. By default, Junos OS supports a maximum packet length of 126 bytes starting with the Ethernet header. The monitored packets are exported to an collector for further processing. The packets are exported in an IPFIX format, which includes information on the original packet size, and incoming or outgoing interface.

Options

instance *instance-name*

Configure parameters for an inline-monitoring instance. See "instance" on page 8750 for more information.

observation-cloud-id *observation-cloud-identifier*

Observation cloud ID—an identifier for a particular observation cloud. An observation cloud is the largest set of observation domains. Per RFC 5101, an observation domain is the largest set of observation points for which flow information can be aggregated by a metering process. For example, a router line card may be an observation domain if it is composed of several interfaces, each of which is an observation point. The observation domain ID is unique per exporting process. Also per RFC 5101, an observation point is a location in the network where IP packets can be observed. Examples include: a line to which a probe is attached, a shared medium, such as an Ethernet-based LAN, a single port of a router, or a set of interfaces (physical or logical) of a router.

- Range: 1 through 255

template *template-name*

Configure templates for inline packet monitoring. See "template" on page 29084 for more information.

traceoptions

(Optional) Configure traceoptions for the inline monitoring process. See "traceoptions" on page 29585 for more information.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

`observation-cloud-id` option introduced in Junos OS Release 21.2R1.

Statement introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

Inline Monitoring Services

Flow-Based Telemetry

inline-services

IN THIS SECTION

- Syntax | 8674
- Hierarchy Level | 8675
- Description | 8675
- Default | 8675
- Options | 8675
- Required Privilege Level | 8675
- Release Information | 8675

Syntax

```
inline-services {
  flex-flow-sizing;
  flow-table-size {
    bridge-flow-table-size bridge-flow-table-size;
    ipv4-flow-table-size ipv4-flow-table-size;
```

```
    ipv6-extended-attrib;
    ipv6-flow-table-size ipv6-flow-table-size;
    mpls-flow-table-size mpls-flow-table-size;
  }
  report-zero-oif-gw-on-discard;
  use-extended-flow-memory;
}
```

Hierarchy Level

```
[edit chassis fpc fpc-number]
```

Description

(MX Series only) Enable inline services for FPCs.

Default

Disabled.

Options

The remaining statements are explained separately. See CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4R1.

inline-services (PIC level)

IN THIS SECTION

- Syntax | 8676
- Hierarchy Level | 8676
- Description | 8676
- Required Privilege Level | 8677
- Release Information | 8677

Syntax

```
inline-services {
                  service-port;
 bandwidth bandwidth;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number]
```

Description

Enable inline services on PICs residing on MPCs and optionally specify a bandwidth for traffic on the inline service interface. Bandwidth values can be 1g, 10g, 20g, 30g, 40g, 50g, 60g, 70g, 80g, 90g, 100g, 200g, 300g, or 400g.

NOTE: For an MPC, such as MPC2, always configure inline-services at the `[chassis fpc slot-number pic number]` hierarchy level. Do not configure inline services for a service card such as MS-MPC.

For MX304, MX10K-LC9600 and MX10K-LC4800:

> The valid bandwidth range is 2Gbps through 400Gbps with increments of 1Gbps. When you configure inline services bandwidth of 1Gbps, the Routing Engine CLI configuration commit does not fail.
>
> If you newly commit an inline tunnel service configuration with the bandwidth of 1Gbps or apply it as part of device reboot, `si-` interfaces are not created. However, no commit error message is displayed. If you change the bandwidth on an existing inline services configuration to 1Gbps, a commit error message will not be displayed. However, the configuration change does not take effect. The previously committed inline services configuration continues to be in effect.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Enabling Inline Service Interfaces

Configuring an L2TP LNS with Inline Service Interfaces

inline-tap

IN THIS SECTION

Syntax | **8678**

- Hierarchy Level | 8678
- Description | 8678
- Required Privilege Level | 8678
- Release Information | 8678

Syntax

```
inline-tap;
```

Hierarchy Level

```
[edit security forwarding-process application-services maximize-idp-sessions]
```

Description

Enable IDP inline tap mode. The inline tap feature provides passive, inline detection of application layer threats for traffic matching security policies which have the IDP application service enabled. When a device is in inline tap mode, packets pass through firewall inspection and are also copied to the independent IDP module.

This statement is supported in SRX1500, SRX 5800, SRX 5600, and SRX 5400 devices and vSRX Virtual Firewall.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

inner-tag-protocol-id

IN THIS SECTION

- Syntax | 8679
- Hierarchy Level | 8679
- Description | 8679
- Default | 8680
- Required Privilege Level | 8680
- Release Information | 8680

Syntax

```
inner-tag-protocol-id tpid;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

Description

Configure the IEEE 802.1Q TPID value to rewrite for the inner tag.

All TPIDs you include in input and output VLAN maps must be among those you specify at the `[edit interfaces interface-name gigether-options ethernet-switch-profile tag-protocol-id [ tpids ]]` hierarchy level.

On MX Series routers, you can use this statement for Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.

Default

If the `inner-tag-protocol-id` statement is not configured, the TPID value is 0x8100.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring Inner and Outer TPIDs and VLAN IDs

inner-tag-protocol-id (Dynamic VLANs)

IN THIS SECTION

- Syntax | 8681
- Hierarchy Level | 8681
- Description | 8681
- Default | 8681
- Options | 8681
- Required Privilege Level | 8681
- Release Information | 8681

Syntax

```
inner-tag-protocol-id tpids;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-numberinput-vlan-
map],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-numberoutput-vlan-
map]
```

Description

For dynamic VLAN interfaces, configure the IEEE 802.1Q TPID value to rewrite for the inner tag. All TPIDs you include in input and output VLAN maps must be among those you specify at the `[edit interfaces interface-name gigether-options ethernet-switch-profile tag-protocol-id tpids ]` hierarchy level.

Default

If the `inner-tag-protocol-id` statement is not configured, the TPID value is 0x8100.

Options

tpids—TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal format.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Inner and Outer TPIDs and VLAN IDs

inner-vlan-id

IN THIS SECTION

- Syntax | 8682
- Hierarchy Level | 8682
- Description | 8682
- Options | 8683
- Required Privilege Level | 8683
- Release Information | 8683

Syntax

```
inner-vlan-id number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

Description

For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers or 100-Gigabit Ethernet Type

5 PIC with CFP, or on Ethernet interfaces on EX Series switches, specify the VLAN ID to rewrite for the inner tag of the final packet.

You cannot include the `inner-vlan-id` statement with the `swap` statement, `swap-push` statement, `push-push` statement, or `push-swap` statement and the `inner-vlan-id` statement at the [`edit interfaces` *`interface-name`* `unit` *`logical-unit-number`*`output-vlan-map`] hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the `inner-vlan-id` statement you include at the [`edit interfaces` *`interface-name`* `unit` *`logical-unit-number`*] hierarchy level.

Options

`number`—VLAN ID number.

- **Range:** 0 through 4094

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring Inner and Outer TPIDs and VLAN IDs

inner-vlan-id (Dynamic VLANs)

IN THIS SECTION

- Syntax | 8684
- Hierarchy Level | 8684
- Description | 8684
- Options | 8684

Required Privilege Level | 8684
Release Information | 8685

Syntax

```
inner-vlan-id number;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number input-vlan-map],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map]
```

Description

For dynamic VLAN interfaces, specify the VLAN ID to rewrite for the inner tag of the final packet.

You cannot include the `inner-vlan-id` statement with the `swap` statement, `swap-push` statement, `push-push` statement, or `push-swap` statement and the `inner-vlan-id` statement at the `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number output-vlan-map]` hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the `inner-vlan-id` statement you include at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level.

Options

number—VLAN ID number. When used for input VLAN maps, you can specify the `$junos-inner-vlan-map-id` predefined variable to dynamically obtain the VLAN identifier.

- **Range:** 0 through 4094

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Inner and Outer TPIDs and VLAN IDs

inner-vlan-id-swap-ranges

IN THIS SECTION

- Syntax | 8685
- Hierarchy Level | 8685
- Description | 8686
- Options | 8686
- Required Privilege Level | 8686
- Release Information | 8686

Syntax

```
inner-vlan-id-swap-ranges low-inner-tag-high-inner-tag;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Define core-facing VLAN ID ranges from which an inner VLAN ID tag can be allocated to replace the outer VLAN tag that was appended by the access node on the upstream packets to the BNG in a Layer 2 wholesale network. The tag swap occurs before the subscriber traffic is forwarded to the network service provider (NSP). You can configure up to 32 non-overlapping inner VLAN ID ranges per core-facing physical interface for VLAN-OOB subscribers.

You can add or remove ranges or increase or decrease the size of existing ranges even while Layer 2 wholesale sessions are assigned to the core-facing interface associated with the ranges. You cannot remove a range from which a VLAN ID has already been allocated. You cannot reduce a range if the new range excludes a VLAN ID that has already been allocated.

Options

low-inner-tag Inner (core-facing) VLAN ID tag representing the lower limit of the swap range.

- **Range:** 1 through 4094.

high-inner-tag Inner (core-facing) VLAN ID tag representing the upper limit of the swap range.

- **Range:** 1 through 4094.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Configuring Multiple Non-Overlapping VLAN Ranges for Core-Facing Physical Interfaces

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

input

IN THIS SECTION

- Syntax | 8687
- Hierarchy Level | 8687
- Description | 8688
- Default | 8688
- Required Privilege Level | 8688
- Release Information | 8688

Syntax

```
input {
     ingress {
         interface (all | interface-name);
         vlan (vlan-id | vlan-name);
    }
     egress {
         interface (all | interface-name);
    }
}
```

Hierarchy Level

For platforms without ELS:

```
[edit ethernet-switching-options analyzer name]
```

For platforms with ELS:

```
[edit forwarding-options analyzer name]
```

Description

Define the traffic to be mirrored. The definition can be a combination of traffic entering or exiting specific ports or VLANs.

The remaining statements are explained separately. See CLI Explorer.

Default

No default.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Understanding Port Mirroring and Analyzers

Configuring Port Mirroring

Examples: Configuring Port Mirroring for Local Analysis

input (Analyzer)

IN THIS SECTION

- Syntax | 8689
- Hierarchy Level | 8689
- Description | 8689
- Default | 8690
- Required Privilege Level | 8690

Release Information | 8690

Syntax

```
input {
    egress {
        bridge-domain bridge-domain-name;
        interface (all | interface-name);
        routing-instance {
            instance-name {
                bridge-domain bridge-domain-name;
            }
        }
    }
    ingress {
        bridge-domain bridge-domain-name;
        interface  (all | interface-name);
        routing-instance {
            instance-name {
                bridge-domain bridge-domain-name;
            }
            vlan (vlan-id | vlan-name);
        }
        vlan (vlan-id | vlan-name);
    }
    maximum-packet-length bytes;
    rate number;
}
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name]
```

Description

Define the traffic to be mirrored in a mirroring configuration—the definition can be a combination of:

- Packets entering or exiting a port
- Packets entering or exiting a VLAN
- Packets entering or exiting a bridge domain

The remaining statements are explained separately. See CLI Explorer.

Native analyzer sessions (that is, the `[edit forwarding-options analyzer analyzer-name input]` hierarchy level for MX Series routers) can be configured without specifying input parameters, which would mean that the instance uses default input values: rate = 1 and maximum-packet-length = 0.

Default

No default.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Example: Configuring Port Mirroring Analyzers for Local Monitoring of Employee Resource Use

Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use

Understanding Port Mirroring Analyzers

input-chain

IN THIS SECTION

Syntax | 8691
Hierarchy Level | 8691

- Description | 8691
- Options | 8691
- Required Privilege Level | 8691
- Release Information | 8692

Syntax

```
input-chain [ filter-name ]
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family filter]
```

Description

The `input-chain` command works like the **next** filter term to allow multiple levels of filtering on the interface in both the ingress and egress direction. For example, you can implement a combination of classification and firewall filter rules for evaluation, wherein the first filter performs a generic filter classification and the subsequent filters go on to perform the actual filtering and actions.

To continue on to the next filter, the terminating action in the current filter must be **accept**.

The `input-chain` command is supported on all interfaces except loopback (lo0 and fxp0) and interfaces under logical systems.

Options

[*filter-names*]—Name of each filter in the chain that is to be evaluated when packets are received on the interface. Supports up to 8 ingress and 8 egress filters. Evaluation occurs in the order in which the names appear, from left to right.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

Statement introduced in Junos OS Evolved Release 21.4R1 (PTX10001-36MR, PTX10004, PTX10008, and PTX10016).

RELATED DOCUMENTATION

Example: Using Firewall Filter Chains

output-chain | **11983**

Applying a Filter to an Interface

input (Congestion Notification)

IN THIS SECTION

- Syntax | **8692**
- Hierarchy Level | **8693**
- Description | **8693**
- Required Privilege Level | **8693**
- Release Information | **8693**

Syntax

```
input {
    ieee-802.1 {
        code-point up-bits pfc ;
    }
}
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name],
[edit class-of-service interfaces interface-name congestion-notification-profile profile-name]
```

Description

Identify the three-bit pattern of the User Priority field that triggers the priority-based congestion notification profile for a specified traffic class.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring an FCoE Transit Switch

Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)

input (Congestion Notification)

IN THIS SECTION

- Syntax | 8694
- Hierarchy Level | 8694
- Description | 8694
- Options | 8694

Required Privilege Level | 8695
Release Information | 8695

Syntax

```
input {
    (dscp | ieee-802.1) {
        code-point [code-point-bits] {
            pfc {
                mru mru-value;
            }
            xon
        }
    }
    cable-length cable-length-value;
}
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name]
```

Description

Configure priority-based flow control (PFC) on incoming traffic based on either IEEE 802.1p priorities in the Layer 2 VLAN header or Differentiated Services code point (DSCP) values in the Layer 3 IP header.

You cannot configure both DSCP-based PFC and IEEE 802.1p PFC under the same congestion notification profile.

Options

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Support for DSCP values introduced in Junos OS Release 17.4R1 for the QFX Series.

RELATED DOCUMENTATION

Example: Configuring CoS PFC for FCoE Traffic

Configuring CoS PFC (Congestion Notification Profiles)

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Understanding PFC Using DSCP at Layer 3 for Untagged Traffic

Configuring DSCP-based PFC for Layer 3 Untagged Traffic

input-current (T4000 Routers)

IN THIS SECTION

- Syntax | 8696
- Hierarchy Level | 8696
- Description | 8696
- Options | 8696
- Required Privilege Level | 8696
- Release Information | 8696

Syntax

```
input-current amps-in-each-feed;
```

Hierarchy Level

```
[edit chassis pem]
```

Description

Configure the amount of input current received in each feed. The value assigned to the `input-current` statement must be equal to the input current capability of each feed.

> **NOTE**: Before configuring input current for your router, see the *T4000 Core Router Hardware Guide* for special considerations.

Options

Values:

-
  - 40–Indicates 40 A of input current is received in each feed.
  - 60–Indicates 60 A of input current is received in each feed.
- **Default:** 60 A

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring the Six-Input DC Power Supply on T Series Routers

input (Dynamic Service Sets)

IN THIS SECTION

- Syntax | 8697
- Hierarchy Level | 8697
- Description | 8697
- Required Privilege Level | 8698
- Release Information | 8698

Syntax

```
input {
     service-set service-set-name {
        service-filter filter-name;
    }
    post-service-filter filter-name;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family service],
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” family family
service]
```

Description

Define the input service sets and filters to be applied to traffic by a dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family service]` hierarchy level introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic Service Sets Overview

Associating Service Sets with Interfaces in a Dynamic Profile

input-excess-bandwidth-share

IN THIS SECTION

- Syntax | 8699
- Hierarchy Level | 8699
- Description | 8699
- Options | 8699
- Required Privilege Level | 8699
- Release Information | 8699

Syntax

```
input-excess-bandwidth-share (proportional value | equal);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit class-of-service interfaces interface-set interface-set-name]
```

Description

Determines the method of sharing excess bandwidth on the ingress interface in a hierarchical scheduler environment. If you do not include this statement, the node shares excess bandwidth proportionally at 32.64 Mbps.

Options

`proportional` *`value`*—(Default) Share ingress excess bandwidth proportionally (default value is 32.64 Mbps).

`equal`—Share ingress excess bandwidth equally.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs

input (Fibre Channel Interfaces)

IN THIS SECTION

- Syntax | 8700
- Hierarchy Level | 8700
- Description | 8700
- Required Privilege Level | 8701

Syntax

```
input {
    ieee-802.1{
        code-point code-point-bits;
    }
}
```

Hierarchy Level

```
[edit class-of-service interfaces fibre-channel-interface-name rewrite-value]
```

Description

Configure the IEEE 802.1p code point value to which all traffic received from the Fibre Channel (FC) network on the specified FC interface is rewritten. After the code point value is rewritten, the interface forwards the traffic to the Ethernet (FCoE) network. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.

FCoE networks typically use priority 3 (IEEE code point 011) for FCoE traffic. The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011).

The system supports only one IEEE 802.1p code point value per FC interface. You cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.

The statements are described separately.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

forwarding-class (Fibre Channel Interfaces) | **7421**

Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

Understanding CoS Classifiers

input-firewall-optimized-profile

IN THIS SECTION

- Syntax | **8702**
- Hierarchy Level | **8702**
- Description | **8702**
- Required Privilege Level | **8702**
- Release Information | **8702**

Syntax

```
input-firewall-optimized-profile
```

Hierarchy Level

```
[edit chassis]
```

Description

Optimize IFP DYN group (IFP dynamic group) by removing less frequently used qualifiers.

> **NOTE**: This configuration statement should be used only when FIP snooping, Fibre Channel over Ethernet (FCoE), and Inband Flow Analyzer (IFA) features are not used. On EX4400, this configuration statement should be used only when FIP snooping, FCoE, IFA, as well as inline monitoring features are not used.

Required Privilege Level

interface-control

Release Information

Statement introduced in Junos OS Release 22.2R1

RELATED DOCUMENTATION

show pfe filter hw summary | 22983

input (Forwarding Table)

IN THIS SECTION

- Syntax | 8703
- Hierarchy Level | 8703
- Description | 8703
- Options | 8703
- Required Privilege Level | 8703
- Release Information | 8704

Syntax

```
input filter-name;
```

Hierarchy Level

```
[edit forwarding-options family (inet | inet6 | mpls | vpls) filter],
[edit routing-instances routing-instance-name forwarding-options family (inet | inet6 | mpls |
vpls) filter]
```

Description

Apply a forwarding table filter to ingress traffic of the forwarding table.

Options

filter-name—Name of the applied filter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Table Filters

input-hierarchical-policer

IN THIS SECTION

- Syntax | 8704
- Hierarchy Level | 8704
- Description | 8705
- Options | 8705
- Required Privilege Level | 8705
- Release Information | 8705

Syntax

```
input-hierarchical-policer policer-name;
```

Hierarchy Level

```
[edit interfaces interface-name layer2-policer],
[edit interfaces interface-name unit logical-unit-number layer2-policer]
```

Description

Apply a hierarchical policer to the Layer 2 input traffic for all protocol families at the physical or logical interface.

Options

policer-name—Name of the hierarchical policer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Hierarchical Policers

layer2-policer (Hierarchical Policer)

input-interface-index

IN THIS SECTION

- Syntax | 8706
- Hierarchy Level | 8706
- Description | 8706
- Options | 8706
- Required Privilege Level | 8706
- Release Information | 8706

Syntax

```
input-interface-index number;
```

Hierarchy Level

```
[edit forwarding-options monitoring name output interface interface-name]
```

Description

Specify a value for the input interface index that overrides the default supplied by SNMP.

Options

number—Input interface index value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Flow Monitoring

input (Interfaces)

IN THIS SECTION

- Syntax | 8707
- Hierarchy Level | 8707
- Description | 8707
- Options | 8707
- Required Privilege Level | 8708
- Release Information | 8708

Syntax

```
input {
    service-set service-set-name <service-filter filter-name>;
    post-service-filter filter-name;
}
```

Hierarchy Level

```
[edit interface interface-name unit logical-unit-number family inet service],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet service]
```

Description

Define the input service sets and filters to be applied to traffic.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Filters and Services to Interfaces

input-list

IN THIS SECTION

- Syntax | 8708
- Hierarchy Level | 8709
- Description | 8709
- Options | 8709
- Required Privilege Level | 8709
- Release Information | 8709

Syntax

```
input-list [ filter-names ];
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family filter],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family filter]
```

Description

Apply a group of filters to evaluate when packets are received on an interface.

Options

[*filter-names*]—Name of a filter to evaluate when packets are received on the interface. Up to 16 filters can be included in a filter input list.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

- *Applying a Filter to an Interface*
- Routing Policies, Firewall Filters, and Traffic Policers User Guide
- Junos OS Administration Library for Routing Devices
- output-list

input-native-vlan-push

IN THIS SECTION

- Syntax | 8710
- Hierarchy Level | 8710
- Description | 8710
- Default | 8710
- Options | 8711
- Required Privilege Level | 8711
- Release Information | 8711

Syntax

```
input-native-vlan-push (disable | enable);
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

On EX2300, EX3400, EX4300, EX4300-MP, EX4600, EX4650, and the QFX5000 line of switches except QFX5220, QFX5130, QFX5230, QFX5240, and QFX5700, for a Q-in-Q tunneling configuration, enable or disable whether the switch inserts a native VLAN identifier in untagged frames received on the C-VLAN interface when the configuration statement "input-vlan-map" on page 8729 with a `push` operation is configured. QFX5220, QFX5130, QFX5230, QFX5240, and QFX5700 switches do not insert a native VLAN identifier in untagged frames received on the C-VLAN interface when the configuration statement "input-vlan-map" on page 8729 with a `push` operation is configured.

Default

If you do not include the `input-native-vlan-push` statement, the option is enabled.

> **NOTE**: Starting in Junos OS Release 21.4R3-S7, the default on EX4300 switches is `enable`, In earlier releases, the default is `disable`.

Options

disable Disable `input-native-vlan-push`.

enable Enable `input-native-vlan-push`.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 19.3R1 on EX Series and QFX Series switches.

RELATED DOCUMENTATION

Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation

input-packet-rate-threshold

IN THIS SECTION

- Syntax | 8712
- Hierarchy Level | 8712
- Description | 8712
- Options | 8712
- Required Privilege Level | 8712
- Release Information | 8712

Syntax

```
input-packet-rate-threshold rate;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Specify a packet rate threshold value that triggers a system log warning message.

Options

rate—Threshold value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring Thresholds*

input-policer

IN THIS SECTION

- Syntax | 8713
- Hierarchy Level | 8713
- Description | 8713
- Options | 8713
- Usage Guidelines | 8714
- Required Privilege Level | 8714
- Release Information | 8714

Syntax

```
input-policer policer-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number layer2-policer]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
layer2-policer]
```

Description

Apply a single-rate two-color policer to the Layer 2 input traffic at the logical interface. The `input-policer` and `input-three-color` statements are mutually exclusive.

Options

policer-name—Name of the single-rate two-color policer that you define at the `[edit firewall]` hierarchy level.

Usage Guidelines

See *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Two-Color and Three-Color Policers at Layer 2

Applying Layer 2 Policers to Gigabit Ethernet Interfaces

Configuring Gigabit Ethernet Policers

input-three-color | **8724**

layer2-policer (EX) | **9654**

logical-interface-policer | **10069**

output-policer | **12001**

output-three-color | **12014**

input (Port Mirroring)

IN THIS SECTION

- Syntax | **8715**
- Hierarchy Level | **8715**
- Description | **8715**
- Required Privilege Level | **8715**
- Release Information | **8715**

Syntax

```
input {
    maximum-packet-length bytes;
    rate number;
    run-length number;
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring],
[edit forwarding-options port-mirroring instance instance-name]
```

Description

Configure input packet properties for port mirroring.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

maximum-packet-length option introduced in Junos OS Release 9.6 for M120 and M320 routers only.

RELATED DOCUMENTATION

Configuring Port Mirroring

input-priority-map

IN THIS SECTION

- Syntax | 8716
- Hierarchy Level | 8716
- Description | 8716
- Required Privilege Level | 8716
- Release Information | 8717

Syntax

```
input-priority-map {
    ieee802.1p premium [ values ];
}
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile]
[edit interfaces interface-name ether-options ethernet-switch-profile ethernet-policer-profile]
```

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the input policer priority map to be applied to incoming frames on this interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

output-priority-map | 12007

input (Routing Options RIB)

IN THIS SECTION

- Syntax | 8717
- Hierarchy Level | 8717
- Description | 8718
- Options | 8718
- Required Privilege Level | 8718
- Release Information | 8718

Syntax

```
input filter-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options rib routing-table-name filter],
[edit routing-options rib routing-table-name filter]
```

Description

Specify the name of the input filter.

Options

filter-name—Name of the input filter.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Forwarding Table Filters

Applying Forwarding Table Filters

input (Sampling)

IN THIS SECTION

- Syntax | 8719
- Hierarchy Level | 8719
- Description | 8719
- Required Privilege Level | 8719
- Release Information | 8719

Syntax

```
input {
    max-packets-per-second number;
    maximum-packet-length bytes;
    rate number;
    run-length number;
}
```

Hierarchy Level

```
[edit forwarding-options sampling],
[edit forwarding-options sampling instance instance-name]
```

Description

Configure traffic sampling on a logical interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for sampling of MPLS traffic introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

input-scheduler-map

IN THIS SECTION

- Syntax | 8720
- Hierarchy Level | 8720
- Description | 8720
- Options | 8721
- Required Privilege Level | 8721
- Release Information | 8721

Syntax

```
input-scheduler-map map-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Associate a scheduler map with a physical or logical input interface. The `input-scheduler-map` and `input-traffic-control-profile` statements are mutually exclusive at the same hierarchy level.

NOTE: On some devices, CoS queuing and scheduling are enabled on the egress side but disabled on the ingress side by default. To enable ingress CoS, you must configure the `traffic-manager` statement with `ingress-and-egress` mode:

```
[edit chassis fpc slot-number pic pic-number]
traffic-manager mode ingress-and-egress;
```

Options

map-name—Name of scheduler map that you define at the `[edit class-of-service scheduler-maps]` hierarchy level.

`default`—The default scheduler mapping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring an Input Scheduler on an Interface

Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs

input-traffic-control-profile | **8726**

input-shaping-rate (Logical Interface)

IN THIS SECTION

- Syntax | **8722**
- Hierarchy Level | **8722**
- Description | **8722**
- Default | **8722**
- Options | **8722**
- Required Privilege Level | **8722**
- Release Information | **8723**

Syntax

```
input-shaping-rate (percent percentage | rate);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Configure input traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface. You can configure hierarchical shaping, meaning you can apply an input shaping rate to both the physical interface and the logical interface.

Default

If you do not include this statement, logical interfaces share a default scheduler. This scheduler has a committed information rate (CIR) that equals 0. (The CIR is the guaranteed rate.) The default scheduler has a peak information rate (PIR) that equals the physical interface shaping rate.

Options

percent *percentage*
: Shaping rate as a percentage of the available interface bandwidth.

- **Range:** 0 through 100 percent

rate
: Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000 bps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Ingress Hierarchical CoS on MIC and MPC Interfaces

Configuring Input Shaping Rates for Both Physical and Logical Interfaces

Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs

input-traffic-control-profile | 8726

input-shaping-rate (Physical Interface)

IN THIS SECTION

- Syntax | 8723
- Hierarchy Level | 8723
- Description | 8724
- Options | 8724
- Required Privilege Level | 8724
- Release Information | 8724

Syntax

```
input-shaping-rate rate;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name]
```

Description

Configure input traffic shaping by specifying the amount of bandwidth to be allocated to the physical interface. You can configure hierarchical shaping, meaning you can apply an input shaping rate to both the physical interface and the logical interface.

Options

rate Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000 bps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Input Shaping Rates for Both Physical and Logical Interfaces

Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs

input-traffic-control-profile | **8726**

input-three-color

IN THIS SECTION

Syntax | **8725**

- Hierarchy Level | 8725
- Description | 8725
- Options | 8725
- Usage Guidelines | 8725
- Required Privilege Level | 8725
- Release Information | 8726

Syntax

```
input-three-color policer-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number layer2-policer]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
layer2-policer]
```

Description

Apply a single-rate or two-rate three-color policer to the Layer 2 input traffic at the logical interface. The `input-three-color` and `input-policer` statements are mutually exclusive.

Options

policer-name—Name of the single-rate or two-rate three-color policer.

Usage Guidelines

See *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Two-Color and Three-Color Policers at Layer 2

Applying Layer 2 Policers to Gigabit Ethernet Interfaces

Configuring Gigabit Ethernet Policers

input-policer | **8713**

layer2-policer (EX) | **9654**

logical-interface-policer | **10069**

output-policer | **12001**

output-three-color | **12014**

input-traffic-control-profile

IN THIS SECTION

- Syntax | **8726**
- Hierarchy Level | **8727**
- Description | **8727**
- Options | **8727**
- Required Privilege Level | **8727**
- Release Information | **8727**

Syntax

```
input-traffic-control-profile profile-name shared-instance instance-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Apply an input traffic scheduling and shaping profile to the physical or logical interface. The `input-traffic-control-profile` and `input-scheduler-map` statements are mutually exclusive at the same hierarchy level.

NOTE: The `shared-instance` statement applies only to Gigabit Ethernet IQ2 and IQ2E PICs.

NOTE: On many platforms, CoS queuing and scheduling are enabled on the egress side but disabled on the ingress side by default. To enable ingress CoS on the ingress side, you must configure the `traffic-manager` statement with `ingress-and-egress` mode:

```
[edit chassis fpc slot-number pic pic-number]
traffic-manager mode ingress-and-egress;
```

Options

profile-name—Name of the traffic-control profile to be applied to this interface.

instance-name—Name of the shared scheduler and shaper to be applied to this interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs

input-shaping-rate (Logical Interface) | 8721

input-scheduler-map | 8720

traffic-control-profiles | 30009

input-traffic-control-profile-remaining

IN THIS SECTION

- Syntax | 8728
- Hierarchy Level | 8728
- Description | 8728
- Options | 8729
- Required Privilege Level | 8729
- Release Information | 8729

Syntax

```
input-traffic-control-profile-remaining profile-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit class-of-service interfaces interface-name interface-set interface-set-name]
```

Description

Apply an input traffic scheduling and shaping profile for the remaining traffic to the logical interface or interface set.

Options

profile-name—Name of the traffic-control profile for the remaining traffic to be applied to this interface or interface set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs

input-traffic-control-profile | 8726

input-vlan-map

IN THIS SECTION

- Syntax | 8729
- Hierarchy Level | 8730
- Description | 8730
- Required Privilege Level | 8730
- Release Information | 8730

Syntax

```
input-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
```

```
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP interfaces, 100-Gigabit Ethernet Type 5 PIC with CFP only as well as Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, define the rewrite profile to be applied to incoming frames on this logical interface.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: Connectivity fault management (CFM) sessions for all interfaces in which `input-vlan-map` is configured are supported only if the interface also has an explicit configuration for `output-vlan-map` as `output-vlan-map pop;`. See "output-vlan-map " on page 12021. This configuration is required for all the interfaces in the topology even when the CFM session is on that interface or on a different interface in the data path of the same topology.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`pop-pop`, `pop-swap`, `push-push`, `swap-push`, and `swap-swap` statements introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Stacking a VLAN Tag

output-vlan-map | 12021

Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support

input-vlan-map (Aggregated Ethernet)

IN THIS SECTION

- Syntax | 8731
- Hierarchy Level | 8731
- Description | 8732
- Required Privilege Level | 8732
- Release Information | 8732

Syntax

```
input-vlan-map {
    (pop | push | swap);
    tag-protocol-id tpid;
    vlan-id number;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Define the rewrite profile to be applied to incoming frames on this logical interface. On MX Series routers, this statement only applies to aggregated Ethernet inferfaces using Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and 100-Gigabit Ethernet Type 5 PIC with CFP.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Starting in Junos OS Release 17.3R1, input-vlan-map for outer vlan is supported for L2 circuit over aggregated Ethernet interfaces for QFX10000 Series switches.

RELATED DOCUMENTATION

Stacking a VLAN Tag

output-vlan-map (Aggregated Ethernet) | 12022

input-vlan-map (Dynamic Interfaces)

IN THIS SECTION

- Syntax | 8733
- Hierarchy Level | 8733
- Description | 8733
- Required Privilege Level | 8733
- Release Information | 8733

Syntax

```
input-vlan-map {
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    (push | swap);
    tag-protocol-id tpid;
    vlan-id number;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

Description

For dynamic interfaces, define the rewrite profile to be applied to incoming frames on this logical interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution

inputs (Analytics)

IN THIS SECTION

- Syntax | 8734
- Hierarchy Level | 8735
- Description | 8735
- Options | 8735
- Required Privilege Level | 8737
- Release Information | 8737

Syntax

```
inputs {
    analytics {
        parameters {
            generate-tags value;
            sample-frequency value;
            sensors path;
        }
    }
    input-ipfix {
        parameters {
            maximum-connections number;
            tcp-port port-number;
            vrf-name name;
        }
    }
    input-jti-ipfix {
        parameters {
            record-group group-name {
                record ipfix-record-name;
                reporting-interval seconds;
            }
        }
```

```
    }
}
```

Hierarchy Level

```
[edit services analytics agent service-agents agent-name]
```

Description

Configure parameters for a Network Telemetry Framework (NTF) service agent input plug-in. For each service agent instance, you can configure more than one input plug-in to push data to the output plug-in.

> **NOTE**: When you modify the input plug-in configuration of a service agent instance, the associated service agent daemon is restarted.

Options

analytics parameters
Configure parameters to collect data from Junos Telemetry Interface (JTI) sensors.

generate-tags *value*
(Optional) Enable tag generation.

- **Default:** Enabled

sample-frequency *value*
Specify the frequency interval (in seconds) at which the JTI sensor generates data to export to the data collector. Range is from 0 to 24 hours.

- **Default:** 5 seconds

sensors *file-path*
Specify the resource string associated with the JTI sensor for collecting JTI data from a specific resource. The format is a file path and must be entered exactly. For a list of available JTI resource string options, see the sensor configuration statement and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface) documentation.

input-ipfix parameters
Configure parameters for the IPFIX mediation service agent to gather and consolidate IPFIX records from downstream devices.

> **NOTE**: Any change you make to an existing `input-ipfix` plug-in configuration restarts the IPFIX service agent daemon to apply the changes.

> **NOTE**: Although each of the parameters has a default value, you must configure at least one of the parameters to enable the plug-in. If you configure only one parameter and want to use the default value, you must specify that value.

maximum-connections *number*

(Optional) Maximum number of TCP connections that the IPFIX mediator can support.

- **Range:** 1 through 500
- **Default:** 100

tcp-port *port-number*

(Optional) TCP port on the IPFIX mediator that receives TCP packets; the listening port.

- **Default:** 4739

vrf-name *name*

(Optional) Name of the VRF (routing instance) in which IPFIX packets are accepted.

- **Default:** default

input-jti-ipfix parameters

Configure parameters for the IPFIX mediation service agent to collect and report local sensor data from the BNG configured as an IPFIX mediator. For each group of records, the plug-in subscribes to the specific sensor data sets associated with each record.

When you remove a record group from the configuration, the sensor sets for the member records are unsubscribed. The template IDs for the associated IPFIX records are returned to the pool for re-use.

record *ipfix-record-name*

One of the following individual IPFIX records associated with a nonconfigurable set of local sensor data. See *Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector* for the sensors collected by each record.

address-pool-utilization	port-statistics

chassis-inventory	resource-utilization
chassis-power	subscriber-statistics
dhcpv4-server-stats	thermal
interface-metadata	uptime
interface-queue-statistics	

BEST PRACTICE: We recommend that you configure the `interface-metadata` record whenever you configure the `interface-queue-statistics` record. The metadata information is essential for understanding details about the subscriber whose queue statistics are being collected.

record-group *group-name*

Name of a group of IPFIX records that subscribes to the sensor data sets associated with the individual records that comprise the record group. You can configure a maximum of 10 record groups.

reporting-interval *seconds*

(Optional) Interval in seconds between reports for the subscribed sensor data. The interval applies to all records (and all sensor sets) in the record group.

- **Range:** 60 through 86,400 seconds
- **Default:** 900 seconds

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.3R1.

`input-jti-ipfix` option added in Junos OS Release 18.4R1 on MX Series routers.

`analytics` option added in Junos OS Release 18.4R1 on MX Series and PTX Series routers.

RELATED DOCUMENTATION

Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data

Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator

Configuring NTF Agent

IPFIX Mediation on the BNG

Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector

insert

IN THIS SECTION

- Syntax | 8738
- Description | 8738
- Options | 8738
- Required Privilege Level | 8739
- Release Information | 8739

Syntax

```
insert <statement-path> identifier1 (before | after) identifier2
```

Description

Insert an identifier in to an existing hierarchy.

Options

statement-path	(Optional) Path to the existing identifier.
identifier1	The new identifier to insert.
after	Place *identifier1* after *identifier2*.

before	Place *identifier1* before *identifier2*.
identifier2	The existing identifier.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Inserting a New Identifier in a Device Configuration

inspection-profile

IN THIS SECTION

- Syntax | 8739
- Hierarchy Level | 8740
- Description | 8740
- Options | 8740
- Required Privilege Level | 8740
- Release Information | 8740

Syntax

```
inspection-profile name
{
```

```
  vxlan name
    {
      policy-set policy-set;
      vni vni-name;
  }
 }
```

Hierarchy Level

```
[edit logical-systems name security tunnel-inspection],
[edit logical-systems name tenants name security tunnel-inspection],
[edit security tunnel-inspection],
[edit tenants name security tunnel-inspection]
```

Description

Configure the `tunnel-inspection` profile.

Options

name Tunnel-inspection profile name

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

install

IN THIS SECTION

- Syntax | 8741
- Hierarchy Level | 8741
- Description | 8741
- Options | 8741
- Required Privilege Level | 8742
- Release Information | 8742

Syntax

```
install {
    ignore-version-check;
    ignore-appid-failure;
}
```

Hierarchy Level

```
[edit security idp security-package]
```

Description

Configures the `install` command with ignore options.

Options

`ignore-version-check`—Ignores version check during the installation of security packages.

`ignore-appid-failure`—Ignores application ID failure during the installation of security packages.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

install-nexthop

IN THIS SECTION

- Syntax | 8742
- Hierarchy Level | 8742
- Description | 8743
- Options | 8743
- Required Privilege Level | 8743
- Release Information | 8743

Syntax

```
install-nexthop (except | fallback | lsp lsp-name | lsp-regex lsp-regular-expression | non-
labelled-nexthop  | non-lsp-nexthop | overlay-vxlan-interfaces | static-lsp lsp-name | static-
lsp-regex lsp-regular-expression)  | strict | strict-named-lsp;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options policy-statement policy-name term term-
name then],
[edit policy-options policy-statement policy-name term term-name then]
```

Description

Select a specific label-switched path (LSP), or select an LSP from a set of similarly named LSPs as the traffic destination for the configured community. Also can prevent the installation of any matching next hops.

Options

`except`—Prevent the installation of any matching next hops.

`fallback`—Backup option

`lsp` *`lsp-name`*—Configure a specific LSP.

`lsp-regex` *`lsp-regular-expression`*—Configure a range of similarly named LSPs. You can use the following wildcard characters when configuring an LSP regular expression:

- Asterisk (*)—Match any characters.
- Period (.)—Match any single digit.

`non-labelled-nexthop`—Next-hop without tag

`non-lsp-nexthop`—Next-hop with non-lsp

`overlay-vxlan-interfaces`—Next-hop for VXLAN interfaces

`static-lsp`—Next-hop static LSP name

`static-lsp-regex`—Next-hop static LSP name regular expression

`strict`—Do not use any other available next hops

`strict-named-lsp`—Do not use any other non-lsp next hops

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`overlay-vxlan-interfaces` statement introduced in Junos OS Release 22.4R1.

`overlay-vxlan-interfaces` statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

Configuring the Policy Statement for the Layer 2 Circuit Community

install (Protocols MPLS)

IN THIS SECTION

- Syntax | 8744
- Hierarchy Level | 8744
- Description | 8745
- Options | 8746
- Required Privilege Level | 8746
- Release Information | 8746

Syntax

```
install {
    destination-prefix <active>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name ingress]
[edit protocols source-packet-routing source-routing-path lsp-name]
```

Description

Associate one or more prefixes with an LSP. When the LSP is up, all the prefixes are installed as entries into the inet.3 or inet6.3 routing table.

The support of the `install` statement at the `[edit protocols source-packet-routing source-routing-path lsp-name]` is applicable for both colored and non-colored segment routing LSPs.

For static colored LSPs, when the install prefix is configured, a route similar to the tunnel ingress route is installed in the `inetcolor.0` or `inet6color.0` routing table.

On the other hand, for static non-colored LSPs, when the install prefix is configured, a route similar to the tunnel `to` route is installed in `inet.3` or `inet6.3` routing table.

You can use the "show route table" on page 23789 command to view the `install` routes for both colored and non-colored segment routing traffic-engineered LSPs.

NOTE: Take the following into consideration when configuring the `install` *destination-prefix* statement at the `[edit protocols source-packet-routing source-routing-path lsp-name]` hierarchy level:

- The install prefixes should be unique in a tunnel, and should not be the same as the tunnel `to` address. A commit check is done to ensure that the prefixes are unique.
- If two install prefixes are same across two different tunnels, then the gateways of the both tunnels are considered only if the segment lists are of the same resolution type. If the first hop resolution types vary, the route is not installed. In such cases, a system log message is generated to record the error.

For example:

```
[edit protocols source-packet-routing]
segment-list path-1 {
    hop-1 ip-address 172.0.12.2;
    hop-2 label 1000012;
    hop-3 label 1000013;
    hop-4 label 1000014;
}
source-routing-path lsp1 {
    to 10.10.10.1;
    install 20.20.20.2;
    install 30.30.30.3;
    primary {
```

```
            path-1;
        }
    }
```

The `inet.3` routing table has two additional routes (to 20.20.20.2 and 30.30.30.3) with next hops derived from the same segment list path-1 with same attributes as the `to` address route.

Options

active (Optional) Install the route into the inet.0 or inet6.0 routing table. This allows you to issue a `ping` or `traceroute` command on this address.

> **NOTE**: The `install` *`destination-prefix`* `active` statement is not supported on static LSPs. When the `install` *`destination-prefix`* `active` statement is configured for a static LSP, the MPLS routes do not get installed into the inet.0 routing table.

destination-prefix IPv4 or IPv6 address to associate with the LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit protocols source-packet-routing source-routing-path` *`lsp-name`*`]` hierarchy level introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table

install (Routing Options)

IN THIS SECTION

- Syntax | 8747
- Hierarchy Level | 8747
- Description | 8748
- Options | 8748
- Required Privilege Level | 8748
- Release Information | 8748

Syntax

```
(install | no-install);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options static (defaults | route)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name static (defaults | route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
(defaults | route)],
[edit logical-systems logical-system-name routing-options static (defaults | route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
(defaults | route)],
[edit routing-instances routing-instance-name routing-options static (defaults | route)],
[edit routing-options rib routing-table-name static (defaults | route)]
[edit routing-options static (defaults | route)]
```

Description

Configure whether Junos OS installs all static routes into the forwarding table. Even if you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols.

Options

install—Explicitly install all static routes into the forwarding table. Include this statement when configuring an individual route in the **route** portion of the `static` statement to override a **no-install** option specified in the **defaults** portion of the statement.

no-install—Do not install the route into the forwarding table, even if it is the route with the lowest preference.

- **Default: install**

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

static (Routing Options) | 28570

instance

IN THIS SECTION

Syntax | 8749
Hierarchy Level | 8749

- Description | 8749
- Options | 8750
- Required Privilege Level | 8750
- Release Information | 8750

Syntax

```
instance {
    instance-name {
        input {
            maximum-packet-length bytes;
            rate number;
            run-length number;
        }
        family (any | ccc| inet | inet6 | mpls | vpls) {
            output {
                interface interface-name {
                    next-hop address;
                }
                no-filter-check;
                server-profile server-profile-name;
            }
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring],
[edit routing-instances routing-instance-name forwarding-options port-mirroring]
```

Description

Configure a port-mirroring instance.

Options

port-mirroring-instance-name—Name of the port-mirroring instance.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3 (MX Series routers only).

Support extended to M120 and M320 routers in Junos OS Release 9.5.

maximum-packet-length and **ccc** options introduced in Junos OS Release 9.6 for M120 and M320 routers only.

`server-profile` option introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers only.

RELATED DOCUMENTATION

Configuring Port Mirroring

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

instance (Inline Monitoring)

IN THIS SECTION

- Syntax (Junos OS) | **8751**
- Syntax (Junos OS Evolved) | **8751**
- Hierarchy Level | **8751**
- Description | **8751**
- Options | **8752**

Required Privilege Level | 8752
Release Information | 8752

Syntax (Junos OS)

```
instance name {
    collector name;
    maximum-clip-length maximum-clip-length;
    template-name template-name;
}
```

Syntax (Junos OS Evolved)

```
instance name {
    collector name;
    controller name;
    maximum-clip-length maximum-clip-length;
    sampling-rate sampling-rate;
    template-name template-name;
}
```

Hierarchy Level

```
[edit services inline-monitoring]
```

Description

Configure inline-monitoring instance parameters. You can use these instances along with firewall filters to monitor different streams of traffic at different sampling rates from the same interface.

You can configure a maximum of sixteen (Junos OS) or seven (Junos OS Evolved) inline-monitoring instances with four collectors each.

Options

name Name of instance.

collector *name* Configure an collector for the inline-monitoring instance. See "collector" on page 4420 for more information.

controller (Junos OS Evolved only) Configure inline-monitoring services for packet redirects to the P4 controller or to the Routing Engine. See "controller" on page 4650 for more information.

maximum-clip-length Maximum packet length.

- **Range:** 64 through 126 bytes (Junos OS)

  Range: 64 through 256 bytes (Junos OS Evolved)

- **Default:** 126 bytes (Junos OS)

  Default: 128 bytes (Junos OS Evolved)

sampling-rate (Junos OS Evolved) Rate at which the traffic is sampled. In Junos OS, you do not specify the sampling rate here; you specify the sampling rate at the `[edit services inline-monitoring instance instance-name collector collector-name]` hierarchy level.

For example, if you specify 1000, then 1 packet out of every 1000 packets is sampled.

- **Range:** 1 through 32000000
- **Default:** 1

template-name Name of the template.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 19.4R1.

Statement introduced in Junos Evolved OS Release 22.2R1.

RELATED DOCUMENTATION

Understanding Inline Monitoring Services

instance (OAM)

IN THIS SECTION

- Syntax | 8753
- Hierarchy Level | 8753
- Description | 8753
- Required Privilege Level | 8753
- Release Information | 8754

Syntax

```
instance vpls-instance-name;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain name]
```

Description

Specify the VPLS instance of the default maintenance domain.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Maintenance Intermediate Points (MIPs)

maintenance-domain | 10287

instance (CASB)

IN THIS SECTION

- Syntax | 8754
- Hierarchy Level | 8755
- Description | 8755
- Options | 8755
- Required Privilege Level | 8756
- Release Information | 8756

Syntax

```
instance instance-name {
 application application-name;
    domain domain-name;
 instance-id instance-id;
 tag tag-value;
 type type;
}
```

Hierarchy Level

```
[edit security casb]
```

Description

Configure Cloud Access Security Broker (CASB) instance.

Configuring the CASB instance with specific parameters such as the application, instance ID, domain, and type allows controlled access by differentiating between corporate and non-corporate usage of cloud applications.

Options

instance *instance-name*

CASB instance name.

application *application-name*

Application name.

domain *domain-name*

An email domain address.

For example, acmecorp07.com is an organization domain. Then, Box, Dropbox, Google Docs, Salesforce, Microsoft OneDrive, SharePoint, Gmail, and Slack uses the same domain for all the users.

instance-id *id*

Each application has a unique web address (URL) to access it. This web address includes a special identifier (such as 'acmecorp07') that points to a specific instance of the application.

- Box—acmecorp07.app.box.com
- Microsoft OneDrive or SharePoint—acmecorp07ms-my.sharepoint.com
- Salesforce URLs—
- • acmecorp07.my.salesforce.com
  - acmecorp07.lightning.force.com
- Slack—acmecorp-zoy8730.slack.com.

Following are generic URLs and instance ID is not applicable.

- Dropbox—dropbox.com
- Google Docs—docs.google.com
- Gmail—mail.google.com

tag Specify tag value. You can specify one of the following options:

- Sanctioned—Application instances sanctioned by your organization.
- Unsanctioned—Application instances unsanctioned by your organization.

type Specify type of instance. You can specify one of the following option.

- Work
- Personal

You must configure the type for Dropbox. For other applications, this configuration is optional.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

instance (Ephemeral Database)

IN THIS SECTION

- Syntax | 8757
- Hierarchy Level | 8757
- Description | 8757
- Options | 8758
- Required Privilege Level | 8758

Release Information | 8759

Syntax

```
instance instance-name {
    cyclic-version version;
    resize {
        database-size-diff size MB;
        database-size-on-disk size MB;
    }
}
```

Hierarchy Level

```
[edit system configuration-database ephemeral]
```

Description

Enable an instance of the ephemeral configuration database.

The order in which the configuration lists the instances determines their priority when merging conflicting configuration statements from different instances into the configuration. The instances are listed in order from highest to lowest priority. In addition, user-defined instances of the ephemeral configuration database have higher priority than the default ephemeral database instance, which has higher priority than the static configuration database.

TIP: When you configure an ephemeral instance, you can specify its placement in the configuration by using the `insert` command instead of the `set` command.

Table 50 on page 8758 summarizes the maximum number of user-defined ephemeral database instances supported for different Junos OS variants and releases.

Table 50: Maximum Ephemeral Database Instances

Junos OS Variant and Release	Maximum Instances
Junos OS Release 18.1 or earlier	8
Junos OS Release 18.2R1 or later	7
Junos OS Evolved	8

Options

`cyclic-version` *`version`*

Number of versions of the ephemeral instance's database for which the system stores deleted configuration objects. After each commit, the device reclaims the space occupied by objects deleted in the previous database version relative to the current database version as determined by the cyclic version value. You can disable cyclic versioning by setting the value to 0.

Thus, starting with the 11th commit operation, the device, by default, reclaims the space occupied by objects deleted in version 1 of the database. The process continues with each subsequent commit operation.

- **Default:** 10
- **Range:** 0 through 4,294,967,295

instance-name

User-defined name for an instance of the ephemeral configuration database.

The instance name must contain only alphanumeric characters, hyphens, and underscores, and it must not exceed 32 characters in length. In addition, starting in Junos OS Release 17.1R3, 17.2R3, 17.3R3, 17.4R2, and 18.1R1, the name of an user-defined instance cannot be `default`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.2R2 and Junos OS Evolved Release 22.1R1.

`cyclic-version` statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Enabling and Configuring Instances of the Ephemeral Configuration Database

Example: Configuring the Ephemeral Configuration Database Using NETCONF

Understanding the Ephemeral Configuration Database

instance-export

IN THIS SECTION

- Syntax | 8759
- Hierarchy Level | 8759
- Description | 8760
- Options | 8760
- Required Privilege Level | 8760
- Release Information | 8760

Syntax

```
instance-export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
```

```
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Apply one or more policies to routes being exported from a routing instance.

Options

policy-names—Name of one or more export policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Routing Policies, Firewall Filters, and Traffic Policers User Guide

instance-import

IN THIS SECTION

- Syntax | 8761
- Hierarchy Level | 8761
- Description | 8761
- Options | 8761
- Required Privilege Level | 8761
- Release Information | 8761

Syntax

```
instance-import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Apply one or more policies to routes being imported into a routing instance.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Routing Policies, Firewall Filters, and Traffic Policers User Guide

instance (Port Mirroring)

IN THIS SECTION

- Syntax: MX, M and T Series Routers and SRX Series Firewalls | 8762
- Syntax: QFX Series (but not QFX10000) and EX4600 Switches and NFX Series Devices | 8762
- Syntax: QFX10000 Switches | 8763
- Hierarchy Level | 8763
- Description | 8764
- Required Privilege Level | 8764
- Release Information | 8764

Syntax: MX, M and T Series Routers and SRX Series Firewalls

```
instance instance-name {
    disable;
    input {
        rate number;
        maximum-packet-length bytes
    }
    family (any | inet | inet6 | vpls) {
        output {
            (next-hop-group group-name | interface interface-name);
        }
    }
}
```

Syntax: QFX Series (but not QFX10000) and EX4600 Switches and NFX Series Devices

```
instance instance-name{
    family
        ethernet-switching {
            output {
```

```
                interface interface-name {
                }
                no-filter-check;
                }
                vlan vlan-name {
                    no-tag;
                }
            }
        inet
            output {
                ip-address address {
                }
                routing-instance instance-name {
                    ip-address address {
                }
            }
        }
    }
```

Syntax: QFX10000 Switches

```
instance instance-name {
    input {
        rate number;
        run-length value;
    }
    output {
        dscp dscp-value;
        forwarding-class class-name;
        ip-destination-address address;
        ip-source-address address;
        policer policer-name;
    }
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring]
```

Description

Specify a port-mirroring configuration (instance). You do not specify an input for this instance. Instead, you, create a firewall filter that specifies the required traffic and directs it to the mirror. This is useful for controlling which types of traffic should be mirrored.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Port Mirroring and Analyzers

Configuring Port Mirroring

Examples: Configuring Port Mirroring for Local Analysis

Example: Mirroring Employee Web Traffic with a Firewall Filter

instance-role

IN THIS SECTION

- Syntax | 8765
- Hierarchy Level | 8765
- Description | 8765
- Options | 8765

Required Privilege Level | 8766
Release Information | 8766

Syntax

```
instance-role (access | nni);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Define the role of the routing instance in a Layer 2 Wholesale network.

Options

`access`—Defines the connectivity role of the routing instance in a Layer 2 Wholesale network as an access routing instance. When defined for this role, the same process occurs as in a Layer 3 Wholesale network—when the first packet is received from a given client, authentication for the client initiates with an external entity (for example, RADIUS). If authentication is successful, a logical interface is created with the appropriate outer and inner VLAN tags for that client.

`nni`—Defines the connectivity role of the routing instance in a Layer 2 Wholesale network as a network to network interface (NNI) routing instance. When defined for this role, only outer VLAN tags are learned. In addition, when the NNI routing instance receives a response from the ISP, the packets are forwarded to the appropriate client, provided the packet has the same two tags that were verified during authentication.

NOTE: If you connect an access node or MSAN device to a router participating in the Layer 2 Wholesale network in an NNI role, you must create a new routing instance of type `l2backhaul-vpn` with an instance role of type `access` for that connection.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers

Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers

Subscriber Management Overview

instance (Sampling)

IN THIS SECTION

- Syntax | 8766
- Hierarchy Level | 8768
- Description | 8768
- Required Privilege Level | 8768
- Release Information | 8768

Syntax

```
instance instance-name {
    disable;
    family (bridge | inet | inet6 | mpls | vpls) {
        disable;
        output {
            aggregate-export-interval seconds;
```

```
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            flow-server hostname {
                aggregation {
                    autonomous-system;
                    destination-prefix;
                    protocol-port;
                    source-destination-prefix {
                        caida-compliant;
                    }
                    source-prefix;
                }
                autonomous-system-type (origin | peer);
                dscp dscp-value;
                forwarding-class class-name;
                (local-dump | no-local-dump);
                port port-number;
                source-address address;
                version format;
                version9 {
                    template template-name;
                }
                version-ipfix {
                    template template-name;
                }
            }
            interface interface-name {
                engine-id number;
                engine-type number;
                source-address address;
            }
            inline-jflow {
                source-address address;
                flow-export-rate rate;
            }
        }
    }
    input {
        rate number;
        run-length number;
        max-packets-per-second number;
        maximum-packet-length bytes;
```

```
    }
}
```

Hierarchy Level

```
[edit forwarding-options sampling]
```

Description

Configure a sampling instance.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.

Support for IPv6 addresses for the `flow-server` *address* and `source-address` *address* statements within sampling instances introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

> *Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches*

instance-shared

IN THIS SECTION

- Syntax | 8769
- Hierarchy Level | 8769
- Description | 8769
- Required Privilege Level | 8770
- Release Information | 8770

Syntax

```
instance-shared;
```

Hierarchy Level

```
[edit firewall family protocol-family-name filter filter-name],
[edit logical-systems logical-system-name firewall family protocol-family-name filter filter-
name]
```

Description

Specify to share the firewall filter across multiple routing instances. By default, firewall filters are not automatically shared across multiple instances. You can configure both shared and nonshared firewall filters on the same routing device. This statement can be used only when network services for the device are configured with enhanced IP mode.

The following protocol families are supported: Bridge, IPv4, IPv6, Layer 2 CCC, MPLS, and VPLS.

NOTE: Only Modular Port Concentrators (MPCs) are supported.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Guidelines for Configuring Firewall Filters

network-services | 11320

instance (Traffic Load Balancer)

IN THIS SECTION

- Syntax | 8770
- Hierarchy Level | 8771
- Description | 8772
- Options | 8772
- Required Privilege Level | 8772
- Release Information | 8772

Syntax

```
instance instance-name {
    client-interface client-interface;
    client-vrf client-vrf;
    group group-name {
        health-check-interface-subunit health-check-interface-subunit;
        network-monitoring-profile profile-name;
```

```
        real-service-rejoin-options no-auto-rejoin;
        real-services [ server-list ];
        <routing-instance routing-instance>;
    }
    interface interface-name;
    real-service real-service  {
        address server-ip-address;
        admin-down;
    }
    server-inet-bypass-filter server-inet-bypass-filter ;
    server-inet6-bypass-filter server-inet6-bypass-filter ;
    server-interface server-interface;
    server-vrf server-vrf-name;
    virtual-service virtual-service-name {
        address virtual-ip-address;
        group group-name;
        load-balance-method {
            hash {
                hash-key method;
            }
            random;
        }
        mode (layer2-direct-server-return | direct-server-return | translated);

        <routing-instance routing-instance-name>;
        <routing-metric route-metric>;
        server-interface server-interface;
        service service-name {
            protocol (udp | tcp);
            server-listening-port port;
            virtual-port virtual-port;
        }
    }
}
```

Hierarchy Level

```
[edit services traffic-load-balance]
```

Description

Configure a Traffic Load Balancer instance.

Options

client-interface *client-interface*	—For translated mode, client interface where the implicit filter is installed to direct the traffic in the forward direction.
client-vrf *client-vrf*	Use the specified name of the routing instance in which the data traffic in the reverse direction is routed to the clients.
instance *instance-name*	Identifier (text string) for a TLB configuration.
server-inet-bypass-filter *server-inet-bypass-filter*	Name of the firewall filter from which the terms are referenced and added to the server-side implicit filters. This enables the operator to bypass reverse (RIP to VIP) translation of IPv4 traffic.
server-inet6-bypass-filter *server-inet6-bypass-filter*	Name of the firewall filter from which the terms are referenced and added to the server-side implicit filters. This enables the operator to bypass reverse (RIP to VIP) translation of IPv6 traffic.
server-interface *server-interface*	For translated mode, specifies the server interfaces where the server filters are implicitly installed to direct the return traffic to the load balancing next hop.
server-vrf *server-vrf-name*	The routing instance in which the data traffic in the forward direction is routed to the servers

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Traffic Load Balancer Overview

Configuring TLB

instance-type

IN THIS SECTION

- Syntax | 8773
- Hierarchy Level | 8773
- Description | 8773
- Options | 8774
- Required Privilege Level | 8776
- Release Information | 8776

Syntax

```
instance-type type;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Define the type of routing instance.

> **CAUTION**: We strongly recommend that if you change an `instance-type` referenced under a firewall filter, for example, from `virtual-router` to `forwarding`, make the change during a maintenance window, as follows:
>
> 1. Deactivate the routing instance.
> 2. Change the `instance-type`.
> 3. Activate the routing instance.
>
> This is not required if you are configuring the `instance-type` for the first time.

Options

type—Can be one of the following:

- `evpn`—Enable an Ethernet VPN (EVPN) on the routing instance.
  - On Junos OS Evolved, `instance-type evpn` is not supported. You can configure an EVPN routing instance using `instance-type mac-vrf` with a routing protocol configuration of `protocols evpn`.

    ```
    set routing-instances $EVPN_INSTANCES$ instance-type mac-vrf protocols evpn
    ```

- `evpn-vpws`—Enable an Ethernet VPN (EVPN) Virtual Private Wire Service (VPWS) on the routing instance.
- `forwarding`—Provide support for filter-based forwarding, where interfaces are not associated with instances. All interfaces belong to the default instance. Other instances are used for populating RPD learned routes. For this instance type, there is no one-to-one mapping between an interface and a routing instance. All interfaces belong to the default instance inet.0.
- `l2backhaul-vpn`—Provide support for Layer 2 wholesale VLAN packets with no existing corresponding logical interface. When using this instance, the router learns both the outer tag and inner tag of the incoming packets, when the `instance-role` statement is defined as `access`, or the outer VLAN tag only, when the `instance-role` statement is defined as `nni`.
- `l2vpn`—Enable a Layer 2 VPN on the routing instance. You must configure the `interface`, `route-distinguisher`, `vrf-import`, and `vrf-export` statements for this type of routing instance.
- `layer2-control` —(MX Series routers only) Provide support for RSTP or MSTP in customer edge interfaces of a VPLS routing instance. This instance type cannot be used if the customer edge interface is multihomed to two provider edge interfaces. If the customer edge interface is multihomed to two provider edge interfaces, use the default BPDU tunneling.

- `mac-vrf`—Enable configuring multiple customer-specific EVPN instances (EVIs) of this type, each of which can support a different EVPN service type. You can have customer-specific virtual routing and forwarding (VRF) tables with MAC addresses on each Juniper Networks device that serves as a virtual tunnel endpoint (VTEP) in the EVPN-VXLAN network. See "mac-vrf" on page 10276 for more on this type of EVPN instance.

- `mpls-forwarding`—(MX Series routers only) Allow filtering and translation of route distinguisher (RD) values in IPv4 and IPv6 VPN address families on both routes received and routes sent for selected BGP sessions. In particular, for Inter-AS VPN Option-B networks, this option can prevent the malicious injection of VPN labels from one peer AS boundary router to another.

- `mpls-internet-multicast`—(EX Series, M Series, MX Series, and T Series routers only) Provide support for ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP or next-generation MVPN.

- `no-forwarding`—This is the default routing instance. Do not create a corresponding forwarding instance. Use this routing instance type when a separation of routing table information is required. There is no corresponding forwarding table. All routes are installed into the default forwarding table. IS-IS instances are strictly nonforwarding instance types.

- `virtual-router`—Enable a virtual router routing instance. This instance type is similar to a VPN routing and forwarding instance type, but used for non-VPN-related applications. You must configure the `interface` statement for this type of routing instance. You do not need to configure the `route-distinguisher`, `vrf-import`, and `vrf-export` statements.

- `virtual-switch`—(Not supported on QFX5xxx switches running either Junos OS or Junos OS Evolved) Provide support for Layer 2 bridging. Use this routing instance type to isolate a LAN segment with its Spanning Tree Protocol (STP) instance and to separate its VLAN identifier space. We also support this routing instance type for EVPN instances.

> **NOTE**: When you want to configure EVPN protocol settings in a `virtual-switch` instance, you must do so at the same time you configure the `virtual-switch` instance. Otherwise the device has problems adding EVPN Type 2 (MAC-IP) route entries in the EVPN routing tables.
>
> If you need to update an existing `virtual-switch` instance in an active configuration to add EVPN protocol settings (`set … protocols evpn`), you must:
>
> 1. Deactivate the `virtual-switch` instance configuration.
> 2. Add the EVPN protocol statements to the `virtual-switch` instance configuration.

3. Reactivate the updated `virtual-switch` instance configuration with the EVPN protocol updates.

- `vpls`—Enable VPLS on the routing instance. Use this routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN. You must configure the `interface`, `route-distinguisher`, `vrf-import`, and `vrf-export` statements for this type of routing instance.
  - On Junos OS Evolved, `instance-type vpls` is not supported. You can configure a VPLS routing instance using `instance-type virtual-switch` with a routing protocol configuration of `protocols vpls`.

```
set routing-instances $VPLS_INSTANCE$ instance-type virtual-switch protocols vpls
```

- `vrf`—VPN routing and forwarding (VRF) instance. Provides support for Layer 3 VPNs, where interface routes for each instance go into the corresponding forwarding table only. Required to create a Layer 3 VPN. Create a VRF table (*instance-name*.`inet.0`) that contains the routes originating from and destined for a particular Layer 3 VPN. For this instance type, there is a one-to-one mapping between an interface and a routing instance. Each VRF instance corresponds with a forwarding table. Routes on an interface go into the corresponding forwarding table. You must configure the `interface`, `route-distinguisher`, `vrf-import`, and `vrf-export` statements for this type of routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`virtual-switch` and `layer2-control` options introduced in Junos OS Release 8.4.

`mpls-internet-multicast` option introduced in Junos OS Release 11.1 for the EX Series, M Series, MX Series, and T Series.

`evpn` option introduced in Junos OS Release 13.2 for MX 3D Series routers.

`evpn` option introduced in Junos OS Release 17.3 for the QFX Series.

`forwarding` option introduced in Junos OS Release 14.2 for the PTX Series.

`mpls-forwarding` option introduced in Junos OS Release 16.1 for the MX Series.

`evpn-vpws` option introduced in Junos OS Release 17.1 for MX Series routers.

`mac-vrf` option introduced in Junos OS Release 20.4 and Junos OS Evolved Release 21.2R1.

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

`evpn-vpws` option introduced in cRPD Release 20.3R1.

RELATED DOCUMENTATION

- *Configuring EVPN Routing Instances*
- *Configuring EVPN Routing Instances on EX9200 Switches*
- Configuring Virtual Router Routing Instances
- Example: Configuring Filter-Based Forwarding on the Source Address
- Example: Configuring Filter-Based Forwarding on Logical Systems
- *MAC-VRF Routing Instance Type Overview*

integer

IN THIS SECTION

- Syntax | **8777**
- Hierarchy Level | **8778**
- Description | **8778**
- Required Privilege Level | **8778**
- Release Information | **8778**

Syntax

```
integer {
    equals value;
    greater-than value;
    less-than value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format]
```

Description

Specify the custom AVP attribute's format as an integer and the value to match for the incoming RADIUS request from the IP-based subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers*
- *Understanding Selection of Properties for an IP-Based TDF Subscriber*
- *Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber*
- *IP-Based Subscriber Setup Overview*

intelligent-prescreening

IN THIS SECTION

- Syntax | 8779
- Hierarchy Level | 8779
- Description | 8779
- Required Privilege Level | 8780
- Release Information | 8780

Syntax

```
intelligent-prescreening;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name scan-
options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name scan-
options]
```

Description

Enable intelligent prescreening.

Intelligent prescreening tells the antivirus module to begin scanning a file much earlier. In this case, the scan engine uses the first packet or the first several packets to determine if a file could possibly contain malicious code. The scan engine does a quick check on these first packets and if the scan engine finds that it is unlikely that the file is infected, it then determines that it is safe to bypass the normal scanning procedure.

You can disable intelligent prescreening with the `no-intelligent-prescreening` statement.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1x49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1x49-D10 onwards.

intelligent-offload-tunable (IDP)

IN THIS SECTION

- Syntax | 8781
- Hierarchy Level | 8781
- Description | 8781
- Options | 8781
- Required Privilege Level | 8781
- Release Information | 8782

Syntax

```
service name {
    offload-limit bytes;
}
```

Hierarchy Level

```
[edit logical-systems name security idp sensor-configuration global intelligent-offload-tunable],
[edit logical-systems name tenants name security idp sensor-configuration global intelligent-
offload-tunable],
[edit security idp sensor-configuration global intelligent-offload-tunable],
[edit tenants name security idp sensor-configuration global intelligent-offload-tunable]
```

Description

Protocol-Specific Intelligent-Offload Configuration feature in Intrusion Detection and Prevention (IDP) systems allows you to tailor inspection depth limits for different protocols, enhancing both performance and security. By configuring individual offload limits for protocols such as SSH and FTP, you can optimize resource usage and ensure more efficient session inspections.

Options

name Service name

offload-limit Service offload limit in bytes (bytes)

- **Range:** 0 through 4294967295.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

inter-area-prefix-export

IN THIS SECTION

- Syntax | 8782
- Hierarchy Level | 8782
- Description | 8783
- Options | 8783
- Required Privilege Level | 8783
- Release Information | 8783

Syntax

```
inter-area-prefix-export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf3 area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ip4-unicast | ipv4-multicast | ipv6-multicast) area area-id],
[edit protocols ospf3 area area-id],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id],
[edit routing-instances routing-instance-name protocols ospf3 area area-id],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
muticast | ipv6-multicast) area area-id]
```

Description

Apply an export policy for OSPFv3 to specify which interarea prefix link-state advertisements (LSAs) are flooded into an area.

Options

policy-name—Name of a policy configured at the `[edit policy-options policy-statement policy-name term term-name]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Import and Export Policies for Network Summaries Overview

definition (Protocols OSPF) | **5234**

Routing Policies, Firewall Filters, and Traffic Policers User Guide

inter-area-prefix-import

IN THIS SECTION

- Syntax | **8784**
- Hierarchy Level | **8784**

- Description | 8784
- Options | 8784
- Required Privilege Level | 8785
- Release Information | 8785

Syntax

```
inter-area-prefix-import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf3 area area-id],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast) area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id],
[edit protocols ospf3 area area-id],
[edit protocols ospf3 realm (ip4-unicast | ipv4-multicast | ipv6-multicast)], area area-id],
[edit routing-instances routing-instance-name protocols ospf3 area area-id],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast) area area-id]
```

Description

Apply an import policy for OSPFv3 to specify which routes learned from an area are used to generate interarea prefixes into other areas.

Options

policy-name—Name of a policy configured at the `[edit policy-options policy-statement policy-name term term-name]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Import and Export Policies for Network Summaries Overview

inter-area-prefix-export | 8782

Routing Policies, Firewall Filters, and Traffic Policers User Guide

inter-as (Routing Instances)

IN THIS SECTION

- Syntax | 8786
- Hierarchy Level | 8786
- Description | 8786
- Options | 8787
- Required Privilege Level | 8787
- Release Information | 8787

Syntax

```
inter-as{
    ingress-replication {
        create-new-ucast-tunnel;
        label-switched-path-template {
            (default-template lsp-template-name);
        }
    }
    inter-region-segmented {
        fan-out <leaf-AD routes>);
        threshold <kilobits>);
    }
    ldp-p2mp;
    rsvp-te {
        label-switched-path-template {
            (default-template lsp-template-name);
        }
    }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name provider-tunnel]
```

Description

These statements add Junos support for segmented RSVP-TE provider tunnels with next-generation Layer 3 multicast VPNs (MVPN), that is, Inter-AS Option B. Inter-AS (autonomous-systems) support is required when an L3VPN spans multiple ASes, which can be under the same or different administrative authority (such as in an inter-provider scenario). Provider-tunnels (p-tunnels) segmentation occurs at the Autonomous System Border Routers (ASBR). The ASBRs are actively involved in BGP-MVPN signaling as well as data-plane setup.

In addition to creating the Intra-AS p-tunnel segment, these Inter-AS configurations are also used for ASBRs to originate the Inter-AS Auto Discovery (AD) route into Exterior Border Gateway Protocol (eBGP).

Options

ingress-replication

Select the ingress replication tunnel for further configuration.

- Choose `create-new-ucast-tunnel` to create a new unicast tunnel for ingress replication.
- Choose `label-switched-path` to create a point-to-point LSP unicast tunnel, and then choose `label-switched-path-template` to use the default template and parameters for dynamic point-to-point LSP.

inter-region-segmented

Select whether Inter-Region Segmented LSPs are triggered by `threshold rate`, or `fan-out`, or both. Supported tunnels are PIM-SSM and PIM-ASM; Inter-region-segmented cannot be set for PIM tunnels.

- Choose `fan-out` and then specify the number (from 1 to 10,000) of remote Leaf-AD routes to use as a trigger point for segmentation.
- Choose `threshold` and then specify a data threshold rate (from 0 to 1,000,000 kilobytes per second) to use to use as a trigger point for segmentation.

ldp-p2mp

Select to use LDP point-to-multipoint LSP for flooding; LDP P2MP must be configured in the master routing instance.

rsvp-te

Select to use RSVP-TE point-to-multipoint LSP for flooding.

- Choose `label-switched-path-template` to use the default template and parameters for dynamic point-to-point LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

BGP-MVPN Inter-AS Option B Overview

inter-domain

IN THIS SECTION

- Syntax | 8788
- Hierarchy Level | 8788
- Description | 8788
- Required Privilege Level | 8788
- Release Information | 8789

Syntax

```
inter-domain;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path label-switched-path-name],
[edit protocols mpls label-switched-path label-switched-path-name]
```

Description

Allows the router to search for routes in the IGP database. You need to configure this statement on routers that might be unable to locate a path using intra-domain CSPF (by looking in the traffic engineering database (TED). When you configure inter-area or inter-AS LSPs, the `inter-domain` statement is required.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring an LSP Across ASs

label-switched-path (Protocols MPLS) | 9585

inter-region

IN THIS SECTION

- Syntax | 8789
- Hierarchy Level | 8789
- Description | 8790
- Options | 8790
- Required Privilege Level | 8790
- Release Information | 8790

Syntax

```
inter-region [
    template template-name;
    no-inter-region-segmentation;
}
```

Hierarchy Level

```
[edit routing-instances instance-name provider-tunnel]
```

Description

Specify template to be used in inter-region segmentation. The tunnels types used will be based on the template mentioned. If a virtual routing and forwarding (VRF) in the area boundary router (ABR) does not want to participate in the inter-region segmentation then `no-inter-region-segmentation` can be specified.

Options

no-inter-region-segmentation	Specify to avoid the template mapping an ABR that is not involved in the inter-region segmentation.
template *template-name*	Specify the template to be used in the inter-region segmentation.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview

Example: Configuring Segmented Inter-Area P2MP LSP

Configuring Segmented Inter-Area P2MP LSP

inter-region-segmented | **8791**

inter-region-template | **8793**

region | **14262**

template(Protocols MVPN) | **29092**

inter-region-segmented

IN THIS SECTION

- Syntax | 8791
- Hierarchy Level | 8791
- Description | 8791
- Options | 8792
- Required Privilege Level | 8792
- Release Information | 8792

Syntax

```
inter-region-segmented {
    fan-out fan-out value;
    threshold rate-value;
}
```

Hierarchy Level

```
[edit routing-instances instance-name provider-tunnel selective group multicast IP address
source source IP-address],
[edit routing-instances instance-name provider-tunnel selective group multicast IP address
wildcard-source],
[edit routing-instances instance-name provider-tunnel selective wildcard-group-inet wildcard-
source],
[edit routing-instances instance-name provider-tunnel selective wildcard-group-inet6 wildcard-
source]
```

Description

Configure segmentation of inter-region on an ingress PE router or transit area boundary routers (ABRs) by specifying only a fan-out factor or threshold rate or both. The inter-region segmentation allows the PE router to signal segmented S-PMSI A-D routes. If the segmentation is based on the threshold, then

the MVPN starts migrating the traffic flow to a segmented S-PMSI when the threshold configured for the selective tunnel reaches the threshold rate. The threshold value can be configured for S-PMSI with LDP-P2MP, IR, and RSVP-TE tunnel types. If the fan-out value is configured, then the traffic flow is moved to segmented S-PMSI when the number of leaf A-D routes per S-PMSI meets the configured fan-out value. The fan-out value can be configured for wildcard S-PMSI, and S-PMSI.

If both threshold and fan-out values are set for S-PMSI with LDP-P2MP, IR, and RSVP-TE tunnel types, then the traffic flow to the segmented inter-region happens only when the traffic rate multiplied by the number of leaf A-D routes exceeds the threshold value.

Options

fan-out *fan-out-value*

Specify the number of remote leaf A-D route values for fan-out. Once the number of leaf A-D routes per S-PMSI meets the configured fan-out value, the inter-region segmentation is triggered and the traffic flow is directed to segmented S-PMSI.

- **Range:** 1 through 10000

threshold *rate-value*

Specify the data threshold rate to trigger inter-region segmentation.

> **NOTE**: This option is available only at the following hierarchy levels:
>
> `[edit routing-instances instance-name provider-tunnel selective group multicast IP address source source IP-address]`,
> `[edit routing-instances instance-name provider-tunnel selective group multicast IP address wildcard-source]`

- **Range:** 0 through 1000000 Kilobits

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview

Example: Configuring Segmented Inter-Area P2MP LSP

Configuring Segmented Inter-Area P2MP LSP

all-regions | **2157**

inter-region | **8789**

inter-region-template | **8793**

region | **14262**

template(Protocols MVPN) | **29092**

inter-region-template

IN THIS SECTION

- Syntax | **8793**
- Hierarchy Level | **8794**
- Description | **8794**
- Required Privilege Level | **8795**
- Release Information | **8795**

Syntax

```
inter-region-template {
    template template-name {
        all-regions {
            incoming;
            ingress-replication {
                create-new-ucast-tunnel;
                label-switched-path {
                    label-switched-path-template (Multicast) {
                        (default-template | lsp-template-name);
                    }
                }
```

```
            }
        ldp-p2mp;
            rsvp-te {
                label-switched-path-template (Multicast) {
                    (default-template | lsp-template-name);
                }
            static-lsp static-lsp;
        region region-name{
            incoming;
            ingress-replication {
                create-new-ucast-tunnel;
                label-switched-path {
                    label-switched-path-template (Multicast){
                        (default-template | lsp-template-name);
                    }
                }
            }
            ldp-p2mp;
            rsvp-te {
                label-switched-path-template (Multicast) {
                    (default-template | lsp-template-name);
                }
                static-lsp static-lsp;
            }
        }
    }
}
```

Hierarchy Level

```
[edit protocols mvpn]
```

Description

Define MVPN inter-region tunnel mapping template. A template is configured on a transit area boundary router (ABR) to define the tunnel type to be used for a given region. There are four types of tunnels that can be used:

- Incoming — Indicates that the tunnel type used across the ABR remains unchanged.

- Ingress-replication — A tunnel type used for passing multicast traffic between routers through the MPLS cloud, or between PE routers when using MVPN. The ingress replication provider tunnel uses MPLS point-to-point LSPs to create the multicast distribution tree.
- Ldp-p2mp — Specify a point-to-multipoint provider tunnel with LDP signaling for an MBGP MVPN.
- Rsvp-te — Configure RSVP tunnels to multicast traffic flooding using point-to-multipoint LSPs.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview

Example: Configuring Segmented Inter-Area P2MP LSP

Configuring Segmented Inter-Area P2MP LSP

all-regions | **2157**

inter-region-segmented | **8791**

region | **14262**

template(Protocols MVPN) | **29092**

inter-switch-link

IN THIS SECTION

Syntax | **8796**

Hierarchy Level | **8796**

- Description | 8796
- Required Privilege Level | 8796
- Release Information | 8796

Syntax

```
inter-switch-link vlan members primary-vlan-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family ethernet-switching interface-mode
trunk]
```

Description

Use this configuration statement when a private VLAN (PVLAN) spans multiple switches. The Inter-Switch Link protocol (ISL) must be configured on a trunk port of the primary VLAN in order to connect the switches composing the PVLAN to each other. You do not need to configure an ISL when a PVLAN is configured on a single switch.

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Creating a Private VLAN Spanning Multiple EX Series Switches with ELS Support (CLI Procedure)

interconnect

IN THIS SECTION

- Syntax | 8797
- Hierarchy Level | 8798
- Description | 8798
- Options | 8800
- Required Privilege Level | 8800
- Release Information | 8800

Syntax

```
interconnect {
    domain-path-id
        domain-id-string
    encapsulation
       (protocols evpn) {
        mpls;
        vxlan;
    }
    esi identifier {
        all-active;
        df-election-type {
            mod;
            preference {
                value value (0..65535);
            }
        }
    }
    interconnected-vlan-list ([vlan-list] | all);
    interconnected-vni-list ([vni-list] | all);
    irb-symmetric-routing vni vni;
    route-distinguisher [number:id | ip-address:id];
    vrf-export (policy-name | [policy-names]);
    vrf-import (policy-name | [policy-names]);
    vrf-target {
```

```
        community;
        import community;
        export community;
    }
}
```

Hierarchy Level

```
[edit protocols evpn],
[edit routing-instances routing-instance-name protocols evpn]
```

Description

Use this statement to enable seamless interconnection of:

- EVPN-VXLAN points of delivery (PODs) in a data center.
- Different EVPN-VXLAN data centers. This use case is also known as data center interconnect (DCI).

You can use this statement on devices that support seamless EVPN-VXLAN stitching to interconnect PODs or data centers through either:

- An EVPN-VXLAN WAN network—Use the `encapsulation vxlan` option.
- An EVPN-MPLS WAN network—Use the `encapsulation mpls` option.

The devices that enable the interconnection are called interconnection gateways (iGW devices). If possible, you should configure redundant iGW devices with all-active multihoming on each side of the interconnection, which we call multihoming peer iGW devices. The iGW devices are usually spine layer devices, or they are separate devices in a super-spine layer.

> **NOTE**: You must pair PTX10000 Series devices functioning as multihomed iGWs with other PTX10000 Series devices in the same DCI, due to their handling of broadcast, unknown unicast, and multicast (BUM) traffic. Remote gateways can be any device.

On each iGW device, you set up an EVPN routing instance using the default switching instance. Alternatively, you can configure the EVPN routing instance using `instance-type virtual-switch` or `evpn`, depending on the instance types the device supports. In the instance, specify these elements:

- Two sets of route distinguishers (RDs) and route targets.

The first set is the local RD and route target for the POD or data center network. The second set, which you specify with this `interconnect` statement, is for the WAN network.

> **NOTE**: The interconnect RD must be different from the local RD in the routing instance. When you configure DCI with seamless stitching for EVPN Type 2 routes, the device throws a commit error if you try to configure the same value for the interconnect RD and the local RD.
>
> To enforce this condition for DCI seamless stitching with EVPN Type 5 routes as well, you also see a commit error with Junos OS and Junos OS Evolved Releases starting in 22.4R2 and 23.1R1.

- An interconnection Ethernet segment identifier (iESI) for all-active multihoming among the peer iGW devices on each side of the interconnection.

  You configure the iESI using the `esi` option of this `interconnect` statement. You must use the same iESI value on all peer iGW devices in the same POD or data center, and unique iESI values for each POD or data center.

  > **NOTE**: When configuring the interconnection ESI, keep in mind that the seamless EVPN-VXLAN stitching feature does not support `single-active` mode, the `auto-derive lacp` option, or the `source-bmac` option.

- (Required for iGW devices with interconnect gateway multihoming) A list of the iGW device's peer multihoming iGW devices.

  On each iGW device, configure the "interconnect-multihoming-peer-gateways" on page 8806 statement at the global [edit protocols evpn] hierarchy level to specify the IP loopback addresses of that device's peer multihoming iGW devices. This option simplifies control plane route handling to optimize forwarding among the peer iGW devices, and helps avoid loops during EVPN configuration changes.

- VXLAN network identifiers (VNIs) that the iGW device translates when different VNIs are configured for the same VLAN in the PODs or data centers, and WAN. For more information about setting up VNI translation in DCI use cases, see "translation-vni" on page 30080 and "interconnected-vni-list" on page 8809.

After you configure these elements, the EVPN control plane stitches the EVPN routes from the POD or data center network and from the WAN network into a single MAC forwarding table.

Options

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Option `interconnected-vni-list` introduced in Junos OS Release 20.3R1.

Options `interconnected-vlan-list` and `encapsulation mpls` introduced in Junos OS Release 21.4R1.

interconnect-device (Aliases)

IN THIS SECTION

- Syntax | 8800
- Hierarchy Level | 8801
- Description | 8801
- Options | 8801
- Required Privilege Level | 8801
- Release Information | 8801

Syntax

```
interconnect-device interconnect-device-name {
    assigned-interconnect-device-name;
}
```

Hierarchy Level

```
[edit fabric aliases]
```

Description

(QFabric systems only) Specify the mapping of user-defined names to QFabric system Interconnect devices.

NOTE: The following rules apply to QFabric component alias naming:

- Alias names must use alphabetic (**A** through **Z** and **a** through **z**), numeric (**0** through **9**), or dash (**-**) characters.
- The maximum length of an alias name is 30 characters.
- Alias names are case sensitive. For example, **MY-NG-1** and **my-ng-1** refer to different components.
- You cannot use the reserved names `all`, `fabric`, or `director-group` as an alias name.

Options

interconnect-device-name—Specify a user-defined name for a QFabric system Interconnect device.

assigned-interconnect-device-name—Specify the Interconnect device identifier or name that has been provided. Identifiers are usually auto-generated by the Interconnect software, and names are usually provided by the administrator of the default partition.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Aliases for the QFabric System

Understanding Interconnect Devices

interconnect-device (Chassis)

IN THIS SECTION

- Syntax | 8802
- Hierarchy Level | 8803
- Description | 8803
- Required Privilege Level | 8803
- Release Information | 8803

Syntax

```
interconnect-device {
    alarm {
        (ethernet (Alarm) | management-ethernet) {
            link-down (red | yellow | ignore);
        }
    }
    container-devices {
        device-count number;
    }
    craft-lockout {
        alarm {
            interface-type {
                link-down (red | yellow | ignore);
            }
        }
        container-devices {
            device-count number;
        }
        fpc slot {
```

```
            power (on | off);
        }
        routing-engine {
            on-disk-failure {
                disk-failure-action (halt | reboot);
            }
        }
    }
    fpc slot {
        power (on | off);
    }
    routing-engine {
        on-disk-failure {
            disk-failure-action (halt | reboot);
        }
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure properties specific to a QFabric system Interconnect device.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Understanding Interconnect Devices

interconnect-device (Flow Groups)

IN THIS SECTION

- Syntax | 8804
- Hierarchy Level | 8804
- Description | 8804
- Options | 8805
- Required Privilege Level | 8805
- Release Information | 8805

Syntax

```
interconnect-device interconnect-device-name;
```

Hierarchy Level

```
[edit fabric flow-groups flow-group-name]
```

Description

In a QFabric system, a hash function is used to select an interconnect device to forward traffic between two Node devices. Since this hash function is performed on all available Interconnect devices, it is possible for redundant multicast streams to flow through one Interconnect device, making that

Interconnect device a potential single point of failure for the redundant flows. You can enforce the use of specific Interconnect devices by using the QFabric system `flow-groups` statement.

```
[edit fabric]
root@qfabric# set flow-groups test1 interconnect-device ?
Possible completions:
        <interconnect-device-name> Serial ID or alias of interconnect-device
        IC-P6601-C  Interconnect-device
        IC-P6602-C  Interconnect device
        IC-P6603-C  Interconnect device
```

Options

interconnect-device-name assigns Interconnect devices to the named flow group

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D15.

RELATED DOCUMENTATION

flow-groups | **7291**

Example: Creating a QFabric Flow Group

interconnect-multihoming-peer-gateways

IN THIS SECTION

- Syntax | 8806
- Hierarchy Level | 8806
- Description | 8806
- Options | 8807
- Required Privilege Level | 8807
- Release Information | 8807

Syntax

```
interconnect-multihoming-peer-gateways (multihoming-peer-gateway-IP-address | [multihoming-peer-gateway-IP-address ... ]);
```

Hierarchy Level

```
[edit logical-systems name protocols evpn],
[edit protocols evpn]
```

Description

Specify Ethernet VPN (EVPN) provider edge (PE) multihoming peer Data Center Interconnect (DCI) gateway devices.

When you configure this option, the DCI gateway devices avoid loops during EVPN configuration changes and can handle control plane routing more simply, which helps optimize forwarding between the peer DCI gateway devices.

We require this statement on all peer multihoming DCI gateway devices that use the `interconnect` statement to configure DCI and seamless stitching with EVPN Type 2 or Type 5 routes.

NOTE: In Junos OS Release 24.2R1, this statement was renamed `multihoming-peer-gateways` to support multiple multihoming feature use cases:

- To identify DCI multihoming peer devices (the original function of this `interconnect-multihoming-peer-gateways` statement).
- To statically identify enhanced OISM multihoming peer leaf devices.

In Junos OS and Junos OS Evolved Release 24.4R1, we've restored the `interconnect-multihoming-peer-gateways` statement for the DCI use case. For the enhanced OISM multihoming peers use case, we provide the "static-multihoming-peer" on page 28589 statement. The `multihoming-peer-gateways` statement isn't available in the Junos CLI anymore.

Options

(*multihoming-peer-gateway-IP-address* | [*multihoming-peer-gateway-IP-address* …])—One or more multihoming peer DCI gateway device loopback IPv4 addresses. Specify a single address or list multiple peer device addresses within square brackets.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

interconnect | 8797

interconnected-vlan-list

IN THIS SECTION

- Syntax | 8808
- Hierarchy Level | 8808
- Description | 8808
- Options | 8809
- Required Privilege Level | 8809
- Release Information | 8809

Syntax

```
interconnected-vlan-list (vlan-list );
```

Hierarchy Level

```
[edit protocols evpn interconnect]
[edit routing-instances routing-instance-name protocols evpn interconnect]
```

Description

Configures seamless stitching between EVPN-VXLAN data centers through an EVPN-MPLS fabric. This interconnect method transfers unicast and BUM traffic using WAN gateways with gateway-redundancy multihoming support. This feature is VLAN-based, and includes VLAN-aware bundle and VLAN bundle support using the VLAN list you configure with this statement.

Options

vlan-list Specifies the VLAN identifiers for interconnection. For example:

```
interconnected-vlan-list [51 52];
encapsulation mpls;
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statements introduced in Junos OS Release 21.4R1.

interconnected-vni-list

IN THIS SECTION

- Syntax | 8809
- Hierarchy Level | 8810
- Description | 8810
- Options | 8810
- Required Privilege Level | 8810
- Release Information | 8810

Syntax

```
interconnected-vni-list (vni-list | all);
```

Hierarchy Level

```
[edit protocols evpn interconnect]
[edit routing-instances routing-instance-name protocols evpn interconnect]
```

Description

Establishes that a Juniper Networks switch translates all VXLAN network identifiers (VNIs) or just a specified list of VNIs. This translation occurs when different VNIs are configured for the same VLAN in EVPN-VXLAN points of delivery (PODs) or data centers, and in the EVPN-VXLAN WAN network that interconnects them.

For information about specifying the VNI to which a Juniper Networks switch translates a VNI, see "translation-vni" on page 30080.

Options

vni-list Translate only the specified VNIs with values that range from 1 through 16777214. For example, `interconnected-vni-list [10-50 60 70]`.

all Translate all VNIs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

interface (services)

IN THIS SECTION

- Syntax | 8811
- Hierarchy Level | 8811
- Description | 8811
- Options | 8811
- Required Privilege Level | 8812
- Release Information | 8812

Syntax

```
interface interface-name.unit-number;
```

Hierarchy Level

```
[edit services service-interface-pools pool pool-name]
```

Description

Add logical service interfaces to the pool of service interfaces.

Options

interface-name.unit-number—Name and logical unit number of the service interface.

- All interfaces in a pool must belong to the same service PIC or DPC.
- All interfaces assigned to the same service must be in the same pool.
- Logical interfaces cannot be in more than one pool.
- All interfaces must have either `family inet` or `family inet6` configured.

- Logical unit 0 cannot be configured in a service interface pool.
- You can configure up to 1000 logical interfaces in a service interface pool.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

interface (VRRP)

IN THIS SECTION

- Syntax | 8812
- Hierarchy Level | 8813
- Description | 8813
- Options | 8813
- Required Privilege Level | 8813
- Release Information | 8813

Syntax

```
interface interface-name {
    bandwidth-threshold bits-per-second priority-cost priority;
    priority-cost priority;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track]
```

Description

Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.

Options

interface-name—Interface to be tracked for this VRRP group.

- **Range:** 1 through 10 interfaces

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

bandwidth-threshold statement added in Junos OS Release 8.1.

RELATED DOCUMENTATION

> *Configuring a Logical Interface to Be Tracked for a VRRP Group*
>
> Junos OS Services Interfaces Library for Routing Devices

interface (Logical Interfaces)

IN THIS SECTION

- Syntax | 8814
- Hierarchy Level | 8814
- Description | 8814
- Options | 8814
- Required Privilege Level | 8815
- Release Information | 8815

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name],
[edit vlans vlan-name]
```

Description

(MX Series routers and EX Series switches only) Specify the logical interfaces to include in the bridge domain, VLAN, VPLS instance, or virtual switch.

Options

interface-name—Name of a logical interface. For more information about how to configure logical interfaces, see the Junos OS Network Interfaces Library for Routing Devices.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for top-level configuration for the `virtual-switch` type of routing instance added in Junos OS Release 9.2.

In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring a Bridge Domain

Configuring a Layer 2 Virtual Switch

Configuring a Layer 2 Virtual Switch on an EX Series Switch

Tunnel Services Overview

Tunnel Interface Configuration on MX Series Routers Overview

interface (Precision Time Protocol)

IN THIS SECTION

- Syntax | 8816
- Hierarchy Level | 8816
- Description | 8816
- Options | 8816
- Required Privilege Level | 8816

Syntax

```
interface interface-name {
    unicast-mode {
        clock-source ip-address {
            local-ip-address local-ip-address;
            }
        }
        transport ipv4;
    }
}
```

Hierarchy Level

```
[edit protocols ptp slave]
```

Description

The interface on which to respond to the upstream PTP timeTransmitter.

Options

unicast-mode Configure upstream unicast PTP timeTransmitter clock sources

clock-source Configure the parameters of the PTP.

local-ip-address Configure the IP address of the local interface acting as the timeReceiver.

asymmetry Specify the asymmetry value between the timeTransmitter and the timeReceiver.

transport Configure the encapsulation type for PTP packet transport.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

IEEE 1588v2 Precision Timing Protocol (PTP)

Precision Time Protocol Overview

Configuring the Precision Time Protocol G.8275.2 Enhanced Profile (Telecom Profile)

interface (LLDP)

IN THIS SECTION

- Syntax | 8817
- Hierarchy Level | 8817
- Description | 8818
- Options | 8818
- Required Privilege Level | 8819
- Release Information | 8819

Syntax

```
interface (all | [interface-name-list]) {
    (disable | enable);
    power-negotiation <(disable | enable)>;
    (tlv-filter | tlv-select);
    trap-notification (disable | enable);
}
```

Hierarchy Level

```
[edit protocols lldp],
[edit routing-instances routing-instance-name protocols lldp]
```

Description

Configure Link Layer Discovery Protocol (LLDP) on all interfaces or on a particular interface.

> **NOTE**: On MX Series and T Series routers, you run LLDP on a physical interface, such as `ge-1/0/0`, and not at the logical interface (unit) level.
>
> Starting with Junos OS Release 14.2, on MX Series devices, you can also configure LLDP on management interfaces, such as `fxp` or `me`.
>
> LLDP is not supported on `vme` interfaces.
>
> For information about interface names, see Interface Naming Overview. For information about interface names for TX Matrix routers, see TX Matrix Router Chassis and Interface Names. For information about FPC numbering on TX Matrix routers, see Routing Matrix with a TX Matrix Router FPC Numbering.
>
> For information about extended port names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .

Options

(all | [*interface-name-list*])

Configure LLDP on all interfaces or on one or more interfaces.

(disable | enable)

Disable or enable LLDP on all interfaces or on the specified interfaces.

- **Default:** Disable

power-negotiation <(disable | enable)>

(EX, QFX Series only) Configure LLDP power negotiation, which negotiates with Power over Ethernet (PoE) powered devices to allocate power.

You must also configure the `management class` statement at the [edit poe] hierarchy level to activate LLDP power negotiation.

- **Values:** Configure one of the following:
  - disable—Disable LLDP power negotiation.
  - enable—Enable LLDP power negotiation.

trap-notification (disable | enable)

Disables or enables the LLDP and physical topology SNMP traps for the specific interface or all the interfaces.

- **Values:** Configure one of the following:

- disable—Disable the LLDP and physical topology SNMP trap notifications.
- enable—Enable the LLDP and physical topology SNMP trap notifications.

- **Default:** disable

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`power-negotiation` introduced in Junos OS Release 12.2 for EX and QFX Series switches.

`trap-notification` introduced in Junos OS Release 15.1R7-S3 for EX3300, EX4200, EX4500, EX4550, EX6200, EX8200 switches.

RELATED DOCUMENTATION

Configuring LLDP (CLI Procedure)

Configuring LLDP

Configuring PoE Interfaces on EX Series Switches

interface (802.1X)

IN THIS SECTION

Syntax | 8820

Hierarchy Level | 8821

Description | 8821

- Options | 8821
- Required Privilege Level | 8828
- Release Information | 8828

Syntax

```
interface (all | [ interface-names ]) {
    authentication-order (captive-portal | dot1x | mac-radius);
    disable;
    eapol-block {
        captive-portal;
        mac-radius;
        server-fail <block-interval>;
    }
    guest-bridge-domain guest-bridge-domain;
    guest-gbp-tag guest-gbp-tag
    guest-vlan guest-vlan (vlan-id | vlan-name;
    ignore-port-bounce;
    mac-radius {
        authentication-protocol {
            eap-md5;
            eap-peap {
                resume;
            }
            pap;
        }
        flap-on-disconnect;
        restrict;
    }
    maximum-requests number;
    multi-domain {
        max-data-session max-data-session;
        packet-action (drop-and-log | shutdown);
        recovery-timeout seconds;
    }
    (no-reauthentication | reauthentication seconds );
    no-tagged-mac-authentication;
    quiet-period seconds;
```

```
    redirect-url  redirect-url;
    retain-mac-aged-session;
    retries number;
    server-fail (bridge-domain bridge-domain | deny | permit | use-cache | vlan-name vlan-name)
gbp-tag gbp-tag;
    server-fail-voip (deny | permit | use-cache | vlan-name vlan-name);
    server-reject-bridge-domain | server-reject-vlan identifier {
        block-interval block-interval;
        eapol-block;
    }
    server-timeout seconds;
    supplicant (single | single-secure | multiple);
    supplicant-timeout  seconds;
    transmit-period seconds;
}
```

Hierarchy Level

```
[edit logical-systems name protocols              dot1x              authenticator],
[edit protocols dot1x                 authenticator]
```

Description

Configure IEEE 802.1X authentication for Port-Based Network Access Control for all interfaces or for specific interfaces.

Options

(all | [*interface-names*])

Configure either a list of interface names or all interfaces for 802.1x authentication.

disable

Disable 802.1X authentication on a specified interface or all interfaces.

- **Default:** 802.1X authentication is disabled on all interfaces.

guest-bridge-domain *guest-bridge-domain*

(MX Series only) Specify the bridge domain tag identifier or the name of the guest bridge domain to which an interface is moved when no 802.1X supplicants are connected on the interface. The bridge domain specified must already exist on the device.

guest-gbp-tag

(EX or QFX Series) Specify the GBP tag to apply when an interface is moved to a guest VLAN. If the `guest-gbp-tag` is configured and the client authenticates in guest VLAN, then the configured `guest-gbp-tag` filter is also installed for the client. You can only configure the `guest-gbp-tag` when the `guest-vlan` *vlan-id* option is configured.

- **Range:** 1 through 65535.

guest-vlan (*vlan-id* | *vlan-name*

(EX, QFX, and SRX Series only) Specify the VLAN tag identifier or the name of the guest VLAN to which an interface is moved. The VLAN specified must already exist on the device. Guest VLANs can be configured on devices that are using 802.1X authentication to provide limited access—typically only to the Internet—for corporate guests. A guest VLAN is not used for supplicants that send incorrect credentials. Those supplicants are directed to the server-reject VLAN instead.

ignore-port-bounce

Ignore the port-bounce command contained in a Change of Authorization (CoA) request. CoA requests are RADIUS messages that are used to dynamically modify an authenticated user session already in progress. CoA requests are sent from the authentication, authorization, and accounting (AAA) server to the device, and are typically used to change the VLAN for the host based on device profiling. End devices such as printers do not have a mechanism to detect the VLAN change, so they do not renew the lease for their DHCP address in the new VLAN. The port-bounce command is used to force the end device to initiate DHCP re-negotiation by causing a link flap on the authenticated port.

- **Default:** The port-bounce command is supported by default. If you do not configure the `ignore-port-bounce` statement, the device responds to a port-bounce command by flapping the link to re-initiate DHCP negotiation for the end device.

maximum-requests *number*

Specify the maximum number of times an EAPoL request packet is retransmitted to the supplicant before the authentication session times out.

- **Range:** 1 through 10
- **Default:** 2

no-reauthentication | reauthentication *seconds*

Either disable reauthentication or configure the number of seconds before the 802.1X authentication session times out and the client must reattempt authentication.

> **NOTE**: If the authentication server sends an authentication session timeout to the client, this takes priority over the value configured locally using the `reauthentication` statement. The session timeout value is sent from the server to the client as an attribute of the RADIUS Access-Accept message.

- **Range:** 1 through 65,535 seconds
- **Default:** Reauthentication is enabled, with 3600 seconds until the client can attempt to authenticate again.

no-tagged-mac-authentication

Don't allow a tagged MAC address for RADIUS authentication.

quiet-period ***seconds***

Specify the number of seconds the interface remains in the wait state following a failed authentication attempt by a supplicant before reattempting authentication.

- **Range:** 0 through 65,535 seconds
- **Default:** 60 seconds

redirect-url ***redirect-url***

Specify a URL that redirects unauthenticated hosts to a central Web authentication (CWA) server. The CWA server provides a web portal where the user can enter a username and password. If these credentials are validated by the CWA server, the user is authenticated and is allowed access to the network.

The redirect URL for central Web authentication can be configured centrally on the AAA server or locally on the switch. Use the `redirect-url` statement to configure the redirect URL locally on the interface connecting the host to the switch.

The redirect URL and a dynamic firewall filter must both be present for the central Web authentication process to be triggered. For more information about configuring the redirect URL and the dynamic firewall filter for central Web authentication, see *Configuring Central Web Authentication*.

> **NOTE**: When the dynamic firewall filter is configured using the special Filter-ID attribute JNPR_RSVD_FILTER_CWA, the CWA redirect URL must include the IP address of the AAA server, for example, **https://10.10.10.10**.

- **Syntax:** The redirect URL must use the HTTP or HTTPS protocol and include an IP address or website name. The following are examples of valid redirect URL formats:
  - http://www.example.com
  - https://www.example.com
  - http://10.10.10.10
  - https://10.10.10.10
  - http://www.example.com/login.html
  - https://www.example.com/login.html
  - http://10.10.10.10/login.html
  - https://10.10.10.10/login.html
- **Default:** Disabled. The redirect URL is not enabled for central Web authentication by default.

request-retry-count *number*

Configure the authentication server to retry sending an EAP request to the supplicant. This can help prevent a timeout of the authentication session due to an unresponsive supplicant. The number of retries is based on the configured value.

- **Range:** 1 through 10 retries
- **Default:** 2 retries

retain-mac-aged-session

Keep the IEEE 802.1X clients active even after their MAC addresses age out, and relearn the MAC addresses again.

- **Default:** Disabled.

retries *number*

Specify the number of times the device attempts to authenticate the port after an initial failure. When the limit is exceeded, the port waits to reattempt authentication for the number of seconds specified with the `quiet-period` option configured at the same hierarchy level.

- **Range:** 1 through 10 retries
- **Default:** 3 retries

server-fail (bridge-domain *bridge-domain* | deny | permit | use-cache | vlan-name *vlan-name*) gbp-tag *gbp-tag*

Specify how end devices connected to a device are supported if the RADIUS authentication server becomes unavailable. Server fail fallback is triggered most often during reauthentication when the already configured and in-use RADIUS server becomes inaccessible. However, server fail fallback can also be triggered by a supplicant's initial attempt at authentication through the RADIUS server.

You must specify an action that the device applies to end devices when the authentication servers are unavailable. The device can accept or deny access to supplicants or maintain the access already granted to supplicants before the RADIUS timeout occurred. You can also configure the switch to move the supplicants to a specific VLAN or bridge domain. The VLAN or bridge domain must already be configured on the device.

> **NOTE**: The `server-fail` statement is specifically for data traffic. For VoIP-tagged traffic, use the `server-fail-voip` statement. The same interface can have a `server-fail` VLAN and a `server-fail-voip` VLAN configured.

- **Values:** *bridge-domain*—(MX Series only) Move the supplicant on the interface to the bridge domain specified by this name or numeric identifier. This action is allowed only if it is the first supplicant connecting to an interface. If an authenticated supplicant is already connected, then the supplicant is not moved to the bridge domain and is not authenticated. The bridge domain must already be configured on the device.

  `deny`—Force the supplicant authentication to fail. No traffic will flow through the interface.

  `permit`—Force the supplicant authentication to succeed. Traffic will flow through the interface as if it were successfully authenticated by the RADIUS server.

  `use-cache`—Force the supplicant authentication to succeed only if it was previously authenticated successfully. This action ensures that already authenticated supplicants are not affected.

  vlan-name—(EX, QFX, or SRX Series only) Move the supplicant on the interface to the VLAN specified by this name or numeric identifier. This action is allowed only if it is the first supplicant connecting to the interface. If an authenticated supplicant is already connected, then the supplicant is not moved to the VLAN and is not authenticated. The VLAN must already be configured on the device.

- **Default:** If the RADIUS authentication server becomes unavailable, the end device is not authenticated and is denied access to the network.

- gbp-tag *gbp-tag*—(EX or QFX Series) Specify the GBP tag to apply on the interface when the server is inaccessible. If you configure the `gbp-tag` *gbp-tag* and the client authenticates in `server-fail` *vlan-name* or `server-fail permit`, then the configured `gbp-tag` *gbp-tag* filter is also installed for the client.

  You can only configure this option when the `server-fail` *vlan-name* or `server-fail permit` option is configured.

  Range: 1 through 65535.

server-fail-voip (deny | permit | use-cache | vlan-name *vlan-name*)

(EX, QFX Series only) Specify how VoIP clients sending voice traffic are supported if the RADIUS authentication server becomes unavailable. Server fail fallback is triggered most often during reauthentication when the already configured and in-use RADIUS server becomes inaccessible. However, server fail fallback can also be triggered by a VoIP client's initial attempt at authentication through the RADIUS server.

You must specify an action that the switch applies to VoIP clients when the authentication servers are unavailable. The switch can accept or deny access to VoIP clients or maintain the access already granted to clients before the RADIUS timeout occurred. You can also configure the switch to move the VoIP clients to a specific VLAN. The VLAN must already be configured on the switch.

The `server-fail-voip` statement is specific to the VoIP-tagged traffic sent by clients. VoIP clients still require that the `server-fail` statement be configured for the un-tagged traffic that they generate. Therefore, when you configure the `server-fail-voip` statement you must also configure the `server-fail` statement.

> **NOTE**: An option other than `server-fail deny` must be configured for `server-fail-voip` to successfully commit.

- **Values:** `deny`—Force the VoIP client authentication to fail. No traffic will flow through the interface.

  `permit`—Force the VoIP client authentication to succeed. Traffic will flow through the interface as if it were successfully authenticated by the RADIUS server.

use-cache—Force the VoIP client authentication to succeed only if it was previously authenticated successfully. This action ensures that already authenticated clients are not affected.

vlan-name—Move the VoIP client on the interface to the VLAN specified by this name or numeric identifier. This action is allowed only if it is the first VoIP client connecting to the interface. If an authenticated VoIP client is already connected, then the VoIP client is not moved to the VLAN and is not authenticated. The VLAN must already be configured on the switch.

- **Default:** If a RADIUS authentication server becomes unavailable, a VoIP client that begins authentication by sending voice traffic is not authenticated, and the voice traffic is dropped.

server-timeout *seconds*

Specify the amount of time a port will wait for a reply when relaying a response from the supplicant to the authentication server before timing out and invoking the server-fail action.

- **Range:** 1 through 60 seconds
- **Default:** 30 seconds

supplicant (single | single-secure | multiple)

Specify the MAC-based method used to authenticate clients.

- **Values:** Specify one of the following:
  - single—Authenticates only the first client that connects to an authenticator port. All other clients connecting to the authenticator port after the first are permitted free access to the port without further authentication. If the first authenticated client logs out, all other supplicants are locked out until a client authenticates again.
  - single-secure—Authenticates only one client to connect to an authenticator port. The host must be directly connected to the switch.
  - multiple—Authenticates multiple clients individually on one authenticator port. You can configure the number of clients per port. If you also configure a maximum number of devices that can be connected to a port through port security settings, the lower of the configured values is used to determine the maximum number of clients allowed per port.
- **Default:** single

supplicant-timeout *seconds*

Specify the number of seconds the port waits for a response when relaying a request from the authentication server to the supplicant before re-sending the request.

- **Range:** 1 through 60 seconds
- **Default:** 30 seconds

transmit-period ***seconds***

Specify the number of seconds the port waits before retransmitting the initial EAPoL PDUs to the supplicant.

- **Range:** 1 through 65,535 seconds
- **Default:** 30 seconds

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`gbp-tag` *`gbp-tag`* and `guest-gbp-tag` *`guest-gbp-tag`* introduced in Junos OS Release 23.4R1 for EX and QFX Series switches.

`server-reject-vlan` introduced in Junos OS Release 9.3 for EX Series switches.

`eapol-block` introduced in Junos OS Release 11.2.

`authentication-order` and `redirect-url` introduced in Junos OS Release 15.1R3.

`server-fail-voip` introduced in Junos OS Releases 14.1X53-D40 and 15.1R4 for EX and QFX Series switches.

`ignore-port-bounce` introduced in Junos OS Release 17.3R1.

`multi-domain` introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

show dot1x | **19548**

Understanding Authentication on Switches

Example: Setting Up 802.1X for Single-Supplicant or Multiple-Supplicant Configurations on an EX Series Switch

Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an EX Series Switch

Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch

Example: Configuring MAC RADIUS Authentication on an EX Series Switch

Configuring 802.1X Interface Settings (CLI Procedure)

Configuring 802.1X Authentication (J-Web Procedure)

Example: Configuring Static MAC Bypass of 802.1X and MAC RADIUS Authentication on an EX Series Switch

Understanding Guest VLANs for 802.1X on Switches

Understanding RADIUS-Initiated Changes to an Authorized User Session

show network-access aaa statistics authentication

Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch

Configuring Central Web Authentication

Configuring RADIUS Server Fail Fallback (CLI Procedure)

Understanding Server Fail Fallback and Authentication on Switches

Example: Configuring Fallback Options on EX Series Switches for EAP-TTLS Authentication and Odyssey Access Clients

Example: Configuring 802.1X Authentication Options When the RADIUS Server Is Unavailable to an EX Series Switch

interface (Access Port Security)

IN THIS SECTION

- Syntax | **8830**
- Hierarchy Level | **8830**
- Description | **8830**
- Options | **8830**
- Required Privilege Level | **8831**
- Release Information | **8831**

Syntax

```
interface (all |  interface-name) {
    allowed-mac {
        mac-address-list;
    }
    (dhcp-trusted | no-dhcp-trusted);
    fcoe-trusted;
    mac-limit limit action (drop | log | none | shutdown);
    no-allowed-mac-log;
    persistent-learning;
    static-ip  ip-address {
        vlan  vlan-name;
        mac mac-address;
    }
    static-ipv6 ip-address {
        vlan  vlan-name;
        mac mac-address;
    }
    }
    vlan vlan-name {
        mac-limit limit action (drop | log | none | shutdown);
    }
}
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port]
```

Description

Apply port security features to all interfaces or to the specified interface.

Options

all—Apply port security features to all interfaces.

interface-name—Apply port security features to the specified interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for the `ipv6-source-guard` statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

RELATED DOCUMENTATION

- Example: Configuring Port Security (non-ELS)
- Example: Protecting Against DHCP Snooping Database Attacks
- Example: Protecting against Ethernet Switching Table Overflow Attacks
- Example: Protecting against DHCP Starvation Attacks
- Example: Protecting against Rogue DHCP Server Attacks
- Configuring MAC Limiting (non-ELS)
- Enabling a Trusted DHCP Server (non-ELS)
- Configuring Static DHCP IP Addresses for DHCP snooping (non-ELS)

interface (Accounting or Sampling)

IN THIS SECTION

- Syntax | 8832
- Hierarchy Level | 8832
- Description | 8832
- Options | 8832
- Required Privilege Level | 8832
- Release Information | 8832

Syntax

```
interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
}
```

Hierarchy Level

```
[edit forwarding-options accounting name output],
[edit forwarding-options sampling family (inet | inet6 | mpls) output],
[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output]
```

Description

Specify the output interface for monitored traffic or for sending copies of packets elsewhere to be analyzed.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Discard Accounting

Configuring Traffic Sampling on MX, M and T Series Routers

Understanding IP-Based Filtering and Selective Port Mirroring of MPLS Traffic

Configuring Discard Accounting

Collecting Traffic Sampling Output in a File

interface (advance-policy-based-routing)

IN THIS SECTION

- Syntax | 8833
- Syntax (Prior Junos OS Release 21.2R1) | 8833
- Hierarchy Level | 8834
- Description | 8834
- Options | 8834
- Required Privilege Level | 8835
- Release Information | 8835

Syntax

```
interface interface-name {
        unit number {
        link-tag link-tag;
        multipath-bandwidth multipath-bandwidth;
        priority priority;
        weight weight;
        }
    }
```

Syntax (Prior Junos OS Release 21.2R1)

```
underlay-interfaces interface-name {
    unit unit-number {
```

```
    link-type (IP | MPLS)
    priority priority-number;
}
```

Hierarchy Level

```
[edit security advance-policy-based-routing]
```

Description

Configure link preference settings for an interface in a APBR profile.

When a list of best paths that meet the SLA requirements are available for the application path, the path selection mechanism selects a path that matches the configured link preference (link type, bandwidth, and priority). The application traffic distribution depends on the link weight configuration. Paths are the WAN links that are used for forwarding the application traffic.

You can define the link type and priority for the underlay links in the APBR profile. The link preference configuration is applied for the application traffic matching the APBR rule.

NOTE: Prior to Junos OS Release 21.2R1, if any of the parameters are not configured (*link-type* or *priority*) then the path selection mechanism follows the existing behavior. That is, applications traffic is assigned to a particular overlay link based on the SLA metrics of that overlay link only.

Options

link-tag

Configure custom tag for the link selection. You can use the tags such as IP, MPLS, VOIP, DSL and so on as per the requirement. Defining a link-tag along with link-priority allows you to prioritize the link selection.

- Default: IP

multipath-bandwidth

Assign the bandwidth for the link in MB. Use this option if you are configuring AMR.

- Range: 1 through 1000000

  Default: 1 GB

priority Assign the priority for the link selection. If there are multiple paths available, AppQOE selects the path with the highest priority. Lower the value, higher the priority for the link.

- **Range:** 1 through 255

weight Configure weight for the application traffic distribution. The weight assigned for a link is based on the factors such as customer preference, CoS, or bandwidth requirement.

- **Range:** 1 through 255

  Default: 1

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1

RELATED DOCUMENTATION

Firewall User Authentication Overview

interface (Analyzer)

IN THIS SECTION

- Syntax | 8836
- Hierarchy Level | 8836
- Description | 8836
- Options | 8836
- Required Privilege Level | 8836

Release Information | 8836

Syntax

```
interface (all | interface-name);
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input egress],
[edit forwarding-options analyzer  analyzer-name input ingress],
[edit forwarding-options analyzer  analyzer-name output]
```

Description

Configure the interfaces for which traffic is mirrored.

Options

all Apply mirroring to all interfaces on the network device. Mirroring a high volume of traffic can be performance intensive for the device. Therefore, you should generally select specific input interfaces in preference to using the `all` keyword, or use the `all` keyword in combination with setting a ratio for statistical sampling. The `all` keyword is not available for the `[edit forwarding-options analyzer analyzer-name output]` hierarchy level.

interface-name Apply mirroring to the specified interface only.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Example: Configuring Port Mirroring Analyzers for Local Monitoring of Employee Resource Use

Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use

Understanding Port Mirroring Analyzers

interface (Backup Selection OSPF)

IN THIS SECTION

- Syntax | 8837
- Hierarchy Level | 8838
- Description | 8838
- Options | 8838
- Required Privilege Level | 8840
- Release Information | 8840

Syntax

```
interface (interface-name | all) {
    admin-group {
        exclude  [ group-name ];
        include-all  [ group-name ];
        include-any  [ group-name ];
        preference  [ group-name ];

    }
    bandwidth-greater-equal-primary;
    dest-metric  (highest | lowest);
    downstream-paths-only ;
    evaluation-order [ admin-group srlg bandwidth protection-type node metric ];
    metric-order  [ root dest ];
    node {
        exclude [ node-address ];
        preference [ node-address ];
```

```
    }
    protection-type  (link | node| node-link);
    root-metric  (highest | lowest);
    srlg (loose |strict);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options backup-selection destination prefix],
[edit logical-systems logical-system-name routing-instances instance-name routing-options backup-
selection destination prefix],
[edit routing-instances instance-name routing-options backup-selection prefix],
[edit routing-options backup-selection destination prefix]
```

Description

Define the backup selection policy for a specific primary next hop.

Options

interface-name	Name of the primary next-hop interface.
all	All the interfaces.
bandwidth-greater-equal-primary	Allow the selection of the backup next hop only if the bandwidth is greater than or equal to the bandwidth of the primary next hop.
dest-metric (highest lowest)	Specifiy the metric from the one-hop neighbor or from the remote router such as an RSVP backup label-switched-path (LSP) tail-end router to the final destination. **highest** Select the backup path that has the highest destination metric. **lowest** Select the backup path that has the lowest destination metric.
downstream-paths-only	Select the backup path that is a downstream path to the destination.
evaluation-order [admin-group srlg bandwidth protection-type node metric]	Control the order and the criteria of evaluating the backup path. The default order of evaluation is admin-group, srlg, bandwidth, protection-type, node and metric.

> **NOTE**: For the explicitly configured evaluation order, only the listed attributes influence the selection of the backup path.

metric-order [root dest]

Specify the order of preference of the root and the destination metric during the backup path selection. The preference order can be:

- [`root dest`] — Backup path selection or preference is first based on the root-metric criteria. If the criteria of all the root-metric is the same, then the selection or preference is based on the dest-metric.
- [`dest root`] — Backup path selection or preference is first based on the dest-metric criteria. If the criteria of all the dest-metric is the same, then the selection is based on the root-metric.

> **NOTE**: Backup path selection or preference is first based on the dest-metric criteria. If the criteria of all the dest-metric is the same, then the selection is based on the root-metric. By default, backup paths with lower destination metric criteria are selected or preferred. If the criteria is the same, then the lowest root metric criteria is preferred or selected.

root The metric to a one-hop neighbor or a remote router.

dest The metric from a one-hop neighbor or remote router to the final destination.

protection-type (link | node | node-link)

Specify the required protection type of the backup path.

> **NOTE**: If no protection-type is configured, then by default the first best path that matches all the other criteria is executed.

link Select the backup path that provides link protection.

node Select the backup path that provides node protection.

node-link Allow either node or link protection LFA where node-protection LFA is preferred over link-protection LFA.

root-metric (highest lowest)

Specify the metric to the one-hop neighbor or to the remote router such as an RSVP backup label-switched-path (LSP) tail-end router.

highest Select the highest root metric.

lowest Select the lowest root metric.

srlg (loose | strict) Define the backup selection to either allow or reject the common shared risk link groups (SRLGs) between the primary link and any link in the backup path.

loose Allow the backup path that has common srlgs between the primary link and any link in the backup path. A backup path with a fewer number of srlg collisions is preferred.

strict Reject the backup path that has common srlgs between the primary next-hop link and each link in the backup path.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Example: Configuring Backup Selection Policy for the OSPF or OSPF3 Protocol

Configuring Backup Selection Policy for the OSPF Protocol

Understanding Backup Selection Policy for OSPF Protocol

interface (BOOTP)

IN THIS SECTION

- Syntax | 8841
- Hierarchy Level | 8841
- Description | 8842
- Options | 8842
- Required Privilege Level | 8842
- Release Information | 8842

Syntax

```
interface (interface-name | interface-group) {
    apply-secondary-as-giaddr;
    broadcast;
    client-response-ttl number;
    description text-description;
    maximum-hop-count number;
    minimum-wait-time seconds;
    no-listen;
    server address {
        logical-system logical-system-name <routing-instance [ <default> routing-instance-
names ]>;
        routing-instance [ <default> routing-instance-names ];
    }
}
```

Hierarchy Level

```
[edit forwarding-options helpers bootp]
```

Description

Specify the interface for a DHCP and BOOTP relay agent.

Options

interface-group	Set a logical interface or group of logical interfaces with a specific DHCP relay configuration.
apply-secondary-as-giaddr	Enable DHCP relay to use secondary gateway IP on this interface.
broadcast	If the layer 2 interface is unknown, then broadcast.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

interface (BPDU Blocking)

IN THIS SECTION

- Syntax | 8843
- Hierarchy Level | 8843
- Description | 8843

- Options | 8844
- Required Privilege Level | 8844
- Release Information | 8844

Syntax

```
interface (all | [interface-name]);
```

Hierarchy Level

```
[edit protocols layer2-control bpdu-block]
```

```
[edit ethernet-switching-options bpdu-block]
```

Description

Applies Bridge Protocol Data Unit (BPDU) protection on all interfaces or on one or more specified interfaces.

Spanning Tree Protocol (STP), Rapid Spanning Tree protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP) provide Layer 2 loop prevention for EX Series switches. The spanning-tree protocols use BPDU frames to communicate. Through their exchange, spanning-tree topologies determine which interfaces block traffic and which interfaces become root ports and forward traffic. User bridge applications running on a PC can also generate BPDUs. If these BPDUs are picked up by STP applications running on the switch, they can trigger STP miscalculations that can lead to network outages.

To block outside BPDUs from reaching a switch interface connected to devices that are not part of a spanning-tree topology, configure BPDU protection on interfaces in the topology.

CAUTION: When configuring BPDU protection on an interface without spanning trees connected to a switch with spanning trees, be careful that you do not configure BPDU protection on all interfaces. Doing so could prevent BPDUs being received on switch

interfaces (such as a trunk interface) that you intended to have receive BPDUs from a switch with spanning trees.

NOTE: Interfaces that are configured as edge interfaces can transition to a forwarding state immediately because they cannot create network loops. As edge ports are connected to end devices, it is imperative that you configure BPDU protection on edge ports to protect the switch from outside BPDUs. When BPDU protection is enabled on an edge interface, the interface shuts down on encountering an outside BPDU thereby preventing any traffic from passing through the interface.

Options

all—All interfaces.

[*interface-name*]—One or more Ethernet interface names.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Understanding BPDU Protection for STP, RSTP, and MSTP

Understanding BPDU Protection for Spanning-Tree Instance Interfaces

Configuring BPDU Protection for Individual Spanning-Tree Instance Interfaces

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches

show spanning-tree bridge | 26548

show spanning-tree interface | 26555

interface (Bridge Domains)

IN THIS SECTION

- Syntax | 8845
- Hierarchy Level | 8845
- Description | 8846
- Options | 8846
- Required Privilege Level | 8846
- Release Information | 8846

Syntax

```
interface interface-name {
    group-limit limit;
    host-only-interface;
    static {
        group ip-address {
            source ip-address;
        }
    }
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
```

```
snooping],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols vlan
vlan-id igmp-snooping]
```

Description

Enable IGMP snooping on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping

interface (Captive Portal)

IN THIS SECTION

- Syntax | 8847
- Hierarchy Level | 8847
- Description | 8847

- Options | 8847
- Required Privilege Level | 8849
- Release Information | 8849

Syntax

```
interface (all | [interface-names]) {
    quiet-period seconds;
    retries number-of-retries;
    server-timeout seconds;
    session-expiry seconds;
    supplicant ( multiple | single | single-secure);
    user-keepalive minutes;
}
```

Hierarchy Level

```
[edit services captive-portal]
```

Description

Configure captive portal authentication for all interfaces or for specific interfaces.

Options

all All interfaces to be configured for captive portal authentication.

[*interface-names*] List of names of interfaces to be configured for captive portal authentication.

quiet-period *seconds* Configure time, in seconds, after a user exceeds the maximum number of retries before they can attempt to authenticate.

- **Range:** 1–65535 seconds
- **Default:** 60 seconds

retries *number-of-tries*

Configure the number of times the user can attempt to submit authentication information.

- **Range:** 1–65535 tries
- **Default:** 3 tries

server-timeout *seconds*

Configure the time in seconds an interface will wait for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action.

- **Range:** 1–65535 seconds
- **Default:** 20 seconds

session-expiry *seconds*

Configure the number of seconds before the captive portal authentication session times out and the client must reattempt authentication.

> **NOTE**: If the authentication server sends an authentication session timeout to the client, this takes priority over the value configured locally using the `session-expiry` statement. The session timeout value is sent from the server to the client as an attribute of the RADIUS Access-Accept message.

- **Range:** 1 through 65535 seconds
- **Default:** 3600 seconds

supplicant (multiple | single | single-secure)

Configure the MAC-based method used to authenticate clients for captive portal authentication.

- **Values:** Configure one of the following:
  - single—Authenticates only the first client that connects to an authenticator port. All other clients connecting to the authenticator port after the first are permitted free access to the port without further authentication. If the first authenticated client logs out, all other supplicants are locked out until a client authenticates again.
  - single-secure—Authenticates only one client to connect to an authenticator port. The host must be directly connected to the switch.

- multiple—Authenticates multiple clients individually on one authenticator port. You can configure the number of clients per port. If you also configure a maximum number of devices that can be connected to a port through port security settings, the lower of the configured values is used to determine the maximum number of clients allowed per port.

- **Default:** single

user-keepalive ***minutes***

Extend a captive portal authentication session after the MAC table aging timer expires, by the configured number of minutes. The keep-alive timer is started when the MAC address of the authenticated host ages out of the Ethernet switching table. If traffic is received within the keep-alive timeout period, the timer is deleted. If there is no traffic within the keep-alive timeout period, the session is deleted, and the host must re-authenticate.

- **Default:** Disabled. The captive portal authentication session ends when the associated MAC address ages out of the Ethernet switching table.

- **Range:** 7 through 65535 minutes

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`user-keepalive` introduced in Junos OS Release 16.1 for EX Series switches.

RELATED DOCUMENTATION

Example: Setting Up Captive Portal Authentication on an EX Series Switch

Configuring Captive Portal Authentication (CLI Procedure)

Access Control and Authentication on Switching Devices

interface (Chassis Cluster)

IN THIS SECTION

- Syntax | 8850
- Hierarchy Level | 8850
- Description | 8850
- Options | 8850
- Required Privilege Level | 8851
- Release Information | 8851

Syntax

```
logical-interface-name secondary-ip-address;
```

Hierarchy Level

```
[edit chassis cluster redundancy-groupgroup-number        ip-monitoring family family-name IP-
address]
```

Description

Specify the redundant Ethernet interface, including its logical-unit-number, through which the monitored IP address must be reachable. The specified redundant Ethernet interface can be in any redundancy group. Likewise specify a secondary IP address to be used as a ping source for monitoring the IP address through the secondary node's redundant Ethernet interface link.

Options

- *logical-interface-name*—Redundant Ethernet interface through which the monitored IP address must be reachable. You must specify the redundant Ethernet interface logical-unit-number. Note that you must also configure a secondary ping source IP address (see below).
- **Range:** `reth0.`*logical-unit-number* through `reth128.`*logical-unit-number* (device dependent)

> **NOTE**: If the redundant Ethernet interface belongs to a VPN routing and forwarding (VRF) routing instance type, then the IP monitoring feature will not work.

- `secondary-ip-address` *IP-address*—Specify the IP address that are used as the source IP address of ping packets for IP monitoring from the secondary child link of the redundant Ethernet interface. An IP address for sourcing the ping packets on the primary link of the redundant Ethernet interface must be configured before you can configure `secondary-ip-address`. For legacy support reasons, monitoring on an IP address without identifying a redundant Ethernet interface and without configuring a secondary ping source IP address is permitted but not recommended.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

interface (DCBX)

IN THIS SECTION

- Syntax | 8852
- Hierarchy Level | 8852
- Description | 8852
- Options | 8852
- Required Privilege Level | 8852
- Release Information | 8853

Syntax

```
interface (interface-name | all) {
    disable;
    application-map application-map-name;
    applications {
        no-auto-negotiation;
    }
    enhanced-transmission-selection {
        no-auto-negotiation;
        no-recommendation-tlv;
        recommendation-tlv {
            no-auto-negotiation;
        }
    }
    dcbx-version (auto-negotiate | ieee-dcbx | dcbx-version-1.01);
    priority-flow-control {
        no-auto-negotiation;
    }
}
```

Hierarchy Level

```
[edit protocols dcbx]
```

Description

Configure DCBX properties on an interface.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

`Mode` and `recommendation-tlv` statements introduced in Junos OS Release 12.2 for the QFX Series.

RELATED DOCUMENTATION

show dcbx neighbors | **19191**

Configuring DCBX Autonegotiation

Example: Configuring DCBX to Support an iSCSI Application

Understanding DCB Features and Requirements

Understanding DCB Features and Requirements on EX Series Switches

Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

interface-delete (Subscriber Management or DHCP Client Management)

IN THIS SECTION

- Syntax | **8853**
- Hierarchy Level | **8854**
- Description | **8854**
- Required Privilege Level | **8854**
- Release Information | **8854**

Syntax

```
interface-delete;
```

Hierarchy Level

```
[edit system services subscriber-management maintain-subscriber]
```

Description

On router—Configure the router to maintain, rather than log out, subscribers when the subscriber interface is deleted. By default, the router logs out subscribers when the subscriber interface is deleted.

On switch—Configure the switch to maintain rather than log out DHCP clients when the client interface is deleted. By default, the switch logs out DHCP clients when the client interface is deleted.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events

interface-description (DHCP Local Server)

IN THIS SECTION

- Syntax | 8855
- Hierarchy Level | 8855
- Description | 8855
- Options | 8855
- Required Privilege Level | 8856

Release Information | 8856

Syntax

```
interface-description (device-interface | logical-interface);
```

Hierarchy Level

```
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Specify that the interface description for the device (physical) interface or the logical interface is concatenated with the username during the subscriber authentication or DHCP client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

The interface description is configured using the `description` statement at the `[edit interfaces interface-name]` hierarchy level. The delimiter is configured with the `delimiter` statement at the `[edit system services dhcp-local-server ... username-include]` hierarchy level.

NOTE: The username delimiter must not be a character that is part of the interface description. For example, if the description includes a forward slash (/), then the delimiter must not be a forward slash. The default delimiter is a period.

Options

device-interface Use the text description configured for the device (physical) interface.

logical-interface Use the text description configured for the logical interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients

interface-description (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 8856
- Hierarchy Level | 8857
- Description | 8857
- Options | 8857
- Required Privilege Level | 8857
- Release Information | 8858

Syntax

```
interface-description (device-interface | logical-interface);
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-
include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify that the interface description for the device (physical) interface or the logical interface is concatenated with the username during the subscriber authentication or DHCP client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

The interface description is configured using the `description` statement at the `[edit interfaces interface-name]` hierarchy level. The delimiter is configured with the `delimiter` statement at the `[edit forwarding-options dhcp-relay ... username-include]` hierarchy level.

> **NOTE**: The username delimiter must not be a character that is part of the interface description. For example, if the description includes a forward slash (/), then the delimiter must not be a forward slash. The default delimiter is a period.

Options

device-interface Description of the physical interface.

logical-interface Description of the logical interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients

interface-description-format

IN THIS SECTION

- Syntax | 8858
- Hierarchy Level | 8858
- Description | 8859
- Options | 8859
- Required Privilege Level | 8859
- Release Information | 8860

Syntax

```
interface-description-format {
    exclude-adapter;
    exclude-channel;
    exclude-sub-interface;
}
```

Hierarchy Level

```
[edit access profile profile-name radius options]
```

Description

Specify the information that is excluded from the interface description that the device passes to RADIUS for inclusion in the RADIUS attributes such as NAS-Port-ID (87) or Calling-Station-ID (31).

The default format for nonchannelized interfaces is as follows:

interface-type-slot/adapter/port.subinterface[:*svlan-vlan*]

For example, consider physical interface ge-1/2/0, with a subinterface of 100 and SVLAN identifier of 100. The interface description used in the NAS-Port-ID is ge-1/2/0.100:100. If you exclude the subinterface, the description becomes ge-1/2/0:100.

The default format for channelized interfaces is as follows:

interface-type-slot/adapter/channel.subinterface[:*svlan-vlan*]

The channel information (logical port number) is determined by this formula:

Logical port number = 100 + (*actual-port-number* x 20) + *channel-number*.

For example, consider a channelized interface 3 on port 2 where the:

- Physical interface is xe-0/1/2:3.
- Subinterface is 4.
- SVLAN is 5.
- VLAN is 6.

Using the formula, the logical port number = 100 + (2 x 20) + 3 = 143. Consequently, the default interface description is xe-0/1/143.4-5.6. If you exclude the channel information, the description becomes xe-0/1/2.4-5.6.

Options

exclude-adapter —(Optional) Exclude the adapter from the interface description.

exclude-channel (Optional) Exclude the channel information from the interface description.

exclude-sub-interface —(Optional) Exclude the subinterface from the interface description.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`exclude-adapter` and `exclude-sub-interface` options added in Junos OS Release 10.4.

`exclude-channel` option added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

Standard and Vendor-Specific RADIUS Attributes

interface (DHCP Local Server)

IN THIS SECTION

- Syntax | **8860**
- Hierarchy Level | **8862**
- Description | **8862**
- Options | **8862**
- Required Privilege Level | **8863**
- Release Information | **8863**

Syntax

```
interface interface-name {
    access-profile profile-name;
    exclude;
    overrides {
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
```

```
        client-discover-match <option60-and-option82 | incoming-interface>;
        client-negotiation-match  incoming-interface;
        delay-advertise {
            based-on (option-15 | option-16 | option-18 | option-37) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        delay-offer {
            based-on (option-60 | option-77 | option-82) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        rapid-commit;
    }
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    trace;
```

```
    upto upto-interface-name;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP local server is enabled. You can repeat the `interface` *`interface-name`* statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group. Also, you cannot use an interface that is being used by the DHCP relay agent.

NOTE: DHCP values are supported in integrated routing and bridging (IRB) configurations. When you configure an IRB interface in a network that is using DHCP, the DHCP information (for example, authentication, address assignment, and so on) is propagated in the associated bridge domain. This enables the DHCP server to configure client IP addresses residing within the bridge domain. IRB currently supports only static DHCP configurations.

Options

`exclude`—Exclude an interface or a range of interfaces from the group. This option and the `overrides` option are mutually exclusive.

`interface-name`—Name of the interface. You can repeat this option multiple times.

`upto-interface-name`—Upper end of the range of interfaces; the lower end of the range is the interface-name entry. The interface device name of the *`upto-interface-name`* must be the same as the device name of the *`interface-name`*.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Options `upto` and `exclude` introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Understanding Differences Between Legacy DHCP and Extended DHCP

Grouping Interfaces with Common DHCP Configurations

Specifying Authentication Support

interface (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 8863
- Hierarchy Level | 8864
- Description | 8864
- Options | 8865
- Required Privilege Level | 8865
- Release Information | 8865

Syntax

```
interface dhcp-interface-name {
    access-profile profile-name;
    exclude;
```

```
    overrides {
        allow-no-end-option
        allow-snooped-clients;
        always-write-giaddr;
        always-write-option-82;
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-discover-match <option60-and-option82 | incoming-interface>;
        client-negotiation-match incoming-interface;
        disable-relay;
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        layer2-unicast-replies;
        no-allow-snooped-clients;
        proxy-mode;
        relay-source
        replace-ip-source-with;
        send-release-on-delete;
        trust-option-82;
    }
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    trace;
    upto upto-interface-name;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP or DHCPv6 relay agent is enabled. You can repeat the `interface` *`interface-name`* statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group.

Also, you cannot use an interface that is being used by the DHCP local server. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

EX Series switches do not support DHCPv6.

> **NOTE**: DHCP values are supported in integrated routing and bridging (IRB) configurations. When you configure an IRB interface in a network that is using DHCP, the DHCP information (for example, authentication, address assignment, and so on) is propagated in the associated bridge domain. This enables the DHCP server to configure client IP addresses residing within the bridge domain. IRB currently only supports static DHCP configurations. .

Options

`exclude`—Exclude an interface or a range of interfaces from the group. This option and the `overrides` option are mutually exclusive.

`interface-name`—Name of the interface. You can repeat this option multiple times.

`overrides`—Override the specified default configuration settings for the interface. The `overrides` statement is described separately.

`upto-interface-name`—Upper end of the range of interfaces; the lower end of the range is the interface-name entry. The interface device name of the *`upto-interface-name`* must be the same as the device name of the *`interface-name`*.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Options `upto` and `exclude` introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Grouping Interfaces with Common DHCP Configurations

Specifying Authentication Support

interface (DHCP Security for MX Series)

IN THIS SECTION

- Syntax | 8866
- Hierarchy Level | 8866
- Description | 8866
- Required Privilege Level | 8867
- Release Information | 8867

Syntax

```
interface interface-name {
    static-ip ip-address {
        mac mac-address;
    }
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security group group-name]
```

Description

Configure an interface for a static IPaddress to MAC address binding (IP-MAC binding) or configure an interface to belong to a group within the bridge domain that has DHCP security attributes that are different from the attributes of other interfaces in the bridge domain.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Relese 14.1.

RELATED DOCUMENTATION

Configuring Static DHCP IP Addresses for DHCP snooping (MX routers)

interface (DHCP Spoofing Prevention)

IN THIS SECTION

- Syntax | 8867
- Hierarchy Level | 8868
- Description | 8868
- Required Privilege Level | 8868
- Release Information | 8868

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name forwarding-options
dhcp-relay group group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name]
```

Description

Configure Dynamic Host Configuration Protocol (DHCP) snooping on the router. When acting as a snooping agent, the MX Series router typically is located between the client and the DHCP relay agent. It creates filters by “snooping” DHCP messages and binding DHCP-issued IP addresses with the MAC address of the client. These filters help prevent DHCP spoofing.

DHCP snooping is configured by including the appropriate interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4 (MX Series routers only).

RELATED DOCUMENTATION

Preventing DHCP Spoofing on MX Series 5G Universal Routing Platforms

interface (DNS, Port, and TFTP Packet Forwarding or Relay Agent)

IN THIS SECTION

- Syntax | 8869
- Hierarchy Level | 8869

- Description | 8869
- Options | 8869
- Required Privilege Level | 8870
- Release Information | 8870

Syntax

```
interface interface-name {
    broadcast;
    description text-description;
    no-listen;
    server address <logical-system logical-system-name> <routing-instance routing-instance-name>;
}
```

Hierarchy Level

```
[edit forwarding-options helpers domain],
[edit forwarding-options helpers port port-number],
[edit forwarding-options helpers tftp]
```

Description

Specify the interface for monitoring and forwarding DNS or TFTP requests, or for forwarding LAN broadcast traffic on a custom UDP port to a particular server as unicast traffic.

When configuring port helpers, in releases prior to Junos OS Release 17.2, only one server can be specified for a given port. For Junos OS Release 17.2 and later, multiple servers can be specified for a given port at the global or interface-specific level. In this case, the same packet, with the originator IP address and port requests, is forwarded to the different configured servers; the payload of the UDP packet is not modified.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for port helpers in Junos OS Release 17.2R1 for EX4300 switches.

Support for multiple server instances for a given port introduced in Junos OS Release 17.2 for MX Series routers.

Support for multiple server instances for a given port introduced in Junos OS Release 17.3R1 for EX9200 switches.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding

Configuring Port-based LAN Broadcast Packet Forwarding

interface-down

IN THIS SECTION

- Syntax | **8871**
- Hierarchy Level | **8871**
- Description | **8871**
- Required Privilege Level | **8871**
- Release Information | **8871**

Syntax

```
interface-down;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile profile-name default-actions]
```

Description

Bring the interface down when a remote MEP connectivity failure is detected.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring a CFM Action Profile to Specify CFM Actions for CFM Events

interface (Dynamic IGMP)

IN THIS SECTION

- Syntax | 8872
- Hierarchy Level | 8872
- Description | 8873

- Options | 8873
- Required Privilege Level | 8873
- Release Information | 8873

Syntax

```
interface interface-name {
    accounting;
    disable;
    distributed;
    group-limit limit;
    group-policy;
    group-threshold value;
    immediate-leave
    log-interval seconds;
    no-accounting;
    oif-map;
    passive;
    promiscuous-mode;
    ssm-map ssm-map-name;
    ssm-map-policy ssm-map-policy-name
    static {
        group group {
            source source;
        }
    }
    version version;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp]
```

Description

Enable IGMP on an interface and configure interface-specific properties.

Options

interface-name—Variable for the interface. Specify the interface variable ($junos-interface-name) to indicate that the dynamic profile chooses an interface for the accessing DHCP client.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Enabling IGMP

interface (Dynamic Interface Sets)

IN THIS SECTION

- Syntax | 8874
- Hierarchy Level | 8874
- Description | 8874
- Options | 8874

Required Privilege Level | 8875
Release Information | 8875

Syntax

```
interface interface-name {
    unit logical unit number {
        advisory-options {
            downstream-rate rate;
            upstream-rate rate;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-set interface-set-name]
```

Description

Add a subscriber interface to a dynamic interface set.

In a dynamic profile that defines an agent circuit identifier (ACI) interface set, observe the following guidelines when you use the `interface` statement:

- Use the predefined dynamic interface variable `$junos-interface-ifd-name` to represent the interface name. Do not use a specific interface name, such as `demux0`, when defining an ACI interface set.
- Do not include the `unit` *`logical-unit-number`* statement.

Options

interface-name–Either the specific name of the interface to include in the interface set, or the predefined dynamic interface variable `$junos-interface-ifd-name`. The interface variable is dynamically replaced with the interface that the DHCP or PPPoE subscriber accesses when connecting to the router.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Defining ACI Interface Sets

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring an Interface Set of Subscribers in a Dynamic Profile

Agent Circuit Identifier-Based Dynamic VLANs Overview

interface (Dynamic MLD)

IN THIS SECTION

- Syntax | 8876
- Hierarchy Level | 8876
- Description | 8876
- Options | 8877
- Required Privilege Level | 8877
- Release Information | 8877

Syntax

```
interface interface-name {
    (accounting | no-accounting);
    disable;
    distributed;
    group-limit limit;
    group-policy;
    group-threshold value;
    immediate-leave;
    log-interval seconds;
    oif-map;
    passive;
    ssm-map ssm-map-name;
    ssm-map-policy ssm-map-policy-name;
    static {
        group multicast-group-address {
            exclude;
            group-count number;
            group-increment increment;
            source ip-address {
                source-count number;
                source-increment increment;
            }
        }
    }
    version version;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld]
```

Description

Enable MLD on a dynamic interface and configure interface-specific properties.

Options

interface-name—Variable for the interface. Specify the interface variable ($junos-interface-name) to indicate that the dynamic profile chooses an interface for the accessing client.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Enabling MLD

interface (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 8878
- Hierarchy Level | 8878
- Description | 8878
- Options | 8878
- Required Privilege Level | 8879
- Release Information | 8879

Syntax

```
interface interface-name {
     current-hop-limitnumber;
    default-lifetime seconds;
    dns-server-address
    (managed-configuration | no-managed-configuration);
    max-advertisement-interval seconds;
    min-advertisement-interval seconds;
    (other-stateful-configuration | no-other-stateful-configuration);
    prefixprefix {
        (autonomous | no-autonomous);
        (on-link | no-on-link);
        preferred-lifetime seconds;
        valid-lifetime seconds;
    }
    reachable-time milliseconds;
    retransmit-timer milliseconds;
}
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement]
```

Description

Dynamically configure router advertisement properties on an interface. To dynamically configure interface properties, include the *$junos-interface-name* dynamic variable for the interface name.

Options

interface-name—Name of an interface. Specify the *$junos-interface-name* dynamic variable or the full, static interface name, including the physical and logical address components.

NOTE: Even though you can specify the static interface name when defining the interface, we recommend using dynamic variable when configuring this statement.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA

Dynamic Router Advertisement Configuration Overview

interface (Dynamic Routing Instances)

IN THIS SECTION

- Syntax | 8879
- Hierarchy Level | 8880
- Description | 8880
- Options | 8880
- Required Privilege Level | 8880
- Release Information | 8880

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances routing-instance-name]
```

Description

Assign the specified interface to the dynamically created routing instance.

Options

interface-name—The interface name variable (*$junos-interface-name*). The interface name variable is dynamically replaced with the interface the accessing client uses when connecting to the router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

interface (Dynamic Routing Options)

IN THIS SECTION

- Syntax | 8881
- Hierarchy Level | 8881
- Description | 8881
- Options | 8881
- Required Privilege Level | 8881
- Release Information | 8881

Syntax

```
interface interface-names {
    no-qos-adjust;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-options multicast],
[edit dynamic-profiles profile-name routing-instances routing-instance-name routing-options
multicast]
```

Description

Define the maximum bandwidth for a dynamic interface on which you want to apply bandwidth management.

Options

interface-name—Names of the physical or logical interface. For details about specifying interfaces, see Types of Interfaces.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

> *Configuring Dynamic Access Routes for Subscriber Management*

interface (EVPN Routing Instances)

IN THIS SECTION

- Syntax | 8882
- Hierarchy Level | 8882
- Description | 8883
- Options | 8883
- Required Privilege Level | 8883
- Release Information | 8883

Syntax

```
interface interface-name {
    ignore-encapsulation-mismatch;
    interface-mac-limit limit {
        packet-action drop;
    }
    mac-pinning
    no-mac-learning;
    protect-interface
    static-mac mac-address;
    vpws-service-id
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols]
[edit routing-instances routing-instance-name protocols evpn]
```

Description

Specify each interface over which the Ethernet VPN (EVPN) traffic travels between the PE device and customer edge (CE) device. The interfaces are bound to the EVPN routing instance.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Statement (mac-pinning) introduced in Junos OS Release 16.2 on MX Series routers.

`vpws-service-id` statement introduced in Junos OS Release 17.1 on MX Series routers.

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

evpn | **6541**

instance-type | **8773**

vpws-service-id | **30954**

interface (Export Profiles)

IN THIS SECTION

- Syntax | 8884
- Hierarchy Level | 8884
- Description | 8884
- Options | 8885
- Required Privilege Level | 8885
- Release Information | 8885

Syntax

```
interface {
    information;
    statistics {
        queue;
        traffic;
    }
    status {
        link;
        queue;
        traffic;
    }
}
```

Hierarchy Level

```
[edit services analytics export-profiles]
```

Description

Configure an export profile for streaming network analytics data for a specific interface to remote servers. Each profile is a template that defines the type of data being streamed for that interface.

> **NOTE**: The `interface` statement is not available in Junos OS Releases prior to 13.2X51-D15.

Options

information Information about the specified interface, including SNMP index, interface index, slot, port number, media type, capability, and port type.

statistics Type of monitoring statistics to be streamed.

- Values:
  - **queue**
  - **traffic**

status Status information about the interface to be streamed.

- Values:
  - **link**
  - **queue**
  - **traffic**

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

interface (Fibre Channel Fabric)

IN THIS SECTION

- Syntax | 8886
- Hierarchy Level | 8886
- Description | 8886
- Options | 8887
- Required Privilege Level | 8887
- Release Information | 8887

Syntax

```
interface {
    interface-name {
        max-login-sessions max-login-sessions;
    }
    interface-name {
        max-login-sessions max-login-sessions;
    }
    <...> {
        max-login-sessions max-login-sessions;
    }
    vlan.interface-name;
}
```

Hierarchy Level

```
[edit fc-fabrics fc-fabric-name]
```

Description

Associate one or more native Fibre Channel (FC) interfaces with an FC fabric and one VLAN interface for FCoE traffic. An FC interface can be associated with only one FC fabric.

Options

interface-name—Name of the native FC interface. You can assign one or more FC interfaces to an FC fabric.

vlan.*vlan-interface-name*—Name of the VLAN interface for FCoE traffic. You can assign one VLAN interface to an FC fabric.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric

Converting an Ethernet Interface To a Fibre Channel Interface

Configuring a Physical Fibre Channel Interface

Configuring an FCoE VLAN Interface on an FCoE-FC Gateway

Understanding Interfaces on an FCoE-FC Gateway

interface (FIP)

IN THIS SECTION

- Syntax | 8888
- Hierarchy Level | 8888
- Description | 8888
- Options | 8888

Required Privilege Level | 8888
Release Information | 8888

Syntax

```
interface {
    interface-name {
        fka-adv-period milliseconds;
        priority priority;
    }
}
```

Hierarchy Level

```
[edit fc-fabrics fc-fabric-name protocols fip]
```

Description

Configure FIP options on a per-interface basis. (Override global FIP configuration for a specified interface.)

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fip | 19938

Configuring FIP on an FCoE-FC Gateway

Overview of FIP

interface (FIP Snooping)

IN THIS SECTION

- Syntax | 8889
- Hierarchy Level | 8889
- Description | 8889
- Required Privilege Level | 8890
- Release Information | 8890

Syntax

```
interface interface-name {
    (fcoe-trusted| no-fcoe-trusted);
}
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options fip-security]
```

Description

NOTE: This statement supports the Enhanced Layer 2 Software (ELS) CLI. If your switch runs the original (non-ELS) software, see interface (Secure Access Port) for how to

specify an interface to configure as FCoE trusted or FCoE untrusted. For ELS details, see Using the Enhanced Layer 2 Software CLI.

Specify an interface to set as FCoE trusted or as FCoE untrusted. Configure interfaces that connect to other switches as trusted interfaces. Configure interfaces that connect directly to FCoE devices as untrusted interfaces and enabled FIP snooping on the untrusted interfaces to prevent unauthorized access to the storage network.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

Understanding FCoE Transit Switch Functionality

interface-group

IN THIS SECTION

- Syntax | 8891
- Hierarchy Level | 8891
- Description | 8891
- Options | 8891
- Required Privilege Level | 8891
- Release Information | 8891

Syntax

```
interface-group {
    interface-device-name
    unit-list
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-association ma-name mep mep-id remote-mep mep-id ]
```

Description

Mark the interface group down for the action profile configured with the action `interface-group-down`. Provides information for the interface-group on which the configured action will be taken when the configured event occur for a specific remote MEP ID.

Options

interface-device name Name of the interface device. Only Ethernet devices are allowed. The device interface name includes **ge, ae, xe and et.** .

unit-list One or more logical interface unit numbers.

- **Range:** A string in the range <0-16385> or <0-16385>-<0-16385>. For example, `unit-list`[12 23-33 44]

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces

interface-group-down | 8893

interface-group (Decapsulate GRE)

IN THIS SECTION

- Syntax | 8892
- Hierarchy Level | 8892
- Description | 8892
- Required Privilege Level | 8893
- Release Information | 8893

Syntax

```
interface-group (0 -255)
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name then decapsulate gre]
```

Description

Allows you to explicitly specify and add an interface group to packets after they have undergone GRE decapsulation. In releases prior to Junos OS Release 14.1, the interface group upon decapsulation was always 0 and could not be changed.

In Junos OS Release 14.1 and later, you can assign any arbitrary value in the range of 1 to 255 to the packets' `interface-group` upon GRE decapsulation. For example, you could use this command to retain the original interface from which the packet was received (if no value is set, the default interface group is 0). You could also use it to ensure that all decapsulated GRE packets are placed in the same group, for

example to trigger additional filtering in the forwarding table on the basis of the this data from the inner packet.

The value used in `interface-group` is set after the GRE packet is decapsulate by a filter action in a filter attached to a given ingress interface.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1 .

RELATED DOCUMENTATION

Example: Configuring a Stateless Firewall Filter on an Interface Group

Applying Forwarding Table Filters

decapsulate (Firewall Filter) | **5160**

interface-group-down

IN THIS SECTION

- Syntax | **8894**
- Hierarchy Level | **8894**
- Description | **8894**
- Required Privilege Level | **8894**
- Release Information | **8894**

Syntax

```
interface-group-down
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
action]
```

Description

Mark the interface group down.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces

interface-group | 8890

interface-group (Routing Options)

IN THIS SECTION

Syntax | 8895

- Hierarchy Level | 8895
- Description | 8895
- Required Privilege Level | 8895
- Release Information | 8896

Syntax

```
interface-group group-id exclude;
```

Hierarchy Level

```
[edit routing-options flow]
```

Description

Exclude applying flowspec filter to traffic received on specific interfaces. Use **exclude** to specify the interface group where you do not want to receive the traffic.

> **CAUTION**: Do not use this statement with BGP flowspec on the QFX10K or Junos based PTX platforms as this might result in erratic behavior. Junos OS does not support the usage of `interface-group` along with BGP flowspec on the QFX10000 Series. Therefore, we do not recommend the use of this statement on these platforms. For details on supported platforms and Junos releases see Feature Explorer.
>
> Note that when you specify exclude groups on Junos based PTX platforms, it is applied to the entire interface or port. Therefore, you cannot configure different filter groups on different logical interfaces. The exclude group is applicable to all families and are exempt from flowspec filtering irrespective to which family you configure this option. Also ensure that you configure the same group here and at [edit inet family] hierarchy level. It does not matter under which family you add the group, either way flows from both families are allowed.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring MPLS-Signaled LSPs to Use GRE Tunnels

Understanding BGP Flow Routes for Traffic Filtering

interface (Hierarchical CoS Schedulers)

IN THIS SECTION

- Syntax | 8896
- Hierarchy Level | 8896
- Description | 8897
- Required Privilege Level | 8897
- Release Information | 8897

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit interfaces interface-set interface-set-name]
```

Description

Specify an interface that is a member of the interface set. Supported on Ethernet interfaces on an MX Series router, Ethernet interfaces on IQ2E PIC on M Series and T Series routers, and IP demux interfaces on an MX Series router.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Junos OS Class of Service User Guide for Routing Devices

interface (Host Fast Reroute)

IN THIS SECTION

- Syntax | 8897
- Hierarchy Level | 8898
- Description | 8898
- Required Privilege Level | 8898
- Release Information | 8898

Syntax

```
interface interface-name {
    arp-prefix-limit number;
```

```
    link-protection;
    supplementary-blackout-timer minutes;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit routing-instances routing-instance-name routing-options]
```

Description

Configure host fast reroute settings, including link protection for the interface and an ARP prefix limit.

The remaining statements are described separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring Link Protection with Host Fast Reroute

interface-id

IN THIS SECTION

Syntax | 8899

- Hierarchy Level | 8899
- Description | 8899
- Options | 8899
- Required Privilege Level | 8899
- Release Information | 8900

Syntax

```
interface-id interface-id;
```

Hierarchy Level

```
[edit access group-profile profile-name l2tp],
[edit access group-profile profile-name ppp],
[edit access profile profile-name client client-name ike],
[edit access profile profile-name client client-name l2tp],
[edit access profile profile-name client client-name ppp]
```

Description

Configure the interface identifier.

Options

interface-id—Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the `[edit interfaces interface-name unit local-unit-number dial-options]` hierarchy level. For more information about the interface ID, see Services Interface Naming Overview.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

Layer 2 Tunneling Protocol (L2TP)

L2TP LNS Inline Service Interfaces

Configuring an IKE Access Profile

interface (IEEE 802.1ag OAM Connectivity-Fault Management)

IN THIS SECTION

- Syntax | 8900
- Hierarchy Level | 8900
- Description | 8901
- Options | 8901
- Required Privilege Level | 8901
- Release Information | 8901

Syntax

```
interface (interface-name | ((ge- | xe-) (fpc/pic/port | fpc/pic/port.unit-number | fpc/pic/
port.unit-number vlan vlan-id)));
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id]
```

Description

For Ethernet interfaces on M320, MX Series, and T Series routers, configure IEEE 802.1ag Operation, Administration, and Management (OAM) support.

For Gigabit Ethernet interfaces and 10-Gigabit Ethernet interfaces on MX Series routers, configure IEEE 802.1ag Connectivity Fault Management (CFM) support on trunk interface ports.

Starting in Junos OS 17.4R1, you can enable support for IEEE 802.1ag CFM on pseudowire service interfaces by configuring maintenance intermediate points (MIPs) on the pseudowire service interfaces.

> **NOTE**: The CFM MIP session is supported only on the pseudowire services interface and not on the pseudowire services tunnel interface.

Options

`interface-name`—Interface to which the MEP is attached. It could be a physical Ethernet interface, logical Ethernet interface, pseudowire services interfaces, or on a specific VLAN of a trunk port interface (MX Series only).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring a MEP to Generate and Respond to CFM Protocol Messages

interface (IGMP Snooping)

IN THIS SECTION

- Syntax | 8902
- Hierarchy Level | 8902
- Description | 8902
- Options | 8903
- Required Privilege Level | 8903
- Release Information | 8903

Syntax

```
interface interface-name {
    group-limit limit;
    host-only-interface;
    immediate-leave;
    multicast-router-interface;
    static {
        group multicast-group-address {
            source ip-address;
        }
    }
}
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan (all | vlan-name)]
```

Description

For IGMP snooping, configure an interface as either a multicast-router interface or as a static member of a multicast group with optional interface-specific properties.

Options

all	All interfaces in the VLAN.
interface-name	Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping on SRX Series Devices

IGMP Snooping Overview

igmp-snooping | 8387

interface-in (Security Forwarding Options)

IN THIS SECTION

- Syntax | 8904
- Hierarchy Level | 8904
- Description | 8904
- Required Privilege Level | 8904
- Release Information | 8904

Syntax

```
interface-in interface-name;
```

Hierarchy Level

```
[edit security forwarding-options mirror-filter filter-name]
```

Description

Specify the incoming logical interface to be matched for mirroring.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

mirror-filter (Security Forwarding Options) | 10768

show security forwarding-options mirror-filter | 24363

interface (Layer 2 Protocol Tunneling)

IN THIS SECTION

- Syntax | 8905
- Hierarchy Level | 8905

- Description | 8905
- Required Privilege Level | 8905
- Release Information | 8905

Syntax

```
interface interface-name {
    enable-all-ifl;
    protocol protocol-name;
}
```

Hierarchy Level

```
[edit logical-systems name protocols layer2-control mac-rewrite],
[edit protocols layer2-control mac-rewrite]
```

Description

Configure an interface for Layer 2 protocol tunneling.

> **NOTE**: The `enable-all-ifl` option is available on EX9200 switches but not on other EX Series switches.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`enable-all-if` statement added in Junos OS Release 13.3.

Support for PVSTP protocol introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Release 14.1X53-D10 and 17.3R1 for EX4300 switches.

RELATED DOCUMENTATION

Understanding Layer 2 Protocol Tunneling

interface (LLDP-MED)

IN THIS SECTION

- Syntax | 8906
- Hierarchy Level | 8907
- Description | 8907
- Default | 8907
- Options | 8907
- Required Privilege Level | 8907
- Release Information | 8908

Syntax

```
interface name {
    (disable | enable);
    location {
        civic-based {
            ca-type  type {
                ca-value value;
            }
            country-code country-code;
            what what;
        }
```

```
        co-ordinate {
            lattitude latitude;
            longitude longitude;
        }
        elin elin;
    }
    tlv-filter;
    tlv-select;
}
```

Hierarchy Level

```
[edit protocols lldp-med]
```

Description

Configure Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) on all interfaces or on a specific interface.

Default

Not enabled.

Options

all | *interface-name* Configure LLDP-MED on all interfaces or on a specific interface.

disable | enable Disable or enable LLDP-MED on all interfaces or on one or more interfaces.

- **Default:** If you do not configure LLDP-MED, it is disabled on the device and on specific interfaces.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show lldp | 21942

Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch

Configuring LLDP-MED (CLI Procedure)

Understanding LLDP and LLDP-MED on EX Series Switches

interface-mac-ip-limit

IN THIS SECTION

- Syntax | 8908
- Hierarchy Level | 8908
- Description | 8909
- Options | 8909
- Required Privilege Level | 8909
- Release Information | 8909

Syntax

```
interface-mac-ip-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name vlans vlan-name
switch-options],
[edit logical-systems logical-system-name vlans vlan-name switch-options],
```

```
[edit routing-instances routing-instance-name protocols evpn],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options],
[edit routing-instances routing-instance-name vlans vlan-name switch-options],
[edit vlans vlan-name switch-options]
```

Description

Limit the number of entries that can be learned on an interface through the MAC-IP bindings database. Depending on the Juniper Networks device, this limit can be applied to EVPN instances, bridge domains configured in a virtual-switch routing instance, or VLANs configured in a virtual-switch routing instance.

To apply this limit systemwide, use the `global-mac-ip-limit` statement at the `[edit protocols l2-learning]` hierarchy level.

Options

number Specify the maximum number of MAC-IP bindings per interface. After that maximum is reached, no additional MAC-IP entries are added to the database.

- **Range:** 1 through 1024.
- **Default:** 1024

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R2.

RELATED DOCUMENTATION

EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression

mac-ip-table-size | **10231**

interface-mac-limit

IN THIS SECTION

- Syntax | 8910
- Hierarchy Level | 8910
- Description | 8911
- Default | 8912
- Options | 8912
- Required Privilege Level | 8912
- Release Information | 8912

Syntax

```
interface-mac-limit {
    limit;
    disable;
    packet-action;
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name bridge-options],
[edit bridge-domains bridge-domain-name bridge-options interface interface-name],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name bridge-options],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name bridge-options
interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name switch-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name switch-options
interface interface-name],
[edit logical-systems logical-system-name switch-options],
```

```
[edit logical-systems logical-system-name switch-options interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options
interface interface-name],
[edit routing-instances routing-instance-name switch-options],
[edit routing-instances routing-instance-name switch-options interface interface-name],
[edit switch-options],
[edit switch-options],
[edit switch-options interface interface-name],
[edit switch-options interface interface-name],
[edit vlans vlan-name switch-options],
[edit vlans vlan-name switch-options interface interface-name]
```

Description

Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.

NOTE: For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the **interface-mac-limit** statement or changing the `interface-mac-limit` configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the **interface-mac-limit** statement or use the **commit at** configuration statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers or switches by using the **clear bridge mac-table** command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

NOTE: In devices that support Enhanced Layer 2 Software (ELS), if you are configuring interface-mac-limit for an interface and the `family ethernet-switching` configuration knob is not enabled on that interface, you must enable it.

In devices that do not support ELS, if you are configuring interface-mac-limit for an interface and the `family bridge` configuration knob is not enabled on that interface, you must enable it.

Default

The default MAC limit varies with the platform.

Options

`disable`—Disables the global interface-mac-limit configuration on an interface and sets the maximum interface-mac-limit that is permitted on the device.

`limit`—Sets the maximum number of MAC addresses learned from an interface.

- **Range:** 1 through <default MAC limit> MAC addresses per interface. Range is platform specific.

If you configure both `disable` and `limit`, disable takes precedence and packet-action is set to `none`. The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for the `switch-options` statement added in Junos OS Release 9.2.

Support for top-level configuration for the `virtual-switch` type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

[edit switch-options], [edit switch-options interface *interface-name*], [edit vlans *vlan-name* switch-options], and [edit vlans *vlan-name* switch-options interface *interface-name*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports

Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port

interface-mac-limit (VPLS)

IN THIS SECTION

- Syntax | 8913
- Hierarchy Level | 8913
- Description | 8914
- Options | 8914
- Required Privilege Level | 8914
- Release Information | 8915

Syntax

```
interface-mac-limit limit {
    packet-action drop;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
site site-name interfaces interface-name],
[edit routing-instances routing-instance-name protocols evpn],
[edit routing-instances routing-instance-name protocols evpn interface interface-name],
[edit routing-instances routing-instance-name protocols vpls site site-name interfaces interface-
name]
```

Description

Specify the maximum number of media access control (MAC) addresses that can be learned by the EVPN or VPLS routing instance. You can configure the same limit for all interfaces configured for a routing instance. You can also configure a limit for a specific interface.

Starting with Junos OS Release 12.3R4, if you do not configure the parameter to limit the number of MAC addresses to be learned by a VPLS instance, the default value is not effective. Instead, if you do not include the `interface-mac-limit` option at the `[edit logical-systems` *`logical-system-name`* `routing-instances` *`routing-instance-name`* `protocols vpls site` *`site-name`* `interfaces` *`interface-name`*`]`, hierarchy level, this setting is not present in the configuration with the default value of 1024 addresses. If you upgrade a router running a Junos OS release earlier than Release 12.3R4 to Release 12.3R4 or later, you must configure the `interface-mac-limit` option with a valid value for it to be saved in the configuration.

> **NOTE**: When you change the number of stored MAC addresses that a device can learned remotely over an ESI interface, the limit does not immediately take effect. To ensure that you are using the configured MAC address limit immediately, apply the following clear MAC address table command on the multihomed PE devices before you change the number of stored MAC addresses.
>
> - On the MX series routers, use **clear bridge mac-table**.
> - On the ACX series routers, EX series and QFX series switches, use **clear ethernet-switching table**.

Options

`limit`—Number of MAC addresses that can be learned from each interface.

- **Range:** 1 through 131,071 MAC addresses

  > **NOTE**: For M120 devices only, the range is 16 through 65,536 MAC addresses.

- **Default:** 1024 addresses

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for EVPNs introduced in Junos OS Release 13.2 on MX 3D Series routers.

Support for EVPNs introduced in Junos OS Release 14.2 on EX Series switches.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

Limiting the Number of MAC Addresses Learned from an Interface

interface (EVPN Routing Instances) | **8882**

mac-table-size (VPLS) | **10269**

interface-map

IN THIS SECTION

- Syntax | **8915**
- Hierarchy Level | **8916**
- Description | **8916**
- Required Privilege Level | **8916**
- Release Information | **8916**

Syntax

```
interface-map {
    collector interface-name;
    file-specification variant-number;
    interface-name {
```

```
        collector interface-name;
        file-specification variant-number;
    }
}
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Match an input interface with a flow collector interface and apply the preset file specifications to the input interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Interface Mappings

interface-mib (Enhanced Subscriber Management)

IN THIS SECTION

Syntax | 8917

- Hierarchy Level | 8917
- Description | 8917
- Default | 8917
- Required Privilege Level | 8917
- Release Information | 8917

Syntax

```
interface-mib;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name]
```

Description

Enable representation of the Interfaces MIB for the specified dynamic interface.

To achieve maximum performance with enhanced subscriber management, we recommend that you *not* enable representation of the Interfaces MIB on all dynamic subscriber interfaces.

Default

If you do not include the `interface-mib` statement, representation of the Interfaces MIB on dynamic subscriber interfaces is disabled by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

Support for the `interface-mib-oids.sh` script added in 18.1R1.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

Junos OS Enhanced Subscriber Management Overview

interface (MLD Snooping)

IN THIS SECTION

- Syntax | 8918
- Hierarchy Level | 8919
- Description | 8919
- Options | 8919
- Required Privilege Level | 8919
- Release Information | 8919

Syntax

```
interface (all | interface-name) {
    group-limit  limit;
    host-only-interface;
    immediate-leave;
    multicast-router-interface;
    static {
        group ip-address {
            source ip-address;
        }
    }
}
```

Hierarchy Level

```
[edit protocols mld-snooping vlan (all | vlan-name)]
[edit routing-instances instance-name protocols mld-snooping vlan vlan-name]
```

Description

For MLD snooping, configure an interface as a static multicast-router interface, a host-side interface, or a static member of a multicast group.

Options

all (All EX Series switches except EX9200) All interfaces in the VLAN.

interface-name Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

Support at the `[edit routing-instances instance-name protocols mld-snooping vlan vlan-name]` hierarchy introduced in Junos OS Release 13.3 for EX Series switches.

Support for the `group-limit`, `host-only-interface`, and the `immediate-leave` statements introduced in Junos OS Release 13.3 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring MLD Snooping on SRX Series Devices

mld-snooping | **10783**

Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)

Understanding MLD Snooping

interface-mode

IN THIS SECTION

- Syntax | 8920
- Hierarchy Level | 8920
- Description | 8920
- Options | 8921
- Required Privilege Level | 8921
- Release Information | 8921

Syntax

```
interface-mode (access | trunk <inter-switch-link>);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family bridge],
[edit interfaces interface-name unit logical-unit-number family ethernet-switching],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family bridge]
```

Description

NOTE: This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see port-mode. For ELS details, see Using the Enhanced Layer 2 Software CLI.

QFX3500 and QFX3600 standalone switches—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the `trunk` option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the `vlan-id` or `vlan-id-list` statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the `access` option to

accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the `vlan-id` statement.

> **NOTE**: On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure `interface-mode` and `irb` for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see Configuring a Trunk Interface on a Bridge Network.

Options

`access`—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the `vlan-id` statement.

`trunk`—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the `vlan-id` or `vlan-id-list` statement.

`trunk inter-switch-link`—For a private VLAN, configure the InterSwitch Link protocol (ISL) on a trunk port of the primary VLAN in order to connect the switches composing the PVLAN to each other. You do not need to configure an ISL when a PVLAN is configured on a single switch. This configuration specifies whether the particular interface assumes the role of interswitch link for the PVLAN domains of which it is a member. This option is supported only on MX240, MX480, and MX960 routers in enhanced LAN mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

`inter-switch-link` option introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

RELATED DOCUMENTATION

- Configuring Access Mode on a Logical Interface
- Configuring a Logical Interface for Trunk Mode

Example: Connecting Access Switches with ELS Support to a Distribution Switch with ELS Support

Tunnel Services Overview

Tunnel Interface Configuration on MX Series Routers Overview

interface-monitor

IN THIS SECTION

- Syntax | 8922
- Hierarchy Level | 8922
- Description | 8922
- Options | 8923
- Required Privilege Level | 8923
- Release Information | 8923

Syntax

```
interface-monitor interface-name {
    weight number;
}
```

Hierarchy Level

```
[edit chassis cluster redundancy-group group-number ]
```

Description

Specify a redundancy group interface to be monitored for failover and the relative weight of the interface.

For a redundancy group to automatically failover to another node, its interfaces must be monitored. Interface monitoring monitors the physical status of an interface.

When you configure a redundancy group, you can specify a set of interfaces that the redundancy group is to monitor for status (or “health”) of interface to determine whether the interface is up or down. A monitored interface can be a child interface of any of its redundant Ethernet interfaces.

When you configure an interface for a redundancy group to monitor, you assign a weight to the interface. Every redundancy group has a threshold tolerance value initially set to 255. When an interface monitored by a redundancy group becomes unavailable, its weight is subtracted from the redundancy group's threshold. When a redundancy group's threshold reaches 0, the redundancy group fails over to the other node in the cluster.

Options

name Name of the interface to monitor

weight Weight assigned to this interface that influences failover

- **Range:** `0-255`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

interface (Monitoring)

IN THIS SECTION

Syntax | 8924

- Hierarchy Level | 8924
- Description | 8924
- Options | 8924
- Required Privilege Level | 8925
- Release Information | 8925

Syntax

```
interface interface-name {
    engine-id number;
    engine-type number;
    input-interface-index number;
    output-interface-index number;
    source-address address;
}
```

Hierarchy Level

```
[edit forwarding-options monitoring group-name inet output]
```

Description

Specify the output interface for monitored traffic.

Options

interface-name—Name of the interface.

engine-id *number*—Identity of the monitoring interface.

engine-type *number*—Type of the monitoring interface.

input-interface-index *number*—Input interface index for records from the interface.

output-interface-index *number*—Output interface index for records from the interface.

source-address *address*—Address used for generating packets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Passive Flow Monitoring

interface (MPLS)

IN THIS SECTION

- Syntax | 8925
- Hierarchy Level | 8926
- Description | 8926
- Default | 8926
- Options | 8926
- Required Privilege Level | 8926
- Release Information | 8926

Syntax

```
interface (all | interface-name);
```

Hierarchy Level

```
[edit protocols mpls]
```

Description

Enable MPLS on all interfaces on the switch or on the specified interface.

Default

MPLS is disabled.

Options

all—All interfaces on the switch.

interface-name—Name of an interface:

- Aggregated Ethernet—**ae*x***
- Gigabit Ethernet—**ge-*fpc/pic/port***

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

> *Example: Configuring MPLS on EX8200 and EX4500 Switches*
>
> *Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS*
>
> *Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect*
>
> *Configuring MPLS on EX8200 and EX4500 Provider Switches*

interface (Multicast Scoping)

IN THIS SECTION

- Syntax | 8927
- Hierarchy Level | 8927
- Description | 8927
- Options | 8927
- Required Privilege Level | 8928
- Release Information | 8928

Syntax

```
interface [ interface-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast scope scope-name],
[edit logical-systems logical-system-name routing-options multicast scope scope-name],
[edit routing-instances routing-instance-name routing-options multicast scope scope-name],
[edit routing-options multicast scope scope-name]
```

Description

Configure the set of interfaces for multicast scoping.

Options

interface-names—Names of the interfaces on which to configure scoping. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify all.

> **NOTE**: You cannot apply a scoping policy to a specific routing instance. All scoping policies are applied to all routing instances. However, you can apply the `scope` statement to a specific routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Examples: Configuring Administrative Scoping

multicast (routing-options) | **11063**

interface (Multicast Static Routes)

IN THIS SECTION

- Syntax | **8929**
- Hierarchy Level | **8929**
- Description | **8929**
- Options | **8930**
- Required Privilege Level | **8930**
- Release Information | **8930**

Syntax

```
interface interface-names {
    disable;
    maximum-bandwidth bps;
    no-qos-adjust;
    reverse-oif-mapping {
        no-qos-adjust;
    }
    subscriber-leave-timer seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Description

Enable multicast traffic on an interface.

By default, multicast packets are forwarded by enabling Protocol Independent Multicast (PIM) on an interface. PIM adds multicast routes into the routing table.

You can also configure multicast packets to be forwarded over a static route, such as a static route associated with an LSP next hop. Multicast packets are accepted on an interface and forwarded over a static route in the forwarding table. This is useful when you want to enable multicast traffic on a specific interface without configuring PIM on the interface.

You cannot enable multicast traffic on an interface and configure PIM on the same interface simultaneously.

Static routes must be configured before you can enable multicast on an interface. Configuring the `interface` statement alone does not install any routes into the routing table. This feature relies on the static route configuration.

Options

interface-names—Name of one or more interfaces on which to enable multicast traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Example: Defining Interface Bandwidth Maximums

Example: Configuring Multicast with Subscriber VLANs

interface (Multichassis Protection)

IN THIS SECTION

- Syntax | 8931
- Hierarchy Level | 8931
- Description | 8931
- Required Privilege Level | 8931
- Release Information | 8931

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit multi-chassis multi-chassis-protection peer]
```

Description

Specify the name of the interface that is being used as an interchassis link-protection link (ICL-PL). The two switches hosting a multichassis link aggregation group (MC-LAG) use this link to pass Inter-Chassis Control Protocol (ICCP) and data traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

interface (MVRP)

IN THIS SECTION

- Syntax | 8932
- Syntax | 8932
- Hierarchy Level | 8933
- Description | 8933
- Default | 8933

- Options | 8933
- Required Privilege Level | 8933
- Release Information | 8934

Syntax

```
interface (all | interface-name)  {
    disable;
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    registration (forbidden | normal);
}
```

Syntax

```
interface (all | interface-name)  {
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    point-to-point;
    registration (forbidden | normal | restricted);
}
```

Hierarchy Level

```
[edit protocols mvrp]
```

```
[edit routing-instances routing-instance-name protocols mvrp] (for virtual switch instance type)
```

```
[edit logical-systems logical-system-name protocols mvrp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
mvrp] (for virtual switch instance type)
```

Description

Specify interfaces on which to configure Multiple VLAN Registration Protocol (MVRP).

> **NOTE**: On QFX Series switches, you must configure specific interfaces—you cannot specify `interface all`. You can enable MVRP on an interface range.

Default

By default, MVRP is disabled.

Options

`all`—All interfaces on the switch.

`interface-name`—Names of interface to be configured for MVRP.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches

Configuring Multiple VLAN Registration Protocol (MVRP) on Switches

Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers

Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration

Verifying That MVRP Is Working Correctly

interface-name

IN THIS SECTION

- Syntax | 8934
- Hierarchy Level | 8935
- Description | 8935
- Options | 8935
- Required Privilege Level | 8935
- Release Information | 8935

Syntax

```
interface-name interface-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim default-vpn-source],
[edit protocols pim default-vpn-source]
```

Description

Specify the primary loopback address configured in the default routing instance to use as the source address when PIM hello messages, join messages, and prune messages are sent over multicast tunnel interfaces for interoperability with other vendors' routers.

Options

interface-name—Primary loopback address configured in the default routing instance to use as the source address when PIM control messages are sent. Typically, the **lo0.0** interface is specified for this purpose.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

interface-name

IN THIS SECTION

- Syntax | 8936
- Hierarchy Level | 8936
- Description | 8936
- Required Privilege Level | 8936

Release Information | 8936

Syntax

```
interface-name;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication username-include],
```

Description

Append the interface name and VLAN ID or stacked VLAN ID to the username string used for authentication. The appended information takes the following format:

- For single VLAN—`<interface-name>:<4-digit-vlan-id>`
- For stack VLANs—`<interface-name>:<4-digit-svlan-id>-<4-digit-vlan-id>`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

interface-name (DHCP Local Server)

IN THIS SECTION

- Syntax | 8937
- Hierarchy Level | 8937
- Description | 8937
- Required Privilege Level | 8937
- Release Information | 8938

Syntax

```
interface-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify that the interface name is concatenated with the username during the subscriber authentication or DHCP client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients

interface-name (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 8938
- Hierarchy Level | 8938
- Description | 8939
- Required Privilege Level | 8939
- Release Information | 8939

Syntax

```
interface-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-
include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
```

```
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify that the interface name is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients

interface (Next-Hop Group)

IN THIS SECTION

- Syntax | 8940
- Hierarchy Level | 8940
- Description | 8940
- Options | 8940

Required Privilege Level | 8940
Release Information | 8940

Syntax

```
interface interface-name {
    next-hop address;
}
```

Hierarchy Level

```
[edit forwarding-options next-hop-group group-name]
```

Description

Specify the output interface for sending copies of packets elsewhere to be analyzed.

The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Next-Hop Groups to Use Multiple Interfaces to Forward Packets Used in Port Mirroring

interface-none

IN THIS SECTION

- Syntax | 8941
- Hierarchy Level | 8941
- Description | 8941
- Required Privilege Level | 8941

Syntax

```
interface-none;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name east-interface]
```

```
[edit protocols protection-group ethernet-ring ring-name west-interface]
```

Description

Designates port as not used for Ethernet ring protection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- Ethernet Ring Protection Switching Overview
- Ethernet Ring Protection Using Ring Instances for Load Balancing
- Example: Configuring Ethernet Ring Protection Switching on EX Series Switches
- Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

interface-notification (programmable-rpd)

IN THIS SECTION

- Syntax | 8942
- Hierarchy Level | 8942
- Description | 8943
- Default | 8943
- Options | 8943
- Required Privilege Level | 8943
- Release Information | 8943

Syntax

```
interface-notification name;
```

Hierarchy Level

```
[edit logical-systems name routing-options programmable-rpd client id],
[edit routing-options programmable-rpd client id]
```

Description

Restrict interface event notifications from the programmable routing protocol process (prpd) to specified JET clients and interfaces. The prpd provides public APIs to program routing systems, making it possible for users to directly access the APIs to customize, create, and modify behavior of their network.

Default

By default, no restrictions are imposed and JET clients are notified of all interfaces.

Options

name	Interface name.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

- show programmable-rpd clients | 23287
- traceoptions (programmable-rpd) | 29651
- purge-timeout (programmable-rpd) | 13870

interface-num

IN THIS SECTION

- Syntax | 8944
- Hierarchy Level | 8944

- Description | 8944
- Options | 8945
- Required Privilege Level | 8945
- Release Information | 8945

Syntax

```
interface-num integer;
```

Hierarchy Level

```
[edit forwarding-options vxlan-routing]
```

Description

Configure the maximum number of physical interfaces reserved for use in an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) overlay network.

Juniper Networks switches support a maximum number of physical interfaces for use in either an underlay network or an EVPN-VXLAN overlay network. By default, a switch allocates 4096 physical interfaces for use in an EVPN-VXLAN network, which leaves the remaining physical interfaces for use in an underlay network. For example, a QFX5110 switch supports a maximum of 12288 physical interfaces. Subtracting the default value of 4096 interfaces from that maximum number leaves 8192 physical interfaces for use in an underlay network. The `interface-num` configuration statement enables you to re-allocate the maximum number of physical interfaces reserved for use in an EVPN-VXLAN network.

If you do not plan to use a switch in an EVPN-VXLAN network, you can set the `interface-num` configuration statement to 0, which allocates all physical interfaces for use in underlay networks.

Conversely, if you plan to use a switch in an EVPN-VXLAN network, you can re-allocate a greater number of the physical interfaces to the EVPN-VXLAN network.

NOTE: Changing the default number of physical interfaces reserved for use in an EVPN-VXLAN network causes the Packet Forwarding Engine to restart. Restarting the Packet

Forwarding Engine interrupts all forwarding operations. Therefore, we strongly recommend using this configuration statement before the EVPN-VXLAN network becomes operational.

Options

integer Maximum number of physical interfaces reserved for use in an EVPN-VXLAN overlay network on a switch.

- **Range:** QFX5110 switches: 0 through 12288. QFX5120 switches: 0 through 14336. The specified value must be a multiple of 2048—for example, 2048, 4096, 6144, and so on.
- **Default:** 4096

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

interface (OAM CFM, for EX Series Switch Only)

IN THIS SECTION

- Syntax | 8946
- Hierarchy Level | 8946
- Description | 8946
- Options | 8946
- Required Privilege Level | 8946
- Release Information | 8946

Syntax

```
interface (interface-name | ((ge- | xe-) (fpc/pic/port | fpc/pic/port.unit-number | fpc/pic/
port.unit-number vlan vlan-id)));
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id]
```

Description

Configure IEEE 802.1ag Operation, Administration, and Management (OAM) Connectivity Fault Management (CFM) support for the specified interface.

Options

interface-name—Interface to which the MEP is attached. It can be a physical Ethernet interface or a logical interface. If the interface is a trunk interface, the VLAN associated with the interface must have a VLAN ID.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)

Junos OS Network Interfaces Configuration Guide

interface (OAM Link-Fault Management)

IN THIS SECTION

- Syntax | 8947
- Hierarchy Level | 8947
- Description | 8948
- Options | 8948
- Required Privilege Level | 8948
- Release Information | 8948

Syntax

```
interface interface-name {
    apply-action-profile profile-name;
    link-discovery (active | passive);
    pdu-interval interval;
    pdu-threshold threshold-value;
    remote-loopback;
    event-thresholds {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
    }
    negotiation-options {
        allow-remote-loopback;
        no-allow-link-events;
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management]
```

Description

For Ethernet interfaces on M320, MX Series, and T Series routers, configure IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options

`interface` *`interface-name`*—Interface to be enabled for IEEE 802.3ah link fault management OAM support.

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

- Enabling IEEE 802.3ah OAM Support

interface-out (Security Forwarding Options)

IN THIS SECTION

- Syntax | 8949
- Hierarchy Level | 8949
- Description | 8949
- Required Privilege Level | 8949
- Release Information | 8949

Syntax

```
interface-out interface-name;
```

Hierarchy Level

```
[edit security forwarding-options mirror-filter filter-name]
```

Description

Specify the outgoing logical interface to be matched for mirroring.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

mirror-filter (Security Forwarding Options) | 10768

show security forwarding-options mirror-filter | 24363

interface (PIC Bundle)

IN THIS SECTION

- Syntax | 8950
- Hierarchy Level | 8950

- Description | 8950
- Options | 8950
- Required Privilege Level | 8950
- Release Information | 8950

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit interfaces pic-set pic-set-name]
```

Description

Sets the PIC bundle and the interface.

Options

- *apply-groups*– Groups from which to inherit configuration data.
- *apply-groups-except*– Do not inherit configuration data from these groups.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Interfaces

interface (PoE)

IN THIS SECTION

- Syntax | 8951
- Hierarchy Level | 8951
- Description | 8952
- Default | 8952
- Options | 8952
- Required Privilege Level | 8952
- Release Information | 8952

Syntax

```
interface  (all | interface-name) {
    disable;
    maximum-power watts;
    priority (high | low);
    telemetries {
        disable;
        duration hours;
        interval minutes;
    }
}
```

Hierarchy Level

```
[edit poe]
```

Description

Enable a PoE interface for a PoE port. The PoE interface must be enabled in order for the port to provide power to a connected powered device.

Default

The PoE interface is enabled by default

Options

- `all`— Apply the configuration to all interfaces on the SRX Series Firewall that have not been explicitly configured otherwise.
- `interface-name`— Explicitly configure a specific interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Power over Ethernet

interface (Port Mirroring)

IN THIS SECTION

Syntax | 8953

- Hierarchy Level | 8953
- Description | 8953
- Options | 8953
- Required Privilege Level | 8954
- Release Information | 8954

Syntax

```
interface (all | interface-name);
```

Hierarchy Level

For platforms without ELS:

```
[edit ethernet-switching-options analyzer name input (egress | ingress)],
[edit ethernet-switching-options analyzer name output]
```

For platforms with ELS:

```
[edit forwarding-options analyzer name input (egress | ingress)]
[edit forwarding-options analyzer name output]
[edit forwarding-options port-mirroring [instance name] family ethernet-switching output]
```

Description

Specify the interfaces for which ingressing traffic is mirrored. Specify the interface that mirrored traffic should be copied to (the output interface).

Options

all—Apply port mirroring to all interfaces on the switch (except the output interface). Mirroring a high volume of traffic can cause performance issues, so you should generally select specific input interfaces.

> **CAUTION**: Configuring **all** in a QFabric system causes all the access interfaces on all the nodes to be mirrored. Be cautious about using this option on a QFabric system.

interface-name—Apply port mirroring to the specified interface only.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Understanding Port Mirroring and Analyzers

Configuring Port Mirroring

Examples: Configuring Port Mirroring for Local Analysis

interface (Port Mirroring)

IN THIS SECTION

- Syntax | 8955
- Hierarchy Level | 8955
- Description | 8955
- Options | 8955
- Usage Guidelines | 8955
- Required Privilege Level | 8955
- Release Information | 8955

Syntax

```
interface interface-name {
    next-hop address;
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family (inet | inet6) output]
```

Description

Specify the output interface for sending copies of packets elsewhere to be analyzed.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Usage Guidelines

See Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

interface (Port Mirroring)

IN THIS SECTION

- Syntax | 8956
- Hierarchy Level | 8956
- Description | 8956
- Options | 8956
- Required Privilege Level | 8957
- Release Information | 8957

Syntax

```
interface interface-name {
    next-hop address;
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring output],
[edit forwarding-options port-mirroring family (inet | inet6) output]
```

Description

Specify the output interface for sending copies of packets elsewhere to be analyzed.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Port Mirroring

interface (Power over Ethernet)

IN THIS SECTION

- Syntax | 8957
- Hierarchy Level | 8958
- Description | 8958
- Options | 8958
- Required Privilege Level | 8958
- Release Information | 8958

Syntax

```
interface (all | all-extended | interface-name) {
    af-mode;
    disable;
    maximum-power watts;
    priority (high | low);
    telemetries {
        disable;
```

```
        duration hours;
        interval minutes;
    }
}
```

Hierarchy Level

```
[edit poe]
```

Description

Specify a PoE interface to be configured.

Options

all—All PoE interfaces on the switch that have not been individually configured for PoE. If a PoE interface has been individually configured, that configuration overrides any settings specified with **all**.

all-extended—(Junos Fusion only) All PoE extended port interfaces in a Junos Fusion that have not been individually configured for PoE. If a PoE interface has been individually configured, that configuration overrides any settings specified with **all-extended**.

interface-name—Name of the specific interface being configured.

If you use the `interface` statement without any substatements, default values are used for the remaining statements.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 17.2R1 for a Junos Fusion Provider Edge.

RELATED DOCUMENTATION

Configuring PoE Interfaces on EX Series Switches

Configuring Power over Ethernet in a Junos Fusion

interface-profile

IN THIS SECTION

- Syntax | 8959
- Hierarchy Level | 8959
- Description | 8960
- Options | 8960
- Required Privilege Level | 8960
- Release Information | 8960

Syntax

```
interface-profile profile-name {
    allow-clear;
    fields {
        field-name;
    }
    file filename;
    interval minutes;
    start-time time;
}
```

Hierarchy Level

```
[edit accounting-options]
```

Description

Create a profile to filter and collect error and packet statistics and write them to a file in the `/var/log` directory. You can specify an interface profile for either a physical or a logical interface.

Options

profile-name Name of the interface profile.

start-time *time* (Optional) Specify the start time in *yyyy-mm-dd.hh:mm* local time format for the accounting profile. Sometimes, if the time at which the records are written to the accounting file is beyond the transfer interval window of the file, drifting of accounting records occurs. To prevent drifting and to have predictable timestamps of when accounting data is written to an accounting file, we can specify the start time option for the accounting profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Interface Profile

Accounting Options Configuration

interface (Protocols DVMRP)

IN THIS SECTION

Syntax | **8961**

- Hierarchy Level | 8961
- Description | 8961
- Options | 8961
- Required Privilege Level | 8961
- Release Information | 8962

Syntax

```
interface interface-name {
    disable;
    hold-time seconds;
    metric metric;
    mode (forwarding | unicast-routing);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols dvmrp],
[edit protocols dvmrp]
```

Description

Enable DVMRP on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

> **NOTE**: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring DVMRP

interface (Protocols IGMP)

IN THIS SECTION

- Syntax | 8962
- Hierarchy Level | 8963
- Description | 8963
- Options | 8963
- Required Privilege Level | 8963
- Release Information | 8964

Syntax

```
interface interface-name {
    (accounting | no-accounting);
    disable;
    distributed;
    group-limit limit;
```

```
    group-policy [ policy-names ];
    immediate-leave;
    oif-map map-name;
    passive;
    promiscuous-mode;
    ssm-map ssm-map-name;
    ssm-map-policy ssm-map-policy-name;
    static {
        group multicast-group-address {
            exclude;
            group-count number;
            group-increment increment;
            source ip-address {
                source-count number;
                source-increment increment;
            }
        }
    }
    version version;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp],
[edit protocols igmp]
```

Description

Enable IGMP on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling IGMP

interface (Protocols IPv6 Neighbor Discovery)

IN THIS SECTION

- Syntax | 8964
- Hierarchy Level | 8965
- Description | 8965
- Options | 8965
- Required Privilege Level | 8965
- Release Information | 8966

Syntax

```
interface interface-name {
    current-hop-limit number;
    default-lifetime seconds;
    (link-mtu | no-link-mtu);
    (managed-configuration | no-managed-configuration);
    max-advertisement-interval seconds;
    min-advertisement-interval seconds;
    (other-stateful-configuration | no-other-stateful-configuration);
    nat-prefix NAT64 IPv6 address prefix;{
         (lifetime seconds);
    }
```

```
    prefix prefix {
        (autonomous | no-autonomous);
        (on-link | no-on-link);
        preferred-lifetime seconds;
        valid-lifetime seconds;
    }
    reachable-time milliseconds;
    retransmit-timer milliseconds;
    solicit-router-advertisement-unicast;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement],
[edit protocols router-advertisement]
```

Description

Configure router advertisement properties on an interface. To configure more than one interface, include the `interface` statement multiple times.

The Junos OS enters the Neighbor Discovery Protocol (NDP) packets into the routing platform cache even if there is no known route to the source.

If you are using Virtual Router Redundancy Protocol (VRRP) for IPv6, you must include the `virtual-router-only` statement on both the primary and backup VRRP on the IPv6 router.

Options

interface-name—Name of an interface. Specify the full interface name, including the physical and logical address components.

nat-prefix—NAT64 IPv6 address prefix. Specify upto 3 NAT64 IPv6 address prefix per interface for router advertisement.

lifetime—Lifetime of the NAT64 IPv6 address prefix in seconds.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

"solicit-router-advertisement-unicast" on page 28106 statement added from 15.1 Release onwards.

The *nat-prefix* and *lifetime* options added from Junos OS 22.4R1 Release onwards.

RELATED DOCUMENTATION

- Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery
- *Secure IPv6 Neighbor Discovery*

interface (Protocols IS-IS)

IN THIS SECTION

- Syntax | 8966
- Hierarchy Level | 8968
- Description | 8968
- Options | 8968
- Required Privilege Level | 8973
- Release Information | 8973

Syntax

```
interface (all | interface-name) {
    disable;
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
```

```
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        transmit-interval {
            threshold milliseconds;
            minimum-interval milliseconds;
        }
        multiplier number;
    }
    checksum;
    flood-group;
    csnp-interval (seconds | disable);
    hello-padding (adaptive | disable | loose | strict);
    interface-group-holddown-delay
    ldp-synchronization {
        disable;
        hold-time seconds;
    }
    lsp-interval milliseconds;
    max-hello-size size;
    mesh-group (value | blocked);
    no-adjacency-holddown;
    no-eligible-backup;
    no-eligible-remote-backup;
    no-ipv4-multicast;
    no-ipv6-multicast;
    no-ipv6-unicast;
    no-unicast-topology;
    node-link-protection;
    passive{
        remote-node-id address;
        remote-node-iso iso-id;
    }
    point-to-point;
    strict-dual-isis
    level level-number {
        disable;
        hello-authentication-key key;
        hello-authentication-key-chain key-chain-name;
```

```
        hello-authentication-type authentication;
        hello-interval seconds;
        hold-time seconds;
        ipv4-multicast-metric metric;
        ipv6-multicast-metric metric;
        ipv6-unicast-metric metric;
        metric metric;
        passive;
        priority number;
        te-metric metric;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Configure interface-specific IS-IS properties. To configure more than one interface, include the `interface` statement multiple times.

Enabling IS-IS on an interface (by including the `interface` statement at the `[edit protocols isis]` or the `[edit routing-instances` *`routing-instance-name`* `protocols isis]` hierarchy level), disabling it (by including the `disable` statement), and not actually having IS-IS run on an interface (by including the `passive` statement) are mutually exclusive states.

Options

all | *interface-name*

Select either automatic interfaces or specify the name of an existing interface.

- `all`—Have Junos OS create IS-IS interfaces automatically. If you include this option, disable IS-IS on the management interface (fxp0).
- *interface-name*—Name of an interface. Specify the full interface name, including the physical and logical address components.

checksum

Enable checksums for packets on this interface. Junos OS supports IS-IS checksums as documented in RFC 3358, *Optional Checksums in Intermediate System to Intermediate System (ISIS)*.

> **NOTE**: The checksum cannot be enabled with MD5 hello authentication on the same interface.

csnp-interval *seconds* | disable

Configure the interval between complete sequence number PDUs (CSNPs) on a LAN interface, or disable it. If the routing device is the designated router on a LAN, IS-IS sends CSN packets every 10 seconds.

> **NOTE**: The `csnp-interval` statement is effective only when configured on LAN interfaces.

seconds Number of seconds between the sending of CSNPs. The range is 1 through 65,535 seconds.

disable Do not send CSNPs on this interface.

disable

Disable IS-IS on the routing device, on an interface, or on a level. Enabling IS-IS on an interface (by including the `interface` statement at the `[edit protocols isis]` or the `[edit routing-instances` *`routing-instance-name`* `protocols isis]` hierarchy level), disabling it (by including the `disable` statement), and not actually having IS-IS run on an interface (by including the `passive` statement) are mutually exclusive states. IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which `family iso` is enabled.

flood-group

Enables IS-IS support for flood-group. This feature limits link-state packet data unit (PDU) flooding over IS-IS interfaces.

hello-padding adaptive | disable | loose | strict

Configure padding on hello packets to accommodate asymmetrical maximum transfer units (MTUs) from different hosts. This helps to prevent a premature adjacency Up state when one routing device's MTU does not meet the requirements to establish the adjacency.

As an OSI Layer 2 protocol, IS-IS does not support data fragmentation. Therefore, maximum packet sizes must be established and supported between two routers. During adjacency establishment, the IS-IS protocol makes sure that the link supports a packet size of 1492 bytes by padding outgoing hello packets up to the maximum packet size of 1492 bytes. This is the default behavior of the Junos OS IS-IS implementation. However, Junos OS provides an option to disable hello padding that can override this behavior.

adaptive On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state type, length, and value (TLV) tuple. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. Adaptive padding has more overhead than loose padding and is able to detect MTU asymmetry from one side of the connection. This one-sided detection can result in generation of extra link-state PDUs that are flooded throughout the network. Specify the `adaptive` option to configure enough padding to establish an adjacency to neighbors.

disable Padding is disabled on all types of interfaces for all adjacency states. Specify the `disable` option to accommodate interfaces that support less than the default packet size of 1492 bytes.

loose The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. Loose padding might not be able to detect certain situations such as asymmetrical MTUs between the routing devices. Specify the `loose` option to configure enough padding to initialize an adjacency to neighbors.

strict Padding is done on all interface types and for all adjacency states, and is continuous. Strict padding has the most overhead. The advantage is that strict padding detects MTU issues on both sides of a link. Specify the `strict` option to configure padding to allow all adjacency states with neighbors.

interface-group-holddown-delay *seconds*

Number of seconds before the routing device replaces the bandwidth based metric. This value configures the time interval that the IS-IS takes before replacing the metric with the new metric value and before flooding the new metric to the link-state PDU. When the bundle changes from a worse bandwidth based metric to a better metric the system waits for the configured time before switching to the new metric. The system also uses this time delay when a link status changes from down to up or when a member link changes status from degrade to non-degrade.

By default, the time delay for a routing device is 20 seconds.

ldp-synchronization

Enable synchronization by advertising the maximum cost metric until LDP is operational on the link. LDP distributes labels in non-traffic-engineered applications. Labels are distributed along the best path determined by IS-IS. If the

synchronization between LDP and IS-IS is lost, the label-switched path (LSP) goes down. Therefore, IS-IS and LDP synchronization is beneficial. When LDP synchronization is configured and when LDP is not fully operational on a given link (a session is not established and labels are not exchanged), IS-IS advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.

LDP synchronization is supported only on point-to-point interfaces and LAN interfaces configured as point-to-point interfaces under IS-IS. LDP synchronization is not supported during graceful restart. To advertise the maximum cost metric until LDP is operational for LDP synchronization, include the `ldp-synchronization` statement.

disable Disable LDP synchronization for IS-IS.

hold-time *seconds* The time period to advertise the maximum cost metric for a link that is not fully operational. The range is 1 through 65,535 seconds. The default is infinity.

> **NOTE**: When an interface has been in the holddown state for more than 3 minutes, a system log message with a warning level is sent. This message appears in both the messages file and the trace file.

lsp-interval *milliseconds* Configure the link-state PDU interval time, in milliseconds. By default, the routing device sends one link-state PDU packet out an interface every 100 milliseconds. The valid range is 0 through 1000 milliseconds.

To disable the transmission of all link-state PDUs, set the interval to 0. Link-state PDU throttling by use of the `lsp-interval` statement controls the flooding pace to neighboring routing devices in order to not overload them.

Also, consider that control traffic (such as link-state PDUs and related packets) might delay user traffic (information packets) because control traffic always has precedence in terms of scheduling on the routing device interface cards. Unfortunately, the control traffic transmission rate is not decreased on low-bandwidth interfaces, such as DS-0 or fractional T1 and E1 interface. Line control traffic stays the same. On a low-bandwidth circuit that is transmitting 30 full-MTU-sized packets, there is not much bandwidth left over for other types of packets.

max-hello-size *size* Modify the maximum size of IS-IS hello packets. IS-IS sends hello packets out of all IS-IS enabled interfaces to discover neighbors and form adjacencies between the

devices. Based on the actual MTU of the physical interface, you can configure upto 16000 bytes as the maximum size for IS-IS packets. The valid *size* range is 512 through 16000 bytes, while the default is 1492 bytes.

mesh-group *value* | blocked

Configure an interface to be part of a mesh group, which is a set of fully connected nodes. A mesh group is a set of routing devices that are fully connected. That is, they have a fully meshed topology. When link-state PDUs are being flooded throughout an area, each router within a mesh group receives only a single copy of a link-state PDU instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of link-state PDUs.

To create a mesh group and designate that an interface be part of the group, assign a mesh-group number to all the routing device interfaces in the group. To prevent an interface in the mesh group from flooding link-state PDUs, configure blocking on that interface.

value The number that identifies the mesh group. The valid range is 1 through 4,294,967,295 (2^{32} – 1; 32 bits are allocated to identify a mesh group).

blocked Configure the interface so that it does not flood link-state PDUs.

no-adjacency-holddown

Disable the hold-down timer for IS-IS adjacencies. A hold-down timer delays the advertising of adjacencies by waiting until a time period has elapsed before labeling adjacencies in the up state. You can disable this hold-down timer, which labels adjacencies up faster. However, disabling the hold-down timer creates more frequent link-state PDU updates and SPF computation.

no-eligible-backup

Exclude the specified interface as a backup interface for IS-IS interfaces on which link protection or node-link protection is enabled.

no-eligible-remote-backup

Disable remote LFA backup calculation for the specified interface. If remote LFA is disabled, Junos OS does not consider the interface for calculating the remote LFA next hop.

no-ipv4-multicast

Exclude an interface from IPv4 multicast topologies.

no-ipv6-multicast

Exclude an interface from IPv6 multicast topologies.

no-ipv6-unicast

Exclude an interface from IPv6 unicast topologies. This enables you to exercise control over the paths that unicast data takes through a network.

no-unicast-topology

Exclude an interface from the IPv4 unicast topologies.

node-link-protection

Enable node-link protection on the specified IS-IS interface. Junos OS creates an alternate loop-free path to the primary next hop for all destination routes that traverse a protected interface. This alternate path avoids the primary next-hop routing device altogether and establishes a path through a different routing device.

point-to-point

Configure an IS-IS interface to behave like a point-to-point connection. You can use the `point-to-point` statement to configure a LAN interface to act like a point-to-point interface for IS-IS. You do not need an unnumbered LAN interface, and it has no effect if configured on an interface that is already point-to-point. The `point-to-point` statement affects only IS-IS protocol procedures on that interface. All other protocols continue to treat the interface as a LAN interface. Only two IS-IS routing devices can be connected to the LAN interface, and both must be configured as point-to-point.

> **BEST PRACTICE**: On QFX10000 switches, we strongly recommend that you configure all IS-IS interfaces, including peer interfaces, as point-to-point interfaces. If you do not, you might experience session flaps, that is, IS-IS sessions that go down and then come back up, when IS-IS is configured in virtual routing and forwarding (VRF) instances. When you scale IS-IS in any scenario, you might also experience scaling issues if you do not configure IS-IS interfaces as point-to-point interfaces.

strict-dual-isis

The `strict-dual-isis` option was introduced in Junos release 18.RR1-S6. This option avoids protocol flap when both IPv4 and IPv6 are configured on the same interface or in the same ISIS instance. The option works by introducing a holdown delay to allow both IPv4 ARP and IPv6 ND to complete before the interface is advertised. The default hold time is 15 seconds. Use the `holddown` argument to specify the desired delay in the range of 1 to 65535 seconds.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`no-eligible-remote-backup` option introduced in Junos OS Release 14.2 for the MX Series.

`interface-group-holddown-delay` option introduced in Junos OS Release 15.2 for the MX Series.

`max-hello-size` option introduced in Junos OS Release 17.1 for the MX Series.

RELATED DOCUMENTATION

Understanding IS-IS Configuration

Example: Configuring IS-IS

Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups

Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding

interface (Protocols Layer 2 Circuit)

IN THIS SECTION

- Syntax | 8974
- Hierarchy Level | 8975
- Description | 8975
- Options | 8975
- Required Privilege Level | 8976
- Release Information | 8976

Syntax

```
interface interface-name {
    backup-neighbor address;
    bandwidth (bandwidth | ctnumber bandwidth);
    community community-name;
    connection-protection;
    (control-word | no-control-word);
    description text;
    egress-protection;
```

```
    encapsulation-type type;
    flow-label-receive;
    flow-label-receive-static;
    flow-label-transmit;
    flow-label-transmit-static;
    ignore-encapsulation-mismatch;
    ignore-mtu-mismatch;
    mtu mtu-number;
    no-revert;
    oam;
    protect-interface interface-name;
    pseudowire-status-tlv hot-standby-vc-on;
    psn-tunnel-endpoint address;
    revert-time seconds;
    send-ip-addr-list-tlv;
    static {
    switchover-delay milliseconds;
    virtual-circuit-id identifier;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit local-switching],
[edit logical-systems logical-system-name protocols l2circuit neighbor address],
[edit protocols l2circuit local-switching],
[edit protocols l2circuit neighbor address]
```

Description

Interface over which Layer 2 circuit traffic travels.

Options

interface-name Name of the interface to configure.

> **NOTE**: The commit operation fails, if the same logical interface is configured for both layer 2 circuit and ccc connection.

connection-protection Enable end-to-end protection through OAM failure detection.

flow-label-receive Advertise capability to pop flow label in receive direction to the remote provider edge (PE) device.

flow-label-receive-static Pop flow label on the pseudowire packets received from the remote PE device. The ingress PE inserts the flow label in the pseudowire packet, irrespective of the information exchanged in the signaling plane. If the egress PE cannot handle the pseudowire packet marked with the flow label, the packet is dropped.

flow-label-transmit Advertise capability to push flow label in transmit direction to the remote PE device.

flow-label-transmit-static Push flow label on the pseudowire packets sent to the remote PE device. If the incoming pseudowire packet is not marked with the flow label, the packet is dropped by the egress PE.

send-ip-addr-list-tlv On a PS service logical interface, signal that the IP address TLV is sent, allowing the access PE router to send the local CE router's address to the service PE router.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`flow-label-receive-static` and `flow-label-transmit-static` options introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring the Neighbor Interface for the Layer 2 Circuit

interface (Protocols Layer 2 VPN)

IN THIS SECTION

- Syntax | 8977
- Hierarchy Level | 8977
- Description | 8977
- Options | 8977
- Required Privilege Level | 8978
- Release Information | 8978

Syntax

```
interface interface-name {
    description text;
    remote-site-id     remote-site-id;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn],
[edit routing-instances routing-instance-name protocols l2vpn]
```

Description

Configure an interface to handle traffic for a circuit configured for the Layer 2 VPN.

Options

interface-name—Name of the interface used for the Layer 2 VPN.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Site

Configuring the Remote Site ID

interface (Protocols LDP)

IN THIS SECTION

- Syntax | 8978
- Hierarchy Level | 8979
- Description | 8979
- Default | 8979
- Options | 8979
- Required Privilege Level | 8979
- Release Information | 8979

Syntax

```
interface interface-name {
    disable;
    hello-interval seconds;
    hold-time seconds;
```

```
    transport-address (interface | loopback);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Enable LDP on one or more router interfaces.

Default

LDP is disabled on all interfaces.

Options

interface-name—Name of an interface. To configure all interfaces, specify `all`.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling and Disabling LDP

interface (Protocols Link Management)

IN THIS SECTION

- Syntax | 8980
- Hierarchy Level | 8980
- Description | 8980
- Options | 8980
- Required Privilege Level | 8980
- Release Information | 8981

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],
[edit protocols link-management te-link te-link-name]
```

Description

Specify the egress router interface.

Options

interface-name—Name of the interface to the egress router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

LMP Configuration Overview

interface (Protocols MLD)

IN THIS SECTION

- Syntax | 8981
- Hierarchy Level | 8982
- Description | 8982
- Options | 8982
- Required Privilege Level | 8982
- Release Information | 8982

Syntax

```
interface interface-name {
    (accounting | no-accounting);
    disable;
    distributed;
    group-limit  limit;
    group-policy [ policy-names ];
    group-threshold value;
    immediate-leave;
    log-interval seconds;
    oif-map [ map-names ];
    passive;
    ssm-map ssm-map-name;
    ssm-map-policy ssm-map-policy-name;
```

```
    static {
        group multicast-group-address {
            exclude;
            group-count number
            group-increment increment
            source ip-address {
                source-count number;
                source-increment increment;
            }
        }
    }
    version version;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld],
[edit protocols mld]
```

Description

Enable MLD on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling MLD

interface (Protocols MPLS)

IN THIS SECTION

- Syntax | 8983
- Hierarchy Level | 8983
- Description | 8983
- Options | 8984
- Required Privilege Level | 8984
- Release Information | 8984

Syntax

```
interface (interface-name | all) {
    disable;
    admin-group [ group-names ];
    srlg srlg-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Enable MPLS on one or more interfaces.

Options

interface-name—Name of the interface on which to configure MPLS. To configure all interfaces, specify `all`. For details about specifying interfaces, see the Junos OS Network Interfaces Library for Routing Devices.

`srlg` *srlg-name*—Name of the SRLG to associate with an interface.

The remaining options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Intermediate (Transit) and Egress Routers for Static LSPs

Example: Configuring SRLG

interface (Protocols OSPF)

IN THIS SECTION

- Syntax | 8985
- Hierarchy Level | 8986
- Description | 8986
- Options | 8986
- Required Privilege Level | 8990
- Release Information | 8990

Syntax

```
interface interface-name {
    disable;
    authentication key <key-id identifier>;
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        transmit-interval {
            threshold milliseconds;
            minimum-interval milliseconds;
        }
        multiplier number;
    }
    dead-interval seconds;
    demand-circuit;
    hello-interval seconds;
    flood-reduction;
    ipsec-sa name;
    interface-type type;
    ldp-synchronization {
        disable;
        hold-time seconds;
    }
    metric metric;
    neighbor address <eligible>;
    no-eligible-backup;
    no-interface-state-traps;
    node-link-protection;
    passive;
    poll-interval seconds;
    priority number;
    retransmit-interval seconds;
    te-metric metric;
```

```
    secondary;
    topology (ipv4-multicast | name) {
        metric metric;
    }
    transit-delay seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast) area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id],
[edit protocols (ospf | ospf3) area area-id],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast) area area-id]
```

Description

Enable OSPF routing on a routing device interface.

You must include at least one `interface` statement in the configuration to enable OSPF on the routing device.

Options

interface-name Specify the interface by IP address or interface name for OSPFv2, or only the interface name for OSPFv3. Using both the interface name and IP address of the same interface produces an invalid configuration. To configure all interfaces, you can specify `all`. Specifying a particular interface and `all` produces an invalid configuration.

> **NOTE**: For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as *interface-name*.

disable

Disable OSPF, an OSPF interface, or an OSPF virtual link. By default, control packets sent to the remote end of a virtual link must be forwarded using the default topology. In addition, the transit area path consists only of links that are in the default topology. You can disable a virtual link for a configured topology, but not for a default topology. Include the `disable` statement at the `[edit protocols ospf area area-id virtual-link neighbor-id router-id transit-area area-id topology name]` hierarchy level.

> **NOTE**: If you disable the virtual link by including the `disable` statement at the `[edit protocols ospf area area-id virtual-link neighbor-id router-id transit-area area-id]` hierarchy level, you disable the virtual link for all topologies, including the default topology. You cannot disable the virtual link only in the default topology.

dead-interval *seconds*

Specify how long OSPF waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor. The interval to wait is in seconds, and can range from 1 through 65,535 seconds. The default is four times the hello interval—40 seconds (broadcast and point-to-point networks); 120 seconds (nonbroadcast multiple access (NBMA) networks).

demand-circuit

Configure an interface as a demand circuit.

flood-reduction

Specify to send self-generated link-state advertisements (LSAs) with the DoNotAge bit set. As a result, self-originated LSAs are not reflooded every 30 minutes, as required by OSPF by default. An LSA is refreshed only when the content of the LSA changes, which reduces OSPF traffic overhead in stable topologies.

hello-interval *seconds*

Specify how often, in *seconds*, the routing device sends hello packets out the interface. The hello interval must be the same for all routing devices on a shared logical IP network. The valid range is 1 through 255 seconds. The default is 10 seconds (broadcast and point-to-point networks); 30 seconds (non-broadcast multiple access [NBMA] networks)

> **NOTE**: For EX Series and QFX Series switches, the hello interval is 10 seconds or longer.

ipsec-sa *name*

Apply the named IPsec authentication to the OSPF interface or virtual link or to an OSPFv2 remote sham link.

ldp-synchronization

Enable synchronization by advertising the maximum cost metric until LDP is operational on the link. LDP distributes labels in non-traffic-engineered applications. Labels are distributed along the best path determined by OSPF. If the synchronization between LDP and OSPF is lost, the label-switched path (LSP) goes down. Therefore, OSPF and LDP synchronization is beneficial. When LDP synchronization is configured and when LDP is not fully operational on a given link (a session is not established and labels are not exchanged), OSPF advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.

LDP synchronization is supported only on point-to-point interfaces and LAN interfaces configured as point-to-point interfaces under OSPF. LDP synchronization is not supported during graceful restart. To advertise the maximum cost metric until LDP is operational for LDP synchronization, include the `ldp-synchronization` statement.

disable

Disable LDP synchronization for IS-IS.

hold-time *seconds*

The time period to advertise the maximum cost metric for a link that is not fully operational. The range is 1 through 65,535 seconds. The default is infinity.

> **NOTE**: When an interface has been in the holddown state for more than 3 minutes, a system log message with a warning level is sent. This message appears in both the messages file and the trace file.

metric *metric*

Specify the cost of an OSPF interface. The cost is a routing metric that is used in the link-state calculation. To set the cost of routes exported into OSPF, configure the appropriate routing policy. Range is 1 through 65,535. By default, the cost of an OSPF route is calculated by dividing the reference-bandwidth value by the bandwidth of the physical interface. Any specific value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the cost of the route for that interface.

neighbor *address* <eligible>

For non-broadcast interfaces only, specify neighboring routers. On a non-broadcast interface, you must specify neighbors explicitly because OSPF does not send broadcast packets to dynamically discover their neighbors. To specify multiple neighbors, include multiple `neighbor` statements.

- *address*—IP address of a neighboring router.
- `eligible`—(Optional) Allow the neighbor to become a designated router. If you omit this option, the neighbor is not considered eligible to become a designated router.

no-eligible-backup

Exclude the specified interface as a backup interface for OSPF interfaces on which link protection or node-link protection is enabled.

no-interface-state-traps

Disable the OSPF traps for interface state changes. This statement is particularly useful for OSPF interfaces in passive mode.

> **NOTE**: The `no-interface-state-traps` statement is supported only for OSPFv2.

node-link-protection

Enable node-link protection on the specified OSPF interface. Junos OS creates an alternate loop-free path to the primary next hop for all destination routes that traverse a protected interface. This alternate path avoids the primary next-hop router altogether and establishes a path through a different router.

> **NOTE**: This feature is not supported for the OSPF IPv4 multicast topology or for the OSPFv3 IPv4 multicast or IPv6 multicast topologies because node-link protection creates alternate next-hop paths only for unicast routes.

poll-interval *seconds*

For non-broadcast interfaces only, specify how often, in *seconds*, the router sends hello packets out of the interface before it establishes adjacency with a neighbor. The valid range is from 1 to 65,535 seconds, and the default is 90 seconds.

priority *number*

Specify the routing device's priority for becoming the designated routing device. The routing device that has the highest priority value on the logical IP network or subnet becomes the network's designated router. You must configure at least one routing device on each logical IP network or subnet to be the designated router. You also should specify a routing device's priority for becoming the designated router on point-to-point interfaces.

The value *number* is the device's priority for becoming the designated router. A priority value of 0 means that the routing device never becomes the designated router. A value of 1 means that the routing device has the least chance of becoming a designated router. The range is 0 through 255, and the default is 128.

retransmit-interval *seconds*

Specify how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements (LSAs) to an interface's neighbors. The range is from 1 through 65,535 seconds, and the default is 5 seconds.

> **NOTE**: You must configure LSA retransmit intervals to be equal to or greater than 3 seconds to avoid triggering a retransmit trap, because Junos OS delays LSA acknowledgments by up to 2 seconds.

secondary

Configure an interface to belong to another OSPF area. A logical interface can be configured as primary interface only for one area. For any other area for which you configure the interface, you must configure it as a secondary interface.

strict-bfd

Enable strict bidirectional forwarding detection over an interface for OSPF.

te-metric *metric*

Metric value used by traffic engineering for information injected into the traffic engineering database. The value of the traffic engineering metric does not affect normal OSPF forwarding. Valid *metric* values can range from 1 through 65,535. The default is the IGP metric value.

transit-delay *seconds*

Set the estimated time required to transmit a link-state update on the interface. When calculating this time, make sure to account for transmission and propagation delays. The valid range is 1 through 65,535 seconds, with a default of 1 second.

> **NOTE**: You should never have to modify the transit delay time.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: You cannot run both OSPF and ethernet-tcc encapsulation between two Juniper Networks routing devices.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `topology` statement introduced in Junos OS Release 9.0.

Support for the `topology` statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

Support for the `no-interface-state-traps` statement introduced in Junos OS Release 10.3. This statement is supported only for OSPFv2.

RELATED DOCUMENTATION

> *Understanding OSPF Configurations*
>
> *Example: Configuring Multiple Address Families for OSPFv3*

interface (Protocols PIM)

IN THIS SECTION

- Syntax | **8991**
- Hierarchy Level | **8993**
- Description | **8993**
- Options | **8993**
- Required Privilege Level | **8993**
- Release Information | **8993**

Syntax

```
interface (Protocols PIM) (all | interface-name) {
    accept-join-always-from [ policy-name ... ];
    accept-remote-source;
    disable;
    multiple-triggered-joins {
    count  number;
```

```
    interval  milliseconds;
    }
    bfd-liveness-detection {
        authentication {
            algorithmalgorithm-name;
            key-chainkey-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (0 | 1 | automatic);
    }
    bidirectional {
        df-election {
            backoff-period milliseconds;
            offer-period milliseconds;
            robustness-count number;
        }
    }
    dual-dr {
        enhanced;
    }
    family (inet | inet6) {
        disable;
    }
    hello-interval seconds;
    mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse | sparse-dense);
    neighbor-policy [ policy-names ];
    override-interval milliseconds;
    priority number;
    propagation-delay milliseconds;
    reset-tracking-bit;
```

```
    version version;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
[edit protocols pim interface interface-name multiple-triggered-joins
```

Description

Enable PIM on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

PIM on Aggregated Interfaces

interface (Protocols RIFT)

IN THIS SECTION

- Syntax | 8994
- Hierarchy Level | 8995
- Description | 8995
- Options | 8995
- Required Privilege Level | 8997
- Release Information | 8997

Syntax

```
interface interface-name {
    allowed-authentication-keys (value | [set of values]);
    bfd-liveness-detection {
        minimum-interval milliseconds;
        multiplier milliseconds;
    }
    (check-common-instance-name | no-check-instance-name);
    (check-common-subnet | no-check-common-subnet);
    disable;
    lie-authentication;
    lie-origination-key;
    lie-transmit-address;
    lie-transmit-port port-number;
    metric metric;
    mode (active | advertise-subnets);
    (relax-three-way-nonce-check | no-relax-three-way-nonce-check);
    tie-receive-port port-number;
}
```

Hierarchy Level

```
[edit protocols rift]
```

Description

Configure the interfaces for the RIFT protocol.

Options

interface-name

Name of the interface on which the RIFT protocol should be configured.

allowed-authentication-keys (value | [set of values])

(Optional) Specify a single or set of values for allowed interface authentication keys (outer keys).

This allows you to set a set of key IDs that are allowed on this interface as outer security keys.

- **Range:** 1 through 255

bfd-liveness-detection

(Optional) Configure Bidirectional Forwarding Detection (BFD) options. The BFD session is automatically brought up if it is configured on both sides of the session.

- Values:

  minimum-interval *milliseconds*

  Specify the minimum transmit and receive interval.

  - **Range:** 1 through 255,000

  multiplier *milliseconds*

  Specify the detection time multiplier.

  - **Range:** 1 through 255

check-common-instance-name

(Optional) Enable check for common instance name advertised by neighboring device.

When multiple routing-instances of RIFT are running using the `routing-instance` statement, then enabling the `check-common-instance-name` option prevents forming mistaken adjacencies across different routing instances. The RIFT protocol declares links that receive a LIE with mismatched instance name as miscabled.

check-common-subnet
(Optional) Enable check for a common subnet on the neighboring device.

disable
(Optional) Disable the RIFT protocol on the specified interface.

lie-authentication (loose | none | permissive | strict)
(Optional) Specify the method to authenticate received LIEs (outer fingerprint).

- Values:

  loose Verify authentication only if present, that is, when the key ID is not 0.

  none Disable authentication checking completely.

  permissive Accept authentication if key identifier is unknown.

  strict Accept authentication only if a key is present and it is valid.

lie-origination-key lie-origination-key
(Optional) Configure the key ID used to protect sent LIEs (outer key). You can configure to set the key used to authenticate LIEs, if required.

- **Range:** 1 through 255

lie-transmit-address
(Optional) Configure the IPv4 or IPv6 IP address on which the link information elements (LIEs) should be sent. See "lie-transmit-address" on page 9744 for more information.

lie-transmit-port *port-number*
(Optional) Port on which the link information elements (LIEs) should be transmitted.

- **Range:** 512 through 65535

metric
(Optional) Specify the advertised cost of the RIFT protocol interface.

- **Range:** 1 through 134217727

mode
(Optional) Specify the mode of RIFT protocol interface.

- Values:
  - active—Run the RIFT protocol without advertising the Gigabit Ethernet interface subnets.
  - advertise-subnets—Run the RIFT protocol and advertise the Gigabit Ethernet interface subnets.

no-check-common-instance-name
(Optional) Disable check for common instance name advertised by neighboring device.

no-check-common-subnet
(Optional) Disable check for a common subnet on the neighboring device.

no-relax-three-way-nonce-check
(Optional) Reject LIEs with undefined remote and local nonce in three-way.

relax-three-way-nonce-check
(Optional) Accept LIEs with undefined remote and local nonce in three-way.

This allows relaxation of the specification to accept undefined nonces in three-way state that allows for faster link bring-up after failures, but opens a security attack possibility (reseting adjacencies through replays).

> **NOTE**: For maximum performance, the `relax-three-way-nonce-check` option should be on.
>
> For maximum security when `lie-origination-key` is used, the `relax-three-way-nonce-check` option should be on.

tie-receive-port *port-number*
(Optional) Port on which the topology information elements (TIEs) should be received.

- **Range:** 512 through 65535

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

interface (Protocols Router Discovery)

IN THIS SECTION

- Syntax | 8998
- Hierarchy Level | 8998
- Description | 8998
- Options | 8998
- Required Privilege Level | 8999
- Release Information | 8999

Syntax

```
interface interface-name {
    lifetime seconds;
    max-advertisement-interval seconds;
    min-advertisement-interval seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery],
[edit protocols router-discovery]
```

Description

Specify physical interfaces on which to configure timers for router advertisement messages.

Options

interface-name—Name of an interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, specify **all**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

interface (Protocols RSVP)

IN THIS SECTION

- Syntax | 8999
- Hierarchy Level | 9001
- Description | 9001
- Default | 9001
- Options | 9001
- Required Privilege Level | 9001
- Release Information | 9001

Syntax

```
interface interface-name {
    disable;
    (aggregate | no-aggregate);
    authentication-key key;
    bandwidth bps;
    hello-interval seconds;
```

```
link-protection {
    disable;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps;
    bypass bypass-name {
        bandwidth bps {
            ct0 bps;
            ct1 bps;
            ct2 bps;
            ct3 bps;
        }
        description text;
        class-of-service cos-value;
        hop-limit number;
        no-cspf;
        path address <strict | loose>;
        priority setup-priority reservation-priority;
        to address;
    }
    class-of-service cos-value;
    hop-limit number;
    max-bypasses number;
    no-cspf;
    no-node-protection;
    optimize-timer seconds;
    path address <strict | loose>;
    priority setup-priority reservation-priority;
    subscription percentage;
}
(reliable | no-reliable);
non-rsvp-bandwdith
    local-bw-override-threshold;
subscription percentage {
    ct0 percentage;
    ct1 percentage;
    ct2 percentage;
    ct3 percentage;
}
update-threshold threshold;
```

```
    update-threshold-max-reservable
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Enable RSVP on one or more router interfaces.

Default

RSVP is disabled on all interfaces.

Options

interface-name—Name of an interface. To configure all interfaces, specify `all`. For details about specifying interfaces, see the Junos OS Network Interfaces Library for Routing Devices.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Minimum RSVP Configuration

interface (RA Guard)

IN THIS SECTION

- Syntax | 9002
- Hierarchy Level | 9002
- Description | 9002
- Options | 9003
- Required Privilege Level | 9003
- Release Information | 9003

Syntax

```
interface interface-name {
    mark-interface (trusted | block);
    policy policy-name (stateful | stateless);
}
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard]
```

Description

Configure IPv6 Router Advertisement (RA) guard on an interface. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in a policy.

Before you can configure RA guard on an interface, you must first configure a policy at the [`edit forwarding-options access-security router-advertisement-guard`] hierarchy level. The policy is then applied to an interface at the [`edit forwarding-options access-security router-advertisement-guard interface` *`interface-name`*] hierarchy level.

> **NOTE**: If you apply an RA guard policy on an interface, you must enable RA guard on the VLAN that is associated with that interface by using the `vlan` statement at the [`edit forwarding-options access-security router-advertisement-guard`] hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Options

interface-name	Configure RA guard parameters on the specified interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

interface-range (QFX Series)

IN THIS SECTION

- Syntax | 9004
- Hierarchy Level | 9005
- Description | 9005
- Options | 9005

Required Privilege Level | 9005
Release Information | 9005

Syntax

The `vlan-id` statement is not supported on OCX Series switches.

```
interface-range  interface-range-name {
    disable;
    description text;
    ether-options {
        802.3ad aex {
            lacp {
                force-up;
            }
        }
         (auto-negotiation| no-auto-negotiation);
        (flow-control | no-flow-control);
        link-mode mode;
        speed (auto-negotiation | speed);
    }
    hold-time milliseconds down milliseconds;
    member interface-name;
    member-range starting-interface-name to  ending-interface-name;
    mtu bytes;
    unit logical-unit-number {
        description text;
        disable;
        family family-name {...}
        (traps | no traps);
        vlan-id vlan-id-number;
    }
}
```

Hierarchy Level

```
[edit interfaces]
```

Description

NOTE: The `vlan-id` statement and Fibre Channel interfaces are not supported on OCX Series switches.

NOTE: The interface range definition is supported only for Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel interfaces. Interface ranges are not supported on channelized interfaces.

Group interfaces that share a common configuration profile.

Options

interface-range-name—Name of the interface range.

NOTE: You can use regular expressions and wildcards to specify the interfaces in the member range configuration. Do not use wildcards for interface types.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Understanding Interface Ranges for Switches

Interfaces Overview for Switches

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

Junos OS Network Interfaces Library for Routing Devices

interface-range

IN THIS SECTION

- Syntax | 9006
- Hierarchy Level | 9007
- Description | 9007
- Options | 9007
- Required Privilege Level | 9008
- Release Information | 9008

Syntax

```
interface-range name {
    member-range interface-name-fpc/pic/port to interface-name-fpc/pic/port;
    member interface-name-fpc/pic/port;
    member interface-name-fpc/[low-high]/*;
    member interface-name-fpc/[pic1,pic2,pic3...picN]/port

    /*Common config is added as part of interface-range definition, as follows*/
    mtu 256;
    hold-time up 10;
    ether-options {
        flow-control;
        speed {
            100m;
        }
        802.3ad primary;
    }
```

```
}
```

Hierarchy Level

```
[edit interfaces (QFX Series, ACX Series)]
```

\

Description

Specify a set of identical interfaces as an interface group, to which you can apply a common configuration to the entire set of interfaces. This group can consist of both lexical member ranges of interfaces specified using the `member-range` *`interface-type-fpc/pic/port`* `to` *`xx-fpc/pic/port`* option (regex not supported), and of individual or non-sequential members using the `member` *`interface-type-fpc/pic/port`* option (with regex support to specify the *`fpc/pic/port`* values).

NOTE:
member-rangeunusedunused

Options

member-range Adds interfaces in lexical order. Regex is not supported.

Format:

```
member-range <start-range> to <end-range>
```

Example:

```
member-range ge-0/0/0 to ge-4/0/40;
```

member To add individual interfaces or multiple interfaces using regex.

Format:

```
member <list of interface names>
```

Example:

```
member ge-0/0/0;

member ge-0/1/1;

member ge-0/*/*;

member ge-0/[1-10]/0;

member ge-1/[1,3,6,10]/12
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Interface Ranges

interfaces (EX Series switches) | 9135

interface (Redundant Trunk Groups)

IN THIS SECTION

- Syntax | 9008
- Hierarchy Level | 9009
- Description | 9009
- Options | 9009
- Required Privilege Level | 9009
- Release Information | 9009

Syntax

```
interface interface-name <primary>;
interface interface-name;
```

Hierarchy Level

For platforms with ELS:

```
[edit switch-options redundant-trunk-group group name]
```

For platforms without ELS:

```
[edit ethernet-switching-options redundant-trunk-group group name]
```

Description

Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over as the primary link without waiting for normal STP convergence.

Options

interface *interface-name*—A logical interface or an aggregated interface containing multiple ports.

primary—(Optional) Specify one of the interfaces in the redundant group as the primary link. The interface without this option is the secondary link in the redundant group. If a link is not specified as **primary**, the software compares the two links and selects the link with the highest port number as the active link. For example, if the two interfaces are **ge-0/1/0** and **ge-0/1/1**, the software assigns **ge-0/1/1** as the active link.

Required Privilege Level

system—To view this statement in the configuration.
system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Hierarchy level `[edit switch-options]` introduced in Junos OS Release 13.2X50-D10 (ELS). (See *Using the Enhanced Layer 2 Software CLI* for information about ELS.)

RELATED DOCUMENTATION

Example: Configuring Redundant Trunk Links for Faster Recovery on EX Series Switches

Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support

Understanding Redundant Trunk Links (Legacy RTG Configuration)

interface-routes

IN THIS SECTION

- Syntax | 9010
- Hierarchy Level | 9010
- Description | 9011
- Options | 9011
- Required Privilege Level | 9011
- Release Information | 9012

Syntax

```
interface-routes {
    family (inet | inet6) {
        export {
            lan;
            point-to-point;
        }
    }
    rib-group group-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
```

```
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Associate a routing table group with the routing device's interfaces, and specify routing table groups into which interface routes are imported.

By default, IPv4 interface routes (also called direct routes) are imported into routing table `inet.0`, and IPv6 interface routes are imported into routing table `inet6.0`. If you are configuring alternate routing tables for use by some routing protocols, it might be necessary to import the interface routes into the alternate routing tables. To define the routing tables into which interface routes are imported, you create a routing table group and associate it with the routing device's interfaces.

To create the routing table groups, include the `passive` statement at the `[edit routing-options]` hierarchy level.

If you have configured a routing table, configure the OSPF primary instance at the `[edit protocols ospf]` hierarchy level with the statements needed for your network so that routes are installed in `inet.0` and in the forwarding table. Make sure to include the routing table group.

To export local routes, include the `export` statement.

To export LAN routes, include the `lan` option. To export point-to-point routes, include the `point-to-point` option.

Only local routes on point-to-point interfaces configured with a destination address are exportable.

Options

`inet`—Specify the IPv4 address family.

inet6—Specify the IPv6 address family.

lan—Export LAN routes.

point-to-point—Export point-to-point routes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

> **NOTE**: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

RELATED DOCUMENTATION

Example: Populating a Routing Table Created by Virtual Router Configuration

Example: Configuring Multiple Routing Instances of OSPF

passive (Routing Options) | **12225**

interface (Routing Instances)

IN THIS SECTION

- Syntax | **9012**
- Hierarchy Level | **9013**
- Description | **9013**
- Options | **9013**
- Required Privilege Level | **9013**
- Release Information | **9013**

Syntax

```
interface interface-name {
    description text;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Specify the interface over which the VPN traffic travels between the PE device and CE device. You configure the interface on the PE device. If the value `vrf` is specified for the `instance-type` statement included in the routing instance configuration, this statement is required.

Options

`interface-name`—Name of the interface.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

interface (VPLS Routing Instances) | **9104**

interface (Routing Options)

IN THIS SECTION

- Syntax | 9014
- Hierarchy Level | 9014
- Description | 9014
- Options | 9015
- Required Privilege Level | 9015
- Release Information | 9015

Syntax

```
interface interface-names {
    maximum-bandwidth bps;
    no-qos-adjust;
    reverse-oif-mapping {
        no-qos-adjust;
    }
    subscriber-leave-timer seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Description

Enable multicast traffic on an interface.

TIP: You cannot enable multicast traffic on an interface by using the **routing-options multicast interface** statement and configure PIM on the interface.

Options

interface-name—Names of the physical or logical interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Example: Defining Interface Bandwidth Maximums

Example: Configuring Multicast with Subscriber VLANs

interface (Scoping)

IN THIS SECTION

- Syntax | 9016
- Hierarchy Level | 9016
- Description | 9016
- Options | 9016
- Required Privilege Level | 9016
- Release Information | 9016

Syntax

```
interface [ interface-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast scope scope-name],
[edit logical-systems logical-system-name routing-options multicast scope scope-name],
[edit routing-instances routing-instance-name routing-options multicast scope scope-name],
[edit routing-options multicast scope scope-name]
```

Description

Configure the set of interfaces for multicast scoping.

Options

interface-names—Names of the interfaces to scope. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Example: Configuring Multicast Snooping*

interface (Secure Access Port)

IN THIS SECTION

- Syntax | 9017
- Hierarchy Level | 9017
- Description | 9017
- Options | 9018
- Required Privilege Level | 9018
- Release Information | 9018

Syntax

```
interface (all |  interface-name) {
    allowed-mac mac-address-list;
    (dhcp-trusted | no-dhcp-trusted);
     mac-limit limit action action;
    no-allowed-mac-log;
    static-ip  ip-address {
        vlan vlan-name;
        mac mac-address;
    }
}
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port]
```

Description

Apply port security features to all interfaces or to the specified interface.

The remaining statements are explained separately. See CLI Explorer.

Options

all—Apply port security features to all interfaces. Does not apply to QFabric systems.

interface-name—Apply port security features to the specified interface.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Understanding How to Protect Access Ports from Common Attacks

Understanding MAC Limiting and MAC Move Limiting for Port Security

Understanding and Using Trusted DHCP Servers

Configuring MAC Limiting (QFX Switches)

interface (Security NAT ARP Proxy)

IN THIS SECTION

- Syntax | 9019
- Hierarchy Level | 9019
- Description | 9019
- Options | 9019
- Required Privilege Level | 9019
- Release Information | 9019

Syntax

```
interface interface-name {
    address ip-address {
        to ip-address;
    }
}
```

Hierarchy Level

```
[edit security nat proxy-arp]
```

Description

Specify the interface on which the ARP proxy is to be configured. It should be a logical interface.

Options

interface-name—Name of the logical interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

interface (Security NAT NDP Proxy)

IN THIS SECTION

- Syntax | 9020
- Hierarchy Level | 9020
- Description | 9020
- Options | 9020
- Required Privilege Level | 9021
- Release Information | 9021

Syntax

```
interface interface-name {
    address ip-address {
        to ip-address;
    }
}
```

Hierarchy Level

```
[edit security nat proxy-ndp]
```

Description

Specify the interface on which the NDP proxy is to be configured. It should be a logical interface.

Options

interface-name—Name of the logical interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

interface (Security Source NAT)

IN THIS SECTION

- Syntax | 9021
- Hierarchy Level | 9021
- Description | 9021
- Options | 9022
- Required Privilege Level | 9022
- Release Information | 9022

Syntax

```
interface (port-overloading off | port-overloading-factor number);
```

Hierarchy Level

```
[edit security nat source]
```

Description

Enable interface NAT with or without port overloading.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10.

interface (Security Source NAT Rule Set)

IN THIS SECTION

- Syntax | 9022
- Hierarchy Level | 9023
- Description | 9023
- Options | 9023
- Required Privilege Level | 9023
- Release Information | 9023

Syntax

```
interface {
    persistent-nat {
        address-mapping;
        inactivity-timeout seconds;
        max-session-number value;
        permit (any-remote-host | target-host | target-host-port);
```

```
    }
}
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name then source-nat]
```

Description

Enable interface NAT with or without port overloading.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

interface (Service Interfaces)

IN THIS SECTION

- Syntax | 9024
- Hierarchy Level | 9024
- Description | 9024
- Options | 9024
- Required Privilege Level | 9024

Release Information | 9024

Syntax

```
interface service-interface-name;
```

Hierarchy Level

```
[edit services service-device-pools pool pool-name]
```

Description

Assign a service interface to a service interface pool. You specify more than one interface for each pool. The interfaces are used to balance traffic loads. For an L2TP tunnel group, the interfaces must be inline service interfaces (si). For a hybrid access tunnel group, the interfaces must be pseudowire service interfaces (ps).

Options

service-interface-name	Name of the service interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support for ps service interfaces added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

interface (Service Redundancy Group Monitor)

IN THIS SECTION

- Syntax | 9025
- Hierarchy Level | 9026
- Description | 9026
- Options | 9026
- Required Privilege Level | 9026
- Release Information | 9026

Syntax

Syntax (monitor)

```
interface interface-name;
```

Syntax (monitor-object)

```
interface {
           interface-name name {
              weight weight;
           }
           threshold threshold;
          }
```

Hierarchy Level

```
[edit chassis high-availability services-redundancy-group name monitor],
[edit chassis high-availability services-redundancy-group name monitor monitor-object name ]
```

Description

Configure interface monitoring for a service redundancy group.

Options

interface *interface-name*

Name of the physical interface to monitor.

weight—

Weight assigned to the monitored interface. A higher weight value indicates a greater importance.

- **Range:** 1 through 1000

threshold (monitor-object) —interface monitoring threshold value.

Within a monitor-object, whenever the particular object fails to trigger a failover as result of interface monitoring, the system considers the event as monitoring failure and adds a count to interface monitoring.

When the count exceeds the threshold value of interface monitoring, the system adds the count to parent monitoring-object's threshold value. Monitoring-object's threshold also includes threshold values of IP monitoring and BFD monitoring.

When the sum of the weights of all monitor objects bound to the SRG reaches a value equal to or greater than the threshold value configured on the SRG, the system triggers a monitor failure for that SRG. SRG fails over to the other node.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

interface-service (Services Interfaces)

IN THIS SECTION

- Syntax | 9027
- Hierarchy Level | 9027
- Description | 9027
- Options | 9027
- Required Privilege Level | 9028
- Release Information | 9028

Syntax

```
interface-service {
    load-balancing-options {
        hash-keys {
            egress-key (destination-ip | source-ip);
            ingress-key (destination-ip | source-ip);
        }
    }
    service-interface name;
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the device name for the interface service Physical Interface Card (PIC).

Options

`service-interface` *name*—Name of the service device associated with the interface-wide service set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Service Sets to be Applied to Services Interfaces

interface (Services)

IN THIS SECTION

- Syntax | 9028
- Hierarchy Level | 9028
- Description | 9029
- Required Privilege Level | 9029
- Release Information | 9029

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit services unified-access-control infranet-controller hostname]
```

Description

Specify the SRX Series interface through which the IC Series device should connect.

This statement is required when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

port (Services) | 12880

password (Services) | 12246

interface (Services Flow Tap)

IN THIS SECTION

- Syntax | 9030
- Hierarchy Level | 9030
- Description | 9030
- Options | 9030
- Required Privilege Level | 9030
- Release Information | 9030

Syntax

```
interface sp-fpc/pic/port.logical-unit-number;
```

Hierarchy Level

```
[edit services flow-tap]
```

Description

Specify the AS PIC interface used with the flow-tap application. Any AS PIC available in the router can be assigned, and any logical interface on the AS PIC can be used.

Options

sp-*fpc/pic/port.logical-unit-number* Use the specified services interface for flow-tap service.

You cannot configure flow-tap services on channelized interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring the Flow-Tap Interface

interface (Services PIC)

IN THIS SECTION

- Syntax | 9031
- Hierarchy Level | 9031
- Description | 9031
- Options | 9032
- Required Privilege Level | 9032
- Release Information | 9032

Syntax

```
[interface interface-name];
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name system service-pics]
```

Description

Specify one or more of the MS-MPC service interfaces that represent the service PICs used for anchoring subscriber-aware services in the MX Series router. The following conditions are applicable to the services PIC interfaces configured here:

- If an aggregated multiservices interface (ams) is specified in this statement, the ams must already be defined at the `[edit interfaces]` hierarchy level.
- The PIC must have the `jservices-hcm`, `jservices-mss`, `jservices-jdpi`, `jservices-pcef`, and `jservices-crypto-base` packages configured at the `[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]` hierarchy level.
- The appropriate services group configuration must be applied to the PIC:

- For each service PIC that requires application identification but not HTTP header enrichment, apply the `tdf-services-xlp-dpi` group.
- For each service PIC that requires both application identification and HTTP header enrichment, configure the `tdf-services-xlp-dpi-with-hcm` group.

- If an MS-MPC service interface is a member of an AMS, then that member interface cannot be specified here. For example, if mams-2/0/0 is a member interface of ams0, then ms-2/0/0/ cannot be directly specified here.

> **NOTE**: If an AMS (for example ams0) is used for the services PIC, then load balancing is performed to distribute subscriber-aware services among the member interfaces. Otherwise, load balancing is not performed.

Options

interface-name

Name of the interface representing the services PIC.

- **Syntax:** The interface must be a valid multiservices interface (ams*n* or ms-*a*/*b*/0, where *n* is the ams number, *a* is the Flexible PIC Concentrator [FPC] slot number, and *b* is the PIC slot number); for example, ams0 or ms-1/0/0.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Service PICs

TDF Gateway Service PICs and Session PICs for Subscriber-Aware Traffic Treatment

interface (Session PICs)

IN THIS SECTION

- Syntax | 9033
- Hierarchy Level | 9033
- Description | 9033
- Options | 9034
- Required Privilege Level | 9034
- Release Information | 9034

Syntax

```
[interface interface-name];
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name system session-pics]
```

Description

Specify one or more of the MS-MPC service interfaces that represent the session PICs used for the control plane in the TDF gateway. The following conditions are applicable to the session PIC interfaces configured here:

- If an aggregated multiservices interface (ams) is specified in this statement, the ams must already be defined at the `[edit interfaces]` hierarchy level.
- The `tdf-session-xlp` group configuration must be applied to the PIC.
- The session PIC must have the `jservices-mobile` package configured at the `[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]` hierarchy level.

- If a session PIC interface is a member of an AMS, then that member interface cannot be specified here. For example, if mams-2/0/0 is a member interface of ams0, then ms-2/0/0/ cannot be directly specified here.

Options

interface-name Name of the interface representing the services PIC.

- **Syntax:** The interface must be a valid multiservices interface (ams*n* or ms-*a*/*b*/0, where *n* is the ams number, *a* is the Flexible PIC Concentrator [FPC] slot number, and *b* is the PIC slot number); for example, ams0 or ms-1/0/0.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Session PICs

TDF Gateway Service PICs and Session PICs for Subscriber-Aware Traffic Treatment

interface-set

IN THIS SECTION

- Syntax | 9035
- Hierarchy Level | 9035
- Description | 9035
- Options | 9035

Required Privilege Level | 9035
Release Information | 9035

Syntax

```
interface-set interface-set-name {
    interface-name;
}
```

Hierarchy Level

```
[edit firewall],
[edit logical-systems logical-system-name firewall]
```

Description

Configure an interface set.

Options

interface-name—Names of each interface to include in the interface set. You must specify more than one name.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Filtering Packets Received on an Interface Set Overview

interface-set (ANCP)

IN THIS SECTION

- Syntax | 9036
- Hierarchy Level | 9036
- Description | 9036
- Options | 9037
- Required Privilege Level | 9037
- Release Information | 9037

Syntax

```
interface-set interface-set-name {
    access-identifier identifier-string;
    underlying-interface underlying-interface-name;
}
```

Hierarchy Level

```
[edit protocols ancp interfaces]
```

Description

Identify a group of VLANs on which traffic is sent to a subscriber identified by the access-loop circuit identifier.

Options

interface-set-name—Name of a group of VLANs that carry traffic to the subscriber identified by the access loop circuit identifier.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent

Associating an Access Node with Subscribers for ANCP Agent Operations

interface-set (Dynamic Profiles)

IN THIS SECTION

- Syntax | 9038
- Hierarchy Level | 9038
- Description | 9038
- Options | 9038
- Required Privilege Level | 9039
- Release Information | 9039

Syntax

```
interface-set interface-set-name {
    interface interface-name {
        unit logical-unit-number;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces]
```

Description

For MX Series routers with enhanced queuing DPCs or MPC/MIC modules, configure an interface set for dynamic CoS.

Options

interface-set-name—Name of the interface set to be configured or one of the following Junos OS predefined variables:

- $junos-interface-set-name—Predefined variable that, when used, is replaced with the interface-set obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.
- $junos-phy-ifd-interface-set-name—Locally generated interface set name associated with the underlying physical interface in a dynamic profile. This predefined variable enables you to group all the subscribers on a specific physical interface so that you can apply services to the entire group of subscribers.

  Another use case for this predefined variable is to conserve CoS resources in a mixed business and residential topology by collecting the residential subscribers into an interface set associated with the physical interface, so that a level 2 node is used for the interface set rather than for each residential interface. Otherwise, because the business and residential subscribers share the same interface and business subscribers require three levels of CoS, then three levels are configured for each residential subscriber. That results in an unnecessary level 2 node being consumed for each residential connection, wasting CoS resources.

- $junos-pon-id-interface-set-name—Locally generated interface set name extracted from the DHCPv4 (Option 82, suboption 2) or DHCPv6 (Option 37) agent remote ID string inserted by an optical line terminal (OLT) in a passive optical network (PON). The OLT must format the agent remote ID string with a pipe symbol (|) as the delimiter between substrings. The substring extracted for the interface set name consists of the characters following the last delimiter in the agent remote ID string.

  The extracted substring identifies individual customer circuits in the PON to be aggregated into the interface set. You determine the format and contents of the substring, and configure your OLT to insert the information. Typically, the substring may include the name and port of the OLT accessed by the CPE optical network terminal (ONT).

- $junos-svlan-interface-set-name—Locally generated interface set name for use by dual-tagged VLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is *physical_interface_name - outer_VLAN_tag*.

- $junos-tagged-vlan-interface-set-name—Locally generated interface set name used for grouping logical interfaces stacked over logical stacked VLAN demux interfaces for either a 1:1 (dual-tagged; individual client) VLAN or N:1 (single tagged; service) VLAN. The format of the generated variable differs with VLAN type. For dual-tagged (client) VLANs, the format of the generated variable is *physical_interface_name - outer_VLAN_tag - inner_VLAN_tag*. For single tagged (service) VLAN, the format of the generated variable is *physical_interface_name - VLAN_tag*.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

$junos-phy-ifd-interface-set-name option added in Junos OS Release 16.1.

$junos-pon-id-interface-set-name option added in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

CoS for Interface Sets of Subscribers Overview

Configuring an Interface Set of Subscribers in a Dynamic Profile

Configuring an Interface Set of Subscribers in a Dynamic Profile

Extracting an Option 82 or Option 37 Substring to Create an Interface Set

interface-set (Dynamic VLAN Interface Sets Association)

IN THIS SECTION

- Syntax | 9040
- Hierarchy Level | 9040
- Description | 9041
- Options | 9041
- Required Privilege Level | 9041
- Release Information | 9041

Syntax

```
interface-set interface-set-name {
    interface interface-name {
        unit logical-unit-number {
            advisory-options {
                downstream-rate rate;
                upstream-rate rate;
            }
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces]
```

Description

For MX Series routers with MPC/MIC modules that face the access side of the network, associate an agent circuit identifier (ACI) or access-line-identifier (ALI) interface set with a dynamic VLAN subscriber interface for DHCP or PPPoE subscribers. To associate the interface set with a dynamic subscriber interface, you must include the `interface-set` stanza in the dynamic profile that defines the logical subscriber interface.

An ACI or ALI interface set is a logical collection of subscriber interfaces that originate at the same household or on the same access-loop port. An ACI set is created based on the receipt only of the ACI for the subscriber access line in a DHCP or PPPoE control packet. An access-line-identifier set is created based on the receipt of a trusted option, which can be the ACI, the ARI, both these identifiers, or the absence of both these identifiers.

You specify the trigger for either interface set type at the `[edit dynamic-profiles` *`profile-name`* `interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" auto-configure]` hierarchy level. For ACI interface sets, use the `agent-circuit-identifier` statement. For ALI interface sets, use the `line-identity` statement.

Options

- *interface-set-name*—Name of the ACI interface set, which is represented in a dynamic profile for a subscriber interface by the predefined variable `$junos-interface-set-name`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Dynamic VLAN Subscriber Interfaces Based on Agent Circuit Identifier Information

Agent Circuit Identifier-Based Dynamic VLANs Overview

Configuring Dynamic VLAN Subscriber Interfaces Based on Access-Line Identifiers

Access-Line-Identifier-Based Dynamic VLANs Overview

interface-set (Dynamic VLAN Interface Sets Definition)

IN THIS SECTION

- Syntax | 9042
- Hierarchy Level | 9042
- Description | 9042
- Options | 9043
- Required Privilege Level | 9043
- Release Information | 9043

Syntax

```
interface-set interface-set-name {
    interface interface-name;
    pppoe-underlying-options {
        max-sessions number;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces]
```

Description

For MX Series routers with MPC/MIC modules that face the access side of the network, configure an agent circuit identifier (ACI) or access-line-identifier (ALI) interface set for the creation of dynamic VLAN subscriber interfaces for DHCP or PPPoE subscribers based on information about the subscriber access line received in DHCP or PPPoE control packets.

An ACI or ALI interface set is a logical collection of subscriber interfaces that originate at the same household or on the same access-loop port. An ACI set is created based on the receipt only of the ACI for the subscriber access line in a DHCP or PPPoE control packet. An access-line-identifier set is created based on the receipt of a trusted option, which can be the ACI, the ARI, both these identifiers, or the absence of both these identifiers.

You specify the trigger for either interface set type at the [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" auto-configure] hierarchy level. For ACI interface sets, use the agent-circuit-identifier statement. For ALI interface sets, use the line-identity statement.

You must associate the interface set with the dynamic subscriber interface by including the interface-set stanza in the dynamic profile that defines the interface set.

Options

- *interface-set-name*—Name of the ACI interface set, which is represented in a dynamic profile by the predefined variable $junos-interface-set-name.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Defining ACI Interface Sets

Clearing Agent Circuit Identifier Interface Sets

Agent Circuit Identifier-Based Dynamic VLANs Overview

Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers

Clearing Access-Line-Identifier Interface Sets

Access-Line-Identifier-Based Dynamic VLANs Overview

interface-set (Ethernet Interfaces)

IN THIS SECTION

- Syntax | 9044
- Hierarchy Level | 9044
- Description | 9044
- Required Privilege Level | 9044
- Release Information | 9045

Syntax

```
interface-set interface-set-name {
    interface ethernet-interface-name {
        (unit unit-number | vlan-tags-outer vlan-tag);
    }
}
```

Hierarchy Level

```
[edit interfaces]
```

Description

The set of interfaces used to configure hierarchical CoS schedulers on Ethernet interfaces on the MX Series router and IQ2E PIC on M Series and T Series routers.

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

interface-set (Hierarchical Schedulers) | 9045

interface-set (Hierarchical Schedulers)

IN THIS SECTION

- Syntax | 9045
- Hierarchy Level | 9045
- Description | 9046
- Options | 9046
- Required Privilege Level | 9046
- Release Information | 9046

Syntax

```
interface-set interface-set-name {
    excess-bandwidth-share (proportional value | equal);
    internal-node;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
}
```

Hierarchy Level

```
[edit class-of-service interfaces]
```

Description

Configure hierarchical schedulers for an interface set.

Options

interface-set-name—Name of the interface set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Interface Sets

Configuring Hierarchical Schedulers for CoS

interface-set (IP Demux Interfaces)

IN THIS SECTION

- Syntax | 9047
- Hierarchy Level | 9047
- Description | 9047
- Required Privilege Level | 9047
- Release Information | 9047

Syntax

```
interface-set interface-set-name {
    interface interface-name {
        unit unit-number;
    }
}
```

Hierarchy Level

```
[edit interfaces]
```

Description

The set of interfaces used to configure hierarchical CoS schedulers for subscribers on IP demux interfaces on the MX Series router.

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Interfaces Fundamentals

interface-shared

IN THIS SECTION

- Syntax | 9048
- Hierarchy Level | 9048
- Description | 9048
- Required Privilege Level | 9048
- Release Information | 9049

Syntax

```
interface-shared;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name filter filter-name],
[edit firewall family family-name filter filter-name]
```

Description

Set the interface-shared attribute for a firewall filter.

NOTE: A firewall filter cannot be both interface-specific and interface-shared.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Interface-Shared Filters Overview

Understanding Dynamic Firewall Filters

Classic Filters Overview

Basic Classic Filter Syntax

interface-shared-with

IN THIS SECTION

- Syntax | 9049
- Hierarchy Level | 9049
- Description | 9050
- Options | 9050
- Required Privilege Level | 9050
- Release Information | 9050

Syntax

```
interface-shared-with psdn;
```

Hierarchy Level

```
[edit interfaces ge-fpc/pic/slot unit logical-unit-number],
[edit interfaces so-fpc/pic/slot unit logical-unit-number],
[edit interfaces xe-fpc/pic/slot unit logical-unit-number]
```

Description

Assign a logical interface under a shared physical interface to a Protected System Domain (PSD).

Options

n—PSD identification as a numeric value.

- **Range: 1** through **31**

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

interface-shutdown-action

IN THIS SECTION

- Syntax | 9051
- Hierarchy Level | 9051
- Description | 9051
- Default | 9051
- Options | 9051
- Required Privilege Level | 9051
- Release Information | 9052

Syntax

```
interface-shutdown-action [soft-shutdown | hard-shutdown]
```

Hierarchy Level

```
[edit switch-options]
```

Description

Configure device to shut down interfaces or temporarily disable interfaces. This action can be done in addition to the default switching device action for storm control (dropping packets).

Storm control enables the switching device to monitor traffic levels and to drop broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level or storm control bandwidth—is exceeded, thus preventing packets from proliferating and degrading the LAN.

When you include the **interface-shutdown-action** statement at the [`edit switch-options`] hierarchy level, the behavior is to temporarily disable interfaces when the storm control level threshold exceeds MAC limit or MAC move limit and packet action shutdown is configured..

When the configuration statement **recovery-timeout** is included under the [`edit interfaces ether-options ethernet-switch-profile`] hierarchy level, a temporarily disabled port will come up again after the specified time interval.

Default

Default behavior for this configuration is soft shutdown.

Options

- `hard-shutdown`—When the storm control level threshold is exceeded, restricts all packets, port will not be operational.
- `soft-shutdown`—When the storm control level threshold is exceeded, restricts data packets while control packets are still received.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D40.

RELATED DOCUMENTATION

Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches

Understanding Storm Control

interface (SLAAC Snooping)

IN THIS SECTION

- Syntax | 9052
- Hierarchy Level | 9053
- Description | 9053
- Options | 9053
- Required Privilege Level | 9054
- Release Information | 9054

Syntax

```
interface (interface-name | all) {
    auto-dad {
        retries retry-count;
        retrans-interval seconds;
    }
    mark-interface trusted;
    maximum-allowed-contentions  {
        count integer;
        duration seconds;
```

```
    }
}
```

Hierarchy Level

```
[edit forwarding-options access-security slaac-snooping]
```

Description

Configure interface-level parameters for IPv6 stateless address auto-configuration (SLAAC) snooping. SLAAC enables an IPv6 client to generate its own local and global addresses using a combination of locally-available information and information advertised by routers through Neighbor Discovery Protocol (NDP). NDP messages are unsecured, which makes SLAAC susceptible to attacks that involve the spoofing (or forging) of link-layer addresses. IPv6 clients using SLAAC for dynamic address assignment are validated against the SLAAC snooping binding table before being allowed access to the network.

The remaining statements are explained separately. See CLI Explorer.

Options

(*interface-name* | all) Configure SLAAC snooping parameters on the specified interface or on all interfaces.

mark-interface trusted Configure the interface as trusted. The binding entry for the trusted interface is added to the SLAAC snooping table using the same process as for untrusted interfaces. When a DAD request is received on a trusted port with an IP/MAC entry that already exists on an untrusted port, SLAAC snooping sends a unicast DAD towards the untrusted port to see whether the host is live. If the host responds with an NA message on the untrusted port, the lease time is renewed for the existing binding entry. If there is no response (NA) on the untrusted port, the corresponding binding entry is deleted.

If the entry for the untrusted port is deleted, the binding for the trusted port is not created immediately. When the trusted port starts to send data traffic, it will send an NS message. At that time, SLAAC snooping adds the new binding on the trusted port.

> **NOTE**: Maximum number of DAD contentions is not applicable to trusted interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

IPv6 Stateless Address Auto-configuration (SLAAC) Snooping

interface (SNMP)

IN THIS SECTION

- Syntax | 9054
- Hierarchy Level | 9055
- Description | 9055
- Default | 9055
- Options | 9055
- Required Privilege Level | 9055
- Release Information | 9055

Syntax

```
interface [ interface-names ];
```

Hierarchy Level

```
[edit snmp]
```

Description

Configure the interfaces on which SNMP requests can be accepted.

Default

If you omit this statement, SNMP requests entering the router or switch through any interface are accepted.

Options

interface-names—Names of one or more logical interfaces.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Interfaces on Which SNMP Requests Can Be Accepted

interface (SNMP RMON History)

IN THIS SECTION

- Syntax | 9056
- Hierarchy Level | 9056
- Description | 9056
- Options | 9056
- Required Privilege Level | 9057
- Release Information | 9057

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit snmp rmon history history-index]
```

Description

Specify the interface to be monitored in the specified RMON history entry.

Only one interface can be specified for a particular RMON history index. There is a one-to-one relationship between the interface and the history index. The interface must be specified in order for the RMON history to be created.

Options

interface-name—Specify the interface to be monitored within the specified entry of the RMON history of Ethernet statistics.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- *Understanding RMON*
- *RMON MIB Event, Alarm, Log, and History Control Tables*
- *Configuring RMON History Sampling*
- *Monitoring RMON MIB Tables*
- Junos OS Network Management Configuration Guide

interface (Spanning Tree)

IN THIS SECTION

- Syntax | 9058
- Syntax | 9058
- Hierarchy Level | 9058
- Description | 9059
- Options | 9059
- Required Privilege Level | 9059
- Release Information | 9059

Syntax

```
interface interface-name {
    bpdu-timeout-action {
        alarm;
        block;
    }
    cost cost;
    edge;
    mode (p2p | shared);
    no-root-port;
    priority interface-priority;
}
```

Syntax

```
interface interface-name {
    arp-on-stp;
     bpdu-timeout-action
         block;
        log;
     cost cost;
     disable;
     edge;
     mode mode;
     no-root-port;
     priority priority;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp | vstp)],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp | vstp)],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id],
```

```
[edit protocols (mstp | rstp | vstp)],
[edit protocols vstp vlan vlan-id],
```

```
[edit protocols (mstp | rstp | vstp)],
[edit protocols vstp vlan vlan-id],
[edit protocols vstp vlan-group group group-name vlan (vlan-id | vlan-range |open-set-of-values)
```

```
[edit routing-instances routing-instance-name protocols (mstp | rstp | vstp)],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id]
```

Description

Configures the interface to participate in the RSTP, MSTP, or VSTP instance.

The **edge**, **mode**, and **no-root-port** options are not available at the [**edit protocols mstp msti** ***msti-id***] hierarchy level.

Options

interface-name—Name of a Gigabit Ethernet or 10-Gigabit Ethernet interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

Statement updated in Junos OS Release 15.1 for EX Series switches to support configuration of spanning tree parameters globally on all interfaces.

NOTE: You cannot disable spanning tree parameters globally on all interfaces.

RELATED DOCUMENTATION

Understanding Spanning-Tree Instance Interfaces

Understanding RSTP

Understanding MSTP

Understanding VSTP

Configuring RSTP on EX Series Switches (CLI Procedure)

Configuring MSTP on Switches

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

interface-specific

IN THIS SECTION

- Syntax | **9061**
- Hierarchy Level | **9061**
- Description | **9061**
- Required Privilege Level | **9061**
- Release Information | **9061**

Syntax

```
interface-specific;
```

Hierarchy Level

```
[edit firewall family family-name filter                filter-name]
```

Description

Configure separate counters for each interface to which a filter is applied.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Firewall Filter Match Conditions and Actions (QFX5100, QFX5110, QFX5120, QFX5200, EX4600, EX4650)

Configuring Firewall Filters

Overview of Firewall Filters (QFX Series)

interface-specific (Dynamic Firewalls)

IN THIS SECTION

Syntax | 9062

- Hierarchy Level | 9062
- Description | 9062
- Required Privilege Level | 9062
- Release Information | 9062

Syntax

```
interface-specific;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family fast-update-filter filter-name]
[edit dynamic-profiles profile-name firewall family family-name filter filter-name]
```

Description

Configure interface-specific names for firewall counters that are based on fast update filters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Fast Update Filters

Interface-Specific Firewall Filter Instances Overview

interface-specific (Firewall Filters)

IN THIS SECTION

- Syntax | 9063
- Hierarchy Level | 9063
- Description | 9063
- Required Privilege Level | 9063
- Release Information | 9064

Syntax

```
interface-specific;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name filter filter-name],
[edit No Link Title family family-name filter filter-name],
[edit logical-systems logical-system-name firewall family family-name filter filter-name]
```

Description

Configure interface-specific names for firewall counters.

NOTE: A firewall filter cannot be both interface-specific and interface-shared.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Firewall Filters and Policers for VPLS

Interface-Specific Firewall Filter Instances Overview

interface-state

IN THIS SECTION

- Syntax | 9064
- Hierarchy Level | 9065
- Description | 9065
- Options | 9065
- Required Privilege Level | 9066
- Release Information | 9066

Syntax

```
interface-state {
    hold-time {
      down seconds;
      up seconds;
    }
    (local | local-interface local-interface | local-remote | remote);
    local-count local-count;
    vtep-count vtep-count;
}
```

Hierarchy Level

```
[edit interfaces name unit num]
```

Description

Consider other local or remote interface state changes when evaluating the state (up or down) of the named interface.

Options

hold-time (down | up) ***seconds***
Time to wait in seconds after other interface or remote reachability states change before considering those changes to evaluate the state of the named interface.

local
Determine the interface state based on the states of the associated local Layer 2 (L2) interface(s).

In the context of an EVPN instance (EVI) where the interface is associated with a bridge domain or VLAN, consider the interface state to be **up** if at least one local L2 interface is **up**.

local-interface ***local-interface***
Determine the interface state based on the state of *local-interface*, which is a loopback aggregated Ethernet (AE) port. Specify *local-interface* by its logical interface name (for example, ae0.100).

You use this option with EVPN-VXLAN support on devices that use a loopback port for routing in and out of VXLAN tunnels (see *Using a RIOT Loopback Port to Route Traffic in an EVPN-VXLAN Network*).

local-remote
Determine the interface state based on:

- The states of the associated local L2 interface(s)
- Remote reachability status, such as the bridge domain, VLAN, or EVI network isolation status and whether the device has routes to the remote VTEPs.

Consider the interface state to be **up** if either of these conditions is true in the context of an EVI where the interface is associated with a bridge domain or VLAN:

- At least one local L2 interface is in **up** state.
- The bridge domain, VLAN, or EVI is not in a network isolation state.

Consider the interface state to be **down** if both of these conditions are true in the context of an EVI where the interface is associated with a bridge domain or VLAN:

- No local L2 interfaces are in **up** state.
- The bridge domain, VLAN or EVI is in a network isolation state.

remote Determine the interface status based on remote reachability status, including the bridge domain or EVPN instance network isolation state.

In the context of an EVPN instance where the interface is associated with a bridge domain, consider the interface state to be:

- **Up** if the bridge domain or EVPN instance is not in a network isolation state.
- **Down** if the bridge domain or EVPN instance is in a network isolation state.

local-count *local-count* Minimum number of local L2 links in **up** state to evaluate the interface state as **up**.

vtep-count *vtep-count* Minimum number of VTEP interfaces in **up** state to evaluate the interface status as **up**.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

network-isolation | **11312**

network-isolation-profile | **11314**

interface (Static MAC Bypass)

IN THIS SECTION

- Syntax | 9067
- Hierarchy Level | 9067
- Description | 9067
- Options | 9067
- Required Privilege Level | 9067
- Release Information | 9068

Syntax

```
interface [interface-names];
```

Hierarchy Level

```
[edit protocols authentication-access-control]
```

Description

Configure interfaces on which the specified MAC addresses are allowed to bypass RADIUS authentication and allowed to connect to the LAN without authentication.

Options

interface-names—List of interfaces.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

interface (Static Subscriber Group)

IN THIS SECTION

- Syntax | 9068
- Hierarchy Level | 9068
- Description | 9069
- Options | 9069
- Required Privilege Level | 9069
- Release Information | 9069

Syntax

```
interface interface-name <exclude> <upto upto-interface-name>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name system services static-subscribers group group-name],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name],
[edit system services static-subscribers group group-name]
```

Description

Specify one or more interfaces, or a range of interfaces, that are within a specified group on which static subscribers are created. You can repeat the `interface` *`interface-name`* statement to specify multiple interfaces within a group. You must configure each interface in only one group.

> **NOTE**: The logical system and routing instance in which the static interfaces are configured must match the logical system and routing instance where the group is configured.

Options

`exclude`—(Optional) Exclude an interface or a range of interfaces from the group.

`interface-name`—Name of the interface on which static subscribers are created. If you do not specify a unit number for the interface, then .0 is assumed. For example, `ge-0/1/0` is interpreted as `ge-0/1/0.0`.

`upto-interface-name`—(Optional) The upper end of the range of interfaces; the lower end of the range is the *interface-name* entry. The interface device name of *upto-interface-name* must be the same as the device name of *interface-name*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support for IPv6 and IPv4 demux static interfaces introduced in Junos OS Release 11.2.

Support for pseudowire interfaces over logical tunnels introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Creating a Static Subscriber Group

interface (Static Subscriber Username)

IN THIS SECTION

- Syntax | 9070
- Hierarchy Level | 9070
- Description | 9071
- Required Privilege Level | 9071
- Release Information | 9071

Syntax

```
interface;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication username-include],
[edit logical-systems logical-system-name system services static-subscribers  authentication
username-include],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers authentication
username-include],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],
[edit system services static-subscribers authentication username-include],
[edit system services static-subscribers group group-name authentication username-include]
```

Description

Specify that a modified version of the interface name is included as part of the username created for all static subscribers or for the static subscribers in a specified group. The group version of the statement takes precedence over the global version. The username is also sent to RADIUS in the Access-Request message. The interface name is modified by replacing the "/" character with the "-" character. For example, ge-0/1/2.50 is converted to ge-0-1-2.50.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Configuring the Static Subscriber Global Username

Configuring the Static Subscriber Group Username

interface-status-send-rdi

IN THIS SECTION

- Syntax | 9072
- Hierarchy Level | 9072
- Description | 9072
- Required Privilege Level | 9072
- Release Information | 9072

Syntax

```
interface-status-send-rdi interface-status-send-rdi;
```

Hierarchy Level

```
[edit logical-systems name protocols oam ethernet connectivity-fault-
management maintenance-domain name maintenance-
association name continuity-check],
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain name maintenance-association
name continuity-check]
```

Description

Configure CFM to propagate the status of the provider edge device via the remote defect indication (RDI) bit in the CC messages when the interface is down. When the status of the EVPN provider edge device is standby, the EVPN VPWS service is notified and it sets the interface status to CCC-down. When the interface status is CCC-down, it indicates that the provider edge service is down. When you enable CFM monitoring, CFM propagates the status of the provider edge device via the remote defect indication (RDI) bit in the CC messages. Thus, the customer edge device is aware that the provider edge device is down.

Usually, when the interface goes down, CFM propagates the status of the provider edge device via the interface status TLV. If the customer edge device does not support the interface status TLV, you can use the RDI bit to propagate the status of the provider edge device.

Required Privilege Level

routing—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

interface-status-tlv | 9073

Configuring Port Status TLV and Interface Status TLV

Understanding CFM Monitoring between CE and PE Devices

interface-status-tlv

IN THIS SECTION

- Syntax | 9073
- Hierarchy Level | 9073
- Description | 9073
- Options | 9074
- Required Privilege Level | 9074
- Release Information | 9074

Syntax

```
interface-status-tlv [ down lower-layer-down ];
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile profile-name event]
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name continuity-check]
```

Description

Defines an `action-profile` consisting of various events and the action. Based on values of `interface-status-tlv` in the received CCM packets, specific action such as *interface-down* can be taken using "action-profile" on page 1778 options.

Options

`down`	When the incoming CCM packet contains interface status TLV with value down, the action will be triggered for this action-profile.
`lower-layer-down`	When the incoming CCM packet contains interface status TLV with value lower-layer-down, the action will be triggered for this action-profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Port Status TLV and Interface Status TLV

interface (Storm Control)

IN THIS SECTION

- Syntax | 9075
- Hierarchy Level | 9075
- Description | 9075
- Default | 9075
- Options | 9075
- Required Privilege Level | 9076
- Release Information | 9076

Syntax

```
interface (all | interface-name) {
    bandwidth bandwidth;
    level level;
    multicast;
    no-broadcast;
    no-multicast;
    no-registered-multicast;
    no-unknown-unicast;
    no-unregistered-multicast;
}
```

Hierarchy Level

```
[edit ethernet-switching-options storm-control]
```

Description

Configure storm control on all interfaces or on the specified interface.

Default

- On EX2200, EX3200, EX3300, EX4200, and EX6200 switches—Storm control does not apply by default to multicast traffic. The factory default configuration enables storm control for broadcast and unknown unicast traffic on all switch interfaces, with the storm control level set to 80 percent of the combined broadcast and unknown unicast streams.
- On EX4500 and EX8200 switches—Storm control applies to broadcast, multicast, and unknown unicast traffic. The factory default configuration enables storm control on all switch interfaces, with the storm control level set to 80 percent of the combined broadcast, multicast, and unknown unicast streams.

Options

`all`—All interfaces. The storm control settings configured with the `all` option affect only those interfaces that have not been individually configured for storm control.

interface-name—Name of an interface. The storm control settings configured with the *interface-name* option override any settings configured with the `all` option.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Configuring Autorecovery for Port Security Events

Enabling and Disabling Storm Control (non-ELS)

interface-switch

IN THIS SECTION

- Syntax | 9077
- Hierarchy Level | 9077
- Description | 9077
- Options | 9077
- Required Privilege Level | 9077
- Release Information | 9077

Syntax

```
interface-switch connection-name {
    interface interface-name.unit-number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols connections],
[edit protocols connections]
```

Description

Configure Layer 2 switching cross-connects. The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first.

For Layer 2 switching cross-connects to work, you must also configure MPLS.

Options

connection-name—Connection name (up to 128 characters in Junos 12.3 and later).

interface *interface-name.unit-number*—Interface name. Include the logical portion of the name, which corresponds to the logical unit number.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring the CCC Connection for Layer 2 Switching Cross-Connects*

Defining the Connection for Switching Cross-Connects

MPLS Applications User Guide

interface (Switching Options)

IN THIS SECTION

- Syntax | 9078
- Hierarchy Level | 9078
- Description | 9078
- Options | 9079
- Required Privilege Level | 9079
- Release Information | 9079

Syntax

```
interface interface-name {
    encapsulation-type;
    ignore-encapsulation-mismatch;
    pseudowire-status-tlv;
    static-mac mac-address {
        vlan-id vlan-id;
    }
}
```

Hierarchy Level

```
[edit vlans vlans-name switch-options]
```

Description

Specify the logical interfaces to include in the VLAN.

Options

- *interface-name*—Name of a logical interface.
- *encapsulation-type*—Encapsulation type for VPN.
- *ignore-encapsulation-mismatch*—Allow different encapsulation types on local and remote devices.
- *pseudowire-status-tlv*—Send pseudowire status.
- *mac-address*—Static MAC address assigned to the logical interface.
- *vlan-id*—VLAN identifier.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.5.

interface (System Services DHCP)

IN THIS SECTION

- Syntax | 9080
- Hierarchy Level | 9080
- Description | 9080
- Options | 9080
- Required Privilege Level | 9080
- Release Information | 9081

Syntax

```
interface interface-name {
    exclude;
    overrides {
        interface-client-limit number;
    }
    trace;
    upto upto-interface-name;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6 group group-name]
```

Description

Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP local server is enabled. You can repeat the interface interface-name statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group.

Options

- *interface-name*—Name of the interface.
- `trace`—Enable tracing of the interface specified by the *interface-name* argument.
- `upto` *upto-interface-name*—The upper end of the range of interfaces; the lower end of the range is the interface-name entry. The interface device name of the *upto-interface-name* must be the same as the device name of the *interface-name*.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

DHCP Server Configuration

DHCP Server Options

interface-tag (DHCP Local Server)

IN THIS SECTION

- Syntax | 9081
- Hierarchy Level | 9082
- Description | 9082
- Options | 9082
- Required Privilege Level | 9082
- Release Information | 9082

Syntax

```
interface-tag name {
   access-profile;
   dynamic-profile (DHCP Local Server);
   overrides (DHCP Local Server);
   service-profile (DHCP Local Server);
   short-cycle-protection (DHCP Local Server and Relay Agent);;
   }
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server group],
[edit logical-systems name system services dhcp-local-server group],
[edit logical-systems name tenants name routing-instances name system services dhcp-local-server
group],
[edit routing-instances name system services dhcp-local-server group],
[edit system services dhcp-local-server group],
[edit tenants name routing-instances name system services dhcp-local-server group]
```

Description

Configure the interface tag name to map auto-configured dynamic VLANs to DHCP groups. Specify an interface tag with a dynamic. You can then use the dynamic profile to map to a subset of incoming VLAN IDs. The maximum number of interface-tags supported is 32. This is also the maximum number of supported auto-configurable VLAN dynamic profiles.

Options

name Interface tag name. The length of the interface tag name cannot exceed 64 characters. The only supported special characters are ^_-*$.

The remaining statements are explained separately. Click the linked statement in the Syntax section for details or search for the statement in CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

group (DHCP Local Server) | 7814

interface-tag (DHCP Relay Server)

IN THIS SECTION

- Syntax | 9083
- Hierarchy Level | 9083
- Description | 9084
- Options | 9084
- Required Privilege Level | 9084
- Release Information | 9085

Syntax

```
interface-tag name {
   access-profile;
   dynamic-profile (DHCP Relay Agent);
   overrides (DHCP Relay Agent);
   service-profile (DHCP Relay Agent);
   short-cycle-protection (DHCP Local Server and Relay Agent);
  }
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group],
[edit forwarding-options dhcp-relay group],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group],
[edit logical-systems name forwarding-options dhcp-relay group],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
```

```
relay group],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay
group],
[edit logical-systems name tenants name routing-instances name bridge-domains name forwarding-
options dhcp-relay group],
[edit logical-systems name tenants name routing-instances name forwarding-options dhcp-relay
group],
[edit logical-systems name tenants name routing-instances name vlans name forwarding-options
dhcp-relay group],
[edit logical-systems name vlans name forwarding-options dhcp-relay group],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group],
[edit routing-instances name forwarding-options dhcp-relay group],
[edit routing-instances name vlans name forwarding-options dhcp-relay group],
[edit tenants name routing-instances name bridge-domains name forwarding-options dhcp-relay
group],
[edit tenants name routing-instances name forwarding-options dhcp-relay group],
[edit tenants name routing-instances name vlans name forwarding-options dhcp-relay group],
[edit vlans name forwarding-options dhcp-relay group]
```

Description

Configure the interface tag name to map auto-configured dynamic VLANs to DHCP groups. Specify an interface tag with a dynamic. You can then use the dynamic profile to map to a subset of incoming VLAN IDs. The maximum number of interface-tags supported is 32. This is also the maximum number of supported auto-configurable VLAN dynamic profiles.

Options

name Interface tag name. The length of the interface tag name cannot exceed 64 characters. The only supported special characters are ^_-*$.

The remaining statements are explained separately. Click the linked statement in the Syntax section for details or search for the statement in CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

group (DHCP Relay Agent) | 7820

interface-traceoptions (DHCP)

IN THIS SECTION

- Syntax | 9085
- Hierarchy Level | 9085
- Description | 9086
- Options | 9086
- Required Privilege Level | 9087
- Release Information | 9087

Syntax

```
interface-traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
    level  (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system processes dhcp-service]
```

Description

Configure extended DHCP tracing operations that can be enabled on a specific interface or group of interfaces.

Replaces deprecated `interface-traceoptions` statements at the `[edit forwarding-options dhcp-relay]` and `[edit system services dhcp-local-server]` hierarchy levels.

To enable the tracing operation on the specific interfaces, you use the `interface` *`interface-name`* `trace` statement.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements

- `all`—Trace all events
- `packet`—Trace packet and option decoding operations
- `state`—Trace changes in state

`level`—Level of tracing to perform; also known as severity level. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- `all`—Match messages of all levels.
- `error`—Match error messages.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure `verbose`, messages at all higher levels are traced. Therefore, the result is the same as when you configure `all`.
- `warning`—Match warning messages.

- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Tracing Extended DHCP Operations for Specific Interfaces

interface-traceoptions (System Services DHCP)

IN THIS SECTION

Syntax | 9088
Hierarchy Level | 9088

- Description | 9088
- Options | 9089
- Required Privilege Level | 9090
- Release Information | 9090

Syntax

```
interface-traceoptions {
    file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name system services dhcp-local-server],
[edit system services dhcp-local-server]
```

Description

Configure extended DHCP local server tracing operations that can be enabled on a specific interface or group of interfaces. You use the `interface` *`interface-name`* `trace` statement at the `[edit system services group` *`group-name`*`]` hierarchy level to enable the tracing operation on the specific interfaces.

Options

file-name—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (“ ”). All files are placed in a file named `jdhcpd` in the directory `/var/log`. If you include the `file` statement, you must specify a filename.

`files` *number*—(Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all events
- `dhcpv6-packet`—Trace DHCPv6 packet decoding operations.
- `dhcpv6-packet-option`—Trace DHCPv6 option decoding operations.
- `dhcpv6-state`—Trace changes in state for DHCPv6 operations.
- `packet`—Trace packet decoding operations
- `packet-option`—Trace DHCP option decoding operations
- `state`—Trace changes in state

`match` *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** `xk` to specify KB, `xm` to specify MB, or `xg` to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

DHCP Server Configuration

DHCP Server Options

interface-transmit-statistics

IN THIS SECTION

- Syntax | 9090
- Hierarchy Level | 9090
- Description | 9091
- Required Privilege Level | 9091
- Release Information | 9091

Syntax

```
interface-transmit-statistics;
```

Hierarchy Level

```
[edit interfaces {interface-name}]
```

Description

Configure the interface to report the transmitted load statistics. If this statement is not included in the configuration, the interface statistics show the offered load on the interface, and not the actual transmitted load.

Aggregated Ethernet interfaces do not support reporting of the transmitted load statistics. You cannot configure aggregated Ethernet interfaces to capture and report the actual transmitted load statistics.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4R3.

RELATED DOCUMENTATION

Improvements to Interface Transmit Statistics Reporting

show interfaces | **20227**

interface-type (Protocols OSPF)

IN THIS SECTION

- Syntax | **9092**
- Hierarchy Level | **9093**
- Description | **9093**
- Default | **9093**
- Options | **9094**
- Required Privilege Level | **9094**
- Release Information | **9094**

Syntax

```
interface-type (nbma | p2mp | p2mp-over-lan | p2p) {
    ipsec-sa;
    ldp-synchronization {
        (disable | enable);
        hold-time;
    }
    metric;
    mtu;
    neighbor address {
        eligible;
    }
    no-advertise-adjacency-segment;
    no-eligible-backup;
    no-eligible-remote-backup;
    no-interface-state-traps;
    no-neighbor-down-notification;
    node-link-protection;
    passive {
        traffic-engineering {
            remote-node-id;
            remote-node-router-id;
        }
    }
    poll-interval;
    priority;
    retransmit-interval;
    secondary;
    te-metric;
    topology (default | ipv4-multicast | name);
    transit-delay;
    bandwidth-based-metrics;
    bfd-liveness-detection;
    dead-interval;
    demand-circuit;
    disable;
    dynamic-neighbors;
    flood-reduction;
    hello-interval;
    link-protection;
```

```
    own-router-lsa;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface
interface-name],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-multicast | ipv4-unicast |
ipv6-multicast) area area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) area area-id interface interface-
name],
[edit protocols (ospf | ospf3) area area-id interface interface-name],
[edit protocols ospf3 realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) area area-id
interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface
interface-name],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-multicast | ipv4-
unicast | ipv6-multicast) area area-id interface interface-name]
```

Description

Specify the type of interface.

By default, the software chooses the correct interface type based on the type of physical interface. Therefore, you should never have to set the interface type. The exception to this is for NBMA interfaces, which default to an interface type of point-to-multipoint. To have these interfaces explicitly run in Nonbroadcast multiaccess (NBMA) mode, configure the nbma interface type, using the IP address of the local ATM interface.

In Junos OS Release 9.3 and later, a point-to-point interface can be an Ethernet interface without a subnet.

Default

The software chooses the correct interface type based on the type of physical interface.

Options

nbma (OSPFv2 only)	Nonbroadcast multiaccess (NBMA) interface
p2mp	Point-to-multipoint interface
p2mp-over-lan	Point-to-multipoint over LAN mode
p2p	Point-to-point interface

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for OSPFv3 for interface type `p2p` only introduced in Junos OS Release 9.4. You cannot configure other interface types for OSPFv3.

Support for OSPFv3 for interface type `p2mp` is introduced in Junos OS Release 18.1R1.

Support for OSPFv3 for interface type `p2p` only introduced in Junos OS Release 9.4 for EX Series switches.

RELATED DOCUMENTATION

About OSPF Interfaces

Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network

interface-type (Protocols RIP)

IN THIS SECTION

- Syntax | 9095
- Hierarchy Level | 9095
- Description | 9095
- Options | 9095
- Required Privilege Level | 9095
- Release Information | 9096

Syntax

```
interface-type p2mp;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit protocols rip group group-name neighbor neighbor-name]
```

Description

Configure the type of interface in a RIP network.

This statement enables a RIP device to have single or multiple peers through an interface.

Options

p2mp Configure an interface in a RIP network as a point-to-multipoint interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring Point-to-Multipoint RIP Networks

interface (Unknown Unicast Forwarding)

IN THIS SECTION

- Syntax | 9096
- Hierarchy Level | 9096
- Description | 9097
- Required Privilege Level | 9097
- Release Information | 9097

Syntax

```
interface  interface-name;
```

Hierarchy Level

- For platforms with ELS:

```
[edit switch-options unknown-unicast-forwarding vlan vlan-name]
```

- For platforms without ELS:

```
[edit ethernet-switching-options unknown-unicast-forwarding vlan vlan-name]
```

Description

Specify the interface to which unknown unicast packets will be forwarded.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Hierarchy level `[edit switch-options]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

show vlans

show ethernet-switching table | 19700

Understanding and Preventing Unknown Unicast Forwarding

interface (Virtual Tunnel in Routing Instances)

IN THIS SECTION

- Syntax | 9098
- Hierarchy Level | 9098
- Description | 9098
- Options | 9098

Required Privilege Level | 9099
Release Information | 9099

Syntax

```
interface vt-fpc/pic/port.unit-number {
    multicast;
    primary;
    unicast;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure a virtual tunnel (VT) interface.

VT interfaces are needed for multicast traffic on routing devices that function as combined provider edge (PE) and provider core (P) routers to optimize bandwidth usage on core links. VT interfaces prevent traffic replication when a P router also acts as a PE router (an exit point for multicast traffic).

In an MBGP MVPN extranet, if there is more than one VRF routing instance on a PE router that has receivers interested in receiving multicast traffic from the same source, VT interfaces must be configured on all instances.

Starting in Junos OS Release 12.3, you can configure multiple VT interfaces in each routing instance. This provides redundancy. A VT interface can be used in only one routing instance.

Options

vt-fpc/pic/port.unit-number—Name of the VT interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

> *Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs*
>
> Example: Configuring MBGP MVPN Extranets

interface (VoIP)

IN THIS SECTION

- Syntax | 9099
- Hierarchy Level | 9100
- Description | 9100
- Options | 9100
- Required Privilege Level | 9101
- Release Information | 9101

Syntax

```
interface (all | [interface-name] | access-ports) {
    forwarding-class forwarding-class;
    vlan (vlan-id  | vlan-name | untagged));
}
```

Hierarchy Level

- For platforms with ELS:

```
[edit switch-options voip]
```

- For platforms without ELS:

```
[edit ethernet-switching-options                voip],
```

Description

(Required) Enable voice over IP (VoIP) on interfaces.

Options

all Enable VoIP on all interfaces.

interface-name Enable VoIP on a specific interface.

access-ports (Switches without ELS only) Enable VoIP on all access ports.

forwarding-class ***forwarding-class*** (Optional) For EX Series switches, configure the forwarding class used to handle packets on the VoIP interface.

> **NOTE**: The `forwarding-class` statement at the [`edit ethernet-switching-options voip interface` *`interface-name`*] hierarchy level is used only by LLDP-MED for advertising the capabilities of VoIP phones. It is not used to classify VoIP traffic.

- **Values:** Specify one of the following:
  - assured-forwarding— Assured forwarding (AF) provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with three drop probabilities: low, medium, and high.

- best-effort—Provides no service profile. For the best effort forwarding class, loss priority is typically not carried in a class-of-service (CoS) value, and random early detection (RED) drop profiles are more aggressive.
- expedited-forwading—Provides a low loss, low latency, low jitter, assured bandwidth, end-to-end service.
- network-control—Provides a typically high priority because it supports protocol control.

- **Default:** Disabled

vlan (*vlan-id* | *vlan-name* | untagged)

(Required) Specify the VLAN name or VLAN tag identifier associated with the VLAN to be sent from the authenticating server to the IP phone or allow untagged VLAN traffic.

- **Syntax:** Specify one of the following:
  - *vlan-name*—Name of a VLAN.
  - *vlan-id*—The VLAN tag identifier.
  - *untagged*—Allow untagged VLAN traffic.
- **Range:** vlan-id range is 1 through 4094. Tags 0 and 4095 are reserved by the Junos OS; do not configure them.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Hierarchy level `[edit switch-options]` introduced in Junos OS Release 13.2X50-D10. (See *Using the Enhanced Layer 2 Software CLI* for information about ELS.)

RELATED DOCUMENTATION

Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch

Example: Configuring VoIP on an EX Series Switch Without Including 802.1X Authentication

Example: Configuring VoIP on an EX Series Switch Without Including LLDP-MED Support

interface (VPLS Mesh-Group)

IN THIS SECTION

- Syntax | 9102
- Hierarchy Level | 9102
- Description | 9102
- Required Privilege Level | 9102
- Release Information | 9103

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]
```

Description

Specify one or more interfaces belonging to the virtual private LAN service (VPLS) flood mesh-group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.2.

RELATED DOCUMENTATION

Configuring VPLS Routing Instances

mesh-group (Protocols VPLS) | 10641

interface (VPLS Multihoming for FEC 129)

IN THIS SECTION

- Syntax | 9103
- Hierarchy Level | 9103
- Description | 9104
- Options | 9104
- Required Privilege Level | 9104
- Release Information | 9104

Syntax

```
interface interface-name {
    preference preference-value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-
homing site site-name],
[edit routing-instances instance-name protocols vpls multi-homing site site-name]
```

Description

Configure the interface that connects this site to the VPN.

The remaining statement is explained separately. See CLI Explorer.

Options

interface-name—Name of the interface (for example, `ge-0/1/0.1`).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring VPLS Multihoming (FEC 129)

interface (VPLS Routing Instances)

IN THIS SECTION

- Syntax | 9105
- Hierarchy Level | 9105
- Description | 9105
- Options | 9105
- Required Privilege Level | 9105
- Release Information | 9105

Syntax

```
interface interface-name {
    mac-pinning;
    interface-mac-limit limit;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
site site-name],
[edit routing-instances routing-instance-name protocols vpls],
[edit routing-instances routing-instance-name protocols vpls site site-name]
```

Description

Configure the interface for a pseudowire to the VPLS customer site. To complete the configuration of interfaces for a VPLS routing instance, you must also configure the interfaces specified for a VPLS site at the `[edit routing-instances routing-instance-name]` hierarchy level as described in *Configuring Routing Instances on PE Routers in VPNs*.

Options

interface-name—Specify the name of the interface used by the VPLS site.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the VPLS Site Interfaces

Configuring Routing Instances on PE Routers in VPNs

interface (Routing Instances) | 9012

interface-client-limit (DHCP Local Server)

IN THIS SECTION

- Syntax | 9106
- Hierarchy Level | 9106
- Description | 9107
- Default | 9107
- Options | 9108
- Required Privilege Level | 9108
- Release Information | 9108

Syntax

```
interface-client-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name overrides],
[edit logical-systems logical-system-name system services dhcp-local-server overrides],
```

```
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 overrides],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name overrides],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
overrides],
[edit system services dhcp-local-server overrides],
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group interface interface-name group-name
overrides],
[edit system services dhcp-local-server group group-name overrides],
[edit system services dhcp-local-server group group-name interface interface-name overrides]
```

Description

Set the maximum number of DHCP subscribers or DHCP clients per interface allowed for a specific group or for all groups. A group specification takes precedence over a global specification for the members of that group.

> **NOTE**: Do not use this statement for dual-stack subscribers. Instead, use the "dual-stack-interface-client-limit" on page 6058 statement for dual-stack subscribers.

Default

No limit

Options

number—Maximum number of clients allowed.

- **Range:** 1 through 500,000

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Specifying the Maximum Number of DHCP Clients Per Interface

Overriding the Default DHCP Local Server Configuration Settings

interface-client-limit (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 9109
- Hierarchy Level | 9109
- Description | 9109
- Default | 9110
- Options | 9110
- Required Privilege Level | 9110
- Release Information | 9110

Syntax

```
interface-client-limit number;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Set the maximum number of DHCP (or DHCPv6) subscribers or clients per interface allowed for a specific group or for all groups. A group specification takes precedence over a global specification for the members of that group. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

> **NOTE**: Do not use this statement for dual-stack subscribers. Instead, use the "dual-stack-interface-client-limit" on page 6058 statement for dual-stack subscribers.

Default

`No limit`

Options

number—Maximum number of clients allowed.

- **Range:** 1 through 500,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

dhcp-relay | **5620**

Extended DHCP Relay Agent Overview

Configuring Group-Specific DHCP Relay Options

Overriding the Default DHCP Relay Configuration Settings

interfaces (VNF)

IN THIS SECTION

- Syntax | 9111
- Hierarchy Level | 9112
- Description | 9112
- Options | 9112
- Required Privilege Level | 9112
- Release Information | 9112

Syntax

```
interfaces interface-name {
    pci-address pci-address;
        mac-address mac-address;
        mtu size;
        mapping {
        hsxe0 {
            virtual-function {
                vlan-id vlan-id;
            }
        }
        hsxe1 {
            virtual-function {
                vlan-id vlan-id;
            }
        }
        vlan {
            members vlan-name;
                mode [access | trunk];
            native-vlan-id vlan-id;
        }
    }
}
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Configure Virtual Network Functions (VNF) interfaces on platforms running disaggregated Junos OS.

Options

interface-name	Name of the VNF interfaces.
mac-address	MAC address of the VNF interfaces.
mtu	Maximum transfer unit (MTU) size of packets in bytes.
pci-address	Target PCI address of the VNF interfaces.
vlan-id	SR-IOV virtual function to use to attach a VNF to a physical interface.
vlan *members*	Membership for this interface.
native-vlan-id	Virtual LAN identifier for untagged frames. For example, 1...4095
vlan-name	Name of the VLAN members.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D50.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

interfaces

IN THIS SECTION

- Syntax (QFX Series) | 9113
- Syntax (EX Series, MX Series and T Series) | 9113
- QFX Series | 9113
- EX Series, MX Series and T Series | 9114
- Description | 9114
- Default | 9114
- Options | 9114
- Required Privilege Level | 9114
- Release Information | 9114

Syntax (QFX Series)

```
interfaces interface-name {
                        no-mac-learning;
}
```

Syntax (EX Series, MX Series and T Series)

```
interfaces { ... }
```

QFX Series

```
[edit ethernet-switching-options]
```

EX Series, MX Series and T Series

```
[edit]
```

Description

Configure settings for interfaces that have been assigned to family `ethernet-switching`.

Default

The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Options

interface-name —Name of an interface that is configured for `family ethernet-switching`.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Physical Interface Properties Overview
- Configuring Aggregated Ethernet Link Protection

interfaces (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 9115
- Hierarchy Level | 9116
- Description | 9116
- Options | 9116
- Required Privilege Level | 9116
- Release Information | 9116

Syntax

```
interfaces interface-name {
    load-balancing-options {
        hash-keys {
            egress-key (source-ip | destination-ip);
            ingress-key  (source-ip | destination-ip);
        }
        high-availability-options {
            (many-to-one | one-to-one)  {
                preferred-backup preferred-backup;
            }
        }
        member-failure-options {
            drop-member-traffic {
                rejoin-timeout rejoin-timeout;
            }
            redistribute-all-traffic {
                enable-rejoin;
            }
        }
        member-interface interface-name;
    }
    redundancy-options {
        primary mams-a/b/0;
        secondary mams-a/b/0;
```

```
    }
    unit interface-unit-number {
        family family;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure the aggregated Multiservices (AMS) interface. The AMS interface provides the infrastructure for load balancing and high availability (HA).

Options

interface-name Name of a valid aggregated multiservices interface (ams)—for example, ams0 or ams1.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Load Balancing on AMS Infrastructure

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

interfaces (Analytics)

IN THIS SECTION

- Syntax | 9117
- Hierarchy Level | 9118
- Description | 9118
- Options | 9118
- Required Privilege Level | 9119
- Release Information | 9119

Syntax

```
interfaces {
    all {
        depth-threshold high number low number;
        latency-threshold high number low number;
        queue-statistics;
        no-queue-statistics;
        traffic-statistics;
        no-traffic-statistics;
    }
    interface-name {
        depth-threshold high number low number;
        latency-threshold high number low number;
        queue-statistics;
        no-queue-statistics;
        traffic-statistics;
        no-traffic-statistics;
    }
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Configure physical interfaces for monitoring traffic and queue statistics by the network analytics processes running on the Packet Forwarding Engine and Routing Engine. You may specify one interface or all interfaces in your configuration.

> **NOTE**: The configuration for a specific interface supersedes the global configuration for all interfaces. You can configure traffic and queue monitoring for physical interfaces only; logical interfaces and Virtual Chassis port (VCP) interfaces are not supported.

> **NOTE**: Disabling the queue or traffic monitoring (using the `no-queue-statistics` or `no-traffic-statistics` configuration statements) supersedes the configuration (enabling) of the feature.

Options

all
Configure all interfaces on the device for high-frequency monitoring.

interface-name
Name of the interface to configure for high-frequency monitoring.

no-queue-statistics
Disable the collection of queue statistics.

> **NOTE**: The `no-queue-statistics` statement supersedes the `queue-statistics` statement.

no-traffic-statistics
Disable the collection of traffic statistics.

> **NOTE**: The `no-traffic-statistics` statement supersedes the `traffic-statistics` statement.

queue-statistics
Enable the collection of queue statistics for a specific interface or all interfaces.

traffic-statistics
Enable the collection of traffic statistics for a specific interface or all interfaces.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

interfaces (Analytics Resource)

IN THIS SECTION

- Syntax | 9119
- Hierarchy Level | 9120
- Description | 9120
- Options | 9120
- Required Privilege Level | 9120
- Release Information | 9120

Syntax

```
interfaces {
    interface-name {
        resource-profile profile-name;
```

```
    }
}
```

Hierarchy Level

```
[edit services analytics resource]
```

Description

Apply the network analytics resource profile to an interface for which you wish to enable queue or traffic statistics monitoring. The resource profile is a template that specifies the parameters for queue and traffic monitoring, as well as for the depth and latency thresholds.

> **NOTE**: The `interfaces` statement in the `[edit services analytics resource]` hierarchy is not available in Junos OS Releases prior to 13.2X51-D15.

Options

interface-name	Name of the interface for which a resource profile has been configured.
resource-profile *profile-name*	Name of a resource profile containing the analytics parameters that have been specified for interfaces. Information contained in a resource profile includes the configuration of queue and traffic monitoring (whether enabled or disabled), and values for the depth and latency thresholds (if applicable).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

interfaces (ANCP)

IN THIS SECTION

- Syntax | 9121
- Hierarchy Level | 9121
- Description | 9122
- Options | 9122
- Required Privilege Level | 9122
- Release Information | 9122

Syntax

```
interfaces {
    interface-set interface-set-name {
        access-identifier identifier-string;
        underlying-interface underlying-interface-name;
    }
    interface-name {
        access-identifier identifier-string
    }
}
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Identify the subscribers whose traffic is reported and shaped by the ANCP agent.

Options

interface-name—Name of a logical interface supporting a single VLAN that carries traffic to the subscriber identified by the access-loop circuit identifier.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent

Associating an Access Node with Subscribers for ANCP Agent Operations

interfaces (Autoinstallation)

IN THIS SECTION

- Syntax | 9123
- Hierarchy Level | 9123
- Description | 9123
- Options | 9123
- Required Privilege Level | 9124

Release Information | 9124

Syntax

```
interfaces {
    interface-name {
        bootp;
        rarp;
    }
}
```

Hierarchy Level

```
[edit system autoinstallation]
```

Description

Configure the interface on which to perform autoinstallation. A request for an IP address is sent from the interface. Specify the IP address procurement protocol.

NOTE: When you run the `system autoinstallation` command, the command will configure unit 0 logical interface for all the active state physical interfaces. However, few commands like `fabric-options` do not allow its physical interface to be configured with a logical interface. If the `system autoinstallation` and the `fabric-options` commands are configured together the following message is displayed `incompatible with 'system autoinstallation'`.

Options

- `bootp`—Enables BOOTP or DHCP during autoinstallation.
- `rarp`—Enables RARP during autoinstallation.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Autoinstallation Overview

Configuring Autoinstallation of Configuration Files

interfaces (Chassis Synchronization Source)

IN THIS SECTION

- Syntax | 9124
- Hierarchy Level | 9125
- Description | 9125
- Options | 9125
- Additional Information | 9125
- Required Privilege Level | 9125
- Release Information | 9125

Syntax

```
interfaces interface-name {
    priority number;
    quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
    request request (force-switch | lockout);
```

```
    wait-to-restore minutes;
}
```

Hierarchy Level

```
[edit chassis synchronization source]
```

Description

Configure an interface with an upstream clock source to be considered one of the candidate sources for the chassis synchronization by the clock source selection algorithm.

Options

interface-name Consider the specified interface as a candidate source.

The remaining statements are explained separately. See CLI Explorer.

Additional Information

A maximum of 8 clock sources can be configured.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

External Clock Synchronization Overview for ACX Series Routers

Configuring External Clock Synchronization for ACX Series Routers

Synchronous Ethernet Overview

show chassis synchronization

interfaces (Class of Service)

IN THIS SECTION

- Syntax | 9126
- Hierarchy Level | 9127
- Description | 9127
- Options | 9127
- Required Privilege Level | 9127
- Release Information | 9127

Syntax

```
interfaces {
    interface-name  {
        congestion-notification-profile profile-name {
            input {
                ieee-802.1 {
                    code-point up-bits pfc;
                    }
                }
            }
        }
        scheduler-map map-name;
        unit logical-unit-number {
            forwarding-class class-name;
            classifiers {
                (dscp | ieee-802.1 | inet-precedence) (classifier-name | default);
            }
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Configure interface-specific class-of-service (CoS) properties for incoming packets.

Options

interface-name Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Example: Configuring CoS on EX Series Switches
- Defining CoS Classifiers (CLI Procedure)
- Defining CoS Classifiers (J-Web Procedure)
- Defining CoS Forwarding Classes (CLI Procedure)
- Defining CoS Forwarding Classes (J-Web Procedure)
- Defining CoS Schedulers and Scheduler Maps (CLI Procedure)
- Defining CoS Schedulers (J-Web Procedure)
- *Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)*

interfaces (Class of Service)

IN THIS SECTION

- Syntax | 9128
- Hierarchy Level | 9129
- Description | 9129
- Options | 9129
- Required Privilege Level | 9129
- Release Information | 9129

Syntax

```
interfaces interface-name {
    classifiers {
        (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
    }
    congestion-notification-profile profile-name;
    dedicated-buffer-profile profile-name;
    forwarding-class forwarding-class-name;
    forwarding-class-set forwarding-class-set-name {
        output-traffic-control-profile                        profile-name;
    }
    rewrite-value {
        input {
            ieee-802.1{
                code-point code-point-bits;
            }
        }
    }
    scheduler-map scheduler-map-name
    unit logical-unit-number {
        classifiers {
            (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
        }
        forwarding-class                                  class-name;
        rewrite-rules {
```

```
            (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Configure interface-specific CoS properties for incoming packets.

> **NOTE**: Only switches that support direct port scheduling also support applying a scheduler map directly to an interface. When using enhanced transmission selection (ETS) hierarchical port scheduling, you cannot apply a scheduler map directly to an interface; instead, you associate the scheduler map with a traffic control profile and apply the traffic control profile to the interface.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Assigning CoS Components to Interfaces

interfaces (Class of Service)

IN THIS SECTION

- Syntax | 9130
- Hierarchy Level | 9132
- Description | 9132
- Options | 9132
- Required Privilege Level | 9132
- Release Information | 9132

Syntax

```
interfaces {
    interface-name {
        classifiers {
            dscp (classifier-name | default) {
        }
            ieee-802.1 (classifier-name | default)  vlan-tag (inner | outer | classifier-name);
            inet-precedence (rewrite-name | default);
        }
        input-scheduler-map map-name;
        input-shaping-rate rate;
        irb {
            unit logical-unit-number {
                classifiers {
                    type (classifier-name | default);
                    no-default;
                }
                rewrite-rules {
                    dscp (rewrite-name | default);
                    dscp-ipv6 (rewrite-name | default);
                    exp (rewrite-name | default)protocol protocol-types;
```

```
                ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                inet-precedence (rewrite-name | default);
            }
        }
    }
    member-link-scheduler (replicate | scale);
    monitoring-profile profile-name shared;
    rewrite-rules {
        dscp (rewrite-name | default);
        ieee-802.1 (rewrite-name | default) vlan-tag (outer);
        inet-precedence (rewrite-name | default);
        }
    }
    scheduler-map map-name;
    scheduler-map-chassis map-name;
    shaping-rate rate;
    slice slice-name {
        output-traffic-control-profile profile-name;
    }
    unit logical-unit-number {
        classifiers {
            type (classifier-name | default) family (mpls | inet);
        }
        forwarding-class class-name;
        fragmentation-map map-name;
        input-shaping-rate (percent percentage | rate);
        input-traffic-control-profile profile-name shared-instance instance-name;
        output-traffic-control-profile profile-name shared-instance instance-name;
        per-session-scheduler;
        policy-map policy-map-name;
        policy-map-marking;
        rewrite-rules {
            dscp (rewrite-name | default);
            dscp-ipv6 (rewrite-name | default);
            exp (rewrite-name | default) protocol protocol-types;
            exp-push-push-push default;
            exp-swap-push-push default;
            ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
            inet-precedence (rewrite-name | default);
        }
        scheduler-map map-name;
        shaping-rate rate;
        translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp |
```

```
to-inet-precedence-from-inet-precedence) table-name;
        }
    }
    interface-set interface-set-name {
        excess-bandwidth-share;
        internal-node;
        output-traffic-control-profile profile-name;
        output-traffic-control-profile-remaining profile-name;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Associate class of service (CoS) configuration elements with an interface.

> **NOTE**: Only devices that support direct port scheduling also support applying a scheduler map directly to an interface. When using hierarchical port scheduling, you cannot apply a scheduler map directly to an interface. Instead, you associate the scheduler map with a traffic control profile and apply the traffic control profile to the interface.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Interface-set level added in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic

Configuring Rewrite Rules

interfaces (Dynamic CoS Definition)

IN THIS SECTION

- Syntax | 9133
- Hierarchy Level | 9134
- Description | 9134
- Options | 9134
- Required Privilege Level | 9134
- Release Information | 9134

Syntax

```
interfaces {
    interface-name {
        unit logical-unit-number {
            classifiers {
                dscp (classifier-name | default);
                dscp-ipv6 (classifier-name | default);
                ieee-802.1 (classifier-name | default) vlan-tag (inner | outer)
                inet-precedence (classifier-name | default);
            }
            output-traffic-control-profile (profile-name  | $junos-cos-traffic-control-profile);
            report-ingress-shaping-rate;
            rewrite-rules {
                dscp (rewrite-name | default);
                dscp-ipv6 (rewrite-name | default);
```

```
                ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                inet-precedence (rewrite-name | default);
            }
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service]
```

Description

Configure interface-specific CoS properties for incoming packets in a dynamic client profile or a dynamic service profile.

Options

interface-name—Either the specific name of the interface you want to assign to the dynamic profile or the interface variable ($junos-interface-ifd-name). The interface variable is dynamically replaced with the interface the client accesses when connecting to the router.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile

interfaces (EX Series switches)

IN THIS SECTION

- Syntax | 9135
- Hierarchy Level | 9142
- Description | 9142
- Options | 9142
- Required Privilege Level | 9143
- Release Information | 9143

Syntax

```
No Link Title
No Link Title
No Link Title
No Link Title
No Link Title
No Link Title
No Link Title
No Link Title
No Link Title
```

interfaces ae

```
ae x {
    accounting-profile name;
                                aggregated-ether-options {
                                    (flow-control | no-flow-control);
        lacp {
            (active | passive);
            admin-key key;
            periodic interval;
            system-id mac-address;
        }
        (link-protection | no-link-protection);
```

```
        link-speed speed;
        (loopback | no-loopback);
        minimum-links number;
    }
                            description                         text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
                           mtu                          bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
                                description                         text;
        disable;
        family family-name {...}
        proxy-arp  (restricted | unrestricted);
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
                            vlan-tagging;
}
```

interfaces ge

```
ge-fpc/pic/port {
    accounting-profile name;
                            description                         text;
    disable;
                                ether-options {
        802.3ad {
            aex;
            (backup | primary);
            lacp {
                force-up;
            }
        }
                                    (auto-negotiation | no-auto-negotiation);
                                    (flow-control | no-flow-control);
```

```
        ieee-802-3az-eee;
                                link-mode                          mode;
        (loopback | no-loopback);
        speed (auto-negotiation | speed);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    media-type;
                            mtu                        bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
                                description                        text;
        disable;
        family family-name {...}
        proxy-arp  (restricted | unrestricted);
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
                            vlan-tagging;
}
```

interfaces interface-range

```
interface-range  name {
    accounting-profile name;
                            description                        text;
    disable;
                                ether-options {
        802.3ad {
            aex;
            (backup | primary);
            lacp {
                force-up;
            }
        }
                                (auto-negotiation | no-auto-negotiation);
```

```
                                    (flow-control | no-flow-control);
        ieee-802-3az-eee;
                                     link-mode                              mode;
        (loopback | no-loopback);
        speed (auto-negotiation | speed);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    member interface-name;
    member-range  starting-interface name to  ending-interface name;
    mtu bytes;
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
                                 description                          text;
        disable;
        family family-name {...}
        proxy-arp  (restricted | unrestricted);
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
                                vlan-tagging;
}
```

interfaces lo0

```
lo0 {
    accounting-profile name;
                                description                          text;
    disable;
    hold-time up milliseconds down milliseconds;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
                                 description                          text;
        disable;
        family family-name {...}
        (traps | no-traps);
```

```
    }
}
```

interfaces me0

```
me0 {
    accounting-profile name;
                            description                         text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
                                description                         text;
        disable;
        family family-name {...}
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
                            vlan-tagging;
}
```

interfaces traceoptions

```
traceoptions {
    file <filename> <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
    flag flag <disable>;
    no-remote-trace;
}
```

interfaces vlan

```
vlan {
    accounting-profile name;
```

```
    description  text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
                            mtu                         bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
                                description                         text;
        disable;
        family family-name {...}
        proxy-arp  (restricted | unrestricted);
        (traps | no-traps);
    }
}
```

interfaces vme

```
vme {
    accounting-profile name;
                            description                         text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
                            mtu                         bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
                                description                         text;
        disable;
        family family-name {...}
        (traps | no-traps);
```

```
        vlan-id vlan-id-number;
    }
                                vlan-tagging;
}
```

interfaces xe

```
xe-fpc/pic/port {
    accounting-profile name;
                                description                         text;
    disable;
                                    ether-options {
        802.3ad {
            aex;
            (backup | primary);
            lacp {
                force-up;
            }
        }
                                        (flow-control | no-flow-control);
                                        link-mode                               mode;
        (loopback | no-loopback);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
                                mtu                             bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
                                    description                         text;
        disable;
        family family-name {...}
        proxy-arp  (restricted | unrestricted);
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
```

```
                                vlan-tagging;
}
```

Hierarchy Level

```
[edit]
```

Description

Configure interfaces on EX Series switches.

Options

See Table 51 on page 9142 for the interface types and protocol-family options supported on the switch. Different protocol families support different subsets of the interface types on the switch. See the `family` statement for syntax of the protocol families supported for switch interfaces.

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

Table 51: Interface Types and Their Supported Protocol Families

Interface Type	Description	Supported Protocol Families					
		ccc	ethernet-switching	inet	inet6	iso	mpls
ae	Aggregated Ethernet interface (also referred to as a link aggregation group [LAG])	✓*	✓	✓	✓	✓	✓
ge	Gigabit Ethernet interface	✓	✓	✓	✓	✓	✓
interface-range	Interface-range configuration	Supported protocol families are the ones supported by the interface types that compose the range.					
lo0	Loopback interface			✓	✓	✓	✓

Table 51: Interface Types and Their Supported Protocol Families *(Continued)*

Interface Type	Description	Supported Protocol Families					
		ccc	ethernet-switching	inet	inet6	iso	mpls
me0	Management Ethernet interface		✓	✓	✓	✓	✓
vlan	Routed VLAN interface (RVI)			✓	✓	✓	
vme	Virtual management Ethernet interface			✓	✓	✓	✓
xe	10-Gigabit Ethernet interface	✓	✓	✓	✓	✓	✓
		*Supported on EX8200 switches only					

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Configuring Gigabit Ethernet Interfaces (CLI Procedure)
- Configuring Aggregated Ethernet Links (CLI Procedure)
- Configuring a Layer 3 Subinterface (CLI Procedure)
- Configuring Routed VLAN Interfaces on Switches (CLI Procedure)
- *Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)*

Interfaces Overview for Switches

Junos OS Interfaces Fundamentals Configuration Guide

Junos OS Ethernet Interfaces Configuration Guide

interfaces external

IN THIS SECTION

- Syntax | 9144
- Hierarchy Level | 9145
- Description | 9145
- Options | 9145
- Required Privilege Level | 9145
- Release Information | 9145

Syntax

```
interfaces {
external/external-0-0/external-1/0 {
    e1-options {
        framing (g704 | g704-no-crc4);
        line-encoding (ami | hdb3);
        sabit bit;
    }
    signal-type (1mhz | 5mhz | 10mhz | 2048khz | t1 | e1);
    t1-options {
        framing (esf | sf);
        line-encoding (ami | b8zs);
    }
    pulse-per-second-enable;
    }
}
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Starting from Junos OS Release 12.3, configure options for the external clock source interface—external—for SCBE.

Starting from Junos OS Release 13.3, configure options for the two external clock source interfaces—external-0/0 and external-1/0—for SCBE2.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

interfaces (Junos Node Slicing)

IN THIS SECTION

Syntax | 9146

- Hierarchy Level | 9146
- Description | 9146
- Options | 9146
- Required Privilege Level | 9147
- Release Information | 9147

Syntax

```
interfaces interface-name{
    unit unit {
        family (inet|inet6) {
            address address;
        }
    }
}
```

Hierarchy Level

[edit]

Description

Configure an IP address for the JDM management interface.

Options

interface-name Name of the interface.

unit *unit* Interface unit number. This is a logical unit number. The only supported value is `0`.

family (inet|inet6) Protocol family.

- inet—Indicates IPv4.
- inet6—Indicates IPv6.

address *address* IP address of the interface.

Required Privilege Level

root

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

server (Junos node slicing) | 16606

virtual-network-functions (Junos node slicing) | 30773

interfaces (MACsec)

IN THIS SECTION

- Syntax | 9147
- Hierarchy Level | 9148
- Description | 9148
- Default | 9148
- Required Privilege Level | 9148
- Release Information | 9148

Syntax

```
interfaces interface-name {
    connectivity-association connectivity-association-name;
}
```

Hierarchy Level

```
[edit security macsec]
```

```
[edit security macsec cluster-data-port]
```

Description

Applies the specified connectivity association to the specified interface to enable MACsec.

One connectivity association can be applied to multiple interfaces.

You must always use this statement to apply a connectivity association to an interface to enable MACsec. You must complete this configuration step regardless of whether MACsec is enabled using static connectivity association key (CAK) security mode or static secure association key (SAK) security mode.

If you are enabling MACsec using static SAK security mode and need to configure MACsec on inbound and outbound traffic on the same interface, you must configure a connectivity association with one secure channel for inbound traffic and a second secure channel for outbound traffic. The connectivity association is then applied to the interface using this statement to enable MACsec for traffic entering and leaving the interface.

Specify chassis cluster fabric interface on which MACsec is enabled. For SRX340, and SRX345 devices, the fabric interface can be any 1 G Ethernet interface. Use this configuration to apply a connectivity association to an interface, which enables Media Access Control Security (MACsec) on that interface.

Default

Interfaces are not associated with any connectivity associations, by default.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced for SRX Series Firewalls in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

interfaces (MACsec for MX Series)

IN THIS SECTION

- Syntax | **9149**
- Hierarchy Level | **9150**
- Description | **9150**
- Default | **9150**
- Options | **9151**
- Required Privilege Level | **9151**
- Release Information | **9151**

Syntax

```
interfaces interface-name {
    connectivity-association connectivity-association-name;
    unit unit-number {
        connectivity-association connectivity-association-name;
    }
}
```

Hierarchy Level

```
[edit security macsec]
```

Description

Applies the specified connectivity association to the specified interface to enable MACsec.

One connectivity association can be applied to multiple interfaces.

You must always use this statement to apply a connectivity association to an interface to enable MACsec. You must complete this configuration step regardless of whether MACsec is enabled using static connectivity association key (CAK) security mode or static secure association key (SAK) security mode.

If you are enabling MACsec using static SAK security mode and need to configure MACsec on inbound and outbound traffic on the same interface, you must configure a connectivity association with one secure channel for inbound traffic and a second secure channel for outbound traffic. The connectivity association is then applied to the interface using this statement to enable MACsec for traffic entering and leaving the interface.

> **NOTE**: Starting in Junos OS Release 16.1R2, when Media Access Control Security (MACsec) is enabled on an interface, the interface flow control capability is enabled by default, regardless of the configuration that you set using the `(flow-control | no-flow-control)` statement at the `[edit interfaces` *`interface- name`* `gigether-options]` hierarchy level. When MACsec is disabled, interface flow control is restored to the configuration that you set using the `flow-control` statement at the `[edit interfaces]` hierarchy level. When MACsec is enabled, additional header bytes are added to the packet by the MACsec PHY. With line rate traffic, when MACsec is enabled and flow control is disabled, the pause frames sent by the MACsec PHY are terminated by the MIC's MAC (enhanced 20-port Gigabit Ethernet MICs on MX Series routers) and not transferred to the Packet Forwarding Engine, causing framing errors. Therefore, when MACsec is enabled on an interface, flow control is also automatically enabled on such an interface.

Default

Interfaces are not associated with any connectivity associations, by default.

Options

`connectivity-association` *`connectivity-association-name`*	Specify the connectivity association to assign to the interface. A connectivity association is a set of MACsec attributes that are used by interfaces to create secure inbound and outbound channels for encrypted traffic.
`unit` *`unit-number`*	Applies the specified connectivity association to a logical interface.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

Support for `unit` option introduced in Junos OS Release 19.3 for MPC7E-10GE line cards.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

interfaces (OVSDB)

IN THIS SECTION

- Syntax | 9152
- Hierarchy Level | 9152
- Description | 9152
- Options | 9152
- Required Privilege Level | 9152
- Release Information | 9152

Syntax

```
interfaces interface-name;
```

Hierarchy Level

```
[edit protocols ovsdb]
```

Description

Specify the physical interfaces on a Juniper Networks device that you want the *Open vSwitch Database (OVSDB)* protocol to manage. Typically, the only interfaces that need to be managed by OVSDB are interfaces that are connected to physical servers.

Options

interface-name	Name of the interface.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

> *Setting Up the OVSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs*
>
> *Setting Up OVSDB on Juniper Networks Devices That Support the Dynamic Configuration of VXLANs*

interfaces (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 9153
- Hierarchy Level | 9153
- Description | 9153
- Options | 9153
- Required Privilege Level | 9154
- Release Information | 9154

Syntax

```
interfaces {
    interface-name port-id port;
}
```

Hierarchy Level

```
[edit protocols openflow switch switch-name]
```

Description

Configure a Layer 2 interface as an OpenFlow-enabled interface.

Options

interface-name Name of the interface, including the logical unit number—for example, ge-1/1/0.0.

port-id *port* (Optional) Unique numeric value specifying the port ID associated with the OpenFlow interface. You can manually configure a port ID in the range 1 through 32640. If you do not specify a port, the system generates a value in the range from 32641 through 65280.

- **Range:** 1 through 32,640

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Support for OpenFlow on Devices Running Junos OS

openflow (Protocols OpenFlow) | **11842**

switch (Protocols OpenFlow) | **28755**

interfaces (Q-in-Q Tunneling)

IN THIS SECTION

- Syntax | **9155**
- Hierarchy Level | **9155**
- Description | **9155**
- Options | **9155**
- Required Privilege Level | **9155**
- Release Information | **9155**

Syntax

```
interfaces interface-name {
                        no-mac-learning;
}
```

Hierarchy Level

```
[edit ethernet-switching-options]
```

Description

Configure settings for interfaces that have been assigned to family **ethernet-switching**.

Options

interface-name --Name of an interface that is configured for **family ethernet-switching**.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Q-in-Q Tunneling and VLAN Translation

interfaces (QFX Series, ACX Series)

IN THIS SECTION

- Syntax | 9156
- Hierarchy Level | 9164
- Description | 9164
- Options | 9164
- Required Privilege Level | 9165
- Release Information | 9165

Syntax

The following statements and their associated substatements are not supported on OCX Series switches: `auto-negotiation`, `speed`, `ethernet-switching`, `fcoe-lag`, `fibre-channel`, `fibrechannel-options`, `mc-ae`, `vlan`, `vlan-id`, and `vlan-tagging`.

```
interfaces {
    aex {
        disable;
            aggregated-ether-options {
            configured-flow-control {
                rx-buffers (on | off);
                tx-buffers (on | off);
            }
            (fcoe-lag | no-fcoe-lag);
            flexible-vlan-tagging;
            (flow-control | no-flow-control);
            lacp mode {
                admin-key key;
                force-up;
                periodic interval;
                system-id mac-address;
            }
            link-speed speed;
            local-bias;
            loopback;
```

```
            no-loopback;
            minimum-links number;
        }
            mc-ae {
                chassis-id chassis-id;
                mc-ae-id mc-ae-id;
                mode (active-active);
                status-control (active | standby);
            }
        description text;
        gratuitous-arp-reply | no-gratuitous-arp-reply)
        hold-time down milliseconds up milliseconds;
        mtu bytes;
        no-gratuitous-arp-request;
        traceoptions;
        (traps | no traps);
        unit logical-unit-number {
            disable;
            description text;
            family {
                ethernet-switching {
                    filter input filter-name;
                    filter output filter-name;
                    native-vlan-id vlan-id;
                    port-mode mode;
                    reflective-relay;
                    vlan {
                        members [ (all | names | vlan-ids) ];
                    }
                }
                inet {
                    address address {
                        primary;
                    }
                    filter input filter-name;
                    filter output filter-name;
                    primary;
                    targeted-broadcast;
                }
            (traps | no traps);
            vlan-id vlan-id-number;
        }
        vlan-tagging;
```

```
    }
    interface-range {
        powersaving {
            member-range starting-interface-name to  ending-interface-name;
            unused;
        }
        interface-range-name {
            disable;
            description text;
            ether-options {
                802.3ad aex {
                    lacp {
                     force-up;

                }
            }
            (auto-negotiation| no-auto-negotiation);
             configured-flow-control {
                 rx-buffers (on | off);
                 tx-buffers (on | off);
             }
             (flow-control | no-flow-control);
             link-mode mode;
             speed (auto-negotiation | speed);
        }
        hold-time milliseconds down milliseconds;
        member interface-name;
        member-range starting-interface-name to  ending-interface-name;
        unused;
        mtu bytes;
        unit logical-unit-number {
            disable;
            description text;
            family family-name {...}
            (traps | no traps);
            vlan-id vlan-id-number;
        }
    }
    lo0 {
        disable;
        description text;
        hold-time milliseconds down milliseconds;
        traceoptions;
```

```
        (traps | no traps);
        unit logical-unit-number {
            disable;
            description text;
            family {
                inet {
                    address address {
                        primary;
                    }
                    filter input filter-name;
                    filter output filter-name;
                    primary;
                    targeted-broadcast;
                }
            (traps | no traps);
        }
    }
    mex {
        disable;
        description text;
        hold-time milliseconds down milliseconds;
        (gratuitous-arp-reply | no-gratuitous-arp-reply);
        no-gratuitous-arp-request;
        traceoptions;
        traps;
        unit logical-unit-number {
            disable;
        description text;
            family {
                ethernet-switching {
                    filter input filter-name;
                    filter output filter-name;
                    native-vlan-id vlan-id;
                    port-mode mode;
                    reflective-relay;
                    vlan {
                        members [ (all | names | vlan-ids) ];
                    }
                }
                inet {
                    address address {
                    primary;
                    filter input filter-name;
```

```
                    filter output filter-name;
                    primary;
                    targeted-broadcast;
                    }
                }
        traps;
        vlan-id vlan-id-number;
        }
    vlan-tagging;
    vlan {
        disable;
        description  text;
        (gratuitous-arp-reply| no-gratuitous-arp-reply);
        hold-time milliseconds down milliseconds;
        mtu bytes;
        no-gratuitous-arp-request;
        traceoptions;
        (traps | no traps);
        unit logical-unit-number {
            description text;
            disable;
            family {
                inet {
                    address address {
                        primary;
                    }
                    filter input filter-name;
                    filter output filter-name;
                    primary;
                    targeted-broadcast;
                }
            (traps | no traps);
        }
    }
    fc-0/0/port {
        fibrechannel-options {
            bb-sc-n;
            (loopback | no-loopback);
            speed (auto-negotiation | 2g | 4g | 8g);
        }
        unit logical-unit-number {
            disable;
            description text;
```

```
            family {
                fibre-channel {
                    port-mode np-port;
                }
            (traps | no traps);
        }
    ge-0/0/port {
        disable;
        description text;
        ether-options {
            802.3ad aex {
                lacp {
                    force-up;
                    primary;
                }
            }
            (auto-negotiation | no-auto-negotiation);
            configured-flow-control {
                rx-buffers (on | off);
                tx-buffers (on | off);
            }
            flexible-vlan-tagging;
            (flow-control | no-flow-control);
            link-mode mode;
            loopback;
            no-loopback;
            speed (auto-negotiation | speed);
        }
        gratuitous-arp-reply| no-gratuitous-arp-reply);
        hold-time milliseconds down milliseconds;
        mac
        mtu bytes;
        no-gratuitous-arp-request;
        traceoptions;
        (traps | no traps);
        unit logical-unit-number {
            description text;
            disable;
            family {
                ethernet-switching {
                    filter input filter-name;
                    filter output filter-name;
                    native-vlan-id vlan-id;
```

```
                    port-mode mode;
                    reflective-relay;
                    vlan {
                        members [ (all | names | vlan-ids) ];
                    }
                }
                inet {
                    address address {
                        primary;
                    }
                    filter input filter-name;
                    filter output filter-name;
                    primary;
                    targeted-broadcast;
                }
            (traps | no traps);
            vlan-id vlan-id-number;
        }
        vlan-tagging;
    }
    vrrp-group group-id {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-key key;
         authentication-type authentication;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bits-per-second priority-cost priority;
                priority-cost priority;
            }
            priority-hold-time seconds;
            route prefix/prefix-length routing-instance instance-name priority-cost priority;
            }
        }
        virtual-address [ addresses ];
    }
    xe-0/0/port {
        disable;
```

```
        description text;
        ether-options {
            802.3ad aex {
                lacp {
                    force-up;
                    (primary | backup);
                }
            }
            configured-flow-control {
                rx-buffers (on | off);
                tx-buffers (on | off);
            }
            flexible-vlan-tagging;
            (flow-control | no-flow-control);
            loopback;
            no-loopback;
        }
        (gratuitous-arp-reply| no-gratuitous-arp-reply
        hold-time milliseconds down milliseconds;
        mac
        mtu bytes;
        no-gratuitous-arp-request;
        traceoptions;
        (traps | no traps);
        unit logical-unit-number {
            disable;
            description text;
            family {
                ethernet-switching {
                    filter input filter-name;
                    filter output filter-name;
                    native-vlan-id vlan-id;
                    port-mode mode;
                    reflective-relay;
                    vlan {
                        members [ (all | names | vlan-ids) ];
                    }
                }
                fibre-channel {
                        port-mode (f-port | np-port);
                }
                inet {
                    address address {
```

```
                        primary;
                    }
                    filter input filter-name;
                    filter output filter-name;
                    primary;
                    targeted-broadcast;
                }
            (traps | no traps);
            vlan-id vlan-id-number;
        }
        vlan-tagging;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure the interfaces on the QFX Series, ACX Series and OCX Series.

The following statements and their associated substatements are not supported on OCX Series switches: `auto-negotiation`, `ethernet-switching`, `fcoe-lag`, `fibre-channel`, `fibrechannel-options`, `mc-ae`, `speed`, `vlan`, `vlan-id`, and `vlan-tagging`

Most standard Junos OS configuration statements are available in the Junos OS for a switch. This topic lists Junos OS statements that you commonly use when configuring a switch as well as statements added to support switches only.

View the Feature Explorer to know the support platforms for the powersaving feature.

Options

ae*x*—Configure an aggregated Ethernet interface.

xe-0/0/*port/*—Configure a 10-Gigabit Ethernet interface.

ge-0/0/*port/*—Configure a Gigabit Ethernet interface.

fc-0/0/*port/*—Configure a Fibre Channel interface.

me*X/*—Configure a management interface.

mc-ae—Configure a multichassis aggregated Ethernet (MC-AE) interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Statement support for QFX series introduced in Junos OS EVO Release 21.3

RELATED DOCUMENTATION

- *Interfaces Overview for Switches*
- *Gigabit Ethernet Interface*
- Configuring a Layer 3 Logical Interface
- interface-range (QFX Series) | **9003**
- *Understanding Interface Ranges for Switches*
- *Configuring Link Aggregation*

interfaces (Security Zones)

IN THIS SECTION

- Syntax | **9166**
- Hierarchy Level | **9166**
- Description | **9166**
- Options | **9166**
- Required Privilege Level | **9166**
- Release Information | **9166**

Syntax

```
interfaces interface-name {
    host-inbound-traffic {
        protocols protocol-name {
                            except;
        }
        system-services service-name {
                            except;
        }
    }
}
```

Hierarchy Level

```
[edit security zones functional-zone management],
[edit security zones security-zone zone-name]
```

Description

Specify the set of interfaces that are part of the zone.

Options

interface-name —Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

interfaces (Services Dynamic Flow Capture)

IN THIS SECTION

- Syntax | 9167
- Hierarchy Level | 9167
- Description | 9167
- Options | 9167
- Required Privilege Level | 9167
- Release Information | 9168

Syntax

```
interfaces interface-name;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Specify the DFC interface used with the control source configured in the same capture group.

Options

interface-name—Name of the DFC interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the DFC PIC Interface

interfaces (sFlow)

IN THIS SECTION

- Syntax: ACX Routers, OCX1100, EX4600 and QFX Switches | 9168
- Syntax: EX Series Switches, PTX Routers | 9168
- Hierarchy Level | 9169
- Description | 9169
- Options | 9169
- Required Privilege Level | 9169
- Release Information | 9169

Syntax: ACX Routers, OCX1100, EX4600 and QFX Switches

```
interfaces interface-name {
    polling-interval seconds;
    sample-rate number;
}
```

Syntax: EX Series Switches, PTX Routers

```
interfaces interface-name {
    polling-interval seconds;
    sample-rate {
```

```
        egress number;
        ingress number;
    }
}
```

Hierarchy Level

```
[edit protocols sflow]
```

Description

Configure sFlow network traffic monitoring on the specified interface on the device. You can configure sFlow parameters such as polling interval and sampling rate with different values on different interfaces, and you can also disable sFlow monitoring on individual interfaces.

The remaining statements are explained separately. See CLI Explorer.

Options

interface-name—Name of the interface on which to configure sFlow parameters.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring sFlow Technology for Network Monitoring (CLI Procedure)

Example: Monitoring Network Traffic Using sFlow Technology

Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches

Configuring sFlow Technology for Network Monitoring (CLI Procedure)

Example: Configuring sFlow Technology to Monitor Network Traffic on MX Series Routers

interfaces (Static and Dynamic Subscribers)

IN THIS SECTION

- Syntax | 9170
- Hierarchy Level | 9175
- Description | 9175
- Options | 9175
- Required Privilege Level | 9176
- Release Information | 9176

Syntax

```
interfaces {
    interface-name {
        unit logical-unit-number {
            actual-transit-statistics;
            auto-configure {
                agent-circuit-identifier {
                    dynamic-profile profile-name;
                }
                line-identity {
                    include {
                        accept-no-ids;
                        circuit-id;
                        remote-id;
                    }
                    dynamic-profile profile-name;
                }
            }
            family family {
                access-concentrator name;
                address address;
                direct-connect;
                duplicate-protection;
                dynamic-profile profile-name;
                filter {
```

```
                    adf {
                        counter;
                        input-precedence precedence;
                        not-mandatory;
                        output-precedence precedence;
                        rule rule-value;
                    }
                    input filter-name {
                        precedence precedence;
                        shared-name  filter-shared-name;
                    }
                    output filter-name {
                        precedence precedence;
                        shared-name  filter-shared-name;
                    }
                }
                max-sessions number;
                max-sessions-vsa-ignore;
                rpf-check {
                    mode loose;
                }
                service {
                     input {
                         service-set service-set-name {
                            service-filter filter-name;
                        }
                        post-service-filter filter-name;
                    }
                    output {
                         service-set service-set-name {
                            service-filter filter-name;
                        }
                    }
                }
                service-name-table table-name
                short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-
seconds>;
                unnumbered-address interface-name <preferred-source-address address>;
            }
            filter {
                input filter-name (
                    precedence precedence;
                    shared-name  filter-shared-name;
```

```
                }
                output filter-name {
                    precedence precedence;
                    shared-name  filter-shared-name;
                }
            }
            host-prefix-only;
            ppp-options {
                chap;
                pap;
            }
            proxy-arp;
            service {
                pcef  pcef-profile-name {
                    activate rule-name | activate-all;
                }
            }
            targeted-options {
                backup backup;
                group group;
                primary primary;
                weight ($junos-interface-target-weight | weight-value);
            }
            vlan-id;
            vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
        }
        vlan-tagging;
    }
    interface-set interface-set-name {
        interface interface-name {
            unit logical unit number {
                advisory-options {
                    downstream-rate rate;
                    upstream-rate rate;
                }
            }
        }
        pppoe-underlying-options {
            max-sessions number;
        }
    }
    demux0 {
        unit logical-unit-number {
```

```
            demux-options {
                underlying-interface interface-name
            }
            family family {
                access-concentrator name;
                address address;
                direct-connect;
                duplicate-protection;
                dynamic-profile profile-name;
                demux-source {
                    source-prefix;
                }
                filter {
                    input filter-name (
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                    output filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                }
                mac-validate (loose | strict):
                max-sessions number;
                max-sessions-vsa-ignore;
                rpf-check {
                    fail-filter filter-name;
                    mode loose;
                }
                service-name-table table-name
                short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-
seconds>;
                unnumbered-address interface-name <preferred-source-address address>;
            }
            filter {
                input filter-name;
                output filter-name;
            }
            vlan-id number;
            vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
        }
    }
    pp0 {
```

```
unit logical-unit-number {
    keepalives interval seconds;
    no-keepalives;
    pppoe-options {
        underlying-interface interface-name;
        server;
    }
    ppp-options {
        aaa-options aaa-options-name;
        authentication [ authentication-protocols ];
        chap {
            challenge-length minimum minimum-length maximum maximum-length;
            local-name name;
        }
        ignore-magic-number-mismatch;
        initiate-ncp (dual-stack-passive | ipv6 | ip)
        ipcp-suggest-dns-option;
        mru size;
        mtu (size | use-lower-layer);
        on-demand-ip-address;
        pap;
        peer-ip-address-optional;
        local-authentication {
            password password;
            username-include {
                circuit-id;
                delimiter character;
                domain-name name;
                mac-address;
                remote-id;
            }
        }
    }
    family inet {
        unnumbered-address interface-name;
        address address;
        service {
            input {
                service-set service-set-name {
                    service-filter filter-name;
                }
                post-service-filter filter-name;
            }
```

```
                    output {
                        service-set service-set-name {
                            service-filter filter-name;
                        }
                    }
                }
                filter {
                    input filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                    output filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                }
            }
        }
    }
    stacked-interface-set {
        interface-set-name interface-set-name {
            interface-set-name interface-set-name;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Define interfaces for dynamic client profiles.

Options

interface-name—The interface variable (`$junos-interface-ifd-name`). The interface variable is dynamically replaced with the interface the DHCP client accesses when connecting to the router.

> **NOTE**: Though we do not recommend it, you can also enter the specific name of the interface you want to assign to the dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
- *Configuring Dynamic PPPoE Subscriber Interfaces*
- *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
- *DHCP Subscriber Interface Overview*
- *Subscribers over Static Interfaces Configuration Overview*
- Demultiplexing Interface Overview

interfaces (Video Monitoring)

IN THIS SECTION

- Syntax | 9177
- Hierarchy Level | 9178
- Description | 9179
- Options | 9179

Required Privilege Level | 9180
Release Information | 9180

Syntax

```
interfaces {
    interface-name {
        family {
            inet {
                input-flows {
                    input-flow-name {
                        destination-address [ address ];
                        destination-port [ port ];
                        source-address [ address ];
                        source-port [ port ];
                        template template-name;
                    }
                }
                output-flows {
                    output-flow-name {
                        destination-address [ address ];
                        destination-port [ port ];
                        source-address [ address ];
                        source-port [ port ];
                        template template-name;
                    }
                }
            }
            inet6 {
                input-flows {
                    input-flow-name {
                        destination-address [ address ];
                        destination-port [ port ];
                        source-address [ address ];
                        source-port [ port ];
                        template template-name;
                    }
                }
```

```
                output-flows {
                    output-flow-name {
                        destination-address [ address ];
                        destination-port [ port ];
                        source-address [ address ];
                        source-port [ port ];
                        template template-name;
                    }
                }
            }
            mpls {
                input-flows {
                    input-flow-name {
                        (destination-address [ address ] | source-address [ address ]);
                        destination-port [ port ];
                        payload-type (ipv4 | ipv6);
                        source-port [ port ];
                        template template-name;
                    }
                }
                output-flows {
                    output-flow-name {
                        (destination-address [ address ] | source-address [ address ]);
                        destination-port [ port ];
                        payload-type (ipv4 | ipv6);
                        source-port [ port ];
                        template template-name;
                    }
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit services video-monitoring]
```

Description

Define video monitoring for specified input or output flows on selected interfaces. You can configure a maximum of 256 flows for an interface.

Options

destination-address *address*

Destination IPv4 or IPv6 address or prefix value of a flow that you want to monitor. You can use up to 32 addresses.

For IPv4-over-MPLS flows, if you configure multiple addresses for both the destination and source, then either all the destination or all the source values must have the same prefix length. For example, the following is allowed, because all the destination addresses have the same prefix length.

```
[edit services video-monitoring interfaces ge-0/2/2.0 family mpls]
user@host# set input-flows input-flow-name destination-address
[203.0.13.0/24 198.51.100.0/24]
user@host# set input-flows input-flow-name source-address [172.16.0.0/12
192.0.2.11/32]
```

For IPv6-over-MPLS flows, if you configure both the destination and source address, you can use multiple addresses for either the destination or the source IP address, but not for both.

destination-port *port*

Destination port number of a flow that you want to monitor. You can use multiple port numbers and port ranges.

- **Range:** 0 through 65,535

input-flows *input-flow-name*

Name of an input flow you are defining.

interface-name

Name of the interface to monitor.

output-flows *output-flow-name*

Name of an output flow you are defining.

payload-type ipv4

Monitor video stream for IPv4-over-MPLS traffic.

payload-type ipv6

Monitor video stream for IPv6-over-MPLS traffic.

source-address *address*

Source IPv4 or IPv6 address or prefix value of a flow that you want to monitor. You can use up to 32 addresses.

For IPv4-over-MPLS flows, if you configure multiple addresses for both the destination and source, then either all the destination or all the source values must have the same prefix length. For example, the following is allowed, because all the destination addresses have the same prefix length.

```
[edit services video-monitoring interfaces ge-0/2/2.0 family mpls]
user@host# set input-flows input-flow-name destination-address
[203.0.13.0/24 198.51.100.0/24]
user@host# set input-flows input-flow-name source-address [172.16.0.0/12
192.0.2.11/32]
```

For IPv6-over-MPLS flows, if you configure both the destination and source address, you can use multiple addresses for either the destination or the source IP address, but not for both.

source-port *port*

Source port number of a flow that you want to monitor. You can use multiple port numbers and port ranges.

- **Range:** 0 through 65,535

template-name

Name of the template used to monitor the input flows or output flows on an interface. The template contains the measurement parameters for video monitoring, and is configured at the `[edit services video-monitoring templates]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

`mpls` option introduced in Junos OS Release 17.2.

`payload-type ipv6` option introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

Configuring Inline Video Monitoring on MX Series Routers

interfaces (Voice Services)

IN THIS SECTION

- Syntax | 9181
- Hierarchy Level | 9181
- Description | 9181
- Default | 9181
- Required Privilege Level | 9181
- Release Information | 9182

Syntax

```
interfaces { ... }
```

Hierarchy Level

```
[edit]
```

Description

Configure interfaces on the router.

Default

The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Network Interfaces Library for Routing Devices

interim-logging-interval

IN THIS SECTION

- Syntax | 9182
- Hierarchy Level | 9182
- Description | 9182
- Options | 9183
- Required Privilege Level | 9183
- Release Information | 9183

Syntax

```
interim-logging-interval timeout-inteval;
```

Hierarchy Level

```
[edit security nat source pool port block-allocation]
```

Description

Specify how often interim system logs are sent for active port blocks and for inactive port blocks with live sessions. Because system logs are UDP-based, they can be lost in the network. For this reason,

configuring interim logging, which triggers re-sending system logs, increases reliability. In a chassis cluster configuration, to limit generation of interim system logs to the primary node only, you must also specify the option `log on_primary_node` at the [`edit security nat source pool port block-allocation`] hierarchy level.

Options

timeout-interval—Number of seconds between interim logging messages.

- **Range:** 1800 through 86400

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D60.

interim-rate (Access)

IN THIS SECTION

- Syntax | 9184
- Hierarchy Level | 9184
- Description | 9184
- Options | 9184
- Required Privilege Level | 9184
- Release Information | 9184

Syntax

```
interim-rate rate;
```

Hierarchy Level

```
[edit access radius-options]
```

Description

Configure the rate at which RADIUS interim update requests are processed.

Options

rate—Maximum number of RADIUS requests sent per second.

- **Range:** 50 through 4000 RADIUS requests per second
- **Default:** 500 RADIUS requests per second

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

interleave-fragments

IN THIS SECTION

- Syntax | 9185
- Hierarchy Level | 9185
- Description | 9185
- Required Privilege Level | 9185
- Release Information | 9186

Syntax

```
interleave-fragments;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel unit logical-unit-number],
[edit logical-systems logical-system-name interfaces ls-fpc/pic/port unit logical-unit-number]
```

Description

For link services and voice services interfaces only, interleave long packets with high-priority packets.

Allows small delay-sensitive packets, such as voice over IP (VoIP) packets, to interleave with long fragmented packets. This minimizes the latency of delay-sensitive packets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces

internal

IN THIS SECTION

- Syntax | 9186
- Hierarchy Level | 9187
- Description | 9187
- Options | 9187
- Required Privilege Level | 9187
- Release Information | 9187

Syntax

```
internal {
    security-association {
        manual {
            direction (bidirectional | inbound | outbound) {
                protocol esp;
                spi spi-value;
                encryption {
                    algorithm 3des-cbc;
                    key ascii-text ascii-text-string;
                }
            }
        }
```

```
    }
}
```

Hierarchy Level

```
[edit security ipsec]
```

Description

(Junos-FIPS only) Define an internal security association (SA) for internal Routing-Engine–to–Routing-Engine communication.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Crypto Officer—To view and add this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec for Enabling Internal Communications Between Routing Engines for Junos OS in FIPS Mode

internal-forwarding-bandwidth

IN THIS SECTION

Syntax | 9188

- Hierarchy Level | 9188
- Description | 9188
- Options | 9188
- Required Privilege Level | 9189
- Release Information | 9189

Syntax

```
internal-forwarding-bandwidth <value>;
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

You can configure the bandwidth of the recycle interface to enable equal distribution of the recycle interface bandwidth at PFE level.

The recycle interface has a default 100Gbps per unit bandwidth allocation, which is used entirely in all ingress-egress-multicast replication.

For example, you can configure each recycle interface with a default 50Gbps (100/2) bandwidth if you enable ingress-egress-multicast.

By default, 100Gbps is the recycle interface bandwidth, which is independent of platform or interface configuration.

Options

value Recycle interface bandwidth value in Gbps.

- **Range:** 1 through 100 Gbps
- **Default:** 100 Gbps

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS EVO Release 21.4R1.

RELATED DOCUMENTATION

Ingress and Egress Multicast Replication using Recycle Replication Model

internal-node

IN THIS SECTION

- Syntax | 9189
- Hierarchy Level | 9189
- Description | 9190
- Default | 9190
- Required Privilege Level | 9190
- Release Information | 9190

Syntax

```
internal-node;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-set interface-set-name]
```

Description

The statement is used to raise the interface set without children to the same level as the other configured interface sets with children, allowing them to compete for the same set of resources.

Default

If you do not include this statement, the node is internal only if its children have a traffic control profile configured.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Internal Scheduler Nodes

internal (Security IPsec)

IN THIS SECTION

- Syntax | 9191
- Hierarchy Level | 9191
- Description | 9191
- Options | 9191
- Required Privilege Level | 9192
- Release Information | 9192

Syntax

```
internal {
    security-association {
        manual {
            encryption {
                algorithm (3des-cbc | aes-128-cbc);
                ike-ha-link-encryption enable;
                key ascii-text;
            }
        }
    }
}
```

Hierarchy Level

```
[edit security ipsec]
```

Description

Enable secure login and to prevent attackers from gaining privileged access through this control port by configuring the internal IP security (IPsec) security association (SA).

When the internal IPsec is configured, IPsec-based `rlogin` and remote command (`rcmd`) are enforced, so an attacker cannot gain unauthorized information.

Options

security-association	Specify an IPsec SA. An SA is a simplex connection that allows two hosts to communicate with each other securely by means of IPsec.
manual encryption	Specify a manual SA. Manual SAs require no negotiation; all values, including the keys, are static and specified in the configuration.
algorithm 3des-cbc	Specify the encryption algorithm for the internal Routing-Engine-to-Routing-Engine IPsec SA configuration.
algorithm aes-128-cbc	Specify the encryption algorithm for high availability encryption link.

iked-ha-link-encryption

Enable encryption for internal messages.

- Values:
  - `enable`—Enable HA link encryption IKE internal messages

key ascii-text

Specify the encryption key. You must ensure that the manual encryption key is in ASCII text and 24 characters long; otherwise, the configuration will result in a commit failure.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

Support for `ike-ha-link-encryption` option added in Junos OS Release 12.1X47-D15.

Support for iked_encryption option added in Junos OS Release 12.1X47-D10.

Support for `aes-128-cbc` option added in Junos OS Release 19.1R1.

Support for `ike-ha-link-encryption` option added for vSRX Virtual Firewall in Junos OS Release 19.4R1

RELATED DOCUMENTATION

Example: Configuring an Active/Passive Chassis Cluster on SRX5800 Devices

show security internal-security-association | **24686**

internet-options

IN THIS SECTION

Syntax (Junos OS) | **9193**

- Syntax (Junos OS Evolved) | 9194
- Hierarchy Level | 9194
- Description | 9194
- Options | 9194
- Required Privilege Level | 9196
- Release Information | 9196

Syntax (Junos OS)

The following section lists all possible options for the `internet-options` command. The options that appear may vary depending on the platform and software release.

```
internet-options {
    (gre-path-mtu-discovery | no-gre-path-mtu-discovery);
    icmpv4-rate-limit bucket-size <bucket-size seconds> <packet-rate packet-rate>;
    icmpv6-rate-limit bucket-size <bucket-size seconds> <packet-rate packet-rate>;
    (ipip-path-mtu-discovery | no-ipip-path-mtu-discovery);
    ipv6-duplicate-addr-detection-transmits ipv6-duplicate-addr-detection-transmits;
    (ipv6-path-mtu-discovery | no-ipv6-path-mtu-discovery);
    (ipv6-reject-zero-hop-limit | no-ipv6-reject-zero-hop-limit);
    ipv6-path-mtu-discovery-timeout minutes;
    no-tcp-reset (drop-all-tcp | drop-tcp-with-syn-only);
    no-tcp-rfc1323;
    no-tcp-rfc1323-paws;
    (path-mtu-discovery | no-path-mtu-discovery);
    source-port {
        upper-limit upper-limit;
    }
    (source-quench | no-source-quench);
    tcp-drop-synfin-set;
    tcp-mss mss-value;
}
```

Syntax (Junos OS Evolved)

```
internet-options {
    ipv6-duplicate-addr-detection-transmits ipv6-duplicate-addr-detection-transmits;
    (path-mtu-discovery | no-path-mtu-discovery);
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure system IP options to protect against certain types of DoS attacks.

Options

gre-path-mtu-discovery

Configure path MTU discovery for outgoing GRE tunnel connections. By default, path MTU discovery is enabled.

- no-gre-path-mtu-discovery—Path MTU discovery is disabled.

icmpv4-rate-limit

Configure rate-limiting parameters for ICMPv4 messages sent.

- Values:
  - bucket-size *seconds*—Number of seconds in the rate-limiting bucket. Range: 0 through 4294967295 seconds. Default: 5.
  - packet-rate *pps*—Rate-limiting packets earned per second. Range: 0 through 4294967295 pps. Default: 1000.

icmpv6-rate-limit

Configure rate-limiting parameters for ICMPv6 messages sent.

- Values:
  - bucket-size *seconds*—Number of seconds in the rate-limiting bucket. Range: 0 through 4294967295 seconds. Default: 5.

- packet-rate *pps*—Rate-limiting packets earned per second. Range: 0 through 4294967295 pps. Default: 1000.

ipip-path-mtu-discovery

Configure path MTU discovery for outgoing IP-IP tunnel connections. By default, path MTU discovery is enabled.

- no-ipip-path-mtu-discovery—Path MTU discovery is disabled.

ipv6-duplicate-addr-detection-transmits

Control the number of attempts for IPv6 duplicate address detection.

- **Range:** 0 to 20
- **Default:** 3

ipv6-path-mtu-discovery

Configure path MTU discovery for IPv6 packets. By default, IPv6 path MTU discovery is enabled.

- no-ipv6-path-mtu-discovery—IPv6 path MTU discovery is disabled.

ipv6-path-mtu-discovery-timeout

Set the IPv6 path MTU discovery time-out interval.

- **Values:** *minutes*—IPv6 path MTU discovery timeout.
- **Default:** 10 minutes.

ipv6-reject-zero-hop-limit

Reject incoming IPv6 packets with a zero hop-limit value in their header. This is enabled by default.

- no-ipv6-reject-zero-hop-limit—Allow incoming IPv6 packets with a zero hop-limit value in their header.

no-tcp-reset

Do not send an RST TCP packet (a packet with the reset flag set) in response to a TCP packet received on a non-listening port.

By default, when a TCP packet is received on a non-listening port, a device sends a TCP packet with the RST flag set and drops the connection. This might lead to a security risk. Configuring this statement prevents the sending of RST TCP packets to non-listening ports.

You must configure this statement with one of two options:

- drop-all-tcp—When a TCP segment is received on a closed port, the device drops the packet and does not send back a RST segment. This helps to protect against stealth port scans.

- drop-tcp-with-syn-only—When a TCP packet with a SYN bit is received on a non-listening port, the device drops the packet and does not send back a RST segment, which makes the device appear as a null route. For all other TCP packets, the device sends back a RST segment and does not drop the packet.

no-tcp-rfc1323 Configure the Junos OS to disable RFC 1323 TCP extensions.

no-tcp-rfc1323-paws Configure the Junos OS to disable the RFC 1323 Protection Against Wrapped Sequence (PAWS) number extension.

path-mtu-discovery Configure path MTU discovery for outgoing Transmission Control Protocol (TCP) connections. By default, path MTU discovery is enabled.

- no-path-mtu-discovery—Disable MTU discovery.

source-port Configure the range of port addresses.

- Values:
  - upper-limit *upper-limit*—(Optional) The range of port addresses can be a value from 5000 through 65,355.

source-quench Configure how the Junos OS handles Internet Control Message Protocol (ICMP) source quench messages. By default, the Junos OS reacts to ICMP source quench messages.

- no-source-quench—Do not react to incoming ICMP source quench messages.

tcp-drop-synfin-set Configure the device to drop packets that have both the SYN and FIN bits set.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`no-tcp-reset` introduced in Junos OS Release 9.4.

`no-tcp-reset` introduced in Junos OS Release 11.1.

`icmpv4-rate-limit` and `source-port` introduced in Junos OS Release 11.1 and Junos OS Release 14.1X53-D20.

`path-mtu-discovery` and `no-path-mtu-discovery` introduced in Junos OS Evolved Release 23.2R2.

RELATED DOCUMENTATION

- *Configure ICMP Features*
- *Configure IPv6 Features*
- *Configure Path MTU Discovery*
- *Configure TCP Options*
- *Configuring Junos OS to Extend the Default Port Address Range*
- Understanding Traffic Processing on Security Devices

interpolate (Drop Profiles)

IN THIS SECTION

- Syntax | **9197**
- Hierarchy Level | **9198**
- Description | **9198**
- Required Privilege Level | **9198**
- Release Information | **9198**

Syntax

```
interpolate {
    drop-probability [values];
    fill-level [values];
}
```

Hierarchy Level

```
[edit class-of-service drop-profiles profile-name]
```

Description

Specify values for interpolating relationship between queue fill level and drop probability.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Defining Packet Drop Behavior by Configuring RED Drop Profiles

interval

IN THIS SECTION

- Syntax | 9199
- Hierarchy Level | 9199
- Description | 9199
- Options | 9199
- Required Privilege Level | 9200
- Release Information | 9200

Syntax

```
interval value;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options signal-degrade]
[edit interfaces interface-name otn-options odu-signal-degrade]
```

Description

Specify the interval for which the BER must stay above the signal degradation threshold—as configured in the `ber-threshold-signal-degrade value` statement—for the alarm to raised. After an alarm is raised, if the BER returns below the clear threshold—as configured in the `ber-threshold-clear value` statement—for the specified interval, the alarm is cleared.

> **NOTE**: Configuring a high BER threshold for signal degradation and a long interval might cause the internal counter register to be saturated. Such a configuration is ignored by the router, and the default values are used instead. A system log message is logged for this error.

If you configure the BER thresholds at the `[edit interfaces interface-name otn-options signal-degrade]` hierarchy level, then the thresholds are calculated using the pre-forward error correction (pre-FEC) BER (the BER before FEC correction). These thresholds are used for pre-FEC BER monitoring. See *Understanding Pre-FEC BER Monitoring and BER Thresholds* for more information about pre-FEC BER monitoring and determining BER threshold settings.

If you configure the BER thresholds at the `[edit interfaces interface-name otn-options odu-signal-degrade]` hierarchy level, then the thresholds are calculated using the post-FEC BER (the BER after FEC correction). This BER is referred to as the optical channel data unit (ODU) BER.

> **NOTE**: You can configure ODU BER thresholds only at the `[edit interfaces interface-name otn-options odu-signal-degrade]` hierarchy level on the P2-100GE-OTN PIC.

Options

value—Time interval in milliseconds.

NOTE: For the P1-PTX-2-100G-WDM PIC, the BER must stay above the signal degradation threshold for ten consecutive intervals for the alarm to be raised and the BER must stay below the clear threshold for ten consecutive intervals for the alarm to be cleared. For example, if the interval is configured as 10 ms, then the BER must stay above the signal degradation threshold for 100 ms (10 ms * 10 intervals) for the alarm to be raised, or below the clear threshold for 100 ms for the alarm to be cleared.

NOTE: For P1-PTX-24-10G-W-SFPP PIC and P2-100GE-OTN PIC, when the router cannot configure BER with the given interval, it selects an optimum interval that is supported for the given BER configuration. If the router is still not able to support the configuration (for example, with a wider gap between the degrade set and clear values), the default values are used and a log is generated.

For the P2-10G-40G-QSFPP PIC, the time interval is supported in multiples of 100 ms. For example, when you configure the interval as 10 ms, then it is rounded off to the nearest multiple of 100 ms.

- **Range:** 1 ms through 1000 ms.
- **Default:** 100 ms.

NOTE: For the P2-100GE-OTN PIC, the default value is 10 ms.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

interval (services)

IN THIS SECTION

- Syntax | 9201
- Hierarchy Level | 9201
- Description | 9201
- Options | 9201
- Required Privilege Level | 9202
- Release Information | 9202

Syntax

```
interval seconds;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then dead-peer-detection]
```

Description

Specify the amount of time that the peer waits for traffic from its destination peer before sending a dead-peer-detection (DPD) request packet. The `interval` value is used for IKEv1 security associations (SAs). Starting in Junos OS Release 17.2R1, the `interval` value is also applicable to IKEv2 SAs. In Junos OS Release 17.1 and earlier, the `interval` option is not applicable to IKEv2 SAs, which use the default value.

Options

seconds Number of seconds that the peer waits before sending a DPD request packet.

- **Range:** 1 through 180 seconds
- **Default:** 10 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

IKEv2 support introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Configuring IPsec Rules

interval (Protocols PIM)

IN THIS SECTION

- Syntax | 9202
- Hierarchy Level | 9203
- Description | 9203
- Options | 9203
- Required Privilege Level | 9203
- Release Information | 9203

Syntax

```
interval milliseconds;
```

Hierarchy Level

```
[edit protocols pim interface (Protocols PIM) interface-name multiple-triggered-joins]
```

Description

Specify the duration between the triggered joins of the PIM neighbors through the PIM interface.

Options

milliseconds—Value for the interval between the triggered joins.

- **Range:** 100 through 1000
- **Default:** 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

interface (Protocols PIM) | **8991**

multiple-triggered-joins | **11119**

interval (OAM)

IN THIS SECTION

Syntax | **9204**

- Hierarchy Level | 9204
- Description | 9204
- Options | 9205
- Required Privilege Level | 9205
- Release Information | 9205

Syntax

```
interval (100ms | 10m | 10ms | 10s | 1m | 1s| 3.3ms);
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name continuity-check]
```

Description

Configure the interval between successive transmissions of continuity check messages (CCMs) as part of the connectivity fault detection strategy. When the receiving maintenance association end point (MEP) does not receive a CCM at the configured interval, the `loss-threshold` statement determines how many CCMs can be lost before the sending MEP is marked as down. The `hold-interval` statement then determines the frequency at which the database of MEPs in the maintenance association (MA) is flushed in the absence of updates.

> **NOTE**: When the continuity-check interval is set to 3.3ms, the actual interval is 4ms (250 pps).

During a unified in-service software upgrade (ISSU), Junos OS connectivity fault management (CFM) works when the peer device is not a Juniper Networks router. Interoperating with the router of another vendor, the Juniper Networks router retains session information and continues to transmit CCM (continuity check message) PDUs during the unified ISSU upgrade. For this feature to work, you must enable Packet Forwarding Engine keepalives with the `hardware-assisted-keepalives` statement, and configure the interval between CCMs to be 1 second with `interval` statement.

NOTE: For the continuity check message interval to be configured for 10 milliseconds, periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine by default. You can disable PPM only on the Packet Forwarding Engine. To disable PPM on the Packet Forwarding Engine, use the `no-delegate-processing` statement at the `[edit routing-options ppm]` hierarchy level.

NOTE: A continuity check interval of 10 milliseconds is not supported for CFM sessions over a label-switched interface (LSI).

Options

`100ms`—100 milliseconds.

`10m`—10 minutes.

`10ms`—10 milliseconds.

`10s`—10 seconds.

`1m`—1 minute.

`1s`—1 second.

`3.3ms`—3.3 milliseconds.

- **Default:** 1m

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Option `10ms` introduced in Junos OS Release 9.1.

Third-party interoperability during a unified in-service software upgrade (ISSU) introduced in Junos OS Release 17.1.

Option `3.3ms` introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Continuity Check Protocol Parameters Overview

Configuring Continuity Check Protocol Parameters for Fault Detection

Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades

interval (RMON)

IN THIS SECTION

- Syntax | 9206
- Hierarchy Level | 9206
- Description | 9206
- Options | 9207
- Required Privilege Level | 9207
- Release Information | 9207

Syntax

```
interval seconds;
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Interval between samples.

Options

seconds—Time between samples, in seconds.

- **Range:** 1 through 2,147,483,647 seconds
- **Default:** 60 seconds

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Interval

interval (Accounting Options)

IN THIS SECTION

- Syntax | 9208
- Hierarchy Level | 9208
- Description | 9208
- Options | 9208
- Required Privilege Level | 9208
- Release Information | 9208

Syntax

```
interval minutes;
```

Hierarchy Level

```
[edit accounting-options class-usage-profile profile-name],
[edit accounting-options filter-profile profile-name],
[edit accounting-options interface-profile profile-name],
[edit accounting-options mib-profile profile-name],
[edit accounting-options routing-engine-profile profile-name]
```

Description

Specify how often statistics are collected for the accounting profile.

Options

minutes—Length of time between each collection of statistics.

- **Range:** 1 through 2880 minutes
- **Default:** 30 minutes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The `[edit accounting-options mib-profile profile-name]` hierarchy level added in Junos OS Release 8.2.

RELATED DOCUMENTATION

> *Configuring the Interface Profile*

Configuring the Filter Profile

Configuring the MIB Profile

Configuring the Routing Engine Profile

interval (CFM MEP)

IN THIS SECTION

- Syntax | 9209
- Hierarchy Level | 9209
- Description | 9209
- Required Privilege Level | 9209
- Release Information | 9210

Syntax

```
interval (1m | 1s);
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
action log-and-generate-ais]
```

Description

Configure the interval between AIS messages that are to be received by the MEP as either 1 minute or 1 second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Alarm Indication Signal (ETH-AIS) Function Overview

Configuring ETH-AIS on a CFM MEP

interval (EX Series Switch Only)

IN THIS SECTION

- Syntax | 9210
- Hierarchy Level | 9210
- Description | 9211
- Options | 9211
- Required Privilege Level | 9211
- Release Information | 9211

Syntax

```
interval (10m | 10s | 1m | 1s | 100ms | 10ms);
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name continuity-check]
```

Description

Configure the time between continuity check messages.

Options

10m—10 minutes.

10s—10 seconds.

1m—1 minute.

1s—1 second.

100ms—100 milliseconds.

10ms—10 milliseconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)

Junos OS Network Interfaces Configuration Guide

interval (Flat-File Accounting Options)

IN THIS SECTION

Syntax | 9212

- Hierarchy Level | 9212
- Description | 9212
- Options | 9212
- Required Privilege Level | 9213
- Release Information | 9213

Syntax

```
interval minutes;
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name]
```

Description

Specify the interval in minutes at which the Packet Forwarding Engine associated with the interface is polled to collect the statistics specified in the flat-file accounting profile. These interim accounting results are recorded in the flat file.

> **NOTE**: The value configured with this statement is superseded by the value configured with the "update-interval" on page 30444 statement at the `[edit access profile profile-name service accounting]` hierarchy level. That access profile interval value is in turn superseded by an update interval value configured in the RADIUS attribute, Service-Interim-Acct-Interval (VSA 26–140).

Options

minutes	Polling interval.

- **Range:** 1 through 2880 minutes
- **Default:** 15 minutes

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Configuring Service Accounting in Local Flat Files

Flat-File Accounting Overview

interval (Health Monitor)

IN THIS SECTION

- Syntax | 9213
- Hierarchy Level | 9214
- Description | 9214
- Options | 9214
- Required Privilege Level | 9214
- Release Information | 9214

Syntax

```
interval seconds;
```

Hierarchy Level

```
[edit snmp health-monitor],
[edit snmp health-monitor idp]
```

Description

Interval between samples.

Options

seconds—Time between samples, in seconds.

- **Range:** `1` through `2147483647` seconds
- **Default:** `300` seconds

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring the Interval

interval (Interfaces)

IN THIS SECTION

Syntax | 9215

- Hierarchy Level | 9215
- Description | 9215
- Options | 9215
- Required Privilege Level | 9215
- Release Information | 9215

Syntax

```
interval seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unitlogical-unit-number radio-router credit]
```

Description

Configure the frequency that the router generates credit announcement messages.

Options

seconds—Interval between PADG credit announcements for each session.

- **Range:** 0 through 60
- **Default:** 1

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.1 of Junos OS.

RELATED DOCUMENTATION

Configuring PPPoE-Based Radio-to-Router Protocols

interval (PoE)

IN THIS SECTION

- Syntax | 9216
- Hierarchy Level | 9216
- Description | 9216
- Options | 9216
- Required Privilege Level | 9217
- Release Information | 9217

Syntax

```
interval minutes;
```

Hierarchy Level

```
[edit poe interface (all | interface-name) telemetries]
```

Description

Modifies the interval for logging telemetries if you are monitoring the per-port power consumption for PoE interfaces.

Options

minutes—Interval at which data is logged.

- **Range:** 1 through 30 minutes

- **Default:** 5 minutes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Interfaces

interval (Power over Ethernet)

IN THIS SECTION

- Syntax | 9217
- Hierarchy Level | 9218
- Description | 9218
- Options | 9218
- Required Privilege Level | 9218
- Release Information | 9218

Syntax

```
interval minutes;
```

Hierarchy Level

```
[edit poe interface (all | interface-name) telemetries]
```

Description

Modify the interval at which data is collected when you are monitoring the power consumption of a PoE interface.

Options

minutes—Frequency of data collection.

- **Range:** 1 through 30
- **Default:** 5

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

> *Configuring PoE Interfaces on EX Series Switches*
>
> Configuring PoE (J-Web Procedure)

interval (Security Antivirus)

IN THIS SECTION

- Syntax | 9219
- Hierarchy Level | 9219
- Description | 9219
- Options | 9220
- Required Privilege Level | 9220
- Release Information | 9220

Syntax

```
interval value;
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update]
[edit security utm feature-profile anti-virus sophos-engine pattern-update]
[edit security utm default-configuration anti-virus avira-engine pattern-update]
```

Description

Set the pattern data files auto-update interval. You can choose to leave the default interval value or you can change it by using this command. You can also force a manual update, if necessary.

NOTE: The data files used with Sophos are not typical virus pattern files; they are small files that help guide virus scanning logic. The full virus pattern database is stored on an external Sophos server called the Sophos Extensible List (SXL) server.

Options

value—Pattern data files auto-update interval in minutes.

- **Range:** 10 through 10,080 minutes (10 minutes through 7 days)
- **Default:** For Juniper Express engine and Kaspersky Lab engine, 60 minutes; for Sophos engine, 1440 minutes (every 24 hours)

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5 for Juniper Express engine and Kaspersky Lab engine. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for Avira engine added in Junos OS Release 18.4R1.

interval (Security IDP)

IN THIS SECTION

- Syntax | 9221
- Hierarchy Level | 9221
- Description | 9221
- Options | 9221
- Required Privilege Level | 9221
- Release Information | 9221

Syntax

```
interval hours;
```

Hierarchy Level

```
[edit security idp security-package automatic]
```

Description

Specify the amount of time that the device waits before updating the signature database. User should insert a default value.

Options

hours—Number of hours that the device waits.

- **Range:** 24 through 336 hours

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

interval (Services)

IN THIS SECTION

- Syntax | 9222
- Hierarchy Level | 9222

- Description | 9222
- Required Privilege Level | 9222
- Release Information | 9222

Syntax

```
interval seconds;
```

Hierarchy Level

```
[edit services unified-access-control ]
```

Description

Specify the value in seconds that the SRX Series Firewall should expect to receive a heartbeat signal from the IC Series device (default 30). This configuration statement is used in conjunction with the `timeout` statement to test active communications with the IC Series device. The value of the `interval` statement must be smaller than the value of `timeout` statement.

Use this statement when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

timeout (Services) | 29395

timeout-action | 29371

interval (SNMP RMON)

IN THIS SECTION

- Syntax | 9223
- Hierarchy Level | 9223
- Description | 9223
- Default | 9224
- Options | 9224
- Required Privilege Level | 9224
- Release Information | 9224

Syntax

```
interval seconds;
```

Hierarchy Level

```
[edit snmp rmon alarm index],
[edit snmp rmon history]
```

Description

Configure the interval over which data is to be sampled for the specified interface.

Default

60 sec for alarm sampling.

1800 sec for Ethernet packet history sampling.

Options

seconds—Interval at which data is to be sampled for the specified interface.

Required Privilege Level

snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Understanding RMON

RMON MIB Event, Alarm, Log, and History Control Tables

Configuring RMON History Sampling

Monitoring RMON MIB Tables

intra-area-prefix

IN THIS SECTION

- Syntax | 9225
- Hierarchy Level | 9225
- Description | 9225
- Required Privilege Level | 9225
- Release Information | 9225

Syntax

```
intra-area-prefix;
```

Hierarchy Level

```
[edit logical-systems name protocols ospf3 overload],
[edit protocols ospf3 overload]
```

Description

Advertise intra-area Prefix with maximum metric.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2.

RELATED DOCUMENTATION

Understanding OSPF Overload Function

allow-route-leaking | 2202

stub-network | 28667

as-external | 2486

intra-as

IN THIS SECTION

- Syntax | 9226
- Hierarchy Level | 9226
- Description | 9226
- Required Privilege Level | 9226
- Release Information | 9227

Syntax

```
intra-as {
    inclusive;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvpn
family (inet / inet6) autodiscovery-only],
[edit routing-instances routing-instance-name protocols mvpn family (inet / inet6) autodiscovery-
only]
```

Description

For Rosen 7, enable the MVPN control plane for autodiscovery only, using intra-AS autodiscovery routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Statement moved to `[..protocols mvpn family inet]` from `[.. protocols mvpn]` in Junos OS Release 13.3.

Support for IPv6 added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs

invalid-authentication-entry-timeout (Services User Identification Active Directory and ClearPass)

IN THIS SECTION

- Syntax | 9227
- Hierarchy Level | 9228
- Description | 9228
- Options | 9229
- Required Privilege Level | 9229
- Release Information | 9229

Syntax

```
invalid-authentication-entry-timeout timeout-value-in-minutes;
```

Hierarchy Level

```
[edit services user-identification active-directory-access]
[edit services user-identification authentication-source (Services User Identification ClearPass)
aruba-clearpass]
```

Description

Configure an independent timeout value to be assigned to invalid user authentication entries in the SRX Series authentication table for either Windows active directory or Aruba ClearPass. The invalid authentication entry timeout setting is different from the general authentication entry timeout setting. It allows you to protect invalid user authentication entries in an authentication table from expiring before the user can be validated.

User authentication entries in an authentication table contain a time-out value after which the entry expires, or is no longer valid. An invalid authentication entry is created with a NULL and INVALID state for a user's IP address and stored in the access directory authentication table when there is no identity information for that user. Prior to implementation of this feature, the current time-out value that applies to all user entries was applied to the invalid entry also.

Separate authentication tables exist for the two authentication sources and you configure separate settings for them, as illustrated in the following examples.

Use the following command to configure the invalid authentication entry timeout for entries in the Windows active directory authentication table. In this example, invalid authentication entries in the active directory authentication table will expire 40 minutes after they were created.

```
user@host# set services user-identification active-directory-access invalid-authentication-entry-
timeout 40
```

Use the following command to configure the invalid authentication entry timeout for entries in the SRX Series ClearPass authentication table. In this example, invalid authentication entries in the SRX Series ClearPass authentication table will expire 22 minutes after they were created.

```
user@host# set services user-identification authentication-source aruba-clearpass invalid-
authentication-entry-timeout 22
```

The following rules govern how the invalid authentication entry timeout setting is used:

- When you initially configure the invalid authentication entry timeout value, it is applied to any invalid authentication entries that are created *after* it was configured.

However, all existing invalid authentication entries retain the default timeout of 30 minutes.

- If you do not configure the invalid authentication entry timeout function, then the default timeout of 30 minutes is applied to all invalid authentication entries.
- If you configure the invalid authentication entry timeout value but later you delete it, the default timeout of 30 minutes is applied to any invalid authentication entries created *after* the deletion.

  However, any invalid authentication entries to which the invalid entry timeout value was applied *before* the deletion retain that setting.
- If you change the setting for the invalid authentication entry timeout value, the new value is applied to all invalid authentication entries that were created *after* the value was changed. However, all existing invalid authentication entries retain the former invalid authentication entry timeout setting, if it applied to them. Those to which the default value of 30 minutes applies retain that setting.
- When the state of an invalid authentication entry changes to Pending or Valid, the invalid authentication entry timeout setting is no longer applicable to it. Therefore, the timeout value assigned to that entry is changed to the value that is set for the general authentication entry timeout.

Options

timeout-value-in-minutes

Expiration time in minutes to be applied to invalid authentication entries in the SRX Series authentication table for either Windows active directory or Aruba ClearPass authentication sources.

- **Range:** 0 through 1440 minutes.
- **Default:** 30 minutes

Required Privilege Level

1. services—To view this statement in the configuration.
2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Understanding the Invalid Authentication Table Entry Timeout Setting

Understanding the Forced Timeout Setting Assigned to Active Directory Authentication Entries for Users Authenticated Through Captive Portal

firewall-authentication-forced-timeout

ioc-npc-connectivity

IN THIS SECTION

- Syntax | 9230
- Hierarchy Level | 9230
- Description | 9230
- Required Privilege Level | 9231
- Release Information | 9231

Syntax

```
ioc-npc-connectivity {
    ioc name {
        npc(none |  npc-slot);
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Connectivity between IOC and NPC

Required Privilege Level

[none specified]

Release Information

Statement introduced in Junos OS Release 10.4R1.

ip-action (Security IDP Rulebase IPS)

IN THIS SECTION

- Syntax | 9231
- Hierarchy Level | 9232
- Description | 9232
- Options | 9232
- Required Privilege Level | 9233
- Release Information | 9233

Syntax

```
ip-action {
    (ip-block | ip-close | ip-notify | ip-connection-rate-limit);
    log;
    log-create;
    refresh-timeout;
    target (destination-address (Security IDP Policy) | service | source-address | source-zone |
source-zone-address | zone-service);
    timeout seconds;
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name then]
```

Description

Specify the actions you want IDP to take against future connections that use the same IP address.

Options

ip-block Block future connections of any session that matches the IP action. If there is an IP action match with multiple rules, then the most severe IP action of all the matched rules is applied. The highest IP action priority (that is, the most severe action) is Drop/Block, then Close, then Notify.

ip-close Close future connections of any new sessions that match the IP action by sending RST packets to the client and server.

ip-notify Do not take any action against future traffic, but do log the event.

ip-connection-rate-limit When a match is made in a rulebase-ddos rule you can set the **then** action to ip-connection-rate-limit, which will limit the rate of future connections based on a connections per second limit that you set. This can be used to reduce the number of attacks from a client.

- **Syntax:** *value*—Defines the connection rate limit per second on the matched host.
- **Range:** 1 to the maximum connections per second capability of the device.

log Log the information about the IP action against the traffic that matches a rule.

log-create Generate a log event on installing the ip-action filter.

refresh-timeout Refresh the ip-action timeout so it does not expire when future connections match the installed ip-action filter.

target Specify the blocking options that you want to set to block the future connections. Blocking options can be based on the following matches of the attack traffic:

- Range:

- `destination-address`—Matches traffic based on the destination address of the attack traffic.
- `service`—For TCP and UDP, matches traffic based on the source address, source port, destination address, and destination port of the attack traffic. This is the default.

  For ICMP flows, the destination port is 0. Any ICMP flow matching source port, source address, and destination address is blocked.
- `source-address`—Matches traffic based on the source address of the attack traffic.
- `source-zone`—Matches traffic based on the source zone of the attack traffic.
- `source-zone-address`—Matches traffic based on the source zone and source address of the attack traffic.
- `zone-service`—Matches traffic based on the source zone, destination address, destination port, and protocol of the attack traffic.

timeout

Specify the number of seconds that you want the IP action to remain in effect after a traffic match.

- **Syntax:** *seconds*—Number of seconds the IP action should remain effective.
- **Range:** 0 through 64,800 seconds
- **Default:** 0 second

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Option `log-create` and `refresh-timeout`, and `ip-connection-rate-limit` introduced in Junos OS Release 10.2.

NOTE: For ICMP flows, the destination port is 0; therefore, any ICMP flow matching source port, source address, and destination address is blocked.

ip-address (Access Profile)

IN THIS SECTION

- Syntax | 9234
- Hierarchy Level | 9234
- Description | 9234
- Required Privilege Level | 9234
- Release Information | 9234

Syntax

```
ip-address address;
```

Hierarchy Level

```
[edit access profile name client name xauth]
```

Description

Specify the IP address for the client.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.4 of Junos OS.

RELATED DOCUMENTATION

Firewall User Authentication Overview

ip-address-first

IN THIS SECTION

- Syntax | 9235
- Hierarchy Level | 9235
- Description | 9235
- Required Privilege Level | 9236
- Release Information | 9236

Syntax

```
ip-address-first;
```

Hierarchy Level

```
[edit logical-systems logical-system-name system services dhcp-local-server pool-match-order],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server pool-match-order],
[edit routing-instances routing-instance-name system services dhcp-local-server pool-match-order],
[edit system services dhcp-local-server pool-match-order]
```

Description

Configure the extended DHCP local server to use the IP address method to determine which address-assignment pool to use. The local server uses the IP address in the gateway IP address if one is present in the DHCP client PDU. If no gateway IP address is present, the local server uses the IP address of the receiving interface to find the address-assignment pool. The DHCP local server uses this method by default when no method is explicitly specified.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use

Understanding Differences Between Legacy DHCP and Extended DHCP

Address-Assignment Pools Overview

ip-address-owner

IN THIS SECTION

- Syntax | 9236
- Hierarchy Level | 9237
- Description | 9237
- Options | 9237
- Required Privilege Level | 9237
- Release Information | 9237

Syntax

```
ip-address-owner owner;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Define the owner for IP addresses hosted on an ms- interface. This statement is used to specify that the steering of control plane packets to the Multiservices PIC be preserved. The provider of your application or its user documentation will tell you when you need to set this configuration.

Options

owner Owner of the IP address for the interface. There are two options:

- Values:

  routing-engine IP address defined on the interface is hosted by the Routing Engine. This option is not used.

  service-plane IP address defined on the interface is hosted by the service plane. This option is used to preserve the packet steering behavior built into your application. The provider of your application or its user documentation will tell you when you need to set this configuration.

Required Privilege Level

view

Release Information

Statement introduced in Junos OS Release 13.2R1.

ip-address (Port Mirroring)

IN THIS SECTION

- Syntax | 9238
- Hierarchy Level | 9238
- Description | 9238
- Required Privilege Level | 9239
- Release Information | 9239

Syntax

```
ip-address ip-address;
```

Hierarchy Level

```
[edit forwarding-options] analyzer name output]
[edit forwarding-options port-mirroring [instance name] family ethernet-switching output interface
name]
```

Description

Specify the IP address to which traffic should be mirrored (the IP address of the analyzer system). The device can be on a remote network. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.) This statement is not supported on QFabric systems.

NOTE: An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.

> **NOTE**: If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53.

ip-can-type (PCRF Partition)

IN THIS SECTION

- Syntax | 9239
- Hierarchy Level | 9240
- Description | 9240
- Options | 9240
- Required Privilege Level | 9240
- Release Information | 9240

Syntax

```
ip-can-type number;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure an IP connectivity access network (IP-CAN) value according to what best fits your operating environment and access network. The Policy and Charging Rules Function (PCRF) partition requires that you configure this statement. The Policy Control and Charging (PCC) enables a centralized control to ensure that the service sessions (IP-CAN sessions) are provided with appropriate bandwidth and QoS.

An IP-CAN bearer is the IP transmission path of defined capacity, delay, and bit error rate. An IP-CAN session incorporates one or more IP-CAN bearers. Support for multiple IP-CAN bearers per IP-CAN session is IP-CAN specific. An IP-CAN session exists as long as the related IPv4 address or IPv6 prefix is assigned and announced to the IP network.

If an IP-CAN session is modified, the Policy and Charging Enforcement Function (PCEF) first uses the event trigger to determine whether to request the PCC rules for the modified IP-CAN session from the PCRF. Then upon reception of updated PCC rules from the PCRF, the PCEF activates, modifies, or removes the PCC rules as indicated by the PCRF.

Options

number Identifier of the IP-CAN value used for your operating environment and access network.

- **Default:** 77

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

Understanding Upstream and Downstream Messages for the PCRF

ip (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 9241
- Hierarchy Level | 9242
- Description | 9242
- Required Privilege Level | 9242
- Release Information | 9242

Syntax

```
ip {
    bad-option;
    block-frag;
    ipv6-extension-header {
        AH-header;
        ESP-header;
        fragment-header;
        hop-by-hop-header {
            CALIPSO-option;
            jumbo-payload-option;
            quick-start-option;
            router-alert-option;
            RPL-option;
            SFM-DPD-option;
            user-defined-option-type <type-low> to <type-high>;
        }
        mobility-header;
```

```
        routing-header;
    }
    loose-source-route-option;
    record-route-option;
    security-option;
    source-route-option;
    stream-option;
    strict-source-route-option;
    tear-drop;
    timestamp-option;
    unknown-protocol;
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name]
```

Description

Configure protection against suspicious IP packet attacks.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

ip-in-udp

IN THIS SECTION

- Syntax | 9243
- Hierarchy Level | 9243
- Description | 9243
- Options | 9243
- Required Privilege Level | 9243
- Release Information | 9244

Syntax

```
ip-in-udp {
    teredo;
}
```

Hierarchy Level

```
[edit security screen ids-option ids-option-name ip tunnel]
```

Description

Configure IP tunnel IPinUDP IDS option.

Options

teredo Enable IP tunnel IPinUDP Teredo IDS option.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D10.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

ip-monitoring

IN THIS SECTION

- Syntax | 9244
- Hierarchy Level | 9245
- Description | 9245
- Options | 9245
- Required Privilege Level | 9246
- Release Information | 9246

Syntax

```
ip-monitoring {
    family {
        inet {
            ipv4-address {
                interface {
                    logical-interface-name;
                    secondary-ip-address ip-address;
                }
                weight number;
            }
        }
```

```
    }
    global-threshold number;
    global-weight number;
    retry-count number;
    retry-interval seconds;
}
```

Hierarchy Level

```
[edit chassis cluster redundancy-group group-number]
```

Description

Specify a global IP address monitoring threshold and weight, and the interval between pings (`retry-interval`) and the number of consecutive ping failures (`retry-count`) permitted before an IP address is considered unreachable for all IP addresses monitored by the redundancy group. Also specify IP addresses, a monitoring weight, a redundant Ethernet interface number, and a secondary IP monitoring ping source for each IP address, for the redundancy group to monitor.

Options

IPv4 address The address to be continually monitored for reachability. You also set up a secondary IP address to allow testing from the secondary node.

> **NOTE**: All monitored object failures, including IP monitoring, are deducted from the redundancy group threshold priority. Other monitored objects include interface monitor, SPU monitor, cold-sync monitor, and NPC monitor (on supported platforms).

interface *interface* secondary-ip-address *ip-address*; Define source address for monitoring packets on secondary link.

global-threshold Define global threshold for IP monitoring. See "global-threshold" on page 7740.

- **Default:** `0`
- **Range:** `0-255`

global-weight Define global weight for IP monitoring. See "global-weight" on page 7743.

- **Default:** 255
- **Range:** 0-255

retry-count Number of retries needed to declare reachablity failure. See "retry-count (Chassis Cluster)" on page 15700.

- **Default:** 5
- **Range:** 5-15

retry-interval Define the time interval in seconds between retries. See "retry-interval (Chassis Cluster)" on page 15707.

- **Default:** 1
- **Range:** 1-30

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement updated in Junos OS Release 10.1.

RELATED DOCUMENTATION

Example: Configuring Chassis Cluster Redundancy Group IP Address Monitoring

ip-monitoring (Services)

IN THIS SECTION

Syntax | 9247

- Hierarchy Level | 9248
- Description | 9248
- Options | 9248
- Required Privilege Level | 9248
- Release Information | 9248

Syntax

```
ip-monitoring {
    policy policy-name {
        match {
            rpm-probe [probe-name];
        }
        no-preempt ;
        then {
            interface interface-name (disable | enable);
            preferred-route {
                withdraw;
                route destination-address {
                    direct-next-hop
                    discard
                    next hop next-hop;
                    preferred-metric metric;
                    }
                routing-instances name;
            }
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
```

```
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure IP monitoring.

Options

withdraw—Remove the primary routes when RPM fails to probe the destination. So, the traffic can choose other routes in the route table. If no other route exists, then the traffic is dropped. When you run the `show services ip-monitoring status` command, in the output, the status of the route `WITHDRAWED` indicates that this route is removed. See "show services ip-monitoring status" on page 25604 for more details.

`direct-next-hop`-This configuration requests the next hop, which must be direct. This configuration ignores the other kind of hops.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support added for the `discard` option starting from Junos OS Release 15.1X49-D60.

Support added for the `withdraw` option starting from Junos OS Release 20.3R1.

`direct-next-hop` is introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

icmp | 8253

ip-options-protocol-queue

IN THIS SECTION

- Syntax | 9249
- Hierarchy Level | 9249
- Description | 9249
- Options | 9250
- Required Privilege Level | 9250
- Release Information | 9250

Syntax

```
ip-options-protocol-queue protocol-name {
    protocol-id protocol-id;
    queue-depth queue-depth;
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure logical queue-depth in the PFE for `ip-options` packets for a given protocol such as TCP, UDP, ICMP, and so on, except IGMP. The queue-depth indicates the number of `ip-options` packets that can be enqueued in the PFE logical queue, beyond which it will start dropping the packets. Currently, IGMP has a default queue-depth of 192 (which is not configurable), and other protocols have a cumulative default

queue-depth of 25. The CLI supports configuration for a maximum of 16 protocols. The sum total of queue-depth for all the protocols should not exceed 1024 packets.

Options

protocol-id *protocol-id*

Identify the protocol.

- **Range:** 1 through 254

queue-depth *queue-depth*

Size of the protocol logical options queue for a given protocol.

- **Range:** 1 through 807 packets

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

ip-prefix-routes

IN THIS SECTION

- Syntax | 9251
- Hierarchy Level | 9251
- Description | 9251
- Options | 9253
- Required Privilege Level | 9255
- Release Information | 9255

Syntax

```
ip-prefix-routes {
    advertise (direct-nexthop | gateway-address);
    encapsulation (mpls | srv6 | vxlan);
    reject-asymmetric-vni;
    donot-advertise-community;
    <export routing-policy-name>;
    gateway-interface interface-name;
    source-packet-routing {
        srv6 {
            locator loc-name {
                (end-dt4/6/46-sid | micro-dt4/6/46-sid);
            }
        }
    }

    vni number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
evpn],
[edit routing-instances routing-instance-name protocols evpn]
```

Description

In an Ethernet VPN (EVPN) environment, enable the device to advertise the IP prefix associated with a specified customer domain as a Type 5 route to remote data centers. You might also enable IP prefix routes in a metro transport network. You use this feature when the Layer 2 domain does not exist at the remote data centers or metro network peering points.

We support two models for implementing Type 5 routes:

- Pure Type 5 route without overlay next hop and Type 2 route
- Type 5 route with gateway IRB interface as overlay next hop and Type 2 route

A pure Type 5 route advertises the summary IP prefix and includes a BGP extended community called a router MAC. The router MAC carries the MAC address of the sending switch and provides next-hop

reachability information for the prefix. In contrast, a standard Type 5 route requires a gateway IP address as a next-hop overlay and a supporting Type 2 route to provide recursive route resolution.

QFX5110 and QFX10000 switches currently support only fully resolved next-hop—that is, EVPN pure Type 5—routes. QFX10000 switches support only EVPN Virtual Extensible LAN (VXLAN) with Type 5 routes (not MPLS encapsulation).

On QFX5110, QFX5120, EX4650, EX4400, and EX4100 switches, when you configure pure Type 5 routing in an overlay EVPN-VXLAN network, you must also configure the "overlay-ecmp" on page 12041 statement at the `[edit forwarding-options vxlan-routing]` hierarchy level. Configuring the `overlay-ecmp` statement causes the Packet Forwarding Engine to restart, which interrupts all forwarding operations. As a result, we strongly recommend you put that configuration in place before the EVPN-VXLAN network becomes operational.

EX9200 switches support both Type 5 routes but only pure Type 5 routes with MPLS encapsulation. EX9200 switches support both MPLS and VXLAN encapsulation with the standard Type 5 route that includes a gateway IP address as the next-hop overlay.

CAUTION: We introduced pure Type 5 routing for EVPN-VXLAN in Junos OS Release 15.1X53-D30 for QFX10002 switches only. In that release, this statement is `ip-prefix-support forwarding-mode symmetric`.

Starting with Junos OS Release 15.1X53-D60, the statement is `ip-prefix-routes advertise direct-nexthop`.

When you upgrade to Junos OS Release 15.1X53-D60 or later, the upgrade process automatically updates any configuration with the original `ip-prefix-support` statement to the new `ip-prefix-routes` statement.

NOTE: Starting in Junos OS Release 15.1X53-D60, QFX10000 switches support pure Type 5 routing.

By default, Juniper Networks devices that support pure EVPN Type 5 routes don't advertise IP prefixes with a mask length of /32 as pure Type 5 routes. To advertise /32 prefixes, you must set up a term in an export policy that explicitly accepts these prefixes into the routing table. The following are excerpts from a sample configuration:

```
user@qfx10002# show policy-options policy-statement slash32
term 1 {
    from {
        protocol bgp;
        route-filter 10.44.44.44/31 orlonger;
    }
```

```
    then accept;
}
user@qfx10002# show routing-instances EVPN-A protocols evpn
ip-prefix-routes {
    advertise direct-nexthop;
    encapsulation vxlan;
    vni 5000;
    export slash32;
...
```

Options

The `advertise`, `encapsulation` and `vni` *`number`* options are required. The `export` *`routing-policy-name`* option is optional.

advertise direct-nexthop

Enable the switch to send IP prefix information using an EVPN pure Type 5 route, which includes a router MAC extended community used to send the MAC address of the switch. This router MAC extended community provides next-hop reachability without requiring an overlay next-hop or supporting Type 2 route.

> **NOTE**: For pure route Type 5, QFX5110 and QFX10000 switches support only VXLAN encapsulation, and EX9200 switches support only MPLS encapsulation.

advertise gateway-address (EX9200 switches only)

Enable the switch to advertise a gateway address in exported IP prefix routes. This gateway address provides overlay next-hop reachability.

> **NOTE**: You must also specify a gateway address by including the `gateway-interface` *`interface-name`* statement.

encapsulation (mpls | srv6 | vxlan)

Specify the encapsulation for IP prefixes: MPLS, SRv6, or VXLAN.

> **NOTE**: The same type of encapsulation must be used end to end. Only VXLAN encapsulation is supported on the QFX10000 line of switches and QFX5110 switches.

reject-asymmetric-vni

Enables the local Junos device to block asymmetric EVPN type 5 traffic. The device examines the incoming EVPN Type 5 route packets and rejects the route when the VNI in the ingress route is different from the locally configured VNI.

> **NOTE**: This option is required when you configure EVPN Type 5 routing on ACX7100 routers, which don't support Type 5 route advertisement with asymmetric VNIs.

donot-advertise-community

Starting with Junos release 19.4R1, for when a Type 5 routes is created from the Type 2 MAC and IP advertisement that was learned from a remote PE device, the default behavior is to include the community information from the Type 2 route (if any) in the Type 5 route. You can prevent behavior this by enabling the `donot-advertise-community` option.

export *routing-policy-name*

(Optional) Specify the name of the routing policy configured at the `[edit policy-options policy-statement policy-statement-name]` hierarchy level to apply to the routes for the specified customer domain. Applying an export policy allows you to further control the IP prefixes to advertise or to suppress through EVPN Type 5 routes for each customer. You can apply a separate export routing policy to one or more customer domains. This allows each customer to each have its own policy.

gateway-interface *interface-name*

Specify the gateway interface to use as a next-hop overlay for a standard Type 5 route. You must use this option in conjunction with the `advertise gateway-address` option.

source-packet-routing

Enable SRv6 source packet routing, identify the SRv6 locator defined at `[edit routing-options source-packet-routing srv6]`, and identify the type of SRv6 segment identifier (SID).

Possible SRv6 SIDs:

- `end-dt4-sid`
- `end-dt46-sid`
- `end-dt6-sid`
- `micro-dt4-sid`
- `micro-dt46-sid`
- `micro-dt6-sid`

vni *number*

Specify the identifier associated with a customer domain. Each customer domain must have a unique identifier.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D60 and Junos OS Release 17.2R1.

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

`reject-asymmetric-vni` option introduced in Junos OS Release 22.2R1 and Junos OS Evolved Release 22.1R1

Support for SRX and vSRX Virtual Firewall Series firewalls added in Junos OS Release 22.4R1.

source-packet-routing option introduced in Junos OS Release 23.4R1 and Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

- *Understanding EVPN Pure Type-5 Routes*
- *EVPN Type 2 and Type 5 Route Coexistence with EVPN-VXLAN*
- *EVPN Overview for Switches*
- policy-statement

ip-prefix-support

IN THIS SECTION

- Syntax | 9256
- Hierarchy Level | 9256
- Description | 9256
- Options | 9257
- Required Privilege Level | 9257
- Release Information | 9257

Syntax

```
ip-prefix-support {
    encapsulation vxlan;
    <export routing-policy-name>;
    forwarding-mode symmetric;
    vni number;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols evpn]
```

Description

In an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) environment, enable the provider edge (PE) switch to forward the IP prefix associated with a specified customer domain in scenarios where the Layer 2 domain does not extend across data centers. Such routes are referred to as EVPN pure type-5 routes. On QFX switches, only fully resolved next-hop routes are supported. The data packets are sent as Layer 2 frames that are encapsulated in the VXLAN header. A unique VXLAN numerical identifier is associated with each customer domain.

When the advertisement of the type-5 route is enabled, you can optionally use an export routing policy configured at the `[edit policy-options policy-statement policy-name]` hierarchy level to further control the IP prefixes to advertise or suppress through EVPN type-5 routes for each customer. The policy is applied to the routes that reside in the customer's inet.0 routing table by specifying a *routing-instance-name* configured at the `[edit routing-instances]` hierarchy level. This *routing-instance-name* refers to a Layer 3 virtual routing and forwarding (VRF) domain for a specific customer. The corresponding name of the table is *instance.name*`.inet.0` This table is created by default when you configure a routing instance.

> **CAUTION**: This statement is deprecated starting with Junos OS Release 15.1x53-D60 and has been replaced with `ip-prefix-routes`. The statement `ip-prefix-support forwarding-mode symmetric` is now `ip-prefix-routes advertise direct-nexthop`. Any configuration with the original `ip-prefix-support` statement is automatically upgraded to the new `ip-prefix-routes` statement when you upgrade to Junos OS Release 15.1x53-D60 or later.

Options

The `forwarding-mode symmetric`, `encapsulation vxlan` and `vni` *`number`* options are required. The `export` *`policy-name`* option is optional.

encapsulation vxlan
Specify to encapsulate forwarded traffic in VXLAN for transmission to the other data center.

export *routing-policy-name*
(Optional) Specify the name of the routing policy configured at the `[edit policy-options policy-statement` *`policy-name`*`]` hierarchy level to apply to the routes for the specified customer domain. Applying an export policy allows you to further control the IP prefixes to advertise or to suppress for each customer. You can apply a separate export routing policy to one or more customer domains. This allows each customer to each have its own policy.

forwarding-mode symmetric
Enable the receiver to resolve the next-hop route using only information carried in the type-5 route.

vni *number*
Specify the identifier associated with a customer domain. Each customer domain must have a unique identifier.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1x53-D30.

RELATED DOCUMENTATION

Understanding EVPN Pure Type-5 Routes

policy-statement

ip-protocol-mapping (Application Identification)

IN THIS SECTION

- Syntax | 9258
- Hierarchy Level | 9258
- Description | 9258
- Options | 9258
- Required Privilege Level | 9259
- Release Information | 9259

Syntax

```
ip-protocol-mapping {
    order order;
    order-priority (high | low);
    protocol  (http | ssl | tcp | udp)
}
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Description

For IP traffic, identify an application by matching the IP protocol. This parameter is used to identify an application based on IP and is intended only for IP traffic.

Options

protocol-number Industry-standard numeric protocol value.

- **Range:** 0 through 254

You can find a complete list of industry standard protocol numbers at the IANA website.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

ip-query (Identity Management Advanced Query)

IN THIS SECTION

- Syntax | 9260
- Hierarchy Level | 9260
- Description | 9260
- Options | 9261
- Required Privilege Level | 9261
- Release Information | 9261

Syntax

```
ip-query {
    no-ip-query;
    query-delay-time seconds;
}
```

Hierarchy Level

```
[edit services user-identification identity-management]
```

Description

Configure the parameters to be used for the IP query function. When this feature is enabled, the SRX Series Firewall queries the Juniper Identity Management Service (JIMS) server for user identity information based on the IP address of a user's device.

For example, if information for a user is missing from a flow, the SRX Series Firewall can issue a query request specifying the IP address of the user's device. Also, If the SRX Series Firewall does not have identity information for a specific user, it can engage captive portal to authenticate the user. After it authenticates the user, the SRX Series Firewall can issue a query request to the Juniper Identity Management Service, specifying the user ID and the IP address of the user's device to obtain additional information, such as the names of the groups that the user belongs to.

If there are many IP query requests in the queue, the SRX Series Firewall can maintain multiple concurrent HTTP/HTTPS connections with the Juniper Identity Management Service to increase throughput. However, the number of concurrent connections are kept at a reasonable level, which is twenty or less, so as not to impose pressure on the Juniper Identity Management Service.

NOTE: IP query is one of three query methods: IP query, batch query, and user query. All three types of queries can occur concurrently. They are not mutually exclusive.

The advanced user identity query feature, to which this configuration statement belongs, relies on the Juniper Identity Management Service that allows you to provision users locally and have their authentication information made available to other sites in your network for policy enforcement and reporting. The feature allows the SRX Series Firewall to query the Juniper Identity Management Service to pull user identity information.

> **WARNING**: Before you use this feature, you must disable active-directory-access and authentication-source options under the user-identification hierarchy. You cannot commit this configuration if active directory authentication or the ClearPass query and webapi functions are configured and committed.

To obtain device information, such as device identity, groups, and the operating system, from the Juniper Identity Management Service server using either the batch-query or ip-query configuration, you must set the device authentication source, as follows.

```
user@host# set services user-identification device-information authentication-source network-access-controller
```

Options

no-ip-query Disable IP query. IP query is enabled by default.

query-delay-time Time after which the SRX Series Firewall sends the query. Rather than allow the SRX Series Firewall to respond automatically by sending a user query *immediately*, you can set a `query-delay-time` parameter, specified in seconds, that allows the SRX Series Firewall to wait for a period of time before sending the query.

- **Default:** `15`
- **Range:** `0-60` seconds

Required Privilege Level

1. services—To view this statement in the configuration.
2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Understanding Advanced Query Feature for Obtaining User Identity Information from JIMS

primary connection (Identity Management Advanced Query) | 13131
secondary connection (Identity Management Advanced Query) | 16409
invalid-authentication-entry-timeout (Services User Identification Active Directory and ClearPass) | 9227

ip-reassembly

IN THIS SECTION

- Syntax | 9262
- Hierarchy Level | 9262
- Description | 9262
- Options | 9263
- Required Privilege Level | 9263
- Release Information | 9263

Syntax

```
ip-reassembly {
    profile profile-name
    rule rule-name{
        match-direction direction
    };
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the IP reassembly parameters to be applied to the L2TP server.

> **NOTE**: Inline IP reassembly configuration does not require you to configure the `profile` statement. The `profile` configuration is used when IP reassembly is configured on services PICs.

Options

profile ***profile-name***	Name of the IP reassembly profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring IP Inline Reassembly for L2TP

IP Packet Fragment Reassembly for L2TP Overview

ip-reassembly (L2TP)

IN THIS SECTION

- Syntax | 9264
- Hierarchy Level | 9264

- Description | 9264
- Options | 9264
- Required Privilege Level | 9264
- Release Information | 9265

Syntax

```
ip-reassembly {
    service-set service-set-name;
}
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Associate the reassembly service-set with the L2TP service.

NOTE: The service set must be defined at the `[edit services]` hierarchy level.

Options

service-set *service-set-name* Identifies the service set to be associated with the L2TP service.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

> *IP Packet Fragment Reassembly for L2TP Overview*
>
> *Configuring IP Inline Reassembly for L2TP*

ip-reassembly-rules (Service Set)

IN THIS SECTION

- Syntax | 9265
- Hierarchy Level | 9265
- Description | 9266
- Options | 9266
- Required Privilege Level | 9266
- Release Information | 9266

Syntax

```
ip-reassembly-rules rule-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify one or more previously configured IP reassembly rules to associate with the service set.

> **NOTE**: The IP reassembly rule must be defined at the `[edit services ip-reassembly rule]` hierarchy level.

Options

rule-name	Name of an IP reassembly rule.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring IP Inline Reassembly for L2TP

IP Packet Fragment Reassembly for L2TP Overview

ip (Security IDP Custom Attack)

IN THIS SECTION

- Syntax | 9267
- Hierarchy Level | 9267

- Description | 9267
- Options | 9267
- Required Privilege Level | 9267
- Release Information | 9268

Syntax

```
ip {
    protocol-number transport-layer-protocol-number;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain protocol-binding],

[edit security idp custom-attack attack-name attack-type signature protocol-binding]
```

Description

Allow IDP to match the attack for a specified IP protocol type.

Options

`protocol-number` *`transport-layer-protocol-number`*—Transport Layer protocol number.

- **Range:** 0 through 139

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

ip (Security Screen)

IN THIS SECTION

- Syntax | 9268
- Hierarchy Level | 9270
- Description | 9270
- Options | 9270
- Required Privilege Level | 9271
- Release Information | 9271

Syntax

```
ip {
    bad-option;
    block-frag {
        white-list name;
        }
    ipv6-extension-header {
        AH-header;
        ESP-header;
        HIP-header;
        destination-header {
            ILNP-nonce-option;
            home-address-option;
            line-identification-option;
            tunnel-encapsulation-limit-option;
            user-defined-option-type <type-low> to <type-high>;
        }
        fragment-header;
        hop-by-hop-header {
```

```
            CALIPSO-option;
            RPL-option;
            SFM-DPD-option;
            jumbo-payload-option;
            quick-start-option;
            router-alert-option;
            user-defined-option-type <type-low> to <type-high>;
        }
        mobility-header;
        no-next-header;
        routing-header;
        shim6-header
        user-defined-option-type <type-low> to <type-high>;
    }
    ipv6-extension-header-limit limit;
    ipv6-malformed-header;
    loose-source-route-option;
    record-route-option;
    security-option;
    source-route-option;
    spoofing;
    stream-option;
    strict-source-route-option;
    tear-drop;
    timestamp-option;
    unknown-protocol;
    tunnel {
        gre {
            gre-4in4;
            gre-4in6;
            gre-6in4;
            gre-6in6;
        }
        ip-in-udp {
            teredo;
        }
        ipip {
            ipip-4in4;
            ipip-4in6;
            ipip-6in4;
            ipip-6in6;
            ipip-6over4;
            ipip-6to4relay;
```

```
            isatap;
            dslite;
        }
        bad-inner-header;
    }
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name]
```

Description

Configure IP layer IDS options.

Options

- `bad-option`—Detect and drop any packet with an incorrectly formatted IP option in the IP packet header. The device records the event in the screen counters list for the ingress interface. This screen option is applicable to IPv4 and IPv6.
- `block-frag`—Enable IP packet fragmentation blocking.
- `loose-source-route-option`—Detect packets where the IP option is 3 (loose source routing), and record the event in the screen counters list for the ingress interface. This option specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified. The type 0 routing header of the loose source route option is the only related header defined in IPv6 .
- `record-route-option`—Detect packets where the IP option is 7 (record route), and record the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- `security-option`—Detect packets where the IP option is 2 (security), and record the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- `source-route-option`—Detect packets, and record the event in the screen counters list for the ingress interface.
- `spoofing`—Prevent spoofing attacks. Spoofing attacks occur when unauthorized agents attempt to bypass firewall security by imitating valid client IP addresses. Using the spoofing option invalidates such false source IP address connections.

The default behavior is to base spoofing decisions on individual interfaces.

- `stream-option`—Detect packets where the IP option is 8 (stream ID), and record the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- `strict-source-route-option`—Detect packets where the IP option is 9 (strict source routing), and record the event in the screen counters list for the ingress interface. This option specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field. Currently, this screen option is applicable only to IPv4.
- `tear-drop`—Block the teardrop attack. Teardrop attacks occur when fragmented IP packets overlap and cause the host attempting to reassemble the packets to crash. The teardrop option directs the device to drop any packets that have such a discrepancy.
- `timestamp-option`—Detect packets where the IP option list includes option 4 (Internet timestamp), and record the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- `unknown-protocol`—Discard all received IP frames with protocol numbers greater than 137 for IPv4 and 139 for IPv6. Such protocol numbers are undefined or reserved.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

ip-source-guard

IN THIS SECTION

- Syntax | 9272
- Hierarchy Level | 9272
- Description | 9272
- Default | 9273
- Required Privilege Level | 9273
- Release Information | 9273

Syntax

```
ip-source-guard;
```

Hierarchy Level

- For platforms with ELS:

```
[edit vlans vlan-name forwarding-options dhcp-security]
```

- For platforms without ELS:

```
[edit ethernet-switching-options  secure-access-port vlan (all | vlan-name)]
```

Description

Perform IP source guard checking on packets sent from access interfaces. Validate source IP addresses and source MAC addresses on all VLANs or on the specified VLAN or VLAN range. Forward packets with valid addresses and drop those with invalid addresses.

- `ip-source-guard`—Enable IP source guard checking.

- `no-ip-source-guard`—(Not available in `[edit vlans vlan-name forwarding-options dhcp-security]`) Disable IP source guard checking.

If you configure IP source guard at the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level:

- IP source guard can be configured only for a specific VLAN, not for a list or a range of VLAN IDs.
- DHCP snooping is automatically enabled.

See *Configuring IP Source Guard (ELS)* for more information about this configuration.

If you configure IP source guard at the `[edit ethernet-switching-options secure-access-port vlan (all | vlan-name]` hierarchy level:

- You must enable DHCP snooping on all VLANs if you configure IP source guard on all VLANs.
- You must enable DHCP snooping for the specific VLAN if you configure IP source guard on that specific VLAN. Otherwise, the default behavior of no DHCP snooping applies to that VLAN.

See *Enabling DHCP Snooping (non-ELS)* for more information about this configuration.

NOTE: On EX9200 switches, IP source guard is not supported in an MC-LAG scenario.

Default

Disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN

Example: Configuring IP Source Guard with Other EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces

Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing

Configuring IP Source Guard (non-ELS)

Configuring IP Source Guard (ELS)

ip-source-guard (MX Series)

IN THIS SECTION

- Syntax | 9274
- Hierarchy Level | 9274
- Description | 9274
- Required Privilege Level | 9275
- Release Information | 9275

Syntax

```
ip-source-guard;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]
```

Description

Perform IP source guard checking on packets sent from access interfaces. Validate source IP addresses and source MAC addresses on all bridge domains or on the specified bridge domain or bridge domain range. Forward packets with valid addresses and drop those with invalid addresses.

- `ip-source-guard`—Enable IP source guard checking.

If you configure IP source guard at the `[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]` hierarchy level:

- IP source guard can be configured only for a specific bridge domain, not for a list or range of bridge domains.
- DHCP snooping is automatically enabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Hierarchy level `[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]` introduced in Junos OS Release 14.1 for the MX Series.

RELATED DOCUMENTATION

> *Configuring IP Source Guard (non-ELS)*
>
> *Example: Configuring IP Source Guard and Dynamic ARP Inspection on a Specified Bridge Domain to Protect the Devices Against Attacks*

ip-subscriber

IN THIS SECTION

- Syntax | 9276
- Hierarchy Level | 9277
- Description | 9277
- Required Privilege Level | 9277
- Release Information | 9277

Syntax

```
ip-subscriber {
    access-interfaces interface-name [interface-name];
    default-local-policy {
        flow-action (drop | forward);
        maximum-bit-rate {
            uplink mbr-uplink-value ;
            downlink mbr-downlink-value;
        }
        burst-size {
            uplink uplink-burst-size;
            downlink downlink-burst-size;
        }
    }
    idle-timeout idle-timout;
    immediate-accounting-response (enabled | disabled);
    maximum-subscribers number;
    subscriber-address {
        inet {
            pool pool-name;
        }
        inet6 {
            pool pool-name;
        }
    }
    subscription-id {
        constant ;
        subscription-id-options {
            entry-name {
                id-components {
                    use-imsi;
                    use-msisdn;
                    use-nai;
                    use-username;
                    use-realm;
                    use-nas-port;
                    use-nas-port-id;
                }
            }
        }
```

```
    }
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
```

Description

Configure TDF domain features that are unique to IP-based subscribers.

The remaining statements are described separately.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

ip-swap

IN THIS SECTION

Syntax | 9278

- Junos OS Hierarchy Level | 9278
- Junos OS Evolved Hierarchy Level | 9278
- Description | 9278
- Required Privilege Level | 9278
- Release Information | 9278

Syntax

```
ip-swap;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Swaps source and destination IPv4 addresses. This statement is applicable only for family `bridge` or family `ethernet-switching`.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

Configuring an RFC 2544-Based Benchmarking Test

rfc2544-benchmarking | **15759**

rfc2544 | **15757**

ip-sweep

IN THIS SECTION

- Syntax | **9279**
- Hierarchy Level | **9279**
- Description | **9280**
- Options | **9280**
- Required Privilege Level | **9280**
- Release Information | **9280**

Syntax

```
ip-sweep {
    threshold number;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name icmp]
```

Description

Configure the device to detect and prevent an IP Sweep attack. An IP Sweep attack occurs when an attacker sends ICMP echo requests (pings) to multiple destination addresses. If a target host replies, the reply reveals the target's IP address to the attacker. If the device receives 10 ICMP echo requests within the number of microseconds specified in this statement, it flags this as an IP Sweep attack, and rejects the 11th and all further ICMP packets from that host for the remainder of the second.

Options

`threshold` *`number`*—Maximum number of microseconds during which up to 10 ICMP echo requests from the same host are allowed into the device. More than 10 requests from a host during this period triggers an IP Sweep attack response on the device during the remainder of the second.

- **Range:** 1000 through 1,000,000 microseconds
- **Default:** 5000 microseconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Attack Detection and Prevention Overview*
>
> *Example: Configuring Multiple Screening Options*

ip-tunnel-rpf-check

IN THIS SECTION

- Syntax | 9281
- Hierarchy Level | 9281
- Description | 9281
- Options | 9282
- Required Privilege Level | 9282
- Release Information | 9282

Syntax

```
ip-tunnel-rpf-check {
    mode (strict | loose);
    fail-filter filter-name;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name routing-options forwarding-table]
```

Description

Configure the system to enable anti-spoofing protection for next-hop-based dynamic tunnels, where reverse path forwarding checks are placed to ensure that the tunnel traffic is received from a legitimate source through designated IP tunnel, where the source is reachable on the same tunnel on which the packet was received.

When a packet comes from a nondesignated source, the reverse path forwarding check fails in the strict mode, and passes in the loose mode. When a packet comes from a nonexistent source, the reverse path forwarding check fails.

By default, the reverse path forwarding check is in strict mode, where the packets are not forwarded if the source of the packet is from a nondesignated tunnel.

Options

mode (*strict* | *loose*)

(Optional) Specify the mode of the reverse path forwarding check to enable anti-spoofing protection for next-hop-based dynamic tunnels.

In the strict mode (default), the reverse path forwarding check fails when the packet is received from a nondesignated tunnel source. The check passes only for packets from designated tunnels.

In the loose mode, the reverse path forwarding check passes even if the packet is received from a nondesignated tunnel source.

When the packet is from a nonexistent tunnel source, the reverse path forwarding check fails in both the strict and loose modes.

- **Default:** If you omit the mode statement, the default behavior is strict mode.

fail-filter *filter-name*

(Optional) Attach a filter to the Layer 3 VPN to log packets that failed the reverse path forwarding check.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview

Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels

ip-user-mapping

IN THIS SECTION

- Syntax | 9283
- Hierarchy Level | 9283
- Description | 9283
- Options | 9284
- Required Privilege Level | 9284
- Release Information | 9284

Syntax

```
ip-user-mapping {
    discovery-method {
        wmi {
            event-log-scanning-interval seconds;
            initial-event-log-timespan hours;
        }
    }
}
```

Hierarchy Level

```
[edit services user-identification active-directory domain]
```

Description

Control how the SRX Series Firewall accesses a domain controller in order to monitor and scan security event logs on the domain controller. By parsing the event log, the SRX Series gets IP address-to-user mappings. This process is part of the integrated user firewall feature. The **ip-user-mapping** statement is optional because WMI is the default discovery method and its properties have default values.

The other available method the SRX Series uses to retrieve address-to-user mapping information is manual (on-demand) probing of a domain PC.

Options

discovery-method Method of discover IP address-to-user mappings.

wmi Windows Management Instrumentation (WMI) is the discovery method used to access the domain controller.

event-log-scanning-interval *seconds* Optional. Interval at which the SRX Series scans the event log on the domain controller.

- **Range:** 5 through 60 seconds
- **Default:** 10 seconds

initial-event-log-timespan *hours* Optional. Time of the earliest event log on the domain controller that the SRX Series will initially scan. This argument applies to the initial deployment only. After WMIC and the user identification start working, the SRX Series scans only the latest event log.

- **Range:** 1 through 168 hours
- **Default:** 1 hour

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

active-directory-access | 1802

clear services user-identification active-directory-access | 4194

request services user-identification active-directory-access ip-user-probe | 14985
user-identification (Services) | 30560
show services user-identification active-directory-access statistics | 26389
traceoptions (Active Directory Access) | 29483

ip-version

IN THIS SECTION

- Syntax | 9285
- Hierarchy Level | 9285
- Description | 9285
- Options | 9286
- Required Privilege Level | 9286
- Release Information | 9286

Syntax

```
ip-version ip-version {
    match-conditions;
    protocol (tcp | udp) {
        match conditions;
    }
}
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name from]
```

Description

Configure the IP version for the firewall filter.

Options

ip-version Version of the IP addressing.

- `ipv4`—IP version 4
- `ipv6`—IP version 6

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1R1.

Option `ipv6` introduced in Junos OS Release 12.2R1.

RELATED DOCUMENTATION

Firewall Filter Match Conditions Based on Address Fields

Firewall Filter Match Conditions for IPv4 Traffic

Firewall Filter Match Conditions for IPv6 Traffic

Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic

ip-version (Family MPLS)

IN THIS SECTION

- Syntax | 9287
- Hierarchy Level | 9287
- Description | 9287
- Required Privilege Level | 9287

Release Information | 9287

Syntax

```
ip-version {
    ipv4;
    ipv6;
}
```

Hierarchy Level

```
[edit firewall family mpls filter name term name from]
```

Description

Specify inner IP version to enable IP-based filtering of MPLS family filter.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Understanding IP-Based Filtering and Selective Port Mirroring of MPLS Traffic

ipc

IN THIS SECTION

- Syntax | 9288
- Hierarchy Level | 9288
- Description | 9288
- Options | 9288
- Required Privilege Level | 9289
- Release Information | 9289

Syntax

```
ipc {
    traceoptions flag flag;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm]
```

Description

Configure trace options for IPC.

Options

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Enable trace for all IPC trace options.
  - `basic`—Trace basic IPC related information.

- `connection-manager`—Trace IPC connection manager information.
- `connection-status`—Trace IPC connection status information.
- `detail`—Trace IPC related detailed information.
- `pfe`—Trace communication with PFE.
- `utm-realtime`—Trace IPC realtime-thread information.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

ipcp-suggest-dns-option

IN THIS SECTION

- Syntax | 9289
- Hierarchy Level | 9290
- Description | 9290
- Required Privilege Level | 9290
- Release Information | 9290

Syntax

```
ipcp-suggest-dns-option;
```

Hierarchy Level

```
[edit access group-profile group-profile-name ppp ppp-options],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" ppp-options],
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit "$junos-
interface-unit" ppp-options],
[edit interfaces pp0 unit logical-unit-number ppp-options],
[edit interfaces si-fpc/pic/port unit logical-unit-number ppp-options]
```

Description

Configure the router to prompt Customer Premises Equipment (CPE) to negotiate both primary and secondary DNS addresses during IPCP negotiation for terminated PPPoE and LNS subscribers. You can configure this for dynamic or static PPPoE subscribers, dynamic or static LNS subscribers, and in an LNS group profile.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Ensuring IPCP Negotiation for Primary and Secondary DNS Addresses

Configuring the PPP Attributes for a Group Profile

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile

Configuring Dynamic Authentication for PPP Subscribers

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

ipfix-sw-mode

IN THIS SECTION

- Syntax | 9291
- Hierarchy Level | 9291
- Description | 9291
- Required Privilege Level | 9291
- Release Information | 9291

Syntax

```
ipfix-sw-mode;
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

(QFX5120 only) Enable software-based IPFIX, used to enable flow-based telemetry for VXLANs. You must reboot the device after you commit this statement to the configuration.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Flow-Based Telemetry for VXLANs

ipip

IN THIS SECTION

- Syntax | 9292
- Hierarchy Level | 9292
- Description | 9293
- Required Privilege Level | 9293
- Release Information | 9293

Syntax

```
ipip {
    destination {
        address address;
    }
    source {
        address address;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit name tunnel encapsulation
        (interfaces-fti)],
[edit dynamic-profiles name logical-systems name interfaces name unit name tunnel encapsulation
        (interfaces-fti)],
[edit interfaces name unit name tunnel encapsulation
        (interfaces-fti)]
```

Description

Configure IP-in-IP tunnel.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.3R1.

ipip (lds-option)

IN THIS SECTION

- Syntax | 9293
- Hierarchy Level | 9294
- Description | 9294
- Options | 9294
- Required Privilege Level | 9294
- Release Information | 9294

Syntax

```
ipip {
    ipip-4in4;
    ipip-4in6;
    ipip-6in4;
    ipip-6in6;
    ipip-6over4;
    ipip-6to4relay;
    isatap;
```

```
        dslite;
    }
```

Hierarchy Level

```
[edit security screen ids-option ids-option-name ip tunnel]
```

Description

Configure IP tunnel IP-IP IDS options.

Options

dslite	Enable IP tunnel IP-IP DS-Lite IDS option.
ipip-4in4	Enable IP tunnel IP-IP 4in4 IDS option.
ipip-4in6	Enable IP tunnel IP-IP 4in6 IDS option.
ipip-6in4	Enable IP tunnel IP-IP 6in4 IDS option.
ipip-6in6	Enable IP tunnel IP-IP 6in6 IDS option.
ipip-6over4	Enable IP tunnel IP-IP 6over4 IDS option.
ipip-6to4relay	Enable IP tunnel IP-IP 6to4 Relay IDS option.
isatap	Enable IP tunnel IP-IP ISATAP IDS option.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D10.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

ips

IN THIS SECTION

- Syntax | 9295
- Hierarchy Level | 9295
- Description | 9296
- Options | 9296
- Required Privilege Level | 9298
- Release Information | 9298

Syntax

```
ips {
    content-decompression-max-memory-kb content-decompression-max-memory-kb;
    content-decompression-max-ratio content-decompression-max-ratio;
    (detect-shellcode | no-detect-shellcode);
    fifo-max-size fifo-max-size;
    (ignore-regular-expressionignore-regular-expression | no-ignore-regular-expression);
    log-supercede-min log-supercede-min;
    (process-ignore-s2c | no-process-ignore-s2c);
    (process-override | no-process-override);
    process-port process-port;
    session-pkt-depth session-pkt-depth;
}
```

Hierarchy Level

```
[edit security idp sensor-configuration]
```

Description

Configure IPS security policy sensor settings. The IPS sensor contains filters, signature entries, or both. These specify which signatures are included in the IPS sensor. You need to create an IPS sensor before specific signatures or filters can be chosen. The signatures can be added to a new sensor before it is saved. However, it is good practice to keep in mind that the sensor and its included filters are separate things, and that they are created separately. While individual signatures can be added to a sensor, a filter allows you to add multiple signatures to a sensor by specifying the characteristics of the signatures to be added.

Options

content-decompression-max-memory-kb

Set the maximum memory allocation in kilobytes for content decompression.

The default memory allocation provides 33 KB per session for an average number of sessions requiring decompression at the same time. To determine if this value is consistent with your environment, analyze values from decompression-related counters and the total number of IDP sessions traversing the device. Estimate the number of sessions requiring decompression at the same time. Assuming that each of these sessions requires 33 KB of memory for decompression, compare your estimated needs to the default value.

> **NOTE**: Because content decompression requires a significant allocation of memory, system performance will be impacted by increasing the maximum memory allocation for decompression.

- **Range:** 50 through 2,000,000 KB

content-decompression-max-ratio

Set the maximum decompression ratio of the size of decompressed data to the size of compressed data.

Some attacks are introduced through compressed content. When the content is decompressed, it can inflate to a very large size taking up valuable system resources resulting in denial of service. This type of attack can be recognized by the ratio of the size of decompressed data to the size of compressed data. Keep in mind, however, that a higher ratio lessens the chance of detecting this type of attack.

- **Range:** 1 through 128

detect-shellcode

Enable to detect the shell code and prevent buffer overflow attacks. By default this setting is enabled.

fifo-max-size

Sets the maximum IPS FIFO size.

- **Range:** 1 through 65,535

ignore-regular-expression

To detect intrusion attempts, you can enable regular expression by issuing the `no-ignore-regular-expression` command. By default, the `no-ignore-regular-expression` command is enabled. If you specify the `ignore-regular-expression` command, regular expression pattern matching will be disabled when detecting intrusion attempts.

- **Default:** Regular expression is enabled by default.

log-supercede-min

Specify the amount of time to supersede the IPS sensor logs.

- **Syntax:** *`minimum-value`*—Minimum time to supersede the log.
- **Default:** 1 second
- **Range:** 0 through 65,535

no-detect-shellcode

Don't detect shellcode

no-ignore-regular-expression

Don't ignore regular expression

no-process-ignore-s2c

Don't process ignore s2c

no-process-override

Don't process override

process-ignore-s2c

Set the command to disable the server-to-client inspection.

process-override

Set the command to forcefully run the IDS inspection module even if there is no policy match.

process-port

Set the command to a specific port to forcefully run the IDS inspection module on that TCP/UDP port even if there is no policy match.

- **Syntax:** *`port-number`*—Port number.
- **Range:** 0 through 65,535

session-pkt-depth

Set the command specify the Session packet scanning depth.

- **Syntax:** *session-pkt-depth*—Session packet depth.
- **Range:** 0 through 1000000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

IDP Sensor Configuration

sensor-configuration | 16560

ipsec

IN THIS SECTION

- Syntax | 9298
- Hierarchy Level | 9299
- Description | 9299
- Required Privilege Level | 9299
- Release Information | 9299

Syntax

```
ipsec {
    proposal proposal-name {
```

```
        authentication-algorithm (hmac-sha-256);
        description description;
        encryption-algorithm algorithm;
        lifetime-seconds seconds;
        protocol (esp | bundle);
    }
    policy policy-name {
        description description;
        perfect-forward-secrecy {
            keys (group1 | group2 |group5 |group14 |group15 | group16 | group24);
        }
        proposals [ proposal-names ];
    }
}
```

Hierarchy Level

```
[edit services ipsec-vpn]
```

Description

Configure IPsec.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Security Associations

ipsec-group-vpn (Security Policies)

IN THIS SECTION

- Syntax | 9300
- Hierarchy Level | 9300
- Description | 9300
- Options | 9300
- Required Privilege Level | 9300
- Release Information | 9301

Syntax

```
ipsec-group-vpn group-vpn;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone  zone-name policy policy-name then permit
tunnel]
```

Description

Specify group VPN tunnel for the traffic configured by the scope policy on a group member.

Options

group-vpn —Name of the group VPN tunnel configured on the group member.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Security Policies Overview

Group VPNv2 Overview

Security Basics Overview

ipsec (High Availability)

IN THIS SECTION

- Syntax | 9301
- Hierarchy Level | 9302
- Description | 9302
- Options | 9302
- Required Privilege Level | 9303
- Release Information | 9304

Syntax

```
ipsec {
    vpn vpn-name {
        ha-link-encryption;
        ike {
            gateway gateway-name;
            ipsec-policy ipsec-policy-name;
        }
    }
    proposal proposal-name {
        description description;
```

```
        encryption-algorithm (aes-256-gcm);
        lifetime-seconds seconds;
        protocol (esp);
    }
    policy policy-name {
        description description;
        proposals proposal-name;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Define IPsec configuration for the multinode high availability feature. A VPN connection can link two LANs (site-to-site VPN) or a remote dial-up user and a LAN. The traffic that flows between these two points passes through shared resources such as routers, switches, and other network equipment that make up the public WAN. An IPsec tunnel is created between two participant devices to secure VPN communication.

Options

vpn-name Configure an IPsec VPN. A VPN provides a means by which remote computers communicate securely across a public WAN such as the Internet.

You must mention the same VPN name for `vpn-profile` in `set chassis high-availability peer-id` *`peer-id`* `vpn-profile` *`profile-name`* configuration.

ha-link-encryption Configure a interchassis link tunnel for secure HA traffic flow between the nodes. Only site-to-site IPsec VPN tunnels are supported for interchassis link tunnels. Both PSK and PKI authentication methods are supported.

gateway-name Name of the remote IKE gateway.

ipsec-policy-name Specify the IPsec policy name.

proposal-name Name of the IPsec proposal. An IPsec proposal lists protocols and algorithms (security services) to be negotiated with the remote IPsec peer.

description Text description of IPsec proposal.

encryption-algorithm Define encryption algorithm. The device deletes existing IPsec SAs when you update the `encryption-algorithm` configuration in the IPsec proposal.

A commit error is thrown if any value other than `aes-256-gcm` is configured.

- Values:
  - `aes-256-gcm`—AES GCM 256-bit encryption algorithm.

    For an IKE proposal, AES 256-bit authenticated encryption algorithm is supported with IKEv2 only. When this option is used, `aes-256-gcm` should be configured at the [`edit security ipsec proposal` *`proposal-name`*] hierarchy level, and the `authentication-algorithm` option should not be configured at the [`edit security ike proposal` *`proposal-name`*] hierarchy level.

lifetime-seconds Lifetime in seconds.

- **Range:** 180 through 86400
- **Default:** 3600 seconds

protocol Define the IPsec protocol for a manual or dynamic security association (SA).

A commit error is thrown if any value other than `esp` is configured.

- Values:
  - esp—Encapsulated Security Payload header

policy-name Define an IPsec policy. An IPsec policy defines a combination of security parameters (IPsec proposals) used during IPsec negotiation. It defines Perfect Forward Secrecy (PFS) and the proposals needed for the connection.

description Enter descriptive text for an IPsec policy.

proposal-name Specify one or more proposals for an IPsec policy.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

- ike (High Availability) | 8424
- High-Availability (Chassis)
- Multinode High Availability
- services-redundancy-group
- local-id
- peer-id
- traceoptions
- show security ike security-associations | 24648
- show security ike active-peer | 24630
- show security ipsec security-associations | 24701
- show security ipsec statistics | 24744
- clear security ike security-associations | 3988
- clear security ipsec security-associations | 3993

ipsec-inside-interface

IN THIS SECTION

- Syntax | 9305
- Hierarchy Level | 9305
- Description | 9305
- Options | 9305
- Required Privilege Level | 9305
- Release Information | 9305

Syntax

```
ipsec-inside-interface interface-name;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name from]
```

Description

Specify the interface name for next-hop-style service sets. This value is also implicitly generated in dynamic endpoint tunneling.

Options

interface-name—Service interface for internal network.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring IPsec Rules*
>
> *Configuring Dynamic Endpoints for IPsec Tunnels*

ipsec-policy

IN THIS SECTION

- Syntax | 9306
- Hierarchy Level | 9306
- Description | 9306
- Options | 9306
- Required Privilege Level | 9307
- Release Information | 9307

Syntax

```
ipsec-policy from-zone zone-name to-zone zone-name ipsec-group-vpn vpn-name;
```

Hierarchy Level

```
[edit security]
```

Description

Specifies that matching traffic is checked against rules associated with the specified Group VPN. Exclude and fail-open rules are configured at the [`edit security group-vpn member ipsec vpn` *`vpn-name`*] hierarchy level.

Options

from-zone *zone-name*
Specify the incoming zone for Group VPN traffic.

to-zone *zone-name*
Specify the outgoing zone for Group VPN traffic.

The `to-zone` zone must include the interface configured with the `group-vpn-external-interface` option at the [`edit security group-vpn member ipsec vpn` *`vpn-name`*] hierarchy level.

ipsec-group-vpn *vpn-name* Specify the Group VPN to which the traffic applies. Only one Group VPN can be referenced by a specific from-zone/to-zone pair.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Group VPNv2 Overview

ipsec-sa

IN THIS SECTION

- Syntax | 9307
- Hierarchy Level | 9308
- Description | 9308
- Options | 9308
- Required Privilege Level | 9308
- Release Information | 9308

Syntax

```
ipsec-sa sa-name;
```

Hierarchy Level

```
[edit interfaces es-fpc/pic/port unit logical-unit-number family inet]
```

Description

Specify the IP Security (IPsec) SA name associated with the interface.

Options

sa-name—IPsec SA name.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Encryption Interfaces

Junos OS Administration Library for Routing Devices

ipsec-sa (Protocols BGP)

IN THIS SECTION

- Syntax | 9309
- Hierarchy Level | 9309
- Description | 9309

- Options | 9309
- Required Privilege Level | 9310
- Release Information | 9310

Syntax

```
ipsec-sa ipsec-sa;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Specify a security association to BGP peers. You can apply the security association globally for all BGP peers, to a group of peers, or to an individual peer.

Options

ipsec-sa—Security association name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Using IPsec to Protect BGP Traffic

ipsec-sa (Security Group VPN)

IN THIS SECTION

- Syntax | 9310
- Hierarchy Level | 9311
- Description | 9311
- Options | 9311
- Required Privilege Level | 9311
- Release Information | 9311

Syntax

```
ipsec-sa name {
    match-policy policy-name {
        destination ip-address/netmask;
        destination-port number;
        protocol number;
        source ip-address/netmask;
        source-port number;
```

```
    }
    proposal proposal-name;
}
```

Hierarchy Level

```
[edit security group-vpn server group name]
```

Description

Configure the group SAs to be downloaded to members. There can be multiple group SAs downloaded to group members.

Options

`ipsec-sa` *`name`*—Define the group SAs to be downloaded to members.

- `match-policy` *`policy-name`*—Configure the group policy with source address, source port, destination address, destination port, and protocol.
  - `destination` *`ip-address/netmask`*—Specify the destination IP address to be matched (0.0.0.0/0 for any).
  - `destination-port` *`number`*—Specify the destination port to be matched (0 for any).
  - `protocol` *`number`*—Specify the protocol number to be matched (0 for any).
  - `source` *`ip-address/netmask`*—Specify the source IP address to be matched (0.0.0.0/0 for any).
  - `source-port` *`number`*—Specify the source port to be matched (0 for any)
- `proposal` *`proposal-name`*—Specify the name of the IPsec proposal configured with the `proposal` configuration statement at the [`edit security group-vpn server ipsec`] hierarchy.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

group (Security Group VPN) | 7894

ipsec (Security)

IN THIS SECTION

- Syntax | 9312
- Hierarchy Level | 9313
- Description | 9313
- Options | 9313
- Required Privilege Level | 9314
- Release Information | 9314

Syntax

```
ipsec {
    anti-replay-window-size anti-replay-window-size;
    internal;
    policy;
    proposal
    security-association sa-name;
    traceoptions;
    vpn vpn-name {
                                    bind-interface st0. ;
                                }
    vpn-monitor-options {
        interval seconds;
        threshold number;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Define IPsec configuration. A VPN connection can link two LANs (site-to-site VPN) or a remote dial-up user and a LAN. The traffic that flows between these two points passes through shared resources such as routers, switches, and other network equipment that make up the public WAN. An IPsec tunnel is created between two participant devices to secure VPN communication.

Options

anti-replay-window-size

Anti-replay window size.

- **Range:** 64 through 8192 bytes
- **Default:** 64 bytes

internal

Configure internal IPsec. When the internal IPsec is configured, IPsec-based `rlogin` and remote command (`rcmd`) are enforced, so an attacker cannot gain unauthorized information.

policy

Define an IPsec policy. An IPsec policy defines a combination of security parameters (IPsec proposals) used during IPsec negotiation. It defines Perfect Forward Secrecy (PFS) and the proposals needed for the connection.

proposal

Name of the IPsec proposal. An IPsec proposal lists protocols and algorithms (security services) to be negotiated with the remote IPsec peer.

security-association

Configure a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link. IPsec can provide authentication and confidentiality to OSPF or OSPFv3 routing packets.

traceoptions

Configure IPsec tracing options. Trace operations track IPsec events and record them in a log file in the /var/log directory.

vpn *vpn-name*

Configure an IPsec VPN. A VPN provides a means by which remote computers communicate securely across a public WAN suchas the Internet

vpn-monitor-options

Configure VPN monitoring options

interval *seconds* — Interval at which to send ICMP requests to the peer.

- **Range:** 2 through 3600 seconds
- **Default:** 10 seconds

threshold *number* — Number of consecutive unsuccessful pings before the peer is declared unreachable.

- **Range:** 1 through 65,535 pings
- **Default:** 10 pings

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

`group15`, `group16`, `group21`, `hmac-sha-512` and `hmac-sha-384` options introduced in Junos OS Release 19.1R1 on SRX Series Firewalls.

RELATED DOCUMENTATION

IPsec Overview

ipsec (Security Group VPN Member)

IN THIS SECTION

- Syntax | 9315
- Hierarchy Level | 9315
- Description | 9315

- Options | 9316
- Required Privilege Level | 9317
- Release Information | 9317

Syntax

```
ipsec {
    vpn vpn-name {
        df-bit (clear | copy | set);
        exclude rule rule-name {
            source-address ip-address/mask;
            destination-address ip-address/mask;
            application application;
        }
        fail-open rule rule-name {
            source-address ip-address/mask;
            destination-address ip-address/mask;
            application application;
        }
        group id;
        group-vpn-external-interface interface;
        ike-gateway gateway-name;
        recovery-probe;
    }
t}
```

Hierarchy Level

```
[edit security group-vpn member]
```

Description

Configure IPsec for Phase 2 exchange on the group member. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 devices and vSRX Virtual Firewall instances.

Options

vpn *vpn-name*
Name of the VPN.

df-bit
Specifies pre-fragmentation and post-fragmentation of IPsec traffic on the group member. One of the following options can be configured:

- clear—Sets the outer IP do not fragment (DF) bit to 0. When the packet size is larger than the path maximum transmission unit (path MTU), pre-fragmentation is done if the DF bit is not set in the inner packet and post-fragmentation is done if the DF bit is set in the inner packet. This is the default.
- copy—Copies the DF bit from the inner header to the outer header. When the packet size is larger than the path PMTU, pre-fragmentation is done if the DF bit is not set in the inner packet. If the DF bit is set in the inner packet, the packet is dropped and an ICMP message is sent back.
- set—Sets the outer IP DF bit to 1. When the packet size is larger than the path MTU, pre-fragmentation is done if the DF bit is not set in the inner packet. If the DF bit is set in the inner packet, the packet is dropped and an ICMP message is sent back

exclude rule
Specifies traffic to be excluded from Group VPN encryption. A maximum of 10 exclude rules can be configured. Source and destination addresses must be specified in *ip-address/mask* format; address books and address sets are not supported. Predefined and user-defined applications are supported, but application sets are not supported.

fail-open rule
Specifies the traffic to be sent in cleartext mode if there is no valid SA key available to protect the traffic. Traffic that is not specified by the fail-open rule is blocked if there is no valid SA key available to protect the traffic. A maximum of 10 fail-open rules can be configured. Source and destination addresses must be specified in *ip-address/mask* format; address books and address sets are not supported. Predefined and user-defined applications are supported, but application sets are not supported.

group *id*
Identifier configured for the Group VPN.

group-vpn-external-interface *interface*
Interface used by the group member to connect to the Group VPN peers. The interface must belong to the same zone as the `to-zone` configured at the [`edit security ipsec-policy`] hierarchy level for Group VPN traffic.

ike-gateway *gateway-name*
Name of the IKE gateway for the Group VPN.

recovery-probe
Enables initiation of `groupkey-pull` exchanges at specific intervals to update the member's SA from the group server if the group member is determined to be out of

synchronization with the group server and other group members. This option is disabled by default.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2. `df-bit`, `exclude rule`, `fail-open rule`, and `recovery-probe` options added in Junos OS Release 15.1X49-D30 for vSRX Virtual Firewall.

RELATED DOCUMENTATION

Group VPNv2 Overview

ipsec (Security Group VPN Server)

IN THIS SECTION

- Syntax | 9317
- Hierarchy Level | 9318
- Description | 9318
- Options | 9318
- Required Privilege Level | 9318
- Release Information | 9318

Syntax

```
ipsec {
    proposal proposal-name {
```

```
        authentication-algorithm (hmac-sha-256-128);
        description description;
        encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
        lifetime-seconds seconds;
    }
}
```

Hierarchy Level

```
[edit security group-vpn server]
```

Description

Configure IPsec proposal for Phase 2 exchange on the group server. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 devices and vSRX Virtual Firewall instances.

Options

`proposal` *`proposal-name`*—Name of the proposal. The proposal name can be up to 32 alphanumeric characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

ipsec (Security Interfaces)

IN THIS SECTION

- Syntax | 9319
- Hierarchy Level | 9320
- Description | 9321
- Options | 9321
- Required Privilege Level | 9321
- Release Information | 9321

Syntax

```
ipsec {
    security-association {
        manual {
            direction (bidirectional | inbound | outbound) {
                protocol esp;
                spi spi-value;
                encryption {
                    algorithm 3des-cbc;
                    key  (ascii-text ascii-text-string | hexadecimal hexadecimal-string);
                }
            }
        }
    }
    policy ipsec-policy-name {
        perfect-forward-secrecy {
            keys (group1 | group2);
        }
        proposals [ proposal-names ];
    }
    proposal ipsec-proposal-name {
        authentication-algorithm (hmac-sha1-96 | hmac-sha2-256);
        encryption-algorithm (3des-cbc | des-cbc);
        lifetime-seconds seconds;
        protocol (ah | esp | bundle);
```

```
    }
    security-association name {
        dynamic {
            ipsec-policy policy-name;
            replay-window-size (32 | 64);
        }
        manual {
            direction (inbound | outbound | bi-directional) {
                authentication {
                    algorithm (hmac-sha1-96 | hmac-sha2-256);
                    key (ascii-text key | hexadecimal key);
                }
                auxiliary-spi auxiliary-spi-value;
                encryption {
                    algorithm (des-cbc | 3des-cbc);
                    key (ascii-text key | hexadecimal key);
                }
                protocol (ah | esp | bundle);
                spi spi-value;
            }
        }
        mode (tunnel | transport);
    }
    traceoptions {
        file <files number> < size size>;
        flag all;
        flag database;
        flag general;
        flag ike;
        flag parse;
        flag policy-manager;
        flag routing-socket;
        flag timer;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure IPsec on encryption interfaces.

> **NOTE**: You must configure the IPsec keys as hexadecimal keys for maximum key strength with Junos OS in FIPS mode.

Options

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Security Associations for IPsec on an ES PIC

ipsec-traffic-selector

IN THIS SECTION

- Syntax | 9322
- Hierarchy Level | 9322
- Description | 9322
- Default | 9322
- Options | 9323

Required Privilege Level | 9323
Release Information | 9323

Syntax

```
ipsec-traffic-selector preference pref-value;
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure the global preference value for the ARI-TS (Auto route insertion for traffic selectors) routes added when a IPsec traffic selector based IPsec VPN has been established.

A new category `ipsec-traffic-selector` for Auto route insertion traffic selectors (ARI-TS) route is added in `[edit protocols]` hierarchy for setting global preference value for ARI-TS route.

When a preference value is updated either in global scope or local scope, the already added secure tunnel routes are updated with new preference value.

The preference value update affects the ARI-TS tunnel routes as follows:

- Any changes to the global values impact all traffic selectors configured without a local value.
- Any changes to the local values impact only those specific tunnels.
- Any routes impacted by the change are deleted and added again with the new value.

Default

The default value for `ipsec-traffic-selector` preference value is `5` if not configured.

Options

preference ***pref-value***

Specify the preferential selection of routes for the same remote IP address or prefix added by different protocols (static, BGP, OSPF, and so on).

- **Range:** 1-4294967295.
- **Default:** 5.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

traffic-selector | 30042

ipsec-vpn-options

IN THIS SECTION

- Syntax | 9324
- Hierarchy Level | 9324
- Description | 9324
- Options | 9324
- Required Privilege Level | 9324
- Release Information | 9324

Syntax

```
ipsec-vpn-options {
    anti-replay-window-size bits;
    clear-dont-fragment-bit;
    ike-access-profile profile-name;
    local-gateway address;
    no-certificate-chain-in-ike;
    no-anti-replay;
    passive-mode-tunneling;
    trusted-ca [ ca-profile-names ];
    tunnel-mtu bytes;
    udp-encapsulation {
        <udp-dest-port destination-port>;
    }
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify IP Security (IPsec) service options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Service Sets

Configuring Service Rules

ipsec-vpn-rules

IN THIS SECTION

- Syntax | 9325
- Hierarchy Level | 9325
- Description | 9325
- Options | 9325
- Required Privilege Level | 9326
- Release Information | 9326

Syntax

```
(ipsec-vpn-rules rule-name | ipsec-vpn-rule-sets rule-set-name);
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the IPsec rules or rule set included in this service set. You can configure multiple rules, but only one rule set for each service.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

rule-set-name—Identifier for the set of rules to be included.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Service Rules

ipsec-vpn (Security Policies)

IN THIS SECTION

- Syntax | 9326
- Hierarchy Level | 9327
- Description | 9327
- Options | 9327
- Required Privilege Level | 9327
- Release Information | 9327

Syntax

```
ipsec-vpn vpn-name;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
tunnel]
```

Description

Define IPsec name for VPN.

Options

vpn-name —Name of the IPsec.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

ipv4-address (Steering Path)

IN THIS SECTION

- Syntax | 9328
- Hierarchy Level | 9328
- Description | 9328

- Options | 9328
- Required Privilege Level | 9328
- Release Information | 9329

Syntax

```
ipv4-address ipv4-address;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name steering path]
[edit services pcef pcc-action-profiles profile-name]
```

Description

Specify the IPv4 address of a third-party server to which the PCC action profile steers HTTP traffic for applying services. If you configure this, the PCC action profile can only be used in PCC rules that match only HTTP-based applications and all flows.

Options

ipv4-address *ipv4-address*	Use the specified IPv4 address of the server.

Required Privilege Level

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

- *Understanding Application-Aware Policy Control for Subscriber Management*
- *Configuring Policy and Charging Control Action Profiles for Subscriber Management*
- Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment
- Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

ipv4-dscp

IN THIS SECTION

- Syntax | 9329
- Hierarchy Level | 9330
- Description | 9330
- Options | 9330
- Required Privilege Level | 9330

Syntax

```
ipv4-dscp number
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

Specify the value used as the DiffServ code point (DSCP) value for all Precision Time Protocol (PTP) IPv4 packets originated by the router.

Options

number—DiffServ code point (DSCP) value for all PTP IPv4 packets originated by the router.

- **Range:** 0 through 63
- **Default:** 56

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

IEEE 1588v2 Precision Timing Protocol (PTP)

Precision Time Protocol Overview

ipv4 (Family MPLS)

IN THIS SECTION

- Syntax | **9331**
- Hierarchy Level | **9331**
- Description | **9331**

- Options | 9332
- Required Privilege Level | 9332
- Release Information | 9332

Syntax

```
ipv4 {
    destination-address ip-address {
        except;
    }
    destination-prefix-list destination-prefix-list-name {
        except;
    }
    protocol protocol {
        (source-port | source-port-except);
        (destination-port | destination-port-except);
    }
    source-address ip-address {
        except;
    }
    source-prefix-list source-prefix-list-name {
        except;
    }
}
```

Hierarchy Level

```
[edit firewall family mpls filter name term name from ip-version]
```

Description

Define Layer 3 and Layer 4 match items to match IPv4 packets for IP-based filtering of MPLS traffic.

Options

destination-address *ip-address*	Match MPLS traffic with the specified IPv4 destination address.
destination-prefix-list *destination-prefix-list-name*	Match MPLS traffic with the specified IPv4 destination prefixes. The prefix-list is defined under the `[edit policy-options prefix-list prefix-list-name]` hierarchy level.
protocol *protocol*	Specify one or a range of inner IPv4 protocols for IP-based filtering of MPLS traffic.
source-address *ip-address*	Match MPLS traffic with the specified IPv4 source address.
source-prefix-list *source-prefix-list-name*	Match MPLS traffic with the specified IPv4 source prefixes. The prefix-list is defined under the `[edit policy-options prefix-list prefix-list-name]` hierarchy level.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Understanding IP-Based Filtering and Selective Port Mirroring of MPLS Traffic

ipv4-flow-table-size

IN THIS SECTION

- Syntax | 9333
- Hierarchy Level | 9333

- Description | 9333
- Options | 9334
- Required Privilege Level | 9334

Syntax

```
ipv4-flow-table-size units;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services flow-table-size]
```

Description

Configure the size of the IPv4 flow table in units of 256K entries.

> **NOTE**: Prior to Junos OS Release 16.1R1 and 15.1F2, any changes in the configured size of the flow table initiates an automatic reboot of the FPC, and we recommend that you run this command in a maintenance window.

Starting with Junos OS Release 16.1R1 and 15.1F2, by default, the software allocates 1K entries for IPv4 flow tables. To allocate fifteen 256K IPv4 flow tables, the former default, you can enter this configuration from the `[edit]` hierarchy level:

```
[edit]
user@router# set chassis fpc slot-number inline-services flow-table-size ipv4-flow-table-size 15
```

Starting with Junos OS Release 17.3R1, the maximum supported flow table size varies by line card type:

Table 52: Maximum Flow Table Size for MX Series Routers

Line Card	Maximum Number of Units (1 unit =256K entries)
MPC2E, MPC3E, MPC4E	15
MPC5E and MPC6E	220
MPC7E , MPC9E, LC2103, and LC480	48
MPC8E	97
MPC10E and MPC11E	48
JNP10K-LC2301, JNP10K-LC4800, JNP10K-LC9600, and MX304-LMIC16BASE	80

Options

units Number of 256K flow entries available for the IPv4 flow table. 1 unit equals 256K flow entries.

- **Range:** 1 through the maximum number of units for the line card. See Table 52 on page 9334 for maximum number of units for line cards.
- **Default:** The absolute number of flow entries: 1024 (1K)—Starting with Junos OS Release 16.1R1 and 15.1F2
- **Default:** The absolute number of flow entries: 3,932,160 (3840K)—Prior to Junos OS Release 16.1R1 and 15.1F2

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.3R1	Starting with Junos OS Release 17.3R1, for LU-based platforms, the maximum number of units is 15. For XL-based platforms, the maximum is 220. For EA-based platforms, the maximum is 48 for MPC7E and MPC9E. For MPC8E, the maximum is 97.
16.1R1	Starting with Junos OS Release 16.1R1 and 15.1F2, by default, the software allocates 1K entries for IPv4 flow tables.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

ipv4-mask (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 9335
- Hierarchy Level | 9336
- Description | 9336
- Options | 9336
- Required Privilege Level | 9336
- Release Information | 9336

Syntax

```
ipv4-mask ipv4-mask;
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
```

Description

Configure the IPv4 mask to identify a byte of the IPv4 subscriber address that you want to modify in the HTTP header. You must also set the `ipv4-or-value` statement at the `[edit services hcm tag-rule rule-name term term-number then tag tag-name]` hierarchy level to specify the new value you want to put in the byte.

Options

ipv4-mask IPv4 mask. Specify **255** in the byte you want to modify and specify **0** in the bytes that you *do not* want to modify.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

ipv4-or-value (HTTP Header Enrichment) | **9337**

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

Configuring HCM Profiles and Assigning Tag Rules

Junos Web Aware HTTP Header Enrichment Overview

ipv4-or-value (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 9337
- Hierarchy Level | 9337
- Description | 9337
- Options | 9337
- Required Privilege Level | 9337
- Release Information | 9338

Syntax

```
ipv4-or-value ipv4-or-value;
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
```

Description

Configure the new IPv4 value for the byte you want to modify in the IPv4 subscriber address in the HTTP header. You must also set the `ipv4-mask` statement at the `[edit services hcm tag-rule rule-name term term-number then tag tag-name]` hierarchy level to clear the existing byte value.

Options

ipv4-or-value IPv4 value. Specify the new value in the byte you are modifying and specify 0 in all other bytes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

ipv4-mask (HTTP Header Enrichment) | 9335

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

Configuring HCM Profiles and Assigning Tag Rules

Junos Web Aware HTTP Header Enrichment Overview

ipv4-prefix

IN THIS SECTION

- Syntax | 9338
- Hierarchy Level | 9339
- Description | 9339
- Options | 9339
- Required Privilege Level | 9339
- Release Information | 9339

Syntax

```
ipv4-prefix {
    as as;
    router-id router-id;
    prefix prefix;
```

```
    system-id system-id;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options policy-statement policy-name term term-name from traffic-engineering],
[edit policy-options policy-statement policy-name term term-name from traffic-engineering]
```

Description

Configure filter options for a traffic engineering policy to filter traffic based on IPv4 prefix addresses. You can specify additional parameters, such as autonomous system (AS), prefix, router ID, and system ID for filtering IPv4 traffic. If you do not specify the additional parameters, the policy matches all IPv4-prefix network layer reachability information (NLRI) subtypes. You cannot apply these filters along with other NLRI filters.

Options

as *as*	Specify an AS to filter traffic.
router-id *router-id*	Specify an IP prefix to match the router-ID against.
prefix	Specify an IPv4 prefix to match against.
system-id *system-id*	Specify an ISO address for the node.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

policy-statement | 12736

show route table | 23789

ipv4 (Security IDP Signature Attack)

IN THIS SECTION

- Syntax | 9340
- Hierarchy Level | 9341
- Description | 9341
- Options | 9341
- Required Privilege Level | 9343
- Release Information | 9344

Syntax

```
ipv4 {
    checksum-validate        {
        match (equal | greater-than | less-than | not-equal | validate);
        value value;
    }
     destination (Security IP Headers Attack) {
        match (equal | greater-than | less-than | not-equal);
        value value;
    }
    identification (Security ICMP Headers) {
        match (equal | greater-than | less-than | not-equal);
        value value;
    }
    ihl {
        match (equal | greater-than | less-than | not-equal);
        value value;
    }
```

```
    ip-flags <(df | no-df)> <(mf | no-mf)> <(rb | no-rb)>;
     protocol  {
        match (equal | greater-than | less-than | not-equal);
        value value;
    }
    source {
        match (equal | greater-than | less-than | not-equal);
        value value;
    }
    tos {
        match (equal | greater-than | less-than | not-equal);
        value value;
    }
    total-length {
        match (equal | greater-than | less-than | not-equal);
        value value;
    }
    ttl  {
        match (equal | greater-than | less-than | not-equal);
        value value;
    }
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name  attack-type signature protocol]
```

Description

An IP header is header information at the beginning of an IP packet which contains information about IP version, source IP address, destination IP address, time-to-live, etc. Allow IDP to match the IP header information for the signature attack.

Options

The remaining statements are explained separately. See CLI Explorer.

identification Specify a unique value used by the destination system to reassemble a fragmented packet.

- Syntax:
  - `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
  - `value` *`identification-value`*—Match a decimal value.
- **Range:** 0 through 65,535

ihl

Specify the IPv4 header length in words.

- Syntax:
  - `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
  - `value` *`identification-value`*—Match a decimal value.
- **Range:** 0 through 15

ip-flags

Specify that IDP looks for a pattern match whether or not the IP flag is set.

- Syntax:
  - `df` | `no-df`—When set, the df (Don't Fragment) indicates that the packet cannot be fragmented for transmission. When unset, it indicates that the packet can be fragmented.
  - `mf` | `no-mf`—When set, the mf (More Fragments) indicates that the packet contains more fragments. When unset, it indicates that no more fragments remain.
  - `rb` | `no-rb`—When set, the rb (Reserved Bit) indicates that the bit is reserved.

protocol

Specify the Transport Layer protocol number.

- Syntax:
  - `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
  - `value` *`transport-layer-protocol-id`*—Match the Transport Layer protocol ID.

source

Specify the IP address or hostname of the attacking device.

- Syntax:
  - `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
  - `value` *`ip-address-or-hostname`*—Match an IP address or a hostname.

tos

Specify the type of service.

- Syntax:
  - `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
  - `value` *`type-of-service-in-decimal`*—The following service types are available:
    - 0000—Default
    - 0001—Minimize Cost
    - 0002—Maximize Reliability
    - 0003—Maximize Throughput
    - 0004—Minimize Delay
    - 0005—Maximize Security

total-length Specify the number of bytes in the packet, including all header fields and the data payload.

- Syntax:
  - `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
  - `value` *`total-length-of-ip-datagram`*—Length of the IP datagram.
- **Range:** 0 through 65,535

ttl Specify the time-to-live (TTL) value of the packet. This value represents the number of routers the packet can pass through. Each router that processes the packet decrements the TTL by 1; when the TTL reaches 0, the packet is discarded.

- Syntax:
  - `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
  - `value` *`time-to-live`*—The time-to-live value.
- **Range:** 0 through 255

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

ipv4-template

IN THIS SECTION

- Syntax | 9344
- Hierarchy Level | 9344
- Description | 9344
- Required Privilege Level | 9345
- Release Information | 9345

Syntax

```
ipv4-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Specify the version 9 or IPFIX template properties for one of the following:

- Template for monitoring IPv4 flows.
- Template for inline monitoring an MPLS-over-UDP flow that is carried between IPv4 endpoints on PTX Series routers. This monitoring looks past the tunnel header to report the inner payload of the packets. To use the template for MPLS-over-UDP flows, you must also configure `tunnel-observation`

`mpls-over-udp` at the `[edit services flow-monitoring (version 9 | version-ipfix) template template-name]` hierarchy level.

> **NOTE**: For an MPLS-over-UDP flow that is encapsulated in an RSVP-TE LSP, configure `mpls-ipvx-template` in Junos OS Release 18.1 or `mpls-template` starting in Junos OS 18.2R1 at the `[edit services flow-monitoring (version 9 | version-ipfix) template template-name]` hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level added in Junos OS Release 10.2.

RELATED DOCUMENTATION

- *Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates*
- *Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers*
- *Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices*
- *Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers*

ipv4-template (Services)

IN THIS SECTION

Syntax | 9346

- Hierarchy Level | 9346
- Description | 9346
- Required Privilege Level | 9346
- Release Information | 9346

Syntax

```
ipv4-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

Specify that the flow monitoring version 9 template is used only for IPv4 records.

Required Privilege Level

services—To view this in the configuration.

services-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

Understanding Interfaces

ipv6-address (Steering Path)

IN THIS SECTION

- Syntax | 9347
- Hierarchy Level | 9347
- Description | 9347
- Options | 9347
- Required Privilege Level | 9347
- Release Information | 9348

Syntax

```
ipv6-address ipv6-address;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name steering path]
[edit services pcef pcc-action-profiles profile-name]
```

Description

Specify the IPv6 address of a third-party server to which the PCC action profile steers HTTP traffic for applying services. If you configure this, the PCC action profile can only be used in PCC rules that match only HTTP-based applications and all flows.

Options

ipv6-address *ipv6-address*	Use the specified IPv6 address of the server.

Required Privilege Level

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Understanding Application-Aware Policy Control for Subscriber Management

Configuring Policy and Charging Control Action Profiles for Subscriber Management

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

ipv6-extended

IN THIS SECTION

- Syntax | 9349
- Hierarchy Level | 9349
- Description | 9349
- Required Privilege Level | 9349
- Release Information | 9349

Syntax

```
ipv6-extended:type-and-subtype value:<IPv6 unicast address>:operator-defined local values
```

Hierarchy Level

```
[edit policy-options community community-name members]
```

Description

Defines the extended format of the IPv6 unicast address in a policy match.

Required Privilege Level

policy-options—To view and add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release and Junos OS Evolved Release 23.1.

ipv6-extended-attrib

IN THIS SECTION

- Syntax | 9350
- Hierarchy Level | 9350
- Description | 9350
- Required Privilege Level | 9350

Syntax

```
ipv6-extended-attrib;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services ipv6 flow-table-size]
```

Description

Enable the inclusion of element ID, 54, fragmentIdentification, and element ID, 64, ipv6ExtensionHeaders, in IPFIX flow templates that are exported to the flow collector

> **NOTE**: Collection of IPv4 fragmentation IDs occurs automatically without having to configure this setting explicitly.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

ipv6-extension-header

IN THIS SECTION

- Syntax | 9351
- Hierarchy Level | 9352

- Description | 9352
- Options | 9352
- Required Privilege Level | 9352
- Release Information | 9352

Syntax

```
ipv6-extension-header {
    AH-header;
    ESP-header
    HIP-header;
    destination-header {
        ILNP-nonce-option;
        home-address-option;
        line-identification-option;
        tunnel-encapsulation-limit-option;
        user-defined-option-type low | <to high>;
    }
    fragment-header;
    hop-by-hop-header {
        CALIPSO-option;
        RPL-option;
        SFM-DPD-option;
        jumbo-payload-option;
        quick-start-option;
        router-alert-option;
        user-defined-option-type low | <to high>;
    }
    mobility-header;
    no-next-header;
    routing-header;
    shim6-header
    user-defined-option-type low | <to high>;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name ip]
```

Description

Define the IPv6 extension header for the intrusion detection service (IDS).

Options

AH-header	Enable the IPv6 Authentication Header screen option.
ESP-header	Enable the IPv6 Encapsulating Security Payload header screen option.
HIP-header	Enable the IPv6 Host Identify Protocol header screen option.
fragment-header	Enable the IPv6 fragment header screen option.
mobility-header	Enable the IPv6 mobility header screen option.
no-next-header	Enable the IPv6 no next header screen option.
routing-header	Enable the IPv6 routing header screen option.
shim6-header	Enable the IPv6 shim header screen option.
user-defined-header-type *low* | <to *high*>	Define the type of header range. • **Range:** 0 through 255.

The remaining statements are explained separately. See the CLI Explorer.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Understanding IPv6 Support for Screens

hop-by-hop-header | 8132

destination-header | 5419

ipv6-extension-header (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 9353
- Hierarchy Level | 9354
- Description | 9354
- Options | 9354
- Required Privilege Level | 9355
- Release Information | 9355

Syntax

```
ipv6-extension-header {
    AH-header;
    ESP-header;
    fragment-header;
    hop-by-hop-header {
        CALIPSO-option;
        jumbo-payload-option;
        quick-start-option;
        router-alert-option;
        RPL-option;
        SFM-DPD-option;
        user-defined-option-type <type-low> to <type-high>;
    }
    mobility-header;
```

```
    routing-header;
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop IP packets that have the configured IPv6 extension header values.

Options

ah-header Authentication Header extension header

esp-header Encapsulating Security Payload extension header

fragment-header Fragment Header extension header

hop-by-hop-header The specified Hop-by-Hop option:

- **CALIPSO-option** Common Architecture Label IPv6 Security Option
- **jumbo-payload-option** IPv6 jumbo payload option
- **quick-start-option** IPv6 quick start option
- **router-alert-option** IPv6 router alert option
- **RPL-option** Routing Protocol for Low-Power and Lossy Networks option
- **SFM-DPD-option** Simplified Muliticast Forwarding IPv6 Duplicate Packet Detection option
- **user-defined-option-type** *type-low* **to** *type-high* A range of header types
  - **Range:** 1 through 255.

mobility-header Mobility Header extension header

routing-header Routing Header extension header

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

ipv6-extension-header-limit

IN THIS SECTION

- Syntax | 9355
- Hierarchy Level | 9356
- Description | 9356
- Options | 9356
- Required Privilege Level | 9356
- Release Information | 9356

Syntax

```
ipv6-extension-header-limit limit;
```

Hierarchy Level

```
[edit security screen ids-option screen-name ip]
```

Description

Define the IPv6 extension header number limit for screen options. The screen blocks packets that have more than the defined number of extension headers.

Options

limit Set the number of IPv6 extension headers that can pass through the screen.

- **Range:** 0 through 32.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Understanding IPv6 Support for Screens

ipv6-extension-header | **9350**

ipv6 (Family MPLS)

IN THIS SECTION

- Syntax | 9357
- Hierarchy Level | 9358
- Description | 9358
- Options | 9358
- Required Privilege Level | 9358
- Release Information | 9358

Syntax

```
ipv6 {
    destination-address destination-ip-address {
        except;
    }
    destination-prefix-list prefix-list-name {
        except;
    }
    protocol protocol {
        (source-port | source-port-except);
        (destination-port | destination-port-except);
    }
    source-address ip-address {
        except;
    }
    source-prefix-list source-prefix-list-name {
        except;
    }
}
```

Hierarchy Level

```
[edit firewall family mpls filter name term name from ip-version]
```

Description

Define Layer 3 and Layer 4 match items to match IPv6 packets for IP-based filtering of MPLS traffic.

Options

destination-address ***ip-address***	Match MPLS traffic with the specified IPv6 destination address.
destination-prefix-list ***destination-prefix-list-name***	Match MPLS traffic with the specified list of IPv6 destination prefixes. The prefix-list is defined under the `[edit policy-options prefix-list prefix-list-name]` hierarchy level. You must configure separate prefix-lists for IPv4 and IPv6 addresses.
protocol ***protocol***	Specify one or a range of inner IPv6 next header for IP-based filtering of MPLS traffic.
source-address ***ip-address***	Match MPLS traffic with the specified IPv6 source address.
source-prefix-list ***source-prefix-list-name***	Match MPLS traffic with the specified IPv6 source prefixes. The prefix-list is defined under the `[edit policy-options prefix-list prefix-list-name]` hierarchy level. You must configure separate prefix-lists for IPv4 and IPv6 addresses.

Required Privilege Level

firewall

Release Information

Statement introduced in Junos OS Release 18.4R1.

ipv6-flow-label

IN THIS SECTION

- Syntax | 9359
- Hierarchy Level | 9359
- Hierarchy Level | 9359
- Description | 9359
- Required Privilege Level | 9360
- Release Information | 9360

Syntax

```
ipv6-flow-label;
```

Hierarchy Level

For QFX Series:

```
[edit forwarding-options enhanced-hash-key inet6]
```

Hierarchy Level

For PTX Series:

```
[edit forwarding-options hash-key family inet6 layer-3]
```

Description

Enable IPv6 packet flow labels for hash calculations and load balancing based on the flow label of the IPv6 header.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D30 on QFX10002 and 10008 switches.

Statement introduced in Junos OS Release 18.3R1 for PTX series routers.

RELATED DOCUMENTATION

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)

show forwarding-options enhanced-hash-key | 20086

ipv6-flow-table-size

IN THIS SECTION

- Syntax | 9360
- Hierarchy Level | 9361
- Description | 9361
- Options | 9361
- Required Privilege Level | 9361

Syntax

```
ipv6-flow-table-size units;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services ipv6 flow-table-size]
```

Description

Configure the size of the IPv6 flow table in units of 256K entries.

NOTE: Prior to Junos OS Release 15.1F2, any changes in the configured size of the flow table initiates an automatic reboot of the FPC.

NOTE: The recommended flow table size is 4 so that it can scale up to 4x256K flows, which is 1M. You can configure more, however, the system will issue a warning message.

NOTE: Starting with Junos OS Release 17.3R1, the maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 256 on MPC5Es and MPC6Es with with 4 GB DDR3 memory or higher. The maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 245 on MPC5Es and MPC6Es with DDR3 memory lower than 4 GB.

Options

units Number of 256K flow entries available for the IPv6 flow table.

- **Range:** 1 through 245
- **Default:** If number of units is not specified, 1024 flow entries are allocated for IPv6.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

ipv6-malformed-header

IN THIS SECTION

- Syntax | 9362
- Hierarchy Level | 9362
- Description | 9362
- Options | 9362
- Required Privilege Level | 9362
- Release Information | 9363

Syntax

```
ipv6-malformed-header;
```

Hierarchy Level

```
[edit security screen ids-option screen-name ip]
```

Description

Enable the IPv6 malformed header intrusion detection service (IDS) option. The screen will not pass any communication using the subnet 2001:db8 :: / 32.

Options

This has no options.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Understanding IPv6 Support for Screens

ipv6-extension-header | 9350

ipv6-mask (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 9363
- Hierarchy Level | 9363
- Description | 9364
- Options | 9364
- Required Privilege Level | 9364
- Release Information | 9364

Syntax

```
ipv6-mask ipv6-mask;
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
```

Description

Configure the IPv6 mask to identify a byte of the IPv6 subscriber address that you want to modify in the HTTP header. You must also set the `ipv6-or-value` statement at the `[edit services hcm tag-rule` *`rule-name`* `term` *`term-number`* `then tag` *`tag-name`*`]` hierarchy level to specify the new value you want to put in the byte.

Options

ipv6-mask IPv6 mask. Specify **ff** in the byte you want to modify and specify **0** in the bytes that you *do not* want to modify.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- ipv6-or-value (HTTP Header Enrichment) | **9366**
- *Configuring HTTP Header Enrichment Overview*
- *Configuring Tag Rules*
- *Configuring HCM Profiles and Assigning Tag Rules*
- *Junos Web Aware HTTP Header Enrichment Overview*

ipv6-multicast-interfaces

IN THIS SECTION

- Syntax | **9365**
- Hierarchy Level | **9365**

- Description | 9365
- Options | 9365
- Required Privilege Level | 9365
- Release Information | 9366

Syntax

```
ipv6-multicast-interfaces (all | interface-name) {
    disable;
}
```

Hierarchy Level

```
[edit services nat],
[edit services softwire]
```

Description

Enable multicast filters on Ethernet interfaces when IPv6 NAT is used for neighbor discovery.

Options

all—Enable filters on all interfaces.

disable—Disable filters on the specified interfaces.

interface-name—Enable filters on a specific interface only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Configuring IPv6 Multicast Interfaces

ipv6-or-value (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 9366
- Hierarchy Level | 9366
- Description | 9367
- Options | 9367
- Required Privilege Level | 9367
- Release Information | 9367

Syntax

```
ipv6-or-value ipv6-or-value;
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
```

Description

Configure the new IPv6 value for the byte you want to modify in the IPv6 subscriber address in the HTTP header. You must also set the `ipv6-mask` statement at the `[edit services hcm tag-rule` *`rule-name`* `term` *`term-number`* `then tag` *`tag-name`*`]` hierarchy level to clear the existing byte value.

Options

ipv6-or-value IPv6 value. Specify the new value in the byte you are modifying and specify 0 in all other bytes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

ipv6-mask (HTTP Header Enrichment) | **9363**

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

Configuring HCM Profiles and Assigning Tag Rules

Junos Web Aware HTTP Header Enrichment Overview

ipv6-origin

IN THIS SECTION

- Syntax | **9368**
- Hierarchy Level | **9368**

- Description | 9368
- Required Privilege Level | 9368
- Release Information | 9368

Syntax

```
ipv6-origin:<IPv6 unicast address>:operator-defined local values
```

Hierarchy Level

```
[edit policy-options community community-name members]
```

Description

Defines the source of the IPv6 unicast address in a policy match and allow routes of the specified community to be advertised.

Required Privilege Level

policy-options—To view and add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release and Junos OS Evolved Release 23.1.

ipv6-prefix

IN THIS SECTION

- Syntax | 9369

- Hierarchy Level | 9369
- Description | 9369
- Options | 9369
- Required Privilege Level | 9370
- Release Information | 9370

Syntax

```
ipv6-prefix {
    as as;
    router-id router-id;
    prefix prefix;
    system-id system-id;
}
```

Hierarchy Level

```
[edit logical-systems name policy-options policy-statement name from traffic-engineering],
[edit policy-options policy-statement name from traffic-engineering]
```

Description

Configure filter options for a traffic engineering policy to filter traffic based on IPv6 prefix addresses. You can specify additional parameters, such as autonomous system (AS), prefix, router ID, and system ID for filtering IPv6 traffic.

Options

as *as*	Specify an AS to filter traffic
router-id *router-id*	Specify an IP prefix to match the router-ID against.

prefix	Specify an IPv6 prefix to match against.
system-id *system-id*	Specify an ISO address for the node.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Link-State Distribution using SRv6

ipv6-source-guard

IN THIS SECTION

- Syntax | 9370
- Hierarchy Level | 9371
- Description | 9371
- Default | 9371
- Required Privilege Level | 9371
- Release Information | 9371

Syntax

```
ipv6-source-guard;
```

Hierarchy Level

- For platforms with Enhanced Layer 2 Software (ELS):

```
[edit vlans vlan-name forwarding-options dhcp-security];
```

- For platforms without ELS:

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name)]
```

Description

Perform IPv6 source guard checking on packets sent from access interfaces. Validate source IP addresses and source MAC addresses on all VLANs or on the specified VLAN. Forward packets with valid addresses and drop those with invalid addresses.

> **NOTE**: If you configure the `ipv6-source-guard` statement at the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level, DHCPv6 snooping is automatically enabled for the specified VLAN.
>
> If you configure the `ipv6-source-guard` statement at the `[edit ethernet-switching-options secure-access-port vlan vlan-name]` hierarchy level, you must also enable DHCPv6 snooping for the specified VLAN.

Default

Disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D20.

Support at the `[edit ethernet-switching-options secure-access-port vlan (all | vlan-name)]` hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing

Configuring IP Source Guard (ELS)

ipv6-source-guard-sessions

IN THIS SECTION

- Syntax | 9372
- Hierarchy Level | 9372
- Description | 9372
- Default | 9373
- Options | 9373
- Required Privilege Level | 9373
- Release Information | 9373

Syntax

```
ipv6-source-guard-sessions {
    max-number max-number;
}
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port]
```

Description

Specify the maximum number of IPv6 source guard sessions for TCAM space provisioning.

> **NOTE**: After setting or changing the maximum number of IPv6 source guard sessions and committing the configuration, you must reboot the switch for the configuration to take effect.

Default

Disabled.

Options

max-number *max-number* The maximum number of IPv6 source guard sessions.

- **Range:** 50 through 300.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing

Configuring IP Source Guard (ELS)

ipv6-target

IN THIS SECTION

- Syntax | 9374
- Hierarchy Level | 9374
- Description | 9374
- Required Privilege Level | 9374
- Release Information | 9374

Syntax

```
ipv6-target:<IPv6 unicast address>:operator-defined local values
```

Hierarchy Level

```
[edit policy-options community community-name members]
```

Description

Defines the VPN IPv6 target unicast address used in a policy match and allow routes of the specified community to be advertised.

Required Privilege Level

policy-options—To view and add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release and Junos OS Evolved Release 23.1.

ipv6-template

IN THIS SECTION

- Syntax | 9375
- Hierarchy Level | 9375
- Description | 9375
- Required Privilege Level | 9375
- Release Information | 9375

Syntax

```
ipv6-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Specify that the flow aggregation version 9 or IPFIX template is used only for IPv6 records.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Support at the [edit services flow-monitoring version-ipfix template *template-name*] hierarchy level added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates

Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers

ipv6-template (Services)

IN THIS SECTION

- Syntax | 9376
- Hierarchy Level | 9376
- Description | 9376
- Required Privilege Level | 9377
- Release Information | 9377

Syntax

```
ipv6-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

Specify that the flow monitoring version 9 template is used only for IPv6 records.

Required Privilege Level

services—To view this in the configuration.

services-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

Understanding Interfaces

ipv6-tunneling

IN THIS SECTION

- Syntax | 9377
- Hierarchy Level | 9378
- Description | 9378
- Required Privilege Level | 9378
- Release Information | 9378

Syntax

```
ipv6-tunneling;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Allow IPv6 routes to be resolved over an MPLS network by converting LDP and RSVP routes stored in the inet.3 routing table to IPv4-mapped IPv6 addresses and then copying them into the inet6.3 routing table. This routing table can be used to resolve next hops for both inet6 and inet6-vpn routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks

ipvacl-firewall-l2-profile

IN THIS SECTION

- Syntax | 9379
- Hierarchy Level | 9379
- Description | 9379
- Required Privilege Level | 9379
- Release Information | 9379

Syntax

```
ipvacl-firewall-l2-profile
```

Hierarchy Level

```
[edit chassis]
```

Description

IPACL group is used when a filter is attached to the input of an L2 interface. IPACL group has both L2 qualifiers and L3 qualifiers. Set this configuration statement to have only L2 qualifiers for IPACL/IVACL.

> **NOTE**: When this configuration statement is enabled, only L2 match conditions should be used for IPACL/IVACL filters.

Required Privilege Level

interface-control

Release Information

Statement introduced in Junos OS Release 22.2R1

RELATED DOCUMENTATION

show pfe filter hw summary | 22983

iq-policing-filter

IN THIS SECTION

- Syntax | 9380
- Hierarchy Level | 9380
- Description | 9380
- Options | 9380
- Required Privilege Level | 9381
- Release Information | 9381

Syntax

```
iq-policing-filter filter-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number family bridge|inet|vpls ]
```

Description

For MPCs that support ingress queuing, attaches a filter to an interface that polices traffic before assigning to ingress queues. This enables rate limiting of traffic before the traffic is assigned to ingress queues.

The `iq-policing-filter` statement is available for the following protocol families: `bridge`, `inet`, and `vpls`.

Options

filter-name Name of the `iq-policing-filter` filter

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

ir-mode

IN THIS SECTION

- Syntax | 9381
- Hierarchy Level | 9381
- Description | 9382
- Default | 9382
- Options | 9382
- Required Privilege Level | 9382
- Release Information | 9383

Syntax

```
ir-mode (IR | R);
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

Configure the license mode of the specified enhanced MPC in an MPC slot as IR or R. Setting the license mode enables you to distinguish between an MPC with an IR license and an MPC with an R license after the MPC is installed on the router.

> **NOTE**: The license mode settings are used only to provide information. You cannot set or alter the license of the MPC when you configure the mode.

Default

The default IR mode is Base. You do not set this configuration statement if Base is the mode of your license. Base mode includes the following features:

- All Layer 2, Layer 2.5, and Layer 3 features.
- Up to 32 Layer 3 routing instances of the virtual routing and forwarding (VRF) instance type.
- Up to 2 million routes in the forwarding information base (FIB), provided there is hardware support. (FIB is also known as forwarding table.)
- Up to 6 million routes in the routing information base (RIB), also known as routing table.

Options

`IR` Configure the license mode IR for an MPC installed in a specified MPC slot. Includes the following features:

- All Layer 2, Layer 2.5, and Layer 3 features.
- Up to 32 Layer 3 routing instances of the virtual routing and forwarding (VRF) instance type.

`R` Configure the license mode R for an MPC installed in a specified MPC slot. Includes full-scale Layer 2, Layer 2.5, and Layer 3 features. Scale is determined by the hardware capabilities.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

- License Modes for Enhanced MPCs Overview
- Configuring the License Mode for Specific Enhanced MPCs on MX Series Routers
- Viewing the License Mode for MPC Cards on MX Series Routers

iracl-firewall-ipv4-profile

IN THIS SECTION

- Syntax | 9383
- Hierarchy Level | 9383
- Description | 9383
- Required Privilege Level | 9384
- Release Information | 9384

Syntax

```
iracl-firewall-ipv4-profile
```

Hierarchy Level

```
[edit chassis]
```

Description

Optimize fixed QSET profiles.

IRACL group has IPv6 source and destination address qualifiers by default which is used for matching IPv4 source and destination addresses as well. Set this configuration statement to have IPv4 profile only.

> **NOTE**: When this configuration statement is enabled, `source-address` and `destination-address` match should not be used for IPv6 IRACL filters.

Required Privilege Level

interface-control

Release Information

Statement introduced in Junos OS Release 22.2R1

RELATED DOCUMENTATION

show pfe filter hw summary | 22983

irb

IN THIS SECTION

- Syntax | 9384
- Hierarchy Level | 9385
- Description | 9385
- Required Privilege Level | 9385
- Release Information | 9385

Syntax

```
irb {
    unit logical-unit-number {
        classifiers {
```

```
            type (classifier-name | default);
        }
        rewrite-rules {
            dscp (rewrite-name | default);
            dscp-ipv6 (rewrite-name | default);
            exp (rewrite-name | default)protocol protocol-types;
            ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
            inet-precedence (rewrite-name | default);
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service interfaces]
```

Description

On the MX Series routers and EX Series switches, you can apply classifiers or rewrite rules to an integrated bridging and routing (IRB) interface. All types of classifiers and rewrite rules are allowed. These classifiers and rewrite rules are independent of others configured on the MX Series router and on EX Series switches.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

CoS Features and Limitations on MX Series Routers

irb (Easy EVPN LAG Configuration)

IN THIS SECTION

- Syntax | 9386
- Hierarchy Level | 9386
- Description | 9387
- Options | 9387
- Required Privilege Level | 9389
- Release Information | 9389

Syntax

```
irb irb-instance {
    apply-config-groups;
    instance {
      mac-vrf-instance instance-id;
      vrf-instance instance-id;
    }
    no-dhcp-relay;
    subnet-address {
      inet (inet-addr | [inet-addr1 inet-addr2 ...]);
      inet6 (inet6-addr | [inet6-addr1 inet6-addr2 ...]);
      no-irb-address-auto-derive;
      virtual-gateway-v4-address virtual-gateway-v4-address;
      virtual-gateway-v6-address virtual-gateway-v6-address;
    }
    use-anycast-address;
    vlan-id vlan-num;
};
```

Hierarchy Level

```
[edit services evpn evpn-vxlan]
```

Description

Specify parameters to generate the configuration statements for the integrated routing and bridging (IRB) instances on a provider edge (PE) device in an Ethernet VPN (EVPN) fabric with Virtual Extensible LAN (VXLAN) encapsulation.

You use statements at the `[edit services evpn]` hierarchy level to simplify configuring an EVPN fabric with Ethernet segment identifier (ESI) link aggregation groups (LAGs) in the topology. The feature you enable with statements in this hierarchy is an easy EVPN LAG configuration feature, also called EZ-LAG.

You only need to configure a few statements to provide the parameters required for an EVPN core fabric with multihomed or single-homed customer edge (CE) devices (or end server devices). When you commit the `[edit services evpn]` configuration statements, the device uses a commit script to generate the corresponding standard EVPN fabric configuration statements.

The commit script uses the parameters at the `[edit services evpn evpn-vxlan irb]` hierarchy level to generate configuration statements for:

- The IRB interfaces for the MAC-VRF and virtual routing and forwarding (VRF) instances in the EVPN fabric, including:
  - A virtual gateway address configuration for the IRB instance

    > **NOTE**: By default, the commit script generates configuration statements that include configuring the `virtual-gateway-accept-data` option at the `[edit routing-instances` *`name`* `interfaces irb unit` *`number`*`]` hierarchy level with the virtual gateway address configuration. With this option, the virtual gateway can accept ping packets.

  - An anycast gateway configuration for the IRB instance
- The VLANs associated with the IRB interfaces
- The aggregated Ethernet interfaces to single-homed or multihomed (ESI-LAG) connected servers

Options

(Required) ***irb-instance***

Specify a string that identifies an IRB instance and the parameters associated with it for commit script processing.

Choose *irb-instance* values that help interpret what you will see in the generated configuration. For example, the command below uses *irb-instance* "irb_10" to provide IRB parameters associated with VLAN ID 10. In this case, because the command doesn't

include a `mac-vrf-instance` option, these parameters apply to the default MAC-VRF instance (`mac-vrf-instance` *`instance-id`* 1):

```
set services evpn evpn-vxlan irb irb_10 vlan-id 10 subnet-address inet 10.1.1.1/24
```

apply-config-groups *config-groups*

Apply the specified configuration groups to the configuration the commit script generates from the statements at this `irb` configuration hierarchy level.

instance (mac-vrf-instance | vrf-instance) *instance-id*

Specify the MAC-VRF instance ID or VRF instance ID associated with this IRB instance.

If you don't configure a MAC-VRF *instance-id* at this level, the commit script associates the IRB instance with the default MAC-VRF instance. The default MAC-VRF instance corresponds to *instance-id* 1 and instance name __SERVICES_EVPN_EVPN_VXLAN_MAC_VRF_1.

If you configure a VRF instance at this hierarchy level, the commit script generates configuration statements to create the VRF instance and associates this IRB instance with that VRF instance. The commit script uses VRF instance name _SERVICES_EVPN_EVPN_VXLAN_VRF_*instance-id* in the generated configuration. If you don't configure a VRF *instance-id* at this level, the commit script associates the IRB instance with the default VRF instance, which corresponds to the default routing table inet.0 (or inet6.0).

- **Range for *instance-id*:** 1 through 255

no-dhcp-relay

Don't include this IRB instance in the generated DHCP Relay configuration when you also configure "dhcp-relay" on page 5638 options at the `[edit services evpn evpn-vxlan]` hierarchy level.

use-anycast-address

Configure the IRB instance with an anycast gateway address.

When you set this option, you must also configure the `[edit services evpn global-parameters] anycast-mac` *`mac-address`* statement. The commit script generates an IRB anycast gateway configuration with that anycast MAC address.

vlan-id *vlan-num*

Specify a VLAN by VLAN ID number associated with the IRB interface. The commit script uses this parameter to generate the statements that configure the IRB interface with a matching logical unit.

You must configure at least one VLAN ID for each IRB instance.

The remaining statements are explained in detail on separate pages.

See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Easy EVPN LAG (EZ-LAG) Configuration

irb (Interfaces)

IN THIS SECTION

- Syntax | 9389
- Hierarchy Level | 9393
- Description | 9393
- Required Privilege Level | 9393
- Release Information | 9394

Syntax

```
irb {
    accounting-profile name;
    arp-l2-validate;
    description  text;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    mtu bytes;
```

```
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        enhanced-convergence;
        disable;
        encapsulation type;
        family inet {
            accounting {
                destination-class-usage;
                source-class-usage {
                    input;
                    output;
                }
            }
                address ipv4-address {
                arp ip-address (mac | multicast-mac) mac-address <publish>;
                broadcast address;
                preferred;
                primary;
                vrrp-group group-number {
                    (accept-data | no-accept-data);
                    advertise-interval seconds;
                    advertisements-threshold number;
                    authentication-key key;
                    authentication-type authentication;
                    fast-interval milliseconds;
                    (preempt | no-preempt) {
                        hold-time seconds;
                    }
                    priority number;
                    track {
                        interface interface-name {
                            bandwidth-threshold bandwidth;
                            priority-cost number;
                        }
                        priority-hold-time seconds;
                        route ip-address/mask routing-instance instance-name priority-cost cost;
```

```
                }
                virtual-address [ addresses ];
                vrrp-inherit-from {
                    active-group group-number;
                    active-interface interface-name;
                }
            }
        }
            filter {
            input filter-name;
            output filter-name;
        }
            mtu bytes;
        no-neighbor-learn;
        no-redirects;
        primary;
        rpf-check {
            fail-filter filter-name;
            mode {
                loose;
            }
        }
        targeted-broadcast {
            forward-and-send-to-re;
            forward-only;
        }
    }
    family inet6 {
        accounting {
            destination-class-usage;
            source-class-usage {
                input;
                output;
            }
        }
            address address {
            eui-64;
            ndp ip-address (mac | multicast-mac) mac-address <publish>;
            preferred;
            primary;
            vrrp-inet6-group group-id {
                accept-data | no-accept-data;
                advertisements-threshold number;
```

```
                authentication-key key;
                authentication-type authentication;
                fast-interval milliseconds;
                inet6-advertise-interval milliseconds;
                preempt | no-preempt {
                    hold-time seconds;
                }
                priority number;
                track {
                    interface interface-name {
                        bandwidth-threshold bandwidth priority-cost number;
                        priority-cost number;
                    }
                    priority-hold-time seconds;
                    route ip-address/mask routing-instance instance-name priority-cost cost;
                }
                virtual-inet6-address [addresses];
                virtual-link-local-address ipv6-address;
                vrrp-inherit-from {
                    active-group group-number;
                    active-interface interface-name;
                }
            }
        }
        (dad-disable | no-dad-disable);
            filter {
            input filter-name;
            output filter-name;
        }
            mtu bytes;
        nd6-stale-time seconds;
        no-neighbor-learn;
        no-redirects;
        policer {
            input policer-name;
            output policer-name;
        }
        rpf-check {
            fail-filter filter-name;
            mode {
                loose;
            }
        }
```

```
        }
        family iso {
            address interface-address;
                mtu bytes;
        }
        family mpls {
                filter {
                input filter-name;
                output filter-name;
            }
                mtu bytes;
            policer {
                input policer-name;
                output policer-name;
            }
        }
        native-inner-vlan-id vlan-id;
        no-auto-virtual-gateway-esi;
        proxy-arp (restricted | unrestricted);
        (traps | no-traps);
         identifier;
        vlan-id-list [vlan-id's];
        vlan-id-range [vlan-id-range];
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure the properties of a specific integrated bridging and routing (IRB) interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R2.

`irb` option introduced in Junos OS Release 13.2 for the QFX Series.

irb-symmetric-routing

IN THIS SECTION

- Syntax | 9394
- Hierarchy Level | 9394
- Description | 9395
- Default | 9395
- Options | 9395
- Required Privilege Level | 9395
- Release Information | 9395

Syntax

```
irb-symmetric-routing {
    vni vni;
}
```

Hierarchy Level

```
[edit routing-instances name protocols evpn]
[edit routing-instances name protocols evpn interconnect]
```

Description

Enable the symmetric IRB model of routing with EVPN Type 2 routes. You must specify the `vni` option and a VXLAN network identifier (VNI) with this statement.

We support this statement in Layer 3 virtual routing and forwarding (VRF) instances on EVPN-VXLAN devices with EVPN instances defined with MAC-VRF routing instances (`instance-type mac-vrf`).

By default, devices use asymmetric Type 2 routing in the routing instance.

Default

Not enabled.

Options

vni (Mandatory) VNI used for EVPN Type 2 symmetric routing IP prefixes.

- **Range:** 1 through 16777214

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.3R1-S1 and 21.4R1 on QFX5210 switches only.

RELATED DOCUMENTATION

Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics

is-ma

IN THIS SECTION

- Syntax | 9396
- Hierarchy Level | 9396
- Description | 9396
- Default | 9396
- Options | 9397
- Required Privilege Level | 9397
- Release Information | 9397

Syntax

```
(is-ma | no-is-ma);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
[edit dynamic-profiles name interfaces name              optics-options],
[edit dynamic-profiles name logical-systems name interfaces name              optics-options],
[edit dynamic-profiles name interfaces name              optics-options],
[edit interfaces name optics-options]
```

Description

Specify whether masked alarms are enabled or disabled.

Default

If you omit the is-ma statement, masked alarms are disabled.

Options

is-ma	Enable masked alarms.
no-is-ma	Do not enable masked alarms.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

isdn-options

IN THIS SECTION

- Syntax | 9398
- Hierarchy Level | 9398
- Description | 9398
- Required Privilege Level | 9398
- Release Information | 9398

Syntax

```
isdn-options {
    bchannel-allocation (ascending | descending);
    calling-number number;
    incoming-called-number number <reject>;
    spid1 spid-string;
    spid2 spid-string;
    static-tei-val value;
    switch-type (att5e | etsi | ni1 | ntdms100 | ntt);
    t310 seconds;
    tei-option (first-call | power-up);
}
```

Hierarchy Level

```
[edit interfaces br-pim/0/port],
[edit interfaces ct1-pim/0/port],
[edit interfaces ce1-pim/0/port]
```

Description

For J Series Services Routers only. Specify the ISDN options for configuring ISDN interfaces for group and user sessions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`bchannel-allocation` option added in Junos OS Release 8.3.

RELATED DOCUMENTATION

- Configuring ISDN Physical Interface Properties
- Allocating B-Channels for Dialout

isid

IN THIS SECTION

- Syntax | 9399
- Hierarchy Level | 9399
- Description | 9399
- Options | 9400
- Required Privilege Level | 9400
- Release Information | 9400

Syntax

```
isid isid-number vlan-id-list vlan-ids {
    source-bmac mac-address length
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name service-groups service-group-name]
```

Description

For IEEE 802.1ah provider backbone bridge (PBB) configurations, configure the service identifier (I-SID) for the customer routing instance (I-component) service group.

Options

isid—Service identifier. Enter an I-SID in the range from **256** through **16777214**.

vlan-id-list [*vlan-ids* **]**—List of service VLANs (S-VLANs).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.0.

isid-list

IN THIS SECTION

- Syntax | **9400**
- Hierarchy Level | **9401**
- Description | **9401**
- Options | **9401**
- Required Privilege Level | **9401**
- Release Information | **9401**

Syntax

```
isid-list  all-service-groups;
```

Hierarchy Level

```
[edit interfaces pseudo-logical-interface-name unit logical-unit-number family bridge]
```

Description

For IEEE 802.1ah provider backbone bridge (PBB) configurations, map all service identifiers (I-SIDs) specified for the service groups.

Options

all-service-groups—Map all service identifiers (I-SIDs) for the specified service groups.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.0.

isis

IN THIS SECTION

- Syntax | 9402
- Hierarchy Level | 9402
- Description | 9402
- Default | 9402
- Required Privilege Level | 9402
- Release Information | 9402

Syntax

```
isis { ... }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable IS-IS routing on the routing device or for a routing instance.

The `isis` statement is the one statement you must include in the configuration to run IS-IS on the routing device or in a routing instance.

Default

IS-IS is disabled on the routing device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Example: Configuring IS-IS*
>
> *Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding*

isis-instance

IN THIS SECTION

- Syntax | 9403
- Hierarchy Level | 9403
- Description | 9403
- Default | 9404
- Required Privilege Level | 9404
- Release Information | 9404

Syntax

```
isis-instance igpinstance-
name{ ... }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols]
```

Description

Enable multiple IGP instances of IS-IS routing on the routing device.

The `isis-instance` statement is the one statement you must include in the configuration to run multiple IS-IS on the routing device.

Enabling IS-IS on an interface (by including the `interface` statement at the `[edit protocols isis-instance igpinstance-name]` hierarchy level), disabling it (by including the `disable` statement), and not actually having IS-IS run on an interface (by including the `passive` statement) are mutually exclusive states.

All the configuration options available under the `[edit protocols isis-instance igpinstance-name]` hierarchy level are similar to the options available under the `[edit protocols isis]` hierarchy level.

> **NOTE**: Junos OS does not support configuring the same logical interface in multiple IGP instances of IS-IS.
>
> We do not support graceful restart for an IS-IS instance. Graceful restart is enabled by default. Therefore, you must explicitly disable graceful restart.

Default

IS-IS is disabled on the routing device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

How to Configure Multiple Independent IGP Instances of IS-IS

iso-vpn

IN THIS SECTION

- Syntax | 9405
- Hierarchy Level | 9405
- Description | 9405
- Default | 9405
- Required Privilege Level | 9405
- Release Information | 9405

Syntax

```
iso-vpn {
    unicast {
        prefix-limit number;
        rib-group group-name;
    }
}
```

Hierarchy Level

```
[edit protocols bgp family],
[edit protocols bgp group group-name family],
[edit protocols bgp group group-name neighbor addressfamily],
[edit routing-instances routing-instance-name protocols bgp family],
[edit routing-instances routing-instance-name protocols bgp group group-name family],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family]
```

Description

Enable BGP to carry ISO VPN NLRI messages between PE routes connecting a VPN.

The remaining statements are explained separately in this chapter.

Default

Disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring BGP for CLNS VPNs

Understanding BGP for CLNS VPNs

isolated

IN THIS SECTION

- Syntax | 9406
- Hierarchy Level | 9406
- Description | 9406
- Required Privilege Level | 9406
- Release Information | 9407

Syntax

```
isolated;
```

Hierarchy Level

```
[edit vlans vlan-name interfaces interface-name]
```

Description

Configure an access or trunk port to be isolated. You configure a trunk port to be isolated so that it can be a secondary VLAN trunk port–that is, it can carry secondary VLAN traffic.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Creating a Private VLAN on a Single QFX Switch without ELS Support

Creating a Private VLAN Spanning Multiple QFX Series Switches without ELS Support

Example: Configuring PVLANs with Secondary VLAN Trunk Ports and Promiscuous Access Ports on a QFX Series Switch

isolated-vlan

IN THIS SECTION

- Syntax | 9407
- Hierarchy Level | 9407
- Description | 9408
- Required Privilege Level | 9408
- Release Information | 9408

Syntax

```
isolated-vlan vlan-name isolated-vlan-name vlan-id isolated-vlan-id;
```

Hierarchy Level

```
[edit vlans primary-vlan-name vlan-id primary-vlan-vlan-id]
```

Description

Configure the specified isolated VLAN to be a secondary VLAN of the specified primary VLAN. An isolated VLAN receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN.

> **NOTE**: Before you specify this configuration statement, you must have already configured an isolated VLAN and assigned a VLAN ID to it. See "private-vlan" on page 13220.

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Creating a Private VLAN on a Single Switch with ELS Support (CLI Procedure)

Creating a Private VLAN Spanning Multiple EX Series Switches with ELS Support (CLI Procedure)

isolated-vlan (MX Series)

IN THIS SECTION

- Syntax | 9409
- Hierarchy Level | 9409
- Description | 9409
- Options | 9409
- Required Privilege Level | 9409

Release Information | 9410

Syntax

```
isolated-vlan vlan-id;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name,
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]
```

Description

Configure the specified isolated VLAN to be a secondary VLAN of the specified primary VLAN. An isolated VLAN receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN.

NOTE: When you specify this configuration statement, the VLAN ID of a logical interface that you associate with a bridge domain that matches with the VLAN ID that you specify using the `isolated-vlan` state is treated as an isolated port.

Options

vlan-id Individual VLAN IDs separated by a space.

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

isolation-id

IN THIS SECTION

- Syntax | 9410
- Hierarchy Level | 9410
- Description | 9410
- Options | 9410
- Required Privilege Level | 9411
- Release Information | 9411

Syntax

```
isolation-id number;
```

Hierarchy Level

```
[edit vlans  vlan-name vlan-id number]
```

Description

Configure an inter-switch isolated VLAN within a private VLAN (PVLAN) that spans multiple switches.

Options

number—VLAN tag identifier.

- **Range:** 0 through 4093

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)

isolation-vlan-id

IN THIS SECTION

- Syntax | 9411
- Hierarchy Level | 9411
- Description | 9412
- Options | 9412
- Required Privilege Level | 9412
- Release Information | 9412

Syntax

```
isolation-vlan-id number;
```

Hierarchy Level

```
[edit vlans vlan-name pvlan]
```

Description

Configure an interswitch isolated VLAN within a private VLAN that spans multiple switches.

Options

number—VLAN tag identifier.

- **Range:** 0 through 4093

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Creating a Private VLAN on a Single QFX Switch without ELS Support

Creating a Private VLAN Spanning Multiple QFX Series Switches without ELS Support

iteration-count

IN THIS SECTION

- Syntax | 9413
- Hierarchy Level | 9413
- Description | 9413
- Options | 9413
- Required Privilege Level | 9413
- Release Information | 9413

Syntax

```
iteration-count count-value;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-
association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile profile-name]
```

Description

Configure the number of iterations for which the connection partakes in the iterator for acquiring SLA measurements.

Options

count-value—Number of iterations for which the connection should partake in the iterator for acquiring SLA measurements.

- **Range:** 1 through 65,535
- **Default:** 0 (or infinite iterations)

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

sla-iterator-profile | **28012**

Configuring a Remote MEP with an Iterator Profile

iteration-period

IN THIS SECTION

- Syntax | 9414
- Hierarchy Level | 9414
- Description | 9414
- Options | 9414
- Required Privilege Level | 9415
- Release Information | 9415

Syntax

```
iteration-period iteration-period-value;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
profile-name]
```

Description

Configure the iteration period, which is the maximum number of cycles per iteration (that is, the number of connections registered to an iterator cannot exceed this value).

Options

iteration-period-value—Maximum number of cycles per iteration.

- **Range:** 1 through 2000
- **Default:** 2000

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring an Iterator Profile

Configuring an Iterator Profile on a Switch (CLI Procedure)

ivacl-firewall-no-portrange-profile

IN THIS SECTION

- Syntax | 9415
- Hierarchy Level | 9415
- Description | 9416
- Required Privilege Level | 9416
- Release Information | 9416

Syntax

```
ivacl-firewall-no-portrange-profile
```

Hierarchy Level

```
[edit chassis]
```

Description

Optimize and merge IVACL VxLAN and IVACL groups.

> **NOTE**: When this configuration statement is enabled, `source-port-range-optimize` and `destination-port-range-optimize` should not be used for IVACL filters.

Required Privilege Level

interface-control

Release Information

Statement introduced in Junos OS Release 22.2R1

RELATED DOCUMENTATION

show pfe filter hw summary | 22983

ivlan-cfi (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 9416
- Hierarchy Level | 9417
- Description | 9417
- Required Privilege Level | 9417
- Release Information | 9417

Syntax

```
ivlan-cfi;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

CFI bit to be used in the inner VLAN header of the frames generated.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

ivlan-id (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | **9418**
- Hierarchy Level | **9418**
- Description | **9418**
- Options | **9418**
- Required Privilege Level | **9418**
- Release Information | **9418**

Syntax

```
ivlan-id number;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Configure the inner VLAN ID for the test frames. This parameter is applicable for dual-tagged packets.

Options

number VLAN ID number.

- **Range:** 0 through 4094

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

ivlan-priority (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 9419
- Hierarchy Level | 9419
- Description | 9419
- Options | 9419
- Required Privilege Level | 9419
- Release Information | 9420

Syntax

```
ivlan-priority value;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Description

Configure the priority value for the IEEE 802.1p bit in the inner VLAN tag. This parameter is valid only for a bridge family for an Ethernet LAN (ELAN) or an Ethernet Line (E-LINE) service.

Options

value IEEE 802.1p priority value in the inner VLAN tag

- **Range:** 0 through 7

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

iwf-params (tdm-options)

IN THIS SECTION

- Syntax | **9420**
- Hierarchy Level | **9421**
- Description | **9421**
- Default | **9421**
- Options | **9421**
- Required Privilege Level | **9421**
- Release Information | **9421**

Syntax

```
iwf-params {
    decap-ecid value;
    encap-ecid value;
}
```

Hierarchy Level

```
[edit interfaces interface-name tdm-options]
```

Description

Configure TDM interworking functionality (IWF) parameters. You can configure these parameters to support multiple streams.

Default

If you do not specify an emulation circuit id, the default emulation circuit id for decapsulation and encapsulation is 0. You cannot configure multiple streams if the emulation circuit id is 0.

Options

decap-ecid *value* Emulation circuit id for de-encapsulation. Possible values: 0 through 1048575.

encap-ecid *value* Emulation circuit id for encapsulation. Possible values: 0 through 1048575.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4.

RELATED DOCUMENTATION

show interfaces smart-sfp-statistics | 21372

show interfaces smart-sfp-defects | 21359

J

IN THIS SECTION

- jdaf | 9422
- jeb | 9423
- jflow-log (Interfaces) | 9426
- jflow-log (Services) | 9427
- jnu-management | 9429
- join-load-balance | 9431
- join-prune-timeout | 9433
- join-timer (MVRP) | 9434
- json | 9436
- jsrc (Access Profile) | 9438
- jsrc (JSRC) | 9439
- jsrc-partition | 9441
- juniper-enhanced | 9442
- juniper-express-engine | 9444
- juniper-local | 9447

jdaf

IN THIS SECTION

- Syntax | 9423
- Hierarchy Level | 9423
- Description | 9423
- Required Privilege Level | 9423
- Release Information | 9423

Syntax

```
jdaf {
    routing-instances [ routing-instance-names ];
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure Juniper distributed application framework (JDAF).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

jeb

IN THIS SECTION

- Syntax | 9424
- Hierarchy Level | 9424
- Description | 9424
- Options | 9424

Required Privilege Level | 9425
Release Information | 9425

Syntax

```
jeb {
    max-seed-size max-seed-size;
    port port;
    rbg (default-rng | hmac-drbg);
    tls {
        cert-bundle cert-bundle;
        certificate certificate;
        key key;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure a Juniper Entropy Beacon (JEB) server. JEB can be used to feed high quality entropy over the network from a SRX345 Firewall to entropy-starved clients.

Options

max-seed-size Maximum allowed size in bytes for a requested entropy seed

- **Default:** 4096
- **Range:** 1 through 65536

port Port to use for JEB service

- **Default:** 57005
- **Range:** 1025 through 65535

rbg Type of random bit generator (RBG) to use for generating entropy seeds

- Values:
  - default-rng—Default cryptographically secure pseudorandom number generator (CSPRNG)
  - hmac-drbg—Deterministic RBG detailed in NIST SP 800-90A

> **NOTE**: If you want to use `hmac-drbg`, it must be configured under the `[edit system rng]` hierarchy before being configured for the JEB server.

TLS Configure TLS attributes for JEB services

cert-bundle Path to the certificate bundle that is used by the server to authenticate the client

certificate Path to the server certificate

key Path to the server key

Required Privilege Level

admin

Release Information

Statement introduced in Junos OS Release 19.1R1.

jflow-log (Interfaces)

IN THIS SECTION

- Syntax | 9426
- Hierarchy Level | 9426
- Description | 9426
- Required Privilege Level | 9426
- Release Information | 9427

Syntax

```
jflow-log {
    message-rate-limit messages-per-second;
}
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Description

Configure generation of log messages or template records in flow monitoring format for NAT error events. These records for NAT error events are generated when addresses for allocation from the NAT pool are not available, when ports for allocation to a subscriber are not available, or when the allocated quota is exceeded for NAT events (more than the configured number of ports is requested).

The remaining statement is described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

jflow-log (Services)

IN THIS SECTION

- Syntax | **9427**
- Hierarchy Level | **9428**
- Description | **9428**
- Required Privilege Level | **9428**
- Release Information | **9428**

Syntax

```
jflow-log {
    collector collector-name {
        source-ip address;
        destination-address address;
        destination-port port-number;
    }
    collector-group collector-group-name {
        [collector-name1 collector-name2];
    }
```

```
    template-profile template-profile-name {
        collector collector-name ;
        collector-group collector-group-name ;
        template-type nat;
        version (ipfix | v9);
        refresh-rate packets packets seconds seconds;
    message-rate-limit messages-per-second
```

Hierarchy Level

```
[edit services]
```

Description

Enable the mechanism to record logging messages in flow monitoring format for NAT events. For this transmission of flow monitoring logs to work properly, the services PIC interface must have an IP address and appropriate logging options configured.

You can configure MX Series routers with MS-MPCs and MS-MICs to log network address translation (NAT) events using the Junos Traffic Vision (previously known as Jflow) version 9 or IPFIX (version 10) template format. This method of generating flow monitoring records for NAT events, such as NAT44 and NAT64 session creation and deletion, and NAT44 and NAT64 binding information base events, enables cohesive and streamlined analysis of NAT traffic and troubleshooting of NAT-related problems.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R2.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

jnu-management

IN THIS SECTION

- Syntax | 9429
- Hierarchy Level | 9430
- Description | 9430
- Options | 9430
- Required Privilege Level | 9431
- Release Information | 9431

Syntax

```
jnu-management {
    authentication authentication;
    controller controller;
    device-group device-group;
    device-group-map name {
        include-groups [ include-groups ... ];
    }
    mode(bng-control-plane | feature-rich);
    other-controller;
    satellite name {
        device-group device-group;
        model model;
        version version;
    }
    satellite-name satellite-name;
```

```
    user user;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

JNU Management configuration.

Options

authentication Authentication string

Option not applicable for Connected Security Distributed Services (CSDS).

controller Controller Name

device-group Satellite device group

Option not applicable for CSDS.

device-group-map Common groups mapping to device groups

Option not applicable for CSDS.

mode JNU Mode

- bng-control-plane—BNG control-plane mode
- feature-rich—Feature Rich Mode

CSDS topology supports only feature-rich mode.

other-controller Name of the dual controller in Junos Node Unifier (JNU) topology for CSDS.

satellite-name Satellite name

user Satellite User Name

CSDS topology supports only jnuadmin user.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

[none specified]

Release Information

Support for `other-controller` option is added for CSDS in Junos OS Release 24.4R1.

Statement introduced in Junos OS Release 21.3R1.

join-load-balance

IN THIS SECTION

- Syntax | 9431
- Hierarchy Level | 9432
- Description | 9432
- Options | 9432
- Required Privilege Level | 9432
- Release Information | 9432

Syntax

```
join-load-balance {
    automatic;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Enable load balancing of PIM join messages across interfaces and routing devices.

Options

automatic Enables automatic load balancing of PIM join messages. When a new interface or neighbor is introduced into the network, ECMP joins are redistributed with minimal disruption to traffic.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring PIM Make-Before-Break Join Load Balancing

Configuring PIM Join Load Balancing

clear pim join-distribution | **3809**

join-prune-timeout

IN THIS SECTION

- Syntax | 9433
- Hierarchy Level | 9433
- Description | 9433
- Options | 9433
- Required Privilege Level | 9434
- Release Information | 9434

Syntax

```
join-prune-timeout seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Configure the timeout for the join state. If the periodic join refresh message is not received before the timeout expires, the join state is removed.

Options

seconds—Number of seconds to wait for the periodic join message to arrive.

- **Range:** 210 through 240 seconds
- **Default:** 210 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Modifying the Join State Timeout

join-timer (MVRP)

IN THIS SECTION

- Syntax | 9434
- Hierarchy Level | 9435
- Description | 9435
- Default | 9435
- Options | 9435
- Required Privilege Level | 9435
- Release Information | 9436

Syntax

```
join-timer milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mvrp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvrp
interface (all | interface-name)] (for virtual switch instance type),
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
mvrp] (for virtual switch instance type),
[edit logical-systems logical-system-name protocols mvrp interface (all | interface-name)]
```

```
[edit protocols mvrp interface (all | interface-name)]
```

```
[edit routing-instances routing-instance-name protocols mvrp] (for virtual switch instance type),
[edit routing-instances routing-instance-name protocols mvrp interface (all | interface-name)]
(for virtual switch instance type)
```

Description

Configure the maximum number of milliseconds interfaces must wait before sending Multiple VLAN Registration Protocol (MVRP) protocol data units (PDUs).

Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Default

200 milliseconds

Options

milliseconds—Interval that the interface must wait before sending MVRP PDUs (range from 100 milliseconds through 500 milliseconds). Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- leave-timer (MVRP) | 9705
- leaveall-timer (MVRP) | 9707
- point-to-point (MVRP) | 12620
- registration | 14268
- *Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers*
- *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches*
- *Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration*
- *Configuring Multiple VLAN Registration Protocol (MVRP) on Switches*
- *Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration*
- *Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration*
- *Verifying That MVRP Is Working Correctly*

json

IN THIS SECTION

- Syntax | 9437
- Hierarchy Level | 9437
- Description | 9437
- Options | 9437
- Required Privilege Level | 9437
- Release Information | 9437

Syntax

```
json {
    reorder-list-keys;
}
```

Hierarchy Level

```
[edit system configuration input format]
```

Description

Configure parameters for Junos configuration data in JavaScript Object Notation (JSON) format.

Options

reorder-list-keys When you load JSON configuration data on the device, reorder list entries so that list keys precede other siblings and appear in the order required by the schema. Configuring this option enables you to load JSON configuration data that contains unordered list entries, that is, entries where the list key is not necessarily the first element.

> **NOTE**: When you configure the `reorder-list-keys` statement, the load operation can take significantly longer to parse the configuration, depending on the size of the configuration and the number of lists. For large configurations or configurations with many lists, we recommend that you instead use the `request system convert-json-configuration` command to convert the configuration before loading it on the device.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1R1 and Junos OS Evolved Release 22.1R1.

RELATED DOCUMENTATION

request system convert-json-configuration | 15060

Load JSON Configuration Data With Unordered List Entries

jsrc (Access Profile)

IN THIS SECTION

- Syntax | 9438
- Hierarchy Level | 9438
- Description | 9438
- Required Privilege Level | 9439
- Release Information | 9439

Syntax

```
jsrc {
    attributes {
        exclude {
            user-name [ authorization-request | provisioning-request ];
            }
        }
    }
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Specify JSRC settings in an access profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

- *JSRC Configuration Overview*
- *Authorizing Subscribers with JSRC*
- *Understanding JSRC-SAE Interactions*
- *Excluding AVPs from Diameter Messages for JSRC*

jsrc (JSRC)

IN THIS SECTION

- Syntax | 9440
- Hierarchy Level | 9440
- Description | 9440
- Required Privilege Level | 9440
- Release Information | 9440

Syntax

```
jsrc {
    dualstack-support;
    partition partition-name  {
        diameter-instance instance-name;
        destination-host hostname;
        destination-realm realm-name;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure JSRC to interact with an SAE in an SRC environment to authorize and provision subscribers.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

JSRC Configuration Overview

jsrc-partition

IN THIS SECTION

- Syntax | 9441
- Hierarchy Level | 9441
- Description | 9441
- Options | 9441
- Required Privilege Level | 9441
- Release Information | 9442

Syntax

```
jsrc-partition partition-name;
```

Hierarchy Level

```
[edit]
```

Description

Specify the JSRC partition to use.

Options

partition-name—Name of the JSRC partition that you want JSRC to use. The name is defined with the `partition` statement at the `[edit jsrc]` hierarchy level.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

JSRC Configuration Overview

Configuring the JSRC Partition

juniper-enhanced

IN THIS SECTION

- Syntax | 9442
- Hierarchy Level | 9443
- Description | 9443
- Options | 9443
- Required Privilege Level | 9443
- Release Information | 9444

Syntax

```
juniper-enhanced {
    cache {
        size value;
        timeout value;
    }
    profile profile-name {
        category customurl-list name {
            action  (block | log-and-permit | permit | quarantine);
        }
        custom-block-message value;
        custom-quarantine-message value;
        default  (block | log-and-permit | permit | quarantine);
```

```
        fallback-settings {
            default  (block | log-and-permit);
            server-connectivity   (block | log-and-permit);
            timeout    (block | log-and-permit);
            too-many-requests  (block | log-and-permit);
        }
        no-safe-search;
        site-reputation-action {
            fairly-safe   (block | log-and-permit | permit | quarantine);
            harmful   (block | log-and-permit | permit | quarantine);
            moderately-safe  (block | log-and-permit | permit | quarantine);
            suspicious  (block | log-and-permit | permit | quarantine);
            very-safe   (block | log-and-permit | permit | quarantine);
            }
        timeout value;
    }
    server {
        host host-name;
        port number;
        proxy-profile proxy profile name;
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering]
```

Description

Configure the Content Security Enhanced Web Filtering feature.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

The `proxy-profile` option is introduced under the `security utm default-configuration web-filtering juniper-enhanced server` hierarchy level in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Web Filtering Overview

juniper-express-engine

IN THIS SECTION

- Syntax | 9444
- Hierarchy Level | 9446
- Description | 9446
- Options | 9446
- Required Privilege Level | 9446
- Release Information | 9446

Syntax

```
juniper-express-engine {
    pattern-update {
        email-notify {
            admin-email email-address;
            custom-message message;
            custom-message-subject message-subject;
```

```
        }
        interval value;
        no-autoupdate;
        proxy {
            password password-string;
            port port-number;
            server address-or-url;
            username name;
        }
        url url;
    }
    profile profile-name {
         fallback-options {
            content-size (block | log-and-permit);
            default (block | log-and-permit);
            engine-not-ready (block | log-and-permit);
            out-of-resources (block | (log-and-permit);
            timeout (block | log-and-permit);
            too-many-requests (block | log-and-permit);
        }
        notification-options {
            fallback-block {
                administrator-email email-address;
                allow-email;
                custom-message message;
                custom-message-subject message-subject;
                display-host;
                (notify-mail-sender | no-notify-mail-sender);
                 type (message | protocol-only);
            }
            fallback-non-block {
                custom-message message;
                custom-message-subject message-subject;
                (notify-mail-recipient | no-notify-mail-recipient);
            }
            virus-detection {
                custom-message message;
                custom-message-subject message-subject;
                (notify-mail-sender | no-notify-mail-sender);
                 type (message | protocol-only);
            }
        }
        scan-options {
```

```
            content-size-limit value;
            (intelligent-prescreening | no-intelligent-prescreening);
            timeout value;
        }
        trickling {
            timeout value;
        }
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus]
```

Description

Configure the Content Security express antivirus feature.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Express Antivirus Configuration Overview

juniper-local

IN THIS SECTION

- Syntax | 9447
- Hierarchy Level | 9447
- Description | 9448
- Options | 9448
- Required Privilege Level | 9448
- Release Information | 9448

Syntax

```
juniper-local {
    profile profile-name {
        block-message (Security UTM) value;
        default (Security Web Filtering) (block | log-and-permit | permit);
        fallback-settings (Security Web Filtering) {
            default (Security Web Filtering) (block | log-and-permit);
            server-connectivity (block | log-and-permit);
            timeout (Security Web Filtering Fallback Settings) (block | log-and-permit);
            too-many-requests (Security Web Filtering Fallback Settings) (block | log-and-permit);
        }
        timeout value;
        no-safe-search;
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering]
```

Description

Configure the Content Security Web-filtering local feature.

Options

no-safe-search Disable the safe search function.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

`no-safe-search` option added in Junos OS Release 20.2R1.

K

IN THIS SECTION

- kaspersky-lab-engine | 9450
- keep | 9452
- keep-address-and-control | 9455
- keep-alive (Protocols MSDP) | 9457
- keep-existing-steering | 9459
- keep-incoming-circuit-id (DHCP Relay Agent) | 9460
- keep-incoming-interface-id (DHCP Relay Agent) | 9462
- keep-incoming-remote-id (DHCP Relay Agent) | 9463
- keep-multiplier | 9465

- keepalive | **9466**
- keepalive-interval | **9468**
- keepalive-retries | **9469**
- keepalive-time (Chassis) | **9471**
- keepalive-time (Protocols OAM) | **9472**
- keepalive-timeout | **9474**
- keepalives | **9475**
- keepalives (Dynamic Profiles) | **9478**
- kernel-replication | **9480**
- key | **9481**
- key (gre) | **9482**
- key (interfaces) | **9484**
- key (Authentication Keychain) | **9485**
- key-chain (Authentication Keychain) | **9488**
- key-chain (Protocols PIM) | **9490**
- key-exchange | **9491**
- key (Junos FIPS) | **9494**
- key-length | **9495**
- key (MACsec) | **9497**
- key (MACsec for MX Series) | **9498**
- key-pair | **9500**
- key-server-priority (MACsec) | **9502**
- key-server-priority (MACsec for MX Series) | **9503**
- krt-nexthop-ack | **9505**
- krt-nexthop-ack-timeout | **9506**

kaspersky-lab-engine

IN THIS SECTION

- Syntax | 9450
- Hierarchy Level | 9452
- Description | 9452
- Options | 9452
- Required Privilege Level | 9452
- Release Information | 9452

Syntax

```
kaspersky-lab-engine {
    pattern-update {
        email-notify {
            admin-email email-address;
            custom-message message;
            custom-message-subject message-subject;
        }
        interval value;
        no-autoupdate;
        proxy {
            password password-string;
            port port-number;
            server address-or-url;
            username name;
        }
        url url;
    }
    profile profile-name {
         fallback-options {
            content-size (block | log-and-permit);
            corrupt-file (block | log-and-permit);
            decompress-layer (block | log-and-permit);
            default (block | log-and-permit);
            engine-not-ready (block | log-and-permit);
```

```
            out-of-resources (block | (log-and-permit);
            password-file (block | (log-and-permit);
            timeout (block | log-and-permit);
            too-many-requests (block | log-and-permit);
        }
        notification-options {
            fallback-block {
                administrator-email email-address;
                allow-email;
                custom-message message;
                custom-message-subject message-subject;
                display-host;
                (notify-mail-sender | no-notify-mail-sender);
                 type (message | protocol-only);
            }
            fallback-non-block {
                custom-message message;
                custom-message-subject message-subject;
                (notify-mail-recipient | no-notify-mail-recipient);
            }
            virus-detection {
                custom-message message;
                custom-message-subject message-subject;
                (notify-mail-sender | no-notify-mail-sender);
                 type (message | protocol-only);
            }
        }
        scan-options {
            content-size-limit value;
            decompress-layer-limit value;
            (intelligent-prescreening | no-intelligent-prescreening);
            scan-extension filename;
            scan-mode  (all | by-extension);
            timeout value;
        }
        trickling {
            timeout value;
        }
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus]
```

Description

Configure the Content Security full file-based antivirus feature.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Kaspersky Antivirus feature is not supported from Junos OS Release 15.1x49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

keep

IN THIS SECTION

- Syntax | 9453
- Hierarchy Level | 9453
- Description | 9453
- Default | 9454
- Options | 9455

Required Privilege Level | 9455
Release Information | 9455

Syntax

```
keep (all | none);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],

[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Control whether or not Junos OS keeps in memory and hides certain routes.

If the `keep none` statement is used, Junos OS does not retain in memory and hide routes that are rejected because of a BGP import policy. Nor does BGP keep in memory and hide routes that are declared unfeasible due to BGP sanity checks. The `keep none` statement causes Junos OS to discard from memory the routes that are rejected due to BGP-specific logic or BGP evaluation. When a route is rejected because of some non-BGP-specific reason, the `keep none` statement has no effect on this route. This rejected route is retained in memory and hidden even though `keep none` is configured. An example of this type of hidden route is a route for which the protocol nexthop is unresolved.

The routing table can retain the route information learned from BGP in one of the following ways:

- Default (omit the `keep` statement)—Keep all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.
- `keep all`—Keep all route information that was learned from BGP.
- `keep none`—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking, such as AS path or next hop. When you configure `keep none` for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.

In an AS path healing situation, routes with looped paths theoretically could become usable during a soft reconfiguration when the AS path loop limit is changed. However, there is a significant memory usage difference between the default and `keep all`.

Consider the following scenarios:

- A peer readvertises routes back to the peer from which it learned them.

  This can happen in the following cases:

  - Another vendor's routing device advertises the routes back to the sending peer.
  - The Junos OS peer’s default behavior of not readvertising routes back to the sending peer is overridden by configuring `advertise-peer-as`.
- A provider edge (PE) routing device discards any VPN route that does not have any of the expected route targets.

When `keep all` is configured, the behavior of discarding routes received in the above scenarios is overridden.

> **CAUTION**: If you add or remove `keep all` or `keep none` and the peer does not support session restart, the associated BGP sessions are restarted (flapped). To determine if a peer supports refresh, check for `Peer supports Refresh capability` in the output of the `show bgp neighbor` command.

Default

By default, BGP retains incoming rejected routes in memory and hides them. If you do not include the `keep` statement, most routes are retained in the routing table. BGP keeps all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.

Options

`all`—Retain all routes.

`none`—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking. When `keep none` is configured for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routing Policies to Control BGP Route Advertisements

out-delay | **11961**

Interprovider VPN Example—MP-EBGP Between ISP Peer Routers

Example: Configuring a Routing Policy for Conditional Advertisement Enabling Conditional Installation of Prefixes in a Routing Table

keep-address-and-control

IN THIS SECTION

- Syntax | **9456**
- Hierarchy Level | **9456**
- Description | **9456**
- Default | **9456**
- Required Privilege Level | **9456**

Release Information | 9456

Syntax

```
keep-address-and-control;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family ccc],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family ccc]
```

Description

For interfaces with encapsulation type PPP CCC, do not remove the address and control bytes before encapsulating the packet into a tunnel.

Default

If you do not include this statement, address and control bytes are removed before encapsulating the packet into a tunnel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Disabling the Removal of Address and Control Bytes

keep-alive (Protocols MSDP)

IN THIS SECTION

- Syntax | 9457
- Hierarchy Level | 9457
- Description | 9458
- Default | 9458
- Options | 9458
- Required Privilege Level | 9458
- Release Information | 9458

Syntax

```
keep-alive seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name routing-instances instance-name protocols msdp],
[edit logical-systems logical-system-name routing-instances instance-name protocols msdp group
group-name peer address],
[edit logical-systems logical-system-name routing-instances instance-name protocols msdp peer
address],
[edit protocols msdp],
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit routing-instances instance-name protocols msdp],
[edit routing-instances instance-name protocols msdp group group-name peer address]
[edit routing-instances instance-name protocols msdp peer address]
```

Description

Specify the keepalive interval to use when maintaining a connection with the MSDP peer. If a keepalive message is not received for the hold-time period, the MSDP peer connection is terminated. According to the RFC 3618, *Multicast Source Discovery Protocol (MSDP)*, the recommended value for the keepalive timer is 60 seconds.

The hold-time period must be longer than the keepalive interval.

You might want to change the keepalive interval and hold-time period for consistency in a multi-vendor environment.

Default

In Junos OS, the default hold-time period is 75 seconds, and the default keepalive interval is 60 seconds.

Options

seconds—Keepalive interval.

- **Range:** 10 through 60 seconds
- **Default:** 60 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Examples: Configuring MSDP

hold-time (Protocols MSDP) | **8112**

sa-hold-time (Protocols MSDP) | **16262**

keep-existing-steering

IN THIS SECTION

- Syntax | 9459
- Hierarchy Level | 9459
- Description | 9459
- Required Privilege Level | 9459
- Release Information | 9460

Syntax

```
keep-existing-steering;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name steering],
[edit services pcef pcc-action-profiles profile-name]
```

Description

Specify that the PCC action profile steering attributes that a PCC rule applies at the start of a data flow will continue to be applied to that data flow when the PCC rule match conditions are modified, deleted, or added to.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

- *Understanding Application-Aware Policy Control for Subscriber Management*
- *Configuring Policy and Charging Control Action Profiles for Subscriber Management*
- Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment
- Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

keep-incoming-circuit-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 9460
- Hierarchy Level | 9461
- Description | 9461
- Required Privilege Level | 9461
- Release Information | 9461

Syntax

```
keep-incoming-circuit-id ;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit logical-systems logical-system-name forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82 circuit-id]
```

Description

Specify that the `jdhcpd` process keeps the incoming circuit ID and prepends the ID with the locally generated ID (in the format, `generated-id + incoming-id`) before sending the leasequery packet to the DHCP server.

This configuration is required for leasequery and bulk leasequery operations when subscriber authentication is based on the circuit ID, and enables leasequery and bulk leasequery to restore the agent circuit identifier/agent remote identifier (ACI/ARI) pair and to use the circuit ID to authenticate subscribers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information

keep-incoming-interface-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 9462
- Hierarchy Level | 9462
- Description | 9462
- Required Privilege Level | 9463
- Release Information | 9463

Syntax

```
keep-incoming-interface-id ;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 relay-agent-
interface-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 relay-agent-interface-id],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 relay-agent-
interface-id],
```

Description

Specify that the `jdhcpd` process keeps the incoming interface ID and prepends the ID with the locally generated ID (in the format, `generated-id + incoming-id`) before sending the leasequery packet to the DHCPv6 server.

This configuration is required for leasequery and bulk leasequery operations when subscriber authentication is based on the interface ID, and enables leasequery and bulk leasequery to restore the agent circuit identifier/agent remote identifier (ACI/ARI) pair and to use the interface ID to authenticate subscribers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets

keep-incoming-remote-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 9463
- Hierarchy Level | 9463
- Description | 9464
- Required Privilege Level | 9464
- Release Information | 9464

Syntax

```
keep-incoming-remote-id ;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-remote-id],
[edit forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 relay-agent-remote-
```

```
id],
[edit logical-systems logical-system-name forwarding-options dhcp-relay relay-option-82
XXXcircuit-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 relay-agent-remote-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay relay-option-82 circuit-id],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 relay-agent-
remote-id],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82
circuit-id],
```

Description

Specify that the `jdhcpd` process keeps the incoming remote ID and prepends the ID with the locally generated ID (in the format, `generated-id + incoming-id`) before sending the leasequery packet to the DHCPv6 server.

This configuration is required for leasequery and bulk leasequery operations when subscriber authentication is based on the remote ID, and enables leasequery and bulk leasequery to restore the agent circuit identifier/agent remote identifier (ACI/ARI) pair and to use the remote ID to authenticate subscribers.

Use the statement at the `[edit ... dhcpv6]` hierarchy level to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Using DHCP Relay Agent Option 82 Information*
>
> *Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets*

keep-multiplier

IN THIS SECTION

- Syntax | 9465
- Hierarchy Level | 9465
- Description | 9465
- Options | 9466
- Required Privilege Level | 9466
- Release Information | 9466

Syntax

```
keep-multiplier number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Specify a value used by RSVP to calculate timer values for network outages, and declare that a reservation or a neighbor is down. It indicates the number of messages that can be lost before a particular state is declared stale and must be deleted. The keep multiplier directly affects the lifetime of an RSVP state.

Starting in Junos OS Release 16.1, for MX Series routers and PTX Series routers, the default RSVP message refresh time, to which this multiplier is applied, has increased from 30 seconds to 20 minutes. The higher message refresh time provides support for RSVP Refresh Reduction Extensions, and improved scaling for MPLS traffic-engineered LSPs, as defined in RFC 2961. The changes are backward compatible so if any nodes in the two-hop neighborhood do not support the higher refresh time, the updated node will automatically fall back to the previous default refresh time to prevent error or tear down messages.

Options

number—Multiplier value.

- **Range:** 1 through 255
- **Default:** 3

> **NOTE**: For MX Series routers and PTX Series routers (running Junos OS release 16.1 or later), this multiplier is applied to a default refresh time of 20 minutes. In earlier Junos OS releases, and for other platforms, the default refresh time remains 30 seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Timers for RSVP Refresh Messages

keepalive

IN THIS SECTION

- Syntax | 9467
- Hierarchy Level | 9467
- Description | 9467
- Options | 9467
- Required Privilege Level | 9467

Release Information | 9467

Syntax

```
keepalive seconds;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp],
[edit access profile profile-name client client-name ppp]
```

Description

Configure the keepalive interval for an L2TP tunnel.

Options

seconds—Time period that must elapse before the Junos OS checks the status of the Point-to-Point Protocol (PPP) session by sending an echo request to the peer.

For L2TP on MX Series routers, the minimum recommended interval is 30 seconds. A value of 0 disables generation of keepalive messages from the LNS.

- **Range:** 0 through 32,767 seconds
- **Default:** 30 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Layer 2 Tunneling Protocol (L2TP)

L2TP LNS Inline Service Interfaces

keepalive-interval

IN THIS SECTION

- Syntax | 9468
- Hierarchy Level | 9468
- Description | 9468
- Options | 9469
- Required Privilege Level | 9469
- Release Information | 9469

Syntax

```
keepalive-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Set the keepalive interval value.

Options

seconds—Keepalive value.

- **Range:** 1 through 65,535
- **Default:** 10 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Interval for LDP Keepalive Messages

keepalive-retries

IN THIS SECTION

- Syntax | 9470
- Hierarchy Level | 9470
- Description | 9470
- Options | 9470
- Required Privilege Level | 9470
- Release Information | 9470

Syntax

```
keepalive-retries number-of-retries;
```

Hierarchy Level

```
[edit access profile profile-name client client-name ppp]
```

Description

Configure the number of retry attempts for checking the keepalive status of a Point-to-Point (PPP) protocol session. Configure this setting to reduce the detection time for PPP client session timeouts or failures if you have configured the keepalive timeout interval (using the `keepalive` statement).

Options

number-of-retries—The maximum number of retries the L2TP network server (LNS) attempts by sending LCP echo requests to the peer to check the keepalive status of the PPP session. If there is no response from the PPP client within the specified number of retries, the PPP session is considered to have timed out.

- **Range:** 3 through 32,767 times
- **Default:** 10 times

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

keepalive | 9466

keepalive-time (Chassis)

IN THIS SECTION

- Syntax | 9471
- Hierarchy Level | 9471
- Description | 9471
- Default | 9471
- Options | 9472
- Required Privilege Level | 9472
- Release Information | 9472

Syntax

```
keepalive-time seconds;
```

Hierarchy Level

```
[edit chassis redundancy]
```

Description

Configure the time period that must elapse before the backup router takes primary role when it detects loss of the keepalive signal.

Default

The `on-loss-of-keepalives` statement at the `[edit chassis redundancy failover]` hierarchy level must be included for failover to occur.

When the `on-loss-of-keepalives` statement is included and graceful Routing Engine switchover *is not* configured, failover occurs after 300 seconds (5 minutes).

When the `on-loss-of-keepalives` statement is included and graceful Routing Engine switchover *is* configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds (4 seconds on M20 routers). You cannot manually reset the keepalive time.

Options

seconds—Time before the backup router takes primary role when it detects loss of the keepalive signal. The range of values is 2 through 10,000.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

On Detection of a Loss of Keepalive Signal from the Primary Routing Engine

failover (Chassis) | **6778**

on-loss-of-keepalives | **11824**

graceful-switchover (Chassis) | **7786**

High Availability Features for EX Series Switches Overview

keepalive-time (Protocols OAM)

IN THIS SECTION

- Syntax | **9473**
- Hierarchy Level | **9473**

- Description | 9473
- Options | 9473
- Required Privilege Level | 9473
- Release Information | 9474

Syntax

```
keepalive-time seconds;
```

Hierarchy Level

```
[edit protocols oam gre-tunnel interface interface-name],
[edit protocols oam gre-tunnel interface interface-name.unit-number]
```

Description

Time difference between consecutive keepalive packets in a GRE tunnel.

> **NOTE**: Support for GRE keepalive packets on MPC line cards became available as of Junos OS Release 11.4.

Options

seconds—Keepalive time value.

- **Default:** 1 second
- **Range:** 1 through 50 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Statement introduced in Junos OS Release 19.3 for MPC10E line card.

RELATED DOCUMENTATION

Configuring Tunnel Interfaces on an MX Series Router with a 16x10GE 3D MPC

Configuring GRE Keepalive Time

hold-time (OAM) | **8102**

keepalive-timeout

IN THIS SECTION

- Syntax | **9474**
- Hierarchy Level | **9474**
- Description | **9475**
- Options | **9475**
- Required Privilege Level | **9475**
- Release Information | **9475**

Syntax

```
keepalive-timeout seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
```

```
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Set the keepalive timeout value. The keepalive timeout defines the amount of time that the neighbor LDP node waits before determining that the session has failed.

Options

seconds—Keepalive timeout value.

- **Range:** 1 through 65,535
- **Default:** 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the LDP Keepalive Timeout

keepalives

IN THIS SECTION

Syntax | 9476
Hierarchy Level | 9476

- Description | 9476
- Default | 9476
- Options | 9477
- Required Privilege Level | 9477
- Release Information | 9477

Syntax

```
keepalives <interval seconds> <down-count number> <up-count number>;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Enable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation.

For ATM2 IQ interfaces only, you can enable keepalives on a logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- `atm-ppp-llc`—PPP over AAL5 LLC encapsulation.
- `atm-ppp-vc-mux`—PPP over AAL5 multiplex encapsulation.

Default

Sending of keepalives is enabled by default. The default keepalive interval is 10 seconds for PPP, Frame Relay, or Cisco HDLC. The default down-count is 3 and the default up-count is 1 for PPP or Cisco HDLC.

Options

`down-count` *`number`*—The number of keepalive packets a destination must fail to receive before the network takes down a link.

- **Range:** 1 through 255
- **Default:** 3

`interval` *`seconds`*—The time in seconds between successive keepalive requests.

- **Range:** 1 through 32767 seconds
- **Default:** 10 seconds

`up-count` *`number`*—The number of keepalive packets a destination must receive to change a link's status from down to up.

- **Range:** 1 through 255
- **Default:** 1

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Keepalives

Configuring Frame Relay Keepalives

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

keepalives (Dynamic Profiles)

IN THIS SECTION

- Syntax | 9478
- Hierarchy Level | 9478
- Description | 9478
- Default | 9479
- Options | 9479
- Required Privilege Level | 9479
- Release Information | 9479

Syntax

```
keepalives {
    interval seconds;
    }
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit logical-unit-number ]
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit”]
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-
interface-unit”]
```

Description

Specify the keepalive interval in a PPP dynamic profile.

Starting in Junos OS Release 15.1R5, you can configure the PPP keepalive interval for subscriber services in the range 1 second through 600 seconds. Subscriber PPP keepalives are handled by the Packet Forwarding Engine. If you configure a value greater than 600 seconds, the number is accepted by the CLI, but the Packet Forwarding Engine limits the interval to 600 seconds.

In earlier Junos OS releases, the range is from 1 second through 60 seconds. The Packet Forwarding Engine limits any higher configured value to an interval of 60 seconds.

PPP keepalives for nonsubscriber services are handled by the Routing Engine with an interval range from 1 second through 32,767 seconds.

Default

Sending of keepalives is enabled by default.

Options

`interval` *`seconds`*—The time in seconds between successive keepalive requests.

- **Range:** 1 through 600 seconds for subscriber services
- **Range:** 1 through 32767 seconds for nonsubscriber services
- **Default:** 30 seconds for LNS-based PPP sessions. 10 seconds for all other PPP sessions.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles` *`profile-name`* `interfaces pp0 unit “$junos-interface-unit”]` hierarchy level introduced in Junos OS Release 10.1.

Support at the `[edit dynamic-profiles` *`profile-name`* `interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]` hierarchy level introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Dynamic Profiles Overview

Configuring Dynamic Authentication for PPP Subscribers

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

kernel-replication

IN THIS SECTION

- Syntax | 9480
- Hierarchy Level | 9480
- Description | 9480
- Options | 9480
- Required Privilege Level | 9481
- Release Information | 9481

Syntax

```
kernel-replication {
    no-multithreading;
    system-reboot recovery-failure;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure kernel replication. Use this configuration statement to debug the kernel synchronization process (ksyncd) and configure automatic recovery from ksyncd initialization errors.

Options

no-multithreading	-(Optional) Run ksyncd in single thread mode for debugging purposes.
system-reboot recovery-failure	-(Optional) Configure the backup RE to automatically reboot if a ksyncd initialization error is detected.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Understanding Graceful Routing Engine Switchover

show system switchover | 27246

key

IN THIS SECTION

- Syntax | 9481
- Hierarchy Level | 9482
- Description | 9482
- Options | 9482
- Required Privilege Level | 9482
- Release Information | 9482

Syntax

```
key number;
```

Hierarchy Level

```
[edit interfaces ftinumber unit logical-unit-number tunnelencapsulation ( gre | ipip | udp )]
```

Description

Key for configuring the tunnel. Identifies an individual traffic flow within a tunnel.

Options

number An identifier for a individual traffic flow within a GRE tunnel.

- **Range:** 1 through 65535.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Evolved Release 20.2R1.

Statement introduced for IPIP tunnels in Junos OS Evolved Release 20.4R1.

Statement introduced for UDP tunnels in Junos OS Evolved Release 21.1R1.

key (gre)

IN THIS SECTION

- Syntax | 9483
- Hierarchy Level | 9483
- Description | 9483
- Options | 9483
- Required Privilege Level | 9483

Release Information | 9483

Syntax

```
key number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number tunnel],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
tunnel]
```

Description

Identify an individual traffic flow within a tunnel, as defined in RFC 2890, *Key and Sequence Number Extensions to GRE*. On M Series and T Series routers, you can configure the GRE interface on an Adaptive Services, Multiservices, or Tunnel PIC. On MX Series routers, configure the interface on a Multiservices DPC.

Options

number—Value of the key.

- **Range:** 0 through 4,294,967,295

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Key Number on GRE Tunnels

key (interfaces)

IN THIS SECTION

- Syntax | 9484
- Hierarchy Level | 9484
- Description | 9484
- Options | 9485
- Required Privilege Level | 9485
- Release Information | 9485

Syntax

```
key number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number tunnel],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
tunnel]
```

Description

For Adaptive Services PICs on M Series routers (except the M320 and M120 routers), identify an individual traffic flow within a tunnel, as defined in RFC 2890, *Key and Sequence Number Extensions to GRE*.

Options

number—Value of the key.

- **Range:** 0 through 4,294,967,295

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Services Interfaces Library for Routing Devices

key (Authentication Keychain)

IN THIS SECTION

- Syntax | 9486
- Hierarchy Level | 9486
- Description | 9486
- Options | 9486
- Required Privilege Level | 9487
- Release Information | 9487

Syntax

```
key key-identifier {
    algorithm (hmac-sha-1  | md5);
    key-name authentication-key-name;
    options (basic | isis-enhanced);
    secret secret-data;
    start-time yyyy-mm-dd.hh:mm:ss;
}
```

Hierarchy Level

```
[edit security authentication-key-chains key-chain key-chain-name]
```

Description

Configure an authentication element (key). You include this statement several times in the configuration, thereby creating a keychain of authentication keys, each with its own identifier, secret (password), and start time. You can have up to 64 keys within a keychain.

Options

key-identifier

(Required) Each key within a keychain is identified by a unique integer value.

- **Range:** 0 through 63

algorithm (hmac-sha-1 | md5)

Configure the authentication algorithm for IS-IS.

- **Values:** Configure one of these authentication algorithms:
  - hmac-sha-1—96-bit hash-based message authentication code (SHA-1).
  - md5—Message digest 5.
- **Default:** md5

key-name *authentication-key-name*

Specify a key name in hexadecimal format, used for MACsec.

options (basic | isis-enhanced)

For IS-IS only, configure the protocol transmission encoding format for encoding the message authentication code in routing protocol packets.

Because this setting is for IS-IS only, the TCP and the BFD protocol ignore the encoding option configured in the key.

- **Values:** Configure one of the following:
  - basic—RFC 5304 based encoding. Junos OS sends and receives RFC 5304-encoded routing protocol packets, and drops 5310-encoded routing protocol packets that are received from other devices.
  - isis-enhanced—RFC 5310 based encoding. Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.
- **Default:** basic

secret ***secret-data***

(Required) Specify a password in encrypted text or plain text format. The secret password always appears in encrypted format. The password can include spaces if the character string is enclosed in quotation marks.

start-time ***yyyy-mm-dd.hh:mm:ss***

(Required) Specify a start time in UTC (Coordinated Universal Time) for key transmission. You do not need to specify an end time for the key. If a new key is present with a new start time, the keychain rolls over to the new one. The start time must be unique within the keychain.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Support for the BFD protocol introduced in Junos OS Release 9.6.

Support for IS-IS introduced in Junos OS Release 11.2.

`algorithm` and `options` introduced in Junos OS Release 11.2.

`key-name` introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols

Example: Configuring BFD Authentication for Securing Static Routes

Example: Configuring Hitless Authentication Key Rollover for IS-IS

Understanding Hitless Authentication Key Rollover for IS-IS

Configuring Media Access Control Security (MACsec) on Routers

key-chain (Authentication Keychain)

IN THIS SECTION

- Syntax | 9488
- Hierarchy Level | 9489
- Description | 9489
- Options | 9489
- Required Privilege Level | 9489
- Release Information | 9489

Syntax

```
key-chain key-chain-name {
    description text-string;
    key key {
        algorithm (md5 | hmac-sha-1);
        options (basic | isis-enhanced);
        key-name authentication-key-name;
        secret secret-data;
        start-time yyyy-mm-dd.hh:mm:ss;
    }
    tolerance seconds;
}
```

Hierarchy Level

```
[edit security authentication-key-chains]
```

Description

Create the key-chain configuration for the Border Gateway Protocol (BGP) and the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol.

Options

key-chain-name Authentication keychain name. It can be up to 32 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

description *text-string* Configure the description for this authentication keychain. Put the text string in quotes (“text description”).

tolerance *seconds* Configure the clock-skew tolerance in seconds for accepting keys for this authentication keychain.

- **Range:** 0 through 4294967295 seconds
- **Default:** 3600 seconds

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Support for the BFD protocol introduced in Junos OS Release 9.6.

Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.

Support for IS-IS introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols

Example: Configuring BFD Authentication for Securing Static Routes

Example: Configuring Hitless Authentication Key Rollover for IS-IS

Configuring Media Access Control Security (MACsec) on Routers

key-chain (Protocols PIM)

IN THIS SECTION

- Syntax | 9490
- Hierarchy Level | 9490
- Description | 9491
- Options | 9491
- Required Privilege Level | 9491
- Release Information | 9491

Syntax

```
key-chain key-chain-name;
```

Hierarchy Level

```
[edit protocols pim interface interface-name family {inet | inet6} bfd-liveness-detection
authentication],
[edit routing-instances routing-instance-name protocols pim interface interface-name family
{inet | inet6} bfd-liveness-detection authentication]
```

Description

Specify the security keychain to use for BFD authentication.

Options

key-chain-name—Name of the security keychain to use for BFD authentication. The name is a unique integer between **0** and **63**. This must match one of the keychains in the `authentication-key-chains` statement at the `[edit security]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Statement modified in Junos OS Release 12.2 to include **family** in the hierarchy level.

RELATED DOCUMENTATION

Configuring BFD Authentication for PIM

Understanding Bidirectional Forwarding Detection Authentication for PIM

authentication (Protocols PIM) | **2647**

key-exchange

IN THIS SECTION

- Syntax | **9492**
- Hierarchy Level | **9492**
- Description | **9492**
- Options | **9492**

Required Privilege Level | 9493
Release Information | 9494

Syntax

```
key-exchange [algorithm1 algorithm2...];
```

Hierarchy Level

```
[edit system services ssh]
```

Description

Specify the set of Diffie-Hellman key exchange methods that the SSH server can use.

Options

Specify one or more of the following Diffie-Hellman key exchange methods:

- `curve25519-sha256`—The EC Diffie-Hellman key exchange method on Curve25519 with SHA2-256.
- `dh-group1-sha1`—The Diffie-Hellman group1 algorithm using SHA-1.
- `dh-group14-sha1`—The Diffie-Hellman group14 algorithm using SHA-1.
- `ecdh-sha2-nistp256`—The ECDH key exchange method with ephemeral keys generated on the nistp256 curve.
- `ecdh-sha2-nistp384`—The ECDH key exchange method with ephemeral keys generated on the nistp384 curve.
- `ecdh-sha2-nistp521`—The ECDH key exchange method with ephemeral keys generated on the nistp521 curve.
- `group-exchange-sha1`—The group exchange algorithm using SHA-1.
- `group-exchange-sha2`—The group exchange algorithm using SHA-2.

NOTE: The key-exchange represents a set. To configure key-exchange:

```
user@host#set system services ssh key-exchange [ecdh-sha2-nistp256 group-exchange-sha1]
```

NOTE: Table 53 on page 9493 shows the supportability of Diffie-Hellman key exchange methods on FIPS mode.

Table 53: Supportability of Diffie-Hellman key exchange methods on FIPS mode

Diffie-Hellman key exchange methods	Supported on FIPS mode
curve25519-sha256	No
dh-group1-sha1	No
dh-group14-sha1	Yes
ecdh-sha2-nistp256	Yes
ecdh-sha2-nistp384	Yes
ecdh-sha2-nistp521	Yes
group-exchange-sha1	No
group-exchange-sha2	No

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2. Support for the curve25519-sha256 option added in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Configuring SSH Service for Remote Access to the Router or Switch

key (Junos FIPS)

IN THIS SECTION

- Syntax | 9494
- Hierarchy Level | 9494
- Description | 9495
- Options | 9495
- Required Privilege Level | 9495
- Release Information | 9495

Syntax

```
key (ascii-text key | hexadecimal key);
```

Hierarchy Level

```
[edit security ipsec internal security-association manual direction encryption]
```

Description

The key used for the internal Routing-Engine–to–Routing-Engine IPsec security association (SA) configuration.

Options

ascii-text-key—The encrypted ASCII text key.

`hexadecimal key`—The encrypted hexadecimal key.

Required Privilege Level

Crypto Officer—To add and view this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

key-length

IN THIS SECTION

- Syntax | 9496
- Hierarchy Level | 9496
- Description | 9496
- Default | 9496
- Options | 9496
- Required Privilege Level | 9496
- Release Information | 9496

Syntax

```
key-length number;
```

Hierarchy Level

```
[edit protocols neighbor-discovery secure cryptographic-address]
```

Description

Specify the length of the RSA key used to generate the public-private key pair for the cryptographic address.

Default

1024

Options

number—RSA key length.

- **Range:** 1024 through 2048

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Secure IPv6 Neighbor Discovery

key (MACsec)

IN THIS SECTION

- Syntax | 9497
- Hierarchy Level | 9497
- Description | 9497
- Default | 9498
- Options | 9498
- Required Privilege Level | 9498
- Release Information | 9498

Syntax

```
key key-string;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel secure-channel-name security-association security-association-number]
```

Description

Specifies the static security key to exchange to enable MACsec using static secure association key (SAK) security mode.

The key string is a 32-digit hexadecimal number. The key string and the security association must match on both sides of an Ethernet connection to secure traffic using MACsec when enabling MACsec using SAK security mode.

You must configure at least two security associations with unique security association numbers and key strings to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec

automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.

Default

This statement does not have a default value.

Options

key-string Specifies the key to exchange with the other end of the link on the secure channel. The *key-string* is a 32-digit hexadecimal string that is created by the user.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

key (MACsec for MX Series)

IN THIS SECTION

- Syntax | 9499
- Hierarchy Level | 9499
- Description | 9499

- Default | 9499
- Options | 9500
- Required Privilege Level | 9500
- Release Information | 9500

Syntax

```
key key-string;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel
secure-channel-name security-association security-association-number]
```

Description

Specifies the static security key to exchange to enable MACsec using static secure association key (SAK) security mode.

The key string is a 32-digit hexadecimal number. The key string and the security association must match on both sides of an Ethernet connection to secure traffic using MACsec when enabling MACsec using SAK security mode.

You must configure at least two security associations with unique security association numbers and key strings to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.

Default

This statement does not have a default value.

Options

key-string Specifies the key to exchange with the other end of the link on the secure channel. The *key-string* is a 32-digit hexadecimal string that is created by the user.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

key-pair

IN THIS SECTION

- Syntax | 9501
- Hierarchy Level | 9501
- Description | 9501
- Default | 9501
- Options | 9501
- Required Privilege Level | 9501
- Release Information | 9501

Syntax

```
key-pair pathname;
```

Hierarchy Level

```
[edit protocols neighbor-discovery secure cryptographic-address]
```

Description

Specify the directory path of the public-private key file generated for the cryptographic address.

A cryptographic address is dynamically generated based on a public key and a subnet prefix.

Default

The default location of the file is the **/var/etc/rsa_key** directory.

Options

pathname—Directory path of the public-private key file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Secure IPv6 Neighbor Discovery

key-server-priority (MACsec)

IN THIS SECTION

- Syntax | 9502
- Hierarchy Level | 9502
- Description | 9502
- Default | 9502
- Options | 9503
- Required Privilege Level | 9503
- Release Information | 9503

Syntax

```
key-server-priority priority-number;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name mka]
```

Description

Specifies the key server priority used by the MACsec Key Agreement (MKA) protocol to select the key server when MACsec is enabled using static connectivity association key (CAK) security mode.

The switch with the lower *priority-number* is selected as the key server.

If the *priority-number* is identical on both sides of a point-to-point link, the MKA protocol selects the device with the lower MAC address as the key server.

Default

The default key server priority number is 16.

Options

priority-number Specifies the MKA server election priority number.

The *priority-number* can be any number between 0 and 255. The lower the number, the higher the priority.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced for SRX Series Firewalls in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

key-server-priority (MACsec for MX Series)

IN THIS SECTION

- Syntax | **9504**
- Hierarchy Level | **9504**
- Description | **9504**
- Default | **9504**
- Options | **9504**

Required Privilege Level | 9504
Release Information | 9505

Syntax

```
key-server-priority priority-number;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name mka]
```

Description

Specifies the key server priority used by the MACsec Key Agreement (MKA) protocol to select the key server when MACsec is enabled using static connectivity association key (CAK) security mode.

The switch with the lower *priority-number* is selected as the key server.

If the *priority-number* is identical on both sides of a point-to-point link, the MKA protocol selects the device with the lower MAC address as the key server.

Default

The default key server priority number is 16.

Options

priority-number

Specifies the MKA server election priority number.

The *priority-number* can be any number between 0 and 255. The lower the number, the higher the priority.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

krt-nexthop-ack

IN THIS SECTION

- Syntax | 9505
- Hierarchy Level | 9505
- Description | 9506
- Required Privilege Level | 9506
- Release Information | 9506

Syntax

```
krt-nexthop-ack;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table],
[edit routing-options forwarding-table]
```

Description

When you configure this option, the routing protocol process (rpd) waits for the next hop to get acknowledged by PFE before using it for a route.

Currently, only BGP-labeled routes and RSVP routes support this option. All other routes will ignore this option.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1.

RELATED DOCUMENTATION

Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine

indirect-next-hop-change-acknowledgements | **8551**

krt-nexthop-ack-timeout | **9506**

krt-nexthop-ack-timeout

IN THIS SECTION

- Syntax | **9507**
- Hierarchy Level | **9507**
- Description | **9507**
- Options | **9507**
- Required Privilege Level | **9507**
- Release Information | **9507**

Syntax

```
krt-nexthop-ack-timeout interval;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table],
[edit routing-options forwarding-table]
```

Description

For indirect next-hop and multicast next-hop change acknowledgements, configure the time interval for which to wait for the next-hop acknowledgement. The routing protocol process (rpd) waits for the specified time period before changing the route to point to the new next hop.

If the acknowledgement is received before the specified time, then rpd updates the route to point to the indirect next-hop and multicast.

If the acknowledgement is not received within the time period, it is assumed to have been received and the route is made to point to the new next hop.

Options

interval—Kernel next-hop acknowledgement timeout interval.

- **Range:** 1 through 100 seconds
- **Default:** 1 second

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine

indirect-next-hop-change-acknowledgements | **8551**

L

IN THIS SECTION

- l2-channel-error-threshold (chassis) | **9520**
- l2-channel-errors (chassis) | **9521**
- l2-domain-id-for-l3 | **9523**
- l2-learning | **9524**
- l2-querier | **9527**
- l2-smart-policy | **9528**
- l2-stats (Flat-File Accounting Options) | **9530**
- l2circuit | **9531**
- l2circuit-control-passthrough | **9534**
- l2ckt | **9535**
- l2pt (Destination Tunnels) | **9537**
- l2tp | **9541**
- l2tp (DDoS Protection) | **9544**
- l2tp-access-profile | **9547**
- l2tp (Group Profile) | **9548**
- l2tp-interface-id | **9550**
- l2tp-maximum-session (Service Interfaces) | **9551**
- l2tp (Profile) | **9553**
- l2tp-tunnel-session-identifier | **9558**
- l2vpn | **9559**
- l2vpn-id | **9562**
- l2vpn (routing-options) | **9564**

- l2vpn-use-bgp-rules-edit-protocols-bgp | **9565**
- l3-interface-ingress-counting | **9567**
- l3-interface (VLAN) | **9568**
- l3vpn (ingress) | **9571**
- l3vpn (transit) | **9573**
- label | **9574**
- label-allocation | **9576**
- label-block-size | **9577**
- label-position | **9579**
- label-switched-path | **9580**
- label-switched-path (Protocols IS-IS) | **9582**
- label-switched-path (Protocols Link Management) | **9583**
- label-switched-path (Protocols MPLS) | **9585**
- label-switched-path (Protocols MVPN) | **9590**
- label-switched-path (Protocols OSPF) | **9592**
- label-switched-path-template | **9593**
- label-switched-path-template (Container LSP) | **9595**
- label-switched-path-template (Multicast) | **9596**
- label-withdrawal-delay | **9598**
- labeled-bgp | **9600**
- labeled-unicast | **9601**
- labeled-unicast (Protocols BGP) | **9603**
- lacp | **9605**
- lacp (802.3ad) | **9607**
- lacp (802.3ad) (MX Series) | **9609**
- lacp (Aggregated Ethernet) | **9610**
- lacp (Interfaces) | **9614**
- lacp (Protocols) | **9615**
- land | **9617**
- land-attack-check | **9619**
- land-attack-check (IDS MS-MPC) | **9620**
- land (IDS Screen Next Gen Services) | **9622**

- lane | **9623**
- lane-all | **9626**
- language (Scripts) | **9630**
- large | **9632**
- large (IDS Screen Next Gen Services) | **9633**
- large-community-count | **9634**
- laser-enable | **9635**
- last-block-recycle-timeout | **9637**
- latency (Interfaces) | **9638**
- latency-threshold | **9639**
- layer2-control | **9641**
- layer2 (enhanced-hash-key) | **9644**
- layer2-liveness-detection (Receive) | **9647**
- layer2-liveness-detection (Send) | **9649**
- layer2-map | **9651**
- layer2-policer | **9653**
- layer2-policer (EX) | **9654**
- layer2-policer (Hierarchical Policer) | **9656**
- layer2-protocol-tunneling | **9658**
- layer2-unicast-replies | **9661**
- layer2-policer (firewall) | **9662**
- layer3-domain-identifier | **9666**
- lcd-menu | **9667**
- lcp-max-conf-req | **9669**
- lcp-renegotiation | **9671**
- lcp-restart-timer | **9672**
- ldap-options | **9673**
- ldap-server | **9676**
- ldap-url | **9678**
- ldp | **9679**
- ldp-p2mp | **9685**
- ldp-stitching (Protocols OSPF) | **9686**

- ldp-tunneling | **9687**
- leaf | **9689**
- leaf-tunnel-limit-inet (MVPN Selective Tunnels) | **9691**
- leaf-tunnel-limit-inet6 (MVPN Selective Tunnels) | **9693**
- learn-sip-register | **9695**
- lease-time | **9696**
- lease-time (dhcp-client) | **9698**
- lease-time-threshold (DHCP Local Server and DHCP Relay Agent) | **9699**
- lease-time-validation (DHCP Local Server and DHCP Relay Agent) | **9701**
- leasequery (DHCP Relay Agent) | **9702**
- least-fill | **9704**
- leave-timer (MVRP) | **9705**
- leaveall-timer (MVRP) | **9707**
- led-beacon | **9710**
- legacy-redirect-ip-action | **9711**
- less-than | **9713**
- level (Connectivity Fault Management) | **9715**
- level (OAM) | **9716**
- level (Disk Partition) | **9718**
- level (App Engine Virtual Machine Management Service) | **9719**
- level (CFM MEP) | **9721**
- level (EX Series Switch Only) | **9722**
- level (Global IS-IS) | **9724**
- level (IS-IS Interfaces) | **9728**
- level (Protocols RIFT) | **9734**
- level (Services) | **9735**
- license | **9737**
- license-mode | **9739**
- license-server | **9740**
- lie-receive-address (Protocols RIFT) | **9742**
- lie-transmit-address (Protocols RIFT) | **9744**
- lifetime (Router Advertisement) | **9745**

- lifetime (Router Discovery) | **9747**
- lifetime-seconds | **9749**
- lifetime-seconds (Security) | **9750**
- light | **9752**
- light-weight-dhcpv6-relay | **9755**
- limit | **9757**
- limit (Content Security Policy) | **9758**
- limit-ports-per-address | **9759**
- limit (Security Log) | **9760**
- limit (Security SCTP) | **9762**
- limit-session | **9764**
- limit-session (IDS Screen Next Gen Services) | **9765**
- limited-ifl-scaling | **9768**
- line-encoding (E1 Options) | **9769**
- line-encoding (T1 Options) | **9770**
- line-identity (Dynamic Access-Line-Identifier VLANs) | **9772**
- line-interval (Dampen Options) | **9773**
- line-loopback | **9775**
- line-protocol | **9776**
- line-rate | **9778**
- line-rate (Interfaces) | **9779**
- linerate-mode | **9780**
- link (Access) | **9781**
- link-adjacency-loss | **9783**
- link-degrade-monitor | **9784**
- link-discovery | **9787**
- link-down (QFX Series) | **9788**
- link-down | **9790**
- link-event-rate | **9791**
- link-fault-management | **9792**
- link-group-protection (Protocols IS-IS) | **9794**
- link-layer-broadcast-inet-check | **9796**

- link-layer-overhead | **9797**
- link-management | **9798**
- link-mode | **9800**
- link-mtu | **9802**
- link-protection | **9803**
- link-protection (Chassis Cluster) | **9805**
- link-protection (Dynamic LSPs) | **9807**
- link-protection (Host Fast Reroute) | **9808**
- link-protection (non-LACP) | **9810**
- link-protection (Protocols IS-IS) | **9811**
- link-protection (Protocols LACP) | **9813**
- link-protection (Protocols OSPF) | **9814**
- link-protection (RSVP) | **9816**
- link-protection (Static LSPs) | **9818**
- link-protection-sub-group (802.3ad) | **9820**
- link-protection-sub-group (aggregated-ether-options) | **9821**
- link-speed | **9823**
- link-speed (Aggregated Ethernet) | **9826**
- link-speed (Aggregated SONET/SDH) | **9828**
- link-speed (Interfaces) | **9830**
- link-to-disable | **9831**
- link-to-monitor | **9832**
- link-tracking (Network Isolation Groups) | **9833**
- linkdown (satellite-policies alarm) | **9835**
- linked-pool-aggregation (Address-Assignment Pools) | **9837**
- linktrace | **9838**
- list | **9839**
- list (Security Content Filtering Block Mime) | **9841**
- listen | **9842**
- listening-mode | **9843**
- liveness-detection (iccp-peer) | **9845**
- liveness-detection | **9846**

- liveness-detection (high availability) | **9848**
- lldp | **9851**
- lldp-med (Ethernet Switching) | **9857**
- lldp-priority | **9860**
- lmi (Ethernet OAM) | **9861**
- lmi-type | **9863**
- lmp-control-channel | **9864**
- lmp-protocol | **9865**
- load | **9867**
- load access-security slaac-snooping persistent-file | **9870**
- load-balance (ecmp) | **9872**
- load-balance | **9874**
- load-balance-algorithm | **9876**
- load-balance (Forwarding Options) | **9878**
- load-balance-group | **9880**
- load-balance-label-capability | **9882**
- load-balance profile1/profile2 | **9883**
- load-balance (Protocols RSVP) | **9885**
- load-balance-stateful (Aggregated Ethernet Interfaces) | **9886**
- load-balancing-options (Aggregated Multiservices) | **9888**
- load-balancing-options (Service Set) | **9890**
- load-interval | **9892**
- load-scripts-from-flash (Scripts) | **9893**
- load-threshold | **9894**
- load-type (Aggregated Ethernet Interfaces) | **9896**
- loc-rib | **9897**
- local | **9899**
- local (certificates) | **9900**
- local-address (Junos Telemetry Interface) | **9902**
- local-address (Protocols AMT) | **9904**
- local-address (Protocols BGP) | **9905**
- local-address (Protocols Link Management) | **9908**

- local-address (Protocols MSDP) | **9910**
- local-address (Protocols PIM) | **9911**
- local-address (Routing Options) | **9913**
- local (Analytics Collector) | **9914**
- local-as | **9916**
- local-authentication (Dynamic PPP Options) | **9919**
- local-bias (ae) | **9921**
- local-bias | **9922**
- local-bias (ae load-balance) | **9923**
- local-bw-override-threshold | **9925**
- local-category (Next Gen Services Service-Set Local System Logging) | **9926**
- local-certificate | **9929**
- local-certificate (System) | **9930**
- local-certificate (Security) | **9932**
- local-chap | **9933**
- local-decision (PCRF Partition) | **9934**
- local-dump (Flow Aggregation) | **9938**
- local-endpoint | **9939**
- local-engine | **9941**
- local (Flat-File Access Profile) | **9942**
- local-gateway (IPSec) | **9944**
- local-gateway (L2TP LNS) | **9945**
- local-id | **9947**
- local-id (High Availability) | **9948**
- local-identity | **9950**
- local-interface (IPv6) | **9952**
- local-ip-addr (ICCP) | **9953**
- local-ipv4-address | **9955**
- local-ipv6-address | **9956**
- local-key-pair | **9958**
- local-log-tag (Next Gen Services Service-Set System Logging) | **9959**
- local-loopback | **9960**

- local-minimum-links-threshold | **9962**
- local-name | **9964**
- local-password | **9966**
- local-port-range | **9967**
- local-ports | **9969**
- local-preference | **9971**
- local-priority | **9973**
- local-replication | **9974**
- local-server-group (DHCP Relay Agent Option) | **9976**
- local-switching (Layer 2 Circuits) | **9977**
- local-switching (VPLS) | **9979**
- local-traffic-selector | **9980**
- locality-bias (MX Series Virtual Chassis) | **9981**
- localized-fib | **9983**
- location (DHCP Snooping Database) | **9984**
- location (DHCP Relay Agent) | **9985**
- location (DHCP Snooping Database) | **9987**
- location (DHCPv6 Relay Agent) | **9988**
- location (LLDP-MED) | **9990**
- location (SNMP) | **9992**
- location (System) | **9993**
- location (Virtual Chassis) | **9995**
- locator | **9997**
- lockout | **9999**
- lockout-timeout (L2TP Destination Lockout) | **10001**
- log | **10002**
- log (aacl) | **10005**
- log-and-generate-ais | **10006**
- log-attacks | **10007**
- log-interval (IGMP Interface) | **10009**
- log-interval (MLD Interface) | **10011**
- log-interval (PIM Entries) | **10012**

- log-interval (Protocols MSDP) | **10014**
- log (Logical Systems and Tenant Systems) | **10016**
- log-prefix (Interfaces) | **10020**
- log-prefix (L2TP) | **10021**
- log-prefix (Services) | **10022**
- log-prefix (System) | **10023**
- log-rotate-frequency | **10025**
- log (Security) | **10026**
- log (Security GTP) | **10032**
- log (Security IDP Sensor Configuration) | **10033**
- log (Security Policies) | **10035**
- log (Security SCTP) | **10037**
- log (Services) | **10038**
- log (Spanning Trees) | **10040**
- log-updown | **10042**
- log-updown (Protocols LDP) | **10043**
- log-updown (Protocols MPLS) | **10045**
- log-warning (Multicast Forwarding Cache) | **10047**
- log-warning (Protocols MSDP) | **10049**
- log-attacks | **10050**
- log-attacks (rulebase-exempt) | **10052**
- logging (IDS MS-DPC) | **10053**
- logging-rule (PCC Action Profile) | **10054**
- logging (Services) | **10056**
- logical-bandwidth-policer | **10058**
- logical-domain-identity-management | **10059**
- logical-interface-aggregate-statistics | **10062**
- logical-interface-chassis-redundancy (MX Series Virtual Chassis) | **10063**
- logical-interface (DDoS Flow Detection) | **10064**
- logical-interface-fpc-redundancy (Aggregated Ethernet Subscriber Interfaces) | **10067**
- logical-interface-fpc-redundancy (PS Interface) | **10068**
- logical-interface-policer | **10069**

- logical-interface-unit-range | **10072**
- logical-system | **10073**
- logical-system (SNMP) | **10075**
- logical-system (Diameter Peer) | **10076**
- logical-system (Diameter Transport) | **10078**
- logical-system-mux | **10079**
- logical-system-name (DHCP Local Server) | **10082**
- logical-system-name (DHCP Relay Agent) | **10083**
- logical-system-name (Static Subscribers) | **10085**
- logical-system (System Security Profile) | **10086**
- logical-system-trap-filter | **10088**
- logical-system (Tunnel Profile) | **10089**
- logical-systems | **10090**
- logical-systems (All) | **10092**
- logical-tunnel-options | **10093**
- login | **10095**
- login (Access) | **10100**
- logout-response-timeout (PCRF Partition) | **10101**
- long-lived (Graceful Restart for BGP Helper) | **10103**
- long-lived (Graceful Restart for BGP Restarter) | **10105**
- longest-match | **10108**
- longest-match (Static Routes) | **10109**
- lookup-intact-on-commit | **10112**
- loop-detect | **10113**
- loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) | **10117**
- loopback-clear-timer | **10119**
- loopback (Fibre Channel Interface) | **10120**
- loopback-firewall-optimization | **10121**
- loopback-firewallv6-optimization | **10123**
- loopback (Local and Remote) | **10125**
- loopback-port | **10126**
- loopback-remote | **10128**

- loopback (Serial) | **10129**
- loopback-tracking | **10131**
- loops (Autonomous System) | **10132**
- loops (BGP Address Family) | **10134**
- loose-authentication-check | **10137**
- loose-check | **10138**
- loose-source-route-option (IDS Screen Next Gen Services) | **10139**
- loss-delay (querier) | **10140**
- loss-priority | **10142**
- loss-priority (Classifiers) | **10143**
- loss-priority (Classifiers) | **10145**
- loss-priority (Classifiers and Rewrite Rules) | **10146**
- loss-priority (Frame Relay Loss Priority) | **10148**
- loss-priority (Drop Profiles) | **10150**
- loss-priority (Dynamic Schedulers) | **10151**
- loss-priority (Firewall Filter) | **10152**
- loss-priority (Firewall Filter Action) | **10154**
- loss-priority high then discard (Three-Color Policer) | **10155**
- loss-priority-maps | **10157**
- loss-priority-maps (Assigning to an Interface) | **10159**
- loss-priority-maps (CoS) | **10160**
- loss-priority-maps (CoS Interfaces) | **10161**
- loss-priority (Rewrite Rules) | **10163**
- loss-priority (Rewrite Rules) | **10164**
- loss-priority-rewrites | **10166**
- loss-priority-rewrites (Assigning to an Interface) | **10167**
- loss-priority (Scheduler Drop Profiles) | **10169**
- loss-priority (Simple Firewall Filter) | **10170**
- loss (querier) | **10171**
- loss (responder) | **10173**
- loss-threshold | **10175**
- low-latency | **10176**

- low-watermark | **10177**
- lowest-priority-defect | **10179**
- lrf-profile (Service Set) | **10180**
- lsa-refresh-interval | **10182**
- lsp-attributes | **10183**
- lsp-cleanup-timer | **10185**
- lsp-external-controller | **10187**
- lsp-lifetime | **10189**
- lsp-next-hop (CoS-Based Forwarding) | **10191**
- lsp-next-hop (Static Routes) | **10192**
- lsp-switch | **10194**
- lsping-channel-type | **10195**
- lsq-failure-options | **10197**
- ltv-syslog-interval (System Process) | **10198**

l2-channel-error-threshold (chassis)

IN THIS SECTION

- Syntax | **9521**
- Hierarchy Level | **9521**
- Description | **9521**
- Required Privilege Level | **9521**
- Release Information | **9521**

Syntax

```
l2-channel-error-threshold threshold; {
            l2-channel-errors (ignore |red | yellow);
    }
```

Hierarchy Level

```
[edit chassis alarm l2-chan-err-ethernet]
```

Description

Use this command to specify the threshold limit for the L2 error count. Range is 1 through 4294967295 per hour and default value is 50,000 per hour.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Chassis Conditions That Trigger Alarms

l2-channel-errors (chassis)

IN THIS SECTION

Syntax | 9522

- Hierarchy Level | 9522
- Description | 9522
- Options | 9522
- Required Privilege Level | 9522
- Release Information | 9523

Syntax

```
l2-channel-errors (ignore |red | yellow);
```

Hierarchy Level

```
[edit chassis alarm l2-chan-err-ethernet]
```

Description

Configure chassis alarm for L2 channel error. When you configure this statement, chassis raises an alarm when the L2 channel error count crosses 50,000 per hour. If you do not configure this statement, chassis do not raise any alarm for L2 channel errors.

Options

`ignore`—This configuration does not raise any alarm for L2 channel error. This is the default configuration.

`red`—This configuration raises a red alarm for L2 channel error.

`yellow`—This configuration raises a red alarm for L2 channel error.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Chassis Conditions That Trigger Alarms

l2-domain-id-for-l3

IN THIS SECTION

- Syntax | 9523
- Hierarchy Level | 9523
- Description | 9523
- Options | 9524
- Required Privilege Level | 9524
- Release Information | 9524

Syntax

```
l2-domain-id-for-l3 id;
```

Hierarchy Level

```
[edit routing-instances instance-name]
```

Description

Specify a Layer 2 domain ID within a routing instance.

Options

`id`—Layer 2 identification number.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R2.

RELATED DOCUMENTATION

Configuring a Layer 2 Virtual Switch on an EX Series Switch

l2-learning

IN THIS SECTION

- Syntax (MX Series, QFX Series, EX Series) | 9525
- Syntax (SRX Series) | 9525
- Hierarchy Level | 9525
- Description | 9525
- Options | 9526
- Required Privilege Level | 9526
- Release Information | 9526

Syntax (MX Series, QFX Series, EX Series)

```
l2-learning {
    global-le-bridge-domain-aging-time;
    global-mac-ip-limit number;
    global-mac-ip-table-aging-time seconds;
    global-mac-limit limit;
    global-mac-statistics;
    global-mac-table-aging-time seconds;
    global-no-mac-learning;
    global-mac-move;
    global-ctx-limit;
}
```

Syntax (SRX Series)

```
l2-learning {
    global-mac-limit limit {
    packet-action-drop
    }
    global-mac-table-aging-time seconds;
    global-mode (switching | transparent-bridge) ;
    global-no-mac-learning;
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure Layer 2 address learning and forwarding properties globally.

The remaining statements are explained separately. See CLI Explorer.

Options

global-le-bridge-domain-aging-time

Specify the aging time of LE bridge-domain. The MAC address is learnt after next hop(NH) and bridge-domain(BD), also called NHBD. This aging time delays the deletion of NHBD. Configuring lesser time, in seconds, results in faster deletion of NHBD.

- **Range:** 120 to 1000000 seconds

global-ctx-limit

Specify the debug context history limit. It is the limit for the number of events that are logged in the process memory for debug purposes. The value '0' to '5000000' is the number of events captured in the process memory. Each event entry takes upto 200 Bytes of memory. By default, 256000 entries are saved to log the events that happen in the Layer 2 process. You may disable the feature by setting `global-ctx-limit` value to '0' and the memory that is currently used for context is freed up and no further event logging happens.

- **Range:** 0 to 5000000 event entries
- **Default:** 256000 event entries

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement modified in Junos OS Release 9.5. Support for global mode added in Junos OS Release 15.1X49-D40.

`global-le-bridge-domain-aging-time` option introduced in Junos OS Release 14.2R5 for the MX Series.

`global-mac-ip-limit` and `global-mac-ip-table-aging-time` options introduced in Junos OS Release 17.4R1 for MX Series routers and EX9200 switches.

`global-ctx-limit` option introduced in Junos OS Release 20.2R3.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding

global-mac-table-aging-time
global-mac-limit (Protocols)
global-no-mac-learning
global-mode (Protocols)

l2-querier

IN THIS SECTION

- Syntax | 9527
- Hierarchy Level | 9527
- Description | 9527
- Options | 9528
- Required Privilege Level | 9528
- Release Information | 9528

Syntax

```
l2-querier {
    source-address ip-address;
}
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan]
```

Description

Configure the device to be an IGMP querier. IGMP querier allows the device to proxy for a multicast router and send out periodic IGMP queries in the network. This action causes the device to consider itself an multicast router port. The remaining devices in the network simply define their respective

multicast router ports as the interface on which they received this IGMP query. Use the `source-address` statement to configure the source address to use for IGMP snooping queries.

Options

seconds—Time interval.

- **Range:** 1 through 1024
- **Default:** 125 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping on SRX Series Devices

IGMP Snooping Overview

igmp-snooping | 8387

l2-smart-policy

IN THIS SECTION

- Syntax | 9529
- Hierarchy Level | 9529
- Description | 9529
- Required Privilege Level | 9529
- Release Information | 9529

Syntax

```
l2-smart-policy;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Prevent LDP from exporting IPv4 FECs over sessions with Layer 2 neighbors only. IPv4 FECs received over such sessions are filtered out.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring LDP IPv4 FEC Filtering

l2-stats (Flat-File Accounting Options)

IN THIS SECTION

- Syntax | 9530
- Hierarchy Level | 9530
- Description | 9530
- Options | 9530
- Required Privilege Level | 9531
- Release Information | 9531

Syntax

```
l2-stats {
    all-fields;
    input-mcast-bytes;
    input-mcast-packets;
}
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name fields]
```

Description

Specify the statistics to collect for the named flat-file-profile field.

Options

`all-fields`—Collect all Layer 2 statistics for the named flat-file profile.

`input-mcast-bytes`—Collect multicast bytes from the input side for the named flat-file profile.

`input-mcast-packets`—Collect multicast packets from the input side for the named flat-file profile.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Flat-File Accounting Overview

l2circuit

IN THIS SECTION

- Syntax | 9531
- Hierarchy Level | 9533
- Description | 9533
- Required Privilege Level | 9533
- Release Information | 9533

Syntax

```
l2circuit {
    auto-sensing{
        password password;
        }
    local-switching {
        interface interface-name {
```

```
            description text;
            end-interface {
                interface interface-name;
                protect-interface interface-name;
            }
            ignore-mtu-mismatch;
            protect-interface interface-name;
        }
    }
    neighbor address {
        interface interface-name {
            backup-neighbor address;
            bandwidth (bandwidth | ctnumber bandwidth);
            community community-name;
            connection-protection;
            (control-word | no-control-word);
            description text;
            egress-protection;
            encapsulation-type type;
            ignore-encapsulation-mismatch;
            ignore-mtu-mismatch;
            mtu mtu-number;
            protect-interface interface-name;
            pseudowire-status-tlv hot-standby-vc-on;
            psn-tunnel-endpoint address;
            virtual-circuit-id identifier;
        }
    }
    resolution {
        preserve-nexthop-heirarchy;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit protocols]
```

Description

Enables a Layer 2 circuit.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Configuring ATM Trunking on Layer 2 Circuits*
- *Configuring Bandwidth Allocation and Call Admission Control in Layer 2 Circuits*
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring LDP for Layer 2 Circuits*
- *Configuring Policies for Layer 2 Circuits*
- *Configuring Static Layer 2 Circuits*
- *Tracing Layer 2 Circuit Operations*

l2circuit-control-passthrough

IN THIS SECTION

- Syntax | 9534
- Hierarchy Level | 9534
- Description | 9534
- Default | 9535
- Required Privilege Level | 9535
- Release Information | 9535

Syntax

```
l2circuit-control-passthrough;
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure the device to allow LACP, LLDP, OAM LFM, and OAM CFM packets to cross the Layer 2 circuit. If the `l2circuit-control-passthrough` statement is not configured, LACP, LLDP, OAM LFM, and OAM CFM packets are classified as control packets and are not transmitted across the Layer 2 circuit.

NOTE: For MX Series routers, the functionality that the `l2circuit-control-passthrough` command provides is performed automatically.

Starting in Junos OS Evolved Release 22.4R1, you do not need to configure the `l2circuit-control-passthrough` statement to allow tunneling of Layer 2 control frames between local and remote provider edge (PE) devices in a Layer 2 VPN on PTX10001-36MR, PTX10004, PTX10008, and PTX10016 routers.

Default

By default, this statement is not configured.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

forwarding-options (Bridge Domains or VLANs) | 7467

Configuring Layer 2 Switching Cross-Connects Using CCC

Configuring an MPLS-Based VLAN CCC with Pop, Push, and Swap and Control Passthrough

l2ckt

IN THIS SECTION

- Syntax | 9536
- Hierarchy Level | 9536
- Description | 9536
- Required Privilege Level | 9536
- Release Information | 9536

Syntax

```
l2ckt {
    l2vpn;
    l3vpn;
    labeled-bgp;
    no-l2vpn;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-
next-hop ingress],
[edit routing-options forwarding-table chained-composite-next-hop ingress]
```

Description

Enable composite chained next hop for ingress Layer 2 circuit label-switched paths (LSPs).

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

l2pt (Destination Tunnels)

IN THIS SECTION

- Syntax | 9537
- Hierarchy Level | 9538
- Description | 9538
- Options | 9538
- Required Privilege Level | 9540
- Release Information | 9540

Syntax

```
l2pt {
    interface name {
        destination {
            vxlan-tunnel;
        }
        enable-all-ifl;
        protocol {
            all;
            cdp;
            elmi;
            gvrp;
            ieee8021x;
            ieee8023ah;
            lacp;
            lldp;
            mmrp;
            mvrp;
            pvstp;
            stp;
            udld;
            vstp;
            vtp;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit protocols layer2-control]
```

Description

Enable Layer 2 (L2) protocol tunneling (L2PT) across the encapsulation tunnels in an EVPN network.

Use this statement to enable a device to receive protocol control frames coming in on an access interface and tunnel those frames across the encapsulation tunnels on the device. Configure this statement with the `interface` *`name`* option (required) and each of the following (required):

1. The `destination` *`tunnel-type`* option.
2. The `protocol` option for each bridge protocol data unit (BPDU) type you want the device to tunnel from the access interface across the encapsulation tunnels on the device. Or, specify the `all` option to enable L2PT for all of the supported protocols.

> **NOTE**: When you enable L2PT over VXLAN for a supported protocol, the device transparently forwards the protocol control frames with that protocol destination MAC address. If you enable L2PT for a protocol that has the same destination MAC address as one or more other supported protocols, the device transparently forwards the control frames for all of those protocols.

Options

`interface` ***`name`***
(Required) Interface name on which to enable L2PT across a destination tunnel type on the device. The interface can be a physical interface or an aggregated Ethernet (AE) interface.

`destination` ***`tunnel-type`***
(Required) Destination tunnel type over which to enable L2PT—you must specify one of the supported encapsulation tunnel types:

- `vxlan-tunnel`—Enable L2PT over VXLAN encapsulation tunnels on the device.

> **NOTE**: The `vxlan-tunnel` tunnel type is currently the only destination encapsulation tunnel type supported with this statement.

enable-all-ifl

Enable L2PT for all the logical interfaces under the specified L2PT-enabled physical interface.

This option is required on QFX5130 and QFX5700 switches to ensure any logical interfaces you configure on an L2PT-enabled physical interface end up with the same L2PT settings when you commit the configuration.

protocol (all | *protocol-name*)

(Required) Protocol for which to transparently tunnel control packets from the specified interface across the destination tunnel type. Some of the available protocol options might not be supported with all encapsulation tunnel types. We support the protocol options listed below with the `vxlan-tunnel` destination tunnel type.

> **NOTE**: EX Series and QFX Series with Junos OS don't include the `pvstp`, `stp`, and `vstp` options listed below for tunneling STP, MSTP, RSTP, PVST/PVST+, and VSTP BPDUs (called xSTP BPDUs here for brevity). These devices flood xSTP control BPDUs into the VXLAN tunnels by default.
>
> If you don't want these devices to tunnel xSTP control BPDUs over VXLAN, you can configure the BPDU protection feature on the ingress access interface. See "bpdu-block" on page 3022 and "bpdu-block-on-edge" on page 3024 for details on the BPDU protection feature.

Specify the `all` option to enable L2PT for all of the supported protocol options, or specify one of the following supported protocol options:

- cdp—Cisco Discovery Protocol (CDP)
- elmi—Ethernet Local Management Interface (E-LMI)
- gvrp—Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP)
- ieee8021x—IEEE 802.1X authentication
- ieee8023ah—IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM)
- lacp—Link Aggregation Control Protocol (LACP)

- lldp—Link Layer Discovery Protocol (LLDP)
- mmrp—Multiple MAC Registration Protocol (MMRP)
- mvrp—MVRP VLAN Registration Protocol (MVRP)
- pvstp—Per-VLAN Spanning Tree (PVST) and Per-VLAN Spanning Tree Plus (PVST+) protocols (Junos OS Evolved QFX Series switches only)
- stp—Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP) (Junos OS Evolved QFX Series switches only)
- udld—Unidirectional Link Detection (UDLD)
- vstp—VLAN Spanning Tree Protocol (VSTP) (Junos OS Evolved QFX Series switches only)
- vtp—VLAN Trunking Protocol (VTP)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Layer 2 Protocol Tunneling over VXLAN Tunnels in EVPN-VXLAN Bridged Overlay Networks

show l2pt interface | **21816**

l2tp

IN THIS SECTION

- Syntax | 9541
- Hierarchy Level | 9543
- Description | 9544
- Required Privilege Level | 9544
- Release Information | 9544

Syntax

```
l2tp {
    access-line-information <connection-speed-update>;
    destination {
        address ip-address {
            access-line-information <connection-speed-update>;
            drain;
            routing-instance routing-instance-name {
                drain;
            }
        }
        lockout-timeout seconds;
        name destination-name {
            drain;
        }
    }
    destination-equal-load-balancing;
    destruct-timeout seconds;
    disable-calling-number-avp;
    disable-failover-protocol;
    drain;
    enable-ipv6-services-for-lac;
    enable-snmp-tunnel-statistics;
    failover-within-preference;
    ip-reassembly;
    maximum-sessions number;
```

```
    rx-connect-speed-when-equal;
    sessions-limit-group limit-group-name {
        maximum-sessions number;
    }
    traceoptions {
        debug-level level;
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        filter {
            protocol name;
            user user@domain;
            user-name username;
        }
        flag flag;
        interfaces interface-name {
            debug-level severity;
            flag flag;
        }
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
    tunnel {
        assignment-id-format (assignment-id | client-server-id);
        failover-resync (failover-protocol | silent-failover);
        idle-timeout seconds;
        maximum-sessions number;
        minimum-retransmission-timeout;
        name name {
            address ip-address {
                drain;
                routing-instance routing-instance-name {
                    drain;
                }
            }
            drain;
        }
        nas-port-method;
        retransmission-count-established count;
        retransmission-count-not-established count;
        rx-window-size packets;
         tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port | reject |
reject-ip-address | reject-udp-port);
    }
```

```
    tunnel-group group-name {
        aaa-access-profile profile-name;
        dynamic-profile profile-name;
        hello-interval seconds;
        hide-avps;
        l2tp-access-profile profile-name;
        local-gateway {
            address address;
            gateway-name gateway-name;
        }
        maximum-send-window packets;
        maximum-sessions number;
        ppp-access-profile profile-name;
        receive-window packets;
        retransmit-interval seconds;
        service-device-pool pool-name;
        service-interface interface-name;
        service-profile profile-name(parameter)&profile-name;
        syslog {
            host hostname {
                facility-override facility-name;
                log-prefix prefix-value;
                services severity-level;
            }
        }
        tos-reflect;
        tunnel-switch-profile profile-name;
        tunnel-timeout seconds;
    }
    tunnel-switch-profile profile-name;
    tx-connect-speed-method method;
    weighted-load-balancing;
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure L2TP services to establish PPP tunnels across a network.

> **NOTE**: L2TP is not supported over GRE tunnels.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

> **NOTE**: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for LAC on MX Series routers introduced in Junos OS Release 10.4.

Support for LNS on MX Series routers introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Layer 2 Tunneling Protocol Overview

L2TP for Subscriber Access Overview

l2tp (DDoS Protection)

IN THIS SECTION

Syntax | 9545

- Hierarchy Level | 9546
- Description | 9546
- Options | 9546
- Required Privilege Level | 9546
- Release Information | 9547

Syntax

```
l2tp {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
```

```
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
```

Hierarchy Level

```
[edit system ddos-protection protocols]
```

Description

Configure L2TP traffic.

Options

Table 54: L2TP Options

> aggregate	Configure aggregate for all L2TP control traffic
+ apply-groups	Groups from which to inherit configuration data
+ apply-groups-except	Don't inherit configuration data from these groups

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos Evolved OS Release 23.4R1.

l2tp-access-profile

IN THIS SECTION

- Syntax | 9547
- Hierarchy Level | 9547
- Description | 9547
- Options | 9548
- Required Privilege Level | 9548
- Release Information | 9548

Syntax

```
l2tp-access-profile profile-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Specify the profile used to validate all L2TP connection requests to the local gateway address.

Options

profile-name—Identifier for the L2TP connection profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> Configuring Access Profiles for L2TP Tunnel Groups
>
> *Configuring an L2TP Access Profile on the LNS*

l2tp (Group Profile)

IN THIS SECTION

- Syntax | 9548
- Hierarchy Level | 9549
- Description | 9549
- Required Privilege Level | 9549
- Release Information | 9549

Syntax

```
l2tp {
    interface-id interface-id;
    lcp-renegotiation;
```

```
    local-chap;
    maximum-sessions-per-tunnel number;
}
```

Hierarchy Level

```
[edit access group-profile profile-name]
```

Description

Configure the Layer 2 Tunneling Protocol for a group profile.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Layer 2 Tunneling Protocol (L2TP)

L2TP LNS Inline Service Interfaces

l2tp-interface-id

IN THIS SECTION

- Syntax | 9550
- Hierarchy Level | 9550
- Description | 9550
- Options | 9550
- Required Privilege Level | 9551
- Release Information | 9551

Syntax

```
l2tp-interface-id name;
(dedicated | shared);
```

Hierarchy Level

```
[edit interfaces sp-fpc/pic/port unit logical-unit-number interface],
[edit logical-systems logical-system-name interfaces sp-fpc/pic/port unit logical-unit-number
interface]
```

Description

Specify the L2TP options for configuring logical interfaces for group and user sessions.

Options

(`dedicated` | `shared`)—Specifies whether a logical interface can host one (dedicated) or multiple (shared) sessions at one time.

name—Interface identifier that must be replicated at the `[edit access profile name]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Services Interfaces Library for Routing Devices

l2tp-maximum-session (Service Interfaces)

IN THIS SECTION

- Syntax | 9551
- Hierarchy Level | 9552
- Description | 9552
- Options | 9552
- Required Privilege Level | 9552
- Release Information | 9552

Syntax

```
l2tp-maximum-session number;
```

Hierarchy Level

```
[edit interfaces si-slot/pic/port],
[edit interfaces asinumber]
```

Description

Specify the maximum number of L2TP sessions allowed on a physical service interface (si) or aggregated service interface (asi).

New session requests on an interface are accepted only when the session count is less than the maximum session limit. If the limit has been reached, subsequent requests are dropped and the LNS responds with a CDN message (Result Code 2, Error Code 4). When a pool of interfaces is configured, interfaces at the maximum limit are ignored in favor of an interface in the pool that has a lower session count. For an asi interface, the configuration applies only to the asi interface. You cannot configure a session limit on the individual member interfaces of an asi bundle.

Configuring the session limit to be less than the current number of sessions on the interface has no effect on existing sessions, but prevents any new sessions from being created until the number of session drops below the new limit.

Options

number Maximum number of L2TP sessions allowed for the interface. A value of 0 prevents the interface from being considered.

- **Default:** 64,000
- **Range:** 0 through 64,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

L2TP LNS Inline Service Interfaces

L2TP LAC Subscriber Configuration

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

L2TP for Subscriber Access Overview

l2tp (Profile)

IN THIS SECTION

- Syntax | 9553
- Hierarchy Level | 9554
- Description | 9554
- Options | 9554
- Required Privilege Level | 9557
- Release Information | 9557

Syntax

```
l2tp {
    interface-id interface-id;
    lcp-renegotiation;
    local-chap;
    maximum-sessions number;
    maximum-sessions-per-tunnel number;
    multilink {
        drop-timeout milliseconds;
        fragment-threshold bytes;
    }
    override-result-code session-out-of-resource;
    ppp-authentication (chap | pap);
```

```
    ppp-profile profile-name;
    sessions-limit-group;
    service-profile profile-name(parameter)&profile-name;
    shared-secret shared-secret;
}
```

Hierarchy Level

```
[edit access profile profile-name client client-name]
```

Description

Configure the L2TP properties for a profile.

> **NOTE**: Only the `interface-id`, `lcp-renegotiation`, `maximum-sessions`, `maximum-sessions-per-tunnel`, `sessions-limit-group` and `shared-secret` statements are supported for L2TP LNS on MX Series routers.

Options

interface-id

Configure the interface identifier.

- Values:
  - *interface-id*—Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the `[edit interfaces interface-name unit local-unit-number dial-options]` hierarchy level. For more information about the interface ID, see Services Interface Naming Overview.

lcp-renegotiation

Configure the L2TP network server (LNS) so it renegotiates the link control protocol (LCP) with the PPP client. When LCP renegotiation is disabled, LNS uses the pre-negotiated LCP parameters between the L2TP access concentrator (LAC) and PPP client to set up the session. When LCP renegotiation is enabled, authentication is also renegotiated.

> **NOTE**: This statement is not supported at the `[edit access group-profile l2tp]` hierarchy level for L2TP LNS on MX Series routers.

local-chap

Configure the Junos OS so that the LNS ignores proxy authentication attribute-value pairs (AVPs) from the L2TP access concentrator (LAC) and reauthenticates the PPP client using a Challenge Handshake Authentication Protocol (CHAP) challenge. When you do this, the LNS directly authenticates the PPP client.

> **NOTE**: This statement is not supported for L2TP LNS on MX Series routers.

maximum-sessions

Specify the maximum number of L2TP sessions for the chassis, all tunnels, a tunnel group, a session limit group, or a client.

- Values:
  - *number*—Number of sessions allowed.
  - Range: (Chassis, tunnel group, session limit group, or client) 1 through the default maximum chassis limit
  - Range: (Tunnel) 1 through 65,536

maximum-sessions-per-tunnel

Configure the maximum sessions for a Layer 2 tunnel.

> **NOTE**: This statement is not supported at the `[edit access group-profile l2tp]` hierarchy level for L2TP LNS on MX Series routers.

- Values:
  - *number*—Maximum number of sessions for a Layer 2 tunnel.

"multilink" on page 11103

Configure Multilink PPP for Layer 2 Tunneling Protocol (L2TP).

The options for this statement are explained separately. Click the linked statement for details.

override-result-code

Configure the LNS to override result codes in Call-Disconnect-Notify (CDN) messages.

- Values:
  - session-out-of-resource—Override result codes 4 and 5 with result code 2. These result codes indicate that the number of L2TP sessions have reached the configured maximum value and the LNS can support no more sessions. When the LAC receives the code, it fails over to another LNS to establish subsequent sessions. Some third-party LACs respond only to result code 2.

ppp-authentication

(T Series only) Configure PPP authentication.

> **NOTE**: This statement is not supported for L2TP LNS on MX Series routers.

- Values:
  - chap—Challenge Handshake Authentication Protocol.
  - pap—Password Authentication Protocol.

ppp-profile

(M Series, T Series only) Specify the profile used to validate PPP session requests through L2TP tunnels.

> **NOTE**: This statement is not supported for L2TP LNS on MX Series routers.

- **Values:** *profile-name*—Identifier for the PPP profile.

sessions-limit-group

(MX Series only) Starting in Junos OS Release 16.1, specify in an L2TP access profile the session limit group to which a client is assigned by the profile.

- **Values:** *limit-group-name*—Identifier of the session-limit group to which a client is assigned.

service-profile

Configure one or more dynamic service profiles to be applied to subscriber sessions at activation for all subscribers in the specified tunnel group or on the specified LAC. Services are typically applied to L2TP sessions with RADIUS VSAs or CoA requests. In multivendor environments, you might use only standard attributes to simplify management of multiple vendor VSAs. This statement enables you to apply services without using an external authority such as RADIUS. The locally configured list of services (service profiles) serves as local authorization that is applied by authd during client session activation. This list of services is subject to the same validation and processing as services originating from an external authority, such as RADIUS.

You can optionally specify parameters that are passed to the corresponding service when it is activated for the session. The parameter might override values configured in the profile itself, such as a downstream shaping rate for a CoS service. This enables you to use the same service profile for multiple situations with different requirements, or to modify a previously applied value for a service.

You can still use RADIUS VSAs or CoA requests together with the service profiles. If services are sourced from an external authority as authorization during

authentication or during subscriber session provisioning (activation), the services from the external authority take strict priority over those in the local configuration. If a service applied with RADIUS is the same as a service applied with a service profile in the CLI, but with different parameters, the RADIUS service is applied with a new session ID and takes precedence over the earlier service profile.

When service profiles are configured on a LAC client and on a tunnel group that uses that LAC client, the LAC configuration overrides the tunnel group configuration. Only the service profile configured on the LAC client is applied to subscribers in the tunnel group.

- Values:
  - *profile-name*—Name of a dynamic service profile that defines a service to be applied to L2TP subscriber sessions. You can specify one or more service profiles, separated by an ampersand (&).
  - *parameter*—(Optional) Value to be passed to the service when it is activated on the subscriber session.

shared-secret

Configure the shared secret.

- Values:
  - *shared-secret*—Shared secret key for authenticating the peer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Layer 2 Tunneling Protocol (L2TP)

L2TP LNS Inline Service Interfaces

L2TP for Subscriber Access Overview

l2tp-tunnel-session-identifier

IN THIS SECTION

- Syntax | 9558
- Hierarchy Level | 9558
- Description | 9558
- Required Privilege Level | 9559
- Release Information | 9559

Syntax

```
l2tp-tunnel-session-identifier
```

Hierarchy Level

```
[edit forwarding-options hash-key family inet
```

Description

For better distribution in load balancing for L2TP tunneled traffic, enable `l2tp-tunnel-session-identifier` to have the hash calculation of IPv4 packets include L2TP header parameters (tunnel ID and session ID). With this option enabled, Junos OS generates different hashes for packets from different tunnels. It will also generate different hashes for packets that belong to the same tunnel, but different sessions. L2TP control traffic is not load balanced based on tunnel ID and session ID.

When the `l2tp-tunnel-session-identifier` option is not enabled, the same hash value is computed for L2TP tunneled traffic using the outer IP header, regardless of tunnel (for tunnels created between the same end points).

To help diagnosis load balancing issues, you can run the `request pfe execute command "show jnh lb" target` *`target`* command from the Junos OS command line.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

hash-key (Forwarding Options) | 7986

Understanding Per-Packet Load Balancing

Configuring Per-Packet Load Balancing

l2vpn

IN THIS SECTION

- Syntax | 9559
- Hierarchy Level | 9562
- Description | 9562
- Required Privilege Level | 9562
- Release Information | 9562

Syntax

```
l2vpn {
    (control-word | no-control-word);
    encapsulation-type type;
    oam {
        bfd-liveness-detection {
            detection-time {
```

```
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                threshold milliseconds;
                minimum-interval milliseconds;
            }
            version (1 | automatic);
        }
        ping-interval seconds;
    }
    site site-name {
        community COMM;
        control-word ;
        encapsulation-type ethernet;
        ignore-encapsulation-mismatch;
        ignore-mtu-mismatch;
        interface interface-name {
            description text;
            community COMM;
            control-word ;
            encapsulation-type ethernet;
            ignore-encapsulation-mismatch;
            ignore-mtu-mismatch;
            mtu 1500;
            no-control-word;
            oam {
                bfd-liveness-detection {
                    detection-time {
                        threshold milliseconds;
                    }
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        threshold milliseconds;
                        minimum-interval milliseconds;
                    }
                    version (1 | automatic);
```

```
                }
                ping-interval seconds; seconds;
            }
            remote-site-id remote-site-id;
            target-attachment-identifier identifier;
        }
        mtu 1500;
        no-control-word;
        oam {
            bfd-liveness-detection {
                detection-time {
                    threshold milliseconds;
                }
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    threshold milliseconds;
                    minimum-interval milliseconds;
                }
                version (1 | automatic);
            }
            ping-interval seconds; seconds;
        }
        site-identifier identifier;
        site-preference preference-value {
            backup;
            primary;
        }
        source-attachment-identifier identifier;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable a Layer 2 VPN routing instance on a PE router or switch.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Layer 2 VPN Routing Instance

Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)

l2vpn-id

IN THIS SECTION

- Syntax | 9563
- Hierarchy Level | 9563
- Description | 9563
- Options | 9563

Required Privilege Level | 9563
Release Information | 9563

Syntax

```
l2vpn-id (as-number:id | ip-address:id);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name],
[edit routing-instances instance-name]
```

Description

Specify a globally unique Layer 2 VPN community identifier for the instance.

Options

as-number:*id*—Autonomous system number (`l2vpn-id:as-number:2-byte-number`. For example: `l2vpn-id l2vpn-id:100:200`. The AS number can be in the range from 1 through 65,535.

ip-address:*id*—IP address (`l2vpn-id:ip-address:2-byte-number`. For example: `l2vpn-id l2vpn-id:10.1.1.1:2`. The IP address can be any globally unique unicast address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4R2.

RELATED DOCUMENTATION

Example: Configuring BGP Autodiscovery for LDP VPLS

Example: Configuring BGP Autodiscovery for LDP VPLS with User-Defined Mesh Groups

Example: Configuring FEC 129 BGP Autodiscovery for VPWS

l2vpn (routing-options)

IN THIS SECTION

- Syntax | 9564
- Hierarchy Level | 9564
- Description | 9565
- Required Privilege Level | 9565
- Release Information | 9565

Syntax

```
l2vpn {
    l2ckt;
    l3vpn;
    labeled-bgp;
    no-l2ckt;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-
next-hop ingress],
[edit routing-options forwarding-table chained-composite-next-hop ingress]
```

Description

Enable composite chained next hop for ingress Layer 2 virtual private network (VPN) label-switched paths (LSPs).

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

l2vpn-use-bgp-rules-edit-protocols-bgp

IN THIS SECTION

- Syntax | 9565
- Hierarchy Level | 9566
- Release Information | 9566
- Description | 9566
- Default | 9566
- Required Privilege Level | 9566

Syntax

```
l2vpn-use-bgp-rules;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp path-selection],
[edit protocols bgp path-selection]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Enable routers to use both the BGP path selection algorithm and the designated forwarder path selection algorithm when selecting the preferred path to each destination in a Layer 2 VPN or VPLS network.When the BGP path selection algorithm is enabled, it is used by all of the Provider routers. The designated forwarder path selection algorithm is used by the PE router participating in the routing instance.

Default

By default, the designated forwarder path selection algorithm is used to select the best path to reach each destination within the Layer 2 VPN and VPLS network.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Enabling BGP Path Selection for Layer 2 VPNs and VPLS

route-distinguisher | **15856**

l3-interface-ingress-counting

IN THIS SECTION

- Syntax | 9567
- Hierarchy Level | 9567
- Description | 9567
- Default | 9567
- Required Privilege Level | 9568
- Release Information | 9568

Syntax

```
l3-interface-ingress-counting layer-3-interface-name;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Description

(EX8200 standalone switch and EX8200 Virtual Chassis) Enable routed VLAN interface (RVI) input counters on an EX8200 switch to collect RVI source statistics for tracking or billing purposes. The input counter is maintained by a firewall filter. The switch can maintain a limited number of firewall filter counters—these counters are allocated on a first-come, first-served basis.

Output (egress) counters for EX8200 switches are always present and cannot be removed.

Reset ingress-counting statistics with the "clear interfaces statistics" on page 3651 command.

Default

The input (ingress) counters (both packets and bytes) are disabled on an RVI by default.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

- show vlans | 27738
- clear interfaces statistics | 3651
- *Configuring Firewall Filters (CLI Procedure)*
- firewall | 7176
- Configuring Routed VLAN Interfaces on Switches (CLI Procedure)

l3-interface (VLAN)

IN THIS SECTION

- Syntax | 9569
- Hierarchy Level | 9569
- Description | 9569
- Default | 9569
- Options | 9569
- Required Privilege Level | 9570
- Release Information | 9570

Syntax

```
l3-interface (vlan logical-interface-number | irb logical-interface-number);
```

```
l3-interface l3-interface-name logical-interface-number {
    l3-interface-ingress-counting;
}
```

```
l3-interface interface-name-logical-unit-number;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

```
[edit interfaces ge-chassis/slot/port unit logical-unit-number                family ethernet-
switching]
[edit vlans vlan-name]
```

Description

Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed. Because traffic between VLANs must be routed, a common Layer 3 interface is required.

Default

No Layer 3 (routing) interface is associated with the VLAN.

Options

interface-name–logical-unit-number—Name of a logical interface.

vlan.*logical-interface-number*—Number of the logical interface. Use the `unit` number that you used when you created the `vlan` interface with a `set interfaces (QFX Series) vlan unit` statement.

> **NOTE**: Use this statement with versions of Junos OS that do not support Enhanced Layer 2 Software (ELS).

irb.*logical-interface-number*—Logical interface defined with a `set interfaces (QFX Series) irb` statement.

> **NOTE**: Use this statement with versions of Junos OS that support Enhanced Layer 2 Software (ELS).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

[edit vlans *vlan-name*] hierarchy level introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.

`irb` option introduced in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

- Configuring Routed VLAN Interfaces on Switches (CLI Procedure)
- *Configuring Integrated Routing and Bridging Interfaces on Switches (CLI Procedure)*
- show ethernet-switching interfaces | **19641**
- show ethernet-switching interface | **19636**
- show vlans | **27738**

l3vpn (ingress)

IN THIS SECTION

- Syntax | 9571
- Hierarchy Level | 9571
- Description | 9571
- Default | 9572
- Required Privilege Level | 9572
- Release Information | 9572

Syntax

```
l3vpn {
    extended-space;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-
next-hop ingress],
[edit routing-options forwarding-table chained-composite-next-hop ingress]
```

Description

(M120 routers, M320 routers with Enhanced III FPCs, and MX Series routers only) Accept larger numbers of Layer 3 VPN BGP updates with unique inner VPN labels (up to one million). The neighboring provider edge (PE) routers are typically from other vendors and are configured to assign a unique inner label to each Layer 3 VPN BGP route.

NOTE:

- On MX Series routers, this statement is not supported when the router has both DPCs and MPCs installed. You must remove one or the other type of card to successfully use this statement.
- On MX Series routers containing both DPC and MPC FPCs, the chassis must be configured with `enhanced-ip` option for enabling chained composite next hops. Only the "l3vpn" on page 9571 option under the `ingress` statement is supported on DPC and DPCE-based MX Series routers. All other statements and functionality under the `ingress` statement are not supported on MX routers with DPC linecards.
- In Junos OS Release 11.4 and earlier, the `l3vpn` statement was titled as `l3vpn-composite-nexthop`, and the statement was available at the `[edit logical-systems` *`logical-system-name`* `routing-options]` and `[edit routing-options]` hierarchy levels.
- On ACX Series routers, the `extended-space` statement is not supported.

NOTE: The `[edit logical-systems]` hierarchy level is not applicable in ACX Series routers.

The remaining statement is explained separately.

Default

This statement is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs

extended-space | 6703

l3vpn (transit)

IN THIS SECTION

- Syntax | 9573
- Hierarchy Level | 9573
- Description | 9573
- Required Privilege Level | 9574
- Release Information | 9574

Syntax

```
transit
    l3vpn;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-next-hop transit]
[edit routing-options forwarding-table chained-composite-next-hop transit]
```

Description

Configure Layer 3 VPN settings on the transit router.

This statement and the associated functionality is available on PTX Packet Transport Routers and QFX10000 switches.

NOTE: You can configure `transit l3vpn` only after configuring the `enhanced-ip` operation mode. See "network-services" on page 11320 to configure the `enhanced-ip` mode.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

chained-composite-next-hop | 3277

label

IN THIS SECTION

- Syntax | 9574
- Hierarchy Level | 9575
- Description | 9575
- Options | 9575
- Required Privilege Level | 9575
- Release Information | 9575

Syntax

```
label {
    allocation label-allocation-policy;
    substitution label-substitution-policy;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options]
[edit routing-instances routing-instance-name routing-options]
```

Description

Specify label allocation and substitution policies on a per-route basis.

Options

`allocation` *`label-allocation-policy`*—Specify a policy to generate labels on a per-route basis.

`substitution` *`label-substituion-policy`*—Specify a policy to substitute labels on a per-route basis. The label substitution policy is used to determine whether or not a label should be substituted on an AS border router. The results of the policy operation are either accept (label substitution is performed) or reject (label substitution is not performed).

- **Default:** accept

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

> *Configuring a Label Allocation and Substitution Policy for VPNs*

label-allocation

IN THIS SECTION

- Syntax | 9576
- Hierarchy Level | 9576
- Description | 9576
- Options | 9576
- Required Privilege Level | 9576
- Release Information | 9577

Syntax

```
label-allocation per-instance | per-bridge-domain;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols evpn]
```

Description

Specifies the MPLS label allocation setting for the EVPN routing instance.

Options

`per-instance`—Allocates a single MPLS label for the EVPN routing instance.

`per-bridge-domain`—Allocates a single MPLS label for each bridge domain. By default, per-bridge-domain is enabled on the PTX10008 Packet Transport Routers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

label-block-size

IN THIS SECTION

- Syntax | 9577
- Hierarchy Level | 9577
- Description | 9578
- Default | 9578
- Options | 9578
- Required Privilege Level | 9578
- Release Information | 9578

Syntax

```
label-block-size size;
```

Hierarchy Level

```
[edit logical-systems profile-name routing-instances instance-name protocols vpls],
[edit routing-instances instance-name protocols vpls]
```

Description

Configure the label block size for VPLS labels.

Default

8

Options

- 2—Allocate the label blocks in increments of 2. For a VPLS domain that has only two sites with no future expansion plans.
- 4—Allocate the label blocks in increments of 4.
- 8 (default)—Allocate the label blocks in increments of 8.
- 16—Allocate the label blocks in increments of 16. A label block size of 16 enables you to minimize the number of routes in the VPLS domain. Use this setting only if the number of routes is the most important concern.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring the Label Block Size for VPLS

label-position

IN THIS SECTION

- Syntax | 9579
- Hierarchy Level | 9579
- Description | 9579
- Default | 9579
- Options | 9579
- Required Privilege Level | 9580
- Release Information | 9580

Syntax

```
label-position [ positions ];
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name mpls-ipv4-template],
[edit services flow-monitoring version9 template template-name mpls-template]
```

Description

Specify positions for up to three labels in the active flow monitoring version 9 template.

Default

[1 2 3]

Options

positions—Numbered positions for the labels.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

> *Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates*

label-switched-path

IN THIS SECTION

- Syntax | 9580
- Hierarchy Level | 9581
- Description | 9581
- Options | 9581
- Required Privilege Level | 9581
- Release Information | 9581

Syntax

```
label-switched-path lsp-name to remote-provider-edge-switch;
```

Hierarchy Level

```
[edit protocols mpls]
```

Description

Define a label-switched path (LSP) to the remote provider edge switch to use for MPLS traffic. You must specify this statement on the provider edge switch.

Options

lsp-name —Name that identifies the LSP. The name can be up to 32 characters and can contain letters, digits, periods, and hyphens. To include other characters, enclose the name in quotation marks. The name must be unique on the ingress switch.

remote-provider-edge-switch —Either the loopback address or the switch address.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Example: Configuring MPLS on EX8200 and EX4500 Switches

Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS

Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect

label-switched-path (Protocols IS-IS)

IN THIS SECTION

- Syntax | 9582
- Hierarchy Level | 9582
- Description | 9582
- Options | 9583
- Required Privilege Level | 9583
- Release Information | 9583

Syntax

```
label-switched-path name level level-number metric metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Advertise label-switched paths (LSPs) into IS-IS as point-to-point links. The LSP is advertised in the appropriate IS-IS levels as a point-to-point link and contains a local address and a remote address.

When you advertise LSPs into IS-IS as point-to-point links, the LSPs are used in SPF calculations. The advertisement contains a local address (the **from** address of the LSP), a remote address (the **to** address of the LSP), and a metric.

Before a single-hop LSP between a multiaccess link can be announced as up and used in SPF calculations, you must configure an LSP in both directions between two label-switched routers.

Options

name—Identifies the LSP.

level-number—IS-IS level number.

- **Values:** 1 or 2

metric—Metric value.

- **Range:** 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)
- **Default:** 0 (for lo0), 10 (for all other interfaces)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Advertising Label-Switched Paths into IS-IS

Understanding Forwarding Adjacencies

label-switched-path (Protocols Link Management)

IN THIS SECTION

- Syntax | 9584
- Hierarchy Level | 9584
- Description | 9584
- Options | 9584
- Required Privilege Level | 9584

Release Information | 9584

Syntax

```
label-switched-path lsp-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],
[edit protocols link-management te-link te-link-name]
```

Description

Specify the label-switched path (LSP) to be used by the forwarding adjacency.

Options

lsp-name—Name of the LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Forwarding Adjacencies

label-switched-path (Protocols MPLS)

IN THIS SECTION

- Syntax | 9585
- Hierarchy Level | 9588
- Description | 9588
- Options | 9588
- Required Privilege Level | 9589
- Release Information | 9590

Syntax

```
label-switched-path lsp-name {
    disable;
    adaptive;
    admin-down;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    auto-bandwidth {
        adjust-interval seconds;
        adjust-threshold percentage;
        maximum-bandwidth bps;
        minimum-bandwidth bps;
        monitor-bandwidth;
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
    class-of-service cos-value;
    cross-credibility-cspf;
```

```
    description text;
    entropy-label;
    fast-reroute {
        (bandwidth bps | bandwidth-percent percentage);
        (exclude [ group-names ] | no-exclude);
        hop-limit number;
        (include-all [ group-names ] | no-include-all);
        (include-any [ group-names ] | no-include-any);
    }
    from address;
    install {
        destination-prefix/prefix-length <active>;
    }
    inter-domain;
    ldp-tunneling;
    link-protection;
    lsp-attributes {
        lsp-external-controller;
        encoding-type (ethernet | packet | pdh | sonet-sdh);
        gpid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 | pos-
scrambling-crc-32 | pos-no-scrambling-crc-32 | ppp);
        signal-bandwidth type;
        switching-type (fiber | lambda | psc-1 | tdm);
    }
    metric metric;
    no-cspf;
    no-decrement-ttl;
    node-link-protection;
    optimize-timer seconds;
    p2mp lsp-name;
    policing {
        filter filter-name;
        no-auto-policing;
    }
    preference preference;
    primary path-name {
        adaptive;
        admin-group {
            exclude [ group-names ];
            include-all [ group-names ];
            include-any [ group-names ];
        }
        bandwidth bps {
```

```
            ct0 bps;
            ct1 bps;
            ct2 bps;
            ct3 bps;
        }
        class-of-service cos-value;
        hop-limit number;
        no-cspf;
        no-decrement-ttl;
        optimize-timer seconds;
        preference preference;
        priority setup-priority reservation-priority;
        (record | no-record);
        select (manual | unconditional);
        standby;
    }
    priority setup-priority reservation-priority;
    (random | least-fill | most-fill);
    (record | no-record);
    retry-limit number;
    retry-timer seconds;
    revert-timer seconds;
    secondary path-name {
        adaptive;
        admin-group {
            exclude[ group-names ];
            include-all [ group-names ];
            include-any [ group-names ];
        }
        bandwidth bps {
            ct0 bps;
            ct1 bps;
            ct2 bps;
            ct3 bps;
        }
        class-of-service cos-value;
        hop-limit number;
        no-cspf;
        no-decrement-ttl;
        optimize-timer seconds;
        preference preference;
        priority setup-priority reservation-priority;
        (record | no-record);
```

```
        select (manual | unconditional);
        standby;
    }
    soft-preemption;
    standby;
    to address;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Configure an LSP to use in dynamic MPLS. When configuring an LSP, you must specify the address of the egress router in the `to` statement. All remaining statements are optional.

Options

`lsp-name` Name that identifies the LSP. The name can be up to 64 characters and can contain letters, digits, periods, and hyphens. To include other characters, enclose the name in quotation marks. The name must be unique within the ingress router.

`cross-credibility-cspf` Enable path computation across credibility levels. The constraint path computation is run across multi-protocol links and nodes, instead of a credibility-by-credibility basis.

`link-protection` Enable protection for LSP from link faults only.

`no-self-ping` Disable self-ping for the LSP.

> **NOTE**: Starting in Junos OS 17.4, you can override the behavior of self-ping running by default using `no-self-ping` statement.

> **NOTE**: Starting with Junos OS 16.1 the MPLS EXP bits transmitted in self ping messages are set based on the DSCP/ToS setting of the corresponding IP packet. In prior Junos OS releases the MPLS EXP bits in self ping packets are set to 0.

`self-ping-duration` *seconds*

Specify the duration (in seconds) for which to run the self-ping mechanism unless the ping succeeds sooner.

LSP self-ping for an LSP starts at the ingress label edge router (LER), once a Resv message for that LSP has been received. The configured self-ping time indicates the duration for which the self-ping mechanism runs once the LSP instance is signaled to be `UP`.

By default, self-ping is enabled. The LSP types like CCC, P2MP, VLAN-based , and non-default instances do not support self-ping .

The self-ping mechanism runs until the self-ping probe is received back (at which point the traffic is immediately switched to it) , or until the configured self-ping duration for the LSP is over (at which point traffic is switched over).

When LSP self-ping-duration is enabled, the LSP behavior reverts back to a timer-based mechanism similar to the `optimize-switchover-delay`, where a specific amount of time is provided for all the downstream nodes to install the forwarding path before switchover. When LSP self-ping-duration is not enabled, the default behavior is to wait for an infinite amount of time for the self-ping to succeed before switching the traffic.

- **Range:** 1 through 65535 seconds

> **NOTE**: Starting in Junos OS 17.4R1, the default time-out duration for which the self-ping runs on an LSP instance is reduced from 65535 (runs until success) to 1800 seconds. You can also configure the self ping duration value between 1 to 65535 (runs until success) seconds.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`cross-credibility-cspf` option introduced in Junos OS Release 14.2.

`self-ping-duration` and `no-self-ping` options introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring the Ingress and Egress Router Addresses for LSPs

Configuring Primary and Secondary LSPs

label-switched-path (Protocols MVPN)

IN THIS SECTION

- Syntax | **9590**
- Hierarchy Level | **9591**
- Description | **9591**
- Required Privilege Level | **9591**
- Release Information | **9591**

Syntax

```
label-switched-path {
    label-switched-path-template (Multicast){
        (default-template | lsp-template-name);
    }
}
```

Hierarchy Level

```
[edit protocols mvpn inter-region-template template template-name region region-name ingress-
replication]
[edit protocols mvpn inter-region-template template template-name all-regions ingress-
replication]
```

Description

Specify point-to-point LSP unicast tunnel at the ABRs. Optionally, you can specify a label-switched path template. If you configure `ingress-replication label-switched-path` and do not include `label-switched-path-template`, ingress replication works with existing LDP or RSVP tunnels. If you include label-switched-path-template, then the RSVP tunnels should be enabled across the ABR regions.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview

all-regions | **2157**

inter-region | **8789**

inter-region-template | **8793**

inter-region-segmented | **8791**

region | **14262**

template(Protocols MVPN) | **29092**

Example: Configuring Segmented Inter-Area P2MP LSP

Configuring Segmented Inter-Area P2MP LSP

label-switched-path (Protocols OSPF)

IN THIS SECTION

- Syntax | 9592
- Hierarchy Level | 9592
- Description | 9592
- Options | 9592
- Required Privilege Level | 9593
- Release Information | 9593

Syntax

```
label-switched-path name metric metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area area-id],
[edit protocols ospf area area-id],
[edit routing-instances routing-instance-name protocols ospf area area-id]
```

Description

Advertise label-switched paths into OSPF as point-to-point links.

The label-switched path is advertised in the appropriate OSPF levels as a point-to-point link and contains a local address and a remote address.

Options

name—Name of the label-switched path.

metric—Metric value.

- **Range:** 1 through 65,535
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Advertising Label-Switched Paths into OSPFv2

label-switched-path-template

IN THIS SECTION

- Syntax | 9594
- Hierarchy Level | 9594
- Description | 9594
- Options | 9594
- Required Privilege Level | 9594
- Release Information | 9594

Syntax

```
label-switched-path-template {
    (default-template | lsp-template-name);
}
```

Hierarchy Level

```
[edit protocols mpls lsp-external-controller lsp-external-controller]
```

Description

Specify the LSP template with parameters for setting up the PCE-initiated LSPs when the PCE initiating the LSP does not provide the PCE-initiated parameters. When `label-switched-path-template` is not configured, the default LSP parameters are used.

Options

`default-template` Specify that the default LSP template be used for the dynamically generated PCE-initiated LSPs.

`lsp-template-name` Specify the name of the LSP template to be used for setting up PCE-initated LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

pcep | 12320

label-switched-path-template (Container LSP)

IN THIS SECTION

- Syntax | 9595
- Hierarchy Level | 9595
- Description | 9595
- Options | 9595
- Required Privilege Level | 9596
- Release Information | 9596

Syntax

```
label-switched-path-template {
    (default-template | lsp-template-name);
}
```

Hierarchy Level

```
[edit protocols mpls container-label-switched-path lsp-name]
[edit routing-instances instance-name provider-tunnel]
```

Description

Specify the LSP template. An LSP template is used as the basis for other dynamically generated LSPs.

Options

default-template Specify that the default LSP template be used for the dynamically generated LSPs.

lsp-template-name Specify the name of an LSP to be used as a template for the dynamically generated LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.

RELATED DOCUMENTATION

container-label-switched-path | 4577

label-switched-path-template (Multicast)

IN THIS SECTION

- Syntax | 9596
- Hierarchy Level | 9597
- Description | 9597
- Options | 9597
- Required Privilege Level | 9598
- Release Information | 9598

Syntax

```
label-switched-path-template {
    (default-template | lsp-template-name);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel rsvp-te],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel ingress-replication label-switched-path],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group address source source-address rsvp-te],
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name rsvp-te
entry-name],
[edit protocols mvpn inter-region-segmented template template-name region region-name ingress-
replication label-switched-path],
[edit protocols mvpn inter-region-segmented template template-name region region-name rsvpe-te],
[edit protocols mvpn inter-region-template template template-name all-regions ingress-
replication label-switched-path],
[edit protocols mvpn inter-region-template template template-name all-regions rsvp-te],
[edit routing-instances routing-instance-name provider-tunnel ingress-replication label-switched-
path],
[edit routing-instances routing-instance-name provider-tunnel rsvp-te],
[edit routing-instances routing-instance-name provider-tunnel selective group address source
source-address rsvp-te],
[edit routing-options dynamic-tunnels tunnel-name rsvp-te entry-name]
[edit routing-instances instance-name provider-tunnel]
```

Description

Specify the LSP template. An LSP template is used as the basis for other dynamically generated LSPs. This feature can be used for a number of applications, including point-to-multipoint LSPs, flooding VPLS traffic, configuring ingress replication for IP multicast using MBGP MVPNs, and to enable RSVP automatic mesh. There is no default setting for the `label-switched-path-template` statement, so you must configure either the default-template using the `default-template` option, or you must specify the name of your preconfigured LSP template.

Options

`default-template`—Specify that the default LSP template be used for the dynamically generated LSPs.

`lsp-template-name`—Specify the name of an LSP to be used as a template for the dynamically generated LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

Configuring Dynamic Point-to-Multipoint Flooding LSPs

Configuring RSVP Automatic Mesh

label-withdrawal-delay

IN THIS SECTION

- Syntax | 9598
- Hierarchy Level | 9599
- Description | 9599
- Options | 9599
- Required Privilege Level | 9599
- Release Information | 9599

Syntax

```
label-withdrawal-delay seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Delay the withdrawal of labels to reduce router workload during IGP convergence.

Options

seconds—Configure the number of seconds to wait before withdrawing labels for the LDP LSPs.

- **Default:** 60 seconds
- **Range:** 0 through 300 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Configuring the Label Withdrawal Timer

labeled-bgp

IN THIS SECTION

- Syntax | 9600
- Hierarchy Level | 9600
- Description | 9600
- Required Privilege Level | 9600
- Release Information | 9601

Syntax

```
labeled-bgp {
    inet6;

inet;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-next-hop ingress],
[edit routing-options forwarding-table chained-composite-next-hop ingress]
```

Description

Allows you to configure chained composite next hops for IPv4 and IPv6 labeled-unicast routes.

The other statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

This statement was introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs

labeled-unicast

IN THIS SECTION

- Syntax | 9601
- Hierarchy Level | 9602
- Description | 9602
- Options | 9602
- Required Privilege Level | 9602
- Release Information | 9602

Syntax

```
labeled-unicast {
    aggregate-label {
        community community-name:
    }
    explicit-null {
        connected-only;
    }
    per-group-label;
    resolve-vpn;
    traffic-statistics {
        file filename <world-readable | no-world-readable>;
        interval seconds;
```

```
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family inet],
[edit logical-systems logical-system-name protocols bgp group group-name family inet],
[edit protocols bgp family inet],
[edit protocols bgp group group-name family inet]
```

Description

Advertise labeled routes from the inet.0 VPN, and place labeled routes into the inet.0 VPN. When the `labeled-unicast` statement is used, the local router automatically performs a next hop to self on all routes advertised into EBGP from IBGP and into IBGP from EBGP.

Options

`resolve-vpn`—(Optional) Store labeled routes in the inet.3 or inet6.3 routing table to resolve routes for a provider edge (PE) router located in a different autonomous system (AS). For a PE router to install a route in the virtual routing and forwarding (*VRF*) table, the next hop must resolve to a route stored in the inet.3 or inet6.3 routing table. This option is also used to configure inter-AS VPLS with MAC operations.

NOTE: VPLS is not supported by QFX Series switches.

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Carrier-of-Carriers VPNs for Customers That Provide Internet Service

Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service

Configuring Inter-AS VPLS with MAC Processing at the ASBR

labeled-unicast (Protocols BGP)

IN THIS SECTION

- Syntax | 9603
- Hierarchy Level | 9604
- Description | 9605
- Required Privilege Level | 9605
- Release Information | 9605

Syntax

```
labeled-unicast {
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
        drop-excess <percentage>;
        hide-excess <percentage>;
      }
    aggregate-label {
        community community-name;
    }
    entropy-label {
        elc-v2-compatible;
        import policy-name;
        no-next-hop-validation;
    }
    explicit-null {
        connected-only;
```

```
    }
    nexthop-resolution {
        preserve-nexthop-hierarchy;
    }
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
        drop-excess <percentage>;
        hide-excess <percentage>;
    }
    protection;
    resolve-vpn;
    rib (inet.3 | inet6.3);
    rib-group group-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family (inet | inet6)],
[edit logical-systems logical-system-name protocols bgp group group-name family (inet | inet6)],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
(inet | inet6)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family (inet | inet6)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family (inet | inet6)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family (inet | inet6)],
[edit protocols bgp family (inet | inet6)],
[edit protocols bgp family prefix-limit (inet | inet6)(any | flow | labeled-unicast | multicast |
unicast)],
[edit protocols bgp group group-name family (inet | inet6)],
[edit protocols bgp group group-name neighbor address family (inet | inet6)],
[edit routing-instances routing-instance-name protocols bgp family (inet | inet6)],
[edit routing-instances routing-instance-name protocols bgp group group-name family (inet |
inet6)],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (inet | inet6)]
```

Description

Configure the family type to be labeled-unicast.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IPv6 BGP Routes over IPv4 Transport

Enabling Layer 2 VPN and VPLS Signaling

Understanding Multiprotocol BGP

lacp

IN THIS SECTION

- Syntax | 9606
- Hierarchy Level | 9606
- Description | 9606
- Options | 9606
- Required Privilege Level | 9606
- Release Information | 9606

Syntax

```
lacp {
    link-protection {
        non-revertive;
    }
    system-priority priority;
}
```

Hierarchy Level

```
[edit chassis aggregated-devices ethernet]
```

Description

For aggregated Ethernet interfaces only, configure Link Aggregation Control Protocol (LACP) parameters at the global level for use by LACP at the interface level.

Options

The statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Junos OS for Supporting Aggregated Devices

lacp (802.3ad)

IN THIS SECTION

- Syntax | 9607
- Hierarchy Level (EX Series) | 9607
- Hierarchy Level (QFX Series) | 9607
- Description | 9608
- Options | 9608
- Required Privilege Level | 9608
- Release Information | 9608

Syntax

```
lacp {
    force-up;
    (primary | backup);

}
```

Hierarchy Level (EX Series)

```
[edit interfaces interface-name ether-options 802.3ad]
[edit interfaces aeX aggregated-ether-options]
[edit chassis aggregated-devices ethernet]
```

Hierarchy Level (QFX Series)

```
[edit interfaces interface-name ether-options 802.3ad]
```

Description

Configure the Link Aggregation Control Protocol (LACP) parameters for aggregated Ethernet interfaces on the global level (for all the aggregated Ethernet interfaces on the switch) or for a specific aggregated Ethernet interface.

Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.

NOTE: The force-up statement is not supported on QFX10002 switches.

Options

- primary—Designate the interface as the primary interface for link-protection mode.
- backup—Designate the interface as the backup interface for link-protection mode.

The remaining statements are described separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support for LACP link protection introduced in Junos OS Release 11.4 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

Configuring Aggregated Ethernet Links (CLI Procedure)

Configuring Aggregated Ethernet LACP (CLI Procedure)

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

Junos OS Ethernet Interfaces Configuration Guide

Configuring Link Aggregation

Configuring Aggregated Ethernet LACP (CLI Procedure)

Understanding Aggregated Ethernet Interfaces and LACP for Switches

lacp (802.3ad) (MX Series)

IN THIS SECTION

- Syntax | 9609
- Hierarchy Level | 9609
- Description | 9609
- Options | 9610
- Required Privilege Level | 9610
- Release Information | 9610

Syntax

```
lacp {
    port-priority port-priority;
    }
```

Hierarchy Level

```
[edit interfaces interface-name fastether-options 802.3ad],
[edit interfaces interface-name gigether-options 802.3ad]
```

Description

Configure the Link Aggregation Control Protocol (LACP) port priority for Ethernet interfaces.

Options

port-priority—Priority for being elected as the active port to collect and distribute traffic. A smaller value indicates a higher priority for selection.

- **Range:** 0 through 65,535
- **Default:** 127

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring LACP for Aggregated Ethernet Interfaces

port-priority | **12856**

lacp (Aggregated Ethernet)

IN THIS SECTION

- Syntax (NFX Series) | **9611**
- Syntax (EX Series) | **9611**
- Hierarchy Level (EX Series) | **9612**
- Hierarchy Level (NFX Series) | **9612**
- Description | **9612**
- Default | **9612**
- Options | **9613**
- Required Privilege Level | **9613**

Release Information | 9613

Syntax (NFX Series)

```
lacp (active | passive) {
    admin-key key;
    fast-failover;
    link-protection {
        disable;
        (revertive | non-revertive);
    }
    periodic interval
    system-ID mac-address;
    system-priority priority;
    force-up;
}
```

Syntax (EX Series)

```
lacp {
    (active | passive);
    admin-key key;
    accept-data;
    fast-failover;

        link-protection {
        disable;
        (revertive |non-revertive);
    }
    periodic interval;
    system-id mac-address;
    system-priority priority;
}
```

Hierarchy Level (EX Series)

```
[edit interfaces aex aggregated-ether-options]
[edit logical-systems logical-system-name interfaces aeX aggregated-ether-options]
```

Hierarchy Level (NFX Series)

```
[edit interfaces interface-name aggregated-ether-options]
```

Description

Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.

For EX Series, when you configure the `accept-data` statement at the `[edit interfaces aeX aggregated-ether-options lacp]` hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.

NOTE: When you configure the `accept-data` statement at the `[edit interfaces aeX aggregated-ether-options lacp]` hierarchy level, this behavior occurs:

- By default, the `accept-data` statement is not configured when LACP is enabled.
- You can configure the `accept-data` statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.
- When LACP is down and a member link receives packets, the router or switch does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the `accept-data` statement is configured.

NOTE: The force-up statement is not supported on QFX10002 switches.

Default

If you do not specify LACP as either `active` or `passive`, LACP remains passive.

Options

active Initiate transmission of LACP packets.

admin-key *number* Specify an administrative key for the router or switch.

> **NOTE**: You must also configure multichassis link aggregation (MC-LAG) when you configure the `admin-key`.

fast-failover Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.

passive Respond to LACP packets.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- *Configuring Link Aggregation*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*
- Configuring LACP for Aggregated Ethernet Interfaces

lacp (Interfaces)

IN THIS SECTION

- Syntax | 9614
- Hierarchy Level | 9614
- Description | 9614
- Options | 9614
- Required Privilege Level | 9615
- Release Information | 9615

Syntax

```
lacp {
    (active | passive);
    periodic (fast | slow);
}
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options]
```

Description

For redundant Ethernet interfaces in a chassis cluster only, configure Link Aggregation Control Protocol (LACP).

Options

- `active`—Initiate transmission of LACP packets.
- `passive`—Respond to LACP packets.
- `Default`—If you do not specify lacp as either active or passive, LACP remains off (the default).

- `periodic`—Interval for periodic transmission of LACP packets. The options are:
  - `fast`—Transmit link aggregation control PDUs every second.
  - `slow`—Transmit link aggregation control PDUs every 30 seconds.
  - `Default`—If you do not specify `periodic` as either `fast` or `fast`, it is set to `fast`. If `lacp` is set to `fast`, LACP is asking the link partner to send an LACP heartbeat every 1-second. If `lacp` is set to `slow`, LACP is asking the link partner to send an LACP heartbeat every 30-seconds.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding LACP on Standalone Devices

periodic (Interfaces)

lacp (Protocols)

IN THIS SECTION

- Syntax | 9616
- Hierarchy Level | 9616
- Description | 9616
- Default | 9616
- Options | 9616
- Required Privilege Level | 9617

Release Information | 9617

Syntax

```
lacp {
    traceoptions {
        file <filename> <files number> <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
    fast-hello-issu;
    ppm (Ethernet Switching) centralized | distributed | inline;
}
```

Hierarchy Level

```
[edit protocols]
```

Description

On MX and T Series routers, you can specify periodic packet management (PPM) as centralized. By default, the PPM is distributed.

MX Series routers support Link Aggregation Control Protocol (LACP) with fast hellos during unified ISSU. This support is disabled by default. You must enable the `fast-hello-issu` option on the main router and on the peer routers before starting unified ISSU. Note that the peer router must also be an MX Series router for this functionality to work.

Default

Distributed PPM processing is enabled for all packets that use PPM.

Options

`ppm`—Specify the type of PPM processing. Possible values: Centralized, distributed, and inline.

`fast-hello-issu`—Enable LACP with fast hellos during unified ISSU.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

The `ppm centralized` option introduced in Junos OS Release 9.4.

The `fast-hello-issu` option introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

- Configuring LACP for Aggregated Ethernet Interfaces
- Unified ISSU System Requirements

land

IN THIS SECTION

- Syntax | 9618
- Hierarchy Level | 9618
- Description | 9618
- Required Privilege Level | 9618
- Release Information | 9618

Syntax

```
land;
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp]
```

Description

Enable prevention of Land attacks by combining the SYN flood defense with IP spoofing protection. Land attacks occur when an attacker sends spoofed IP packets with headers containing the target's IP address for the source and destination IP addresses. The attacker sends these packets with the SYN flag set to any available port. The packets induce the target to create empty sessions with itself, filling its session table and overwhelming its resources.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

land-attack-check

IN THIS SECTION

- Syntax | 9619
- Hierarchy Level | 9619
- Description | 9619
- Options | 9619
- Required Privilege Level | 9620
- Release Information | 9620

Syntax

```
land-attack-check (ip-only | ip-port );
```

Hierarchy Level

```
[edit services service-set service-set-name nat-options]
```

Description

Enable land attack checks based on either IP address only or both IP address and IP port number.

NOTE: If you do not configure this statement, there is no land attack check for hairpinning NAT packets.

Options

ip-only—Land attack check is based on IP address only.

ip-port—Land attack check is based on IP address and IP port number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced with Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring Service Rules

max-sessions-per-subscriber | 10474

land-attack-check (IDS MS-MPC)

IN THIS SECTION

- Syntax | 9620
- Hierarchy Level | 9621
- Description | 9621
- Options | 9621
- Required Privilege Level | 9621
- Release Information | 9621

Syntax

```
land-attack-check (ip-only | ip-port);
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Identify and drop SYN packets that have the same source and destination address or port, which provides protection against land attacks. This statement can only be used in IDS rules assigned to a service set on an MS-MPC.

Options

ip-only Identify and drop SYN packets that have the same source and destination address.

ip-port Identify and drop SYN packets that have the same source and destination address and port.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

land (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 9622
- Hierarchy Level | 9622
- Description | 9622
- Required Privilege Level | 9622
- Release Information | 9622

Syntax

```
land;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Description

Identify and drop SYN packets that have the same source and destination address or port, which protects against land attacks. In a land attack, the target using up its resources as it repeatedly replies to itself.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

lane

IN THIS SECTION

- Syntax | 9623
- Hierarchy Level | 9624
- Description | 9624
- Options | 9624
- Required Privilege Level | 9626
- Release Information | 9626

Syntax

```
lane [lane-number] {
        rx-cdr (off | on);
        rx-output (off | on);
        rx-output-amplitude (1200mV | 400mV | 600mV | 800mV);
        rx-output-emphasis rx-output-e;
        rx-rate-select rx-rate-select;
        rx-squelch (off | on);
        tx-cdr (off| no);
        tx-input-equalization-adaptive (off | on);
        tx-input-equalization-manual tx-input-manual;
        tx-laser (off | on);
        tx-rate-select tx-rate-select;
        tx-squelch (off | on);
        }
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Description

Configure the optical module parameters of QSFP (Quad Small Form-Factor Pluggable) modules for a specific media lane number. The following QSFP modules are supported:

- QSFP+ : 4 host lanes, supports 10g per lane, total bandwidth of 40g
- QSFP28 : 4 host lanes, supports 25g per lane, total bandwidth of 100g
- QSFP56 : 4 host lanes, supports 50g per lane, total bandwidth of 200g
- QSFP28-DD : 8 host lanes, supports 25g per lane, total bandwidth of 200g
- QSFP56-DD : 8 host lanes, supports 50g per lane, total bandwidth of 400g

The following specifications standardize how the host device interacts with the QSFP optical modules:

- The 4-lane QSFP modules adhere to SFF-8636 management interface specification.
- The 8-lane QSFP-DD modules adhere to CMIS (Common Management Interface Specification).

You can configure optical module parameters for QSFP28 and QSFP28-DD optical modules for Juniper qualified optics as well as third-party optics from Junos OS. You cannot configure optical module parameters for QSFP copper DAC (direct attach cables) or QSFP load modules.

Options

lane-number Specifies the media lane number.

- Range: 0 through 7

rx-cdr Controls receive clock and data recovery block of the optics.

- `Off`—Receive clock and data recovery is disabled.
- `On`—Receive clock and data recovery is enabled.

rx-output Controls receive output from the optics towards the host device.

- `Off`—Receive output for optics is disabled.

- `On`—Receive output for optics is enabled.

rx-output-amplitude Controls the amplitude of receive output for which equalization is to be applied.

- `1200mV`—No equalization for 600mV - 1200mV amplitude.
- `400mV`—No equalization for 100mV - 400mV amplitude.
- `600mV`—No equalization for 300mV - 600mV amplitude.
- `800mV`—No equalization for 400mV - 800mV amplitude.

rx-output-emphasis Equalization value (in dB) for receive output signal of the optics.

- Range: 0 through 10

rx-rate-select Controls baud rate for receive output signal.

- Range: 0 through 3

rx-squelch Squelch the receive output signal of optics.

- `Off`—Receive output signal for the optics is disabled.
- `On`—Receive output signal for the optics is enabled.

tx-cdr Controls transmit clock and data recovery block of the optics.

- `Off`—Transmit clock and data recovery is disabled.
- `On`—Transmit clock and data recovery is enabled.

tx-input-equalization-adaptive Controls adaptive equalization for optics transmit input.

- `Off`—Adaptive equalization is disabled.
- `On`—Adaptive equalization is enabled.

tx-input-equalization-manual Equalization value (in dB) for optics transmit input.

- Range: 0 through 15

tx-laser Controls transmit laser block of the optics.

- `Off`—Transmit laser is disabled.
- `On`—Transmit laser is enabled.

tx-rate-select Controls baud rate for transmit input signal.

- Range: 0 through 3

tx-squelch Squelch the transmit input signal of the optics.

- `Off`—Transmit input signal for the optics is disabled.
- `On`—Transmit input signal for the optics is enabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

optics-options | 11847

lane-all | 9626

lane-all

IN THIS SECTION

- Syntax | 9627
- Hierarchy Level | 9627
- Description | 9627
- Options | 9628
- Required Privilege Level | 9629
- Release Information | 9629

Syntax

```
lane-all  {
        rx-cdr (off | on);
        rx-output (off | on);
        rx-output-amplitude (1200mV | 400mV | 600mV | 800mV);
        rx-output-emphasis rx-output-e;
        rx-rate-select rx-rate-select;
        rx-squelch (off | on);
        tx-cdr (off | no);
        tx-input-equalization-adaptive (off | on);
        tx-input-equalization-manual tx-input-manual;
        tx-laser (off | on);
        tx-rate-select tx-rate-select;
        tx-squelch (off | on);
        }
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Description

Configure the optical module parameters of QSFP (Quad Small Form-Factor Pluggable) modules for all the media lanes. The following QSFP modules are supported:

- QSFP+ : 4 host lanes, supports 10g per lane, total bandwidth of 40g
- QSFP28 : 4 host lanes, supports 25g per lane, total bandwidth of 100g
- QSFP56 : 4 host lanes, supports 50g per lane, total bandwidth of 200g
- QSFP28-DD : 8 host lanes, supports 25g per lane, total bandwidth of 200g
- QSFP56-DD : 8 host lanes, supports 50g per lane, total bandwidth of 400g

The following specifications standardize how the host device interacts with the QSFP optical modules:

- The 4-lane QSFP modules adhere to SFF-8636 management interface specification.
- The 8-lane QSFP-DD modules adhere to CMIS (Common Management Interface Specification).

You can configure optical module parameters for QSFP28 and QSFP28-DD optical modules for Juniper qualified optics as well as third-party optics from Junos OS. You cannot configure optical module parameters for QSFP copper DAC (direct attach cables) or QSFP load modules.

Options

rx-cdr Controls receive clock and data recovery block of the optics.

- `Off`—Receive clock and data recovery is disabled.
- `On`—Receive clock and data recovery is enabled.

rx-output Controls receive output from the optics towards the host device.

- `Off`—Output for optics is disabled.
- `On`—Output for optics is enabled.

rx-output-amplitude Controls the amplitude of receive output for which equalization is to be applied.

- `1200mV`—No equalization for 600mV - 1200mV amplitude.
- `400mV`—No equalization for 100mV - 400mV amplitude.
- `600mV`—No equalization for 300mV - 600mV amplitude.
- `800mV`—No equalization for 400mV - 800mV amplitude.

rx-output-emphasis Equalization value (in dB) for receive output signal of the optics.

- Range: 0 through 15

rx-rate-select Controls baud rate for receive output signal.

- Range: 0 through 3

rx-squelch Squelches the receive output signal of optics.

- `Off`—Receive output signal for the optics is disabled.
- `On`—Receive output signal for the optics is enabled.

tx-cdr Control transmit clock and data recovery block of the optics.

- `Off`—Transmit clock and data recovery is disabled.
- `On`—Transmit clock and data recovery is enabled.

tx-input-equalization-adaptive Controls adaptive equalization for optics transmit input.

- `Off`—Transmit adaptive equalization is disabled.
- `On`—Transmit adaptive equalization is enabled.

tx-input-equalization-manual Specifies the manual transmit equalization status of the optics.

- Range: 0 through 15

tx-laser Specifies the transmit laser status of the optics.

- `Off`—Transmit laser is disabled.
- `On`—Transmit laser is enabled.

tx-rate-select Specifies if the optics can transmit rate selection.

- Range: 0 through 3

tx-squelch Squelches the transmit input signal of the optics.

- `Off`—Transmit input signal for the optics is disabled.
- `On`—Transmit input signal for the optics is enabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

optics-options | **11847**

lane | **9623**

language (Scripts)

IN THIS SECTION

- Syntax | 9630
- Hierarchy Level | 9630
- Description | 9630
- Default | 9631
- Options | 9631
- Required Privilege Level | 9631
- Release Information | 9632

Syntax

```
language (python | python3);
```

Hierarchy Level

```
[edit system scripts]
```

Description

Enable the execution of unsigned Python scripts on Junos devices, and specify the version of Python used to execute the scripts.

Configuring this statement enables you to execute unsigned Python commit, event, op, and SNMP automation scripts; unsigned Python action and translation scripts for custom YANG modules; and unsigned Python scripts developed using the Juniper Extension Toolkit (JET) on Junos devices. To prevent the execution of unauthorized Python code, unsigned Python scripts must be owned by either the root user or a user in the Junos OS `super-user` login class, and only the file owner can have write permission for the file.

NOTE: Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS `super-user` login class, and only the file owner can have write permission for the file. Prior to Junos OS Release 16.1R3, unsigned Python scripts must be owned by the root user.

NOTE: Starting in Junos OS Evolved Release 21.2R1, the `junos-defaults` configuration group includes the `language python` statement by default.

Local Python commit, event, op, and SNMP scripts must be configured under the hierarchy level appropriate to the script type, and the filename must include the **.py** extension. To execute Python op scripts from a remote site, you must configure the `allow-url-for-python` statement at the `[edit system scripts op]` hierarchy level.

Default

If you do not include the `language (python | python3)` statement, and the device does not include the statement in the `junos-defaults` configuration group, you cannot execute unsigned Python scripts on the device.

Options

(python | python3) Specify the Python version used to execute Python scripts. Configure `python` to use Python 2.7, or configure `python3` to use Python 3.

If you configure the `language python3` statement, the device uses Python 3 to execute scripts that support this Python version, and it uses Python 2.7 to execute scripts that do not support Python 3 in the given Junos OS release. Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute all YANG action and translation scripts. For more information, see *Understanding Python Automation Scripts for Devices Running Junos OS*.

For JET scripts, you must configure the version of the statement that corresponds to the Python version the script supports. For example, if you configure the `language python3` statement for a JET script that only supports Python 2.7, you will receive a runtime error when you execute the script.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

`python3` option added in Junos OS Release 19.4R1 for commit, event, op, and SNMP scripts.

`python3` option support for JET and YANG scripts added in Junos OS Release 20.2R1.

`python3` option added in Junos OS Evolved Release 21.1R1.

`python` option deprecated in Junos OS Release 21.1R1 and Junos OS Evolved Release 22.3R1.

RELATED DOCUMENTATION

> *Understanding Python Automation Scripts for Devices Running Junos OS*
>
> file (JET) | **7059**

large

IN THIS SECTION

- Syntax | **9632**
- Hierarchy Level | **9633**
- Description | **9633**
- Required Privilege Level | **9633**
- Release Information | **9633**

Syntax

```
large;
```

Hierarchy Level

```
[edit security screen ids-option screen-name icmp]
```

Description

Configure the device to detect and drop any ICMP frame with an IP length greater than 1024 bytes.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

large (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 9633
- Hierarchy Level | 9634
- Description | 9634
- Required Privilege Level | 9634
- Release Information | 9634

Syntax

```
large;
```

Hierarchy Level

```
[edit services screen ids-option screen-name icmp]
```

Description

Identify and drop any ICMP frame with an IP length greater than 1024 bytes, which protects against ICMP large packet attacks.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

large-community-count

IN THIS SECTION

- Syntax | 9635
- Hierarchy Level | 9635
- Description | 9635
- Required Privilege Level | 9635
- Release Information | 9635

Syntax

```
large-community-count
```

Hierarchy Level

```
[edit policy-options policy-statement <statement-name>],
[edit policy-options policy-statement <statement-name> term <term-name>]
```

Description

Count number of BGP large communities (RFC 8092).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.3R1.

laser-enable

IN THIS SECTION

- Syntax | 9636
- Hierarchy Level | 9636
- Description | 9636
- Default | 9636
- Options | 9636
- Required Privilege Level | 9636

Release Information | 9636

Syntax

```
(laser-enable | no-laser-enable);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify whether lasers are enabled or disabled.

Default

If you omit the laser-enable statement, lasers are disabled.

Options

laser-enable	Enable lasers.
no-laser-enable	Do not enable lasers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

last-block-recycle-timeout

IN THIS SECTION

- Syntax | 9637
- Hierarchy Level | 9637
- Description | 9637
- Options | 9638
- Required Privilege Level | 9638
- Release Information | 9638

Syntax

```
last-block-recycle-timeout timeout-inteval;
```

Hierarchy Level

```
[edit security nat source pool port block-allocation]
```

Description

Specify the amount of time before the last active port block is released. This option is used with the `active-block-timeout` option at [`edit security nat source pool port block-allocation`] hierarchy level. When the `active-block-timeout` option is set to 0 (zero), port blocks are filled completely before a new port block is allocated. However, the last port block remains active indefinitely. The `last-block-recycle-timeout` option allows you to release the last active block when there are no live sessions remaining. If the `active-block-timeout` option is set to anything but 0, the `last-block-recycle-timeout` option is not necessary.

Options

timeout-interval—Number of seconds before the active block is released.

- **Range:** 120 through 864000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D60.

latency (Interfaces)

IN THIS SECTION

- Syntax | 9638
- Hierarchy Level | 9638
- Description | 9639
- Required Privilege Level | 9639
- Release Information | 9639

Syntax

```
latency number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number radio-router]
```

Description

This option controls the latency weight (value 0–100).

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

latency-threshold

IN THIS SECTION

- Syntax | 9639
- Hierarchy Level | 9640
- Description | 9640
- Options | 9640
- Required Privilege Level | 9641
- Release Information | 9641

Syntax

```
latency-threshold {
    high number;
```

```
    low number;
}
```

Hierarchy Level

```
[edit services analytics interfaces]
[edit services analytics resource-profiles]
```

Description

If network analytics queue statistics monitoring is enabled, specify the high and low values (in microseconds) of the latency threshold of the queue. If you configure a latency threshold, you cannot configure the depth threshold. You can configure the latency threshold for one interface or all interfaces. Specify the high and low latency threshold numbers:

> **NOTE**: The configuration for a specific interface supersedes the global configuration for all interfaces.

Options

high *number*

Specify the maximum value for the latency threshold.

- Range:
  - Junos OS Release 13.2X51-D15 or later—1 to 100,000,000 nanoseconds (0.001 to 100,000 microseconds)
  - Junos OS Release 13.2X51-D10 or earlier—1 to 100,000 microseconds
- Default:
  - Junos OS Release 13.2X51-D15 or later—1,000,000 nanoseconds (1000 microseconds or 1 millisecond)
  - Junos OS Release 13.2X51-D10—1000 microseconds
  - Junos OS Release 13.2X50-D15—900 microseconds

low *number*

Specify the minimum value for the latency threshold.

- Range:
  - Junos OS Release 13.2X51-D15 or later—1 to 100,000,000 nanoseconds
  - Junos OS Release 13.2X51-D10 or earlier—1 to 100,000 microseconds
- Default:
  - Junos OS Release 13.2X51-D15 or later—100 nanoseconds (0.1 microseconds)
  - Junos OS Release 13.2X51-D10—50 microseconds
  - Junos OS Release 13.2X50-D15—300 microseconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Statement in the `[edit services analytics resource-profiles]` hierarchy level introduced in Junos OS Release 13.2X51-D15.

RELATED DOCUMENTATION

Network Analytics Overview

depth-threshold | 5319

layer2-control

IN THIS SECTION

Syntax | 9642

Hierarchy Level | 9643

- Description | 9643
- Options | 9643
- Required Privilege Level | 9643
- Release Information | 9644

Syntax

```
layer2-control {
    bpdu-block {
        disable-timeout seconds;
        interface interface-name;
    }
    ephemeral-db-support;
    l2pt {
        interface name {
            destination {
                vxlan-tunnel;
            }
            enable-all-ifl;
            protocol {
                (all | protocol-name);
        }
    }
    mac-rewrite {
        interface interface-name {
            enable-all-ifl;
            protocol protocol-name;

        }
        tunnel-destination-mac destination mac address;
    }
    nonstop-bridging;
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable | no-world-
readable>;
        flag flag <disable>;
```

```
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure Layer 2 control protocols to enable features such as Layer 2 protocol tunneling (L2PT) and nonstop bridging.

You can't configure the `l2pt` and `mac-rewrite` child statements for the same interface at the same time.

- The `l2pt` child statement enables L2PT transparency over destination encapsulation tunnel types on the device (such as VXLAN tunnels).
- The `mac-rewrite` child statement enables standard L2PT for interfaces that are not encapsulation tunnel interfaces, and uses MAC rewrite operations. You can configure a destination MAC address using the `tunnel-destination-mac` *`destination-mac-address`* option.

Options

ephemeral-db-support	(Junos OS only) Enable users to configure Multiple Spanning Tree Protocol (MSTP), Rapid Spanning Tree Protocol (RSTP), or VLAN Spanning Tree Protocol (VSTP) in the ephemeral configuration database.
nonstop-bridging	For details on configuring `nonstop-bridging`, see the Junos OS High Availability User Guide. When you configure this statement on routing platforms with two Routing Engines, a primary Routing Engine switches over gracefully to a backup Routing Engine and preserves Layer 2 Control Protocol (L2CP) information.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

`bpdu-block` statement added in Junos OS Release 9.4.

Statement introduced in Junos OS Release 14.1X53-D10 for EX4300 switches.

`tunnel-destination-mac` statement added in Junos OS Evolved Release 21.3R1.

`ephemeral-db-support` statement added in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Understanding Layer 2 Protocol Tunneling

Configuring Layer 2 Protocol Tunneling

instance-type | **8773**

Layer 2 Protocol Tunneling over VXLAN Tunnels in EVPN-VXLAN Bridged Overlay Networks

layer2 (enhanced-hash-key)

IN THIS SECTION

- Syntax (EX Series and QFX5100 Switch) | **9645**
- Syntax (QFX10002 Switch) | **9645**
- Hierarchy Level | **9645**
- Description | **9645**
- Default | **9646**
- Options | **9646**
- Required Privilege Level | **9646**
- Release Information | **9646**

Syntax (EX Series and QFX5100 Switch)

```
layer2 {
    no-destination-mac-address;
    no-ether-type;
    no-source-mac-address;
    vlan-id;
}
```

Syntax (QFX10002 Switch)

```
layer2 {
    no-incoming-port;
    no-destination-mac-address;
    no-ether-type;
    no-source-mac-address;
    source-mac-address;
    vlan-id;
    no-vlan-id;
    inner-vlan-id;
}
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key]
```

Description

Select the fields in the Layer 2 header that are used by the hashing algorithm to make hashing decisions.

When traffic enters a link aggregation group (LAG) bundle, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links. The hashing algorithm always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When traffic is exiting a device that has enabled ECMP, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to forward traffic to the next hop device.

The hashing algorithm only inspects the fields in the Layer 2 header when the hash mode is set to Layer 2 header. You can set the hash mode to Layer 2 header using the `set forwarding-options enhanced-hash-key hash-mode layer2-header` statement.

Default

The hash mode of the hashing algorithm is set to Layer 2 payload, by default. When the hash mode is set to Layer 2 payload, the hashing algorithm does not use fields in the Layer 2 header to make hashing decisions.

The following fields are used by the hashing algorithm when the hash mode of the hashing algorithm is set to Layer 2 header, by default:

- Destination MAC address
- Ethertype
- Source MAC address

Options

`no-destination-mac-address`	Exclude the destination MAC address field from the hashing algorithm.
`no-ether-type`	Exclude the Ethertype field from the hashing algorithm.
`no-source-mac-address`	Exclude the source MAC address field from the hashing algorithm.
`vlan-id`	Include the VLAN ID field in the hashing algorithm.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D15.

RELATED DOCUMENTATION

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic

enhanced-hash-key (EX) | **6378**

hash-mode | **7997**

layer2-liveness-detection (Receive)

IN THIS SECTION

- Syntax | **9647**
- Hierarchy Level | **9647**
- Description | **9647**
- Required Privilege Level | **9648**
- Release Information | **9648**

Syntax

```
layer2-liveness-detection;
```

Hierarchy Level

```
[edit system services subscriber-management overrides interfaces family (inet | inet6)]
```

Description

Enable a DHCP client host to determine the state of the DHCPv4 or DHCPv6 client session from the perspective of a router acting as a broadband network gateway (BNG). This statement causes the BNG to conduct a host connectivity check on its directly connected DHCPv4 and DHCPv6 clients when it receives ARP or Neighbor Discovery (ND) packets.

When the BNG receives either of these packets, it does the following:

1. Checks whether Layer 2 liveness detection for subscriber management is enabled globally for the relevant address family, inet or inet6.

2. If liveness detection is not enabled, then the BNG responds as usual to the received packets without checking the state of the client session.

   If liveness detection is enabled for the family, then the BNG checks whether the client session is still in the bound state.

3. If the client session is bound, the BNG responds to the client with the appropriate ARP or ND packet.

   If the session is not bound, the BNG drops the received packet. It does not send an ARP or ND response packet to the host, enabling the host to determine that the BNG considers the session to be down.

This behavior can be referred to as the *receive* functionality for BNG Layer 2 liveness detection, as opposed to the *send* functionality configured with the `layer2-liveness-detection (Send)` statement for DHCP relay or DHCP local server.

The usefulness of the receive functionality depends on the ability of the DHCP client host to reclaim resources from the stale client based on the absence of a response packet from the BNG for an unbound client session. If this capability requires a change in the client implementation, you may want to use the send functionality.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

- DHCP Liveness Detection Using ARP and Neighbor Discovery Packets
- DHCP Liveness Detection Overview
- *Configuring Junos OS Enhanced Subscriber Management*

layer2-liveness-detection (Send)

IN THIS SECTION

- Syntax | 9649
- Hierarchy Level | 9649
- Description | 9650
- Options | 9650
- Required Privilege Level | 9651
- Release Information | 9651

Syntax

```
layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval seconds;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name liveness-detection
method],
[edit forwarding-options dhcp-relay group group-name liveness-detection method],
[edit forwarding-options dhcp-relay liveness-detection method],


[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method],
[edit system services dhcp-local-server dhcpv6 liveness-detection method],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name liveness-detection
method],
[edit system services dhcp-local-server group group-name liveness-detection method],
[edit system services dhcp-local-server liveness-detection method].
```

Description

Configure a router acting as a broadband network gateway (BNG) to conduct a host connectivity check on its directly connected DHCPv4 and DHCPv6 clients to determine the validity and state of the DHCP client session, and to clean up inactive sessions.

The BNG sends ARP or ND request packets to the each DHCP client at a configurable interval, then waits for a response. If it receives a response from a client before the interval times out, it sends another request to the client when the timer expires.

If the BNG does not receive a response before the interval times out, it sets the timer to 30 seconds and sends another request. This is the first retry attempt.

If it receives a response from a client before the 30-second interval times out, it sends another request to the client when the timer expires. If the 30-second timer expires before a response is received, the BNG sets the timer to 10 seconds and sends another request. This is the second retry attempt. If the BNG does not receive a response within this interval it resets the timer to 10 seconds and sends another request. The BNG continues to send requests at 10-second intervals until it either receives a response from the client before the interval times out or exhausts the number of retry attempts.

The first retry attempt uses a 30-second interval. Subsequent retries occur at 10-second intervals. The number of possible 10-second retries is therefore the total number of retries minus 1. For example, if you configure 5 retries, there is one 30-second retry and up to four 10-second retries.

If the BNG attempts all the retries and never receives a response from a client within the interval, the client session is declared to be down.

> **NOTE**: The only option to the `failure-action` statement supported by Layer 2 liveness detection is `clear-binding`.

Options

max-consecutive-retries *number*

Maximum number of consecutive times that the router sends an ARP request packet in the absence of an ARP response packet.

- **Range:** 3 through 6 retries
- **Default:** 3 retries

transmit-interval *seconds*

Initial interval that the router waits for an ARP response after sending an ARP request packet to the client or waits for an ND response packet after sending an NG request packet to the client.

- **Range:** 300 through 1800 seconds
- **Default:** 300 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

DHCP Liveness Detection Overview

layer2-map

IN THIS SECTION

- Syntax | 9652
- Hierarchy Level | 9652
- Description | 9652
- Default | 9652
- Required Privilege Level | 9652
- Release Information | 9653

Syntax

```
layer2-map;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name protocols isis interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit protocols isis interface interface-name],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Enable Layer 2 mapping of ARP or neighbor discovery next hops in the kernel for the specified instance or interface. IS-IS Layer 2 mapping ensures that forwarding next-hop resolution is topology-driven rather than traffic-driven, which results in minimal traffic loss while activating an Ethernet link.

IS-IS LAN and point-to-point Hellos supply all relevant Layer 2 and Layer 3 binding address information for address resolution. The device at the receiving end can extract the information and populate the ARP or neighbor discovery cache of the kernel even before the route installation time. When Layer 2 mapping is enabled, IS-IS installs ARP or neighbor discovery next-hop entries into the forwarding table. Because this provides Layer 2 next-hop bindings ahead of time, IS-IS networks do not experience traffic loss while bringing up a link.

Default

The default setting is `no-layer2-map` for both instances and interfaces.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Layer 2 Mapping for IS-IS

Example: Configuring Layer 2 Mapping for IS-IS

layer2-policer

IN THIS SECTION

- Syntax | 9653
- Hierarchy Level | 9653
- Description | 9654
- Options | 9654
- Required Privilege Level | 9654
- Release Information | 9654

Syntax

```
layer2-policer {
    output-policer policer-name;
    output-three-color policer-name;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

Description

Specify the output policers to be used in the dynamic profile.

Options

output-policer ***policer-name***—Two-color output policer to associate with the interface. You define this policer at the `[edit firewall policer]` hierarchy level.

output-three-color ***policer-name***—Tricolor output policer to associate with the interface. You define this policer at the `[edit firewall]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Applying the Policers to Dynamic Profile Interfaces

Creating a Dynamic Profile for the Complex Configuration

layer2-policer (EX)

IN THIS SECTION

- Syntax | 9655
- Hierarchy Level | 9655
- Description | 9655
- Options | 9655
- Required Privilege Level | 9656

Release Information | 9656

Syntax

```
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series, MX Series, and T Series routers, and for aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces on EX Series switches, apply Layer 2 logical interface policers. The following policers are supported:

- Two-color
- Single-rate tricolor marking (srTCM)
- Two-rate tricolor marking (trTCM)

Two-color and tricolor policers are configured at the `[edit firewall]` hierarchy level.

Options

`input-policer` *`policer-name`*—Two-color input policer to associate with the interface. This statement is mutually exclusive with the `input-three-color` statement.

`input-three-color` *`policer-name`*—Tricolor input policer to associate with the interface. This statement is mutually exclusive with the `input-policer` statement.

`output-policer` *`policer-name`*—Two-color output policer to associate with the interface. This statement is mutually exclusive with the `output-three-color` statement.

`output-three-color` *`policer-name`*—Tricolor output policer to associate with the interface. This statement is mutually exclusive with the `output-policer` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Applying Layer 2 Policers to Gigabit Ethernet Interfaces

Configuring Gigabit Ethernet Two-Color and Tricolor Policers

layer2-policer (Hierarchical Policer)

IN THIS SECTION

- Syntax | 9657
- Hierarchy Level | 9657
- Description | 9657
- Options | 9657
- Required Privilege Level | 9657
- Release Information | 9657

Syntax

```
layer2-policer {
    input-hierarchical-policer policer-name

}
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number]
```

Description

Apply a hierarchical policer to the Layer 2 input traffic for all protocol families at the physical or logical interface. The following interfaces are supported:

- SONET interfaces hosted on M40e, M120, and M320 edge routers with incoming Flexible PIC Concentrators (FPCs) as SFPC and outgoing FPCs as FFPC
- Interfaces on MX Series, T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Hierarchical Policers

input-hierarchical-policer | 8704

Two-Color and Three-Color Policers at Layer 2

layer2-protocol-tunneling

IN THIS SECTION

- Syntax | 9658
- Hierarchy Level | 9658
- Description | 9659
- Default | 9659
- Options | 9659
- Required Privilege Level | 9660
- Release Information | 9660

Syntax

```
layer2-protocol-tunneling all | protocol-name {
 drop-threshold number;
    shutdown-threshold number;
}
```

Hierarchy Level

```
[edit vlans vlan-name dot1q-tunneling]
```

Description

Enable Layer 2 protocol tunneling (L2PT) on a VLAN on switches that do not use the the Enhanced Layer 2 Software (ELS) configuration style (which includes EX2200, EX3300, EX4200, EX4500, and EX4450 switches).

> **NOTE**: This command is not available on switches that use ELS configuration style (including EX2300, EX3400, EX4300, EX4600, EX9200, and any QFX Series switches that support L2PT). On those switches, you enable L2PT using the `protocol` statement at the `[edit protocols layer2-control mac-rewrite interface interface-name]` hierarchy level. See *Configuring Layer 2 Protocol Tunneling*.
>
> For details on ELS, see *Using the Enhanced Layer 2 Software CLI*.

The remaining statements are explained separately. See CLI Explorer.

Default

L2PT is not enabled.

Options

`all`—Enable all supported Layer 2 protocols.

`protocol-name`—Name of the Layer 2 protocol. Values are:

- `802.1x`—IEEE 802.1X authentication
- `802.3ah`—IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM)

  > **NOTE**: If you enable L2PT for untagged OAM LFM packets, do not configure LFM on the corresponding access interface.

- `cdp`—Cisco Discovery Protocol
- `e-lmi`—Ethernet local management interface
- `gvrp`—GARP VLAN Registration Protocol
- `lacp`—Link Aggregation Control Protocol

> **NOTE:** If you enable L2PT for untagged LACP packets, do not configure LACP on the corresponding access interface.

- `lldp`—Link Layer Discovery Protocol
- `mmrp`—Multiple MAC Registration Protocol
- `mvrp`—Multiple VLAN Registration Protocol
- `stp`—Spanning Tree Protocol, Rapid Spanning Tree Protocol, and Multiple Spanning Tree Protocol
- `udld`—Unidirectional Link Detection (UDLD)
- `vstp`—VLAN Spanning Tree Protocol
- `vtp`—VLAN Trunking Protocol

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0 for EX Series switches.

RELATED DOCUMENTATION

show ethernet-switching layer2-protocol-tunneling interface | **19652**

show ethernet-switching layer2-protocol-tunneling statistics | **19654**

show ethernet-switching layer2-protocol-tunneling vlan | **19658**

Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

layer2-unicast-replies

IN THIS SECTION

- Syntax | 9661
- Hierarchy Level | 9661
- Description | 9662
- Required Privilege Level | 9662
- Release Information | 9662

Syntax

```
layer2-unicast-replies;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Override the setting of the broadcast bit in DHCP request packets and instead use the Layer 2 unicast transmission method to transmit DHCP Offer reply packets and DHCP ACK reply packets from the DHCP server to DHCP clients during the discovery process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

dhcp-relay | 5620

layer2-policer (firewall)

IN THIS SECTION

- Syntax | 9663
- Hierarchy Level | 9663
- Description | 9663
- Default | 9665
- Required Privilege Level | 9665
- Release Information | 9665

Syntax

```
layer2-policer
```

Hierarchy Level

```
[edit firewall family <firewall filter family> filter <firewall filter name>]
```

Description

By default, a packet gets policed by a policer on the packet's layer 3 header length along with the packet's payload length. Applying the `layer2-policer` configuration on a firewall filter allows the policer to account layer 2 overhead, which changes the policer behavior as well as the firewall filter counter behavior and will calculate filter count bytes to full frame length.

- **Default policer overhead calculation**
  - layer 3 header length + payload length
- **Policer overhead calculation after applying the `layer2-policer` configuration**
  - layer 2 header length + layer 3 header length + payload length

Example

The following configurations show the `layer2-policer` configuration applied to the `f_inet`, `f_inet6`, `f_mpls`, and `f_any` firewall filters configured for the supported firewall filter families - INET, INET6, MPLS, and ANY, respectively.

Setting the `layer2-policer` configuration in the the `f_inet` firewall filter configuration for INET firewall filter family.

```
set firewall family inet filter f_inet layer2-policer
set firewall family inet filter f_inet term t1 then policer pc1
set firewall family inet filter f_inet term t1 then count c1
set firewall family inet filter f_inet term t1 then accept
```

Setting the `layer2-policer` configuration in the `f_inet6` firewall filter configuration for INET6 firewall filter family.

```
set firewall family inet6 filter f_inet6 layer2-policer
set firewall family inet6 filter f_inet6 term t1 then policer pc1
set firewall family inet6 filter f_inet6 term t1 then count c1
set firewall family inet6 filter f_inet6 term t1 then accept
```

Setting the `layer2-policer` configuration in the `f_mpls` firewall filter configuration for MPLS firewall filter family.

```
set firewall family mpls filter f_mpls layer2-policer
set firewall family mpls filter f_mpls term t1 then policer pc1
set firewall family mpls filter f_mpls term t1 then count c1
set firewall family mpls filter f_mpls term t1 then accept
```

Setting the `layer2-policer` configuration in the `f_any` firewall filter configuration for ANY firewall filter family.

```
set firewall family any filter f_any layer2-policer
set firewall family any filter f_any term t1 then policer pc1
set firewall family any filter f_any term t1 then count c1
set firewall family any filter f_any term t1 then accept
```

Configuring the policer in the firewall filter.

```
set firewall policer pc1 if-exceeding bandwidth-limit 20k
set firewall policer pc1 if-exceeding burst-size-limit 2k
set firewall policer pc1 then discard
```

Applying the firewall filter to the interface. In the following configuration, the firewall filter for INET firewall filter family is applied to the interface.

```
set interfaces et-2/0/13:2 unit 0 family inet filter input f_inet
```

Viewing statistics. In the following output, statistics of the firewall filter for the INET firewall family is shown.

```
Filter: f_inet
Counters:
Name                          Bytes
Packets
c1                            872960
6820
Policers:
Name                          Bytes                  Packets
pc1-t1                        370688                  2896
```

Default

Default policer overhead calculation is - layer 3 header length + payload length.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Evolved Release 24.3R1

Supported on:

- PTX10001-36MR
- PTX10003
- PTX10004 - LC1201 & LC1202

- PTX10008 - LC1201 & LC1201
- PTX10016 - LC1201 & LC1201
- PTX10002-36QDD

layer3-domain-identifier

IN THIS SECTION

- Syntax | 9666
- Hierarchy Level | 9666
- Description | 9666
- Options | 9666
- Required Privilege Level | 9667
- Release Information | 9667

Syntax

```
layer3-domain-identifier identifier;
```

Hierarchy Level

```
[edit routing-instances instance-name]
```

Description

Specify a Layer 3 domain ID within a routing instance.

Options

`id`—Layer 3 identification number.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R2.

RELATED DOCUMENTATION

Configuring a Layer 2 Virtual Switch on an EX Series Switch

lcd-menu

IN THIS SECTION

- Syntax | 9667
- Hierarchy Level | 9668
- Description | 9668
- Options | 9668
- Required Privilege Level | 9668
- Release Information | 9669

Syntax

EX3200, EX3300, EX4200, or EX4500 switch:

```
lcd-menu fpc slot-number {
    menu-item (menu-name | menu-option) <disable>;
}
```

EX6200 or EX8200 switch or XRE200 External Routing Engine:

```
lcd-menu {
    menu-item (menu-name | menu-option) <disable>;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Disable or enable the Maintenance menu or the Status menu in the LCD panel.

Options

none—(EX6200 and EX8200 switches and XRE200 External Routing Engines only) Disable or enable the specified menu or menu options.

fpc ***slot-number***—(EX3200, EX3300, EX4200, and EX4500 switches only) Disable or enable the specified menu or menu options, where ***slot-number*** is:

- 0—On standalone switches.
- 0–9—On a device in a Virtual Chassis. The value is the member ID of the device.

> **NOTE**: This option is not available on an EX8200 Virtual Chassis. The LCD panel on an XRE200 External Routing Engine provides information for the XRE200 External Routing Engine only.

disable—(Optional) Disable the specified menu.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the LCD Panel on EX Series Switches (CLI Procedure)

LCD Panel in EX3200 Switches

LCD Panel in EX3300 Switches

LCD Panel in EX4200 Switches

LCD Panel in EX4500 Switches

LCD Panel in an EX6200 Switch

LCD Panel in an EX8200 Switch

LCD Panel in an XRE200 External Routing Engine

lcp-max-conf-req

IN THIS SECTION

- Syntax | 9669
- Hierarchy Level | 9670
- Description | 9670
- Options | 9670
- Required Privilege Level | 9670
- Release Information | 9670

Syntax

```
lcp-max-conf-req number
```

Hierarchy Level

```
[edit interfaces so-fpc/pic/port unit number ppp-options]
```

Description

Set the maximum number of LCP Configure-Requests to be sent, after which the router goes to LCP down state.

Options

number From 0 to 65,535, where 0 means send infinite LCP Configure-Requests, and any other value specifies the maximum number LCP Configure-Requests to send and then stop sending.

Default 254

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring PPP

ppp-options | **12949**

lcp-renegotiation

IN THIS SECTION

- Syntax | 9671
- Hierarchy Level | 9671
- Description | 9671
- Required Privilege Level | 9671
- Release Information | 9672

Syntax

```
lcp-renegotiation;
```

Hierarchy Level

```
[edit access group-profile profile-name l2tp],
[edit access profile profile-name client client-name l2tp]
```

Description

Configure the L2TP network server (LNS) so it renegotiates the link control protocol (LCP) with the PPP client. When LCP renegotiation is disabled, LNS uses the pre-negotiated LCP parameters between the L2TP access concentrator (LAC) and PPP client to set up the session. When LCP renegotiation is enabled, authentication is also renegotiated.

NOTE: This statement is not supported at the `[edit access group-profile l2tp]` hierarchy level for L2TP LNS on MX Series routers.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Layer 2 Tunneling Protocol (L2TP)

L2TP LNS Inline Service Interfaces

lcp-restart-timer

IN THIS SECTION

- Syntax | 9672
- Hierarchy Level | 9672
- Description | 9673
- Options | 9673
- Required Privilege Level | 9673
- Release Information | 9673

Syntax

```
lcp-restart-timer milliseconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number ppp-options],
[edit logical-systems logical-system-name interfaces interface-nameunit logical-unit-number ppp-
options]
```

Description

For interfaces with PPP, PPP TCC, PPP over Ethernet, PPP over ATM, and PPP over Frame Relay encapsulations, configure a restart timer for the Link Control Protocol (LCP) component of a PPP session.

Options

milliseconds—The time, in milliseconds, between successive LCP configuration requests.

- **Range:** 20 through 10000 milliseconds
- **Default:** 3 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring PPP

ldap-options

IN THIS SECTION

- Syntax | 9674
- Hierarchy Level | 9675
- Description | 9675
- Options | 9675

Required Privilege Level | 9676
Release Information | 9676

Syntax

```
ldap-options {
        revert-interval seconds;
        base-distinguished-name base-distinguished-name;
        search {
            admin-search;
        }
        allowed-groups {
            group-name {
                address-assignment {
                    pool pool-name;
                }
        }
    }
    ldap-server {
        ip address;
    }
}
address-assignment {
    pool pool-name1 {
        family inet {
            network 100.127.255.255/10;
            xauth-attributes {
                primary-dns 100.127.255.255/12;
                secondary-dns 110.127.255.255/12;
                primary-wins 100.127.255.255/12;
                secondary-wins 110.127.255.255/12;
            }
        }
    }
    pool pool-name2 {
        family inet {
            network 120.127.255.255/10;
            xauth-attributes {
```

```
                primary-dns 120.127.255.255/12;
                secondary-dns 130.127.255.255/12;
                primary-wins 120.127.255.255/12;
                secondary-wins 130.127.255.255/12;
            }
        }
    }
}
firewall-authentication {
    web-authentication {
        default-profile default-profile-name;
    }
}
```

Hierarchy Level

```
[edit access],
[edit access profile profile-name authentication-order ldap]
```

Description

Configure LDAP authentication options.

You can configure user groups using `ldap-options` command for the user groups that are user authenticated. You can authenticate users that are assigned roles according to their LDAP group memberships. The `allowed-groups` attribute authenticates users that are assigned according to their group memberships. If none of the user groups match a user group, then the user cannot access the system.

Membership characteristics are queried from the LDAP server as per configuration. After firewall authentication, a user can be assigned IP addresses from the associated pools with the authenticated group.

Options

allowed-groups Allow members of only specific groups to sign in. Group lists are limited to 255 bytes. The order in which the membership attribute is received from the LDAP server determines how a user is associated with the configured (allowed) groups. To match the user, the first group in the list received from the LDAP server that matches any of the configured groups is used.

Any user who is a member of more than one group can obtain resources from either group, depending on the order of the LDAP server's response. To ensure that the user is assigned the intended resource with certainty, it is recommended that the user belong to only one group.

group-name Name of the group which should be allowed.

name Address pool name

The remaining options are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

`allowed-groups` option introduced in Release 21.4R1 of Junos OS.

ldap-server

IN THIS SECTION

- Syntax | 9677
- Hierarchy Level | 9677
- Description | 9677
- Options | 9677
- Required Privilege Level | 9677
- Release Information | 9678

Syntax

```
ldap-server server-address {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
    no-tls-certificate-check;
    tls-min-version (v1.1 | v1.2);
    tls-peer-name;
    tls-timeout;
    tls-type {
        start-tls;
    }
}
```

Hierarchy Level

```
[edit access]
[edit access profile profile-name]
```

Description

Specify that the device uses a Lightweight Directory Access Protocol (LDAP) server for authentication.

Options

server-address—Address of the LDAP authentication server.

The remaining statements are explained separately.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Obtaining Username and Role Information Through Firewall Authentication

LDAP Functionality in Integrated User Firewall

ldap-url

IN THIS SECTION

- Syntax | 9678
- Hierarchy Level | 9678
- Description | 9679
- Options | 9679
- Required Privilege Level | 9679
- Release Information | 9679

Syntax

```
<ldap-url url-name>;
```

Hierarchy Level

```
[edit security certificates certification-authority ca-profile-name]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) (Optional) Specify the Lightweight Directory Access Protocol (LDAP) URL for digital certificates.

Options

url-name—Name of the LDAP URL.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Specifying an LDAP URL

ldp

IN THIS SECTION

- Syntax | 9680
- Hierarchy Level | 9683
- Description | 9683
- Default | 9683
- Options | 9684
- Required Privilege Level | 9684
- Release Information | 9684

Syntax

```
ldp {
    auto-targeted-session {
         maximum-sessions;
         no-dynamic-tunnels;
         no-rlfa;
         teardown-delay;
        {
        color-import-policy [ policy-names ]
    (deaggregate | no-deaggregate);
        dod-request-policy
    dual-transport {
        inet-lsr-id ;
        inet6-lsr-id;
        }
    egress-policy [ policy-names ];
    entropy-label {
        ingress-policy [ policy-names ]
    explicit-null;
    export [ policy-names ];
    family {
        inet;
        inet6;
      }
    graceful-restart {
        disable;
        helper-disable;
        maximum-neighbor-reconnect-time;
        maximum-neighbor-recovery-time;
        reconnect-time;
        recovery-time;
      }
    igp-synchronization {
        holddown-interval;
       }
    import [ policy-names];
    interface (interface-name | all);
    keepalive-interval seconds;
    keepalive-timeout seconds;
        l2-smart-policy;
        label-withdrawal-delay;
```

```
    log-updown {
        trap {
            disable;
        }
     }
    longest-match {
        policy {
            [ policy-names];
          }
      }
    make-before-break {
        switchover-delay ;
        timeout;
      }
        mtu-discovery;
    neighbor [ neighbor];
        next-hop {
            merged {
              policy [ policy names];
            }
            no-rsvp-tunneling {
              policy [ policy names];
             }
        }
        no-deaggregate;
    no-forwarding;
        no-mtu-discovery;
        no-propagate-ttl;
    oam {
        bfd-liveness-detection {
            detection-time {
                threshold;
                }
            }
            ecmp;
            failure-action {
                remove-nexthop;
                remove-route;
            }
            holddown-interval ;
            minimum-interval;
            minimum-receive-interval;
            multiplier;
```

```
            no-adaptation;
            no-router-alert-option;
            transmit-interval {
                 minimum-interval;
                 threshold;
               }
        bfd-port-egress-policy [policy names];
        fec [fec address];
        ingress-policy [policy names];
        lsp-ping-interval;
        lsp-ping-multiplier;
        periodic-traceroute {
            exp;
            fanout;
            frequency;
            paths;
            retries;
            source;
            ttl;
            wait;
        }
      }
    p2mp {
        no-rsvp-tunneling;
        recursive {
            route;
          }
        root-address [address];
      }
    policing {
        fec [address];
      }
    preference preference;
        rib-group;
        rib-group6;
    session [address];
    session-group session prefix;
    session-protection {
        timeoute;
       }
        sr-mapping-client {
            policy policy;
           }
```

```
        }
    strict-targeted-hellos;
        targeted-hello {
            hello-interval;
            hold-time;
       }
    traceoptions {
            flag [trace flags];
            file [file name];
       }
    track-igp-metric;
    traffic-statistics {
            file [file name];
            interval;
            no-penultimate-hop;
            sensor-based-stats;
       }
    transport-address (address | interface | router-id);
transport-preference [ipv4 | ipv6];
        upstream-label-assignment;
  }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable LDP routing on the router or switch.

You must include the `ldp` statement in the configuration to enable LDP on the router or switch.

Default

LDP is disabled on the router.

Options

transport-preference ipv4 | ipv6

Select the preferred transport for TCP connection when both IPv4 and IPv6 are enabled. If `transport-preference ipv4` is configured, LDP will attempt to establish the TCP connection using IPv4. If `transport-preference ipv6` is configured, LDP will attempt to establish the TCP connection using IPv6.

- **Default:** ipv6

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`dual-transport` statement introduced in Junos OS Release 16.1 for the M320 Series, MX Series, and PTX Series.

`family` statement introduced in Junos OS Release 16.1 for the M320 Series, MX Series, and PTX Series.

`transport-preference` option introduced in Junos OS Release 16.1 for the M320 Series, MX Series, and PTX Series.

RELATED DOCUMENTATION

Minimum LDP Configuration

Enabling and Disabling LDP

ldp-p2mp

IN THIS SECTION

- Syntax | 9685
- Hierarchy Level | 9685
- Description | 9686
- Required Privilege Level | 9686
- Release Information | 9686

Syntax

```
ldp-p2mp;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name provider-tunnel],
[edit logical-systems logical-system-name routing-instances instance-name provider-tunnel
selective wildcard-group-inet wildcard-source],
[edit logical-systems logical-system-name routing-instances instance-name provider-tunnel
selective wildcard-group-inet6 wildcard-source],
[edit logical-systems logical-system-name routing-instances instance-name provider-tunnel
selective group group-prefix wildcard-source],
[edit logical-systems logical-system-name routing-instances instance-name provider-tunnel
selective group group-prefix source source-prefix],
[edit protocols mvpn inter-region-template template template-name all-regions],
[edit protocols mvpn inter-region-template template template-name region region-name],
[edit routing-instances instance-name provider-tunnel]
[edit routing-instances instance-name provider-tunnel inter-as],
[edit routing-instances instance-name provider-tunnel selective wildcard-group-inet wildcard-
source],
[edit routing-instances instance-name provider-tunnel selective wildcard-group-inet6 wildcard-
source],
[edit routing-instances instance-name provider-tunnel selective group group-prefix wildcard-
source],
```

```
[edit routing-instances instance-name provider-tunnel selective group group-prefix source source-prefix]
```

Description

Specify a point-to-multipoint provider tunnel with LDP signalling for an MBGP MVPN.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs

ldp-stitching (Protocols OSPF)

IN THIS SECTION

- Syntax | 9687
- Hierarchy Level | 9687
- Description | 9687
- Required Privilege Level | 9687
- Release Information | 9687

Syntax

```
ldp-stitching;
```

Hierarchy Level

```
[edit protocols ospf source-packet-routing],
[edit routing-instances name protocols ospf source-packet-routing]
```

Description

Enable segment routing to LDP stitching.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

LDP Mapping Server for Interoperability of Segment Routing with LDP Overview

ldp-tunneling

IN THIS SECTION

Syntax | 9688
Hierarchy Level | 9688

- Description | 9688
- Required Privilege Level | 9688
- Release Information | 9688

Syntax

```
ldp-tunneling;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
[edit protocols source-packet-routing source-routing-path lsp-name]
```

Description

Enable the LSP to be used for LDP tunneling.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling LDP over RSVP-Established LSPs

leaf

IN THIS SECTION

- Syntax | 9689
- Hierarchy Level | 9689
- Description | 9689
- Options | 9690
- Required Privilege Level | 9690
- Release Information | 9690

Syntax

```
leaf {
    deterministic-ar-policy [ deterministic-ar-policy-name ... ];
    replicator-activation-delay seconds;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols evpn assisted-replication],
[edit routing-instances name protocols evpn assisted-replication],
[edit protocols evpn assisted-replication]
```

Description

Configure a device in an EVPN network into the assisted replication (AR) leaf role. AR leaf devices might also be called AR client devices.

AR helps to optimize broadcast, unknown unicast, and multicast (BUM) traffic in EVPN networks. To enable AR, you configure devices in the EVPN network to operate as AR replicator and AR leaf devices.

AR replicator devices help perform BUM traffic replication and forwarding tasks for AR leaf devices. When a leaf device receives BUM traffic, by default is uses ingress replication to replicate the traffic into the EVPN core. When you configure a device as an AR leaf device, the device forwards the traffic on an

AR tunnel to an AR replicator device that can better handle the replication load. The AR replicator device then replicates and forwards the traffic to other overlay tunnels. You configure a secondary loopback interface IP address on each AR replicator device for the AR tunnel.

You can configure multiple AR replicator devices in an EVPN network. In that case, AR leaf devices load-balance among the available AR replicator devices as follows:

- AR leaf devices use the following platform-specific load balancing methods by default (see *Default AR Leaf Load-Balancing When Detecting Multiple AR Replicators*):
  - AR leaf devices that are EX4650 switches or switches in the QFX5000 line automatically designate a particular AR replicator device for a VLAN or VXLAN network identifier (VNI) to load-balance among the available AR replicators.
  - AR leaf devices in the QFX10000 line of switches actively load-balance among the available AR replicators based on traffic flow levels within a VLAN or VNI.
- Supported AR leaf devices in the QFX5000 line can steer particular multicast flows to specified preferred AR replicators if you explicitly configure deterministic AR replicator routing policies and assign those policies to an EVPN instance on the AR leaf device. (See *Deterministic Load Balancing and Traffic Steering to AR Replicators*.)

Options

replicator-activation-delay

Interval (in seconds) to wait after receiving an AR replicator route from an AR replicator before starting to forward traffic to it. You might want to adjust the delay to make sure AR replicator devices have fully learned the current EVPN state from the network after an AR replicator goes down and comes back up again.

- **Default:** 10 seconds
- **Range:** 0 through 180 seconds

The remaining statements are explained separately. Click on a statement link in the Syntax section, or see CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.4R2 and 19.4R1.

`deterministic-ar-policy` option added in Junos OS Release 23.4R2.

RELATED DOCUMENTATION

Assisted Replication Multicast Optimization in EVPN Networks

Deterministic Load Balancing and Traffic Steering to AR Replicators

replicator | 14387

show evpn multicast-snooping assisted-replication next-hops | 19807

show evpn multicast-snooping assisted-replication multihomed-peers | 19804

show evpn multicast-snooping assisted-replication replicators | 19811

leaf-tunnel-limit-inet (MVPN Selective Tunnels)

IN THIS SECTION

- Syntax | 9691
- Hierarchy Level | 9691
- Description | 9692
- Default | 9692
- Options | 9692
- Required Privilege Level | 9692
- Release Information | 9692

Syntax

```
leaf-tunnel-limit-inet number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name provider-tunnel
selective],
[edit routing-instances instance-name provider-tunnel selective]
```

Description

Configure the maximum number of selective leaf tunnels for IPv4 control-plane routes.

The purpose of the `leaf-tunnel-limit-inet` statement is to supplement the multicast forwarding-cache limit when the MVPN `rpt-spt` mode is configured and when traffic is flowing through selective service provider multicast service inteface (S-PMSI) tunnels and is forwarded by way of the (*,G) entry, even though the forwarding cache limit has already blocked the forwarding entries from being created.

The `leaf-tunnel-limit-inet` statement limits the number of Type-4 leaf autodiscovery (AD) route messages that can be originated by receiver provider edge (PE) routers in response to receiving from the sender PE router S-PMSI AD routes with the leaf-information-required flag set. Thus, this statement limits the number of leaf nodes that are created when a selective tunnel is formed.

You can configure the statement only when the MVPN mode is `rpt-spt`.

This statement is independent of the `cmcast-joins-limit-inet` statement and of the `forwarding-cache threshold` statement.

Setting the `leaf-tunnel-limit-inet` statement or reducing the value of the limit does not alter or delete the already existing and installed routes. If needed, you can run the `clear pim join` command to force the limit to take effect. Those routes that cannot be processed because of the limit are added to a queue, and this queue is processed when the limit is removed or increased and when existing routes are deleted.

Default

Unlimited

Options

number Maximum number of selective leaf tunnels for IPv4.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Examples: Configuring the Multicast Forwarding Cache

Example: Configuring MBGP Multicast VPN Topology Variations

leaf-tunnel-limit-inet6 (MVPN Selective Tunnels)

IN THIS SECTION

- Syntax | 9693
- Hierarchy Level | 9693
- Description | 9693
- Default | 9694
- Options | 9694
- Required Privilege Level | 9694
- Release Information | 9694

Syntax

```
leaf-tunnel-limit-inet6 number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name provider-tunnel
selective],
[edit routing-instances instance-name provider-tunnel selective]
```

Description

Configure the maximum number of selective leaf tunnels for IPv6 control-plane routes.

The purpose of the `leaf-tunnel-limit-inet6` statement is to supplement the multicast forwarding-cache limit when the MVPN `rpt-spt` mode is configured and when traffic is flowing through selective service

provider multicast service inteface (S-PMSI) tunnels and is forwarded by way of the (*,G) entry, even though the forwarding cache limit has already blocked the forwarding entries from being created.

The `leaf-tunnel-limit-inet6` statement limits the number of Type-4 leaf autodiscovery (AD) route messages that can be originated by receiver provider edge (PE) routers in response to receiving from the sender PE router S-PMSI AD routes with the leaf-information-required flag set. Thus, this statement limits the number of leaf nodes that are created when a selective tunnel is formed.

You can configure the statement only when the MVPN mode is `rpt-spt`.

This statement is independent of the `cmcast-joins-limit-inet6` statement and of the `forwarding-cache threshold` statement.

Setting the `leaf-tunnel-limit-inet6` statement or reducing the value of the limit does not alter or delete the already existing and installed routes. If needed, you can run the `clear pim join` command to force the limit to take effect. Those routes that cannot be processed because of the limit are added to a queue, and this queue is processed when the limit is removed or increased and when existing routes are deleted.

Default

Unlimited

Options

number Maximum number of selective leaf tunnels for IPv6.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Examples: Configuring the Multicast Forwarding Cache

Example: Configuring MBGP Multicast VPN Topology Variations

learn-sip-register

IN THIS SECTION

- Syntax | 9695
- Hierarchy Level | 9695
- Description | 9695
- Required Privilege Level | 9696
- Release Information | 9696

Syntax

```
learn-sip-register;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Learn potential incoming SIP calls by inspecting the SIP register method.

More information: You can enable the router to accept any incoming SIP calls for the endpoint devices that are behind the NAT firewall. When a device behind the firewall registers with the proxy that is outside the firewall, the AS or Multiservices PIC maintains the registration state. When the `learn-sip-register` statement is enabled, the router can use this information to accept inbound calls. If this statement is not configured, no inbound calls are accepted; only the devices behind the firewall can call devices outside the firewall.

NOTE: You can manually inspect the SIP register by running the `show services stateful-firewall sip-register` command. This command is not supported for the MX240.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

- *ALG Descriptions*
- *Configuring Application Sets*
- *Configuring SIP*
- *Examples: Configuring Application Protocols*
- *Verifying the Output of ALG Sessions*

lease-time

IN THIS SECTION

- Syntax | 9697
- Hierarchy Level | 9697
- Description | 9697
- Default | 9697
- Options | 9697
- Required Privilege Level | 9697
- Release Information | 9698

Syntax

```
lease-time (length | infinite);
```

Hierarchy Level

```
[edit interfaces interface-name                unit                logical-unit-
number            family     inet dhcp]
```

Description

Request a specific lease time for the IP address. The lease time is the length of time in seconds that a client holds the lease for an IP address assigned by a DHCP server.

Default

If no lease time is requested by client, then the server sends the lease time. The default lease time on a Junos OS DHCP server is one day.

Options

seconds Request a lease time of a specific duration.

> **NOTE**: Starting in Junos OS Release 23.4R1, the DHCP client silently discards the DHCP OFFER which has a lease-time of less than 15 seconds.

- **Range:** 60 through 2147483647 seconds

infinite Request that the lease never expire.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

DHCP Client

DHCPv6 Client

lease-time (dhcp-client)

IN THIS SECTION

- Syntax | 9698
- Hierarchy Level | 9698
- Description | 9698
- Options | 9699
- Required Privilege Level | 9699
- Release Information | 9699

Syntax

```
lease-time seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcp-client]
```

Description

Specify the time to negotiate and exchange Dynamic Host Configuration Protocol (DHCP) information.

Options

seconds Request time to negotiate and exchange information.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

Starting with Junos OS Release 17.3R1, on all SRX Series Firewalls and vSRX Virtual Firewall instances, the CLI option `dhcp-client` at `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy is changed to `dhcp` to keep consistent with other Junos platforms. There is no change in the functionality with the changed CLI option `dhcp`.

RELATED DOCUMENTATION

DHCP Overview

lease-time-threshold (DHCP Local Server and DHCP Relay Agent)

IN THIS SECTION

- Syntax | 9700
- Hierarchy Level | 9700
- Description | 9700
- Options | 9700
- Required Privilege Level | 9700
- Release Information | 9701

Syntax

```
lease-time-threshold seconds;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay lease-time-validation],
[edit forwarding-options dhcp-relay dhcpv6 lease-time-validation],
[edit forwarding-options dhcp-relay group group-name lease-time-validation],
[edit forwarding-options dhcp-relay dhcpv6 group group-name lease-time-validation],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server lease-time-validation],
[edit system services dhcp-local-server dhcpv6 lease-time-validation],
[edit system services dhcp-local-server group group-name lease-time-validation],
[edit system services dhcp-local-server dhcpv6 group group-name lease-time-validation]
```

Description

Configure the minimum DHCP lease time allowed in your subscriber access network. If a third-party DHCP server or address pool provides a client lease that is less than the configured threshold, the router performs the action specified by the `violation-action` statement.

Options

seconds Minimum client lease time allowed.

- **Range:** 60 through 2,147,483,647 seconds

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring a DHCP Lease-Time Threshold

lease-time-validation (DHCP Local Server and DHCP Relay Agent)

IN THIS SECTION

- Syntax | 9701
- Hierarchy Level | 9701
- Description | 9702
- Required Privilege Level | 9702
- Release Information | 9702

Syntax

```
lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name ...],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name]
```

Description

Enable the lease-time validation feature on the router. You can then configure the lease-time threshold and an optional action to take when a lease-time violation occurs.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring a DHCP Lease-Time Threshold

leasequery (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 9703
- Hierarchy Level | 9703
- Description | 9703

- Options | 9703
- Required Privilege Level | 9704
- Release Information | 9704

Syntax

```
leasequery {
    attempts number-of-attempts;
    timeout seconds;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name ...],
[edit routing-instances routing-instance-name ...]
```

Description

Enable support for leasequery on a DHCPv4 or DHCPv6 relay agent. You can also configure parameters that the DHCP relay agent uses when sending DHCP leasequery messages to obtain lease information from the DHCP local servers in the logical system/routing instance.

> **NOTE**: You must also configure support on the relevant DHCP local servers with the "allow-leasequery" on page 2191 statement.

Options

attempts *number-of-attempts*

Specify the number of times the DHCP relay agent attempts to send DHCP leasequery messages to the configured DHCP servers in the logical system/routing instance. DHCP relay agent resends the query message if the configured `timeout` value is reached, and

either a confirmed reply or a reply from all configured DHCP servers has not been received. DHCP relay agent sends the subsequent messages only to the DHCP servers that have not replied to previous queries.

- **Range:** 1 through 10
- **Default:** 6

timeout *seconds*

Specify the number of seconds that DHCP relay agent waits before resending a leasequery message to the DHCP servers when all servers have not responded to a previous message.

- **Range:** 1 through 10
- **Default:** 10

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring and Using DHCP Individual Leasequery

least-fill

IN THIS SECTION

- See | 9705

See

"random" on page 14043

leave-timer (MVRP)

IN THIS SECTION

- Syntax | 9705
- EX Series, QFX Series, QFabric | 9705
- M Series, SRX Series, MX Series, T Series | 9705
- Description | 9706
- Default | 9706
- Options | 9706
- Required Privilege Level | 9706
- Release Information | 9706

Syntax

```
leave-timer milliseconds;
```

EX Series, QFX Series, QFabric

```
[edit protocols mvrp interface (all | interface-name)]
```

M Series, SRX Series, MX Series, T Series

```
[edit logical-systems logical-system-name protocols mvrp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvrp
interface (all | interface-name)] (for virtual switch instance type),
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
```

```
mvrp] (for virtual switch instance type),
[edit logical-systems logical-system-name protocols mvrp interface (all | interface-name)],
```

```
[edit protocols mvrp interface (all | interface-name)],
```

```
[edit routing-instances routing-instance-name protocols mvrp] (for virtual switch instance type),
[edit routing-instances routing-instance-name protocols mvrp interface (all | interface-name)]
(for virtual switch instance type)
```

Description

For Multiple VLAN Registration Protocol (MVRP), configure the number of milliseconds the switch retains a VLAN in the Leave state before the VLAN is unregistered. If the interface receives a join message before this timer expires, the VLAN remains registered.

Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Default

1000 milliseconds

Options

milliseconds—Interval that the switch retains a VLAN in the Leave state before the VLAN is unregistered. At a minimum, set the `leave-timer` interval at twice the join-timer interval (range from 300 milliseconds through 1000 milliseconds). Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers

Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches

Configuring Multiple VLAN Registration Protocol (MVRP) on Switches

Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration

Verifying That MVRP Is Working Correctly

Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

join-timer (MVRP) | **9434**

leaveall-timer (MVRP) | **9707**

point-to-point (MVRP) | **12620**

registration | **14268**

leaveall-timer (MVRP)

IN THIS SECTION

- Syntax | **9707**
- EX Series and QFX Series | **9708**
- SRX 1500, SRX 300, SRX 550M, SRX 345, SRX 340, SRX 320 | **9708**
- EX Series, M Series, SRX Series, T Series, MX Series | **9708**
- Description | **9708**
- Options | **9709**
- Required Privilege Level | **9709**
- Release Information | **9709**

Syntax

```
leaveall-timer interval;
```

EX Series and QFX Series

- For platforms with ELS:

```
[edit protocols mvrp],
[edit protocols mvrp interface interface-name]
```

- For platforms without ELS:

```
[edit protocols mvrp interface (all | interface-name)]
```

SRX 1500, SRX 300, SRX 550M, SRX 345, SRX 340, SRX 320

```
[edit routing-instances routing-instance-name protocols mvrp] (for virtual switch instance type),
[edit routing-instances routing-instance-name protocols mvrp interface (all | interface-name)]
(for virtual switch instance type)
```

EX Series, M Series, SRX Series, T Series, MX Series

```
[edit logical-systems logical-system-name protocols mvrp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvrp
interface (all | interface-name)] (for virtual switch instance type),
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
mvrp] (for virtual switch instance type),
[edit logical-systems logical-system-name protocols mvrp interface (all | interface-name)],
[edit protocols mvrp interface (all | interface-name)],
[edit routing-instances routing-instance-name protocols mvrp] (for virtual switch instance type),
[edit routing-instances routing-instance-name protocols mvrp interface (all | interface-name)]
(for virtual switch instance type)
```

Description

For Multiple VLAN Registration Protocol (MVRP), configure the interval at which the LeaveAll state operates on the interface.

Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP. However, if you

choose to change the default values, keep in mind that on an EX Series switch that uses Junos OS with support for ELS, if the timer value set on an interface level is different from the value set on a switch level, then the value on the interface level takes precedence.

Options

EX Series and QFX Series:

interval—Number of seconds or milliseconds between the sending of Leave All messages.

- **Default:** 10 seconds, or 10,000 milliseconds

SRX 1500, SRX 300, SRX 550M, SRX 345, SRX 340, SRX 320:

seconds—Interval between the sending of Leave All messages (range from 10 seconds through 60 seconds. Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

- **Default:** 60 seconds

EX Series, M Series, SRX Series, T Series, MX Series:

milliseconds—Interval between the sending of Leave All messages. Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

- **Default:** 10000 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Hierarchy level `[edit protocols mvrp]` introduced in Junos OS Release 13.2X50-D10 (ELS). (See *Using the Enhanced Layer 2 Software CLI* for information about ELS.)

RELATED DOCUMENTATION

> *Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration*
>
> *Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration*

Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration

Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches

Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers

Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches with ELS Support

Configuring Multiple VLAN Registration Protocol (MVRP) on Switches

Verifying That MVRP Is Working Correctly

join-timer (MVRP) | **9434**

leave-timer (MVRP) | **9705**

point-to-point (MVRP) | **12620**

registration | **14268**

led-beacon

IN THIS SECTION

- Syntax | **9710**
- Hierarchy Level | **9710**
- Description | **9711**
- Required Privilege Level | **9711**
- Release Information | **9711**

Syntax

```
led-beacon
```

Hierarchy Level

```
[edit interfaces interface-name (with port number)]
```

Description

This command causes the LED for the specified port to flash green. You can use the command to physically locate a specific optic port on the PIC.

> **NOTE**: At the [`edit interfaces` *`interface-name (with port number)`*] hierarchy level, you must include the port number as part of the interface name. For example, `et-x/y/z(:n)`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F3 and 16.1R2.

RELATED DOCUMENTATION

show interfaces detail | 20521

legacy-redirect-ip-action

IN THIS SECTION

- Syntax | 9712
- Hierarchy Level | 9712
- Description | 9712
- Options | 9712
- Required Privilege Level | 9713
- Release Information | 9713

Syntax

```
legacy-redirect-ip-action {
    send;
    receive;
}
```

Hierarchy Level

```
[edit protocols bgp group group-name family (inet | inet-vpn) flow],
[edit protocols bgp group group-name neighbor address family (inet | inet-vpn) flow],
[edit routing-instances routing-instance-name protocols bgp group group-name family (inet | inet-
vpn) flow],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (inet | inet-vpn) flow]
```

Description

Configure the legacy redirect to IP action for flow specification routes to provide traffic filtering options for DDos mitigation in service provider networks. You can choose to accept legacy redirect to IP action as specified in the Internet draft draft-ietf-idr-flowspec-redirect-ip-00.txt, *BGP Flow-Spec Extended Community for Traffic Redirect to IP Next Hop,* . You can also configure BGP to advertise the redirect to IP action as a legacy redirect attribute.

> **NOTE**: If legacy encoding configuration is modified, then use the `clear bgp neighbor soft` command to reevaluate the routes and for legacy encoding to take effect.

Options

send Advertise the redirect action as legacy redirect attribute. Specify this option to encode redirect to IP action as flow spec redirect to IP next-hop attribute and advertise the next-hop attribute with the redirect address.

receive Accept legacy encoded redirect-to-ip action attribute. The legacy encoded redirect to IP action is ignored. Specify this option to accept and process the legacy encoded redirect to IP and to generate the redirect-to-ip community for sending to peers that support only new encoding of redirect to IP action.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

- discard-action-for-unresolved-redir-addr | 5818
- *Configuring BGP Flow Specification Action Redirect to IP to Filter DDoS Traffic*
- *Understanding BGP Flow Routes for Traffic Filtering*

less-than

IN THIS SECTION

- Syntax | 9713
- Hierarchy Level | 9714
- Description | 9714
- Options | 9714
- Required Privilege Level | 9714
- Release Information | 9714

Syntax

```
less-than value;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format integer],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format time]
```

Description

Specify a value for the custom AVP attribute below which the incoming RADIUS request from the subscriber must match.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

Options

value Value that the attribute must be less than.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

level (Connectivity Fault Management)

IN THIS SECTION

- Syntax | 9715
- Hierarchy Level | 9715
- Description | 9715
- Options | 9715
- Required Privilege Level | 9716
- Release Information | 9716

Syntax

```
level number;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]
```

Description

A number used in connectivity fault management (CFM) messages to identify the maintenance association. The number is embedded in each of the CFM frames. CFM messages within a given level are processed by maintenance end points (MEPs) at the same level. For example, the operator domain can be level 0, the provider domain can be level 3, and the customer domain can be level 7.

Options

number—A number used to identify the maintenance domain to which the CFM message belongs.

- **Range:** 0 through 7

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Creating a Maintenance Domain

level (OAM)

IN THIS SECTION

- Syntax | 9716
- Hierarchy Level | 9717
- Description | 9717
- Default | 9717
- Options | 9717
- Required Privilege Level | 9717
- Release Information | 9717

Syntax

```
level level;
```

Hierarchy Level

```
[edit ethernet-switching-options storm-control interface (all | interface-name)]
```

Description

For interfaces that are enabled for storm control, configure the storm control level as a percentage of the combined traffic streams that are subject to storm control on that interface.

Default

When storm control is enabled on an interface, the default storm control level is 80 percent of the combined traffic streams that are subject to storm control on that interface.

Options

level—Percentage of the combined traffic streams that are subject to storm control on that interface.

- **Range:** 0 through 100 percent
- **Default:** 80 percent

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 9.1 .

Statement deprecated in JUNOS Release 9.5 for EX Series switches.

Statement reinstated in JUNOS Release 11.4 for EX Series switches.

RELATED DOCUMENTATION

bandwidth (storm control) | **2833**

Configuring Autorecovery for Port Security Events

Understanding Storm Control

level (Disk Partition)

IN THIS SECTION

- Syntax | 9718
- Hierarchy Level | 9718
- Description | 9718
- Description | 9718
- Options | 9719
- Required Privilege Level | 9719
- Release Information | 9719
- RELATED INFORMATION | 9719

Syntax

```
level (full | high)
{
  free-space threshold-value (mb | percent);
}
```

Hierarchy Level

```
[edit chassis disk-partition partition-name]
```

Description

Set the level for monitoring disk usage for the selected disk partition.

Description

Select the desired threshold state for triggering the alarm.

Options

state — Select the desired threshold state for triggering the alarm. Note: You cannot monitor more than one level on a single partition.

- **full** — Set the threshold for indicating that the disk usage is full. Full disk usage is indicated when the free space remaining for the partition is 10%, provided no value has been explicitly specified for the free space remaining.
- **high** — Set the threshold for indicating that the disk usage is high. High disk usage is indicated when the free space remaining for the partition is 25%, provided no value has been explicitly specified for the free space remaining.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3 for EX Series switches.

RELATED INFORMATION

- request system partition hard-disk
- show system storage partitions

level (App Engine Virtual Machine Management Service)

IN THIS SECTION

- Syntax | 9720
- Hierarchy Level | 9720
- Description | 9720
- Default | 9720

- Options | 9720
- Required Privilege Level | 9721
- Release Information | 9721

Syntax

```
level (all | error| info | notice | verbose | warning)
```

Hierarchy Level

```
[edit system processes app-engine-virtual-machine-management-service traceoptions]
```

Description

Set level of debugging output.

Default

info

Options

all	Match all levels.
error	Match error conditions.
info	Match informational messages.
notice	Match conditions that must be handled specially.
verbose	Match verbose messages.
warning	Match warning messages.

Required Privilege Level

system-control

Release Information

Statement introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

traceoptions (App Engine Virtual Machine Management Service) | 29497
flag (App Engine Virtual Machine Management Service) | 7204
file (App Engine Virtual Machine Management Service) | 7030

level (CFM MEP)

IN THIS SECTION

- Syntax | 9721
- Hierarchy Level | 9722
- Description | 9722
- Options | 9722
- Required Privilege Level | 9722
- Release Information | 9722

Syntax

```
level value;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
action log-and-generate-ais]
```

Description

Configure the server maintenance domain level for the MEP.

Options

value Maintenance domain level.

- **Range:** 1 through 7

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Alarm Indication Signal (ETH-AIS) Function Overview

Configuring ETH-AIS on a CFM MEP

level (EX Series Switch Only)

IN THIS SECTION

Syntax | 9723

- Hierarchy Level | 9723
- Description | 9723
- Options | 9723
- Required Privilege Level | 9723
- Release Information | 9723

Syntax

```
level number;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]
```

Description

Configure a number to be used in CFM messages to identify the maintenance association.

Options

number—Number used to identify the maintenance domain to which the CFM message belongs.

- **Range:** 0 through 7

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)

Junos OS Network Interfaces Configuration Guide

level (Global IS-IS)

IN THIS SECTION

- Syntax | 9724
- Hierarchy Level | 9725
- Description | 9725
- Options | 9725
- Required Privilege Level | 9728
- Release Information | 9728

Syntax

```
level level-number {
    authentication-key key;
    authentication-key-chain (Protocols IS-IS) key-chain-name;
    authentication-type type;
    disable;
    external-preference preference;
    no-authentication-check;
    no-csnp-authentication;
    no-hello-authentication;
    no-psnp-authentication;
    preference preference;
    purge-originator {
        empty|self;
    }
    wide-metrics-only;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Configure the global-level properties.

You can administratively divide a single AS into smaller groups called areas. You configure each routing device interface to be in an area. Any interface can be in any area. The area address applies to the entire routing device. You cannot specify one interface to be in one area and another interface in a different area. To route between areas, you must have two adjacent Level 2 routers that communicate with each other.

Level 1 routers can only route within their IS-IS area. To send traffic outside their area, Level 1 routers must send packets to the nearest intra-area Level 2 router. A routing device can be a Level 1 router, a Level 2 router, or both. You specify the router level on a per-interface basis, and a routing device becomes adjacent to other routing devices on the same level on that link only.

You can configure one Level 1 routing process and one Level 2 routing process on each interface, and you can configure the two levels differently.

Options

level-number The IS-IS device level number, which can be 1 or 2. The default is for the device to operate as both a Level 1 and 2 device.

authentication-key *key* Authentication key (password). Neighboring devices use the password to verify the authenticity of packets sent from this interface. For the key to work, you also must include the `authentication-type` statement. All devices must use the same password. If you are using the Junos OS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces adjacent to the Juniper Networks device.

> **NOTE**: If you do not include this statement and the `authentication-type` statement, IS-IS authentication is disabled.

- *key*—Authentication password. The password can be up to 1024 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

> **CAUTION**: A simple password for authentication is truncated if it exceeds 254 characters.

authentication-key-chain *key-chain-name*

Apply and enable an authentication key-chain to the routing device. The variable *key-chain-name* is the authentication key-chain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

authentication-type *authentication*

Enable authentication and specify the authentication scheme for IS-IS. If you enable authentication, you must specify a password by including the `authentication-key` statement. If you do not include this statement and the `authentication-key` statement, IS-IS authentication is disabled. Following are the valid values for *authentication*:

-
- `md5`—Use HMAC authentication in combination with MD5. HMAC-MD5 authentication is defined in RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*.
- `simple`—Use a simple password for authentication. The password is included in the transmitted packet, making this method of authentication relatively insecure.

> **NOTE**: Juniper Networks does not recommend using the `simple` authentication method, because it is not as secure as MD5.

disable

Disable IS-IS on the routing device, on an interface, or on a level. Enabling IS-IS on an interface (by including the `interface` statement at the `[edit protocols isis]` or the `[edit routing-instances` *`routing-instance-name`* `protocols isis]` hierarchy level), disabling it (by including the `disable` statement), and not actually having IS-IS run on an interface (by including the `passive` statement) are mutually exclusive states. IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which `family iso` is enabled.

external-preference *preference*

Configure the preference of external routes. The value for *preference* can be 0 through 4,294,967,295 (2^{32} – 1). The default is 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes).

no-authentication-check
Generate authenticated packets and check the authentication on received packets, but do not reject packets that cannot be authenticated.

no-csnp-authentication
Suppress authentication check on complete sequence number PDU (CSNP) packets.

no-hello-authentication
Suppress authentication check on complete sequence number hello packets.

no-psnp-authentication
Suppress authentication check on partial sequence number PDU (PSNP) packets.

preference *preference*
Configure the preference of internal routes. Route preferences (also known as administrative distances) are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected. To change the preference values, include the `preference` statement (for internal routes) or the `external-preference` statement. The value for *preference* can be 0 through 4,294,967,295 (2^{32} – 1). The default is 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes).

purge-originator (empty | self)
Enable purge originator identification (POI) by adding the type, length and value (TLV) with the Intermediate System (IS) identification to the LSPs that do not contain POI information. If an IS generates a purge, Junos adds this TLV with the system ID of the IS to the purge. If an IS receives a purge that does not include this TLV, it adds this TLV with both its own system ID and the system ID of the IS from which it received the purge. This allows the IS that receives this purge to log the system ID of the originator, or the upstream source of the purge and makes it easier to locate the origin of the purge.

- `empty`—(Optional) Add POI to a purge received from an IS that does not contain POI information.
- `self`—(Optional) (Optional) Add POI only to self-originated purge LSPs.

wide-metrics-only
Configure IS-IS to generate metric values greater than 63 on a per IS-IS level basis. Normally, IS-IS metrics can have values up to 63, and IS-IS generates two type, length, and value (TLV) tuples, one for an IS-IS adjacency and the second for an IP prefix. To allow IS-IS to support traffic engineering, a second pair of TLVs has been added to IS-IS, one for IP prefixes and the second for IS-IS adjacency and traffic engineering information. With these TLVs, IS-IS metrics can have values up to 16,777,215 (2^{24} – 1).

To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, include the `wide-metrics-only` statement.

> **NOTE**: By default, Junos OS supports the sending and receiving of wide metrics. Junos OS allows a maximum metric value of 63 and generates both pairs of TLVs.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Understanding IS-IS Configuration*
- *Example: Configuring IS-IS*
- *Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups*
- *Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding*

level (IS-IS Interfaces)

IN THIS SECTION

- Syntax | 9729
- Hierarchy Level | 9730
- Description | 9730
- Options | 9730
- Required Privilege Level | 9733
- Release Information | 9733

Syntax

```
level level-number {

    application-specific {
       attribute-group name {
          admin-group name;
          advertise-delay-metric
          application flex-algorithm;
          srlg name1, name2;
          te-metric metric;
       }
          }
    hello-authentication-key key;
    hello-authentication-key-chain key-chain-name;
    hello-authentication-type authentication;
    hello-interval seconds;
    hold-time seconds;
    ipv4-multicast-metric number;
    ipv6-multicast-metric number;
    ipv6-unicast-metric number;
    metric metric;
    passive;
    priority number;
    srv6-adjacency-segment{
        protected | unprotected
        locator name {
            end-x-sid sid | dynamic-end-x-sid
                flavor {
                    psp;
                    usp;
                    usd;
                }
            micro-adjacency-sid}
            }
        }

    }
    te-metric metric;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis
interface interface-name],
[edit protocols isis interface interface-name],
[edit routing-instances routing-instance-name protocols isis interface interface-name]
```

Description

Configure the IS-IS level. You can configure one instance of Level 1 routing and one instance of Level 2 routing on each interface, and you can configure the two levels differently.

Options

level-number
The IS-IS level number, which can be 1 or 2. The default is for the device to operate as both a Level 1 and 2 device.

application-specific
Configure and advertise application (flexible algorithm) specific TE attributes such as delay-metric, te-metric, srlg or admin-groups.

hello-authentication-key *password*
Configure an authentication key (password) for hello packets. Neighboring routing devices use the password to verify the authenticity of packets sent from an interface. For the key to work, you also must include the `hello-authentication-type` statement. By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password. The authentication *password* can be up to 255 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

hello-authentication-key-chain *key-chain-name*
Apply an authentication keychain to the IS-IS interface. The authentication keychain name *key-chain-name* can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

hello-authentication-type (md5 | simple)
Enable authentication on an interface for hello packets. If you enable authentication on hello packets, you must specify a password by including the `hello-authentication-key` statement. You can configure authentication for a given IS-IS level on an interface. On a point-to-point link, if you enable hello authentication for both IS-IS levels, the password configured for Level 1 is used for both levels.

> **CAUTION**: If no authentication is configured for Level 1 on a point-to-point link with both levels enabled, the hello packets are sent without any password, regardless of the Level 2 authentication configurations.

By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.

- `md5`— Specifies Message Digest 5 as the packet verification type.
- `simple`—Specifies simple authentication as the packet verification type.

hello-interval *seconds*

Modify the frequency with which the routing device sends hello packets out of an interface, in seconds. Routing devices send hello packets at a fixed interval on all interfaces to establish and maintain neighbor relationships. This interval is advertised in the hello interval field in the hello packet. You can send out hello packets in subsecond intervals. To send out hello packets every 333 milliseconds, set the hold-time value to `1`.

The valid range is from 1 to 20,000 seconds. The default is 3 seconds (for designated intermediate system [DIS] routers), 9 seconds (for non-DIS routers).

> **NOTE**: The default hello intervals on EX 9200 switches are not sufficient to avoid protocol flapping during an ISSU Linux upgrade. To avoid protocol flapping during the upgrade, configure the hello intervals for 35 seconds or more.

> **NOTE**: When elected as a designated intermediate system [DIS] router on any LAN adjacency, the hello and hold-timer intervals are scaled down by a factor of 3. This means the default values of 9 and 27 seconds of ISIS hello and hold time intervals are scaled down to 3 and 9 seconds for LAN adjacencies. During switchovers, this hello interval is too short to form LAN adjacencies. Therefore, you can configure one of the following solutions:
>
> - Set the hello and hold time interval for LAN adjacencies to 30 seconds and 90 seconds respectively on both the DIS and a neighboring router.

- Convert the LAN interfaces to point-to-point IS-IS interfaces.

On QFX10000 switches, we strongly recommend that you configure all IS-IS interfaces, including peer interfaces, as point-to-point interfaces. If you do not, you might experience session flaps, that is, IS-IS sessions that go down and then come back up, when IS-IS is configured in virtual routing and forwarding (VRF) instances. When you scale IS-IS in any scenario, you might also experience scaling issues if you do not configure IS-IS interfaces as point-to-point interfaces.

ipv4-multicast-metric *metric*

Specify the IPv4 multicast topology metric value for the level. The value for *metric* can range from 0 through 16,777,215.

ipv6-multicast-metric *metric*

Specify the IPv6 alternate multicast topology metric value for the level. The value for *metric* can range from 0 through 16,777,215.

ipv6-unicast-metric *metric*

Specify the IPv6 unicast topology metric value for the level. The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. The value for *metric* can range from 0 through 16,777,215.

priority *number*

Configure the interface's priority for becoming the designated router. The interface with the highest priority value becomes that level's designated router. The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.

A routing device advertises its priority to become a designated router in its hello packets. On all multiaccess networks, IS-IS uses the advertised priorities to elect a designated router for the network. This routing device is responsible for sending network link-state advertisements, which describe all the routing devices attached to the network. These advertisements are flooded throughout a single area. A routing device's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127. Routing devices with a higher value are more likely to become the designated router. The default priority value is 64.

srv6-adjacency-segment

Configure the SRv6 adjacency segment options.

protected | unprotected

Enable or disable TI-LFA backup protection.

locator

Specify the address to which the end-x-sid function is assigned. Ensure that the locator address and interface address belong to the same prefix to avoid unexpected network behavior.

end-x-sid *sid*

Specify the function for the SRv6 adjacency.

dynamic-end-x-sid You can configure this option in addition to the existing `end-x-sid`. IS-IS dynamically allocates one end-x-sid address for a locator on an interface. Note that you cannot configure this option for micro SIDs.

flavor Specify the behavior of the end-x-sid function such as penultimate segment pop (psp), ultimate segment pop (usp) or ultimate segment decapsulation (usd).

micro-adjacency-sid Configure micro adjacency parameters. If you do not specify a SID, IS-IS allocates it from the dynamic pool and the size is as per the defaul block configuration. Configure a value that is compatible with the default locator parameters.

te-metric *metric* Set the metric value used by traffic engineering for information injected into the traffic engineering database. The value of the traffic engineering metric does not affect normal IS-IS forwarding. When traffic engineering is enabled on the routing device, you can use this statement to configure an IS-IS metric that is used exclusively for traffic engineering. The value for metric can range from 1 through 16,777,215. The default is the value of the IGP metric.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding IS-IS Configuration

Example: Configuring IS-IS

Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups

Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding

level (Protocols RIFT)

IN THIS SECTION

- Syntax | 9734
- Hierarchy Level | 9734
- Description | 9734
- Options | 9735
- Required Privilege Level | 9735
- Release Information | 9735

Syntax

```
level {
    auto;
    configured-value value;
    leaf;
    top-of-fabric;
}
```

Hierarchy Level

```
[edit protocols rift]
```

Description

Configure the level to identify the role of node in the RIFT topology.

If you do not specify any options under the `level` statement, the default option is `auto` level.

Options

auto	(Optional) Enable zero touch provisioning to identify the level of the node automatically.
configured-value *value*	(Optional) Enter the configured value of the level. • **Range:** 1 through 23
leaf	(Optional) Identify node as leaf in the RIFT topology. Automatically, this is level 0.
top-of-fabric	(Optional) Identify node as top-of-fabric in the RIFT topology. In automatic zero touch provisioning, this is level 24.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

level (Services)

IN THIS SECTION

- Syntax | 9736
- Hierarchy Level | 9736
- Description | 9736

- Options | 9736
- Required Privilege Level | 9736
- Release Information | 9736

Syntax

```
level [brief | detail | extensive | verbose];
```

Hierarchy Level

```
[edit services ssl traceoptions]
```

Description

Specify the level of debugging the output. This statement is supported on the SRX550M, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall.

Options

- *brief*—Specify brief debugging output.
- *detail*—Specify detailed debugging output.
- *extensive*—Specify extensive debugging output.
- *verbose*—Specify verbose debugging output.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

license

IN THIS SECTION

- Syntax | 9737
- Hierarchy Level | 9738
- Description | 9738
- Options | 9738
- Required Privilege Level | 9739
- Release Information | 9739

Syntax

```
license {
    autoupdate {
        url url <password password>;
    }
    keys {
        key key
    }
    renew {
        before-expiration number;
        interval interval-hours;
    }
    traceoptions {
         file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
```

```
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Specify license information for the device.

Options

autoupdate
Autoupdate license keys from license servers.

before-expiration *number*
License renewal lead time before expiration, in days.

- **Range:** 0 through 60 days

interval *interval-hours*
License checking interval, in hours.

- **Range:** 1 through 336 hours

keys key *key*
Configure one or more license keys. For example,

```
[edit]
user@device# set system license keys key "key_1"
user@device# set system license keys key "key_2"
user@device# set system license keys key "key_3"
user@device# set system license keys key "key_4"
user@device# commit
commit complete
```

renew
License renewal lead time and checking interval.

url
URL of a license server.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Options `keys` introduced in Junos OS Release 14.1X53-D10.

license-mode

IN THIS SECTION

- Syntax | 9739
- Hierarchy Level | 9739
- Description | 9740
- Options | 9740
- Required Privilege Level | 9740
- Release Information | 9740

Syntax

```
license-mode IR | R
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

Configures the FPC's license mode. This can be used to track the number of PICs configured for a specific mode.

Options

IR Sets the PIC to LSR mode.

R Sets the PIC to full IP mode.

> **NOTE**: Starting in Junos OS Release 16.1R3 for PTX Series routers, the `IR` and `R` options are used. For previous releases (starting in Junos OS Release 15.1F3) the `Ip` (full IP mode) and `lsr` (LSR mode) options are used.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F3.

RELATED DOCUMENTATION

fpc (M320, T320 and T640 Devices) | **7516**

license-server

IN THIS SECTION

- Syntax | **9741**
- Hierarchy Level | **9741**

- Description | 9741
- Options | 9741
- Required Privilege Level | 9742
- Release Information | 9742

Syntax

```
license-server {
    ip-address address;
    log-interval seconds;
    services (jflow | cgnat | firewall);
}
```

Hierarchy Level

```
[edit]
```

Description

On MX Series routers with MS-MICs and MS-MPCs, configure the capability to transmit the throughput details per service for the Junos Address Aware, Junos Traffic Vision, and Junos Network Secure services in the last time interval to an external log collector.

Options

ip-address *address*
: Use the specified IP address of the license log server.

log-interval *seconds*
: Use the specified frequency at which throughput data must be sent from the router to the log collector.

- **Range:** 60 through 86,400 seconds

services
: Specify the services for which throughput data must be exported.

- `jflow`—Use inline flow monitoring service or Junos Traffic Vision.
- `cgnat`—Use carrier-grade NAT service or Junos Address Aware.
- `firewall`—Use stateful firewall or Junos Network Secure.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

> *License Server Management for Throughput Data Export on MX Series Routers for NAT, Firewall, and Inline Flow Monitoring Services*

lie-receive-address (Protocols RIFT)

IN THIS SECTION

- Syntax | **9743**
- Hierarchy Level | **9743**
- Description | **9743**
- Options | **9743**
- Required Privilege Level | **9743**
- Release Information | **9743**

Syntax

```
lie-receive-address {
    family {
        (inet ipv4-address | inet6 ipv6-address);
    }
}
```

Hierarchy Level

```
[edit protocols rift]
```

Description

Configure the address on which the link information elements (LIEs) should be received.

Options

family Specify the family type for the RIFT protocol interface on which the LIEs should be received.

inet *ipv4-address* Specify the IPv4 IP address of the RIFT protocol interface on which the LIEs should be recieved.

inet6 *ipv6-address* Specify the IPv6 IP address of the RIFT protocol interface on which the LIEs should be recieved.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

lie-transmit-address (Protocols RIFT)

IN THIS SECTION

- Syntax | 9744
- Hierarchy Level | 9744
- Description | 9744
- Options | 9745
- Required Privilege Level | 9745
- Release Information | 9745

Syntax

```
lie-transmit-address {
    family {
        (inet ipv4-address | inet6 ipv6-address);
    }
}
```

Hierarchy Level

```
[edit protocols rift interface interface-name]
```

Description

Configure the IPv4 or IPv6 IP address on which the link information elements (LIEs) should be sent to discover neighbors on the other side of the link.

Options

family	Specify the family type for the RIFT protocol interface on which the LIEs should be sent.
inet *ipv4-address*	Specify the IPv4 IP address of the RIFT protocol interface on which the LIEs should be sent.
inet6 *ipv6-address*	Specify the IPv6 IP address of the RIFT protocol interface on which the LIEs should be sent.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

interface (Protocols RIFT) | **8994**

rift | **15794**

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

lifetime (Router Advertisement)

IN THIS SECTION

- Syntax | **9746**
- Hierarchy Level | **9746**
- Description | **9746**
- Options | **9746**

Required Privilege Level | 9746
Release Information | 9747

Syntax

```
lifetime seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name
dns-server-address address],
[edit protocols router-advertisement interface interface-name dns-server-address address]
```

Description

Specify the maximum time in seconds for which the recursive DNS server address remains valid. The device can use the specified recursive DNS server address for DNS name resolution until the time specified by this statement.

seconds– Maximum time for which the recursive DNS server address remains valid.

Options

- **Range:** 0 through 4294967295 seconds
- **Default:** 1800 seconds
- **Values:** 0 indicates that the advertised recursive DNS server address is no longer valid and that this recursive DNS server address entry can be deleted.

  4294967295 seconds indicates an infinite lifetime and a persistent entry in the device for this recursive DNS server address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Recursive DNS Servers for IPv6

Example: Configuring Recursive DNS Server Addresses for IPv6 Hosts

dns-server-address | 5864

lifetime (Router Discovery)

IN THIS SECTION

- Syntax | 9747
- Hierarchy Level | 9747
- Description | 9748
- Options | 9748
- Required Privilege Level | 9748
- Release Information | 9748

Syntax

```
lifetime seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery interface interface-name],
[edit protocols router-discovery interface interface-name]
```

Description

Specify how long the addresses sent by the server in its router advertisement packets are valid. This time must be long enough so that another router advertisement packet is sent before the lifetime has expired. The lifetime value is placed in the advertisement lifetime field of the router advertisement packet. If this amount of time passes and the host has not received a router advertisement from the server, the router marks the advertised addresses as invalid.

Options

seconds—Lifetime value. A value of 0 indicates that one or more addresses are no longer valid.

- **Range:** Three times the value set by the `max-advertisement-interval` statement through 2 hours, 30 minutes (9000 seconds)
- **Default:** 1800 seconds (30 minutes, which is three times the default value for the `max-advertisement-interval` statement)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

max-advertisement-interval (Protocols Router Discovery) | **10416**

lifetime-seconds

IN THIS SECTION

- Syntax | 9749
- Hierarchy Level | 9749
- Description | 9749
- Options | 9749
- Required Privilege Level | 9750
- Release Information | 9750

Syntax

```
lifetime-seconds seconds;
```

Hierarchy Level

```
[edit services ipsec-vpn ike proposal proposal-name],
[edit services ipsec-vpn ipsec proposal proposal-name]
```

Description

Configure the lifetime of an IKE or IPsec SA. This statement is optional.

Options

seconds—Lifetime

- **Default:** 3600 seconds (IKE); 28,800 seconds (IPsec)
- **Range:** 180 through 86,400

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IKE Proposals

Configuring IPsec Proposals

Configuring Security Associations

lifetime-seconds (Security)

IN THIS SECTION

- Syntax | 9750
- Hierarchy Level | 9751
- Description | 9751
- Options | 9751
- Required Privilege Level | 9751
- Release Information | 9751

Syntax

```
<lifetime-seconds seconds>;
```

Hierarchy Level

```
[edit security ike proposal ike-proposal-name],
[edit security ipsec proposal ipsec-proposal-name]
```

Description

(Optional) Configure the lifetime of IKE or IPsec SA. When the SA expires, it is replaced by a new SA (and SPI) or terminated.

Options

seconds—Lifetime, in seconds.

- **Range:** 180 through 86,400

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Lifetime for an IKE SA

Configuring the Lifetime for an IPsec SA

light

IN THIS SECTION

- Junos OS Syntax | 9752
- Junos OS Evolved Syntax | 9752
- Junos OS Hierarchy Level | 9752
- Junos OS Evolved Hierarchy Level | 9753
- Description | 9753
- Options | 9753
- Required Privilege Level | 9754
- Release Information | 9754

Junos OS Syntax

```
light {
    port number;
}
```

Junos OS Evolved Syntax

```
light {
    offload-type (none | inline-timestamp | pfe-timestamp);
    port number;
    stateful-sequence cache-size size;
}
```

Junos OS Hierarchy Level

```
[edit services rpm twamp server]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring twamp server]
```

Description

Enable the Two-Way Active Measurement Protocol (TWAMP) server for light control on the UDP port that reflects the test packets. If you want the default IANA port for TWAMP (862), you do not need to configure the port option. To complete the configuration for light control, you also need to configure the `control-type light` statement at either the Junos OS `[edit services rpm twamp client control-connection` *`control-connection-name`*`]` hierarchy level or the Junos OS Evolved `[edit services monitoring twamp client control-connection` *`control-connection-name`*`]` hierarchy level.

You configure the light statement because you want a stateless version of TWAMP where test parameters are predefined instead of negotiated. All test packets received by the server on a test port are reflected back and forgotten right away. If you want to have a stateful TWAMP server for Junos OS Evolved, you should configure the `managed` statement at the [edit services monitoring twamp server] hierarchy level instead. For Junos OS, if you want to have a stateful TWAMP server, do not configure the `light` statement for the TWAMP server, because managed mode is the default. Because TWAMP light servers are stateless, information about them is not included in the output of the `show services rpm twamp server connection` (Junos OS) or the `show services monitoring twamp server control-info` (Junos OS Evolved) operational mode command; only information about managed servers is included.

Options

offload-type (none | inline-timestamp | pfe-timestamp)

Enable timestamping of TWAMP probe messages in the Routing Engine, inline in the hardware, or Packet Forwarding Engine host processor. Prior to Junos OS Evolved 22.3R1, for IPv6 traffic, you must configure `none`, as only Routing Engine timestamping is supported. Starting in Junos OS Evolved 22.3R1, you can configure the `pfe-timestamping` option for IPv6 traffic.

- **none**—Timestamping in the Routing Engine.
- **inline-timestamping**—Timestamping inline, in the hardware at the generator, and reflection in the hardware at the reflector. When this option is configured for an ACX router, the ACX can act as either a server or client, but not both. Also for an ACX router, `inline-timestamping` is not supported on PM50 ports if FEC74 is enabled on that port. You can configure the `inline-timestamping` option for up to 256 test sessions.

- **pfe-timestamping**—Timestamping in the Packet Forwarding Engine.
- **Default:** pfe-timestamp

port *number*

Specify the UDP port that reflects the TWAMP test packets.

- **Range:** You can specify any port from 862 through 65535.
- **Default:** 862 (IANA port for TWAMP)

stateful-sequence cache-size *number-of-entries*

Configure stateful sequencing to generate the sequence number in the reflected payload independently, instead of copying the sequence number from the client frame. With independently generated sequence numbers, the software is able to detect in which direction drops have occurred. Unidirectional drop values are inferred minimum values, and are not guaranteed to be exact. The software relies on an actual packet to be reflected successfully to be able to compute differences. If such a packet does not occur, then the software is unable to tell whether the drop happened in the egress or ingress direction.

Because TWAMP light (just like STAMP) is by nature stateless, stateful sequencing can only be done using a cache system. Currently, the software uses a FIFO system to manage the cache; the entry without an update for the longest time is the first one out to make space for a new entry.

- **Range:**1000-1000000 entries (number of client sessions kept in the cache)
- **Default:** 10000 entries

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 20.3R1.

Support added in Junos OS Release 21.1R1.

`offload-type` and `stateful-sequence` options introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

light-weight-dhcpv6-relay

IN THIS SECTION

- Syntax | 9755
- Hierarchy Level | 9755
- Description | 9755
- Required Privilege Level | 9756
- Release Information | 9756

Syntax

```
lightweight-dhcpv6-relay;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options]
```

Description

Configure a Lightweight DHCPv6 Relay Agent (LDRA) to insert relay agent information in messages sent from a DHCPv6 client to a server or other relay agent on the same IPv6 link. The LDRA acts as a relay agent, but without performing the routing function necessary to forward messages to a server or relay agent that resides on a different IPv6 link.

When the LDRA receives a DHCPv6 Solicit message from a client, it encapsulates that message within a DHCPv6 Relay-Forward message, which it then forwards to the server or to another relay agent. Before it forwards the Relay-Forward message, the LDRA can also insert DHCPv6 options in the message. These options contain information that the server uses to assign IP addresses, prefixes, and other configuration parameters for the client.

You must configure LDRA if you configure the following DHCPv6 options at the [`edit vlan` *`vlan-name`* `forwarding-options dhcp-security dhcpv6-options`] hierarchy level:

- `option-16` (Vendor ID)—Option 16 provides the server with information about the vendor that manufactured the hardware on which the DHCPv6 client is running. Option 16 is the DHCPv6 equivalent of the `vendor-id` suboption of DHCP option 82.
- `option-18` (Interface ID)—A unique identifier for the interface on which the client DHCPv6 packet is received. Suboptions can be configured to include a prefix with the interface ID or to change the type of information used to identify the interface. Option 18 is the DHCPv6 equivalent of the `circuit-id` suboption of DHCP option 82.
- `option-37` (Remote ID)—A unique identifier for the remote host. Suboptions can be configured to include a prefix with the remote ID or to change the interface portion of the ID. Option 37 is the DHCPv6 equivalent of the `remote-id` suboption of DHCP option 82.

> **NOTE**: Option 18 is mandatory in Relay-Forward messages and is included even if it is not explicitly configured. However, suboptions of option 18 are included in Relay-Foward messages only if they are configured using the `option-18` CLI statement at the [`edit vlan` *`vlan-name`* `forwarding-options dhcp-security dhcpv6-options`] hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R3.

RELATED DOCUMENTATION

Using Lightweight DHCPv6 Relay Agent (LDRA)

no-option18 | **11538**

no-dhcpv6-options | **11440**

limit

IN THIS SECTION

- Syntax | 9757
- Hierarchy Level | 9757
- Description | 9757
- Options | 9757
- Required Privilege Level | 9757
- Release Information | 9758

Syntax

```
limit max-sub-sessions;
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name max-data-sessions-
per-subscriber]
```

Description

Specify the limit for the maximum number of subscriber sessions.

Options

max-sub-sessions—Maximum number of subscriber sessions.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

limit (Content Security Policy)

IN THIS SECTION

- Syntax | 9758
- Hierarchy Level | 9758
- Description | 9758
- Required Privilege Level | 9758
- Release Information | 9759

Syntax

```
limit value;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name traffic-options sessions-per-client]
```

Description

In an attempt to consume all available resources and hinder the ability of the device, a malicious user might generate a large amount of traffic all at once. To prevent such activity from succeeding, you can impose a session throttle to limit sessions.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

limit-ports-per-address

IN THIS SECTION

- Syntax | 9759
- Hierarchy Level | 9759
- Description | 9759
- Options | 9760
- Required Privilege Level | 9760
- Release Information | 9760

Syntax

```
limit-ports-per-address number;
```

Hierarchy Level

```
[edit services nat pool nat-pool-name]
```

Description

Specify the limit for number of ports allocated per host (IP address).

Options

number Number of ports allocated per host (IP address).

- **Range:** 2 through 65,435

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Pools of Addresses and Ports for Network Address Translation Overview

limit (Security Log)

IN THIS SECTION

- Syntax | 9761
- Hierarchy Level | 9761
- Description | 9761
- Options | 9761
- Required Privilege Level | 9761
- Release Information | 9761

Syntax

```
limit value;
```

Hierarchy Level

```
[edit security log cache]
[edit logical-systems name security log cache]
[edit tenants tenant-name security log cache]
```

Description

Specify the number of security log entries to be kept in memory.

Options

Once the maximum value limit is reached, new entries will not be added until the cache size drops.

- **Range:** 0 through 4,294,967,295
- **Default:** 10,000 security log entries.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

The [edit `logical-systems` *name* security log cache] and [edit `tenants` *tenant-name* security log cache] hierarchy levels introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

syslog (System) | **28836**

limit (Security SCTP)

IN THIS SECTION

- Syntax | 9762
- Hierarchy Level | 9763
- Description | 9763
- Options | 9763
- Required Privilege Level | 9763
- Release Information | 9763

Syntax

```
limit {
    address name {
        payload-protocol (asap | bicc | ddp-segment | ddp-stream | diameter-dtls | diameter-sctp
| dua | enrp | h248 | h323 | id | iua | m2pa | m2ua | m3ua | others | qipc | reserved | s1ap |
simco | sua | tali | v5ua | x2ap) {
            rate rate;
        }
    }
    payload-protocol (asap | bicc | ddp-segment | ddp-stream | diameter-dtls | diameter-sctp |
dua | enrp | h248 | h323 | id | iua | m2pa | m2ua | m3ua | others | qipc | reserved | s1ap |
simco | sua | tali | v5ua | x2ap) {
        rate rate;
    }
    rate {
        address name {
            sccp sccp;
            ssp ssp;
            sst sst;
        }
        sccp sccp;
        ssp ssp;
        sst sst;
```

```
    }
}
```

Hierarchy Level

```
[edit security sctp profile profile-name]
```

Description

Use the limit option to set the rate limit per association for local Services Processing Unit (SPU) packets. You can configure a GPRS SCTP profile by setting the limit rate parameter and the payload protocol parameter for SCTP inspection.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs sctp profile profile-name]` hierarchy level is replaced by `[edit security sctp profile profile-name]`.

Options

address ***ip-address*** To set Signalling Connection Control Part (SCCP), Subsystem-Prohibited (SSP), and Subsystem Status Test (SST) messages rate limit to an IP address. The IP address can accept either an IPv4 address or an IPv6 address.

sccp ***rate-limit*** To set the SCCP messages rate limit.

ssp ***rate-limit*** To set the SSP messages rate limit.

sst ***rate-limit*** To set the SSP messages rate limit.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2. Statement is modified in Junos OS Release 12.1X46-D10. Support for **address** option accepting both IPv4 and IPv6 formats added in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Understanding Stream Control Transmission Protocol

SCTP Packet Structure Overview

limit-session

IN THIS SECTION

- Syntax | 9764
- Hierarchy Level | 9764
- Description | 9764
- Options | 9765
- Required Privilege Level | 9765
- Release Information | 9765

Syntax

```
limit-session {
    destination-ip-based number;
    source-ip-based number;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name]
```

Description

Limit the number of concurrent sessions the device can initiate from a single source IP address or the number of sessions it can direct to a single destination IP address.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

limit-session (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 9765
- Hierarchy Level | 9767
- Description | 9767
- Required Privilege Level | 9767
- Release Information | 9767

Syntax

```
limit-session {
    by-destination{
        by-protocol {
```

```
            icmp {
                maximum-sessions number;
                packet-rate number;
                session-rate number;
            }
            tcp {
                maximum-sessions number;
                packet-rate number;
                session-rate number;
            }
            udp {
                maximum-sessions number;
                packet-rate number;
                session-rate number;
            }
        }
        maximum-sessions number;
        packet-rate number;
        session-rate number;
    }
    by-source {
        by-protocol {
            icmp {
                maximum-sessions number;
                packet-rate number;
                session-rate number;
            }
            tcp {
                maximum-sessions number;
                packet-rate number;
                session-rate number;
            }
            udp {
                maximum-sessions number;
                packet-rate number;
                session-rate number;
            }
        }
        maximum-sessions number;
        packet-rate number;
        session-rate number;
```

```
    }
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name]
```

Description

Configure session limits for individual destination or source addresses, or for individual destination or source subnets. This protects against network probing attacks and network flooding attacks. You can specify limits for specific protocols (ICMP, TCP, and UDP), or specify limits independent of a protocol. When a session limit is exceeded for a source or destination, packets from the source or to the destination are dropped until the session limit is no longer exceeded.

To specify limits for destination or source subnets rather than individual addresses, include the `aggregations` statement at the `[edit services screen ids-option screen-name]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

limited-ifl-scaling

IN THIS SECTION

- Syntax | 9768
- Hierarchy Level | 9768
- Description | 9768
- Required Privilege Level | 9768
- Release Information | 9768

Syntax

```
limited-ifl-scaling;
```

Hierarchy Level

```
[edit chassis network-services enhanced-ip]
```

Description

Limits the maximum number of logical interfaces on MX Series routers with MS-DPCs to 64,000 for enhanced IP network services mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

Configuring Network Services Mode

Configuring Network Services Mode

Configuring Enhanced IP Network Services for a Virtual Chassis

line-encoding (E1 Options)

IN THIS SECTION

- Syntax | 9769
- Hierarchy Level | 9769
- Description | 9769
- Options | 9770
- Required Privilege Level | 9770
- Release Information | 9770

Syntax

```
line-encoding (ami | hdb3);
```

Hierarchy Level

```
[edit chassis synchronization interfaces external e1-options]
[edit chassis synchronization interfaces (external-0/0 | external-1/0) e1-options]
```

Description

Configure the line encoding format on the E1 interface.

Options

ami Set the line encoding format as automatic mark inversion.

hdb3 Set the line encoding format as high-density bipolar 3 code.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

line-encoding (T1 Options)

IN THIS SECTION

- Syntax | 9771
- Hierarchy Level | 9771
- Description | 9771
- Options | 9771
- Required Privilege Level | 9771
- Release Information | 9771

Syntax

```
line-encoding (ami | b8zs);
```

Hierarchy Level

```
[edit chassis synchronization interfaces external t1-options]
[edit chassis synchronization interfaces (external-0/0 | external-1/0) t1-options]
```

Description

Configure the line encoding format on the T1 interface.

Options

ami Set the line encoding format as automatic mark inversion.

b8zs Set the line encoding format as 8-bit zero suppression.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

line-identity (Dynamic Access-Line-Identifier VLANs)

IN THIS SECTION

- Syntax | 9772
- Hierarchy Level | 9772
- Description | 9772
- Required Privilege Level | 9773
- Release Information | 9773

Syntax

```
line-identity {
    dynamic-profile profile-name;
    include {
        accept-no-ids;
        circuit-id;
        remote-id;

    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces “$junos-interface-ifd-name” unit “$junos-
interface-unit” auto-configure],
[edit interfaces interface-name unit logical-unit-number auto-configure]
```

Description

Configure the access-line identifier received in DHCP or PPPoE discovery packets as a trusted option that is accepted for dynamically creating a VLAN on a static or dynamic underlying VLAN interface according to the specified dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 17.1.

RELATED DOCUMENTATION

Configuring Dynamic Underlying VLAN Interfaces to Use Access-Line Identifiers

Configuring Dynamic VLAN Subscriber Interfaces Based on Access-Line Identifiers

line-interval (Dampen Options)

IN THIS SECTION

- Syntax | 9773
- Hierarchy Level | 9774
- Description | 9774
- Options | 9774
- Required Privilege Level | 9774
- Release Information | 9774

Syntax

```
line-interval line-interval;
```

Hierarchy Level

```
[edit event-options event-script dampen dampen-options],
[edit event-options event-script file filename dampen dampen-options],
[edit system scripts commit dampen dampen-options],
[edit system scripts op dampen dampen-options],
[edit system scripts op file filename dampen dampen-options]
```

Description

Specify the number of lines of scripts to execute before pausing.

Script dampening is initiated when CPU use exceeds the value specified using the `cpu-factor` statement. Junos OS then executes the number of lines of script specified using the `line-interval` statement and then pauses script execution for the time specified using the `time-interval` statement. If you do not configure values for these statements, the default values are used.

Options

`line-interval` Specify the line interval at which to pause.

- **Range:** 1 through 4,294,967,295
- **Default:** 10

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Dampen Script Execution

line-loopback

IN THIS SECTION

- Syntax | 9775
- Hierarchy Level | 9775
- Description | 9775
- Default | 9775
- Options | 9775
- Required Privilege Level | 9776
- Release Information | 9776

Syntax

```
(line-loopback-enable | no-line-loopback);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify whether line-loopback is enabled or disabled.

Default

If you omit the line-loopbback-enable statement, line-loopback is disabled.

Options

line-loopback-enable	Enable line-loopback.
no-line-loopback	Disable line-loopback.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

line-protocol

IN THIS SECTION

- Syntax | 9776
- Hierarchy Level | 9777
- Description | 9777
- Options | 9777
- Required Privilege Level | 9777
- Release Information | 9777

Syntax

```
line-protocol protocol;
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

For serial interfaces only, configure the line protocol.

Options

protocol—You can specify the one of the following line protocols:

- `eia530`—Line protocol EIA-530
- `v.35`—Line protocol V.35
- `x.21`—Line protocol X.21

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Line Protocol

line-rate

IN THIS SECTION

- Syntax | 9778
- Hierarchy Level | 9778
- Description | 9778
- Options | 9778
- Required Privilege Level | 9779
- Release Information | 9779

Syntax

```
line-rate line-rate;
```

Hierarchy Level

```
[edit interfaces interface-name shdsl-options],
[edit logical-systems logical-system-name interfaces interface-name shdsl-options]
```

Description

For J Series Services Routers only, configure the SHDSL line rate.

Options

line-rate—SHDSL line rate, in Kbps. Possible values are:

2-wire (Kbps): `192`, `256`, `320`, `384`, `448`, `512`, `576`, `640`, `704`, `768`, `832`, `896`, `960`, `1024`, `1088`, `1152`, `1216`, `1280`, `1344`, `1408`, `1472`, `1536`, `1600`, `1664`, `1728`, `1792`, `1856`, `1920`, `1984`, `2048`, `2112`, `2176`, `2240`, `2304`, `auto`

4-wire (Kbps): `384`, `512`, `640`, `768`, `896`, `1024`, `1152`, `1280`, `1408`, `1536`, `1664`, `1792`, `1920`, `2048`, `2176`, `2304`, `2432`, `2560`, `2688`, `2816`, `2944`, `3072`, `3200`, `3328`, `3456`, `3584`, `3712`, `3840`, `3968`, `4096`, `4224`, `4352`, `4480`, `4608`

- **Default:** For 2-wire mode, `auto`; for 4-wire mode, `4608` Kbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

line-rate (Interfaces)

IN THIS SECTION

- Syntax | 9779
- Hierarchy Level | 9779
- Description | 9779
- Options | 9780
- Required Privilege Level | 9780
- Release Information | 9780

Syntax

```
line-rate
```

Hierarchy Level

```
[edit interfaces interfaces name shdsl-options]
```

Description

Specify a line rate for an G.SHDSL interface.

Options

- `auto`— Automatically selects a line rate.
- `value` — Select the values between 192 kbps and 22784 kbps for the speed of transmission of data on the G.SHDSL connection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.0.

linerate-mode

IN THIS SECTION

- Syntax | 9780
- Hierarchy Level | 9781
- Description | 9781
- Required Privilege Level | 9781
- Release Information | 9781

Syntax

```
linerate-mode;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number linerate-mode],
[edit chassis lcc number fpc slot-number pic pic-number linerate-mode]
```

Description

For 10-port 10-Gigabit Oversubscribed Ethernet (OSE) PICs only, configure the line rate operation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Junos OS Network Interfaces Library for Routing Devices

link (Access)

IN THIS SECTION

- Syntax | 9782
- Hierarchy Level | 9782
- Description | 9782
- Options | 9782
- Required Privilege Level | 9782
- Release Information | 9782

Syntax

```
link pool-name;
```

Hierarchy Level

```
[edit access address-assignment pool]
```

Description

Configure the name of the secondary address-assignment pool that is linked to a primary address-assignment pool. The secondary pool provides a backup pool for local address assignment.

Options

pool-name—Name of the address assignment pool.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.4 of Junos OS.

RELATED DOCUMENTATION

> *Firewall User Authentication Overview*
>
> *Obtaining Username and Role Information Through Firewall Authentication*

link-adjacency-loss

IN THIS SECTION

- Syntax | 9783
- Hierarchy Level | 9783
- Description | 9783
- Required Privilege Level | 9783
- Release Information | 9783

Syntax

```
link-adjacency-loss;
```

Hierarchy Level

```
[edit protocols oam                 ethernet                 link-fault-management
action-profile                 event]
```

Description

Configure **loss of adjacency** event with the IEEE 802.3ah link fault management (LFM) peer. When included, the loss of adjacency event triggers the action specified under the "action" on page 1764 statement.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Example: Configuring Ethernet OAM Link Fault Management

Configuring Ethernet OAM Link Fault Management

link-degrade-monitor

IN THIS SECTION

- Syntax | 9784
- Hierarchy Level | 9785
- Description | 9785
- Options | 9786
- Required Privilege Level | 9786
- Release Information | 9786

Syntax

```
link-degrade-monitor {
    (link-degrade-monitor-enable | no-link-degrade-monitor-enable);
    actions media-based;
    recovery {
        (auto | manual);
        timer timer;
    }
    thresholds {
        clear clear-value;
        interval interval-value;
        set set-value;
        warning-clear warning-clear-value;
        warning-set warning-set-value;
    }
}
```

Hierarchy Level

```
[edit interfaces  interfaces- name]
```

Description

Configure link degrade monitoring on an interface and when a link degrade event is detected, specify the corrective action to be triggered. When configured, the feature monitors the quality of physical links on Ethernet interfaces (10-Gigabit, 25-Gigabit, 40-Gigabit, 50-Gigabit, 100-Gigabit, 200-Gigabit and 400-Gigabit) and triggers the user-configured action when the link's bit error rate (BER) value breaches the configured threshold. This feature can detect a BER value as low as 10^{-15} to 10^{-1}.

Starting in Junos OS Evolved Release 20.2R1, you can monitor physical link degradation based on Bit Error Rate (BER)) on 10-Gigabit, 25-Gigabit, 40-Gigabit, 50-Gigabit, 100-Gigabit, 200-Gigabit and 400-Gigabit Ethernet interfaces on PTX1001-36MR. The guidelines to configure `link-degrade-monitor` on Gigabit Ethernet interfaces are:

- Whenever there is a local degrade event, in any of the interfaces, a `LOCAL-DEGRADE` flag is set and a `REMOTE-DEGRADE` flag is advertised on the peer interface of the connected router. When you enable the `link-degrade` option, Junos OS will maintain statistics such as, number of times the link degrade event has occurred and if the link is degraded, then seconds elapsed since the link degrade event. You can view these details by using the "show interfaces " on page 20227 command. This command also shows the `Link Degrade Status` (if the link degrade `Defect`) is active or cleared. See "show interfaces " on page 20227 for output information.

- On 400-Gigabit Ethernet interfaces:

  - The permissible range of BER is between 1E-5 to 1E-10. If you set any threshold value beyond the permissible range the value will be auto adjusted between 1E-5 to 1E-10. If you configure the set threshold value higher than 1E-5 (i.e between 1E-1 to 1E-4), the value will be adjusted to 1E-5. If you configure the clear threshold value lesser than 1E-10 (i.e between 1E-11 and 1E-15), the value will be adjusted to 1E-10.

  - The option `threshold intervals` and `recover timer` are ignored. The error rates for 400-Gigabit Ethernet are always calculated per second. As soon as the BER per second becomes lower than clear threshold error rate, the link will be considered as recovered from degradation.

  - The bit errors detected before applying FEC (Pre-FEC) correction is considered for determining the degrade state. The `set` and `clear` thresholds apply for pre FEC bit errors. For information FEC support see, Port Speed on PTX10001-36MR Router Overview.

Options

link-degrade-monitor-enable	Enable link degrade monitoring on an interface.
no-link-degrade-monitor-enable	Disable link degrade monitoring on an interface.
media-based	Action to be taken when a link degrade event is detected. A media-based action brings down the physical link at both local and remote ends of the interface, and stops BER monitoring at the local end until an autorecovery is triggered.

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

The options `link-degrade-monitor-enable` and `no-link-degrade-monitor-enable` are introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

Link Degrade Monitoring Overview

show interfaces | **20227**

thresholds | **29346**

recovery | **14146**

request interface link-degrade-recover | **14642**

Common Output Fields Description

link-discovery

IN THIS SECTION

- Syntax | 9787
- Hierarchy Level | 9787
- Description | 9787
- Options | 9788
- Required Privilege Level | 9788
- Release Information | 9788

Syntax

```
link-discovery (active | passive);
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management
interface interface-name]
```

Description

Specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on an interface. Link monitoring is done when the interface sends periodic OAM PDUs.

For Ethernet interfaces on EX Series switches, and M320, M120, MX Series, and T Series routers, specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Management (OAM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs.

Options

active—In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality.

passive—In passive mode, the peer initiates the discovery process.

Once the discovery process is initiated, both sides participate in discovery.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Ethernet OAM Link Fault Management

link-down (QFX Series)

IN THIS SECTION

- Syntax | 9789
- Hierarchy Level | 9789
- Description | 9789
- Options | 9789
- Required Privilege Level | 9789
- Release Information | 9789

Syntax

```
link-down (red | yellow | ignore);
```

Hierarchy Level

```
[edit chassis alarm ethernet (Alarm)],
[edit chassis alarm fibre-channel],
[edit chassis interconnect-device name alarm ethernet (Alarm)],
[edit chassis node-group name alarm fibre-channel]
```

Description

Specify either red, yellow, or ignore to display when the link is down.

Options

red Indicates that one or more hardware components have failed or exceeded temperature thresholds, or an alarm condition configured on an interface has triggered a critical warning.

yellow Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.

ignore Suppresses or ignores the alarm.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

link-down

IN THIS SECTION

- Syntax | 9790
- Hierarchy Level | 9790
- Description | 9790
- Required Privilege Level | 9790
- Release Information | 9790

Syntax

```
link-down;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management
action-profile action]
```

Description

Mark the interface as down for transit traffic.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Ethernet OAM Link Fault Management

link-event-rate

IN THIS SECTION

- Syntax | 9791
- Hierarchy Level | 9791
- Description | 9791
- Required Privilege Level | 9792
- Release Information | 9792

Syntax

```
link-event-rate {
                        frame-error                     count;
                        frame-period                    count;
                        frame-period-summary                    count;
                        symbol-period                    count;
}
```

Hierarchy Level

```
[edit protocols oam                 ethernet                 link-fault-management
action-profile                 event]
```

Description

Configure the number of link fault management (LFM) events per second.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Ethernet OAM Link Fault Management

link-fault-management

IN THIS SECTION

- Syntax | 9792
- Hierarchy Level | 9793
- Description | 9793
- Required Privilege Level | 9793
- Release Information | 9794

Syntax

```
link-fault-management {
    action-profile profile-name;
        action {
            syslog;
            link-down;
        }
        event {
            link-adjacency-loss;
            link-event-rate {
```

```
                frame-error  count;
                frame-period count;
                frame-period-summary  count;
                symbol-period count;
            }
        }
    interface interface-name {
        link-discovery (active | passive);
        pdu-interval interval;
        pdu-threshold threshold-value;
        remote-loopback;
        event-thresholds  {
            frame-error count;
             frame-period count;
             frame-period-summary  count;
             symbol-period count;
        }
        negotiation-options {
            allow-remote-loopback;
             no-allow-link-events;
        }
    }
}
```

Hierarchy Level

```
[edit protocols oam                ethernet]
```

Description

Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Example: Configuring Ethernet OAM Link Fault Management

Configuring Ethernet OAM Link Fault Management

link-group-protection (Protocols IS-IS)

IN THIS SECTION

- Syntax | 9794
- Hierarchy Level | 9794
- Description | 9795
- Options | 9795
- Required Privilege Level | 9795
- Release Information | 9795

Syntax

```
link-group-protection {
    minimum-bandwidth rate;
    revert-bandwidth rate;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis interface-group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis
interface-group group-name],
```

```
[edit protocols isis interface-group group-name],
[edit routing-instances routing-instance-name protocols isis interface-group group-name]
```

Description

Configure link group protection. This enables Packet Forwarding Engine-based local repair by pre-downloading suboptimal backup paths with appropriate weights in the Packet Forwarding Engine.

Options

minimum-bandwidth *rate*

Minimum bandwidth to carry traffic. This is the minimum bandwidth at which the outgoing traffic can no longer be carried on this link group. When the bandwidth becomes less than or equal to this bandwidth, traffic will get load-balanced between primary and backup path.

- **Range:** 9600 through 1,000,000,000,000

revert-bandwidth *rate*

Revert bandwidth to carry traffic. This is the bandwidth above which the outgoing traffic can be carried over the members of this link group. There is no need for the Packet Forwarding Engine to load-balance across member links of other suboptimal link groups if it provides a path with the shortest metric.

- **Range:** 9600 through 1,000,000,000,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1 for the MX Series.

link-layer-broadcast-inet-check

IN THIS SECTION

- Syntax | 9796
- Hierarchy Level | 9796
- Description | 9796
- Required Privilege Level | 9796
- Release Information | 9796

Syntax

```
link-layer-broadcast-inet-check;
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Enable destination MAC and IP address check to prevent the Router from forwarding IPV4 packets, which have link layer destination address set to `broadcast` or `multicast`, unless it is directed to an *IPV4 multicast* address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

link-layer-overhead

IN THIS SECTION

- Syntax | 9797
- Hierarchy Level | 9797
- Description | 9797
- Options | 9797
- Required Privilege Level | 9798
- Release Information | 9798

Syntax

```
link-layer-overhead percent;
```

Hierarchy Level

```
[edit interfaces interface-name mlfr-uni-nni-bundle-options],
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For link services IQ (`lsq`) interfaces only, configure the percentage of total bundle bandwidth to be set aside for link-layer overhead. Link-layer overhead accounts for the bit stuffing on serial links. Bit stuffing is used to prevent data from being interpreted as control information. Overhead resulting from link-layer encapsulation and framing is computed automatically.

Options

percent—Percentage of total bundle bandwidth to be set aside for link-layer overhead.

- **Range:** 0 through 50 percent

- **Default:** 0 percent

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring CoS Scheduling Queues on Logical LSQ Interfaces*

link-management

IN THIS SECTION

- Syntax | 9798
- Hierarchy Level | 9799
- Description | 9799
- Required Privilege Level | 9799
- Release Information | 9800

Syntax

```
link-management {
    peer peer-name {
        address ip-address;
        control-channel control-channel-interface;
        lmp-control-channel control-channel-interface {
            remote-address ip-address;
```

```
        }
        lmp-protocol {
            hello-dead-interval milliseconds;
            hello-interval milliseconds;
            passive;
            retransmission-interval milliseconds;
            retry-limit number;
        }
        te-link te-link-name;
    }
    te-link te-link-name {
        disable;
        interface interface-name {
            disable;
            local-address ip-address;
            remote-address ip-address;
            remote-id id-number;
        }
        local-address ip-address;
        remote-address ip-address;
        remote-id id-number;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit protocols]
```

Description

Enable Link Management Protocol (LMP) on the router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- LMP Configuration Overview

link-mode

IN THIS SECTION

- Syntax | 9800
- Hierarchy Level | 9800
- Description | 9801
- Default | 9801
- Options | 9801
- Required Privilege Level | 9801
- Release Information | 9801

Syntax

```
link-mode mode;
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
```

Description

Set the device's link-connection characteristic.

> **NOTE**: Some QFX Series switch interfaces are full-duplex only. In such cases, the CLI shows only "link-mode full-duplex" or does not include the "link-mode" statement at all. You cannot configure the link mode on these interfaces; it is set to the full-duplex default.

Default

The **full-duplex** mode is enabled.

Options

mode —Link characteristic:

- **full-duplex**—Connection is full duplex.
- **half-duplex**—Connection is half duplex.
- **automatic**—Link mode is negotiated.

If **no-auto-negotiation** is specified in the **ether-options** option, you can select only **full-duplex** or **half-duplex**. If "auto-negotiation" on page 2718 is specified in the **ether-options** option, you can select any mode.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

> *Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches*
>
> Junos OS Network Interfaces Library for Routing Devices

link-mtu

IN THIS SECTION

- Syntax | 9802
- Hierarchy Level | 9802
- Description | 9802
- Default | 9803
- Required Privilege Level | 9803
- Release Information | 9803

Syntax

```
(link-mtu | no-link-mtu);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Description

Specify whether to include the maximum transmission unit (MTU) option in router advertisement messages:

- **link-mtu**–Includes the MTU option in router advertisements.
- **no-link-mtu**–Does not include the MTU option in router advertisements.

The MTU option included in router advertisement messages ensures that all nodes on a link use the same MTU value in situations where the link MTU is not well known.

Default

Router advertisement messages do not include the MTU option.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 10.3.

RELATED DOCUMENTATION

- Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

link-protection

IN THIS SECTION

- Syntax | 9803
- Hierarchy Level | 9804
- Description | 9804
- Options | 9805
- Required Privilege Level | 9805
- Release Information | 9805

Syntax

```
    aggregated-ether-options {
        lacp {
+           link-protection {
```

```
+               revertive | non-revertive;
+         disable
            }
        }
    }
}
link-protection {
    disable;
    (revertive |non-revertive);
    disable
}
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
[edit interfaces aex aggregated-ether-options lacp]
```

Description

On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the `primary` and `backup` statements at the `[edit interfaces ge-fpc/pic/port gigether-options 802.3ad aex]` hierarchy level or the `[edit interfaces fe-fpc/pic/port fastether-options 802.3ad aex]` hierarchy level.

On the switch, you can configure either Junos OS link protection for aggregated Ethernet interfaces or the LACP standards link protection for aggregated Ethernet interfaces.

For Junos OS link protection, specify `link-protection` at the following hierarchy levels:

- `[edit interfaces ge-fpc/pic/port ether-options 802.3ad aex]`
- `[edit interfaces xe-fpc/pic/port ether-options 802.3ad aex]` hierarchy level or at the `[edit interfaces xe-fpc/pic/port ether-options 802.3ad aex]` hierarchy level.

  To disable link protection, use the `delete interface ae aggregate-ether-options link-protection` statement at the `[edit interfaces aex aggregated-ether-options]` hierarchy level or the `[edit interfaces aex aggregated-ether-options lacp]]` hierarchy level.

NOTE: Junos Evolved series versions do not support N:N link-protection.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support for `disable`, `revertive`, and `non-revertive` statements added in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Aggregated Ethernet Link Protection

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

link-protection (Chassis Cluster)

IN THIS SECTION

- Syntax | 9806
- Hierarchy Level | 9806
- Description | 9806
- Options | 9806
- Required Privilege Level | 9806
- Release Information | 9806

Syntax

```
link-protection {
    non-revertive;
}
```

Hierarchy Level

```
[edit chassis aggregated-devices ethernet lacp]
```

Description

Enable Link Aggregation Control Protocol (LACP) link protection at the global (chassis) level.

By default LACP link protection reverts to a higher-priority (lower-numbered) link when the higher-priority link becomes operational or when a higher-priority link is added to the aggregated Ethernet bundle.

You can suppress link calculation by adding the non-revertive statement to the link protection configuration. In nonrevertive mode, when a link is active in sending and receiving LACP packets, adding a higher-priority link to the bundle does not change the status of the currently active link. It remains active.

Options

non-revertive Disable the ability to switch to a better priority link (if one is available) after a link is established as active and a collection or distribution is enabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

Example: Configuring Aggregated Ethernet Device with LAG and LACP on a Security Device (CLI Procedure)

link-protection (Dynamic LSPs)

IN THIS SECTION

- Syntax | 9807
- Hierarchy Level | 9807
- Description | 9807
- Default | 9808
- Required Privilege Level | 9808
- Release Information | 9808

Syntax

```
link-protection;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Enable link protection on the specified LSP, which helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails. For point-to-multipoint LSPs, including this statement extends link protection to all of the paths used by the LSP.

To fully enable link protection, you must also include the `link-protection` statement at the `[edit protocols rsvp interface interface-name]` or `[edit logical-systems logical-system-name protocols rsvp interface interface-name]` hierarchy level.

Default

Link protection is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Link Protection for Point-to-Multipoint LSPs

Configuring Node Protection or Link Protection for LSPs

link-protection (RSVP) | **9816**

link-protection (Host Fast Reroute)

IN THIS SECTION

- Syntax | **9809**
- Hierarchy Level | **9809**
- Description | **9809**
- Required Privilege Level | **9809**
- Release Information | **9809**

Syntax

```
link-protection;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options interface interface-name],
[edit routing-instances routing-instance-name routing-options interface interface-name]
```

Description

Enable link protection on a specified interface in a routing instance.

Configuring link protection on an interface ensures that traffic is rerouted to the destination device through alternate loop-free paths that traverse the protected interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring Link Protection with Host Fast Reroute

link-protection (non-LACP)

IN THIS SECTION

- Syntax | 9810
- Hierarchy Level | 9810
- Description | 9810
- Options | 9811
- Required Privilege Level | 9811
- Release Information | 9811

Syntax

```
link-protection {
    revertive | non-revertive;
}
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

Description

User can specify the link-protection revertive statement in the link protection configuration at the aggregated Ethernet interface level to set revertive mode. In revertive mode, adding a higher-priority link to the aggregated Ethernet bundle results in recalculation of the priorities and traffic will switch from the currently active link to the newly added, higher-priority link. Recalculation of priorities is performed only while link event such as addition\deletion and UP/Down operation on link, that is, configuration of this option will not result in any recalculation immediately until next link-event occurs.

In addition to enabling static link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the `primary` and `backup` statements at the `[edit interfaces ge-`*fpc*`/`*pic*`/`*port*` gigether-options 802.3ad]` hierarchy level or the `[edit interfaces fe-`*fpc*`/`*pic*`/`*port*` fastether-options 802.3ad]` hierarchy level.

For static link protection, specify `link-protection` at the hierarchy level:

```
[edit interfaces ae0 aggregated-ether-options link-protection]
```

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring Aggregated Ethernet Link Protection

link-protection (Protocols IS-IS)

IN THIS SECTION

- Syntax | 9812
- Hierarchy Level | 9812
- Description | 9812
- Options | 9812
- Required Privilege Level | 9812
- Release Information | 9812

Syntax

```
link-protection;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis
interface interface-name],
[edit protocols isis interface interface-name],
[edit routing-instances routing-instance-name protocols isis interface interface-name]
```

Description

Enable link protection on the specified IS-IS interface. Junos OS creates a backup loop-free alternate path to the primary next hop for all destination routes that traverse the protected interface.

Options

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

> *Understanding Loop-Free Alternate Routes for IS-IS*
>
> *Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN*

link-protection (Protocols LACP)

IN THIS SECTION

- Syntax | 9813
- Hierarchy Level | 9813
- Description | 9813
- Options | 9813
- Required Privilege Level | 9813
- Release Information | 9814

Syntax

```
link-protection {
    non-revertive;
}
```

Hierarchy Level

```
[edit chassis aggregated-devices ethernet lacp]
```

Description

Enable LACP link protection at the global (chassis) level.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Junos OS for Supporting Aggregated Devices

link-protection (Protocols OSPF)

IN THIS SECTION

- Syntax | 9814
- Hierarchy Level | 9814
- Description | 9815
- Required Privilege Level | 9815
- Release Information | 9815

Syntax

```
link-protection;
```

Hierarchy Level

```
[edit protocols (ospf | ospf3) area area-name interface interface-name],
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-name interface
interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-name interface
interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-name interface interface-name],
[edit protocols ospf3 realm ipv4-unicast area area-id],
```

```
[edit logical-systems logical-system-name protocols ospf3 realm ipv4-unicast area area-id],
[edit routing-instances routing-instance-name protocols ospf3 realm ipv4-unicast area area-id],
[edit protocols ospf area area-id interface interface-name topology (default | name)],
[edit logical-systems logical-system-name protocols ospf area area-id interface interface-name
topology (default | name)],
[edit routing-instances routing-instance-name protocols ospf area area-id interface interface-
name topology (default | name)]
```

Description

Enable link protection on the specified OSPF interface. Junos OS creates a backup loop-free alternate path to the primary next hop for all destination routes that traverse the protected interface.

> **NOTE**: This feature calculates alternate next hop paths for unicast routes only. Therefore, this statement is not supported with the OSPF IPv4 multicast topology or with the OSPFv3 IPv4 multicast and IPv6 multicast realms.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring Link Protection for OSPF

link-protection (RSVP)

IN THIS SECTION

- Syntax | 9816
- Hierarchy Level | 9817
- Description | 9817
- Default | 9817
- Options | 9817
- Required Privilege Level | 9817
- Release Information | 9818

Syntax

```
link-protection {
    disable;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps;
    bypass bypass-name {
        bandwidth bps {
            ct0 bps;
            ct1 bps;
            ct2 bps;
            ct3 bps;
        }
        description text;
        class-of-service cos-value;
        hop-limit number;
        no-cspf;
        path address <strict | loose>;
        priority setup-priority reservation-priority;
        to address;
```

```
    }
    class-of-service cos-value;
    hop-limit number;
    max-bypasses number;
    no-cspf;
    no-node-protection;
    optimize-timer seconds;
    path address <strict | loose>;
    priority setup-priority reservation-priority;
    subscription percentage;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit protocols rsvp interface interface-name]
```

Description

Enable link protection on the specified interface. Using link protection, you can configure a network to reroute traffic quickly around broken links. To fully enable link protection, you also need to configure the `link-protection` statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level. You can configure single or multiple bypasses for protected interface.

Default

Link protection is disabled.

Options

`no-node-protection`—Disable node-link protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

RELATED DOCUMENTATION

Configuring Link Protection on Interfaces Used by LSPs

link-protection (Dynamic LSPs) | 9807

link-protection (Static LSPs)

IN THIS SECTION

- Syntax | 9818
- Hierarchy Level | 9819
- Description | 9819
- Default | 9819
- Options | 9819
- Required Privilege Level | 9819
- Release Information | 9819

Syntax

```
link-protection bypass-name name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
transit incoming-label],
[edit protocols mpls static-label-switched-path lsp-name ingress],
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Description

Enable link protection on the specified static LSP. Link protection helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails.

Default

Link protection is disabled.

Options

`bypass-name` *name*—Bypass LSP name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *Configuring Static LSPs*
>
> *Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP*

link-protection-sub-group (802.3ad)

IN THIS SECTION

- Syntax | 9820
- Hierarchy Level | 9820
- Description | 9820
- Options | 9820
- Required Privilege Level | 9821
- Release Information | 9821

Syntax

```
link-protection-sub-group group-name;
```

Hierarchy Level

```
[edit interfaces interface-name ether-options 802.3ad]
```

Description

Add an interface in an aggregated Ethernet bundle into a link-protection subgroup.

A link protection subgroup is created and named using the `link-protection-sub-group` statement in the `[edit interfaces ae*x* aggregated-ether-options]` hierarchy.

Options

group-name Name of the link protection subgroup that will include this interface after this statement is entered. The link protection subgroup is named when it is created using the `link-protection-sub-group` statement in the `[edit interfaces ae*x* aggregated-ether-options]` hierarchy.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

- Configuring Aggregated Ethernet Links (CLI Procedure)
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*
- Q-in-Q Support on Redundant Trunk Links Using LAGs with Link Protection

link-protection-sub-group (aggregated-ether-options)

IN THIS SECTION

- Syntax | 9821
- Hierarchy Level | 9822
- Description | 9822
- Options | 9822
- Required Privilege Level | 9822
- Release Information | 9822

Syntax

```
link-protection-sub-group group-name {
    [primary | backup];
}
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

Description

Create and name a link protection subgroup.

Link protection subgroups allow you to provide link protection to a collection of Ethernet links within a link aggregation group (LAG). If you need to provide link protection to a single link in a LAG, you do not need to configure link protection subgroups.

A link protection subgroup includes multiple links within the LAG. If one link in the primary link protection subgroup fails, traffic is forwarded over the links in the backup link protection subgroup.

Links within the LAG are added to the link protection subgroup using the `link-protection-sub-group` statement in the [`edit interfaces` *`interface-name`* `ether-options 802.3ad`] hierarchy.

Options

`group-name`	User-provided name of the link protection subgroup.
`primary`	Subgroup is the primary subgroup.
`backup`	Subgroup is the backup subgroup.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

- Configuring Aggregated Ethernet Links (CLI Procedure)
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

Q-in-Q Support on Redundant Trunk Links Using LAGs with Link Protection

link-speed

IN THIS SECTION

- Syntax | 9823
- Hierarchy Level (QFX, NFX, EX Series, QFabric System, OCX1100, EX4600) | 9823
- Hierarchy Level (EX Series) | 9823
- Description | 9824
- Options | 9824
- Required Privilege Level | 9825
- Release Information | 9825

Syntax

```
link-speed speed;
```

Hierarchy Level (QFX, NFX, EX Series, QFabric System, OCX1100, EX4600)

```
[edit interfaces aex aggregated-ether-options]
```

Hierarchy Level (EX Series)

```
[edit interfaces aex aggregated-ether-options],
[edit interfaces interface-range name aggregated-ether-options],
[edit interfaces interface-range name aggregated-sonet-options]
```

Description

For aggregated Ethernet interfaces only, set the required link speed.

Options

speed—For aggregated Ethernet links, you can specify *speed* in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

Aggregated Ethernet links on the M120 router can have one of the following speeds:

- `100m`—Links are 100 Mbps.
- `10g`—Links are 10 Gbps.
- `1g`—Links are 1 Gbps.
- `oc192`—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- `10m`—Links are 10 Mbps.
- `100m`—Links are 100 Mbps.
- `1g`—Links are 1 Gbps.
- `10g`—Links are 10 Gbps.

speed—For aggregated Ethernet links, you can specify the speed in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

On QFX5100 and EX4600 standalone switches and on a QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.

Aggregated Ethernet links on T Series, MX Series, PTX Series routers, and QFX5100, QFX10002, QFX10008, and QFX10016 switches can be configured to operate at one of the following speeds:

- `100g`—Links are 100 Gbps.
- `100m`—Links are 100 Mbps.
- `10g`—Links are 10 Gbps.

- `1g`—Links are 1 Gbps.
- `40g`—Links are 40 Gbps.
- `50g`—Links are 50 Gbps.
- `80g`—Links are 80 Gbps.
- `8g`—Links are 8 Gbps.
- `mixed`—Links are of various speeds.
- `oc192`—Links are OC192.

`mixed`—Enables bundling of different Ethernet rate links in the same Aggregated Ethernet interface.

> **NOTE**: OCX Series switches only support 10g and 40g interfaces. Mixed rate aggregated Ethernet interfaces are not support on the OCX Series.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring Link Aggregation

Configuring Aggregated Ethernet Links (CLI Procedure)

Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

link-speed (Aggregated Ethernet)

IN THIS SECTION

- Syntax | 9826
- Hierarchy Level | 9826
- Description | 9826
- Options | 9826
- Required Privilege Level | 9827
- Release Information | 9828

Syntax

```
link-speed (100m | 10g| 10m | 1g| 2.5g| 25g| 400g| 40g| 50g| 5g| 80g| 8g| mixed| oc192);
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options],
[edit interfaces interface-range name aggregated-ether-options],
[edit interfaces interface-range name aggregated-sonet-options]
```

Description

For aggregated Ethernet interfaces only, set the required link speed. You can specify *speed* in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000). By default, the speed of the first aggregated interface that you add to the aggregated ethernet bundle determines the link speed of the aggregated ethernet interface bundle. If you specify the link speed, all the aggregated ethernet interfaces of the bundle operate at the specified speed.

Options

Aggregated Ethernet links on devices can have one of the following speeds:

- `100m`—(EX Series, QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Links are 100 Mbps.
- `10g`—(EX Series, QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Links are 10 Gbps.
- `10m`—(EX Series, QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Links are 10 Mbps.
- `1g`—(EX Series, QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Links are 1 Gbps.
- `2.5g`—(QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Links are 2.5 Gbps.
- `25g`—(QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Links are 25 Gbps.
- `400g`—Links are 400 Gbps.
- `40g`—(QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Links are 40 Gbps.
- `50g`—(QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Links are 50 Gbps.
- `80g`—(QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Links are 80 Gbps.
- `8g`—(QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Links are 8 Gbps.
- `mixed`—(QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Enables bundling of different Ethernet rate links in the same Aggregated Ethernet interface.
- `oc192`—(QFX5100, QFX100002, QFX100008, QFX100016 switches, MX Series, PTX Series routers) Links are OC192.
- `oc192`—Links are OC192 or STM64c.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`mixed` option added in Junos OS Release 15.1F3 and 16.1R2 for PTX5000 routers and 15.1F6 and 16.1R2 for PTX3000 routers.

RELATED DOCUMENTATION

- Aggregated Ethernet Interfaces Overview
- Configuring Aggregated Ethernet Link Speed
- Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles
- Configuring Aggregated Ethernet Links (CLI Procedure)
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

link-speed (Aggregated SONET/SDH)

IN THIS SECTION

- Syntax | 9828
- Hierarchy Level | 9829
- Description | 9829
- Options | 9829
- Required Privilege Level | 9829
- Release Information | 9829

Syntax

```
link-speed (speed | mixed);
```

Hierarchy Level

```
[edit interfaces asx aggregated-sonet-options]
```

Description

For aggregated SONET/SDH interfaces only, set the required link speed.

Options

speed—Aggregated SONET/SDH links can have one of the following speed values.

- `oc3`—Links are OC3c or STM1c.
- `oc12`—Links are OC12c or STM4c.
- `oc48`—Links are OC48c or STM16c.
- `oc192`—Links are OC192c or STM64c.
- `oc768`—Links are OC768c or STM256c.

`mixed`—For aggregated SONET/SDH links on T Series routers, you can mix interface speeds in SONET/SDH aggregation bundles. Interface speeds from OC3 through OC768 are supported.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`mixed` option added in Release 8.0.

RELATED DOCUMENTATION

Configuring Aggregated Ethernet Link Speed

Configuring Aggregated SONET/SDH Interfaces

link-speed (Interfaces)

IN THIS SECTION

- Syntax | 9830
- Hierarchy Level | 9830
- Description | 9830
- Options | 9830
- Required Privilege Level | 9830
- Release Information | 9831

Syntax

```
link-speed speed;
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options]
```

Description

For redundant Ethernet interfaces in a chassis cluster only, set the required link speed.

Options

speed —For redundant Ethernet links, you can specify *speed* in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Release 9.0 of Junos OS.

link-to-disable

IN THIS SECTION

- Syntax | 9831
- Hierarchy Level | 9831
- Description | 9831
- Options | 9832
- Required Privilege Level | 9832
- Release Information | 9832

Syntax

```
link-to-disable interface-name;
```

Hierarchy Level

```
[edit protocols uplink-failure-detection group group-name]
```

Description

Configure the downlink interfaces to be disabled when the switch detects an uplink failure. The switch can monitor a maximum of 48 downlink interfaces in a group.

NOTE: In QFX5120 and QFX5210 Series switches, you can configure a maximum of 128 downlink interfaces.

Options

interface-name Name of the downlink interface or interface range in the group. The interface can be a physical interface or a logical interface.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

link-to-monitor

IN THIS SECTION

- Syntax | 9832
- Hierarchy Level | 9833
- Description | 9833
- Options | 9833
- Required Privilege Level | 9833
- Release Information | 9833

Syntax

```
link-to-monitor interface-name;
```

Hierarchy Level

```
[edit protocols uplink-failure-detection group group-name]
```

Description

Configure the uplink interfaces to be monitored for uplink failure detection. The switch can monitor a maximum of 48 uplink interfaces in a group.

> **NOTE**: In QFX5120 and QFX5210 Series switches, you can configure a maximum of 128 uplink interfaces.

An interface can be configured as link-to-monitor in multiple groups.

Options

interface-name Name of the uplink interface or interface range in the group. The interface can be a physical interface or a logical interface.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

link-tracking (Network Isolation Groups)

IN THIS SECTION

Syntax | 9834

- Hierarchy Level | 9834
- Description | 9834
- Options | 9834
- Required Privilege Level | 9835
- Release Information | 9835

Syntax

```
link-tracking {
    failure-type <link>;
    interface name;
    minimum-links count;
}
```

Hierarchy Level

```
[edit protocols network-isolation group name detection]
```

Description

Track the link status of the configured interface or interfaces to detect network isolation conditions. You configure this statement when you define a network isolation group that you apply to a bridge domain, VLAN, or EVPN instance (EVI).

Options

(Required) failure-type link	Type of failure. The only supported failure type currently is `link` for link failures, which is also the default failure type.
interface {*interface-name* | [*interface-name* ...]}	Track the interface state of the listed logical Layer 3 (L3) interfaces to detect network isolation conditions. Enclose multiple interface names in square brackets separated by spaces, or specify a single interface name without square brackets.

minimum-links *count*

Minimum number of L2 links that must be **up** to indicate the device isn't in a network isolation state. These are the L2 links associated with the L3 interfaces you specify with the `interface` option.

- **Range:** 1 through 10
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Determine IRB Interface State Changes from Local L2 Interface or Remote Connectivity Status in EVPN Fabrics

interface-state | **9064**

network-isolation-profile | **11314**

linkdown (satellite-policies alarm)

IN THIS SECTION

- Syntax | **9836**
- Hierarchy Level | **9836**
- Description | **9836**
- Options | **9836**
- Required Privilege Level | **9836**
- Release Information | **9836**

Syntax

```
linkdown [ignore | red | yellow]
```

Hierarchy Level

```
[edit policy-options satellite-policies environment-monitoring-policy policy-name alarm]
```

Description

Configure the alarm behavior when an Ethernet link goes down on a satellite device in a Junos Fusion.

The configured alarm behavior can be applied to any satellite device in the Junos Fusion. The alarm behavior is applied to satellite devices using environment monitoring policies.

Options

ignore Do not signal an alarm when an Ethernet link-down event occurs.

red Raise a major alarm when an Ethernet link-down event occurs.

yellow Raise a minor alarm when an Ethernet link-down event occurs.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Satellite Device Alarm Handling Using an Environment Monitoring Satellite Policy in a Junos Fusion

Understanding Satellite Policies in a Junos Fusion

linked-pool-aggregation (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 9837
- Hierarchy Level | 9837
- Description | 9837
- Required Privilege Level | 9838
- Release Information | 9838

Syntax

```
linked-pool-aggregation;
```

Hierarchy Level

```
[edit access]
```

Description

Change how the search for an available address is performed. When you include this statement, all ranges in the matching pool are searched, but only from the last-saved next address to the highest address in the range. The same search is performed as necessary in a linked pool, through the end of the chain of pools. Then the search begins again at the first pool in the chain; all ranges are searched in each pool, but this time the entire range is searched, from the lowest to the highest address in the range.

This nondefault behavior enables addresses to be assigned non-contiguously, meaning that a free address can be allocated from farther down the chain of linked pools even when free addresses are available in the matching pool or the first pool in the chain. This may be desirable if you configure your RADIUS server to use the IP address alone to identify subscribers.

For example, without this statement, a subscriber can disconnect and that address can be assigned to the next subscriber. The Acct-Start for the second subscriber is sent before the Acct-Stop is sent for the disconnected subscriber. When the Acct-Stop is received, the new subscriber, identified only by the IP address, may be disconnected.

You can avoid this situation by either including the `linked-pool-aggregation` statement or configuring your RADIUS server to use the subscriber session ID (instead of the IP address) for identification.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Configuring Address-Assignment Pool Linking

Address Allocation from Linked Address Pools

Address-Assignment Pools Overview

Address-Assignment Pool Configuration Overview

linktrace

IN THIS SECTION

- Syntax | 9839
- Hierarchy Level | 9839
- Description | 9839
- Required Privilege Level | 9839
- Release Information | 9839

Syntax

```
linktrace {
    age (30m | 10m | 1m | 30s | 10s);
    path-database-size path-database-size;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

Description

Configure connectivity fault management linktrace parameters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Linktrace Protocol in CFM

list

IN THIS SECTION

Syntax | 9840

- Hierarchy Level | 9840
- Description | 9840
- Options | 9840
- Required Privilege Level | 9840
- Release Information | 9841

Syntax

```
list listname {
    exception listname;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus mime-whitelist]
```

Description

Configure the antivirus scanner to use MIME bypass lists (custom objects). If you want to have exceptions to the allowlist, create a mime-pattern list with the `exception` statement in addition to the `list` statement.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

list (Security Content Filtering Block Mime)

IN THIS SECTION

- Syntax | 9841
- Hierarchy Level | 9841
- Description | 9841
- Required Privilege Level | 9841
- Release Information | 9842

Syntax

```
list list-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name block-mime]
```

Description

Configure the content filter to use MIME block lists (custom objects).

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

listen

IN THIS SECTION

- Syntax | 9842
- Hierarchy Level | 9842
- Description | 9843
- Options | 9843
- Required Privilege Level | 9843
- Release Information | 9843

Syntax

```
listen address <port port>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols sap],
[edit protocols sap]
```

Description

Specify an address and optionally a port on which SAP and SDP listen, in addition to the default SAP address and port on which they always listen, 224.2.127.254:9875. To specify multiple additional addresses or pairs of address and port, include multiple **listen** statements.

Options

address—(Optional) Address on which SAP listens for session advertisements.

- **Default:** 224.2.127.254

port *port*—(Optional) Port on which SAP listens for session advertisements.

- **Default:** 9875

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Session Announcement Protocol

listening-mode

IN THIS SECTION

- Syntax | 9844
- Hierarchy Level | 9844
- Description | 9844

Required Privilege Level | 9844
Release Information | 9844

Syntax

```
listening-mode;
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

Use listening-mode to enable listening-mode for GTP inspection.

The SRX Series firewall enables GPRS tunneling protocol (GTP) listening mode on the S5, S8, or S11 interface. In listening mode, the firewall performs sanity check and stateful inspection on GTP packets. Any failure in the GTP packet processing will stop the sanity check and stateful inspection, but the GTP packets are forwarded directly.

In this case, SRX Series Firewall does not provide protection to GTP traffic in listening mode. If a GTP packet containing invalid content arrives at the destination, it can harm the operator's network.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show security gtp profile | 24448

liveness-detection (iccp-peer)

IN THIS SECTION

- Syntax | 9845
- Hierarchy Level | 9845
- Description | 9846
- Required Privilege Level | 9846
- Release Information | 9846

Syntax

```
liveness-detection {
    detection-time {
        threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
```

Hierarchy Level

```
[edit protocols iccp peer]
```

Description

Enable Bidirectional Forwarding Detection (BFD). BFD enables rapid detection of communication failures between peers.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

liveness-detection

IN THIS SECTION

- Syntax | 9846
- Hierarchy Level | 9847
- Description | 9847
- Required Privilege Level | 9848
- Release Information | 9848

Syntax

```
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
```

```
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
        layer2-liveness-detection {
            max-consecutive-retries number;
            transmit-interval interval;
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name ],
[edit forwarding-options dhcp-relay group group-name],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server group group-name]
```

Description

Configure bidirectional failure detection timers and authentication criteria for static routes.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- DHCP Liveness Detection Overview
- Configuring Detection of DHCP Local Server Client Connectivity with BFD
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity with BFD
- Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients
- Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients
- DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

liveness-detection (high availability)

IN THIS SECTION

- Syntax | 9849
- Hierarchy Level | 9849
- Description | 9849
- Options | 9849
- Required Privilege Level | 9851
- Release Information | 9851

Syntax

```
liveness-detection {
    detection-time {
       threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier multiplier;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (0 | 1 | automatic);
}
```

Hierarchy Level

```
[edit chassis high-availability peer-id]
```

```
[edit chassis high-availability peer-id peer-forwarding-ip]
```

Description

Configure Bidirectional Forwarding Detection (BFD) options for the peer node.

Options

detection-time

Specify BFD failure detection time.

- threshold—Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

  Range: 1 through 255,000 milliseconds

minimum-interval

Configure the minimum interval at which the device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session

- Range: 1 through 255,000 milliseconds

minimum-receive-interval

(Optional) Configure the minimum interval after which the local device must receive a reply from a neighbor with which it has established a BFD session.

- Range: 1 through 255,000 milliseconds

multiplier

(Optional) Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

- Range: 1 through 255
- Default: 3

transmit-interval

The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum time that it requires between packets sent from its peer; the receive interval is not negotiated between peers.

- minimum-interval *milliseconds*—Configure the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.
  - Range: 1 through 255
- threshold *milliseconds*—Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.
  - Range: 0 through 4,294,967,295 (2^32 – 1)

  The threshold value specified in the threshold statement must be greater than the value specified in the minimum-interval statement for the transmit-interval statement.

version

BFD protocol version number

- 0—BFD version 0 (deprecated)
- 1—BFD version 1
- automatic—Choose BFD version automatically

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Multinode High Availability

Example: Configure Multinode High Availability in a Layer 3 Network

lldp

IN THIS SECTION

- Syntax | 9851
- Hierarchy Level | 9852
- Description | 9852
- Default | 9853
- Options | 9853
- Required Privilege Level | 9856
- Release Information | 9856

Syntax

```
lldp {
    advertisement-interval seconds;
    (disable | enable);
    hold-multiplier number;
{
    chassis-id (chassis-id-type | chassis-id-value)
```

```
    }
interface (all | [interface-name]) {
        (disable | enable);
        power-negotiation <(disable | enable)>;
        tlv-filter;
        tlv-select;
        trap-notification (disable | enable);
    }
    lldp-configuration-notification-interval seconds;
    management-address ip-management-address;;
    mau-type;
    netbios-snooping;
    no-tagging;
    neighbour-port-info-display (port-description | port-id);
    port-description-type (interface-alias | interface-description);
    port-id-subtype (interface-name | locally-assigned);
    ptopo-configuration-maximum-hold-time seconds;
    ptopo-configuration-trap-interval seconds;
    tlv-filter;
    tlv-select;
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable | no-world-
readable>;
        flag flag <disable>;
    }
    transmit-delay (LLDP) seconds;
    vlan-name-tlv-option (name | vlan-id);
}
```

Hierarchy Level

```
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Configure Link Layer Discovery Protocol (LLDP). The switch uses LLDP to advertise its identity and capabilities on a LAN, as well as to receive information about other network devices. LLDP is defined in the IEEE standard 802.1AB-2005.

NOTE: The transmit-delay and netbios-snooping options are not available on QFabric systems.

NOTE: On EX4300 switches, LLDP cannot be configured on the me0 or vme interface. Issuing the command `set protocols lldp interface me0` generates the following error message:

```
error: name: 'me0': Invalid interface
error: statement creation failed: interface
```

Issuing the command `set protocols lldp interface vme` generates the following error message:

```
error: name: 'vme': Invalid interface
error: statement creation failed: interface
```

Default

LLDP is disabled. If you configure LLDP for all interfaces, you can later disable a particular interface.

NOTE: The `interface-name` must be the physical interface and not a logical interface (unit).

Options

advertisement-interval *seconds*

Specify the frequency at which LLDP advertisements are sent. This value is also used in combination with the hold-multiplier value to determine the length of time LLDP information is held before it is discarded.

The `advertisement-interval` value must be greater than or equal to four times the `transmit-delay` value, or an error will be returned when you attempt to commit the configuration.

NOTE: The default value of `transmit-delay` is 2 seconds. If you configure `advertisement-interval` as less than 8 seconds and you do not configure a value for `transmit-delay`, the value of `transmit-delay` is automatically changed to 1 second to satisfy the requirement that

the `advertisement-interval` value be greater than or equal to four times the `transmit-delay` value.

- **Default:** 30 seconds
- **Range:** 5 through 32768 seconds

disable | enable

Disable or enable LLDP on the device.

- **Default:** If you do not configure LLDP, it is disabled on the device.

hold-multiplier *number*

Specify the multiplier used in combination with the advertisement-interval value to determine the length of time LLDP information is held before it is discarded.

- **Range:** 2 through 10
- **Default:** 4 (or 120 seconds with the default of 30 seconds for advertisement-interval)

lldp-configuration-notification-interval *seconds*

Specify how often SNMP trap notifications are generated as a result of LLDP database changes.

- **Range:** 5 through 3600 seconds
- **Default:** Disabled

management-address *ip-management-address;*

Specify the management address to be used in LLDP Management Address type, length, and value (TLV) messages. The Management Address TLV typically contains the IPv4 or IPv6 management addresses of the local system. Only out-of-band management addresses can be used for the management-address. Other remote managers can use this address to obtain information related to the local device.

- **Default:** The LLDP Management Address TLV uses the IP address of the switch's management Ethernet interface (me0), or the IP address of the virtual management Ethernet (VME) interface if the switch is a Virtual Chassis member.

mau-type

(EX4300, EX9200, and EX9250 switches only) Configure the switch to advertise information about the medium attachment unit (MAU) type. The MAU is a transceiver that interconnects the attachment unit interface (AUI) port on an attached host computer to an Ethernet cable. MAU types are defined in the IEEE 802.3 standard.

The MAU type is included in the MAC/PHY Configuration Status type, length, and value (TLV) message. TLVs are used by LLDP-capable devices to transmit information to neighbor devices. The MAC/PHY Configuration Status TLV is an

organizationally defined TLV that advertises information about the physical interface. In addition to the MAU type, the MAC/PHY Configuration Status TLV also includes information such as autonegotiation status, support, and advertised capabilities.

The MAU type cannot be changed by configuration; however, you must configure the `mau-type` statement to include the MAU type value in the MAC/PHY Configuration Status TLV.

- **Default:** If the `mau-type` statement is not configured, the MAU type field of the MAC/PHY Configuration Status TLV contains the value `Unknown`.

netbios-snooping

(EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, EX6210, EX8208, and EX8216 switches only) Enable NetBIOS snooping to learn information about NetBIOS hosts that are connected to the switch.

no-tagging

(EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, EX6210, EX8208, and EX8216 switches only) Configure the switch to send LLDPDUs without including VLAN tags on the interfaces on which VLAN tagging is enabled (tagged interfaces).

- **Default:** Interfaces for which VLAN tagging is enabled include a VLAN tag (tag 0) in LLDPDUs if the `no-tagging` option is not configured.

neighbour-port-info-display (port-description | port-id)

Configure the type of LLDP neighbor port information that the device displays in the `Port info` field in the output of the `show lldp neighbors` CLI command.

Devices in a network use LLDP to learn about and identify neighbor devices. LLDP-capable devices transmit information in type, length, and value (TLV) messages to neighbor devices.

The `Port info` field of the `show lldp neighbors` command displays the port information received from LLDP neighbors. This information is sent from the LLDP neighbor to the device in a type, length, and value (TLV) message. You can use the `neighbor-port-info-display` statement to configure the device to display the information contained in either the Port Description TLV or the Port Identification TLV.

- **Values:** Configure one of the following:

  - port-description—Display the information from the Port Description TLV in the `Port info` field of the `show lldp neighbors` CLI command.

    The Port Description TLV contains the textual description of the logical unit or the port. The description for the logical unit is used, if available; otherwise,

the description for the physical interface (port) is used. For example, LAG member interfaces do not contain a logical unit; therefore, only the description configured on the physical interface is used.

- port-id—Display the information from the Port Identification TLV in the `Port info` field of the `show lldp neighbors` CLI command.

  The Port Identification TLV contains the identifier for the neighbor port. The SNMP index of the interface is used as the port identifier.

- **Default:** port-description—The information contained in the Port Description TLV is displayed in the `Port info` field.

port-description-type (interface-alias | interface-description)

Configure the value to be used to generate the Port Description TLV that the device advertises to neighbors.

- **Values:** Configure one of the following:

  - interface-alias—Use the *ifAlias* MIB object value to generate the port description TLV. The LLDP MIB variable *lldpLocPortDesc* then contains the same value as *ifAlias*, which is the same as the description of the interface.

  - interface-description—Use the *ifDescr* MIB object value to generate the port description TLV. The LLDP MIB variable *lldpLocPortDesc* then contains the same value as *ifDescr*, which is the same as the interface name.

- **Default:** interface-alias—The `interface-alias` value is same as the description of an interface configured with `set interface` *`name`* `description` *`description`* `command`.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`management-address` introduced in Junos OS Release 9.5.

`netbios-snooping` introduced in Junos OS Release 11.1 for EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, EX6210, EX8208, and EX8216 switches.

`port-description-type` introduced in Junos OS Release 13.3R5, 14.2R2, 14.1R4, and 12.3R9.

`no-tagging` introduced in Junos OS Release 14.1X53-D10 for EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, EX6210, EX8208, and EX8216 switches.

`neighbour-port-info-display` introduced in Junos OS Release 14.1X53-D40 and Release 15.1R5 and Release 16.1R3.

`mau-type` introduced in Junos OS Release 15.1 for EX4300, EX9200, and EX9250 switches.

`chassis-id` introduced in Junos OS Evolved, Release 23.3R1.

RELATED DOCUMENTATION

Configuring LLDP

show lldp | 21942

Configuring LLDP (CLI Procedure)

Understanding LLDP

Understanding LLDP and LLDP-MED on EX Series Switches

Configuring NetBIOS Snooping (CLI Procedure)

lldp-med (Ethernet Switching)

IN THIS SECTION

- Syntax | 9858
- Hierarchy Level | 9858
- Description | 9858
- Default | 9859
- Options | 9859
- Required Privilege Level | 9859
- Release Information | 9859

Syntax

```
lldp-med {
    fast-start fast-start;
    interface                    name {
        (disable | enable);
        location {
            civic-based {
                ca-type  name {
                    ca-value  ca-value;
                }
                country-code country-code;
                what what;
            }
            co-ordinate {
                lattitude latitude;
                longitude longitude;
            }
            elin elin;
        }
        tlv-filter;
        tlv-select;
    }
    tlv-filter;
    tlv-select;
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure Link Layer Discovery Protocol–Media Endpoint Discovery. LLDP-MED is an extension of LLDP. The device uses LLDP-MED to support device discovery of VoIP telephones and to create location databases for these telephone locations for emergency services. LLDP-MED is defined in the standard ANSI/TIA-1057 by the Telecommunications Industry Association (TIA).

Default

Disabled.

Options

fast-start *number*

Configure the number of LLDP-MED advertisements sent from the device in the first second after it has detected an LLDP-MED device (such as an IP telephone).

- **Range:** 1 through 10 advertisements
- **Default:** 3 advertisements

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show lldp | 21942

Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch

Configuring LLDP-MED (CLI Procedure)

Understanding LLDP and LLDP-MED on EX Series Switches

lldp-priority

IN THIS SECTION

- Syntax | 9860
- Hierarchy Level | 9860
- Description | 9860
- Required Privilege Level | 9860
- Release Information | 9860

Syntax

```
lldp-priority;
```

Hierarchy Level

```
[edit poe],
[edit poe fpc (all | slot-number)]
```

Description

Configure the switch to assign interfaces the power priority provided by the powered device by using Link Layer Discovery Protocol (LLDP) power negotiation rather than the power priority configured on the switch interface.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring PoE Interfaces on EX Series Switches

lmi (Ethernet OAM)

IN THIS SECTION

- Syntax | 9861
- Hierarchy Level | 9862
- Description | 9862
- Options | 9862
- Required Privilege Level | 9862
- Release Information | 9862

Syntax

```
lmi {
    status-counter count;
    polling-verification-timer value;
    interface name {
        uni-id uni-name;
        status-counter number;
        polling-verification-timer value;
        evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
        evc evc-name {
            default-evc;
            vlan-list vlan-id-list;
        }
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet]
```

Description

On routers with `ge`, `xe`, or `ae` interfaces, configure an OAM Ethernet Local Management Interface (E-LMI).

> **NOTE**: On MX Series routers, E-LMI is supported on Gigabit Ethernet (`ge`), 10-Gigabit Ethernet (`xe`), and Aggregated Ethernet (`ae`) interfaces configured on MX Series routers with DPC only.

Options

`status-counter` *`count`*—Status counter (N393), defaults to 4.

`interface` *`name`*—Polling verification timer (T392), defaults to 15 seconds.

`uni-id` *`uni-name`*—(Optional) Defaults to the physical interface name.

`status-counter` *`number`*—(Optional) Defaults to a global value.

`polling-verification-timer` *`value`*—(Optional) Defaults to a global value.

`evc-map-type (all-to-one-bundling | bundling | service-multiplexing)`—Specify the Ethernet virtual connection (EVC) map type.

`evc` *`evc-name`*—Specify the name of the EVC.

`default-evc`—Set the specified EVC as the default EVC.

`vlan-list` *`vlan-id-list`*—Specify a group of VLANs to assign to the EVC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Ethernet Local Management Interface

evcs | 6518

lmi-type

IN THIS SECTION

- Syntax | 9863
- Hierarchy Level | 9863
- Description | 9863
- Options | 9863
- Required Privilege Level | 9864
- Release Information | 9864

Syntax

```
lmi-type (ansi | itu);
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

Set the Frame Relay Local Management Interface (LMI) type.

Options

`ansi`—Use American National Standards Institute (ANSI) T1.167 Annex D LMIs.

`itu`—Use ITU Q933 Annex A LMIs.

- **Default:** `itu`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring Keepalives on Link Services Physical Interfaces*

lmp-control-channel

IN THIS SECTION

- Syntax | 9864
- Hierarchy Level | 9865
- Description | 9865
- Options | 9865
- Required Privilege Level | 9865
- Release Information | 9865

Syntax

```
lmp-control-channel control-channel-interface {
    remote-address ip-address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name],
[edit protocols link-management peer peer-name]
```

Description

Specify the Link Management Protocol (LMP) control channel interface for the peer.

Options

control-channel-interface—Name of the control channel interface.

The remaining statement is described separately in this chapter.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring the LMP Control Channel Interface for the Peer

lmp-protocol

IN THIS SECTION

- Syntax | 9866
- Hierarchy Level | 9866

- Description | 9866
- Options | 9866
- Required Privilege Level | 9866
- Release Information | 9867

Syntax

```
lmp-protocol {
    hello-dead-interval milliseconds;
    hello-interval milliseconds;
    passive;
    retransmission-interval milliseconds;
    retry-limit number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name],
[edit protocols link-management peer peer-name]
```

Description

Configure attributes of Link Management Protocol (LMP) to establish and maintain the LMP control channel for the peer.

Options

The statements are described separately in this chapter.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring LMP Peers

load

IN THIS SECTION

- Syntax | 9867
- QFX Series | 9867
- VMX Series | 9868
- Description | 9868
- Options | 9868
- Required Privilege Level | 9869
- Sample Output | 9869
- Release Information | 9870

Syntax

```
load (factory-default | merge | override | patch | replace | set | update) (filename | terminal)
<json>
<relative>
```

QFX Series

```
load (dhcp-snooping filename)
```

VMX Series

```
load (<dhcp-security-snoop | dhcpv6-security-snoop>filename)
```

Description

Load a configuration from an ASCII configuration file, from terminal input, or from the factory default. Your current location in the configuration hierarchy is ignored when the load operation occurs.

For information on valid filename and URL formats, see Format for Specifying Filenames and URLs in Junos OS CLI Commands.

NOTE: `load` can be run from configuration mode only.

Options

`dhcp-snooping`—(QFX Series switches) Loads DHCP snooping entries.

`dhcp-security-snoop`—(VMX Series switches) Loads DHCP snooping entries.

`dhcpv6-security-snoop`—(VMX Series switches) Loads DHCPv6 snooping entries.

`factory-default`—Loads the factory configuration. The factory configuration contains the manufacturer's suggested configuration settings. The factory configuration is the first configuration for the router or switch and is loaded when the router or switch is first installed and powered on. The `factory-default` option cannot be combined with other options.

NOTE: To load the factory default configuration, you must first "unprotect" on page 30428 any protected hierarchies in the configuration.

`filename`—Name of the file to load. For information about specifying the filename, see *Viewing Files and Directories on a Juniper Networks Device*.

`json`—(Optional) Load configuration data that uses JavaScript Object Notation (JSON) format. This option can be used with the `merge`, `override`, or `update` options.

`merge`—Combine the configuration that is currently shown in the CLI with the configuration.

`override`—Discard the entire configuration that is currently shown in the CLI and load the entire configuration. Marks every object as changed.

`patch`—Change part of the configuration and mark only those parts as changed.

`relative`—(Optional) Load the new configuration data relative to the current edit point in the configuration hierarchy.

`replace`—Look for a `replace` tag in *`filename`*, delete the existing statement of the same name, and replace it with the configuration.

`set`—Merge a set of commands with an existing configuration. This option executes the configuration instructions line by line as they are stored in a file or from a terminal. The instructions can contain any configuration mode command, such as `set`, `edit`, `exit`, and `top`.

`terminal`—Use the text you type at the terminal as input to the configuration. Type Ctrl+d to end terminal input.

`update`—Discard the entire configuration that is currently shown in the CLI, and load the entire configuration. Marks changed objects only.

NOTE: If you are using Junos OS in a Common Criteria environment, system log messages are created whenever a `secret` attribute is changed (for example, password changes or changes to the RADIUS shared secret). These changes are logged during the following configuration load operations:

```
load merge
load replace
load override
load update
```

For more information, see the *Secure Configuration Guide for Common Criteria and Junos-FIPS*.

Required Privilege Level

NOTE: The `load` commands for `dhcp-security-snoop` and `dhcpv6-security-snoop` are restricted to admin only.

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Sample Output

The following is an example of a load scenario using Secure Copy (scp).

To Load a Configuration File Using Secure Copy Protocol (scp) with 'source-address' and 'routing-instance' options

To load a configuration file using the scp command with the `source-address` and `routing-instance` options, enter the following command:

```
root@host# load merge scp://user@hostname/path/filename source-address address routing-instance
instance-name
```

The scp options `source-address` and `routing-instance` are supported for **load override**, **load patch**, **load replace**, **load set**, and **load update** options also.

Release Information

Command introduced before Junos OS Release 7.4.

`json` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Loading a Configuration from a File or the Terminal

load access-security slaac-snooping persistent-file

IN THIS SECTION

- Syntax | 9871
- Description | 9871
- Options | 9871
- Required Privilege Level | 9871
- Output Fields | 9871
- Sample Output | 9871
- Release Information | 9871

Syntax

```
load access-security slaac-snooping persistent-file (local_pathname | remote_URL)
```

Description

Load a SLAAC snooping database file to ensure that IPv6-MAC bindings persist through the device reboots by specifying a local pathname or a remote URL.

Options

(*local_pathname* | *remote_URL*)	Specify either a local pathname or a remote URL for the SLAAC snooping database file.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Output Fields

This command produces no output.

Sample Output

load access-security slaac-snooping persistent-file

```
user@switch> load access-security slaac-snooping persistent-file
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show access-security slaac-snooping binding | 17003

load-balance (ecmp)

IN THIS SECTION

- Syntax | 9872
- Hierarchy Level | 9872
- Description | 9872
- Options | 9873
- Required Privilege Level | 9873
- Release Information | 9873

Syntax

```
load-balance {
                        adaptive;
                        consistent-hash;
                        destination-ip-only;
                        per-packet;
                        profile1;
                        profile2;
                        random;
                        source-ip-only;
                       }
```

Hierarchy Level

```
[edit policy-options policy-statement policy-name then]
```

Description

Specify the type of load balancing of an equal-cost multipath (ECMP) in the forwarding table.

By default, Junos OS uses a hashing method based only on the destination address to elect a forwarding next hop when multiple equal-cost paths are available. As a result, when multiple routers or switches

share the same set of forwarding next hops for a given destination, they can elect the same forwarding next hop.

Options

adaptive Adaptive load balancing. Dynamically manages and distributes traffic loads across multiple ECMP paths by monitoring packet and byte counts for each hash bucket and employing a feedback mechanism to detect and correct overload conditions.

consistent-hash Prefix consistent load-balancing. Consistent load balancing ensures that if an ECMP member (that is, a path) fails, only flows flowing through the failed member are redistributed to other active ECMP members.

destination-ip-only Destination based IP load-balancing. In equal-cost multipath, (ECMP) per-subscriber aware environments such as content service providers who service residential customers, traffic in both directions within the service provider network should always pass through the content servers that maintain the subscriber state information for a given subscriber. This is accomplished by calculating the load balancing hash based solely on source address for traffic coming into the service provider network and calculating the load balancing hash based solely on the destination address for traffic leaving the service provider network.

per-packet Load-balance on a per-packet basis.

profile1 Per-packet load balance with hash keys in profile1.

profile2 Per-packet load balance with hash keys in profile2.

random Load-balance using packet random spray.

source-ip-only Source based IP load-balancing

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

load-balance

IN THIS SECTION

- Syntax | 9874
- Hierarchy Level | 9874
- Description | 9875
- Options | 9875
- Required Privilege Level | 9875
- Release Information | 9875

Syntax

```
load-balance {
    adaptive{
        pps;
        scan-interval multiple;
        tolerance percentage;
    }
    no-adaptive;
    per-packet;
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options],
[edit dynamic-profiles name interfaces name logical-tunnel-options],
[edit dynamic-profiles name interfaces interface-range name aggregated-ether-options],
[edit dynamic-profiles name interfaces interface-range name logical-tunnel-options],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-ether-options],
[edit dynamic-profiles name logical-systems name interfaces name logical-tunnel-options],
[edit dynamic-profiles name logical-systems name interfaces interface-range name aggregated-
ether-options],
[edit dynamic-profiles name logical-systems name interfaces interface-range name logical-tunnel-
options],
```

```
[edit interfaces name aggregated-ether-options],
[edit interfaces name logical-tunnel-options],
[edit interfaces interface-range              name aggregated-ether-options],
[edit interfaces interface-range              name logical-tunnel-options]
```

Description

Load-balances the received traffic across all the available paths of aggregated Ethernet bundles for better link utilization.

Options

adaptive (MX Series and PTX Series) Corrects a genuine traffic imbalance by using a feedback mechanism to distribute the traffic across the links of an Aggregated Ethernet bundle.

no-adaptive (MX Series and PTX Series) Disables the adaptive load-balancing solution configured to distribute traffic by using a feedback mechanism.

per-packet (MX Series only) Randomly sprays packets to the aggregate next hops in a round-robin manner to avoid traffic imbalance.

Required Privilege Level

interface - To view statement in the configuration.

interface-control - To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Aggregated Ethernet Load Balancing

Example: Configuring Aggregated Ethernet Load Balancing

load-balance-algorithm

IN THIS SECTION

- Syntax | 9876
- Hierarchy Level | 9876
- Description | 9876
- Options | 9876
- Required Privilege Level | 9877
- Release Information | 9877

Syntax

```
load-balance-algorithm (simple | enode-based | flogi-based);
```

Hierarchy Level

```
[edit fc-fabrics fabric-name proxy]
```

Description

Set the load-balancing algorithm that the QFX Series uses to distribute FCoE sessions (FLOGI and FDISC sessions from the FCoE devices in the Ethernet network) among the NP_Port links to the FC switch.

NOTE: Changing the load-balancing algorithm when FCoE sessions are running forces the FCoE sessions to log out, then log in again.

Options

simple—Load balancing is based on the weighted utilization (load) of the NP_Ports connected to an FC fabric. Each new FLOGI or FDISC is assigned to the least-loaded link. When a link load rebalance occurs, the system minimizes disruption by using an algorithm to log out only the sessions that need to be moved to other links to balance the link load. To further minimize disruption, the algorithm logs out the

sessions with the fewest dependencies (for example, FDISC sessions are logged out before FLOGI sessions). When the sessions log in again, they are placed on NP_Port interfaces in a manner that balances the link loads. This is the default load-balancing algorithm.

enode-based—Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions (VN_Port sessions) associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. When a link load rebalance occurs, the system logs off all sessions. The sessions log in again and are placed on NP_Port interfaces in a balanced manner.

flogi-based—FLOGI-based load balancing is similar to ENode-based load balancing, but the behavior when the loads are rebalanced is different. Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. When a link load rebalance occurs, the system minimizes disruption by using an algorithm to log out only the sessions that need to be moved to other links to balance the link load. When the logged out sessions log back in, they are placed on NP_Port interfaces in a manner that balances the link loads.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Defining the Proxy Load-Balancing Algorithm

Example: Configuring Automated Fibre Channel Interface Load Rebalancing

Simulating On-Demand Fibre Channel Link Load Rebalancing (Dry Run Test)

Monitoring Fibre Channel Interface Load Balancing

Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric

load-balance (Forwarding Options)

IN THIS SECTION

- Syntax | 9878
- Hierarchy Level | 9878
- Description | 9879
- Options | 9880
- Required Privilege Level | 9880
- Release Information | 9880

Syntax

```
load-balance {
    indexed-load-balance;
    per-flow {
        hash-seed;
    }
    per-prefix {
        hash-seed number;
    }
}
```

Hierarchy Level

```
[edit forwarding-options],
[edit logical-systems logical-system-name forwarding-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options],
[edit routing-instances routing-instance-name forwarding-options]
```

Description

Enable per-prefix or per-flow load balancing so that the router or switch elects a next hop independently of the route selected by other routers or switches.

For the active route, when there are multiple equal-cost paths to the same destination, by default, Junos OS chooses in a random fashion one of the next-hop addresses to install into the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is chosen again, also in a random fashion.

You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This is called per-packet load balancing. You can use load balancing to spread traffic across multiple paths between routing devices. The behavior of the per-packet load-balancing function varies according to the version of the Internet Processor ASIC in the routing device.

On routing devices with an Internet Processor I ASIC, when per-packet load balancing is configured, traffic between routing devices with multiple paths is spread in a random fashion across the available interfaces. The forwarding table balances the traffic headed to a destination, transmitting packets in round-robin fashion among the multiple next hops (up to a maximum of eight equal-cost load-balanced paths). The traffic is load-balanced on a per-packet basis.

Per-packet load distribution uses a hashing algorithm that distributes packets over equal-cost links. The algorithm is designed to distribute packets to prevent any single link from being saturated. However, per-packet load balancing offers no guarantee of equal distribution of traffic over equal-cost links, nor does it guarantee that increasing the number of Internet flows creates a better hash distribution.

On routing devices with the Internet Processor II ASIC and T Series Internet Processor II ASIC, when per-packet load balancing is configured, traffic between routing devices with multiple paths is divided into individual traffic flows (up to a maximum of 16 equal-cost load-balanced paths). On some platforms, you can increase the number of paths by using the `chassis maximum-ecmp` statement. With this statement, you can change the maximum number of equal-cost load-balanced paths to 32 or 64. Packets for each individual flow are kept on a single interface. To recognize individual flows in the transit traffic, the routing device examines each of the following:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Source interface index
- Type of service (ToS)

The routing device recognizes packets in which all of these parameters are identical, and it ensures that these packets are sent out through the same interface. This prevents problems that might otherwise occur with packets arriving at their destination out of their original sequence.

Load balancing is not supported on management and internal Ethernet (**fxo**) interfaces because this type of interface cannot handle the routing process. On **fxp** interfaces, you cannot configure multiple next hops and enable load balancing.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for per-flow load balancing introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

- Example: Load Balancing BGP Traffic
- *Configuring Per-Flow Load Balancing Based on Hash Values*
- *Configuring Per-Prefix Load Balancing*

load-balance-group

IN THIS SECTION

- Syntax | **9881**
- Hierarchy Level | **9881**
- Description | **9881**

- Options | 9881
- Required Privilege Level | 9881
- Release Information | 9881

Syntax

```
load-balance-group group-name {
    next-hop-group [ group-names ];
}
```

Hierarchy Level

```
[edit firewall]
```

Description

Configure a load-balance group.

Options

group-name—Name of load-balance group.

group-names—Name of next-hop groups to include in the load-balance group set.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Load-Balance Groups

Routing Policies, Firewall Filters, and Traffic Policers User Guide

load-balance-label-capability

IN THIS SECTION

- Syntax | 9882
- Hierarchy Level | 9882
- Description | 9882
- Required Privilege Level | 9883
- Release Information | 9883

Syntax

```
load-balance-label-capability;
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Enables the router to push and pop the load balancing label and causes LDP and RSVP to advertise the entropy label TLV to neighboring routers.

The `load-balance-label-capability` and `no-load-balance-label-capability` statements at the `[edit forwarding-options]` hierarchy level are mutually exclusive, and at a given point in time, configuring one statement overrides the other.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring the Entropy Label for LSPs

no-load-balance-label-capability | 11495

load-balance profile1/profile2

IN THIS SECTION

- Syntax | 9883
- Hierarchy Level | 9884
- Description | 9884
- Default | 9884
- Options | 9884
- Required Privilege Level | 9884
- Release Information | 9884

Syntax

```
load-balance profile1/profile2
```

Hierarchy Level

```
[edit policy-options policy-statement policy name then]
```

Description

Apply any one of two predefined load balance profiles (profile1 or profile2) for the matching route.

Default

- Default profile or global profile - a route without a policy-options specified load-balance profile will be grouped to the default profile or global profile.

Options

- profile1
- profile2

Required Privilege Level

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos 23.2

RELATED DOCUMENTATION

- enhanced-hash-key profile1/profile2
- hash-seed (profile)

load-balance (Protocols RSVP)

IN THIS SECTION

- Syntax | 9885
- Hierarchy Level | 9885
- Description | 9885
- Options | 9885
- Required Privilege Level | 9885
- Release Information | 9886

Syntax

```
load-balance {
    bandwidth;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Load-balance traffic between RSVP LSPs.

Options

`bandwidth`—Load-balance traffic between RSVP LSPs based on the bandwidth configured for each LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Load Balancing Across RSVP LSPs

load-balance-stateful (Aggregated Ethernet Interfaces)

IN THIS SECTION

- Syntax | 9886
- Hierarchy Level | 9887
- Description | 9887
- Options | 9887
- Required Privilege Level | 9887
- Release Information | 9887

Syntax

```
load-balance-stateful {
    per-flow;
        rebalance interval;
    load-type (low | medium | large);
}
```

Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options]
```

Description

Define the capability to perform uniform load balancing and also perform rebalancing is introduced on MX Series routers with MPCs, except MPC3Es and MPC4Es. Rebalancing is not supported when load-balancing is skewed or distorted owing to a change in the number of flows. The mechanism to record and retain states for the flows and distribute the traffic load accordingly is added. As a result, for m number of flows, they are distributed among n member links of a LAG bundle or among the unilist of next-hops in an ECMP link. This method of splitting the load among member links is called stateful load balancing and it uses 5-tuple information (source and destination addresses, protocol, source and destination ports). Such a method can be mapped directly to the flows, or to a precompute hash based on certain fields in the flow. As a result, the deviation observed on each child link is reduced.

Options

stateful Define the stateful load-distribution mechanism for traffic flows on aggregated Ethernet interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces

load-balancing-options (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 9888
- Hierarchy Level | 9889
- Description | 9889
- Required Privilege Level | 9890
- Release Information | 9890

Syntax

```
load-balancing-options {
    high-availability-options {
        (many-to-one | one-to-one) {
            preferred-backup preferred-backup;
        }
    }
    member-failure-options {
        drop-member-traffic {
            rejoin-timeout rejoin-timeout;
        }
        redistribute-all-traffic {
            enable-rejoin;
        }
    }
    hash-keys {
        egress-key (destination-ip | source-ip);
        ingress-key (destination-ip | source-ip);
    }
    member-interface interface-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure the high availability (HA) options for the aggregated multiservices (AMS) interface.

Many-to-one (N:1) high availability mode for service applications like Network Address Translation (NAT) is supported. In the case of N:1 high availability mode, one services PIC is the backup (in hot standby mode) for one or more (N) active services PICs. If one of the active services PICs goes down, then the backup replaces it as the active services PIC. When the failed PIC comes back online, it becomes the new backup. This is called *floating backup mode*. In an N:1 (stateless) configuration, traffic states and data structures are not synchronized between active PICs and the backup PIC.

You can also configure a one-to-one (1:1) high availability mode. In the 1:1 configuration, a single interface is configured as the backup for another single active interface. If the active interface goes down, the backup interface replaces it as the active interface. A 1:1 (stateful) configuration synchronizes traffic states and data structures between the active services PIC and the backup services PIC. This is required for IPsec connections. One-to-one high availability is supported on the MS-MPC but it is not supported for MX-SPC3 in this release.

Load-balancing might not be uniform among member interfaces in certain network deployments. The variance can be because of a misconfiguration, which causes the traffic itself not to be sufficiently randomly distributed, causing the hash keys to be ineffective (for example, the hash key is destination IP but all sessions have only source IP address). The variation can be within the expected range and the load balancing depends on the IP addresses chosen. The hash calculation performs a checksum on several bits of the IP address and not only on the last few lower significant bits of the IP address. In such a scenario, the load-balancing ratio can change, for instance, if the source IP address is changed from 20.0.0.0/24 to 20.0.1.0/24.

The distribution of traffic across member interfaces of an AMS interface is static load-balancing. Flows are load balanced based on a packet hash on parameters such as source IP or destination IP. Load-balancing effectiveness depends on the IP address or protocol diversity. For example, if the hash key is destination IP and all packets have the same destination, then all flows are directed to the same member. This is flow-level load balancing and not per packet. As a result, traffic between a pair of addresses may be 10,000 pps, whereas another pair of addresses may have 1 pps. The load of the former is not distributed among members. High availability is limited to stateless HA. When a backup interface takes over as an active interface, all flows are reestablished (for example. packets may undergo NAT processing differently after failover).

With a stateful firewall, static NAT as basic-nat44 or destination-nat44, and dynamic NAT as nat64, napt-44, dynamic-nat44, and with application layer gateways (ALGs) configured, NAT hairpinning is not

supported. Input direction for rule match to be applied is supported only for dynamic NAT types (NAT64, NAT44, and dynamic-NAT44). Service-set policies need to have input or input-output direction only. Flows on all active members are reset when the number of actives changes. The resetting of flows can be avoided at the cost of failed-member's traffic loss using certain options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

load-balancing-options (Service Set)

IN THIS SECTION

- Syntax | 9891
- Hierarchy Level | 9891
- Description | 9891
- Required Privilege Level | 9891
- Release Information | 9891

Syntax

```
load-balancing-options {
    hash-keys {
        egress-key (destination-ip | source-ip);
        ingress-key (destination-ip | source-ip);
    }
}
```

Hierarchy Level

```
[edit services service-set service-set-name interface-service]
```

Description

Configure the load-balancing options for aggregated multiservices (AMS) in service applications (Network Address Translation [NAT], stateful firewall, application-level gateway [ALG], HTTP header enrichment, and mobility). AMS for service applications can be used for load balancing with or without high availability (HA). Currently, load balancing is based on the configured hash keys.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 11.4.

RELATED DOCUMENTATION

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

load-interval

IN THIS SECTION

- Syntax | 9892
- Hierarchy Level | 9892
- Description | 9892
- Options | 9892
- Required Privilege Level | 9893
- Release Information | 9893

Syntax

```
load-interval seconds;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN logical interfaces, specify the interval used to calculate the average load on the network. By default, the average interface load is calculated every 60 seconds.

Options

seconds—Number of seconds at which the average load calculation is triggered.

- **Range:** 20 through 180, in 10-second intervals
- **Default:** 60 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

load-scripts-from-flash (Scripts)

IN THIS SECTION

- Syntax | 9893
- Hierarchy Level | 9893
- Description | 9894
- Default | 9894
- Required Privilege Level | 9894
- Release Information | 9894

Syntax

```
load-scripts-from-flash;
```

Hierarchy Level

```
[edit system scripts]
```

Description

Load commit, event, op, SNMP, and library scripts from the device's flash memory instead of the hard disk.

When you configure this statement, all commit, event, op, SNMP, and script library scripts are loaded from the **/config/scripts** directory in flash memory under the subdirectory appropriate to the script type. To ensure that the current versions of the scripts are executed, you must manually move scripts from the hard disk to flash memory. They are not moved automatically.

Default

If you do not include the `load-scripts-from-flash` statement, scripts are executed from the **/var/db/scripts** directory on the device's hard disk.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Storing Scripts in Flash Memory

scripts | **16392**

load-threshold

IN THIS SECTION

- Syntax | **9895**
- Hierarchy Level | **9895**
- Description | **9895**
- Options | **9895**

Required Privilege Level | 9895
Release Information | 9895

Syntax

```
load-threshold percent;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN logical interfaces, specify the bandwidth threshold percentage used for adding interfaces. Another link is added to the multilink bundle when the load reaches the threshold value you set. Specify a percentage between 0 and 100.

Options

percent—Bandwidth threshold percentage used for adding interfaces. When set to 0, all available channels are dialed.

- **Range:** 0 through 100 seconds
- **Default:** 100 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

load-type (Aggregated Ethernet Interfaces)

IN THIS SECTION

- Syntax | 9896
- Hierarchy Level | 9896
- Description | 9896
- Options | 9896
- Required Privilege Level | 9897
- Release Information | 9897

Syntax

```
load-type (low | medium | large);
```

Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options load-balance-stateful]
```

Description

Define the load-balancing type to inform the Packet Forwarding Engine regarding the appropriate memory pattern to be used for traffic flows. The approximate number of flows for effective load-balancing for each keyword is a derivative.

Options

low Define a low load-balancing method if the number of flows that flow on the specified aggregated Ethernet interface is less or minimal (between 1 and 100 flows).

medium Define a medium or moderate load-balancing method if the number of flows that flow on the specified aggregated Ethernet interface is relatively higher (between 100 and 1000 flows).

large Define a high load-balancing method if the number of flows that flow on the specified aggregated Ethernet interface is excessive or reaches the maximum supported flows (between 1000 and 10,000 flows).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces

loc-rib

IN THIS SECTION

- Syntax | 9897
- Hierarchy Level | 9898
- Description | 9898
- Options | 9898
- Required Privilege Level | 9898
- Release Information | 9899

Syntax

```
loc-rib {
```

```
        rib-filter <X>
        route-filter <Y>
        }
```

Hierarchy Level

```
[edit logical-systems name routing-instancesname routing-options bmp route-monitoring],
[edit logical-systems name routing-options bmp route-monitoring],
[edit logical-systems name tenants name routing-instances name routing-options bmp route-monitoring],
[edit routing-instances name routing-options bmp route-monitoring],
[edit routing-options bmp route-monitoring],
[edit tenants name routing-instances name routing-options bmp route-monitoring]
```

Description

Enable local routing information base (RIB) route monitoring messages.

Options

rib-filter *x* Specify a policy to accept or reject a set of RIBs. You can specify multiple rib lists to filter in or filter them out.

route-filter *y* Specify a policy with a set of match conditions for route filtering such as neighbor, prefixes etc.

> **NOTE**: A match condition specified in a policy might apply to all or a subset of RIBs. For example, the match condition neighbor is applicable to both adj-RIB-in and adj-RIB-out. However, the match condition does not apply to the local RIB.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

rib-list | 15783

local

IN THIS SECTION

- Syntax | 9899
- Hierarchy Level | 9900
- Description | 9900
- Required Privilege Level | 9900
- Release Information | 9900

Syntax

```
local {
    address address;
    disable;
    family (inet | inet6) anycast-pim;
    }
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
    hold-time seconds;
    override;
    priority number;
    process-non-null-as-null-register  ;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Configure the routing device's RP properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Local PIM RPs

Example: Configuring PIM Anycast With or Without MSDP

local (certificates)

IN THIS SECTION

- Syntax | 9901
- Hierarchy Level | 9901

- Description | 9901
- Options | 9901
- Required Privilege Level | 9902
- Release Information | 9902

Syntax

```
local certificate-name {
    certificate-key-string;
    load-key-file URL filename;
}
```

Hierarchy Level

```
[edit security certificates]
```

Description

Import a paired X.509 private key and authentication certificate, to enable Junos XML protocol client applications to establish Secure Sockets Layer (SSL) connections to the router or switch.

NOTE: For FIPS mode, the digital security certificates must be compliant with the National Institute of Standards and Technology (NIST) SP 800-131A standard.

Options

certificate-key-string	String of alphanumeric characters that constitute the private key and certificate.
certificate-name	Name that uniquely identifies the certificate.
load-key-file *URL filename*	File that contains the private key and certificate. It can be one of two types of values:

- Pathname of a file on the local disk (assuming you have already used another method to copy the certificate file to the router's or switch's local disk)
- URL to the certificate file location (for instance, on the computer where the Junos XML protocol client application runs)

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Importing SSL Certificates for Junos XML Protocol Support

local-address (Junos Telemetry Interface)

IN THIS SECTION

- Syntax | 9903
- Hierarchy Level | 9903
- Description | 9903
- Options | 9903
- Required Privilege Level | 9903
- Release Information | 9903

Syntax

```
local-address local-address;
```

Hierarchy Level

```
[edit services analytics export-profile export-profile-name]
```

Description

For legacy gRPC dialout connections, configure the local IPv4 or IPv6 address in an export profile.

Configuring the local address lets you bypass the current default way of addressing, which is that the outgoing interface IP address is used as the source address. You can configure the local address to be your router's specified IP address, interface address, or a loopback address.

You can only configure an IPv6 address as the local address for legacy remote procedure call (gRPC) dial-out connections,. If you try to commit the configuration with an export profile that contains `transport` type as `udp` instead of `grpc` an error is displayed.

Options

local-address	An IPv4 or IPv6 address.

Required Privilege Level

Edit

Release Information

Statement introduced in Junos OS Evolved 24.2R1 (for IPv6) on ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7900, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, and QFX5700E.

RELATED DOCUMENTATION

Using gRPC Dial-Out for Secure Telemetry Collection

export-profile (Junos Telemetry Interface) | 6669

routing-instance (Junos Telemetry Interface) | 15937

local-address (Protocols AMT)

IN THIS SECTION

- Syntax | 9904
- Hierarchy Level | 9904
- Description | 9905
- Default | 9905
- Options | 9905
- Required Privilege Level | 9905
- Release Information | 9905

Syntax

```
local-address ip-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols amt relay family inet],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols amt
relay family inet],
[edit protocols amt relay family inet],
[edit routing-instances routing-instance-name protocols amt relay family inet]
```

Description

Specify the local unique IP address to send in Automatic Multicast Tunneling (AMT) relay advertisement messages, for use as the IP source of AMT control messages, and as the source of the data tunnel encapsulation. The address can be configured on any interface in the system. Typically, the router's **lo0.0** loopback address is used for configuring the AMT local address in the default routing instance, and the router's **lo0.*n*** loopback address is used for configuring the AMT local address in VPN routing instances.

Default

None. The local address must be configured.

Options

ip-address—Unique unicast IP address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

local-address (Protocols BGP)

IN THIS SECTION

- Syntax | 9906
- Hierarchy Level | 9906

- Description | 9906
- Default | 9907
- Options | 9907
- Required Privilege Level | 9908
- Release Information | 9908

Syntax

```
local-address address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Specify the address of the local end of a BGP session. This address is used to accept incoming connections to the peer and to establish connections to the remote peer. When none of the operational interfaces are configured with the specified local address, a session with a BGP peer is placed in the idle state.

You generally configure a local address to explicitly configure the system's IP address from BGP's point of view. This IP address can be either an IPv6 or IPv4 address. Typically, an IP address is assigned to a loopback interface, and that IP address is configured here.

For internal BGP (IBGP) peering sessions, generally the loopback interface (lo0) is used to establish connections between the IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address, the IBGP peering session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peering session also goes up and down. Thus, the loopback interface provides fault tolerance in case the physical interface or the link goes down, if the device has link redundancy.

When a device peers with a remote device's loopback interface address, the local device expects BGP update messages to come from (be sourced by) the remote device's loopback interface address. The `local-address` statement enables you to specify the source information in BGP update messages. If you omit the `local-address` statement, the expected source of BGP update messages is based on the device's source address selection rules, which normally result in the egress interface address being the expected source of update messages. When this happens, the peering session is not established because a mismatch exists between the expected source address (the egress interface of the peer) and the actual source (the loopback interface of the peer). To ensure that the expected source address matches the actual source address, specify the loopback interface address in the `local-address` statement.

> **NOTE**: Although a BGP session can be established when only one of the paired routing devices has `local-address` configured, we strongly recommend that you configure `local-address` on both paired routing devices for IBGP and multihop EBGP sessions. The `local-address` statement ensures that deterministic fixed addresses are used for the BGP session end-points.

If you include the `default-address-selection` statement in the configuration, the software chooses the system default address as the source for most locally generated IP packets. For protocols in which the local address is unconstrained by the protocol specification, for example IBGP and multihop EBGP, if you do not configure a specific local address when configuring the protocol, the local address is chosen using the same methods as other locally generated IP packets.

Default

If you do not configure a local address, BGP uses the routing device's source address selection rules to set the local address.

Options

address—IPv6 or IPv4 address of the local end of the connection.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring Internal BGP Peering Sessions on Logical Systems

Example: Configuring Internal BGP Peer Sessions

Understanding Internal BGP Peering Sessions

router-id | **15902**

local-address (Protocols Link Management)

IN THIS SECTION

- Syntax | **9908**
- Hierarchy Level | **9909**
- Description | **9909**
- Options | **9909**
- Required Privilege Level | **9909**
- Release Information | **9909**

Syntax

```
local-address ip-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],
[edit logical-systems logical-system-name protocols link-management te-link te-link-name
interface interface-name],
[edit protocols link-management te-link te-link-name],
[edit protocols link-management te-link te-link-name interface interface-name]
```

Description

Specify the local IP address associated with the traffic engineering link. If you configure the local IP address, you must also configure the "remote-address" on page 14314 statement.

Options

local-address—Local IP address of the traffic engineering link.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Local IP Address for Traffic Engineering Links

Configuring the Local IP Address for Forwarding Adjacencies

remote-address (for LMP Traffic Engineering) | **14314**

local-address (Protocols MSDP)

IN THIS SECTION

- Syntax | 9910
- Hierarchy Level | 9910
- Description | 9911
- Options | 9911
- Required Privilege Level | 9911
- Release Information | 9911

Syntax

```
local-address address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name],
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
peer address],
[edit protocols msdp],
[edit protocols msdp group group-name],
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit routing-instances routing-instance-name protocols msdp],
[edit routing-instances routing-instance-name protocols msdp group group-name],
```

```
[edit routing-instances routing-instance-name protocols msdp group group-name peer address],
[edit routing-instances routing-instance-name protocols msdp peer address]
```

Description

Configure the local end of an MSDP session. You must configure at least one peer for MSDP to function. When configuring a peer, you must include this statement. This address is used to accept incoming connections to the peer and to establish connections to the remote peer.

Options

address—IP address of the local end of the connection.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP in a Routing Instance

local-address (Protocols PIM)

IN THIS SECTION

- Syntax | 9912
- Hierarchy Level | 9912
- Description | 9912
- Options | 9912

Required Privilege Level | 9912
Release Information | 9912

Syntax

```
local-address address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp local family (inet | inet6) anycast-pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp local family (inet | inet6) anycast-pim],
[edit protocols pim rp local family (inet | inet6) anycast-pim],
[edit routing-instances routing-instance-name protocols pim rp local family (inet | inet6) anycast-
pim]
```

Description

Configure the routing device local address for the anycast rendezvous point (RP). If this statement is omitted, the router ID is used as this address.

Options

address—Anycast RP IPv4 or IPv6 address, depending on **family** configuration.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring PIM Anycast With or Without MSDP

local-address (Routing Options)

IN THIS SECTION

- Syntax | 9913
- Hierarchy Level | 9913
- Description | 9913
- Options | 9914
- Required Privilege Level | 9914
- Release Information | 9914

Syntax

```
local-address address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast backup-pe-group group-name],
[edit logical-systems logical-system-name routing-options multicast backup-pe-group group-name],
[edit routing-instances routing-instance-name routing-options multicast backup-pe-group group-
name],
[edit routing-options multicast backup-pe-group group-name]
```

Description

Configure the address of the local PE for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

Options

address—Address of local PEs in the backup group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement added to the multicast hierarchy in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring Ingress PE Redundancy

local (Analytics Collector)

IN THIS SECTION

- Syntax | 9914
- Hierarchy Level | 9915
- Description | 9915
- Required Privilege Level | 9915
- Release Information | 9915

Syntax

```
local {
    file filename {
```

```
        size size;
        files number;
    }
}
```

Hierarchy Level

```
[edit services analytics collector]
```

Description

Configure a local file for logging network analytics (queue and traffic) statistics.

NOTE: The `local` statement is not available in Junos OS Releases prior to 13.2X51-D15.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

collector (Analytics) | 4412

local-as

IN THIS SECTION

- Syntax | 9916
- Hierarchy Level | 9916
- Description | 9917
- Options | 9917
- Required Privilege Level | 9918
- Release Information | 9919

Syntax

```
local-as autonomous-system <loops number> <private | alias> <no-prepend-global-as>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Specify the local autonomous system (AS) number. An AS is a set of routing devices that are under a single technical administration and generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices.

Internet service providers (ISPs) sometimes acquire networks that belong to a different AS. When this occurs, there is no seamless method for moving the BGP peers of the acquired network to the AS of the acquiring ISP. The process of configuring the BGP peers with the new AS number can be time-consuming and cumbersome. In this case, it might not be desirable to modify peer arrangements or configuration. During this kind of transition period, it can be useful to configure BGP-enabled devices in the new AS to use the former AS number in BGP updates. This former AS number is called a *local AS*.

NOTE: If you are using BGP on the routing device, you must configure an AS number before you specify the `local-as` number.

In Junos OS Release 9.1 and later, the AS numeric range in plain-number format is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <*16-bit high-order value in decimal*>.<*16-bit low-order value in decimal*>. For example, the 4-byte AS number of 65546 in plain-number format is represented as 1.10 in the AS-dot notation format.

The auto route target feature does not support the local AS number for BGP neighbors associated with the Ethernet Virtual Private Network Instance (EVI).

Options

`alias`

(Optional) Configure the local AS as an alias of the global AS number configured for the router at the `[edit routing-options]` hierarchy level. As a result, a BGP peer considers any local AS to which it is assigned as equivalent to the primary AS number configured for the routing device. When you use the `alias` option, only the AS (global or local) used to establish the BGP session is prepended in the AS path sent to the BGP neighbor.

NOTE: The `private` and `alias` options are mutually exclusive. You cannot configure both options with the same `local-as` statement.

`autonomous-system`

AS number.

- **Range:** 1 through 4,294,967,295 (2^{32} – 1) in plain-number format
- **Range:** 0.0 through 65535.65535 in AS-dot notation format

loops *number* (Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure `loops 1`, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure `loops 2`, the route is hidden if the AS number is detected in the path two or more times.

The behavior of this statement is slightly different from the "loops (BGP Address Family)" on page 10134 statement.

> **NOTE**: If you configure the local AS values for any BGP group, the detection of routing loops is performed using both the AS and the local AS values for all BGP groups.
>
> If the local AS for the EBGP or IBGP peer is the same as the current AS, do not use the `local-as` statement to specify the local AS number.
>
> When you configure the local AS within a VRF, this impacts the AS path loop-detection mechanism. All of the `local-as` statements configured on the device are part of a single AS domain. The AS path loop-detection mechanism is based on looking for a matching AS present in the domain.

- **Range:** 1 through 10
- **Default:** 1

no-prepend-global-as (Optional) Specify to strip the global AS and to prepend only the local AS in AS paths sent to external peers.

private (Optional) Configure to use the local AS only during the establishment of the BGP session with a BGP neighbor but to hide it in the AS path sent to external BGP peers. Only the global AS is included in the AS path sent to external peers.

> **NOTE**: The `private` and `alias` options are mutually exclusive. You cannot configure both options with the same `local-as` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`alias` option introduced in Junos OS Release 9.5.

`loops` and `no-prepend-global-as` options introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Examples: Configuring BGP Local AS

Example: Configuring a Local AS for EBGP Sessions

autonomous-system | **2764**

family (Protocols BGP) | **6900**

local-authentication (Dynamic PPP Options)

IN THIS SECTION

- Syntax | **9919**
- Hierarchy Level | **9920**
- Description | **9920**
- Options | **9920**
- Required Privilege Level | **9920**
- Release Information | **9920**

Syntax

```
local-authentication {
    password password;
    username-include {
        circuit-id;
```

```
        delimiter character;
        domain-name name;
        mac-address;
        remote-id;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces $junos-interface-ifd-name unit $junos-interface-unit ppp-options]
```

Description

Configure local authentication for terminated PPP subscribers. This enables the external RADIUS server to pass implementation-specific configuration for successfully authenticated subscribers. Local authentication enables the same dynamic profile to support both CPEs that do not negotiate authentication protocols and CPEs that use PAP or CHAP authentication.

Options

password *password* Specify the local authentication password

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

PPP Subscriber Access Networks Overview

local-bias (ae)

IN THIS SECTION

- Syntax | 9921
- Hierarchy Level | 9921
- Description | 9921
- Required Privilege Level | 9921
- Release Information | 9922

Syntax

```
local-bias;
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

Description

Enable local link bias for all links in the aggregated Ethernet interface.

Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic.

You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced as it exits the Virtual Chassis or VCF.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D20.

RELATED DOCUMENTATION

Configuring Local Link Bias

Understanding Local Link Bias

local-bias

IN THIS SECTION

- Syntax | 9922
- Hierarchy Level | 9922
- Description | 9923
- Required Privilege Level | 9923
- Release Information | 9923

Syntax

```
local-bias percent bias;
```

Hierarchy Level

```
[edit interfaces rlt x logical-tunnel-options load-balance]
```

Description

Next hop addresses may be local or remote, and absent any other load balancing scheme (except adaptive load balancing), traffic can be expected to be more-or-less evenly distributed among the available next-hop addresses whether they are local or remote. You can skew distribution to favor local addresses, however, by setting a value for local bias (local relative to the packet forwarding engine (PFE) performing the packet look up).

For example, a value of 100 would exclude remote next-hop addresses from the traffic distribution by forcing 100% of next-hop traffic flows to use local addresses. A value of 50, on the other hand, would skew 50% of the flows that would otherwise use remote links so they use local links instead. That is, for a value set to 50, given four next-hop links, two of which are local and two of which are remote, each of the remote links could be expected to get one eighth of the flows (25% / 2) = 12.5%. Likewise, each of the local links could also be expected to receive about a third of the flows (25% + 12.5%) = 37.5%.

In contrast, with no value set for local bias, each of the four links would be expected to receive 25% of the total flows.

Required Privilege Level

interface - To view statement in the configuration.

interface-control - To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

local-bias (ae load-balance)

IN THIS SECTION

- Syntax | 9924
- Hierarchy Level | 9924
- Description | 9924
- Required Privilege Level | 9925
- Release Information | 9925

Syntax

```
local-bias percent bias;
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options load-balance]
```

Description

Next hop addresses may be local or remote, and traffic can be expected to be more-or-less evenly distributed among the available next-hop addresses whether they are local or remote. You can skew distribution to favor local addresses, however, by setting a value for local bias (local relative to the packet forwarding engine (PFE) performing the packet look up).

For example, a value of 100 would exclude remote next-hop addresses from the traffic distribution by forcing 100% of next-hop traffic flows to use local addresses. A value of 50, on the other hand, would skew 50% of the flows that would otherwise use remote links so they use local links instead. That is, for a value set to 50, given four next-hop links, two of which are local and two of which are remote, each of the remote links could be expected to get one eighth of the flows (25% / 2) = 12.5%. Likewise, each of the local links could also be expected to receive about a third of the flows (25% + 12.5%) = 37.5%.

In contrast, with no value set for local bias, each of the four links would be expected to receive 25% of the total flows.

You can use `local-bias` with adaptive load balancing, which uses a feedback mechanism to automatically correct load imbalance by adjusting the bandwidth and packet streams traversing links within an AE bundle. In this case, `local-bias` also employs a combination of link-saturation prediction and random restart delay to govern link utilization in a way that prevents oscillation of the load and load balancing schemes in effect.

> **NOTE**: MPC5 and MPC6 line cards include XM and XL-based packet forwarding engines, or PFEs, and locality is decided on the basis of the XL chip, not the XM chip. Therefore, when an AE bundle has child links hosted on two different XMs that are connected (in the chip architecture) to a single XL, they are considered local to the XL PFE. In practice, what this means is that if a single AE interface includes member links that happen to be spread over two XMs but are actually served by the same XL, `local-bias` may not work as expected because links are considered local to the XL PFE.

Required Privilege Level

interface - To view statement in the configuration.

interface-control - To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Understanding Aggregated Ethernet Load Balancing

Example: Configuring Aggregated Ethernet Load Balancing

local-bw-override-threshold

IN THIS SECTION

- Syntax | 9925
- Hierarchy Level | 9926
- Release Information | 9926
- Description | 9926
- Required Privilege Level | 9926

Syntax

```
local-bw-override-threshold;
```

Hierarchy Level

```
[edit logical-systems name protocols rsvp interface name non-rsvp-bandwdith],
[edit logical-systems name routing-instances name protocols rsvp interface name non-rsvp-bandwdith],
[edit protocols rsvp interface name non-rsvp-bandwdith],
[edit routing-instances name protocols rsvp interface name non-rsvp-bandwdith]
```

Release Information

Statement introduced in Junos OS Release 21.3R1 and Junos OS Evolved Release 21.3R1.

Description

RSVP local unreserved bandwidth is updated irrespective of the threshold configured on the `update-threshold-max-reservable` is crossed or not. As this statement only changes RSVP local bandwidth behavior, IGP bandwidth update behavior using the `update-threshold-max-reservable` statement remains same.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

update-threshold-max-reservable | 30458

local-category (Next Gen Services Service-Set Local System Logging)

IN THIS SECTION

- Syntax | 9927
- Hierarchy Level | 9927
- Description | 9927

- Options | 9927
- Required Privilege Level | 9928
- Release Information | 9928

Syntax

```
local-category category, category....category;
```

Hierarchy Level

```
[edit services service-set name  syslog
```

Description

Specify the category for which you want to collect local logs.

Options

all	All events are logged
content-security	Content security events are logged
fw-auth	Fw-auth events are logged
screen	Screen events are logged
alg	Alg events are logged
nat	NAT events are logged
flow	Flow events are logged
sctp	Sctp events are logged
gtp	Gtp events are logged

ipsec	Ipsec events are logged
idp	Idp events are logged
rtlog	Rtlog events are logged
pst-ds-lite	Pst-ds-lite events are logged
appqos	Appqos events are logged
secintel	Secintel events are logged
aamw	AAMW events are logged
sfw	Stateful Firewall events are logged
session	Session open and close events are logged
session-open	Session open events are logged
session-close	Session close events are logged
urlf	DNS request filtering events are logged
ha	Stateful High-Availability open and close events are logged
ha-open	Stateful High-Availability open events are logged
ha-close	Stateful High-Availability close events are logged
pcp	PCP logs

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

local-certificate

IN THIS SECTION

- Syntax | 9929
- Hierarchy Level | 9929
- Description | 9929
- Options | 9929
- Required Privilege Level | 9930
- Release Information | 9930

Syntax

```
local-certificate identifier;
```

Hierarchy Level

```
[edit services ipsec-vpn ike policy policy-name]
```

Description

Name of the certificate that needs to be sent to the peer during the IKE authentication phase.

Options

identifier—Name of certificate.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configuring IKE Policies

local-certificate (System)

IN THIS SECTION

- Syntax | 9930
- Hierarchy Level | 9931
- Description | 9931
- Required Privilege Level | 9931
- Release Information | 9931

Syntax

```
local-certificate name;
```

Hierarchy Level

```
[edit system services service-deployment],
[edit system services grpc request-response grpc ssl],
[edit system ntp nts]
```

Description

Specify the SSL certificate to use to authenticate the local device for system services, gRPC Remote Procedure Calls (gRPC), and Network Time Security (NTS) using Network Time Protocol (NTP).

Specify the name of a local certificate. There is no default for `local-certificate`. The value for `local-certificate` should be the same as the name provided during the import of the certificate. You can import the certificate using the CLI configuration statement `local` at the `[edit security certificates]` hierarchy level. Alternatively, on supported devices, you can import the certificate into the Junos public key infrastructure (PKI) by using the `request security pki local-certificate load` operational mode command.

Use `set system ntp nts local-certificate nts-cert` to configure local NTS certificate. The server provides a local certificate and the certificate ID to the client for verification, that is loaded on to the device during certificate enrollment. The server must configure the local-certificate. The client verifies the certificate sent by the server and proceeds with time synchronization if server certificate is valid.

NOTE: If you load the local certificate in the Junos PKI on supported devices, you must configure both the `local-certificate` statement and the `use-pki` statement at the same hierarchy level. The `use-pki` statement instructs the device to search the PKI database for local certificates.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for the `[edit system services extension-service request-response thrift]` hierarchy level in Junos OS Release 16.1 for MX80, MX480, MX960, MX2010, MX2020, vMX, and PTX Series.

RELATED DOCUMENTATION

Configuring clear-text or SSL Service for Junos XML Protocol Client Applications

Importing SSL Certificates for Junos XML Protocol Support

local

local-certificate (Security)

IN THIS SECTION

- Syntax | 9932
- Hierarchy Level | 9932
- Description | 9932
- Options | 9932
- Required Privilege Level | 9933
- Release Information | 9933

Syntax

```
local-certificate certificate-filename;
```

Hierarchy Level

```
[edit security ike policy ike-peer-address]
```

Description

Configure the certificate filename from which to read the local certificate.

Options

certificate-filename—File from which to read the local certificate.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an IKE Policy for Digital Certificates for an ES PIC

local-chap

IN THIS SECTION

- Syntax | 9933
- Hierarchy Level | 9933
- Description | 9934
- Required Privilege Level | 9934
- Release Information | 9934

Syntax

```
local-chap;
```

Hierarchy Level

```
[edit access group-profile profile-name l2tp],
[edit access profile profile-name client client-name l2tp]
```

Description

Configure the Junos OS so that the LNS ignores proxy authentication attribute-value pairs (AVPs) from the L2TP access concentrator (LAC) and reauthenticates the PPP client using a Challenge Handshake Authentication Protocol (CHAP) challenge. When you do this, the LNS directly authenticates the PPP client.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Layer 2 Tunneling Protocol (L2TP)

local-decision (PCRF Partition)

IN THIS SECTION

- Syntax | 9935
- Hierarchy Level | 9935
- Description | 9935
- Options | 9936
- Required Privilege Level | 9937
- Release Information | 9937

Syntax

```
local-decision {
    deny;
    grant;
    reinit-on-failure;
    reinit-on-rar;
    reinit-timeout seconds;
    timeout seconds;
}
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure a non-standard extension that allows the subscriber login to proceed even if the Policy and Charging Rule Function (PCRF) is not available or if the PCRF does not respond within the corresponding logout response timeout. This statement determines how long to wait before using the local decision to log in.

You can also use this statement to configure the router to reinitialize the PCRF session if either of two PCRF errors result when the following sequence occurs:

1. The router sends a CCR-GX-I to the PCRF.
2. The PCRF responds with a CCA-GX-I, but the message is lost.
3. The router retries sending a CCR-GX-I to the PCRF.

Table 55 on page 9935 lists the errors and the resolution.

Table 55: Reinitialization Options for PCRF Errors

Error	Resolution
The PCRF server responds to the router's CCR-GX-I with a CCA-GX-I that contains an unable-to-comply error code (5012) in AVP 268.	Include the `reinit-on-failure` option.

Table 55: Reinitialization Options for PCRF Errors *(Continued)*

Error	Resolution
The PCRF responds with a RAR message of any type.	Include the `reinit-on-rar` option.

The reinitialization operation consists of the following steps in both cases:

1. The router sends a session termination request, CCR-GX-T, to the PCRF.
2. The router waits for a CCA-GX-T for the length of the `reinit-timeout` period.
3. If the router either receives a CCA-GX-T within the timeout or the timeout expires without receiving a CCA-GX-T, the router regenerates the session ID (conveyed in Diameter AVP 263) and appends a session stamp that consists of the UTC time when the router creates the CCR-GX-I.

The `reinit-on-failure` and `reinit-on-rar` options affect existing subscribers only when the PCRF client (local) state is local-active, local-grant, or started.

Options

deny (Optional) Prevent subscriber logins from occurring when the PCRF is not available or not responding.

grant (Optional) Allow subscriber logins to occur when the PCRF is not available or not responding.

- **Default:** deny

reinit-on-failure (Optional) Reinitialize the PCRF session if PCRF responds with an unable-to-comply error code instead of sending another CCA-GX-I, after the router retries sending a CCR-GX-I.

> **NOTE**: You must also configure the following:
>
> - The `grant` option
> - The `use-session-stamp` option with the "partition" on page 12187 statement

reinit-on-rar (Optional) Reinitialize the PCRF session if the PCRF responds with a RAR of any type after the router retries sending a CCR-GX-I.

> **NOTE**: You must also configure the following:.
>
> - The `grant` option
> - The `use-session-stamp` option with the "partition" on page 12187 statement

reinit-timeout *seconds*

(Optional) Sets the number of seconds that the router waits for a CCA-GX-T during a local reinitialization operation. If the timeout expires before the CCA-GX-T is received, then the router performs the same actions as if it had received the CCA-GX-T within the timeout period. The router sends another CCR-GX-I to the PCRF with a new extended session ID (conveyed in Diameter AVP 263). The extended session ID has an appended session stamp that consists of the UTC time when the router creates the CCR-GX-I.

- **Default:** 10
- **Range:** 0 through 90

timeout *seconds*

Sets the number of seconds to wait before using the local decision to log in.

- **Default:** 90
- **Range:** 0 through 86,400 seconds (24 hours)

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

`reinit-on-failure`, `reinit-on-rar`, and `reinit timeout` options added in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

Understanding Upstream and Downstream Messages for the PCRF

local-dump (Flow Aggregation)

IN THIS SECTION

- Syntax | 9938
- Hierarchy Level | 9938
- Description | 9938
- Options | 9938
- Required Privilege Level | 9939
- Release Information | 9939

Syntax

```
(local-dump | no-local-dump);
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server hostname],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server hostname]
```

Description

Enable collection of cflowd records in a log file.

Options

`no-local-dump`—Do not dump cflowd records to a log file before exporting.

`local-dump`—Dump cflowd records to a log file before exporting.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Flow Aggregation

Debugging cflowd Flow Aggregation

local-endpoint

IN THIS SECTION

- Syntax | 9939
- Hierarchy Level | 9940
- Description | 9940
- Options | 9940
- Required Privilege Level | 9940
- Release Information | 9940

Syntax

```
local-endpoint {
 <initial-tnla-weight-factor initial-tnla-weight-factor>;
 ip-address address;
```

```
    <port port-number>;
  }
```

Hierarchy Level

```
[edit services agf amf]
```

Description

Configure the local endpoint.options on the AGF. AGF uses the local endpoint values to initiate an initial static TNLA with the AMF.

Options

initial-tnla-weight-factor (Optional) Specify the initial weight factor that AGF assigns to the static TNLA with the AMF.

- **Default:** 128
- **Range:** 1 through 255

ip-address *address* Specify the IP address that the AGF uses as the local SCTP endpoint

port *port-number* (Optional) Specifies the port number that AGF binds to the local SCTP endpoint. If the port is not configured, the operating system selects an available ephemeral port number.

- Range: 37100 through 73164

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 22.3R1.

local-engine

IN THIS SECTION

- Syntax | 9941
- Hierarchy Level | 9942
- Description | 9942
- Required Privilege Level | 9942
- Release Information | 9942

Syntax

```
local-engine {
    user username {
        authentication-md5 {
            (authentication-key authentication-key | authentication-password authentication-
password);
        }
        authentication-none;
        authentication-sha {
            (authentication-key authentication-key | authentication-password authentication-
password);
        }
        privacy-aes128 {
            (privacy-key privacy-key | privacy-password privacy-password);
        }
        privacy-des {
            (privacy-key privacy-key | privacy-password privacy-password);
        }
        privacy-3des {
            (privacy-key privacy-key | privacy-password privacy-password);
        }
        privacy-none {
        }
    }
}
```

Hierarchy Level

```
[edit snmp v3 usm]
```

Description

Configure local engine information for the user-based security model (USM).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Creating SNMPv3 Users

local (Flat-File Access Profile)

IN THIS SECTION

- Syntax | 9943
- Hierarchy Level | 9943
- Description | 9943
- Options | 9943
- Required Privilege Level | 9943
- Release Information | 9944

Syntax

```
local {
    flat-file-profile profile-name;
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure local reporting of service accounting statistics according to the specified flat-file profile, which determines the accounting statistics and nonstatistical parameters that are collected and reported in the local flat file. This configuration collects the running total service statistics per interface family. Because the statistics are maintained in the Routing Engine in a statistics database, they are not affected by a line-card restart, a graceful Routing Engine switchover, or a unified in-service software upgrade (ISSU). The statistics counters are reset when the router reboots.

Collecting service statistics into a local flat file is useful when you do not need to send the information to RADIUS, such as when your are using the information for internal monitoring and need an accurate record for analytics.

> **NOTE**: Starting in Junos OS Release 18.4R1, this statement is no longer supported. If included in a configuration, it is ignored.

Options

profile-name	Name of the flat-file profile.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

This statement is no longer supported as of Junos OS Release 18.4R1

RELATED DOCUMENTATION

Configuring Service Accounting in Local Flat Files

Flat-File Accounting Overview

local-gateway (IPSec)

IN THIS SECTION

- Syntax | 9944
- Hierarchy Level | 9944
- Description | 9945
- Options | 9945
- Required Privilege Level | 9945
- Release Information | 9945

Syntax

```
local-gateway address <gw-interface interface-name.logical-unit-number>;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Define the local IPv4 or IPv6 address for the IPsec traffic.

> **NOTE**: You cannot use a VRRP "virtual-address" on page 30739 for defining the local-gateway address.

Options

address	Local address.

The remaining statement is explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Service Rules

local-gateway (L2TP LNS)

IN THIS SECTION

- Syntax | 9946
- Hierarchy Level | 9946
- Description | 9946
- Options | 9946

Required Privilege Level | 9946
Release Information | 9946

Syntax

```
local-gateway {
    address address;
    gateway-name gateway-name;
}
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Specify the IP address or name for the local (LNS) gateway for L2TP tunnel.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

address—Local IP address; corresponds to the IP address that is used by LACs to identify the LNS. When the LAC is an MX Series router, this address matches the remote gateway address configured in the LAC tunnel profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring the Local Gateway Address and PIC
- Configuring L2TP Tunnel Groups
- *Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces*

local-id

IN THIS SECTION

- Syntax | 9947
- Hierarchy Level | 9947
- Description | 9947
- Options | 9948
- Required Privilege Level | 9948
- Release Information | 9948

Syntax

```
local-id (distinguished-name  container container-string-values |wildcard wildcard-string-values
ipv4_addr ipv4-address | ipv6-addr ipv6-address | key-id identifier fqdn fqdn);
```

Hierarchy Level

```
[edit services ipsec-vpn ike policy policy-name]
```

Description

Specify local identifiers for IKE Phase 1 negotiation. This statement is optional.

Options

`ipv4_addr` *`ipv4-address`*—IPv4 address identification value.

`ipv6_addr` *`ipv6-address`*—IPv6 address identification value.

`key_id` *`identifier`*—Key identification value.

`fqdn` *`fqdn`*—Fully-qualified domain name.

`distinguished-name container` *`container-string-values`* `|wildcard` *`wildcard-string-values`*—One or more distinguished name values.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`ipv6_addr` option added in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Security Associations

local-id (High Availability)

IN THIS SECTION

- Syntax | 9949
- Hierarchy Level | 9949
- Description | 9949
- Options | 9949
- Required Privilege Level | 9949

Release Information | 9950

Syntax

```
local-id id-number {
local-forwarding-ip loopback-address;
local-ip local-ip-address;
}
```

Hierarchy Level

```
[edit chassis high-availability]
```

Description

Configure the local node identifier for Multinode High Availability.

Options

id Local identifier number

- Range: 1 through 10

local-ip Local IPv4 address

local-forwarding-ip IP address used for local forwarding IP. This IP address must be associated to loopback interface only. You must configure this parameter to support asymmetric traffic flows in Multinode High Availability.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Multinode High Availability

Example: Configure Multinode High Availability in a Layer 3 Network

local-identity

IN THIS SECTION

- Syntax | 9950
- Hierarchy Level | 9950
- Description | 9951
- Options | 9951
- Required Privilege Level | 9951
- Release Information | 9951

Syntax

```
local-identity (distinguished-name | hostname identity-hostname | inet identity-ipv4 | inet6
identity-ipv6 | key-id | user-at-hostname identity-user);
```

Hierarchy Level

```
[edit security ike gateway gateway-name]
```

Description

Specify the local IKE identity to send in the exchange with the destination peer to establish communication. If you do not configure a local-identity, the device uses the IPv4 or IPv6 address corresponding to the local endpoint by default.

For Network Address Translation Traversal (NAT-T), both local identity and remote identity must be configured.

Options

- `distinguished-name` *`distinguished name`*—Specify a distinguished name as the identifier for the remote gateway.
- `hostname` *`hostname`*—Specify identity as a fully qualified domain name (FQDN).
- `inet` *`ip-address`*—Specify identity as an IPv4 address.
- `inet6` *`ip-address`*—Specify identity as an IPv6 address.
- `user-at-hostname` *`e-mail-address`*—Specify identity as an e-mail address.
- `key-id` *`sring-key-id`*—Specify key ID in ASCII string.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. The `inet6` option added in Junos OS Release 11.1.

RELATED DOCUMENTATION

IPsec Overview

local-interface (IPv6)

IN THIS SECTION

- Syntax | 9952
- Hierarchy Level | 9952
- Description | 9952
- Options | 9953
- Required Privilege Level | 9953
- Release Information | 9953

Syntax

```
local-interface interface-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp group group-name neighbor ipv6-link-local-address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name neighbor ipv6-link-local-address],
[edit protocols bgp group group-name neighbor ipv6-link-local-address],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor ipv6-link-local-address]
```

Description

Specify the interface name of the EBGP peer that uses IPv6 link-local addresses. This peer is link-local in scope.

TIP: Configure a local interface only if you need to use the IPv6 link-local addresses as BGP endpoints for an IPv6 BGP session.

> **NOTE**: The local interface option does not work if you have configured the `allow` option at the `[edit protocols bgp group group-name]` hierarchy level. You need to configure a BGP neighbor with an IPv6 link-local address if you have implicitly allowed peer connections from specified networks or hosts.

Options

`interface-name`—Interface name of the EBGP IPv6 peer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring Internal BGP Peering Sessions on Logical Systems

Example: Configuring Internal BGP Peer Sessions

Example: Configuring External BGP on Logical Systems with IPv6 Interfaces

Understanding Internal BGP Peering Sessions

local-ip-addr (ICCP)

IN THIS SECTION

- Syntax | 9954
- Hierarchy Level | 9954
- Description | 9954

- Options | 9954
- Required Privilege Level | 9954
- Release Information | 9954

Syntax

```
local-ip-addr local-ip-address;
```

Hierarchy Level

```
[edit protocols iccp],
[edit protocols iccp peer peer-IP-address]
```

Description

Specify the local IP address of the interchassis link (ICL) interface that Inter-Chassis Control Protocol (ICCP) uses to communicate to the peers that host a multichassis link aggregation group (MC-LAG).

Options

local-ip-address—Default local IP address to be used by all peers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

local-ipv4-address

IN THIS SECTION

- Syntax | 9955
- Hierarchy Level | 9955
- Description | 9955
- Required Privilege Level | 9956
- Release Information | 9956

Syntax

```
local-ipv4-address local ipv4 address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family address-family],
[edit logical-systems logical-system-name protocols bgp group group-name family address-family],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
address-family],
[edit protocols bgp family address-family],
[edit protocols bgp group group-name family address-family],
[edit protocols bgp group group-name neighbor address family address-family]
```

Description

Specify the local IPv4 address of a device that is also configured with an IPv6 address in a dual-stack environment. This enables BGP to advertise IPv4 unicast reachability with IPv4 next hop to remote IPv4 hosts over an IPv6 BGP session.

BGP advertises IPv4 unicast reachability with IPv4 next hop over an IPv6 BGP session only where IPv4 is configured at both endpoints. You cannot configure this feature for the inet6 unicast, inet6 multicast, or inet6 labeled-unicast address families because BGP already has the capability to advertise these address families over an IPv6 BGP session.

The configured `local-ipv4-address` is used only when BGP advertises routes with self-next hop. When IBGP advertises routes learned from EBGP peers, or the route reflector advertises BGP routes to its clients, BGP does not change the route next hop and ignores the configured `local-ipv4-address` and uses the original IPv4 next hop.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Advertising IPv4 Routes over BGP IPv6 Sessions Overview

Example: Advertising IPv4 Routes over IPv6 BGP Sessions

local-ipv6-address

IN THIS SECTION

- Syntax | **9957**
- Hierarchy Level | **9957**
- Release Information | **9957**
- Description | **9957**
- Required Privilege Level | **9957**

Syntax

```
local-ipv6-address local-ipv6-address;
```

Hierarchy Level

```
[edit logical-systems name protocols pcep pce],
[edit protocols pcep pce]
```

Release Information

Statement introduced in Junos OS Release 21.3R1.

Description

Specify the IPv6 address of the local end of Path Computation Element Protocol (PCEP) session.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

destination-ipv6-address | 5443

local-key-pair

IN THIS SECTION

- Syntax | 9958
- Hierarchy Level | 9958
- Description | 9958
- Options | 9958
- Required Privilege Level | 9958
- Release Information | 9959

Syntax

```
local-key-pair private-public-key-file;
```

Hierarchy Level

```
[edit security ike policy ike-peer-address]
```

Description

Specify private and public keys.

Options

private-public-key-file—Specify the file from which to read the private and public key pair.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos 7.4.

RELATED DOCUMENTATION

Configuring an IKE Policy for Digital Certificates for an ES PIC

local-log-tag (Next Gen Services Service-Set System Logging)

IN THIS SECTION

- Syntax | 9959
- Hierarchy Level | 9959
- Description | 9959
- Required Privilege Level | 9960
- Release Information | 9960

Syntax

```
local-log-tag tag-stamp;
```

Hierarchy Level

```
[edit services service-set name  syslog
edit services service-set name syslog stream stream-name
```

Description

Each log message is stamped with this tag.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

local-loopback

IN THIS SECTION

- Syntax | 9960
- Hierarchy Level | 9961
- Description | 9961
- Default | 9961
- Options | 9961
- Required Privilege Level | 9961
- Release Information | 9961

Syntax

```
(local-loopback-enable | no-local-loopback);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify whether local-loopback is enabled or disabled.

Default

If you omit the local-loopbback-enable statement, local-loopback is disabled.

Options

local-loopback-enable	Enable local-loopback.
no-local-loopback	Disable local-loopback.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

> *100-Gigabit Ethernet OTN Options Configuration Overview*
>
> *Configuring OTN Interfaces on P1-PTX-2-100G-WDM*

local-minimum-links-threshold

IN THIS SECTION

- Syntax | 9962
- Hierarchy Level | 9962
- Description | 9962
- Default | 9963
- Options | 9963
- Required Privilege Level | 9963
- Release Information | 9963

Syntax

```
local-minimum-links-threshold threshold-value
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

Description

For an aggregated Ethernet interface (LAG bundle) with member links spanning multiple chassis (member switches) in a Virtual Chassis or Virtual Chassis Fabric (VCF), set a threshold for the percentage of member links local to any particular chassis that must be up for that chassis to continue to be active in the aggregated Ethernet bundle.

> **NOTE**: This statement is available only on member switches in a non-mixed mode QFX5100 Virtual Chassis or VCF.

You configure a threshold for a particular aggregated Ethernet interface (ae*x*). When set, the threshold applies locally to any chassis with links in the specified LAG bundle. The threshold value represents a percentage of active local member links out of the total number of local member links for the chassis.

When one or more member links on a chassis go down, the system compares the percentage of local member links that are “up” to the threshold. If the percentage of local member links that are “up” is less than the threshold, any remaining active local links are forced down as well, to prevent forwarding traffic for the aggregated Ethernet interface through any member links on that chassis. If the percentage of “up” links is greater than or equal to the threshold, the status of the active links remains unchanged and they can forward LAG traffic.

For example, consider a case where the threshold is set to 52, and one local LAG member link goes down on one switch in a Virtual Chassis Fabric that has a total of four local LAG member links. In this case, 75 percent of the links are still up (greater than the threshold, 52 percent), so the remaining local member links stay up. However, if two local member links go down, only 50 percent of the links are up, so the local minimum links feature forces the remaining two active local member links down as well.

This feature also adjusts local member link status accordingly when failed links come up again, if you reconfigure the threshold value, or if you add or remove local member links in the LAG bundle.

The local minimum links feature helps avoid traffic loss due to asymmetric bandwidth on the forwarding paths across a chassis when some local aggregated Ethernet member links fail and some remain active. Enable this feature only if you want to closely manage ingress and egress traffic forwarding paths on aggregated Ethernet interfaces, especially where local link bias is also enabled.

Default

The local minimum links feature is disabled by default.

Options

threshold-value—Percentage of member links in an aggregated Ethernet bundle local to a chassis that must be up for any local member links on that chassis to be active in the aggregated Ethernet bundle.

- **Range:** 1 through 100 (decimal)
- **Default:** none—This option is not enabled by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D40.

RELATED DOCUMENTATION

Understanding Local Minimum Links

local-name

IN THIS SECTION

- Syntax | 9964
- Hierarchy Level | 9964
- Description | 9965
- Options | 9965
- Required Privilege Level | 9965
- Release Information | 9965

Syntax

```
local-name name;
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options chap],
[edit interfaces interface-name ppp-options pap],
[edit interfaces interface-name unit logical-unit-number ppp-options chap],
[edit interfaces interface-name unit logical-unit-number ppp-options pap],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options chap],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options pap]
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options],
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-
interface-unit” ppp-options]
```

Description

Specify the name of the interface used for CHAP or PAP authentication. Dynamic interfaces are supported only for CHAP authentication.

For ATM2 IQ interfaces only, you can configure a CHAP local name on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- `atm-ppp-llc`—PPP over AAL5 LLC encapsulation.
- `atm-ppp-vc-mux`—PPP over AAL5 multiplex encapsulation.

Options

name Name of the interface used as an identifier in CHAP challenge and response packets or PAP request and response packets.

- **Default:** When you do not include the `local-name` statement in the configuration, the interface sends the router's system hostname in CHAP challenge and response packets or PAP request and response packets.
- **Range:** For CHAP authentication, a string of 1 through 32 characters. For PAP authentication, a string of 1 through 8 characters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for PAP added in Junos OS Release 8.3.

Support at the `[edit dynamic-profiles` *`profile-name`* `interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options]` hierarchy level introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring the PPP Challenge Handshake Authentication Protocol

Configuring the PPP Password Authentication Protocol On a Physical Interface

local-password

IN THIS SECTION

- Syntax | 9966
- Hierarchy Level | 9966
- Description | 9966
- Required Privilege Level | 9966
- Release Information | 9966

Syntax

```
local-password password;
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options pap],
[edit interfaces interface-name unit logical-unit-number ppp-options pap],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options pap]
```

Description

Configure the host password for sending PAP requests.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring the PPP Password Authentication Protocol

local-port-range

IN THIS SECTION

- Syntax | 9967
- Hierarchy Level | 9967
- Description | 9967
- Default | 9968
- Options | 9968
- Required Privilege Level | 9968
- Release Information | 9968

Syntax

```
local-port-range {
    low low-value;
    high high-value;
}
```

Hierarchy Level

```
[edit unified-edge pcef flow-descriptions flow-identifier],
[edit services pcef flow-descriptions flow-identifier]
```

Description

Specify the port range to identify the subscriber traffic that you want the service data flow (SDF) filter to detect.

> **NOTE**: You can specify either `local-port-range` or a list of ports with `local-ports`, but not both.

If you are using Junos OS Subscriber Aware, specify the port range at the `[edit unified-edge pcef flow-descriptions flow-identifier]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the port range at the `[edit services pcef flow-descriptions flow-identifier]` hierarchy level.

Default

If the `local-port-range` statement is not configured, the default is any range of local ports.

Options

low-value

Lower boundary for the port range.

- **Range:** 1 through 65,535

high-value

Upper boundary for the port range.

- **Range:** 1 through 65,535

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef flow-descriptions flow-identifier]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

Configuring Service Data Flow Filters

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Understanding Application-Aware Policy Control for Subscriber Management

local-ports

IN THIS SECTION

- Syntax | 9969
- Hierarchy Level | 9969
- Description | 9969
- Default | 9970
- Options | 9970
- Required Privilege Level | 9970
- Release Information | 9970

Syntax

```
local-ports [number];
```

Hierarchy Level

```
[edit unified-edge pcef flow-description flow-identifier],
[edit services pcef flow-description flow-identifier]
```

Description

Specify a port number or list of port numbers to identify the subscriber traffic that you want the service data flow (SDF) filter to detect.

NOTE: You can specify either a list of ports or a port range, but not both.

If you are using Junos OS Subscriber Aware, specify the port numbers at the `[edit unified-edge pcef flow-description flow-identifier]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the port numbers at the `[edit services pcef flow-description flow-identifier]` hierarchy level.

Default

If the `local-ports` statement is not configured, the default is any local ports.

Options

number Number of a port or list of port numbers. You can specify a maximum of three port numbers (separated by a space) in a list.

- **Range:** 1 through 65,535

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef flow-description flow-identifier]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

Configuring Service Data Flow Filters

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Understanding Application-Aware Policy Control for Subscriber Management

local-preference

IN THIS SECTION

- Syntax | 9971
- Hierarchy Level | 9971
- Description | 9972
- Default | 9972
- Options | 9972
- Required Privilege Level | 9972
- Release Information | 9972

Syntax

```
local-preference local-preference;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
```

```
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Modify the value of the LOCAL_PREF path attribute, which is a metric used by IBGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.

The LOCAL_PREF path attribute always is advertised to internal BGP peers and to neighboring confederations. It is never advertised to external BGP peers.

Default

If you omit this statement, the LOCAL_PREF path attribute, if present, is not modified.

Options

local-preference—Preference to assign to routes that are advertised to the group or peer.

- **Range:** 0 through 4,294,967,295 (2^{32} – 1)
- **Default:** If the LOCAL_PREF path attribute is present, do not modify its value. If a BGP route is received without a LOCAL_PREF attribute, the route is handled locally (it is stored in the routing table and advertised by BGP) as if it were received with a LOCAL_PREF value of 100. By default, non-BGP routes that are advertised by BGP are advertised with a LOCAL_PREF value of 100.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring the Local Preference Value for BGP Routes

Understanding Internal BGP Peering Sessions

preference (Protocols BGP) | **13038**

local-priority

IN THIS SECTION

- Syntax | **9973**
- Hierarchy Level | **9973**
- Description | **9974**
- Options | **9974**
- Required Privilege Level | **9974**
- Release Information | **9974**

Syntax

```
local-priority local-priority;
```

Hierarchy Level

For MX Series:

```
[edit protocols ptp]
[edit protocols ptp master interface interface-name multicast-mode]
[edit protocols ptp master interface interface-name unicast-mode]
[edit protocols ptp slave interface interface-name multicast-mode]
[edit protocols ptp slave interface interface-name unicast-mode]
[edit protocols ptp stateful interface interface-name multicast-mode]
[edit protocols ptp stateful interface interface-name unicast-mode]
```

Description

Priority assigned to the local clock.

> **NOTE**: The stateful statement is not supported on QFX Series switches that support PTP

Configure a clock's local priority to be used as a tie-breaker in the dataset comparison algorithm, in the event that all other previous attributes of the datasets being compared are equal. The dataset comparison algorithm compares one clock with another by using the datasets representing those clocks, appended with the local-priority attribute. The local priority is assigned to the local clock and is used if needed when the data associated with the local clock is compared with data on another potential grandmaster (or the timeTransmitter) clock.

Options

local-priority *value*

The priority value of the clock.

- **Range:** 1 through 255
- **Default:** 128

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4

local-replication

IN THIS SECTION

- Syntax | 9975
- Hierarchy Level | 9975

- Description | 9975
- Required Privilege Level | 9975
- Release Information | 9975

Syntax

```
local-replication
```

Hierarchy Level

```
[edit forwarding-options satellite]
```

Description

Enables multicast replication on all the satellite devices that are connected to the aggregation device. You cannot selectively enable local replication for specific satellite devices, bridge domains, or route prefixes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Egress (Local) Replication on a Junos Fusion

local-server-group (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 9976
- Hierarchy Level | 9976
- Description | 9976
- Options | 9977
- Required Privilege Level | 9977
- Release Information | 9977

Syntax

```
local-server-group local-server-group;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option (default-action | equals | starts-with)],
[edit forwarding-options dhcp-relay group group-name relay-option (default-action | equals | starts-
with)],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Forward DHCP client packets to the specified group of DHCP local servers when you use the DHCP relay selective processing feature. You can configure the forwarding operation globally or for a group of interfaces.

The `local-server-group` option is not supported for DHCPv6 relay agent.

Options

local-server-group	Name of DHCP local server group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic

local-switching (Layer 2 Circuits)

IN THIS SECTION

- Syntax | 9977
- Hierarchy Level | 9978
- Description | 9978
- Required Privilege Level | 9978
- Release Information | 9978

Syntax

```
local-switching {
    interface interface-name {
        description text;
```

```
        encapsulation-type;
        end-interface {
            interface interface-name;
            no-revert;
            protect-interface interface-name;
        }
        ignore-encapsulation-mismatch;
        ignore-mtu-mismatch;
        no-revert;
        protect-interface interface-name;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit],
[edit protocols l2circuit]
```

Description

Configure a local switching interface. A local switching interface allows you to terminate a virtual circuit on the local router.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Local Interface Switching in Layer 2 Circuits

local-switching (VPLS)

IN THIS SECTION

- Syntax | 9979
- Hierarchy Level | 9979
- Description | 9979
- Required Privilege Level | 9979
- Release Information | 9980

Syntax

```
local-switching;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]
```

Description

Allows you to terminate multiple Layer 2 circuit pseudowires at a single VPLS mesh group. Note that when using `vpls-id-list`, you can enable local-switching to allow local switching of traffic (including BUM traffic) between multiple pseudo wires. Enabling local-switching also eliminates the need to have a dedicated mesh-group for each LDP spoke pseudo wire (which has a limit of 14).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Switching Between Pseudowires Using VPLS Mesh Groups

local-traffic-selector

IN THIS SECTION

- Syntax | 9980
- Hierarchy Level | 9980
- Description | 9981
- Options | 9981
- Required Privilege Level | 9981
- Release Information | 9981

Syntax

```
local-traffic-selector {
    (ipv4-prefix ipv4-prefix | ipv6-prefix ipv6-prefix);
    port port;
    protocol protocol;
}
```

Hierarchy Level

```
[edit security host-vpn connections connection-name children child-name]
```

Description

Configure the local IPsec traffic to be protected by the child security association. A traffic selector is a traffic filter that defines and identifies the traffic flow permitted between two systems (a specified pair of local and remote addresses) that have IPsec protection.

Options

(ipv4-prefix *ipv4-prefix* | ipv6-prefix *ipv6-prefix*)	Specify traffic to be protected by the child security association using either IPv4 or IPv6 with a prefix. The prefix allows for specifying more general traffic.
port *port*	Specify the port to protect by number or name. For example, port 21 and port ftp refer to the same port.
protocol *protocol*	Specify the protocol to protect by number or name. For example, protocol 6 and protocol tcp refer to the same protocol.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 18.3R1.

locality-bias (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 9982
- Hierarchy Level | 9982
- Description | 9982
- Required Privilege Level | 9982
- Release Information | 9982

Syntax

```
locality-bias;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Configure unicast transit traffic for equal-cost multipath (ECMP) groups and aggregated Ethernet bundles to egress links in the same (local) member router in an MX Series Virtual Chassis rather than to egress links in the remote member router, provided that the local member router has an equal or larger number of available egress links than the remote member router.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring Locality Bias for a Virtual Chassis

Locality Bias in a Virtual Chassis

Guidelines for Configuring Locality Bias in a Virtual Chassis

localized-fib

IN THIS SECTION

- Syntax | 9983
- Hierarchy Level | 9983
- Description | 9983
- Options | 9983
- Required Privilege Level | 9984
- Release Information | 9984

Syntax

```
localized-fib {
    fpc-slot fpc-slot-number;
}
```

Hierarchy Level

```
[edit routing-instances instance-name routing-options]
```

Description

Localize the VRF routing instance routes to a specific FPC of CE-facing physical interfaces, or specify an FPC slot number if the CE-facing interfaces are logical interfaces like AE or RSQL or IRB.

Options

fpc-slot *fpc-slot-number*	Specify the FPC slot number to localize the VRF routing instance for logical interfaces like AE or RSQL or IRB.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Understanding VRF Localization in Layer 3 VPNs

Example: Improving Scalability Using VRF Localization for Layer 3 VPNs

vpn-localization | **30942**

location (DHCP Snooping Database)

IN THIS SECTION

- Syntax | **9984**
- Hierarchy Level | **9985**
- Description | **9985**
- Options | **9985**
- Required Privilege Level | **9985**
- Release Information | **9985**

Syntax

```
location local_pathname | remote_URL;
```

Hierarchy Level

```
[edit ethernet-switching-options                secure-access-port                dhcp-snooping-file]
```

Description

Specify either a local pathname or a remote URL as the location in which to store the DHCP snooping database.

Options

local_pathname* | *remote_URL —Location for storing the DHCP snooping database.

- ***local_pathname*** —Use ***/path*** to store the database on a local switch.
- ***remote_URL*** —Use **ftp:*// ip-address*** or **ftp:*// hostname/ path*** to store the database at a remote location.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

location (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 9986
- Hierarchy Level | 9986
- Description | 9986
- Required Privilege Level | 9986

Release Information | 9987

Syntax

```
location name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82 vendor-specific],
[edit forwarding-options dhcp-relay group group-name relay-option-82 vendor-specific]
```

Description

Supports the addition of a vendor-specific location in the option 82, suboption 9 field of DHCPv4 control messages on server-facing interfaces. The location should be specified as interface, vlan ID, and if applicable, svlan ID. For example, **<ifd-name>:<vlan>** (ae0:100) or **<ifd-name>:<svlan> -<vlan>** (ae0:100-10).

Junos OS automatically adds the remaining vendor-specific information as per RFC 4243, *Vendor-Specific Information Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option*. The Juniper Networks enterprise ID is 2636. The the location is option-data 2 (the hostname is option-data 1). The DHCPv4 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query the hostname to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 16.2.

location (DHCP Snooping Database)

IN THIS SECTION

- Syntax | 9987
- Hierarchy Level | 9987
- Description | 9987
- Options | 9988
- Required Privilege Level | 9988
- Release Information | 9988

Syntax

```
location (local_pathname | remote_url);
                        timeout                     seconds;
                        write-interval                      seconds;
}
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port dhcp-snooping-file],
[edit ethernet-switching-options secure-access-port dhcpv6-snooping-file]
```

Description

Configure IP-MAC address bindings to persist through switch reboots by specifying a location in which to store the DHCP snooping database. When specifying the location for the DHCP snooping database,

you must also specify how frequently the switch writes (`write-interval`) the database entries into the DHCP snooping database file.

If you choose to store the DHCP snooping database on a remote FTP site, you might want to specify the time (`timeout`) that the switch waits for a remote system to respond when the DHCP snooping database is stored on a remote FTP site. This is optional.

Options

local_pathname | *remote_url*

- *local_pathname*—Use */path* to store the database file on the local switch.
- *remote_url*—Use `ftp://`*ip-address* or `ftp://` *hostname*/*path* to store the database on a remote FTP site.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Support at the `[edit ethernet-switching-options secure-access-port dhcpv6-snooping-file]` hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

RELATED DOCUMENTATION

Configuring Persistent Bindings in the DHCP or DHCPv6 (non-ELS)

Understanding DHCP Snooping (non-ELS)

location (DHCPv6 Relay Agent)

IN THIS SECTION

Syntax | **9989**

- Hierarchy Level | 9989
- Description | 9989
- Required Privilege Level | 9989
- Release Information | 9990

Syntax

```
location name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 relay-option-vendor-specific],
[edit forwarding-options dhcp-relay group group-name dhcpv6 relay-option-vendor-specific]
```

Description

Supports the addition of a vendor-specific location in the vendor-specific option (17) of DHCPv6 control messages on server-facing interfaces. The location should be specified as interface, vlan ID, and if applicable, svlan ID. For example, **<ifd-name>:<vlan>** (ae0:100) or **<ifd-name>:<svlan> -<vlan>** (ae0:100-10).

Junos automatically adds the remaining vendor-specific information as per RFC 4243, *Vendor-Specific Information Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option*. The Juniper Networks enterprise ID is 2636. The the location is option-data 2 (the hostname is option-data 1). The DHCPv6 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query the hostname to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 16.2.

location (LLDP-MED)

IN THIS SECTION

- Syntax | 9990
- Hierarchy Level | 9991
- Description | 9991
- Default | 9991
- Options | 9991
- Required Privilege Level | 9991
- Release Information | 9992

Syntax

```
location {
    civic-based {
        ca-type  type {
            ca-value value;
        }
        country-code country-code;
        what what;
    }
    co-ordinate {
        lattitude latitude;
        longitude longitude;
    }
```

```
    elin elin;
}
```

Hierarchy Level

```
[edit protocols lldp-med interface (all | interface-name)]
```

Description

For Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED), configure the location information. Location information is advertised from the switch to the MED. This information is used during emergency calls to identify the location of the MED.

The remaining statements are explained separately. See CLI Explorer.

Default

Disabled.

Options

co-ordinate Geographical coordinates for the location of the MED.

- **Values:** Specify these values:
  - lattitude *latitude*—Latitude value for the location.
  - longitude *longitude*—Longitude value for the location.
- **Range:** 0 through 360 degrees for both the latitude and longitude values.

elin *elin* Configure the Emergency Line Identification Number (ELIN) as part of the location information. The ELIN is a 10-digit telephone number, including the area code.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show lldp | 21942

Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch

Configuring LLDP-MED (CLI Procedure)

location (SNMP)

IN THIS SECTION

- Syntax | 9992
- Hierarchy Level | 9992
- Description | 9993
- Options | 9993
- Required Privilege Level | 9993
- Release Information | 9993

Syntax

```
location location;
```

Hierarchy Level

```
[edit snmp]
```

Description

Define the value of the MIB II **sysLocation** object, which is the physical location of the managed system.

Options

location—Location of the local system. You must enclose the name within quotation marks (" ").

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the System Location for a Device Running Junos OS

location (System)

IN THIS SECTION

- Syntax | 9994
- Hierarchy Level | 9994
- Description | 9994
- Options | 9994
- Required Privilege Level | 9995
- Release Information | 9995

Syntax

```
location {
    altitude feet;
    building name;
    country-code code;
    floor number;
    hcoord horizontal-coordinate;
    lata transport-area;
    latitude degrees;
    longitude degrees;
    npa-nxx number;
    postal-code postal-code;
    rack number;
    vcoord vertical-coordinate;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the system location in various formats.

Options

altitude *feet*	Number of feet above sea level.
building *name*	Name of the building. The name of the building can be 1 to 28 characters in length. If the string contains spaces, enclose it in quotation marks (" ").
country-code *code*	Two-letter country code.
floor *number*	Floor number in the building.
hcoord *horizontal-coordinate*	Bellcore Horizontal Coordinate.

lata *transport-area*	Local Access Transport Area.
latitude *degrees*	Latitude in degree format.
longitude *degrees*	Longitude in degree format.
npa-nxx *number*	First six digits of the phone number (area code and exchange).
postal-code *postal-code*	Postal code.
rack *number*	Rack number.
vcoord *vertical-coordinate*	Bellcore Vertical Coordinate.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

System Settings

location (Virtual Chassis)

IN THIS SECTION

- Syntax | 9996
- Hierarchy Level | 9996
- Description | 9996
- Options | 9996
- Required Privilege Level | 9996

Release Information | 9996

Syntax

```
location location;
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

Set a description of the location of the Virtual Chassis or VCF member switch or external Routing Engine.

The `Location` field is visible to users who enter the `show virtual-chassis status detail` command.

Setting this description has no effect on the operation of the member device.

Options

location—Location of the current member switch or external Routing Engine. The `location` can be any single word.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric

Configuring an EX4650 or a QFX Series Virtual Chassis

Configuring an EX2300, EX3400, or EX4300 Virtual Chassis

Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis

locator

IN THIS SECTION

- Syntax | **9997**
- Hierarchy Level | **9998**
- Description | **9998**
- Options | **9998**
- Required Privilege Level | **9999**
- Release Information | **9999**

Syntax

```
locator name {
    anycast
    end-sid sid | dynamic-end-sid
        flavor {
            psp;
            usp;
            usd;
        }
        algorithm algorithm
    }
    micro-sid
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options source-packet-routing srv6],
[edit logical-systems name routing-options source-packet-routing srv6],
[edit logical-systems name tenants name routing-instances name routing-options source-packet-
routing srv6],
[edit routing-instances name routing-options source-packet-routing srv6],
[edit routing-options source-packet-routing srv6],
[edit tenants name routing-instances name routing-options source-packet-routing srv6]
```

Description

Configure the locator for an SRv6 node in an IPv6 domain to indicate the SRv6 capability of the router. A Segment Identifier (SID) is a 128-bit IPv6 address that represents a specific segment in a segment routing domain. A SID consists of two parts, the first part is called the locator and consists of the most significant bits that represent the address of a particular SRv6 node. Routing protocols such as IGP and BGP that typically advertise routing prefixes, advertise the locator in an SRv6 node. The size of the locator is flexible that you can customize the size as per your requirements. However, we recommend configuring locator with prefix length of 64, which is sufficient for various services and transport purposes.

NOTE: Make sure that the locator and SIDs belong to the same subnet to avoid commit errors.

NOTE: Note: From 23.4R1 release, the existing `routing-options source-packet-routing srv6 locator` configuration commit check is removed to facilitate configuration of algorithms for micro-sids under the same locator.

Options

name Specify the locator entry name.

anycast Set A flag in the prefix-attribute sub-TLV through a policy to enable anycast locators. By default, the A flag is preserved while leaking IS-IS routes.

end-sid Configure the address of the locator service for SRv6. Ensure that the locator address and interface address belong to the same prefix to avoid unexpected network behavior.

flavor Specify end sid behavior such as Penultimate Segment Pop (PSP), Ultimate Segment Pop (USP) or Ultimate Segment Decapsulation (USD).

algorithm Specify the algorithm associated with this locator. You can define a flexible algorithm as per your network requirements.

- **Range:** 0 through 255

dynamic-end-sid You can configure this option in addition to the existing `end-sid`. IS-IS dynamically allocates one end-sid address for a locator. Note that you cannot configure this option for micro SIDs.

micro-sid Configure a micro node SID with the same value as that of the locator. Note that the micro node SID has the same flavors as that of the specified locator. If you configure this option, IS-IS advertises micro-node-SID in addition to the locator.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.3R1.

`micro-sid` option introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

srv6 (Routing Options) | **28415**

How to Enable SRv6 Network Programming in IS-IS Networks

Example: Configuring SRv6 Network Programming in IS-IS Networks

Configuring Flexible Algorithm for Segment Routing Traffic Engineering

lockout

IN THIS SECTION

Syntax | **10000**

- Hierarchy Level | **10000**
- Description | **10000**
- Required Privilege Level | **10000**

Syntax

```
request protection-group ethernet-aps lockout md <md> ma <ma>
```

Hierarchy Level

```
[edit protocols protection-group ethernet-aps]
```

Description

Configure a lockout of the protection path, forcing the use of the working path and locking out the protect path regardless of anything else.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Ethernet Automatic Protection Switching Overview

lockout-timeout (L2TP Destination Lockout)

IN THIS SECTION

- Syntax | 10001
- Hierarchy Level | 10001
- Description | 10001
- Options | 10002
- Required Privilege Level | 10002
- Release Information | 10002

Syntax

```
lockout-timeout seconds;
```

Hierarchy Level

```
[edit services l2tp destination lockout-result-code lockout-result-code-name lockout-error-code
lockout-error-code- name]
```

Description

Set the duration of the timeout period for which all future destinations are locked out, meaning that they are not considered for selection when a new tunnel is created. Destinations are locked out when L2TP cannot connect to the destination during the tunnel selection process. This statement does not affect destinations that are currently locked out.

NOTE: The *ip-address* option for the `destination` statement does not apply to the `lockout-timeout` statement.

Options

seconds Length of the period during which the destination is locked out.

Range: 60 through 3600 seconds

Default: 300 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

L2TP LNS Inline Service Interfaces

log

IN THIS SECTION

- Syntax | 10003
- Hierarchy Level | 10003
- Description | 10003
- Options | 10004
- Required Privilege Level | 10004
- Release Information | 10004

Syntax

```
log {
    session {
        client;
        all;
        dhcpv6 {
            client;
            server;
            relay;
            dynamic-server;
            all;
        }
        server;
        relay;
    }
}
```

Hierarchy Level

```
[edit system processes dhcp-service]
```

Description

Enable DHCP session log on the device. Session logs include the information on the session creation, deletion and renew events. You can use the session logs for monitoring and troubleshooting purposes.

DHCP logs are written to system syslog that you define using the following command:

```
set system syslog ( file filename | host destination) user logging-level
```

Example:

```
syslog {
        host 10.10.10.10 {
        user any;
    }
    file file-1 {
        user any;
```

```
    }
}
```

In the example:

Syslog messages are stored on the host with IP address 10.10.10.10 for all severity levels.

Syslog messages are stored in the file file-1 at /var/log/<file-name> for all severity levels.

Options

- `session`—Logs of the DHCP sessions.
- `client`—Log sessions of the DHCP client.
- `all`—Log sessions of the DHCP client, server and relay.
- `dhcpv6`—Log sessions of the DHCPv6.
- `client`—Log sessions of the DHCPv6 client.
- `dynamic-server`—Log sessions of the DHCPv6 dynamic server.
- `all`—Log sessions of the DHCPv6 client, server, relay and dynamic server.
- `server`—Log sessions of the DHCP server.
- `relay`—Log sessions of the DHCP relay.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

dhcp-service | 5646

log (aacl)

IN THIS SECTION

- Syntax | 10005
- Hierarchy Level | 10005
- Description | 10005
- Options | 10005
- Required Privilege Level | 10006
- Release Information | 10006

Syntax

```
log event-type
```

Hierarchy Level

```
[edit services aacl rule rule-name No Link Title term-name No Link Title ]
```

Description

Enable AACL logging of flows for known or unknown applications.

Options

event-type—Enable logging of the specified *event-type*:

- session-start
- session-end
- session-start-end-no-stats
- session-start-interim-end

- session-interim end
- session-end

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

log-and-generate-ais

IN THIS SECTION

- Syntax | 10006
- Hierarchy Level | 10007
- Description | 10007
- Options | 10007
- Required Privilege Level | 10007
- Release Information | 10007

Syntax

```
log-and-generate-ais {
    interval (1m | 1s);
    level level;
    priority level;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
action]
```

Description

Configure the action be taken when an AIS alarm is detected. The action includes generating and logging the AIS statistics along with the interval between AIS messages, the server maintenance domain level, and the priority of the AIS message.

Options

The other statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Alarm Indication Signal (ETH-AIS) Function Overview

Configuring ETH-AIS on a CFM MEP

log-attacks

IN THIS SECTION

Syntax | 10008

- Hierarchy Level | **10008**
- Description | **10008**
- Options | **10008**
- Required Privilege Level | **10009**
- Release Information | **10009**

Syntax

```
log-attacks {
    alert;
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name then notification]
```

Description

Enable the log attacks to create a log record that appears in the log viewer.

In addition to the regular system log messages, IDP generates event logs for attacks. IDP generates event logs when an event matches an IDP policy rule in which logging is enabled. When you configure a rule for logging, the device creates a log entry for each event that matches that rule.

Specify an alert option for a rule in the IDP policy. When the rule is matched, the corresponding log record displays an alert in the alert column of the Log Viewer. Security administrators use alerts to become aware of and react to important security events.

Options

alert Set an alert flag in the Alert column of the Log Viewer for the matching log record.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

log-interval (IGMP Interface)

IN THIS SECTION

- Syntax | 10009
- Hierarchy Level | 10009
- Description | 10010
- Default | 10010
- Options | 10010
- Required Privilege Level | 10010
- Release Information | 10010

Syntax

```
log-interval seconds;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name]
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Description

Specify the minimum time interval (in seconds) between sending consecutive log messages to the system log for multicast groups on static or dynamic IGMP interfaces. To configure the time interval, you must specify the maximum number of multicast groups allowed on the interface. You must configure the `group-limit` statement before you configure the `log-interval` statement.

To confirm the configured log interval on the interface, use the `show igmp interface` command.

Default

By default, there is no configured time interval.

Options

seconds—Minimum time interval (in seconds) between log messages. You must explicitly configure the **group-limit** to configure a time interval to send log messages.

- **Range:** 6 through 32,767 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

group-limit (IGMP) | **7844**

group-threshold (Protocols IGMP Interface) | **7903**

log-interval (MLD Interface)

IN THIS SECTION

- Syntax | 10011
- Hierarchy Level | 10011
- Description | 10011
- Default | 10011
- Options | 10012
- Required Privilege Level | 10012
- Release Information | 10012

Syntax

```
log-interval seconds;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

Specify the minimum time interval (in seconds) between sending consecutive log messages to the system log for multicast groups on static or dynamic MLD interfaces. To configure the time interval, you must specify the maximum number of multicast groups allowed on the interface.

To confirm the configured log interval on the interface, use the `show mld interface` command.

Default

By default, there is no configured time interval.

Options

seconds—Minimum time interval (in seconds) between log messages. You must explicitly configure the **group-limit** to configure a time interval to send log messages.

- **Range:** 6 through 32,767 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring the Number of MLD Multicast Group Joins on Logical Interfaces

group-limit (Protocols MLD) | **7847**

group-threshold (Protocols MLD Interface) | **7905**

log-interval (PIM Entries)

IN THIS SECTION

- Syntax | **10013**
- Hierarchy Level | **10013**
- Description | **10014**
- Options | **10014**
- Required Privilege Level | **10014**
- Release Information | **10014**

Syntax

```
log-interval value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim sglimit],
[edit logical-systems logical-system-name protocols pim sglimit family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
sglimit],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
sglimit family],
[edit protocols pim sglimit],
[edit protocols pim sglimit family],
[edit routing-instances routing-instance-name protocols pim sglimit],
[edit routing-instances routing-instance-name protocols pim sglimit family],
[edit logical-systems logical-system-name protocols pim rp group-rp-mapping],
[edit logical-systems logical-system-name protocols pim rp group-rp-mapping family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp group-rp-mapping],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp group-rp-mapping family],
[edit protocols pim rp group-rp-mapping],
[edit protocols pim rp group-rp-mapping family],
[edit routing-instances routing-instance-name protocols pim rp group-rp-mapping],
[edit routing-instances routing-instance-name protocols pim rp group-rp-mapping family],
[edit logical-systems logical-system-name protocols pim rp register-limit],
[edit logical-systems logical-system-name protocols pim rp register-limit family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp register-limit],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp register-limit family],
[edit protocols pim rp register-limit],
[edit protocols pim rp register-limit family],
[edit routing-instances routing-instance-name protocols pim rp register-limit],
[edit routing-instances routing-instance-name protocols pim rp register-limit family]
```

Description

Configure the amount of time between log messages.

Options

seconds—Minimum time interval (in seconds) between log messages. To configure the time interval, you must explicitly configure the maximum number of entries received with the `maximum` statement. You can apply the log interval to incoming PIM join messages, PIM register messages, and group-to-RP mappings.

- **Range:** 1 through 65,535

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

clear pim join | 3806

log-interval (Protocols MSDP)

IN THIS SECTION

- Syntax | 10015
- Hierarchy Level | 10015
- Description | 10015
- Options | 10015
- Required Privilege Level | 10015

Release Information | 10016

Syntax

```
log-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp active-source-limit],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
active-source-limit],
[edit protocols msdp active-source-limit],
[edit routing-instances routing-instance-name protocols msdp active-source-limit]
```

Description

Specify the minimum time interval (in seconds) between sending consecutive log messages to the system log for MSDP active source messages. To configure the time interval, you must specify the maximum number of MSDP active source messages received by the device.

To confirm the configured log interval, use the `show msdp source-active` command.

Options

seconds—Minimum time interval (in seconds) between log messages. You must explicitly configure the **maximum** value to configure a time interval to send log messages.

- **Range:** 6 through 32,767 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2

RELATED DOCUMENTATION

Example: Configuring MSDP with Active Source Limits and Mesh Groups

maximum (MSDP Active Source Messages) | **10536**

log (Logical Systems and Tenant Systems)

IN THIS SECTION

- Syntax | **10016**
- Hierarchy Level | **10017**
- Description | **10017**
- Options | **10017**
- Required Privilege Level | **10019**
- Release Information | **10019**

Syntax

```
        log {
            disable;
            facility-override (authorization | daemon | ftp | kernel | local0 | local1 | local2
| local3 | local4 | local5 | local6 | local7 | user);
            mode (event | stream);
            format (binary | sd-syslog | syslog);
            source-address;
            source-interface;
            stream name {
                category name;
                filter name;
                host {
```

```
                    ipaddr;
                    port port;
                    routing-instance routing-instance;
                }
                rate-limit rate;
            }
            transport {
                protocol (tcp | tls |udp);
                tls-profile;
            }
        }
    }
```

Hierarchy Level

```
[edit logical-systems name security]
[edit tenants name security]
```

Description

Configure security log for logical systems and tenant systems. Set the mode of logging (event for traditional system logging or stream for streaming security logs through a revenue port to a server). You can also specify all the other parameters for security logging.

Options

disable Disable the security logging for the device.

facility-override Alternate facility for logging to remote host.

- Values:
  - `authorization` —Authorization system
  - `daemon` —Various system processes
  - `ftp` —FTP process
  - `kernel` —Kernel

- `local0` —Local logging option number 0
- `local1` —Local logging option number 1
- `local2` —Local logging option number 2
- `local3` —Local logging option number 3
- `local4` —Local logging option number 4
- `local5` —Local logging option number 5
- `local6` —Local logging option number 6
- `local7` —Local logging option number 7
- `user` —User processes

`format` Set security log format for the device.

- Values:
  - `binary` —Binary log
  - `binarysd-syslog` —Structured syslog
  - `syslog` —Traditional syslog

`mode` Controls how security logs are processed and exported.

- Values:
  - `event` —Process security logs in the control plane
  - `stream` —Process security logs directly in the forwarding plane

`source-address` Specify a source IP address or IP address used when exporting security logs, which is mandatory to configure *stream host*.

`source-interface` Specify a source interface name, which is mandatory to configure *stream host*.

`stream` Set security log stream settings.

`transport` Set security log transport settings.

- Values:

- tcp—TCP transfer for log
- tls—TLS transfer for log
- udp—UDP transfer for log

utc-timestamp Specify to use UTC time for security log timestamps.

The following options are not supported under logical system and tenant system:

- `event-rate` and `rate-cap`— Use to limit the log rate between Packet Forwarding Engine (PFE) and Routing Engine (RE).
- `file`— Use to store binary log with event mode.
- `max-database-record` and `report`— Use to enable SQLite Version 3 (sqlite3) database for local log management daemon (llmd).
- `traceoptions`— Specify security log daemon trace options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

Statement introduced in Junos OS Release 18.3R1 for tenant systems.

The `routing-instance` option introduced in Junos OS Release 18.3R1 for tenant systems.

RELATED DOCUMENTATION

Troubleshooting Logical Systems

log-prefix (Interfaces)

IN THIS SECTION

- Syntax | 10020
- Hierarchy Level | 10020
- Description | 10020
- Options | 10020
- Required Privilege Level | 10020
- Release Information | 10021

Syntax

```
log-prefix prefix-value;
```

Hierarchy Level

```
[edit interfaces interface-name services-options syslog host hostname]
```

Description

Set the system logging prefix value.

Options

prefix-value—System logging prefix value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Services Interfaces Library for Routing Devices

Configuring System Logging for Services Interfaces

log-prefix (L2TP)

IN THIS SECTION

- Syntax | 10021
- Hierarchy Level | 10021
- Description | 10021
- Options | 10022
- Required Privilege Level | 10022
- Release Information | 10022

Syntax

```
log-prefix prefix-value;
```

Hierarchy Level

```
[edit services l2tp tunnel-group group-name syslog host hostname]
```

Description

Set the system logging prefix value.

Options

prefix-value—System logging prefix value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring System Logging of L2TP Tunnel Activity

log-prefix (Services)

IN THIS SECTION

- Syntax | 10022
- Hierarchy Level | 10023
- Description | 10023
- Options | 10023
- Required Privilege Level | 10023
- Release Information | 10023

Syntax

```
log-prefix prefix-value;
```

Hierarchy Level

```
[edit services service-set service-set-name syslog host hostname]
```

Description

Set the system logging prefix value.

Options

prefix-value—System logging prefix value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring System Logging for Service Sets

log-prefix (System)

IN THIS SECTION

- Syntax | 10024
- Hierarchy Level | 10024
- Description | 10024
- Options | 10024

Required Privilege Level | 10024
Release Information | 10024

Syntax

```
log-prefix string;
```

Hierarchy Level

```
[edit system syslog host]
```

Description

Include a text string in each message directed to a remote destination.

Options

string—Text string to include in each message.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Adding a Text String to System Log Messages Directed to a Remote Destination

log-rotate-frequency

IN THIS SECTION

- Syntax | 10025
- Hierarchy Level | 10025
- Description | 10025
- Options | 10025
- Required Privilege Level | 10026
- Release Information | 10026

Syntax

```
log-rotate-frequency frequency;
```

Hierarchy Level

```
[edit system syslog]
```

Description

Configure the system log file rotation frequency by configuring the time interval for checking the log file size.

When the log file size has exceeded the configured limit, the old log file is archived and a new log file is created.

Options

frequency Frequency of rotation of the system log file.

- **Range:** 1 minute through 59 minutes
- **Default:** 15 minutes

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Specifying Log File Size, Number, and Archiving Properties

syslog (System) | 28836

log (Security)

IN THIS SECTION

- Syntax | 10026
- Hierarchy Level | 10029
- Description | 10029
- Options | 10029
- Required Privilege Level | 10031
- Release Information | 10031

Syntax

```
log {
    (source-address source-address | source-interface source-interface);
    cache {
        exclude name {
            destination-address destination-address;
            destination-port destination-port;
```

```
            event-id event-id;
            failure;
            interface-name interface-name;
            policy-name policy-name;
            process process;
            protocol protocol;
            source-address source-address;
            source-port source-port;
            success;
            username username;
        }
        limit limit;
    }
    disable;
    escape;
    time-format (year | millisecond);
    event-rate logs per second;
    facility-override (authorization | daemon | ftp | kernel | local0 | local1 | local2 | local3 |
local4 | local5 | local6 | local7 | user);
    file {
        files files;
        name name;
        path path;
        size size;
    }
    format (binary | protobuf | sd-syslog | syslog);
    max-database-record max-database-record;
    message-rate-limit messages per second;
    mode (event | stream | stream-event);
    rate-cap logs per second;
    report {
        logs-per-table {
            idp idp;
            ipsec-vpn ipsec-vpn;
            screen screen;
            session-all session-all;
            sky sky;
            utm utm;
        }
        table-lifetime table-lifetime;
        table-mode {
            dense;
        }
```

```
    }
    retry-count retry-count;
    root-streaming;
    stream stream-name {
        category (all | content-security | fw-auth | screen | alg | nat | flow | sctp | gtp |
ipsec | idp | rtlog |pst-ds-lite | appqos |secintel |aamw);
        filter {
            threat-attack;
            }
        format (binary | protobuf | sd-syslog | syslog | welf);
        host {
            ip-address;
            port port-number;
            routing-instanceinstance-name;
        }
        radius {
            ipaddr {
                backup backup;
                primary primary;
            }
            port port number;
            secret secret;
        }
        rate-limit {
            log-rate;
        }
        severity (alert | critical | debug | emergency | error | info | notice | warning);
        source-address {
            ip-address;
        }
        time-format (year | millisecond);
        transport {
            protocol (tcp | tls | udp);
            tcp-connections tcp-connections;
            tls-profile tls-profile;
        }
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace;
    }
```

```
    transport {
        protocol (tcp | tls | udp);
        tcp-connections tcp-connections;
        tls-profile tls-profile;
    }
    utc-timestamp;
}
```

Hierarchy Level

```
[edit security]
[edit logical-systems name security]
[edit tenants tenant-name security]
```

Description

Configure security log. Set the mode of logging (event for traditional system logging or stream for streaming security logs through a revenue port to a server). You can also specify all the other parameters for security logging.

Options

cache Cache security log events in the audit log buffer.

disable Disable the security logging for the device.

escape Escapes the stream log forwarding to avoid parsing errors. Stream mode supports escape in `sd-syslog` and `binary` format. Event mode supports escape only in `binary` format.

time-format Specify the year, the millisecond, or both in the timestamp.

event-rate *rate* Limit the rate at which logs are streamed per second.

- **Range:** 0 through 1500
- **Default:** 1500

facility-override Alternate facility for logging to remote host.

file Specify the security log file options for logs in binary format.

- Values:
  - *max-file-number*—Maximum number of binary log files.
    - The range is 2 through 10 and the default value is 10.
  - *file-name*—Name of binary log file.
  - *binary-log-file-path*—Path to binary log files.
  - *maximum-file-size*—Maximum size of binary log file in megabytes.
    - The range is 1 through 10 and the default value is 10.

format Set the security log format for the device.

max-database-record The following are the disk usage range limits for the database:

- Range:
  - SRX1500, SRX4100, and SRX4200: 0 through 15,000,000
  - vSRX: 0 through 1,000,000
- Default:
  - SRX1500, SRX4100, and SRX4200: 15,000,000
  - vSRX: 1,000,000

  Be sure there is enough free space in `/var/log/hostlogs/`, otherwise logs might be dropped when written into the database.

mode Control how security logs are processed and exported.

rate-cap ***rate-cap-value*** Work with event mode only. This option limits the rate at which data plane logs are generated per second.

- **Range:** 0 through 5000 logs per second
- **Default:** 5000 logs per second

retry-count ***retry-count*** Specify the maximum number attempts to send accounting message to the RADIUS accounting server in case of a transmission failure.

root-streaming Allows the user logical systems to generate the logs using the root logical system's stream configuration.

source-address *source-address*	Specify a source IP address or IP address used when exporting security logs, which is mandatory to configure *stream host*.
source-interface *interface-name*	Specify a source interface name, which is mandatory to configure *stream host*. The `source-address` and `source-interface` are alternate values. Using one of the options is mandatory.
stream	Every stream can configure file or host.
traceoptions	Specify security log daemon trace options.
transport	Set security log transport settings.
utc-timestamp	Specify to use UTC time for security log timestamps.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

The [edit `logical-systems` *name* security] and [edit `tenants` *tenant-name* security] hierarchy levels introduced in Junos OS Release 19.1R1.

`escape` option added in Junos OS Release 20.2R1.

`root-streaming` option added in Junos OS Release 20.3R1.

`protobuf` option is introduced in Junos OS Release 23.2R1.

`retry-count` option is introduced in Junos OS Release 24.2R1.

log (Security GTP)

IN THIS SECTION

- Syntax | 10032
- Hierarchy Level | 10032
- Description | 10032
- Options | 10033
- Required Privilege Level | 10033
- Release Information | 10033

Syntax

```
log {
    forwarded (basic | detail);
    gtp-u name;
    prohibited (basic | detail);
    rate-limited (basic | detail);
    event (basic | detail);
}
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

Use the console or syslog to view GPRS tunneling protocol version 2 (GTPv2) traffic logs. You can configure the device to log GTPv2 packets based on their status. You can use the log configuration statement to enable GTPv2 logging on the device. By default, all logs are disabled on the device.

From Junos OS Release 20.4R1 onwards, the `edit security gprs gtp profile` *`profile-name`* hierarchy level is replaced by `edit security gtp profile` *`profile-name`*.

The `state-invalid (basic | detail)` option is replaced by `event (basic | detail)`.

Options

- **forwarded**—A packet that the security device transmitted because it was valid.
- **GTP-U**—A packet that the security devices have for gtp-u.
- **prohibited**—A packet that the security device dropped because it was invalid.
- **rate-limited**—A packet that the security device dropped because it exceeded the maximum rate limit of the destination GSN.
- **event**—A packet that the security device dropped because it failed stateful inspection, log alarms and tunnel management events.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. Support. The `state-invalid (basic | detail)` option is replaced by `event (basic | detail)` in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Understanding GTPv2 Traffic Logging

Example: Enabling GTPv2 Traffic Logging

log (Security IDP Sensor Configuration)

IN THIS SECTION

Syntax | 10034

Hierarchy Level | 10034

- Description | 10034
- Options | 10034
- Required Privilege Level | 10035
- Release Information | 10035

Syntax

```
log {
    cache-size size;
    suppression {
        disable;
        (include-destination-address | no-include-destination-address);
        max-logs-operate value;
        max-time-report value;
        start-log value;
    }
}
```

Hierarchy Level

```
[edit security idp sensor-configuration]
```

Description

Configure IDP security policy logs.

Options

The remaining statements are explained separately. See CLI Explorer.

cache-size Specify the size in bytes for each user's log cache.

- **Syntax:** `size`—Cache size.

- **Range:** 1 through 65,535 bytes
- **Default:** 12,800 bytes

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

log (Security Policies)

IN THIS SECTION

- Syntax | 10035
- Hierarchy Level | 10035
- Description | 10036
- Options | 10036
- Required Privilege Level | 10036
- Release Information | 10036

Syntax

```
log (session-close | session-init);
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then]
```

Description

Log traffic information for a specific policy. Traffic information is logged when a session begins (session-init) or closes (session-close).

> **NOTE**: You must configure security policy for the session using the `set security policies from-zone zone-name to-zone zone-name policy policy-name then log session-close` command to list all the applications and nested applications in Application Tracking on J-Web using the on-box reporting feature.

Options

session-close	Start logging traffic when the session ends.
session-init	Start logging traffic when the session begins.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

- *Security Policies Overview*
- *Understanding On-Box Logging and Reporting*
- *Example: Configuring Application Tracking*

log (Security SCTP)

IN THIS SECTION

- Syntax | 10037
- Hierarchy Level | 10037
- Description | 10037
- Options | 10037
- Required Privilege Level | 10038
- Release Information | 10038

Syntax

```
log name;
```

Hierarchy Level

```
[edit security sctp]
```

Description

Use the console or syslog to view Stream Control Transmission Protocol (SCTP) logs.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs sctp]` hierarchy level is replaced by `[edit security sctp]`.

Options

association	To log association events.
configuration	To log the CLI configuration.
control-message-all	To log both dropped and passed control messages.

control-message-drop	To log the dropped control messages.
data-message-drop	To log the dropped data messages.
decoding-error	To log the decoding errors.
dropped-packet	To log dropped packets and out of resource errors.
rate-limit	To log the rate limit.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2. The options `association`, `control-message-all`, `control-message-drop`, and `data-message-drop` added in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding Stream Control Transmission Protocol

SCTP Packet Structure Overview

log (Services)

IN THIS SECTION

- Syntax | 10039
- Hierarchy Level | 10039
- Description | 10039
- Options | 10039
- Required Privilege Level | 10040
- Release Information | 10040

Syntax

```
log {
    all;
    errors;
    info;
    sessions-allowed;
    sessions-dropped;
    sessions-ignored;
    sessions-whitelisted;
    warning;
}
```

Hierarchy Level

```
[edit services ssl proxy profile profile-name actions]
```

Description

Specify the logging actions. When configuring SSL proxy, you can choose to set the option to receive some or all of the logs.

SSL proxy logs contain the logical system name, SSL proxy allowlists, policy information, SSL proxy information, and other information that helps you troubleshoot when there is an error.

You can configure logging of all or specific events, such as error, warning, and information events. You can also configure logging of sessions that are allowlisted, dropped, ignored, or allowed after an error occurs.

Options

- `all`—Log all events.
- `errors`—Log all error events.
- `info`—Log all information events.
- `sessions-allowed`—Log SSL session allowed events after an error.
- `sessions-dropped`—Log only SSL session dropped events.

- `sessions-ignored`—Log session ignored events.
- `sessions-whitelisted`—Log SSL session allowlisted events.
- `warning`—Log all warning events.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

log (Spanning Trees)

IN THIS SECTION

- Syntax | 10040
- Hierarchy Level | 10041
- Description | 10041
- Required Privilege Level | 10041
- Release Information | 10041

Syntax

```
log;
```

Hierarchy Level

```
[edit protocols mstp interface (all | interface-name) bpdu-timeout-action],
[edit protocols rstp interface (all | interface-name) bpdu-timeout-action],
[edit protocols stp interface (all | interface-name) bpdu-timeout-action],
[edit protocols vstp vlan vlan-id interface (all | interface-name) bpdu-timeout-action]
```

Description

For interfaces configured for loop protection, configure the software to generate a message to be sent to the system log file **/var/log/messages** to record the loop-protection event.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

RELATED DOCUMENTATION

show spanning-tree bridge | 26548

show spanning-tree interface | 26555

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on non-ELS EX Series Switches

log-updown

IN THIS SECTION

- Syntax | 10042
- Hierarchy Level | 10042
- Description | 10043
- Required Privilege Level | 10043
- Release Information | 10043

Syntax

```
log-updown;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],

[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],

[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],

[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Specify to generate a log message whenever a BGP peer makes a state transition. Messages are logged using the system logging mechanism located at the `[edit system syslog]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Preventing BGP Session Resets

traceoptions (Protocols BGP) | 29662

log-updown (Protocols LDP)

IN THIS SECTION

- Syntax | 10044
- Hierarchy Level | 10044
- Description | 10044
- Options | 10044
- Required Privilege Level | 10044
- Release Information | 10044

Syntax

```
log-updown {
    trap disable;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Disable LDP traps on the router, logical system, or routing instance.

Options

`trap disable`—Disable LDP traps.

- **Default:** LDP traps are enabled on the router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Disabling SNMP Traps for LDP

log-updown (Protocols MPLS)

IN THIS SECTION

- Syntax | 10045
- Hierarchy Level | 10045
- Description | 10046
- Default | 10046
- Options | 10046
- Required Privilege Level | 10046
- Release Information | 10046

Syntax

```
log-updown {
    no-trap {
        mpls-lsp-traps;
        rfc3812-traps;
    }
    (syslog | no-syslog);
    trap;
    trap-path-down;
    trap-path-up;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Log a message or send an SNMP trap whenever an LSP makes a transition from up to down, or vice versa, and whenever an LSP switches from one active path to another. Only the ingress router performs these operations.

> **NOTE**: System log messages for LSPs are generated by default. To disable the default logging of messages for LSPs, configure the `no-syslog` option under the `log-updown` statement.

Default

There is no default behavior for this statement. If you do not specify the options, the configuration cannot be committed.

Options

`no-syslog`—Do not log a message to the system log file.

`no-trap`—Do not send an SNMP trap.

`syslog`—Log a message to the system log file.

`trap`—Send an SNMP trap.

`trap-path-down`—Send an SNMP trap when an LSP path goes down.

`trap-path-up`—Send an SNMP trap when an LSP path comes up.

The `no-trap` statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.
The `mpls-lsp-traps` and `rfc-3812-traps` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

System Log Messages and SNMP Traps for MPLS

Junos OS Network Management Administration Guide for Routing Devices

no-trap | **11604**

traceoptions (Protocols MPLS) | **29699**

log-warning (Multicast Forwarding Cache)

IN THIS SECTION

- Syntax | **10047**
- Hierarchy Level | **10047**
- Description | **10048**
- Options | **10048**
- Required Privilege Level | **10048**
- Release Information | **10048**

Syntax

```
log-warning value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast forwarding-cache threshold],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast forwarding-cache family (inet | inet6)threshold],
[edit logical-systems logical-system-name routing-options multicast forwarding-cache threshold],
[edit logical-systems logical-system-name routing-options multicast forwarding-cache family (inet
| inet6) threshold],
[edit routing-instances routing-instance-name routing-options multicast forwarding-cache threshold],
[edit routing-instances routing-instance-name routing-options multicast forwarding-cache family
```

```
(inet | inet6) threshold],
[edit routing-options multicast forwarding-cache threshold],
[edit routing-options multicast forwarding-cache family (inet | inet6) threshold]
```

Description

Specify the threshold at which the device logs a warning message in the system log for multicast forwarding cache entries. This threshold is a percentage of the maximum number of multicast forwarding cache entries received by the device. Configuring the `threshold` statement globally for the multicast forwarding cache or including the `family` statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.

To confirm the configured warning threshold, use the `show multicast forwarding-cache statistics` command.

Options

value—Percentage of the number of multicast forwarding cache entries that can be added to the cache that starts triggering the warning. You must explicitly configure the **suppress** value to configure a warning threshold value.

- **Range:** 1 through 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

> *Example: Configuring the Multicast Forwarding Cache*

log-warning (Protocols MSDP)

IN THIS SECTION

- Syntax | 10049
- Hierarchy Level | 10049
- Description | 10049
- Options | 10050
- Required Privilege Level | 10050
- Release Information | 10050

Syntax

```
log-warning value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp active-source-limit],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
active-source-limit],
[edit protocols msdp active-source-limit],
[edit routing-instances routing-instance-name protocols msdp active-source-limit]
```

Description

Specify the threshold at which the device logs a warning message in the system log for received MSDP active source messages. This threshold is a percentage of the maximum number of MSDP active source messages received by the device.

To confirm the configured warning threshold, use the `show msdp source-active` command.

Options

value—Percentage of the number of active source messages that starts triggering the warnings. You must explicitly configure the **maximum** value to configure a warning threshold value.

- **Range:** 1 through 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2

RELATED DOCUMENTATION

Example: Configuring MSDP with Active Source Limits and Mesh Groups

maximum (MSDP Active Source Messages) | **10536**

log-attacks

IN THIS SECTION

- Syntax | **10051**
- Hierarchy Level | **10051**
- Description | **10051**
- Options | **10051**
- Required Privilege Level | **10051**
- Release Information | **10051**

Syntax

```
log-attacks {
    alert;
}
```

Hierarchy Level

```
[edit logical-systems name security idp idp-policy name rulebase-exempt rule name then
notification],
[edit logical-systems name tenants name security idp idp-policy name rulebase-exempt rule name
then notification],
[edit security idp idp-policy name rulebase-exempt rule name then notification],
[edit tenants name security idp idp-policy name rulebase-exempt rule name then notification]
```

Description

Enable attack logging

Options

alert Set alert flag in attack log

Required Privilege Level

[none specified]

Release Information

log-attacks (rulebase-exempt)

IN THIS SECTION

- Syntax | 10052
- Hierarchy Level | 10052
- Description | 10052
- Options | 10053
- Required Privilege Level | 10053
- Release Information | 10053

Syntax

```
log-attacks {
    alert;
}
```

Hierarchy Level

```
[edit logical-systems name security idp idp-policy name rulebase-exempt rule name then
notification],
[edit logical-systems name tenants name security idp idp-policy name rulebase-exempt rule name
then notification],
[edit security idp idp-policy name rulebase-exempt rule name then notification],
[edit tenants name security idp idp-policy name rulebase-exempt rule name then notification]
```

Description

Enable exempt rule logging in the IDP system, which can be used to monitor and analyze traffic patterns, detect potential security threats, and troubleshoot network issues.

Options

alert	Set alert flag in attack log

Required Privilege Level

security:

To view this statement in the configuration.

security-control:

To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

logging (IDS MS-DPC)

IN THIS SECTION

- Syntax | 10053
- Hierarchy Level | 10054
- Description | 10054
- Options | 10054
- Required Privilege Level | 10054
- Release Information | 10054

Syntax

```
logging {
    syslog;
```

```
    threshold rate;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Set logging values for this IDS term when using the MS-DPC.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in IDS Rules

logging-rule (PCC Action Profile)

IN THIS SECTION

Syntax | 10055

- Hierarchy Level | 10055
- Description | 10055
- Options | 10055
- Required Privilege Level | 10055
- Release Information | 10056

Syntax

```
logging-rule lrf-rule-name;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],
[edit services pcef pcc-action-profiles profile-name]
```

Description

Assign the LRF rule to the PCC action profile of a static PCC rule. When the matching conditions in the PCC rule are met, the LRF rule is activated.

If you are using Junos OS Subscriber Aware, specify the name of the LRF rule at the `[edit unified-edge pcef pcc-action-profiles profile-name]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the LRF rule at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level.

Options

lrf-rule-name LRF rule name. The referenced LRF rule must be configured in an LRF profile.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Configuring the Activation of an LRF Rule by a PCC Rule

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring Logging and Reporting for Subscriber Management

Configuring an LRF Profile for Subscribers

logging (Services)

IN THIS SECTION

- Syntax | **10057**
- Hierarchy Level | **10057**
- Description | **10057**
- Options | **10057**
- Required Privilege Level | **10057**
- Release Information | **10057**

Syntax

```
logging {
    traceoptions {
        file filename <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Define global services properties.

Options

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Tracing Services PIC Operations

logical-bandwidth-policer

IN THIS SECTION

- Syntax | 10058
- Hierarchy Level | 10058
- Description | 10058
- Required Privilege Level | 10058
- Release Information | 10059

Syntax

```
logical-bandwidth-policer;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name],
[edit firewall policer policer-name],
[edit logical-systems logical-system-name firewall policer policer-name]
```

Description

For a policer with a bandwidth limit configured as a percentage (using the `bandwidth-percent` statement), specify that the percentage be based on the shaping rate defined on the logical interface, rather than on the media rate of the physical interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Logical systems support introduced in Junos OS Release 9.3.

Support at the [`edit dynamic-profiles ... policer` *`policer-name`*] hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Bandwidth Policers
- Configuring Policers Based on Logical Interface Bandwidth
- bandwidth-percent | **2866**
- interface-specific

logical-domain-identity-management

IN THIS SECTION

- Syntax | **10059**
- Hierarchy Level | **10061**
- Description | **10061**
- Options | **10061**
- Required Privilege Level | **10061**
- Release Information | **10061**

Syntax

```
logical-domain-identity-management {
    active {
        authentication-entry-timeout minutes;
        filter {
            domain name;
```

```
            exclude-ip {
                address-book book-name;
                address-set address-set;
            }
            include-ip {
                address-book book-name;
                address-set address-set;
            }
        }
        invalid-authentication-entry-timeout minutes;
        ip-query {
            query-delay-time seconds;
        }
        query-server name {
            batch-query {
                items-per-batch items-per-batch;
                query-interval seconds;
            }
            connection (Identity Management Advanced Query){
                connect-method (http | https);
                port port;
                primary {
                    address address;
                    ca-certificate ca-certificate;
                    client-id client-id;
                    client-secret client-secret;
                }
                query-api query-api;
                secondary {
                    address address;
                    ca-certificate ca-certificate;
                    client-id client-id;
                    client-secret client-secret;
                }
                token-api token-api;
            }
        }
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        level (all | error | info | notice | verbose | warning);
```

```
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Specify the logical domain identity management configuration for both logical systems and tenant systems.

Options

active Displays the active mode for logical domain identity management module.

traceoptions Displays the tracing options.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

> *User Authentication for Logical Systems*
>
> *Understanding Integrated User Firewall Support in a Tenant System*

logical-interface-aggregate-statistics

IN THIS SECTION

- Syntax | 10062
- Hierarchy Level | 10062
- Description | 10062
- Required Privilege Level | 10063
- Release Information | 10063

Syntax

```
logical-interface-aggregate-statistics;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name]
```

Description

Enable the collection of aggregate queue statistics for all underlying logical interfaces on a particular physical or aggregate (for example, ae0) interface. By default to preserve memory resources, aggregate queue statistics are not collected for underlying logical interfaces (Level 2 interfaces); queue statistics are only collected for the upper-level logical interfaces (Level 3 and above) and the physical interface (Level 1). The command `show interfaces queue` *`interface-name`* shows all zeroes for underlying logical interfaces by default.

CAUTION: To enable a hierarchical CoS feature on a physical interface with an already active interface hierarchy:

1. "deactivate" on page 5142 `class-of-service` on all hierarchy levels of the physical interface.

2. Enable the hierarchical CoS setting on the physical interface.
3. Re-"activate" on page 1795 `class-of-service` on all hierarchy levels of the physical interface.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

- show interfaces queue | 21208
- *Hierarchical Class of Service Overview*
- show class-of-service scheduler-hierarchy interface | 19119

logical-interface-chassis-redundancy (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 10063
- Hierarchy Level | 10064
- Description | 10064
- Required Privilege Level | 10064
- Release Information | 10064

Syntax

```
logical-interface-chassis-redundancy;
```

Hierarchy Level

```
[edit interfaces aenumber aggregated-ether-options]
```

Description

For member routers in an MX Series Virtual Chassis, provide chassis redundancy for IP demultiplexing (demux) and VLAN demux subscriber interfaces in aggregated Ethernet bundles configured with targeted distribution. With chassis redundancy, the router assigns the backup link in the aggregated Ethernet bundle to an MPC/MIC module in a member router *other* than the router on which the primary link resides.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Chassis Redundancy for a Virtual Chassis

logical-interface (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 10065
- Hierarchy Level | 10065
- Description | 10065
- Options | 10065

Required Privilege Level | 10066
Release Information | 10066

Syntax

```
logical-interface (flow-bandwidth | flow-control-mode | flow-detection-mode)
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type flow-level-bandwidth],
[edit system ddos-protection protocols protocol-group packet-type flow-level-control],
[edit system ddos-protection protocols protocol-group packet-type flow-level-detection]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure flow bandwidth, flow control mode, or flow detection mode for flow detection at the logical interface flow aggregation level for the packet type.

Options

flow-bandwidth—Bandwidth for the flow at the logical interface level. Available only at the `[edit system ddos-protection protocols protocol-group packet-type flow-level-bandwidth]` hierarchy level.

- **Default:** 200 packets per second
- **Range:** 1 through 30,000 packets per second

flow-control-mode—Mode for how traffic in the detected flow is controlled at the logical interface level. Available only at the `[edit system ddos-protection protocols protocol-group packet-type flow-level-control]` hierarchy level.

> **NOTE**: The configuration at this level overrides the global configuration using the `flow-level-control` statement at the `[edit system ddos-protection global]` hierarchy level.

- `drop`—Drop all traffic in flow.

- `keep`—Keep all traffic in flow.
- `police`—Police the traffic to within its allowed bandwidth.
- **Default:** `drop`

`flow-detection-mode`—Mode for how flow detection operates at the logical interface level when a policer has been violated. Available only at the `[edit system ddos-protection protocols` *`protocol-group packet-type`* `flow-level-detection]` hierarchy level.

> **NOTE**: The configuration at this level overrides the global configuration using the `flow-detection-mode` statement at the `[edit system ddos-protection global]` hierarchy level.

- `automatic`—Search flows at the logical interface level only when a DDoS policer is being violated and only when the flow causing the policer violation is not discovered at the finer flow aggregation level, subscriber. When the suspicious flow is not found at this level, then the search moves to a coarser level of flow aggregation (physical interface). Flows at the logical interface level are subsequently not searched again until the policer is no longer violated at the coarser level, and a subsequent violation occurs that cannot be found at the subscriber level.
- `off`—Disable flow detection at the logical interface level so that flows are never searched at this level.
- `on`—Search flows at the logical interface level, even when no DDoS protection policer is currently being violated. Monitoring continues at this level regardless of whether a suspect flow is identified at this level.
- **Default:** `automatic`

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring the Maximum Flow Bandwidth at Each Flow Aggregation Level

Configuring How Traffic in a Culprit Flow Is Controlled at Each Flow Aggregation Level

Configuring How Flow Detection Operates at Each Flow Aggregation Level

Setting Up and Using Flow Detection

logical-interface-fpc-redundancy (Aggregated Ethernet Subscriber Interfaces)

IN THIS SECTION

- Syntax | 10067
- Hierarchy Level | 10067
- Description | 10067
- Required Privilege Level | 10068
- Release Information | 10068

Syntax

```
logical-interface-fpc-redundancy;
```

Hierarchy Level

```
[edit interfaces aenumber aggregated-ether-options]
```

Description

Provide module redundancy for demux subscribers on aggregated Ethernet bundles configured with targeted distribution. Backup links for a subscriber are chosen on a different EQ DPC or MPC from the primary link, based on the link with the fewest number of subscribers among the links on different modules. If all links are on a single module when this is configured, backup links are not provisioned.

By default, link redundancy is provided for the aggregated Ethernet bundle.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

- Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface
- Configuring Module Redundancy for a Virtual Chassis

logical-interface-fpc-redundancy (PS Interface)

IN THIS SECTION

- Syntax | 10068
- Hierarchy Level | 10069
- Description | 10069
- Required Privilege Level | 10069
- Release Information | 10069

Syntax

```
logical-interface-fpc-redundancy;
```

Hierarchy Level

```
[edit interfaces ps number targeted-options]
```

Description

Provide module redundancy for demux subscribers on redundant logical tunnel (RLT) bundles configured with targeted distribution. Backup links for a subscriber are chosen on a different Packet Forwarding Engine from the primary link, based on the link with the fewest number of subscribers among the links on different modules. If all links are on a single module when this is configured, backup links are not provisioned.

By default, link redundancy is provided for the RLT bundle.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

MPLS Pseudowire Subscriber Logical Interfaces

logical-interface-policer

IN THIS SECTION

- Syntax | 10070
- Hierarchy Level | 10070
- Description | 10070

Required Privilege Level | 10071
Release Information | 10071

Syntax

```
logical-interface-policer;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name],
[edit dynamic-profiles profile-name firewall three-color-policer name],
[edit firewall atm-policeratm-policer-name],
[edit firewall policer policer-name],
[edit firewall policer policer-template-name],
[edit firewall three-color-policer policer-name],
[edit logical-systems logical-system-name firewall policer policer-name],
[edit logical-systems logical-system-name firewall three-color-policer name]
```

Description

Configure a logical interface policer. For PTX series routers running Junos OS Release 18.3R1 or later, you can use this command to configure separate firewall filters for different family address types (IPv4 and IPv6) that share the same interface, and configure the same policer as an action for the filter.

To configure the aggregate policer, configure the firewall policer you want to use as `logical-interface-policer`. And at the `firewall family` *`family-name`* `filter` *`filter-name`* hierarchy level where you will reference the policer, make the policer an `interface-specific` firewall filter action.

The sample configuration shows the relationship.

```
firewall {
    policer Shared_Policer {
      logical-interface-policer;
        if-exceeding {
            bandwidth-limit 100m;
            burst-size-limit 500k;
```

```
        }
        then {
            discard;
        }
    }
}
```

```
family inet {
    filter filter_name{
       interface-specific;
        term term_name {
            then {
                policer Shared_Policer;
                count cinet;
            }
        }
    }
}
```

NOTE: Starting in Junos OS Release 12.2R2, on T Series Core Routers only, you can configure an MPLS LSP policer for a specific LSP to be shared across different protocol family types. You must include the `logical-interface-policer` statement to do so.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit firewall three-color-policer policer-name]` hierarchy level introduced in Junos OS Release 8.2.

Logical systems support introduced in Junos OS Release 9.3.

Support at the `[edit dynamic-profiles ... policer policer-name]` and `[edit dynamic-profiles ... three-color-policer name]`hierarchy levels introduced in Junos OS Release 11.4.

Support for PTX series routers with third-generation FPCs added in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

- Two-Color and Three-Color Logical Interface Policers
- Traffic Policer Types
- Configuring and Applying Tricolor Marking Policers
- action | 1759
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers
- action

logical-interface-unit-range

IN THIS SECTION

- Syntax | 10072
- Hierarchy Level | 10072
- Description | 10073
- Options | 10073
- Required Privilege Level | 10073
- Release Information | 10073

Syntax

```
logical-interface-unit-range (high | low)
```

Hierarchy Level

```
[edit system services extensible-subscriber-services]
```

Description

Configure the range from which the unit number is selected for the logical interface service that is created by Extensible Subscriber Services Manager by using an op script. Extensible Subscriber Services Manager assigns the first available unit number in the specified range.

Options

`high` Upper limit of the logical interface unit range.

- **Range:** 1 through 16,385
- **Default:** 16,385

`low` Lower limit of the logical interface unit range.

- **Range:** 1 through 16,385
- **Default:** 1

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

logical-system

IN THIS SECTION

- Syntax | 10074
- Hierarchy Level | 10074
- Description | 10074
- Options | 10074

Required Privilege Level | 10074
Release Information | 10074

Syntax

```
logical-system logical-system-name {
    [ routing-instances instance-name ];
}
```

Hierarchy Level

```
[edit services rpm bgp]
```

Description

Specify the logical system used by the probes.

Options

logical-system-name—Logical system name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring BGP Neighbor Discovery Through RPM

logical-system (SNMP)

IN THIS SECTION

- Syntax | 10075
- Hierarchy Level | 10075
- Description | 10075
- Options | 10076
- Required Privilege Level | 10076
- Release Information | 10076

Syntax

```
logical-system logical-system-name {
    routing-instance routing-instance-name {
        source-address address;
    }
}
```

Hierarchy Level

```
[edit snmp community community-name],
[edit snmp trap-group],
[edit snmp trap-options]
[edit snmp v3target-address target-address-name]
```

Description

Specify a logical system name for SNMP v1 and v2c clients.

Include at the `[edit snmp trap-options]` hierarchy level to specify a logical-system address as the source address of an SNMP trap.

Include at the `[edit snmp v3 target-address]` hierarchy level to specify a logical-system name as the destination address for an SNMPv3 trap or inform.

Options

`logical-system-name`–Name of the logical system.

`routing-instance` *`routing-instance-name`*–Statement to specify a routing instance associated with the logical system.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3

> **NOTE**: The `logical-system` statement replaces the `logical-router` statement, and is backward-compatible with Junos OS Release 8.3 and later.

RELATED DOCUMENTATION

Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community

Configuring the Trap Target Address

logical-system (Diameter Peer)

IN THIS SECTION

Syntax | 10077

- Hierarchy Level | 10077
- Description | 10077
- Options | 10077
- Required Privilege Level | 10077
- Release Information | 10078

Syntax

```
logical-system  logical-system-name  <routing-instance routing-instance-name > ;
```

Hierarchy Level

```
[edit diameter peer peer-name]
```

Description

Specify a logical system and optionally a routing instance for a Diameter peer. Alternatively, you can include the `routing-instance` statement at the `[edit diameter peer peer-name]` hierarchy level to configure only a routing instance.

Options

`logical-system-name`— Name of the logical system.

- **Default:** Default logical system

`routing-instance` *`routing-instance-name`*—(Optional) Name of the routing instance.

- **Default:** Master routing instance

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Peers

logical-system (Diameter Transport)

IN THIS SECTION

- Syntax | 10078
- Hierarchy Level | 10078
- Description | 10078
- Options | 10079
- Required Privilege Level | 10079
- Release Information | 10079

Syntax

```
logical-system logical-system-name <routing-instance routing-instance-name >;
```

Hierarchy Level

```
[edit diameter transport transport-name]
```

Description

Specify a logical system and optionally a routing instance for the transport layer connection.

> **NOTE**: The logical system and routing instance must match that for the peer or a configuration error is reported.

Options

logical-system-name—Name of the logical system.

- **Default:** Default logical system

routing-instance *routing-instance-name*—(Optional) Name of the routing instance.

- **Default:** Master routing instance

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Diameter

Configuring the Diameter Transport

logical-system-mux

IN THIS SECTION

- Syntax | 10080
- Hierarchy Level | 10080
- Description | 10081

Options | 10081
Required Privilege Level | 10081

Syntax

```
logical-system-mux {
    traceoptions {
        file {
            <file name>;
            files;
            no-world-readable;
            size;
            world-readable;
            }

        flag {
            all;
            debug;
            general;
            normal;
            parse;
            policy;
            route;
            state;
            task;
            timer;
            }

        }
    }
```

Hierarchy Level

```
[edit routing-options]
```

Description

Display the logical system control daemon information.

Options

<filename>— Specify the filename, in which trace information should be written.

files—Specify the maximum number of trace files. Range being from 2 to 1000.

no-world-readable—Do not allow any user to read the log file.

size—Specify the maximum size of trace file. Range being from 10240 to 4294967295 bytes.

world-readable—Allow any user to read the log file.

all—Trace everything

debug—Trace debug-level lrmuxd activities.

general—Trace general events.

normal—Trace normal events.

parse—Trace configuration parsing.

policy—Trace policy parsing.

route—Trace routing information.

state—Trace state transitions.

task—Trace routing protocol task processing.

timer—Trace routing protocol timer processing.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

logical-system-name (DHCP Local Server)

IN THIS SECTION

- Syntax | 10082
- Hierarchy Level | 10082
- Description | 10082
- Required Privilege Level | 10082
- Release Information | 10083

Syntax

```
logical-system-name;
```

Hierarchy Level

```
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...]
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Specify that the logical system name be concatenated with the username during the subscriber authentication or DHCP client process. No logical system name is concatenated if the configuration is in the default logical system.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

logical-system-name (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 10083
- Hierarchy Level | 10083
- Description | 10084
- Required Privilege Level | 10084
- Release Information | 10084

Syntax

```
logical-system-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-
include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
```

```
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-
include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication
username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
authentication username-include]
```

Description

Specify that the logical system name is concatenated with the username during the subscriber authentication or client authentication process. No logical system name is concatenated if the configuration is in the default logical system. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.

Support at the [edit ... dual-stack-group *dual-stack-group-name*] hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Specifying Authentication Support

Creating Unique Usernames for DHCP Clients

logical-system-name (Static Subscribers)

IN THIS SECTION

- Syntax | 10085
- Hierarchy Level | 10085
- Description | 10086
- Required Privilege Level | 10086
- Release Information | 10086

Syntax

```
logical-system-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication username-include],
[edit logical-systems logical-system-name system services static-subscribers  authentication
username-include],
[edit logical-systems logical-system-name system services static-subscribers group group-name
```

```
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],
[edit system services static-subscribers authentication username-include],
[edit system services static-subscribers group group-name authentication username-include]
```

Description

Specify that the name of the logical system is included as part of the username created for all static subscribers or for the static subscribers in a specified group. The group version of the statement takes precedence over the global version. The username is also sent to RADIUS in the Access-Request message.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Configuring the Static Subscriber Global Username

Configuring the Static Subscriber Group Username

logical-system (System Security Profile)

IN THIS SECTION

Syntax | 10087

- Hierarchy Level | 10087
- Description | 10087
- Options | 10087
- Required Privilege Level | 10087
- Release Information | 10088

Syntax

```
logical-system logical-system;
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the user logical system to bind the security profile.

The primary administrator uses security profiles to provision logical systems with resources. You can bind security profiles to user logical systems and the primary logical system. The primary administrator can configure more than one security profile allocating different amounts of a resource in various ones.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

logical-system-name–Name of the logical system.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways

logical-system-trap-filter

IN THIS SECTION

- Syntax | 10088
- Hierarchy Level | 10088
- Description | 10088
- Required Privilege Level | 10089
- Release Information | 10089

Syntax

```
logical-system-trap-filter;
```

Hierarchy Level

```
[edit snmp]
```

Description

Restrict the routing instances from receiving traps that are not related to the logical system networks to which they belong.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

SNMP Traps Supported for Routing Instances

logical-system (Tunnel Profile)

IN THIS SECTION

- Syntax | 10089
- Hierarchy Level | 10089
- Description | 10090
- Options | 10090
- Required Privilege Level | 10090
- Release Information | 10090

Syntax

```
logical-system  logical-system-name;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Description

Specify a logical system for a tunnel. When you specify a logical system, you must also specify a routing instance.

Options

logical-system-name— Name of the logical system.

- **Default:** Logical system *default*

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

logical-systems

IN THIS SECTION

- Syntax | 10091
- Hierarchy Level | 10091
- Description | 10091
- Options | 10091
- Required Privilege Level | 10091
- Release Information | 10091

Syntax

```
logical-systems {
    logical-system-name {
        ...logical-system-configuration...
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure a logical system.

Options

logical-system-name—Name of the logical system.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement name changed from `logical-routers` in Junos OS Release 9.3.

RELATED DOCUMENTATION

Logical Systems User Guide for Routers and Switches

logical-systems (All)

IN THIS SECTION

- Syntax | 10092
- Hierarchy Level | 10092
- Description | 10092
- Options | 10093
- Required Privilege Level | 10093
- Release Information | 10093

Syntax

```
logical-systems {
    logical-system-name {
        ...logical-system-configuration...
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure a logical system. Only the primary administrator can configure a logical system at [`edit`] hierarchy level.

You can include several of the hierarchies that can be included at the [`edit`] hierarchy level. For descriptions of the applicable statements, see the appropriate hierarchies.

NOTE: The `logical-systems` configuration statement can be used only by the primary administrator.

Options

logical-system-name—Name of the logical system.

Required Privilege Level

all—To view this statement in the configuration.

all—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways

logical-tunnel-options

IN THIS SECTION

- Syntax | 10093
- Hierarchy Level | 10094
- Description | 10094
- Options | 10094
- Required Privilege Level | 10095
- Release Information | 10095

Syntax

```
logical-tunnel-options {
    link-protection {
        non-revertive;
```

```
        revertive;
    }
    load-balance{
        adaptive{
            pps;
            scan-interval scan-interval;
            tolerance percent;
        }
        no-adaptive;
        per-packet;
        local-biaspercent;
    }
    per-unit-mac-disable;
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces],
[edit dynamic-profiles name logical-systems name interfaces],
[edit interfaces]
```

Description

For redundant logical tunnels, specifies the logical tunnel interface-specific options for load balancing and link protection. The remaining statements are explained separately. See CLI Explorer.

Options

link-protection Enables link protection for redundant logical tunnel interfaces. In addition to enabling static link protection, you must configure a primary and secondary (backup) link for egress traffic.

- Values:
  - non-revertive—Do not revert back from active backup link to primary, if primary is UP.
  - revertive—Revert back from active backup link to primary, if primary is UP.
- **Default:** revertive

per-unit-mac-disable Disable the creation of per unit mac address on LT IFLs for VPLS/CCC encaps

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 19.2R1.

login

IN THIS SECTION

- Syntax | 10095
- Hierarchy Level | 10098
- Description | 10098
- Options | 10098
- Required Privilege Level | 10100
- Release Information | 10100

Syntax

```
login {
    announcement text;
    class class-name {
        allow-hidden-commands;
        no-hidden-commands {
            except ["regular expression or command 1" "regular expression or command
2" ...];
        }
        access-end hh:mm;
        access-start hh:mm;
```

```
        ( allow-commands "(regular-expression1)|(regular-expression2)..." | allow-commands-
regexps ["regular expression 1" "regular expression 2 " ... ]);
        ( allow-configuration "(regular-expression1)|(regular-expression2)..." | allow-
configuration-regexps ["regular expression 1" "regular expression 2 " ... ]);
        allow-sources [ source-addresses ... ];
        allow-times [ times ... ];
        allowed-days [ days of the week ];
        cli {
            prompt prompt;
        }
        configuration-breadcrumbs;
        confirm-commands ["regular expression or command 1" "regular expression or command
2" ...] {
            confirmation-message;
        }
        ( deny-commands "(regular-expression1)|(regular-expression2)..." | deny-commands-
regexps ["regular expression 1" "regular expression 2 " ... ]);
        ( deny-configuration "(regular-expression1)|(regular-expression2)..." | deny-
configuration-regexps ["regular expression 1" "regular expression 2 " ... ]);
        deny-sources [ source-addresses ... ];
        deny-times [ times ... ];
        idle-timeout minutes;
        logical-system logical-system-name;
        login-alarms;
        login-script login-script;
        login-tip;
        no-scp-server;
        no-sftp-server;
        permissions [ permissions ];
        satellite all;
        security-role (audit-administrator | crypto-administrator | ids-administrator | security-
administrator);
        tenant tenant-system-name;
    }
    deny-sources {
        address [ source-addresses ... ];
    }
    idle-timeout minutes;
    message text;
    password {
        change-type (character-sets | set-transitions);
        format (sha256 | sha512);
        maximum-length length;
```

```
        maximum-lifetime days

        minimum-changes number;
        minimum-character-changes number
        minimum-length length;
        minimum-lifetime days
        minimum-lower-cases number;
        minimum-numerics number;
        minimum-punctuations number;
        minimum-reuse number;
        minimum-upper-cases number;
    }
    retry-options {
        backoff-factor seconds;
        backoff-threshold number;
        lockout-period minutes;
        maximum-time seconds;
        minimum-time seconds;
        tries-before-disconnect number;
    }
    user username {
        authentication {
            encrypted-password encrypted-password;
            no-public-keys;
            ssh-ecdsa name {
                from from;
            }
            ssh-ed25519 name {
                from from;
            }
            ssh-rsa name {
                from from;
            }
        }
        cli {
            prompt prompt;
        }
        class class-name;
        full-name full-name;
        uid uid-value;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure user access to the device.

Options

announcement ***text***

Configure a system login announcement. This announcement appears after a user logs in. Sometimes you want to make announcements to authorized users only after they have logged in. For example, you might want to announce an upcoming maintenance event.

To display a message before the user logs in, configure a system login message using the `message` statement rather than configuring a system login announcement.

You can format the announcement using the following special characters:

- \n—New line
- \t—Horizontal tab
- \'—Single quotation mark
- \"—Double quotation mark
- \\—Backslash

If the text of the announcement contains any spaces, enclose the text in quotation marks.

- **Default:** No login announcement is displayed.

deny-sources

(Mandatory) Never allow access from these hosts. The source addresses can be IPv4 or IPv6 addresses, prefix lengths, or hostnames.

- **Syntax:** address [*source-addresses*]

idle-timeout ***minutes***

For a login class, configure the maximum time in minutes that a session can be idle before the session times out and the user is logged out of the device. The session times out after remaining at the CLI operational mode prompt for the specified time.

> **NOTE**: After the user logs in to a device from a shell prompt such as csh, if the user starts another program to run in the foreground of the CLI, the idle-timer control is stopped from being computed. The calculation of the idle time of the CLI session is restarted only after the foreground process exits and the control is returned to the shell prompt. When the restart of the idle-timer control occurs, if no interaction from the user occurs on the shell, the user is automatically logged out after the time set on this statement.

- **Default:** If you omit this statement, a user is never forced off the system after extended idle times.
- **Range:** Range: 0 through 4294967295 minutes

> **NOTE**: The idle-timeout feature is disabled if the value of *minutes* is set to 0.

message *text*

Configure a system login message. A login message displays a banner to users when they access the device, before they log in. To display a message only after the user logs in, configure a system login announcement using the `announcement` statement instead of configuring a system login message.

Before you create any user accounts, it's a good idea to configure an initial login message.

You can format the message using the following special characters:

- \n—New line
- \t—Horizontal tab
- \'—Single quotation mark
- \"—Double quotation mark
- \\—Backslash

If the text of the message contains any spaces, enclose the text in quotation marks.

- **Default:** No login message is displayed.

The remaining statements are explained separately. See CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`deny-sources` option introduced in Junos OS Release 11.2.

All of the statements and options introduced previously were introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

RELATED DOCUMENTATION

- Defining Junos OS Login Classes
- Configuring Junos OS to Display a System Login Announcement

login (Access)

IN THIS SECTION

- Syntax | 10100
- Hierarchy Level | 10101
- Description | 10101
- Options | 10101
- Required Privilege Level | 10101
- Release Information | 10101

Syntax

```
login string;
```

Hierarchy Level

```
[edit access firewall-authentication pass-through default-profile profile-name (ftp | http |
telnet) banner]
```

Description

Specify the login banner for users using FTP, HTTP, and Telnet during the authentication process.

Options

string—Banner text. Maximum length of the message text is 250 characters. Enclose the banner text within spaces or special characters—for example quotation marks (” “).

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Obtaining Username and Role Information Through Firewall Authentication

logout-response-timeout (PCRF Partition)

IN THIS SECTION

- Syntax | 10102
- Hierarchy Level | 10102

- Description | 10102
- Options | 10102
- Required Privilege Level | 10103
- Release Information | 10103

Syntax

```
logout-response-timeout seconds;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the amount of time in seconds before a Policy and Charging Rule Function (PCRF) partition stops attempting to send a subscriber logout.

If you set `draining` and the `draining-response-timeout` statements for the PCRF partition, any new subscriber logins are denied, and the time limit you set in the `draining-response-timeout` statement is used instead of the `logout-response-timeout` time limit.

> **NOTE**: Any configuration changes made to this statement apply to all subscribers currently waiting to log out within a 60 second period.

Options

seconds Number of seconds to wait before a PCRF partition stops attempting to send a subscriber logout.

- **Default:** 7200
- **Range:** 0 through 86,400 seconds (24 hours)

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

Understanding Upstream and Downstream Messages for the PCRF

long-lived (Graceful Restart for BGP Helper)

IN THIS SECTION

- Syntax | 10103
- Hierarchy Level | 10104
- Description | 10104
- Required Privilege Level | 10105
- Release Information | 10105

Syntax

```
long-lived {
    receiver {
        enable:
```

```
        disable;
    }
    advertise-to-non-llgr-neighbor {
        omit-no-export;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
graceful-restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp graceful-
restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name graceful-restart],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name neighbor address graceful-restart],
[edit protocols bgp graceful-restart],
[edit protocols bgp group group-name graceful-restart],
[edit protocols bgp group group-name neighbor address graceful-restart]
```

Description

Configure the long-lived graceful restart mechanism to preserve BGP routing details for a longer period from a failed BGP peer than the duration for which such routing information is maintained using the BGP graceful restart functionality. Long-lived graceful restart receiver or helper mode is enabled by default, unless ordinary graceful restart receiver or helper mode is disabled.

The `long-lived receiver enable` overrides a disable option inherited from a higher level in the configuration. It does not enable long-lived graceful restart restarter mode for all families—restarter mode must be configured explicitly for each family. When the LLGR receiver or helper mode is enabled or disabled, the session is reset. This behavior enables the new capability value to be sent to the neighbor.

> **NOTE**: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.

Configure graceful restart globally at the `[edit routing-options]` or `[edit routing-instances instance-name routing-options]` hierarchy level to enable the feature. You cannot enable graceful restart for specific

protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.

We have enhanced the long-lived graceful restart (LLGR) capabilities for BGP Helper and introduced the configuration statement `extended-route-retention` at `[edit protocols bgp group neighbor graceful-restart long-lived]` hierarchy level. When you enable this feature, `extended-route-retention` supports LLGR helper mode regardless of the BGP peer LLGR capabilities.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

long-lived (Graceful Restart for BGP Restarter)

IN THIS SECTION

- Syntax | 10106
- Hierarchy Level | 10106
- Description | 10107
- Required Privilege Level | 10107
- Release Information | 10107

Syntax

```
long-lived {
    restarter {
        disable;
        stale-time interval;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family inet (labeled-unicast | unicast |
multicast) graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name family inet (labeled-
unicast | unicast | multicast) graceful-restart],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
inet (labeled-unicast | unicast | multicast) graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family inet (labeled-unicast | unicast | multicast) graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family inet (labeled-unicast | unicast | multicast) graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family inet (labeled-unicast | unicast | multicast) graceful-
restart],
[edit routing-instances routing-instance-name protocols bgp family inet (labeled-unicast |
unicast | multicast) graceful-restart],
[edit routing-instances routing-instance-name protocols bgp group group-name family inet
(labeled-unicast | unicast | multicast) graceful-restart],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family inet (labeled-unicast | unicast | multicast) graceful-restart],
[edit protocols bgp family inet (labeled-unicast | flow) graceful-restart long-lived restarter],
[edit protocols bgp group group-name family inet (labeled-unicast | flow) graceful-restart long-
lived restarter],
[edit protocols bgp group group-name neighbor address family inet (labeled-unicast | flow)
graceful-restart long-lived restarter]
```

Description

Configure the long-lived graceful restart mechanism on the restarter router to preserve BGP routing details for a longer period from a failed BGP peer than the duration for which such routing information is maintained using the BGP graceful restart functionality. Long-lived graceful restart restarter mode is enabled by default, unless ordinary graceful restart on the restarter router is disabled.

> **NOTE**: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.

Configure graceful restart globally at the `[edit routing-options]` or `[edit routing-instances instance-name routing-options]` hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

> *Configuring Graceful Restart Options for BGP*
>
> Junos OS High Availability User Guide

longest-match

IN THIS SECTION

- Syntax | 10108
- Hierarchy Level | 10108
- Description | 10108
- Options | 10108
- Required Privilege Level | 10108
- Release Information | 10109

Syntax

```
longest-match {
    policy value |(expression) | [values ];
}
```

Hierarchy Level

```
[edit protocols ldp]
```

Description

Enable longest match to allow LDP to learn the routes aggregated or summarized across OSPF areas or IS-IS levels in inter-domain.

Options

policy *value* **|(***expression***)	[***values***]**	Specify policy to provide per prefix granularity.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Longest Match Support for LDP Overview

Example: Configuring Longest Match for LDP

Configuring Longest Match for LDP

longest-match (Static Routes)

IN THIS SECTION

- Syntax | 10109
- Hierarchy Level | 10109
- Description | 10110
- Required Privilege Level | 10111
- Release Information | 10111

Syntax

```
longest-match;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options static route destination-prefix next-hop
address resolve]
[edit routing-instances routing-instance-name routing-options static route destination-prefix next-
```

```
hop address resolve],
[edit routing-options static route destination-prefix next-hop address resolve]
```

Description

Specify the static route on the device to resolve and determine the packet's next-hop interface using the Longest Match Routing Rule (most specific entry), sometimes referred to as the longest prefix match or maximum prefix length match. The Longest Match Routing Rule is an algorithm used by IP routers to select an entry from a routing table. The router uses the longest (prefix) match to determine the egress (outbound) interface and the address of the next device to which to send a packet. Typically, the static route prefers the directly connected subnet route for resolving the next hop rather than performing a longest prefix match with any other available routes.

> **NOTE**: (Required) You must include the `resolve` next-hop option to specify the `longest-match` statement. Next-hop options define additional information about static routes that are included with the route when it is installed in the routing table. You alter the default next-hop resolution behavior using the `resolve` next-hop option.

The router implements the Longest Match Routing Rule as follows:

1. The router receives a packet.
2. While processing the header, the router compares the destination IP address, bit-by-bit, with the entries in the routing table.

The entry that has the longest number of network bits that match the IP destination address is always the best match (or best path) as shown in the following example:

Longest Match Example

- The router receives a packet with a destination IP address of 192.168.1.33.
- The routing table contains the following possible matches:
  - 192.168.1.32/28
  - 192.168.1.0/24
  - 192.168.0.0/16

To determine the longest match, it's easiest to convert the IP addresses in Table 56 on page 10111 to binary and compare them.

Table 56: Converted IP Addresses

Address	Converted Binary Address
192.168.1.33 (destination IP address)	**11000000.10101000.00000001.001**00001
192.168.1.32/28	**11000000.10101000.00000001.001**00000 (<–Best match)
192.168.1.0/24	**11000000.10101000.00000001.**00000000
192.168.0.0/16	**11000000.10101000.**00000000.00000000

NOTE: When determining the next-hop interface for customer deployments, setting the `longest-match` statement results in traffic loss.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Static Route Preferences and Qualified Next Hops

lookup-intact-on-commit

IN THIS SECTION

- Syntax | 10112
- Hierarchy Level | 10112
- Description | 10112
- Required Privilege Level | 10113
- Release Information | 10113

Syntax

```
lookup-intact-on-commit;
```

Hierarchy Level

```
[edit security policies]
```

Description

Activate to prevent interruptions in the session flow due to modifications in the security policy. Once you commit the configuration, restart the forwarding plane to activate the functionality.

Use the following command to check the status and eligibility of the device before enabling the `lookup-intact-on-commit` option.

```
[edit]
user@host> show security policies lookup-intact-on-commit
```

The command output displays if the `lookup-intact-on-commit` option is already configured on the device and displays eligibility of the device in terms of available memory storage for activating `lookup-intact-on-commit` option.

Required Privilege Level

security—To view this statement in the configuration

security-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 24.4R1

loop-detect

IN THIS SECTION

- Syntax | 10113
- Hierarchy Level | 10114
- Description | 10114
- Options | 10115
- Required Privilege Level | 10116
- Release Information | 10116

Syntax

```
loop-detect {
    enhanced {
        interface (name | all) {
            loop-detect-action (interface-down | laser-off);
            revert-interval seconds;
            transmit-interval (10m | 1m | 10s | 1s);
            vlan-id vlan-id;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name protocols],
[edit protocols]
```

Description

Configure lightweight loop detection on server-facing logical interfaces in a Layer 2 (L2) network. This feature can detect the following types of Ethernet loops:

- A loop between two interfaces in different Ethernet segments (ESs). This loop is typically caused by miswiring fabric components.
- A loop between two interfaces with the same Ethernet segment identifier (ESI). This loop is typically caused by miswiring a third-party switch to the fabric.

You enable this feature on logical interface names; we don't support this feature on physical interfaces.

> **NOTE**: With EVPN-VXLAN, we recommend that you enable loop detection initially *before* you configure EVPN, so you can detect any loops and take corrective actions before EVPN traffic is flowing through the network. Otherwise, if you bring up an EVPN fabric with loops, even with this feature enabled, traffic can continue to flow through the loops for some time while the network stabilizes.
> *Supported Interface ConfigurationsEVPN-VXLAN Lightweight Leaf to Server Loop Detection*

After you’ve enabled loop detection, the interfaces periodically send multicast loop-detection protocol data units (PDUs). If a loop detection-enabled interface receives a PDU, it detects a loop and triggers the configured action to break the loop. For example, if you configured the interface-down action, the device brings the interface down. After you have repaired the loop condition and the `revert-interval` timer expires, the device reverts the loop-detection action and brings the interface up again.

Use the following command to manually revert (clear) the status on one or all logical interfaces where the device detected a loop:

```
clear loop-detect enhanced interface {interface-name}
```

If you don't specify an interface name, the `clear loop-detect enhanced` command reverts the interface status for all logical interfaces where the device performed a loop detection action. You can use the clear command in the following cases:

- When you didn't configure a `revert-interval` value. In that case, the interface loop detection status doesn't clear until you clear it manually.
- If you configured a `revert-interval`, but you want to revert the status before the interval expires.

Use the following command to see loop detection status information for an interface or all interfaces with loop detection enabled:

```
show loop-detect enhanced interface {interface-name}
```

Options

interface (*name* | all)

Specify the name of a server-facing L2 logical interface on which you want to enable lightweight loop detection, or use the `all` keyword to enable this feature on all such logical interfaces.

loop-detect-action

If the feature detects a loop, you can configure either of the following actions:

- Values:
  - `interface-down`—The device brings down the interface on which the loop was detected.
  - `laser-off`—The device turns off the optics transmit laser.

> **NOTE**: Platforms that support the laser-off port capability support this action.

revert-interval

(Optional) Specify the interval, in seconds, in which the interface will revert to its previous state after the loop condition is fixed. For example, if the `loop-detect-action` is `interface-down`, after the loop condition is repaired and the interval expires, the device brings the interface up again.

- **Default:** 0

> **NOTE**: If you retain this default setting (0) and a loop detection event triggers the `interface-down` action, you must enter the `clear loop-detect enhanced interface <interface-name>` command to bring the interface back up. Until you enter the clear command, the interface remains down.

- **Range:** 0 through 300 seconds

transmit-interval

(Optional) The interval in seconds or minutes at which the interface transmits loop detection PDUs.

- **Default:** 1 second (1s)
- Values:
  - 10 minutes (10m)
  - 1 minute (1m)
  - 10 seconds (10s)
  - 1 second (1s)

vlan-id ***vlan-id***

(Required for trunk interfaces, enterprise style interface configurations or service provider style configurations) Specify the VLAN identifier for the interface.

- **Range:** 1 through 4094

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

Statement support added to ACX-Series (ACX7100-32C, ACX7100-48L) devices in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

EVPN-VXLAN Lightweight Leaf to Server Loop Detection

clear loop-detect enhanced interface | **3693**

show loop-detect enhanced interface | **21987**

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)

IN THIS SECTION

- Syntax | 10117
- Hierarchy Level | 10117
- Description | 10118
- Default | 10118
- Required Privilege Level | 10118
- Release Information | 10118

Syntax

```
(loopback | no-loopback);
```

Hierarchy Level

```
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name ether-options],
[edit interfaces interface-name fastether-options],
[edit interfaces interface-name gigether-options],
[edit interfaces interface-range name ether-options]
```

For QFX Series and EX Series:

```
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name ether-options],
```

For SRX Series Firewalls and vSRX Virtual Firewall:

```
[edit interfaces interface-name redundant-ether-options]
```

Description

For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.

NOTE:

- By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system.
- IPv6 Neighbor Discovery Protocol (NDP) addresses are not supported on Gigabit Ethernet interfaces when loopback mode is enabled on the interface. That is, if the `loopback` statement is configured at the `[edit interfaces ge-`*`fpc/pic/port`*` gigether-options]` hierarchy level, an NDP address cannot be configured at the `[edit interfaces ge-`*`fpc/pic/port`*` unit `*`logical-unit-number`*` family inet6 address]` hierarchy level.

Default

By default, loopback is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Statement modified in Junos OS Release 9.2 for the SRX Series.

RELATED DOCUMENTATION

Configuring Ethernet Loopback Capability

Understanding Interfaces

loopback-clear-timer

IN THIS SECTION

- Syntax | 10119
- Hierarchy Level | 10119
- Description | 10119
- Options | 10119
- Required Privilege Level | 10120
- Release Information | 10120

Syntax

```
loopback-clear-timer seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number ppp-options],
[edit logical-systems logical-system-name interfaces interface-nameunit logical-unit-number ppp-
options]
```

Description

For interfaces with PPP, PPP TCC, PPP over Ethernet, PPP over ATM, and PPP over Frame Relay encapsulations, configure a loop detection clear timer for the Link Control Protocol (LCP) component of a PPP session.

Options

seconds—The time in seconds to wait before the loop detection flag is cleared if it is not cleared by the protocol.

- **Range:** 1 through 60 seconds

- **Default:** 9 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring PPP

loopback (Fibre Channel Interface)

IN THIS SECTION

- Syntax | **10120**
- Hierarchy Level | **10121**
- Description | **10121**
- Default | **10121**
- Required Privilege Level | **10121**
- Release Information | **10121**

Syntax

```
(loopback | no-loopback);
```

Hierarchy Level

```
[edit interfaces interface-name                fibrechannel-options]
```

Description

Enable or disable loopback mode for FC interfaces.

Default

By default, loopback mode is disabled on FC interfaces.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel interfaces | 19980

Converting an Ethernet Interface To a Fibre Channel Interface

loopback-firewall-optimization

IN THIS SECTION

- Syntax | 10122
- Hierarchy Level | 10122
- Description | 10122
- Required Privilege Level | 10123

Release Information | 10123

Syntax

```
loopback-firewall-optimization;
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable this setting to increase the system limit for the number of loopback filter terms that can be configured. When enabled, you can configure up to 768 loopback filter terms for IPv6, and up to 1152 terms for IPv4. The packet forwarding engine (PFE) will restart upon commit for the new system limits to take effect.

TTL (time-to-live) values such as 0/1 are not matched unless explicitly mentioned in the firewall filter term. Terms that include a reserved multicast destination such as 224.0.0.x/24 are not directly supported as match condition in filters used with the loopback address (lo0). Instead, to count packets destined to the reserved multicast address of 224.0. 0.6, you would need to create a filter that specifies **protocol OSPF** as the match term. An example showing such a configuration is provided below.

```
[edit interfaces]
lo0 {
    unit 0 {
        family <inet | inet6> {
            filter {
                input f1;
            }
             address 127.0.0.1/32;
        }
    }
}
[edit firewall]
    family <inet | inet6> {
```

```
    filter f1 {
        term t1 {
            from {
                protocol ospf;
            }
            then {
                count c1;
                accept;
            }
        }
    }
}
```

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Planning the Number of Firewall Filters to Create

loopback-firewallv6-optimization

IN THIS SECTION

- Syntax | 10124
- Hierarchy Level | 10124
- Description | 10124
- Required Privilege Level | 10124
- Release Information | 10124

Syntax

```
loopback-firewallv6-optimization
```

Hierarchy Level

```
[edit chassis]
```

Description

Move IPv6 loopback filters to a separate single wide group and IPv6 address match to VFP. Can be used to increase IPv6 loopback filter scale when scale of IPv4 loopback filters are used in EVPN VXLAN deployments.

With EVPN/VXLAN, we can scale either 512 terms for IPv4 loopback filter or 512 terms for IPv6 loopback filter using `set chassis loopback-firewall-optimization` alone. If `set chassis loopback-firewallv6-optimization` is also enabled then we can achieve 512 terms each for IPv4 and IPv6 loopback filters.

> **NOTE**: Only supported on QFX5210.
> When enabled, the CLI knob will lead to PFE restart.
>
> Packet-length match condition will not be supported for IPv6 loopback filter which has any source/destination IPv6 address match.

Required Privilege Level

interface

Release Information

Statement introduced in Junos 23.2

loopback (Local and Remote)

IN THIS SECTION

- Syntax | 10125
- Hierarchy Level | 10125
- Description | 10125
- Options | 10125
- Required Privilege Level | 10125
- Release Information | 10126

Syntax

```
loopback (local | remote);
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options]
```

Description

Enables local loopback and enables remote loopback. This allows you to test the transceiver cable connection from the far end to the retimer interface without changing the cable.

Options

local	Enables local loopback
remote	Enables remote loopback

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F3 and 16.1R2.

RELATED DOCUMENTATION

Configuring Ethernet Loopback Capability

loopback-port

IN THIS SECTION

- Syntax | 10126
- Hierarchy Level | 10126
- Description | 10127
- Required Privilege Level | 10127
- Release Information | 10128

Syntax

```
loopback-port loopback-port;
```

Hierarchy Level

```
[edit forwarding-options vxlan-routing ]
```

Description

Define the global loopback aggregated Ethernet interface for VXLAN routing on QFX5210 leaf devices in EVPN-VXLAN edge-routed bridging (ERB) overlay fabrics.

Specify *loopback-port* using its physical interface name, without a specific unit number (such as ae0 rather than ae0.100, for example).

> **NOTE**: The Packet Forwarding Engine restarts when you configure or update this setting.

QFX5210 switches don't have native VXLAN routing hardware support for Layer 3 VXLAN gateway functions. Instead, the switch uses an intermediary loopback port in a two-pass process for routing in and out of VXLAN tunnels (RIOT). You configure a loopback link aggregation group (LAG) and use this statement to assign it as the RIOT loopback LAG port on the device.

Keep the following in mind when you configure the RIOT loopback LAG:

- When you configure the loopback LAG members, you can include any network ports on the switch that you're not using for another purpose. Based on your VXLAN routing traffic load, you can adjust:
  - The number of links in the LAG.
  - The minimum number of links for the LAG to be considered "up".
- You must configure the loopback LAG interface as a member of all VLANs with IRB interfaces used for VXLAN routing.
- Configure the loopback LAG interface:
  - As a trunk interface
  - With `encapsulation flexible-ethernet-services` and `flexible-vlan-tagging`
- The device automatically turns off MAC learning on the interface when you set it as the loopback LAG port. In this way, upon receiving packets on the loopback port, the device avoids learning the port's own MAC address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.3R1 for QFX5210 switches.

RELATED DOCUMENTATION

Using a RIOT Loopback Port to Route Traffic in an EVPN Network

loopback-remote

IN THIS SECTION

- Syntax | 10128
- Hierarchy Level | 10128
- Description | 10129
- Required Privilege Level | 10129
- Release Information | 10129

Syntax

```
loopback-remote;
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name gigether-options],
[edit dynamic-profiles name logical-systems name interfaces name gigether-options],
[edit interfaces name gigether-options]
```

Description

Configure `loopback-remote` on the remote end of the link. When you configure on an interface `loopback-remote`, the local end goes down and remote end goes up.

> **NOTE**: The `loopback-remote` option is supported on MX routers, starting from the MPC7E, MPC8E, and MPC9E linecards.

Required Privilege Level

interface

Release Information

RELATED DOCUMENTATION

Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test

loopback (Serial)

IN THIS SECTION

- Syntax | **10130**
- Hierarchy Level | **10130**
- Description | **10130**
- Default | **10130**
- Options | **10130**
- Required Privilege Level | **10130**
- Release Information | **10130**

Syntax

```
loopback mode;
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

Configure a loopback connection.

Default

If you do not include this statement, there is no loopback connection.

Options

mode—You can specify the one of the following loopback modes:

- `dce-local`—For EIA-530 interfaces only, loop packets back on the local DCE.
- `dce-remote`—For EIA-530 interfaces only, loop packets back on the remote DCE.
- `local`—Loop packets back on the local router's PIC.
- `remote`—Loop packets back on the line interface unit (LIU).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Serial Loopback Capability

loopback-tracking

IN THIS SECTION

- Syntax | 10131
- Hierarchy Level | 10131
- Description | 10131
- Required Privilege Level | 10131
- Release Information | 10132

Syntax

```
loopback-tracking;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management]
```

Description

Enables loopback tracking on Ethernet interfaces. When loopback tracking is enabled and the Ethernet Operation, Administration, and Management (OAM) link-fault management process (lfmd) detects its own generated packets on an interface, it marks the interface as down. When the loopback issue resolves, the interface is brought back up.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

IEEE 802.3ah OAM Link-Fault Management Overview

Enabling IEEE 802.3ah OAM Support

loops (Autonomous System)

IN THIS SECTION

- Syntax | 10132
- Hierarchy Level | 10132
- Description | 10133
- Options | 10133
- Required Privilege Level | 10133
- Release Information | 10133

Syntax

```
loops number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp local-as],
[edit logical-systems logical-system-name protocols bgp group group-name local-as],
[edit logical-systems logical-system-nameprotocols bgp group group-name neighbor address local-
as],
```

```
[edit logical-systems logical-system-name routing-options autonomous-system as-number],
[edit protocols bgp local-as],
[edit protocols bgp group group-name local-as],
[edit protocols bgp group group-name neighbor address local-as],
[edit routing-options autonomous-system as-number]
```

Description

Globally, for the local-AS BGP attribute, allow the local device's AS number to be in the received AS paths, and specify the number of times detection of the local device's AS number in the AS_PATH attribute causes the route to hidden. For example, if you configure `loops 1`, the route is hidden if the local device's AS number is detected in the path one or more times. This prevents routing loops and is the default behavior. If you configure `loops 2`, the route is hidden if the local device's AS number is detected in the path two or more times.

NOTE: The behavior of this statement is slightly different from the "loops (BGP Address Family)" on page 10134 statement.

Options

number—Number of times detection of the AS number in the AS_PATH attribute causes the route to be hidden.

- **Range:** 1 through 10
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Enabling BGP Route Advertisements

autonomous-system | 2764
family (Protocols BGP) | 6900
local-as | 9916
loops (BGP Address Family) | 10134

loops (BGP Address Family)

IN THIS SECTION

- Syntax | 10134
- Hierarchy Level | 10134
- Description | 10135
- Options | 10136
- Required Privilege Level | 10136
- Release Information | 10136

Syntax

```
loops number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family address-family],
[edit logical-systems logical-system-name protocols bgp group group-name family address-
family],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
address-family],
[edit protocols bgp family address-family],
[edit protocols bgp group group-name family address-family],
[edit protocols bgp group group-name neighbor address family address-family]
```

Description

For the specified BGP address family, allow the local device's AS number in the received AS paths and specify the number of times the detection of the local device's AS in the AS_PATH attribute is allowed. If the count exceeds the specified loop count, the system discards this route. For example, if you configure `loops 1`, the route is discarded if the local device's AS number is detected in the path more than once. This prevents routing loops and is the default behavior. If you configure `loops 2`, the route is discarded if the local device's AS number is detected more than 2 times.

For debugging, you can configure the `keep all` option If you want to hide this route.

Some examples of BGP address families are as follows:

- `inet unicast`
- `inet-vpn multicast`
- `inet6 any`
- `l2vpn auto-discovery-only`
- ...

This list is truncated for brevity. For a complete list of protocol families for which you can specify the `loops` statement, enter the `help apropos loops` configuration command at the `[edit protocols bgp]` hierarchy level on your device.

```
[edit protocols bgp]
user@host# help apropos loops
set family inet unicast loops
    Allow local AS in received AS paths
set family inet unicast loops <loops>
    AS-Path loop count
set family inet multicast loops
    Allow local AS in received AS paths
set family inet multicast loops <loops>
    AS-Path loop count
set family inet flow loops
    Allow local AS in received AS paths
set family inet flow loops <loops>
    AS-Path loop count
set family inet any loops
    Allow local AS in received AS paths
set family inet any loops <loops>
    AS-Path loop count
```

```
set family inet labeled-unicast loops
    Allow local AS in received AS paths
...
```

NOTE: When you configure the `loops` statement for a specific BGP address family, that value is used to evaluate the AS path for routes received by a BGP peer for the specified address family, rather than the loops value configured for the global AS number with the `loops` statement at the `[edit routing-options autonomous-system as-number]` hierarchy level.

Options

number—Maximum number of times that the local device's AS number is allowed in the AS_PATH attribute to accept the route.

- **Range:** 1 through 10
- **Default:** None. The system does not take any action unless the `loops (BGP Address Family)` statement is configured.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Enabling BGP Route Advertisements

autonomous-system | **2764**

family (Protocols BGP) | **6900**

local-as | **9916**

loops (Autonomous System) | **10132**

loose-authentication-check

IN THIS SECTION

- Syntax | 10137
- Hierarchy Level | 10137
- Description | 10137
- Required Privilege Level | 10137
- Release Information | 10138

Syntax

```
loose-authentication-check;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Allow the use of MD5 authentication without requiring network-wide deployment.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding Hitless Authentication Key Rollover for IS-IS

Example: Configuring Hitless Authentication Key Rollover for IS-IS

loose-check

IN THIS SECTION

- Syntax | 10138
- Hierarchy Level | 10138
- Description | 10139
- Required Privilege Level | 10139
- Release Information | 10139

Syntax

```
loose-check;
```

Hierarchy Level

```
[edit protocols pim interface interface-name bfd-liveness-detection authentication],
[edit routing-instances routing-instance-name protocols pim interface interface-name bfd-
liveness-detection authentication]
```

Description

Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD Authentication for PIM

Understanding Bidirectional Forwarding Detection Authentication for PIM

authentication (Protocols PIM) | **2647**

loose-source-route-option (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | **10140**
- Hierarchy Level | **10140**
- Description | **10140**
- Required Privilege Level | **10140**
- Release Information | **10140**

Syntax

```
loose-source-route-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop IPv4 packets that have the IP option of 3 (Loose Source Routing).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

loss-delay (querier)

IN THIS SECTION

- Syntax | 10141
- Hierarchy Level | 10141
- Description | 10141

- Required Privilege Level | **10141**
- Release Information | **10142**

Syntax

```
loss-delay {
    traffic-class tc-value {
        average-sample-size sample size;
        loss-threshold loss threshold value;
        loss-threshold-window number of samples for loss threshold;
        measurement-quantity bytes|packets;
        padding-size size;
        query-interval milliseconds;
        rtt-delay-threshold rtt threshold value;
        twcd-delay-threshold twcd threshold value;
    }
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring
querier],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring
querier]
```

Description

Configure combined loss-delay measurement options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Pro-Active Loss and Delay Measurements

On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview

performance-monitoring (Protocols MPLS) | 12434

loss-priority

IN THIS SECTION

- Syntax | 10142
- Hierarchy Level | 10142
- Description | 10143
- Options | 10143
- Required Privilege Level | 10143
- Release Information | 10143

Syntax

```
loss-priority (high | low);
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile
output-priority-map classifier premium forwarding-class class-name]
```

Description

Specify the packet loss priority value.

Options

`high`—Packet has high loss priority.

`low`—Packet has low loss priority.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

loss-priority (Classifiers)

IN THIS SECTION

- Syntax | 10144
- Hierarchy Level | 10144
- Description | 10144
- Options | 10144
- Required Privilege Level | 10144
- Release Information | 10144

Syntax

```
loss-priority level;
```

Hierarchy Level

```
[edit class-of-service classifiers type classifier-name forwarding-class class-name]
```

Description

Specify packet loss priority value for a specific set of code-point aliases and bit patterns.

Options

level can be one of the following:

- `high`—Packet has high loss priority.
- `medium-high`—Packet has medium-high loss priority.
- `medium-low`—Packet has medium-low loss priority.
- `low`—Packet has low loss priority.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic*
>
> *Configuring and Applying Tricolor Marking Policers*

loss-priority (Classifiers)

IN THIS SECTION

- Syntax | 10145
- Hierarchy Level | 10145
- Description | 10145
- Options | 10145
- Required Privilege Level | 10146
- Release Information | 10146

Syntax

```
loss-priority level {
    code-points [aliases] [bit-patterns];
}
```

Hierarchy Level

```
[edit class-of-service classifiers classifier-type classifier-name forwarding-class class-name]
```

Description

Configure packet loss priority value for a specific set of code-point aliases and bit patterns.

Options

level—Can be one of the following:

- `low`—Packet has low loss priority.
- `medium-low`—Packet has medium-low loss priority.
- `medium-high`—Packet has medium-high loss priority.

- `high`—Packet has high loss priority.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Unicast Classifiers

Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

Understanding CoS Classifiers

loss-priority (Classifiers and Rewrite Rules)

IN THIS SECTION

- Syntax | 10147
- Hierarchy Level | 10147
- Description | 10147
- Options | 10147
- Required Privilege Level | 10147
- Release Information | 10148

Syntax

```
loss-priority level {
    code-points [aliases] [6-bit-patterns / 3-bit-patterns];
}
```

Hierarchy Level

```
[edit class-of-service classifiers (dscp | ieee-802.1 | inet-precedence | exp) classifier-name
forwarding-class class-name],
[edit class-of-service rewrite-rules (dscp | ieee-802.1 | inet-precedence | exp) rewrite-rule-name
forwarding-class class-name]
```

Description

Specify packet loss priority value for a specific set of code-point aliases and bit patterns.

Options

level—Can be one of the following:

- `high`—Packet has high loss priority.
- `medium-high`— Code points to classify to loss priority medium-high.
- `low`—Packet has low loss priority.
- `medium-low` — Code points to classify to loss priority medium-low.

> **NOTE**: The EX4300 hardware supports only three levels of loss priority — `high`, `medium-high`, and `low`. You can configure a loss priority of `medium-low` on the EX4300, but the hardware will convert it to `medium-high`.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement expanded to apply to EXP classifiers in Junos OS Release 10.1 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches

Defining CoS Classifiers (CLI Procedure)

Defining CoS Classifiers (J-Web Procedure)

Defining CoS Rewrite Rules (CLI Procedure)

Defining CoS Rewrite Rules (J-Web Procedure)

Understanding CoS Classifiers

Understanding CoS Rewrite Rules

loss-priority (Frame Relay Loss Priority)

IN THIS SECTION

- Syntax | 10148
- Hierarchy Level | 10149
- Description | 10149
- Options | 10149
- Required Privilege Level | 10149
- Release Information | 10149

Syntax

```
loss-priority level code-points [values ];
```

Hierarchy Level

```
[edit class-of-service loss-priority-maps frame-relay-de map-name]
```

Description

Map CoS values to a packet loss priority (PLP). In Junos OS, classifiers associate incoming packets with a forwarding class (FC) and PLP. PLPs allow you to set the priority for dropping packets. Typically, you mark packets exceeding some service level with a high loss priority—that is, a greater likelihood of being dropped.

Options

level can be one of the following:

- `high`—Packet has high loss priority.
- `medium-high`—Packet has medium-high loss priority.
- `medium-low`—Packet has medium-low loss priority.
- `low`—Packet has low loss priority.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

> Understanding Interfaces
>
> *Understanding Packet Loss Priorities*

loss-priority (Drop Profiles)

IN THIS SECTION

- Syntax | 10150
- Hierarchy Level | 10150
- Description | 10150
- Options | 10150
- Required Privilege Level | 10151
- Release Information | 10151

Syntax

```
loss-priority level protocol protocol drop-profile profile-name;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name drop-profile-map]
```

Description

Configure packet loss priority value for a weighted random early detection (WRED) drop profile mapped to a system drop profile.

Options

level—Can be one of the following:

- `low`—Packet has low loss priority.
- `medium-high`—Packet has medium-high loss priority.
- `high`—Packet has high loss priority.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

loss-priority (Dynamic Schedulers)

IN THIS SECTION

- Syntax | 10151
- Hierarchy Level | 10151
- Description | 10152
- Options | 10152
- Required Privilege Level | 10152
- Release Information | 10152

Syntax

```
loss-priority (any | low | medium-low | medium-high | high);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name drop-profile-map]
```

Description

Specify a loss priority to which to apply a drop profile in a dynamic profile. The drop profile map sets the drop profile for a specific PLP and protocol type. The inputs for the map are the PLP designation and the protocol type. The output is the drop profile.

Options

`any`—The drop profile applies to packets with any PLP.

`high`—The drop profile applies to packets with high PLP.

`medium-high`—The drop profile applies to packets with medium-high PLP.

`medium-low`—The drop profile applies to packets with medium-low PLP.

`low`—The drop profile applies to packets with low PLP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Schedulers in a Dynamic Profile for Subscriber Access

loss-priority (Firewall Filter)

IN THIS SECTION

Syntax | 10153

- Hierarchy Level | 10153
- Description | 10153
- Options | 10153
- Required Privilege Level | 10154
- Release Information | 10154

Syntax

```
loss-priority (high | low | medium-high | medium-low);
```

Hierarchy Level

```
[edit dynamic-profiles name firewall family inet filter name term name from],
[edit dynamic-profiles name firewall filter name term name from],
[edit firewall family inet filter name term name from],
[edit firewall filter name term name from],
[edit logical-systems name firewall family inet filter name term name from],
[edit logical-systems name firewall filtername term name from]
```

Description

Set the loss priority of incoming packets, which governs the likelihood of the system dropping packets in the event of congestion. For example, to ensure delivery of critical traffic, you might want to set the loss priority of non-critical flows to high or medium-high to intentionally sacrifice those packets in favor of the preferred traffic whenever there is contention of resources.

Options

high	Highest probability of being dropped at times of congestion
medium-high	Second highest probability of being dropped at times of congestion
medium-low	Third highest probability of being dropped at times of congestion

low Lowest probability of being dropped at times of congestion

Required Privilege Level

firewall

Release Information

Statement introduced before Junos OS Release 7.4.

For EX9200 series switches, support for `medium-high` and `medium-low` loss priorities was added in Junos OS and Junos OS Evolved releases 19.x and 20.x for the `ethernet-switching` firewall family.

RELATED DOCUMENTATION

Configuring Multifield Classifiers

loss-priority (Firewall Filter Action)

IN THIS SECTION

- Syntax | 10154
- Hierarchy Level | 10155
- Description | 10155
- Required Privilege Level | 10155
- Release Information | 10155

Syntax

```
loss-priority (high | low);
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name then],
[edit logical-systems logical-system-name firewall family family-name filter filter-name term term-name then]
```

Description

Set the loss priority of incoming packets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Firewall Filter Nonterminating Actions*
- *Policer Color-Marking and Actions*
- *Multifield Classification Overview*

loss-priority high then discard (Three-Color Policer)

IN THIS SECTION

- Syntax | 10156
- Hierarchy Level | 10156
- Description | 10156
- Required Privilege Level | 10156

Release Information | 10156

Syntax

```
loss-priority high then discard;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name action],
[edit firewall three-color-policer policer-name action],
[edit logical-systems logical-system-name firewall three-color-policer policer-name action]
```

Description

For packets with high loss priority, discard the packets. The loss priority setting is implicit and is not configurable. Include this statement if you do not want the local router to forward packets that have high packet loss priority.

For single-rate three-color policers, the Junos OS assigns high loss priority to packets that exceed the committed information rate and the excess burst size.

For two-rate three-color policers, the Junos OS assigns high loss priority to packets that exceed the peak information rate and the peak burst size.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 8.2.

Logical systems support introduced in Junos OS Release 9.3.

Support at the **[edit dynamic-profiles ... action]** hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Three-Color Policer Configuration Overview
- Basic Single-Rate Three-Color Policers
- Basic Two-Rate Three-Color Policers
- Two-Color and Three-Color Logical Interface Policers
- Two-Color and Three-Color Physical Interface Policers
- Two-Color and Three-Color Policers at Layer 2
- action | 1759

loss-priority-maps

IN THIS SECTION

- Syntax | 10157
- Hierarchy Level | 10158
- Description | 10158
- Options | 10158
- Required Privilege Level | 10158
- Release Information | 10158

Syntax

```
loss-priority-maps {
    frame-relay-de rewrite-name {
        loss-priority level {
        code-points [ aliases] [ bit-patterns ];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Map the loss priority of incoming packets based on the CoS values.

Options

`frame-relay-de` *`rewrite-name`*—Name of the Frame Relay DE bit loss priority map.

`loss-priority` *`level`*—The loss priority level can be one of the following:

- `high`—Packet has high loss priority.
- `low`—Packet has low loss priority.
- `medium-high`—Packet has medium-high loss priority.
- `medium-low`—Packet has medium-low loss priority.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 11.4.

RELATED DOCUMENTATION

> *Assigning the Default Frame Relay Discard Eligibility Loss Priority Map to an Interface*

loss-priority-maps (Assigning to an Interface)

IN THIS SECTION

- Syntax | 10159
- Hierarchy Level | 10159
- Description | 10159
- Options | 10159
- Required Privilege Level | 10160
- Release Information | 10160

Syntax

```
loss-priority-maps {
    frame-relay-de (loss-priority-rewrite-name | default);
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Assign the loss priority map to a logical interface.

Options

`default`—Apply the default loss priority map. The default map includes the following configuration:

```
loss-priority low code-point 0;
loss-priority high code-point 1;
```

`map-name`—Name of loss priority map to be applied.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 11.4.

RELATED DOCUMENTATION

> *Assigning the Default Frame Relay Discard Eligibility Loss Priority Map to an Interface*
>
> unit (CoS Interfaces) | **30382**

loss-priority-maps (CoS)

IN THIS SECTION

- Syntax | **10160**
- Hierarchy Level | **10161**
- Description | **10161**
- Required Privilege Level | **10161**
- Release Information | **10161**

Syntax

```
loss-priority-maps {
    frame-relay-de loss-priority-map-name {
        loss-priority (high  | low  | medium-high  | medium-low) {
            code-points [bit-string];
        }
```

```
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Map the loss priority of incoming packets based on CoS values.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Interfaces

loss-priority-maps (CoS Interfaces)

IN THIS SECTION

- Syntax | 10162
- Hierarchy Level | 10162
- Description | 10162
- Options | 10162
- Required Privilege Level | 10162

Release Information | 10162

Syntax

```
loss-priority-maps {
    frame-relay-de (map-name | default);
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Assign the loss priority map to a logical interface.

Options

- `default`—Apply default loss priority map. The default map contains the following:

  ```
  loss-priority low code-point 0;
  loss-priority high code-point 1;
  ```

- *`map-name`*—Name of loss priority map to be applied.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Interfaces

loss-priority (Rewrite Rules)

IN THIS SECTION

- Syntax | 10163
- Hierarchy Level | 10163
- Description | 10163
- Options | 10163
- Required Privilege Level | 10164
- Release Information | 10164

Syntax

```
loss-priority level;
```

Hierarchy Level

```
[edit class-of-service rewrite-rules type rewrite-name forwarding-class class-name]
```

Description

Specify a loss priority to which to apply a rewrite rule. The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and packet loss priority (PLP). The inputs for the map are the forwarding class and the PLP. The output of the map is the code-point alias or bit pattern.

Options

level can be one of the following:

- `high`—The rewrite rule applies to packets with high loss priority.
- `low`—The rewrite rule applies to packets with low loss priority.
- `medium-high`—The rewrite rule applies to packets with medium-high loss priority.
- `medium-low`—The rewrite rule applies to packets with medium-low loss priority.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Class of Service User Guide (Security Devices)

loss-priority (Rewrite Rules)

IN THIS SECTION

- Syntax | 10165
- Hierarchy Level | 10165
- Description | 10165
- Options | 10165
- Required Privilege Level | 10165
- Release Information | 10165

Syntax

```
loss-priority level {
    code-point (alias | bit-pattern);
}
```

Hierarchy Level

```
[edit class-of-service rewrite-rules type rewrite-name forwarding-class class-name]
```

Description

Specify a loss priority to which to apply a rewrite rule. The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and loss priority. Packets that match the forwarding class and loss priority are rewritten with the rewrite code-point alias or bit pattern.

Options

level can be one of the following:

- `high`—The rewrite rule applies to packets with high loss priority.
- `low`—The rewrite rule applies to packets with low loss priority.
- `medium-high`—The rewrite rule applies to packets with medium-high loss priority.
- `medium-low`—The rewrite rule applies to packets with medium-low loss priority.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Rewrite Rules

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic

Understanding CoS Rewrite Rules

Defining CoS Rewrite Rules

loss-priority-rewrites

IN THIS SECTION

- Syntax | 10166
- Hierarchy Level | 10166
- Description | 10167
- Options | 10167
- Required Privilege Level | 10167
- Release Information | 10167

Syntax

```
loss-priority-rewrites {
    frame-relay-de rewrite-name {
        loss-priority level {
        code-points [ bit-patterns ];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify the Frame Relay discard eligibility (DE) bit rewrite rule on the enhanced IQ PIC.

Options

`frame-relay-de` *`rewrite-name`*—Name of the Frame Relay DE bit loss priority rewrite rule.

`loss-priority` *`level`*—The loss priority level can be one of the following:

- `high`—Packet has high loss priority.
- `low`—Packet has low loss priority.
- `medium-high`—Packet has medium-high loss priority.
- `medium-low`—Packet has medium-low loss priority.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

loss-priority-rewrites (Assigning to an Interface)

IN THIS SECTION

- Syntax | 10168
- Hierarchy Level | 10168
- Description | 10168
- Options | 10168
- Required Privilege Level | 10168

Release Information | 10168

Syntax

```
loss-priority-rewrites {
    frame-relay-de (loss-priority-rewrite-name | default);
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Associate the loss priority rewrites to an outgoing packet.

Options

loss-priority-rewrite-name—Name of the loss priority rewrite to be applied to an interface.

default—Default loss priority rewrite to be applied to an interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

loss-priority (Scheduler Drop Profiles)

IN THIS SECTION

- Syntax | 10169
- Hierarchy Level | 10169
- Description | 10169
- Options | 10169
- Required Privilege Level | 10170
- Release Information | 10170

Syntax

```
loss-priority (any | high | low | medium-high | medium-low);
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name drop-profile-map]
```

Description

Specify a loss priority to which to apply a drop profile. The drop profile map sets the drop profile for a specific PLP and protocol type. The inputs for the map are the PLP designation and the protocol type. The output is the drop profile.

Options

`any`—The drop profile applies to packets with any PLP.

`high`—The drop profile applies to packets with high PLP.

`low`—The drop profile applies to packets with low PLP.

`medium-high`—The drop profile applies to packets with medium-high PLP.

`medium-low`—The drop profile applies to packets with medium-low PLP.

> **NOTE**: On ACX Series Routers, if you configure the "protocol" on page 13459 as `tcp`, then the loss-priority (`any | high | low | medium-high`) values are supported. If you configure the "protocol" on page 13459 with either `non-tcp` or `any` option, then irrespective of the loss-priority value, only one drop-profile can be specified.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Default Schedulers Overview*
- *Determining Packet Drop Behavior by Configuring Drop Profile Maps for Schedulers*
- *Configuring Schedulers for Priority Scheduling*
- *Configuring and Applying Tricolor Marking Policers*
- protocol (Schedulers) | **13459**

loss-priority (Simple Firewall Filter)

IN THIS SECTION

- Syntax | **10171**
- Hierarchy Level | **10171**
- Description | **10171**
- Required Privilege Level | **10171**
- Release Information | **10171**

Syntax

```
loss-priority (high | low | medium);
```

Hierarchy Level

```
[edit firewall family family-name simple-filter filter-name term term-name then]
```

Description

Set the loss priority of incoming packets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Multifield Classifiers

loss (querier)

IN THIS SECTION

- Syntax | 10172
- Hierarchy Level | 10172
- Description | 10172

- Required Privilege Level | **10172**
- Release Information | **10173**

Syntax

```
loss {
    traffic-class tc-value {
        average-sample-size sample size;
        loss-threshold loss threshold value;
        loss-threshold-window number of samples for loss threshold;
        measurement-quantity bytes|packets;
        query-interval milliseconds;
    }
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring
querier],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring
querier]
```

Description

Configure loss measurement options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Pro-Active Loss and Delay Measurements

On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview

performance-monitoring (Protocols MPLS) | **12434**

loss (responder)

IN THIS SECTION

- Syntax | **10173**
- Hierarchy Level | **10173**
- Description | **10174**
- Options | **10174**
- Required Privilege Level | **10174**
- Release Information | **10174**

Syntax

```
loss {
    min-query-interval milliseconds;
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring responder],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring responder],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring
```

```
responder],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring
responder]
```

Description

Configure loss measurement options.

Options

min-query-interval *milliseconds*

(Optional) Specify the minimum query interval that the responder supports. If the minimum query interval of the responder is greater than the query interval configured at the querier, the effective message query rate is the minimum query interval configured for the responder.

- **Default:** 10 seconds
- **Range:** 1000 through 4294967295 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Pro-Active Loss and Delay Measurements

On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview

performance-monitoring (Protocols MPLS) | **12434**

loss-threshold

IN THIS SECTION

- Syntax | 10175
- Hierarchy Level | 10175
- Description | 10175
- Options | 10175
- Required Privilege Level | 10175
- Release Information | 10176

Syntax

```
loss-threshold number;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name continuity-check]
```

Description

Specify the number of continuity check messages lost before marking the remote MEP as down. The value can be from 3 to 256 protocol data units (PDUs). The default value is 3 PDUs.

Options

number—The number of continuity check messages that can be lost before the remote MEP is considered down.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Continuity Check Protocol Parameters Overview

Configuring Continuity Check Protocol Parameters for Fault Detection

low-latency

IN THIS SECTION

- Syntax | 10176
- Hierarchy Level | 10176
- Description | 10177
- Required Privilege Level | 10177
- Release Information | 10177

Syntax

```
low-latency
```

Hierarchy Level

```
[edit chassis fpc fpc-slot-number pic pic-slot-number services-offload]
```

Description

Enables the low-latency mode on the selected NP-IOC. Low-latency is not enabled by default. The low latency allows you to configure the mode of the network processor's traffic manager (TM) on the egress path. If low latency is enabled, the network processor is initialized without the traffic manager, thus reducing the overall latency in the Express Path (formerly known as services offloading).

> **NOTE**: Because all SRX Series CoS functions are supported by the traffic manager, CoS functions are not supported when low latency is enabled.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

low-watermark

IN THIS SECTION

- Syntax | **10178**
- Hierarchy Level | **10178**
- Description | **10178**
- Options | **10178**
- Required Privilege Level | **10178**
- Release Information | **10178**

Syntax

```
low-watermark percent;
```

Hierarchy Level

```
[edit security flow aging]
```

Description

Set the point at which the aggressive aging-out process ends.

Options

percent —Percentage of session-table capacity at which aggressive aging-out ends.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Understanding Traffic Processing on Security Devices*

lowest-priority-defect

IN THIS SECTION

- Syntax | 10179
- Hierarchy Level | 10179
- Description | 10179
- Options | 10179
- Required Privilege Level | 10180
- Release Information | 10180

Syntax

```
lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect | rem-err-xcon |
xcon)
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id]
```

Description

Specify the lowest priority defect that is allowed to generate a Fault Alarm whenever CFM detects a defect. This configuration is done at the MEP level.

Options

Specify one of the following lowest priority defect options:

all-defects Allows all defects.

err-xcon Allows only erroneous CCM and cross-connect CCM defects.

`mac-rem-err-xcon`	Allows only MAC, not receiving CCM, erroneous CCM, and cross-connect defects.
`no-defect`	Allows no defects.
`rem-err-xcon`	Allows only not receiving CCM, erroneous CCM, and cross-connect CCM defects.
`xcon`	Allows only cross-connect CCM defects.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring a MEP to Generate and Respond to CFM Protocol Messages

lrf-profile (Service Set)

IN THIS SECTION

- Syntax | **10181**
- Hierarchy Level | **10181**
- Description | **10181**
- Options | **10181**
- Required Privilege Level | **10181**
- Release Information | **10181**

Syntax

```
lrf-profile profile-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Assign the LRF profile to the service set that is that is configured for application-aware policy control.

Options

profile-name LRF profile name. The referenced LRF profile must be configured.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Assigning an LRF Profile to Subscribers

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Junos OS Subscriber Aware

Applying Logging and Reporting Configuration to a Subscriber Management Service Set

Configuring Logging and Reporting for Subscriber Management

lsa-refresh-interval

IN THIS SECTION

- Syntax | 10182
- Hierarchy Level | 10182
- Description | 10182
- Options | 10183
- Required Privilege Level | 10183
- Release Information | 10183

Syntax

```
lsa-refresh-interval minutes;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Configure the refresh interval for all self-generated link-state advertisement (LSAs). The OSPF standard requires that every LSA be refreshed every 30 minutes. The Juniper Networks implementation refreshes

LSAs every 50 minutes. By default, any LSA that is not refreshed expires after 60 minutes. By using this configuration, you can specify when self-originated LSAs are refreshed.

You can override the default behavior by globally configuring the OSPF LSA refresh interval at the `[edit protocols ospf | ospf3]` hierarchy level. However, if you also have OSPF flood reduction configured for a specific interface in an OSPF area at the `[edit protocols ospf | ospf3 area` *`area-id`* `interface` *`interface-name`*`]` hierarchy level, the flood reduction configuration takes precedence for that specific interface.

Options

minutes—Time between an LSA refresh, in minutes.

- **Range:** 25 through 50 minutes (1,500 through 3,000 seconds)
- **Default:** 50 minutes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring OSPF Refresh and Flooding Reduction in Stable Topologies

lsp-attributes

IN THIS SECTION

- Syntax | **10184**
- Hierarchy Level | **10184**
- Description | **10184**

Required Privilege Level | 10184
Release Information | 10184

Syntax

```
lsp-attributes {
    encoding-type (ethernet | packet | pdh | sonet-sdh);
    gpid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 | pos-
scrambling-crc-32 | pos-no-scrambling-crc-32 | ppp);
    signal-bandwidth type;
    switching-type (fiber | lambda | psc-1 | tdm);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Define the parameters signaled during LSP setup. These usually determine the nature of the resource (label) allocated for the LSP.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`pos-scrambling-crc-16`, `pos-no-scrambling-crc-16`, `pos-scrambling-crc-32`, and `pos-no-scrambling-crc-32` options added in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring MPLS LSPs for GMPLS

lsp-cleanup-timer

IN THIS SECTION

- Syntax | 10185
- Hierarchy Level | 10185
- Description | 10186
- Options | 10186
- Required Privilege Level | 10186
- Release Information | 10186

Syntax

```
lsp-cleanup-timer seconds;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
[edit protocols pcep pce-group group-id]
```

Description

Specify the amount of time (in seconds) that the Path Computation Client (PCC) must wait before deleting any non-delegated Path Computation Element (PCE)-initiated LSPs from the failed PCE after a PCEP session terminates.

> **NOTE**: In compliance with *draft-ietf-pce-stateful-pce-09*, revoking of PCE-initiated LSP delegations by a PCC happens in a make-before-break fashion before the LSPs are redelegated to an alternate PCE. Starting in Junos OS Release 18.1R1, the `lsp-cleanup-timer` must be greater than or equal to the `delegation-cleanup-timeout` for the PCC to revoke the LSP delegations. If not, the redelegation timeout interval for the PCC can be set to infinity, where the LSP delegations to that PCE remain intact until specific action is taken by the PCC to change the parameters set by the PCE.

Options

seconds Time (in seconds) that the PCC must wait before deleting any non-delegated PCE-initiated LSPs from the failed PCE after a PCEP session terminates. Non-delegated PCE-initiated LSPs are deleted immediately.

- **Range:** 0 through 2147483647 seconds
- **Default:** 60

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

pcep | 12320

lsp-external-controller

IN THIS SECTION

- Syntax | 10187
- Hierarchy Level | 10187
- Description | 10187
- Options | 10187
- Required Privilege Level | 10188
- Release Information | 10188

Syntax

```
lsp-external-controller controller-name{
 (delay-metric-bound | igp-metric-bound | te-metric-bound);
}
```

Hierarchy Level

```
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name]
[edit protocols spring-traffic-engineering]
```

Description

Enable external path computing capability for the device.

Options

controller-name Name of the external path computing entity. By default, *pccd* is the only allowed LSP external controller.

- **Values:** *pccd*

delay-metric-bound Indicates the maximum metric value, a path can have, for the metric-type delay variation.

- **Values:** *1 through 4294967295*

igp-metric-bound Indicates the maximum metric value, a path can have, for the metric-type IGP.

- **Values:** *1 through 4294967295*

te-metric-bound Indicates the maximum metric value, a path can have, for the metric-type TE.

- **Values:** *1 through 4294967295*

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for ACX Series added in Junos OS Release 17.1.

Support at the `[edit protocols spring-traffic-engineering]` hierarchy level introduced in Junos OS Release 17.2.

Support at the following hierarchy levels introduced in Junos OS Release 20.1R1: `[edit protocols source-packet-routing source-routing-path` *`name`* `, [edit protocols source-packet-routing source-routing-path` *`name`* `primary` *`name`*, and `[edit protocols source-packet-routing source-routing-path-template` *`name`*.

RELATED DOCUMENTATION

pcep | **12320**

PCEP Configuration

lsp-lifetime

IN THIS SECTION

- Syntax | 10189
- Hierarchy Level | 10189
- Description | 10189
- Default | 10190
- Options | 10190
- Required Privilege Level | 10190
- Release Information | 10190

Syntax

```
lsp-lifetime seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Specify how long a link-state PDU originating from the routing device should persist in the network. The routing device sends link-state PDUs often enough so that the link-state PDU lifetime never expires.

Because link-state PDUs have a maximum lifetime, they need to be refreshed. Refreshing means that a routing device needs to re-originate its link-state PDUs periodically. The re-origination interval must be less than the link-state PDU's lifetime. For example, if the link-state PDU is valid for 1200 seconds, the routing device needs to refresh the link-state PDU in less than 1200 seconds to avoid removal of the link-state PDU from the link-state database by other routing devices. The recommended maximum link-

state PDU origination interval is the lifetime minus 300 seconds. So, in a default environment this would be 900 seconds. In Junos OS, the refresh interval is derived from the lifetime and is equal to the lifetime minus 317 seconds. You can change the lifetime to a higher value to reduce the number of refreshes in the network. (You would rarely want to increase the number of refreshes.) Often theses periodic link-state PDU refreshes are referred to as refresh noise, and network administrators want to reduce this noise as much as possible.

The `show isis overview` command displays the link-state PDU lifetime.

Default

By default, link-state PDUs are maintained in network databases for 1200 seconds (20 minutes) before being considered invalid. This length of time, called the *LSP lifetime*, normally is sufficient to guarantee that link-state PDUs never expire.

Options

seconds—link-state PDU lifetime, in seconds.

- **Range:** 350 through 65,535 seconds
- **Default:** 1200 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces

Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces

Understanding Link-State PDU Throttling for IS-IS Interfaces

Understanding the Transmission Frequency for CSNPs on IS-IS Interfaces

https://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

lsp-next-hop (CoS-Based Forwarding)

IN THIS SECTION

- Syntax | 10191
- Hierarchy Level | 10191
- Description | 10191
- Options | 10191
- Required Privilege Level | 10191
- Release Information | 10192

Syntax

```
lsp-next-hop [ lsp-regular-expression ];
```

Hierarchy Level

```
[edit class-of-service forwarding-policy next-hop-map map-name forwarding-class class-name]
```

Description

Specify the LSP regular expression to which to map forwarded traffic.

Options

lsp-regular-expression—Next-hop LSP label.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring CoS-Based Forwarding

lsp-next-hop (Static Routes)

IN THIS SECTION

- Syntax | 10192
- Hierarchy Level | 10192
- Description | 10193
- Options | 10193
- Required Privilege Level | 10193
- Release Information | 10193

Syntax

```
lsp-next-hop lsp-name {
    metric metric;
    preference preference;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options
static route destination-prefix],
[edit logical-systems logical-system-name routing-options static route destination-prefix],
```

```
[edit routing-instances routing-instance-name routing-options static route destination-prefix]
[edit routing-options static route destination-prefix]
```

Description

Specify an LSP as the next hop for a static route, and configure an independent metric or preference on that next-hop LSP.

> **NOTE**: The preference and metric configured by means of the `lsp-next-hop` statement only apply to the LSP next hops. The LSP next-hop preference and metric override the route preference and metric (for that specific LSP next hop), similar to how the route preference overrides the default preference and metric (for that specific route).

Options

lsp-name—Name of the next-hop LSP.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP

lsp-switch

IN THIS SECTION

- Syntax | 10194
- Hierarchy Level | 10194
- Description | 10194
- Options | 10194
- Required Privilege Level | 10195
- Release Information | 10195

Syntax

```
lsp-switch connection-name {
    transmit-lsp label-switched-path;
    receive-lsp label-switched-path;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols connections],
[edit protocols connections]
```

Description

Configure Layer 2 switching cross-connects.

Options

connection-name—Connection name.

receive-lsp *label-switched-path*—Name of the LSP from the connection's source.

transmit-lsp *label-switched-path*—Name of the LSP to the connection's destination.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Connection for Layer 2 Switching TCCs

lsping-channel-type

IN THIS SECTION

- Syntax | 10195
- Hierarchy Level | 10196
- Description | 10196
- Options | 10196
- Required Privilege Level | 10196
- Release Information | 10196

Syntax

```
lsping-channel-type {
    ipv4;
    on-demand-cv;
}
```

Hierarchy Level

```
[edit protocols mpls label-switched-path  lsp-name oam mpls-tp-mode],
[edit protocols mpls oam mpls-tp-mode]
```

Description

Specify the control-channel types for MPLS-TP mode. By default, LSPING (0x0008) is used, and the GACH-TLV is used along with this channel type.

As per RFC 7026, GACH-TLV is deprecated for ipv4 and on-demand-cv channel types.

Options

ipv4 Channel type 0x0021. This channel type uses the IP/UDP encapsulation and provides interoperability support with other vendor devices using IP addressing.

on-demand-cv Channel type 0x0025. This is a new pseudowire channel type and is used for on-demand CV without IP/UDP encapsulation, where IP addressing is not available or non-IP encapsulation is preferred.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

mpls-tp-mode | 10988

lsq-failure-options

IN THIS SECTION

- Syntax | 10197
- Hierarchy Level | 10197
- Description | 10197
- Options | 10197
- Required Privilege Level | 10197
- Release Information | 10198

Syntax

```
lsq-failure-options {
    no-termination-request;
    trigger-link-failure interface-name;
}
```

Hierarchy Level

```
[edit interfaces lsq-fpc/pic/port]
```

Description

For link services IQ (`lsq`) interfaces only, define the failure recovery option settings.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Association between LSQ and SONET Interfaces

ltv-syslog-interval (System Process)

IN THIS SECTION

- Syntax | 10198
- Hierarchy Level | 10198
- Description | 10199
- Options | 10199
- Required Privilege Level | 10199
- Release Information | 10199

Syntax

```
ltv-syslog-interval seconds;
```

Hierarchy Level

```
[edit system processes dhcp-service]
```

Description

Configure how often the router logs consolidated syslog messages for DHCP lease-time violations.

Options

seconds Time interval that specifies how often the router logs syslog messages.

- **Range:** 600 through 86,400 seconds

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring a DHCP Lease-Time Threshold

processes

M

IN THIS SECTION

- mac | **10215**
- mac (MPLS) | **10217**
- mac-address | **10218**
- mac-address (Accept Source Mac) | **10219**
- mac-address (DHCP Local Server) | **10221**

- mac-address (DHCP Relay Agent) | **10223**
- mac-address (Dynamic Access-Internal Routes) | **10224**
- mac-address (MACsec) | **10226**
- mac-address (VLAN and Stacked VLAN Interfaces) | **10227**
- mac-flush | **10229**
- mac-ip-table-size | **10231**
- mac-learn-enable | **10233**
- mac-learning-priority | **10234**
- mac-limit | **10236**
- mac-limit (VLAN) | **10238**
- mac-limit (Access Port Security) | **10241**
- mac-list | **10243**
- mac-lookup-length | **10244**
- mac-move-limit | **10246**
- mac-notification | **10249**
- mac (Option 82) | **10250**
- mac-persistence-timer | **10252**
- mac-pinning | **10254**
- mac-pinning (EVPN Routing Instances) | **10256**
- mac-radius | **10257**
- mac-rewrite | **10259**
- mac (Static MAC-Based VLANs) | **10261**
- mac-statistics | **10263**
- mac-table-aging-time (MX Series) | **10265**
- mac-table-aging-time | **10267**
- mac-table-size (VPLS) | **10269**
- mac-table-size | **10271**
- mac-validate | **10273**
- mac-validate (Dynamic IP Demux Interface) | **10275**
- mac-vrf | **10276**
- macsec | **10278**
- macsec (MX Series) | **10281**

- maintain-subscriber (Subscriber Management) | **10284**
- maintenance-association | **10285**
- maintenance-domain | **10287**
- major | **10290**
- major-event (Process) | **10291**
- major-event (System) | **10293**
- major-ring-name | **10295**
- make-before-break (LDP) | **10296**
- managed | **10298**
- managed-configuration | **10301**
- managed-configuration (Dynamic Router Advertisement) | **10302**
- managed-services | **10304**
- management | **10305**
- management-ethernet (Alarm) | **10307**
- management-instance | **10309**
- management (PoE) | **10310**
- management (Security Zones) | **10311**
- mandatory | **10313**
- manual | **10315**
- manual (Junos-FIPS Software) | **10316**
- manual (Junos OS) | **10318**
- manual (Security IPsec) | **10320**
- manual switch | **10323**
- many-to-one (Aggregated Multiservices) | **10324**
- map-dest-bmac-to-dest-cmac | **10326**
- map (Domain Map) | **10327**
- map-e | **10329**
- map-entry-timeout | **10333**
- map-t | **10334**
- mapped-port (Security Static NAT) | **10336**
- mapping | **10338**
- mapping (VNF) | **10340**

- mapping-agent-election | **10342**
- mapping-range | **10344**
- mapping-refresh | **10345**
- mapping-server-entry | **10346**
- mapping-timeout | **10348**
- mapping-timeout (Source NAT Next Gen Services) | **10349**
- mapping-type | **10351**
- mapping-type (Source NAT Next Gen Services) | **10352**
- mark-interface (RA Guard) | **10353**
- mark-local-ip-host-routes | **10355**
- martians | **10356**
- mask uli | **10358**
- master | **10360**
- master-only | **10364**
- master-password | **10365**
- mastership-priority | **10367**
- match (Syslog) | **10369**
- match-direction | **10371**
- match-direction (ipsec-vpn) | **10372**
- match-direction (IDS) | **10373**
- match-direction (IP Reassembly Rule) | **10375**
- match-direction (NAT Next Gen Services) | **10376**
- match-direction (PCP) | **10377**
- match-direction (Services CoS) | **10378**
- match-direction (Services NAT) | **10380**
- match-direction (Services Stateful Firewall) | **10381**
- match (Forwarding-Options) | **10382**
- match-list | **10383**
- match (Next Gen Services Global System Logging) | **10386**
- match-option | **10387**
- match-order (Dynamic Firewalls) | **10389**
- match-rules-on-reverse-flow | **10391**

- match-rules-on-reverse-flow (Next Gen Services) | **10392**
- match (Security Destination NAT) | **10393**
- match (Security IDP Policy) | **10395**
- match (Security Policies) | **10396**
- match (Security Policies Global) | **10398**
- match (Security Source NAT) | **10401**
- match (Security Static NAT) | **10402**
- match (Services CoS Next Gen Services) | **10404**
- match (Stateful Firewall Rule Next Gen Services) | **10406**
- match-strings | **10407**
- matches | **10409**
- matches (Event Policy) | **10412**
- max-advertisement-interval (Dynamic Router Advertisement) | **10413**
- max-advertisement-interval (Protocols IPv6 Neighbor Discovery) | **10415**
- max-advertisement-interval (Protocols Router Discovery) | **10416**
- max-age | **10418**
- max-areas | **10420**
- max-bypasses | **10422**
- max-configurations-on-flash (System) | **10423**
- max-connection-duration | **10424**
- max-connections | **10426**
- max-data-sessions-per-subscriber | **10427**
- max-datasize | **10429**
- max-db-size (Enhanced Subscriber Management) | **10432**
- max-drop-flows | **10435**
- max-duplicates | **10436**
- max-failures | **10438**
- max-flows | **10439**
- max-hops | **10441**
- max-login-sessions | **10442**
- max-login-sessions-per-node | **10444**
- max-lsp-size | **10446**

- max-message-length | **10448**
- max-outstanding-requests (Diameter Applications) | **10449**
- max-packet-memory-ratio | **10451**
- max-packets-per-second (Sampling) | **10452**
- max-pending-accounting-stops (Access Profile) | **10454**
- max-policies (Event Policy) | **10455**
- max-power | **10456**
- max-queues | **10458**
- max-queues-per-interface | **10461**
- max-reass-packet-memory-ratio | **10463**
- max-retrans-time | **10464**
- max-session-number | **10466**
- max-session-setup-rate (Service Set) | **10467**
- max-sessions (Dynamic PPPoE) | **10469**
- max-sessions-per-connection | **10471**
- max-sessions-per-enode | **10472**
- max-sessions-per-subscriber | **10474**
- max-sessions-per-subscriber (Service Set Next Gen Services) | **10475**
- max-sessions (PPPoE Service Name Tables) | **10477**
- max-sessions (Security Packet Log) | **10478**
- max-sessions (Tunnel Profile) | **10480**
- max-sessions-vsa-ignore (Static and Dynamic Subscribers) | **10481**
- max-snp-size | **10483**
- max-tcp-session-packet-memory | **10484**
- max-time-report | **10486**
- max-transmit-quality-level | **10487**
- max-udp-session-packet-memory | **10489**
- max-unknown-messages | **10491**
- max-unknown-requests | **10492**
- max-withhold-time (Access Profile) | **10494**
- maximize-idp-sessions | **10495**
- maximum | **10497**

- maximum-age | **10498**
- maximum-allowed-contentions | **10499**
- maximum-backoff-duration | **10501**
- maximum-bandwidth | **10502**
- maximum-bandwidth (Protocols MPLS) | **10504**
- maximum-bit-rate (Default Local Policy) | **10505**
- maximum-bit-rate (PCC Action Profiles) | **10507**
- maximum-bit-rate (TDF Domain) | **10509**
- maximum-call-duration (Security) | **10510**
- maximum-capture-size (Datapath Debug) | **10512**
- maximum-certificates | **10513**
- maximum-connections | **10514**
- maximum-connections-per-client | **10516**
- maximum-contexts | **10517**
- maximum-discovery-table-entries | **10519**
- maximum-ecmp | **10520**
- maximum-helper-recovery-time | **10522**
- maximum-helper-restart-time | **10524**
- maximum-helper-restart-time (RSVP) | **10525**
- maximum-hop-count | **10527**
- maximum-labels | **10528**
- maximum-lease-time (DHCP) | **10530**
- maximum-length (Origin Validation for BGP) | **10531**
- maximum-links | **10532**
- maximum-message-length | **10534**
- maximum (MSDP Active Source Messages) | **10536**
- maximum-neighbor-reconnect-time | **10537**
- maximum-neighbor-recovery-time | **10538**
- maximum-packet-length | **10540**
- maximum-paths | **10542**
- maximum-pending-reqs-limit | **10544**
- maximum-pending-requests (Diameter) | **10546**

- maximum (PIM Entries) | **10547**
- maximum-power (Interface) | **10549**
- maximum-power (PoE) | **10552**
- maximum-prefixes | **10554**
- maximum-rps | **10556**
- maximum-send-window | **10558**
- maximum-sessions | **10559**
- maximum-sessions (L2TP) | **10560**
- maximum-sessions-per-connection | **10562**
- maximum-sessions-per-tunnel | **10563**
- maximum-sessions (TDF Gateway) | **10565**
- maximum-sessions-trap-percentage (TDF Gateway) | **10566**
- maximum-subscribers | **10568**
- maximum-subscribers | **10569**
- maximum-transactions | **10570**
- maximum-transmit-rate (Protocols IGMP) | **10572**
- maximum-transmit-rate (Protocols MLD) | **10574**
- mc-ae | **10575**
- mc-ae-id | **10581**
- mcast stats-enable | **10582**
- mclag | **10584**
- mclag-arp-nd-sync | **10585**
- mclag-arpreq-sync | **10586**
- mdi-mode | **10587**
- mdt | **10590**
- measurement-interval | **10591**
- measurement-type | **10593**
- med-igp-update-interval | **10594**
- media-loss-rate | **10596**
- media-rate-variation | **10597**
- media-source-port-any | **10599**
- media-type (Dual-Purpose Uplink Ports) | **10600**

- media-type (Interfaces) | **10602**
- medium (Tunnel Profile) | **10603**
- member | **10604**
- member (Application Identification) | **10606**
- member-failure-options (Aggregated Multiservices) | **10608**
- member-id (Junos Fusion Satellite Device Cluster) | **10611**
- member-interface (Aggregated Multiservices) | **10612**
- member-link-scheduler | **10614**
- member (MX Series Virtual Chassis) | **10615**
- member (NSSU Upgrade Groups) | **10617**
- member-range | **10618**
- member (Redundant Power System) | **10620**
- member (Security Group VPN) | **10621**
- member (Security IDP) | **10624**
- members | **10625**
- memory | **10628**
- memory-enhanced | **10629**
- memory (System) | **10630**
- memory (TDF Gateway) | **10633**
- menu-item | **10634**
- mep | **10639**
- mesh-group (Protocols VPLS) | **10641**
- message-flood (Security H323) | **10644**
- message-flood (Security MGCP) | **10645**
- message-ie-profile-v1 | **10647**
- message-ie-profile-v2 | **10648**
- message-list | **10650**
- message-processing-model | **10652**
- message-rate-limit (interfaces) | **10653**
- message-rate-limit | **10655**
- message-rate-limit (Flow Monitoring Logs for NAT) | **10656**
- message-size | **10658**

- message-type | **10659**
- saml (Access Profile) | **10662**
- method | **10663**
- metric | **10665**
- metric (Aggregate, Generated, or Static Route) | **10667**
- metric (Diameter Route) | **10668**
- metric (Dynamic Access-Internal Routes) | **10670**
- metric-in (Protocols RIP) | **10671**
- metric-in (Protocols RIPng) | **10673**
- metric-out (Protocols RIP) | **10674**
- metric-out | **10676**
- metric-out (Protocols RIPng) | **10679**
- metric (Policy Action) | **10681**
- metric (Protocols DVMRP) | **10682**
- metric (Protocols IS-IS) | **10684**
- metric (Protocols MPLS) | **10686**
- metric (Protocols OSPF Sham Link) | **10688**
- metric (Qualified Next Hop on Static Route) | **10689**
- metrics-profile | **10691**
- mgcp | **10693**
- mib-profile | **10695**
- micro-sid | **10697**
- microloop-avoidance | **10700**
- microloop-avoidance (Protocols OSPF) | **10702**
- mif (TDF Interface) | **10703**
- mime-pattern | **10705**
- mime-whitelist | **10706**
- min-advertisement-interval (Dynamic Router Advertisement) | **10708**
- min-advertisement-interval (Protocols IPv6 Neighbor Discovery) | **10709**
- min-advertisement-interval (Protocols Router Discovery) | **10711**
- min-message-length | **10713**
- min-rate | **10714**

- min-rate (source-active-advertisement) | **10718**
- mine-dhcp-client-options | **10719**
- mine-dhcpv6-client-options | **10721**
- minimum-bandwidth | **10722**
- minimum-bandwidth-adjust-interval | **10724**
- minimum-bandwidth-adjust-threshold-change | **10725**
- minimum-bandwidth-adjust-threshold-value | **10726**
- minimum-bandwidth (aggregated Ethernet) | **10728**
- minimum-effective-aigp | **10730**
- minimum-hold-time | **10731**
- minimum-interval | **10733**
- minimum-interval (BFD Liveness Detection) | **10735**
- minimum-interval (Liveness Detection) | **10737**
- minimum-interval (PIM BFD Liveness Detection) | **10738**
- minimum-interval (PIM BFD Transmit Interval) | **10739**
- minimum-links | **10741**
- minimum-links (candidate-uplink-port-profile) | **10744**
- minimum-links (Interfaces) | **10746**
- minimum-priority | **10747**
- minimum-quality | **10748**
- minimum-receive-interval (Protocols PIM) | **10751**
- minimum-receive-interval | **10752**
- minimum-receive-interval (BFD Liveness Detection) | **10754**
- minimum-receive-interval (Liveness Detection) | **10756**
- minimum-retransmission-timeout (L2TP Tunnel) | **10757**
- minimum-wait-time | **10758**
- minor | **10760**
- minor-event (Process) | **10761**
- minor-event (System) | **10763**
- mip-half-function | **10765**
- mirror-decrypt-traffic | **10766**
- mirror-filter (Security Forwarding Options) | **10768**

- mirror-flash-on-disk | **10771**
- mirror-once | **10772**
- mixed-rate-mode | **10774**
- mka | **10775**
- mka (MX Series) | **10777**
- mld | **10778**
- mld (Dynamic Profiles) | **10781**
- mld-snooping | **10783**
- mldp-inband-signalling (Protocols Multipoint LDP) | **10787**
- mlfr-uni-nni-bundle-options | **10789**
- mlfr-uni-nni-bundles-inline | **10790**
- mobile-edge session-scale | **10792**
- mobility | **10793**
- mode (RFC 2544 Benchmarking Tests) | **10795**
- mode (ike) | **10797**
- mode | **10798**
- mode (Dynamic Profiles) | **10801**
- mode (IKE) | **10802**
- mode (Interfaces) | **10804**
- mode (IPsec) | **10805**
- mode (Traffic Manager) | **10807**
- mode (Multicast VLAN Registration) | **10809**
- mode (Next Gen Services Service-Set System Logging) | **10811**
- mode (Protocols DVMRP) | **10813**
- mode (Protocols MSDP) | **10814**
- mode (Protocols PIM) | **10816**
- mode (QFX Series) | **10818**
- mode (Security Forwarding Options) | **10819**
- mode (Security Group VPN) | **10821**
- mode (Security Log) | **10823**
- modem-options | **10825**
- modulation-format | **10826**

- mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain) | **10827**
- mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain) | **10829**
- mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain) | **10830**
- mofrr-primary-path-selection-by-routing (Multicast-Only Fast Reroute) | **10832**
- monitor-bandwidth | **10834**
- monitor-end-point | **10835**
- monitor ethernet delay-measurement | **10836**
- monitor ethernet loss-measurement | **10843**
- monitor ethernet synthetic-loss-measurement | **10850**
- monitor interface | **10856**
- monitor interface error | **10874**
- monitor label-switched-path | **10877**
- monitor list | **10881**
- monitor list (JDM) | **10883**
- monitor mpls delay rsvp | **10884**
- monitor mpls loss-delay rsvp | **10890**
- monitor mpls loss rsvp | **10896**
- monitor (Multinode High Availability) | **10903**
- monitor-object | **10906**
- monitor pfe traffic interface | **10908**
- monitor-profile | **10912**
- monitor security flow file | **10914**
- monitor security flow filter | **10916**
- monitor security flow start | **10918**
- monitor security flow stop | **10920**
- monitor security packet-drop | **10921**
- monitor-session | **10926**
- monitor start | **10927**
- monitor start (JDM) | **10930**
- monitor stop | **10932**
- monitor stop (JDM) | **10934**
- monitor traffic | **10935**

- monitor-timer | **10953**
- monitoring | **10954**
- monitoring (Forwarding Options) | **10956**
- monitoring-key (PCC Action Profile) | **10957**
- monitoring-profile | **10959**
- monitoring (Services) | **10961**
- most-fill | **10967**
- moving-average-size | **10967**
- mpls | **10969**
- mpls-flow-table-size | **10972**
- mpls (security) | **10974**
- mpls (forwarding-options) | **10975**
- mpls-inet6 | **10977**
- mpls-internet-multicast | **10978**
- mpls-ipv4-template | **10979**
- mpls-ipvx-template | **10981**
- mpls (packet-forwarding-options) | **10982**
- mpls (Protocols) | **10984**
- mpls (Security Forwarding Options) | **10985**
- mpls-template | **10986**
- mpls-tp-mode | **10988**
- mrru | **10989**
- mru (Gigabit Ethernet Options) | **10991**
- mru (Congestion Notification Profiles) | **10992**
- mru (Dynamic and Static PPPoE) | **10994**
- msdp | **10996**
- msrpc | **10999**
- mss (IDS MS-DPC) | **11001**
- mss (Security IDP) | **11002**
- msti | **11003**
- mstp | **11006**
- mtrace | **11010**

- mtrace from-source | **11013**
- mtrace monitor | **11017**
- mtrace to-gateway | **11019**
- mtu (QFX Series) | **11023**
- mtu (Multilink and Link Services Logical Interface) | **11025**
- mtu | **11026**
- mtu (interfaces) | **11027**
- mtu-discovery | **11032**
- mtu (Dynamic and Static PPPoE) | **11034**
- mtu (Dynamic Profiles) | **11035**
- mtu-signaling | **11037**
- mtu (TDF Interface) | **11038**
- multi-address-embedded-option-response (DHCP Local Server) | **11039**
- multi-bgp-path | **11041**
- multi-chassis | **11042**
- multi-chassis-protection | **11044**
- multi-chassis-protection (Interfaces) | **11045**
- multi-destination (Class of Service) | **11047**
- multi-domain | **11048**
- multi-homing (VPLS Multihoming for BGP VPLS) | **11051**
- multi-homing (VPLS Multihoming for FEC 129) | **11052**
- multi-link-layer-2-inline | **11054**
- multi-sa | **11055**
- multi-tenant-hash | **11057**
- multi-tenant-support | **11059**
- multi-bgp-path | **11060**
- multicast | **11062**
- multicast (routing-options) | **11063**
- multicast-dlci | **11066**
- multicast (Dynamic Routing Options) | **11067**
- multicast-group | **11069**
- multicast-ipv4 | **11070**

- multicast-loadbalance | **11072**
- multicast-mode (EVPN) | **11074**
- multicast-only | **11076**
- multicast (Protocols Router Discovery) | **11077**
- multicast (QFabric Routing Options) | **11078**
- multicast-replication | **11080**
- multicast-router-interface (IGMP Snooping) | **11084**
- multicast-router-interface (MLD Snooping) | **11086**
- multicast-snooping-options | **11087**
- multicast-statistics | **11089**
- multicast-statistics (packet-forwarding-options) | **11091**
- multicast (Virtual Tunnel in Routing Instances) | **11092**
- multicast-vlan-flood | **11094**
- multichassis-lag-replicate-state | **11095**
- multichunk-inspection | **11097**
- multihop | **11098**
- multilink | **11103**
- multilink-class | **11104**
- multilink-max-classes | **11106**
- multipath-build-priority | **11107**
- multipath (OSPF SPF Options) | **11109**
- multipath (Protocols BGP) | **11110**
- multipath (Routing Options) | **11113**
- multipath-rule | **11115**
- multipath (SPF Options) | **11117**
- multiple-triggered-joins | **11119**
- multiplier (Protocols PIM) | **11120**
- multiplier | **11122**
- multiplier (BFD Liveness Detection) | **11123**
- multiplier (Liveness Detection) | **11125**
- multipoint | **11127**
- multiservice | **11128**

- multiservice-options (Flow Monitoring) | **11130**
- multiservice-options | **11131**
- must-secure | **11133**
- mutual-authentication (REST API) | **11134**
- mvpn | **11135**
- mvpn (Draft-Rosen MVPN) | **11138**
- mvpn-iana-rt-import | **11140**
- mvpn-mode | **11142**
- mvpn (NG-MVPN) | **11143**
- mvrp | **11145**

mac

IN THIS SECTION

- Syntax | **10215**
- Hierarchy Level | **10216**
- Description | **10216**
- Options | **10216**
- Required Privilege Level | **10216**
- Release Information | **10216**

Syntax

```
mac mac-address;
```

Hierarchy Level

- For platforms with Enhanced Layer 2 Software (ELS):

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name interface interface-name static-ip  ip-address]
```

- For platforms without ELS:

```
[edit ethernet-switching-options secure-access-port interface (all | interface-name) static-ip  ip-address vlan  vlan-name]
```

- For MX Series platforms:

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security group group-name interface interface-name static-ip  ip-address]
```

Description

Configure the media access control (MAC) address or hardware address of the device connected to the specified interface.

Options

mac-address—Value (in hexadecimal format) of the address assigned to this device.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

Hierarchy level `[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]` introduced in Junos OS Release 14.1 for the MX Series.

RELATED DOCUMENTATION

> *Configuring Static DHCP IP Addresses for DHCP snooping (non-ELS)*
>
> *Configuring Static DHCP IP Addresses for DHCP snooping (ELS)*
>
> *Configuring Static DHCP IP Addresses for DHCP snooping (MX routers)*

mac (MPLS)

IN THIS SECTION

- Syntax | **10217**
- Hierarchy Level | **10217**
- Description | **10217**
- Options | **10218**
- Required Privilege Level | **10218**
- Release Information | **10218**

Syntax

```
mac mac-address;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Set the MAC address of the interface.

Use this statement at the `[edit interfaces ... ps0]` hierarchy level to configure the MAC address for a pseudowire logical device that is used for subscriber interfaces over point-to-point MPLS pseudowires.

Options

mac-address—MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn*.*nnnn*.*nnnn* or *nn*:*nn*:*nn*:*nn*:*nn*:*nn*. For example, `0000.5e00.5355` or `00:00:5e:00:53:55`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> Configuring the MAC Address on the Management Ethernet Interface
>
> *Configuring a Pseudowire Subscriber Logical Interface Device*

mac-address

IN THIS SECTION

- Syntax | **10219**
- Hierarchy Level | **10219**
- Description | **10219**
- Required Privilege Level | **10219**
- Release Information | **10219**

Syntax

```
mac-address  mac-address;
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

MAC address for the VNF interfaces.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D47.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | **27723**

mac-address (Accept Source Mac)

IN THIS SECTION

- Syntax | **10220**
- Hierarchy Level | **10220**

- Description | 10220
- Options | 10220
- Required Privilege Level | 10220
- Release Information | 10221

Syntax

```
mac-address mac-address policer;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number accept-source-mac],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
accept-source-mac]
```

Description

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), for Gigabit Ethernet DPCs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP, specify a remote MAC address on which to count incoming and outgoing packets.

Options

mac-address—MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn*.*nnnn*.*nnnn* or *nn*:*nn*:*nn*:*nn*:*nn*:*nn*. For example, `0011.2233.4455` or `00:11:22:33:44:55`.

`policer`—MAC policer. For more information, see "policer (MAC)" on page 12641.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

mac-address (DHCP Local Server)

IN THIS SECTION

- Syntax | 10221
- Hierarchy Level | 10221
- Description | 10222
- Required Privilege Level | 10222
- Release Information | 10222

Syntax

```
mac-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication username-
include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
```

```
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify that the MAC address from the client PDU be concatenated with the username during the subscriber authentication or DHCP client authentication process.

For DHCPv6 clients, because the DHCPv6 packet format has no specific field for the client MAC address, the MAC address is derived from among several sources with the following priority:

- Client DUID Type 1 or Type 3.
- Option 79 (client link-layer address), if present.
- The packet source address if the client is directly connected.
- The link local address.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support for DHCPv6 added in Junos OS Release 17.2 for MX Series Routers.

RELATED DOCUMENTATION

Specifying Authentication Support

mac-address (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 10223
- Hierarchy Level | 10223
- Description | 10223
- Required Privilege Level | 10224
- Release Information | 10224

Syntax

```
mac-address;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify that the MAC address from the client PDU be concatenated with the username during the subscriber authentication or client authentication process.

For DHCPv6 clients, because the DHCPv6 packet format has no specific field for the client MAC address, the MAC address is derived from among several sources with the following priority:

- Client DUID Type 1 or Type 3.

- Option 79 (client link-layer address), if present.
- The packet source address if the client is directly connected.
- The link local address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support for DHCPv6 added in Junos OS Release 17.2 for MX Series Routers.

RELATED DOCUMENTATION

Specifying Authentication Support

mac-address (Dynamic Access-Internal Routes)

IN THIS SECTION

- Syntax | 10225
- Hierarchy Level | 10225
- Description | 10225
- Options | 10225
- Required Privilege Level | 10225
- Release Information | 10225

Syntax

```
mac-address address;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
access-internal route subscriber-ip-address qualified-next-hop underlying-interface],
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
rib routing-table-name access-internal route subscriber-ip-address qualified-next-hop underlying-
interface],
[edit dynamic-profiles routing-options access-internal route subscriber-ip-address qualified-next-
hop underlying-interface].
```

Description

Dynamically configure the MAC address variable for an access-internal route for unnumbered interfaces such as DHCP subscriber interfaces.

Options

address—Either the specific MAC address you want to assign to the access-internal route or the MAC address variable ($junos-subscriber-mac-address). The MAC address variable is dynamically replaced with the value supplied by DHCP when a subscriber logs in.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options route subscriber-ip-address qualified-next-hop underlying-interface]` and `[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options rib routing-table-name route subscriber-ip-address qualified-next-hop underlying-interface]` hierarchy levels introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Routes for DHCP and PPP Subscriber Access Networks

mac-address (MACsec)

IN THIS SECTION

- Syntax | 10226
- Hierarchy Level | 10226
- Description | 10226
- Default | 10227
- Options | 10227
- Required Privilege Level | 10227
- Release Information | 10227

Syntax

```
mac-address mac-address;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel secure-channel-name id]
```

Description

Specify a MAC address to enable MACsec using static secure association key (SAK) security mode. The `mac-address` variables must match on the sending and receiving ends of a link to enable MACsec using static SAK security mode.

If you are configuring a MAC address on a secure channel in the outbound direction, you should specify the MAC address of the interface as the `mac-address`.

If you are configuring a MAC address on a secure channel in the inbound direction, you should specify the MAC address of the interface at the other end of the link as the `mac-address`.

You only use this configuration option when you are configuring MACsec using static SAK security mode. This option does not need to be specified when you are enabling MACsec using static connectivity association key (CAK) security mode.

Default

No MAC address is specified in the secure channel, by default.

Options

mac-address	The MAC address, in six groups of two hexadecimal digits.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

mac-address (VLAN and Stacked VLAN Interfaces)

IN THIS SECTION

- Syntax | 10228
- Hierarchy Level | 10228
- Description | 10228

Required Privilege Level | 10228
Release Information | 10228

Syntax

```
mac-address;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication username-include],
```

Description

Specify that the client hardware address (chaddr) from the incoming DHCP discover packet be concatenated with the username during the subscriber authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

mac-flush

IN THIS SECTION

- Syntax | 10229
- Hierarchy Level | 10229
- Description | 10229
- Options | 10230
- Required Privilege Level | 10231
- Release Information | 10231

Syntax

```
mac-flush [ explicit-mac-flush-message-options ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]
```

Description

Enable media access control (MAC) flush processing for the virtual private LAN service (VPLS) routing instance or for the mesh group under a VPLS routing instance. MAC flush processing removes MAC addresses from the MAC address database that have been learned dynamically. With the dynamically learned MAC addresses removed, MAC address convergence requires less time to complete.

For certain cases where MAC flush processing is not initiated by default, you can also specify *explicit-mac-flush-message-options* that additionally configure the router to send explicit MAC flush messages. To

configure the router to send explicit MAC flush messages under specific conditions, include ***explicit-mac-flush-message-options*** with the statement.

> **NOTE**: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

In certain cases, BGP updates sent by the provider edge (PE) device are delayed for 1 to 5 seconds.

This happens when all of the following condtions are true:

- BGP-based VPLS multihoming sites are configured.
- The `mac-flush` statement is included in the configuration.
- a non-minimum designated-forwarder site (site-x, for example) transitions to non-designated-forwarder status

The BGP update being delayed corresponds to the explicit-MAC flush notification message sent by site-x's PE device (PE2, for example). This BGP update message is not deferred if the designated-forwarder status is lost due to a locally-triggered event (for example, a local attachment-circuit interface going down). In other words, BGP update messages are deferred (in Device PE2) only when the designated-forwarder state is lost due to external events taking place in remote PE devices that also hold site-x (for example, in PE1). Suppose, for example, that Device PE1 is the default designated-forwarder with site-x's local interface in the DOWN state. Device PE2 defers BGP update message after Device PE1's local interface comes back to the UP state.

Options

`explicit-mac-flush-message-options`—(Optional) You can specify one or more of the following explicit MAC flush message options:

- `any-interface`—(Optional) Send a MAC flush message when any customer-facing attachment circuit interface goes down.
- `any-spoke`—(Optional) Send a MAC FLUSH-FROM-ME flush message to all provider edge (PE) routers in the core when one of the spoke pseudowires between the multitenant unit switch and the other network-facing provider edge (NPE) router goes down, causing the multitenant unit switch to switch to this NPE router.

  > **NOTE**: This option has a similar effect in a VPLS multihoming environment with multiple multitenant unit switches connected to NPE routers, where both multitenant unit switches have pseudowires that terminate in a mesh group with local-switching

configured. If the `any-spoke` option is enabled, then both PE routers send MAC FLUSH-FROM-ME flush messages to all PEs in the core.

- `propagate`—(Optional) Propagate MAC flush to the core.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring VPLS Routing Instances

Configuring Interoperability Between BGP Signaling and LDP Signaling in VPLS

mac-ip-table-size

IN THIS SECTION

- Syntax | 10232
- Hierarchy Level | 10232
- Description | 10232
- Options | 10232
- Required Privilege Level | 10232
- Release Information | 10232

Syntax

```
mac-ip-table-size number;
```

Hierarchy Level

```
[edit logical-systems name routing-instances routng-instance-name protocols evpn],
[edit routing-instances routing-instance-name protocols evpn],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options],
[edit routing-instances routing-instance-name vlans vlan-name switch-options]
```

Description

Limit the number of entries that can be added to the MAC-IP address bindings database. Depending on the Juniper Networks device, this limit can be applied to EVPN routing instances, bridge domains configured in a virtual-switch routing instance, or VLANs configured in a virtual-switch routing instance. When the specified maximum number of entries is reached, no additional entries are added to the MAC-IP bindings database.

To apply this limit systemwide, use the `global-mac-ip-limit` statement at the `[edit protocols l2-learning]` hierarchy level.

Options

number Maximum number of entries that can be added to the MAC-IP address bindings database.

- **Range:** 16 through 1,048,575.
- **Default:** 8192

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R2.

Statement introduced in Junos OS Release 18.1R1 for QFX Series switches.

RELATED DOCUMENTATION

EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression

interface-mac-ip-limit | 8908

mac-learn-enable

IN THIS SECTION

- Syntax | 10233
- Hierarchy Level | 10233
- Description | 10233
- Required Privilege Level | 10234
- Release Information | 10234

Syntax

```
mac-learn-enable;
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]
[edit interfaces aex aggregated-ether-options ethernet-switch-profile]
```

Description

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), for Gigabit Ethernet DPCs on MX Series routers, for 100-Gigabit Ethernet Type 5 PIC with CFP, and for MPC3E, MPC4E, MPC5E, MPC5EQ, and

MPC6E MPCs, configure dynamic learning of the source and destination MAC addresses. By default, the interface is not allowed to dynamically learn source and destination MAC addresses.

To disable dynamic learning of the source and destination MAC addresses after it has been configured, you must delete `mac-learn-enable` from the configuration.

MPCs support MAC address accounting for an individual interface or an aggregated Ethernet interface member link only after the interface has received traffic from the MAC source. If traffic is only exiting an interface, the MAC address is not learned and MAC address accounting does not occur.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for statement under the `[edit interfaces ae*x* aggregated-ether-options ethernet-switch-profile]` hierarchy introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

Configuring MAC Address Accounting

mac-learning-priority

IN THIS SECTION

- Syntax | 10235
- Hierarchy Level | 10235
- Description | 10235
- Options | 10235
- Required Privilege Level | 10235

Release Information | 10236

Syntax

```
mac-learning-priority value packet-action action;
```

Hierarchy Level

```
[edit switch-options]
```

Description

Configure MAC learning priority on the interfaces so that MAC addresses are always learnt on the high priority interface.

Options

value—Specify the MAC learning priority for the interface.

- **Range:** 1 through 65535
- **Default:** 4

action—Discard the packets and do not learn.

- **Action:** discard
- **Default:** Forward

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1.

mac-limit

IN THIS SECTION

- Syntax | 10236
- Hierarchy Level | 10236
- Description | 10236
- Default | 10237
- Options | 10237
- Required Privilege Level | 10237
- Release Information | 10237

Syntax

```
mac-limit limit {
    <action action>;
}
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port interface (all | interface-name)]
```

Description

Specify the number of MAC addresses that can be dynamically added to the MAC address cache for this access interface (port) and the action to be taken if the limit is exceeded.

Default

The default action is **drop**.

Options

limit—Maximum number of MAC addresses.

action *action*—(Optional) Action to take when the MAC address limit is exceeded:

- **drop**—Drop the packet and generate a system log entry. This is the default.
- **log**—Do not drop the packet but generate a system log entry.
- **none**—No action.
- **shutdown**—Disable the interface and generate an alarm. If you configure the switch with the `port-error-disable` statement, the disabled interface recovers automatically upon expiration of the specified timeout. If this statement is not configured, you can bring up the disabled interfaces by running the `clear ethernet-switching port-error` command.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Understanding MAC Limiting and MAC Move Limiting for Port Security

Configuring MAC Limiting (QFX Switches)

allowed-mac

mac-limit (VLAN)

IN THIS SECTION

- Syntax (QFX Series and EX4600) | 10238
- Syntax (SRX Series and EX Series) | 10238
- Hierarchy Level | 10238
- Description | 10238
- Default | 10239
- Options | 10239
- Required Privilege Level | 10240
- Release Information | 10240

Syntax (QFX Series and EX4600)

```
mac-limit number;
```

Syntax (SRX Series and EX Series)

```
mac-limit limit action action;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Description

Specify the maximum number of MAC addresses to be associated with a VLAN—the default is **unlimited**, which can leave the network vulnerable to flooding. Change **unlimited** to any number from 2 to the switch's maximum VLAN MAC limit. The maximum number of MAC addresses allowed in a switching table per VLAN varies depending on the EX Series switch. To see the maximum number of MAC

addresses per VLAN allowed on your switch, issue the `set vlans vlan-name mac-limit ?` configuration-mode command.

> **NOTE**: Do not set the **mac-limit** value to 1. The first learned MAC address is often inserted into the forwarding database automatically—for instance, for a routed VLAN interface (RVI), the first MAC address inserted into the forwarding database is the MAC address of the RVI. For aggregated Ethernet bundles (LAGs) using LACP, the first MAC address inserted into the forwarding database in the Ethernet switching table is the source address of the protocol packet. In these cases, the switch does not learn MAC addresses other than the automatic address when **mac-limit** is set to 1, and this causes problems with MAC learning and forwarding.

When the MAC limit set by this statement is reached, no more MAC addresses are added to the Ethernet switching table. You can also, optionally, have a system log entry generated when the limit is exceeded by adding the option **action log**.

> **NOTE**: When you reconfigure the number of MAC addresses, the Ethernet switching table is not automatically cleared. Therefore, if you reduce the number of addresses from the default (unlimited) or a previously set limit, you could already have more entries in the table than the new limit allows. Previous entries remain in the table after you reduce the number of addresses, so you should clear the Ethernet switching table for a specified interface, MAC address, or VLAN when you reduce the MAC limit. Use the command **clear ethernet-switching table** to clear existing MAC addresses from the table before using the **mac-limit** configuration statement.

Default

The MAC limit is disabled, so entries are unlimited.

Options

QFX Series and EX4600:

number—Maximum number of MAC addresses.

- **Range:** 1 through 32768

> **NOTE**: This statement is not supported on QFabric systems.

EX Series:

limit—Maximum number of MAC addresses.

- **Range:** 1 through *switch maximum*

SRX Series:

number—Maximum number of MAC addresses.

- **Range:** 1 through 5120

action—**Log** is the only action available. Configure **action log** to add a message to the system log when the mac-limit value is exceeded. A typical logged message looks like this:

May 5 06:18:31 bmp-199p1-dev edwd[5665]: ESWD_VLAN_MAC_LIMIT_EXCEEDED: vlan default mac 00:1f:12:37:af:5b (tag 40). vlan limit exceeded

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

The short description of `interface-mac-limit` at the CLI command hierarchy is changed from `Maximum number of MAC addresses per interface (1..16383)` to `Maximum number of MAC addresses per interface (1..5120)` at the `[edit vlans vlan-name switch-options]` hierarchy level from Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show vlans | **27738**

Example: Setting Up Basic Bridging and a VLAN on Switches

Configuring MAC Table Aging on Switches

Understanding Bridging and VLANs on Switches

mac-limit (Access Port Security)

IN THIS SECTION

- Syntax | 10241
- Hierarchy Level | 10241
- Description | 10241
- Default | 10242
- Options | 10242
- Required Privilege Level | 10242
- Release Information | 10243

Syntax

```
mac-limit limit action action;
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port interface (all | interface-name)],
[edit ethernet-switching-options secure-access-port interface interface-name vlan vlan-name]
```

Description

Set a limit on the number of MAC addresses that can be added to the Ethernet switching table.

- [edit ethernet-switching options secure-access-port interface]—Set the MAC address learning limit for a specific interface, for a range of interfaces, or for all interfaces on the switch.
- [edit ethernet-switching options secure-access-port interface *interface-name* vlan *vlan-name*]—Set the MAC address learning limit for a specific interface as a member of a specific VLAN (VLAN membership MAC limit).

> **NOTE:** If you set the MAC address limit on a specific interface as a member of a specific VLAN (VLAN membership MAC limit), the switch drops any additional packets when the VLAN membership MAC limit is exceeded and logs the MAC addresses of those packets. You cannot specify a different action for this specific configuration. If a single interface belongs to more than one VLAN, you can set separate VLAN membership MAC limits for the same interface.

When you reset the number of MAC addresses, the MAC address table is not automatically cleared. Previous entries remain in the table after you reduce the number of addresses, so you should clear the forwarding table for the specified interface or MAC address. Use the `clear ethernet-switching table` command to clear the existing MAC addresses from the table.

Default

The default action is `drop`.

Options

action *action*—(Optional) Action to take when the MAC address limit for an interface or for all interfaces is exceeded:

- **drop**—Drop the packet and generate a system log entry.
- **log**—Do not drop the packet but generate a system log entry.
- **none**—No action.
- **shutdown**—Disable the interface and generate a system log entry. If you have configured the switch with the `port-error-disable` statement, the disabled interface recovers automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the `clear ethernet-switching port-error` command.

limit—Maximum number of MAC addresses.

Required Privilege Level

system—To view this statement in the configuration.
system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- allowed-mac | 2214
- clear ethernet-switching table persistent-learning | 3586
- *Example: Configuring Port Security (non-ELS)*
- *Example: Protecting against Ethernet Switching Table Overflow Attacks*
- *Example: Protecting against DHCP Starvation Attacks*
- *Configuring MAC Limiting (non-ELS)*
- *Configuring MAC Limiting (J-Web Procedure)*
- *Configuring Autorecovery for Port Security Events*

mac-list

IN THIS SECTION

- Syntax | 10243
- Hierarchy Level | 10244
- Description | 10244
- Options | 10244
- Required Privilege Level | 10244
- Release Information | 10244

Syntax

```
mac-list name {
    mac-addresses;
}
```

Hierarchy Level

```
[edit policy-options]
```

Description

Define a list of MAC addresses for use in an IPv6 Router Advertisement (RA) guard policy.

Options

mac-addresses—List of MAC addresses, one MAC address per line in the configuration.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding Prefix Lists for Use in Routing Policy Match Conditions

mac-lookup-length

IN THIS SECTION

- Syntax | 10245
- Hierarchy Level | 10245
- Description | 10245
- Default | 10245

- Options | 10246
- Required Privilege Level | 10246
- Release Information | 10246

Syntax

```
mac-lookup-length number-of-entries;
```

Hierarchy Level

```
[edit ethernet-switching-options]
```

Description

Increase the maximum number of searchable hash indexes to mitigate situations in which hash index collisions are causing problems with the learning of MAC addresses in the forwarding database (FDB).

The FDB on EX3200, EX3300, EX4200, EX4500, EX4550, and EX6210 switches is a hash table with 8192 hash indexes (rows) of MAC addresses and four entries per hash index. When the FDB is searched, a configured hash function calculates the hash index at which to start the search. By default, after the search starts at the determined hash index, the maximum number of hash indexes that can be searched is one hash index, or four entries

> **NOTE**: Increasing the number of hash indexes increases the chances of finding an open entry in which to add a newly learned MAC address. However, searching more hash indexes requires more bandwidth and may impact the FDB performance and line-rate traffic.

Default

4

Options

number-of-entries—Maximum number of searchable hash indexes in the FDB.

- **Range:** 4, 8, 12

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Bridging and VLANs on Switches

mac-move-limit

IN THIS SECTION

- Syntax | 10247
- Hierarchy Level | 10247
- Description | 10247
- Default | 10247
- Options | 10247
- Required Privilege Level | 10248
- Release Information | 10248

Syntax

```
mac-move-limit {
    limit;
    <action action | packet-action action>;
}
```

Hierarchy Level

- For platforms with ELS:

```
[edit vlans vlan-name  switch-options]
```

- For platforms without ELS:

```
[edit ethernet-switching-options                secure-access-port                vlan (all |
vlan-name)]
```

Description

Specify the number of times a MAC address can move to a new interface (port) in one second and the action to be taken by the switch if the MAC address move limit is exceeded.

Starting in Junos OS 14.1X53-D51, do not configure both the `no-mac-learning` statement at `[edit ethernet-switching-options interfaces interface-name]` and the `mac-move-limit` statement. Because MAC move limiting requires that the device learns MAC addresses, you cannot disable MAC learning.

Default

If you do not specify `mac-move-limit`, the default MAC address move limit is unlimited.

Options

`limit` *limit*—Maximum number of moves to a new interface per second.

- `action` *action*—(Optional) (Available *only* under the hierarchy level `[edit ethernet-switching-options secure-access-port vlan (all | vlan-name) mac-move-limit]`) Action to take when the MAC address move limit is reached:

- `drop`—Drop the packet and generate a system log entry. This is the default.
- `log`—Do not drop the packet but generate a system log entry.
- `none`—No action.
- `shutdown`—Logically disable the interface and generate a system log entry. If you have configured the switch with the `port-error-disable` statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the `clear ethernet-switching port-error` command.

- `packet-action` *`action`*—(Optional) (Available *only* under the hierarchy level, `[edit vlans` *`vlan-name`* `switch-options mac-move-limit]`) Action to take when the MAC address move limit is reached:

  NOTE: There is no default action.

  - `drop`—Drop the packet and do not generate an alarm.
  - `drop and log`—Drop the packet and generate an alarm, an SNMP trap, or system log entry.
  - `log`— Do not drop the packet, but generate an alarm, an SNMP trap, or a system log entry.
  - `none`—No action.
  - `shutdown`—Logically disable the interface and generate an alarm or an SNMP trap. If you have configured the interface with the `recovery-timeout` statement, the disabled interface recovers automatically upon expiration of the specified timeout. If you have not configured the interface for a recovery timeout, you can bring up the disabled interface by running the operational command `clear ethernet-switching recovery-timeout`.

Required Privilege Level

system—To view this statement in the configuration.
system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Hierarchy level `[edit vlans` *`vlan-name`* `switch-options]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

Example: Configuring Port Security (non-ELS)

Configuring MAC Move Limiting (ELS)

Configuring Persistent MAC Learning (ELS)

Configuring Autorecovery for Port Security Events

Configuring Autorecovery for Port Security Events

mac-notification

IN THIS SECTION

- Syntax | 10249
- Hierarchy Level | 10249
- Description | 10250
- Default | 10250
- Required Privilege Level | 10250
- Release Information | 10250

Syntax

```
mac-notification {
    notification-interval seconds;
}
```

Hierarchy Level

```
[edit ethernet-switching-options]
[edit switch-options]
```

Description

Enable MAC notification for a switch. If you configure this statement without setting a notification interval, MAC notification is enabled with the default MAC notification interval of 30 seconds.

The remaining statement is explained separately. See CLI Explorer.

Default

MAC notification is disabled by default.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Hierarchy level `[edit switch-options]` added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.

RELATED DOCUMENTATION

Configuring Non-ELS MAC Notification

Configuring MAC Notification on Switches with ELS Support

mac (Option 82)

IN THIS SECTION

- Syntax | **10251**
- Hierarchy Level | **10251**

- Description | 10251
- Required Privilege Level | 10251
- Release Information | 10251

Syntax

```
mac;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security option-82 remote-id]
```

Description

Use the MAC address of the port connected to the DHCP client as the `remote-id` suboption of DHCP option 82 (also known as the DHCP relay agent information option) in DHCP request packet headers before forwarding or relaying requests to a DHCP server. This suboption provides a trusted identifier for the host system that has forwarded or relayed requests to the server.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

> *Example: Setting Up DHCP Option 82 (No Relay)*
>
> http://tools.ietf.org/html/rfc3046

mac-persistence-timer

IN THIS SECTION

- Syntax | 10252
- Hierarchy Level | 10252
- Description | 10252
- Default | 10252
- Options | 10253
- Required Privilege Level | 10253
- Release Information | 10253

Syntax

```
mac-persistence-timer (minutes | disable);
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Specify how long the Virtual Chassis or VCF continues to use the MAC address of the switch that was originally configured in the primary role as the system MAC base address after the original primary switch is removed from the Virtual Chassis or VCF. The system MAC base address does not change in the event of a switchover provided the switch originally configured in the primary role remains a member of the Virtual Chassis or VCF.

The maximum timer limit is 60 minutes starting in Junos OS Release 12.2. There are no minimum or maximum timer limits in prior Junos OS releases.

Default

The MAC persistence timer is set to 10 minutes by default.

Options

minutes Time in minutes that the member switch in the backup role continues to use the system MAC base address of the old primary before using its own system MAC base address after the switch in the primary role is physically disconnected or removed from the Virtual Chassis or VCF.

`disable` Disable the MAC persistence timer. The system MAC base address never changes when the MAC persistence timer is disabled, even when the switch in the primary role is physically disconnected or removed from the Virtual Chassis or VCF.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Option `disable` introduced in Junos OS Release 12.2 for EX Series switches.

The maximum timer limit changed from no maximum timer limit to 60 minutes in Junos OS Release 12.2 for EX Series switches.

RELATED DOCUMENTATION

- Configuring the Timer for the Backup Member to Start Using Its Own MAC Address as Primary of a Virtual Chassis
- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*

mac-pinning

IN THIS SECTION

- Syntax | 10254
- Hierarchy Level | 10254
- Description | 10255
- Default | 10255
- Required Privilege Level | 10255
- Release Information | 10255

Syntax

```
mac-pinning;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name bridge-options interface interface-name],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name bridge-options
interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name switch-options
interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options
interface interface-name],
[edit routing-instances routing-instance-name protocols evpn],
[edit routing-instances routing-instance-name switch-options interface interface-name],
[edit switch-options interface interface-name],
[edit switch-options interface interface-name interface-mac-limit  limit]
```

Description

Enable MAC pinning on an interface. When you enable MAC pinning on an interface in a bridge domain or VPLS domain, MAC addresses learned over that interface cannot be relearned on any other interface in the same bridge domain or VPLS domain until the MAC address is cleared from the MAC table. If the same MAC address is learned over any other interface in the same bridge domain, it is discarded. This, effectively, controls MAC address moves and prevents creation of loops in Layer 2 bridges.

> **NOTE**: If the timeout interval for the MAC addresses is not specified by setting the `mac-table-aging-time` statement, the MAC addresses learned over the MAC pinning interface are pinned to the interface until the default timeout period.

Default

MAC pinning is not enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support at `[edit routing-instances routing-instance-name protocols evpn]` hierarchy introduced in Junos OS Release 17.2R1 on MX Series Routers with MPC and MIC interfaces.

RELATED DOCUMENTATION

- Understanding Layer 2 Learning and Forwarding for Bridge Domains
- Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports
- *Understanding MAC Pinning*

mac-pinning (EVPN Routing Instances)

IN THIS SECTION

- Syntax | 10256
- Hierarchy Level | 10256
- Description | 10256
- Required Privilege Level | 10256
- Release Information | 10257

Syntax

```
mac-pinning;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols protocol interface interface-name]
[edit logical-systems logical-system-name routing-instances routing-instance-name vlans vlan-namebridge-domains bridge-name bridge-options-interface interface-name]
[edit logical-systems logical-system-name routing-instances routing-instance-name vlans vlan-namerouting-instances routing-instance-name switch-options interface interface-name]
```

Description

Sets mac-pinning on the interface, which makes it static. That is, the MAC-pinned address cannot be moved to any other interface in the bridge domain.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

EVPN MAC Pinning Overview

Configuring EVPN MAC Pinning

Configuring MAC Pinning for PBB-EVPN

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports

mac-radius

IN THIS SECTION

- Syntax | 10257
- Hierarchy Level | 10258
- Description | 10258
- Options | 10258
- Required Privilege Level | 10259
- Release Information | 10259

Syntax

```
mac-radius {
    authentication-protocol {
        eap-md5;
        eap-peap {
            resume;
        }
        pap;
```

```
    }
    flap-on-disconnect;
    restrict;
}
```

Hierarchy Level

```
[edit protocols dot1x authenticator                interface interface-name]
```

Description

Configure MAC RADIUS authentication for specific interfaces. MAC RADIUS authentication allows LAN access to permitted MAC addresses. When a new MAC address appears on an interface, the device consults the RADIUS server to check whether the MAC address is a permitted address. If the MAC address is configured on the RADIUS server, the device is allowed access to the LAN.

If MAC RADIUS is configured, the device first tries to get a response from the host for 802.1X authentication. If the host is unresponsive, the device attempts to authenticate using MAC RADIUS.

To restrict authentication to MAC RADIUS only, use the `restrict` option. In restrictive mode, all 802.1X packets are eliminated and the attached device on the interface is considered a nonresponsive host.

Options

flap-on-disconnect (Optional) When the RADIUS server sends a disconnect message to a supplicant, the device resets the interface on which the supplicant is authenticated. If the interface is configured for multiple supplicant mode, the device resets all the supplicants on the specified interface. This option takes effect only when the `restrict` option is also set.

restrict (Optional) Restricts authentication to MAC RADIUS only. When `mac-radius restrict` is configured, the device drops all 802.1X packets. This option is useful when no other 802.1X authentication methods, such as guest VLAN, are needed on the interface, and eliminates the delay that occurs while the switch determines that a connected device is a non-802.1X-enabled host.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

`flap-on-disconnect` introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

- show dot1x | **19548**
- password (MAC RADIUS Authentication) | **12241**
- *Example: Configuring MAC RADIUS Authentication on an EX Series Switch*
- *Example: Setting Up 802.1X for Single-Supplicant or Multiple-Supplicant Configurations on an EX Series Switch*
- *Configuring MAC RADIUS Authentication (CLI Procedure)*
- *Configuring 802.1X Interface Settings (CLI Procedure)*
- Understanding Authentication on Switches

mac-rewrite

IN THIS SECTION

- Syntax | **10260**
- Hierarchy Level | **10260**
- Description | **10260**
- Required Privilege Level | **10261**
- Release Information | **10261**

Syntax

```
mac-rewrite {
    interface interface-name {
        enable-all-ifl;
        protocol protocol-name;
        }
    }
}
```

Hierarchy Level

```
[edit protocols layer2-control]
```

Description

Enable rewriting of the MAC address for Layer 2 protocol tunneling (L2PT).

When a service provider edge port configured for L2PT receives a control packet for a supported protocol, the device rewrites the multicast destination MAC address with the predefined multicast tunneling MAC address 01:00:0c:cd:cd:d0 and sends the packet on the Q-in-Q tunnel interface. The packet travels across the provider network transparently to the other end of the tunnel, and the destination device restores the original multicast destination MAC address to deliver the packet at its destination.

Refer to "protocol" on page 13424 for the list of protocols that you can configure for L2PT on different devices.

To see the protocols for which you enabled L2PT on an interface, enter the "show mac-rewrite interface" on page 21993 command.

To see the MAC addresses installed by L2PT on products except MX Series routers, see "show ethernet-switching table" on page 19700. To see the MAC addresses installed by L2PT on MX Series routers, see "show bridge mac-table" on page 17264.

On MX Series and ACX Series routers and EX9200 switches with L2PT configured, customer-facing ports should not receive packets with the L2PT MAC address as the destination address unless there is a network tolopogy or configuration error. Any such interface receiving an L2PT packet becomes “Disabled”, and you must subsequently re-enable it using the "clear error mac-rewrite" on page 3568 command.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`enable-all-if` statement added in Junos OS Release 13.3.

Support for PVSTP protocol introduced in Junos OS Release 13.3 for MX Series routers and EX9200 switches.

Statement introduced in Junos OS Release 14.1X53-D10 and 17.3R1 for EX4300 switches.

RELATED DOCUMENTATION

Understanding Layer 2 Protocol Tunneling

Configuring Layer 2 Protocol Tunneling

show mac-rewrite interface | **21993**

clear error mac-rewrite | **3568**

mac (Static MAC-Based VLANs)

IN THIS SECTION

- Syntax | **10262**
- Hierarchy Level | **10262**
- Description | **10262**
- Options | **10262**
- Required Privilege Level | **10262**

Syntax

```
mac  mac-address {
        next-hop interface-name;
}
```

Hierarchy Level

```
[edit ethernet-switching-options static vlan vlan-name]
```

Description

Specify the MAC address to add to the Ethernet switching table.

The remaining statement is explained separately. See CLI Explorer.

Options

mac-address	MAC address

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

> *Adding a Static MAC Address Entry to the Ethernet Switching Table*

mac-statistics

IN THIS SECTION

- Syntax | 10263
- Hierarchy Level | 10263
- Description | 10264
- Default | 10264
- Required Privilege Level | 10264
- Release Information | 10264

Syntax

```
mac-statistics;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name bridge-options],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name switch-options],
[edit logical-systems logical-system-name switch-options],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options],
[edit routing-instances routing-instance-name switch-options],
[edit routing-instances routing-instance-name protocols evpn],
[edit switch-options],
[edit switch-options],
[edit vlans vlan-name switch-options]
```

Description

(MX Series routers, EX Series switches, and QFX Series only) For bridge domains or VLANs, enable MAC accounting either for a specific bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port.

Default

disabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for the `switch-options` statement added in Junos OS Release 9.2.

Support for top-level configuration for the `virtual-switch` type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

[edit switch-options] and [edit vlans *vlan-name* switch-options] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.

[edit switch-options] and [edit vlans *vlan-name* switch-options] hierarchy levels introduced in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

- Understanding Layer 2 Learning and Forwarding for Bridge Domains
- Layer 2 Learning and Forwarding for VLANs Overview
- Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports
- Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

mac-table-aging-time (MX Series)

IN THIS SECTION

- Syntax | 10265
- Hierarchy Level | 10265
- Description | 10266
- Options | 10266
- Required Privilege Level | 10266
- Release Information | 10266

Syntax

```
mac-table-aging-time time;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit routing-instances routing-instance-name protocols vpls]
[edit bridge-options],
[edit routing-instances routing-instance-name protocols evpn]
```

NOTE: For MX Series routers, the configuration statement is supported at the `[bridge-options]`, `[protocols vpls]`, and `[protocols evpn]` hierarchy levels only.

Description

Modify the timeout interval for the MAC table.

For MX Series routers, you can use the `global-mac-table-aging-time` statement at the `[edit protocols l2-learning]` hierarchy level to configure the timeout interval at the global level or use the `mac-table-aging-time` to configure the timeout interval for a bridge domain or for a specific VPLS or EVPN instance. If multiple timeout interval values are configured on a router, the router determines the timeout interval value in the following order of priority:

- Timeout interval configured at the VPLS or EVPN instance
- Timeout interval configured for the bridge domain
- Global timeout interval configured on the router

> **NOTE**: For MX Series routers, the timeout interval configuration feature is supported on routers with MPCs only.

Options

`time`—Specify the number of seconds to wait between MAC table clearings.

- **Range:** 10 through 1,000,000 seconds
- **Default:** 300 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the VPLS MAC Table Timeout Interval

Configuring the MAC Table Timeout Interval

mac-table-aging-time

IN THIS SECTION

- Syntax | 10267
- Hierarchy Level | 10267
- Description | 10267
- Default | 10268
- Options | 10268
- Required Privilege Level | 10268
- Release Information | 10268

Syntax

```
mac-table-aging-time (seconds | unlimited);
```

Hierarchy Level

```
[edit ethernet-switching-options],
[edit vlans vlan-name]
```

Description

You configure how long MAC addresses remain in the Ethernet switching table using the `mac-table-aging-time` statement in either the `[edit ethernet-switching-options]` or the`[edit vlans]` hierarchy, depending on whether you want to configure it for the entire device or only for specific VLANs.

If you specify the time as `unlimited`, entries are never removed from the table. Generally, use this setting only if the device or the VLAN has a fairly static number of end devices; otherwise the table will eventually fill up. You can use this setting to minimize traffic loss and flooding that might occur when traffic arrives for MAC addresses that have been removed from the table.

> **NOTE**: The `mac-table-aging-time` statement appears in the Junos OS CLI for devices that do *not* support the Enhanced Layer 2 Software (ELS) configuration style. If your device runs software that supports ELS, use the "global-mac-table-aging-time" on page 7719 statement, which appears in the `[edit protocols l2-learning]` hierarchy. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

Default

Entries remain in the Ethernet switching table for 300 seconds.

Options

seconds—Time that entries remain in the Ethernet switching table before being removed.

- **Range:** 60 through 1,000,000 seconds
- **Default:** 300 seconds

`unlimited`—Entries remain in the Ethernet switching table.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement updated in Junos OS Release 9.4 for EX Series switches to include `[edit ethernet-switching-options]` hierarchy level.

RELATED DOCUMENTATION

show ethernet-switching statistics aging | **19692**

Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch

Configuring MAC Table Aging on Switches

Controlling Authentication Session Timeouts (CLI Procedure)

mac-table-size (VPLS)

IN THIS SECTION

- Syntax | 10269
- Hierarchy Level | 10269
- Description | 10269
- Options | 10270
- Required Privilege Level | 10270
- Release Information | 10270

Syntax

```
mac-table-size size {
    packet-action drop;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit routing-instances routing-instance-name protocols evpn],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Specify the size of the MAC address table.

Options

size—Size of the MAC address table.

Range:

- - (M Series and T Series routers only) 16 through 65,536 MAC addresses
  - (MX Series routers only) 16 through 1,048,575 MAC addresses
  - (T4000 routers with Type 5 FPCs only) 16 through 262,143 MAC addresses

    > **NOTE**: Before modifying the size of the MAC address table (to 262,143 addresses), you must enable network services mode by including the `enhanced-mode` statement at the `[edit chassis network-services]` hierarchy level and then reboot the router.

- **Default:** 5120 MAC addresses

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

Configuring the Size of the VPLS MAC Address Table

Configuring Improved VPLS MAC Address Learning on T4000 Routers with Type 5 FPCs

enhanced-mode

evpn | **6541**

mac-table-size

IN THIS SECTION

- Syntax | 10271
- Hierarchy Level | 10271
- Description | 10272
- Options | 10272
- Required Privilege Level | 10272
- Release Information | 10272

Syntax

```
mac-table-size limit {
    packet-action drop;
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name bridge-options],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name switch-options],
[edit logical-systems logical-system-name switch-options],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options],
[edit routing-instances routing-instance-name switch-options],
[edit switch-options],
[edit switch-options],
[edit vlans vlan-name switch-options]
```

Description

Modify the size of the MAC address table for the bridge domain or VLAN, a set of bridge domains or VLANs associated with a trunk port, or a virtual switch. The default is 5120 MAC addresses.

> **NOTE**: For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the **mac-table-size** statement or changing the `mac-table-size` configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the **mac-table-size** statement or use the **commit at** configuration statement to commit the changes at the same time in both the peer nodes.
>
> Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers by using the **clear bridge mac-table** command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

Options

`limit`—Specify the maximum number of addresses in the MAC address table.

- **Range:** 16 through 1,048,575 MAC addresses
- **Default:** 5120 MAC addresses

  There is no default MAC address limit for the `mac-table-size` statement at the `[edit switch-options]` hierarchy level. The number of MAC addresses that can be learned is only limited by the platform, 65,535 MAC addresses for EX Series switches and 1,048,575 MAC addresses for other devices.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for the `switch-options` statement added in Junos OS Release 9.2.

Support for top-level configuration for the `virtual-switch` type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

[edit switch-options] and **[edit vlans *vlan-name* switch-options]** hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Support at the `[edit vlans vlan-name switch-options` hierarchy level introduced in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports

Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port

mac-validate

IN THIS SECTION

- Syntax | **10273**
- Hierarchy Level | **10274**
- Description | **10274**
- Options | **10274**
- Required Privilege Level | **10274**
- Release Information | **10274**

Syntax

```
mac-validate (loose | strict);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family]
```

Description

Enable IP and MAC address validation for static Ethernet and IP demux interfaces.

NOTE: When you configure `mac-validate`, irb interface does not allow traffic.

Options

`loose`—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not match the MAC address of the tuple. Continues to forward incoming packets when the source address of the incoming packet does not match any of the trusted IP addresses.

`strict`—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

- MAC Address Validation on Static Ethernet Interfaces Overview
- Configuring an IP Demultiplexing Interface
- Configuring a VLAN Demultiplexing Interface

mac-validate (Dynamic IP Demux Interface)

IN THIS SECTION

- Syntax | 10275
- Hierarchy Level | 10275
- Description | 10275
- Options | 10275
- Required Privilege Level | 10276
- Release Information | 10276

Syntax

```
mac-validate (loose | strict);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family inet]
```

Description

Enable IP and MAC address validation for dynamic IP demux interfaces in a dynamic profile.

Options

`loose`—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not match the MAC address of the tuple. Continues to forward incoming packets when the source address of the incoming packet does not match any of the trusted IP addresses.

`strict`—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the MAC address does not match the tuple's

MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring MAC Address Validation for Subscriber Interfaces

mac-vrf

IN THIS SECTION

- Syntax | 10276
- Hierarchy Level | 10277
- Description | 10277
- Required Privilege Level | 10277
- Release Information | 10278

Syntax

```
mac-vrf;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name instance-type],
[edit logical-systems name tenants name routing-instances name instance-type ],
[edit routing-instances name instance-type],
[edit tenants name routing-instances name instance-type]
```

Description

Specify a routing instance of type MAC-VRF. A MAC-VRF routing instance is a type of virtual-switch routing instance.

You can use MAC-VRF routing instances in Ethernet VPN (EVPN) environments to create MAC-based virtual routing and forwarding (VRF) instances. With MAC-VRF instances, you can manage multiple tenants with customer-specific VRF tables to isolate or group customer or tenant workloads.

In routing instances of type `mac-vrf`, configure basic parameters such as the following :

- `bridge-domain | vlans`—Specify the bridge domains or the VLANs associated with the MAC-VRF routing instance. You must specify VLANs in the MAC-VRF routing instance or specify a bridge domain that you configure elsewhere in the configuration. If you do not specify either parameter, the configuration editor displays the error message: "Warning: bridge-domains must be configured for mac-vrf instance".The device prevents you from committing the configuration.

- `protocol evpn`—Set the protocol to use in the MAC-VRF instance. In this case, you must select EVPN as the protocol.

- `service-type (vlan-aware | vlan-based | vlan-bundle)`—Select an EVPN service type for this MAC-VRF instance.

  See "service-type (MAC-VRF Routing Instances)" on page 16764 and *Understanding VLAN-Aware Bundle and VLAN-Based Service for EVPN* for more on these service types.

- `forwarding-instance identifier`—Map the MAC-VRF instance to a specific forwarding instance. Only some platforms support this option. See "forwarding-instance identifier" on page 7457 and *MAC-VRF Routing Instance Type Overview* for more on how MAC-VRF instances might map to forwarding instances.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

MAC-VRF Routing Instance Type Overview

macsec

IN THIS SECTION

- Syntax | 10278
- Syntax (SRX Series Firewalls) | 10279
- Hierarchy Level | 10280
- Description | 10280
- Options | 10280
- Required Privilege Level | 10281
- Release Information | 10281

Syntax

```
macsec {
    connectivity-association connectivity-association-name {
        exclude-protocol protocol-name;
        include-sci;
        mka {
            must-secure;
            key-server-priority priority-number;
            transmit-interval interval;
        }
        no-encryption;
        offset (0|30|50);
        pre-shared-key {
```

```
            cak hexadecimal-number;
            ckn hexadecimal-number;
        }
        replay-protect{
            replay-window-size number-of-packets;
        }
        secure-channel secure-channel-name {
            direction (inbound | outbound);
            encryption (MACsec);
            id {
                mac-address mac-address;
                port-id port-id-number;
            }
            offset (0|30|50);
            security-association security-association-number {
                key key-string;
            }
        }
        security-mode security-mode;
    }
    interfaces interface-name {
        connectivity-association connectivity-association-name;
    }
}
```

Syntax (SRX Series Firewalls)

```
macsec {
    cluster-control-port <idx> {
        connectivity-association connectivity-association-name;
    }
    cluster-data-port interface-name {
        connectivity-association connectivity-association-name;
    }
    connectivity-associationconnectivity-association-name {
        exclude-protocol protocol-name;
        include-sci;
        mka {
            key-server-priority priority-number;
            must-secure;
            transmit-interval milliseconds;
```

```
        }
        no-encryption;
        offset (0|30|50);
        pre-shared-key {
            cak hexadecimal-number;
            ckn hexadecimal-number;
        }
        replay-protect {
            replay-window-size number-of-packets;
        }
        security-mode security-mode;
    }
    traceoptions (Chassis Cluster){
        file {
            filename;
            files number;
            match regular-expression;
            (world-readable | no-world-readable);
            size maximum-file-size;
        }
        flag flag;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure Media Access Control Security (MACsec). Media Access Control Security(MACsec) is supported on control and fabric ports of SRX340, SRX345, and SRX4600 devices in chassis cluster mode to secure point-to-point Ethernet links between the peer devices in a cluster. Each point-to-point Ethernet link must be configured independently to secure using MACsec. You can enable MACsec encryption on device-to-device links using static connectivity association key (CAK) security mode.

Options

cluster-control-port <idx> Specify chassis cluster control interface on which MACsec is enabled.

- **Values:** 0.

cluster-data-port ***interface-name***	Specify chassis cluster fabric interface on which MACsec is enabled.
connectivity-association	Create or configure a MACsec connectivity association.
traceoptions	Define MACsec configuration tracing operations.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced in SRX Series Firewalls in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Configuring Media Access Control Security (MACsec)

Understanding Media Access Control Security (MACsec)

macsec (MX Series)

IN THIS SECTION

- Syntax | **10282**
- Hierarchy Level | **10283**
- Description | **10283**
- Options | **10283**

Required Privilege Level | 10283
Release Information | 10283

Syntax

```
macsec {
    connectivity-association connectivity-association-name {
        cipher-suite encryption-algorithm-name;
            exclude-protocol protocol-name;
        pre-shared-key-chain macsec-pre-shared-key-chain-name
        include-sci;
        mka {
            key-server-priority priority-number;
            must-secure;
            should-secure;
            transmit-interval interval;
        }
        no-encryption;
        offset (0|30|50);
        pre-shared-key {
            cak hexadecimal-number;
            ckn hexadecimal-number;
        }
        replay-protect{
            replay-window-size number-of-packets;
        }
        secure-channel secure-channel-name {
            direction (inbound | outbound);
            encryption ;
            id {
                mac-address mac-address;
                port-id port-id-number;
            }
            offset (0|30|50);
            security-association security-association-number {
                key key-string;
            }
        }
```

```
        security-mode security-mode;
    }
    interfaces interface-name {
        connectivity-association connectivity-association-name;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure Media Access Control Security (MACsec) on MX Series routers.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

maintain-subscriber (Subscriber Management)

IN THIS SECTION

- Syntax | 10284
- Hierarchy Level | 10284
- Description | 10284
- Required Privilege Level | 10285
- Release Information | 10285

Syntax

```
maintain-subscriber {
    interface-delete;
}
```

Hierarchy Level

```
[edit system services subscriber-management]
```

Description

Configure the router to maintain, rather than log out, DHCP relay and DHCP local server based subscribers when the specified type of event occurs.

For example, by default, the router logs out DHCP subscribers when an interface delete event occurs, such as a line card reboot or failure. You would specify the `interface-delete` option to ensure that the router maintains subscribers during line card reboots or failures. However, this option does not maintain subscribers during router reboots or failures.

This statement provides a global configuration for the router, which applies to all DHCP local server and DHCP relay clients in all routing instances.

> **NOTE**: The `maintain-subscriber` statement and `remove-when-no-subscribers` statement are mutually exclusive. You cannot specify that dynamically configured VLAN interfaces are removed when no subscribers exist when the router is also configured to maintain subscribers.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Maintaining Subscribers During Interface Delete Events

Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events

maintenance-association

IN THIS SECTION

- Syntax | 10286
- Hierarchy Level | 10286
- Description | 10286
- Options | 10287
- Required Privilege Level | 10287
- Release Information | 10287

Syntax

```
maintenance-association ma-name {
    short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
    protect-maintenance-association protect-ma-name;
    remote-maintenance-association remote-ma-name;
    continuity-check {
        hold-interval minutes;
        interval (10m | 10s | 1m | 1s| 100ms);
        loss-threshold number;
    }
    mep mep-id {
        auto-discovery;
        direction (up | down);
        interface interface-name (protect | working);
        lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect | rem-err-
xcon | xcon );
        priority number;
        remote-mep mep-id {
            action-profile profile-name;
            sla-iterator-profile profile-name {
                data-tlv-size size;
                iteration-count count-value;
                priority priority-value;
            }
        }
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]
```

Description

Configure the name of the maintenance association in IEEE-compliant format.

Options

`ma-name`—The name of the maintenance association within the maintenance domain.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Creating a Maintenance Association

Configuring a MEP to Generate and Respond to CFM Protocol Messages

maintenance-domain

IN THIS SECTION

- Syntax | 10288
- Hierarchy Level | 10289
- Description | 10289
- Options | 10289
- Required Privilege Level | 10289
- Release Information | 10289

Syntax

```
maintenance-domain domain-name {
    bridge-domain name <vlan-id [ vlan-ids ]>;
    instance vpls-instance-name;
    level number;
    maintenance-association ma-name {
        protect-maintenance-association  protect-ma-name;
        remote-maintenance-association  remote-ma-name;
        short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
        continuity-check {
            hold-interval minutes;
            interval (10m | 10s | 1m | 1s| 100ms);
            loss-threshold number
        }
        mep mep-id {
            auto-discovery;
            direction (up | down);
            interface interface-name (protect | working);
            lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect | rem-
err-xcon | xcon );
            priority number;
            remote-mep mep-id {
                action-profile profile-name;
                sla-iterator-profile profile-name {
                    data-tlv-size size;
                    iteration-count count-value;
                    priority priority-value;
                }
            }
        }
        mip-half-function(none | default | explicit);
        name-format (character-string | none | dns | mac+2oct);
    }
    virtual-switch name {
         bridge-domain name <vlan-id [ vlan-ids ]>;
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

Description

Configure the name of the maintenance domain in IEEE-compliant format.

> **NOTE**: For MX Series Routers, you can configure multiple down MEPs for a single instance of maintenance domain identifier and maintenance association name to monitor services provided on Virtual Private LAN Service (VPLS), bridge, circuit cross-connect (CCC), and IPv4 domains.

Options

domain-name—Name of the maintenance domain.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for multiple down MEP introduced in Junos OS Release 15.1R1 for MX Series Routers.

RELATED DOCUMENTATION

Creating a Maintenance Domain

Configuring a MEP to Generate and Respond to CFM Protocol Messages

major

IN THIS SECTION

- Syntax | 10290
- Hierarchy Level | 10290
- Description | 10290
- Required Privilege Level | 10291
- Release Information | 10291

Syntax

```
major {
    threshold threshold-value;
    action (chassis error) {
        alarm;
        disable-pfe;
        get-state;
        log;
        offline;
        reset;
    }
}
```

Hierarchy Level

[edit chassis "fpc" on page 7516 *slot-number* "error " on page 6441]

[edit chassis]

Description

Severity level of the error. An error that results in continuing loss of packet traffic but does not affect other modules is a major error. The severity level of an error cannot be configured by a user.

The other statements are explained separately.

Required Privilege Level

interface	To view this statement in the configuration.
interface-control	To add this statement to the configuration.

Release Information

Statement introduced for PTX Series routers in Junos OS Release 13.3.

Statement introduced for MX240, MX480, MX960, and MX2020 routers in Junos OS Release 14.2.

RELATED DOCUMENTATION

Fabric Resiliency

show chassis fpc errors | **18387**

Managing Errors

major-event (Process)

IN THIS SECTION

- Syntax | **10291**
- Hierarchy Level | **10292**
- Description | **10292**
- Options | **10292**
- Required Privilege Level | **10292**
- Release Information | **10292**

Syntax

```
major-event {
    threshold process name memory-limit value ;
```

```
    action [snmp-trap | live-core | (terminate-process| terminate-process-force)];
}
```

Hierarchy Level

```
[edit system monitor memory process]
```

Description

Indicates an event that, if left unchecked, may result in significant loss of service or interruption of traffic. You can specify the process threshold value for specific processes and their memory limit. When the threshold is crossed, you can configure the action to be performed. For instance, you can specify the threshold for the `rpd` process and the memory limit for that process which you would classify as a a major event.

Options

threshold Specifies the process threshold level for a major event. You must specify the process name and the memory limit for the process.

memory Specifies the permissible memory limit for the process after which the major event is triggered.

action Specifies the action to be performed when the threshold is crossed. For major events, you can raise an SNMP trap or terminate the process. You can also display a list of system core files created when the system failed which you can use to debug the issue.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1.

RELATED DOCUMENTATION

System Log Overview

System Log Explorer

major-event (System)

IN THIS SECTION

- Syntax | 10293
- Hierarchy Level | 10293
- Description | 10293
- Options | 10294
- Required Privilege Level | 10294
- Release Information | 10294

Syntax

```
major-event {
    threshold swap-usage value ;
    action [snmp-trap | live-vmcore];
}
```

Hierarchy Level

```
[edit system monitor memory system]
```

Description

Indicates an event that, if left unchecked, may result in significant loss of service or interruption of traffic. At the system level you can monitor the overall free pages (cache and free pages), kernel memory allocations and compare them with threshold values. You can also monitor the swap usage and specify the threshold for swap space usage. Most devices have a swap partition that you can configure to

provide additional transient memory. Ideally, the swap partition is used only under exceptional circumstances. In current dynamic systems, you can sporadically use the swap partition as long as the usage is not high or persistent. For instance, you could categorise a major event as one where the swap usage is greater than 50%. If the threshold value is crossed, then you can raise an SNMP trap or display a list of system core files created when the system failed which can be used to debug the issue.

You can only configure swap usage threshold for memory monitoring. If you do not configure the swap usage threshold, memory monitoring for swap usage is based on the default threshold. Memory monitoring for other features is based on the default threshold.

Options

threshold Specifies the process threshold level for a major event. You must specify the process name and the memory limit for the process.

swap-usage Specifies the swap usage in percentage.

action Specifies the action to be performed when the threshold is crossed. You can either

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1.

RELATED DOCUMENTATION

System Log Overview

System Log Explorer

major-ring-name

IN THIS SECTION

- Syntax | 10295
- Hierarchy Level | 10295
- Description | 10295
- Required Privilege Level | 10295
- Release Information | 10295

Syntax

```
major-ring-name name;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Specify the name of major ring to which the sub-ring node is interconnected.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

make-before-break (LDP)

IN THIS SECTION

- Syntax | 10296
- Hierarchy Level | 10296
- Description | 10296
- Options | 10297
- Required Privilege Level | 10297
- Release Information | 10297

Syntax

```
make-before-break {
    timeout seconds;
    switchover-delay seconds;
}
```

Hierarchy Level

```
[edit protocols ldp]
```

Description

Configures make before break (MBB) for multicast LDP (MLDP) link protection to ensure minimum packet loss when attempting to signal a new label-switched path (LSP) before tearing down the old LSP path.

Options

timeout *seconds*

Specify a value to change a make -before-break timeout for point-to-multipoint LSPs. Even if an MBB acknowledgment is not received for a point-to-multipoint LSP before the specified timeout period expires, the label-switching router (LSR) performs an MBB switchover from the old LSR to the new upstream LSR.

- **Range:** 1 through 300 seconds
- **Default:** 30 seconds

switchover-delay *seconds*

Specify a value to change switchover delay for a point-to-multipoint LSP from the old LSR to the new upstream LSR. If an MBB acknowledgment is received on a point of local repair (PLR) router, the PLR waits for the specified seconds to switch its upstream LSR from the old LSR to the new LSR.

- **Range:** 1 through 300 seconds
- **Default:** 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring LDP Link Protection

managed

IN THIS SECTION

- Syntax | 10298
- Hierarchy Level | 10298
- Description | 10299
- Options | 10299
- Required Privilege Level | 10300
- Release Information | 10300

Syntax

```
managed {
    client-limit limit;
    client-list list-name {
        address address <routing-instance [instance-name...]>;
    }
    control-inactivity-timeout seconds;
    control-maximum-duration seconds;
    control-per-client-limit number;
    offload-type (none | inline-timestamp | pfe-timestamp);
    port number;
    stateful-sequence;
    test-per-client-limit limit;
}
```

Hierarchy Level

```
[edit services monitoring twamp server]
```

Description

Enable the Two-Way Active Measurement Protocol (TWAMP) server for managed control on the TCP port that reflects the test packets.

Options

client-limit *limit*

Specify the maximum number of TWAMP clients.

- **Range:** 0 to 1000
- **Default:** 0 (disabled, which means there is no limit to the number of clients)

client-list *list-name* { address *address* <routing-instance [*instance-name*...]>}

Specify a list of clients and their corresponding addresses and routing instances. If you do not specify a routing instance, the default routing table is used (inet.0).

- **Syntax:** address *address* <routing-instance [*instance-name*...]>—Specify one address statement for each client, configuring an IPv4 address and, optionally, a list of routing instances for each one, to filter packets accordingly.

control-inactivity-timeout *seconds*

Specify the number of seconds to wait after the control connection becomes inactive before timing out.

- **Range:** 0 through 86400 seconds; specify 0 to disable.
- **Default:** 900 seconds

control-maximum-duration *seconds*

Specify how long the control connection remains up.

- **Range:** 0 through 86400 seconds
- **Default:** 0 (disabled, which means there is no limit to the length of time the control connection remains up)

control-per-client-limit *number*

Specify the maximum number of control connections allowed per client.

- **Range:** 0 to 1000
- **Default:** 0 (disabled, which means there is no limit to the number of control connections allowed)

offload-type (none | inline-timestamp | pfe-timestamp)

Enable timestamping of TWAMP probe messages in the Routing Engine, inline in the hardware, or Packet Forwarding Engine host processor. Prior to Junos OS Evolved 22.3R1, for IPv6 traffic, you must configure `none`, as only Routing Engine timestamping is supported. Starting in Junos OS Evolved 22.3R1, you can configure the `pfe-timestamping` option for IPv6 traffic.

- **none**—Timestamping in the Routing Engine.
- **inline-timestamping**—Timestamping inline, in the hardware at the generator, and reflection in the hardware at the reflector. When this option is configured for an ACX router, the ACX can act as either a server or client, but not both. Also for an ACX router, `inline-timestamping` is not supported on PM50 ports if FEC74 is enabled on that port. You can configure the `inline-timestamping` option for up to 256 test sessions.
- **pfe-timestamping**—Timestamping in the Packet Forwarding Engine.
- **Default:** pfe-timestamp

port *number*

Specify which port can accept TWAMP control connections.

- **Range:** You can specify port 862, or any port from 49152 through 65535.
- **Default:** 862 (IANA port for TWAMP)

stateful-sequence

Configure stateful sequencing to generate the sequence number in the reflected payload independently, instead of copying the sequence number from the client frame. With independently generated sequence numbers, the software is able to detect in which direction drops have occurred. Unidirectional drop values are inferred minimum values, and are not guaranteed to be exact. The software relies on an actual packet to be reflected successfully to be able to compute differences. If such a packet does not occur, then the software is unable to tell whether the drop happened in the egress or ingress direction.

test-per-client-limit *limit*

Specify the maximum number of test sessions allowed per client.

- **Range:** 0 to 1000
- **Default:** 0 (disabled, which means there is no limit to the number of test sessions allowed per client)

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 20.3R1.

`offload-type` and `stateful-sequence` options introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

managed-configuration

IN THIS SECTION

- Syntax | 10301
- Hierarchy Level | 10301
- Description | 10301
- Default | 10302
- Required Privilege Level | 10302
- Release Information | 10302

Syntax

```
(managed-configuration | no-managed-configuration);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Description

Specify whether to enable the host to use a stateful autoconfiguration protocol for address autoconfiguration, along with any stateless autoconfiguration already configured:

- **managed-configuration**—Enable host to use stateful autoconfiguration.
- **no-managed-configuration**—Disable host from using stateful autoconfiguration.

You can set two fields in the router advertisement message to enable stateful autoconfiguration on a host: the managed configuration field and the other stateful configuration field. Setting the managed configuration field enables the host to use a stateful autoconfiguration protocol for address autoconfiguration, along with any stateless autoconfiguration already configured. Setting the other stateful configuration field enables autoconfiguration of other nonaddress-related information.

Default

Stateful autoconfiguration is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

other-stateful-configuration | 11953

managed-configuration (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 10303
- Hierarchy Level | 10303
- Description | 10303
- Default | 10303

Required Privilege Level | 10303
Release Information | 10303

Syntax

```
(managed-configuration | no-managed-configuration);
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Specify whether to enable the dynamic host to use a stateful autoconfiguration protocol for address autoconfiguration, along with any stateless autoconfiguration already configured:

- `managed-configuration`—Enable host to use stateful autoconfiguration.
- `no-managed-configuration`—Disable host from using stateful autoconfiguration.

Default

The configured object is disabled unless explicitly enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA

Dynamic Router Advertisement Configuration Overview

managed-services

IN THIS SECTION

- Syntax | 10304
- Hierarchy Level | 10304
- Description | 10304
- Options | 10305
- Required Privilege Level | 10305
- Release Information | 10305

Syntax

```
managed-services name;
```

Hierarchy Level

```
[edit chassis high-availability services-redundancy-group]
```

Description

Enable set of services for the specified services redundancy group (SRG).

You can selectively and flexibly associate a service such as IPsec VPN to one of multiple SRGs configured on SRX Series firewall in a Multinode High Availability setup.

Options

name

- ipsec—IPsec VPN Service

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1.

management

IN THIS SECTION

- Syntax | 10305
- Hierarchy Level | 10306
- Description | 10306
- Default | 10306
- Options | 10306
- Required Privilege Level | 10307
- Release Information | 10307

Syntax

```
management (class | static | high-power);
```

Hierarchy Level

```
[edit poe],
[edit poe (all | fpc slot-number)]
```

Description

Designate how the PoE controller allocates power to the PoE interfaces.

Default

class

Options

- `class`—The amount of power allocated to the interface is based on the class of the connected powered device. If LLDP power negotiation is enabled, the powered device can request more or less power. LLDP power negotiation is enabled by default in class management mode.

  If LLDP power negotiation is disabled, the power allocation is determined solely by the class of the connected powered device. If there is no powered device connected, standard 15.4W power is allocated on the interface. For more information about disabling LLDP power negotiation, see Device Discovery Using LLDP and LLDP-MED on Switches.

- `static`—The amount of power allocated to the interface is determined by the value of the `maximum-power` statement, not the class of the connected powered device. This amount is allocated even when a powered device is not connected to the interface, ensuring that power is available when needed.

  NOTE: Static mode is not supported in PoE-bt.

- `high-power`—(ACX2000 routers only) ACX2000 PoE interfaces support power delivery of up to 65 W per port using all four pairs of Ethernet RJ45 cables. Traditional PoE ports use only two pairs of Ethernet cable for power delivery. According to the IEEE 802.3af standard, each port can deliver a maximum power of up to 32 W. With `high-power` mode of power delivery over all four pairs, the power sourcing equipment (PSE) has an option to deliver up to 65 W per port, provided the powered devices request this high power over all four pairs of the Ethernet cable. By default, `high-power` mode is not enabled and has to be explicitly enabled. When the PoE controller is configured for `high-power` mode, the PoE controller does not deliver power to normal powered devices that request power over two pairs.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 17.2R1 for a Junos Fusion Provider Edge.

RELATED DOCUMENTATION

- Understanding PoE on EX Series Switches
- Configuring PoE Interfaces on EX Series Switches
- *Configuring Power over Ethernet in a Junos Fusion*

management-ethernet (Alarm)

IN THIS SECTION

- Syntax | 10308
- Hierarchy Level | 10308
- Description | 10308
- Options | 10308
- Required Privilege Level | 10308
- Release Information | 10308

Syntax

```
management-ethernet {
    link-down  (red | yellow | ignore);
}
```

Hierarchy Level

```
[edit chassis alarm],
[edit chassis interconnect-device name alarm],
[edit chassis node-group name alarm]
```

Description

Configure alarms for a management Ethernet interface.

> **NOTE**: If you configure a yellow alarm on the Interconnect device, it will be handled as a red alarm.

Options

The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

management-instance

IN THIS SECTION

- Syntax | 10309
- Hierarchy Level | 10309
- Description | 10309
- Options | 10310
- Required Privilege Level | 10310
- Release Information | 10310

Syntax

```
management-instance;
```

Hierarchy Level

```
[edit system]
```

Description

Enable a dedicated management virtual routing and forwarding (VRF) instance. The name of the dedicated management instance is reserved and hardcoded as `mgmt_junos`; you cannot configure any other routing instance by the name `mgmt_junos`.

The management Ethernet interface (usually named `fxp0` or `em0`) provides the out-of-band management network for the router. There is no clear separation between either out-of-band management traffic and in-band protocol control traffic, or user traffic at the routing-instance or routing table level. The `management-instance` configuration statement confines the management interface in the dedicated management instance `mgmt_junos`, and it enables an administrative routing table dedicated to management tasks for the network device.

For more information about configuring `management-instance` and the `mgmt_junos` routing instance, see *Management Interface in a Non-Default Instance*.

Options

This command has no options.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

management (PoE)

IN THIS SECTION

- Syntax | 10310
- Hierarchy Level | 10311
- Description | 10311
- Default | 10311
- Options | 10311
- Required Privilege Level | 10311
- Release Information | 10311

Syntax

```
management (class | static);
```

Hierarchy Level

```
[edit poe]
```

Description

Designates how the SRX Series Firewall allocates power to the PoE ports.

Default

static

Options

- `static`—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power configured for the port.
- `class`—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power for the class as defined by the IEEE 802.3 AF standard.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

management (Security Zones)

IN THIS SECTION

- Syntax | 10312
- Hierarchy Level | 10312

- Description | 10313
- Options | 10313
- Required Privilege Level | 10313
- Release Information | 10313

Syntax

```
management {
    description text;
    host-inbound-traffic {
        protocols protocol-name {
                            except;
        }
        system-services service-name {
                            except;
        }
    }
    interfaces interface-name  {
        host-inbound-traffic {
            protocols protocol-name {
                                except;
            }
            system-services service-name {
                                except;
            }
        }
    }
    screen screen-name;
}
```

Hierarchy Level

```
[edit security zones functional-zone]
```

Description

Specify the host for out-of-band management interfaces. You can set firewall options in this zone to protect the management interface from different types of attacks. Because this zone cannot be specified in policies, traffic entering from this zone can only be traffic originating from the device itself and cannot originate from any other zone.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.The `description` option added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Security Zones Overview

Understanding Functional Zones

mandatory

IN THIS SECTION

- Syntax | 10314
- Hierarchy Level | 10314
- Description | 10314
- Required Privilege Level | 10314
- Release Information | 10314

Syntax

```
mandatory;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name variables variable-name]
```

Description

Specify that the external server (for example, RADIUS) must return a value for a user-defined variable. If the external server does notreturn a value for the variable, the dynamic profile fails.

> **NOTE**: When a dynamic profile has mandatory and non-mandatory variables, configure mandatory variables first in the profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

User-Defined Variables

Using Variable Expressions in User-Defined Variables

Configuring User-Defined Dynamic Variables in Dynamic Profiles

manual

IN THIS SECTION

- Syntax | 10315
- Hierarchy Level | 10315
- Description | 10316
- Required Privilege Level | 10316
- Release Information | 10316

Syntax

```
manual {
    direction (inbound | outbound | bidirectional) {
        authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96);
            key (ascii-text key | hexadecimal key);
        }
        auxiliary-spi spi-value;
        encryption {
            algorithm algorithm;
            key (ascii-text key | hexadecimal key);
        }
        spi spi-value;
        protocol (ah | esp | bundle);
    }
}
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Define a manual IPsec SA.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Security Associations

manual (Junos-FIPS Software)

IN THIS SECTION

- Syntax | 10317
- Hierarchy Level | 10317
- Description | 10317
- Options | 10317
- Required Privilege Level | 10317
- Release Information | 10318

Syntax

```
manual {
    direction (bidirectional | inbound | outbound) {
        protocol esp;
        spi spi-value;
        encryption {
            algorithm 3des-cbc;
            key ascii-text ascii-text-string;
        }
        auxiliary-spi auxiliary-spi-value;
        encryption {
            algorithm 3des-cbc;
            key (ascii-text key | hexadecimal key);
        }
        protocol (esp | bundle);
        spi spi-value;
    }
}
```

Hierarchy Level

[edit security ipsec internal security-association]

Description

Define a manual security association (SA) for internal Routing Engine-to-Routing Engine communication.

Options

The remaining statements are explained separately.

Required Privilege Level

Crypto Officer—To view and add this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring IPsec for Enabling Internal Communications Between Routing Engines for Junos OS in FIPS Mode*

manual (Junos OS)

IN THIS SECTION

- Syntax | 10318
- Hierarchy Level | 10319
- Description | 10319
- Options | 10319
- Required Privilege Level | 10319
- Release Information | 10319

Syntax

```
manual {
    direction (inbound | outbound | bi-directional) {
        authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96);
            key (ascii-text key | hexadecimal key);
        }
        auxiliary-spi auxiliary-spi-value;
        encryption {
            algorithm (des-cbc | 3des-cbc);
            key (ascii-text key | hexadecimal key);
        }
        protocol (ah | esp | bundle);
```

```
        spi spi-value;
    }
}
```

Hierarchy Level

```
[edit security ipsec security-association]
```

Description

Define a manual IPsec SA.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring Manual IPsec Security Associations for an ES PIC*

manual (Security IPsec)

IN THIS SECTION

- Syntax | 10320
- Hierarchy Level | 10320
- Description | 10321
- Options | 10321
- Required Privilege Level | 10322
- Release Information | 10322

Syntax

```
manual {
    authentication {
        algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
        key (ascii-text  key | hexadecimal  key );
    }
    encryption {
        algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-256-cbc | aes-256-
gcm | des-cbc);
        key (ascii-text key  | hexadecimal  key );
    }
    external-interface external-interface-name;
    gateway ip-address;
    protocol (ah | esp);
    spi spi-value;
}
```

Hierarchy Level

```
[edit security ipsec vpn vpn-name]
```

Description

Define a manual IPsec security association (SA).

Options

authentication algorithm

Hash algorithm that authenticates packet data. It can be one of the following

- `hmac-md5-96`—Produces a 128-bit digest.
- `hmac-sha-256-128`—Provides data origin authentication and integrity protection. This version of the hmac-sha-256 authenticator produces a 256-bit digest and specifies truncation to 128 bits.
- `hmac-sha1-96`—Hash algorithm that authenticates packet data. It produces a 160-bit digest. Only 96 bits are used for authentication.
- `authentication key`—Type of authentication key. It can be one of the following:
  - `ascii-text` *key*—ASCII text key. For `hmac-md5-96`, the key is 16 ASCII characters; for `hmac-sha1-96`, the key is 20 ASCII characters.
  - `hexadecimal` *key*—Hexadecimal key. For `hmac-md5-96`, the key is 32 hexadecimal characters; for `hmac-sha1-96`, the key is 40 hexadecimal characters.

encryption algorithm

Select the encryption algorithm for the internal Routing-Engine-to-Routing-Engine IPsec security association (SA) configuration. It can be one of the following:

- `des-cbc`—Encryption algorithm with block size of 8 bytes (64 bits) and key size 48 bits.
- `3des-cbc`—Encryption algorithm with block size of 8 bytes (64 bits) and key size of 192 bits.

  For `3des-cbc`, we recommend that the first 8 bytes be different from the second 8 bytes, and the second 8 bytes be the same as the third 8 bytes.
- `aes-128-cbc`—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- `aes-128-gcm`—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- `aes-192-cbc`—Advanced Encryption Standard (AES) 192-bit encryption algorithm.
- `aes-256-cbc`—Advanced Encryption Standard (AES) 256-bit encryption algorithm.

- `aes-256-gcm`—Advanced Encryption Standard (AES) 256-bit encryption algorithm.
- `encryption key`—Type of encryption key. It can be one of the following:
  - `ascii-text key`—ASCII text key. For the `des-cbc` option, the key contains 8 ASCII characters; for `3des-cbc`, the key contains 24 ASCII characters.
  - `hexadecimal key`—Hexadecimal key. For the `des-cbc` option, the key contains 16 hexadecimal characters; for the `3des-cbc` option, the key contains 48 hexadecimal characters.

external-interface Specify the outgoing interface for the manual security association

gateway For a manual security association, specify the IPv4 or IPv6 address of the peer

protocol Define an IPsec protocol for the manual security association

- Values:
  - ah—Authentication Header protocol
  - esp—ESP protocol (To use the ESP protocol, you must also use the tunnel statement at the [edit security ipsec security-association sa-name mode] hierarchy level)

spi Configure a security parameter index (SPI) for a security association (SA). An arbitrary value that uniquely identifies which security association (SA) to use at the receiving host (the destination address in the packet).

- **Range:** 256 through 16,639

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1.

Support for `hmac-sha-256-128` added to SRX5400, SRX5600, and SRX5800 devices in Junos OS Release 12.1X46-D20. Support for authentication algorithms (SHA1: hmac-sha1-96 and SHA2: hmac-sha-256-128) in PowerMode IPsec (PMI) mode is introduced for SRX4100, SRX4200, and vSRX Virtual

Firewall in Junos OS Release 19.3R1. Support for vSRX Virtual Firewall 3.0 is introduced in Junos OS Release 20.1R1.

Support for cipher algorithms aes-128-cbc, aes-192-cbc, and aes-256-cbc in PowerMode IPsec (PMI) mode is introduced for SRX4100, SRX4200, and vSRX Virtual Firewall in Junos OS Release 19.3R1. Support for vSRX Virtual Firewall 3.0 is introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

IPsec Overview

vpn (Security) | **30946**

manual switch

IN THIS SECTION

- Syntax | **10323**
- Hierarchy Level | **10323**
- Description | **10324**
- Required Privilege Level | **10324**

Syntax

```
request protection-group ethernet-aps manual-switch md <md> ma <ma>
```

Hierarchy Level

```
[edit protocols protection-group ethernet-aps]
```

Description

Forces traffic to switch from the active path to the alternate path, even in the absence of a failure on the working path. If the working path is the active path, traffic will be switched to the protection path. If the protection path is the active path, traffic will be switched to the protection path.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Ethernet Automatic Protection Switching Overview

many-to-one (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 10324
- Hierarchy Level | 10325
- Description | 10325
- Options | 10325
- Required Privilege Level | 10325
- Release Information | 10325

Syntax

```
many-to-one {
    preferred-backup preferred-backup;
}
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options high-availability-options]
```

Description

Configure the many-to-one (N:1) preferred backup for the aggregated multiservices (AMS) interface.

> **NOTE**: The preferred backup must be one of the member interfaces (mams–) that have already been configured at the [`edit interfaces interface-name load-balancing-options`] hierarchy level. Even in the case of mobile control plane redundancy, which is one-to-one (1:1), the initial preferred backup is configured at this hierarchy level.

Options

`preferred-backup` *`preferred-backup`*
: Use the specified interface as the preferred backup member interface. The member interface format is mams-*a*/*b*/0, where *a* is the FPC slot number and *b* is the PIC slot number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

high-availability-options (Aggregated Multiservices) | **8068**

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

map-dest-bmac-to-dest-cmac

IN THIS SECTION

- Syntax | 10326
- Hierarchy Level | 10326
- Description | 10326
- Required Privilege Level | 10326
- Release Information | 10326

Syntax

```
map-dest-bmac-to-dest-cmac <b-mac-address> <c-mac-address>;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name service-groups service-group-name]
```

Description

Specify the destination customer MAC (C-MAC) address to Provider Backbone Bridging (PBB) MAC (B-MAC) address mapping. The address mapping prevents unnecessary flooding of packets towards the core.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

routing-instances (Multiple Routing Entities) | 15976

map (Domain Map)

IN THIS SECTION

- Syntax | 10327
- Hierarchy Level | 10328
- Description | 10328
- Options | 10328
- Required Privilege Level | 10329
- Release Information | 10329

Syntax

```
map domain-map-name {
    aaa-logical-system logical-system-name {
        aaa-routing-instance routing-instance-name;
    }
    aaa-routing-instance routing-instance-name;
    access-profile profile-name;
    address-pool pool-name;
    dynamic-profile profile-name;
    strip-domain;
    strip-username (left-to-right | right-to-left);
    sub-domain name {
   (
   aaa-logical-system name {aaa-routing-instance (default | name)
   } | aaa-routing-instance (default | name));
   (
   target-logical-system name {target-routing-instance (default | name)
   } | target-routing-instance (default | name));
   access-profile access-profile;
   address-pool address-pool;
```

```
    dynamic-profile dynamic-profile;
    override-chap-password override-chap-password;
    override-password override-password;
    qualifier {
     vlan-id-list [ vlan-id-list ... ];
    }
    strip-domain;
    strip-username (left-to-right | right-to-left);
    tunnel-profile tunnel-profile;
    using-user-password;
   }
     override-password password;
     target-logical-system logical-system-name {
         target-routing-instance routing-instance-name;
     }
     target-routing-instance routing-instance-name;
     tunnel-profile profile-name;
     tunnel-switch-profile profile-name;
}
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the domain map to use to map options and parameters to subscriber sessions based on the subscriber domain.

Options

domain-map-name—Name of the domain map. The name is the same as the subscriber domain to which it will apply. For example, for the username `user1@example.com`, the domain map name is `example.com`.

- `*` —Use the asterisk wildcard character in the *domain-map-name* to specify a wildcard domain map, which enables mapping based on a partial match (for example, `xyz*northern.example.com`).The router performs the wildcard lookup when there is no exact match for the subscriber domain name. The wildcard can appear anywhere within the domain name string, and can match zero or more characters. The asterisk is the only wildcard character, and only one wildcard is supported in a domain map name. If

you include multiple asterisks, the first asterisk is treated as the wildcard character and the others are treated as non-wildcard characters.

- `default`—Use a domain map name of `default` to specify the domain map that the router uses when there is no exact or wildcard match for the domain or realm name in the subscriber username.
- `none`—Use a domain map name of `none` to specify the domain map the router uses when a subscriber username does not have a domain or realm name.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`override-password` and `strip-username` options introduced in Junos OS Release 15.1.

wildcard character introduced in Junos OS Release 16.1.

`sub-domain` option introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

Configuring a Domain Map

map-e

IN THIS SECTION

- Syntax | 10330
- Hierarchy Level | 10330
- Description | 10330

- Options | 10331
- Required Privilege Level | 10332
- Release Information | 10332

Syntax

```
map-e name {
    confidentiality;
    disable-auto-route;
    ea-bits-len ea-bits-len;
    ipv4-prefix ipv4-prefix;
    mape-prefix mape-prefix;
    mtu-v6 mtu-v6;
    psid-length psid-length;
    psid-offset psid-offset;
    softwire-address softwire-address;
    v4-partial-reassembly
    v4-reassembly;
    v6-reassembly;
    version-03;
}
```

Hierarchy Level

```
[edit services softwire softwire-concentrator],
[edit services softwires softwire-types],
[edit security softwires]
```

Description

Configure Mapping of Address and port – Encapsulation (MAP-E) as an inline service on devices that use MPC and MIC interfaces. MAP-E is an automatic tunneling mechanism that encapsulates IPv4 packets within an IPv6 address. The IPv4 packets are carried in an IPV4-over-IPV6 tunnel from the MAP-E Customer Edge (CE) devices to the MAP-E Provider Edge (PE) devices (also called as Border Relay (BR) devices) through an IPV6 routing topology, where they are de-tunneled for further processing.

Options

confidentiality

Configure Junos MAP-E confidentiality. This helps to hide MAP-E rule parameters in CLI show commands and logs.

disable-auto-route

Disable auto-routes and enable static routes to facilitate ECMP load balancing.

> **NOTE**: When you enable the `disable-auto-route` option, you must configure static routes.

name

Name of the MAP-E softwire concentrator.

ea-bits-len

Configure rule for Embedded Address (EA) length for the MAP-E domain.

> **NOTE**:
>
> - If `v4-prefix-len` is 0 then `ea-bits-len` must be non-zero, and vice versa.
> - It is possible that `ea-bits-len` is equal to 0, but `psid-len` is non-zero.
> - If the sum of `v4-prefix-len` and `ea-bits-len` is less than 32, then the `psid-len` must be equal to the difference between 32 and the sum total of `v4-prefix-len` and `ea-bits-len`.

- **Range:** 0 through 48

ipv4-prefix

Configure rule for IPv4 prefix and length of the MAP-E domain.

- **Range:** 0 through 32

mape-prefix

Configure rule for IPV6 prefix and length for the MAP-E domain. The MAP-E IPv4 and IPv6 prefix must be unique per softwire concentrator.

mtu-v6

(Optional) Specify the Maximum transmission unit (MTU) for the MAP-E softwire tunnel.

- **Default:** 9192
- **Range:** 1280 through 9192

psid-length

Configure Port Set ID (PSID) length value for the MAP-E domain.

> **NOTE:**
>
> - If the sum of `v4-prefix-len` and `ea-bits-len` is less than 32, then the `psid-len` must be equal to the difference between 32 and the sum total of `v4-prefix-len` and `ea-bits-len`.

- **Range:** 0 through 16

psid-offset (Optional) Configure PSID offset value for the MAP-E domain.

- **Default:** 4
- **Range:** 0 through 16

softwire-address Specify the Border Relay device unicast IPv6 address as the softwire concentrator IPV6 address.

v4-partial-reassembly (Optional) Enable IPv4 partial reassembly for MAP-E.

v4-reassembly | v6-reassembly (Optional) Enable IPv4 and IPv6 reassembly for MAP-E.

version-03 (Optional) Configure version number to distinguish between currently supported version of the Internet draft draft-ietf-softwire-map-03 (expires on July 28, 2013), *Mapping of Address and Port with Encapsulation (MAP)* and the latest available version.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.2R1.

map-entry-timeout

IN THIS SECTION

- Syntax | 10333
- Hierarchy Level | 10333
- Description | 10333
- Options | 10333
- Required Privilege Level | 10334
- Release Information | 10334

Syntax

```
map-entry-timeout map-entry-timeout;
```

Hierarchy Level

```
[edit security alg msrpc]
[edit security alg sunrpc]
[edit tenants tenant-name security alg msrpc]
[edit tenants tenant-name security alg sunrpc]
```

Description

Configure the mapping entry timeout value. When the incoming traffic hits the mapping entry, the timeout value has been reset to configured value. The mapping entry is removed from the table when the timeout value expires. The lifetime of the mapping entry is global and applies to all entries in the table.

Options

map-entry-timeout	Specify the Microsoft remote procedure call Application Layer Gateway (MS-RPC ALG) or Sun RPC ALG mapping entry timeout value in hours.

- **Range:** 1 through 72 hours.
- **Default:** 32 hours.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D10.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

msrpc | 10999

sunrpc | 28722

map-t

IN THIS SECTION

- Syntax | 10335
- Hierarchy Level | 10335
- Description | 10335
- Options | 10335
- Required Privilege Level | 10336
- Release Information | 10336

Syntax

```
map-t name {
 dmr-prefix dmr-prefix;
 ea-bits-len ea-bits-len;
 ipv4-prefix ipv4-prefix;
 mapt-prefix mapt-prefix;
 mtu-v6 mtu-v6;
 psid-length psid-length;
 psid-offset psid-offset;
 rule rule-name{
 v4-reassembly;
 v6-reassembly;
 }
}
```

Hierarchy Level

```
[edit services softwire softwire-concentrator]
```

Description

Configure Mapping of Address and Port using Translation (MAP-T) as an inline service on MX Series routers that use MPC and MIC interfaces. MAP-T is double stateless NAT64-based solution. The solution uses IPv4-IPv6 translation as the form of IPv6 domain transport.

Options

name MAP-T softwire concentrator name

dmr-prefix 64-bit IPv6 prefix advertised by BR used to form IPv4-embedded IPv6 address

ea-bits-len MAP-T domain's rule EA (Embedded Address) length

- **Range:** 0 through 48

ipv4-prefix IPv4 prefix assigned by service provider for a MAP rule

mapt-prefix IPv6 prefix assigned by service provider for a MAP rule

mtu-v6 MTU for the MAP-T softwire tunnel

- **Default:** 9192
- **Range:** 1280 through 9192

psid-length MAP-T domain's port set length

- **Range:** 0 through 16

psid-offset MAP-T domain's port set offset

- **Default:** 4
- **Range:** 0 through 16

v4-reassembly (Optional) Enable IPv4 full reassembly for MAP-T.

v6-reassembly (Optional) Enable IPv6 full reassembly for MAP-T.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1

mapped-port (Security Static NAT)

IN THIS SECTION

- Syntax | 10337
- Hierarchy Level | 10337

- Description | 10337
- Options | 10337
- Required Privilege Level | 10337
- Release Information | 10337

Syntax

```
mapped-port lower-port-range to upper-port-range;
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name then static-nat prefix]
[edit security nat static rule-set rule-set-name rule rule-name then static-nat prefix-name]
```

Description

Specify a destination port or port range to allow static NAT to map ports.

Options

- `lower-port-range`—Specify the lower limit of the port range.
- `upper-port-range`—Specify the upper limit of the port range.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

mapping

IN THIS SECTION

- Syntax (EX Series) | 10338
- Syntax (QFX Series) | 10338
- Hierarchy Level | 10338
- Description | 10339
- Options | 10339
- Required Privilege Level | 10340
- Release Information | 10340

Syntax (EX Series)

```
mapping (native (push | swap) | policy | tag (push | swap));
```

Syntax (QFX Series)

```
mapping (native (push | swap) | tag (push | swap));
mapping native inner-tag tag push;
mapping native push inner-tag tag;
```

Hierarchy Level

```
[edit vlans vlan-name interface interface-name egress],
[edit vlans vlan-name interface interface-name ingress],
[edit vlans vlan-name interface interface-name]
```

Description

Map a specific C-VLAN to an S-VLAN. By default, the received incoming or outgoing tag is replaced with the new tag.

This statement is also required if you are configuring firewall filters to map traffic from an interface to a VLAN. If you are configuring firewall filters to map traffic from an interface to a VLAN, the **mapping policy** option must be configured using this command. The firewall filter also has to be configured using the **vlan** action for a match condition in the firewall filter stanza for firewall filters to map traffic from an interface for a VLAN.

Options

For EX Series:

native—Maps untagged and priority-tagged packets to an S-VLAN.

policy—Maps the interface to a firewall filter policy to an S-VLAN.

push—Retains the incoming tag and add an additional VLAN tag instead of replacing the original tag.

swap—Swaps the incoming VLAN tag with the VLAN ID tag of the S-VLAN. Use of this option is also referred to as VLAN ID translation.

tag—Retains the incoming 802.1Q tag on the interface.

For QFX Series:

`inner-tag` (QFabric systems only)—apply the specified tag as an inner tag to packets that are received as untagged on an access interface.

`native`—Map untagged and priority-tagged packets to an S-VLAN.

`push`—Retain the incoming tag (as an inner tag) and adds an additional VLAN tag. When you use this option, the TPID of the outer tag is set as follows:

- If Q-in-Q tunneling is not enabled in the VLAN, then the Ethertype for outer tag is set to 0x8100.
- If Q-in-Q tunneling is enabled in the VLAN and a packet is egressing from a trunk port, then the Ethertype is set to 0x88a8 (or as configured by an `ether-type` statement).

`swap`—Replaces the incoming VLAN tag with the VLAN ID tag of the S-VLAN. Using this option is also referred to as VLAN ID translation. When you use this option on a trunk port for which Q-in-Q tunneling is enabled, use the `ether-type` statement to set the Ethertype.

`tag`—Original VLAN tag that will be replaced (with `swap`) or that will become an inner tag (with `push`).

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Option **swap** introduced in Junos OS Release 10.0 for EX Series switches.

RELATED DOCUMENTATION

Configuring Q-in-Q Tunneling on QFX Series Switches

Example: Setting Up Q-in-Q Tunneling on QFX Series Switches

Configuring VLANs for EX Series Switches

Understanding Q-in-Q Tunneling and VLAN Translation

Understanding Bridging and VLANs on Switches

mapping (VNF)

IN THIS SECTION

- Syntax | **10341**
- Hierarchy Level | **10341**
- Description | **10341**
- Options | **10341**
- Required Privilege Level | **10342**
- Release Information | **10342**

Syntax

```
mapping {
    hsxe0 {
        virtual-function {
            vlan-id vlan-id;
        }
    }
    hsxe1 {
        virtual-function {
            vlan-id vlan-id;
        }
    }
    vlan {
        members vlan-name;
            mode [access | trunk];
        native-vlan-id vlan-id;
    }
    peer-interfaces peer-interface-name;
}
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Mapping Virtual Network Functions (VNF) interfaces on platforms running disaggregated Junos OS.

Options

vlan-id	SR-IOV virtual function to use to attach a VNF to a physical interface.
vlan *members*	Membership for this interface.
native-vlan-id	Virtual LAN identifier for untagged frames. For example, 1...4095
vlan-name	Name of the VLAN members.

peer-interface-name Name of the virtual peer interfaces that is mapped to a physical interface.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D50.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

mapping-agent-election

IN THIS SECTION

- Syntax | 10342
- Hierarchy Level | 10343
- Description | 10343
- Options | 10343
- Required Privilege Level | 10343
- Release Information | 10343

Syntax

```
(mapping-agent-election | no-mapping-agent-election);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp auto-rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp auto-rp],
[edit protocols pim rp auto-rp],
[edit routing-instances routing-instance-name protocols pim rp auto-rp]
```

Description

Configure the routing device mapping announcements as a mapping agent.

Options

mapping-agent-election—Mapping agents do not announce mappings when receiving mapping messages from a higher-addressed mapping agent.

no-mapping-agent-election—Mapping agents always announce mappings and do not perform mapping agent election.

- **Default: mapping-agent-election**

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configuring PIM Auto-RP

mapping-range

IN THIS SECTION

- Syntax | 10344
- Hierarchy Level | 10344
- Description | 10344
- Options | 10344
- Required Privilege Level | 10345
- Release Information | 10345

Syntax

```
mapping-range C-VLAN-range (push | swap) <vlan-id-start S-VLAN-ID>;
```

Hierarchy Level

```
[edit vlans vlan-name vlans interface-name]
```

Description

Configure an access interface to map a range of C-VLANs to a range of S-VLANs. Use this statement instead of using multiple `set vlans VLAN-name interfaces interface-name mapping (push | swap)` statements to configure Q-in-Q tunneling or VLAN translation on a per-VLAN basis. This statement is particularly useful if you have used the `vlan-range` statement to create multiple VLANs.

Options

push—Retain the incoming tag and adds an additional VLAN tag (Q-in-Q tunneling).

swap—Swap the incoming VLAN tag with the VLAN ID tag of the S-VLAN (VLAN translation).

vlan-ID-start *S-VLAN-ID*—(Optional) Set the start of the S-VLAN range that the C-VLANs will be mapped to. If you omit this option, mapping begins with the first ID in the range of S-VLAN IDs (which you configure using the **set vlans vlan-range** statement).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

mapping-refresh

IN THIS SECTION

- Syntax | 10345
- Hierarchy Level | 10345
- Description | 10346
- Options | 10346
- Required Privilege Level | 10346
- Release Information | 10346

Syntax

```
mapping-refresh (inbound | outbound | inbound-outbound);
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated secure-nat-mapping]
```

Description

Specify how the flow timer should be refreshed based on the mapping refresh configured for all types of fwnat flows.

When configured, `tcp-tickles` sends tickles to both directions irrespective of mapping-refresh direction.

Starting in Junos OS Release 15.1R3, `mapping-refresh` is also supported on the MS-MPC and MS-MIC.

Options

inbound	Refresh the flow timer for inbound flows only.
inbound-outbound	Refresh the flow timer for all flows.
outbound	Refresh the flow timer for outbound flows only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3

RELATED DOCUMENTATION

Protecting CGN Devices Against Denial of Service (DOS) Attacks

mapping-server-entry

IN THIS SECTION

- Syntax | **10347**
- Hierarchy Level | **10347**

- Description | 10347
- Options | 10347
- Required Privilege Level | 10348
- Release Information | 10348

Syntax

```
mapping-server-entry mapping-server-name {
    prefix-segment prefix index number;
    prefix-segment-range prefix-segment-range-name;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options source-packet-routing],
[edit logical-systems name routing-options source-packet-routing],
[edit routing-instances name routing-options source-packet-routing],
[edit routing-options source-packet-routing]
```

Description

Configure an LDP mapping server to enable interoperability between islands of devices supporting only segment routing and only LDP in an LDP network domain.

The mapping server configuration can be included an on any device in the segment routing network.

> **NOTE**: The operating system accepts range sizes that are less than or equal to 1024. However, in cases where a range exceeds 1024, it will be automatically truncated to a maximum size of 1024.

Options

mapping-server-entry-name	Name of the LDP mapping server.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.2R1.

mapping-timeout

IN THIS SECTION

- Syntax | 10348
- Hierarchy Level | 10348
- Description | 10348
- Options | 10349
- Required Privilege Level | 10349
- Release Information | 10349

Syntax

```
mapping-timeout seconds;
```

Hierarchy Level

```
[edit services nat pool nat-pool-name]
```

Description

Specify the duration for mappings that use the specified NAT pool.

Options

seconds—Lifetime of mappings in seconds.

- **Default:** 300
- **Range:** 120 through 864,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

`mapping-timeout` statement introduced in JUNOS Release 10.1.

> **NOTE**: This configuration option has been replaced by "app-mapping-timeout" on page 2293. This option is currently retained only for backward compatibility.

RELATED DOCUMENTATION

Configuring Source and Destination Addresses Network Address Translation Overview

mapping-timeout (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | **10350**
- Hierarchy Level | **10350**
- Description | **10350**
- Options | **10350**
- Required Privilege Level | **10350**
- Release Information | **10350**

Syntax

```
mapping-timeout mapping-timeout;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name]
```

Description

Specify the timeout period for address-pooling paired mappings that use the specified NAT pool. Mappings that are inactive for this amount of time are dropped.

If you do not configure `ei-mapping-timeout` for endpoint independent translations, then the `mapping-timeout` value is used for endpoint independent translations.

Options

mapping-timeout *mapping-timeout*

Length of timeout period in seconds.

- **Range:** 120 through 86,400
- **Default:** 300

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

mapping-type

IN THIS SECTION

- Syntax | 10351
- Hierarchy Level | 10351
- Description | 10351
- Options | 10351
- Required Privilege Level | 10351
- Release Information | 10352

Syntax

```
mapping-type endpoint-independent;
```

Hierarchy Level

```
[edit services nat rule (Services NAT) rule-name term (Services NAT) term-name then (Services NAT)
translated]
```

Description

Specify the source NAT mapping type.

Options

endpoint-independent—Currently, the only valid setting specifies endpoint-independent mapping behavior.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.1.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

mapping-type (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 10352
- Hierarchy Level | 10352
- Description | 10353
- Options | 10353
- Required Privilege Level | 10353
- Release Information | 10353

Syntax

```
mapping-type {
    address-pooling-paired;
    endpoint-independent;
}
```

Hierarchy Level

```
[edit services nat source rule-set rule-set rule rule-name then source-nat]
```

Description

Configure the source NAT mapping type.

Options

endpoint-independent	Mapping to ensure that the same external address and port are assigned to all connections from a given host.
address-pooling-paired	Mapping to ensure assignment of the same external IP address for all sessions originating from the same internal host.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

mark-interface (RA Guard)

IN THIS SECTION

- Syntax | 10354
- Hierarchy Level | 10354
- Description | 10354
- Options | 10354
- Required Privilege Level | 10354
- Release Information | 10355

Syntax

```
mark-interface (trusted | block);
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard interface  interface-name]
```

Description

Configure an interface as blocked or trusted for IPv6 Router Advertisement (RA) guard. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. RA guard performs checks on incoming RA messages to make sure that they are sent from legitimate routers. If the sender of the RA message cannot be validated, the RA message is dropped.

You can configure the `mark-interface` statement on an interface to bypass RA guard policy checks on that interface. If an interface is configured as either a trusted interface or a blocked interface, RA messages received on the interface are not subject to inspection by RA guard, even if the interface or VLAN is enabled for RA guard. If the interface is trusted, it forwards all RA messages. If the interface is blocked, it drops all RA messages.

Options

`block` Configure an interface as blocked for bypassing inspection of RA messages received on that interface by RA guard. When you configure an interface as blocked, all RA messages received on the interface are dropped.

`trusted` Configure an interface as trusted for bypassing inspection of RA messages received on that interface by RA guard. When you configure an interface as trusted, all RA messages received on the interface are forwarded.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

mark-local-ip-host-routes

IN THIS SECTION

- Syntax | 10355
- Hierarchy Level | 10355
- Description | 10356
- Options | 10356
- Required Privilege Level | 10356
- Release Information | 10356

Syntax

```
mark-local-ip-host-routes export [policy-name];
```

Hierarchy Level

```
[edit routing-instances name protocols evpn]
```

Description

Configure `mark-local-ip-host-routes` along with the required "policy" on page 12736 and "community" on page 4465 to enable the Type 5 host-route auto-generated community. Leaf devices configured with the auto-generated community add a community to locally learned MAC-IP ARP/NDP based Type 5 host routes and Type 2 MAC-IP routes.

Options

export Policy to mark local IP host routes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1 and Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

EVPN-VXLAN Pure Type 5 Host-Route Auto-Generated Community

Routing Policies, Firewall Filters, and Traffic Policers User Guide

martians

IN THIS SECTION

- Syntax | 10357
- Hierarchy Level | 10357
- Description | 10357
- Options | 10357
- Required Privilege Level | 10358

Release Information | 10358

Syntax

```
martians {
    destination-prefix match-type <allow>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name],
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-options rib routing-table-name],
[edit routing-instances routing-instance-name routing-options],
[edit routing-instances routing-instance-name routing-options rib routing-table-name],
[edit routing-options],
[edit routing-options rib routing-table-name]
```

Description

Configure martian addresses.

Options

allow—(Optional) Explicitly allow a subset of a range of addresses that has been disallowed. The **allow** option is the only supported action.

destination-prefix—Destination route you are configuring:

- ***destination-prefix/ prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.
- **default**—Default route to use when routing packets do not match a network or host in the routing table. This is equivalent to specifying the IP address **0.0.0.0/0**.

match-type—Criteria that the destination must match:

- **exact**—Exactly match the route's mask length.
- **longer**—The route's mask length is greater than the specified mask length.
- **orlonger**—The route's mask length is equal to or greater than the specified mask length.
- **through *destination-prefix***—The route matches the first prefix, the route matches the second prefix for the number of bits in the route, and the number of bits in the route is less than or equal to the number of bits in the second prefix.
- **upto *prefix-length***—The route's mask length falls between the two destination prefix lengths, inclusive.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Removing the Class E Prefix on Martian Addresses

mask uli

IN THIS SECTION

- Syntax | 10359
- Hierarchy Level | 10359
- Description | 10359
- Options | 10359
- Required Privilege Level | 10359

Release Information | 10360

Syntax

```
mask-uli (1-65535);
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

This feature allows the user/operator to configure a GTP profile that masks the ULI. It provides control to the user by taking input on the mask value for the ULI. The value should lie between 1-65535 [00 01 to FF FF].

Attach this GTP profile to a policy which matches the operating conditions necessary to mask the ULI.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name]` hierarchy level is replaced by `[edit security gtp profile profile-name]`.

NOTE: The uli masking feature is only applicable for GTPv1 and GTPv2 only. It is not available for GTPv0.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.1R1.

master

IN THIS SECTION

- MX Series | 10360
- QFX Series | 10361
- ACX Series (ACX5448, ACX5448-D, and ACX5448-M) | 10361
- Hierarchy Level | 10362
- Description | 10362
- Options | 10362
- Required Privilege Level | 10363
- Release Information | 10363

MX Series

```
master {
    announce-interval announce-interval-value;
    clock-step (one-step | two-step);
    sync-interval sync-interval-value;
    interface interface-name {
        unicast-mode {
            transport ipv4;
            clock-client ip-address
           {
            local-ip-address local-ip-address;
            primary <primary-interface>
            secondary <secondary-interface>
            }
        }
    }
```

```
    multicast-mode {
        transport 802.3 link-local;
    }
}
```

QFX Series

```
master {
    interface interface-name {
        unicast-mode {
            transport ipv4;
            clock-client ip-address {
                local-ip-address local-ip-address;
            }
        }
    }
    multicast-mode {
        transport (ipv4 | ieee-802.3)
        local-ip-address local-ip-address;
    }
max-announce-interval max-announce-interval;
max-delay-response-interval max-delay-response-interval;
max-sync-interval max-sync-interval;
min-announce-interval min-announce-interval;
min-delay-response-interval min-delay-response-interval;
min-sync-interval min-sync-interval;
sync-interval sync-interval;
}
```

ACX Series (ACX5448, ACX5448-D, and ACX5448-M)

```
master {
    announce-interval announce-interval-value;
    interface interface-name {
        unicast-mode {
            clock-client ip-address local-ip-address local-ip-address {
                transport (ipv4 | ipv6)
            }
        }
```

```
    }
    sync-interval sync-interval;
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

Configure the timeTransmitter with parameters.

Options

master	Configure the timeTransmitter with parameers
announce-interval	Configure the mean interval for the announce messages from the timeTransmitter.
clock-step	Configures the clock-step that determines timing information on messages
sync-interval	Configure the mean interval for sync interval messages from the timeTransmitter.
unicast-mode	Configure upstream unicast PTP timeTransmitter clock sources. IPv4 and IPv6 support is available.
clock-client	Configure the IP address of the timeReceiver.
local-priority	Configure a clock's local priority to be used as a tie-breaker in the dataset comparison algorithm.
transport 802.3	Configure Ethernet as the encapsulation type for transport of PTP packets.
transport	Configure the encapsulation type for PTP packet transport.
local-ip-address	Configure the IP address of the interface acting as the timeReceiver or the timeTransmitter.
<primary-interface>	Configure the timeTransmitter interface name for the aggregated Ethernet bundle. (MX10004 Device)

secondary *<secondary-interface>*	Configure secondary interface name for the aggregated Ethernet bundle. (MX10004 Device)
max-announce-interval	Configure the maximum log mean interval between announce messages.
max-delay-response-interval	Configure the maximum log mean interval between delay-response messages.
max-sync-interval	Configure the maximum log mean interval between synchronization messages.
min-announce-interval	Configure the minimum log mean interval between announce messages.
min-delay-response-interval	Configure the minimum log mean interval between delay-response messages.
min-sync-interval	Configure the minimum log mean interval between synchronization messages.
sync-interval	Configure the log mean interval between synchronization messages.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

Configuring the Precision Time Protocol G.8275.2 Enhanced Profile (Telecom Profile)

master-only

IN THIS SECTION

- Syntax | 10364
- Hierarchy Level | 10364
- Description | 10364
- Required Privilege Level | 10364
- Release Information | 10365

Syntax

```
master-only;
```

Hierarchy Level

```
[edit interfaces (fxp0 | em0) unit  logical-unit-number family family address],
[edit logical-systems logical-system-name interfaces fxp0 unit  logical-unit-number family
family address]
```

Description

Configure the IP address to be used when the Routing Engine is the current primary.

NOTE: When you use the `master-only` cli command, you cannot use the management interface (fxp/em0) IP as the source address for services like system logging. If you use the management interface, the error **eventd: could not bind to address x.x.x.x: Can't assign requested address** is displayed from backup routing engine. This error occurs because the backup routing engine does not have access to master-only IP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Consistent Management IP Address

Junos OS CLI User Guide

master-password

IN THIS SECTION

- Syntax | 10365
- Hierarchy Level | 10366
- Description | 10366
- Options | 10366
- Required Privilege Level | 10366
- Release Information | 10366

Syntax

```
master-password {
    plain-text-password
    iteration-count iteration-count;
    pseudorandom-function (hmac-sha1 | hmac-sha2-256 | hmac-sha2-512);
}
```

Hierarchy Level

```
[edit system]
```

Description

Master password for 8-based password-encryption. The master password is used as input to the password-based key derivation function (PBKDF2) to generate an encryption key. The key is used as input to the Advanced Encryption Standard in Galois/Counter Mode (AES256-GCM). The plain text that the user enters is processed by the encryption algorithm (with key) to produce the encrypted text (cipher text).

Options

plain-text-password Set the master password with plain text. The password quality is evaluated for strength, and the device gives feedback if weak passwords are used.

iteration-count The number of iterations to use for the PBKDF2 hash function. The iteration count slows the hashing count, thus slowing attacker guesses.

- **Default:** 100
- **Range:** 10-10000

pseudorandom-function Choose the algorithm to use for unpredictable number generation.

- Values:
  - hmac-sha1—Hash-based MAC using secure hash algorithm-1 (SHA-1)
  - hmac-sha2-256—256-bits of hash-based MAC using SHA-2
  - hmac-sha2-512—512-bits of hash-based MAC using SHA-2

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

Hardening Shared Secrets in Junos OS

request system decrypt password | 15066

mastership-priority

IN THIS SECTION

- Syntax | 10367
- Hierarchy Level | 10367
- Description | 10368
- Default | 10368
- Options | 10368
- Required Privilege Level | 10368
- Release Information | 10369

Syntax

```
mastership-priority number;
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

Assign a primary-role priority value to a member of a Virtual Chassis or VCF. The primary-role priority is the most important factor that determines which member assumes the primary Routing Engine role in the Virtual Chassis or VCF.

A Virtual Chassis or VCF configuration has two member switches running in the Routing Engine role, one as the primary and one as the backup. The remaining members operate in the linecard role. In a preprovisioned Virtual Chassis or VCF, you configure the "role" on page 15824 for each member by associating it with the member switch's serial number, and the Virtual Chassis assigns default primary-role priorities based on the assigned role. In a non-provisioned Virtual Chassis or VCF, you can configure roles explicitly, or the primary-role priorities of the members implicitly designate the role of each member based on the primary-role election algorithm.

The Virtual Chassis or VCF elects the member switch with the highest primary-role priority value as the primary Routing Engine. When two members have the same highest primary-role priority value among all the members, the primary-role election algorithm uses other factors to determine which of the two members functions as the primary and which is the backup. See Understanding How the Primary in a Virtual Chassis Is Elected for all the conditions considered in primary-role election.

You should assign the *same* (highest) primary-role priority value to the members you want to be the primary and backup Routing Engine members to ensure reliable graceful Routing Engine switchover (GRES) operation. This and other primary-role election considerations also help keep primary role from switching back and forth rapidly between the two members under failover conditions.

This statement is not used for the EX8200 Virtual Chassis, which determines primary role by external Routing Engine uptime.

A switch with a primary-role priority of `0` never takes the primary or backup role.

Default

128

Options

number—Primary-role priority value.

- **Range:** 0 through 255

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric

Configuring an EX4650 or a QFX Series Virtual Chassis

Configuring an EX2300, EX3400, or EX4300 Virtual Chassis

Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis

match (Syslog)

IN THIS SECTION

- Syntax | 10369
- Hierarchy Level | 10369
- Description | 10370
- Options | 10370
- Required Privilege Level | 10370
- Release Information | 10370

Syntax

```
match "regular-expression";
```

Hierarchy Level

```
[edit logical-systems logical-system-name system syslog file filename],
[edit logical-systems logical-system-name system syslog user (username | *)],
```

```
[edit system syslog file filename],
[edit system syslog host hostname | other-routing-engine| scc-master)],
[edit system syslog user (username | *)]
```

Description

Specify a text string that must (or must not) appear in a message for the message to be logged to a destination.

If you configure both the `match` and `match-strings` statements for the same destination, Junos OS evaluates the `match-strings` condition first, and if the condition is satisfied, then the message is logged and the `match` condition is not evaluated. If the condition in the `match-strings` statement is not satisfied, then the system evaluates the regular expression in the `match` configuration statement.

Options

regular-expression

Regular expression against which to match messages to log.

The regular expressions must use the notation defined in POSIX Standard 1003.2 for extended (modern) UNIX regular expressions and must not contain a slash (/) or a percent sign (%).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Using Strings and Regular Expressions to Refine the Set of Logged Messages

match-strings | **10407**

match-direction

IN THIS SECTION

- Syntax | 10371
- Hierarchy Level | 10371
- Description | 10371
- Options | 10371
- Required Privilege Level | 10371
- Release Information | 10372

Syntax

```
match-direction (input | output | input-output);
```

Hierarchy Level

```
[edit services aacl rule rule-name]
```

Description

Specify the direction in which the rule match is applied.

Options

`input`—Apply the rule match on the input side of the interface.

`output`—Apply the rule match on the output side of the interface.

`input-output`—Apply the rule match bidirectionally.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Match Direction for AACL Rules

match-direction (ipsec-vpn)

IN THIS SECTION

- Syntax | 10372
- Hierarchy Level | 10372
- Description | 10373
- Options | 10373
- Required Privilege Level | 10373
- Release Information | 10373

Syntax

```
match-direction (input | output);
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name]
```

Description

Specify the direction in which the rule match is applied.

Options

input—Apply the rule match on input.

output—Apply the rule match on output.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

match-direction (IDS)

IN THIS SECTION

- Syntax | 10374
- Hierarchy Level | 10374
- Description | 10374
- Options | 10374
- Required Privilege Level | 10374
- Release Information | 10374

Syntax

```
match-direction (input | output | input-output);
```

Hierarchy Level

```
[edit services ids rule rule-name]
```

Description

Specify the direction in which the rule match is applied.

Options

`input`—Apply the rule match on input.

`output`—Apply the rule match on output.

`input-output`—Apply the rule match bidirectionally.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring Match Conditions in IDS Rules*
>
> *Configuring Protection Against Network Attacks on an MS-MPC*

match-direction (IP Reassembly Rule)

IN THIS SECTION

- Syntax | 10375
- Hierarchy Level | 10375
- Description | 10375
- Options | 10375
- Required Privilege Level | 10375
- Release Information | 10376

Syntax

```
match-direction direction
```

Hierarchy Level

```
[edit services ip-reassemblyrule rule-name]
```

Description

Configure the direction in which the IP reassembly rule matching is applied. The match direction is used with respect to the traffic flow through the inline services interface. You must configure a match direction for an IP reassembly rule.

Options

direction Match direction. For inline IP reassembly, `input` is the only match direction supported.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

> *Configuring IP Inline Reassembly for L2TP*
>
> *IP Packet Fragment Reassembly for L2TP Overview*

match-direction (NAT Next Gen Services)

IN THIS SECTION

- Syntax | **10376**
- Hierarchy Level | **10376**
- Description | **10377**
- Required Privilege Level | **10377**
- Release Information | **10377**

Syntax

Hierarchy Level

```
[edit services nat source rule-set rule-set],
[edit services nat destination rule-set rule-set]
```

Description

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

match-direction (PCP)

IN THIS SECTION

- Syntax | 10377
- Hierarchy Level | 10377
- Description | 10378
- Options | 10378
- Required Privilege Level | 10378
- Release Information | 10378

Syntax

```
match-direction (input | output);
```

Hierarchy Level

```
[edit services pcp rule rule-name]
```

Description

Specify the direction in which the rule match is applied. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

input Apply the rule match on input.

output Apply the rule match on output.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

match-direction (Services CoS)

IN THIS SECTION

- Syntax | 10379
- Hierarchy Level | 10379
- Description | 10379
- Options | 10379

Required Privilege Level | 10379
Release Information | 10379

Syntax

```
match-direction (input | output | input-output);
```

Hierarchy Level

```
[edit services cos rule rule-name]
```

Description

Specify the direction in which the rule match is applied.

Options

`input`—Apply the rule match on the input side of the interface.

`output`—Apply the rule match on the output side of the interface.

`input-output`—Apply the rule match bidirectionally.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

match-direction (Services NAT)

IN THIS SECTION

- Syntax | 10380
- Hierarchy Level | 10380
- Description | 10380
- Options | 10380
- Required Privilege Level | 10381
- Release Information | 10381

Syntax

```
match-direction (input | output);
```

Hierarchy Level

```
[edit services nat rule rule-name]
```

Description

Specify the direction in which the rule match is applied.

Options

input—Apply the rule match on input.

output—Apply the rule match on output.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

match-direction (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 10381
- Hierarchy Level | 10381
- Description | 10382
- Options | 10382
- Required Privilege Level | 10382
- Release Information | 10382

Syntax

```
match-direction (input | output | input-output);
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name]
```

Description

Specify the direction in which the rule match is applied.

Options

`input`—Apply the rule match on the input side of the interface.

`output`—Apply the rule match on the output side of the interface.

`input-output`—Apply the rule match bidirectionally.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Stateful Firewall Rules

match (Forwarding-Options)

IN THIS SECTION

- Syntax | 10383
- Hierarchy Level | 10383
- Description | 10383
- Required Privilege Level | 10383
- Release Information | 10383

Syntax

```
match expression;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling traceoptions file]
```

Description

Specify the regular expression for lines to be logged for tracing.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

Configuring Traffic Sampling on MX, M and T Series Routers

match-list

IN THIS SECTION

Syntax | 10384

- Hierarchy Level | 10384
- Description | 10384
- Options | 10385
- Required Privilege Level | 10385
- Release Information | 10385

Syntax

```
match-list {
    match-criteria {
        (match-all | match-any);
    }
    prefix-list-name prefix-list-name;
    source-ip-address-list address-list-name;
    source-mac-address-list address-list-name;
}
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard policy policy-name accept (Router
Advertisement Guard Policy)]
```

Description

Configure one or more lists of IPv6 addresses, MAC addresses, or IPv6 address prefixes to be associated with an IPv6 Router Advertisement (RA) guard *accept* policy.

RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.

You can configure match lists in either an accept policy or a discard policy. When RA guard is enabled by using an accept policy, RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped. When RA guard is enabled by using a

discard policy, RA messages that match the conditions are dropped, and RA messages that do not match the conditions are forwarded.

You can associate match lists or match conditions (see `match-option`) with an accept policy. You can configure match lists that be associated with an accept policy by using the match-list statement. The lists configured by using the `match-list` statement can contain IPv6 addresses, MAC addresses, or IPv6 address prefixes. RA guard examines the source address or address prefix. You configure the lists at the [`edit policy-options`] hierarchy level by using the `prefix-list` option for an IPv6 address or address prefix list, and `mac-list` for a MAC address list.

Options

`match-all` Configure the RA guard policy so that a received RA message is accepted only if it matches criteria in all of the lists configured under `match-list`; otherwise, the message is discarded.

`match-any` Configure the RA guard policy so that a received RA message is accepted if it matches criteria in any of the lists configured under `match-list`; otherwise, the message is discarded.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

> *Understanding IPv6 Router Advertisement Guard*
>
> *Configuring Stateful IPv6 Router Advertisement Guard*

match (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | 10386
- Hierarchy Level | 10386
- Description | 10386
- Options | 10386
- Required Privilege Level | 10386
- Release Information | 10387

Syntax

```
match match;
```

Hierarchy Level

```
[edit services rtlog traceoptions file]
```

Description

Regular expression for lines to be logged

Options

match Regular expression for lines to be logged

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

match-option

IN THIS SECTION

- Syntax | 10387
- Hierarchy Level | 10388
- Description | 10388
- Options | 10388
- Required Privilege Level | 10389
- Release Information | 10389

Syntax

```
match-option {
    hop-limit {
        (maximum | minimum) value;
    }
    managed-config-flag;
    other-config-flag;
    router-preference maximum (high | low | medium);
}
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard policy policy-name accept (Router
Advertisement Guard Policy)]
```

Description

Configure one or more parameters such as hop-count limit, managed configuration flag, other configuration flag, or router preference priority as the match condition to be associated with an IPv6 Router Advertisement (RA) guard *accept* policy.

RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.

You can associate match lists (see `match-list`) or match conditions with an accept policy. You can configure match conditions by using the `match-option` statement in an RA guard accept policy. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped.

Options

`hop-limit` Configure the RA guard policy to verify the minimum or maximum hop count for an incoming RA message. Use `maximum` to set a maximum hop count, or `minimum` to set a minimum hop count.

`managed-config-flag` Configure the RA guard policy to verify that the managed address configuration flag of an incoming RA message is set. When the managed address configuration flag is set, it indicates that addresses are available for allocation by Dynamic Host Configuration Protocol version 6 (DHCPv6).

`other-config-flag` Configure the RA guard policy to verify that the other configuration flag of an incoming RA message is set. When this flag is set, it indicates that other configuration information is available through DHCPv6. Examples of such information are DNS-related information or information on other servers within the network.

`router-preference-maximum` Configure the RA guard policy to verify that the default router preference parameter value of an incoming RA message is lower than or equal to a specified limit. The default router preference value improves the ability of IPv6 hosts to select a default router to

reach a remote destination when the host has multiple routers on its default router list. Use `high`, `medium`, or `low` to set the maximum preference.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

match-order (Dynamic Firewalls)

IN THIS SECTION

- Syntax | 10389
- Hierarchy Level | 10390
- Description | 10390
- Options | 10390
- Required Privilege Level | 10390
- Release Information | 10390

Syntax

```
match-order [match-order];
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family fast-update-filter filter-name]
```

Description

Specify the match conditions and the order in which the conditions are examined. Enclose a string of multiple conditions in brackets. The router examines only the conditions you specify, and examines them in the specified order.

Options

match-order—One or more of the following conditions. *Fast Update Filter Match Conditions* describes the match conditions.

- destination-address
- destination-port
- dscp (IPv4 only)
- protocol (IPv4 only)
- source-address
- source-port

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Fast Update Filters

Configuring the Match Order for Fast Update Filters

match-rules-on-reverse-flow

IN THIS SECTION

- Syntax | 10391
- Hierarchy Level | 10391
- Description | 10391
- Required Privilege Level | 10391
- Release Information | 10392

Syntax

```
match-rules-on-reverse-flow;
```

Hierarchy Level

```
[edit services service-set service-set-name cos-options]
```

Description

Configure the service set to create a CoS session even if a packet is first received in the reverse direction of the matching direction of the CoS rule. The CoS rule values are then applied as soon as a packet in the correct match direction is received.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R5 and 17.4R1.

RELATED DOCUMENTATION

Configuring CoS Rules

match-rules-on-reverse-flow (Next Gen Services)

IN THIS SECTION

- Syntax | 10392
- Hierarchy Level | 10392
- Description | 10392
- Required Privilege Level | 10393
- Release Information | 10393

Syntax

```
match-rules-on-reverse-flow;
```

Hierarchy Level

```
[edit services service-set service-set-name cos-options]
```

Description

Configure the service set to create a CoS session even if a packet is first received in the reverse direction of the matching direction of the CoS rule. The CoS rule values are then applied as soon as a packet in the correct match direction is received.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

RELATED DOCUMENTATION

Class of Service for Services PICs (Next Gen Services)

match (Security Destination NAT)

IN THIS SECTION

- Syntax | 10393
- Hierarchy Level | 10394
- Description | 10394
- Options | 10394
- Required Privilege Level | 10394
- Release Information | 10394

Syntax

```
match {
        application {
            [application];
            any;
        }
    (destination-address ip-address | destination-address-name address-name);
```

```
    destination-port (port-or-low <to high>);
    protocol [protocol-name-or-number];
    source-address [ip-address];
    source-address-name [address-name];
}
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name rule rule-name]
```

Description

Specify the destination rules to be used as match criteria.

> **NOTE**: If the options `destination-port` and `protocol` are configured as match conditions, then you cannot also configure the `application` option as a match condition. The reverse is also true: if you configure the `application` option as a match condition for a rule, you cannot also configure the `destination-port` and `protocol options`.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X47-D10.

match (Security IDP Policy)

IN THIS SECTION

- Syntax | 10395
- Hierarchy Level | 10395
- Description | 10396
- Options | 10396
- Required Privilege Level | 10396
- Release Information | 10396

Syntax

```
match {
    attacks (Security Exempt Rulebase) {
        custom-attack-groups [attack-group-name];
        custom-attacks [attack-name];
        dynamic-attack-groups [attack-group-name];
        predefined-attack-groups [attack-group-name];
        predefined-attacks [attack-name];
    }
    destination-address (Security IDP Policy) ([address-name] | any | any-ipv4 | any-ipv6);
    destination-except [address-name];
    from-zone (Security IDP Policy) (zone-name | any );
    source-address (Security IDP Sensor Configuration) ([address-name] | any | any-ipv4 | any-ipv6);
    source-except [address-name];
    to-zone (Security IDP Policy) (zone-name | any);
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name]
[edit security idp idp-policy policy-name rulebase-ips rule rule-name]
```

Description

Specify the rules to be used as match criteria.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

match (Security Policies)

IN THIS SECTION

- Syntax | 10396
- Hierarchy Level | 10397
- Description | 10397
- Options | 10398
- Required Privilege Level | 10398
- Release Information | 10398

Syntax

```
match {
    application {
        [application];
```

```
        any;
    }
    destination-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
    }
    source-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
    }
    source-identity {
        [role-name];
        any;
        authenticated-user;
        unauthenticated-user;
        unknown-user;
    }
    destination-identity-context {
        [user-or-role-name];
        any;
        identity-context-authenticated;
        identity-context-unauthenticated;
        identity-context-unavailable;
    }
    destination-identity-context-profile device-identity-profile-name;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name]
```

Description

Configure security policy match criteria.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Statement updated with the `source-identity` option in Junos OS Release 12.1. Statement updated with the `destination-identity-context` option in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Security Policies Overview

match (Security Policies Global)

IN THIS SECTION

- Syntax | 10399
- Hierarchy Level | 10400
- Description | 10400
- Options | 10400
- Required Privilege Level | 10400
- Release Information | 10400

Syntax

```
match {
    application {
        [application];
        any;
    }
    destination-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
    }
    destination-identity-feed feed-name
    from-zone {
        [zone-name];
        any;
    }
    source-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
    }
    source-identity {
        [role-name];
        any;
        authenticated-user;
        unauthenticated-user;
        unknown-user;
    }
    source-identity-feed feed-name;
    to-zone {
        [zone-name];
        any;
    }
}
```

Hierarchy Level

```
[edit security policies global policy policy-name]
```

Description

Configure security global policy match criteria.

> **NOTE**: We recommend that, for security reasons and to avoid spoofing traffic, when you create a multizone policy you use identical matching criteria (source address, destination address, application) and an identical action. For more information see *Global Policy Overview*.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5. Statement updated with `source-identity` option in Junos OS Release 12.1. Statement updated with `to-zone` and `from-zone` options in Junos OS Release 12.1X47-D10.

Statement updated with `source-identity-feed` and `destination-identity-feed` options in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Security Policies Overview

match (Security Source NAT)

IN THIS SECTION

- Syntax | 10401
- Hierarchy Level | 10401
- Description | 10401
- Options | 10402
- Required Privilege Level | 10402
- Release Information | 10402

Syntax

```
match {
    application {
        [application];
        any;
    }
    (destination-address <ip-address> | destination-address-name <address-name>);
    destination-port (port-or-low <to high>);
    protocol [protocol-name | protocol-number];
    source-address [ip-address];
    source-address-name [address-name];
    source-port (port-or-low <to high>);
}
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name]
```

Description

Specify the source rules to be used as match criteria.

> **NOTE**: If the options `source-port`, `destination-port`, and `protocol` are configured as match conditions, then you cannot also configure the `application` option as a match condition. The reverse is also true: if you configure the `application` option as a match condition for a rule, you cannot also configure the `source-port`, `destination-port`, and `protocol` options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

match (Security Static NAT)

IN THIS SECTION

- Syntax | 10403
- Hierarchy Level | 10403
- Description | 10403
- Options | 10403
- Required Privilege Level | 10403
- Release Information | 10403

Syntax

```
match {
    (destination-address <ip-address> | destination-address-name <address-name>);
    destination-port (port-or-low | <to high>);
    source-address [ip-address];
    source-address-name [ip-address-name];
    source-port (port-or-low <to high>);
}
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name]
```

Description

Specify the static rules to be used as match criteria.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10.

match (Services CoS Next Gen Services)

IN THIS SECTION

- Syntax | 10404
- Hierarchy Level | 10404
- Description | 10404
- Options | 10405
- Required Privilege Level | 10405
- Release Information | 10405

Syntax

```
match {
    application [ application-names ];
    destination-address address;
    destination-address-range low minimum-value high maximum-value;
    destination-port port-number;
    destination-prefix-list list-name;
    source-address address;
    source-address-range low minimum-value high maximum-value;
    source-prefix-list list-name;
}
```

Hierarchy Level

```
[edit services cos rule rule-name policy policy-name]
```

Description

Configure the matching conditions for a policy in a services CoS rule. Matching conditions include packet source and destination addresses and packet applications. Packets that are processed by a service set and that match the conditions are assigned the Differentiated Services (DiffServ) code point (DSCP) marking and forwarding-class assignments specified in the policy.

The service set that the CoS rule is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

application [*application-names*]	One or more port-based applications.	
destination-address *address*	Destination address of the packet.	
destination-address-range low *minimum-value* high *maximum-value*	Range of destination addresses of the packet.	
	minimum-value	Lower boundary of address range.
	maximum-value	Upper boundary of address range.
destination-port *port-number*	Destination port number of the packet.	
source-address *address*	Source address of the packet.	
source-address-range low *minimum-value* high *maximum-value*	Range of source addresses of the packet.	
	minimum-value	Lower boundary of address range.
	maximum-value	Upper boundary of address range.
source-prefix-list *list-name*	Name of a prefix list for matching the source address prefix. You configure the prefix list by using the `prefix-list` statement at the `[edit policy-options]` hierarchy level.	

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Class of Service for Services PICs (Next Gen Services)

match (Stateful Firewall Rule Next Gen Services)

IN THIS SECTION

- Syntax | 10406
- Hierarchy Level | 10406
- Description | 10406
- Options | 10407
- Required Privilege Level | 10407
- Release Information | 10407

Syntax

```
match {
    application [application-name];
    destination-address (address | any);
    destination-address-excluded address;
    source-address (address | any);
    source-address-excluded address;
}
```

Hierarchy Level

```
[edit services policies stateful-firewall-rule rule-name policy policy-name]
```

Description

Specify the matching properties for a stateful firewall rule policy. When a flow matches these properties, the policy actions are applied to the flow.

Options

application [*application-name*]	One or more application protocols of flows to which the stateful firewall policy applies. The application protocol definition is configured at the `[edit applications]` hierarchy level.
destination-address (*address* | any)	The destination address of the flows to which the stateful firewall rule policy applies. The option `any` matches all destination addresses.
destination-address-excluded *address*	The destination address of the flows to which the stateful firewall rule policy does not apply.
source-address (*address* | any)	The source address of the flows to which the stateful firewall rule policy applies. The option `any` matches all source addresses.
source-address-excluded *address*	The source address of the flows to which the stateful firewall rule policy does not apply.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Stateful Firewalls for Next Gen Services

match-strings

IN THIS SECTION

Syntax | **10408**

- Hierarchy Level | 10408
- Description | 10408
- Options | 10408
- Required Privilege Level | 10409
- Release Information | 10409

Syntax

```
match-strings [match-strings];
```

```
match-strings [string1 string2];
```

Hierarchy Level

```
[edit system syslog file name],
[edit system syslog host name],
[edit system syslog user name]
```

Description

Specify a text substring that must appear in a message that is logged to a destination. If multiple strings are configured, the message is logged to the given destination if any of the substrings are matched.

The `match-strings` statement performs a simple string comparison, and as a result, it is less CPU-intensive than using the `match` statement to match against complex regular expressions. If you configure both the `match` and `match-strings` statements for the same destination, Junos OS evaluates the `match-strings` condition first; if the message includes any of the configured substrings, then the message is logged and the `match` condition is not evaluated. If the `match-strings` condition is not satisfied, then the system evaluates the message against the regular expression in the `match` configuration statement.

Options

string String or list of strings against which to match messages to log.

- **Range:** 1 through 50 strings
- **Range:** 1 through 1020 characters per string

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Use Strings and Regular Expressions to Refine the Set of Logged Messages

matches

IN THIS SECTION

- Syntax | 10409
- Hierarchy Level | 10410
- Description | 10410
- Options | 10411
- Required Privilege Level | 10411
- Release Information | 10411

Syntax

```
matches {
    apply-groups [group-names];
```

```
    apply-groups-except [group-names];
    value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from called-station-id],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from calling-station-id],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from class],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from 3gpp-imsi],
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name format string]
```

Description

Specify the regular expression that the attribute must match.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

Table 57: Regular Expression Operators for the matches Statement

Operator	Matches
. (period)	One instance of any character except the space.
* (asterisk)	Zero or more instances of the immediately preceding term.
+ (plus sign)	One or more instances of the immediately preceding term.
? (question mark)	Zero or one instance of the immediately preceding term.
| (pipe)	One of the terms that appears on either side of the pipe operator.

Table 57: Regular Expression Operators for the matches Statement *(Continued)*

Operator	Matches
! (exclamation point)	Any string except the one specified by the expression when the exclamation point appears at the start of the expression. Use of the exclamation point is specific to Junos OS.
^ (caret)	Start of a line when the caret appears outside square brackets. One instance of any character that does not follow it within square brackets when the caret is the first character inside square brackets.
$ (dollar sign)	End of a line.
[ ] (paired square brackets)	One instance of one of the enclosed alphanumeric characters. To indicate a range of characters, use a hyphen (-) to separate the beginning and ending characters of the range. For example, [a-z0-9] matches any letter or number.
( ) (paired parentheses)	One instance of the evaluated value of the enclosed term. Parentheses are used to indicate the order of evaluation in the regular expression.

Options

value Regular expression to match.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

matches (Event Policy)

IN THIS SECTION

- Syntax | 10412
- Hierarchy Level | 10412
- Description | 10412
- Options | 10413
- Required Privilege Level | 10413
- Release Information | 10413

Syntax

```
event.attribute-name matches regular-expression;
```

Hierarchy Level

```
[edit event-options policy policy-name attributes-match]
```

Description

Configure an event policy to trigger only if the specified attribute of *event* matches the regular expression.

Options

event.attribute-name	Event attribute to compare to a regular expression.
regular-expression	Regular expression to compare.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Use Correlated Events to Trigger an Event Policy

Use Regular Expressions to Refine the Set of Events That Trigger a Policy

max-advertisement-interval (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 10414
- Hierarchy Level | 10414
- Description | 10414
- Options | 10414
- Required Privilege Level | 10414
- Release Information | 10414

Syntax

```
max-advertisement-interval seconds;
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Maximum interval between each router advertisement message.

Options

seconds—Maximum interval.

- **Range:** 4 through 1800 seconds
- **Default:** 600 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *Configuring an Interval Range for Unsolicited Router Advertisements to IPv6 Neighbors*
>
> *IPv6 WAN Link Addressing with NDRA*
>
> *Dynamic Router Advertisement Configuration Overview*

max-advertisement-interval (Protocols IPv6 Neighbor Discovery)

IN THIS SECTION

- Syntax | 10415
- Hierarchy Level | 10415
- Description | 10415
- Options | 10415
- Required Privilege Level | 10416
- Release Information | 10416

Syntax

```
max-advertisement-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Description

Set the maximum interval between each router advertisement message.

The router sends router advertisements on each interface configured to transmit messages. The advertisements include route information and indicate to network hosts that the router is operational. The router sends these messages periodically, with a time range defined by minimum and maximum values.

Options

seconds—Maximum interval.

- **Range:** 4 through 1800 seconds
- **Default:** 600 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

min-advertisement-interval (Protocols IPv6 Neighbor Discovery) | 10709

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

max-advertisement-interval (Protocols Router Discovery)

IN THIS SECTION

- Syntax | 10416
- Hierarchy Level | 10417
- Description | 10417
- Options | 10417
- Required Privilege Level | 10417
- Release Information | 10417

Syntax

```
max-advertisement-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery interface interface-name],
[edit protocols router-discovery interface interface-name]
```

Description

Specify the maximum time the router waits before sending periodic router advertisement packets out the interface. These packets are broadcast or multicast, depending on how the address corresponding to this physical interface is configured.

Options

seconds—Maximum time between router advertisement packets.

- **Range:** 4 through 1800 seconds
- **Default:** 600 seconds (10 minutes)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

broadcast (Router Discovery) | **3043**

lifetime (Router Discovery) | **9747**

min-advertisement-interval (Protocols Router Discovery) | **10711**

multicast (Protocols Router Discovery) | **11077**

max-age

IN THIS SECTION

- Syntax | 10418
- Hierarchy Level | 10418
- Description | 10419
- Default | 10419
- Options | 10419
- Required Privilege Level | 10419
- Release Information | 10419

Syntax

```
max-age seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp)],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id],
```

```
[edit protocols mstp],
[edit protocols rstp],
```

```
[edit protocols stp],
[edit protocols vstp vlan vlan-id]
```

```
[edit routing-instances routing-instance-name protocols (mstp | rstp)],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id]
```

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specifies the maximum expected arrival time of hello BPDUs.

Default

20 seconds

Options

seconds—(Optional) Number of seconds expected between hello BPDUs.

seconds—The maximum age of received protocol BPDUs.

- **Range:** 6 through 40
- **Default:** 20 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding STP

Understanding MSTP

Understanding VSTP

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

show spanning-tree bridge | 26548

show spanning-tree interface | 26555

max-areas

IN THIS SECTION

- Syntax | 10420
- Hierarchy Level | 10420
- Description | 10421
- Options | 10421
- Required Privilege Level | 10421
- Release Information | 10421

Syntax

```
max-areas number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis]
```

```
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Modify the maximum number of IS-IS areas advertised.

This value is included in the Maximum Address Area field of the IS-IS common PDU header included in all outgoing PDUs.

The maximum number of areas you can advertise is restricted to 36 to ensure that the IIH PDUs have enough space to include other type, length, and value (TLV) fields, such as the Authentication and IPv4 and IPv6 Interface Address TLVs.

Options

number—Maximum number of areas to include in the IS-IS hello (IIH) PDUs and link-state PDUs.

- **Range:** 3 through 36
- **Default:** 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

> *Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups*
>
> *Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding*

max-bypasses

IN THIS SECTION

- Syntax | 10422
- Hierarchy Level | 10422
- Description | 10422
- Options | 10422
- Required Privilege Level | 10423
- Release Information | 10423

Syntax

```
max-bypasses number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit protocols rsvp interface interface-name link-protection]
```

Description

Specify the maximum number of dynamic bypass LSPs permitted for protecting this interface. When this option is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled. The limit on bypasses configured applies only to dynamically generated bypass LSPs. By default, this option is disabled and only one dynamic bypass LSP is enabled for each interface. If you configure `max-bypasses`, you must also configure the "bandwidth" on page 2872 statement.

Options

number—Configure the maximum number of bypass LSPs. If you configure a value of 0, no dynamic bypass LSPs are allowed to be established for the interface. Only static bypass LSPs can be configured.

- **Range:** 0 through 99
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Range modified in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring the Maximum Number of Bypass LSPs

max-configurations-on-flash (System)

IN THIS SECTION

- Syntax | 10423
- Hierarchy Level | 10424
- Description | 10424
- Options | 10424
- Required Privilege Level | 10424
- Release Information | 10424

Syntax

```
max-configurations-on-flash number;
```

Hierarchy Level

```
[edit system]
```

Description

Specify the number of configurations stored on the CompactFlash card.

Options

number The number of configurations stored on the CompactFlash card.

- **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Using Junos OS to Specify the Number of Configurations Stored on the CompactFlash Card

max-connection-duration

IN THIS SECTION

Syntax | 10425

- Hierarchy Level | 10425
- Description | 10425
- Options | 10425
- Required Privilege Level | 10425
- Release Information | 10426

Syntax

```
max-connection-duration hours;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Specify the maximum time a connection can exist between a client and the server.

Options

hours Number of hours a connection can exist between a client and the server.

- **Default:** `24`
- **Range:** `0-120`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

max-connections

IN THIS SECTION

- Syntax | 10426
- Hierarchy Level | 10426
- Description | 10427
- Options (Junos OS) | 10427
- Options (Junos OS Evolved) | 10427
- Required Privilege Level | 10427
- Release Information | 10427

Syntax

```
max-connections max-connections;
```

Hierarchy Level

```
[edit system services extension-service request-response grpc]
```

Description

Number of simultaneous connections for request-response that can be attached to jsd. The higher the number, the higher the impact on clients performance.

Options (Junos OS)

- **Range:** 1 through 8
- **Default:** 8

Options (Junos OS Evolved)

Junos OS Evolved Release 22.2R1 and later:

- **Range:** 1 through 300

Releases before Junos OS Evolved Release 22.2R1:

- **Range:** 1 through 30

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

max-data-sessions-per-subscriber

IN THIS SECTION

- Syntax | 10428
- Hierarchy Level | 10428
- Description | 10428

Required Privilege Level | 10428
Release Information | 10428

Syntax

```
max-data-sessions-per-subscriber {
    limit max-sub-sessions;
    exceed-action {
        drop;
        syslog;
    }
}
```

Hierarchy Level

```
[edit services service-set services-set-name subscriber-profile profile-name]
```

Description

Specify the maximum number of sessions that are concurrently enabled for the named service. The system randomly selects a number of sessions and enables the named service for them. To limit the service's use of resources, other sessions cannot access these named services.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

max-datasize

IN THIS SECTION

- Syntax | 10429
- Hierarchy Level | 10429
- Description | 10429
- Default | 10430
- Options | 10430
- Required Privilege Level | 10431
- Release Information | 10431

Syntax

```
max-datasize size;
```

Hierarchy Level

```
[edit event-options event-script],
[edit system extension extension-service application],
[edit system scripts commit],
[edit system scripts op],
[edit system scripts snmp],
[edit system scripts translation]
```

Description

Maximum amount of memory allocated for the data segment during execution of a script of the configured type. The device sets the maximum memory limit for the executing script to the configured value irrespective of the total memory available on the system at the time of execution. If the executing script exceeds the specified maximum memory limit for that script type, it exits gracefully.

NOTE: For op scripts, the `max-datasize` statement is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the `op url` command, the device uses the default memory allocation settings.

NOTE: For op scripts run with the `max-datasize` statement configured for the minimum, an error occurs. In Junos OS, the error is "Memory allocation failed." In Junos OS Evolved, the error is "Out of memory."

Default

If you do not include the `max-datasize` statement, the default memory allocated to the data segment portion of the executed script depends on the operating system and release, which is as follows:

- Junos OS—Allocates half of the total available memory of the system up to a maximum value of 128 MB.
- Junos OS Evolved Release 21.4R1 and later—Allocates 1024 MB.
- Junos OS Evolved Release 21.3 and earlier—Allocates 128 MB.

Options

size Maximum amount of memory allocated for the data segment during execution of a script of the given type. If you do not specify a unit of measure, the default is bytes. The listed limits apply to the following types of scripts: commit, event, op, SNMP, translation, and extension service scripts.

- Syntax: *size* to specify bytes, *size*k to specify KB, *size*m to specify MB, or *size*g to specify GB
- Range: Table 58 on page 10430 lists the minimum and maximum configurable values for the statement for the different releases and script types.

Table 58: max-datasize Range

OS	Scripts	Releases	Minimum (bytes)	Maximum (bytes)
Junos OS (32-bit)	All	–	23,068,672 (22 MB)	1,073,741,824 (1 GB)

Table 58: max-datasize Range *(Continued)*

OS	Scripts	Releases	Minimum (bytes)	Maximum (bytes)
Junos OS (64-bit)	Commit, event, op, translation	-	23,068,672 (22 MB)	3,221,225,472 (3 GB)
	Extension service	16.1R3 and earlier	23,068,672 (22 MB)	1,073,741,824 (1 GB)
		16.1R4 and later 16.2R2 and later 17.1R1 and later	23,068,672 (22 MB)	3,221,225,472 (3 GB)
	SNMP	21.3 and earlier	23,068,672 (22 MB)	1,073,741,824 (1 GB)
		21.4R1 and later	23,068,672 (22 MB)	3,221,225,472 (3 GB)
Junos OS Evolved (64-bit)	All	21.3 and earlier	23,068,672 (22 MB)	1,073,741,824 (1 GB)
		21.4R1 and later	268,435,456 (256 MB)	2,147,483,648 (2 GB)

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support at the `[edit system extension extension-service application]` hierarchy level introduced in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, and vMX.

Support at the `[edit system scripts translation]` hierarchy level introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

max-policies (Event Policy) | 10455

Understanding Limits on Executed Event Policies and Memory Allocation for Scripts

Example: Configure Limits on Executed Event Policies and Memory Allocation for Scripts

max-db-size (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 10432
- Hierarchy Level | 10432
- Description | 10433
- Options | 10434
- Required Privilege Level | 10434
- Release Information | 10434

Syntax

```
max-db-size size;
```

Hierarchy Level

```
[edit system configuration-database]
```

Description

Specify the maximum amount of system shared memory that is available for the Junos OS configuration database and the schema database, together. Starting in Junos OS Release 20.1R1, Junos OS uses a single memory map for both the configuration and schema databases.

In Junos OS Release 19.4Rx and lower-numbered releases, this statement specifies the maximum amount of system shared memory that is available for only the Junos OS configuration database

JUNOS OS processes map shared memory into their process space. For example, on MX240 through MX10003 routers, processes can map up to 1GB of shared memory. Enhanced subscriber management processes contend for shared memory with the JUNOS OS configuration database. Shared memory that is not assigned to the configuration database is automatically available to enhanced subscriber management. By default, the configuration database tries to reserve 80 percent of the shared memory map, leaving insufficient space for subscriber management to function.

The majority of configurations require much less than 300MB of mapped space. An appropriate database size enables subscriber management to operate and scale optimally. In some circumstances, you must limit the size of the database to increase the amount of shared memory available to subscriber management. In other circumstances, we recommend that you allow the router to determine the appropriate size and that you do not configure a maximum size.

- For MX5, MX10, MX40, MX80, and MX104 routers, you must always configure the maximum size to be no more than 100MB, regardless of the which Junos OS release is running and regardless of Routing Engine RAM.

- For MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, and MX10003 routers, the decision whether to explicitly configure a maximum size and what that size is, depends on the Junos OS release and the amount of RAM in the Routing Engines.

  The `max-db-size` is an optional configuration command on routers having >=32 GB DRAM, for example, on MX960 platform. To enable subscriber-management, use the command `set chassis network-services enhanced-ip` and `set system services subscriber-management enable`. The router reboots and comes-up with subscriber-management enabled without `max-db-size` (optional) configuration and requires only 1 reboot.

-

Junos OS Release	Routing Engine RAM	Recommendation
Release 17.4R1 and earlier releases Release 18.1R1	Any	Configure maximum size to no more than 300MB.

(Continued)

Junos OS Release	Routing Engine RAM	Recommendation
Release 17.4R2 and higher 17.4x releases Release 18.1R2 and higher releases	Routing Engines have at least 32GB each	Allow the router to determine the appropriate size. Do not configure a maximum size.
Release 17.4R2 and higher 17.4x releases Release 18.1R2 and higher releases	Routing Engines have less than 32GB each	Configure maximum size to no more than 300MB.

BEST PRACTICE: You must reboot the device before a change in `max-db-size` value takes effect. By setting this value as part of your initial configuration, you can avoid multiple reboots.

Options

size Specifies the portion of system shared memory, in megabytes (MB), that is allocated for the Junos configuration database.

- **Syntax:** *`size`* `M` to specify MB
- **Default:** 698,343,424 bytes, only when running Junos OS Release 17.4R1, enhanced IP network services and enhanced subscriber management are enabled, and both Routing Engines in the chassis have at least 32 GB of RAM. Otherwise, there is no default value.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

Junos OS Enhanced Subscriber Management Overview

max-drop-flows

IN THIS SECTION

- Syntax | 10435
- Hierarchy Level | 10435
- Description | 10435
- Options | 10436
- Required Privilege Level | 10436
- Release Information | 10436

Syntax

```
max-drop-flows {
    ingress ingress-flows;
    egress egress-flows;
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Configure the maximum drop flows allowed per ingress and egress direction. The configuration is per service set. The configured limits indicate the maximum number of drop flows that can be created at a given instance of time in both directions. If max drop flows ingress is 10 and egress is 5 then at a given instance of time maximum of 10 ingress drop flows and 5 egress drop flows can be present. Two

counters, one for each direction ingress and egress, are to be added to service set stateful-firewall statistics to track the number of drop flows not created due to the drop flow limits exceeded. These limits applies to all types of drop flows i.e., TCP, UDP, ICMP etc. Ingress drop flows are forward flows for match-direction input rules and reverse flows for match-direction output rules. Similarly egress drop flows are reverse flows for match-direction input and forward flows for match-direction output rules. The limits are applied cumulatively on all the nat rules associated with the service-set.

If you specify the maximum drop flows to be zero, it indicates that the configuration is not effective. You must specify a value higher than zero for maximum drop flows.

Options

ingress-flows	Maximum number of drop flows on the ingress interface.
egress-flows	Maximum number of drop flows on the egress interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3

RELATED DOCUMENTATION

Configuring Service Set Limitations

max-duplicates

IN THIS SECTION

- Syntax | **10437**
- Hierarchy Level | **10437**

- Description | 10437
- Default | 10437
- Options | 10437
- Required Privilege Level | 10437
- Release Information | 10438

Syntax

```
max-duplicates number;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Specify the maximum number of content destinations to which the DFC PIC can send data from a single input set of packets. Limiting the number of duplicates reduces the load on the PIC. This setting overrides the globally applied `g-max-duplicates` setting.

Default

If no value is configured, a default setting of 3 is used.

Options

number—Maximum number of content destinations.

- **Range:** 1 through 64

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

g-max-duplicates | 7646

Limiting the Number of Duplicates of a Packet

max-failures

IN THIS SECTION

- Syntax | 10438
- Hierarchy Level | 10438
- Description | 10439
- Options | 10439
- Required Privilege Level | 10439
- Release Information | 10439

Syntax

```
max-failures max-failures;
```

Hierarchy Level

```
[edit logical-systems name protocols ppp-service],
[edit protocols ppp-service]
```

Description

The `max-failures` statement enables you to define the maximum number of failure attempts allowed while establishing the Point-to-Point Protocol (PPP) service.

Options

max-failures Specify the maximum number of failure attempts allowed while establishing the PPP service.

- **Range:** 1 through 10
- **Default:** 5

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

max-flows

IN THIS SECTION

- Syntax | 10440
- Hierarchy Level | 10440
- Description | 10440
- Options | 10440
- Required Privilege Level | 10440
- Release Information | 10440

Syntax

```
max-flows number;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Maximum number of flows allowed for the service set.

Options

number—Maximum number of flows.

> **NOTE**: When an aggregated multiservices (AMS) interface is configured as the service interface for a service set, the `max-flow` value configured for the service set is applied to each of the member interfaces in the AMS interface. That is, if you have configured 1000 as the `max-flow` value for a service set that uses an AMS interface with four active member interfaces, each of the member interfaces can handle 1000 flows each, resulting in an effective `max-flow` value of 4000.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Service Set Limitations

max-hops

IN THIS SECTION

- Syntax | 10441
- Hierarchy Level | 10441
- Description | 10441
- Options | 10442
- Required Privilege Level | 10442
- Release Information | 10442

Syntax

```
max-hops hops;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mstp],
[edit logical-systems logical-system-name routing-instances routing-instance-name
protocols mstp],
```

```
[edit protocols mstp],
```

```
[edit routing-instances routing-instance-name protocols mstp]
```

Description

Configure the maximum number of hops a BPDU can be forwarded in the MSTP region.

Options

hops—(Optional) Number of hops the BPDU can be forwarded.

- **Range:** 1 through 255
- **Default:** 19 hops

  20 hops

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding MSTP

Configuring Multiple Spanning Tree Protocol

Configuring MSTP on Switches

Example: Configuring Network Regions for VLANs with MSTP on Switches

Understanding MSTP

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

max-login-sessions

IN THIS SECTION

Syntax | **10443**

- Hierarchy Level | 10443
- Description | 10443
- Options | 10443
- Required Privilege Level | 10444
- Release Information | 10444

Syntax

```
max-login-sessions max-login-sessions;
```

Hierarchy Level

```
[edit fc-fabrics fc-fabric-name], [edit fc-fabrics fc-fabric-name interface interface-name]
```

Description

Set the maximum number of FCoE initialization protocol (FIP) session logins permitted for an individual NP_Port interface in an FCoE-FC gateway fabric (FC fabric) or for the entire FCoE-FC gateway fabric. You can set a maximum FIP session limit for each NP_Port interface connected to an FC switch. You can also set a maximum FIP session limit for the entire FC fabric. The sum of the maximum login sessions permitted on the NP_Port interfaces in an FC fabric should not exceed the maximum login sessions configured for that FC fabric.

The maximum number of FIP sessions (the combined total of all VN2VF_Port and VN2VN_Port sessions on the system) is 2500 sessions.

Options

max-login-sessions—Maximum number of FIP login sessions.

- **Range:** 128 through 2500
- **Default:** 2500

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

max-login-sessions-per-node | 10444
Setting the Maximum Number of FIP Login Sessions per FC Interface
Setting the Maximum Number of FIP Login Sessions per FC Fabric
Understanding Interfaces on an FCoE-FC Gateway

max-login-sessions-per-node

IN THIS SECTION

- Syntax | 10444
- Hierarchy Level | 10445
- Description | 10445
- Options | 10445
- Required Privilege Level | 10445
- Release Information | 10445

Syntax

```
max-login-sessions-per-node max-login-sessions-per-node;
```

Hierarchy Level

```
[edit fc-options]
```

Description

Set the maximum number of FCoE initialization protocol (FIP) session logins permitted on a Node device. (This is the combined total of all VN2VF_Port and VN2VN_Port sessions on the Node device.)

On a QFX3500 switch, the **max-login-sessions-per-node** command sets the maximum FIP session login limit for all of the FC fabrics configured on the device. The combined number of FIP sessions on all FC fabrics on the device should not exceed this limit.

On a QFabric system, the **max-login-sessions-per-node** command globally sets the maximum FIP session login limit for each QFabric system Node device in the QFabric system. For example, if you set the Node limit to 2000 login sessions, then each QFabric Node device supports up to 2000 FIP login sessions. The total configured maximum number of login sessions of all of the FC fabrics on a Node device should not exceed the Node session limit.

> **NOTE**: FIP login session limits configured at the FC fabric level or at the FC fabric interface level might limit a Node device to fewer total sessions than the configured Node limit.

Options

max-login-sessions-per-node—Maximum number of FIP login sessions.

- **Range:** 128 through 2500
- **Default:** 2500

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

max-login-sessions | 10442

Setting the Maximum Number of FIP Login Sessions per Node Device

Understanding Interfaces on an FCoE-FC Gateway

max-lsp-size

IN THIS SECTION

- Syntax | 10446
- Hierarchy Level | 10446
- Description | 10447
- Options | 10447
- Required Privilege Level | 10447
- Release Information | 10447

Syntax

```
max-lsp-size size;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
[edit logical-systems logical-system-name protocols isis level],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis
level]
```

Description

Modify the maximum size of IS-IS link-state PDUs. IS-IS sends link-state PDUs out of IS-IS enabled interfaces to distribute routing information between the IS-IS nodes. You can configure the maximum lsp size per level or at the global level. If configured at both global and per level, the per level setting takes precedence.

> **NOTE**: Set the max lsp size to a value less than or equal to the MTU of all the interfaces within the ISIS domain. Otherwise the bigger size LSP will not get flooded to all the nodes in the domain, which might lead to traffic loss and unforeseen events.
>
> Junos uses ethertype ox8870 to send IS-IS packets with maximum LSP size greater than 1500 bytes, which is not standardized across networks. This can sometimes cause packet loss in networks that do not recognize this value leading to network issues. Therefore, we recommend that you exercise caution when you configure this value to greater than 1500 bytes.

Options

size Maximum size allocated for IS-IS link-state PDUs.

- **Range:** 512 through 9192 bytes
- **Default:** 1492 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

Per Level statement introduced in Junos OS Release 23.1

RELATED DOCUMENTATION

Understanding IS-IS Configuration

Example: Configuring IS-IS

max-message-length

IN THIS SECTION

- Syntax | 10448
- Hierarchy Level | 10448
- Description | 10448
- Options | 10448
- Required Privilege Level | 10449
- Release Information | 10449

Syntax

```
max-message-length number;
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

In the GPRS tunneling protocol (GTP) header, the message length field indicates the length, in octets, of the GTP payload. Use the max-message-length option to set the maximum message payload length (in bytes). The maximum GTP message length is 65,535 bytes. The message length range is from 1 through 65,535 bytes.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name]` hierarchy level is replaced by `[edit security gtp profile profile-name]`.

Options

number—To set the maximum message payload length in bytes.

- **Range:** 1 through 65,535 bytes.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. Support.

RELATED DOCUMENTATION

Understanding GTP Message Filtering

Example: Setting the GTP Message-Length Filtering

max-outstanding-requests (Diameter Applications)

IN THIS SECTION

- Syntax | 10449
- Hierarchy Level | 10450
- Description | 10450
- Options | 10450
- Required Privilege Level | 10450
- Release Information | 10450

Syntax

```
max-outstanding-requests number;
```

Hierarchy Level

```
[edit access gx-plus global],
[edit access ocs partition partition-name],
[edit access pcrf partition partition-name]
```

Description

Limit the number of outstanding requests that the Diameter-based application (function) can retry to a remote server when the requests are improperly answered. Too many requests risks overloading the server and increases the chance of losing messages.

The `gx-plus` statement limits retries from the Gx-Plus function to the Gx-Plus server using the Gx and JSRC protocols. The `ocs` statement limits retries from the OCS function to the OCS server using the Gy protocol. The `pcrf` statement limits retries from the PCRF function to the PCRF server using the Gx protocol.

Options

number—Number of outstanding requests from the function to the server that can exist at any time.

- **Default:** 40
- **Range:** 2 through 40

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support for OCS and PCRF partitions introduced in Junos OS Release 16.2 for MX Series routers.

RELATED DOCUMENTATION

> *Configuring Gx-Plus Global Attributes*
>
> *Configuring Gx-Plus*

Configuring the OCS Partition

Configuring the PCRF Partition

Understanding Gx Interactions Between the Router and the PCRF

Understanding Gy Interactions Between the Router and the OCS

max-packet-memory-ratio

IN THIS SECTION

- Syntax | 10451
- Hierarchy Level | 10451
- Description | 10451
- Required Privilege Level | 10452
- Release Information | 10452

Syntax

```
max-packet-memory-ratio percentage-value;
```

Hierarchy Level

```
[edit security idp sensor-configuration application-identification]
```

Description

By default, the amount of IDP memory used for application identification packet memory is established as a percentage of all IDP memory. In most cases, the default value is adequate.

If a deployment exhibits an excessive number of ignored IDP sessions due to application identification memory allocation failures, use the `max-packet-memory-ratio` option to set application identification packet memory limit at a higher percentage of available IDP memory. This memory is only used by IDP in cases

where application identification delays identifying an application. Acceptable values are between 5 percent and 40 percent.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D20.

max-packets-per-second (Sampling)

IN THIS SECTION

- Syntax | 10452
- Hierarchy Level | 10452
- Description | 10453
- Options | 10453
- Required Privilege Level | 10453
- Release Information | 10453

Syntax

```
max-packets-per-second number;
```

Hierarchy Level

```
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance instance-name input]
```

Description

Specify the traffic threshold that must be exceeded before packets are dropped. A value of `0` instructs the Packet Forwarding Engine not to sample any traffic.

> **NOTE**: The `max-packets-per-second` statement is not supported when you configure inline flow monitoring (by including the `inline-jflow` statement at the `[edit forwarding-options sampling instance` *`instance-name`* `family (inet | inet6) output]` hierarchy level).

> **NOTE**: When you configure active monitoring and specify a Monitoring Services, Adaptive Services, or Multiservices PIC in the `output` statement, the `max-packets-per-second` value is ignored.

Options

`number`—Maximum number of packets per second.

- **Range:** 0 through 65,535
- **Default:** 1000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring Traffic Sampling on MX, M and T Series Routers*
>
> *Configuring Traffic Sampling*

max-pending-accounting-stops (Access Profile)

IN THIS SECTION

- Syntax | 10454
- Hierarchy Level | 10454
- Description | 10454
- Options | 10454
- Required Privilege Level | 10455
- Release Information | 10455

Syntax

```
max-pending-accounting-stops number;
```

Hierarchy Level

```
[edit access accounting-backup-options]
```

Description

Set the maximum number of pending accounting stop requests that the router backs up while all the RADIUS accounting servers in the profile are offline.

Options

number Number of stops to hold.

- **Range:** 1 through 168,000
- **Default:** 168,000

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Configuring Back-up Options for RADIUS Accounting

max-policies (Event Policy)

IN THIS SECTION

- Syntax | 10455
- Hierarchy Level | 10455
- Description | 10456
- Options | 10456
- Required Privilege Level | 10456
- Release Information | 10456

Syntax

```
max-policies policies;
```

Hierarchy Level

```
[edit event-options]
```

Description

Define the maximum number of event policies that can run concurrently on the device.

Options

policies Maximum number of event policies that can run concurrently.

- **Range:** 0 through 20
- **Default:** 15

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

max-datasize | **10429**

Understanding Limits on Executed Event Policies and Memory Allocation for Scripts

Example: Configure Limits on Executed Event Policies and Memory Allocation for Scripts

max-power

IN THIS SECTION

- Syntax | **10457**
- PTX10008 | **10457**
- PTX10001-36MR | **10457**
- Description | **10457**
- Options | **10457**

Required Privilege Level | 10458
Release Information | 10458

Syntax

```
max-power watts;
```

PTX10008

```
[edit chassis fpc slot-number]
```

PTX10001-36MR

```
[edit chassis psm]
```

Description

For the PTX10001-36MR router— Override the maximum power value of a power supply module (PSM) by specifying a lesser power value (in watts). You can use this configuration to monitor the power consumed in real time against the configured value. The system raises a chassis alarm when the consumption exceeds the configured power.

For the PTX10008 router—Override the default power budget allocated to the line card by specifying a power value (in watts). You can use the command `show chassis fpc detail` to view the maximum power consumption by a line card.

Options

watts—Specify power value (in watts) with which you want to override the default/maximum power.

NOTE: On the PTX10008 router, if you configure a power value (in watts) that is greater than the maximum value supported by the line card, the power budget for the line card is

set to its supported maximum. On the PTX10001-36MR router, if you configure a power value (in watts) that is greater than the maximum value supported by the PSM, the configuration does not work.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

The `max-power` *`watts`* statement introduced in Junos OS Evolved Release 20.1R2 for PTX10008 routers.

The `max-power` *`watts`* statement introduced in Junos OS Evolved Release 20.2R1 for PTX10001-36MR routers.

RELATED DOCUMENTATION

show chassis fpc | 18311

max-queues

IN THIS SECTION

- Syntax | 10459
- Hierarchy Level | 10459
- Description | 10459
- Options | 10459
- Required Privilege Level | 10460
- Release Information | 10460

Syntax

```
max-queues queues-per-line-card;
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

Configure the maximum number of queues allowed per MPC7E, MPC8E, MPC9E, MPC10E, and MPC11E line cards in an MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10001, MX10002 router or fixed MIC with XFP in an MX5, MX10, MX40, MX80, MX104, MX304 modular chassis router.

Reducing the number of queues allowed in a hierarchical scheduling environment in turn reduces the degree of jitter in the queues.

CAUTION: Committing a change to the maximum number of queues on a line card causes a reset of the line card.

Options

queues-per-line-card

Maximum number of queues allowed for the line card. Only the following keywords are valid: **8k**, **16k**, **32k**, **48k**, **64k**, **96k**, **128k**, **192k**, **256k**, **384k**, **512k**, **768k**, **1M**, **1536k**, or **2M**.

- Built-in 10-Gigabit Ethernet MICs and MPC1 Q line cards support up to 128 K queues. You can use this statement to configure the single Packet Forwarding Engine to support a lower maximum number of queues.
- MPC2 Q and MPC2 EQ line cards support up to 256 K queues. You can use this statement to configure the two Packet Forwarding Engines to support a lower maximum number of queues.

If you configure a keyword for a value that exceeds the number of queues supported by the line card hardware, the system uses the maximum number of queues supported by the line card.

If the `max-queues` statement is *not* configured on MPC7E, MPC8E, and MPC9E, which is the default mode, the MPC starts with a message similar to the following:

```
FPC 0 supports only port based queuing. A license is required for per-VLAN and hierarchical
features.
```

If the `max-queues` statement is configured on MPC7E, MPC8E, and MPC9E and the value is less than or equal to 32,000, the MPC starts with a message similar to the following:

```
FPC 0 supports port based queuing and is configured in 16384 queue mode. A limited per-VLAN queuing
license is required for per VLAN and hierarchical queuing features.
```

If the `max-queues` statement is configured on MPC7E, MPC8E, and MPC9E and the value is greater than 32,000, the MPC starts with a message similar to the following:

```
FPC 0 supports port based queuing and is configured in 524288 queue mode. A full scale per-VLAN
queuing license is required for per VLAN and hierarchical queuing features.
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Support for MPC7E, MPC8E, and MPC9E from Junos OS Release 15.1F4.

RELATED DOCUMENTATION

Jitter Reduction in Hierarchical CoS Queues

MPC1 Q

MPC2 Q

MPC2 EQ

MIC/MPC Compatibility

10-Gigabit Ethernet MICs with XFP

max-queues-per-interface

IN THIS SECTION

- Syntax | 10461
- Hierarchy Level | 10461
- Description | 10461
- Required Privilege Level | 10462
- Release Information | 10462

Syntax

```
max-queues-per-interface (8 | 4);
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number],
[edit chassis lcc number fpc slot-number pic pic-number] (Routing Matrix)
```

Description

On IQ, MPC, and DPC interfaces on M120, T320, T640, T1600, TX Matrix, and TX Matrix Plus routers, or on MIC or MPC interfaces on MX Series routers, set the number of egress queues per port to four or eight.

NOTE: If you include the `max-queues-per-interface 8` statement, the configuration at the `[edit class-of-service]` hierarchy level must also support eight queues per interface.

NOTE: When you include the `max-queues-per-interface` statement and commit the configuration, all physical interfaces on the line card are deleted and readded. Also, the FPC/MPC restarts automatically. If the FPC/MPC is a fixed-configuration line card, that

does not contain a MIC or PIC, you should restart the FPC/MPC to apply the configuration. You should change modes between four queues and eight queues only when there is no active traffic going to the FPC/MPC or to the PIC/MIC.

By default, IQ PICs on T Series and M320 routers are restricted to a maximum of four egress queues per interface. If you include the `max-queues-per-interface 4` statement, you can configure all four ports and configure up to four queues per port.

For Quad T3 and Quad E3 PICs and for 4-port OC3c/STM1 Type I and Type II PICs on M320 and T Series routers, when you include the `max-queues-per-interface 8` statement, you can configure up to eight queues on ports 0 and 2. After you commit the configuration, the PIC goes offline and comes back online with only ports 0 and 2 operational. No interfaces can be configured on ports 1 and 3.

NOTE: Starting from Junos OS Release 14.1R8, 14.2R6, 15.1F6, 15.1R3, 15.1R4, and 16.1R1, the restricted queue PICs without the `max-queues-per-interface` configuration boot up with a maximum of eight queues per port and two operational ports (port 0 and 2). PICs with restricted queues include Quad T3 PIC, Quad E3 PIC, 4-port SONET/SDH OC3c/STM1 PIC, and 4-Port OC3 and 1-port OC12 PICs with SFP.

On certain older MPCS (MPC1 Q, MPC1E Q, MPC2 Q, MPC2E Q), you can include the `max-queues-per-interface` statement to set the number of queues per logical interface to four or eight. Setting `max-queues-per-interface 4` sets the MPC to have four queues per logical interface and provides twice as many logical interfaces on the MPC as setting `max-queues-per-interface 8`.

NOTE: For consistency, `max-queues-per-interface` should not be set on MPCs starting from Junos OS 14.1X51.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for TX Matrix and TX Matrix Plus added in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Configuring the Junos OS to Support Eight Queues on IQ Interfaces for T Series and M320 Routers
- Configuring Up to 16 Custom Forwarding Classes
- Enabling Eight Queues on ATM Interfaces
- *Configuring the Maximum Number of Queues for Trio MPC/MIC Interfaces*
- Example: Configuring CoS on SRX5000 Devices with an MPC
- Example: Enabling Eight-Queue Class of Service on Redundant Ethernet Interfaces on SRX Series Devices in a Chassis Cluster

max-reass-packet-memory-ratio

IN THIS SECTION

- Syntax | 10463
- Hierarchy Level | 10463
- Description | 10464
- Required Privilege Level | 10464
- Release Information | 10464

Syntax

```
max-reass-packet-memory-ratio percentage-value;
```

Hierarchy Level

```
[edit security idp sensor-configuration application-identification]
```

Description

By default, the amount of IDP memory used for packet memory by the application identification reassembler is established as a percentage of all IDP memory. In most cases, the default value is adequate.

If a deployment exhibits an excessive number of ignored IDP sessions due to packet memory limitations of the application identification reassembler, use the `max-reass-packet-memory-ratio` option to set the reassembler packet memory limit to a higher percentage of available IDP memory. Acceptable values are between 5% and 40%.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D20.

max-retrans-time

IN THIS SECTION

- Syntax | 10465
- Hierarchy Level | 10465
- Description | 10465
- Options | 10465
- Required Privilege Level | 10466
- Release Information | 10466

Syntax

```
max-retrans-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip group group-name],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip group group-name],
[edit protocols rip group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip group group-name],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

RIP demand circuits send update messages every 5 seconds to an unresponsive peer. Configure the retransmission timer to limit the number of times the demand circuit resends update messages to an unresponsive peer. If the configured retransmission threshold is reached, routes from the next hop router are marked as unreachable and the hold-down timer starts. You must configure a pair of RIP demand circuits for this timer to take effect.

To determine the number of times to resend the update message, use the following calculation:

```
5 seconds x number of retransmissions = retransmission seconds
```

Options

seconds—The total amount of time the demand circuit resends update messages to an unresponsive peer. The seconds range corresponds to sending an update message a minimum of 1 time (5 seconds) and a maximum of 36 times (180 seconds).

- **Range:** 5 through 180 seconds
- **Default:** 5 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring RIP Demand Circuits

RIP Demand Circuits Overview

demand-circuit (Protocols RIP) | 5294

max-session-number

IN THIS SECTION

- Syntax | 10466
- Hierarchy Level | 10467
- Description | 10467
- Options | 10467
- Required Privilege Level | 10467
- Release Information | 10467

Syntax

```
max-session-number number;
```

Hierarchy Level

```
[edit security nat source rule-set ruleset rule rule then source-nat interface persistent-nat]
[edit security nat source rule-set ruleset rule rule then source-nat pool persistent-nat]
```

Description

The maximum number of the sessions with which a persistent NAT binding can be associated. For example, if the `max-session-number` of the persistent NAT rule is 65,536, then a 65,537th session cannot be established if that session uses the persistent NAT binding created from the persistent NAT rule.

Options

number—Maximum number of sessions.

- **Range:** 8 through 65,536
- **Default:** 30 sessions

Required Privilege Level

`security`—To view this statement in the configuration.

`security-control`—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 9.6.

max-session-setup-rate (Service Set)

IN THIS SECTION

- Syntax | 10468
- Hierarchy Level | 10468
- Description | 10468

- Options | 10468
- Required Privilege Level | 10468
- Release Information | 10469

Syntax

```
max-session-setup-rate (number | numberk);
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Set the maximum number of session setups allowed per second for the service set. After this setup rate is reached, any additional session setup attempts are dropped. If you do not include the `max-session-setup-rate` statement, the session setup rate is not limited.

Options

max-session-setup-rate *number*

Use the specified maximum number of session setups per second.

- **Range:** 1 through 429,496,729
- **Default:** 0 (The session setup rate is not limited.)

numberk

Maximum number of sessions, expressed in thousands. Starting in Junos OS Release 18.4R1, 1k=1000. Prior to Junos OS Release 18.4R1, 1k=1024.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring Service Set Limitations

max-sessions (Dynamic PPPoE)

IN THIS SECTION

- Syntax | 10469
- Hierarchy Level | 10469
- Description | 10470
- Options | 10470
- Required Privilege Level | 10470
- Release Information | 10470

Syntax

```
max-sessions number;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
pppoe],
[edit dynamic-profiles profile-name interfaces interface-set interface-set-name pppoe-underlying-
options]
```

```
[edit interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
pppoe-underlying-options]
```

Description

Configure the maximum number of dynamic PPPoE logical interfaces that the router can activate on the underlying interface. The `max-sessions` value does not affect the maximum number of static PPPoE logical interfaces that can be configured on the underlying interface.

> **NOTE**: The `[edit ... family pppoe]` hierarchies and the `[edit dynamic-profiles ... interfaces interface-set ... pppoe-underlying-options]` hierarchy level are supported only on MX Series routers with MPCs/MICs.

Options

number—Maximum number of dynamic PPPoE logical interfaces (sessions) that the router can activate on the underlying interface. The default value is equal to the maximum number of PPPoE sessions supported on your routing platform. You can configure from 1 to the platform-specific default for your routing platform. Changing the `max-sessions` value has no effect on dynamic PPPoE logical interfaces that are already active.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

Support for the `[edit ... family pppoe]` hierarchies introduced in Junos OS Release 11.2.

Support at the `[edit dynamic-profiles ... interfaces interface-set ... pppoe-underlying-options]` hierarchy level introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface

Defining ACI Interface Sets

PPPoE Maximum Session Limit Overview

Guidelines for Using PPPoE Maximum Session Limit from RADIUS

Juniper Networks VSAs Supported by the AAA Service Framework

Configuring an Interface Set of Subscribers in a Dynamic Profile

Subscriber Interfaces and PPPoE Overview

max-sessions-per-connection

IN THIS SECTION

- Syntax | 10471
- Hierarchy Level | 10471
- Description | 10472
- Options | 10472
- Required Privilege Level | 10472
- Release Information | 10472

Syntax

```
max-sessions-per-connection number
```

Hierarchy Level

```
[edit system services ssh]
```

Description

Specify the maximum number of ssh sessions allowed per single SSH connection.

Options

Default: 10

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configure SSH Service for Remote Access to the Router or Switch

ssh (System Services) | **28430**

Security Services Administration Guide

max-sessions-per-enode

IN THIS SECTION

- Syntax | **10473**
- Hierarchy Level | **10473**
- Description | **10473**
- Options | **10473**
- Required Privilege Level | **10473**
- Release Information | **10473**

Syntax

```
max-sessions-per-enode max-sessions-per-enode;
```

Hierarchy Level

```
[edit fc-fabrics fc-fabric-name protocols fip]
```

Description

Set the maximum number of FCoE login sessions (FLOGI plus FDISC) from a single ENode allowed on the gateway FC fabric (the fabric configured on the QFabric system). The maximum number of logins per ENode is 2000 sessions.

> **NOTE**: A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions. There is no limit to the number of end-to-end storage sessions.

Options

max-sessions-per-enode—Maximum number of FCoE sessions a single ENode can establish on the switch.

- **Range:** 32 through 2000
- **Default:** 32

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

fcoe-trusted | 6965

show fibre-channel fip | 19938

Configuring FIP on an FCoE-FC Gateway

Understanding FIP Parameters on an FCoE-FC Gateway

max-sessions-per-subscriber

IN THIS SECTION

- Syntax | 10474
- Hierarchy Level | 10474
- Description | 10474
- Options | 10475
- Required Privilege Level | 10475
- Release Information | 10475

Syntax

```
max-sessions-per-subscriber session-number;
```

Hierarchy Level

```
[edit services service-set service-set-name nat-options]
```

Description

Set the maximum number of sessions from a single subscriber allowed for NAPT-44. This statement does not apply to other types of NAT. The maximum number of sessions per subscriber is 32,000 sessions.

> **NOTE**: If you do not configure this statement, there is no limit to the number of sessions a subscriber can have.

Options

session-number—Maximum number of sessions a single subscriber can establish for NAPT-44.

- **Range:** 1 through 32000
- **Default:** None

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring Service Rules

land-attack-check | 9619

max-sessions-per-subscriber (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 10476
- Hierarchy Level | 10476
- Description | 10476
- Options | 10476

Required Privilege Level | 10476
Release Information | 10476

Syntax

```
max-sessions-per-subscriber session-number;
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Description

Set the maximum number of sessions allowed from a single subscriber.

Options

session-number Maximum number of sessions.

> **NOTE**: There is no default value. You must configure a value for the configuration to take effect.

- **Range:** 1 through 32000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

max-sessions (PPPoE Service Name Tables)

IN THIS SECTION

- Syntax | 10477
- Hierarchy Level | 10477
- Description | 10477
- Options | 10478
- Required Privilege Level | 10478
- Release Information | 10478

Syntax

```
max-sessions number;
```

Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name]
```

Description

Configure the maximum number of active PPPoE sessions using either static or dynamic PPPoE interfaces that the router can establish with the specified named service, `empty` service, or `any` service entry in a PPPoE service name table. The router maintains a count of active PPPoE sessions for each service entry to determine when the maximum sessions limit has been reached.

The router uses the `max-sessions` value for a PPPoE service name table entry in conjunction with the `max-sessions` value configured for the PPPoE underlying interface, and with the maximum number of PPPoE sessions supported on your router. If your configuration exceeds any of these maximum session limits, the router is unable to establish the PPPoE session.

Options

number—Maximum number of active PPPoE sessions that the router can establish with the specified PPPoE service name table entry, in the range 1 to the platform-specific maximum PPPoE sessions supported for your router. The default value is equal to the maximum number of PPPoE sessions supported on your routing platform.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name

Configuring PPPoE Service Name Tables

PPPoE Maximum Session Limit Overview

Configuring an Interface Set of Subscribers in a Dynamic Profile

Subscriber Interfaces and PPPoE Overview

max-sessions (Security Packet Log)

IN THIS SECTION

- Syntax | 10479
- Hierarchy Level | 10479
- Description | 10479
- Options | 10479
- Required Privilege Level | 10479
- Release Information | 10479

Syntax

```
max-sessions <percentage>;
per-logical-domain <number>
```

Hierarchy Level

```
[edit security idp sensor-configuration packet-log]
```

Description

Configure the maximum number of sessions actively conducting pre-attack packet captures on a device at one time. This value is expressed as a percentage of the maximum number of IDP sessions for the device.

Options

percentage—Maximum number of packet capture sessions expressed as a percentage of the IDP session capacity for the device.

- **Range:** 1 through 100 percent
- **Default:** 10

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

max-sessions (Tunnel Profile)

IN THIS SECTION

- Syntax | 10480
- Hierarchy Level | 10480
- Description | 10480
- Options | 10480
- Required Privilege Level | 10481
- Release Information | 10481

Syntax

```
max-sessions number;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Description

Specify the maximum number of sessions allowed in the tunnel.

Options

number—Maximum number of sessions allowed in the tunnel. A value of 0 means that the maximum configurable number of sessions is allowed.

- **Range:** 0 through 60,000
- **Default:** 0

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

max-sessions-vsa-ignore (Static and Dynamic Subscribers)

IN THIS SECTION

- Syntax | 10481
- Hierarchy Level | 10482
- Description | 10482
- Default | 10482
- Required Privilege Level | 10482
- Release Information | 10482

Syntax

```
max-sessions-vsa-ignore;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number pppoe-
underlying-options],
[edit interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
pppoe-underlying-options]
```

Description

Configure the router to ignore (clear) the value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks vendor-specific attribute (VSA) [26-143], and restore the PPPoE maximum session value on the underlying interface to the value configured in the CLI with the `max-sessions` statement. The PPPoE maximum session value specifies the maximum number of concurrent static or dynamic PPPoE logical interfaces (sessions) that the router can activate on the PPPoE underlying interface, or the maximum number of active static or dynamic PPPoE sessions that the router can establish with a particular service entry in a PPPoE service name table.

Default

If you do not include the `max-sessions-vsa-ignore` statement, the maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA takes precedence over the PPPoE maximum session value configured with the `max-sessions` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface

PPPoE Maximum Session Limit Overview

Guidelines for Using PPPoE Maximum Session Limit from RADIUS

Juniper Networks VSAs Supported by the AAA Service Framework

Configuring an Interface Set of Subscribers in a Dynamic Profile

Subscriber Interfaces and PPPoE Overview

max-snp-size

IN THIS SECTION

- Syntax | 10483
- Hierarchy Level | 10483
- Description | 10484
- Options | 10484
- Required Privilege Level | 10484
- Release Information | 10484

Syntax

```
max-snp-size size;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Modify the maximum size of partial or complete IS-IS sequence number PDUs. IS-IS sends sequence number packets out of IS-IS enabled interfaces to control the distribution of link-state PDUs between the IS-IS nodes. Sequence number packets provide a mechanism to synchronize the link-state databases of routers in the same area.

Options

size Maximum size allocated for sequence number of partial or complete IS-IS packets.

- **Range:** 512 through 1400 bytes
- **Default:** 1400 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding IS-IS Configuration

Example: Configuring IS-IS

max-tcp-session-packet-memory

IN THIS SECTION

Syntax | 10485

Hierarchy Level | 10485

- Description | 10485
- Options | 10485
- Required Privilege Level | 10485
- Release Information | 10485

Syntax

```
max-tcp-session-packet-memory value;
```

Hierarchy Level

```
[edit security idp sensor-configuration application-identification]
```

Description

Specify the maximum number of TCP sessions that IDP maintains. If the sensor reaches the maximum, it drops all new TCP sessions.

Options

value—Maximum number of TCP sessions.

- **Range:** 0 through 60,000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

max-time-report

IN THIS SECTION

- Syntax | 10486
- Hierarchy Level | 10486
- Description | 10486
- Options | 10486
- Required Privilege Level | 10487
- Release Information | 10487

Syntax

```
max-time-report value;
```

Hierarchy Level

```
[edit security idp sensor-configuration log suppression]
```

Description

When log suppression is enabled, IDP maintains a count of multiple occurrences of the same event. After the specified number of seconds has passed, IDP writes a single log entry containing the count of occurrences.

Options

value—Time after which IDP writes a single log entry containing the count of occurrences.

- **Range:** 1 through 60 seconds
- **Default:** 5 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

max-transmit-quality-level

IN THIS SECTION

- Syntax | 10487
- Hierarchy Level | 10487
- Description | 10488
- Options | 10488
- Required Privilege Level | 10489
- Release Information | 10489

Syntax

```
max-transmit-quality-level;
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Configure the threshold quality level for the router. If the received quality level is below the threshold quality level, then the router will send out a received quality level of SEC. The available quality levels are listed in Table 59 on page 10488.

> **NOTE**: For GPS external output, when you configure the maximum transmit quality level as PRC and the router is rebooted, no valid output is obtained from SCBE. However, when the maximum transmit quality level is configured to any other quality level other than PRC and the router gets rebooted, then the SCBE works normally.

Options

quality-level

The available quality levels are as given in Table 59 on page 10488.

Table 59: Quality Levels

Quality Level	Description
prc	Timing quality of a primary reference clock (option-1 only).
prs	Clock traceable to a primary reference source (option-2 only).
sec	Timing quality of an SDH equipment clock (option-1 only).
smc	Clock traceable to a self-timed SONET (option-2 only).
ssu-a	Timing quality of a type I or IV timeReceiver clock (option-1 only).
ssu-b	Timing quality of a type VI timeReceiver clock (option-1 only).
st2	Clock traceable to Stratum 2 (option-2 only).
st3	Clock traceable to Stratum 3 (option-2 only).
st3e	Clock traceable to Stratum 3E (option-2 only).

Table 59: Quality Levels *(Continued)*

Quality Level	Description
st4	Clock traceable to Stratum 4 free-run (option-2 only).
stu	Clock traceable to an unknown quality (option-2 only).
tnc	Clock traceable to a transit node clock (option-2 only).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (ACX Series) | 28787
synchronization (MX Series) | 28795
synchronization (PTX Series) | 28802

max-udp-session-packet-memory

IN THIS SECTION

- Syntax | 10490
- Hierarchy Level | 10490
- Description | 10490

- Options | 10490
- Required Privilege Level | 10490
- Release Information | 10490

Syntax

```
max-udp-session-packet-memory value;
```

Hierarchy Level

```
[edit security idp sensor-configuration application-identification]
```

Description

Specify the maximum number of UDP sessions that IDP maintains. If the sensor reaches the maximum, it drops all new UDP sessions.

Options

value—Maximum number of UDP sessions.

- **Range:** 0 through 20,000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

max-unknown-messages

IN THIS SECTION

- Syntax | 10491
- Hierarchy Level | 10491
- Description | 10491
- Options | 10491
- Required Privilege Level | 10492
- Release Information | 10492

Syntax

```
max-unknown-messages messages-per-minute;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Description

Specify the number of unknown messages per minute that the Path Computation Client (PCC) can receive at maximum after which the PCEP session is closed.

Options

messages-per-minute

Number of unknown messages per minute that the PCC can receive at maximum after which the PCEP session is closed. Recommended value is **5**. If the number of unknown messages received by the PCC is greater than or equal to the maximum number, the PCEP session is closed.

- **Range:** 1 through 16384

- **Default:** 0 (disabled or no limit)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

pce | 12303

max-unknown-requests

IN THIS SECTION

- Syntax | 10493
- Hierarchy Level | 10493
- Description | 10493
- Options | 10493
- Required Privilege Level | 10493
- Release Information | 10493

Syntax

```
max-unknown-requests requests-per-minute;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
[edit protocols pcep pce-group group-id]
```

Description

Specifies the number of unknown requests per minute that the Path Computation Client (PCC) can receive at maximum after which the PCEP session is terminated.

Options

requests-per-minute Number of unknown requests per minute that the PCC can receive at maximum after which the PCEP session is terminated.

- **Range:** 0 through 16384 (0 disables this statement)
- **Default:** 5

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

pce | 12303

max-withhold-time (Access Profile)

IN THIS SECTION

- Syntax | 10494
- Hierarchy Level | 10494
- Description | 10494
- Options | 10494
- Required Privilege Level | 10495
- Release Information | 10495

Syntax

```
max-withhold-time hold-time;
```

Hierarchy Level

```
[edit access accounting-backup-options]
```

Description

Set the timer that determines how long the router holds pending accounting stop requests. Any remaining accounting stop messages are flushed when the timer expires, even if the accounting server is again online.

Options

hold-time

Number of minutes.

- **Range:** 1 through 1440
- **Default:** 60

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Configuring Back-up Options for RADIUS Accounting

maximize-idp-sessions

IN THIS SECTION

- Syntax | 10495
- Hierarchy Level | 10496
- Description | 10496
- Options | 10496
- Required Privilege Level | 10496
- Release Information | 10496

Syntax

```
maximize-idp-sessions {
    weight (Security)  (equal | firewall | idp);
}
```

Hierarchy Level

```
[edit security forwarding-process application-services]
```

Description

If you are deploying IDP policies, you can tune the device to increase IDP session capacity. By using the provided commands to change the way the system allocates resources, you can achieve a higher IDP session capacity. See weight for information about the options provided.

This statement is supported on SRX1500, SRX 5800, SRX 5600, and SRX 5400 devices and vSRX Virtual Firewall.

NOTE: The IDP session capacity is restricted to 100,000 sessions per SPU.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

maximum

IN THIS SECTION

- Syntax | 10497
- Hierarchy Level | 10497
- Description | 10497
- Options | 10497
- Required Privilege Level | 10498
- Release Information | 10498

Syntax

```
maximum number;
```

Hierarchy Level

```
[edit interfaces interface-name services-options session-limit]
```

Description

Specify the maximum number of sessions allowed simultaneously on services cards. If you specify the maximum number of sessions to be zero, it indicates that the configuration is not effective. You must specify a value higher than zero for the maximum number of sessions.

Options

number Maximum number of sessions.

- **Range:** 1 through 4,294,967,295

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

maximum-age

IN THIS SECTION

- Syntax | 10498
- Hierarchy Level | 10498
- Description | 10499
- Options | 10499
- Required Privilege Level | 10499
- Release Information | 10499

Syntax

```
maximum-age minutes;
```

Hierarchy Level

```
[edit services flow-collector transfer-log-archive]
```

Description

Maximum age of transfer log file.

Options

minutes—Transfer log file age.

- **Range:** 1 through 360

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Transfer Logs

maximum-allowed-contentions

IN THIS SECTION

- Syntax | 10500
- Hierarchy Level | 10500
- Description | 10500
- Options | 10500
- Required Privilege Level | 10500
- Release Information | 10501

Syntax

```
maximum-allowed-contentions {
    count integer;
    duration seconds;
}
```

Hierarchy Level

```
[edit forwarding-options access-security slaac-snooping interface  (interface-name | all)]
```

Description

Configure the maximum number of Duplicate Address Detection (DAD) contentions for an interface. DAD is used by IPv6 clients to verify the uniqueness of addresses obtained through stateless address auto-configuration (SLAAC). DAD sends a Neighbor Solicitation message destined to the new address. If there is a reply, then the address is a duplicate; if the address is unique, it is assigned to the interface.

DAD contentions can be either Neighbor Solicitation or Neighbor Advertisement messages. If the maximum number of contentions is exceeded during the allowed time interval, the interface is considered invalid and the SLAAC snooping table is not updated with any bindings for that client.

> **NOTE**: The maximum allowed contentions configuration is not applicable on trusted ports. However, the CLI does not restrict the configuration of `max-allowed-contentions` on an interface that is configured with `mark-interface trusted`.

Options

count *integer* — Configure the number of DAD contentions permitted on the interface.

duration *seconds* — Configure the length of the interval during which the maximum allowed contentions can not be exceeded.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

IPv6 Stateless Address Auto-configuration (SLAAC) Snooping

maximum-backoff-duration

IN THIS SECTION

- Syntax | 10501
- Hierarchy Level | 10501
- Description | 10502
- Options | 10502
- Required Privilege Level | 10502
- Release Information | 10502

Syntax

```
maximum-backoff-duration milliseconds;
```

Hierarchy Level

```
[edit protocols ovsdb controller]
```

Description

Specify (in milliseconds) how long a Juniper Networks device that supports the *Open vSwitch Database (OVSDB)* management protocol waits before it tries again to connect with a software-defined networking (SDN) controller after a previous attempt has failed.

Options

milliseconds—Number of milliseconds a Juniper Networks device waits before it tries again to connect with an SDN controller.

- **Range:** 1000 through 4,294,967,295
- **Default:** 1000

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Understanding How to Set Up OVSDB Connections on a Juniper Networks Device

maximum-bandwidth

IN THIS SECTION

- Syntax | **10503**
- Hierarchy Level | **10503**
- Description | **10503**
- Options | **10503**

- Required Privilege Level | **10503**
- Release Information | **10503**

Syntax

```
maximum-bandwidth bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast interface interface-name],
[edit logical-systems logical-system-name routing-options multicast interface interface-name],
[edit routing-instances routing-instance-name routing-options multicast interface interface-name],
[edit routing-options multicast interface interface-name]
```

Description

Configure the multicast bandwidth for the interface.

Options

bps—Bandwidth rate, in bits per second, for the multicast interface.

- **Range:** 0 through any amount of bandwidth

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Example: Defining Interface Bandwidth Maximums

maximum-bandwidth (Protocols MPLS)

IN THIS SECTION

- Syntax | 10504
- Hierarchy Level | 10504
- Description | 10504
- Options | 10504
- Required Privilege Level | 10505
- Release Information | 10505

Syntax

```
maximum-bandwidth bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Specify the maximum amount of bandwidth in bits per second (bps).

Options

bps—Maximum amount of bandwidth.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth

maximum-bit-rate (Default Local Policy)

IN THIS SECTION

- Syntax | 10505
- Hierarchy Level | 10506
- Description | 10506
- Options | 10506
- Required Privilege Level | 10506
- Release Information | 10506

Syntax

```
maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber default-local-policy]
```

Description

Configure the maximum bit rate (MBR) for a subscriber's uplink and downlink traffic entering or exiting the access interface of the TDF domain when a TDF IP-based subscriber session does not exist. Uplink traffic originates from the subscriber towards the public data network (PDN), and downlink traffic comes from the PDN and is destined for the subscriber.

Options

mbr-uplink-value

MBR value for the uplink direction.

- **Range:** 0 through 6144000 Kbps.

mbr-downlink-value

MBR value for the downlink direction.

- **Range:** 0 through 6144000 Kbps.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

maximum-bit-rate (PCC Action Profiles)

IN THIS SECTION

- Syntax | 10507
- Hierarchy Level | 10507
- Description | 10507
- Default | 10508
- Options | 10508
- Required Privilege Level | 10508
- Release Information | 10508

Syntax

```
maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],
[edit services pcef pcc-action-profiles profile-name]
```

Description

Specify the maximum bit rate (MBR) that you want a PCC action profile to use for uplink and downlink traffic.

If you are using Junos OS Subscriber Aware, specify the MBR at the `[edit unified-edge pcef pcc-action-profiles profile-name` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the MBR at the `[edit services pcef pcc-action-profiles profile-name` hierarchy level.

Default

If you configure the `maximum-bit-rate` statement but do not specify MBR values for `uplink` and `downlink`, the default value is 0.

Options

mbr-uplink-value

MBR value for the uplink direction.

- **Range:** 1 through 6144000 Kbps.

mbr-downlink-value

MBR value for the downlink direction.

- **Range:** 1 through 6144000 Kbps.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-action-profiles` *`profile-name`* hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Configuring Policy and Charging Control Action Profiles for Subscriber Management

Understanding Application-Aware Policy Control for Subscriber Management

maximum-bit-rate (TDF Domain)

IN THIS SECTION

- Syntax | 10509
- Hierarchy Level | 10509
- Description | 10509
- Options | 10510
- Required Privilege Level | 10510
- Release Information | 10510

Syntax

```
maximum-bit-rate {
    apply-groups [group-names];
    apply-groups-except [group-names];
    downlink mbr-downlink-value;
    uplink mbr-uplink-value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
```

Description

Configure the TDF domain's default TDF subscriber maximum bit rate (MBR) for uplink and downlink traffic. Uplink traffic originates from the subscriber towards the public data network (PDN), and downlink traffic comes from the PDN and is destined for the subscriber.

Options

mbr-downlink-value MBR value for the downlink direction.

- **Range:** 0 through 1,048,000 Kbps.

mbr-uplink-value MBR value for the uplink direction.

- **Range:** 0 through 6,144,000 Kbps.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based and IFL-Based TDF Subscribers Overview

IP-Based Subscriber Setup Overview

maximum-call-duration (Security)

IN THIS SECTION

Syntax | 10511

Hierarchy Level | 10511

- Description | 10511
- Options | 10511
- Required Privilege Level | 10511
- Release Information | 10511

Syntax

```
maximum-call-duration value-in-minutes;
```

Hierarchy Level

```
[edit security alg mgcp],
[edit security alg sip]
```

Description

Specify the time at which the SIP call ends. The media session is released after the call has ended.

Options

value-in-minutes—Maximum amount of time at which the call ends and releases the media sessions.

- **Range:** 3 through 720 minutes
- **Default:** 720 minutes

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

mgcp | 10693

sip (Security) | 27977

maximum-capture-size (Datapath Debug)

IN THIS SECTION

- Syntax | 10512
- Hierarchy Level | 10512
- Description | 10512
- Options | 10512
- Required Privilege Level | 10513
- Release Information | 10513

Syntax

```
maximum-capture-size maximum-capture-size;
```

Hierarchy Level

```
[edit security datapath-debug]
```

Description

Specifies maximum packet capture length.

Options

- `maximum-capture-size` *`maximum-capture-size`*—Specify the maximum packet capture length.
- **Range:** 68 through 10,000 bytes

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

System Log Messages

maximum-certificates

IN THIS SECTION

- Syntax | 10513
- Hierarchy Level | 10513
- Description | 10514
- Options | 10514
- Required Privilege Level | 10514
- Release Information | 10514

Syntax

```
maximum-certificates number;
```

Hierarchy Level

```
[edit security certificates]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the maximum number of peer digital certificates to be cached.

Options

number—Maximum number of peer digital certificates to be cached.

- **Range:** 64 through 4,294,967,295 peer certificates
- **Default:** 1024 peer certificates

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Maximum Number of Peer Certificates

maximum-connections

IN THIS SECTION

- Syntax | **10515**
- Hierarchy Level | **10515**
- Description | **10515**
- Options | **10515**
- Required Privilege Level | **10515**

Release Information | 10515

Syntax

```
maximum-connections count;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Configure the maximum number of allowed connections between the server and all control client hosts.

> **NOTE**: The maximum number of connections between the server and all control client hosts must be greater than or equal to the number of connections between the server and a single controlled client host.

Options

count—Maximum number of connections.

- **Default:** 64
- **Range:** 1 through 1000

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

maximum-connections-per-client

IN THIS SECTION

- Syntax | 10516
- Hierarchy Level | 10516
- Description | 10516
- Options | 10517
- Required Privilege Level | 10517
- Release Information | 10517

Syntax

```
maximum-connections-per-client count;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Configure the maximum number of allowed connections between the server and a single control client host.

NOTE: The maximum number of connections between the server and all control client hosts must be greater than or equal to the number of connections between the server and a single controlled client host.

Options

count—Maximum number of connections.

- **Default:** 64
- **Range:** 1 through 500

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

maximum-contexts

IN THIS SECTION

- Syntax | **10518**
- Hierarchy Level | **10518**
- Description | **10518**
- Options | **10518**
- Required Privilege Level | **10518**
- Release Information | **10518**

Syntax

```
maximum-contexts number <force>;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number compression rtp],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
compression rtp]
```

Description

Specify the maximum number of RTP contexts to accept during negotiation.

Options

number—Maximum number of contexts.

force—(Optional) Requires the PIC to use the value specified for maximum RTP contexts, regardless of the negotiated value. This option allows the software to interoperate with Junos OS Releases that base the RTP context value on link speed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configuring Compression of Voice Traffic

maximum-discovery-table-entries

IN THIS SECTION

- Syntax | 10519
- Hierarchy Level | 10519
- Description | 10519
- Default | 10519
- Options | 10520
- Required Privilege Level | 10520
- Release Information | 10520

Syntax

```
maximum-discovery-table-entries entry-number;
```

Hierarchy Level

```
[edit protocols ancp],
[edit protocols ancp neighbor ip-address]
```

Description

Specify the maximum number of discovery table entries accepted from all ANCP neighbors or from a particular ANCP neighbor. The number of entries configured for an individual neighbor supersedes the global value. The neighbor can continue to update previously created entries when the maximum has been exceeded, but no new entries are accepted.

Default

No limit on the number of table entries

Options

entry-number—Maximum number of discovery table entries.

- **Range:** 1 through 100,000
- **Default:** 100,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring the ANCP Agent

Configuring ANCP Neighbors

maximum-ecmp

IN THIS SECTION

- Syntax | **10521**
- Hierarchy Level | **10521**
- Description | **10521**
- Default | **10521**
- Options | **10521**
- Required Privilege Level | **10522**
- Release Information | **10522**

Syntax

```
maximum-ecmp next-hops;
```

Hierarchy Level

```
[edit chassis]
```

Description

MX Series) Configure 16, 32, or 64, and 128 ECMP next hops for RSVP or LDP LSPs, or MPLS static LSPs that are configured using `set protocols mpls static-label-switched-path`.

This command is used to control the maximum number of ECMP legs in an NH. This command applies to all protocols, and the maximum configurable ECMP next hops are chassis dependent.

> **NOTE**: MX Series routers with one or more Modular Port Concentrator (MPC) cards and with Junos OS 11.4 or earlier installed, support the configuration of the `maximum-ecmp` statement with only 16 next hops. You should *not* configure the `maximum-ecmp` statement with 32 or 64 next hops. When you commit the configuration with 32 or 64 next hops, the following warning message appears:
>
> `Error: Number of members in Unilist NH exceeds the maximum supported 16 on Trio.`

Default

16

Options

next-hops—Specify the number of next hops (16, 32, 64, or 128) for RSVP or LDP LSPs, or MPLS static LSPs.

The following types of routes support next hops with up to 128 ECMP gateways:

- RSVP routes in inet tables where multiple RSVP LSPs are created to the same destination. In the case where LSP path protection or FRR is configured, the combination of active, backup and FRR next-hops is a maximum of 128.
- LDP routes in inet.3 and mpls.0 where the associated IGP route contains 128 next-hop gateways.

-
- ISIS, OSPF, iBGP, eBGP and Static routes in inet and inet6 tables.

> **NOTE**: These routes also include routes in routing-instances (foo.inet.0).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing

maximum-helper-recovery-time

IN THIS SECTION

- Syntax | 10523
- Hierarchy Level | 10523
- Description | 10523
- Options | 10523
- Required Privilege Level | 10523
- Release Information | 10523

Syntax

```
maximum-helper-recovery-time seconds;
```

Hierarchy Level

```
[edit protocols rsvp graceful-restart],
[edit logical-systems logical-system-name protocols rsvp graceful-restart]
```

Description

Specify the length of time the router or switch retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.

Options

seconds—Legnth of time that the router retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.

- **Range:** 1 through 3600
- **Default:** 180

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for RSVP, CCC, and TCC

maximum-helper-restart-time (RSVP) | **10525**

maximum-helper-restart-time

IN THIS SECTION

- Syntax | 10524
- Hierarchy Level | 10524
- Description | 10524
- Options | 10524
- Required Privilege Level | 10525
- Release Information | 10525

Syntax

```
maximum-helper-restart-time seconds;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Specify how long other router processes wait for the ANCP agent to restart before considering it to be down.

Options

seconds—Number of seconds other processes wait for ANCP to restart.

- **Range:** 45 through 600 seconds
- **Default:** 45 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent

Specifying How Long Processes Wait for the ANCP Agent Restart to Complete

maximum-helper-restart-time (RSVP)

IN THIS SECTION

- Syntax | 10525
- Hierarchy Level | 10526
- Description | 10526
- Options | 10526
- Required Privilege Level | 10526
- Release Information | 10526

Syntax

```
maximum-helper-restart-time seconds;
```

Hierarchy Level

```
[edit protocols rsvp graceful-restart],
[edit logical-systems logical-system-name protocols rsvp graceful-restart]
```

Description

Specify the length of time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. This value is applied to all RSVP neighbor routers and should be based on the time that the slowest RSVP neighbor requires for restart.

Options

seconds—The time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down.

- **Range:** 1 through 1800
- **Default:** 60

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

> *Configuring Graceful Restart Options for RSVP, CCC, and TCC*
>
> maximum-helper-recovery-time | **10522**

maximum-hop-count

IN THIS SECTION

- Syntax | 10527
- Hierarchy Level | 10527
- Description | 10527
- Options | 10527
- Required Privilege Level | 10528
- Release Information | 10528

Syntax

```
maximum-hop-count number;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface (interface-name | interface-group)]
```

Description

Set the maximum allowed number of hops. This value is compared against the hops field in the BOOTP request message. BOOTP request messages that have a number in the hops field that exceeds `maximum-hop-count` are not forwarded. If you omit the `maximum-hop-count` statement, the default value is four hops.

Options

number—Maximum number of hops for BOOTP request messages.

- **Range:** 1 through 16
- **Default:** 4

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

maximum-labels

IN THIS SECTION

- Syntax | 10528
- Hierarchy Level | 10529
- Description | 10529
- Options | 10529
- Required Privilege Level | 10529
- Release Information | 10529

Syntax

```
maximum-labels maximum-labels;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family mpls],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family mpls]
```

Description

On the logical interface, specify the maximum number of MPLS labels upon which MPLS can operate.

Options

maximum-labels—Maximum number of labels for the protocol family.

> **NOTE**: On PTX Series routers with third-generation FPCs, the maximum labels that can be pushed cannot exceed 8 labels.

- **Range:** 3 through 16
- **Default:** 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *Configuring the Maximum Number of MPLS Labels*
>
> Junos OS VPNs Library for Routing Devices

maximum-lease-time (DHCP)

IN THIS SECTION

- Syntax | 10530
- Hierarchy Level | 10530
- Description | 10530
- Options | 10530
- Required Privilege Level | 10530
- Release Information | 10531

Syntax

```
maximum-lease-time seconds;
```

Hierarchy Level

```
[edit system services dhcp]
```

Description

For J Series Services Routers and EX Series switches only. Specify the maximum length of time in seconds for which a client can request and hold a lease on a DHCP server.

An exception is that the dynamic BOOTP lease length can exceed the maximum lease length specified.

Options

seconds—The maximum number of seconds the lease can be held.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration

Release Information

Statement introduced before Junos OS Release 7.4.

maximum-length (Origin Validation for BGP)

IN THIS SECTION

- Syntax | 10531
- Hierarchy Level | 10531
- Description | 10532
- Options | 10532
- Required Privilege Level | 10532
- Release Information | 10532

Syntax

```
maximum-length prefix-length {
    origin-autonomous-system as-number {
        validation-state (invalid | valid);
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation static record destination],
[edit logical-systems logical-system-name routing-options validation static record destination],
[edit routing-instances instance-name routing-options validation static record destination],
[edit routing-options validation static record destination]
```

Description

Configure the maximum prefix-length for a route validation (RV) record. This is a required statement.

Options

prefix-length—Maximum prefix-length range for a given RV entry.

- **Range:** 1 through 128

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

maximum-links

IN THIS SECTION

- Syntax | 10533
- Hierarchy Level | 10533
- Description | 10533
- Options | 10533

Required Privilege Level | 10534
Release Information | 10534

Syntax

```
maximum-links maximum-links-limit;
```

Hierarchy Level

```
[edit chassis aggregated-devices]
```

Description

Configure the maximum links limit for aggregated devices. Note that for MX Series routers, to set a range of 32 or 64 the router must be running in Enhanced IP mode, which is only supported for Trio-based MPCs and multiservice DPCs (MS-DPCs). For more information on Enhanced IP mode, Network Services Mode Overview.

For MX series routers and PTX series switches, the option for 64 links is only supported for Junos OS release 12.3 and later.

NOTE: This statement is not supported on the MX80, MX104, and PTX1000 routers.

Options

maximum-links-limit Maximum links limit for aggregated devices.

- **Range:** 16, 32, 64

NOTE: On T-Series routers, the maximum-links supported is 32 in an aggregated Ethernet link.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- Network Services Mode Overview
- Configuring Junos OS for Supporting Aggregated Devices
- Configuring an Aggregated Ethernet Interface
- network-services | 11320

maximum-message-length

IN THIS SECTION

- Syntax | 10534
- Hierarchy Level | 10535
- Description | 10535
- Options | 10535
- Required Privilege Level | 10535
- Release Information | 10535

Syntax

```
maximum-message-length number;
```

Hierarchy Level

```
[edit logical-systems name security alg dns],
[edit logical-systems name tenants name security alg dns],
[edit security alg dns],
[edit services alg dns],
[edit tenants name security alg dns]
```

Description

Configure the maximum-message-length parameter to specify the maximum DNS message length.

Options

`number`—Maximum length in bytes of a single DNS message.

- **Range:** 512 through 8192 bytes.
- **Default:** 512 bytes.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

dns | 5847

maximum (MSDP Active Source Messages)

IN THIS SECTION

- Syntax | 10536
- Hierarchy Level | 10536
- Description | 10536
- Options | 10536
- Required Privilege Level | 10537
- Release Information | 10537

Syntax

```
maximum number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp active-source-limit],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
active-source-limit],
[edit protocols msdp active-source-limit],
[edit routing-instances routing-instance-name protocols msdp active-source-limit]
```

Description

Configure the maximum number of MSDP active source messages the router accepts.

Options

number—Maximum number of active source messages.

- **Range:** 1 through 1,000,000
- **Default:** 25,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP with Active Source Limits and Mesh Groups

threshold (MSDP Active Source Messages) | **29327**

maximum-neighbor-reconnect-time

IN THIS SECTION

- Syntax | **10537**
- Hierarchy Level | **10538**
- Description | **10538**
- Options | **10538**
- Required Privilege Level | **10538**
- Release Information | **10538**

Syntax

```
maximum-neighbor-reconnect-time seconds;
```

Hierarchy Level

```
[edit protocols ldp graceful-restart],
[edit logical-systems logical-system-name protocols ldp graceful-restart],
[edit routing-instances routing-instance-name protocols ldp graceful-restart]
```

Description

Specify the maximum length of time allowed to reestablish connection from a restarting neighbor.

Options

seconds—Maximum time allowed for reconnection.

- **Range:** 30 through 300

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for LDP

maximum-neighbor-recovery-time

IN THIS SECTION

Syntax | 10539

- Hierarchy Level | 10539
- Description | 10539
- Options | 10539
- Required Privilege Level | 10539
- Release Information | 10540

Syntax

```
maximum-neighbor-recovery-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ldp graceful-restart],
[edit protocols ldp graceful-restart],
[edit routing-instances routing-instance-name protocols ldp graceful-restart]
```

Description

Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.

Options

seconds—Configure the maximum recovery time, in seconds.

- **Range:** 120 through 1800 seconds
- **Default:** 140 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4. Statement changed from `maximum-recovery-time` to `maximum-neighbor-recovery-time` in Junos OS Release 9.1.

RELATED DOCUMENTATION

Configuring Recovery Time and Maximum Recovery Time

Configuring Graceful Restart Options for LDP

recovery-time (LDP) | **14148**

maximum-packet-length

IN THIS SECTION

- Syntax | **10540**
- Hierarchy Level | **10540**
- Description | **10541**
- Options | **10541**
- Required Privilege Level | **10541**
- Release Information | **10542**

Syntax

```
maximum-packet-length bytes;
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input],
[edit forwarding-options port-mirroring input],
[edit forwarding-options port-mirroring instance instance-name input],
```

```
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance instance-name input]
```

Description

Set the maximum packet length to be used for port mirroring or traffic sampling. Packets longer than the maximum are truncated. This statement cannot be used with inline flow monitoring (`[edit forwarding-options sampling instance instance-name family (inet | inet6) output inline-jflow]`.

> **NOTE**: For MX Series routers with Modular Port Interface Concentrators (MPCs), when `maximum-packet-length` (clip length) is configured for port-mirrored packets and the mirror-destination interface is a next-hop-group, the clip length is effective only for the first member interface of the next-hop-group. The mirrored packet copy sent to the rest of the interfaces is not clipped.
>
> In addition, native analyzer sessions (that is, the `[edit forwarding-options analyzer analyzer-name input]` hierarchy level for MX Series routers) can be configured without specifying input parameters. As such, these instances use the following input values by default: rate = 1, and maximum-packet-length = 0.

Options

bytes—Maximum length (in bytes) of the mirrored packet or the sampled packet.

Set the maximum-packet-length value to zero to disable truncation; that is, to mirror or sample the entire packet. Otherwise, Juniper recommends that you configure the packet length to be equal to, or greater than, the IP header length. For IPv4, set the maximum length to at least 20, and for IPv6, set the maximum length to at least 40.

- **Range:** 0 through 9216. For MX Series routers with MPCs, and for EX9200 switches, the range is 1 through 255 bytes.
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

For MX Series routers except the MX 80, support at the `[edit forwarding-options analyzer analyzer-name input]` hierarchy level was introduced in Junos OS Release 14.1

RELATED DOCUMENTATION

Configuring Port Mirroring

Configuring Traffic Sampling on MX, M and T Series Routers

maximum-paths

IN THIS SECTION

- Syntax | **10542**
- Hierarchy Level | **10542**
- Description | **10543**
- Options | **10543**
- Required Privilege Level | **10544**
- Release Information | **10544**

Syntax

```
maximum-paths path-limit <log-interval seconds> <log-only | threshold value>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
```

```
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Configure a limit for the number of routes installed in a routing table based upon the route path.

> **NOTE**: The `maximum-paths` statement is similar to the `maximum-prefixes` statement. The `maximum-prefixes` statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:
>
> ```
> OSPF 10.10.10.0/24
> ISIS 10.10.10.0/24
> ```
>
> These are two routes, but only one destination (prefix). The `maximum-paths` limit applies the total number of routes (two). The `maximum-prefixes` limit applies to the total number of unique prefixes (one).
>
> Local and direct routes are not considered as part of `maximum-path` and `maximum-prefixes` limits.

Options

`log-interval` *`seconds`*—(Optional) Minimum time interval (in seconds) between log messages.

- **Range:** 5 through 86,400

`log-only`—(Optional) Sets the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.

`path-limit`—Maximum number of routes. If this limit is reached, a warning is triggered and additional routes are rejected.

- **Range:** 1 through 4,294,967,295 ($2^{32} - 1$)
- **Default:** No default

`threshold` *`value`*—(Optional) Percentage of the maximum number of routes that starts triggering a warning. You can configure a percentage of the *`path-limit`* value that starts triggering the warnings.

- **Range:** 1 through 100

> **NOTE**: When the number of routes reaches the `threshold` value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *`path-limit`* value, then additional routes are rejected.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs

maximum-pending-reqs-limit

IN THIS SECTION

- Syntax | 10545
- Hierarchy Level | 10545
- Description | 10545
- Options | 10545
- Required Privilege Level | 10545
- Release Information | 10545

Syntax

```
maximum-pending-reqs-limit number;
```

Hierarchy Level

```
[edit access radius network-element name]
```

Description

Configure the maximum number of requests that can be queued to the network element. When the pending-request queue is full, any additional requests are dropped.

Options

number Maximum number of pending requests.

- **Range:** 512 through 8192

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Network Elements

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

maximum-pending-requests (Diameter)

IN THIS SECTION

- Syntax | 10546
- Hierarchy Level | 10546
- Description | 10546
- Options | 10546
- Required Privilege Level | 10547
- Release Information | 10547

Syntax

```
maximum-pending-requests requests;
```

Hierarchy Level

```
[edit access diameter applications pcc-gx]
```

Description

Configure the maximum number of pending requests parameter for the Diameter application.

Options

requests Maximum number of pending requests.

- **Range:** 1000 through 65,535
- **Default:** 20,000

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

maximum (PIM Entries)

IN THIS SECTION

- Syntax | 10547
- Hierarchy Level | 10548
- Description | 10548
- Options | 10549
- Required Privilege Level | 10549
- Release Information | 10549

Syntax

```
maximum limit;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim sglimit],
[edit logical-systems logical-system-name protocols pim sglimit family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
sglimit],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
sglimit family],
[edit protocols pim sglimit],
[edit protocols pim sglimit family],
[edit routing-instances routing-instance-name protocols pim sglimit],
[edit routing-instances routing-instance-name protocols pim sglimit family],
[edit logical-systems logical-system-name protocols pim rp group-rp-mapping],
[edit logical-systems logical-system-name protocols pim rp group-rp-mapping family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp group-rp-mapping],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp group-rp-mapping family],
[edit protocols pim rp group-rp-mapping],
[edit protocols pim rp group-rp-mapping family],
[edit routing-instances routing-instance-name protocols pim rp group-rp-mapping],
[edit routing-instances routing-instance-name protocols pim rp group-rp-mapping family],
[edit logical-systems logical-system-name protocols pim rp register-limit],
[edit logical-systems logical-system-name protocols pim rp register-limit family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp register-limit],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp register-limit family],
[edit protocols pim rp register-limit],
[edit protocols pim rp register-limit family],
[edit routing-instances routing-instance-name protocols pim rp register-limit],
[edit routing-instances routing-instance-name protocols pim rp register-limit family]
```

Description

Configure the maximum number of specified PIM entries received by the device. If the device reaches the configured limit, no new entries are received.

> **NOTE**: The maximum limit settings that you configure with the **maximum** and the `family (inet | inet6) maximum` statements are mutually exclusive. For example, if you configure a global maximum PIM join state limit, you cannot configure a limit at the family level for IPv4 or IPv6 joins. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.

Options

limit—Maximum number of PIM entries received by the device. If you configure both the **log-interval** and the `maximum` statements, a warning is triggered when the maximum limit is reached.

Depending on your configuration, this limit specifies the maximum number of PIM joins, PIM register messages, or group-to-RP mappings received by the device.

- **Range:** 1 through 65,535

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

clear pim join | 3806

maximum-power (Interface)

IN THIS SECTION

Syntax | 10550

- Hierarchy Level | 10550
- Description | 10550
- Options | 10550
- Required Privilege Level | 10552
- Release Information | 10552

Syntax

```
maximum-power watts;
```

Hierarchy Level

```
[edit poe interface (all | all-extended | interface-name)]
```

Description

Configure the maximum amount of power that the switch can supply to the PoE port. The maximum power configuration is valid when the PoE power management is in static mode. If PoE power management is in class mode, which is the default, the maximum power configuration will have no effect on the power allocation. For more information on power management configuration options, see `management`.

NOTE: Maximum power configuration is not supported in PoE-bt (IEEE 802.3bt).

NOTE: A standalone switch's default setting and range for maximum power does not change if the switch is configured as a satellite device in a Junos Fusion. For instance, an EX4300 switch has a 30W default and a range of 0.0 through 30.0 W when configured as a standalone switch and when it is configured into a satellite device in a Junos Fusion.

Options

watts—The maximum power in watts that can be supplied to the ports.

For EX2200, EX3300, EX4200, EX4300, EX4600, EX6200, and EX8200 switches:

- **Range:** 0.0 through 30.0
- **Default:** 15.4 W for ports that support IEEE 802.3af and 30 W for ports that support IEEE 802.3at

For EX3200 switches:

- **Range:** 0.0 through 18.6
- **Default:** 15.4 W

> **NOTE**: EX4600 switches support PoE only when operating in a mixed Virtual Chassis with EX4300 switches.

For ACX2000 routers:

- **Range:** 1 through 65 W
- **Default:** 32 W

> **NOTE**: The **maximum-power** setting permitted by the CLI might be greater than the maximum power a given PoE port can deliver. For example, the CLI permits you to set any PoE port on an EX8200 line card to 30 W; however, only ports 0 through 11 support 30 W. Similarly, the CLI permits you to set any PoE port on an EX4200 switch to 30 W, but some models of EX4200 switch support only 18.6 W per port. If you configure a **maximum-power** value that is greater than the maximum power supported by a port, the power allocated to the port will be the maximum supported.
>
> If you use the **all** option to set **maximum-power** to a value greater than 15.4 W on all interfaces on an EX8200 line card, the maximum power allocated to all ports is 15.4 W.

> **NOTE**: Support for a maximum of 18.6 W per port instead of 15.4 W per port on EX3200 switches and P and T models of EX4200 switch requires Junos OS Release 11.1 or later. In addition to requiring an upgrade of Junos OS to Release 11.1 or later, switches that are running an earlier release of Junos OS release require the PoE controller software be upgraded as described in Upgrading the PoE Controller Software. If the controller software is not upgraded and you set **maximum-power** to a value greater than 15.4 W, the configuration is accepted when you commit it, but the actual power allocated to the port will be 15.4 W.

> **NOTE**: On ACX2000 routers, the power sourcing equipment (PSE) delivers up to 65 W per port, provided the management mode is set to high-power mode, by using the `high-power` option at the [`edit poe management`] hierarchy level. By default, the management mode is set to `static`. In the static mode, the PSE can deliver power up to 32 W.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring PoE Interfaces on EX Series Switches

Configuring Power over Ethernet in a Junos Fusion

maximum-power (PoE)

IN THIS SECTION

- Syntax | 10553
- Hierarchy Level | 10553
- Description | 10553
- Default | 10553
- Options | 10553
- Required Privilege Level | 10553
- Release Information | 10553

Syntax

```
maximum-power watts;
```

Hierarchy Level

```
[edit poe interface (all | interface-name)]
```

Description

Maximum amount of power that can be supplied to the port.

Default

15.4 W

Options

`Watts`—The maximum number of watts that can be supplied to the port.

`Range` —0 through 15.4

`Default`—15.4 W

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

maximum-prefixes

IN THIS SECTION

- Syntax | 10554
- Hierarchy Level | 10554
- Description | 10554
- Options | 10555
- Required Privilege Level | 10556
- Release Information | 10556

Syntax

```
maximum-prefixes prefix-limit <log-interval seconds> <log-only | threshold percentage>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Configure a limit for the number of routes installed in a routing table based upon the route prefix.

Using a prefix limit, you can curtail the number of prefixes received from a CE router in a VPN. Prefix limits apply only to dynamic routing protocols and are not applicable to static, programmed, or interface routes.

NOTE: Static, programmed [that is, programmable routing protocol process (prpd) configured], and interface routes contribute to the overall route total count used in the

maximum-prefix calculation. The number of such routes affects the number of dynamic routes that are accepted in to the virtual routing instance. In the configuration example shown below, if `VRF1` contains 98 static routes and 2 interface routes, then a maximum of 900 BGP routes (that is, 1000 - 98 - 2 = 900) are allowed, with a syslog warning shown at a total of 800 dynamic routes, which is 80%.

```
set routing-instances VRF1 routing-options maximum-prefixes 1000
set routing-instances VRF1 routing-options maximum-prefixes threshold 80
```

If the number of static, programmed, or interface routes increases, then it can affect the maximum number of dynamic routes that are accepted. Note that the static, programmed, or interface routes are not subject to the configured maximum-prefixes limit and can have a negative or undesired effect on the dynamic routes.

The `maximum-prefixes` statement is similar to the `maximum-paths` statement. The `maximum-prefixes` statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The `maximum-paths` limit applies the total number of routes (two). The `maximum-prefixes` limit applies to the total number of unique prefixes (one).

Options

`log-interval` *`seconds`*—(Optional) Minimum time interval (in seconds) between log messages.

- **Range:** 5 through 86,400

`log-only`—(Optional) Sets the prefix limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.

`prefix-limit`—Maximum number of route prefixes. If this limit is reached, a warning is triggered and any additional routes are rejected.

- **Range:** 1 through 4,294,967,295
- **Default:** No default

`threshold` *`value`*—(Optional) Percentage of the maximum number of prefixes that starts triggering a warning. You can configure a percentage of the *`prefix-limit`* value that starts triggering the warnings.

- **Range:** 1 through 100

> **NOTE**: When the number of routes reaches the `threshold` value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *`prefix-limit`* value, then additional routes are rejected.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs

maximum-rps

IN THIS SECTION

- Syntax | 10557
- Hierarchy Level | 10557
- Description | 10557
- Options | 10557
- Required Privilege Level | 10557
- Release Information | 10557

Syntax

```
maximum-rps limit;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp embedded-rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp embedded-rp],
[edit protocols pim rp embedded-rp],
[edit routing-instances routing-instance-name protocols pim rp embedded-rp]
```

Description

Limit the number of RPs that the routing device acknowledges.

Options

limit—Number of RPs.

- **Range:** 1 through 500
- **Default:** 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring PIM Embedded RP for IPv6

maximum-send-window

IN THIS SECTION

- Syntax | 10558
- Hierarchy Level | 10558
- Description | 10558
- Options | 10558
- Required Privilege Level | 10559
- Release Information | 10559

Syntax

```
maximum-send-window packets;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Specify the size of the send window for L2TP tunnels, which limits the remote end's receive window size.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options

packets—Maximum number of packets the send window can hold at one time.

- **Default:** 32

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Window Size for L2TP Tunnels

maximum-sessions

IN THIS SECTION

- Syntax | 10559
- Hierarchy Level | 10559
- Description | 10560
- Options | 10560
- Required Privilege Level | 10560
- Release Information | 10560

Syntax

```
maximum-sessions count;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Configure the maximum number of allowed test sessions the server can have running at one time.

> **NOTE**: The maximum number of test sessions running on the server at one time must be greater than or equal to the maximum number of sessions the server can open on a single client connection.

Options

count—Maximum number of sessions.

- **Range:** 1 through 2048
- **Default:** 64

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

maximum-sessions (L2TP)

IN THIS SECTION

Syntax | **10561**
Hierarchy Level | **10561**

- Description | 10561
- Options | 10561
- Required Privilege Level | 10561
- Release Information | 10562

Syntax

```
maximum-sessions number;
```

Hierarchy Level

```
[edit access profile profile-name client client-name],
[edit services l2tp],
[edit services l2tp sessions-limit-group],
[edit services l2tp tunnel],
[edit services l2tp tunnel-group group-name],
```

Description

Specify the maximum number of L2TP sessions for the chassis, all tunnels, a tunnel group, a session limit group, or a client.

Options

number—Number of sessions allowed.

- **Range:** (Chassis, tunnel group, session limit group, or client) 1 through the default maximum chassis limit
- **Range:** (Tunnel) 1 through 65,536

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *L2TP LAC Tunneling for Subscribers*
- *L2TP LAC Subscriber Configuration*
- *L2TP LNS Inline Service Interfaces*
- *L2TP LNS Inline Service Interfaces*
- *L2TP for Subscriber Access Overview*

maximum-sessions-per-connection

IN THIS SECTION

- Syntax | 10562
- Hierarchy Level | 10562
- Description | 10563
- Options | 10563
- Required Privilege Level | 10563
- Release Information | 10563

Syntax

```
maximum-sessions-per-connection count;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Configure the maximum number of allowed sessions the server can open on a single client connection.

> **NOTE**: The maximum number of test sessions running on the server at one time must be greater than or equal to the maximum number of sessions the server can open on a single client connection.

Options

count—Maximum number of sessions.

- **Default:** 64
- **Range:** `1-1024`

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

maximum-sessions-per-tunnel

IN THIS SECTION

Syntax | 10564
Hierarchy Level | 10564

- Description | 10564
- Options | 10564
- Required Privilege Level | 10564
- Release Information | 10564

Syntax

```
maximum-sessions-per-tunnel number;
```

Hierarchy Level

```
[edit access group-profile l2tp],
[edit access profile profile-name client client-name l2tp]
```

Description

Configure the maximum sessions for a Layer 2 tunnel.

NOTE: This statement is not supported at the `[edit access group-profile l2tp]` hierarchy level for L2TP LNS on MX Series routers.

Options

number—Maximum number of sessions for a Layer 2 tunnel.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Layer 2 Tunneling Protocol (L2TP)

L2TP LNS Inline Service Interfaces

maximum-sessions (TDF Gateway)

IN THIS SECTION

- Syntax | 10565
- Hierarchy Level | 10565
- Description | 10565
- Options | 10565
- Required Privilege Level | 10566
- Release Information | 10566

Syntax

```
maximum-sessions max-sessions;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name cac]
```

Description

Configure the maximum number of TDF subscriber sessions that may be running.

Options

max-sessions Maximum number of TDF subscriber sessions, expressed in thousands.

- **Range:** 10 thousands through 5000 thousands

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a TDF Gateway

maximum-sessions-trap-percentage (TDF Gateway)

IN THIS SECTION

- Syntax | 10566
- Hierarchy Level | 10567
- Description | 10567
- Options | 10567
- Required Privilege Level | 10567
- Release Information | 10567

Syntax

```
maximum-sessions-trap-percentage max-sessions-pct;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name cac]
```

Description

Configure the trap threshold for the number of TDF subscriber sessions as a percentage of the maximum number of sessions (`maximum-sessions`) that was configured at the `[edit unified-edge gateways tdf gateway-name cac]` hierarchy level. If the number of subscriber sessions reaches the threshold, the SNMP trap **jnxScgSMSessionThreshHigh** is generated.

Options

max-sessions-pct Percentage of the maximum number of TDF subscriber sessions.

- **Range:** 1 through 90

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a TDF Gateway

maximum-subscribers

IN THIS SECTION

- Syntax | 10568
- Hierarchy Level | 10568
- Description | 10568
- Options | 10568
- Required Privilege Level | 10568
- Release Information | 10569

Syntax

```
maximum-subscribers number;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
```

Description

Specify the maximum number of IP-based subscriber sessions that the TDF domain can support.

Options

number Maximum number of subscriber sessions allowed.

- **Range:** 100 thousands through 5000 thousands.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

maximum-subscribers

IN THIS SECTION

- Syntax | 10569
- Hierarchy Level | 10569
- Description | 10570
- Options | 10570
- Required Privilege Level | 10570
- Release Information | 10570

Syntax

```
maximum-subscribers limit;
```

Hierarchy Level

```
[edit system services extensible-subscriber-services]
```

Description

Configure the maximum number of subscriber sessions supported at a time.

Options

limit Maximum number of subscribers.

- **Range:** 1 through 2000
- **Default:** 1000

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

show subscribers | 26609

show subscribers summary | 26663

maximum-transactions

IN THIS SECTION

- Syntax | 10571
- Hierarchy Level | 10571
- Description | 10571
- Options | 10571
- Required Privilege Level | 10572

Release Information | 10572

Syntax

```
maximum-transactions transactions-number;
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Configure the maximum number of transactions matched by application identification for finalizing the application.

Application classification does not terminate for applications that are transaction based such as Facebook applications. To terminate the application classifications for such applications, you can choose to consider the results from multiple transaction as the final classification. You can configure the number of transactions before concluding the final result for the identified application.

For example, when you configure the maximum number of transactions as 10, the following sequence is applied for identifying the final application:

- In the first and second transactions, application-1 and application-2 are identified respectively.
- The identification process continues till the 10th transaction is reached.
- Since 10th transaction is equal to the configured value of the maximum number of transactions, the application identified in this transaction is considered as the final match.

Options

maximum-transactions *transactions-number*

Number of transaction results that can be considered before concluding the final result for application identification.

- **Range:** 0 through 25

- **Default:** 5

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1

RELATED DOCUMENTATION

Application Identification

maximum-transmit-rate (Protocols IGMP)

IN THIS SECTION

- Syntax | **10572**
- Hierarchy Level | **10573**
- Description | **10573**
- Options | **10573**
- Required Privilege Level | **10573**
- Release Information | **10573**

Syntax

```
maximum-transmit-rate packets-per-second;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp],
[edit protocols igmp]
```

Description

Limit the transmission rate of IGMP packets

Options

packets-per-second—Maximum number of IGMP packets transmitted in one second by the routing device.

- **Range:** 1 through 10000
- **Default:** 500 packets

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Limiting the Maximum IGMP Message Rate

maximum-transmit-rate (Protocols MLD)

IN THIS SECTION

- Syntax | 10574
- Hierarchy Level | 10574
- Description | 10574
- Options | 10574
- Required Privilege Level | 10575
- Release Information | 10575

Syntax

```
maximum-transmit-rate packets-per-second;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld],
[edit protocols mld]
```

Description

Limit the transmission rate of MLD packets.

Options

packets-per-second—Maximum number of MLD packets transmitted in one second by the routing device.

- **Range:** 1 through 10000
- **Default:** 500 packets

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Limiting the Maximum MLD Message Rate

mc-ae

IN THIS SECTION

- Syntax | 10575
- Hierarchy Level | 10576
- Description | 10576
- Options | 10576
- Required Privilege Level | 10580
- Release Information | 10580

Syntax

```
mc-ae {
    chassis-id chassis-id;
        events {
            iccp-peer-down;
                force-icl-down;
            prefer-status-control-active;
            }
```

```
            init-delay-time seconds;
            mc-ae-id mc-ae-id;
            mode (active-active | active-standby);
            redundancy-group group-id;
            revert-time revert-time;
            status-control (active | standby);
            switchover-mode (non-revertive |revertive);
        }
```

Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options],
[edit logical-systems logical-system-name interfaces aeX aggregated-ether-options]
```

Description

Enable multichassis link aggregation groups (MC-LAG), which enables one device to form a logical LAG interface with two or more other devices.

Options

chassis-id Specify the chassis ID for Link Aggregation Control Protocol (LACP) to calculate the port number of MC-LAG physical member links. Each MC-LAG peer should have a unique chassis ID.

- **Values:** 0 or 1

events Specify an action if a specific MC-LAG event occurs.

iccp-peer-down Specify an action if the ICCP peer of this node goes down.

force-icl-down If the node's ICCP peer goes down, bring down the interchassis-link logical interface.

prefer-status-control-active Specify that the node configured as `status-control active` become the active node if the peer of this node goes down.

When ICCP goes down, you can use this keyword to make a mc-lag PE to become the active PE. For example, if you want mc-lag PE1 to be Active on ICCP down, then configure this keyword in PE1. It is not recommended to configure this keyword in both the mc-lag PEs.

> **NOTE**: The **prefer-status-control-active** statement can be configured with the **status-control standby** configuration to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure. Use this configuration only if you can ensure that ICCP will not go down unless the router or switch is down. You must also configure the **hold-time down** value (at the **[edit interfaces *interface-name*]** hierarchy level) for the interchassis link with the **status-control standby** configuration to be higher than the ICCP BFD timeout. This configuration prevents data traffic loss by ensuring that when the router or switch with the **status-control active** configuration goes down, the router or switch with the **status-control standby** configuration does not go into standby mode.
>
> To make the **prefer-status-control-active** configuration work with the **status-control standby** configuration when an interchassis-link logical interface is configured on aggregate Ethernet interface, you must either configure the **lacp periodic *interval*** statement at the **[edit interface *interface-name* aggregated-ether-options]** hierarchy level as `slow` or configure the **detection-time threshold** statement at the **[edit protocols iccp peer liveness-detection]** hierarchy level as less than 3 seconds.

`init-delay-time` *seconds*

To minimize traffic loss, specify the number of seconds in which to delay bringing the multichassis aggregated Ethernet interface back to the up state when you reboot an MC-LAG peer. By delaying the startup of the interface until after protocol convergence, you can prevent packet loss during the recovery of failed links and devices.

> **NOTE**: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

`mc-ae-id` *mc-ae-id*

Specify the identification number of the MC-LAG device. The two MC-LAG network devices that manage a given MC-LAG must have the same identification number.

- **Range:** 1 through 65,535

mode (active-active | active-standby)

Specify whether the MC-LAG is in active-active or active-standby mode. Chassis that are in the same group must be in the same mode.

> **NOTE**: You can configure IPv4 (inet) and IPv6 (inet6) addresses on `mc-ae` interfaces when the `active-standby` mode is configured.

In active-active mode, all member links are active on the MC-LAG. In this mode, media access control (MAC) addresses learned on one MC-LAG peer are propagated to the other MC-LAG peer. Active-active mode is a simple and deterministic design and is easier to troubleshoot than active-standby mode.

> **NOTE**: Active-active mode is not supported on Dense Port Concentrator (DPC) line cards. Instead, use active-standby mode.

In active-active MC-LAG topologies, network interfaces are categorized into three interface types, as follows:

- S-Link—Single-homed link (S-Link) terminating on an MC-LAG peer device
- MC-Link—MC-LAG link
- ICL—Inter-chassis link
- Mode

  Indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

  In active-active mode, all member links are active on the MC-LAG. In this mode, media access control (MAC) addresses learned on one MC-LAG peer are propagated to the other MC-LAG peer. Active-active mode is a simple and deterministic design and is easier to troubleshoot than active-standby mode.

  > **NOTE**: Active-active mode is not supported on Dense Port Concentrator (DPC) line cards. Instead, use active-standby mode.

  Depending on the incoming and outgoing interface types, some constraints are added to the Layer 2 forwarding rules for MC-LAG configurations. The following data traffic forwarding rules apply.

> **NOTE**: If only one MC-LAG member link is in the up state, it is considered an S-Link.

- When an MC-LAG network receives a packet from a local MC-Link or S-Link, the packet is forwarded to other local interfaces, including S-Links and MC-Links based on the normal Layer 2 forwarding rules and on the configuration of the `mesh-group` and `no-local-switching` statements. If MC-Links and S-Links are in the same mesh group and their `no-local-switching` statements are enabled, the received packets are only forwarded upstream and not sent to MC-Links and S-Links.
- The following circumstances determine whether or not an ICL receives a packet from a local MC-Link or S-Link:
  - If the peer MC-LAG network device has S-Links or MC-LAGs that do not reside on the local MC-LAG network device
  - Whether or not interfaces on two peering MC-LAG network devices are allowed to talk to each other
- When an MC-LAG network receives a packet from the ICL, the packet is forwarded to all local S-Links and active MC-LAGs that do not exist in the MC-LAG network from which the packet was sent.

In active-standby mode, only one of the MC-LAG peers is active at any given time. The other MC-LAG peer is in backup (standby) mode. The active MC-LAG peer uses Link Aggregation Control Protocol (LACP) to advertise to client devices that its child link is available for forwarding data traffic. Active-standby mode should be used if you are interested in redundancy only. If you require both redundancy and load sharing across member links, use active-active mode.

> **NOTE**: Active-standby mode is not supported on EX4300 and QFX Series switches.

redundancy-group *group-id*

Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions. The value you assign to this option must match the value you assign to the `redundancy-group-id-list` option that you configure on the ICCP peer. If the value differs, when you commit the changes, the system reports a commit check error.

> **BEST PRACTICE**: We recommend that you configure only one redundancy group between MC-LAG nodes. The redundancy group represents the domain of high availability between the MC-LAG nodes. One redundancy group is sufficient between a pair of MC-LAG nodes. If you are using logical systems, then configure one redundancy group between MC-LAG nodes in each logical system.

- **Range:** 1 through 4,294,967,294

revert-time

Wait interval (in minutes) before the switchover to the preferred node is performed when the `switchover-mode` is configured as revertive.

- **Range:** 1 through 10

status-control (active | standby)

Specify whether the chassis becomes active or remains in standby mode when an interchassis link failure occurs.

- Events ICCP-Peer-Down Force-ICL-Down

  Forces the ICL to be down if the peer of this node goes down.

- Events ICCP-Peer-Down Prefer-Status-Control-Active

  Allows the LACP system ID to be retained during a reboot, which provides better convergence after a failover.

switchover-mode (non-revertive | revertive)

Specify whether Junos OS should trigger a link switchover to the preferred node when the active node is available.

> **NOTE**: For `revertive` mode to automatically switch over to the preferred node, the `status-control` statement should be configured as `active`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

`events` statement introduced in Junos OS Release 11.4R4 for MX Series routers.

`prefer-status-control-active` statement introduced in Junos OS Release 13.2R1 for EX Series switches.

`init-delay-time` *seconds* statement introduced in Junos OS Release 13.2R3 for EX Series switches.

`switchover-mode` and `revert-time` statements introduced in Junos OS Release 13.3.

Support for logical systems introduced in Junos OS Release 14.1.

mc-ae-id

IN THIS SECTION

- Syntax | 10581
- Hierarchy Level | 10581
- Description | 10581
- Options | 10582
- Required Privilege Level | 10582
- Release Information | 10582

Syntax

```
mc-ae-id mc-ae-id;
```

Hierarchy Level

```
[edit interfaces aggregated-ether-options mc-ae]
```

Description

Specify the multichassis aggregated Ethernet (MC-AE) identification number of the MC-AE that a given aggregated Ethernet interface belongs to. The two peers that host a given multichassis link aggregation group (MC-LAG) must have the same multichassis aggregated Ethernet ID.

Options

- **Range:** 1 through 65535.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

mcast stats-enable

IN THIS SECTION

- Syntax | 10582
- Hierarchy Level | 10583
- Description | 10583
- Required Privilege Level | 10583
- Release Information | 10583

Syntax

```
mcast {
    stats-enable;
}
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

Enable L3 Multicast Statistics to track the total number of packets and bytes that hit a multicast route.

Statistics are attached and collected on a FCFS (first-come, first-served) basis for the first 16K (S,G) routes that require statistics. The routes can be a combination of IPv4 and IPv6 (S,G) routes. Both ASM and SSM routes are supported. Statistics for each multicast (S,G) route is calculated independently.

> **NOTE**: Multicast statistics can be attached only to (S,G) multicast routes and not (*,G) routes.

When you commit the `set system packet-forwarding-options mcast stats-enable` configuration statement, a PFE restart is initiated.

To disable multicast statistics, use the configuration statement `delete system packet-forwarding-options mcast stats-enable`.

To view multicast statistics, use the `show multicast route extensive` command.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

show multicast route | 22306

mclag

IN THIS SECTION

- Syntax | 10584
- Hierarchy Level | 10584
- Description | 10584
- Required Privilege Level | 10584
- Release Information | 10585

Syntax

```
mclag {
    bgp-peer ip-address;
}
```

Hierarchy Level

```
[edit routing-instances name protocols evpn]
```

Description

Configure parameters that enable the interworking of Ethernet VPN-MPLS (EVPN-MPLS) with a Junos Fusion Enterprise or a multichassis link aggregation group (MC-LAG) topology.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG

mclag-arp-nd-sync

IN THIS SECTION

- Syntax | 10585
- Hierarchy Level | 10585
- Description | 10585
- Required Privilege Level | 10586
- Release Information | 10586

Syntax

```
mclag-arp-nd-sync;
```

Hierarchy Level

```
[edit logical-systems name protocols l2-learning],
[edit protocols l2-learning]
```

Description

When MC-LAG peers power up or reboot, Address Resolution Protocol (ARP) and Neighbor Discovery (ND) entries are learned and synchronized by the peers. Because the MC-LAG peers synchronize and

learn each other's ARP and ND entries, relearning these entries is not needed on the device. Only the ARP and ND entries that are pointing to either an MC-AE or corresponding ICL are synchronized.

This same behavior is applicable when an IRB comes up and goes down. In this case, only those ARP and ND entries corresponding to an IRB that is part of an MC-LAG (for example, bridge domain, or VLAN) are synchronized by the peers.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Understanding Multichassis Link Aggregation Groups

mclag-arpreq-sync

IN THIS SECTION

- Syntax | 10586
- Hierarchy Level | 10587
- Description | 10587
- Required Privilege Level | 10587
- Release Information | 10587

Syntax

```
mclag-arpreq-sync;
```

Hierarchy Level

```
[edit logical-systems name protocols l2-learning],
[edit protocols l2-learning]
```

Description

In an active-active MC-LAG configuration, Address Resolution Protocol (ARP) requests are synchronized between the MC-LAG peers.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Understanding Multichassis Link Aggregation Groups

mdi-mode

IN THIS SECTION

- Syntax | 10588
- Hierarchy Level | 10588
- Description | 10588
- Options | 10589
- Required Privilege Level | 10589
- Release Information | 10589

Syntax

```
mdi-mode mode;
```

Hierarchy Level

```
[edit interfaces interface-name ether-options],
[edit interfaces interface-range range ether-options]
```

Description

You must configure media dependent interface (MDI) properties for a 10-Gigabit Ethernet interface on a copper network port of an EX4550 switch to ensure that both sides of the link are compatible.

MDI refers to the IEEE standard for the interface to an unshielded twisted pair (UTP) cable. Twisted-pair Ethernet standards are such that the majority of cables can be wired "straight-through" (pin 1 to pin 1, pin 2 to pin 2 and so on), but others may need to be wired in the "crossover" form (receive to transmit and transmit to receive).

For most ports, the switch can automatically detect the required connection type and can therefore configure the interface appropriately. However, the switch cannot automatically detect whether the connection type of a 10-Gigabit Ethenet interface on a copper network port is straight-through or crossover.

Therefore, you must set the MDI properties of the local interface of a 10-Gigabit Ethernet interface on a copper network port to ensure that it will work correctly with the other side of the link. When you set this configuration on an interface, you must also disable "auto-negotiation" on page 2719 and set the "speed" on page 28349 to 100m.

> **NOTE**: This configuration does not apply to management Ethernet or console interfaces and it does not apply to 1-Gigabit copper ports.

The proper setting depends both on the type of cable and the setting that is being used on the other side of the link:

- For crossover cables—Set the polarity to match the other side of the link. Specify `mdi` for the switch interface if `mdi` is being used on the other side of the link; specify `mdix` if `mdix` is being used on the other side of the link.

- For straight cables—Set the polarity to be the opposite of the other side link. Specify `mdi` for the switch interface if `mdix` is being used on the other side of the link; specify `mdix` if `mdi` is being used on the other side of the link.

Options

One of the following modes:

`auto` Set the MDI properties to automatic. This setting should *not* be used with 10-Gigabit Ethernet interfaces on a copper network port of an EX4550 switch.

`mdi` Set the MDI properties of the interface to straight through mode. The selection of the mode depends on whether crossover or straight cables are being used and on the setting used on the other side of the link.

`mdix` Set the MDI properties of the interface to crossover mode. The selection of the mode depends on whether crossover or straight cables are being used and on the setting used on the other side of the link.

`force` For SRX Series Firewalls this option enables the MDI properties to auto-mdix always.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

Support for SRX Series Firewalls introduced in Junos OS Release 15.1x49.D80.

RELATED DOCUMENTATION

- Configuring Gigabit Ethernet Interfaces (CLI Procedure)
- *Interfaces Overview for Switches*
- Junos OS Ethernet Interfaces Configuration Guide
- Interfaces User Guide for Security Devices

mdt

IN THIS SECTION

- Syntax | 10590
- Hierarchy Level | 10590
- Description | 10591
- Required Privilege Level | 10591
- Release Information | 10591

Syntax

```
mdt {
    data-mdt-reuse;
    group-range multicast-prefix;
    threshold {
        group group-address {
            source source-address {
                rate threshold-rate;
            }
        }
        tunnel-limit limit;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel family inet | inet6],
[edit routing-instances routing-instance-name protocols pim],
[edit routing-instances routing-instance-name provider-tunnel family inet | inet6]
```

Description

Establish the group address range for data MDTs, the threshold for the creation of data MDTs, and tunnel limits for a multicast group and source. A multicast group can have more than one source of traffic.

The remaining statements are explained separately. .

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4. In Junos OS Release 17.3R1, the `mdt` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default MDT in Rosen 7, and data MDT for Rosen 6 and Rosen 7. The `provider-tunnel mdt` hierarchy is now hidden for backward compatibility with existing scripts.

RELATED DOCUMENTATION

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

measurement-interval

IN THIS SECTION

- Syntax | 10592
- Hierarchy Level | 10592
- Description | 10592
- Default | 10592
- Required Privilege Level | 10592
- Release Information | 10592

Syntax

```
measurement-interval (5|15|30|60)
```

Hierarchy Level

```
[edit protocols oam ethernet cfm performance-monitoring]
```

```
[edit protocols oam ethernet cfm performance-monitoring sla-iterator-profiles profile-name]
```

Description

Configure measurement interval to be used for a performance monitoring session. You must configure the measurement-interval at the `[edit protocols oam ethernet cfm performance-monitoring]` hierarchy level, which is a global level parameter. You can override the configured value by specifying a measurement-interval for the iterator profile at the `[edit protocols oam ethernet cfm performance-monitoring sla-taerator-profiles profile-name]` hierarchy level.

> **NOTE**: When you configure when MEF-36-compliant performance monitoring, you must also configure an `enhanced-sla-iterator` at the `[edit protocols oam ethernet cfm performance-monitoring]` hierarchy level.

Default

15 minutes

Required Privilege Level

configure—To enter configuration mode.

control—To modify any configuration.

Release Information

RELATED DOCUMENTATION

IEEE 802.1ag OAM Connectivity Fault Management Overview

Junos OS Support for Performance Monitoring Compliant with Technical Specification MEF 36

measurement-type

IN THIS SECTION

- Syntax | 10593
- Hierarchy Level | 10593
- Description | 10593
- Options | 10593
- Required Privilege Level | 10594
- Release Information | 10594

Syntax

```
measurement-type  (loss | statistical-loss-measurement | two-way-delay);
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
profile-name]
```

Description

Configure the measurement type for the service level agreement (SLA) frames. An SLA frame is a type of packet used to measure frame loss in Ethernet connections.

Options

`loss`—Use Y.1731-compliant line module (LM) frames to measure frame loss.

`statistical-loss-measurement`— Use Y.1731-compliant two-way data module (DM) frames to statistically measure frame loss.

`two-way-delay`—Use Y.1731-compliant two-way DM frames to measure frame loss.

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `statistical-loss-measurement` option introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring an Iterator Profile

med-igp-update-interval

IN THIS SECTION

- Syntax | 10594
- Hierarchy Level | 10595
- Description | 10595
- Options | 10595
- Required Privilege Level | 10595
- Release Information | 10595

Syntax

```
med-igp-update-interval minutes;
```

Hierarchy Level

```
[edit routing-options]
```

Description

Configure a timer for how long to delay updates for the multiple exit discriminator (MED) path attribute for BGP groups and peers configured with the `metric-out igp` *`offset`* `delay-med-update` statement. The timer delays MED updates for the interval configured unless the MED is lower than the previously advertised attribute or another attribute associated with the route has changed or if the BGP peer is responding to a refresh route request.

Options

minutes—Interval to delay MED updates.

- **Range:** 10 through 600
- **Default:** 10 minutes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0

RELATED DOCUMENTATION

> *Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates*
>
> metric-out | **10676**

media-loss-rate

IN THIS SECTION

- Syntax | 10596
- Hierarchy Level | 10596
- Description | 10596
- Required Privilege Level | 10597
- Release Information | 10597

Syntax

```
media-loss-rate {
    no-syslog-generation;
    generate-snmp-traps;
    storm-control {
        count number;
        interval number;
    }
    alarm-mode {
        immediate;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the media loss rate. The media loss rate is the number of media packets lost over a configurable time interval (interval-duration) where the flow packets are packets carrying streaming application information. A single IP packet can contain zero or more streaming packets.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Inline Video Monitoring on MX Series Routers

alarms (Video Monitoring) | 2125

media-rate-variation

IN THIS SECTION

- Syntax | 10597
- Hierarchy Level | 10598
- Description | 10598
- Required Privilege Level | 10598
- Release Information | 10598

Syntax

```
media-rate-variation {
    no-syslog-generation;
    generate-snmp-traps;
    storm-control {
        count number;
        interval number;
    }
```

```
    alarm-mode {
        mdi-records-count number;
        average;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the media rate variation. The media rate variation is the difference between the expected packet rate and actual packet rate expressed as a percentage of the expected packet rate.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Inline Video Monitoring on MX Series Routers

alarms (Video Monitoring) | **2125**

media-source-port-any

IN THIS SECTION

- Syntax | 10599
- Hierarchy Level | 10599
- Description | 10599
- Required Privilege Level | 10599
- Release Information | 10600

Syntax

```
media-source-port-any;
```

Hierarchy Level

```
[edit logical-systems name security alg h323],
[edit logical-systems name tenants name security alg h323],
[edit security alg h323],
[edit services alg h323],
[edit tenants name security alg h323]
```

Description

Configure media-source-port-any to allow media traffic from any port number. By default, this feature is disabled, which allows a temporary opening in the firewall (pinhole) to open media traffic.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

h323 | 7961

media-type (Dual-Purpose Uplink Ports)

IN THIS SECTION

- Syntax | 10600
- Hierarchy Level | 10600
- Description | 10601
- Default | 10601
- Options | 10601
- Required Privilege Level | 10601
- Release Information | 10601

Syntax

```
media-type (copper | fiber);
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

(EX2200-C switch only) Configure the media type for a dual-purpose uplink port (one RJ-45 port and one SFP port) on an EX2200 switch. If you use the media-type for a dual-purpose uplink port, the alternate media type cannot be used with the port.

(ACX1000 routers only) Select the media type (copper or fiber) for the 1-Gigabit Ethernet interfaces.

Default

When **media-type** is not set, the port accepts either type of connection. The default media type is fiber.

Options

copper The dual-purpose uplink port accepts only a 10/100/1000BASE-T copper connection.

fiber The dual-purpose uplink port accepts only an SFP fiber connection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring the Media Type on Dual-Purpose Uplink Ports

media-type (Interfaces)

IN THIS SECTION

- Syntax | 10602
- Hierarchy Level | 10602
- Description | 10602
- Options | 10602
- Required Privilege Level | 10602
- Release Information | 10603

Syntax

```
media-type
```

Hierarchy Level

```
[edit interfaces interface-name media-type]
```

Description

Configure the operating modes for the 2-Port 10 Gigabit Ethernet XPIM.

Options

- copper
- fiber

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Interfaces

medium (Tunnel Profile)

IN THIS SECTION

- Syntax | 10603
- Hierarchy Level | 10603
- Description | 10603
- Default | 10604
- Options | 10604
- Required Privilege Level | 10604
- Release Information | 10604

Syntax

```
medium type;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Description

Specify the medium type for the tunnel.

Default

`ipv4`

Options

type—Medium type for the tunnel. The only value currently available is `ipv4`.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

member

IN THIS SECTION

- Syntax | 10605
- Hierarchy Level | 10605
- Description | 10605
- Default | 10605
- Options | 10605
- Required Privilege Level | 10606
- Release Information | 10606

Syntax

```
member member-id {
    fabric-tree-root;
    location location;
    mastership-priority number;
    no-management-vlan;
    serial-number serial-number;
    role role;
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Configure a switch or an XRE200 External Routing Engine as a member of a Virtual Chassis or a Virtual Chassis Fabric (VCF) with characteristics specified by the available options.

Default

When an EX Series switch or a QFX Series devices configured in standalone mode is powered on but not interconnected through its Virtual Chassis ports (VCPs) with other member switches, its default member ID is `0`.

There is no default member ID in an EX8200 or EX9200 Virtual Chassis. An EX8200 or EX9200 Virtual Chassis must be preprovisioned, and that process configures the member IDs.

Options

member-id—Identifies a specific member switch of a Virtual Chassis or VCF configuration.

The exact range for a specific Virtual Chassis or VCF depends on the number of switches allowed in the Virtual Chassis or VCF.

In an EX8200 Virtual Chassis, member IDs 0 through 7 are reserved for EX8200 member switches and member IDs 8 and 9 are reserved for the primary and backup external Routing Engines.

The remaining statement options set characteristics of the Virtual Chassis or VCF member, and are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- *Adding a Device to a Virtual Chassis Fabric*
- Configuring an EX4650 or a QFX Series Virtual Chassis
- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis
- Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis
- Configuring an EX9200 Virtual Chassis
- Configuring an EX4650 or a QFX Series Virtual Chassis

member (Application Identification)

IN THIS SECTION

- Syntax | **10607**
- Hierarchy Level | **10607**
- Description | **10607**
- Options | **10607**
- Required Privilege Level | **10607**
- Release Information | **10607**

Syntax

```
[member member-name];
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type
signature l4-l7-signature-name]
```

Description

Define a member name for a custom application definition. Custom definitions can contain multiple members that define attributes for an application. You can define a maximum of four member names.

Options

member-name Name of a member for a custom application definition. You can define a maximum of four member names.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

member-failure-options (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 10608
- Hierarchy Level | 10608
- Description | 10608
- Default | 10610
- Required Privilege Level | 10610
- Release Information | 10610

Syntax

```
member-failure-options {
    drop-member-traffic {
        rejoin-timeout rejoin-timeout;
    }
    redistribute-all-traffic {
        enable-rejoin;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options]
```

Description

Configure the possible behavior for the aggregated Multiservices (AMS) interface in case of failure of more than one active member.

> **NOTE**: The `drop-member-traffic` configuration and the `redistribute-all-traffic` configuration are mutually exclusive.

Table 60 on page 10609 displays the behavior of the member interface after the failure of the first services PIC. Table 61 on page 10609 displays the behavior of the member interface after the failure of two services PICs.

> **NOTE**: The AMS infrastructure has been designed to handle one failure automatically. However, in the unlikely event that more than one services PIC fails, the AMS infrastructure provides configuration options to minimize the impact on existing traffic flows.

Table 60: Behavior of Member Interface After One Multiservices PIC Fails

High Availability Mode	Member Interface Behavior
Many-to-one (N:1) high availability support for service applications	Automatically handled by the AMS infrastructure

Table 61: Behavior of Member Interface After Two Multiservices PICs Fail

High Availability Mode	Configuration	`rejoin-timeout`	Behavior when member rejoins before `rejoin-timeout` expires	Behavior when member rejoins after `rejoin-timeout` expires
Many-to-one (N:1) high availability support for service applications	`drop-member-traffic`	Configured	The existing traffic for the second failed member will *not* be redistributed to the other members. The first member to rejoin becomes an active member. The second member to rejoin becomes the backup. This behavior is handled automatically by the AMS infrastructure.	The existing traffic for the second failed member will *not* be redistributed to the other members. The first member will rejoin the AMS automatically. However, the other members who are rejoining will be moved to the discard state.

Table 61: Behavior of Member Interface After Two Multiservices PICs Fail *(Continued)*

High Availability Mode	Configuration	`rejoin-timeout`	Behavior when member rejoins before `rejoin-timeout` expires	Behavior when member rejoins after `rejoin-timeout` expires
Many-to-one (N:1) high availability support for service applications	`redistribute-all-traffic`	Not applicable	Before rejoin, the traffic is redistributed to existing active members. After a failed member rejoins, the traffic is load-balanced afresh. This may impact existing traffic flows.	

The remaining statements are explained separately. See CLI Explorer.

Default

If `member-failure-options` are not configured, then the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

load-balancing-options (Aggregated Multiservices) | **9888**

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

member-id (Junos Fusion Satellite Device Cluster)

IN THIS SECTION

- Syntax | 10611
- Hierarchy Level | 10611
- Description | 10611
- Default | 10611
- Options | 10611
- Required Privilege Level | 10612
- Release Information | 10612

Syntax

```
member-id member-ID-number;
```

Hierarchy Level

```
[edit chassis satellite-management cluster cluster-name fpc slot-id ]
```

Description

Assign a member ID number to a satellite device in a satellite device cluster.

The member ID is used to identify the satellite device within the satellite device cluster.

Default

Member ID numbers are not assigned in a satellite device cluster, by default.

Options

member-ID-number Specifies the member ID of the satellite device in the satellite device cluster.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

member-interface (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 10612
- Hierarchy Level | 10612
- Description | 10613
- Options | 10613
- Required Privilege Level | 10613
- Release Information | 10613

Syntax

```
member-interface interface-name;
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options]
```

Description

Specify the member interfaces for the aggregated multiservices (AMS) interface. You can configure multiple interfaces by specifying each interface in a separate statement.

Starting with Junos OS Release 16.2, an AMS interface can have up to 32 member interfaces. In Junos OS Release 16.1 and earlier, an AMS interface can have a maximum of 24 member interfaces. If you configure more than 24 member interfaces, you must set the "pic-boot-timeout " on page 12484 value to 240 or 300 seconds at the `[edit interfaces interface-name multiservice-options]` hierarchy level for every services PIC interface on the MX Series router.

For high availability service applications like Network Address Translation (NAT) that support many-to-one (N:1) redundancy, you can specify two or more interfaces.

On an MS-MPC, you can configure one-to-one (1:1) redundancy. In a 1:1 (stateful) configuration, a single backup interface provides redundancy for a single active interface. A 1:1 configuration is required for IPsec. 1:1 redundancy is not supported on the MX-SPC3 in this release.

> **NOTE**: The member interfaces that you specify must be members of aggregated multiservices interfaces (mams-).

Options

interface-name Name of the member interface. The member interface format is mams-*a*/*b*/0, where *a* is the FPC slot number and *b* is the PIC slot number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Understanding Aggregated Multiservices Interfaces for Next Gen Services

Configuring Aggregated Multiservices Interfaces

load-balancing-options (Aggregated Multiservices) | **9888**

member-link-scheduler

IN THIS SECTION

- Syntax | **10614**
- Hierarchy Level | **10614**
- Description | **10614**
- Default | **10615**
- Options | **10615**
- Required Privilege Level | **10615**
- Release Information | **10615**

Syntax

```
member-link-scheduler (replicate | scale);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit logical-systems logical-system-name class-of-service interfaces interface-name],
[edit routing-instances routing-instance-name class-of-service interfaces interface-name]
```

Description

Determines whether scheduler parameters for aggregated interface member links are applied in a replicated or scaled manner.

Default

By default, scheduler parameters are scaled (in “equal division mode”) among aggregated interface member links.

Options

`replicate`—Scheduler parameters are copied to each level of the aggregated interface member links.

`scale`—Scheduler parameters are scaled based on number of member links and applied each level of the aggregated interface member links.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Hierarchical Schedulers for CoS

member (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 10616
- Hierarchy Level | 10616
- Description | 10616
- Options | 10616
- Required Privilege Level | 10616
- Release Information | 10616

Syntax

```
member member-id {
    role (routing-engine | line-card);
    serial-number serial-number;
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Configure an MX Series router as a member of a Virtual Chassis configuration. You can configure a maximum of two member routers in an MX Series Virtual Chassis.

Options

member-id—Numeric value that identifies a member router in a Virtual Chassis configuration.

- **Values:** 0 or 1

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

> *Configuring Preprovisioned Member Information for a Virtual Chassis*
>
> *Example: Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis*

member (NSSU Upgrade Groups)

IN THIS SECTION

- Syntax | 10617
- Hierarchy Level | 10617
- Description | 10617
- Options | 10618
- Required Privilege Level | 10618
- Release Information | 10618

Syntax

```
member member-id {
    fpcs (slot-number | [list-of-slot-numbers]);
}
```

Hierarchy Level

```
[edit chassis nssu upgrade-group group-name]
```

Description

Specify the Virtual Chassis member whose line-card slot numbers you are assigning to an NSSU upgrade group.

NOTE: This statement is not applicable to Virtual Chassis or VCF member switches that do not support separate line cards. To configure Virtual Chassis or VCF member switches that do not have separate line cards into an NSSU upgrade group, use the `fpcs` statement alone, and specify the Virtual Chassis or VCF member IDs to include in the upgrade group in place of line card slot numbers.

To reduce the time an NSSU takes, you can configure NSSU line-card upgrade groups on an EX6200 or EX8200 switch with redundant Routing Engines; EX8200 Virtual Chassis; QFX3500, QFX3600, and QFX5100 Virtual Chassis; and Virtual Chassis Fabric (VCF).

To configure an upgrade group that includes line cards on different switches that support multiple line cards and comprise a Virtual Chassis, use this statement hierarchy with the `fpcs` option to first specify the Virtual Chassis member ID and then desired line card slot number or numbers on that member switch to include in the upgrade group. Use multiple statements to add line cards from different Virtual Chassis members to the upgrade group.

Options

member-id The ID of the Virtual Chassis or VCF member switch containing one or more line cards to include in an NSSU upgrade group.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches

Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade

member-range

IN THIS SECTION

Syntax | **10619**

- Hierarchy Level | 10619
- Description | 10619
- Options | 10619
- Required Privilege Level | 10619
- Release Information | 10619

Syntax

```
member-range  starting-interface-name  ending-interface-name;
```

Hierarchy Level

```
[edit interfaces interface-range interface-range-name]
```

Description

Specify the names of the first and last members of a sequence of interfaces belonging to an interface range.

Options

starting interface-name ending interface-name—Name of the first member and the name of the last member in the interface sequence.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Understanding Interface Ranges for Switches

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches

Interfaces Overview for Switches

Interfaces Overview for Switches

Junos OS Network Interfaces Library for Routing Devices

member (Redundant Power System)

IN THIS SECTION

- Syntax | 10620
- Hierarchy Level | 10620
- Description | 10621
- Options | 10621
- Required Privilege Level | 10621
- Release Information | 10621

Syntax

```
member vc-member-number {
    priority (0|1|2|3|4|5|6);
}
```

Hierarchy Level

```
[edit redundant-power-system]
```

Description

Specify the Virtual Chassis member ID of a switch connected to the Redundant Power System (RPS) for backup power supply. The member ID is required only for switches that can be configured in a Virtual Chassis. If the switch has never been configured in a Virtual Chassis, the value is always 0.

Options

member-number—Member ID of a switch that has Virtual Chassis capability that is connected to the RPS.

- **Range:** 0 through maximum members in the Virtual Chassis
- **Default:** 0

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Determining and Setting Priority for Switches Connected to an EX Series RPS

member (Security Group VPN)

IN THIS SECTION

- Syntax | **10622**
- Hierarchy Level | **10623**
- Description | **10623**

- Options | 10623
- Required Privilege Level | 10623
- Release Information | 10623

Syntax

```
member {
    ike {
        gateway gateway-name;
            policy;
            proposal;
            traceoptions;
        }
    ipsec {
        vpn vpn-name {
            df-bit (clear | copy | set);
            exclude rule rule-name {
                source-address ip-address/mask;
                destination-address ip-address/mask;
                application application;
            }
            fail-open rule rule-name {
                source-address ip-address/mask;
                destination-address ip-address/mask;
                application application;
            }
            group id;
            group-vpn-external-interface interface;
            ike-gateway gateway-name;
            recovery-probe;
        }
    }
}
```

Hierarchy Level

```
[edit security group-vpn]
```

Description

Configure group VPN member. A group member encrypts the traffic and is responsible for the actual encryption and decryption of data traffic. A group member is configured with IKE Phase 1 parameters and GC/KS information.

Options

ike*gateway-name*	Configure IKE gateway for group VPN member.
policy *policy-name*	Configure an IKE policy.
proposal*proposal-name*	Define an IKE proposal.
traceoptions	Configure group VPN tracing options to aid in troubleshooting the IKE issues.
ipsec	Configure IPsec for Phase 2 exchange on the group member.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

member (Security IDP)

IN THIS SECTION

- Syntax | 10624
- Hierarchy Level | 10624
- Description | 10624
- Options | 10624
- Required Privilege Level | 10625
- Release Information | 10625

Syntax

```
member member-name {
    attack-type (Security Anomaly) {
        (anomaly ...same statements as in [edit security idp custom-attack attack-name attack-type anomaly] hierarchy level | signature ...same statements as in [edit security idp custom-attack attack-name attack-type signature] hierarchy level);
    }
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain]
```

Description

Create the list of member attacks.

Options

member-name—Name of the member list.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

members

IN THIS SECTION

- Syntax | 10625
- Hierarchy Level | 10625
- Description | 10625
- Options | 10626
- Required Privilege Level | 10627
- Release Information | 10627

Syntax

```
members [(all | names | vlan-ids)];
```

Hierarchy Level

```
[edit interfaces (QFX Series) interface-name unit 0 family ethernet-switching vlan]
```

Description

For trunk interfaces, configure the VLANs for which the interface can carry traffic.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type **?** after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

NOTE: The number of VLANs supported per switch varies for each model. Use the configuration-mode command `set vlans id vlan-id ?` to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum.

On an EX Series switch that runs Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style, the maximum number of VLAN members allowed on the switch is 8 times the maximum number of VLANs the switch supports (vmember limit = vlan max * 8). If the switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (**eswd**) due to memory allocation failure.

On an EX Series switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs the switch supports (vmember limit = vlan max * 24). If the configuration of one of these switches exceeds the recommended VLAN member maximum, a warning message appears in the system log (syslog).

Options

all—Specifies that this trunk interface is a member of all VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.

NOTE: Since VLAN members are limited, specifying **all** could cause the number of VLAN members to exceed the limit at some point.

NOTE: Each VLAN that is configured must have a specified VLAN ID when you attempt to commit the configuration; otherwise, the configuration commit fails. Also, **all** cannot be the name of a VLAN on the switch.

names—Name of one or more VLANs. VLAN IDs are applied automatically in this case.

vlan-ids—Numeric identifier of one or more VLANs. For a series of tagged VLANs, specify a range; for example, **10–20** or **10–20 23 27–30**.

> **NOTE**: Each configured VLAN must have a specified VLAN ID to successfully commit the configuration; otherwise, the configuration commit fails.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement updated with enhanced **?** (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

RELATED DOCUMENTATION

- Configuring Gigabit Ethernet Interfaces (CLI Procedure)
- *Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support*
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure)
- *Configuring VLANs for EX Series Switches*
- Configuring VLANs for EX Series Switches with ELS Support (CLI Procedure)
- *Example: Setting Up Basic Bridging and a VLAN on Switches*
- *Understanding Bridging and VLANs on Switches*
- show ethernet-switching interfaces | **19641**
- show vlans | **27738**

memory

IN THIS SECTION

- Syntax | 10628
- Hierarchy Level | 10628
- Description | 10628
- Options | 10628
- Required Privilege Level | 10629
- Release Information | 10629

Syntax

```
memory {
    size size;
    features {
    hugepages;
    }
}
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Configure memory parameters for VNFs on a platform that is running a disaggregated Junos OS.

Options

memory *size* Amount of memory allocated to a VNF in kilobytes. The default size is 1 GB.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D40.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

memory-enhanced

IN THIS SECTION

- Syntax | 10629
- Hierarchy Level | 10630
- Description | 10630
- Required Privilege Level | 10630
- Release Information | 10630

Syntax

```
memory-enhanced {
    filter;
    route;
    vpn-label;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Allocate more jtree memory for routing tables and Layer 3 VPNs.

> **NOTE**: The `memory-enhanced` statement supports MX Series routers with DPC (I-chip based) line cards only.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Allocating More Memory for Routing Tables, Firewall Filters, and Layer 3 VPN Labels

memory (System)

IN THIS SECTION

- Syntax | 10631
- Hierarchy Level | 10632
- Description | 10632

- Options | 10632
- Required Privilege Level | 10632
- Release Information | 10633

Syntax

```
memory {
    system {
        alarm;
        frequency (minutes | hours| days) <value>;
        minor event {
            threshold swap usage <value in percentage>;
            action snmp-trap;
        }
        major event {
            threshold swap usage <value in percentage>;
            action [snmp-trap | live-vmcore];
        }
        critical event {
            threshold swap usage <value in percentage>;
            action [snmp-trap | live-vmcore| (RE-restart| RE-switchover) ];
        }
    }
    process (Memory) {
    frequency (minutes | hours| days) <value>;
        minor event {
            threshold <process-name> memory-limit <memory-limit-in MBs>;
            action snmp-trap;
        }
        major event {
            threshold <process-name> memory-limit <memory-limit-in MBs>;
            action [ snmp-trap |live-core | (terminate-process | terminate-process-force)] ;
        }
        critical event {
            threshold <process-name> memory-limit <memory-limit-in MBs>;
            action [ snmp-trap |live-core | (terminate-process | terminate-process-force)] ;
```

```
        }
}
```

Hierarchy Level

```
[edit system monitor]
```

Description

Configure the memory monitoring system on your PTX Series and MX Series devices. The memory monitoring system on your device monitors anomalies with respect to the memory utilization of Junos applications and overall system in the control plane of your device. Apart from monitoring the system memory, you can also configure the device to monitor the per-process memory footprint as well to help detect possible leaks and take configured action.

Memory can be exhausted when there is over scaling, for instance when you test a high scale on a device that does support the scale because of physical resource limitations. Memory can also be exhausted when there are temporary spikes for example when a new FIB object is added before the older FIB objects are deleted. These temporary spikes are usually transient and resolve after a while. Also, a memory leak in any one of the applications or even in the kernel can lead to memory exhaustion. So, it is necessary that you monitor memory usage across applications and the kernel.

Starting in Junos OS Release 23.1R1, you can configure the memory monitoring on the SRX Series, QFX Series, and EX Series devices to raise an alarm using the `set system monitor memory system alarm` command statement. The alarm is raised when the device is running low on memory.

The remaining statements are explained separately.

Options

alarm

Enable system monitor alarm.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1R1.

The `alarm` option is introduced in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

System Log Overview

memory (TDF Gateway)

IN THIS SECTION

- Syntax | 10633
- Hierarchy Level | 10633
- Description | 10634
- Options | 10634
- Required Privilege Level | 10634
- Release Information | 10634

Syntax

```
memory memory-pct;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name cac]
```

Description

Configure the threshold for the maximum amount of memory that the TDF gateway may use. If the amount of memory that the TDF gateway uses reaches the threshold, the SNMP trap **jnxScgSMMemoryThreshHigh** is generated.

Options

memory-pct Maximum percentage of memory that can be used.

- **Range:** 1 through 90.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a TDF Gateway

menu-item

IN THIS SECTION

- Syntax | **10635**
- Hierarchy Level | **10635**
- Description | **10635**
- Options | **10635**
- Required Privilege Level | **10638**

Release Information | 10638

Syntax

```
menu-item (menu-name | menu-option) <disable>;
```

Hierarchy Level

```
[edit chassis lcd-menu],
[edit chassis lcd-menu fpc slot-number]
```

Description

Disable or enable the Maintenance menu, the Status menu, or an individual option in one of those menus in the LCD panel.

On EX3200, EX3300, EX4200, and EX4500 switches, you use **menu-item** at the `[edit chassis lcd-menu fpc slot-number]` hierarchy level.

On EX6200 and EX8200 switches, and on XRE200 External Routing Engines, you use **menu-item** at the `[edit chassis lcd-menu]` hierarchy level.

Options

menu-name—Name of the LCD menu:

- **maintenance-menu**
- **status-menu**

menu-option—Specific option on one of the LCD menus. You must include the quotation marks when you type the option. Table 62 on page 10636 describes the different menu options of the LCD menus supported on the switches.

Table 62: Menu Options of the LCD Menus Supported on the Switches

Menu	Menu Options	Option Descriptions	Platforms Supported
maintenance-menu	**"maintenance-menu halt-menu"**	System halt option	All switches except EX2200
	"maintenance-menu system-reboot"	System reboot option	All switches except EX2200
	"maintenance-menu rescue-config"	Load rescue option	All switches except EX2200
	"maintenance-menu vc-uplink-config"	Request VC port option for a device in a Virtual Chassis configuration	EX3300, EX4200, and EX4500 switches and XRE200 External Routing Engines only
	"maintenance-menu factory-default"	Factory default option	All switches except EX2200
status-menu	**"status-menu vcp-status"**	Virtual Chassis port (VCP) status for a device in a Virtual Chassis configuration	EX3300, EX4200, and EX4500 switches and XRE200 External Routing Engines only

Table 62: Menu Options of the LCD Menus Supported on the Switches *(Continued)*

Menu	Menu Options	Option Descriptions	Platforms Supported
	"status-menu sf-status1-menu"	Status of the switch fabric on the Switch Fabric and Routing Engine (SRE) module in slot SRE0 on EX8208 switches Status of the switch fabric on the Switch Fabric (SF) modules in slots SF0 and SF1 on EX8216 switches	EX8208 and EX8216 switches only
	"status-menu sf-status2-menu"	Status of the switch fabric on the SRE module in slot SRE1 on EX8208 switches Status of the switch fabric on the SF modules in slots SF2–SF5 on EX8216 switches	EX8208 and EX8216 switches only
	"status-menu sf-status3-menu"	Status of the switch fabric on the SF modules in slots SF6 and SF7 on EX8216 switches	EX8216 switches only
	"status-menu power-status"	Status of the power supply or power supplies	EX3200, EX3300, EX4200, and EX4500 switches and XRE200 External Routing Engines only
	"status-menu psu-status1-menu"	Status of the power supplies in slots P0 and P1	EX8208 and EX8216 switches only

Table 62: Menu Options of the LCD Menus Supported on the Switches *(Continued)*

Menu	Menu Options	Option Descriptions	Platforms Supported
	"status-menu psu-status2-menu"	Status of the power supplies in slots P2–P5	EX8208 and EX8216 switches only
	"status-menu environ-menu"	Status of the fan; current chassis temperature	All switches (except EX2200) and XRE200 External Routing Engine
	"status-menu show-version"	The version of Junos OS loaded on the switch	All switches except EX2200

disable—(Optional) Disable the specified menu.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the LCD Panel on EX Series Switches (CLI Procedure)

LCD Panel in EX3200 Switches

LCD Panel in EX3300 Switches

LCD Panel in EX4200 Switches

LCD Panel in EX4500 Switches

LCD Panel in EX4550 Switches

LCD Panel in an EX6200 Switch

LCD Panel in an EX8200 Switch

LCD Panel in an XRE200 External Routing Engine

mep

IN THIS SECTION

- Syntax | 10639
- Hierarchy Level | 10640
- Description | 10640
- Options | 10640
- Required Privilege Level | 10640
- Release Information | 10640

Syntax

```
mep mep-id {
    action-profile action-profile-name
    auto-discovery;
    direction (up | down);
    interface interface-name (protect | working);
    priority number;
    remote-mep mep-id {
        action-profile profile-name;
        sla-iterator-profile profile-name {
            data-tlv-size size;
            iteration-count count-value;
            priority priority-value;
        }
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-association ma-name]
```

Description

The numeric identifier of the maintenance association end point (MEP) within the maintenance association.

Options

`mep` *`mep-id`*—Specify the numeric identifier of the MEP.

- **Range:** 1 through 8191

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Ethernet Alarm Indication Signal (ETH-AIS) Function Overview

Configuring ETH-AIS on a CFM MEP

Configuring a MEP to Generate and Respond to CFM Protocol Messages

mesh-group (Protocols VPLS)

IN THIS SECTION

- Syntax | 10641
- Hierarchy Level | 10642
- Description | 10642
- Options | 10643
- Required Privilege Level | 10643
- Release Information | 10643

Syntax

```
mesh-group mesh-group-name {
    interface interface-name;
    l2vpn-id (as-number:id | ip-address:id);
    local-switching;
    mac-flush [ explicit-mac-flush-message-options ];
    neighbor address {...};
    peer-as all;
    pseudowire-status-tlv hot-standby-vc-on;
    route-distinguisher (as-number:id | ip-address:id);
    vpls-id number;
    vrf-export [ policy-names ];
    vrf-import [ policy-names ];
    vrf-target {
        community;
        import community-name;
        export community-name;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
site site-name],
[edit routing-instances routing-instance-name protocols vpls],
[edit routing-instances routing-instance-name protocols vpls site site-name]
```

Description

Specify the virtual private LAN service (VPLS) mesh group. The statement options allow you to specify each provider edge (PE) device that is a member of the mesh group. This statement is also used in the configuration of inter-autonomous system (AS) VPLS with media access control (MAC) operations.

When you configure mesh groups on the ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X routers, you must also enable the system profile by including the `set system packet-forwarding-options system-profile vpls-meshgroups` statement on the ACX router. Upon issue of this command, the packet forwarding engine (PFE) restarts.

The ACX7000 routers have the following limitations when you configure hierarchical VPLS (H-VPLS) mesh groups:

- Control word for EVPN MPLS multihoming is not supported.
- Penultimate Segment Pop (PSP) on SRv6 is not supported.
- VXLAN stitching is not supported.
- Up to 3 mesh-groups with no-local switching in a single VPLS routing instance.
- A total maximum of 14 mesh-groups with either local switching or no-local switching per routing instance.

When you configure hot-standby redundant pseudowires in an H-VPLS hub-and-spoke network, we recommend that you also enable `vpls-hot-standby-convergence`. With VPLS hot-standby convergence, the device advertises two different labels (L1 and L2) to its neighbors. The device sends the L1 advertisement to the primary neighbor and the L2 advertisement to the backup neighbor. This allows a spoke device to differentiate between messages sent by the primary pseudowire node from messages sent by the backup pseudowire node. This reduces traffic sent and improves network convergence when there is a transition from the primary pseudowire to the backup pseudowire. To enable VPLS hot-standby convergence, include `vpls-hot-standby-convergence` at the `[edit routing-options forwarding-table]` hierarchy level on all the spoke devices.

> **NOTE**: You can configure the `no-local-switching` option at the `[edit routing-instances routing-instance-name]` or the `[edit logical-systems logical-system-name routing-instances routing-instance-name]` hierarchies to disable local switching in a routing instance with VPLS.
>
> If you set this option, you can't set the `local-switching` option under a VPLS mesh group to enable local switching only for a particular VPLS mesh group. The instance-level `no-local-switching` statement overrides the per-mesh-group `local-switching` setting.

Options

mesh-group-name—Name of the VPLS mesh group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`local-switching`, `mac-tlv-receive`, `mac-tlv-send`, and `peer-as` options introduced in Junos OS Release 9.3.

`pseudowire-status-tlv` and `mac-flush` options introduced in Junos OS Release 10.0.

`route-distinguisher`, `vrf-export`, `vrf-import`, and `vrf-target` options introduced in Junos OS Release 11.2.

`hot-standby-vc-on` and `interface interface-name` options introduced in Junos OS Release 15.1R2.

Support for this statement was introduced in Junos OS Evolved Release 24.1R1 for ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024 Routers.

RELATED DOCUMENTATION

Configuring VPLS Routing Instances

Configuring Interoperability Between BGP Signaling and LDP Signaling in VPLS

message-flood (Security H323)

IN THIS SECTION

- Syntax | 10644
- Hierarchy Level | 10644
- Description | 10644
- Options | 10645
- Required Privilege Level | 10645
- Release Information | 10645

Syntax

```
message-flood {
    gatekeeper {
        threshold rate;
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg h323 application-screen],
[edit logical-systems name tenants name security alg h323 application-screen],
[edit security alg h323 application-screen],
[edit services alg h323 application-screen],
[edit tenants name security alg h323 application-screen]
```

Description

Configure the message-flood parameter to limit the number of message requests allowed per second from Remote Access Server (RAS) to the gatekeeper. Messages exceeding the threshold are dropped.

Options

`gatekeeper threshold` *rate*—Maximum number of RAS connection requests per second allowed per gateway.

- **Range:** 1 through 65,535
- **Default:** 1000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

h323 | **7961**

message-flood (Security MGCP)

IN THIS SECTION

- Syntax | **10646**
- Hierarchy Level | **10646**
- Description | **10646**
- Options | **10646**
- Required Privilege Level | **10646**
- Release Information | **10646**

Syntax

```
message-flood {
    threshold rate;
}
```

Hierarchy Level

```
[edit logical-systems name security alg mgcp application-screen],
[edit logical-systems name tenants name security alg mgcp application-screen],
[edit security alg mgcp application-screen],
[edit services alg mgcp application-screen],
[edit tenants name security alg mgcp application-screen]
```

Description

Configure the message-flood parameter to limit the number of Media Gateway messages processed per second. Messages exceeding the threshold are dropped by the Media Gateway Control Protocol (MGCP) Application Layer Gateway (ALG). This feature is disabled by default.

Options

`threshold` *`rate`* —Number of connection requests per second allowed per Media Gateway.

- **Range:** 2 through 50000
- **Default:** 1000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

message-ie-profile-v1

IN THIS SECTION

- Syntax | 10647
- Hierarchy Level | 10647
- Description | 10647
- Options | 10648
- Required Privilege Level | 10648
- Release Information | 10648

Syntax

```
message-ie-profile-v1 name {
    message msg-type {
        ie ie-number;
    }
}
```

Hierarchy Level

```
[edit security gtp]
```

Description

Specify a GTPv1 message IE profile. You can implement Must-IE check for GTPv1 using message profile configurations, which helps you to define must IEs of interested messages. Along with appropriate message profile configurations, Must-IE check can easily accommodate any GTP releases, message format, or IE status.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp]` hierarchy level is replaced by `[edit security gtp]`.

Options

name Name of the GTPv1 message IE profile.

msg-type Message type.

- **Range:** 1 to 255

ie-number IE number.

- **Range:** 1 to 255

Required Privilege Level

security—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Example: Configure Must-IE check for GTPv1 and GTPv2

message-ie-profile-v2

IN THIS SECTION

- Syntax | 10649
- Hierarchy Level | 10649
- Description | 10649
- Options | 10649
- Required Privilege Level | 10650
- Release Information | 10650

Syntax

```
message-ie-profile-v2 name {
    message msg-type  {
        ie ie-number  {
            instance inst-number  {
                grouped-ie-profile grpie-prf-name;
            }
        }
    }
}
```

Hierarchy Level

```
[edit security gtp]
```

Description

Specify a GTPv2 message IE profile. You can implement Must-IE check for GTPv2 using message profile configurations, which helps you to define must IEs of interested messages. Along with appropriate message profile configurations, Must-IE check can easily accommodate any GTP releases, message format, or IE status.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp]` hierarchy level is replaced by `[edit security gtp]`.

Options

name Name of the message IE filter profile.

msg-type Message type.

- **Range:** 1 to 255

ie-number IE number.

- **Range:** 1 to 255

inst-number Instance number.

- **Range:** 0 to 11
- **Default:** 0

grpie-prf-name GTP grouped IE filter profile name.

Required Privilege Level

security—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Example: Configure Must-IE check for GTPv1 and GTPv2

message-list

IN THIS SECTION

- Syntax | **10651**
- Hierarchy Level | **10651**
- Description | **10651**
- Options | **10651**
- Required Privilege Level | **10651**
- Release Information | **10651**

Syntax

```
message-list msg-list-name  {
message msg-number;
}
```

Hierarchy Level

```
[edit security gtp]
```

Description

Use the message-list function to configure a list of messages.

Options

msg-list-name Specifies the message list name.

msg-number Specifies the message type.

- **Range** 1 to 255

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

show security gtp message-list

message-processing-model

IN THIS SECTION

- Syntax | 10652
- Hierarchy Level | 10652
- Description | 10652
- Options | 10652
- Required Privilege Level | 10652
- Release Information | 10653

Syntax

```
message-processing-model (v1 | v2c | v3);
```

Hierarchy Level

```
[edit snmp v3 target-parameters target-parameter-name parameters]
```

Description

Configure the message processing model to be used when generating SNMP notifications.

Options

`v1`—SNMPv1 message process model.

`v2c`—SNMPv2c message process model.

`v3`—SNMPv3 message process model.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Message Processing Model

message-rate-limit (interfaces)

IN THIS SECTION

- Syntax | 10653
- Hierarchy Level | 10653
- Description | 10654
- Options | 10654
- Required Privilege Level | 10654
- Release Information | 10654

Syntax

```
message-rate-limit messages-per-second
```

Hierarchy Level

```
interfaces interface-name {
    services-options {
        cgn-pic;
        disable-global-timeout-override;
        ignore-errors <alg> <tcp>;
```

```
        inactivity-non-tcp-timeout seconds;
        inactivity-tcp-timeout seconds;
        inactivity-timeout seconds;
        open-timeout seconds;
        session-limit {
            maximum number;
            rate (Interface Services) new-sessions-per-second;
        }
        session-timeout seconds;
        syslog {
        }
    }
}
```

Description

Maximum system log messages per second allowed from this interface.

> **NOTE**: The message-rate-limit command can be configured only for physical service interfaces (sp-*x*/*x*/*x*) and not for redundancy services PIC interfaces (rsp*x*).

Options

messages-per-second —This option configures the maximum number of system log messages per second that can be formatted and sent from the PIC to either the Routing Engine (local) or to an external server (remote). The default rates are 10,000 for the Routing Engine and 800,000 for an external server.

- **Range:** 0 through 2147483647

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring System Logging for Service Sets

message-rate-limit

IN THIS SECTION

- Syntax | 10655
- Hierarchy Level | 10655
- Description | 10655
- Options | 10655
- Required Privilege Level | 10656
- Release Information | 10656

Syntax

```
message-rate-limit messages-per-minute;
```

Hierarchy Level

```
[edit protocols pcep]
```

Description

Specify the number of messages per minute that the Path Computation Client (PCC) can receive at maximum.

Options

messages-per-minute Number of messages per minute that the PCC can receive at maximum.

- **Range:** 1 through 16384
- **Default:** 0 (disabled or no limit)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

pcep | 12320

message-rate-limit (Flow Monitoring Logs for NAT)

IN THIS SECTION

- Syntax | 10657
- Hierarchy Level | 10657
- Description | 10657
- Options | 10657
- Required Privilege Level | 10657
- Release Information | 10657

Syntax

```
message-rate-limit messages-per-second
```

Hierarchy Level

```
[edit interfaces interface-name services-options jflow-log]
```

Description

Define the maximum number of logs or template records in flow monitoring format to be generated for NAT error events per second from the specified interface. These records for NAT error events are generated when addresses for allocation from the NAT pool are not available, when ports for allocation to a subscriber are not available, or when the allocated quota is exceeded for NAT events (more than the configured number of ports is requested).

NOTE: The `message-rate-limit` option can be configured only for multiservices interfaces (ms-*x*/*x*/*x*) and not with other interface types.

Options

messages-per-second —Maximum number of flow monitoring log messages per second for NAT error events that can be formatted and sent from the PIC to an external collector. The default rate is 10,000 for an external collector.

- **Range:** 1 through 2,147,483,647

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

message-size

IN THIS SECTION

- Syntax | **10658**
- Hierarchy Level | **10658**
- Description | **10659**
- Options | **10659**
- Required Privilege Level | **10659**
- Release Information | **10659**

Syntax

```
message-size number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip],
[edit protocols rip group group-name neighbor neighbor-name],
```

```
[edit routing-instances routing-instance-name protocols rip],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-name]
```

Description

Specify the number of route entries to be included in every RIP update message. To ensure interoperability with other vendors' equipment, use the standard of 25 route entries per message.

Options

number—Number of route entries per update message.

- **Range:** 25 through 255 entries
- **Default:** 25 entries

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

message-type

IN THIS SECTION

- Syntax | 10660
- Hierarchy Level | 10661
- Description | 10661
- Options | 10661
- Required Privilege Level | 10661

Release Information | 10661

Syntax

```
message-type {
    create-req {
        alarm-threshold (Security GPRS) {
            forward number;
            reverse number;
        }
        drop-threshold (Security GPRS) {
            forward number;
            reverse number;
        }
    }
    delete-req {
        drop-threshold (Security GPRS) {
            forward number;
            reverse number;
        }
        drop-threshold (Security GPRS) {
            forward number;
            reverse number;
        }
    }
    echo-req {
        drop-threshold (Security GPRS) {
            forward number;
            reverse number;
        }
        drop-threshold (Security GPRS) {
            forward number;
            reverse number;
        }
    }
    other {
        drop-threshold (Security GPRS) {
            forward number;
            reverse number;
```

```
        }
        drop-threshold (Security GPRS) {
            forward number;
            reverse number;
        }
    }
}
```

Hierarchy Level

```
[edit security gtp profile profile-name path-rate-limit]
```

Description

Use the message-type function to specify the group of control messages.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name path-rate-limit]` hierarchy level is replaced by `[edit security gtp profile profile-name path-rate-limit]`.

Options

`create-req`—To limit packet-per-second of GTP create request.

`delete-req`—To limit packet-per-second of GTP delete request.

`echo-req`—To limit packet-per-minute of GTP echo request.

`other`—To limit packet-per-second of all other GTP control messages.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

GTPv1 Message Filtering

saml (Access Profile)

IN THIS SECTION

- Syntax | 10662
- Hierarchy Level | 10662
- Description | 10662
- Options | 10663
- Required Privilege Level | 10663
- Release Information | 10663

Syntax

```
saml {
    service-provider service-provider;
    identity-provider domain {
        idp-name idp-name;
    }
}
```

Hierarchy Level

```
[edit access profile name]
```

Description

Describe Security Assertion Markup Language (SAML) configuration for the access profile.

Options

identity-provider ***domain***

Specify the identity provider (IdP) parameters for the profile.

- Values: `any` or specify user's *domain* such as example.net.

idp-name ***idp-name***

Specify the IDP name.

service-provider ***service-provider***

Specify the service provider name defined in the access.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

saml | 16281

method

IN THIS SECTION

- Syntax | 10664
- Hierarchy Level | 10664
- Description | 10665
- Required Privilege Level | 10665
- Release Information | 10665

Syntax

```
method {
    bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        detection-time {
            threshold milliseconds;
        }
        session-mode (automatic | multihop | singlehop);
        holddown-interval milliseconds;
    }
    layer2-liveness-detection {
        max-consecutive-retries number;
        transmit-interval interval;
    }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name liveness-detection],
[edit forwarding-options dhcp-relay group group-name liveness-detection],
[edit forwarding-options dhcp-relay liveness-detection],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection],
[edit system services dhcp-local-server dhcpv6 liveness-detection],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name liveness-detection],
[edit system services dhcp-local-server group group-name liveness-detection],
[edit system services dhcp-local-server liveness-detection]
```

Description

Configure the liveness detection method.

> **NOTE**: The "bfd" on page 2909 stanza is not available at the `[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name liveness-detection method]` or `[edit system services dhcp-local-server dual-stack-group dual-stack-group-name liveness-detection` hierarchy levels.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- DHCP Liveness Detection Overview
- Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients
- Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients
- DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

metric

IN THIS SECTION

- Syntax | **10666**
- Hierarchy Level | **10666**
- Description | **10666**

- Options | 10666
- Required Privilege Level | 10666
- Release Information | 10666

Syntax

```
metric route-cost;
```

Hierarchy Level

```
[edit routing-options access route ip-prefix</prefix-length>]
```

Description

Configure the cost for an access route.

Options

route-cost—Specific cost you want to assign to the access route.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

metric (Aggregate, Generated, or Static Route)

IN THIS SECTION

- Syntax | 10667
- Hierarchy Level | 10667
- Description | 10667
- Options | 10667
- Required Privilege Level | 10668
- Release Information | 10668

Syntax

```
(metric | metric2 | metric3 | metric4) metric <type type>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)]
```

Description

Specify the metric value for an aggregate, generated, or static route. You can specify up to four metric values, starting with **metric** (for the first metric value) and continuing with **metric2**, **metric3**, and **metric4**.

Options

metric—Metric value.

- **Range:** 0 through 4,294,967,295 ($2^{32} - 1$)

type ***type***—(Optional) Type of route.

When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric.

- **Range:** 1 through 16

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

aggregate (Routing) | **2070**

generate | **7672**

static (Routing Options) | **28570**

Example: Summarizing Static Routes Through Route Aggregation

Understanding Route Aggregation

metric (Diameter Route)

IN THIS SECTION

- Syntax | **10669**
- Hierarchy Level | **10669**
- Description | **10669**
- Options | **10669**
- Required Privilege Level | **10669**
- Release Information | **10669**

Syntax

```
metric route-metric;
```

Hierarchy Level

```
[edit diameter network-statement element-name forwarding route dne-route-name]
```

Description

Specify the metric associated with a destination and function. Together, these three elements define a route reachable through a Diameter network element. A lower metric makes a route more preferred.

Options

route-metric—Metric assigned to the route.

- **Range:** 0 through 255

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

> *Configuring Diameter*
>
> *Configuring Diameter Network Elements*

metric (Dynamic Access-Internal Routes)

IN THIS SECTION

- Syntax | 10670
- Hierarchy Level | 10670
- Description | 10670
- Options | 10670
- Required Privilege Level | 10671
- Release Information | 10671

Syntax

```
metric route-cost;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
access route prefix],
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
rib routing-table-name access route prefix],
[edit dynamic-profiles profile-name routing-options access route prefix]
```

Description

Dynamically configure the cost for an access route.

Options

route-cost—Either the specific cost you want to assign to the access route or either of the following cost variables:

- `$junos-framed-route-cost`—Cost of an IPv4 access route; the variable is dynamically replaced with the metric value (Subattribute 3) from the RADIUS Framed-Route attribute [22].

- `$junos-framed-route-ipv6-cost`—Cost of an IPv6 access route; the variable is dynamically replaced with the metric value (Subattribute 3) from the RADIUS Framed-IPv6-Route attribute [99].

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*] and [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*] hierarchy levels introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Routes for DHCP and PPP Subscriber Access Networks

metric-in (Protocols RIP)

IN THIS SECTION

- Syntax | 10672
- Hierarchy Level | 10672
- Description | 10672
- Options | 10672
- Required Privilege Level | 10672
- Release Information | 10672

Syntax

```
metric-in metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip],
[edit protocols rip group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

Specify the metric to add to incoming routes when the routing device advertises into RIP routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIP routes learned through a specific neighbor.

Options

metric—Metric value.

- **Range:** 1 through 16
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring the Metric Value Added to Imported RIP Routes

metric-in (Protocols RIPng)

IN THIS SECTION

- Syntax | 10673
- Hierarchy Level | 10673
- Description | 10674
- Options | 10674
- Required Privilege Level | 10674
- Release Information | 10674

Syntax

```
metric-in metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng],
[edit logical-systems logical-system-name protocols ripng group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ripng],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ripng group group-name neighbor neighbor-name],
[edit protocols ripng],
[edit protocols ripng group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols ripng],
[edit routing-instances routing-instance-name protocols ripng group group-name neighbor neighbor-name]
```

Description

Specify the metric to add to incoming routes when advertising into RIPng routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIPng routes learned through a specific neighbor.

Options

metric—Metric value.

- **Range:** 1 through 16
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

> *Example: Configuring the Metric Value Added to Imported RIPng Routes to Control the Route Selection Process*

metric-out (Protocols RIP)

IN THIS SECTION

- Syntax | **10675**
- Hierarchy Level | **10675**
- Description | **10675**

- Options | 10675
- Required Privilege Level | 10675
- Release Information | 10676

Syntax

```
metric-out metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIP routes sent from this neighbor.

Options

metric—Metric value.

- **Range:** 1 through 15
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

metric-out

IN THIS SECTION

- Syntax | 10676
- Hierarchy Level | 10676
- Description | 10677
- Options | 10677
- Required Privilege Level | 10678
- Release Information | 10678

Syntax

```
metric-out (metric | minimum-igp offset | igp (delay-med-update | offset);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],

[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
```

```
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Specify the metric for all routes sent using the multiple exit discriminator (MED, or **MULTI_EXIT_DISC**) path attribute in update messages. This path attribute is used to discriminate among multiple exit points to a neighboring AS. If all other factors are equal, the exit point with the lowest metric is preferred.

You can specify a constant metric value by including the ***metric*** option. For configurations in which a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the `multihop` command—you can specify a variable metric by including the **minimum-igp** or **igp** option.

You can increase or decrease the variable metric calculated from the IGP metric (either from the **igp** or `minimum-igp` statement) by specifying a value for ***offset***. The metric is increased by specifying a positive value for ***offset***, and decreased by specifying a negative value for ***offset***.

In Junos OS Release 9.0 and later, you can specify that a BGP group or peer not advertise updates for the MED path attributes used to calculate IGP costs for BGP next hops unless the MED is lower. You can also configure an interval to delay when MED updates are sent by including the `med-igp-update-interval` *minutes* statement at the `[edit routing-options]` hierarchy level.

Options

delay-med-update—Specify that a BGP group or peer configured with the `metric-out igp` statement not advertise MED updates unless the current MED value is lower than the previously advertised MED value, or another attribute associated with the route has changed, or the BGP peer is responding to a refresh route request.

> **NOTE**: You cannot configure the `delay-med-update` statement at the global BGP level.

igp—Set the metric to the most recent metric value calculated in the IGP to get to the BGP next hop. Routes learned from an EBGP peer usually have a next hop on a directly connected interface and thus the IGP value is equal to zero. This is the value advertised.

metric—Primary metric on all routes sent to peers.

- **Range:** 0 through 4,294,967,295 ($2^{32} - 1$)
- **Default:** No metric is sent.

minimum-igp—Set the metric to the minimum metric value calculated in the IGP to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value. When you change a neighbor's export policy from any configuration to a configuration that sets the minimum IGP offset on an exported route, the advertised MED is not updated if the value would increase as a result, even if the previous configuration does not use a minimum IGP-based MED value. This behavior helps to prevent unnecessary route flapping when an IGP cost changes, by not forcing a route update if the metric value increases past the previous lowest known value.

offset—Increases or decreases the metric by this value.

- **Range:** -2^{31} through $2^{31} - 1$
- **Default:** None

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Option **delay-med-update** introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates

Examples: Configuring BGP MED

Example: Configuring the MED Attribute That Determines the Exit Point in an AS

Understanding the MED Attribute That Determines the Exit Point in an AS

med-igp-update-interval | **10594**

metric-out (Protocols RIPng)

IN THIS SECTION

- Syntax | 10679
- Hierarchy Level | 10679
- Description | 10679
- Options | 10680
- Required Privilege Level | 10680
- Release Information | 10680

Syntax

```
metric-out metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ripng group group-name neighbor neighbor-name],
[edit protocols ripng group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols ripng group group-name neighbor neighbor-name]
```

Description

Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIPng routes sent from this neighbor.

When an export policy is configured, RIPng exports all learned routes to neighbors configured with the `neighbor` statement.

If a route being exported was learned from a member of the same RIPng group, the metric associated with that route (unless modified by an export policy) is the normal RIPng metric. For example, a RIPng

route with a metric of 5 learned from a neighbor configured with a **metric-in** value of 2 is advertised with a combined metric of 7 when advertised to RIPng neighbors in the same group. However, if this route was learned from a RIPng neighbor in a different group or from a different protocol, the route is advertised with the metric value configured for that group with the `metric-out` statement. The default value for **metric-out** is 1.

To modify the metric for routes advertised outside a group, include the `metric-out` statement.

Options

metric—Metric value.

- **Range:** 1 through 16
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

> *Example: Configuring the Metric Value Added to Imported RIPng Routes to Control the Route Selection Process*
>
> *Understanding RIPng Traffic Control with Metrics for Optimizing the Path Cost*

metric (Policy Action)

IN THIS SECTION

- Syntax | 10681
- Hierarchy Level | 10681
- Description | 10681
- Options | 10681
- Required Privilege Level | 10682
- Release Information | 10682

Syntax

```
metric (add add | aigp | expression  expression | igp metric_offset  | minimum-igp metric_offset
| subtract subtract)
```

Hierarchy Level

```
[edit policy-options policy-statement policy-name term term-name then ],
[edit policy-options policy-statement policy-name then ]
```

Description

Specify this CLI policy action in an import or export policy to set the metric value to one of the following options as per your network requirement.

Options

add *add*
Specify a constant that should be added to the metric value. Use this option if you want to increase the metric of a route.

expression
Calculate the metric value based on the route metric and metric2.

igp *metric_offset*	Configure this option to set the IGP metric for BGP. Specify a metric offset to increase or decrease the calculated IGP metric.
minimum-igp *metric_offset*	Configure this option to set the minimum IGP metric for BGP.
aigp	Configure this option to set IGP metric to the accumulated interior gateway protocol (AIGP) metric value if an AIGP attribute has been configured. Specify this value in an export policy to redistribute BGP routes to an IGP, such as the OSPF protocol.
subtract *subtract*	Specify a constant that must be subtracted from the metric value. Use this option if you want to decrease the metric of a route.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

aigp-adjust (Policy Action) | 2095

show policy | 23199

metric (Protocols DVMRP)

IN THIS SECTION

- Syntax | 10683
- Hierarchy Level | 10683
- Description | 10683
- Options | 10683

Required Privilege Level | 10683
Release Information | 10683

Syntax

```
metric metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols dvmrp interface interface-name],
[edit protocols dvmrp interface interface-name]
```

Description

Define the DVMRP metric value.

Options

metric—Metric value.

- **Range:** 1 through 31
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

NOTE: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable

in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring DVMRP

metric (Protocols IS-IS)

IN THIS SECTION

- Syntax | 10684
- Hierarchy Level | 10684
- Description | 10685
- Options | 10686
- Required Privilege Level | 10686
- Release Information | 10686

Syntax

```
metric metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis interface interface-name level level-number],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis interface interface-name level level-number],
[edit protocols isis interface interface-name level level-number],
```

```
[edit routing-instances routing-instance-name protocols isis interface interface-name level
level-number]
```

Description

Specify the IPv4 unicast topology metric value for the level. The IS-IS interface metrics for the IPv6 topology can be configured independently of the IPv4 metrics.

All IS-IS routes have a cost, which is a routing metric that is used in the IS-IS link-state calculation. The cost is an arbitrary, dimensionless integer that can be from 1 through 63, or from 1 through 16,777,215 (2^{24} – 1) if you are using wide metrics.

Similar to other routing protocols, IS-IS provides a way of exporting routes from the routing table into the IS-IS network. When a route is exported into the IS-IS network without a specified metric, IS-IS uses default metric values for the route, depending on the protocol that was used to learn the route.

Table 63 on page 10685 depicts IS-IS route export metric default values.

Table 63: Default Metric Values for Routes Exported into IS-IS

Protocol Used for Learning the Route	Default Metric Value
Direct	10
Static	Same as reported by the protocol used for exporting the route
Aggregate	10
Generate	10
RIP	Same as reported by the protocol used for exporting the route
OSPF	Same as reported by the protocol used for exporting the route
BGP	10

The default metric values behavior can be customized by using routing policies.

Options

metric—Metric value.

- **Range:** 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)
- **Default:** 10 (for all interfaces except lo0), 0 (for the lo0 interface)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Enabling Wide IS-IS Metrics for Traffic Engineering

Understanding Wide IS-IS Metrics for Traffic Engineering

metric (Protocols MPLS)

IN THIS SECTION

- Syntax | 10687
- Hierarchy Level | 10687
- Description | 10687
- Options | 10687
- Required Privilege Level | 10687
- Release Information | 10687

Syntax

```
metric metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Description

Compare against another LSP or against an IGP route. To disable dynamic metric tracking, assign a fixed metric value to an LSP. If no metric is assigned, the LSP metric is dynamic and automatically tracks underlying IGP metrics.

Options

metric—LSP metric value.

- **Default:** No metric assigned (dynamic)
- **Range:** 1 through 16,777,215

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring Static LSP Metrics*

metric (Protocols OSPF Sham Link)

IN THIS SECTION

- Syntax | 10688
- Hierarchy Level | 10688
- Description | 10688
- Default | 10688
- Options | 10689
- Required Privilege Level | 10689
- Release Information | 10689

Syntax

```
metric number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area area-id sham-link-remote address],
[edit routing-instances routing-instance-name protocols ospf area area-id sham-link-remote
address]
```

Description

Specify the cost of using the OSPF sham link.

Default

number—1

Options

number—1 through 65,535

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Giving VPN Backbones Priority with OSPFv2 Sham Links for Layer 3 VPNs

metric (Qualified Next Hop on Static Route)

IN THIS SECTION

- Syntax | 10689
- Hierarchy Level | 10690
- Description | 10690
- Options | 10690
- Required Privilege Level | 10690
- Release Information | 10690

Syntax

```
metric metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options static route destination-prefix
qualified-next-hop],
[edit routing-options static route destination-prefix qualified-next-hop]
```

Description

Metric value for a static route.

Options

metric—Metric value.

- **Range:** 0 through 4,294,967,295 (2^{32} – 1)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

qualified-next-hop (Static Routes) | **13910**

static (Routing Options) | **28570**

metrics-profile

IN THIS SECTION

- Syntax | 10691
- Hierarchy Level | 10692
- Description | 10692
- Options | 10692
- Required Privilege Level | 10693
- Release Information | 10693

Syntax

```
metrics-profile metrics-profile-name {
    sla-threshold {
        delay-round-trip {
            delay-value;
        }
         jitter {
            jitter-value;
        }
        jitter-type {
            egress-jitter ;
            ingress-jitter;
             two-way-jitter;
        }
        match {
            [all | any] ;
        }
        packet-loss {
            loss-value;
        }
    }
```

Hierarchy Level

[edit security advance-policy-based-routing]

Description

Create a set of metrics, which can be used by AppQoE to evaluate the SLA of the link.

A metrics profile defines the performance metrics for delay round trip, one-way jitter or two-way jitter, and packet loss.

To ensure compliance with the SLA, metrics are required to measure and monitor the network performance. This measurement capability provides a greater visibility into the performance characteristics of the links and helps in network performance evaluation.

Options

metrics-profile-name Metrics profile name.

delay-round-trip *delay-value* Sets the total round-trip time (in microseconds), from the device to the remote server, that triggers a probe failure.

jitter *jitter-value* Total jitter (in microseconds) for a test, which, if exceeded, triggers a probe failure

jitter-type Jitter type.

- **Values:** Ingress jitter, egress jitter, and two-way jitter.
- **Default:** Two-way jitter

match Matching SLA metrics.

all The path selection mechanism attempts to find a path that satisfies all the metrics. If no such path exists, then the next best path (based on number of metrics satisfied) is used. If there are more than one path that satisfy the metric, a random path among the available paths will be selected. Also, SLA violation will be detected and raised even if any one of the metrics is violated.

This is the default match option.

any Path selection mechanism attempts to find a path which satisfies the maximum number of metrics. For example, if there is a path available that conforms to more than one metric, then the path is chosen over another path which satisfies less number of metrics. In this case, SLA violation is detected only when none of

the metrics meets the requirement. If either one of the metric is meets the requirement, then violation is not triggered.

packet-loss *loss-value* — Percentage of number of packets that must be lost successively to trigger a probe failure.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Application Quality of Experience

Advanced Policy-Based Routing

mgcp

IN THIS SECTION

- Syntax | 10694
- Hierarchy Level | 10694
- Description | 10695
- Options | 10695
- Required Privilege Level | 10695
- Release Information | 10695

Syntax

```
mgcp {
    application-screen {
        connection-flood {
            threshold rate;
        }
        message-flood {
            threshold rate;
        }
        unknown-message {
            permit-nat-applied;
            permit-routed;
        }
    }
    disable;
    dscp-rewrite {
        code-point string;
    }
    inactive-media-timeout value-in-seconds;
    maximum-call-duration value-in-minutes;
    traceoptions {
        flag {
            all {
                extensive
            }
        }
    }
    transaction-timeout value-in-seconds;
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Specify the Media Gateway Control Protocol (MGCP) ALG on the device. MGCP is a text-based Application Layer protocol that can be used for call setup and call control.

Options

`disable`—Disable the MGCP ALG. By default, the MGCP ALG is enabled.

> **NOTE**: By default, the MGCP ALG is disabled for NFX Series and SRX Series Firewalls.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding the MGCP ALG

mib-profile

IN THIS SECTION

- Syntax | 10696
- Hierarchy Level | 10696
- Description | 10696

- Options | 10696
- Required Privilege Level | 10697
- Release Information | 10697

Syntax

```
mib-profile profile-name {
    file filename;
    interval minutes;
    start-time time;
    object-names {
        mib-object-name;
    }
    operation operation-name;
}
```

Hierarchy Level

```
[edit accounting-options]
```

Description

Create a MIB profile to collect selected MIB statistics and write them to a file in the `/var/log` directory.

> **NOTE**: Do not configure MIB objects related to interface octets or packets for a MIB profile, because it can cause the SNMP walk or a CLI show command to time out.

Options

profile-name Name of the MIB statistics profile.

start-time ***time*** (Optional) Specify the start time in *yyyy-mm-dd.hh:mm* local time format for the accounting profile. Sometimes, if the time at which the records are written to the accounting file is

beyond the transfer interval window of the file, drifting of accounting records occurs. To prevent drifting and to have predictable timestamps of when accounting data is written to an accounting file, we can specify the start time option for the accounting profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring the MIB Profile

micro-sid

IN THIS SECTION

- Syntax | 10698
- Hierarchy Level | 10698
- Description | 10698
- Options | 10699
- Required Privilege Level | 10699
- Release Information | 10699

Syntax

```
micro-sid {
            anycast
            block-name block-name
            flavor {
                none;
                psp;
                usd;
                usp;
            }
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options source-packet-routing srv6
locator],
[edit logical-systems name routing-options source-packet-routing srv6 locator],
[edit logical-systems name tenants name routing-instances name routing-options source-packet-
routing srv6 locator],
[edit routing-instances name routing-options source-packet-routing srv6 locator],
[edit routing-options source-packet-routing srv6 locator],
[edit tenants name routing-instances name routing-options source-packet-routing srv6 locator]
```

Description

Enable compression of SIDs for the specified locator and configure parameters for the micro SID locator. All micro-SIDs derived from this locator take the default micro-SID-length of 16. If you need to include more than six SRv6 SIDs, micro-SIDs help with compressing multiple IPv6 addresses. Compression of the SID list to a single destination address (the micro-SID) reduces the bandwidth overhead.

NOTE: Currently we do not support flexible algorithm with micro SIDs extension. If you configure micro-SID, you cannot configure flexible algorithm.

Options

anycast Indicates that SIDs derived from this locator are anycast SIDs. SRv6 advertises the locator as anycast locator in IGP advertisements.

block-name Specify a block name for this micro-SID locator.

flavor Specify micro-SID locator behavior such as Penultimate Segment Pop (PSP), Ultimate Segment Pop (USP) or Ultimate Segment Decapsulation (USD). You can configure any combination of PSP, USP, USD flavors. For example, you can specify just PSP, or both (PSP and USP), or all three (PSP, USP, USD). The flavor is applicable for all SIDs derived from this locator.

- none—Use this option if you do not want to confugure a flavor
- psp—Penultimate segment pop of the SRH
- usd—Ultimate segment decapsulation
- usp—Ultimate segment pop of the SRH

- Default: If you do not specify a flavor, IS-IS uses PSP and USD by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.2R1 and Junos OS 23.4R1.

RELATED DOCUMENTATION

block | **2980**

locator | **9997**

How to Enable SRv6 Network Programming in IS-IS Networks

microloop-avoidance

IN THIS SECTION

- Syntax | 10700
- Hierarchy Level | 10700
- Description | 10701
- Options | 10701
- Required Privilege Level | 10701
- Release Information | 10701

Syntax

```
microloop-avoidance {
    post-convergence-path {
        delay milliseconds;
        maximum-labels maximum-labels-value;
        maximum-srv6-sids maximum-srv6-sids-value;
    }
}
```

Hierarchy Level

```
[edit logical-systems name protocols isis spf-options],
[edit logical-systems name routing-instances name protocols isis spf-options],
[edit logical-systems name tenants name routing-instances name protocols isis spf-options],
[edit protocols isis spf-options],
[edit routing-instances name protocols isis spf-options],
[edit tenants name routing-instances name protocols isis spf-options]
```

Description

Configure microloop avoidance mechanisms that temporarily install a post-convergence path for routes potentially affected by microloops. The temporary path is used for the specified delay time period and reverts to the node SIDs of the destination after the delay timer expires.

> **NOTE**: Micro-loop avoidance is not a replacement for local repair mechanisms like TI-LFA which detects local failure very fast and activates a precomputed loop-free-alternative path.

Options

post-convergence-path delay *milliseconds*;	Time in milliseconds after which temporary post-convergence paths are removed. • Range: 500 through 60000 • Default: 5000ms
maximum-labels (SR-MPLS)	Maximum number of labels in a microloop avoidance path. Number of labels must be less than or equal to the maximum labels configured. Range: 2 through 8
maximum-srv6-sids (SR-MPLS)	Maximum number of SRv6 SIDs in a microloop avoidance path. Number of SIDs must be less than or equal to the maximum SRv6 SIDs configured. Range: 2 through 12

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 21.1R1.

`maximum-labels` and `maximum-srv6-sids` options introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

Understanding IS-IS Microloop Avoidance

microloop-avoidance (Protocols OSPF)

IN THIS SECTION

- Syntax | 10702
- Hierarchy Level | 10702
- Description | 10703
- Options | 10703
- Required Privilege Level | 10703
- Release Information | 10703

Syntax

```
microloop-avoidance {
  post-convergence-path {
   delay delay;
    maximum-labels maximum-labels;
    }
  }
```

Hierarchy Level

```
[edit logical-systems name protocols ospf spf-options],
[edit logical-systems name routing-instances name protocols ospf spf-options],
[edit logical-systems name tenants name routing-instances name protocols ospf spf-options],
[edit routing-instances name protocols ospf spf-options],
[edit tenants name routing-instances name protocols ospf spf-options],
[edit protocols ospf spf-options]
```

Description

Configure microloop avoidance mechanisms that temporarily install a post-convergence path for routes potentially affected by microloops. The temporary path is used for the specified delay time period and reverts to the node SIDs of the destination after the delay timer expires.

> **NOTE**: Micro-loop avoidance is not a replacement for local repair mechanisms like TI-LFA which detects local failure very fast and activates a precomputed loop-free-alternative path.

Options

post-convergence-path delay *delay*;	Time in milliseconds after which temporary post-convergence paths are removed. • Range: 500 through 60000 • Default: 5000ms
maximum-labels *maximum-labels*;	Maximum number of labels in a microloop avoidance path. Number of labels must be less than or equal to the maximum labels configured. Range: 2 through 8

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 22.1R1.

mif (TDF Interface)

IN THIS SECTION

Syntax | 10704

- Hierarchy Level | 10704
- Description | 10704
- Required Privilege Level | 10704
- Release Information | 10705

Syntax

```
mif {
    mtu;
    unit interface-unit-number {
        family family-name {
            service {
                input service-set;
                output service-set;
            }
        }
    }
}
```

Hierarchy Level

```
[edit interfaces]
```

Description

Configure the TDF interfaces for the TDF domains. A TDF interface is distinct from other types of interfaces and is used to associate a TDF domain's subscribers with an access interface in a virtual routing and forwarding table (VRF). You need to configure one TDF interface logical interface (unit) for every TDF domain.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a TDF Logical Interface

mime-pattern

IN THIS SECTION

- Syntax | 10705
- Hierarchy Level | 10705
- Description | 10706
- Options | 10706
- Required Privilege Level | 10706
- Release Information | 10706

Syntax

```
mime-pattern object-name {
    value [value];
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm custom-objects]
```

Description

The gateway device uses MIME (Multipurpose Internet Mail Extension) types to decide which traffic is allowed to bypass various types of scanning.

Options

- *object-name*—Name of the MIME object.
- value *value*—Value of the MIME object. You can configure multiple values separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The [edit security utm default-configuration] hierarchy level is introduced in Junos OS Release 18.2R1.

mime-whitelist

IN THIS SECTION

- Syntax | 10707
- Hierarchy Level | 10707
- Description | 10707
- Options | 10707
- Required Privilege Level | 10708
- Release Information | 10708

Syntax

```
mime-whitelist {
    exception listname;
    list listname {
        exception listname;
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus]
```

Description

Configure the antivirus scanner to use MIME bypass lists and exception lists. You can use your own custom object lists, or you can use the default list that ships with the device called junos-default-bypass-mime.

WARNING: When you configure the MIME allowlist feature, be aware that, because header information in HTTP traffic can be spoofed, you cannot always trust HTTP headers to be legitimate. When a Web browser is determining the appropriate action for a given file type, it detects the file type without checking the MIME header contents. However, the MIME allowlist feature does refer to the MIME encoding in the HTTP header. For these reasons, it is possible in certain cases for a malicious website to provide an invalid HTTP header. For example, a network administrator might inadvertently add a malicious website to a MIME allowlist, and, because the site is in the allowlist, it will not be blocked by Sophos even though Sophos has identified the site as malicious in its database. Internal hosts would then be able to reach this site and could become infected.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5. Statement updated.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

min-advertisement-interval (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 10708
- Hierarchy Level | 10708
- Description | 10709
- Options | 10709
- Required Privilege Level | 10709
- Release Information | 10709

Syntax

```
min-advertisement-interval seconds;
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Minimum interval between each router advertisement message.

Options

seconds—Minimum interval.

- **Range:** 3 seconds through three-quarter times the maximum advertisement interval value
- **Default:** One-third the maximum advertisement interval value

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA

Dynamic Router Advertisement Configuration Overview

min-advertisement-interval (Protocols IPv6 Neighbor Discovery)

IN THIS SECTION

- Syntax | 10710
- Hierarchy Level | 10710
- Description | 10710
- Options | 10710
- Required Privilege Level | 10710

Release Information | 10711

Syntax

```
min-advertisement-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Description

Set the minimum interval between each router advertisement message.

The router sends router advertisements on each interface configured to transmit messages. The advertisements include route information and indicate to network hosts that the router is operational. The router sends these messages periodically, with a time range defined by minimum and maximum values.

Options

seconds—Minimum interval.

- **Range:** 3 seconds through three-quarter times the maximum advertisement interval value
- **Default:** One-third the maximum advertisement interval value

  By default, the maximum advertisement interval is 600 seconds and the minimum advertisement interval is one-third the maximum interval, or 200 seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

max-advertisement-interval (Protocols IPv6 Neighbor Discovery) | 10415

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

min-advertisement-interval (Protocols Router Discovery)

IN THIS SECTION

- Syntax | 10711
- Hierarchy Level | 10711
- Description | 10712
- Options | 10712
- Required Privilege Level | 10712
- Release Information | 10712

Syntax

```
min-advertisement-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery interface interface-name],
[edit protocols router-discovery interface interface-name]
```

Description

Specify the minimum time the router waits before sending router advertisement packets out the interface in response to router solicitation packets it receives from a client. These packets are broadcast or multicast, depending on how the address corresponding to this physical interface is configured.

Options

seconds—Minimum time between router advertisement packets.

- **Range:** 3 seconds through 1800 seconds
- **Default:** 400 seconds (0.75 times the default value for the `max-advertisement-interval` statement)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

broadcast (Router Discovery) | **3043**

max-advertisement-interval (Protocols Router Discovery) | **10416**

multicast (Protocols Router Discovery) | **11077**

min-message-length

IN THIS SECTION

- Syntax | 10713
- Hierarchy Level | 10713
- Description | 10713
- Options | 10713
- Required Privilege Level | 10714
- Release Information | 10714

Syntax

```
min-message-length number;
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

In the GPRS tunneling protocol (GTP) header, the message length field indicates the length, in octets, of the GTP payload. Use the min-message-length option to set the minimum message payload length (in bytes). The minimum GTP message length is 0 bytes. The message length range is from 1 through 65,535 bytes.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name]` hierarchy level is replaced by `[edit security gtp profile profile-name]`.

Options

number—To set the minimum message payload length in bytes.

- **Range:** 0 through 65,535 bytes.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. Support.

RELATED DOCUMENTATION

Understanding GTP Message Filtering

Example: Setting the GTP Message-Length Filtering

min-rate

IN THIS SECTION

- Syntax | 10714
- Hierarchy Level | 10715
- Description | 10715
- Options | 10716
- Required Privilege Level | 10717
- Release Information | 10717

Syntax

```
min-rate {
    rate bps;
    revert-delay seconds;
```

```
    ipmsi-min-rate rate;
    }
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols mvpn hot-root-standby]
```

Description

Fast failover (that is, sub-50ms switch over for C-multicast streams as defined in *Draft Morin L3VPN Fast Failover 05,*) is supported for MPC cards operating in `enhanced-ip` mode that are running next generation (NG) MVPNs with `hot-root-standby` enabled.

Live-live NG MVPN traffic is available by enabling both sender-based reverse path forwarding (RPF) and hot-root standby. In this scenario, any upstream failure in the network can be repaired locally at the egress PE, and fast failover is triggered if the flow rate of monitored traffic falls below the threshold configured for `min-rate`.

In earlier Junos releases, after the switchover, when the primary upstream is restored and is above the `min-rate` threshold for that path, traffic reverted to the primary path from the backup path. In certain scale scenarios, the PFE may switch to the primary path even before all the flows converged in the inclusive tunnel, resulting in a slow start which in turn leads to loss of traffic.

Junos OS Release 22.4R3-S1 introduced the no-revert behavior. When there is a primary upstream failure, the traffic is switched from the primary path to the backup path. Traffic will continue to flow through the backup path as long as the traffic flow rate on the backup path does not go below the configured `min-rate` threshold.

On the egress PE, redundant multicast streams are received from a source that has been multihomed to two or more senders (upstream PEs). Only one stream is forwarded to the customer network, however, because the sender-based RPF running on the egress PE prevents any duplication.

Note that fast failover only supports VRF configured with a virtual tunnel (VT) interface, that is, anchored to a tunnel PIC to provide upstream tunnel termination. Label switched interfaces (LSI) are not supported.

Hot root standby can handle the failure or restoration of one upstream UMH at a time. Multiple network events or failures cannot be handled within 50ms. In scenarios with multiple network events, convergence will normally follow the timeframe of unicast convergence. However if the `revert-delay` timer is active, the router might wait for the timer to expire before installing the updated route. Though the switchover from one path to another is based on the monitored traffic rate, UMH selection is based

on the unicast reachability of the upstream PE. There can be a maximum of two active UMHs at any time from which traffic can be forwarded.

> **NOTE**: `min-rate` is not strictly supported for MPC3 and MPC4 line cards (these cards have multiple lookup chips and an aggregate value is not calculated across chips). So, when setting the rate, choose a value that is high enough to ensure that lookup will be triggered at least once on each chip every 10 milliseconds or less. As a result, for line cards with multiple look up chips, a small percentage of duplicate multicast packets may be observed being leaked to the to the egress interface. This is normal behavior. The re-route is triggered when traffic rate on the primary tunnel hits zero. Likewise, if no packets are detected on any of the lookup chips during the configured interval, the tunnel will go down.

Options

`rate`—Specify a rate to represent the typical flow rate of aggregate multicast traffic from the provider tunnel (P tunnel). Aggregate multicast traffic from the P tunnel is monitored, and if it falls below the threshold set here a failover to the hot-root standby is triggered.

- **Range:** 3 Mb through 100 Gb

`revert-delay` *seconds*—Use the specified interval to allow time for the network to converge when and if the original link comes back online. You can specify a time, in seconds, for the router to wait before updating its multicast routes. For example, if the original link goes down and triggers the switchover to an alternative link, and then the original link comes back up, the update of multicast routes reflecting the new path can be delayed to accommodate the time it may take to for the network to converge back on the original link.

The `revert-delay` timer is triggered when the old primary upstream path is designated as the new backup path and a new primary path is coming up. During this time, the PFE maintains the forwarding states of the old primary path from which the traffic is forwarded until the new primary path is established or until the revert-delay timer expires. UMH selection is based on the unicast reachability of the upstream PE and is the only criterion for triggering the revert-delay timer.

- **Range:** 0 through 600 seconds

`ipmsi-min-rate` (optional)—Configure a separate min-rate threshold for IPMSI traffic. If not configured, the existing configuration for `min-rate` will be applicable to both IPMSI and SPMSI traffic.

Selective traffic is per-flow and would be much lesser than inclusive traffic which is the combined traffic of multiple flows. In a transient state, when the primary ingress PE's forwarding states are being updated to remove the core facing provider tunnel (ptnl) from the multicast S,G flows, traffic continues to be forwarded over the core for the flows which are not yet updated. If the threshold configured for `min-rate` is greatly lesser than the actual inclusive tunnel traffic, even with flows updated to not forward traffic

over the core on the ingress PE, the remaining traffic would still be greater than the threshold, therefore the failover to the backup path is not triggered. You will observe traffic loss for the flows which are already updated on the ingress PE until the overall inclusive traffic rate goes below the configured `min-rate` and the traffic is switched to the backup path. This situation can be avoided with the introduction of `ipmsi-min-rate`, with a separately configurable threshold for inclusive traffic, exclusive of the `min-rate` for all other traffic.

- **Range:** 3 Mb through 100 Gb

> **NOTE**: Configuring `rate` too close to the actual traffic rate, may result in session flaps or cause the session to remain down.

> **NOTE**: `ipmsi-min-rate` is supported only on uKern based FPCs (Flexible PIC Concentrators) on the egress PE.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Non-revertive behavior introduced in Junos OS Release 22.4R3-S1.

`ipmsi-min-rate` introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Understanding Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

Example: Configuring Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

hot-root-standby (MBGP MVPN) | **8184**

min-rate (source-active-advertisement)

IN THIS SECTION

- Syntax | 10718
- Hierarchy Level | 10718
- Description | 10718
- Required Privilege Level | 10719
- Release Information | 10719

Syntax

```
min-rate bps
```

Hierarchy Level

```
[edit logical-systems logical-system--name routing-instances instance-name protocols mvpn mvpn-mode spt-only source-active-advertisement],
[edit routing-instances instance-name protocols mvpn mvpn-mode spt-only source-active-advertisement],
[edit routing-instances instance-name protocols mvpn mvpn-mode spt-only source-active-advertisement]
```

Description

Minimum traffic rate required to advertise Source-Active route (1 to 1000000 bits per second), set on the ingress PEs.

Use the command, for example, to ensure that the egress PEs only receive Source-Active A-D route advertisements from ingress PEs that are receiving traffic at or above a minimum rate, regardless of how many ingress PEs there may be. Only one of the ingress PEs is chosen as the upstream multicast hop (UMH). Traffic flow continues because the egress PE removes its Type 7 advertisements to the old UMH and re-advertises a Type 7 to the new UMH.

The `min-rate` command works by polling traffic stats to determine the traffic rate of each flow on the ingress PE. Rather than advertising the Source-Active A-D route immediately upon learning of the S,G,

the ingress PE waits until the traffic rate reaches the threshold set for `min-rate` before sending the Source-Active A-D route. If the rate then drops below the threshold, the Source-Active A-D route is withdrawn.

To verify that the value is set as expected, you can check whether the Type 5 (Source-Active route) has been advertised using the `show route table` *`vrf`*`.mvpn.0` command. It may take several minutes before you can see the changes in the Source-Active A-D route advertisement after making changes to the `min-rate`.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs

dampen (Protocols MVPN) | **4813**

mine-dhcp-client-options

IN THIS SECTION

- Syntax | **10720**
- Hierarchy Level (EX Series, QFX Series) | **10720**
- Description | **10720**
- Default | **10720**
- Required Privilege Level | **10720**
- Release Information | **10720**

Syntax

```
mine-dhcp-client-options;
```

Hierarchy Level (EX Series, QFX Series)

```
[edit vlans vlan-name forwarding-options dhcp-security]
```

Description

Enables DHCP snooping to collect various DHCP client options for analytics.

Default

Disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.1.

RELATED DOCUMENTATION

dhcp-security | 5640

show dhcp-security binding | 19344

mine-dhcpv6-client-options

IN THIS SECTION

- Syntax | 10721
- Hierarchy Level (EX Series, QFX Series) | 10721
- Description | 10721
- Default | 10721
- Required Privilege Level | 10721
- Release Information | 10722

Syntax

```
mine-dhcpv6-client-options;
```

Hierarchy Level (EX Series, QFX Series)

```
[edit vlans vlan-name forwarding-options dhcp-security]
```

Description

Enables DHCP snooping to collect various DHCPv6 client options for analytics.

Default

Disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.1.

RELATED DOCUMENTATION

dhcp-security | 5640

show dhcp-security ipv6 binding | 19351

minimum-bandwidth

IN THIS SECTION

- Syntax | 10722
- Hierarchy Level | 10722
- Description | 10723
- Options | 10723
- Required Privilege Level | 10723
- Release Information | 10723

Syntax

```
minimum-bandwidth bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Set the minimum bandwidth in bps for an LSP with automatic bandwidth allocation enabled.

> **NOTE**: For a label-swtiched path (LSP) that has both `bandwidth` and `minimum-bandwidth` for autobandwidth configured under the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level, the LSP bandwidth is adjusted differently.
>
> The LSP is initiated with the bandwidth value configured under the `bandwidth` statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level. At the expiry of the `adjust-interval` timer, the LSP bandwidth gets adjusted based on the traffic flow.
>
> If the bandwidth to be signaled is less than the value configured under the `minimum-bandwidth` statement at the `[edit protocols mpls label-switched-path lsp-name autobandwidth]` hierarchy level, then the LSP is signaled only using the minimum bandwidth.
>
> If the bandwidth to be signaled is greater than the value configured under the `maximum-bandwidth` statement at the `[edit protocols mpls label-switched-path lsp-name autobandwidth]` hierarchy level, then the LSP is signaled only using the maximum bandwidth.

Options

bps—Miniminum bandwidth for the LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth*

minimum-bandwidth-adjust-interval

IN THIS SECTION

- Syntax | 10724
- Hierarchy Level | 10724
- Description | 10724
- Options | 10724
- Required Privilege Level | 10724
- Release Information | 10725

Syntax

```
minimum-bandwidth-adjust-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Specify the duration (in seconds) for which minimum bandwidth is frozen.

Options

seconds—Minimum bandwidth reallocation interval, in seconds.

- **Range:** 300 through 31,536,000 seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth

minimum-bandwidth-adjust-threshold-change

IN THIS SECTION

- Syntax | 10725
- Hierarchy Level | 10725
- Description | 10726
- Options | 10726
- Required Privilege Level | 10726
- Release Information | 10726

Syntax

```
minimum-bandwidth-adjust-threshold-change percentage;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Specify the percentage change in maximum average bandwidth to freeze the minimum bandwidth.

Options

percentage—Percentage change in maximum average bandwidth.

- **Range:** Range: 0 through 100 percent.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

> *Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth*

minimum-bandwidth-adjust-threshold-value

IN THIS SECTION

- Syntax | **10727**
- Hierarchy Level | **10727**
- Description | **10727**
- Options | **10727**
- Required Privilege Level | **10727**
- Release Information | **10727**

Syntax

```
minimum-bandwidth-adjust-threshold-value bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Specify the value in bits per second (bps) to freeze the minimum bandwidth if the maximum average bandwidth falls below this value.

Options

bps—Threshold value for minimum bandwidth if the maximum average bandwidth falls below the specified value.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

> *Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth*

minimum-bandwidth (aggregated Ethernet)

IN THIS SECTION

- Syntax | 10728
- Hierarchy Level | 10728
- Description | 10728
- Options | 10729
- Required Privilege Level | 10729
- Release Information | 10729

Syntax

```
minimum-bandwidth bw-unit unit bw-value value;
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

Description

Configure the minimum bandwidth unit for an aggregated Ethernet bundle as bps, Gbps, Kbps, or Mbps and the bandwidth value from 1 through 128,000.

(T Series, MX Series, PTX Series routers) You cannot configure the minimum number of links and the minimum bandwidth for an aggregated Ethernet bundle at the same time. They are mutually exclusive. To determine the status of the bundle, the device compares the value configured for minimum links and the value for minimum bandwidth. Because both cannot be configured at the same time, the device compares the configured value of the parameter with the default value of the other parameter. The device picks the higher value of the two parameters to determine the status of the Ethernet bundle. Table 64 on page 10729 describes how the device determines the bundle status based on sample values assigned to both parameters

Table 64: Determination of Bundle Status based on Minimum links and Minimum bandwidth Parameters

Current Bandwidth	Minimum Links	Minimum Bandwidth	Bundle Status
100G (10x10G)	2 Bandwidth=20G (2x10G)	N/A Default (1 bps)	Up
50G (5x10G)	N/A Default (1 link) Bandwidth = 10G	20G	Up
50G (5x10G)	N/A Default (1 link) Bandwidth = 10G	60G	Down

Options

unit—Minimum bandwidth unit for the aggregated Ethernet bundle as bps, Gbps, Kbps, or Mbps.

value—Minimum bandwidth value from 1 through 128,000.

- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 14.1R1 and 14.2.

RELATED DOCUMENTATION

- Aggregated Ethernet Interfaces Overview
- Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles
- Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles

minimum-effective-aigp

IN THIS SECTION

- Syntax | 10730
- Hierarchy Level | 10730
- Description | 10730
- Options | 10731
- Required Privilege Level | 10731
- Release Information | 10731

Syntax

```
minimum-effective-aigp <metric-offset>;
```

Hierarchy Level

```
[edit fabric protocols bgp group name metric-out],
[edit logical-systems name protocols bgp group name metric-out],
[edit logical-systems name routing-instances name protocols bgp group name metric-out],
[edit logical-systems name tenants name routing-instances name protocols bgp group name metric-out],
[edit protocols bgp group name metric-out],
[edit routing-instances name protocols bgp group name metric-out],
[edit tenants name routing-instances name protocols bgp group name metric-out]
```

Description

To track the minimum effective AIGP metric. Effective AIGP is the AIGP value advertised with the route plus the IGP cost to reach the nexthop. You can configure a minimum-aigp to prevent unnecessary updation of route when effective-aigp changes past the previously known lowest value.

Options

metric-offset	Metric offset for MED.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Example: Configuring the Accumulated IGP Attribute for BGP

minimum-hold-time

IN THIS SECTION

- Syntax | 10731
- Hierarchy Level | 10732
- Description | 10732
- Options | 10732
- Required Privilege Level | 10732
- Release Information | 10732

Syntax

```
minimum-hold-time time-in-seconds;
```

Hierarchy Level

```
[edit protocols bgp],
[edit protocols bgp group group-name  ],
[edit protocols bgp group group-name  neighborip-address  ]
```

Description

Prevent BGP session establishment toward BGP peers that attempt to negotiate a lower BGP session hold-time than the configured `minimum-hold-time`. Such BGP peers will be rejected. BGP session establishment attempt will be dropped and a notification message with `unacceptable hold time` error will be sent to the neighbor. This helps reduce the load on the router by avoid the sending of constant keepalive messages at a high frequency.

> **NOTE**: We recommend using Bidirectional Forwarding Detection (BFD) rather than lowering BGP hold timers and also recommend configuring a meaningful `minimum-hold-time` value (for example, 20 seconds or higher) for all BGP peers (at the BGP group level).
>
> If a BGP remote node does not support BFD, and therefore a lower BGP hold-time is desired for the quicker discovery of a BGP neighbor failure, you can configure a smaller `minimum-hold-time` value. However, use it with caution and only for a limited number of BGP peers.

Options

time-in-seconds	Specify hold time in seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 19.3.

RELATED DOCUMENTATION

hold-time (Protocols BGP)

minimum-interval

IN THIS SECTION

- Syntax | 10733
- Hierarchy Level | 10733
- Description | 10734
- Options | 10734
- Required Privilege Level | 10734
- Release Information | 10734

Syntax

```
minimum-interval milliseconds;
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server group group-name liveness-detection method bfd],
[edit system services dhcp-local-server group group-name liveness-detection method bfd transmit-
interval],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd
```

```
transmit-interval],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd transmit-
interval],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd
transmit-interval]
```

Description

Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the `transmit-interval minimal-interval` and `minimum-receive-interval` statements.

Options

milliseconds — Specify the minimum interval value for BFD liveliness detection.

- **Range:** 1 through 255,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients

Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

minimum-interval (BFD Liveness Detection)

IN THIS SECTION

- Syntax | 10735
- Hierarchy Level | 10735
- Description | 10736
- Options | 10736
- Required Privilege Level | 10736
- Release Information | 10736

Syntax

```
minimum-interval milliseconds;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-
```

```
detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam
bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-
name neighbor neighbor-id oam bfd-liveness-detection]
```

Description

Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the `minimum-interval` (specified under the `transmit-interval` statement) and `minimum-receive-interval` statements.

Options

milliseconds—Specify the minimum interval value for BFD liveliness detection.

- **Range:** 1 through 255,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring BFD for Layer 2 VPN and VPLS

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

bfd-liveness-detection (Layer 2 VPN and VPLS) | **2929**

minimum-receive-interval (BFD Liveness Detection) | **10754**

transmit-interval (BFD Liveness Detection) | **30087**

minimum-interval (Liveness Detection)

IN THIS SECTION

- Syntax | 10737
- Hierarchy Level | 10737
- Description | 10737
- Options | 10737
- Required Privilege Level | 10738
- Release Information | 10738

Syntax

```
minimum-interval milliseconds;
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection]
```

Description

Configure simultaneously the minimum interval at which the peer transmits liveness detection requests and the minimum interval at which the peer expects to receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately by using the `transmit-interval minimal-interval` and `minimum-receive-interval` statements, respectively.

Options

milliseconds—Specify the minimum interval value for Bidirectional Forwarding Detection (BFD).

- **Range:** 1 through 255,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

minimum-interval (PIM BFD Liveness Detection)

IN THIS SECTION

- Syntax | 10738
- Hierarchy Level | 10738
- Description | 10739
- Options | 10739
- Required Privilege Level | 10739
- Release Information | 10739

Syntax

```
minimum-interval milliseconds;
```

Hierarchy Level

```
[edit protocols pim interface (Protocols PIM) interface-name bfd-liveness-detection],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM) interface-
name bfd-liveness-detection]
```

Description

Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and `minimum-receive-interval` statements.

Options

milliseconds—Minimum transmit and receive interval.

- **Range:** 1 through 255,000 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring BFD for PIM

minimum-interval (PIM BFD Transmit Interval)

IN THIS SECTION

- Syntax | **10740**
- Hierarchy Level | **10740**
- Description | **10740**
- Options | **10740**
- Required Privilege Level | **10740**

Release Information | 10741

Syntax

```
minimum-interval milliseconds;
```

Hierarchy Level

```
[edit protocols pim interface interface-name bfd-liveness-detection transmit-interval],
[edit routing-instances routing-instance-name protocols pim interface interface-name bfd-
liveness-detection transmit-interval]
```

Description

Configure the minimum interval after which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the `minimum-interval` statement at the `[edit protocols pim interface interface-name bfd-liveness-detection]` hierarchy level.

Options

milliseconds—Minimum transmit interval value.

- **Range:** 1 through 255,000

> **NOTE**: The threshold value specified in the `threshold` statement must be greater than the value specified in the `minimum-interval` statement for the `transmit-interval` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Support for BFD authentication introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD for PIM

bfd-liveness-detection (Protocols PIM) | **2945**

minimum-interval (PIM BFD Liveness Detection) | **10738**

threshold (PIM BFD Transmit Interval) | **29333**

minimum-links

IN THIS SECTION

- Syntax (SRX, MX, T, M, EX, QFX Series, EX4600, Qfabric System) | **10741**
- Hierarchy Level (EX Series) | **10742**
- Hierarchy Level (QFX Series) | **10742**
- Description | **10742**
- Options | **10743**
- Required Privilege Level | **10743**
- Release Information | **10744**

Syntax (SRX, MX, T, M, EX, QFX Series, EX4600, Qfabric System)

```
minimum-links number;
```

Hierarchy Level (EX Series)

```
[edit interfaces aex aggregated-ether-options],
[edit interfaces aex aggregated-sonet-options],
[edit interfaces interface-name mlfr-uni-nni-bundle-options],
[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-range range aggregated-ether-options],
[edit interfaces interface-range range aggregated-sonet-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Hierarchy Level (QFX Series)

```
[edit interfaces aex aggregated-ether-options]
```

Description

For aggregated Ethernet, SONET/SDH, multilink, link services, and voice services interfaces only, set the minimum number of links that must be up for the bundle to be labeled up.

(T Series, MX Series, PTX Series routers) You cannot configure the minimum number of links and the minimum bandwidth of an aggregated Ethernet bundle at the same time. They are mutually exclusive. To determine the status of the bundle, the device compares the value configured for minimum links and the value for minimum bandwidth. Because both cannot be configured at the same time, the device compares the configured value of the parameter with the default value of the other parameter. The device picks the higher value of the two parameters to determine the status of the Ethernet bundle. Table 65 on page 10742 describes how the device determines the bundle status based on sample values assigned to both parameters

Table 65: Determination of Bundle Status based on Minimum links and Minimum bandwidth Parameters

Current Bandwidth	Minimum Links	Minimum Bandwidth	Bundle Status
100G (10x10G)	2 Bandwidth=20G (2x10G)	N/A Default (1 bps)	Up

Table 65: Determination of Bundle Status based on Minimum links and Minimum bandwidth Parameters *(Continued)*

Current Bandwidth	Minimum Links	Minimum Bandwidth	Bundle Status
50G (5x10G)	N/A Default (1 link) Bandwidth = 10G	20G	Up
50G (5x10G)	N/A Default (1 link) Bandwidth = 10G	60G	Down

Options

number—Number of links.

- **Range:** On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, the valid range for minimum-links number is 1 through 64. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

  On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for minimum-links number is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

  On EX8200 switches, the range of valid values for minimum-links number is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

  On EX4600, QFX Series and Q Fabric Systems , the range of valid values for minimum-links number is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

- **Default:** 1

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Configuring Aggregated Ethernet Minimum Links*
- Configuring Aggregated SONET/SDH Interfaces
- Configuring Aggregated Ethernet Links (CLI Procedure)
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- Junos OS Services Interfaces Library for Routing Devices
- *Configuring Link Aggregation*

minimum-links (candidate-uplink-port-profile)

IN THIS SECTION

- Syntax | **10744**
- Hierarchy Level | **10745**
- Description | **10745**
- Default | **10745**
- Options | **10745**
- Required Privilege Level | **10745**
- Release Information | **10745**

Syntax

```
minimum-links link-count;
```

Hierarchy Level

```
[edit policy-options satellite-policies candidate-uplink-port-profile profile-name]
[edit policy-options satellite-policies candidate-uplink-port-profile profile-name term term-name from]
```

Description

Configures the minimum number of active uplink ports a satellite device must have. If a satellite device has fewer than this number of active uplink ports, uplink failure detection shuts down its extended ports.

Default

The default number of minimum links is 1.

Options

link-count—Specifies the minimum number of active uplink ports a satellite device must have. Valid values are 1-32 links.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

minimum-links (Interfaces)

IN THIS SECTION

- Syntax | 10746
- Hierarchy Level | 10746
- Description | 10746
- Options | 10747
- Required Privilege Level | 10747
- Release Information | 10747

Syntax

```
minimum-links                    number;
```

Hierarchy Level

```
[edit interfaces                    interface-name                    redundant-ether-
options]
```

Description

For redundant Ethernet interfaces configured as 802.3ad redundant Ethernet interface link aggregation groups (LAGs) in a chassis cluster only, set the required minimum number of physical child links on the primary node that must be working to prevent the interface from being down. Interfaces configured as redundant Ethernet interface LAGs typically have between 4 and 16 physical interfaces, but only half, those on the primary node, are relevant to the minimum-links setting.

If the number of operating interfaces on the primary node falls below the configured value, it will cause the interface to be down even if some of the interfaces are still working.

For an aggregated ethernet interface, you cannot configure all three configuration options, `bfd-liveness-detection`, `minimum-links`, and `sync-reset` at the same time.

Options

number—For redundant Ethernet interface link aggregation group links, specify the number of physical child links on the primary node in the redundant Ethernet interface that must be working. The default **minimum-links** value is 1. The maximum value is half of the total number of physical child interfaces bound to the redundant Ethernet interface being configured or 8, whichever is smaller.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement added in Release 10.1 of Junos OS.

minimum-priority

IN THIS SECTION

- Syntax | 10747
- Hierarchy Level | 10748
- Description | 10748
- Options | 10748
- Required Privilege Level | 10748
- Release Information | 10748

Syntax

```
minimum-priority value;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

Specify the minimum priority for the control source.

Options

value—Minimum priority value; if not specified, defaults to 0.

- **Range:** 0 through 254

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring the Control Source

minimum-quality

IN THIS SECTION

- Syntax | 10749
- Hierarchy Level | 10749
- Description | 10749

Required Privilege Level | 10750
Release Information | 10750

Syntax

```
minimum-quality;
```

Hierarchy Level

```
[edit chassis synchronization output interfaces external]
[edit chassis synchronization output interfaces (external-0/0 | external-1/0)]
[edit chassis synchronization]
```

Hierarchy level `[edit chassis synchronization]` added for ACX710 in Junos OS Release 20.2R1.

Description

Configure the minimum quality level threshold to select a clock source (see Table 66 on page 10749). If the quality level of the output clock source drops below the configured minimum quality level threshold, the external output clock is suppressed on the external output interface—external—on SCBE and on the external output interfaces—external-0/0 and external-1/0—on SCBE2.

Table 66: Quality Levels

Quality Level	Description
prc	Timing quality of a primary reference clock (option-1 only).
prs	Clock traceable to a primary reference source (option-2 only).
sec	Timing quality of an SDH equipment clock (option-1 only).
smc	Clock traceable to a self-timed SONET (option-2 only).

Table 66: Quality Levels *(Continued)*

Quality Level	Description
ssu-a	Timing quality of a type I or IV timeReceiver clock (option-1 only).
ssu-b	Timing quality of a type VI timeReceiver clock (option-1 only).
st2	Clock traceable to Stratum 2 (option-2 only).
st3	Clock traceable to Stratum 3 (option-2 only).
st3e	Clock traceable to Stratum 3E (option-2 only).
st4	Clock traceable to Stratum 4 free-run (option-2 only).
stu	Clock traceable to an unknown quality (option-2 only).
tnc	Clock traceable to a transit node clock (option-2 only).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (ACX Series) | **28787**

synchronization (MX Series) | **28795**

minimum-receive-interval (Protocols PIM)

IN THIS SECTION

- Syntax | 10751
- Hierarchy Level | 10751
- Description | 10751
- Options | 10751
- Required Privilege Level | 10752
- Release Information | 10752

Syntax

```
minimum-receive-interval milliseconds;
```

Hierarchy Level

```
[edit protocols pim interface (Protocols PIM) interface-name bfd-liveness-detection],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM) interface-
name bfd-liveness-detection]
```

Description

Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the `minimum-interval` statement at the `[edit protocols pim interface interface-name bfd-liveness-detection]` hierarchy level.

Options

milliseconds—Minimum receive interval.

- **Range:** 1 through 255,000 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring BFD for PIM

minimum-receive-interval

IN THIS SECTION

- Syntax | 10752
- Hierarchy Level | 10753
- Description | 10753
- Options | 10753
- Required Privilege Level | 10753
- Release Information | 10753

Syntax

```
minimum-receive-interval milliseconds;
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options
dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group group-name liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd]
```

Description

Configure the minimum interval at which the local routing device (or switch) must receive a reply from a neighbor with which it has established a BFD session.

Options

milliseconds — Specify the minimum receive interval value.

- **Range:** 1 through 255,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients

Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

minimum-receive-interval (BFD Liveness Detection)

IN THIS SECTION

- Syntax | 10754
- Hierarchy Level | 10754
- Description | 10755
- Options | 10755
- Required Privilege Level | 10755
- Release Information | 10755

Syntax

```
minimum-receive-interval milliseconds;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls) mesh-group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols evpn | l2vpn | vpls) oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols evpn | l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols evpn | l2vpn | vpls) mesh-group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls)  oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls)  neighbor neighbor-id  oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls)  mesh-group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-
```

```
detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam
bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) mesh-group mesh-group-
name neighbor neighbor-id oam bfd-liveness-detection]
```

Description

Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the `minimum-interval` statement.

Options

milliseconds—Specify the minimum receive interval value.

- **Range:** 1 through 255,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Support for BFD authentication introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD for Layer 2 VPN and VPLS

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

bfd-liveness-detection (Layer 2 VPN and VPLS) | **2929**

minimum-interval (BFD Liveness Detection) | **10735**

transmit-interval (BFD Liveness Detection) | **30087**

minimum-receive-interval (Liveness Detection)

IN THIS SECTION

- Syntax | 10756
- Hierarchy Level | 10756
- Description | 10756
- Options | 10756
- Required Privilege Level | 10756
- Release Information | 10757

Syntax

```
minimum-receive-interval milliseconds;
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection]
```

Description

Configure the minimum interval at which the peer must receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session.

Options

milliseconds—Specify the minimum interval value.

- **Range:** 1 through 255,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

minimum-retransmission-timeout (L2TP Tunnel)

IN THIS SECTION

- Syntax | 10757
- Hierarchy Level | 10757
- Description | 10757
- Options | 10758
- Required Privilege Level | 10758
- Release Information | 10758

Syntax

```
minimum-retransmission-timeout seconds;
```

Hierarchy Level

```
[edit services l2tp tunnel]
```

Description

Configure the minimum (initial) interval that the LAC or the LNS waits for a response after transmitting an L2TP control message to a peer. If no response has been received by the time the period expires, the message is retransmitted. The timeout period is doubled for each retransmission until the maximum of 16 seconds is reached.

Options

seconds—Minimum interval before initial retransmission.

- **Range:** 1, 2, 4, 8, or 16 seconds
- **Default:** 1

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

L2TP for Subscriber Access Overview

L2TP LAC Subscriber Configuration

L2TP LNS Inline Service Interfaces

minimum-wait-time

IN THIS SECTION

- Syntax | 10759
- Hierarchy Level | 10759
- Description | 10759
- Options | 10759
- Required Privilege Level | 10759
- Release Information | 10759

Syntax

```
minimum-wait-time seconds;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface (interface-name | interface-group)]
```

Description

To set the minimum allowed number of seconds in the secs field of the BOOTP message, include the minimum-wait-time statement. This setting configures a minimum number of seconds since the client sent its first BOOTP request. BOOTP messages that have a smaller number in the secs field than the allowed minimum are not forwarded. The default value for the minimum wait time is zero (0).

The default value for the minimum wait time is zero (0) seconds. If the minimum wait time is 0 and the secs field in the BOOTP request message is 0, the device forwards the packet.

Options

seconds Minimum wait time the BOOTP client has waited before packets are forwarded.

- **Range:** 0 to 30,000
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

minor

IN THIS SECTION

- Syntax | 10760
- Hierarchy Level | 10760
- Description | 10761
- Required Privilege Level | 10761
- Release Information | 10761

Syntax

```
minor {
    threshold threshold-value;
    action (chassis error) {
        alarm;
        disable-pfe;
        get-state;
        log;
        offline;
        reset;
    }
}
```

Hierarchy Level

[edit chassis "fpc" on page 7516 *slot-number* "error " on page 6441]

[edit chassis]

Description

Severity level of the error. An error that results in the loss of a single packet but is fully recoverable is a minor error. The severity level of an error cannot be configured by a user.

The other statements are explained separately.

Required Privilege Level

interface	To view this statement in the configuration.
interface-control	To add this statement to the configuration.

Release Information

Statement introduced for PTX Series routers in Junos OS Release 13.3.

Statement introduced for MX240, MX480, MX960, and MX2020 routers in Junos OS Release 14.2.

RELATED DOCUMENTATION

Fabric Resiliency

show chassis fpc errors | **18387**

Managing Errors

minor-event (Process)

IN THIS SECTION

- Syntax | **10762**
- Hierarchy Level | **10762**
- Description | **10762**
- Options | **10762**
- Required Privilege Level | **10762**
- Release Information | **10763**

Syntax

```
minor-event {
    threshold process name memory-limit value ;
    action snmp-trap;
}
```

Hierarchy Level

```
[edit system monitor memory process]
```

Description

Indicates a non-critical event that, if left unchecked, may cause an interruption in service or degradation in performance but does not significantly affect the device. You can specify the process threshold value for specific processes and their memory limit. When the threshold is crossed, you can configure the action to be performed. For instance, you can specify the process threshold for the `rpd` process and the memory limit for that process. If the system detects that the memory utilized by the process is greater than the threshold memory limit, based on the value, a minor event is triggered. Usually, for minor events, an SNMP trap is raised.

Options

threshold Specifies the process threshold level for a minor event. You must specify the process name and the memory limit for the process.

memory Specifies the permissible memory limit for the process after which the minor event is triggered.

action Specifies the action to be performed when the threshold is crossed. Usually, in case of minor events, an SNMP trap is raised.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1.

RELATED DOCUMENTATION

System Log Overview

System Log Explorer

minor-event (System)

IN THIS SECTION

- Syntax | 10763
- Hierarchy Level | 10763
- Description | 10764
- Options | 10764
- Required Privilege Level | 10764
- Release Information | 10764

Syntax

```
minor-event {
    threshold swap-usage value ;
    action snmp-trap ;
}
```

Hierarchy Level

```
[edit system monitor memory system]
```

Description

Indicates a non-critical event that, if left unchecked, may cause an interruption in service or degradation in performance but does not significantly affect the device. You can monitor the swap usage and specify the threshold for swap space usage. Most devices have a swap partition that you can configure to provide additional transient memory. Ideally, the swap partition is used only under exceptional circumstances. In current dynamic systems, you can sporadically use the swap partition as long as the usage is not high or persistent. For instance, you could categorise a minor event as one where the swap usage is greater than 20%.

You can only configure swap usage threshold for memory monitoring. If you do not configure the swap usage threshold, memory monitoring for swap usage is based on the default threshold. Memory monitoring for other features is based on the default threshold.

Options

threshold Specifies the process threshold level for a minor event. You must specify the process name and the memory limit for the process.

swap-usage Specifies the swap usage in percentage.

action Specifies the action to be performed when the threshold is crossed.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1.

RELATED DOCUMENTATION

System Log Overview

System Log Explorer

mip-half-function

IN THIS SECTION

- Syntax | 10765
- Hierarchy Level | 10765
- Description | 10765
- Options | 10765
- Required Privilege Level | 10766
- Release Information | 10766

Syntax

```
mip-half-function (none | default | explicit);
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-managementmaintenance-domain md-name],
[edit protocols oam ethernet connectivity-fault-managementmaintenance-association ma-name]
```

Description

Specify the OAM Ethernet CFM maintenance domain MIP half functions.

NOTE: Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the `mip-half-function` value for all maintenance domains and maintenance associations are the same.

Options

`none`—Specify to not use the mip-half-function.

`default`—Specify to use the default mip-half-function.

`explicit`—Specify an explicit mip-half-function.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Creating a Maintenance Domain

maintenance-domain | 10287

mirror-decrypt-traffic

IN THIS SECTION

- Syntax | 10766
- Hierarchy Level | 10767
- Description | 10767
- Options | 10767
- Required Privilege Level | 10768
- Release Information | 10768

Syntax

```
mirror-decrypt-traffic {
    interface interface-name;
    only-after-secruity-policies-enforcement;
```

```
    destination-mac-address mac-address;
}
```

Hierarchy Level

```
[edit services ssl proxy profile profile-name]
```

Description

Specify SSL decryption mirroring options to forward the copy of SSL decrypted traffic to an external traffic collection device.

To use SSL decryption mirroring, configure the SSL decryption port mirroring interface on SRX Series Firewall and MAC address of the of the external mirror traffic collector port in an SSL proxy profile. Next, apply the SSL proxy profile as application services in the security policy. The SSL traffic matching the security policy rule is decrypted and a copy of the decrypted traffic is forwarded to an external traffic collection device through the SSL decryption port mirroring interface.

Options

interface

SSL decryption port mirroring interface on SRX Series Firewall. This is an Ethernet interface on SRX Series Firewall through which the copy of the SSL decrypted traffic is forwarded to a mirror port.

only-after-secruity-policies-enforcement

Enables forwarding the copy of the decrypted traffic to the external mirror traffic collector after enforcing the Layer 7 security services through a security policy.

By default, forwarding of the SSL decrypted payload to the external mirror traffic collector port occurs before enforcing Layer 7 security services including IDP, Juniper ATP Cloud, and Content Security. When you select to forward the copy of the decrypted traffic after security policies enforcement, and if the decrypted payload is modified while enforcing the security policy, the modified decrypted payload is forwarded to external traffic collection device. Similarly, if the decrypted traffic is dropped because of policy enforcement (for example, a threat is detected in the decrypted traffic), that particular decrypted traffic is not forwarded.

destination-mac-address

MAC address of the of the external mirror traffic collector port.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

Firewall User Authentication Overview

mirror-filter (Security Forwarding Options)

IN THIS SECTION

- Syntax | 10768
- Hierarchy Level | 10769
- Description | 10769
- Options | 10770
- Required Privilege Level | 10770
- Release Information | 10770

Syntax

```
mirror-filter filter-name {
    destination-port port-number;
    destination-prefix destination-prefix;
    interface-in interface-name;
    interface-out interface-name;
    output (Security Forwarding Options) {
```

```
        destination-mac mac-address;
        interface interface-name;
    }
    protocol protocol;
    source-port port-number;
    source-prefix set source-prefix;
}
```

Hierarchy Level

```
[edit security forwarding-options]
```

Description

Configure a mirror filter for filtering X2 packets to be mirrored and sent to a packet analyzer.

As a network operator, you need a way to monitor X2 traffic to debug any handover issues across eNodeBs. The mirror filter feature allows you to do that. Traffic coming out of an IPsec tunnel is decrypted, mirrored and analyzed, and then encrypted again to go into the outbound IPsec tunnel.

To use the mirror filter feature to monitor X2 traffic, you configure mirror filters. You can configure up to 15 different mirror filters to be used concurrently to filter for various kinds of traffic. Each mirror filter contains a set of parameters and their values against which traffic is matched.

> **NOTE**: The SRX Series mirror filter feature is bidirectional, much like a session. X2 traffic flowing through an IPSec VPN from devices that match the configured filter conditions is mirrored and analyzed.

Starting in Junos OS Release 18.4R1, if the output X2 interface of a mirror filter is configured for an st0 interface to filter traffic that you want to analyze, the packet is duplicated and encrypted by the IPsec tunnel bound to the st0 interface. This enhancement supports the SRX Series Firewalls to send traffic mirrored from a port on an IPsec tunnel.

In addition to the following parameters for a mirror filter, you specify the output interface and the MAC address of the packet analyzer as part of the configuration.

> **NOTE**: Although there is no minimum required number of parameters for a mirror filter, please be mindful that if you specify too few criteria or accidentally commit an

incomplete filter, an over-proportional amount of traffic flow through the system could be mirrored.

- destination IP address prefix
- destination port
- IP protocol
- source IP address prefix
- source port
- incoming and outgoing interfaces

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

show security forwarding-options mirror-filter | 24363

clear security forward-options mirror filter | 3935

mirror-flash-on-disk

IN THIS SECTION

- Syntax | 10771
- Hierarchy Level | 10771
- Description | 10771
- Options | 10772
- Required Privilege Level | 10772
- Release Information | 10772

Syntax

```
mirror-flash-on-disk;
```

Hierarchy Level

```
[edit system]
```

Description

Configure the hard disk to automatically mirror the contents of the CompactFlash card. The hard disk maintains a synchronized mirror copy of the CompactFlash card contents. Data written to the CompactFlash card is simultaneously updated in the mirrored copy of the hard disk. If the CompactFlash card fails to read data, the hard disk automatically retrieves its mirrored copy of the CompactFlash card.

CAUTION: We recommend that you disable flash disk mirroring when you upgrade or downgrade the router.

You cannot issue the `request system snapshot` command while the `mirror-flash-on-disk` statement is enabled.

> **NOTE**: After you have enabled or disabled the `mirror-flash-on-disk` statement, you must reboot the router for your changes to take effect. To reboot, issue the `request system reboot` command.

Options

This command has no options.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement deprecated for Junos OS with Upgraded FreeBSD in Junos OS Release 15.1.

> **NOTE**: To determine which platforms run Junos OS with Upgraded FreeBSD, see the table listing the platforms currently running Junos OS with upgraded FreeBSD in Release Information for Junos OS with Upgraded FreeBSD.

RELATED DOCUMENTATION

Configuring Automatic Mirroring of the CompactFlash Card on the Hard Drive

mirror-once

IN THIS SECTION

- Syntax | **10773**
- Hierarchy Level | **10773**

- Description | 10773
- Required Privilege Level | 10773
- Release Information | 10773

Syntax

```
mirror-once;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring]
```

Description

Configure the router to mirror packets only once. This feature is useful if you configure port mirroring on both ingress and egress interfaces, which could result in the same packet being mirrored twice.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3 (MX Series routers only).

Support extended to M120 routers in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Port Mirroring

mixed-rate-mode

IN THIS SECTION

- Syntax | 10774
- Hierarchy Level | 10774
- Description | 10774
- Required Privilege Level | 10774
- Release Information | 10774

Syntax

mixed-rate-mode;

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number mixed-rate-mode],
[edit chassis lcc number fpc slot-number pic pic-number mixed-rate-mode](Routing Matrix)
```

Description

Configure the mixed-rate mode for the 24-port 10 Gigabit Ethernet PIC (PF-24XGE-SFPP) only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Modes of Operation

Configuring Mixed-Rate Mode Operation

mka

IN THIS SECTION

- Syntax | 10775
- Syntax (SRX Series Firewalls) | 10775
- Hierarchy Level | 10776
- Description | 10776
- Options | 10776
- Required Privilege Level | 10776
- Release Information | 10776

Syntax

```
mka {
    must-secure;
    key-server-priority priority-number;
    transmit-interval interval;
}
```

Syntax (SRX Series Firewalls)

```
mka {
    bounded-delay;
    disable-icv-indicator;
    eapol-address(lldp-multicast | pae | provider-bridge | unicast-address);
    key-server-priority key-server-priority;
    sak-rekey-interval <varname>seconds</varname>;
```

```
    should-secure;
    suspend-for;
    suspend-on-request;
    transmit-interval milliseconds;
}
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Specify parameters for the MACsec Key Agreement (MKA) protocol. You initially establish a MACsec-secured link using a pre-shared key when you are using static CAK security mode to enable MACsec. Once matching pre-shared keys are successfully exchanged, the MACsec Key Agreement (MKA) protocol is enabled. The MKA protocol is responsible for maintaining MACsec on the link, and decides which switch on the point-to-point link becomes the key server. The key server then creates an SAK that is shared with the switch at the other end of the point-to-point link only, and that SAK is used to secure all data traffic traversing the link.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced for SRX Series Firewalls in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | 10278

mka (MX Series)

IN THIS SECTION

- Syntax | 10777
- Hierarchy Level | 10777
- Description | 10778
- Options | 10778
- Required Privilege Level | 10778
- Release Information | 10778

Syntax

```
mka {
    bounded-delay;
    key-server-priority priority-number;
    must-secure;
    should-secure;
    transmit-interval interval;
    eapol-address (pae | provider-bridge | lldp-multicast | destination unicast-address);
}
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Specify parameters for the MACsec Key Agreement (MKA) protocol.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

Option `eapol-address` introduced in Junos OS Release 18.3R1 for MX Series routers.

Option `bounded-delay` introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

eapol-address (MACSec) | **6181**

mld

IN THIS SECTION

- Syntax | **10779**
- Hierarchy Level | **10780**
- Description | **10780**
- Default | **10780**
- Options | **10780**

- Required Privilege Level | **10780**
- Release Information | **10780**

Syntax

```
mld {
    accounting;
    interface interface-name {
        (accounting | no-accounting);
        disable;
        distributed;
        group-limit  limit;
        group-policy [ policy-names ];
        immediate-leave;
        oif-map [ map-names ];
        passive;
        ssm-map ssm-map-name;
        ssm-map-policy ssm-map-policy-name;
        static {
            group multicast-group-address {
                exclude;
                group-count number;
                group-increment increment;
                source ip-address {
                    source-count number;
                    source-increment increment;
                }
            }
        }
        version version;
    }
    maximum-transmit-rate packets-per-second;
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit protocols]
```

Description

Enable MLD on the router. MLD must be enabled for the router to receive multicast packets.

Default

MLD is disabled on the router. MLD is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling MLD

show mld group | 22066

show mld interface | 22074

show mld statistics | 22104

clear mld membership | 3702

clear mld statistics | 3707

mld (Dynamic Profiles)

IN THIS SECTION

- Syntax | 10781
- Hierarchy Level | 10782
- Description | 10782
- Required Privilege Level | 10782
- Release Information | 10782

Syntax

```
mld {
     interfaceinterface-name {
        (accounting | no-accounting);
        disable;
        group-limit limit;
        group-policy;
        group-threshold value;
        immediate-leave;
        log-interval seconds;
        oif-map;
        passive;
        ssm-map ssm-map-name;
        ssm-map-policy ssm-map-policy-name;
        static {
            group multicast-group-address {
                exclude;
                group-count number;
                group-increment increment;
                source ip-address {
                    source-count number;
                    source-increment increment;
                }
            }
        }
        version version;
```

```
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols]
```

Description

Configure interface-specific MLD values on dynamic interfaces.

You can configure MLD in a dynamic client profile or a dynamic service profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Enabling MLD

mld-snooping

IN THIS SECTION

- Syntax (ACX Series, EX9200, and MX Series) | 10783
- Syntax (EX Series and SRX Series) | 10784
- Syntax (QFX Series) | 10785
- Hierarchy Level | 10786
- Description | 10786
- Default | 10787
- Required Privilege Level | 10787
- Release Information | 10787

Syntax (ACX Series, EX9200, and MX Series)

```
mld-snooping {
    evpn-ssm-reports-only;
    immediate-leave;
    interface interface-name {
        group-limit limit;
        host-only-interface;
        immediate-leave;
        multicast-router-interface;
        static {
            group ip-address {
                source ip-address;
            }
        }
    }
    proxy {
        source-address ip-address;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
```

```
    }
    vlan vlan-id {
        immediate-leave;
        interface interface-name {
            group-limit limit;
            host-only-interface;
            immediate-leave;
            multicast-router-interface;
            static {
                group ip-address {
                    source ip-address;
                }
            }
        }
        proxy {
            source-address ip-address;
        }
        query-interval seconds;
        query-last-member-interval seconds;
        query-response-interval seconds;
        robust-count number;
    }
}
```

Syntax (EX Series and SRX Series)

```
mld-snooping   {
    vlan (all | vlan-name) {
        immediate-leave;
        interface interface-name {
            group-limit;
            host-only-interface;
            immediate-leave;
            multicast-router-interface;
            static {
                group ip-address {
                source ip-address;
                }
            }
        }
        qualified-vlan vlan-id;
```

```
        query-interval;
        query-last-member-interval;
        query-response-interval;
        robust-count number;
        trace-options {
            file (files | no-word-readable | size | word-readable):
            flag (all | client-notification | general | group | host-notification | leave |
noraml | packest | policy | query | report | route | report | state | task |timer):
        }
    }
}
```

Syntax (QFX Series)

```
mld-snooping   {
    vlan (vlan-name) {
        evpn-ssm-reports-only;
        immediate-leave;
        interface (all | interface-name) {
            group-limit  limit;
            host-only-interface;
            immediate-leave;
            multicast-router-interface;
            static {
                group ip-address {
                    source ip-address;
                }
            }
        }
        proxy {
            source-address ip-address;
        }
        qualified-vlan vlan-id;
        query-interval seconds;
        query-last-member-interval seconds;
        query-response-interval seconds;
        robust-count number;
        traceoptions  {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier>;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable and configure Multicast Listener Discovery (MLD) snooping. MLD snooping constrains IPv6 multicast traffic at Layer 2 by configuring Layer 2 LAN ports dynamically to forward IPv6 multicast traffic only to those ports that want to receive it.

MLD is a protocol built on ICMPv6 and used by IPv6 routers and hosts to discover and indicate interest in a multicast group, similar to how IGMP manages multicast group membership for IPv4 multicast traffic. There are two versions, MLDv1 (RFC 2710), which is equivalent to IGMP version 2 (IGMPv2), and MLDv2 (RFC 3810), which is equivalent to IGMP version 3 (IGMPv3). Like IGMP, both MLDv1 and MLDv2 support Query, Report and Done messages. MLDv2 further supports source-specific Query messages (reports) and multi-record reports. MLD configuration options are similar to those for IGMP snooping.

MLD restricts forwarding IPv6 multicast traffic to only those interfaces in a bridge-domain, VLAN, or VPLS that have interested listeners, rather than flooding the traffic to all interfaces in the bridge-domain, VLAN, or VPLS. The device finds the interfaces with interested listeners using the following steps:

- Snoops or monitors MLD control packets.
- Identifies the set of outgoing interfaces for a multicast stream.
- Builds the forwarding state accordingly.

The device snoops Query messages and floods them to all ports. The device snoops Report and Done messages and selectively forwards them only to multicast router ports.

> **NOTE**: For MX Series devices, MLD snooping is not supported on DPC linecards. The operational commands for MLD snooping, including defaults, functionality, logging, and tracing are similar to those for IGMP snooping.

The remaining statements are explained separately. See CLI Explorer.

Default

MLD snooping is disabled.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- *Understanding MLD Snooping*
- *Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment*
- *Configuring MLD*
- *Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)*
- *Configuring MLD Snooping on a Switch VLAN with ELS Support (CLI Procedure)*
- *Configure Multicast Forwarding with MLD Snooping in an EVPN-MPLS Environment*
- *Example: Configuring MLD Snooping on EX Series Switches*
- *Example: Configuring MLD Snooping on Switches with ELS Support*

mldp-inband-signalling (Protocols Multipoint LDP)

IN THIS SECTION

- Syntax | 10788
- Hierarchy Level | 10788
- Description | 10788
- Required Privilege Level | 10788
- Release Information | 10788

Syntax

```
mldp-inband-signalling {
    policy policy-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit protocols pim]
```

Description

Multipoint LDP (mLDP) in-band signalling lets you carry multicast traffic across an existing IP/MPLS backbone, while avoiding the use of PIM in the provider core.

On the label-edge router (LER), enable PIM to use mLDP in-band signaling for the upstream neighbors when the LER does not detect a PIM upstream neighbor. On the egress nodes, configure the MPLS LSP root in the PIM configuration, using the `policy` statement.

When used in conjunction with distributed MLD or distributed IGMP, mLDP inband signalling supports interconnecting separate PIM domains via a MPLS-based core. To enable the inter-working, `chassis network-services enhanced-ip` must be enabled and you need to set the `dynamic-profiles profile-name protocols` *`igmp|mld`* `interface` *`interface-name`* to `distributed`. Enabling this command, `mldp-inband-signalling`, has PIM act as a multipoint LDP inband edge router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Support added in Junos OS Release 18.2R1 for using this command in conjunction with distributed MLD or distributed IGMP.

RELATED DOCUMENTATION

No Link Title

mlfr-uni-nni-bundle-options

IN THIS SECTION

- Syntax | 10789
- Hierarchy Level | 10790
- Description | 10790
- Required Privilege Level | 10790
- Release Information | 10790

Syntax

```
mlfr-uni-nni-bundle-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    cisco-interoperability send-lip-remove-link-for-link-reject;
    drop-timeout milliseconds;
    fragment-threshold bytes;
    hello-timer milliseconds;
    lmi-type (ansi | itu | c-lmi);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 number;
    t392 number;
    yellow-differential-delay milliseconds;
}
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel]
```

Description

Configure link services interface management properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Encapsulation for Link Services Physical Interfaces

mlfr-uni-nni-bundles-inline

IN THIS SECTION

- Syntax | 10791
- Hierarchy Level | 10791
- Description | 10791
- Options | 10791
- Required Privilege Level | 10791
- Release Information | 10791

Syntax

```
mlfr-uni-nni-bundles-inline number;
```

Hierarchy Level

```
[edit chassis fpc number pic number]
```

Description

Specify the number of inline multilink frame relay user-to-network interface (*UNI*) and network-to-network interface (*NNI*) bundles.

Options

number Specify the number of inline multilink frame relay UNI NNI bundles.

- **Range:** 1 through 255

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Inline MLPPP for WAN Interfaces Overview

Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces

Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces

mobile-edge session-scale

IN THIS SECTION

- Syntax | 10792
- Hierarchy Level | 10792
- Description | 10792
- Options | 10792
- Required Privilege Level | 10793
- Release Information | 10793

Syntax

```
session-scale {
    1m
    250k
    500k
    750k
}
```

Hierarchy Level

```
[edit services mobile-edge session-scale scale-option]
```

Description

Set the session scaling profile for your mobile-edge configuration.

Options

1m Set the session scaling profile to 1 million sessions.

250k Set the session scaling profile to 250,000 sessions.

500k Set the session scaling profile to 500,000 sessions.

750k Set the session scaling profile to 750,000 sessions.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Session Maintenance and Optimization

mobility

IN THIS SECTION

- Syntax | 10793
- Hierarchy Level | 10794
- Description | 10794
- Options | 10794
- Required Privilege Level | 10794
- Release Information | 10795

Syntax

```
mobility {
    user-plane;
    sgw;
```

```
}
```

Hierarchy Level

```
[edit chassis fpc fpc-slot pfe pfe-id forwarding-packages]
```

Description

Specify the forwarding package that the Packet Forwarding Engines associated with mobility must use.

NOTE:

- You must include every Packet Forwarding Engine configured with the `user-plane` forwarding package at the `[edit unified-edge gateways user-plane gateway-name system anchor-pfes]` hierarchy level on the broadband gateway. If you do not specify the Packet Forwarding Engine as an anchor interface, then the Packet Forwarding Engine will not be used by the broadband gateway.
- You must include every Packet Forwarding Engine configured with the `sgw` forwarding package at the `[edit unified-edge gateways sgw gateway-name system anchor-pfes]` hierarchy level on the broadband gateway. If you do not specify the Packet Forwarding Engine as an anchor interface, then the Packet Forwarding Engine will not be used by the broadband gateway.

Options

user-plane Configure the router as a gateway GPRS support node (GGSN) or as a Packet Data Network Gateway (P-GW) or as an SAEGW-U.

sgw Configure the router as a Serving Gateway (S-GW).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced for Junos Multi-Acess User Plane in Junos OS Release 19.4R1.

`ggsn-pgw` option deprecated and replaced with `user-plane` option in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

forwarding-packages | 7478

mode (RFC 2544 Benchmarking Tests)

IN THIS SECTION

- Syntax | 10795
- Junos OS Hierarchy Level | 10795
- Junos OS Evolved Hierarchy Level | 10796
- Description | 10796
- Options | 10796
- Required Privilege Level | 10796
- Release Information | 10796

Syntax

```
mode (initiate-and-terminate | reflect);
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the test mode for the packets that are sent during the benchmarking test.

Options

initiate-and-terminate (Most ACX routers, except for the ACX5448, ACX5048, and ACX5096) Initiate and terminate the test on the chosen service (IPv4 or Ethernet). The ACX7100, ACX7509, ACX7024, and ACX7024X devices support this mode only for an IPv4 service.

reflect (MX routers and the ACX7100, ACX7509, ACX7024, ACX5448, ACX5048, and ACX5096 routers) Reflect the test frames on the chosen service (IPv4 or Ethernet).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Evolved Release 21.1R1 for reflect mode only.

`initiate-and-terminate` mode introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

RFC 2544-Based Benchmarking Tests Overview

rfc2544-benchmarking | **15759**

rfc2544 | **15757**

mode (ike)

IN THIS SECTION

- Syntax | 10797
- Hierarchy Level | 10797
- Description | 10797
- Default | 10797
- Options | 10797
- Required Privilege Level | 10798
- Release Information | 10798

Syntax

```
mode (aggressive | main);
```

Hierarchy Level

```
[edit services ipsec-vpn ike policy policy-name]
```

Description

Define an IKE policy mode.

Default

main

Options

aggressive—Takes half the number of messages of main mode, has less negotiation power, and does not provide identity protection.

NOTE: In Junos FIPS mode, the `aggressive` option is not supported.

`main`—Uses six messages, in three peer-to-peer exchanges, to establish the IKE SA. These three steps include the IKE SA negotiation, a Diffie-Hellman exchange, and authentication of the peer. Also provides identity protection.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IKE Policies

mode

IN THIS SECTION

- Syntax | 10799
- Hierarchy Level | 10799
- Description | 10799
- Default | 10800
- Options | 10800
- Required Privilege Level | 10800
- Release Information | 10800

Syntax

```
mode (p2p | shared);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp | vstp) interface interface-name],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp | vstp) interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id interface interface-name],
```

```
[edit protocols (mstp | rstp | vstp) interface interface-name],
[edit protocols vstp vlan vlan-id interface interface-name],
[edit protocols mstp interface (all | interface-name) arp-on-stp],
[edit protocols mstpmsti msti-id interface interface-name) arp-on-stp],
[edit protocols rstp interface (all | interface-name) arp-on-stp],
```

```
[edit routing-instances routing-instance-name protocols (mstp | rstp | vstp) interface interface-
name],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id interface interface-name]
[edit protocols stp interface (all | interface-name)],
[edit protocols stp interface (all | interface-name) arp-on-stp],
[edit protocols vstp vlan vlan-id interface (all | interface-name) arp-on-stp]
```

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configures the link mode to identify point-to-point links.

NOTE: For EX4300 switches, the interfaces operate in full-duplex mode only.

Default

When the link is configured as full-duplex, the default link mode is **p2p**. When the link is configured half-duplex, the default link mode is **shared**.

Options

p2p—The link is point to point.

shared—The link is shared media.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Spanning-Tree Instance Interfaces

Understanding STP

Understanding VSTP

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Configuring Faster Convergence and Network Stability on ELS Switches with RSTP

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

mode (Dynamic Profiles)

IN THIS SECTION

- Syntax | 10801
- Hierarchy Level | 10801
- Description | 10801
- Default | 10801
- Required Privilege Level | 10802
- Release Information | 10802

Syntax

```
mode loose;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
(inet) rpf-check]
```

Description

Check whether the packet has a source address with a corresponding prefix in the routing table. If a corresponding prefix is not found, unicast reverse path forwarding (RPF) loose mode does not accept the packet. Unlike strict mode, loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.

Default

If you do not include this statement, unicast RPF is in strict mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Unicast RPF Strict Mode

Unicast RPF in Dynamic Profiles for Subscriber Interfaces

mode (IKE)

IN THIS SECTION

- Syntax | 10802
- Hierarchy Level | 10803
- Description | 10803
- Default | 10803
- Options | 10803
- Required Privilege Level | 10803
- Release Information | 10803

Syntax

```
mode (aggressive | main);
```

Hierarchy Level

```
[edit security ike policy ike-peer-address]
```

Description

Define the IKE policy mode.

NOTE: IKEv2 protocol does not negotiate using mode configuration.

Default

main

Options

`aggressive`—Take half the number of messages of main mode, has less negotiation power, and does not provide identity protection.

`main`—Use six messages, in three peer-to-peer exchanges, to establish the IKE SA. These three steps include the IKE SA negotiation, a Diffie-Hellman exchange, and authentication of the peer. Also provides identity protection.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an IKE Policy for Preshared Keys

mode (Interfaces)

IN THIS SECTION

- Syntax | 10804
- Hierarchy Level | 10804
- Description | 10804
- Default | 10805
- Required Privilege Level | 10805
- Release Information | 10805

Syntax

```
mode (strict | loose);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family (inet | inet6) rpf-check],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family (inet | inet6) rpf-check]
```

Description

Check whether the packet has a source address with a corresponding prefix in the routing table. If a corresponding prefix is not found, unicast reverse path forwarding (RPF) loose mode does not accept the packet. Unlike strict mode, loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.

In strict mode, unicast RPF checks whether the incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix. If the incoming packet fails the unicast RPF check, the packet is not accepted on the interface.

Default

If you do not include this statement, unicast RPF is in strict mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Unicast RPF Strict Mode

mode (IPsec)

IN THIS SECTION

- Syntax | 10805
- Hierarchy Level | 10806
- Description | 10806
- Default | 10806
- Options | 10806
- Required Privilege Level | 10806
- Release Information | 10806

Syntax

```
mode (transport | tunnel);
```

Hierarchy Level

```
[edit security ipsec security-association name]
```

Description

Define the mode for the IPsec security association.

Default

tunnel

Options

transport—Protect traffic when the communication endpoint and cryptographic endpoint are the same. The data portion of the IP packet is encrypted, but the IP header is not. Virtual Private Network (VPN) gateways that provide encryption and decryption services for protected hosts cannot use transport mode for protected VPN communications.

tunnel—Protect traffic using preshared keys with IKE to authenticate peers or digital certificates with IKE to authenticate peers.

> **NOTE**: The Junos OS supports only encapsulating security payload (ESP) when you use tunnel mode.
>
> In transport mode, the Junos OS does not support authentication header (AH) and ESP header bundles.
>
> In transport mode, the Junos OS supports only Border Gateway Protocol (BGP).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Using IPsec to Protect BGP Traffic

Configuring IPsec Tunnel Mode

mode (Traffic Manager)

IN THIS SECTION

- Syntax | 10807
- Hierarchy Level | 10807
- Description | 10807
- Options | 10808
- Required Privilege Level | 10808
- Release Information | 10808

Syntax

```
mode traffic-manager-mode;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number traffic-manager],
[edit chassis lcc number fpc slot-number pic pic-number traffic-manager]
[edit chassis fpc slot-number pic pic-number port port-number traffic-manager], (MX80 and MX104
routers only)
```

Description

Enable shaping on an L2TP session.

Options

traffic-manager-mode—Configure CoS traffic manager mode of operation on this interface. This option has the following suboptions:

egress-only Enable CoS queuing and scheduling on the egress side for the PIC that houses the interface. This is the default mode.

If ingress packet drops are observed at a high rate for a PIC, configure the `traffic-manager` statement to work in the `egress-only` mode.

ingress-and-egress Enable CoS queuing and scheduling on both the egress and ingress sides for the PIC.

Junos does not support `ingress-and-egress` mode on label-switched interfaces (LSIs) configured with VPLS.

For EQ DPCs, you must configure the `traffic-manager` statement with `ingress-and-egress` mode to enable ingress CoS on the EQ DPC. EQ DPCs have 250 ms of buffering, with only egress queuing (default mode). When `ingress-and-egress` is configured, the buffer is partitioned as 50 ms for the ingress direction and 200 ms for the egress direction.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring Ingress Hierarchical CoS on MIC and MPC Interfaces

egress-shaping-overhead | **6223**

ingress-shaping-overhead | **8649**

traffic-manager | **30036**

mode (Multicast VLAN Registration)

IN THIS SECTION

- Syntax | 10809
- Hierarchy Level | 10809
- Description | 10809
- Default | 10810
- Options | 10810
- Required Privilege Level | 10811
- Release Information | 10811

Syntax

```
mode (proxy | transparent);
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan vlan-name data-forwarding receiver]
```

Description

Configure the operating mode for a Multicast VLAN Registration (MVR) receiver VLAN.

A multicast VLAN (MVLAN) forwards multicast streams to interfaces on other VLANs that are configured as MVR receiver VLANs for that MVLAN, and can operate in either of two modes, transparent or proxy. The mode setting affects how IGMP reports are sent to the upstream multicast router. In transparent mode, the device sends IGMP reports out of the MVR receiver VLAN, and in proxy mode, the device sends IGMP reports out of the MVLAN.

We recommend that you configure proxy mode on devices that are closest to the upstream multicast router, because in transparent mode, IGMP reports are only sent out on the MVR receiver VLAN. As a result, MVR receiver ports receiving an IGMP query from an upstream router on the MVLAN will only reply on MVR receiver VLAN multicast router ports, the upstream router will not receive the replies, and

the upstream router will not continue to forward traffic. In proxy mode, IGMP reports are sent out on the MVLAN for its MVR receiver VLANs, so the upstream multicast router receives IGMP replies on the MVLAN and continues to forward the multicast traffic on the MVLAN.

In either mode, the device forms multicast group memberships on the MVLAN, and IGMP queries and forwards multicast traffic received on the MVLAN to subscribers in MVR receiver VLANs tagged with the MVLAN tag by default. If you also configure the translate option at the `[edit protocols igmp-snooping vlans vlan-name data-forwarding receiver]` hierarchy level for hosts on trunk ports in MVR receiver VLANs, then upon egress, the device translates MVLAN tags into the MVR receiver VLAN tags instead.

> **NOTE**: This statement is available to configure the MVR mode only on devices that support the Enhanced Layer 2 Software (ELS) configuration style. Devices with software that does not support ELS operate in transparent mode by default, or operate in proxy mode if you configure the `proxy` statement at the `[edit protocols igmp-snooping vlan vlan-name]` hierarchy level for a VLAN configured as a data-forwarding VLAN.

Default

Transparent mode

Options

transparent MVR operates in transparent mode if this option is configured (and is also the default if no mode is configured). In transparent mode, IGMP reports are sent out from the device in the context of the MVR receiver VLAN. IGMP join and leave messages received on MVR receiver VLAN interfaces are forwarded to the multicast router ports on the MVR receiver VLAN. IGMP queries received on the MVR receiver VLAN are forwarded to all MVR receiver ports. IGMP queries received on the MVLAN are forwarded to the MVR receiver ports that are in the receiver VLANs belonging to the MVLAN, even though those ports might not be on the MVLAN itself.

When a host on an MVR receiver VLAN joins a multicast group, the device installs a bridging entry on the MVLAN and forwards MVLAN traffic for that group to the host, even though the host is not in the MVLAN. You can also configure the device to install the bridging entries on the MVR receiver VLAN (see the `install` option at the `[edit protocols igmp-snooping vlans vlan-name data-forwarding receiver]` hierarchy level).

proxy When you configure proxy mode for an MVR receiver VLAN, the device acts as a proxy to the IGMP multicast router for MVR group membership requests received on MVR receiver VLANs. The device forwards IGMP reports from hosts on MVR receiver VLANs in the context of the MVLAN and forwards them to the multicast router ports on the MVLAN

only, so the multicast router receives IGMP reports only on the MVLAN for those MVR receiver hosts. IGMP queries are handled in the same way as in transparent mode; IGMP queries received on either the MVR receiver VLAN or the MVLAN are forwarded to all MVR receiver ports in receiver VLANs belonging to the MVLAN (even though those ports are not on the MVLAN itself).

When a host on an MVR receiver VLAN joins a multicast group, the device installs a bridging entry on the MVLAN, and subsequently forwards MVLAN traffic for that group to the host although the host is not in the MVLAN. You cannot configure the `install` option to install the bridging entries on the MVR receiver VLAN for a data-forwarding MVR receiver VLAN that is configured in proxy mode.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 18.4R1 for EX2300 and EX3400 switches.

RELATED DOCUMENTATION

Understanding Multicast VLAN Registration

Configuring Multicast VLAN Registration on EX Series Switches

mode (Next Gen Services Service-Set System Logging)

IN THIS SECTION

- Syntax | 10812
- Hierarchy Level | 10812
- Description | 10812
- Options | 10812

Required Privilege Level | 10812
Release Information | 10812

Syntax

```
mode {
    event ;
    stream stream-name;
}
```

Hierarchy Level

```
[edit services services-set name syslog]
```

Description

Mode in which the system message logger sends messages

Options

event Send messages to a file on the local routing engine

stream Send messages to one or more remote log servers. Each remote server requires its own stream.

Required Privilege Level

system

Release Information

Support introduced in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

mode (Protocols DVMRP)

IN THIS SECTION

- Syntax | 10813
- Hierarchy Level | 10813
- Description | 10813
- Options | 10814
- Required Privilege Level | 10814
- Release Information | 10814

Syntax

```
mode (forwarding | unicast-routing);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols dvmrp interface interface-name],
[edit protocols dvmrp interface interface-name]
```

Description

Configure DVMRP for multicast traffic forwarding or unicast routing.

Options

forwarding—DVMRP performs unicast routing as well as multicast data forwarding.

unicast-routing—DVMRP performs unicast routing only. To forward multicast data, you must configure Protocol Independent Multicast (PIM) on the interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

> **NOTE**: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring DVMRP to Announce Unicast Routes

mode (Protocols MSDP)

IN THIS SECTION

- Syntax | 10815
- Hierarchy Level | 10815
- Description | 10815
- Default | 10815
- Options | 10815

Required Privilege Level | 10815
Release Information | 10816

Syntax

```
mode (mesh-group | standard);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name],
[edit protocols msdp group group-name],
[edit routing-instances routing-instance-name protocols msdp group group-name]
```

Description

Configure groups of peers in a full mesh topology to limit excessive flooding of source-active messages to neighboring peers. The default flooding mode is **standard**.

Default

If you do not include this statement, default flooding is applied.

Options

mesh-group—Group of peers that are mesh group members.

standard—Use standard MSDP source-active flooding rules.

- **Default: standard**

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP with Active Source Limits and Mesh Groups

mode (Protocols PIM)

IN THIS SECTION

- Syntax | 10816
- Hierarchy Level | 10816
- Description | 10817
- Options | 10817
- Required Privilege Level | 10817
- Release Information | 10817

Syntax

```
mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse | sparse-dense);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name],
[edit protocols pim interface (Protocols PIM)  interface-name],
```

```
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name]
```

Description

Configure the PIM mode on the interface.

Options

The choice of PIM mode is closely tied to controlling how groups are mapped to PIM modes, as follows:

- **bidirectional-sparse**—Use if all multicast groups are operating in bidirectional, sparse, or SSM mode.
- **bidirectional-sparse-dense**—Use if multicast groups, except those that are specified in the `dense-groups` statement, are operating in bidirectional, sparse, or SSM mode.
- **dense**—Use if all multicast groups are operating in dense mode.
- **sparse**—Use if all multicast groups are operating in sparse mode or SSM mode.
- **sparse-dense**—Use if multicast groups, except those that are specified in the `dense-groups` statement, are operating in sparse mode or SSM mode.
- **Default:** Sparse mode

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

bidirectional-sparse and **bidirectional-sparse-dense** options introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring PIM Dense Mode Properties

Configuring PIM Sparse-Dense Mode Properties

Example: Configuring Bidirectional PIM

mode (QFX Series)

IN THIS SECTION

- Syntax | 10818
- Hierarchy Level | 10818
- Description | 10818
- Options | 10818
- Required Privilege Level | 10818
- Release Information | 10819

Syntax

```
mode active-active ;
```

Hierarchy Level

```
[edit interfaces aggregated-ether-options mc-ae]
```

Description

Configure the multichassis link aggregation group (MC-LAG) to be in active-active mode. In active-active mode, all of the members of the MC-LAG will be active on both routing or switching devices. Only active-active mode is supported at this time.

Options

active-active Specify that all of the members of the MC-LAG will be active on both routing or switching devices.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

mode (Security Forwarding Options)

IN THIS SECTION

- Syntax | 10819
- Hierarchy Level | 10819
- Description | 10819
- Options | 10819
- Required Privilege Level | 10820
- Release Information | 10821

Syntax

```
mode (drop |  flow-based | packet-based);
```

Hierarchy Level

```
[edit security forwarding-options family inet6]
```

Description

Specify forwarding options for IPv6 traffic.

Options

- `drop`—Drop IPv6 packets. This is the default setting.

- `flow-based`—Perform flow-based packet forwarding.
- `packet-based`—Perform simple packet forwarding.

> **NOTE**: Packet-based processing is not supported on the following SRX Series Firewalls: SRX5400, SRX5600, and SRX5800.

- Starting with Junos OS Release 15.1X49-D70, on SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, SRX5800 devices, the default mode is changed from drop mode to flow based mode.
- If you change the forwarding option mode for IPv6, you might have to perform a reboot to initialize the configuration change. Table 67 on page 10820 summarizes device status upon configuration change.
- Starting with Junos OS Release 15.1X49-D70, on SRX5400, SRX5600, and SRX5800 devices, reboot is not required when you change the modes between flow-based mode and drop mode.

Table 67: Device Status Upon Configuration Change

Configuration Change	Commit Warning	Reboot Required	Impact on Existing Traffic Before Reboot	Impact on New Traffic Before Reboot
Drop to flow-based	Yes	Yes	Dropped	Dropped
Drop to packet-based	No	No	Packet-based	Packet-based
Flow-based to packet-based	Yes	Yes	None	Flow sessions created
Flow-based to drop	Yes	Yes	None	Flow sessions created
Packet-based to flow-based	Yes	Yes	Packet-based	Packet-based
Packet-based to drop	No	No	Dropped	Dropped

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Support on SRX Series Firewalls for flow-based mode for family inet6 added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

mode (Security Group VPN)

IN THIS SECTION

- Syntax | 10821
- Hierarchy Level | 10821
- Description | 10822
- Options | 10822
- Required Privilege Level | 10822
- Release Information | 10822

Syntax

```
mode (aggressive | main);
```

Hierarchy Level

```
[edit security group-vpn member ike policy policy-name]
[edit security group-vpn server ike policy policy-name]
```

Description

Define the mode used for Internet Key Exchange (IKE) Phase 1 negotiations. Use aggressive mode only when you need to initiate an IKE key exchange without ID protection, as when a peer unit has a dynamically assigned IP address. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 devices and vSRX Virtual Firewall instances.

- IKEv2 protocol does not negotiate using mode configuration.
- The device deletes existing IKE and IPsec SAs when you update the `mode` configuration in the IKE policy.

Options

- `aggressive`—Aggressive mode.
- `main`—Main mode. Main mode is the recommended key-exchange method because it conceals the identities of the parties during the key exchange.

  Configuring `mode main` for group VPN servers or members is not supported when the remote gateway has a dynamic address and the authentication method is `pre-shared-keys`.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for `group-vpn` hierarchies added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

policy (Security Group VPN IKE) | **12711**

mode (Security Log)

IN THIS SECTION

- Syntax | 10823
- Hierarchy Level | 10823
- Description | 10823
- Options | 10823
- Required Privilege Level | 10824
- Release Information | 10824

Syntax

```
mode (event | stream | stream-event)
```

Hierarchy Level

```
[edit security log]
[edit logical-systems name security log]
[edit tenants tenant-name security log]
```

Description

Set the mode of logging (event for traditional system logging or stream for streaming security logs through a revenue port to a server).

Options

- `event`— Process security logs in the control plane.
- `stream`— Process security logs directly in the forwarding plane.
- `stream-event`— Process security logs in the control and forwarding plane.

Default:

- event is a default mode on SRX100, SRX110, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, and SRX650, devices.

  Apart from Junos OS Release 18.4R1, 18,4R2, 19.1 and 19.2R1, in all the other releases starting in Junos OS Release 18.3R3, the default logging mode for SRX300, SRX320, SRX340, SRX345, SRX550, and SRX550M devices is stream mode.

- stream is a default mode on SRX1400, SRX1500, SRX3400, SRX3600, SRX5400, SRX5600, SRX5800, SRX4100, and SRX4200 devices.

  When a device is configured in stream mode, logs are sent from the PFE to an external syslog server through a revenue port other than fxp0 interface. You cannot use fxp0 interface for stream mode irrespective of whether the fxp0 interface is part of the default routing instance or mgmt_junos routing instance.

  Starting with Junos OS Release 15.1X49-D140, the default mode for SRX1500 device is stream mode. Prior to Junos OS Release 15.1X49-D140, the default mode for SRX1500 device was event mode.

- event is a default mode on vSRX Virtual Firewall and vSRX3.0 devices.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

The [edit `logical-systems` *name* security log] hierarchy level introduced in Junos OS Release 18.2R1.

The [edit `tenants` *tenant-name* security log] hierarchy level introduced in Junos OS Release 18.3R1.

`stream-event` option is introduced in Junos OS Release 20.2R1 for SRX Series Firewalls.

RELATED DOCUMENTATION

syslog (System) | **28836**

modem-options

IN THIS SECTION

- Syntax | 10825
- Hierarchy Level | 10825
- Description | 10825
- Required Privilege Level | 10825
- Release Information | 10826

Syntax

```
modem-options {
    dialin (console | routable);
    init-command-string initialization-command-string;
}
```

Hierarchy Level

```
[edit interfaces umd0]
```

Description

For J Series Services Routers, configure a USB port to act as a USB modem.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

modulation-format

IN THIS SECTION

- Syntax | 10826
- Hierarchy Level | 10826
- Description | 10826
- Options | 10826
- Required Privilege Level | 10827
- Release Information | 10827

Syntax

```
modulation-format (qpsk | 8qam | 16qam);
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Description

Specify the modulation format.

Options

`qpsk`—Quadrature Phase Shift Keying modulation format.

`8qam`—8 quadrature amplitude modulation format.

`16qam`—16 quadrature amplitude modulation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Configuring Link Down Notification for Optics Options Alarm or Warning

100-Gigabit Ethernet OTN Options Configuration Overview

mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain)

IN THIS SECTION

- Syntax | 10827
- Hierarchy Level | 10828
- Description | 10828
- Required Privilege Level | 10828
- Release Information | 10828

Syntax

```
mofrr-asm-starg;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast stream-protection],
[edit logical-systems logical-system-name routing-options multicast stream-protection],
[edit routing-instances routing-instance-name routing-options multicast stream-protection],
[edit routing-options multicast stream-protection]
```

Description

Enable `mofrr-asm-starg` to include any-source multicast (ASM) for (*,G) joins in the Multicast-only fast reroute (MoFRR).

> **NOTE**: `mofrr-asm-starg` applies to IP-PIM only. When enabled for group G, *,G will undergo MoFRR as long as there is no S#,G for Group G. In other words, *,G MoFRR will cease and any old states will be torn down when S#,G is created. Note too, that `mofrr-asm-starg` is not supported for mLDP (since mLDP itself does not support *,G).

In a PIM domain with MoFRR enabled, the default for `stream-protection` is S,G routes only.

Context: Multicast-only fast reroute (MoFRR) can be used to reduce traffic loss in a multicast distribution tree in the event of link down. To employ MoFRR, a downstream router is configured with an alternative path back towards the source, over which it receives a backup live stream of the same multicast traffic. That router propagates the same (S,G) join toward both upstream neighbors in order to create duplicate multicast trees. If a failure is detected on the primary tree, the router switches to the backup tree to prevent packet loss.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Multicast-Only Fast Reroute

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain

mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain)

IN THIS SECTION

- Syntax | 10829
- Hierarchy Level | 10829
- Description | 10830
- Required Privilege Level | 10830
- Release Information | 10830

Syntax

```
mofrr-disjoint-upstream-only;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast stream-protection],
[edit logical-systems logical-system-name routing-options multicast stream-protection],
[edit routing-instances routing-instance-name routing-options multicast stream-protection],
[edit routing-options multicast stream-protection]
```

Description

When you configure multicast-only fast reroute (MoFRR) in a PIM domain, allow only a disjoint RPF (an RPF on a separate plane) to be selected as the backup RPF path.

In a multipoint LDP MoFRR domain, the same label is shared between parallel links to the same upstream neighbor. This is not the case in a PIM domain, where each link forms a neighbor. The `mofrr-disjoint-upstream-only` statement does not allow a backup RPF path to be selected if the path goes to the same upstream neighbor as that of the primary RPF path. This ensures that MoFRR is triggered only on a topology that has multiple RPF upstream neighbors.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Multicast-Only Fast Reroute

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain

mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain)

IN THIS SECTION

- Syntax | 10831
- Hierarchy Level | 10831
- Description | 10831
- Required Privilege Level | 10831

Release Information | 10831

Syntax

```
mofrr-no-backup-join;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast stream-protection],
[edit logical-systems logical-system-name routing-options multicast stream-protection],
[edit routing-instances routing-instance-name routing-options multicast stream-protection],
[edit routing-options multicast stream-protection]
```

Description

When you configure multicast-only fast reroute (MoFRR) in a PIM domain, prevent sending join messages on the backup path, but retain all other MoFRR functionality.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Multicast-Only Fast Reroute

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain

mofrr-primary-path-selection-by-routing (Multicast-Only Fast Reroute)

IN THIS SECTION

- Syntax | 10832
- Hierarchy Level | 10832
- Description | 10832
- Default | 10833
- Required Privilege Level | 10833
- Release Information | 10833

Syntax

```
mofrr-primary-path-selection-by-routing;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast stream-protection],
[edit logical-systems logical-system-name routing-options multicast stream-protection],
[edit routing-instances routing-instance-name routing-options multicast stream-protection],
[edit routing-options multicast stream-protection]
```

Description

MoFRR is supported on both equal-cost multipath (ECMP) paths and non-ECMP paths. Unicast loop-free alternate (LFA) routes need to be enabled to support MoFRR on non-ECMP paths. LFA routes are enabled with the link-protection statement in the interior gateway protocol (IGP) configuration. When you enable link protection on an OSPF or IS-IS interface, Junos OS creates a backup LFA path to the primary next hop for all destination routes that traverse the protected interface.

In the context of load balancing, MoFRR prioritizes the disjoint backup in favor of load balancing the available paths.

For Junos OS releases before 15.1R7, for both ECMP and Non-ECMP scenarios, the default MoFRR behavior was *sticky* , that is, if the Active link went down, the Active Path selection would give preference to Backup Path for the transition. The Active Path would not follow the unicast selected gateway

Starting in Junos OS Release 15.1R7 however, the default behavior for non-EMCP scenarios is to be *non*sticky, that is, the selection of Active Path strictly follows unicast selected gateway. MoFRR no longer chooses a unicast LFA path to become the MoFRR Active path; only a unicast LFA path can be selected to become MoFRR Backup.

Default

By default, the backup path gets promoted to be the primary path when MoFRR is configured in a PIM domain.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Multicast-Only Fast Reroute

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain

monitor-bandwidth

IN THIS SECTION

- Syntax | 10834
- Hierarchy Level | 10834
- Description | 10834
- Required Privilege Level | 10834
- Release Information | 10834

Syntax

```
monitor-bandwidth;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Do not automatically adjust bandwidth allocation. However, the maximum average bandwidth utilization is monitored on the LSP, and the information is recorded in the MPLS statistics file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Passive Bandwidth Utilization Monitoring

monitor-end-point

IN THIS SECTION

- Syntax | 10835
- Hierarchy Level | 10835
- Description | 10835
- Default | 10835
- Options | 10836
- Required Privilege Level | 10836
- Release Information | 10836

Syntax

```
(monitor-end-point | no-monitor-end-point);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

Description

Originate or do not originate the connection monitor end point.

Default

By default, do not originate the connection monitor end point.

Options

monitor-end-point	Originate the connection monitor end point.
no-monitor-end-point	Do not originate the connection monitor end point.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

monitor ethernet delay-measurement

IN THIS SECTION

- Syntax | **10837**
- Description | **10837**
- Options | **10838**
- Additional Information | **10839**
- Required Privilege Level | **10839**
- Output Fields | **10840**
- Sample Output | **10842**
- Release Information | **10843**

Syntax

```
monitor ethernet delay-measurement
maintenance-domain md-name
maintenance-association ma-name
(one-way | two-way)
(remote-mac-address | mep remote-mep-id)
<count frame-count>
<local-mep mep-id>
<wait interval-seconds>
<priority 802.1p value>
<size>
<no-session-id-tlv>
<xml>
```

Description

Start an ITU-T Y.1731 Ethernet frame delay measurement session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a summary of the frames exchanged in the measurement session. Frame delay measurement statistics are stored at one of the MEPs for later retrieval.

> **NOTE**: If you attempt to monitor delays to a nonexistent MAC address, you must type Ctrl +c to explicitly quit the `monitor ethernet delay-measurement` command and return to the CLI command prompt.

To start an Ethernet frame delay measurement session, the router initiates an exchange of frames carrying one-way or two-way frame delay measurement protocol data units (PDUs) between the local and remote MEPs. The frame counts—the types of and number of Ethernet frame delay measurement PDU frames exchanged to measure frame delay times—are displayed as the runtime output of the `monitor ethernet delay-measurement` command and are also stored at both the initiator and receiver MEPs for later retrieval. Ethernet frame delay measurement statistics, described below, are measured and stored at only one of the MEPs:

Frame delay — The difference, in microseconds, between the time a frame is sent and when it is received.

Frame delay variation — The difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.”

For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet frame delay measurement, only the initiator MEP (on the local system) collects statistics.

Options

maintenance-domain *md-name*	Name of an existing CFM maintenance domain.
maintenance-association *ma-name*	Name of an existing CFM maintenance association.
one-way	Measurement type is one-way Ethernet frame delay measurement, which is based on the difference between the time at which the initiator MEP sends a one-way delay measurement request (1DM) frame and the time at which the receiver MEP receives the frame.
two-way	Measurement type is two-way Ethernet frame delay measurement, which is based on the difference between the time at which the initiator MEP sends a two-way delay measurement message (DMM) frame and the time at which the initiator MEP receives an associated two-way delay measurement reply (DMR) frame from the responder MEP, subtracting the time elapsed at the responder MEP.
mep *remote-mep-id*	Numeric identifier of the peer MEP with which to perform Ethernet frame delay measurement. The discovered MAC address of the peer MEP is used. The range of values is 1 through 8191.
remote-mac-address	Unicast MAC address of the peer MEP with which to perform Ethernet frame delay measurement. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn***. For example, **0011.2233.4455** or **00:11:22:33:44:55**. Multicast MAC addresses are not supported.
count *frame-count*	(Optional) Number of frames to send to the specified peer MEP. The range of values is 1 through 65,535 frames. The default value is 10 frames.
local-mep *mep-id*	(Required when multiple MEPs are configured) Identifier for the local maintenance association end point.
wait *interval-seconds*	(Optional) Number of seconds to wait between sending frames. The range of values is from 1 through 255 seconds. The default value is 1 second.

priority ***802.1p value***	(Optional) Priority of the delay measurement request frame supported by both one-way delay measurement and two-way delay measurement. The range of values is from 0 through 7. The default value is zero.
size	(Optional) Size of the data TLV to be included in the request frame. The range of values is from 1 through 1400 bytes.
no-session-id-tlv	(Optional) Prevent insertion of the session ID TLV in the request frame.
xml	(Optional) Allow the output of the command to be displayed in XML format supported by both one-way delay measurement and two-way delay measurement. Note that the only way to get output in XML format is to use the **xml** argument. The `display xml` command does not work.

Additional Information

To display the frame counts collected at an MEP as the result of this command, see the following command descriptions in the CLI Explorer:

- `show oam ethernet connectivity-fault-management interfaces detail`
- `show oam ethernet connectivity-fault-management mep-database`
- `show oam ethernet connectivity-fault-management mep-statistics`

To display the statistics collected at an MEP as the result of this command, see the following command descriptions in the CLI Explorer.

- `show oam ethernet connectivity-fault-management delay-statistics`
- `show oam ethernet connectivity-fault-management mep-statistics`

To clear both the frame counts and the statistics collected for MEPs, use the `clear oam ethernet connectivity-fault-management statistics` command, described in the CLI Explorer.

For a complete description of Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the Junos OS Network Interfaces Library for Routing Devices.

Required Privilege Level

trace and maintenance

Output Fields

The `monitor ethernet delay-measurement` command displays different output at the CLI, depending on whether you start a one-way or two-way frame delay measurement:

- Table 68 on page 10840 lists the run-time output fields for the `monitor ethernet delay-measurement one-way` command.
- Table 69 on page 10841 lists the run-time output fields for the `monitor ethernet delay-measurement two-way` command.

Output fields are listed in the approximate order in which they appear.

Table 68: monitor ethernet delay-measurement one-way Output Fields

Output Field Name	Output Field Description
One-way ETH-DM request to	Unicast MAC address of the remote peer MEP.
Interface	Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.
1DM Frames sent	PDU frames sent to the remote MEP in this ETH-DM session.
Packets transmitted	Total number of 1DM PDU frames sent to the remote MEP during this measurement session.
Average delay	Average two-way frame delay measured in this session.
Average delay variation	Average frame jitter measured in this session.
Best case delay	Lowest two-way frame delay measured in this session.
Worst case delay	Highest two-way frame delay measured in this session.

NOTE: For one-way delay measurement, these CLI output fields display **NA** (“not applicable”) at the initiator MEP because one-way frame delay measurements occur at the receiver MEP.

Table 69: monitor ethernet delay-measurement two-way Output Fields

Output Field Name	Output Field Description
Two-way Ethernet frame delay measurement request to	Unicast MAC address of the remote peer MEP.
Interface	Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.
DMR received from	Unicast MAC address of the remote MEP that transmitted this DMR frame in response to a DMM frame.
Delay	Two-way delay, in microseconds, for the initiator-transmitted DMM frame.
Delay variation	Difference, in microseconds, between the current and previous delay values. This is also known as frame jitter.
Packets transmitted	Total number of DMM PDU frames sent to the remote MEP in this measurement session.
Valid packets received	Total number of DMR PDU frames received from the remote MEP in this measurement session.
Average delay	Average two-way frame delay measured in this session.
Average delay variation	Average frame jitter measured in this session.
Best case delay	Lowest two-way frame delay measured in this session.
Worst case delay	Highest two-way frame delay measured in this session.

Sample Output

monitor ethernet delay-measurement one-way

```
user@host> monitor ethernet delay-measurement one-way 00:05:85:73:39:4a maintenance-domain md6
maintenance-association ma6 count 10
One-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA
```

monitor ethernet delay-measurement two-way

```
user@host> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a maintenance-domain md6
maintenance-association ma6 count 10
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a  Delay: 100 usec  Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a  Delay: 111 usec  Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a  Delay: 110 usec  Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a  Delay: 119 usec  Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a  Delay: 122 usec  Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a  Delay: 108 usec  Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 10
Average delay: 103 usec, Average delay variation: 8 usec
Best case delay: 92 usec, Worst case delay: 122 usec
```

monitor ethernet delay-measurement two-way (Invalid DMR Frames Received)

```
user@host> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a maintenance-domain md6
maintenance-association ma6 count 10
```

```
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a  Delay: 100 usec  Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a  Delay: 111 usec  Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a  Delay: 110 usec  Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a  Delay: 119 usec  Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a  Delay: 122 usec  Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a with invalid timestamp(s).
DMR received from 00:05:85:73:39:4a  Delay: 108 usec  Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 9, Invalid packets received: 1
Average delay: 105 usec, Average delay variation: 9 usec
Best case delay: 92 usec, Worst case delay: 122 usec
```

Release Information

Command introduced in Junos OS Release 9.5.

`local-mep` option introduced in Junos OS Release 15.1.

monitor ethernet loss-measurement

IN THIS SECTION

- Syntax | 10844
- Description | 10844
- Options | 10844
- Additional Information | 10845
- Required Privilege Level | 10846
- Output Fields | 10846
- Sample Output | 10848
- Release Information | 10850

Syntax

```
monitor ethernet loss-measurement
maintenance-domain md-name
maintenance-association ma-name
(remote-mac-address | mep remote-mep-id)
<count frame-count>
<local-mep mep-id>
<wait interval-seconds>
<priority 802.1p value>
<no-session-id-tlv>
<xml>
```

Description

Start an ITU-T Y.1731 Ethernet frame loss measurement session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a count of transmitted and received data frames between the pair of MEPs. Frame loss measurement statistics are stored at one of the MEPs for later retrieval. For MX Series routers, supports point-to-point down MEPs for Ethernet interfaces (as per IEEE 802.1ag over VPWS).

> **NOTE**: If you attempt to monitor loss to a nonexistent MAC address, you must type Ctrl + c to explicitly quit the `monitor ethernet loss-measurement` command and return to the CLI command prompt.

To start an Ethernet frame loss measurement session, the router first sends frames with ETH-LM information to a peer MEP and similarly receives frames with ETH-LM information from the peer MEP. Frame loss is calculated by collecting the counter values applicable for ingress and egress service frames where the counters maintain a count of transmitted and received data frames between a pair of MEPs. The loss measurement statistics are retrieved as the output of the `monitor ethernet loss-measurement` command and are also stored at the initiator. The frames counts are stored at both the initiator and the receiver MEPs for later retrieval.

Options

maintenance-domain *md-name*	Name of an existing CFM maintenance domain.
maintenance-association *ma-name*	Name of an existing CFM maintenance association.

mep *remote-mep-id*	Numeric identifier of the peer MEP with which to perform Ethernet frame loss measurement. The discovered MAC address of the peer MEP is used. The range of values is from 1 through 8192.
remote-mac-address	Unicast MAC address of the peer MEP with which to perform Ethernet frame loss measurement. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn* (for example, **0011.2233.4455** or **00:11:22:33:44:55**). Multicast MAC addresses are not supported.
count *frame-count*	(Optional) Number of frames to send to the specified peer MEP. The range of values is from 1 through 65535 frames. The default value is 10 frames.
local-mep *mep-id*	(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.
wait *interval-seconds*	(Optional) Number of seconds to wait between sending frames. The range of values is from 1 through 255 seconds. The default value is 1 second.
priority *802.1p value*	(Optional) Priority of the delay measurement request frame. The range of values is from 0 through 7. The default value is 1 second.
no-session-id-tlv	(Optional) Disable the **session id TLV** argument set in the request frame.
xml	(Optional) Allow the output of the command to be displayed in XML format.

Additional Information

To display the interator output for an LM session, run the following command:

- `show oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator <profile> maintenance-association <MA> maintenance-domain <MD> local-mep <MEP> remote-mep <RMEP>`

To display the frame counts collected at an MEP as the result of this command, see the following command descriptions in the CLI Explorer:

- `show oam ethernet connectivity-fault-management loss-statistics`
- `show oam ethernet connectivity-fault-management interfaces detail`
- `show oam ethernet connectivity-fault-management mep-database`
- `show oam ethernet connectivity-fault-management mep-statistics`

To display the statistics collected at an MEP as the result of this command, see the following command descriptions in the CLI Explorer:

- `show oam ethernet connectivity-fault-management delay-statistics`
- `show oam ethernet connectivity-fault-management mep-statistics`

To clear both the frame counts and the statistics collected for MEPs, use the `clear oam ethernet connectivity-fault-management loss-statistics maintenance-domain` *`md-name`* `maintenance-association` *`ma-name`* command, as described in the CLI Explorer.

For a complete description of Ethernet frame loss measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the Junos OS Network Interfaces Library for Routing Devices.

Required Privilege Level

trace and maintenance

Output Fields

Table 70 on page 10846 lists the output fields for the `monitor ethernet loss-measurement` command and their descriptions. Output fields are listed in the approximate order in which they appear.

Table 70: monitor ethernet loss-measurement output fields

Output Field Name	Output Field Description
Ethernet loss delay measurement request to	Unicast MAC address of the remote peer MEP.
Interface	Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.
LMR received from	Unicast MAC address of the remote MEP that transmitted this LMR frame in response to a loss measurement message (LMM) frame.
Near-end frame loss	Count of frame loss associated with ingress data frames.
Far-end frame loss	Count of frame loss associated with egress data frames.

Table 70: monitor ethernet loss-measurement output fields *(Continued)*

Output Field Name	Output Field Description
Near-end loss ratio	Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the ingress interface.
Far-end loss ratio	Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the egress interface.
LMM packets transmitted	Total number of LMM PDU frames sent to the remote MEP in this measurement session.
LMR packets received	Total number of LMR PDU frames received from the remote MEP in this measurement session.
Average near-end frame loss	Average frame loss measured in this session associated with ingress data frames.
Average near-end loss ratio	Average frame loss ratio measured in this session associated with ingress data frames.
Average far-end frame loss	Average frame loss measured in this session associated with egress data frames.
Average far-end loss ratio	Average frame loss ratio measured in this session associated with egress data frames.
Near-end best case frame loss	Lowest frame loss measured in this session associated with ingress data frames.
Near-end best case loss ratio	Lowest frame loss ratio measured in this session associated with ingress data frames.
Near-end worst case frame loss	Highest frame loss measured in this session associated with ingress data frames.

Table 70: monitor ethernet loss-measurement output fields *(Continued)*

Output Field Name	Output Field Description
Near-end worst case loss ratio	Highest frame loss ratio measured in this session associated with ingress data frames.
Far-end best case frame loss	Lowest frame loss measured in this session associated with egress data frames.
Far-end best case loss ratio	Lowest frame loss ratio measured in this session associated with egress data frames.
Far-end worst case frame loss	Highest frame loss measured in this session associated with egress data frames.
Far-end worst case loss ratio	Highest frame loss ratio measured in this session associated with egress data frames.

Note that in the preceding table, the term *number of service frames not delivered* is the difference between the number of service frames arriving at the ingress Ethernet flow point and the number of service frames delivered at the egress Ethernet flow point in a point-to-point Ethernet connection.

Sample Output

monitor ethernet loss- measurement

```
user@host> monitor ethernet loss-measurement mep 2 64:87:88:6a:da:94 maintenance-domain md
maintenance-association ma count 10
ETH-LM request to 64:87:88:6a:da:94, Interface ge-2/3/2.0

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):260
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.43537%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):1
```

```
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.51546%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):257
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.31615%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):261
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.77551%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):260
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.43537%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):1
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.51020%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):259
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.09524%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):519
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.71795%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):1
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.51020%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):259
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.09524%

--- Loss measurement statistics ---
LMM packets transmitted: 10, Valid LMR packets received: 8
Average near-end loss(CIR)         : 0.00000
Average near-end loss ratio(CIR)   : 0.00000%
Average far-end loss(CIR)          : 0.42857
Average far-end loss ratio(CIR)    : 0.21941%
Near-end best case loss(CIR)       : 0
Near-end best case loss ratio(CIR) : 0.00000%
Near-end worst case loss(CIR)      : 0
```

```
Near-end worst case loss ratio(CIR): 0.00000%
Far-end best case loss(CIR)        : 0
Far-end best case loss ratio(CIR)  : 0.00000%
Far-end worst case loss(CIR)       : 1
Far-end worst case loss ratio(CIR) : 0.51546%
Average near-end loss(EIR)         : 0.00000
Average near-end loss ratio(EIR)   : 0.00000%
Average far-end loss(EIR)          : 296.42857
Average far-end loss ratio(EIR)    : 88.41011%
Near-end best case loss(EIR)       : 0
Near-end best case loss ratio(EIR) : 0.00000%
Near-end worst case loss(EIR)      : 0
Near-end worst case loss ratio(EIR): 0.00000%
Far-end best case loss(EIR)        : 257
Far-end best case loss ratio(EIR)  : 88.09524%
Far-end worst case loss(EIR)       : 519
Far-end worst case loss ratio(EIR) : 88.77551%
```

Release Information

Command introduced in Junos OS Release 11.1.

`local-mep` option introduced in Junos OS Release 15.1

RELATED DOCUMENTATION

- *ITU-T Y.1731 Ethernet Service OAM Overview*
- Junos OS Network Interfaces Library for Routing Devices
- CLI Explorer

monitor ethernet synthetic-loss-measurement

IN THIS SECTION

- Syntax | 10851
- Description | 10851

- Options | 10852
- Additional Information | 10853
- Required Privilege Level | 10853
- Output Fields | 10853
- Sample Output | 10855
- Release Information | 10855

Syntax

```
monitor ethernet synthetic-loss-measurement
maintenance-domain md-name
maintenance-association ma-name
(remote-mac-address | mep remote-mep-id)
<count frame-count>
<local-mep mep-id>
<wait interval-milliseconds>
<priority 802.1p value>
<size>
<xml>
```

Description

(MX Series routers) Start an ITU-T Y.1731 Ethernet synthetic loss measurement (ETH-SLM) session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a summary of the frames exchanged in the measurement session. ETH-SLM (also called synthetic frame loss measurement) statistics are stored at one of the MEPs for later retrieval.

> **NOTE**: If you attempt to monitor delays to a nonexistent MAC address, you must press Ctrl +c to explicitly quit the `monitor ethernet synthetic-loss-measurement` command and return to the CLI command prompt.

To start an Ethernet synthetic frame loss measurement session, the router initiates an exchange of frames carrying synthetic frame loss measurement protocol data units (PDUs) between the local and remote MEPs. The frame counts—the types of and number of Ethernet synthetic frame loss measurement PDU frames exchanged to measure frame delay times—are displayed as the run-time

output of the `monitor ethernet synthetic-loss-measurement` command and are also stored at both the initiator and receiver MEPs for later retrieval. Ethernet synthetic frame loss measurement statistics, described below, are measured and stored at only one of the MEPs:

Frame delay—The difference, in microseconds, between the time a frame is sent and when it is received.

Frame delay variation—The difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.”

Options

maintenance-domain *md-name*	Name of an existing CFM maintenance domain.
maintenance-association *ma-name*	Name of an existing CFM maintenance association.
mep *remote-mep-id*	Numeric identifier of the peer MEP with which to perform Ethernet synthetic frame loss measurement. The discovered MAC address of the peer MEP is used. The range of values is from 1 through 8191.
remote-mac-address	Unicast MAC address of the peer MEP with which to perform Ethernet synthetic frame loss measurement. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn*.*nnnn*.*nnnn* or *nn*:*nn*:*nn*:*nn*:*nn*:*nn*. For example, `0011.2233.4455` or `00:11:22:33:44:55`. Multicast MAC addresses are not supported.
count *frame-count*	(Optional) Number of frames to send to the specified peer MEP. The range of values is from 1 through 65,535 frames. The default value is 10 frames.
local-mep *mep-id*	(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.
wait *interval-milliseconds*	(Optional) Number of milliseconds to wait between sending frames. You must specify this value in multiples of 100 milliseconds. The range of values is from 100 through 50000 milliseconds. The default value is 100 milliseconds.
priority *802.1p value*	(Optional) Priority of the ETH-SLM request frame supported. The range of values is from 0 through 7. The default value is zero.
size	(Optional) Size of the data TLV to be included in the request frame. The range of values is from 1 through 1400 bytes.
xml	(Optional) Allow the output of the command to be displayed in XML format for ETH-SLM. Note that the only way to get output in XML format is to use the `xml` argument. The `display xml` command does not work.

Additional Information

To display the frame counts collected at a MEP as the result of this command, use the following commands as described in the CLI Explorer:

- `show oam ethernet connectivity-fault-management interfaces detail`
- `show oam ethernet connectivity-fault-management mep-database`
- `show oam ethernet connectivity-fault-management mep-statistics`

To display the statistics collected at a MEP as the result of this command, use the following commands as described in the CLI Explorer.

- `show oam ethernet connectivity-fault-management synthetic-loss-measurement`
- `show oam ethernet connectivity-fault-management mep-statistics`

To clear both the frame counts and the statistics collected for MEPs, use the `clear oam ethernet connectivity-fault-management statistics` command, described in the CLI Explorer.

For a complete description of Ethernet synthetic frame loss measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the Junos OS Network Interfaces Library for Routing Devices.

Required Privilege Level

trace and maintenance

Output Fields

The `monitor ethernet synthetic-loss-measurement` command displays different output at the CLI, depending on when you start a synthetic frame loss measurement:

- Table 71 on page 10853 lists the run-time output fields for the `monitor ethernet synthetic-loss-measurement` command.

Output fields are listed in the approximate order in which they appear.

Table 71: monitor ethernet synthetic-loss-measurement Output Fields

Output Field Name	Output Field Description
`ETH-SLM request to`	Unicast MAC address of the remote peer MEP.

Table 71: monitor ethernet synthetic-loss-measurement Output Fields *(Continued)*

Output Field Name	Output Field Description
`Interface`	Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.
`SLM packets sent`	Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.
`SLR packets received`	Total number of synthetic loss reply (SLR) PDU frames received by the source MEP from the remote MEP during this measurement session.
`Local TXFC1 value`	Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.
`Local RXFC1 value`	Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet.
`Last Received SLR frame TXFCf(tc)`	Value of the local counter TxFCl at the time of SLM frame transmission.
`Last Received SLR frame TXFCb(t`	Value of the local counter RxFCl at the time of SLR frame transmission.
`Frame loss (near-end)`	Count of frame loss associated with ingress data frames.
`Frame loss (far-end)`	Count of frame loss associated with egress data frames.

Sample Output

monitor ethernet synthetic-loss-measurement

```
user@host> monitor ethernet synthetic-loss-measurement 00:05:85:73:39:4a maintenance-domain md6
maintenance-association ma6 count 10
ETH-SLM request to 00:05:85:73:39:4a, interface ge-1/0/0.0
      Synthetic Loss measurement statistics:
       SLM packets sent                            : 100
       SLR packets received                        : 100
       Accumulated SLM statistics:
       Local TXFC1 value                           : 100
       Local RXFC1 value                           : 100
       Last Received SLR frame TXFCf(tc)           : 100
       Last Received SLR frame TXFCb(tc)           : 100
       SLM Frame Loss:
       Frame Loss (far-end)                        : 0 (0.00 %)
       Frame Loss (near-end)                       : 0 (0.00 %)
```

Release Information

Command introduced in Junos OS Release 13.2.

`local-mep` option introduced in Junos OS Release 15.1

RELATED DOCUMENTATION

clear oam ethernet connectivity-fault-management synthetic-loss-measurement | 3777

monitor ethernet synthetic-loss-measurement | 10850

show oam ethernet connectivity-fault-management loss-statistics | 22615

monitor interface

IN THIS SECTION

- Syntax | 10856
- Description | 10856
- Options | 10856
- Additional Information | 10857
- Required Privilege Level | 10858
- Output Fields | 10858
- Sample Output | 10861
- Release Information | 10873

Syntax

```
monitor interface
<interface-name | traffic <detail>>
```

Description

Display real-time statistics about interfaces, updating the statistics every second. Check for and display common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors.

NOTE: On Junos OS Evolved, you can use the `monitor interface` command over SSH sessions, but console and Telnet sessions are not supported.

NOTE: This command is not supported on the QFX3000 QFabric switch.

Options

none Display real-time statistics for all interfaces.

detail (Optional) With traffic option only, display detailed output.

interface-name (Optional) Display real-time statistics for the specified interface. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified line-card chassis (LCC) only.

traffic (Optional) Display traffic data for all active interfaces. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified LCC only.

Additional Information

The output of this command shows how much each field has changed since you started the command or since you cleared the counters by pressing the c key. For a description of the statistical information provided in the output of this command, see the `show interfaces extensive` command for a particular interface type in the CLI Explorer. To control the output of the `monitor interface` command while it is running, use the keys listed in Table 72 on page 10857. The keys are not case-sensitive.

Table 72: Output Control Keys for the monitor interface interface-name Command

Key	Action
c	Clears (returns to zero) the delta counters since `monitor interface` was started. This does not clear the accumulative counter. To clear the accumulative counter, use the `clear interfaces interval` command.
f	Freezes the display, halting the display of updated statistics and delta counters.
i	Displays information about a different interface. The command prompts you for the name of a specific interface.
n	Displays information about the next interface. The `monitor interface` command displays the physical or logical interfaces in the same order as the `show interfaces terse` command.
q or Esc	Quits the command and returns to the command prompt.
t	Thaws the display, resuming the update of the statistics and delta counters.

To control the output of the `monitor interface traffic` command while it is running, use the keys listed in Table 73 on page 10858. The keys are not case-sensitive.

Table 73: Output Control Keys for the monitor interface traffic Command

Key	Action
b	Displays the statistics in units of bytes and bits per second (bps).
c	Clears (return to 0) the delta counters in the `Current Delta` column. The statistics counters are not cleared.
d	Displays the `Current Delta` column (instead of the rate column) in bps or packets per second (pps).
p	Displays the statistics in units of packets and packets per second (pps).
q or Esc	Quits the command and returns to the command prompt.
r	Displays the rate column (instead of the `Current Delta` column) in bps and pps.

Required Privilege Level

trace

Output Fields

Table 74 on page 10858 describes the output fields for the `monitor interface` command. Output fields are listed in the approximate order in which they appear.

Table 74: monitor interface Output Fields

Field Name	Field Description	Level of Output
`router1`	Hostname of the router.	All levels
`Seconds`	How long the monitor interface command has been running or how long since you last cleared the counters.	All levels
`Time`	Current time (UTC).	All levels

Table 74: monitor interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Delay` *x*/*y*/*z*	Time difference between when the statistics were displayed and the actual clock time. • *x*—Time taken for the last polling (in milliseconds). • *y*—Minimum time taken across all pollings (in milliseconds). • *z*—Maximum time taken across all pollings (in milliseconds).	All levels
`Interface`	Short description of the interface, including its name, status, and encapsulation.	All levels
`Link`	State of the link: `Up`, `Down`, or `Test`.	All levels
`Current delta`	Cumulative number for the counter in question since the time shown in the Seconds field, which is the time since you started the command or last cleared the counters.	All levels
`Local Statistics`	(Logical interfaces only) Number and rate of bytes and packets destined to the router or switch through the specified interface. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	All levels

Table 74: monitor interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Remote Statistics`	(Logical interfaces only) Statistics for traffic transiting the router or switch. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	All levels
Traffic statistics	Total number of bytes and packets received and transmitted on the interface. These statistics are the sum of the local and remote statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	All levels
`Description`	With the `traffic` option, displays the interface description configured at the [`edit interfaces` *`interface-name`*] hierarchy level.	`detail`
Input Errors	Sum of incoming frame aborts and FCS errors.	All levels
Input Drops	Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.	All levels
Input Framing errors	Number of packets received with an invalid frame checksum (FCS).	All levels

Table 74: monitor interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Policed discards	Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.	All levels
L3 incompletes	Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` statement.	All levels
L2 channel errors	Number of times the software did not find a valid logical interface for an incoming frame.	All levels
L2 mismatch timeouts	Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.	All levels

Sample Output

monitor interface (Physical)

```
user@host> monitor interface so-0/0/0
router1                         Seconds: 19                 Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC48
Traffic statistics:                                        Current Delta
  Input packets:                     6045 (0 pps)                   [11]
  Input bytes:                    6290065 (0 bps)                [13882]
  Output packets:                   10376 (0 pps)                   [10]
  Output bytes:                  10365540 (0 bps)                [9418]
Encapsulation statistics:
  Input keepalives:                  1901                            [2]
  Output keepalives:                 1901                            [2]
  NCP state: Opened
  LCP state: Opened
```

```
Error statistics:
  Input errors:                          0                              [0]
  Input drops:                           0                              [0]
  Input framing errors:                  0                              [0]
  Policed discards:                      0                              [0]
  L3 incompletes:                        0                              [0]
  L2 channel errors:                     0                              [0]
  L2 mismatch timeouts:                  0                              [0]
  Carrier transitions:                   1                              [0]
  Output errors:                         0                              [0]
  Output drops:                          0                              [0]
  Aged packets:                          0                              [0]
Active alarms : None
Active defects: None
SONET error counts/seconds:
  LOS count                              1                              [0]
  LOF count                              1                              [0]
  SEF count                              1                              [0]
  ES-S                                   0                              [0]
  SES-S                                  0                              [0]
SONET statistics:
  BIP-B1                            458871                              [0]
  BIP-B2                            460072                              [0]
  REI-L                             465610                              [0]
  BIP-B3                            458978                              [0]
  REI-P                             458773                              [0]
Received SONET overhead:
  F1      : 0x00  J0        : 0x00  K1        : 0x00
  K2      : 0x00  S1        : 0x00  C2        : 0x00
  C2(cmp) : 0x00  F2        : 0x00  Z3        : 0x00
  Z4      : 0x00  S1(cmp)   : 0x00
Transmitted SONET overhead:
  F1      : 0x00  J0        : 0x01  K1        : 0x00
  K2      : 0x00  S1        : 0x00  C2        : 0xcf
  F2      : 0x00  Z3        : 0x00  Z4        : 0x00

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (OTN Interface)

```
user@host> monitor interface ge-7/0/0


Interface: ge-7/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
  Input bytes:                        0 (0 bps)
  Output bytes:                       0 (0 bps)
  Input packets:                      0 (0 pps)
  Output packets:                     0 (0 pps)
Error statistics:
  Input errors:                       0
  Input drops:                        0
  Input framing errors:               0
  Policed discards:                   0
  L3 incompletes:                     0
  L2 channel errors:                  0
  L2 mismatch timeouts:               0
  Carrier transitions:                5
  Output errors:                      0
  Output drops:                       0
  Aged packets:                       0
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
  Unicast packets                     0
  Broadcast packets                   0
  Multicast packets                   0
  Oversized frames                    0
  Packet reject count                 0
  DA rejects                          0
  SA rejects                          0
Output MAC/Filter Statistics:
  Unicast packets                     0
  Broadcast packets                   0
  Multicast packets                   0
  Packet pad count                    0
  Packet error count                  0
OTN Link 0
  OTN Alarms: OTU_BDI, OTU_TTIM, ODU_BDI
```

```
  OTN Defects: OTU_BDI, OTU_TTIM, ODU_BDI, ODU_TTIM
  OTN OC - Seconds
    LOS                                 2
    LOF                                 9
  OTN OTU - FEC Statistics
    Corr err ratio                    N/A
    Corr bytes                          0
    Uncorr words                        0
  OTN OTU - Counters
    BIP                                 0
    BBE                                 0
    ES                                  0
    SES                                 0
    UAS                               422
  OTN ODU - Counters
    BIP                                 0
    BBE                                 0
    ES                                  0
    SES                                 0
    UAS                               422
  OTN ODU - Received Overhead    APSPCC 0-3:              0
```

monitor interface (MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface)

```
user@host> monitor interface xe-0/0/3
Interface: xe-0/0/3, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:                                          Current delta
  Input bytes:                        0 (0 bps)                         [0]
  Output bytes:                       0 (0 bps)                         [0]
  Input packets:                      0 (0 pps)                         [0]
  Output packets:                     0 (0 pps)                         [0]
Error statistics:
  Input errors:                       0                                 [0]
  Input drops:                        0                                 [0]
  Input framing errors:               0                                 [0]
  Policed discards:                   0                                 [0]
  L3 incompletes:                     0                                 [0]
  L2 channel errors:                  0                                 [0]
  L2 mismatch timeouts:               0                                 [0]
  Carrier transitions:                5                                 [0]
```

```
  Output errors:                      0                                   [0]
  Output drops:                       0                                   [0]
  Aged packets:                       0                                   [0]
Active alarms : None
Active defects: None
PCS statistics:                     Seconds
  Bit Errors                          0                                   [0]
  Errored blocks                      4                                   [0]
Input MAC/Filter statistics:
  Unicast packets                     0                                   [0]
  Broadcast packets                   0                                   [0]
  Multicast packets                   0                                   [0]
  Oversized frames                    0                                   [0]
  Packet reject count                 0                                   [0]
  DA rejects                          0                                   [0]
  SA rejects                          0                                   [0]
Output MAC/Filter Statistics:
  Unicast packets                     0                                   [0]
  Broadcast packets                   0                                   [0]
  Multicast packets                   0                                   [0]
  Packet pad count                    0                                   [0]
  Packet error count                  0                                   [0]



Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (MX480 Router with MPC5E and 100-Gigabit Ethernet Interface)

```
user@host> monitor interface et-2/1/0
Interface: et-2/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100000mbps
Traffic statistics:                                           Current delta
  Input bytes:                        0 (0 bps)                           [0]
  Output bytes:                       0 (0 bps)                           [0]
  Input packets:                      0 (0 pps)                           [0]
  Output packets:                     0 (0 pps)                           [0]
Error statistics:
  Input errors:                       0                                   [0]
  Input drops:                        0                                   [0]
  Input framing errors:               0                                   [0]
```

```
  Policed discards:                    0                                    [0]
  L3 incompletes:                      0                                    [0]
  L2 channel errors:                   0                                    [0]
  L2 mismatch timeouts:                0                                    [0]
  Carrier transitions:               263                                    [0]
  Output errors:                       0                                    [0]
  Output drops:                        0                                    [0]
  Aged packets:                        0                                    [0]
OTN Link 0
  OTN Alarms:
  OTN Defects:
  OTN OC - Seconds
    LOS                              129                                    [0]
    LOF                                2                                    [0]
  OTN OTU - FEC Statistics
    Corr err ratio                 <8E-5
    Corr bytes              169828399453                                    [0]
    Uncorr words             28939961456                                    [0]
  OTN OTU - Counters                                                        [0]
    BIP                                0
    BBE                                0                                    [0]
    ES                                24                                    [0]
    SES                                0                                    [0]
    UAS                             1255                                    [0]
  OTN ODU - Counters                                                        [0]
    BIP                                0
    BBE                                0                                    [0]
    ES                                24                                    [0]
    SES                                0                                    [0]
    UAS                             1256                                    [0]
  OTN ODU - Received Overhead                                               [0]
    APSPCC 0-3:              00 00 00 00


Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface)

```
user@host> monitor interface xe-6/1/0
Interface: xe-6/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
```

```
Traffic statistics:                                          Current delta
  Input bytes:                        0 (0 bps)                        [0]
  Output bytes:                       0 (0 bps)                        [0]
  Input packets:                      0 (0 pps)                        [0]
  Output packets:                     0 (0 pps)                        [0]
Error statistics:
  Input errors:                       0                                [0]
  Input drops:                        0                                [0]
  Input framing errors:               0                                [0]
  Policed discards:                   0                                [0]
  L3 incompletes:                     0                                [0]
  L2 channel errors:                  0                                [0]
  L2 mismatch timeouts:               0                                [0]
  Carrier transitions:                1                                [0]
  Output errors:                      0                                [0]
  Output drops:                       0                                [0]
  Aged packets:                       0                                [0]
Active alarms : None
Active defects: None
PCS statistics:                     Seconds
  Bit Errors                          0                                [0]
  Errored blocks                      1                                [0]
Input MAC/Filter statistics:
  Unicast packets                     0                                [0]
  Broadcast packets                   0                                [0]
  Multicast packets                   0                                [0]
  Oversized frames                    0                                [0]
  Packet reject count                 0                                [0]
  DA rejects                          0                                [0]
  SA rejects                          0                                [0]
Output MAC/Filter Statistics:
  Unicast packets                     0                                [0]
  Broadcast packets                   0                                [0]
  Multicast packets                   0                                [0]
  Packet pad count                    0                                [0]
  Packet error count                  0                                [0]



Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface)

```
user@host> monitor interface et-9/0/0
Interface: et-9/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100000mbps
Traffic statistics:                                          Current delta
  Input bytes:                        0 (0 bps)                         [0]
  Output bytes:                       0 (0 bps)                         [0]
  Input packets:                      0 (0 pps)                         [0]
  Output packets:                     0 (0 pps)                         [0]
Error statistics:
  Input errors:                       0                                 [0]
  Input drops:                        0                                 [0]
  Input framing errors:               0                                 [0]
  Policed discards:                   0                                 [0]
  L3 incompletes:                     0                                 [0]
  L2 channel errors:                  0                                 [0]
  L2 mismatch timeouts:               0                                 [0]
  Carrier transitions:                1                                 [0]
  Output errors:                      0                                 [0]
  Output drops:                       0                                 [0]
  Aged packets:                       0                                 [0]




Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface)

```
user@host> monitor interface xe-3/0/0
host name                           Seconds: 67                  Time: 23:46:46
                                                            Delay: 0/0/13
Interface: xe-3/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:                                          Current delta
  Input bytes:                        0 (0 bps)                         [0]
  Output bytes:                       0 (0 bps)                         [0]
  Input packets:                      0 (0 pps)                         [0]
  Output packets:                     0 (0 pps)                         [0]
```

```
Error statistics:
  Input errors:                        0                                    [0]
  Input drops:                         0                                    [0]
  Input framing errors:                0                                    [0]
  Policed discards:                    0                                    [0]
  L3 incompletes:                      0                                    [0]
  L2 channel errors:                   0                                    [0]
  L2 mismatch timeouts:                0                                    [0]
  Carrier transitions:                 3                                    [0]
  Output errors:                       0                                    [0]
  Output drops:                        0                                    [0]
  Aged packets:                        0                                    [0]
OTN Link 0
  OTN Alarms:
  OTN Defects:
  OTN OC - Seconds
    LOS                                0                                    [0]
    LOF                                0                                    [0]
  OTN OTU - FEC Statistics
    Corr err ratio                   N/A
    Corr bytes                         0                                    [0]
    Uncorr words                       0                                    [0]
  OTN OTU - Counters                                                        [0]
    BIP                                0
    BBE                                0                                    [0]
    ES                                 0                                    [0]
    SES                                0                                    [0]
    UAS                                0                                    [0]
  OTN ODU - Counters                                                        [0]
    BIP                                0
    BBE                                0                                    [0]
    ES                                 0                                    [0]
    SES                                0                                    [0]
    UAS                                0                                    [0]
  OTN ODU - Received Overhead                                               [0]
    APSPCC 0-3:              00 00 00 00


Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (Logical)

```
user@host> monitor interface so-1/0/0.0
host name                Seconds: 16                  Time: 15:33:39
                                                               Delay: 0/0/1
Interface: so-1/0/0.0, Enabled, Link is Down
Flags: Hardware-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Local statistics:                                          Current delta
  Input bytes:                          0                             [0]
  Output bytes:                         0                             [0]
  Input packets:                        0                             [0]
  Output packets:                       0                             [0]
Remote statistics:
  Input bytes:                          0 (0 bps)                     [0]
  Output bytes:                         0 (0 bps)                     [0]
  Input packets:                        0 (0 pps)                     [0]
  Output packets:                       0 (0 pps)                     [0]
Traffic statistics:
  Destination address: 192.168.8.193, Local: 192.168.8.21

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (QFX3500 Switch)

```
user@switch> monitor interface ge-0/0/0
Interface: ge-0/0/0, Enabled, Link is Down
Encapsulation: Ethernet, Speed: Unspecified
Traffic statistics:                                        Current delta
  Input bytes:                        0 (0 bps)                       [0]
  Output bytes:                       0 (0 bps)                       [0]
  Input packets:                      0 (0 pps)                       [0]
  Output packets:                     0 (0 pps)                       [0]
Error statistics:
  Input errors:                       0                               [0]
  Input drops:                        0                               [0]
  Input framing errors:               0                               [0]
  Policed discards:                   0                               [0]
  L3 incompletes:                     0                               [0]
  L2 channel errors:                  0                               [0]
```

```
  L2 mismatch timeouts:                0                                [0]
  Carrier transitions:                 0                                [0]
  Output errors:                       0                                [0]
  Output drops:                        0                                [0]
  Aged packets:                        0                                [0]
Active alarms : LINK
Active defects: LINK
Input MAC/Filter statistics:
  Unicast packets                      0                                [0]
  Broadcast packets                    0  Multicast packet              [0]

Interface warnings:
  o Outstanding LINK alarm
```

monitor interface traffic

```
user@host> monitor interface traffic
host name                Seconds: 15                  Time: 12:31:09

Interface    Link  Input packets        (pps)     Output packets        (pps)
 so-1/0/0    Down              0          (0)                  0          (0)
 so-1/1/0    Down              0          (0)                  0          (0)
 so-1/1/1    Down              0          (0)                  0          (0)
 so-1/1/2    Down              0          (0)                  0          (0)
 so-1/1/3    Down              0          (0)                  0          (0)
 t3-1/2/0    Down              0          (0)                  0          (0)
 t3-1/2/1    Down              0          (0)                  0          (0)
 t3-1/2/2    Down              0          (0)                  0          (0)
 t3-1/2/3    Down              0          (0)                  0          (0)
 so-2/0/0      Up         211035          (1)              36778          (0)
 so-2/0/1      Up         192753          (1)              36782          (0)
 so-2/0/2      Up         211020          (1)              36779          (0)
 so-2/0/3      Up         211029          (1)              36776          (0)
 so-2/1/0      Up         189378          (1)              36349          (0)
 so-2/1/1    Down              0          (0)              18747          (0)
 so-2/1/2    Down              0          (0)              16078          (0)
 so-2/1/3      Up              0          (0)              80338          (0)
 at-2/3/0      Up              0          (0)                  0          (0)
 at-2/3/1    Down              0          (0)                  0          (0)
```

```
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

monitor interface traffic (QFX3500 Switch)

```
user@switch> monitor interface traffic
switch                              Seconds: 7                  Time: 16:04:37

Interface    Link  Input packets        (pps)     Output packets        (pps)
 ge-0/0/0    Down              0          (0)                  0          (0)
 ge-0/0/1      Up         392187          (0)             392170          (0)
 ge-0/0/2    Down              0          (0)                  0          (0)
 ge-0/0/3    Down              0          (0)                  0          (0)
 ge-0/0/4    Down              0          (0)                  0          (0)
 ge-0/0/5    Down              0          (0)                  0          (0)
 ge-0/0/6    Down              0          (0)                  0          (0)
 ge-0/0/7    Down              0          (0)                  0          (0)
 ge-0/0/8    Down              0          (0)                  0          (0)
 ge-0/0/9      Up         392184          (0)             392171          (0)
 ge-0/0/10   Down              0          (0)                  0          (0)
 ge-0/0/11   Down              0          (0)                  0          (0)
 ge-0/0/12   Down              0          (0)                  0          (0)
 ge-0/0/13   Down              0          (0)                  0          (0)
 ge-0/0/14   Down              0          (0)                  0          (0)
 ge-0/0/15   Down              0          (0)                  0          (0)
 ge-0/0/16   Down              0          (0)                  0          (0)
 ge-0/0/17   Down              0          (0)                  0          (0)
 ge-0/0/18   Down              0          (0)                  0          (0)
 ge-0/0/19   Down              0          (0)                  0          (0)
 ge-0/0/20   Down              0          (0)                  0          (0)
 ge-0/0/21   Down              0          (0)                  0          (0)
 ge-0/0/22     Up         392172          (0)             392187          (0)
 ge-0/0/23     Up         392185          (0)             392173          (0)
 vcp-0       Down              0                               0
 vcp-1       Down              0                               0
 ae0         Down              0          (0)                  0          (0)
 bme0          Up              0                         1568706
```

monitor interface traffic detail (QFX3500 Switch)

```
user@switch> monitor interface traffic detail
switch                                                  Seconds: 74                 Time:
16:03:02

Interface    Link  Input packets        (pps)     Output packets        (pps) Description
 ge-0/0/0    Down              0          (0)                  0          (0)
 ge-0/0/1      Up         392183          (0)             392166          (0)
 ge-0/0/2    Down              0          (0)                  0          (0)
 ge-0/0/3    Down              0          (0)                  0          (0)
 ge-0/0/4    Down              0          (0)                  0          (0)
 ge-0/0/5    Down              0          (0)                  0          (0)
 ge-0/0/6    Down              0          (0)                  0          (0)
 ge-0/0/7    Down              0          (0)                  0          (0)
 ge-0/0/8    Down              0          (0)                  0          (0)
 ge-0/0/9      Up         392181          (0)             392168          (0)
 ge-0/0/10   Down              0          (0)                  0          (0)
 ge-0/0/11   Down              0          (0)                  0          (0)
 ge-0/0/12   Down              0          (0)                  0          (0)
 ge-0/0/13   Down              0          (0)                  0          (0)
 ge-0/0/14   Down              0          (0)                  0          (0)
 ge-0/0/15   Down              0          (0)                  0          (0)
 ge-0/0/16   Down              0          (0)                  0          (0)
 ge-0/0/17   Down              0          (0)                  0          (0)
 ge-0/0/18   Down              0          (0)                  0          (0)
 ge-0/0/19   Down              0          (0)                  0          (0)
 ge-0/0/20   Down              0          (0)                  0          (0)
 ge-0/0/21   Down              0          (0)                  0          (0)
 ge-0/0/22     Up         392169          (0)             392184          (1)
 ge-0/0/23     Up         392182          (0)             392170          (0)
 vcp-0       Down              0                               0
 vcp-1       Down              0                               0
 ae0         Down              0          (0)                  0          (0)
 bme0          Up              0                         1568693
```

Release Information

Command introduced before Junos OS Release 7.4.

monitor interface error

IN THIS SECTION

- Syntax | 10874
- Description | 10874
- Options | 10874
- Required Privilege Level | 10874
- Output Fields | 10874
- Sample Output | 10876
- Release Information | 10876

Syntax

```
monitor interface error
```

Description

You can use the `monitor interface error` command to display errors in packet headers that are transmitted through all network interfaces. Use the following options at the end of the command for more specific information.

Options

`interface-name` (Optional) Displays packet header errors for a specified interface or interfaces. If an interface is not specified, all interfaces are monitored.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `monitor interface error` command.

Table 75: monitor interface error Output Fields

Field Name	Field Description
`CRC Errors`	Displays the number of CRC errors on the selected interfaces since you started monitoring.
`Corr errors`	Displays the number of FEC corrected errors since monitoring began.
`FEC Mode`	Displays the FEC mode configured on the interface.
`Interface Name`	Displays the interface name.
`Link Flaps`	Displays the number of link flaps observed during the current monitoring session.
`Link Status`	Displays the interface link status.
`PCS bit err sec`	Displays the PCS bit error count.
`PCS err blk sec`	Displays the number of PCS error blocks found.
`Pre-FEC BER`	Displays the pre-FEC bit error rate (BER). This is used to monitor the condition of an optical transport network (OTN) link.
`Speed`	Displays the speed of the interface.
`Since`	Displays the time since the error table was cleared/updated.
`Uncorr errors`	Displays the number of FEC uncorrected errors since monitoring began.

Sample Output

monitor interface error

```
user@host> monitor interface error
```

```
EVOvSCAPA1-RE0-re0                  Seconds: 6          Time: 11:01:51

Interface  Speed  Link   Link   PCS bit  PCS err  CRC     FEC   Corr   Uncorr  Pre-FEC   Since
                  status flaps  err sec  blk sec  errors  mode  errors errors  BER
 et-1/0/1  400G    Up      0       0        0     0        -      -      -      -      00:25:51
ago
 et-1/0/2  400G    Up      1       0        0     0        -      -      -      -         Never
 et-1/0/3  400G    Up      1       0        0     0        -      -      -      -         Never
 et-1/0/4  400G    Up      1       0        0     0        -      -      -      -         Never
et-1/0/11  400G    Up      1       0        0     0        -      -      -      -         Never
et-1/0/12  400G    Up      1       0        0     0        -      -      -      -         Never
et-1/0/13  400G    Up      1       0        0     0        -      -      -      -         Never
et-1/0/14  400G    Up      1       0        0     0        -      -      -      -         Never

Clear='c', Next Page='d', Prev Page='u', Quit='q' or ESC
```

monitor interface error et-1/0/1

```
user@host>monitor interface error et-1/0/1
```

```
EVOvSCAPA1-RE0-re0

Interface  Speed  Link   Link   PCS bit  PCS err  CRC     FEC   Corr    Uncorr  Pre-FEC   Since
                  status flaps  err sec  blk sec  errors  mode  errors  errors  BER
 et-1/0/1  400G    Up     0        0        0       0      -      -       -       -       Never
```

Release Information

Command introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

Monitor Interfaces (overview)

monitor interface | 10856

monitor label-switched-path

IN THIS SECTION

- Syntax | 10877
- Description | 10877
- Options | 10877
- Additional Information | 10878
- Required Privilege Level | 10878
- Output Fields | 10879
- Sample Output | 10880
- Release Information | 10881

Syntax

```
monitor label-switched-path lsp-name
<logical-system (logical-system-name)>
```

Description

Display the real-time status of the specified RSVP label-switched path (LSP). You can also use this command to monitor LSPs configured within logical systems.

Options

logical-system (*logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-name Name of the LSP.

Additional Information

You can track the amount of traffic traversing an RSVP LSP and observe its essential parameters, such as uptime, ingress and egress addresses, labels, routes, and ports. Values are typically sampled every second. The display also allows you to scroll to other currently running LSPs. You cannot use this command to display information about static LSPs or LDP-signaled LSPs.

The output of this command shows how much each field has changed since you started the command or since you cleared the counters by using the c key. To control the output of the `monitor label-switched-path` command while it is running, use the keys listed in Table 76 on page 10878. The keys are not case-sensitive.

Table 76: Output Control Keys for the monitor label-switched-path Command

Key	Action
c	Clears the screen and refreshes the display for this LSP.
f	Freezes the display, preventing new information from being displayed.
l	Monitors a different LSP. After you type l, you can type the new LSP name.
n	Displays information about the next LSP (whose name is alphabetically higher than the current LSP name) configured on the router.
p	Goes to the previous LSP (whose name is alphabetically lower than the current LSP name) configured on the router.
q or Esc	Quits the command and returns to the command prompt.
t	Thaws, or restarts, the data display for this LSP.

Required Privilege Level

trace

Output Fields

Table 77 on page 10879 describes the output fields for the `monitor label-switched-path` command. Output fields are listed in the approximate order in which they appear.

Table 77: monitor label-switched-path Output Fields

Field Name	Field Description
(1)	Displays the following information: • *hostname*—Name of the router. • Seconds—Time elapsed since this display was started. • Time—Current local time.
(2)	Delay—Length of the time delay, in milliseconds, required to obtain the information in the monitor display. The first number shows the current sampling delay. The second number shows the shortest delay recorded to date. The third number shows the worst delay recorded to date. This delay can vary substantially depending on the system load.
(3)	Displays the following: • To—Destination address of the LSP. • From—Originating address of the LSP. • State—Current state of the LSP: Up or Down.
(4)	Displays the following: • LSPName—Name of the LSP. • Type—Type of LSP: Ingress, Egress, or Transit.
(5)	Displays the following: • Label in—Incoming label of the LSP. • Label out—Outgoing label of the LSP.

Table 77: monitor label-switched-path Output Fields *(Continued)*

Field Name	Field Description
(6)	`Port number`—Port number for the sending router, the port number for the receiving router, and the protocol ID. For MPLS traffic engineering applications, the protocol ID is always `0`.
(7/8)	`Record route`—All intermediate and egress router addresses for this LSP.
(9/10/11)	Displays traffic statistics: • `Output packets`—Number of packets that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago. • `Output bytes`—Number of bytes that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago.
(12)	Displays any errors the router encountered while attempting to retrieve information on the LSP.
(13)	Lists the keyboard commands you can use to navigate to other LSPs. For a description of the keyboard commands, see Table 76 on page 10878.

Sample Output

monitor label-switched-path

```
user@host> monitor label-switched-path
(1) host                              Seconds: 112              Time: 15:32:22
(2)                                                              Delay: 0/0/0
(3) To 10.10.10.16, From 10.10.10.17, state: Up
(4)   LSPname: k, type: Ingress
(5)   Label in: -, Label out: 126000
(6)   Port number: sender 1, receiver 45583, protocol 0
(7)   Record Route: <self> 192.168.224.196
(8)     192.168.224.202 192.168.224.179
(9)   Traffic statistics:                                        Current delta
(10)    Output packets:                       0                            [0]
(11)    Output bytes:                         0                            [0]
```

```
(12)
(13)Next='n', Prev='p', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c',     LSP='l'
```

Release Information

Command introduced before Junos OS Release 7.4.

Logical system support introduced in Junos OS Release 9.4.

monitor list

IN THIS SECTION

- Syntax | 10881
- Description | 10881
- Options | 10881
- Additional Information | 10882
- Required Privilege Level | 10882
- Output Fields | 10882
- Sample Output | 10882
- Release Information | 10883

Syntax

```
monitor list
```

Description

Display the status of monitored log and trace files.

Options

This command has no options.

Additional Information

Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are configured with the `syslog` statement at the `[edit system]` hierarchy level and the `options` statement at the `[edit routing-options]` hierarchy level. The trace files generated by the routing protocol process are those configured with `traceoptions` statements at the `[edit routing-options]`, `[edit interfaces]`, and `[edit protocols` *`protocol`*`]` hierarchy levels.

Required Privilege Level

trace

Output Fields

Table 78 on page 10882 describes the output fields for the `monitor list` command. Output fields are listed in the approximate order in which they appear.

Table 78: monitor list Output Fields

Field Name	Field Description
`monitor start`	Indicates the file is being monitored.
"*`filename`*"	Name of the file that is being monitored.
`Last changed`	Date and time at which the file was last modified.

Sample Output

monitor list

```
user@host> monitor list
monitor start "vrrpd" (Last changed Dec 03:11:06 20)
monitor start "cli-commands" (Last changed Nov 07:3)
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

monitor start | 10927

monitor stop | 10932

monitor list (JDM)

IN THIS SECTION

- Syntax | 10883
- Description | 10883
- Required Privilege Level | 10883
- Output Fields | 10883
- Sample Output | 10884
- Release Information | 10884

Syntax

```
monitor list
```

Description

Show status of monitored files.

Required Privilege Level

View

Output Fields

Sample Output

monitor list

```
user@jdm> monitor list
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

monitor start (JDM) | 10930

monitor mpls delay rsvp

IN THIS SECTION

- Syntax | 10884
- Description | 10885
- Options | 10885
- Required Privilege Level | 10885
- Output Fields | 10886
- Sample Output | 10888
- Release Information | 10890

Syntax

```
monitor mpls delay rsvp lsp-name
<detail>
<count count>
<interval seconds>
```

```
<padding-size padding-size>
<traffic-class traffic-class>
```

Description

Perform an on-demand delay measurement and display the measured values for associated bidirectional MPLS ultimate hop popping (UHP) point-to-point label-switched paths (LSPs).

Options

lsp-name Name of the associated bidirectional MPLS UHP LSP for which the delay measurement is performed.

detail (Optional) Display detailed output of the LSP delay measurement.

count *count* (Optional) Specify the number of delay measurements to be carried out for the MPLS UHP LSP. For LSP delay measurement, the number of queries sent is the specified count number plus one additional query, because the LSP delay is measured using successive messages.

- **Default:** 10
- **Range:** 1 through 1000000

interval *seconds* (Optional) Specify in seconds the interval between two successive query messages.

- **Range:** 1 through 255 seconds

padding-size *padding-size* (Optional) Specify the length of padding TLV to be included in the query message.

- **Range:** 0 through 1500

traffic-class *traffic-class* (Optional) Specify the traffic class for the LSP delay measurement. When the traffic-class value is not specified, the default traffic-class code-point of 111 is used.

- **Range:** 0 though 7

Required Privilege Level

view

Output Fields

Table 79 on page 10886 describes the output fields for the `monitor mpls delay rsvp` command. Output fields are listed in the approximate order in which they appear.

Table 79: monitor mpls delay rsvp Output Fields

Field Name	Field Description	Level of Output
Current two-way channel delay	Sum of packet delays, excluding the processing time of the remote provider edge (PE) router.	All Levels
Current round-trip-time	Total time taken for completing round-trip of packet.	All Levels
Best two-way channel delay	Best available two-way channel delay count.	All Levels
Worst two-way channel delay	Worst available two-way channel delay count.	All Levels
Average two-way channel delay	Average of the available two-way channel delay counts.	All Levels
Best round-trip-time	Best available round-trip-time count.	All Levels
Worst round-trip-time	Worst available round-trip-time count.	All Levels
Average round-trip-time	Average of the available round-trip-time counts.	All Levels
Average forward delay variation	Average of the variation in forward delay.	All Levels
Average reverse delay variation	Average of the variation in reverse delay.	All Levels
DM queries sent	Number of queries sent for delay measurement.	All Levels

Table 79: monitor mpls delay rsvp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
DM responses received	Number of responses received for delay measurement queries.	All Levels
DM queries timedout	Number of timed out queries sent for delay measurement.	All Levels
DM responses dropped due to errors	Number of loss measurement responses dropped due to errors.	All Levels
Response code	Status of the messages used for delay measurement. Response code can be one of the following: • `Success`—Successful response code. • `Failed`—Failed response code.	`detail`
Querier transmit timestamp	Timestamp on the query message when the message is sent out the ingress PE router (querier). This is done in the hardware before packet is sent out of an interface.	`detail`
Responder receive timestamp	Timestamp on the response message when the message is received by the egress PE router (responder). This is done in the hardware before packet is received by an interface.	`detail`
Responder transmit timestamp	Timestamp on the query message when the message is sent out the egress PE router (responder). This is done in the hardware before packet is sent out of an interface.	`detail`
Querier receive timestamp	Timestamp on the response message when the message is received by the ingress PE router (querier). This is done in the hardware before packet is received by an interface.	`detail`

Sample Output

monitor mpls lsp delay rsvp count

```
user@host> monitor mpls lsp delay rsvp LSP-A count 2

(1)
Current two-way channel delay            : 44 usecs
Current round-trip-time                  : 3243 usecs
(2)
Current two-way channel delay            : 45 usecs
Current round-trip-time                  : 1752 usecs

Best two-way channel delay               : 44 usecs
Worst two-way channel delay              : 45 usecs
Average two-way channel delay            : 45 usecs
Best round-trip-time                     : 1752 usecs
Worst round-trip-time                    : 3243 usecs
Average round-trip-time                  : 2498 usecs
Average forward delay variation          : 1 usecs
Average reverse delay variation          : 1 usecs

DM queries sent                          : 2
DM responses received                    : 2
DM queries timedout                      : 0
DM responses dropped due to errors       : 0
```

monitor mpls lsp delay rsvp count detail

```
user@host> monitor mpls lsp delay rsvp LSP-A count 2 detail

(1)
Response code                            : Success
Querier transmit timestamp               : 1404129122 secs, 479955401 nsecs
Responder receive timestamp              : 1404129122 secs, 468519022 nsecs
Responder transmit timestamp             : 1404129122 secs, 470255123 nsecs
Querier receive timestamp                : 1404129122 secs, 481736403 nsecs
Current two-way channel delay            : 44 usecs
Current round-trip-time                  : 1781 usecs
(2)
```

```
Response code                           : Success
Querier transmit timestamp              : 1404129123 secs, 480926210 nsecs
Responder receive timestamp             : 1404129123 secs, 469488696 nsecs
Responder transmit timestamp            : 1404129123 secs, 471130706 nsecs
Querier receive timestamp               : 1404129123 secs, 482613911 nsecs
Current two-way channel delay           : 45 usecs
Current round-trip-time                 : 1687 usecs

Best two-way channel delay              : 44 usecs
Worst two-way channel delay             : 45 usecs
Average two-way channel delay           : 45 usecs
Best round-trip-time                    : 1687 usecs
Worst round-trip-time                   : 1781 usecs
Average round-trip-time                 : 1734 usecs
Average forward delay variation         : 1 usecs
Average reverse delay variation         : 1 usecs

DM queries sent                         : 2
DM responses received                   : 2
DM queries timedout                     : 0
DM responses dropped due to errors      : 0
user@host> monitor mpls loss-delay-measurement lsp LSP1_A_to_B count 2
(1)
Current forward loss                    : 0 packets
Current forward loss ratio              : 0.000000
Current forward throughput              : 0.957 kpps
Current reverse loss                    : 0 packets
Current reverse loss ratio              : 0.000000
Current reverse throughput              : 0.962 kpps
Current two-way channel delay           : 48 usecs
Current round-trip-time                 : 3476 usecs
(2)
Current forward loss                    : 0 packets
Current forward loss ratio              : 0.000000
Current forward throughput              : 0.599 kpps
Current reverse loss                    : 0 packets
Current reverse loss ratio              : 0.000000
Current reverse throughput              : 0.599 kpps
Current two-way channel delay           : 50 usecs
Current round-trip-time                 : 1856 usecs

Cumulative forward transmit count       : 1557
Cumulative forward loss                 : 0 packets
```

```
Average forward loss ratio              : 0.000000
Average forward throughput              : 0.778 kpps
Cumulative reverse transmit count       : 1562
Cumulative reverse loss                 : 0 packets
Average reverse loss ratio              : 0.000000
Average reverse throughput              : 0.780 kpps

Best two-way channel delay              : 48 usecs
Worst two-way channel delay             : 50 usecs
Average two-way channel delay           : 49 usecs
Best round-trip-time                    : 1856 usecs
Worst round-trip-time                   : 3476 usecs
Average round-trip-time                 : 2445 usecs
Average forward delay variation         : 1 usecs
Average reverse delay variation         : 1 usecs

LDM queries sent                        : 3
LDM responses received                  : 3
LDM queries timedout                    : 0
LDM responses dropped due to errors     : 0
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

monitor mpls loss rsvp | 10896

monitor mpls loss-delay rsvp | 10890

Example: Configuring On-Demand Loss and Delay Measurement

monitor mpls loss-delay rsvp

IN THIS SECTION

Syntax | 10891

- Description | 10891
- Options | 10891
- Required Privilege Level | 10892
- Output Fields | 10892
- Sample Output | 10892
- Release Information | 10896

Syntax

```
monitor mpls loss-delay rsvp lsp-name
<detail>
<bytes>
<count count>
<interval seconds>
<padding-size padding-size>
<traffic-class traffic-class>
```

Description

Perform a simultaneous on-demand loss and delay measurement using combined loss and delay messages, and display the measured values for associated bidirectional MPLS ultimate hop popping (UHP) point-to-point label-switched paths (LSPs).

Options

lsp-name	Name of the associated bidirectional MPLS UHP LSP for which the delay measurement is performed.
detail	(Optional) Display detailed output of the LSP delay measurement.
bytes	(Optional) Specify the measurement quantity for the LSP loss measurement as bytes. By default, LSP loss is measured in packets.

> **NOTE**: The byte count of a packet sent or received over a channel counts only the payload, including the total length of that packet and excluding the headers, labels, and framing of the channel itself.

count *count*

(Optional) Specify the number of delay measurements to be carried out for the MPLS UHP LSP. For LSP delay measurement, the number of queries sent is the specified count number plus one additional query, because the LSP delay is measured using successive messages.

- **Default:** 10
- **Range:** 1 through 1000000

interval *seconds*

(Optional) Specify in seconds the interval between two successive query messages.

- **Range:** 1 through 255 seconds

padding-size *padding-size*

(Optional) Specify the length of padding TLV to be included in the query message.

- **Range:** 0 through 1500

traffic-class *traffic-class*

(Optional) Specify the traffic class for the LSP delay measurement. When the traffic-class value is not specified, the default traffic-class code-point of 111 is used.

- **Range:** 0 though 7

Required Privilege Level

view

Output Fields

For output field descriptions, see the "monitor mpls loss rsvp" on page 10896 and "monitor mpls delay rsvp" on page 10884 commands.

Sample Output

monitor mpls loss-delay rsvp count

```
user@host> monitor mpls loss-delay rsvp LSP-A count 2
```

```
(1)
Current forward loss                     : 0 packets
Current forward loss ratio               : 0.000000
Current forward throughput               : 0.957 kpps
Current reverse loss                     : 0 packets
Current reverse loss ratio               : 0.000000
Current reverse throughput               : 0.962 kpps
Current two-way channel delay            : 48 usecs
Current round-trip-time                  : 3476 usecs
(2)
Current forward loss                     : 0 packets
Current forward loss ratio               : 0.000000
Current forward throughput               : 0.599 kpps
Current reverse loss                     : 0 packets
Current reverse loss ratio               : 0.000000
Current reverse throughput               : 0.599 kpps
Current two-way channel delay            : 50 usecs
Current round-trip-time                  : 1856 usecs

Cumulative forward transmit count        : 1557
Cumulative forward loss                  : 0 packets
Average forward loss ratio               : 0.000000
Average forward throughput               : 0.778 kpps
Cumulative reverse transmit count        : 1562
Cumulative reverse loss                  : 0 packets
Average reverse loss ratio               : 0.000000
Average reverse throughput               : 0.780 kpps

Best two-way channel delay               : 48 usecs
Worst two-way channel delay              : 50 usecs
Average two-way channel delay            : 49 usecs
Best round-trip-time                     : 1856 usecs
Worst round-trip-time                    : 3476 usecs
Average round-trip-time                  : 2445 usecs
Average forward delay variation          : 1 usecs
Average reverse delay variation          : 1 usecs

LDM queries sent                         : 3
LDM responses received                   : 3
LDM queries timedout                     : 0
LDM responses dropped due to errors      : 0
```

monitor mpls loss-delay rsvp count detail

```
user@host> monitor mpls loss-delay rsvp LSP-A count 2 detail

(0)
Response code                           : Success
Forward transmit count                  : 142049
Forward receive count                   : 142049
Reverse transmit count                  : 142167
Reverse receive count                   : 142167
Querier transmit timestamp              : 1404129161 secs, 554422723 nsecs
Responder receive timestamp             : 1404129161 secs, 542877570 nsecs
Responder transmit timestamp            : 1404129161 secs, 546004545 nsecs
Querier receive timestamp               : 1404129161 secs, 557599327 nsecs
(1)
Response code                           : Success
Forward transmit count                  : 143049
Forward receive count                   : 143049
Reverse transmit count                  : 143168
Reverse receive count                   : 143168
Current forward transmit count          : 1000
Current forward receive count           : 1000
Current forward loss                    : 0 packets
Current forward loss ratio              : 0.000000
Current forward throughput              : 1.000 kpps
Current reverse transmit count          : 1001
Current reverse receive count           : 1001
Current reverse loss                    : 0 packets
Current reverse loss ratio              : 0.000000
Current reverse throughput              : 1.001 kpps
Querier transmit timestamp              : 1404129162 secs, 554465742 nsecs
Responder receive timestamp             : 1404129162 secs, 542919166 nsecs
Responder transmit timestamp            : 1404129162 secs, 545812736 nsecs
Querier receive timestamp               : 1404129162 secs, 557409175 nsecs
Current two-way channel delay           : 49 usecs
Current round-trip-time                 : 2943 usecs
(2)
Response code                           : Success
Forward transmit count                  : 143677
Forward receive count                   : 143677
Reverse transmit count                  : 143799
Reverse receive count                   : 143799
```

```
Current forward transmit count          : 628
Current forward receive count           : 628
Current forward loss                    : 0 packets
Current forward loss ratio              : 0.000000
Current forward throughput              : 0.627 kpps
Current reverse transmit count          : 631
Current reverse receive count           : 631
Current reverse loss                    : 0 packets
Current reverse loss ratio              : 0.000000
Current reverse throughput              : 0.630 kpps
Querier transmit timestamp              : 1404129163 secs, 556698575 nsecs
Responder receive timestamp             : 1404129163 secs, 545150128 nsecs
Responder transmit timestamp            : 1404129163 secs, 546918408 nsecs
Querier receive timestamp               : 1404129163 secs, 558515047 nsecs
Current two-way channel delay           : 48 usecs
Current round-trip-time                 : 1816 usecs

Cumulative forward transmit count       : 1628
Cumulative forward loss                 : 0 packets
Average forward loss ratio              : 0.000000
Average forward throughput              : 0.813 kpps
Cumulative reverse transmit count       : 1632
Cumulative reverse loss                 : 0 packets
Average reverse loss ratio              : 0.000000
Average reverse throughput              : 0.815 kpps

Best two-way channel delay              : 48 usecs
Worst two-way channel delay             : 49 usecs
Average two-way channel delay           : 49 usecs
Best round-trip-time                    : 1816 usecs
Worst round-trip-time                   : 3176 usecs
Average round-trip-time                 : 2645 usecs
Average forward delay variation         : 1 usecs
Average reverse delay variation         : 0 usecs

LDM queries sent                        : 3
LDM responses received                  : 3
LDM queries timedout                    : 0
LDM responses dropped due to errors     : 0
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

monitor mpls loss rsvp | 10896

monitor mpls delay rsvp | 10884

Example: Configuring On-Demand Loss and Delay Measurement

monitor mpls loss rsvp

IN THIS SECTION

- Syntax | 10896
- Description | 10897
- Options | 10897
- Required Privilege Level | 10897
- Output Fields | 10898
- Sample Output | 10900
- Release Information | 10903

Syntax

```
monitor mpls loss rsvp lsp-name
<detail>
<bytes>
<count count>
<interval seconds>
<traffic-class traffic-class>
```

Description

Perform an on-demand loss measurement and display the measured values for associated bidirectional MPLS ultimate hop popping (UHP) point-to-point label-switched paths (LSPs).

Options

lsp-name Name of the associated bidirectional MPLS UHP LSP for which the loss measurement is performed.

detail (Optional) Display detailed output of the LSP loss measurement.

bytes (Optional) Specify the measurement quantity for the LSP loss measurement as bytes. By default, LSP loss is measured in packets.

> **NOTE**: The byte count of a packet sent or received over a channel counts only the payload, including the total length of that packet and excluding the headers, labels, and framing of the channel itself.

count *count* (Optional) Specify the number of loss measurements to be carried out for the MPLS UHP LSP. For LSP loss measurement, the number of queries sent is the specified count number plus one additional query, because the LSP loss is measured using successive messages.

- **Default:** 10
- **Range:** 1 through 1000000

interval *seconds* (Optional) Specify in seconds the interval between two successive query messages.

- **Range:** 1 through 255 seconds

traffic-class *traffic-class* (Optional) Specify the traffic class and enable traffic-class-statistics for the LSP loss measurement.

- **Range:** 0 though 7

Required Privilege Level

view

Output Fields

Table 80 on page 10898 describes the output fields for the `monitor mpls loss rsvp` command. Output fields are listed in the approximate order in which they appear.

Table 80: monitor mpls loss rsvp Output Fields

Field Name	Field Description	Level of Output
Current forward loss	Difference between the current forward transmit count and the current forward receive count.	All Levels
Current forward loss ratio	Total packet loss (current forward loss divided by current forward transmit count).	All Levels
Current forward throughput	Current forward transmit count divided by 1000.	All Levels
Current reverse loss	Difference between the current reverse transmit count and the current reverse receive count.	All Levels
Current reverse loss ratio	Total packet loss (current reverse loss divided by current reverse transmit count).	All Levels
Current reverse throughput	Current reverse transmit count divided by 1000.	All Levels
Cumulative forward transmit count	Cumulative forward transmit counter value at the time the loss measurement message was originated.	All Levels
Cumulative forward loss	Cumulative forward loss counter value at the time the loss measurement message was originated.	All Levels
Average forward loss ratio	Average packet loss (current forward loss divided by current forward transmit count).	All Levels
Average forward throughput	Average forward transmit count divided by 1000.	All Levels

Table 80: monitor mpls loss rsvp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Cumulative reverse transmit count	Cumulative reverse transmit counter value at the time the loss measurement message was originated.	All Levels
Cumulative reverse loss	Difference between the cumulative reverse transmit count and the cumulative reverse receive count.	All Levels
Average reverse loss ratio	Average packet loss (average reverse loss divided by average reverse transmit count).	All Levels
Average reverse throughput	Average reverse transmit count divided by 1000.	All Levels
LM queries sent	Number of queries sent for loss measurement.	All Levels
LM responses received	Number of responses received for loss measurement queries.	All Levels
LM queries timedout	Number of timed out queries sent for loss measurement.	All Levels
LM responses dropped due to errors	Number of loss measurement responses dropped due to errors.	All Levels
Response code	Status of the messages used for loss measurement. Response code can be one of the following: • `Success`—Successful response code. • `Failed`—Failed response code.	`detail`
Origin timestamp	Time and date the loss measurement message is originated without any specific format (NTP and PTP).	`detail`
Forward transmit count	Forward transmit counter value at the time the loss measurement message was originated.	`detail`

Table 80: monitor mpls loss rsvp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Forward receive count	Forward receive counter value at the time the loss measurement message was originated.	`detail`
Reverse transmit count	Reverse transmit counter value at the time the loss measurement message was originated.	`detail`
Reverse receive count	Reverse receive counter value at the time the loss measurement message was originated.	`detail`
Current forward transmit count	Difference between the current forward transit count and the previous forward transit count.	`detail`
Current forward receive count	Difference between the current forward receive count and the previous forward receive count.	`detail`
Current reverse transmit count	Difference between the current reverse transit count and the previous reverse transit count.	`detail`
Current reverse receive count	Difference between the current reverse receive count and the previous reverse receive count.	`detail`

Sample Output

monitor mpls lsp loss rsvp count

```
user@host> monitor mpls lsp loss rsvp count 2

(1)
Current forward loss                    : 0 packets
Current forward loss ratio              : 0.000000
Current forward throughput              : 1.006 kpps
Current reverse loss                    : 0 packets
Current reverse loss ratio              : 0.000000
Current reverse throughput              : 1.007 kpps
```

```
(2)
Current forward loss                    : 0 packets
Current forward loss ratio              : 0.000000
Current forward throughput              : 0.559 kpps
Current reverse loss                    : 0 packets
Current reverse loss ratio              : 0.000000
Current reverse throughput              : 0.562 kpps

Cumulative forward transmit count       : 1559
Cumulative forward loss                 : 0 packets
Average forward loss ratio              : 0.000000
Average forward throughput              : 0.782 kpps
Cumulative reverse transmit count       : 1563
Cumulative reverse loss                 : 0 packets
Average reverse loss ratio              : 0.000000
Average reverse throughput              : 0.784 kpps

LM queries sent                         : 3
LM responses received                   : 3
LM queries timedout                     : 0
LM responses dropped due to errors      : 0
```

monitor mpls lsp loss rsvp detail

```
user@host> monitor mpls lsp loss rsvp detail
(0)
Response code                           : Success
Origin timestamp                        : 1404129082 secs, 905571890 nsecs
Forward transmit count                  : 83040
Forward receive count                   : 83040
Reverse transmit count                  : 83100
Reverse receive count                   : 83100
(1)
Response code                           : Success
Origin timestamp                        : 1404129083 secs, 905048410 nsecs
Forward transmit count                  : 83841
Forward receive count                   : 83841
Reverse transmit count                  : 83904
Reverse receive count                   : 83904
Current forward transmit count          : 801
Current forward receive count           : 801
```

```
Current forward loss                    : 0 packets
Current forward loss ratio              : 0.000000
Current forward throughput              : 0.801 kpps
Current reverse transmit count          : 804
Current reverse receive count           : 804
Current reverse loss                    : 0 packets
Current reverse loss ratio              : 0.000000
Current reverse throughput              : 0.804 kpps
(2)
Response code                           : Success
Origin timestamp                        : 1404129084 secs, 904828715 nsecs
Forward transmit count                  : 84423
Forward receive count                   : 84423
Reverse transmit count                  : 84487
Reverse receive count                   : 84487
Current forward transmit count          : 582
Current forward receive count           : 582
Current forward loss                    : 0 packets
Current forward loss ratio              : 0.000000
Current forward throughput              : 0.582 kpps
Current reverse transmit count          : 583
Current reverse receive count           : 583
Current reverse loss                    : 0 packets
Current reverse loss ratio              : 0.000000
Current reverse throughput              : 0.583 kpps

Cumulative forward transmit count       : 1383
Cumulative forward loss                 : 0 packets
Average forward loss ratio              : 0.000000
Average forward throughput              : 0.692 kpps
Cumulative reverse transmit count       : 1387
Cumulative reverse loss                 : 0 packets
Average reverse loss ratio              : 0.000000
Average reverse throughput              : 0.694 kpps

LM queries sent                         : 3
LM responses received                   : 3
LM queries timedout                     : 0
LM responses dropped due to errors      : 0
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

monitor mpls delay rsvp | 10884

monitor mpls loss-delay rsvp | 10890

Example: Configuring On-Demand Loss and Delay Measurement

monitor (Multinode High Availability)

IN THIS SECTION

- Syntax | 10903
- Hierarchy Level | 10904
- Description | 10904
- Options | 10905
- Required Privilege Level | 10905
- Release Information | 10905

Syntax

```
monitor {
    bfd-liveliness destination-ip-address {
        interface interface-name;
        routing-instance routing-instance-name;
        session-type (multihop | singlehop);
        src-ip src-ip;
    }
    interface name;
    ip destination-ip-address {
```

```
            routing-instance routing-instance-name;
        }
        monitor-object {
           bfd-liveliness {
              destination-ip name {
               interface interface;
               routing-instance routing-instance;
               session-type(multihop | singlehop);
               src-ip src-ip;
               weight weight;
              }
             threshold threshold;
             }
          interface {
              interface-name name {
                 weight weight;
              }
              threshold threshold;
             }
          ip {
              destination-ip name {
              routing-instance routing-instance;
               weight weight;
              }
          threshold threshold;
        }
          object-threshold object-threshold;
            }
        srg-threshold;
        }
```

Hierarchy Level

```
[edit chassis high-availability services-redundancy-group id-number]
```

Description

Configure monitoring options for the Multinode High Availability setup.

Options

bfd-liveliness

Configure Bidirectional Forwarding Detection (BFD) monitoring to detect failures in a network. You can configure Multinode High Availability to monitor one or more links using BFD. This configuration triggers a failover in the event of BFD failure. Configure BFD liveliness by specifying source and destination IP and the interface where the peer device is directly connected to.

- interface—Name of the interface for single-hop sessions
- routing-instance—Routing instance to locate the route
- session-type —(Optional) Configure BFD session mode to be single-hop or multihop. By default, BGP uses single-hop BFD sessions if the peer is directly connected to the router's interface.
- src-ip-—Source IPv4 or IPv6 address for activeness probe.

"interface" on page 9025 ***interface-name***

Configure interface monitoring for SRG. The node which detects the interface monitoring failure transitions to ineligible state for the corresponding SRG and the other node (if healthy) takes over the active role or that SRG and the subsequent GARP ensures traffic switching and recovery.

ip ***destination-ip-address***

Configure IP monitoring. You can configure IP monitoring by specifying destination IP address (of upstream router) for probing.

"monitor-object" on page 10906 ***object name***

SRG monitoring object. You can define one or more monitor objects (up to 10) to group monitor-specific attributes associated with an SRG.

srg-threshold ***object name***

SRG monitoring object.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Example: Configure Multinode High Availability in a Layer 3 Network

monitor-object

IN THIS SECTION

- Syntax | 10906
- Hierarchy Level | 10907
- Description | 10907
- Options | 10907
- Required Privilege Level | 10908
- Release Information | 10908

Syntax

```
monitor-object name {
 bfd-liveliness {
  destination-ip name {
   interface interface;
   routing-instance routing-instance;
   session-type(multihop | singlehop);
   src-ip src-ip;
   weight weight;
  }
  threshold threshold;
 }
 interface {
  interface-name name {
   weight weight;
  }
  threshold threshold;
 }
 ip {
  destination-ip name {
   routing-instance routing-instance;
   weight weight;
  }
  threshold threshold;
```

```
    }
    object-threshold object-threshold;
   }
```

Hierarchy Level

```
[edit chassis high-availability services-redundancy-group name monitor]
```

Description

SRG monitoring object

Options

name SRG monitoring object name.

- Range - 1 through 32.

bfd-liveliness Configure BFD monitoring for the specific monitoring object.

interface Interface to monitor for the specific monitoring object.

- interface-name *interface-name*—Destination IP address to monitor.
  - weight—Weight assigned to the interface. A higher weight value indicates a greater importance.
- threshold—Threshould value assigned to the interface that influences failover.

ip *ip-address* IP address to monitor for the specific monitoring object.

- destination-ip *destination-ip-address*—Destination IP address to monitor.
  - weight—Weight assigned to the interface or monitored IP address. A higher weight value indicates a greater importance.
  - routing-instance—Routing instance to locate the destination route.
- threshold—Threshould value assigned to the IP address that influences failover.

object-threshold SRG monitoring object threshold value. This is the threshold value assigned to the monitoring-object that influences SRG failover.

- **Range:** 1 through 1000

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

monitor pfe traffic interface

IN THIS SECTION

- Syntax | 10908
- Description | 10909
- Options | 10909
- Required Privilege Level | 10910
- Output Fields | 10910
- Sample Output | 10911
- Release Information | 10912

Syntax

```
monitor pfe traffic interface interface-name
<absolute-sequence>
<brief | detail | extensive>
<byte-offset>
<count count>
<egress>
```

```
<egress-interface interface-name>
<ingress>
<inner-dmac | inner-dst-ipv4 | inner-smac | inner-src-ipv4>
<layer2-headers>
<no-domain-names>
<no-promiscuous>
<no-resolve>
<no-timestamp>
<outer-dmac | outer-dst-ipv4 | outer-smac | outer-src-ipv4>
<print-ascii>
<print-hex>
<read-file filename>
<resolve-timeout>
<size size>
<vlan-id>
<vnid>
<write-file filename>
```

Description

Display packet headers or packets that are received and sent from the packet forwarding engine for monitoring and verifying the functionality of on-box packet sniffing capabilities.

Options

absolute-sequence	Display absolute TCP sequence numbers.
brief | detail | extensive	Display the specified level of output.
byte-offset	Display the specified level of output.
count *count*	Specify the number of packet headers to display (0 through 1,000,000).
egress	Egress data traffic monitoring.
egress-interface	Specify the egress interface where loopback is enabled.
ingress	Ingress data traffic monitoring.
inner-dmac | inner-dst-ipv4 | inner-smac | inner-src-ipv4	Display the inner header addresses to monitor traffic and packet headers.

layer2-headers	Display the link-level header on each line.
no-domain-names	Suppress the display of the domain portion of hostnames.
no-promiscuous	Do not put the interface into promiscuous mode.
no-resolve	Do not attempt to convert host addresses to names.
no-timestamp	Suppress timestamps on displayed packets.
outer-dmac | outer-dst-ipv4 | outer-smac | outer-src-ipv4	Display the outer header addresses to monitor traffic and packet headers.
print-ascii	Display each packet in ASCII format.
print-hex	Display each packet in hexadecimal format.
resolve-timeout ***timeout***	Amount of time the router or switch waits for each reverse lookup before timing out. You can set the timeout for 1 through 4,294,967,295 seconds.
size ***size***	Snap length of the frame to be captured, in bytes.
vlan-id	VLAN identifier to monitor traffic.
vnid	VXLAN identifier to monitor traffic.
write-file ***filename***	Write packets to the file specified.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To view this statement in the configuration.

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

monitor pfe traffic interface *interface-name* ingress outer-dmac *destination MAC address*

```
user@host> monitor pfe traffic interface xe-0/0/0:0 ingress outer-dmac 00:00:00:00:00:0C
verbose output suppressed, use  or  for full protocol decode
Address resolution is ON. Use  to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on xe-0/0/0:0, capture size 96 bytes

Reverse lookup for 192.85.1.2 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use  to avoid reverse lookups on IP addresses.

10:19:07.579828  In IP truncated-ip - 50 bytes missing! 192.85.1.2 > 192.85.1.3:  ip-proto-253 86
10:19:07.579833  In IP truncated-ip - 50 bytes missing! 192.85.1.2 > 192.85.1.3:  ip-proto-253 86
```

monitor pfe traffic interface *interface-name* egress egress-interface *interface-name* vnid 1000 byte-offset 12

```
user@host> monitor pfe traffic interface xe-0/0/0:0 egress egress-interface xe-0/0/30:0 vnid
1000 byte-offset 12

verbose output suppressed, use  or  for full protocol decode
Address resolution is ON. Use  to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on xe-0/0/30:0, capture size 96 bytes

Reverse lookup for 204.1.1.1 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use  to avoid reverse lookups on IP addresses.

14:30:36.718089  In IP truncated-ip - 110 bytes missing! 204.1.1.1.36912 > 203.1.1.1.vxlan:
VXLAN(VNI:1000 Flags:8000000)
14:30:36.720244  In IP truncated-ip - 110 bytes missing! 204.1.1.1.36912 > 203.1.1.1.vxlan:
VXLAN(VNI:1000 Flags:8000000)
```

Release Information

Command introduced before Junos OS Release 24.4R1.

monitor-profile

IN THIS SECTION

- Syntax | 10912
- Hierarchy Level | 10912
- Description | 10913
- Options | 10913
- Required Privilege Level | 10914
- Release Information | 10914

Syntax

```
monitor-profile monitor-profile-name {
    lower-threshold lower-threshold;
    notification-type(alarm | none | syslog);
    upper-threshold upper-threshold;
}
```

Hierarchy Level

```
[edit system packet-forwarding-options hw-resource-monitor]
```

Description

Map a monitor profile to a configured resource list (optional) to monitor hardware resource utilization and set thresholds to trigger notifications as well as the type of notification you would like to receive when a resource breach occurs. Monitoring resource utilization for a hardware resource as it nears or crosses the maximum capacity threshold that you set can enable you to take appropriate and timely actions to prevent unintended consequences in your network, such as traffic loss.

You can use the same monitor profile multiple resource lists, but a single resource list cannot be mapped to multiple monitor profiles. A resource list can only be mapped only to a single monitor-profile. Configure a monitor profile prior to mapping it to a resource list.

To associate a monitor profile with the resource list, in configuration mode, go to the `[edit system packet-forwarding-options hw-resource-monitor resource-list]` hierarchy level and map a monitor profile:

```
[edit]
user@host# edit system packet-forwarding-options hw-resource-monitor resource-list <resource-list-name> monitor-profile <monitor-profile-name>
```

Options

monitor-profile-name The hardware resource monitor profile name.

lower-threshold The lower threshold, by percent. A lower threshold breach raises a minor notification

- **Range:** 1 through 100 percent
- **Default:** 50 percent

notification-type The notification type for threshold breaches.

- alarm—System sends an alarm for threshold breaches.
- none—No notification issued for threshold breaches.
- syslog—System sends a syslog for threshold breaches.
- **Default:** none

upper-threshold The upper threshold, by percent. An upper threshold breach raises a major notification

- **Range:** 1 through 100 percent
- **Default:** 90 percent

Required Privilege Level

Edit

Release Information

Statement introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

resource-list (Network Management) | 15610
polling-interval (Network Management) | 12752
monitor-profile | 10912
No Link Title

monitor security flow file

IN THIS SECTION

- Syntax | 10915
- Description | 10915
- Options | 10915
- Required Privilege Level | 10916
- Output Fields | 10916
- Release Information | 10916

Syntax

```
monitor security flow file
<file-name>
<files number>
<match regular-expression>
<size maximum-file-size>
<(world-readable | no-world-readable)>
```

Description

Configure options for the security flow monitoring output.

Options

filename

Name of the file to receive the output of the monitoring operation. All output is saved in the `/var/log/` directory.

files *number*

Maximum number of output files. If you specify a maximum number of files, you must also specify a maximum file size with the `size` option.

Range: 2 through 1000 files

Default: 10 files

match *regular-expression*

Refine the output to include lines that contain the regular expression.

size *maximum-file-size*

Maximum size of each output file. When an output file named *`output`* reaches this size, it is renamed *`output`*`.0`. When the output file again reaches its maximum size, *`output`*`.0` is renamed *`output`*`.1` and `output` file is renamed *`output`*`.0`. This renaming scheme continues until the maximum number of output files is reached. Then the oldest output file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of output files with the `files` option.

Range: 10 KB through 1 GB

Default: 128 KB

(world-readable | no-world-readable) By default, the output files can be accessed only by the user who configures the monitoring operation. The `world-readable` option enables all users to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

Required Privilege Level

trace

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Monitoring Security Flow Sessions Overview

monitor security flow filter | **10916**

monitor security flow start | **10918**

show monitor security flow | **22122**

monitor security flow filter

IN THIS SECTION

- Syntax | **10917**
- Description | **10917**
- Options | **10917**
- Required Privilege Level | **10918**
- Release Information | **10918**

Syntax

```
monitor security flow filter filter-name
<conn-tag session-connection-tag>
<destination-port (port-range | protocol-name)>
<destination-prefix destination-prefix>
<interface interface-name>
<logical-system logical-system-name>
<protocol (protocol name | protocol number)>
<root-logical-system>
<source-port  (port-range | protocol-name)>
<source-prefix source-prefix>
```

Description

Set security flow filters to define flow sessions that you want to monitor. A maximum of 64 filters is supported at a time.

Defining the filters themselves does not trigger monitoring. You must explicitly use the `monitor security flow start` command to enable monitoring. Once monitoring starts, any traffic that matches the specified filters is saved in an output file in the `/var/log/` directory.

> **NOTE**: Unlike filters defined in the configuration mode, these filters defined using operational mode commands are cleared when you reboot your system. They are used expressly for debugging purposes.

Options

filter *filter-name*	Specify a name for the filter. The filter name can contain letters, numbers, underscores (_) and hyphens (-) and can be up to 64 characters long.
conn-tag	Specify the session connection tag. The session connection tag uniquely identifies a session.
destination-port (*port-range* | *protocol-name*)	Specify the TCP or UDP destination port to match. You can also specify a range of TCP or UDP destination ports and monitor all traffic in this group.
destination-prefix *destination-prefix*	Specify the destination IPv4 or IPv6 address prefix to match.

interface *interface-name*	Specify the logical interface name to match.
logical-system *logical-system-name*	Specify the logical system name to match.
protocol (*protocol name* | *protocol number*)	Specify the IP protocol type to match.
root-logical-system	(Default) Specify the root logical system to match.
source-port (*port-range* | *protocol-name*)	Specify the TCP or UDP source port to match. You can also specify a range of TCP or UDP source ports and monitor all traffic in this group.
source-prefix *source-prefix*	Specify the source IP address prefix to match.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 12.1X46-D10. The was updated in Junos OS Release 15.1X49-D70 with the addition of the conn-tag filter parameter.

RELATED DOCUMENTATION

Monitoring Security Flow Sessions Overview

monitor security flow file | **10914**

monitor security flow start | **10918**

monitor security flow stop | **10920**

monitor security flow start

IN THIS SECTION

- Syntax | **10919**
- Description | **10919**

- Options | 10919
- Required Privilege Level | 10919
- Output Fields | 10919
- Release Information | 10919

Syntax

```
monitor security flow start
```

Description

Start the monitoring of security flow session. Once monitoring starts, any traffic that matches the specified filters is saved in an output file in the `var/log/` directory. At least one filter must be defined for the monitoring to start.

Use the `monitor security flow stop` command to stop the monitoring of flow sessions.

Options

This command has no options.

Required Privilege Level

trace

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Monitoring Security Flow Sessions Overview

show monitor security flow | 22122

monitor security flow filter | 10916

monitor security flow stop | 10920

monitor security flow stop

IN THIS SECTION

- Syntax | 10920
- Description | 10920
- Options | 10920
- Required Privilege Level | 10921
- Output Fields | 10921
- Release Information | 10921

Syntax

```
monitor security flow stop
```

Description

Stop monitoring the security flow session. Use the `monitor security flow start` command to start the monitoring of flow sessions.

Options

This command has no options.

Required Privilege Level

trace

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Monitoring Security Flow Sessions Overview

monitor security flow start | 10918

monitor security packet-drop

IN THIS SECTION

- Syntax | 10921
- Description | 10922
- Options | 10922
- Required Privilege Level | 10922
- Sample Output | 10923
- Release Information | 10926

Syntax

```
monitor security packet-drop
<source-prefix>
<destination-prefix>
```

```
<source-port>
<destination-port>
<ingress-interface-name>
<protocol>
<logical-system-name>
<count-number>
```

Description

Displays the packet-drop information without committing the configuration, which allows you to trace and monitor the traffic flow. This command output is displayed on the screen until you press **Ctrl+c** or until the security device collects the requested number of packet drops. The command includes various filters to generate the output fields per your requirement.

You can save the packet pocket records to a file using the `monitor security packet-drop | append /var/log/filename.log` command.

Options

source-prefix Display the packet drop information for the given source IP or source prefix address.

destination-prefix Display the packet drop information for the given destination IP or source prefix address.

source-port Display the packet drop information for the given source port.

destination-port Display the packet drop information for the given destination port.

ingress-interface-name Display the packet drop information for the given ingress interface address.

protocol Display the packet drop information for the given protocol number.

logical-system-name Display the packet drop information for the logical system name.

count-number Display the packet drop information for the given count.

Range: 1 thru 8000

Default: 50

Required Privilege Level

view

Sample Output

monitor security packet-drop

When configuration is set to default IDP security policy, using the command `set security idp idp-policy IDP_Default rulebase-ips rule 1 then action drop-packet`, the following type of output is displayed for **monitor security packet-drop** command.

```
user@host> monitor security packet-drop
```

```
Starting packet drop:
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by IDP:IDP
drop SLL packet
```

When configuration is set to drop connection action, the following type of output is displayed for **monitor security packet-drop** command.

```
user@host> monitor security packet-drop
```

```
Starting packet drop:
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by
IDP:Malformed IPV6 header in IPv4 tunnel
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by IDP:Bad
TCP headers
```

Output for packet drops independent of IPS policies.

```
user@host> monitor security packet-drop
```

```
Starting packet drop:
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by IDP:IDP
drop SLL packet
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by
IDP:Malformed IPV6 header in IPv4 tunnel
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by IDP:Bad
TCP headers
```

```
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by
IDP:Overflow drops
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by
IDP:Sequence number wrap around errors
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by IDP:IDP
Policy Initn failed
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by IDP:Flow
Rejected
```

Output for packet drops dependent on IPS policies.

```
user@host> monitor security packet-drop
```

```
Starting packet drop:
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by IDP:IDP
Rule Action Drop Packet ICMP:INFO:ECHO-REQUEST
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by IDP:IDP
Rule Action Close-Client-And-Server ICMP:INFO:ECHO-REQUEST
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by IDP:IDP
Rule Action Close-Server ICMP:INFO:ECHO-REQUEST
14:46:45.511471:LSYS-ID-00 4.0.0.1/19895-->5.0.0.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by IDP:IDP
Rule Action Close-Client ICMP:INFO:ECHO-REQUEST
```

monitor security packet-drop with filters

```
user@host> monitor security packet-drop source-prefix 192.0.2.1 destination-prefix 192.151.100.1
proto icmp
```

```
14:46:45.511471:LSYS-ID-00 192.0.2.1/19895-->192.151.100.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by
IDP:IDP Rule Action Close-Server ICMP:INFO:ECHO-REQUEST
14:46:45.511471:LSYS-ID-00 192.0.2.1/19895-->192.151.100.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by
IDP:IDP Rule Action Close-Client ICMP:INFO:ECHO-REQUEST
```

monitor security packet-drop count 2

```
user@host> monitor security packet-drop count 2
```

```
14:46:45.511471:LSYS-ID-00 192.0.2.1/19895-->192.151.100.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by
IDP:IDP Rule Action Close-Server ICMP:INFO:ECHO-REQUEST
14:46:45.511471:LSYS-ID-00 192.0.2.1/19895-->192.151.100.1/1;icmp,ipid-0,ge-0/0/0.0,Dropped by
IDP:IDP Rule Action Close-Server ICMP:INFO:ECHO-REQUEST
```

monitor security packet-drop | append /var/log/abcd.log

```
user@host> monitor security packet-drop | append /var/log/abcd.log
```

```
^C[abort]
Wrote 7 lines of output to '/var/log/abcd.log'
```

show log abcd.log

```
user@host> show log abcd.log
```

```
Starting packet drop:
07:35:36.742809:LSYS-ID-00 192.0.2.1/2198-->192.151.100.1/1;icmp,ipid-16088,ge-0/0/2.0,Dropped
by POLICY:Denied by Policy: default-policy-logical-system-00
07:35:37.640858:LSYS-ID-00 192.0.2.1/2198-->192.151.100.1/2;icmp,ipid-52440,ge-0/0/2.0,Dropped
by POLICY:Denied by Policy: default-policy-logical-system-00
07:35:38.665155:LSYS-ID-00 192.0.2.1/2198-->192.151.100.1/3;icmp,ipid-28633,ge-0/0/2.0,Dropped
by POLICY:Denied by Policy: default-policy-logical-system-00
07:35:39.689185:LSYS-ID-00 192.0.2.1/2198-->192.151.100.1/4;icmp,ipid-47577,ge-0/0/2.0,Dropped
by POLICY:Denied by Policy: default-policy-logical-system-00
07:35:40.712870:LSYS-ID-00 192.0.2.1/2198-->192.151.100.1/5;icmp,ipid-44762,ge-0/0/2.0,Dropped
by POLICY:Denied by Policy: default-policy-logical-system-00
07:35:41.797742:LSYS-ID-00 192.0.2.1/2198-->192.151.100.1/6;icmp,ipid-16859,ge-0/0/2.0,Dropped
by POLICY:Denied by Policy: default-policy-logical-system-00
```

Release Information

Command introduced in Junos OS Release 21.1R1.

Signature added to packet drop reason in Junos OS Release 21.2R2.

monitor-session

IN THIS SECTION

- Syntax | 10926
- Hierarchy Level | 10926
- Description | 10926
- Default | 10927
- Options | 10927
- Required Privilege Level | 10927
- Release Information | 10927

Syntax

```
monitor-session (interface-name | all);
```

Hierarchy Level

```
[edit protocols ppp]
```

Description

Monitor PPP packet exchanges. When monitoring is enabled, packets exchanged during a session are logged to the default log of `/var/log/pppd`.

Default

If you do not include this statement, no PPPD-specific monitoring operations are performed.

Options

`all`—Monitor PPP packet exchanges on all sessions.

`interface-name`—Logical interface name on which to enable session monitoring.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configuring PPP

monitor start

IN THIS SECTION

- Syntax | 10928
- Description | 10928
- Options | 10928
- Additional Information | 10928
- Required Privilege Level | 10928
- Output Fields | 10928
- Sample Output | 10929

Release Information | 10929

Syntax

```
monitor start filename
```

Description

Start displaying the system log or trace file and additional entries being added to those files.

Options

filename	Specific log or trace file.

Additional Information

Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are configured with the `syslog` statement at the `[edit system]` hierarchy level and the `options` statement at the `[edit routing-options]` hierarchy level. The trace files generated by the routing protocol process are configured with `traceoptions` statements at the `[edit routing-options]`, `[edit interfaces]`, and `[edit protocols protocol]` hierarchy levels.

NOTE: To monitor a log file within a logical system, issue the `monitor start logical-system-name/filename` command.

Required Privilege Level

trace

Output Fields

Table 81 on page 10929 describes the output fields for the `monitor start` command. Output fields are listed in the approximate order in which they appear.

Table 81: monitor start Output Fields

Field Name	Field Description
****filename ****	Name of the file from which entries are being displayed. This line is displayed initially and when the command switches between log files.
Date and time	Timestamp for the log entry.

Sample Output

monitor start

```
user@host> monitor start system-log
*** system-log***
Jul 20 15:07:34 hang sshd[5845]: log: Generating 768 bit RSA key.
Jul 20 15:07:35 hang sshd[5845]: log: RSA key generation complete.
Jul 20 15:07:35 hang sshd[5845]: log: Connection from 204.69.248.180 port 912
Jul 20 15:07:37 hang sshd[5845]: log: RSA authentication for root accepted.
Jul 20 15:07:37 hang sshd[5845]: log: ROOT LOGIN as 'root' from host.example.com
Jul 20 15:07:37 hang sshd[5845]: log: Closing connection to 204.69.248.180
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

monitor list

monitor stop | **10932**

monitor start (JDM)

IN THIS SECTION

- Syntax | 10930
- Description | 10930
- Additional Information | 10930
- Required Privilege Level | 10930
- Output Fields | 10930
- Sample Output | 10931
- Release Information | 10932

Syntax

```
monitor start file-name
```

Description

Start displaying the system log or trace file and additional entries being added to those files.

Additional Information

Log files are generated by the routing protocol process or by system logging.

Required Privilege Level

View

Output Fields

Table 82 on page 10931 describes the output fields for the `monitor start` command. Output fields are listed in the approximate order in which they appear.

Table 82: monitor start Output Fields

Field Name	Field Description
***filename ***	Name of the file from which entries are being displayed.
Date and time	Timestamp for the log entry.

Sample Output

monitor start

```
user@jdm> monitor start syslog

*** syslog ***
Oct 19 19:44:21 jdm mgd[4069]: UI_CMDLINE_READ_LINE: User 'root', command
'monit                                                                    or
start syslog '
Oct 19 19:44:29 jdm mgd[3268]: UI_CMDLINE_READ_LINE: User 'root', command
'delet                                                                    e
groups server0 server interfaces cb1 '
Oct 19 19:44:29 jdm mgd[3268]: UI_CFG_AUDIT_OTHER: User 'root' delete: [groups
s                                                                     erver0
server interfaces cb1]
Oct 19 19:44:35 jdm mgd[3268]: UI_CMDLINE_READ_LINE: User 'root', command
'delet                                                                    e
groups server1 server interfaces cb1 '
Oct 19 19:44:35 jdm mgd[3268]: UI_CFG_AUDIT_OTHER: User 'root' delete: [groups
s                                                                     erver1
server interfaces cb1]
Oct 19 19:44:36 jdm mgd[3268]: UI_CMDLINE_READ_LINE: User 'root', command
'commi                                                                    t '
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT: User 'root' requested 'commit'
operati
on (comment: none)
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT_PROGRESS: Commit operation in
progress:
 Obtaining lock for commit
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT_PROGRESS: Commit operation in
```

```
progress:
 updating commit revision
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT_PROGRESS: Commit operation in
progress:
 obtaining db lock on  server1
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT_PROGRESS: Commit operation in
progress:
 UI extensions feature is not configured
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT_PROGRESS: Commit operation in
progress:
 Started running translation script
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

monitor stop (JDM) | 10934

monitor stop

IN THIS SECTION

- Syntax | 10933
- Description | 10933
- Options | 10933
- Additional Information | 10933
- Required Privilege Level | 10933
- Output Fields | 10933
- Sample Output | 10933
- Release Information | 10933

Syntax

```
monitor stop filename
```

Description

Stop displaying the system log or trace file.

Options

filename	Specific log or trace file.

Additional Information

Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are those configured with the `syslog` statement at the `[edit system]` hierarchy level and the `options` statement at the `[edit routing-options]` hierarchy level. The trace files generated by the routing protocol process are those configured with `traceoptions` statements at the `[edit routing-options]`, `[edit interfaces]`, and `[edit protocols protocol]` hierarchy levels.

Required Privilege Level

trace

Output Fields

This command produces no output.

Sample Output

monitor stop

```
user@host> monitor stop
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

monitor list

monitor start | 10927

monitor stop (JDM)

IN THIS SECTION

- Syntax | 10934
- Description | 10934
- Additional Information | 10934
- Required Privilege Level | 10934
- Output Fields | 10934
- Release Information | 10935

Syntax

```
monitor stop file-name
```

Description

Stop displaying the system log or trace file.

Additional Information

Log files are generated by the routing protocol process or by system logging.

Required Privilege Level

View

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

monitor start (JDM) | 10930

monitor traffic

IN THIS SECTION

- Syntax | 10935
- Description | 10936
- Options | 10936
- Additional Information | 10938
- Required Privilege Level | 10944
- Output Fields | 10944
- Sample Output | 10944
- Release Information | 10952

Syntax

```
monitor traffic
<brief | detail | extensive>
<absolute-sequence>
<count count>
<interface interface-name>
<layer2-headers>
<matching matching>
<no-domain-names>
<no-promiscuous>
<no-resolve>
```

```
<no-timestamp>
<print-ascii>
<print-hex>
<read-file filename>
<resolve-timeout>
<size size>
<write-file filename>
```

Description

Display packet headers or packets received and sent from the Routing Engine.

NOTE:

- Using the `monitor-traffic` command can degrade router or switch performance.
- Delays from DNS resolution can be eliminated by using the `no-resolve` option.

NOTE: This command is not supported on the QFabric system.

Options

none	(Optional) Display packet headers transmitted through fxp0. On a TX Matrix Plus router, display packet headers transmitted through em0.
brief | detail | extensive	(Optional) Display the specified level of output.
absolute-sequence	(Optional) Display absolute TCP sequence numbers.
count *count*	(Optional) Specify the number of packet headers to display (0 through 1,000,000). The monitor traffic command quits automatically after displaying the number of packets specified.
interface *interface-name*	(Optional) Specify the interface on which the `monitor traffic` command displays packet data. If no interface is specified, the `monitor traffic` command displays packet data arriving on the lowest-numbered interface. In Junos OS Evolved:

- If you modify an interface that you are monitoring with the `monitor traffic interface` command, the monitoring session ends with the message: pcap_loop: read: Device not configured. To continue monitoring the interface, rerun the `monitor traffic interface` command. However, if the monitored interface is removed, the command session continues, but there will be no packets or errors reported.
- When you use the command `monitor traffic interface` *`interface-name`* on a logical interface, the output displays all packets received or transmitted on that interface, including Layer 2 traffic. When you use this command on a physical interface, the output only displays packets received and transmitted on the physical interface and does not include traffic from the logical interface.

layer2-headers (Optional) Display the link-level header on each line.

matching *matching* (Optional) Display packet headers that match a regular expression. Use matching expressions to define the level of detail with which the `monitor traffic` command filters and displays packet data.

no-domain-names (Optional) Suppress the display of the domain portion of hostnames. With the `no-domain-names` option enabled, the `monitor traffic` command displays only `team` for the hostname **team.company.net**.

no-promiscuous (Optional) Do not put the interface into promiscuous mode.

no-resolve (Optional) Suppress reverse lookup of the IP addresses.

no-timestamp (Optional) Suppress timestamps on displayed packets.

print-ascii (Optional) Display each packet in ASCII format.

print-hex (Optional) Display each packet, except the link-level header, in hexadecimal format.

read-file *filename* Read packets from the file specified.

resolve-timeout *timeout* (Optional) Amount of time the router or switch waits for each reverse lookup before timing out. You can set the timeout for 1 through 4,294,967,295 seconds. The default is 4 seconds. To display each packet, use the `print-ascii`, `print-hex`, or `extensive` option.

size *size* (Optional) Read but do not display up to the specified number of bytes for each packet. When set to `brief` output, the default packet size is 96 bytes and

is adequate for capturing IP, ICMP, UDP, and TCP packet data. When set to `detail` and `extensive` output, the default packet size is 1514. The `monitor traffic` command truncates displayed packets if the matched data exceeds the configured size.

write-file *filename* Write packets to the file specified.

> **NOTE**: Starting in Junos OS Evolved 20.4R1, the `write-file` option at the `monitor traffic interface` hierarchy level takes precedence over the `extensive` option when you configure them simultaneously. If you try to configure these options at the same time, Junos OS Evolved gives you a warning message that the options are not compatible, and it only runs the `monitor traffic interface write-file` command.

Additional Information

In the `monitor traffic` command, you can specify an expression to match by using the `matching` option and including the expression in quotation marks:

```
monitor traffic matching "expression"
```

Replace *expression* with one or more of the match conditions listed in Table 83 on page 10938.

Table 83: Match Conditions for the monitor traffic Command

Match Type	Condition	Description
Entity	`host [address | hostname]`	Matches packets that contain the specified address or hostname. The protocol match conditions `arp`, `ip`, or `rarp`, or any of the directional match conditions can be prepended to the `host` match condition.
	`net address`	Matches packets with source or destination addresses containing the specified network address.

Table 83: Match Conditions for the monitor traffic Command *(Continued)*

Match Type	Condition	Description
	`net` *`address`* `mask` *`mask`*	Matches packets containing the specified network address and subnet mask.
	`port` (*`port-number`* | *`port-name`*)	Matches packets containing the specified source or destination TCP or UDP port number or port name. In place of the numeric port address, you can specify a text synonym, such as `bgp` (179), `dhcp` (67), or `domain` (53) (the port numbers are also listed).
Directional	`dst`	Matches packets going to the specified destination. This match condition can be prepended to any of the entity type match conditions.
	`src`	Matches packets from a specified source. This match condition can be prepended to any of the entity type match conditions.
	`src` and `dst`	Matches packets that contain the specified source and destination addresses. This match condition can be prepended to any of the entity type match conditions.
	`src` or `dst`	Matches packets containing either of the specified addresses. This match condition can be prepended to any of the entity type match conditions.
Packet Length	`less` *`value`*	Matches packets shorter than or equal to the specified value, in bytes.

Table 83: Match Conditions for the monitor traffic Command *(Continued)*

Match Type	Condition	Description
	`greater` *`value`*	Matches packets longer than or equal to the specified value, in bytes.
Protocol	`amt`	Matches all AMT packets. Use the extensive level of output to decode the inner IGMP packets in addition to the AMT outer packet.
	`arp`	Matches all ARP packets.
	`ether`	Matches all Ethernet packets.
	`ether (broadcast | multicast)`	Matches broadcast or multicast Ethernet frames. This match condition can be prepended with`src` and `dst`.
	`ether` *`protocol`* `(`*`address`* `| (arp | ip | rarp))`	Matches packets with the specified Ethernet address or Ethernet packets of the specified protocol type. The `ether protocol` arguments `arp`, `ip`, and `rarp` are also independent match conditions, so they must be preceded by a backslash (\) when used in the `ether protocol` match condition.
	`icmp`	Matches all ICMP packets.
	`ip`	Matches all IP packets.
	`ip (broadcast | multicast)`	Matches broadcast or multicast IP packets.

Table 83: Match Conditions for the monitor traffic Command *(Continued)*

Match Type	Condition	Description
	`ip protocol (`*`address`*` | (icmp | igrp | tcp | udp))`	Matches packets with the specified address or protocol type. The `ip protocol` arguments `icmp`, `tcp`, and `udp` are also independent match conditions, so they must be preceded by a backslash (\) when used in the `ip protocol` match condition.
	`isis`	Matches all IS-IS routing messages.
	`proto` *`ip-protocol-number`*	Matches packets whose headers contain the specified IP protocol number.
	`rarp`	Matches all RARP packets.
	`tcp`	Matches all TCP datagrams.
	`udp`	Matches all UDP datagrams.

To combine expressions, use the logical operators listed in Table 84 on page 10941.

Table 84: Logical Operators for the monitor traffic Command

Logical Operator (Highest to Lowest Precedence)	Description
!	Logical NOT. If the first condition does not match, the next condition is evaluated.
&&	Logical AND. If the first condition matches, the next condition is evaluated. If the first condition does not match, the next condition is skipped.

Table 84: Logical Operators for the monitor traffic Command *(Continued)*

Logical Operator (Highest to Lowest Precedence)	Description
`||`	Logical OR. If the first condition matches, the next condition is skipped. If the first condition does not match, the next condition is evaluated.
`( )`	Group operators to override default precedence order. Parentheses are special characters, each of which must be preceded by a backslash (\).

You can use relational operators to compare arithmetic expressions composed of integer constants, binary operators, a length operator, and special packet data accessors. The arithmetic expression matching condition uses the following syntax:

```
monitor traffic matching "ether[0] & 1 != 0" "arithmetic_expression relational_operator
arithmetic_expression"
```

The packet data accessor uses the following syntax:

```
protocol [byte-offset <size>]
```

The optional *size* field represents the number of bytes examined in the packet header. The available values are`1`, `2`, or `4` bytes. The following sample command captures all multicast traffic:

```
user@host> monitor traffic matching "ether[0] & 1 != 0"
```

To specify match conditions that have a numeric value, use the arithmetic and relational operators listed in Table 85 on page 10943.

NOTE: Because the Packet Forwarding Engine removes Layer 2 header information before sending packets to the Routing Engine:

- The `monitor traffic` command cannot apply match conditions to inbound traffic.

- The `monitor traffic interface` command also cannot apply match conditions for Layer 3 and Layer 4 packet data, resulting in the `match` pipe option (`| match`) for this command for Layer 3 and Layer 4 packets not working either. Therefore, ensure that you specify match conditions as described in this command summary. For more information about match conditions, see Table 83 on page 10938.

- The 802.1Q VLAN tag information included in the Layer 2 header is removed from all inbound traffic packets. Because the `monitor traffic interface ae[x]` command for aggregated Ethernet interfaces (such as) only shows inbound traffic data, the command does not show VLAN tag information in the output.

Table 85: Arithmetic and Relational Operators for the monitor traffic Command

Arithmetic or Relational Operator	Description
Arithmetic Operator	
+	Addition operator.
-	Subtraction operator.
/	Division operator.
&	Bitwise AND.
*	Bitwise exclusive OR.
|	Bitwise inclusive OR.
Relational Operator (Highest to Lowest Precedence)	
<=	If the first expression is less than or equal to the second, the packet matches.
>=	If the first expression is greater than or equal to the second, the packet matches.

Table 85: Arithmetic and Relational Operators for the monitor traffic Command *(Continued)*

Arithmetic or Relational Operator	Description
<	If the first expression is less than the second, the packet matches.
>	If the first expression is greater than the second, the packet matches.
=	If the compared expressions are equal, the packet matches.
!=	If the compared expressions are unequal, the packet matches.

Required Privilege Level

trace

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

monitor traffic count

```
user@host> monitor traffic count 2
listening on fxp0
04:35:49.814125  In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529478 win 16798 (DF)
04:35:49.814185
Out my-server.work.net.telnet > my-server.home.net.1295: P
1:38(37) ack 0 win 17680 (DF) [tos 0x10]
```

monitor traffic detail count

```
user@host> monitor traffic detail count 2
listening on fxp0
04:38:16.265864  In my-server.home.net.1295 > my-server.work.net.telnet: . ack 4122529971 win
17678 (DF) (ttl 121, id 6812)
04:38:16.265926
Out my-server.work.net.telnet.telnet > my-server.home.net.1295: P 1:38(37) ack 0 win 17680 (DF)
[tos 0x10]  (ttl 6)
```

monitor traffic extensive (Absolute Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20 matching
"tcp" absolute-sequence
listening on fxp0
 In 203.0.113.193.179 > 192.168.4.227.1024: . 4042780859:4042780859(0)
ack 1845421797 win 16384 <nop,nop,timestamp 4935628 965951> [tos 0xc0]  (ttl )
 In 203.0.113.193.179 > 192.168.4.227.1024: P 4042780859:4042780912(53)
ack 1845421797 win 16384
<nop,nop,timestamp 4935628 965951>:
BGP [|BGP UPDAT)
In 192.168.4.227.1024 > 203.0.113.193.179:
P 1845421797:1845421852(55) ack 4042780912 win 16384 <nop,nop,timestamp 965951 4935628>: BGP [|
BGP UPDAT)
...
```

monitor traffic extensive (Relative Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20 matching
"tcp"
listening on fxp0
 In 172.24.248.221.1680 > 192.168.4.210.23: . 396159737:396159737(0)
ack 1664980689 win 17574 (DF) (ttl 121, id 50003)
Out 192.168.4.210.23 > 172.24.248.221.1680: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10]  (ttl 64, id 5394)
 In 203.0.113.193.179 > 192.168.4.227.1024: P 4042775817:4042775874(57)
```

```
ack 1845416593 win 16384 <nop,nop,timestamp 4935379 965690>: BGP [|BGP UPDAT)
...
```

monitor traffic extensive count

```
                                monitor traffic extensive count 5 no-domain-names no-
resolve
listening on fxp013:18:17.406933
In 192.168.4.206.2723610880 > 172.17.28.8.2049:
40 null (ttl 64, id 38367)13:18:17.407577
In 172.17.28.8.2049 > 192.168.4.206.2723610880:
reply ok 28 null (ttl 61, id 35495)13:18:17.541140
In 0:e0:1e:42:9c:e0 0:e0:1e:42:9c:e0 9000 60:
0000 0100 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 000013:18:17.591513
In 172.24.248.156.4139 > 192.168.4.210.23:
3556964918:3556964918(0)
ack 295526518 win 17601 (DF)
(ttl 121, id 14)13:18:17.591568
Out 192.168.4.210.23 >
172.24.248.156.4139: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10]
(ttl 64, id 52376)
```

monitor traffic interface

```
user@host> monitor traffic interface fxp0
listening on fxp0.0
18:17:28.800650  In server.home.net.723 > host1-0.lab.home.net.log
18:17:28.800733 Out host2-0.lab.home.net.login > server.home.net.7
18:17:28.817813  In host30.lab.home.net.syslog > host40.home0
18:17:28.817846  In host30.lab.home.net.syslog > host40.home0
...
```

monitor traffic interface (Junos OS Evolved)

In this example, ae0 is a physical interface and ae0.1 is a logical interface.

```
user@host> monitor traffic interface ae0
reading from file -, link-type EN10MB (Ethernet)
17:51:30.691523 LLDP, length 441: host.example.com
17:51:32.296133 LLDP, length 445: host.example.com
17:51:33.029399 LLDP, length 445: host.example.com
17:51:33.523333 LLDP, length 445: host.example.com
...
user@host> monitor traffic interface ae0.1
reading from file -, link-type EN10MB (Ethernet)
17:51:20.932958 IP 10.1.1.2 > 10.1.1.1: ICMP echo request, id 33378, seq 4, length 64
17:51:20.933273 IP 10.1.1.1 > 10.1.1.2: ICMP echo reply, id 33378, seq 4, length 64
17:51:21.933840 IP 10.1.1.2 > 10.1.1.1: ICMP echo request, id 33378, seq 5, length 64
17:51:21.934147 IP 10.1.1.1 > 10.1.1.2: ICMP echo reply, id 33378, seq 5, length 64
...
```

monitor traffic matching

```
user@host> monitor traffic matching "net 192.168.1.0/24"
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on fxp0, capture size 96 bytes

Reverse lookup for 192.168.1.255 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use no-resolve to avoid reverse lookups on IP addresses.

21:55:54.003511  In IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003585 Out IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003864  In arp who-has 192.168.1.17 tell 192.168.1.9
...
```

monitor traffic (TX Matrix Plus Router)

```
user@host> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on em0, capture size 96 bytes

04:11:59.862121 Out IP truncated-ip - 25 bytes missing!
summit-em0.example.net.syslog > sv-log-01.example.net.syslog:
SYSLOG kernel.info, length: 57
04:11:59.862303
Out IP truncated-ip - 25 bytes missing!
summit-em0.example.net.syslog >
sv-log-02.example.net.syslog: SYSLOG kernel.info, length: 57
04:11:59.923948
In IP aj-em0.example.net.65235 >
summit-em0.example.net.telnet: .
ack 1087492766 win 33304 <nop,nop,timestamp 42366734 993490>
04:11:59.923983 Out IP truncated-ip - 232 bytes missing!
summit-em0.example.net.telnet > aj-em0.example.net.65235: P 1:241(240) ack 0 win 33304
<nop,nop,timestamp 993590 42366734>
04:12:00.022900
In IP aj-em0.exmaple.net.65235 >
summit-em0.example.net.telnet: . ack 241 win 33304 <nop,nop,timestamp 42366834 993590>
04:12:00.141204
In IP truncated-ip - 40 bytes missing!
ipg-lnx-shell1.example.net.46182 > summit-em0.example.net.telnet: P 2950530356:2950530404(48)
ack 485494987 win 63712
<nop,nop,timestamp 1308555294 987086>
04:12:00.141345
Out IP summit-em0.example.net.telnet >
ipg-lnx-shell1.example.net.46182: P 1:6(5)
ack 48 win 33304
<nop,nop,timestamp 993809 1308555294>
04:12:00.141572
In IP ipg-lnx-shell1.example.net.46182 >
summit-em0.example.net.telnet: .
ack 6 win 63712
<nop,nop,timestamp 1308555294 993809>
04:12:00.141597
Out IP summit-em0.example.net.telnet >
```

```
ipg-lnx-shell1.example.net.46182: P 6:10(4) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.141821
In IP ipg-lnx-shell1.example.net.46182 >
summit-em0.exmaple.net.telnet: .
ack 10 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.141837 Out IP truncated-ip - 2 bytes missing!
summit-em0.example.net.telnet >
ipg-lnx-shell1.example.net.46182: P 10:20(10) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.142072
In IP ipg-lnx-shell1.example.net.46182 >
summit-em0.example.net.telnet: . ack 20 win 63712
<nop,nop,timestamp 1308555294 993810>
04:12:00.142089 Out IP summit-em0.example.net.telnet >
ipg-lnx-shell1.example.net.46182: P 20:28(8) ack 48 win 33304 <nop,nop,timestamp 993810
1308555294>
04:12:00.142321
In IP ipg-lnx-shell1.exmample.net.46182 >
summit-em0.englab.example.net.telnet: .
ack 28 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.142337
Out IP truncated-ip - 1 bytes missing!
summit-em0.example.net.telnet >
ipg-lnx-shell.example.net.46182: P 28:37(9) ack 48 win 33304 <nop,nop,timestamp 993810
1308555294>
...
```

monitor traffic (QFX3500 Switch)

```
user@switch> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on me4, capture size 96 bytes
Reverse lookup for 172.22.16.246 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use <no-resolve> to avoid reverse lookups on IP addresses.
16:35:32.240873 Out IP truncated-ip - 112 bytes missing! labqfx-me0.example.net.ssh >
172.22.16.246.telefinder: P 4200727624:4200727756(132) ack 2889954831 win 65535
16:35:32.240900 Out IP truncated-ip - 176 bytes missing! labqfx-me0.example.net.ssh >
```

```
172.22.16.246.telefinder: P 132:328(196) ack 1 win 65535
...
```

monitor traffic matching icmp

```
user@host> monitor traffic matching "icmp" no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

09:23:17.728737  In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq 322, length 40
09:23:17.728780 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq 322, length 40
09:23:18.735848  In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq 323, length 40
09:23:18.735891 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq 323, length 40
09:23:19.749732  In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq 324, length 40
09:23:19.749775 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq 324, length 40
09:23:20.749747  In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq 325, length 40
09:23:20.749791 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq 325, length 40
...
```

monitor traffic matching IP protocol number

```
user@host> monitor traffic matching "proto 89" no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

13:06:14.700311  In IP truncated-ip - 16 bytes missing! 10.94.211.254 > 224.0.0.
5: OSPFv2, Hello, length 56
13:06:16.067010  In IP truncated-ip - 20 bytes missing! 10.94.211.102 > 224.0.0.
5: OSPFv2, Hello, length 60
13:06:16.287566  In IP truncated-ip - 20 bytes missing! 10.94.211.142 > 224.0.0.
5: OSPFv2, Hello, length 60
13:06:20.758500  In IP truncated-ip - 16 bytes missing! 10.200.211.254 > 224.0.0
.5: OSPFv2, Hello, length 56
13:06:24.309882  In IP truncated-ip - 20 bytes missing! 10.94.211.102 > 224.0.0.
5: OSPFv2, Hello, length 60
13:06:24.396699  In IP truncated-ip - 16 bytes missing! 10.94.211.254 > 224.0.0.
5: OSPFv2, Hello, length 56
13:06:25.067386  In IP truncated-ip - 20 bytes missing! 10.94.211.142 > 224.0.0.
```

```
5: OSPFv2, Hello, length 60
13:06:29.499988  In IP truncated-ip - 16 bytes missing! 10.200.211.254 > 224.0.0
.5: OSPFv2, Hello, length 56
13:06:32.858753  In IP truncated-ip - 20 bytes missing! 10.94.211.102 > 224.0.0.
5: OSPFv2, Hello, length 60
...
```

monitor traffic matching arp

```
user@host> monitor traffic matching “arp” no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

11:57:54.664501  In arp who-has 10.10.213.109 (00:1f:d5:f3:28:30) tell 10.10.213.31
11:57:56.828387  In arp who-has 10.10.213.233 (00:24:9d:06:77:4f) tell 10.10.213.31
11:58:01.735803  In arp who-has 10.10.213.251 (88:e0:f4:1d:41:40) tell 10.10.213.31
11:58:04.663241  In arp who-has 10.10.213.254 tell 10.94.211.170
11:58:28.488191  In arp who-has 10.10.213.149 (00:e0:91:c2:ff:8d) tell 10.10.213.31
11:58:41.858612  In arp who-has 10.10.213.148 tell 10.94.211.254
11:58:42.621533  In arp who-has 10.10.213.254 (5f:5e:ac:79:49:81) tell 10.10.213.31
11:58:44.533391  In arp who-has 10.10.213.186 tell 10.94.211.254
11:58:45.170405  In arp who-has 10.10.213.186 tell 10.94.211.254
11:58:45.770512  In arp who-has 10.10.213.186 tell 10.94.211.254
```

monitor traffic matching port

```
user@host> monitor traffic matching “port 22” no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

13:14:19.108089  In IP 192.0.2.22.56714 > 10.19.300.05.22: S 2210742342:2210742342(0) win 65535
<mss 1360,nop,wscale 7,nop,nop,sackOK>
13:14:19.108165 Out IP 10.19.300.05.22 > 192.0.2.22.56714: S 23075150:23075150(0) ack 2210742343
win 65535 <mss 1460,nop,wscale 1,sackOK,eol>
13:14:19.136883  In IP 192.0.2.22.56714 > 10.19.300.05.22: . ack 1 win 32768
13:14:19.231364 Out IP truncated-ip - 1 bytes missing! 10.19.300.05.22 > 172.29.102.9.56714: P
1:22(21) ack 1 win 33320
13:14:19.260174  In IP truncated-ip - 10 bytes missing! 192.0.2.22.56714 > 10.94.211.93.22: P
```

```
1:31(30) ack 22 win 32767
13:14:19.284865 Out IP truncated-ip - 964 bytes missing! 10.19.300.05.22 > 172.29.102.9.56714: P
22:1006(984) ack 31 win 33320
13:14:19.314549  In IP truncated-ip - 652 bytes missing! 192.0.2.22.56714 > 10.94.211.93.22: P
31:703(672) ack 1006 win 32760
13:14:19.414135 Out IP 10.19.300.05.22 > 192.0.2.22.56714: . ack 703 win 33320
13:14:19.443858  In IP 192.0.2.22.56714 > 10.19.300.05.22: P 703:719(16) ack 1006 win 32760
13:14:19.467379 Out IP truncated-ip - 516 bytes missing! 10.19.300.05.22 > 172.29.102.9.56714: P
1006:1542(536) ack 719 win 33320
13:14:19.734097  In IP 192.0.2.22.56714 > 10.19.300.05.22: . ack 1542 win 32768
13:14:19.843574  In IP truncated-ip - 508 bytes missing! 192.0.2.22.56714 > 10.94.211.93.22: P
719:1247(528) ack 1542 win 32768
...
```

monitor traffic read-files

```
user@host> monitor traffic read-file tcpdump_20_7_18.pcap
15:20:42.597413 Out IP 128.0.0.1.6234 > 128.0.0.17.37217: . ack 1416364513 win 65535
<nop,nop,timestamp 2494269906 347794433>
15:20:42.597424 Out IP 128.0.0.1.6234 > 128.0.0.16.49400: . ack 3549610340 win 65535
<nop,nop,timestamp 2494269906 347799892>
15:20:42.598214 Out IP truncated-ip - 32 bytes missing! 128.0.0.1.6234 > 128.0.0.16.49400: P
0:40(40) ack 1 win 65535 <nop,nop,timestamp 2494269907 347799892>
                         0001 0000 0020 0000
```

monitor traffic write-file

```
user@host> monitor traffic write-file filename
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on em1, capture size 96 bytes

^C
955 packets received by filter
0 packets dropped by kernel
```

Release Information

Command introduced before Junos OS Release 7.4.

Options `read-file` and `write-file` introduced in Junos OS Release 19.1R1.

monitor-timer

IN THIS SECTION

- Syntax | 10953
- Hierarchy Level | 10953
- Description | 10953
- Options | 10954
- Required Privilege Level | 10954
- Release Information | 10954

Syntax

```
monitor-timer {
    time seconds;
}
```

Hierarchy Level

```
[edit services pfe traffic]
```

Description

Configure timer for Packet Forwarding engine traffic to uninstall the filter and to increase the default timer expiry from 300 seconds. This setting configures a maximum number of 3600 seconds and delete the trace log entries upon timer expiry.

Options

time Default is 5 mins, max configurable time is 60 mins (seconds).

- **Default:** 300 seconds
- **Range:** 300 through 3600 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

monitoring

IN THIS SECTION

- Syntax | 10954
- Hierarchy Level | 10955
- Description | 10955
- Required Privilege Level | 10955
- Release Information | 10955

Syntax

```
monitoring group-name {
    family inet {
        output {
            cflowd hostname {
                port port-number;
```

```
            }
            export-format cflowd-version-5;
            flow-active-timeout seconds;
            flow-export-destination {
                (cflowd-collector | collector-pic);
            }
            flow-inactive-timeout seconds;
            interface interface-name {
                engine-id number;
                engine-type number;
                input-interface-index number;
                output-interface-index number;
                source-address address;
            }
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Specify flow monitoring instance name and properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Passive Flow Monitoring

monitoring (Forwarding Options)

IN THIS SECTION

- Syntax | 10956
- Hierarchy Level | 10957
- Description | 10957
- Required Privilege Level | 10957
- Release Information | 10957

Syntax

```
monitoring name {
    family inet {
        output {
            cflowd hostname port-number;
            export-format cflowd-version-5;
            flow-active-timeout seconds;
            flow-export-destination {
                (cflowd-collector | collector-pic);
            }
            flow-inactive-timeout seconds;
            interface interface-name {
                number;
                engine-type number;
                input-interface-index number;
                output-interface-index number;
                source-address address;
            }
        }
```

```
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Specify the flow monitoring instance name and properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Flow Monitoring

monitoring-key (PCC Action Profile)

IN THIS SECTION

- Syntax | 10958
- Hierarchy Level | 10958
- Description | 10958

- Options | 10958
- Required Privilege Level | 10958
- Release Information | 10958

Syntax

```
monitoring-key key_string
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name]
```

Description

Specify the monitoring key that controls TDF subscriber usage monitoring for traffic that matches the data flows or applications identified in the predefined PCC rules containing the PCC action profile. The monitoring key is defined by the PCRF.

Options

key_string	Identifier for the monitoring key.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring TDF Subscriber Usage Monitoring for Traffic That Matches Predefined PCC Rules

Understanding Usage Monitoring for TDF Subscribers

monitoring-profile

IN THIS SECTION

- Syntax | 10959
- Hierarchy Level | 10959
- Description | 10960
- Options | 10960
- Required Privilege Level | 10961
- Release Information | 10961

Syntax

```
monitoring-profile profile-name {
    export-filters filter-name {
        peak-queue-length {
            percent percent;
        }
        queue [queue(s)];
    }
}
```

Hierarchy Level

```
[edit class-of-service],
[edit dynamic-profiles name class-of-service]
```

Description

Define a queue-length monitor to measure peak queue occupancy of a virtual output queue (VOQ).

In a typical deployment scenario that uses strict-priority scheduling, a `HIGH` priority queue can starve `LOW` priority queues. Thus the packets in such `LOW` priority queues can remain longer than desired. Use this VOQ queue-length monitoring feature, along with queue transmit statistics, to detect such stalled conditions.

Options

monitoring-profile

(Required) Name of the monitoring profile.

Use the monitoring profile name to attach the monitoring profile to one or more interfaces. Each monitoring profile contains one or more export filters.

export-filters

(Required) Name of the export filter.

Each export filter consists of an optional peak queue length percentage filter and one or more queues.

peak-queue-length percent

(Optional) The percentage of the peak queue length filter.

When configured, Junos exports VOQ telemetry only when the hardware-captured peak queue length exceeds the configured value. When not configured, Junos always exports VOQ telemetry.

- **Range:** 0 through 100
- **Default:** 0

queue

(Required) The queue, or queues, for the export filter to monitor.

- **Range:** 0 through 7

Attach the monitoring profile at the `[edit class-of-service interfaces interface-name]` hierarchy level. For example:

```
[edit class-of-service]
user@host# show
interfaces {

    ae0 {
        monitoring-profile mp1 shared;
    }
```

```
    et-0/0/1 {
        monitoring-profile mp1;
    }
}
```

By default, a monitoring profile that you assign to an aggregated Ethernet (ae-) interface replicates across all members of the ae- interface. The monitoring profile also reports virtual output queue (VOQ) depth individually on each interface. On large systems, this process can quickly consume the maximum supported hardware monitoring profile IDs. To conserve monitoring profile IDs, include the `shared` option at the [`set class-of-service interfaces` *`ae-interface`* `monitoring-profile` *`profile-name`*] hierarchy level. The configured `shared` option creates only one monitoring profile ID to share across all member interfaces. The option also reports the largest peak on a member interface as the common peak for the ae-interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

VOQ Queue-depth Monitoring

show class-of-service monitoring-profile | **19108**

monitoring (Services)

IN THIS SECTION

- Syntax | **10962**
- Hierarchy Level | **10965**
- Description | **10965**

Required Privilege Level | 10965
Release Information | 10966

Syntax

```
monitoring {
rfc2544 {
    test-profile profile-name {
        bandwidth bandwidth;
        packet-size size;
        step-percent percent;
        test-type (throughput | latency | frame-loss | back-back-frames);
    }
    tests{
        test-name test-name {
            destination-ipv4-address address;
            destination-mac-address mac-address;
            destination-udp-port port-number;
            direction (egress | ingress);
            disable-signature-check;
            dscp-code-points codepoint;
            family (ccc | ethernet-switching |inet);
            halt-on-prefix-down;
            in-service;
            ip-swap;
            mode (initiate-and-terminate | reflect);initiate RLI 45755/48989
            reflect-etype;
            reflect-mode (mac-swap | no-mac-swap);
            service-type service-type;
            skip-arp-iteration;
            source-ipv4-address address;
            source-mac-address mac-address;
            source-udp-port port-number;
            test-interface interface-name;
            test-iterator-duration seconds;
            test-profile profile-name;
            udp-tcp-port-swap;
        }
```

```
    }
}
rpm {
    owner name {
        test test-name {
            data-fill data;
            data-size size;
            destination-port port;
            dscp-code-points dscp-bits;
            history-size size;
            moving-average-size number;
            offload-type {
                none;
                pfe-timestamp;
            }
            probe-count count;
            probe-interval seconds;
            probe-server {
                icmp;
                tcp {
                    port number;
                }
                udp {
                    port number;
                }
            }
            probe-type type;
            routing-instance instance-name;
            source-address address;
            target (url url | address address);
            test-interval interval;
            thresholds thresholds;
            traps traps;
            ttl hop-count;
        }
    }
}
twamp {
    client {
        control-connection  name {
            control-type (managed | light);
            destination-port  destination-port;
            routing-instance routing-instance;
```

```
            source-address source-address;
            target target-address;
            test-start (auto | manual);
            test-interval seconds;
            traps {
                control-connection-closed;
                test-iteration-done;
            }
            test-session name {
                data-size data-size;
                destination-port destination-port;
                dscp-code-points dscp-code-points;
                history-size history-size;
                moving-average-size moving-average-size;
                offload-type (none | pfe-timestamp);
                probe-count  probe-count;
                probe-interval  seconds;
                source-address source-address;
                target target-address local-link IPv6-link-local-interface-
name ;
                thresholds {
                    control-failure (on | off);
                    successive-loss number;
                    total-loss number;
                    threshold-type (microseconds | average);
                }
                traps {
                    egress-jitter-exceeded
                    egress-time-exceeded;
                    ingress-jitter-exceeded;
                    ingress-time-exceeded;
                    jitter-exceeded;
                    probe-failure;
                    rtt-exceeded;
                    test-completion;
                    test-failure;
                }
                ttl hop-count;
            }
        }
    }
    server {
        managed {
```

```
            client-limit limit;
            client-list  {
                address address <routing-instance [instance-name...]>;
            }
            control-inactivity-timeout seconds;
            control-per-client-limit limit;
            control-maximum-duration seconds;
            offload-type (none | inline-timestamp | pfe-timestamp);
            port  port;
            stateful-sequence;
            test-per-client-limit limit;
        }
        light {
            offload-type (none | inline-timestamp | pfe-timestamp);
            port  port;
            stateful-sequence cache-size size;
        }
    }
}
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure monitoring services.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.1R1.

`twamp` statement introduced in Junos OS Evolved Release 20.3R1.

`rfc2544` statement introduced in Junos OS Evolved Release 21.1R1.

`traps` option for RPM introduced in Junos OS Evolved Release 21.2R1.

`traps` option for TWAMP added in Junos OS Evolved Release 21.3R1.

The `local-link` sub-option of the `target` option for TWAMP Light test sessions introduced in Junos OS Evolved Release 22.3R1.

Support for the `destination-mac-address`, `direction`, `disable-signature-check`, `in-service`, `ip-swap`, `reflect-etype`, `reflect-mode`, `service-type`, `source-mac-address`, and `udp-tcp-port-swap` configuration statements introduced in Junos OS Evolved 22.4R1.

The `inline-timestamping` suboption of the `offload-type` option introduced in Junos OS Evolved 22.4R1.

Support for `bandwidth`, `dscp-code-points`, `halt-on-prefix-down`, `mode initiate-and-terminate`, `packet-size`, `step-percent`, `test-iterator-duration`, `test-profile`, and `test-type` configuration statements introduced in Junos OS Evolved Release 23.4R1.

`tcp` sub-option for the `probe-server` option for RPM introduced in Junos OS Evolved Release 23.4R1.

`tcp-ping`, `http-get`, and `http-metadata-get` probe types for RPM introduced in Junos OS Evolved Release 23.4R1.

`offload-type` and `stateful-sequence` options for light and managed TWAMP servers introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Understanding Two-Way Active Measurement Protocol on Routers

most-fill

IN THIS SECTION

- See | 10967

See

"random" on page 14043

moving-average-size

IN THIS SECTION

- Syntax | 10967
- Junos OS Hierarchy Levels | 10968
- Junos OS Evolved Hierarchy Level | 10968
- Description | 10968
- Options | 10968
- Required Privilege Level | 10968
- Release Information | 10968

Syntax

```
moving-average-size number;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Enable statistical calculation operations to be performed across a configurable number of the most recent samples.

Options

number—Number of samples to be used in calculations.

- **Range:** 0 through 255

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Statement at the [`edit services rpm twamp client control-connection` *`control-client-name`*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

mpls

IN THIS SECTION

- Syntax | 10969
- Hierarchy Level | 10971
- Description | 10971
- Default | 10972
- Required Privilege Level | 10972
- Release Information | 10972

Syntax

```
mpls {
    disable;
    class-of-service cos-value;
    no-cspf;
    no-decrement-ttl;

    advertisement-hold-time seconds;
    explicit-null;
    icmp-tunneling;
    interface (interface-name | all) {
        disable;
    }
    ipv6-tunneling;
    no-propagate-ttl;
    path path-name {
        (address | hostname) <loose | strict>;
```

```
    }
    label-switched-path lsp-name {
        disable;
        auto-bandwidth {
            adjust-interval seconds;
            adjust-threshold percentage;
            adjust-threshold-overflow-limit count;
            adjust-threshold-underflow-limit
            maximum-bandwidth bps;
            minimum-bandwidth bps;
            monitor-bandwidth;
        }
        description text-string;
        from address;
        install destination-prefix</prefix-length> <active>;
        ldp-tunneling;
        no-cspf;
        no-decrement-ttl;
        primary path-name {
            adaptive;
            select (manual | unconditional);
        }
        secondary path-name {
            adaptive;
            select (manual | unconditional);
        }
        to address;
        traceoptions {
            file filename <files number> <size maximum-file-size> <world-readable | no-world-
readable>;
            flag flag;
        }
    }
    static-label-switched-path lsp-name {
        bypass bypass-name {
            description text-string;
            next-hop (address | interface-name | address/interface-name);
            to address;
        }
        ingress {
            description string;
            install {
                destination-prefix <active>;
```

```
            }
            link-protection bypass-name name;
            next-hop (address | interface-name | address/interface-name);
            to address;
        }
        transit incoming-label {
            bandwidth bps;
            description text-string;
            link-protection bypass-name name;
            next-hop (address | interface-name | address/interface-name);
            pop;
            swap out-label;
        }
    statistics {
        auto-bandwidth;
        file filename <files number> <size maximum-file-size> <world-readable | no-world-
readable>;
        interval seconds;
    }
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable | no-world-
readable>;
        flag flag;
    }
    traffic-engineering (bgp | bgp-igp);
}
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Enable MPLS on the switch.

The remaining statements are explained separately.

Default

MPLS is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

- *Example: Configuring MPLS on EX8200 and EX4500 Switches*
- *Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS*
- *Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect*
- *Configuring MPLS on EX8200 and EX4500 Provider Switches*
- Junos OS MPLS Applications Configuration Guide

mpls-flow-table-size

IN THIS SECTION

- Syntax | **10973**
- Hierarchy Level | **10973**
- Description | **10973**
- Options | **10973**
- Required Privilege Level | **10973**
- Release Information | **10974**

Syntax

```
mpls-flow-table-size units;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services flow-table-size]
```

Description

Configure the size of the MPLS flow table in units of 256,000 entries.

> **NOTE**: Starting with Junos OS Release 17.3R1, the maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 256 on MPC5Es and MPC6Es with with 4 GB DDR3 memory or higher. The maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 245 on MPC5Es and MPC6Es with DDR3 memory lower than 4 GB.

> **NOTE**: The recommended flow table size is 4 so that it can scale up to 4x256K flows, which is 1M. You can configure more, however, the system will issue a warning message.

Options

units Number of 256,000 flow entries available for the MPLS flow table.

- **Range:** 1 through 245
- **Default:** 15 (3,840,000)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

mpls (security)

IN THIS SECTION

- Syntax | 10974
- Hierarchy Level | 10974
- Description | 10975
- Options | 10975
- Required Privilege Level | 10975
- Release Information | 10975

Syntax

```
mpls {
    mode(drop | packet-based);
}
```

Hierarchy Level

```
[edit security forwarding-options family]
```

Description

The mpls has two modes, drop and packet mode. The default mode for mpls is drop.

Family MPLS

Options

mode Forwarding mode

- drop—Disable forwarding
- packet-based—Enable packet-based forwarding

Required Privilege Level

[none specified]

Release Information

Statement supported in Junos 24.2R1 for SRX Series Firewalls

mpls (forwarding-options)

IN THIS SECTION

- Syntax | 10976
- Hierarchy Level | 10976
- Description | 10976
- Default | 10976
- Required Privilege Level | 10976
- Release Information | 10976

Syntax

```
mpls port-number
```

Hierarchy Level

```
[edit security forwarding-options tunnel udp payload-port-profile profile-name]
```

Description

Enable the forwarding of MPLS traffic.

Default

If port number is not configured for MPLS, the default value of the port number is 6635.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos Evolved Release 21.4R1 for the PTX Series routers.

RELATED DOCUMENTATION

MPLS Overview

mpls-inet6

IN THIS SECTION

- Syntax | 10977
- Hierarchy Level | 10977
- Description | 10977
- Required Privilege Level | 10977
- Release Information | 10978

Syntax

```
mpls-inet6;
```

Hierarchy Level

```
[edit routing-instances name protocols evpn encapsulation]
```

Description

Enables the use of the EVPN-VXLAN and point-to-point (P2P) services over an MPLS-based IPv6 underlay on the PE routers using Segment Routing with Multiprotocol Label Switching (SR-MPLS). The services also operate on the segment routing–traffic engineering (SR-TE) addresses that are responsible for the path calculation between the EVPN PE devices. You can use the EVPN-MPLS commands with the MPLS-based IPv6 underlays.

To enable EVPN-MPLS over IPv6 functionality, include the `protocols evpn encapsulation mpls-inet6` configuration statement at the `[edit routing-instance name]` hierarchy level for each EVPN routing instance.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 21.2.

mpls-internet-multicast

IN THIS SECTION

- Syntax | 10978
- Hierarchy Level | 10978
- Description | 10978
- Required Privilege Level | 10979
- Release Information | 10979

Syntax

```
mpls-internet-multicast;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name instance-type]
[edit protocols pim]
```

Description

A nonforwarding routing instance type that supports Internet multicast over an MPLS network for the default master instance. No interfaces can be configured for it. Only one `mpls-internet-multicast` instance can be configured for each logical system.

The `mpls-internet-multicast` configuration statement is also explicitly required under PIM in the master instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

> *Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs*
>
> ingress-replication | **8647**

mpls-ipv4-template

IN THIS SECTION

- Syntax | **10979**
- Hierarchy Level | **10980**
- Description | **10980**
- Required Privilege Level | **10980**
- Release Information | **10980**

Syntax

```
mpls-ipv4-template {
    label-position [ positions ];
}
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Specify the flow aggregation version 9 or IPFIX properties for templates that combine IPv4 and MPLS records. The remaining statement is explained separately.

Starting in Junos OS Release 18.4R1, use `mpls-template` instead of `mpls-ipv4-template` for inline flow monitoring of MPLS-IPv4 flows on the MX Series. You must also configure `tunnel-observation ipv4` at the `[edit services flow-monitoring (version-ipfix | version9) template template-name]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates*
>
> *Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers*
>
> *Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers*

mpls-ipvx-template

IN THIS SECTION

- Syntax | 10981
- Hierarchy Level | 10981
- Description | 10981
- Required Privilege Level | 10982
- Release Information | 10982

Syntax

```
mpls-ipvx-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

In Junos OS Release 18.1, specify the version 9 or IPFIX template for inline monitoring an MPLS-over-UDP flow that is encapsulated in an RSVP-TE LSP on PTX Series routers. This monitoring looks past the tunnel header to report the inner payload of the packets. To use the template for MPLS-over-UDP flows, you must also configure `tunnel-observation mpls-over-udp` at the `[edit services flow-monitoring (version 9 | version-ipfix) template template-name]` hierarchy level.

Starting in Junos OS Release 18.2R1, use `mpls-template` instead of `mpls-ipvx-template`.

> **NOTE**: For an MPLS-over-UDP flow that is carried between IPv4 endpoints, configure `ipv4-template` at the `[edit services flow-monitoring (version9 | version-ipfix) template template-name]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers

mpls (packet-forwarding-options)

IN THIS SECTION

- Syntax | 10982
- Hierarchy Level | 10983
- Description | 10983
- Default | 10983
- Options | 10983
- Required Privilege Level | 10983
- Release Information | 10983

Syntax

```
mpls {
 lsp-ingress-stats-enable;
 lsp-transit-stats-enable;
}
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

Enable or disable MPLS LSP statistics at PFE level. When you enable or disable MPLS LSP statistics at PFE level, the PFE automatically restarts.

Default

By default, LSP statistics are disabled.

Options

lsp-ingress-stats-enable	Enable ingress LSP statistics.
lsp-transit-stats-enable	Enable transit LSP statistics.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Options `lsp-ingress-stats-enable` and `lsp-transit-stats-enable` are introduced in Junos OS Evolved Release 22.4R1 for ACX7100-32C, ACX7100-48L, ACX7024, and ACX7509 routers.

mpls (Protocols)

IN THIS SECTION

- Syntax | 10984
- Hierarchy Level | 10984
- Description | 10984
- Required Privilege Level | 10984
- Release Information | 10984

Syntax

```
mpls { ... }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit protocols]
```

Description

Enable MPLS on the router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring MPLS

mpls (Security Forwarding Options)

IN THIS SECTION

- Syntax | 10985
- Hierarchy Level | 10985
- Description | 10985
- Required Privilege Level | 10986
- Release Information | 10986

Syntax

```
mpls {
    mode packet-based;
}
```

Hierarchy Level

```
[edit security forwarding-options family]
```

Description

Enable the forwarding of MPLS traffic. By default, the device drops MPLS traffic.

CAUTION: Because MPLS operates in packet mode, security services are not available.

> **NOTE**: Packet-based processing is not supported on the following SRX Series Firewalls: SRX1500, SRX5600, and SRX5800.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

MPLS Overview

mpls-template

IN THIS SECTION

- Syntax | 10987
- Hierarchy Level | 10987
- Description | 10987
- Required Privilege Level | 10987
- Release Information | 10988

Syntax

```
mpls-template {
    label-position [positions ];
}
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Specify the flow aggregation IPFIX or version 9 properties for templates used only for MPLS records.

Starting in Junos OS Release 18.2R1, you can also use `mpls-template` to specify the version 9 or IPFIX template for inline monitoring an MPLS-over-UDP flow that is encapsulated in an RSVP-TE LSP on PTX Series routers. (In Junos OS Release 18.1, use `mpls-ipvx-template` instead of `mpls-template`.) This monitoring looks past the tunnel header to report the inner payload of the packets. To use the template for MPLS-over-UDP flows, you must also configure `tunnel-observation mpls-over-udp` at the `[edit services flow-monitoring (version 9 | version-ipfix) template template-name]` hierarchy level.

> **NOTE**: For an MPLS-over-UDP flow that is carried between IPv4 endpoints, configure `ipv4-template` at the `[edit services flow-monitoring (version9 | version-ipfix) template template-name]` hierarchy level.

Starting in Junos OS Release 18.4R1, use `mpls-template` instead of `mpls-ipv4-template` for inline flow monitoring of MPLS-IPv4 flows on the MX Series. You must also configure `tunnel-observation ipv4` at the `[edit services flow-monitoring (version-ipfix | version9) template template-name]` hierarchy level.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the [edit services flow-monitoring version-ipfix template *template-name*] hierarchy level introduced in Junos OS Release 16.1.

Statement introduced in Junos OS Release 18.2R1 on PTX Series routers.

RELATED DOCUMENTATION

> *Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates*

mpls-tp-mode

IN THIS SECTION

- Syntax | 10988
- Hierarchy Level | 10988
- Description | 10989
- Required Privilege Level | 10989
- Release Information | 10989

Syntax

```
mpls-tp-mode;
lsping-channel-type;
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name oam],
[edit protocols mpls oam]
```

Description

Enable GAL or G-Ach OAM operation without IP encapsulation on a label-switched path (LSP).

Include this statement at the `[edit protocols mpls oam]` hierarchy level to enable GAL or G-Ach OAM operation without IP encapsulation on all LSPs in the MPLS network. Include this statement at the `[edit protocols mpls label-switched-path` *`lsp-name`* `oam]` hierarchy level to enable GAL and G-Ach OAM operation without IP encapsulation on a specific LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

`lsping-channel-type` statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Example: Configuring the MPLS Transport Profile for OAM

mrru

IN THIS SECTION

- Syntax | 10990
- Hierarchy Level | 10990
- Description | 10990
- Options | 10990
- Required Privilege Level | 10990
- Release Information | 10990

Syntax

```
mrru bytes;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options],
[edit interfaces (ml-fpc/pic/port| ls-fpc/pic/port) unit logical-unit-number],
[edit logical-systems logical-system-name interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-
options],
[edit logical-systems logical-system-name interfaces (ml-fpc/pic/port| ls-fpc/pic/port) unit
logical-unit-number]
```

Description

For multilink or link services interfaces only, set the maximum received reconstructed unit (MRRU). The MRRU is similar to the maximum transmission unit (MTU), but is specific to multilink interfaces.

Options

bytes—MRRU size.

- **Range:** 1500 through 4500 bytes
- **Default:** 1500 bytes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring MRRU on Multilink and Link Services Logical Interfaces*

mru (Gigabit Ethernet Options)

IN THIS SECTION

- Syntax | 10991
- Hierarchy Level | 10991
- Hierarchy Level: SRX Series | 10991
- Description | 10991
- Options | 10992
- Required Privilege Level | 10992
- Release Information | 10992

Syntax

```
mru mru;
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name gigether-options],
[edit dynamic-profiles name logical-systems name interfaces name gigether-options],
[edit interfaces name gigether-options]
```

Hierarchy Level: SRX Series

```
[edit interfaces name gigether-options]
```

Description

Configure the maximum receive unit (MRU) of the interface in bytes. The maximum receive unit of an interface indicates the largest size of a packet that the interface can accept. You can configure the parameters so that the value of MRU equals the value of MTU. You can also configure different values

for MRU and MTU. When a device receives packets whose size is greater than the interface MRU, those packets are dropped by the device's forwarding plane.

Options

mru MRU size in bytes.

- **Range:** 256 through 16008 bytes.
- **Range (SRX Series) :** 256 through 9200 bytes.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Ethernet Interfaces Overview

gigether-options | 7685

mru (Congestion Notification Profiles)

IN THIS SECTION

- Syntax | 10993
- Hierarchy Level | 10993
- Description | 10993
- Default | 10993
- Options | 10993
- Required Privilege Level | 10993
- Release Information | 10994

Syntax

```
mru mru-value;
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name input dscp code-point code-point-bits pfc]
[edit class-of-service congestion-notification-profile profile-name input ieee-802.1 code-point code-point-bits pfc]
```

Description

Configure the maximum receive unit (MRU) of the interface in bytes. The system uses the MRU and the cable length to calculate the amount of buffer headroom reserved to support priority-based flow control (PFC). The lower the MRU and the shorter the cable length, the less headroom buffer space is required for PFC.

> **NOTE**: You can also set a maximum transmission unit (MTU) value (the largest packet size the interface sends) for interfaces by including the `mtu` statement at the `[edit interfaces interface-name]` hierarchy level.

Default

For priority 3 traffic, the default MRU value is 2500 bytes.

For priority 4 traffic, the default MRU value is 9216 bytes.

For user-configured priorities, the default MRU value is 2500 bytes.

Options

mru-value—Value of the maximum packet receive unit size in bytes (generally from 1500 to 9216 bytes, but there is no configuration restriction).

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support added in `dscp` statement hierarchy in Junos OS Release 17.4R1 for the QFX Series.

RELATED DOCUMENTATION

- *Configuring CoS PFC (Congestion Notification Profiles)*
- *Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces*
- *Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface*
- *Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)*
- *Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)*
- *Understanding CoS Flow Control (Ethernet PAUSE and PFC)*
- *Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows*
- *Understanding PFC Using DSCP at Layer 3 for Untagged Traffic*
- *Configuring DSCP-based PFC for Layer 3 Untagged Traffic*

mru (Dynamic and Static PPPoE)

IN THIS SECTION

- Syntax | 10995
- Hierarchy Level | 10995
- Description | 10995
- Options | 10995
- Required Privilege Level | 10995
- Release Information | 10995

Syntax

```
mru size;
```

Hierarchy Level

```
[edit access group-profile group-profile-name ppp ppp-options]
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options],
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name"  unit “$junos-
interface-unit” ppp-options],
[edit interfaces pp0 unit unit-number ppp-options]
[edit interfaces si interface-id unit unit-number ppp-options]
```

Description

Specify the size of maximum receive unit (MRU) that the router uses during link control protocol (LCP) negotiation for dynamic and static PPP subscribers and L2TP tunneled subscribers.

Options

size MRU size in bytes that is used during LCP negotiation.

- **Range:** 64–65,535

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring MTU and MRU for PPP Subscribers

msdp

IN THIS SECTION

- Syntax | 10996
- Hierarchy Level | 10998
- Description | 10998
- Default | 10998
- Options | 10998
- Required Privilege Level | 10998
- Release Information | 10998

Syntax

```
msdp {
    disable;
    active-source-limit {
        log-interval seconds;
        log-warning value;
        maximum number;
        threshold number;
    }
    data-encapsulation (disable | enable);
    export [ policy-names ];
    group group-name {
        ... group-configuration ...
    }
    hold-time seconds;
    import [ policy-names ];
    local-address address;
    keep-alive seconds;
    peer address {
        ... peer-configuration ...
```

```
    }
    rib-group group-name;
    source ip-prefix</prefix-length> {
        active-source-limit {
            maximum number;
            threshold number;
        }
    }
    sa-hold-time seconds;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    group group-name {
        disable;
        export [ policy-names ];
        import [ policy-names ];
        local-address address;
        mode (mesh-group | standard);
        peer address {
            ... same statements as at the  [edit protocols msdp peer address] hierarchy level
shown just following ...
        }
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
    }
    peer address {
        disable;
        active-source-limit {
            maximum number;
            threshold number;
        }
        authentication-key peer-key;
        default-peer;
        export [ policy-names ];
        import [ policy-names ];
        local-address address;
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable MSDP on the router or switch. You must also configure at least one peer for MSDP to function.

Default

MSDP is disabled on the router or switch.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP in a Routing Instance

msrpc

IN THIS SECTION

- Syntax | 10999
- Hierarchy Level | 10999
- Description | 11000
- Options | 11000
- Required Privilege Level | 11000
- Release Information | 11000

Syntax

```
msrpc {
    disable;
    map-entry-timeout map-entry-timeout;
    traceoptions {
        flag {
            all {
                extensive
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Specify the Microsoft remote procedure call Application Layer Gateway (MS-RPC ALG) on the device. MS-RPC provides a way for a program running on one host to call procedures in a program running on another host.

Options

- `disable`—Disable the MS-RPC ALG. By default, the MS-RPC ALG is enabled.
- `map-entry-timeout` *`map-entry-timeout`*—Specify the MS-RPC ALG mapping entry timeout value in hours.

  Range: 1 through 72 hours.

  Default: 32 hours.
- `traceoptions`—Configure MS-RPC ALG tracing options.
  - `flag`—Trace operation to perform.
    - `all`—Trace all events.
    - *`extensive`*—Display extensive amount of data.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding Microsoft RPC ALGs

mss (IDS MS-DPC)

IN THIS SECTION

- Syntax | 11001
- Hierarchy Level | 11001
- Description | 11001
- Options | 11001
- Required Privilege Level | 11002
- Release Information | 11002

Syntax

```
mss value;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then syn-cookie]
```

Description

Specify the maximum segment size (MSS) value used in Transmission Control Protocol (TCP) delayed binding when using the MS-DPC.

Options

value—MSS value.

- **Default:** 1500
- **Range:** 128 through 8192

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in IDS Rules

mss (Security IDP)

IN THIS SECTION

- Syntax | 11002
- Hierarchy Level | 11003
- Description | 11003
- Options | 11003
- Required Privilege Level | 11003
- Release Information | 11003

Syntax

```
mss {
    match (equal | greater-than | less-than | not-equal);
    value maximum-segment-size;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify the maximum segment size (MSS) in the TCP header.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`maximum-segment-size`*—Match the maximum segment size value.
- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

msti

IN THIS SECTION

- Syntax | 11004
- Syntax | 11004
- Hierarchy Level | 11005
- Description | 11005
- Default | 11005

- Options | 11005
- Required Privilege Level | 11005
- Release Information | 11005

Syntax

```
msti msti-id {
    bridge-priority priority;
    vlan (vlan-id | vlan-range|open-set-of-values);
    interface (interface-name | all) {
        cost cost;
        edge;
        priority interface-priority;
    }
}
```

Syntax

```
msti msti-id {
     vlan (vlan-id | vlan-name);
     interface interface-name {
         disable-timeout;
         cost cost;
         priority priority;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mstp],
[edit logical-systems logical-system-name routing-instances routing-instance-name
protocols mstp],
```

```
[edit protocols mstp],
```

```
[edit routing-instances routing-instance-name protocols mstp]
```

Description

Configures the Multiple Spanning Tree Instance (MSTI) identifier for Multiple Spanning Tree Protocol (MSTP). MSTI IDs are local to each region, so you can reuse the same MSTI ID in different regions.

Default

MSTI is disabled.

Options

msti-id—MSTI instance identifier.

- **Range:** 1 through 64
- **Range:** 1 through 4094. The Common Instance Spanning Tree (CIST) is always MSTI 0.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

Statement updated in Junos OS Release 15.1 for EX Series switches to support configuration of spanning tree parameters globally on all interfaces.

NOTE: You cannot disable spanning tree parameters globally on all interfaces.

RELATED DOCUMENTATION

Understanding MSTP

Configuring Multiple Spanning Tree Protocol

Configuring MSTP on Switches

Configuring MSTP Instances on a Physical Interface

Example: Configuring Network Regions for VLANs with MSTP on Switches

Understanding MSTP

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

mstp

IN THIS SECTION

- MX Series, QFX Series, EX Series (Enhanced Layer 2 Software (ELS) Configuration Style) | **11007**
- EX Series (Without Enhanced Layer 2 Software (ELS) Configuration Style) | **11008**
- Hierarchy Level | **11009**
- Description | **11009**
- Required Privilege Level | **11009**
- Release Information | **11009**

MX Series, QFX Series, EX Series (Enhanced Layer 2 Software (ELS) Configuration Style)

```
mstp {
    backup-bridge-priority priority;
    bpdu-block-on-edge;
    bpdu-destination-mac-address provider-bridge-group;
    bridge-priority priority;
    configuration-name configuration-name;
    disable;
    forward-delay seconds;
    hello-time seconds;
    interface ( all | interface-name ) {
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode (p2p | shared);
        no-root-port;
        priority interface-priority;
    }
    max-age seconds;
    max-hops hops;
    msti msti-id {
        backup-bridge-priority priority;
        bridge-priority priority;
        interface interface-name {
            cost cost;
            priority interface-priority;
        }
        vlan vlan-id;
    }
    priority-hold-time seconds;
    revision-level revision-level;
    system-id system-id-value {
        ip-address(es);
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
```

```
    }
    vpls-flush-on-topology-change;
}
```

EX Series (Without Enhanced Layer 2 Software (ELS) Configuration Style)

```
mstp  {
                  bpdu-block-on-edge;
                          bridge-priority                    priority;
                          configuration-name                    name;
                          disable;
                          forward-delay                     seconds;
                          hello-time                      seconds;
                              interface ( all | interface-name  ){
                                      bpdu-timeout-action {
                                          block;
                                          log;
        }
                                  cost                          cost;
                                  disable;
                                  edge;
                                  mode                         mode;
                                  no-root-port;
                                  priority                         priority;
    }
                          max-age                       seconds;
                          max-hops                      hops;
                              msti                          msti-id {
        bridge-priority priority;
        interface                              interface-name {
            cost                             cost;
            disable;
            priority                             priority;
        }
        vlan (vlan-id | vlan-name);
    }
    revision-level                     revision-level;
    traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable | no-world-
readable>;
        flag flag;
```

```
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Configures Multiple Spanning Tree Protocol (MSTP). MSTP is defined in the IEEE 802.1Q-2003 specification and is used to create a loop-free topology in networks with multiple spanning tree regions.

> **NOTE**: MX Series routers and QFX5100 switches do not support the `interface all` statement to enable MSTP on all interfaces with one command.
>
> On MX Series routers, you must enable MSTP on interfaces individually using the `set ... protocols mstp interface` *`interface-name`* statement.
>
> On QFX5100 switches, to apply MSTP configuration to more than one interface, you must first configure one or more interface ranges for the interfaces on which you want to configure MSTP, and then issue the `set protocols mstp interface` *`interface-name`* command using each interface range as the *interface-name* parameter.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

`bpdu-block-on-edge` statement added in Junos OS Release 9.4.

`bpdu-timeout-action` statement added in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding MSTP

Configuring Multiple Spanning Tree Protocol

Configuring MSTP on Switches

Example: Configuring Network Regions for VLANs with MSTP on Switches

show spanning-tree bridge | 26548

show spanning-tree interface | 26555

mtrace

IN THIS SECTION

- Syntax | 11010
- Description | 11011
- Options | 11011
- Additional Information | 11011
- Required Privilege Level | 11011
- Output Fields | 11011
- Sample Output | 11012
- Release Information | 11012

Syntax

```
mtrace source
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display trace information about an IP multicast path.

Options

source	Source hostname or address.
logical-system (*logical-system-name*)	(Optional) Perform this operation on a logical system.
routing-instance *routing-instance-name*	(Optional) Trace a particular routing instance.

Additional Information

The **mtrace** command for multicast traffic is similar to the `traceroute` command used for unicast traffic. Unlike **traceroute**, **mtrace** traces traffic backwards, from the receiver to the source.

Required Privilege Level

view

Output Fields

Table 86 on page 11011 describes the output fields for the `mtrace` command. Output fields are listed in the approximate order in which they appear.

Table 86: mtrace Output Fields

Field Name	Field Description
Mtrace from	IP address of the receiver.
to	IP address of the source.
via group	IP address of the multicast group (if any).
Querying full reverse path	Indicates the full reverse path query has begun.

Table 86: mtrace Output Fields *(Continued)*

Field Name	Field Description
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
protocol	Protocol used (for example, PIM).
Round trip time	Average round-trip time, in milliseconds (ms).
total ttl of	Time-to-live (TTL) threshold.

Sample Output

mtrace source

```
user@host> mtrace 192.168.4.2
Mtrace from 192.168.4.2 to 192.168.1.2 via group 0.0.0.0
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.168.1.2)
 -1  routerB.lab.mycompany.net (192.168.2.2)  PIM  thresh^ 1
 -2  routerC.lab.mycompany.net (192.168.3.2)  PIM  thresh^ 1
 -3  hostA.lab.mycompany.net (192.168.4.2)
Round trip time 2 ms; total ttl of 2 required.
```

Release Information

Command introduced before Junos OS Release 7.4.

mtrace from-source

IN THIS SECTION

- Syntax | 11013
- Description | 11013
- Options | 11014
- Required Privilege Level | 11015
- Output Fields | 11015
- Sample Output | 11016
- Release Information | 11017

Syntax

```
mtrace from-source source source
<brief | detail>
<extra-hops extra-hops>
<group group>
<interval interval>
<loop>
<max-hops max-hops>
<max-queries max-queries>
<multicast-response | unicast-response>
<no-resolve>
<no-router-alert>
<response response>
<routing-instance routing-instance-name>
<ttl ttl>
<wait-time wait-time>
```

Description

Display trace information about an IP multicast path from a source to this router or switch. If you specify a group address with this command, Junos OS returns additional information, such as packet rates and losses.

Options

brief | detail	(Optional) Display the specified level of output.
extra-hops *extra-hops*	(Optional) Number of hops to take after reaching a nonresponsive router. You can specify a number between **0** and **255**.
group *group*	(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.
interval *interval*	(Optional) Number of seconds to wait before gathering statistics again. The default value is **10** seconds.
loop	(Optional) Loop indefinitely, displaying rate and loss statistics.
max-hops *max-hops*	(Optional) Maximum hops to trace toward the source. The range of values is **0** through **255**. The default value is **32** hops.
max-queries *max-queries*	(Optional) Maximum number of query attempts for any hop. The range of values is **1** through **32**. The default is **3**.
multicast-response	(Optional) Always request the response using multicast.
no-resolve	(Optional) Do not attempt to display addresses symbolically.
no-router-alert	(Optional) Do not use the router-alert IP option.
response *response*	(Optional) Send trace response to a host or multicast address.
routing-instance *routing-instance-name*	(Optional) Trace a particular routing instance.
source *source*	Source hostname or address.
ttl *ttl*	(Optional) IP time-to-live (TTL) value. You can specify a number between **0** and **255**. Local queries to the multicast group use a value of **1**. Otherwise, the default value is **127**.
unicast-response	(Optional) Always request the response using unicast.
wait-time *wait-time*	(Optional) Number of seconds to wait for a response. The default value is **3**.

Required Privilege Level

view

Output Fields

Table 87 on page 11015 describes the output fields for the `mtrace from-source` command. Output fields are listed in the approximate order in which they appear.

Table 87: mtrace from-source Output Fields

Field Name	Field Description
Mtrace from	IP address of the receiver.
to	IP address of the source.
via group	IP address of the multicast group (if any).
Querying full reverse path	Indicates the full reverse path query has begun.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
protocol	Protocol used (for example, PIM).
Round trip time	Average round-trip time, in milliseconds (ms).
total ttl of	Time-to-live (TTL) threshold.
source	Source address.
Response Dest	Response destination address.

Table 87: mtrace from-source Output Fields *(Continued)*

Field Name	Field Description
Overall	Average packet rate for all traffic at each hop.
Packet Statistics for Traffic From	Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop.
Receiver	IP address receiving the multicast.
Query source	IP address sending the mtrace query.

Sample Output

mtrace from-source

```
user@host> mtrace from-source source 192.168.4.2 group 233.252.0.1
Mtrace from 192.168.4.2 to 192.168.1.2 via group 233.252.0.1
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.168.1.2)
 -1  routerB.lab.mycompany.net (192.168.2.2)  PIM  thresh^ 1
 -2  routerC.lab.mycompany.net (192.168.3.2)  PIM  thresh^ 1
 -3  hostA.lab.mycompany.net (192.168.4.2)
Round trip time 2 ms; total ttl of 2 required.

Waiting to accumulate statistics...Results after 10 seconds:

Source        Response Dest    Overall     Packet Statistics For Traffic From
192.168.4.2 192.168.1.2    Packet      192.168.4.2 To 233.252.0.1
     v       __/  rtt    2 ms     Rate       Lost/Sent = Pct  Rate
192.168.2.1
192.168.3.2  routerC.lab.mycompany.net
     v     ^      ttl    2                     0/0    = --    0 pps
192.168.4.1
192.168.2.2   routerB.lab.mycompany.net
     v      \__   ttl    3                     ?/0            0 pps
```

```
192.168.1.2   192.168.1.2
  Receiver      Query Source
```

Release Information

Command introduced before Junos OS Release 7.4.

mtrace monitor

IN THIS SECTION

- Syntax | 11017
- Description | 11017
- Options | 11017
- Required Privilege Level | 11018
- Output Fields | 11018
- Sample Output | 11018
- Release Information | 11019

Syntax

```
mtrace  monitor
```

Description

Listen passively for IP multicast responses. To exit the `mtrace monitor` command, type Ctrl+c.

Options

none Trace the master instance.

Required Privilege Level

view

Output Fields

Table 88 on page 11018 describes the output fields for the `mtrace monitor` command. Output fields are listed in the approximate order in which they appear.

Table 88: mtrace monitor Output Fields

Field Name	Field Description
`Mtrace query at`	Date and time of the query.
`by`	Address of the host issuing the query.
`resp to`	Response destination.
`qid`	Query ID number.
`packet from...to`	IP address of the query source and default group destination.
`from...to`	IP address of the multicast source and the response address.
`via group`	IP address of the group to trace.
`mxhop`	Maximum hop setting.

Sample Output

mtrace monitor

```
user@host> mtrace monitor
Mtrace query at Oct 22 13:36:14 by 192.168.3.2, resp to 233.252.0.32, qid 74a5b8
packet from 192.168.3.2 to 233.252.0.2
```

```
from 192.168.3.2 to 192.168.3.38 via group 233.252.0.1 (mxhop=60)

Mtrace query at Oct 22 13:36:17 by 192.681.3.2, resp to 233.252.0.32, qid 1d07ba
packet from 192.168.3.2 to 233.252.0.2
from 192.168.3.2 to 192.168.3.38 via group 233.252.0.1 (mxhop=60)

Mtrace query at Oct 22 13:36:20 by 192.681.3.2, resp to same, qid 2fea1d
packet from 192.168.3.2 to 233.252.0.2
from 192.168.3.2 to 192.168.3.38 via group 233.252.0.1 (mxhop=60)

Mtrace query at Oct 22 13:36:30 by 192.168.3.2, resp to same, qid 7c88ad
packet from 192.168.3.2 to 233.252.0.2
from 192.168.3.2 to 192.168.3.38 via group 233.252.0.1 (mxhop=60)
```

Release Information

Command introduced before Junos OS Release 7.4.

mtrace to-gateway

IN THIS SECTION

- Syntax | 11019
- Description | 11020
- Options | 11020
- Required Privilege Level | 11021
- Output Fields | 11021
- Sample Output | 11022
- Release Information | 11022

Syntax

```
mtrace to-gateway gateway gateway
<brief | detail>
```

```
<extra-hops extra-hops>
<group group>
<interface interface-name>
<interval interval>
<loop>
<max-hops max-hops>
<max-queries max-queries>
<multicast-response | unicast-response>
<no-resolve>
<no-router-alert>
<response response>
<routing-instance routing-instance-name>
<ttl ttl>
<unicast-response>
<wait-time wait-time>
```

Description

Display trace information about a multicast path from this router or switch to a gateway router or switch.

Options

gateway *gateway*	Send the trace query to a gateway multicast address.
brief | detail	(Optional) Display the specified level of output.
extra-hops *extra-hops*	(Optional) Number of hops to take after reaching a nonresponsive router or switch. You can specify a number between **0** and **255**.
group *group*	(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.
interface *interface-name*	(Optional) Source address for sending the trace query.
interval *interval*	(Optional) Number of seconds to wait before gathering statistics again. The default value is **10**.
loop	(Optional) Loop indefinitely, displaying rate and loss statistics.
max-hops *max-hops*	(Optional) Maximum hops to trace toward the source. You can specify a number between **0** and **255**.. The default value is **32**.

max-queries *max-queries*	(Optional) Maximum number of query attempts for any hop. You can specify a number between **0** and **255**. The default value is **3**.
multicast-response	(Optional) Always request the response using multicast.
no-resolve	(Optional) Do not attempt to display addresses symbolically.
no-router-alert	(Optional) Do not use the router-alert IP option.
response *response*	(Optional) Send trace response to a host or multicast address.
routing-instance *routing-instance-name*	(Optional) Trace a particular routing instance.
ttl *ttl*	(Optional) IP time-to-live value. You can specify a number between **0** and **225**. Local queries to the multicast group use TTL **1**. Otherwise, the default value is **127**.
unicast-response	(Optional) Always request the response using unicast.
wait-time *wait-time*	(Optional) Number of seconds to wait for a response. The default value is **3**.

Required Privilege Level

view

Output Fields

Table 89 on page 11021 describes the output fields for the `mtrace to-gateway` command. Output fields are listed in the approximate order in which they appear.

Table 89: mtrace to-gateway Output Fields

Field Name	Field Description
Mtrace from	IP address of the receiver.
to	IP address of the source.
via group	IP address of the multicast group (if any).

Table 89: mtrace to-gateway Output Fields *(Continued)*

Field Name	Field Description
Querying full reverse path	Indicates the full reverse path query has begun.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
protocol	Protocol used (for example, PIM).
Round trip time	Average round-trip time, in milliseconds (ms).
total ttl of	Time-to-live (TTL) threshold.

Sample Output

mtrace to-gateway

```
user@host> mtrace to-gateway gateway 192.168.3.2 group 233.252.0.1 interface 192.168.1.73
brief

Mtrace from 192.168.1.73 to 192.168.1.2 via group 233.252.0.1
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.1.1.2)
 -1  routerA.lab.mycompany.net (192.1.1.2)  PIM  thresh^ 1
 -2  routerB.lab.mycompany.net (192.1.2.2)  PIM  thresh^ 1
 -3  routerC.lab.mycompany.net (192.1.3.2)  PIM  thresh^ 1
Round trip time 2 ms; total ttl of 3 required.
```

Release Information

Command introduced before Junos OS Release 7.4.

mtu (QFX Series)

IN THIS SECTION

- Syntax | 11023
- Hierarchy Level | 11023
- Description | 11023
- Options | 11024
- Required Privilege Level | 11024
- Release Information | 11024

Syntax

```
mtu bytes;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-range interface-name]
```

Description

Specify the maximum transmission unit (MTU) size for the media. Changing the media MTU size causes an interface to be deleted and added again. On QFX3500, QFX3600, QFX5100, and OCX Series switches, either standalone or as part of the QFabric system, the maximum MTU value on an untagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus four, whereas the maximum MTU value on a tagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight. The maximum MTU value on an untagged or tagged packet transiting through an ingress 10-Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight.

Keep the following points in mind if you are configuring MTU size for jumbo frames on these special types of interfaces:

- For LAG interfaces—Configuring the jumbo MTU size on a link aggregation group (LAG) interface (**ae***x*) automatically configures the jumbo MTU size on the member links.
- For RVIs—Jumbo frames of up to 9216 bytes are supported on the routed VLAN interface (RVI), which is named **vlan**. The RVI functions as a logical router. To route jumbo data packets on the RVI, you must configure the jumbo MTU size on the member physical interfaces of the RVI and not on the RVI itself (the **vlan** interface). However, for jumbo control packets—for example, to ping the RVI with a packet size of 6000 bytes or more—you must explicitly configure the jumbo MTU size on the interface named **vlan** (the RVI). On a QFX5100 switch jumbo frames on the RVI are configured on the basis of the interface MTU.

> **NOTE**: RVIs are not supported on OCX Series switches.

> **CAUTION**: Setting or deleting the jumbo MTU size on the RVI (the **vlan** interface) while the switch is transmitting packets might result in dropped packets.

Options

bytes —MTU size.

- **Range:** 64 through 9216 bytes
- **Default:** 1514 bytes

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

> *Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches*
>
> Junos OS Network Interfaces Library for Routing Devices

mtu (Multilink and Link Services Logical Interface)

IN THIS SECTION

- Syntax | 11025
- Hierarchy Level | 11025
- Description | 11025
- Options | 11025
- Required Privilege Level | 11026
- Release Information | 11026

Syntax

```
mtu bytes;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family]
```

Description

Maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values.

Options

bytes—MTU size.

- **Range:** 0 through 5012 bytes

- **Default:** 1500 bytes (`inet`, `inet6`, and `iso` families), 1448 bytes (`mpls`)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring MRRU on Multilink and Link Services Logical Interfaces*
>
> Junos OS Network Interfaces Library for Routing Devices

mtu

IN THIS SECTION

- Syntax | **11026**
- Hierarchy Level | **11027**
- Description | **11027**
- Options | **11027**
- Required Privilege Level | **11027**
- Release Information | **11027**

Syntax

```
mtu size;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Specify the maximum transmission unit (MTU) size for the media in bytes. MTU size can be either 1500 bytes or 2048 bytes.

Options

size	Size of the MTU

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D47.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

mtu (interfaces)

IN THIS SECTION

Syntax | 11028

- Hierarchy Level | 11028
- Description | 11029
- Options | 11030
- Required Privilege Level | 11031
- Release Information | 11031

Syntax

```
mtu bytes;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number family family],
[edit interfaces interface-range name],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family],
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name backup-neighbor address],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name backup-neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn interface interface-name],
[edit protocols l2circuit local-switching interface interface-name backup-neighbor address],
[edit protocols l2circuit neighbor address interface interface-name]
[edit protocols l2circuit neighbor address interface interface-name backup-neighbor address],
[edit routing-instances routing-instance-name protocols l2vpn interface interface-name],
[edit logical-systems name protocols ospf area name interface],
[edit logical-systems name routing-instances name protocols ospf area
name interface],
[edit protocols ospf area name interface],
[edit routing-instances name protocols ospf area name interface]
```

Description

Specify the maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

If a packet whose size is larger than the configured MTU size is received on the receiving interface, the packet is eventually dropped. The value considered for MRU (maximum receive unit) size is also the same as the MTU size configured on that interface.

Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for:

- Management Ethernet interfaces (fxp0, em0, or me0)
- Service interfaces (ms, vms, or ams)
- Loopback, multilink, and multicast tunnel devices

MTU for IRB interfaces is calculated by removing the Ethernet header overhead `(6(DMAC)+6(SMAC)+2(EtherType))`, and the MTU is a minimum of the two values:

- Configured MTU
- Associated bridge domain's physical or logical interface MTU
  - For Layer 2 logical interfaces configured with `flexible-vlan-tagging`, IRB MTU is calculated by including 8 bytes overhead (`SVLAN+CVLAN`).
  - For Layer 2 logical interfaces configured with `vlan-tagging`, IRB MTU is calculated by including single VLAN 4 bytes overhead.

> **NOTE**: Changing the Layer 2 logical interface option from `vlan-tagging` to `flexible-vlan-tagging` or vice versa adjusts the logical interface MTU by 4 bytes with the existing MTU size. As a result, the Layer 2 logical interface is deleted and re-added, and the IRB MTU is re-computed appropriately.

On EX Series switches, to route jumbo data packets on an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI), you must configure the jumbo MTU size on the member physical interfaces of the VLAN that you have associated with the IRB interface or RVI, as well as on the IRB interface or RVI itself (the interface named `irb` or `vlan`, respectively).

> **CAUTION**: For EX Series switches, setting or deleting the jumbo MTU size on an IRB interface or RVI while the switch is transmitting packets might cause packets to be dropped.

On ACX Series routers, you can configure the protocol MTU by including the `mtu` statement at the `[edit interfaces` *`interface-name`* `unit` *`logical-unit-number`* `family inet]` or `[edit interfaces` *`interface-name`* `unit` *`logical-unit-number`* `family inet6]` hierarchy level.

- If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.
- If you are configuring the protocol MTU for both `inet` and `inet6` families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for `inet` and `inet6` families that are configured on the same logical interface.

For more information about configuring MTU for specific interfaces and router or switch combinations, see *Configure the Media MTU*.

Options

`bytes`—MTU size.

- **Range:** 256 through 9192 bytes
  - ACX710-48L, ACX710-32C, ACX7509, and ACX7024: 256 through 9216 bytes
  - EX Series switch interfaces: 256 through 9216
  - MX Series routers: 256 through 9500 bytes
  - PTX Series routers: 256 through 9500 bytes

  Starting in Junos OS Evolved Release 21.2R1, MTU 16KB is only for transiting traffic of WAN interfaces on the PTX10001-36MR, PTX10008 and PTX10004 routers. MTU is 9500B for protocols and 16KB for transit traffic.

  > **NOTE**: Starting in Junos OS Release 16.1R1, the MTU size for a media or protocol is increased from 9192 to 9500 for Ethernet interfaces on the following MX Series MPCs:
  >
  > - MPC1
  > - MPC2

    - MPC2E
    - MPC3E
    - MPC4E
    - MPC5E
    - MPC6E

- **Default:**
    - INET, INET6, and ISO families: 1500 bytes
    - EX Series switch interfaces: 1514 bytes
    - MPLS: 1448 bytes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for Layer 2 VPNs introduced in Junos OS Release 10.4.

Support at the`[set interfaces interface-name unit logical-unit-number family ccc]` hierarchy level introduced in Junos OS Release 12.3R3 for MX Series routers.

RELATED DOCUMENTATION

> *Configure the Media MTU*
>
> *Protocol MTU*
>
> *Configuring the MTU for Layer 2 Interfaces*

mtu-discovery

IN THIS SECTION

- Syntax | 11032
- Hierarchy Level | 11032
- Description | 11033
- Required Privilege Level | 11033
- Release Information | 11033

Syntax

```
mtu-discovery;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Configure TCP path maximum transmission unit (MTU) discovery.

TCP path MTU discovery enables BGP to automatically discover the best TCP path MTU for each BGP session. In Junos OS, TCP path MTU discovery is disabled by default for all BGP neighbor sessions.

When MTU discovery is disabled, TCP sessions that are not directly connected transmit packets of 512-byte maximum segment size (MSS). These small packets minimize the chances of packet fragmentation at a device along the path to the destination. However, because most links use an MTU of at least 1500 bytes, 512-byte packets do not result in the most efficient use of link bandwidth. For directly connected EBGP sessions, MTU mismatches prevent the BGP session from being established. As a workaround, enable path MTU discovery within the EBGP group.

Path MTU discovery dynamically determines the MTU size on the network path between the source and the destination, with the goal of avoiding IP fragmentation. Path MTU discovery works by setting the Don't Fragment (DF) bit in the IP headers of outgoing packets. When a device along the path has an MTU that is smaller than the packet, the device drops the packet. The device also sends back an ICMP Fragmentation Needed (Type 3, Code 4) message that contains the device's MTU, thus allowing the source to reduce its path MTU appropriately. The process repeats until the MTU is small enough to traverse the entire path without fragmentation.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Limiting TCP Segment Size for BGP

Configure Path MTU Discovery

mtu (Dynamic and Static PPPoE)

IN THIS SECTION

- Syntax | 11034
- Hierarchy Level | 11034
- Description | 11034
- Options | 11034
- Required Privilege Level | 11035
- Release Information | 11035

Syntax

```
mtu (size | use-lower-layer);
```

Hierarchy Level

```
[edit access group-profile group-profile-name ppp ppp-options]
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options],
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name"  unit “$junos-
interface-unit” ppp-options],
[edit interfaces pp0 unit unit-number ppp-options]
[edit interfaces si interface-id unit unit-number ppp-options]
```

Description

Specify the size of maximum transmission unit (MTU) for the PPP connection. For a PPP connection, the MTU size defines the largest data unit that can be forwarded without fragmentation. This size does not include the overhead of the lower layers.

Options

size	MTU size in bytes for a PPP connection.

- **Range:** 64–65,535

use-lower-layer Set the PPP MTU size to the interface MTU size excluding the overhead of the lower layers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring MTU and MRU for PPP Subscribers

Understanding MTU and MRU Configuration for PPP Subscribers

mtu (Dynamic Profiles)

IN THIS SECTION

- Syntax | 11036
- Hierarchy Level | 11036
- Description | 11036
- Options | 11036
- Required Privilege Level | 11036
- Release Information | 11036

Syntax

```
mtu mtu;
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit name family inet ],
[edit dynamic-profiles name logical-systems name interfaces name unit name family inet ],
[edit interfaces name unit name family inet ]
```

Description

Configure a per-subscriber maximum transmission unit (MTU) on a logical interface. A subscriber logical interface can specify MTU through its dynamic profile for the same underlying physical interfaces. The dynamic logical interface MTU must be no greater than the physical interface MTU minus the dynamic VLAN header size). You can configure either a static value or the predefined variable, $junos-interface-mtu. The value of the variable is returned by RADIUS during authentication in the Framed-MTU attribute (12). By default, the variable value is the MTU of the payload.

If you specify the predefined variable but RADIUS does not return a value in the Framed-MTU attribute,

Options

mtu MTU for the subscriber logical interface. You can specify a static value or the predefined variable, $junos-interface-mtu.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 18.1R1 on MX Series routers.

RELATED DOCUMENTATION

Understanding Per-subscriber Support of Maximum Transmission Unit for Dynamic Profiles

mtu-signaling

IN THIS SECTION

- Syntax | 11037
- Hierarchy Level | 11037
- Description | 11037
- Required Privilege Level | 11037
- Release Information | 11037

Syntax

```
mtu-signaling;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls path-mtu rsvp],
[edit protocols mpls path-mtu rsvp]
```

Description

Enable MTU signaling in RSVP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling MTU Signaling in RSVP

mtu (TDF Interface)

IN THIS SECTION

- Syntax | 11038
- Hierarchy Level | 11038
- Description | 11038
- Options | 11038
- Required Privilege Level | 11039
- Release Information | 11039

Syntax

```
mtu mtu-size;
```

Hierarchy Level

```
[edit interfaces mif]
```

Description

Configure the maximum transmission unit (MTU) size for the TDF interface.

Options

mtu-size MTU size.

- **Range:** 256 through 9192 bytes
- **Default:** 500 bytes (inet, inet6, and ISO families), 1448 bytes (MPLS)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a TDF Logical Interface

multi-address-embedded-option-response (DHCP Local Server)

IN THIS SECTION

- Syntax | 11039
- Hierarchy Level | 11040
- Description | 11040
- Default | 11040
- Required Privilege Level | 11040
- Release Information | 11040

Syntax

```
multi-address-embedded-option-response;
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 ...],
[edit logical-systems logical-system-name system services system services dhcp-local-server
dhcpv6 ...],
[edit routing-instances routing-instance-name system services system services dhcp-local-server
dhcpv6 ...]
```

Description

Configure DHCPv6 local server to return the DNS server address (DHCPv6 attribute 23) as a suboption in the respective IA_NA or IA_PD headers.

Default

DHCPv6 local server returns the DNS server address as a global DHCPv6 option.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.

Statement supported in Junos OS Release 13.3 and later releases. (Not supported in Junos OS Release 13.1 and Release 13.2.)

RELATED DOCUMENTATION

> *Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment*
>
> *Overriding the Default DHCP Local Server Configuration Settings*

multi-bgp-path

IN THIS SECTION

- Syntax | 11041
- Hierarchy Level | 11041
- Description | 11041
- Default | 11042
- Required Privilege Level | 11042
- Release Information | 11042

Syntax

```
multi-bgp-path;
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix template template-name]
[edit services flow-monitoring version9 template template-name]
```

Description

(Junos OS only) Enable multiple BGP path support for reporting an accurate BGP next-hop address for traffic load-balanced over multiple BGP peers in the ingress direction only. Without this statement, we report only the first in a list of next hops. To contain this accurate BGP next-hop address, we use the IPv4 BGP Nexthop Address (IE 18) field in the IPv4 and MPLS-IPv4 templates and the IPv6 BGP Nexthop Address (IE 63) field in the IPv6 and MPLS-IPv6 templates, for both the IPFIX and the version 9 formats. For IPv6 and MPLS-IPv6 flows, you also need to configure the `ipv6-extended-attrib` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level. When this feature is enabled, the `fragmentIdentification` (IE 54) field reports a value of 0.

For this feature to work, you must also configure the the `nexthop-learning enable` statement at the `[edit services flow-monitoring version version template template-name]` hierarchy level.

Existing flows must be manually cleared using the `clear services accounting flow inline-jflow fpc-slot` *`slot-number`* command when this feature is enabled or disabled, because common memory is used both for the `fragmentIdentification` (IE 54) field and for this feature. Otherwise, the export thread may wrongly misunderstand the fragment identifier of an existing flow as a next-hop address when enabled, or wrongly misunderstand the next-hop address of an existing flow as a fragment identifier when disabled. To avoid this situation, existing flows must be cleared from all FPC slots on which inline active flow monitoring is enabled.

When the number of BGP peers is 4000 or above, we cannot allocate an accurate BGP next-hop address for a new BGP peer. In such cases, a BGP next-hop address of zero is used for these BGP peers, and all flows destined to these BGP peers report the default BGP next-hop address, which is the first BGP next-hop address in a list of next hops. That is, this feature is not supported for these BGP peers with a BGP next-hop address of zero.

Default

Disabled. Without this statement, we report only the first in a list of BGP next hops.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

multi-chassis

IN THIS SECTION

- Syntax | 11043
- Hierarchy Level | 11043
- Description | 11043
- Required Privilege Level | 11043
- Release Information | 11043

Syntax

```
multi-chassis {
    multi-chassis-protection peer-ip-address {
        interface interface-name;
    }
    mc-lag consistency-check;
}
```

Hierarchy Level

```
[edit]
```

Description

Configure an interchassis link-protection link (ICL-PL) between the two peers that host a multichassis link aggregation group (MC-LAG). You can configure either an aggregated Ethernet interface or a 10-Gigabit Ethernet interface to be an ICL-PL.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

multi-chassis-protection

IN THIS SECTION

- Syntax | 11044
- Hierarchy Level | 11044
- Description | 11044
- Required Privilege Level | 11044
- Release Information | 11045

Syntax

```
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
```

Hierarchy Level

```
[edit multi-chassis]
```

Description

Configure multichassis link protection between the two peers that host a multichassis link aggregation group (MC-LAG). If the Interchassis Control Protocol (ICCP) connection is up and the interchassis link (ICL) comes up, the peer configured as standby brings up the multichassis aggregated Ethernet interfaces shared with the peer. Multichassis protection must be configured on one interface for each peer.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

multi-chassis-protection (Interfaces)

IN THIS SECTION

- Syntax | 11045
- Hierarchy Level | 11045
- Description | 11045
- Options | 11046
- Required Privilege Level | 11046
- Release Information | 11046

Syntax

```
multi-chassis-protection {
    peer a.b.c.d {
        interface interface-name;
        }
    }
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

For MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces, you can use this statement under the physical interface level to reduce the configuration at the logical interface level if the following assumption exists:

If there are n +1 logical interfaces under `ae0`, from `ae0.0` through `ae0.n`, there will be n + 1 logical interfaces under `ge-0/0/0` as well, from `ge-0/0/0.0` through `ge-0/0/0.n`, and each `ge-0/0/0` logical interface will be a protection link for the `ae0` logical interface.

> **NOTE**: A bridge domain cannot have MC-AE logical interfaces which belong to different redundancy groups.

If the Inter-Chassis Control Protocol (ICCP) connection is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer.

The remaining statements are explained separately. See CLI Explorer.

Options

`interface` ***`interface-name`***	Specify the interface: `interface` *`interface-name-fpc/pic/port`*

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers and QFX Series Switches
- *Configuring Aggregated Ethernet Link Protection*
- *Example: Configuring Aggregated Ethernet Link Protection*
- peer | **12339**

multi-destination (Class of Service)

IN THIS SECTION

- Syntax | 11047
- Hierarchy Level | 11047
- Description | 11047
- Options | 11048
- Required Privilege Level | 11048
- Release Information | 11048

Syntax

```
multi-destination {
    classifiers {
        (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) classifier-name;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
[edit class-of-service interfaces interface-name]
```

Description

Define a CoS behavior aggregate (BA) classifier for multi-destination, or BUM (broadcast, unknown unicast, and multicast), traffic.

There is no default multi-destination classifier. Without explicit multi-destination classifier configuration, all BUM traffic is sent to the default multicast queue (queue 8), which receives 20 percent of the bandwidth by default. If you apply a multi-destination classifier to send BUM traffic to other multicast queues (queue 9 or queue 10), you must ensure the egress scheduler allocates bandwidth for those queues.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers

Assigning CoS Components to Interfaces

Understanding CoS Classifiers

multi-domain

IN THIS SECTION

- Syntax | 11049
- Hierarchy Level | 11049
- Description | 11049
- Options | 11050
- Required Privilege Level | 11050
- Release Information | 11050

Syntax

```
multi-domain {
    max-data-session max-data-sessions;
    packet-action (drop-and-log | shutdown);
    recovery-timeout seconds;
}
```

Hierarchy Level

```
[edit logical-systems name protocols          dot1x              authenticator
interface],
[edit protocols dot1x   authenticator              interface]
```

Description

Configure multi-domain authentication to restrict the number of authenticated data and VoIP sessions on the port. Multi-domain authentication is an extension of multiple supplicant mode for 802.1X authentication, and is designed to support VoIP and data clients on the same interface. The interface is divided into two domains; one is the data domain and the other is the voice domain.

In multiple supplicant mode, any number of VoIP or data sessions can be authenticated; the number of sessions can be restricted using MAC limiting, but there is no way to apply the limit specifically to either data or VoIP sessions. Multi-domain authentication maintains separate session counts based on the domain type.

The data device can be authenticated using 802.1X authentication or MAC RADIUS authentication. Multi-domain authentication does not enforce the order of authentication. For best results, the VoIP device should be authenticated before the data device.

You can configure the maximum number of authenticated data sessions allowed on the interface using the `max-data-session` statement. The number of VoIP sessions is not configurable; only one authenticated VoIP session is allowed.

If a new client attempts to authenticate on the interface after the maximum session count has been reached, the default action is to drop the packet and generate an error log message. You can also configure the action to shut down the interface. The port can be manually recovered from the down state by issuing the `clear dot1x recovery-timeout` command, or can recover automatically after a recovery timeout period. To configure automatic recovery, use the `recovery-timeout` option.

Options

max-data-session *max-data-sessions*

The maximum number of authenticated data sessions allowed in the data domain on the 802.1X-enabled interface.

- **Range:** 1 through 1,000 sessions
- **Default:** 1

packet-action (drop-and-log | shutdown)

Specify the action the device should take on packets that exceed the limit of authenticated sessions allowed on the interface. The limit for data sessions is configured using the `max-data-session` option. The number of VoIP sessions is not configurable; only one authenticated VoIP session is allowed.

- **Values:** Specify one of the following:
  - drop-and-log—Drop the packet and generate an error syslog message.
  - shutdown—Shut down the interface.
- **Default:** drop-and-log

recovery-timeout *seconds*

If you configure the packet action with the shutdown option and you configure the recovery timeout, the interface is temporarily disabled when the maximum number of authenticated sessions is reached. The interface will recover automatically after the number of seconds specified.

- **Range:** 60 through 3600 seconds
- **Default:** none

> **NOTE**: The default values of these optional parameters will be enforced when multi-domain is configured.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

authenticator | 2673

dot1x | 5917

interface (802.1X) | 8819

multi-homing (VPLS Multihoming for BGP VPLS)

IN THIS SECTION

- Syntax | 11051
- Hierarchy Level | 11051
- Description | 11051
- Required Privilege Level | 11052
- Release Information | 11052

Syntax

```
multi-homing;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
site site-name],
[edit routing-instances routing-instance-name protocols vpls site site-name]
```

Description

Configuration of this statement tracks BGP peers, such as to prevent isolation of the PE router from the core or split brain.. If no BGP peer is available, all active interfaces for a site are deactivated.

When identical site IDs are used without configuring multihoming, a collision log message is generated at each signaling: `RPD_L2VPN_SITE_COLLISION: Same site ID 2 configured on remote PE (RD 8.8.8.1:1013:) and local PE in VPN 1013 (non-multihomed site 2)`. This is expected behavior.

> **NOTE**: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configuring Multihoming on the PE Router

multi-homing (VPLS Multihoming for FEC 129)

IN THIS SECTION

- Syntax | 11053
- Hierarchy Level | 11053
- Description | 11053
- Required Privilege Level | 11053
- Release Information | 11054

Syntax

```
multi-homing {
    peer-active;
    site site-name {
        active-interface interface-name {
            any;
            primary interface-name;
        }
        identifier identifier;
        interface interface-name {
            preference preference-value;
        }
        peer-active;
        preference (preference-value | backup | primary);
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls],
[edit routing-instances instance-name protocols vpls]
```

Description

For VPLS autodiscovery (FEC 129), specify the parameters for multihoming to two or more provider edge (PE) routers.

> **NOTE**: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring VPLS Multihoming (FEC 129)

multi-link-layer-2-inline

IN THIS SECTION

- Syntax | 11054
- Hierarchy Level | 11054
- Description | 11054
- Required Privilege Level | 11055
- Release Information | 11055

Syntax

```
multi-link-layer-2-inline;
```

Hierarchy Level

```
[edit chassis fpc number pic number]
```

Description

Enable inline Layer 2 bundling services.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Inline MLPPP for WAN Interfaces Overview

Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces

Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces

multi-sa

IN THIS SECTION

- Syntax | 11055
- Hierarchy Level | 11056
- Description | 11056
- Options | 11056
- Required Privilege Level | 11056
- Release Information | 11056

Syntax

```
multi-sa {
    forwarding-class expedited-forwarding | assured-forwarding | best-effort | network-control;
}
```

Hierarchy Level

```
[edit security ipsec vpn]
```

Description

Negotiate multiple security association (SAs) based on configuration choice. Multiple SAs negotiates with the same traffic selector on the same IKE SA. By negotiating multiple SAs, the peer gateways have more replay windows. If the peer gateways create separate multiple SAs for the configured Forwarding-Classes (FC), then potentially a separate anti-replay window is available for each FC value. With this mapping, even if CoS can reorder packets, reordering is done with in a given multiple SA, thus avoiding packets drop due to the anti-replay checks.

Options

forwarding-class Forwarding classes (FCs) allow you to group packets for transmission and to assign packets to output queues.

- Values:
  - `expedited-forwarding`—Provides a low-loss, low-latency, low-jitter, assured-bandwidth, end-to-end service.
  - `assured-forwarding`—Provides a group of values you can define and includes four subclasses—AF1, AF2, AF3, and AF4—each with three drop probabilities (low, medium, and high).
  - `best-effort`—Provides no service profile. For the BE forwarding class, loss priority is typically not carried in a class-of-service (CoS) value, and random early detection (RED) drop profiles are more aggressive.
  - `network-control`—This class is typically high priority because it supports protocol control.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show security ipsec security-associations | 24701

Understanding CoS-Based IPsec VPNs with Multiple IPsec SAs

Example: Configuring CoS-Based IPsec VPNs

vpn (Security) | 30946

multi-tenant-hash

IN THIS SECTION

- Syntax | 11057
- Hierarchy Level | 11058
- Description | 11058
- Options | 11058
- Required Privilege Level | 11058
- Release Information | 11058

Syntax

```
multi-tenant-hash {
file-hash-key {
ascii-text;
hexa-decimal;
}
hash-method {

ascii-text;
hexa-decimal;
}
}
```

Hierarchy Level

```
[edit services web-filter]
```

Description

Configure the hash key and hash method to validate the file-hash present at the end of the hashed indicator file.

When `multi-tenant-hash` is configured, It indicates that the global DNS feed file consists of only encrypted feeds. When `multi-tenant-hash` is not configured, it indicates that the feeds present in the global DNS feed file are in plain text format.

Options

file-hash-key Define the web-filter global file key.

hash-method Define the authentication algorithm.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

Support added in Junos OS Release 21.1R1 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

multi-tenant-support

IN THIS SECTION

- Syntax | 11059
- Hierarchy Level | 11059
- Description | 11059
- Required Privilege Level | 11059
- Release Information | 11059

Syntax

```
multi-tenant-support
```

Hierarchy Level

```
[edit services service-set service-set-name web-filter
```

Description

Configure domain-names and actions for multiple tenants such that domain feeds can be managed per tenant basis. Thus, you can achieve hierarchical management of the domain feeds per customer (Profile), partitioning the customer into multiple entities (Templates), and further partitioning IP-addresses within that entity into subnets/CIDR/Single IP (Term).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

Support added in Junos OS Release 21.1R1 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

multi-bgp-path

IN THIS SECTION

- Syntax | 11060
- Hierarchy Level | 11060
- Description | 11060
- Default | 11061
- Required Privilege Level | 11061
- Release Information | 11061

Syntax

```
multi-bgp-path;
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix template template-name]
[edit services flow-monitoring version9 template template-name]
```

Description

(Junos OS only) Enable multiple BGP path support for reporting an accurate BGP next-hop address for traffic load-balanced over multiple BGP peers in the ingress direction only. Without this statement, we

report only the first in a list of next hops. To contain this accurate BGP next-hop address, we use the IPv4 BGP Nexthop Address (IE 18) field in the IPv4 and MPLS-IPv4 templates and the IPv6 BGP Nexthop Address (IE 63) field in the IPv6 and MPLS-IPv6 templates, for both the IPFIX and the version 9 formats. For IPv6 and MPLS-IPv6 flows, you also need to configure the `ipv6-extended-attrib` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level. When this feature is enabled, the `fragmentIdentification` (IE 54) field reports a value of 0.

For this feature to work, you must also configure the the `nexthop-learning enable` statement at the `[edit services flow-monitoring version version template template-name]` hierarchy level.

Existing flows must be manually cleared using the `clear services accounting flow inline-jflow fpc-slot slot-number` command when this feature is enabled or disabled, because common memory is used both for the `fragmentIdentification` (IE 54) field and for this feature. Otherwise, the export thread may wrongly misunderstand the fragment identifier of an existing flow as a next-hop address when enabled, or wrongly misunderstand the next-hop address of an existing flow as a fragment identifier when disabled. To avoid this situation, existing flows must be cleared from all FPC slots on which inline active flow monitoring is enabled.

When the number of BGP peers is 4000 or above, we cannot allocate an accurate BGP next-hop address for a new BGP peer. In such cases, a BGP next-hop address of zero is used for these BGP peers, and all flows destined to these BGP peers report the default BGP next-hop address, which is the first BGP next-hop address in a list of next hops. That is, this feature is not supported for these BGP peers with a BGP next-hop address of zero.

Default

Disabled. Without this statement, we report only the first in a list of BGP next hops.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

multicast

IN THIS SECTION

- Syntax | 11062
- Hierarchy Level | 11062
- Description | 11062
- Default | 11062
- Required Privilege Level | 11063
- Release Information | 11063

Syntax

```
multicast;
```

Hierarchy Level

```
[edit ethernet-switching-options storm-control (switch) interface (all | interface-name)]
```

Description

Enable storm control for all multicast traffic (both registered multicast and unregistered multicast) for the specified interface or for all interfaces.

Default

- On EX2200, EX3200, and EX4200 switches—Storm control does not apply to multicast traffic by default.
- On EX4500 and EX8200 switches—Storm control is enabled for multicast traffic.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Enabling and Disabling Storm Control (non-ELS)

multicast (routing-options)

IN THIS SECTION

- Syntax | 11063
- Hierarchy Level | 11065
- Description | 11065
- Required Privilege Level | 11065
- Release Information | 11065

Syntax

```
multicast {
    asm-override-ssm;
    backup-pe-group group-name {
        backups [ addresses ];
        local-address address;
    }
    cont-stats-collection-interval interval;
    flow-map flow-map-name {
        bandwidth (bps | adaptive);
```

```
        forwarding-cache {
            timeout (never non-discard-entry-only | minutes);
        }
        policy [ policy-names ];
        redundant-sources [ addresses ];
    }
    forwarding-cache {
        threshold suppress value <reuse value>;
        timeout minutes;
    }
    interface interface-name {
        enable;
        maximum-bandwidth bps;
        no-qos-adjust;
        reverse-oif-mapping {
            no-qos-adjust;
        }
        subscriber-leave-timer seconds;
    }
    local-address address
    omit-wildcard-address
    pim-to-igmp-proxy {
        upstream-interface [ interface-names ];
    }
    pim-to-mld-proxy {
        upstream-interface [ interface-names ];
    }
    rpf-check-policy [ policy-names ];
    scope scope-name {
        interface [ interface-names ];
        prefix destination-prefix;
    }
    scope-policy [ policy-names ];
    ssm-groups [ addresses ];
    ssm-map ssm-map-name {
        policy [ policy-names ];
        source [ addresses ];
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <disable>;
```

```
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-options],
[edit dynamic-profiles profile-name routing-instances routing-instance-name routing-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Configure multicast routing options properties. Note that you cannot apply a scope policy to a specific routing instance. That is, all scoping policies are applied to all routing instances. However, the `scope` statement does apply individually to a specific routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

interface and **maximum-bandwidth** statements introduced in Junos OS Release 8.3.

interface and **maximum-bandwidth** statements introduced in Junos OS Release 9.0 for EX Series switches.

Statement added to **[edit dynamic-profiles routing-options]** and **[edit dynamic-profiles *profile-name* routing-instances *routing-instance-name* routing-options]** hierarchy levels in Junos OS Release 9.6.

RELATED DOCUMENTATION

Examples: Configuring Administrative Scoping

Example: Configuring the Multicast Forwarding Cache

Example: Configuring a Multicast Flow Map

Example: Configuring Source-Specific Multicast Groups with Any-Source Override

indirect-next-hop

multicast-dlci

IN THIS SECTION

- Syntax | **11066**
- Hierarchy Level | **11066**
- Description | **11066**
- Options | **11067**
- Required Privilege Level | **11067**
- Release Information | **11067**

Syntax

```
multicast-dlci dlci-identifier;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For point-to-multipoint link services interfaces only, enable multicast support on the interface. You can configure multicast support on the interface if the Frame Relay switch performs multicast replication.

Options

dlci-identifier—DLCI identifier, a number from 16 through 1022 that defines the Frame Relay DLCI over which the switch expects to receive multicast packets for replication.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Configuring Multicast-Capable DLCIs for MLFR FRF.16 Bundles*
- Configuring a Multicast-Capable Frame Relay Connection
- dlci
- multipoint-destination
- Junos OS Services Interfaces Library for Routing Devices

multicast (Dynamic Routing Options)

IN THIS SECTION

- Syntax | 11068
- Hierarchy Level | 11068
- Description | 11068
- Required Privilege Level | 11068
- Release Information | 11068

Syntax

```
multicast {
    interface interface-name {
        no-qos-adjust;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-options],
[edit dynamic-profiles profile-name routing-instances routing-instance-name routing-options]
```

NOTE: You cannot apply a scope policy to a specific routing instance. That is, all scoping policies are applied to all routing instances. However, the `scope` statement does apply individually to a specific routing instance.

Description

Dynamically configure interface-specific multicast routing options properties in a dynamic client profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring the Multicast Forwarding Cache

Example: Configuring a Multicast Flow Map

Example: Configuring Source-Specific Multicast Groups with Any-Source Override

multicast-group

IN THIS SECTION

- Syntax | 11069
- Hierarchy Level | 11069
- Description | 11069
- Required Privilege Level | 11069
- Release Information | 11070

Syntax

```
multicast-group address;
```

Hierarchy Level

```
[edit vlans vlan-name vxlan]
```

Description

Assign a multicast group address to a Virtual Extensible LAN (VXLAN). All members of a VXLAN must use the same multicast group address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Understanding VXLANs

Manually Configuring VXLANs on QFX Series and EX4600 Switches

Examples: Manually Configuring VXLANs on QFX Series and EX4600 Switches

multicast-ipv4

IN THIS SECTION

- Syntax | 11070
- Hierarchy Level | 11070
- Description | 11071
- Options | 11071
- Required Privilege Level | 11071
- Release Information | 11071

Syntax

```
multicast-ipv4 {
    source-ip source-ip-address;
    dest-ip-multicast-group ipv4-mcast-address;
}
```

Hierarchy Level

```
[edit protocols oam fabric flow-specs flow-specification-name]
```

Description

Configure multicast parameters for the IPv4 Ethernet protocol type in the internal fabric monitoring flow specifications.

Options

dest-ip-multicast-group ***ipv4-mcast-address***

IPv4 multicast group address of the destination fabric maintenance endpoint (FMEP) that is configured in the flow specification. Configuration of this parameter is mandatory.

- **Syntax:** Dotted decimal notation without the prefix

source-ip ***source-ip-address***

(Optional) IPv4 multicast address of the source fabric maintenance endpoint (FMEP) that is configured in the flow specification.

- **Syntax:** Dotted decimal notation without the prefix

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | **29924**

multicast-loadbalance

IN THIS SECTION

- Syntax | 11072
- Hierarchy Level | 11072
- Description | 11072
- Default | 11073
- Options | 11073
- Required Privilege Level | 11073
- Release Information | 11073

Syntax

```
multicast-loadbalance {
    disable (Multicast Load Balancing);
    hash-mode mode;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

(EX8200 switches only) Enable multicast load balancing to evenly distribute the Layer 3 routed multicast traffic on aggregated 10-gigabit Ethernet interfaces. When multicast load balancing is enabled on the EX8200 switch, it takes effect on aggregated 10-gigabit Ethernet interfaces that are currently configured as well as those that are configured later.

Crc-sgip mode is the default. If the default mode is implemented and the Layer 3 routed multicast traffic on the aggregated interfaces is not evenly distributed, you can try the other hashing algorithms. Use the algorithm that best balances your Layer 3 routed multicast traffic.

> **NOTE**: If you change the hashing algorithm when multicast load balancing is disabled, the new algorithm takes effect after you reenable multicast load balancing.

Default

Multicast load balancing is enabled.

Options

hash-mode ***mode***—Specify one of these hashing algorithms:

balanced Use a round-robin method to determine the aggregated Ethernet links with the least amount of traffic.

crc-gip Use multicast packets' group IP address as a key to a CRC algorithm to get the hash index. The hashing algorithm used is CRC.

crc-sgip Use multicast packets' source IP address and group IP address for hashing. The hashing algorithm used is CRC.

crc-sip Use multicast packets' source IP address for hashing. The hashing algorithm used is CRC.

simple-gip Use multicast packets' group IP address for hashing. The hashing algorithm used is XOR.

simple-sgip Use multicast packets' GIP bits to calculate the hash index. The hashing algorithm used is XOR.

simple-sip Use multicast packets' source IP address bits for hashing. The hashing algorithm used is XOR.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches

Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure)

multicast-mode (EVPN)

IN THIS SECTION

- Syntax | 11074
- Hierarchy Level | 11074
- Description | 11074
- Options | 11075
- Required Privilege Level | 11075
- Release Information | 11075

Syntax

```
multicast-mode client | ingress-replication;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols evpn],
[edit routing-instances routing-instance-name protocols evpn]
```

Description

Configure the multicast server mode for delivering traffic and packets for Ethernet VPN (EVPN). This statement is required for a VXLAN EVPN instance.

> **NOTE**: If you configure the `multicast-mode` statement, then you must also configure the `encapsulation (protocols evpn) vxlan` statement.

Options

client Use the client as the multicast mode for delivering traffic and multicast packets across routers and switches.

ingress-replication Use ingress replication as the multicast mode for delivering broadcast, unknown unicast, and multicast (BUM) traffic and multicast packets across routers and switches.

- **Default:** ingress-replication

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D30.

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

Support for primary and secondary asymmetry on MX Series routers added in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Understanding EVPN with VXLAN Data Plane Encapsulation

EVPN Over VXLAN Encapsulation Configuration Overview for QFX Series and EX4600 Switches

Using a Default Layer 3 Gateway to Route Traffic in an EVPN-VXLAN Overlay Network

Example: Configuring an EVPN Control Plane and VXLAN Data Plane

multicast-only

IN THIS SECTION

- Syntax | 11076
- Hierarchy Level | 11076
- Description | 11076
- Required Privilege Level | 11076
- Release Information | 11076

Syntax

```
multicast-only;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet]
```

Description

Configure the unit and family so that the interface can transmit and receive multicast traffic only. You can configure this property on the IP family only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Restricting Tunnels to Multicast Traffic

tunnel | 30192

multicast (Protocols Router Discovery)

IN THIS SECTION

- Syntax | 11077
- Hierarchy Level | 11077
- Description | 11077
- Default | 11078
- Required Privilege Level | 11078
- Release Information | 11078

Syntax

```
(multicast | broadcast);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery address address],
[edit protocols router-discovery address address]
```

Description

Specify when the server should include the IP addresses in router advertisement packets. On the same physical interfaces, some addresses might be included only in multicast packets, while others might be included only in broadcast packets.

If you specify **multicast**, the server includes the addresses in router advertisement packets only if the packets are multicast. If the router supports IP multicast, and if the interface supports IP multicast,

multicast is the default. Otherwise, the addresses are included in broadcast router advertisement packets. If the router does not support IP multicast, the addresses are not included.

Default

multicast if the router supports IP multicast; **broadcast** if the router does not support IP multicast.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

broadcast (Router Discovery) | **3043**

multicast (QFabric Routing Options)

IN THIS SECTION

- Syntax | **11079**
- Hierarchy Level | **11079**
- Description | **11079**
- Required Privilege Level | **11079**
- Release Information | **11079**

Syntax

```
multicast {
    fabric-optimized-distribution
    no-make-before-break;
}
```

Hierarchy Level

```
[edit fabric routing-options]
```

Description

(QFabric systems only) Set multicast routing options in a QFabric system, such as "no-make-before-break" on page 11508.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

> *Understanding the QFabric System Data Plane*

multicast-replication

IN THIS SECTION

- Syntax | 11080
- Hierarchy Level | 11080
- Description | 11080
- Default | 11081
- Options | 11081
- Required Privilege Level | 11084
- Release Information | 11084

Syntax

```
multicast-replication {
    evpn {
        irb (local-only | local-remote | oism | enhanced-oism);
        smet-forwarding;
        smet-nexthop-limit smet-nexthop-limit;
    }
    ingress;
    local-latency-fairness;
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure the mode of multicast replication that helps to optimize multicast latency.

> **NOTE**: The `multicast-replication` statement is supported only on platforms with the `enhanced-ip` mode enabled.

Default

This statement is disabled by default.

Options

> **NOTE**: The options in the evpn hierarchy are specific to EVPN configurations. The `ingress` and the `local-latency-fairness` options do not apply to EVPN configurations.

`evpn irb (local-only | local-remote | oism | enhanced-oism)`

Use the `evpn irb` option to enable inter-VLAN multicast forwarding in an EVPN-VXLAN network in one of the following modes (these options are mutually exclusive):

- `local-only`—Use this mode with an edge-routed bridging (ERB) overlay, also known as a collapsed IP fabric. In this mode, the PFE on the leaf devices in the fabric performs local multicast routing at the fabric edge. The spine devices, also called *lean spines*, primarily act as IP transit devices for the fabric.

- `local-remote`—Use this mode with a centrally-routed bridging (CRB) overlay, also known as a two-layer IP fabric. In this mode, the spine devices in the fabric centrally route the multicast traffic between VLANS. The spine devices forward the routed VLAN traffic into the EVPN core toward interested receivers. The spine devices use a PIM designated router (DR) to avoid duplicating packets into the core. The leaf devices forward multicast traffic received on a VLAN to their receivers on that VLAN.

- `oism`—Use this mode with an ERB overlay to support routing multicast traffic inside the fabric as well as to and from external devices. This mode implements *optimized inter-subnet multicast* (OISM) according to the IETF draft https://tools.ietf.org/html/draft-ietf-bess-evpn-irb-mcast. In this mode, the leaf devices operate in local-remote mode like in a CRB overlay design, while at the same time, the device performs local multicast routing. This mode enables the original regular OISM mode that uses a symmetric bridge domains (VLANs) OISM model. With regular OISM mode, you must configure all revenue bridge domains (VLANs) symmetrically on all OISM devices.

OISM also supports:

- Selective forwarding on the source VLAN.
- Routing multicast traffic through an external PIM domain.

> **NOTE**: You can alternatively enable the enhanced version of OISM, if all of the OISM devices in your network support the enhanced OISM mode option. See the `enhanced-oism` option description (next) for information on when you might use that option instead of this original OISM mode option.

- `enhanced-oism`—This mode implements an enhanced version of OISM for EVPN multicast with ERB overlays. Enhanced OISM doesn't require you to configure all revenue bridge domains (VLANs) in the network on all OISM devices. On each device, you can configure only the revenue VLANs the device hosts. As a result, we say this mode uses an *asymmetric bridge domains (VLANs)* model. On supported devices, this mode enables OISM to scale well when your network has leaf devices that host larger numbers of different VLANs.

  This mode has some operational differences and small configuration differences to support the enhanced OISM asymmetric bridge domains model. As a result, if you enable OISM, all OISM devices in the network must use the same OISM mode option, either `oism` or `enhanced-oism`.

  You can choose to use the `enhanced-oism` option instead of the regular `oism` option if all OISM devices in the network support enhanced OISM. In that case, you might want to use enhanced OISM when:

  - Your network has a large number of revenue bridge domains (VLANs), and resources might be strained on some devices to configure all the VLANs there.
  - Your network has a large number of disjointed bridge domains (VLANs) in the network (different devices host different sets of VLANs).
  - On OISM devices in your network, you don't have policies configured that are based on the source MAC address of the packets. If you do have source MAC address policies, use the `oism` option in your network instead.

  > **NOTE**: You should use regular OISM instead of enhanced OISM if your network needs to pass multicast packets with stringent

requirements for decrementing the time-to-live (TTL) field. The enhanced OISM model inherently has a limitation where packets with TTL=1 will not reach receivers on devices that are not multihoming peers of the source device. Regular OISM forwards source traffic on the source VLAN and doesn't decrement the TTL value for destinations on the same VLAN.

See *Summary of Enhanced OISM Differences* for details.

Enhanced OISM also supports:

- Selective forwarding on the source VLAN.
- Routing multicast traffic through an external PIM domain.

Default EVPN multicast IRB mode:`evpn irb local-remote`

You can enable only one of these modes at a time, and you must configure the same mode on all devices in the network that process multicast traffic.

> **NOTE**: Only `local-remote`, `oism`, and `enhanced-oism` modes support forwarding with selective multicast Ethernet tag (SMET) routes (EVPN Type 6 routes).

`evpn smet-forwarding`

Enables SMET forwarding in an EVPN-MPLS network.

`evpn smet-nexthop-limit` *`smet-nexthop-limit`*

Configures a limit for the number of SMET next hops for selective multicast forwarding in an EVPN network. A PE device uses the SMET next hops list of outgoing interfaces to selectively replicate and forward multicast traffic. When the list of SMET next hops reaches this limit, the PE device stops adding new SMET next hops. At that point, the PE device sends new multicast group traffic to all egress devices.

- **Range:** 10,000 through 40,000
- **Default:** 10,000

`ingress`

Complete ingress replication of the multicast data packets where all the egress Packet Forwarding Engines receive packets from the ingress Packet Forwarding Engines directly.

`local-latency-fairness`

Complete parallel replication of the multicast data packets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

`evpn` stanza introduced in Junos OS Release 17.3R3 for QFX Series switches.

`oism` option in `evpn irb stanza` introduced in Junos OS Release 21.2R1.

`smet-forwarding` option in `evpn` stanza introduced in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

forwarding-options (Routing Instance) | **7463**

IPv4 Inter-VLAN Multicast Forwarding Modes for EVPN-VXLAN Overlay Networks

Optimized Inter-Subnet Multicast in EVPN Networks

multicast-router-interface (IGMP Snooping)

IN THIS SECTION

- Syntax | **11085**
- Hierarchy Level | **11085**
- Description | **11085**
- Default | **11085**
- Required Privilege Level | **11085**
- Release Information | **11086**

Syntax

```
multicast-router-interface;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-
name],
[edit protocols igmp-snooping vlan (all | vlan-name) interface (all | interface-name)],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
snooping interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols vlan
vlan-id igmp-snooping interface interface-name]
[edit protocols igmp-snooping vlan vlan-name interface interface-name]
```

Description

Statically configure the interface as an IGMP snooping multicast-router interface—that is, an interface that faces toward a multicast router or other IGMP querier.

> **NOTE**: If the specified interface is a trunk port, the interface becomes a multicast-routing device interface for all VLANs configured on the trunk port. In addition, all unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded to the multicast routing device interface, even if the interface is configured as a multicast routing device interface only for IGMP snooping.
>
> Configure an interface as a bridge interface toward other multicast routing devices.

Default

Disabled. If this statement is disabled, the interface drops IGMP messages it receives.

The interface can either be a host-side or multicast-routing device interface.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping

IGMP Snooping in MC-LAG Active-Active Mode

host-only-interface | **8154**

multicast-router-interface (MLD Snooping)

IN THIS SECTION

- Syntax | **11086**
- Hierarchy Level | **11086**
- Description | **11087**
- Required Privilege Level | **11087**
- Release Information | **11087**

Syntax

```
multicast-router-interface;
```

Hierarchy Level

```
[edit protocols mld-snooping vlan  (all | vlan-name) interface (all | interface-name)],
[edit routing-instances instance-name protocols mld-snooping vlan vlan-name interface interface-
name]
```

Description

Statically configure the interface as a multicast-router interface—that is, an interface that faces towards a multicast router or other MLD querier.

> **NOTE**: If the specified interface is a trunk port, the interface becomes a multicast-router interface for all VLANs configured on the trunk port. In addition, all unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded to the multicast router interface, even if the interface is configured as a multicast-router interface only for MLD snooping.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

Support at the `[edit routing-instances instance-name protocols mld-snooping vlan vlan-name interface interface-name]` hierarchy level introduced in Junos OS Release 13.3 for EX Series switches.

RELATED DOCUMENTATION

Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)

multicast-snooping-options

IN THIS SECTION

- Syntax | 11088
- Hierarchy Level | 11088
- Description | 11088
- Options | 11089

Required Privilege Level | 11089
Release Information | 11089

Syntax

```
multicast-snooping-options {
    flood-groups [ ip-addresses ];
    forwarding-cache {
        threshold suppress value <reuse value>;
    }
    host-outbound-traffic (Multicast Snooping) {
        forwarding-class class-name;
        dot1p number;
    }
    graceful-restart <restart-duration seconds>;
    ignore-stp-topology-change;

    nexthop-hold-time milliseconds;
    oism {
        install-star-g-routes;
        conserve-mcast-routes-in-pfe;
     }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Establish multicast snooping option values.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Multicast Snooping

Enabling Bulk Updates for Multicast Snooping

Example: Configuring Multicast Snooping

multicast-statistics

IN THIS SECTION

- Syntax | 11090
- Hierarchy Level | 11090
- Description | 11090
- Default | 11090
- Required Privilege Level | 11090
- Release Information | 11090

Syntax

```
multicast-statistics;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

For Ethernet, SONET, aggregated Ethernet, and aggregated SONET interfaces in T Series or TX Matrix routers, specify support for multicast statistics on a physical interface to enable multicast accounting for all the logical interfaces below the physical interface.

Default

not enabled—must be configured to enable

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 10.2.

RELATED DOCUMENTATION

- Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces
- Configuring Multicast Statistics Collection on Aggregated SONET Interfaces
- *Configuring Multicast Statistics Collection on Ethernet Interfaces*
- Configuring Multicast Statistics Collection on SONET Interfaces

multicast-statistics (packet-forwarding-options)

IN THIS SECTION

- Syntax | 11091
- Hierarchy Level | 11091
- Description | 11091
- Required Privilege Level | 11092
- Release Information | 11092

Syntax

```
multicast-statistics;
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

Counts packets and checks the bandwidth of IPv4 and IPv6 multicast traffic received from a host and group in a routing instance by using firewall filters.

With `multicast-statistics` enabled, route statistics are updated by a firewall counter for the next 512 multicast routes. Statistics are attached and collected on a first-come, first-served basis. To count the packets and bandwidth, the switch uses ingress filters to match on the source IP, destination IP and VRF ID fields. These filters reside in an ingress filter processor (IFP) group that contains a list of routes and their corresponding filter IDs.

When using this command, consider the following:

- You cannot configure filters for reserved multicast addresses.
- The multicast statistic group is the group with the least priority. If there's a rule conflict in another group, the action for the group with the higher priority takes effect.

- Each route takes up one entry in the IFP ternary content-addressable memory (TCAM). If no TCAM space is available, the filter installation fails.
- If you delete this command, any installed firewall rules for multicast statistics are deleted. If you delete a route, the corresponding filter entry is also deleted. When you delete the last entry, the group is automatically removed.

To check the rate and bandwidth per route, enter the "show multicast route" on page 22306 *extensive* command. To see how many filters are on the switch, enter the VTY command `show filter hw groups`. To clear the route counters, enter the "clear multicast statistics" on page 3730 command.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Firewall Filters for EX Series Switches Overview

multicast (Virtual Tunnel in Routing Instances)

IN THIS SECTION

- Syntax | **11093**
- Hierarchy Level | **11093**
- Description | **11093**
- Default | **11093**
- Required Privilege Level | **11093**
- Release Information | **11093**

Syntax

```
multicast;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name interface vt-
fpc/pic/port.unit-number],
[edit routing-instances routing-instance-name interface vt-fpc/pic/port.unit-number]
```

Description

In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure the virtual tunnel (VT) interface to be used for multicast traffic only.

Default

If you omit this statement, the VT interface can be used for both multicast and unicast traffic.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

> *Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs*
>
> Example: Configuring MBGP MVPN Extranets

multicast-vlan-flood

IN THIS SECTION

- Syntax | 11094
- Hierarchy Level | 11094
- Description | 11094
- Options | 11094
- Required Privilege Level | 11094
- Release Information | 11095

Syntax

```
multicast-vlan-flood
```

Hierarchy Level

```
[edit protocols oam fabric flow-specs flow-specification-name]
```

Description

Configure parameters of the multicast VLAN flood protocol type in the internal fabric monitoring flow specification.

Options

There are no configuration options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | **29924**

multichassis-lag-replicate-state

IN THIS SECTION

- Syntax | **11095**
- Hierarchy Level | **11096**
- Description | **11096**
- Default | **11096**
- Required Privilege Level | **11096**
- Release Information | **11096**

Syntax

```
multichassis-lag-replicate-state;
```

Hierarchy Level

```
[edit multicast-snooping-options]
```

Description

Provide multicast snooping for multichassis link aggregation group interfaces. Replicate IGMP join and leave messages from the active link to the standby link of a dual-link multichassis link aggregation group interface, enabling faster recovery of membership information after failover.

This command is configured globally and effects all routing instances.

Default

If not included, membership information is recovered using a standard IGMP network query.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Multicast Snooping

multicast-snooping-options

multichunk-inspection

IN THIS SECTION

- Syntax | 11097
- Hierarchy Level | 11097
- Description | 11097
- Options | 11097
- Required Privilege Level | 11098
- Release Information | 11098

Syntax

```
multichunk-inspection (enable | disable);
```

Hierarchy Level

```
[edit security sctp]
```

Description

The Stream Control Transmission Protocol (SCTP) firewall checks all chunks in a message and then permits or drops the packet based on the policy. Use the multichunk-inspection enable command to enable SCTP multichunk inspection. It checks all the chunks in a message. Use the multichunk-inspection disable command to disable SCTP multichunk inspection. It checks only the first chunk.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs sctp]` hierarchy level is replaced by `[edit security sctp]`.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Understanding SCTP Multichunk Inspection

multihop

IN THIS SECTION

- Syntax | 11098
- Hierarchy Level | 11099
- Description | 11099
- Default | 11099
- Options | 11100
- Required Privilege Level | 11102
- Release Information | 11102

Syntax

```
multihop {
    no-nexthop-change;
    ttl ttl-value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Configure an EBGP multihop session.

For Layer 3 VPNs, you configure the EBGP multihop session between the *PE* and *CE* routing devices. This allows you to configure one or more routing devices between the PE and CE routing devices.

An external confederation peer is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case because multihop behavior is implied.

If you have external BGP confederation peer-to-loopback addresses, you still need the multihop configuration.

NOTE: You cannot configure the `accept-remote-nexthop` statement at the same time.

Default

If you omit this statement, all EBGP peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or “regular,” BGP session), and the default time-to-live (TTL) value is 1.

NOTE: In Junos, you can configure TTL for both IPv4 and IPv6 multi-hop eBGP sessions. TTL for IPv6 multi-hop eBGP session refers to hop-limit setting.

Options

no-nexthop-change

Specify that the BGP next-hop value not be changed.

An external confederation peer is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case; multihop behavior is implied.

If you have external BGP confederation peer-to-loopback addresses, you still need the multihop configuration.

NOTE: You cannot configure the `accept-remote-nexthop` statement at the same time.

- **Default:** If you omit this statement, all EBGP peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or "regular," BGP session), and the default time-to-live (TTL) value is 1.

ttl *ttl-value*

Configure the maximum time-to-live (TTL) value for the TTL in the IP header of BGP packets.

Configure the maximum time-to-live (TTL) value for the TTL in the IP header of BGP packets.

For BGP multihop scenarios, in which EBGP peers are not directly connected to each other, setting a TTL is optional. The default setting is 64.

For BGP single-hop scenarios, in which external EBGP peers are directly connected to each other, you can, optionally, set the TTL to 255 and configure an inbound firewall filter to allow only BGP control packets with the TTL set to 255. This is in accordance with RFC 3682, *The Generalized TTL Security Mechanism (GTSM)*. For example:

Send all BGP control packets with the TTL set to 255:

```
user@host# show protocols
bgp {
    group toAS2 {
        type external;
```

```
        peer-as 2;
        ttl 255;
        neighbor 10.1.2.3;
        neighbor 10.3.4.5;
        neighbor 10.5.6.7;
    }
}
```

Accept only BGP control packets that have the TTL set to 255:

```
user@host# show firewall
filter ttl-security {
    term gtsm {
        from {
            source-address {
                10.1.2.3/32;
                10.3.4.5/32;
                10.5.6.7/32;
            }
            protocol tcp;
            ttl-except 255;
            port 179;
        }
        then {
            discard;
        }
    }
    term else {
        then {
            accept;
        }
    }
}
```

Apply the firewall filter to the inbound interface for the EBGP single-hop peer:

```
user@host# show interfaces
ge-1/0/0 {
    unit 0 {
        family inet {
```

```
            filter {
                input ttl-security;
            }
        }
    }
}
```

- **Range:** 1 through 255, for multihop peers
- **Default:** 64 (for multihop EBGP sessions, confederations, and IBGP sessions)
- **Range:** 1 or 255, for single-hop peers
- **Default:** 1 (for single-hop EBGP sessions)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for setting the TTL on single-hop external BGP (EBGP) peers introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Example: Configuring EBGP Multihop Sessions

Configuring EBGP Multihop Sessions Between PE and CE Routers in Layer 3 VPNs

accept-remote-nexthop

https://forums.juniper.net/t5/Automation/Scripting-How-To-Use-the-ttl-security-script-to-turn-on-TTL/ta-p/279358

multilink

IN THIS SECTION

- Syntax | 11103
- Hierarchy Level | 11103
- Description | 11103
- Options | 11103
- Required Privilege Level | 11104
- Release Information | 11104

Syntax

```
multilink {
    drop-timeout milliseconds;
    fragment-threshold bytes;
}
```

Hierarchy Level

```
[edit access profile profile-name client client-name l2tp]
```

Description

Configure Multilink PPP for Layer 2 Tunneling Protocol (L2TP).

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options

drop-timeout (M Series, MX Series, PTX Series, T Series only) Configure the drop timeout for a multilink bundle.

- Values:
  - *milliseconds*—Number of milliseconds for the timeout that is associated with the first fragment on the reassembly queue. If the timeout expires before all the fragments have been collected, the fragments at the beginning of the reassembly queue are dropped. If the drop timeout is not specified, the Junos OS holds on to the fragments. (Fragments may still be dropped if the multilink reassembly algorithm determines that another fragment belonging to the packet on a reassembly queue has been lost.)

fragment-threshold (M Series, MX Series, PTX Series, T Series only) Configure the fragmentation threshold for a multilink bundle.

- Values:
  - *bytes*—The maximum number of bytes in a packet. If a packet exceeds the fragmentation threshold, the Junos OS fragments it into two or more multilink fragments.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Layer 2 Tunneling Protocol (L2TP)

multilink-class

IN THIS SECTION

Syntax | 11105

- Hierarchy Level | 11105
- Description | 11105
- Options | 11105
- Required Privilege Level | 11105
- Release Information | 11106

Syntax

```
multilink-class number;
```

Hierarchy Level

```
[edit class-of-service fragmentation-maps map-name forwarding-class class-name]
```

Description

For link services IQ (`lsq`) interfaces only, map a forwarding class into a multiclass MLPPP (MCML).

The `multilink-class` statement and `no-fragmentation` statements are mutually exclusive.

Options

number—The multilink class assigned to this forwarding class.

- **Range:** 0 through 7
- **Default:** None

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces

Configuring Fragmentation by Forwarding Class

Junos OS Services Interfaces Library for Routing Devices

multilink-max-classes | **11106**

multilink-max-classes

IN THIS SECTION

- Syntax | **11106**
- Hierarchy Level | **11106**
- Description | **11107**
- Options | **11107**
- Required Privilege Level | **11107**
- Release Information | **11107**

Syntax

```
multilink-max-classes number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For link services IQ (`lsq`) interfaces only, configure the number of multilink classes to be negotiated when a link joins the bundle.

Options

number—The number of multilink classes to be negotiated when a link joins the bundle.

- **Range:** 1 through 8
- **Default:** None

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

multipath-build-priority

IN THIS SECTION

- Syntax | 11108
- Hierarchy Level | 11108
- Description | 11108
- Required Privilege Level | 11108
- Release Information | 11108

Syntax

```
multipath-build-priority {
    ( low | medium);
}
```

Hierarchy Level

```
[edit fabric protocols bgp],
[edit logical-systems name protocols bgp],
[edit logical-systems name routing-instances name protocols bgp],
[edit protocols bgp],
[edit routing-instances name protocols bgp]
```

Description

Configure a priority for multipath resolution during load balancing. When multipath is enabled on a route reflector, BGP calculates the multipaths each time a new route is added or whenever an existing route is changed, which uses up system resources and slows down the BGP RIB resolution. Configure BGP multipath job priority to delay the multipath calculation and to improve the RIB, also known as the routing table learning rate.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

defer-initial-multipath-build | 5229

Understanding BGP Multipath

multipath (OSPF SPF Options)

IN THIS SECTION

- Syntax | 11109
- Hierarchy Level | 11109
- Description | 11109
- Options | 11110
- Required Privilege Level | 11110
- Release Information | 11110

Syntax

```
multipath [
    weighted {
        one-hop;
    }
}
```

Hierarchy Level

```
[edit protocols ospf spf-options]
```

Description

Enable weighted equal-cost multipath (ECMP) for load-sharing data based on the bandwidth for incoming traffic destined to OSPFv2 neighbors that are at a one-hop distance. The OSPFv2 protocol gets the logical interface bandwidth information associated with the gateways of ECMP next hop and shares this information with the routing protocol process (rpd) and the Packet Forwarding Engine.

Instead of distributing traffic equally over multiple paths, weighted equal-cost multipath (ECMP) routing distributes traffic based on the total available bandwidth for each gateway of a next hop. It thereby distributes traffic unequally over multiple paths and utilizes available bandwidth more efficiently for better load balancing.

NOTE: Weighted equal-cost multipath routing displays unequal traffic distribution prominently when the traffic flow is as heavy as 10000 flows.

Options

weighted Enable weighted ECMP for efficient load sharing.

one-hop Enable weighted ECMP load balancing on IOSPFv2 neighbors that are just one hop away based on the interface bandwidth.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

show route | 23512

multipath (Protocols BGP)

IN THIS SECTION

- Syntax | 11111
- Hierarchy Level | 11111
- Description | 11112
- Options | 11112
- Required Privilege Level | 11113

Release Information | 11113

Syntax

```
multipath {
    allow-external-internal;
    allow-protection;
    disable;
    multiple-as;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp]
[edit logical-systems logical-system-name protocols bgp group group-name]
[edit logical-systems logical-system-name protocols bgp group group-name neighbor]
[edit logical-systems logical-system-name protocols bgp multipath allow-external-internal]
[edit logical-systems name routing-instances name protocols bgp],
[edit logical-systems name routing-instances name protocols bgp group
group-name],
[edit logical-systems name routing-instances name protocols bgp group
group-name neighbor],
[edit logical-systems name tenants name routing-instances name protocols bgp],
[edit logical-systems name tenants name routing-instances name protocols bgp group
group-name],
[edit logical-systems name tenants name routing-instances name protocols bgp group
group-name neighbor],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor],
[edit routing-instances name protocols bgp],
[edit routing-instances name protocols bgp group group-name],
[edit routing-instances name protocols bgp group group-name neighbor],
[edit tenants name routing-instances name protocols bgp]
[edit tenants name routing-instances name protocols bgp group group-name]
[edit tenants name routing-instances name protocols bgp group group-name neighbor]
```

Description

Allow load sharing among multiple EBGP paths and multiple IBGP paths. A path is considered a BGP equal-cost path (and will be used for forwarding) if a tie-break is performed. The tie-break is performed after the BGP route path selection step that chooses the next-hop path that is resolved through the IGP route with the lowest metric. All paths with the same neighboring AS, learned by a multipath-enabled BGP neighbor, are considered.

From 23.4R1 release Junos supports EBGP and IBGP (EIBGP) multipath. In the existing BGP multipath, EBGP routes take priority over IBGP routes because both have different metrics. After you enable EIBGP multipath and there is equal load sharing between EBGP and IBGP routes, Junos OS initiates ECMP using a blend of both EBGP and IBGP.

Feature-specific policies specify prefixes that support EIBGP multipath. You can configure the policy to choose the prefixes based on any match condition.

> **NOTE**: BGP multiple path options must be consistent for all routes forming a BGP multiple path. If BGP multiple path options differ, the multiple path feature chooses a preference, and the multiple path feature might not function as intended.

> **NOTE**: Starting in Junos OS Release 18.1, BGP multipath is supported globally at the `[edit protocols bgp]` hierarchy level. You can selectively disable multipath on some BGP groups and neighbors. Include `disable` at `[edit protocols bgp group` *`group-name`* `multipath]` hierarchy level to disable the multipath option for a group or a specific BGP neighbor.

Options

allow-external-internal Allows IBGP paths to coexist with EBGP paths, and considers both EBGP and IBGP paths equal when path selection between them reaches a point where EBGP paths are preferred.

`allow-protection` Allows BGP multipath and protection to co-exist. When `allow-protection` is configured there may be a change in `show route` output: N+1 nexthop formation (N ecmp plus 1 backup paths).

`disable` Disable the multipath option either at the global level for BGP or for a specific group or neighbor and still allow multipath for other groups or neighbors. Before disabling multipath for specific groups or neighbors, ensure that you have enabled multipath at the global level, by configuring the `multipath` statement at the `[edit protocols bgp]` hierarchy level.

`multiple-as` Disable the default check requiring that paths accepted by BGP multipath must have the same neighboring AS.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`disable` option introduced in Junos OS Release 18.1R1.

`allow-protection` option introduced in Junos OS Release 18.4R1.

`allow-external-internal` option introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Understanding BGP Path Selection

Example: Load Balancing BGP Traffic

multipath (Routing Options)

IN THIS SECTION

- Syntax | 11114
- Hierarchy Level | 11114
- Description | 11114
- Required Privilege Level | 11114
- Release Information | 11115

Syntax

```
multipath {
    vpn-unequal-cost equal-external-internal;
    as-path-compare;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name],
[edit routing-instances routing-instance-name routing-options],
[edit routing-instances routing-instance-name routing-options rib routing-table-name]
```

Description

Enable protocol-independent load balancing for Layer 3 VPNs. This allows the forwarding next hops for both the active route and alternative paths to be used for load balancing. For IPv6, you must configure the `multipath` statement at both the `[edit routing-instances routing-instance-name routing-options]` hierarchy level and the `[edit routing-instances routing-instance-name routing-options rib routing-table-name]` hierarchy level.

Protocol-independent load balancing is applied to VPN routes that are equal up to their router identifiers (`router-id`) with regard to route selection.

The options are explained separately.

NOTE: ACX Series routers do not support `vpn-unequal-cost` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`equal-external-internal` option added in Junos OS Release 8.4.

`as-path-compare` option introduced in Junos OS Release 10.1

RELATED DOCUMENTATION

Configuring Protocol-Independent Load Balancing in Layer 3 VPNs

multipath-rule

IN THIS SECTION

- Syntax | **11115**
- Hierarchy Level | **11116**
- Description | **11116**
- Options | **11116**
- Required Privilege Level | **11117**
- Release Information | **11117**

Syntax

```
multipath-rule name {
      application;
      application-group;
      bandwidth-limit bandwidth-limit;
      enable-reverse-wing;
      link-preferences;
      link-type (IP | MPLS);
```

```
        number-of-paths number-of-paths;
        }
```

Hierarchy Level

```
[edit security advance-policy-based-routing]
```

Description

Configure application-based multipath routing options.

A multipath rule is attached to an SLA rule. The SLA rule is an action part of the APBR rule in a APBR profile. So, whenever a traffic matches the APBR profile, the SLA rule and the multipath rule gets applied and hence multipathing is enabled for that traffic.

Options

name Name of the multipath rule.

application Dynamic application type to match.

application-group Dynamic application group type to match.

bandwidth-limit Bandwidth limit on slowest link to avoid saturation in percentage.

- Range: 1 through 100
- Default: 50% of 1 Gbps

enable-reverse-wing Enable packet copy in the reverse direction flow.

link-preferences Specify link preference for overlay interfaces. You can specify interfaces such as generic routing encapsulation (GRE) and secure tunnel (st)

If you have not configured the `link-preference`, then the AMR selects links from the first two available links from the configured paths.

link-type Specify the underlay interface user defined link types. If you do not configure `link-type`, then AMR selects the first two links from the configured paths.

number-of-path Allocate the number of paths for multipath routing.

- Range: 2 to 16
- Default: 2

Required Privilege Level

Security

Release Information

Statement introduced in Junos OS Release 19.2R1.

multipath (SPF Options)

IN THIS SECTION

- Syntax | 11117
- Hierarchy Level | 11118
- Description | 11118
- Options | 11118
- Required Privilege Level | 11118
- Release Information | 11118

Syntax

```
multipath [
    weighted {
        one-hop;
    }
}
```

Hierarchy Level

```
[edit protocols isis spf-options]
```

Description

Enable weighted equal-cost multipath (WECMP) routing to load-balance traffic in proportion to the changing aggregated ethernet link bandwidth on an equal cost path (ECMP) between directly connected routers. Distributing traffic equally on all equal cost path (ECMP) might lead to imbalanced link utilization if a change occurs in the capacity of any of the local links. In this case, weighted ECMP facilitates load balancing of traffic between equal cost paths only in proportion to the capacity of the local links. After you enable the new configuration statement, whenever an interface change happens, route installation gets scheduled with the updated bandwidth value to distribute traffic accordingly.

Options

weighted Enable weighted ECMP for efficient load sharing.

one-hop Enable weighted ECMP load balancing on IS-IS neighbors that are just one hop away based on the interface bandwidth.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F4.

RELATED DOCUMENTATION

show route | 23512

Example: Weighted ECMP Traffic Distribution on One-Hop IS-IS Neighbors

Understanding Weighted ECMP Traffic Distribution on One-Hop IS-IS Neighbors

multiple-triggered-joins

IN THIS SECTION

- Syntax | 11119
- Hierarchy Level | 11119
- Description | 11119
- Options | 11119
- Required Privilege Level | 11120
- Release Information | 11120

Syntax

```
multiple-triggered-joins {
count (Protocols PIM)number;
interval (Protocols PIM)milliseconds;
}
```

Hierarchy Level

```
[edit protocols piminterface (Protocols PIM)interface-name ]
```

Description

Enable PIM which emits multiple triggered joins between PIM neighbors at configured or default short intervals.

The remaining statements are explained separately. See CLI Explorer.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

count — Number of triggered joins.

- **Range:** 5 through 15
- **Default:** 5

interval — Interval between multiple triggered joins in milliseconds.

- **Range:** 100 through 1000
- **Default:** 100

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

PIM on Aggregated Interfaces

count (Protocols PIM) | **4684**

interval (Protocols PIM) | **9202**

interface (Protocols PIM) | **8991**

multiplier (Protocols PIM)

IN THIS SECTION

- Syntax | **11121**
- Hierarchy Level | **11121**
- Description | **11121**

- Options | 11121
- Required Privilege Level | 11121
- Release Information | 11121

Syntax

```
multiplier number;
```

Hierarchy Level

```
[edit protocols pim interface (Protocols PIM) interface-name bfd-liveness-detection],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM) interface-
name bfd-liveness-detection]
```

Description

Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Options

number—Number of hello packets.

- **Range:** 1 through 255
- **Default:** 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring BFD for PIM

multiplier

IN THIS SECTION

- Syntax | 11122
- Hierarchy Level | 11122
- Description | 11123
- Options | 11123
- Required Privilege Level | 11123
- Release Information | 11123

Syntax

```
multiplier number;
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group group-name liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd]
```

Description

Configure the number of hello packets not received by the neighbor before Bidirectional Forwarding Detection (BFD) declares the neighbor down.

Options

number Maximum allowable number of hello packets missed by the neighbor.

- **Range:** 1 through 255
- **Default:** 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients
- Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

multiplier (BFD Liveness Detection)

IN THIS SECTION

- Syntax | 11124
- Hierarchy Level | 11124
- Description | 11125
- Options | 11125

Required Privilege Level | 11125
Release Information | 11125

Syntax

```
multiplier number;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls)  oam
bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-
id oam bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-
group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)oam
bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn |
vpls)neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-
group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-
detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-name neighbor
neighbor-id oam bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-
detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam
bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-
name neighbor neighbor-id oam bfd-liveness-detection]
```

Description

Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Options

number—Number of hello packets.

- **Range:** 1 through 255
- **Default:** 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Support for BFD authentication introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD for Layer 2 VPN and VPLS

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

bfd-liveness-detection (Layer 2 VPN and VPLS) | **2929**

multiplier (Liveness Detection)

IN THIS SECTION

- Syntax | **11126**
- Hierarchy Level | **11126**

- Description | 11126
- Options | 11126
- Required Privilege Level | 11126
- Release Information | 11126

Syntax

```
multiplier number;
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection]
```

Description

Configure the number of liveness detection requests not received by the peer before Bidirectional Forwarding Detection (BFD) declares the peer is down.

Options

number Maximum allowable number of liveness detection requests missed by the peer.

- **Range:** 1 through 255
- **Default:** 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

multipoint

IN THIS SECTION

- Syntax | 11127
- Hierarchy Level | 11127
- Description | 11127
- Default | 11127
- Required Privilege Level | 11127
- Release Information | 11128

Syntax

```
multipoint;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure the interface unit as a multipoint connection.

Default

If you omit this statement, the interface unit is configured as a point-to-point connection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Multipoint Connection

point-to-point | 12619

multiservice

IN THIS SECTION

- Syntax | 11128
- Hierarchy Level | 11129
- Description | 11129
- Options | 11129
- Required Privilege Level | 11129
- Release Information | 11129

Syntax

```
multiservice {
    source-mac;
    destination-mac;
    payload {
        ip {
            layer3-only;
            layer-3 (source-ip-only | destination-ip-only);
            layer-4;
            inner-vlan-id;
            outer-vlan-id;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number hash-key family]
```

Description

(MX Series 5G Universal Routing Platforms and QFX10000 switches only) Configure data used in a hash key for the **multiservice** protocol family when configuring PIC-level hashing for load balancing on an 802.3ad Link Aggregation Group.

Options

destination-mac Include destination MAC address in the hash key.

payload Include payload data in the hash key. This option has the following suboptions:

- **ip**—Include the IP address of the IPv4 or IPv6 payload into the hash key.
- **layer-3**—Include Layer 3 IP information in the hash key.
- **layer-4**—Include Layer 4 IP information in the hash key.
- **outer-vlan-id**—Include outer VLAN ID information in the hash key.
- **inner-vlan-id**—Include inner VLAN ID information in the hash key.

source-mac Include source MAC address in the hash key.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D10.

RELATED DOCUMENTATION

Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers

multiservice-options (Flow Monitoring)

IN THIS SECTION

- Syntax | 11130
- Hierarchy Level | 11130
- Description | 11131
- Required Privilege Level | 11131
- Release Information | 11131

Syntax

```
multiservice-options {
    (core-dump | no-core-dump);
    (syslog | no-syslog);
    flow-control-options {
        down-on-flow-control;
        dump-on-flow-control;
        reset-on-flow-control;
    }
}
```

Hierarchy Level

```
[edit interfaces mo-fpc/pic/port]
```

Description

For flow-monitoring interfaces only, configure multiservice-specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Flow Monitoring

multiservice-options

IN THIS SECTION

- Syntax | 11131
- Hierarchy Level | 11132
- Description | 11132
- Required Privilege Level | 11132
- Release Information | 11132

Syntax

```
multiservice-options {
    (syslog | no-syslog);
```

```
    (core-dump | no-core-dump);
    (dump-on-flow-control);
    flow-control-options {
        down-on-flow-control;
        dump-on-flow-control;
        reset-on-flow-control;
        up-on-flow-control;
    }
}
```

Hierarchy Level

```
[edit interfaces mo-fpc/pic/port]
```

Description

For monitoring services interfaces only, configure multiservice-specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring Multiservice Physical Interface Properties
- Junos OS Services Interfaces Library for Routing Devices
- passive-monitor-mode | 12212

must-secure

IN THIS SECTION

- Syntax | 11133
- Hierarchy Level | 11133
- Description | 11133
- Default | 11134
- Required Privilege Level | 11134
- Release Information | 11134

Syntax

```
must-secure;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name mka]
```

Description

Specifies that all traffic travelling on the MACsec-secured link must be MACsec-secured to be forwarded onward.

When the `must-secure` option is enabled, all traffic that is not MACsec-secured that is received on the interface is dropped.

When the `must-secure` option is disabled, all traffic from devices that support MACsec is MACsec-secured while traffic received from devices that do no support MACsec is forwarded through the network.

The `must-secure` option is particularly useful in scenarios where multiple devices, such as a phone and a PC, are accessing the network through the same Ethernet interface. If one of the devices supports MACsec while the other device does not support MACsec, the device that doesn't support MACsec can continue to send and receive traffic over the network—provided the `must-secure` option is disabled—while

traffic to and from the device that supports MACsec is MACsec-secured. In this scenario, traffic to the device that is not MACsec-secured must be VLAN-tagged.

Default

The `must-secure` option is disabled.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Configuring Media Access Control Security (MACsec) on Routers

mutual-authentication (REST API)

IN THIS SECTION

- Syntax | 11135
- Hierarchy Level | 11135
- Description | 11135
- Options | 11135
- Required Privilege Level | 11135
- Release Information | 11135

Syntax

```
mutual-authentication {
    certificate-authority certificate-authority-profile-name;
}
```

Hierarchy Level

```
[edit system services rest https]
```

Description

Enable SSL/TLS mutual authentication. The server certificate must be set when configuring mutual authentication.

The remaining statement is explained separately. See CLI Explorer.

Options

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

mvpn

IN THIS SECTION

Syntax | 11136

- Hierarchy Level | **11137**
- Description | **11138**
- Options | **11138**
- Required Privilege Level | **11138**
- Release Information | **11138**

Syntax

```
mvpn {
    inter-region-template{
        template template-name {
            all-regions {
                incoming;
                ingress-replication {
                    create-new-ucast-tunnel;
                    label-switched-path {
                        label-switched-path-template (Multicast) {
                            (default-template | lsp-template-name);
                        }
                    }
                }
            ldp-p2mp;
                rsvp-te {
                    label-switched-path-template (Multicast) {
                        (default-template | lsp-template-name);
                    }
                static-lsp static-lsp;
            region region-name{
                incoming;
                ingress-replication {
                    create-new-ucast-tunnel;
                    label-switched-path {
                        label-switched-path-template (Multicast){
                            (default-template | lsp-template-name);
                        }
                    }
                }
                ldp-p2mp;
```

```
                rsvp-te {
                    label-switched-path-template (Multicast) {
                        (default-template | lsp-template-name);
                    }
                    static-lsp static-lsp;
                }
            }
        }
    }
    mvpn-mode (rpt-spt | spt-only);
    receiver-site;
    sender-site;
    route-target {
        export-target {
            target target-community;
            unicast;
        }
        import-target {
            target {
                target-value;
                receiver target-value;
                sender target-value;
            }
            unicast {
                receiver;
                sender;
            }
        }
    }
}
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable next-generation multicast VPNs in a routing instance.

Options

`receiver-site`—Allow sites with multicast receivers.

`sender-site`—Allow sites with multicast senders.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for the `traceoptions` statement at the `[edit protocols mvpn]` hierarchy level introduced in Junos OS Release 13.3.

Support for the `inter-region-template` statement at the `[edit protocols mvpn]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Routing Instances for an MBGP MVPN

mvpn (Draft-Rosen MVPN)

IN THIS SECTION

- Syntax | 11139
- Hierarchy Level | 11139

- Description | 11139
- Options | 11139
- Required Privilege Level | 11140
- Release Information | 11140

Syntax

```
mvpn {
    family {
        inet {
            autodiscovery {
                inet-mdt;
                }
            disable
        }
        inet6 {
            disable
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Configure the control plane to be used for PE routers in the VPN to discover one another automatically. From here, you can also disable IPv6 draft-rosen multicast VPN at this hierarchy by using the `disable` command at the `protocols pim mvpn family inet6` hierarchy.

Options

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

The `autodiscovery` statement was moved from `[.. protocols pim mvpn]` to `[..protocols pim mvpn family inet]` in Junos OS Release 13.3.

RELATED DOCUMENTATION

> *Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs*

mvpn-iana-rt-import

IN THIS SECTION

- Syntax | 11140
- Hierarchy Level | 11141
- Description | 11141
- Default | 11141
- Required Privilege Level | 11141
- Release Information | 11141

Syntax

```
mvpn-iana-rt-import;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit protocols bgp group group-name]
```

Description

Enables the use of *IANA* assigned rt-import type values (0x010b) for mutlicast VPNs. You can configure this statement on ingress PE routers only.

> **NOTE**: If you configure the `mvpn-iana-rt-import` statement in Junos OS release 10.4R2 and later, the Juniper Networks router can inter-operate with other vendors routers for multicast VPNs. However, the Juniper Networks router cannot inter-operate with Juniper Networks routers running Junos OS release 10.4R1 and earlier.
>
> If you do not configure the `mvpn-iana-rt-import` statement in Junos OS release 10.4R2 and later, the Juniper Networks router cannot inter-operate with other vendors routers for multicast VPNs. However, the Juniper Networks router can inter-operate with Juniper Networks routers running Junos OS release 10.4R1 and earlier.

Default

The default rt-import type value is 0x010a.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS release 10.4R2.

Statement deprecated in Junos OS release 17.3, which means it no longer appears in the CLI but can be accessed by scripts or by typing the command name until it is finally removed.

mvpn-mode

IN THIS SECTION

- Syntax | 11142
- Hierarchy Level | 11142
- Description | 11142
- Default | 11142
- Required Privilege Level | 11142
- Release Information | 11143

Syntax

```
mvpn-mode (rpt-spt | spt-only);
```

Hierarchy Level

```
[edit logical-systems profile-name routing-instances instance-name protocols mvpn],
[edit routing-instances instance-name protocols mvpn]
```

Description

Configure the mode for customer PIM (C-PIM) join messages. Mixing MVPN modes within the same VPN is not supported. For example, you cannot have spt-only mode on a source PE and rpt-spt mode on the receiver PE.

Default

`spt-only`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs

Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs

mvpn (NG-MVPN)

IN THIS SECTION

- Syntax | 11143
- Hierarchy Level | 11144
- Description | 11144
- Required Privilege Level | 11144
- Release Information | 11144

Syntax

```
mvpn {
    autodiscovery-only {
        intra-as {
            inclusive;
        }
    }
    receiver-site;
    route-target {
        export-target {
            target target-community;
            unicast;
```

```
        }
        import-target {
            target {
                target <target:number:number> <receiver | sender>;
                unicast <receiver | sender>;
            }
            unicast {
                receiver;
                sender;
            }
        }
    }
    sender-site;
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable | no-world-
readable>;
        flag flag <flag-modifier> <disable>;
    }
    unicast-umh-election;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable the MVPN control plane for autodiscovery only.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs

mvrp

IN THIS SECTION

- Syntax (EX Series with ELS Support) | **11145**
- Syntax (EX Series) | **11146**
- Syntax (MX Series, EX Series, SRX 1500, SRX 300, SRX 550M, SRX 345, SRX 340, SRX 320) | **11146**
- EX Series with ELS Support | **11147**
- EX Series and MX Series | **11147**
- SRX 1500, SRX 300, SRX 550M, SRX 345, SRX 340, SRX 320 | **11147**
- Description | **11147**
- Default | **11148**
- Required Privilege Level | **11148**
- Release Information | **11148**

Syntax (EX Series with ELS Support)

```
mvrp {
    interface interface-name {
        join-timer milliseconds;
        leave-timer milliseconds;
        leaveall-timer seconds;
        registration (forbidden | normal);
    }
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer seconds;
    no-attribute-length-in-pdu
    no-dynamic-vlan;
    traceoptions  {
```

```
        file filename <files number > <size size> < world-readable | no-world-readable>;
        flag <flag> <disable>;
    }
}
```

Syntax (EX Series)

```
mvrp {
    add-attribute-length-in-pdu;
    disable (MVRP);
    interface (MVRP) (all | interface-name) {
        disable (MVRP);
        join-timer (MVRP) milliseconds;
        leave-timer (MVRP) milliseconds;
        leaveall-timer milliseconds;
        registration (forbidden | normal);
    }
    no-dynamic-vlan;
    traceoptions  {
        file filename <files number > <size size> <no-stamp | world-readable | no-world-
readable>;
        flag flag;
    }
}
```

Syntax (MX Series, EX Series, SRX 1500, SRX 300, SRX 550M, SRX 345, SRX 340, SRX 320)

```
mvrp {
    bpdu-destination-mac-address provider-bridge-group;
    join-timer (MVRP) milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    interface (all | interface-name)  {
        join-timer (MVRP) milliseconds;
        leave-timer milliseconds;
        leaveall-timer milliseconds;
        point-to-point;
         registration (forbidden | normal | restricted);
```

```
    }
    no-attribute-length-in-pdu
    no-dynamic-vlan;
    traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable | no-world-
readable>;
        flag flag;
    }
}
```

EX Series with ELS Support

```
[edit protocols]
```

EX Series and MX Series

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols]
(for virtual switch instance type),
```

```
[edit protocols],
```

```
[edit routing-instances routing-instance-name protocols] (for virtual switch instance type),
```

SRX 1500, SRX 300, SRX 550M, SRX 345, SRX 340, SRX 320

```
[edit protocols],
[edit routing-instances routing-instance-name protocols] (for virtual switch instance type),
```

Description

For Layer 2 networks, configure Multiple VLAN Registration Protocol (MVRP) to dynamically share VLAN information and dynamically configure needed VLANs. Maintaining VLAN configurations based

on active VLANs reduces the amount of traffic traveling in the network, saving network resources. MVRP is configured on trunk interfaces.

Configure Multiple VLAN Registration Protocol (MVRP) on a trunk interface to ensure that the VLAN membership information on the trunk interface is updated as the switch's access interfaces become active or inactive in the configured VLANs.

> **NOTE**: At Junos OS Release 11.3, MVRP was updated to conform to the IEEE standard 802.1ak. This update might result in compatibility issues in mixed release networks. For details, see *Configuring Multiple VLAN Registration Protocol (MVRP) on Switches* .

The remaining statements are explained separately. See CLI Explorer.

Default

MVRP is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- *Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers*
- *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches*
- *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches with ELS Support*
- *Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration*
- *Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration*
- *Configuring Multiple VLAN Registration Protocol (MVRP) on Switches*
- *Verifying That MVRP Is Working Correctly*
- *Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration*

N

IN THIS SECTION

- n-plus-n (Power Management) | **11160**
- n-plus-n (satellite-management) | **11162**
- n391 | **11163**
- n392 | **11164**
- n393 | **11166**
- name (SNMP) | **11167**
- name (early-aging) | **11168**
- name-format (Flow Collection) | **11170**
- name-format | **11172**
- name (L2TP Destination) | **11173**
- name (L2TP Tunnel Destination) | **11175**
- name (Next Gen Services Global System Logging) | **11176**
- name-server (Access) | **11178**
- name-server (System Services) | **11179**
- name (Transport-class) | **11181**
- nas-ip-address | **11182**
- nas-port-extended-format | **11184**
- nas-port-extended-format (Interfaces) | **11186**
- nas-port-id-format (Subscriber Management) | **11188**
- nas-port-method (L2TP LAC) | **11191**
- nas-port-method (Tunnel Profile) | **11192**
- nas-port-options (RADIUS Options) | **11193**
- nas-port-type | **11195**
- nas-port-type (RADIUS Options) | **11197**
- nas-port-type (Subscriber Management) | **11199**
- nasreq (Diameter Application) | **11202**
- nat | **11205**
- nat-cone-binding | **11210**

- nat-destination-pool | **11212**
- nat-destination-rule | **11214**
- nat-interface-port-ol (System) | **11216**
- nat-keepalive (Services IPsec VPN) | **11218**
- nat-nopat-address | **11219**
- nat-options | **11221**
- nat-options (Next Gen Services) | **11222**
- nat-pat-address | **11223**
- nat-pat-portnum | **11225**
- nat-port-ol-ipnumber | **11227**
- nat-rule-referenced-prefix (System) | **11229**
- nat-rule-sets (Service Set) | **11230**
- nat-rule-sets (Service Set Next Gen Services) | **11231**
- nat-rules | **11232**
- nat-source-pool | **11234**
- nat-source-rule | **11236**
- nat-static-rule | **11238**
- native-vlan-id | **11240**
- natt-install-interval | **11242**
- ncp-max-conf-req | **11243**
- ncp-restart-timer | **11245**
- nd-maxucast-retry | **11246**
- nd-override-preferred-src | **11248**
- nd-retransmit-timer | **11249**
- nd-system-cache-limit | **11251**
- nd6-max-cache | **11252**
- nd6-new-hold-limit | **11254**
- nd6-stale-time | **11256**
- ndp | **11257**
- ndp-proxy | **11259**
- negate | **11260**
- negotiate-address | **11261**

- negotiation-options | **11263**
- neighbor | **11264**
- neighbor (Define ANCP) | **11266**
- neighbor-discovery | **11268**
- neighbor-discovery-inspection | **11270**
- neighbor-policy | **11272**
- neighbor (Protocols BGP) | **11273**
- neighbor (Protocols Layer 2 Circuit) | **11278**
- neighbor (Protocols LDP) | **11281**
- neighbor (Protocols RIPng) | **11282**
- neighbor (Protocols VPLS) | **11284**
- nested-application (Security IDP) | **11286**
- nested-application-unknown | **11287**
- nested-applications | **11288**
- netconf | **11290**
- netconf-monitoring (NETCONF) | **11293**
- netconf (Remote Device Management) | **11295**
- network | **11298**
- network (Access) | **11300**
- network (Address Pools) | **11301**
- network-domain | **11302**
- network-element (AAA Profile) | **11304**
- network-element (Diameter Base Protocol) | **11305**
- network-element (Diameter Base Protocol) | **11307**
- network-element (Subscriber Aware Policy Control) | **11308**
- network-elements (RADIUS) | **11310**
- network-isolation | **11312**
- network-isolation-profile | **11314**
- network-management | **11317**
- network-option | **11318**
- network-services | **11320**
- network-slices | **11322**

- network-slicing | **11323**
- network-summary-export | **11325**
- network-summary-import | **11326**
- network (TDF Domain) | **11328**
- next-hop (Access) | **11329**
- next-hop (Access Internal) | **11330**
- next-hop-based-tunnel (Routing Options) | **11332**
- next-hop (Class-Of-Service) | **11333**
- next-hop (Dynamic Access Routes) | **11334**
- next-hop (Forwarding Options) | **11336**
- next-hop-group | **11337**
- next-hop-group (Analyzer) | **11339**
- next-hop-group (Forwarding Options) | **11340**
- next-hop-group (Port Mirroring) | **11343**
- next-hop-group (Unknown Unicast Forwarding) | **11344**
- next-hop-map | **11346**
- next-hop (PIM RPF Selection) | **11347**
- next-hop (Protocols MPLS) | **11349**
- next-hop (RPM) | **11350**
- next-hop-service | **11352**
- next-hop (Static MAC-Based VLANs) | **11353**
- next-hop-tunnel | **11355**
- next-hop (VXLAN Routing) | **11356**
- next-server | **11358**
- next-table | **11359**
- nextboot | **11361**
- nexthop-hold-time | **11362**
- nexthop-learning | **11363**
- ng-juniper | **11366**
- ngfw | **11369**
- no-accounting | **11371**
- no-adaptation | **11373**

- no-adaptation (BFD Liveness Detection) | **11374**
- no-adaptation (Liveness Detection) | **11376**
- no-adaptation (PIM BFD Liveness Detection) | **11377**
- no-adaptive | **11378**
- no-adjacency-down-notification (Protocols IS-IS) | **11380**
- no-advertise-adjacency-segment (Protocols IS-IS) | **11382**
- no-advertise-adjacency-segment (Protocols OSPF) | **11383**
- no advertise-peer-as | **11385**
- no-aggregate-delegate-processing | **11386**
- no-aggregator-id | **11387**
- no-allow-link-events | **11389**
- no-allow-snooped-clients | **11390**
- no-allow-url (Op Scripts) | **11392**
- no-allowed-mac-log | **11393**
- no-anti-replay | **11395**
- no-anti-replay (Services Service Set) | **11396**
- no-application-identification (Services) | **11397**
- no-application-system-cache | **11399**
- no-application-system-cache (Services) | **11400**
- no-attribute-length-in-pdu | **11401**
- no-authentication-check (Protocols RSVP) | **11403**
- no-auto-conversion | **11404**
- no-auto-failover | **11406**
- no-auto-iccp-provisioning (Junos Fusion Redundancy Group) | **11407**
- no-auto-mdix | **11409**
- no-auto-virtual-gateway-esi | **11410**
- no-autostart | **11412**
- no-autoupdate | **11413**
- no-bfd-triggered-local-repair | **11414**
- no-bidirectional-mode | **11416**
- no-bind-on-request (DHCP Relay Agent) | **11418**
- no-broadcast | **11420**

- no-bundle-flap | **11422**
- no-certificate-chain-in-ike | **11423**
- no-client-reflect | **11424**
- no-control-word | **11426**
- no-control-word (Protocols Layer 2 VPN) | **11427**
- no-cspf | **11429**
- no-cspf (Protocols RSVP) | **11431**
- no-decrement-ttl | **11432**
- no-default-gateway-ext-comm | **11434**
- no-delegate-processing | **11435**
- no-dhcp-snooping | **11436**
- no-dhcp-trusted | **11438**
- no-dhcpv6-options | **11440**
- no-dhcpv6-snooping | **11441**
- no-dns-propagation | **11442**
- no-domain-vpn-tag | **11443**
- no-dr-flood (PIM Snooping) | **11445**
- no-dynamic-vlan | **11446**
- no-encryption (MACsec) | **11447**
- no-encryption (MACsec for MX Series) | **11449**
- no-enhanced-frr-bypass (Protocols RSVP) | **11451**
- no-examine-dhcpv6 | **11452**
- no-external-export | **11454**
- no-fabric-wwn-verify | **11456**
- no-fcoe-lag | **11457**
- no-fcoe-trusted | **11458**
- no-filter-check | **11460**
- no-fip-snooping-scaling | **11461**
- no-flow-logging (DDoS Flow Detection) | **11463**
- no-forwarding | **11464**
- no-fragmentation | **11466**
- no-gratuitous-arp-request | **11467**

- no-hidden-commands | **11469**
- no-hierarchical-ecmp | **11470**
- no-icmp-packet-too-big | **11472**
- no-incremental-update | **11473**
- no-inner-payload | **11474**
- no-install | **11475**
- no-install-to-address | **11476**
- no-intelligent-prescreening | **11477**
- no-interface-hello | **11479**
- no-ip-tos-rewrite | **11480**
- no-ipsec-tunnel-in-traceroute | **11481**
- no-ipv4-routing | **11483**
- no-ipv6-routing | **11484**
- no-irb-layer-2-copy | **11486**
- no-issu-timer-negotiation | **11487**
- no-keepalives | **11488**
- no-keepalives (Dynamic Profiles) | **11490**
- no-l2ckt | **11491**
- no-l2vpn | **11492**
- no-listen | **11494**
- no-load-balance-label-capability | **11495**
- no-load-throttle (Resource Monitor) | **11496**
- no-local-reversion | **11497**
- no-local-switching | **11499**
- no-local-switching (VLAN) | **11500**
- no-local-switching (VPLS) | **11502**
- no-logging (Resource Monitor) | **11503**
- no-mac-learning | **11504**
- no-make-before-break | **11508**
- no-management-vlan | **11509**
- no-mcast-replication | **11510**
- no-multicast | **11512**

- no-multicast-echo | **11514**
- no-nat-traversal (Services IPsec VPN) | **11515**
- no-native-vlan-insert | **11517**
- no-neighbor-down-notification | **11519**
- no-neighbor-learn | **11520**
- no-next-header-to-payload-protocol-mapping | **11522**
- no-node-hello-on-bypass | **11523**
- no-node-id-subobject | **11524**
- no-normalization | **11526**
- no-notify-mail-recipient | **11528**
- no-notify-mail-sender (Security Content Filtering Notification Options) | **11529**
- no-notify-mail-sender (Security Fallback Block) | **11530**
- no-notify-mail-sender (Security Virus Detection) | **11531**
- no-nssa-abr | **11533**
- no-odu-backward-frr-enable | **11534**
- no-odu-signal-degrade-monitor-enable | **11536**
- no-option16 | **11537**
- no-option18 | **11538**
- no-option37 | **11540**
- no-option82 | **11541**
- no-overload-on-prefix-export-limit | **11542**
- no-p2mp-sublsp | **11544**
- no-per-unit-scheduler | **11545**
- no-ping-record-route | **11546**
- no-ping-time-stamp | **11547**
- no-policy-cold-synchronization | **11548**
- no-power-budget | **11550**
- no-pre-classifier | **11551**
- no-propagate-ttl | **11553**
- no-qos-adjust | **11554**
- no-qos-adjust (Dynamic Routing Options) | **11556**
- no-recommendation-tlv | **11557**

- no-recommended | **11558**
- no-redirects | **11559**
- no-redirects (IPv4 Traffic) | **11560**
- no-registered-multicast | **11562**
- no-remote-trace (Next Gen Services Global System Logging) | **11564**
- no-remote-trace (System) | **11565**
- no-remote-trace (Trace Options) | **11566**
- no-revert (Local Switching) | **11568**
- no-revert (Neighbor Interface) | **11569**
- no-rfc-1583 | **11571**
- no-root-port | **11572**
- no-route-localize | **11574**
- no-sbl-default-server | **11575**
- no-send-pads-ac-info | **11577**
- no-send-pads-error | **11578**
- no-send-to-ue | **11579**
- no-sequence-check | **11580**
- no-snoop (DHCP Local Server and Relay Agent) | **11582**
- no-source-packet-routing (Protocols OSPF) | **11583**
- no-split-detection | **11585**
- no-split-detection (MX Series Virtual Chassis) | **11586**
- no-syn-check | **11588**
- no-syn-check-in-tunnel | **11589**
- no-syslog | **11590**
- no-syslog-generation | **11591**
- no-tag | **11592**
- no-tag | **11594**
- no-temperature-monitoring | **11595**
- no-termination-request | **11596**
- no-throttle (Resource Monitor) | **11598**
- no-tls-certificate-check | **11599**
- no-transit-statistics | **11600**

- no-translation | **11601**
- no-translation (Source NAT Next Gen Services) | **11603**
- no-trap | **11604**
- no-tunnel-services | **11605**
- no-unknown-unicast | **11607**
- no-unregistered-multicast | **11609**
- no-unsolicited-ra (Enhanced Subscriber Management) | **11611**
- no-uri-check | **11612**
- no-user-query (Services User Identification) | **11614**
- no-validate | **11615**
- no-vlan-interface-name | **11617**
- no-vrf-advertise | **11620**
- no-vrf-propagate-ttl | **11621**
- no-walkup | **11623**
- no-world-readable (Next Gen Services Global System Logging) | **11624**
- no-zero-suppression | **11626**
- no-dhcp-flood | **11627**
- no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping) | **11628**
- no-incremental-update | **11630**
- node (Chassis Cluster Redundancy Group) | **11631**
- node-detection (EVPN-VXLAN) | **11633**
- node-device (Aliases) | **11635**
- node-device (Chassis) | **11637**
- node-device (Flow Groups) | **11638**
- node-device (Resources) | **11640**
- node-group (Chassis) | **11641**
- node-group (OxID Hash Control) | **11643**
- node-group (Resources) | **11645**
- node-hello | **11646**
- node-id | **11648**
- node-link-protection (Protocols MPLS) | **11649**
- node-protection | **11651**

- node-protection (Static LSP) | **11652**
- node-segment (Protocols OSPF) | **11654**
- node (System) | **11656**
- non-active-hold-priority | **11657**
- non-lsp-next-hop | **11659**
- non-revertive | **11660**
- non-revertive (Chassis) | **11661**
- non-revertive (Interfaces) | **11662**
- non-strict-priority-scheduling | **11664**
- non-subscriber-no-reply | **11665**
- non-vc-mode | **11666**
- nonpersistent | **11668**
- nonstop-bridging | **11669**
- nonstop-bridging (Ethernet Switching) | **11670**
- nonstop-routing | **11671**
- nonstop-routing-options | **11673**
- nonvolatile | **11675**
- normalization | **11676**
- normalization (Protocols EVPN) | **11678**
- not (Event Policy) | **11680**
- not-on-disk-underperform | **11681**
- not-present (DHCP Relay Agent) | **11682**
- notification | **11686**
- notification (Commit) | **11687**
- notification-control | **11688**
- notification-interval | **11690**
- notification (NETCONF) | **11691**
- notification-options (Security Antivirus) | **11693**
- notification-options (Security Content Filtering) | **11695**
- notification (System Services) | **11697**
- notification-targets | **11699**
- notification-time | **11700**

- notify | 11701
- notify-filter (Applying to the Management Target) | 11703
- notify-filter (Configuring the Profile Name) | 11704
- notify-mail-recipient | 11705
- notify-mail-sender (Security Content Filtering Notification Options) | 11707
- notify-mail-sender (Security Fallback Block) | 11708
- notify-mail-sender (Security Virus Detection) | 11709
- notify-view | 11710
- np-cache (Flexible PIC Concentrator) | 11711
- nptv6-prefix | 11713
- nptv6-prefix-name | 11714
- nsr-phantom-holdtime | 11715
- nssa | 11717
- nssu | 11719
- ntp | 11721
- ntp (QFabric) | 11729
- ntp threshold | 11731
- nullpdu | 11733
- num-65-127-prefix | 11734
- number-of-frames | 11736
- number-of-ports | 11737
- number-of-sub-ports (interface) | 11740

n-plus-n (Power Management)

IN THIS SECTION

- Syntax | 11161
- Hierarchy Level | 11161
- Description | 11161

Required Privilege Level | 11161
Release Information | 11161

Syntax

```
n-plus-n;
```

Hierarchy Level

```
[edit chassis (EX Series) psu redundancy]
```

Description

Configure *N*+*N* power supply redundancy for power management on an EX6200 or EX8200 switch.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Power Supply Redundancy (CLI Procedure)

n-plus-n (satellite-management)

IN THIS SECTION

- Syntax | 11162
- Hierarchy Level | 11162
- Description | 11162
- Required Privilege Level | 11162
- Release Information | 11162

Syntax

```
n-plus-n
```

Hierarchy Level

```
[edit chassis satellite-management psu redundancy]
```

Description

Configure *N*+*N* power supply redundancy for the satellite devices in a Junos Fusion.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Understanding Power over Ethernet in a Junos Fusion

n391

IN THIS SECTION

- Syntax | 11163
- Hierarchy Level | 11163
- Description | 11163
- Options | 11163
- Required Privilege Level | 11164
- Release Information | 11164

Syntax

```
n391 number;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

For link services interfaces only, set the Frame Relay full status polling interval.

Options

number—Polling interval.

- **Range:** 1 through 255

- **Default:** 6

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Keepalives on Link Services Physical Interfaces

n392 | **11164**

n393 | **11166**

t391 | **28890**

t392 | **28892**

n392

IN THIS SECTION

- Syntax | **11165**
- Hierarchy Level | **11165**
- Description | **11165**
- Options | **11165**
- Required Privilege Level | **11165**
- Release Information | **11165**

Syntax

```
n392 number;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

For link services interfaces only, set the Frame Relay error threshold, in number of errors.

Options

number—Error threshold.

- **Range:** 1 through 10
- **Default:** 3

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Keepalives on Link Services Physical Interfaces

n391 | **11163**

n393 | **11166**

t391 | **28890**

t392 | **28892**

n393

IN THIS SECTION

- Syntax | 11166
- Hierarchy Level | 11166
- Description | 11166
- Options | 11166
- Required Privilege Level | 11166
- Release Information | 11167

Syntax

```
n393 number;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

For link services interfaces only, set the Frame Relay monitored event count.

Options

number—Event count.

- **Range:** 1 through 10
- **Default:** 4

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Keepalives on Link Services Physical Interfaces

n391 | **11163**

n392 | **11164**

t391 | **28890**

t392 | **28892**

name (SNMP)

IN THIS SECTION

- Syntax | **11167**
- Hierarchy Level | **11168**
- Description | **11168**
- Options | **11168**
- Required Privilege Level | **11168**
- Release Information | **11168**

Syntax

```
name name;
```

Hierarchy Level

```
[edit snmp]
```

Description

Set the system name from the command-line interface.

Options

name—System name override.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Different System Name

name (early-aging)

IN THIS SECTION

- Syntax | 11169
- Hierarchy Level | 11169
- Description | 11169
- Required Privilege Level | 11169

Release Information | 11169

Syntax

```
            name;
```

Hierarchy Level

```
[edit security flow aging early-aging application]
```

Description

Application name - Allows configuration cutoff time based on application.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

This statement is introduced in Junos OS Release 23.4R1.

name-format (Flow Collection)

IN THIS SECTION

- Syntax | 11170
- Hierarchy Level | 11170
- Description | 11170
- Options | 11170
- Required Privilege Level | 11171
- Release Information | 11172

Syntax

```
name-format "format";
```

Hierarchy Level

```
[edit services flow-collector file-specification variant variant-number]
```

Description

Specify the name format for a specific file format. The files can include supported macros. Use macros to organize files on the external machine to which they are exported from the collector PIC.

Options

format—Specify the filename format, within quotation marks. The name format can include the following macros, typed in braces:

- `{%D}`—Date
- `{%T)`—Time when the file is created
- `{%I}`—Description string for the logical interface configured using the `collector` statement at the `[edit services flow-collector interface-map]` hierarchy level

- {%N}—Unique, sequential number for each new file created
- {am_pm}—AM or PM
- {date}—Current date using the{year} {month} {day} macros
- {day}—From01 through 31
- {day_abbr}—Sun through Sat
- {day_full}—Sunday through Saturday
- {generation number}—Unique, sequential number for each new file created
- {hour_12}—From 01 through 12
- {hour_24}—From 00 through 23
- {ifalias}—Description string for the logical interface configured using the collector statement at the [edit services flow-collector interface-map] hierarchy level
- {minute}—From 00 through 59
- {month}—From 01through 12
- {month_abbr}—Jan through Dec
- {month_full}—January through December
- {num_zone}—From -2359 through +2359; this macro is not supported
- {second}—From 00 through 60
- {time}—Time the file is created, using the {hour_24} {minute} {second} macros
- {time_zone}—Time zone code name of the locale; for example,gmt (this macro is not supported).
- {year}—In the format YYYY; for example, 1970
- {year_abbr}—From 00 through 99

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring File Formats

name-format

IN THIS SECTION

- Syntax | 11172
- Hierarchy Level | 11172
- Description | 11172
- Options | 11173
- Required Privilege Level | 11173
- Release Information | 11173

Syntax

```
name-format (character-string | none | dns | mac+2oct);
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]
```

Description

Specify the format of the maintenance domain name.

Options

`character-string`—The name is an ASCII character string.

`none`—The maintenance domain name is not used.

`dns`—The name is in domain name service (DNS) format. For example: www.juniper.net.

`mac+2oct`—Name is the MAC address plus a two-octet maintenance association identifier. For example: 08:00:22:33:44:55.100.

- **Default:** `character-string`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Creating a Maintenance Association

Creating a Maintenance Domain

name (L2TP Destination)

IN THIS SECTION

- Syntax | 11174
- Hierarchy Level | 11174
- Description | 11174
- Options | 11174
- Required Privilege Level | 11174

Release Information | 11174

Syntax

```
name destination-name {
    drain;
}
```

Hierarchy Level

```
[edit services l2tp destination]
```

Description

Specify the name of the L2TP destination for the tunnel.

Options

destination-name—Locally assigned name of the tunnel destination.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

name (L2TP Tunnel Destination)

IN THIS SECTION

- Syntax | 11175
- Hierarchy Level | 11175
- Description | 11176
- Options | 11176
- Required Privilege Level | 11176
- Release Information | 11176

Syntax

```
name name {
    address ip-address {
        drain;
        routing-instance routing-instance-name {
            drain;
        }
    }
    drain;
}
```

Hierarchy Level

```
[edit services l2tp tunnel]
```

Description

Specify the local name and other attributes of the L2TP tunnel.

Options

name —Locally assigned name of the tunnel; in the format *destination-name*/*tunnel-name* or *tunnel-name*.

> **NOTE**: When only the tunnel name is provided, then you must identify the destination for the tunnel by including the `address` *ip-address* statement at the `[edit services l2tp tunnel name` *name*`]` hierarchy level.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

name (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | **11177**
- Hierarchy Level | **11177**
- Description | **11177**

- Options | 11177
- Required Privilege Level | 11177
- Release Information | 11177

Syntax

```
name;
```

Hierarchy Level

```
[edit services rtlog traceoptions flag]
```

Description

Specify what to flag in the trace information.

Options

all	Everything
configuration	Reading of configuration
hpl	Trace HPL logging
report	Trace report
source	Communication with security log forwarder

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

name-server (Access)

IN THIS SECTION

- Syntax | 11178
- Hierarchy Level | 11178
- Description | 11178
- Required Privilege Level | 11178
- Release Information | 11179

Syntax

```
name-server address
```

Hierarchy Level

```
[edit access address-assignment pool <name> family (inet | inet6) dhcp-attributes]
```

Description

Specify the DNS server IP address for an address-assignment pool.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

address-assignment (Access) | 1858

name-server (System Services)

IN THIS SECTION

- Syntax | 11179
- Hierarchy Level | 11180
- Description | 11180
- Options | 11180
- Required Privilege Level | 11180
- Release Information | 11180

Syntax

```
name-server {
    address {
        routing-instance routing-instance;
    }
}
```

Hierarchy Level

```
[edit system],
[edit system services dhcp],
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]
```

Description

Configure one or more Domain Name System (DNS) name servers.

Options

address Address of the name server. To configure multiple name servers, include a maximum of three *address* options.

routing-instance *routing-instance* Configure name of the routing instance through which the name server is reachable.

> **NOTE**: The only routing instance supported is `mgmt_junos`. Also, this routing instance command is not supported on SRX Series devices.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`routing-instance` options introduced in Junos OS Release 19.2R1 under the `[edit system]` hierarchy level only.

RELATED DOCUMENTATION

- Configuring a DNS Name Server for Resolving Hostnames into Addresses

name (Transport-class)

IN THIS SECTION

- Syntax | 11181
- Hierarchy Level | 11181
- Description | 11182
- Options | 11182
- Required Privilege Level | 11182
- Release Information | 11182

Syntax

```
nametc-name {
    tunnel-egress {
        end-point address;
    }
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options transport-class],
[edit logical-systems name routing-options transport-class],
[edit logical-systems name tenants name routing-instances name routing-options transport-class],
[edit routing-instances name routing-options transport-class],
[edit routing-options transport-class],
[edit tenants name routing-instances name routing-options transport-class]
```

Description

Configure a tunnel end-point to advertise in BGP CT family, which installs a static route in the `tc-<n>.inet.3 RIB` routing table. Originating BGP-CT routes at the egress-node allows downstream border nodes to recognize them as redundant path to the same PE.

Options

tunnel-egress Specify egress tunnel endpoint origination.

- `endpoint-addr`—Specify the address of the endpoint of the egress tunnel.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1

RELATED DOCUMENTATION

BGP Classful Transport Planes Overview

nas-ip-address

IN THIS SECTION

- Syntax | 11183
- Hierarchy Level | 11183
- Description | 11183
- Required Privilege Level | 11183
- Release Information | 11183

Syntax

```
nas-ip-address {
    equals value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from]
```

Description

Specify the RADIUS AVP NAS-IP-Address for the incoming RADIUS request from the subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers*
- *Understanding Selection of Properties for an IP-Based TDF Subscriber*
- *Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber*
- *IP-Based Subscriber Setup Overview*

nas-port-extended-format

IN THIS SECTION

- Syntax | 11184
- Hierarchy Level | 11184
- Description | 11185
- Options | 11185
- Required Privilege Level | 11185
- Release Information | 11186

Syntax

```
nas-port-extended-format {
    adapter-width bits;
    ae-width bits;
    atm {
        adapter-width bits;
        port-width bits;
        slot-width bits;
        vci-width bits;
        vpi-width bits;
    }
    port-width bits;
    pw-width bits;
    slot-width bits;
    stacked-vlan-width bits;
    vlan-width bits;
}
```

Hierarchy Level

```
[edit access profile profile-name radius options]
```

Description

Configure the RADIUS client to use the extended format for RADIUS attribute 5 (NAS-Port) and specify the width in bits of the fields in the NAS-Port attribute.

The NAS-Port attribute specifies the physical port number of the NAS that is authenticating the user, and is formed by a combination of the physical port's slot number, port number, adapter number, VLAN ID, and S-VLAN ID. The NAS-Port extended format specifies the number of bits (bit width) for each field in the NAS-Port attribute: slot, adapter, port, aggregated, Ethernet, VLAN, and S-VLAN.

> **NOTE**: The combined total of the widths of all fields for a subscriber must not exceed 32 bits, or the configuration fails. The router may truncate the values of individual fields depending on the bit width you specify.

Options

`adapter-width` *`width`*—Number of bits in the adapter field.

`ae-width` *`width`*—(Ethernet subscribers only) Number of bits in the aggregated Ethernet identifier field.

`atm`—Specify width for fields for ATM subscribers.

`port-width` *`width`*—Number of bits in the port field.

`pw-width` *`width`*—(Ethernet subscribers only) Number of bits in the pseudowire field. Appears in the Cisco NAS-Port-Info AVP (100).

`slot-width` *`width`*—Number of bits in the slot field.

`stacked-vlan-width` *`width`*—Number of bits in the SVLAN ID field.

`vci-width` *`width`*—(ATM subscribers only) Number of bits in the ATM virtual circuit identifier (VCI) field.

`vlan-width` *`width`*—Number of bits in the VLAN ID field.

`vpi-width` *`width`*—(ATM subscribers only) Number of bits in the ATM virtual path identifier (VPI) field.

> **NOTE**: The total of the widths must not exceed 32 bits, or the configuration will fail.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`ae-width` option added in Junos OS Release 12.1.

`atm` option added in Junos OS Release 12.3R3 and supported in later 12.3R*x* releases.

`atm` option supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

`pw-width` option added in Junos OS Release 15.1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

nas-port-extended-format (Interfaces)

IN THIS SECTION

- Syntax | **11186**
- Hierarchy Level | **11187**
- Description | **11187**
- Options | **11187**
- Required Privilege Level | **11188**
- Release Information | **11188**

Syntax

```
nas-port-extended-format {
    adapter-width bits;
    ae-width bits;
    port-width bits;
    slot-width bits;
```

```
    stacked;
    stacked-vlan-width bits;
    vci-width bits:
    vlan-width bits;
    vpi-width bits;
}
```

Hierarchy Level

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
```

Description

Configure the RADIUS client to use the extended format for RADIUS attribute 5 (NAS-Port) and specify the width in bits of the fields in the NAS-Port attribute.

Options

`adapter-width` *width*—Number of bits in the adapter field.

`ae-width` *width*—Number of bits in the aggregated Ethernet identifier field.

`port-width` *width*—Number of bits in the port field.

`slot-width` *width*—Number of bits in the slot field.

`stacked`—Include stacked VLAN IDs, in addition to VLAN IDs, in the NAS-Port extended format.

`stacked-vlan-width` *width*—Number of bits in the SVLAN ID field.

`vci-width` *width*—Number of bits in the ATM virtual circuit identifier (VCI) field.

`vlan-width` *width*—Number of bits in the VLAN ID field.

`vpi-width` *width*—Number of bits in the ATM virtual path identifier (VPI) field.

NOTE: Each field can be 0 through 32 bits wide; however, the total of the widths of all fields must not exceed 32 bits, or the configuration fails.

The router may truncate the values of individual fields depending on the bit width you specify.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Options `vci-width` and `vpi-width` introduced in Junos OS Release 12.3R3 and supported in later 12.3R*x* releases.

Options `vci-width` and `vpi-width` supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

RELATED DOCUMENTATION

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN

nas-port-id-format (Subscriber Management)

IN THIS SECTION

- Syntax | 11189
- Hierarchy Level | 11189
- Description | 11189
- Default | 11190
- Options | 11190
- Required Privilege Level | 11190
- Release Information | 11191

Syntax

```
nas-port-id-format {
    agent-circuit-id;
    agent-remote-id;
    concatenated-vlan-tags (fixed-size-inner-tag | fixed-size-outer-tag)
    interface-description;
    interface-text-description;
    nas-identifier;
    order (agent-circuit-id | agent-remote-id | interface-description | interface-text-
description | nas-identifier | postpend-vlan-tags);
    postpend-vlan-tags;
}
```

Hierarchy Level

```
[edit access profile profile-name radius options]
```

Description

Specify the optional information that the router includes in the NAS-Port-ID (RADIUS attribute 87) that is passed to the RADIUS server during authentication and accounting. You can include any combination of the optional values.

When you specify the values for the NAS-Port-ID, you can configure the values to appear in either the default order or a custom order of your choice.

> **NOTE**: The default and custom order methods are mutually exclusive. The configuration fails if you attempt to configure a NAS-Port-ID that includes values in both types of orders.

To specify that the optional values appear in the default order in the NAS-Port-ID, configure the values directly under the `nas-port-id-format` statement. The default order is as follows, in which the `#` character is the delimiter:

```
nas-identifier # interface-description # interface-text-description # agent-circuit-id # agent-remote-id #
postpend-vlan-tags
```

To specify a custom order for the NAS-Port-ID string, you use the `order` option. Include the `order` option before each optional value you want to include in the string, in the order in which you want the options

to appear. For example, the configuration, `order interface-text-description order nas-identifier order agent-remote-id` produces the following NAS-Port-ID, in which the `#` character is the delimiter:

```
interface-text-description # nas-identifier # agent-remote-id
```

Starting in Junos OS Release 21.3R1, we have introduced a new NAS-Port-ID format for RADIUS server access request. The NAS-Port-ID format is `S-VLAN<concatenated 0's>C-VLAN:S-VLAN-C-VLAN`. For customer-VLAN (C-VLAN), if the number of digits is less than four, prepend it with zeroes. For example, NAS-Port-ID for an S-VLAN 72 and C-VLAN 82 is `720082:72-82`.

Default

The router includes the interface description in the NAS-Port-ID when no optional values are specified.

Options

`agent-circuit-id`—Include the agent circuit ID from either DHCP option 82 or the DSL forum VSAs.

`agent-remote-id`—Include the agent remote ID from either DHCP option 82 or the DSL forum VSAs.

`concatenated-vlan-tags`—Include the vlan tags as a concatenated string.

`fixed-size-inner-tag`—Fixed size inner VLAN tag value of 4 octets.

`fixed-size-outer-tag`—Fixed size outer VLAN tag value of 4 octets.

`interface-description`—Include the interface description (interface identifier).

`interface-text-description`—Include the textual interface description (the text description that is statically configured in the CLI).

`nas-identifier`—Include the NAS identifier value (RADIUS attribute 32).

`order`—Specify the optional values you want to include in the NAS-Port-ID and the customized order in which you want the values to appear. You must include the `order` option before each optional value (for example, `order agent-circuit-id order interface-description`).

`postpend-vlan-tags`—Include the VLAN tags. The router includes the tags in the format `:<outer-tag>-<inner-tag>` for a double-tagged VLAN, or `:<outer-tag>` for a single-tagged VLAN.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Options `interface-text-description`, `order`, and `postpend-vlan-tags` introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Access Profile Options for Interactions with RADIUS Servers

Configuring a NAS-Port-ID with Additional Options

RADIUS Servers and Parameters for Subscriber Access

nas-port-method (L2TP LAC)

IN THIS SECTION

- Syntax | 11191
- Hierarchy Level | 11191
- Description | 11192
- Required Privilege Level | 11192
- Release Information | 11192

Syntax

```
nas-port-method cisco-avp;
```

Hierarchy Level

```
[edit services l2tp tunnel]
```

Description

Globally configure the LAC to interoperate with Cisco LNS devices by including the Cisco NAS Port Info AVP (100) in the ICRQ to the LNS. This AVP conveys the physical NAS port number identifier and the type of the physical port, such as Ethernet or ATM.

> **NOTE**: This global configuration can be overridden by the configuration in a tunnel profile or by the RADIUS configuration.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

L2TP LAC Subscriber Configuration

L2TP LAC Tunneling for Subscribers

L2TP LAC Subscriber Configuration

nas-port-method (Tunnel Profile)

IN THIS SECTION

- Syntax | 11193
- Hierarchy Level | 11193
- Description | 11193
- Required Privilege Level | 11193
- Release Information | 11193

Syntax

```
nas-port-method cisco-avp;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Description

Configure the LAC to interoperate with Cisco LNS devices by including the Cisco NAS Port Info AVP (100) in the ICRQ to the LNS. This AVP conveys the physical NAS port number identifier and the type of the physical port, such as Ethernet or ATM.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

nas-port-options (RADIUS Options)

IN THIS SECTION

Syntax | 11194

- Hierarchy Level | 11194
- Description | 11194
- Options | 11195
- Required Privilege Level | 11195
- Release Information | 11195

Syntax

```
nas-port-options nas-port-options-name {
    nas-port-extended-format {
        adapter-width width;
        ae-width width;
        port-width width;
        slot-width width;
        stacked;
        stacked-vlan-width width;
        vci-width width:
        vlan-width width;
        vpi-width width;
    }
    nas-port-type port-type;
    stacked-vlan-ranges (any | low-outer-tag-high-outer-tag),any;
    vlan-ranges (any | low-tag-high-tag);
}
```

Hierarchy Level

```
[edit interfaces interface-name radius-options]
```

Description

Create a NAS-Port options definition to configure the NAS-Port-Type (61) RADIUS IETF attribute, and an extended format for the NAS-Port (5) RADIUS IETF attribute, on a per-physical interface, per-VLAN, or per-stacked VLAN (S-VLAN) basis. Each NAS-Port options definition includes the NAS-Port extended

format, the NAS-Port-Type, and either the VLAN range of subscribers or the S-VLAN range of subscribers to which the definition applies.

> **NOTE**: You can configure a maximum of 16 NAS-Port options definitions per physical interface. Each definition can include a maximum of 32 VLAN ranges or 32 S-VLAN ranges, but cannot include a combination of VLAN ranges and S-VLAN ranges.

Options

nas-port-options-name	Name of the NAS-Port options definition.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN

nas-port-type

IN THIS SECTION

Syntax | 11196

- Hierarchy Level | 11196
- Description | 11196
- Options | 11196
- Required Privilege Level | 11196
- Release Information | 11197

Syntax

```
nas-port-type {
    ethernet (ethernet);
}
```

Hierarchy Level

```
[edit access profile name radius options]
```

Description

RADIUS is an authentication method for validating users trying to access the device using Telnet. Using the `nas-port-type` configuration statement, you can define the type of physical port to authenticate the user.

Options

ethernet Translation mechanism for changing the Ethernet value.

- Values:
  - ethernet—Configure the NAS port type as Ethernet

Required Privilege Level

access—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D110.

nas-port-type (RADIUS Options)

IN THIS SECTION

- Syntax | 11197
- Hierarchy Level | 11197
- Description | 11197
- Default | 11198
- Options | 11198
- Required Privilege Level | 11199
- Release Information | 11199

Syntax

```
nas-port-type port-type;
```

Hierarchy Level

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
```

Description

Specify the port type used to authenticate subscribers. The router includes the port type in the NAS-Port-Type (61) RADIUS IETF attribute.

Default

If you do not include the `nas-port-type` statement at the `[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]` hierarchy level, the global value configured for `nas-port-type` at the `[edit access profile profile-name radius options]` hierarchy level takes effect.

Options

port-type One of the following port types:

- *value*—A value from 0 through 65535
- `adsl-cap`—Asymmetric DSL, carrierless amplitude phase (CAP) modulation
- `adsl-dmt`— Asymmetric DSL, discrete mutilating (DOT)
- `async`—Asynchronous
- `cable`—Cable
- `ethernet`—Ethernet
- `fddi`—Fiber Distributed Data Interface
- `g3-fax`—G.3 Fax
- `hdlc-clear-channel`—HDLC Clear Channel
- `iapp`—Inter-Access Point Protocol (IAPP)
- `idsl`—ISDN DSL
- `isdn-sync`—ISDN Synchronous
- `isdn-v110`—ISDN Async V.110
- `isdn-v120`—ISDN Async V.120
- `piafs`—Personal Handyphone System (PHS) Internet Access Forum Standard
- `sdsl`—Symmetric DSL
- `sync`—Synchronous
- `token-ring`—Token Ring
- `virtual`—Virtual

- `wireless`—Other wireless
- `wireless-1x-ev`—Wireless 1xEV
- `wireless-cdma2000`—Wireless code division multiple access (CDMA) 2000
- `wireless-ieee80211`—Wireless 802.11
- `wireless-umts`—Wireless universal mobile telecommunications system (UMTS)
- `x25`—X.25
- `x75`—X.75
- `xdsl`—DSL of unknown type

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN

nas-port-type (Subscriber Management)

IN THIS SECTION

Syntax | 11200

- Hierarchy Level | 11200
- Description | 11200
- Default | 11200
- Options | 11201
- Required Privilege Level | 11202
- Release Information | 11202

Syntax

```
nas-port-type {
    ethernet {
        port-type;
    }
}
```

Hierarchy Level

```
[edit access profile profile-name radius options]
```

Description

Specify the port type used to authenticate subscribers. The router includes the port type in RADIUS attribute 61 (NAS-Port-Type attribute).

NOTE: This statement is ignored if the `ethernet-port-type-virtual` statement is included in the same access profile.

Default

The router uses a port type of `ethernet`.

Options

port-type—One of the following port types:

- *value*—A value from 0-65535
- `adsl-cap`—Asymmetric DSL, carrierless amplitude phase (CAP) modulation
- `adsl-dmt`— Asymmetric DSL, discrete multitone (DMT)
- `async`—Asynchronous
- `cable`—Cable
- `ethernet`—Ethernet
- `fddi`—Fiber Distributed Data Interface
- `g3-fax`—G.3 Fax
- `hdlc-clear-channel`—HDLC Clear Channel
- `iapp`—Inter-Access Point Protocol (IAPP)
- `idsl`—ISDN DSL
- `isdn-sync`—ISDN Synchronous
- `isdn-v110`—ISDN Async V.110
- `isdn-v120`—ISDN Async V.120
- `piafs`—Personal Handyphone System (PHS) Internet Access Forum Standard
- `sdsl`—Symmetric DSL
- `sync`—Synchronous
- `token-ring`—Token Ring
- `virtual`—Virtual
- `wireless`—Other wireless
- `wireless-1x-ev`—Wireless 1xEV
- `wireless-cdma2000`—Wireless code division multiple access (CDMA) 2000
- `wireless-ieee80211`—Wireless 802.11

- `wireless-umts`—Wireless universal mobile telecommunications system (UMTS)
- `x25`—X.25
- `x75`—X.75
- `xdsl`—DSL of unknown type

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Access Profile Options for Interactions with RADIUS Servers

RADIUS Servers and Parameters for Subscriber Access

nasreq (Diameter Application)

IN THIS SECTION

- Syntax | 11203
- Hierarchy Level | 11203
- Description | 11203
- Options | 11204
- Required Privilege Level | 11204
- Release Information | 11204

Syntax

```
nasreq
    max-outstanding-requests number;
    partition partition-name {
        destination-host hostname;
        destination-realm realm-name;
        diameter-instance master;
    }
    request-retry retries;
    timeout seconds;
}
```

Hierarchy Level

```
[edit access]
```

Description

Specify the destination and transmission parameters for the Diameter Network Access Server Requirements (NASREQ) protocol.

NASREQ is a Diameter-based authentication and authorization protocol. The NASREQ client has two queues, the transmit queue and response queue. The transmit queue stores outbound packets until sent to Diameter, and includes requests and responses. The response queue stores request packets until Diameter responds to the request, and includes only requests waiting for a response.

The following configuration options control transmission flow and use of the queues:

- `max-outstanding-requests` option specifies the maximum number of requests (includes AAR and STR) that the NASREQ client sends to Diameter for wireline transmissions. Effectively, this is the maximum count of requests on the response-queue (the maximum number of in-flight requests for which there has not been a response or timeout); it does not include sent responses.
- The `request-retry` option specifies how many times the NASREQ client retries transmitting a packet to the Diameter server when a timeout is received from the Diameter server for the request.
- The `timeout` option specifies the number of seconds that an outbound packet remains in the transmit queue before it is declared to have timed out. The NASREQ client does not transmit packets that

have timed out. The timeout value applies to all packets in the transmit queue, including both requests and responses to be sent. Diameter manages packets that time out after transmission.

Options

number Maximum number of requests that the NASREQ client sends to Diameter.

- **Range:** 20 through 100

retries Number of times to re-send a timeout request to Diameter after the initial request. This value applies only to requests in the response queue. A value of 0 indicates to try transmission once and then do not retry.

- **Range:** 0 through 3

seconds Number of seconds that an outbound packet remains in the transmit queue before it is declared to have timed out.

- **Range:** 5 through 30

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Diameter Network Access Server Application (NASREQ)

Configuring the Diameter Network Access Server Application (NASREQ)

nat

IN THIS SECTION

- Syntax | 11205
- Hierarchy Level | 11209
- Description | 11209
- Options | 11209
- Required Privilege Level | 11210
- Release Information | 11210

Syntax

```
nat {
    destination {
        pool pool-name {
            address ip-address {
                (port port-number | to ip-address);
            }
            description text;
            routing-instance routing-instance-name;
        }
        rule-set rule-set-name {
            description text;
            from {
                interface [interface-name];
                routing-instance [routing-instance-name];
                zone [zone-name];
            }
            rule rule-name {
                description text;
                match {
                    (destination-address <ip-address> | destination-address-name <address-name>);
                    destination-port port-number;
                    protocol [protocol-name-or-number];
                    source-address [ip-address];
                    source-address-name [address-name];
```

```
                }
                then {
                    destination-nat (off | pool pool-name);
                }
            }
        }
    }
    proxy-arp {
        interface interface-name {
            address ip-address {
                to ip-address;
            }
        }
    }
    proxy-ndp {
        interface interface-name {
            address ip-address {
                to ip-address;
            }
        }
    }
    natv6v4 {
        no-v6-frag-header;
    }
    source {
        address-persistent;
        interface {
            port-overloading {
                off;
            }
        }
        pool pool-name {
            address ip-address {
                to ip-address;
            }
            description text;
            host-address-base ip-address;
            overflow-pool (interface | pool-name);
            port {
                (no-translation | port-overloading-factor number | range port-low <to port-
high>);
            }
            routing-instance routing-instance-name;
```

```
        }
        pool-default-port-range lower-port-range to upper-port-range;
        pool-utilization-alarm {
            clear-threshold value;
            raise-threshold value;
        }
        port-randomization {
            disable;
        }
        port-round-robin {
            disable;
        }
        port-scaling-enlargement;
        radius-accounting {
            session-drop;
        }
        rule-set rule-set-name {
            description text;
            from {
                interface [interface-name];
                routing-instance [routing-instance-name];
                zone [zone-name];
            }
            rule rule-name {
                description text;
                match {
                    (destination-address <ip-address> | destination-address-name <address-name>);
                    destination-port port-number;
                    protocol [protocol-name-or-number];
                    source-address [ip-address];
                    source-address-name [address-name];
                }
                then {
                    source-nat {
                        interface {
                            persistent-nat {
                                address-mapping;
                                inactivity-timeout seconds;
                                max-session-number value;
                                permit (any-remote-host | target-host | target-host-port);
                            }
                        }
                        off;
```

```
                        pool {
                             persistent-nat {
                                 address-mapping;
                                inactivity-timeout seconds;
                                max-session-number number;
                                permit (any-remote-host | target-host | target-host-port);
                             }
                             pool-name;
                        }
                    }
                }
            }
            to {
                interface [interface-name];
                routing-instance [routing-instance-name];
                zone [zone-name];
            }
        }
        session-persistence-scan;
        subscriber-extension;
    }
    static {
        rule-set rule-set-name {
            description text;
            from {
                interface [interface-name];
                routing-instance [routing-instance-name];
                zone [zone-name];
            }
            rule rule-name {
                description text;
                match {
                    (destination-address ip-address | destination-address-name address-name);
                }
                then {
                    static-nat {
                        inet {
                            routing-instance (default | routing-instance-name);
                        }
                        prefix {
                            address-prefix;
                            routing-instance (default | routing-instance-name);
                        }
```

```
                        prefix-name {
                            address-prefix-name;
                            routing-instance (default | routing-instance-name);
                        }
                    }
                }
            }
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security]
[edit tenants tenant-name security]
```

Description

Configure Network Address Translation (NAT) for NFX Series and SRX Series Firewalls.

Options

destination Configure Destination NAT.

natv6v4 Configure NAT between IPv6 and IPv4 options.

no-v6-frag-header Configure to disable adding fragment header in non-fragmented IPv6 packets when performing IPv4 to IPv6 translation.

proxy-arp	Configure Proxy ARP.
proxy-ndp	Configure Proxy NDP.
source	Configure Source NAT.
static	Configure Static NAT.
traceoptions	Configure NAT traceoptions.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

The `description` option added in Junos OS Release 12.1.

The tenant option is introduced in Junos OS Release 18.3R1.

The `radius-accounting` and `subscriber-extension` options are introduced in Junnos OS Release 24.2R1.

RELATED DOCUMENTATION

Introduction to NAT

Understanding Logical Systems Network Address Translation

nat-cone-binding

IN THIS SECTION

- Syntax | **11211**
- Hierarchy Level | **11211**
- Description | **11211**

- Options | 11211
- Required Privilege Level | 11212
- Release Information | 11212

Syntax

```
nat-cone-binding {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of NAT cone binding configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are

available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 2097152

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Introduction to NAT

nat-destination-pool

IN THIS SECTION

- Syntax | 11213
- Hierarchy Level | 11213
- Description | 11213
- Options | 11213
- Required Privilege Level | 11214
- Release Information | 11214

Syntax

```
nat-destination-pool {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile],
[edit tenants tenant-name security-profile]
```

Description

Specify the number of NAT destination pool configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 8192

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Introduction to NAT

nat-destination-rule

IN THIS SECTION

- Syntax | 11215
- Hierarchy Level | 11215
- Description | 11215
- Options | 11215
- Required Privilege Level | 11216
- Release Information | 11216

Syntax

```
nat-destination-rule {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile],
[edit tenants tenant-name system security-profile]
```

Description

Specify the number of NAT destination rule configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 8192

reserved A reserved quota that guarantees that the resource amount specified is always available.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Introduction to NAT

nat-interface-port-ol (System)

IN THIS SECTION

- Syntax | 11216
- Hierarchy Level | 11217
- Description | 11217
- Options | 11217
- Required Privilege Level | 11217
- Release Information | 11217

Syntax

```
nat-interface-port-ol {
    maximum maximum;
```

```
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the security NAT interface port overloading the quota of a logical system.

Options

maximum *amount* Specify the maximum allowed quota value.

- **Range:** 0 through 64

reserved *amount* Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Introduction to NAT

nat-keepalive (Services IPsec VPN)

IN THIS SECTION

- Syntax | 11218
- Hierarchy Level | 11218
- Description | 11218
- Options | 11218
- Required Privilege Level | 11218
- Release Information | 11219

Syntax

```
nat-keepalive seconds;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Specify the interval at which NAT keepalive packets can be sent so that NAT translation continues.

Options

seconds —Maximum interval in seconds at which NAT keepalive packets can be sent.

- **Range:** 1 through 300 seconds.
- **Default:** 20 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

Configuring Junos VPN Site Secure or IPSec VPN

nat-nopat-address

IN THIS SECTION

- Syntax | 11219
- Hierarchy Level | 11219
- Description | 11220
- Options | 11220
- Required Privilege Level | 11220
- Release Information | 11220

Syntax

```
nat-nopat-address {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of NAT without port address translation configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 4194304

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Introduction to NAT

nat-options

IN THIS SECTION

- Syntax | 11221
- Hierarchy Level | 11221
- Description | 11221
- Required Privilege Level | 11222
- Release Information | 11222

Syntax

```
nat-options {
    land-attack-check (ip-only | ip-port);
    max-sessions-per-subscriber session-number;
    stateful-nat64 {
        clear-dont-fragment-bit;
    }
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify parameters for NAT operation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced with Junos OS Release 12.1.

`land-attack-check` and `max-sessions-per-subscriber` statements added in 13.3.

RELATED DOCUMENTATION

Configuring Service Rules

clear-dont-fragment-bit (Services NAT Options) | **3561**

land-attack-check | **9619**

max-sessions-per-subscriber | **10474**

stateful-nat64 | **28523**

nat-options (Next Gen Services)

IN THIS SECTION

- Syntax | **11222**
- Hierarchy Level | **11223**
- Description | **11223**
- Required Privilege Level | **11223**
- Release Information | **11223**

Syntax

```
nat-options {
    nptv6 {
```

```
        icmpv6-error-messages;
    }
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Send ICMP error messages if NPTv6 address translation fails.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

nat-pat-address

IN THIS SECTION

- Syntax | 11224
- Hierarchy Level | 11224
- Description | 11224
- Options | 11224
- Required Privilege Level | 11225
- Release Information | 11225

Syntax

```
nat-pat-address {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of NAT with port address translation (PAT) configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 2048

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

> *Understanding Logical Systems Network Address Translation*
>
> Introduction to NAT

nat-pat-portnum

IN THIS SECTION

- Syntax | **11225**
- Hierarchy Level | **11226**
- Description | **11226**
- Options | **11226**
- Required Privilege Level | **11226**
- Release Information | **11226**

Syntax

```
nat-pat-portnum {
    maximum maximum;
```

```
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the maximum quantity and the reserved quantity of ports for the logical system as part of its security profile. The total number of PAT pools must not exceed the configured maximum ports for the logical system.

Options

maximum Specify the maximum number of ports allowed for a logical system. The maximum quantity is not guaranteed and is shared among multiple logical systems.

- **Range:** 0 through 134217728

reserved Specify the number of resources guaranteed for a logical system.

- **Range:** For SRX5600 and SRX5800 devices, up to 402,653,184 ports are supported. Pool specifications for logical systems can be viewed using the **show security nat source summary logical-system all** command.
- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways

nat-port-ol-ipnumber

IN THIS SECTION

- Syntax | 11227
- Hierarchy Level | 11227
- Description | 11227
- Options | 11228
- Required Privilege Level | 11228
- Release Information | 11228

Syntax

```
nat-port-ol-ipnumber {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of NAT port overloading IP number configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 16

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways

nat-rule-referenced-prefix (System)

IN THIS SECTION

- Syntax | 11229
- Hierarchy Level | 11229
- Description | 11229
- Options | 11229
- Required Privilege Level | 11230
- Release Information | 11230

Syntax

```
nat-rule-referenced-prefix {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the security NAT rule referenced IP prefix quota of a logical system.

Options

maximum Specify the maximum allowed quota value.

- **Range:** 0 through 1048576

reserved Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

> *Understanding Logical Systems for SRX Series Services Gateways*

nat-rule-sets (Service Set)

IN THIS SECTION

- Syntax | 11230
- Hierarchy Level | 11231
- Description | 11231
- Options | 11231
- Required Privilege Level | 11231
- Release Information | 11231

Syntax

```
nat-rule-sets rule-set-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the Network Address Translation (NAT) rule set included in the service set. You can configure only one NAT rule set. If you specify a NAT rule set, you cannot specify a NAT rule.

Options

rule-set-name	Name of the NAT rule set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Services to Subscriber-Aware Traffic with a Service Set

nat-rule-sets (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 11232
- Hierarchy Level | 11232
- Description | 11232

- Required Privilege Level | 11232
- Release Information | 11232

Syntax

```
nat-rule-sets rule-set-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the NAT rules set included in the service set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

nat-rules

IN THIS SECTION

- Syntax | 11233
- Hierarchy Level | 11233

- Description | 11233
- Options | 11233
- Required Privilege Level | 11233
- Release Information | 11233

Syntax

```
(nat-rules rule-name | nat-rule-sets rule-set-name);
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the Network Address Translation (NAT) rules or rule set included in this service set. You can configure multiple rules, but only one rule set for each service.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

rule-set-name—Identifier for the set of rules to be included.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Service Rules

Applying Services to Subscriber-Aware Traffic with a Service Set

nat-source-pool

IN THIS SECTION

- Syntax | 11234
- Hierarchy Level | 11234
- Description | 11234
- Options | 11235
- Required Privilege Level | 11235
- Release Information | 11235

Syntax

```
nat-source-pool {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile],
[edit tenants tenant-name system security-profile]
```

Description

Specify the NAT source pool configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 2048

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways

nat-source-rule

IN THIS SECTION

- Syntax | 11236
- Hierarchy Level | 11236
- Description | 11236
- Options | 11237
- Required Privilege Level | 11237
- Release Information | 11237

Syntax

```
nat-source-rule {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile],
[edit tenants tenant-name system security-profile]
```

Description

Specify the NAT source rule configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 8192

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways

nat-static-rule

IN THIS SECTION

- Syntax | 11238
- Hierarchy Level | 11238
- Description | 11238
- Options | 11239
- Required Privilege Level | 11239
- Release Information | 11239

Syntax

```
nat-static-rule {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile],
[edit tenants tenant-name system security-profile]
```

Description

Specify the number of NAT static rule configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 8192

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways

native-vlan-id

IN THIS SECTION

- Syntax | 11240
- Hierarchy Level | 11240
- Description | 11240
- Default | 11241
- Options | 11241
- Required Privilege Level | 11241
- Release Information | 11242

Syntax

```
native-vlan-id vlan-id;
```

Hierarchy Level

For platforms without ELS:

```
[edit interfaces interface-name unit 0 family ethernet-switching]
```

For MX Series routers and platforms with ELS:

```
[edit interfaces interface-name]
```

Description

Configure the VLAN identifier to associate untagged packets received on a trunk interface.

You must configure a logical interface with the VLAN ID that is same as the native VLAN ID confiured on the interface, to receive the untagged packets on the interface.

To configure the logical interface, include the `vlan-id` statement (matching the `native-vlan-id` statement on the physical interface) at the `[edit interfaces interface-name unit logical-unit-number] hierarchy level`.

When the `native-vlan-id` statement is included with the "flexible-vlan-tagging" on page 7235 statement, untagged packets are accepted on the same mixed VLAN-tagged port and on the interfaces that are configured for Q-in-Q tunneling.

When the `native-vlan-id` statement is combined with the "interface-mode" on page 8920 statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.

> **NOTE**: Starting in Junos OS Release 17.1R1, you can send untagged traffic without a native VLAN ID to the remote end of the network. To do this, remove the native VLAN ID from the untagged traffic configuration by setting the `no-native-vlan-insert` statement. If you do not configure this statement, the native VLAN ID is added to the untagged traffic.

Default

By default, the untagged packets are dropped on trunk interfaces. That is, if you do not configure the `native-vlan-id` option on trunk interfaces, the untagged packets are dropped.

Options

vlan-id—Numeric identifier of the VLAN.

- **Range:** 1 through 4094

number—VLAN ID number.

- **Range:** 0 through 4094.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

- Configuring Gigabit Ethernet Interfaces (CLI Procedure)
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure)
- *Understanding Bridging and VLANs on Switches*
- *Enabling VLAN Tagging*
- *Configuring Access Mode on a Logical Interface*
- *Configuring the Native VLAN Identifier on Switches With ELS Support*
- *Understanding Interfaces*
- *Understanding Q-in-Q Tunneling and VLAN Translation*
- no-native-vlan-insert | **11517**
- *Sending Untagged Traffic Without VLAN ID to Remote End*
- show ethernet-switching interfaces | **19641**
- show vlans | **27738**
- flexible-vlan-tagging | **7235**
- Junos OS Network Interfaces Configuration Guide

natt-install-interval

IN THIS SECTION

- Syntax | **11243**
- Hierarchy Level | **11243**
- Description | **11243**
- Required Privilege Level | **11243**
- Release Information | **11243**

Syntax

```
natt-install-interval seconds ;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then dynamic]
```

Description

Specify the duration of delay in installing IPSec SA in a NAT-T scenario soon after the IPSec SA negotiation is complete. The default value is 0 seconds.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2.

RELATED DOCUMENTATION

ipsec-vpn-rules | 9325

ncp-max-conf-req

IN THIS SECTION

Syntax | 11244
Hierarchy Level | 11244

- Description | 11244
- Options | 11244
- Required Privilege Level | 11244
- Release Information | 11245

Syntax

```
ncp-max-conf-req number
```

Hierarchy Level

```
[edit interfaces so-fpc/pic/port unit number ppp-options]
```

Description

Set the maximum number of NCP Configure-Requests to be sent, after which the router goes to NCP down state.

Options

number Ranges from 0 to 65535, where 0 means send infinite NCP Configure-Requests and any other value specifies the maximum number NCP Configure-Requests to send and then stop sending.

Default 254

- **Range:** 0 through 65,535

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring PPP

ppp-options | 12949

ncp-restart-timer

IN THIS SECTION

- Syntax | 11245
- Hierarchy Level | 11245
- Description | 11246
- Options | 11246
- Required Privilege Level | 11246
- Release Information | 11246

Syntax

```
ncp-restart-timer milliseconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number ppp-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options]
```

Description

For interfaces with PPP and PPP TCC encapsulations and on multilink PPP bundle interfaces, configure a restart timer for the Network Control Protocol (NCP) component of a PPP session.

Options

milliseconds—The time in milliseconds between successive NCP configuration requests.

- **Range:** 500 through 10,000 milliseconds
- **Default:** 3 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring PPP

nd-maxucast-retry

IN THIS SECTION

- Syntax | 11247
- Hierarchy Level | 11247
- Description | 11247
- Default | 11247
- Options | 11247

Required Privilege Level | 11247
Release Information | 11247

Syntax

```
nd-maxucast-retry count
```

Hierarchy Level

```
[edit system]
```

Description

Configures how many times the device attempts to reach the IPv6 neighbor entries.

Default

The default number of retry attempts is 3.

Options

count Indicates the maximum number of retries for unicast neighbor solicitation.

Required Privilege Level

admin-control

Release Information

Statement introduced in Junos OS Evolved Release 19.4R1.

RELATED DOCUMENTATION

Example: Configuring ARP Cache Protection

arp | 2458

nd-override-preferred-src

IN THIS SECTION

- Syntax | 11248
- Hierarchy Level | 11248
- Description | 11248
- Default | 11249
- Required Privilege Level | 11249
- Release Information | 11249

Syntax

```
nd-override-preferred-src;
```

Hierarchy Level

```
[edit system]
```

Description

Configure the router to override the default configuration and use the appropriate address based on destination address scope for the source address for Neighbor Solicitation/Neighbor Advertisement (NS/NA) for unnumbered interfaces.

Default

The router uses the preferred source address, if configured, as source for NS/NA for unnumbered interfaces. If no preferred source address is configured, the router uses the appropriate address based on destination address scope.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3

RELATED DOCUMENTATION

unnumbered-address (Dynamic Profiles) | 30422

nd-retransmit-timer

IN THIS SECTION

- Syntax | 11250
- Hierarchy Level | 11250
- Description | 11250
- Options | 11250
- Required Privilege Level | 11250
- Release Information | 11250

Syntax

```
nd-retransmit-timer milliseconds;
```

Hierarchy Level

```
[edit system]
```

Description

Set the retransmit timer for neighbor discovery messages. Whenever the state of a neighbor during the Neighbor Discovery (ND) process changes from stale to probe, the value of the retransmit timer controls the interval between the neighbor solicitation messages that are sent out. Also, the retransmit timer controls the time for which the neighbor is in the probe state. A device sends a neighbor solicitation message after the specified number of milliseconds in the nd-retransmit-timer statement, until a reachability confirmation is received. If a solicited neighbor advertisement (NA) message is not received from the neighbor in response to the soliciation message sent from the device, the neighbor remains in the probe state.

Options

milliseconds—Retransmission frequency.

- **Default:** 0 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

> *Using NDRA to Provide IPv6 WAN Link Addressing Overview*

nd-system-cache-limit

IN THIS SECTION

- Syntax | 11251
- Hierarchy Level | 11251
- Description | 11251
- Default | 11251
- Options | 11252
- Required Privilege Level | 11252
- Release Information | 11252

Syntax

```
nd-system-cache-limit number;
```

Hierarchy Level

```
[edit system]
```

Description

Specify the maximum system cache size for IPv6 next-hop addresses. This limit enforces a systemwide cap on the neighbor discovery cache entries for all interfaces, including the loopback interface for the internal routing instance, management interfaces, and the public interfaces.

Default

100,000

Options

number Maximum system cache size for IPv6 next-hop addresses.

- **Range:** 200 through 2,000,000

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Example: Configuring Neighbor Discovery Cache Protection to Prevent Denial-of-Service Attacks

nd6-max-cache | **11252**

nd6-new-hold-limit | **11254**

nd6-max-cache

IN THIS SECTION

- Syntax | **11253**
- Hierarchy Level | **11253**
- Description | **11253**
- Default | **11253**
- Options | **11253**
- Required Privilege Level | **11253**
- Release Information | **11253**

Syntax

```
nd6-max-cache nd6-max-cache;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6]
```

Description

Specify the maximum number of entries that can be added to the Neighbor Discovery Protocol (NDP) IPv6 neighbor discovery cache for an interface. When this maximum is reached, no new entries are allowed.

Default

- 100,000 for M Series and PTX Series.
- 75,000 for MX Series and QFX Series.
- 20,000 for EX Series.

Options

nd6-max-cache Maximum size of the neighbor discovery next-hop cache for an interface.

- **Range:** 1 through 2,000,000 for MX Series, PTX Series or QFX Series.
- **Range:** 1 through 700,000 for M Series.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Neighbor Discovery Cache Protection

Example: Configuring Neighbor Discovery Cache Protection to Prevent Denial-of-Service Attacks

IPv6 Neighbor Discovery Overview

nd6-new-hold-limit | **11254**

nd-system-cache-limit | **11251**

nd6-new-hold-limit

IN THIS SECTION

- Syntax | **11254**
- Hierarchy Level | **11254**
- Description | **11254**
- Default | **11255**
- Options | **11255**
- Required Privilege Level | **11255**
- Release Information | **11255**

Syntax

```
nd6-new-hold-limit nd6-new-hold-limit;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6]
```

Description

Specify the maximum number of entries for unresolved next-hop addresses that can be added to the Neighbor Discovery Protocol (NDP) IPv6 neighbor discovery cache for an interface.

Default

- 100,000 for M Series.
- 75,000 for MX Series and QFX Series.
- 20,000 for EX Series.

Options

nd6-new-hold-limit	Maximum number of new unresolved next-hop addresses that can be added to the IPv6 neighbor discovery cache.

- **Range:** 1 through 2,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Neighbor Discovery Cache Protection

Example: Configuring Neighbor Discovery Cache Protection to Prevent Denial-of-Service Attacks

IPv6 Neighbor Discovery Overview

nd-system-cache-limit | **11251**

nd6-max-cache | **11252**

nd6-stale-time

nd6-stale-time

IN THIS SECTION

- Syntax | 11256
- Hierarchy Level | 11256
- Description | 11256
- Default | 11257
- Options | 11257
- Required Privilege Level | 11257
- Release Information | 11257

Syntax

```
nd6-stale-time seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6]
```

Description

Set the stale timer for IPv6 neighbor reachability confirmation. Reachability of the IPv6 neighbors is confirmed only after the stale timer has expired. For example, by setting the stale timer to 180 seconds, users can specify that IPv6 neighbor reachability be confirmed every 180 seconds.

NOTE: When the Routing Engine sends a control packet to an IPv6 neighbor, the stale timer is the maximum interval in which neighbor reachability is confirmed. In such cases, IPv6 neighbor reachability is confirmed before the stale timer expires.

Default

Default is 20 minutes (1200 seconds)

Options

seconds Duration in seconds.

- **Range:** 1 to 18000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

> *IPv6 Neighbor Discovery Overview*
>
> show ipv6 neighbors | **21680**

ndp

IN THIS SECTION

- Syntax | **11258**
- Hierarchy Level | **11258**
- Release Information | **11258**
- Description | **11258**
- Options | **11258**
- Required Privilege Level | **11259**

Syntax

```
ndp name {
 mac mac;
 multicast-mac multicast-mac;
 l2-interface l2-interface;
 publish (mac mac | multicast-mac multicast-mac;
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit name family inet6 address],
[edit dynamic-profiles name logical-systems name interfaces name unit name family inet6 address],
[edit interfaces name unit name family inet6 address]
```

Release Information

Description

Configures static Neighbor Discovery Protocol (NDP) entries.

Options

name	Specify the destination IP address.
mac	Specify the MAC address.
multicast-mac	Specify the multicast MAC address.
l2-interface	Specify the Layer 2 interface name for NDP entry.
publish	Use this to reply to NDP requests for this entry.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface

ndp-proxy

IN THIS SECTION

- Syntax | 11259
- Hierarchy Level | 11259
- Description | 11260
- Options | 11260
- Required Privilege Level | 11260
- Release Information | 11260

Syntax

```
ndp-proxy {
 interface-restricted;
 interface-unrestricted;
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit name family inet6],
[edit dynamic-profiles name logical-systems name interfaces name unit name family inet6],
[edit interfaces name unit name family inet6]
```

Description

Enable NDP proxy on interface

Options

interface-restricted	Enable NDP interface proxy restricted to interface
interface-unrestricted	Enable NDP interface proxy unrestricted to interface

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1

Option `interface-unrestricted` introduced in Junos OS Release 22.1R1 (on SRX Series and vSRX Virtual Firewall 3.0 platform).

> **NOTE**: You cannot configure both interface restricted and interface unrestricted modes on the same interface simultaneously.

negate

IN THIS SECTION

- Syntax | **11261**
- Hierarchy Level | **11261**
- Description | **11261**
- Required Privilege Level | **11261**
- Release Information | **11261**

Syntax

```
negate;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature]
```

Description

Select negate to exclude the specified pattern from being matched.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

negotiate-address

IN THIS SECTION

- Syntax | 11262
- Hierarchy Level | 11262
- Description | 11262
- Required Privilege Level | 11262
- Release Information | 11262

Syntax

```
negotiate-address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet]
```

Description

For interfaces with PPP encapsulation, enable the interface to be assigned an IP address by the remote end.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring IPCP Options for Interfaces with PPP Encapsulation
- address | **1846**
- unnumbered-address (PPP) | **30426**
- Junos OS Administration Library for Routing Devices

negotiation-options

IN THIS SECTION

- Syntax | 11263
- Hierarchy Level | 11263
- Description | 11263
- Required Privilege Level | 11263
- Release Information | 11264

Syntax

```
negotiation-options {
                        allow-remote-loopback;
                        no-allow-link-events;
}
```

Hierarchy Level

```
[edit protocols oam                ethernet                link-fault-management
interface                interface-name]
```

Description

Enable and disable IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) features for Ethernet interfaces.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Ethernet OAM Link Fault Management

neighbor

IN THIS SECTION

- Syntax | 11264
- Hierarchy Level | 11265
- Description | 11265
- Options | 11265
- Required Privilege Level | 11266
- Release Information | 11266

Syntax

```
neighbor neighbor-name {
    authentication-key password;
    authentication-type type;
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
```

```
        minimum-receive-interval milliseconds;
        transmit-interval {
            threshold milliseconds;
            minimum-interval milliseconds;
        }
        multiplier number;
        version (0 | 1 | automatic);
    }
    (check-zero | no-check-zero);
    demand-circuit;
    import policy-name;
    max-retrans-time seconds;
    message-size number;
    metric-in metric;
    metric-out metric;
    receive receive-options;
    route-timeout seconds;
    send send-options;
    update-interval seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name],
[edit protocols rip group group-name],
[edit routing-instances routing-instance-name protocols rip group group-name]
```

Description

Configure neighbor-specific RIP parameters, thereby overriding the defaults set for the routing device.

Options

neighbor-name—Name of an interface over which a routing device communicates to its neighbors.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

neighbor (Define ANCP)

IN THIS SECTION

- Syntax | 11266
- Hierarchy Level | 11267
- Description | 11267
- Options | 11267
- Required Privilege Level | 11267
- Release Information | 11267

Syntax

```
neighbor ip-address {
    adjacency-loss-hold-time seconds;
    adjacency-timer;
    auto-configure-trigger interface interface-name;
    ietf-mode;
    maximum-discovery-table-entries entry-number;
    pre-ietf-mode;
}
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an ANCP neighbor with which the ANCP agent on the router forms an adjacency for reporting and shaping traffic.

Options

ip-address—IP address of the ANCP neighbor.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent

Configuring ANCP Neighbors

neighbor-discovery

IN THIS SECTION

- Syntax | 11268
- Hierarchy Level | 11269
- Description | 11269
- Options | 11269
- Default | 11269
- Required Privilege Level | 11270
- Release Information | 11270

Syntax

```
neighbor-discovery
 dad-proxy {
  no-proxy-on-resolve;
 }
 ndp-proxy {
  no-proxy-on-resolve;
  proxy-force-resolve;
 }
 no-dad-on-state-change;
    onlink-subnet-only;
    secure {
        security-level {
            (default | secure-messages-only);
        }
        cryptographic-address {
            key-length number;
            key-pair pathname;
        }
        timestamp {
            clock-drift number;
            known-peer-window number;
            new-peer-window number;
```

```
        }
        traceoptions {
            file filename <files number> <match regular-expression> <size size> <world-
readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Enable Secure Neighbor Discovery.

Options

dad-proxy Configure DAD proxy behaviour.

- no-proxy-on-resolve—Disable proxy on unresolved address

ndp-proxy Configure NDP proxy behaviour.

- no-proxy-on-resolve—Disable proxy on unresolved address
- proxy-force-resolve—Enable Neighbor Solicitation for already learned address

no-dad-on-state-change Disable DAD on interface state change.

The remaining statements are explained separately. See CLI Explorer.

Default

Disabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Option `proxy-force-resolve` introduced onn SRX Series and vSRX Virtual Firewall 3.0 platforms in Junos OS Release 22.1R1.

RELATED DOCUMENTATION

Secure IPv6 Neighbor Discovery

neighbor-discovery-inspection

IN THIS SECTION

- Syntax | 11270
- Hierarchy Level | 11271
- Description | 11271
- Default | 11271
- Required Privilege Level | 11271
- Release Information | 11271

Syntax

```
neighbor-discovery-inspection;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security],
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name)]
```

Description

Perform dynamic IPv6 neighbor discovery inspection on the specified VLAN.

When neighbor discovery inspection is configured, the switch inspects IPv6 packets with neighbor discovery messages and validates them against the DHCPv6 binding table. The source IP address and source MAC address of each packet are checked against the table, and if a valid match is not found, the packet is dropped.

> **NOTE**: If you configure the `neighbor-discovery-inspection` statement at the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level, DHCPv6 snooping is automatically enabled for the specified VLAN.
> See *IPv6 Neighbor Discovery Inspection* for more information about this configuration.
>
> If you configure the `neighbor-discovery-inspection` statement at the `[edit ethernet-switching-options secure-access-port vlan (all | vlan-name)]` hierarchy level, you must also enable DHCPv6 snooping for the specified VLAN or VLANs.

Default

Disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D20.

Support at the `[edit ethernet-switching-options secure-access-port vlan (all | vlan-name)]` hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

RELATED DOCUMENTATION

Configuring Port Security (ELS)

Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing

neighbor-policy

IN THIS SECTION

- Syntax | 11272
- Hierarchy Level | 11272
- Description | 11272
- Options | 11273
- Required Privilege Level | 11273
- Release Information | 11273

Syntax

```
neighbor-policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name],
[edit protocols pim interface (Protocols PIM)  interface-name],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name]
```

Description

Apply a PIM interface-level policy to filter neighbor IP addresses.

Options

policy-name—Name of the policy that filters neighbor IP addresses.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring Interface-Level PIM Neighbor Policies

neighbor (Protocols BGP)

IN THIS SECTION

- Syntax | 11273
- Hierarchy Level | 11277
- Description | 11277
- Options | 11277
- Required Privilege Level | 11278
- Release Information | 11278

Syntax

```
neighbor address {
    accept-remote-nexthop;
    advertise-bgp-static
```

```
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    as-override;
    authentication-algorithm algorithm;
    authentication-key key;
    authentication-key-chain key-chain;
    cluster cluster-identifier;
    damping;
    description text-description;
    enforce-first-as;
    export [ policy-names ];
    family {
        (inet | inet6 | inet-mvpn | inet6-mpvn | inet-vpn | inet6-vpn | iso-vpn | l2-vpn) {
            (any | flow | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                    drop-excess <percentage>;
                    hide-excess <percentage>;
                }
                damping;
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                    drop-excess <percentage>;
                    hide-excess <percentage>;
                }
                rib-group group-name;
                topology name {
                    community {
                        target identifier;
                    }
                }
            }
            flow {
                no-validate policy-name;
            }
            labeled-unicast {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                    drop-excess <percentage>;
```

```
                hide-excess <percentage>;
            }
            aggregate-label {
                community community-name:
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
                drop-excess <percentage>;
                hide-excess <percentage>;
            }
            resolve-vpn;
            rib inet.3;
            rib-group group-name;
            topology name {
                community {
                    target identifier;
                }
            }
        }
    }
    forwarding-context
    route-target {
        advertise-default;
        external-paths number;
        accepted-prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
            drop-excess <percentage>;
            hide-excess <percentage>;
            }
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
            drop-excess <percentage>;
            hide-excess <percentage>;
            }
    }
    signaling {
        prefix-limit {
```

```
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
                drop-excess <percentage>;
                hide-excess <percentage>;
            }
        }
    }
    forwarding-context rti-name;
    graceful-restart {
        disable;
        restart-time seconds;
        stale-routes-time seconds;
    }
    hold-time seconds;
    import [ policy-names ];
    ipsec-sa ipsec-sa;
    keep (all | none);
    local-address address;
    local-as autonomous-system <private>;
    local-interface interface-name;
    local-preference preference;
    log-updown;
    metric-out (metric | minimum-igp <offset> | igp <offset>);
    mtu-discovery;
    multihop <ttl-value>;
    multipath {
        multiple-as;
    }
    no-aggregator-id;
    no-client-reflect;
    out-delay seconds;
    passive;
    peer-as autonomous-system;
    preference preference;
    rfc6514-compliant-safi129;
    tcp-aggressive-transmission;
    tcp-mss segment-size;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
```

```
    vpn-apply-export;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit protocols bgp group group-name],
[edit protocols bgp family prefix-limit (inet | inet6)(any | flow | labeled-unicast | multicast |
unicast)],
[edit routing-instances routing-instance-name protocols bgp group group-name]
```

Description

Explicitly configure a neighbor (peer). To configure multiple BGP peers, include multiple `neighbor` statements.

By default, the peer's options are identical to those of the group. You can override these options by including peer-specific option statements within the `neighbor` statement.

The `neighbor` statement is one of the statements you can include in the configuration to define a minimal BGP configuration on the routing device. (You can include an `allow all` statement in place of a `neighbor` statement.)

> **NOTE**: On MX Series routers configured with enhanced subscriber management, you can use this statement to statically provision a subscriber's client IP address as the BGP neighbor IP address. This is supported for only LNS subscribers. With enhanced subscriber management, you must also configure the `routing-services` statement at the `[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]` hierarchy level.

Options

address—IPv6 or IPv4 address of a single peer.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

BGP User Guide

neighbor (Protocols Layer 2 Circuit)

IN THIS SECTION

- Syntax | 11278
- Hierarchy Level | 11280
- Description | 11280
- Options | 11280
- Required Privilege Level | 11280
- Release Information | 11280

Syntax

```
neighbor address {
    interface interface-name {
        backup-neighbor address {
             community name;
            hot-standby;
            psn-tunnel-endpoint address;
            standby;
```

```
            virtual-circuit-id number;
        }
        bandwidth (bandwidth | ctnumber bandwidth);
        community community-name;
        (control-word | no-control-word);
        description text;
        egress-protection {
            protected-l2circuit {
                egress-pe address;
                ingress-pe address;
                virtual-circuit-id identifier;
            }
            protector-interface interface-name;
            protector-pe address {
                context-identifier identifier;
                lsp lsp-name;
                }
            }
        }
        encapsulation-type type;
        ignore-encapsulation-mismatch;
        ignore-mtu-mismatch;
        mtu mtu-number;
        no-revert;
        protect-interface interface-name;
        pseudowire-status-tlv hot-standby-vc-on;
        psn-tunnel-endpoint address;
        revert-time seconds;
        static {
            incoming-label label;
            outgoing-label label;
            send-oam;
        }
        switchover-delay milliseconds;
        virtual-circuit-id identifier;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit],
[edit protocols l2circuit]
```

Description

Each Layer 2 circuit is represented by the logical interface connecting the local provider edge (PE) router or switch to the local customer edge (CE) router or switch. All the Layer 2 circuits using a particular remote PE router or switch designated for remote CE routers or switches are listed under the `neighbor` statement (neighbor designates the PE router or switch). Each neighbor is identified by its IP address and is usually the end-point destination for the LSP tunnel (transporting the Layer 2 circuit).

Options

address—IP address of a neighboring router or switch.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Neighbor Interface for the Layer 2 Circuit

neighbor (Protocols LDP)

IN THIS SECTION

- Syntax | 11281
- Hierarchy Level | 11281
- Description | 11281
- Options | 11281
- Required Privilege Level | 11282
- Release Information | 11282

Syntax

```
neighbor neighbor-address;
```

Hierarchy Level

```
[edit protocols ldp]
```

Description

Configure an LDP targeted neighbor.

LDP sends a targeted hello message to the configured remote neighbor with a targeted request-send-targeted hello message (T) bit set. If the remote neighbor allows receipt of asymmetric hello messages, or if it is configured with the source address as the targeted neighbor, it responds with a targeted hello message. The receipt of a targeted hello message establishes a targeted adjacency with the remote neighbor as described in RFC 5036. Subsequently, a targeted session is established to the remote neighbor.

Options

neighbor-address	IP address of the remote LDP neighbor.

Required Privilege Level

routing

Release Information

Statement introduced before Junos OS Release 14.2R1.

RELATED DOCUMENTATION

Minimum LDP Configuration

Enabling and Disabling LDP

neighbor (Protocols RIPng)

IN THIS SECTION

- Syntax | 11282
- Hierarchy Level | 11283
- Description | 11283
- Options | 11283
- Required Privilege Level | 11283
- Release Information | 11283

Syntax

```
neighbor neighbor-name {
    import [ policy-names ];
    metric-in metric;
    receive <none>;
    route-timeout seconds;
    send <none>;
```

```
    update-interval seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng group group-name],
[edit protocols ripng group group-name],
[edit routing-instances routing-instance-name protocols ripng group group-name]
```

Description

Configure neighbor-specific RIPng parameters, thereby overriding the defaults set for the routing device.

Options

neighbor-name—Name of an interface over which a routing device communicates to its neighbors.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

neighbor (Protocols VPLS)

IN THIS SECTION

- Syntax | 11284
- Hierarchy Level | 11285
- Description | 11285
- Options | 11285
- Required Privilege Level | 11286
- Release Information | 11286

Syntax

```
neighbor neighbor-id {
    mac-pinning;
    associate-profile{
        dynamic-profile-name;
        profile-variable-set profile-variable-set-name;
    }
    backup-neighbor {...}
    community community-name;
    connection-protection;
    encapsulation-type type;
    ignore-encapsulation-mismatch;
    oam {
        bfd-liveness-detection {
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
```

```
            version (1 | automatic);
        }
        ping-interval;
    }
    pseudowire-status-tlv;
    psn-tunnel-endpoint address;
    revert-time seconds;

    static {
        incoming-label label;
        outgoing-label label;
        control-word;
        no-control-word;
    }
    switchover-delay milliseconds;
    vpls-id-list vc-id-numbers;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls mesh-
group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls],
[edit routing-instances instance-name protocols vpls mesh-group mesh-group-name]
```

Description

Specify each of the PE routers participating in the VPLS domain. Configuring this statement enables LDP for signaling VPLS.

> **NOTE**: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Options

neighbor-id—Specify the neighbor identifier for each PE router participating in the VPLS domain.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

The `pseudowire-status-tlv` option was added in Junos OS Release 10.0.

The `vpls-id-list` option was added in Junos OS Release 14.2 for MX Series routers to provide support for multiple pseudowires between the same pair of PEs in LDP-VPLS.

RELATED DOCUMENTATION

Configuring LDP Signaling for VPLS

nested-application (Security IDP)

IN THIS SECTION

- Syntax | 11287
- Hierarchy Level | 11287
- Description | 11287
- Options | 11287
- Required Privilege Level | 11287
- Release Information | 11287

Syntax

```
nested-application nested-application-name;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain protocol-binding]
[edit security idp custom-attack attack-name attack-type signature protocol-binding]
```

Description

Specify the nested application name.

Options

nested-application-name—Name of the nested application.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

nested-application-unknown

IN THIS SECTION

- Syntax | 11288
- Hierarchy Level | 11288
- Description | 11288

Required Privilege Level | 11288
Release Information | 11288

Syntax

```
nested-application-unknown
```

Hierarchy Level

```
[edit services aacl rule rule-name No Link Title term-name No Link Title]
```

Description

Enable AACL logging of flows for unknown nested applications.

Required Privilege Level

interface	To view this statement in the configuration.
interface-control	To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

nested-applications

IN THIS SECTION

Syntax | 11289

- Hierarchy Level | 11289
- Description | 11289
- Options | 11289
- Required Privilege Level | 11289
- Release Information | 11289

Syntax

```
nested-applications [ nested-application-names ];
```

Hierarchy Level

```
[edit services aacl rule rule-name term (Application Aware Access List) term-name from (Application
Aware Access List)]
```

Description

Identify one or more nested applications defined in the application identification configuration for inclusion as a match condition.

Options

nested-application-names—Identifiers of the nested applications.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

netconf

IN THIS SECTION

- Syntax | 11290
- Hierarchy Level | 11292
- Description | 11292
- Default | 11292
- Options | 11292
- Required Privilege Level | 11292
- Release Information | 11293

Syntax

```
netconf {
    flatten-commit-results;
    hello-message {
        yang-module-capabilities {
            advertise-native-yang-modules;
            advertise-custom-yang-modules;
            advertise-standard-yang-modules;
        }
    }
    netconf-monitoring {
        netconf-state-schemas {
            retrieve-custom-yang-modules;
            retrieve-standard-yang-modules;
        }
    }
    notification {
        interleave;
    }
    rfc-compliant;
    ssh {
        client-alive-count-max number;
        client-alive-interval seconds;
```

```
        connection-limit limit;
        port port;
        rate-limit limit;
    }
    tls {
        client-identity client-id {
            fingerprint fingerprint;
            map-type (san-dirname-cn | specified);
            username username;
        }
        default-client-identity {
            map-type (san-dirname-cn | specified);
            username username;
        }
        local-certificate local-certificate;
        traceoptions {
            file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
            flag name;
            level (all | error | info | notice | verbose | warning);
            no-remote-trace;
        }
    }
    traceoptions {
        file <filename> <files number> <match regular-expression> <size size> <world-readable |
 no-world-readable>;
        flag flag;
        no-remote-trace;
        on-demand;
    }
    version-1.1;
    yang-compliant;
    yang-modules {
        device-specific;
        emit-anyxml-in-rpc-output;
        emit-extensions;
        emit-family-ns-and-module-name;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure the NETCONF XML management protocol.

Default

If you do not include the `netconf` statement, NETCONF connections are not permitted.

Options

`flatten-commit-results` Suppress the `<commit-results>` XML subtree in the NETCONF server's response for `<commit>` operations. This statement must be configured in conjunction with the `rfc-compliant` statement.

> **NOTE**: You should configure the `flatten-commit-results` statement outside of a NETCONF session, for example, in the CLI. If you configure the statement in a NETCONF session and commit and synchronize the configuration, the RPC reply might return unclosed or mismatched XML tags.

version-1.1 Make NETCONF sessions compliant to RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*. This statement must be configured in conjunction with the `rfc-compliant` statement.

When you enable RFC 6242-compliant NETCONF sessions and both peers advertise the `:base:1.1` capability, the NETCONF session uses chunked framing instead of the end-of-message character sequence (]]>]]>) to define message boundaries.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

`flatten-commit-results` option added in Junos OS Release 21.2R1.

`version-1.1` statement added in Junos OS Release 24.4R1 and Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

traceoptions (NETCONF and Junos XML Protocol) | 29613

netconf-monitoring (NETCONF)

IN THIS SECTION

- Syntax | 11293
- Hierarchy Level | 11294
- Description | 11294
- Options | 11294
- Required Privilege Level | 11294
- Release Information | 11294

Syntax

```
netconf-monitoring {
    netconf-state-schemas {
        retrieve-custom-yang-modules;
        retrieve-standard-yang-modules;
    }
}
```

Hierarchy Level

```
[edit system services netconf]
```

Description

Configure NETCONF monitoring options.

Options

netconf-state-schemas Specify the schemas that the NETCONF server should return when a client application retrieves the list of supported schemas, which by default includes only the Junos OS native schema. The following statements are supported:

- `retrieve-custom-yang-modules`—Include third-party YANG modules installed on the device.
- `retrieve-standard-yang-modules`—Include standard YANG modules supported by the device, for example, OpenConfig modules.

Client applications can request the list of supported schemas by using the following RPC:

```
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get>
    <filter type="subtree">
      <netconf-state xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-monitoring">
        <schemas/>
      </netconf-state>
    </filter>
  </get>
</rpc>
```

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1 and Junos OS Evolved Release 21.1R1.

netconf (Remote Device Management)

IN THIS SECTION

- Syntax | 11295
- Hierarchy Level | 11295
- Description | 11296
- Options | 11296
- Required Privilege Level | 11298
- Release Information | 11298

Syntax

```
netconf {
    bulk-interval milliseconds;
    bulk-limit number;
    connection-retry-interval seconds;
    password password;
    port port-number;
    reconfigure-bulk-limit number;
    response-timeout seconds;
    response-timeout-count number;
    user-name name;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name system services remote-device-management service-
device device-name provisioning-method],
[edit system services remote-device-management service-device device-name provisioning-method]
```

Description

Configure the NETCONF XML management protocol as the method for provisioning and deprovisioning services on the remote device.

Options

bulk-interval *milliseconds*

(Optional) Interval during which multiple services are provisioned or deprovisioned based on the assigned dictionary before the configuration is committed to the service device. When the interval times out, the service actions are committed in bulk before additional actions for the device can take place. You can use the interval (together with the `bulk-limit` option) to optimize your service device configuration during scaled subscriber negotiation and service provisioning or subscriber termination and service deprovisioning.

> **NOTE**: The `bulk-interval` configuration is ignored when the `bulk-limit` is set to 1.

- **Range:** 500 through 5000
- **Default:** 1

bulk-limit *number*

(Optional) Maximum number of services provisioned or deprovisioned based on the assigned dictionary during the bulk interval before the configuration is committed to the service device. When the limit is reached, the service actions are committed in bulk before additional actions for the device can take place. You can use the limit (together with the `bulk-interval` option) to optimize your service device configuration during scaled subscriber negotiation and service provisioning or subscriber termination and service deprovisioning.

- **Range:** 1 through 1000
- **Default:** 1

connection-retry-interval *seconds*

(Optional) The interval between successive attempts to establish a NETCONF session with the remote device.

- **Range:** 1 through 10
- **Default:** 3

password *password*

Password used by the NETCONF protocol to access the remote device during service management. The maximum length of the password is 64 bytes.

> **NOTE**: If you change the password when any active subscriber services are mapped to the device, the change takes effect only when the device reconnects.

port *port-number*

(Optional) TCP port number for the NETCONF session.

> **NOTE**: You cannot change the port number for the device when any active subscriber services are mapped to it.

- **Range:** 1 through 65,535
- **Default:** 830 (NETCONF over SSH)

reconfigure-bulk-limit *number*

(Optional) When the device is reconfigured, this is the maximum number of services provisioned or deprovisioned on the service device for the access domain based on the assigned dictionary before the configuration is committed to the service device. When the limit is reached, the service actions are committed in bulk before additional actions for the device can take place.

If the access domain is not yet configured, this option has no effect, because there are no matching services to install.

- **Range:** 1 through 1000
- **Default:** 100

response-timeout *seconds*

(Optional) Period during which the device must respond to an attempt to provision or deprovision a service. For a subscriber service over a NETCONF connection, the timeout is a failure equivalent to an explicit failure response received from the device.

- **Range:** 1 through 180
- **Default:** 3

response-timeout-count *number*

(Optional) Number of consecutive response timeouts that can occur before the BNG takes action. The default action is to close and reopen the NETCONF connection.

- **Range:** 1 through 10
- **Default:** 3

user-name *name*

Name used to access the remote device during service management. The maximum length of the name is 64 bytes.

> **NOTE**: If you change the username when any active subscriber services are mapped to the device, the change takes effect only when the device reconnects.

Required Privilege Level

system—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Configuring Remote Device Management for Service Provisioning

Remote Device Services Manager (RDSM) Overview

network

IN THIS SECTION

- Syntax | 11299
- Hierarchy Level | 11299
- Description | 11299
- Options | 11299
- Required Privilege Level | 11299
- Release Information | 11299

Syntax

```
network ip-prefix</prefix-length>;
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet]
```

Description

Configure subnet information for an IPv4 address-assignment pool.

Options

ip-prefix—IP version 4 address or prefix value.

prefix-length—(Optional) Subnet mask.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management

network (Access)

IN THIS SECTION

- Syntax | 11300
- Hierarchy Level | 11300
- Description | 11300
- Required Privilege Level | 11300
- Release Information | 11300

Syntax

```
network
```

Hierarchy Level

```
[edit access address-assignment pool <name> family (inet | inet6)]
```

Description

Specify the IPv4 network address for the pool. This attribute is mandatory. For an IPv6 pool, you will set the IPv6 network prefix.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.4 of Junos OS.

RELATED DOCUMENTATION

Understanding Pass-Through Authentication

network (Address Pools)

IN THIS SECTION

- Syntax | 11301
- Hierarchy Level | 11301
- Description | 11301
- Options | 11302
- Required Privilege Level | 11302
- Release Information | 11302

Syntax

```
network {
    [network-prefix] {
        external-assigned;
    }
}
```

Hierarchy Level

```
[edit access address-assignment address-pools name family inet],
[edit access address-assignment address-pools name family inet6]
```

Description

Specify the network prefix for the address pool for IPv4 or IPv6 addresses.

> **NOTE**: At least one network prefix must be configured but you can configure more than one prefix.

Options

network-prefix	Network prefix (IPv4 or IPv6).

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Understanding Source IP Filtering with Address Pools in TDF Domains for IP-Based Subscribers

Configuring Address Pools for Source-IP Filtering of IP-Based Subscribers

network-domain

IN THIS SECTION

- Syntax | 11303
- Hierarchy Level | 11303
- Description | 11303
- Required Privilege Level | 11303
- Release Information | 11303

Syntax

```
network-domain;
```

Hierarchy Level

```
[edit fabric resources node-group node-group-name]
```

Description

(QFabric systems only) Designate a Node group to become a *network Node group*, which is used to route traffic between a QFabric system and external networks. The absence of the `network-domain` configuration statement implies that the Node group is a *server Node group*, which is used to group sets of Node devices that are connected to servers or storage devices.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Node Groups for the QFabric System

Understanding Node Groups

network-element (AAA Profile)

IN THIS SECTION

- Syntax | 11304
- Hierarchy Level | 11304
- Description | 11304
- Options | 11304
- Required Privilege Level | 11304
- Release Information | 11305

Syntax

```
network-element network-element-name;
```

Hierarchy Level

```
[edit unified-edge aaa profiles aaa-profile-name radius authentication],
[edit unified-edge aaa profiles aaa-profile-name radius accounting]
```

Description

Specify the network element providing policy management for TDF subscribers. The network element must already be defined at the `[edit access radius]` hierarchy level.

Options

network-element-name	Name of the network element.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring an AAA Profile

Understanding AAA Profiles

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

Configuring RADIUS Network Elements

network-element (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 11305
- Hierarchy Level | 11306
- Description | 11306
- Default | 11306
- Options | 11306
- Required Privilege Level | 11306
- Release Information | 11306

Syntax

```
network-element element-name {
    function function-name;
    peer peer-name {
        priority priority-value;
        <timeout seconds>;
```

```
    }
}
```

Hierarchy Level

```
[edit access diameter]
```

Description

Configure the Diameter network element, which is similar to a peer group that provides function-specific features including failover and load balancing. Specify the associated function that the network element supports. You can prioritize the peers to support failover or load balancing.

Default

By default, all network elements are available on every session PIC unless Diameter bindings are configured.

Options

element-name—Name of the network element.

- **Range:** Up to 32 alphanumeric characters

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | **5699**

network-element (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 11307
- Hierarchy Level | 11307
- Description | 11308
- Options | 11308
- Required Privilege Level | 11308
- Release Information | 11308

Syntax

```
network-element element-name {
    dne-origin realm realm-name <host hostname>
    forwarding {
        route dne-route-name {
            destination realm realm-name <host hostname> ;
            function [function-name] <partition partition-name>;
            metric route-metric;
        }
    }
    function function-name;
    peer peer-name {
        priority priority-number;
    }
}
```

Hierarchy Level

```
[edit diameter]
```

Description

Specify the transport layer Diameter configuration. The Diameter network element includes a list of routes reachable through the Diameter instance, associated functions, and prioritized Diameter peers.

Options

element-name—Name of the network element.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Network Elements

network-element (Subscriber Aware Policy Control)

IN THIS SECTION

- Syntax | 11309
- Hierarchy Level | 11309
- Description | 11309
- Options | 11309
- Required Privilege Level | 11310

Release Information | 11310

Syntax

```
network-element {
    element-name {
        session-pics {
            group {
                group-name {
                    [session-pic interface-name];
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit unified-edge tdf gateway-name diameter]
```

Description

Configure the Diameter network element associated with Diameter bindings for this TDF gateway.

NOTE: If you want to set up Diameter bindings for session PICs on the TDF gateway, contact Juniper Networks Professional Services for assistance.

Options

element-name

Name of the network element.

- **Range:** Up to 32 alphanumeric characters

> **NOTE**: The specified network element must already be configured on the TDF gateway at the `[edit access diameter network-element]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (TDF Gateway) | 5702

network-elements (RADIUS)

IN THIS SECTION

- Syntax | 11311
- Hierarchy Level | 11311
- Description | 11311
- Options | 11311
- Required Privilege Level | 11311
- Release Information | 11311

Syntax

```
network-elements name {
    server name {
        priority priority;
    }
    maximum-pending-reqs-limit number;
    pending-queue-watermark watermark;
    pending-queue-watermark-abate abate-watermark;
}
```

Hierarchy Level

```
[edit access radius]
```

Description

Configure a network element, which is a load-balanced group of RADIUS servers providing policy management for TDF subscribers.

Options

name Name of the network element.

- **Range:** Up to 31 alphanumeric characters.

The remaining statements are described separately.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Network Elements

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

network-isolation

IN THIS SECTION

- Syntax | **11312**
- Hierarchy Level | **11313**
- Description | **11313**
- Options | **11313**
- Required Privilege Level | **11314**
- Release Information | **11314**

Syntax

```
network-isolation {
  group name {
    detection {
      hold-time <down milliseconds><up milliseconds>;
      link-tracking {
        failure-type link;
        interface name;
        minimum-links minimum-links;
      }
      service-tracking {
        core-isolation;
      }
    }
    service-tracking-action (lacp-out-of-sync | link-down);
```

```
  }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure a network isolation group in an EVPN network to define the parameters the device uses to determine network isolation condition changes and act upon them.

A network isolation condition means a bridge domain or EVPN instance (EVI) becomes isolated from the EVPN core network, usually for one or more of the following reasons:

- The associated interconnecting physical or logical links are down.
- The device doesn't have routes to the remote endpoints in the EVI, bridge domain, or VLAN.

The parameters you define in a network isolation group comprise a network isolation group profile. You use the `network-isolation-profile group` *`name`* statement to apply the profile as follows:

- When you define the group with the `link-tracking` option, you can assign the network isolation profile to an EVPN instance (EVI) or to a particular bridge domain or VLAN in an EVI. See "link-tracking (Network Isolation Groups)" on page 9833 for details.
- When you define the group with the `service-tracking` option, you can assign the network isolation profile to an interface. See "service-tracking (Network Isolation Groups)" on page 16760 for details.

You can configure either the `link-tracking` option or the `service-tracking` option in a particular network isolation group, and assign that group as a network isolation profile. You can't define link tracking and service tracking parameters in the same network isolation group.

Options

group ***name*** (Required) Network isolation group name. You can define up to 64 network isolation groups.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 23.1R1.

`service-tracking` and `service-tracking-action` options added in Junos OS and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

network-isolation-profile | 11314

network-isolation-profile

IN THIS SECTION

- Syntax | 11314
- Hierarchy Level | 11315
- Description | 11315
- Options | 11316
- Required Privilege Level | 11316
- Release Information | 11316

Syntax

```
network-isolation-profile {
  group name;
}
```

Hierarchy Level

```
[edit bridge-domains name bridge-options],
[edit interfaces name],
[edit routing-instances name bridge-domains name bridge-options],
[edit routing-instances name switch-options],
[edit routing-instances name vlans name switch-options],
[edit switch-options],
[edit vlans name switch-options]
```

Description

On a provider edge (PE) device in an EVPN network, assign a network isolation group profile to an EVPN instance (EVI), to a particular bridge domain or VLAN in an EVI, or to an interface.

A network isolation group encompasses the parameters that the device uses to:

- Detect when an EVPN instance. a bridge domain, or a VLAN is in a network isolation state.
- Take actions on the associated interfaces after the device detects a network isolation status change.

You use the `network-isolation group name` statement at the `[edit protocols]` hierarchy level to configure the parameters that define a network isolation group. You can define link tracking or service tracking parameters with that statement.

Then configure the `network-isolation-profile` statement to assign the network isolation group as follows:

- If the network isolation group uses the `link-tracking` option, you can assign the network isolation profile to:
  - An EVPN instance (EVI).
  - A particular bridge domain or VLAN in an EVI.
- If the network isolation group uses the `service-tracking` option, you can assign the network isolation profile to an interface. The interface can be a Layer 2 (L2) interface that is:
  - A logical single link or aggregated Ethernet (AE) link bundle (LAG) to a single-homed customer edge device in the network.
  - An Ethernet segment identifier (ESI) LAG interface to a multihomed customer edge device in the network.
  - A physical interface.

> **NOTE**: You can configure either the `link-tracking` option or the `service-tracking` option in a particular network isolation group, and assign that group as a network isolation profile. You can't define link tracking and service tracking parameters in the same network isolation group.

Options

group ***name*** Group name of the network isolation profile you want to assign to the desired EVI, bridge domain, VLAN, or interface (depending on the configuration hierarchy level you use with this statement).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Determine IRB Interface State Changes from Local L2 Interface or Remote Connectivity Status in EVPN Fabrics

Layer 2 Interface Status Tracking and Shutdown Actions for EVPN Core Isolation Conditions

interface-state | **9064**

network-isolation | **11312**

network-management

IN THIS SECTION

- Syntax | 11317
- Hierarchy Level | 11317
- Description | 11317
- Options | 11317
- Required Privilege Level | 11318
- Release Information | 11318

Syntax

```
network-management {
    cluster-master;
}
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Define parameters for network management. To manage an SRX Series Services Gateway cluster through a non-fxp0 interface, use this command to define the node as a virtual chassis in NSM. This command establishes a single DMI connection from the primary node to the NSM server. This connection is used to manage both nodes in the cluster. Note that the non-fxp0 interface (regardless of which node it is present on) is always controlled by the primary node in the cluster. The output of a *<get-system-information>* RPC returns a *<chassis-cluster>* tag in all SRX Series Firewalls. When NSM receives this tag, it models SRX Series clusters as devices with autonomous control planes.

Options

`cluster-master`—Enable in-band management on the primary cluster node through NSM.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

network-option

IN THIS SECTION

- Syntax | 11318
- Hierarchy Level | 11318
- Description | 11319
- Options | 11319
- Required Privilege Level | 11319
- Release Information | 11319

Syntax

```
network-option (option-1 | option-2);
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Configure the Ethernet equipment clock (EEC) synchronization networking type.

NOTE: For SCB, this option is set with the `set chassis synchronization network-type (option-1 | option-2)` configuration command at the `[edit]` hierarchy level.

NOTE: For Junos OS releases 11.2R4 through 13.3R3 for MX240, MX480, and MX960 with SCB, SCBE, and SCBE2; and MX2010 and MX2020 with SCB or SCBE; you must execute the following commands after you change the network option at the `[edit chassis synchronization]` hierarchy level. This is because the loop bandwidth does not change automatically when you change the network option.

```
user@host# deactivate chassis synchronization
user@host# activate chassis synchronization
```

Options

Depending on the configuration of the Synchronization Status Messages (SSM) quality level, the network option functions in the following ways:

- `option-1`—Maps to the G.813 option 1 (EEC1).
- `option-2`—Maps to the G.812 type IV clock (EEC1).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

synchronization (MX Series) | **28795**

network-services

IN THIS SECTION

- Syntax | 11320
- Hierarchy Level | 11320
- Description | 11320
- Default | 11320
- Options | 11321
- Required Privilege Level | 11321
- Release Information | 11321

Syntax

```
network-services (ethernet | enhanced-ethernet | ip | enhanced-ip | lan);
```

Hierarchy Level

```
[edit chassis]
```

Description

Set the router's network services to a specific mode of operation. On MX240, MX480, and MX960 routers, MPC5E and MPC7E power on only if the network services mode configured is `enhanced-ip` or `enhanced-ethernet`.

MX2010 and MX2020 support only `enhanced-ip` and `enhanced-ethernet` network services modes.

Default

- MX80, MX104, MX2010, MX2020—`enhanced-ip`
- MX240, MX480, MX960—`ip`

Options

`ethernet`—Set the router's network services to Ethernet and use standard, compiled firewall filter format.

`enhanced-ethernet`—Set the router's network services to enhanced Ethernet and use enhanced mode capabilities. Only MPCs and MS-DPCs are powered on in the chassis.

`ip`—Set the router's network services to Internet Protocol and use standard, compiled firewall filter format.

`enhanced-ip`—Set the router's network services to enhanced Internet Protocol and use enhanced mode capabilities. Only MPCs and MS-DPCs are powered on in the chassis. Non-service DPCs do not work with enhanced network services mode options. This feature is enabled by default on MX80, MX104, MX2010, and MX2020 Universal Routing Platforms. For MX960 platform, `enhanced-ip` configuration must be enabled on both REs together. After committing the configuration, GRES configuration must be disabled and both REs must be rebooted to ensure that all the FPCs reboot and have the same network-services as REs. Any mismatch in network services between RE0, RE1, FPCs will lead to unexpected results.

`lan`—Set the router's network services to LAN and use standard, compiled firewall filter format. Reboot the system after setting the router's network services to LAN.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 8.5.

`enhanced-ethernet` and `enhanced-ip` options introduced in Junos OS Release 11.4.

`limited-ifl-scaling` option introduced in Junos OS Release 15.1R3 for MX Series routers.

RELATED DOCUMENTATION

- Network Services Mode Overview
- Firewall Filters and Enhanced Network Services Mode Overview
- Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers
- Configuring Enhanced IP Network Services for a Virtual Chassis

Limiting the Maximum Number of Logical Interfaces on MX Series Routers With MS-DPCs in Enhanced IP Network Services Mode

network-slices

IN THIS SECTION

- Syntax | 11322
- Hierarchy Level | 11322
- Description | 11323
- Required Privilege Level | 11323
- Release Information | 11323

Syntax

```
network-slices {
    guest-network-functions{
        gnf  id {
            control-plane-bandwidth-percent percent;
            description description;
            fpcs fpcs;
            af-name
            collapsed-forward
        }
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure Junos node slicing.

Junos node slicing enables a single MX Series router to be partitioned to appear as multiple, independent routers. Each partition has its own Junos control plane, which runs as a virtual machine (VM), and a dedicated set of line cards. Each partition is called a guest network function (GNF).

The remaining statements are explained separately.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Understanding Junos Node Slicing

Setting Up Junos Node Slicing

Setting Up Junos Node Slicing

network-slicing

IN THIS SECTION

- Syntax | 11324
- Hierarchy Level | 11324
- Description | 11324
- Required Privilege Level | 11324
- Release Information | 11324

Syntax

```
network-slicing {
    slice slice-name;
}
```

Hierarchy Level

```
[edit services]
```

Description

Specify the network slice(s) on the device.

Network slicing is the partitioning of a physical network into multiple logical networks. Each logical network is called a slice. By virtue of being a logical network, a slice is a designated set of network resources, such as interfaces, firewall filters, policers, virtual output queues, schedulers, shapers, traffic control profiles, etc. See *Hierarchical Class of Service for Network Slicing* for more information.

Required Privilege Level

- interface—To view this statement in the configuration
- interface-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

> *Hierarchical Class of Service for Network Slicing*
>
> slice (CoS Interfaces) | **28039**

network-summary-export

IN THIS SECTION

- Syntax | 11325
- Hierarchy Level | 11325
- Description | 11325
- Options | 11325
- Required Privilege Level | 11326
- Release Information | 11326

Syntax

```
network-summary-export policy-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area area-id],
[edit protocols ospf area area-id],
[edit routing-instances routing-instance-name protocols ospf area area-id]
```

Description

Apply an export policy that specifies which network-summary link-state advertisements (LSAs) are flooded into an OSPFv2 area.

Options

policy-name—Name of a policy configured at the `[edit policy-options policy-statement policy-name term term-name]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Import and Export Policies for Network Summaries Overview

Example: Configuring an OSPF Export Policy for Network Summaries

network-summary-import | **11326**

Routing Policies, Firewall Filters, and Traffic Policers User Guide

network-summary-import

IN THIS SECTION

- Syntax | **11326**
- Hierarchy Level | **11327**
- Description | **11327**
- Options | **11327**
- Required Privilege Level | **11327**
- Release Information | **11327**

Syntax

```
network-summary-import policy-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area area-id],
[edit protocols ospf area area-id],
[edit routing-instances routing-instance-name protocols ospf area area-id]
```

Description

Apply an import policy that specifies which routes learned from an OSPFv2 area are used to generate network-summary link-state advertisements to other areas.

Options

policy-name—Name of a policy configured at the `[edit policy-options policy-statement policy-name term term-name]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Import and Export Policies for Network Summaries Overview

Example: Configuring an OSPF Import Policy for Network Summaries

network-summary-export | **11325**

Routing Policies, Firewall Filters, and Traffic Policers User Guide

network (TDF Domain)

IN THIS SECTION

- Syntax | 11328
- Hierarchy Level | 11328
- Description | 11328
- Options | 11328
- Required Privilege Level | 11329
- Release Information | 11329

Syntax

```
network address net-mask;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name subscriber-exclude-prefix
family inet],
[edit unified-edge gateways tdf gateway-name domains domain-name subscriber-exclude-prefix
family inet6]
```

Description

Specify the network prefix of source IP addresses for uplink packets and destination IP addresses for downlink packets that do not undergo TDF processing.

Options

address	Network address for the network prefix to exclude.
net-mask	Netmask for the network prefix.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based and IFL-Based TDF Subscribers Overview

IP-Based Subscriber Setup Overview

next-hop (Access)

IN THIS SECTION

- Syntax | 11329
- Hierarchy Level | 11330
- Description | 11330
- Options | 11330
- Required Privilege Level | 11330
- Release Information | 11330

Syntax

```
next-hop next-hop;
```

Hierarchy Level

```
[edit routing-options access route ip-prefix</prefix-length>]
```

Description

Configure the next-hop address for an access route. Access routes are typically unnumbered interfaces.

Options

next-hop—Specific next-hop address you want to assign to the access route.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

next-hop (Access Internal)

IN THIS SECTION

- Syntax | 11331
- Hierarchy Level | 11331
- Description | 11331
- Options | 11331
- Required Privilege Level | 11331
- Release Information | 11331

Syntax

```
next-hop next-hop;
```

Hierarchy Level

```
[edit routing-options access-internal route ip-prefix</prefix-length>]
```

Description

Configure the next-hop address for an internal access route. Access routes are typically unnumbered interfaces.

> **NOTE**: Starting in Junos OS Release 15.1R4, the router no longer supports a configuration where a static route points to a next hop that is tied to a subscriber. Typically, this might occur when RADIUS assigns the next hop with the Framed-IP-Address attribute. An alternative to this misconfiguration is to have the RADIUS server provide a Framed-Route attribute that matches the static route.

Options

next-hop Specific next-hop address you want to assign to the internal access route.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

next-hop-based-tunnel (Routing Options)

IN THIS SECTION

- Syntax | 11332
- Hierarchy Level | 11332
- Description | 11332
- Required Privilege Level | 11333
- Release Information | 11333

Syntax

```
next-hop-based-tunnel;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options dynamic-tunnels gre next-hop-based-
tunnel],
[edit routing-options dynamic-tunnels gre next-hop-based-tunnel]
```

Description

Enable dynamic next hop based GRE tunnels.

Create a tunnel composite next hop for every dynamic GRE tunnel configured. The tunnel composite next hop includes the dynamic tunnel's encapsulation data and a VPN label (when chained composite next hop is not enabled). When this option is not configured, the default interface-based tunnel mode is enabled. By configuring this option, a device can scale up to 32,000 IP tunnels, which is otherwise restricted to the system limit on the number of interfaces supported.

At a given point in time, either the next-hop-based dynamic tunnel or the default interface-based dynamic GRE tunnel can exist on a device. A switchover from one tunnel mode to another deletes the existing tunnels and creates new tunnels in the new tunnel mode. As a result, a tunnel mode switchover can impact network performance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring MPLS-Signaled LSPs to Use GRE Tunnels

next-hop (Class-Of-Service)

IN THIS SECTION

- Syntax | 11333
- Hierarchy Level | 11333
- Description | 11334
- Options | 11334
- Required Privilege Level | 11334
- Release Information | 11334

Syntax

```
next-hop [ next-hop-name ];
```

Hierarchy Level

```
[edit class-of-service forwarding-policy next-hop-map map-name forwarding-class class-name]
```

Description

Specify the next-hop name or address to which to map forwarded traffic.

Options

next-hop-name—Next-hop alias or IP address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring CoS-Based Forwarding

next-hop (Dynamic Access Routes)

IN THIS SECTION

- Syntax | 11335
- Hierarchy Level | 11335
- Description | 11335
- Options | 11335
- Required Privilege Level | 11335
- Release Information | 11336

Syntax

```
next-hop next-hop;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
access route prefix],
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
rib routing-table-name access route prefix],
[edit dynamic-profiles profile-name routing-options access route prefix]
```

Description

Dynamically configure the next-hop address for an access route. Access routes are typically unnumbered interfaces.

The next-hop gateway can be specified explicitly in the framed route, as either the subscriber's fixed address (common for business subscribers) or 0.0.0.0. Alternatively, the absence of the gateway address implies address 0.0.0.0. The address 0.0.0.0, whether implicit or explicitly configured, resolves to the subscriber's assigned address (host route).

If the RADIUS Framed-Route attribute (22) or Framed-IPv6-Route attribute [99] does not specify the next-hop gateway—as is common—the variable representing the next-hop automatically resolves to the subscriber's IP address.

Options

next-hop—Either the specific next-hop address you want to assign to the access route or one of the following next-hop address predefined variables.

- For IPv4 access routes, use the variable, `$junos-framed-route-nexthop`. The route prefix variable is dynamically replaced with the value in Framed-Route RADIUS attribute [22].
- For IPv6 access routes, use the variable, `$junos-framed-route-ipv6-nexthop`. The variable is dynamically replaced with the value in Framed-IPv6-Route RADIUS attribute [99].

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*] and [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*] hierarchy levels introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring Dynamic Access Routes for Subscriber Management

next-hop (Forwarding Options)

IN THIS SECTION

- Syntax | 11336
- Hierarchy Level | 11336
- Description | 11337
- Options | 11337
- Required Privilege Level | 11337
- Release Information | 11337

Syntax

```
next-hop address;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family (inet | inet6) output interface interface-name]
```

Description

Specify the next-hop address for sending copies of packets to an analyzer.

> **NOTE**: You need not configure the option for QFX5220-128C and QFX5220-32CD routers.

Options

address—IP address of the next-hop router.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

Configuring Port Mirroring

next-hop-group

IN THIS SECTION

- Syntax | 11338
- Hierarchy Level | 11338
- Description | 11338
- Options | 11338

Required Privilege Level | 11339
Release Information | 11339

Syntax

```
next-hop-group group-name{
    group-type inet6;
    interface interface-name {
        next-hop ipv6-address;
    }
    next-hop-subgroup group-name{
        interface interface-name {
            next-hop ipv6-address;
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family inet6 output]
```

Description

Specify the next-hop group through which to send port-mirror traffic to an analyzer. This configuration enables multipacket port mirroring on MX Series routers with or without the use of a Tunnel PIC. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.

Options

group-name—Name of the next-hop group. Up to 30 next-hop groups are supported for the router. Each next-hop group must have at least two next-hop addresses.

interface-name—Name of the interface used to reach the next-hop destination.

ipv6-address—IPv6 address of the next-hop router. Each next-hop group supports up to 16 next-hop addresses. Each next-hop subgroup can have up to 16 next-hop groups.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

next-hop-group (Analyzer)

IN THIS SECTION

- Syntax | 11339
- Hierarchy Level | 11339
- Description | 11340
- Options | 11340
- Required Privilege Level | 11340
- Release Information | 11340

Syntax

```
next-hop-group next-hop-group-name;
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name ouput]
```

Description

Configure next-hop group through which the port-mirrored traffic is sent.

Options

next-hop-group-name Name of the next-hop group through which the port-mirrored traffic is sent.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

> *Example: Configuring Mirroring to Multiple Interfaces for Remote Monitoring of Employee Resource Use on EX9200 Switches*

next-hop-group (Forwarding Options)

IN THIS SECTION

- Syntax for Routers | **11341**
- Syntax for EX3400, EX4100, EX4100-F, EX4300-MP, and EX4400 Switches | **11341**
- Hierarchy Level | **11341**
- Description | **11341**
- Options | **11342**
- Required Privilege Level | **11342**
- Release Information | **11342**

Syntax for Routers

```
next-hop-group group-name {
    interface interface-name {
        next-hop address;
    }
    next-hop-subgroup subgroup-name {
        interface interface-name {
            next-hop address;
        }
    }
}
```

Syntax for EX3400, EX4100, EX4100-F, EX4300-MP, and EX4400 Switches

```
next-hop-group group-name {
  group-type layer-2;
  interface interface-name;
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Specify the next-hop address for sending copies of packets to an analyzer.

(This paragraph applies only to routers.) The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.

NOTE: (This Note applies only to routers.) In Junos OS releases earlier through Release 14.2, the `next-hop-group` statement is present in the `forwarding-options` stanza for a routing instance, but the `next-hop-group` statement is not allowed in a routing instance. In other words, in a routing instance, `[edit routing-instances routing-instance-name forwarding-options`

next-hop-group] is not supported. You will get an error message if you try to commit this type of configuration. Starting in Junos OS Release 14.2, the next-hop-group statement is not present in [edit routing-instances *routing-instance-name* forwarding-options].

Options

interface *interface-name*
Name of interface used to reach the next-hop destination.

group-type layer-2
(Applies only to EX switches) Group type of the next-hop group. The allowed value for group-type is layer-2.

next-hop *address*
(Applies only to routers) IP address of the next-hop router. Each next-hop group supports up to 16 next-hop addresses. Up to 30 next-hop groups are supported. Each next-hop group must have at least two next-hop addresses.

next-hop-group *group-name*
Name of the next-hop group. The next-hop group can contain Layer 2 members, Layer 3 members, and subgroups that are either unit list (mirroring packets to each interface) or load-balanced (mirroring packets to one of several interfaces).

- Routers—Up to 30 next-hop groups are supported. Each next-hop group must have at least two next-hop addresses.
- Switches—Up to 4 next-hop groups are supported.

next-hop-subgroup *subgroup-name*
Name of a next-hop subgroup. Subgroups are either unit list (mirroring packets to each interface) or load-balanced (mirroring packets to one of several interfaces)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 22.4R1 for EX3400, EX4100, EX4100-F, EX4300-MP, and EX4400 switches.

RELATED DOCUMENTATION

Configuring Next-Hop Groups to Use Multiple Interfaces to Forward Packets Used in Port Mirroring

Configuring Port Mirroring for Multiple Destinations

1:N Port Mirroring to Multiple Destinations on Switches

next-hop-group (Port Mirroring)

IN THIS SECTION

- Syntax | 11343
- Hierarchy Level | 11343
- Description | 11343
- Options | 11344
- Required Privilege Level | 11344
- Release Information | 11344

Syntax

```
next-hop-group group-name;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family (inet | vpls) output],
[edit forwarding-options port-mirroring instance instance-name family (inet | vpls) output]
```

Description

Specify the next-hop address for sending copies of packets to an analyzer. This configuration enables multipacket port mirroring on MX Series routers and EX Series switches without the use of a Tunnel PIC.

The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.

Options

group-name—Name of next-hop group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Port Mirroring with Next-Hop Groups

next-hop-group (Unknown Unicast Forwarding)

IN THIS SECTION

- Syntax | 11344
- Hierarchy Level | 11345
- Description | 11345
- Options | 11345
- Required Privilege Level | 11345
- Release Information | 11345

Syntax

```
next-hop-group group-name {
    group-type {
        layer-2;
```

```
    }
    interface interface-name {
        next-hop address;
    }
    next-hop-subgroup subgroup-name {
        interface interface-name;
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure a next-hop group to forward unknown unicast packets to a specific interface or interfaces.

Options

group-name—Name of the next-hop group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring Unknown Unicast Forwarding (ELS)

Configuring Mirroring on EX9200 Switches to Analyze Traffic (CLI Procedure)

next-hop-map

IN THIS SECTION

- Syntax | 11346
- Hierarchy Level | 11346
- Description | 11347
- Options | 11347
- Required Privilege Level | 11347
- Release Information | 11347

Syntax

```
next-hop-map map-name {
    forwarding-class class-name {
        discard;
        lsp-next-hop [ lsp-regular-expression ];
        next-hop [next-hop-name];
        non-lsp-next-hop;
    }
    forwarding-class-default {
        discard;
        lsp-next-hop [ lsp-regular-expression ];
        next-hop [next-hop-name];
        non-lsp-next-hop;
    }
}
```

Hierarchy Level

```
[edit class-of-service forwarding-policy]
```

Description

Specify the map for CoS forwarding routes.

Options

map-name—Map that defines next-hop routes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for QFX10000 Series switches in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

Configuring CoS-Based Forwarding

next-hop (PIM RPF Selection)

IN THIS SECTION

- Syntax | 11348
- Hierarchy Level | 11348
- Description | 11348
- Options | 11348
- Required Privilege Level | 11348
- Release Information | 11348

Syntax

```
next-hop next-hop-address;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address
source source-address],
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address
wildcard-source],
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-
list-addresses source source-address],
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-
list-addresses wildcard-source]
```

Description

Configure the specific next-hop address for the PIM group source.

Options

next-hop-address—Specific next-hop address for the PIM group source.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring PIM RPF Selection

next-hop (Protocols MPLS)

IN THIS SECTION

- Syntax | 11349
- Hierarchy Level | 11349
- Description | 11349
- Options | 11350
- Required Privilege Level | 11350
- Release Information | 11350

Syntax

```
next-hop (address | interface-name | address/interface-name);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
bypass],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
transit incoming-label],
[edit protocols mpls static-label-switched-path lsp-name bypass],
[edit protocols mpls static-label-switched-path lsp-name ingress],
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Description

Location of the next hop to the destination, specified as the IPv4 or IPv6 address of the next hop, the interface name (for point-to-point interfaces only), or the *address/interface-name* to specify an IP address on an operational interface.

Options

address—IPv4 or IPv6 address of the next-hop router.

> **NOTE**: IPv6 static LSPs are not supported at the `[edit protocols mpls static-label-switched-path lsp-name ingress]` hierarchy level.

interface-name—IP address of the outgoing interface. It must be a point-to-point interface. The name can be a simple or fully qualified domain name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for IPv6 addresses in static LSP configurations added in Junos OS Release 17.2R1.

Support for IPv4 or IPv6 lookup after poping the label added in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Configuring the Ingress Router for Static LSPs

next-hop (RPM)

IN THIS SECTION

- Syntax | 11351
- Hierarchy Level | 11351
- Description | 11351
- Required Privilege Level | 11351

Release Information | 11351

Syntax

```
next-hop next-hop;
```

Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Description

Specify the next-hop address to which the probe should be sent.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

probe | 13228

next-hop-service

IN THIS SECTION

- Syntax | 11352
- Hierarchy Level | 11352
- Description | 11352
- Options | 11352
- Required Privilege Level | 11353
- Release Information | 11353

Syntax

```
next-hop-service {
    inside-service-interface interface-name.unit-number;
    outside-service-interface interface-name.unit-number;
    outside-service-interface-type interface-type;
    service-interface-pool name;
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify interface names or a service interface pool for the forwarding next-hop service set. You cannot specify both a service interface pool and an inside or outside interface.

Options

`inside-service-interface` *`interface-name.unit-number`*—Name and logical unit number of the service interface associated with the service set applied inside the network.

`outside-service-interface` *`interface-name.unit-number`*—Name and logical unit number of the service interface associated with the service set applied outside the network.

outside-service-interface-type *interface-type*—Identifies the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local.

`service-interface-pool` *`name`*—Name of the pool of logical interfaces configured at the `[edit services service-interface-pools pool` *`pool-name`*`]` hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.

> **NOTE**: **service-interface-pool** is not applicable for IP reassembly configuration on L2TP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`service-interface-pool` option added in Junos OS Release 9.3.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

- Configuring Service Sets to be Applied to Services Interfaces

next-hop (Static MAC-Based VLANs)

IN THIS SECTION

- Syntax | 11354

- Hierarchy Level | 11354
- Description | 11354
- Options | 11354
- Required Privilege Level | 11354
- Release Information | 11354

Syntax

```
next-hop interface-name;
```

Hierarchy Level

```
[edit ethernet-switching-options static vlan vlan-name mac mac-address]
```

Description

Specify the next hop for the indicated Ethernet node.

Options

interface-name	Name of the next-hop interface.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Adding a Static MAC Address Entry to the Ethernet Switching Table

next-hop-tunnel

IN THIS SECTION

- Syntax | 11355
- Hierarchy Level | 11355
- Description | 11355
- Options | 11356
- Required Privilege Level | 11356
- Release Information | 11356

Syntax

```
next-hop-tunnel gateway-address ipsec-vpn vpn-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family-name]
```

Description

For the secure tunnel (st) interface, create entries in the Next-Hop Tunnel Binding (NHTB) table, which is used to map the next-hop gateway IP address to a particular IP Security (IPsec) Virtual Private Network (VPN) tunnel. NHTB allows the binding of multiple IPsec VPN tunnels to a single IPsec tunnel interface.

Options

- *gateway-address*—Next-hop gateway IP address.
- ipsec-vpn *vpn-name* —VPN to which the next-hop gateway IP address is mapped.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Interfaces

next-hop (VXLAN Routing)

IN THIS SECTION

- Syntax | 11357
- Hierarchy Level | 11357
- Description | 11357
- Options | 11357
- Required Privilege Level | 11358
- Release Information | 11358

Syntax

```
next-hop integer;
```

Hierarchy Level

```
[edit forwarding-options vxlan-routing]
```

Description

Configure the maximum number of next hops reserved for use in an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) overlay network.

Juniper Networks switches can store a maximum number of next hops for use in either an underlay network or an EVPN-VXLAN overlay network. By default, a switch allocates 8192 next hops for use in an EVPN-VXLAN network, which leaves the remaining next hops for use in an underlay network. For example, a QFX5110 switch can store a maximum of 45056 next hops. Subtracting the default setting of 8192 next hops from that maximum number leaves 36864 next hops for use in an underlay network. The `next-hop` configuration statement enables you to re-allocate the maximum number of next hops reserved for use in an EVPN-VXLAN network.

If you do not plan to use a switch in an EVPN-VXLAN network, you can set the `next-hop` configuration statement to 0, which allocates all next hops for use in an underlay network.

Conversely, if you plan to use a switch in an EVPN-VXLAN network, you can re-allocate a greater number of the next hops to the EVPN-VXLAN network.

> **NOTE**: Changing the default number of next hops reserved for use in an EVPN-VXLAN network causes the Packet Forwarding Engine to restart. Restarting the Packet Forwarding Engine interrupts all forwarding operations. Therefore, we strongly recommend using this configuration statement before the EVPN-VXLAN network becomes operational.

Options

integer Maximum number of next hops reserved for use in an EVPN-VXLAN overlay network on a switch.

- **Range:** QFX5110 switches: 0 through 45056. QFX5120 switches: 0 through 61440. The integer you specify must be a multiple of 4096—for example, 4096, 8192, 12288, and so on.
- **Default:** 8192

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

next-server

IN THIS SECTION

- Syntax | 11358
- Hierarchy Level | 11359
- Description | 11359
- Options | 11359
- Required Privilege Level | 11359
- Release Information | 11359

Syntax

```
next-server next-server;
```

Hierarchy Level

```
[edit system services dhcp],
[edit system services dhcp pool pool-id],
[edit system services dhcp static-binding mac-address]
```

Description

(J Series Services Routers only) Specify the IP address for the next DHCP server used for communication after a DHCP boot client establishes initial contact.

Options

next-server—The IP address of the DHCP server that is used as the “siaddr” in a DHCP protocol packet.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

next-table

IN THIS SECTION

- Syntax | 11360
- Hierarchy Level | 11360
- Description | 11360
- Required Privilege Level | 11360
- Release Information | 11360

Syntax

```
next-table table-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
bypass],
[edit protocols mpls static-label-switched-path lsp-name bypass]
```

Description

Specify the routing table used to resolve the IPv4 or IPv6 next hop for a bypass LSP.

NOTE:

- The transit LSP that's associated with a bypass LSP that uses `next-table` must use a label in the range of *1000000-1048575*
- The transit LSP that's associated with a bypass LSP that uses `next-table` must be configured with a `pop` function to allow for an IPv4 or IPv6 lookup in the specified routing table.
- The transit LSP that's associated with a bypass LSP that uses `next-table` cannot specify a `next-hop`. The `next-hop` and the `next-table` options are mutually exclusive.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Configuring the Ingress Router for Static LSPs

nextboot

IN THIS SECTION

- Syntax | 11361
- Hierarchy Level | 11361
- Description | 11361
- Required Privilege Level | 11361
- Release Information | 11362

Syntax

```
nextboot {
    media-device;
}
```

Hierarchy Level

```
[edit]
```

Description

You can use the `nextboot` function to check the current bootup device. This function lists the available boot device choice for the media devices and current settings.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Veriexec Overview

nexthop-hold-time

IN THIS SECTION

- Syntax | 11362
- Hierarchy Level | 11362
- Description | 11363
- Options | 11363
- Required Privilege Level | 11363
- Release Information | 11363

Syntax

```
nexthop-hold-time milliseconds;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name multicast-snooping-options]
```

Description

Accumulate outgoing interface changes in order to perform bulk updates to the forwarding table and the routing table. Delete the statement to turn off bulk updates.

Options

milliseconds—Set the hold time duration from 1 through 1000 milliseconds.

- **Range:** 1 through 1000 milliseconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Enabling Bulk Updates for Multicast Snooping

nexthop-learning

IN THIS SECTION

- Syntax | 11364
- Hierarchy Level | 11364
- Description | 11364
- Options | 11365
- Required Privilege Level | 11365
- Release Information | 11365

Syntax

```
nexthop-learning (disable | enable);
```

Hierarchy Level

```
[edit services flow-monitoring (version-ipfix |version9) template template-name]
```

Description

Enable learning of next-hop addresses to correctly report the next hop address, output SNMP, destination IP address, destination IP mask values, the packet loss priority, and the first two characters of the configured forwarding class name in the flow records when a destination is reachable through multiple paths. By default, this behavior of learning the next-hop addresses is disabled for inline active flow monitoring.

For PTX Series, starting in Junos OS Evolved 21.2R1 and in Junos OS Release 21.3R1, every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. The effect of the `nexthop-learning` statement on this behavior varies depending upon the operating system. For Junos OS Evolved, we do not recommend that you configure the `nexthop-learning` statement, as it reduces the number of packets that can be processed. For Junos OS, you can configure the `nexthop-learning` statement to change this default no-flow behavior and once again create and maintain flows, then attach the template to all sampling instances associated with FPCs that require the previous behavior.

When learning of next-hop addresses is disabled, data is reported as follows:

- If the destination address of the sampled IPv4 flow is reachable through multiple paths, the ipNextHopIPv4Address (Element ID 15) and egressInterface (Element ID 14) in the IPv4 flow record are set to the gateway IP address and SNMP index of the first path seen in the forwarding table.
- If the destination address of the sampled IPv6 flow is reachable through multiple paths, the ipNextHopIPv6Address (Element ID 62) and egressInterface (Element ID 14) in the IPv6 flow records are set to 0.
- The Incoming Interface (IIF) and Outgoing Interface (OIF) should be part of the same VRF. If the OIF is in a different VRF, destinationIPv4PrefixLength (Element ID 13), bgpDestinationAsNumber (Element ID 17), ipNextHopIPv4Address (Element ID 15), and egressInterface (Element ID 14) are set to 0 in IPv4 flow records and destinationIPv6PrefixLength (Element ID 30), bgpDestinationAsNumber (Element ID 17), ipNextHopIPv6Address (Element ID 62), and egressInterface (Element ID 14) are set to 0 in IPv6 flow records.

- The packet loss priority and forwarding class information is not reported for the PTX1000, PTX10008 (without the JNP10008-SF3), and PTX10016 routers.

When learning of next-hop addresses is enabled, output SNMP, destination IP address, destination IP mask values, the packet loss priority, and the first two characters of the configured forwarding class in the flow records are reported correctly. In addition, when enabled, mplsTopLabelIPv4Address (Element ID 47) in IPv4 flow records reports correctly when MPLS ingress sampling is enabled.

> **NOTE**: Nexthop learning is supported only when sampling is implemented on the PFE. This is known as inline active flow monitoring. Nexthop learning does not work when sampling is configured on the MS-DPC/MS-MPC/MS-MIC service cards.

Options

`disable`—Disable the learning of next-hop information required for inline active flow monitoring.

`enable`—Enable the learning of next-hop information required for inline active flow monitoring.

- **Default:** disable

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F2.

`nexthop-learning` is supported for the `bridge-template` statement in Junos OS Release 18.2R1 for MX Series routers.

Support introduced in Junos OS Release 20.3R1 to correctly report the packet loss priority and the first two characters of the configured forwarding class in the IPv4 and IPv6 IPFIX templates for the PTX1000, PTX10008 (without the JNP10008-SF3), and PTX10016 routers.

RELATED DOCUMENTATION

> *Configuring Next-Hop Address Learning on MX Series and PTX Series Routers for Destinations Accessible Over Multiple Paths*

ng-juniper

IN THIS SECTION

- Syntax | 11366
- Hierarchy Level | 11367
- Description | 11367
- Options | 11367
- Required Privilege Level | 11368
- Release Information | 11368

Syntax

```
base-filter {
    base-filter;
    ng-default-filter;
    }
category category-name {
    action (block | log-and-permit | permit | quarantine);
    custom-message custom-message;
    }
custom-message custom-message;
default (block | log-and-permit | permit | quarantine);
fallback-settings {
    default (block | permit);
    server-connectivity (block | log-and-permit);
    timeout (block | log-and-permit);
    too-many-requests (block | log-and-permit);
    }
no-safe-search;
server {
    host host;
    port port;
    proxy-profile proxy-profile;
    routing-instance routing-instance;
    source-address source-address;
    tls-profile tls-profile;
```

```
        }
    site-reputation-action (very-safe | moderately-safe | fairly-safe | suspicious | harmful);
    timeout timeout;
```

Hierarchy Level

```
[edit security utm default-configuration web-filtering ng-juniper]
[edit security utm feature-profile web-filtering ng-juniper profile name]
```

Description

Configure the Juniper NextGen Web filtering engine. Juniper NextGen Web filtering acts as a gateway for the SRX Series devices seeking URL reputation or category from the Juniper NextGen Web Filtering (NGWF) cloud.

Options

base-filter Juniper base filter.

ng-default-filter Juniper default base filter.

category *name* Juniper NextGen category name.

action Action to perform when web traffic matches category. The possible options are block, log and permit, permit, and quarantine.

custom-message Custom message for the action taken when web traffic matches category.

default Juniper NextGen default profile.

fallback-settings Juniper NextGen fallback settings.

no-safe-search Do not perform safe-search for Juniper NextGen protocol.

server Configure Juniper NextGen server.

host Server host IP address or string host name.

port Server port number.

- Range: 1 thru 65535

proxy-profile Proxy profile name.

routing-instance Routing instance name.

source-address Source IP address used to connect the server.

tls-profile SSL initiation profile.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

> **NOTE**: You must configure SSL initiation profile for NGWF to communicate through HTTPS connections.

```
user@host#set security utm default-configuration web-filtering ng-juniper server ?
Possible completions:
  <[Enter]>            Execute this command
+ apply-groups         Groups from which to inherit configuration data
+ apply-groups-except  Don't inherit configuration data from these groups host Server host IP
address or string host name
  port                 Server port (1..65535)
  proxy-profile        Proxy profile
  routing-instance     Routing instance name
  source-address       Source IP address used to connect server tls-profile SSL initiation
profile
  tls-profile       SSL initiation profile
```

The following is a configuration sample for SSL initiation profile:

1. Create a self-signed certificate that can be used for SSL handshake with the Content Security server.

```
request security pki generate-key-pair certificate-id utmcert size 1024 type rsa
request security pki local-certificate generate-self-signed certificate-id utmcert subject
"DC=Domain_component,CN=utmcert,OU=SLT_QA,O=Juniper,L=Sunnyvale,ST=CA,C=US" ip-address 0.0.0.0 domain-
name juniper.net
```

2. 2. Configure the SSL initiation profile by calling the certificate that you created.

```
set services ssl initiation profile ssl_init_prof client-certificate utmcert
set services ssl initiation profile ssl_init_prof actions ignore-server-auth-failure
set services ssl initiation profile ssl_init_prof trusted-ca all
```

The following is a sample of the configurations that are required for NGWF to work:

```
set security policies from-zone trust to-zone untrust policy fw_policy match source-address any
set security policies from-zone trust to-zone untrust policy fw_policy match destination-address any
set security policies from-zone trust to-zone untrust policy fw_policy match application any
set security policies from-zone trust to-zone untrust policy fw_policy then permit application-services ssl-
proxy profile-name ssl-profile
set security policies from-zone trust to-zone untrust policy fw_policy then permit application-services utm-
policy WF
```

ngfw

IN THIS SECTION

- Syntax | **11370**
- Hierarchy Level | **11370**
- Description | **11370**
- Options | **11371**

Required Privilege Level | 11371
Release Information | 11371

Syntax

```
ngfw {
    default-profile {
        application-traffic-control {
            rule-set rule-set;
        }
        ssl-proxy {
            profile-name profile-name;
        }
    }
}
```

Hierarchy Level

```
[edit security],
[edit security logical-systems logical-system-name]
```

Description

Specify a default profile to manage conflicts when a security policy lookup returns a list of policies before the final application is identified.

The initial policy lookup phase occurs prior to identifying a dynamic application. If there are multiple policies present in the potential policy list that contain different SSL proxy profiles, then the SRX Series Firewall applies the default profile until a more explicit match has occurred.

You can configure a default profile for an SSL proxy and for an application quality of service (AppQoS) under the `[edit security ngfw]` hierarchy level.

You can configure an SSL proxy profile under the `[edit services ssl proxy]` hierarchy level, which can be applied as the default SSL proxy profile under the `[edit security ngfw]` hierarchy level. Similarly, you can configure application traffic rule sets under the `[edit class-of-service]` hierarchy level, and apply the rule set under the `[edit security ngfw]` hierarchy level as the default AppQoS rule set.

Options

application-traffic-control Specify the application traffic control rule as the default rule.

- **rule-set *rule-set*** Rule set name of the application traffic control.

ssl-proxy Specify the SSL forward proxy profile or the SSL reverse proxy profile as the default profile.

- **profile-name *profile-name*** Name of the SSL forward proxy profile or the SSL reverse proxy profile.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

Firewall User Authentication Overview

no-accounting

IN THIS SECTION

- Syntax | 11372
- Hierarchy Level | 11372
- Description | 11372
- Required Privilege Level | 11372

Release Information | 11372

Syntax

```
no-accounting;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name]
```

Description

Disable the collection of IGMP join and leave event statistics on a per-interface basis.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Recording IGMP Join and Leave Events

no-adaptation

IN THIS SECTION

- Syntax | 11373
- Hierarchy Level | 11373
- Description | 11373
- Required Privilege Level | 11373
- Release Information | 11374

Syntax

```
no-adaptation;
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group group-name liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd]
```

Description

Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients
- Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

no-adaptation (BFD Liveness Detection)

IN THIS SECTION

- Syntax | 11374
- Hierarchy Level | 11374
- Description | 11375
- Required Privilege Level | 11375
- Release Information | 11376

Syntax

```
no-adaptation;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls)  oam bfd-liveness-
detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-
id oam bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-
```

```
group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)oam
bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn |
vpls)neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-
group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-
detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-name neighbor
neighbor-id oam bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-
detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam
bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-
name neighbor neighbor-id oam bfd-liveness-detection]
```

Description

Configure BFD sessions not to adapt to changing network conditions. We recommend that you *do not* disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. However, include the `no-adaptation` statement in the configuration if you do not want BFD sessions to adapt to changing network conditions.

You can use the `clear bfd adaptation` command to return BFD interval timers to their configured values. The `clear bfd adaptation` command does not affect traffic flow on the routing device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0

Support for BFD authentication introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD for Layer 2 VPN and VPLS

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

bfd-liveness-detection (Layer 2 VPN and VPLS) | **2929**

no-adaptation (Liveness Detection)

IN THIS SECTION

- Syntax | **11376**
- Hierarchy Level | **11376**
- Description | **11377**
- Required Privilege Level | **11377**
- Release Information | **11377**

Syntax

```
no-adaptation;
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection]
```

Description

Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

no-adaptation (PIM BFD Liveness Detection)

IN THIS SECTION

- Syntax | 11377
- Hierarchy Level | 11378
- Description | 11378
- Required Privilege Level | 11378
- Release Information | 11378

Syntax

```
no-adaptation;
```

Hierarchy Level

```
[edit protocols pim interface interface-name bfd-liveness-detection],
[edit routing-instances routing-instance-name protocols pim interface interface-name bfd-
liveness-detection]
```

Description

Configure BFD sessions not to adapt to changing network conditions. We recommend that you *do not* disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0

Support for BFD authentication introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD for PIM

bfd-liveness-detection (Protocols PIM) | **2945**

no-adaptive

IN THIS SECTION

- Syntax | **11379**
- Hierarchy Level | **11379**
- Description | **11379**

- Required Privilege Level | 11379
- Release Information | 11380

Syntax

```
no-adaptive;
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options load-balance],
[edit dynamic-profiles name interfaces name logical-tunnel-options load-balance],
[edit dynamic-profiles name interfaces interface-range name aggregated-ether-options load-
balance],
[edit dynamic-profiles name interfaces interface-range name logical-tunnel-options load-balance],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-ether-options load-
balance],
[edit dynamic-profiles name logical-systems name interfaces name logical-tunnel-options load-
balance],
[edit dynamic-profiles name logical-systems name interfaces interface-range name aggregated-
ether-options load-balance],
[edit dynamic-profiles name logical-systems name interfaces interface-range name logical-tunnel-
options load-balance],
[edit interfaces name aggregated-ether-options load-balance],
[edit interfaces name logical-tunnel-options load-balance],
[edit interfaces interface-range              name aggregated-ether-options load-balance],
[edit interfaces interface-range              name logical-tunnel-options load-balance]
```

Description

Disables the adaptive load-balancing solution configured on the aggregated Ethernet bundle to distribute traffic by using a feedback mechanism.

Required Privilege Level

interface - To view statement in the configuration.

interface-control - To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R3.

RELATED DOCUMENTATION

Understanding Aggregated Ethernet Load Balancing

no-adjacency-down-notification (Protocols IS-IS)

IN THIS SECTION

- Syntax | 11380
- Hierarchy Level | 11380
- Description | 11381
- Required Privilege Level | 11381
- Release Information | 11381

Syntax

```
no-adjacency-down-notification;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis interface interface-name],
[edit protocols isis interface interface-name]
```

Description

Disable adjacency down notification for IS-IS to allow for migration from IS-IS to OSPF without disruption of the RSVP neighbors and associated RSVP-signaled label-switched paths (LSPs).

Whenever IS-IS is deactivated, the IS-IS adjacencies are brought down. IS-IS signals to RSVP to bring down any RSVP neighbors associated with the IS-IS adjacencies, and this further causes the associated LSPs signaled by RSVP to go down as well.

A similar process occurs whenever OSPF is deactivated. The OSPF neighbors are brought down. OSPF signals to RSVP to bring down any of the RSVP neighbors associated with the OSPF neighbors, and this further causes the associated LSPs signaled by RSVP to go down as well.

If you need to migrate from IS-IS to OSPF or from OSPF to IS-IS, the internal gateway protocol (IGP) notification to RSVP for an adjacency or neighbor down event needs to be ignored. Using the `no-adjacency-down-notification` or `no-neighbor-down-notification` statements, you can disable IS-IS adjacency down notification or OSPF neighbor down notification, respectively, until the migration is complete. The network administrator is responsible for configuring the statements before the migration, and then removing them from the configuration afterward, so that IGP notification can function normally.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

no-neighbor-down-notification | **11519**

no-advertise-adjacency-segment (Protocols IS-IS)

IN THIS SECTION

- Syntax | 11382
- Hierarchy Level | 11382
- Description | 11382
- Default | 11382
- Required Privilege Level | 11383
- Release Information | 11383

Syntax

```
no-advertise-adjacency-segment;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis
interface interface-name],
[edit protocols isis interface interface-name],
[edit routing-instances routing-instance-name protocols isis interface interface-name]
```

Description

Disable advertising of the adjacency segment on all levels for the specified interface.

Default

Enabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

isis | 9401

source-packet-routing (Protocols IS-IS) | 28266

no-advertise-adjacency-segment (Protocols OSPF)

IN THIS SECTION

- Syntax | 11383
- Hierarchy Level | 11384
- Description | 11384
- Default | 11384
- Required Privilege Level | 11384
- Release Information | 11384

Syntax

```
no-advertise-adjacency-segment;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area name interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area name interface interface-name],
[edit protocols ospf area name interface interface-name],
[edit routing-instances routing-instance-name protocols ospf area name interface interface-name]
```

Description

Disable advertising of the adjacency segment for the specified interface.

Default

Enabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

OSPF Overview

source-packet-routing

no advertise-peer-as

IN THIS SECTION

- Syntax | 11385
- Hierarchy Level | 11385
- Description | 11385
- Required Privilege Level | 11386
- Release Information | 11386

Syntax

```
no-advertise-peer-as;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Enable the default behavior of suppressing AS routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring BGP Route Advertisement

Configuring Routing Policies to Control BGP Route Advertisements

advertise-peer-as | **2020**

no-aggregate-delegate-processing

IN THIS SECTION

- Syntax | **11386**
- Hierarchy Level | **11387**
- Description | **11387**
- Default | **11387**
- Required Privilege Level | **11387**
- Release Information | **11387**

Syntax

```
no-aggregate-delegate-processing;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

Description

Disable distribution of connectivity fault management (CFM) sessions on aggregated Ethernet interfaces.

Default

CFM sessions on aggregated Ethernet interfaces are distributed by default.

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

IEEE 802.1ag OAM Connectivity Fault Management Overview

no-aggregator-id

IN THIS SECTION

- Syntax | 11388
- Hierarchy Level | 11388
- Description | 11388

- Default | 11389
- Required Privilege Level | 11389
- Release Information | 11389

Syntax

```
no-aggregator-id;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],

[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Prevent different routing devices within an AS from creating aggregate routes that contain different AS paths.

Junos OS performs route aggregation, which is the process of combining the characteristics of different routes so that only a single route is advertised. Aggregation reduces the amount of information that BGP must store and exchange with other BGP systems. When aggregation occurs, the local routing device adds the local AS number and the router ID to the aggregator path attiribute. The `no-aggregator-id`

statement causes Junos OS to place a 0 in the router ID field and thus eliminate the possibility of having multiple aggregate advertisements in the network, each with different path information.

Default

If you omit this statement, the router ID is included in the BGP aggregator path attribute.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

BGP Messages Overview

no-allow-link-events

IN THIS SECTION

- Syntax | 11390
- Hierarchy Level | 11390
- Description | 11390
- Required Privilege Level | 11390
- Release Information | 11390

Syntax

```
no-allow-link-events;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management
interface interface-name negotiation-options]
```

Description

Disable the sending of link event TLVs.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Ethernet OAM Link Fault Management

no-allow-snooped-clients

IN THIS SECTION

- Syntax | 11391
- Hierarchy Level | 11391
- Description | 11391

Required Privilege Level | 11391
Release Information | 11392

Syntax

```
no-allow-snooped-clients;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Explicitly disable DHCP snooping support on DHCP relay agent.

Use the statement at the `[edit ... dhcpv6]` hierarchy levels to explicitly disable snooping support for DHCPv6 relay agent.

> **NOTE**: In Junos OS Release 10.0 and earlier, DHCP snooping is *enabled* by default. In Release 10.1 and later, DHCP snooping is *disabled* by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- *Extended DHCP Relay Agent Overview*
- *Overriding the Default DHCP Relay Configuration Settings*
- DHCP Snooping Support
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent

no-allow-url (Op Scripts)

IN THIS SECTION

- Syntax | 11392
- Hierarchy Level | 11392
- Description | 11393
- Default | 11393
- Required Privilege Level | 11393
- Release Information | 11393

Syntax

```
no-allow-url;
```

Hierarchy Level

```
[edit system scripts op]
```

Description

Prohibit the remote execution of Junos OS op scripts. When you include this configuration statement, the `op url` operational mode command does not support executing op scripts from a remote site and generates an error.

Default

If you do not include the `no-allow-url` statement, authorized users can issue the `op url` command to execute op scripts from a remote site.

> **NOTE**: You must configure the `allow-url-for-python` statement in order to execute Python automation scripts from a remote site. The `no-allow-url` statement takes precedence over the `allow-url-for-python` statement; if both statements are configured, you cannot execute Python op scripts from a remote site on the device.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

file (Op Scripts) | 7067

Executing an Op Script from a Remote Site

no-allowed-mac-log

IN THIS SECTION

- Syntax | 11394
- Hierarchy Level | 11394

- Description | 11394
- Default | 11394
- Required Privilege Level | 11394
- Release Information | 11394

Syntax

```
no-allowed-mac-log;
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port interface (all | interface-name)]
```

Description

Specify that the switch does not log messages when it receives packets from invalid MAC addresses on an interface that has been configured for particular (allowed) MAC addresses.

Default

The switch logs messages when it receives packets from invalid MAC addresses on an interface that has been configured for particular (allowed) MAC addresses.

Required Privilege Level

system—To view this statement in the configuration.
system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

allowed-mac | 2214

Example: Configuring Port Security (non-ELS)

Example: Protecting Against DHCP Snooping Database Attacks

Example: Protecting against DHCP Starvation Attacks

Configuring MAC Limiting (non-ELS)

no-anti-replay

IN THIS SECTION

- Syntax | 11395
- Hierarchy Level | 11395
- Description | 11395
- Required Privilege Level | 11396
- Release Information | 11396

Syntax

```
no-anti-replay;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Disable IPsec antireplay service, which occasionally causes interoperability issues for security associations.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

no-anti-replay (Services Service Set)

IN THIS SECTION

- Syntax | 11396
- Hierarchy Level | 11396
- Description | 11397
- Required Privilege Level | 11397
- Release Information | 11397

Syntax

```
no-anti-replay;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Disable IPsec antireplay service for this service set, which occasionally causes interoperability issues for security associations. This statement is useful for dynamic endpoint tunnels for which you cannot configure the `no-anti-reply` statement at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level.

For static IPsec tunnels, this statement disables the antireplay check for all the tunnels within this service set. If antireplay check has to be enabled for a particular tunnel, then set the `anti-replay-window-size` statement at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level.

> **NOTE**: Setting the `anti-replay-window-size` and `no-anti-replay` statements at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level overrides the settings specified at the `[edit services service-set service-set-name ipsec-vpn-options]` hierarchy level.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring IPsec Service Sets

Configuring IPsec Rules

no-application-identification (Services)

IN THIS SECTION

- Syntax | 11398
- Hierarchy Level | 11398

- Description | 11398
- Required Privilege Level | 11398
- Release Information | 11398

Syntax

```
no-application-identification;
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Disable the application identification of applications running on nonstandard ports. By default, application identification is enabled on the device. You can disable application identification by using the following command:

```
user@host# set services application-identification no-application-identification
```

If you want to reenable application identification, delete the configuration statement that specifies disabling of application identification by using the following command:

```
user@host# delete services application-identification no-application-identification
```

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Disabling and Reenabling Junos OS Application Identification

no-application-system-cache

IN THIS SECTION

- Syntax | 11399
- Hierarchy Level | 11399
- Description | 11399
- Required Privilege Level | 11400
- Release Information | 11400

Syntax

```
no-application-system-cache;
```

Hierarchy Level

```
[edit services application-identification],
[edit services application-identification nested-application-settings]
```

Description

Application identification information is saved in the application system cache to improve performance. This cache is updated when a different application is identified. This caching is turned on by default. Use the `no-application-system-cache` statement to turn it off.

ASC is enabled by default when a session is created. You can manually turn this caching off using the `set services application-identification no-application-system-cache` command. You can re-enable the ASC by using the `set services application-identification application-system-cache` command.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support for Next Gen Services introduced in Junos OS Release 19.3R2 and 19.4R1 on MX Series routers MX240, MX480 and MX960.

Statement deprecated in Junos OS Release 12.3X48-D10 at the [edit services application-identification nested-application-settings] hierarchy level.

RELATED DOCUMENTATION

Configuring Global APPID Properties

Application Identification for Nested Applications

no-application-system-cache (Services)

IN THIS SECTION

- Syntax | 11400
- Hierarchy Level | 11401
- Description | 11401
- Required Privilege Level | 11401
- Release Information | 11401

Syntax

```
no-application-system-cache;
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Application identification information is saved in the application system cache to improve performance. This cache is updated when a different application is identified. This caching is turned on by default. Use the `no-application-system-cache` statement to turn it off.

ASC is enabled by default when a session is created. You can manually turn this caching off using the `set services application-identification no-application-system-cache` command. You can re-enable the ASC by using the `set services application-identification application-system-cache` command.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Deactivating Application System Cache Information for Application Identification (CLI Procedure)

no-attribute-length-in-pdu

IN THIS SECTION

- Syntax | 11402
- Hierarchy Level | 11402
- Description | 11402

Required Privilege Level | 11402
Release Information | 11403

Syntax

```
no-attribute-length-in-pdu;
```

Hierarchy Level

```
[edit protocols mvrp]
```

Description

Include an extra byte in protocol data units (PDUs) sent by the Multiple VLAN Registration Protocol (MVRP). You can disable the extra byte to address a compatibility issue between MVRP in Junos OS Releases 13.2 and later for EX Series switches with support for the Enhanced Layer 2 Software (ELS), which includes the extra byte, and MVRP in Junos OS Releases 11.3 and later for EX Series switches that do not support ELS, which does not include the extra byte. If this compatibility issue arises, the ELS version of MVRP does not recognize PDUs without the extra byte sent by the non-ELS version of MVRP.

You can recognize an MVRP version compatibility issue by observing the switch running the ELS version of MVRP. Because a switch running the ELS version of MVRP cannot interpret an unmodified PDU from a switch running the non-ELS version of MVRP, the switch will not add VLANs from the non-ELS version of MVRP. When you execute the command `show mvrp statistics` in the ELS version of MVRP, the values for `Received Join Empty` and `Received Join In` will incorrectly display zero, even though the value for the `Received MVRP PDUs without error` has been increased. Another indication that MVRP is having a version compatibility issue is that unexpected VLAN activity, such as multiple VLAN creation, takes place on the switch running the ELS version of MVRP.

Required Privilege Level

routing—To view this statement in the configuration.
routing control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Configuring Multiple VLAN Registration Protocol (MVRP) on Switches

Understanding Multiple VLAN Registration Protocol (MVRP)

Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

no-authentication-check (Protocols RSVP)

IN THIS SECTION

- Syntax | 11403
- Hierarchy Level | 11403
- Description | 11403
- Required Privilege Level | 11404
- Release Information | 11404

Syntax

```
no-authentication-check;
```

Hierarchy Level

```
[edit protocols rsvp]
```

Description

Skip authentication check for received messages.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

RSVP Authentication

no-auto-conversion

IN THIS SECTION

- Syntax | 11404
- Hierarchy Level | 11404
- Description | 11405
- Required Privilege Level | 11405
- Release Information | 11405

Syntax

```
no-auto-conversion;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Disable automatic Virtual Chassis port (VCP) conversion in a Virtual Chassis.

> **NOTE**: This statement is available as of Junos OS Releases 14.1X53-D47, 17.4R2, 18.1R3, 18.2R2, and 18.3R1 only on EX Series and QFX Series switches in a Virtual Chassis that have automatic VCP conversion enabled by default, which include all EX4300, EX4600, EX4650, and QFX Series switches that support Virtual Chassis. See for details.

Configuring this statement disables the feature for the entire Virtual Chassis, and subsequently removing this statement returns the Virtual Chassis to the default behavior with automatic VCP conversion enabled.

VCP links connect members in a Virtual Chassis using supported VCP ports on the member switches. Ports that are supported as VCPs and are not dedicated VCPs or configured as VCPs by default must be converted into VCPs to use them to interconnect Virtual Chassis members. You can manually convert a supported port into a VCP port using the `request virtual-chassis vc-port` command, or when automatic VCP conversion is enabled, a port will be automatically converted into a VCP under the following conditions:

- LLDP is enabled on the interfaces for the members on both sides of the link. The two interfaces exchange LLDP packets to accomplish the port conversion.
- The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the `set virtual-chassis member` command.
- The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

If you want to deterministically control VCP port conversion during a preprovisioned Virtual Chassis configuration or expansion, you might want to disable this feature and set up all VCPs manually. Otherwise, you can use the default automatic VCP conversion behavior in a “plug and play” approach to simplify adding a new switch to an existing Virtual Chassis or adding a redundant VCP link between two existing members of a Virtual Chassis.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 14.1X53-D47, 17.4R2, 18.1R3, 18.2R2, and 18.3R1.

RELATED DOCUMENTATION

Understanding Virtual Chassis Components

Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port

Adding a New Switch to an Existing EX2300, EX3400, or EX4300 Virtual Chassis

Adding an EX4600 Switch to a Mixed or Non-mixed Virtual Chassis

Adding a New Switch to an Existing EX4650 or QFX Series Virtual Chassis

no-auto-failover

IN THIS SECTION

- Syntax | 11406
- Hierarchy Level | 11406
- Description | 11406
- Default | 11407
- Required Privilege Level | 11407
- Release Information | 11407

Syntax

```
no-auto-failover;
```

Hierarchy Level

```
[edit chassis redundancy feb redundancy-group group-name]
```

Description

Disable automatic failover to a backup FEB when an active FEB in a redundancy group fails.

Default

Automatic failover is enabled by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring FEB Redundancy on the M120 Router

no-auto-iccp-provisioning (Junos Fusion Redundancy Group)

IN THIS SECTION

- Syntax | 11407
- Hierarchy Level | 11408
- Description | 11408
- Default | 11408
- Required Privilege Level | 11408
- Release Information | 11408

Syntax

```
no-auto-iccp-provisioning;
```

Hierarchy Level

```
[edit chassis satellite-management redundancy-groups redundancy-group-name peer-chassis-id peer-chassis-id-number]
```

Description

Disable automatic ICCP provisioning of an interchassis link (ICL) for the redundancy group used to configure dual aggregation devices in a Junos Fusion topology.

Automatic ICCP provisioning of an interchassis link (ICL) simplifies configuration of a Junos Fusion with dual aggregation devices by automatically provisioning the ICCP configuration within the Junos Fusion, instead of requiring the user to manually configure all ICCP parameters. Automatic ICCP Provisioning of an interchassis link (ICL) is enabled by default for a Junos Fusion using a dual aggregation device topology; this statement disables automatic ICCP provisioning.

If this statement is entered, the user has to manually configure ICCP in the redundancy group used to configure dual aggregation devices in a Junos Fusion topology.

This statement is optional. You can manually configure any available ICCP configuration parameters in the redundancy group when automatic ICCP provisioning is enabled. In cases where a user configures an ICCP parameter when automatic ICCP provisioning is enabled, the user-configured ICCP configuration is used over the automatically-provisioned ICCP configuration parameter.

User configuration of ICCP in a Junos Fusion is not required. This statement is recommended for use by expert users only.

Default

Automatic ICCP Provisioning is enabled by default in redundancy groups used to configure dual aggregation devices in a Junos Fusion topology.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

no-auto-mdix

IN THIS SECTION

- Syntax | 11409
- Hierarchy Level | 11409
- Description | 11409
- Default | 11410
- Options | 11410
- Required Privilege Level | 11410
- Release Information | 11410

Syntax

```
no-auto-mdix;
```

Hierarchy Level

```
[edit interface interface-name gigether-options]
```

Description

Disable the Auto MDI/MDIX feature.

MX Series routers with Gigabit Ethernet interfaces automatically detect MDI and MDIX port connections. Use this statement to override the default setting. Remove this statement to return to the default setting.

The `show interfaces interface-name` command displays incorrect auto MDI-X status when you disable auto MDI-X.

Default

Auto MDI/MDIX is enabled by default.

Options

There are no options for this statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

Ethernet Interfaces Overview

gigether-options | **7685**

no-auto-virtual-gateway-esi

IN THIS SECTION

- Syntax | **11411**
- Hierarchy Level | **11411**
- Description | **11411**
- Required Privilege Level | **11411**

Release Information | 11411

Syntax

```
no-auto-virtual-gateway-esi;
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit],
[edit dynamic-profiles name logical-systems name interfaces name unit],
[edit interfaces name unit]
```

Description

Disable auto ESI generation for virtual gateway address. Automatic ESI generation is enabled by default on devices in EVPN-VXLAN networks with EVPN multihoming for virtual gateway redundancy (see Understanding the Redundant Default Gateway).

Starting in Junos OS Release 22.1R1, MX960, MX2020, and MX10008 routers also enable automatic ESI generation by default for EVPN Layer 3 gateway IRB interface ESIs. However, the `no-auto-virtual-gateway-esi` statement is not supported with EVPN-MPLS networks. As a result, you will always see auto-generated ESIs for IRB interfaces in that case.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

no-autostart

IN THIS SECTION

- Syntax | 11412
- Hierarchy Level | 11412
- Description | 11412
- Required Privilege Level | 11412
- Release Information | 11412

Syntax

```
no-autostart;
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Disable auto-start of VNF on the VNF configuration commit.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D40.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

no-autoupdate

IN THIS SECTION

- Syntax | 11413
- Hierarchy Level | 11413
- Description | 11413
- Required Privilege Level | 11414
- Release Information | 11414

Syntax

```
no-autoupdate;
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update]
[edit security utm feature-profile anti-virus sophos-engine pattern-update]
[edit security utm default-configuration anti-virus avira-engine pattern-update]
```

Description

Turn off automatic data file (pattern file) update for the Kaspersky Lab, Juniper Express, or Sophos engines.

> **NOTE**: The data files used with Sophos are not typical virus pattern files; they are small files that help guide virus scanning logic. The full virus pattern database is stored on an external Sophos server called the Sophos Extensible List (SXL) server.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5 for Juniper Express engine and Kaspersky Lab engine. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for Avira engine added in Junos OS Release 18.4R1.

no-bfd-triggered-local-repair

IN THIS SECTION

- Syntax | 11415
- Hierarchy Level | 11415
- Description | 11415
- Default | 11415
- Required Privilege Level | 11415
- Release Information | 11415

Syntax

```
no-bfd-triggered-local-repair;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit routing-options]
```

Description

Disable Bidirectional Forwarding Detection (BFD) sessions to trigger fast reroute (FRR) using MPLS-FRR and loop-free alternates (LFAs). When this statement is configured, no BFD-triggered local repair is supported. However, logical interface down-based local repair is in force.

When using this statement to disable local repair, you also must restart routing to ensure proper behavior. To restart routing, include the **graceful-restart** command for the interior gateway protocol (IGP) used in your configuration. For example, if your IGP is OSPF, include the **graceful-restart** statement at the [edit protocols ospf] hierarchy level.

Default

BFD-triggered local repair is the default behavior. The loss of a neighbor results in BFD local repair for all next hops that derive themselves from the base next hop with which the BFD session is established.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

BFD-Triggered Local Repair for Rapid Convergence

graceful-restart (Enabling Globally) | 7760

no-bidirectional-mode

IN THIS SECTION

- Syntax | 11416
- Hierarchy Level | 11416
- Description | 11416
- Default | 11417
- Required Privilege Level | 11417
- Release Information | 11417

Syntax

```
no-birectional-mode;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
graceful-restart],
[edit protocols pim graceful-restart],
[edit routing-instances routing-instance-name protocols pim graceful-restart]
```

Description

Disable forwarding for bidirectional PIM routes during graceful restart recovery, both in cases of a routing protocol process (rpd) restart and graceful Routing Engine switchover.

Bidirectional PIM accepts packets for a bidirectional route on multiple interfaces. This means that some topologies might develop multicast routing loops if all PIM neighbors are not synchronized with regard to the identity of the designated forwarder (DF) on each link. If one router is forwarding without actively

participating in DF elections, particularly after unicast routing changes, multicast routing loops might occur.

If graceful restart for PIM is enabled and the forwarding of packets on bidirectional routes is disallowed (by including the **no-bidirectional-mode** statement in the configuration), PIM behaves conservatively to avoid multicast routing loops during the recovery period. When the routing protocol process (rpd) restarts, all bidirectional routes are deleted. After graceful restart has completed, the routes are re-added, based on the converged unicast and bidirectional PIM state. While graceful restart is active, bidirectional multicast flows drop packets.

Default

If graceful restart for PIM is enabled and the bidirectional PIM is enabled, the default graceful restart behavior is to continue forwarding packets on bidirectional routes. If the gracefully restarting router was serving as a DF for some interfaces to rendezvous points, the restarting router sends a DF Winner message with a metric of 0 on each of these RP interfaces. This ensures that a neighbor router does not become the DF due to unicast topology changes that might occur during the graceful restart period. Sending a DF Winner message with a metric of 0 prevents another PIM neighbor from assuming the DF role until after graceful restart completes. When graceful restart completes, the gracefully restarted router sends another DF Winner message with the actual converged unicast metric.

NOTE: Graceful Routing Engine switchover operates independently of the graceful restart behavior. If graceful Routing Engine switchover is configured without graceful restart, all PIM routes for all modes are deleted when the rpd process restarts. If graceful Routing Engine switchover is configured with graceful restart, the behavior is the same as described here, except that the recovery happens on the Routing Engine that assumes primary role.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring PIM Sparse Mode Graceful Restart

Understanding Bidirectional PIM

Example: Configuring Bidirectional PIM

no-bind-on-request (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 11418
- Hierarchy Level | 11418
- Description | 11419
- Required Privilege Level | 11419
- Release Information | 11419

Syntax

```
no-bind-on-request;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Explicitly disable automatic binding of received DHCP request messages that have no entry in the database (*stray* requests). Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

> **NOTE**: Beginning with Junos OS Release 10.4, automatic binding of stray requests is enabled by default. In Junos OS Release 10.3 and earlier releases, automatic binding of stray requests is disabled by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Overriding the Default DHCP Relay Configuration Settings

Disabling Automatic Binding of Stray DHCP Requests

no-broadcast

IN THIS SECTION

- Syntax | 11420
- Hierarchy Level | 11420
- Description | 11420
- Default | 11421
- Required Privilege Level | 11421
- Release Information | 11421

Syntax

```
no-broadcast;
```

Hierarchy Level

- For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers):

```
[edit forwarding-options storm-control-profiles profile-name all]
```

Description

Disable storm control for broadcast traffic for the specified interface or for all interfaces.

Default

- On EX2200, EX3200, EX3300, and EX4200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams. You can selectively disable storm control on broadcast, multicast, or unknown-unicast traffic.
- On EX4300 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. You can selectively disable storm control on any type of traffic.
- On EX4500 and EX8200 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. On EX8200 switches, you can selectively disable storm control on registered multicast traffic, on unregistered multicast traffic, or on both types of multicast traffic.
- On EX6200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams. You can selectively disable storm control for each type of traffic individually.
- On EX9200 switches—Storm control is not enabled by default.
- On MX Series routers—Storm control is not enabled by default.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Hierarchy level `[edit forwarding-options]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

> *Configuring Autorecovery for Port Security Events*
>
> Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches

Example: Using Storm Control to Prevent Network (MX Routers)

Enabling and Disabling Storm Control (non-ELS)

Enabling and Disabling Storm Control (ELS)

no-bundle-flap

IN THIS SECTION

- Syntax | 11422
- Hierarchy Level | 11422
- Description | 11422
- Required Privilege Level | 11423
- Release Information | 11423

Syntax

```
no-bundle-flap;
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name load-balancing-options]
```

Description

When you add a new member to an existing AMS bundle, all the existing members and the newly added member of the AMS bundle go for reboot and disrupts the traffic. To overcome this problem for IPsec services, configure the `no-bundle-flap` statement before adding a new member to the AMS bundle. When you configure `no-bundle-flap` command and add a new member to the AMS bundle, the existing members of AMS bundle will not reboot, only the newly added member reboot avoiding the traffic disruption.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1

no-certificate-chain-in-ike

IN THIS SECTION

- Syntax | 11423
- Hierarchy Level | 11423
- Description | 11423
- Required Privilege Level | 11424
- Release Information | 11424

Syntax

```
no-certificate-chain-in-ike;
```

Hierarchy Level

```
[edit services service-set name ipsec-vpn-options]
```

Description

To avoid IKE fragmentation, send only the end-entity certificate for certificate-based IKE authentication instead of the full certificate chain.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Configuring IPsec Service Sets

no-client-reflect

IN THIS SECTION

- Syntax | 11424
- Hierarchy Level | 11424
- Description | 11425
- Required Privilege Level | 11425
- Release Information | 11425

Syntax

```
no-client-reflect;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],

[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Disable intracluster route redistribution by the system acting as the route reflector. Include this statement when the client cluster is fully meshed to prevent the sending of redundant route advertisements. Route reflection provides a way to decrease BGP control traffic and minimizing the number of update messages sent within the AS.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring BGP Route Reflectors

cluster | 4367

no-control-word

IN THIS SECTION

- Syntax | 11426
- Hierarchy Level | 11426
- Description | 11426
- Required Privilege Level | 11427
- Release Information | 11427

Syntax

```
no-control-word;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
evpn],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name neighbor address static],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor address static],
[edit routing-instances routing-instance-name protocols evpn],
[edit routing-instances routing-instance-name protocols vpls],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name neighbor
address static],
[edit routing-instances routing-instance-name protocols vpls neighbor address static]
```

Description

Set `no-control-word` to request that other routers *not* insert a control word between the label stack and the MPLS payload. This is the default setting for BGP VPLS.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced for static pseudowires in mesh-group and in VPLS instance in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Control Word for BGP VPLS Overview

Configuring a Control Word for BGP VPLS

control-word | 4644

no-control-word (Protocols Layer 2 VPN)

IN THIS SECTION

- Syntax | 11427
- Hierarchy Level | 11428
- Description | 11428
- Default | 11428
- Required Privilege Level | 11428
- Release Information | 11428

Syntax

```
no-control-word;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn],
[edit routing-instances routing-instance-name protocols l2vpn]
```

Description

Disable the control word. This might be necessary on networks with equipment that does not support the control word.

> **NOTE**: The following configuration statements are ignored for time-division multiplexing pseudowires at the `[edit protocols l2vpn]` hierarchy level:
>
> - control-word
> - no-control-word

Default

The control word is enabled by default. Use the `no-control-word` statement to disable the control word.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Disabling the Control Word for Layer 2 VPNs

control-word (Protocols Layer 2 VPN) | **4648**

no-cspf

IN THIS SECTION

- Syntax | 11429
- Hierarchy Level | 11429
- Description | 11429
- Default | 11430
- Required Privilege Level | 11430
- Release Information | 11430

Syntax

```
no-cspf;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mplslabel-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

Disable constrained-path LSP computation.

An explicit-path LSP is completely configured through operator action. Once configured, it is initiated only along the explicitly specified path.

A constrained-path LSP relies on an ingress router to compute the complete path. The ingress router takes into account the following information during the computation:

- Interior gateway protocol (IGP) topology database
- Link utilization information from extensions in the IGP link-state database
- Administrative group information from extensions in the IGP link-state database
- LSP requirements, including bandwidth, hop count, and administrative group

Constrained-path LSPs can generally avoid link failures and congested links. They also permit recomputation (therefore, a new path) during topology changes or unsuccessful setup.

Default

Constrained-path LSP computation enabled.

> **NOTE**: On PTX and MX routers the `no-local-reversion` option is enabled by default starting with Junos OS release 16.1. If you use the `no-cspf` option on these platforms with either link protection or node-link protection enabled, you must also configure the `local-reversion` option to enure proper behavior,

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Disabling Constrained-Path LSP Computation

Configuring Explicit-Path LSPs

no-cspf (Protocols RSVP)

IN THIS SECTION

- Syntax | 11431
- Hierarchy Level | 11431
- Description | 11431
- Default | 11431
- Required Privilege Level | 11432
- Release Information | 11432

Syntax

```
no-cspf;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection bypass bypass-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Description

Disable CSPF computation on all bypass LSPs or on a specific bypass LSP. You need to disable CSPF for link protection to function properly on interarea paths.

Default

CSPF is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Disabling CSPF for Bypass LSPs

no-decrement-ttl

IN THIS SECTION

- Syntax | 11432
- Hierarchy Level | 11433
- Description | 11433
- Default | 11433
- Required Privilege Level | 11433
- Release Information | 11433

Syntax

```
no-decrement-ttl;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

Disable normal time-to-live (TTL) decrementing, which decrements the TTL field in the IP header by 1. This statement decrements the IP TTL by 1 before encapsulating the IP packet within an MPLS packet. When the penultimate router pops off the top label, it does not use the standard write-back procedure of writing the MPLS TTL into the IP TTL field. Therefore, the IP packet is decremented by 1. The ultimate router then decrements the packet by one more for a total cloud appearance of 2, thus hiding the network topology.

Default

Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1 on PTX10008.

RELATED DOCUMENTATION

Disabling Normal TTL Decrementing

no-propagate-ttl | **11553**

no-default-gateway-ext-comm

IN THIS SECTION

- Syntax | 11434
- Hierarchy Level | 11434
- Description | 11434
- Required Privilege Level | 11434
- Release Information | 11434

Syntax

```
no-default-gateway-ext-comm;
```

Hierarchy Level

```
[edit routing-instances name protocols evpn]
```

Description

Configure a routing instance to suppress the advertisement of the extended community on the default gateway. An extended community is an 8-octet community used by Ethernet VPNs (EVPNs).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Implementing EVPN-VXLAN for Data Centers

EVPN Multihoming Overview

no-delegate-processing

IN THIS SECTION

- Syntax | 11435
- Hierarchy Level | 11435
- Description | 11435
- Default | 11436
- Required Privilege Level | 11436
- Release Information | 11436

Syntax

```
no-delegate-processing;
```

Hierarchy Level

```
[edit routing-options ppm]
```

Description

Disable distributed periodic packet management (PPM) processing and run all PPM processing on the Routing Engine.

PPM processing cannot be completely disabled on EX Series switches. You can only configure whether PPM processing is distributed between the access ports (EX3200 and EX4200 switches) or the line cards (EX8200 switches) and the Routing Engine or is handled just on the Routing Engine.

BEST PRACTICE: Generally, you should only disable distributed PPM if Juniper Networks Customer Service advised you to do so. You should only disable distributed PPM if you have a compelling reason to disable it.

Default

Distributed PPM processing is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

no-dhcp-snooping

IN THIS SECTION

- Syntax | 11437
- Hierarchy Level (EX Series, QFX Series) | 11437
- Hierarchy Level (MX Series) | 11437
- Description | 11437
- Default | 11437
- Required Privilege Level | 11438
- Release Information | 11438

Syntax

```
no-dhcp-snooping;
```

Hierarchy Level (EX Series, QFX Series)

```
[edit vlans vlan-name forwarding-options dhcp-security]
```

Hierarchy Level (MX Series)

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]
```

Description

Disable DHCP snooping for the specified VLAN or bridge domain.

> **NOTE**: Explicitly disabling DHCP snooping also disables any other port security features that you have enabled under `[edit vlans vlan-name forwarding-options dhcp-security]`, including dynamic ARP inspection (DAI) and IP source guard for the specified VLAN or bridge domain.
>
> There is no configuration statement that explicitly enables DHCP snooping.

Default

DHCP snooping is not enabled.

> **NOTE**: Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style and MX Series routers do not have a configuration statement that explicitly enables DHCP snooping.
>
> DHCP snooping is enabled automatically by Junos OS if any of the following is configured at the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level for EX Series and QFX Series switches or at the `[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]` for MX Series routers:

- DAI
- IP source guard
- Static IP
- DHCP option 82

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

Hierarchy level `[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]` introduced in Junos OS Release 14.1 for the MX Series.

RELATED DOCUMENTATION

Understanding DHCP Snooping (non-ELS)

no-dhcp-trusted

IN THIS SECTION

- Syntax | 11439
- Hierarchy Level | 11439
- Description | 11439
- Default | 11439
- Required Privilege Level | 11439
- Release Information | 11439

Syntax

```
(dhcp-trusted | no-dhcp-trusted);
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port  interface (Access Port Security) (all | interface-name)]
```

Description

Port security features, such as DHCP snooping and dynamic ARP inspection inspect packets only on untrusted interfaces.

Allow or deny DHCP responses from the specified interfaces (ports) or all interfaces.

- **dhcp-trusted**—Allow DHCP responses.
- **no-dhcp-trusted**—Deny DHCP responses.

Default

Trusted for trunk ports, untrusted for access ports.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Understanding How to Protect Access Ports from Common Attacks

Understanding and Using Trusted DHCP Servers

no-dhcpv6-options

IN THIS SECTION

- Syntax | 11440
- Hierarchy Level | 11440
- Description | 11440
- Required Privilege Level | 11440
- Release Information | 11440

Syntax

```
no-dhcpv6-options;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name overrides]
```

Description

Configure a specific group of one or more access interfaces within the VLAN not to add any DHCPv6 options, even if the VLAN is configured to perform DHCPv6 snooping. DHCPv6 options include option 16, option 18, and option 37.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

dhcpv6-options | 5673

Understanding DHCP Snooping (ELS)

Understanding DHCP Option 82

no-dhcpv6-snooping

IN THIS SECTION

- Syntax | 11441
- Hierarchy Level | 11441
- Description | 11441
- Default | 11441
- Required Privilege Level | 11442
- Release Information | 11442

Syntax

```
no-dhcpv6-snooping;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security]
```

Description

Disable DHCPv6 snooping for the specified VLAN.

Default

DHCPv6 snooping is not enabled by default.

There is no configuration statement that explicitly enables DHCPv6 snooping. DHCPv6 snooping is enabled automatically by Junos OS if any port security feature, such as IPv6 neighbor discovery inspection or IPv6 source guard, is configured at the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D20.

RELATED DOCUMENTATION

Understanding DHCP Snooping (non-ELS)

no-dns-propagation

IN THIS SECTION

- Syntax | 11442
- Hierarchy Level | 11443
- Description | 11443
- Required Privilege Level | 11443
- Release Information | 11443

Syntax

```
no-dns-propagation;
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number family inet | inet6 dhcp-client]
```

Description

Disable the propagation of DNS information to the kernel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D35.

RELATED DOCUMENTATION

Understanding Interfaces

no-domain-vpn-tag

IN THIS SECTION

- Syntax | 11444
- Hierarchy Level | 11444
- Description | 11444
- Options | 11444
- Required Privilege Level | 11444
- Release Information | 11444

Syntax

```
no-domain-vpn-tag;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)]
```

Description

Disable the virtual private network (VPN) tag for OSPFv2 and OSPFv3 external routes generated by the provider edge (PE) router when the VPN tag is no longer needed.

Options

None.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.3.

RELATED DOCUMENTATION

Configuring Routing Between PE and CE Routers in Layer 3 VPNs

no-dr-flood (PIM Snooping)

IN THIS SECTION

- Syntax | 11445
- Hierarchy Level | 11445
- Description | 11445
- Required Privilege Level | 11445
- Release Information | 11446

Syntax

```
no-dr-flood;
```

Hierarchy Level

```
[edit routing-instances <instance-name> protocols pim-snooping traceoptions],
[edit logical-systems <logical-system-name> routing-instances <instance-name> protocols pim-
snooping traceoptions],
[edit routing-instances <instance-name> protocols pim-snooping vlan <vlan-id>],
[edit logical-systems <logical-system-name> routing-instances <instance-name> protocols pim-
snooping  vlan<vlan-id>]
```

Description

Disable default flooding of multicast data on the PIM designated router port.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

no-dynamic-vlan

IN THIS SECTION

- Syntax | 11446
- Hierarchy Level | 11446
- Description | 11446
- Default | 11447
- Required Privilege Level | 11447
- Release Information | 11447

Syntax

```
no-dynamic-vlan;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mvrp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
mvrp],
[edit protocols mvrp],
[edit routing-instances routing-instance-name protocols mvrp]
```

Description

Disable the dynamic creation of VLANs using Multiple VLAN Registration Protocol (MVRP) for interfaces participating in MVRP.

Dynamic VLAN configuration can be enabled on an interface independent of MVRP. The MVRP dynamic VLAN configuration setting does not override the interface configuration dynamic VLAN configuration setting. If dynamic VLAN creation is disabled on the interface in the interface configuration, no dynamic VLANs are created on the interface, including dynamic VLANs created using MVRP.

This option can only be applied globally; it cannot be applied per interface.

Default

If MVRP is enabled, the dynamic creation of VLANs as a result of MVRP protocol exchange messages is enabled.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers

Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration

Controlling the Management State of a VLAN in MVRP Configurations

Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

no-encryption (MACsec)

IN THIS SECTION

- Syntax | 11448
- Hierarchy Level | 11448
- Description | 11448

- Default | 11449
- Required Privilege Level | 11449
- Release Information | 11449

Syntax

```
no-encryption;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Disables MACsec encryption for a connectivity association that is configured to enable MACsec using static connectivity association key (CAK) or dynamic security mode.

MACsec encryption is disabled when MACsec is configured, by default.

You can enable MACsec without enabling encryption. If a connectivity association that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the packet, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic does not represent a security threat.

This command is used to disable encryption when MACsec is configured using static CAK or dynamic security mode only. When MACsec is configuring using static secure association key (SAK) security mode, the encryption setting is managed in the secure channel using the "encryption" on page 6334 configuration statement.

Enable MACsec encryption within a secure channel.

Traffic traversing a MAC-enabled point-to-point Ethernet link traverses the link at the same speed regardless of whether encryption is enabled or disabled. You cannot increase the speed of traffic traversing a MACsec-enabled Ethernet link by disabling encryption.

When MACsec is configuring using static connectivity association key (CAK) security mode, the encryption setting is configured outside of the secure channel using the `no-encryption` configuration statement.

Default

MACsec encryption is enabled if MACsec is enabled using static CAK or dynamic security mode.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced for SRX Series Firewalls in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

no-encryption (MACsec for MX Series)

IN THIS SECTION

- Syntax | **11450**
- Hierarchy Level | **11450**
- Description | **11450**
- Default | **11450**

Required Privilege Level | 11450
Release Information | 11451

Syntax

```
no-encryption;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Disables MACsec encryption for a connectivity association that is configured to enable MACsec using static connectivity association key (CAK) or dynamic security mode.

You can enable MACsec without enabling encryption. If a connectivity association that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the packet, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic does not represent a security threat.

This command is used to disable encryption when MACsec is configured using static CAK or dynamic security mode only. When MACsec is configuring using static secure association key (SAK) security mode, the encryption setting is managed in the secure channel using the `encryption` configuration statement.

Default

MACsec encryption is enabled if MACsec is enabled using static CAK or dynamic security mode.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

no-enhanced-frr-bypass (Protocols RSVP)

IN THIS SECTION

- Syntax | 11451
- Hierarchy Level | 11451
- Description | 11452
- Default | 11452
- Required Privilege Level | 11452
- Release Information | 11452

Syntax

```
no-enhanced-frr-bypass;
```

Hierarchy Level

```
[edit protocols rsvp ]
```

Description

Enable `no-enhanced-frr-bypass` to turn off all Fast reroute (FRR) facility protection enhancements, which includes improved LSP scaling and enhanced RSVP message handling, and reduce the default refresh time to 30 seconds.

FRR is enabled by default for MX Series and PTX Series routers starting in Junos OS Release 15.2R1.

Default

This feature, `no-enhanced-frr-bypass`, is disabled by default. That is, enhanced FRR is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.2R1.

RELATED DOCUMENTATION

RSVP Refresh Reduction

no-examine-dhcpv6

IN THIS SECTION

- Syntax | 11453
- Hierarchy Level | 11453
- Description | 11453
- Default | 11453
- Required Privilege Level | 11453
- Release Information | 11453

Syntax

```
no-examine-dhcpv6 {
    forwarding-class class-name;
}
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name)]
```

Description

Disable DHCPv6 snooping on all VLANs or on the specified VLAN.

The remaining statement is explained separately.

Default

Disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

examine-dhcpv6 | 6563

Example: Configuring Port Security (non-ELS)

Example: Protecting Against Address Spoofing and Layer 2 DoS Attacks

Example: Protecting Against ARP Spoofing Attacks

Example: Prioritizing Snooped and Inspected Packet

Enabling DHCP Snooping (non-ELS)

Enabling DHCP Snooping (J-Web Procedure)

no-external-export

IN THIS SECTION

- Syntax | 11454
- Hierarchy Level | 11454
- Description | 11455
- Options | 11455
- Required Privilege Level | 11455
- Release Information | 11455

Syntax

```
no-external-export;
```

Hierarchy Level

```
[edit logical-systems name protocols isis no-external-export protocol],
[edit logical-systems name routing-instances name protocols isis no-external-export protocol],
[edit logical-systems name tenants name routing-instances name protocols isis no-external-export
protocol],
[edit protocols isis no-external-export protocol],
[edit routing-instances name protocols isis no-external-export protocol],
[edit tenants name routing-instances name protocols isis no-external-export protocol]
```

Description

Block the export policy from redistributing the routes from other protocols into IS-IS.

Disabling external export has the following results

- IS-IS discards the routes from the configured protocol
- IS-IS withdraws any prefixes from the configured protocol that were already advertised.

Options

bgp	BGP
isis	ISIS
ospf	OSPF
static	Static

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.1

RELATED DOCUMENTATION

export (ISIS) | 6658

no-fabric-wwn-verify

IN THIS SECTION

- Syntax | 11456
- Hierarchy Level | 11456
- Description | 11456
- Default | 11456
- Required Privilege Level | 11457
- Release Information | 11457

Syntax

```
no-fabric-wwn-verfiy;
```

Hierarchy Level

```
[edit fc-fabrics fabric-name proxy]
```

Description

Disable the fabric worldwide name (WWN) verification check in the fabric login accept message (FLOGI-ACC) for implicit FLOGIs. If you enable this option, when a QFX Series NP_Port performs a FLOGI to the FC fabric, the QFX Series does not verify the fabric WWN in the FLOGI-ACC against the current fabric WWN.

NOTE: Disabling or enabling the fabric WWN verification check logs out all FCoE sessions.

Default

Disabled. By default, all implicit FLOGIs from the QFX Series NP_Ports to the FC fabric are verified against the current fabric WWN.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

show fibre-channel proxy fabric-state | 19987

Understanding FCoE-FC Gateway Functions

no-fcoe-lag

IN THIS SECTION

- Syntax | 11457
- Hierarchy Level | 11457
- Description | 11458
- Required Privilege Level | 11458
- Release Information | 11458

Syntax

```
no-fcoe-lag;
```

Hierarchy Level

```
[edit interfaces lag-interface-name aggregated-ether-options]
```

Description

Convert an FCoE LAG into a standard LAG. When you convert an FCoE LAG into a standard LAG, the standard LAG no longer works reliably for FCoE traffic. This is because FCoE traffic must use the same physical link within a LAG interface for communication between the FCoE device and the Fibre Channel SAN across a QFabric system Node device. A standard LAG uses a hashing algorithm to determine the link to use for each transmission, so there is no way to guarantee that a response will use the same link on which a device receives a request. An FCoE LAG guarantees that the same physical LAG link is used for communication between an FCoE device and the QFabric system Node device.

If you convert an FCoE LAG into a standard LAG, do not use the standard LAG for FCoE traffic.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X52-D10.

RELATED DOCUMENTATION

- fcoe-lag | **6963**
- no-fip-snooping-scaling | **11461**
- *Configuring an FCoE LAG*
- *Configuring VLANs for FCoE Traffic on an FCoE Transit Switch*
- *Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch*
- *Understanding FCoE LAGs*

no-fcoe-trusted

IN THIS SECTION

- Syntax | **11459**

- Hierarchy Level | 11459
- Description | 11459
- Required Privilege Level | 11460
- Release Information | 11460

Syntax

```
no-fcoe-trusted;
```

Hierarchy Level

Original CLI

```
[edit ethernet-switching-options secure-access-port interface interface-name]
```

ELS CLI for Platforms that Support FCoE

```
[edit vlans vlan-name forwarding-options fip-security interface interface-name]
```

NOTE: The `no-fcoe-trusted` configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.

Description

Configure the specified 10-Gigabit Ethernet interface not to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is directly connected to an FCoE device, the interface should not be configured as an FCoE trusted interface. If an interface that you want to connect to an FCoE device has been configured as an FCoE trusted interface, use the `no-fcoe-trusted` statement to convert the interface to an untrusted interface. Untrusted interfaces can perform FIP snooping to provide access security for FCoE traffic.

However, if an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

show fip snooping | 20008

Example: Configuring an FCoE Transit Switch

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

no-filter-check

IN THIS SECTION

- Syntax | 11460
- Hierarchy Level | 11461
- Description | 11461
- Required Privilege Level | 11461
- Release Information | 11461

Syntax

```
no-filter-check;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family (inet | inet6) output]
[edit forwarding-options port-mirroring family (inet | inet6 | ccc | vpls) output]
```

Description

Disable filter checking on the port-mirroring interface.

This statement is required when you send port-mirrored traffic to a Tunnel PIC that has a filter applied to it.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

no-fip-snooping-scaling

IN THIS SECTION

- Syntax | 11462
- Hierarchy Level (FCoE-FC gateway) | 11462
- Hierarchy Level (FCoE Transit Switch) | 11462
- Description | 11462
- Default | 11462

Required Privilege Level | 11463
Release Information | 11463

Syntax

```
no-fip-snooping-scaling
```

Hierarchy Level (FCoE-FC gateway)

```
[edit fc-options]
```

Hierarchy Level (FCoE Transit Switch)

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name) examine-fip]
```

Description

Disable FIP snooping scaling on all FCoE VLANs on an FCoE-FC gateway, or disable FIP snooping scaling on the specified FCoE VLAN on an FCoE transit switch.

Disabling FIP snooping scaling reduces the maximum number of FIP snooping sessions from 2,500 sessions (the maximum with FIP snooping scaling enabled) to 376 sessions. FIP snooping scaling is enabled by default.

Use this statement to disable FIP snooping scaling if you want to configure an FCoE LAG on an FCoE-FC gateway that contains one or more untrusted FC fabrics. Untrusted FC fabrics do not support FIP snooping scaling.

On an FCoE transit switch, you can use this statement to disable FIP snooping scaling on a specified FCoE VLAN.

Default

FIP snooping scaling is enabled by default.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X52-D10.

RELATED DOCUMENTATION

fcoe-lag | 6963

Configuring an FCoE LAG

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

Configuring VLANs for FCoE Traffic on an FCoE Transit Switch

Understanding FCoE LAGs

no-flow-logging (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 11463
- Hierarchy Level | 11464
- Description | 11464
- Required Privilege Level | 11464
- Release Information | 11464

Syntax

```
no-flow-logging;
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Disable automatic logging of flow detection culprit flow events (flow reports) for the packet type.

NOTE: You can disable logging of suspicious flow events (violation reports) with the `disable-logging` statement at the `[edit system ddos-protection global` hierarchy level.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Disabling Automatic Logging of Culprit Flow Events for a Packet Type

Setting Up and Using Flow Detection

no-forwarding

IN THIS SECTION

Syntax | 11465

- Hierarchy Level | 11465
- Description | 11465
- Default | 11465
- Required Privilege Level | 11465
- Release Information | 11466

Syntax

```
no-forwarding;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Do not add ingress routes to the inet.0 routing table even if `traffic-engineering bgp-igp` (configured at the `[edit protocols mpls]` hierarchy level) is enabled.

Default

The `no-forwarding` statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when `traffic-engineering bgp-igp` is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Preventing Addition of Ingress Routes to the inet.0 Routing Table

Configuring Virtual-Router Routing Instances in VPNs

no-fragmentation

IN THIS SECTION

- Syntax | 11466
- Hierarchy Level | 11466
- Description | 11467
- Default | 11467
- Required Privilege Level | 11467
- Release Information | 11467

Syntax

```
no-fragmentation;
```

Hierarchy Level

```
[edit class-of-service fragmentation-maps forwarding-class class-name]
```

Description

For link services IQ (`lsq`) interfaces only, set traffic on a particular forwarding class to be interleaved, rather than fragmented. This statement specifies that no extra fragmentation header is prepended to the packets received on this queue and that static-link load balancing is used to ensure in-order packet delivery.

Static-link load balancing is done based on packet payload. For IP version 4 (IPv4) and IP version 6 (IPv6) traffic, the link is chosen based on a hash computed from the source address, destination address, and protocol. If the IP payload is Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) traffic, the hash also includes source port and destination port. For MPLS traffic, the hash includes all MPLS labels and fields in the payload, whether the MPLS payload is IPv4 or IPv6.

Default

If you do not include this statement, the traffic in forwarding class *`class-name`* is fragmented.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces

no-gratuitous-arp-request

IN THIS SECTION

- Syntax | **11468**
- Hierarchy Level | **11468**

- Description | 11468
- Default | 11468
- Required Privilege Level | 11468
- Release Information | 11468

Syntax

```
no-gratuitous-arp-request;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on Layer 2 Ethernet switching interfaces, and integrated routing and bridging (IRB) interfaces or routed VLAN interfaces (RVIs). (On EX Series switches that use Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style, the feature is known as an IRB interface. On EX Series switches that use Junos OS that does not support ELS, the feature is known as an RVI.)

Default

Gratuitous ARP responses are enabled on all Ethernet switching interfaces, and IRB interfaces or RVIs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

- Example: Configuring Proxy ARP on an EX Series Switch
- Configuring Proxy ARP on Switches
- Configuring Proxy ARP on Devices with ELS Support

no-hidden-commands

IN THIS SECTION

- Syntax | 11469
- Hierarchy Level | 11469
- Description | 11470
- Default | 11470
- Options | 11470
- Required Privilege Level | 11470
- Release Information | 11470

Syntax

```
no-hidden-commands;
```

Hierarchy Level

```
[edit system]
```

Description

Hidden commands are software commands that are not published but could be run on a router. Hidden commands serve a specific purpose, but for most part are not expected to be used, and as such are not actively supported. The `no-hidden-commands` statement allows you to block all hidden commands to all users except the root users.

Default

Hidden commands are enabled by default.

Options

This command has no options.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

no-hierarchical-ecmp

IN THIS SECTION

- Syntax | 11471
- Hierarchy Level | 11471
- Description | 11471
- Required Privilege Level | 11471
- Release Information | 11471

Syntax

```
no-hierarchical-ecmp;
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Disable hierarchical equal-cost multipath (ECMP) groups at system start time. Hierarchical ECMP is enabled by default and provides for two-level route resolution. Disabling hierarchical ECMP effectively increases the number of ECMP groups available for route resolution because hierarchical ECMP allocates separate groups for overlay and underlay routes. Increasing the number of ECMP groups available is potentially useful in a virtual extensible LAN (VXLAN) environment or when MPLS LDP forwarding equivalence classes (FECs) are configured.

CAUTION: Disabling hierarchical ECMP causes the Packet Forwarding Engine to restart.

To reenable hierarchical ECMP, use the following command: `delete forwarding-options no-hierarchical-ecmp`. Reenabling hierarhical ECMP also causes the Packet Forwarding Engine to restart.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D210.

RELATED DOCUMENTATION

- Overview of Hierarchical ECMP Groups on QFX5200 Switches

no-icmp-packet-too-big

IN THIS SECTION

- Syntax | 11472
- Hierarchy Level | 11472
- Description | 11472
- Required Privilege Level | 11472
- Release Information | 11472

Syntax

```
no-icmp-packet-too-big;
```

Hierarchy Level

```
[edit security ipsec vpn hub-to-spoke-vpn no-icmp-packet-too-big]
```

Description

For IPv6, the `no-icmp-packet-too-big` option disables sending the ICMP `Packet Too Big` message.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.3R1

no-incremental-update

IN THIS SECTION

- Syntax | 11473
- Hierarchy Level | 11473
- Description | 11473
- Required Privilege Level | 11473
- Release Information | 11473

Syntax

```
no-incremental-update
```

Hierarchy Level

```
[edit firewall family filter]
```

Description

Prevent the filter from undergoing incremental update. The filter undergoes MBB.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement released in Junos OS Evolved Release 24.1

no-inner-payload

IN THIS SECTION

- Syntax | 11474
- Hierarchy Level | 11474
- Description | 11474
- Required Privilege Level | 11474
- Release Information | 11474

Syntax

```
no-inner-payload;
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key vxlan]
```

Description

Enable load-balancing on VXLAN traffic based on outer IP/UDP headers.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

flex-hashing | 7227

conditional-match | 4504

no-install

IN THIS SECTION

- Syntax | 11475
- Hierarchy Level | 11475
- Description | 11475
- Required Privilege Level | 11475
- Release Information | 11476

Syntax

```
no-install;
```

Hierarchy Level

```
[edit protocols bgp family (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) (any | flow | labeled-
unicast | multicast | unicast)]
```

Description

Prohibit installing received routes in the forwarding table. This statement can be set per family.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding BGP Route Reflectors

Understanding BGP Flow Routes for Traffic Filtering

no-install-to-address

IN THIS SECTION

- Syntax | 11476
- Hierarchy Level | 11476
- Description | 11477
- Default | 11477
- Required Privilege Level | 11477
- Release Information | 11477

Syntax

```
no-install-to-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Description

Prevent the egress router address configured using the `to` statement from being installed into the inet.3 and inet.0 routing tables.

Default

The egress router address for an LSP is installed into the inet.3 and inet.0 routing tables.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Preventing the Addition of Egress Router Addresses to Routing Tables

to | **29419**

no-intelligent-prescreening

IN THIS SECTION

- Syntax | **11478**
- Hierarchy Level | **11478**
- Description | **11478**
- Required Privilege Level | **11478**
- Release Information | **11478**

Syntax

```
no-intelligent-prescreening;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name scan-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name scan-options]
```

Description

Disables intelligent prescreening.

Intelligent prescreening tells the antivirus module to begin scanning a file much earlier. In this case, the scan engine uses the first packet or the first several packets to determine if a file could possibly contain malicious code. The scan engine does a quick check on these first packets and if the scan engine finds that it is unlikely that the file is infected, it then determines that it is safe to bypass the normal scanning procedure.

You can enable intelligent prescreening with the `intelligent-prescreening` statement.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

no-interface-hello

IN THIS SECTION

- Syntax | 11479
- Hierarchy Level | 11479
- Description | 11479
- Required Privilege Level | 11480
- Release Information | 11480

Syntax

```
no-interface-hello;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Explicitly disable RSVP interface hellos globally on the router.

> **NOTE**: For Junos OS Release 16.1 running on MX Series or PTX Series routers, the behavior of this statement has changed. On these platforms, rather than disabling RSVP interface hellos globally, the `no-interface-hello` command triggers a switch back to the previous profile for all label-switched paths (LSPs.)

This type of configuration might be necessary in networks where the Juniper Networks router has numerous RSVP connections with equipment from other vendors. However, if you disable RSVP interface hellos globally, you can also configure a hello interval on an RSVP interface using the "hello-interval" on page 8038 statement. This configuration disables RSVP interface hellos globally but enables RSVP interface hellos on the specified interface. This configuration might be necessary in a

heterogeneous network where some devices support RSVP node-ID hellos and other devices support RSVP interface hellos.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring RSVP Node-ID Hellos

hello-interval (Protocols RSVP) | **8038**

no-ip-tos-rewrite

IN THIS SECTION

- Syntax | **11480**
- Hierarchy Level | **11481**
- Description | **11481**
- Options | **11481**
- Required Privilege Level | **11481**
- Release Information | **11481**

Syntax

```
no-ip-tos-rewrite
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

Use the option of this command to allow DSCP propagation. When this CLI command is set or deleted, `evo-pfemand` process restarts.

Options

no-ip-tos-rewrite helpstring: Disable L3 code point rewrite at system level, restart PFE.

VLAN Priority Code Point (PCP) propagation is currently not supported and is set to zero. The new configuration statement can handle IP code point propagation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 22.1R1 (ACX7100-32 and ACX7100-48 routers).

RELATED DOCUMENTATION

DSCP Propagation and Default CoS on ACX Series Routers

no-ipsec-tunnel-in-traceroute

IN THIS SECTION

Syntax | 11482

- Hierarchy Level | 11482
- Description | 11482
- Required Privilege Level | 11482
- Release Information | 11482

Syntax

```
no-ipsec-tunnel-in-traceroute;
```

Hierarchy Level

```
[edit services ipsec-vpn]
```

Description

Disables displaying the IPsec tunnel endpoint in the trace route output. The IPsec tunnel is not treated as a next hop and TTL is not decremented. If the TTL becomes zero, the ICMP time exceeded message will not be generated.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring Security Associations

no-ipv4-routing

IN THIS SECTION

- Syntax | 11483
- Hierarchy Level | 11483
- Description | 11483
- Required Privilege Level | 11484
- Release Information | 11484

Syntax

```
no-ipv4-routing;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Disable IP version 4 (IPv4) routing.

Disabling IPv4 routing has the following results:

- The routing device does not advertise the network layer protocol identifier (NLPID) for IPv4 in the Junos OS link-state PDU fragment zero.
- The routing device does not advertise any IPv4 prefixes in Junos OS link-state PDUs.
- The routing device does not advertise the NLPID for IPv4 in Junos OS hello packets.
- The routing device does not advertise any IPv4 addresses in Junos OS hello packets.

- The routing device does not calculate any IPv4 routes.

> **NOTE**: Note: Even when **no-ipv4-routing** is configured, an IS-IS traceoptions log can list rejected IPv4 addresses. When a configuration is committed, IS-IS schedules a scan of the routing table to determine whether any routes need to be exported into the IS-IS link state database. The implicit default export policy action is to reject everything. IPv4 addresses from the routing table are examined for export, rejected by the default policy, and the rejections are logged.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

Understanding IS-IS IPv4 and IPv6 Unicast Topologies

no-ipv6-routing

IN THIS SECTION

- Syntax | **11485**
- Hierarchy Level | **11485**
- Description | **11485**
- Required Privilege Level | **11485**
- Release Information | **11485**

Syntax

```
no-ipv6-routing;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Disable IP version 6 (IPv6) routing.

Disabling IPv6 routing has the following results:

- The routing device does not advertise the network layer protocol identifier (NLPID) for IPv6 in the Junos OS link-state PDU fragment zero.
- The routing device does not advertise any IPv6 prefixes in Junos OS link-state PDUs.
- The routing device does not advertise the NLPID for IPv6 in Junos OS hello packets.
- The routing device does not advertise any IPv6 addresses in Junos OS hello packets.
- The routing device does not calculate any IPv6 routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

Understanding IS-IS IPv4 and IPv6 Unicast Topologies

no-irb-layer-2-copy

IN THIS SECTION

- Syntax | 11486
- Hierarchy Level | 11486
- Description | 11486
- Usage Guidelines | 11487
- Required Privilege Level | 11487
- Release Information | 11487

Syntax

```
no-irb-layer-2-copy;
```

Hierarchy Level

```
[edit bridge-domains],
[edit logical-routers logical-router-name bridge-domains],
[edit routing-instances routing-instance-name bridge-domains]
```

Description

If you include this statement when using port mirroring with Integrated Routing and Bridging (IRB), then the packet is mirrored as a Layer 3 packet. By default, the packet is mirrored as a Layer 2 packet. This statement is also supported if a routing instance is set to type VPLS.

Usage Guidelines

See *Configuring a Bridge Domain*

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring a Layer 2 Virtual Switch

no-issu-timer-negotiation

IN THIS SECTION

- Syntax | 11487
- Hierarchy Level | 11488
- Description | 11488
- Required Privilege Level | 11488
- Release Information | 11488

Syntax

```
no-issu-timer-negotiation;
```

Hierarchy Level

```
[edit protocols bfd],
[edit logical-systems logical-system-name protocols bfd],
[edit routing-instances routing-instance-name protocols bfd]
```

Description

Disable unified ISSU timer negotiation for Bidirectional Forwarding Detection (BFD) sessions.

> **CAUTION**: The sessions might flap during unified ISSU or Routing Engine switchover, depending on the detection intervals.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Managing and Tracing BFD Sessions During Unified ISSU Procedures

Junos OS Routing Protocols Library for Routing Devices

no-keepalives

IN THIS SECTION

- Syntax | 11489
- Hierarchy Level | 11489

- Description | 11489
- Required Privilege Level | 11489
- Release Information | 11489

Syntax

```
no-keepalives;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Disable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation. The default keepalive interval is 10 seconds.

For ATM2 IQ interfaces only, you can disable keepalives on a logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- `atm-ppp-llc`—PPP over AAL5 LLC encapsulation.
- `atm-ppp-vc-mux`—PPP over AAL5 multiplex encapsulation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring Keepalives
- *Disabling the Sending of PPPoE Keepalive Messages*
- Configuring Frame Relay Keepalives

no-keepalives (Dynamic Profiles)

IN THIS SECTION

- Syntax | 11490
- Hierarchy Level | 11490
- Description | 11490
- Required Privilege Level | 11490
- Release Information | 11491

Syntax

```
no-keepalives;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number],
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit”]
```

Description

Disable the sending of keepalives.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support of the `[edit dynamic-profiles profile-name]` hierarchy level introduced in Junos OS Release 9.5.

Support of the `[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"]` hierarchy level introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic Profiles Overview

Configuring Dynamic Authentication for PPP Subscribers

no-l2ckt

IN THIS SECTION

- Syntax | 11491
- Hierarchy Level | 11492
- Description | 11492
- Required Privilege Level | 11492
- Release Information | 11492

Syntax

```
no-l2ckt;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-
next-hop ingress],
[edit routing-options forwarding-table chained-composite-next-hop ingress]
```

Description

Disable composite chained next hop for ingress Layer 2 circuit label-switched paths (LSPs).

The remaining statements are explained separately.

> **NOTE**: For PTX Series routers, it is recommended that you do not use this command and disable composite chained next hop for ingress LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

no-l2vpn

IN THIS SECTION

- Syntax | 11493
- Hierarchy Level | 11493
- Description | 11493
- Required Privilege Level | 11493
- Release Information | 11493

Syntax

```
no-l2vpn {
    l2ckt;
    l3vpn;
    labeled-bgp;
    no-l2ckt;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-
next-hop ingress],
[edit routing-options forwarding-table chained-composite-next-hop ingress]
```

Description

Disable composite chained next hop for ingress Layer 2 virtual private networkk (VPN) label-switched paths (LSPs).

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

no-listen

IN THIS SECTION

- Syntax | 11494
- Hierarchy Level | 11494
- Description | 11494
- Required Privilege Level | 11494
- Release Information | 11495

Syntax

```
no-listen;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp interface (interface-name | interface-group)],
[edit forwarding-options helpers domain interface interface-name],
[edit forwarding-options helpers port port-number interface interface-name],
[edit forwarding-options helpers tftp interface interface-name]
```

Description

Disable recognition of DNS requests or stop packets from being forwarded on a logical interface, a group of logical interfaces, a router, or a switch.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding

Configuring Port-based LAN Broadcast Packet Forwarding

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

no-load-balance-label-capability

IN THIS SECTION

- Syntax | 11495
- Hierarchy Level | 11495
- Description | 11495
- Required Privilege Level | 11496
- Release Information | 11496

Syntax

```
no-load-balance-label-capability;
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Disables advertisement of entropy label capability in LDP and RSVP.

When you configure the `no-load-balance-label-capability` statement, it also disables the flow-aware transport of pseudowires (FAT) flow label for FEC 128.

The `load-balance-label-capability` and `no-load-balance-label-capability` statements at the `[edit forwarding-options]` hierarchy level are mutually exclusive, and at a given point in time, configuring one statement overrides the other.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

load-balance-label-capability | **9882**

entropy-label (protocols mpls) | **6416**

Configuring the Entropy Label for LSPs

no-load-throttle (Resource Monitor)

IN THIS SECTION

- Syntax | **11497**
- Hierarchy Level | **11497**
- Description | **11497**
- Required Privilege Level | **11497**
- Release Information | **11497**

Syntax

```
no-load-throttle;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

The no-load-throttle statement disables line card load-based throttling. Load-based throttling is also disabled when you configure the no-throttle statement.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

no-local-reversion

IN THIS SECTION

- Syntax | 11498
- Hierarchy Level | 11498
- Description | 11498
- Required Privilege Level | 11499
- Release Information | 11499

Syntax

```
local-reversion;
no-local-reversion;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Disable RSVP local revertive mode as specified in RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*.

> **NOTE**: For Junos OS Release 16.1 running on MX Series or PTX Series routers, `no-local-reversion` is enabled by default, that is, local reversion is not running, and the statement has been deprecated. To enable local reversion, use the `local-reversion` statement.

RSVP local revertive mode is supported on all Juniper Networks routers running Junos OS. It is the default behavior. If you include this statement, the Juniper Networks router uses global revertive mode instead. You might need to disable RSVP local revertive mode on Juniper Networks routers if your network includes equipment that does not support this mode.

The following information can also be found in RFC 4090. Refer to the RFC for additional information. When an LSP fails, the connection can be repaired locally using a traffic protection mechanism such as fast reroute. To restore the LSP to a full working path, RFC 4090 specifies the following strategies:

- Local revertive mode—Upon detecting that the path is restored, the point of local repair (PLR) resignals each of the LSPs that were formerly routed over the restored path. Every LSP successfully resignaled along the restored path is switched back.
- Global revertive mode—The ingress router of each tunnel is responsible for reoptimizing the LSPs that used the failed path. There are several potential reoptimization triggers: RSVP error messages, inspection of OSPF LSAs or IS-IS LSPs, and timers. This re-optimization process can proceed as soon as the failure is detected. It is not tied to the restoration of the failed path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

no-local-switching

IN THIS SECTION

- Syntax | 11499
- Hierarchy Level | 11499
- Description | 11499
- Required Privilege Level | 11500
- Release Information | 11500

Syntax

```
no-local-switching;
```

Hierarchy Level

```
[edit routing-instances instance-name]
```

Description

Specify that access ports in this routing instance do not forward packets to each other.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R2.

RELATED DOCUMENTATION

Configuring a Layer 2 Virtual Switch on an EX Series Switch

no-local-switching (VLAN)

IN THIS SECTION

- Syntax | 11500
- Hierarchy Level | 11500
- Description | 11501
- Required Privilege Level | 11501
- Release Information | 11501

Syntax

```
no-local-switching
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Description

Specify that access ports in this VLAN domain do not forward packets to each other. You use this statement with primary VLANs and isolated secondary VLANs. You can also disable local switching on both customer edge (CE) and VPLS edge (VE) mesh-groups. Access and core-facing interfaces are included in the system-generated CE mesh-group and VE mesh-group, respectively. Traffic originating on access interfaces, including EVPN tunnel traffic, will be flooded only to core-facing interfaces. Traffic originating on core-facing interfaces, including EVPN tunnel traffic, will be flooded only to access interfaces.

> **NOTE**: (MX80, MX104, and the 16x10GE MPC, MPC1, or MPC2 on MX240, MX480, MX960, MX2010, and MX2020 only) If you configure the `no-local-switching` command at the `[edit bridge-domains` *`bridge-domain-name`*`]` hierarchy level, it might not prevent multicast traffic from being forwarded between the CE-facing interfaces of the bridge domain. Broadcast, unknown unicast, and known multicast traffic does not exhibit this behavior.

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

routing—To view this statement in the configuration.

routing–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Support for EVPN-MPLS on PTX Series routers added in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

- Creating a Private VLAN on a Single QFX Switch
- Creating a Private VLAN Spanning Multiple QFX Series Switches

no-local-switching (VPLS)

IN THIS SECTION

- Syntax | 11502
- Hierarchy Level | 11502
- Description | 11502
- Required Privilege Level | 11502
- Release Information | 11503

Syntax

```
no-local-switching;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name]
```

Description

Prevent CE devices from communicating directly with each other. If the `no-local-switching` statement is configured, frames arriving on a CE interface are sent to a VPLS edge (VE) device or core-facing interfaces only.

> **NOTE**: (MX80, MX104, and the 16x10GE MPC, MPC1, or MPC2 on MX240, MX480, MX960, MX2010, and MX2020 only) If you configure the `no-local-switching` command at the `[edit bridge-domains bridge-domain-name]` hierarchy level, it might not prevent multicast traffic from being forwarded between the CE-facing interfaces of the bridge domain. Broadcast, unknown unicast, and known multicast traffic does not exhibit this behavior.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring VPLS and Integrated Routing and Bridging

no-logging (Resource Monitor)

IN THIS SECTION

- Syntax | 11503
- Hierarchy Level | 11503
- Description | 11504
- Options | 11504
- Required Privilege Level | 11504
- Release Information | 11504

Syntax

```
no-logging;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Disable the generation of error log messages when the utilization of memory resources exceeds the threshold or checkpoint levels. By default, messages are written to /var/log/rsmonlog.

Options

`no-logging` Disable the generation of error log messages when the utilization of memory resources exceeds the configured level. By default, error logs are recorded when the resource level utilization is exceeded.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

no-mac-learning

IN THIS SECTION

- Syntax | **11505**
- QFX Series and EX4600 | **11505**
- QFX Series per VLAN | **11505**
- EX Series Q-in-Q Interfaces | **11505**
- EX Series and SRX Series Q-inQ Vlans | **11505**
- ACX Series, MX Series, EX Series with ELS support, M Series, PTX Series, T Series | **11505**
- Description | **11506**
- Default | **11507**
- Required Privilege Level | **11507**
- Release Information | **11507**

Syntax

```
no-mac-learning;
```

QFX Series and EX4600

For QFX Series and EX4600 platforms without ELS:

```
[edit ethernet-switching-options interfaces interface-name]
```

For QFX Series and EX4600 platforms with ELS:

```
[edit vlans vlan-name switch-options]
```

QFX Series per VLAN

```
[edit vlans vlan-name]
```

```
[edit vlans vlan-name switch-options]
```

EX Series Q-in-Q Interfaces

```
[edit ethernet-switching-options interfaces interface-name]
```

EX Series and SRX Series Q-inQ Vlans

```
[edit vlans                 vlan-name]
```

ACX Series, MX Series, EX Series with ELS support, M Series, PTX Series, T Series

```
[edit bridge-domains bridge-domain-name bridge-options],
[edit bridge-domains bridge-domain-name bridge-options interface interface-name],
```

```
[edit logical-systems logical-system-name bridge-domains bridge-domain-name bridge-options],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name bridge-options
interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name switch-options],
[edit logical-systems logical-system-name switch-options],
[edit bridge-domains bridge-domain-name bridge-options interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options
interface interface-name],
[edit routing-instances routing-instance-name protocols evpn],
[edit routing-instances routing-instance-name protocols evpn interface interface-name],
[edit routing-instances routing-instance-name switch-options],
[edit switch-options],
[edit switch-options],
[edit switch-options interface interface-name],
[set vlans vlan-name switch-options]
```

Description

For QFX Series, EX Series switches and SRX Series Firewalls, disables MAC address learning for the specified VLAN.

For QFX Series and EX4600, disable MAC address learning for the specified interface. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member.

For EX Series switches' Q-in-Q interfaces, disables MAC address learning for the specified interface. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member.

For MX Series routers and EX Series switches with ELS support, disables MAC learning for a virtual switch, for a bridge domain or VLAN, for a specific logical interface in a bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port. On platforms that support EVPNs, you can disable MAC learning on an EVPN.

NOTE: When MAC learning is disabled for a VPLS routing instance, traffic is not load-balanced and only one of the equal-cost next hops is used.

Default

MAC learning is enabled.

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for the `switch-options` statement added in Junos OS Release 9.2.

Support for top-level configuration for the `virtual-switch` type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

`[edit switch-options]`, `[edit switch-options interface interface-name]`, `[edit vlans vlan-name switch-options]`, and `[edit vlans vlan-name switch-options interface interface-name]` hierarchy levels introduced in Junos OS Release 12.3 R2 for EX Series switches.

Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.

Hierarchy levels `[edit switch-options interface interface-name]` and `[edit vlans vlan-name switch-options]` introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Support for EVPN-MPLS on PTX Series routers added in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

Understanding Layer 2 Learning and Forwarding for Bridge Domains

Layer 2 Learning and Forwarding for VLANs Overview

Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports
Understanding Bridging and VLANs on Switches
Understanding Q-in-Q Tunneling and VLAN Translation
Understanding Q-in-Q Tunneling and VLAN Translation
Configuring Q-in-Q Tunneling on EX Series Switches

no-make-before-break

IN THIS SECTION

- Syntax | 11508
- Hierarchy Level | 11508
- Description | 11508
- Required Privilege Level | 11509
- Release Information | 11509

Syntax

```
no-make-before-break;
```

Hierarchy Level

```
[edit fabric routing-options multicast]
```

Description

(QFabric systems only) Disable the default *make-before-break* multicast feature on QFabric systems to prevent the duplication of system traffic. This feature increases the speed of data plane traffic, but carries the risk of minor traffic losses when compared with the default make-before-break method of creating new multicast fabric paths before tearing down the old paths. The absence of the `no-make-before-break` configuration statement implies that the make-before-break default system behavior is in effect.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Understanding the QFabric System Data Plane

no-management-vlan

IN THIS SECTION

- Syntax | 11509
- Hierarchy Level | 11509
- Description | 11510
- Required Privilege Level | 11510
- Release Information | 11510

Syntax

```
no-management-vlan;
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

Remove the specified member's out-of-band management port from the virtual management Ethernet (VME) global management VLAN of the Virtual Chassis or VCF configuration.

For a member that is functioning in a linecard role, you can use this configuration to reserve the member's management Ethernet port for local troubleshooting:

```
virtual-chassis {
    member 2 {
        no-management-vlan;
        }
    }
```

You cannot configure the IP address for a local management Ethernet port using the CLI or the J-Web interface. To do this, you need to use the shell `ifconfig` command.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Understanding Global Management of a Virtual Chassis

Understanding Virtual Chassis Fabric Configuration

no-mcast-replication

IN THIS SECTION

Syntax | 11511

- Hierarchy Level | 11511
- Description | 11511
- Default | 11511
- Required Privilege Level | 11511
- Release Information | 11512

Syntax

```
no-mcast-replication;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number],
[edit chassis lcc number fpc slot-number pic pic-number]
```

Description

For point-to-multipoint LSPs configured on T Series routers, protect the Packet Forwarding Engine (PFE) from bandwidth saturation. When a PFE does not need to replicate traffic, the PFE's bandwidth is less likely to become saturated. When you include the `no-mcast-replication` statement, the PFE is forced to be a leaf node in the binary tree. Leaf nodes, unlike branch nodes, do not replicate traffic in the process of forwarding traffic. Because leaf nodes have no children, they do not need to replicate traffic, and thus are less likely to become saturated with traffic.

Default

If you omit the `no-mcast-replication` statement, the PFE can become a branch node or a leaf node. When the PFE becomes a branch node, the PFE must replicate traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Point-to-Multipoint LSPs Overview

no-multicast

IN THIS SECTION

- Syntax | 11512
- Hierarchy Level | 11512
- Description | 11513
- Default | 11513
- Required Privilege Level | 11513
- Release Information | 11513

Syntax

```
no-multicast;
```

Hierarchy Level

- For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers):

```
[edit forwarding-options storm-control-profiles profile-name all]
```

- For platforms without ELS:

```
[edit ethernet-switching-options storm-control interface (all | interface-name)]
```

Description

Disable storm control for all multicast traffic (both registered multicast and unregistered multicast) for the specified interface or for all interfaces.

Default

- On EX2200, EX3200, EX3300, and EX4200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams.
- On EX4300 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. You can selectively disable storm control on any type of traffic.
- On EX4500 and EX8200 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. On EX8200 switches, you can selectively disable storm control on registered multicast traffic, on unregistered multicast traffic, or on both types of multicast traffic.
- On EX6200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams. You can selectively disable storm control for each type of traffic individually.
- On EX9200 switches—Storm control is not enabled by default.
- On MX Series routers—Storm control is not enabled by default.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.3.

Hierarchy level `[edit forwarding-options]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

- no-registered-multicast | 11562
- no-unregistered-multicast | 11609
- *Enabling and Disabling Storm Control (non-ELS)*
- *Enabling and Disabling Storm Control (ELS)*

no-multicast-echo

IN THIS SECTION

- Syntax | 11514
- Hierarchy Level | 11514
- Description | 11515
- Default | 11515
- Required Privilege Level | 11515
- Release Information | 11515

Syntax

```
no-multicast-echo
```

Hierarchy Level

```
[edit system]
```

Description

Disable the Routing Engine from responding to ICMP echo requests sent to multicast group addresses.

Default

The Routing Engine responds to ICMP echo requests sent to multicast group addresses.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configure ICMP Features

no-nat-traversal (Services IPsec VPN)

IN THIS SECTION

- Syntax | **11516**
- Hierarchy Level | **11516**
- Description | **11516**
- Required Privilege Level | **11516**
- Release Information | **11516**

Syntax

```
no-nat-traversal;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Configure to disable NAT-T at the services-set level (tunnel level). NAT-T is enabled by default therefore you must use the `no-nat-traversal` for disabling the NAT-T.

> **NOTE**: The global disable NAT-T setting at [edit services ipsec-vpn] hierarchy level overrides the default NAT-T setting at [edit service-set] hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Disabling NAT-T on MX Series Routers for Handling NAT with IPsec-Protected Packets

no-native-vlan-insert

IN THIS SECTION

- Syntax | 11517
- Hierarchy Level | 11517
- Description | 11517
- Default | 11518
- Required Privilege Level | 11518
- Release Information | 11518

Syntax

```
no-native-vlan-insert;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Send traffic without the native VLAN ID ("native-vlan-id" on page 11240) to the remote end of the network if untagged traffic is received.

If this statement is not configured, then "native-vlan-id" on page 11240 is added to untagged traffic. But if this statement is configured, then "native-vlan-id" on page 11240 is not added to untagged traffic.

NOTE:

- This feature works only on MX Series routers with MPCs/MICs. Configuring this statement on MX Series routers with DPCs results is no behavioral change. However,

if you configure the statement on aggregated Ethernet (ae) interfaces with logical interfaces across MPCs/MICs and DPCs, then the MPCs/MICs and DPCs behave differently.

- In the egress direction, this feature is disrupted by VLAN normalization. Because of normalization, the egress interface cannot distinguish between untagged traffic and tagged traffic. And untagged traffic is sent out with "native-vlan-id" on page 11240. Consider this while configuring both VLAN normalization and new "native-vlan-id" on page 11240 statement.

  There will be a problem with ingress firewall filter if filter term includes "native-vlan-id" on page 11240. With `no-native-vlan-insert` statement configured, "native-vlan-id" on page 11240 will not be inserted to untagged traffic. So, firewall filter term will not match with untagged traffic. But if incoming traffic have VLAN ID which is equal to "native-vlan-id" on page 11240, then firewall filter term will match and firewall will work.

- When this feature is used with AE, all sub-interfaces of AE should be in same type of FPC.

Default

By default, "native-vlan-id" on page 11240 is inserted to untagged traffic. That is, if this statement is not configured, then "native-vlan-id" on page 11240 is inserted to untagged traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

Configuring Mixed Tagging Support for Untagged Packets

Configuring Access Mode on a Logical Interface

Configuring the Native VLAN Identifier on Switches With ELS Support

Understanding Bridging and VLANs on Switches

flexible-vlan-tagging | **7235**

native-vlan-id | **11240**

Understanding Q-in-Q Tunneling and VLAN Translation

Sending Untagged Traffic Without VLAN ID to Remote End

no-neighbor-down-notification

IN THIS SECTION

- Syntax | **11519**
- Hierarchy Level | **11519**
- Description | **11519**
- Required Privilege Level | **11520**
- Release Information | **11520**

Syntax

```
no-neighbor-down-notification;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id interface interface-name],
[edit protocols ospf area area-id interface interface-name]
```

Description

Disable neighbor down notification for OSPF to allow for migration from OSPF to IS-IS without disruption of the RSVP neighbors and associated RSVP-signaled LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

no-neighbor-learn

IN THIS SECTION

- Syntax | 11520
- Hierarchy Level | 11521
- Description | 11521
- Required Privilege Level | 11521

Syntax

```
no-neighbor-learn {
    accounting  {
        destination-class-usage;
        source-class-usage direction;
    }
    address address;
    dhcp{
        client-identifier (ascii ascii | hexadecimal hexadecimal);
        lease-time (seconds | infinte);
        retransmission-attempt (DHCP Client) number;
        retransmission-interval (DHCP Client) seconds;
        server-address ip-address;
        update-server;
        vendor-id vendor-id;
    }
```

```
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    no-redirects;
    primary;
    rpf-check;
    targeted-broadcast;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

Disable neighbor address learning on this interface (for both control plane and transit traffic) for all addresses or for the named address.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- Configuring Gigabit Ethernet Interfaces (CLI Procedure)
- Configuring Routed VLAN Interfaces on Switches (CLI Procedure)
- *Configure ARP Learning and Aging Options*

no-next-header-to-payload-protocol-mapping

IN THIS SECTION

- Syntax | 11522
- Hierarchy Level | 11522
- Description | 11522
- Required Privilege Level | 11523
- Release Information | 11523

Syntax

```
no-next-header-to-payload-protocol-mapping
```

Hierarchy Level

```
[edit firewall]
```

Description

Is used to control the behavior of the next-header match condition. When enabled, next-header matches the first extension header and if disabled next-header matches value of the payload/upper-layer protocol.

The next-header and payload-protocol match conditions were introduced in Junos releases 21.4R2-S1-EVO, 21.4R3-S1-EVO, 22.2R2-S1-EVO, 22.3R1-EVO and onwards. Previous to these releases, such as 21.4R3-EVO and 22.1R1-EVO, payload-protocol match condition did not exist; next-header was used to match the value of the payload/upper-layer protocol.

Currently, on Junos releases that support next-header and payload-protocol, the next-header match condition cannot be used when the filter using this match condition is bound to a WAN egress interface and/or the lo0 ingress interface. Customers upgrading to such Junos releases could have next-header match conditions attached to a WAN egress interface and/or the lo0 ingress interface. The upgrade fails in such scenarios.

Even if the filter is bound to supported interfaces, customers might get unexpected behavior. It is because of the change in semantics of the next-header match condition, where on previous releases it was expected to match the payload-protocol, but on newer releases, Junos automatically does an internal conversion to make it match the first extension header.

Customers are expected to modify their configurations before an upgrade to avoid such issues. The `no-next-header-to-payload-protocol-mapping` allows customers time to adapt to the new behavior of next-header to prevent a failure during upgrade or prevent unexpected behavior. This configuration command helps the user to control whether the next-header match condition should exhibit its previous behavior of matching the payload-protocol or current behavior of matching the first extension header.

Required Privilege Level

- interface—To view this statement in the configuration.
- interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2

no-node-hello-on-bypass

IN THIS SECTION

- Syntax | 11523
- Hierarchy Level | 11524
- Description | 11524
- Required Privilege Level | 11524
- Release Information | 11524

Syntax

```
no-node-hello-on-bypass;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Disable the sending of RSVP hellos over a bypass LSP. Starting with Junos release 16.1, Junos routers send RSVP hello packets over a bypass LSP (when one is present), instead of the IGP next hop. To return to the original behavior specify the `no-node-hello-on-bypass` option.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring RSVP Node-ID Hellos

no-node-id-subobject

IN THIS SECTION

- Syntax | 11525
- Hierarchy Level | 11525
- Description | 11525
- Required Privilege Level | 11525
- Release Information | 11525

Syntax

```
no-node-id-subobject;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Disable the record route object (RRO) node-ID subobject for compatibility with earlier versions of Junos OS.

> **NOTE**: For Junos OS Release 16.1 running on MX Series or PTX Series routers, the behavior of this statement has changed. On these platforms, rather than disabling the record route object (RRO) node ID sub-object, the `no-node-id-subobject` command triggers a switch back to the previous profile for all label-switched paths (LSPs).

To interoperate with other vendors' equipment, Junos OS supports the RRO node-ID subobject for use in inter-AS link and node protection configurations.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Inter-AS Node and Link Protection

no-normalization

IN THIS SECTION

- Syntax | 11526
- Hierarchy Level | 11526
- Description | 11526
- Required Privilege Level | 11527
- Release Information | 11527

Syntax

```
no-normalization;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name],
[edit logical-systems name routing-instances name bridge-domains name bridge-options],
[edit routing-instances name ],
[edit routing-instances name bridge-domains name bridge-options]
```

Description

The `no-normalization` statement disables VLAN ID (VID) normalization for interfaces. This allows the configuration of VLAN-based service for EVPN with VID translation. This means the customer can have a different VID per Ethernet segment per PE with a strict compliance of RFC 7432. Use of this statement sets the Ethernet tag ID to zero.

Configuration of `vlan-id none` and `no-normalization` statements together enable transportation of dual-tagged frames across the EVPN core network in a virtual switch instance type.

Following are examples of the prerequisites for and use of the `no-normalization` statement on routing instance r1:

- `no-normalization` statement can be configured Only for EVPN instances:

```
set routing-instances r1 instance-type evpn
set routing-instances r1 no-normalization
```

- `no-normalization` of VLAN can be configured Only when `instance-type virtual-switch` along with `protocols evpn` is preconfigured:

```
set routing-instances r1 instance-type virtual-switch
set routing-instances r1 protocols evpn
set routing-instances r1 bridge-domains xyz bridge-options no-normalization
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

routing-instances (Logical Systems) | **15966**

instance-type | **8773**

VLAN-Based Service for EVPN

Virtual Switch Support for EVPN Overview

no-notify-mail-recipient

IN THIS SECTION

- Syntax | 11528
- Hierarchy Level | 11528
- Description | 11528
- Required Privilege Level | 11528
- Release Information | 11529

Syntax

```
no-notify-mail-recipient;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-non-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-non-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-non-block]
```

Description

Do not notify the e-mail recipient about errors returned by the antivirus scan engine when a fallback nonblocking action occurs.

You can specify that the e-mail recipient is to be notified with the `notify-mail-recipient` statement.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

no-notify-mail-sender (Security Content Filtering Notification Options)

IN THIS SECTION

- Syntax | 11529
- Hierarchy Level | 11529
- Description | 11529
- Required Privilege Level | 11530
- Release Information | 11530

Syntax

```
no-notify-mail-sender;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name notification-options]
```

Description

Do not notify the e-mail sender.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

no-notify-mail-sender (Security Fallback Block)

IN THIS SECTION

- Syntax | 11530
- Hierarchy Level | 11531
- Description | 11531
- Required Privilege Level | 11531
- Release Information | 11531

Syntax

```
no-notify-mail-sender;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-block]
```

Description

Do not notify the e-mail sender about errors returned by the antivirus scan engine when a fallback action occurs.

You can specify that the e-mail sender is to be notified with the `notify-mail-sender` statement.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

no-notify-mail-sender (Security Virus Detection)

IN THIS SECTION

- Syntax | 11532
- Hierarchy Level | 11532

- Description | 11532
- Required Privilege Level | 11532
- Release Information | 11532

Syntax

```
no-notify-mail-sender;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options virus-detection]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options virus-detection]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options virus-detection]
```

Description

Do not notify the e-mail sender when a virus is detected by the antivirus engine.

You can specify that the e-mail sender is to be notified with the `notify-mail-sender` statement.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The [edit security utm default-configuration] hierarchy level is introduced in Junos OS Release 18.2R1.

no-nssa-abr

IN THIS SECTION

- Syntax | 11533
- Hierarchy Level | 11533
- Description | 11534
- Required Privilege Level | 11534
- Release Information | 11534

Syntax

```
no-nssa-abr;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Disable exporting Type 7 link-state advertisements into not-so-stubby-areas (NSSAs) for an autonomous system boundary router (ASBR) or an area border router (ABR).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring OSPF Not-So-Stubby Areas

no-odu-backward-frr-enable

IN THIS SECTION

- Syntax | 11535
- Hierarchy Level | 11535
- Description | 11535
- Default | 11535
- Required Privilege Level | 11535
- Release Information | 11535

Syntax

```
no-odu-backward-frr-enable;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

Description

Disable preemptive fast reroute (FRR) ODU backward FRR insertion.

Default

By default, FRR ODU backward FRR insertion is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

> *100-Gigabit Ethernet OTN Options Configuration Overview*
>
> *Configuring OTN Interfaces on P1-PTX-2-100G-WDM*

no-odu-signal-degrade-monitor-enable

IN THIS SECTION

- Syntax | 11536
- Hierarchy Level | 11536
- Description | 11536
- Default | 11536
- Required Privilege Level | 11536
- Release Information | 11537

Syntax

```
no-odu-signal-degrade-monitor-enable;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

Description

Disable monitoring of signal degradation of ODU BER in the received OTN frames.

Default

By default, FRR signal degrade monitoring disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

no-option16

IN THIS SECTION

- Syntax | 11537
- Hierarchy Level | 11537
- Description | 11538
- Required Privilege Level | 11538
- Release Information | 11538

Syntax

```
no-option16;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name overrides]
```

Description

Configure a specific group of one or more access interfaces within the VLAN not to transmit DHCPv6 option 16 information, even if the VLAN is configured to perform DHCPv6 snooping. Option 16 information that has already been added by a DHCPv6 client will be forwarded as is.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

- option-16 (DHCPv6 Snooping) | **11868**
- *Understanding DHCP Snooping (ELS)*
- *Understanding DHCP Option 82*

no-option18

IN THIS SECTION

- Syntax | **11539**
- Hierarchy Level | **11539**
- Description | **11539**
- Required Privilege Level | **11539**
- Release Information | **11539**

Syntax

```
no-option18;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name overrides]
```

Description

Configure a specific group of one or more access interfaces within the VLAN *not* to transmit DHCP option 18 information, even if the VLAN is configured to perform DHCPv6 snooping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

option-18 (DHCPv6 Snooping) | **11869**

Understanding DHCP Option 82

Understanding DHCP Snooping (ELS)

no-option37

IN THIS SECTION

- Syntax | 11540
- Hierarchy Level | 11540
- Description | 11540
- Required Privilege Level | 11540
- Release Information | 11540

Syntax

```
no-option37;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name overrides]
```

Description

Configure a specific group of one or more access interfaces within the VLAN *not* to transmit DHCP option 37 information, even if the VLAN is configured to perform DHCPv6 snooping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D20.

RELATED DOCUMENTATION

option-37 (DHCPv6 Snooping) | 11872

Understanding DHCP Option 82

Understanding DHCP Snooping (non-ELS)

no-option82

IN THIS SECTION

- Syntax | 11541
- Hierarchy Level (EX Series, QFX Series) | 11541
- Hierarchy Level (MX Series) | 11541
- Description | 11542
- Required Privilege Level | 11542
- Release Information | 11542

Syntax

```
no-option82;
```

Hierarchy Level (EX Series, QFX Series)

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name overrides]
```

Hierarchy Level (MX Series)

```
[edit bridge-domains bridge-domain-name forwarding-options group group group-name overrides]
```

Description

Configure a specific group of one or more access interfaces within the VLAN or bridge domain *not* to transmit DHCP option 82 information, even if the VLAN or bridge domain is configured to use option 82.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

- option-82 | **11884**
- *Setting Up DHCP Option 82 on the Switch with No Relay (ELS)*
- *Understanding DHCP Option 82*
- *Understanding DHCP Snooping (non-ELS)*

no-overload-on-prefix-export-limit

IN THIS SECTION

- Syntax | **11543**
- Hierarchy Level | **11543**
- Description | **11543**
- Required Privilege Level | **11543**
- Release Information | **11543**

Syntax

```
no-overload-on-prefix-export-limit;
```

Hierarchy Level

```
[edit logical-systems name protocols isis dynamic-overload],
[edit logical-systems name routing-instances name protocols isis dynamic-overload],
[edit logical-systems name tenants name routing-instances name protocols isis dynamic-overload],
[edit protocols isis dynamic-overload],
[edit routing-instances name protocols isis dynamic-overload],
[edit tenants name routing-instances name protocols isis dynamic-overload]
```

Description

Prevent ISIS entering the Overload state even when the number of prefixes to advertise are more than the configured prefix limit.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.1

RELATED DOCUMENTATION

overload (Protocols IS-IS) | **12048**

prefix-export-limit (Protocols IS-IS) | **13083**

no-p2mp-sublsp

IN THIS SECTION

- Syntax | 11544
- Hierarchy Level | 11544
- Description | 11544
- Default | 11544
- Required Privilege Level | 11545
- Release Information | 11545

Syntax

```
no-p2mp-sublsp;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Reject Resv messages that include the S2L_SUB_LSP object. By default, Resv messages that include the S2L_SUB_LSP object are accepted. However, in a network which includes Juniper Networks devices running both Junos OS Release 9.2 and later and Junos OS Release 9.1 and earlier, it is necessary to configure the `no-p2mp-sublsp` statement on devices running Junos OS Release 9.2 and later to ensure that point-to-multipoint LSPs function properly.

Default

Resv messages that include the S2L_SUB_LSP object are accepted.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases

no-per-unit-scheduler

IN THIS SECTION

- Syntax | 11545
- Hierarchy Level | 11545
- Description | 11546
- Required Privilege Level | 11546
- Release Information | 11546

Syntax

```
no-per-unit-scheduler;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

To enable traffic control profiles to be applied at FRF.16 bundle (physical) interface level, disable the per-unit scheduler, which is enabled by default. This statement and the `shared-scheduler` statement are mutually exclusive.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 11.4.

RELATED DOCUMENTATION

Oversubscribing Interface Bandwidth

no-ping-record-route

IN THIS SECTION

- Syntax | 11546
- Hierarchy Level | 11547
- Description | 11547
- Required Privilege Level | 11547
- Release Information | 11547

Syntax

```
no-ping-record-route;
```

Hierarchy Level

```
[edit system]
```

Description

Configure the Junos OS to disable the reporting of the IP address in ping responses.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configure ICMP Features

no-ping-time-stamp

IN THIS SECTION

- Syntax | 11548
- Hierarchy Level | 11548
- Description | 11548
- Required Privilege Level | 11548
- Release Information | 11548

Syntax

```
no-ping-time-stamp;
```

Hierarchy Level

```
[edit system]
```

Description

Configure the Junos OS to disable the recording of timestamps in ping responses.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configure ICMP Features

no-policy-cold-synchronization

IN THIS SECTION

- Syntax | 11549
- Hierarchy Level | 11549
- Description | 11549

Required Privilege Level | 11549
Release Information | 11549

Syntax

no-policy-cold-synchronization

Hierarchy Level

```
[edit security idp sensor-configuration high-availability]
```

Description

Disable policy cold synchronization functionality. This prevents the SRX Series Firewalls from waiting for the IPS policy to be loaded on all service PICs, and there by forfeits IPS service protection during the ISSU (In-service software upgrade) operation/window.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Security Policies Overview

no-power-budget

IN THIS SECTION

- Syntax | 11550
- Hierarchy Level | 11550
- Description | 11550
- Required Privilege Level | 11551
- Release Information | 11551

Syntax

```
no-power-budget;
```

Hierarchy Level

```
[edit chassis]
```

Description

Disable power budgeting. If you disable the power management on the PTX10008, the router keeps all the FRUs powered on by default. The router does not move any of the FRUs to offline state even if it does not have sufficient power. However, in case of a power shortage, a power redundancy alarm is raised as shown in the following example.

```
user@router> show system alarms

1 alarm currently active

Alarm time Class Description

2019-07-25 21:16:25 UTC Major chassis No Redundant Power
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.1R2.

no-pre-classifier

IN THIS SECTION

- Syntax | 11551
- Hierarchy Level | 11551
- Description | 11552
- Default | 11552
- Required Privilege Level | 11552
- Release Information | 11552

Syntax

```
no-pre-classifier;
```

Hierarchy Level

```
[edit chassis fpc n pic n]
```

Description

Specify disabling the control queue for all ports on the 10-Gigabit Ethernet LAN/WAN PIC. Deleting this configuration re-enables the control queue feature on all ports of the 10-Gigabit Ethernet LAN/WAN PIC.

> **NOTE**: For the 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP), the control queue has a rate limiter to limit the control traffic to 2 Mbps (fixed, not user-configurable) per port. If the transit control traffic crosses this limit, then it can cause drops on locally terminating control traffic, causing flap of protocols such as BGP and OSPF. To avoid the control traffic being dropped, configure the `no-pre-classifier` statement to disable the control queue.

Default

The `no-pre-classifier` statement is not configured and the control queue is operational.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

10-port 10-Gigabit Ethernet LAN/WAN PIC Overview

Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC

no-propagate-ttl

IN THIS SECTION

- Syntax | 11553
- Hierarchy Level | 11553
- Description | 11553
- Default | 11553
- Required Privilege Level | 11554
- Release Information | 11554

Syntax

```
no-propagate-ttl;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Disable normal time-to-live (TTL) decrementing. You configure this statement once per router, and it affects all RSVP-signaled or LDP-signaled LSPs. When this router acts as an ingress router for an LSP, it pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When the router acts as the penultimate router, it pops the MPLS header without writing the MPLS TTL into the IP packet.

When you add the `no-propagate-ttl` statement to the configuration or delete it from the configuration, the effect takes place immediately. There is no need to clear existing RSVP LSPs or LDP sessions.

Default

Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1 on PTX10008.

RELATED DOCUMENTATION

Disabling Normal TTL Decrementing

Example: Diagnosing Networking Problems Related to Layer 3 VPNs by Disabling TTL Decrementing

Layer 3 VPNs User Guide for Routing Devices

no-decrement-ttl | **11432**

no-qos-adjust

IN THIS SECTION

- Syntax | **11554**
- Hierarchy Level | **11555**
- Description | **11555**
- Required Privilege Level | **11555**
- Release Information | **11555**

Syntax

```
no-qos-adjust;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast interface interface-name reverse-oif-mapping],
[edit logical-systems logical-system-name routing-options multicast interface interface-name],
[edit logical-systems logical-system-name routing-options multicast interface interface-name
reverse-oif-mapping],
[edit routing-instances routing-instance-name routing-options multicast interface interface-name],
[edit routing-instances routing-instance-name routing-options multicast interface interface-name
reverse-oif-mapping],
[edit routing-options multicast interface interface-name],
[edit routing-options multicast interface interface-name reverse-oif-mapping]
```

Description

Disable hierarchical bandwidth adjustment for all subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Statement added to **[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*]**, and **[edit routing-options multicast interface *interface-name*]** hierarchy levels in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring Multicast with Subscriber VLANs

no-qos-adjust (Dynamic Routing Options)

IN THIS SECTION

- Syntax | 11556
- Hierarchy Level | 11556
- Description | 11556
- Required Privilege Level | 11556
- Release Information | 11556

Syntax

```
no-qos-adjust;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-options multicast interface interface-name]
```

Description

Disable hierarchical bandwidth adjustment for all dynamically created subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring Multicast with Subscriber VLANs

no-recommendation-tlv

IN THIS SECTION

- Syntax | 11557
- Hierarchy Level | 11557
- Description | 11557
- Default | 11557
- Required Privilege Level | 11558
- Release Information | 11558

Syntax

```
no-recommendation-tlv;
```

Hierarchy Level

```
[edit protocols dcbx interface interface-name enhanced-transmission-selection]
```

Description

Disable DCBX to send the ETS Recommendation TLV (also known as the Information TLV) on egress. This feature is valid only if the interface DCBX mode is IEEE DCBX. If the interface DCBX mode is DCBX version 1.01, this statement has no effect. (DCBX version 1.01 does not advertise separate TLVs for individual attributes.)

Default

DCBX-enabled interfaces send the ETS recommendation TLV unless it is disabled.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

show dcbx neighbors | 19191

Configuring DCBX Autonegotiation

no-recommended

IN THIS SECTION

- Syntax | 11558
- Hierarchy Level | 11558
- Description | 11559
- Required Privilege Level | 11559
- Release Information | 11559

Syntax

```
no-recommended;
```

Hierarchy Level

```
[edit security idp dynamic-attack-group dynamic-attack-group-name filters]
```

Description

Specify non recommended attack objects in the dynamic attack group.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4R6.

RELATED DOCUMENTATION

Understanding IDP Policy Rules

no-redirects

IN THIS SECTION

- Syntax | 11559
- Hierarchy Level | 11560
- Description | 11560
- Default | 11560
- Required Privilege Level | 11560
- Release Information | 11560

Syntax

```
no-redirects;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family]
```

Description

Do not send protocol redirect messages on the interface.

To disable the sending of protocol redirect messages for the entire router or switch, include the `no-redirects` statement at the `[edit system]` hierarchy level.

Default

Interfaces send protocol redirect messages.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Disabling the Transmission of Redirect Messages on an Interface
- Junos OS Administration Library for Routing Devices

no-redirects (IPv4 Traffic)

IN THIS SECTION

Syntax | 11561

- Hierarchy Level | 11561
- Description | 11561
- Default | 11561
- Required Privilege Level | 11562
- Release Information | 11562

Syntax

```
no-redirects;
```

Hierarchy Level

```
[edit system]
[edit interfaces interface-name unit logical-unit-number family family]
```

Description

Stop protocol redirect messages for IPv4 traffic from being sent on the entire device or on an interface on the device.

To disable the sending of protocol redirect messages for the entire device, include the `no-redirects` statement at the `[edit system]` hierarchy level.

To disable the sending of protocol redirect messages on a specific interface, include the `no-redirects` statement at the `[edit interfaces interface-name unit logical-unit-number family family]` hierarchy level.

> **NOTE**: Hyper-mode is the default forwarding mode on the SCBE3-MX. In Junos versions prior to 19.2R1, ICMP redirects are disabled by default and cannot be re-enabled in hyper mode.

Default

For EX Series Switches: The interface sends redirect messages.

For other devices: The device sends redirect messages.

Required Privilege Level

For EX Series Switches:

- interface—To view this statement in the configuration.
- interface-control—To add this statement to the configuration.

For other devices:

- system—To view this statement in the configuration.
- system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Protocol Redirect Messages

Junos OS Network Interfaces Library for Routing Devices

no-registered-multicast

IN THIS SECTION

- Syntax | 11563
- Hierarchy Level | 11563
- Description | 11563
- Default | 11563
- Required Privilege Level | 11564
- Release Information | 11564

Syntax

```
no-registered-multicast;
```

Hierarchy Level

- For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers):

  ```
  [edit forwarding-options storm-control-profiles profile-name all]
  ```

- For platforms without ELS:

  ```
  [edit ethernet-switching-options storm-control interface (all | interface-name)]
  ```

Description

(EX8200 switches only) Disable storm control for registered multicast traffic for the specified interface or for all interfaces.

(EX4300 and EX9200 switches only) Exclude storm control for registered multicast traffic from a storm control profile.

(MX Series routers only) Exclude storm control for registered multicast traffic from a storm control profile.

Default

EX4300 and EX8200 switches—Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic. The default storm control level is 80 percent of the available bandwidth used by the combined applicable traffic streams.

EX9200 switches—Storm control is not enabled by default.

MX Series routers—Storm control is not enabled by default.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.3.

Hierarchy level `[edit forwarding-options]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

- no-multicast | **11512**
- no-unregistered-multicast | **11609**
- *Understanding Storm Control*
- *Understanding Storm Control*

no-remote-trace (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | **11564**
- Hierarchy Level | **11565**
- Description | **11565**
- Required Privilege Level | **11565**
- Release Information | **11565**

Syntax

```
no-remote-trace;
```

Hierarchy Level

```
[edit services rtlog traceoptions]
```

Description

Disable remote tracing

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

no-remote-trace (System)

IN THIS SECTION

- Syntax | 11566
- Hierarchy Level | 11566
- Description | 11566
- Required Privilege Level | 11566
- Release Information | 11566

Syntax

```
no-remote-trace;
```

Hierarchy Level

```
[edit system scripts commit traceoptions]
```

Description

Disable remote tracing.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

traceoptions (Security Datapath Debug) | 29783

no-remote-trace (Trace Options)

IN THIS SECTION

- Syntax | 11567
- Hierarchy Level | 11567
- Description | 11567

Required Privilege Level | 11567
Release Information | 11567

Syntax

```
no-remote-trace;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions],
[edit forwarding-options sampling traceoptions]
```

Description

Disable remote tracing.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing Traffic Sampling Operations

no-revert (Local Switching)

IN THIS SECTION

- Syntax | 11568
- Hierarchy Level | 11568
- Description | 11568
- Required Privilege Level | 11569
- Release Information | 11569

Syntax

```
no-revert;
```

Hierarchy Level

```
[edit protocols l2circuit local-switching interfaces interface-name]
[edit protocols l2circuit local-switching interfaces interface-name end-interface interface-name]
```

Description

(Optional) Prevent the local switching interface from reverting to the primary interface.

NOTE: The protect interface must be configured prior to configuring the `no-revert` statement.

Typically, when the primary interface goes down, the pseudowire starts using the protect interface. By default, when the primary interface comes back online, the interface is switched-over back from the protect interface to the primary interface. To prevent the switchover back to the primary interface, unless the primary interface goes down, include the `no-revert` statement. This prevents loss of traffic during the switchover.

> **NOTE**: If the protect interface fails, the interface is switched-over back to the primary interface, irrespective of whether or not the `no-revert` statement is included in the configuration.

This statement can be configured both for the starting interface and the ending interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Local Interface Switching in Layer 2 Circuits

no-revert (Neighbor Interface)

IN THIS SECTION

- Syntax | 11570
- Hierarchy Level | 11570
- Description | 11570
- Required Privilege Level | 11570
- Release Information | 11570

Syntax

```
no-revert;
```

Hierarchy Level

```
[edit protocols l2circuit neighbor address interfaces interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interfaces
interface-name]
```

Description

(Optional) Prevent the protect interface from reverting to the primary interface.

> **NOTE**: The protect interface must be configured prior to configuring the `no-revert` statement.

Typically, when the primary interface goes down, the pseudowire starts using the protect interface. By default, when the primary interface comes back online, the interface is switched-over back from the protect interface to the primary interface. To prevent the switchover back to the primary interface, unless the protect interface goes down, include the `no-revert` statement. This prevents loss of traffic during the switchover.

> **NOTE**: If the protect interface fails, the interface is switched-over back to the primary interface, irrespective of whether or not the `no-revert` statement is included in the configuration.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Interfaces for Layer 2 Circuits

no-rfc-1583

IN THIS SECTION

- Syntax | 11571
- Hierarchy Level | 11571
- Description | 11572
- Default | 11572
- Required Privilege Level | 11572
- Release Information | 11572

Syntax

```
no-rfc-1583;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Disable compatibility with RFC 1583, *OSPF Version 2*. If the same external destination is advertised by AS boundary routers that belong to different OSPF areas, disabling compatibility with RFC 1583 can prevent routing loops.

Default

Compatibility with RFC 1583 is enabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Example: Disabling OSPFv2 Compatibility with RFC 1583

no-root-port

IN THIS SECTION

- Syntax | **11573**
- Hierarchy Level | **11573**
- Description | **11573**
- Required Privilege Level | **11574**

Release Information | 11574

Syntax

```
no-root-port;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp | vstp) interface interface-name],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp | vstp) interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id interface interface-name],
```

```
[edit protocols (mstp | rstp | vstp) interface interface-name],
[edit protocols vstp vlan vlan-id interface interface-name],
[edit protocols mstp interface (all | interface-name) arp-on-stp],
[edit protocols rstp interface (all | interface-name) arp-on-stp],
[edit protocols stp interface (all | interface-name) arp-on-stp],
```

```
[edit routing-instances routing-instance-name protocols (mstp | rstp | vstp) interface interface-
name],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id interface interface-name]
[edit protocols vstp vlan vlan-id interface (all | interface-name) arp-on-stp]
```

Description

Configures an interface to be a spanning-tree designated port. If the bridge receives superior STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding VSTP

Understanding Root Protection for Spanning-Tree Instance Interfaces in a Layer 2 Switched Network

Enabling Root Protection for a Spanning-Tree Instance Interface

Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on non-ELS EX Series Switches

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

no-route-localize

IN THIS SECTION

- Syntax | **11575**
- Hierarchy Level | **11575**
- Description | **11575**
- Required Privilege Level | **11575**
- Release Information | **11575**

Syntax

```
no-route-localize;
```

Hierarchy Level

```
[edit policy-options policy-statement policy-name term term-name then]
```

Description

Enforce installation of routes on all FIB-remote Packet Forwarding Engines.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring Packet Forwarding Engine FIB Localization

no-sbl-default-server

IN THIS SECTION

- Syntax | 11576
- Hierarchy Level | 11576
- Description | 11576

Required Privilege Level | 11576
Release Information | 11576

Syntax

```
no-sbl-default-server;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-spam sbl profile profile-name]
```

Description

Disable the default SBL server lookup.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Antispam Filtering Overview

no-send-pads-ac-info

IN THIS SECTION

- Syntax | 11577
- Hierarchy Level | 11577
- Description | 11577
- Required Privilege Level | 11577
- Release Information | 11578

Syntax

```
no-send-pads-ac-info;
```

Hierarchy Level

```
[edit protocols pppoe]
```

Description

Prevent the router from sending the AC-Name and AC-Cookie tags in the PPPoE Active Discovery Session (PADS) packet. When you configure this statement, it affects PADS packets sent on all PPPoE interfaces configured on the router after the command is issued; it has no effect on previously created PPPoE interfaces. By default, the AC-Name and AC-Cookie tags are transmitted in the PADS packet, along with the Service-Name, Host-Uniq, Relay-Session-Id, and PPP-Max-Payload tags.

NOTE: In Junos OS Release 12.1 and earlier, only the Service-Name, Host-Uniq, Relay-Session-Id, and PPP-Max-Payload tags are contained in the PADS packet by default. The AC-Name and AC-Cookie tags are not transmitted in the PADS packet by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Disabling the Sending of PPPoE Access Concentrator Tags in PADS Packets

no-send-pads-error

IN THIS SECTION

- Syntax | 11578
- Hierarchy Level | 11578
- Description | 11578
- Required Privilege Level | 11579
- Release Information | 11579

Syntax

```
no-send-pads-error;
```

Hierarchy Level

```
[edit protocols pppoe]
```

Description

Discard PADR messages to prevent transmission of PADS control packets with AC-System-Error tags.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Discarding PADR Messages to Accommodate Abnormal CPE Behavior

no-send-to-ue

IN THIS SECTION

- Syntax | 11579
- Hierarchy Level | 11579
- Description | 11580
- Default | 11580
- Required Privilege Level | 11580
- Release Information | 11580

Syntax

```
no-send-to-ue;
```

Hierarchy Level

```
[edit unified-edge pcef flow-description flow-identifier]
```

Description

Specify that signaling information about the service data flow (SDF) filter is not sent to the user equipment.

Default

By default, if this statement is not configured, signaling information about the SDF filter is sent to the user equipment.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Service Data Flow Filters

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

no-sequence-check

IN THIS SECTION

- Syntax | 11581
- Hierarchy Level | 11581
- Description | 11581
- Required Privilege Level | 11581
- Release Information | 11581

Syntax

```
no-sequence-check;
```

Hierarchy Level

```
[edit security flow tcp-session]
```

Description

Specify that the device does not check sequence numbers in TCP segments during stateful inspection. By default, the device monitors the sequence numbers in TCP segments. The device detects the window scale specified by source and destination hosts in a session and adjusts a window for an acceptable range of sequence numbers according to their specified parameters. The device then monitors the sequence numbers in packets sent between these hosts. If the device detects a sequence number outside this range, it drops the packet.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

no-snoop (DHCP Local Server and Relay Agent)

IN THIS SECTION

- Syntax | 11582
- Hierarchy Level | 11582
- Description | 11582
- Required Privilege Level | 11583
- Release Information | 11583

Syntax

```
no-snoop;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6]
```

Description

Disable DHCP snooping filters.

DHCP snooping provides DHCP security by identifying incoming DHCP packets. In the default DHCP snooping configuration, all traffic is snooped. You can optionally use the `forward-snooped-clients` statement to evaluate the snooped traffic and to determine if the traffic is forwarded or dropped, based on whether or not the interface is configured as part of a group.

> **NOTE**: Starting in Junos OS Release 22.4R1, 'no-snoop' (DHCP Local Server and Relay Agent) is supported for DHCPV6 on QFX10K platforms.
>
> The "no-snoop" option is not supported on the PTX10002-60C devices.

In both the default configuration and in configurations using the `forward-snooped-clients` statement, all DHCP traffic is forwarded from the hardware control plane to the routing plane of the routing instance to ensure that all DHCP packets are intercepted. In certain topologies, such as a Metropolitan Routing Ring topology, forwarding all DHCP traffic to the control plane can result in excessive traffic. The `no-snoop` configuration statement disables the snooping filter for DHCP traffic that can be directly forwarded on the hardware control plane, such as Layer 3 unicast packets with a valid route, causing those DHCP packets to bypass the slower routing plane.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R2.

RELATED DOCUMENTATION

Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent

no-source-packet-routing (Protocols OSPF)

IN THIS SECTION

- Syntax | 11584
- Hierarchy Level | 11584
- Description | 11584
- Required Privilege Level | 11584

Release Information | 11584

Syntax

```
no-source-packet-routing;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf backup-spf-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
backup-spf-options],
[edit protocols ospf backup-spf-options],
[edit routing-instances routing-instance-name protocols ospf backup-spf-options]
```

Description

Disables use of source packet routing node segment labels for computing backup paths for normal IPv4 OSPF prefixes and OSPF source packet routing node segments.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

OSPF Overview

no-split-detection

IN THIS SECTION

- Syntax | 11585
- Hierarchy Level | 11585
- Description | 11585
- Default | 11585
- Required Privilege Level | 11586
- Release Information | 11586

Syntax

```
no-split-detection;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Disable the split and merge feature in a Virtual Chassis or VCF configuration.

The split and merge feature is enabled by default when you initially set up a Virtual Chassis. If you have only two members in your Virtual Chassis, we strongly recommend that you configure `no-split-detection` to disable the split and merge feature. This feature makes sure both switches remain in the correct Virtual Chassis roles in the event of a Virtual Chassis split. If you expand the two-member Virtual Chassis later to add more members, delete the `no-split-detection` configuration item to re-enable the split and merge feature again.

Default

The split and merge feature is enabled.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

- *Understanding Split and Merge in a Virtual Chassis*
- *Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge*
- *Disabling Split and Merge in a Virtual Chassis*
- *Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge*

no-split-detection (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 11586
- Hierarchy Level | 11587
- Description | 11587
- Required Privilege Level | 11587
- Release Information | 11587

Syntax

```
no-split-detection;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

As part of the preprovisioned configuration for an MX Series Virtual Chassis, disable detection of a split in the Virtual Chassis configuration. By default, split detection in the Virtual Chassis is enabled. To maintain the Virtual Chassis configuration in the event of a failure of one of the two member routers, we recommend that you use the `no-split-detection` statement to disable split detection in Virtual Chassis configurations in which you think the backup router is more likely to fail than the link to the backup router.

> **BEST PRACTICE**: We recommend that you use the `no-split-detection` statement to disable split detection for a two-member MX Series Virtual Chassis configuration if you think the backup router is more likely to fail than the Virtual Chassis port links to the backup router. Configuring redundant Virtual Chassis ports on different line cards in each member router reduces the likelihood that all Virtual Chassis port interfaces to the backup router can fail.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

- *Configuring Preprovisioned Member Information for a Virtual Chassis*
- *Disabling Split Detection in a Virtual Chassis Configuration*
- *Split Detection Behavior in a Virtual Chassis*
- *Global Roles and Local Roles in a Virtual Chassis*
- *Example: Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis*

no-syn-check

IN THIS SECTION

- Syntax | 11588
- Hierarchy Level | 11588
- Description | 11588
- Required Privilege Level | 11588
- Release Information | 11588

Syntax

```
no-syn-check;
```

Hierarchy Level

```
[edit security flow tcp-session]
```

Description

Disable checking of the TCP SYN bit before creating a session. By default, the device checks that the SYN bit is set in the first packet of a session. If the bit is not set, the device drops the packet.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

no-syn-check-in-tunnel

IN THIS SECTION

- Syntax | 11589
- Hierarchy Level | 11589
- Description | 11589
- Required Privilege Level | 11589
- Release Information | 11590

Syntax

```
no-syn-check-in-tunnel;
```

Hierarchy Level

```
[edit security flow tcp-session]
```

Description

Disable checking of the TCP SYN bit before creating a session for tunneled packets. By default, the device checks that the SYN bit is set in the first packet of a VPN session. If the bit is not set, the device drops the packet.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

no-syslog

IN THIS SECTION

- Syntax | 11590
- Hierarchy Level | 11590
- Description | 11591
- Required Privilege Level | 11591
- Release Information | 11591

Syntax

```
no-syslog;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

Disable system logging of control protocol requests and responses. By default, these messages are logged.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring System Logging

no-syslog-generation

IN THIS SECTION

- Syntax | 11591
- Hierarchy Level | 11592
- Description | 11592
- Required Privilege Level | 11592
- Release Information | 11592

Syntax

```
no-syslog-generation;
```

Hierarchy Level

```
[edit services]
```

Description

Disable system log generation.

> **NOTE**: If this statement is not configured, `edit services` generates a system log with respective severity level for values not within the configured range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Inline Video Monitoring on MX Series Routers

alarms (Video Monitoring) | 2125

no-tag

IN THIS SECTION

- Syntax | 11593
- Hierarchy Level | 11593
- Description | 11593

Required Privilege Level | 11593
Release Information | 11593

Syntax

```
no-tag;
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name  output vlan (vlan-id | vlan-name)]
```

Description

Specify that remote mirroring packets are not tagged.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

Hierarchy level `[edit forwarding-options]` introduced in Junos OS Release 13.2X50-D10 (ELS).

RELATED DOCUMENTATION

> *Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches*

no-tag

IN THIS SECTION

- Syntax | 11594
- Hierarchy Level | 11594
- Description | 11594
- Required Privilege Level | 11594
- Release Information | 11595

Syntax

```
no-tag;
```

Hierarchy Level

```
[edit ethernet-switching-options analyzer name  output vlan (vlan-id | vlan-name)]
[edit forwarding-options port-mirroring family ethernet-switching output vlan (vlan-name | vlan-
id) ]
[edit forwarding-options port-mirroring instance instance-name family ethernet-switching output
vlan (vlan-name | vlan-id)]
```

Description

Specify that remote port-mirroring packets are not tagged with the tag of the output (analyzer) VLAN.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

Hierarchy `[edit forwarding-options port-mirroring family ethernet-switching output vlan]` introduced in Junos OS Release 13.2.

Hierarchy `[edit forwarding-options port-mirroring instance instance-name family ethernet-switching output vlan]` introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches

no-temperature-monitoring

IN THIS SECTION

- Syntax (PTX10004, PTX10008, PTX10001-36MR, and PTX100016 Devices) | **11595**
- Hierarchy Level | **11595**
- Description | **11596**
- Required Privilege Level | **11596**
- Release Information | **11596**

Syntax (PTX10004, PTX10008, PTX10001-36MR, and PTX100016 Devices)

```
no-temperature-monitoring;
```

Hierarchy Level

```
[edit chassis fpc name pic name port],
```

```
[edit chassis lcc name fpc name pic name port],
[edit chassis member name fpc name pic name port]
```

Description

Disable EM policy. Disable temperature monitoring to explicitly disable optics temperature sensors for EM policy monitoring and management, on specific WAN ports.

NOTE:

```
no-temperature-monitoringset system syslog file messages any anyshow log
messagesPICD_OPTICS_EMPOLICY_NOTIFICATION: et-1/0/29: Optics Temperature monitoring is disabled.
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

- delete chassis fpc | 5281
- fpc (Chassis) | 7501
- fpc (PTX Series Devices) | 7525
- show log | 21980

no-termination-request

IN THIS SECTION

- Syntax | 11597

- Hierarchy Level | 11597
- Description | 11597
- Required Privilege Level | 11597
- Release Information | 11597

Syntax

```
no-termination-request;
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options],
[edit interfaces lsq-fpc/pic/port lsq-failure-options]
```

Description

Inhibit PPP termination-request messages to the remote host if the primary circuit fails.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

Support at the `[edit interfaces interface-name ppp-options]` hierarchy level added in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring the Association between LSQ and SONET Interfaces

no-throttle (Resource Monitor)

IN THIS SECTION

- Syntax | 11598
- Hierarchy Level | 11598
- Description | 11598
- Options | 11599
- Required Privilege Level | 11599
- Release Information | 11599

Syntax

```
no-throttle;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Disable the throttling of subscriber services and sessions when the utilization of memory resources exceeds the threshold levels. For the subscriber service management, sessions are throttled based on line card resources, such as heap, counter memory, expansion memory, and firewall counter memory. If the resource has been used above a certain threshold, the subscribers and services are throttled to prevent the system from being overloaded and resulting in a breakdown. This feature gathers input from each of the line cards and transfers this statistical detail to the Routing Engine process using a well-known internal port. This information is scanned by the daemon on the Routine Engine and using the shared memory space built into the the session database, the existing active subscribers and sessions are throttled. You can configure the resource-monitoring capability on MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.

Options

no-throttle Disable the throttling of subscriber services and sessions when the utilization of memory resources exceeds the threshold levels. The throttling capability is enabled by default.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

no-tls-certificate-check

IN THIS SECTION

- Syntax | 11599
- Hierarchy Level | 11599
- Description | 11600
- Required Privilege Level | 11600
- Release Information | 11600

Syntax

```
no-tls-certificate-check;
```

Hierarchy Level

```
[edit access profile profile-name ldap-server ip-address]
```

Description

Specify validation of the server certificate not required. SRX Series Firewalls support an additional check on the Lightweight Directory Access Protocol (LDAP) server's certificate during the Transport Layer Security (TLS) handshake for LDAP authentication. If the validation of the server certificate is not required, you can use this option to ignore the validation and accept the certificate without checking. By default, this option is disabled.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Example: Configuring Integrated User Firewall on SRX Series

no-transit-statistics

IN THIS SECTION

- Syntax | 11601
- Hierarchy Level | 11601
- Description | 11601
- Required Privilege Level | 11601
- Release Information | 11601

Syntax

```
no-transit-statistics;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls statistics],
[edit protocols mpls statistics]
```

Description

(PTX Series only) Disables the collection of MPLS statistics for LSPs transiting the router.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2 .

RELATED DOCUMENTATION

Configuring MPLS to Gather Statistics

statistics (Protocols MPLS) | **28594**

no-translation

IN THIS SECTION

Syntax | **11602**

- Hierarchy Level | 11602
- Description | 11602
- Options | 11602
- Required Privilege Level | 11602
- Release Information | 11602

Syntax

```
no-translation;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then]
```

Description

Specify that traffic is not to be translated.

The `no-translation` statement is supported on MX series routers with MS-DPCs and on M Series routers with MS-100, MS-400, and MS-500 MultiServices PICS. The `no-translation` statement is supported on MX series routers with MS-MPCs and MS-MICs starting in Junos OS release 15.1R1.

Options

none

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

no-translation (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 11603
- Hierarchy Level | 11603
- Description | 11603
- Required Privilege Level | 11603
- Release Information | 11604

Syntax

```
no-translation;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Description

Disable port translation for NAT. By default, port translation is enabled for NAT.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

no-trap

IN THIS SECTION

- Syntax | 11604
- Hierarchy Level | 11604
- Description | 11604
- Options | 11605
- Required Privilege Level | 11605
- Release Information | 11605

Syntax

```
no-trap {
    mpls-lsp-traps;
    rfc-3812-traps;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls log-updown],
[edit protocols mpls log-updown]
```

Description

Prevent the transmission of SNMP traps.

Options

`mpls-lsp-traps`—Block the MPLS LSP traps defined in the**rfc-3812-traps**, but allows the **rfc3812.mib** traps.

`rfc-3812-traps`—Block the traps defined in the **rfc3812.mib**, but allows the MPLS LSP traps defined in the **jnx-mpls.mib**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.
The `mpls-lsp-traps` and `rfc-3812-traps` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

System Log Messages and SNMP Traps for MPLS

Junos OS Network Management Administration Guide for Routing Devices

traceoptions (Protocols MPLS) | **29699**

no-tunnel-services

IN THIS SECTION

- Syntax | **11606**
- Hierarchy Level | **11606**
- Description | **11606**
- Required Privilege Level | **11606**
- Release Information | **11607**

Syntax

```
no-tunnel-services;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols vpls static-vpls],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit protocols vpls static-vpls],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Configure VPLS on a router without a Tunnel Services PIC. Configuring the `no-tunnel-services` statement creates a label-switched interface (LSI) to provide VPLS functionality. An LSI MPLS label is used as the inner label for VPLS. This label maps to a VPLS routing instance. On the PE router, the LSI label is stripped and then mapped to a logical LSI interface. The Layer 2 Ethernet frame is then forwarded using the LSI interface to the correct VPLS routing instance.

NOTE: In VPLS documentation, the word *Router* in terms such as *PR Router* is used to refer to any device that provides routing functions.

NOTE: On MX Series routers, label-switched interfaces configured with the `no-tunnel-services` statement are not supported with GRE tunnels when the GRE interface resides on a DPC.

NOTE: Although visible in the CLI, the `no-tunnel-services` statement is not supported on DPC cards at the `[edit logical-systems logical-system-name protocols vpls static-vpls]` and the `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls]` hierarchy levels.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Support for static VPLS added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring VPLS Without a Tunnel Services PIC

Configuring Static Pseudowires for VPLS

Configuring EXP-Based Traffic Classification for VPLS

no-unknown-unicast

IN THIS SECTION

- Syntax | 11607
- Hierarchy Level | 11608
- Description | 11608
- Default | 11608
- Required Privilege Level | 11609
- Release Information | 11609

Syntax

```
no-unknown-unicast;
```

Hierarchy Level

- For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers):

```
[edit forwarding-options storm-control-profiles profile-name all]
```

- For platforms without ELS:

```
[edit ethernet-switching-options storm-control interface (all | interface-name)]
```

Description

Disable storm control for unknown unicast traffic for the specified interface or for all interfaces.

Default

- On EX2200, EX3200, EX3300, and EX4200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams. You can selectively disable storm control on broadcast, multicast, or unknown-unicast traffic.
- On EX4300 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. You can selectively disable storm control on any type of traffic.
- On EX4500 and EX8200 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. On EX8200 switches, you can selectively disable storm control on registered multicast traffic, on unregistered multicast traffic, or on both types of multicast traffic.
- On EX6200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams. You can selectively disable storm control for each type of traffic individually.
- On EX9200 switches—Storm control is not enabled by default.
- MX Series routers—Storm control is not enabled by default.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Hierarchy level `[edit forwarding-options]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

- *Configuring Autorecovery for Port Security Events*
- Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches
- *Example: Using Storm Control to Prevent Network (MX Routers)*
- *Enabling and Disabling Storm Control (non-ELS)*
- *Enabling and Disabling Storm Control (ELS)*

no-unregistered-multicast

IN THIS SECTION

- Syntax | **11610**
- Hierarchy Level | **11610**
- Description | **11610**
- Default | **11610**
- Required Privilege Level | **11611**
- Release Information | **11611**

Syntax

```
no-unregistered-multicast;
```

Hierarchy Level

- For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers):

  ```
  [edit forwarding-options storm-control-profiles profile-name all]
  ```

- For platforms without ELS:

  ```
  [edit ethernet-switching-options storm-control interface (all | interface-name)]
  ```

Description

(EX8200 switches only) Disable storm control for unregistered multicast traffic for the specified interface or for all interfaces.

(EX4300 and EX9200 switches only) Exclude storm control for unregistered multicast traffic from a storm control profile.

(MX Series routers) Exclude storm control for unregistered multicast traffic from a storm control profile.

Default

EX4300 and EX8200 switches—Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic. The default storm control level is 80 percent of the available bandwidth used by the combined applicable traffic streams.

EX9200 switches—Storm control is not enabled by default.

MX Series routers—Storm control is not enabled by default.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.3.

Hierarchy level `[edit forwarding-options]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

- no-multicast | 11512
- no-registered-multicast | 11562
- *Understanding Storm Control*
- *Understanding Storm Control*

no-unsolicited-ra (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 11611
- Hierarchy Level | 11612
- Description | 11612
- Required Privilege Level | 11612
- Release Information | 11612

Syntax

```
no-unsolicited-ra;
```

Hierarchy Level

```
[edit system services subscriber-management overrides]
```

Description

Disable the default transmission and periodic refresh of unsolicited Router Advertisement messages by the router when the subscriber interface is created, and at configured periodic intervals thereafter. When you include the `no-unsolicited-ra` statement, the router sends Router Advertisement messages and associated periodic refresh messages only when it receives a Router Solicitation message from the subscriber.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

no-uri-check

IN THIS SECTION

- Syntax | 11613
- Hierarchy Level | 11613
- Description | 11613
- Required Privilege Level | 11613

Release Information | 11613

Syntax

```
no-uri-check;
```

Hierarchy Level

```
[edit security utm default-configuration anti-virus scan-options]
```

Description

Do not perform Sophos antivirus Uniform Resource Identifier (URI) checking. URI checking is performed by analyzing HTTP traffic URI content against a remote Sophos database server to identify malware or malicious content. URI checking is on by default.

NOTE: Starting in Junos OS release 18.4R1, the URI checking is off by default.

You can enable Sophos antivirus URI checking with the `uri-check` statement.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

no-user-query (Services User Identification)

IN THIS SECTION

- Syntax | 11614
- Hierarchy Level | 11614
- Description | 11614
- Required Privilege Level | 11614
- Release Information | 11615

Syntax

```
no-user-query;
```

Hierarchy Level

```
[edit services user-identification authentication-source (Services User Identification ClearPass)
aruba-clearpass]
```

Description

Disable the integrated ClearPass authentication and enforcement user query function, if it is configured. You can use the no-user-query statement to turn off the user query function without having to delete the configuration.

The user query function allows the SRX Series Firewall to query the ClearPass webserver for authentication and identity information for an individual user whose information was not posted to the SRX Series Firewall by ClearPass.

Required Privilege Level

services	To view this statement in the configuration.
services-control	To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

no-validate

IN THIS SECTION

- Syntax | 11615
- Hierarchy Level | 11615
- Description | 11615
- Options | 11616
- Required Privilege Level | 11616
- Release Information | 11616

Syntax

```
no-validate policy-name;
```

Hierarchy Level

```
[edit protocols bgp group group-name family (inet | inet flow)],
[edit protocols bgp group group-name neighbor address family (inet | inet flow)],
[edit routing-instances routing-instance-name protocols bgp group group-name family (inet | inet
flow)],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (inet | inet flow)]
```

Description

When BGP is carrying flow-specification network layer reachability information (NLRI) messages, the `no-validate` statement omits the flow route validation procedure after packets are accepted by a policy.

The receiving BGP-enabled device accepts a flow route if it passes the following criteria:

- The originator of a flow route matches the originator of the best match unicast route for the destination address that is embedded in the route.
- There are no more specific unicast routes, when compared to the destination address of the flow route, for which the active route has been received from a different next-hop autonomous system.

The first criterion ensures that the filter is being advertised by the next-hop used by unicast forwarding for the destination address embedded in the flow route. For example, if a flow route is given as 10.1.1.1, proto=6, port=80, the receiving BGP-enabled device selects the more specific unicast route in the unicast routing table that matches the destination prefix 10.1.1.1/32. On a unicast routing table containing 10.1/16 and 10.1.1/24, the latter is chosen as the unicast route to compare against. Only the active unicast route entry is considered. This follows the concept that a flow route is valid if advertised by the originator of the best unicast route.

The second criterion addresses situations in which a given address block is allocated to different entities. Flows that resolve to a best-match unicast route that is an aggregate route are only accepted if they do not cover more specific routes that are being routed to different next-hop autonomous systems.

You can bypass the validation process and use your own specific import policy. To disable the validation procedure and use an import policy instead, include the `no-validate` statement in the configuration.

Flow routes configured for VPNs with family **inet-vpn** are not automatically validated, so the `no-validate` statement is not supported at the `[edit protocols bgp group group-name family inet-vpn]` hierarchy level. No validation is needed if the flow routes are configured locally between devices in a single AS.

Options

policy-name—Import policy to match NLRI messages.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Enabling BGP to Carry Flow-Specification Routes

Understanding BGP Flow Routes for Traffic Filtering

no-vlan-interface-name

IN THIS SECTION

- Syntax | 11617
- Hierarchy Level | 11617
- Description | 11618
- Required Privilege Level | 11619
- Release Information | 11619

Syntax

```
no-vlan-interface-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name (relay-agent-interface-id | relay-agent-
remote-id)],
[edit forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)],
[edit forwarding-options dhcp-relay group group-name relay-option-82 (circuit-id | remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay dhcpv6 (relay-agent-
interface-id | relay-agent-remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay ... relay-option-82
(circuit-id | remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6  (relay-agent-
interface-id | relay-agent-remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ... relay-option-82
(circuit-id | remote-id)],
```

```
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-18],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-37]
```

Description

When you do not want bridge domain or VLAN tag information, do not include the VLAN ID nor the VLAN interface name (the default) in the circuit or remote ID value in the DHCP option 82 information.

> **NOTE**: The `no-vlan-interface-name` statement is mutually exclusive with the `use-interface-description` and `use-vlan-id` statements.

When you configure the `no-vlan-interface-name` statement only, the format displays only the Layer 3 interface:

```
irb.subunit
```

> **NOTE**: The *subunit* is required and used to differentiate the interface for remote systems.

When you configure the `no-vlan-interface-name` and `use-interface-description` statements, the format displays the IRB interface description:

```
irb_descr
```

If you configure the `no-vlan-interface-name` and `use-interface-description` statements, and no description for the IRB interface is found, the format displays the IRB interface name:

```
irb.subunit
```

When you configure the `no-vlan-interface-name` and `include-irb-and-l2` statements, the format displays the Layer 2 logical interface name and the IRB interface name:

```
(fe | ge)-fpc/pic/port.subunit+irb.subunit
```

When you configure the `no-vlan-interface-name`, `include-irb-and-l2` and `use-interface-name` statements, the format displays the Layer 2 interface description and the IRB interface name:

```
l2_descr+irb.subunit
```

If you configure the `no-vlan-interface-name`, `include-irb-and-l2` and `use-interface-name` statements, and no description for the Layer 2 interface is found, the format displays the Layer 2 logical interface name and the IRB interface name:

```
(fe | ge)-fpc/pic/port.subunit+irb.subunit
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

> **NOTE**: The EX Series switches that support the `no-vlan-interface-name` statement are the EX4300, EX4600, and EX9200 switches.

RELATED DOCUMENTATION

Including a Textual Description in DHCP Options

Using DHCP Relay Agent Option 82 Information

Configuring DHCPv6 Relay Agent Options

no-vrf-advertise

IN THIS SECTION

- Syntax | 11620
- Hierarchy Level | 11620
- Description | 11620
- Required Privilege Level | 11621
- Release Information | 11621

Syntax

```
no-vrf-advertise;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Prevent advertising VPN routes from a VRF routing instance to remote PE routers.

Within a hub-and-spoke configuration, you can configure a PE router to not advertise VPN routes from the primary (hub) routing instance. Instead, these routes are advertised from the secondary (downstream) routing instance. You can do this without configuring routing table groups, by using the `no-vrf-advertise` statement.

NOTE: This statement does not prevent the exportation of VPN routes to other VRF routing instances on the same router by configuring the `[edit routing-options auto-export]` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Policies for the VRF Table on PE Routers in VPNs

Configuring Hub-and-Spoke VPN Topologies: One Interface

Limiting Routes to Be Advertised by an MVPN VRF Instance

vrf-advertise-selective | 30957

no-vrf-propagate-ttl

IN THIS SECTION

- Syntax | 11621
- Hierarchy Level | 11622
- Description | 11622
- Default | 11622
- Required Privilege Level | 11622
- Release Information | 11622

Syntax

```
no-vrf-propagate-ttl;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Disable normal time-to-live (TTL) decrementing in a VRF routing instance. You configure this statement once per routing instance, and it affects only RSVP-signaled or LDP-signaled LSPs in the routing instance. When this router acts as an ingress router for an LSP, it pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When the router acts as the penultimate router, it pops the MPLS header without writing the MPLS TTL into the IP packet.

Default

Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

> *Example: Diagnosing Networking Problems Related to Layer 3 VPNs by Disabling TTL Decrementing*
>
> Disabling Normal TTL Decrementing

no-walkup

IN THIS SECTION

- Syntax | 11623
- Hierarchy Level | 11623
- Description | 11623
- Default | 11623
- Required Privilege Level | 11624
- Release Information | 11624

Syntax

```
no-walkup;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options defaults route-filter],
[edit logical-systems logical-system-name policy-options policy-statement policy-statement-name
defaults route-filter]
```

Description

Override route filter walkup globally or locally for a particular policy statement. The walkup feature examines more than the longest match route filters in a policy statement term with more than one route filter, allowing consolidation of terms and a potential performance enhancement.

Default

By default, the policy statement performs the type of route filter processing that is enabled at the global level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

- walkup | 31019
- route-filter | 15861
- *Walkup for Route Filters Overview*
- *Configuring Walkup for Route Filters to Improve Operational Efficiency*
- *Example: Configuring Walkup for Route Filters Globally to Improve Operational Efficiency*
- *Example: Configuring Walkup for Route Filters Locally to Improve Operational Efficiency*

no-world-readable (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | 11625
- Hierarchy Level | 11625
- Description | 11625
- Default | 11625
- Options | 11625
- Required Privilege Level | 11625
- Release Information | 11625

Syntax

```
no-world-readable
```

Hierarchy Level

```
[edit services rtlog traceoptions file]
```

Description

Do not allow any user to read the log file. Use this option to revert to `no-world-readable` configuration from `world-readable` setting.

Default

By default, `no-world-readable` option is set. No user is allowed to read the log file.

Options

world-readable	Do not allow any user to read the log file

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

world-readable (Next Gen Services Global System Logging)

no-zero-suppression

IN THIS SECTION

- Syntax | 11626
- Hierarchy Level | 11626
- Description | 11626
- Required Privilege Level | 11627

Syntax

```
no-zero-suppression
```

Hierarchy Level

```
[edit services analytics zero-suppression]
```

Description

Disable zero suppression for gRPC-based sensors. When this feature is enabled, data for a sensor is sent to the collector if the sensor value is zero. You can expect to see all counters, even ones with a value of zero. All key value pair updates will be streamed to a collector without performing any zero suppression. The default behavior without zero suppression is that a device will suppress zero values from streamed data.

To enable zero suppression again (the default), delete the configuration by entering **#delete services analytics zero-suppression no-zero-suppression**.

Whenever this feature is set or deleted, any existing collector is disconnected.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
21.4R1	Statement introduced in Junos OS Release 21.4R1 for MX240, MX280, and MX960 operating with MPC10E-10-MRATE, MPC10E-10C-MRATE, MPC10E-15C-MRATE line cards and MX2010 and MX2020 operating with the MX2K-MPC11E line card.
21.1R1	Statement introduced in Junos Evolved OS Release 21.1R1 for PTX10003-80C and PTX10003-160C and PTX10008.
20.4R2	Statement introduced in Junos Evolved OS Release 20.4R2 for PTX10008.

no-dhcp-flood

IN THIS SECTION

- Syntax | 11627
- Hierarchy Level | 11628
- Description | 11628
- Required Privilege Level | 11628
- Release Information | 11628

Syntax

```
no-dhcp-flood;
```

Hierarchy Level

```
[edit interfaces name unit]
```

Description

In an EVPN-VXLAN fabric where DHCP relay is configured with integrated routing and bridging (IRB) interfaces on leaf devices, it is possible for DHCP Discover packets to be flooded across a VXLAN attached to an IRB. This scenario leads to multiple copies of DHCP packets flooded across the fabric. To counter this behavior, configure the `no-dhcp-flood` knob at the `[edit interfaces irb unit <unit-number>]` hierarchy.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1 and Junos OS Evolved Release 22.4R1 for the QFX5xxx Series switches, EX4100 Series switches, and EX4400 Series switches.

RELATED DOCUMENTATION

DHCP Relay Agent over EVPN-VXLAN

no-flood-to-multicast-router-interfaces (IGMP and MLD Snooping)

IN THIS SECTION

- Syntax | 11629
- Hierarchy Level | 11629

Description | 11629
Release Information | 11629

Syntax

```
no-flood-to-multicast-router-interfaces;
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan vlan-name no-flood-to-multicast-router-interfaces],
[edit protocols mld-snooping vlan vlan-name no-flood-to-multicast-router-interfaces],
[edit routing-instances routing-instance-name protocols igmp-snooping vlan vlan-name no-flood-to-
multicast-router-interfaces],
[edit routing-instances routing-instance-name protocols mld-snooping vlan vlan-name no-flood-to-
multicast-router-interfaces],
[edit bridge-domains bridge-domain-name protocols igmp-snooping no-flood-to-multicast-router-
interfaces],
[edit bridge-domains bridge-domain-name protocols mld-snooping no-flood-to-multicast-router-
interfaces],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
snooping no-flood-to-multicast-router-interfaces],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols mld-
snooping no-flood-to-multicast-router-interfaces]
```

Description

Limit the flooding of multicast streams to multicast-router interfaces on a per VLAN or bridge domain level. Based on the receipt of IGMP or MLD join messages from interested receivers from behind the multicast-router interface, only known multicast traffic is sent to the multicast-router interface resulting in improved bandwidth utilization.

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

no-incremental-update

IN THIS SECTION

- Syntax | 11630
- Hierarchy Level | 11630
- Description | 11630
- Required Privilege Level | 11630
- Release Information | 11631

Syntax

```
no-incremental-update
```

Hierarchy Level

```
[edit firewall family filter]
```

Description

Prevent the filter from undergoing incremental update. The filter undergoes MBB.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement released in Junos OS Evolved Release 24.1

node (Chassis Cluster Redundancy Group)

IN THIS SECTION

- Syntax | 11631
- Hierarchy Level | 11631
- Description | 11631
- Options | 11632
- Required Privilege Level | 11632
- Release Information | 11632

Syntax

```
node (0 | 1 ) {
    priority number;
}
```

Hierarchy Level

```
[edit chassis cluster redundancy-group          group-number        ]
```

Description

Identify each cluster node in a redundancy group and set its relative priority for primary role.

A redundancy group is collection of objects. A redundancy group contains objects on both nodes. A redundancy group is primary on one node and backup on the other at any time. When a redundancy group is said to be primary on a node, its objects on that node are active. The primacy of a redundancy

group is dependent on the priority configured for the node, the node ID (in case of tied priorities), and the order in which the node comes up.

Options

node

Cluster node number, You can set with the `set chassis cluster node` *`node-number`* statement.

- Values:
  - `0`—Node identifier 0
  - `1`—Node identifier 1

priority ***number***

Priority value of the node. Each node is given a priority within a redundancy group. The eligible node with the highest priority is elected primary. Initiating a failover with the `request chassis cluster failover node` or `request chassis cluster failover redundancy-group` command overrides the priority settings.

- **Range:** `1-254`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

redundancy-group (Chassis Cluster) | **14173**

node-detection (EVPN-VXLAN)

IN THIS SECTION

- Syntax | 11633
- Hierarchy Level | 11634
- Description | 11634
- Options | 11634
- Required Privilege Level | 11635
- Release Information | 11635

Syntax

```
node-detection {
  action (laser-off | trigger-node-isolation);
  bfd-liveness-detection {
   no-adaptation;
   detection-time {
    threshold milliseconds;
   }
   inline-disable;
   minimum-interval milliseconds;
   minimum-receive-interval milliseconds;
   multiplier multiplier;
   neighbor neighbor;
   pdu-size pdu-size;
   transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
   }
   version(0 | 1 | automatic);
  }
  next-hop next-hop interface;
 }
```

Hierarchy Level

```
[edit logical-systems name protocols evpn],
[edit protocols evpn]
```

Description

Configure node detection to enable "bfd-liveness-detection" on page 2923 in a multihomed environment. Node detection monitors the liveness of EVPN BGP peers using Bidirectional Forwarding Detection (BFD) protocol. Configure one of the peers to take `action` to forward traffic through the network based on the status of the BGP and BFD sessions. The two options for the `action` configuration are turning the interface lasers off and triggering node isolation. When the `action` is turning the lasers off, the lasers are off for the ESI LAG interfaces during core isolation. The trigger node isolation `action` puts the ESI LAG interfaces in a link-down state.

Each option for the `action` configuration also requires additional configuration.

- The `laser-off` action requires that "asynchronous-notification" on page 2522 be configured on the ESI-LAG interfaces.
- The `trigger-node-isolation` action requires a "network-isolation" on page 11312 profile and a "service-tracking-action link-down" on page 16758 configuration that brings the interface down on core-isolation.

The BFD session can run over standard interfaces or over the management interfaces between the EVPN BGP peers.

NOTE: Junos OS Evolved platforms do not support BFD over management ports.

Options

action Configure the isolation action only on one device and not both. The isolation action determines how to implement the core isolation.

- laser-off—Will trigger laser-off on core-isolation
- trigger-node-isolation—Will trigger link-down on core-isolation

next-hop Next-hop interface connected to the node detection peer

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.3R1 and Junos OS Evolved Release 23.3R1.

RELATED DOCUMENTATION

Backup Liveness Detection on EVPN Dual Homing

node-device (Aliases)

IN THIS SECTION

- Syntax | 11635
- Hierarchy Level | 11636
- Description | 11636
- Options | 11636
- Required Privilege Level | 11636
- Release Information | 11636

Syntax

```
node-device node-device-name {
    assigned-node-device-name;
}
```

Hierarchy Level

```
[edit fabric aliases]
```

Description

(QFabric systems only) Specify the mapping of user-defined names to QFabric system Node devices.

NOTE: The following rules apply to QFabric component alias naming:

- Alias names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
- The maximum length of an alias name is 30 characters.
- Alias names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.
- You cannot use the reserved names `all`, `fabric`, or `director-group` as an alias name.

Options

node-device-name	Specify a user-defined name for a QFabric system Node device.
assigned-node-device-name	Specify the Node device identifier or name that has been provided. Identifiers are usually autogenerated by the Director software, and names are usually provided by the administrator of the default partition.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Aliases for the QFabric System

Understanding Node Devices

node-device (Chassis)

IN THIS SECTION

- Syntax | 11637
- Hierarchy Level | 11638
- Description | 11638
- Required Privilege Level | 11638
- Release Information | 11638

Syntax

```
node-device name {
    fibre-channel {
        port-range {
            port-range-low port-range-high;
        }
    }
    pic pic-number {
        fte {
            port port-number;
            port-range port-range-low port-range-high;
        }
        xe {
            port port-number;
            port-range port-range-low port-range-high;
        }
    }
}
```

Hierarchy Level

```
[edit chassis node-group]
```

Description

Configure properties specific to a Node device in a QFabric system.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Link Aggregation

node-device (Flow Groups)

IN THIS SECTION

- Syntax | 11639
- Hierarchy Level | 11639
- Description | 11639
- Options | 11639
- Required Privilege Level | 11639
- Release Information | 11640

Syntax

```
node-device (node-device-name | serial-ID | alias );
```

Hierarchy Level

```
[edit fabric flow-groups flow-group-name]
```

Description

In a QFabric system, a hash function is used to select an Interconnect device to forward traffic between two Node devices. Since this hash function is performed on all available Interconnect devices, it is possible for redundant multicast streams to flow through one Interconnect device, making that Interconnect device a potential single point of failure for the redundant flows. You can enforce the use of specific Interconnect devices by using the QFabric system `flow-groups` statement.

NOTE: A Node device can belong to only one flow group.

```
root@qfabric# set fabric flow-groups test1 node-device ?
Possible completions:
 <node-device-name> Serial ID or alias of node-device
P3686-C                 Node device
TA3713480203      Node device
TA3713500165     Node device
TA3713500185     Node device
TA3713500307     Node device
```

Options

node-device-name **(or *serial-ID* or *alias*)** selects the indicated ingress Node device switch for a flow group

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D15.

RELATED DOCUMENTATION

flow-groups | 7291

Segregating QFabric Traffic Flows With Flow Groups

Example: Creating a QFabric Flow Group

Understanding QFabric Multicast Flow Groups

node-device (Resources)

IN THIS SECTION

- Syntax | 11640
- Hierarchy Level | 11640
- Description | 11641
- Required Privilege Level | 11641
- Release Information | 11641

Syntax

```
node-device node-device-name;
```

Hierarchy Level

```
[edit fabric resources node-group node-group-name]
```

Description

(QFabric systems only) Assign Node devices to Node groups within a QFabric system.

> **NOTE**: Ensure that the Node devices you assign to a redundant server Node group are of the same type, either two QFX3500 Node devices, two QFX3600 Node devices, or two QFX5100 Node devices. You cannot mix and match different Node device types in the same redundant server Node group.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Node Groups for the QFabric System

Understanding Node Groups

node-group (Chassis)

IN THIS SECTION

- Syntax | 11642
- Hierarchy Level | 11643
- Description | 11643
- Required Privilege Level | 11643
- Release Information | 11643

Syntax

```
node-group name {
    aggregated-devices {
        ethernet {
            device-count number;
        }
    }
    alarm {
        interface-type {
            link-down (ignore | red | yellow);
        }
    }
    container-devices {
        device-count number;
    }
    node-device name {
        fibre-channel {
            port-range {
                port-range-low port-range-high;
            }
        }
    }
    pic pic-number {
        fte {
            port port-number;
            port-range port-range-low port-range-high;
        }
        xe {
            port port-number;
            port-range port-range-low port-range-high;
        }
    }
    routing-engine {
        on-disk-failure {
            disk-failure-action (halt | reboot);
        }
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure properties specific to a Node group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Link Aggregation

node-group (OxID Hash Control)

IN THIS SECTION

- Syntax | 11644
- Hierarchy Level | 11644
- Description | 11644
- Options | 11644
- Required Privilege Level | 11644
- Release Information | 11644

Syntax

```
node-group (node-group-name | all) {
    oxid (enable | disable);
}
```

Hierarchy Level

```
[edit forwarding-options hash-key family fcoe ethernet-interfaces]
[edit forwarding-options hash-key family fcoe fabric-interfaces]
```

Description

Specify a QFabric system Node group on which to enable or disable OxID hash control. OxID hash control is enabled or disabled on the fabric ports or on the Ethernet (FCoE) LAG ports that face an FCoE forwarder (FCF).

Options

node-group Name of the Node group on which you want to enable OxID hash control.

all All Node groups on the QFabric system (OxID hash control will be enabled or disabled on all Node groups).

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X52-D10.

RELATED DOCUMENTATION

Enabling and Disabling CoS OxID Hash Control for FCoE Traffic on QFabric Systems

Understanding OxID Hash Control for FCoE Traffic Load Balancing on QFabric Systems

node-group (Resources)

IN THIS SECTION

- Syntax | 11645
- Hierarchy Level | 11645
- Description | 11645
- Required Privilege Level | 11646
- Release Information | 11646

Syntax

```
node-group node-group-name {
    network-domain;
    node-device node-device-name;
}
```

Hierarchy Level

```
[edit fabric resources]
```

Description

(QFabric systems only) Define Node groups within a QFabric system.

NOTE: The following rules apply to QFabric Node group naming:

- The network Node group must use the predefined name *NW-NG-0*. You must use this name when adding Node devices to the network Node group. You cannot specify a different name. Also, you can configure only one network Node group per partition.
- Node group names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters.
- The maximum length of a Node group name is 30 characters.
- Node group names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components.
- You cannot use the reserved names `all`, `fabric`, or `director-group` as a Node group name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Node Groups for the QFabric System

Understanding Node Groups

node-hello

IN THIS SECTION

Syntax | 11647

Hierarchy Level | 11647

- Description | 11647
- Required Privilege Level | 11647
- Release Information | 11648

Syntax

```
(node-hello | no-node-hello);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Enable node-ID based RSVP hellos globally on all of the RSVP interfaces on the router to allow Juniper Networks routers to interoperate with the equipment of other vendors. By default, node-ID based RSVP hellos are enabled. To disable node-id based hello sessions, use the `no-node-hello` statement. If you have disabled RSVP node IDs on the router, Junos OS does not accept any node-ID hello packets.

NOTE: For Junos OS Release 15.1 or earlier, node-ID based hellos are disabled by default. To enable node-id based hello sessions on routers running Junos OS Release 15.1 or earlier, use the **node-hello** statement.

NOTE: If `link-protection` is enabled, remote node hellos that are initiated by the Point of Local Repair (PLR) to Node Protecting Merge Point (NP-MP) are enabled automatically. However, if `no-enhanced-frr-bypass` is enabled (that is, enhanced FRR is off), remote node hellos are automatically disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring RSVP Node-ID Hellos

no-enhanced-frr-bypass (Protocols RSVP) | **11451**

node-id

IN THIS SECTION

- Syntax | **11648**
- Hierarchy Level | **11648**
- Description | **11648**
- Required Privilege Level | **11649**
- Release Information | **11649**

Syntax

```
node-id mac-address;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

For EX Series switches and QFX Series switches, node-id is not configurable.

For MX Series routers, optionally specify the MAC address of a node in the protection group. If this statement is not included, the router assigns the node's MAC address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Example: Configuring Ethernet Ring Protection Switching on EX Series Switches

Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS

node-link-protection (Protocols MPLS)

IN THIS SECTION

- Syntax | 11650
- Hierarchy Level | 11650
- Description | 11650
- Default | 11650
- Required Privilege Level | 11650
- Release Information | 11650

Syntax

```
node-link-protection;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Enable node and link protection on the specified LSP. To fully enable node and link protection, you also need to include the `link-protection` statement at the `[edit protocols rsvp interface interface-name]` hierarchy level.

Default

Node and link protection is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

RELATED DOCUMENTATION

Configuring Node Protection or Link Protection for LSPs

MPLS Feature Support on QFX Series and EX4600 Switches

Understanding Interprovider and Carrier-of-Carriers VPNs

node-protection

IN THIS SECTION

- Syntax | 11651
- Hierarchy Level | 11651
- Description | 11651
- Options | 11651
- Required Privilege Level | 11652
- Release Information | 11652

Syntax

```
node-protection cost cost;
```

Hierarchy Level

```
[edit logical-systems name protocols isis interface name level number post-convergence-lfa],
[edit logical-systems name routing-instances name protocols isis interface name level number post-
convergence-lfa],
[edit protocols isis interface name level number post-convergence-lfa],
[edit routing-instances name protocols isis interface name level  number post-convergence-lfa]
```

Description

Enable node protection mode for topology-independent loop-free alternate (TI-LFA) routes for IS-IS. You can configure the cost of all the links used for calculating the TI-LFA post-convergence failure path cost. If node protection is enabled without configuring a cost value, then the cost is set to the maximum cost.

Options

cost Cost for node protection.

- **Range:** `0` through `16,777,215`

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 17.4R1 for MX Series, PTX Series, and QFX Series.

RELATED DOCUMENTATION

Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS

Configuring Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS

post-convergence-lfa | **12903**

use-post-convergence-lfa | **30531**

use-for-post-convergence-lfa | **30517**

node-protection (Static LSP)

IN THIS SECTION

- Syntax | **11653**
- Hierarchy Level | **11653**
- Description | **11653**
- Default | **11653**
- Options | **11653**
- Required Privilege Level | **11653**
- Release Information | **11653**

Syntax

```
node-protection bypass-name name next-next-label label;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
transit incoming-label],
[edit protocols mpls static-label-switched-path lsp-name ingress],
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Description

Enable node protection on the specified static bypass LSP. Node protection ensures that traffic from an LSP traversing a neighboring router can continue to reach its destination even if the neighboring router fails.

Default

Node protection is disabled.

Options

`bypass-name name`—Bypass LSP name.

`next-next-label label`—Bypass LSP name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.1.

RELATED DOCUMENTATION

Configuring Static LSPs

node-segment (Protocols OSPF)

IN THIS SECTION

- Syntax | 11654
- Hierarchy Level | 11654
- Description | 11655
- Options | 11655
- Required Privilege Level | 11655
- Release Information | 11655

Syntax

```
node-segment {
    ipv4-index index;
    index-range index range;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf source-packet-routing],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
source-packet-routing],
[edit protocols ospf source-packet-routing],
[edit routing-instances routing-instance-name protocols ospf source-packet-routing]
```

Description

Enable source packet routing in networking (SPRING) at all levels. SPRING, or segment routing, is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to determine the actual path it should take.

> **NOTE**: You can provision an IPv4 node segment index for a routing instance, not for a specific OSPF area. A node segment index is attached to the IPv4 router-id, if the router-ids are configured on the loopback interface. Otherwise, the lowest IP address on the loopback interface is chosen to attach the node segment identifier..

Options

index-range ***index range***

Range of node segment indices allowed.

- **Default:** 4096
- **Range:** 32 through 4096

ipv4-index ***index***

IPv4 node segment index.

- **Range:** 0 through 4095

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

OSPF Overview

node (System)

IN THIS SECTION

- Syntax | 11656
- Hierarchy Level | 11656
- Description | 11656
- Options | 11656
- Required Privilege Level | 11656
- Release Information | 11657

Syntax

```
node {
    offline node-name;
}
```

Hierarchy Level

```
[edit system]
```

Description

Specify a node by name and set node attributes. A node is a computational resource: it has a general-purpose CPU and can be used to do useful work. An attribute is a specific detail about a node.

Options

offline *node-name* Configure a particular node to be offline.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 18.4R1.

RELATED DOCUMENTATION

show system nodes | 26939

non-active-hold-priority

IN THIS SECTION

- Syntax | 11657
- Hierarchy Level | 11658
- Release Information | 11658
- Description | 11658
- Default | 11658
- Range | 11658
- Required Privilege Level | 11658

Syntax

```
non-active-hold-priority non-active-hold-priority-value;
```

Hierarchy Level

```
[edit logical-systems name protocols mpls label-switched-path name secondary name non-active-hold-priority],
[edit logical-systems name routing-instances name protocols mpls label-switched-path name secondary name non-active-hold-priority],
[edit logical-systems name tenants name routing-instances name protocols mpls label-switched-path name secondary name non-active-hold-priority],
[edit protocols mpls label-switched-path name secondary name non-active-hold-priority],
[edit routing-instances name protocols mpls label-switched-path name secondary name non-active-hold-priority],
[edit tenants name routing-instances name protocols mpls label-switched-path name secondary name non-active-hold-priority]
```

Release Information

Statement introduced in Junos OS Release 21.2R1.

Description

Configure the non-active hold priority value for secondary standby path.

Default

There is no default value for non-active-hold-priority.

Range

0 through 7, where 0 is the highest and 7 is the lowest priority.

Required Privilege Level

routing

non-lsp-next-hop

IN THIS SECTION

- Syntax | 11659
- Hierarchy Level | 11659
- Description | 11659
- Required Privilege Level | 11659
- Release Information | 11659

Syntax

```
non-lsp-next-hop;
```

Hierarchy Level

```
[edit class-of-service forwarding-policy next-hop-map map-name forwarding-class class-name]
```

Description

Use a non-LSP next hop for traffic sent to this forwarding class next-hop map of this forwarding policy.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring CoS-Based Forwarding

non-revertive

IN THIS SECTION

- Syntax | 11660
- Hierarchy Level | 11660
- Description | 11660
- Required Privilege Level | 11660
- Release Information | 11661

Syntax

```
non-revertive;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Enable nonrevertive operation where traffic is allowed to use the RPL if it has not failed, even after a switch condition has cleared. The default mode of operation is revertive.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

non-revertive (Chassis)

IN THIS SECTION

- Syntax | 11661
- Hierarchy Level | 11661
- Description | 11661
- Required Privilege Level | 11662
- Release Information | 11662

Syntax

```
non-revertive;
```

Hierarchy Level

```
[edit chassis aggregated-devices ethernet lacp link-protection]
```

Description

Disable the ability to switch to a better priority link (if one is available) once a link is established as active and a collection or distribution is enabled.

BEST PRACTICE: (MX Series) By default, Link Aggregation Control Protocol link protection is revertive. This means that after the current link becomes active, the router switches to a higher-priority link if one becomes operational or is added to the aggregated Ethernet bundle. In a highly scaled configuration over aggregated Ethernet, we recommend that you prevent the router from performing such a switch by including the `non-revertive` statement. Failure to do so may result in some traffic loss if a MIC on which a member interface is located reboots. Using the `non-revertive` statement for this purpose is not effective if both the primary and secondary interfaces are on the MIC that reboots.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Junos OS for Supporting Aggregated Devices

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

non-revertive (Interfaces)

IN THIS SECTION

- Syntax | **11663**
- Hierarchy Level | **11663**
- Description | **11663**
- Required Privilege Level | **11663**
- Release Information | **11663**

Syntax

```
non-revertive;
```

Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
```

Description

Disable the ability to switch to a better priority link (if one is available) once a link is established as active and collection distribution is enabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

link-protection | **9803**

Configuring Aggregated Ethernet Link Protection

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

non-strict-priority-scheduling

IN THIS SECTION

- Syntax | 11664
- Hierarchy Level | 11664
- Description | 11664
- Required Privilege Level | 11664
- Release Information | 11664

Syntax

```
non-strict-priority-scheduling;
```

Hierarchy Level

```
[edit class-of-service],
[edit dynamic-profiles name class-of-service]
```

Description

Configure non-strict priority scheduling to avoid starvation of lower-priority queues on SRX300, SRX320, SRX340, SRX345, SRX1500, SRX1600, SRX550M, and vSRX Virtual Firewall 2.0 devices.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D80.

NOTE: This statement is supported only on SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX1600, and vSRX2.0 firewalls.

RELATED DOCUMENTATION

Example: Configuring CoS Non-Strict Priority Scheduling

non-subscriber-no-reply

IN THIS SECTION

- Syntax | 11665
- Hierarchy Level | 11665
- Description | 11666
- Options | 11666
- Required Privilege Level | 11666
- Release Information | 11666

Syntax

```
non-subscriber-no-reply;
```

Hierarchy Level

```
[edit system arp]
```

Description

Enable this option to drop ARP requests from non-subscribers when a user route is dynamically added for a subscriber. Configuring this statement suppresses the ARP response from the kernel when there is an ARP request for a loopback interface from static DHCP subscribers using a common LAN segment between two devices. However, this configuration might not be effective if the subscriber configuration has suppressed either a destination Layer 2 route or an access Layer 3 route.

Options

This command has no options.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3R9.

RELATED DOCUMENTATION

autoinstallation | 2753

non-vc-mode

IN THIS SECTION

- Syntax | 11667
- Hierarchy Level | 11667
- Description | 11667
- Required Privilege Level | 11667
- Release Information | 11667

Syntax

```
non-vc-mode;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Configure a node on the sub-ring to operate in non-virtual channel mode. If this option is enabled then all the nodes in the sub-ring are configured with this option. Also, the `non-vc-mode` option should be used with care and only for open rings. Using this option for closed rings creates loops for RAPS control messages.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

> *Ethernet Ring Protection Switching Overview*
>
> Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

nonpersistent

IN THIS SECTION

- Syntax | 11668
- Hierarchy Level | 11668
- Description | 11668
- Required Privilege Level | 11668
- Release Information | 11669

Syntax

```
nonpersistent;
```

Hierarchy Level

```
[edit accounting-options file filename]
```

Description

Store log files used for accounting data in the **mfs/var/log** directory (located on DRAM) instead of the **cf/var/log** directory (located on the compact flash drive). This feature is useful for minimizing read/write traffic on the router's compact flash drive.

NOTE: If log files for accounting data are stored on DRAM, these files are lost when you reboot the router. Therefore, you should back up these files periodically.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring Accounting Options, Source Class Usage and Destination Class Usage Options

nonstop-bridging

IN THIS SECTION

- Syntax | 11669
- Hierarchy Level | 11669
- Description | 11669
- Required Privilege Level | 11670
- Release Information | 11670

Syntax

```
nonstop-bridging;
```

Hierarchy Level

```
[edit protocols layer2-control]
```

Description

For platforms with two Routing Engines, configure a primary Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 Control Protocol (L2CP) information.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Synchronizing the Routing Engine Configuration

Configuring Nonstop Bridging

Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)

Configuring Nonstop Bridging on Switches (CLI Procedure)

nonstop-bridging (Ethernet Switching)

IN THIS SECTION

- Syntax | 11670
- Hierarchy Level | 11671
- Description | 11671
- Required Privilege Level | 11671
- Release Information | 11671

Syntax

```
nonstop-bridging;
```

Hierarchy Level

```
[edit ethernet-switching-options]
```

Description

For switches with two Routing Engines or for Virtual Chassis, configure a primary Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 protocol information for the Layer 2 protocols that support nonstop bridging (NSB). For a list of the EX Series switches and Layer 2 protocols that support nonstop bridging, see EX Series Switch Software Features Overview.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)

nonstop-routing

IN THIS SECTION

- Syntax | **11672**
- Hierarchy Level | **11672**
- Description | **11672**
- Default | **11672**
- Required Privilege Level | **11672**

Release Information | 11672

Syntax

```
nonstop-routing;
```

Hierarchy Level

```
[edit routing-options]
```

NOTE: Although `nonstop-routing` is also a valid keyword at the `logical-systems` hierarchy level, it is not supported.

Description

For routing platforms with two Routing Engines, configure a primary Routing Engine to switch over gracefully to a backup Routing Engine and to preserve routing protocol information.

Default

disabled

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring Nonstop Active Routing

nonstop-routing-options

IN THIS SECTION

- Syntax | 11673
- Hierarchy Level | 11673
- Description | 11673
- Options | 11674
- Required Privilege Level | 11674
- Release Information | 11674

Syntax

```
nonstop-routing-options {
    precision-timers-max-period precision-timers-max-period;
}
```

Hierarchy Level

```
[edit routing-options]
```

Description

For routing platforms with two Routing Engines, a primary Routing Engine is configured to switch over gracefully to a backup Routing Engine. This allows the routing protocol information to be preserved even after failover. Support of precision-timers in the kernel is a feature where the kernel takes over autogeneration of BGP keepalives right after the switchover from backup to primary event occurs. The kernel in the Routing Engine continues this autogeneration until RPD is able to take over the session or until a maximum period has elapsed since the switchover event occurred.

> **NOTE**: This maximum period configuration applies only when at least one client protocol such as BGP registers for the automatic keepalive generation service provided by the kernel, and the kernel timer generates control plane session keepalives on behalf of that protocol after a switchover event.

Options

precision-timers-max-period ***precision-timers-max-period***

The maximum period for which the kernel auto generates keepalives on behalf of BGP after a switchover event from backup to primary.

> **NOTE**: You can verify the `precision-timers-max-period` using "show nonstop-routing" on page 22571 command.

- **Default:** 600 seconds
- **Range:** 60 seconds to 1800 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.2.

RELATED DOCUMENTATION

> *Understanding Maximum Period Configuration for Automatic Generation of BGP Keepalives by Kernel Timers After Switchover*

nonvolatile

IN THIS SECTION

- Syntax | 11675
- Hierarchy Level | 11675
- Description | 11675
- Required Privilege Level | 11675
- Release Information | 11676

Syntax

```
nonvolatile {
    commit-delay seconds;
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Configure options for SNMP **Set** requests.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The `commit-delay` statement introduced in Junos OS Release 9.0 for EX Series switches.

RELATED DOCUMENTATION

commit-delay | 4446

normalization

IN THIS SECTION

- Syntax | 11676
- Hierarchy Level | 11677
- Description | 11677
- Options | 11677
- Required Privilege Level | 11678
- Release Information | 11678

Syntax

```
normalization {
    failover-normalization;
    no-incremental-normalize;
    normalization-retry-duration seconds;
    normalization-retry-limits number;
    normalize-interval seconds;
}
```

Hierarchy Level

```
[edit protocols mpls container-label-switched-path lsp-name splitting-merging]
```

Description

Perform normalization.

Options

failover-normalization

Enable the ingress router to pro-actively normalize or re-distribute traffic when a link or node failure happens on a member LSP. A member LSP can go down between two scheduled normalization events because of a link-failure or pre-emption.

- **Default:** Disabled

no-incremental-normalize

Disables automatic switchover by the ingress router to a new instance of the container LSP until the desired demand is satisfied, although the given number of LSPs can be successfully signaled such that the new aggregate bandwidth value exceeds the old aggregate bandwidth value.

- **Default:** False (disabled)

normalization-retry-duration *seconds*

Specifies the duration before which the ingress router performs a normalization reattempt when the previous normalization has not been successful. Normalization is done until a sufficient number of LSPs come up with an aggregate bandwidth that is more than the current aggregate or desired bandwidth.

- **Default:** 30 seconds

normalization-retry-limits *number*

Specifies the maximum number of times the ingress router performs normalization reattempts until a sufficient number of LSPs come up successfully with new bandwidth values.

- **Default:** 1

normalize-interval *seconds*

Specifies the duration between two normalization events.

- **Range:** 300 seconds through 315360000 seconds
- **Default:** 86400 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

splitting-merging | 28391

normalization (Protocols EVPN)

IN THIS SECTION

- Syntax | 11678
- Hierarchy Level | 11679
- Description | 11679
- Default | 11679
- Required Privilege Level | 11679
- Release Information | 11679

Syntax

```
normalization;
```

Hierarchy Level

```
[edit routing-instances name protocols evpn]
```

Description

Enable normalization, which translates ingress or egress VLANs to another VLAN ID.

Default

Routing instances configure as `instance-type evpn` have normalization enabled by default.

Routing instances configured as `instance-type mac-vrf` have normalization disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 21.3R1

RELATED DOCUMENTATION

no-normalization | 11526

not (Event Policy)

IN THIS SECTION

- Syntax | 11680
- Hierarchy Level | 11680
- Description | 11680
- Options | 11680
- Required Privilege Level | 11680
- Release Information | 11681

Syntax

```
not events [ events ];
```

Hierarchy Level

```
[edit event-options policy policy-name within seconds]
```

Description

Configure an event policy to trigger only if the events in the `not` statement do not occur within a specified time interval before one of the trigger events.

Options

events [*events*] List of events. Events can be generated by Junos OS processes, or they can be internally generated.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Use Correlated Events to Trigger an Event Policy

not-on-disk-underperform

IN THIS SECTION

- Syntax | 11681
- Hierarchy Level | 11681
- Description | 11681
- Required Privilege Level | 11682
- Release Information | 11682

Syntax

```
not-on-disk-underperform;
```

Hierarchy Level

```
[edit chassis redundancy failover]
```

Description

Prevent gstatd from causing failovers in dual Routing Engines set for graceful Routing Engine switchover (GRES). The gstatd log message is still generated. This is an optional configuration.

> **NOTE**: Configure the `disk-write-threshold` and `disk-read-threshold` statements to customize the gstatd timeout threshold.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3R6.

RELATED DOCUMENTATION

Preventing Graceful Routing Engine Switchover in the Case of Slow Disks

not-present (DHCP Relay Agent)

IN THIS SECTION

- Syntax | **11682**
- Hierarchy Level | **11683**
- Description | **11685**
- Required Privilege Level | **11685**
- Release Information | **11685**

Syntax

```
not-present {
    drop drop;
    forward-only forward-only;
```

```
    local-server-group local-server-group;
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-77],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-77],
[edit forwarding-options dhcp-relay group name relay-option],
[edit forwarding-options dhcp-relay group name relay-option option-60],
[edit forwarding-options dhcp-relay group name relay-option option-77],
[edit forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay relay-option option-60],
[edit forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-77],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option option-60],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option option-77],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
```

```
relay relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option option-60],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option option-77],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option option-60],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option option-77],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option
option-60],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option
option-77],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option option-60],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option option-77],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option option-60],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option option-77],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option
option-60],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option
option-77],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option option-77],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
```

```
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-77],
[edit routing-instances name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-77],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-60],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-77],
[edit vlans name forwarding-options dhcp-relay relay-option],
[edit vlans name forwarding-options dhcp-relay relay-option option-60],
[edit vlans name forwarding-options dhcp-relay relay-option option-77]
```

Description

Option not present action

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address

notification

IN THIS SECTION

- Syntax | 11686
- Hierarchy Level | 11686
- Description | 11687
- Options | 11687
- Required Privilege Level | 11687
- Release Information | 11687

Syntax

```
notification {
    log-attacks {
        alert;
    }
    packet-log (Security IDP Policy) {
        post-attack number;
        post-attack-timeout seconds;
        pre-attack number;
    }
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name then]
```

Description

Configure the logging options against the action. When attacks are detected, you can choose to log an attack and create log records with attack information and send that information to the log server.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. Added packet capture support in Junos OS Release 10.2.

notification (Commit)

IN THIS SECTION

- Syntax | 11687
- Hierarchy Level | 11688
- Description | 11688
- Options | 11688
- Required Privilege Level | 11688
- Release Information | 11688

Syntax

```
notification;
```

Hierarchy Level

```
[edit system commit]
```

Description

Notify applications upon commit completion.

Options

There are no options for this configuration statement.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R3.

RELATED DOCUMENTATION

commit (System) | 4455

notification-control

IN THIS SECTION

- Syntax | 11689
- Hierarchy Level | 11689
- Description | 11689
- Required Privilege Level | 11689

Release Information | 11689

Syntax

```
notification-control {
    fpc slot-number {
        disable;
    }
}
```

Hierarchy Level

```
[edit poe]
```

Description

Enable or disable the generation of PoE SNMP traps. If PoE SNMP traps are enabled, an SNMP trap is sent whenever a PoE interface is enabled or disabled.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch

Configuring PoE Interfaces on EX Series Switches

notification-interval

IN THIS SECTION

- Syntax | 11690
- Hierarchy Level | 11690
- Description | 11690
- Options | 11690
- Required Privilege Level | 11691
- Release Information | 11691

Syntax

```
notification-interval seconds;
```

Hierarchy Level

```
[edit ethernet-switching-options mac-notification]
[edit switch-options mac-notification]
```

Description

Configure the MAC notification interval for a switch.

The MAC notification interval is the amount of time the switch waits before sending learned or unlearned MAC address SNMP notifications to the network management server. For instance, if the MAC notification interval is set to 10, all of the MAC address addition and removal SNMP notifications will be sent to the network management system every 10 seconds.

Options

seconds The MAC notification interval, in seconds.

- **Range:** 1 through 60

- **Default:** 30

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Hierarchy level `[edit switch-options]` added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.

RELATED DOCUMENTATION

Configuring Non-ELS MAC Notification

Configuring MAC Notification on Switches with ELS Support

notification (NETCONF)

IN THIS SECTION

- Syntax | 11692
- Hierarchy Level | 11692
- Description | 11692
- Default | 11692
- Options | 11692
- Required Privilege Level | 11692
- Release Information | 11692

Syntax

```
notification {
    interleave;
}
```

Hierarchy Level

```
[edit system services netconf]
```

Description

Enable the NETCONF event notification service. When this service is enabled on supported devices, the NETCONF server sends asynchronous event notifications to clients that subscribe to notifications within a NETCONF session.

Default

If you do not include the `notification` statement, NETCONF clients cannot subscribe to event notifications in a NETCONF session.

Options

`interleave` Enable the NETCONF client and server to exchange RPCs and RPC replies within a NETCONF session that subscribes to event notifications.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 21.2R1.

`interleave` statement introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

NETCONF Event Notifications

notification-options (Security Antivirus)

IN THIS SECTION

- Syntax | 11693
- Hierarchy Level | 11694
- Description | 11694
- Options | 11694
- Required Privilege Level | 11695
- Release Information | 11695

Syntax

```
notification-options {
    fallback-block {
        custom-message custom-message;
        custom-message-subject custom-message-subject;
        (notify-mail-sender | no-notify-mail-sender);
        type (message | protocol-only);
    }
    fallback-non-block {
        custom-message custom-message;
        custom-message-subject custom-message-subject;
        (notify-mail-recipient | no-notify-mail-recipient);
    }
    reputation-block {
        custom-message custom-message;
        custom-message-subject custom-message-subject;
        (notify-mail-sender | no-notify-mail-sender);
        type (message | protocol-only);
    }
```

```
                virus-detection {
                    custom-message custom-message;
                    custom-message-subject custom-message-subject;
                    (notify-mail-sender | no-notify-mail-sender);
                    type (message | protocol-only);
                }
            }
```

Hierarchy Level

```
[edit security utm default-configuration anti-virus],
[edit security utm default-configuration anti-virus sophos-engine],
[edit security utm feature-profile anti-virus profile],
[edit security utm feature-profile anti-virus sophos-engine profile],
[edit logical-systems name security utm default-configuration anti-virus],
[edit logical-systems name security utm default-configuration anti-virus sophos-engine],
[edit logical-systems name security utm feature-profile anti-virus profile],
[edit logical-systems name security utm feature-profile anti-virus sophos-engine profile],
[edit logical-systems name tenants name security utm default-configuration anti-virus],
[edit logical-systems name tenants name security utm default-configuration anti-virus sophos-
engine],
[edit logical-systems name tenants name security utm feature-profile anti-virus profile],
[edit logical-systems name tenants name security utm feature-profile anti-virus sophos-engine
profile],
[edit tenants name security utm default-configuration anti-virus],
[edit tenants name security utm default-configuration anti-virus sophos-engine],
[edit tenants name security utm feature-profile anti-virus profile],
[edit tenants name security utm feature-profile anti-virus sophos-engine profile]
```

Description

There are multiple notification options you can configure to trigger when a virus is detected.

Options

reputation-block Specify the Sophos file reputation block notofication.

`custom-message` *`custom-message`*	Specify the custom message for the reputation block notication.
`custom-message-subject` *`custom-message-subject`*	Specify the custom message subject for the reputation block notication.
`notify-mail-sender`	Notify mail sender.
`no-notify-mail-sender`	Do not notify mail sender.
`type`	Specify virus detection notification type.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

The `reputation-block` option is introduced in Junos OS Release 24.4R1.

notification-options (Security Content Filtering)

IN THIS SECTION

- Syntax | 11696
- Hierarchy Level | 11696
- Description | 11696
- Options | 11696
- Required Privilege Level | 11696

Release Information | 11696

Syntax

```
notification-options {
    custom-message message;
    (notify-mail-sender | no-notify-mail-sender);
     type (message | protocol-only);
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name]
```

Description

You can configure a message notification to trigger when a content filter is matched.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

notification (System Services)

IN THIS SECTION

- Syntax | 11697
- Hierarchy Level | 11697
- Description | 11697
- Options | 11698
- Required Privilege Level | 11698
- Release Information | 11698

Syntax

```
notification {
    allow-clients {
        address ip-address;
    }
    broker-socket-send-buffer-size broker-socket-send-buffer-size;
    max-connections max-connections;
    port port;
}
```

Hierarchy Level

```
[edit system services extension-service]
```

Description

Enable notification services for applications.

Options

`allow-clients address` *`ip-address`*	(Required) Specify IPv4 or IPv6 addresses (prefix length optional) or host names of clients that are allowed to request notifications from the application. You can specify a set of values using square brackets ([ ]). If this option is not configured, the client cannot subscribe to notifications.
`broker-socket-send-buffer-size` *`broker-socket-send-buffer-size`*	Socket send buffer size for the broker to publish the messages
`max-connections` *`max-connections`*	Specify the maximum number of clients that can connect to the application. • **Range:** 1 through 20 • **Default:** 20
`port` *`port`*	Specify the number of the port to accept incoming connections. • **Range:** 1 through 65535 • **Default:** 1883

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

grpc | 7920

notification-targets

IN THIS SECTION

- Syntax | 11699
- Hierarchy Level | 11699
- Description | 11699
- Options | 11699
- Required Privilege Level | 11699
- Release Information | 11700

Syntax

```
notification-targets address port port-number;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

List the destination IP addresses and User Datagram Protocol (UDP) ports to which DFC PICs log exception information and control protocol state transitions, such as timeout values.

Options

address—Allowed destination IP address.

`port` *port-number*—Allowed destination UDP port number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Control Source

notification-time

IN THIS SECTION

- Syntax | 11700
- Hierarchy Level | 11700
- Description | 11701
- Default | 11701
- Options | 11701
- Required Privilege Level | 11701
- Release Information | 11701

Syntax

```
notification-time seconds;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

(MX Series routers only) Configure the notification time value for MAC move reports that a MAC address moves before counting against the threshold values.

Default

1 second

Options

seconds—Time elapsed before MAC move reports are generated.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring MAC Move Parameters

notify

IN THIS SECTION

- Syntax | **11702**
- Hierarchy Level | **11702**
- Description | **11702**

- Options | 11702
- Required Privilege Level | 11702
- Release Information | 11703

Syntax

```
notify name {
    tag tag-name;
    type (trap | inform);
}
```

Hierarchy Level

```
[edit snmp v3]
```

Description

Select management targets for SNMPv3 notifications as well as the type of notifications. Notifications can be either traps or informs.

Options

name—Name assigned to the notification.

tag-name—Notifications are sent to all targets configured with this tag.

`type`—Notification type is `trap` or `inform`. Traps are unconfirmed notifications. Informs are confirmed notifications.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`type inform` option added in Junos OS Release 7.4.

notify-filter (Applying to the Management Target)

IN THIS SECTION

- Syntax | 11703
- Hierarchy Level | 11703
- Description | 11703
- Options | 11703
- Required Privilege Level | 11704
- Release Information | 11704

Syntax

```
notify-filter profile-name;
```

Hierarchy Level

```
[edit snmp v3 target-parameters target-parameters-name]
```

Description

Specify the notify filter to be used by a specific set of target parameters.

Options

profile-name—Name of the notify filter to apply to notifications.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

notify-filter (Configuring the Profile Name)

IN THIS SECTION

- Syntax | 11704
- Hierarchy Level | 11704
- Description | 11705
- Options | 11705
- Required Privilege Level | 11705
- Release Information | 11705

Syntax

```
notify-filter profile-name {
    oid oid (include | exclude);
}
```

Hierarchy Level

```
[edit snmp v3]
```

Description

Specify a group of MIB objects for which you define access. The notify filter limits the type of traps or informs sent to the network management system.

Options

profile-name—Name assigned to the notify filter.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

oid (SNMPv3) | 11799

notify-mail-recipient

IN THIS SECTION

- Syntax | 11706
- Hierarchy Level | 11706
- Description | 11706
- Required Privilege Level | 11706
- Release Information | 11706

Syntax

```
notify-mail-recipient;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-non-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-non-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-non-block]
```

Description

Notify the e-mail recipient about errors returned by the antivirus scan engine when a fallback nonblocking action occurs.

You can specify that the e-mail recipient is not to be notified with the `no-notify-mail-recipient` statement.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

notify-mail-sender (Security Content Filtering Notification Options)

IN THIS SECTION

- Syntax | 11707
- Hierarchy Level | 11707
- Description | 11707
- Required Privilege Level | 11707
- Release Information | 11707

Syntax

```
notify-mail-sender;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name notification-options]
```

Description

Notify the e-mail sender.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

notify-mail-sender (Security Fallback Block)

IN THIS SECTION

- Syntax | 11708
- Hierarchy Level | 11708
- Description | 11708
- Required Privilege Level | 11709
- Release Information | 11709

Syntax

```
notify-mail-sender;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-block]
```

Description

E-mail notification is used to notify the sender or the recipient about the errors returned by either the scan engine or the scan manager when a fallback action occurs.

You can specify that the sender is not to be notified with the `no-notify-mail-sender` statement.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

notify-mail-sender (Security Virus Detection)

IN THIS SECTION

- Syntax | 11709
- Hierarchy Level | 11709
- Description | 11710
- Required Privilege Level | 11710
- Release Information | 11710

Syntax

```
notify-mail-sender;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options virus-detection]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
```

```
notification-options virus-detection]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options virus-detection]
```

Description

E-mail notification is used to notify the sender or the recipient about the detected viruses or the scanning errors. When a virus is detected, an e-mail is sent to the sender upon virus detection.

You can specify that the sender is not to be notified with the `no-notify-mail-sender` statement.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

notify-view

IN THIS SECTION

- Syntax | **11711**
- Hierarchy Level | **11711**
- Description | **11711**
- Options | **11711**
- Required Privilege Level | **11711**
- Release Information | **11711**

Syntax

```
notify-view view-name;
```

Hierarchy Level

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix context-
prefix) security-model (any | usm | v1 | v2c) security-level (authentication | none | privacy)]
```

Description

Associate the notify view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).

Options

view-name—Name of the view to which the SNMP user group has access.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

np-cache (Flexible PIC Concentrator)

IN THIS SECTION

- Syntax | 11712
- Hierarchy Level | 11712

- Description | 11712
- Options | 11713
- Required Privilege Level | 11713
- Release Information | 11713

Syntax

```
np-cache;
```

Hierarchy Level

```
[edit chassis fpc fpc-slot-number]
```

Description

Enable session cache table on IOC.

Starting with Junos OS Release 15.1X49-D10 and and Junos OS Release 17.3R1, NP cache is supported on the SRX5K-MPC (IOC2), SRX5K-MPC3-100G10G (IOC3), and SRX5K-MPC3-40G10G (IOC3) for SRX5400, SRX5600, and SRX5800 devices.

From Junos OS Release 21.2R1, `np-cache` is enabled by default.

The security policy determines whether a session is for Express Path (formerly known as *services offloading*) mode on the selected Flexible PIC Concentrator (FPC).

> **NOTE**: The IOC2 and the IOC3 utilize the delay sessions delete mechanism. The same sessions (sessions with the same five tuples) that are deleted and then reinstalled immediately are not cached on the IOCs.

> **NOTE**: In a cluster deployment, when you configure `np-cache` for an IOC card from node0, the IOC cards on both cluster nodes are activated.

> **NOTE**: To enable IPsec VPN affinity, you must also enable the session cache on IOCs (IOC2 and IOC3) by using the `set chassis fpc <fpc-slot> np-cache` command.

Options

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D10.

nptv6-prefix

IN THIS SECTION

- Syntax | 11713
- Hierarchy Level | 11714
- Description | 11714
- Options | 11714
- Required Privilege Level | 11714
- Release Information | 11714

Syntax

```
nptv6-prefix {
    address-prefix;
        routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name then static-nat]
```

Description

Specify a static IPv6 address prefix. The longest prefix supported is /64.

Options

- *address-prefix*—Specify the address prefix.
- routing-instance—Use the user-defined static NAT routing instance to perform static NAT.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D25.

nptv6-prefix-name

IN THIS SECTION

- Syntax | 11715
- Hierarchy Level | 11715
- Description | 11715
- Options | 11715
- Required Privilege Level | 11715
- Release Information | 11715

Syntax

```
nptv6-prefix-name {
    address-prefix-name;
        routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name then static-nat]
```

Description

Specify an address prefix name from an address book. The longest prefix name supported is /64.

Options

- *address-prefix-name*—Specify an address prefix name from an address book.
- `routing-instance` —Use the user-defined static NAT routing instance to perform static NAT.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D25.

nsr-phantom-holdtime

IN THIS SECTION

Syntax | 11716

- Hierarchy Level | 11716
- Description | 11716
- Options | 11716
- Required Privilege Level | 11717
- Release Information | 11717

Syntax

```
nsr-phantom-holdtime seconds;
```

Hierarchy Level

```
[edit routing-options]
```

Description

Specify to hold phantom IP addresses, that is, prevent these routes from being added to routing tables, for a specified period of time. During this hold-time interval, these routes are maintained in the kernel. After the hold time expires, the routes are added to the routing tables.

You use this option with the Graceful Routing Engine switchover (GRES) feature with nonstop routing (NSR). We strongly recommend that you configure this statement on switches in the QFX10000 line with redundant Routing Engines before you initiate a GRES operation. Doing so prevents traffic loss because routes are added to the routing tables after the hold-time interval expires, but before they are deleted from the kernel during a switchover.

NOTE: This statement doesn't apply to QFX10002 switches because those devices don't have redundant Routing Engines and don't support GRES.

Options

seconds Specify the interval of time to prevent phantom IP addresses from being added to the routing tables. After this interval expires, the routes are added to the routing tables.

> **BEST PRACTICE**: In an EVPN/VXLAN environment with NSR and GRES enabled, the recommended hold-time value is 300 (5 minutes)

- **Range:** 0 through 10000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1x53-D60.

RELATED DOCUMENTATION

- Understanding Graceful Routing Engine Switchover
- *Preventing Traffic Loss in an EVPN/VXLAN Environment With GRES and NSR*
- Example: Configuring Nonstop Active Routing on Switches

nssa

IN THIS SECTION

- Syntax | **11718**
- Hierarchy Level | **11718**
- Description | **11718**
- Options | **11719**
- Required Privilege Level | **11719**
- Release Information | **11719**

Syntax

```
nssa {
    area-range network/mask-length <restrict> <exact> <override-metric metric>;
    default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
    }
    (no-summaries | summaries);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3) area area-id],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Configure a not-so-stubby area (NSSA). An NSSA allows external routes to be flooded within the area. These routes are then leaked into other areas.

You cannot configure an area as being both a stub area and an NSSA.

The remaining statements are explained separately. See CLI Explorer.

Options

summaries | no-summaries Configure whether or not area border routers advertise summary routes into an not-so-stubby area (NSSA):

- `summaries`—Flood summary link-state advertisements (LSAs) into the NSSA.
- `no-summaries`—Prevent area border routers from advertising summaries into an NSSA. If `default-metric` is configured for an NSSA, a Type 3 LSA is injected into the area by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Understanding OSPF Areas

Example: Configuring OSPF Not-So-Stubby Areas

stub | **28665**

nssu

IN THIS SECTION

Syntax | **11720**

Hierarchy Level | **11720**

- Description | 11720
- Default | 11721
- Required Privilege Level | 11721
- Release Information | 11721

Syntax

```
nssu {
    rcp-count number;
    upgrade-group group-name {
        fpcs (slot-number | [list-of-slot-numbers]);
        member member-id {
            fpcs (slot-number | [list-of-slot-numbers]);
        }
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure parameters that affect the nonstop software upgrade (NSSU) process.

> **NOTE**: You use the "request system software nonstop-upgrade" on page 15319 command to initiate NSSU.

The `rcp-count` option (available only on QFX5100 switches) sets the number of parallel `rcp` sessions that NSSU uses to copy the new software to multiple Virtual Chassis or VCF member switches at a time.

The `upgrade-group` options define line-card upgrade groups for NSSU. When you initiate NSSU with at least one upgrade group configured, NSSU upgrades the line cards or Virtual Chassis or VCF members in each upgrade group to the new software version at the same time, in the order in which you configured

them. Upgrade groups reduce the time required to complete an NSSU operation and control the order in which the line cards or members are upgraded.

Line-card upgrade groups are supported on some EX Series switches and EX Series Virtual Chassis that support NSSU and on a QFX5100 VCF.

These statements are all explained separately. You can also consult CLI Explorer.

Default

If you do not configure `rcp-count`, NSSU uses a default algorithm to determine the number of parallel `rcp` sessions to use based on the number of members in the Virtual Chassis or VCF.

If you do not define any line-card upgrade groups, NSSU upgrades line cards or members of a Virtual Chassis or VCF one at a time in ascending order by slot or member number.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`rcp-count` statement introduced in Junos OS Release 14.1X53-D40 for QFX5100 switches only.

RELATED DOCUMENTATION

Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade

Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches

ntp

IN THIS SECTION

Syntax | 11722

- Hierarchy Level | 11723
- Description | 11723
- Options | 11723
- Required Privilege Level | 11728
- Release Information | 11728

Syntax

```
ntp {
    authentication-key key-number type (md5 | sha1 | sha256) value password;
    boot-server  (address | hostname);
    broadcast <address> <key key-number> <routing-instance-name routing-instance-name> <ttl
value>  <version value>;
    broadcast-client;
    interval-range value;
    multicast-client <address>;
    nts <local-certificate local-certificate><trusted-ca (trusted-ca-group trusted-ca-group |
trusted-ca-profile trusted-ca-profile)>;
    peer address <key key-number> <prefer> <version value>;
    restrict address {
        mask network-mask;
        noquery;
    }
    server name {
       key key;
       nts <remote-identity distinguished-name(container container | wildcard wildcard) hostname
hostname>;
       prefer;
       routing-instance routing-instance;
       version version;
   }
     source-address source-address <routing-instance routing-instance-name>;
     threshold value action (accept | reject);
     trusted-key [ key-numbers ];
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure NTP on the device. In both standalone and chassis cluster modes, the primary Routing Engine runs the NTP process to get the time from the external NTP server. Although the secondary Routing Engine runs the NTP process in an attempt to get the time from the external NTP server, this attempt fails because of network issues. For this reason, the secondary Routing Engine uses NTP to get the time from the primary Routing Engine.

When configuring the NTP service in the management VRF (`mgmt_junos`), you must configure at least one IP address on a physical or logical interface within the default routing instance and ensure that this interface is up in order for the NTP service to work with the mgmt_junos VRF.

Options

authentication-key ***key_number***

Configure key (key ID, key type, and key value) to authenticate NTP packets with the devices (servers and clients). The authentication key has two fields:

- `type`—When authentication is specified, the key identifier (key ID) followed by the message digest is appended to the NTP packet header. The supported message digest formats are md5, sha1, sha256.

  > **NOTE**: EX4600 does not support SHA256 authentication for NTP and supports MD5 only.

- `value`—If the key value is available in ASCII format and without special characters, it can be entered directly. If the key value contains special characters or is available in hex format, consider the following:

  For specifying the keys in hex format, prepend a "\x" for each two characters. For hex key example, af60112f...39af4ced, `set system ntp authentication-key <ID> value "\xaf\x60\x11\x2f\....\x39\xaf\x4c \xed"`.

  If the key contains one of the characters from `(null) 0x00, (space) 0x20, " 0x22, & 0x26, ( 0x28 ) 0x29` prepend a "\\x" . For example, `\\x22`.

- **Range:** 1 to 65534

boot-server (*address* | *hostname*)

Configure the server that NTP queries when the device boots to determine the local date and time.

When you boot the device, it issues an ntpdate request, which polls a network server to determine the local date and time. You must configure an NTP boot server that the device uses to determine the time when the device boots. Otherwise, NTP cannot synchronize to a time server if the server time significantly differs from the local device's time.

If you configure an NTP boot server, then when the device boots, it immediately synchronizes with the boot server even if the NTP process is explicitly disabled or if the time difference between the client and the boot server exceeds the threshold value of 1000 seconds.

- **Values:** Configure one of the following:
  - *address*—IP address of an NTP boot server.
  - *hostname*—Hostname of an NTP boot server. If you configure a hostname instead of an IP address, the ntpdate request resolves the hostname to an IP address when the device boots up.

> **NOTE**: This option is deprecated starting in Junos OS Release 20.4R1. NTP boot-server is not supported in Junos OS Evolved. Configure the server using *set system ntp server* under the *edit system ntp server* hierarchy.

broadcast <*address*> <key *key-number*> <routing-instance-name *routing-instance-name*> <ttl *value*> <version *value*>

Configure the device to operate in broadcast mode with the remote system at the specified address. In this mode, the device sends periodic broadcast messages to a client population at the specified broadcast or multicast address. Normally, you include this statement only when the device is operating as a transmitter.

address

Configure the broadcast address on one of the local networks or a multicast address assigned to NTP. You must specify an address, not a hostname. If the multicast address is used, it must be `224.0.1.1`. Configure the multicast protocols PIM and IGMP in order to facilitate the device to transmit NTP packets over multicast address 224.0.1.1.

> **NOTE**: NTP over multicast is not supported within the routing instance on the device.

key *key-number*

(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number (any unsigned 32-bit integer except 0). The key corresponds to the key number you specified in the authentication-key statement.

routing-instance-name *routing-instance-name*

(Optional) Configure the routing instance name in which the interface has an address in the broadcast subnet.

- **Default:** The default routing instance is used to broadcast packets.

ttl *value*

(Optional) Configure the time-to-live (TTL) value.

- **Range:** 1 through 255
- **Default:** 1

version *value*

(Optional) Specify the version number to be used in outgoing NTP packets.

- **Range:** 1 through 4
- **Default:** 4

broadcast-client

Configure the local device to listen for broadcast messages on the local network to discover other servers on the same subnet. To avoid accidental or malicious disruption in this mode, both the local and remote systems must use authentication and the same trusted key and key identifier.

interval-range *value*

Configure the poll interval range.

- **Range:** 0 through 3

multicast-client <*address*>

Configure the local device to listen for multicast messages on the local network. To avoid accidental or malicious disruption in this mode, both the local and remote systems must use authentication and the same trusted key and key identifier.

- **Syntax:** <*address*>—(Optional) Specify one or more IP addresses. If you specify addresses, the device joins those multicast groups.
- **Default:** 224.0.1.1

nts <local-certificate *local-*

Configure the Network Time Security (NTS) features for NTP on your device.

certificate> <trusted-ca (trusted-ca-group *trusted-ca-group* | trusted-ca-profile *trusted-ca-profile*)

local-certificate *local-certificate* — Specify the certificate loaded in your device during certificate enrollment or loaded manually by specifying certificate file.

trusted-ca — Specify a trusted CA group or a CA profile. This configuration is optional. If you do not specify a trusted CA profile the NTP trust all CA profiles configured for NTS.

trusted-ca-group *trusted-ca-group* — Specify the trusted CA group defined under [`set security pki trusted-ca-group`].

trusted-ca-profile *trusted-ca-profile* — Specify the trusted CA profile.

peer *address* <**key** *key-number*> <**prefer**> <**version** *value*>

Configure the local device to operate in symmetric active mode with the remote system at the specified address. In this mode, the local device and the remote system can synchronize with each other. This configuration is useful in a network in which either the local device or the remote system might be a better source of time.

address — Address of the remote system. You must specify an address, not a hostname.

key *key-number* — (Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number (any unsigned 32-bit integer except 0). The key corresponds to the key number you specified in the authentication-key statement.

prefer — (Optional) Mark the remote system as the preferred host, which means that if all other factors are equal, this remote system is chosen for synchronization among a set of correctly operating systems.

version *value* — (Optional) Specify the NTP version number to be used in outgoing NTP packets.

- **Range:** 1 through 4
- **Default:** 4

restrict *address* **mask** *network-mask* **noquery**

Restrict packets from hosts (including remote time servers) and subnets.

- **Syntax:**
  - *address*—Specify the IP address for a host or network.
  - mask *network-mask*—Specify the network mask for a host or network.
  - noquery—Deny ntpq and ntpdc queries from hosts and subnets. These queries can be used in amplification attacks.

server

Configure the local device to operate in client mode with the remote system at the specified address. In this mode, the device can be synchronized with the remote system, but the remote system can never be synchronized with the device.

If the NTP client time drifts so that the difference in time from the NTP server exceeds 128 milliseconds, the client is automatically stepped back into synchronization. If the offset between the NTP client and server exceeds the 1000-second threshold, the client still synchronizes with the server, but it also generates a system log message noting that the threshold was exceeded.

address	Address of the remote system. You must specify an address, not a hostname.
key *key-number*	(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number (any unsigned 32-bit integer except 0). The key corresponds to the key number you specified in the authentication-key statement.
prefer	(Optional) Mark the remote system as the preferred host, which means that if all other factors are equal, this remote system is chosen for synchronization among a set of correctly operating systems.
routing-instance *routing-instance*	(Optional) Routing instance through which the server is reachable.
nts	Enables NTS, which uses Transport Layer Security (TLS) protocol and Authenticated Encryption with Associated Data (AEAD) to obtain network time in an authenticated manner to the users. Specified server must also support the NTS feature when you enable NTS on a client device.

version *value*

(Optional) Specify the NTP version number to be used in outgoing NTP packets.

- **Range:** 1 through 4
- **Default:** 4

source-address *source-address* **<routing-instance** *[routing-instance-name]***>**

A valid IP address configured on one of the device's interfaces to be used as the source address for messages sent to the NTP server, and optionally, the routing instance in which the source address is configured.

- **Default:** The primary address of the interface

threshold *seconds* **action (accept | reject)**

Configure the maximum threshold in seconds allowed for NTP adjustment and specify the mode for NTP abnormal adjustment.

- **Range:** 1 through 600 seconds
- **Values:** Configure one of the following:
  - accept—Enable log mode for abnormal NTP adjustment.
  - reject—Enable reject mode for abnormal NTP adjustment.

trusted-key [*key-numbers* **]**

Configure one or more keys you are allowed to use to authenticate other time servers, when you configure the local device to synchronize its time with other systems on the network. Each key can be any 32-bit unsigned integer except 0. The key corresponds to the key number you specify in the authentication-key statement.

By default, network time synchronization is unauthenticated. The device synchronizes to whatever system appears to have the most accurate time. We strongly encourage you to configure authentication of network time services.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`routing-instance` option for the `server` statement introduced in Junos OS Release 18.1.

`restrict` statement introduced in Junos OS Release 20.1.

RELATED DOCUMENTATION

- Synchronizing and Coordinating Time Distribution Using NTP
- Understanding NTP Time Servers
- Configuring NTP Authentication Keys
- NTP Time Synchronization on SRX Series Devices
- Configuring the NTP Time Server and Time Services
- Configuring the Switch to Listen for Broadcast Messages Using NTP
- Configuring the Switch to Listen for Multicast Messages Using NTP

ntp (QFabric)

IN THIS SECTION

- Syntax | **11729**
- Hierarchy Level | **11730**
- Description | **11730**
- Options | **11730**
- Required Privilege Level | **11730**
- Release Information | **11730**

Syntax

```
ntp {
    authentication-key number type type value password;
     server address <key key-number> <version value> <prefer>;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure Network Time Protocol (NTP) on the switch.

Options

authentication-key Configure NTP authentication keys so authenticated packets can be sent.

server Configure the switch to operate in client mode with the remote system at the specified server address.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

> *Configuring NTP Authentication Keys (QFabric System)*
>
> *Configuring the NTP Time Server and Time Services (QFabric System)*

ntp threshold

IN THIS SECTION

- Syntax | 11731
- Hierarchy Level | 11731
- Description | 11731
- Options | 11732
- Required Privilege Level | 11732
- Release Information | 11732

Syntax

```
threshold value<action (accept | reject)>;
```

Hierarchy Level

```
[edit system ntp]
```

Description

Assign a threshold value for Network Time Protocol (NTP) adjustment that is outside of the acceptable NTP update and specify whether to accept or reject NTP synchronization when the proposed time from the NTP server exceeds the configured threshold value. If **accept** is the specified action, the system synchronizes the device time with the NTP server, but logs the time difference between the configured threshold and the time proposed by the NTP server; if **reject** is the specified action, synchronization with the time proposed by the NTP server is rejected, but the system provides the option of manually synchronizing the device time with the time proposed by the NTP server and logs the time difference between the configured threshold and the time proposed by the NTP server. By logging the time difference and rejecting synchronization when the configured threshold is exceeded, this feature helps improve security on the NTP service.

If this command is not configured or by default, the NTP will allow time adjustments for up to 1000 seconds except for first time adjustment. After NTP synchronization starts, it will allow first time

adjustment to happen without any time limit. After first time adjustment happens, 1000 seconds time limit will be enforced for future time adjustments.

Options

value Specify the maximum value in seconds allowed for NTP adjustment.

- **Range:** 1 through 600.
- **Default:** The default value is 400.

action Specify the actions for NTP abnormal adjustment.

- `accept`—Enable log mode for abnormal NTP adjustment. When the proposed time from the NTP server is outside of the configured threshold value, the device time synchronizes with the NTP server, but the system logs the time difference between the configured threshold and the time proposed by the NTP server.
- `reject`—Enable log and reject mode for abnormal NTP adjustment. When the proposed time from the NTP server is outside of the configured threshold value, the system rejects synchronization, but provides the option for manually synchronizing the time and logs the time difference between the configured threshold and the time proposed by the NTP server.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

ntp | **11721**

set date ntp | **16888**

show system ntp threshold | **26941**

NTP Time Synchronization on SRX Series Devices

nullpdu

IN THIS SECTION

- Syntax | 11733
- Hierarchy Level | 11733
- Description | 11733
- Options | 11734
- Required Privilege Level | 11734
- Release Information | 11734

Syntax

```
nullpdu {
    protocol (ID-0x0000 | ID-0xFFFF);
}
```

Hierarchy Level

```
[edit security sctp]
```

Description

Use the nullpdu function to configure the Stream Control Transmission Protocol (SCTP) null protocol data unit (PDU) value. If you set the null PDU value to 0xFFFF, then the payload protocol identifier value is replaced with 0xFFFF and the user data field is not modified. If you set the null PDU value to 0x0000, then the payload protocol identifier value is replaced with 0x0000 and the first four bytes of the user data field is replaced with zeroes. If all chunks in a packet are null PDUs, the SCTP firewall drops the packet.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs sctp]` hierarchy level is replaced by `[edit security sctp]`.

Options

`protocol`—To specify the SCTP null PDU payload protocol identifier.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Understanding SCTP Multichunk Inspection

num-65-127-prefix

IN THIS SECTION

- Syntax | 11734
- Hierarchy Level | 11735
- Description | 11735
- Options | 11735
- Required Privilege Level | 11736
- Release Information | 11736

Syntax

```
num-65-127-prefix number;
```

Hierarchy Level

```
[edit chassis (QFX Series) forwarding-options profile-name]
```

Description

For the Unified Forwarding Table (UFT) feature, specify how much forwarding table memory to allocate for IPv6 entries with prefix lengths in the range of /65 through /127. The ability to allocate flexibly the memory for IPv6 entries with prefixes in this range extends the use of this memory space to accommodate the appropriate mix of longest-prefix match (LPM) entries that best suits your network. The LPM table stores IPv4 unicast prefixes, IPv6 prefixes with lengths equal to or less than 64, and IPv6 prefixes with lengths from 65 through 127. With this option, you can increase, decrease, or allocate no memory for IPv6 prefixes with lengths from 65 through 127, depending on which version of Junos OS you are using.

NOTE: This statement is supported only for the following forwarding table memory profiles: `l2-profile-one`, `l2-profile-three`, `l2-profile-two`, and `l3-profile`. Do not use this statement with the `custom-profile` or the `lpm-profile` statements.

NOTE: The values you can configure are different depending on the version of Junos OS you are using.

Options

number Specify a numerical value.

- **Range:** (Junos OS Release 13.2x51-D10 only) 1 through 128. Each increment represents 16 IPv6 prefixes with lengths in the range of /65 through /127, for a total maximum of 2,058 prefixes (16 x 128 = 2,048).
- **Default:** 1 (16 IPv6 prefixes with lengths in the range of /65 through /127).
- **Range:** (Junos OS Release 13.2X51-D15 or later) 0 through 4. Each increment allocates memory for 1,000 IPv6 prefixes with lengths in the range of /65 through /127, for a maximum of 4,000 such IPv6 prefixes.
- **Default:** 1 (1,000 IPv6 prefixes with lengths in the range of /65 through /127).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Support for QFX5200 Series switches introduced in Junos OS Release 15.1X53-D30.

RELATED DOCUMENTATION

Configuring the Unified Forwarding Table on Switches

number-of-frames

IN THIS SECTION

- Syntax | 11736
- Hierarchy Level | 11737
- Description | 11737
- Options | 11737
- Required Privilege Level | 11737
- Release Information | 11737

Syntax

```
number-of-frames value;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

Description

Specify the number of consequent frames to declare a delay measurement (DM) session completed.

Options

value—Number of consequent frames to declare DM completed.

- **Range:** 0 through 255 frames.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

number-of-ports

IN THIS SECTION

- Syntax | **11738**
- Hierarchy Level | **11738**

- Hierarchy Level | **11738**
- Description | **11738**
- Options | **11739**
- Required Privilege Level | **11739**
- Release Information | **11739**

Syntax

```
number-of-ports number-of-active-physical-ports;
```

Hierarchy Level

```
[edit chassis fpc slot-number]
[edit chassis fpc slot-number  pic pic-number]
```

Hierarchy Level

```
[edit chassis fpc fpc-slot pic pic-number pic-mode pic-speed]
```

Description

Administratively enable physical ports, for example, to prevent oversubscription of the line card fabric interface. By default, all available ports are enabled. When disabled, the LED on the affected line card will appear yellow on capable line cards.

(MX Series with 16x10GE MPC, MPC3, MPC4, MPC5, and MPC6) You can disable a subset of the physical ports available on the Packet Forwarding Engines of the 16x10GE MPC, and for MICs installed in MPC3, MPC4, MPC5, and MPC6. Specify either 8 or 12 ports by using this statement. When eight active ports are configured, two ports per Packet Forwarding Engine are disabled, and the LEDs on the MPC appear **yellow**. When you specify 12 active ports, one port per Packet Forwarding Engine is disabled and the corresponding LED appear **yellow**. When you do not include this statement in the configuration, all 16 default ports on the MPC are active.

(MX Series with MPC7E-MRATE, MPC8E, and MPC9E) To ensure guaranteed bandwidth by preventing fabric oversubscription, you can disable a subset of the physical ports available on MPC7E-MRATE, MPC8E, and MPC9E. For information about the active ports for MPC7E-MRATE, MPC8E, and MPC9E, see Supported Active Physical Ports on MPC7E-MRATE to Prevent Oversubscription, Supported Active Physical Ports on MIC-MRATE to Prevent Oversubscription.

(MX240, MX480, MX960, MX2010, and MX2020 routers with MIC 10x10GE SFPP and 10x1GE SFP on MPC2E-NG and MPC3E-NG) To ensure guaranteed bandwidth by preventing fabric oversubscription, you can disable a subset of the physical ports available on MPC2E-NG and MPC3E-NG.

(MX204 routers) To ensure guaranteed bandwidth by preventing oversubscription, you can disable a subset of the physical ports available on MX204 routers. For information about the active ports for MX204 routers, see Supported Active Physical Ports on MX204 to Prevent Oversubscription.

Options

number-of-active-physical-ports—Specify the number of physical ports to enable on PICs or MICs on an MPC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

Support for MPC3, MPC4, MPC5, and MPC6 introduced in Junos OS Release 13.3R2.

Support for MPC7E-MRATE MPC introduced in Junos OS Release 15.1F4.

Support for MPC8E and MPC9E introduced in Junos OS Release 15.1F5.

Support for MX10003 MPC introduced in Junos OS Release 17.3R1

Support for MX204 routers introduced in Junos OS Release 17.4R1

Statement introduced in Junos OS Release 16.1 for EX9200 switches.

RELATED DOCUMENTATION

Configuring the Number of Active Ports on 16x10GE 3D MPC

Port Speed on MIC-MRATE Overview

Port Speed on MPC7E-MRATE Overview

Port Speed on MX204 Overview

number-of-sub-ports (interface)

IN THIS SECTION

- Syntax | 11740
- Hierarchy Level | 11740
- Description | 11740
- Options | 11740
- Required Privilege Level | 11741
- Release Information | 11741

Syntax

```
number-of-sub-ports <number-of-sub-ports>;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Specifies the number of channelized interfaces that you can configure on a physical port with the specified speed.

Options

number-of-sub-ports *number-of-subports*	Specify the number of sub-ports per physical port.

- (PTX10003-80C and PTX10003-160C routers)—1, 2, and 4.
- (JNP10K-LC1201)—0 through 7.
- (JNP10K-LC1202)—1 through 8.
- (PTX10001-36MR)—1, 2, 3, 4, and 8.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

Interface hierarchy introduced in Junos OS Evolved Release 20.1R2 for JNP10K-LC1201 line cards on PTX10008 routers.

RELATED DOCUMENTATION

speed (Ethernet) | **28349**

Port Speed

O

IN THIS SECTION

- oam (Layer 2 Circuit Protocols) | **11748**
- oam (MX Series) | **11749**
- oam | **11753**
- oam detection-disable | **11756**
- oam-on-svlan (Ethernet Interfaces) | **11758**

- oam (Protocols LDP) | **11759**
- oam (Protocols MPLS) | **11762**
- oam runtime-disable | **11764**
- object-cache-size | **11765**
- object-names | **11767**
- observation-domain-id | **11768**
- oc192 | **11770**
- ocs (Diameter Applications) | **11771**
- ocsp (Security PKI) | **11774**
- odu-backward-frr-enable | **11776**
- odu-delay-management | **11777**
- odu-signal-degrade | **11779**
- odu-signal-degrade-monitor-enable | **11781**
- odu-ttim-action-enable | **11782**
- off (Destination NAT Next Gen Services) | **11784**
- offer-period | **11785**
- offline-on-fabric-bandwidth-reduction | **11786**
- offload-type | **11788**
- offloads | **11789**
- offset | **11791**
- offset (MX Series) | **11793**
- oid | **11796**
- oid (SNMP Scripts) | **11797**
- oid (SNMPv3) | **11799**
- oif-map (Dynamic IGMP Interface) | **11800**
- oif-map (Dynamic MLD Interface) | **11801**
- oif-map (IGMP Interface) | **11803**
- oif-map (MLD Interface) | **11804**
- oism | **11805**
- oism (Multicast Snooping Options) | **11807**
- omit-no-export (Graceful Restart for BGP Helper) | **11810**
- omit-wildcard-address | **11812**

- on-demand-address-allocation | **11813**
- on-demand-ip-address | **11815**
- on-disk-failure | **11817**
- on-disk-failure (Chassis Routing Engine) | **11818**
- on-error | **11819**
- on-link | **11821**
- on-link (Dynamic Router Advertisement) | **11823**
- on-loss-of-keepalives | **11824**
- on-loss-of-keepalives | **11826**
- on-re-to-fpc-stale | **11827**
- one-to-one (Aggregated Multiservices) | **11829**
- one-way-hardware-timestamp | **11830**
- onlink-subnet-only | **11832**
- op | **11833**
- op invoke-debugger cli | **11836**
- op (Op Scripts) | **11838**
- open-timeout | **11841**
- openflow (Protocols OpenFlow) | **11842**
- operating-mode | **11844**
- operation | **11846**
- optics-options | **11847**
- optimal-route-reflection | **11850**
- optimize-adaptive-teardown | **11852**
- optimize-aggregate-frr | **11853**
- optimize-aggressive | **11855**
- optimize-hold-dead-delay | **11856**
- optimize-switchover-delay | **11857**
- optimize-timer (Protocols MPLS) | **11859**
- optimize-timer (Protocols RSVP) | **11861**
- option-15 (DHCP Relay Agent) | **11862**
- option-16 (DHCP Relay Agent) | **11865**
- option-16 (DHCPv6 Snooping) | **11868**

- option-18 (DHCPv6 Snooping) | **11869**
- option-18 (Interface-ID for DHCPv6 Autosense VLANs) | **11871**
- option-37 (DHCPv6 Snooping) | **11872**
- option-37 (Relay Agent Remote-ID for DHCPv6 Autosense VLANs) | **11874**
- option-60 (DHCP Local Server) | **11876**
- option-60 (DHCP Relay Agent) | **11877**
- option-77 (DHCP Relay Agent) | **11880**
- option-79 | **11883**
- option-82 | **11884**
- option-82 (interfaces) | **11886**
- option-82 (Address-Assignment Pools) | **11888**
- option-82 (DHCP Local Server Authentication) | **11890**
- option-82 (DHCP Local Server Pool Matching) | **11892**
- option-82 (DHCP Relay Agent) | **11893**
- option (DHCP server) | **11895**
- option-match | **11896**
- option-number (DHCP Relay Agent Option) | **11898**
- option-order (DHCP Relay Agent) | **11899**
- option-refresh-rate | **11902**
- option-refresh-rate (Services) | **11903**
- option (Security IDP) | **11905**
- option-type | **11906**
- optional (Commit Scripts) | **11907**
- optional (Event Scripts) | **11909**
- optional-parameters | **11910**
- options (Access Profile) | **11912**
- options (Routing Options) | **11922**
- options-template-id | **11924**
- order | **11926**
- order (Application Identification) | **11927**
- order-priority (Application Identification) | **11929**
- order (Security IDP) | **11930**

- origin-autonomous-system (Origin Validation for BGP) | **11932**
- origin (Diameter Base Protocol) | **11934**
- origin (Diameter Base Protocol) | **11935**
- origin-validation | **11937**
- ospf | **11938**
- ospf-instance | **11940**
- ospf3 | **11942**
- oss-map | **11944**
- other | **11945**
- other-bytes | **11949**
- other-overhead-adjust | **11951**
- other-stateful-configuration | **11953**
- other-stateful-configuration (Dynamic Router Advertisement) | **11954**
- otn-options | **11955**
- otu-ttim-action-enable | **11958**
- otu4 | **11960**
- out-delay | **11961**
- out-of-resources | **11963**
- out-of-resources (Security Antivirus Sophos Engine) | **11965**
- outbound-https | **11966**
- outbound-route-filter | **11970**
- outbound-ssh | **11972**
- outer-tag-protocol-id (RFC 2544 Benchmarking) | **11975**
- outgoing-queue | **11976**
- output (Service Set) | **11978**
- output | **11980**
- output (Accounting) | **11981**
- output-chain | **11983**
- output (Congestion Notification) | **11984**
- output (Dynamic Service Sets) | **11986**
- output-filename (Event Policy) | **11988**
- output-format (Event Policy) | **11989**

- output-forwarding-class-map | **11991**
- output (Forwarding Table) | **11992**
- output-interface (CCC) | **11994**
- output-interface-index | **11995**
- output interfaces (Chassis Synchronization) | **11996**
- output (Mirroring) | **11998**
- output (Monitoring) | **12000**
- output-policer | **12001**
- output (Port Mirroring) | **12003**
- output (Port Mirroring) | **12004**
- output-priority-map | **12007**
- output-queue-priority | **12008**
- output (Sampling) | **12010**
- output (Security Forwarding Options) | **12013**
- output-three-color | **12014**
- output-traffic-control-profile | **12016**
- output-traffic-control-profile (Dynamic CoS Definition) | **12018**
- output-traffic-control-profile-remaining | **12019**
- output-vlan-map | **12021**
- output-vlan-map (Aggregated Ethernet) | **12022**
- output-vlan-map (Dynamic Interfaces) | **12023**
- outputs (Analytics) | **12025**
- over (Application Identification) | **12028**
- over-limit | **12031**
- overall-packet (Flat-File Accounting Options) | **12032**
- overflow-pool | **12035**
- overhead-accounting | **12036**
- overhead-accounting (ANCP) | **12038**
- overhead-accounting (Dynamic Traffic Shaping) | **12039**
- overlay-ecmp | **12041**
- overlay (Packet Forwarding Options) | **12042**
- overlay-path | **12044**

- overload-pool | **12046**
- overload-prefix | **12047**
- overload (Protocols IS-IS) | **12048**
- overload (Protocols OSPF) | **12051**
- override | **12054**
- override-chap-password | **12055**
- override-firewall | **12057**
- override-interval | **12058**
- override-password (Domain Map) | **12060**
- override (PIM Static RP) | **12061**
- override (RADIUS Options) | **12063**
- overrides (DHCP Local Server) | **12065**
- overrides (DHCP Relay Agent) | **12068**
- overrides (DHCP Security) | **12071**
- overrides (DHCP Security for MX Series) | **12072**
- overrides (Enhanced Subscriber Management) | **12074**
- overrides (System Services DHCP) | **12077**
- oversize-message-drop | **12081**
- ovlan-cfi (RFC 2544 Benchmarking) | **12082**
- ovlan-id (RFC 2544 Benchmarking) | **12084**
- ovlan-priority (RFC 2544 Benchmarking) | **12085**
- ovsdb | **12086**
- ovsdb-managed | **12088**
- owner | **12090**
- owner (RMON) | **12092**
- oxid | **12093**

oam (Layer 2 Circuit Protocols)

IN THIS SECTION

- Syntax | 11748
- Hierarchy Level | 11748
- Description | 11748
- Options | 11749
- Required Privilege Level | 11749
- Release Information | 11749

Syntax

```
oam {
    bfd-liveness-detection;
    ping-interval;
    ping-multiplier ping-count;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols l2vpn],
[edit routing-instances routing-instance-name protocols vpls],
[edit routing-instances routing-instance-name protocols vpls neighbor address],
[edit protocols l2circuit neighbor address interface interface-name]
```

Description

Allows you to configure bidirectional forwarding detection (BFD) and a control channel for a pseudowire, in addition to the corresponding operations and management functions to be used over that control channel. BFD provides a low resource fault detection mechanism for the continuous monitoring of the pseudowire data path and for detecting data plane failures. The `control-channel` statement is not applicable to Layer 2 circuit pseudowires.

Options

bfd-liveness-detection The `bfd-liveness-detection` statement and substatements are described in the Junos OS Routing Protocols Library.

ping-multiplier *ping-count* Specify the number of LSP ping packets to delay the virtual circuit connectivity verification (VCCV) Bidirectional Forwarding Detection (BFD) session from going down. The VCCV BFD session is signaled down only after the specified number of LSP ping packets are lost. This feature is supported for Layer 2 Circuit, Layer 2 VPN, and VPLS technologies.

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support for VPLS FEC 129 introduced in Junos OS Release 12.2.

`ping-multiplier` statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring BFD for VCCV for Layer 2 VPNs, Layer 2 Circuits, and VPLS

oam (MX Series)

IN THIS SECTION

- Syntax | 11750
- Hierarchy Level | 11752

- Description | 11752
- Required Privilege Level | 11753
- Release Information | 11753

Syntax

```
oam {
    ethernet {
        connectivity-fault-management {
            action-profile profile-name {
                default-actions {
                    interface-down;
                }
            }
            performance-monitoring {
                delegate-server-processing;
                hardware-assisted-timestamping;
                hardware-assisted-keepalives;
                sla-iterator-profiles {
                    profile-name {
                        avg-fd-twoway-threshold;
                        avg-ifdv-twoway-threshold;
                        avg-flr-forward-threshold;
                        avg-flr-backward-threshold;
                        disable;
                        calculation-weight {
                            delay delay-weight;
                            delay-variation delay-variation-weight;
                        }
                        cycle-time milliseconds;
                        iteration-period connections;
                        measurement-type (loss | statistical-frame-loss | two-way-delay);
                    }
                }
            }
            linktrace {
                age (30m | 10m | 1m | 30s | 10s);
                path-database-size path-database-size;
```

```
            }
            maintenance-domain domain-name {
                level number;
                name-format (character-string | none | dns | mac+2octet);
                maintenance-association ma-name {
                    short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
                    protect-maintenance-association protect-ma-name;
                    remote-maintenance-association remote-ma-name;
                    continuity-check {
                        convey-loss-threshold;
                        hold-interval minutes;
                        interface-status-tlv;
                        interval (100ms | 10m | 10ms | 10s | 1m | 1s);
                        loss-threshold number;
                        port-status-tlv;
                    }
                    mep mep-id {
                        auto-discovery;
                        direction (up | down);
                        interface interface-name (protect | working);
                        lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-
defect | rem-err-xcon | xcon );
                        priority number;
                        remote-mep mep-id {
                            action-profile profile-name;
                            sla-iterator-profile profile-name {
                                data-tlv-size size;
                                iteration-count count-value;
                                priority priority-value;
                            }
                        }
                    }
                }
            }
        }
        link-fault-management {
            action-profile profile-name {
                action {
                    link-down;
                    send-critical-event;
                    syslog;
                }
                event {
```

```
                    link-adjacency-loss;
                    link-event-rate {
                        frame-error count;
                        frame-period count;
                        frame-period-summary count;
                        symbol-period count;
                    }
                    protocol-down;
                }
            }
            interface interface-name {
                apply-action-profile
                link-discovery (active | passive);
                loopback-tracking;
                pdu-interval interval;
                pdu-threshold threshold-value;
                remote-loopback;
                event-thresholds {
                    frame-error count;
                    frame-period count;
                    frame-period-summary count;
                    symbol-period count;
                }
                negotiation-options {
                    allow-remote-loopback;
                    no-allow-link-events;
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

For Ethernet interfaces on M320, M120, MX Series, and T Series routers and PTX Series Packet Transport Routers, provide IEEE 802.3ah Operation, Administration, and Maitenance (OAM) support.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

IEEE 802.3ah OAM Link-Fault Management Overview

oam

IN THIS SECTION

- Syntax | 11753
- Hierarchy Level | 11756
- Description | 11756
- Required Privilege Level | 11756
- Release Information | 11756

Syntax

```
oam {
    ethernet{
        connectivity-fault-management {
            action-profile profile-name {
                action {
                    interface-down;
```

```
            }
            default-actions {
                interface-down;
            }
            event {
                adjacency-loss;
            }
        }
        linktrace {
            age (30m | 10m | 1m | 30s | 10s);
            path-database-size path-database-size;
        }
        maintenance-domain domain-name {
            level number;
            mip-half-function (none | default |explicit);
            name-format (character-string | none | dns | mac+2oct);
            maintenance-association ma-name {
                continuity-check {
                    hold-interval minutes;
                    interface-status-tlv;
                    interval (10m | 10s | 1m | 1s| 100ms);
                    loss-threshold number;
                    port-status-tlv;
                }
                mep mep-id {
                    auto-discovery;
                    direction down;
                    interface interface-name;
                    remote-mep mep-id {
                        action-profile profile-name;
                    }
                }
            }
        }
        performance-monitoring {
            sla-iterator-profiles {
                profile-name {
                    calculation-weight {
                        delay delay-value;
                        delay-variation delay-variation-value;
                    }
                    cycle-time cycle-time-value;
                    iteration-period iteration-period-value;
```

```
                        measurement-type two-way-delay;
                        passive;
                    }
                }
            }
        }
        link-fault-management {
            action-profile profile-name;
                action {
                    syslog;
                    link-down;
                }
                event {
                    link-adjacency-loss;
                    link-event-rate {
                        No Link Title count;
                        No Link Title count;
                        No Link Title count;
                        symbol-period count;
                    }
                }
            interface interface-name {
                link-discovery (active | passive);
                pdu-interval interval;
                pdu-threshold threshold-value;
                remote-loopback;
                No Link Title {
                    No Link Titlecount;
                     No Link Title count;
                     No Link Title count;
                     symbol-period count;
                }
                negotiation-options {
                    No Link Title;
                     no-allow-link-events;
                }
            }
        }
    }
```

Hierarchy Level

```
[edit protocols]
```

Description

Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support for Ethernet interfaces on EX Series switches or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support on the switches.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

connectivity-fault-management introduced in Junos OS Release 10.2 for EX Series switches.

RELATED DOCUMENTATION

Example: Configure Ethernet OAM Link Fault Management

Example: Configure Ethernet OAM Connectivity Fault Management on EX Series Switches

Configure Ethernet OAM Link Fault Management

Configure Ethernet OAM Connectivity Fault Management (CLI Procedure)

oam detection-disable

IN THIS SECTION

Syntax | 11757

- Hierarchy Level | 11757
- Description | 11757
- Default | 11757
- Required Privilege Level | 11757
- Release Information | 11758

Syntax

```
oam detection-disable
```

Hierarchy Level

```
[edit chassis fabric]
```

Description

Disable the fabric Operation, Administration, Maintenance (OAM) feature, which helps you detect failures in fabric paths when a switch fabric card (switch interface board or SIB) or an FPC transitions from online to offline state or vice versa.

> **NOTE**: This command is not applicable to Junos OS Evolved Releases 20.4R2 and 21.1R1.

Default

The fabric OAM monitoring feature is enabled by default. However, in Junos OS Evolved Releases 20.4R2 and 21.1R1, this feature is disabled by default.

Required Privilege Level

Interface—To view this statement in the configuration.

Release Information

Interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Error Handling by Fabric OAM

oam-on-svlan (Ethernet Interfaces)

IN THIS SECTION

- Syntax | 11758
- Hierarchy Level | 11758
- Description | 11758
- Required Privilege Level | 11759
- Release Information | 11759

Syntax

```
oam-on-svlan;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Enable propagation of the Ethernet IEEE 802.1ag Operation, Administration, and Maintenance (OAM) state of a static single-tagged service VLAN (S-VLAN) logical interface to the dynamic or static double-tagged customer VLAN (C-VLAN) logical interface and associated subscriber interfaces configured on the S-VLAN. The static S-VLAN logical interface must be configured with Ethernet OAM connectivity

fault management (CFM) on a Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet physical interface. The C-VLAN logical interface must have the same S-VLAN (outer) tag as the S-VLAN logical interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Configuring Ethernet OAM Support for Service VLANs with Double-Tagged Customer VLANs

Ethernet OAM Support for Service VLANs Overview

oam (Protocols LDP)

IN THIS SECTION

- Syntax | 11759
- Hierarchy Level | 11760
- Description | 11761
- Options | 11761
- Required Privilege Level | 11761
- Release Information | 11761

Syntax

```
oam {
    bfd-liveness-detection {
```

```
        detection-time threshold milliseconds;
        ecmp;
        failure-action {
            remove-nexthop;
            remove-route;
        }
        holddown-interval milliseconds;
        ingress-policy ingress-policy-name;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-transmit-interval milliseconds;
        multiplier detection-time-multiplier;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (0 | 1 | automatic);
    }
    fec fec-address;
    ingress-policy ingress-policy-name;
    lsp-ping-interval seconds;
    periodic-traceroute {
        disable;
        exp exp-value;
        fanout fanout-value;
        frequency minutes;
        paths number-of-paths;
        retries retry-attempts;
        source address;
        ttl ttl-value;
        wait seconds;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp]
[edit protocols ldp]
```

Description

Configure Operation, Administration, and Maintenance (OAM) and Bidirectional Forwarding Detection (BFD) protocol for LDP.

Options

`fec` *`fec-address`*—Specify the forwarding equivalence class (FEC) address. You must either specify a FEC address or configure an OAM ingress policy to ensure that the BFD session comes up.

`lsp-ping-interval` *`seconds`*—Specify the duration of the LSP ping interval in seconds. To issue a ping on an LDP-signaled LSP, use the `ping mpls ldp` command.

- **Default:** 60 seconds
- **Range:** 30 through 3,600 seconds

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

`lsp-ping-interval` option introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring BFD for LDP LSPs

oam (Protocols MPLS)

IN THIS SECTION

- Syntax | 11762
- Hierarchy Level | 11763
- Description | 11763
- Options | 11764
- Required Privilege Level | 11764
- Release Information | 11764

Syntax

```
oam {
    bfd-liveness-detection{
        failure-action teardown;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-transmit-interval milliseconds;
        multiplier detection-time-multiplier;
    }
    lsp-ping-interval seconds;
    mpls-tp-mode;
    performance-monitoring {
        querier {
            loss {
                traffic-class tc-value {
                    query-interval milliseconds;
                    measurement-quantity bytes|packets;
                    average-sample-size sample size;
                    loss-threshold loss threshold value;
                    loss-threshold-window number of samples for loss threshold;
                }
            }
            delay {
                traffic-class tc-value {
                    query-interval milliseconds;
```

```
                    padding-size size;
                    average-sample-size sample size;
                    rtt-delay-threshold rtt threshold value;
                    twcd-delay-threshold twcd threshold value;
                }
            }
            loss-delay {
                traffic-class tc-value {
                    query-interval milliseconds;
                    measurement-quantity bytes|packets;
                    padding-size size;
                    average-sample-size sample size;
                    loss-threshold loss threshold value;
                    loss-threshold-window number of samples for loss threshold;
                    rtt-delay-threshold rtt threshold value;
                    twcd-delay-threshold twcd threshold value;
                }
            }
        }
        responder {
            loss {
                min-query-interval milliseconds;
            }
            delay {
                min-query-interval milliseconds;
            }
        }
    }
}
```

Hierarchy Level

```
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name]
[edit protocols mpls label-switched-path lsp-name primary path-name]
```

Description

Enable Operation, Administration, and Maintenance (OAM) for RSVP-signaled LSPs.

Options

`lsp-ping-interval` *seconds*—Specify the duration of the LSP ping interval in seconds. To issue a ping on an RSVP-signaled LSP, use the `ping mpls rsvp` command.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

`lsp-ping-interval` option introduced in Junos OS Release 9.4.

`performance-monitoring` configuration statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring BFD for MPLS IPv4 LSPs

oam runtime-disable

IN THIS SECTION

- Syntax | 11765
- Hierarchy Level | 11765
- Description | 11765
- Default | 11765
- Required Privilege Level | 11765
- Release Information | 11765

Syntax

```
oam runtime-disable
```

Hierarchy Level

```
[edit chassis fabric]
```

Description

Disable the runtime fabric Operation, Administration, Maintenance (OAM) feature, which helps you detect and report failures in fabric planes during system runtime.

Default

The runtime Fabric OAM feature is enabled by default.

Required Privilege Level

Interface—To view this statement in the configuration.

Release Information

Interface-control—To add this statement to the configuration.

object-cache-size

IN THIS SECTION

- Syntax | 11766
- Hierarchy Level | 11766
- Description | 11766
- Options | 11766

Required Privilege Level | 11766
Release Information | 11767

Syntax

```
object-cache-size value;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-
provider]
```

Description

Configure the size of the object cache. When this setting is changed, the PIC reboots.

Options

value—Amount of object cache, in MB. Only values in increments of 128 MB are allowed.

- **Range:** For Multiservices 100 PIC, range is 128 MB through 512 MB. If the `wired-process-mem-size` statement at the same hierarchy level has a value of 512 MB, the maximum value for this statement is 128 MB.
- **Range:** For Multiservices 400 PIC, range is 128 MB through 1280 MB. If the `wired-process-mem-size` statement at the same hierarchy level has a value of 512 MB, the maximum value for this statement is 512 MB.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

forwarding-db-size | 7454
policy-db-size | 12674
wired-process-mem-size | 31133

object-names

IN THIS SECTION

- Syntax | 11767
- Hierarchy Level | 11767
- Description | 11768
- Options | 11768
- Required Privilege Level | 11768
- Release Information | 11768

Syntax

```
object-names {
    mib-object-name;
}
```

Hierarchy Level

```
[edit accounting-options mib-profile profile-name]
```

Description

Specify the name of each MIB object for which MIB statistics are collected for an accounting-data log file.

Options

mib-object-name—Name of a MIB object. You can specify more than one MIB object name.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configure the MIB Profile

observation-domain-id

IN THIS SECTION

- Syntax | 11769
- Hierarchy Level | 11769
- Description | 11769
- Options | 11769
- Required Privilege Level | 11770
- Release Information | 11770

Syntax

```
observation-domain-id domain-id;
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix template template-name]
```

Description

For IPFIX flows, an identifier of an observation domain is locally unique to an exporting process of the templates. The export process uses the observation domain ID to uniquely identify to the collection process in which the flows were metered. We recommend that you configure this ID to be unique for each IPFIX flow. A value of 0 indicates that no specific observation domain is identified by this information element. The ID configured here is contained within Information Element 149.

MX and QFX Series: Typically, this attribute is used to limit the scope of other information elements. If the observation domain is not unique, the collector cannot uniquely identify an IPFIX device. If you configure the same observation domain ID for different template types, such as for IPv4 and IPv6, it does not impact flow monitoring because the actual or the base observation domain ID is transmitted in the flow. The actual observation domain ID is derived from the value you configure and plus other parameters such as the slot number, lookup chip (LU) instance, and Packet Forwarding Engine instance. Such a method of computation of the observation domain ID ensures that this ID is not the same for two IPFIX devices.

PTX Series: When you configure the observation domain ID, the software attaches the ID to a particular template type. If you configure the same observation domain ID for two different template types, such as for IPv4 and IPv6, this does not impact flow monitoring, because the configured ID is not what is being sent. The value sent in the packets is derived from that configured value and the FPC slot value. This method ensures two IPFIX devices can never have the same value of observation domain ID.

Options

domain-id—Identifier for the observation domain for IPFIX flows.

- **Range:** 0 through 255

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Evolved 22.4R1.

RELATED DOCUMENTATION

Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows

Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows

oc192

IN THIS SECTION

- Syntax | 11770
- Hierarchy Level | 11771
- Description | 11771
- Options | 11771
- Required Privilege Level | 11771
- Release Information | 11771

Syntax

```
oc192;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

Description

Set the line rate or speed of the OTN signal to optical channel transport unit 2 (OTU2).

Options

oc192 OTU2 line rate or 10 Gbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3 for MX Series routers.

RELATED DOCUMENTATION

10-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

ocs (Diameter Applications)

IN THIS SECTION

- Syntax | 11772
- Hierarchy Level | 11773
- Description | 11773

- Required Privilege Level | 11774
- Release Information | 11774

Syntax

```
ocs {
    global {
        service-context-id service-context;
    }
    partition partition-name  {
        called-station-id station-name;
        charging-id number;
        destination-host ocs-hostname;
        destination-realm ocs-realm-name;
        diameter-instance;
        draining;
        draining-response-timeout seconds;
        final-response-timeout seconds;
        force-continue;
        ggsn-address address;
        ggsn-mcc-mncggsn-mcc-mnc;
        max-outstanding-requests number;
        send-origin-state-id number;
        user-name-include  {
            base-interface-name;
            delimiter delimiter-character;
            domain-name my-domain;
            interface-name;
            mac-address;
            nas-port-id;
            origin-host;
            origin-realm;
            user-name;
            user-prefix pref;
        }
            backup {
                limit;
                timeout seconds;
```

```
                overflow { deny-login | drop-oldest };
            }
            sftp-backup {
                accumulation-timeout seconds;
                accumulation-count messages;
                accumulation-size bytes;
                logical-system;
                retry-interval seconds;
                response-timeout seconds;
                routing-instance;
                address ipv4 | ipv6;
                port number;
                directory filename;
                file-name-prefix filename-prefix;
                node-ipv4-address;
                node-ipv6-address;
                ssh-login;
                ssh-connection-linger;
                ssh-log-verbose;
                ssh-passphrase;
            }
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure the Online Charging System (OCS) global attributes and partition. The OCS interacts with the Policy and Charging Enforcement Function (PCEF). The PCEF optionally reports usage and receives additional authorizations from the OCS using the 3rd Generation Partnership Project (3GPP) Gy protocol. Broadband PCEF (BPCEF) interactions with the OCS use online session charging with centralized unit determination and centralized rating.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

ocsp (Security PKI)

IN THIS SECTION

- Syntax | **11774**
- Hierarchy Level | **11775**
- Description | **11775**
- Options | **11775**
- Required Privilege Level | **11776**
- Release Information | **11776**

Syntax

```
ocsp {
    connection-failure (disable | fallback-crl);
```

```
    disable-responder-revocation-check;
    nonce-payload (enable | disable);
    url ocsp-url;
}
```

Hierarchy Level

```
[edit security pki ca-profile ca-profile-name revocation-check]
```

Description

Configure Online Certificate Status Protocol (OCSP) to check the revocation status of a certificate.

Options

connection-failure (Optional) Specify action to take if there is a connection failure to the OCSP responder. If this option is not configured and there is no response from the OCSP responder, certificate validation will fail.

- **disable** Skip the revocation check if the OCSP responder is not reachable.
- **fallback-crl** Use CRL to check the revocation status of the certificate.

disable-responder-revocation-check (Optional) Disable revocation check for the CA certificate received in an OCSP response. The certificates received in an OCSP response generally have shorter lifetimes and revocation check is not required.

nonce-payload (Optional) Send a nonce payload to prevent replay attack. A nonce payload is sent by default unless it is explicitly disabled. If enabled, the SRX Series Firewall expects OCSP responses to contain a nonce payload, otherwise the revocation check will fail. If OCSP responders are not capable of responding with a nonce payload, disable this option.

- **disable** Explicitly disable the sending of a nonce payload.
- **enable** Enable the sending of a nonce payload. This is the default.

url *ocsp-url* Specify HTTP addresses for OCSP responders. A maximum of two HTTP URL addresses can be configured. If the configured URLs are not reachable, or URLs are not configured, the URL from the certificate being verified is checked.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D20.

RELATED DOCUMENTATION

Basic Elements of PKI in Junos OS

revocation-check (Security PKI) | **15739**

odu-backward-frr-enable

IN THIS SECTION

- Syntax | **11776**
- Hierarchy Level | **11777**
- Description | **11777**
- Default | **11777**
- Required Privilege Level | **11777**
- Release Information | **11777**

Syntax

```
odu-backward-frr-enable;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

Description

Insert the ODU status into the transmitted OTN frames and monitor the received OTN frames for the ODU BER status.

Default

By default, FRR ODU backward FRR insertion is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

odu-delay-management

IN THIS SECTION

- Syntax | 11778
- Hierarchy Level | 11778

- Description | 11778
- Default | 11778
- Options | 11778
- Required Privilege Level | 11778
- Release Information | 11779

Syntax

```
odu-delay-management {
    (bypass | no-bypass);
    (monitor-end-point | no-monitor-end-point);
    number-of-frames value;
    (no-start-measurement | start-measurement;
}
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify Optical Channel Data Unit (ODU) delay management options.

Default

If you omit the odu-delay-management statement, the ODU delay management options are disabled.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

odu-signal-degrade

IN THIS SECTION

- Syntax | 11779
- Hierarchy Level | 11780
- Description | 11780
- Default | 11780
- Options | 11780
- Required Privilege Level | 11780
- Release Information | 11780

Syntax

```
odu-signal-degrade {
    ber-threshold-clear;
    ber-threshold-signal-degrade;
    interval
}
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify optical channel data unit (ODU) signal degradation threshold-related values.

Default

If you omit the `odu-signal-degrade` statement, the default threshold values are used.

The following are the default threshold values for optical channel data unit (ODU) signal degradation for the P2-100GE-OTN PIC:

- `ber-threshold-clear`—1E-09
- `ber-threshold-signal-degrade`—1E-06
- `interval`—10 ms

The following are the default threshold values for optical channel data unit (ODU) signal degradation for the MPC5E and the MIC6-100G-CFP2 MIC:

- `ber-threshold-clear`—1.14E-5
- `ber-threshold-signal-degrade`—1.0E-6
- `interval`—10 ms

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

odu-signal-degrade-monitor-enable

IN THIS SECTION

- Syntax | 11781
- Hierarchy Level | 11781
- Description | 11781
- Default | 11782
- Required Privilege Level | 11782
- Release Information | 11782

Syntax

```
odu-signal-degrade-monitor-enable;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

Description

Enable monitoring of signal degradation of ODU BER in the received OTN frames.

Default

By default, FRR signal degrade monitoring disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

odu-ttim-action-enable

IN THIS SECTION

- Syntax | 11783
- Hierarchy Level | 11783
- Description | 11783
- Default | 11783
- Options | 11783
- Required Privilege Level | 11783
- Release Information | 11783

Syntax

```
(odu-ttim-action-enable | no-odu-ttim-action-enable);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify whether consequent action for Optical Channel Data Unit (ODU) TTIM is enabled or disabled.

Default

If you omit the odu-ttim-action-enable statement, consequent action for ODU TTIM is disabled.

Options

odu-ttim-action-enable	Enable consequent action for ODU TTIM.
no-odu-ttim-action-enable	Disable consequent action for ODU TTIM.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

off (Destination NAT Next Gen Services)

IN THIS SECTION

- Syntax | 11784
- Hierarchy Level | 11784
- Description | 11784
- Required Privilege Level | 11784

Syntax

```
off;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set-name rule rule-name then destination-nat]
```

Description

Tun off destination address translation for the rule. Use this statement when configuring port forwarding without destination address translation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

offer-period

IN THIS SECTION

- Syntax | 11785
- Hierarchy Level | 11785
- Description | 11785
- Options | 11786
- Required Privilege Level | 11786
- Release Information | 11786

Syntax

```
offer-period milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name
bidirectional df-election],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name bidirectional df-election],
[edit protocols pim interface (Protocols PIM)  interface-name bidirectional df-election],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name bidirectional df-election]
```

Description

Configure the designated forwarder (DF) election offer period for bidirectional PIM. When a DF election Offer or Winner message fails to be received, the message is retransmitted. The `offer-period` statement modifies the interval between repeated DF election messages. The `robustness-count` statement determines the minimum number of DF election messages that must fail to be received for DF election to fail. To prevent routing loops, all routing devices on the link must have a consistent view of the DF. When the DF election fails because DF election messages are not received, forwarding on bidirectional PIM routes is suspended.

If a router receives from a neighbor a better offer than its own, the router stops participating in the election for a period of `robustness-count` * `offer-period`. Eventually, all routers except the best candidate stop sending Offer messages.

Options

milliseconds—Interval to wait before retransmitting DF Offer and Winner messages.

- **Range:** 100 through 10,000 milliseconds
- **Default:** 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Bidirectional PIM

Example: Configuring Bidirectional PIM

robustness-count | **15823**

offline-on-fabric-bandwidth-reduction

IN THIS SECTION

- Syntax | **11787**
- Hierarchy Level | **11787**
- Description | **11787**
- Required Privilege Level | **11787**

Release Information | 11787

Syntax

```
offline-on-fabric-bandwidth-reduction;
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

Configure an FPC with degraded fabric bandwidth offline, to avoid causing a null route in the chassis for an extended time.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Disabling an FPC with Degraded Fabric Bandwidth

Managing Bandwidth Degradation

offload-type

IN THIS SECTION

- Syntax | 11788
- Hierarchy Level | 11788
- Description | 11788
- Options | 11788
- Required Privilege Level | 11789
- Release Information | 11789

Syntax

```
offload-type {
    none;
    pfe-timestamp;
}
```

Hierarchy Level

```
[edit services monitoring rpm owner name test test-name]
```

Description

Enable timestamping of RPM probe messages in the Packet Forwarding Engine (PFE) host processor, by offloading the processing of RPM probes to the PFE. This feature is supported only with `icmp-ping`, `icmp-ping-timestamp`, `udp-ping`, and `udp-ping-timestamp` probe types.

Options

none	Timestamping performed on the Routing Engine
pfe-timestamp	Timestamping performed on the PFE

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.1.

RELATED DOCUMENTATION

> *Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers*

offloads

IN THIS SECTION

- Syntax | 11789
- Hierarchy Level | 11790
- Description | 11790
- Options | 11790
- Required Privilege Level | 11790
- Release Information | 11790

Syntax

```
offloads {
    disable;
}
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Offloads configuration for the VNF interface.

NOTE: This feature is available only in enhanced-orchestration mode.

Options

disable Disable the **offloads** option that offloads the configuration for the VNF interface.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D471.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | **27723**

offset

IN THIS SECTION

- Syntax | 11791
- Hierarchy Level | 11791
- Description | 11791
- Default | 11792
- Options | 11792
- Required Privilege Level | 11793
- Release Information | 11793

Syntax

```
offset (0 |30 | 50);
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
[edit security macsec connectivity-association connectivity-association-name secure-channel secure-channel-name]
```

Description

Specifies the number of octets in an Ethernet frame that are sent in unencrypted plain-text when encryption is enabled for MACsec.

Configure a confidentiality offset for MACsec. When MACsec is enabled with encryption, the confidentiality offset specifies a number of octets in an Ethernet frame that are sent in unencrypted plain-text.

Per 802.1AE-2006, confidentiality offset is relevant for switch-to-host connections by allowing a system that is incapable of terminating the secure association before distributing the load to perform load balancing across multiple processors based on the first few bytes of packets. Additionally, confidentiality

offset can be used to expose IPv4 or IPv6 headers to bump-in-the-wire devices, such as transparent firewalls or monitoring devices.

Setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the remaining traffic. Setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the remaining traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

You configure the `offset` in the [`edit security macsec connectivity-association` *`connectivity-association-name`*] hierarchy when you are enabling MACsec using static connectivity association key (CAK) or dynamic security mode.

You configure the `offset` in the [`edit security macsec connectivity-association` *`connectivity-association-name`* `secure-channel` *`secure-channel-name`*] hierarchy when you are enabling MACsec using static secure association key (SAK) security mode.

Default

0

Options

0 Specifies that no octets are unencrypted. When you set the offset to 0, all traffic on the interface where the connectivity association or secure channel is applied is encrypted.

30 Specifies that the first 30 octets of each Ethernet frame are unencrypted.

> **NOTE**: In IPv4 traffic, setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the rest of the traffic. An offset of 30, therefore, is typically used when a feature needs this information to perform a task on IPv4 traffic.

50 Specified that the first 50 octets of each Ethernet frame are unencrypted.

> **NOTE**: In IPv6 traffic, setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the rest of the traffic. An offset

of 50, therefore, is typically used when a feature needs this information to perform a task on IPv6 traffic.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced on SRX Series Firewalls in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | 10278

offset (MX Series)

IN THIS SECTION

- Syntax | 11794
- Hierarchy Level | 11794
- Description | 11794
- Default | 11794
- Options | 11795
- Required Privilege Level | 11795
- Release Information | 11795

Syntax

```
offset (0 |30 | 50);
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name],
[edit security macsec connectivity-association connectivity-association-name secure-channel
secure-channel-name]
```

Description

Specifies the number of octets in an Ethernet frame that are sent in unencrypted plain-text when encryption is enabled for MACsec.

Setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the remaining traffic. Setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the remaining traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

You configure the `offset` in the `[edit security macsec connectivity-association` *`connectivity-association-name`*`]` hierarchy when you are enabling MACsec using static connectivity association key (CAK) or dynamic security mode.

You configure the `offset` in the `[edit security macsec connectivity-association` *`connectivity-association-name`* `secure-channel`*`secure-channel-name`*`]` hierarchy when you are enabling MACsec using static secure association key (SAK) security mode.

Default

0

Options

0 Specifies that no octets are unencrypted. When you set the offset to 0, all traffic on the interface where the connectivity association or secure channel is applied is encrypted.

30 Specifies that the first 30 octets of each Ethernet frame are unencrypted.

> **NOTE**: In IPv4 traffic, setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the rest of the traffic. An offset of 30, therefore, is typically used when a feature needs this information to perform a task on IPv4 traffic.

50 Specified that the first 50 octets of each Ethernet frame are unencrypted.

> **NOTE**: In IPv6 traffic, setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the rest of the traffic. An offset of 50, therefore, is typically used when a feature needs this information to perform a task on IPv6 traffic.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

oid

IN THIS SECTION

- Syntax | 11796
- Hierarchy Level | 11796
- Description | 11796
- Default | 11796
- Options | 11796
- Required Privilege Level | 11797
- Release Information | 11797

Syntax

```
oid object-identifier <exclude | include>;
```

Hierarchy Level

```
[edit snmp view view-name]
```

Description

Specify an object identifier (OID) used to represent a subtree of MIB objects.

Default

`include`. The subtree of MIB objects represented by the specified OID is included.

Options

`exclude`—Exclude the subtree of MIB objects represented by the specified OID.

`include`—Include the subtree of MIB objects represented by the specified OID.

object-identifier—OID used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

oid (SNMP Scripts)

IN THIS SECTION

- Syntax | 11797
- Hierarchy Level | 11798
- Description | 11798
- Options | 11798
- Required Privilege Level | 11798
- Release Information | 11798

Syntax

```
oid id {
    priority;
}
```

Hierarchy Level

```
[edit system scripts snmp file filename]
```

Description

For Junos OS SNMP scripts, specify the global object identifier (*OID*), associated with a subtree of object in the management information base (*MIB*), to be mapped to the SNMP script. Specify the OID using either a sequence of dotted integers or a subtree name.

Options

priority (Optional) Override the default registration priority of the OID. The highest priority is 1; the lowest priority is 255.

- **Default:** 127
- **Range:** 1 through 255

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Example: Processing Unsupported OIDs with an SNMP Script

oid (SNMPv3)

IN THIS SECTION

- Syntax | 11799
- Hierarchy Level | 11799
- Description | 11799
- Default | 11799
- Options | 11799
- Required Privilege Level | 11800
- Release Information | 11800

Syntax

```
oid oid <include | exclude>;
```

Hierarchy Level

```
[edit snmp v3 notify-filter profile-name]
```

Description

Specify an object identifier (OID) used to represent a subtree of MIB objects. This OID is a prefix that the represented MIB objects have in common.

Default

`include`. The subtree of MIB objects represented by the specified OID is included.

Options

`exclude`—Exclude the subtree of MIB objects represented by the specified OID.

`include`—Include the subtree of MIB objects represented by the specified OID.

oid—Object identifier used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

oif-map (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | 11800
- Hierarchy Level | 11801
- Description | 11801
- Options | 11801
- Required Privilege Level | 11801
- Release Information | 11801

Syntax

```
oif-map map-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name]
```

Description

Associates an OIF map to the IGMP interface using a dynamic profile. The OIF map is a routing policy statement that can contain multiple terms.

Options

map-name—Name of the OIF map.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

oif-map (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 11802
- Hierarchy Level | 11802
- Description | 11802
- Options | 11802
- Required Privilege Level | 11802
- Release Information | 11802

Syntax

```
oif-map map-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Description

Associate an outgoing interface (OIF) map to a dynamic MLD logical interface. The OIF map is a routing policy statement that can contain multiple terms.

Options

map-name—Name of the OIF map.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *Dynamic MLD Configuration Overview*
>
> Example: Configuring Multicast with Subscriber VLANs

oif-map (IGMP Interface)

IN THIS SECTION

- Syntax | 11803
- Hierarchy Level | 11803
- Description | 11803
- Required Privilege Level | 11803
- Release Information | 11803

Syntax

```
oif-map map-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Description

Associates an outgoing interface (OIF) map to the IGMP interface. The OIF map is a routing policy statement that can contain multiple terms.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring Multicast with Subscriber VLANs

oif-map (MLD Interface)

IN THIS SECTION

- Syntax | 11804
- Hierarchy Level | 11804
- Description | 11804
- Required Privilege Level | 11804
- Release Information | 11805

Syntax

```
oif-map map-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

Associate an outgoing interface (OIF) map to an MLD logical interface. The OIF map is a routing policy statement that can contain multiple terms.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring Multicast with Subscriber VLANs

oism

IN THIS SECTION

- Syntax | 11805
- Hierarchy Level | 11806
- Description | 11806
- Options | 11806
- Required Privilege Level | 11806
- Release Information | 11807

Syntax

```
oism {
    originate-smet-on-revenue-vlan-too;
    pim-evpn-gateway {
        external-irb irb-interface-name;
        peg-df-election {
            delay-time seconds;
            mod;
            preference {
                use-least;
                value value;
            }
        }
    }
```

```
    supplemental-bridge-domain-irb irb-interface-name;
}
```

Hierarchy Level

```
[edit routing-instances name protocols evpn]
```

Description

Set optimized intersubnet multicast (OISM) options in a Layer 3 virtual routing and forwarding (VRF) routing instance on an EVPN fabric device.

Options

originate-smet-on-revenue-vlan-too

Originate selective multicast Ethernet tag (SMET) routes on revenue VLANs (or bridge domains) upon receiving a local IGMP report. SMET routes are EVPN Type 6 routes. The revenue VLANs are the customer VLANs that serve the hosts on the access side.

By default, when an OISM device receives a local IGMP join request, it originates a Type 6 route on the OISM supplemental bridge domain (SBD). With this option, the device also sends Type 6 routes on the revenue VLANs.

This option is helpful for compatibility with other vendor devices that don't support OISM. Those devices can't create the right states for the revenue bridge domain VLANs upon receiving Type 6 routes on the SBD.

supplemental-bridge-domain-irb *irb-interface-name*

Associate an IRB interface with the OISM SBD on a device where you have enabled OISM.

The remaining statements are explained separately. See CLI Explorer for all options.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

IPv4 Inter-VLAN Multicast Forwarding Modes for EVPN-VXLAN Overlay Networks

Optimized Inter-Subnet Multicast in EVPN Networks

PEG DF Election

oism (Multicast Snooping Options)

IN THIS SECTION

- Syntax | 11807
- Hierarchy Level | 11807
- Description | 11808
- Options | 11808
- Required Privilege Level | 11809
- Release Information | 11809

Syntax

```
oism {
    (conserve-mcast-routes-in-pfe | install-star-g-routes);
}
```

Hierarchy Level

```
[edit multicast-snooping-options],
[edit routing-instances name multicast-snooping-options]
```

Description

Set optimized intersubnet multicast (OISM) options. You can use OISM to optimize multicast traffic flow in an EVPN-VXLAN fabric.

NOTE: The functions of the `conserve-mcast-routes-in-pfe` option and the `install-star-g-routes` option are mutually exclusive, so you can use only one or the other of these options in a routing instance. See the descriptions of these options below for details on when to use each option.

Options

conserve-mcast-routes-in-pfe

(ACX Series routers, QFX5130 switches, and QFX5700 switches) Conserve Packet Forwarding Engine (PFE) multicast route table space on the device by not installing Layer 2 (L2) multicast snooping (*,G) and (S,G) states in the PFE when the device receives EVPN Type 6 routes. The device uses only the Layer 3 (L3) multicast routes in the PFE to forward multicast traffic. Without this option, these devices install L2 multicast group entries in the PFE for all OISM revenue bridge domain VLANs and the SBD.

We require you configure this option on ACX Series routers, QFX5130 switches, and QFX5700 switches in EVPN-VXLAN networks running OISM in the server leaf or border leaf role. With this option, these devices conserve PFE table space by installing only the L3 multicast routes; they avoid installing L2 multicast snooping routes.

QFX5130 or QFX5700 switches running as standalone AR replicators with OISM conserve PFE multicast snooping route table space by default. As a result, don't use this option when you configure those switches in the assisted replication (AR) replicator role with OISM.

See *OISM and AR Scaling with Many VLANs* for more information on using this option.

install-star-g-routes

Set the Routing Engine to immediately install (*,G) multicast routes (upon receiving an EVPN Type 6 route) on the PFE for all of the revenue VLANs in the routing instance.

Without this option, the device prioritizes saving resources on the PFE by not installing multicast routes until the multicast traffic arrives. The default behavior avoids issues with scaling on lower-end devices.

If you configure this option, the device ensures the PFE has multicast routing information available before any multicast traffic arrives. This option minimizes packet loss at the onset of a multicast flow. However, the device consumes extra PFE resources for routes that are not even being used yet (and might never be used). As a result, when you configure this option, you trade off taking up the extra PFE resources to improve network latency.

This option acts only on VLANs associated with a Layer 3 virtual routing and forwarding (VRF) instance with OISM enabled.

Set this option:

- Globally if you configure EVPN in the default-switch instance.
- In the MAC-VRF instances if you configure EVPN in instances of type `mac-vrf`.

We require that you configure `install-star-g-routes` with OISM on the QFX10000 line of switches, QFX5130-32CD switches, and QFX700 switches when you configure those devices with AR in the AR replicator role.

In releases prior to Junos OS and Junos OS Evolved Release 23.4R1, you must also configure the `install-star-g-routes` option on the following devices when you configure them as OISM server leaf or border leaf devices:

- The QFX10000 line of switches.
- PTX10001-36MR, PTX10004, PTX10008, and PTX10016 routers.

Starting in Junos OS and Junos OS Evolved Release 23.4R1, we no longer require this option on any platform when you configure the device as an OISM server leaf or border leaf device.

We don't recommend setting this option other than in the use cases mentioned above. Consider this option only if you have very stringent latency requirements and can trade off higher scaling to achieve better network latency.

For more information about how this option works, see *Latency and Scaling Trade-Offs for Installing Multicast Routes with OISM (install-star-g-routes Option)* .

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 and Junos OS Evolved 22.1R1.

`conserve-multicast-routes-in-pfe` option introduced in Junos OS Evolved Releases 22.4R2 and 23.1R1.

RELATED DOCUMENTATION

Optimized Inter-Subnet Multicast in EVPN Networks

Assisted Replication Multicast Optimization in EVPN Networks

omit-no-export (Graceful Restart for BGP Helper)

IN THIS SECTION

- Syntax | 11810
- Hierarchy Level | 11810
- Description | 11811
- Required Privilege Level | 11812
- Release Information | 11812

Syntax

```
omit-no-export;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp graceful-restart long-lived advertise-to-
non-llgr-neighbor],
[edit logical-systems logical-system-name protocols bgp group group-name graceful-restart long-
lived advertise-to-non-llgr-neighbor],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
graceful-restart long-lived advertise-to-non-llgr-neighbor],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp graceful-
restart long-lived advertise-to-non-llgr-neighbor],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
```

```
group-name graceful-restart long-lived advertise-to-non-llgr-neighbor],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name neighbor address graceful-restart long-lived advertise-to-non-llgr-neighbor],
[edit routing-instances instance-name protocols bgp graceful-restart long-lived advertise-to-non-
llgr-neighbor],
[edit routing-instances instance-name protocols bgp group group-name graceful-restart long-lived
advertise-to-non-llgr-neighbor],
[edit routing-instances instance-name protocols bgp group group-name neighbor address graceful-
restart long-lived advertise-to-non-llgr-neighbor],
[edit protocols bgp graceful-restart long-lived advertise-to-non-llgr-neighbor],
[edit protocols bgp group group-name graceful-restart long-lived advertise-to-non-llgr-neighbor],
[edit protocols bgp group group-name neighbor address graceful-restart long-lived advertise-to-
non-llgr-neighbor]
```

Description

Cause the no-export BGP community to be prevented from being automatically added to routes advertised to external BGP neighbors (presumed to be CE routers). The no-export routes containing this community name are not advertised outside a BGP confederation boundary. In VPN deployments, for example, BGP is often used as a PE-CE protocol. It might be a practical necessity in such deployments to accommodate interoperation with CEs that cannot easily be upgraded to support specifications such as this one. This requirement causes a problem while ensuring that "stale" routing information does not leak beyond the perimeter of routers that support these procedures where one or more IBGP routers are not upgraded. In the VPN PE-CE case, the protocol in use is EBGP, and the LOCAL_PREF, an IBGP-only path attribute, is used.

The principal motivation for restricting the propagation of "stale" routing information is the reason to prevent it from spreading without limit once it exits the BGP confederation boundary. VPN deployments are typically topologically constrained, removing this concern. For this reason, an implementation might advertise stale routes over a PE-CE session, when explicitly configured. In such a scenario, the implementation must attach the NO_EXPORT community to the routes in question by default, as an additional protection against stale routes spreading without limit. Attachment of the NO_EXPORT community can be disabled explicitly to accommodate exceptional cases. It might be necessary to advertise stale routes to a CE in some VPN deployments, even if the CE does not support this specification. In that case, if you configure the PE routers to advertise such routes, you must notify the operator of the CE receiving the routes, and the CE must be configured to depreference the routes. Typical BGP implementations perform this operation by matching on the LLGR_STALE community, and setting the LOCAL_PREF for matching routes to zero.

When the `omit-no-export` option is added or removed, the session is reset. This rest of a session enables LLGR stale routes to be readvertised with or without the no- export community (which is added outside of the export policy).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

> *Configuring Graceful Restart Options for BGP*
>
> Junos OS High Availability User Guide

omit-wildcard-address

IN THIS SECTION

- Syntax | 11812
- Hierarchy Level | 11813
- Description | 11813
- Default | 11813
- Required Privilege Level | 11813
- Release Information | 11813

Syntax

```
omit-wildcard-address;
```

Hierarchy Level

```
[edit dynamic-profiles name routing-options (Dynamic Profiles) multicast]
```

Description

Omit wildcard source and group fields in S-PMSI AD / Leaf AD NLRI updates sent by Junos devices. This helps to avoid format mismatches on wildcard addresses thereby maximizing MVPN interoperability with devices from other vendors.

Default

omit-wildcard-address is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1R2

on-demand-address-allocation

IN THIS SECTION

- Syntax | 11814
- Hierarchy Level | 11814
- Description | 11814
- Required Privilege Level | 11814
- Release Information | 11814

Syntax

```
on-demand-address-allocation;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server dual-stack-group
dual-stack-group-name],
[edit logical-systems name system services dhcp-local-server dual-stack-group dual-stack-group-
name],
[edit routing-instances name system services dhcp-local-server dual-stack-group dual-stack-group-
name],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name]
```

Description

Enables an address or prefix to be allocated on demand when a dual-stack subscriber session is established. Starting in Junos OS Release 18.1R1, when reauthentication is configured, enables per-family address allocation as each family's DHCP session is established. In earlier released, applies only to the second family of the dual stack.

> **NOTE**: You must configure `on-demand-address-allocation` if you also configure reauthentication for dual-stack, single-session DHCP subscribers. This is true whether you enable reauthentication with the `reauthenticate` statement or the Reauthenticate-On-Renew VSA (26-206).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring RADIUS Reauthentication for DHCP Subscribers

RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers

Single-Session DHCP Dual-Stack Overview

Configuring Access Profile Options for Interactions with RADIUS Servers

Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution

Configuring Access Components for the PPPoE Wholesale Network Solution

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

on-demand-ip-address

IN THIS SECTION

- Syntax | **11815**
- Hierarchy Level | **11815**
- Description | **11816**
- Default | **11816**
- Required Privilege Level | **11816**
- Release Information | **11816**

Syntax

```
on-demand-ip-address;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-
interface-unit”].
```

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" ppp-options],
[edit interfaces pp0 unit unit-number ppp-options], [edit protocols ppp-service]
```

Description

For IPv4 and IPv6 dual-stack PPP subscribers, enables on-demand allocation and de-allocation of an IPv4 address after initial PPP authentication for a subscriber who does not have an existing IPv4 address.

Configuration changes take effect as follows:

- When you change this setting for a dynamic PPP interface (at the `[edit dynamic-profiles]` hierarchy level), the change takes effect only for new subscriber logins.
- When you change this setting for a static PPP interface (at the `[edit interfaces pp0]` hierarchy level, the subscribers on the interface are logged out.
- When you change this setting globally (at the `[edit protocols ppp-service]` hierarchy level), the change takes effect only for new subscriber logins.

If you enable on-demand allocation at both the interface and global levels, the global configuration takes precedence and changes take effect for new subscriber logins.

Default

This functionality is disabled by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Conserving IPv4 Addresses for Dual-Stack PPP Subscribers Using On-Demand IPv4 Address Allocation

Configuring Static On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers

Configuring Dynamic On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers

Configuring Global On-Demand IPv4 Address Allocation for Dual-Stack PPP Subscribers

on-disk-failure

IN THIS SECTION

- Syntax | 11817
- Hierarchy Level | 11817
- Description | 11817
- Options | 11817
- Required Privilege Level | 11818
- Release Information | 11818

Syntax

```
on-disk-failure {
    disk-failure-action (halt | reboot);
}
```

Hierarchy Level

```
[edit chassis redundancy]
[edit chassis routing-engine]
```

Description

Instruct the router to halt or reboot if it detects hard disk errors on the Routing Engine.

Options

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

graceful-switchover (Chassis) | 7786

Enabling a Routing Engine to Reboot on Hard Disk Errors

Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)

High Availability Features for EX Series Switches Overview

on-disk-failure (Chassis Routing Engine)

IN THIS SECTION

- Syntax | 11818
- Hierarchy Level | 11819
- Description | 11819
- Options | 11819
- Required Privilege Level | 11819
- Release Information | 11819

Syntax

```
on-disk-failure {
    disk-failure-action (halt | reboot);
}
```

Hierarchy Level

```
[edit chassis routing-engine]
```

Description

Instruct the router to halt or reboot if it detects hard disk errors on the Routing Engine.

Options

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before JUNOS Release 7.4.

The `disk-failure-action` statement added in JUNOS Release 9.0.

RELATED DOCUMENTATION

Enabling a Routing Engine to Reboot on Hard Disk Errors

on-error

IN THIS SECTION

- Syntax | 11820
- Hierarchy Level | 11820
- Description | 11820

- Options | 11820
- Required Privilege Level | 11821
- Release Information | 11821

Syntax

```
on-error {
    raise-alarm;
    power (cycle | off);
    write-coredump;
}
```

Hierarchy Level

```
[edit chassis cfeb slot-number],
[edit chassis feb slot-number],
[edit chassis fpc slot-number sanity-poll],
[edit chassis lcc number fpc number sanity-poll]
```

Description

Instruct the FPC or FEB or CFEB to perform actions during an error condition.

Options

`raise-alarm`—Generate and display a chassis alarm in case of an error.

`power cycle`—Reboot the FPC or FEB or CFEB after generating a core file. This statement is useful in case of temporary software errors that are eliminated after reboot.

`power off`—Halt the FPC or FEB or CFEB and keep it offline. This statement is useful in case of permanent hardware failures.

CAUTION: The `power off` statement halts the FPC or FEB or CFEB. Ensure that you have backup paths through different FPC or FEB or CFEB to avoid service outage.

NOTE: The `power cycle` and `power off` statements are mutually exclusive: You can configure either the `power cycle` or the `power off` statement for an error.

`write-coredump`—Trigger the core file in case of an error.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Sanity Polling

sanity-poll | **16302**

retry-count (FPC Error Condition) | **15698**

on-link

IN THIS SECTION

- Syntax | **11822**
- Hierarchy Level | **11822**
- Description | **11822**
- Default | **11822**
- Required Privilege Level | **11822**

Release Information | 11823

Syntax

```
(on-link | no-on-link);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-
name prefix prefix],
[edit protocols router-advertisement interface interface-name prefix prefix]
```

Description

Specify whether to enable prefixes to be used for onlink determination:

- **no-on-link**—Disable prefixes from being used for onlink determination.
- **on-link**—Enable prefixes to be used for onlink determination.

Router advertisement messages carry prefixes and information about them. A prefix is onlink when it is assigned to an interface on a specified link. The prefixes specify whether they are onlink or not onlink. A node considers a prefix to be onlink if it is represented by one of the link's prefixes, a neighboring router specifies the address as the target of a redirect message, a neighbor advertisement message is received for the (target) address, or any neighbor discovery message is received from the address. These prefixes are also used for address autoconfiguration. The information about the prefixes specifies the lifetime of the prefixes, whether the prefix is autonomous, and whether the prefix is onlink.

Default

Prefixes are onlink unless explicitly disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

on-link (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 11823
- Hierarchy Level | 11823
- Description | 11823
- Default | 11824
- Required Privilege Level | 11824
- Release Information | 11824

Syntax

```
(on-link | no-on-link);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols router-advertisement interface interface-nameprefix prefix]
```

Description

Specify whether to enable prefixes to be used for onlink determination:

- `no-on-link`—Disable prefixes from being used for onlink determination.
- `on-link`—Enable prefixes to be used for onlink determination.

Default

The configured object is enabled unless explicitly disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA

Dynamic Router Advertisement Configuration Overview

on-loss-of-keepalives

IN THIS SECTION

- Syntax | 11825
- Hierarchy Level | 11825
- Description | 11825
- Default | 11825
- Required Privilege Level | 11825
- Release Information | 11825

Syntax

```
on-loss-of-keepalives;
```

Hierarchy Level

```
[edit chassis redundancy failover]
```

Description

Instruct the backup router to take primary role if it detects a loss of keepalive signal from the primary Routing Engine.

Default

The `on-loss-of-keepalives` statement must be included at the `[edit chassis redundancy failover]` hierarchy level for failover to occur.

When the `on-loss-of-keepalives` statement is included but graceful Routing Engine switchover *is not* configured, failover occurs after 300 seconds (5 minutes).

When the `on-loss-of-keepalives` statement is included and graceful Routing Engine switchover *is* configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds (4 seconds on M20 routers) . The keepalive time is not configurable.

> **NOTE**: The `on-loss-of-keepalives` statement at the `[edit chassis redundancy]` hierarchy is not supported on PTX platforms running Junos Evolved. These platforms default to a switchover when keepalive messages are not detected.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

On Detection of a Loss of Keepalive Signal from the Primary Routing Engine

keepalive-time (Chassis) | **9471**

on-loss-of-keepalives

IN THIS SECTION

- Syntax | **11826**
- Hierarchy Level | **11826**
- Description | **11826**
- Default | **11827**
- Required Privilege Level | **11827**
- Release Information | **11827**

Syntax

```
on-loss-of-keepalives;
```

Hierarchy Level

```
[edit chassis redundancy failover]
```

Description

Instruct the backup router to take primary role if it detects a loss of keepalive signal from the primary Routing Engine.

Default

The `on-loss-of-keepalives` statement must be included at the `[edit chassis redundancy failover]` hierarchy level for failover to occur.

When the `on-loss-of-keepalives` statement is included but graceful Routing Engine switchover *is not* configured, failover occurs after 300 seconds (5 minutes).

When the `on-loss-of-keepalives` statement is included and graceful Routing Engine switchover *is* configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

- graceful-switchover (Chassis) | **7786**
- keepalive-time (Chassis) | **9471**
- *Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)*
- *High Availability Features for EX Series Switches Overview*

on-re-to-fpc-stale

IN THIS SECTION

- Syntax | **11828**
- Hierarchy Level | **11828**
- Release Information | **11828**
- Description | **11828**

Required Privilege Level | 11828

Required Privilege Level

Syntax

```
on-re-to-fpc-stale;
```

Hierarchy Level

```
[edit chassis redundancy failover]
```

Release Information

Statement introduced in Junos OS Release 15.2 on the MX240, MX480, MX960, MX2010, and MX2020.

Statement added in Junos OS Release 22.2R1 on the MX10008, MX10016, PTX10008, PTX10016, QFX10008, and QFX10016.

Description

Instruct the backup router to take mastership if the em0 (MX240, MX480, MX960, MX2010, and MX2020) or management Ethernet Interface (MX10008, MX10016, PTX10008, PTX10016, QFX10008, and QFX10016) fails on the master Routing Engine.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

failover (Chassis) | 6778

redundancy | 14166

redundancy (Graceful Switchover) | 14170

one-to-one (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 11829
- Hierarchy Level | 11829
- Description | 11829
- Options | 11830
- Required Privilege Level | 11830
- Release Information | 11830

Syntax

```
one-to-one {
    preferred-backup preferred-backup;
}
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options high-availability-options]
```

Description

Configure a single interface to be the preferred one-to-one (1:1) backup for an active aggregated multiservices (AMS) interface. If the active interface fails, the backup interface takes over. A 1:1 (stateful) configuration synchronizes traffic states and data structures between the active and backup PICs. This is required for high availability of IPsec connections. 1:1 configuration is supported only on MS-MPCs but is not currently supported on MX-SPC3s.

NOTE: The preferred backup must be one of the member interfaces (mams–) that have already been configured at the [`edit interfaces interface-name load-balancing-options`] hierarchy level.

Options

preferred-backup *preferred-backup*

Use the specified interface as the preferred backup member interface. The member interface format is mams-*a*/*b*/0, where *a* is the Flexible PIC Concentrator (FPC) slot number and *b* is the PIC slot number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.2.

RELATED DOCUMENTATION

high-availability-options (Aggregated Multiservices) | **8068**

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

one-way-hardware-timestamp

IN THIS SECTION

- Syntax | **11831**
- Hierarchy Level | **11831**
- Description | **11831**
- Required Privilege Level | **11831**
- Release Information | **11831**

Syntax

```
one-way-hardware-timestamp;
```

Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Description

Enable timestamping of RPM probe messages for one-way delay and jitter measurements. You must configure this statement along with the `destination-interface` statement to invoke timestamping. This feature is supported only with `icmp-ping`, `icmp-ping-timestamp`, `udp-ping`, and `udp-ping-timestamp` probe types.

> **NOTE**: Starting in Junos OS Evolved Release 20.1R1, the function provided by this statement has been replaced by the `offload-type (none|pfe-timestamp)` statement.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Statement introduced in Junos OS Release 19.1 for PTX Series routers.

RELATED DOCUMENTATION

Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches

destination-interface | 5430

hardware-timestamp | 7975

onlink-subnet-only

IN THIS SECTION

- Syntax | 11832
- Hierarchy Level | 11832
- Description | 11832
- Default | 11832
- Required Privilege Level | 11833
- Release Information | 11833

Syntax

```
onlink-subnet-only;
```

Hierarchy Level

```
[edit protocols neighbor-discovery]
```

Description

Enable this option to prevent the device from responding to a neighbor solicitation (NS) from a prefix that is not included as one of the device interface prefixes.

After configuring the `onlink-subnet-only` statement, the Routing Engine needs to be restarted using the `request system reboot both-routing-engines` command. If the attacker's IPv6 destination address is already in the forwarding-table, it is not removed after you configure the `onlink-subnet-only` statement, and therefore the device continues to respond to ping NSs. Restarting the Routing Engine removes the entry from the forwarding table.

Default

Disabled

Required Privilege Level

admin— To view this statement in the configuration.

admin-control— To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Understanding How to Control Inbound Traffic Based on Protocols

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-routing/config-guide-routing-neighbor-discovery.html

op

IN THIS SECTION

- Syntax | 11833
- Description | 11834
- Options | 11834
- Additional Information | 11835
- Required Privilege Level | 11835
- Output Fields | 11835
- Sample Output | 11835
- Release Information | 11835

Syntax

```
op filename
<detail>
```

```
<argument-name argument-value>
<key (md5 | sha-256 | sha1) key-value
<url url>
```

Description

Execute an op script. To execute local op scripts, you must configure the op script filename at the `[edit system scripts op file]` hierarchy level.

Op scripts can be stored in one of the following locations:

- On the local device in the **/var/db/scripts/op** directory on the hard disk
- On the local device in the **/config/scripts/op** directory on the flash drive.

  To execute op scripts from the flash drive, you must configure the `load-scripts-from-flash` statement at the `[edit system scripts]` hierarchy level.
- At a remote URL

Op scripts are executed with the access privileges of the user who invokes the command to execute the script.

> **NOTE**: If you include the `op url` command under the `[edit event-options policy` *`policy-name`* `then execute-commands commands]` hierarchy to execute a remote op script as an event policy action, Python scripts are executed using the access privileges of the generic, unprivileged user and group `nobody`. If you do not configure the `user-name` statement, SLAX and XSLT scripts are executed with root privileges.

Options

detail	(Optional) Display detailed output.
argument-name argument-value	(Optional) Specify one or more arguments to the script. For each argument you include on the command line, you must specify a corresponding value for the argument.
key (md5 | sha-256 | sha1) ***key-value***	(Optional) With the `url` option, specify a checksum hash to verify the integrity of the remote op script. Starting in Junos OS Release 18.2R2 and 18.3R1, Junos OS supports only the SHA-256 hash function for script checksum hashes. Earlier releases support the *MD5*, *SHA-1*, and SHA-256 hash functions.

url *url* (Optional) Specify the URL for a remote script.

> **NOTE**: Statements configured under the `[edit system scripts op]` hierarchy level are only enforced for op scripts that are local to the device. If you configure memory allocation, script dampening, script start options, traceoptions, or other op script-specific statements within that hierarchy, the device does not apply the configuration when you execute a remote script using the `op url` command.

Additional Information

For more information about op scripts, see the Automation Scripting User Guide.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

op

```
user@host> op script1 interface ge-0/2/0.0 protocol inet
```

op url

```
user@host> op url https://www.example.net/opscripts/script1.slax key sha-256
150bf53383769f3bfedd41fe73320777f208d4fda81230cb27b8738 interface ge-0/2/0.0 protocol inet
```

Release Information

Command introduced in Junos OS Release 7.6.

`key` option introduced in Junos OS Release 10.0.

`url` option introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Executing an Op Script on the Local Device

Executing an Op Script from a Remote Site

checksum (Scripts) | **3311**

file checksum md5 | **7037**

file checksum sha-256 | **7039**

file checksum sha1 | **7041**

op invoke-debugger cli

IN THIS SECTION

- Syntax | **11836**
- Description | **11837**
- Options | **11837**
- Required Privilege Level | **11837**
- Output Fields | **11837**
- Sample Output | **11837**
- Release Information | **11838**

Syntax

```
op invoke-debugger cli filename
<detail>
<argument-name argument-value>
```

Description

Invoke an interactive debugger to trace the execution of an op script that is enabled in the configuration.

When you execute this command for a SLAX script, the device invokes the SLAX debugger and displays the (sdb) prompt. When you execute this command for a Python script, the device invokes the `pdb` interactive debugger for Python programs and displays the (Pdb) prompt. For more information about the Python debugger, see the official Python documentation.

Options

detail	Display detailed output.
argument-name argument-value	(Optional) Specify one or more arguments for the script. For each argument you include on the command line, you must specify a corresponding value for the argument.
filename	Script filename

Required Privilege Level

maintenance

Output Fields

When you enter the `op invoke-debugger cli` command, depending on the script type, the CLI displays either the (sdb) or (Pdb) prompt where you can enter debugging commands.

Sample Output

op invoke-debugger cli (SLAX script)

```
user@host> op invoke-debugger cli script1.slax
sdb: The SLAX Debugger (version )
Type 'help' for help
(sdb)
```

op invoke-debugger cli (Python script)

```
user@host> op invoke-debugger cli script1.py
(Pdb)
```

op invoke-debugger cli script argument1 value1

```
user@host> op invoke-debugger cli script2.slax interface ge-0/2/0.0
sdb: The SLAX Debugger (version )
Type 'help' for help
(sdb)
```

Release Information

Command introduced in Junos OS Release 13.1.

Support for the Python debugger added in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *SLAX Debugger, Profiler, and callflow*

op (Op Scripts)

IN THIS SECTION

- Syntax | 11839
- Hierarchy Level | 11840
- Description | 11840
- Options | 11840
- Required Privilege Level | 11840
- Release Information | 11840

Syntax

```
op {
    allow-url-for-python;
    dampen {
        dampen-options {
            cpu-factor cpu-factor;
            line-interval line-interval;
            time-interval time-interval;
        }
    }
    file filename {
        allow-commands "regular-expression";
        arguments {
            argument-name {
                description descriptive-text;
            }
        }
        cert-file path;
        checksum hash-algorithm hash-value;
        command filename-alias;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        description descriptive-text;
        refresh;
        refresh-from url;
        routing-instance routing-instance-name;
        source url;
        start {
            start-options {
                mem-factor percentage;
            }
        }
    }
    max-datasize size;
    no-allow-url;
    refresh;
```

```
        refresh-from url;
        start {
            start-options {
                mem-factor percentage;
            }
        }
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
```

Hierarchy Level

```
[edit system scripts]
```

Description

Configure an operation scripting mechanism.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Storing and Enabling Scripts

open-timeout

IN THIS SECTION

- Syntax | 11841
- Hierarchy Level | 11841
- Description | 11841
- Options | 11841
- Required Privilege Level | 11842
- Release Information | 11842

Syntax

```
open-timeout seconds;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

```
[edit services service-set service-set-name service-set-options tcp-session]
```

Description

Configure a timeout period for Transmission Control Protocol (TCP) session establishment.

Options

seconds—Timeout period.

- **Default:** 5 seconds
- **Range:** 4 through 224 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring Default Timeout Settings for Services Interfaces

openflow (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 11842
- Hierarchy Level | 11843
- Description | 11843
- Default | 11843
- Options | 11843
- Required Privilege Level | 11844
- Release Information | 11844

Syntax

```
openflow {
    switch switch-name {
        controller {
```

```
            address address;
            id id;
            protocol tcp {
                port port;
            }
            role equal;
        }
        default-action (drop | packet-in);
        interfaces {
            interface-name port-id port;
        }
        purge-flow-timer seconds;
    }
    traceoptions {
        file <filename> <files number> <match regular-expression> <size size> <world-readable |
 no-world-readable>;
        flag flag;
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure support for OpenFlow on a device running Junos OS. To configure OpenFlow, the device must be running a Junos OS release that supports OpenFlow and have the OpenFlow software package installed. The OpenFlow software package release must match the Junos OS release of the device on which the software is installed.

Default

OpenFlow is disabled on the device.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Support on Juniper Networks Devices

Understanding Support for OpenFlow on Devices Running Junos OS

OpenFlow Operational Mode Commands

operating-mode

IN THIS SECTION

- Syntax | 11844
- Hierarchy Level | 11845
- Description | 11845
- Options | 11845
- Required Privilege Level | 11845
- Release Information | 11845

Syntax

```
operating-mode mode;
```

Hierarchy Level

```
[edit interfaces at-fpc/pic/port dsl-options]
```

Description

For J Series Services Routers only, modify the operating mode of the digital subscriber line for an ATM interface.

Options

mode—Operating mode for ATM-over-ADSL interfaces. The mode can be one of the following:

- `adsl2plus`—Set the ADSL line to train in the ITU G.992.5 mode.
- `ansi-dmt`—Set the ADSL line to train in the ANSI T1.413 Issue 2 mode.
- `auto`—Set the ADSL line to autonegotiate the setting to match the setting of the DSL access multiplexer (DSLAM) located at the central office. The ADSL line trains in the ANSI T1.413 Issue 2 (`ansi-dmt`) or ITU G.992.1 (`itu-dmt`) mode.
- `etsi`—Set the ADSL line to train in the ETSI TS 101 388 V1.3.1 mode.
- `itu-annexb-ur2`—Set the ADSL line to train in the ITU G.992.1 UR-2 mode.
- `itu-annexb-non-ur2`—Set the ADSL line to train in the ITU G.992.1 non-UR-2 mode.
- `itu-dmt`—Set the ADSL line to train in the ITU G.992.1 mode.
- `itu-dmt-bis`—Set the ADSL line to train in the ITU G.992.3 mode.
- **Default:** `auto`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

ATM-over-ADSL Overview

operation

IN THIS SECTION

- Syntax | 11846
- Hierarchy Level | 11846
- Description | 11846
- Options | 11846
- Required Privilege Level | 11847
- Release Information | 11847

Syntax

```
operation operation-name;
```

Hierarchy Level

```
[edit accounting-options mib-profile profile-name]
```

Description

Specify the name of the operation used to collect MIB statistics for an accounting-data log file.

Options

operation-name—Name of the operation to use. You can specify a `get`, `get-next`, or `walk` operation.

- **Default:** `walk`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configure the MIB Profile

optics-options

IN THIS SECTION

- Syntax | 11847
- Hierarchy Level | 11848
- Description | 11849
- Options | 11849
- Required Privilege Level | 11849
- Release Information | 11850

Syntax

```
optics-options

        application
        { hostid id mediaid id [ domainid id ];
        }
 {
    alarm low-light-alarm {
```

```
        (link-down | syslog);
        }
        encoding (differential | non-differential);
        fec (hgfec | sdfec | sdfec15 | sdfec25);
        high-polarization;
        host-side-fec (off | on | not-supported);
        high-power-mode;
        (is-ma| no-is-ma);
        ( lane| lane-all);
        (laser-enable | no-laser-enable);
        loopback loopbacktype media-side-out;
     loopback loopbacktype media-side-in;
     loopback loopbacktype host-side-in;
     loopback loopbacktype host-side-out;

        los-alarm-threshold dBm;
        los-warning-threshold dBm;
        low-power-mode;
        media-side-fec (off | on |unsupported);
        modulation-format (16qam | 8qam | qpsk);
        signal-degrade {
            ber-threshold-clear ber-threshold-clear;
            ber-threshold-signal-degrade ber-threshold-signal-degrade;
            interval milliseconds;
            q-threshold-signal-degrade dB;
            q-threshold-signal-degrade-clear dB;
        }
    tca tca-identifier (enable-tca | no-enable-tca) (threshold number | threshold-24hrs number);
    tx-power dbm;
    warning low-light-warning {
        (link-down | syslog);
    }
    wavelength nm;
    loopback (local);
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

For 10-Gigabit Ethernet, 100-Gigabit Ethernet, or 400-Gigabit Ethernet ZR/ZR+ dense wavelength-division multiplexing (DWDM) interfaces only, configure full C-band International Telecommunication Union (ITU)-Grid tunable optics.

Options

- `application hostid id mediaid id [ domainid id ]` —Displays host ID, media ID of a selected application.
- `host-side-fec`—Enable or disable the FEC on the host side optics.
- `high-power-mode`

  — Configures a port in high-power mode.
- `loopback` —Configures the media side output loopback by default.
- `loopback loopbacktype <loopbacktype>` —Configures the loopback with an option. You can choose any of the following loopback options:
  - media-side-out
  - media-side-in
  - host-side-in
  - host-side-out
- `los-alarm-threshold`—Specify the Loss of Signal (LOS) threshold above which an alarm is raised.
- `los-warning-threshold`—Specify the Loss of Signal (LOS) threshold above which a warning is issued.
- `low-power-mode`—Forces the optics to be on low power mode. By default, it is disabled. On channelized interfaces, you can configure the `low-power-mode` on channel 0 only.
- `media-side-fec`—Enable of disable the FEC on media side optics.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`alarm` option and `warning` options introduced in Junos OS Release 10.0.

Statement and `tx-power` option introduced in Junos OS Release 13.2 for PTX Series routers.

`tca` option introduced in Junos OS Release 14.2 for PTX Series routers.

`los-alarm-threshold` and `los-alarm-warning` options introduced in Junos OS Release 14.2.

`loopback` option introduced in Junos OS Release 19.2R1 for QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 routers.

`lane`, `lane-allhost-side-fec`, and `media-side-fec` options introduced in Junos OS Evolved Release 21.1R1

`local` loopback option introduced in Junos OS Evolved Release 21.2R1 for PTX10001-36MR.

`media-side-in`, `host-side-in`, `host-side-out` loopback types introduced in 23.1R1 for PTX10001-36MR.

`application hostid <> mediaid <> [ domainid <> ]` introduced in 23.1R1 for PTX10001-36MR.

optimal-route-reflection

IN THIS SECTION

- Syntax | **11850**
- Hierarchy Level | **11851**
- Description | **11851**
- Options | **11851**
- Required Privilege Level | **11851**
- Release Information | **11851**

Syntax

```
optimal-route-reflection {
```

```
    export policy;
    igp-backup igp-backup;
    igp-primary igp-primary;
   }
```

Hierarchy Level

```
[edit protocols bgp group]
```

Description

Enable optimal route reflection for BGP group.

Options

export	Optimal route reflection export policy.
igp-backup	Backup node identifier for this client group.
igp-primary	Primary node identifier for this client group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

show bgp group | 17153

optimize-adaptive-teardown

IN THIS SECTION

- Syntax | 11852
- Hierarchy Level | 11852
- Description | 11852
- Options | 11852
- Required Privilege Level | 11853
- Release Information | 11853

Syntax

```
optimize-adaptive-teardown {
    p2p;
    delay;
    timeout;
}
```

Hierarchy Level

```
[edit protocols mpls]
```

Description

Uses RPD infrastructure to determine when all the routes using the old LSP instance have fully shifted to the new LSP instance after MBB switchover. When this statement is configured, the `optimize-hold-dead-delay` statement, which delays the teardown of the old LSP instance after MBB switchover, is ignored.

Options

p2p Only point-to-point LSPs configured in the system will be affected.

`delay` *delay-value* Delays the tearing down of old optimized LSP paths based on the configured value.

- **Range:** from 3 through 65535 seconds
- **Default:** 10 seconds

> **NOTE**: The delay CLI statement is introduced from Junos OS Release 17.2 onwards.

`timeout` *timeout* Timeout for adaptive teardown to clean up LSP

- **Range:** from 1 through 65535 seconds
- **Default:** 600 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R1.

RELATED DOCUMENTATION

Optimizing Signaled LSPs

Achieving a Make-Before-Break, Hitless Switchover for LSPs

optimize-aggregate-frr

IN THIS SECTION

Syntax | 11854

Hierarchy Level | 11854

Description | 11854
Required Privilege Level | 11854

Syntax

```
aggregated-ether-options {
   minimum-links 2;
   optimize-aggregate-frr;
   link-speed 10g;
   lacp {
      active;
   }
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-ether-options],
[edit interfaces name aggregated-ether-options]
```

Description

Enable Aggregate fast reroute optimization

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

optimize-aggressive

IN THIS SECTION

- Syntax | 11855
- Hierarchy Level | 11855
- Description | 11855
- Default | 11855
- Required Privilege Level | 11855
- Release Information | 11856

Syntax

```
optimize-aggressive;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

If enabled, the LSP reoptimization is based solely on the IGP metric. The reoptimization process ignores the available bandwidth ratio calculations, the least-fill 10 percent congestion improvement rule, and the hop-counts rule. This statement makes reoptimization more aggressive than the default.

Default

Aggressive optimization is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Optimizing Signaled LSPs

optimize-hold-dead-delay

IN THIS SECTION

- Syntax | 11856
- Hierarchy Level | 11856
- Description | 11857
- Options | 11857
- Required Privilege Level | 11857

Syntax

```
optimize-hold-dead-delay seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switch-path lsp-name],
[edit protocols mpls],
[edit protocols mpls label-switch-path lsp-name]
```

Description

Allows you to specify the amount of time to delay the tear down of old paths after the router has switched traffic to new optimized paths. This delay timer starts when the timer specified by the `optimize-switchover-delay` statement has elapsed, which is typically 30 seconds, and at the start of the next retry sequence (in other words, the delay is not an absolute countdown of the seconds configured here).

You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The specified delay helps to ensure that old paths are not torn down before all routes have been switched over to the new optimized paths.

Options

seconds Configure the time in seconds to wait before tearing down the old paths that were in use prior to the last LSP optimization.

- **Default:** 60 to 90 seconds
- **Range:** 0 through 65,535 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Optimizing Signaled LSPs

optimize-switchover-delay | **11857**

optimize-timer (Protocols MPLS) | **11859**

optimize-switchover-delay

IN THIS SECTION

Syntax | **11858**

- Hierarchy Level | 11858
- Description | 11858
- Options | 11858
- Required Privilege Level | 11858
- Release Information | 11859

Syntax

```
optimize-switchover-delay seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Delays the switch over of LSPs to newly optimized paths. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The specified delay helps to ensure that the new optimized paths have been established before traffic is switched over from the old paths.

Options

seconds Configure the time in seconds to wait before switching LSPs to newly optimized paths.

- **Default:** 1 second
- **Range:** 1 through 900 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1R1.

RELATED DOCUMENTATION

Optimizing Signaled LSPs

optimize-hold-dead-delay | **11856**

optimize-timer (Protocols MPLS) | **11859**

optimize-timer (Protocols MPLS)

IN THIS SECTION

- Syntax | **11859**
- Hierarchy Level | **11859**
- Description | **11860**
- Default | **11860**
- Options | **11860**
- Required Privilege Level | **11860**
- Release Information | **11860**

Syntax

```
optimize-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
```

```
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

Enable periodic reoptimization of an LSP that is already set up. If topology changes occur, an existing path might become suboptimal, and a subsequent recomputation might be able to determine a better path.

> **NOTE**: This feature is useful only on LSPs for which constrained-path computation is enabled; that is, for which the `no-cspf` statement is not configured. Also, you only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers).

To avoid extensive resource consumption that might result because of frequent path recomputations, or to avoid destabilizing the network as a result of constantly changing LSPs, we recommend that you either leave the timer value sufficiently large or disable the timer value.

Default

The optimize timer is disabled.

Options

seconds—Length of the optimize timer, in seconds.

- **Range:** 0 through 65,535 seconds
- **Default:** 0 seconds (the optimize timer is disabled)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Optimizing Signaled LSPs

optimize-timer (Protocols RSVP)

IN THIS SECTION

- Syntax | 11861
- Hierarchy Level | 11861
- Description | 11861
- Options | 11862
- Required Privilege Level | 11862
- Release Information | 11862

Syntax

```
optimize-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit protocols rsvp interface interface-name link-protection]
```

Description

Configure an optimize timer for a bypass LSP. The optimize timer initiates a periodic optimization process that reshuffles data LSPs among bypass LSPs to achieve the most efficient use of network resources. The optimization process attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all bypasses, or both.

Options

seconds—Specify the number of seconds between optimizations.

- **Range:** 0 through 65,535 seconds
- **Default:** 0 (disabled)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Optimization Interval for Bypass LSPs

option-15 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 11863
- Hierarchy Level | 11863
- Description | 11864
- Required Privilege Level | 11864
- Release Information | 11865

Syntax

```
option-15 {
    default-action {
        drop drop;
        forward-only forward-only;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
        }
    }
    not-present {
        drop drop;
        forward-only forward-only;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
        }
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit forwarding-options dhcp-relay dhcpv6 relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name
```

```
relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit logical-systems name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay dhcpv6 group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay dhcpv6 relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dhcpv6 group name
relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dhcpv6
group name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dhcpv6
relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay dhcpv6 group name relay-
option],
[edit logical-systems name vlans name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name
relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit routing-instances name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 group name relay-
option],
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit vlans name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit vlans name forwarding-options dhcp-relay dhcpv6 relay-option]
```

Description

Specify that the payload of the Option 15 (Vendor Class Identifier) from the client PDU is concatenated with the username during the subscriber authentication or client authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address

option-16 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 11865
- Hierarchy Level | 11866
- Description | 11867
- Required Privilege Level | 11867
- Release Information | 11867

Syntax

```
option-16 {
    default-action {
        drop drop;
        forward-only forward-only;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
        }
```

```
    }
    not-present {
        drop drop;
        forward-only forward-only;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
        }
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit forwarding-options dhcp-relay dhcpv6 relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name
relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit logical-systems name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay dhcpv6 group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay dhcpv6 relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dhcpv6 group name
relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dhcpv6
group name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay dhcpv6
relay-option],
```

```
[edit logical-systems name vlans name forwarding-options dhcp-relay dhcpv6 group name relay-
option],
[edit logical-systems name vlans name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dhcpv6 group name
relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay dhcpv6 relay-
option],
[edit routing-instances name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 group name relay-
option],
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 relay-option],
[edit vlans name forwarding-options dhcp-relay dhcpv6 group name relay-option],
[edit vlans name forwarding-options dhcp-relay dhcpv6 relay-option]
```

Description

Specify that the payload of the Option 16 (Vendor Class Identifier) from the client PDU is concatenated with the username during the subscriber authentication or client authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address

option-16 (DHCPv6 Snooping)

IN THIS SECTION

- Syntax | 11868
- Hierarchy Level | 11868
- Description | 11868
- Options | 11868
- Required Privilege Level | 11869
- Release Information | 11869

Syntax

```
option-16 {
    use-string string;
}
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options]
```

Description

Configure the DHCPv6 Vendor ID option (option 16) to be included in DHCPv6 requests from clients before forwarding them to a DHCPv6 server. Option 16 provides the server with information about the vendor that manufactured the hardware on which the DHCP client is running. When configured, the switch will overwrite any existing option 16 information sent by clients in the DHCPv6 packets.

Option 16 is the DHCPv6 equivalent of the `vendor-id` sub-option of DHCP option 82.

Options

use-string *string*	Define a custom string to be used as the DHCPv6 vendor identifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10 .

RELATED DOCUMENTATION

Enabling Dynamic ARP Inspection (ELS)

Configuring IP Source Guard (ELS)

Setting Up DHCP Option 82 on the Switch with No Relay (ELS)

Configuring Static DHCP IP Addresses for DHCP snooping (ELS)

option-18 (DHCPv6 Snooping)

IN THIS SECTION

- Syntax | 11869
- Hierarchy Level | 11870
- Description | 11870
- Options | 11870
- Required Privilege Level | 11871
- Release Information | 11871

Syntax

```
option-18 {
    prefix {
        host-name;
```

```
        logical-system-name;
        routing-instance-name;
        vlan-id;
        vlan-name;
    }
    use-interface-index (device | logical);
    use-interface-description (device | logical);
    use-interface-mac;
    use-interface-name (device | logical);
    use-string string;
}
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options]
```

Description

Configure the DHCPv6 Relay Agent Interface ID option (option 18) to insert information in DHCPv6 requests from clients before forwarding them to a DHCPv6 server. Option 18 provides information about the port on which the request was received, which the server can use to assign IP addresses, prefixes, and other configuration parameters for the client.

When option 18 is configured, a unique interface ID is inserted into the DHCPv6 packet headers. Suboptions can be configured to include a prefix with the interface ID or to change the type of information used to identify the interface. The default fields included in option 18 are the Juniper Enterprise ID, VLAN, and MAC address of the interface.

Option 18 is the DHCPv6 equivalent of the `circuit-id` sub-option of DHCP option 82.

> **NOTE**: DHCPv6 packets that already contain option 18 information when received from a client are dropped by the switch.

Options

use-interface-mac	Use the MAC address of the interface in the DHCPv6 interface ID.
use-string *string*	Use a custom string in the DHCPv6 interface ID.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

no-option18 | 11538

no-dhcpv6-options | 11440

option-18 (Interface-ID for DHCPv6 Autosense VLANs)

IN THIS SECTION

- Syntax | 11871
- Hierarchy Level | 11872
- Description | 11872
- Required Privilege Level | 11872
- Release Information | 11872

Syntax

```
option-18;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication username-include]
```

Description

Specify that Option 18 (Interface-ID) information received in the innermost DHCPv6 Relay-Forward message header is concatenated with the username during the subscriber authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets

Creating Unique Usernames for DHCP Clients

Using DHCP Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs

option-37 (DHCPv6 Snooping)

IN THIS SECTION

- Syntax | 11873
- Hierarchy Level | 11873
- Description | 11873

- Options | 11874
- Required Privilege Level | 11874
- Release Information | 11874

Syntax

```
option-37 {
    prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
        vlan-id;
        vlan-name;
    }
    use-interface-index (device | logical);
    use-interface-description (device | logical);
    use-interface-mac;
    use-interface-name (device | logical);
    use-string string;
}
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options]
```

Description

Configure the DHCPv6 Relay Agent Remote ID option (option 37) to insert information in DHCPv6 requests from clients before forwarding them to a DHCPv6 server. Option 37 provides information about the remote host, which the server can use to assign IP addresses, prefixes, and other configuration parameters for the client.

When option 37 is configured, a unique remote ID is inserted into the DHCPv6 packet headers. Suboptions can be configured to include a prefix with the remote ID or to change the interface portion of the ID. The default fields included in option 37 are the Juniper Enterprise ID, VLAN, and MAC address of the interface.

Option 37 is the DHCPv6 equivalent of the `remote-id` sub-option of DHCP option 82.

> **NOTE**: DHCPv6 packets that already contain option 37 information when received from a client are dropped by the switch.

Options

`use-interface-mac`	Use the MAC address of the interface in the DHCPv6 remote ID.
`use-string` *`string`*	Use a custom string in the DHCPv6 remote ID.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10 .

RELATED DOCUMENTATION

- no-option37 | **11540**
- no-dhcpv6-options | **11440**
- *Setting Up DHCP Option 82 on the Switch with No Relay (ELS)*
- *Configuring Static DHCP IP Addresses for DHCP snooping (ELS)*

option-37 (Relay Agent Remote-ID for DHCPv6 Autosense VLANs)

IN THIS SECTION

- Syntax | **11875**

- Hierarchy Level | 11875
- Description | 11875
- Required Privilege Level | 11875
- Release Information | 11875

Syntax

```
option-37;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication username-include]
```

Description

Specify that Option 37 (DHCPv6 Relay Agent Remote-ID) information, received in the innermost DHCPv6 Relay-Forward message header, is concatenated with the username during the subscriber authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

> *Configuring VLAN Interface Username Information for AAA Authentication*
>
> *Creating Unique Usernames for DHCP Clients*

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets

Using DHCP Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs

option-60 (DHCP Local Server)

IN THIS SECTION

- Syntax | 11876
- Hierarchy Level | 11876
- Description | 11877
- Required Privilege Level | 11877
- Release Information | 11877

Syntax

```
option-60;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server authentication
username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
```

```
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify that the payload of Option 60 (Vendor Class Identifier) from the client PDU be concatenated with the username during the subscriber authentication or DHCP client authentication process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

option-60 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 11878
- Hierarchy Level | 11879
- Description | 11879
- Required Privilege Level | 11880
- Release Information | 11880

Syntax

```
option-60 {
    default-action {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
    }
    not-present {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay group name relay-option],
[edit forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay relay-option]
```

Description

Specify that the payload of the Option 60 (Vendor Class Identifier) from the client PDU is concatenated with the username during the subscriber authentication or client authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Statement updated in Junos OS Release 17.4R1 for MX Series.

RELATED DOCUMENTATION

Specifying Authentication Support

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address

option-77 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 11880
- Hierarchy Level | 11881
- Description | 11882
- Required Privilege Level | 11882
- Release Information | 11882

Syntax

```
option-77 {
    default-action {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
```

```
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
    }
    not-present {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
    equals {
        ascii name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
        hexadecimal name {
            drop drop;
            forward-only forward-only;
            local-server-group local-server-group;
        }
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay group name relay-option],
[edit forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
```

```
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay relay-option]
```

Description

Specify that the payload of the Option 77 (Vendor Class Identifier) from the client PDU is concatenated with the username during the subscriber authentication or client authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address

option-79

IN THIS SECTION

- Syntax | 11883
- Hierarchy Level | 11883
- Description | 11883
- Required Privilege Level | 11884
- Release Information | 11884

Syntax

```
option-79;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options]
```

Description

Configure the DHCPv6 client link layer option (option 79) to include the source MAC address in DHCPv6 requests from clients before the messages are forwarded to a DHCPv6 server. In a dual-stack network environment in which devices act as both DHCPv4 and DHCPv6 clients, the client MAC address can be used to associate DHCPv4 and DHCPv6 messages with the same client interface. The client MAC address also provides network operators with additional information for event debugging and logging related to the client at the relay agent and the server.

You can configure `option-79` to enable DHCPv6 option 79 for a lightweight DHCPv6 relay agent (LDRA). When DHCPv6 option 79 is enabled, the LDRA reads the source MAC address of DHCPv6 Solicit and DHCPv6 Request messages that it receives from a client. The LDRA encapsulates the Solicit and

Request messages within a DHCPv6 Relay-Forward message, and inserts the client MAC address as option 79 in the Relay-Forward header before relaying the message to the server.

You can also configure DHCPv6 option 79 for a layer 3 DHCPv6 relay agent using the `relay-agent-option-79` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Inserting the DHCPv6 Client MAC Address Option (Option 79) In DHCPv6 Packets

option-82

IN THIS SECTION

- Syntax | **11885**
- Hierarchy Level (EX Series, QFX Series) | **11885**
- Hierarchy Level (MX Series) | **11885**
- Description | **11885**
- Default | **11886**
- Required Privilege Level | **11886**
- Release Information | **11886**

Syntax

```
option-82 {
    circuit-id {
        prefix (host-name | routing-instance-name);
        use-interface-description;
        use-vlan-id;
    }
    remote-id {
        host-name;
        mac (Option 82);
        use-interface-description;
        use-string                          string;
    }
    vendor-id {
        use-string                          string;
    }
}
```

Hierarchy Level (EX Series, QFX Series)

```
[edit vlans vlan-name forwarding-options dhcp-security]
```

Hierarchy Level (MX Series)

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]
```

Description

Have the device insert DHCP option 82 (also known as the DHCP relay agent information option) information in the DHCP request packet header of a DHCP request that it receives from a DHCP client connected to one of its interfaces before it forwards or relays that DHCP request to a DHCP server. The server uses the option 82 information, which provides details about the circuit and host the request came from. However, in formulating the reply, the server does not make any changes to the option 82 information in the packet header. The device receives the reply and then removes the DHCP option 82 information before forwarding the reply to the client.

The remaining statements are explained separately. See CLI Explorer.

Default

Insertion of DHCP option 82 information is not enabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Setting Up DHCP Option 82 on the Switch with No Relay (ELS)

no-option82 | **11541**

Understanding DHCP Option 82

Understanding DHCP Snooping (non-ELS)

http://tools.ietf.org/html/rfc3046

option-82 (interfaces)

IN THIS SECTION

- Syntax | **11887**
- Hierarchy Level | **11887**
- Description | **11887**
- Options | **11887**
- Required Privilege Level | **11887**
- Release Information | **11888**

Syntax

```
option-82 <circuit-id> <remote-id>;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication username-include]
```

Description

Specify that the option 82 information from the client PDU is concatenated with the username during the subscriber authentication process.

For autosense VLANs, you can additionally specify Option 82 suboption information that is concatenated with the username. You can specify either both or neither of the Agent Circuit ID (suboption 1) and Agent Remote ID (suboption 1). If you specify both, the Agent Circuit ID is supplied first, followed by a delimiter, and then the Agent Remote ID. If you specify that neither suboption is supplied, the raw payload of Option 82 from the PDU is concatenated to the username.

NOTE: The option 82 value used in creating the username is based on the option 82 value that is encoded in the incoming DHCP discover packet. The use of suboptions is supported for DHCPv4 only.

Options

none	Use the raw payload of Option 82 from the PDU.
circuit-id	(Optional) Use the Agent Circuit ID suboption (suboption 1).
remote-id	(Optional) Use the Agent Remote ID suboption (suboption 2).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Options `circuit-id` and `remote-id` introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

Using DHCP Option 82 Suboptions in Authentication Usernames for Autosense VLANs

option-82 (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 11888
- Hierarchy Level | 11889
- Description | 11889
- Options | 11889
- Required Privilege Level | 11889
- Release Information | 11889

Syntax

```
option-82 {
    circuit-id value range named-range;
    remote-id value range named-range;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet dhcp-attributes option-match],
[edit access protocol-attributes attribute-set-name option-match]
```

Description

Specify the list of option 82 suboption match criteria used to select the named address range used for the client. The server matches the option 82 value in the user PDU to the specified option 82 match criteria and uses the named address range associated with the string.

Options

circuit-id Configure the address-assignment pool named-range to use for a particular option 82 Agent Circuit ID value.

- Values:
  - *value*—String for the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option (option 82) in DHCP packets.
  - range *named-range*—Name of the address-assignment pool range to use.

remote-id Specify the address-assignment pool named range to use based on the particular option 82 Agent Remote ID value.

- Values:
  - range *named-range*—Name of the address-assignment pool range to use.
  - *value*—String for Agent Remote ID suboption (suboption 2) of the DHCP relay agent information option (option 82) in DHCP packets.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management

option-82 (DHCP Local Server Authentication)

IN THIS SECTION

- Syntax | 11890
- Hierarchy Level | 11890
- Description | 11891
- Options | 11891
- Required Privilege Level | 11891
- Release Information | 11891

Syntax

```
option-82 <circuit-id> <remote-id>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server authentication
username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
```

```
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify the type of Option 82 information from the client PDU that is concatenated with the username during the subscriber authentication or DHCP client authentication process. You can specify either, both, or neither of the Agent Circuit ID and Agent Remote ID suboptions. If you specify both, the Agent Circuit ID is supplied first, followed by a delimiter, and then the Agent Remote ID. If you specify that neither suboption is supplied, the raw payload of Option 82 from the PDU is concatenated to the username.

Options

`circuit-id`—(Optional) Agent Circuit ID suboption (suboption 1).

`remote-id`—(Optional) Agent Remote ID suboption (suboption 2).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

option-82 (DHCP Local Server Pool Matching)

IN THIS SECTION

- Syntax | 11892
- Hierarchy Level | 11892
- Description | 11892
- Required Privilege Level | 11892
- Release Information | 11893

Syntax

```
option-82;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server pool-match-order],
[edit logical-systems logical-system-name system services dhcp-local-server pool-match-order],
[edit routing-instances routing-instance-name system services dhcp-local-server pool-match-order],
[edit system services dhcp-local-server pool-match-order]
```

Description

Configure the extended DHCP local server to use the option 82 value in the DHCP client DHCP PDU together with the ip-address-first method to determine which address-assignment pool to use. You must configure the `ip-address-first` statement before configuring the `option-82` statement. The DHCP local server first determines which address-assignment pool to use based on the ip-address-first method. Then, the local server matches the option 82 value in the client PDU with the option 82 configuration in the address-assignment pool. This statement is supported for IPv4 address-assignment pools only.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use

Understanding Differences Between Legacy DHCP and Extended DHCP

Address-Assignment Pools Overview

option-82 (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 11893
- Hierarchy Level | 11893
- Description | 11894
- Options | 11894
- Required Privilege Level | 11894
- Release Information | 11894

Syntax

```
option-82 <circuit-id> <remote-id>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
```

```
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name authentication username-include]
```

Description

Specify the option 82 that is concatenated with the username during the subscriber authentication or client authentication process. You can specify either, both, or neither the Agent Circuit ID and the Agent Remote ID suboptions. If you specify both, the Agent Circuit ID is supplied first, followed by a delimiter, and then the Agent Remote ID. If neither suboption is supplied, the raw payload of option 82 is concatenated to the username.

> **NOTE**: The option 82 value used in creating the username is based on the option 82 value that is encoded in the outgoing (relayed) PDU.

Options

`circuit-id`—(Optional) The string for the Agent Circuit ID suboption (suboption 1).

`remote-id`—(Optional) The string for the Agent Remote ID suboption (suboption 2).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

option (DHCP server)

IN THIS SECTION

- Syntax | 11895
- Hierarchy Level | 11895
- Description | 11895
- Options | 11896
- Required Privilege Level | 11896
- Release Information | 11896

Syntax

```
option {
    [ (id-number option-type option-value) | (id-number array option-type option-value) ];
}
```

Hierarchy Level

```
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]
```

Description

Configure one or more user-defined options that are not included in the Junos default implementation of the DHCP server. For example, if a client requests a DHCP option that is not included in the DHCP server, you can create a user-defined option that enables the server to respond to the client's request.

Options

- *id-number*—Any whole number. The ID number is used to index the option and must be unique across a DHCP server.
- *option-type*—Any of the following types: `byte`, `byte-stream`, `flag`, `integer`, `ip-address`, `short`, `string`, `unsigned-integer`, `unsigned-short`.
- `array`—An option can include an array of values.
- *option-value*—Value associated with an option. The option value must be compatible with the option type (for example, an `On` or `Off` value for a `flag` type).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 .

option-match

IN THIS SECTION

- Syntax | 11897
- Hierarchy Level | 11897
- Description | 11897
- Required Privilege Level | 11897
- Release Information | 11897

Syntax

```
option-match {
    option-82 {
        circuit-id value range named-range;
        remote-id value range named-range;
    }
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet dhcp-attributes],
[edit access protocol-attributes attribute-set-name]
```

Description

Specify a list of match criteria used to determine which named address range in the address-assignment pool to use. The extended DHCP local server matches this information to the match criteria specified in the client PDUs. For example, for option 82 match criteria, the server matches the option 82 value in the user PDU to the specified option 82 string and uses the named range associated with the string.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management

option-number (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 11898
- Hierarchy Level | 11898
- Description | 11898
- Options | 11899
- Required Privilege Level | 11899
- Release Information | 11899

Syntax

```
option-number option-number;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay dhcpv6 relay-option],
[edit forwarding-options dhcp-relay group group-name relay-option],
[edit forwarding-options dhcp-relay dhcpv6 group group-name relay-option],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify the DHCP option DHCP relay agent uses for selective processing of client traffic. You can configure support globally or for a named group of interfaces. You can also configure support for the extended DHCP relay agent on a per logical system and per routing instance basis.

Use the `[edit forwarding-options dhcp-relay dhcpv6]` hierarchy level to configure the DHCPv6 relay agent support.

Options

option-number—The DHCP or DHCPv6 option in the incoming traffic.

> **NOTE**: EX Series switches do not support the User Class Options.

- 15 (DHCPv6 only)—Use DHCPv6 option 15 (User Class Option) in packets
- 16 (DHCPv6 only)—(MX Series routers and EX Series switches only) Use DHCPv6 option 16 (Vendor Class Option) in packets
- 60 (DHCPv4 only)—(MX Series routers and EX Series switches only) Use DHCP option 60 (Vendor Class Identifier) in DHCP packets
- 77 (DHCPv4 only)—Use DHCP option 77 (User Class Identifier) in packets

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic

option-order (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 11900
- Hierarchy Level | 11900
- Description | 11901

- Options | 11901
- Required Privilege Level | 11901
- Release Information | 11901

Syntax

```
option-order name;
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay group name relay-option],
[edit forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
```

```
[edit routing-instances name forwarding-options dhcp-relay relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay relay-option]
```

Description

Options precedence order

Options

name

Option number

- Values:
  - `60`—Option 60
  - `77`—Option 77

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

DHCPv4 and DHCPv6 Forward-Only Action for Relay Traffic with Unknown DHCP Server Address

option-refresh-rate

IN THIS SECTION

- Syntax | 11902
- Hierarchy Level | 11902
- Description | 11902
- Options | 11902
- Required Privilege Level | 11903
- Release Information | 11903

Syntax

```
option-refresh-rate packets packets seconds seconds;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name],
```

Description

Specify the frequency at which the flow generator sends updates about template options, such as sampling rate, to the flow collector. Specify the refresh rate, in either the number of packets or seconds.

Options

packets—Refresh rate, in number of packets.

- **Range:** 1 through 480,000
- **Default:** 4800

seconds—Refresh rate, in number of seconds.

- **Range:** 10 through 600
- **Default:** 600

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level added in Junos OS Release 10.2.

Support at the `[edit services flow-monitoring version9 template template-name]` hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.

RELATED DOCUMENTATION

Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates

Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices

option-refresh-rate (Services)

IN THIS SECTION

- Syntax | 11904
- Hierarchy Level | 11904
- Description | 11904
- Options | 11904
- Required Privilege Level | 11904
- Release Information | 11904

Syntax

```
option-refresh-rate
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

Specifies the frequency at which the flow generator needs to send the Jflow options template (both template and system scope) to the collector.

Option refresh rate can be configured as number of packets. If the number of packets configured is “n”, it indicates that after every export of “n” data packets to the collector, the options template (both template and system scope) should be exported to the collector.

Option refresh rate can also be configured in seconds. If the seconds configured is “n”, it indicates that after every “n” seconds, the options template (both template and system scope) should be exported to the collector.

The options template provides information such as flow active timeout, flow inactive timeout, sampling interval, and sampling algorithm.

Options

- `packets`—Specify the number of packets. The range is from 1 through 480,000.
- `seconds`—Specify the number of seconds. The range is from 10 through 600.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Flow Aggregation to Use Version 9 Flow Templates

option (Security IDP)

IN THIS SECTION

- Syntax | 11905
- Hierarchy Level | 11905
- Description | 11905
- Options | 11905
- Required Privilege Level | 11906
- Release Information | 11906

Syntax

```
option {
    match (equal | greater-than | less-than | not-equal);
    value tcp-option;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify the TCP option type (kind field in the TCP header).

Options

`match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.

`value` *`tcp-option`*—Match the option value.

- **Range:** 0 through 255

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

option-type

IN THIS SECTION

- Syntax | 11906
- Hierarchy Level | 11907
- Description | 11907
- Options | 11907
- Required Privilege Level | 11907
- Release Information | 11907

Syntax

```
option-type {
    match (equal | greater-than | less-than | not-equal);
    value header-value;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol ipv6 extension-
header destination-option]
```

Description

Specify the type of option for destination header type.

Options

match (equal | greater-than | less-than | not-equal)	Match an operand.
value	Match a decimal value. • **Range:** 0 through 255

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

optional (Commit Scripts)

IN THIS SECTION

- Syntax | 11908
- Hierarchy Level | 11908
- Description | 11908
- Default | 11908

- Required Privilege Level | 11908
- Release Information | 11908

Syntax

```
optional;
```

Hierarchy Level

```
[edit system scripts commit file filename]
```

Description

Allow a commit operation to succeed even if the commit script specified in the `file` statement is missing from the appropriate commit script directory on the device.

> **NOTE**: Commit scripts are stored on a device's hard disk in the **/var/db/scripts/commit** directory or on the flash drive in the **/config/scripts/commit** directory. On the QFabric system, commit scripts are stored in the **/pbdata/mgd_shared/*partition-ip*/var/db/scripts/commit/** directory on the Director device.

Default

If you do not include the `optional` statement for a commit script, the commit operation fails if that script is absent from the commit script directory.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Controlling Execution of Commit Scripts During Commit Operations

optional (Event Scripts)

IN THIS SECTION

- Syntax | 11909
- Hierarchy Level | 11909
- Description | 11909
- Default | 11910
- Required Privilege Level | 11910
- Release Information | 11910

Syntax

```
optional;
```

Hierarchy Level

```
[edit event-options event-script]
```

Description

For Junos OS event scripts, globally indicate the physical presence of the scripts as optional to allow commits to proceed if they are missing. Scripts can be missing, for example, if the primary storage on the device running Junos OS becomes corrupted.

If you do not configure the `optional` statement, the event process (eventd) validates the physical existence of event scripts during the commit process. If a configured script is not physically accessible, the commit fails.

If you configure the `optional` statement, the event process (eventd) ignores missing event scripts during the commit process and proceeds with the commit. If an event script is present but invalid, the commit fails.

Default

By default, optional is not set. Event scripts are required to be accessible during commit operations.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

source (Event Scripts) | 28206

event-script (Event Options) | 6533

optional-parameters

IN THIS SECTION

- Syntax | 11911
- Hierarchy Level | 11911
- Description | 11911
- Options | 11911
- Required Privilege Level | 11911
- Release Information | 11911

Syntax

```
optional-parameters parameter-name parameter-value;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature]
```

Description

Configure the Hyperscan optional parameters to enhance the pattern matching process.

Options

max-offset Maximum offset in the data stream at which the pattern match ends.

min-length Minimum match length required to successfully match the pattern.

min-offset Minimum offset in the data stream at which the pattern match ends.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Attack Objects and Object Groups for IDP Policies

IDP Signature Database Overview

options (Access Profile)

IN THIS SECTION

- Syntax | 11912
- Hierarchy Level | 11914
- Description | 11914
- Options | 11914
- Required Privilege Level | 11921
- Release Information | 11921

Syntax

```
options {
    accounting-session-id-format (decimal | description);
    calling-station-id-delimiter delimiter-character;
    calling-station-id-format {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        nas-identifier;
    }
    chap-challenge-in-request-authenticator;
    client-accounting-algorithm (direct | round-robin);
    client-authentication-algorithm  (direct | round-robin);
    coa-dynamic-variable-validation;
    ethernet-port-type-virtual;
    interface-description-format {
        exclude-adapter;
        exclude-channel;
        exclude-sub-interface;
    }
    ip-address-change-notify message;
    juniper-access-line-attributes;
    nas-identifier identifier-value;
    nas-port-extended-format {
        adapter-width width;
```

```
        ae-width width;
        port-width width;
        slot-width width;
        stacked-vlan-width width;
        vlan-width width;
        atm {
            adapter-width width;
            port-width width;
            pw-width width;
            slot-width width;
            vci-width width;
            vpi-width width;
        }
    }
    nas-port-id-delimiter delimiter-character;
    nas-port-id-format {
        agent-circuit-id;
        agent-remote-id;
        concatenated-vlan-tags  {
            fixed-size-inner-tag;
            fixed-size-outer-tag;
        }
        interface-description;
        interface-text-description;
        nas-identifier;
        order {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            interface-text-description;
            nas-identifier;
            postpend-vlan-tags;
        }
        postpend-vlan-tags;
    }
    nas-port-type {
        ethernet {
            port-type;
        }
    }
    override {
        calling-station-id remote-circuit-id;
        nas-ip-address tunnel-client-gateway-address;
```

```
        nas-port tunnel-client-nas-port;
        nas-port-type tunnel-client-nas-port-type;
    }
    remote-circuit-id-delimiter;
    remote-circuit-id-fallback;
    remote-circuit-id-format {
        agent-circuit-id;
        agent-remote-id;
    }
    revert-interval interval;
    service-activation {
        dynamic-profile (optional-at-login | required-at-login);
        extensible-service (optional-at-login | required-at-login);
    }
    vlan-nas-port-stacked-format;
}
```

Hierarchy Level

```
[edit access profile profile-name radius]
```

Description

Configure the options used by RADIUS authentication and accounting servers.

Options

accounting-session-id-format

(EX Series, MX Series only) Configure the format the router or switch uses to identify the accounting session. The default is `decimal`.

- Values:
  - decimal—Use the decimal format.
  - description—Use the generic format, in the form: `jnpr` *`interface-specifier`*:*`subscriber-session-id`*.

calling-station-id-delimiter

(MX Series, T Series only) Starting in Junos OS Release 13.1, specify the character that the router uses as a separator between the concatenated values in the Calling-Station-ID (RADIUS IETF attribute 31) string. The router uses the delimiter when

you configure more than one value in the `calling-station-id-format` statement. The default is the hash (#) character.

- Values:
  - *delimiter-character*—Character to use for the delimiter. You must enclose the delimiter character in quotation marks (“ ”).

chap-challenge-in-request-authenticator

(MX Series only) Starting in Junos OS Release 15.1, configure the `authd` process to insert the random challenge generated by the NAS into the Request Authenticator field of Access-Request packets, if the challenge value is 16 bytes long. If you enable the `chap-challenge-in -request-authenticator` statement and the random challenge is not 16 bytes long, `authd` ignores the statement and uses the default behavior, which inserts the random challenge as the CHAP-Challenge attribute (RADIUS attribute 60) in Access-Request packets.

client-accounting-algorithm

(EX Series, MX Series, SRX3xx and SRX550HMonly) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, configure the access method the router uses to access RADIUS accounting servers. The default is the `direct` option. The default behaviour applies for devices which do not support this configuration option.

- Values:
  - direct—Use the direct method.
  - round-robin—Use the round-robin method.

client-authentication-algorithm

(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, configure the method that the authenticator uses to access RADIUS authentication servers when there are multiple servers configured. Initially, a RADIUS client sends a request to a RADIUS authentication or accounting server. The router or switch, acting as the authenticator, waits for a response from the server before sending another request.

When there are multiple RADIUS server connections configured for a client, the authenticator attempts to reach the different servers in the order that they are configured. If there is no response from the first RADIUS server, the authenticator attempts to reach the next RADIUS server. This process repeats until the client is either granted access or there are no more configured servers.

If the `direct` method is configured, the authenticator always treats the first server in the list as the primary server. The authenticator moves on to the second server only if the attempt to reach the first server fails. If the `round-robin` method is configured, the server chosen first will be rotated based on which server was used last. The first server in the list is treated as a primary for the first authentication request, but for

the second request, the second server configured is treated as primary, and so on. With this method, all of the configured servers receive roughly the same number of requests on average so that no single server has to handle all of the requests.

> **NOTE**: The `round-robin` access method is not supported on EX Series switches.

- **Default:** The default is the `direct` option.
- Values:
  - direct—Use the direct access method. The authenticator contacts the first RADIUS server on the list for each request, the second server if the first one fails, and so on.
  - round-robin—Use the round-robin method. The authenticator contacts the first RADIUS server for the first request, the second server for the second request, and so on.

coa-dynamic-variable-validation

(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, specify that when a CoA operation includes a change to a client profile dynamic variable that cannot be applied (such as an update to a non-existent filter), the router does not apply any changes to client profile dynamic variables in the request, and responds with a NACK message.

- **Default:** If you do not configure this statement, the router does not apply any incorrect variable updates, but does make any other changes to the client profile dynamic variables, and responds with an ACK message.

ethernet-port-type-virtual

(EX Series, M Series, MX Seriesonly) Specify the physical port type the router or switch uses to authenticate clients. The router or switch passes a port type of `ethernet` in RADIUS attribute 61 (NAS-Port-Type) by default. This statement specifies a port type of `virtual`.

> **NOTE**: This statement takes precedence over the `nas-port-type` statement if you include both statements in the same access profile. `client-authentication-algorithm`

access-loop-id-local

Specify that the Agent-Remote-Id and Agent-Circuit-Id are generated locally when these values are not present in the client database.

ip-address-change-notify

(MX Series only) Starting in Junos OS Release 13.1, for on-demand address allocation for dual-stack PPP subscribers, specify that the BNG includes the IPv4-Release-Control VSA (26–164) in the Access-Request that is sent during on-demand IP address allocation and in the Interim-Accounting messages that are sent to report an address change. The configuration of this statement has no effect when on-demand IP address allocation or deallocation is not configured.

Optionally, configure a message that is included in the VSA when it is sent to the RADIUS server.

- **Default:** This functionality is disabled by default.
- **Values:** ***message***—VSA message.
- **Range:** Up to 32 characters.

juniper-access-line-attributes

Configure AAA to add Juniper Networks access line VSAs to the RADIUS authentication and accounting request messages for subscribers. If the router has not received and processed the corresponding ANCP attributes from the access node, then AAA provides only the following in these RADIUS messages:

- Downstream-Calculated-QoS-Rate (IANA 4874, 26-141)—Default configured advisory transmit speed.
- Upstream-Calculated-QoS-Rate (IANA 4874, 26-142)—Default configured advisory receive speed.

> **NOTE**: Starting in Junos OS Release 19.2R1, the `juniper-access-line-attributes` option replaces the `juniper-dsl-attributes` option. The difference between these options is that `juniper-dsl-attributes` supported only DSL TLVs received in the ANCP Port Status message. The `juniper-access-line-attributes` option supports PON TLVs in addition to DSL TLVs, and will be extensible to future access technologies.
>
> For backward compatibility with existing scripts, the `juniper-dsl-attributes` option redirects to the new `juniper-access-line-attributes` option. We recommend that you use `juniper-access-line-attributes`.

> **NOTE**: The `juniper-access-line-attributes` option is not backward compatible with Junos OS Release 19.1 or earlier releases. This means that if you have configured `juniper-access-line-attributes` option in Junos OS Release 19.2 or higher releases, you must

perform the following steps to downgrade to Junos OS Release 19.1 or earlier releases:

1. Delete the `juniper-access-line-attributes` option from all access profiles that include it.
2. Perform the software downgrade.
3. Add the `juniper-dsl-attributes` option to the affected access profiles.

- **Default:** The Juniper Networks access line VSAs are not added to the RADIUS authentication and accounting request messages. However, the DSL Forum VSA —if available—is added to RADIUS messages by default.

nas-identifier

(EX Series, MX Series, SRX Series only) Configure the value for the client RADIUS attribute 32 (NAS-Identifier). This attribute is used for authentication and accounting requests. This statement was introduced in Junos OS Release 15.1X49-D110 for SRX300, SRX320, SRX340, SRX345, and SRX550M Series devices.

- **Values:** *identifier-value*—String to use for authentication and accounting requests.
- **Range:** 1 through 64 characters.

nas-port-id-delimiter

(MX Series only) Starting in Junos OS Release 11.4, specify the character that the router uses as a separator between the concatenated values in the NAS-Port-ID string. The router uses the delimiter when you configure more than one value in the `nas-port-id-format` statement. The default is the hash (`#`) character. This statement was introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

- **Values:** *delimiter-character*—Character used for the delimiter.

remote-circuit-id-delimiter

(MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure a delimiter character for the remote circuit ID string when you use the `remote-circuit-id-format` statement to configure the string to use instead of the Calling-Station ID in L2TP Calling Number AVP 22. If more than one value is configured for the remote circuit ID format, the delimiter character is used as a separator between the concatenated values in the resulting remote circuit ID string. The default is the hash (`#`) character.

- **Values:** *delimiter*—Delimiter character to be used between components of the remote circuit ID string.

remote-circuit-id-fallback

(MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure the fallback value for the LAC to send in L2TP Calling Number AVP 22, either the configured Calling-Station-ID or the default underlying interface. Use of the fallback value is triggered when the components of the override string you configured with the `remote-circuit-id-format` statement—the ACI, the ARI, or both ACI and ARI—are not received by the LAC in the PPPoE Active Discovery Request (PADR) packet.

- Values:
  - configured-calling-station-id—Send the configured Calling-Station-ID in the Calling Number AVP.
  - default—Send the underlying interface value in the Calling Number AVP.

remote-circuit-id-format

(MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure the format of the string that overrides the Calling-Station-ID format in the Calling Number AVP 22 sent by the LAC to the LNS in the ICRQ packet when an L2TP session is being established. You can specify the ACI, the ARI, or both the ACI and ARI. This statement enables you to decouple the AVP 22 value from the RADIUS Calling-Station-ID attribute (31); the values for AVP 22 and the Calling-Station-ID attribute are the same when you use the `calling-station-id-format` statement to configure AVP 22.

> **NOTE**: You must configure the `override calling-circuit-id remote-circuit-id` statement for the remote circuit ID format to be used in the calling number AVP.

- Values:
  - agent-circuit-id—Specifies use of the ACI string that uniquely identifies the subscriber's access node and the digital subscriber line (DSL) on the access node. For PPPoE traffic, the ACI string is in the DSL Forum Agent-Circuit-ID VSA [26-1] of PPPoE Active Discovery Initiation (PADI) and PPPoE Active Discovery Request (PADR) control packets.
  - agent-remote-id—Specifies use of the ARI string that identifies the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request. The agent remote identifier (ARI) string is stored in the DSL Forum Agent-Remote-ID VSA [26-2] for PPPoE traffic.

service-activation

(MX Series only) Starting in Junos OS Release 16.2, specify whether subscribers are allowed to log in even when service activation failures related to configuration errors occur during family activation request processing by authd for a newly

authenticated subscriber. Configuration errors include missing or incorrect syntax, missing or incomplete references to dynamic profiles, and missing or incomplete variables.

> **NOTE**: This configuration does not apply to services activated by means of RADIUS CoA requests, JSRC Push-Profile-Request (PPR) messages, or subscriber secure policies.

You can enable separate configurations for subscriber login services for two `service-activation` types: `dynamic-profile` and `extensible-service`. You configure the `dynamic-profile` type services in the dynamic profile at the `[edit dynamic-profiles]` hierarchy level; the profile is used to provide dynamic subscriber access and services for broadband applications. The `extensible-service` type is for business services configured in an operation script and provisioned by the Extensible Subscriber Services Manager daemon (essmd).

- Default:

  Default behavior depends on the service type:

  - For `extensible-service` services: `optional-at-login`.
  - For `dynamic-profile` services: `required-at-login`.

- Values:

  - optional-at-login—Service activation is optional. Failure due to configuration errors does not prevent activation of the address family; it allows subscriber access. Failure for any other reason causes network family activation to fail. If no other network family is already active for the subscriber, then the client application logs out the subscriber.
  - required-at-login—Service activation is required. Failure for any reason causes the Network-Family-Activate-Request for that network family to fail. If no other network family is already active for the subscriber, then the client application logs out the subscriber.

vlan-nas-port-stacked-format (MX Series only) Configure RADIUS attribute 5 (NAS-Port) to include the S-VLAN ID, in addition to the VLAN ID, for subscribers on Ethernet interfaces.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`juniper-dsl-attributes` introduced in Junos OS Release 11.4.

`nas-port-id-delimiter` introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

`calling-station-id-delimiter` introduced in Junos OS Release 13.1.

`ip-address-change-notify` introduced in Junos OS Release 13.1.

`coa-dynamic-variable-validation`, `client-authentication-algorithm`, and `client-accounting-algorithm` introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

`remote-circuit-id-delimiter`, `remote-circuit-id-fallback`, and `remote-circuit-id-format` introduced in Junos OS Release 13.3R1 on MX Series.

`chap-challenge-in-request-authenticator` introduced in Junos OS Release 15.1.

`nas-identifier` introduced in Junos OS Release 15.1X49-D110 for SRX300, SRX320, SRX340, SRX345, and SRX550M Series devices.

`service-activation` introduced in Junos OS Release 16.2.

`juniper-access-line-attributes` introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

RADIUS Authentication and Accounting Basic Configuration

options (Routing Options)

IN THIS SECTION

- Syntax | 11922
- Hierarchy Level | 11922
- Description | 11922
- Options | 11923
- Required Privilege Level | 11923
- Release Information | 11923

Syntax

```
options {
    syslog (level level | upto level level);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],

[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Configure the types of system logging messages sent about the routing protocols process to the system message logging file. These messages are also displayed on the system console. You can log messages at a particular level, or up to and including a particular level.

Options

level ***level***—Severity of the message. It can be one or more of the following levels, in order of decreasing urgency:

- **alert**—Conditions that should be corrected immediately, such as a corrupted system database.
- **critical**—Critical conditions, such as hard drive errors.
- **debug**—Software debugging messages.
- **emergency**—Panic or other conditions that cause the system to become unusable.
- **error**—Standard error conditions.
- **info**—Informational messages.
- **notice**—Conditions that are not error conditions, but might warrant special handling.
- **warning**—System warning messages.

upto level ***level***—Log all messages up to a particular level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Administration Library for Routing Devices

options-template-id

IN THIS SECTION

- Syntax | 11924
- Hierarchy Level | 11924
- Description | 11924
- Options | 11925
- Required Privilege Level | 11925
- Release Information | 11925

Syntax

```
options-template-id id;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Define a unique options template ID to be used for flow aggregation of version 9 and IPFIX flows. After a re-start of the exporting process, the software may re-assign option template IDs.

MX and QFX Series

The options template ID can be a value in the range of 1024 through 65535. If you do not configure values for the options template ID, a default value is assumed for this ID, which is different for the various address families. If you configure the same options template ID value for different address families, such a setting is not processed properly and might cause unexpected behavior. For example, if you configure the same options template ID value for both IPv4 and IPv6, the collector validates the export data based on the options template ID value that it last receives. In this case, if IPv6 is configured after IPv4, the value is effective for IPv6 and the default value is used for IPv4.

PTX Series

If you choose to configure the options template ID, the range is 1024 to 65520. If you do not configure this ID, the default value that is set is in the range 256-511 and is different for each template.

You can configure the `option-template-id` statement for family `inet`, `inet6`, and `mpls` only.

You must not configure the same ID for different templates (option or data template). The operating system does not check to ensure that you do not configure the same ID value for different templates. If you configure the same ID value, such a setting is not processed properly and might cause unexpected behavior.

The options template ID range [configured `options-template-id` value + 20) is reserved and you must not configure any another ID in this range. The difference between configured options template IDs across families should be at least 20; for example, if `options-template-id 1056` is configured for family `inet`, you should not configure template IDs in the range of 1056 to 1075 for any other families.

For Junos OS, if you change the options template ID, all flows are inactively timed out. New flows are learned afresh.

For Junos OS Evolved, if you change the options template ID, this change does not impact the flows.

Options

id—Unique identifier for the options template to be used for version 9 or IPFIX flows. The range 0-255 is reserved and is used for template sets, options template sets, and for future expansion. If you do not configure an options template ID, the software generates a template ID for you.

- **Range:** 1024 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows

Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows

order

IN THIS SECTION

- Syntax | 11926
- Hierarchy Level | 11926
- Description | 11926
- Options | 11927
- Required Privilege Level | 11927
- Release Information | 11927

Syntax

```
order [ accounting-method ];
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Set the order in which the Junos OS tries different accounting methods for client activity. When a client logs in, the software tries the accounting methods in the specified order.

Options

accounting-method—One or more accounting methods. When a client logs in, the software tries the accounting methods in the following order, from first to last. The only valid value is `radius` for RADIUS accounting.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Authentication and Accounting Basic Configuration

order (Application Identification)

IN THIS SECTION

- Syntax | 11928
- Hierarchy Level | 11928
- Description | 11928
- Options | 11928
- Required Privilege Level | 11928
- Release Information | 11928

Syntax

```
order order;
```

Hierarchy Level

```
[edit services application-identification application name address-mapping name],
[edit services application-identification application application-name icmp-mapping],
[edit services application-identification application application-name ip-protocol-mapping],
[edit services application-identification application application-name over protocol-type
signature l4-l7-signature-name],
[edit services application-identification application application-name
```

Description

Define application matching priority. For address configurations, the order number resolves the conflict when multiple address entries are matched for a specific session. The lower number has a higher priority.

Options

order Order sequence number. This value is mandatory and must be unique.

- **Default:** 0
- **Range:** 0 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Support at the [edit services application-identification application *application-name*] hierarchy level introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

order-priority (Application Identification)

IN THIS SECTION

- Syntax | **11929**
- Hierarchy Level | **11929**
- Description | **11930**
- Options | **11930**
- Required Privilege Level | **11930**
- Release Information | **11930**

Syntax

```
order-priority (high | low);
```

Hierarchy Level

```
[edit services application-identification application application-name address-mapping name],
[edit services application-identification application application-name icmp-mapping],
[edit services application-identification application application-name ip-protocol-mapping],
[edit services application-identification application application-name over protocol-type
signature l4-l7-signature-name]
```

Description

Define the priority of signatures when both a custom signature and predefined signature apply to a protocol bundle.

Options

high Custom signatures have priority over predefined signatures.

low Predefined signatures have priority over custom signatures.

- **Default:** high

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

order (Security IDP)

IN THIS SECTION

Syntax | 11931

- Hierarchy Level | 11931
- Description | 11931
- Required Privilege Level | 11931
- Release Information | 11931

Syntax

```
order;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain]
```

Description

Create a compound attack object that must match each member signature or protocol anomaly in the order you specify. If you do not specify an ordered match, the compound attack object still must match all members, but the attacks or protocol anomalies can appear in random order.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

origin-autonomous-system (Origin Validation for BGP)

IN THIS SECTION

- Syntax | 11932
- Hierarchy Level | 11932
- Description | 11932
- Options | 11933
- Required Privilege Level | 11933
- Release Information | 11933

Syntax

```
origin-autonomous-system as-number {
    validation-state (invalid | valid);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation static record destination maximum-length prefix-length],
[edit logical-systems logical-system-name routing-options validation static record destination
maximum-length prefix-length],
[edit routing-instances instance-name routing-options validation static record destination
maximum-length prefix-length],
[edit routing-options validation static record destination maximum-length prefix-length]
```

Description

Configure the legitimate originator autonomous system (AS). This is a required statement.

Options

as-number Configure a legitimate originator AS number.

validation-state (invalid | valid) Configure the validation state for a route validation record.

- **Values:** Configure one of two mutually exclusive validation states:
  - invalid—A negative (invalid) validation state. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.
  - valid—A positive (valid) validation state. Indicates that the prefix and AS pair are found in the database.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Origin Validation for BGP

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

origin (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 11934
- Hierarchy Level | 11934
- Description | 11934
- Required Privilege Level | 11934
- Release Information | 11935

Syntax

```
origin {
    host hostname;
    realm realm-name;
}
```

Hierarchy Level

```
[edit access diameter]
```

Description

Specify values of the Origin-Realm AVP and the Origin-Host AVP used in all messages sent by the Diameter instance. These values must be unique for each session PIC.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

origin (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 11935
- Hierarchy Level | 11935
- Description | 11936
- Options | 11936
- Required Privilege Level | 11936
- Release Information | 11936

Syntax

```
origin realm realm-name host hostname
```

Hierarchy Level

```
[edit diameter]
```

Description

Specify the hostname and realm of the endpoint node that originates Diameter messages for the Diameter instance. The Diameter instance supplies these values to be conveyed by the Origin-Realm-AVP and Origin-Host-AVP for all messages sent by the Diameter instance.

NOTE: Both the host and realm are mandatory.

Options

host *hostname*	Name of the message origin host that is supplied as the value of the Origin-Host AVP for all Diameter messages in the Diameter master instance.
realm *realm-name*	Name of the message origin realm, that is supplied as the value of the Origin-Realm AVP for all Diameter messages in the Diameter master instance.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter

Configuring the Origin Attributes of the Diameter Instance

origin-validation

IN THIS SECTION

- Syntax | 11937
- Hierarchy Level | 11937
- Description | 11937
- Default | 11937
- Options | 11938
- Release Information | 11938

Syntax

```
origin-validation accept | reject;
```

Hierarchy Level

```
[edit protocols bgp ebgp-community-cleanup]
[edit protocols bgp group <group-name> ebgp-community-cleanup]
[edit protocols bgp group <group-name> neighbor <address> ebgp-community-cleanup]
```

Description

You can choose to accept or reject the origin validation extended communities received from an eBGP peer. The default behavior of Origin Validation State Extended Community (OVS EC) is to reject any eBGP peer received. You can configure it to accept the community when needed by choosing the origin-validation accept option.

Default

The default behavior is to reject the origin validation extended communities received from an eBGP peer.

Options

Accept Use this option to accept the origin validation extended communities received from an eBGP peer.

Reject Use this option to reject the origin validation extended communities received from an eBGP peer.

Release Information

Statement ontroduced in Junos OS Release 22.2R1.

ospf

IN THIS SECTION

- Syntax | 11938
- Hierarchy Level | 11939
- Description | 11939
- Options | 11939
- Required Privilege Level | 11940
- Release Information | 11940

Syntax

```
ospf {
    domain-id domain-id;
    domain-vpn-tag number;
    route-type-community (iana | vendor);
    traffic-engineering {
        <advertise-unnumbered-interfaces>;
        <credibility-protocol-preference>;
        ignore-lsp-metrics;
        multicast-rpf-routes;
```

```
        no-topology;
        shortcuts {
            lsp-metric-into-summary;
        }
    }
    ... ospf-configuration ...
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable OSPF routing on the routing device. You must include the `ospf` statement to enable OSPF on the routing device. By default, OSPF is disabled.

Options

domain-id *domain-id*

The domain ID identifies the OSPF domain from which the route originated. If the router ID is not configured in the routing instance, the router ID is derived from an interface address belonging to the routing instance. The default OSPF domain ID is the null value 0.0.0.0.

domain-vpn-tag *number*

Set a virtual private network (VPN) tag for OSPFv2 external routes generated by the provider edge (PE) routing device. The *number* corresponds to the VPN tag.

route-type-community (iana | vendor)

Specify an extended community value to encode the OSPF route type. Each extended community is coded as an eight-octet value. This statement sets the most significant bit to either an IANA or vendor-specific route type.

- `iana`—Encode a route type with the value 0x0306. This is the default value.
- `vendor`—Encode the route type with the value 0x8000.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding OSPF Configurations

ospf-instance

IN THIS SECTION

- Syntax | 11940
- Hierarchy Level | 11941
- Description | 11941
- Default | 11941
- Required Privilege Level | 11941
- Release Information | 11942

Syntax

```
ospf-instance igpinstance-
name{ ... }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols]
```

Description

Enable multiple IGP instances of OSPFv2 routing on the routing device. The OSPFv2 IGP instance is associated with the default routing-instance and the default RIB. An OSPFv2 IGP instance exports routes installed in the RIB by another OSPFv2 IGP instance using standard Junos OS routing policy configuration. By default, the routes installed by different OSPFv2 IGP instances have the same route preference.

All the configuration options available under the `[edit protocols ospf-instance igpinstance-name]` hierarchy level are similar to the options available under the `[edit protocols ospf]` hierarchy level.

NOTE: The `ospf-instance` configuration statement is not supported at the [`edit routing-instances` *`routing-instance-name`* `protocols`] hierarchy level.

NOTE: Junos OS does not support configuring the same logical interface in multiple IGP instances of OSPFv2.

Starting in Junos OS and Junos OS Evolved Release 24.4R1, you can configure and run multiple independent interior gateway protocol (IGP) instances of OSPFv2 with segment routing (SR) on a router. You can create two or more OSPF instances and apply SR-MPLS on each instance. Multiple instances of OSPF can advertise different prefix-segment identifiers (prefix-SIDs). Other instances can use these SIDs for making routing decisions.

Default

OSPF is disabled on the routing device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

ospf3

IN THIS SECTION

- Syntax | 11942
- Hierarchy Level | 11943
- Description | 11943
- Options | 11943
- Required Privilege Level | 11943
- Release Information | 11943

Syntax

```
ospf3 {
    domain-id domain-id;
    domain-vpn-tag number;
    route-type-community (iana | vendor);
    traffic-engineering {
        <advertise-unnumbered-interfaces>;
        <credibility-protocol-preference>;
        ignore-lsp-metrics;
        multicast-rpf-routes;
        no-topology;
        shortcuts {
            lsp-metric-into-summary;
        }
    }
    ... ospf3-configuration ...
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable OSPFv3 routing on the routing device. You must include the `ospf3` statement to enable OSPFv3. By default, OSPFv3 is disabled.

Options

domain-id *domain-id*

The domain ID identifies the OSPF domain from which the route originated. If the router ID is not configured in the routing instance, the router ID is derived from an interface address belonging to the routing instance. The default OSPF domain ID is the null value 0.0.0.0.

domain-vpn-tag *number*

Set a virtual private network (VPN) tag for OSPFv2 external routes generated by the provider edge (PE) routing device. The *number* corresponds to the VPN tag.

route-type-community (iana | vendor)

Specify an extended community value to encode the OSPF route type. Each extended community is coded as an eight-octet value. This statement sets the most significant bit to either an IANA or vendor-specific route type.

- `iana`—Encode a route type with the value 0x0306. This is the default value.
- `vendor`—Encode the route type with the value 0x8000.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding OSPF Configurations

oss-map

IN THIS SECTION

- Syntax | 11944
- Hierarchy Level | 11944
- Description | 11944
- Options | 11945
- Required Privilege Level | 11945
- Release Information | 11945

Syntax

```
oss-map {
    model-name t640|t1600;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the operations support systems (OSS) mapping feature to map a T4000 chassis to a T1600 chassis or a T640 chassis, so that the T4000 chassis is represented as a T1600 chassis or a T640 chassis, respectively. The configuration helps prevent requalifying the T1600 chassis and T640 chassis as a new chassis on the OSS.

Options

`model-name t640`—Perform OSS mapping on a T4000 chassis to represent it as a T640 chassis, thereby overriding the chassis model name as displayed in the output of the `show chassis hardware`, the `show snmp mib walk system`, and the `show snmp mib walk jnxBoxAnatomy` operational commands.

`model-name t1600`—Perform OSS mapping on a T4000 chassis to represent it as a T1600 chassis, thereby overriding the chassis model name as displayed in the output of the `show chassis hardware`, the `show snmp mib walk system`, and the `show snmp mib walk jnxBoxAnatomy` operational commands.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R3, 13.1R2, and 13.2R1.

RELATED DOCUMENTATION

Configuring OSS Mapping to Represent a T4000 Chassis as a T1600 or a T640 Chassis

Example: Configuring a T4000 Chassis to Represent a T640 Chassis

show chassis oss-map | **18587**

other

IN THIS SECTION

- Syntax | **11946**
- Hierarchy Level | **11946**
- Description | **11946**
- Options | **11948**
- Required Privilege Level | **11948**

Release Information | 11948

Syntax

```
other {
    alarm-threshold (Security GPRS) {
        forward number;
        reverse number;
    }
    drop-threshold (Security GPRS) {
        forward number;
        reverse number;
    }
}
```

Hierarchy Level

```
[edit security gtp profile profile-name path-rate-limit message-type]
```

Description

Use other parameter to limit the number of packets per second for all the other GTPv0/GTPv1-C/GTPv2-C messages.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name path-rate-limit message-type]` hierarchy level is replaced by `[edit security gtp profile profile-name path-rate-limit message-type]`.

The message types not included in the other GTPv0 messages are listed in Table 90 on page 11946.

Table 90: GTPv0 Messages

Message Type	Message
1	Echo Request

Table 90: GTPv0 Messages *(Continued)*

Message Type	Message
16	Create PDP Context Request
20	Delete PDP Context Request
22	Create AA PDP Context Request
255	T-PDU

The message types not included in the other GTPv1 messages are listed in Table 91 on page 11947

Table 91: GTPv1 Messages

Message Type	Message
1	Echo Request
16	Create PDP Context Request
20	Delete PDP Context Request
255	G-PDU

The message types not included in the other GTPv2 messages are listed in Table 92 on page 11947.

Table 92: GTPv2 Messages

Message Type	Message
1	Echo Request
32	Create Session Request

Table 92: GTPv2 Messages *(Continued)*

Message Type	Message
36	Delete Session Request
95	Create Bearer Request
99	Delete Bearer Request

Options

`alarm-threshold`—To set alarm threshold for path rate limiting.

`drop-threshold`—To set drop threshold for path rate limiting.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding Path Rate Limiting for GTP Control Messages

Example: Limiting the Message Rate and Path Rate for GTP Control Messages

other-bytes

IN THIS SECTION

- Syntax | 11949
- Hierarchy Level | 11949
- Description | 11949
- Options | 11950
- Required Privilege Level | 11950
- Release Information | 11950

Syntax

```
other-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of frame overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for an access line of DSL type OTHER. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

The router reports some access technology types—such as Gigabit passive optical network (GPON) lines—as DSL type OTHER.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `other-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream frame overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS

Traffic Rate Reporting and Adjustment by the ANCP Agent

Configuring the ANCP Agent

other-overhead-adjust

IN THIS SECTION

- Syntax | 11951
- Hierarchy Level | 11951
- Description | 11951
- Options | 11952
- Required Privilege Level | 11952
- Release Information | 11952

Syntax

```
other-overhead-adjust percentage;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the actual downstream rate for an access line of DSL type OTHER received in the ANCP Port Up message by multiplying the rate by the specified percentage. The ANCP agent reports the adjusted rate to CoS.

The router reports some access technology types—such as Gigabit passive optical network (GPON) lines—as DSL type OTHER.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `other-overhead-adjust` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with

an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

percentage Percentage by which to multiply the rate.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS

Traffic Rate Reporting and Adjustment by the ANCP Agent

Configuring the ANCP Agent

other-stateful-configuration

IN THIS SECTION

- Syntax | 11953
- Hierarchy Level | 11953
- Description | 11953
- Default | 11953
- Required Privilege Level | 11954
- Release Information | 11954

Syntax

```
(other-stateful-configuration | no-other-stateful-configuration);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Description

Specify whether to enable autoconfiguration of other nonaddress-related information:

- **no-other-stateful-configuration**—Disable autoconfiguration of other nonaddress-related information.
- **other-stateful-configuration**—Enable autoconfiguration of other nonaddress-related information.

Default

By default, stateful autoconfiguration is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

Secure IPv6 Neighbor Discovery

other-stateful-configuration (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 11954
- Hierarchy Level | 11955
- Description | 11955
- Default | 11955
- Required Privilege Level | 11955
- Release Information | 11955

Syntax

```
(other-stateful-configuration | no-other-stateful-configuration);
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Specify whether to enable autoconfiguration of other nonaddress-related information:

- `no-other-stateful-configuration`—Disable autoconfiguration of other nonaddress-related information.
- `other-stateful-configuration`—Enable autoconfiguration of other nonaddress-related information.

Default

The configured object is disabled unless explicitly enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA

Dynamic Router Advertisement Configuration Overview

otn-options

IN THIS SECTION

Syntax | 11956

- Hierarchy Level | 11957
- Description | 11957
- Options | 11957
- Required Privilege Level | 11957
- Release Information | 11957

Syntax

```
otn-options {
    bytes (otn-options) transmit-payload-type value;
    fec (efec | gfec | gfec-sdfec | none );
    (is-ma | no-is-ma);
    (laser-enable | no-laser-enable);
    (line-loopback | no-line-loopback);
    (local-loopback | no-local-loopback);
    (odu-ttim-action-enable | no-odu-ttim-action-enable);
    (otu-ttim-action-enable | no-otu-ttim-action-enable);
    odu-delay-management {
        (bypass | no-bypass);
        (monitor-end-point | no-monitor-end-point);
        number-of-frames value;
        (no-start-measurement | start-measurement;
    }
    odu-signal-degrade {
        ber-threshold-clear value;
        ber-threshold-signal-degrade value;
        interval value;
    }
    (prbs | no-prbs);
    preemptive-fast-reroute {
        (backward-frr-enable | no-backward-frr-enable);
        (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
        odu-backward-frr-enable | no-odu-backward-frr-enable;
        odu-signal-degrade-monitor-enable | no-odu-signal-degrade-monitor-enable;
    }
    rate {
        (fixed-stuff-bytes | no-fixed-stuff-bytes);
        oc192;
```

```
        otu4;
        (pass-through | no-pass-through);
    }
    signal-degrade {
        ber-threshold-clear value;
        ber-threshold-signal-degrade value;
        interval value;
    }
    tca tca-identifier (enable-tca | no-enable-tca) (threshold number | threshold-24hrs number);
    transport-monitoring;
    trigger trigger-identifier;
    tti tti-identifier;
}
```

Hierarchy Level

```
[edit interfaces ge-fpc/pic/port]
[edit interfaces xe-fpc/pic/port]
[edit interfaces et-fpc/pic/port]
```

Description

Specify the Ethernet optical transport network (OTN) interface and options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interfaces-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

bytes, is-ma, local-loopback, no-is-ma, no-local-loopback, no-odu-ttim-action-enable, no-otu-ttim-action-enable, no-prbs, odu-delay-management, odu-ttim-action-enable, otu-ttim-action-enable, prbs, preemptive-fast-reroute, and signal-degrade statements introduced in Junos OS Release 13.2 for PTX Series routers.

oc192 statement introduced in Junos OS Release 13.3R3 for MX Series routers.

odu-signal-degrade, odu-backward-frr-enable | no-odu-backward-frr-enable, odu-signal-degrade-monitor-enable | no-odu-signal-degrade-monitor-enable statements introduced in Junos OS Release 14.1R2 and 14.2 for P2-100GE-OTN PIC in PTX5000 routers.

tca option introduced in Junos OS Release 14.2 for PTX Series routers.

bytes, line-loopback, local-loopback, preemptive-fast-reroute, tca, trigger, prbs, and tti statements introduced in 18.3R1 for ACX6360 routers.

RELATED DOCUMENTATION

10-Gigabit Ethernet OTN Options Configuration Overview

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

otu-ttim-action-enable

IN THIS SECTION

- Syntax | **11959**
- Hierarchy Level | **11959**
- Description | **11959**
- Default | **11959**
- Options | **11959**
- Required Privilege Level | **11959**
- Release Information | **11959**

Syntax

```
(otu-ttim-action-enable | no-otu-ttim-action-enable);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify whether consequent action for Optical Channel Transport Unit (OTU) TTIM is enabled or disabled.

Default

If you omit the otu-ttim-action-enable statement, consequent action for OTU TTIM is disabled.

Options

otu-ttim-action-enable	Enable consequent action for OTU TTIM.
no-otu-ttim-action-enable	Disable consequent action for OTU TTIM.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

otu4

IN THIS SECTION

- Syntax | 11960
- Hierarchy Level | 11960
- Description | 11960
- Default | 11960
- Options | 11960
- Required Privilege Level | 11961
- Release Information | 11961

Syntax

```
otu4;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

Description

Sets the line rate or speed of the OTN signal to optical channel transport unit 4 (OTU4).

Default

By default, the rate is OTU4 on PTX Series routers.

Options

otu4	OTU4 line rate or 100 Gbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

out-delay

IN THIS SECTION

- Syntax | 11961
- Hierarchy Level | 11962
- Description | 11962
- Default | 11963
- Options | 11963
- Required Privilege Level | 11963
- Release Information | 11963

Syntax

```
out-delay seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Control how often BGP and the routing table exchange route information by specifying how long a route must be present in the Junos OS routing table before it is exported to BGP. Use this time delay to help bundle routing updates and to avoid sending updates too often.

Alternatively or in addition, external BGP (EBGP) sessions can also use the route-flap damping mechanism upon the reception of BGP messages coming from an external neighbor.

BGP stores the route information it receives from update messages in the routing table, and the routing table exports active routes from the routing table into BGP. BGP then advertises the exported routes to its peers. The `out-delay` statement enables a form of rate limiting. The delay is added to each update for each prefix individually. When a routing device changes its best path to a destination prefix, the device does not inform its peer about the change unless the route has been present in its routing table for the specified `out-delay`. If you use `out-delay` to perform rate-limiting, you can expect a less bursty pattern of updates. You will see a pattern in which updates arrive in a steady flow, and two updates for the same prefix are always spaced by at least the `out-delay` timer value (for example, 30 seconds). Thus, the `out-delay` setting is useful for limiting oscillation (sometimes called *churn*) in a network. Keep in mind that, regardless of the `out-delay` setting, BGP peers exchange routes immediately after neighbor establishment. The `out-delay` setting is only designed to delay the exchange of routes between BGP and the local routing table.

Caution is warranted because an `out-delay` can delay convergence. If your network is configured in a way that avoids oscillation, setting an `out-delay` is not necessary.

When configured, the `out-delay` value displays as `Outbound Timer` when using `show bgp group` or `show bgp group neighbor` commands.

Default

By default, the exchange of route information between BGP and the routing table occurs immediately after the routes are received. This immediate exchange of route information might cause instabilities in the network reachability information. If you omit this statement, routes are exported to BGP immediately after they have been added to the routing table.

Options

seconds—Output delay time.

- **Range:** 0 through 65,535 seconds
- **Default:** 0 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

keep | 9452

out-of-resources

IN THIS SECTION

Syntax | 11964

- Hierarchy Level | 11964
- Description | 11964
- Options | 11964
- Required Privilege Level | 11964
- Release Information | 11965

Syntax

```
out-of-resources (block | (log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name fallback-
options]
```

Description

Virus scanning requires a great deal of memory and CPU resources. Due to resource constraints, memory allocation requests can be denied by the system. When out-of-resources occurs, scanning is terminated. The default action is block.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

out-of-resources (Security Antivirus Sophos Engine)

IN THIS SECTION

- Syntax | **11965**
- Hierarchy Level | **11965**
- Description | **11965**
- Options | **11966**
- Required Privilege Level | **11966**
- Release Information | **11966**

Syntax

```
default (block | log-and-permit | permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name fallback-
options]
```

Description

Virus scanning requires a great deal of memory and CPU resources. Due to resource constraints, memory allocation requests can be denied by the system. When out-of-resources occurs, scanning is terminated.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic
- `permit`—Permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Sophos Antivirus Configuration Overview

outbound-https

IN THIS SECTION

- Syntax | **11967**
- Hierarchy Level | **11967**
- Description | **11967**
- Options | **11968**
- Required Privilege Level | **11969**
- Release Information | **11969**

Syntax

```
outbound-https {
    client client-id {
        address {
            port port;
            trusted-cert trusted-cert;
        }
        device-id device-id;
        reconnect-strategy (in-order | sticky);
        secret password;
        waittime seconds;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure a Junos device that's behind a firewall to initiate outbound HTTPS connections to communicate with client management applications on the other side of the firewall. The `outbound-https` configuration is consumed by the outbound HTTPS extension service. You must configure this service at the `[edit system extensions extension-service application file nc_grpc_app.pyc]` hierarchy level in order to initiate the outbound HTTPS connections.

When you configure and start the outbound HTTPS extension service on supported Junos devices, the extension service uses the `outbound-https` configuration to connect to and authenticate each configured client, which corresponds to a gRPC server running on a network management system. The device and gRPC server establish a persistent HTTPS connection over a TLS-encrypted gRPC session. The device authenticates the gRPC server using an X.509 digital certificate, and the gRPC server uses the `device-id` and `shared-secret` values to authenticate the Junos device. An outbound HTTPS client can establish multiple NETCONF or shell sessions with the device.

You can configure multiple outbound HTTPS clients, and you can configure one or more backup gRPC servers for each client. The device connects to only one gRPC server in the client's server list at any one time.

Options

client *client-id*

Define a device-initiated outbound HTTPS connection.

This value serves to uniquely identify the `outbound-https` configuration stanza. Each stanza represents a connection to a single outbound HTTPS client. Thus, the administrator is free to assign the `client-id` any meaningful unique value. This attribute is not sent to the client management application.

address

Hostname or IPv4 address of the gRPC server running on the network management system.

The hostname or IP address must match the value of the Common Name (CN) field or the SubjectAltName IP Address field, respectively, in that gRPC server's X.509 certificate. You can configure multiple backup gRPC servers, but the device connects to only one server in the list at any given time.

You must configure the following connection parameters for each server:

- port *port*—Port on which the gRPC server is listening for outbound HTTPS connection requests.
- trusted-cert *trusted-cert*—Certificate information used to authenticate the gRPC server's X.509 certificate.

  If the server's certificate is self-signed, configure the contents of the gRPC server's certificate, omitting any newlines.

  If the server's certificate is authenticated using a certificate chain, concatenate any intermediate CA and root CA certificates in that order, remove all newlines, and configure the resulting single string.

device-id *device-id*

Identifies the Junos device to the management application. Each time the device establishes an outbound HTTPS connection, it sends its device identifier and shared secret to the management application, and the management application uses the values to authenticate the device.

reconnect-strategy (in-order | sticky)

(Optional) Method used to reestablish a disconnected outbound HTTPS connection.

- Values:
  - in-order—Attempt to reconnect to the first server in the list. If the server is unavailable, attempt to connect to the next server in the list, and so on, until the device establishes a connection.

- sticky—Attempt to reconnect to the server to which the device was last connected. If the server is unavailable, attempt to connect to the next server in the list, and so on, until the device establishes a connection.

- **Default:** in-order

secret *password*

Shared secret between the Junos device and the management application. Each time the device establishes an outbound HTTPS connection, it sends its device identifier and shared secret to the management application, and the management application uses the values to authenticate the device.

waittime *seconds*

Number of seconds that the device waits before attempting to connect or reconnect to the servers in the list if none of the servers are available. That is, if the device reaches the end of the configured server list and cannot establish a connection, it waits the specified number of seconds before again attempting to connect to each server in the list, starting from the top.

- **Default:** 30 seconds

- **Range:** 0 through 4,294,967,295 seconds

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

NETCONF and Shell Sessions over Enhanced Outbound HTTPS

outbound-route-filter

IN THIS SECTION

- Syntax | 11970
- Hierarchy Level | 11970
- Description | 11971
- Options | 11971
- Required Privilege Level | 11971
- Release Information | 11971

Syntax

```
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            (inet | inet6);
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
```

```
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Configure a BGP peer to accept outbound route filters from a remote peer.

Options

accept—Specify that outbound route filters from a BGP peer be accepted.

inet—Specify that IPv4 prefix-based outbound route filters be accepted.

inet6—Specify that IPv6 prefix-based outbound route filters be accepted.

> **NOTE**: You can specify that both IPv4 and IPv6 outbound route filters be accepted.

prefix-based—Specify that prefix-based filters be accepted.

The `bgp-orf-cisco-mode` statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Example: Configuring BGP Prefix-Based Outbound Route Filtering

outbound-ssh

IN THIS SECTION

- Syntax | 11972
- Hierarchy Level | 11973
- Description | 11973
- Options | 11973
- Required Privilege Level | 11975
- Release Information | 11975

Syntax

```
outbound-ssh {
    client client-id {
        address {
            port port-number;
            retry number;
            timeout seconds;
        }
        device-id device-id;
        keep-alive {
            retry number;
            timeout seconds;
        }
        reconnect-strategy (in-order | sticky);
        routing-instance routing-instance-name;
        secret password;
        services  netconf;
    }
    traceoptions {
        file <filename> <files number> <match regular-expression> <size size> <(world-readable |
no-world-readable)>;
        flag flag;
        no-remote-trace;
    }
```

```
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure a device running the Junos OS behind a firewall to initiate outbound SSH connections to communicate with client management applications on the other side of the firewall.

Options

client *client-id*

Defines a device-initiated connection. This value serves to uniquely identify the outbound-ssh configuration stanza. Each outbound-ssh stanza represents a single outbound SSH connection. Thus, the administrator is free to assign the client-id any meaningful unique value. This attribute is not sent to the client management application.

address

Hostname, IPv4 address, or IPv6 address of the management application server.

> **NOTE**: Starting in Release 15.1, Junos OS supports outbound SSH connections with devices having IPv6 addresses.

- **Syntax:** You can list multiple servers by adding each server's IP address or hostname along with the following connection parameters:
  - port *port-number*—Specifies the port number at which a server listens for outbound SSH connection requests.

    Default: port 22
  - retry *number*—Specifies the maximum number of times the device attempts to establish an outbound SSH connection before giving up.

    Default: 3 attempts

- timeout *seconds*—Specifies how long the device waits between attempts to reconnect to the specified IP address to establish an outbound SSH connection before giving up.

  Default: 15 seconds

device *device-id*

(Required) Identifies the device to the management application. Each time the device establishes an outbound SSH connection, it first sends an initiation sequence (device-id) to the management application.

keep-alive

(Optional) When configured, specifies that the device should send SSH protocol keepalive messages to the management application.

- **Syntax:** To configure keepalive messages, you must set both the retry and timeout attributes:

  - retry *number*—specifies how many keepalive messages the device sends without receiving a response from the application. When that number is exceeded, the device disconnects from the application, ending the outbound SSH connection.

    Default: 3 attempts

  - timeout *seconds*—specifies how long the device waits to receive data before sending a request for acknowledgment from the application.

    Default: 15 seconds

reconnect-strategy (in-order|sticky)

(Optional) Specify the method the device uses to reestablish a disconnected outbound SSH connection.

- **Values:** Two methods are available:

  - in-order—Configures the device to reconnect to the first configured server. If this server is unavailable, the device tries to connect to the next configured server. The device keeps trying each server in the configured list until the device can establish a connection.

  - sticky—Specify that the device should first attempt to reconnect to the management server from which it disconnected. If that server is unavailable, the device then attempts to connect to the next configured server. The device keeps trying each server in the configured list until the device can establish a connection.

routing-instance *routing-*

(SRX Series and MX Series only) Specify the name of the routing instance on which the outbound SSH connection needs to be established. If you do not specify a routing

instance-name	instance, your device will establish the outbound SSH connection using the default routing table.
secret *password*	Configures the device to send the device's public SSH host key when the device connects to the management server. This is the recommended method of maintaining a current copy of the device's public key.
services *netconf*	Configures the management application to accept NETCONF as an available service.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for IPv6 addresses added in Junos OS Release 12.1X47-D15.

`routing-instance` option introduced in Junos OS Release 19.3R1 for SRX Series and MX Series devices.

outer-tag-protocol-id (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 11976
- Hierarchy Level | 11976
- Description | 11976
- Required Privilege Level | 11976
- Release Information | 11976

Syntax

```
outer-tag-protocol-id;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

TPID to be used in the outer VLAN tag. Supported values are 0x8100, 0x88a8, 0x9100, 0x9200

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

outgoing-queue

IN THIS SECTION

Syntax | **11977**

- Hierarchy Level | **11977**
- Description | **11977**
- Options | **11977**
- Required Privilege Level | **11978**
- Release Information | **11978**

Syntax

```
outgoing-queue {
    <high-watermark high-watermark>;
    <low-watermark low-watermark>;
    size size;
}
```

Hierarchy Level

```
[edit access diameter peer peer-name]
```

Description

Configure the outgoing queue properties for this peer. When the queue size reaches the high watermark, the peer is marked unavailable, any new messages to the Diameter network element are not sent to this peer, and the SNMP trap **Diameter_PeerOutQHiWMarkNotif** is generated. When the queue size descends below the low watermark after reaching the high watermark, the peer becomes available and the SNMP trap **Diameter_PeerLowQHiWMarkNotif** is generated.

Options

`high-watermark` *`high-watermark`*

(Optional) Use the specified high watermark for this peer.

- **Range:** 1 through 100 percent
- **Default:** 80

`low-watermark` *`low-watermark`*

(Optional) Use the specified low watermark for this peer.

- **Range:** 1 through 100 percent
- **Default:** 60

`size` *size*

Use the specified size of the queue. The default is 6000.

- **Range:** 1 through 65,535 packets

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Diameter Peers

diameter (Subscriber Aware Policy Control) | **5699**

output (Service Set)

IN THIS SECTION

- Syntax | **11979**
- Hierarchy Level | **11979**
- Description | **11979**
- Options | **11979**
- Required Privilege Level | **11979**
- Release Information | **11979**

Syntax

```
output {
    [ service-set service-set-name <service-filter filter-name> ];
}
```

Hierarchy Level

```
[edit interface interface-name unit logical-unit-number family inet service],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet service]
```

Description

Define the output service sets and filters to be applied to traffic.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Filters and Services to Interfaces

output

IN THIS SECTION

- Syntax | 11980
- Hierarchy Level | 11980
- Description | 11981
- Required Privilege Level | 11981
- Release Information | 11981

Syntax

```
output {
    interface interface-name;
    ip-address ip-address;
    vlan (vlan-id | vlan-name);
    routing-instance instance-name {
       ip-address address;
    }
}
```

Hierarchy Level

For platforms without ELS:

```
[edit ethernet-switching-options analyzer name]
```

For platforms with ELS:

```
[edit forwarding-options analyzer name],
[edit forwarding-options port-mirroring instance name family ethernet-switching]
```

Description

Configure the destination for mirrored traffic, either an interface on the switch (for local monitoring) or a VLAN (for remote monitoring).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Option **output vlan** added in Junos OS Release 12.1 for the QFX Series.

output (Accounting)

IN THIS SECTION

- Syntax | 11981
- Hierarchy Level | 11982
- Description | 11982
- Required Privilege Level | 11982
- Release Information | 11982

Syntax

```
output {
    aggregate-export-interval seconds;
    cflowd hostname {
        aggregation {
            autonomous-system;
```

```
            destination-prefix;
            protocol-port;
            source-destination-prefix {
                caida-compliant;
            }
            source-prefix;
        }
        autonomous-system-type (origin | peer);
        (local-dump | no-local-dump);
        port port-number;
        version format;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface interface-name {
        engine-id number;
        engine-type number;
        source-address address;
    }
}
```

Hierarchy Level

```
[edit forwarding-options accounting name]
```

Description

Configure cflowd, output interfaces, and flow properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Discard Accounting

Configuring Discard Accounting

output-chain

IN THIS SECTION

- Syntax | 11983
- Hierarchy Level | 11983
- Description | 11983
- Options | 11984
- Required Privilege Level | 11984
- Release Information | 11984

Syntax

```
output-chain [ filter-name ]
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family filter]
```

Description

The `output-chain` command works like the **next** filter term to allow multiple levels of filtering on the interface in both the ingress and egress direction. For example, you can implement a combination of classification and firewall filter rules for evaluation, wherein the first filter performs a generic filter classification and the subsequent filters go on to perform the actual filtering and actions.

To continue on to the next filter, the terminating action in the current filter must be **accept**.

The `output-chain` command is supported on all interfaces except loopback (lo0 and fxp0) and interfaces under logical systems.

Options

[*filter-names*]—Name of each filter in the chain that is to be evaluated when packets are received on the interface. Supports up to 8 ingress and 8 egress filters. Evaluation occurs in the order in which the names appear, from left to right.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

Statement introduced in Junos OS Evolved Release 21.4R1 (PTX10001-36MR, PTX10004, PTX10008, and PTX10016).

RELATED DOCUMENTATION

Example: Using Firewall Filter Chains

input-chain | **8690**

Understanding Multiple Firewall Filters Applied as a List

Applying a Filter to an Interface

output (Congestion Notification)

IN THIS SECTION

- Syntax | **11985**
- Hierarchy Level | **11985**
- Description | **11985**

- Options | 11985
- Required Privilege Level | 11985
- Release Information | 11985

Syntax

```
output {
    ieee-802.1 {
        code-point [code-point-bits] {
            flow-control-queue [queue | list-of-queues];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name]
```

Description

Configure priority-based flow control (PFC) on output queues.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring CoS PFC (Congestion Notification Profiles)

Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces

Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface

Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)

Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)

Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows

output (Dynamic Service Sets)

IN THIS SECTION

- Syntax | **11986**
- Hierarchy Level | **11987**
- Description | **11987**
- Options | **11987**
- Required Privilege Level | **11987**
- Release Information | **11987**

Syntax

```
output {
    service-set service-set-name {
    service-filter filter-name;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family service],
[edit dynamic-profiles profile-name interfaces pp0 unit  "$junos-interface-unit" family family
service]
```

Description

Define the output service sets and filters to be applied to traffic by a dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

service-set-name—Name of the service set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support of the `[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family service]` hierarchy level introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *Dynamic Service Sets Overview*
>
> *Associating Service Sets with Interfaces in a Dynamic Profile*

output-filename (Event Policy)

IN THIS SECTION

- Syntax | 11988
- Hierarchy Level | 11988
- Description | 11988
- Options | 11989
- Required Privilege Level | 11989
- Release Information | 11989

Syntax

```
output-filename filename;
```

Hierarchy Level

```
[edit event-options policy policy-name then event-script filename],
[edit event-options policy policy-name then execute-commands]
```

Description

Define a descriptive string that will be included in the filename of the output file in which to write command or script output. For op scripts, this statement is optional.

The name of the output file also includes the hostname and timestamp to ensure that the each filename is unique. If an event policy is triggered multiple times in a 1-second period, an index number is appended to the filename. The index number range is 001 through 999.

Prior to Junos OS Release 14.1R3, the filename has the following naming convention:

```
hostname_filename_YYYYMMDD_HHMMSS_index-number
```

Starting in Junos OS Release 14.1R3, the ***filename*** string is placed after the timestamp.

```
hostname_YYYYMMDD_HHMMSS_filename_index-number
```

Options

filename String included in the filename of the output file.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

Support at the `[edit event-options policy policy-name then event-script filename]` hierarchy level introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configure an Event Policy to Execute Operational Mode Commands

Execute Event Scripts in an Event Policy

output-format (Event Policy)

IN THIS SECTION

- Syntax | 11990
- Hierarchy Level | 11990
- Description | 11990
- Default | 11990
- Options | 11990

Required Privilege Level | 11991
Release Information | 11991

Syntax

```
output-format (text | xml);
```

Hierarchy Level

```
[edit event-options policy policy-name then event-script filename],
[edit event-options policy policy-name then execute-commands]
```

Description

Specify the format (*ASCII* text or XML) for the output of event scripts or commands executed by an event policy.

> **NOTE**: Event script output is always emitted in XML format even if you configure the `output-format text` statement under the `[edit event-options policy policy-name then event-script filename]` hierarchy.

Default

If you do not include the `output-format` statement, the output is emitted in XML format.

Options

text Formatted ASCII text.

xml Junos Extensible Markup Language (XML) tags.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

Support at the `[edit event-options policy policy-name then event-script filename]` hierarchy level introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

> *Configure an Event Policy to Execute Operational Mode Commands*
>
> *Execute Event Scripts in an Event Policy*

output-forwarding-class-map

IN THIS SECTION

- Syntax | 11991
- Hierarchy Level | 11992
- Description | 11992
- Options | 11992
- Usage Guidelines | 11992
- Required Privilege Level | 11992
- Release Information | 11992

Syntax

```
output-forwarding-class-map forwarding-class-map-name;
```

Hierarchy Level

```
[edit class-of-service forwarding-classes-interface-specific]
```

Description

Apply a configured forwarding class map to a logical interface.

Options

forwarding-class-map-name—Name of a forwarding class mapping configured at the **[edit class-of-service forwarding-classes-interface-specific]** hierarchy level.

Usage Guidelines

Classifying Packets by Egress Interface

Required Privilege Level

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

output (Forwarding Table)

IN THIS SECTION

- Syntax | **11993**
- Hierarchy Level | **11993**
- Description | **11993**
- Options | **11993**
- Required Privilege Level | **11993**

Release Information | 11993

Syntax

```
output filter-name;
```

Hierarchy Level

```
[edit forwarding-options family (inet | inet6 | mpls) filter],
[edit routing-instances routing-instance-name forwarding-options family (inet | inet6 | mpls)
filter]
```

Description

Configure filtering on the egress traffic of the forwarding table.

Options

filter-name—Name of the applied filter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Applying Forwarding Table Filters

output-interface (CCC)

IN THIS SECTION

- Syntax | 11994
- Hierarchy Level | 11994
- Description | 11994
- Required Privilege Level | 11994
- Release Information | 11994

Syntax

```
output-interface [interface-name 1 interface-name n];
```

Hierarchy Level

```
[edit protocols connections p2mp-transmit-switch p2mp-transmit-switch-name]
```

Description

Specify one or more output interfaces to switch traffic on an incoming CCC interface to one or more outgoing CCC interfaces.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring CCC Switching for Point-to-Multipoint LSPs

output-interface-index

IN THIS SECTION

- Syntax | 11995
- Hierarchy Level | 11995
- Description | 11995
- Options | 11995
- Required Privilege Level | 11996
- Release Information | 11996

Syntax

```
output-interface-index number;
```

Hierarchy Level

```
[edit forwarding-options monitoring name output interface interface-name]
```

Description

Specify a value for the output interface index that overrides the default supplied by SNMP.

Options

number—Output interface index value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Flow Monitoring

output interfaces (Chassis Synchronization)

IN THIS SECTION

- Syntax (ACX710) | 11996
- Syntax (MX Series) | 11997
- Hierarchy Level | 11997
- Description | 11997
- Options | 11997
- Required Privilege Level | 11998
- Release Information | 11998

Syntax (ACX710)

```
output {
    interfaces {
        bits {
            source-mode (chassis | ptp-clock);
            squelch {
                method (ais | ql-dnu | rai | shut);
```

```
            }
            ssm option-2  (GEN-1 | GEN-2);
        }
    }
}
```

Syntax (MX Series)

```
output {
    interfaces {
        external {
            holdover-mode-disable;
            minimum-quality (prc | prs | sec | smc | ssu-a | ssu-b  | st2 | st3 | st3e | st4 |
stu | tnc);
            source-mode (chassis | line);
            tx-dnu-to-line-source-enable;
            wander-filter-disable;
        }
    }
}
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Configure the options for clock output interfaces.

Options

external (SCBE only) Specify the properties of the external output interface.

(external-0/0 | external-1/0) (SCBE2 only) Specify the properties of the external output interface.

holdover-mode-disable Disable holdover.

tx-dnu-to-line-source-enable	Set Tx quality level to DNU/DUS on line source interface that has been selected as the external output source.
wander-filter-disable	Disable wander filtering.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (ACX Series) | **28787**

synchronization (MX Series) | **28795**

output (Mirroring)

IN THIS SECTION

- Syntax | **11999**
- Hierarchy Level | **11999**
- Description | **11999**
- Required Privilege Level | **11999**
- Release Information | **11999**

Syntax

```
output {
    bridge-domain bridge-domain-name;
    interface interface-name;
    next-hop-group next-hop-group-name;
    routing-instance {
        instance-name {
            bridge-domain {
                bridge-domain-name;
            }
        }
        vlan (vlan-id | vlan-name);
    }
    vlan (vlan-id | vlan-name);
}
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name]
```

Description

Configure the destination for mirrored traffic, either an interface on the network device for local monitoring, or a VLAN or bridge domain, for remote monitoring.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Example: Configuring Port Mirroring Analyzers for Local Monitoring of Employee Resource Use

Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use

output (Monitoring)

IN THIS SECTION

- Syntax | 12000
- Hierarchy Level | 12001
- Description | 12001
- Required Privilege Level | 12001
- Release Information | 12001

Syntax

```
output {
    cflowd hostname port port-number;
    export-format format;
    flow-active-timeout seconds;
    flow-export-destination {
        (cflowd-collector | collector-pic);
    }
    flow-inactive-timeout seconds;
    interface interface-name {
        engine-id number;
        engine-type number;
        input-interface-index number;
        output-interface-index number;
        source-address address;
    }
}
```

Hierarchy Level

```
[edit forwarding-options monitoring name family inet]
```

Description

Configure cflowd, output interfaces, and flow properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Flow Monitoring

output-policer

IN THIS SECTION

- Syntax | 12002
- Hierarchy Level | 12002
- Description | 12002
- Options | 12002
- Required Privilege Level | 12002
- Release Information | 12002

Syntax

```
output-policer policer-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number layer2-policer],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
layer2-policer]
```

Description

Apply a single-rate two-color policer to the Layer 2 output traffic at the logical interface. The `output-policer` and `output-three-color` statements are mutually exclusive.

Options

policer-name—Name of the single-rate two-color policer that you define at the `[edit firewall]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Two-Color and Three-Color Policers at Layer 2

Applying Layer 2 Policers to Gigabit Ethernet Interfaces

Configuring Gigabit Ethernet Policers

input-policer | **8713**

input-three-color | **8724**

layer2-policer (EX) | 9654
logical-interface-policer | 10069
output-three-color | 12014

output (Port Mirroring)

IN THIS SECTION

- Syntax | 12003
- Syntax: QFX10000 Switches | 12003
- Hierarchy Level | 12004
- Description | 12004
- Required Privilege Level | 12004
- Release Information | 12004

Syntax

```
output {
    interface interface-name {
        next-hop address;
    }
    no-filter-check;
}
```

Syntax: QFX10000 Switches

```
output {
        dscp dscp-value;
        forwarding-class class-name;
        ip-destination-address address;
        ip-source-address address;
```

```
        policer policer-name;
    }
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family (inet | inet6)],
[edit forwarding-options port-mirroring instance instance-name]
```

Description

Configure output interfaces and flow properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

output (Port Mirroring)

IN THIS SECTION

- Syntax | 12005
- Hierarchy Level | 12005

- Description | 12006
- Options | 12006
- Required Privilege Level | 12006
- Release Information | 12006

Syntax

```
output {
      file file-name {
          linear;
          max-size value;
        }
    interface interface-name {
        next-hop address;
    }
    next-hop-group group-name{
        group-type inet6;
        interface interface-name {
            next-hop ipv6-address;
        }
        next-hop-subgroup group-name{
            interface interface-name {
                next-hop ipv6-address;
            }
        }
    }
    no-filter-check;
    server-profile server-profile-name;
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family (any | ccc | inet | inet6 | mpls | vpls)],
[edit forwarding-options port-mirroring instance instance-name family (any | ccc | inet | inet6 |
mpls | vpls)]
```

Description

Configure the port-mirroring destination properties.

Options

`file` *`file-name`*—PCAP file to save self-mirrored packets to.

`linear`—Write mode of the mirrored-packet file. (The default write mode is circular.)

`max-size` *`value`*—Maximum size of the mirrored-packet file. Up to 5MB.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`vpls` option introduced in Junos OS Release 9.3 for MX Series routers only; support extended to M7i, M10i, M120, and M320 routers in Junos OS Release 9.5.

`ccc` option introduced in Junos OS Release 9.6 for M120 and M320 routers only.

`server-profile` option introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers only.

`next-hop-group` option introduced for `family inet6` in Junos OS Release 14.2 for MX Series routers only.

RELATED DOCUMENTATION

Configuring Port Mirroring

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

output-priority-map

IN THIS SECTION

- Syntax | 12007
- Hierarchy Level | 12007
- Description | 12007
- Required Privilege Level | 12008
- Release Information | 12008

Syntax

```
output-priority-map {
    classifier {
        premium {
            forwarding-class class-name {
                loss-priority (high | low);
            }
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile]
[edit interfaces interface-name ether-options ethernet-switch-profile ethernet-policer-profile]
```

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the output policer priority map to be applied to outgoing frames on this interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

input-priority-map | 8716

output-queue-priority

IN THIS SECTION

- output-queue-priority (System Configuration) | 12008
- Hierarchy Level (System Configuration) | 12009
- output-queue-priority (Implementation) | 12009
- Hierarchy Level (Implementation) | 12009
- Description | 12009
- Default | 12009
- Required Privilege Level | 12010
- Release Information | 12010

output-queue-priority (System Configuration)

```
output-queue-priority {
    expedited update-tokens number-of-tokens;
```

```
    priority priority-queue-number (1-16) update-tokens number-of-tokens;
}
```

Hierarchy Level (System Configuration)

```
[edit logical-systems logical-system-name protocols bgp],
[edit protocols bgp]
```

output-queue-priority (Implementation)

```
output-queue-priority {
    (expedited | priority priority-queue-number (1-16));
}
```

Hierarchy Level (Implementation)

```
[edit logical-systems logical-system-name protocols bgp family family-name sub-family],
[edit logical-systems logical-system-name protocols bgp group group-name family family-name sub-
family],
[edit protocols bgp family family-name sub-family],
[edit protocols bgp group group-name neighbor neighbor-id family family-name]
```

Description

When configuring the queues for BGP route prioritization, the `output-queue-priority` statement allows you to specify the number of tokens that are available within each of the 17 BGP output priority queues. This allows you to balance the amount of work that can be done within the route prioritization queues.

When implementing BGP route prioritization, the `output-queue-priority` statement allows you to specify the priority at which a given BGP route or route type is serviced. BGP route prioritization can also be specified per BGP neighbor during BGP configuration within the `inet` address family.

Default

By default, each of the 17 BGP route priority queues (or buckets) is assigned 1 work token so that misconfigured queues do not result in starvation.

When implementing BGP route prioritization, the following types of update messages are assigned to the lowest priority queue (1) by default: route refresh, topology change, and route withdraw.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding BGP Route Prioritization

withdraw-priority | 31134

output (Sampling)

IN THIS SECTION

- Syntax | 12010
- Hierarchy Level | 12012
- Description | 12012
- Required Privilege Level | 12012
- Release Information | 12012

Syntax

```
output {
    aggregate-export-interval seconds;
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
```

```
    flow-server hostname {
        aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
                caida-compliant;
            }
            source-prefix;
        }
        autonomous-system-type (origin | peer);
        dscp dscp-value;
        forwarding-class class-name;
        (local-dump | no-local-dump);
        port port-number;
        version format;
        version9 {
            template template-name;
        }
        version-ipfix {
            template template-name;
        }
    }
    interface interface-name {
        engine-id number;
        engine-type number;
        source-address address;
    }
    file {
        disable;
        filename filename;
        files number;
        size bytes;
        (stamp | no-stamp);
        (world-readable | no-world-readable);
    }
    inline-jflow {
        source-address address;
        flow-export-rate rate;
    }
}
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls)],
[edit forwarding-options sampling family (inet |inet6 |mpls | vpls)]
```

Description

Configure cflowd or flow monitoring, output files and interfaces, and flow properties.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: The `inline-jflow` statement is valid only under the `[edit forwarding-options sampling instance` *`instance-name`* `family` *`family`* `output]` hierarchy level. The `file` statement is valid only under the `[edit forwarding-options sampling family inet output]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for IPv6 addresses for the `flow-server` *`address`* and `source-address` *`address`* statements within sampling instances introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

> *Configuring Traffic Sampling on MX, M and T Series Routers*
>
> *Collecting Traffic Sampling Output in a File*

output (Security Forwarding Options)

IN THIS SECTION

- Syntax | 12013
- Hierarchy Level | 12013
- Description | 12013
- Options | 12014
- Required Privilege Level | 12014
- Release Information | 12014

Syntax

```
output {
    destination-mac mac-address;
    interface interface-name;
}
```

Hierarchy Level

```
[edit security forwarding-options mirror-filter filter-name]
```

Description

Specify the MAC address or interface for mirrored traffic.

Starting in Junos OS Release 18.4R1, if the output X2 interface of a mirror filter is configured for an st0 interface to filter traffic that you want to analyze, the packet is duplicated and encrypted by the IPsec tunnel bound to the st0 interface. This enhancement supports the SRX Series Firewalls to send traffic mirrored from a port on an IPsec tunnel.

Options

destination-mac ***mac-address***	Specify the MAC address for the mirrored traffic.
interface ***interface-name***	Specify the logical interface for the mirrored traffic.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

mirror-filter (Security Forwarding Options) | **10768**

show security forwarding-options mirror-filter | **24363**

output-three-color

IN THIS SECTION

- Syntax | **12015**
- Hierarchy Level | **12015**
- Description | **12015**
- Options | **12015**
- Required Privilege Level | **12015**
- Release Information | **12015**

Syntax

```
output-three-color policer-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number layer2-policer]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
layer2-policer]
```

Description

Apply a single-rate or two-rate three-color policer to the Layer 2 output traffic at the logical interface. The `output-three-color` and `output-policer` statements are mutually exclusive.

Options

policer-name—Name of the single-rate or two-rate three-color policer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Two-Color and Three-Color Policers at Layer 2

Applying Layer 2 Policers to Gigabit Ethernet Interfaces

Configuring Gigabit Ethernet Policers

input-three-color | **8724**

input-policer | **8713**

layer2-policer (EX) | 9654
logical-interface-policer | 10069
output-policer | 12001

output-traffic-control-profile

IN THIS SECTION

- Syntax | 12016
- Hierarchy Level | 12016
- Description | 12017
- Options | 12017
- Required Privilege Level | 12017
- Release Information | 12017

Syntax

```
output-traffic-control-profile profile-name;
```

```
output-traffic-control-profile profile-name shared-instance instance-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name ],
[edit class-of-service interfaces interface-name unit logical-unit-number],
[edit class-of-service interfaces interface-name interface-set interface-set-name]
[edit class-of-service interfaces interface-name slice slice-name]
[edit class-of-service interfaces interface-name forwarding-class-set forwarding-class-set-name]
```

Description

Apply the specified CoS traffic control profile (traffic scheduling and shaping configuration objects) to the output traffic at the physical interface, logical interface, interface set, slice, or forwarding class set (priority group).

Options

profile-name—Name of the traffic-control profile to be applied to this interface.

shared-instance *instance-name*—For platforms that support this option, name of the shared scheduler and shaper to be applied to this interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

interface-set option added in Junos OS Release 8.5.

RELATED DOCUMENTATION

Oversubscribing Interface Bandwidth

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

Configuring Hierarchical Schedulers for CoS

Configuring Interface Sets

output-traffic-control-profile-remaining | **12019**

traffic-control-profiles | **30009**

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Assigning CoS Components to Interfaces

Understanding CoS Traffic Control Profiles

output-traffic-control-profile (Dynamic CoS Definition)

IN THIS SECTION

- Syntax | 12018
- Hierarchy Level | 12018
- Description | 12018
- Options | 12018
- Required Privilege Level | 12019
- Release Information | 12019

Syntax

```
output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-
number]
```

Description

Apply an output traffic scheduling and shaping profile to the logical interface.

Options

profile-name—Name of the traffic-control profile to be applied to this interface

`$junos-cos-traffic-control-profile`—Variable for the traffic-control profile that is specified for the logical interface. The variable is replaced with the traffic-control profile when the subscriber is authenticated at login.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Variable `$junos-cos-traffic-control-profile` introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile

Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers

output-traffic-control-profile-remaining

IN THIS SECTION

- Syntax | 12019
- Hierarchy Level | 12020
- Description | 12020
- Options | 12020
- Required Privilege Level | 12020
- Release Information | 12020

Syntax

```
output-traffic-control-profile-remaining profile-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit class-of-service interfaces interface-name interface-set interface-set-name]
```

Description

Apply the specified traffic control profile (traffic scheduling and shaping configuration objects) to the remaining output traffic at the physical interface or interface set. The remaining traffic is transmitted by the default interface or interface set.

You can map the TCP to the interface or interface set by using the `output-traffic-control-profile-remaining` statement to explicitly configure the queues of the default interface or interface set scheduler that transmits the remaining traffic.

Options

profile-name—Name of the traffic-control profile for remaining traffic to be applied to this interface or interface set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Hierarchical Schedulers for CoS

Configuring Remaining Common Queues on MIC and MPC Interfaces

output-traffic-control-profile | **12016**

output-vlan-map

IN THIS SECTION

- Syntax | 12021
- Hierarchy Level | 12021
- Description | 12021
- Required Privilege Level | 12022
- Release Information | 12022

Syntax

```
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For EX Series switches, defines the rewrite operation to be applied to outgoing frames.

For MX Series routers and NFX Series devices' Gigabit Ethernet IQ and 10-Port 10-Gigabit Ethernet SFPP interfaces only, defines the rewrite operation to be applied to outgoing frames on this logical interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

pop-pop, **pop-swap**, **push-push**, **swap-push**, and **swap-swap** statements added in Junos OS Release 8.1.

RELATED DOCUMENTATION

Stacking and Rewriting Gigabit Ethernet VLAN Tags

input-vlan-map | 8729

output-vlan-map (Aggregated Ethernet)

IN THIS SECTION

- Syntax | 12022
- Hierarchy Level | 12023
- Description | 12023
- Required Privilege Level | 12023
- Release Information | 12023

Syntax

```
output-vlan-map {
    (pop | push | swap);
    tag-protocol-id tpid;
    vlan-id number;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Define the rewrite profile to be applied to outgoing frames on this logical interface. On MX Series routers, this statement only applies to aggregated Ethernet inferfaces using Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and 100-Gigabit Ethernet Type 5 PIC with CFP..

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Starting in Junos OS Release 17.3R1, input-vlan-map for outer vlan is supported for L2 circuit over aggregated Ethernet interfaces for QFX10000 Series switches.

RELATED DOCUMENTATION

Stacking and Rewriting Gigabit Ethernet VLAN Tags

input-vlan-map (Aggregated Ethernet) | **8731**

output-vlan-map (Dynamic Interfaces)

IN THIS SECTION

- Syntax | **12024**

- Hierarchy Level | 12024
- Description | 12024
- Required Privilege Level | 12024
- Release Information | 12025

Syntax

```
output-vlan-map {
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    (pop | swap);
    tag-protocol-id tpid;
    vlan-id number;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

Description

For dynamic interfaces, define the rewrite profile to be applied to outgoing frames on this logical interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution

outputs (Analytics)

IN THIS SECTION

- Syntax | **12025**
- Hierarchy Level | **12026**
- Description | **12026**
- Options | **12026**
- Required Privilege Level | **12028**
- Release Information | **12028**

Syntax

```
outputs {
    file {
        parameters {
            path file-path;
        }
    }
    kafka {
        parameters {
            server ip-address;
            topic topic-name;
            encoding encoding-type;
        }
    }
```

```
    output-ipfix {
        parameters {
            collector-address ip-address;
            collector-ca-certificate file-path;
            collector-certificate file-path;
            collector-certificate-key file-path;
            collector-connection-retry-interval seconds;
            collector-tcp-port port-number;
            collector-vrf-name vrf-name;
        }
    }
}
```

Hierarchy Level

```
[edit services analytics agent service-agents agent-name]
```

Description

Configure parameters for the Network Telemetry Framework (NTF) agent output plug-in.

> **NOTE**: When you modify the output plug-in configuration of a service agent instance, the associated service agent daemon is restarted.

Options

file parameters	Configure parameters for sending data in a log file to a data collector.
path *pathname*	Path for the log file to which to save the data. For example, `path /tmp/example_file.log`
kafka parameters	Configure parameters for sending data to a Kafka data collector.
server *ip-address*	IP address of the Kafka server.
topic *filename*	Kafka topic name. The naming convention of the topic is *server-name*.jti.*encoding-type*. The encoding type options are `avro`, `json`, or `msgpack`.
encoding *encoding-type*	Encoding type. Options are `avro`, `json`, or `msgpack`.

output-ipfix parameters

Configure parameters for the IPFIX mediation service agent to send the IPFIX records that have been consolidated on the router to the IPFIX collector.

You must configure the IP address of the upstream IPFIX collector. When you optionally configure at least one of the collector certificate options (`collector-ca-certificate`, `collector-certificate`, and `collector-certificate-key`), the IPFIX mediator attempts to use TLS to connect with the collector. Otherwise, the mediator uses a TCP connection.

> **NOTE**: Any change you make to an existing `output-ipfix` output plug-in configuration restarts the IPFIX service agent daemon to apply the changes.

collector-address *ip-address*

IP address of the upstream IPFIX collector.

collector-ca-certificate *file-path*

(Optional) Path for the certificate, provided by a trusted certificate authority (CA), that is used to sign the peer certificate at the peer (IPFIX collector) level. The certificate is expected to be in .pem container format.

collector-certificate *file-path*

(Optional) Path for the client certificate that the server (IPFIX collector) uses to authenticate the client and enable mutual authentication. The fully-qualified domain name (FQDN) of both the client and the server are stored in the certificate's Subject Alternative Name field when the client and server certificates are generated. The certificate is expected to be in .pem container format.

collector-certificate-key *file-path*

(Optional) Private key file that is loaded to decrypt the encrypted message sent from the peer.

collector-connection-retry-interval *seconds*

(Optional) Interval in seconds at which the output plug-in retries connecting to the IPFIX collector.

- **Range:** 1 through 25
- **Default:** 20

collector-tcp-port *port-number*

(Optional) Number of the TCP port used to connect to the IPFIX collector.

- **Default:** 4740

collector-vrf-name *vrf-name*

(Optional) Name of the VRF (routing instance) in which IPFIX packets are routed.

- **Default:** default

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.3R1.

`kafka` and `file` options added in Junos OS Release 18.4R1 on MX Series and PTX Series routers.

RELATED DOCUMENTATION

- *Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data*
- *Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator*
- Configuring NTF Agent
- *IPFIX Mediation on the BNG*
- *Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector*

over (Application Identification)

IN THIS SECTION

- Syntax | **12029**
- Hierarchy Level | **12029**
- Hierarchy Level | **12029**
- Description | **12029**
- Options | **12030**
- Required Privilege Level | **12030**
- Release Information | **12030**

Syntax

```
over protocol-type  {
    chain-order;
    signature name  {
        member name  {
            context {
                context;
            }
            direction {
                any;
                client-to-server;
                server-to-client;
            }
            pattern pattern;
            depth byte-number;
        }
        port-range value;
```

Hierarchy Level

```
            [edit services application-identification application]
```

Hierarchy Level

[edit services application-identification application *application-name*]

Description

Specify set of L4/L7 application that carries given application

Configure a custom signature based on Layer 4/Layer 7 applications. You create Layer 7-based custom application signatures for the identification of multiple applications running on the same Layer 7 protocols. For example, applications such as Facebook and Yahoo Messenger can both run over HTTP, but there is a need to identify them as two different applications running on the same Layer 7 protocol.

Options

protocol-type	Application protocol
"chain-order" on page 3276	The order of members is used to match the pattern (Next Gen Services)
signature *name*	Name of the custom application signature. Must be a unique name with a maximum length of 63 characters.
member *name*	Member name for a custom application signature. Custom signatures can contain multiple members that define attributes for an application. (The supported member name range is m01 through m15.)
context	Service-specific context, such as http-header-content-type.
direction	Connection direction of the packets to match pattern
patterns	(Optional) Deterministic finite automaton (DFA) pattern matched on the context. The DFA pattern specifies the pattern to be matched for the signature. Maximum length is 128.
depth	Maximum number of bytes to check for context match. Use the byte limit for AppID to identify custom application pattern for applications running over TCP or UDP or Layer 7 applications.
port-range	Port range. This option is applicable for TCP or UDP-based applications only.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D40.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Understanding Junos OS Application Identification Custom Application Signatures

over-limit

IN THIS SECTION

- Syntax | 12031
- Hierarchy Level | 12031
- Description | 12031
- Options | 12031
- Required Privilege Level | 12032
- Release Information | 12032

Syntax

```
over-limit (block | log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name traffic-options sessions-per-client]
```

Description

In an attempt to consume all available resources and hinder the ability of the device, a malicious user might generate a large amount of traffic all at once. To prevent such activity from succeeding, you can impose a session throttle to limit sessions and configure an action to occur when the limit is exceeded.

Options

- `block`—Log the error and deny the traffic

- `log-and-permit`—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Security | 4589

overall-packet (Flat-File Accounting Options)

IN THIS SECTION

- Syntax | 12032
- Hierarchy Level | 12033
- Description | 12033
- Options | 12033
- Required Privilege Level | 12034
- Release Information | 12034

Syntax

```
overall-packet {
    all-fields;
    input-bytes;
```

```
    input-discards;
    input-errors;
    input-packets;
    inputv6-bytes;
    inputv6-packets;
    output-bytes;
    output-errors;
    output-packets;
    outputv6-bytes;
    outputv6-packets;
    input-v4-bytes;
    input-v4-packets;
    output-v4-bytes;
    output-v4-packets;
    input-bytes-per-sec;
    input-packets-per-sec;
}
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name fields]
```

Description

Specify overall packet statistics to be collected for the interface.

Options

all-fields Collect all overall packet statistics available for the interface context, logical or physical.

input-bytes Collect the number of octets received on the interface.

input-discards (Physical interfaces only) Collect the number of received packets that were discarded on the interface.

input-errors (Physical interfaces only) Collect the number of frames with errors received on the interface.

input-packets Collect the number of packets received on the interface.

input-v6-bytes	Collect the number of IPv6 octets received on the interface.
input-v6-packets	Collect the number of IPv6 packets received on the interface.
output-bytes	Collect the number of octets transmitted on the interface.
output-errors	(Physical interfaces only) Collect the number of frames that could not be transmitted on the interface because of errors.
output-packets	Collect the number of packets transmitted on the interface.
output-v6-bytes	Collect the number of IPv6 octets transmitted on the interface.
output-v6-packets	Collect the number of IPv6 packets transmitted on the interface.
input-v4-bytes	Collect the number of IPv4 octets received on the interface.
input-v4-packets	Collect the number of IPv4 packets received on the interface.
output-v4-bytes	Collect the number of IPv4 octets transmitted on the interface.
output-v4-packets	Collect the number of IPv4 packets transmitted on the interface.
input-bytes-per-sec	Collect the total number of bytes received per second.
input-packets-per-sec	Collect the total number of packets received per second.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Flat-File Accounting Overview

overflow-pool

IN THIS SECTION

- Syntax | 12035
- Hierarchy Level | 12035
- Description | 12035
- Options | 12035
- Required Privilege Level | 12036
- Release Information | 12036

Syntax

```
overflow-pool (interface | pool-name);
```

Hierarchy Level

```
[edit security nat source pool pool-name]
```

Description

Specify a source pool to use when the current address pool is exhausted. Currently the statement is applicable for IPv4 addresses only.

NOTE: The length of the IPv6 prefix must be 96 when the pool is used for NAT-PT.

Options

- `interface` — Allow the interface pool to support overflow.
- *`pool-name`* — Name of the source address pool.

> **NOTE**: The source pool must have Port Address Translation (PAT) enabled. PAT is not supported when the address is an IPv6 prefix address.

Required Privilege Level

security — To view this statement in the configuration.

security-control — To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

overhead-accounting

IN THIS SECTION

- Syntax | 12036
- Hierarchy Level | 12037
- Description | 12037
- Options | 12037
- Required Privilege Level | 12037
- Release Information | 12037

Syntax

```
overhead-accounting {
    bytes bytes;
    cell-mode cell-mode-bytes cell-mode-bytes;
    frame-mode frame-mode-bytes frame-mode-bytes;
}
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name]
```

Description

Configure the mode to shape downstream traffic.

Options

NOTE: Available options can vary depending on your platform and software release.

bytes Shaping overhead adjustment value, in bytes. This options allows you to add/subtract the number of overhead bytes to/from the packet to determine the shaped session packet length.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

`bytes` option introduced for PTX10000 Series routers in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Configuring Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates

Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates

egress-shaping-overhead | **6223**

overhead-accounting (ANCP)

IN THIS SECTION

- Syntax | 12038
- Hierarchy Level | 12038
- Description | 12038
- Required Privilege Level | 12039
- Release Information | 12039

Syntax

```
overhead-accounting;
```

Hierarchy Level

```
[edit protocols ancp interfaces interface-name]
```

Description

Prevent ANCP from performing an adjustment on the actual downstream data rate that ANCP receives from the DSLAM for the difference between the customer premise equipment (CPE) protocol overhead and the B-RAS protocol overhead. You include this statement when you want CoS to perform the adjustment on the data rate from the DSLAM according to the overhead accounting configuration in a CoS traffic control profile.

When this statement is not configured (the default condition), ANCP makes the traffic rate adjustment according to the configuration of the `qos-adjust-line-type` statements and reports that rate to CoS. CoS then applies (if configured) the adjustment set by the `overhead-accounting` statement in the CoS traffic profile.

NOTE: ANCP reports a traffic rate to CoS only if the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level has been configured.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent

overhead-accounting (Dynamic Traffic Shaping)

IN THIS SECTION

- Syntax | 12039
- Hierarchy Level | 12040
- Description | 12040
- Default | 12040
- Options | 12040
- Required Privilege Level | 12040
- Release Information | 12040

Syntax

```
overhead-accounting {
    bytes bytes;
    cell-mode cell-mode-bytes cell-mode-bytes;
    frame-mode frame-mode-bytes frame-mode-bytes;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

Configure the mode to shape downstream ATM traffic based on either frames or cells.

Default

The default is `frame-mode`.

Options

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

- *CoS Adjustment Control Profiles Overview*
- *Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates*
- *Bandwidth Management for Downstream Traffic in Edge Networks Overview*
- egress-shaping-overhead

overlay-ecmp

IN THIS SECTION

- Syntax | 12041
- Hierarchy Level | 12041
- Description | 12041
- Required Privilege Level | 12042
- Release Information | 12042

Syntax

```
overlay-ecmp;
```

Hierarchy Level

```
[edit forwarding-options vxlan-routing]
```

Description

Enable two-level equal-cost multipath next hops.

On QFX5110, QFX5120, EX4650, EX4400, and EX4100 switches, you must configure this statement for Ethernet VPN Virtual Extensible LAN (EVPN-VXLAN) overlay networks when overlay routing is configured (e.g. Overlay routing using pure Type 5 or IRB). For more information about configuring pure type-5 routes, see "ip-prefix-routes" on page 9250. For more information on Overlay Routing using an IRB interface, see *Using a Default Layer 3 Gateway to Route Traffic in an EVPN-VXLAN Overlay Network*.

NOTE: Configuring this statement causes the Packet Forwarding Engine to restart. Restarting the Packet Forwarding Engine interrupts all forwarding operations. Therefore, we strongly recommend using this configuration statement before the EVPN-VXLAN network becomes operational.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

vxlan-routing | 31007

overlay (Packet Forwarding Options)

IN THIS SECTION

- Syntax | 12042
- Hierarchy Level | 12043
- Description | 12043
- Options | 12043
- Required Privilege Level | 12043
- Release Information | 12043

Syntax

```
overlay {
    interface-num interface-num;
    next-hop next-hop;
}
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

Configure overlay interface and next-hop size. Configuration will be followed by a PFE restart for the configuration to take effect.

Options

interface-num Number of L3 Interfaces. Must be multiple of 1024. This will restart PFE.

- **Range:** 1024 through 7168

  Default is 4096

next-hop Number of L3 Next-hops. Must be multiple of 4096. This will restart PFE.

- **Range:** 4096 through 28762

  Default is 16384

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 24.4R1

RELATED DOCUMENTATION

packet-forwarding-options | 12145

overlay-path

IN THIS SECTION

- Syntax | 12044
- Hierarchy Level | 12044
- Description | 12044
- Options | 12045
- Required Privilege Level | 12045
- Release Information | 12045

Syntax

```
overlay-path overlay-path-name {
    probe-path {
        local ip-address;
        remote ip-address
    }
    tunnel-path {
        local ip-address;
        remote ip-address
    }
}
```

Hierarchy Level

[edit security advance-policy-based-routing]

Description

Configure overlay path to specify the destinations to which the active probe data needs to be sent. Overlay paths are configured for all overlay endpoints. Overlay path configuration includes two set of IP addresses—tunnel IP addresses and probe IP addresses.

You need to create the overlay setup between local and remote endpoints on both ends of the overlay (spoke device and hub device).

Options

overlay-path-name	Overlay path name.
probe-path	Probe IP addresses are used as probes' start and end addresses to send over the corresponding tunnel paths. Probe IP addresses must be unique across individual overlay paths.
	local *ip-address* — IP address of the local device.
	remote *ip-address* — IP address of the remote device.
tunnel-path	Start and end IP addresses of a tunnel. Tunnel IP addresses must be unique across individual overlay paths.
	local *ip-address* — IP address of the local device.
	remote *ip-address* — IP address of the remote device.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Application Quality of Experience

Advanced Policy-Based Routing

overload-pool

IN THIS SECTION

- Syntax | 12046
- Hierarchy Level | 12046
- Description | 12046
- Options | 12046
- Required Privilege Level | 12046
- Release Information | 12047

Syntax

```
overload-pool overload-pool-name;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated]
```

Description

Specify an address pool that can be used if the source pool becomes exhausted.

Options

overload-pool-name—Name of the overload pool.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

overload-prefix

IN THIS SECTION

- Syntax | 12047
- Hierarchy Level | 12047
- Description | 12047
- Options | 12048
- Required Privilege Level | 12048
- Release Information | 12048

Syntax

```
overload-prefix overload-prefix;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated]
```

Description

Specify the prefix that can be used if the source pool becomes exhausted.

Options

overload-prefix—Prefix value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

overload (Protocols IS-IS)

IN THIS SECTION

- Syntax | 12048
- Hierarchy Level | 12049
- Description | 12049
- Options | 12050
- Required Privilege Level | 12051
- Release Information | 12051

Syntax

```
overload {
    advertise-high-metrics;
    allow-route-leaking;
```

```
    external-prefixes;
    internal-prefixes;
    timeout seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Configure the local routing device so that it appears to be overloaded. This statement causes the routing device to continue participating in IS-IS routing, but prevents it from being used for transit traffic. Traffic destined to immediately attached subnets continues to transit the routing device.

You can also advertise maximum link metrics in network layer reachability information (NLRI) instead of setting the overload bit.

You configure or disable overload mode in IS-IS with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the IS-IS instance started is less than the specified timeout.

A timer is started for the difference between the timeout and the time elapsed since the instance started. If the time elapsed after the IS-IS instance is enabled is less than the specified timeout, overload mode is set. When the timer expires, overload mode is cleared. In overload mode, the routing device IS-IS advertisements are originated with the overload bit set. This causes the transit traffic to take paths around the routing device. However, the overloaded routing device's own links are still accessible.

The value of the overload bit depends on these three scenarios:

1. When the overload bit has already been set to a given value and the routing process is restarted: Link-state PDUs are regenerated with the overload bit cleared.

2. When the overload bit is reset to a lesser value while the routing process is running: Link-state PDUs are regenerated with the overload bit cleared.

3. When the overload bit is reset to a greater value while the routing process is running: Link-state PDUs are regenerated with the overload bit set to the difference between the old and new value.

To understand the reason for setting the overload bit, consider that BGP converges slowly. It is not very good at detecting that a neighbor is down because it has slow-paced keepalive timers. Once the BGP neighbor is determined to be down, it can take up to 2 minutes for a BGP router to declare the neighbor down. IS-IS is much quicker. IS-IS only takes 10-30 seconds to detect absent peers. It is the slowness of BGP, more precisely the slowness of internal BGP (IBGP), that necessitates the use of the overload bit. IS-IS and BGP routing are mutually dependent on each other. If both do not converge at the same time, traffic is dropped without notification (black holed).

You might want to configure the routing device so that it appears to be overloaded when you are restarting routing on the device. Setting the overload bit for a fixed amount of time right after a restart of the routing protocol process (rpd) ensures that the router does not receive transit traffic while the routing protocols (especially IBGP) are still converging.

Setting the overload bit is useful when performing hardware or software maintenance work on a routing device. After the maintenance work, clear the overload bit to carry on forwarding transit traffic. Manual clearing of the overload bit is not always possible. What is needed is an automated way of clearing the overload bit after some amount of time. Most networks use a time value of 300 seconds. This 5-minute value provides a good balance, allowing time to bring up even large internal IBGP meshes, while still relatively quick.

Another appropriate application for setting for the overload bit is on dedicated devices such as BGP route reflectors, which are intentionally not meant to carry any transit traffic. In this case, you would not use the timer.

You can verify that the overload bit is set by running the `show isis database` command.

Options

`advertise-high-metrics`—Advertise maximum link metrics in NLRIs instead of setting the overload bit.

The `advertise-high-metric` setting is only valid while the routing device is in overload mode. When `advertise-high-metric` is configured, IS-IS does not set the overload bit. Rather, it sets the metric to 63 or 16,777,214, depending whether wide metrics are enabled. This allows the overloaded routing device to be used for transit as a last resort.

An L1-L2 router in overload mode stops leaking route information between L1 and L2 levels and clears its attached bit. This is also true when `advertise-high-metrics` is configured.

`allow-route-leaking`—Enable leaking of route information into the network even if the overload bit is set.

> **NOTE**: The `allow-route-leaking` option does not work if the routing device is in dynamic overload mode. Dynamic overload can occur if the device has exceeded its resource

limits, such as the prefix limit. You can use the `auto-recovery` command to automatically exit from the IS-IS dynamic overload state.

`external-prefixes`—Enable this option to stop receiving local traffic destined to external prefixes unless this is the only router in the network that hosts the prefix. External prefixes are prefixes that IS-IS receives from other protocols. The overloaded router advertises these external prefixes with a high metric to stop receiving traffic. After the configured `timeout` lapses regular metric is restored to resume traffic reception.

`internal-prefixes`—Enable this option to stop receiving local traffic destined to internal prefixes unless this is the only router in the network that hosts the prefix. Internal prefixes are local IS-IS prefixes that the overloaded router advertises with a high metric to stop receiving traffic. After the configured `timeout` lapses regular metric is restored to resume traffic reception.

`timeout` *seconds*—Number of seconds at which the overloading is reset.

- **Range:** 60 through 1800 seconds
- **Default:** 0 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`external-prefixes` and `internal-prefixes` options introduced in Junos OS Release 18.3R1.

overload (Protocols OSPF)

IN THIS SECTION

- Syntax | 12052
- Hierarchy Level | 12052
- Description | 12053

- Options | 12053
- Required Privilege Level | 12053
- Release Information | 12053

Syntax

```
overload {
    timeout seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (oospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf topology (default | ipv4-multicast |
name)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
topology (default | ipv4-multicast | name)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf topology (default | ipv4-multicast | name)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf topology (default | ipv4-multicast |
name],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Configure the local routing device so that it appears to be overloaded. You might do this when you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic.

> **NOTE**: Traffic destined to directly attached interfaces continues to reach the routing device.

Options

`timeout` *seconds*—(Optional) Number of seconds at which the overloading is reset. If no timeout interval is specified, the routing device remains in overload state until the `overload` statement is deleted or a timeout is set.

- **Range:** 60 through 1800 seconds
- **Default:** 0 seconds

The timeout is configured with a prefix-limit. If the number of prefixes exceeds the configured limit, the overload state is reached. The routing device remains in the overload state even though the prefixes have been reduced under the limit. Therefore, you need to clear the overload state using the "clear (ospf | ospf3) overload" on page 3793 command.

> **NOTE**: Multitopology Routing does not support the `timeout` option.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for Multitopology Routing introduced in Junos OS Release 9.0.

Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring OSPF to Make Routing Devices Appear Overloaded

Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths

override

IN THIS SECTION

- Syntax | 12054
- Hierarchy Level | 12054
- Description | 12054
- Options | 12055
- Required Privilege Level | 12055
- Release Information | 12055

Syntax

```
override tag vlan-tag dynamic-profile profile name;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges],
[edit interfaces interface-name auto-configure stacked-vlan-ranges]
```

Description

Override dynamic profile assignment to individual VLANs that are already part of a previously defined VLAN range and dynamic profile.

Options

vlan-tag—VLAN tag that you want to override.

profile-name—Name of the dynamic profile that you want to use when overriding the specified VLAN tag.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Overriding the Dynamic Profile Used for an Individual VLAN

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

override-chap-password

IN THIS SECTION

- Syntax | 12056
- Hierarchy Level | 12056
- Release Information | 12056
- Description | 12056
- Required Privilege Level | 12056

Syntax

```
override-chap-password password;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Release Information

Statement introduced in Junos OS Release 21.3R1.

Description

Configure a CHAP password to override the existing password for authenticating any subscriber associated with the domain map. The override CHAP password replaces the CHAP challenge response that the PPP client sends to the server for authentication.

The CHAP challenge response is normally sent in the RADIUS Access-Request message as the CHAP-Password attribute (attribute 3). When you configure an override CHAP password, it is sent in the Access-Request as the User-Password attribute (attribute 2), and the CHAP-Password attribute is not included in the message.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement in the configuration.

RELATED DOCUMENTATION

Changing the Username and Password to Simplify Off-Chassis Provisioning

override-firewall

IN THIS SECTION

- Syntax | 12057
- Hierarchy Level | 12057
- Description | 12057
- Required Privilege Level | 12057
- Release Information | 12058

Syntax

```
override-firewall;
```

Hierarchy Level

```
[edit class-of-service host-outbound-traffic],
[edit dynamic-profiles name class-of-service host-outbound-traffic]
```

Description

Junos OS release 16.1 and earlier support `forwarding-class` and `dscp-code-point` firewall filters. These settings would take precedence over the same configuration if set at the `class-of-service host-outbound-traffic` level of the hierarchy, as applied to Routing Engine output traffic.

Starting with Junos OS release 16.1, you can set `override-firewall` to have `forwarding-class` and `dscp-code-point` firewall filters set at the `class-of-service host-outbound-traffic` level of the hierarchy take precedence.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1

RELATED DOCUMENTATION

dscp-code-point (CoS Host Outbound Traffic) | 5991
forwarding-class (CoS Host Outbound Traffic) | 7417

override-interval

IN THIS SECTION

- Syntax | 12058
- Hierarchy Level | 12058
- Description | 12059
- Options | 12059
- Required Privilege Level | 12059
- Release Information | 12059

Syntax

```
override-interval milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name],
```

```
[edit protocols pim],
[edit protocols pim interface (Protocols PIM)  interface-name],
[edit routing-instances routing-instance-name protocols pim]
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name]
```

Description

Set the maximum time in milliseconds to delay sending override join messages for a multicast network that has join suppression enabled. When a router or switch sees a prune message for a join it is currently suppressing, it waits for the interval specified by the override timer before it sends an override join message.

Options

This is a random timer with a value in milliseconds.

- **Range:** 0 through maximum override value
- **Default:** 2000 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Example: Enabling Join Suppression

propagation-delay | **13381**

reset-tracking-bit | **15583**

override-password (Domain Map)

IN THIS SECTION

- Syntax | 12060
- Hierarchy Level | 12060
- Description | 12060
- Options | 12060
- Required Privilege Level | 12061
- Release Information | 12061

Syntax

```
override-password password
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Specify a password to be used as the authenticating password for all outgoing authentication requests for subscribers who match the domain map.

NOTE: This override works only when PAP is the authentication method. It does not work for CHAP.

Options

password	Name of the password.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Changing the Username and Password to Simplify Off-Chassis Provisioning

override (PIM Static RP)

IN THIS SECTION

- Syntax | 12061
- Hierarchy Level | 12062
- Description | 12062
- Default | 12062
- Required Privilege Level | 12062
- Release Information | 12062

Syntax

```
override;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp local],
[edit logical-systems logical-system-name protocols pim rp local family inet],
[edit logical-systems logical-system-name protocols pim rp local family inet6],
[edit logical-systems logical-system-name protocols pim rp static address address],
[edit logical-systems logical-system-name routing-instances instance-name protocols pim rp
local],
[edit logical-systems logical-system-name routing-instances instance-name protocols pim rp local
family inet],
[edit logical-systems logical-system-name routing-instances instance-name protocols pim rp local
family inet6],
[edit logical-systems logical-system-name routing-instances instance-name protocols pim rp
static address address],
[edit protocols pim rp local],
[edit protocols pim rp local family inet],
[edit protocols pim rp local family inet6],
[edit protocols pim rp static address address],
[edit routing-instances instance-name protocols pim rp local],
[edit routing-instances instance-name protocols pim rp local family inet],
[edit routing-instances instance-name protocols pim rp local family inet6],
[edit routing-instances instance-name protocols pim rp static address address]
```

Description

When you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for a given group range, and allow dynamic RP mapping for all other groups.

Default

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Static RP

Configuring PIM Auto-RP

override (RADIUS Options)

IN THIS SECTION

- Syntax | 12063
- Hierarchy Level | 12063
- Description | 12063
- Options | 12064
- Required Privilege Level | 12064
- Release Information | 12064

Syntax

```
override {
    calling-station-id remote-circuit-id;
    nas-ip-address tunnel-client-gateway-address;
    nas-port tunnel-client-nas-port;
    nas-port-type tunnel-client-nas-port-type;
}
```

Hierarchy Level

```
[edit access profile profile-name radius options]
```

Description

Override the values for certain RADIUS options.

Options

calling-station-id remote-circuit-id

Override the Calling-Station-ID format for the string that the LAC sends to the LNS in the Calling Number AVP 22. AVP 22 is included in the incoming-call request (ICRQ) packet when the L2TP session is being established.

If you have configured the `calling-station-id-format` statement, the LAC sends AVP 22 in a different format that can include other information in addition to the ACI or ARI. In this case, the LNS sends the value of AVP 22 to the RADIUS server as the Calling-Station-ID attribute (31).

If you do not want the AVP 22 value to be used as the Calling-Station-ID attribute (31) for some subscribers, you can override that format in the access profile by including the `calling-station-id remote-circuit-id` override option. You must also configure the new format to be used with the `remote-circuit-id-format` statement; you can specify: the ACI, ARI, or both the ACI and ARI received from the PADR packet.

nas-ip-address tunnel-client-gateway-address

Check the SDB to determine whether the session's LAC endpoint IP address is available. If so, use that value in the RADIUS NAS-IP-Address attribute (4). The LNS subsequently sends the attribute to the RADIUS server in Access-Request and accounting messages. If the value is not present, the attribute is not sent.

nas-port tunnel-client-nas-port

Check the SDB to determine whether the LAC NAS port information is conveyed to the LNS in the Cisco Systems NAS Port Info AVP (100). If so, use that value in the RADIUS NAS-Port attribute (5). The LNS subsequently sends the attribute to the RADIUS server in Access-Request and accounting messages. If the value is not present, the LNS sends its local IP address.

nas-port-type tunnel-client-nas-port-type

Check the SDB to determine whether the LAC NAS port information is conveyed to the LNS in the Cisco Systems NAS Port Info AVP (100). If so, override the value of the RADIUS NAS-Port-Type attribute (61) at the LNS with that value. Otherwise, use the original NAS-Port-Type in attribute 61.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3R1.

`nas-ip-address tunnel-client-gateway-address` and `nas-port tunnel-client-nas-port` options added in Junos OS Release 17.3R1 on MX Series.

`calling-station-id remote-circuit-id` option added in Junos OS Release 13.3R1 on MX Series.

RELATED DOCUMENTATION

L2TP Subscriber Access Lines and Connection Speeds

RADIUS Servers and Parameters for Subscriber Access

overrides (DHCP Local Server)

IN THIS SECTION

- Syntax | 12065
- Hierarchy Level | 12067
- Description | 12067
- Required Privilege Level | 12068
- Release Information | 12068

Syntax

```
asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match <option60-and-option82 | incoming-interface>;
    client-negotiation-match incoming-interface;
    delay-advertise {
        based-on (option-15 | option-16 | option-18 | option-37) {
            equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            not-equals {
                ascii ascii-string;
```

```
                hexadecimal hexadecimal-string;
            }
            starts-with {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    delay-offer {
        based-on (option-60 | option-77 | option-82) {
            equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            not-equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            starts-with {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    delegated-pool;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
    include-option-82 {
        forcerenew;
        nak;
    }
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    protocol-attributes attribute-set-name;
    rapid-commit;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server group group-name interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Override the default configuration settings for the extended DHCP local server. Specifying the `overrides` statement with no subordinate statements removes all DHCP local server overrides at that hierarchy level.

- To override global DHCP local server configuration options, include the `overrides` statement and its subordinate statements at the `[edit system services dhcp-local-server]` hierarchy level.
- To override configuration options for a named group of interfaces, include the statements at the `[edit system services dhcp-local-server group group-name]` hierarchy level.
- To override configuration options for a specific interface within a named group of interfaces, include the statements at the `[edit system services dhcp-local-server group group-name interface interface-name]` hierarchy level.
- Use the `[edit system services dhcp-local-server dhcpv6]` hierarchy level to override DHCPv6 configuration options.

> **NOTE**: By default, `jdhcp` does not process DHCPINFORM message. Only after you enable the overrides command using the `set system services dhcp-local-server overrides process-inform` statement, `jdhcp` starts processing the DHCPINFORM message.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

The `interface-client-limit` statement is not supported in the `[edit system services dhcp-local-server dhcpv6]` hierarchy level.

The `asymmetric-prefix-lease-time`, `delegated-pool`, `multi-address-embedded-option-response`, and `rapid-commit` statements are supported in the `[edit system services dhcp-local-server dhcpv6 ...]` hierarchy level only.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support for the `allow-no-end` option introduced in Junos OS Release 14.1X53-D15 for EX Series switches.

RELATED DOCUMENTATION

Understanding Differences Between Legacy DHCP and Extended DHCP

Overriding the Default DHCP Local Server Configuration Settings

Configuring a DHCP Server on Switches

overrides (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 12069
- Hierarchy Level | 12069
- Description | 12070
- Required Privilege Level | 12070
- Release Information | 12071

Syntax

```
overrides {
    allow-no-end-option;
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    bootp-support;
    based-on {
     (option-60 {..} | option-77 {..} | option-82 {..})

    client-discover-match <option60-and-option82 | incoming-interface>;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    disable-relay;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    relay-source
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
```

```
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Override the default configuration settings for the extended DHCP relay agent. Specifying the `overrides` statement with no subordinate statements removes all DHCP relay agent overrides at that hierarchy level. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

The following statements are supported at both the `[edit ... dhcp-relay]` and `[edit ... dhcpv6]` hierarchy levels.

- `allow-snooped-clients`
- `asymmetric-lease-time`
- `delete-binding-on-renegotiation`
- `dual-stack`
- `interface-client-limit`
- `no-allow-snooped-clients`
- `no-bind-on-request`
- `relay-source`
- `send-release-on-delete`

The following statements are supported at the `[edit ... dhcpv6]` hierarchy levels only.

- `asymmetric-prefix-lease-time`

All other statements are supported at the `[edit ... dhcp-relay]` hierarchy levels only.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Overriding the Default DHCP Relay Configuration Settings

overrides (DHCP Security)

IN THIS SECTION

- Syntax | 12071
- Hierarchy Level | 12072
- Description | 12072
- Required Privilege Level | 12072
- Release Information | 12072

Syntax

```
overrides {
    no-dhcpv6-options;
    no-option16;
    no-option18;
    no-option37;
    no-option82;
    trusted;
    untrusted;
}
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name]
```

Description

Modify selected DHCP attributes for a group of interfaces that is configured within a specified VLAN.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

Support for the `no-option37` option introduced in Junos OS Release 13.2X51-D20 for EX Series switches.

Support for the `no-dhcpv6-options`, `no-option16` and `no-option18` options introduced in Junos OS Release 14.2 for EX Series switches.

RELATED DOCUMENTATION

Enabling a Trusted DHCP Server (ELS)

Understanding DHCP Snooping (non-ELS)

Understanding DHCP Option 82

overrides (DHCP Security for MX Series)

IN THIS SECTION

- Syntax | 12073
- Hierarchy Level | 12073

- Description | 12073
- Options | 12073
- Required Privilege Level | 12073
- Release Information | 12074

Syntax

```
overrides (trusted | untrusted | no-option82);
```

Hierarchy Level

```
[edit bridge-domains  bridge-domain-names forwarding-options dhcp-security group group-name]
```

Description

Modify selected attributes of a specific interface within a group of interfaces configured within a specified bridge domain.

Options

no-option 82 The interface specified in this group does not support DHCP option 82.

trusted The interface specified in this group is trusted. DHCP snooping does not apply to the trusted interface. Likewise, DAI and IP source guard—even if they are enabled for the VLAN or bridge domain—do not apply to the interface that is configured with the `overrides` and the `trusted` options. Trunk interfaces are trusted by default. Access interfaces are untrusted by default.

untrusted The interface specified in this group is untrusted. Trunk interfaces are trusted by default. Access interfaces are untrusted by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring DHCP Option 82 on a Router with Bridge Domain

Understanding DHCP Option 82

overrides (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 12074
- Hierarchy Level | 12075
- Description | 12075
- Options | 12075
- Required Privilege Level | 12076
- Release Information | 12077

Syntax

```
overrides {
    event {
        catastrophic-failure {
            reboot (master | standby);
        }
    }
    force-show-arp-resolve;
    interfaces {
        family (inet | inet6) {
```

```
            layer2-liveness-detection;
            ipoe-dynamic-arp-enable;
            receive-gratuitous-arp;
        }
    }
    no-unsolicited-ra;
    ra-initial-interval-max seconds;
    ra-initial-interval-min seconds;
    shmlog {
        disable;
        file filename <files maximum-no-files> <size maximum-file-size>;
        filtering enable;
        log-name {
            all;
            logname {
                <brief | detail | extensive | none | terse>;
                <file-logging |no-file-logging>;
            }
            }
        log-type (debug | info | notice);
    |
}
```

Hierarchy Level

```
[edit system services subscriber-management]
```

Description

Override the default configuration settings for the Junos OS enhanced subscriber management software for subscriber management.

Options

force-show-arp-resolve
Force the output of the `show arp` command to display the IP address by the hostname of each device.

ra-initial-interval-max *seconds*
Specify the high end of the range from which the router randomly selects an interval for sending the first three unsolicited IPv6 router advertisement messages. You must also configure the `ra-initial-interval-min` option.

- **Range:** 1 through 16

ra-initial-interval-min *seconds*

Specify the low end of the range from which the router randomly selects an interval for sending the first three unsolicited IPv6 router advertisement messages. You must also configure the `ra-initial-interval-max` option.

> **BEST PRACTICE**: Always configure the value of `ra-initial-interval-min` to be less than or equal to the value of `ra-initial-interval-max`. If you configure the values to be the same, the initial router advertisement intervals are constant and not randomized.

- **Range:** 1 through 16

ipoe-dynamic-arp-enable

Enable dynamic ARP to resolve the MAC address for IPv4 framed host (32-bit) routes. By default the framed route is permanently associated with the source MAC address received in the packet that triggered creation of the dynamic VLAN.

receive-gratuitous-arp

Enable the router to compare the source MAC address received in a gratuitous ARP request or reply packet with the value in the ARP cache. The router updates the cache with the received MAC address when it determines this address is different from the cache entry.

This situation occurs when an IPv4 address is moved to a different device. The device broadcasts a gratuitous ARP reply packet with its MAC address as the source MAC address. When the `receive-gratuitous-arp` option is configured, the router compares the MAC addresses and updates the cache to associate the IPv4 address with the new MAC address.

If the `receive-gratuitous-arp` option is not configured, the router does not accept the gratuitous ARP request or reply packet and cannot quickly learn about the new address. Instead, the original dynamic ARP entry in the cache eventually times out. Before deleting the entry, the router sends an ARP request for the target IP address. The client responds with the new MAC address. This delay in learning about the new address means there is a period during which the MAC address in the ARP cache does not match the address in the new device's NIC.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

`ra-initial-interval-max` and `ra-initial-interval-min` options added in Junos OS Release 18.2R1.

`ipoe-dynamic-arp-enable` and `receive-gratuitous-arp` options added in Junos OS Release 18.4R1.

`force-show-arp-resolve` option added in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

Junos OS Enhanced Subscriber Management Overview

DHCP Liveness Detection Using ARP and Neighbor Discovery Packets

Configuring an Interval Range for Unsolicited Router Advertisements to IPv6 Neighbors

overrides (System Services DHCP)

IN THIS SECTION

- Syntax | 12077
- Hierarchy Level | 12080
- Description | 12080
- Options | 12080
- Required Privilege Level | 12080
- Release Information | 12081

Syntax

```
overrides
{
```

```
interface-client-limit number;
always-add-option-dns-server;
always-process-option-request-option;
asymmetric-lease-time seconds;
asymmetric-prefix-lease-time seconds;
client-negotiation-matchincoming-interface;
delay-advertise {
 based-on {
  option-15 {
   equals {
    ascii name;

    hexadecimal name;
   }
   not-equals {
    ascii name;
    hexadecimal name;
   }
   starts-with {
    ascii name;
    hexadecimal name;
   }
  }
  option-16 {
   equals {
    ascii name;

    hexadecimal name;
   }
   not-equals {
    ascii name;
    hexadecimal name;
   }
   starts-with {
    ascii name;
    hexadecimal name;
   }
  }
  option-18 {
   equals {
    ascii name;

    hexadecimal name;
```

```
      }
      not-equals {
       ascii name;

       hexadecimal name;

      }
      starts-with {
       ascii name;

       hexadecimal name;

      }
     }
     option-37 {
      equals {
       ascii name;

       hexadecimal name;

      }
      not-equals {
       ascii name;

       hexadecimal name;

      }
      starts-with {
       ascii name;

       hexadecimal name;

      }
     }
    }
    delay-time seconds;
   }
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6]
[edit system services dhcp-local-server dhcpv6 group group-name]
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name]
```

Description

Override the default configuration settings for the extended DHCP local server. Specifying the `overrides` statement with no subordinate statements removes all DHCP local server overrides at that hierarchy level.

- To override global DHCP local server configuration options, include the `overrides` statement and its subordinate statements at the `[edit system services dhcp-local-server]` hierarchy level.
- To override configuration options for a named group of interfaces, include the statements at the `[edit system services dhcp-local-server dhcpv6 group group-name]` hierarchy level.
- To override configuration options for a specific interface within a named group of interfaces, include the statements at the `[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name]` hierarchy level.
- Use the DHCPv6 hierarchy levels to override DHCPv6 configuration options.

Options

`interface-client-limit number`—Sets the maximum number of DHCP clients per interface allowed for a specific group or for all groups. A group specification takes precedence over a global specification for the members of that group.

- **Range:** 1 through 500,000
- **Default:** No limit

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Additional Configurations for DHCP Clients

oversize-message-drop

IN THIS SECTION

- Syntax | 12081
- Hierarchy Level | 12081
- Description | 12082
- Options | 12082
- Required Privilege Level | 12082
- Release Information | 12082

Syntax

```
oversize-message-drop;
```

Hierarchy Level

```
[edit logical-systems name security alg dns],
[edit logical-systems name tenants name security alg dns],
[edit security alg dns],
[edit services alg dns],
[edit tenants name security alg dns]
```

Description

Enable or disable oversize message drop.

When you enable the `oversize-message-drop` command, Domain Name System (DNS) packets that exceed the threshold (range from 512 through 8192) are dropped.

When you disable the `oversize-message-drop` command, DNS Application Layer Gateway (ALG) will only record and forward the oversized DNS packets.

Options

`oversize-message-drop`—Drop the oversized DNS packets for which the packet length exceed the threshold value (range from 512 through 8192).

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D55.

RELATED DOCUMENTATION

dns | 5847

ovlan-cfi (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 12083
- Hierarchy Level | 12083
- Description | 12083
- Required Privilege Level | 12083

Release Information | 12083

Syntax

```
ovlan-cfi;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

CFI bit to be used in the outer VLAN tag.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | 15759

ovlan-id (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 12084
- Hierarchy Level | 12084
- Description | 12084
- Options | 12084
- Required Privilege Level | 12084
- Release Information | 12085

Syntax

```
ovlan-id number;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Configure the outer VLAN ID for the test frames. This parameter is applicable for single-tagged and dual-tagged packets.

Options

number VLAN ID number.

- **Range:** 0 through 4094

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

ovlan-priority (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | **12085**
- Hierarchy Level | **12085**
- Description | **12086**
- Options | **12086**
- Required Privilege Level | **12086**
- Release Information | **12086**

Syntax

```
ovlan-priority value;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Description

Configure the priority value for the IEEE 802.1p bit in the outer VLAN tag. This parameter is valid only for a bridge family for an Ethernet LAN (ELAN) or an Ethernet Line (E-LINE) service.

Options

value IEEE 802.1p priority value in the outer VLAN tag

- **Range:** 0 through 7

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

ovsdb

IN THIS SECTION

- Syntax | **12087**
- Hierarchy Level | **12087**
- Description | **12087**
- Default | **12087**

Required Privilege Level | 12088
Release Information | 12088

Syntax

```
ovsdb {
    controller  ip-address {
        inactivity-probe-duration milliseconds;
        maximum-backoff-duration milliseconds;
        protocol  protocol {
            port number;
        }
    }
    interfaces interface-name;
    traceoptions {
        file <filename> <files number> <match regular-expression> <no-world-readable | world-
readable> <size size>;
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure support for the *Open vSwitch Database (OVSDB)* management protocol on a Juniper Networks device.

The remaining statements are explained separately. See CLI Explorer.

Default

The OVSDB management protocol is disabled on Juniper Networks devices.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Understanding the OVSDB Protocol Running on Juniper Networks Devices

ovsdb-managed

IN THIS SECTION

- Syntax | 12088
- Hierarchy Level | 12088
- Description | 12089
- Required Privilege Level | 12089
- Release Information | 12089

Syntax

```
ovsdb-managed;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name vxlan],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name vxlan],
[edit routing-instances routing-instance-name switch-options],
```

```
[edit routing-instances routing-instance-name vlans vlan-name vxlan],
[edit routing-instances routing-instance-name vxlan],
[edit switch-options],
[edit vlans vlan-name vxlan]
```

Description

Disable a Juniper Networks device from learning about other Juniper Networks devices that function as hardware *virtual tunnel endpoints (VTEPs)* in a specified Virtual Extensible LAN (VXLAN) and the media access control (MAC) addresses learned by the hardware VTEPs. Instead, the Juniper Networks device uses the *Open vSwitch Database (OVSDB)* management protocol to learn about the hardware VTEPs in the VXLAN and the MAC addresses learned by the hardware VTEPs.

The specified VXLAN must have a VXLAN network identifier (VNI) configured, using the `vni` statement in the `[edit bridge-domains bridge-domain-name vxlan]`, `[edit routing-instance routing-instance-name vxlan]`, or `[edit vlans vlan-name vxlan]` hierarchy.

Also, for OVSDB-managed VXLANs, the multicast scheme described in *Understanding How Layer 2 BUM and Layer 3 Routed Multicast Traffic Are Handled with OVSDB* is used. Therefore, specifying the `multicast-group` statement in the `[edit bridge-domains bridge-domain-name vxlan]`, `[edit routing-instances routing-instance-name vxlan]`, or `[edit vlans vlan-name vxlan]` hierarchy has no effect.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Understanding Dynamically Configured VXLANs in an OVSDB Environment

Configuring OVSDB-Managed VXLANs

owner

IN THIS SECTION

- Syntax | 12090
- Hierarchy Level | 12091
- Description | 12091
- Options | 12091
- Required Privilege Level | 12091
- Release Information | 12091

Syntax

```
owner name {
    test test-name {
        data-fill data;
        data-size size;
        destination-port port;
        dscp-code-points (RPM) dscp-bits;
        history-size size;
        moving-average-size number;
        offload-type {
            none;
            pfe-timestamp;
        }
        probe-count count;
        probe-interval seconds;
        probe-type type;
        routing-instance instance-name;
        source-address address;
        target (url url | address address);
        test-interval interval;
        thresholds thresholds;
        traps traps;
        ttl hop-count
```

```
    }
}
```

Hierarchy Level

```
[edit services monitoring rpm]
```

Description

Specify an owner name and test options. The owner name combined with the test options for a particular test name represents a single RPM configuration instance.

Options

owner *name*

You can configure any name, up to 32 characters in length, or specify one of the pre-defined owner names:

- icmp-evo
- icmp-junos
- udp-evo
- udp-junos

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.1R1.

`traps` option introduced in Junos OS Evolved Release 21.2R1.

RELATED DOCUMENTATION

> *Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers*

owner (RMON)

IN THIS SECTION

- Syntax | 12092
- Hierarchy Level | 12092
- Description | 12092
- Options | 12092
- Required Privilege Level | 12093
- Release Information | 12093

Syntax

```
owner owner-name;
```

Hierarchy Level

```
[edit snmp rmon history]
```

Description

Specify the user or group responsible for this RMON history configuration.

Options

owner-name—The user or group responsible for this configuration.

- **Range:** 0 through 32 alphanumeric characters

Required Privilege Level

snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

RMON Overview

RMON MIB Event, Alarm, Log, and History Control Tables

oxid

IN THIS SECTION

- Syntax | 12093
- Hierarchy Level | 12094
- Description | 12094
- Default | 12094
- Options | 12094
- Required Privilege Level | 12094
- Release Information | 12094

Syntax

```
oxid (enable | disable)
```

Hierarchy Level

QFX Series Standalone Switches

```
[edit forwarding-options hash-key family fcoe]
```

QFabric Systems

```
[edit forwarding-options hash-key family fcoe ethernet-interfaces node-group (node-group-name |
all)], [edit forwarding-options hash-key family fcoe fabric-interfaces node-group (node-group-name |
all)]
```

Description

Enable or disable whether the switch uses the originator exchange identifier (OxID) field for hash control for FCoE traffic load balancing.

Default

OxID hash control is enabled by default.

Options

oxid (enable | disable) Enable or disable whether the switch uses the OxID hash control field for FCoE traffic load balancing.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Enabling and Disabling CoS OxID Hash Control for FCoE Traffic on Standalone Switches

Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches

P

IN THIS SECTION

- p2mp-ldp-next-hop | **12119**
- p2mp-lsp-next-hop | **12120**
- p2mp (Protocols LDP) | **12122**
- p2mp (Protocols MPLS) | **12124**
- p2mp-receive-switch | **12125**
- p2mp-transmit-switch | **12127**
- package (Loading on PIC) | **12128**
- packet-action | **12129**
- packet-burst (Hierarchical Policer) | **12133**
- packet-burst (Policer) | **12135**
- packet-capture | **12137**
- packet-filter | **12141**
- packet-format-match | **12143**
- packet-forwarding-options | **12145**
- packet-log (Security Flow) | **12150**
- packet-log (Security IDP Policy) | **12152**
- packet-log (Security IDP Sensor Configuration) | **12153**
- packet-loss-priority (RFC 2544 Benchmarking) | **12155**
- packet-ordering-mode (Application Services) | **12157**
- packet-scheduling | **12158**
- packet-size (RFC 2544 Benchmarking) | **12160**
- packet-timestamp | **12162**
- packet-types (Dynamic VLAN Authentication) | **12163**

- packet-timestamp | **12165**
- pado-advertise | **12166**
- pair-policy | **12168**
- pap | **12169**
- pap (Dynamic PPP) | **12171**
- pap (L2TP) | **12172**
- parameter-preference | **12173**
- parameters | **12175**
- parse-direction (Access Profile) | **12176**
- parse-direction (Domain Map) | **12178**
- parse-order (Domain Map) | **12179**
- partition | **12181**
- partition (Gx-Plus) | **12182**
- partition (NASREQ Diameter Application) | **12183**
- partition (OCS) | **12185**
- partition (PCRF) | **12187**
- partition (s6a) | **12190**
- pass-through (interfaces) | **12192**
- pass-through (Security Policy) | **12194**
- pass-through (Firewall-authentication) | **12196**
- passive (CHAP) | **12198**
- passive (Dynamic IGMP Interface) | **12199**
- passive (Dynamic MLD Interface) | **12201**
- passive (IGMP) | **12202**
- passive (MLD) | **12204**
- passive-mode | **12205**
- passive-mode-tunneling | **12207**
- passive-mode-tunneling (Security IPsec) | **12208**
- passive-monitor-mode | **12212**
- passive (MVRP) | **12213**
- passive (PAP) | **12214**
- passive-probe-params | **12216**

- passive (Protocols BGP) | **12218**
- passive (Protocols IS-IS) | **12219**
- passive (Protocols Link Management) | **12222**
- passive (Protocols OSPF) | **12223**
- passive (Routing Options) | **12225**
- password (Access) | **12225**
- password (DHCP Local Server) | **12226**
- password (DHCP Relay Agent) | **12229**
- password-file | **12231**
- password (Flow Collector File Servers) | **12232**
- password (Interfaces) | **12233**
- password (Login) | **12235**
- password (MAC RADIUS Authentication) | **12241**
- password-options | **12243**
- password (Security Antivirus) | **12245**
- password (Services) | **12246**
- password (Static Subscribers) | **12247**
- password (Transfer Log File Servers) | **12249**
- path | **12250**
- path-database-size | **12252**
- path-database-size (EX Series Switch Only) | **12253**
- path-length | **12255**
- path-mtu | **12256**
- path-mtu-discovery (Tunnel) | **12257**
- path (Protocols MPLS) | **12259**
- path (Protocols RSVP) | **12261**
- path-rate-limit | **12263**
- path-selection | **12266**
- path (Steering) | **12269**
- pattern (Application Identification) | **12271**
- pattern (Class Attribute) | **12272**
- pattern-pcre (Security IDP) | **12274**

- pattern (Security IDP) | **12275**
- pattern-update (Security Antivirus) | **12276**
- pause-computation-during-churn | **12277**
- payload | **12279**
- pba-interim-logging-interval | **12280**
- pbb-evpn-core | **12282**
- pbb-service-options | **12283**
- pcc-action-profile (PCC Rules) | **12284**
- pcc-action-profiles | **12286**
- pcc-context | **12289**
- pcc-rule | **12291**
- pcc-rulebases (PCEF) | **12293**
- pcc-rulebases (PCEF Profile) | **12294**
- pcc-rules (PCEF) | **12297**
- pcc-rules (PCEF Profile) | **12299**
- pcc-time-of-day-profiles | **12301**
- pce | **12303**
- pce-group (PCE) | **12306**
- pce-group (Protocols PCEP) | **12307**
- pce-type | **12309**
- pcef | **12311**
- pcef (Dynamic Profiles) | **12314**
- pcef-profile (Service Set) | **12315**
- pcef-profile (TDF Domain) | **12317**
- pcef-profile (TDF Domain Selection) | **12318**
- pcep | **12320**
- pci-address | **12322**
- pcp-rules | **12323**
- pcp-server | **12324**
- pcrf (Diameter Applications) | **12325**
- pdu-interval | **12328**
- pdu-threshold | **12329**

- peak-burst-size (QFX) | **12331**
- peak-burst-size | **12332**
- peak-information-rate (QFX) | **12334**
- peak-information-rate | **12336**
- peak-rate | **12338**
- peer | **12339**
- peer-active (VPLS Multihoming for FEC 129) | **12341**
- peer-as-billing-template | **12343**
- peer-as-list | **12345**
- peer-as (Protocols BGP) | **12346**
- peer-as (VPLS) | **12348**
- peer-certificate-type | **12350**
- peer (Diameter Base Protocol) | **12351**
- peer (Diameter Base Protocol) | **12353**
- peer (Diameter Network Element) | **12354**
- peer (Diameter Network Element) | **12356**
- peer-forwarding-ip | **12357**
- peer-gnf | **12359**
- peer-groups (access network peers) | **12361**
- peer-groups (control plane peers) | **12362**
- peer-groups (core network peers) | **12364**
- peer (ICCP) | **12366**
- peer-id | **12367**
- peer-interface (Protocols OSPF) | **12370**
- peer-interface (Protocols RSVP) | **12372**
- peer-ip-address-optional | **12373**
- peer-limit | **12375**
- peer (Multichassis) | **12376**
- peer-origin (Diameter Peer) | **12377**
- peer (Protocols LMP) | **12379**
- peer (Protocols MSDP) | **12380**
- peer (Protocols RIP) | **12382**

- peer-selection-service | **12384**
- peer-subnet (Easy EVPN LAG Configuration) | **12385**
- peer-to-peer (Easy EVPN LAG Configuration) | **12388**
- peer-unit | **12390**
- peer-unit (Routing Policy Condition) | **12391**
- peers (Commit) | **12393**
- peers-synchronize | **12394**
- peg-df-election | **12395**
- pem-absence | **12397**
- pem (MX Series Devices, ACX7024 Devices) | **12399**
- pem (T640, T1600, and T4000 Routers with Six-Input DC Power Supply) | **12400**
- pending-queue-watermark | **12402**
- pending-queue-watermark-abate | **12403**
- pending-sess-queue-length | **12405**
- per-esi | **12406**
- per-esi-vlan | **12408**
- per-evi | **12409**
- per-flow | **12413**
- per-flow (Aggregated Ethernet Interfaces) | **12414**
- per-group-label | **12416**
- per-prefix | **12417**
- per-prefix-label | **12419**
- per-route-accounting | **12420**
- per-session-scheduler | **12422**
- per-unit-scheduler | **12423**
- perfect-forward-secrecy (Security) | **12424**
- perfect-forward-secrecy (Services) | **12426**
- performance | **12428**
- performance-monitor | **12429**
- performance-monitoring | **12431**
- performance-monitoring (OAM LFM) | **12432**
- performance-monitoring (Protocols MPLS) | **12434**

- periodic (Chassis) | **12436**
- periodic | **12438**
- periodic (Interfaces) | **12439**
- periodic-traceroute | **12441**
- permit-command | **12443**
- permit (Security Persistent NAT) | **12444**
- permit (Security Policies) | **12446**
- permit (Security SCTP) | **12449**
- perpetual-poe | **12450**
- persistent-learning | **12451**
- persistent-nat | **12453**
- persistent-results | **12454**
- pfc | **12455**
- pfc (Input Congestion Notification) | **12456**
- pfc-priority | **12458**
- pfe | **12460**
- pfe power off | **12463**
- pgcp | **12465**
- pgcp-rules | **12466**
- pgcp-service | **12468**
- phase-adjust-threshold | **12469**
- phase-lock-threshold | **12471**
- phy-timestamping | **12472**
- physical-interface (DDoS Flow Detection) | **12474**
- physical-interface-filter | **12476**
- physical-interface-policer | **12479**
- pic | **12480**
- pic-boot-timeout | **12484**
- pic-console-authentication | **12485**
- pic (M Series and T Series Routers) | **12487**
- pic-memory-threshold | **12489**
- pic-mode | **12490**

- pic-mode (MX Series) | **12491**
- pic-mode (Chassis T1 Mode) | **12494**
- pic (MX Series Routers) | **12496**
- pic (Port) | **12498**
- pic (PTX Series Routers) | **12500**
- pic (satellite-policies port-group-alias) | **12502**
- pic-services-logging | **12504**
- pic (Subscriber Limits) | **12506**
- pim | **12507**
- pim-asm | **12514**
- pim-evpn-gateway | **12515**
- pim-snooping | **12517**
- pim-ssm (Provider Tunnel) | **12519**
- pim-ssm (Selective Tunnel) | **12520**
- pim-to-igmp-proxy | **12522**
- pim-to-mld-proxy | **12523**
- ping | **12525**
- ping clns | **12534**
- ping-death | **12537**
- ping-death (IDS Screen Next Gen Services) | **12538**
- ping fabric multicast-flow | **12539**
- ping fabric unicast-flow | **12542**
- ping-interval | **12545**
- ping mpls bgp | **12547**
- ping mpls l2circuit | **12550**
- ping mpls l2vpn | **12554**
- ping mpls l3vpn | **12558**
- ping mpls ldp | **12561**
- ping mpls lsp-end-point | **12565**
- ping mpls rsvp | **12568**
- ping mpls segment routing isis | **12575**
- ping mpls segment routing ospf | **12579**

- ping mpls segment routing spring-te | **12583**
- ping mpls segment routing static egress-ip | **12592**
- ping overlay | **12596**
- ping srv6 | **12603**
- ping vpls instance | **12608**
- pki | **12610**
- pki (Services) | **12612**
- plmn | **12615**
- poe | **12616**
- point-to-point | **12619**
- point-to-point (MVRP) | **12620**
- policer | **12622**
- policer (Applying to a Logical Interface) | **12624**
- policer (CFM Firewall) | **12626**
- policer (CFM Global) | **12627**
- policer (CFM Session) | **12629**
- policer (Configuring) | **12630**
- policer (CoS) | **12633**
- policer-drop-probability-low | **12635**
- policer (Firewall Filter Action) | **12636**
- policer (Interface) | **12638**
- policer (Layer 2 VPN) | **12639**
- policer (MAC) | **12641**
- policer-overhead | **12643**
- policer-overhead-adjustment | **12644**
- policer-charge | **12646**
- policer-charge | **12648**
- policies | **12650**
- policing | **12661**
- policing (Protocols LDP) | **12662**
- policing (Protocols MPLS) | **12664**
- policy | **12665**

- policy (advanced-policy-based-routing) | **12667**
- policy (Aggregate and Generated Routes) | **12670**
- policy-based-logging (LRF Profile) | **12672**
- policy-db-size | **12674**
- policy (Event Policy) | **12675**
- policy (Explicit-Proxy) | **12678**
- policy (Flow Maps) | **12683**
- policy (IPsec) | **12684**
- policy-lookup-cache | **12685**
- policy-map (Class of Service) | **12686**
- policy-match | **12689**
- policy (Multicast-Only Fast Reroute) | **12691**
- policy-multipath | **12694**
- policy-oids | **12696**
- policy-options | **12697**
- policy-options (QFX) | **12700**
- policy-options (Dynamic Profiles) | **12702**
- policy (PIM rpf-vector) | **12703**
- policy (Protocols Multipoint LDP) | **12706**
- policy-rematch | **12707**
- policy (Security Alarms) | **12709**
- policy (Security Group VPN IKE) | **12711**
- policy (Security IKE) | **12713**
- policy (Security Interfaces) | **12716**
- policy (Security IPsec) | **12718**
- policy (Security IPsec Interfaces) | **12722**
- policy (Security Policies) | **12724**
- policy (Services CoS Next Gen Services) | **12728**
- policy (Services IKE) | **12730**
- policy-set | **12731**
- policy (SSM Maps) | **12733**
- policy (Stateful Firewall Rules Next Gen Services) | **12734**

- policy-statement | **12736**
- policy-stats | **12743**
- policy (Subscriber Secure Policy) | **12745**
- policy (System Security Profile) | **12747**
- policy-with-count | **12749**
- policy-map-marking | **12751**
- polling-interval (Network Management) | **12752**
- polling-interval | **12753**
- pon (Access-Line Rate Adjustment) | **12755**
- pool (NAT) | **12760**
- pool | **12762**
- pool (Access) | **12763**
- pool (Address-Assignment Pools) | **12766**
- pool-default-port-range | **12768**
- pool-default-port-range (Source NAT Next Gen Services) | **12770**
- pool-default-twin-port-range | **12771**
- pool (Destination NAT Next Gen Services) | **12773**
- pool (DHCP Local Server Overrides) | **12774**
- pool-match-order | **12777**
- pool (NAT Rule Next Gen Services) | **12778**
- pool (Security Destination NAT) | **12779**
- pool (Security Source NAT) | **12781**
- pool (Security Source NAT Rule Set) | **12783**
- pool (Service Interface) | **12784**
- pool (Service Interfaces) | **12786**
- pool (Source NAT Next Gen Services) | **12787**
- pool (System) | **12789**
- pool (TDF Domain) | **12791**
- pool-utilization-alarm | **12793**
- pool-utilization-alarm (Security Source NAT Pool) | **12794**
- pool-utilization-alarm (Source NAT Next Gen Services) | **12796**
- pop (protocols mpls) | **12797**

- pop | **12799**
- pop-all-labels | **12800**
- pop-and-forward (Protocols MPLS) | **12802**
- pop-and-forward (Protocols RSVP) | **12803**
- pop (Dynamic VLANs) | **12804**
- pop-pop | **12806**
- pop-swap | **12807**
- pop3-profile (Security Content Security Policy Antivirus) | **12808**
- pop3-profile (Security Content Security Policy Content Filtering) | **12809**
- port | **12811**
- port (Access LDAP) | **12812**
- port (cflowd) | **12813**
- port (Chassis) | **12815**
- port (Chassis Synchronization) | **12818**
- port (Diameter Peer) | **12819**
- port-error-disable | **12821**
- port (Flow Monitoring) | **12823**
- port-forwarding | **12824**
- port-forwarding (Destination NAT Next Gen Services) | **12826**
- port-forwarding-mappings | **12827**
- port-forwarding-mappings (Destination NAT Rule Next Gen Services) | **12828**
- port-group-alias (satellite-policies) | **12829**
- port-id | **12831**
- port-id (MACsec for MX Series) | **12833**
- port (LRF Profile) | **12834**
- port-mirroring | **12836**
- port-mode | **12841**
- port-mode (Fibre Channel Interfaces) | **12843**
- port (NETCONF) | **12844**
- port (Origin Validation for BGP) | **12846**
- port-overloading-factor (Security Source NAT Interface) | **12847**
- port-overloading-factor (Security Source NAT Pool) | **12849**

- port-overloading (Security Source NAT Interface) | **12850**
- port (OVSDB) | **12852**
- port (Packet Forwarding) | **12853**
- port-priority | **12856**
- port (Protocols OpenFlow) | **12857**
- port (RADIUS Server) | **12859**
- port-randomization | **12860**
- port-range | **12861**
- port-range (Address Mapping) | **12863**
- port-range (Application Identification) | **12865**
- port (REST API) | **12866**
- port-round-robin | **12867**
- port-round-robin (Source NAT Next Gen Services) | **12869**
- port (RPM) | **12870**
- port (satellite-policies port-group-alias) | **12871**
- port-scaling-enlargement | **12873**
- port-scan | **12874**
- port (Security Antivirus) | **12875**
- port (Security Source NAT) | **12877**
- port (Security Web Filtering Server) | **12879**
- port (Services) | **12880**
- port (Services NAT) | **12881**
- port (Services Voice) | **12884**
- port (Source NAT Next Gen Services) | **12885**
- port (SRC Server) | **12887**
- port-status-tlv | **12888**
- port (Subscriber Limits) | **12890**
- port (Syslog) | **12891**
- port (System Log Messsages) | **12893**
- port (TWAMP) | **12894**
- ports (System) | **12896**
- ports-per-session | **12897**

- post-attack | **12898**
- post-attack-timeout | **12900**
- post-cli-implicit-firewall | **12901**
- post-convergence-lfa | **12903**
- post-convergence-lfa (Protocols OSPF) | **12904**
- post-service-filter | **12907**
- post-service-filter (Dynamic Service Sets) | **12908**
- post-policy | **12910**
- potential-violation | **12912**
- power | **12915**
- power | **12918**
- power-budget-priority | **12919**
- power-mode-disable | **12921**
- power-mode-ipsec | **12922**
- power-mode-ipsec-qat | **12924**
- pp0 (Dynamic PPPoE) | **12925**
- ppm | **12927**
- ppm (Ethernet Switching) | **12930**
- ppp | **12931**
- ppp-access-profile | **12944**
- ppp (Group Profile) | **12945**
- ppp-options | **12949**
- ppp-options (Dynamic PPP) | **12951**
- ppp-options (L2TP) | **12954**
- ppp-over-ether | **12956**
- ppp (Profile) | **12957**
- ppp-subscriber-services | **12959**
- pppoe (System Processes) | **12960**
- pppoe (DDoS Protection) | **12962**
- pppoe-options (SRX Series) | **12972**
- pppoe-options | **12974**
- pppoe-options (Dynamic PPPoE) | **12976**

- pppoe-tags (Adjustment Control Profiles) | **12977**
- pppoe-underlying-options (Dynamic VLAN Interface Sets) | **12979**
- pppoe-underlying-options (Static and Dynamic Subscribers) | **12980**
- pps-limit (Hierarchical Policer) | **12982**
- pps-limit (Policer) | **12984**
- pptp | **12986**
- prbs | **12987**
- prbs-test-start | **12989**
- prbs-test-stop | **12991**
- pre-attack | **12993**
- pre-filter-shellcode | **12994**
- pre-id-default-policy | **12995**
- pre-ietf-mode | **12998**
- pre-rewrite-tos | **13000**
- pre-shared-key | **13001**
- pre-shared-key (MX Series) | **13003**
- pre-shared-key (Security) | **13004**
- pre-shared-key (Services IKE) | **13006**
- preauthentication-order (Access Profile) | **13007**
- precedence | **13008**
- precision-timers | **13010**
- precision-timers-max-period | **13012**
- predefined-attack-groups | **13013**
- predefined-attacks | **13015**
- predefined-variable-defaults (Dynamic Client Profiles) | **13016**
- preempt | **13019**
- preempt (Chassis Cluster) | **13021**
- preempt-cutover-timer | **13023**
- preempt (VRRP) | **13025**
- preemption | **13026**
- preemptive-fast-reroute | **13028**
- prefer-framed-ip-address (RADIUS Clients) | **13030**

- prefer-framed-ipv6-prefix (RADIUS Clients) | **13031**
- preference (Access) | **13032**
- preference (Interface-Level Preference for VPLS Multihoming for FEC 129) | **13033**
- preference (IPv6 Router Advertisement) | **13035**
- preference (PEG DF Election) | **13037**
- preference (Protocols BGP) | **13038**
- preference (Protocols LDP) | **13040**
- preference (Protocols MPLS) | **13042**
- preference (Protocols RIP) | **13043**
- preference (Protocols RIPng) | **13045**
- preference (Routing Options) | **13046**
- preference (Site-Level Preference for VPLS Multihoming for FEC 129) | **13049**
- preference (Subscriber Management) | **13050**
- preference (Tunnel Profile) | **13052**
- preferred | **13053**
- preferred-ciphers | **13055**
- preferred-lifetime | **13056**
- preferred-lifetime (Dynamic Router Advertisement) | **13058**
- preferred-prefix-length | **13059**
- preferred-source-address | **13060**
- prefix | **13062**
- prefix-65-127-disable | **13064**
- prefix (Access IPv6) | **13067**
- prefix-action (Configuring) | **13068**
- prefix-action (Firewall Filter Action) | **13070**
- prefix (Address-Assignment Pools) | **13071**
- prefix-attribute-flags | **13072**
- prefix (Circuit ID for Option 82) | **13074**
- prefix (DHCP Client) | **13076**
- prefix (DHCP Relay Agent) | **13077**
- prefix (DHCPv6 Options) | **13079**
- prefix (Dynamic Router Advertisement) | **13081**

- prefix-export-limit (Protocols IS-IS) | **13083**
- prefix-export-limit (Protocols OSPF) | **13085**
- prefix-limit | **13087**
- prefix-list | **13090**
- prefix-list-filter | **13092**
- prefix-list (Junos Telemetry Interface) | **13093**
- prefix-list-name | **13095**
- prefix-list (PIM RPF Selection) | **13097**
- prefix-name (Security Static NAT) | **13098**
- prefix (Protocols IPv6 Neighbor Discovery) | **13100**
- prefix (Remote ID for Option 82) | **13101**
- prefix (Security Static NAT) | **13103**
- prefix-segment | **13105**
- prefix-segment-range | **13107**
- prefix-segment (Routing Options) | **13109**
- premium (Hierarchical Policer) | **13110**
- premium (Output Priority Map) | **13112**
- premium (Policer) | **13113**
- preprovisioned | **13115**
- preprovisioned (MX Series Virtual Chassis) | **13116**
- preserve-fpc-poweron-sequence | **13118**
- preserve-incoming-fragment-size | **13119**
- preserve-interface | **13121**
- preserve-nexthop-hierarchy (SRv6 TE) | **13122**
- preserve-parity (Source NAT Next Gen Services) | **13124**
- preserve-range (Source NAT Next Gen Services) | **13125**
- primary | **13126**
- primary | **13127**
- primary (Adaptive Services Interfaces) | **13129**
- primary (Address on Interface) | **13130**
- primary connection (Identity Management Advanced Query) | **13131**
- primary-data-record-fields | **13134**

- primary-dns | **13137**
- primary-interface (Aggregated Inline Services) | **13138**
- primary (Interface for Router) | **13140**
- primary (Link Services IQ PIC Interfaces) | **13142**
- primary (MBGP MVPN) | **13143**
- primary (Protocols MPLS) | **13144**
- primary-server | **13146**
- primary (Virtual Tunnel in Routing Instances) | **13148**
- primary-vlan | **13150**
- primary (VPLS Multihoming) | **13151**
- primary-wins | **13153**
- priority | **13154**
- priority (Application Identification With Next Gen Services) | **13156**
- priority (Bootstrap) | **13157**
- priority (CFM MEP) | **13159**
- priority (Chassis Synchronization Source) | **13160**
- priority (Clock Synchronization) | **13162**
- priority-cost (VRRP) | **13163**
- priority (DDoS) | **13165**
- priority (Diameter Network Element) | **13166**
- priority (Diameter Peer) | **13168**
- priority (Dynamic Schedulers) | **13169**
- priority (Fabric Priority) | **13171**
- priority (Fabric Queues, Schedulers) | **13172**
- priority (FIP) | **13174**
- priority-flow-control | **13176**
- priority-hold-time | **13177**
- priority-hold-time | **13179**
- priority (Management Traffic) | **13180**
- priority (OAM Connectivity-Fault Management) | **13182**
- priority-override (Event Policy) | **13183**
- priority (PIM Interfaces) | **13184**

- priority (PIM RPs) | **13186**
- priority (PoE) | **13188**
- priority (policy-options) | **13189**
- priority (Power over Ethernet) | **13191**
- priority (Protocols MPLS) | **13193**
- priority (Protocols OAM) | **13194**
- priority (Protocols Router Discovery) | **13196**
- priority (Protocols RSVP) | **13197**
- priority (Protocols STP) | **13199**
- priority (Protocols VRRP) | **13201**
- priority (RADIUS Network Elements) | **13203**
- priority (Redundant Power System) | **13204**
- priority (Schedulers) | **13206**
- priority1 | **13209**
- priority2 | **13211**
- privacy-3des | **13212**
- privacy-aes128 | **13213**
- privacy-des | **13215**
- privacy-key | **13216**
- privacy-none | **13217**
- privacy-password | **13218**
- private-vlan | **13220**
- proactive-arp-detection | **13221**
- probe (ICMP Router Discovery) | **13222**
- probe | **13228**
- probe-count | **13231**
- probe-interval | **13233**
- probe-limit | **13234**
- probe-server | **13236**
- probe-type | **13238**
- process-inform | **13240**
- process (Memory) | **13242**

- process-monitor | **13245**
- process-non-null-as-null-register | **13246**
- processes | **13248**
- product-model (Junos Fusion) | **13250**
- product-name | **13252**
- products | **13253**
- profile (Application Identification) | **13254**
- profile | **13256**
- profile | **13257**
- profile (Access) | **13259**
- profile (Access) | **13263**
- profile (Application Firewall) | **13271**
- profile (Captive Portal Content Delivery) | **13273**
- profile(dynamic-application) | **13275**
- profile (Explicit Proxy) | **13277**
- profile (HTTP Header Enrichment) | **13282**
- profile (icap-redirect) | **13283**
- profile (Juniper Secure Connect) | **13286**
- profile (LRF) | **13289**
- profile (Rule Sets) | **13291**
- profile (security) | **13292**
- profile (Security Antispam SBL) | **13310**
- profile (Security Antivirus Juniper Express Engine) | **13312**
- profile (Security Antivirus Kaspersky Lab Engine) | **13314**
- profile (Security Content Filtering) | **13317**
- profile (Security GTP) | **13318**
- profile (Security SCTP) | **13323**
- profile (Security Sophos Engine Antivirus) | **13326**
- profile (Security Web Filtering Juniper Enhanced) | **13328**
- profile (Security Web Filtering Juniper Local) | **13330**
- profile (Security Web Filtering Surf Control Integrated) | **13332**
- profile (Security Web Filtering Websense Redirect) | **13333**

- profile (Services Application Identification) | **13335**
- profile (Services PCEF) | **13336**
- profile (Services Proxy) | **13338**
- profile (Services SSL Proxy) | **13339**
- profile (SSL Initiation) | **13343**
- profile (SSL Termination) | **13347**
- profile (Traffic Load Balancer) | **13349**
- profile-type (Dynamic Service Profiles) | **13353**
- profile-variable-set (Dynamic Profiles) | **13354**
- profile-variable-set (Routing Instances) | **13356**
- profile (Web Filter) | **13357**
- profilerd | **13360**
- profiles | **13362**
- profiles (AAA) | **13363**
- profiles (PCEF) | **13365**
- profiles (RFC 2544 Benchmarking) | **13368**
- programmable-rpd | **13369**
- promiscuous | **13372**
- promiscuous-mode (Dynamic IGMP Interface) | **13373**
- promiscuous-mode (Interfaces) | **13374**
- promiscuous-mode (Protocols IGMP) | **13375**
- promote | **13377**
- propagate-settings | **13379**
- propagate-tc | **13380**
- propagation-delay | **13381**
- proposal (Security Group VPN Member IKE) | **13383**
- proposal (Security Group VPN Server IKE) | **13385**
- proposal (Security Group VPN Server IPsec) | **13387**
- proposal (Security IKE) | **13389**
- proposal (Security IKE Dynamic SA) | **13393**
- proposal (Security IPsec) | **13394**
- proposal (Services IKE) | **13399**

- proposal (Services IPsec VPN) | **13400**
- proposals (Services) | **13402**
- proposals | **13403**
- protect | **13404**
- protect (Security) | **13406**
- protect core | **13407**
- protect-interface | **13409**
- protect-maintenance-association (OAM) | **13410**
- protected-l2circuit | **13412**
- protection-group | **13413**
- protection (Protocols BGP) | **13416**
- protection (Protocols MPLS) | **13417**
- protection-revert-time | **13419**
- protector-interface | **13420**
- protector-pe | **13421**
- protocol (Firewall Filter) | **13423**
- protocol | **13424**
- protocol (early-aging) | **13428**
- protocol (Application Identification) | **13429**
- protocol (Applications) | **13431**
- protocol-binding | **13433**
- protocol-command | **13434**
- protocol-down | **13436**
- protocol (Drop Profile Map) | **13437**
- protocol (Drop Profiles) | **13438**
- protocol (Dynamic Schedulers) | **13439**
- protocol (Flow Descriptions) | **13441**
- protocol (Host Outbound Traffic) | **13443**
- protocol (IPsec) | **13444**
- protocol (Junos-FIPS Software) | **13445**
- protocol (Junos OS) | **13446**
- protocol-master | **13448**

- protocol-name | **13451**
- protocol-origin | **13452**
- protocol (OVSDB) | **13454**
- protocol (Protocols OpenFlow) | **13455**
- protocol (Rewrite Rules) | **13457**
- protocol (Schedulers) | **13459**
- protocol (Security Destination NAT) | **13460**
- protocol (Security Forwarding Options) | **13461**
- protocol (Security IDP Signature Attack) | **13463**
- protocol (Security Source NAT) | **13469**
- protocol (Server MEP) | **13471**
- protocol (Services Proxy) | **13472**
- protocol (Subscriber Secure Policy) | **13474**
- protocol-version | **13475**
- protocols | **13477**
- protocols | **13748**
- protocols (DDoS) | **13751**
- protocols (DDoS) (ACX Series, PTX Series, and QFX Series) | **13765**
- protocols (Dynamic Profiles) | **13791**
- protocols (Fabric) | **13794**
- protocols (FIP) | **13795**
- protocols rsvp expand-flood-reflector-hop | **13797**
- protocols (Security Zones Host Inbound Traffic) | **13798**
- protocols (Security Zones Interfaces) | **13800**
- protocols (STP Type) | **13802**
- protocols (Tenant Systems) | **13804**
- provider-tunnel | **13806**
- providers | **13812**
- provisioning-method (Remote Device Management) | **13814**
- provisioning-order (Diameter Applications) | **13815**
- proxy (interfaces) | **13817**
- proxy (Bridge Domains) | **13819**

- proxy-arp | **13820**
- proxy-arp (Dynamic Profiles) | **13822**
- proxy-arp (Security NAT) | **13823**
- proxy (Fibre Channel) | **13825**
- proxy-generate | **13826**
- proxy (Interfaces) | **13828**
- proxy-mac | **13829**
- proxy-macip-advertisement | **13831**
- proxy-mode | **13833**
- proxy (Multicast VLAN Registration) | **13834**
- proxy-ndp (Security NAT) | **13836**
- proxy-profile | **13838**
- proxy (Security Antivirus) | **13839**
- proxy (Services) | **13841**
- proxy (snmp) | **13843**
- proxy (System) | **13845**
- ps0 (Pseudowire Subscriber Interfaces) | **13847**
- pseudowire-service (Pseudowire Subscriber Interfaces) | **13848**
- pseudowire-status-tlv | **13849**
- psn-tunnel-endpoint | **13851**
- psu | **13852**
- psu (satellite-management) | **13853**
- ptp | **13855**
- ptp-mode | **13865**
- ptsp-rules | **13866**
- pulse-per-second-enable | **13868**
- purge-flow-timer (Protocols OpenFlow) | **13869**
- purge-timeout (programmable-rpd) | **13870**
- purging | **13872**
- push (protocols mpls) | **13873**
- push | **13875**
- push-backup-to-master (Accounting Options) | **13876**

- push (Dynamic VLANs) | 13878
- push-push | 13879
- push-to-identity-management | 13880
- pvlan | 13882
- pvlan-trunk | 13883
- python-script-user (Event and SNMP Scripts) | 13884

p2mp-ldp-next-hop

IN THIS SECTION

- Syntax | 12119
- Hierarchy Level | 12119
- Description | 12120
- Options | 12120
- Required Privilege Level | 12120
- Release Information | 12120

Syntax

```
p2mp-ldp-next-hop {
root-address root-address{
    lsp-id id;
        }
    }
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options
static route destination-prefix],
```

```
[edit logical-systems logical-system-name routing-options static route destination-prefix],
[edit routing-instances routing-instance-name routing-options static route destination-prefix].
[edit routing-options static route destination-prefix]
```

Description

Specify a point-to-multipoint LDP label-switched path (LSP) as the next hop for a static route, and configure a root and provide an lsp-id on that LDP-signalled label-switched path.

Options

root-address *root address* Specify the root address of the point-to-multipoint LSP.

lsp-id *id* Specify the generic LSP identifier. The range is 1 through 65535.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

p2mp-lsp-next-hop

IN THIS SECTION

- Syntax | 12121
- Hierarchy Level | 12121
- Description | 12121
- Required Privilege Level | 12121
- Release Information | 12121

Syntax

```
p2mp-lsp-next-hop {
    metric metric;
    preference preference;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options
static route destination-prefix],
[edit logical-systems logical-system-name routing-options static route destination-prefix],
[edit routing-instances routing-instance-name routing-options static route destination-prefix].
[edit routing-options static route destination-prefix]
```

Description

Specify a point-to-multipoint LSP as the next hop for a static route, and configure an independent metric or preference on that next-hop LSP.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Static Unicast Routes for Point-to-Multipoint LSPs

Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP

Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems

p2mp (Protocols LDP)

IN THIS SECTION

- Syntax | 12122
- Hierarchy Level | 12122
- Description | 12122
- Options | 12123
- Required Privilege Level | 12123
- Release Information | 12123

Syntax

```
p2mp {
    recursive {
        fec;
        route;
        }
    }
    root-address root-address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Enable point-to-multipoint MPLS LSPs in an LDP-signaled LSP.

Options

no-rsvp-tunneling

(Optional) Disable LDP point-to-multipoint LSPs from using RSVP-TE LSPs for tunneling, and use LDP paths instead.

> **NOTE**: The `no-rsvp-tunneling` option is introduced in Junos OS Release 16.1R5, 17.3R1, 17.2R2, 16.2R3, and later releases.

Starting in Junos OS Release 12.3R1, Junos OS provides support for Multipoint LDP (M-LDP) for Targeted LDP (T-LDP) sessions with unicast replication, in addition to link sessions. As a result, the default behavior of M-LDP over RSVP tunneling is similar to unicast LDP. However, because T-LDP is chosen over LDP and link sessions to signal point-to-multipoint LSPs, the `no-rsvp-tunelling` option enables LDP natively throughout the network.

recursive

(Optional) Configure point-to-multipoint recursive parameters, including `route`.

route

(Optional) Allow recursive route resolution to signal point-to-multipoint Forwarding Equivalence Class (FEC).

fec

(Optional) Allow recursive opaque value to signal P2MP FEC.

root-address ***root-address***

(Optional) Specify the root address of the point-to-multipoint LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

`no-rsvp-tunneling` option added in Junos OS Release 16.1R5.

`fec` option added in Junos OS and Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs

Point-to-Multipoint LSPs Overview

p2mp (Protocols MPLS)

IN THIS SECTION

- Syntax | 12124
- Hierarchy Level | 12124
- Description | 12124
- Options | 12125
- Required Privilege Level | 12125
- Release Information | 12125

Syntax

```
p2mp p2mp-lsp-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Specify an LSP as either a point-to-multipoint LSP or as a branch LSP of a point-to-multipoint LSP by specifying the point-to-multipoint LSP path name.

Options

p2mp-lsp-name—Name of the point-to-multipoint LSP path that identifies the sequence of nodes that form the point-to-multipoint LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Primary Point-to-Multipoint LSP

p2mp-receive-switch

IN THIS SECTION

- Syntax | 12126
- Hierarchy Level | 12126
- Description | 12126
- Options | 12126
- Required Privilege Level | 12126
- Release Information | 12126

Syntax

```
p2mp-receive-switch point-to-multipoint-switch-name {
    output-interface [ interface-name.unit-number ];
    receive-p2mp-lsp receiving-point-to-multipoint-lsp;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols connections],
[edit protocols connections]
```

Description

Configure the CCC switch for a point-to-multipoint LSP on the egress PE router.

Options

point-to-multipoint-switch-name—Point-to-multipoint CCC receive switch name.

`output-interface` *interface-name.unit-number*—Name of the egress interfaces for the point-to-multipoint LSP traffic. You can configure multiple output interfaces.

`receive-p2mp-lsp` *receiving-point-to-multipoint-lsp*—Name of the point-to-multipoint LSP that is switched to the output interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers*

p2mp-transmit-switch

IN THIS SECTION

- Syntax | 12127
- Hierarchy Level | 12127
- Description | 12127
- Options | 12127
- Required Privilege Level | 12128
- Release Information | 12128

Syntax

```
p2mp-transmit-switch point-to-multipoint-transmit-switch-name {
    input-interface interface-name.unit-number;
    transmit-p2mp-lsp transmitting-point-to-multipoint-lsp;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols connections],
[edit protocols connections]
```

Description

Configure the CCC switch for the point-to-multipoint LSP on the ingress PE router.

Options

point-to-multipoint-transmit-switch-name—Point-to-multipoint CCC transmit switch name.

`input-interface` *input-interface-name.unit-number*—Specify the name of the interface carrying incoming traffic to be switched to the point-to-multipoint LSP.

`transmit-p2mp-lsp` *`transmitting-point-to-multipoint-lsp`*—Specify the name of the point-to-multipoint LSP carrying traffic to the CCC switch on the egress PE router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers

package (Loading on PIC)

IN THIS SECTION

- Syntax | 12128
- Hierarchy Level | 12129
- Description | 12129
- Options | 12129
- Required Privilege Level | 12129
- Release Information | 12129

Syntax

```
package package-name;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-
provider]
```

Description

Identify a package to be loaded on the PIC. When a package is added or removed, the PIC reboots.

Options

package-name—Name of the package to be loaded on the PIC. There can be up to eight packages loaded on a PIC; however, only one data package is allowed per PIC. An error message is displayed if more than eight packages are specified.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

packet-action

IN THIS SECTION

- Syntax | 12130
- Hierarchy Level | 12130
- Description | 12131
- Default | 12131
- Options | 12131
- Required Privilege Level | 12132

Release Information | 12132

Syntax

```
packet-action action;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name bridge-options interface interface-name interface-mac-limit
limit],
[edit bridge-domains bridge-domain-name bridge-options interface-mac-limit limit],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name bridge-options
interface interface-name interface-mac-limit limit],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name bridge-options interface-
mac-limit limit],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options interface interface-name interface-mac-limit limit],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options interface-mac-limit limit],
[edit logical-systems logical-system-name routing-instances routing-instance-name switch-options
interface interface-name interface-mac-limit limit],
[edit logical-systems logical-system-name routing-instances routing-instance-name switch-options
interface-mac-limit limit],
[edit logical-systems logical-system-name switch-options interface-mac-limit limit],
[edit protocols l2-learning global-mac-limit limit],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options
interface interface-name interface-mac-limit limit],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options
interface-mac-limit limit],
[edit routing-instances routing-instance-name protocols evpn interface-mac-limit (VPLS)],
[edit routing-instances routing-instance-name protocols evpn interface interface-name interface-
mac-limit (VPLS)],
[edit routing-instances routing-instance-name protocols evpn mac-table-size limit],
[edit routing-instances routing-instance-name switch-options interface interface-name interface-mac-
limit limit],
[edit routing-instances routing-instance-name switch-options interface-mac-limit limit],
[edit switch-options interface-mac-limit limit],
```

```
[edit switch-options mac-table-size limit],
[edit switch-options interface interface-name interface-mac-limit limit],
[edit vlans vlan-name switch-options mac-table-size limit][edit vlans vlan-name switch-options
interface-mac-limit limit],
[edit vlans vlan-name switch-options interface interface-name interface-mac-limit limit]
```

Description

Specify the action taken when packets with new source MAC addresses are received after the MAC address limit is reached. If this statement is not configured, packets with new source MAC addresses are forwarded by default.

NOTE: The `packet-action` statement is not supported on the QFX10002-60C switch.

NOTE: In ACX710 and ACX5448 Series devices, the VPLS mac-table-size and interface-mac-limit packet-action drop is not supported.

Default

NOTE: On a QFX Series Virtual Chassis, if you include the `shutdown` option at the `[edit vlans vlan-name switch-options interface interface-name interface-mac-limit packet-action]` hierarchy level and issue the `commit` operation, the system generates a commit error. The system does not generate an error if you include the `shutdown` option at the `[edit switch-options interface interface-name interface-mac-limit packet-action]` hierarchy level.

Disabled. The default is for packets for new source MAC addresses to be forwarded after the MAC address limit is reached.

Options

drop (EX Series switches, QFX Series switches, and PTX Series routers in which Junos OS Evolved is installed) Drop packets with new source MAC addresses, and do not learn the new source MAC addresses.

NOTE: On QFX10000 switches, if you include the drop option, you cannot configure unicast reverse-path forwarding (URFP) on integrated routing

and bridging (IRB) and MAC limiting on the same interface. If you have an MC-LAG configuration, you cannot configure MAC limiting on the interchassis link (ICL) interface.

drop-and-log (EX Series switches, QFX Series switches, and PTX Series routers in which Junos OS Evolved is installed) Drop packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

log (EX Series switches and QFX Series switches) Hold packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

(PTX Series routers in which Junos OS Evolved is installed) Forward packets with the new source MAC addresses, and generate an alarm, an SNMP trap, and a system log entry when the MAC limit exceeds the configured value.

none (EX Series switches, QFX Series switches, and PTX Series routers in which Junos OS Evolved is installed) Forward packets with new source MAC addresses.

shutdown (EX Series switches, QFX Series switches, and PTX Series routers in which Junos OS Evolved is installed) Disable the specified interface, and generate an alarm, an SNMP trap, or a system log entry.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for the `switch-options` statement added in Junos OS Release 9.2.

Support for top-level configuration for the `virtual-switch` type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

[edit switch-options interface *interface-name* interface-mac-limit *limit*], [edit switch-options interface-mac-limit *limit*], [edit switch-options mac-table-size *limit*], [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*], [edit vlans *vlan-name* switch-options interface-mac-limit *limit*],

and [edit vlans *vlan-name* switch-options mac-table-size *limit*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Support for EVPNs introduced in Junos OS Release 13.2 on MX Series 5G Universal Routing Platforms.

Support at the `[edit switch-options interface interface-name interface-mac-limit limit]` hierarchy level and hierarchy levels under `[edit vlans vlan-name]` introduced in Junos OS Release 13.2X50-D10 for EX Series switches and Junos OS Release 13.2 for the QFX Series.

Support for the `none` option introduced in Junos OS Evolved Release 20.4R1 for PTX Series routers.

Support for the `drop`, `drop-and-log`, `log`, and `shutdown` options introduced in Junos OS Evolved Release 22.4R1 for PTX Series routers.

RELATED DOCUMENTATION

- *Configuring EVPN Routing Instances*
- *Configuring EVPN Routing Instances on EX9200 Switches*
- Configuring MAC Limiting (ELS)
- Configuring Persistent MAC Learning (ELS)
- Understanding Layer 2 Learning and Forwarding for Bridge Domains
- Layer 2 Learning and Forwarding for VLANs Overview
- Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports
- Layer 2 Learning and Forwarding for VLANs Overview
- Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port

packet-burst (Hierarchical Policer)

IN THIS SECTION

- Syntax | **12134**
- Hierarchy Level | **12134**
- Description | **12134**
- Options | **12134**
- Required Privilege Level | **12134**

Release Information | 12135

Syntax

```
packet-burst packets;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer hierarchical-policer-name aggregate
if-exceeding-pps],
[edit dynamic-profiles profile-name firewall hierarchical-policer hierarchical-policer-name premium
if-exceeding-pps],
[edit No Link Title hierarchical-policer hierarchical-policer-name aggregate if-exceeding-pps],
[edit No Link Title hierarchical-policer hierarchical-policer-name premium if-exceeding-pps]
```

Description

On MPCs hosted on MX Series routers, configure the packet burst limit for premium or aggregate traffic in a hierarchical policer. When used in combination with the if-exceeding-ppsand "pps-limit" on page 12984 statements, you can control the number of packets that will be allowed over a configured packets-per-second limit when traffic is in burst state.

Options

packets—Packet burst limit in packets. You can specify the number of packets either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), or `m` (1000000).

- **Range:** 1 through 24414062

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Hierarchical Policer Configuration Overview

Policer Color-Marking and Actions

Single Token Bucket Algorithm

Hierarchical Policers

aggregate (Hierarchical Policer) | **2064**

bandwidth-limit (Hierarchical Policer) | **2855**

premium (Hierarchical Policer) | **13110**

packet-burst (Policer)

IN THIS SECTION

- Syntax | **12135**
- Hierarchy Level | **12136**
- Description | **12136**
- Options | **12136**
- Required Privilege Level | **12137**
- Release Information | **12137**

Syntax

```
packet-burst packets;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name if-exceeding-pps],
[edit firewall policer policer-name if-exceeding-pps],
[edit logical-systems logical-system-name firewall policer policer-name if-exceeding-pps]
```

Description

For a single-rate two-color policer, configure the `packet-burst` as a number of packets. Single-rate two-color policing uses the single token bucket algorithm to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the pps-limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with `low` packet loss priority (PLP) and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

NOTE: This statement specifies the packet burst limit as an absolute number of packets.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allows bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options

packets—Specify the number of packets either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), or `m` (1000000).

- **Range:** 1 through 24414062
- **Default:** None

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Two-Color Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Single Token Bucket Algorithm*
- *Determining Proper Burst Size for Traffic Policers*
- bandwidth-percent | **2866**
- burst-size-limit (Policer) | **3086**

packet-capture

IN THIS SECTION

- Syntax | **12138**
- Hierarchy Level | **12138**
- Description | **12138**
- Options | **12138**
- Required Privilege Level | **12140**
- Release Information | **12140**

Syntax

```
packet-capture {
    aggressive-mode;
    buffer-packets-limit bytes;
    capture-interval capture-interval;
    capture-limit capture-limit;
    global;
    max-bytes bytes;
    max-files max-files;
    max-packets max-packets;
    no-inconclusive;
    storage-limit bytes;
}
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Specify packet capture options to capture unknown application traffic.

You can use the packet capture of unknown applications functionality to gather more details about an unknown application on your security device. Once you've configured packet capture options on your security device, the unknown application traffic is gathered and stored on the device in a packet capture file (**.pcap**) at **/var/log/pcap/** location.

Options

aggressive-mode

Capture all traffic before AppID classifies the applications. In this mode, the system captures all application traffic irrespective of the application system cache (ASC) entry. Packet capture starts for the first packet of the first session.

buffer-packets-limit

Maximum memory to buffer packets (bytes). Use this option to limit the memory available in the Packet Forwarding Engine for packet capture functionality.

- **Default:** 1% of available data in shared memory

- **Range:** 40 bytes to 5% of available data in shared memory
- **Default:** 1 MB (for cSRX)
- **Range:** 40 bytes through 5 MB

capture-interval

Timeout value in minutes to avoid repetitive capture of the same traffic. After this interval, the system continues to capture newer packet details for unknown applications until the capture limit is reached.

- **Default:** 1440 minutes (24 hours).
- **Range:** 1 through 525,600 seconds

capture-limit

Number of repetitive captures of the same traffic. Use this option to limit the number of times the same traffic can be repeatedly captured before the cache entry times out.

- **Default:** 5
- **Range:** 1 through 1000

global

Enable packet capture globally to capture all unknown application traffic. Another option is to enable capturing of unknown application traffic specific to a security policy.

max-bytes

Maximum number of TCP bytes per session (bytes). For TCP sessions, the count includes the actual payload data length and excludes IP/TCP headers for the maximum bytes limit.

If you are setting the packet capture at the security policy level, the packet capture concludes only after the final policy is applied even if the configured limit is reached.

Limitation—Jumbo frames can have up to 1500 bytes of the payload saved in the capture file.

- **Default:** 6000 bytes
- **Range:** 40 through 1,073,741,824

max-files

Maximum number of unique packet capture files to create before the oldest file is overwritten by a new file created.

- **Default:** 100

  (Previously 25)

- **Range:** 1 through 2500

max-packets Maximum number of UDP packets per session.

- **Default:** 10 packets
- **Range:** 1 through 1000

no-inconclusive Disable packet capturing of inconclusive traffic. This option disables the packet capture for the following sessions:

- Sessions that are closed before the application identification or classification completes.
- Sessions that are not getting classified even whn they reach the maximum packet capture limit.

If you do not configure this option, by default, the system captures packets for inconclusive sessions.

storage-limit Maximum disk space (bytes) that can be used in the Routing Engine for packet capture files.

- **Default:** 50 MB
- **Range:** 1,048,576 through 4,294,967,295 bytes

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Configure Packet Capture For Unknown Application Traffic

show services application-identification packet-capture counters | **25443**

packet-filter

IN THIS SECTION

- Syntax | 12141
- Hierarchy Level | 12141
- Description | 12141
- Options | 12142
- Required Privilege Level | 12142
- Release Information | 12142

Syntax

```
packet-filter packet-filter-name {
    action-profile profile-name {
    destination-port (Security Forwarding Options) (port-range | protocol-name);
    destination-prefix destination-prefix;
    interface logical-interface-name;
    protocol (Security Forwarding Options) (protocol-number | protocol-name;
    source-port (Security Forwarding Options) (port-range | protocol- name);
    source-prefix source-prefix;
    }
```

Hierarchy Level

```
[edit security datapath-debug]
```

Description

Set packet filter for taking the datapath-debug action. A filter is defined to filter traffic, then an action profile is applied to the filtered traffic. Be sure to configure multiple packet filters to capture the traffic. One packet filter only captures the traffic as specified in it, such as from one source to one destination. The same packet filter will not capture the traffic in the reverse direction. You need to configure another packet filter to capture the traffic in reverse direction and specify the source and destination according

to the response packet in it. The action profile specifies a variety of actions on the processing unit. A maximum of four filters are supported at the same time. Packet filters can be configured with source and destination prefix and port (including ranges), and protocol.

Action-profile settings have no specific minimum setting, it is based on trace, count, packet summary and packet-dump. Enabling end-to-end debugging without or with a very broad filter is not recommended. This could result in a high PFE CPU usage. Therefore when selecting what to capture through a filter care must be taken. List as many and specific criteria which then results in the minimum amount of traffic to be captured.

NOTE: Packet filter is supported on SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800 devices.

Options

- `action-profile` *`profile-name`*—Identify the action profile to use. You can specify the name of the action profile to use. Using the request security action-profile command, you can set the action for the packet match for a specified filter. Action-profile must be defined.
- `destination-port (`*`port-range`* `|` *`protocol name`*`)`—Specify a destination port to match TCP/UDP destination port.
- `destination-prefix` *`destination-prefix`*—Specify a destination IPv4/IPv6 address prefix.
- `interface` *`logical-interface-name`*—Specify a logical interface name.
- `protocol (`*`protocol-number`* `|` *`protocol-name`*—Match IP protocol type.
- `source-port (`*`port-range`* `|` *`protocol-name`*—Match TCP/UDP source port.
- `source-prefix` *`source-prefix`*—Specify a source IP address prefix.

Required Privilege Level

security—To view this in the configuration

security-control—To add this to the configuration.

Release Information

Command introduced in Junos OS Release 9.6 ; Support for IPv6 addresses for the `destination-prefix` and `source-prefix` options added in Junos OS Release 10.4.

packet-format-match

IN THIS SECTION

- Syntax | 12143
- Hierarchy Level | 12143
- Description | 12143
- Default | 12144
- Options | 12144
- Required Privilege Level | 12144
- Release Information | 12144

Syntax

```
packet-format-match {
    (mpls-packet-format-1 | mpls-packet-format-2 | mpls-packet-format-3 | mpls-packet-format-4 |
mpls-packet-format-5 | mpls-packet-format-6 | mpls-packet-format-7 | mpls-packet-format-8);
}
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

Types of packet formats for MPLS labeled traffic.

NOTE: You must restart the PFE every time you configure this command. A PFE reboot forces disruptions to the data plane, a hit that can last several minutes while the component reboots, and is then repopulated with the current routing state.

Default

The default packet format is `mpls-packet-format-1`.

Options

You can specify a value from `1` through `8`.

mpls-packet-format-1	Untagged packet with one or two labels.
mpls-packet-format-2	Untagged packet with one label.
mpls-packet-format-3	Untagged packet with two labels.
mpls-packet-format-4	Single tagged packet with one or two labels.
mpls-packet-format-5	Single tagged packet with one label.
mpls-packet-format-6	Single tagged packet with two labels
mpls-packet-format-7	Untagged and single tagged packets with one label.
mpls-packet-format-8	Untagged and single tagged packets with two labels.

Required Privilege Level

admin

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

family mpls | 6886

Overview of MPLS Firewall Filters on Loopback Interface

Configuring MPLS Firewall Filters and Policers on Switches

packet-forwarding-options

IN THIS SECTION

- Syntax | 12145
- Hierarchy Level | 12147
- Description | 12147
- Options | 12147
- Required Privilege Level | 12150
- Release Information | 12150

Syntax

```
packet-forwarding-options {
      eracl-ip6-match {
          (srcip6-and-destip6 | srcip6-only);
      }
       firewall {
        eracl-profile {
         eracl-scale;
        }
        profiles {
         ethernet-switching {
          (profile1 | profile2);
         }
         inet {
          (profile1 | profile2);
         }
        }
       }
      firewall-profile {
          (default-profile | evpn-mh-profile | mef-profile);
          (profile-one | profile-two);
          egress-l2-extended-match;
          filter-modify {
               (full-scale | make-before-break);
```

```
            }
        }
        forwarding-profile (default-profile | host-acl-profile | host-profile | lpm-profile);
        hw-db-profile {
          (balanced | balanced-exem | l2-xl | l3-xl);
          lpm-distribution (1 | 2 | 200 | 3 | 4);
        }
        internal-forwarding-bandwidth internal-forwarding-bandwidth-value {
          reserve ingress-egress-mcast percentage;
        }
        ipfix-sw-mode;
        l2-profile {
          fast-flush-profile;
        }
        lag-profile {
          (default-profile | lag-mode-high | lag-mode-medium);
        }
        mcast-repl-type {
           ingress-egress-recycle;
           ingress-only;
        }
        mef-forwarding-profile;
        mpls {
         lsp-ingress-stats-enable;
         lsp-transit-stats-enable;
         }
        multicast-profile {
           strict-ttl-check-profile;
        }
        multicast-statistics;
        no-ip-tos-rewrite;
        overlay {
            interface-num interface-num;
            next-hop next-hop;
        }
        packet-format-match (mpls-packet-format-2 | mpls-packet-format-3 | mpls-packet-format-4 |
mpls-packet-format-5 | mpls-packet-format-6 | mpls-packet-format-7 | mpls-packet-format-8);
        recycle-bandwidth-profiles {
           name;
           evpn-irb evpn-irb;
           fti-virtual-loop fti-virtual-loop;
           ingress-egress-mcast ingress-egress-mcast;
           srv6 srv6;
```

```
        srv6-ext srv6-ext;
        tunnel-services tunnel-services;
        vrf-leak vrf-leak;
      }
      system-profile {
       (default-profile | vxlan-extended | vxlan-stitching | vpls-meshgroups);
      }
      warm-boot;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure packet forwarding options. The packet forwarding option specifies that data is forwarded to the specified forwarding destination.

Options

eracl-ip6-match IPv6 match for ERACL, source or destination. The default is destination only.

- `srcip6-and-destip6`—Match on both the source and destination.
- `srcip6-only`—Match on source only.

firewall-profile Set the firewall profile.

- `default-profile`—Set the profile to support default services.
- `evpn-mh-profile`—Set the profile to support EVPN-MH.
- `mef-profile`—Set the profile to support MEF services.
- `profile-one`—Selects profile one feature set.
- `profile-two`—Selects profile two feature set.

forwarding-profile (default-profile | host-acl-profile | Configure the profile for the unified forwarding table.

host-profile | lpm-profile)

- `default-profile`—MAC: 32K, L3-host: 32K, LPM: 720K, FP-Compression: 18K. Restarts PFE.
- `host-acl-profile`—MAC: 160K, L3-host: 160K, LPM: 65K, FP-Compression: 18K. Restarts PFE.
- `host-profile`—MAC: 160K, L3-host: 160K, LPM: 72K, FP-Compression: 0. Restarts PFE.
- `lpm-profile`—MAC: 32K, L3-host: 32K, LPM: 1.24M, ARP: 61K, FP-Compression: 0, Tunnels: 0. Restarts PFE.

hw-db-profile

Set the hardware database profile.

- `(balanced | balanced-exem | l2-xl | l3-xl)`—Select the profile type. Restarts PFE.
- (Mandatory) `lpm-distribution (1 | 2 | 200 | 3 | 4)`—Specify route distribution between public and private (VRF/VPN) routes in longest prefix match (LPM). Default is 1.
  - `1`—Set the LPM distribution to 1
  - `2`—Set the LPM distribution to 2. Not valid for the `balanced` profile.
  - `200`—Set the LPM distribution to 200. Valid only for the `balanced` profile.
  - `3`—Set the LPM distribution to 3. Not valid for the `balanced` profile.
  - `4`—Set the LPM distribution to 4. Valid only for the `l2-xl` profile.

internal-forwarding-bandwidth

Set the internal forwarding bandwidth.

ipfix-sw-mode

(QFX5120 only) Enable software-based IPFIX processing, used to enable flow-based telemetry for VXLANs. Must reboot for this option to take effect.

l2-profile

Set the L2 profile.

- `fast-flush-profile`—Set the profile to support fast flush.

lag-profile

Set link aggregation group (LAG) configuration mode.

- `default-profile`—Allow to create 256 LAGs and maximum 64 members per LAG.
- `lag-mode-high`—Allow to create 1024 LAGs and maximum 16 members per LAG.
- `lag-mode-medium`—Allow to create 512 LAGs and maximum 32 members per LAG.

mcast-repl-type Set multicast replication type. Restarts PFE.

mef-forwarding-profile Enable MEF Forwarding Profile, restart PFE

multicast-profile Set the multicast profile.

- `strict-ttl-check-profile`—Set the profile to support strict TTL check for multicast traffic.

multicast-statistics Enable multicast statistics support.

no-ip-tos-rewrite (ACX7100 only) Disable L3 code point rewrite at system level, restart PFE.

packet-format-match Packet format match conditions based on VLAN tags and MPLS labels. The default match is an untagged packet with one or two labels.

- `mpls-packet-format-2`—Untagged packet with one label.
- `mpls-packet-format-3`—Untagged packet with two labels.
- `mpls-packet-format-4`—Single tagged packet with one or two labels.
- `mpls-packet-format-5`—Single tagged packet with one label.
- `mpls-packet-format-6`—Single tagged packet with two labels.
- `mpls-packet-format-7`—Untagged or single tagged packet with one label.
- `mpls-packet-format-8`—Untagged or single tagged packet with two labels.

recycle-bandwidth-profiles Configure a profile that maps a feature to the recycle port bandwidth.

system-profile Specify the system-profile

- default-profile—Select the default system profile.
- vpls-meshgroups—Select the VPLS mesh group system profile.
- vxlan-extended—Select the vxlan-extended system profile.
- vxlan-stitching—Select the vxlan-stitching system profile.

> **NOTE**: You are required to configure `set system packet-forwarding-options system-profile vxlan-stitching` on ACX series routers to enable Layer 2 stitching. The PFE will restart once this configuration is applied.

warm-boot Enable warm boot support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced prior to Junos OS Release 21.1R1 and Junos OS Evolved Release 21.1R1.

`no-ip-tos-rewrite` option introduced in Junos OS Evolved Release 22.1R1.

`ipfix-sw-mode` option introduced in Junos OS Release 22.2R1.

`vpls-meshgroups` option introduced in Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

Ingress and Egress Multicast Replication using Recycle Replication Model

Flow-Based Telemetry for VXLANs (QFX5120)

packet-log (Security Flow)

IN THIS SECTION

- Syntax | 12151
- Hierarchy Level | 12151
- Description | 12151

- Options | 12151
- Required Privilege Level | 12152
- Release Information | 12152

Syntax

```
packet-log {
    enable;
    throttle-interval;
    packet-filter <filter-name>;
}
```

Hierarchy Level

```
[edit security flow ]
```

Description

Configure flow packet log.

Starting in Junos OS Release 17.3R1, SRX Series Firewalls support logging of unsynchronized out-of-state TCP packets that are dropped by the flow module.

The logging functionality will allow notifications to the administrators (or automated systems) about any of TCP out-of-state packets drops.

Options

enable	Enable log for dropped packets.
packet-filter	Configure packet log filter.
throttle-interval	Configure the interval as a power of two (0..32768 milliseconds).

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

packet-log (Security IDP Policy)

IN THIS SECTION

- Syntax | 12152
- Hierarchy Level | 12152
- Description | 12153
- Options | 12153
- Required Privilege Level | 12153
- Release Information | 12153

Syntax

```
packet-log {
    post-attack number;
    post-attack-timeout seconds;
    pre-attack number;
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name then notification]
```

Description

In response to a rule match, capture the packets received before and after the attack for further offline analysis of attacker behavior. You can configure the number of pre-attack and post-attack packets to be captured for this attack, and limit the duration of post-attack packet capture by specifying a timeout value.

In case of multiple attacks that are hit for the same session, packets capture occurs for the first attack that was hit on that specific session ID.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

packet-log (Security IDP Sensor Configuration)

IN THIS SECTION

- Syntax | 12154
- Syntax (Logical Systems and Tenant Systems) | 12154
- Hierarchy Level | 12154
- Description | 12154
- Options | 12154
- Required Privilege Level | 12155
- Release Information | 12155

Syntax

```
packet-log {
    host (Security IDP Sensor Configuration) <ip-address>;
    port <port number>;
    max-sessions (Security Packet Log) <percentage>;
    threshold-logging-interval <minutes>;
    source-address (Security IDP Sensor Configuration) <ip-address>;
    ssl-profile-name <profile-name>;
total-memory <percentage>}
```

Syntax (Logical Systems and Tenant Systems)

```
packet-log {
    host (Security IDP Sensor Configuration) <ip-address>;
    port <port number>;
    source-address (Security IDP Sensor Configuration) <ip-address>;
    ssl-profile-name <profile-name>;
    threshold-logging-interval <minutes>}
```

Hierarchy Level

```
[edit security idp sensor-configuration]
[edit tenants tenant name security idp sensor-configuration]
[edit logical-systems logical system name security idp sensor-configuration]
```

Description

Configure the sensor for packet capture. This configuration defines the amount of memory to be allocated for packet capture and the maximum number of sessions that can generate packet capture data for the device at one time. The configuration also identifies the source address and host address for transmission of the completed packet capture object.

Options

- `ssl-profile-name` *`profile-name`* —SSL initiation profile name to be used for encrypted packet log transmission.

Starting in Junos OS Release 22.1R1, you can enable a secure SSL connection and send encrypted IDP packet capture log to the packet capture receiver. To establish the SSL connection, you must specify the SSL initiation profile that you want to use in the IDP packet log configuration.

The SSL profile name must be configured in the SSL initiation profile configuration. SSL initiation profile configuration performs the required SSL certificates and SSL handshake operations to establish a secure connection. SSL versions are chosen based on the SSL initiation configuration. Also, SSL profile needs to be configured in each logical system separately.

If SSL profile name is not configured in SSL initiation profile configuration, then the following message is displayed **Referenced SSL initiation profile is not defined**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

IDP Security Packet Capture

packet-loss-priority (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 12156
- Hierarchy Level | 12156
- Description | 12156
- Required Privilege Level | 12156
- Release Information | 12156

Syntax

```
packet-loss-priority;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specify the packet loss priority (PLP) value to be used for the test frames generated in the packet forwarding pipeline. If a value is not configured, then the default value of `low` is used.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

packet-ordering-mode (Application Services)

IN THIS SECTION

- Syntax | 12157
- Hierarchy Level | 12157
- Description | 12157
- Options | 12158
- Required Privilege Level | 12158
- Release Information | 12158

Syntax

```
packet-ordering-mode {
    (hardware | software);
}
```

Hierarchy Level

[edit security forwarding-process application-services]

Description

Enables or disables the packet-ordering functionality using the Packet Ordering Engine. By default, packet-ordering functionality using the Packet Ordering Engine (hardware) is enabled.

A system reboot is required when this feature is enabled or disabled, and a warning message is displayed during the commit.

NOTE: Packet-ordering functionality using Packet Ordering Engine is supported on SRX5800 and SRX5600 devices with next-generation SPCs. Starting from Junos OS release 12.1X46-D10, SRX5400 device with next-generation SPCs also supports this feature.

Options

hardware Enables packet-ordering functionality using the Packet Ordering Engine.

software Disables packet-ordering functionality using the Packet Ordering Engine.

Required Privilege Level

- security—To view this in the configuration.
- security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10. This statement is supported.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

packet-scheduling

IN THIS SECTION

- Syntax | 12159
- Hierarchy Level | 12159
- Description | 12159
- Default | 12159
- Options | 12159
- Required Privilege Level | 12159
- Release Information | 12159

Syntax

```
(packet-scheduling | no-packet-scheduling);
```

Hierarchy Level

```
[edit chassis]
```

Description

(M 160 routers only) Enable packet-scheduling mode, in which the Packet Director application-specific integrated circuit (ASIC) schedules packet dispatches to compensate for transport delay differences. This preserves the interpacket gaps as the packets are distributed from the Packet Director ASIC to the Packet Forwarding Engine.

Default

no-packet-scheduling

> **NOTE**: The packet-scheduling feature is available on M160 routers only.

Options

no-packet-scheduling—Do not schedule packets.

packet-scheduling—Schedule packets to preserve interpacket gaps.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling an M160 Router to Operate in Packet Scheduling Mode

packet-size (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 12160
- Junos OS Hierarchy Level | 12160
- Junos OS Evolved Hierarchy Level | 12160
- Description | 12160
- Options | 12161
- Required Privilege Level | 12161
- Release Information | 12161

Syntax

```
packet-size [bytes];
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking profiles test-profile profile-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 test-profile profile-name]
```

Description

Define up to 10 packet sizes that are used sequentially for the test.

Options

bytes Size of the test packet. If you enter multiple packet sizes, you must separate each number with a space. You should base the test packet size on the configured MTU.

- **Range:** 64 through 9136 bytes
- **Minimum Packet Sizes (Junos OS Evolved):**
  - Untagged egress logical interface: 64 bytes
  - Single tagged interface: 66 bytes
  - Dual-tagged egress interface: 70 bytes

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved 23.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

rfc2544

packet-timestamp

IN THIS SECTION

- Syntax | 12162
- Hierarchy Level | 12162
- Description | 12162
- Options | 12162
- Required Privilege Level | 12163
- Release Information | 12163

Syntax

```
packet-timestamp {
    enable;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number traffic-manager]
```

Description

On MX Series routers, enable timestamping of class-of-service (CoS) queues for a configured Flexible PIC Concentrator (FPC).

Options

enable Display the last time a packet was enqueued as day, date, time, and year in the format *day-of-the-week month day-date hh:mm:ss yyyy*, for example, `Mon Jul 7 07:36:46 2014`. The Packet Forwarding Engine collects the timestamp for all inbound and outbound queue counters for all subscribers that are configured on the FPC and, when requested, also returns statistics corresponding to data traffic on the router.

If the timestamp is not enabled for the CoS queue, the `Last-packet enqueued` field is not displayed in statistics.

> **NOTE**: When you commit the change after enabling or disabling the timestamp feature for a particular FPC *after* that is already up, the FPC is automatically restarted and the change takes effect when it is back up.

The Routing Engine generates a system log message in the following situations:

- When `packet-timestamp` is enabled and committed for an FPC slot, but the FPC inserted into that slot does not support timestamping.
- When `packet-timestamp` is enabled and committed for an FPC slot, but an FPC that does not support timestamping is later inserted into that slot.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Enabling a Timestamp for Ingress and Egress Queue Packets

show interfaces queue | **21208**

packet-types (Dynamic VLAN Authentication)

IN THIS SECTION

- Syntax | **12164**
- Hierarchy Level | **12164**

- Description | 12164
- Options | 12164
- Required Privilege Level | 12165
- Release Information | 12165

Syntax

```
packet-types [packet-types]
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication]
```

Description

Specify one or more packet types to trigger authentication of an auto-configured dynamic VLAN. The packet types must be a subset of the packet types configured in the VLAN dynamic profile to trigger creation of the dynamic VLAN.

Options

packet-type—One or more of the following packet types that triggers VLAN authentication:

- `any`—Any packet type.
- `dhcp-v4`—IPv4 DHCP packet type.
- `dhcp-v6`—IPv6 DHCP packet type.
- `inet`—IPv4 Ethernet and ARP packet type.
- `inet6`—IPv6 Ethernet packet type.
- `pppoe`—Point-to-Point Protocol over Ethernet packet type.

NOTE: The pppoe VLAN Ethernet packet type option is supported only for MIC and MPC interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring Subscriber Packet Types to Trigger VLAN Authentication

Subscriber Packet Type Authentication Triggers for Dynamic VLANs

packet-timestamp

IN THIS SECTION

- Syntax | 12166
- Hierarchy Level | 12166
- Description | 12166
- Required Privilege Level | 12166
- Release Information | 12166

Syntax

```
packet-timestamp;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring]
```

Description

Apply 64-bit nanosecond EPOCH timestamp on mirrored packets in family `any`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

pado-advertise

IN THIS SECTION

- Syntax | 12167
- Hierarchy Level | 12167

- Description | 12167
- Required Privilege Level | 12167
- Release Information | 12167

Syntax

```
pado-advertise;
```

Hierarchy Level

```
[edit protocols pppoe]
```

Description

Enable named services configured in PPPoE service name tables to be advertised in PPPoE Active Discovery Offer (PADO) control packets. By default, advertisement of named services in PADO packets is disabled.

> **NOTE**: If you enable advertisement of named services in PADO packets, make sure the number and length of of all advertised service entries does not exceed the maximum transmission unit (MTU) size of the PPPoE underlying interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

Enabling Advertisement of Named Services in PADO Control Packets

pair-policy

IN THIS SECTION

- Syntax | 12168
- Hierarchy Level | 12168
- Description | 12168
- Required Privilege Level | 12169
- Release Information | 12169

Syntax

```
pair-policy pair-policy ;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
tunnel]
```

Description

Link the policy that you are configuring with another policy that references the same VPN tunnel so that both policies share one proxy ID and one security association (SA). Policy pairing is useful when you want to allow bidirectional traffic over a policy-based VPN that is using source or destination address translation with a dynamic IP address pool or destination address translation with a mapped IP (MIP) or dynamic IP (DIP) address pool.

Without policy pairing, the device derives a different proxy ID from the outbound and inbound policies. Two proxy IDs causes a problem for the remote peer with a single proxy ID for the VPN tunnel.

Pairing two policies solves the proxy ID problem for the remote peer and conserves SA resources. The single proxy ID is derived from the policy you configured last.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

pap

IN THIS SECTION

- Syntax | 12169
- Hierarchy Level | 12170
- Description | 12170
- Required Privilege Level | 12170
- Release Information | 12170

Syntax

```
pap {
    access-profile name;
    default-pap-password password;
```

```
    local-name name;
    local-password password;
    passive;
}
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options],
[edit interfaces interface-name unit logical-unit-number ppp-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options]
```

Description

Configure the Password Authentication Protocol (PAP). Use PAP authentication as a means to provide a simple method for the peer to establish its identity using a two-way handshake. This is done only upon initial link establishment.

After the link is established, an ID and password pair is repeatedly sent by the peer to the authenticator until authentication is acknowledged or the connection is terminated.

> **BEST PRACTICE**: On inline service (si) interfaces for L2TP, only the `pap` statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring the PPP Challenge Handshake Authentication Protocol

Configuring the PPP Password Authentication Protocol

Configuring PPP

traceoptions (PPP Process)

Configuring L2TP

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

pap (Dynamic PPP)

IN THIS SECTION

- Syntax | 12171
- Hierarchy Level | 12171
- Description | 12171
- Required Privilege Level | 12172
- Release Information | 12172

Syntax

```
pap;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options],
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-
interface-unit” ppp-options]
```

Description

Specify PAP authentication in a PPP dynamic profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit” ppp-options]` hierarchy level introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Dynamic Profiles Overview

Configuring Dynamic Authentication for PPP Subscribers

Attaching Dynamic Profiles to Static PPP Subscriber Interfaces

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

pap (L2TP)

IN THIS SECTION

- Syntax | 12172
- Hierarchy Level | 12173
- Description | 12173
- Required Privilege Level | 12173
- Release Information | 12173

Syntax

```
pap;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp ppp-options]
```

Description

(MX Series routers only) Specify PAP authentication for PPP subscribers in an L2TP LNS user group profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

parameter-preference

IN THIS SECTION

- Syntax | 12174
- Hierarchy Level | 12174
- Description | 12174
- Options | 12174
- Required Privilege Level | 12174
- Release Information | 12174

Syntax

```
parameter-preference(configured | proxied);
```

Hierarchy Level

```
[edit logical-systems name protocols router-advertisement interface],
[edit protocols router-advertisement interface]
```

Description

Preference to select configured or proxied parameters for downstream interface.

Options

configured Send configured parameters on downstream interface

proxied Send proxied parameters on downstream interface

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1R1

parameters

IN THIS SECTION

- Syntax | 12175
- Hierarchy Level | 12175
- Description | 12175
- Required Privilege Level | 12175
- Release Information | 12176

Syntax

```
parameters {
    message-processing-model (v1 | v2c | v3);
    security-level (none | authentication | privacy);
    security-model (usm | v1 | v2c);
    security-name security-name;
}
```

Hierarchy Level

```
[edit snmp v3 target-parameters target-parameters-name]
```

Description

Configure a set of target parameters for message processing and security.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

parse-direction (Access Profile)

IN THIS SECTION

- Syntax | 12176
- Hierarchy Level | 12176
- Description | 12176
- Default | 12177
- Options | 12177
- Required Privilege Level | 12177
- Release Information | 12177

Syntax

```
parse-direction (left-to-right | right-to-left);
```

Hierarchy Level

```
[edit access profile profile-name session-options strip-user-name]
```

Description

Specify the direction in which a subscriber login string is parsed to identify the first delimiter that matches one configured with the "delimiter" on page 5286 statement. When subscriber username stripping is configured in a subscriber access profile, the characters to the right of the identified delimiter are stripped and discarded along with the delimiter. characters become the new, modified username.

Default

`left-to-right`

Options

left-to-right Parse the subscriber login string from left to right up to the delimiter.

For example, when the direction is `left-to-right`, the characters `/@$%#` are configured as the delimiters, and the login string is `drgt21@example.com$84`, the `@` is reached before the `$`, so the username is modified to `drgt21`.

right-to-left Parse the subscriber login string from right to left up to the delimiter.

For example, when the direction is `right-to-left`, the characters `/@$%#` are configured as the delimiters, and the login string is `drgt21@example.com$84`, the `$` is reached before the `@`, so the username is modified to `drgt21@example.com`.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access

Configuring Username Modification for Subscriber Sessions

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

parse-direction (Domain Map)

IN THIS SECTION

- Syntax | 12178
- Hierarchy Level | 12178
- Description | 12178
- Default | 12178
- Options | 12178
- Required Privilege Level | 12179
- Release Information | 12179

Syntax

```
parse-direction (left-to-right | right-to-left);
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the direction in which the router searches for the domain name in a username.

Default

`left-to-right`

Options

`left-to-right`—The router searches starting at the left-most character. When the router reaches a domain delimiter, it uses anything to the right of the delimiter as the domain name.

`right-to-left`—The router searches starting at the right-most character. When the router reaches a domain delimiter, it uses anything to the right of the delimiter as the domain name.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Specifying the Parsing Direction for Domain and Realm Names

Configuring Domain and Realm Name Usage for Domain Maps

parse-order (Domain Map)

IN THIS SECTION

- Syntax | 12179
- Hierarchy Level | 12180
- Description | 12180
- Default | 12180
- Options | 12180
- Required Privilege Level | 12180
- Release Information | 12180

Syntax

```
parse-order (domain-first | realm-first);
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the order in which the router searches for a domain name, either the domain first, or the realm first.

Default

`domain-first`

Options

`domain-first`—The router searches for a domain name first. When the router reaches a domain delimiter, it uses anything to the right of the delimiter as the domain name.. If no domain is found, then the router searches for a realm. If the router does not find either a domain or realm, then there is no domain.

`realm-first`—The router searches for a realm name first. When the router reaches a realm delimiter, it uses anything to the left of the delimiter as the domain name. If no realm is found, then the router searches for a domain. If the router does not find either a domain or realm, then there is no domain.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Specifying the Parsing Direction for Domain and Realm Names

Configuring Domain and Realm Name Usage for Domain Maps

partition

IN THIS SECTION

- Syntax | 12181
- Hierarchy Level | 12181
- Description | 12181
- Options | 12181
- Required Privilege Level | 12182
- Release Information | 12182

Syntax

```
partition partition-name  {
    diameter-instance instance-name;
    destination-host hostname;
    destination-realm realm;
}
```

Hierarchy Level

```
[edit jsrc]
```

Description

Configure a JSRC partition.

Options

partition-name—Name of the JSRC partition.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

> *JSRC Configuration Overview*
>
> *Configuring the JSRC Partition*

partition (Gx-Plus)

IN THIS SECTION

- Syntax | 12182
- Hierarchy Level | 12183
- Description | 12183
- Options | 12183
- Required Privilege Level | 12183
- Release Information | 12183

Syntax

```
partition partition-name  {
    diameter-instance instance-name;
    destination-host hostname;
    destination-realm realm;
}
```

Hierarchy Level

```
[edit access gx-plus]
```

Description

Configure a Gx-Plus partition.

Options

partition-name—Name of the Gx-Plus partition.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Gx-Plus

Configuring the Gx-Plus Partition

partition (NASREQ Diameter Application)

IN THIS SECTION

Syntax | 12184

- Hierarchy Level | 12184
- Description | 12184
- Options | 12184
- Required Privilege Level | 12185
- Release Information | 12185

Syntax

```
partition partition-name {
    destination-host hostname;
    destination-realm realm-name;
    diameter-instance master;
}
```

Hierarchy Level

```
[edit access nasreq]
```

Description

Define the NASREQ partition by specifying the destination host and realm where the application resides. Only the`master` routing instance is supported for the application partition.

Options

diameter-instance master	Specifies that the master Diameter instance is used. `master` is the only supported value.
hostname	Name of the host where the NASREQ server application resides; generally, this value is not set unless needed.
partition-name	Name of the NASREQ partition.
realm-name	Name of the realm where the NASREQ server application resides.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Diameter Network Access Server Application (NASREQ)*
>
> *Configuring the Diameter Network Access Server Application (NASREQ)*

partition (OCS)

IN THIS SECTION

- Syntax | 12185
- Hierarchy Level | 12186
- Description | 12186
- Options | 12186
- Required Privilege Level | 12187
- Release Information | 12187

Syntax

```
partition partition-name  {
    called-station-id station-name;
    charging-id number;
    destination-host ocs-hostname;
    destination-realm ocs-realm-name;
    diameter-instance;
```

```
    draining;
    draining-response-timeout seconds;
    final-response-timeout seconds;
    force-continue;
    ggsn-address address;
    ggsn-mcc-mncggsn-mcc-mnc;
    max-outstanding-requests number;
    send-origin-state-id number;
    user-name-include  {
        base-interface-name;
        delimiter delimiter-character;
        domain-name domain-name;
        interface-name;
        mac-address;
        nas-port-id;
        origin-host;
        origin-realm;
        user-name;
        user-prefix prefix;
    }
}
```

Hierarchy Level

```
[edit access ocs]
```

Description

Configure an Online Charging System (OCS) partition (a specific logical system:routing instance context) and its parameters to define user ID information, destination host and realm names, and number of outstanding requests to the OCS, and control subscriber traffic flow before the first interrogation with the OCS. The OCS that interacts with the Policy and Charging Enforcement Function (PCEF). Broadband PCEF (BPCEF) interactions with the OCS use online session charging with centralized unit determination and centralized rating.

Options

partition-name—Name of the OCS partition. You can define only one partition for the OCS.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the OCS Partition

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

partition (PCRF)

IN THIS SECTION

- Syntax | **12188**
- Hierarchy Level | **12189**
- Description | **12189**
- Options | **12189**
- Required Privilege Level | **12190**
- Release Information | **12190**

Syntax

```
partition partition-name  {
    destination-host pcrf-hostname;
    destination-realm pcrf-realm-name;
    diameter-instance;
    draining;
    draining-response-timeout seconds;
    ip-can-type number;
    local-decision {
        deny;
        grant;
        reinit-on-failure;
        reinit-on-rar;
        reinit-timeout seconds;
        timeout seconds;
    }
    logout-response-timeout seconds;
    max-outstanding-requests number;
    report-local-rule;
    report-resource-allocation;
    report-successful-resource-allocation;
    send-dyn-subscription-indicator;
    send-network-family-indicator;
    send-origin-state-id;
    subscription-id-data-include  {
        base-interface-name;
        delimiter delimiter-character;
        domain-name name;
        interface-name;
        mac-address;
        nas-port-id;
        origin-host;
        origin-realm;
        user-name;
        user-prefix prefix;
        vlan-tags;
    }
    subscription-id-type number;
    update-response-timeout seconds;
```

```
    use-session-stamp;
}
```

Hierarchy Level

```
[edit access pcrf]
```

Description

Configure a Policy and Charging Rules Function (PCRF) partition (a specific logical system:routing instance context) and its parameters to define subscription ID information, destination host and realm names, and reporting for rules and resources. The PCRF is a centralized policy decision point that deploys business policy rules to allocate broadband network resources and manages flow-based charges for subscribers and services.

Options

partition-name Name of the PCRF partition. You can define only one partition for the PCRF.

use-session-stamp (Optional) Adds a 32-bit string to the end of the PCRF session ID for new subscribers. Existing subscribers are not affected when you configure this option. The string, called the session stamp or the session timestamp, consists of the UTC time when the router creates the CCR-GX-I. Appending the timestamp is considered to be an extended format for the session ID. You can use the extended-format session ID when you need an eternally unique ID or when you configure local reinitialization.

> **NOTE**: You must configure the `use-session-stamp` option when you configure local reinitialization with the `reinit-on-failure` or `reinit-on-rar` options with the "local-decision" on page 9934 statement.

> **NOTE**: This configuration also affects OCS sessions without any further configuration. The session ID for a given subscriber is the same for both Gx and Gy sessions.

Click a linked statement in the Syntax section for more information about that statement.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

`use-session-stamp` option added in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

Understanding Upstream and Downstream Messages for the PCRF

partition (s6a)

IN THIS SECTION

- Syntax | 12191
- Hierarchy Level | 12191
- Description | 12191
- Options | 12191
- Required Privilege Level | 12192
- Release Information | 12192

Syntax

```
partition name {
    destination-host destination-host;
    destination-realm destination-realm;
    diameter-instance diameter-instance;
    max-outstanding-requests max-outstanding-requests;
    response-timeout seconds;
}
```

Hierarchy Level

```
[edit access s6a]
```

Description

Define the s6a partition by specifying the destination host and realm where the application resides and specify the endpoint origin, the remote peers, and the network elements that associate routes with peers to control diameter forwarding of S6a messages. Only the `master` routing instance is supported for the application partition.

Options

name Name of the S6a partition.

destination-host Name of the host where the S6A server application resides; generally, this value is not set unless needed.

destination-realm Name of the realm where the S6a server application resides.

diameter-instance Specifies that the master Diameter instance is used. `master` is the only supported value.

max-outstanding-requests Configure the maximum number of outstanding requests for the S6a application.

- **Default:** 40
- **Range:** 2 through 1024

response-timeout Configure the amount of time in seconds that the Mobility Management Entity (MME) waits to receive a response from the Home Subscriber Server (HSS).

- **Default:** 15
- **Range:** 1 through 30

Required Privilege Level

access

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Example: Configure S6a Application

authentication-order | 2633

Diameter Base Protocol Overview

pass-through (interfaces)

IN THIS SECTION

- Syntax | 12193
- Hierarchy Level | 12193
- Description | 12193
- Default | 12193
- Options | 12193
- Required Privilege Level | 12193
- Release Information | 12193

Syntax

```
(pass-through | no-pass-through);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

Description

Enable or disable OTN pass-through mode.

Default

By default, OTN pass-through mode is disabled.

Options

no-pass-through	Do not enable OTN pass-through mode.
pass-through	Enable OTN pass-through mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

> *100-Gigabit Ethernet OTN Options Configuration Overview*
>
> *Configuring OTN Interfaces on P1-PTX-2-100G-WDM*

pass-through (Security Policy)

IN THIS SECTION

- Syntax | 12194
- Hierarchy Level | 12194
- Description | 12194
- Options | 12195
- Required Privilege Level | 12195
- Release Information | 12195

Syntax

```
pass-through {
    access-profile profile-name;
    client-match user-or-group-name;
    ssl-termination-profile profile-name;
    web-redirect;
    web-redirect-to-https;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication]
```

Description

Configure pass-through firewall user authentication. The user needs to use an FTP, Telnet, or HTTP client to access the IP address of the protected resource in another zone. Subsequent traffic from the user or host is allowed or denied based on the result of this authentication. Once authenticated, the firewall proxies the connection.

Options

- `access-profile` *`profile-name`* —(Optional) Specify the name of the access profile.
- `client-match` *`user-or-group`* —(Optional) Specify the name of the users or user groups in a profile who are allowed access by this policy. If you do not specify any users or user groups, any user who is successfully authenticated is allowed access.
- `ssl-termination-profile` *`profile-name`* —(Optional) Specify the SSL termination profile used for SSL offloading.
- `web-redirect`—(Optional) Enable redirecting an HTTP request to the device and redirecting the client system to a webpage for authentication. Including this statement allows users an easier authentication process because they need to know only the name or IP address of the resource they are trying to access.
- `web-redirect-to-https`—(Optional) Redirect unauthenticated HTTP requests to the internal HTTPS Web server of the device.

> **NOTE**: If `web-redirect-to-https` is set, then you must specify the SSL termination profile used for SSL offloading.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for `ssl-termination-profile` and `web-redirect-to-https` options added in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Security Policies Overview

pass-through (Firewall-authentication)

IN THIS SECTION

- Syntax | 12196
- Hierarchy Level | 12197
- Description | 12197
- Options | 12197
- Required Privilege Level | 12197
- Release Information | 12197

Syntax

```
pass-through {
    default-profile profile-name;
    ftp {
         banner {
            fail string;
            login string;
            success string;
        }
    }
    http {
         banner {
            fail string;
            login string;
             success string;
        }
    telnet {
        banner {
            fail string;
            login string;
            success string;
        }
}
```

Hierarchy Level

```
[edit access firewall-authentication]
[edit logical-systems name tenants name access firewall-authentication],
[edit tenants name access firewall-authentication]
```

Description

Configure pass-through , when a host or user from one zone needs to access a protected resource in another zone. A user must use an FTP, Telnet, or HTTP client to access the IP address of the protected resource and get authenticated by the firewall. The device uses FTP, Telnet, and HTTP to collect username and password information. Subsequent traffic from the user or host is allowed or denied based on the result of this authentication. After the user is authenticated, the firewall proxies the connection.

Options

default-profile Name of profile to use if not specified in policy

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Statement introduced for logical-systems and tenant in Junos OS Release 18.3R1.

HTTPS for pass-through authentication is supported on SRX5400, SRX5600, and SRX5800 devices starting from Junos OS Release 12.1X44-D10 and on vSRX Virtual Firewall, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Obtaining Username and Role Information Through Firewall Authentication

passive (CHAP)

IN THIS SECTION

- Syntax | 12198
- Hierarchy Level | 12198
- Description | 12198
- Required Privilege Level | 12199
- Release Information | 12199

Syntax

```
passive;
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options chap],
[edit interfaces interface-name unit logical-unit-number ppp-options chap],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options chap]
```

Description

Do not challenge the peer, but respond if challenged. If you omit this statement from the configuration, the interface always challenges its peer.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- `atm-ppp-llc`—PPP over AAL5 LLC encapsulation.
- `atm-ppp-vc-mux`—PPP over AAL5 multiplex encapsulation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Passive Mode

passive (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | 12199
- Hierarchy Level | 12200
- Description | 12200
- Options | 12200
- Required Privilege Level | 12200
- Release Information | 12200

Syntax

```
passive <allow-receive> <send-general-query> <send-group-query>;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name]
```

Description

Dynamically specify that IGMP run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as IGMP reports, queries, and leaves.

> **NOTE**: You can selectively activate up to two out of the three available options for the `passive` statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the `passive` statement.

Options

`allow-receive`—Enables IGMP to receive control traffic on the interface.

`send-general-query`—Enables IGMP to send general queries on the interface.

`send-group-query`—Enables IGMP to send group-specific and group-source-specific queries on the interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

`allow-receive`, `send-general-query`, and `send-group-query` options were introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- Example: Configuring Multicast with Subscriber VLANs
- Multicast Protocols User Guide

passive (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 12201
- Hierarchy Level | 12201
- Description | 12201
- Options | 12201
- Required Privilege Level | 12202
- Release Information | 12202

Syntax

```
passive <allow-receive> <send-general-query> <send-group-query>;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Description

Specify that MLD run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as MLD reports, queries, and leaves.

NOTE: You can selectively activate up to two out of the three available options for the `passive` statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the `passive` statement.

Options

`allow-receive`—(Optional) Enables MLD to receive control traffic on the interface.

`send-general-query`—(Optional) Enables MLD to send general queries on the interface.

`send-group-query`—(Optional) Enables MLD to send group-specific and group-source-specific queries on the interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Example: Configuring Multicast with Subscriber VLANs

passive (IGMP)

IN THIS SECTION

- Syntax | 12202
- Hierarchy Level | 12203
- Description | 12203
- Options | 12203
- Required Privilege Level | 12203
- Release Information | 12203

Syntax

```
passive <allow-receive> <send-general-query> <send-group-query>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Description

When configured for passive IGMP mode, the interface listens for IGMP reports but it will not send or receive IGMP control traffic such as IGMP reports, queries and leaves. You can, however, configure exceptions to allow the interface to receive certain control traffic or queries.

> **NOTE**: When an interface is configured for IGMP passive mode, Junos no longer processes static IGMP group membership on the interface.

Options

You can selectively activate up to two out of the three available options for the `passive` statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the `passive` statement.

allow-receive—Enables IGMP to receive control traffic on the interface.

send-general-query—Enables IGMP to send general queries on the interface.

send-group-query—Enables IGMP to send group-specific and group-source-specific queries on the interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

allow-receive, **send-general-query**, and **send-group-query** options were added in Junos OS Release 10.0.

RELATED DOCUMENTATION

Example: Configuring Multicast with Subscriber VLANs

Enabling IGMP

passive (MLD)

IN THIS SECTION

- Syntax | 12204
- Hierarchy Level | 12204
- Description | 12204
- Options | 12205
- Required Privilege Level | 12205
- Release Information | 12205

Syntax

```
passive <allow-receive> <send-general-query> <send-group-query>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

Specify that MLD run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as MLD reports, queries, and leaves.

> **NOTE**: You can selectively activate up to two out of the three available options for the `passive` statement while keeping the other functions passive (inactive). Activating all three options is equivalent to not using the `passive` statement.

Options

allow-receive—Enables MLD to receive control traffic on the interface.

send-general-query—Enables MLD to send general queries on the interface.

send-group-query—Enables MLD to send group-specific and group-source-specific queries on the interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

allow-receive, **send-general-query**, and **send-group-query** options added in Junos OS Release 10.0.

RELATED DOCUMENTATION

Example: Configuring Multicast with Subscriber VLANs

passive-mode

IN THIS SECTION

- Syntax | **12206**
- Hierarchy Level | **12206**
- Description | **12206**

Required Privilege Level | 12206
Release Information | 12206

Syntax

```
passive-mode;
```

Hierarchy Level

```
[edit logical-systems name protocols router-advertisement interface],
[edit protocols router-advertisement interface]
```

Description

Configure passive mode on an interface. When you configure this option on the interface, the interface only reieves and processes the RA packets and does not send the RAs (receive-only mode). The commit fails if the interface has both downstream and passive-mode option configured.

Configure RA receive only mode

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1R1.

passive-mode-tunneling

IN THIS SECTION

- Syntax | 12207
- Hierarchy Level | 12207
- Description | 12207
- Required Privilege Level | 12208
- Release Information | 12208

Syntax

```
passive-mode-tunneling;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Allows tunneling of malformed packets. When this statement is enabled, traffic bypasses the usual active IP checks. The IPsec tunnel is not treated as a next hop and TTL is not decremented. If the packet size exceeds the tunnel MTU value, an ICMP error is not generated. Starting with Junos OS Release 13.3R4 and 14.2R1, passive mode tunneling is supported on MS-MICs and MS-MPCs.

NOTE: The `header-integrity-check` option that is supported on MS-MICs and MS-MPCs to verify the packet header for anomalies in IP, TCP, UDP, and ICMP information and flag such anomalies and errors has a functionality that is opposite to the functionality caused by passive mode tunneling. If you configure both the `header-integrity-check` statement and the `passive-mode tunneling` statement on MS-MICs and MS-MPCs, and attempt to commit such a configuration, an error is displayed during commit.

The passive mode tunneling functionality (by including the `passive-mode-tunnelin` statement at the `[edit services service-set *service-set-name* ipsec-vpn-options]` hierarchy level) is a superset of the capability to disable IPsec tunnel endpoint in the traceroute output (by including `no-ipsec-tunnel-in-traceroute` statement at the `[edit services ipsec-vpn]` hierarchy level). Passive mode tunneling also bypasses the active IP checks and tunnel MTU check in addition to not treating an IPsec tunnel as a next-hop as configured by the `no-ipsec-tunnel-in-traceroute` statement.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring IPsec Service Sets

passive-mode-tunneling (Security IPsec)

IN THIS SECTION

- Syntax | 12209
- Hierarchy Level | 12209
- Description | 12209
- Options | 12209
- Required Privilege Level | 12209
- Platform-Specific Passive Mode Tunneling Behavior | 12209
- Release Information | 12211

Syntax

```
passive-mode-tunneling;
```

Hierarchy Level

```
[edit security ipsec vpn vpn-name]
```

Description

Allows tunneling of malformed packets. By default, the feature is disabled. When you enable this feature using the configuration statement `set security ipsec vpn vpn-name passive-mode-tunneling`:

- Traffic bypasses the usual active IP checks such as the version, TTL, protocol, options, and address.
- There is no effect on the TTL value (decrement) as IPsec tunnel is not treated as the next hop.
- Even if the packet size exceeds the tunnel MTU value, it doesn't generate ICMP error message or perform fragmentation.

Options

No specific options are needed. If you configure the statement, the feature is enabled.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Platform-Specific Passive Mode Tunneling Behavior

Use Feature Explorer to confirm platform and release support for specific features.

Use the following table to review platform-specific behavior for your platforms.

Platform	Difference
MX Series	• MX-SPC3 services card that support this feature allows the packets to bypass the flow processing and directly undergo encapsulation and decapsulation when passing through the passive mode tunnel. As a result, the device does not create session flow output. You'll not see session details in the output of the `show security flow session` command with packet based processing of IPsec traffic in passive mode tunnels. • You must configure `passive-monitor-mode` statement before enabling `passive-mode-tunneling` statement so that the malformed packets can reach the MX-SPC3 services card from the Packet Forwarding Engine (PFE). See passive-monitor-mode. • MX-SPC3 services card supports `header-integrity-check` statement in the `service-set` configuration to verify the packet header for anomalies in IP, TCP, UDP, and ICMP information. It is important to note that this functionality is opposite to that of the `passive-mode-tunneling` statement. So configuring both the `header-integrity-check` statement and the `passive-mode-tunneling` statement, results in an error during the configuration commit. See header-integrity-check. • In MX-SPC3 services card with `passive-mode-tunneling` enabled and `header-integrity-check` configured in the `service-set`, the following implications apply to multiple VPN configurations: • If the `header-integrity-check` statement is enabled, `passive-mode-tunneling` statement must be disabled for all VPNs. Additionally, `service-set` cannot have two or more IPsec VPNs with different `passive-mode-tunneling` values. This means that if `header-integrity-check` statement is enabled, a `service-set` can have only one type

(Continued)

Platform	Difference
	of VPN configured with either `passive-mode-tunneling` enabled or disabled. • If the `header-integrity-check` statement is disabled, then a `service-set` can have two or more IPsec VPNs with different `passive-mode-tunneling` values.
SRX Series	• SRX Series Firewalls that support this feature allows the firewall to process malformed packets and create a flow session. You can view the session flow details using the `show security flow session` command. The flow process checks the passive mode tunneling state and decides whether to bypass IP checks and prevent packet drops. • We do not support the feature for mixed mode configurations on your firewall. This implies that when you enable the feature on one IPsec VPN tunnel, it gets enabled for all the tunnels on your firewall. • Enabling passive mode tunneling on your firewall restricts the configuration of security features such as the screen options for attack detection and prevention, as well as the security policies (excluding the default permit all policy). Conversely, if you configure security features such as screens options and security policies, passive mode tunneling cannot be enabled. • We recommend you to configure the default permit all policy to avoid packet drop due to a policy match.

Release Information

Statement introduced in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

passive-mode-tunneling

passive-monitor-mode

IN THIS SECTION

- Syntax | 12212
- Hierarchy Level | 12212
- Description | 12212
- Required Privilege Level | 12212
- Release Information | 12213

Syntax

```
passive-monitor-mode;
```

Hierarchy Level

```
[edit interfaces ATM, Fast Ethernet, or Gigabit Ethernet interface-name],
[edit interfaces SONET/SDH interface-name unit logical-unit-number]
```

Description

For Asynchronous Transfer Mode (ATM), SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, monitor packet flows from another router. If you include this statement in the configuration, the SONET/SDH interface does not send keepalives or alarms, and does not participate actively on the network.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers

multiservice-options (Flow Monitoring) | 11130

passive (MVRP)

IN THIS SECTION

- Syntax | 12213
- Hierarchy Level | 12213
- Description | 12214
- Default | 12214
- Required Privilege Level | 12214
- Release Information | 12214

Syntax

```
passive;
```

Hierarchy Level

```
[edit protocols mvrp],

[edit protocols mvrp interface(all | interface-name)]
```

Description

Configure an MVRP interface to not announce its membership in a VLAN or send any VLAN declarations (updates) unless it receives registration for that VLAN from a peer (server).

Default

Passive mode is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Multiple VLAN Registration Protocol

passive (PAP)

IN THIS SECTION

- Syntax | 12215
- Hierarchy Level | 12215
- Description | 12215
- Required Privilege Level | 12215
- Release Information | 12215

Syntax

```
passive;
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options pap],
[edit interfaces interface-name unit logical-unit-number ppp-options pap],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options pap]
```

Description

Initiate an authentication request when the PAP option is received from a peer. If you omit this statement from the configuration, the interface requires the peer to initiate an authentication request.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

> *Configuring Passive Mode*
>
> Junos OS Administration Library for Routing Devices

passive-probe-params

IN THIS SECTION

- Syntax | 12216
- Hierarchy Level | 12216
- Description | 12216
- Options | 12217
- Required Privilege Level | 12217
- Release Information | 12217

Syntax

```
passive-probe-params {
    sampling-percentage {
        percentage;
    }
    sampling-period {
        period;
    }
    type {
        book-ended;
    }
```

Hierarchy Level

[edit security advance-policy-based-routing]

Description

Configure the passive probe parameters with the SLA rule.

Passive probes measure the service quality of an application by inserting a custom probe header in the live traffic between the spoke and hub points and measuring the RTT, jitter and packet loss between the points of installation of the probes.

SLA violation is determined through passive probing of live application or application group traffic.

Options

sampling-percentage *percentage*

Indicates the percentage of sessions that are selected for a book-ended SLA measurement.

Example: If 18 sessions are available for a particular application are available, and if you have configured 25%, then 25% of the 18 sessions—that is 5 sessions out of 18 sessions, are evaluated.

- **Range:** 1-100

sampling-period *period*

Indicates a defined sampling period (in milliseconds) in which the number of violations are collected. Once this period is expired, the collected sampling data is purged and a new data is collected.

- **Range:** 2000-60,000
- **Default:** 5000 milliseconds

type

Indicates the type of probe measurement, only p-encap or book-ended supported.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Application Quality of Experience

Advanced Policy-Based Routing

passive (Protocols BGP)

IN THIS SECTION

- Syntax | 12218
- Hierarchy Level | 12218
- Description | 12219
- Default | 12219
- Required Privilege Level | 12219
- Release Information | 12219

Syntax

```
passive;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Configure the routing device so that active open messages are not sent to the peer. Once you configure the routing device to be passive, the routing device will wait for the peer to issue an open request before a message is sent.

Default

If you omit this statement, all explicitly configured peers are active, and each peer periodically sends open requests until its peer responds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Preventing BGP Session Flaps When VPN Families Are Configured

passive (Protocols IS-IS)

IN THIS SECTION

- Syntax | 12220
- Hierarchy Level | 12220
- Description | 12220
- Default | 12221
- Options | 12221
- Required Privilege Level | 12221

Release Information | 12221

Syntax

```
passive {
    remote-node-id address;
    remote-node-iso iso-id;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis interface interface-name],
[edit logical-systems logical-system-name protocols isis interface interface-name level level-number],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis interface interface-name level level-number],
[edit protocols isis interface interface-name],
[edit protocols isis interface interface-name level level-number],
[edit routing-instances routing-instance-name protocols isis interface interface-name],
[edit routing-instances routing-instance-name protocols isis interface interface-name level level-number]
```

Description

Advertise the direct interface addresses on an interface or into a level on the interface without actually running IS-IS on that interface or level.

This statement effectively prevents IS-IS from running on the interface. To enable IS-IS on an interface, include the `interface` statement at the `[edit protocols isis]` or the `[edit routing-instances routing-instance-name protocols isis]` hierarchy level. To disable it, include the `disable` statement at those hierarchy levels. The three states—enabling, disabling, or not running IS-IS on an interface—are mutually exclusive.

> **NOTE**: Configuring IS-IS on a loopback interface automatically renders it as a passive interface, irrespective of whether the `passive` statement was used in the configuration of the interface.

If neither passive mode nor the `family iso` option is configured on the IS-IS interface, then the routing device treats the interface as not being operational, and no direct IPv4/IPv6 routes are exported into IS-IS. (You configure the `family iso` option at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level.)

Default

By default, IS-IS must be configured on an interface or a level for direct interface addresses to be advertised into that level.

Options

`remote-node-id address`	IP address of the remote link.
`remote-node-iso iso-id`	ISO ID of the remote node.

> **NOTE**: The options `remote-node-id address` and `remote-node-iso iso-id` do not apply under the `[edit routing-instances routing-instance-name protocols isis]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`remote-node-id address` option introduced in Junos OS Release 14.2.

`remote-node-iso iso-id` option introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding

passive (Protocols Link Management)

IN THIS SECTION

- Syntax | 12222
- Hierarchy Level | 12222
- Description | 12222
- Required Privilege Level | 12222
- Release Information | 12223

Syntax

```
passive;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-protocol],
[edit protocols link-management peer peer-name lmp-protocol]
```

Description

Specify that the router not configure the Link Management Protocol (LMP) control channels but wait for the remote peer to configure the LMP control channels.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Preventing the Local Peer from Initiating LMP Negotiation

passive (Protocols OSPF)

IN THIS SECTION

- Syntax | 12223
- Hierarchy Level | 12223
- Description | 12224
- Required Privilege Level | 12224
- Release Information | 12224

Syntax

```
passive {
    traffic-engineering {
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface
interface-name],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast) area area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-id interface interface-name],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-
name],
[edit protocols (ospf | ospf3) area area-id interface interface-name],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id
interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface
interface-name],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast) area area-id interface interface-name]
```

Description

Advertise the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run.

To configure an interface in OSPF passive traffic engineering mode, include the `traffic-engineering` statement. Configuring OSPF passive traffic engineering mode enables the dynamic discovery of OSPF AS boundary routers.

Enable OSPF on an interface by including the `interface` statement at the `[edit protocols (ospf | ospf3) area` *`area-id`*`]` or the `[edit routing-instances` *`routing-instance-name`* `protocols ospf area` *`area-id`*`]` hierarchy levels. Disable it by including the `disable` statement, To prevent OSPF from running on an interface, include the `passive` statement. These three states are mutually exclusive.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

traffic-engineering and `remote-node-id` *`address`* statements introduced in Junos OS Release 8.0.

traffic-engineering and `remote-node-id` *`address`* statements introduced in Junos OS Release 8.0 for EX Series switches.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring a Passive OSPF Interface

Example: Configuring OSPF Passive Traffic Engineering Mode

passive (Routing Options)

IN THIS SECTION

- See | 12225

See

"active" on page 1800

password (Access)

IN THIS SECTION

- Syntax | 12225
- Hierarchy Level | 12226
- Description | 12226
- Options | 12226
- Required Privilege Level | 12226
- Release Information | 12226

Syntax

```
password password;
```

Hierarchy Level

```
[edit access ldap-options search admin-search],
[edit access profile profile-name ldap-options search admin-search]
```

Description

Configure the plain-text password for the administrative user. This password is used in the bind for performing the LDAP search.

Options

password—Administrative user password.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring Integrated User Firewall on SRX Series

password (DHCP Local Server)

IN THIS SECTION

- Syntax | 12227
- Hierarchy Level | 12227
- Description | 12228

- Options | 12228
- Required Privilege Level | 12228
- Release Information | 12228

Syntax

```
password password-string;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6
authentication],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name authentication],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication],
```

```
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication],
[edit system services dhcp-local-server authentication],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name authentication],
[edit system services dhcp-local-server group group-name authentication]
```

Description

Configure the password that is sent to the external AAA authentication server for subscriber authentication or DHCP client authentication.

Options

password-string—Authentication password.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

password (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 12229
- Hierarchy Level | 12229
- Description | 12230
- Options | 12230
- Required Privilege Level | 12230
- Release Information | 12230

Syntax

```
password password-string;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication],
[edit forwarding-options dhcp-relay dhcpv6 authentication],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication],
[edit forwarding-options dhcp-relay group group-name authentication],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
authentication],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name authentication],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
```

```
options dhcp-relay dhcpv6 group group-name authentication],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
authentication],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication]
```

Description

Configure the password that is sent to the external AAA authentication server for subscriber authentication or client authentication. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Options

password-string—Authentication password.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name authentication]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

> *Specifying Authentication Support*
>
> *Example-Configuring DHCP with External Authentication Server*

password-file

IN THIS SECTION

- Syntax | 12231
- Hierarchy Level | 12231
- Description | 12231
- Options | 12231
- Required Privilege Level | 12232
- Release Information | 12232

Syntax

```
password-file (block | (log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name fallback-
options]
```

Description

Password protected file is the error returned by the scan engine when the scanned file is protected by a password. The default action is log-and-permit.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Full Antivirus Configuration Overview

password (Flow Collector File Servers)

IN THIS SECTION

- Syntax | 12232
- Hierarchy Level | 12232
- Description | 12233
- Options | 12233
- Required Privilege Level | 12233
- Release Information | 12233

Syntax

```
password "password";
```

Hierarchy Level

```
[edit services flow-collector destination ftp:url]
```

Description

Specify the primary and secondary destination FTP server password.

Options

password—FTP server password.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Destination FTP Servers for Flow Records

password (Interfaces)

IN THIS SECTION

- Syntax | 12234
- Hierarchy Level | 12234
- Description | 12234
- Options | 12234
- Required Privilege Level | 12234
- Release Information | 12234

Syntax

```
password password-string;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication]
```

Description

Configure the password that is sent to the external AAA authentication server for subscriber VLAN or stacked VLAN interface authentication.

Options

password-string—Authentication password.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring an Authentication Password for VLAN or Stacked VLAN Ranges

password (Login)

IN THIS SECTION

- Syntax | 12235
- Hierarchy Level | 12235
- Description | 12236
- Options | 12236
- Required Privilege Level | 12240
- Release Information | 12240

Syntax

```
password {
    change-type (character-sets | set-transitions);
    format (sha256 | sha512);
    maximum-length length;
    maximum-lifetime days;
    minimum-changes number;
    minimum-character-changes number;
    minimum-length length;
    minimum-lifetime days;
    minimum-lower-cases number;
    minimum-numerics number;
    minimum-punctuations number;
    minimum-reuse number;
    minimum-upper-cases number;
}
```

Hierarchy Level

```
[edit system login]
```

Description

Configure special requirements such as character length and encryption format for plain-text passwords. Newly created passwords must meet these requirements.

Using several password minimum requirement options will cause the `minimum-length` to be reset if the total sum of the required minimums exceeds the `minimum-length` setting.

Options

change-type

Set requirements for using character sets in plain-text passwords. When you combine this statement with the `minimum-changes` statement, you can check for the total number of character sets included in the password or for the total number of character-set changes in the password. Newly created passwords must meet these requirements.

- **Values:** Specify one of the following:
  - `character-sets`—The number of character sets in the password. Valid character sets include uppercase letters, lowercase letters, numbers, punctuation, and other special characters.
  - `set-transitions`—The number of transitions between character sets.

format

Configure the authentication algorithm for plain-text passwords. The hash algorithm that authenticates the password can be one of these algorithms:

- Values:
  - sha256—Secure Hash Algorithm 256. Produces a 256-bit digest. The encrypted password starts with 5.
  - sha512—Secure Hash Algorithm 512. Produces a 512-bit digest. The encrypted password starts with 6.
- **Default:** For Junos OS, the default encryption format is `sha512`.

maximum-length *length*

Specify the maximum number of characters allowed in plain-text passwords. Newly created passwords must meet this requirement.

- **Range:** 20 through 128 characters
- **Default:** For Junos-FIPS software, the maximum number of characters for plain-text passwords is 20. For Junos OS, no maximum is set.

maximum-lifetime *days*

Specify the maximum duration of a password in days, where the password expires after the maximum duration is reached. If you have the required permissions, you are able to control the maximum duration of a password. If the age of the password reaches the maximum time configured, the password expires and must be changed. If your password has expired, you cannot commit a configuration until you change your password. Only passwords for local user accounts can expire based on time configured on the maximum lifetime statement.

> **NOTE**: You cannot reuse the same password when the password expires, unless you also configure the number of times the password can be reused on the `minimum-reuse` statement. Older passwords cannot be re-configured on password expiry. Therefore, if you want to reuse an old password, you must configure the `minimum-reuse` statement as well as the `maximum-lifetime` statement in the new configuration, otherwise the commit fails.
>
> If the `maximum-lifetime` statement is configured, a validation check for an expired password is performed at the time of login and at the time of commit based on the password timestamp. For passwords configured before the `minimum-reuse` configuration statement is committed, the timestamp of the passwords is the time at which any configuration under the [edit system login] hierarchy is committed following the commit for the `minimum-reuse` statement. For passwords configured after `minimum-reuse` configuration statement is committed, the timestamp of the passwords is the time at which those passwords are committed.

- **Range:** 30 through 365 days

minimum-changes *number*

Specify the minimum number of character sets (or character set changes) required for plain-text passwords. Newly created passwords must meet this requirement.

This statement is used in combination with the `change-type` statement. If the change-type is `character-sets`, then the number of character sets included in the password is checked against the specified minimum. If change-type is `set-transitions`, then the number of character set changes in the password is checked against the specified minimum.

- **Default:** For Junos OS, the minimum number of changes is 1. For Junos-FIPS Software, the minimum number of changes is 3.

minimum-character-changes *number*

Specify the minimum number of character changes between old and new passwords. Newly created passwords must meet this requirement. If you have the required permissions, you are able to configure the number of character changes between

passwords. If the number of character changes between the old password and new password is greater than or equal to the configured value for minimum number of character changes, the new password is accepted. If the number of character changes is less than the configured value, the new password is rejected.

- **Range:** 4 through 15 characters

minimum-length *length*

Specify the minimum number of characters required in plain-text passwords. Newly created passwords must meet this requirement.

This statement can be used in combination with all of the other requirement options for plain-text passwords, such as `minimum-upper-cases`, `minimum-punctuations`, `minimum-lower-cases`, and so on.

Using several password minimum requirement options will cause the minimum password length to be reset if the total sum of the required minimums exceeds the setting configured on the `minimum-length` statement.

- **Default:** For Junos OS, the minimum number of characters for plain-text passwords is six. For Junos-FIPS software, the minimum number of characters for plain-text passwords is 10.

- **Range:** 6 through 20 characters

minimum-lifetime *days*

Specify in days the minimum duration of a password before the password can be changed. If you have the required permissions, you are able to control the minimum lifetime of a password. You cannot change the password if the age of the password does not exceed the duration configured on the `minimum-lifetime` statement. When you change a password, the age of the existing password is retrieved based on the time at which the password was configured and the current time is fetched. If the age of the password is less than or equal to the configured value for the `minimum-lifetime` statement, the new password is not accepted and an error message is displayed. If the age of the password is more than the configured value for the `minimum-lifetime` statement, the new password is accepted.

> **NOTE**: The `minimum-lifetime` statement can be committed only after configuring the `minimum-reuse` statement. The `minimum lifetime` statement works in coordination with password history requirements, else the commit fails and an error message is displayed.
>
> If `minimum-lifetime` is configured, password change for a user is accepted or rejected based on the timestamp of the current password for that user. For passwords configured before the `minimum-reuse` configuration

statement is committed, the timestamp of the passwords is the time at which any configuration under the [edit system login] hierarchy is committed following the commit for the `minimum-reuse` statement. For passwords configured after `minimum-reuse` configuration statement is committed, the timestamp of the passwords is the time at which those passwords are committed.

- **Range:** 1 through 30 days

minimum-lower-cases *number*

Specify the minimum number of lower-case letters required in plain-text passwords. Newly created passwords must meet this requirement.

This statement can be used in combination with all of the other requirement options for plain-text passwords, such as `minimum-length`, `minimum-punctuations`, `minimum-upper-cases`, and so on.

Using several password minimum requirement options will cause the minimum password length to be reset if the total sum of the required minimums exceeds the setting configured on the `minimum-length` statement.

- **Range:** 1 through 128 lower-case letters

minimum-numerics *number*

Specify the minimum number of numeric-class characters required in plain-text passwords. Newly created passwords must meet this requirement.

This statement can be used in combination with all of the other requirement options for plain-text passwords, such as `minimum-length`, `minimum-punctuations`, `minimum-lower-cases`, and so on.

Using several password minimum requirement options will cause the minimum password length to be reset if the total sum of the required minimums exceeds the setting configured on the `minimum-length` statement.

- **Range:** 1 through 128 numeric-class characters

minimum-punctuations *number*

Specify the minimum number of punctuation-class characters required in plain-text passwords. Newly created passwords must meet this requirement.

This statement can be used in combination with all of the other requirement options for plain-text passwords, such as `minimum-length`, `minimum-upper-cases`, `minimum-lower-cases`, and so on.

Using several password minimum requirement options will cause the minimum password length to be reset if the total sum of the required minimums exceeds the setting configured on the `minimum-length` statement.

- **Range:** 1 through 128 punctuation-class characters

minimum-reuse *number*

Specify the number of old passwords which should not match the new password. Newly created passwords must meet this requirement. If you have the required permissions, you are able to control the number of old passwords that need to be compared. The number of old passwords to compare with the new password depends on the value configured. If a match is found between the new password and any of the old passwords, the device rejects the new password and terminates. If the new password is different from the configured number of old passwords, the new password is accepted.

- **Range:** 1 through 20 passwords

minimum-upper-cases

Specify the minimum number of upper-case letters required in plain-text passwords. Newly created passwords must meet this requirement.

This statement can be used in combination with all of the other requirement options for plain-text passwords, such as `minimum-length`, `minimum-punctuations`, `minimum-lower-cases`, and so on.

Using several password minimum requirement options will cause the minimum password length to be reset if the total sum of the required minimums exceeds the setting configured on the `minimum-length` statement.

- **Range:** 1 through 128 upper-case letters

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

Statements `minimum-lower-cases`, `minimum-numerics`, `minimum-punctuations`, and `minimum-upper-cases` introduced in Junos OS Release 12.1.

All of the previously mentioned statements were introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Statements `minimum-reuse` and `minimum-character-changes` introduced in Junos OS Release 18.3.

Statements `maximum-lifetime` and `minimum-lifetime` introduced in Junos OS Release 18.4.

Option sha1 is not supported in Junos OS Evolved.

sha1 option deprecated in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Example: Changing the Requirements for Junos OS Plain-Text Passwords

Special Requirements for Junos OS Plain-Text Passwords

password (MAC RADIUS Authentication)

IN THIS SECTION

- Syntax | 12241
- Hierarchy Level | 12241
- Description | 12242
- Default | 12242
- Options | 12242
- Required Privilege Level | 12242
- Release Information | 12243

Syntax

```
password password-string;
```

Hierarchy Level

```
[edit protocols dot1x authenticator mac-radius]
```

Description

Configure a global password to be used as the password (instead of the MAC address) for all sessions that use MAC RADIUS authentication. Use of a global password is supported on all three MAC RADIUS-based authentication protocols (eap-md5, eap-peap, pap).

> **NOTE**: When enabled (or disabled), currently active dot1x sessions are not affected. MAC RADIUS authentication will use the configured password from the time of reauthentication.

Default

Disabled.

Options

password-string

Authentication password.

Password can be between 1 and 128 characters in length.

Password can include most character classes (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters), excluding space.

If you use any of the following character(s), enclose the character(s) within double quotation marks ("):

- [
-]
- &
- &&
- ||
- (
-)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Configuring MAC RADIUS Authentication (CLI Procedure)

password-options

IN THIS SECTION

- Syntax | 12243
- Hierarchy Level | 12244
- Description | 12244
- Options | 12244
- Required Privilege Level | 12244
- Release Information | 12244

Syntax

```
password-options {
    apply-groups;
    apply-groups-except;
    tacplus-authorization;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure options for local authentication.

> **NOTE**: To enable password-options, you must configure password under [edit system authentication-order]. The feature does not work in a local fallback scenario because password is not configured under authentication-order for a local fallback scenario.

Options

apply-groups—Choose the groups from which to inherit configuration data.

apply-groups-except—Choose the groups to be excluded from configuration data being inherited.

tacplus-authorization—Configure remote authorization on the TACACS+ server for locally authenticated users. Authorization parameters configured remotely for that user are combined with the authorization parameters configured locally.

- **Default:** If a user is authenticated locally, authorization for that user is done locally.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

Option tacplus-authorization introduced in Junos OS Release 19.3R1 on MX Series routers.

password (Security Antivirus)

IN THIS SECTION

- Syntax | 12245
- Hierarchy Level | 12245
- Description | 12245
- Required Privilege Level | 12245
- Release Information | 12246

Syntax

```
password password-string;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update proxy]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update proxy]
[edit security utm feature-profile anti-virus sophos-engine pattern-update proxy]
```

Description

Set the password for the proxy server.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Release 11.2.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Security | 4589

password (Services)

IN THIS SECTION

- Syntax | 12246
- Hierarchy Level | 12246
- Description | 12247
- Required Privilege Level | 12247
- Release Information | 12247

Syntax

```
password password;
```

Hierarchy Level

```
[edit services unified-access-control infranet-controller hostname]
```

Description

Specify the password that the SRX Series Firewall should send to the IC Series device to establish communications. The SRX Series Firewall sends the password in its first message to the IC Series device.

This statement is required when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

ca-profile (Services) | 3160

server-certificate-subject | 16618

password (Static Subscribers)

IN THIS SECTION

- Syntax | 12248
- Hierarchy Level | 12248
- Description | 12248
- Options | 12248
- Required Privilege Level | 12248
- Release Information | 12248

Syntax

```
password password-string;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers authentication],
[edit logical-systems logical-system-name system services static-subscribers authentication],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication],
[edit routing-instances routing-instances-name system services static-subscribers authentication],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],  authentication],
[edit system services static-subscribers authentication],
[edit system services static-subscribers group group-name authentication]
```

Description

Specify the password that is sent to AAA for user login for all static subscribers on interfaces configured at the `[edit system services static-subscribers interface]` hierarchy level, or for the subscribers in a specified group. The group version of the statement takes precedence over the global version.

Options

password-string—String that defines the password.

Required Privilege Level

system-level—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Configuring the Static Subscriber Global Authentication Password

Configuring the Static Subscriber Group Authentication Password

password (Transfer Log File Servers)

IN THIS SECTION

- Syntax | 12249
- Hierarchy Level | 12249
- Description | 12249
- Options | 12249
- Required Privilege Level | 12250
- Release Information | 12250

Syntax

```
password "password";
```

Hierarchy Level

```
[edit services flow-collector transfer-log-archive archive-sites]
```

Description

Specify the primary and secondary destination FTP server password.

Options

password—FTP server password.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Transfer Logs

path

IN THIS SECTION

- Syntax | 12250
- Hierarchy Level | 12251
- Description | 12251
- Options | 12251
- Required Privilege Level | 12251
- Release Information | 12252

Syntax

```
path destination {
    <address | hostname> <strict | loose>
}
```

Hierarchy Level

```
[edit protocols mpls]
```

Description

Configure path protection on your MPLS network.

Options

destination —Name of a label switched path (LSP). In addition to specifying the name of the configured LSP, you can include some other designation such as **primary-path**.

address —(Optional) IP address of each transit switch (or the IP address of the loopback interface on the switch) in the LSP. If you want to control exactly which switches are selected for the LSP, specify the address or hostname of each transit switch. Specify the addresses in order, starting with the first provider (transit) switch, and continuing sequentially along the path until reaching the egress provider edge switch.

- **Default:** If you do not specify the addresses or hostnames of any switches, the LSP is calculated by the switch.

hostname —(Optional) See ***address*** .

- **Default:** If you do not specify the addresses or hostnames of any switches, the LSP is calculated by the switch.

loose—(Optional) Indicates that the next address in the `path` statement is a loose link. This means that the LSP can traverse through other switches before reaching this switch.

- **Default: strict**

strict—(Optional) Indicates that the LSP must go to the next address specified in the `path` statement without traversing other switches. This is the default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Path Protection in an MPLS Network (CLI Procedure)

path-database-size

IN THIS SECTION

- Syntax | 12252
- Hierarchy Level | 12252
- Description | 12252
- Options | 12253
- Required Privilege Level | 12253
- Release Information | 12253

Syntax

```
path-database-size path-database-size;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management linktrace]
```

Description

Number of linktrace reply entries to be stored per linktrace request.

Options

`path-database-size`—Database size.

- **Range:** 1 through 255
- **Default:** 64

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Connectivity Fault Management (CFM)

path-database-size (EX Series Switch Only)

IN THIS SECTION

- Syntax | 12254
- Hierarchy Level | 12254
- Description | 12254
- Options | 12254
- Required Privilege Level | 12254
- Release Information | 12254

Syntax

```
path-database-size path-database-size;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management linktrace]
```

Description

Specify the number of linktrace reply entries to be stored per linktrace request.

Options

path-database-size—Database size (number of entries stored per request).

- **Range:** 1 through 500
- **Default:** 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

> *Configure Ethernet OAM Connectivity Fault Management (CLI Procedure)*
>
> Junos OS Network Interfaces Configuration Guide

path-length

IN THIS SECTION

- Syntax | 12255
- Hierarchy Level | 12255
- Description | 12255
- Options | 12255
- Required Privilege Level | 12256
- Release Information | 12256

Syntax

```
path-length certificate-path-length;
```

Hierarchy Level

```
[edit security certificates]
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the digital certificate path length.

Options

certificate-path-length—Digital certificate path length.

- **Range:** 2 through 15 certificates
- **Default:** 15 certificates

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Path Length for the Certificate Hierarchy

path-mtu

IN THIS SECTION

- Syntax | 12256
- Hierarchy Level | 12257
- Description | 12257
- Required Privilege Level | 12257
- Release Information | 12257

Syntax

```
path-mtu {
    allow-fragmentation;
    rsvp {
        mtu-signaling;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Configure MTU options for MPLS paths, including packet fragmentation and MTU signaling.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring MTU Signaling in RSVP

path-mtu-discovery (Tunnel)

IN THIS SECTION

- Syntax | 12258
- Hierarchy Level | 12258
- Description | 12258
- Default | 12258
- Required Privilege Level | 12258

Release Information | 12259

Syntax

```
(path-mtu-discovery | no-path-mtu-discovery);
```

Hierarchy Level

```
[edit interfaces ip-fpc/pic/port unit logical-unit-number tunnel]
[edit interfaces gr-fpc/pic/port unit logical-unit-number tunnel]
```

Description

Configure path MTU discovery for outgoing tunnel connections:

- `path-mtu-discovery`—Path MTU discovery is enabled.
- `no-path-mtu-discovery`—Path MTU discovery is disabled.

> **NOTE**: Starting in Junos OS Release 17.2R1, the `no-path-mtu-discovery` configuration statement in the `[edit interfaces ip-fpc/pic/port unit logical-unit-number tunnel]` and `[edit interfaces gr-fpc/pic/port unit logical-unit-number tunnel]` hierarchies is no longer available for `ipip6` tunnels.

Default

Path MTU discovery is enabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configure Path MTU Discovery

path (Protocols MPLS)

IN THIS SECTION

- Syntax | 12259
- Hierarchy Level | 12259
- Description | 12260
- Options | 12260
- Required Privilege Level | 12261
- Release Information | 12261

Syntax

```
path path-name {
    abstract-hop-name (abstract | loose | loose-link | strict);
    (address | hostname) <strict | loose>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Create a named path and optionally specify the sequence of explicit routers that form the path.

You must include this statement when configuring explicit LSPs.

Options

`abstract-hop-name`—Name of the predefined abstract hop. The abstract hop can be used in combination with real IP next hops. An abstract hop is traversed by traversing the member nodes. This traversal can be done by either links that satisfy the logical combination of defined constituent attributes, or by any kind of link. This choice is controlled by the use of abstract hop qualifiers – abstract, loose, loose-link, and strict.

`abstract`—Indicate that the next hop configured in the `path` statement is an abstract hop..

`loose-link`—Indicate that the next hop in the `path` statement is a loose-link abstract hop. This means that the LSP cannot traverse other routers before reaching this router. In other words, the abstract hop of type loose-link is processed only if any of the viable routers is reached in constraint through a link of associated abstract hop membership.

`loose`—Indicate that the next hop in the `path` statement is a loose abstract hop. The path can traverse any real nodes that do not have abstract hop membership, before reaching a node with abstract hop membership, which is a feasible starting point for processing the next abstract hop.

`strict`—Indicate that the next hop in the `path` statement is a strict abstract hop. After the last processed hop in the constraint list, the path can traverse any real nodes that do not have abstract hop membership, before reaching a node with abstract hop membership, which is a feasible starting point for processing the next abstract hop.

address—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is `loose`. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.

- **Default:** If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.

hostname—See *address*.

- **Default:** If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.

`loose`—(Optional) Indicate that the next address in the `path` statement is a loose link. This means that the LSP can traverse through other routers before reaching this router.

- **Default:** `strict`

path-name—Name that identifies the sequence of nodes that form an LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router.

`strict`—(Optional) Indicate that the LSP must go to the next address specified in the `path` statement without traversing other nodes. This is the default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`abstract-hop-name` option introduced in Junos OS Release 17.1 for all platforms.

RELATED DOCUMENTATION

Creating Named Paths

abstract-hop | **1625**

path (Protocols RSVP)

IN THIS SECTION

- Syntax | **12262**
- Hierarchy Level | **12262**
- Description | **12262**
- Default | **12262**
- Options | **12262**
- Required Privilege Level | **12263**
- Release Information | **12263**

Syntax

```
path address <strict | loose>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection bypass bypass-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Description

Configure an explicit path (a sequence of strict or loose routes) to control where and how a bypass LSP is established. If multiple bypasses are configured, they all will use the same explicit path.

Default

No path is configured. CSPF automatically calculates the path the bypass LSP takes.

Options

address—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is `loose`. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.

- **Default:** If you do not specify any routers explicitly, no routing limitations are imposed on the bypass LSP.

`loose`—(Optional) The next address in the `path` statement is loose. The LSP can traverse other routers before reaching this router.

- **Default:** `strict`

`strict`—(Optional) The LSP must go to the next address specified in the `path` statement without traversing other nodes. This is the default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an Explicit Path for Bypass LSPs

path-rate-limit

IN THIS SECTION

- Syntax | 12263
- Hierarchy Level | 12265
- Description | 12265
- Options | 12265
- Required Privilege Level | 12265
- Release Information | 12265

Syntax

```
path-rate-limit {
    message-type {
        create-req {
            alarm-threshold (Security GPRS) {
                forward number;
                reverse number;
            }
```

```
            drop-threshold (Security GPRS) {
                forward number;
                reverse number;
            }
        }
        delete-req {
            alarm-threshold (Security GPRS) {
                forward number;
                reverse number;
            }
            drop-threshold (Security GPRS) {
                forward number;
                reverse number;
            }
        }
        echo-req {
            alarm-threshold (Security GPRS) {
                forward number;
                reverse number;
            }
            drop-threshold (Security GPRS) {
                forward number;
                reverse number;
            }
        }
        other {
            alarm-threshold (Security GPRS) {
                forward number;
                reverse number;
            }
            drop-threshold (Security GPRS) {
                forward number;
                reverse number;
            }
        }
    }
}
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

Use the path-rate-limit function to control specific GPRS tunneling protocol (GTP) messages in both the forward and reverse directions. A drop threshold and an alarm threshold parameter can be configured for each control message in the forward and reverse direction for one path.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name]` hierarchy level is replaced by `[edit security gtp profile profile-name]`.

Options

`message-type`—To specific group of control messages

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

> *Understanding Path Rate Limiting for GTP Control Messages*
>
> *Example: Limiting the Message Rate and Path Rate for GTP Control Messages*

path-selection

IN THIS SECTION

- Syntax | 12266
- Hierarchy Level | 12266
- Description | 12266
- Default | 12267
- Options | 12267
- Required Privilege Level | 12268
- Release Information | 12269

Syntax

```
path-selection {
    (always-compare-med | cisco-non-deterministic | external-router-id);
    as-path-ignore;
     l2vpn-use-bgp-rules;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit protocols bgp],
[edit routing-instances routing-instance-name protocols bgp]
```

Description

Configure BGP path selection.

Default

If the `path-selection` statement is not included in the configuration, only the multiple exit discriminators (MEDs) of routes that have the same peer ASs are compared.

Options

`always-compare-med`

Always compare MEDs whether or not the peer ASs of the compared routes are the same.

> **NOTE**: We recommend that you configure the `always-compare-med` option.

`as-path-ignore`

In the best-path algorithm, skip the step that compares the autonomous system (AS) path lengths. By default, the best-path algorithm evaluates the length of the AS paths and prefers the route with the shortest AS path length.

> **NOTE**: Starting with Junos OS Release 14.1R8, 14.2R7, 15.1R4, 15.1F6, and 16.1R1, the **as-path-ignore** option is supported for routing instances.

`cisco-non-deterministic`

Emulate the Cisco IOS default behavior. This mode evaluates routes in the order that they are received and does not group them according to their neighboring AS. With `cisco-non-deterministic` mode, the active path is always first. All inactive, but eligible, paths follow the active path and are maintained in the order in which they were received, with the most recent path first. Ineligible paths remain at the end of the list.

As an example, suppose you have three path advertisements for the 192.168.1.0 /24 route:

- Path 1—learned through EBGP; AS Path of 65010; MED of 200
- Path 2—learned through IBGP; AS Path of 65020; MED of 150; IGP cost of 5
- Path 3—learned through IBGP; AS Path of 65010; MED of 100; IGP cost of 10

These advertisements are received in quick succession, within a second, in the order listed. Path 3 is received most recently, so the routing device compares it against path 2, the next most recent advertisement. The cost to the IBGP peer is better for path 2, so the routing device eliminates path 3 from contention. When comparing paths 1 and 2, the routing device prefers path 1 because it is received from an EBGP peer. This allows the routing device to install path 1 as the active path for the route.

> **NOTE**: We do not recommend using this configuration option in your network. It is provided solely for interoperability to allow all routing devices in the network to make consistent route selections.

`external-router-id`
: Compare the router ID between external BGP paths to determine the active path.

`igp-multiplier` *number*
: The multiplier value for the IGP cost to a next-hop address. This option is useful for making the MED and IGP cost comparable.

- **Range:** 1 through 1000
- **Default:** 1

`l2vpn-use-bgp-rules`
: Enable routers to use both the BGP path selection algorithm and the designated forwarder path selection algorithm when selecting the preferred path to each destination in a Layer 2 VPN or VPLS routing instance. The BGP path selection algorithm is used by all of the Provider routers participating in the routing instance. The designated forwarder path selection algorithm is used by the PE router participating in the routing instance.

- **Default:** By default, the designated forwarder path selection algorithm is used to select the best path to reach each destination within Layer 2 VPN and VPLS routing instances.

`med-multiplier` *number*
: The multiplier value for the MED calculation. This option is useful for making the MED and IGP cost comparable.

- **Range:** 1 through 1000
- **Default:** 1

`med-plus-igp`
: Add the IGP cost to the indirect next-hop destination to the MED before comparing MED values for path selection. This statement only affects best-path selection. It does not affect the advertised MED.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`med-plus-igp` option introduced in Junos OS Release 8.1.

`as-path-ignore` and `l2vpn-use-bgp-rules` options introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding BGP Path Selection

Enabling BGP Path Selection for Layer 2 VPNs and VPLS

route-distinguisher | **15856**

Example: Ignoring the AS Path Attribute When Selecting the Best Path

path (Steering)

IN THIS SECTION

- Syntax | **12269**
- Hierarchy Level | **12270**
- Description | **12270**
- Required Privilege Level | **12270**
- Release Information | **12270**

Syntax

```
path {
    ipv4-address ipv4-address;
    ipv6-address ipv6-address;
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name steering],
[edit services pcef pcc-action-profiles profile-name]
```

Description

Specify the IP address of a third-party server to which the PCC action profile steers HTTP traffic for applying services. If you configure this, the PCC action profile can only be used in PCC rules that match only HTTP-based applications and all flows.

The remaining statements are explained separately.

Required Privilege Level

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Understanding Application-Aware Policy Control for Subscriber Management

Configuring Policy and Charging Control Action Profiles for Subscriber Management

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

pattern (Application Identification)

IN THIS SECTION

- Syntax | 12271
- Hierarchy Level | 12271
- Description | 12271
- Options | 12271
- Required Privilege Level | 12272
- Release Information | 12272

Syntax

```
pattern pattern;
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type
signature l4-l7-signature-name member member-name]
```

Description

Define an attack pattern to be detected.

Options

pattern User-defined pattern of attack to match, using a regular expression.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

pattern (Class Attribute)

IN THIS SECTION

- Syntax | 12272
- Hierarchy Level | 12273
- Description | 12273
- Options | 12273
- Required Privilege Level | 12273
- Release Information | 12273

Syntax

```
pattern "pattern";
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name subscription-id use-class]
```

Description

Configure characters to insert between the resulting regular expression groups that are generated from parsing the Class attribute contents of the accounting request from the BNG, PGW, or GGSN. Regular expression groups are identified with **\n** for a group number.

Options

pattern Characters to insert between regular expression groups. A regular expression group number "n" is identified as **\n**.

For example, the following configuration generates " `000118191129|ALICE:DRAV3:`" out of "`000118191129#000118191129#ALICE:DRAV3:#7168#nflat#ADSL##`":

```
[edit unified-edge gateways tdf TDF1 domains domain1 subscription-id]
user@host# set use-class regex "[^#]*#\([^#]*\)\#\([^#]*\)"
user@host# set use-class pattern "\1|\2"
```

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

pattern-pcre (Security IDP)

IN THIS SECTION

- Syntax | 12274
- Hierarchy Level | 12274
- Description | 12274
- Options | 12274
- Required Privilege Level | 12274
- Release Information | 12275

Syntax

```
pattern-pcre  signature-pattern-pcre;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature]
```

Description

Specify the pattern in standard PCRE format. You construct the attack pattern in PCRE format just as you would when creating a new signature attack object. This is an optional field. The pattern field is unused under this configuration.

Options

signature-pattern-pcre —Specify the signature pattern in standard PCRE format.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1x49-D40.

pattern (Security IDP)

IN THIS SECTION

- Syntax | 12275
- Hierarchy Level | 12275
- Description | 12275
- Options | 12275
- Required Privilege Level | 12276
- Release Information | 12276

Syntax

```
pattern signature-pattern;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature]
```

Description

Specify the pattern IDP should match. You construct the attack pattern just as you would when creating a new signature attack object.

Options

signature-pattern—Specify the signature pattern.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

pattern-update (Security Antivirus)

IN THIS SECTION

- Syntax | 12276
- Hierarchy Level | 12277
- Description | 12277
- Required Privilege Level | 12277
- Release Information | 12277

Syntax

```
pattern-update {
    email-notify {
        admin-email email-address;
        custom-message message;
        custom-message-subject message-subject;
    }
    interval value;
    no-autoupdate;
    proxy-profile proxy profile name;
    routing-instance name;
    start-time start-time;
    url url;
}
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus juniper-express-engine]
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
[edit security utm default-configuration anti-virus avira-engine]
```

Description

Updates to the pattern file are added as new viruses are discovered. You can configure the security device to regularly update the pattern file automatically, or you can update the file manually.

Starting in Junos OS Release 23.1R1, we've deprecated the `pattern-update` statement at the [`edit security utm feature-profile anti-virus sophos-engine`] and [`edit security utm default-configuration anti-virus sophos-engine`] hierarchy levels.

Required Privilege Level

security— To view this statement in the configuration.

security-control— To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The [`edit security utm default-configuration`] hierarchy level is introduced in Junos OS Release 18.2R1.

The `proxy-profile` option is introduced under the `security utm feature-profile anti-virus sophos-engine` hierarchy level in Junos OS Release 18.3R1.

Support for Avira engine added in Junos OS Release 18.4R1.

pause-computation-during-churn

IN THIS SECTION

- Syntax | 12278
- Hierarchy Level | 12278

- Description | 12278
- Required Privilege Level | 12278
- Release Information | 12278

Syntax

```
pause-computation-during-churn;
```

Hierarchy Level

```
[edit fabric protocols bgp multipath],
[edit logical-systems name protocols bgp multipath],
[edit logical-systems name routing-instances name protocols bgp multipath],
[edit logical-systems name tenants name routing-instances name protocols bgp multipath],
[edit protocols bgp multipath],
[edit routing-instances name protocols bgp multipath],
[edit tenants name routing-instances name protocols bgp multipath]
```

Description

Defer multipath computation for all families during a BGP peering churn. In a very large-scale network, there is a temporary spike in multipath computation during BGP peering, which takes a toll on the Packet Forwarding Engine resources. You can pause the multipath computation and resume after the peering churn settles down. Note that if there is no BGP peering churn, then multipath computation is not paused.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.3R1.

payload

IN THIS SECTION

- Syntax | 12279
- Hierarchy Level | 12279
- Description | 12279
- Options | 12279
- Required Privilege Level | 12280
- Release Information | 12280

Syntax

```
payload {
    ip {
        layer-3;
        layer-4;
    }
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number hash-key family multiservice]
```

Description

(MX Series 5G Universal Routing Platforms only) Include payload data in a hash key for the **multiservice** protocol family when configuring PIC-level symmetrical load balancing on an 802.3ad Link Aggregation Group.

Options

ip—Include IPv4 payload data in the hash key. This option has the following suboptions:

- **layer-3**—Include Layer 3 IP information in the hash key.
- **layer-4**—Include Layer 4 IP information in the hash key.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

> *Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers*

pba-interim-logging-interval

IN THIS SECTION

- Syntax | 12280
- Hierarchy Level | 12281
- Description | 12281
- Options | 12281
- Required Privilege Level | 12281
- Release Information | 12281

Syntax

```
pba-interim-logging-interval seconds
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Description

Port block allocation (PBA) generates one syslog entry per set of ports allocated to a subscriber. These logs are UDP based and can be lost in the network,especially for long running flows. Interim logging resends the above logs at a configured interval for all active blocks that have traffic on at least one block. For the MS-MIC and MS-MPC, log messages are generated for sessions which have a port in a block, even if the block has no traffic.

Options

seconds Interval, in seconds, for re-sending of session logs

- **Default:** 0—This indicates that interim logging is not used.
- **Range:** 1800 to 86,400seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Configuring NAT Session Logs

pbb-evpn-core

IN THIS SECTION

- Syntax | 12282
- Hierarchy Level | 12282
- Description | 12282
- Required Privilege Level | 12282
- Release Information | 12282

Syntax

```
pbb-evpn-core;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols]
[edit routing-instances routing-instance-name protocol evpn]
```

Description

Specify that Ethernet VPN (EVPN) is running for Provider Backbone Bridging (PBB).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

evpn | 6541

pbb-service-options

IN THIS SECTION

- Syntax | 12283
- Hierarchy Level | 12283
- Description | 12283
- Options | 12284
- Required Privilege Level | 12284
- Release Information | 12284

Syntax

```
pbb-service-options {
    default-isid isid-number;
    isid isid-number  vlan-id-list [ vlan-ids ];
    mac-address mac-address;
    source-bmac <mac-address>
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name service-groups service-group-name]
```

Description

For IEEE 802.1ah provider backbone bridge (PBB) configurations, configure PBB service options for the customer routing instance (I-component) service group.

NOTE: The `mac-address` statement is not supported for PBB-EVPN.

The remaining statements are explained separately. See CLI Explorer.

Options

service-group-name—Name of a service group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.0.

RELATED DOCUMENTATION

Example: Configuring E-LINE and E-LAN Services for a PBB Network on MX Series Routers

pcc-action-profile (PCC Rules)

IN THIS SECTION

- Syntax | 12285
- Hierarchy Level | 12285
- Description | 12285
- Options | 12285
- Required Privilege Level | 12285
- Release Information | 12286

Syntax

```
pcc-action-profile profile-name;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rules-name then],
[edit services pcef pcc-rules rules-name then]
```

Description

Specify the name of the action profile to include in a policy and charging control (PCC) rule configuration. The action profile defines the treatment to be applied to specific service data flows or to packets associated with specific applications.

If you are using Junos OS Subscriber Aware, specify the name of the action profile at the `[edit unified-edge pcef pcc-rules rules-name then]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the action profile at the `[edit services pcef pcc-rules rules-name then]` hierarchy level.

Options

profile-name Name of the PCC action profile that the PCC rule references. The referenced action profile must be configured.

- **Range:** 1 through 63 characters.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-rules rules-name then]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

- *Configuring Policy and Charging Control Rules*
- Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware
- Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment
- *Configuring Policy and Charging Control Action Profiles for Subscriber Management*
- *Understanding Application-Aware Policy Control for Subscriber Management*

pcc-action-profiles

IN THIS SECTION

- Syntax | **12286**
- Hierarchy Level | **12287**
- Description | **12287**
- Options | **12288**
- Required Privilege Level | **12288**
- Release Information | **12288**

Syntax

```
pcc-action-profiles profile-name {
    forwarding-class class-name;
    gate-status (uplink | downlink | uplink-downlink | disable-both);
```

```
    hcm-profile hcm-profile-name;
    logging-rule lrf-rule-name;
    maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value;
    monitoring-key key_string;
    redirect {
        url url-name;
    }
    steering {
        keep-existing-steering;
        path {
            ipv4-address ipv4-address;
            ipv6-address ipv6-address;
        }
        routing-instance {
            downlink downlink-vrf-name;
            uplink uplink-vrf-name;
        }
    }
}
```

Hierarchy Level

```
[edit unified-edge pcef],
[edit services pcef]
```

Description

Configure a PCC action profile. A PCC action profile defines the treatment to be applied to specific service data flows or to packets associated with specific applications. A PCC action profile is specified in the `then` clause of a PCC rule.

If you are using Junos OS Subscriber Aware, configure the PCC action profile at the `[edit unified-edge pcef]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the PCC action profile at the `[edit services pcef]` hierarchy level. The following options are not applicable to subscriber management:

- `hcm-profile`
- `monitoring-key`

Options

profile-name Name of the PCC action profile.

- **Range:** 1 through 63 characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Understanding Application-Aware Policy Control for Subscriber Management

Configuring Policy and Charging Control Action Profiles for Subscriber Management

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Configuring TDF Subscriber Usage Monitoring for Traffic That Matches Predefined PCC Rules

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

pcc-context

IN THIS SECTION

- Syntax | 12289
- Hierarchy Level | 12289
- Description | 12289
- Options | 12290
- Required Privilege Level | 12290
- Release Information | 12290

Syntax

```
pcc-context {
    input-service-filter-name filter-name;
    input-service-set-name service-set-name;
    ipv6-input-service-filter-name filter-name;
    ipv6-input-service-set-name service-set-name;
    ipv6-output-service-filter-name filter-name;
    ipv6-output-service-set-name service-set-name;
    output-service-filter-name filter-name;
    output-service-set-name service-set-name;
    profile-name pcef-profile-name;
}
```

Hierarchy Level

```
[edit access profile profile-name session-options]
```

Description

Specify the PCEF profile that contains the policy and charging control (PCC) rules that a PCRF can directly activate for a subscriber, and specify the input and output service sets that process the PCC rules. You can optionally specify service filters for the service sets. PCC rules specify the policies to

apply to traffic based on the application being used by the subscriber or the Layer 3 and Layer 4 service data flow information.

Options

input-service-filter-name *filter-name*	(Optional) Input IPv4 service filter name. The service filter identifies conditions for which you want to skip application-aware policy control.
input-service-set-name *service-set-name*	Input IPv4 service-set name. Use a service set that is enabled for application-aware policy control.
ipv6-input-service-filter-name *filter-name*	(Optional) Input IPv6 service filter name. The service filter identifies conditions for which you want to skip application-aware policy control.
ipv6-input-service-set-name *service-set-name*	Input IPv6 service set name. Use a service set that is enabled for application-aware policy control.
ipv6-output-service-filter-name *filter-name*	(Optional) Output IPv6 service filter name. The service filter identifies conditions for which you want to skip application-aware policy control.
ipv6-output-service-set-name *service-set-name*	Output IPv6 service set name. Use a service set that is enabled for application-aware policy control.
output-service-filter-name *filter-name*	(Optional) Output IPv4 service filter name. The service filter identifies conditions for which you want to skip application-aware policy control.
output-service-set-name *service-set-name*	Output IPv4 service-set name. Use a service set that is enabled for application-aware policy control.
profile-name *pcef-profile-name*	PCEF profile name. The PCEF profile contains PCC rules or PCC rulesets that the PCRF can directly activate. The PCEF profile must be configured under `dynamic-policy-control` at the `[edit services pcef profiles profile-name]` hierarchy level.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1 on MX Series.

RELATED DOCUMENTATION

Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management

pcc-rule

IN THIS SECTION

- Syntax | 12291
- Hierarchy Level | 12291
- Description | 12291
- Options | 12292
- Required Privilege Level | 12292
- Release Information | 12292

Syntax

```
[pcc-rule rule-name  precedence number];
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rule-bases rulebase-name],
[edit services pcef pcc-rule-bases rulebase-name]
```

Description

Specify one or more policy and charging control (PCC) rules and the rules precedence in a PCC rulebase.

If you are using Junos OS Subscriber Aware, configure the PCC rules at the `[edit unified-edge pcef pcc-rule-bases rulebase-name]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the PCC rules at the `[edit services pcef pcc-rule-bases rulebase-name]` hierarchy level.

Options

rule-name Name of the PCC rule. The referenced PCC rule must be configured.

- **Range:** 1 through 63 characters.

number Precedence value assigned to the PCC rule. The precedence assigned must be unique among the configured PCC rules.

- **Range:** 1 through 65,535

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-rule-bases rulebase-name]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

Configuring a Policy and Charging Control Rulebase

Configuring Policy and Charging Control Rules

pcc-rulebases (PCEF)

IN THIS SECTION

- Syntax | 12293
- Hierarchy Level | 12293
- Description | 12293
- Options | 12294
- Required Privilege Level | 12294
- Release Information | 12294

Syntax

```
pcc-rulebases rulebase-name {
    [pcc-rule rule-name precedence number];
}
```

Hierarchy Level

```
[edit unified-edge pcef],
[edit services pcef]
```

Description

Configure a policy and charging control (PCC) rulebase. You can specify from 1 through 4000 rules in a rulebase.

If you are using Junos OS Subscriber Aware, configure the PCC rulebase at the `[edit unified-edge pcef]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the PCC rulebase at the `[edit services pcef]` hierarchy level.

Options

rulebase-name Name of the PCC rulebase.

- **Range:** 1 through 63 characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

Configuring a Policy and Charging Control Rulebase

pcc-rulebases (PCEF Profile)

IN THIS SECTION

Syntax | 12295

- Hierarchy Level | 12295
- Description | 12295
- Options | 12295
- Required Privilege Level | 12296
- Release Information | 12296

Syntax

```
[pcc-rulebases rulebase-name <time-of-day-profile profile-name>];
```

Hierarchy Level

```
[edit unified-edge pcef profiles profile-name aaa-policy-control],
[edit unified-edge pcef profiles profile-name dynamic-policy-control],
[edit unified-edge pcef profiles profile-name static-policy-control],
[edit services pcef profiles profile-name static-policy-control],
[edit services pcef profiles profile-name dynamic-policy-control]
```

Description

Specify a policy and charging control (PCC) rulebase for a policy control profile.

If you are using Junos OS Subscriber Aware, specify the PCC rulebase at the `[edit unified-edge pcef profiles profile-name (aaa-policy-control | dynamic-policy-control | static-policy-control)]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the PCC rulebase at the `[edit services pcef profiles profile-name (static-policy-control | dynamic-policy-control)]` hierarchy level.

Options

rulebase-name	Name of the PCC rulebase. The referenced PCC rulebase must be configured.
time-of-day-profile *profile-name*	(Optional; only applies to rulebases in static PCEF profiles for Junos OS Subscriber Aware) Use the specified name of the time-of-day profile to apply to the PCC rulebase. The referenced profile must already be defined at the `[edit unified-edge pcef]` hierarchy

level. The time-of-day profile specifies the time of day, day of the week, or day of the month to activate or deactivate the PCC rulebase for subscribers assigned to the PCEF profile.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef profiles` *`profile-name`* `static-policy-control]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support at the `[edit services pcef profiles` *`profile-name`* `dynamic-policy-control]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 18.2R1 on MX Series.

RELATED DOCUMENTATION

- Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Dynamic Policies
- Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Static Policies
- Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Policies That a RADIUS Server Controls
- Configuring Static Time-of-Day PCC Rule Activation and Deactivation in a Junos OS Subscriber Aware PCEF Profile
- *Configuring a Policy and Charging Control Rulebase*
- *Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management*

pcc-rules (PCEF)

IN THIS SECTION

- Syntax | 12297
- Hierarchy Level | 12297
- Description | 12297
- Options | 12298
- Required Privilege Level | 12298
- Release Information | 12298

Syntax

```
pcc-rules rule-name {
    from {
        <application-groups [application-group-name]>;
        <applications [application-name]>;
        flows ([flow-identifier | any)];
    }
    then {
        pcc-action-profile  profile-name;
    }
}
```

Hierarchy Level

```
[edit unified-edge pcef],
[edit services pcef]
```

Description

Configure the PCC rules. A PCC rule identifies the subscriber IP packets that are associated with a service data flow (SDF) or application and defines the treatment to be applied to the packets.

If you are using Junos OS Subscriber Aware, configure the PCC rule at the `[edit unified-edge pcef]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the PCC rule at the `[edit services pcef]` hierarchy level.

Options

rule-name Name of the PCC rule.

- **Range:** 1 through 63 characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Rules

Configuring TDF Subscriber Usage Monitoring for Traffic That Matches Predefined PCC Rules

pcc-rules (PCEF Profile)

IN THIS SECTION

- Syntax | 12299
- Hierarchy Level | 12299
- Description | 12299
- Options | 12300
- Required Privilege Level | 12300
- Release Information | 12300

Syntax

```
pcc-rules [rule-name precedence number <time-of-day-profile profile-name>];
```

Hierarchy Level

```
[edit unified-edge pcef profiles profile-name aaa-policy-control],
[edit unified-edge pcef profiles profile-name dynamic-policy-control],
[edit unified-edge pcef profiles profile-name static-policy-control],
[edit services pcef profiles profile-name static-policy-control],
[edit services pcef profiles profile-name dynamic-policy-control]
```

Description

Specify the policy and charging control (PCC) rules for a policy and charging enforcement function (PCEF) profile and assign a precedence to each PCC rule. You can configure up to 32 PCC rules in a PCEF profile.

If you are using Junos OS Subscriber Aware, specify the PCC rules at the `[edit unified-edge pcef profiles profile-name (aaa-policy-control | dynamic-policy-control | static-policy-control)]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the PCC rules at the `[edit services pcef profiles profile-name (static-policy-control | dynamic-policy-control)]` hierarchy level.

Options

rule-name — Name of the PCC rule. The referenced PCC rule must be configured.

precedence *number* — Use the specified precedence value assigned to a PCC rule. A lower precedence value indicates a higher precedence.

- **Range:** 1 through 65,535

time-of-day-profile *profile-name* — (Optional; only applies to rules in static PCEF profiles for Junos OS Subscriber Aware) Use the specified name of the time-of-day profile to apply to the PCC rule. The referenced profile must already be defined at the `[edit unified-edge pcef]` hierarchy level. The time-of-day profile specifies the time of day, day of the week, or day of the month to activate or deactivate the PCC rule for subscribers assigned to the PCEF profile.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef profiles` *`profile-name`* `static-policy-control]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support at the `[edit services pcef profiles` *`profile-name`* `dynamic-policy-control]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 18.2R1 on MX Series.

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Dynamic Policies

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Static Policies

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Policies That a RADIUS Server Controls

Configuring Static Time-of-Day PCC Rule Activation and Deactivation in a Junos OS Subscriber Aware PCEF Profile

Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management

Configuring Policy and Charging Control Rules

pcc-time-of-day-profiles

IN THIS SECTION

- Syntax | **12301**
- Hierarchy Level | **12302**
- Description | **12302**
- Options | **12302**
- Required Privilege Level | **12302**
- Release Information | **12302**

Syntax

```
pcc-time-of-day-profiles profile-name {
    rule-activation-time {
        <day-of-week | day-of-month month>;
        <hour:min>;
    }
    rule-deactivation-time {
        <day-of-week | day-of-month month>;
        <hour:min>;
    }
}
```

Hierarchy Level

```
[edit unified-edge pcef]
```

Description

Configure a PCC time-of-day profile to specify the time of day, day of the week, or day of the month to activate and deactivate a PCC rule or rulebase. A PCC time-of-day profile is applied to a PCC rule or PCC rulebase within a static PCEF profile. If a time zone is configured on the router, the time-of-day settings apply to the configured time zone.

Options

profile-name Name of the PCC time-of-day profile.

- **Range:** 1 through 63 characters.

The remaining statements are explained separately.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> *Configuring Static Time-of-Day PCC Rule Activation and Deactivation in a Junos OS Subscriber Aware PCEF Profile*
>
> *Understanding Static Time-of-Day PCC Rule Activation and Deactivation*

pce

IN THIS SECTION

- Syntax | 12303
- Hierarchy Level | 12304
- Description | 12304
- Options | 12304
- Required Privilege Level | 12305
- Release Information | 12305

Syntax

```
pce pce-id {
    authentication-key key;
    authentication-key-chain key-chain;
    delegation-cleanup-timeout seconds;
    delegation-priority priority-number;
    destination-ipv4-address ipv4-address;
    destination-port port-number;
    local-address ip-address;
    lsp-cleanup-timer seconds;
    lsp-provisioning;
    lsp-retry-delegation;
     lsp-retry-delegation-timer seconds;
     max-sid-depth max-sid-depth;
    max-unknown-messages messages-per-minute;
    max-unknown-requests requests-per-minute;
    p2mp-lsp-init-capability;
    p2mp-lsp-report-capability;
    p2mp-lsp-update-capability;
    pce-group pce-group-name;
    pce_traffic_steering;
    pce-type ;
    request-timer seconds;
    request-priority priority;
    spring-capability;
```

```
    srv6-capability;
traceoptions ;
```

Hierarchy Level

```
[edit protocols pcep]
```

Description

Configure per-Path Computation Element (PCE) parameters.

Options

pce-id	IP address of the PCE.
authentication-key *key*	(Optional) Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”). It is recommended to define and bind an authentication key for securing a PCEP session, as opposed to binding an authentication keychain.
authentication-key-chain *key-chain*	(Optional) Authentication keychain password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).
local-address *ip-address*	(Optional) IP address of the local end of the PCEP session, the PCC.
lsp-retry-delegation	(Optional) Enable retry LSP delegation process.
lsp-retry-delegation-timer	(Optional) Specify the amount of time (in seconds) that the Path Computation Client (PCC) must wait before retrying delegation of Path Computation Element (PCE)-initiated LSPs in case of delegation failure or re-delegation. • **Default:** 3600 seconds • **Range:** 0 through 4294967294 seconds
max-sid-depth	(Optional) Specify the maximum value for service identifier (SID) depth.

- **Default:** 5
- **Range:** 1 through 5

p2mp-lsp-init-capability (Optional) Capability to provision point-to-multipoint RSVP-TE LSPs by a PCE. By default, this capability is not supported on a PCC, and should be explicitly configured to enable PCE-initated point-to-multipoint LSPs.

p2mp-lsp-report-capability (Optional) Capability to report point-to-multipoint RSVP-TE LSPs to a PCE. By default, this capability is not supported on a PCC, and should be explicitly configured to enable reporting of point-to-multipoint LSPs to a PCE.

p2mp-lsp-update-capability (Optional) Capability to update point-to-multipoint RSVP-TE LSP parameters by a PCE. By default, this capability is not supported on a PCC, and should be explicitly configured to enable updating of PCE-initated point-to-multipoint LSPs.

pce_traffic_steering (Optional) Configure the flow specification capability (also called traffic steering functionality) for enabling the mapping of PCE-initiated point-to-multipoint LSPs to an MVPN routing-instance.

spring-capability (Optional) Enable SPRING-based provisioning for the PCE.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

`lsp-cleanup-timer`, `lsp-provisioning`, `max-unknown-requests`, `request-timer`, and `request-priority` options introduced in Junos OS Release 13.3.

`authentication-key` *`key`*, `authentication-key-chain` *`key-chain`*, and `p2mp-lsp-report-capability` options introduced in Junos OS Release 16.1.

`max-sid-depth` and `spring-capability` options introduced in Junos OS Release 17.2.

`p2mp-lsp-init-capability` and `p2mp-lsp-update-capability` options introduced in Junos OS Release 18.3R1 on all platforms.

`pce_traffic_steering` option introduced in Junos OS Release 19.4R1 on all platforms.

RELATED DOCUMENTATION

pcep | 12320

PCEP Configuration

pce-group (PCE)

IN THIS SECTION

- Syntax | 12306
- Hierarchy Level | 12306
- Description | 12307
- Options | 12307
- Required Privilege Level | 12307
- Release Information | 12307

Syntax

```
pce-group pce-group-name;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Description

Specify the Path Computation Element (PCE) group to which the configured PCE belongs.

Options

pce-group-name	Name of the PCE group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

pce | 12303

pce-group (Protocols PCEP)

IN THIS SECTION

- Syntax | 12308
- Hierarchy Level | 12308
- Description | 12308
- Required Privilege Level | 12308

Release Information | 12309

Syntax

```
pce-group pce-group-id {
    delegation-cleanup-timeout seconds;
    max-unknown-messages messages-per-minute;
    pce-type {
        active stateful;
        srv6-capability;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag (all | pcep);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit protocols pcep]
```

Description

Configure the Path Computation Element (PCE) group parameters. A maximum of 10 PCE groups can be configured at any given point in time.
The remaining statements are explained separately.

NOTE: A PCE group can include PCEs that are either only stateful or only active stateful. A combination of stateful PCEs and active stateful PCEs in one PCE group is not supported.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

pcep | 12320

pce-type

IN THIS SECTION

- Syntax | 12309
- Hierarchy Level | 12310
- Description | 12310
- Required Privilege Level | 12310
- Release Information | 12310

Syntax

```
pce-type {
    active stateful;
}
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Description

Configure the path computation element (PCE) type:

- active—Uses LSP state information learned from PCCs to optimize path computations, and actively updates LSP parameters in those PCCs that delegated control over their LSPs to the PCE.
- stateful—Uses LSP state information learned from PCCs to optimize path computations, but does not actively update the LSP state. A PCC maintains synchronization with the PCE.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

pce | 12303

pcef

IN THIS SECTION

- Syntax | 12311
- Hierarchy Level | 12313
- Description | 12313
- Required Privilege Level | 12313
- Release Information | 12313

Syntax

```
pcef {
    flow-descriptions flow-identifier {
        direction (uplink | downlink | both);
        local-port-range {
            low low-value high high-value;
        }
        local-ports number;
        no-send-to-ue;
        protocol number;
        remote-address (ipv4-address ipv4-address | ipv6-address ipv6-address);
        remote-port-range {
            low low-value high high-value;
        }
        remote-ports number;
    }
    pcc-action-profiles profile-name {
        forwarding-class class-name;
        gate-status (uplink | downlink | uplink-downlink | disable-both);
        hcm-profile hcm-profile-name;
        logging-rule lrf-rule-name;
        maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value;
        monitoring-key key_string;
        redirect {
            url url-name;
        }
```

```
        steering {
            keep-existing-steering;
            path {
                ipv4-address ipv4-address;
                ipv6-address ipv6-address;
            }
            routing-instance {
                downlink downlink-vrf-name;
                uplink uplink-vrf-name;
            }
        }
    }
    pcc-rulebases rulebase-name {
        [pcc-rule rule-name precedence number];
    }
    pcc-rules rule-name {
        from {
            <application-groups [application-group-name]>;
            <applications [application-name]>;
            flows ([flow-identifier ] | any);
        }
        then {
            pcc-action-profile  profile-name;
        }
    }
    pcc-time-of-day-profiles profile-name {
        rule-activation-time {
            <day-of-week | day-of-month month>;
            <hour:min>;
        }
        rule-deactivation-time {
            <day-of-week | day-of-month month>;
            <hour:min>;
        }
    }
    profiles profile-name {
        aaa-policy-control {
            aaa-profile aaa-profile-name;
            pcc-rulebases [rulebase-name];
            user-password password;
        }
        dynamic-policy-control {
            pcc-rules {
```

```
                [rule-name  number];
            }
            pcc-rulebases {
                [rulebase-name];
            }
            diameter-profile gx-profile-name;
        }
        static-policy-control {
            pcc-rules {
                [rule-name precedence number <time-of-day-profile profile-name>];
            }
            pcc-rulebases {
                [rulebase-name <time-of-day-profile profile-name>];
            }
        }
    }
}
```

Hierarchy Level

```
[edit unified-edge]
```

Description

Set up the overall policy and control enforcement function (PCEF) configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Understanding Junos Subscriber Aware Policy and Charging Enforcement Function (PCEF)

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

pcef (Dynamic Profiles)

IN THIS SECTION

- Syntax | 12314
- Hierarchy Level | 12314
- Description | 12314
- Options | 12315
- Required Privilege Level | 12315
- Release Information | 12315

Syntax

```
pcef pcef-profile-name {
    activate rule-name | activate-all;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number service]
```

Description

Assign a policy and charging enforcement function (PCEF profile) to the dynamic profile. The PCEF profile specifies a set of PCC rules and rulebases to assign to a subscriber for application-aware policy control, and assigns a precedence value to each predefined rule. Also specify which of the PCC rules to activate for the dynamic profile.

Options

activate *rule-name* Name of the PCC rule to activate or the predefined dynamic interface variable `$junos-pcef-rule`. To specify more than one rules, include this line multiple times.

activate-all Activates all of the PCC rules.

pcef-profile-name Name of the PCEF profile or the predefined dynamic interface variable `$junos-pcef-profile`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Configuring Application-Aware Policy Control for Subscriber Management

pcef-profile (Service Set)

IN THIS SECTION

- Syntax | 12316
- Hierarchy Level | 12316
- Description | 12316
- Options | 12316
- Required Privilege Level | 12316
- Release Information | 12316

Syntax

```
pcef-profile pcef-profile-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the dummy PCEF profile that you configured at the `[edit services pcef]` hierarchy level. This profile is a placeholder profile with no configuration options, but it must be specified to enable PCEF functionality on the services plane.

Options

pcef-profile-name	Name of the PCEF profile.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Applying Services to Subscriber-Aware Traffic with a Service Set

Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control

Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management

pcef-profile (TDF Domain)

IN THIS SECTION

- Syntax | 12317
- Hierarchy Level | 12317
- Description | 12317
- Options | 12318
- Required Privilege Level | 12318
- Release Information | 12318

Syntax

```
pcef-profile name;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
```

Description

Specify the policy and charging enforcement function (PCEF) profile to be applied to subscribers in the TDF domain. This is required for IFL-based subscribers, and optional for IP-based subscribers. If you do not identify a PCEF profile, then the PCEF profile must be assigned under the `[edit unified-edge gateways tdf gateway-name domain-selection term]` hierarchy level.

Options

name Name of the PCEF profile.

> **NOTE**: The PCEF profile must have been previously configured at the `[edit unified-edge pcef]` hierarchy level.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain*
- *Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain*
- *Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain*
- *IP-Based and IFL-Based TDF Subscribers Overview*
- *IP-Based Subscriber Setup Overview*
- *Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber*
- *Understanding PCEF Profiles*

pcef-profile (TDF Domain Selection)

IN THIS SECTION

- Syntax | 12319
- Hierarchy Level | 12319

- Description | 12319
- Options | 12319
- Required Privilege Level | 12319
- Release Information | 12320

Syntax

```
pcef-profile pcef-profile-name;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name then]
```

Description

Specify the policy and charging enforcement function (PCEF) profile to be selected for the IP-based TDF subscriber when the criteria specified in the `domain-selection term` *`term-name`* `from` statement are matched. This PCEF profile is required if the TDF domain selected for a subscriber does not specify a PCEF profile or you want to allow different members of the same TDF domain to have different PCEF profiles.

Options

pcef-profile-name Name of the PCEF profile.

> **NOTE**: The PCEF profile must have been previously configured at the `[edit unified-edge pcef]` hierarchy level.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

pcep

IN THIS SECTION

- Syntax | 12320
- Hierarchy Level | 12321
- Description | 12321
- Required Privilege Level | 12321
- Release Information | 12321

Syntax

```
pcep {
    message-rate-limit messages-per-minute;
    pce pce-id;
    pce-group pce-group-id;
    traceoptions;
    update-rate-limit updates-per-minute;
    maximum-srmpls-segment-list-depth maximum-srmpls-segment-list-depth;
    maximum-srv6-segment-list-depth
```

```
    propagate-lsp-max-segment-list-depth
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure the Path Computation Client (PCC) parameters.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for ACX Series added in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

Support of the Path Computation Element Protocol for RSVP-TE Overview

Example: Configuring the Path Computation Element Protocol for MPLS RSVP-TE

Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs

Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Controlled Point-to-Multipoint LSPs

pci-address

IN THIS SECTION

- Syntax | 12322
- Hierarchy Level | 12322
- Description | 12322
- Required Privilege Level | 12322
- Release Information | 12322

Syntax

```
pci-address  pci-address;
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

PCI address for the VNF interfaces.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D50.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

pcp-rules

IN THIS SECTION

- Syntax | 12323
- Hierarchy Level | 12323
- Description | 12323
- Options | 12324
- Required Privilege Level | 12324
- Release Information | 12324

Syntax

```
pcp-rules rule-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the PCP rule to apply to the service set. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC. Starting in Junos OS Release 20.1R1, PCP is also supported for Next Gen Services.

Options

rule-name The PCP rule to apply to the service set.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

pcp-server

IN THIS SECTION

- Syntax | 12324
- Hierarchy Level | 12324
- Description | 12325
- Options | 12325
- Required Privilege Level | 12325
- Release Information | 12325

Syntax

```
pcp-server server-name;
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name then]
```

Description

Specify the PCP server that handles the traffic that matches the PCP rule term.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

server-name	Name of the PCP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

pcrf (Diameter Applications)

IN THIS SECTION

- Syntax | 12326
- Hierarchy Level | 12327
- Description | 12327
- Required Privilege Level | 12327
- Release Information | 12327

Syntax

```
pcrf {
    global {
        rule-param avp-code;
    }
    partition partition-name  {
        destination-host pcrf-hostname;
        destination-realm pcrf-realm-name;
        diameter-instance;
        draining;
        draining-response-timeout seconds;
        ip-can-type number;
        local-decision {
            deny;
            grant;
            reinit-on-failure;
            reinit-on-rar;
            reinit-timeout seconds;
            timeout seconds;
        }
        logout-response-timeout seconds;
        max-outstanding-requests number;
        report-local-rule;
        report-resource-allocation;
        report-successful-resource-allocation;
        send-dyn-subscription-indicator;
        send-network-family-indicator;
        send-origin-state-id;
        subscription-id-data-include  {
            base-interface-name;
            delimiter delimiter-character;
            domain-name name;
            interface-name;
            mac-address;
            nas-port-id;
            origin-host;
            origin-realm;
            user-name;
            user-prefix prefix;
            vlan-tags;
        }
```

```
        subscription-id-type number;
        update-response-timeout seconds;
        use-session-stamp;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure the Policy and Charging Rules Function (PCRF) global attributes, rules, parameters, and partition to authorize and provision subscribers. PCRF is a centralized policy decision point that deploys business policy and charging rules to allocate broadband network resources and manages flow-based charges for subscribers and services. PCRF pushes the rules down to the Policy and Charging Enforcement Function (PCEF) using the 3GPP Gx protocol and online policy interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- *Configuring the PCRF Partition*
- *Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers*
- *Understanding Gx Interactions Between the Router and the PCRF*
- *Understanding Interactions Between the PCRF, PCEF, and OCS*
- *Understanding Upstream and Downstream Messages for the PCRF*

pdu-interval

IN THIS SECTION

- Syntax | 12328
- Hierarchy Level | 12328
- Description | 12328
- Options | 12328
- Required Privilege Level | 12329
- Release Information | 12329

Syntax

```
pdu-interval interval;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface interface-name]
```

Description

For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the periodic OAM PDU sending interval for fault detection. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options

interval—Periodic OAM PDU sending interval.

- **Range:** For MX, M, T, ACX, Series routers, SRX Series firewalls and EX Series switches – 100 through 1000 milliseconds
- **Default:** 1000 milliseconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring the OAM PDU Interval

Example: Configuring Ethernet OAM Link Fault Management

Configuring Ethernet OAM Link Fault Management

pdu-threshold

IN THIS SECTION

- Syntax | 12329
- Hierarchy Level | 12330
- Description | 12330
- Options | 12330
- Required Privilege Level | 12330
- Release Information | 12330

Syntax

```
pdu-threshold threshold-value;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface interface-name]
```

Description

Configure how many protocol data units (PDUs) are missed before declaring the peer lost in Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the number of OAM PDUs to miss before an error is logged. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options

threshold-value—The number of PDUs missed before declaring the peer lost.

- **Range:** 3 through 10 PDUs
- **Default:** 3 PDUs

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

> *Configuring the OAM PDU Threshold*
>
> *Configuring Ethernet OAM Link Fault Management*

peak-burst-size (QFX)

IN THIS SECTION

- Syntax | 12331
- Hierarchy Level | 12331
- Description | 12331
- Options | 12331
- Required Privilege Level | 12332
- Release Information | 12332

Syntax

```
peak-burst-size bytes;
```

Hierarchy Level

```
[edit firewall three-color-policer policer-name two-rate]
```

Description

Configure the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red).

> **NOTE**: When you include the `peak-burst-size` statement in the configuration, you must also include the **committed-burst-size** and `peak-information-rate` statements at the same hierarchy level.

Options

bytes—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

- **Range:** 1500 bytes through 100,000,000,000 bytes (100 GB)

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Two-Color and Three-Color Policers to Control Traffic Rates

Overview of Policers

peak-burst-size

IN THIS SECTION

- Syntax | 12332
- Hierarchy Level | 12333
- Description | 12333
- Options | 12333
- Required Privilege Level | 12333
- Release Information | 12334

Syntax

```
peak-burst-size bytes;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name two-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

For a two-rate three-color policer, configure the peak burst size (PBS) as a number of bytes. The PBS defines the maximum number of bytes of unused peak bandwidth capacity that can be accumulated. The accumulated bandwidth allows for moderate periods of bursting traffic that exceeds the peak information rate (PIR) and the committed burst size (CBS).

> **NOTE**: When you include the `peak-burst-size` statement in the configuration, you must also include the `committed-burst-size` and `peak-information-rate` statements at the same hierarchy level.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits.

- A traffic flow is categorized green if it conforms to both the committed information rate (CIR) and the CBS-bounded accumulation of available committed bandwidth capacity.
- A traffic flow is categorized yellow if exceeds the CIR and CBS but conforms to the PIR. Packets in a yellow flow are marked with `medium-high` packet loss priority (PLP) and then passed through the interface.
- A traffic flow is categorized red if exceeds the PIR and the PBS-bounded accumulation of available peak bandwidth capacity. Packets in a red traffic flow are marked with `high` PLP and then either passed through the interface or optionally discarded.

Options

bytes—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1500 through 100,000,000,000 bytes

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

Support at the `[edit dynamic-profiles ... two-rate]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Three-Color Policer Configuration Overview
- Policer Bandwidth and Burst-Size Limits
- Policer Color-Marking and Actions
- Dual Token Bucket Algorithms
- Determining Proper Burst Size for Traffic Policers
- committed-burst-size | 4458
- committed-information-rate | 4460
- excess-burst-size | 6580
- peak-information-rate | 12336

peak-information-rate (QFX)

IN THIS SECTION

- Syntax | 12335
- Hierarchy Level | 12335
- Description | 12335
- Options | 12335
- Required Privilege Level | 12335
- Release Information | 12335

Syntax

```
peak-information-rate bits-per-second;
```

Hierarchy Level

```
[edit firewall three-color-policer policer-name two-rate]
```

Description

Configure the maximum achievable rate. Packets that exceed the committed information rate (CIR) but are below the peak information rate (PIR) are marked with medium-high packet loss priority (yellow). Packets that exceed the PIR are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR.

> **NOTE**: When you include the `peak-information-rate` statement in the configuration, you must also include the **committed-information-rate** and `peak-burst-size` statements at the same hierarchy level.

Options

bits-per-second—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

- **Range:** 32,000 bps through 10,000,000,000 bps (10 gbps)

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Two-Color and Three-Color Policers to Control Traffic Rates

Overview of Policers

peak-information-rate

IN THIS SECTION

- Syntax | 12336
- Hierarchy Level | 12336
- Description | 12336
- Options | 12337
- Required Privilege Level | 12337
- Release Information | 12337

Syntax

```
peak-information-rate bps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name two-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

For a two-rate three-color policer, configure the peak information rate (PIR) as a number of bits per second. The PIR is the maximum rate for traffic arriving at or departing from the interface under peak line conditions. Traffic that exceeds the committed information rate (CIR) and the committed burst size (CBS) is metered to the PIR.

> **NOTE**: When you include the `peak-information-rate` statement in the configuration, you must also include the `committed-information-rate` and `peak-burst-size` statements at the same hierarchy level.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits.

- A traffic flow is categorized green if it conforms to both the CIR and the CBS-bounded accumulation of available committed bandwidth capacity.
- A traffic flow is categorized yellow if exceeds the CIR and CBS but conforms to the PIR. Packets in a yellow flow are marked with `medium-high` packet loss priority (PLP) and then passed through the interface.
- A traffic flow is categorized red if exceeds the PIR and the PBS-bounded accumulation of available peak bandwidth capacity. Packets in a red traffic flow are marked with `high` PLP and then either passed through the interface or optionally discarded.

Options

bps—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

Range:

- - 1500 through 100,000,000,000 bps on EX, M, and T Series routers
  - 1500 through 18,446,744,073,709,551,615 bps on Mx Series routers

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

Support at the `[edit dynamic-profiles ... two-rate]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Three-Color Policer Configuration Overview
- Policer Bandwidth and Burst-Size Limits
- Policer Color-Marking and Actions
- Dual Token Bucket Algorithms
- Determining Proper Burst Size for Traffic Policers
- committed-burst-size | 4458
- committed-information-rate | 4460
- excess-burst-size | 6580
- peak-burst-size | 12332

peak-rate

IN THIS SECTION

- Syntax | 12338
- Hierarchy Level | 12338
- Description | 12339
- Options | 12339
- Required Privilege Level | 12339
- Release Information | 12339

Syntax

```
peak-rate rate;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
```

Description

(MX Series routers) Define ATM peak cell rate on ATM MICs in cells per second by entering a decimal number followed by the abbreviation c; where 1 cps = 384 bps.

Options

rate—ATM peak rate in cells per second.

- **Range:** 61 cps through 353,206 cps.
- **Range:** (MX Series with MPCs and ATM MICs with SFP) 61 cps through 1,412,829 cps

> **NOTE**: Beginning with Junos OS Release 14.2, to configure OC12 CBR bandwidth speed per virtual circuit (VC) on an ATM MIC with SFP (MIC-3D-8OC3-2OC12-ATM), specify up to 1,412,829 cps as the ATM peak cell rate.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

show class-of-service traffic-control-profile | 19145

peer

IN THIS SECTION

Syntax | 12340

- Hierarchy Level | 12340
- Description | 12340
- Options | 12340
- Required Privilege Level | 12340
- Release Information | 12341

Syntax

```
peer a.b.c.d {
    interface interface-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name multi-chassis-protection]
```

Description

For MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces, use the `multi-chassis-protection` statement under the physical interface level to reduce the configuration at the logical interface level. If the interchassis control protocol connection (ICCP) is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the `peer` statement. You must also specify the peer's physical interface.

Options

a.b.c.d	Specify the IP address of the peer.
interface *interface-name*	Specify the peer's physical interface: interface *interface-name-fpc/pic/port*

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers and QFX Series Switches
- *Configuring Aggregated Ethernet Link Protection*
- *Example: Configuring Aggregated Ethernet Link Protection*
- multi-chassis-protection (Interfaces) | **11045**

peer-active (VPLS Multihoming for FEC 129)

IN THIS SECTION

- Syntax | **12341**
- Hierarchy Level | **12342**
- Description | **12342**
- Default | **12343**
- Required Privilege Level | **12343**
- Release Information | **12343**

Syntax

```
peer-active;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-homing],
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-homing site site-name],
[edit routing-instances instance-name protocols vpls multi-homing],
[edit routing-instances instance-name protocols vpls multi-homing site site-name]
```

Description

Keep customer edge (CE) interfaces in the up state when all BGP peers go down.

> **NOTE**: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Consider a scenario in which two provider edge (PE) routers are sharing two multihomed sites under one routing instance, with two CE devices, CE1 and CE2.

Figure 2: Two Provider Edge Topology



If the BGP peering session drops between Router PE1 and Router PE2, each one would consider itself to be the designated forwarder (DF) for Device CE1 and Device CE2. This creates a loop through the two CE devices, in which traffic loops from one CE device to the other then back to the first.

Junos OS overcomes this scenario by dropping all multihomed CE interface traffic on all multihoming PE routers when the BGP session drops between the PE routers. This functionality is enabled by default for all sites in a routing instance.

The `peer-active` statement disables the default functionality, so that PE routers keep their multihomed CE interfaces in the up state, even though the BGP peering session is down.

If you configure this statement in the `multi-homing` hierarchy, the default functionality is disabled for all sites. If you configure this statement for a site, the default functionality is disabled only for that particular site.

Default

If you omit this statement, Junos OS drops all multihomed CE interface traffic on all multihoming PE routers when the BGP session drops between the PE routers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring VPLS Multihoming (FEC 129)

peer-as-billing-template

IN THIS SECTION

- Syntax | 12344
- Hierarchy Level | 12344
- Description | 12344

- Required Privilege Level | 12344
- Release Information | 12344

Syntax

```
peer-as-billing-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

Enables the extraction of bandwidth usage information for billing purposes in PIC-based sampling configurations. This capability is supported on routers and applies only to IPv4 and IPv6 traffic.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

> *Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates*

peer-as-list

IN THIS SECTION

- Syntax | 12345
- Hierarchy Level | 12345
- Description | 12345
- Required Privilege Level | 12345
- Release Information | 12346

Syntax

```
peer-as-list peer-as-list;
```

Hierarchy Level

```
[edit fabric protocols bgp group],
[edit logical-systems name protocols bgp group],
[edit logical-systems name routing-instances name protocols bgp group],
[edit logical-systems name tenants name routing-instances name protocols bgp group],
[edit protocols bgp group],
[edit routing-instances name protocols bgp group],
[edit tenants name routing-instances name protocols bgp group]
```

Description

Specify a list of peer autonomous systems for neighbors whose IP addresses are not known. This supports dynamic peering and allows you to control BGP peering when the neighbor's exact IP address is not known.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

as-list | 2487

peer-as (Protocols BGP)

IN THIS SECTION

- Syntax | 12346
- Hierarchy Level | 12346
- Description | 12347
- Options | 12347
- Required Privilege Level | 12348
- Release Information | 12348

Syntax

```
peer-as autonomous-system;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
```

```
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Specify the neighbor (peer) autonomous system (AS) number.

For EBGP, the peer is in another AS, so the AS number you specify in the `peer-as` statement must be different from the local router's AS number, which you specify in the `autonomous-system` statement. For IBGP, the peer is in the same AS, so the two AS numbers that you specify in the **autonomous-system** and `peer-as` statements must be the same.

The AS numeric range in plain-number format has been extended in Junos OS Release 9.1 to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of the Junos OS support 2-byte AS numbers.

In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <*16-bit high-order value in decimal*>.<*16-bit low-order value in decimal*>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.

With the introduction of 4-byte AS numbers, you might have a combination of routers that support 4-byte AS numbers and 2-byte AS numbers. For more information about what happens when establishing BGP peer relationships between 4-byte and 2-byte capable routers, see the following topics:

- *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview.*

Options

autonomous-system—AS number.

- **Range:** 1 through 4,294,967,295 (2^{32} – 1) in plain-number format for 4-byte AS numbers
- **Range:** 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

- **Range:** 0.0 through 65535.65535 in AS-dot notation format for 4-byte AS numbers

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

peer-as (VPLS)

IN THIS SECTION

- Syntax | 12348
- Hierarchy Level | 12349
- Description | 12349
- Options | 12349
- Required Privilege Level | 12349
- Release Information | 12349

Syntax

```
peer-as {
    all;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]
```

Description

Enable the autonomous system border router (ASBR) to establish a single pseudowire to each of the other ASBRs interconnected using inter-AS VPLS with MAC processing at the ASBR.

NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Options

`all`—This option is required. All peer routers, the ASBRs, are placed within the same VPLS mesh group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Inter-AS VPLS with MAC Processing at the ASBR

peer-certificate-type

IN THIS SECTION

- Syntax | 12350
- Hierarchy Level | 12350
- Description | 12350
- Options | 12350
- Required Privilege Level | 12351
- Release Information | 12351

Syntax

```
peer-certificate-type (pkcs7 | x509-signature);
```

Hierarchy Level

```
[edit services ipsec-vpn ike policy policy-name]
```

Description

(MX Series routers only) Specify a preferred type of certificate (PKCS7 or X509). By default, X509 encoding format is used. With the flexibility to configure the encoding format in which certificate requests are sent to the peer, you can determine the type of certificate to be used depending on the type supported by the peer. For example, if the peer does not support PKCS7, certificate authentication cannot occur unless you configure the same type on MX Series routers as the initiator or sender.

Options

- `pkcs7`—Public-Key Cryptography Standard #7.
- `x509-signature`—X509 is an ITU-T standard for public key infrastructure.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 15.1.

RELATED DOCUMENTATION

Configuring IKE Policies

peer (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 12351
- Hierarchy Level | 12352
- Description | 12352
- Options | 12352
- Required Privilege Level | 12352
- Release Information | 12352

Syntax

```
peer peer-name {
    address ip-address;
    connect-actively {
        <capabilities-exchange-timeout seconds>;
        <port port-number>;
        <repeat-timeout seconds>;
        <retry-timeout seconds>;
```

```
        <timeout seconds>;
        transport transport-name;
    }
    <disconnect-peer-timeout seconds>;
    <incoming-queue> {
        size size;
    }
    <outgoing-queue> {
        <high-watermark high-watermark>;
        <low-watermark low-watermark>;
        size size;
    }
    <watchdog-timeout seconds>;
}
```

Hierarchy Level

```
[edit access diameter]
```

Description

Configure a remote peer for the Diameter instance. You can configure up to 31 peers.

Options

peer-name—Name of the peer.

- **Range:** 1 through 32 alphanumeric characters

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

peer (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 12353
- Hierarchy Level | 12353
- Description | 12354
- Options | 12354
- Required Privilege Level | 12354
- Release Information | 12354

Syntax

```
peer  peer-name {
    address ip-address;
    connect-actively {
        port port-number;
        transport transport-name;
    }
    logical-system logical-system-name  <routing-instance routing-instance-name>;
    peer-origin realm realm-name host hostname;
    routing-instance routing-instance-name;
    send-origin-state-id;
}
```

Hierarchy Level

```
[edit diameter]
```

Description

Configure peer-specific Origin-Realm and Origin-Name AVP values to include in the CER, DWR, DWA, DPR, and DPA messages of this peer.

Options

peer-name—Name of the peer.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Peers

peer (Diameter Network Element)

IN THIS SECTION

- Syntax | 12355
- Hierarchy Level | 12355
- Description | 12355
- Options | 12355
- Required Privilege Level | 12355

Release Information | 12356

Syntax

```
peer peer-name {
    priority priority-value;
    <timeout seconds>;
}
```

Hierarchy Level

```
[edit access diameter network-element element-name]
```

Description

Define and prioritize a peer associated with a Diameter network element. You must prioritize the associated peer by including the `priority` statement.

Options

peer-name—Name of the peer.

- **Range:** 1 through 32 alphanumeric characters

> **NOTE**: The specified peer must already be configured at the `[edit access diameter peer]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

peer (Diameter Network Element)

IN THIS SECTION

- Syntax | 12356
- Hierarchy Level | 12356
- Description | 12357
- Options | 12357
- Required Privilege Level | 12357
- Release Information | 12357

Syntax

```
peer peer-name {
    priority priority-value;
}
```

Hierarchy Level

```
[edit diameter network-element element-name]
```

Description

Assigns a peer to a Diameter network element and sets a priority value for that peer within the DNE.

Options

peer-name—Name of the peer.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Network Elements

peer-forwarding-ip

IN THIS SECTION

- Syntax | 12358
- Hierarchy Level | 12358
- Description | 12358
- Options | 12358
- Required Privilege Level | 12359
- Release Information | 12359

Syntax

```
peer-forwarding-ip loopback-address {
        interface interface-name;
        liveness-detection {
            no-adaptation;
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier multiplier;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
        version (0 | 1 | automatic);
  {
```

Hierarchy Level

```
[edit chassis high-availability peer-id]
```


```
[edit chassis high-availability peer-id]
```


Description

Configure peer node for managing asymmetric traffic handling.

Options

IP-address	Specify peer forwarding IP. This forwarding address used to forward data packet over ICD link.
Interface	Specify interface for forwarding data packets for ICD.

"liveness-detection" on page 9848 Specify BFD failure detection time.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Multinode High Availability

Example: Configure Multinode High Availability in a Layer 3 Network

peer-gnf

IN THIS SECTION

- Syntax | 12359
- Hierarchy Level | 12360
- Description | 12360
- Options | 12360
- Required Privilege Level | 12360
- Release Information | 12360

Syntax

```
peer-gnf {
    id peer-gnf-id;
```

```
    remote-af-name;
}
```

Hierarchy Level

```
[edit chassis network-slices guest-network-functions gnf id af-name]
```

Description

Configure the details of the GNF peer connected using the abstracted fabric (af) interface.

Options

id*peer-gnf -id*

Name of the GNF peer connected using the abstracted fabric (af) interface.

- **Range:** 1 through 10

remote-af-name

Name of the af interface on the peer GNF.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Setting Up Junos Node Slicing

peer-groups (access network peers)

IN THIS SECTION

- Syntax | 12361
- Hierarchy Level | 12361
- Description | 12361
- Options | 12361
- Required Privilege Level | 12362
- Release Information | 12362

Syntax

```
peer-groups name {
    peer {
        address [ address ... ];
        hostname hostname;
    }
    routing-instance routing-instance;
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw name access-network-peers]
```

Description

Define a group of eNodeBs for an SAEGW-U or SGW-U.

Options

name Peer group name.

peer IPv4 address or prefix value (required) of the GTP-U peer and hostname (optional) of the GTP-U peer.

> **NOTE**: The maximum number of peers that can appear in a peer group is 4000.

routing-instance Routing-instance of a GPT-U peer.

> **NOTE**: If at least one peer group is used, eNodeBs matching only this address/prefix are accepted by the user plane function during session establishment. The eNodeB is bound to the routing instance within the `peer-groups` stanza, if available. Otherwise, the user plane function uses the routing instance under "access-network-peers" on page 16270.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.4R1.

peer-groups (control plane peers)

IN THIS SECTION

- Syntax | 12363
- Hierarchy Level | 12363
- Description | 12363
- Options | 12363
- Required Privilege Level | 12364
- Release Information | 12364

Syntax

```
peer-groups name {
    heartbeat-interval seconds;
    n3-requests n3-requests;
    path-management (disable | enable);
    peer {
        address [ address ... ];
        hostname hostname;
    }
    routing-instance routing-instance;
    t3-response seconds;
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw name control-plane-peers]
```

Description

If at least one peer group is configured, PFCP packets from control plane peers matching this address/prefix alone are accepted. The control plane peers are bound to the routing instance defined within this peer group, if configured. Otherwise the control plane function is bound to the routing instance listed under "control-plane-peers" on page 16272.

Options

name Peer group name

heartbeat-interval Time between origination of two successive heartbeat requests (seconds).

- **Range:** 60 through 255
- **Default:** 60

n3-requests Maximim number of retries of PFCP request messages upon t3-response timeout.

- **Range:** 1 through 5
- **Default:** 3

path-management Enable/disable origination of heartbeat message requests to control peers.

- **Default:** Disabled

peer IPv4 or IPv6 address or prefix value (required) of the PFCP peer and hostname (optional) of the PFCP peer.

routing-instance Local routing instance of the PFCP peer.

t3-response Waiting time of gateway before retrying a PFCP signaling-request upon response timeout (seconds).

- **Range:** 1 through 5
- **Default:** 5

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.4R1.

peer-groups (core network peers)

IN THIS SECTION

- Syntax | 12365
- Hierarchy Level | 12365
- Description | 12365
- Options | 12365
- Required Privilege Level | 12366
- Release Information | 12366

Syntax

```
peer-groups name {
    peer {
        address [ address ... ];
        hostname hostname;
    }
    routing-instance routing-instance;
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw name core-network-peers]
```

Description

Define a peer group of PGW-Us for an SGW-U or a group of SGW-Us for a PGW-U.

Options

name Peer group name.

peer IPv4 address or prefix value (required) of the GTP-U peer and hostname (optional) of the GTP-U peer.

> **NOTE**: The maximum number of peers that can appear in a peer group is 4000.

routing-instance Routing-instance of a GPT-U peer.

> **NOTE**: If at least one peer group is used, peers matching only this address/prefix are accepted by the user plane function during session establishment. The peer is bound to the routing instance within the `peer-groups` stanza, if available. Otherwise, the user plane function uses the routing instance under "saegw-core-network-peers" on page 16276.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 20.4R1.

peer (ICCP)

IN THIS SECTION

- Syntax | 12366
- Hierarchy Level | 12367
- Description | 12367
- Required Privilege Level | 12367
- Release Information | 12367

Syntax

```
peer ip-address {
    authentication-key string;
    backup-liveness-detection {
        backup-peer-ip ip-address;
    }
    liveness-detection {
        detection-time {
            threshold milliseconds;
        }
         minimum-intervalmilliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
```

```
            threshold milliseconds;
        }
        version (1 | automatic);
    }
    local-ip-addr ipv4-address;
    session-establishment-hold-time seconds;
}
```

Hierarchy Level

```
[edit protocols iccp]
```

Description

Configure the peers that host a multichassis link aggregation group (MC-LAG). You must configure Inter-Chassis Control Protocol (ICCP) for both peers that host the MC-LAG.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

peer-id

IN THIS SECTION

- Syntax | 12368
- Hierarchy Level | 12369

- Description | 12369
- Options | 12369
- Required Privilege Level | 12369
- Release Information | 12370

Syntax

```
peer-id name {
    desc desc;
    interface interface;
    liveness-detection {
        no-adaptation;
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier multiplier;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (0 | 1 | automatic);
    }
    peer-forwarding-ip loopback-address {
            interface interface-name;
            liveness-detection {
                no-adaptation;
                detection-time {
                    threshold milliseconds;
                }
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier multiplier;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
```

```
            version (0 | 1 | automatic);
    peer-ip peer-ip;
    routing-instance routing-instance;
    vpn-profile vpn-profile;
}
```

Hierarchy Level

```
[edit chassis high-availability]
```

Description

Configure the other node related information in a Multinode High Availability setup.

Options

ID	Peer node identifier.
desc	Peer node description.
interface	Name of the interface used for communicating with the peer node using the interchassis link (ICL.
"liveness-detection" on page 9848	Liveness detection options for the peer node.
"peer-forwarding-ip" on page 12357	IP address used as peer forwarding IP. This IP address must be associated to loopback interface only. You must configure this parameter to support asymmetric traffic flows in Multinode High Availability.
peer-ip	IP address of the peer node.
routing-instance	Routing instance to locate the peer node route.
vpn-profile	VPN profile name for ICL encryption. You must configure a VPN profile for the HA traffic and apply the profile for both the nodes.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Multinode High Availability

Example: Configure Multinode High Availability in a Layer 3 Network

peer-interface (Protocols OSPF)

IN THIS SECTION

- Syntax | 12370
- Hierarchy Level | 12371
- Description | 12371
- Options | 12371
- Required Privilege Level | 12371
- Release Information | 12371

Syntax

```
peer-interface interface-name {
    disable;
    dead-interval seconds;
    hello-interval seconds;
    retransmit-interval seconds;
    transit-delay seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id],
[edit protocols ospf area area-id]
```

Description

Configure a peer interface.

Options

interface-name—Name of the peer interface. To configure all interfaces, you can specify `all`.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring OSPFv2 Peer interfaces

Configuring RSVP and OSPF for LMP Peer Interfaces

Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP

peer-interface (Protocols RSVP)

IN THIS SECTION

- Syntax | 12372
- Hierarchy Level | 12372
- Description | 12372
- Required Privilege Level | 12372
- Release Information | 12373

Syntax

```
peer-interface peer-interface-name {
    disable;
    (aggregate | no-aggregate);
    authentication-key key;
    dynamic-bidirectional-transport template template;
    hello-interval seconds;
    (reliable | no-reliable);
}
```

Hierarchy Level

```
[edit protocols rsvp]
```

Description

Configure the name of the LMP peer device.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`dynamic-bidirectional-transport template` *`template`* option introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring RSVP and OSPF for LMP Peer Interfaces

peer-ip-address-optional

IN THIS SECTION

- Syntax | 12373
- Hierarchy Level | 12373
- Description | 12374
- Required Privilege Level | 12374
- Release Information | 12374

Syntax

```
peer-ip-address-optional;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” ppp-options],
[edit interfaces interface-name pp0 unit unit-number ppp-options],
[edit access group-profile profile-name ppp ppp-options]
```

Description

Enable the Internet Protocol Control Protocol (IPCP) negotiation to succeed even though the peer does not include the IP address option in an IPCP configuration request for static and dynamic, and terminated and tunneled, Point-to-Point Protocol over Ethernet (PPPoE) subscribers. By default, this statement is disabled.

If the client's provisioned IP address on the customer premises equipment (CPE):

- Matches the Framed-Route RADIUS attribute, then you must configure the access route at the `[edit dynamic-profiles routing-options]` hierarchy level.
- Matches the Framed-IP-Address RADIUS attribute, then no access route configuration is required.

> **NOTE**: You must assign an IP address by configuring the Framed-IP-Address RADIUS attribute or the Framed-Pool RADIUS attribute, or by allocating an IP address from the local address pool without a RADIUS-specified pool name, with an *optional* Framed-Route RADIUS attribute returned from the RADIUS Server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

IPCP Negotiation with Optional Peer IP Address

Dynamic Profiles Overview

peer-limit

IN THIS SECTION

- Syntax | 12375
- Hierarchy Level | 12375
- Description | 12375
- Required Privilege Level | 12376
- Release Information | 12376

Syntax

```
peer-limit peer-limit;
```

Hierarchy Level

```
[edit fabric protocols bgp group name dynamic-neighbor],
[edit logical-systems name protocols bgp group name dynamic-neighbor],
[edit logical-systems name routing-instances name protocols bgp group name dynamic-neighbor],
[edit logical-systems name tenants name routing-instances name protocols bgp group name dynamic-
neighbor],
[edit protocols bgp group name dynamic-neighbor],
[edit routing-instances name protocols bgp group name dynamic-neighbor],
[edit tenants name routing-instances name protocols bgp group name dynamic-neighbor]
```

Description

Limit the number of peers for each subnet configured via dynamic-neighbor

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4.

RELATED DOCUMENTATION

neighbor (Protocols BGP) | 11273

peer (Multichassis)

IN THIS SECTION

- Syntax | 12376
- Hierarchy Level | 12377
- Description | 12377
- Required Privilege Level | 12377
- Release Information | 12377

Syntax

```
peer ip-address {
    interface interface-name;
}
```

Hierarchy Level

```
[edit multi-chassis]
```

Description

Configure the IP address of the peer that is part of the interchassis link-protection link (ICL-PL). If Inter-Chassis Control Protocol (ICCP) is up and the interchassis link (ICL) comes up, the peer configured as standby will bring up the multichassis aggregated Ethernet interfaces shared with the active peer specified by the `peer` statement. You must specify the physical interface of the peer.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

peer-origin (Diameter Peer)

IN THIS SECTION

- Syntax | 12378
- Hierarchy Level | 12378
- Description | 12378
- Options | 12378
- Required Privilege Level | 12378
- Release Information | 12378

Syntax

```
peer-origin realm realm-name host hostname;
```

Hierarchy Level

```
[edit diameter peer peer-name]
```

Description

Specify values of Origin-Realm-AVP and Origin-Host-AVP used in messages sent for the specified peer by the Diameter instance.

NOTE: Both the host and realm are mandatory for the peer origin.

Options

host ***hostname***	Name of the message origin host that is supplied as the value of the Origin-Host AVP for Diameter messages to the peer.
realm ***realm-name***	Name of the message origin realm, that is supplied as the value of the Origin-Realm AVP for Diameter messages to the peer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring Diameter Peers

peer (Protocols LMP)

IN THIS SECTION

- Syntax | 12379
- Hierarchy Level | 12379
- Description | 12380
- Options | 12380
- Required Privilege Level | 12380
- Release Information | 12380

Syntax

```
peer peer-name {
    address ip-address;
    control-channel control-channel-interface;
    lmp-control-channel control-channel-interface;
    lmp-protocol {
        hello-dead-interval milliseconds;
        hello-interval milliseconds;
        passive;
        retransmission-interval milliseconds;
        retry-limit number;
    }
    te-link te-link-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management],
[edit protocols link-management]
```

Description

Configure a network peer.

Options

peer-name—Name of the network peer.

The remaining statements are described separately in this chapter.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`lmp-protocol` statement and substatements added in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring LMP Peers

peer (Protocols MSDP)

IN THIS SECTION

- Syntax | 12381
- Hierarchy Level | 12381
- Description | 12381
- Options | 12382
- Required Privilege Level | 12382
- Release Information | 12382

Syntax

```
peer address {
    disable;
    active-source-limit {
        maximum number;
        threshold number;
    }
    authentication-key peer-key;
    default-peer;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name],
[edit protocols msdp],
[edit protocols msdp group group-name],
[edit routing-instances routing-instance-name protocols msdp],
[edit routing-instances routing-instance-name protocols msdp group group-name]
```

Description

Define an MSDP peering relationship. An MSDP routing device must know which routing devices are its peers. You define the peer relationships explicitly by configuring the neighboring routing devices that are the MSDP peers of the local routing device. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. To configure multiple MSDP peers, include multiple `peer` statements.

By default, the peer's options are identical to the global or group-level MSDP options. To override the global or group-level options, include peer-specific options within the "peer (Protocols MSDP)" on page 12380 statement.

At least one peer must be configured for MSDP to function. You must configure ***address*** and **local-address**.

Options

address—Name of the MSDP peer.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP in a Routing Instance

peer (Protocols RIP)

IN THIS SECTION

- Syntax | 12383
- Hierarchy Level | 12383
- Description | 12383
- Options | 12383
- Required Privilege Level | 12383

Release Information | 12383

Syntax

```
peer IP address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit protocols rip group group-name neighbor neighbor-name]
```

Description

Configure a static peer for an interface in a point-to-multipoint RIP network.

Options

address IP address of the static peer to be configured.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring Point-to-Multipoint RIP Networks

peer-selection-service

IN THIS SECTION

- Syntax | 12384
- Hierarchy Level | 12384
- Description | 12384
- Options | 12384
- Required Privilege Level | 12385
- Release Information | 12385

Syntax

```
peer-selection-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Enable the peer selection service process.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the peer selection service process.

- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Interfaces User Guide for Security Devices

peer-subnet (Easy EVPN LAG Configuration)

IN THIS SECTION

- Syntax | 12386
- Hierarchy Level | 12386
- Description | 12386
- Options | 12387
- Required Privilege Level | 12387
- Release Information | 12387

Syntax

```
peer-subnet {
    inet inet-address;
    inet6 inet6-address;
    interface-name [interface-name ... ];
}
```

Hierarchy Level

```
[edit services evpn device-attribute peer-to-peer]
```

Description

Specify interface parameters for the peer-to-peer connections between Ethernet VPN (EVPN) fabric provider edge (PE) devices as part of a generated EVPN fabric configuration.

You use the statements in the `[edit services evpn]` hierarchy level to simplify configuring an EVPN fabric with Ethernet segment identifier (ESI) link aggregation groups (LAGs) in the topology. The feature you enable with statements in this hierarchy is an easy EVPN LAG configuration feature, also called EZ-LAG.

You only need to configure a few statements to provide the parameters required for an EVPN core fabric with multihomed or single-homed customer edge (CE) devices (or end server devices). When you commit the `[edit services evpn]` configuration statements, the device uses a commit script to generate the corresponding standard EVPN fabric configuration statements.

The statements at the `[edit services evpn device-attribute peer-to-peer peer-subnet]` hierarchy level provide the interface names and subnet addresses for the interfaces that connect the peer PE devices in the fabric. The commit script uses these parameters to generate the required statements to configure the aggregated Ethernet interfaces and underlay peering in the EVPN fabric.

You must provide an `inet` or an `inet6` subnet address (or both) for the aggregated Ethernet interface links between the peer PE devices.

For configuration simplicity, you configure the same subnet address on both peer PE devices. By default, the commit script derives a different interface address on each peer PE device based on that device's peer ID, as follows:

- If the peer ID is 1, the commit script uses the subnet address provided.
- Otherwise, the commit script uses the subnet address + 1 (adds 1 to the low order segment of the subnet range).

For example, if you specify *inet-address* as 10.0.1.0/31:

- Peer ID 1 peer-to-peer link address = 10.0.1.0
- Peer ID 2 peer-to-peer link address = 10.0.1.1

> **NOTE**: The commit process doesn't do a commit check to enforce setting the same subnet address across both devices. As a result, we strongly recommend that you make sure when you configure this `subnet-address` statement, you use the same base value on both peer PE devices. That way the default address derivation works as expected for the peer-to-peer interface addresses on each peer PE device.

Options

inet ***inet-address***	The IPv4 subnet address for the interfaces or aggregated Ethernet interfaces that connect the PE device to a peer PE device. You should configure the same subnet address for this option on both peer PE devices.
inet6 ***inet6-address***	The IPv6 subnet address for the interfaces or aggregated Ethernet interfaces that connect the PE device to a peer PE device. You should configure the same subnet address for this option on both peer PE devices.
(Required) interface-name [*interface-name* ...]	The interface or interfaces that connect the PE device to a peer PE device.

See CLI Explorer for more on other related statements.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Easy EVPN LAG (EZ-LAG) Configuration

peer-to-peer (Easy EVPN LAG Configuration)

IN THIS SECTION

- Syntax | 12388
- Hierarchy Level | 12388
- Description | 12389
- Options | 12389
- Required Privilege Level | 12389
- Release Information | 12390

Syntax

```
peer-to-peer {
  overlay-connectivity {
    ibgp;
   }
  peer-subnet {
    inet inet-address;
    inet6 inet6-address;
    interface-name [ interface-name ... ];
  }
  underlay-connectivity {
    ospf;
  }
}
```

Hierarchy Level

```
[edit services evpn device-attribute]
```

Description

Specify parameters for the peer-to-peer interface connections between Ethernet VPN (EVPN) fabric provider edge (PE) devices as part of a generated EVPN fabric configuration.

You use the statements in the `[edit services evpn]` hierarchy level to simplify configuring an EVPN fabric with Ethernet segment identifier (ESI) link aggregation groups (LAGs) in the topology. The feature you enable with statements in this hierarchy is an easy EVPN LAG configuration feature, also called EZ-LAG.

You only need to configure a few statements to provide the parameters required for an EVPN core fabric with multihomed or single-homed customer edge (CE) devices (or end server devices). When you commit the `[edit services evpn]` configuration statements, the device uses a commit script to generate the corresponding standard EVPN fabric configuration statements.

The statements at the `[edit services evpn device-attribute peer-to-peer]` hierarchy level provide parameters to:

- Specify the required interface names and subnet addresses for the connections between peer PE devices.
- Customize the underlay and overlay peering in the generated configuration.

Options

overlay-connectivity

Use the specified protocol for the EVPN fabric overlay peering instead of EBGP, which is the default overlay protocol the `[edit services evpn]` commit script configures.

Available options:

- `ibgp`—Use IBGP for the overlay peering in the generated configuration.

underlay-connectivity

Use the specified protocol for the EVPN fabric underlay peering instead of EBGP, which is the default underlay protocol the `[edit services evpn]` commit script configures.

Available options:

- `ospf`—Use OSPF for the underlay peering in the generated configuration.

The remaining statements are explained in detail on separate pages.

You can also consult CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Easy EVPN LAG (EZ-LAG) Configuration

peer-unit

IN THIS SECTION

- Syntax | 12390
- Hierarchy Level | 12390
- Description | 12391
- Options | 12391
- Required Privilege Level | 12391
- Release Information | 12391

Syntax

```
peer-unit unit-number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure a peer relationship between two logical systems.

Options

unit-number—Peering logical system unit number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 24.1R1 for ACX7K Series Routers.

RELATED DOCUMENTATION

Configuring Logical Tunnel Interfaces

peer-unit (Routing Policy Condition)

IN THIS SECTION

- Syntax | 12392
- Hierarchy Level | 12392
- Description | 12392
- Options | 12392
- Required Privilege Level | 12392
- Release Information | 12392

Syntax

```
peer-unit unit-number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options condition if-route-exists address-family
ccc],
[edit policy-options condition if-route-exists address-family ccc]
```

Description

Specify the associated logical tunnel interface's peer-unit. This is required for logical-tunnel-based routes.

Options

unit-number—Logical unit number of the logical tunnel peer interface.

- **Range:** 0 through 8192

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

> *Example: Configuring Pseudowire Redundancy in a Mobile Backhaul Scenario*

peers (Commit)

IN THIS SECTION

- Syntax | 12393
- Hierarchy Level | 12393
- Description | 12393
- Options | 12394
- Required Privilege Level | 12394
- Release Information | 12394

Syntax

```
peers {
    name of peer {
        user name of user;
        authentication string;
        routing-instance name ;
    }
}
```

Hierarchy Level

```
[edit system commit]
```

Description

Configure options for the peers participating in commit synchronization.

NOTE: In configuration sync across SRX Series Firewalls, routing instance need be specified and configured to let local node reach peer nodes through a configurable routing instance instead of default routing instance.

Options

name of peer	Hostname or IP address of the peer participating in commit synchronization.
`user`	Name of administrator configuring commit synchronization.
`authentication`	Plain-text password string that is stored as an encrypted password string.
`routing-instance`	Specific routing instance for connectivity to remote peer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R6.

peers-synchronize

IN THIS SECTION

- Syntax | 12394
- Hierarchy Level | 12395
- Description | 12395
- Required Privilege Level | 12395
- Release Information | 12395

Syntax

```
peers-synchronize;
```

Hierarchy Level

```
[edit system commit]
```

Description

Configure the `commit` command to automatically perform a `peers-synchronize` action between peers. The local peer (or requesting peer) on which you enable the `peers-synchronize` statement copies and loads its configuration to the remote (or responding) peer. Each peer then performs a syntax check on the configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both peers.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R6.

RELATED DOCUMENTATION

server (Batch Commits) | 16613

synchronize | 28804

peg-df-election

IN THIS SECTION

- Syntax | 12396
- Hierarchy Level | 12396
- Description | 12396

- Options | 12397
- Required Privilege Level | 12397
- Release Information | 12397

Syntax

```
peg-df-election {
    delay-time seconds;
    mod;
    preference {
      use-least;
      value value;
    }
}
```

Hierarchy Level

```
[edit routing-instances name protocols evpn oism pim-evpn-gateway]
```

Description

Use one of the available PIM EVPN gateway (PEG) designated forwarder (DF) election methods with optimized intersubnet multicast (OISM) for a virtual routing and forwarding (VRF) instance in an EVPN fabric.

Otherwise, if you don't enable PEG DF election, by default the OISM border devices you configure as PEG devices use native PIM DF election.

> **NOTE**: When PEG devices use PEG DF election, they don't use the PIM protocol (they don't exchange PIM protocol packets) within the data center. As a result, we recommend you have external L3 redundancy on the PEG devices.

All peer PEG devices must use the same DF election mechanism. As a result, if you enable PEG DF election, configure the same DF election method symmetrically on all peer PEG devices. If the configured PEG DF election methods don't match, the peer PEG devices all fall back to using the default

mod-based PEG DF election method. If you enable PEG DF election on only some of the peer PEG devices, all of the devices fall back to using PIM-based DF election.

Options

delay-time *seconds*

Time to wait (in seconds) before applying the selected method to elect a PEG DF.

- **Range:** 1 through 1800

mod

Use mod-based PEG DF election. This method is the default if you enable PEG DF election but don't specify this option or the `preference` option.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1 and Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

PEG DF Election

show evpn oism peg-df-status | **19826**

pem-absence

IN THIS SECTION

- Syntax | **12398**
- Hierarchy Level | **12398**
- Description | **12398**

- Options | 12398
- Required Privilege Level | 12398
- Release Information | 12398

Syntax

```
pem-absence ignore;
```

Hierarchy Level

```
[edit chassis alarm]
```

Description

Configure `pem-absence ignore` to ignore the red alarm of missing power supply unit. To delete the configuration use `[# delete chassis alarm pem-absence]` command.

Options

`ignore`—Ignores the power supply module absence alarm.

Required Privilege Level

Release Information

Statement introduced in Junos OS Release 20.3R1 for SRX4100 and SRX4200.

Statement introduced in Junos OS Release 20.4R1 for SRX1500 and SRX4600.

RELATED DOCUMENTATION

> *Understanding Chassis Alarms*
>
> show chassis alarms | 17304

pem (MX Series Devices, ACX7024 Devices)

IN THIS SECTION

- Syntax | 12399
- Hierarchy Level | 12399
- Description | 12399
- Options | 12399
- Required Privilege Level | 12400
- Release Information | 12400

Syntax

```
pem {
    minimum number
    redundancy;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the minimum number of Power Entry Modules (PEMs) on an M320, MX10004, MX10008 and ACX7024 devices. With this configuration, PEM absent alarms are generated only if the PEM count falls below the minimum specified.

Options

minimum ***number***—Minimum number of PEMs on the router. ACX7024 device supports dual PSU (1+1 redundancy) mode by default, and also single PSU mode with one PEM or PSM.

- **Range:** 1 through 4 for M320 device. 1, 2 for ACX7024 device.

redundancy— Specify power redundancy. This option is only applicable on the MX10004 and MX10008 platforms with JNP10K-PWR-AC3 PSM.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Junos OS to Support Entry-Level Configuration on an M320 Router with a Minimum Number of SIBs and PIMs

sib | **27939**

show chassis power | **18633**

redundancy (Junos OS) | **14183**

pem (T640, T1600, and T4000 Routers with Six-Input DC Power Supply)

IN THIS SECTION

- Syntax | **12401**
- Hierarchy Level | **12401**
- Description | **12401**
- Options | **12401**
- Required Privilege Level | **12401**
- Release Information | **12401**

Syntax

```
pem {
    feeds number-of-input-feeds;
    input-current amps-in-each-feed;
}
```

Hierarchy Level

```
[edit chassis],
[edit chassis lcc lcc-number]
```

Description

Configure the power supply parameters of the six-input DC power supply on T640, T1600, or T4000 routers.

Options

`feeds` *`number-of-input-feeds`*—Number of input feeds connected to the six-input DC power supply.

(For T4000 routers only) `input-current` *`amps-in-each-feed`*—Input current (in amperes) in each feed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

Option `feeds` introduced in Junos OS Release 12.1.

Option `input-current` introduced for T4000 routers in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring the Six-Input DC Power Supply on T Series Routers

pending-queue-watermark

IN THIS SECTION

- Syntax | 12402
- Hierarchy Level | 12402
- Description | 12402
- Options | 12403
- Required Privilege Level | 12403
- Release Information | 12403

Syntax

```
pending-queue-watermark watermark;
```

Hierarchy Level

```
[edit access radius network-element name]
```

Description

Configure the pending-request queue high watermark for the network element. This is a percentage of the maximum number of requests that can be queued to the network element, which is configured in the `maximum-pending-reqs-limit number` statement at the `[edit access radius network-element name]` hierarchy level. When the queue size reaches the high watermark, a `flow control on` message is generated.

Options

watermark High watermark for the network element pending request queue.

- **Range:** 1 through 100 percent.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Network Elements

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

pending-queue-watermark-abate

IN THIS SECTION

- Syntax | 12404
- Hierarchy Level | 12404
- Description | 12404
- Options | 12404
- Required Privilege Level | 12404
- Release Information | 12404

Syntax

```
pending-queue-watermark-abate abate-watermark;
```

Hierarchy Level

```
[edit access radius network-element name]
```

Description

Configure the low watermark of the pending-request queue for the network element. This is a percentage of the maximum size of the pending-request queue, which is configured in the `maximum-pending-reqs-limit` *`watermark`* statement at the `[edit access radius network-element` *`name`*`]` hierarchy level. When the number of pending requests drops below this low watermark value after having exceeded the high watermark configured in the `pending-queue-watermark` *`watermark`* statement, a `flow control off` message is generated.

Options

abate-watermark Low watermark for the network element pending request queue.

- **Range:** 1 through 100 percent.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Network Elements

Understanding Network Elements

pending-sess-queue-length

IN THIS SECTION

- Syntax | 12405
- Hierarchy Level | 12405
- Description | 12405
- Options | 12405
- Required Privilege Level | 12406
- Release Information | 12406

Syntax

```
pending-sess-queue-length (high | moderate | normal);
```

Hierarchy Level

```
[edit security flow]
```

Description

Configure the maximum queued length per pending session.

Options

- `high`— Allow the maximum number of queued sessions.
- `moderate`—Allow more queued sessions than the normal number.
- `normal`—Allow the normal number of queued session.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

per-esi

IN THIS SECTION

- Syntax | 12406
- Hierarchy Level | 12407
- Description | 12407
- Options | 12407
- Required Privilege Level | 12408
- Release Information | 12408

Syntax

```
per-esi {
    lacp-oos-on-ndf;
}
```

Hierarchy Level

```
[edit interfaces name esi df-election-granularity]
```

Description

Ensure that all the member links in an aggregated Ethernet (AE) interface between a multihomed customer edge (CE) device and its multihoming peer provider edge (PE) devices take the same designated forwarder (DF) election role per Ethernet segment (ES). The behavior with this option is called port-based DF election with EVPN multihoming.

You can't configure this statement and the `df-election-granularity per-esi-vlan` statement at the same time. If you configure one of these statements when the other is already set, the CLI replaces the existing setting with the new one.

With active-standby multihoming, if the CE device has an AE interface that shares member links with both the active and standby peer PE devices, you must include the `lacp-oos-on-ndf` option with this statement. By default, the multihomed CE device load balances traffic toward the peer PE devices among all the AE links on the ES. The `lacp-oos-on-ndf` option ensures that the member link to the PE device that isn't the DF (the non-DF) remains in LACP Waiting state so the device only sends traffic on the link to the DF PE device. Otherwise, if the CE device sends part of the traffic to the non-DF, the non-DF doesn't forward it and the network loses some traffic.

> **NOTE**: You cannot enable the `sync-reset` statement at the `[edit interfaces name aggregated-ether-options lacp]` level when you have the `lacp-oos-on-ndf` statement enabled.

Options

lacp-oos-on-ndf In an EVPN single-active configuration, the PE device that is the non-DF will send LACP out-of-sync packets to the multihomed CE device. As a result, LACP goes down on the CE device, and the CE device does not send traffic on the links connected to the non-DF.

If the connection between a CE device and the DF PE device fails, the DF election process selects a new DF. If the connection between the CE device and a non-DF PE device fails, the current DF assignment remains the same.

This option only applies to single-active multihoming where the CE device has an AE interface with member links to multiple peer PE devices associated with an ESI. The CLI doesn't allow you to configure this option if:

- You configure an ES with active-active (all-active) multihoming.
- The interface is not an AE interface.
- You didn't enable LACP on the AE interface.

Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1

RELATED DOCUMENTATION

per-esi-vlan | 12408

per-esi-vlan

IN THIS SECTION

- Syntax | 12408
- Hierarchy Level | 12409
- Description | 12409
- Required Privilege Level | 12409
- Release Information | 12409

Syntax

```
per-esi-vlan;
```

Hierarchy Level

```
[edit interfaces name esi df-election-granularity]
```

Description

Enable aggregated Ethernet interface member links between a multihomed CE device and its multihoming peer provider edge (PE) devices to take different designated forwarder (DF) election roles per Ethernet segment (ES) and per VLAN. This is the default behavior with EVPN multihoming.

If you want to ensure all the member links take the same DF election role for the Ethernet segment, use the `df-election-granularity per-esi` option instead of this option.

You can't configure this statement and the `per-esi` statement at the same time. If you configure one of these statements when the other is already set, the CLI replaces the existing setting with the new one.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1

RELATED DOCUMENTATION

per-esi | 12406

per-evi

IN THIS SECTION

- Syntax | 12410
- Hierarchy Level | 12410
- Description | 12410

Required Privilege Level | 12412
Release Information | 12412

Syntax

```
per-evi;
```

Hierarchy Level

```
[edit interfaces interface-name esi df-election-granularity],
[edit interfaces interface-name unit logical-unit-number esi df-election-granularity]
```

Description

Enable the designated forwarder (DF) election granularity to support Ethernet segment ID (ESI) at the EVPN instance (EVI) level.

You can configure multiple access interfaces to support ESI at the EVI level between a multihomed customer edge (CE) device and its multihoming provider edge (PE) peer. When you configure the access interfaces in the same EVI with the same ESI and the `per-evi` statement, these interfaces will share the same designated forwarder (DF) election role. They will transition to the backup designated forwarder (BDF) role only if all the access interfaces within the EVI fail. This configuration avoids traffic disruptions caused by role changes due to a single interface failure in the EVI.

To enable the ESI at the EVI level support:

- Configure a unique ESI at the EVI level.
  - [edit routing-instances *routing-instance-name* protocols "evpn" on page 6541 esi]
- Configure the access logical interfaces (IFLs) with the same ESI value as the EVI.
  - [edit interfaces *interface-name* unit *logical-unit-number* "esi" on page 6453]

> **NOTE**: Different IFLs under different physical interfaces (IFDs) can have the same ESI configured, but they must be part of the same bridge domain.

- Configure the IFLs with the "df-election-granularity" on page 5574 option `per-evi`.
- Configure the IFLs with the `esi single-active` option.
- Configure the IFLs with the "df-election-type" on page 5576 option `mod`.

> **NOTE**: If you configure the the `df-election-granularity per-evi` option, you must also configure the "df-election-type" on page 5576 `mod` option without configuring a `df-election-type preference` option to prevent an error when committing the configuration.

The following shows snippets from a sample configuration:

```
user@MX960> show configuration interfaces
ge-0/0/1 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 100;
        esi {
            00:11:11:11:11:11:11:11:11:11;
            single-active;
            df-election-granularity {
                per-evi;
            }
            df-election-type {
                mod;
            }
        }
    }
}
ge-0/0/2 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 100;
```

```
        esi {
            00:11:11:11:11:11:11:11:11:11;
            single-active;
            df-election-granularity {
                per-evi;
            }
            df-election-type {
                mod;
            }
        }
    }
}

user@MX960> show configuration routing-instances
EVPN {
    instance-type evpn;
    protocols {
        evpn {
            esi {
                00:11:11:11:11:11:11:11:11:11;
            }
        }
    }
    vlan-id 100;
    interface ge-0/0/1.0;
    interface ge-0/0/2.0;
    route-distinguisher 10.255.0.1:100;
    vrf-target target:100:100;
}
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1

RELATED DOCUMENTATION

EVPN Multihoming Overview

per-flow

IN THIS SECTION

- Syntax | 12413
- Hierarchy Level | 12413
- Description | 12413
- Options | 12414
- Required Privilege Level | 12414
- Release Information | 12414

Syntax

```
per-flow {
    hash-seed;
}
```

Hierarchy Level

```
[edit forwarding-options load-balance],
[edit logical-systems logical-system-name forwarding-options load-balance],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options load-balance],
[edit routing-instances routing-instance-name forwarding-options load-balance]
```

Description

Enable per-flow load balancing based on hash values.

Options

hash-seed—Configure the hash value. Junos OS automatically chooses a value for the hashing algorithm used. You cannot configure a specific hash value for per-flow load balancing.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3 (M120, M320, and MX Series routers only).

RELATED DOCUMENTATION

Configuring Per-Flow Load Balancing Based on Hash Values

load-balance (Forwarding Options) | **9878**

per-flow (Aggregated Ethernet Interfaces)

IN THIS SECTION

- Syntax | **12415**
- Hierarchy Level | **12415**
- Description | **12415**
- Options | **12415**
- Required Privilege Level | **12415**
- Release Information | **12415**

Syntax

```
per-flow;
```

Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options load-balance-stateful]
```

Description

Enable the mechanism to perform an even, effective distribution of traffic flows across member links of an aggregated Ethernet interface (ae) bundle on MX Series routers with MPCs, except MPC3Es and MPC4Es. When multiple flows are transmitted out of an ae interface, the flows must be distributed across the different member links evenly to enable an effective and optimal load-balancing behavior. To obtain a streamlined and robust method of load-balancing, the member link of the aggregated Ethernet interface bundle that is selected each time for load balancing plays a significant part.

Options

per-flow Enable the stateful load-distribution mechanism per traffic flow on an aggregated Ethernet interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces

per-group-label

IN THIS SECTION

- Syntax | 12416
- Hierarchy Level | 12416
- Description | 12416
- Required Privilege Level | 12416
- Release Information | 12417

Syntax

```
per-group-label;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family inet labeled-unicast],
[edit logical-systems logical-system-name protocols bgp group group-name family inet labeled-
unicast],
[edit protocols bgp family inet labeled-unicast],
[edit protocols bgp group group-name family inet labeled-unicast]
```

Description

Account for traffic from each customer separately by advertising separate labels for the same prefix to the peer routers in the BGP groups.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics

Understanding Interprovider and Carrier-of-Carriers VPNs

per-prefix

IN THIS SECTION

- Syntax | 12417
- Hierarchy Level | 12417
- Description | 12418
- Options | 12418
- Required Privilege Level | 12418
- Release Information | 12418

Syntax

```
per-prefix {
    hash-seed number;
}
```

Hierarchy Level

```
[edit forwarding-options load-balance],
[edit logical-systems logical-system-name forwarding-options load-balance],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
```

```
options load-balance],
[edit routing-instances routing-instance-name forwarding-options load-balance]
```

Description

Configure the hash parameter for per-prefix load balancing.

Options

hash-seed—Per-prefix load-balancing hash function.

number—Hash value.

- **Range:** 0 through 65,534
- **Default:** 0

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring Per-Prefix Load Balancing

load-balance (Forwarding Options) | **9878**

per-prefix-label

IN THIS SECTION

- Syntax | 12419
- Hierarchy Level | 12419
- Description | 12419
- Default | 12420
- Required Privilege Level | 12420
- Release Information | 12420

Syntax

```
per-prefix-label;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family inet labeled-unicast],
[edit logical-systems logical-system-name protocols bgp group group-name family inet labeled-
unicast],
[edit protocols bgp family inet labeled-unicast],
[edit protocols bgp group group-name family inet labeled-unicast],
[edit routing-instances instance-name logical-systems logical-system-name protocols bgp family
inet labeled-unicast],
[edit routing-instances instance-name logical-systems logical-system-name protocols bgp group
group-name family inet labeled-unicast],
[edit routing-instances instance-name protocols bgp family inet labeled-unicast],
[edit routing-instances instance-name protocols bgp group group-name family inet labeled-unicast]
```

Description

Allocate a unique label for each prefix. The `per-prefix-label` statement helps minimize packet loss in most deployments.

Although allocating a label for each prefix is not generally ideal for scaling, it is assumed that a small number of labels are used for BGP labeled-unicast. When labeled BGP is used to set up transport label-switched paths (LSPs), the common case is that each prefix has a unique next hop. Thus, the use of per-prefix labels does not have an adverse scaling impact. On the contrary, the use of per-prefix labels reduces churn in the network when multipath load balancing is enabled for IPv4 labeled-unicast, and a subset of the paths are withdrawn for some reason.

The advantage of per-prefix labeling is that the advertised upstream label is more stable during network changes. That is, if the downstream label changes, the advertised upstream label remains the same under most scenarios. This way, the upstream router is isolated from the downstream network change, and the overall network is more stable. The greater stability of the advertised upstream label helps to reduce traffic loss during many different network change scenarios.

Default

By default, label allocation is per next-hop router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1x48.

RELATED DOCUMENTATION

MPLS Label Allocation

per-route-accounting

IN THIS SECTION

- Syntax | **12421**
- Hierarchy Level | **12421**

- Description | **12421**
- Default | **12421**
- Required Privilege Level | **12421**

Syntax

```
per-route-accounting
```

Hierarchy Level

```
[edit routing-options flow],
[edit routing-instances flow],
[edit routing-options-rib-inet6.0 flow],
[edit routing-instance-rib-inet6.0 flow]
```

Description

Earlier to Junos OS Evolved Release 21.1R1, by default, in the usual BGP FlowSpec filter, an individual counter was added for each FlowSpec term configured. As part of optimization of FlowSpec, starting from Junos OS Evolved Release 21.1R1 this behavior is changed, and counter action is not added by default. This will create more room for filter terms compression. When the term counters are not programmed, the traffic matching against a specific flow term is not accounted. You can use the new configuration statement `per-route-accounting` under the `[edit routing-options flow]` hierarchy level to enable creation of term counters for flowspec at global level on PTX10001-36MR, PTX10004, and PTX10008.

Default

The term counters are disabled by default.

Required Privilege Level

routing-options flow-To enable the new configuration statement.

per-session-scheduler

IN THIS SECTION

- Syntax | 12422
- Hierarchy Level | 12422
- Description | 12422
- Required Privilege Level | 12422
- Release Information | 12422

Syntax

```
per-session-scheduler;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Enable session-aware CoS shaping on this L2TP interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

ingress-shaping-overhead | 8649

mode (Traffic Manager) | 10807

per-unit-scheduler

IN THIS SECTION

- Syntax | 12423
- Hierarchy Level | 12423
- Description | 12423
- Required Privilege Level | 12424
- Release Information | 12424

Syntax

```
per-unit-scheduler;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Enable the association of scheduler maps or traffic control profiles with logical interfaces.

CAUTION: Turning on per-unit scheduling causes the interface to reinitialize, which means all logical interfaces (units) on the interface are deleted and recreated.

When including the `per-unit-scheduler` statement, you must also include the `vlan-tagging` statement or the `flexible-vlan-tagging` statement (to apply scheduling to VLANs) or the `encapsulation frame-relay` statement (to apply scheduling to DLCIs) at the [`edit interfaces` *`interface-name`*] hierarchy level.

> **NOTE**: To enable per-unit scheduling on MX80 and MX104 routers, configure the `per-unit-scheduler` statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs*
- *Example: Applying Scheduling and Shaping to VLANs*
- *Configuring Virtual LAN Queuing and Shaping on PTX Series Routers*
- *Providing a Guaranteed Minimum Rate*

perfect-forward-secrecy (Security)

IN THIS SECTION

- Syntax | 12425
- Hierarchy Level | 12425
- Description | 12425

- Options | 12425
- Required Privilege Level | 12425
- Release Information | 12426

Syntax

```
perfect-forward-secrecy {
    keys (group1 | group2);
}
```

Hierarchy Level

```
[edit security ipsec policy ipsec-policy-name]
```

Description

Define the Perfect Forward Secrecy (PFS) protocol. Create single-use keys.

Options

`keys`—Type of Diffie-Hellman prime modulus group that IKE uses when performing the new Diffie-Hellman exchange.

The key can be one of the following:

- `group1`—768-bit.
- `group2`—1024-bit.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the IPsec Policy for an ES PIC

perfect-forward-secrecy (Services)

IN THIS SECTION

- Syntax | 12426
- Hierarchy Level | 12426
- Description | 12427
- Options | 12427
- Required Privilege Level | 12427
- Release Information | 12427

Syntax

```
perfect-forward-secrecy {
    keys (group1 | group2 |group5 |group14 |group15 | group16 | group24);
}
```

Hierarchy Level

```
[edit services ipsec-vpn ipsec policy policy-name]
```

Description

Define Perfect Forward Secrecy (PFS). Creates single-use keys. This statement is optional.

Options

keys—Type of Diffie-Hellman prime modulus group that IKE uses when performing the new Diffie-Hellman exchange. The key can be one of the following:

group1	768-bit.
group2	1024-bit.
group5	1536-bit.
group14	2048-bit.
group15	3072-bit.
group16	4096-bit.
group24	2048-bit with 256-bit Prime Order Subgroup.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`group15`, `group16`, and `group24` options added in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Configuring IPsec Policies

performance

IN THIS SECTION

- Syntax | 12428
- Hierarchy Level | 12428
- Description | 12428
- Options | 12428
- Required Privilege Level | 12429
- Release Information | 12429

Syntax

```
performance {
    values [fast normal slow unknown];
}
```

Hierarchy Level

```
[edit security idp dynamic-attack-group  dynamic-attack-group-name filters]
```

Description

Specify a performance filter to add attack objects based on the performance level that is vulnerable to the attack.

Options

`values`—Name of the performance filter. You can select from the following performance levels:

- `fast`—Fast track performance level.
- `normal`—Normal track performance level.
- `slow`—Slow track performance level.

- `unknown`—By default, all compound attack objects are set to Unknown. As you fine-tune IDP to your network traffic, you can change this setting to help you track performance level.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

performance-monitor

IN THIS SECTION

- Syntax | 12429
- Hierarchy Level | 12430
- Description | 12430
- Options | 12430
- Required Privilege Level | 12430
- Release Information | 12430

Syntax

Supported for PTX10004 and PTX10008 devices only.

```
performance-monitor {
 passive-port {
  delay-request-rate delay-request-rate;
 }
```

```
  slave-port;
 }
```

Hierarchy Level

```
[edit logical-systems name protocols ptp],
[edit protocols ptp]
```

Description

PTP packet delay metrics

Options

passive-port Enable PTP performance monitoring for passive ports

slave-port Enable IEEE 1588v2.1 (2019) Annex J1 & J3 performance monitoring

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.4R1.

performance-monitoring

IN THIS SECTION

- Syntax | 12431
- Hierarchy Level | 12432
- Description | 12432
- Required Privilege Level | 12432
- Release Information | 12432

Syntax

```
performance-monitoring {
    delegate-server-processing;
    hardware-assisted-timestamping;
    hardware-assisted-keepalives;
    sla-iterator-profiles {
        profile-name {
            avg-fd-twoway-threshold;
            avg-ifdv-twoway-threshold;
            avg-flr-forward-threshold;
            avg-flr-backward-threshold;
            disable;
            calculation-weight {
                delay delay-weight;
                delay-variation delay-variation-weight;
            }
            cycle-time milliseconds;
            iteration-period connections;
            measurement-type (loss | statistical-frame-loss | two-way-delay);
        }
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

Description

Specify performance monitoring support for Ethernet frame delay measurement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

ITU-T Y.1731 Ethernet Service OAM Overview

Configuring Ethernet Frame Delay Measurement Sessions

performance-monitoring (OAM LFM)

IN THIS SECTION

- Syntax | 12433
- Hierarchy Level | 12433
- Description | 12433
- Required Privilege Level | 12433
- Release Information | 12433

Syntax

```
performance-monitoring {
    sla-iterator-profiles {
        profile-name {
            calculation-weight {
                delay delay-value;
                delay-variation delay-variation-value;
            }
            cycle-time cycle-time-value;
            iteration-period iteration-period-value;
            measurement-type two-way-delay;
            passive;
        }
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

Description

Specify performance monitoring support for Ethernet frame delay measurement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configure MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements (CLI Procedure)

Configure One-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure)

Configure Two-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure)

performance-monitoring (Protocols MPLS)

IN THIS SECTION

- Syntax | 12434
- Hierarchy Level | 12435
- Description | 12435
- Required Privilege Level | 12436
- Release Information | 12436

Syntax

```
performance-monitoring {
    querier {
        delay {
            traffic-class tc-value {
                average-sample-size sample size;
                padding-size size;
                query-interval milliseconds;
                rtt-delay-threshold rtt threshold value;
                twcd-delay-threshold twcd threshold value;
            }
        }
        loss {
            traffic-class tc-value {
                average-sample-size sample size;
                loss-threshold loss threshold value;
                loss-threshold-window number of samples for loss threshold;
```

```
                measurement-quantity bytes|packets;
                query-interval milliseconds;
            }
        }
        loss-delay {
            traffic-class tc-value {
                average-sample-size sample size;
                loss-threshold loss threshold value;
                loss-threshold-window number of samples for loss threshold;
                measurement-quantity bytes|packets;
                padding-size size;
                query-interval milliseconds;
                rtt-delay-threshold rtt threshold value;
                twcd-delay-threshold twcd threshold value;
            }
        }
    }
    responder {
        delay {
            min-query-interval milliseconds;
        }
        loss {
            min-query-interval milliseconds;
        }
    }
}
```

Hierarchy Level

```
[edit protocols mpls oam],
[edit protocols mpls label-switched-path lsp-name oam],
[edit protocols mpls label-switched-path lsp-name primary path-name oam],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam]
```

Description

Configure performance monitoring options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

periodic (Chassis)

IN THIS SECTION

- Syntax | 12436
- Hierarchy Level | 12437
- Description | 12437
- Options | 12437
- Required Privilege Level | 12437
- Release Information | 12437

Syntax

```
periodic {
    ch-power-check-interval seconds;
    no-sib-asic-stats;
    one-sib-per-check;
    re-check-interval seconds;
    slow-fpc-check;
    yield-cpu;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Chassisd periodic options

Options

ch-power-check-interval

Chassis power check interval, 0=default (seconds)

- **Default:** 0
- **Range:** 0 through 60

no-sib-asic-stats

Disable asic statis for all SIBs in periodic

one-sib-per-check

Check one SIB per periodic

re-check-interval

Chassis RE check interval, 0=default (seconds)

- **Default:** 0
- **Range:** 0 through 60

slow-fpc-check

Slow FPC check in periodic

yield-cpu

Yield cpu in FRU periodic

Required Privilege Level

[none specified]

Release Information

Statement introduced in Junos OS Release 16.2R1.

periodic

IN THIS SECTION

- Syntax | 12438
- Hierarchy Level | 12438
- Description | 12438
- Options | 12438
- Required Privilege Level | 12439
- Release Information | 12439

Syntax

```
periodic (fast | slow);
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options lacp],
[edit interfaces interface-range name aggregated-ether-options lacp]
```

Description

For aggregated Ethernet interfaces only, configure the interval for periodic transmission of LACP packets.

Options

- `fast`—Transmit packets every second.
- `slow`—Transmit packets every 30 seconds.
- **Default:** `fast`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring LACP for Aggregated Ethernet Interfaces
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*
- Junos OS Network Interfaces Library for Routing Devices

periodic (Interfaces)

IN THIS SECTION

- Syntax | 12440
- Hierarchy Level | 12440
- Description | 12440
- Options | 12440
- Required Privilege Level | 12440
- Release Information | 12440

Syntax

```
periodic (fast | slow);
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options lacp]
```

Description

For redundant Ethernet interfaces in a chassis cluster only, configure the interval at which the interfaces on the remote side of the link transmit link aggregation control protocol data units (PDUs) by configuring the `periodic` statement on the interfaces on the local side. It is the configuration on the local side that specifies the behavior of the remote side. That is, the remote side transmits link aggregation control PDUs at the specified interval.

Options

- `fast`—Transmit link aggregation control PDUs every second.
- `slow`—Transmit link aggregation control PDUs every 30 seconds.
- **Default:** `fast`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

periodic-traceroute

IN THIS SECTION

- Syntax | 12441
- Hierarchy Level | 12441
- Description | 12442
- Options | 12442
- Required Privilege Level | 12443
- Release Information | 12443

Syntax

```
periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp oam],
[edit logical-systems logical-system-name protocols ldp oam fec fec-address],
[edit protocols ldp oam],
[edit protocols ldp oam fec fec-address]
```

Description

Enable tracing of forwarding equivalence classes (FECs) for LDP LSPs.

Options

`disable`—(Optional) Disable tracing for a specific FEC. This option is available at the `[edit protocols ldp oam fec` *`fec-address`* `periodic-traceroute]` and `[edit logical-systems` *`logical-system-name`* `protocols ldp oam fec` *`fec-address`* `periodic-traceroute]` hierarchy levels only.

`exp` *`exp-value`*—(Optional) Specify the class of service to use when sending probes.

- **Default:** 7
- **Range:** 0 through 7

`fanout` *`fanout-value`*—(Optional) Specify the maximum number of next hops to search per node.

- **Default:** 16
- **Range:** 1 though 16

`frequency` *`minutes`*—(Optional) Specify the interval between traceroute attempts.

- **Default:** 60 minutes
- **Range:** 15 through 120 minutes

`paths` *`number-of-paths`*—(Optional) Specify the maximum number of paths to search.

- **Default:** 3
- **Range:** 1 through 255

`retries` *`retry-attempts`*—(Optional) Specify the number of attempts to send a probe to a specific node before giving up.

- **Default:** 3
- **Range:** 1 through 9

`source` *`address`*—(Optional) Specify the IPv4 source address to use when sending probes.

`ttl` *`value`*—(Optional) Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.

- **Default:** 64
- **Range:** 1 through 255

`wait` *`seconds`*—(Optional) Specify the wait interval before resending a probe packet.

- **Default:** 10 seconds
- **Range:** 5 though 15 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.
Support added at the `[edit protocols ldp oam]` and `[edit logical-systems` *`logical-system-name`* `protocols ldp oam]` hierarchy levels in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring LDP LSP Traceroute

permit-command

IN THIS SECTION

- Syntax | 12444
- Hierarchy Level | 12444
- Description | 12444
- Required Privilege Level | 12444
- Release Information | 12444

Syntax

```
permit-command protocol-command-list;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name]
```

Description

Apply protocol permit command custom-objects to the content-filtering profile.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

permit (Security Persistent NAT)

IN THIS SECTION

Syntax | 12445

- Hierarchy Level | 12445
- Description | 12445
- Options | 12445
- Required Privilege Level | 12446
- Release Information | 12446

Syntax

```
permit ( any-remote-host | target-host | target-host-port );
```

Hierarchy Level

```
[edit security nat source rule-set ruleset rule rule then source-nat interface persistent-nat]
[edit security nat source rule-set ruleset rule rule then source-nat pool persistent-nat]
```

Description

Configure persistent NAT mappings.

Options

- `any-remote-host`—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. (The reflexive transport address is the public IP address and port created by the NAT device closest to the STUN server.) Any external host can send a packet to the internal host by sending the packet to the reflexive transport address.
- `target-host`—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address.
- `target-host-port`—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address and port.

> **NOTE**: The target-host-port configuration is not supported for NAT64 when configured with IPv6 address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 9.6. Support.

permit (Security Policies)

IN THIS SECTION

- Syntax | 12446
- Hierarchy Level | 12448
- Description | 12448
- Options | 12448
- Required Privilege Level | 12448
- Release Information | 12448

Syntax

```
permit {
    advanced-connection-tracking;
    application-services (Security Policies) {
        application-firewall {
            rule-set rule-set-name;
        }
        application-traffic-control {
```

```
            rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
        ssl-proxy {
            profile-name profile-name;
        }
        uac-policy {
            captive-portal captive-portal;
        }
        utm-policy policy-name;
    }
    destination-address {
        drop-translated;
        drop-untranslated;
    }
    firewall-authentication {
        pass-through {
            access-profile profile-name;
            client-match user-or-group-name;
            ssl-termination-profile profile-name;
            web-redirect;
            web-redirect-to-https;
        }
        user-firewall {
            access-profile profile-name;
            domain domain-name
            ssl-termination-profile profile-name;
        }
        web-authentication {
            client-match user-or-group-name;
        }
    }
    services-offload;
    tcp-options {
        sequence-check-required;
        syn-check-required;
    }
    tunnel {
        ipsec-group-vpn group-vpn;
        ipsec-vpn vpn-name;
```

```
        pair-policy pair-policy;
    }
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then]
```

Description

Specify the policy action to perform when packets match the defined criteria.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for the `tcp-options` added in Junos OS Release 10.4. Support for the `services-offload` option added in Junos OS Release 11.4. Support for the `ssl-termination-profile` and `web-redirect-to-https` options added in Junos OS Release 12.1X44-D10. Support for the `user-firewall` option added in Junos OS Release 12.1X45-D10. Support for the `advanced-connection-tracking` option is added in Junos OS Release 20.2R1.

You can configure the `advanced-connection-tracking` option under `[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit]` to mandate that traffic matching given policy do a lookup in the *to-zone*'s connection track mapping table using the new session's key information. If there is no match, a new connection is not created.

permit (Security SCTP)

IN THIS SECTION

- Syntax | 12449
- Hierarchy Level | 12449
- Description | 12449
- Options | 12449
- Required Privilege Level | 12450
- Release Information | 12450

Syntax

```
permit {
    payload-protocol name;
}
```

Hierarchy Level

```
[edit security sctp profile profile-name]
```

Description

Use the permit option to display information about the configuration of the current Stream Control Transmission Protocol (SCTP) inspection.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs sctp profile profile-name]` hierarchy level is replaced by `[edit security sctp profile profile-name]`.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Understanding Stream Control Transmission Protocol

Example: Configuring a GPRS SCTP Profile for Policy-Based Inspection to Reduce Security Risks

perpetual-poe

IN THIS SECTION

- Syntax | 12450
- Hierarchy Level | 12451
- Description | 12451
- Required Privilege Level | 12451
- Release Information | 12451

Syntax

```
perpetual-poe;
```

Hierarchy Level

```
[edit poe]
```

Description

Configure perpetual power over Ethernet (PoE) to prevent interruptions to power during a reboot of the power sourcing equipment (PSE). The PSE is a switch that provides PoE to connected powered devices.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1R1.

RELATED DOCUMENTATION

Perpetual PoE and Fast PoE

persistent-learning

IN THIS SECTION

- Syntax | 12452
- Hierarchy Level | 12452
- Description | 12452
- Required Privilege Level | 12452
- Release Information | 12452

Syntax

```
persistent-learning;
```

Hierarchy Level

- For platforms without ELS:

```
[edit ethernet-switching-options secure-access-port interface (all | interface-name)]
```

- For platforms with ELS:

```
[edit switch-options interface interface-name]
```

Description

Specify that learned MAC addresses persist on the specified interfaces across restarts of the switch and link-down conditions. This feature is also known as sticky MAC.

Required Privilege Level

system—To view this statement in the configuration.
system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Hierarchy level [edit switch-options interface interface-name] introduced in Junos OS Release 13.2X50-D10

RELATED DOCUMENTATION

Example: Configuring Port Security (non-ELS)

Configuring Persistent MAC Learning (non-ELS)

Configuring Persistent MAC Learning (ELS)

persistent-nat

IN THIS SECTION

- Syntax | 12453
- Hierarchy Level | 12453
- Description | 12453
- Options | 12454
- Required Privilege Level | 12454
- Release Information | 12454

Syntax

```
persistent-nat {
    address-mapping;
    inactivity-timeout seconds;
    max-session-number value;
    permit (any-remote-host | target-host | target-host-port);
}
```

Hierarchy Level

```
[edit security nat source rule-set ruleset rule rule then source-nat interface]
[edit security nat source rule-set ruleset rule rule then source-nat pool]
```

Description

Use the `persistent-nat` feature to ensure that all requests from the same internal transport address are mapped to the same reflexive transport address (the public IP address and port created by the NAT device closest to the *STUN* server). The source NAT rule action can use a source NAT pool (with or without port translation) or an egress interface.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

`security`—To view this statement in the configuration.

`security-control`—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 9.6. Support.

RELATED DOCUMENTATION

Understanding Persistent NAT and NAT64

persistent-results

IN THIS SECTION

- Syntax | 12454
- Hierarchy Level | 12455
- Description | 12455
- Required Privilege Level | 12455
- Release Information | 12455

Syntax

```
persistent-results;
```

Hierarchy Level

```
[edit services rpm twamp client control-connection control-client-name]
```

Description

When enabled, allows to display the old and current tests' results after a network recovery or after TWAMP server reachability when you execute the `show services rpm twamp client probe-results` command. By default, the option is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

Understanding Two-Way Active Measurement Protocol on Routers

pfc

IN THIS SECTION

- Syntax | 12456
- Hierarchy Level | 12456
- Description | 12456
- Required Privilege Level | 12456
- Release Information | 12456

Syntax

```
pfc;
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options compression],
[edit interfaces interface-name unit logical-unit-number ppp-options compression],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-
options compression]
```

Description

For interfaces with PPP encapsulation, configure the router to compress the protocol field to one byte.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring PPP

pfc (Input Congestion Notification)

IN THIS SECTION

Syntax | 12457

- Hierarchy Level | 12457
- Description | 12457
- Options | 12457
- Required Privilege Level | 12457
- Release Information | 12458

Syntax

```
pfc {
    mru mru-value;
}
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name input dscp code-point code-point-bits]
[edit class-of-service congestion-notification-profile profile-name input ieee-802.1 code-point code-point-bits]
```

Description

Enable and configure ingress interface priority-based flow control (PFC) for a configured IEEE 802.1p code point in the VLAN-tagged packet header at Layer 2, or on a DSCP value from the Layer 3 IP header to support PFC with untagged traffic.

Options

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for DSCP-based PFC introduced in Junos OS Release 17.4R1 for the QFX Series.

RELATED DOCUMENTATION

Configuring CoS PFC (Congestion Notification Profiles)

Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces

Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface

Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)

Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows

Understanding PFC Using DSCP at Layer 3 for Untagged Traffic

Configuring DSCP-based PFC for Layer 3 Untagged Traffic

pfc-priority

IN THIS SECTION

- Syntax | **12459**
- Hierarchy Level | **12459**
- Description | **12459**
- Options | **12459**
- Required Privilege Level | **12459**
- Release Information | **12460**

Syntax

```
pfc-priority pfc-priority;
```

Hierarchy Level

```
[edit class-of-service forwarding-classes class class-name]
```

Description

Explicitly map a forwarding class to a priority-based flow control (PFC) priority value.

To support lossless behavior for untagged traffic across Layer 3 connections to Layer 2 subnetworks, you can configure PFC to be invoked based on a configured 6-bit Distributed Services code point (DSCP) value in the Layer 3 IP header, rather than an IEEE 802.1p code point in a VLAN header at Layer 2. However, because PFC sends Layer 2 pause frames specifying a PFC priority on which to notify the peer about congestion, this statement defines the PFC priority to use in pause frames when PFC is triggered by a configured DSCP value.

DSCP-based PFC is used to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2).

To enable DSCP-based PFC:

- Use this statement to map a lossless forwarding class to a PFC priority value to use in the PFC pause frames.
- Use the `[edit class-of-service congestion-notification-profile name input] dscp` statement to define an input congestion notification profile to enable PFC on traffic specified by a desired DSCP value.
- Use the `[edit class-of-service classifiers] dscp` statement to set up a DSCP classifier for the desired DSCP value and forwarding class mapped to a PFC priority above.

Options

pfc-priority—3-bit PFC priority value in decimal form (0-7).

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Understanding PFC Using DSCP at Layer 3 for Untagged Traffic

Configuring DSCP-based PFC for Layer 3 Untagged Traffic

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

pfe

IN THIS SECTION

- Syntax | 12460
- Syntax | 12461
- Syntax (PTX-Series Devices) | 12461
- Hierarchy Level | 12462
- Description | 12462
- Options | 12462
- Required Privilege Level | 12463
- Release Information | 12463

Syntax

```
pfe pfe-instance {
  exception-reporting {
     category category-name; {
     inline-monitoring-instance inline-monitoring-instance;
     }
  }
  forwarding-packages {
```

```
    mobility;
  }
  power (off | on);
  tunnel-services;
}
```

Syntax

```
  pfepfe-id{
        core number {
            channel number {
                tunnel-services {
                    bandwidth(100g | 10g | 1g | 2.5g | 200g | 20g | 300g | 30g | 400g | 40g |
480g | 50g | 60g | 70g | 80g | 90g | bandwidth);
                }
            }
        }
}
```

Syntax (PTX-Series Devices)

```
pfe pfe-instance {
  exception-reporting {
     category category-name; {
     inline-monitoring-instance inline-monitoring-instance;
     }
  }
  forwarding-packages {
     mobility;
  }
  power (off | on);
  temp-perf-throttle-disable;
  temp-volt-reduction-disable;
  tunnel-services;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number],
[edit chassis lcc name fpc slot-number],
[edit chassis member name fpc slot-number]
```

Description

Configure options for the Packet Forwarding Engine (PFE). Click here to see the platforms supporting `pfe` configuration options.

Options

pfe-instance* or *pfe-id

PFE identifier or PFE number that represents the PFE ASIC ID ranges from 0-4 for PTX10004 and 0-7 for PTX10008 devices.

PFE identifier or PFE number that represents the PFE ASIC ID. The value is 0 for ACX7024, and ACX7024X devices.

PFE identifier or PFE number that represents the PFE ASIC ID. The value is 0 for ACX7100-48L and ACX7100-32C devices.

PFE identifier or PFE number that represents the PFE ASIC ID. The value is 0 or 1 for ACX7509 devices.

PFE identifier or PFE number that represents the PFE ASIC ID. The value is 0 for ACX7348 and ACX7332 devices.

power

Power PFEs on or off. See "pfe power off" on page 12463 for more details.

temp-perf-throttle-disable

Disable temperature based PFE performance throttling. By default, the system detects overheat condition at individual PFE level and gradually reduces the performance of the affected PFE. This results in reduction of power consumption, heat dissipation, PFE operating temperature, and prevents line card shutdown.

- **Default**

  By default, automatic performance throttling is enabled.

temp-volt-reduction-disable

Disable temperature based PFE voltage reduction. By default, the system detects the ASIC temperature at individual PFE level that reaches the safe operating limit and dynamically manages the voltage of the DC-DC converters. This results in reduction of power consumption.

- **Default**

  By default, temperature voltage reduction is enabled.

tunnel-services Tunnel services configuration

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

`exception-reporting` option added in Junos OS Release 21.2R1.

`temp-perf-throttle-disable` option added in Junos OS Evolved Release 21.4R1.

`temp-volt-reduction-disable` option added in Junos OS Evolved Release 22.2R1.

`exception-reporting` option added in Junos OS Evolved Release 22.2R1.

Statement introduced in Junos OS Evolved Release 24.1R1 for ACX7K Series Routers.

pfe power off

IN THIS SECTION

- Syntax | 12464
- Hierarchy Level | 12464
- Description | 12464
- Default | 12464
- Options | 12464
- Required Privilege Level | 12464
- Release Information | 12465

Syntax

```
pfe pfe-instance power (off | on);
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

Configure a Packet Forwarding Engine (PFE) to remain powered off when an FPC comes online or power off or power on a PFE during runtime. Click here to see the platforms supporting `pfe` configuration options.

Default

`on`

Options

pfe *pfe-id*

The ID of the Packet Forwarding Engine to be powered off.

> **NOTE:**
>
> - Powering off a supported PFE using this configuration command triggers an automatic FPC restart on an MX series router.
> - For the MPC10E-15C-MRATE line card, only the Packet Forwarding Engine 2 is applicable.

power (off | on)

Power off or on the selected Packet Forwarding Engine.

- **Default:** on

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Powering Off Packet Forwarding Engines

MPC10E-15C-MRATE

pgcp

IN THIS SECTION

- Syntax | 12465
- Hierarchy Level | 12466
- Description | 12466
- Required Privilege Level | 12466
- Release Information | 12466

Syntax

```
pgcp {
    hint [ hint-strings ];
    ports-per-session ports;
    remotely-controlled;
    transport [ transport-protocols ];
}
```

Hierarchy Level

```
[edit services nat pool nat-pool-name]
```

Description

Specify that the NAT pool is used exclusively by the BGF.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

remotely-controlled and `ports-per-session` statements added in Junos OS Release 8.5.

`hint` statement added in Junos OS Release 9.0.

pgcp-rules

IN THIS SECTION

- Syntax | 12467
- Hierarchy Level | 12467
- Description | 12467
- Options | 12467
- Required Privilege Level | 12467
- Release Information | 12467

Syntax

```
(pgcp-rules rule-name | pgcp-rules-sets rule-set-name);
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the Packet Gateway Control Protocol (PGCP) rules or rule set included in this service set. You can configure multiple rules but only one rule set for each service.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

rule-set-name—Identifier for the set of rules to be included.

Required Privilege Level

interface—To view this statement in the configuration.

interface–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring Service Sets to be Applied to Services Interfaces

pgcp-service

IN THIS SECTION

- Syntax | 12468
- Hierarchy Level | 12468
- Description | 12468
- Options | 12468
- Required Privilege Level | 12469
- Release Information | 12469

Syntax

```
pgcp-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the Packet Gateway Control Protocol (PGCP) that is required for the border gateway function (BGF) feature.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the Packet Gateway Control Protocol (PGCP) process.

- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

phase-adjust-threshold

IN THIS SECTION

- Syntax | 12469
- Hierarchy Level | 12470
- Description | 12470
- Options | 12470
- Required Privilege Level | 12470
- Release Information | 12470

Syntax

```
phase-adjust-threshold phase-offset
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

The phase adjust threshold option is the maximum phase offset that will be adjusted by the G.8275.2 enhanced servo in a phase-aligned state. A higher phase adjust threshold permits a clock to tolerate phase offset changes because of higher network packet delay variation (PDV). This option can be used for PTP applications which do not demand the strict requirements that are typically needed in the case of telecom applications.

Options

phase-offset The maximum phase offset value in nanoseconds (ns).

- Range: 1100 ns through 100000 ns
- Default: 1100 ns

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1 for MX10008.

Supported on ACX7100-32C, ACX7100-48L, and ACX7024 in Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

Configuring the Precision Time Protocol G.8275.2 Enhanced Profile (Telecom Profile)

phase-lock-threshold

IN THIS SECTION

- Syntax | 12471
- Hierarchy Level | 12471
- Description | 12471
- Options | 12471
- Required Privilege Level | 12472
- Release Information | 12472

Syntax

```
phase-lock-threshold phase-offset
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

The phase lock threshold option controls the maximum phase offset with which G.8275.2 enhanced servo enters the phase-aligned state. Phase locking with a higher phase lock threshold leads to the servo and time receiver getting into the phase-aligned state quickly. This higher offset is subsequently adjusted in the phase-aligned state. This option can be used for PTP applications which do not demand the strict requirements that are typically needed in the case of telecom applications.

Options

phase-offset The maximum phase offset value in nanoseconds (ns).

- Range: 100 ns through 1000 ns
- Default: 200 ns

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1 for MX10008.

Supported on ACX7100-32C, ACX7100-48L, and ACX7024 in Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

Configuring the Precision Time Protocol G.8275.2 Enhanced Profile (Telecom Profile)

phy-timestamping

IN THIS SECTION

- Syntax | 12472
- Hierarchy Level | 12473
- Description | 12473
- Required Privilege Level | 12473
- Release Information | 12473

Syntax

```
phy-timestamping
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

Configure timestamping of the IEEE 1588 event packets at the physical layer. Timestamping the packet at the physical layer, also known as PHY timestamping, eliminates the noise or the packet delay variation (PDV) that is introduced by the Packet Forwarding Engine.

By default PHY timestamping is disabled.

On Junos EVO platforms, PHY timestamping is enabled by default and the `phy-timestamping` configuration option is unavailable.

NOTE: MX104 routers and the MPC5E and MPC6E line cards do not support PHY timestamping in WAN PHY mode. If the interface framing is changed from LAN PHY mode to WAN PHY mode or WAN PHY mode to LAN PHY mode on MX104, MPC5E, or MPC6E, you must reboot the MX104 router or the MPC5E and MPC6E line cards. For PHY timestamping to work for MX104 routers and the MPC5E and MPC6E line cards, the packet forwarding engine must come up in LAN PHY mode after a reboot.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

Precision Time Protocol Overview

physical-interface (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 12474
- Hierarchy Level | 12474
- Description | 12474
- Options | 12474
- Required Privilege Level | 12475
- Release Information | 12476

Syntax

```
physical-interface  (flow-bandwidth | flow-control-mode | flow-detection-mode)
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type flow-level-bandwidth],
[edit system ddos-protection protocols protocol-group packet-type flow-level-control],
[edit system ddos-protection protocols protocol-group packet-type flow-level-detection]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure flow bandwidth, flow control mode, or flow detection mode at the physical interface flow aggregation level for the packet type.

Options

flow-bandwidth—Bandwidth for the flow at the physical interface level. Available only at the `[edit system ddos-protection protocols protocol-group packet-type flow-level-bandwidth]` hierarchy level.

- **Default:** 20,000 packets per second

- **Range:** 1 through 50,000 packets per second

flow-control-mode—Mode for how traffic in the detected flow is controlled at the physical interface level. Available only at the `[edit system ddos-protection protocols` *`protocol-group packet-type`* `flow-level-control]` hierarchy level.

> **NOTE**: The configuration at this level overrides the global configuration using the `flow-level-control` statement at the `[edit system ddos-protection global]` hierarchy level.

- `drop`—Drop all traffic in flow.
- `keep`—Keep all traffic in flow.
- `police`—Police the traffic to within its allowed bandwidth.
- **Default:** `drop`

flow-detection-mode—Mode for how flow detection operates at the physical interface level when a policer has been violated. Available only at the `[edit system ddos-protection protocols` *`protocol-group packet-type`* `flow-level-detection]` hierarchy level.

> **NOTE**: The configuration at this level overrides the global configuration using the `flow-detection-mode` statement at the `[edit system ddos-protection global]` hierarchy level.

- `automatic`—Search flows at the physical interface level only when a DDoS policer is being violated and only when the policer violation is not discovered at the finer aggregation levels, logical interface or subscriber. Flows at the physical interface level are subsequently not searched again until a subsequent violation occurs that cannot be found at the subscriber or logical interface levels.
- `off`—Disable flow detection at the physical interface level so that flows are never searched at this level.
- `on`—Search flows at the physical interface level, even when no DDoS protection policer is currently being violated. Monitoring continues at this level regardless of whether a suspect flow is identified at this level.
- **Default:** `automatic`

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring the Maximum Flow Bandwidth at Each Flow Aggregation Level

Configuring How Traffic in a Culprit Flow Is Controlled at Each Flow Aggregation Level

Configuring How Flow Detection Operates at Each Flow Aggregation Level

Setting Up and Using Flow Detection

physical-interface-filter

IN THIS SECTION

- Syntax | 12476
- Hierarchy Level | 12476
- Description | 12477
- Required Privilege Level | 12478
- Release Information | 12478

Syntax

```
physical-interface-filter;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name],
[edit logical-systems logical-system-name firewall family family-name filter filter-name],
[edit routing-instances routing-instance-name firewall family family-name filter filter-name],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name firewall family
family-name filter filter-name]
```

Description

Configure a physical interface filter. Use this statement to reference a physical interface policer for the specified protocol family.

For PTX series routers running Junos OS Release 18.3R1 or later, you can use this command to configure separate firewall filters for different family address types (IPv4 and IPv6) that share the same interface, and configure the same policer as an action for the filter.

To use the aggregate policer, configure the firewall policer you want as `physical-interface-policer`. In addition, at the `firewall family` *`family-name`* `filter` *`filter-name`* hierarchy level where you will reference the policer, make the policer a `physical-interface-filter` firewall filter action. This creates a unique instance of the filter on the physical interface.

The sample configuration shows the settings and relationship between them.

```
firewall {
    policer Shared_Policer {
      physical-interface-policer;
        if-exceeding {
            bandwidth-limit 100m;
            burst-size-limit 500k;
        }
        then {
            discard;
        }
    }
}
```

```
firewall {
    filter Filter_Name {
        physical-interface-filter;
        term term_name {
            then {
                policer Shared_Policer;
                count cinet;
            }
        }
```

```
    }
}
```

```
family inet {
    filter filter_name {
       physical-interface-filter;
        term term_name {
            then {
                policer Shared_Policer;
                count cinet;
            }
        }
    }
}
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support for PTX series routers with third-generation FPCs added in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Two-Color and Three-Color Physical Interface Policers

physical-interface-policer | **12479**

policer (Configuring) | **12630**

physical-interface-policer

IN THIS SECTION

- Syntax | 12479
- Hierarchy Level | 12479
- Description | 12479
- Required Privilege Level | 12480
- Release Information | 12480

Syntax

```
physical-interface-policer;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name],
[edit firewall policer policer-name],
[edit firewall three-color-policer policer-name],
[edit logical-systems logical-system-name firewall policer policer-name],
[edit logical-systems logical-system-name three-color-policer policer-name],
[edit routing-instances routing-instance-name firewall policer policer-name],
[edit routing-instances routing-instance-name firewall three-color-policer policer-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name firewall
policer policer-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name firewall three-
color-policer policer-name]
```

Description

Configure an aggregate policer for a physical interface.

A physical interface policer can be a two-color or three-color policer. When you apply physical interface policer, to different protocol families on the same logical interface, the protocol families share the same policer instance. This means that rate limiting is performed in aggregate for the protocol families for

which the policer is applied. This feature enables you to use a single policer instance to perform aggregate policing for different protocol families on the same physical interface. If you want a policer instance to be associated with a protocol family, the corresponding physical interface filter needs to be applied to that protocol family. The policer is not automatically applied to all protocol families configured on the physical interface.

In contrast, with logical interface policers there are multiple separate policer instances.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support at the `[edit dynamic-profiles ... policer policer-name]` hierarchy level introduced in Junos Release OS 11.4.

Support for PTX series routers with third-generation FPCs added in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

- Two-Color and Three-Color Physical Interface Policers
- physical-interface-filter

pic

IN THIS SECTION

- Syntax (EX4600, OCX1100, QFX Series, QFabric system with ELS) | **12481**
- Syntax (EX4600, QFX Series, QFabric system) | **12481**
- Hierarchy Level (EX4600, OCX1100, QFX Series, QFabric system with ELS) | **12482**
- Hierarchy Level (EX4600, QFX Series, QFabric system) | **12482**
- Description | **12482**

- Options | 12482
- Required Privilege Level | 12483
- Release Information | 12483

Syntax (EX4600, OCX1100, QFX Series, QFabric system with ELS)

```
pic pic-number{
    tunnel-port port-number tunnel-services;
    port port-number{
        channel-speed (speed|disable-auto-speed-detection) ;
    }
    port-range port-range-low port-range-high {
        channel-speed (speed|disable-auto-speed-detection) ;
    }
}
```

Syntax (EX4600, QFX Series, QFabric system)

```
pic pic-number {
    fibre-channel {
        port-range {
            port-range-low port-range-high;
        }
    }
    fte{
        port port-number;
        (port-range port-range-low | port-range-high);
    }
    xe {
        (port port-number | port-range port-range-low port-range-high);
    }
    xle {
        (port port-number | port-range port-range-low port-range-high);
    }
```

Hierarchy Level (EX4600, OCX1100, QFX Series, QFabric system with ELS)

```
[edit chassis fpc slot]
```

Hierarchy Level (EX4600, QFX Series, QFabric system)

```
[edit chassis fpc slot]
```

```
[edit chassis node-group name node-device name]
```

Description

(QFX3500, QFX3600, and QFX5100 standalone switches running Enhanced Layer 2 Software only)—Configure a specific port or a range of ports to operate as 10-Gigabit Ethernet ports or 40-Gigabit Ethernet ports.

Enable the specified port on the Physical Interface Card (PIC) to perform in the specified operating mode.

Options

pic ***pic-number***	(QFX3500 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify **1** to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports. (QFX3600 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify **0** to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports.
port ***physical-port-number***	Port number on which you want to configure the port type.
port-range*port-range-low*	Lowest-numbered port in the range of ports.
port-range*port-range-high*	Highest-numbered port in the range of ports.
channel-speed (*speed* **|disable-auto-speed-detection)**	Configure *10g* for 10-Gigabit Ethernet type ports, and configure disable-auto-speed-detection) to disable auto-channelization.

> **NOTE**: This statement is not supported on the OCX Series.

- On a QFX3500 switch, specify **0** if the port type is `fiber-channel`, and **2** if the port type is `xle`.
- On a QFX3600 switch, specify **0** if the port type is `xe`, and **1** if the port type is `xle`.
- On a QFX5100 switch, specify **0** if the port type is `xe`, and **1** if the port type is `xle` and `fte`.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Options `xe` and `xle` introduced in Junos OS 12.2X50-D20 for the QFX Series.

Option `channel-speed` introduced in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches

Configuring the QSFP+ Port Type on QFX3500 Standalone Switches

Configuring the Port Type on QFX3600 Standalone Switches

Configuring the QSFP+ Port Type on QFX5100 Devices

pic-boot-timeout

IN THIS SECTION

- Syntax | 12484
- Hierarchy Level | 12484
- Description | 12484
- Options | 12484
- Required Privilege Level | 12485
- Release Information | 12485

Syntax

```
pic-boot-timeout (240 | 300);
```

Hierarchy Level

```
[edit interfaces interface-name multiservice-options]
```

Description

Set the service PIC boot timeout value to 240 or 300 seconds for every services PIC on the MX Series router if you are configuring more than 24 members in an aggregated multiservices (AMS) interface. We recommend setting the boot timeout value to 240. Starting with Junos OS Release 16.2, an AMS interface can have up to 32 member interfaces. In Junos OS Release 16.1 and earlier, an AMS interface can have a maximum of 24 member interfaces.

If you do not set the timeout value to 240 or 300 and you configure an AMS interface with more than 24 members, the default value of 180 causes some of the service PICs to go offline.

Options

240 Service PIC boot timeout is set to 240 seconds.

300 Service PIC boot timeout is set to 300 seconds.

- **Default:** 180

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring Aggregated Multiservices Interfaces

pic-console-authentication

IN THIS SECTION

- Syntax | 12486
- Hierarchy Level | 12486
- Description | 12486
- Default | 12486
- Options | 12486
- Required Privilege Level | 12486
- Release Information | 12486

Syntax

```
pic-console authentication {
    (encrypted-password "password");
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure console access to Physical Interface Cards (PICs).

Default

Disabled. By default, there is no password setting for console access.

Options

encrypted-password "*password*"

Use MD5 or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password.

You cannot configure a blank password for `encrypted-password` using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Junos OS to Set Console and Auxiliary Port Properties

Configuring Password Authentication for Console Access to PICs

pic (M Series and T Series Routers)

IN THIS SECTION

- Syntax | 12487
- Hierarchy Level | 12488
- Description | 12488
- Options | 12488
- Required Privilege Level | 12488
- Release Information | 12488

Syntax

```
pic pic-number {
    ce1 {
        e1 port-number {
            channel-group group-number timeslots slot-number;
        }
    }
    ct3 {
        port port-number {
            t1 link-number {
                channel-group group-number timeslots slot-number;
            }
        }
    }
    framing (sdh | sonet);
    idle-cell format {
        itu-t;
        payload-pattern payload-pattern-byte;
```

```
    }
     inline-services {
        bandwidth  (1g | 10g);
    }
    max-queues-per-interface (8 | 4);
    no-concatenate;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

Configure properties for an individual PIC.

Options

pic-number—Slot number in which the PIC is installed.

- **Range:** 0 through 3

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring the Junos OS to Enable SONET/SDH Framing for SONET/SDH PICs
- Configuring the Junos OS to Enable a SONET PIC to Operate in Channelized (Multiplexed) Mode

Configuring the Junos OS to Support Channelized DS3-to-DS0 Naming for Channel Groups and Time Slots

Configuring the Junos OS to Support Channel Groups and Time Slots for Channelized E1 PICs

pic-memory-threshold

IN THIS SECTION

- Syntax | 12489
- Hierarchy Level | 12489
- Description | 12489
- Options | 12489
- Required Privilege Level | 12490
- Release Information | 12490

Syntax

```
pic-memory-threshold percentage percentage;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Specify a PIC memory usage percentage that triggers a system log warning message.

Options

percentage—PIC memory threshold value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Thresholds

pic-mode

IN THIS SECTION

- Syntax | 12490
- Hierarchy Level | 12490
- Description | 12491
- Options | 12491
- Required Privilege Level | 12491
- Release Information | 12491

Syntax

```
pic-mode  mode;
```

Hierarchy Level

```
[edit chassis fpc slot pic pic-number sfpplus]
```

Description

Configure the operating mode for the specified port on the SFP+ uplink module on an EX3200 or EX4200 switch.

On a switch using the SFP+ or SFP+ MACSec uplink module, the `pic-mode` setting defines the speeds of the SFP+ interfaces. When the PIC mode is set to `10g` on the SFP+ or SFP+ MACSec uplink module, uplink ports 0 and 2 support MACSec at 10-Gbps speeds while ports 1 and 3 cannot be used to send any traffic. When the PIC mode is set to `1g`, all four SFP+ ports on the uplink module support MACSec at 1-Gbps speeds.

Options

mode—Operating mode of the SFP+ uplink module:

- **1G**—1-gigabit operating mode
- **10G**—10-gigabit operating mode

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module

pic-mode (MX Series)

IN THIS SECTION

Syntax | 12492

- Hierarchy Level | 12492
- Description | 12492
- Options | 12493
- Required Privilege Level | 12494
- Release Information | 12494

Syntax

```
pic-mode pic-speed;
```

Hierarchy Level

```
[edit chassis fpc slot pic pic-number ]
```

Description

Configure the operating speed of all ports on the MPC7E-MRATE MPC, MIC-MRATE MIC, MX10003 MPC, and MX204 routers.

(MX240, MX480, MX960, MX2010, and MX2020 routers with MPC7E-MRATE) To configure 100 Gbps, 10 Gbps, and 40 Gbps speed on all supported ports, specify `100G`, `10G`, or `40G`, respectively, as the speed for the specified PIC. All the six ports of PIC 0 and PIC 1 of an MPC7E-MRATE MPC support 10-Gbps and 40-Gbps speeds. However, only ports 2 and 5 of PIC 0 and PIC 1 support 100-Gbps speed. Therefore, if you configure `100G` as the operating speed of a PIC, the PIC is rebooted and the ports 0, 1, 3, and 4 are disabled.

(MX2010 and MX2020 routers with MIC-MRATE on MPC8E and MPC9E) To configure the port speed at MIC level or on all supported ports, specify `100G`, `10G`, or `40G`, respectively, as the speed for the MIC-MRATE MIC on MPC8E and MPC9E. All the twelve ports of MIC-MRATE MIC support 10 Gbps and 40 Gbps speeds. When you configure the port speed as 100 Gbps at the PIC level for MPC8E, you can configure only 4 ports of the 12 MIC-MRATE ports on MPC8E to operate at 100 Gbps port speed. The other ports are disabled. Therefore, if you configure `100G` as the operating speed for ports 0, 1, 6, and 7, the other ports are disabled on MPC8E. Similarly, when you configure the port speed as 100 Gbps at the PIC level you can configure only 8 ports of the 12 MIC-MRATE ports on MPC9E with 100 Gbps port speed. Therefore, if you configure `100G` as the operating speed for ports 0, 1, 2, 3, 6,7, 8, and 9, the other

ports can support only 40 Gbps or 10 Gbps. However, enabling port speed of 40 Gbps or 10 Gbps at the PIC level, enables all ports and sets the desired port speed on all ports.

(MX10003 routers with MX10003 MPC) To configure 100 Gbps, 10 Gbps, and 40 Gbps speed on all supported ports, specify 100G, 10G, or 40G, respectively, as the speed for the specified PIC. All the six ports of the fixed port PIC support 10-Gbps and 40-Gbps speeds. All the 12 ports of the Multi-rate MIC support 100-Gbps, 10-Gbps and 40-Gbps speeds. To configure all ports to operate at the same speed, configure rate selectability at the PIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the PIC level, use the `pic-mode` statement and specify the port speed. For more information see MX10003 MPC (Multi-Rate) and Supported Active Physical Ports on MX10003 MPC to Prevent Oversubscription.

(MX204 routers) To configure all ports to operate at the same speed, configure rate selectability at the PIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the PIC level, use the `pic-mode` statement and specify the port speed. The MX204 has four rate-selectable ports (referred to a PIC 0 ports) that can be configured as 100-Gigabit Ethernet ports or 40-Gigabit Ethernet port, or each port can be configured as four 10-Gigabit Ethernet ports (by using a breakout cable). The MX204 also has eight 10-Gigabit Ethernet ports (referred to as PIC 1 ports).

(MX240, MX480, MX960, MX2010, and MX2020 routers with 10x10GE SFPP and 10x1GE SFP MIC on MPC2E-NG and MPC3E-NG) 10x10GE SFPP and 10x1GE SFP MICs consist of 10 ports that support 10-Gbps and 1-Gbps speeds with PTP and MACsec capabilities. To configure all ports to operate at the same speed, configure rate selectability at the PIC level by using the `pic-mode` statement and specifying the port speed, in which case you cannot configure the speed of individual ports.

> **NOTE**: When you insert the 10x10GE SFPP or 10x1GE SFP MIC into slot 0 of the MPC and run the show commands, it appears under PIC 0. But when you insert the MIC into slot 1 of MPC, it appears under PIC 2 even though only one logical PIC is mapped. This indexing is made to accommodate other MICs that have two logical PICs mapped and prevent physical interface (IFD) name collisions.

The MX204 router does not support heterogeneous mode. That is, in PIC mode if 40-Gbps or 100-Gbps speed is configured on PIC 0, then the "number-of-ports" on page 11737 on PIC 1 must be configured to 0 only. For more information, see MX204 Router Overview.

Use the port-checker tool to check whether the combination of ports you want to use is valid or not.

Options

pic-speed—Operating speed of the interfaces configured on the ports of an MPC7E-MRATE MPC.

`100G`—Supported ports operate at 100 Gbps speed.

`40G`—Supported ports operate at 40 Gbps speed.

`10G`—Supported ports operate at 10 Gbps speed.

`1G`—Supported ports operate at 1 Gbps speed.

- **Default:** `10G`

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F4.

Statement introduced in Junos OS Release 17.3R1 for MX10003 routers with the MX10003 MPC.

RELATED DOCUMENTATION

Port Speed

pic-mode (Chassis T1 Mode)

IN THIS SECTION

- Syntax | 12495
- Hierarchy Level | 12495
- Description | 12495
- Options | 12495
- Required Privilege Level | 12495
- Release Information | 12495

Syntax

```
pic-mode (clear-channel);
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number ethernet]
```

Description

Configure normal T1 mode or channelized T1 mode.

Options

- `clear-channel`—(default) Normal T1 mode.
- `ct1`—Channelized T1 mode.

> **NOTE**: When chassis clustering is enabled, it is necessary to indicate in the command which node is being configured. In such circumstances, the `edit chassis fpc command` becomes `edit chassis node` *`node-id`* `fpc`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Interfaces

pic (MX Series Routers)

IN THIS SECTION

- Syntax | 12496
- Hierarchy Level | 12497
- Description | 12497
- Options | 12497
- Required Privilege Level | 12497
- Release Information | 12498

Syntax

```
pic pic-number {
    account-layer2-overhead;
    adaptive-services {
        (layer-2 | layer-3);
    }
    aggregate-ports;
    ce1 {
        e1 port-number {
            channel-group group-number timeslots slot-number;
        }
    }
    channelization;
    ct3 {
        port port-number {
            t1 link-number {
                channel-group group-number timeslots slot-number;
            }
        }
    }
    egress-policer-overhead bytes;
    framing (sdh | sonet);
    idle-cell format {
        itu-t;
        payload-pattern payload-pattern-byte;
```

```
    }
    ingress-policer-overhead bytes;
     inline-services {
        bandwidth  (1g | 10g);
    }
    max-queues-per-interface (8 | 4);
    mlfr-uni-nni-bundles number;
    mlfr-uni-nni-bundles-inline number;
    multi-link-layer-2-inline;
    no-concatenate;
    no-multi-rate;
    pic-type OID of PIC type;
    sparse-dlcis;
    tunnel-services (Chassis) {
        bandwidth (1g | 10g | 20g | 40g);
        tunnel-only;
    }
    vtmapping (klm | itu-t);
}
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

Configure properties for an individual PIC.

Options

pic-number—Slot number in which the PIC is installed.

- **Range:** 0 through 3

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`multi-link-layer-2-inline` and `mlfr-uni-nni-bundles-inline` options introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

- Configuring the Junos OS to Enable SONET/SDH Framing for SONET/SDH PICs
- Configuring the Junos OS to Enable a SONET PIC to Operate in Channelized (Multiplexed) Mode
- Configuring the Junos OS to Support Channelized DS3-to-DS0 Naming for Channel Groups and Time Slots
- Configuring the Junos OS to Support Channel Groups and Time Slots for Channelized E1 PICs
- *Enabling Inline Service Interfaces*

pic (Port)

IN THIS SECTION

- Syntax | **12498**
- Hierarchy Level | **12499**
- Description | **12499**
- Options | **12499**
- Required Privilege Level | **12500**
- Release Information | **12500**

Syntax

```
pic pic-number {
    fte {
        (port port-number | port-range port-range-low port-range-high);
    }
    xe {
```

```
        (port port-number | port-range port-range-low port-range-high);
    }
}
```

Hierarchy Level

```
[edit chassis interconnect-device name fpc slot]
```

Description

(QFX3600 Node device only) Configure port types on QFX3600 Node devices.

> **CAUTION**: The Packet Forwarding Engine on the QFX3600 Node device is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the Node device.

Options

pic *pic-number*	Number of the physical interface card (PIC) on which you want to configure port types. Specify **0** to configure *xe* (10-Gigabit Ethernet) type ports. Specify **1** to configure *fte* (40-gigabit data plane uplink) type ports.
fte	Configure a specific port or a range of ports to operate as 40-gigabit data plane uplink ports.
xe	Configure a specific port or a range of ports to operate as four 10-Gigabit Ethernet ports.
port-number	Port number on which you want to configure the port type. Valid values are 0 through 7 if the port type is `fte`, and 2 through 15 if the port type is `xe`.
port-range-low	Lowest-numbered port in the range of ports. The lowest possible value is 0 if the port type is `fte`. The lowest possible value is 2 if the port type is `xe`.
port-range-high	Highest-numbered port in the range of ports. The highest possible value is 7 if the port type is `fte`. The highest possible value is 15 if the port type is `xe`.

NOTE:

- By default, ports Q0 through Q3 operate as fte type ports, and ports Q4 through Q15 operate as xe type ports.
- Only ports Q0 through Q7 can be configured as fte type ports.
- Only ports Q2 through Q15 can be configured as xe type ports.

NOTE: When you delete the port type configuration for an individual port or a block of ports, the ports return to operating in their default port type.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring the Port Type on QFX3600 Node Devices

pic (PTX Series Routers)

IN THIS SECTION

- Syntax | **12501**
- Hierarchy Level | **12501**
- Description | **12502**
- Options | **12502**
- Required Privilege Level | **12502**

- Release Information | **12502**

Syntax

```
pic
   {
        ce1 {
                e1 port-number
                {
                    channel-group group-number timeslots slot-number;
                }

        ct3 {
            port port-number {
                t1 link-number {
                    channel-group group-number timeslots slot-number;
                }
            }
        }
        framing  (sdh | sonet);
        idle-cell-format {
            itu-t;
            payload-pattern payload-pattern-byte; }

        max-queues-per-interface (8 | 4);
        no-concatenate
        q-pic-large-buffer (large-scale | small-scale);
        no-temperature-monitoring

    }
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

Configure properties for an individual PIC.

Options

pic-number—Slot number in which the PIC is installed.

- **Range:** 0 through 3

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring the Junos OS to Enable SONET/SDH Framing for SONET/SDH PICs
- Configuring the Junos OS to Enable a SONET PIC to Operate in Channelized (Multiplexed) Mode
- Configuring the Junos OS to Support Channelized DS3-to-DS0 Naming for Channel Groups and Time Slots
- Configuring the Junos OS to Support Channel Groups and Time Slots for Channelized E1 PICs

| 7527

pic (satellite-policies port-group-alias)

IN THIS SECTION

Syntax | 12503

- Hierarchy Level | 12503
- Description | 12503
- Options | 12503
- Required Privilege Level | 12503
- Release Information | 12504

Syntax

```
pic pic-number {
    port [port-number | port-number-range | all];
}
```

Hierarchy Level

```
[edit policy-options satellite-policies port-group-alias port-group-alias-name]
```

Description

Specify the PIC number to apply a port group alias for satellite policies in a Junos Fusion.

You must also specify the ports on the PIC when you use this statement.

Options

pic-number The PIC number on the satellite device.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

pic-services-logging

IN THIS SECTION

- Syntax | 12504
- Hierarchy Level | 12504
- Description | 12505
- Options | 12505
- Required Privilege Level | 12505
- Release Information | 12505

Syntax

```
pic-services-logging {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Enable PICs to send special logging information to the Routing Engine for archiving on a hard disk.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the PIC services logging process.
- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

syslog (System) | **28836**

pic (Subscriber Limits)

IN THIS SECTION

- Syntax | 12506
- Hierarchy Level | 12506
- Description | 12506
- Options | 12507
- Required Privilege Level | 12507
- Release Information | 12507

Syntax

```
pic number {
    limit limit;
    port number {
        limit limit;
    }
}
```

Hierarchy Level

```
[edit system services resource-monitor subscribers-limit client-type name fpc slot-number]
```

Description

Configure the maximum number of subscribers of a client type allowed to be logged in on the specified MIC. When that number is reached, subsequent logins on the MIC are denied until the current number of subscribers drops below the maximum allowed. You can also specify the maximum number of subscribers of a client type allowed per port, per MPC, and per chassis.

Options

number MIC number.

- **Range:** 0 through 3

limit Maximum number of subscribers.

- **Range:** 1 through 256,000

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Limiting Subscribers by Client Type and Hardware Element with Resource Monitor

Resource Monitoring for Subscriber Management and Services

pim

IN THIS SECTION

- Syntax | 12508
- Hierarchy Level | 12513
- Description | 12513
- Default | 12513

- Required Privilege Level | 12513
- Release Information | 12513

Syntax

```
pim {
    disable;
    assert-timeout seconds;
    dense-groups {
        addresses;
    }
    dr-election-on-p2p;
    export;
    family (inet | inet6) {
        disable;
    }
    graceful-restart {
        disable;
        no-bidirectional-mode;
        restart-duration seconds;
    }
    import [ policy-names ];
    interface (Protocols PIM)  interface-name {
        family  (inet | inet6) {
            disable;
            }
            bfd-liveness-detection {
                    authentication {
                        algorithm algorithm-name;
                        key-chain key-chain-name;
                    loose-check;
                    detection-time {
                        threshold milliseconds;
                    }
                minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
```

```
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                version (0 | 1 | automatic);
            }

            accept-join-always-from [ policy-name ... ];
            accept-remote-source;
            disable;
            bidirectional {
                df-election {
                    backoff-period milliseconds;
                    offer-period milliseconds;
                    robustness-count number;
                }
            }
            family (inet | inet6) {
                disable;
            }
            hello-interval seconds;
            mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse | sparse-
dense);
            neighbor-policy [ policy-names ];
            override-interval milliseconds;
            priority number;
            propagation-delay milliseconds;
            reset-tracking-bit;
            version version;
        }
        join-load-balance;
        join-prune-timeout;
        mdt {
            data-mdt-reuse;
            group-range multicast-prefix;
            threshold {
                group group-address {
                    source source-address {
                        rate threshold-rate;
                    }
                }
                tunnel-limit limit;
            }
```

```
        }
        mvpn {
            autodiscovery {
                inet-mdt;
            }
        }
        nonstop-routing;
        override-interval milliseconds;
        propagation-delay milliseconds;
        reset-tracking-bit;
        rib-group group-name;
        rp {
            auto-rp {
                (announce | discovery | mapping);
                (mapping-agent-election | no-mapping-agent-election);
            }
            bidirectional {
                address address {
                    group-ranges {
                        destination-ip-prefix</prefix-length>;
                    }
                    hold-time seconds;
                    priority number;
                    }
                }
            bootstrap {
                family (inet | inet6) {
                    export [ policy-names ];
                    import [ policy-names ];
                    priority number;
                }
            }
            bootstrap-import [ policy-names ];
            bootstrap-export [ policy-names ];
            bootstrap-priority number;
            dr-register-policy [ policy-names ];
            embedded-rp {
                group-ranges {
                    destination-ip-prefix</prefix-length>;
                }
                maximum-rps limit;
            }
            group-rp-mapping {
```

```
                family (inet | inet6) {
                    log-interval seconds;
                    maximum limit;
                    threshold value;
                    }
                }
                log-interval seconds;
                maximum limit;
                threshold value;
                }
            }
            local {
                family (inet | inet6) {
                    address address;
                    anycast-pim {
                        rp-set {
                            address address <forward-msdp-sa>;
                        }
                        disable;
                        local-address address;
                    }
                    group-ranges {
                        destination-ip-prefix</prefix-length>;
                    }
                    hold-time seconds;
                    override;
                    priority number;
                }
            }
            register-limit {
                family (inet | inet6) {
                    log-interval seconds;
                    maximum limit;
                    threshold value;
                    }
                }
                log-interval seconds;
                maximum limit;
                threshold value;
                }
            }
            rp-register-policy [ policy-names ];
            spt-threshold {
```

```
                infinity [ policy-names ];
            }
            static {
                address address {
                    override;
                    version version;
                    group-ranges {
                        destination-ip-prefix</prefix-length>;
                    }
                }
            }
        }
        rpf-selection {
            group group-address{
                sourcesource-address{
                    next-hop next-hop-address;
                }
                wildcard-source {
                    next-hop next-hop-address;
                }
            }
            prefix-list prefix-list-addresses {
                source source-address {
                    next-hop next-hop-address;
                }
                wildcard-source {
                    next-hop next-hop-address;
                }
            }
        sglimit {
            family (inet | inet6) {
                log-interval seconds;
                maximum limit;
                threshold value;
                }
            }
            log-interval seconds;
            maximum limit;
            threshold value;
            }
        }
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
```

```
            flag flag <flag-modifier> <disable>;
        }
        tunnel-devices [ mt-fpc/pic/port ];
    }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable PIM on the routing device.

The remaining statements are explained separately. See CLI Explorer.

Default

PIM is disabled on the routing device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

family statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

Configuring PIM Dense Mode Properties

Configuring PIM Sparse-Dense Mode Properties

pim-asm

IN THIS SECTION

- Syntax | 12514
- Hierarchy Level | 12514
- Description | 12514
- Required Privilege Level | 12514
- Release Information | 12515

Syntax

```
pim-asm {
    group-address (Routing Instances) address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel],
[edit routing-instances routing-instance-name provider-tunnel]
```

Description

Specify a Protocol Independent Multicast (PIM) sparse mode provider tunnel for an MBGP MVPN or for a draft-rosen MVPN.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

pim-evpn-gateway

IN THIS SECTION

- Syntax | 12515
- Hierarchy Level | 12515
- Description | 12516
- Options | 12516
- Required Privilege Level | 12516
- Release Information | 12516

Syntax

```
pim-evpn-gateway {
    external-irb irb-interface-name;
    peg-df-election {
      delay-time seconds;
      mod;
      preference {
        use-least;
        value value;
      }
    }
}
```

Hierarchy Level

```
[edit routing-instances name protocols evpn oism]
```

Description

Configure a border leaf device to act as a PIM EVPN gateway (PEG) for optimized intersubnet multicast (OISM).

The PEG device connects an EVPN fabric to an external PIM domain to route multicast traffic:

- From external sources to receivers in the EVPN fabric.
- From sources in the EVPN fabric to external receivers.

You set this option in the Layer 3 virtual routing and forwarding (VRF) routing instances on the EVPN fabric devices.

Options

`external-irb` *`irb-interface-name`*

Specify an IRB interface the device uses in PEG mode to route multicast traffic to and from the external PIM domain.

Include this option on OISM border leaf devices that use the M-VLAN IRB method for external multicast. You don't configure this option for the classic L3 interface method or the non-EVPN IRB method for external multicast.

See *Supported Methods for Multicast Data Transfer to or from an External PIM Domain* for a list of the methods you can use with OISM on a PEG device to connect the EVPN fabric to external multicast sources and receivers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 and Junos OS Evolved Release 22.1R1.

RELATED DOCUMENTATION

Optimized Inter-Subnet Multicast in EVPN Networks

PEG DF Election

pim-snooping

IN THIS SECTION

- Syntax | 12517
- Hierarchy Level | 12517
- Description | 12518
- Default | 12518
- Options | 12518
- Required Privilege Level | 12518
- Release Information | 12518

Syntax

```
pim-snooping {
    no-dr-flood;
        traceoptions{
            file [filename  files | no-word-readable | size | word-readable];
            flag [all | general | hello | join | normal | packets | policy | prune | route |
state | task | timer];
        }
    vlan<vlan-id>{
        no-dr-flood;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name instance-type vpls
protocols],
[edit logical-systems logical-system-name routing-instances instance-name protocols],
[edit routing-instances instance-name protocols]
```

Description

PIM snooping snoops PIM hello and join/prune packets on each interface to find interested multicast receivers and then populates the multicast forwarding tree with the information. PIM snooping is configured on PE routers connected using pseudowires and ensures that no new PIM packets are generated in the VPLS (with the exception of PIM messages sent through LDP on pseudowires). PIM snooping differs from PIM proxying in that PIM snooping floods both the PIM hello and join/prune packets in the VPLS, whereas PIM proxying only floods hello packets.

Default

PIM snooping is disabled on the device.

Options

no-dr-flood Disable default flooding of multicast data on the PIM-designated router port.

traceoptions Configure tracing options for PIM snooping.

vlan *<vlan-id>* Configure PIM snooping parameters for a VLAN.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

PIM Snooping for VPLS

pim-ssm (Provider Tunnel)

IN THIS SECTION

- Syntax | 12519
- Hierarchy Level | 12519
- Description | 12519
- Required Privilege Level | 12519
- Release Information | 12520

Syntax

```
pim-ssm {
    group-address (Routing Instances) address;
    tunnel-source address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel family inet | inet6],
[edit routing-instances routing-instance-name provider-tunnel family inet | inet6]
```

Description

Configure the PIM source-specific multicast (SSM) provider tunnel. Use `family inet6 pim-ssm` for Rosen 7 running on IPv6 . For Rosen 7 on IPv4, use `family inet pim-ssm`. The customer data-MDT can be configured on IPv4 or IPv6, but not both (the provider space always runs on IPv4). Enable Rosen IPv4 before enabling Rosen IPv6.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

In Junos OS Release 17.3R1, the `pim-ssm` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default multicast distribution tree (MDT) in Rosen 7, and data MDT for Rosen 6 and Rosen 7.

RELATED DOCUMENTATION

Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs

pim-ssm (Selective Tunnel)

IN THIS SECTION

- Syntax | 12520
- Hierarchy Level | 12521
- Description | 12521
- Required Privilege Level | 12521
- Release Information | 12521

Syntax

```
pim-ssm {
    group-range multicast-prefix;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group group-address source source-address],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group group-address wildcard-source],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet wildcard-source],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet6 wildcard-source],
[edit routing-instances routing-instance-name provider-tunnel selective group group-address
source source-address],
[edit routing-instances routing-instance-name provider-tunnel selective group group-address
wildcard-source],
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet
wildcard-source],
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet6
wildcard-source]
```

Description

Establish the multicast group address range to use for creating MBGP MVPN source-specific multicast selective PMSI tunnels.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

pim-to-igmp-proxy

IN THIS SECTION

- Syntax | 12522
- Hierarchy Level | 12522
- Description | 12522
- Required Privilege Level | 12523
- Release Information | 12523

Syntax

```
pim-to-igmp-proxy {
    upstream-interface [ interface-names ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-options multicast]
```

Description

Use the **pim-to-igmp-proxy** statement to have Internet Group Management Protocol (IGMP) forward IPv4 multicast traffic across Protocol Independent Multicast (PIM) sparse mode domains.

Configure the rendezvous point (RP) routing device that resides between a customer edge-facing PIM domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding IGMP report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device.

On the IGMP upstream interface(s) used to send proxied PIM traffic, set the IP address so it is the lowest IP on the network to ensure that the proxying router is always the IGMP querier.

Note too that you should not enable PIM on the IGMP upstream interface(s).

The `pim-to-igmp-proxy` statement is not supported for routing instances configured with multicast VPNs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring PIM-to-IGMP Message Translation

pim-to-mld-proxy

IN THIS SECTION

- Syntax | 12523
- Hierarchy Level | 12524
- Description | 12524
- Required Privilege Level | 12524
- Release Information | 12524

Syntax

```
pim-to-mld-proxy {
    upstream-interface [ interface-names ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Description

Configure the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Multicast Listener Discovery (MLD) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the `pim-to-mld-proxy` statement enables you to use MLD to forward IPv6 multicast traffic across the PIM sparse mode domains.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring PIM-to-MLD Message Translation

ping

IN THIS SECTION

- Syntax | 12525
- Syntax (QFX Series) | 12526
- Syntax (Junos OS Evolved) | 12526
- Description | 12527
- Options | 12527
- Required Privilege Level | 12530
- Output Fields | 12530
- Sample Output | 12530
- Release Information | 12533

Syntax

```
ping host
<bypass-routing>
<ce-ip destination-ip-address instance routing-instance-name source-ip source-ip-address>
<count requests>
<do-not-fragment>
<inet | inet6>
<interface source-interface>
<interval seconds>
<no-resolve>
<pattern string>
<rapid>
<record-route>
<routing-instance routing-instance-name>
<logical-system logical-system-name>
<tenant tenant-name>
<size bytes>
<source source-address>
<tos type-of-service>
<ttl value>
```

```
<verbose>
<wait seconds>
```

Syntax (QFX Series)

```
ping host
<bypass-routing>
<count requests>
<detail>
<do-not-fragment>
<inet>
<interface source-interface>
<interval seconds>
<logical-system logical-system-name>
<loose-source value>
<mac-address mac-address>
<no-resolve>
<pattern string>
<rapid>
<record-route>
<routing-instance routing-instance-name>
<size bytes>
<source source-address>
<strict>
< strict-source value>
<tos type-of-service>
<ttl value>
<verbose>
<wait seconds>
```

Syntax (Junos OS Evolved)

```
ping host
<bypass-routing>
<ce-ip destination-ip-address instance routing-instance-name source-ip source-ip-address>
<count requests>
<do-not-fragment>
<inet | inet6>
<interface source-interface>
```

```
<interval seconds>
<no-resolve>
<pattern string>
<rapid>
<record-route>
<routing-instance routing-instance-name>
<size bytes>
<source source-address>
<tos type-of-service>
<ttl value>
<verbose>
<wait seconds>
```

Description

Check host reachability and network connectivity. The `ping` command sends Internet Control Message Protocol (ICMP) ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host. Press Ctrl+c to interrupt a ping command.

Options

host	IP address or hostname of the remote system to ping.
bypass-routing	(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.
ce-ip ***destination-ip-address*** **instance** ***routing-instance-name*** **source-ip** ***source-ip-address***	(MX Series routers with MPC and MIC interfaces only) (Optional) Check the connectivity information of customer edge (CE) devices, such as reachability, attachment points, and MAC addresses, from a provider edge (PE) device in a virtual private LAN service (VPLS), hierarchical VPLS (H-VPLS), and Ethernet VPN (EVPN) network. The `ce-ip` option is based on the LSP ping infrastructure, where the `ping` utility is extended to use the CE device IP address as the target host and the PE device loopback address as the source for a specific VPLS or EVPN routing instance.
destination-ip-address	IPv4 address of the CE device to ping.

instance *routing-instance-name*

Name of the VPLS or EVPN routing instance. The command output displays the connectivity information of the CE device based on the configured routing instance type.

source-ip *source-ip-address*

Loopback address of the PE device.

count *requests*

(Optional) Number of ping requests to send. The range of values is `1` through `2,000,000,000`. The default value is an unlimited number of requests.

detail

(Optional) This option is not supported for Junos OS Evolved Release 18.3R1. Include in the output the interface on which the ping reply was received.

do-not-fragment

(Optional) Set the do-not-fragment (DF) flag in the IP header of the ping packets.

For Junos OS Evolved Release 18.3R1, IPv6 `ping` does not have `do-not-fragment` support. The `ping` command is identified as IPv6 Ping when destination is IPv6 address or `inet6` option is used.

For Junos OS IPv6 packets, this option disables fragmentation.

> **NOTE**: In Junos OS Release 11.1 and later, when issuing the `ping` command for an IPv6 route with the `do-not-fragment` option, the maximum ping packet size is calculated by subtracting 48 bytes (40 bytes for the IPV6 header and 8 bytes for the ICMP header) from the MTU. Therefore, if the ping packet size (including the 48-byte header) is greater than the MTU, the ping operation might fail.

inet

(Optional) Ping Packet Forwarding Engine IPv4 routes.

inet6

(Optional) Ping Packet Forwarding Engine IPv6 routes.

interface *source-interface*

(Optional) Interface to use to send the ping requests.

interval *seconds*

(Optional) How often to send ping requests. The range of values, in seconds, is `1` through infinity. The default value is `1`.

logical-system *logical-system-name*

(Optional) Name of logical system from which to send the ping requests.

Alternatively, enter the `set cli logical-system` *`logical-system-name`* command and then run the `ping` command. To return to the main router or switch, enter the `clear cli logical-system` command.

tenant *tenant-name* (Optional) Name of tenant system from which to send the ping requests.

loose-source *value* (Optional) Intermediate loose source route entry (IPv4). Open a set of values.

mac-address *mac-address* (Optional) Ping the physical or hardware address of the remote system you are trying to reach.

no-resolve (Optional) Do not attempt to determine the hostname that corresponds to the IP address.

pattern *string* (Optional) Specify a hexadecimal fill pattern to include in the ping packet.

rapid (Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the `count` option.

record-route (Optional) Record and report the packet's path (IPv4).

routing-instance *routing-instance-name* (Optional) Name of the routing instance for the ping attempt. For Junos OS Evolved, the `routing-instance` option supports only `mgmt_junos`.

size *bytes* (Optional) Size of ping request packets. The range of values, in bytes, is `0` through `65,468`. The default value is `56`, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

source *source-address*

(Optional) IP address to be set in source ip address field of outgoing echo request. If this option is not specified, the outgoing interface IP is selected as source ip , lo0 otherwise.

strict (Optional) Use the strict source route option (IPv4).

strict-source *value* (Optional) Intermediate strict source route entry (IPv4). Open a set of values.

tos *type-of-service* (Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is `0` through `255`.

If the device configuration includes the `dscp-code-point` *value* statement at the `[edit class-of-service host-outbound-traffic]` hierarchy level, the configured DSCP value overrides the value specified in this command option. In this case, the ToS field of ICMP echo request packets sent on behalf of this command carries the DSCP value specified in the `dscp-code-point` configuration statement instead of the value you specify in this command option.

ttl *value*	(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is `0` through `255`.
verbose	(Optional) Display detailed output.
vpls *instance-name*	(Optional) Ping the instance to which this VPLS belongs.
wait *seconds*	(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is `10` seconds. If this option is used without the count option, a default count of `5` packets is used.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (`!`) indicates that an echo reply was received. A period (`.`) indicates that an echo reply was not received within the timeout period. An `x` indicates that an echo reply was received with an error code. These packets are not counted in the received packets count. They are accounted for separately.

When pinging a nonexistant route, the display output of `ping` command does not print the number of packets sent or received or the packet loss.

Sample Output

ping ce-ip *destination-ip-address* instance *routing-instance-name* source-ip *source-ip-address* (EVPN)

```
user@host> ping ce-ip 10.0.0.4 instance foo source-ip 127.0.0.1
! -> PE5|foo|evpn|ge-0/0/2.100, 00:11:22:33:44:55:66:77:88:99|12:23:ab:98:34:05
! -> PE5|foo|evpn|ge-0/0/2.100, 00:11:22:33:44:55:66:77:88:99|12:23:ab:98:34:05
! -> PE5|foo|evpn|ge-0/0/2.100, 00:11:22:33:44:55:66:77:88:99|12:23:ab:98:34:05
! -> PE5|foo|evpn|ge-0/0/2.100, 00:11:22:33:44:55:66:77:88:99|12:23:ab:98:34:05
! -> PE5|foo|evpn|ge-0/0/2.100, 00:11:22:33:44:55:66:77:88:99|12:23:ab:98:34:05
--- ce-ip ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping ce-ip *destination-ip-address* instance *routing-instance-name* source-ip *source-ip-address* (VPLS)

```
user@host> ping ce-ip 10.0.0.4 instance foo source-ip 127.0.0.1
! -> PE2|foo|vpls|ge-0/0/2.100|12:23:ab:98:34:02
! -> PE2|foo|vpls|ge-0/0/2.100|12:23:ab:98:34:02
! -> PE2|foo|vpls|ge-0/0/2.100|12:23:ab:98:34:02
! -> PE2|foo|vpls|ge-0/0/2.100|12:23:ab:98:34:02
! -> PE2|foo|vpls|ge-0/0/2.100|12:23:ab:98:34:02
--- ce-ip ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping *hostname*

Output for Junos OS Evolved:

```
user@host> ping device1.example.com
PING device1.example.com (192.0.2.0): 56(84) bytes of data.
64 bytes from 192.0.2.0: icmp_seq=1 ttl=64 time=44.7 ms
64 bytes from 192.0.2.0: icmp_seq=2 ttl=64 time=3.76 ms
^C
--- device1.example.com ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 3.765/24.235/44.705/20.470 ms
```

Output for Junos OS:

```
user@host> ping device1.example.com
PING device1.example.com (192.0.2.0): 56 data bytes
64 bytes from 192.0.2.0: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.0.2.0: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.0.2.0: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 1192.0.2.0: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.0.2.0: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.0.2.0: icmp_seq=5 ttl=253 time=1.044 ms
^C --- device1.example.com ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.492/0.641/0.789/0.148 ms
```

ping *hostname* rapid

Output for Junos OS Evolved:

```
user@host> ping device1.example.com rapid
PING device1.example.com (192.0.2.0): 56(84) bytes of data.
..
--- device1.example.com ping statistics ---
5 packets transmitted, 3 received, 40% packet loss, time 505ms
rtt min/avg/max/mdev = 0.026/0.081/0.137/0.044 ms, ipg/ewma 126.258/0.112 ms
```

Output for Junos OS:

```
user@host> ping device1.example.com rapid
PING device1.example.com (192.0.2.0): 56 data bytes
!!!!!
--- device1.example.com ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

ping *hostname* size count

```
user@host> ping device1.example.com size 200 count 5
PING device1.example.com (192.0.2.0): 200 data bytes
208 bytes from 192.0.2.0: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.0.2.0: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.0.2.0: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.0.2.0: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.0.2.0: icmp_seq=4 ttl=253 time=17.898 ms

--- device1.example.com ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

ping *hostname* count size (No route to host)

Output for Junos OS Evolved:

```
user@host> ping 40.0.0.2 count 20 size 500
connect: No route to host
```

Output for Junos OS:

```
user@host> ping 40.0.0.2 count 20 size 500
Aug 02 12:56:56 [INFO ] Step 2: Host and Transit ping has to fail
Aug 02 12:56:56 [TRACE] [R0 evo-ptx-b] [cmd] run ping 40.0.0.2 rapid count 50 size 500
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] PING 40.0.0.2 (40.0.0.2): 500 data bytes
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] ping: sendto: No route to host
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] .ping: sendto: No route to host
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] ping: sendto: No route to host
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] .ping: .sendto: No route to host
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] ping: sendto: No route to host
ug 02 12:57:21 [TRACE] [R0 evo-ptx-b] ..
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] --- 40.0.0.2 ping statistics ---
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] 50 packets transmitted, 0 packets received, 100% packet
loss
```

Release Information

Command introduced before Junos OS Release 7.4.

`ce-ip` option introduced in Junos OS Release 17.3 for MX Series routers with MPC and MIC interfaces.

The following options are deprecated for Junos OS Evolved Release 18.3R1: `detail`, `logical-system`, `loose-source`, `mac-address`, `strict`, `strict-source`, and `vpls`.

The command `tenant` option is introduced in Junos OS Release 19.2R1 for SRX Series.

RELATED DOCUMENTATION

Rate Limit ICMPv4 and ICMPv6 Traffic

Pinging Customer Edge Device IP Address

ping clns

IN THIS SECTION

- Syntax | 12534
- Description | 12534
- Options | 12535
- Required Privilege Level | 12536
- Output Fields | 12536
- Sample Output | 12536
- Release Information | 12536

Syntax

```
ping clns host
<brief>
<count requests>
<detail>
<do-not-fragment>
<interval seconds>
<logical-system logical-system-name>
<no-resolve>
<pattern string>
<rapid>
<routing-instance routing-instance-name>
<size bytes>
<source source-address>
<ttl value>
<verbose>
<wait seconds>
```

Description

Check the reachability of a remote Connectionless Network Service (CLNS) node. Enter **Ctrl+c** to interrupt a `ping clns` command.

Options

host	IP address or hostname of the remote system to ping.
brief	(Optional) Display brief information.
count *requests*	(Optional) Number of ping requests to send. The range of values is 1 through 2,000,000,000. The default is an unlimited number of requests.
detail	(Optional) Include in the output the interface on which the ping reply was received.
do-not-fragment	(Optional) Set the do-not-fragment (DF) bit in the IP header of the ping packets.
interval *seconds*	(Optional) How often to send ping requests. The range of values, in seconds, is 1 through infinity. The default value is 1.
logical-system *logical-system-name*	(Optional) Run the command from the specified logical system.
no-resolve	(Optional) Do not attempt to determine the hostname that corresponds to the IP address.
pattern *string*	(Optional) Specify a hexadecimal fill pattern to include in the ping packet.
rapid	(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of request, include the count option.
routing-instance *routing-instance-name*	(Optional) Name of the routing instance for the ping attempt.
size *bytes*	(Optional) Size of ping request packets. The range of values, in bytes, is 0 through 65,468. The default value is 56, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.
source *source-address*	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface.
ttl *value*	(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is 0 through 255.
verbose	(Optional) Display detailed output.

wait *seconds* (Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is 10 seconds. If this option is used without the count option, a default count of 5 packets is used.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping clns

```
user@host> ping clns 47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00
PING 47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00
(47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00): 55 data bytes
64 bytes from 47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00: seq=0 ttl=30 time=15.051 ms
64 bytes from 47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00: seq=1 ttl=30 time=10.370 ms
64 bytes from 47.0005.9000.f800.0000.0108.0001.1921.6812.4058.00: seq=2 ttl=30 time=10.367 ms
--- ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 10.367/11.929/15.051/2.207 ms
```

Release Information

Command introduced before Junos OS Release 7.4.

ping-death

IN THIS SECTION

- Syntax | 12537
- Hierarchy Level | 12537
- Description | 12537
- Required Privilege Level | 12537
- Release Information | 12537

Syntax

```
ping-death;
```

Hierarchy Level

```
[edit security screen ids-option          screen-name        icmp]
```

Description

Configure the device to detect and reject oversized and irregular ICMP packets. Although the TCP/IP specification requires a specific packet size, many ping implementations allow larger packet sizes. Larger packets can trigger a range of adverse system reactions, including crashing, freezing, and restarting.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

ping-death (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 12538
- Hierarchy Level | 12538
- Description | 12538
- Required Privilege Level | 12539
- Release Information | 12539

Syntax

```
ping-death;
```

Hierarchy Level

```
[edit services screen ids-option screen-name icmp]
```

Description

Identify and drop oversized and irregular ICMP packets, which protects against the ping of death attack. In the ping of death attack, the attacker sends the target ping packets whose IP datagram length (ip_len) exceeds the maximum legal length (65,535 bytes) for IP packets, and the packets are fragmented. When the target attempts to reassemble the IP packets, a buffer overflow might occur, resulting in system crashing, freezing, and restarting.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

ping fabric multicast-flow

IN THIS SECTION

- Syntax | 12539
- Description | 12540
- Options | 12540
- Required Privilege Level | 12540
- Output Fields | 12540
- Sample Output | 12541
- Release Information | 12541

Syntax

```
ping fabric multicast-flow source-fmep-id source-fmep-id fma-name fma-name flow-spec-name flow-
specification-name
<verbose>
```

Description

Send a fabric multicast ping command to the multicast group destination fabric maintenance endpoints (FMEPs) configured in the flow specification. Send a single PDU with a timeout after 1 second.

Options

flow-spec-name *flow-specification-name*	Name of the flow specification that defines the protocol parameters (unicast or multicast) for the fabric ping operation.
fma-name *fma-name*	Name of the FMA.
source-fmep-id *source-fmep-id*	FMEP ID that is the source of the fabric ping operation.
verbose	(Optional) Detailed version of the output display.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Output Fields

Table 93 on page 12540 lists the output fields for the `ping fabric-multicast-flow` command. Output fields are listed in the approximate order in which they appear.

Table 93: ping fabric multicast-flow Output Fields

Field Name	Field Description
`Using fabric flow-specification`	Name of the flow specification.
`Ethernet frame-size`	Ethernet frame size in bytes.
`Source-IP`	MAC address of the source FMEP interface.
`Destination-IP`	MAC address of the destination FMEP interface.

Table 93: ping fabric multicast-flow Output Fields *(Continued)*

Field Name	Field Description
`Protocol`	IP protocol.
`received response from fmep-id`	FMEP ID from which responses are received.

Sample Output

ping fabric multicast-flow

```
user@host> ping fabric multicast-flow source-fmep-id 3 fma-name fma1 flow-spec-name fspec2
verbose

Using fabric flow-specification: fspec2
 Ethernet frame-size: 256
 Source-IP: 87.238.205.21
 Destination-IP: 225.0.0.1 IP
 Protocol: 17
 received response from fmep-id 32775...
 received response from fmep-id 32794...
 received response from fmep-id 32795...
 sent 1 requests, received 3 responses
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | 29924

ping fabric unicast-flow

IN THIS SECTION

- Syntax | 12542
- Description | 12542
- Options | 12542
- Required Privilege Level | 12543
- Output Fields | 12543
- Sample Output | 12544
- Release Information | 12545

Syntax

```
ping fabric unicast-flow fma-name fma-name source-fmep-id source-fmep-id dest-fmep-id flow-spec-name flow-specification-name
<count count>
<forced-fte-interface fte-interface>
<verbose>
```

Description

Send a unicast fabric ping command to the destination fabric maintenance endpoints (FMEPs) specified.

Options

count *count*	(Optional) Number of fabric ping PDUs to send. The maximum PDU count is 5. The default PDU count is 1.
dest-fmep-id *destination-fmep-id*	ID of the destination FMEP of the fabric ping operation.

flow-spec-name *flow-specification-name*	Name of the flow specification that defines the protocol parameters (unicast or multicast) for the fabric ping operation.
fma-name *fma-name*	Name of the FMA.
forced-fte-interface *fte-interface*	(Optional) Forces the fabric ping operation to use the specified FTE interface to inject the PDU into the fabric instead of using the internal forwarding path lookup table to determine the FTE interface.
source-fmep-id *source-fmep-id*	ID of the FMEP that is the source of the fabric ping operation.
verbose	(Optional) Detailed version of the output display.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Output Fields

Table 94 on page 12543 lists the output fields for the `ping fabric-unicast-flow` command. Output fields are listed in the approximate order in which they appear.

Table 94: ping fabric unicast-flow Output Fields

Field Name	Field Description
`Fabric flow ping between source` *`source-fmep-name`* `destination` *`destination-fmep-name`*	Path of the unicast-flow ping between the source and detination FMEPs.
`Using fabric flow-specification`	Name of the flow specification.
`Ethernet frame-size`	Ethernet frame size in bytes.
`Source-MAC`	MAC address of the source FMEP interface.
`Destination-MAC`	MAC address of the destination FMEP interface.

Table 94: ping fabric unicast-flow Output Fields *(Continued)*

Field Name	Field Description
Ethertype	EtherType protocol.
received response from fmep-id	ID of the FMEP sending the response.

Sample Output

ping fabric unicast-flow

```
user@host> ping fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2 flow-spec-name
fspec1 verbose

Fabric flow ping between source ED1494 destination ED1497
Using fabric flow-specification: fspec1
Ethernet frame-size: 256
Source-MAC: 0:2:B1:61:3B:96
Destination-MAC: 0:26:78:90:70:21
Ethertype: 8295
received response from fmep-id 2...
sent 1 requests, received 1 responses
```

ping fabric unicast-flow (With Masking Enabled on the Flow Specification)

```
user@host> ping fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2 flow-spec-name
fspec1 verbose

Fabric flow ping between source ED1494 destination ED1497
Using fabric flow-specification: fspec1
Ethernet frame-size: 256
Source-MAC: 0:1:20:A0:0:1
Destination-MAC: 0:8A:3A:B3:7D:67
Source-MAC Mask: FF:F8:FF:FF:FF:FF
Ethertype: 26785
Current Source-MAC value: 0:1:20:a0:0:1
received response from fmep-id 2...
```

```
Current Source-MAC value: 0:2:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:3:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:4:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:5:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:6:20:a0:0:1
received response from fmep-id 2...
Current Source-MAC value: 0:7:20:a0:0:1
received response from fmep-id 2...
sent 7 requests, received 7 responses
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

traceroute fabric unicast-flow | **29924**

ping-interval

IN THIS SECTION

- Syntax | **12546**
- Hierarchy Level | **12546**
- Description | **12546**

- Options | 12546
- Required Privilege Level | 12547
- Release Information | 12547

Syntax

```
ping-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name oam],
[edit logical-systems logical-system-name routing-instances instance-name protocols l2vpn oam],
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls
neighbor address oam],
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls mesh-
group mesh-group-name neighbor address oam],
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls oam],
[edit protocols l2circuit neighbor address interface interface-name oam],
[edit routing-instances instance-name protocols l2vpn oam],
[edit routing-instances instance-name protocols vpls neighbor address oam],
[edit routing-instances instance-name protocols vpls mesh-group mesh-group-name neighbor address
oam],
[edit routing-instances instance-name protocols vpls oam]
```

Description

Configure the time interval between ping messages for bidirectional forwarding detection (BFD) sessions enabled over pseudowires inside a VPN.

Options

seconds—Time interval between ping messages.

- **Range:** 30 through 3600

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support for FEC 129 VPLS added in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring BFD for VCCV for Layer 2 VPNs, Layer 2 Circuits, and VPLS

Junos OS VPNs Library for Routing Devices

ping mpls bgp

IN THIS SECTION

- Syntax | 12547
- Description | 12548
- Options | 12548
- Additional Information | 12549
- Required Privilege Level | 12550
- Output Fields | 12550
- Sample Output | 12550
- Release Information | 12550

Syntax

```
ping mpls bgp fec
<bottom-label-ttl>
```

```
<count count>
<destination address>
<detail>
<exp forwarding-class>
<instance routing-instance-name>
<logical-system (all | logical-system-name)>
<size bytes>
<source source-address>
<sweep>
```

Description

Check the operability of MPLS BGP-signaled label-switched path (LSP) connections. Press Ctrl+c to interrupt a `ping mpls bgp` command.

NOTE: The `ping mpls bgp` *fec* command only supports single paths.

Options

bottom-label-ttl (Optional) Time-to-live (TTL) value for the bottom label in the label stack. The range of values is 1 through 255. The default value is **255**.

count *count* (Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination *address* (Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail (Optional) Display detailed information about the echo requests sent and received.

exp *forwarding-class* (Optional) Value of the forwarding class for the MPLS ping packets.

fec Ping a BGP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.

instance *routing-instance-name* (Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on the specified logical system.
size *bytes*	(Optional) Size of the LSP ping request packet (88 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.
source *source-address*	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).
stitched-protocol	(Optional) Used for BGP to LDP stitched LSPs to specify the stitching protocol on an intermediate node. In the context of a `ping mpls bgp` command only LDP is supported as the stitched protocol.
sweep	(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the `ping` command continues to be used. You must configure MPLS at the `[edit protocols mpls]` hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only BGP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

NOTE: In a Juniper-Cisco interoperation network scenario, a point-to-multipoint LSP ping echo reply message from a Cisco device in a different IGP area is dropped on the Juniper device when the source address of the reply message is an interface address other than the loopback address or router ID. Starting in Junos OS Release 13.3X8, 14.2R6, 15.1R4, 15.1F6, 15.1F5-S8, 16.1R1, and later releases, such point-to-multipoint LSP ping echo reply messages are accepted by the Juniper device and the messages get logged as uncorrelated responses.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately. To display the error codes, use the **detail** option (for example, **ping mpls bgp 10.255.245.222 detail**).

Sample Output

ping mpls bgp fec count

```
user@host> ping mpls bgp 10.255.245.222 count 10
!!!xxx...x--- lsping statistics ---10 packets transmitted, 3 packets received, 70% packet loss 4
packets received with error status, not counted as received.
```

Release Information

Command introduced in Junos OS Release 11.1.

ping mpls l2circuit

IN THIS SECTION

- Syntax | 12551
- Description | 12551
- Options | 12551
- Additional Information | 12553
- Required Privilege Level | 12553

- Output Fields | 12553
- Sample Output | 12553
- Release Information | 12554

Syntax

```
ping mpls  l2circuit (interface interface-name | virtual-circuit virtual-circuit-id
neighbor address)
<count count>
<destination address>
<detail>
<exp forwarding-class>
<logical-system (all | logical-system-name)>
reply-mode (application-level-control-channel | ip-udp | no-reply)
<size bytes>
<source source-address>
<sweep>
<v1>
```

Description

Check the operability of the MPLS Layer 2 circuit connections. Type Ctrl+c to interrupt a `ping mpls l2circuit` command.

NOTE: This command is not supported on EX4500 and EX4550 switches.

Options

count *count*

(Optional) Number of ping requests to send. If `count` is not specified, five ping requests are sent. The range of values is `1` through `1,000,000`. The default value is `5`.

destination *address*

(Optional) Specify an address other than the default (`127.0.0.1/32`) for the ping echo requests. The address can be anything within the `127/8` subnet.

detail

(Optional) Display detailed information about the echo requests sent and received.

exp *forwarding-class*

(Optional) Value of the forwarding class for the MPLS ping packets.

interface *interface-name*

Ping an interface configured for the Layer 2 circuit on the egress provider edge (PE) router.

logical-system (all | *logical-system-name***)**

(Optional) Perform this operation on all logical systems or on the specified logical system.

reply-mode

(Optional) Reply mode for the ping request. This option has the following suboptions:

application-level-control-channel	Reply using an application level control channel.
ip-udp	Reply using an IPv4 or IPv6 UDP packet.
no-reply	Do not reply to the ping request.

> **NOTE**: The `reply-mode` option and its suboptions `application-level-control-channel`, `ip-udp`, and `no-reply` are also available in Junos OS Release 10.2R4 and 10.3R2.

size *bytes*

(Optional) Size of the label-switched path (LSP) ping request packet (`96` through `65468` bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

source *source-address*

(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (`lo.0`).

sweep

(Optional) Automatically determine the size of the maximum transmission unit (MTU).

v1

(Optional) Use the type 9 Layer 2 circuit type, length, and value (TLV).

virtual-circuit *virtual-circuit-id* **neighbor** *address*

Ping the virtual circuit identifier on the egress PE router or switch and the specified neighbor, testing the integrity of the Layer 2 circuit between the ingress and egress PE routers or switches.

Additional Information

You must configure MPLS at the `[edit protocols mpls]` hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An `x` indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls l2circuit interface

```
user@host> ping mpls l2circuit interface so-1/0/0.1
Request for seq 1, to interface 69, labels <100000, 100208>, packet size 100
Reply for seq 1, return code: Egress-ok, time: 0.439 ms
```

ping mpls l2circuit virtual-circuit detail

```
user@host> ping mpls l2circuit virtual-circuit 200 neighbor 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100048, 100128>, packet size 100

Reply for seq 1, return code: Egress-ok time: 0.539 ms
```

ping mpls l2circuit interface <interface-name> reply-mode

```
user@host> ping mpls l2circuit interface lt-1/2/0.21 reply-mode application-level-control-channel
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Release Information

Command introduced before Junos OS Release 7.4.

The `size` and `sweep` options were introduced in Junos OS Release 9.6.

The `reply-mode` option and its suboptions are introduced in Junos OS Release 10.4R1.

ping mpls l2vpn

IN THIS SECTION

- Syntax | 12554
- Description | 12555
- Options | 12555
- Additional Information | 12556
- Required Privilege Level | 12556
- Output Fields | 12557
- Sample Output | 12557
- Release Information | 12558

Syntax

```
ping mpls l2vpn (instance instance-name  local-site-id local-site-id-number remote-site-id
remote-site-id-number | interface interface-name)
<bottom-label-ttl>
```

```
<count count>
<destination address>
<detail>
<exp forwarding-class>
<logical-system (all | logical-system-name)>
reply-mode (application-level-control-channel | ip-udp | no-reply)
<size bytes>
<source source-address>
<sweep>
```

Description

Check the operability of MPLS Layer 2 virtual private network (VPN) connections. Type Ctrl+c to interrupt a `ping mpls l2vpn` command.

Options

bottom-label-ttl	(Optional) Display the time-to-live value for the bottom label in the label stack.
count *count*	(Optional) Number of ping requests to send. If `count` is not specified, five ping requests are sent. The range of values is `1` through 1,000,000. The default value is `5`.
destination *address*	(Optional) Specify an address other than the default (`127.0.0.1/32`) for the ping echo requests. The address can be anything within the `127/8` subnet.
detail	(Optional) Display detailed information about the echo requests sent and received.
exp *forwarding-class*	(Optional) Value of the forwarding class for the MPLS ping packets.
instance *instance-name* local-site-id *local-site-id-number* remote-site-id *remote-site-id-number*	Ping a combination of the Layer 2 VPN routing instance name, the local site identifier, and the remote site identifier, testing the integrity of the Layer 2 VPN circuit (specified by the identifiers) between the ingress and egress provider edge (PE) routers or switches.
interface *interface-name*	Ping an interface configured for the Layer 2 VPN on the egress PE router or switch.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on the specified logical system.

reply-mode (Optional) Reply mode for the ping request. This option has the following suboptions:

application-level-control-channel	Reply using an application level control channel.
ip-udp	Reply using an IPv4 or IPv6 UDP packet.
no-reply	Do not reply to the ping request.

The `reply-mode` option and its suboptions `application-level-control-channel`, `ip-udp`, and `no-reply` are also available in Junos OS Release 10.2R4 and 10.3R2.

size *bytes* (Optional) Size of the label-switched path (LSP) ping request packet (`96` through `65468` bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

source *source-address* (Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (`lo.0`).

sweep (Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

You must configure MPLS at the `[edit protocols mpls]` hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls l2vpn instance

```
user@host> ping mpls l2vpn instance vpn1 remote-site-id 1 local-site-id 2
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l2vpn instance detail

```
user@host> ping mpls l2vpn instance vpn1 remote-site-id 1 local-site-id 2 detail
Request for seq 1, to interface 68, labels <800001, 100176>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <800001, 100176>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <800001, 100176>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <800001, 100176>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <800001, 100176>
Reply for seq 5, return code: Egress-ok

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l2vpn interface <interface-name> reply-mode

```
user@host> ping mpls l2vpn interface lt-1/2/0.21 reply-mode ip-udp
!!!!!
```

```
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Release Information

Command introduced before Junos OS Release 7.4.

The `size` and `sweep` options were introduced in Junos OS Release 9.6.

The `reply-mode` option and its suboptions are introduced in Junos OS Release 10.4R1.

ping mpls l3vpn

IN THIS SECTION

- Syntax | 12558
- Description | 12559
- Options | 12559
- Additional Information | 12560
- Required Privilege Level | 12560
- Output Fields | 12560
- Sample Output | 12560
- Release Information | 12561

Syntax

```
ping mpls l3vpn prefix prefix-name
<l3vpn-name>
<bottom-label-ttl>
<count count>
<destination address>
<detail>
<exp forwarding-class>
<logical-system (all | logical-system-name)>
```

```
<size bytes>
<source source-address>
<sweep>
```

Description

Check the operability of a MPLS Layer 3 virtual private network (VPN) connection. Type Ctrl+c to interrupt a `ping mpls l3vpn` command.

Options

bottom-label-ttl	(Optional) Display the time-to-live value for the bottom label in the label stack.
count *count*	(Optional) Number of ping requests to send. If `count` is not specified, five ping requests are sent. The range of values is `1` through `1,000,000`. The default value is `5`.
destination *address*	(Optional) Specify an address other than the default (`127.0.0.1/32`) for the ping echo requests. The address can be anything within the `127/8` subnet.
detail	(Optional) Display detailed information about the echo requests sent and received.
exp *forwarding-class*	(Optional) Value of the forwarding class for the MPLS ping packets.
l3vpn-name	(Optional) Layer 3 VPN name.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on the specified logical system.
prefix *prefix-name*	Ping to test whether a prefix is present in a provider edge (PE) router's or switch's VPN routing and forwarding (VRF) table, by means of a Layer 3 VPN destination prefix. This option does not test the connection between a PE router or switch and a customer edge (CE) router or switch.
size *bytes*	(Optional) Size of the label-switched path (LSP) ping request packet (`96` through `65468` bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

source ***source-address***	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (`lo.0`).
sweep	(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

You must configure MPLS at the `[edit protocols mpls]` hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

If the Layer 3 VPN traffic transits a route reflector within the network, the `ping mpls l3vpn` command does not work.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (`!`) indicates that an echo reply was received. A period (`.`) indicates that an echo reply was not received within the timeout period. An `x` indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls l3vpn

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l3vpn detail

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100128, 100112>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <100128, 100112>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <100128, 100112>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <100128, 100112>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <100128, 100112>
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Release Information

Command introduced before Junos OS Release 7.4.

The `size` and `sweep` options were introduced in Junos OS Release 9.6.

ping mpls ldp

IN THIS SECTION

- Syntax | 12562
- Description | 12562
- Options | 12562
- Additional Information | 12563
- Required Privilege Level | 12564
- Output Fields | 12564
- Sample Output | 12564
- Release Information | 12565

Syntax

```
ping mpls ldp fec
<count count>
<destination address>
<detail>
<exp forwarding-class>
<instance routing-instance-name>
<logical-system (all | logical-system-name)>
<p2mp root-addr ip-address lsp-id identifier>
<size bytes>
<source source-address>
<stitched-protocol>
<sweep>
```

Description

Check the operability of MPLS LDP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a `ping mpls` command.

Options

count *count*	(Optional) Number of ping requests to send. If `count` is not specified, five ping requests are sent. The range of values is `1` through `1,000,000`. The default value is `5`.
destination *address*	(Optional) Specify an address other than the default (`127.0.0.1/32`) for the ping echo requests. The address can be anything within the `127/8` subnet.
detail	(Optional) Display detailed information about the echo requests sent and received.
exp *forwarding-class*	(Optional) Value of the forwarding class for the MPLS ping packets.
fec	Ping an LDP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.
instance *routing-instance-name*	(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on the specified logical system.
p2mp root-addr *ip-address* lsp-id *identifier*	(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.
size *bytes*	(Optional) Size of the LSP ping request packet (`88` through `65468` bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.
source *source-address*	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (`lo.0`).
stitched-protocol	(Optional) Used for LDP to BGP stitched LSPs to specify the stitching protocol on an intermediate node. In the context of a `ping mpls ldp` command specify the stitching protocol as BGP, ISIS, or OSPF based on the configuration.
sweep	(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the `ping` command continues to be used. You must configure MPLS at the `[edit protocols mpls]` hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

You can configure the ping interval for the `ping mpls ldp` command by specifying a new time in seconds using the `lsp-ping-interval` statement at the `[edit protocols ldp oam]` hierarchy level. For more information, see the MPLS Applications User Guide.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

> **NOTE**: In a Juniper-Cisco interoperation network scenario, a point-to-multipoint LSP ping echo reply message from a Cisco device in a different IGP area is dropped on the Juniper device when the source address of the reply message is an interface address other than the loopback address or router ID. Starting in Junos OS Release 13.3X8, 14.2R6, 15.1R4, 15.1F6, 15.1F5-S8, 16.1R1, and later releases, such point-to-multipoint LSP ping echo reply messages are accepted by the Juniper device and the messages get logged as uncorrelated responses.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls ldp fec count

```
user@host> ping mpls ldp 10.255.245.222 count 10
!!!xxx...x--- lsping statistics ---10 packets transmitted, 3 packets received, 70% packet loss 4
packets received with error status, not counted as received.
```

ping mpls ldp p2mp root-addr lsp-id

```
user@host>ping mpls ldp p2mp root-addr 10.1.1.1/32 lsp-id 1 count 1
Request for seq 1, to interface 71, no label stack.
Request for seq 1, to interface 70, label 299786
Reply for seq 1, egress 10.1.1.3, return code: Egress-ok, time: 18.936 ms
        Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
        Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
Reply for seq 1, egress 10.1.1.4, return code: Egress-ok, time: 18.936 ms
```

```
        Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
        Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
Reply for seq 1, egress 10.1.1.5, return code: Egress-ok, time: 18.936 ms
        Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
        Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```

Release Information

Command introduced before Junos OS Release 7.4.

`size` and `sweep` options introduced in Junos OS Release 9.6.

`instance` option introduced in Junos OS Release 10.0.

`p2mp`, `root-address`, and `lsp-id` options introduced in Junos OS Release 11.2.

ping mpls lsp-end-point

IN THIS SECTION

- Syntax | 12565
- Description | 12566
- Options | 12566
- Additional Information | 12567
- Required Privilege Level | 12567
- Output Fields | 12567
- Sample Output | 12568
- Release Information | 12568

Syntax

```
ping mpls lsp-end-point prefix-name
<count count>
<destination address>
<detail>
```

```
<exp forwarding-class>
<instance routing-instance-name>
<logical-system (all | logical-system-name)>
<size bytes>
<source source-address>
<sweep>
```

Description

Check the operability of MPLS label-switched path (LSP) endpoint connections. Type Ctrl+c to interrupt a `ping mpls` command.

Options

count *count*	(Optional) Number of ping requests to send. If `count` is not specified, five ping requests are sent. The range of values is `1` through `1,000,000`. The default value is `5`.
destination *address*	(Optional) Specify an address other than the default (`127.0.0.1/32`) for the ping echo requests. The address can be anything within the `127/8` subnet.
detail	(Optional) Display detailed information about the echo requests sent and received.
exp *forwarding-class*	(Optional) Value of the forwarding class for the MPLS ping packets.
instance *routing-instance-name*	(Optional) Ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP connection.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on the specified logical system.
prefix-name	LDP forwarding equivalence class (FEC) prefix or RSVP LSP endpoint address.
size *bytes*	(Optional) Size of the LSP ping request packet. If the endpoint is LDP-based, the minimum size of the packet is `88` bytes. If the endpoint is RSVP-based, the minimum size of the packet is `100` bytes. The maximum size in either case is `65468` bytes.
source *source-address*	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (`lo.0`).

sweep (Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the `ping` command continues to be used. You must configure MPLS at the `[edit protocols mpls]` hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

> **NOTE**: In a Juniper-Cisco interoperation network scenario, a point-to-multipoint LSP ping echo reply message from a Cisco device in a different IGP area is dropped on the Juniper device when the source address of the reply message is an interface address other than the loopback address or router ID. Starting in Junos OS Release 13.3X8, 14.2R6, 15.1R4, 15.1F6, 15.1F5-S8, 16.1R1, and later releases, such point-to-multipoint LSP ping echo reply messages are accepted by the Juniper device and the messages get logged as uncorrelated responses.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (`!`) indicates that an echo reply was received. A period (`.`) indicates that an echo reply was not received within the timeout period. An `x` indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls lsp-end-point detail

```
user@host> ping mpls lsp-end-point 10.255.245.119 detail
Route to end point address is via LDP FEC
Request for seq 1, to interface 67, label 100032
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, label 100032
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 67, label 100032
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 67, label 100032
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 67, label 100032
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Release Information

Command introduced before Junos OS Release 7.4.

The `size` and `sweep` options were introduced in Junos OS Release 9.6.

The `instance` option was introduced in Junos OS Release 10.0.

ping mpls rsvp

IN THIS SECTION

- Syntax | 12569
- Description | 12569
- Options | 12569
- Additional Information | 12571
- Required Privilege Level | 12571

- Output Fields | 12571
- Sample Output | 12572
- Release Information | 12575

Syntax

```
ping mpls rsvp
<lsp-name>
<count count>
<destination address>
<detail>
<dynamic-bypass>
<egress egress-address>
<exp forwarding-class>
<interface interface-name>
<logical-system (all | logical-system-name)>
<manual-bypass>
<multipoint>
<size bytes>
<source source-address>
<standby standby-path-name>
<sweep>
```

Description

Check the operability of MPLS RSVP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a `ping mpls` command.

Options

count *count*

(Optional) Number of ping requests to send. If `count` is not specified, five ping requests are sent. The range of values is `1` through `1,000,000`. The default value is `5`.

destination *address*

(Optional) Specify an address other than the default (`127.0.0.1/32`) for the ping echo requests. The address can be anything within the `127/8` subnet.

detail

(Optional) Display detailed information about the echo requests sent and received.

> **NOTE**: When using the `detail` option, the reported time is based on the system time configured on the local and remote routers. Differences in these system times can result in inaccurate one way ping trip times being reported.
>
> In practice, it is difficult to synchronize the system times of independent Juniper Networks routers with sufficient accuracy to provide a meaningful time value for the `detail` option (even when synchronized using NTP).

dynamic-bypass

(Optional) Ping dynamically generated bypass LSPs, used for protecting other LSPs.

egress *egress-address*

(Optional) Only the specified egress router or switch responds to the ping request.

exp *forwarding-class*

(Optional) Value of the forwarding class for the MPLS ping packets.

interface

(Optional) Specify the name of the interface protected by the manual bypass LSP. This option is only available when you have also used the `manual-bypass` option.

logical-system (all | *logical-system-name***)**

(Optional) Perform this operation on all logical systems or on the specified logical system.

lsp-name

Ping an RSVP-signaled LSP using an LSP name.

manual-bypass

(Optional) Ping manually configured bypass LSPs, used for protecting other LSPs. For this option, you must also specify the interface protected by the manual bypass LSP using the `interface` option.

multipoint

(Optional) Send ping requests to each of the egress routers or switches participating in a point-to-multipoint LSP. You can also include the `egress` option to ping a specific egress router or switch participating in a point-to-multipoint LSP.

size *bytes*

(Optional) Size of the LSP ping request packet (`100` through `65468` bytes). Packets are 4-byte aligned. For example, if you enter a size of 101, 102, 103, or 104, the router or switch uses a size value of 104 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 100-byte minimum.

source *source-address*	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface.
standby *standby-path-name*	(Optional) Name of the standby path.
sweep	(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the `ping` command continues to be used. You must configure MPLS at the `[edit protocols mpls]` hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

> **NOTE**: In a Juniper-Cisco interoperation network scenario, a point-to-multipoint LSP ping echo reply message from a Cisco device in a different IGP area is dropped on the Juniper device when the source address of the reply message is an interface address other than the loopback address or router ID. Starting in Junos OS Release 13.3X8, 14.2R6, 15.1R4, 15.1F6, 15.1F5-S8, 16.1R1, and later releases, such point-to-multipoint LSP ping echo reply messages are accepted by the Juniper device and the messages get logged as uncorrelated responses.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (`!`) indicates that an echo reply was received. A period (`.`) indicates that an echo reply was not received within the timeout period. An `x` indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls rsvp (Echo Reply Received)

```
user@host> ping mpls rsvp test1
!!!!!--- lsping statistics ---5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls rsvp (Echo Reply with Error Code)

```
user@host> ping mpls rsvp test2
!!xxx--- lsping statistics ---5 packets transmitted, 2 packets received, 60% packet loss3
packets received with error status, not counted as received.
```

ping mpls rsvp detail

```
user@host> ping mpls rsvp to-green detail
Request for seq 1, to interface 67, labels <100095, 0, 0>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, labels <100095, 0, 0>
Reply for seq 2, return code: Egress-ok
```

ping mpls rsvp multipoint egress detail count

```
user@host>ping mpls rsvp sample-lsp multipoint egress 192.168.1.3 detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.242 ms
    Local transmit time: 1205310695s 215737us
    Remote receive time: 1205310695s 215979us

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

ping mpls rsvp multipoint detail count

```
user@host>ping mpls rsvp sample-lsp multipoint detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, return code: Unknown TLV, time: 9.877 m Local transmit time: 1205310615s
347317us
 Remote receive time: 1205310615s 357194us
Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.351 ms
 Local transmit time: 1205310615s 347262us
 Remote receive time: 1205310615s 347613us
Reply for seq 1, egress 192.168.1.13, return code: Egress-ok, time: 0.301 ms
 Local transmit time: 1205310615s 347167us
 Remote receive time: 1205310615s 347468us
Timeout for seq 1, egress 192.168.1.1
Timeout for seq 1, egress 192.168.1.4
Timeout for seq 1, egress 192.168.1.14

--- lsping, egress 192.168.1.1 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.4 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.13 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.14 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss
```

ping mpls rsvp destination detail count size

```
user@host>ping mpls rsvp chaser-access destination 192.168.0.1 detail count 1 size 4468

Request for seq 1, to interface 88, label 299984, packet size 4468
```

```
Reply for seq 1, return code: Egress-ok, time: 44.804 ms
        Local transmit time: 2009-03-30 22:05:02 CEST 408.629 ms
        Remote receive time: 2009-03-30 22:05:02 CEST 453.433 ms

--- lsping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

ping mpls rsvp destination detail sweep size

```
user@router> ping mpls rsvp chaser-access destination 192.168.0.1 detail sweep size 4500
Request for seq 1, to interface 86, no label stack., packet size 100
Reply for seq 1, return code: Egress-ok, time: -39.264 ms
        Local transmit time: 2009-04-24 14:05:40 CEST 541.423 ms
        Remote receive time: 2009-04-24 14:05:40 CEST 502.159 ms
Request for seq 2, to interface 86, no label stack., packet size 2300
Reply for seq 2, return code: Egress-ok, time: -38.179 ms
        Local transmit time: 2009-04-24 14:05:41 CEST 544.240 ms
        Remote receive time: 2009-04-24 14:05:41 CEST 506.061 ms
Request for seq 3, to interface 86, no label stack., packet size 4500
Timeout for seq 3
Request for seq 4, to interface 86, no label stack., packet size 3400
Reply for seq 4, return code: Egress-ok, time: -37.545 ms
        Local transmit time: 2009-04-24 14:05:45 CEST 549.953 ms
        Remote receive time: 2009-04-24 14:05:45 CEST 512.408 ms
Request for seq 5, to interface 86, no label stack., packet size 3952
Reply for seq 5, return code: Egress-ok, time: -37.176 ms
        Local transmit time: 2009-04-24 14:05:46 CEST 555.881 ms
        Remote receive time: 2009-04-24 14:05:46 CEST 518.705 ms
Request for seq 6, to interface 86, no label stack., packet size 4228
Reply for seq 6, return code: Egress-ok, time: -36.962 ms
        Local transmit time: 2009-04-24 14:05:47 CEST 561.809 ms
        Remote receive time: 2009-04-24 14:05:47 CEST 524.847 ms
Request for seq 7, to interface 86, no label stack., packet size 4368
Reply for seq 7, return code: Egress-ok, time: -36.922 ms
        Local transmit time: 2009-04-24 14:05:48 CEST 568.738 ms
        Remote receive time: 2009-04-24 14:05:48 CEST 531.816 ms
Request for seq 8, to interface 86, no label stack., packet size 4440
Reply for seq 8, return code: Egress-ok, time: -36.855 ms
        Local transmit time: 2009-04-24 14:05:49 CEST 575.669 ms
        Remote receive time: 2009-04-24 14:05:49 CEST 538.814 ms
Request for seq 9, to interface 86, no label stack., packet size 4476
```

```
Timeout for seq 9
Request for seq 10, to interface 86, no label stack., packet size 4460
Reply for seq 10, return code: Egress-ok, time: -36.906 ms
        Local transmit time: 2009-04-24 14:05:53 CEST 584.382 ms
        Remote receive time: 2009-04-24 14:05:53 CEST 547.476 ms
Request for seq 11, to interface 86, no label stack., packet size 4480
Timeout for seq 11
Request for seq 12, to interface 86, no label stack., packet size 4472
Timeout for seq 12
Request for seq 13, to interface 86, no label stack., packet size 4468
Reply for seq 13, return code: Egress-ok, time: -36.943 ms
        Local transmit time: 2009-04-24 14:06:00 CEST 594.884 ms
        Remote receive time: 2009-04-24 14:06:00 CEST 557.941 ms
Request for seq 14, to interface 86, no label stack., packet size 4476
Timeout for seq 14
Request for seq 15, to interface 86, no label stack., packet size 4472
Timeout for seq 15

--- lsp ping sweep result---
Maximum Transmission Unit (MTU) is 4468 bytes
```

Release Information

Command introduced before Junos OS Release 7.4.

The `egress` and `multipoint` options were introduced in Junos OS Release 9.2.

The `size` and `sweep` options were introduced in Junos OS Release 9.6.

The `dynamic-bypass` and `manual-bypass` options were introduced in Junos OS Release 10.2.

ping mpls segment routing isis

IN THIS SECTION

- Syntax | 12576
- Description | 12576
- Options | 12576

- Required Privilege Level | 12577
- Output Fields | 12578
- Sample Output | 12578
- Release Information | 12579

Syntax

```
ping mpls segment routing isis fec
<algorithm algorithm-id>
<count count>
<destination address>
<detail>
<exp forwarding-class>
<instance routing-instance-name>
<logical-system (all | logical-system-name)>
<p2mp root-addr ip-address lsp-id identifier>
<size bytes>
<source source-address>
<stitched-protocol>
<sweep>
```

Description

Check the operability of MPLS segment routing label-switched path (LSP) connections added by ISIS protocol. Type Ctrl+c to interrupt a `ping mpls segment routing isis` command.

Options

algorithm *algorithm-id*
(Optional) Specify the flexible algorithm identifier. The range of values is `0` through `255`. The default value is `0` which means no algorithm is configured. The *algorithm-id* in conjunction with the prefix is used to distinguish labelled SR paths.

count *count*
(Optional) Number of ping requests to send. If `count` is not specified, five ping requests are sent. The range of values is `1` through `1,000,000`. The default value is `5`.

destination *address*	(Optional) Specify an address other than the default (`127.0.0.1/32`) for the ping echo requests. The address can be anything within the `127/8` subnet.
detail	(Optional) Display detailed information about the echo requests sent and received.
exp *forwarding-class*	(Optional) Value of the forwarding class for the MPLS ping packets.
fec	Ping an LDP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.
instance *routing-instance-name*	(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on the specified logical system.
p2mp root-addr *ip-address* **lsp-id** *identifier*	(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.
size *bytes*	(Optional) Size of the LSP ping request packet (`88` through `65468` bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.
source *source-address*	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (`lo.0`).
stitched-protocol	(Optional) Protocol stitched on intermediate node.
sweep	(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls segment-routing isis

```
user@host>ping mpls segment-routing isis 6.6.6.6
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing isis

```
user@host>ping mpls segment-routing isis 6.6.6.6 detail
Request for seq 1, to interface 333, label 402006, packet size 80
Reply for seq 1, return code: Egress-ok, time: -7539.106 ms
        Local transmit time: 2019-03-06 11:54:15 IST 966.883 ms
        Remote receive time: 2019-03-06 11:54:08 IST 427.777 ms
Request for seq 2, to interface 333, label 402006, packet size 80
Reply for seq 2, return code: Egress-ok, time: -7535.925 ms
        Local transmit time: 2019-03-06 11:54:16 IST 978.890 ms
        Remote receive time: 2019-03-06 11:54:09 IST 442.965 ms
Request for seq 3, to interface 333, label 402006, packet size 80
Reply for seq 3, return code: Egress-ok, time: -7526.134 ms
        Local transmit time: 2019-03-06 11:54:17 IST 972.870 ms
        Remote receive time: 2019-03-06 11:54:10 IST 446.736 ms
Request for seq 4, to interface 333, label 402006, packet size 80
Reply for seq 4, return code: Egress-ok, time: -7539.500 ms
        Local transmit time: 2019-03-06 11:54:18 IST 967.014 ms
        Remote receive time: 2019-03-06 11:54:11 IST 427.514 ms
Request for seq 5, to interface 333, label 402006, packet size 80
Reply for seq 5, return code: Egress-ok, time: -7539.605 ms
        Local transmit time: 2019-03-06 11:54:19 IST 969.893 ms
        Remote receive time: 2019-03-06 11:54:12 IST 430.288 ms
```

```
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing isis (Labelled IS-IS flexible-algorithm IPv4 ping)

```
user@host>ping mpls segment-routing isis 10.10.10.10 algorithm 130
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing isis (Labelled IS-IS flexible-algorithm IPv6 ping)

```
user@host>ping mpls segment-routing isis 10:10:10::10 algorithm 130
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Release Information

Command introduced in Junos OS Release 19.1R1

`algorithm` option introduced in Junos OS Release 23.1R1 and Junos OS Evolved Release 23.1R1.

ping mpls segment routing ospf

IN THIS SECTION

- Syntax | 12580
- Description | 12580
- Options | 12580
- Required Privilege Level | 12581
- Output Fields | 12581

Sample Output | 12582
Release Information | 12583

Syntax

```
ping mpls segment routing ospf fec
algorithm algorithm-id
<count count>
<destination address>
<detail>
<exp forwarding-class>
<instance routing-instance-name>
<logical-system (all | logical-system-name)>
<p2mp root-addr ip-address lsp-id identifier>
<size bytes>
<source source-address>
<stitched-protocol>
<sweep>
```

Description

Check the operability of MPLS segment routing label-switched path (LSP) connections added by ospf protocol. Type Ctrl+c to interrupt a `ping mpls segment routing ospf` command.

Options

algorithm *algorithm-id*
(Optional) Specify the flexible algorithm identifier. The range of values is `0` through `255`. The default value is `0` which means no algorithm is configured. The *algorithm-id* in conjunction with the prefix is used to distinguish labelled SR paths.

count *count*
(Optional) Number of ping requests to send. If `count` is not specified, five ping requests are sent. The range of values is `1` through `1,000,000`. The default value is `5`.

destination *address*
(Optional) Specify an address other than the default (`127.0.0.1/32`) for the ping echo requests. The address can be anything within the `127/8` subnet.

detail	(Optional) Display detailed information about the echo requests sent and received.
exp *forwarding-class*	(Optional) Value of the forwarding class for the MPLS ping packets.
fec	Ping an LDP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.
instance *routing-instance-name*	(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on the specified logical system.
p2mp root-addr *ip-address* **lsp-id** *identifier*	(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.
size *bytes*	(Optional) Size of the LSP ping request packet (`88` through `65468` bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.
source *source-address*	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (`lo.0`).
stitched-protocol	(Optional) Protocol stitched on intermediate node.
sweep	(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (`!`) indicates that an echo reply was received. A period (`.`) indicates that an echo reply was not received within the timeout period. An `x` indicates that an echo reply was received with an error

code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls segment-routing ospf

```
user@host>ping mpls segment-routing ospf 6.6.6.6
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing isis

```
user@host>ping mpls segment-routing ospf 6.6.6.6 detail
Request for seq 1, to interface 333, label 22, packet size 80
Reply for seq 1, return code: Egress-ok, time: -7542.667 ms
        Local transmit time: 2019-03-06 11:55:03 IST 639.914 ms
        Remote receive time: 2019-03-06 11:54:56 IST 97.247 ms
Request for seq 2, to interface 333, label 22, packet size 80
Reply for seq 2, return code: Egress-ok, time: -7543.680 ms
        Local transmit time: 2019-03-06 11:55:04 IST 641.965 ms
        Remote receive time: 2019-03-06 11:54:57 IST 98.285 ms
Request for seq 3, to interface 333, label 22, packet size 80
Reply for seq 3, return code: Egress-ok, time: -7530.457 ms
        Local transmit time: 2019-03-06 11:55:05 IST 639.923 ms
        Remote receive time: 2019-03-06 11:54:58 IST 109.466 ms
Request for seq 4, to interface 333, label 22, packet size 80
Reply for seq 4, return code: Egress-ok, time: -7540.548 ms
        Local transmit time: 2019-03-06 11:55:06 IST 642.870 ms
        Remote receive time: 2019-03-06 11:54:59 IST 102.322 ms
Request for seq 5, to interface 333, label 22, packet size 80
Reply for seq 5, return code: Egress-ok, time: -7540.672 ms
        Local transmit time: 2019-03-06 11:55:07 IST 646.870 ms
        Remote receive time: 2019-03-06 11:55:00 IST 106.198 ms

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing ospf (Labelled OSPF flexible-algorithm IPv4 ping)

```
user@host>ping mpls segment-routing ospf 10.10.10.10 algorithm 140
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Release Information

Command introduced in Junos OS Release 19.1R1

`algorithm` option introduced in Junos OS Release 23.1R1 and Junos OS Evolved Release 23.1R1.

ping mpls segment routing spring-te

IN THIS SECTION

- Syntax | 12583
- Description | 12584
- Options | 12584
- Additional Information | 12587
- Required Privilege Level | 12588
- Output Fields | 12588
- Sample Output | 12588
- Release Information | 12592

Syntax

```
ping mpls segment-routing spring-te
<egress-ip (active | color |  count | detail | exp | instance | logical-system | secondary |
segment-list | size | skip-fec-validation  | source | sweep | tunnel-source) >
<label-stack ( detail | egress | labels | logical-system | nexthop-address | nexthop-interface )
>
```

```
<source-routing-path ( active | color | count | detail | egress-ip | exp | instance | logical-
system | secondary | segment-list | size | skip-fec-validation | source | sweep | tunnel-source>
```

Description

Check the operability of MPLS segment routing label-switched path (LSP) connections added by segment routing traffic-engineered (SPRING-TE) protocol. Type Ctrl+c to interrupt a `ping mpls segment routing spring-te` command.

Options

egress-ip Specify to ping to or install IP address to use when sending probes.

active Specify to use the forwarding path or nexthops from the RIB table.

color Specify the color identifier for the tunnel end-point.

- **Range:** 1 through 4294967295

count *count* (Optional) Number of ping requests to send. If `count` is not specified, five ping requests are sent.

- **Range:** 1 through 1000000
- **Default:** 5

detail (Optional) Display detailed output.

exp *exp* (Optional) Specify the class of service to use when sending probes.

- **Range:** 0 through 7

instance *routing-instance-name* (Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

`logical-system` *`logical-system-name`* (Optional) Specify the name of the logical system.

`secondary` Specify to use configured secondary segment list for the given segment routing path.

`segment-list` Specify the segment list to use when sending probes.

size *bytes*	(Optional) Size of the LSP ping request packet (`88` through `65468` bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.
skip-fec-validation	Specify to skip forwarding equivalence class (FEC) validation or use NIL FEC.
source *source-address*	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (`lo.0`).
sweep	(Optional) Automatically determine the size of the maximum transmission unit (MTU).
tunnel-source	Specify the source protocol used to create tunnel.

label-stack

Specify label stack for ping packets. This option works only if you do a minimum configuration of `source-packet-routing` at the `[edit protocols]` hierarchy level.

detail	(Optional) Display detailed output.
egress	Specify the egress IP address.
labels	Specify the labels in label stack. • **Range:** 16 through 1048575
logical-system *logical-system-name*	(Optional) Specify the name of the logical system.
nexthop-address *address*	Specify the nexthop IP address for the ping packet.
nexthop-interface *interface*	Specify the outgoing interface for the ping packet.

source-routing-path

Specify ping source path routing to use when sending probes.

active	Specify to use the forwarding path or nexthops from the RIB table.
color	Specify the color identifier for the tunnel end-point. • **Range:** 1 through 4294967295

count *count*	(Optional) Number of ping requests to send. If `count` is not specified, five ping requests are sent. • **Range:** 1 through 1000000 • **Default:** 5
detail	(Optional) Display detailed output.
egress-ip	
exp *exp*	(Optional) Specify the class of service to use when sending probes. • **Range:** 0 through 7
instance *routing-instance-name*	(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.
`logical-system` *`logical-system-name`*	(Optional) Specify the name of the logical system.
`secondary`	Specify to use configured secondary segment list for the given segment routing path.
`segment-list`	Specify the segment list to use when sending probes.
size *bytes*	(Optional) Size of the LSP ping request packet (`88` through `65468` bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.
`skip-fec-validation`	Specify to skip forwarding equivalence class (FEC) validation or use NIL FEC.
source *source-address*	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (`lo.0`).
sweep	(Optional) Automatically determine the size of the maximum transmission unit (MTU).
tunnel-source	Specify the source protocol used to create tunnel.

Additional Information

NOTE: With `segment-list` option, valid `tunnel-source` (source protocol used to create tunnel) is only static.

NOTE: `segment-list` option is not present when `tunnel-source` is BGP SR-TE since BGP SR-TE does not have named segment list.

NOTE: With `source-routing-path` name, valid `tunnel-source` are PCEP and static.

NOTE: With `tunnel-source` as PCEP, `secondary` option is not valid as PCEP always sends one primary LSP.

NOTE: FEC validation is performed for `ping mpls segment-routing spring-te` by default.

NOTE: FEC validation is supported when segment-routing traffic-engineering (SR-TE) tunnel has IGP labels (OSPF and IS-IS) only. For all the other labels (for example, static, LDP, RSVP, etc.), you can use `skip-fec-validation` option.

NOTE: For a combination of `egress-ip` and `color` options, valid `tunnel-source` are static and BGP SR-TE.

NOTE: `ping mpls segment-routing spring-te` command with `label-stack` option is supported in operational mode only. It is not supported in configuration mode.

NOTE: Parallel SID, binding SID, LDP over SR-TE is not supported.

NOTE: `ping mpls segment-routing spring-te` command with `label-stack` option adds explicit null label at the bottom of the label stack for interoperability. You can include NIL FEC for label “0” or “router alert label” per OAM standard.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls segment-routing spring-te

```
user@host> ping mpls segment-routing spring-te
Possible completions:
  source-routing-path Source Path routing to use when sending probes
  egress-ip                  To/Install IP address to use when sending probes
  label-stack                Label stack for traceroute packets
```

ping mpls segment-routing spring-te (label-stack)

```
user@host>ping mpls segment-routing spring-te label-stack labels [10153 10152] nexthop-address
23.0.1.2 nexthop-interface ge-0/0/0.0 egress 5.5.5.5 detail

Request for seq 1, to interface 333, labels <0, 10153, 10152>, packet size 88
Reply for seq 1, return code: Egress-ok, time: 4.009 ms
        Local transmit time: 2020-02-06 12:02:52 IST 530.885 ms
        Remote receive time: 2020-02-06 12:02:52 IST 534.894 ms
Request for seq 2, to interface 333, labels <0, 10153, 10152>, packet size 88
Reply for seq 2, return code: Egress-ok, time: 4.117 ms
        Local transmit time: 2020-02-06 12:02:53 IST 530.870 ms
        Remote receive time: 2020-02-06 12:02:53 IST 534.987 ms
Request for seq 3, to interface 333, labels <0, 10153, 10152>, packet size 88
Reply for seq 3, return code: Egress-ok, time: 4.121 ms
        Local transmit time: 2020-02-06 12:02:54 IST 530.915 ms
        Remote receive time: 2020-02-06 12:02:54 IST 535.036 ms
```

```
Request for seq 4, to interface 333, labels <0, 10153, 10152>, packet size 88
Reply for seq 4, return code: Egress-ok, time: 3.979 ms
        Local transmit time: 2020-02-06 12:02:55 IST 531.869 ms
        Remote receive time: 2020-02-06 12:02:55 IST 535.848 ms
Request for seq 5, to interface 333, labels <0, 10153, 10152>, packet size 88
Reply for seq 5, return code: Egress-ok, time: 3.707 ms
        Local transmit time: 2020-02-06 12:02:56 IST 530.960 ms
        Remote receive time: 2020-02-06 12:02:56 IST 534.667 ms

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing spring-te (source-routing-path) (No FEC validation)

```
user@host>ping mpls segment-routing spring-te source-routing-path path_1 egress-ip 5.5.5.5 color
6 skip-fec-validation detail

Request for seq 1, to interface 334, labels <900005, 900041>, packet size 76
Reply for seq 1, return code: Egress-ok, time: -1207.605 ms
        Local transmit time: 2019-04-02 09:29:55 IST 548.038 ms
        Remote receive time: 2019-04-02 09:29:54 IST 340.433 ms
Request for seq 2, to interface 334, labels <900005, 900041>, packet size 76
Reply for seq 2, return code: Egress-ok, time: -1161.555 ms
        Local transmit time: 2019-04-02 09:29:56 IST 544.832 ms
        Remote receive time: 2019-04-02 09:29:55 IST 383.277 ms
Request for seq 3, to interface 334, labels <900005, 900041>, packet size 76
Reply for seq 3, return code: Egress-ok, time: -1206.463 ms
        Local transmit time: 2019-04-02 09:29:57 IST 542.845 ms
        Remote receive time: 2019-04-02 09:29:56 IST 336.382 ms
Request for seq 4, to interface 334, labels <900005, 900041>, packet size 76
Reply for seq 4, return code: Egress-ok, time: -1223.260 ms
        Local transmit time: 2019-04-02 09:29:58 IST 543.824 ms
        Remote receive time: 2019-04-02 09:29:57 IST 320.564 ms
Request for seq 5, to interface 334, labels <900005, 900041>, packet size 76
Reply for seq 5, return code: Egress-ok, time: -1025.098 ms
        Local transmit time: 2019-04-02 09:29:59 IST 548.127 ms
        Remote receive time: 2019-04-02 09:29:58 IST 523.029 ms

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing spring-te (source-routing-path) (FEC validation)

```
user@host>ping mpls segment-routing spring-te source-routing-path path_1 detail

Request for seq 1, to interface 342, labels <800007, 800006, 800005>, packet size 80
Reply for seq 1, return code: Egress-ok, time: 1223.132 ms
        Local transmit time: 2019-05-15 13:19:55 IST 233.274 ms
        Remote receive time: 2019-05-15 13:19:56 IST 456.406 ms
Request for seq 2, to interface 342, labels <800007, 800006, 800005>, packet size 80
Reply for seq 2, return code: Egress-ok, time: 1240.785 ms
        Local transmit time: 2019-05-15 13:19:56 IST 238.590 ms
        Remote receive time: 2019-05-15 13:19:57 IST 479.375 ms
Request for seq 3, to interface 342, labels <800007, 800006, 800005>, packet size 80
Reply for seq 3, return code: Egress-ok, time: 1247.976 ms
        Local transmit time: 2019-05-15 13:19:57 IST 240.568 ms
        Remote receive time: 2019-05-15 13:19:58 IST 488.544 ms
Request for seq 4, to interface 342, labels <800007, 800006, 800005>, packet size 80
Reply for seq 4, return code: Egress-ok, time: 1408.181 ms
        Local transmit time: 2019-05-15 13:19:58 IST 234.108 ms
        Remote receive time: 2019-05-15 13:19:59 IST 642.289 ms
Request for seq 5, to interface 342, labels <800007, 800006, 800005>, packet size 80
Reply for seq 5, return code: Egress-ok, time: 1237.968 ms
        Local transmit time: 2019-05-15 13:19:59 IST 234.587 ms
        Remote receive time: 2019-05-15 13:20:00 IST 472.555 ms

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing spring-te (egress-ip) (No FEC validation)

```
user@host>ping mpls segment-routing spring-te egress-ip 5.5.5.5 color 6 skip-fec-validation
detail

Request for seq 1, to interface 334, labels <900005, 900041>, packet size 76
Reply for seq 1, return code: Egress-ok, time: -1224.444 ms
        Local transmit time: 2019-04-02 09:43:35 IST 618.804 ms
        Remote receive time: 2019-04-02 09:43:34 IST 394.360 ms
Request for seq 2, to interface 334, labels <900005, 900041>, packet size 76
Reply for seq 2, return code: Egress-ok, time: -1222.231 ms
        Local transmit time: 2019-04-02 09:43:36 IST 617.763 ms
```

```
        Remote receive time: 2019-04-02 09:43:35 IST 395.532 ms
Request for seq 3, to interface 334, labels <900005, 900041>, packet size 76
Reply for seq 3, return code: Egress-ok, time: -1091.469 ms
        Local transmit time: 2019-04-02 09:43:37 IST 617.897 ms
        Remote receive time: 2019-04-02 09:43:36 IST 526.428 ms
Request for seq 4, to interface 334, labels <900005, 900041>, packet size 76
Reply for seq 4, return code: Egress-ok, time: -1212.647 ms
        Local transmit time: 2019-04-02 09:43:38 IST 617.791 ms
        Remote receive time: 2019-04-02 09:43:37 IST 405.144 ms
Request for seq 5, to interface 334, labels <900005, 900041>, packet size 76
Reply for seq 5, return code: Egress-ok, time: -1199.348 ms
        Local transmit time: 2019-04-02 09:43:39 IST 617.984 ms
        Remote receive time: 2019-04-02 09:43:38 IST 418.636 ms

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing spring-te (egress-ip) (FEC validation)

```
user@host>ping mpls segment-routing spring-te egress-ip 81.7.7.7 color 5 tunnel-source static
detail
Request for seq 1, to interface 342, labels <800007, 800006, 800005>, packet size 80
Reply for seq 1, return code: Egress-ok, time: 1223.339 ms
        Local transmit time: 2019-05-15 13:29:21 IST 108.718 ms
        Remote receive time: 2019-05-15 13:29:22 IST 332.057 ms
Request for seq 2, to interface 342, labels <800007, 800006, 800005>, packet size 80
Reply for seq 2, return code: Egress-ok, time: 1237.089 ms
        Local transmit time: 2019-05-15 13:29:22 IST 113.581 ms
        Remote receive time: 2019-05-15 13:29:23 IST 350.670 ms
Request for seq 3, to interface 342, labels <800007, 800006, 800005>, packet size 80
Reply for seq 3, return code: Egress-ok, time: 1224.890 ms
        Local transmit time: 2019-05-15 13:29:23 IST 111.563 ms
        Remote receive time: 2019-05-15 13:29:24 IST 336.453 ms
Request for seq 4, to interface 342, labels <800007, 800006, 800005>, packet size 80
Reply for seq 4, return code: Egress-ok, time: 1228.785 ms

        Local transmit time: 2019-05-15 13:29:24 IST 108.765 ms
        Remote receive time: 2019-05-15 13:29:25 IST 337.550 ms
Request for seq 5, to interface 342, labels <800007, 800006, 800005>, packet size 80
Reply for seq 5, return code: Egress-ok, time: 1224.228 ms
        Local transmit time: 2019-05-15 13:29:25 IST 107.678 ms
```

```
        Remote receive time: 2019-05-15 13:29:26 IST 331.906 ms

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Release Information

Command introduced in Junos OS Release 20.2R1.

ping mpls segment routing static egress-ip

IN THIS SECTION

- Syntax | 12592
- Description | 12593
- Options | 12593
- Additional Information | 12595
- Required Privilege Level | 12595
- Output Fields | 12595
- Sample Output | 12595
- Release Information | 12596

Syntax

```
ping mpls segment-routing static egress-ip
<egress-ip (fec | count | destination | detail | exp | instance | logical-system | size | source
| sweep) >
```

```
ping mpls segment-routing spring-te label-stack
<label-stack (count | destination | detail | egress | exp | labels | logical-system | nexthop-
address | nexthop-interface | size | source | sweep) >
```

Description

Use `ping mpls segment-routing static` to check the MPLS segment routing label-switched path (LSP) connections on the prpd programmed static routes. Type Ctrl+c to interrupt a `ping mpls segment-routing static` command.

Options

egress-ip Specify to ping PRPD-programmed labelled switched paths (LSP).

<fec>	Specify to use the IP prefix/length of forwarding equivalence class.
count *count*	(Optional) Specify number of ping requests to send. If `count` is not specified, five ping requests are sent. • **Range:** 1 through 1000000 • **Default:** 5
destination *destination*	(Optional) Specify IP address of destination for echo request.
detail	(Optional) Display detailed output.
exp *exp*	(Optional) Specify the class of service to use when sending probes. • **Range:** 0 through 7
instance *routing-instance-name*	(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.
`logical-system` *`logical-system-name`*	(Optional) Specify the name of the logical system.
size *bytes*	(Optional) Size of the LSP ping request packet (`88` through `65468` bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message displays reminding you of the 88-byte minimum.
source *source-address*	(Optional) IP address of the echo request (outgoing interface). This address is sent in the IP source address field of the ping

	request. If this option is not specified, the default address is usually the loopback interface (`lo.0`).
sweep	(Optional) Automatically determine the size of the maximum transmission unit (MTU).

`label-stack` Specify label stack for ping packets.

count *count*	(Optional) Specify number of ping requests to send. If `count` is not specified, five ping requests are sent. • **Range:** 1 through 1000000 • **Default:** 5
destination *destination*	(Optional) Specify IP address of destination for echo request.
detail	(Optional) Display detailed output.
egress	(Optional) Specify the egress IP address.
exp	(Optional) Specify the forwarding class entry. • **Range:**0 through 7
labels	Specify the labels in label stack. • **Range:** 16 through 1048575
logical-system *logical-system-name*	(Optional) Specify the name of the logical system.
nexthop-address *address*	Specify the nexthop IP address for the ping packet.
nexthop-interface *interface*	Specify the outgoing interface for the ping packet.
size *bytes*	(Optional) Size of the LSP ping request packet (`88` through `65468` bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.
source *source-address*	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this

	option is not specified, the default address is usually the loopback interface (`lo.0`).
sweep	(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

> **NOTE**: NIL FEC validation is supported. The FEC validation for SWAN ping and traceroute command is not supported.

> **NOTE**: SWAN ping command takes label-stack, nexthop-address, and nexthop-interface as input will work only if all the three inputs are given by user in the label-stack command.

> **NOTE**: `ping mpls segment-routing spring-te` command with `label-stack`, nexthop-address and nexthop-interface as input will have vendor private tlv in mpls echo request only if egress-ip is provided by user.

Required Privilege Level

Network

Output Fields

When you enter `ping mpls segment-routing static` command, you can see the feedback on the status of your request. An exclamation point (`!`) indicates that an echo reply. A period (`.`) indicates that an echo reply is not received within the timeout period. An `x` indicates that an echo reply with an error code. Packets with error codes are not counted in the received packets count. They are accounted separately.

Sample Output

ping mpls segment-routing static ?

```
user@host> ping mpls segment-routing static ?
Possible completions:
```

```
  egress-ip                 To install IP address to use when sending probes
```

ping mpls segment-routing static egress-ip

```
user@router>ping mpls segment-routing static egress-ip ::100.100.100.4
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing spring-te (label-stack)

```
user@host>ping mpls segment-routing spring-te label-stack labels 600040 nexthop-address 21.1.2.2
nexthop-interface ge-0/0/2.0
Warning: License key missing; requires 'sr' license
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Release Information

Command introduced in Junos OS Release 22.4R1.

ping overlay

IN THIS SECTION

- Syntax | 12597
- Description | 12597
- Options | 12598
- Required Privilege Level | 12600
- Output Fields | 12600
- Sample Output | 12601

Release Information | 12603

Syntax

```
ping overlay
<tunnel-type>
vni vni
tunnel-src ip-source-address
tunnel-dst ip-destination-address
<mac mac-address>
<count requests>
<ttl value>
<hash source-udp-port>
<hash-source-mac source-mac-address>
<hash-destination-mac destination-mac-address>
<hash-source-address source-IP-address>
<hash-destination-address destination-IP-address>
<hash-vlan vlan-id>
<hash-input-interface input-interface>
<hash-protocol protocol-id>
<hash-source-port source-layer4-port>
<hash-destination-port destination-layer4-port>
```

Description

Verify the presence of the Virtual Extensible LAN (VXLAN) tunnel endpoints (VTEPs), which can originate and terminate VXLAN tunnels, and service connectivity within the context of the overlay VXLAN segment. Use `ping overlay` as a fault detection tool to determine failure within an overlay VXLAN tunnel. Type **Ctrl+c** to interrupt a `ping overlay` command.

NOTE: The `ping overlay` command is not supported for IPv6.

NOTE: The `ping overlay` command is not supported when there are multiple virtual-switch routing instances.

Options

tunnel-type

(Optional) Specify the overlay tunnel type used in a virtualized environment such as: VXLAN, Network Virtualization using Generic Routing Encapsulation (NVGRE), MPLS over User Datagram Protocol (UDP), and MPLS over General Routing Encapsulation (GRE) tunnels.

> **NOTE**: Only VXLAN overlay tunnel types are supported.

vni *vni*

Specify the VNI of the VXLAN overlay segment.

tunnel-src *ip-source-address*

Specify the IP address of the source entity at the end of the tunnel, such as the source VTEP.

tunnel-dst *ip-destination-address*

Specify the IP address of the destination entity at the end of the tunnel, such as a remote VTEP.

mac *mac-address*

(Optional) Include the physical or hardware address on the end host system you are trying to reach.

count *requests*

(Optional) Number of ping requests to send.

For QFX and EX9200 switches, the range of values is **1** through **65,535**. The default value is **10**.

For MX Series routers, the range of values is **1** through **2,000,000,000**. The default value is **5**.

ttl *value*

(Optional) Time-to-live (TTL) value to include in the ping request.

For QFX and EX9200 switches, the range of values is **1** through **255**. The default value is **255**.

For MX Series routers, the range of values is **0** through **255**. The default value is **255**.

hash *source-udp-port*

(Optional) Specify the UDP source port. The range of values is **1** through **65,535**. Usage of this parameter will override all other hash-* parameters used to calculate the UDP source port.

hash-source-mac *source-mac-address*

(Optional) Specify the MAC address of the source end host system.

> **NOTE**: The hash parameters provide values that correspond to a particular data flow that the `ping overlay` command debugs. Based on the values that you specify, the system calculates a VXLAN UDP header source port hash, which is included in the VXLAN UDP header of the overlay ping and traceroute packets. Including the calculated hash in the VXLAN header enables the overlay ping and traceroute packets to emulate data packets in the flow that you are troubleshooting.
>
> When using the hash parameters, we recommend that you specify a value for each parameter. The exception to this guideline is the hash-vlan parameter, which you do not have to use if the source endpoint is not a member of a VLAN. This practice ensures that the overlay ping and traceroute processes are successful and that the output for each command is accurate. If you do not specify a value for one or more of the hash parameters, the system sends an OAM request that might include incorrect hash values and generates a warning message.
>
> Hash computation supports TCP and UDP protocols only.

hash-destination-mac *destination-mac-address*	(Optional) Specify the MAC address of the destination end host system.
hash-source-address *source-IP-address*	(Optional) Specify the IP address of the source end host system.
hash-destination-address *destination-IP-address*	(Optional) Specify the IP address of the destination end host system.
hash-vlan *vlan-id*	(Optional, QFX switches only) Specify the VLAN ID of the end host system.
hash-input-interface *input-interface*	(Optional, QFX switches only) Specify the ingress interface of the flow on the Juniper Networks device.
hash-protocol *protocol-id*	(Optional) Specify the TCP/UDP IP protocol ID. The range of values is **1** through **255**.
hash-source-port *source-layer4-port*	(Optional) Specify the Layer 4 source port. The range of values is **1** through **65,535**.

hash-destination-port *destination-layer4-port* (Optional) Specify the Layer 4 destination port. The range of values is **1** through **65,535**.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. These packets are not counted in the received packets count. They are accounted for separately.

Table 95 on page 12600 lists the output fields for the `ping overlay` command. Output fields are listed in the approximate order in which they appear.

Table 95: ping overlay Output Fields

Field Name	Field Description
vni	The VNI of the VXLAN overlay segment.
tunnel src ip	The IP address of the source end of the tunnel.
tunnel dst ip	The IP address of the destination end of the tunnel.
mac address	The physical or hardware address on the end host system you are trying to reach.
count	Number of ping requests sent.
ttl	TTL value for maximum number of pings.

Table 95: ping overlay Output Fields *(Continued)*

Field Name	Field Description
hash-parameters	The hash parameters provide the input-interface, source MAC address, destination MAC address, source IP address, destination IP address, and the VLAN of the two end hosts within an overlay segment. Hash parameters enable platform-specific hash computation to use as the source port in the outer UDP header.
Request/Response for seq *x* to/ from *address* at *timestamp*	Number of ping request and response counts for determining overlay segments in tunnel.

Sample Output

run ping overlay

```
user@host> ping overlay tunnel-type vxlan vni 100 tunnel-src 192.0.2.10 tunnel-dst 192.0.2.20
mac 00:00:5E:00:53:cc count 5 hash-source-mac 00:00:5E:00:53:aa hash-destination-mac
00:00:5E:00:53:cc hash-source-address 198.51.100.1 hash-destination-address 198.51.100.3 hash-
vlan 150 hash-input-interface xe-0/0/2 hash-protocol 17 hash-source-port 4456 hash-destination-
port 4540
vni 100
        tunnel src ip 192.0.2.10
        tunnel dst ip 192.0.2.20
        mac address 00:00:5E:00:53:cc
        count 5
        ttl 255

        hash-parameters:
                input-ifd-idx 653
                end-host smac 00:00:5E:00:53:aa
                end-host dmac 00:00:5E:00:53:cc
                end-host src ip 198.51.100.1
                end-host dst ip 198.51.100.3
                end-host protocol 17
                end-host l4-src-port 4456
                end-host l4-dst-port 4540
                end-host vlan 150
```

```
Request for seq 1, to 192.0.2.20, at 09-24 19:15:33 PDT.352 msecs

Response for seq 1, from 192.0.2.20, at 09-24 19:15:33 PDT.359 msecs, rtt 11 msecs

  Overlay-segment present at RVTEP 192.0.2.20

      End-System Present


Request for seq 2, to 192.0.2.20, at 09-24 19:15:33 PDT.363 msecs

Response for seq 2, from 192.0.2.20, at 09-24 19:15:33 PDT.370 msecs, rtt 10 msecs

  Overlay-segment present at RVTEP 192.0.2.20

      End-System  Present

Request for seq 3, to 192.0.2.20, at 09-24 19:15:33 PDT.374 msecs

Response for seq 3, from 192.0.2.20, at 09-24 19:15:33 PDT.381 msecs, rtt 10 msecs

  Overlay-segment present at RVTEP 192.0.2.20

      End-System Present


Request for seq 4, to 192.0.2.20, at 09-24 19:15:33 PDT.385 msecs

Response for seq 4, from 192.0.2.20, at 09-24 19:15:33 PDT.392 msecs, rtt 10 msecs

  Overlay-segment present at RVTEP 192.0.2.20

      End-System Present


Request for seq 5, to 192.0.2.20, at 09-24 19:15:33 PDT.396 msecs

Response for seq 5, from 192.0.2.20, at 09-24 19:15:33 PDT.403 msecs, rtt 11 msecs

  Overlay-segment present at RVTEP 192.0.2.20

      End-System Present
```

Release Information

Command introduced in Junos OS Release 14.1X53-D30.

RELATED DOCUMENTATION

Understanding Overlay ping and traceroute Packet Support

Example: Troubleshoot a VXLAN Overlay Network with Overlay Ping and Overlay Traceroute on QFX Series Switches

traceroute overlay | 29976

ping srv6

IN THIS SECTION

- Syntax | 12603
- Syntax | 12604
- Description | 12604
- Options | 12604
- Required Privilege Level | 12605
- Output Fields | 12605
- Sample Output | 12606
- Release Information | 12608

Syntax

```
ping srv6 isis host
<count requests>
<detail>
<instance routing-instance-name>
<interval seconds>
<logical-system logical-system-name>
<pattern hex-string>
```

```
<size bytes>
<source source-address>
<sweep>
<tos type-of-service>
<ttl value>
```

Syntax

```
ping srv6 spring-te sids-stack
<count requests>
<detail>
<instance routing-instance-name>
<interval seconds>
<logical-system logical-system-name>
<nexthop-address IPv6-address>
<nexthop interface interface-name>
<pattern hex-string>
<sids value>
<size bytes>
<source source-address>
<sweep>
<tos type-of-service>
<ttl value>
```

Description

Check the reachability of a Segment Routing with IPv6 data plane (SRv6) network.

Options

host	IPv6 loopback address or SRv6 Segment Identifier (SID) of the remote system to ping.
count *requests*	(Optional) Number of ping requests to send. The range of values is 1 through 1,000,000. The default value is 5.
detail	(Optional) Include in the output the interface on which the ping reply was received.

instance *routing-instance-name*	(Optional) Name of the routing instance for the ping attempt.
interval *seconds*	(Optional) How often to send ping requests. The range of values, in seconds, is 0 through 2,147,483,647.
logical-system *logical-system-name*	(Optional) Run the command from the specified logical system.
nexthop-address *IPv6-address*	Nexthop IPv6 address of the outgoing interface.
nexthop-interface *interface-name*	Interface name of the outgoing interface.
pattern *hex-string*	(Optional) Specify a hexadecimal fill pattern to include in the ping packet.
sids *value*	Specify the SRv6 SIDs in the SRv6 stack. Add a set of values.
size *bytes*	(Optional) Size of ping request packets. The range of values, in bytes, is 0 through 65,468.
source *source-address*	(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface.
sweep	(Optional) Automatically determine the size of the maximum transmission unit (MTU).
tos *type-of-service*	Type-of-Service (TOS) value to include in the ping request (IPv6). The range of values is 0 through 255.
ttl *value*	(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is 0 through 255.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping srv6 is-is

```
user@host> ping srv6 isis 6000::1 count 3
!!!
--- lsping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
```

```
user@host> ping srv6 isis 6000::1 detail count 3
Request for icmp_seq=1, via 333, packet size 88
Reply for icmp_seq=1 from ::21.5.6.2 via 333 hlim: 60, trip-time: 13.132 ms
        Local transmit time: 2020-09-10 10:05:16 IST 414.986 ms
        Remote receive time: 2020-09-10 10:05:16 IST 444.730 ms

--- lsping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

```
user@host> ping srv6 isis 7777::7 detail count 1
Request for icmp_seq=3, via 335, packet size 88
Reply for icmp_seq=3 from 7777::7 via 335 hlim: 60, trip-time: 150.522 ms
        Local transmit time: 2020-09-10 10:06:09 IST 117.285 ms

--- lsping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

ping srv6 spring-te sids-stack

```
user@host> ping srv6 spring-te sids-stack sids [ 6000:1::1 4000:1::1 2000::2 ] nexthop-
address ::21.0.1.2 nexthop-interface ge-0/0/0.0 count 1 detail
Request for icmp_seq=1, via interface 333, packet size 120
Reply for icmp_seq=1 from ::21.5.6.2 via interface 333 hops: 4, trip-time: 160.159 ms
        Local transmit time: 2020-09-22 11:13:06 IST 645.842 ms
        Remote receive time: 2020-09-22 11:13:06 IST 827.070 ms

--- lsping statistics ---
```

```
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max = 160.159/160.159/160.159 ms
```

Ping failure outputs

```
user@host> ping srv6 isis 8000::1
No route to host, exiting.
```

```
user@host> ping srv6 isis 6000::1 detail count 1 ttl 1
Request for icmp_seq=1, via 333, packet size 88
Reply for icmp_seq=1 from ::21.1.2.2: Time to live exceeded

--- lsping statistics ---
1 packets transmitted, 0 packets received, 100% packet loss
```

```
user@host> ping srv6 spring-te sids-stack sids [ 6000:1::1 4000:1::1 2000::2 ] nexthop-
address ::21.0.1.2 nexthop-interface ge-0/0/0.0 count 1 detail
Request for icmp_seq=1, via interface 333, packet size 120
Timeout for seq 1

--- lsping statistics ---
1 packets transmitted, 0 packets received, 100% packet loss
round-trip min/avg/max = 0.000/0.000/0.000 ms
```

```
user@host> ping srv6 spring-te sids-stack sids [ 6000:1::1 4000:1::1 2000::2 ] nexthop-
address ::21.0.1.2 nexthop-interface ge-0/0/0.0 count 1 detail xml
<lsping-results xmlns="http://xml.juniper.net/junos/20.4I0/junos-probe-tests">
<lsping-probe-request>
<lsping-sequence-number>1</lsping-sequence-number>
<lsping-interface-index>333</lsping-interface-index>
<lsping-packet-size>120</lsping-packet-size>
</lsping-probe-request>
<lsping-probe-reply>
<lsping-sequence-number>1</lsping-sequence-number>
<lsping-return-code>Timeout</lsping-return-code>
</lsping-probe-reply>
```

```
<lsping-results-summary>
<lsping-packets-transmitted>1</lsping-packets-transmitted>
<lsping-packets-received>0</lsping-packets-received>
<lsping-packet-loss>100</lsping-packet-loss>
<lsping-error-packets>0</lsping-error-packets>
</lsping-results-summary>
</lsping-results>
```

Release Information

Command introduced in Junos OS Release 21.2.

ping vpls instance

IN THIS SECTION

- Syntax | 12608
- Description | 12609
- Options | 12609
- Required Privilege Level | 12610
- Output Fields | 12610
- Sample Output | 12610
- Release Information | 12610

Syntax

```
ping vpls instance instance-name destination-mac address source-ip address
<bd-name name>
<control-plane-response>
<count number>
<detail>
<learning-vlan-id number>
<logical-system logical-system-name>
```

Description

Check the operability of virtual private LAN service (VPLS) connections. Type Ctrl+c to interrupt a `ping vpls instance` command.

When you issue a `ping vpls instance` command, a chassis MAC address is drawn from the ingress PE router's pool of MAC addresses and used to create the VPLS ping packet. The ping packet is then forwarded to the egress PE router. When the egress PE router receives the ping packet, it learns the MAC address from the VPLS ping packet. The MAC address is added to the egress PE router's MAC table.

The `ping vpls instance` command relies on the LSP ping and trace infrastructure defined in RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures* and further enhancements defined in Internet draft draft-stokes-vkompella-ppvpn-hvpls-oam-02, *Testing Hierarchical Virtual Private LAN Services*.

Options

This statement is supported on the MX Series routers. Feature support is determined by the platform and release that you are using. Use Feature Explorer to determine if a feature is supported on your platform.

instance *instance-name*	Specify the name of the VPLS routing instance.
destination-mac *address*	Specify a destination MAC address for the ping echo requests.
source ip *address*	IP address of the outgoing interface.
bd-name *name*	(Optional) Name of the bridge domain.
control-plane-response	(Optional) Request VPLS OAM responses using the control plane.
count *number*	(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is **1** through **1,000,000**. The default value is **5**.
detail	(Optional) Display detailed information about the echo requests sent and received.
learning-vlan-id *number*	(Optional) Specify a learning VLAN identifier for the ping echo requests. The range of values is **0** through **4094**.
logical-system *logical-system-name*	(Optional) Specify a logical system name for the ping echo requests.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping vpls instance

```
user@host> ping vpls instance red destination-mac 00:89:67:1a:23:6f source-ip 10.255.17.138
! -> sample-router:red:ge-4/1/1.0
! -> sample-router:red:ge-4/1/1.0
! -> sample-router:red:ge-4/1/1.0
! -> sample-router:red:ge-4/1/1.0
! -> sample-router:red:ge-4/1/1.0

--- vpls ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Release Information

Command introduced in Junos OS Release 9.1.

pki

IN THIS SECTION

Syntax | 12611

- Hierarchy Level | 12611
- Description | 12611
- Options | 12612
- Required Privilege Level | 12612
- Release Information | 12612

Syntax

```
pki {
    auto-re-enrollment;
    ca-profile ca-profile-name;
    default-trusted-ca-certs (Security);
    traceoptions;
    trusted-ca-group trusted-ca-group-name {
        ca-profiles ca-profiles;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure an IPsec profile and related options to request digital certificates. The Public Key Infrastructure (PKI) provides an infrastructure for digital certificate management.

You can use ‘trusted-ca’ option to specify ca-profiles that are trusted by the NTS clients. You can specify a trusted ca-group (defined under PKI) or ca-profile by name. This configuration is optional and if it is not specified then, NTP trusts all loaded ca-profiles for NTS. Only client can configure trusted-ca options.

Options

auto-re-enrollment Configure the automatic reenrollment of a local end-entity (EE) certificate.

ca-profile *ca-profile-name* Configure certificate authority (CA) profile.

default-trusted-ca-certs Configure automatic download of default trusted CA certificates.

traceoptions Configure public key infrastructure (PKI) tracing options.

trusted-ca-group *trusted-ca-group-name* Configure trusted certificate authority group.

ca-profiles Name of the CA profiles. You can configure maximum of 20 CA profiles.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

`default-trusted-ca-certs` option is added in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

Basic Elements of PKI in Junos OS

pki (Services)

IN THIS SECTION

Syntax | 12613

- Hierarchy Level | 12614
- Description | 12614
- Options | 12614
- Required Privilege Level | 12614
- Release Information | 12614

Syntax

```
pki {
    auto-re-enrollment {
        certificate-id {
            ca-profile ca-profile-name;
            challenge-password password;
            re-enroll-trigger-time-percentage percentage;
            re-generate-keypair;
            validity-period days;
        }
    }
    ca-profile ca-profile-name {
        ca-identity ca-identity;
        enrollment {
            url url-name;
            retry number-of-enrollment-attempts;
            retry-interval seconds;
        }
        revocation-check {
            disable;
            crl {
                disable on-download-failure;
                refresh-interval hours;
                url {
                    url-name;
                    password;
                }
            }
        }
    }
    traceoptions {
```

```
        file filename <files number> <match regular-expression> <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure an IPsec profile to request digital certificates for J Series Services Routers and Adaptive Services (AS) and MultiServices PICs installed in M Series and T Series routers.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

`revocation-check` and `crl` statements added in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring Digital Certificates for Adaptive Services Interfaces

CLI Explorer

plmn

IN THIS SECTION

- Syntax | 12615
- Hierarchy Level | 12615
- Description | 12615
- Options | 12615
- Required Privilege Level | 12616
- Release Information | 12616

Syntax

```
plmn name {
 mcc mcc;
 mnc mnc;
}
```

Hierarchy Level

```
[edit services agf]
```

Description

Configure the public land mobile networks (PLMNs) that the Access Gateway Function (AGF) supports.

Options

name Public land mobile network (PLMN) name

mcc *mcc*	Mobile country code
mnc *mnc*	Mobile network code

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3R1.

poe

IN THIS SECTION

- Syntax | 12616
- Hierarchy Level | 12618
- Description | 12618
- Required Privilege Level | 12618
- Release Information | 12618

Syntax

For switches other than EX6200 and EX8200 switches:

```
poe {
    guard-band watts;
    interface (all | interface-name) {
        disable;
        maximum-power watts;
        priority (high | low);
        telemetries {
```

```
            disable;
            duration hours;
            interval minutes;
        }
    }
    lldp-priority;
    management (class | static);
    notification-control {
        fpc slot-number {
            disable;
        }
    }
}
```

For a Junos Fusion:

```
poe {
    guard-band watts;
    interface (all-extended | interface-name) {
        disable;
        maximum-power watts;
        priority (high | low);
    }
    management (class | static);
}
```

For EX6200 and EX8200 switches:

```
poe {
    fpc ( all | slot-number) {
        guard-band watts;
        lldp-priority;
        management (class | static);
        maximum-power watts;
    }
    interface (all | interface-name) {
        af-mode;
        disable;
        maximum-power watts;
        priority (high | low);
        telemetries {
```

```
            disable;
            duration hours;
            interval minutes;
        }
    }
    notification-control {
        fpc slot-number {
            disable;
        }
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure PoE options. PoE ports on Juniper network switches provide power to PoE-enabled devices only when straight-through cables are used. Power is not provided when crossover cables are used.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 17.2R1 for a Junos Fusion Provider Edge.

RELATED DOCUMENTATION

Configuring PoE Interfaces on EX Series Switches

Configuring Power over Ethernet in a Junos Fusion

point-to-point

IN THIS SECTION

- Syntax | 12619
- Hierarchy Level | 12619
- Description | 12619
- Default | 12619
- Required Privilege Level | 12619
- Release Information | 12620

Syntax

```
point-to-point;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, configure the interface unit as a point-to-point connection. This is the default connection type.

Default

If you omit this statement, the interface unit is configured as a point-to-point connection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Point-to-Point Connection

multipoint | 11127

point-to-point (MVRP)

IN THIS SECTION

- Syntax | 12620
- Hierarchy Level | 12621
- Description | 12621
- Default | 12621
- Required Privilege Level | 12621
- Release Information | 12621

Syntax

```
point-to-point;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mvrp interface (all | interface-name)],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvrp
interface (all | interface-name)],
```

```
[edit protocols mvrp interface (all | interface-name)],
```

```
[edit routing-instances routing-instance-name protocols mvrp interface (all | interface-name)]
```

Description

(Optional) For Multiple VLAN Registration Protocol (MVRP) configurations, configure an interface to be recognized as a point-to-point connection. If specified, a point-to-point subset of the MRP state machine is used to provide a simpler and more efficient method to accelerate convergence on the network. Point-to-point must be enabled after enabling MVRP for the interface to be recognized as a point-to-point connection.

Default

MVRP is disabled by default.

point-to-point is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers

Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration

Verifying That MVRP Is Working Correctly

Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

Understanding Multiple VLAN Registration Protocol (MVRP)

join-timer (MVRP) | **9434**

leaveall-timer (MVRP) | **9707**

leave-timer (MVRP) | **9705**

registration | **14268**

policer

IN THIS SECTION

- Syntax | **12622**
- Hierarchy Level | **12623**
- Description | **12623**
- Options | **12623**
- Required Privilege Level | **12623**
- Release Information | **12623**

Syntax

```
policer policer-name {
    filter-specific;
    if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
    then {
```

```
        policer-action;
    }
}
```

Hierarchy Level

```
[edit firewall]
```

Description

Configure policer rate limits and actions. To activate a policer, you must include the **policer** action modifier in the `then` statement in a firewall filter term.

Each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or across filters—the policer's implicit counter is used to count packets that are policed in all of these terms. If you want to obtain separate packet counts for each term, use these approaches:

- Configure a unique policer for each term.
- Configure only one policer, but use a unique, explicit counter in each term.

Options

policer-name—Name that identifies the policer. The name can contain letters, numbers, hyphens (-), and can be up to 64 characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Two-Color and Three-Color Policers to Control Traffic Rates

Configuring Firewall Filters

Overview of Policers

policer (Applying to a Logical Interface)

IN THIS SECTION

- Syntax | 12624
- Hierarchy Level | 12624
- Description | 12625
- Options | 12625
- Required Privilege Level | 12625

Syntax

```
policer {
    input policer-name;
    output policer-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number],
[edit interfaces interface-name unit unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit unit-number
family family]
```

Description

Apply a single-rate two-color policer—except for a physical interface policer—to Layer 3 input or output traffic at a logical interface.

- To rate-limit all traffic types, regardless of the protocol family, you can apply a logical interface policer at the logical unit level of a supported interface.
- To rate-limit traffic of a specific protocol family, you can apply a basic two-color policer, a bandwidth policer, or a logical interface policer at the protocol family level of a supported interface.

> **NOTE**: You cannot apply a physical interface policer as part of the interface configuration. You can apply a physical interface policer by referencing the policer from a physical interface filter term.

Options

`input` *`policer-name`*—Name of one policer to evaluate packets received on the interface.

`output` *`policer-name`*—Name of one policer to evaluate packets transmitted on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Single-Rate Two-Color Policer Overview

Bandwidth Policer Overview

Logical Interface (Aggregate) Policer Overview

policer (CFM Firewall)

IN THIS SECTION

- Syntax | 12626
- Hierarchy Level | 12626
- Description | 12626
- Required Privilege Level | 12626
- Release Information | 12627

Syntax

```
policer cfm-policer {
    if-exceeding {
        bandwidth-limit 8k;
        burst-size-limit 2k;
    }
    then discard;
}
```

Hierarchy Level

```
[edit firewall]
```

Description

Attach an explicit policer to CFM sessions.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring Rate Limiting of Ethernet OAM Messages

policer (CFM Global) | 12627

policer (CFM Session) | 12629

policer (CFM Global)

IN THIS SECTION

- Syntax | 12627
- Hierarchy Level | 12628
- Description | 12628
- Options | 12628
- Required Privilege Level | 12628
- Release Information | 12628

Syntax

```
policer {
    all cfm-policer-name;
    continuity-check cfm-policer-name;
    other cfm-policer-name;
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

Description

Specify a policer at the global level to police the CFM traffic belonging to all sessions.

Options

continuity-check *cfm-policer-name*	Police all continuity check packets with the policer specified.
other *cfm-policer-name*	Police all non-continuity check packets with the policer specified.
all *cfm-policer-name*	Police all CFM packets with policer specified. If the `all` option is used, then you cannot specify above two options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring Connectivity Fault Management (CFM)

policer (CFM Session) | **12629**

policer (CFM Session)

IN THIS SECTION

- Syntax | 12629
- Hierarchy Level | 12629
- Description | 12629
- Options | 12629
- Required Privilege Level | 12630
- Release Information | 12630

Syntax

```
policer {
    all cfm-policer-name;
    continuity-check cfm-policer-name;
    other cfm-policer-name;
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain name level number
maintenance-association name]
```

Description

Specify a separate policer to rate-limit packets specific to that session.

Options

- **continuity-check** *cfm-policer-name* Police continuity check packets belonging to this session.

`other` *`cfm-policer-name`*	Police all non-continuity check packets belonging to this session.
`all` *`cfm-policer-name`*	Police all CFM packets belonging to this session. If the `all` option is used, then you cannot specify the above two options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring Connectivity Fault Management (CFM)

policer (CFM Global) | 12627

policer (Configuring)

IN THIS SECTION

- Syntax | 12631
- Hierarchy Level | 12631
- Description | 12631
- Options | 12632
- Required Privilege Level | 12632
- Release Information | 12632

Syntax

```
policer policer-name {
    filter-specific;
    counter {
        counter-id counter-index;}
    if-exceeding {
        bandwidth-limit bps;
        bandwidth-percent number;
        burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
    physical-interface-policer;
    shared-bandwidth-policer;
    then {
        policer-action;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall],
[edit firewall],
[edit logical-systems logical-system-name firewall]
```

Description

Configure policer rate limits and actions. When included at the `[edit firewall]` hierarchy level, the `policer` statement creates a template, and you do not have to configure a policer individually for every firewall filter or interface. To activate a policer, you must include the `policer-action` modifier in the `then` statement in a firewall filter term or on an interface.

You can configure the policer in static firewall filters or dynamic firewall filters in a dynamic client profile or a dynamic service profile.

Except for EX8200 switches, each policer that you configure includes an implicit counter. To obtain term-specific packet counts, configure a policer for each term in the filter that requires policing. For EX8200 switches, configure a policer and associate it with a global management counter using the "counter" on page 4690 option.

Options

policer-action One or more actions to take:

- `discard`—Discard traffic that exceeds the rate limits.
- `forwarding-class` *`class-name`*—Specify the particular forwarding class.
- `loss-priority`—Set the packet loss priority (PLP) to `low`, `medium-low`, `medium-high`, or `high`.

policer-name Name that identifies the policer. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”). Policer names cannot begin with an underscore in the form `__.*`.

then Actions to take on matching packets.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The `out-of-profile` policer action added in Junos OS Release 8.1.

The `logical-bandwidth-policer` statement added in Junos OS Release 8.2.

Logical systems support introduced in Junos OS Release 9.3.

The `physical-interface-policer` statement introduced in Junos OS Release 9.6.

The `shared-bandwidth-policer` statement added in Junos OS Release 11.2.

Support at the `[edit dynamic-profiles ... firewall]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Bandwidth Policer Overview

Configuring Firewall Filters and Policers for VPLS

Configuring Multifield Classifiers
Logical Interface (Aggregate) Policer Overview
Physical Interface Policer Overview
Single-Rate Two-Color Policer Overview
Using Multifield Classifiers to Set Packet Loss Priority
filter (Configuring) | 7127
priority (Schedulers)

policer (CoS)

IN THIS SECTION

- Syntax | 12633
- Hierarchy Level | 12634
- Description | 12634
- Options | 12634
- Required Privilege Level | 12634
- Release Information | 12634

Syntax

```
policer cos-policer-name {
    aggregate {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
    premium {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile]
```

Description

For Gigabit Ethernet IQ , Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, define a CoS policer template to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits. The premium policer is not supported on MX Series routers or for Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

Options

cos-policer-name—Name of one policer to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

policer-drop-probability-low

IN THIS SECTION

- Syntax | 12635
- Hierarchy Level | 12635
- Description | 12635
- Required Privilege Level | 12636
- Release Information | 12636

Syntax

```
policer-drop-probability-low;
```

Hierarchy Level

```
[edit chassis]
```

Description

Reduces the possibility that policers configured on the router might drop packets. For some Juniper Networks routers, policers can mark packets as out-of-specification in accordance with TCP. By default, these policers begin to randomly drop packets when the current credit exceeds the credit limit. In the context of TCP, this random drop mechanism helps to smooth the flow of traffic. The `policer-drop-probability-low` statement causes the policers to operate as strict rate limiters and to ignore the standard TCP behavior.

The `policer-drop-probability-low` statement is applicable to the following routing platforms:

- M7i
- M10i
- M120
- M320

- MX Series

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4R1.

RELATED DOCUMENTATION

show pfe cfeb

show pfe feb

show pfe fpc

policer (Firewall Filter Action)

IN THIS SECTION

- Syntax | 12636
- Hierarchy Level | 12637
- Description | 12637
- Options | 12637
- Required Privilege Level | 12637
- Release Information | 12637

Syntax

```
policer policer-name;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name then],
[edit logical-systems logical-system-name firewall family family-name filter filter-name term term-name then]
```

Description

For T Series routers and M320 routers with Enhanced II Flexible PIC Concentrators (FPCs) and the T640 Core Router with Enhanced Scaling FPC4, apply a tricolor marking policer.

Options

policer-name—Name of a single-rate two-color policer to use to rate-limit traffic.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Firewall Filter Nonterminating Actions

Two-Color Policer Configuration Overview

policer (Interface)

IN THIS SECTION

- Syntax | 12638
- Hierarchy Level | 12638
- Description | 12638
- Options | 12638
- Required Privilege Level | 12639
- Release Information | 12639

Syntax

```
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family]
```

Description

Apply a policer to an interface.

Options

`arp` *`policer-template-name`*—For `inet` family only, name of one policer to evaluate when ARP packets are received on the interface.

`input` *`policer-template-name`*—Name of one policer to evaluate when packets are received on the interface.

`output` *`policer-template-name`*—Name of one policer to evaluate when packets are transmitted on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for MX Series routers added in Junos OS Release 20.2R1

RELATED DOCUMENTATION

Applying Policers

Configuring Firewall Filters and Policers for VPLS

Routing Policies, Firewall Filters, and Traffic Policers User Guide

Junos OS Services Interfaces Library for Routing Devices

policer (Layer 2 VPN)

IN THIS SECTION

- Syntax | 12640
- Hierarchy Level | 12640
- Description | 12640
- Options | 12640
- Required Privilege Level | 12640
- Release Information | 12640

Syntax

```
policer {
    input policer-template-name;
    output policer-template-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family (ccc | inet | tcc)],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family (ccc | inet | tcc)]
```

Description

Use policing to control the amount of traffic flowing over the interfaces servicing a Layer 2 VPN.

Options

`input` *`policer-template-name`*—Name of one policer to evaluate when packets are received on the interface.

`output` *`policer-template-name`*—Name of one policer to evaluate when packets are transmitted on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Traffic Policing in Layer 2 VPNs

Routing Policies, Firewall Filters, and Traffic Policers User Guide

policer (MAC)

IN THIS SECTION

- Syntax | 12641
- Hierarchy Level | 12641
- Description | 12641
- Options | 12642
- Required Privilege Level | 12642
- Release Information | 12642

Syntax

```
policer {
    input cos-policer-name;
    output cos-policer-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number accept-source-mac mac-address mac-address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
accept-source-mac mac-address mac-address]
```

Description

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, configure MAC policing.

NOTE: On MX Series routers with Gigabit Ethernet or Fast Ethernet PICs, the following considerations apply:

- Interface counters do not count the 7-byte preamble and 1-byte frame delimiter in Ethernet frames.
- In MAC statistics, the frame size includes MAC header and CRC before any VLAN rewrite/imposition rules are applied.
- In traffic statistics, the frame size encompasses the L2 header without CRC after any VLAN rewrite/imposition rule.

Options

`input` *`cos-policer-name`*—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

`output` *`cos-policer-name`*—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

policer-overhead

IN THIS SECTION

- Syntax | 12643
- Hierarchy Level | 12643
- Description | 12643
- Options | 12644
- Required Privilege Level | 12644
- Release Information | 12644

Syntax

```
policer-overhead {
    egress bytes;
    ingress bytes;
    policer-overhead-value;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit number]
```

Description

Policer overhead adjustment for the logical interface. The adjustment can be made to ingress policers, egress policers, or both.

NOTE: The policer overhead adjustment only works on the logical-interface-policer.

Options

egress Optional. Egress overhead adjustment in bytes.

- **Range:** -64 through 64

ingress Optional. Ingress overhead adjustment in bytes.

- **Range:** -64 through 64

policer-overhead-value Policer overhead adjustment to be made in both ingress and egress directions.

- **Range:** -64 through 64

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

logical-interface-policer

shaping-rate (CoS Interfaces) | **16933**

policer-overhead-adjustment

IN THIS SECTION

- Syntax | **12645**
- Hierarchy Level | **12645**
- Description | **12645**
- Options | **12645**
- Required Privilege Level | **12645**

Release Information | 12646

Syntax

```
interfaces interface-name
    unit unit-number
policer-overhead <ingress | egress> <adjustment in-bytes>
policer-overhead <adjustment in-bytes>
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Add the configured number of bytes to the length of a packet entering the interface or leaving the interface. The policer overhead adjustment at per IFL or direction granularity is in the range of -64 bytes to +64 bytes.

NOTE: If you configure the `policer-overhead` adjustment statement without the ingress or egress option and only with values in the supported range, it is considered for both ingress and egress.

For example: set interfaces xe-1/0/1 policer-overhead <*value*>

Options

adjustment-in-bytes—Policer overhead bytes to be accounted in ingress and egress. Number of bytes used to adjust the length of a packet used for policing purposes entering an interface and leaving an interface.

- **Range:** -64 to +64 bytes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Accounting of the Policer Overhead Attribute at the Interface Level

Configuring the Accounting of Policer Overhead in Interface Statistics

policer-charge

IN THIS SECTION

- Syntax | 12646
- Hierarchy Level | 12646
- Description | 12647
- Default | 12648
- Required Privilege Level | 12648
- Release Information | 12648

Syntax

```
policer-charge
```

Hierarchy Level

```
[edit firewall family firewall filter family filter filter name term term name then]
```

Description

Hierarchical policers support policing of premium and non-premium traffic. The former is policed by a premium policer and the latter by an aggregate policer. Premium traffic is subject to expedited forwarding. Non-premium traffic is subject to non-expedited forwarding. Out of the aggregate bandwidth available to both types of traffic, a user-selectable bandwidth can be assigned to premium traffic. Assuming the aggregate traffic is 10 Mbps, and assured bandwidth for premium traffic is 5 mbps, the following scenarios are possible:

- When premium traffic is present, then premium traffic gets an assured bandwidth of 5 mbps. Remaining bandwidth is available to non-premium traffic.
- When premium traffic is absent, non-premium traffic can use upto 10 Mbps.
- When premium traffic is present, and non-premium traffic is absent, then premium traffic gets an assured bandwith of not more than 5 Mpbs.

The `policer-charge` firewall filter action enables this type of hierarchical policing on supported platforms. The following configuration is representative of a typical firewall filter configuration with the `policer-charge` action. Here, the `policer-charge` substracts the available credits, and makes it available to the aggregate policer. As an example, if aggregate bandwidth is 10 Mpbs, and assured premium bandwidth is 5 Mbps, and premium traffic is consuming 3 Mbps, then `policer-charge` substracts 3 Mbps from 10 Mbps. The remaining bandwith, 7 Mpbs, is made available to the aggregate policer.

```
set firewall family inet filter f1 term t1 from Premium traffic
set firewall family inet filter f1 term t1 then policer premium
set firewall family inet filter f1 term t1 then next term
set firewall family inet filter f1 term t2 from Premium traffic
set firewall family inet filter f1 term t2 then policer-charge aggregate
set firewall family inet filter f1 term t2 then accept
set firewall family inet filter f1 term t3 then policer aggregate
set firewall family inet filter f1 term t3 then accept
```

NOTE: All the terms need to be configured in same order.

The term that has the `policer-charge` and the preceding term that has the next term action, must have same match conditions.

All the policers, be it premium or aggregate, must be of same type.

The premium and aggregate policers must have discard actions.

The aggregate bandwidth should always be higher than the premium bandwidth.

Policer overhead adjustment is supported.

Policer accuracy - To support higher policer accuracy, a new packet size is added, which enables higher policer support. Policer rate configurable is also increased to 14.4 tbps. Policer accuracy is provided as close as possible to configured rate.

Default

`policer-charge` *`aggregate policer`*

Required Privilege Level

firewall—To view this statement in the configuration.

firewall -control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.4R1

policer-charge

IN THIS SECTION

- Syntax | 12649
- Hierarchy Level | 12649
- Description | 12649
- Default | 12650
- Required Privilege Level | 12650
- Release Information | 12650

Syntax

```
policer-charge
```

Hierarchy Level

```
[edit firewall family firewall filter family filter filter name term term name then]
```

Description

Hierarchical policers support policing of premium and non-premium traffic. The former is policed by a premium policer and the latter by an aggregate policer. Premium traffic is subject to expedited forwarding. Non-premium traffic is subject to non-expedited forwarding. Out of the aggregate bandwidth available to both types of traffic, a user-selectable bandwidth can be assigned to premium traffic. Assuming the aggregate traffic is 10 Mbps, and assured bandwidth for premium traffic is 5 mbps, the following scenarios are possible:

- When premium traffic is present, then premium traffic gets an assured bandwidth of 5 mbps. Remaining bandwidth is available to non-premium traffic.
- When premium traffic is absent, non-premium traffic can use upto 10 Mbps.
- When premium traffic is present, and non-premium traffic is absent, then premium traffic gets an assured bandwith of not more than 5 Mpbs.

The `policer-charge` firewall filter action enables this type of hierarchical policing on supported platforms. The following configuration is representative of a typical firewall filter configuration with the `policer-charge` action. Here, the `policer-charge` substracts the available credits, and makes it available to the aggregate policer. As an example, if aggregate bandwidth is 10 Mpbs, and assured premium bandwidth is 5 Mbps, and premium traffic is consuming 3 Mbps, then `policer-charge` substracts 3 Mbps from 10 Mbps. The remaining bandwith, 7 Mpbs, is made available to the aggregate policer.

```
set firewall family inet filter f1 term t1 from Premium traffic
set firewall family inet filter f1 term t1 then policer premium
set firewall family inet filter f1 term t1 then next term
set firewall family inet filter f1 term t2 from Premium traffic
set firewall family inet filter f1 term t2 then policer-charge aggregate
set firewall family inet filter f1 term t2 then accept
```

```
set firewall family inet filter f1 term t3 then policer aggregate
set firewall family inet filter f1 term t3 then accept
```

NOTE: All the terms need to be configured in same order.

The term that has the `policer-charge` and the preceding term that has the next term action, must have same match conditions.

All the policers, be it premium or aggregate, must be of same type.

The premium and aggregate policers must have discard actions.

The aggregate bandwidth should always be higher than the premium bandwidth.

Policer overhead adjustment is supported.

Policer accuracy - To support higher policer accuracy, a new packet size is added, which enables higher policer support. Policer rate configurable is also increased to 14.4 tbps. Policer accuracy is provided as close as possible to configured rate.

Default

`policer-charge` *`aggregate policer`*

Required Privilege Level

firewall—To view this statement in the configuration.

firewall -control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.4R1

policies

IN THIS SECTION

Syntax | 12651

- Hierarchy Level | 12659
- Description | 12659
- Options | 12659
- Required Privilege Level | 12659
- Release Information | 12659

Syntax

```
policies {
    default-policy (deny-all | permit-all);
    from-zone from-zone-name {
        to-zone;
        policy name {
            description description;
            match (Security Policies Global) {
                source-address (Security Policies);
                destination-address (Security Policies);
                application (Security Policies);
                source-identity;
                source-end-user-profile <source-end-user-profile-name>;
                dynamic-application (Security Policies);
                url-category;
                from-zone (Security Policies Global);
                to-zone (Security Policies Global);
                source-l3vpn-vrf-group [ source-l3vpn-vrf-group ... ];
                destination-l3vpn-vrf-group [ destination-l3vpn-vrf-group ... ];
                destination-address-excluded;
                source-address-excluded;
            }
            scheduler-name scheduler-name;
            then {
                deny;
                permit {
                    application-services {
                        (redirect-wx | reverse-redirect-wx);
                        advanced-anti-malware-policy advanced-anti-malware-policy;
                        application-traffic-control {
                            rule-set rule-set;
```

```
                    }
                    gprs-gtp-profile gprs-gtp-profile;
                    gprs-sctp-profile gprs-sctp-profile;
                    icap-redirect icap-redirect;
                    idp;
                    idp-policy idp-policy;
                    security-intelligence-policy security-intelligence-policy;
                    ssl-proxy {
                        profile-name profile-name;
                    }
                    uac-policy {
                        captive-portal captive-portal;
                    }
                    utm-policy utm-policy;
                    web-proxy {
                        profile-name profile-name;
                    }
                }
                destination-address (Security IDP Policy) {
                    (drop-translated | drop-untranslated);
                }
                firewall-authentication {
                    pass-through {
                        access-profile access-profile;
                        auth-only-browser;
                        auth-user-agent name;
                        client-match [ client-match ... ];
                        ssl-termination-profile ssl-termination-profile;
                        web-redirect;
                        web-redirect-to-https;
                    }
                    user-firewall {
                        access-profile access-profile;
                        auth-only-browser;
                        auth-user-agent name;
                        domain domain;
                        ssl-termination-profile ssl-termination-profile;
                        web-redirect;
                        web-redirect-to-https;
                    }
                    web-authentication {
                        client-match [ client-match ... ];
                    }
```

```
                    push-to-identity-management;
                }
                services-offload;
                tcp-options {
                    initial-tcp-mss initial-tcp-mss;
                    reverse-tcp-mss reverse-tcp-mss;
                    sequence-check-required;
                    syn-check-required;
                    window-scale;
                }
                tunnel {
                    ipsec-vpn ipsec-vpn;
                    pair-policy pair-policy;
                }
            }
            reject {
                profile profile;
                ssl-proxy {
                    profile-name profile-name;
                }
            }
            count {
            }
            log {
                session-close;
                session-init;
            }
        }
    }
}
global {
    policy name {
        description description;
        match (Security Policies Global) {
            source-address (Security Policies);
            destination-address (Security Policies);
            application (Security Policies);
            source-identity;
            source-end-user-profile <source-end-user-profile-name>;
            dynamic-application (Security Policies);
            url-category;
            from-zone (Security Policies Global);
            to-zone (Security Policies Global);
```

```
                source-l3vpn-vrf-group [ source-l3vpn-vrf-group ... ];
                destination-l3vpn-vrf-group [ destination-l3vpn-vrf-group ... ];
                destination-address-excluded;
                source-address-excluded;
            }
            scheduler-name scheduler-name;
            then {
                deny;
                permit {
                    application-services {
                        (redirect-wx | reverse-redirect-wx);
                        advanced-anti-malware-policy advanced-anti-malware-policy;
                        application-traffic-control {
                            rule-set rule-set;
                        }
                        gprs-gtp-profile gprs-gtp-profile;
                        gprs-sctp-profile gprs-sctp-profile;
                        icap-redirect icap-redirect;
                        idp;
                        idp-policy idp-policy;
                        security-intelligence-policy security-intelligence-policy;
                        ssl-proxy {
                            profile-name profile-name;
                        }
                        uac-policy {
                            captive-portal captive-portal;
                        }
                        utm-policy utm-policy;
                        web-proxy {
                            profile-name profile-name;
                        }
                    }
                    destination-address {
                        (drop-translated | drop-untranslated);
                    }
                    firewall-authentication {
                        pass-through {
                            access-profile access-profile;
                            auth-only-browser;
                            auth-user-agent name;
                            client-match [ client-match ... ];
                            ssl-termination-profile ssl-termination-profile;
                            web-redirect;
```

```
                        web-redirect-to-https;
                    }
                    user-firewall {
                        access-profile access-profile;
                        auth-only-browser;
                        auth-user-agent name;
                        domain domain;
                        ssl-termination-profile ssl-termination-profile;
                        web-redirect;
                        web-redirect-to-https;
                    }
                    web-authentication {
                        client-match [ client-match ... ];
                    }
                    push-to-identity-management;
                }
                services-offload;
                tcp-options {
                    initial-tcp-mss initial-tcp-mss;
                    reverse-tcp-mss reverse-tcp-mss;
                    sequence-check-required;
                    syn-check-required;
                    window-scale;
                }
                tunnel {
                    ipsec-vpn ipsec-vpn;
                    pair-policy pair-policy;
                }
            }
            reject {
                profile profile;
                ssl-proxy {
                    profile-name profile-name;
                }
            }
            count {
            }
            log {
                session-close;
                session-init;
            }
        }
    }
```

```
    }
    policy-rematch <extensive>;
    policy-stats {
        system-wide (disable | enable);
    }
    pre-id-default-policy {
        then {
            log {
                session-close;
                session-init;
            }
            session-timeout {
                icmp seconds;
                icmp6 seconds;
                ospf seconds;
                others seconds;
                tcp seconds;
                udp seconds;
            }
        }
    }
    stateful-firewall-rule name {
        match-direction (input | input-output | output);
        policy name {
            description description;
            match (Security Policies Global) {
                source-address (Security Policies);
                destination-address (Security Policies);
                application (Security Policies);
                source-identity;
                source-end-user-profile <source-end-user-profile-name>;
                dynamic-application (Security Policies);
                url-category;
                from-zone (Security Policies Global);
                to-zone (Security Policies Global);
                source-l3vpn-vrf-group [ source-l3vpn-vrf-group ... ];
                destination-l3vpn-vrf-group [ destination-l3vpn-vrf-group ... ];
                destination-address-excluded;
                source-address-excluded;
            }
            scheduler-name scheduler-name;
            then {
                deny;
```

```
                permit {
                    application-services {
                        (redirect-wx | reverse-redirect-wx);
                        advanced-anti-malware-policy advanced-anti-malware-policy;
                        application-traffic-control {
                            rule-set rule-set;
                        }
                        gprs-gtp-profile gprs-gtp-profile;
                        gprs-sctp-profile gprs-sctp-profile;
                        icap-redirect icap-redirect;
                        idp;
                        idp-policy idp-policy;
                        security-intelligence-policy security-intelligence-policy;
                        ssl-proxy {
                            profile-name profile-name;
                        }
                        uac-policy {
                            captive-portal captive-portal;
                        }
                        utm-policy utm-policy;
                        web-proxy {
                            profile-name profile-name;
                        }
                    }
                    destination-address {
                        (drop-translated | drop-untranslated);
                    }
                    firewall-authentication {
                        pass-through {
                            access-profile access-profile;
                            auth-only-browser;
                            auth-user-agent name;
                            client-match [ client-match ... ];
                            ssl-termination-profile ssl-termination-profile;
                            web-redirect;
                            web-redirect-to-https;
                        }
                        user-firewall {
                            access-profile access-profile;
                            auth-only-browser;
                            auth-user-agent name;
                            domain domain;
                            ssl-termination-profile ssl-termination-profile;
```

```
                        web-redirect;
                        web-redirect-to-https;
                    }
                    web-authentication {
                        client-match [ client-match ... ];
                    }
                    push-to-identity-management;
                }
                services-offload;
                tcp-options {
                    initial-tcp-mss initial-tcp-mss;
                    reverse-tcp-mss reverse-tcp-mss;
                    sequence-check-required;
                    syn-check-required;
                    window-scale;
                }
                tunnel {
                    ipsec-vpn ipsec-vpn;
                    pair-policy pair-policy;
                }
            }
            reject {
                profile profile;
                ssl-proxy {
                    profile-name profile-name;
                }
            }
            count {
            }
            log {
                session-close;
                session-init;
            }
        }
    }
}
stateful-firewall-rule-set name {
    stateful-firewall-rule name;
}
traceoptions (Security Policies) {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
    flag name;
```

```
        no-remote-trace;
    }
    unified-policy {
        max-lookups max-lookups;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure a network security policies with IPv6 addresses only if flow support for IPv6 traffic is enabled on the device.

Options

default-policy

Configure a default action when no user-defined policy match.

- Values:
  - deny-all—Deny all traffic if no policy match
  - permit-all—Permit all traffic if no policy match

policy-rematch

Re-evaluate the policy when changed.

- Values:
  - extensive—Perform policy extensive rematch

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Support for the `services-offload` option added in Junos OS Release 11.4.

Support for the `source-identitiy` option added in Junos OS Release 12.1.

Support for the `description` option added in Junos OS Release 12.1.

Support for the `ssl-termination-profile` and `web-redirect-to-https` options are added starting from Junos OS Release 12.1X44-D10 and Junos OS Release 15.1X49-D40.

Support for the `user-firewall` option added in Junos OS Release 12.1X45-D10.

Support for the `domain` option, and for the `from-zone` and `to-zone` global policy match options, added in Junos OS Release 12.1X47-D10.

Support for the `initial-tcp-mss` and `reverse-tcp-mss` options added in Junos OS Release 12.3X48-D20. Support for the `extensive` option for `policy-rematch` added in Junos OS Release 15.1X49-D20.

Starting in Junos OS Release 18.2R1, an IDP policy is available within unified security policy. The IDP policy access is simplified and made available under the unified policy as one of the policy. When an IDP policy is available within a unified security policy, configuring source or destination address, source and destination-except, from and to zone, or application is not required, because the match happens in the security policy itself.

Starting in Junos OS Release 18.3R1, when an SRX Series Firewall is configured with a unified policies, you can configure multiple IDP policies and set one of those policies as the default IDP policy. If multiple IDP policies are configured for a session and when policy conflict occurs, the device applies the default IDP policy for that session and thus resolves any policy conflicts.

NOTE: If you have configured two or more IDP policies in a unified security policy, then you must configure the default IDP policy.

RELATED DOCUMENTATION

Security Policies Overview

policing

IN THIS SECTION

- Syntax | 12661
- Hierarchy Level | 12661
- Description | 12661
- Options | 12661
- Required Privilege Level | 12662
- Release Information | 12662

Syntax

```
policing (filter filter-name | no-automatic-policing);
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name]
[edit interfaces interface-id unit number-of-logical-unit family inet address ip-address]
```

Description

Apply a rate-limiting policer as the specified policing filter:

- To the LSP for MPLS over CCC.
- To the customer-edge interface for IP over MPLS.

Options

filter *filter-name*—Specify the name of the policing filter.

no-automatic-policing—Disable automatic policing on this LSP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1 .

RELATED DOCUMENTATION

Configuring Policers to Control Traffic Rates (CLI Procedure)

Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect

Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS

policing (Protocols LDP)

IN THIS SECTION

- Syntax | 12662
- Hierarchy Level | 12663
- Description | 12663
- Options | 12663
- Required Privilege Level | 12663
- Release Information | 12663

Syntax

```
policing {
    fec fec-address {
        ingress-traffic filter-name;
        transit-traffic filter-name;
```

```
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Enable policing of forwarding equivalence classes (FECs) for LDP.

Options

`fec` *`fec-address`*—Specify the address for the FEC.

`ingress-traffic` *`filter-name`*—Specify the name of the filter for policing ingress FEC traffic.

`transit-traffic` *`filter-name`*—Specify the name of the filter for policing transit FEC traffic.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Policers for LDP FECs

policing (Protocols MPLS)

IN THIS SECTION

- Syntax | 12664
- Hierarchy Level | 12664
- Description | 12664
- Options | 12664
- Required Privilege Level | 12665
- Release Information | 12665

Syntax

```
policing {
    filter filter-name;
    no-auto-policing;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Description

Specify the policing filter for the LSP.

Options

`filter` *`filter-name`*—Specify the name of the policing filter.

`no-auto-policing`—Disable automatic policing on this LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

NOTE: Filter based LSP policing is not supported for point-to-multipont LSPs.

RELATED DOCUMENTATION

Configuring Policers for LSPs

auto-policing | 2724

policy

IN THIS SECTION

- Syntax | 12665
- Hierarchy Level | 12666
- Description | 12666
- Required Privilege Level | 12667
- Release Information | 12667

Syntax

```
policy policy-name {
    accept (Router Advertisement Guard Policy) {
```

```
        match-list {
            match-criteria {
                (match-all | match-any);
            }
            prefix-list-name prefix-list-name;
            source-ip-address-list address-list-name;
            source-mac-address-list address-list-name;
        }
        match-option {
            hop-limit {
                (maximum | minimum) value;
            }
            managed-config-flag;
            other-config-flag;
            router-preference (high | low | medium);
        }
    }
    discard (forwarding-options) {
        prefix-list-name prefix-list-name;
        source-ip-address-list address-list-name;
        source-mac-address-list address-list-name;
    }
}
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard]
[edit forwarding-options access-security router-advertisement-guard interface interface-name]
[edit forwarding-options access-security router-advertisement-guard vlans (vlan-name| all)]
```

Description

Configure the policy for an IPv6 Router Advertisement (RA) guard. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages based on whether they match the conditions defined in the policy.

RA guard compares the information contained in attributes of RA messages to the information contained in the policy. You must configure the policy before you can enable RA guard. You can configure either an accept policy or a discard policy and enable it on an interface or on a VLAN. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are

forwarded, and RA messages that do not match the conditions are dropped. When RA guard is enabled by using a discard policy, any RA messages that match the conditions defined in the policy are dropped, and RA messages that do not match the conditions are forwarded.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

policy (advanced-policy-based-routing)

IN THIS SECTION

- Syntax | 12668
- Hierarchy Level | 12668
- Description | 12668
- Options | 12669
- Required Privilege Level | 12670
- Release Information | 12670

Syntax

```
policy name {
 description description;
 match {
  source-address;
  destination-address;
  application [ application ... ];
  source-identity;
  destination-address-excluded;
  source-address-excluded;
 }
 scheduler-name scheduler-name;
 then {
  application-services {
   advance-policy-based-routing-profile advance-policy-based-routing-profile;
  }
 }
}
```

Hierarchy Level

```
[edit security advanced-policy-based-routing from-zone name]
```

Description

Configure advanced policy-based routing (APBR) policies.

You can create APBR policies for a security zone and apply advanced policy-based routing (APBR) profiles on the traffic that matches the policy.

In the APBR policy, you can define source addresses, destination addresses, and applications as match conditions; and after a successful match, the configured APBR profile is applied as an application services for the session.

The routing instance associated with APBR profile includes a static route and next hop configured. The matching traffic arriving at the trust zone is forwarded to a specific device or interface as specified by the next-hop IP address.

> **NOTE**: When using specific address or address set in the APBR policy rule, we recommend to use the global address book. Because, zone specific rules might not be applicable for destination address, as the destination zone is not known at time of policy evaluation.

Options

policy *policy-name*
Specify the name of the APBR policy.

description
Specify descriptive text for the APBR policy.

match
Specify an APBR policy match-criteria.

- **source-address** — Define the source address as the matching criteria.
- **destination-address** — Define the destination address as the matching criteria.
- **application** — Name of the predefined or custom application or application set used as match criteria.
- **destination-address-excluded** — Exclude destination addresses.
- **source-address-excluded** — Exclude source addresses.
- **source-identity** — Specify users and roles to be used as the match criteria.

scheduler-name *scheduler-name*
Specify the name of the scheduler to associate with APBR policy.

Schedulers allow you to activate a policy for a specified duration. To define a scheduler for a APBR policy, you must first create a scheduler and then refer the scheduler in your policy configuration.

When a scheduler times out, the associated policy is deactivated. All sessions associated with the policy are subsequently timed out only if policy-rematch is used.

then
Specify the policy action to be performed when packets match the defined criteria.

- **application-services** — Enable application services within a security policy. the following application services is supported:
  - advance-policy-based-routing-profile *apbr-profile-name*—Specify the advanced policy-based routing (APBR) profile.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1

RELATED DOCUMENTATION

Advanced Policy-Based Routing

policy (Aggregate and Generated Routes)

IN THIS SECTION

- Syntax | 12670
- Hierarchy Level | 12671
- Description | 12671
- Options | 12672
- Required Privilege Level | 12672
- Release Information | 12672

Syntax

```
policy policy-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options (aggregate | generate) (defaults | route)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name (aggregate | generate) (defaults | route)],
[edit logical-systems logical-system-name routing-options (aggregate | generate) (defaults |
route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name (aggregate |
generate) (defaults | route)],
[edit routing-instances routing-instance-name routing-options (aggregate | generate) (defaults |
route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name (aggregate |
generate) (defaults | route)],
[edit routing-options (aggregate | generate) (defaults | route)],
[edit routing-options rib routing-table-name (aggregate | generate) (defaults | route)]
```

Description

Associate a routing policy when configuring an aggregate or generated route's destination prefix in the **routes** part of the **aggregate** or `generate` statement. This provides the equivalent of an import routing policy filter for the destination prefix. That is, each potential contributor to an aggregate route, along with any aggregate options, is passed through the policy filter. The policy then can accept or reject the route as a contributor to the aggregate route.

If the contributor is accepted, the policy can modify the default preferences. The contributor with the numerically smallest prefix becomes the most preferred, or *primary*, contributor. A rejected contributor still can contribute to a less specific aggregate route. If you do not specify a policy filter, all candidate routes contribute to an aggregate route.

The following algorithm is used to compare two generated contributing routes in order to determine which one is the primary or preferred contributor:

1. Compare the protocol's **preference** of the contributing routes. The lower the preference, the better the route. This is similar to the comparison that is done while determining the best route for the routing table.

2. Compare the protocol's **preference2** of the contributing routes. The lower **preference2** value is better. If only one route has **preference2**, then this route is preferred.

3. The preference values are the same. Proceed with a numerical comparison of the prefixes' values.

   a. The primary contributor is the numerically smallest prefix value.

b. If the two prefixes are numerically equal, the primary contributor is the route that has the smallest prefix length value.

At this point, the two routes are the same. The primary contributor does not change. An additional next hop is available for the existing primary contributor.

A rejected contributor still can contribute to less specific generated route. If you do not specify a policy filter, all candidate routes contribute to a generated route.

Options

policy-name—Name of a routing policy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation

aggregate (Routing) | 2070

generate | 7672

policy-based-logging (LRF Profile)

IN THIS SECTION

- Syntax | 12673
- Hierarchy Level | 12673
- Description | 12673

Required Privilege Level | 12673
Release Information | 12673

Syntax

```
policy-based-logging;
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Description

Configure policy-based logging, which causes the LRF rules to be activated by PCC rules.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring Logging and Reporting for Subscriber Management

policy-db-size

IN THIS SECTION

- Syntax | 12674
- Hierarchy Level | 12674
- Description | 12674
- Options | 12674
- Required Privilege Level | 12675
- Release Information | 12675

Syntax

```
policy-db-size size;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-
provider]
```

Description

Configure the size of the policy database. When this setting is changed, the PIC reboots.

NOTE: At least one data core must be configured to configure the size of the policy database.

Options

size—Size of the policy database, in megabytes (MB). The size of the forwarding database and the size of the policy database together must be smaller than the size of the object cache.

- **Range:** 0 through 1279 MB

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

forwarding-db-size | 7454
object-cache-size | 11765
wired-process-mem-size | 31133

policy (Event Policy)

IN THIS SECTION

- Syntax | 12675
- Hierarchy Level | 12677
- Description | 12677
- Default | 12678
- Options | 12678
- Required Privilege Level | 12678
- Release Information | 12678

Syntax

```
policy policy-name {
    attributes-match {
        event1.attribute-name equals event2.attribute-name;
```

```
        event.attribute-name matches regular-expression;
        event1.attribute-name starts-with event2.attribute-name;
    }
    events [ events ];
    then {
        ... the then subhierarchy appears at the end of the [edit event-options policy policy-
name] hierarchy level ...
    }
    within seconds {
        events [ events ];
        not events [ events ];
        trigger (on | after | until) event-count;
    }

    then {
        change-configuration {
            commands {
                "command";
            }
            commit-options {
                check <synchronize>;
                force;
                log "comment-string";
                synchronize;
            }
            retry count number interval seconds;
            user-name username;
        }
        event-script filename {
            arguments {
                argument-name argument-value;
            }
            destination destination-name {
                retry-count count retry-interval seconds;
                transfer-delay seconds;
            }
            output-filename filename;
            output-format (text | xml);
            user-name username;
        }
        execute-commands {
            commands {
                "command";
```

```
            }
            destination destination-name {
                retry-count count retry-interval seconds;
                transfer-delay seconds;
            }
            output-filename filename;
            output-format (text | xml);
            user-name username;
        }
        ignore;
        priority-override {
            facility facility-type;
            severity severity-level;
        }
        raise-trap;
        upload filename (filename | committed) destination destination-name {
            retry-count count retry-interval seconds;
            transfer-delay seconds;
            user-name username;
        }
    }
}
```

Hierarchy Level

```
[edit event-options]
```

Description

Define an event policy to be processed by the eventd process. If you configure a policy, the `events` and `then` statements are mandatory.

You can configure multiple policies to be processed for an event. The policies are executed in the order in which they appear in the configuration. If you configure more than one policy for an event, and if one of the policies is configured to ignore the event, no policies that follow the policy with the `ignore` statement are executed.

Default

If you do not configure a policy for an event, no actions are taken upon receipt of the event, which is recorded in the system log.

Options

policy-name	Name of an event policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

policy (Explicit-Proxy)

IN THIS SECTION

- Syntax | 12678
- Hierarchy Level | 12682
- Description | 12682
- Options | 12682
- Required Privilege Level | 12682
- Release Information | 12682

Syntax

```
policy name {
```

```
description description;
match {
 destination-address;
 destination-address-excluded;
 destination-identity-feed [ destination-identity-feed ... ];
 dynamic-application;
 application;
 source-address;
 source-address-excluded;
 source-identity;
 source-identity-feed [ source-identity-feed ... ];

 source-end-user-profile< source-end-user-profile-name>;

 url-category;
}
report-skip;
scheduler-name scheduler-name;
then {
 deny {
  application-services {
   security-intelligence {
    add-destination-identity-to-feed name-of-feed;
    add-destination-ip-to-feed name-of-feed;
    add-source-identity-to-feed name-of-feed;
    add-source-ip-to-feed name-of-feed;
   }
  }
 }
 log {
  profile profile;
  session-close;
  session-init;
  session-update minutes;
 }
 permit {
  advanced-connection-tracking;
  application-services {
   advanced-anti-malware-policy advanced-anti-malware-policy;
   anti-virus-policy anti-virus-policy;
   application-traffic-control {
    rule-set rule-set;
   }
```

```
        captive-portal {
         profile profile;
        }
        icap-redirect icap-redirect;
        idp-policy idp-policy;
        packet-capture;
        security-intelligence {
         add-destination-identity-to-feed name-of-feed;
         add-destination-ip-to-feed name-of-feed;
         add-source-identity-to-feed name-of-feed;
         add-source-ip-to-feed name-of-feed;
        }
        security-intelligence-policy security-intelligence-policy;
        ssl-proxy {
         profile-name profile-name;
        }
        utm-policy utm-policy;
       destination-address {
        (drop-translated | drop-untranslated);
       }
       firewall-authentication {
        pass-through {
         access-profile access-profile;
         auth-only-browser;
         auth-user-agent name;

         client-match [ client-match ... ];
         ssl-termination-profile ssl-termination-profile;
         web-authentication-server web-authentication-server;
         web-redirect;
         web-redirect-to-https;
        }
        user-firewall {
         access-profile access-profile;
         auth-only-browser;
         auth-user-agent name;

         domain domain;
         ssl-termination-profile ssl-termination-profile;
         web-authentication-server web-authentication-server;
         web-redirect;
         web-redirect-to-https;
        }
```

```
        web-authentication {
         client-match [ client-match ... ];
        }
        push-to-identity-management;
       }
       no-services-offload;
       tcp-options {
        initial-tcp-mss initial-tcp-mss;
        reverse-tcp-mss reverse-tcp-mss;
        sequence-check-required;
        syn-check-required;
        window-scale;
       }
       tunnel {
        ipsec-vpn ipsec-vpn;
        pair-policy pair-policy;
       }
      tunnel-inspection profile-name;
      }
      reject {
       application-services {
        security-intelligence {
         add-destination-identity-to-feed name-of-feed;
         add-destination-ip-to-feed name-of-feed;
         add-source-identity-to-feed name-of-feed;
         add-source-ip-to-feed name-of-feed;
        }
       }
       profile profile;
       ssl-proxy {
        profile-name profile-name;
       }
      }

     }
     }
```

Hierarchy Level

```
[edit security policies explicit-proxy profile]
```

Description

Define a explicit proxy profiles policies.

Options

name Security policy name.

Range: 1 through 63 characters

description Text description of policy.

The descriptive text should not include characters “<”, “>”, “&”, or “\n”.
The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 900 characters, the configuration might fail to take effect.

Range: 1 through 900 characters

report-skip Skip report for this policy

scheduler-name Name of scheduler

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

policy (Flow Maps)

IN THIS SECTION

- Syntax | 12683
- Hierarchy Level | 12683
- Description | 12683
- Options | 12683
- Required Privilege Level | 12684
- Release Information | 12684

Syntax

```
policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast flow-map flow-map-name],
[edit logical-systems logical-system-name routing-options multicast flow-map flow-map-name],
[edit routing-instances routing-instance-name routing-options multicast flow-map flow-map-name],
[edit routing-options multicast flow-map flow-map-name]
```

Description

Configure a flow map policy.

Options

policy-names—Name of one or more policies for flow mapping.

Required Privilege Level

routing—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

policy (IPsec)

IN THIS SECTION

- Syntax | 12684
- Hierarchy Level | 12684
- Description | 12685
- Options | 12685
- Required Privilege Level | 12685
- Release Information | 12685

Syntax

```
policy policy-name {
    description description;
    perfect-forward-secrecy {
        keys (group1 | group2 |group5 |group14 |group15 | group16 | group24);
    }
    proposals [ proposal-names ];
}
```

Hierarchy Level

```
[edit services ipsec-vpn ipsec]
```

Description

Define an IPsec policy.

Options

policy-name—IPsec policy name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Policies

policy-lookup-cache

IN THIS SECTION

- Syntax | 12686
- Hierarchy Level | 12686
- Description | 12686
- Default | 12686
- Required Privilege Level | 12686
- Release Information | 12686

Syntax

```
(policy-lookup-cache | no-policy-lookup-cache);
```

Hierarchy Level

```
[edit security idp sensor-configuration global]
```

Description

Enable cache to accelerate IDP policy lookup which improves IDP performance.

Default

`policy-lookup-cache` is enabled by default.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

policy-map (Class of Service)

IN THIS SECTION

- Syntax | 12687
- Hierarchy Level | 12687
- Description | 12687
- Options | 12689

Required Privilege Level | 12689
Release Information | 12689

Syntax

```
policy-map policy-map-name {
    inet-precedence proto-ip code-point [alias | bits];
    inet-precedence proto-mpls code-point [alias | bits];
    inet-precedence proto-all code-point [alias | bits];
    inet6-precedence proto-ip code-point [alias | bits];
    inet6-precedence proto-mpls code-point [alias | bits];
    dscp proto-ip code-point [alias | bits];
    dscp proto-mpls code-point [alias | bits];
    dscp proto-all code-point [alias | bits];
    dscp-ipv6 proto-ip code-point [alias | bits];
    dscp-ipv6 proto-mpls code-point [alias | bits];
    dscp-ipv6 proto-all code-point [alias | bits];
    exp all-label code-point [alias | bits];
    exp outer-label code-point [alias | bits];
    ieee-802.1 outer code-point [alias | bits];
    ieee-802.1 outer-and-inner code-point [alias | bits];
    ieee-802.1ad outer code-point [alias | bits];
    ieee-802.1ad outer-and-inner code-point [alias | bits];
    user-attribute integer;
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Define the packet-marking scheme (rewrite rules) for a customer. Traditionally, packet marking in Junos uses the forwarding class and loss priority determined from a behavior aggregate (BA) classifier or multifield classifier. This approach does not allow rewrite rules to be directly assigned for each customer

because of the limited number of combinations of forwarding class and loss priority. `policy-map` enables you to define rewrite rules on a per-customer basis.

NOTE: To enable the policy map feature:

- On most devices, enable `enhanced-ip`, `enhanced-ethernet`, or `enhanced-mode` under [`edit chassis network-services`].
- On PTX10002-36QDD routers, enable `policy-map-marking` on each egress logical interface that you want to do policy map marking. You enable `policy-map-marking` at the [`edit class-of-service interfaces` *`interface-name`* `unit` *`unit-number`*] hierarchy level.

NOTE: Policy maps have the following configuration restrictions:

- When configuring both `proto-ip` and `proto-mpls` options for `inet-precedence`, `inet6-precedence`, `dscp`, or `dscp-ipv6`, you must configure both options with the same code point or code point alias.
- You cannot configure `inet-precedence` and `dscp` in the same policy map.
- You cannot configure `inet6-precedence` and `dscp-ipv6` in the same policy map.
- In case of MPLS SWAP/PUSH operation, only the new labels are marked on all label-switching routers (LSRs), except the penultimate hop case where if it exposes the next label in the stack, then the exposed label is marked. Therefore, with the penultimate hop, the service label is changed.
- You cannot configure `ieee-802.1` and `ieee-802.1ad` in the same policy map.
- You cannot configure both `outer` and `outer-and-inner` options for `ieee-802.1` and `ieee-802.1ad` code points in the same policy map.
- For IEEE 802.1ad with the `outer-and-inner` option, the discard eligibility (DE) bit is marked only for the outer VLAN header. For the inner VLAN header, only the three CoS Bits are marked.

You assign the policy map to a customer through a firewall action on an ingress or egress firewall filter (where the match conditions identify the customer). Alternatively, you can also assign a policy map to an ingress interface or a routing instance. A policy map is executed on a packet just before it is queued, so it overrides any other packet-marking scheme that was previously applied to the packet.

Options

alias The predefined CoS code point alias on the device.

bits The bit pattern for the CoS code point.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

`inet6-precedence` options added in Junos OS Release 24.2R1.

Refer to Feature Explorer for platform and release support information.

RELATED DOCUMENTATION

interfaces (Class of Service) | 9130

routing-instances (CoS) | 15971

policy-match

IN THIS SECTION

- Syntax | 12690
- Hierarchy Level | 12690
- Description | 12690
- Options | 12690
- Required Privilege Level | 12690
- Release Information | 12691

Syntax

```
policy match {
    duration interval;
    size count;
    threshold value;
}
```

Hierarchy Level

```
[edit security alarms potential-violation policy]
```

Description

Configure alarms for a policy rule or a group of rule violations within a specified time period.

Options

- `duration` *`interval`*—Indicate the duration of counters.
- **Range:** 1 through 3600 seconds.
- **Default:** 1 second.
- `size` *`count`*—Indicate the number of policies for which policy violation checks can be done concurrently.
- **Range:** 1 through 10240.
- **Default:** 1024.
- `threshold` *`value`*—Indicate the number of policies for which policy violation checks can be done concurrently.
- **Range:** 1 through 4294967295.
- **Default:** 1000.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Security Policies Overview

policy (Multicast-Only Fast Reroute)

IN THIS SECTION

- Syntax | 12691
- Hierarchy Level | 12691
- Description | 12692
- Required Privilege Level | 12693
- Release Information | 12693

Syntax

```
policy policy-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast stream-protection],
[edit logical-systems logical-system-name routing-options multicast stream-protection],
[edit routing-instances routing-instance-name routing-options multicast stream-protection],
[edit routing-options multicast stream-protection]
```

Description

When you configure multicast-only fast reroute (MoFRR), apply a routing policy that filters for a restricted set of multicast streams to be affected by your MoFRR configuration. You can apply filters that are based on source or group addresses.

For example:

```
routing-options {
    multicast {
        stream-protection {
            policy mofrr-select;
        }
    }
}
policy-statement mofrr-select {
    term A {
        from {
            source-address-filter 225.1.1.1/32 exact;
        }
        then {
            accept;
        }
    }
    term B {
        from {
            source-address-filter 226.0.0.0/8 orlonger;
        }
        then {
            accept;
        }
    }
    term C {
        from {
            source-address-filter 227.1.1.0/24 orlonger;
            source-address-filter 227.4.1.0/24 orlonger;
            source-address-filter 227.16.1.0/24 orlonger;
        }
        then {
            accept;
        }
    }
```

```
    term D {
        from {
            source-address-filter 227.1.1.1/32 exact;
        }
        then {
            reject; #MoFRR disabled
        }
    }
    term E {
        from {
            route-filter 227.1.1.0/24 orlonger;
        }
        then accept;
    }
    ...
}
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Multicast-Only Fast Reroute

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain

policy-multipath

IN THIS SECTION

- Syntax | 12694
- Hierarchy Level | 12694
- Description | 12694
- Options | 12695
- Required Privilege Level | 12695
- Release Information | 12695

Syntax

```
policy-multipath policy [ policy }
    traceoptions <file filename <files files>   <size size> <(world-readable | no-world-
readable)>> name detail disable receive send
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options rib],
[edit logical-systems name routing-options rib],
[edit logical-systems name tenants name routing-instances name routing-options rib],
[edit routing-instances name routing-options rib],
[edit routing-options rib],
[edit tenants name routing-instances name routing-options rib]
```

Description

Create policy-based multipath route using a combination of segment routing traffic-engineered (SR-TE) LDP or RSVP routes and SR-TE IP routes. You can resolve BGP service routes over the mutlipath route, and apply export policies to steer traffic differently for different prefixes. The policy-based multipath feature is supported for both IP and IPv6 protocols.

Options

policy Import policy to create policy-based multipath.

> **NOTE**: This statement is supported only at the `[edit routing-option policy-multipath]` hierarchy level.

Any action commands configured in the policy, such as apply, is evaluated using the active route. For non-active routes, the policy is applied to check if the routes can participate in the multipath route or not. Multipath routes inherit all attributes of the active route. These attributes can be modified using the multipath policy configuration.

preserve-nexthop-hierarchy Enable expanded nh hierarchy support for policy-multipath routes for RIBs. For any inet[6].3 RIBs, the preserve-nexthop-hierarchy is enabled by default.

> **NOTE**: This statement is supported only at the `[edit routing-options rib <rib-name> policy-multipath]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Policy-Based Multipath Routes Overview

policy-oids

IN THIS SECTION

- Syntax | 12696
- Hierarchy Level | 12696
- Description | 12696
- Options | 12696
- Required Privilege Level | 12697
- Release Information | 12697

Syntax

```
policy-oids [ oid ];
```

Hierarchy Level

```
[edit services ipsec-vpn ike policy policy-name certificate]
```

Description

Configure policy object identifiers (OIDs). This configuration is optional.

Options

oid Policy OID contained in a peer's certificate or certificate chain. Up to five policy OIDs can be configured. Each OID can be up to 63 bytes long.

> **NOTE**: You must ensure that at least one of the configured policy OIDs is included in a peer's certificate or certificate chain. Note that the **policy-oids** field in a peer's certificate is optional. If you configure policy OIDs in an IKE policy and the peer's

certificate chain does not contain any policy OIDs, certificate validation for the peer fails.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding Digital Certificate Validation

Example: Improving Digital Certificate Validation by Configuring Policy OIDs on an MX Series Device

policy-options

IN THIS SECTION

- Syntax | 12697
- Hierarchy Level | 12699
- Description | 12699
- Required Privilege Level | 12699
- Release Information | 12699

Syntax

```
policy-options {
     as-list-group group-name {
```

```
        as-list name members [as-start-as-finish., as, as-start-as-finish];
    }
    as-path name regular-expression;
    as-path-group group-name;
    community name {
        invert-match;
        members [ community-ids ];
    }
    condition condition-name {
        if-route-exists address table table-name;
    }
    damping name {
        disable;
        half-life minutes;
        max-suppress minutes;
        reuse number;
        suppress number;
    }
    rib-list <name> {
        <rib-name>;
    }

    policy-statement policy-name {
        term term-name {
            from {
                as-path-neighbors (as-list | as-list-group);
                as-path-origins (as-list | as-list-group);
                as-path-transits (as-list | as-list-group);
                family;
                fpc-pfes-offline pfes-offline-per-fpc;
                match-conditions;
                policy subroutine-policy-name;
                prefix-list name;
                route-filter destination-prefix match-type <actions>;
                rib-list name
                source-address-filter source-prefix match-type <actions>;
            }
            to {
                match-conditions;
                policy subroutine-policy-name;
            }
            then actions;
            default-action (accept | reject);
```

```
            prefix-segment {
                index index;
                node-segment;
            }
        }
        then {
            no-entropy-label-capability;
            priority (high | medium | low);
        }
    }
    prefix-list name {
        ip-addresses;
    }
}
```

Hierarchy Level

```
[edit],
[edit dynamic-profiles profile-name]
```

Description

Configure routing policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit dynamic-profiles]` hierarchy level introduced in Junos OS Release 11.4.

`as-path-neighbors`, `as-path-origins`, and `as-path-transits` statements introduced in Junos OS release 21.3R1.

RELATED DOCUMENTATION

Example: Using Routing Policy in an ISP Network

policy-options (QFX)

IN THIS SECTION

- Syntax | 12700
- Hierarchy Level | 12701
- Description | 12701
- Required Privilege Level | 12701
- Release Information | 12701

Syntax

```
policy-options
    application-maps application-map-name {
        application application-name {
            code-points [ aliases ] [ bit-patterns ];
        }
    }
    policy-statement policy-name {
        term term-name {
            from {
                family family-name;
                match-conditions;
                policy subroutine-policy-name;
                prefix-list prefix-list-name;
                prefix-list-filter prefix-list-name match-type <actions>;
                route-filter destination-prefix match-type <actions>;
                source-address-filter source-prefix match-type <actions>;
            }
            to {
                match-conditions;
                policy subroutine-policy-name;
```

```
            }
            then actions;
        }
    }
```

Hierarchy Level

```
[edit]
```

Description

Configure options such as application maps for DCBX application protocol exchange and policy statements.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- *Defining an Application for DCBX Application Protocol TLV Exchange*
- *Example: Configuring DCBX Application Protocol TLV Exchange*
- Example: Configuring DCBX to Support an iSCSI Application
- *Understanding DCBX Application Protocol TLV Exchange*
- Understanding DCBX Application Protocol TLV Exchange on EX Series Switches

policy-options (Dynamic Profiles)

IN THIS SECTION

- Syntax | 12702
- Hierarchy Level | 12702
- Description | 12702
- Options | 12703
- Required Privilege Level | 12703
- Release Information | 12703

Syntax

```
policy-options {
    prefix-list uid {
        ip-addresses;
        dynamic-db;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Define a list of IPv4 or IPv6 address prefixes for use in a dynamic firewall filter or in an HTTP redirect configuration.

You can configure up to 85,325 prefixes in each prefix list. To configure more than 85,325 prefixes, configure multiple prefix lists and apply them to multiple firewall filter terms.

You can configure policy options in a dynamic client profile or a dynamic service profile.

Options

uid Unique identifier of the prefix list. You must assign a UID as the prefix list name.

ip-addresses List of IPv4 or IPv6 address prefixes, one IP address per line in the configuration.

dynamic-db Specify that the routing policy and policy objects reference policies configured in the dynamic database at the `[edit dynamic]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Using Routing Policy in an ISP Network

policy (PIM rpf-vector)

IN THIS SECTION

- Syntax | 12704
- Hierarchy Level | 12704
- Description | 12704
- Required Privilege Level | 12706
- Release Information | 12706

Syntax

```
policy [policy-name];
```

Hierarchy Level

```
[edit dynamic-profiles name protocols pim rp rpf-vector],
[edit logical-systems name protocols pim rprpf-vector],
[edit logical-systems name routing-instances name protocols pim rp rpf-vector],
[edit protocols pim rp rpf-vector],
[edit routing-instances name protocols pim rp rpf-vector]
```

Description

Create a filter policy. The configured device checks the policy configuration to determine whether or not to apply "rpf-vector" on page 16017 to (S,G).

RPF Vector Policy Example

This example policy shows Source and Group, using Source, using Group.

```
policy-statement pim-rpf-vector-example {
        term A {
            from {
               source-address-filter <filter A>;
            }
            then {
               accept;
            }
        }
        term B {
            from {
               source-address-filter <filter A>;
               route-filter <filter D>;
            }
            then {
               p2mp-lsp-root {
                   address root address;
```

```
                }
                accept;
            }
        }
        term C {
            from {
                route-filter <filter D>;
            }
            then {
                accept;
            }
        }
        ...
    }
```

RPF Vector Policy Configuration statements

This example policy using Source, Group.

```
set protocols pim rpf-vector policy rpf-vector-policy
set policy-options policy-statement rpf-vector-policy term 1 from route-filter 232.0.0.1/32 exact
set policy-options policy-statement rpf-vector-policy term 1 from source-address-filter
22.1.1.2/32 exact
set policy-options policy-statement rpf-vector-policy term 1 then p2mp-lsp-root address 200.1.1.2
set policy-options policy-statement rpf-vector-policy term 1 then accept
```

RPF Vector Policy Configuration statements

This example policy using Group, Source wildcard.

```
set protocols pim rpf-vector policy rpf-vector-policy
set policy-options policy-statement rpf-vector-policy term 1 from source-address-filter
22.1.1.2/32 exact
set policy-options policy-statement rpf-vector-policy term 1 from route-filter 0.0.0.0/0 longer
set policy-options policy-statement rpf-vector-policy term 1 then p2mp-lsp-root address 200.1.1.2
set policy-options policy-statement rpf-vector-policy term 1 then accept
```

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

show pim join | 23070

policy (Protocols Multipoint LDP)

IN THIS SECTION

- Syntax | 12706
- Hierarchy Level | 12706
- Description | 12707
- Options | 12707
- Required Privilege Level | 12707
- Release Information | 12707

Syntax

```
policy policy-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim mldp-inband-signalling],
[edit protocols pim mldp-inband-signalling]
```

Description

Multipoint LDP (M-LDP) in-band signaling enables you to carry multicast traffic across an existing IP/MPLS backbone, while avoiding the use of PIM in the provider core.

On the egress nodes of the point-to-multipoint LSP, specify an M-LDP join translation filter policy where PIM messages are translated into M-LDP FEC bindings. The policy statement is needed when internal BGP (IBGP) is not available in the core site or to override IBGP-based LSP root detection.

The filter policy is configured at the `[edit policy-options]` hierarchy level. The policy generally specifies one or more source-address filters and the point-to-multipoint LDP root IP address using the `p2mp-lsp-root` policy action.

Options

`policy-name`—Name of a policy configured at the `[edit policy-options]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

No Link Title

policy-rematch

IN THIS SECTION

- Syntax | 12708
- Hierarchy Level | 12708

- Description | 12708
- Options | 12708
- Required Privilege Level | 12709
- Release Information | 12709

Syntax

```
policy-rematch <extensive>;
```

Hierarchy Level

```
[edit security policies]
```

Description

Enable the device to reevaluate an active session when its associated security policy is modified. The session remains open if it still matches the policy that allowed the session initially.

The session is closed if its associated policy is renamed, deactivated, or deleted. However, you can use the `extensive` option to reevaluate an active session when its associated security policy is renamed, deactivated, or deleted.

The policy rematch feature is disabled by default.

Options

extensive When a policy is modified or deleted, extensive option checks if any suitable policy permit to keep these sessions alive. This check is done through a fully new policy lookup for the session to see if any policy can still permit it.

> **NOTE**: The `extensive` option does not apply to ALG data sessions or to policies that specify a source-identity, application-services, destination-address (drop-untranslated or drop-translated), firewall-authentication, or a tunnel.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for the `extensive` option added in Junos OS Release 15.1X49-D20.

RELATED DOCUMENTATION

> *Security Policies Overview*
>
> *Unified Policies Configuration Overview*

policy (Security Alarms)

IN THIS SECTION

- Syntax | 12709
- Hierarchy Level | 12710
- Description | 12710
- Options | 12710
- Required Privilege Level | 12710
- Release Information | 12710

Syntax

```
policy {
    application {
        duration interval;
        size count;
        threshold value;
```

```
    }
    destination-ip {
        duration interval;
        size count;
        threshold value;
    }
    policy match {
        duration interval;
        size count;
        threshold value;
    }
    source-ip {
        duration interval;
        size count;
        threshold value;
    }
}
```

Hierarchy Level

```
[edit security alarms potential-violation]
```

Description

Configure alarms for policy violation, based on source IP, destination IP, application, and policy rule.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Security Policies Overview

policy (Security Group VPN IKE)

IN THIS SECTION

- Syntax | 12711
- Hierarchy Level | 12711
- Description | 12711
- Options | 12712
- Required Privilege Level | 12712
- Release Information | 12712

Syntax

```
policy policy-name {
    description description;
    mode2 (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposals proposal-name;
}
```

Hierarchy Level

```
[edit security group-vpn member ike]
[edit security group-vpn server ike]
```

Description

Configure an IKE policy. An IKE policy defines a combination of security parameters (IKE proposals) to be used during IKE negotiation. It defines a peer address, the preshared key for the given peer, and the

proposals needed for that connection. During the IKE negotiation, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation sends all its policies to the remote peer, and the remote peer tries to find a match.

Options

policy *policy-name*	Name of the IKE policy. The policy name can be up to 32 alphanumeric characters long.
description *description*	Specify descriptive text for an IKE policy.
mode	Define the mode used for Internet Key Exchange (IKE) Phase 1 negotiations.
pre-shared-key	Define a preshared key for an IKE policy. Preshared keys are used to secure the Phase 1 SAs between the root-server and the sub-servers and between the sub-servers and the group members. Ensure that the preshared keys used are strong keys. On the sub-servers, the preshared key configured for the IKEpolicy RootSrv must match the preshared key configured on the root-server, and the preshared key configured for the IKE policy GMs must match the preshared key configured on the group members.
proposals *proposal-name*	Specify up to four Phase 1 proposals for an IKE policy. If you include multiple proposals, use the same Diffie-Hellman group in all of the proposals.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

policy (Security IKE)

IN THIS SECTION

- Syntax | 12713
- Hierarchy Level | 12714
- Description | 12714
- Options | 12714
- Required Privilege Level | 12716
- Release Information | 12716

Syntax

```
policy policy-name {
    blocklist blocklist-name;
    certificate {
        local-certificate certificate-id;
        peer-certificate-type (pkcs7 | x509-signature);
        policy-oids [ oid ];
        trusted-ca {
            ca-profile ca-profile-name;
            trusted-ca-group trusted-ca-group-name;
        }
    }
    description description;
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    seeded-pre-shared-key (ascii-text key | hexadecimal key);
    proposal-set (basic | compatible | prime-128 | prime-256 | standard | suiteb-gcm-128 | suiteb-
gcm-256);
    proposals proposal-name;
    reauth-frequency number;
}
```

Hierarchy Level

```
[edit security ike]
```

Description

IKE policies define a combination of security parameters (IKE proposals) to be used during IKE negotiation, including peer address, the preshared key for the given peer, and the proposals needed for that connection. During the IKE negotiation, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation sends all its policies to the remote peer, and the remote peer tries to find a match.

IKE proposals in the `policy` statement are evaluated in list order, from top to bottom, so when creating the policy, specify the highest priority proposal first, followed by the next highest priority, and so on.

Options

policy-name—Name of the IKE policy. The policy name can be up to 32 alphanumeric characters long.

`blocklist` *blocklist-name*—Specify the name of the corresponding remote peer's IKE identity blocklist. The blocklist is used for blocking IKE-IDs during IKE SA negotiation authentication phase.

`certificate`—Specify usage of a digital certificate to authenticate the virtual private network (VPN) initiator and recipient. For more information, See certificate.

`description` *description*—Specify the description of IKE policy.

`mode`—Define the mode used for Internet Key Exchange (IKE) Phase 1 negotiations. Use aggressive mode only when you need to initiate an IKE key exchange without ID protection, as when a peer unit has a dynamically assigned IP address. IKEv2 protocol does not negotiate using mode configuration. The device deletes existing IKE and IPsec SAs when you update the `mode` configuration in the IKE policy.

- `aggressive`—Aggressive mode.
- `main`—Main mode. Main mode is the recommended key-exchange method because it conceals the identities of the parties during the key exchange.

  Configuring `mode main` for group VPN servers or members is not supported when the remote gateway has a dynamic address and the authentication method is `pre-shared-keys`.

`pre-shared-key`—Define a preshared key for an IKE policy. The device deletes existing IKE and IPsec SAs when you update the `pre-shared-key` configuration in the IKE policy.

- `ascii-text` *`key`*—Specify a string of 1 to 255 ASCII text characters for the key. To include the special characters `( ) [ ] ! & ? |` enclose either the entire key string or the special character in quotation marks; for example `“str)ng”` or `str”)”ng`. Other use of quotation marks within the string is not allowed. With `des-cbc` encryption, the key contains 8 ASCII characters. With `3des-cbc` encryption, the key contains 24 ASCII characters.

- `hexadecimal` *`key`*—Specify a string of 1 to 255 hexadecimal characters for the key. Characters must be hexadecimal digits `0` through `9`, or letters `a` through `f` or `A` through `F`. With `des-cbc` encryption, the key contains 16 hexadecimal characters. With `3des-cbc` encryption, the key contains 48 hexadecimal characters.

`seeded-pre-shared-key`—Define a seeded preshared key in ASCII or hexadecimal format for an IKE policy. The `seeded-pre-shared-key` is a master key that is used to generate the `pre-shared-key` for the peers. Thus each peer will have different `pre-shared-key`. The advantage of this option is that each peer connection to gateway will have different pre-shared key, so if one of the peer's `pre-shared-key` is compromised, then the other peers are not impacted.

The peer preshared keys are generated using the master key configured as `seeded-pre-shared-key` and shared across the peers. To view the peer's pre-shared-key, execute the `show security ike pre-shared-key` command, share and configure the displayed pre-shared key in peer's device as pre-shared-key (in ASCII format). Master key is only configured in the gateway device and not shared to any peer.

You can retrieve the peer preshared key using the `show security ike pre-shared-key user-id` *`peer ike-id`* `master-key` *`master key`* or `show security ike pre-shared-key user-id` *`peer ike-id`* `gateway` *`gateway name`* command.

- `ascii-text` *`key`*—Configure a string of 1 to 255 ASCII text characters for the key. To include the special characters `( ) [ ] ! & ? |` enclose either the entire key string or the special character in quotation marks; for example `“str)ng”` or `str”)”ng`. Other use of quotation marks within the string is not allowed.

- `hexadecimal` *`key`*—Specify a string of 1 to 255 hexadecimal characters for the key. Characters must be hexadecimal digits `0` through `9`, or letters `a` through `f` or `A` through `F`.

`proposal-set`—Specify a set of default Internet Key Exchange (IKE) proposals.

`proposals` *`proposal-name`*—Specify up to four Phase 1 proposals for an IKE policy. If you include multiple proposals, use the same Diffie-Hellman group in all of the proposals.

`reauth-frequency` *`number`*—Configure the reauthentication frequency to trigger a new IKEv2 reauthentication. Reauthentication creates a new IKE SA, creates new child SAs within the IKE SA, and then deletes the old IKE SA. This option is disabled by default. umber of IKE rekeys that occurs before reauthentication occurs. If `reauth-frequency` is `1`, reauthentication occurs every time there is an IKE rekey. If `reauth-frequency` is `2`, reauthentication occurs at every other IKE rekey. If `reauth-frequency` is `3`, reauthentication occurs at every third IKE rekey.

- **Default:** 0 (disable)

- **Range:** 0-100

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

Support for `suiteb-gcm-128` and `suiteb-gcm-256` options added in Junos OS Release 12.1X45-D10.

Support for `policy-oids` option added in Junos OS Release 12.3X48-D10.

Support for `trusted-ca` option added in Junos OS Release 18.1R1.

Support for `reauth-frequency` option added in Junos OS Release 15.1X49-D60.

Support for `seeded-pre-shared-key` option added in Junos OS Release 21.1R1.

Support for `blocklist` option added in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

IPsec Overview

Example: Configuring AutoVPN with Pre-Shared Key

policy (Security Interfaces)

IN THIS SECTION

- Syntax | 12717
- Hierarchy Level | 12717
- Description | 12717
- Options | 12717
- Required Privilege Level | 12717

Release Information | 12718

Syntax

```
policy ike-peer-address {
    description policy-description;
    encoding (binary | pem);
    identity identity-name;
    local-certificate certificate-filename;
    local-key-pair private-public-key-file;
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposals [ proposal-names ];
}
```

Hierarchy Level

```
[edit security ike]
```

Description

Define an IKE policy.

Options

ike-peer-address—A tunnel address configured at the `[edit interfaces es]` hierarchy level.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an IKE Policy for Preshared Keys

Configuring an IKE Policy for Digital Certificates for an ES PIC

policy (Security IPsec)

IN THIS SECTION

- Syntax | 12718
- Hierarchy Level | 12719
- Description | 12719
- Options | 12719
- Required Privilege Level | 12721
- Release Information | 12721

Syntax

```
policy policy-name {
    description description;
    perfect-forward-secrecy keys (group1 | group14 | group19 | group2 | group20 | group24 |
group5 | group15 | group16 | group21);
    proposal-set (basic | compatible | prime-128 | prime-256 | standard | suiteb-gcm-128 |
suiteb-gcm-256);
    proposals proposal-name;
}
```

Hierarchy Level

```
[edit security ipsec]
```

Description

Define an IPsec policy. An IPsec policy defines a combination of security parameters (IPsec proposals) used during IPsec negotiation. It defines Perfect Forward Secrecy (PFS) and the proposals needed for the connection.

Options

name Name of the IPsec policy.

description Enter descriptive text for an IPsec policy.

perfect-forward-secrecy keys Specify Perfect Forward Secrecy (PFS) as the method that the device uses to generate the encryption key. PFS generates each new encryption key independently from the previous key. The device deletes existing IPsec SAs when you update the `perfect-forward-secrecy` configuration in the IPsec policy.

- Values:
  - `group1`—768-bit Modular Exponential (MODP) algorithm.
  - `group2`—1024-bit MODP algorithm.
  - `group5`—1536-bit MODP algorithm.
  - `group14`—2048-bit MODP group.
  - `group15`—3072-bit MODP algorithm.
  - `group16`—4096-bit MODP algorithm.
  - `group19`—256-bit random Elliptic Curve Groups modulo a Prime (ECP groups) algorithm.
  - `group20`—384-bit random ECP groups algorithm.
  - `group21`—521-bit random ECP groups algorithm.
  - `group24`—2048-bit MODP Group with 256-bit prime order subgroup.

proposal-set Define a set of default IPsec proposals.

- Values:
  - basic—IPsec basic proposal set. esp-des-sha and esp-des-md5.
    - Encapsulating Security Payload (ESP) protocol
    - Encryption algorithm—DES-CBC encryption algorithm
    - Authentication algorithm—SHA1 or MD5 authentication algorithm
  - compatible—IPsec compatible proposal set. esp-3des-sha, esp-3des-md5, esp-des-sha, and esp-des-md5.
    - ESP protocol
    - Encryption algorithm—3DES-CBC or DES-CBC encryption algorithm
    - Authentication algorithm—SHA1 or MD5 authentication algorithm
  - `prime-128`—Provides the following proposal set:
    - Encapsulating Security Payload (ESP) protocol
    - Encryption algorithm—Advanced Encryption Standard Galois/Counter mode (AES-GCM)128-bit
    - Authentication algorithm—None (AES-GCM provides both encryption and authentication)

    This option is not supported on Group VPNv2.
  - `prime-256`—Provides the following proposal set:
    - ESP protocol
    - Encryption algorithm—AES-GCM 256-bit
    - Authentication algorithm—None (AES-GCM provides both encryption and authentication)

    This option is not supported on Group VPNv2.
  - `standard`—esp-3des-sha and esp-aes128-sha

- ESP protocol
- Encryption algorithm—3DES-CBC or AES-CBC 128-bit encryption algorithm
- Authentication algorithm—SHA1 authentication algorithm

- `suiteb-gcm-128`—Provides the following proposal set:
  - ESP protocol
  - Encryption algorithm—AES-GCM 128-bit
  - Authentication algorithm—None (AES-GCM provides both encryption and authentication)

  This option is not supported on Group VPNv2.

- `suiteb-gcm-256`—Provides the following proposal set:
  - ESP protocol
  - Encryption algorithm—AES-GCM 256-bit
  - Authentication algorithm—None (AES-GCM provides both encryption and authentication)

  This option is not supported on Group VPNv2.

proposals ***proposal-name***

Specify up to four Phase 2 proposals for an IPsec policy. If you include multiple proposals, use the same Diffie-Hellman group in all of the proposals.

Proposals are evaluated in the order they appear on the list, from top down, so specify the highest priority first, followed by the next highest priority, and so on.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

Support for group 14 is added in Junos OS Release 11.1.

Support for `group14` options added in Junos OS Release 11.1.

Support for `group19`, `group20`, and `group24` options added in Junos OS Release 12.1X45-D10.

`group15`, `group16`, and `group21` options introduced in Junos OS Release 19.1R1 on SR5000 line of devices with `junos-ike` package installed.

Support for `suiteb-gcm-128` and `suiteb-gcm-256` options added in Junos OS Release 12.1X45-D10. Support for `prime-128` and `prime-256` options added in Junos OS Release 15.1X49-D40.

Starting in Junos OS Release 20.2R1, we've changed the help text description as `NOT RECOMMENDED` for the CLI options `group1`, `group2`, and `group5` for devices running IKED with `junos-ike` package installed.

Support for `group15`, `group16`, and `group21` options added in Junos OS Release 20.3R1 on vSRX Virtual Firewall instances with `junos-ike` package installed.

Support for `group15`, `group16`, and `group21` options added in Junos OS Release 21.1R1 on vSRX Virtual Firewall 3.0 instances with `junos-ike` package installed.

RELATED DOCUMENTATION

IPsec Overview

Installing Junos IKE package

policy (Security IPsec Interfaces)

IN THIS SECTION

- Syntax | **12723**
- Hierarchy Level | **12723**
- Description | **12723**
- Options | **12723**
- Required Privilege Level | **12723**
- Release Information | **12723**

Syntax

```
policy ipsec-policy-name {
    description description;
        perfect-forward-secrecy {
        keys (group1 | group 14 | group2 | group 5);
    }
    proposals [ proposal-names ];
}
```

Hierarchy Level

```
[edit security ipsec]
```

Description

Define an IPsec policy.

Options

ipsec-policy-name—Specify an IPsec policy name.

The remaining statements are explained separately.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the IPsec Policy for an ES PIC

policy (Security Policies)

IN THIS SECTION

- Syntax | 12724
- Hierarchy Level | 12726
- Description | 12726
- Options | 12727
- Required Privilege Level | 12727
- Release Information | 12727

Syntax

```
policy policy-name {
    description description;
    match {
        application {
            [application];
            any;
            junos-twamp;
        }
        destination-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
        }
        source-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
        }
        source-identity {
            [role-name];
            any;
            authenticated-user;
```

```
            unauthenticated-user;
            unknown-user;
        }
    }
    scheduler-name scheduler-name;
    then {
                            count {
            alarm {
                per-minute-threshold number;
                per-second-threshold number;
            }
        }
        deny;
        log {
            session-close;
            session-init;
        }
        permit {
            application-services {
                application-firewall {
                    rule-set rule-set-name;
                }
                application-traffic-control {
                    rule-set rule-set-name;
                }
                gprs-gtp-profile profile-name;
                gprs-sctp-profile profile-name;
                idp;
                redirect-wx | reverse-redirect-wx;
                ssl-proxy {
                    profile-name profile-name;
                }
                uac-policy {
                    captive-portal captive-portal;
                }
                utm-policy policy-name;
            }
            destination-address {
                drop-translated;
                drop-untranslated;
            }
            firewall-authentication {
                pass-through {
```

```
                    access-profile profile-name;
                    client-match user-or-group-name;
                    web-redirect;
                }
                user-firewall {
                    access-profile profile-name;
                    domain domain-name
                    ssl-termination-profile profile-name;
                }
                web-authentication {
                    client-match user-or-group-name;
                }
            }
            services-offload;
            tcp-options {
                initial-tcp-mss mss-value;
                reverse-tcp-mss mss-value;
                sequence-check-required;
                syn-check-required;
            }
            tunnel {
                ipsec-group-vpn group-vpn;
                ipsec-vpn vpn-name;
                pair-policy pair-policy;
            }
        }
        reject;
    }
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name]
```

Description

Define a security policy.

Options

policy-name Name of the security policy.

description Descriptive text about the security policy.

Range: 1 through 900 characters

The descriptive text should not include characters “<”, “>”, “&”, or “\n”.
The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 900 characters, the configuration might fail to take effect.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. The `services-offload` option added in Junos OS Release 11.4. Statement updated with the `source-identity` option and the `description` option added in Junos OS Release 12.1. Support for the `user-firewall` option added in Junos OS Release 12.1X45-D10. Support for the `initial-tcp-mss` and `reverse-tcp-mss` options added in Junos OS Release 12.3X48-D20.

The `junos-twamp` application is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

Security Policies Overview

policy (Services CoS Next Gen Services)

IN THIS SECTION

- Syntax | 12728
- Hierarchy Level | 12729
- Description | 12729
- Options | 12729
- Required Privilege Level | 12729
- Release Information | 12729

Syntax

```
policy policy-name {
    match {
        application [ application-names ];
        destination-address address;
        destination-address-range low minimum-value high maximum-value;
        destination-port port-number;
        destination-prefix-list list-name;
        source-address address;
        source-address-range low minimum-value high maximum-value;
        source-prefix-list list-name;
    }
    then {
        application-profile profile-name;
        dscp (alias | bits);
        forwarding-class class-name;
        reflexive; | revert; | reverse {
            application-profile profile-name;
            dscp (alias | bits);
            forwarding-class class-name;
        }
    }
}
```

Hierarchy Level

```
[edit services cos rule rule-name]
```

Description

Configure a policy in a services CoS rule. The policy specifies Differentiated Services (DiffServ) code point (DSCP) marking and forwarding-class assignment for packets that are processed by a service set. The policy identifies the matching conditions for packet source and destination addresses and for packet applications, and the actions to take on those packets. A CoS rule can include multiple policies.

The service set that the CoS rule is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

policy-name	Name of the policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Class of Service for Services PICs (Next Gen Services)

policy (Services IKE)

IN THIS SECTION

- Syntax | 12730
- Hierarchy Level | 12730
- Description | 12731
- Options | 12731
- Required Privilege Level | 12731
- Release Information | 12731

Syntax

```
policy policy-name {
    description description;
    local-certificate identifier;
    local-id (ipv4_addr ipv4-address | ipv6-addr ipv6-address | key-id identifier);
    version (1 | 2);
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposals [ proposal-names ];
    remote-id {
        any-remote-id;
        ipv4_addr [ values ];
        ipv6_addr [ values ];
        key_id [ values ];
    }
    respond-bad-spi max-responses
}
```

Hierarchy Level

```
[edit services ipsec-vpn ike]
```

Description

Define an IKE policy.

Options

policy-name—IKE policy name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IKE Policies

policy-set

IN THIS SECTION

- Syntax | 12732
- Hierarchy Level | 12732
- Description | 12732
- Options | 12732
- Required Privilege Level | 12733
- Release Information | 12733

Syntax

```
policy-set policy-set-name {
            policy policy-name {
              match {
                       source-address any;
                       destination-address any;
                       application any;
                       from-zone trust
                       to-zone untrust;
              }
           then permit;
            }
        }
    }
}
```

Hierarchy Level

```
[edit security policies]
```

Description

Specify policy that applies for the inner session created by VXLAN inner header

Options

source-address Source address of the traffic as matching criteria. You can specify one or more IP addresses, address sets, or wildcard addresses.

destination-address Destination address of the traffic as matching criteria. You can specify one or more IP addresses, address sets, or wildcard addresses.

application Name of the predefined or custom application or application set used as match criteria.

from-zone Name of the source zone used as match criteria.

to-zone Name of the destination zone used as match criteria.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1R1.

policy (SSM Maps)

IN THIS SECTION

- Syntax | 12733
- Hierarchy Level | 12733
- Description | 12734
- Options | 12734
- Required Privilege Level | 12734
- Release Information | 12734

Syntax

```
policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast ssm-map ssm-map-name],
[edit logical-systems logical-system-name routing-options multicast ssm-map ssm-map-name],
[edit routing-instances routing-instance-name routing-options multicast ssm-map ssm-map-name],
[edit routing-options multicast ssm-map ssm-map-name]
```

Description

Apply one or more policies to an SSM map.

Options

policy-names—Name of one or more policies for SSM mapping.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring SSM Mapping

policy (Stateful Firewall Rules Next Gen Services)

IN THIS SECTION

- Syntax | 12735
- Hierarchy Level | 12735
- Description | 12735
- Options | 12735
- Required Privilege Level | 12735
- Release Information | 12736

Syntax

```
policy policy-name {
    match {
        application [application-name];
        destination-address (address | any);
        destination-address-excluded address;
        source-address (address | any);
        source-address-excluded address;
    }
    then {
        count;
        deny;
        permit;
        reject;
    }
}
```

Hierarchy Level

```
[edit services policies stateful-firewall-rule rule-name]
```

Description

Configure one or more policies in a stateful firewall rule. Each policy identifies the matching conditions for a flow, and whether or not to allow the flow. Once a policy in the rule matches a flow, that policy is applied and no other policies in the rule are processed.

Options

policy-name	Name of the policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Stateful Firewalls for Next Gen Services

policy-statement

IN THIS SECTION

- Syntax (with terms) | 12736
- Syntax (without terms) | 12738
- Hierarchy Level | 12738
- Description | 12738
- Options | 12739
- Required Privilege Level | 12742
- Release Information | 12742

Syntax (with terms)

```
policy-statement policy-name {
    term term-name {
        from {
            as-path-neighbors (as-list | as-list-group);
            as-path-origins (as-list | as-list-group);
            as-path-transits (as-list | as-list-group);
            as-path-unique-count count (equal | orhigher | orlower);
            as-path-calc-length count (equal | orhigher | orlower);
            community BGP community;
            family family-name;
            match-conditions;
            neighbor address;
```

```
            policy subroutine-policy-name;
            prefix-list prefix-list-name;
            prefix-list-filter prefix-list-name match-type <actions>;
            programmed;
            protocol protocol-name;
            route-filter destination-prefix match-type <actions> invert-match;
            validation-database-instance {
                                 database <database-name> state (valid|invalid|unknown);
                                 state (valid|invalid|unknown);
                                 }
            source-address-filter source-prefix match-type <actions>;
            tag value;
            traffic-engineering;
        }
        to {
            match-conditions;
            policy subroutine-policy-name;
        }
        then actions;
    }
    then {
        advertise-locator;

        aggregate-bandwidth;
        dynamic-tunnel-attributes dynamic-tunnel-attributes;
        limit-bandwidth limit-bandwidth;
        load-balance (adaptive | consistent-hash | destination-ip-only | per-flow | per-packet |
per-prefix | profile1 | profile2 | random | source-ip-only | symmetric-consistent-hash);
        multipath-resolve;
        no-entropy-label-capability;
        prefix-attribute-flags;
        prefix-segment {
            index index;
            node-segment;
        }
        priority (high | medium | low);
        resolution-map map-name;
        set-down-bit
        srv6 sid-value;
    }
}
```

Syntax (without terms)

The policy statement configuration can be used without terms. An example configuration is as below.

```
policy-statement policy-name from <match-condition> then <action>
```

Hierarchy Level

```
[edit dynamic-profiles profile-name policy-options],
[edit logical-systems logical-system-name policy-options],
[edit policy-options]
```

Description

Define a routing policy, including subroutine policies.

A *term* is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of `accept` or `reject` is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either `accept` or `reject`) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the `[edit policy-options]` hierarchy level by `policy-statement` *`policy-name`* in alphabetical order, enter the `show policy-options` configuration command.

The statements are explained separately.

Options

`actions`—(Optional) One or more actions to take if the conditions match. The actions are described in Configuring Flow Control Actions.

`family` *`family-name`*—(Optional) Specify an address family protocol. Specify `inet` for IPv4. Specify `inet6` for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify `iso`. For IPv4 multicast VPN traffic, specify `inet-mvpn`. For IPv6 multicast VPN traffic, specify `inet6-mvpn`. For multicast-distribution-tree (MDT) IPv4 traffic, specify `inet-mdt`. For BGP route target VPN traffic, specify `route-target`. For traffic engineering, specify `traffic-engineering`.

> **NOTE**: When `family` is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the "family" on page 6900 statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

`from`—(Optional) Match a route based on its source address.

`as-path-origins (as-list | as-list-group)`—Compares the AS that originated the route. Evaluates if the right most AS number on the AS path belongs to the `as-list` or `as-list-group` specified in the `as-path-origins` configuration statement. In the case where the route has been aggregated, and the location of the originating AS contains an AS-set, the `as-path-origins` operator evaluates to true if any AS contained in the AS-set belongs to the `as-list` or `as-list-group` specified in the `as-path-origins` configuration statement.

`as-path-neighbors (as-list | as-list-group)`—Compares the neighbor AS in the AS path. Evaluates if the first AS number on the AS path matches the `as-list` or `as-list-group` specified in the `as-path-neighbors` configuration statement. If the neighboring AS location happens to be an AS-set, the `as-path-neighbors` operator evaluates to true if any AS contained in the AS-set belongs to the `as-list` or `as-list-group` specified in the `as-path-neighbors` configuration statement.

`as-path-transits (as-list | as-list-group)`—Compares any AS in the AS-Path. Evaluates when any AS belongs to the `as-list` or `as-list-group` specified in the `as-path-transit` configuration statement. In the case of AS-set, the `as-path-transit` operator compares all the ASes in the AS-set.

`as-path-calc-length` *`count`* `(equal | orhigher | orlower)`—(Optional) Specify a number from 0 through 1024 to filter routes based on the number of calculated autonomous systems (ASs) in the AS path.

NOTE:

- ASs in a sequence count as 1.
- AS sets count as 1.
- BGP confederation segments count as 0.

`as-path-unique-count` *`count`* `(equal | orhigher | orlower)`—(Optional) Specify a number from 0 through 1024 to filter routes based on the total number of unique non-BGP confederation autonomous systems (ASs) in the AS path.

NOTE: Duplicate AS numbers are ignored for the count.

`advertise-locator`—(Optional) Enable IS-IS to summarize and advertise locator prefixes.

Range: *0-255*

`aggregate-bandwidth [transitive | non-transitive]`—(Optional) Enable BGP to advertise aggregate outbound link bandwidth for load balancing. By default, the aggregate link-bandwidth community is transitive. You can choose to use the BGP link-bandwidth community as a non-transitive attribute.

`aggregate-bandwidth divide-equal`—(Optional) Enable equal division of total link-bandwidth by number of peers in the advertising group. If you modify the BGP group, the updated aggregate bandwidth is sent to all peers.

`dynamic-tunnel-attributes` *`dynamic-tunnel-attributes`*—(Optional) Choose a set of defined dynamic tunnel attributes for forwarding traffic over V4oV6 tunnels.

`match-conditions`—(Optional in `from` statement; required in `to` statement) One or more conditions to use to make a match. The qualifiers are described in Routing Policy Match Conditions.

`multipath-resolve` *`multipath-resolve`*–(Optional) Enable the use of all paths for resolution over the specified prefix.

`limit-bandwidth` *`limit-bandwidth`*—(Optional) Specify the limit for advertised aggregate outbound link bandwidth for load balancing.

- **Range:** 0 through 4,294,967,295 bytes

`no-entropy-label-capability`—(Optional) Disable the entropy label capability advertisement at egress or transit routes specified in the policy.

`neighbor`—(Optional) Specify a neighbor for route filtering.

`priority (high | medium | low)`—(Optional) Configure the priority for an IS-IS route to change the default order in which the routes are installed in the routing table, in the event of a network topology change.

`policy` *`subroutine-policy-name`*—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”). Policy names cannot take the form `__.*-internal__`, as this form is reserved. For information about how to configure subroutines, see Understanding Policy Subroutines in Routing Policy Match Conditions.

`policy-name`—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

`prefix-list` *`prefix-list-name`*—Name of a list of IPv4 or IPv6 prefixes.

`prefix-list-filter` *`prefix-list-name`*—Name of a prefix list to evaluate using qualifiers; *`match-type`* is the type of match, and *`actions`* is the action to take if the prefixes match.

`programmed`—(Optional) Allow policy matches for routes injected by JET APIs.

`protocol` *`protocol-name`*—Name of the protocol used to control traffic engineering database import at the originating point. For example,

`protocol` *`srv6`* matches SRv6 routes. Note that SRv6 routes are added to the routing infrastructure by BGP. To view all SRv6 routes, run `show route protocol srv6`.

Starting in Junos OS Release 19.1R1, you can specify options to match label IS-IS and label OSPF routes using the `l-isis` and `l-ospf` options, respectively. The `isis` options matches all IS-IS routes, excluding labelled IS-IS routes. The `ospf` option matches all OSPF routes, including OSPFv2, OSPFv3 and labelled OSPF routes.

`resolution-map`—(Optional) Set resolution map modes. A given resolution-map can be shared across multiple policy-statements.

`route-filter` *`destination-prefix match-type <actions>`*—(Optional) List of routes on which to perform an immediate match; *`destination-prefix`* is the IPv4 or IPv6 route prefix to match, *`match-type`* is the type of match (see Configuring Route Lists), and *`actions`* is the action to take if the *`destination-prefix`* matches.

When `invert-match` *match-type* is configured, it will return `true` if the route doesn't pass any prefixes defined in the `route-filter-list`. When you add the `invert-match` *match-type*, you only need to provide a list of prefixes it *should not* match instead of a list of prefixes it should.

`source-address-filter` *`source-prefix match-type <actions>`*—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. *`source-prefix`* is the IPv4 or IPv6 route prefix to match, *`match-type`* is the type of match (see Configuring Route Lists), and *`actions`* is the action to take if the *`source-prefix`* matches.

`tag` *`value`*—(Optional) A numeric value that identifies a route. You can tag certain routes to prioritize them over other routes. In the event of a network topology change, Junos OS updates these routes in the routing table before updating other routes with lower priority. You can also tag some routes to identify and reject them based on your requirement.

`term` *`term-name`*—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (“ ”). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the `term` statement when defining match conditions and actions.

`to`—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

`then`—(Optional) Actions to take on matching routes. The actions are described in Configuring Flow Control Actions and Configuring Actions That Manipulate Route Characteristics.

`set-down-bit`—(Optional) Configure this option to aggregate leaked locator routes using routing policies.

`srv6` *`SID value`* - Enter the Segment Routing over IPv6 (SRv6) SID value. You use this configuration to define SID values for routes, which can also be SRv6 routes.

`validation-database-instance`—(Optional) Name to identify a validation-state with database name.`database-name <database-name>`—(Optional) Route Validation Database name to be looked at. `state (valid|invalid|unknown)`—(Optional) Name to identify a validation-state

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.

`inet-mdt` option introduced in Junos OS Release 10.0R2.

`route-target` option introduced in Junos OS Release 12.2.

`protocol` and `traffic-engineering` options introduced in Junos OS Release 14.2.

`no-entropy-label-capability` option introduced in Junos OS Release 15.1.

`priority` and `tag` *`value`* options introduced in Junos OS Release 17.1.

`as-path-unique-count` option introduced in Junos OS Release 17.2R1.

`prefix-segment` option introduced in Junos OS Release 17.2R1 for MX Series routers, PTX Series routers, QFX5100 switches, and QFX10000 switches.

`multipath-resolve` and `dynamic-tunnel-attributes` options introduced in Junos OS Release 17.3R1.

`aggregate-bandwidth` and `limit-bandwidth` *`limit-bandwidth`* options introduced in Junos OS Release 17.4R1 for MX Series, PTX Series, and QFX Series.

l-isis and *l-ospf* keywords at the `protocol` option is introduced in Junos OS Release 19.1R1.

`resolution-map` statement introduced in Junos OS Release 19.2R1-S1 on MX and PTX Series routers.

`lsp` and `lsp-regex` options introduced in Junos OS Release 19.4R1.

`as-path-neighbors`, `as-path-origins`, and `as-path-transits` statements introduced in Junos OS Release 21.3R1.

`advertise-locator` and `set-down-bit` options introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

dynamic-db

Understanding Source Packet Routing in Networking (SPRING)

policy-stats

IN THIS SECTION

- Syntax | 12744
- Hierarchy Level | 12744
- Description | 12744
- Options | 12744
- Required Privilege Level | 12744
- Release Information | 12744

Syntax

```
policy-stats {
    system-wide (disable | enable) ;
}
```

Hierarchy Level

```
[edit security policies]
```

Description

Configure systemwide policies statistics. The systemwide policies statistics are disabled by default.

Options

disable Disable systemwide policy statistics.

enable Enable systemwide policy statistics.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

show security policies | 25028

policy (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 12745
- Hierarchy Level | 12746
- Description | 12746
- Options | 12746
- Required Privilege Level | 12746
- Release Information | 12746

Syntax

```
policy policy-name {
    inet {
        drop-policy rule-name {
            from {
                apply-groups group-name;
                apply-groups-except group-name;
                destination-address address;
                destination-port port-number;
                dscp dscp-value;
                protocol protocol;
                source-address address;
                source-port port-number;
            }
        }
    }
    inet6 {
        drop-policy rule-name {
            from {
                apply-groups group-name;
                apply-groups-except group-name;
                destination-address address;
                destination-port port-number;
                dscp dscp-value;
                protocol protocol;
```

```
                source-address address;
                source-port port-number;
            }
        }
    }
}
```

Hierarchy Level

```
[edit services radius-flow-tap]
```

Description

Specify the policy that is applied to mirrored packets sent to a mediation device.

Options

policy-name—Name of the policy from which to drop traffic.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

> *Subscriber Secure Policy Overview*
>
> *Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview*

policy (System Security Profile)

IN THIS SECTION

- Syntax | 12747
- Hierarchy Level | 12747
- Description | 12747
- Options | 12748
- Required Privilege Level | 12748
- Release Information | 12748

Syntax

```
policy {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of security policies that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 80000

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways

policy-with-count

IN THIS SECTION

- Syntax | 12749
- Hierarchy Level | 12749
- Description | 12749
- Options | 12750
- Required Privilege Level | 12750
- Release Information | 12750

Syntax

```
policy-with-count {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of security policies with a count that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 1024

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways

policy-map-marking

IN THIS SECTION

- Syntax | 12751
- Hierarchy Level | 12751
- Description | 12751
- Required Privilege Level | 12751
- Release Information | 12751

Syntax

```
policy-map-marking;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Enable policy maps on a logical interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

policy-map (Class of Service) | 12686

Assigning Rewrite Rules on a Per-Customer Basis Using Policy Maps

polling-interval (Network Management)

IN THIS SECTION

- Syntax | 12752
- Hierarchy Level | 12752
- Description | 12752
- Options | 12753
- Required Privilege Level | 12753
- Release Information | 12753

Syntax

```
polling-interval polling-interval
```

Hierarchy Level

```
[edit system packet-forwarding-options hw-resource-monitor]
```

Description

Configure a single polling interval for all hardware resources. The same polling interval is used for all the resources.

Options

polling-interval

- **Range:** 10 through 86400000 milliseconds
- **Default:** 1 second.
- This is an optional configuration.

Required Privilege Level

Edit

Release Information

Statement introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

No Link Title

resource-list (Network Management) | **15610**

monitor-profile | **10912**

show system packet-forwarding-options hw-resource-utilization-info | **26952**

show system packet-forwarding-options hw-resource-monitor utilization-info | **26948**

polling-interval

IN THIS SECTION

- Syntax | **12754**
- Hierarchy Level | **12754**
- Description | **12754**
- Default | **12754**
- Options | **12754**
- Required Privilege Level | **12754**

Release Information | 12755

Syntax

```
polling-interval seconds;
```

Hierarchy Level

```
[edit protocols sflow],
[edit protocols sflow interfaces interface-name]
```

Description

Configure the interval (in seconds) that the device waits between port statistics update messages. *Polling* refers to the device successively gathering various statistics for the network interfaces configured for sFlow monitoring and exporting the statistics to the configured sFlow collector.

Default

If no polling interval is configured for a particular interface, the device uses the global polling interval configured at the `[edit protocols sflow]` hierarchy level. If no global interval is configured, the device uses the default polling interval of 20 seconds.

Options

seconds—Number of seconds between successive samples of interface statistics. Specifying a value of 0 (zero) disables the polling.

- **Range:** 0 through 3600 seconds
- **Default:** 20 seconds

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring sFlow Technology for Network Monitoring (CLI Procedure)

Example: Monitoring Network Traffic Using sFlow Technology

Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches

Example: Configuring sFlow Technology to Monitor Network Traffic on MX Series Routers

pon (Access-Line Rate Adjustment)

IN THIS SECTION

- Syntax | 12755
- Hierarchy Level | 12756
- Description | 12756
- Options | 12757
- Required Privilege Level | 12759
- Release Information | 12759

Syntax

```
pon {
    gpon {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    other {
        overhead-adjust percent;
        overhead-bytes bytes;
```

```
        total-adjust percent;
    }
    twdm-pon {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    type tlv-value {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    wdm-pon {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    xg-pon1 {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
    xgs-pon {
        overhead-adjust percent;
        overhead-bytes bytes;
        total-adjust percent;
    }
}
```

Hierarchy Level

```
[edit system access-line]
```

Description

Configure adjustments to actual PON data rates as follows:

- Multiply the actual data rate by a percentage for only downstream rates or for both upstream an downstream rates

- Adjust the encapsulation overhead by adding to or subtracting from the total cell or frame bytes a specified number of bytes

The actual (unadjusted) downstream and upstream data rates, PON line type, and encapsulation mode are received from the access node by the ANCP agent in ANCP port messages, or by the PPPoE daemon from the PPPoE intermediate agent (PPPoE-IA) in PADI or PADR messages. The ANCP agent or PPPoE daemon subsequently adjusts rates and bytes based on the configuration.

If the PON-Access-Type TLV (0x92) is not received in either the ANCP Port Status message or PPPoE-IA tags, the default adjustment leaves the rates and bytes unchanged.

Adjustments are applied to all subscribers using access lines of the specific PON line type. Depending on the value, it may be reported to AAA, CoS, or both:

- Adjusted and unadjusted downstream and upstream rates are always reported to AAA in response to an AAA request.
- Adjusted and unadjusted downstream rates and overhead byte adjustments are reported to CoS, but only when you include the `qos-adjust` statement at the `[edit protocols ancp]` hierarchy level.
- Overhead byte adjustments are not reported to AAA.

AAA reports the adjusted values to the RADIUS server in the Access-Request and Accounting-Request messages through Juniper Networks VSAs 26-141, Downstream-Calculated-Qos-Rate Rate, and 26-142, Upstream-Calculated-Qos-Rate. A change in value triggers an immediate Interim-Accounting message to the RADIUS server if you have configured the "ancp-speed-change-immediate-update" on page 2251 statement.

The ANCP agent reports these values to the LAC in an L2TP network. The LAC passes the rates to the LNS in the following AVPs and messages:

- AVP 24, Tx Connect Speed (ICCN message)—For the initial total rate adjustment to ONT/ONU-Peak-Data-Rate-Downstream TLV (0x94).
- AVP 38, Rx Connect Speed (ICCN message)—For the initial total rate adjustments to ONT/ONU-Peak-Data-Rate-Upstream TLV (0x95).
- AVP 97, Connect Speed Update (CSUN message)—For subsequent changes to the initial rates reported in AVP 24 and AVP 38.

Options

gpon Sets attributes for GPON access lines.

other Sets attributes for access lines of type OTHER. This might used, for example, when an OLT uses a new PON type that is not yet formally supported by the Broadband Forum.

overhead-adjust *percentage*

Adjusts the actual downstream rates for all subscribers on an access line or PON tree of the specified types by multiplying the rate by the specified percentage. This adjustment accounts for the Layer 1 overhead for the PON type. The rates are adjusted as follows:

- For the subscriber access line, the adjustment is made to the downstream rate reported in the ONT/ONU-Peak-Data-Rate-Downstream TLV (0x94). The adjusted rate is applied to the logical interface or the (child) interface set for the subscriber access line, as determined by the CoS adjustment control profile. This adjustment applies for FTTH connections.
- For the PON tree, the adjustment is made to the downstream rate reported in the PON-Tree-Maximum-Data-Rate-Downstream TLV (0x98). The adjusted rate is applied to the PON tree (parent) interface set, as determined by the CoS adjustment control profile. This adjustment applies for both FTTH and FTTB connections. For FTTB, this is the rate on the OLT to DPU-P access line.
- **Range:** 80 through 100 percent
- **Default:** 100 percent

overhead-bytes *bytes*

Adjusts the actual downstream cell overhead for all subscribers on the specified access line or PON tree by adding or subtracting the specified number of bytes. The adjustment accounts for the traffic encapsulation overhead. It shapes the logical interface or the (child) interface set for the subscriber access line or the PON tree (parent) interface set, as determined by the CoS adjustment control profile. The overhead is adjusted as follows:

- For the subscriber access line, the adjustment is made to the overhead reported in the ONT/ONU-Peak-Data-Rate-Downstream TLV (0x94). The adjusted overhead is applied to the logical interface or the (child) interface set for the subscriber access line, as determined by the CoS adjustment control profile. This adjustment applies for FTTH connections.
- For the PON tree, the adjustment is made to the overhead reported in the PON-Tree-Maximum-Data-Rate-Downstream TLV (0x98). The adjusted overhead is applied to the PON tree (parent) interface set, as determined by the CoS adjustment control profile. This adjustment applies for both FTTH and FTTB connections. For FTTB, this is the overhead on the OLT to DPU-P access line.
- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

twdm-pon

Sets attributes for TWDM-PON access lines.

total-adjust *percentage*

Adjusts the downstream and upstream data rates for all subscribers on an access line of the specified type by multiplying the rate by the specified percentage. This adjustment accounts for the total Layer 1 and encapsulation overhead for the PON type. The adjustment is made to the ONT/ONU-Peak-Data-Rate-Downstream TLV (0x94) and the ONT/ONU-Maximum-Data-Rate-Upstream TLV (0x95). The adjustments apply to FTTH connections.

- **Range:** 1 through 100 percent
- **Default:** 100 percent

type *tlv-value*

Sets attributes for access lines by specifying the unsigned integer value of the PON-Access-Type TLV (0x92) to reference the PON type. This option enables the `pon` statement to be used for PON access line types that might be introduced in the future.

- **Range:** 6 through 4294967295

wdm-pon

Sets attributes for WDM-PON access lines.

xg-pon1

Sets attributes for XG-PON access lines.

xgs-pon

Sets attributes for XGS-PON access lines.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Neighbors and Operations

pool (NAT)

IN THIS SECTION

- Syntax | 12760
- Hierarchy Level | 12761
- Description | 12761
- Options | 12761
- Required Privilege Level | 12761
- Release Information | 12761

Syntax

```
pool nat-pool-name {
    address ip-prefix</prefix-length>;
    address-allocation round-robin;
    address-range low minimum-value high maximum-value;
    app- mapping-timeout app-mapping-timeout;
    ei-mapping-timeout ei-mapping-timeout;
    limit-ports-per-address number;
    mapping-timeout mapping-timeout;
    pgcp {
        hint [ hint-strings ];
        ports-per-session ports;
        remotely-controlled:
    }
    port {
        automatic (sequential | random-allocation);
        range low minimum-value high maximum-value random-allocation;
        preserve-parity;
        preserve-range;
        secured-port-block-allocation {
            active-block-timeout timeout-seconds;
            block-size block-size;
            max-blocks-per-user max-blocks;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit services nat]
```

Description

Specify the NAT name and properties.

Options

nat-pool-name—Identifier for the NAT address pool.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: If you make any configuration changes to a NAT pool that has secured port block allocation configured, you must deactivate the service set first, do required pool attribute changes, and activate the service set. We also strongly recommend that you perform this procedure if you make any changes to the NAT pool configuration, even when secured port block allocation is not configured.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`pgcp` statement added in Junos OS Release 8.4.

`remotely-controlled` and `ports-per-session` statements added in Junos OS Release 8.5.

`hint` statement added in Junos OS Release 9.0.

`address-allocation` statement added in Junos OS Release 11.2.

`sequential` statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring Pools of Addresses and Ports for Network Address Translation Overview

pool

IN THIS SECTION

- Syntax | 12762
- Hierarchy Level | 12762
- Description | 12763
- Options | 12763
- Required Privilege Level | 12763
- Release Information | 12763

Syntax

```
pool pool-name <priority priority>;
```

Hierarchy Level

```
[edit interfaces br-pim/0/port dialer-options],
[edit interfaces umd0 dialer-options],
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers, for logical and physical ISDN interfaces, specify the dial pool. The dial pool allows logical (dialer) and physical (`br-pim/0/port`) interfaces to be bound together dynamically on a per-call basis. On a dialer interface, `pool` directs the dialer interface which dial pool to use. On `br-pim/0/port` interface, `pool` defines the pool to which the interface belongs.

Options

pool-name—Pool identifier.

`priority` *priority*—(Physical `br-pim/0/port` interfaces only) Specify a priority value of 0 (lowest) to 255 (highest) for the interface within the pool.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

pool (Access)

IN THIS SECTION

- Syntax | **12764**
- Hierarchy Level | **12765**
- Description | **12765**
- Options | **12765**
- Required Privilege Level | **12765**
- Release Information | **12765**

Syntax

```
pool pool-name {
    family {
        inet {
            dhcp-attributes {
                boot-file boot file name;
                boot-server boot server name;
                domain-name domain name;
                grace-period seconds;
                maximum-lease-time (seconds | infinite);
                name-server ipv4-address;
                name-server address;
                netbios-node-type (b-node | h-node | m-node | p-node);
                option dhcp option-identifier-code;
                option-match {
                    option-82 {
                        circuit-id match-value;
                        remote-id match-value;
                    }
                }
                router IPv4 address;
                server-identifier IP address;
                tftp-server server name;
                wins-server IPv4 address;
            }
            host hostname;
            network network address;
            range range-name {
                high upper-limit;
                low lower-limit;
            }
            xauth-attributes {
                primary-dns IP address;
                primary-wins IP address;
                secondary-dns IP address;
                secondary-wins IP address;
            }
        }
        inet6 {
            dhcp-attributes {
                dns-server IPv6-address;
```

```
                grace-period seconds;
                maximum-lease-time seconds;
                option dhcp-option-identifier-code;
                sip-server-address IPv6-address;
                sip-server-domain-name domain-name;
            }
            prefix IPv6-network-prefix;
            range range-name {
                high upper-limit;
                low lower-limit;
                prefix-length delegated-prefix-length;
            }
    link pool-name;
}
```

Hierarchy Level

```
[edit access address-assignment]
```

Description

Configure the name of an address assignment pool. The remaining statements are explained separately.

Options

pool-name—Name assigned to the address-assignment pool.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.4 of Junos OS.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Obtaining Username and Role Information Through Firewall Authentication

pool (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 12766
- Hierarchy Level | 12767
- Description | 12767
- Options | 12767
- Required Privilege Level | 12768
- Release Information | 12768

Syntax

```
pool pool-name {
    active-drain;
    family family {
        dhcp-attributes {
            [ protocol-specific attributes ]
        }
        excluded-address ip-address;
        excluded-range name low minimum-value high maximum-value;
        host hostname {
            hardware-address mac-address;
            ip-address ip-address;
        }
        network ip-prefixprefix-length>;
        prefix ipv6-prefix;
        range range-name {
            high upper-limit;
            low lower-limit;
```

```
            prefix-length prefix-length;
        }
    }
    hold-down;
    link pool-name;
}
```

Hierarchy Level

```
[edit access address-assignment]
[edit routing-instances routing-instances-name access address-assignment]
```

Description

Configure the name of an address-assignment pool.

> **NOTE**: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

pool-name Name assigned to the address-assignment pool.

active-drain Configure the DHCP local server to stop allocating addresses from this pool. When this is configured, the DHCP local server gracefully shifts clients from this address pool to an alternative pool for which active drain is not configured. When existing clients with an address from this pool submit a DHCPv4 request or DHCPv6 renew, they receive a NAK, forcing them to renegotiate. The server responds with a DHCPv4 offer or DHCPv6 advertise message with an address from a different pool.

family Configure the protocol family for the address-assignment pool.

The options for this statement are explained separately. Click the linked statement for details.

hold-down Configure an address-assignment pool that is currently in use to be unavailable for further address allocation. When a pool is in the hold-down state, the pool is no longer used to allocate IP addresses for subscribers. Current subscribers who previously obtained an address from the pool are not affected; they can continue to renew their leases. As each of

these users disconnects, their address is not reallocated. The pool becomes inactive when all subscribers have disconnected and their addresses are returned to the pool.

link Designate a secondary address-assignment pool that is linked to the pool being configured. When the pool being configured has no addresses available for allocation, the secondary pool can be searched for a free address. You can configure a chain of linked pools, but you cannot directly link more than one pool to or from any other pool. Each linked pool in the chain serves as a backup pool for the pool immediately before it in the chain.

- **Values:** *pool-name*—Name assigned to the secondary address-assignment pool.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support at the `[edit routing-instances routing-instances-name access address-assignment]` hierarchy level at tenant system level introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management

DHCP Client Attribute and Address Assignment

pool-default-port-range

IN THIS SECTION

Syntax | 12769

- Hierarchy Level | 12769
- Description | 12769
- Options | 12769
- Required Privilege Level | 12770
- Release Information | 12770

Syntax

```
pool-default-port-range lower-port-range to upper-port-range;
```

Hierarchy Level

```
[edit security nat source]
```

Description

Set the global default single port range for source NAT pools with port translation. If the port range in source NAT pools is not specified, the configured default port range is used. If neither the port range in source NAT pools nor the default port range are configured, the default single port range is 1024 through 63,487.

To set the global twin port range for source NAT pools with port translation, use the `pool-default-twin-port-range` statement at the [edit security nat source] hierarchy. The twin port range is 63,488 through 65,535.

To set the single port range for a specific pool, use the `port range` *`port-low`* `(to` *`port-high`*`)` statement at the [edit security nat source pool] hierarchy level.

Options

- *`lower-port-range`*—Specify the lower limit of the port range.
- *`upper-port-range`*—Specify the upper limit of the port range.
- **Range:** 1024 through 63,487. To view pool information, use the **show security nat source pool** command.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

pool (Security Source NAT) | 12781

pool-default-twin-port-range | 12771

pool-default-port-range (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 12770
- Hierarchy Level | 12771
- Description | 12771
- Options | 12771
- Required Privilege Level | 12771
- Release Information | 12771

Syntax

```
pool-default-port-range port-low to port-high;
```

Hierarchy Level

```
[edit services nat source]
```

Description

Configure a global default port range for NAT pools that use port translation. This port range is used when a NAT pool does not specify a port range and does not specify automatic port assignment.

Options

port-low	The lower end of the port range.
port-high	The upper end of the port range.

- **Range:** 1024 through 65,535

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

pool-default-twin-port-range

IN THIS SECTION

- Syntax | 12772
- Hierarchy Level | 12772
- Description | 12772
- Options | 12772

Required Privilege Level | 12772
Release Information | 12773

Syntax

```
pool-default-twin-port-range lower-port-range to upper-port-range;
```

Hierarchy Level

```
[edit security nat source]
```

Description

Specify the global default twin port range for all source pools. Two ports within range (63488, 65535) are allocated at a time for RTP/RTCP applications such as SIP, H.323, and RTSP for source pools with PAT.

The default twin port range is 2048. If you have an SRX5400, SRX5600, or SRX5800 device that supports a maximum of 1 million IP addresses, use this option to limit the twin port range and avoid exceeding the port capacity of 384 million.

To set the twin port range for a specific pool, use the `port range twin-port` *`port-low`* `(to` *`port-high`*`)` statement at the [edit security nat source pool] hierarchy level.

Options

- *`lower-twin-port-range`*—Specify the lower limit of the port range.
- *`upper-twin-port-range`*—Specify the upper limit of the port range.
- **Range:** 63,488 through 65,535.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

pool (Security Source NAT) | 12781

Understanding Source NAT Pool Capacities

pool (Destination NAT Next Gen Services)

IN THIS SECTION

- Syntax | 12773
- Hierarchy Level | 12773
- Description | 12774
- Options | 12774
- Required Privilege Level | 12774
- Release Information | 12774

Syntax

```
pool nat-pool-name{
    address address-prefix;
}
```

Hierarchy Level

```
[edit services nat destination]
```

Description

Configure a set of addresses used for Network Address Translation (NAT) of destination addresses.

Options

nat-pool-name Name of the NAT pool.

If you are configuring twice NAT, do not use the same name that you use for the source pool.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

pool (DHCP Local Server Overrides)

IN THIS SECTION

- Syntax | 12775
- Hierarchy Level | 12775
- Description | 12776
- Options | 12776
- Required Privilege Level | 12776
- Release Information | 12776

Syntax

```
pool pool-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server overrides process-inform],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 overrides process-inform],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name overrides process-inform],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name interface interface-name overrides process-
inform],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name overrides process-inform],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name interface interface-name overrides process-inform],
[edit logical-systems logical-system-name system services dhcp-local-server overrides process-
inform],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 overrides
process-inform],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name overrides process-inform],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name interface interface-name overrides process-inform],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
overrides process-inform],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
interface interface-name overrides process-inform],
[edit routing-instances routing-instance-name system services dhcp-local-server overrides process-
inform],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 overrides
process-inform],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name overrides process-inform],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name interface interface-name overrides process-inform],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
```

```
overrides process-inform],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
interface interface-name overrides process-inform],
[edit system services dhcp-local-server overrides process-inform],
[edit system services dhcp-local-server dhcpv6 overrides process-inform],
[edit system services dhcp-local-server dhcpv6 group group-name overrides process-inform],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name
overrides process-inform],
[edit system services dhcp-local-server group group-name overrides process-inform],
[edit system services dhcp-local-server group group-name interface interface-name overrides
process-inform]
```

Description

Configure DHCP or DHCPv6 local server to reply to DHCP information request messages (DHCPINFORM for DHCPv4 and INFORMATION-REQUEST for DHCPv6) with information taken from the specified pool without interacting with AAA.

Options

pool-name Name of the address pool, which must be configured within `family inet` for DHCP local server and within `family inet6` for DHCPv6 local server.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Enabling Processing of Client Information Requests

Overriding the Default DHCP Local Server Configuration Settings

pool-match-order

IN THIS SECTION

- Syntax | 12777
- Hierarchy Level | 12777
- Description | 12777
- Default | 12778
- Required Privilege Level | 12778
- Release Information | 12778

Syntax

```
pool-match-order {
    external-authority;
    ip-address-first;
    option-82;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server],
[edit logical-systems logical-system-name system services dhcp-local-server],
[edit routing-instances routing-instance-name system services dhcp-local-server],
[edit system services dhcp-local-server]
```

Description

Configure the order in which the DHCP local server uses information in the DHCP client PDU to determine how to obtain an address for the client.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Default

DHCP local server uses the `ip-address-first` method to determine which address pool to use.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use

Understanding Differences Between Legacy DHCP and Extended DHCP

pool (NAT Rule Next Gen Services)

IN THIS SECTION

- Syntax | 12778
- Hierarchy Level | 12779
- Description | 12779
- Required Privilege Level | 12779
- Release Information | 12779

Syntax

```
pool nat-pool-name;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name then source-nat],
[edit services nat source rule-set rule-set rule rule-name then source-nat]
```

Description

Specify the name of the NAT pool that contains the addresses or subnets to which addresses are translated.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

pool (Security Destination NAT)

IN THIS SECTION

- Syntax | 12780
- Hierarchy Level | 12780
- Description | 12780
- Options | 12780
- Required Privilege Level | 12780
- Release Information | 12781

Syntax

```
pool pool-name {
    address <ip-address> {
        (port port-number | to ip-address);
    }
    description text;
    routing-instance (routing-instance-name | default);
    session-retain;
}
```

Hierarchy Level

```
[edit security nat destination]
```

Description

Define a destination NAT pool to identify the pool uniquely.

Options

- *pool-name*—Name of the pool.
- description—Description of the pool.
- session-retain —To support destination FQDN NAT session continue to use existing IP from DNS resolution up to session lifetime. By default, the `session-retain` is disabled.

  When enabled, the retained sessions remain in the session table, and cleared only upon connection termination from clients, or timeout due to sessions being inactive, or when sessions are explicitly cleared by the user through CLI.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. The `description` option added in Junos OS Release 12.1.

The `session-retain` option added in Junos OS Release 23.1R1.

pool (Security Source NAT)

IN THIS SECTION

- Syntax | 12781
- Hierarchy Level | 12782
- Description | 12782
- Options | 12782
- Required Privilege Level | 12782
- Release Information | 12782

Syntax

```
pool pool-name {
    address ip-address {
         to ip-address;
    }
    burst-pool burst-pool-name;
    address-persistent subscriber ipv6-prefix-length prefix-length;
    address-pooling (paired | no-paired);
    address-shared;
    description text;
    host-address-base ip-address;
    overflow-pool (interface | pool-name);
    pool-utilization-alarm (clear-threshold value | raise-threshold value);
    port (no-translation | port-overloading-factor number | port-overloading-usage-alarm (raise-
threshold | clear-threshold) | preserve-range | enhanced-port-overloading-algorithm | range port-
low (to port-high));
```

```
    routing-instance routing-instance-name;
}
```

Hierarchy Level

```
[edit security nat source]
```

Description

Define a source NAT pool to identify the pool uniquely.

Options

pool-name	Name of the pool.
description	Description of the pool.
enhanced-port-overloading-algorithm	The hash algorithm uses destination IP and port to calculate the port overloading index.
preserve-range	Assign a port within the same range as the incoming port—either 0 through 1023 or 1024 through 65,535. You cannot configure the preserve range option with port no-translation and secure port block allocation.
port-overloading-usage-alarm (*raise-threshold* | *clear-threshold*)	Through the *raise-threshold* option, you can raise threshold for alarm on port usage per overload index value through 50 to 100. Through the *clear-threshold* option, you can clear threshold for alarm on port usage per overload index value through 40 to 100.
burst-pool	Burst mode allows you to use the ports beyond the allocated port blocks.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. The `description` option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release

12.3X48-D10. The `enhanced-port-overloading-algorithm` and `preserve-range` option added in Junos OS Release 22.4R1.

The `burst-pool` and `port-overloading-usage-alarm` options are added in Junos OS Release 23.4R1.

pool (Security Source NAT Rule Set)

IN THIS SECTION

- Syntax | 12783
- Hierarchy Level | 12783
- Description | 12784
- Options | 12784
- Required Privilege Level | 12784
- Release Information | 12784

Syntax

```
pool (pool-name) {
    persistent-nat {
         address-mapping;
        inactivity-timeout seconds;
        max-session-number number;
        permit (any-remote-host | target-host | target-host-port);
     }
}
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name  then source-nat]
```

Description

Specify to use source NAT pool.

Options

pool-name—Name of the source NAT pool.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

pool (Service Interface)

IN THIS SECTION

- Syntax | 12785
- Hierarchy Level | 12785
- Description | 12785
- Options | 12785
- Required Privilege Level | 12785
- Release Information | 12785

Syntax

```
pool pool-name {
    interface interface-name.unit-number;
}
```

Hierarchy Level

```
[edit services service-interface-pools]
```

Description

Configure a service interface pool for VPN aggregation for the BGF feature.

Options

pool-name—Name of the service interface pool.

The remaining options are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Service Interface Pools

pool (Service Interfaces)

IN THIS SECTION

- Syntax | 12786
- Hierarchy Level | 12786
- Description | 12786
- Options | 12786
- Required Privilege Level | 12787
- Release Information | 12787

Syntax

```
pool pool-name {
    interface service-interface-name;
}
```

Hierarchy Level

```
[edit services service-device-pools]
```

Description

Define a pool of service interfaces that can be assigned to an L2TP tunnel group for traffic load-balancing. The interfaces in the pool must be inline service interfaces (si). The service device pool is required for dynamic LNS sessions.

Options

pool-name	Name of the service interface pool.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support for ps service interfaces added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

pool (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 12787
- Hierarchy Level | 12788
- Description | 12789
- Options | 12789
- Required Privilege Level | 12789
- Release Information | 12789

Syntax

```
pool nat-pool-name {
    address address-prefix | address address-prefix to address address-prefix;
    address-pooling {
    }
    ei-mapping-timeout ei-mapping-timeout;
```

```
    host-address-base ip-address;
    mapping-timeout mapping-timeout;
    pool-utilization-alarm {
        clear-threshold value;
        raise-threshold value;
    }
    port {
        automatic (random-allocation | round-robin);
        block-allocation {
            active-block-timeout timeout-interval;
            block-size block-size;
            interim-logging-interval timeout-interval;
            maximum-blocks-per-host maximum-block-number
        }
        deterministic {
            block-size block-size;
            host {
                address address;
            }
            include-boundary-addresses;
        }
        deterministic-nat-configuration-log-interval seconds;
        no-translation;
        preserve-range;
        preserve-parity;
        range {
            port-low to port-high;
            (random-allocation | round-robin);
        }
        port-overloading-factor value;
        enhanced-port-overloading-algorithm;
    }
}
```

Hierarchy Level

```
[edit services nat source]
```

Description

Configure a set of addresses (or prefixes), address ranges, and ports used for Network Address Translation (NAT) of source addresses. Port-overloading factor is configurable between 2 - 32.

Options

nat-pool-name Name of the NAT pool.

If you are configuring twice NAT, do not use the same name that you use for the destination pool.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

pool (System)

IN THIS SECTION

- Syntax | 12790
- Hierarchy Level | 12790
- Description | 12790
- Options | 12790
- Required Privilege Level | 12790
- Release Information | 12791

Syntax

```
pool address/prefix-length {
    address-range {
        low address;
        high address;
    }
    exclude-address {
        address;
    }
}
```

Hierarchy Level

```
[edit system services dhcp]
```

Description

For J Series Services Routers and EX Series switches only. Configure a pool of IP addresses for DHCP clients on a subnet. When a client joins the network, the DHCP server dynamically allocates an IP address from this pool.

Options

`address-range`—Lowest and highest IP addresses in the pool that are available for dynamic address assignment. If no range is specified, the pool will use all available addresses within the subnet specified. (Broadcast addresses, interface addresses, and excluded addresses are not available.)

`exclude-address`—Addresses within the range that are not used for dynamic address assignment. You can exclude one or more addresses within the range.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

address-assignment (Address-Assignment Pools) | 1862

pool (TDF Domain)

IN THIS SECTION

- Syntax | 12791
- Hierarchy Level | 12791
- Description | 12792
- Options | 12792
- Required Privilege Level | 12792
- Release Information | 12792

Syntax

```
pool pool-name;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber subscriber-
address (inet | inet6)]
```

Description

Specify the address pool that contains the source IP addresses for IP-based subscriber packets that undergo TDF processing.

You can specify only one address pool.

Options

pool-name	Name of the address pool.

> **NOTE**: The address pool must have been previously configured at the `[edit access address-assignment]` hierarchy level.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring Address Pools for Source-IP Filtering of IP-Based Subscribers*
- *Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain*
- *Understanding Source IP Filtering with Address Pools in TDF Domains for IP-Based Subscribers*
- *Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain*

pool-utilization-alarm

IN THIS SECTION

- Syntax | 12793
- Hierarchy Level | 12793
- Description | 12793
- Options | 12793
- Required Privilege Level | 12794
- Release Information | 12794

Syntax

```
pool-utilization-alarm (clear-threshold value | raise-threshold value);
```

Hierarchy Level

```
[edit security nat source]
```

Description

Define the global pool utilization alarm thresholds for Network Address Translation (NAT) source IP address pools without Port Address Translation (PAT). When the pool utilization exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered.

Options

clear-threshold *value*—Lower threshold at which an SNMP trap is triggered.

- **Range:** 40 through 100.

raise-threshold *value*—Upper threshold at which an SNMP trap is triggered.

- **Range:** 50 through 100.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

pool-utilization-alarm (Security Source NAT Pool) | 12794

pool-utilization-alarm (Security Source NAT Pool)

IN THIS SECTION

- Syntax | 12794
- Hierarchy Level | 12794
- Description | 12795
- Options | 12795
- Required Privilege Level | 12795
- Release Information | 12795

Syntax

```
pool-utilization-alarm (clear-threshold value | raise-threshold value);
```

Hierarchy Level

```
[edit security nat source pool pool-name]
```

Description

Define utilization alarm thresholds for a specific Network Address Translation (NAT) source pool. When pool utilization exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered. Threshold settings that use this statement take precedence over thresholds that are set using the global `pool-utilization-alarm` statement in the `[security nat source]` hierarchy.

Options

clear-threshold *value*

Lower threshold at which an SNMP trap is triggered.

- **Range:** 40 through 100.

raise-threshold *value*

Upper threshold at which an SNMP trap is triggered.

- **Range:** 50 through 100.

> **NOTE**: If you enter a value for `raise-threshold` but not for `clear-threshold`, `clear-threshold` is automatically set to 80 percent of `raise-threshold`.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

pool-utilization-alarm | 12793

pool-utilization-alarm (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 12796
- Hierarchy Level | 12796
- Description | 12796
- Options | 12797
- Required Privilege Level | 12797
- Release Information | 12797

Syntax

```
pool-utilization-alarm {
    clear-threshold value;
    raise-threshold value;
}
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name]
```

Description

Define the NAT pool utilization level that triggers SNMP traps and the pool utilization level that clears SNMP traps. For pools that use port-block allocation, the utilization is based on the number of ports that are used; for pools that do not use port-block allocation, the utilization is based on the number of addresses that are used.

If you do not configure `pool-utilization-alarm`, traps are not created.

Options

clear-threshold *value*

NAT pool utilization percentage that clears the trap.

- **Range:** 40 through 100
- **Default:** 0 (traps are not created)

raise-threshold *value*

NAT pool utilization percentage that triggers the trap.

- **Range:** 50 through 100
- **Default:** There is not default value. Traps are not raised if you do not configure a value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

pop (protocols mpls)

IN THIS SECTION

- Syntax | 12798
- Hierarchy Level | 12798
- Description | 12798
- Required Privilege Level | 12798
- Release Information | 12798

Syntax

```
pop;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
transit incoming-label],
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Description

Remove the label from the top of the label stack. If there is another label in the stack, that label becomes the label at the top of the label stack. Otherwise, the packet is forwarded as a native protocol packet (typically, as an IP packet).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Intermediate and Egress Routers for Static LSPs

swap (protocols mpls) | **28738**

pop

IN THIS SECTION

- Syntax | 12799
- Hierarchy Level | 12799
- Description | 12799
- Required Privilege Level | 12800
- Release Information | 12800

Syntax

```
pop;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

Description

NOTE: On EX4300 switches, pop is not supported at the [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map] hierarchy level.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2, and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Removing a VLAN Tag*
>
> *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support*

pop-all-labels

IN THIS SECTION

- Syntax | 12800
- Hierarchy Level | 12801
- Description | 12801
- Default | 12801
- Required Privilege Level | 12801
- Release Information | 12801

Syntax

```
pop-all-labels {
    required-depth number;
}
```

Hierarchy Level

```
[edit interfaces interface-name atm-options mpls],
[edit interfaces interface-name fastether-options mpls],
[edit interfaces interface-name gigether-options mpls],
[edit interfaces interface-name sonet-options mpls]
```

Description

For passive monitoring on ATM, SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, removes up to two MPLS labels from incoming IP packets. For passive monitoring on T Series devices, removes up to five MPLS labels from incoming IP packets. For passive monitoring on MX Series routers with MPCs, all labels are popped by default and the `required-depth` statement is ignored.

Except for MX Series routers with MPCs, this statement has no effect on IP packets with more than two MPLS labels, or IP packets with more than five MPLS labels on T Series devices. Packets with MPLS labels cannot be processed by the monitoring PIC; if packets with MPLS labels are forwarded to the monitoring PIC, they are discarded.

The remaining statement is explained separately. See CLI Explorer.

Default

If you omit this statement, the MPLS labels are not removed, and the packet is not processed by the monitoring PIC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Passive Flow Monitoring for MPLS Encapsulated Packets*
>
> Junos OS Network Interfaces Library for Routing Devices

pop-and-forward (Protocols MPLS)

IN THIS SECTION

- Syntax | 12802
- Hierarchy Level | 12802
- Description | 12802
- Required Privilege Level | 12802
- Release Information | 12802

Syntax

```
pop-and-forward;
```

Hierarchy Level

```
[edit logical-systems name protocols mpls label-switched-path  ],
[edit logical-systems name routing-instances name protocols mpls label-switched-path  ],
[edit protocols mpls label-switched-path  ],
[edit routing-instances name protocols mpls label-switched-path  ]
```

Description

Enable LSP as pop-and-forward with auto-delegation signaling enabled by default.

The LSP undergoes a make-before-break from a regular point-to-point LSP to a pop-and-forward LSP.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Pop-and-Forward LSP Configuration

show rsvp pop-and-forward | 23854

pop-and-forward (Protocols RSVP) | 12803

pop-and-forward (Protocols RSVP)

IN THIS SECTION

- Syntax | 12803
- Hierarchy Level | 12803
- Description | 12804
- Options | 12804
- Required Privilege Level | 12804
- Release Information | 12804

Syntax

```
pop-and-forward {
    application-label depth depth;
}
```

Hierarchy Level

```
[edit logical-systems name protocols rsvp],
[edit logical-systems name routing-instances name protocols rsvp],
[edit protocols rsvp],
[edit routing-instances name protocols rsvp]
```

Description

Specify RSVP pop-and-forward LSP tunnel-specific global parameters. The application label depth (AppLD) value must be configured uniformly across the RSVP-TE network.

Options

application-label depth *depth*

Specify the maximum number of service labels.

- **Range:** 0 through 3
- **Default:** 1

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Pop-and-Forward LSP Configuration

show rsvp pop-and-forward | **23854**

pop-and-forward (Protocols MPLS) | **12802**

pop (Dynamic VLANs)

IN THIS SECTION

- Syntax | **12805**
- Hierarchy Level | **12805**
- Description | **12805**
- Required Privilege Level | **12805**

Release Information | 12805

Syntax

```
pop;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map]
```

Description

For dynamic VLAN interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Removing a VLAN Tag

Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution

pop-pop

IN THIS SECTION

- Syntax | 12806
- Hierarchy Level | 12806
- Description | 12806
- Required Privilege Level | 12806
- Release Information | 12807

Syntax

```
pop-pop;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

Description

For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP, and for 10-Gigabit Ethernet SFP interfaces on EX Series switches, specify the VLAN rewrite operation to remove both the outer and inner VLAN tags of the frame.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Removing the Outer and Inner VLAN Tags

pop-swap

IN THIS SECTION

- Syntax | 12807
- Hierarchy Level | 12807
- Description | 12808
- Required Privilege Level | 12808
- Release Information | 12808

Syntax

```
pop-swap;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

Description

Specify the VLAN rewrite operation to remove the outer VLAN tag of the frame, and replace the inner VLAN tag of the frame with a user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.

You can use this statement on Gigabit Ethernet IQ, IQ2, IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag

pop3-profile (Security Content Security Policy Antivirus)

IN THIS SECTION

- Syntax | 12809
- Hierarchy Level | 12809
- Description | 12809
- Required Privilege Level | 12809
- Release Information | 12809

Syntax

```
pop3-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name anti-virus]
```

Description

Configure a Content Security policy for the antivirus POP3 protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

pop3-profile (Security Content Security Policy Content Filtering)

IN THIS SECTION

- Syntax | 12810
- Hierarchy Level | 12810
- Description | 12810
- Required Privilege Level | 12810

Release Information | 12810

Syntax

```
pop3-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name content-filtering]
```

Description

Configure a Content Security policy for the content filtering POP3 protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

port

IN THIS SECTION

- Syntax | 12811
- Hierarchy Level | 12811
- Description | 12811
- Default | 12811
- Options | 12811
- Required Privilege Level | 12812
- Release Information | 12812

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit snmp v3 target-address target-address-name]
```

Description

Configure a UDP port number for an SNMP target.

Default

If you omit this statement, the default port is 162.

Options

port-number—Port number for the SNMP target.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

port (Access LDAP)

IN THIS SECTION

- Syntax | 12812
- Hierarchy Level | 12812
- Description | 12813
- Options | 12813
- Required Privilege Level | 12813
- Release Information | 12813

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit access ldap-server server-address],
[edit access profile profile-name ldap-server server-address]
```

Description

Configure the port number on which to contact the LDAP server.

Options

port-number—Port number on which to contact the LDAP server.

- **Default:** 389

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

LDAP Functionality in Integrated User Firewall

port (cflowd)

IN THIS SECTION

- Syntax | 12814
- Hierarchy Level | 12814
- Description | 12814
- Options | 12814
- Required Privilege Level | 12814
- Release Information | 12814

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit forwarding-options accounting group-name output cflowd hostname],
[edit forwarding-options monitoring group-name family inet output flow-server hostname],
[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server hostname]
```

Description

Specify the UDP port number on the cflowd host system.

Options

port-number—Any valid UDP port number on the host system.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Directing Traffic Sampling Output to a Server Running the cflowd Application*

port (Chassis)

IN THIS SECTION

- Syntax | 12815
- Hierarchy Level | 12815
- Description | 12816
- Options | 12816
- Required Privilege Level | 12817
- Release Information | 12817

Syntax

```
port name {
    channel-speed(10g | 25g | 50g | disable-auto-speed-detection);
    forwarding-mode {
        sa-multicast;
    }
    framing(e1 | e3 | sdh | sonet | t1 | t3);
    no-mcast-replication;
    no-temperature-monitoring;
    number-of-sub-ports number-of-sub-ports;
    short-reach-mode(disable | enable);
    speed(100g | 10g | 1G | 200g | 25g | 400g | 40g | 50g | oc12-stm4 | oc3-stm1 | oc48-stm16);
    traffic-manager {
        mode ingress-and-egress;
    }
    unused;
}
```

Hierarchy Level

```
[edit chassis fpc name pic],
```

```
[edit chassis lcc name fpc name pic],
[edit chassis member name fpc name pic]
```

Description

Configure properties for a specific port number.

Options

name

Port number

- **Range:** Depends on the device.

channel-speed

Port channel speed

- 10g—Set the port speed to 10G. This will restart PFE on some platforms.
- 25g—Set the port speed to 25G.
- 50g—Set the port speed to 50G.
- disable-auto-speed-detection—Disables automatic speed detection

forwarding-mode

PIC packet distribution mode - Brooklyn interop mode

- sa-multicast—SA multicast mode - interop with 100G Brooklyn PIC

framing

Framing mode

- e1—E1 mode
- e3—E3 mode
- sdh—SDH mode
- sonet—SONET mode
- t1—T1 mode
- t3—T3 mode

no-mcast-replication

No multicast replication

no-temperature-monitoring

Disable EM Policy support. Please check Junos OS Release notes and Junos OS Evolved Release Notes to verify if your device supports EM Policy.

number-of-sub-ports

Number of subports per physical port

- **Range:** 1 through 8

short-reach-mode

Short reach mode (For ports 0...47)

- disable—Disable
- enable—Enable

speed

Port speed

- 100g—Sets the interface mode to 100Gbps
- 10g—Sets the interface mode to 10Gbps
- 1G—1GE-Gigabit Ethernet
- 200g—Sets the interface mode to 200 Gbps
- 25g—Sets the interface mode to 25Gbps
- 400g—Sets the interface mode to 400Gbps
- 40g—Sets the interface mode to 40Gbps
- 50g—Sets the interface mode to 50Gbps
- oc12-stm4—OC12 or STM4
- oc3-stm1—OC3 or STM1
- oc48-stm16—OC48 or STM16

unused

Keep the port unused.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

port (Chassis Synchronization)

IN THIS SECTION

- Syntax (MX Series) | 12818
- Hierarchy Level | 12818
- Description | 12819
- Options | 12819
- Required Privilege Level | 12819
- Release Information | 12819

Syntax (MX Series)

```
port {
    auxiliary {
        client {
            time-of-day-format {
                ascii string;
            }
        }
    }
}
```

Hierarchy Level

```
[edit chassis synchronization]
[edit chassis synchronization source interfaces (external-a | external-b | interface interface-name)]
[edit chassis synchronization source interfaces (external | interface interface-name)]
[edit chassis synchronization source interfaces (external-0/0 | external-1/0 | interface interface-name)]
```

Description

Configure the clock synchronization ports.

Options

time-of-day-format Specify the time of day (TOD) message format.

- ascii—Configure the TOD message format in ASCII characters on the auxiliary port that receives the external clock signals.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Integrated Global Navigation Satellite System (GNSS) on ACX500 Series Routers

synchronization (MX Series) | **28795**

port (Diameter Peer)

IN THIS SECTION

- Syntax | **12820**
- Hierarchy Level | **12820**
- Description | **12820**
- Options | **12820**
- Required Privilege Level | **12820**

Release Information | 12820

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit diameter peer peer-name connect-actively]
```

Description

Specify the destination TCP port used by the active connection to peer.

Options

port-number—Number of the TCP port.

- **Default:** 3868

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Peers

port-error-disable

IN THIS SECTION

- Syntax | 12821
- Hierarchy Level | 12821
- Description | 12821
- Default | 12822
- Required Privilege Level | 12822
- Release Information | 12822

Syntax

```
port-error-disable {
    disable-timeout timeout ;
}
```

Hierarchy Level

```
[edit ethernet-switching-options]
```

Description

Disable rather than block an interface when enforcing MAC limiting, MAC move limiting, and rate-limiting configuration options for shutting down the interface, and allow the interface to recover automatically from the error condition after a specified period of time:

- If you have enabled MAC limiting with the `shutdown` option and you enable `port-error-disable`, the switch disables (rather than shuts down) the interface when the MAC address limit is reached.
- If you have enabled MAC move limiting with the `shutdown` option and you enable `port-error-disable`, the switch disables (rather than shuts down) the interface when the maximum number of moves to a new interface is reached.

- If you have enabled storm control with the `action-shutdown` option and you enable `port-error-disable`, the switch disables (rather than shuts down) the interface when applicable traffic exceeds the specified levels. Depending upon the configuration, applicable traffic could include broadcast, unknown unicast, and multicast traffic.

> **NOTE**: The `port-error-disable` configuration does not apply to pre-existing error conditions. It impacts only error conditions that are detected after `port-error-disable` has been enabled and committed. To clear a pre-existing error condition and restore the interface to service, use the operational command that appears in your CLI:
>
> - `clear ethernet-switching port-error`

The remaining statement is explained separately. See CLI Explorer.

Default

Not enabled.

Required Privilege Level

system—To view this statement in the configuration.
system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

action-shutdown | 1787

Configuring MAC Move Limiting (non-ELS)

port (Flow Monitoring)

IN THIS SECTION

- Syntax | 12823
- Hierarchy Level | 12823
- Description | 12823
- Options | 12823
- Required Privilege Level | 12824
- Release Information | 12824

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit forwarding-options accounting name output cflowd hostname],
[edit forwarding-options monitoring name family inet output cflowd hostname],
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-
server hostname],
[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server hostname]
```

Description

Specify the User Datagram Protocol (UDP) port number on the cflowd host system or flow server.

Options

port-number—Any valid UDP port number on the host system.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Flow Aggregation

port-forwarding

IN THIS SECTION

- Syntax | 12824
- Hierarchy Level | 12825
- Description | 12825
- Options | 12825
- Required Privilege Level | 12825
- Release Information | 12825

Syntax

```
port-forwarding map-name {
    destined-port;
    translated-port;
}
```

Hierarchy Level

```
[edit services nat]
```

Description

Specify the mapping for port forwarding.

The `port-forwarding` statement is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, the `port-forwarding` statement is also supported on the MS-MPC and MS-MIC.

Options

map-name Identifier for the port forwarding map.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Port Forwarding for Static Destination Address Translation

Configuring Port Forwarding Without Destination Address Translation

port-forwarding (Destination NAT Next Gen Services)

IN THIS SECTION

- Syntax | 12826
- Hierarchy Level | 12826
- Description | 12826
- Options | 12826
- Required Privilege Level | 12827

Syntax

```
port-forwarding map-name {
    destined-port port-id translated-port port-id;
}
```

Hierarchy Level

```
[edit services nat destination]
```

Description

Configure a port forwarding map, which translates the original destination port of a packet to a different port. This translation is a static, one-to-one mapping.

Port forwarding allows a packet to reach a host within a masqueraded, typically private, network, based on the port number on which the packet was received from the originating host. An example of this type of destination is the host of a public HTTP server within a private network.

Options

map-name	Name of the port forwarding map.
destined-port *port-id*	Original destination port number.

translated-port *port-id* Port number to which the original port is mapped.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

port-forwarding-mappings

IN THIS SECTION

- Syntax | 12827
- Hierarchy Level | 12827
- Description | 12827
- Options | 12828
- Required Privilege Level | 12828
- Release Information | 12828

Syntax

```
port-forwarding-mappings map-name;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then]
```

Description

Specify the name of the mapping for port forwarding in a Network Address Translation rule.

The `port-forwarding-mappings` statement is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Port forwarding on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS supports only the `dnat-44` and `twice-napt-44` translation types in NAT rules, and supports only IPv4 networks. Starting in Junos OS Release 17.4R1, the `port-forwarding-mappings` statement is also supported on the MS-MPC and MS-MIC.

Options

map-name	Identifier for the port forwarding mapping.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Port Forwarding for Static Destination Address Translation

Configuring Port Forwarding Without Destination Address Translation

port-forwarding-mappings (Destination NAT Rule Next Gen Services)

IN THIS SECTION

- Syntax | 12829
- Hierarchy Level | 12829
- Description | 12829
- Required Privilege Level | 12829

Syntax

```
port-forwarding-mappings map-name;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set-name rule rule-name then]
```

Description

Specify the name of the port-forwarding map that the NAT rule uses to translate the original destination port of a packet to a different port.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

port-group-alias (satellite-policies)

IN THIS SECTION

- Syntax | 12830
- Hierarchy Level | 12830
- Description | 12830
- Default | 12830
- Options | 12830
- Required Privilege Level | 12831
- Release Information | 12831

Syntax

```
port-group-alias port-group-alias-name {
    pic pic-number {
        port [port-number | port-number-range | all];
    }
}
```

Hierarchy Level

```
[edit policy-options satellite-policies]
```

Description

Configure a port group alias for satellite policies in a Junos Fusion.

A port group alias is used to define the candidate uplink ports on satellite devices that use the satellite policy.

> **CAUTION**: The physically connected uplink ports on a satellite device must be defined as candidate uplink ports in the Junos Fusion configuration. If the uplink ports on a satellite device are not configured as candidate uplink ports, uplink failure detection cannot be enabled on the device, and a system log message is generated.

A port group alias is associated with a satellite policy using the `set uplink-port-group` *`uplink-port-group-name`* statement in the [`edit policy-options satellite-policies candidate-uplink-policy` *`policy-name`*] hierarchy.

Default

Each satellite device model has a set of default uplink ports (see *Overview of Uplink Failure Detection on a Junos Fusion*). You only need to assign an uplink port group to a satellite device if you do not want to use the default uplink ports.

Options

port-group-alias-name	The user-defined name of the port group alias.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

port-id

IN THIS SECTION

- Syntax | 12831
- Hierarchy Level | 12832
- Description | 12832
- Default | 12832
- Options | 12832
- Required Privilege Level | 12832
- Release Information | 12832

Syntax

```
port-id port-id-number;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel secure-channel-name id]
```

Description

Specify a port ID in a secure channel when enabling MACsec using static secure association key (SAK) security mode. The port IDs must match on a sending and receiving secure channel on each side of a link to enable MACsec.

Once the port numbers match, MACsec is enabled for all traffic on the connection.

You only use this configuration option when you are configuring MACsec using static SAK security mode. This option does not need to be specified when you are enabling MACsec using static connectivity association key (CAK) security mode.

Default

No port ID is specified.

Options

port-id-number	The port ID number.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

port-id (MACsec for MX Series)

IN THIS SECTION

- Syntax | 12833
- Hierarchy Level | 12833
- Description | 12833
- Default | 12834
- Options | 12834
- Required Privilege Level | 12834
- Release Information | 12834

Syntax

```
port-id port-id-number;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel
secure-channel-name id]
```

Description

Specify a port ID in a secure channel when enabling MACsec using static secure association key (SAK) security mode. The port IDs must match on a sending and receiving secure channel on each side of a link to enable MACsec.

Once the port numbers match, MACsec is enabled for all traffic on the connection.

You only use this configuration option when you are configuring MACsec using static SAK security mode. This option does not need to be specified when you are enabling MACsec using static connectivity association key (CAK) security mode.

Default

No port ID is specified.

Options

port-id-number	The port ID number.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

port (LRF Profile)

IN THIS SECTION

- Syntax | 12835
- Hierarchy Level | 12835
- Description | 12835
- Options | 12835
- Required Privilege Level | 12835
- Release Information | 12835

Syntax

```
port collector-port-number;
```

Hierarchy Level

```
[edit services lrf profile profile-name collector collector-name destination]
```

Description

Specify the destination port of the collector.

Options

collector-port-number Port number for the destination address of the collector.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

port-mirroring

IN THIS SECTION

- Syntax: MX Series and PTX Series Routers, M120 and M320 | 12836
- Syntax: QFX Series (but not QFX10000) and EX4600 Switches and NFX Series Devices | 12837
- Syntax: OCX1100 | 12838
- Syntax: QFX10000 Switches | 12839
- Hierarchy Level | 12840
- Description | 12840
- Required Privilege Level | 12840
- Release Information | 12840

Syntax: MX Series and PTX Series Routers, M120 and M320

```
port-mirroring {
    input {
        maximum-packet-length bytes;
        rate number;
        run-length number;
    }
    family (any | ccc | inet | inet6 | vpls) {
        output {
            interface interface-name {
                next-hop address;
            }
            next-hop-group group-name {
                group-type inet6;
                interface interface-name {
                    next-hop ipv6-address;
                }
                next-hop-subgroup group-name {
                    interface interface-name {
                        next-hop ipv6-address;
                    }
                }
```

```
            }
            no-filter-check;
        }
    }
    instance {
        instance-name {
            input {
                maximum-packet-length bytes;
                rate number;
                run-length number;
            }
            family (any | ccc | inet | inet6 | vpls) {
                output {
                    interface interface-name {
                        next-hop address;
                    }
                    no-filter-check;
                    server-profile server-profile-name;
                }
            }
        }
    }
    mirror-once;
    packet-timestamp;
    traceoptions {
        file filename  <files number> <size bytes> <world-readable | no-world-readable>;
    }
}
```

Syntax: QFX Series (but not QFX10000) and EX4600 Switches and NFX Series Devices

```
port-mirroring {
    family {
        ethernet-switching {
            output {
                interface interface-name;
                no-filter-check;
                vlan vlan-name {
                    no-tag;
                }
```

```
            }
        }
        inet {
            output {
                ip-address address;
                routing-instance instance-name {
                    ip-address address;
                }
            }
        }
    }
    instance instance-name {
        family {
            ethernet-switching {
                output {
                    interface interface-name;
                    no-filter-check;
                    vlan vlan-name {
                        no-tag;
                    }
                }
            }
            inet {
                output {
                    ip-address address;
                    routing-instance instance-name {
                        ip-address address;
                    }
                }
            }
        }
    }
}
```

Syntax: OCX1100

```
port-mirroring {
    family {
        inet {
            output {
                ip-address address;
```

```
                routing-instance instance-name {
                    ip-address address;
                }
            }
        }
    }
    instance instance-name {
        family {
            inet {
                output {
                    ip-address address;
                    routing-instance instance-name {
                        ip-address address;
                    }
                }
            }
        }
    }
}
```

Syntax: QFX10000 Switches

```
port-mirroring {
     instance instance-name {
             input {
                 rate number;
                 run-length number;
            }
             output {
                dscp value;
                forwarding-class class-name;
                ip-destination-address address;
                ip-source-address address;
                policer policer-name;
            }
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Create a port-mirroring configuration. Specify the address family, rate, run length, interface, and next-hop address for sending copies of packets to an analyzer.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`family vpls` statement introduced in Junos OS Release 9.3.

`instance` *`instance-name`* statement introduced in Junos OS Release 9.3 .

`mirror-once` statement introduced in Junos OS Release 9.3 .

`family ccc` statement introduced in Junos OS Release 9.6.

`family inet6` and `next-hop-group` statements introduced in Junos OS Release 14.2 .

RELATED DOCUMENTATION

Configuring Port Mirroring

Configuring Port Mirroring

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

Understanding Port Mirroring and Analyzers

Understanding Port Mirroring

Examples: Configuring Port Mirroring for Local Analysis

Example: Mirroring Employee Web Traffic with a Firewall Filter

port-mode

IN THIS SECTION

- Syntax | 12841
- Hierarchy Level | 12841
- Description | 12841
- Default | 12842
- Options | 12842
- Required Privilege Level | 12842
- Release Information | 12842

Syntax

```
port-mode (access | tagged-access | trunk);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family ethernet-switching]
```

Description

NOTE: This statement does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see "interface-mode" on page 8920. For ELS details, see Using the Enhanced Layer 2 Software CLI.

Configure whether an interface on the switch operates in access, tagged access, or trunk mode.

Default

All switch interfaces are in access mode.

Options

access—Have the interface operate in access mode. In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to network devices, such as PCs, printers, IP telephones, and IP cameras.

tagged-access—Have the interface operate in tagged-access mode. In this mode, the interface can be in multiple VLANs. Tagged access interfaces typically connect to network devices, such as PCs, printers, IP telephones, and IP cameras.

trunk—Have the interface operate in trunk mode. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN.

> **NOTE**: The number of VLANs supported per switch varies for each model. Use the configuration-mode command `set vlans id vlan-id ?` to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum. To determine the maximum number of VLAN members allowed on a switch, multiply the VLAN maximum for the switch times 8 (vmember limit = vlan max * 8).
>
> If a switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (**eswd**) due to memory allocation failure.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Example: Connecting an EX Series Access Switch to a Distribution Switch
- Configuring Gigabit Ethernet Interfaces (CLI Procedure)
- Configuring VLANs for EX Series Switches
- Junos OS Ethernet Interfaces Configuration Guide

port-mode (Fibre Channel Interfaces)

IN THIS SECTION

- Syntax | 12843
- Hierarchy Level | 12843
- Description | 12843
- Options | 12844
- Required Privilege Level | 12844
- Release Information | 12844

Syntax

```
port-mode (f-port | np-port);
```

Hierarchy Level

```
[edit interfaces vlan unit unit family fibre-channel],
[edit interfaces interface-name unit logical-unit-number family fibre-channel]
```

Description

Configure the FCoE VLAN interface port mode to F_Port to connect the switch to FCoE initiators, or configure the native FC interface port mode to proxy N_Port (NP_Port) to connect the switch to an FC switch fabric port (F_Port).

Options

f-port—Configure an FCoE VLAN interface to connect to FCoE initiator Virtual N_Ports (VN_Ports).

np-port—Configure a native FC port to connect to an FC switch F_Port.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel interfaces | **19980**

Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric

Configuring an FCoE VLAN Interface on an FCoE-FC Gateway

Converting an Ethernet Interface To a Fibre Channel Interface

Understanding Interfaces on an FCoE-FC Gateway

port (NETCONF)

IN THIS SECTION

- Syntax | **12845**
- Hierarchy Level | **12845**
- Description | **12845**
- Options | **12845**
- Required Privilege Level | **12845**
- Release Information | **12846**

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit system services netconf ssh]
```

Description

Configure the TCP port used for NETCONF-over-SSH connections.

NOTE:

- The configured port accepts only NETCONF-over-SSH connections. Regular SSH session requests for this port are rejected.
- The default SSH port (22) continues to accept NETCONF sessions even with a configured NETCONF server port. To disable the SSH port from accepting NETCONF sessions, you can specify this in the login event script.
- We do not recommend configuring the default ports for FTP (21) and Telnet (23) services for configuring NETCONF-over-SSH connections.

Options

`port` *`port-number`*—Port number on which to enable incoming NETCONF connections over SSH.

- **Default:** 830 (as specified in RFC 4742, *Using the NETCONF Configuration Protocol over Secure Shell (SSH)*)
- **Range:** 1 through 65535

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring NETCONF-Over-SSH Connections on a Specified TCP Port

port (Origin Validation for BGP)

IN THIS SECTION

- Syntax | 12846
- Hierarchy Level | 12846
- Release Information | 12847
- Description | 12847
- Options | 12847
- Required Privilege Level | 12847

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation group group-name session server-ip-address],
[edit logical-systems logical-system-name routing-options validation group group-name session
server-ip-address],
[edit routing-instances instance-name routing-options validation group group-name session server-
ip-address],
[edit routing-options validation group group-name session server-ip-address]
```

Release Information

Statement introduced in Junos OS Release 12.2.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure an alternative TCP port number to be used for the outgoing connection with the cache server. The well-known resource public key infrastructure (RPKI) port is TCP port 2222. For a given deployment, an RPKI cache server might listen on some other TCP port number. If so, configure the alternative port number with this statement.

Options

port-number TCP port number to be used for the outgoing connection to the cache server.

- Default: 8282

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

BGP Origin Validation

port-overloading-factor (Security Source NAT Interface)

IN THIS SECTION

- Syntax | **12848**
- Hierarchy Level | **12848**
- Description | **12848**

- Options | 12848
- Required Privilege Level | 12849
- Release Information | 12849

Syntax

```
port-overloading-factor number;
```

Hierarchy Level

```
[edit security nat source interface]
```

Description

Configure the port overloading capacity for the source NAT interface. If `port-overloading-factor` is set to *x*(1 up to the maximum port capacity), then x times the maximum port capacity is allocated for interface-based NAT.

> **NOTE**: There is also a `port-overloading` option, but it is not supported for logical systems, and should not be used in conjunction with the `port-overloading-factor` option because the statements can overwrite each other. For example, if `port-overloading` has been set to **off** to disable interface port overloading, and subsequently `port-overloading-factor` is configured with any value greater than 1, the `port-overloading-factor` setting will override the `port-overloading` setting. (Configuring **port-overloading-factor 1** is equivalent to configuring **port-overloading off**.)

Options

number A number ranging from 1 through the maximum port capacity.

For example, if `port-overloading-factor` is set to 2, and it is multiplied by a maximum port capacity of 63,486, the port overloading threshold is 126,972. If the configured `port-overloading-factor` setting exceeds the maximum port capacity of the interface, an error message is generated during the configuration commit.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

port-overloading (Security Source NAT Interface) | 12850

port-overloading-factor (Security Source NAT Pool)

IN THIS SECTION

- Syntax | 12849
- Hierarchy Level | 12849
- Description | 12850
- Options | 12850
- Required Privilege Level | 12850
- Release Information | 12850

Syntax

port-overloading-factor

Hierarchy Level

[edit security nat source pool *source-pool-name* port]

Description

Configures the port overloading capacity in source NAT. If the port-overloading-factor is set to *x*, each translated IP address will have *x* number of ports available.

> **NOTE**: The port-overloading-factor statement cannot be configured with port no-translation (source NAT pool without PAT).

Options

- **Range:** 2 through 32

For example, If you set `port-overloading-factor` to 2 for a source pool with two IP addresses, each with the single port range of 1024 through 2047, the ports are multiplied by 2, increasing the port capacity for each from 1024 to 2048. If the configured port-overloading-factor setting exceeds the maximum port capacity of the pool, an error message is generated during the configuration commit.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2

port-overloading (Security Source NAT Interface)

IN THIS SECTION

- Syntax | 12851
- Hierarchy Level | 12851
- Description | 12851
- Options | 12851
- Required Privilege Level | 12851

Release Information | 12851

Syntax

```
port-overloading off
```

Hierarchy Level

```
[edit security nat source interface]
```

Description

Enable interface NAT with or without port overloading.

Options

off Specify off to disable interface port overloading.

> **NOTE**: The `port-overloading` option should not be used in conjunction with the `port-overloading-factor` option because they can override each other. For example, if `port-overloading` has been set to `off` to disable interface port overloading, and subsequently the `port-overloading-factor` is configured with any value greater than 1, the `port-overloading-factor` setting will override the `port-overloading` setting. (Configuring `port-overloading-factor 1` is equivalent to configuring `port-overloading off`.)

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

port-overloading-factor (Security Source NAT Interface) | 12847

port (OVSDB)

IN THIS SECTION

- Syntax | 12852
- Hierarchy Level | 12852
- Description | 12852
- Options | 12852
- Required Privilege Level | 12853
- Release Information | 12853

Syntax

```
port number;
```

Hierarchy Level

```
[edit protocols ovsdb controller protocol]
```

Description

Specify the software-defined networking (SDN) controller port to which a Juniper Networks device that supports the *Open vSwitch Database (OVSDB)* management protocol connects.

Options

number Number of SDN controller port.

- **Range:** 1024 through 65,535

- **Default:** 6632

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

> *Understanding How to Set Up OVSDB Connections on a Juniper Networks Device*

port (Packet Forwarding)

IN THIS SECTION

- Syntax | 12853
- Hierarchy Level | 12854
- Description | 12854
- Options | 12855
- Required Privilege Level | 12855
- Release Information | 12855

Syntax

```
port port-number {
    description text-description;
    interface interface-name {
        broadcast;
        description text-description;
```

```
        no-listen;
        server address <logical-system logical-system-name> <routing-instance routing-instance-
name>;
    }
    server address <logical-system logical-system-name> <routing-instance routing-instance-name>;
}
```

Hierarchy Level

```
[edit forwarding-options helpers]
```

Description

Configure a port helper on the router or switch, which listens for LAN broadcast traffic on a custom UDP port number and forwards traffic to particular destination servers as unicast traffic.

To set up a port helper, configure the UDP port number and optionally an interface on which to listen for broadcast traffic, and the destination server address to receive that traffic, as shown in either of the following sample configurations:

```
user@ host# show forwarding-options
helpers {
    port 1200 {
        server 10.20.30.40;
        }
    }
```

```
user@ host# show forwarding-options
helpers {
    port 3000 {
        interface {
            fe-0/0/1.0 {
                server 192.0.2.2;
            }
        }
    }
    port 3001 {
        interface {
            fe-0/0/0.0 {
```

```
                server 192.0.2.2;
            }
        }
    }
}
```

Starting in Junos OS Release 17.2R1, you can configure forwarding traffic to multiple destination servers for a given port number by specifying multiple port configuration statements with the same port number and different server addresses.

You cannot configure port helpers for standard ports used by services such as BOOTP, DNS and TFTP; instead, use the `helpers` configuration statements specifically for forwarding packets for those protocols.

The remaining statements are explained separately. See CLI Explorer.

Options

port-number UDP port number for listening.

- **Range:** 1 through 65535

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for multiple server instances for a given port introduced in Junos OS Release 17.2 for MX Series routers.

Support for multiple servers on a given port introduced in Junos OS Release 17.3R1 for EX9200 switches.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding

Configuring Port-based LAN Broadcast Packet Forwarding

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

port-priority

IN THIS SECTION

- Syntax | 12856
- Hierarchy Level | 12856
- Description | 12856
- Options | 12856
- Required Privilege Level | 12857
- Release Information | 12857

Syntax

```
port-priority priority;
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options 802.3ad lacp]
```

Description

Define LACP port priority at the interface level.

Options

priority—Priority for being elected to be the active port and both collect and distribute traffic. A smaller value indicates a higher priority for being elected.

- **Range:** 0 through 65535
- **Default:** 127

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

Configuring Aggregated Ethernet LACP (CLI Procedure)

port (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 12857
- Hierarchy Level | 12858
- Description | 12858
- Options | 12858
- Required Privilege Level | 12858
- Release Information | 12858

Syntax

```
port port;
```

Hierarchy Level

```
[edit protocols openflow switch switch-name controller protocol protocol]
```

Description

Specify the OpenFlow controller port to which the OpenFlow virtual switch connects.

Options

port Numeric value specifying the OpenFlow controller port to which the device should connect.

- **Range:** 1024 through 65,535
- **Default:** 6633

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Support for OpenFlow on Devices Running Junos OS

address (Protocols OpenFlow) | **1912**

controller (Protocols OpenFlow) | **4654**

protocol (Protocols OpenFlow) | **13455**

port (RADIUS Server)

IN THIS SECTION

- Syntax | 12859
- Hierarchy Level | 12859
- Description | 12859
- Options | 12859
- Required Privilege Level | 12859
- Release Information | 12860

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure the port number to which the RADIUS requests are sent.

Options

port-number Port number to which the RADIUS requests are sent.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

port-randomization

IN THIS SECTION

- Syntax | 12860
- Hierarchy Level | 12860
- Description | 12861
- Options | 12861
- Required Privilege Level | 12861
- Release Information | 12861

Syntax

```
port-randomization disable;
```

Hierarchy Level

```
[edit security nat source]
```

Description

Disable random port allocation for pool-based and interface source NAT.

Options

`disable`—Disables random port allocation for pool-based and interface source NAT. For pool-based source NAT and interface NAT, port numbers are allocated randomly by default. Although randomized port number allocation can provide protection from security threats such as DNS poison attacks, it can also affect performance and memory usage for pool-based source NAT.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

port-range

IN THIS SECTION

- Syntax | 12862
- Hierarchy Level | 12862
- Description | 12862
- Options | 12862
- Required Privilege Level | 12862
- Release Information | 12863

Syntax

```
port-range port-range-low port-range-high;
```

Hierarchy Level

```
[edit chassis (QFX Series) fpc fpc-id pic pic-id fibre-channel]
```

Description

Configure a contiguous block of ports as FC ports. You can configure the FC-capable ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5, and ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47 to create blocks of native FC interfaces. You cannot individually configure a single port as a native FC interface. Within these port blocks, you cannot mix FC interfaces with Ethernet interfaces. All of the ports in a block must be either native FC interfaces or Ethernet interfaces.

You can configure:

- Six native FC interfaces by configuring either ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5, or ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47.
- Twelve native FC interfaces by configuring ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5 and ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47.
- No native FC interfaces by leaving ports xe-0/0/0 through xe-0/0/5 and ports xe-0/0/42 through xe-0/0/47 in their default state as Ethernet interfaces.

Options

port-range-low—Lowest-numbered port in the block of native FC interfaces, either **0** or **42**.

port-range-high—Highest-numbered port in the block of native FC interfaces. The value is **5** if the ***port-range-low*** value is **0**. The value is **47** if the ***port-range-low*** value is **42**.

> **NOTE**: Only a complete block of ports, xe-0/0/0 through xe-0/0/5, xe-0/0/42 through xe0/0/47, or both, can be configured as FC ports.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel interfaces | 19980

Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric

Configuring a Physical Fibre Channel Interface

Understanding Interfaces on an FCoE-FC Gateway

port-range (Address Mapping)

IN THIS SECTION

- Syntax | 12863
- Hierarchy Level | 12864
- Description | 12864
- Options | 12864
- Required Privilege Level | 12864
- Release Information | 12864

Syntax

```
port-range {
    tcp [port];
    udp [port];
}
```

Hierarchy Level

[]

```
edit services application-identification application application-name address-mapping address-name filter
```

Description

Specify a port to match a TCP or UDP destination port for Layer 3 and Layer 4 address-based custom applications.

.

Layer 3 and Layer 4 address-based custom applications, you can match the IP address and port range to destination IP address and port. When both IP address and port are configured, both should match destination tuples (IP address and port range) of the packet. The format for numeric port ranges is in the format *minimum-value–maximum-value*.

Options

- **tcp [*port*]**—Define the TCP port range for the application.
- **udp [*port*]**—Define the UDP port range for the application.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

Understanding Junos OS Application Identification Custom Application Signatures

port-range (Application Identification)

IN THIS SECTION

- Syntax | 12865
- Hierarchy Level | 12865
- Description | 12865
- Options | 12865
- Required Privilege Level | 12866
- Release Information | 12866

Syntax

```
port-range {
    tcp [port-range];
    udp [port-range];
}
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type
signature l4-l7-signature-name]
```

Description

Define TCP or UDP port number range.

Options

port-range Numeric port ranges. The format for numeric port ranges is in the format *minimum-value–maximum-value*.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

port (REST API)

IN THIS SECTION

- Syntax | 12866
- Hierarchy Level | 12867
- Description | 12867
- Options | 12867
- Required Privilege Level | 12867
- Release Information | 12867

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit system services rest http],
[edit system services rest https]
```

Description

Specify the port number.

Options

port-number Port number on which to accept HTTP or HTTPS connections.

- **Range:** 1024 through 65535
- **Default:** 3000 for HTTP, 3443 for HTTPS

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

port-round-robin

IN THIS SECTION

- Syntax | 12868
- Hierarchy Level | 12868
- Description | 12868
- Options | 12868

Required Privilege Level | 12868
Release Information | 12868

Syntax

```
port-round-robin disable;
```

Hierarchy Level

```
[edit security nat source]
```

Description

Disable round-robin port allocation for pool-based and interface source NAT on SRX5400, SRX5600, and SRX5800 devices.

Options

`disable`—Disables round-robin port allocation for pool-based and interface source NAT.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D30.

port-round-robin (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 12869
- Hierarchy Level | 12869
- Description | 12869
- Required Privilege Level | 12869
- Release Information | 12870

Syntax

```
port-round-robin {
    disable;
}
```

Hierarchy Level

```
[edit services nat source]
```

Description

Disable round-robin port allocation for any NAT pools that do not specify an `automatic (random-allocation | round-robin)` setting at the `[edit services nat source pool` *`nat-pool-name`* `port]` hierarchy level. The `automatic (random-allocation | round-robin)` setting for a pool overrides the `port-round-robin disable` setting.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

port (RPM)

IN THIS SECTION

- Syntax | 12870
- Junos OS Hierarchy Level | 12870
- Junos OS Evolved Hierarchy Level | 12870
- Description | 12871
- Options | 12871
- Required Privilege Level | 12871
- Release Information | 12871

Syntax

```
port number;
```

Junos OS Hierarchy Level

```
[edit services rpm probe-server (tcp | udp)]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name udp]
```

Description

Specify the port number for the RPM probe server.

Options

number—Port number for the probe server. The value can be `7` or `49,160` through `65,535`.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

Configuring RPM Receiver Servers

port (satellite-policies port-group-alias)

IN THIS SECTION

- Syntax | 12872
- Hierarchy Level | 12872
- Description | 12872
- Options | 12872
- Required Privilege Level | 12872
- Release Information | 12872

Syntax

```
port [port-number | port-number-range | all];
```

Hierarchy Level

```
[edit policy-options satellite-policies port-group-alias port-group-alias-name pic pic-number]
```

Description

Specify the port or ports to apply a port group alias for satellite policies in a Junos Fusion.

You must also specify the PIC when you use this statement.

Options

port-number	The port number on the PIC on the satellite device.
port-number-range	A range of port numbers on the PIC.
all	All ports on the PIC.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

> *Configuring Junos Fusion Provider Edge*
>
> Configuring or Expanding a Junos Fusion Enterprise

port-scaling-enlargement

IN THIS SECTION

- Syntax | 12873
- Hierarchy Level | 12873
- Description | 12873
- Required Privilege Level | 12873
- Release Information | 12873

Syntax

```
port-scaling-enlargement;
```

Hierarchy Level

```
[edit security nat source]
```

Description

Increase the source NAT port capacity on SRX5400, SRX5600, and SRX5800 devices with next-generation Services Processing Cards (SPCs).

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D60.

port-scan

IN THIS SECTION

- Syntax | 12874
- Hierarchy Level | 12874
- Description | 12874
- Options | 12875
- Required Privilege Level | 12875
- Release Information | 12875

Syntax

```
port-scan {
    threshold number;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp]
```

```
[edit security screen ids-option screen-name udp]
```

Description

Prevent port scan attacks. A port scan attack occurs when an attacker sends packets with different port numbers to scan available services. The attack succeeds if a port responds. To prevent this attack, the device internally logs the number of different ports scanned from a single remote source. For example, if a remote host scans 10 ports in 0.005 seconds (equivalent to 5000 microseconds, the default threshold setting), the device flags this behavior as a port scan attack, and rejects further packets from the remote source.

Options

`threshold` *`number`* —Number of microseconds during which the device accepts packets from the same remote source with up to 10 different port numbers. If the number of ports during the threshold period reaches 10 or more, the device rejects additional packets from the source.

- **Range:** 1000 through 1,000,000 microseconds
- **Default:** 5000 microseconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

port (Security Antivirus)

IN THIS SECTION

- Syntax | 12876
- Hierarchy Level | 12876
- Description | 12876
- Options | 12876
- Required Privilege Level | 12876
- Release Information | 12876

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update proxy]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update proxy]
[edit security utm feature-profile anti-virus sophos-engine pattern-update proxy]
```

Description

Set the port number for the proxy server.

Options

- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

port (Security Source NAT)

IN THIS SECTION

- Syntax | 12877
- Hierarchy Level | 12878
- Description | 12878
- Options | 12878
- Required Privilege Level | 12878
- Release Information | 12879

Syntax

```
port {
    block-allocation {
        active-block-timeout timeout-interval;
        block-size block-size;
        interim-logging-interval timeout-interval;
        last-block-recycle-timeout timeout-interval;
        log disable;
        maximum-blocks-per-host maximum-block-number;
    }
    deterministic {
        block-size block-size;
        host {
            address ip-address;
            address-name address-name ;
        }
    }
    no-translation;
    port-overloading-factor number;
    range {
        port-low <to port-high>;
        to port-high;
        twin-port port-low <to port-high>;
```

```
    }
}
```

Hierarchy Level

```
[edit security nat source pool pool-name]
```

Description

Specify the Port Address Translation (PAT) for a source pool.

Options

- `block-allocation`—Allocates a block of ports for translation, instead of allocating individual ports.
- `deterministic`—Maps an incoming (source) IP address and port to the specific destination address and port block, based on a predefined deterministic NAT algorithm.
- `no-translation`—Specifies that no PAT is required. This option cannot be configured with the `port-overloading-factor` or `range` options.
- `port-overloading-factor` *`number`*—Configures the port overloading capacity in source NAT. This option cannot be configured with the `no-translation` option.
- `range` *`port-low`* `<to` *`port-high`*`>`—Specifies the port number range attached to each address in the pool. This option cannot be configured with the `no-translation` option.
- `twin port`—Configures the twin port range for source NAT pools to avoid port overloading.

The remaining statements are explained separately.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. Statement updated with `block-allocation`, `deterministic`, and **twin-port** options in Junos OS Release 12.1X47-D10. Statement updated with `interim-logging-interval` and `last-block-recycle-timeout` options in Junos OS Release 15.1X49-D60.

port (Security Web Filtering Server)

IN THIS SECTION

- Syntax | 12879
- Hierarchy Level | 12879
- Description | 12879
- Required Privilege Level | 12880
- Release Information | 12880

Syntax

```
port number;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated server]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name server]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name server]
```

Description

Enter the port number for communicating with the server. (Default ports are 80, 8080, and 8081.)

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

port (Services)

IN THIS SECTION

- Syntax | 12880
- Hierarchy Level | 12880
- Description | 12881
- Required Privilege Level | 12881
- Release Information | 12881

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit services unified-access-control infranet-controller hostname]
```

Description

Specify the port on the IC Series device through which the SRX Series Firewall should establish connections (default 11123). Possible values for this statement range from 1 through 65,535.

Use this statement when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

interface (Services) | 9028

password (Services) | 12246

port (Services NAT)

IN THIS SECTION

- Syntax | 12882
- Hierarchy Level | 12882
- Description | 12882
- Options | 12883
- Required Privilege Level | 12883
- Release Information | 12883

Syntax

```
port {
    automatic (sequential | random-allocation);
    range low minimum-value high maximum-value random-allocation;
    preserve-parity;
    preserve-range;
    deterministic-port-block-allocation <block-size block-size> <include-boundary-addresses>;
    secured-port-block-allocation {
        active-block-timeout timeout-seconds;
        block-size block-size;
        max-blocks-per-user max-blocks;
    }
}
```

Hierarchy Level

```
[edit services nat pool nat-pool-name]
```

Description

Specify the NAT pool port or range. You can configure an automatically assigned port or specify a range with minimum and maximum values.

> **NOTE**: Starting in Junos OS Release 14.2R1, the `sequential` option is introduced to enable you to configure sequential allocation of ports. The `sequential` and `random-allocation` options available with the `port automatic` statement at the `[edit services nat pool nat-pool-name]` hierarchy level are mutually exclusive. You can include the `sequential` option for sequential allocation and the `random-allocation` option for random delegation of ports. By default, sequential allocation of ports takes place if you include only the `port automatic` statement at the `[edit services nat pool nat-pool- name]` hierarchy level.
>
> For releases earlier than Junos OS Release 14.2R1, configure automatic sequential port assignment by using the `auto` option at the `[edit services nat pool nat-pool-name port automatic]` hierarchy level.
>
> If you upgrade a router running a Junos OS release earlier than Release 14.2R1 to Release 14.2 and if the router contains the `port automatic` statement defined without the

`auto` option included with the configuration, the router validates the auto option present in the configuration for sequential allocation of ports.

Options

NOTE: Starting with Junos OS Release 15.1R1, the preserve-port and preserve-range functionalities are supported on MX Series routers with MS-MPCs and MS-MICs.

automatic	Cause the port assignment type to be automatically performed by the router.
sequential	Allocate ports in a sequential manner. With sequential allocation, the next available address in the NAT pool is selected only when all the ports available from an address are exhausted.
minimum-value	Lower boundary for the port range.
maximum-value	Upper boundary for the port range.
preserve-parity	Allocate ports with same parity as the original port.
preserve-range	Preserve privileged port range after translation.
random-allocation	Allocate ports within a specified range randomly.

Other options are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

`port` statement introduced before Junos OS Release 7.4.

`random-allocation` statement introduced in Junos OS Release 9.3.

`secured-port-block-allocation` statement introduced in Junos OS Release 11.2.

`deterministic-port-block-allocation` statement introduced in Junos OS Release 12.1.

`sequential` statement introduced in Junos OS Release 14.2R1.

RELATED DOCUMENTATION

Configuring Source and Destination Addresses Network Address Translation Overview

Configuring Address Pools for Network Address Port Translation (NAPT) Overview

port (Services Voice)

IN THIS SECTION

- Syntax | 12884
- Hierarchy Level | 12884
- Description | 12884
- Options | 12885
- Required Privilege Level | 12885
- Release Information | 12885

Syntax

```
port {
    minimum port-number;
    maximum port-number;
}
```

Hierarchy Level

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number compression rtp],
[edit logical-systems logical-system-name interfaces lsq-fpc/pic/port unit logical-unit-number
compression rtp]
```

Description

For voice services interfaces only, specify a range of User Datagram Protocol (UDP) destination port numbers in which RTP compression takes place.

Options

`minimum` *`port-number`*—Specify the minimum port number.

- **Range:** 0 through 65,535

`maximum` *`port-number`*—Specify the maximum port number.

- **Range:** 0 through 65,535

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Compression of Voice Traffic

port (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 12886
- Hierarchy Level | 12886
- Description | 12886
- Required Privilege Level | 12887
- Release Information | 12887

Syntax

```
port {
    automatic (random-allocation | round-robin);
    block-allocation {
        active-block-timeout timeout-interval;
        block-size block-size;
        interim-logging-interval timeout-interval;
        maximum-blocks-per-host maximum-block-number
    }
    deterministic {
        block-size block-size;
        host {
            address address;
        }
        include-boundary-addresses;
    }
    deterministic-nat-configuration-log-interval seconds;
    no-translation;
    preserve-range;
    preserve-parity;
    range {
        port-low to port-high;
        (random-allocation | round-robin);
    }
        port-overloading-factor value;
        port-overloading-usage-alarm (raise-threshold | clear-threshold)
        enhanced-port-overloading-algorithm;
}
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name]
```

Description

Configure port assignment for a source NAT pool. Port-overloading factor is configurable between 2 - 32

Through the *raise-threshold* option, you can raise threshold for alarm on port usage per overload index value through 50 to 100.

Through the *clear-threshold* option, you can clear threshold for alarm on port usage per overload index value through 40 to 100.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

The `port-overloading-usage-alarm` option added in Junos OS Release 23.4R1.

port (SRC Server)

IN THIS SECTION

- Syntax | 12887
- Hierarchy Level | 12888
- Description | 12888
- Options | 12888
- Required Privilege Level | 12888
- Release Information | 12888

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit system services service-deployment servers server-address]
```

Description

Configure the port number on which to contact the SRC server.

Options

port-number—(Optional) The TCP port number for the SRC server.

- **Default:** 3333

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Junos OS to Work with SRC Software

port-status-tlv

IN THIS SECTION

- Syntax | 12889
- Hierarchy Level | 12889
- Description | 12889

- Options | 12889
- Required Privilege Level | 12889
- Release Information | 12889

Syntax

```
port-status-tlv blocked;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile tlv-action event]
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name continuity-check]
```

Description

Define an `action-profile` consisting of various events and the action. Based on values of `port-status-tlv` in the received CCM packets, specific action such as *interface-down* can be taken using "action-profile" on page 1778 options.

Options

blocked When the incoming CCM packet contains port status TLV with value blocked, the action will be triggered for this action-profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

CFM Monitoring between CE and PE Devices

port (Subscriber Limits)

IN THIS SECTION

- Syntax | 12890
- Hierarchy Level | 12890
- Description | 12890
- Options | 12891
- Required Privilege Level | 12891
- Release Information | 12891

Syntax

```
port number {
    limit limit;
}
```

Hierarchy Level

```
[edit system services resource-monitor subscribers-limit client-type name fpc slot-number pic
number]
```

Description

Configure the maximum number of subscribers of a client type allowed to be logged in on the specified port. When that number is reached, subsequent logins on the port are denied until the current number of subscribers drops below the maximum allowed. You can also specify the maximum number of subscribers of a client type allowed per MIC, per MPC, and per chassis.

Options

number	Port number.
limit	Maximum number of subscribers. • **Range:** `1` through `256,000`

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Limiting Subscribers by Client Type and Hardware Element with Resource Monitor

Resource Monitoring for Subscriber Management and Services

port (Syslog)

IN THIS SECTION

- Syntax | 12892
- Hierarchy Level | 12892
- Description | 12892
- Options | 12892
- Required Privilege Level | 12892
- Release Information | 12892

Syntax

```
port port number;
```

Hierarchy Level

```
[edit system syslog host hostname | other-routing-engine| scc-master)]
```

Description

Specify the port number for the remote syslog server.

Options

port number Port number of the remote syslog server.

- **Range:** 0 through 65535
- **Default:** 514

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

syslog (System) | **28836**

host (System) | **8171**

port (System Log Messsages)

IN THIS SECTION

- Syntax | 12893
- Hierarchy Level | 12893
- Description | 12893
- Options | 12893
- Required Privilege Level | 12893
- Release Information | 12894

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit interfaces interface-name services-options syslog host hostname]
```

Description

Specify the UDP port for system log messages on the host. The default port is 514.

Options

port-number	Port number for system log messages.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring System Logging for Services Interfaces

port (TWAMP)

IN THIS SECTION

- Syntax | 12894
- Junos OS Hierarchy Level | 12894
- Junos OS Evolved Hierarchy Level | 12895
- Description | 12895
- Options | 12895
- Required Privilege Level | 12895
- Release Information | 12895

Syntax

```
port number;
```

Junos OS Hierarchy Level

```
[edit services rpm twamp server]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring twamp server]
```

Description

TWAMP server listening port.

Options

number—Port number.

- **Default:** 862
- **Range:** 1 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Statement introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

ports (System)

IN THIS SECTION

- Syntax | 12896
- Hierarchy Level | 12896
- Description | 12897
- Options | 12897
- Required Privilege Level | 12897
- Release Information | 12897

Syntax

```
ports {
    auxiliary {
        disable;
        insecure;
        type terminal-type;
        port-type  (mini-usb | rj45);
    }
    console {
        disable;
        insecure;
        log-out-on-disconnect;
        type terminal-type;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the properties of the console and auxiliary ports.

The ports are located on the device's craft interface. See the device's hardware documentation for port locations.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Junos OS to Set Console and Auxiliary Port Properties

ports-per-session

IN THIS SECTION

- Syntax | 12898
- Hierarchy Level | 12898
- Description | 12898
- Options | 12898
- Required Privilege Level | 12898
- Release Information | 12898

Syntax

```
ports-per-session ports;
```

Hierarchy Level

```
[edit services nat pool nat-pool-name pgcp]
```

Description

Configure the number of ports required to support Real-Time Transport Protocol (RTP), Real-Time Control Protocol (RTCP), Real-Time Streaming Protocol (RTSP), and forward error correction (FEC) for voice and video flows on the Multiservices PIC.

Options

number-of-ports—Number of ports to enable: 2 or 4 for combined voice and video services.

- **Default:** 2

Required Privilege Level

interface—To view this statement in the configuration.

interface–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

post-attack

IN THIS SECTION

Syntax | 12899

- Hierarchy Level | 12899
- Description | 12899
- Options | 12899
- Required Privilege Level | 12899
- Release Information | 12900

Syntax

```
post-attack number;
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name then notification packet-log]
```

Description

Specify the number of packets received after an attack that should be captured for further analysis of attacker behavior. If post-attack packets are not significant to your analysis or the configured attack response ends packet transfer, you can set the post-attack option to 0.

Options

number—Number of post-attack packets to be captured.

- **Range:** 0 through 255
- **Default:** 5
- **NOTE**: On SRX3400, SRX3600, and SRX3800 devices, the default value is 1.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

post-attack-timeout

IN THIS SECTION

- Syntax | 12900
- Hierarchy Level | 12900
- Description | 12900
- Options | 12901
- Required Privilege Level | 12901
- Release Information | 12901

Syntax

```
post-attack-timeout seconds;
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name then notification packet-
log]
```

Description

Specify a time limit for capturing post-attack packets for a session. No packet capture is conducted after the timeout has expired.

Options

seconds—Maximum number of seconds for post-attack packet capture.

- **Range:** 0 through 1800 seconds
- **Default:** 5

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

post-cli-implicit-firewall

IN THIS SECTION

- Syntax | 12901
- Hierarchy Level | 12902
- Description | 12902
- Default | 12902
- Required Privilege Level | 12902
- Release Information | 12902

Syntax

```
post-cli-implicit-firewall;
```

Hierarchy Level

```
[edit services rpm twamp]
```

Description

Ensure that the CLI configured (`explicit firewall`) takes precedence over the implicit firewall. The inline TWAMP client or server uses implicit firewall to achieve its functionality.

> **NOTE**: Wrong configuration of CLI firewall can lead to improper functioning of inline TWAMP client or server. After you enable or disable this configuration statement, you must restart the router, or restart remote operation using the command `restart remote-operations`, for the operation to be effective.
>
> When you issue the command `restart remote-operations`, all TWAMP sessions (both client and server) are forced to stop. You must restart all the RPM sessions and all TWAMP sessions (both client and server).

Default

The default for this configuration statement is in disabled status.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

post-convergence-lfa

IN THIS SECTION

- Syntax | 12903
- Hierarchy Level | 12903
- Description | 12903
- Options | 12903
- Required Privilege Level | 12904
- Release Information | 12904

Syntax

```
post-convergence-lfa <fate-sharing-protection fate-sharing-protection> <node-protection <cost
cost> <srlg-protection>;
```

Hierarchy Level

```
[edit logical-systems name protocols isis interface name level number],
[edit logical-systems name routing-instances name protocols isis interface name level number],
[edit protocols isis interface name level number ],
[edit routing-instances name protocols isis interface name level number]
```

Description

Configure backup paths along post-convergence on this interface.

Options

fate-sharing-protection Enable fate-sharing protection. A list of fate-sharing groups are configured on each point of local repair (PLR) with the links in each fate-sharing group identified by their respective IP addresses. The PLR associates a cost with each fate-sharing group. The fate-sharing-aware post-convergence path is computed by assuming that the cost of each link in the

same fate-sharing group as the failed link has increased the cost associated with that group.

srlg-protection Enable Shared Risk Link Group (SRLG) protection in an IS-IS network if you want IS-IS to choose a fast reroute path that does not include SRLG links in the topology-independent loop-free alternate (TI-LFA) backup paths. If you have configured fate-sharing-protection in addition to srlg-protection then both costs are added to the link metric to calculate the final TI-LFA backup path. These links have a higher metric cost and therefore TI-LFA backup computation enables IS-IS to avoid these links.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 17.4R1.

`srlg-protection` option introduced in Junos OS release 20.2R1 for MX Series and PTX Series.

RELATED DOCUMENTATION

Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS

Configuring Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS

use-for-post-convergence-lfa | **30517**

use-post-convergence-lfa | **30531**

node-protection | **11651**

post-convergence-lfa (Protocols OSPF)

IN THIS SECTION

Syntax | **12905**

- Hierarchy Level | 12905
- Description | 12905
- Options | 12905
- Required Privilege Level | 12906
- Release Information | 12906

Syntax

```
post-convergence-lfa <fate-sharing-protection fate-sharing-protection<node-protection <cost
cost> <srlg-protection>;
```

Hierarchy Level

```
[edit logical-systems name protocols ospf area interface interface-name ],
[edit logical-systems name routing-instances name protocols ospf area interface interface-name],
[edit protocols ospf area interface interface-name ],
[edit routing-instances name protocols ospf area interface interface-name]
```

Description

Configure the installation of backup-paths that follow the post-convergence paths corresponding to the failure of this interface.

Options

fate-sharing-protection Enable fate-sharing protection. A list of fate-sharing groups are configured on each point of local repair (PLR) with the links in each fate-sharing group identified by their respective IP addresses. The PLR associates a cost with each fate-sharing group. The fate-sharing-aware post-convergence path is computed by assuming that the cost of each link in the same fate-sharing group as the failed link has increased the cost associated with that group.

node-protection Enable node protection mode for topology-independent loop-free alternate (TI-LFA) routes for OSPF.

cost

Enable a node-protecting post-convergence backup path to be computed for all primary next-hops using this interface. Configure the cost of all the links used for calculating the TI-LFA post-convergence failure path cost. If node protection is enabled without configuring a cost value, then the cost is set to the maximum cost or default value of 65535.

- **Default:** 65535
- **Range:** `1` through `65535`

srlg-protection

Enable Shared Risk Link Group (SRLG) protection in an OSPFv2 network if you want OSPFv2 to choose a fast reroute path that does not include SRLG links in the topology-independent loop-free alternate (TI-LFA) backup paths. If you have configured fate-sharing-protection in addition to srlg-protection then both costs are added to the link metric to calculate the final TI-LFA backup path. These links have a higher metric cost and therefore TI-LFA backup computation enables OSPFv2 to avoid these links.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.2R1 for MX Series, PTX Series, and QFX Series.

`fate-sharing-protection` option introduced in Junos OS release 20.3R1 for MX Series and PTX Series.

`srlg-protection` option introduced in Junos OS release 20.3R1 for MX Series and PTX Series.

RELATED DOCUMENTATION

use-post-convergence-lfa (Protocols OSPF) | **30533**

Topology-Independent Loop-Free Alternate with Segment Routing for OSPF

post-service-filter

IN THIS SECTION

- Syntax | 12907
- Hierarchy Level | 12907
- Description | 12907
- Options | 12907
- Required Privilege Level | 12908
- Release Information | 12908

Syntax

```
post-service-filter filter-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet service input],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet service input]
```

Description

Define the filter to be applied to traffic after service processing. The filter is applied only if a service set is configured and selected. You can configure a postservice filter on the input side of the interface only.

The `post-service-filter` statement is not supported when the service interface is on an MS-MIC or MS-MPC.

Options

`filter-name`—Identifier for the post-service filter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Filters and Services to Interfaces

post-service-filter (Dynamic Service Sets)

IN THIS SECTION

- Syntax | 12908
- Hierarchy Level | 12909
- Description | 12909
- Options | 12909
- Required Privilege Level | 12909
- Release Information | 12909

Syntax

```
post-service-filter filter-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family service input],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family
service input]
```

Description

Define the filter to be applied to traffic after service processing. The filter is applied only if a service set is configured and selected. You can configure a post-service filter on the input side of the interface only.

Options

filter-name—Identifier for the post-service filter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family service input]` hierarchy level introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *Dynamic Service Sets Overview*
>
> *Associating Service Sets with Interfaces in a Dynamic Profile*

post-policy

IN THIS SECTION

- Syntax | 12910
- Hierarchy Level | 12910
- Release Information | 12910
- Description | 12911
- Options | 12911
- Required Privilege Level | 12912

Syntax

```
post-policy {
      exclude-non-eligible;
}
```

Hierarchy Level

```
[edit protocols bgp bmp route-monitoring],
[edit protocols bgp group group-name bmp route-monitoring],
[edit protocols bgp group neighbor group-name neighbor address bmp route-monitoring],
[edit routing-options bmp route-monitoring],
[edit routing-options bmp station station-name route-monitoring]
```

Release Information

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

For BMP route monitoring, allows you to excludes routes that are non-eligible for the decision process (for example, protocol nexthop not resolved). This represents the view of the BGP routes after running the import policy. If the import policy has rejected the BGP route, the route does not exist in the post policy view.

> **NOTE**: If post-policy is configured the router applies the import policies on the routes received so that BMP can display the routes as follows:
>
> - If an import policy modifies any parameters in the route, the route is displayed by BMP with the new values.
> - If an import policy modifies the NH and the NH is unresolved by the local router, the route is displayed by BMP as reachable.
> - If an import policy rejects a route, it is displayed as unreachable in BMP update.
> - If a route is received with unreachable NH, but the import policy doesn't reject the route, then BMP displays the route as reachable.

Using post-policy with **exclude-non-eligible** causes an unreachable BMP update in the following:

- If an import policy modifies the NH and the NH is unresolved by the local router, the route is displayed by BMP as unreachable.
- if a route is received with a NH that is unreachable, and the import policy does not reject this route, it is displayed by BMP as unreachable.

> **NOTE**: If you have initially configured post-policy route monitoring, but later delete the initial configuration and configure pre-policy route monitoring, then the previously advertised post-policy routes are withdrawn. Conversely, if you have initially configured pre-policy route monitoring and later modify it to post-policy route monitoring, then the previously advertised pre-policy routes are withdrawn.

> **NOTE**: When `post-policy` is not configured, BMP post policy feature is turned off for all routes, regardless of whether their nexthops are reachable or resolved or not.

Options

exclude-non-eligible—Exclude routes that are non-eligible for the decision process.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

BGP Monitoring Protocol

potential-violation

IN THIS SECTION

- Syntax | 12912
- Hierarchy Level | 12913
- Description | 12914
- Options | 12914
- Required Privilege Level | 12915
- Release Information | 12915

Syntax

```
potential-violation {
    authentication failures;
    cryptographic-self-test;
    decryption-failures {
        threshold value;
    }
    encryption-failures {
        threshold value;
    }
    idp;
    ike-phase1-failures {
```

```
        threshold value;
    }
    ike-phase2-failures {
        threshold value;
    }
    key-generation-self-test;
    non-cryptographic-self-test;
    policy {
        application {
            duration interval;
            size count;
            threshold value;
        }
        destination-ip {
            duration interval;
            size count;
            threshold value;
        }
        policy match {
            duration interval;
            size count;
            threshold value;
        }
        source-ip {
            duration interval;
            size count;
            threshold value;
        }
    }
    replay-attacks {
         threshold value;
    }
    security-log-percent-full percentage;
}
```

Hierarchy Level

```
[edit security alarms]
```

Description

Configure alarms for potential violation.

Options

authentication	Raise a security alarm when the device or switch detects a specified number of authentication failures (bad password attempts) before an alarm is raised.
cryptographic-self-test	Raise a security alarm when the device or switch detects a cryptographic self-test failure. Cryptographic self-tests are a set of preoperational tests that are performed after the device or switch is powered on. The self-tests run without operator intervention. No alarm is raised upon failure of a cryptographic self-test.
decryption-failures	Raise a security alarm after exceeding a specified number of decryption failures.
encryption-failures	Raise a security alarm after exceeding a specified number of encryption failures.
ike-phase1-failures	Raise a security alarm after exceeding a specified number of Internet Key Exchange (IKE) Phase 1 failures.
ike-phase2-failures	Raise a security alarm after exceeding a specified number of Internet Key Exchange (IKE) phase 2 failures.
key-generation-self-test	Raise a security alarm when the device or switch detects a key generation self-test failure. Key generation is the process of generating keys for cryptography. A key is used to encrypt and decrypt data. The self-tests run without operator intervention. No alarm is raised upon failure of a key generation self-test.
non-cryptographic-self-test	Raise a security alarm when the device or switch detects a noncryptographic self-test failure. The self-tests run without operator intervention. No alarm is raised upon failure of a noncryptographic self-test.
non-cryptographic-self-test	Raise a security alarm when the device or switch detects a noncryptographic self-test failure. The self-tests run without operator intervention. No alarm is raised upon failure of a noncryptographic self-test.
policy	Configure alarms for policy violation, based on source IP, destination IP, application, and policy rule.
replay-attacks	Raise a security alarm when the device detects a replay attack.

security-log-percent-full	Raise a security alarm when security log exceeds a specified percent of total capacity.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

power

IN THIS SECTION

- Syntax | 12915
- Hierarchy Level | 12916
- Hierarchy Level (EX Series Switches) | 12916
- Description | 12916
- Default | 12917
- Options | 12917
- Required Privilege Level | 12917
- Release Information | 12917

Syntax

```
power (off | on);
```

Hierarchy Level

```
[edit chassis fpc slot-number]
[edit chassis fpc slot  pic],
[edit chassis lcc name fpc name  pic  slot],
[edit chassis member name fpc slot  pic],
[edit chassis fpc slot slc slc-id]
[edit chassis feb slot slot-number]
```

Hierarchy Level (EX Series Switches)

```
[edit chassis fpc slot]
```

Description

Turn a Flexible PIC Concentrator (FPC), a Physical Interface Card (PIC) in the specified FPC, a sub line card (SLC), or a Forwarding Engine Board (FEB) on or off.

The `power (off | on)` command at the `edit chassis fpc slot-number` hierarchy configures the Flexible PIC Concentrator (FPC) to stay offline or to come online automatically.

The `edit chassis fpc slot-number pic pic-number power off` command turns off the power to the PIC in the specified FPC.

NOTE:

- The `power off` command is applicable only to the fixed-configuration MPC with six 40-Gigabit Ethernet ports and twenty-four 10-Gigabit Ethernet ports (MPC5E-40G10G). For other PICs, it is ignored with a syslog message.
- On an EX6200 switch, the power statement has no effect when you configure it for an uplink port FPC on the Switch Fabric and Routing Engine (SRE) module. If you configure the statement for those FPCs, the configuration will be committed, but a message that informs you that the configuration has no effect is logged in the system log. You cannot turn the power on or off for these FPCs.

- The PTX10003-80C and PTX10003-160C routers running Junos OS Evolved do not support the `set chassis fpc fpc-slot power on` command. Executing this command on an FPC which is offline could cause unintended reboots of the router.
- The `set chassis feb slot slot-number power (off | on)` command is supported on ACX7509 routers.

In Junos Node Slicing, issuing the command `set chassis fpc slot-number power off` on the BSYS (base system) powers off even those FPCs that are assigned to the Guest Network Functions (GNFs) in which unified in-service software upgrade (ISSU) is in progress. Learn more about Junos Node Slicing.

Default

`on`

Options

`off`—Take the FPC (or a PIC in the specified FPC), or SLC offline, and configure it to stay offline, as, for example, after a system reboot.

`on`—Bring the FPC (or a PIC in the specified FPC) online, and configure it to come online automatically, as, for example, after a system reboot.

`slc slc-id`—The ID of the SLC to be powered off.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

The `edit chassis fpc slot-number power (off | on)` statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.4.

The `edit chassis fpc slot-number pic pic-number power off` introduced in Junos OS Release 13.3R2.

The `edit chassis fpc slot-number pfe slot-number power (off | on)` statement introduced in Junos OS Release 19.1R1.

The `edit chassis fpc slot-number slc slc-id power (off | on)` statement introduced in Junos OS Release 21.1R1.

The `set chassis feb slot slot-number power (off | on)` command for ACX7509 routers introduced in Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

Managing Power

Removing a Line Card from an EX6200 Switch

EX8208 Switch Hardware and CLI Terminology Mapping

EX8216 Switch Hardware and CLI Terminology Mapping

EX6210 Switch Hardware and CLI Terminology Mapping

power

IN THIS SECTION

- Syntax | 12918
- Hierarchy Level | 12918
- Description | 12919
- Default | 12919
- Options | 12919
- Required Privilege Level | 12919
- Release Information | 12919

Syntax

```
power (off | on);
```

Hierarchy Level

```
[edit chassis fpc slot-number slc slc-id]
```

Description

Power on or power off sub line cards (SLCs).

Default

on

Options

slc *slc-id*	The ID of the SLC to be powered off.
power (off | on)	Power off or on the selected SLC. • **Default:** on

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Sub Line Card Overview

power-budget-priority

IN THIS SECTION

Syntax | 12920

- Hierarchy Level | 12920
- Description | 12920
- Default | 12920
- Options | 12920
- Required Privilege Level | 12921
- Release Information | 12921

Syntax

```
power-budget-priority priority;
```

Hierarchy Level

```
[edit chassis fpc slot]
```

Description

Assign a power priority to the specified line card slot on an EX6200 or EX8200 switch.

> **NOTE**: On an EX6200 switch, you cannot change the power priority of a slot containing a Switch Fabric and Routing Engine (SRE) module. Although the CLI allows you to set a different power priority for the slot, your change does not go into effect, and the power priority remains 0. A message is sent to the system log to inform you that changing the power priority of the slot is unsupported.

Default

All line card slots are initially assigned the lowest priority, with the exception of slot 4 and slot 5 on the EX6200 switch, which always are assigned a priority of 0.

Options

priority—Assigned power priority for the slot, with 0 being the highest priority:

- 0 through 9 for an EX6200 switch
- 0 through 7 for an EX8208 switch
- 0 through 15 for an EX8216 switch

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the Power Priority of Line Cards (CLI Procedure)

power-mode-disable

IN THIS SECTION

- Syntax | 12921
- Hierarchy Level | 12922
- Release Information | 12922
- Description | 12922
- Required Privilege Level | 12922

Syntax

power-mode-disable;

Hierarchy Level

```
[edit security flow]
```

Release Information

Statement introduced in Junos OS Release 21.3R1.

Description

PowerMode is enabled by default. To disable the feature, use `power-mode-disable` statement at the `[edit security flow]` hierarchy level.

Required Privilege Level

security

RELATED DOCUMENTATION

PowerMode

power-mode-ipsec

IN THIS SECTION

- Syntax | 12923
- Hierarchy Level | 12923
- Description | 12923
- Required Privilege Level | 12923
- Release Information | 12923

Syntax

```
power-mode-ipsec;
```

Hierarchy Level

```
[edit security flow (Security Flow)]
```

Description

Enable PowerMode IPsec. processing. PMI is a new mode of operation that provides IPsec performance improvements.

For SRX4100, SRX4200 devices running Junos OS Release 18.4R1, SRX4600 devices running Junos OS Release 20.4R1, and vSRX Virtual Firewall instances running Junos OS Release 18.3R1, you can enable or disable the PMI. Starting in Junos OS Release 21.1R1, you can enable or disable the PMI on MX-SPC3 services card.

If you use Junos OS Release 18.3R1, you must reboot the device for the configuration to take effect.

From Junos OS Release 18.4R1 and later, you don't need reboot the device after enabling or disabling this feature.

Packets cannot go through the PMI when firewall or advanced security services are combined with IPsec. Hence, PMI must not be used when firewall or advanced security services are combined with IPsec.

Required Privilege Level

flow-tap

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Improving IPsec Performance with PowerMode IPsec

power-mode-ipsec-qat

IN THIS SECTION

- Syntax | 12924
- Hierarchy Level | 12924
- Description | 12924
- Required Privilege Level | 12924
- Release Information | 12924

Syntax

```
power-mode-ipsec-qat;
```

Hierarchy Level

```
[edit security flow]
```

Description

Enable QuickAssist (QAT) for encryption of cryptographic operation.

PowerMode IPsec (PMI) uses AES-NI and QAT for encryption and FPGA for decryption of cryptographic operation to enhance IPSec VPN performance.

PMI with QAT is supported on SRX5400, SRX5600, and SRX5800 with SPC3 cards.

Required Privilege Level

flow-tap

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

inline-fpga-crypto | 8668

Improving IPsec Performance with PowerMode IPsec

application-services (Security Forwarding Process)

show security flow status

pp0 (Dynamic PPPoE)

IN THIS SECTION

- Syntax | 12925
- Hierarchy Level | 12926
- Description | 12927
- Required Privilege Level | 12927
- Release Information | 12927

Syntax

```
pp0 {
    unit logical-unit-number {
        keepalives interval seconds;
        no-keepalives;
        pppoe-options {
            underlying-interface interface-name;
            server;
        }
        ppp-options {
            aaa-options aaa-options-name;
            authentication [ authentication-protocols ];
            chap {
                challenge-length minimum minimum-length maximum maximum-length;
            }
            ignore-magic-number-mismatch;
            initiate-ncp (ip | ipv6 | dual-stack-passive)
```

```
            ipcp-suggest-dns-option;
            mru size;
            mtu (size | use-lower-layer);
            on-demand-ip-address;
            pap;
            peer-ip-address-optional;
        }
        family inet {
            unnumbered-address interface-name;
            address address;
            service {
                input {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                    post-service-filter filter-name;
                }
                output {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                }
            }
            filter {
                input filter-name {
                    precedence precedence;
                }
                output filter-name {
                    precedence precedence;
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces]
```

Description

Configure the dynamic PPPoE logical interface in a dynamic profile. When the router creates a dynamic PPPoE logical interface on an underlying Ethernet interface configured with PPPoE (`ppp-over-ether`) encapsulation, it uses the information in the dynamic profile to determine the properties of the dynamic PPPoE logical interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

- *Configuring a PPPoE Dynamic Profile*
- *Configuring Dynamic Authentication for PPP Subscribers*
- Configuring PPPoE

ppm

IN THIS SECTION

- Syntax | **12928**
- Syntax (QFX5110, QFX5120, QFX5200, and QFX5210) | **12928**
- Hierarchy Level | **12928**
- Description | **12928**
- Default | **12929**
- Options | **12929**

Required Privilege Level | 12929
Release Information | 12929

Syntax

```
ppm {
    no-delegate-processing;
    redistribution-timer;
}
```

Syntax (QFX5110, QFX5120, QFX5200, and QFX5210)

```
ppm {
    no-delegate-processing;
    inline-processing-enable;
}
```

Hierarchy Level

```
[edit routing-options]
```

Description

Disable distributed periodic packet management (PPM) to the Packet Forwarding Engine (on routers), to access ports (on EX3200 and EX4200 switches, and QFX Series), or to line cards (on EX6200 and EX8200 switches).

After you disable PPM, PPM processing continues to run on the Routing Engine.

In Junos OS Release 8.2, PPM was moved from the Routing Engine to the Packet Forwarding Engine, access ports, or line cards. The `no-delegate-processing` statement disables the default behavior and restores the legacy behavior.

Default

Distributed PPM processing is enabled for all protocols that use PPM such as:

- Bidirectional Forwarding Detection (BFD)
- Connectivity Fault Management (CFM)
- Link Aggregation Control Protocol (LACP)
- Link Fault Management (LFM)
- Real-time Performance Monitoring (RPM)
- Spanning Tree Protocol (STP)
- Synchronous Ethernet (SYNCE)
- Virtual Router Redundancy Protocol (VRRP)

> **NOTE**: The protocols BFD and CFM are supported on SRX300, SRX320, SRX340, SRX345 and SRX1500 devices.
>
> The protocol LACP is supported on SRX5400, SRX 5600, and SRX5800 devices.

Options

`no-delegate-processing`	Disable PPM to the Packet Forwarding Engine, access ports, or line cards. Distributed PPM is enabled by default.
`redistribution-timer`	Ensures that link aggregation and spanning tree protocols work seamlessly on chassis or virtual chassis based devices during ISSU/NSSU, NSR and GRES. Recommended value is 120 seconds.
`inline-processing-enable`	Enable inline processing. Inline processing allows you to set BFD timers to less than one second. On platforms that support hardware-assisted inline BFD, use this option to enable it.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

`inline-processing-enable` option introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

Ensuring That Distributed ppm Is Not Disabled

ppm (Ethernet Switching)

IN THIS SECTION

- Syntax | 12930
- Hierarchy Level | 12930
- Description | 12930
- Default | 12931
- Options | 12931
- Required Privilege Level | 12931
- Release Information | 12931

Syntax

```
ppm {
    centralized | distributed | inline;
}
```

Hierarchy Level

```
[edit protocols lacp]
```

Description

Configure PPM processing options for Link Aggregation Control Protocol (LACP) packets.

This command configures the PPM processing options for LACP packets only. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the **no-delegate-processing** configuration statement in the `[edit routing-options ppm]` statement hierarchy.

Default

Distributed PPM processing is enabled for all packets that use PPM.

Options

`distributed`—Specifies distributed non-inline PPM processing.

`inline`—Specifies distributed inline PPM processing.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

- Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)
- Configuring Distributed Periodic Packet Management

ppp

IN THIS SECTION

- Syntax | 12932
- Hierarchy Level | 12942

- Description | 12942
- Options | 12942
- Required Privilege Level | 12944
- Release Information | 12944

Syntax

```
ppp {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
  }
  recover-time seconds;
```

```
}
authentication {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
    }
    echo-rep {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    echo-req {
     bandwidth packets per second;
```

```
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ipcp {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
```

```
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ipv6cp {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
```

```
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   isis {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
```

```
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    lcp {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   mlppp-lcp {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
```

```
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    mplscp {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
```

```
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    unclassified {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
```

Hierarchy Level

```
[edit system ddos-protection protocols]
```

Description

Configure PPP control traffic.

Options

Table 96: PPP Options

> aggregate	Configure aggregate for all PPP control traffic
+ apply-groups	Groups from which to inherit configuration data

+ apply-groups-except	Don't inherit configuration data from these groups
> padi	Configure PPPoE PADI
> padm	Configure PPPoE PADM
> padn	Configure PPPoE PADN
> pado	Configure PPPoE PADO
> padr	Configure PPPoE PADR
> pads	Configure PPPoE PADS
> padt	Configure PPPoE PADT
authentication	Configure Authentication Protocol
> echo-rep	Configure LCP Echo Reply
> echo-req	Configure LCP Echo Request
> ipcp	Configure IP Control Protocol
> ipv6cp	Configure IPv6 Control Protocol
> isis	Configure ISIS Protocol
> lcp	Configure Link Control Protocol
> mlppp-lcp	Configure MLPPP LCP
> mplscp	Configure MPLS Control Protocol
> unclassified	Configure unclassified PPP control traffic

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos Evolved OS Release 23.4R1.

ppp-access-profile

IN THIS SECTION

- Syntax | 12944
- Hierarchy Level | 12944
- Description | 12945
- Options | 12945
- Required Privilege Level | 12945
- Release Information | 12945

Syntax

```
ppp-access-profile profile-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Specify the profile used to validate all Point-to-Point Protocol (PPP) session requests through L2TP tunnels established to the local gateway address.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options

profile-name—Identifier for the PPP profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Access Profiles for L2TP Tunnel Groups

ppp (Group Profile)

IN THIS SECTION

- Syntax | 12946
- Hierarchy Level | 12946
- Description | 12946
- Options | 12947
- Required Privilege Level | 12948

Release Information | 12948

Syntax

```
ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-pool framed-pool;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    ppp-options {
        aaa-options aaa-options-name;
        chap;
        ignore-magic-number-mismatch;
        initiate-ncp (ip | ipv6 | dual-stack-passive)
        ipcp-suggest-dns-option;
        mru;
        mtu;
        pap;
        peer-ip-address-optional;
    }
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
}
```

Hierarchy Level

```
[edit access group-profile profile-name]
```

Description

Configure PPP properties for a group profile.

Options

cell-overhead

Configure the session to use Asynchronous Transfer Mode (ATM)-aware egress shaping on the IQ2 PIC.

encapsulation-overhead

Configure the encapsulation overhead for class-of-service calculations.

- **Values:** *bytes*—The number of bytes used as encapsulation overhead for the session.

framed-pool

Configure the address pool.

- **Values:** *framed-pool*—References a configured address pool.

idle-timeout

Configure the idle timeout for a user. The router might consider a PPP session to be idle because of the following reasons:

- There is no ingress traffic on the PPP session.
- There is no egress traffic.
- There is neither ingress or egress traffic on the PPP session.
- There is no ingress or egress PPP control traffic. This is applicable only if keepalives are enabled.
- **Values:** *seconds*—Number of seconds a user can remain idle before the session is terminated.
- **Range:** 0 through 4,294,967,295 seconds
- **Default:** 0

interface-id ***interface-id***

Configure the interface identifier.

- **Values:** *interface-id*—Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the `[edit interfaces interface-name unit local-unit-number dial-options]` hierarchy level. For more information about the interface ID, see Services Interface Naming Overview.

keepalive

Configure the keepalive interval for an L2TP tunnel.

- **Values:** *seconds*—Time period that must elapse before the Junos OS checks the status of the Point-to-Point Protocol (PPP) session by sending an echo request to the peer.

For L2TP on MX Series routers, the minimum recommended interval is 30 seconds. A value of 0 disables generation of keepalive messages from the LNS.

- **Range:** 0 through 32,767 seconds
- **Default:** 30 seconds

primary-dns Configure the primary Domain Name System (DNS) server.

- **Values:** *primary-dns*—An IPv4 address.

primary-wins Configure the primary Windows Internet name server.

- **Values:** *primary-wins*—An IPv4 address.

secondary-dns Configure the secondary DNS server.

- **Values:** *secondary-dns*—An IPv4 address.

secondary-wins Configure the secondary Windows Internet name server.

- **Values:** *secondary-wins*—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

L2TP LNS Inline Service Interfaces

ppp-options

IN THIS SECTION

- Syntax | 12949
- Hierarchy Level | 12950
- Description | 12950
- Required Privilege Level | 12950
- Release Information | 12951

Syntax

```
ppp-options {
    authentication [ authentication-protocols ];
        mru size;
    mtu (size | use-lower-layer);
    chap {
        access-profile name;
        challenge-length minimum minimum-length maximum maximum-length;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile profile-name;
    initiate-ncp (ip | ipv6 | dual-stack-passive)
    ipcp-suggest-dns-option;
    lcp-max-conf-req number
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-max-conf-req number
    ncp-restart-timer milliseconds;
    on-demand-ip-address
    pap {
```

```
        access-profile name;
        default-pap-password password;
        local-name name;
        local-password password;
        passive;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

On interfaces with PPP encapsulation, configure PPP-specific interface properties.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

> **BEST PRACTICE**: On inline service (si) interfaces for L2TP, only the **chap** and **pap** statements are typically used for subscriber management. We recommend that you leave the other statements subordinate to **ppp-options**—including those subordinate to **chap** and **pap**—at their default values.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the PPP Challenge Handshake Authentication Protocol

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

ppp-options (Dynamic PPP)

IN THIS SECTION

- Syntax | 12951
- Hierarchy Level | 12952
- Description | 12952
- Options | 12953
- Required Privilege Level | 12953
- Release Information | 12953

Syntax

```
ppp-options {
    aaa-options aaa-options-name;
    authentication [ authentication-protocols ];
    chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
    }
    ignore-magic-number-mismatch;
    initiate-ncp (dual-stack-passive | ipv6 | ip)
    ipcp-suggest-dns-option;
    lcp-connection-update;
    mru size;
```

```
    mtu (size | use-lower-layer);
    on-demand-ip-address;
    pap;
    peer-ip-address-optional;
    local-authentication {
        password password;
        username-include {
            circuit-id;
            delimiter character;
            domain-name name;
            mac-address;
            remote-id;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-
interface-unit”].
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit”]
```

Description

Configure PPP-specific interface properties in a dynamic profile.

NOTE: PPP options can also be configured in a group profile with the `ppp-options (L2TP)` statement. The following behavior determines the interaction between the PPP options configured in a group profile and the PPP options configured in a dynamic profile:

- When PPP options are configured only in the group profile, the group profile options are applied to the subscriber.
- When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options

between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

Options

lcp-connection-update

Enable PPP to act on a Connection-Status-Message VSA (26–218) received by authd in either a RADIUS Access-Accept message or a CoA message. PPP conveys the contents of the VSA in an LCP Connection-Update-Request message to the remote peer, such as a home gateway. This action requires the following to be true:

- At least the first address family has been successfully negotiated and the session is active.
- The router LCP is in the Opened state.

Otherwise PPP takes no action on the VSA. If you do not enable the `lcp-connection-update` option, PPP processes the notification from authd, but takes no action.

- **Default:** Disabled

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]` hierarchy level introduced in Junos OS Release 12.2.

`lcp-connection-update` option added in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Dynamic Profiles Overview

Configuring a PPPoE Dynamic Profile

Configuring Dynamic Authentication for PPP Subscribers

Attaching Dynamic Profiles to Static PPP Subscriber Interfaces

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

How to Configure RADIUS-Sourced Connection Status Updates to CPE Devices

ppp-options (L2TP)

IN THIS SECTION

- Syntax | 12954
- Hierarchy Level | 12954
- Description | 12955
- Required Privilege Level | 12955
- Release Information | 12955

Syntax

```
ppp-options {
    aaa-options aaa-options-name;
    chap;
    ignore-magic-number-mismatch;
    initiate-ncp (ip | ipv6 | dual-stack-passive)
    ipcp-suggest-dns-option;
    mru;
    mtu;
    pap;
    peer-ip-address-optional;
}
```

Hierarchy Level

```
[edit access group-profile profile-name ppp]
```

Description

Configure PPP-specific properties in a group profile that applies to tunneled PPP subscribers at the LNS.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

NOTE: PPP options can also be configured for an inline service interface within a dynamic profile with the `ppp-options (Dynamic PPP)` statement. The following behavior determines the interaction between the PPP options configured in a group profile and the PPP options configured in a dynamic profile:

- When PPP options are configured only in the group profile, the group profile options are applied to the subscriber.
- When PPP options are configured in both the dynamic profile and the group profile, the group profile options are applied to the subscriber only when the dynamic profile PPP options do not include any of the following attributes: "aaa-options" on page 1614, "chap" on page 3291, "ipcp-suggest-dns-option" on page 9289, "mru" on page 10994, "mtu" on page 11034, "pap" on page 12171, and "peer-ip-address-optional" on page 12373. When any of these attributes is present, the dynamic profile is applied to the subscriber.

  When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

`mtu` statement introduced in Junos OS Release 14.2

RELATED DOCUMENTATION

Configuring the PPP Attributes for a Group Profile

L2TP LNS Inline Service Interfaces

ppp-over-ether

IN THIS SECTION

- Syntax | 12956
- Hierarchy Level | 12956
- Description | 12956
- Required Privilege Level | 12956
- Release Information | 12957

Syntax

```
ppp-over-ether;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number encapsulation]
```

Description

This encapsulation is used for underlying interfaces of pp0 interfaces. This encapsulation is supported on Fast Ethernet interface, Gigabit Ethernet interface, and Redundant Ethernet interface. When Redundant Ethernet interface is used as underlying interface, an existing pppoe session can be continued in case of failover.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 11.2.

This encapsulation is supported for Redundant Ethernet interface in Junos OS Release 11.2.

RELATED DOCUMENTATION

> *Understanding Ethernet Interfaces*

ppp (Profile)

IN THIS SECTION

- Syntax | 12957
- Hierarchy Level | 12958
- Description | 12958
- Options | 12958
- Required Privilege Level | 12958
- Release Information | 12958

Syntax

```
ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-ip-address address;
    framed-pool framed-pool;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    primary-dns primary-dns;
```

```
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
}
```

Hierarchy Level

```
[edit access profile profile-name client client-name]
```

Description

Configure PPP properties for a client profile.

Options

framed-ip-address (SRX Series, T Series only) Specify a framed IP address.

- **Values:** *address*—The IP version 4 (IPv4) prefix.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement `idle-timeout` introduced in Junos OS Release 11.1 for the QFX Series.

Statement `idle-timeout` introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

ppp-subscriber-services

IN THIS SECTION

- Syntax | 12959
- Hierarchy Level | 12959
- Description | 12959
- Options | 12959
- Required Privilege Level | 12960
- Release Information | 12960

Syntax

```
ppp-subscriber-services (disable | enable);
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable dynamic PPP subscriber services on non-PPPoE interfaces on certain PICs.

NOTE: When you include this statement, the relevant PICs restart. This action disrupts subscribers already logged in through those PICs. You can confirm completion of the restart by issuing the **show chassis pic fpc-slot *slot-number* pic-slot *slot-number*** command.

Options

`disable`—Disable subscriber services.

`enable`—Enable subscriber services.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

> *Attaching Dynamic Profiles to MLPPP Bundles*
>
> *Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces*

pppoe (System Processes)

IN THIS SECTION

- Syntax | 12960
- Hierarchy Level | 12961
- Description | 12961
- Options | 12961
- Required Privilege Level | 12961
- Release Information | 12961

Syntax

```
pppoe {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Enable users to connect to a network of hosts over a bridge or access concentrator.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the Point-to-Point Protocol over Ethernet process.
- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

pppoe (DDoS Protection)

IN THIS SECTION

- Syntax | 12962
- Hierarchy Level | 12970
- Description | 12971
- Options | 12971
- Required Privilege Level | 12971
- Release Information | 12971

Syntax

```
pppoe {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
```

```
  fpc name {
   bandwidth-scale percent;
```

```
      burst-scale percent;
      disable-fpc;
     }


     recover-time seconds;





    }
    padi {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
```

```
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   padm {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
```

```
    recover-time seconds;
    timeout-active-flows;


  }
  padn {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;
```

```
}
pado {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
padr {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



   }
   pads {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   padse {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   padt {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
```

Hierarchy Level

```
[edit system ddos-protection protocols]
```

Description

Configure PPPoE control traffic.

Options

Table 97: PPPoE Options

> aggregate	Configure aggregate for all PPPoE control traffic
+ apply-groups	Groups from which to inherit configuration data
+ apply-groups-except	Don't inherit configuration data from these groups
> padi	Configure PPPoE PADI
> padm	Configure PPPoE PADM
> padn	Configure PPPoE PADN
> pado	Configure PPPoE PADO
> padr	Configure PPPoE PADR
> pads	Configure PPPoE PADS
> padt	Configure PPPoE PADT

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos Evolved OS Release 23.4R1.

pppoe-options (SRX Series)

IN THIS SECTION

- Syntax | 12972
- Hierarchy Level | 12972
- Description | 12972
- Options | 12973
- Required Privilege Level | 12973
- Release Information | 12973

Syntax

```
pppoe-options {
    access-concentrator name ;
    auto-reconnect seconds;
    (client | server);
    ignore-eol-tag;
    service-name name;
    underlying-interface interface-name;
}
```

Hierarchy Level

```
[edit interfaces pp0 unit logical-unit-number],
[edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number]
```

Description

Configure PPP over Ethernet-specific interface properties.

Options

access-concentrator *name*	(SRX Series devices with Point-to-Point Protocol over Ethernet (PPPoE) interfaces) Configure the name of the access concentrator. If you configure a specific access concentrator name on the client and the same access concentrator name server is available, then a PPPoE session is established. If there is a mismatch between the access concentrator names of the client and the server, the PPPoE session gets closed.
auto-reconnect *seconds*	Configure the amount of time to wait before reconnecting after a session has terminated.
client	Configure the device to operate in the PPPoE client mode.
idle-timeout *seconds*	Configure the maximum time that a session can be idle.
ignore-eol-tag	Disable the `End-of-List` tag to process the tags after the `End-of-List` tag in a PPPoE Active Discovery Offer (PADO) packet.
service-name *name*	Configure the service to be requested from the PPP over Ethernet server; that is, the access concentrator. For example, you can use this statement to indicate an Internet service provider (ISP) name or a class of service.
server	Configure the device to operate in the PPPoE server mode.
underlying-interface *interface-name*	Configure the interface on which PPP over Ethernet is running.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 12.3X48 to include `ignore-eol-tag` statement.

RELATED DOCUMENTATION

Example: Configuring PPPoE Interfaces

pppoe-options

IN THIS SECTION

- Syntax | 12974
- Hierarchy Level | 12974
- Description | 12974
- Required Privilege Level | 12975
- Release Information | 12975

Syntax

```
pppoe-options {
    access-concentrator name;
    auto-reconnect seconds;
    (client | server);
    service-name name;
    underlying-interface interface-name;
    ppp-max-payload ppp-max-payload
}
```

Hierarchy Level

```
[edit interfaces pp0 unit logical-unit-number],
[edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number],
[edit interface ppp interface unit logical-unit-number ppp-max-payload ppp-max-payload]
```

Description

Configure PPP over Ethernet-specific interface properties.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

The maximum payload allowed on an Ethernet frame is 1500 bytes. For a PPPoE interface, the PPPoE header uses 6 bytes and the PPP protocol ID uses 2 bytes. This restricts the maximum MTU size on a PPPoE interface to 1492 bytes, which can cause frequent fragmentation and reassembly of larger PPP packets received over the PPPoE interface. To prevent frequent fragmentation and reassembly for PPP packets over Ethernet, you can configure the maximum transmission unit (MTU) and MRU sizes for PPP subscribers.

For PPPoE subscribers, the PPP MRU or PPP MTU size can be greater than 1492 bytes if the PPP-Max-Payload tag is received in the PPPoE Active Discovery Request (PADR) packets.

The PPP-Max-Payload option allows you to override the default behavior of the PPPoE client by providing a maximum size that the PPP payload can support in both sending and receiving directions. The PPPoE server might allow the negotiation of an MRU larger than 1492 octets and the ability to use an MTU larger than 1500 octets.

It is important to set an appropriate value for the MTU size of the physical interface before setting `ppp-max-payload`. The value of `mtu` must be greater than the value of `ppp-max-payload`.

To enable Jumbo frames refer Understanding Jumbo Frames Support for Ethernet Interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`client` Statement introduced in Junos OS Release 8.5.

`server` Statement introduced in Junos OS Release 8.5.

`ppp-max-payload` Statement introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Configuring a PPPoE Interface

pppoe-options (Dynamic PPPoE)

IN THIS SECTION

- Syntax | 12976
- Hierarchy Level | 12976
- Description | 12976
- Required Privilege Level | 12976
- Release Information | 12977

Syntax

```
pppoe-options {
    underlying-interface interface-name;
    server;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"]
```

Description

Configure the underlying interface and PPPoE server mode for a dynamic PPPoE logical interface in a dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring a PPPoE Dynamic Profile

Configuring Dynamic PPPoE Subscriber Interfaces

pppoe-tags (Adjustment Control Profiles)

IN THIS SECTION

- Syntax | 12977
- Hierarchy Level | 12977
- Description | 12978
- Options | 12978
- Required Privilege Level | 12978
- Release Information | 12978

Syntax

```
pppoe-tags {
    priority priority;
    algorithm algorithm;
}
```

Hierarchy Level

```
[edit class-of-service adjustment-control-profiles profile-name application]
```

Description

Configure the shaping rate adjustment controls for the Point-to-Point Protocol over Ethernet (PPPoE) Tags application.

Options

priority Priority of the Point to Point Protocol over Ethernet IA Tags application in the adjustment control profile.

- **Range:** 1 through 10; 1 being the highest priority.
- **Default:** 2

algorithm Rate adjustment algorithm used by the Point to Point Protocol over Ethernet (PPPoE) IA Tags application.

- Values:
  - adjust-never—Do not perform rate adjustments.
  - adjust-always—Adjust the shaping rate unconditionally.
  - adjust-less—Adjust the shaping rate if it is less than the configured value.
  - adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
  - adjust-greater—Adjust the shaping rate if it is greater than the configured value.
  - adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.
- **Default:** adjust-less

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

CoS Adjustment Control Profiles Overview

Configuring CoS Adjustment Control Profiles

Verifying the CoS Adjustment Control Profile Configuration

adjustment-control-profiles | 1961

application (Adjustment Control Profiles) | 2301

pppoe-underlying-options (Dynamic VLAN Interface Sets)

IN THIS SECTION

- Syntax | 12979
- Hierarchy Level | 12979
- Description | 12979
- Required Privilege Level | 12980
- Release Information | 12980

Syntax

```
pppoe-underlying-options {
    max-sessions number;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-set "$junos-interface-set-name"]
```

Description

Configure PPPoE-specific interface properties in the dynamic profile that defines the agent circuit identifier (ACI) interface set. An ACI interface set is a logical collection of subscriber interfaces that

originate at the same household or on the same access-loop port. Configuring PPPoE-specific interface properties for an ACI interface set enables you to apply these attributes to all subscribers on a per-household basis.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

> **NOTE**: When you configure PPPoE-specific interface properties for an ACI interface set, only the `max-sessions` statement is currently supported.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Dynamic VLANs Based on Agent Circuit Identifier Information

Agent Circuit Identifier-Based Dynamic VLANs Overview

pppoe-underlying-options (Static and Dynamic Subscribers)

IN THIS SECTION

- Syntax | 12981
- Hierarchy Level | 12981
- Description | 12981
- Required Privilege Level | 12981
- Release Information | 12981

Syntax

```
pppoe-underlying-options {
    access-concentrator name;
    dynamic-profile profile-name;
    direct-connect
    duplicate-protection;
    max-sessions number;
    max-sessions-vsa-ignore;
    service-name-table table-name;
    short-cycle-protection <lockout-time-min minimum-seconds> <lockout-time-max maximum-seconds>
<filter [aci]>;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure PPPoE-specific interface properties for the underlying interface on which the router creates a static or dynamic PPPoE logical interface. The underlying interface must be configured with PPPoE (`ppp-over-ether`) encapsulation.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring PPPoE

Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces

Assigning a Service Name Table to a PPPoE Underlying Interface

pps-limit (Hierarchical Policer)

IN THIS SECTION

- Syntax | 12982
- Hierarchy Level | 12982
- Description | 12983
- Options | 12983
- Required Privilege Level | 12983
- Release Information | 12983

Syntax

```
pps-limit pps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer hierarchical-policer-name aggregate
if-exceeding],
[edit dynamic-profiles profile-name firewall hierarchical-policer hierarchical-policer-name premium
if-exceeding],
[edit No Link Title hierarchical-policer hierarchical-policer-name aggregate if-exceeding],
[edit No Link Title hierarchical-policer hierarchical-policer-name premium if-exceeding]
```

Description

Configure the maximum bandwidth in packets per second (pps) for premium or aggregate traffic in a hierarchical policer.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allow bursts of EF traffic for short periods and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options

pps—Specify the number of packets per second either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), or `m` (1000000).

- **Range:** 2 through 24414062
- **Default:** None

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Hierarchical Policer Configuration Overview

Policer Bandwidth and Burst-Size Limits

Policer Color-Marking and Actions

Single Token Bucket Algorithm

Determining Proper Burst Size for Traffic Policers

aggregate (Hierarchical Policer) | **2064**

burst-size-limit (Hierarchical Policer) | **3084**

premium (Hierarchical Policer) | **13110**

pps-limit (Policer)

IN THIS SECTION

- Syntax | 12984
- Hierarchy Level | 12984
- Description | 12984
- Options | 12985
- Required Privilege Level | 12985
- Release Information | 12985

Syntax

```
pps-limit pps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name if-exceeding-pps],
[edit firewall policer policer-name if-exceeding-pps],
[edit logical-systems logical-system-name firewallpolicer policer-name if-exceeding-pps]
```

Description

For a single-rate two-color policer, configure the packets-per-second (pps) limit as a number of packets per second. Single-rate two-color policing uses the single token bucket algorithm to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the pps limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with `low` packet loss priority (PLP) and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow

might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

> **NOTE**: This statement specifies the pps limit as an absolute number of packets per second. You cannot use the pps limit as a percentage of interface bandwidth.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allow more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allow bursts of EF traffic for short periods and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options

pps—Specify the number of packets per second either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), or `m` (1000000).

- **Range:** 2 through 24414062
- **Default:** None

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Two-Color Policer Configuration Overview

Policer Bandwidth and Burst-Size Limits

Policer Color-Marking and Actions

Single Token Bucket Algorithm

Determining Proper Burst Size for Traffic Policers

bandwidth-percent | **2866**

burst-size-limit (Policer) | 3086

pptp

IN THIS SECTION

- Syntax | 12986
- Hierarchy Level | 12986
- Description | 12987
- Options | 12987
- Required Privilege Level | 12987
- Release Information | 12987

Syntax

```
pptp {
    disable;
    traceoptions {
        flag {
            all {
                extensive
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Configure pptp to specify the Point-to-Point Tunneling Protocol (PPTP) ALG on the device. PPTP ALG is a TCP-based ALG.

Options

- `disable`—Disable the PPTP ALG. By default, the PPTP ALG is enabled.
- `traceoptions`—Configure PPTP ALG tracing options.
  - `flag`—Trace operation to perform.
    - `all`—Trace all events.
    - *`extensive`*—Display extensive amount of data.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Security Policies User Guide for Security Devices

prbs

IN THIS SECTION

Syntax | 12988

- Hierarchy Level | 12988
- Description | 12988
- Default | 12988
- Options | 12988
- Required Privilege Level | 12988
- Release Information | 12989

Syntax

```
(prbs | no-prbs);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify whether OTN payload Pseudo-Random Binary Sequence (PBRS) is enabled or disabled.

Default

By default, OTN payload prbs is disabled.

Options

prbs	Enable OTN payload PBRS.
no-prbs	Disable OTN payload PBRS.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

prbs-test-start

IN THIS SECTION

- Syntax | 12989
- Description | 12989
- Options | 12990
- Required Privilege Level | 12991
- Release Information | 12991

Syntax

```
test interfaces  ifd-name prbs-test-start pattern-type type direction 0/1 flip 0/1
```

Description

You can check the physical link connectivity by issuing this command that starts collecting the Pseudo Random Binary Sequence (PRBS) statistics. The PRBS pattern generation and verification validates the physical link connectivity in a routers. If the PRBS test passes with PRBS31 pattern type consistently, it indicates that the quality of signal received is good.

This command provides the PRBS test statistics while test is in progress. Use the `clear interfaces statistics` command to clear the collected statistics.

Use the `show interfaces interface-name prbs-stats` command to view the collected statistics.

NOTE:

- While running PRBS statistics, the link will be down.
- The interface link status goes down when PRBS is enabled and the interface state is back to the original state when PRBS is disabled.
- On MX10003 routers, the PRBS58 pattern is supported only on JNP-MIC1-MACSEC MIC.

Issue the "prbs-test-stop" on page 12991 command to stop collecting the PRBS statistics .

For the step-by-step procedure on how to collect and view the PRBS statistics, refer *Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test*.

Options

ifd-name—Name of the interface.

type—Specifies the pattern type, that is in the range from 7 to 58.

Following pattern types are supported:

Pattern Type	Pattern Name
7	PRBS7
9	PRBS9
11	PRBS11
15	PRBS15
23	PRBS23
31	PRBS31 **NOTE**: Recommended pattern to check the quality of the received link.

(Continued)

Pattern Type	Pattern Name
58	PRBS58

`direction`—Specifies to configure transmit or receive PRBS pattern.

`flip`—Specifies if the pattern bits needs to be flipped or not.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

prbs-test-stop | 12991

show interfaces prbs-stats | 21170

clear interfaces statistics | 3651

Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test

prbs-test-stop

IN THIS SECTION

- Syntax | 12992
- Description | 12992
- Options | 12992
- Required Privilege Level | 12992
- Sample Output | 12992

Release Information | 12992

Syntax

```
test interfaces  ifd-name prbs-test-stop direction (0|1)
```

Description

Use this command to stop collecting the Pseudo Random Binary Sequence (PRBS) statistics that is initiated using "prbs-test-start" on page 12989 command. This command only disables the statistics collection and does not clear the statistics collected. To clear the collected statistics, issue the `clear interfaces statistics` command.

For the step-by-step procedure on how to collect and view the PRBS statistics, refer *Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test*.

Options

ifd-name	Name of the interface.
direction	Direction to transmit or receive PRBS pattern.

Required Privilege Level

view

Sample Output

Stopping PRBS test statistics collection

```
user@host> test interface et-0/1/2 prbs-test-stop direction 1
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

prbs-test-start | 12989

show interfaces prbs-stats | 21170

clear interfaces statistics | 3651

Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test

pre-attack

IN THIS SECTION

- Syntax | 12993
- Hierarchy Level | 12993
- Description | 12993
- Options | 12994
- Required Privilege Level | 12994
- Release Information | 12994

Syntax

```
pre-attack number;
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name then notification packet-
log]
```

Description

Specify the number of packets received before an attack that should be captured for further analysis of attacker behavior.

Options

number—Number of pre-attack packets.

- **Range:** 1 through 255
- **Default:** 1

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

pre-filter-shellcode

IN THIS SECTION

- Syntax | 12994
- Hierarchy Level | 12995
- Description | 12995
- Required Privilege Level | 12995
- Release Information | 12995

Syntax

```
pre-filter-shellcode;
```

Hierarchy Level

```
[edit security idp sensor-configuration ips]
```

Description

Enable to pre-filter the shell code and protects it from buffer overflow attacks. By default this setting is enabled.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

pre-id-default-policy

IN THIS SECTION

- Syntax | 12996
- Hierarchy Level | 12996
- Description | 12996
- Options | 12997
- Required Privilege Level | 12998
- Release Information | 12998

Syntax

```
pre-id-default-policy {
    then {
        log {
            session-init;
            session-close;
        }
        session-timeout {
            icmp seconds;
            icmp6 seconds;
            ospf seconds;
            others seconds;
            tcp seconds;
            udp seconds;
        }
    }
}
```

Hierarchy Level

```
[edit security policies]
```

Description

Configure pre-ID default policy settings.

Pre-ID default policy denies the flow before performing application identification (AppID) when there are no potential policies to permit the flow.

When the device receives the first packet of a traffic flow, it performs basic 5-tuple matching and checks the defined potential policies to determine how to treat the packet. If all potential policies have action as `deny`, and the default policy action is also set to `deny`, then the device denies the traffic and does not perform application identification.

If any policy has action as other than `deny` then the device performs DPI to identify the application.

The device checks for potential policies on both zone context and global context.

In the Junos OS Release prior to 23.4R1, the pre-ID default policy temporarily allows the session to get created so that DPI can get the packet and perform application identification (AppID).

> **CAUTION**: Configuring `session-init` logging for the `pre-id-default-policy` can generate a large amount of logs. Each session that enters the SRX that initially matches the `pre-id-default-policy` will generate an event. We recommend only using this option for troubleshooting purposes.

Options

then Specifies the policy action that has to be taken when the packet matches the criteria.

log Specifies the log details at session close time and session initialization time.

- Values:
  - session-init—Log at the beginning of a session
  - session-close—Log at the closure of a session

> **NOTE**: In recent versions of Junos OS, the factory-default configuration of an SRX includes the `session-close` knob. If not already present, we recommend that customers implement `session-close logging` within their `pre-id-default-policy`. This will ensure security logs are generated by the SRX if a flow is unable to leave the `pre-id-default-policy`. These events are generally a result of JDPI being unable to properly classify traffic, although they may also indicate potential attempts at evading the APPID engine.

session-timeout When you update a session, the session timeout is configured, which specifies the session timeout details in seconds.

- **Values:** icmp—Timeout value for ICMP sessions (seconds)
- **Range:** 4 through 86,400
- **Values:** icmp6—Timeout value for ICMP6 sessions (seconds)
- **Range:** 4 through 86,400
- **Values:** ospf—Timeout value for OSPF sessions (seconds)
- **Range:** 4 through 86,400
- **Values:** others—Timeout value for other sessions (seconds)

- **Range:** 4 through 86,400
- **Values:** tcp—Timeout value for TCP sessions (seconds)
- **Range:** 4 through 86,400
- **Values:** udp—Timeout value for UDP sessions (seconds)
- **Range:** 4 through 86,400

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Understanding Advanced Policy-Based Routing

pre-ietf-mode

IN THIS SECTION

- Syntax | 12999
- Hierarchy Level | 12999
- Description | 12999
- Required Privilege Level | 12999
- Release Information | 12999

Syntax

```
pre-ietf-mode
```

Hierarchy Level

```
[edit protocols ancp],
[edit protocols ancp neighbor ip-address]
```

Description

Configure the ANCP agent to run in a mode that is backward compatible with Internet draft draft-wadhwa-gsmp-l2control-configuration-00.txt, *GSMP extensions for layer2 control (L2C)* for all neighbors or for a specific neighbor.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent

Configuring the ANCP Agent for Backward Compatibility

Configuring ANCP Neighbors

pre-rewrite-tos

IN THIS SECTION

- Syntax | 13000
- Hierarchy Level | 13000
- Description | 13000
- Required Privilege Level | 13000
- Release Information | 13000

Syntax

```
pre-rewrite-tos;
```

Hierarchy Level

```
[edit forwarding-options sampling]
```

Description

Preserve prenormalized type-of-service (ToS) value for egress sampled or mirrored packets. This configuration preserves the prerewrite ToS value for all forms of sampling, such as Routing Engine-based sampling, port mirroring, flow monitoring, and so on. This statement is effective for egress sampling only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

pre-shared-key

IN THIS SECTION

- Syntax | 13001
- Hierarchy Level | 13001
- Description | 13001
- Default | 13002
- Options | 13002
- Required Privilege Level | 13002
- Release Information | 13002

Syntax

```
pre-shared-key {
    cak hexadecimal-number;
    ckn hexadecimal-number;
}
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Specifies the pre-shared key used to enable MACsec using static connectivity association key (CAK) security mode.

A pre-shared key includes a connectivity association key name (CKN) and a connectivity association key (CAK). A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using static CAK security mode. The MACsec Key Agreement (MKA) protocol is enabled after the pre-shared keys are successfully verified and exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link.

Default

No pre-shared keys exist, by default.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced on SRX Series Firewalls in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

pre-shared-key (MX Series)

IN THIS SECTION

- Syntax | 13003
- Hierarchy Level | 13003
- Description | 13003
- Default | 13004
- Options | 13004
- Required Privilege Level | 13004
- Release Information | 13004

Syntax

```
pre-shared-key {
    cak hexadecimal-number;
    ckn hexadecimal-number;
}
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Specifies the pre-shared key used to enable MACsec using static connectivity association key (CAK) security mode.

A pre-shared key includes a connectivity association key name (CKN) and a connectivity association key (CAK). A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using static CAK security mode. The MACsec Key Agreement (MKA) protocol is enabled after the pre-shared keys are successfully verified and exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link.

Default

No pre-shared keys exist, by default.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

pre-shared-key (Security)

IN THIS SECTION

- Syntax | 13005
- Hierarchy Level | 13005
- Description | 13005
- Options | 13005
- Required Privilege Level | 13005
- Release Information | 13005

Syntax

```
pre-shared-key (ascii-text key | hexadecimal key);
```

Hierarchy Level

```
[edit security ike policy ike-peer-address]
```

Description

Configure the key used to authenticate a dynamic peer during IKE phase 1 negotiation. Specify the key in either ASCII or hexadecimal format.

Options

`ascii-text` *key*—Authentication key in ASCII format.

`hexadecimal` *key*—Authentication key in hexadecimal format.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring an IKE Policy for Preshared Keys*

pre-shared-key (Services IKE)

IN THIS SECTION

- Syntax | 13006
- Hierarchy Level | 13006
- Description | 13006
- Options | 13006
- Required Privilege Level | 13006
- Release Information | 13007

Syntax

```
pre-shared-key (ascii-text key | hexadecimal key);
```

Hierarchy Level

```
[edit services ike policy policy-name]
```

Description

Define a preshared key for an IKE policy.

Options

key—Value of preshared key. The key can be one of the following:

- `ascii-text`—ASCII text key.
- `hexadecimal`—Hexadecimal key.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IKE Policies

preauthentication-order (Access Profile)

IN THIS SECTION

- Syntax | 13007
- Hierarchy Level | 13007
- Description | 13008
- Options | 13008
- Required Privilege Level | 13008
- Release Information | 13008

Syntax

```
preauthentication-order [ preauthentication-method ];
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Set the order in which the Junos OS uses preauthentication methods for the LLID service when multiple methods are configured. Junos OS supports only the `radius` method.

> **NOTE**: You cannot configure this statement if the Calling-Station-ID attribute is excluded from RADIUS Access-Request messages by the `exclude` statement.

Options

`preauthentication-method`

- `radius`—Verify the client using RADIUS.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

RADIUS Logical Line Identifier (LLID) Overview

Configuring Logical Line Identification (LLID) Preauthentication

precedence

IN THIS SECTION

- Syntax | 13009
- Hierarchy Level | 13009

- Description | 13009
- Options | 13009
- Required Privilege Level | 13010
- Release Information | 13010

Syntax

```
precedence precedence;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family filter input filter-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family filter output filter-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number filter input
filter-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number filter
output filter-name],
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family filter
input filter-name],
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family filter
output filter-name],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family
filter input filter-name],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family  family
filter output filter-name]
```

Description

Apply a precedence to a dynamic filter.

Options

precedence—Precedence value for the filter. The lower the precedence value, the higher the precedence.

- **Range:** 0 through 250
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

The `[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family inet filter input filter-name]` hierarchy level and `[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family inet filter output filter-name]` hierarchy level introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

- Firewall Filters Overview
- *Understanding Dynamic Firewall Filters*
- *Classic Filters Overview*
- *Fast Update Filters Overview*
- *Basic Classic Filter Syntax*
- *Basic Fast Update Filter Syntax*

precision-timers

IN THIS SECTION

- Syntax | 13011
- Hierarchy Level | 13011
- Description | 13011
- Required Privilege Level | 13012

Release Information | 13012

Syntax

```
precision-timers;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit protocols bgp]
```

Description

Enable BGP sessions to send frequent keepalive messages with a hold time as short as 10 seconds.

NOTE: The hold time is three times the interval at which keepalive messages are sent, and the hold time is the maximum number of seconds allowed to elapse between successive keepalive messages that BGP receives from a peer. When establishing a BGP connection with the local routing device, a peer sends an open message, which contains a hold-time value. BGP on the local routing device uses the smaller of either the local hold-time value or the peer's hold-time value as the hold time for the BGP connection between the two peers.

The default hold-time is 90 seconds, meaning that the default frequency for keepalive messages is 30 seconds. More frequent keepalive messages and shorter hold times might be desirable in large-scale deployments with many active sessions (such as edge or large VPN deployments). To configure the hold time and the frequency of keepalive messages, include the `hold-time` statement at the `[edit protocols bgp]` hierarchy level. You can configure the hold time at a logical system, routing instance, global, group, or neighbor level. When you set a hold time value to less than 20 seconds, we recommend that you also configure the BGP `precision-timers` statement. The `precision-timers` statement ensures that if scheduler slip messages occur, the routing device continues to send keepalive messages. When the `precision-timers` statement is included, keepalive message generation is performed in a dedicated kernel thread, which helps to prevent BGP session flaps.

> **NOTE**: Starting with Junos OS Release 15.2, you can register or unregister keepalives of BGP with the automated keepalive precision timer service of the kernel. This service ensures a reliable generation of keepalives for some configurable maximum period after a switchover of the routing engine from backup to primary until BGP is able to take over the keepalive generation.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

hold-time (Protocols BGP) | **8104**

precision-timers-max-period

IN THIS SECTION

- Syntax | **13013**
- Hierarchy Level | **13013**
- Description | **13013**
- Required Privilege Level | **13013**
- Release Information | **13013**

Syntax

```
precision-timers-max-period precision-timers-max-period;
```

Hierarchy Level

```
[edit routing-options nonstop-routing-options]
```

Description

Support of precision-timers in the kernel is a feature where the kernel takes over auto-generation of BGP keepalives right after the switchover from standby to primary event occurs. The kernel in the RE continues this auto-generation until the BGP protocol is able to take over the session or until a maximum period has elapsed since the switchover event occurred. The maximum period for which the kernel auto-generates keepalives on behalf of BGP after a switchover event from standby to primary ranges from 60 seconds to 1800 seconds. The default value is 600 seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

predefined-attack-groups

IN THIS SECTION

- Syntax | 13014
- Hierarchy Level | 13014
- Description | 13014

- Options | 13014
- Required Privilege Level | 13014
- Release Information | 13014

Syntax

```
predefined-attack-groups [attack-group-name];
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match attacks],
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match attacks]
```

Description

Specify predefined attack groups that you can use to match the traffic against known attack objects. You can update only the list of attack objects.

Options

attack-name —Name of the predefined attack object group.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

predefined-attacks

IN THIS SECTION

- Syntax | 13015
- Hierarchy Level | 13015
- Description | 13015
- Options | 13015
- Required Privilege Level | 13016
- Release Information | 13016

Syntax

```
predefined-attacks [attack-name];
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match attacks],
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match attacks]
```

Description

Specify predefined attack objects that you can use to match the traffic against known attacks. You can update only the list of attack objects.

Options

attack-name—Name of the predefined attack objects.

You can view the IPS Policy in predefined attacks and attack groups under IPS/Exempt rule in tenant systems. Use the commands, `show security idp attack attack-group-entries`, `show security idp attack attack-list`, and `show security idp attack group-list` to view the IPS policies for logical system and tenant.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

predefined-variable-defaults (Dynamic Client Profiles)

IN THIS SECTION

- Syntax | 13016
- Hierarchy Level | 13017
- Description | 13017
- Options | 13017
- Required Privilege Level | 13019
- Release Information | 13019

Syntax

```
predefined-variable-defaults {
    cos-excess-rate <percent percentage> <proportion number>;
    cos-excess-rate-low <percent percentage> <proportion number>;
    cos-excess-rate-high <percent percentage> <proportion number>;
    cos-scheduler-bs <percent percentage> <temporal microseconds>;
    cos-scheduler-shaping-rate <percent percentage> <rate bps>;
    cos-scheduler-tx <percent percentage> <rate bps>;
    predefined-variable default-value;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Configure default values for the predefined variables that are configured in a dynamic client profile. These default values are used when RADIUS does not supply a value. Omit the predefined variable prefix, $junos-, when you configure a default. Most predefined variables support only a single default value; these predefined variables are not listed separately here. The listed predefined variables support default values for two attributes; for these you can specify the default value for either attribute or for both attributes.

NOTE: Not all predefined variables support default values. For a list of predefined variables and options for which you can configure default values, see *Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs*

Defining default values for these predefined variables enables you to determine whether to source values locally from the profile instead of only from RADIUS. This enables you to use RADIUS as a way to selectively override predefined variable values, instead of being the sole source of those values.

NOTE: The `proportion` option provides greater granularity than the `percent` option, enabling you to specify the equivalent of a tenth of a percent. Configuring `proportion 121` applies a value equivalent to 12.1 percent, which you cannot achieve with the `percent` option.

NOTE: Do not use this statement in a dynamic service profile.

Options

cos-excess-rate percent *percentage*

Specify the percentage of excess bandwidth to share from all traffic.

- **Range:** 1 through 100

cos-excess-rate proportion *number*

Specify the proportion of excess bandwidth to share from all traffic.

- **Range:** 0 through 1000

cos-excess-rate-low percent ***percentage***

Specify the percentage of excess bandwidth to share from low-priority traffic.

- **Range:** 1 through 100

cos-excess-rate-low proportion ***number***

Specify the proportion of excess bandwidth to share from low-priority traffic.

- **Range:** 0 through 1000

cos-excess-rate-high percent ***percentage***

Specify the percentage of excess bandwidth to share from high-priority traffic.

- **Range:** 1 through 100

cos-excess-rate-high proportion ***number***

Specify the proportion of excess bandwidth to share from high-priority traffic.

- **Range:** 0 through 1000

cos-scheduler-bs percent ***percentage***

Specify the buffer size as a percentage of the total buffer.

- **Range:** 0 through 100

cos-scheduler-bs temporal ***microseconds***

Specify the temporal value that determines the buffer size. The temporal value multiplied by the logical interface speed determines the size.

- **Range:** 0 through 4294967295

cos-scheduler-shaping-rate percent ***number***

Specify the shaping rate as a percentage of the available interface bandwidth.

- **Range:** 1 through 100

cos-scheduler-shaping-rate rate ***bps***

Specify the shaping rate in bits per second.

- **Range:** 3200 through 160,000,000,000

cos-scheduler-tx percent ***number***

Specify the transmit rate as a percentage of the total available bandwidth.

- **Range:** 0 through 100

cos-scheduler-tx rate ***bps***

Specify the transmit rate in bits per second.

- **Range:** 3200 through 160,000,000,000

`default-value`—Default value that you want to assign to the predefined variable. Use this for predefined variables that support only a single value.

predefined-variable—Name of the Junos OS predefined variable to which you want to assign a default value. Do not include the `$junos-` prefix.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Support for the $junos-interface-mtu predefined variable (`interface-mtu`) introduced in Junos OS Release 18.1R1 on MX Series.

RELATED DOCUMENTATION

Configuring Default Values for Predefined Variables in a Dynamic Profile

Per-Subscriber Support of Maximum Transmission Unit for Dynamic Profiles

preempt

IN THIS SECTION

- Syntax | 13020
- Hierarchy Level | 13020
- Description | 13020
- Required Privilege Level | 13020
- Release Information | 13020

Syntax

```
(preempt | no-preempt) {
    hold-time seconds;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
```

Description

Configure whether a backup router can preempt a primary router:

- **preempt**—Allow the primary router to be preempted.
- **no-preempt**—Prohibit the preemption of the primary router.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring VRRP for IPv6 (CLI Procedure)

preempt (Chassis Cluster)

IN THIS SECTION

- Syntax | 13021
- Hierarchy Level | 13021
- Description | 13021
- Options | 13022
- Required Privilege Level | 13022
- Release Information | 13022

Syntax

```
preempt {
    delay seconds;
    limit limit;
    period seconds;
}
```

Hierarchy Level

```
[edit chassis cluster redundancy-group group-number]
```

Description

Allow preemption of the primary system based on the priority within a redundancy group.

Configuring preemption for RG0 is not allowed and is blocked when committed.

By configuring the preemptive delay timer and failover limit, you can limit the flapping of the redundancy group state between the secondary and the primary in a preemptive failover.

By default, preemption is disabled.

Example: Consider the following scenario where you have configured a preemptive period as 300 seconds and preemptive limit as 50.

When the preemptive limit is configured as 50, the count starts at 0 and increments with a first preemptive failover; this process continues until the count reaches the configured preemptive limit, that is 50, before the preemptive period expires. When the preemptive limit (50) is exceeded, you must manually reset the preempt count to allow preemptive failovers to occur again.

When you have configured the preemptive period as 300 seconds, and if the time difference between the first preemptive failover and the current failover has already exceeded 300 seconds, and the preemptive limit (50) is not yet reached, then the preemptive period will be reset. After resetting, the last failover is considered as the first preemptive failover of the new preemptive period and the process starts all over again.

Options

delay Time to wait before the node in secondary state transitions to primary state in a preemptive failover.

- **Range:** 1 to 21,600 seconds
- **Default:** 1

limit Maximum number of preemptive failovers allowed in a configured preemptive period.

- **Range:** 1 to 50

period Time period during which the preemptive limit is applied.

- **Range:** 1 to 1400 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0. Support for `delay`, `limit`, and `period` options are added in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

redundancy-group (Chassis Cluster) | **14173**

preempt-cutover-timer

IN THIS SECTION

- Syntax | 13023
- Hierarchy Level | 13023
- Description | 13024
- Default | 13024
- Options | 13024
- Required Privilege Level | 13024
- Release Information | 13024

Syntax

```
preempt-cutover-timer seconds;
```

Hierarchy Level

- For platforms with ELS:

```
[edit switch-options redundant-trunk-group group name]
[edit interfaces name aggregated-ether-options lacp link-protection rtg-config]
[edit interfaces name aggregated-ether-options link-protection rtg-config]
```

- For platforms without ELS:

```
[edit ethernet-switching-optionsredundant-trunk-group group name]
```

Description

Change the length of time that a re-enabled primary link waits to take over from an active secondary link in a redundant trunk group (RTG).

Default

If you do not change the time with the `preempt-cutover-timer` statement, a re-enabled primary link takes over from the active secondary link after 1 second.

Options

seconds—Number of seconds that the primary link waits to take over from the active secondary link.

- **Range:** 1 through 600 seconds

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Hierarchy level `[edit switch-options]` introduced in Junos OS Release 13.2X50-D10 (ELS). (See *Using the Enhanced Layer 2 Software CLI* for information about ELS.)

RELATED DOCUMENTATION

> *Example: Configuring Redundant Trunk Links for Faster Recovery on EX Series Switches*
>
> *Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support*
>
> Q-in-Q Support on Redundant Trunk Links Using LAGs with Link Protection

preempt (VRRP)

IN THIS SECTION

- Syntax | 13025
- Hierarchy Level | 13025
- Description | 13025
- Default | 13026
- Required Privilege Level | 13026
- Release Information | 13026

Syntax

```
(preempt | no-preempt) {
    hold-time seconds;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
```

Description

In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a backup router can preempt a primary router:

- `preempt`—Allow the primary router to be preempted.

> **NOTE:** By default, a higher-priority backup router can preempt a lower-priority primary router.

- `no-preempt`—Prohibit the preemption of the primary router. When `no-preempt` is configured, the backup router cannot preempt the primary router even if the backup router has a higher priority.

The remaining statement is explained separately. See CLI Explorer.

Default

By default the `preempt` statement is enabled, and a higher-priority backup router preempts a lower-priority primary router even if the `preempt` statement is not explicitly configured.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Backup Router to Preempt the VRRP Primary Router

Configuring VRRP Preemption and Hold Time

preemption

IN THIS SECTION

- Syntax | 13027
- Hierarchy Level | 13027
- Description | 13027

- Default | 13027
- Options | 13027
- Required Privilege Level | 13028
- Release Information | 13028

Syntax

```
preemption {
    (aggressive | disabled | normal);
    soft-preemption {
        cleanup-timer seconds;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Control RSVP session preemption.

Default

`normal`

Options

`aggressive`—Preempt RSVP sessions whenever bandwidth is insufficient to handle all sessions. A session is preempted whenever bandwidth is lowered or a new higher-priority session is established.

`disabled`—Do not preempt RSVP sessions.

`normal`—Preempt RSVP sessions to accommodate new higher-priority sessions when bandwidth is insufficient to handle all sessions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Preempting RSVP Sessions

preemptive-fast-reroute

IN THIS SECTION

- Syntax | 13029
- Hierarchy Level | 13029
- Description | 13029
- Default | 13029
- Options | 13029
- Required Privilege Level | 13029
- Release Information | 13029

Syntax

```
preemptive-fast-reroute {
    (backward-frr-enable | no-backward-frr-enable);
    (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
    (odu-backward-frr-enable | no-odu-backward-frr-enable);
    (odu-signal-degrade-monitor-enable |no-odu-signal-degrade-monitor-enable);
}
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Enable or disable preemptive fast reroute options.

Default

By default, backward fast reroute insertion and signal degradation monitoring are disabled.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

10-Gigabit Ethernet OTN Options Configuration Overview

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

prefer-framed-ip-address (RADIUS Clients)

IN THIS SECTION

- Syntax | 13030
- Hierarchy Level | 13030
- Description | 13030
- Required Privilege Level | 13031
- Release Information | 13031

Syntax

```
prefer-framed-ip-address;
```

Hierarchy Level

```
[edit access radius clients client-name]
```

Description

Specify that the framed-ip-address is used for subscriber creation when both the framed-route and framed-ip-address attributes are in the RADIUS accounting request from the RADIUS client. The framed-ip-netmask is also used for subscriber creation if it is in the request.

By default, the framed-route attribute is used for subscriber creation when both the framed-route and framed-ip-address attributes are in the RADIUS accounting request.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers

prefer-framed-ipv6-prefix (RADIUS Clients)

IN THIS SECTION

- Syntax | 13031
- Hierarchy Level | 13031
- Description | 13032
- Required Privilege Level | 13032
- Release Information | 13032

Syntax

```
prefer-framed-ipv6-prefix;
```

Hierarchy Level

```
[edit access radius clients client-name]
```

Description

Specify that the framed-ipv6-prefix is used for subscriber creation when both the delegated-ipv6-prefix and framed-ipv6-prefix attributes are in the RADIUS accounting request from the RADIUS client.

By default, the delegated-ipv6-prefix attribute is used for subscriber creation when both the delegated-ipv6-prefix and framed-ipv6-prefix attributes are in the RADIUS accounting request.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers

preference (Access)

IN THIS SECTION

- Syntax | 13033
- Hierarchy Level | 13033
- Description | 13033
- Options | 13033
- Required Privilege Level | 13033
- Release Information | 13033

Syntax

```
preference route-distance;
```

Hierarchy Level

```
[edit routing-options access route ip-prefix</prefix-length>]
```

Description

Configure the distance for an access route.

Options

route-distance—Specific distance you want to assign to the access route.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

preference (Interface-Level Preference for VPLS Multihoming for FEC 129)

IN THIS SECTION

- Syntax | 13034
- Hierarchy Level | 13034
- Description | 13034

- Options | 13034
- Required Privilege Level | 13035
- Release Information | 13035

Syntax

```
preference preference-value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-
homing site site-name interface interface-name],
[edit routing-instances instance-name protocols vpls multi-homing site site-name interface
interface-name]
```

Description

Specify a preference for the interface to become the designated forwarder (DF) for a multihomed VPLS site. This `preference` statement can be useful when you want the interface preference for a site to change dynamically so that the DF election can be influenced depending on the interface state. Among the list of interface preferences, Junos OS advertises the best preference as the VPLS site's preference value. For example, if the site has three interfaces configured with preference values 12, 10, and 9, respectively, 12 is advertised as the site preference. If that interface goes down, 10 is advertised as the site preference.

If you configure interface-level preference, you cannot configure site-level preference.

Options

preference-value

Preference value for the interface.

- **Range:** 1 through 65535

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring VPLS Multihoming (FEC 129)

preference (IPv6 Router Advertisement)

IN THIS SECTION

- Syntax | 13035
- Hierarchy Level | 13036
- Description | 13036
- Options | 13036
- Required Privilege Level | 13036
- Release Information | 13036

Syntax

```
preference (high | low | medium);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Description

Specify the router preference that is communicated to IPv6 hosts through router advertisements. The `preference` value in the router advertisements enables IPv6 hosts to select a default router to reach a remote destination.

The `preference` can be configured when there are multiple devices that route to distinct sets of prefixes and where one of the devices would lead to considerably fewer redirects. You can indicate a lower preference for a new device that is not completely configured yet, so that hosts do not adopt this new device as the default device and thus avoid traffic loss.

Options

You can specify different levels of preference depending on your requirements:

high Specify a high preference for a device.

low Specify a low preference for a device.

medium Specify a medium preference for a device. This is the default preference.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

preference (PEG DF Election)

IN THIS SECTION

- Syntax | 13037
- Hierarchy Level | 13037
- Description | 13037
- Options | 13037
- Required Privilege Level | 13038
- Release Information | 13038

Syntax

```
preference {
    use-least;
    value value;
}
```

Hierarchy Level

```
[edit routing-instances name protocols evpn oism pim-evpn-gateway peg-df-election]
```

Description

Use preference-based PIM EVPN gateway (PEG) designated forwarder (DF) election with optimized intersubnet multicast (OISM) for a virtual routing and forwarding (VRF) instance in an EVPN fabric.

Specify the preference value and (optionally) the preference order that peer PEG DF candidates use to elect the PEG DF.

Options

use-least Elect the OISM PEG DF based on the lowest preference value.

By default, the DF election preference-based algorithm uses the highest preference value.

value *value* Preference value for OISM PEG DF election.

- **Default:** 32767
- **Range:** 0 through 65535

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

PEG DF Election

preference (Protocols BGP)

IN THIS SECTION

- Syntax | 13039
- Hierarchy Level | 13039
- Description | 13039
- Options | 13039
- Required Privilege Level | 13040
- Release Information | 13040

Syntax

```
preference preference;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Specify the preference for routes learned from BGP.

At the BGP global level, the preference statement sets the preference for routes learned from BGP. You can override this preference in a BGP group or peer preference statement.

At the group or peer level, the preference statement sets the preference for routes learned from the group or peer. Use this statement to override the preference set in the BGP global preference statement when you want to favor routes from one group or peer over those of another.

NOTE: Do not set `preference2` for BGP route-policy.

Options

preference—Preference to assign to routes learned from BGP or from the group or peer.

- **Range:** 0 through 4,294,967,295 ($2^{32} - 1$)

- **Default:** 170 for the primary preference

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

local-preference | 9971

Example: Configuring the Preference Value for BGP Routes

preference (Protocols LDP)

IN THIS SECTION

- Syntax | 13040
- Hierarchy Level | 13041
- Description | 13041
- Options | 13041
- Required Privilege Level | 13041
- Release Information | 13041

Syntax

```
preference preference;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit protocols ldp interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit protocols ldp interface interface-name],
[edit routing-instances routing-instance-name protocols ldp interface interface-name]
```

Description

Set the route preference level for LDP routes.

Options

preference—Preferred value.

- **Range:** 0 through 255
- **Default:** 9

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring LDP Route Preferences

preference (Protocols MPLS)

IN THIS SECTION

- Syntax | 13042
- Hierarchy Level | 13042
- Description | 13042
- Options | 13043
- Required Privilege Level | 13043
- Release Information | 13043

Syntax

```
preference preference;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name],
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Description

Preference for the route.

You can optionally configure multiple LSPs between the same pair of ingress and egress routers. This is useful for balancing the load among the LSPs because all LSPs, by default, have the same preference

level. To prefer one LSP over another, set different preference levels for individual LSPs. The LSP with the lowest preference value is used. The default preference for LSPs is lower (more preferred) than all learned routes except direct interface routes.

Options

preference—Preference to assign to the route. A route with a lower preference value is preferred.

- **Range:** 1 through 255
- **Default:** 5 for static MPLS LSPs, 7 for RSVP MPLS LSPs, 9 for LDP MPLS LSPs

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Preference Values for LSPs

Configuring the Ingress Router for Static LSPs

Configuring the Intermediate (Transit) and Egress Routers for Static LSPs

preference (Protocols RIP)

IN THIS SECTION

- Syntax | 13044
- Hierarchy Level | 13044
- Description | 13044
- Options | 13044

Required Privilege Level | 13044
Release Information | 13045

Syntax

```
preference preference;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name],
[edit protocols rip group group-name],
[edit routing-instances routing-instance-name protocols rip group group-name]
```

Description

Specify the preference of external routes learned by RIP as compared to those learned from other routing protocols.

By default, Junos OS assigns a preference of 100 to routes that originate from RIP. When Junos OS determines a route's preference to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.

Options

preference—Preference value. A lower value indicates a more preferred route.

- **Range:** 0 through 4,294,967,295 (2^{32} – 1)
- **Default:** 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Route Preferences Overview

preference (Protocols RIPng)

IN THIS SECTION

- Syntax | 13045
- Hierarchy Level | 13045
- Description | 13046
- Options | 13046
- Required Privilege Level | 13046
- Release Information | 13046

Syntax

```
preference preference;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng group group-name],
[edit protocols ripng group group-name],
[edit routing-instances routing-instance-name protocols ripng group group-name]
```

Description

Specify the preference of external routes learned by RIPng as compared to those learned from other routing protocols.

By default, Junos OS assigns a preference of 100 to routes that originate from RIPng. When Junos OS determines that a route is to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.

To modify the default RIPng preference value, include the `preference` statement.

Options

preference—Preference value. A lower value indicates a more preferred route.

- **Range:** 0 through 4,294,967,295 ($2^{32} - 1$)
- **Default:** 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

preference (Routing Options)

IN THIS SECTION

- Syntax | 13047
- Hierarchy Level | 13047
- Description | 13047
- Options | 13048

Required Privilege Level | 13048
Release Information | 13048

Syntax

```
preference {
    metric-value;
    <type metric_type>
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options (aggregate | generate | static) (defaults | route)],

[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name (aggregate | generate | static) (defaults | route)],

[edit logical-systems logical-system-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-instances routing-instance-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib routing-table-name (aggregate | generate | static) (defaults | route)]
```

Description

Preference value for a static, aggregate, or generated route. You also can specify a secondary preference value, as well as color values, which are even finer-grained preference values.

You can specify a primary route preference (by including the **preference** statement in the configuration), and a secondary preference that is used as a tiebreaker (by including the **preference2** statement). You

can also mark route preferences with additional route tiebreaker information by specifying a color and a tiebreaker color (by including the **color** and **color2** statements in the configuration).

If the Junos OS routing table contains a dynamic route to a destination that has a better (lower) preference value than the static, aggregate, or generated route, the dynamic route is chosen as the active route and is installed in the forwarding table.

Options

metric_value—The metric value for an aggregate, a generated, or a static route to determine the best route among multiple routes to a destination

- **Range:** 0 through 4,294,967,295 (2^{32} – 1)
- **Default:** 5 (for static routes), 130 (for aggregate and generated routes)

type ***metric_type***—(Optional) External metric type for routes exported by OSPF. When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric.

- **Range:** 1 through 16

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation

aggregate (Routing) | **2070**

generate | **7672**

static (Routing Options) | **28570**

color | **4425**

preference (Site-Level Preference for VPLS Multihoming for FEC 129)

IN THIS SECTION

- Syntax | 13049
- Hierarchy Level | 13049
- Description | 13049
- Default | 13049
- Options | 13050
- Required Privilege Level | 13050
- Release Information | 13050

Syntax

```
preference (preference-value | backup | primary);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-
homing site site-name],
[edit routing-instances instance-name protocols vpls multi-homing site site-name]
```

Description

Influence the designated forwarder (DF) selection for a multihomed VPLS site. Configure the preference in terms of keywords `primary` and `backup`, or configure the preference value explicitly.

Default

If this statement is omitted, the default preference value for the site is 100.

Options

preference-value	Preference value for the DF. • **Range:** 1 through 65535
backup	Less likely to become the DF.
primary	Most likely to become the DF.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring VPLS Multihoming (FEC 129)

preference (Subscriber Management)

IN THIS SECTION

- Syntax | 13051
- Hierarchy Level | 13051
- Description | 13051
- Options | 13051
- Required Privilege Level | 13051
- Release Information | 13051

Syntax

```
preference route-distance;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
access route prefix],
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
rib routing-table-name access route prefix],
[edit dynamic-profiles profile-name routing-options access route prefix]
```

Description

Dynamically configure the distance for an access route.

Options

route-distance—Either the specific distance you want to assign to the access route or either of the following distance variables:

- `$junos-framed-route-distance`—Distance of an IPv4 access route; the variable is dynamically replaced with the preference value (Subattribute 5) from the RADIUS Framed-Route attribute [22].
- `$junos-framed-route-ipv6-distance`—Distance of an IPv6 access route; the variable is dynamically replaced with the preference value (Subattribute 5) from the RADIUS Framed-IPv6-Route attribute [99].

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at `[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options access route prefix]` and `[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options rib routing-table-name access route prefix]` hierarchy levels introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring Dynamic Access Routes for Subscriber Management

preference (Tunnel Profile)

IN THIS SECTION

- Syntax | 13052
- Hierarchy Level | 13052
- Description | 13052
- Options | 13053
- Required Privilege Level | 13053
- Release Information | 13053

Syntax

```
preference number;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Description

Specify the preference for a tunnel. You can specify up to 8 levels of preference, and you can assign the same preference to a maximum of 31 tunnels. When you define multiple preferences for a destination, you increase the probability of a successful connection.

This value can be overridden by RADIUS attribute Tunnel-Preference [83].

Options

number—Number that indicates the order in which the router attempts to connect to the destination. Zero is the highest level of preference.

- **Range:** 0 through 2000
- **Default:** 2000

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

preferred

IN THIS SECTION

- Syntax | 13054
- Hierarchy Level | 13054
- Description | 13054
- Default | 13054
- Required Privilege Level | 13054
- Release Information | 13054

Syntax

```
preferred;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family address address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family  family address address]
```

Description

Configure this address to be the preferred address on the interface. If you configure more than one address on the same subnet, the preferred source address is chosen by default as the source address when you initiate frame transfers to destinations on the subnet.

NOTE: The `edit logical-systems` hierarchy is not available on QFabric systems.

Default

The lowest-numbered address on the subnet is the preferred address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Interface Address

preferred-ciphers

IN THIS SECTION

- Syntax | 13055
- Hierarchy Level | 13055
- Description | 13055
- Options | 13055
- Required Privilege Level | 13056
- Release Information | 13056

Syntax

```
preferred-ciphers (custom | medium | strong | weak);
```

Hierarchy Level

```
[edit services ssl proxy profile profile-name ]
[edit services ssl termination profile profile-name ]
[edit services ssl initiation profile profile-name]
```

Description

Select preferred ciphers. Preferred ciphers allow you to define an SSL cipher that can be used with acceptable key strength. Ciphers are divided in three categories depending on their key strength: strong, medium, or weak.

Custom ciphers allow you to define your own cipher list. If you do not want to use one of the three categories, you can select ciphers from each of the categories to form a custom cipher set. To configure custom ciphers, you must set `preferred-ciphers` to `custom`.

Options

- `custom`—Configure custom cipher suite and order of preference.

- `medium`—Use ciphers with key strength of 128 bits or greater.
- `strong`—Use ciphers with key strength of 168 bits or greater.
- `weak`—Use ciphers with key strength of 40 bits or greater.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Firewall User Authentication Overview

SSL Proxy Overview

preferred-lifetime

IN THIS SECTION

- Syntax | 13057
- Hierarchy Level | 13057
- Description | 13057
- Options | 13057
- Required Privilege Level | 13057
- Release Information | 13057

Syntax

```
preferred-lifetime seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name prefix prefix],
[edit protocols router-advertisement interface interface-name prefix prefix]
```

Description

Specify how long the prefix generated by stateless autoconfiguration remains preferred.

Options

seconds—Preferred lifetime, in seconds. If you set the preferred lifetime to **0xffffffff**, the lifetime is infinite. The preferred lifetime is never greater than the valid lifetime.

- **Default:** 604,800 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

valid-lifetime | **30644**

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

Secure IPv6 Neighbor Discovery

preferred-lifetime (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 13058
- Hierarchy Level | 13058
- Description | 13058
- Options | 13058
- Required Privilege Level | 13058
- Release Information | 13059

Syntax

```
preferred-lifetime seconds;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols router-advertisement interface interface-nameprefix prefix]
```

Description

Specify how long the prefix generated by stateless autoconfiguration remains preferred.

Options

seconds—Preferred lifetime, in seconds. If you set the preferred lifetime to `0xffffffff`, the lifetime is infinite. The preferred lifetime is never greater than the valid lifetime.

- **Default:** 604,800 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA

Dynamic Router Advertisement Configuration Overview

preferred-prefix-length

IN THIS SECTION

- Syntax | 13059
- Hierarchy Level | 13059
- Description | 13060
- Required Privilege Level | 13060
- Release Information | 13060

Syntax

```
preferred-prefix-length preferred-prefix-length;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcpv6-client prefix-
delegating]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet6 dhcpv6-client prefix-delegating]
```

```
[edit tenants tenant-name interfaces interface-name unit logical-unit-number family inet6 dhcpv6-
client prefix-delegating]
```

Description

Allows you to configure DHCPv6 client preferred prefix length. If preferred-prefix-length is configured, the DHCPv6 client checks the prefix length in the ADVERTISE packet and if the check fails, a sysolg is created.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30 and in Junos OS Release 15.1X49-D100.

The `logical-systems` and `tenants` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

sub-prefix-length | 28674

preferred-source-address

IN THIS SECTION

- Syntax | 13061
- Hierarchy Level | 13061
- Description | 13061
- Options | 13061
- Required Privilege Level | 13062
- Release Information | 13062

Syntax

```
preferred-source-address address;
```

Hierarchy Level

```
[edit dynamic-profiles interfaces interface-name unit logical-unit-number family family unnumbered-
address interface-name],
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family]
```

Description

For unnumbered Ethernet interfaces configured with a loopback interface as the donor interface, specify one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network. To configure the preferred source address dynamically, instead of using this statement, you must include the `$junos-preferred-source-address` predefined variable for IPv4 (family inet) addresses or the `$junos-preferred-source-ipv6-address` predefined variable for IPv6 (family inet6) addresses.

Configuration of a preferred source address for unnumbered Ethernet interfaces is supported for IPv4 and IPv6 address families.

> **NOTE**: When you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the `$junos-routing-instance` predefined variable, you must not configure a preferred source address, whether with the `$junos-preferred-source-address` predefined variable, the `$junos-preferred-source-ipv6-address` predefined variable, or the `preferred-source-address` statement. Configuring the preferred source address in this circumstance causes a commit failure.

Options

address—Secondary IP address of the donor loopback interface. Alternatively, use the `$junos-preferred-source-address` or the `$junos-preferred-source-ipv6-address` predefined variable to dynamically apply a preferred source address to the unnumbered Ethernet interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support for the `$junos-preferred-source-address` and `$junos-preferred-source-ipv6-address` predefined variables introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Configuring an Unnumbered Interface
- Junos OS Network Interfaces Library for Routing Devices
- Junos OS Administration Library for Routing Devices

prefix

IN THIS SECTION

- Syntax | 13062
- Hierarchy Level | 13063
- Description | 13063
- Options | 13063
- Required Privilege Level | 13063
- Release Information | 13063

Syntax

```
prefix destination-prefix;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast scope scope-name],
[edit logical-systems logical-system-name routing-options multicast scope scope-name],
[edit routing-instances routing-instance-name routing-options multicast scope scope-name],
[edit routing-options multicast scope scope-name]
```

Description

Configure the prefix for multicast scopes.

Options

destination-prefix—Address range for the multicast scope.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Examples: Configuring Administrative Scoping

Example: Creating a Named Scope for Multicast Scoping

multicast (routing-options) | **11063**

prefix-65-127-disable

IN THIS SECTION

- Syntax | 13064
- Hierarchy Level | 13064
- Description | 13064
- Options | 13066
- Required Privilege Level | 13066
- Release Information | 13066

Syntax

```
prefix-65-127-disable;
```

Hierarchy Level

```
[edit chassis (QFX Series) forwarding-options lpm-profile]
```

Description

For the Unified Forwarding Table (UFT) feature, specify not to allocate any memory for IPv6 prefixes with lengths in the range /65 through /127 for longest-prefix-match (LPM) entries. Doing so increases the memory available for LPM entries for IPv4 unicast prefixes and IPv6 prefixes with lengths equal to or less than 64. The maximum default value for LPM entries is 16,000 IPv6 prefixes of all lengths.

In an environment where the switch is being used in the core of the network, for example, it might not need to store IPv6 prefixes with lengths in the range /65 through /127. IPv6 prefixes of this type are not typically used in the core.

NOTE: When using this statement, IPv6 prefixes within the range /65 through /127 will still appear in the routing table, but will *not* be installed in the forwarding table; therefore, matching traffic will be dropped. Note further that if a default route is

configured, traffic will be forwarded, though it will be sent through the RE and rate-limited.

NOTE: On QFX5100 switches, when you configure this statement, the maximum number of LPM IPv6 entries with prefix lengths equal to or less than 64 increases to 128,000. On the QFX5200 switch, when you configure this statement, the maximum number of IPv6 entries with prefix lengths equal to or less 64 that are allocated in the LPM table increases to 98,0000.

NOTE: This statement is supported only with the `lpm-profile`. No other profile is supported.

The effects of this statement can be seen on a QFX5100 as follows:

```
[edit]
user@host# set chassis forwarding-options lpm-profile prefix-65-127-disable


[edit]
user@host# commit
configuration check succeeds
commit complete


[edit]
user@host# run show chassis forwarding-options
fpc0:
--------------------------------------------------------------------------
Current UFT Configuration:
lpm-profile. (MAC: 32K L3-host: 16K LPM: 128K)
prefix-65-127 = disable


[edit]
user@host# run show pfe route summary hw
Slot 0
================ fpc0 ================

Unit: 0
```

```
Profile active: lpm-profile
Type            Max       Used      Free      % free
----------------------------------------------------
IPv4 Host       16384     20        16354     99.82
IPv4 LPM        131072    5         131065    99.99
IPv4 Mcast      8192      0         8177      99.82

IPv6 Host       8192      5         8177      99.82
IPv6 LPM(< 64)  131072    2         131065    99.99
IPv6 LPM(> 64)  0         0         0         0.00
IPv6 Mcast      4096      0         4089      99.83
```

Options

None This statement has no options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D15.

Support introduced in Junos OS Release 15.1X53-D30 for QFX5200 Series switches.

Support introduced in Junos OS Release 18.1R1 for QFX5200-48C and QFX5210 switches.

RELATED DOCUMENTATION

Configuring the Unified Forwarding Table on Switches

Understanding the Unified Forwarding Table

prefix (Access IPv6)

IN THIS SECTION

- Syntax | 13067
- Hierarchy Level | 13067
- Description | 13067
- Options | 13067
- Required Privilege Level | 13067
- Release Information | 13068

Syntax

```
prefix IPv6-network prefix;
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet6]
```

Description

Specify the IPv6 prefix for the IPv6 address-assignment pool. This statement is mandatory for IPv6 address-assignment pools.

Options

IPv6-network-prefix—IPv6 prefix.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.4 of Junos OS.

RELATED DOCUMENTATION

Firewall User Authentication Overview

prefix-action (Configuring)

IN THIS SECTION

- Syntax | 13068
- Hierarchy Level | 13069
- Description | 13069
- Options | 13069
- Required Privilege Level | 13069
- Release Information | 13069

Syntax

```
prefix-action prefix-action-name {
    count;
    destination-prefix-length prefix-length;
    filter-specific;
    policer policer-name;
    source-prefix-length prefix-length;
    subnet-prefix-length prefix-length;
}
```

Hierarchy Level

```
[edit firewall family inet],
[edit logical-systems logical-system-name firewall family inet]
```

Description

Configure a prefix-specific action.

Options

`count`—Enable counter.

`destination-prefix-length` *`prefix-length`*—Destination prefix length.

- **Range:** 0 through 32

`filter-specific`—Create the prefix-specific set of policers and counters as a filter-specific set. If this option is not specified, the prefix-specific set of policers and counters are created as term-specific.

`policer` *`policer-name`*—Policer name.

`source-prefix-length` *`prefix-length`*—Source prefix length.

- **Range:** 0 through 32

`subnet-prefix-length` *`prefix-length`*—Subnet prefix length.

- **Range:** 0 through 32

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Prefix-Specific Counting and Policing Actions

prefix-action (Firewall Filter Action)

IN THIS SECTION

- Syntax | 13070
- Hierarchy Level | 13070
- Description | 13070
- Options | 13070
- Required Privilege Level | 13071
- Release Information | 13071

Syntax

```
prefix-action prefix-action-name;
```

Hierarchy Level

```
[edit firewall family inet  filter filter-name term term-name then],
[edit logical-systems logical-system-name firewall family inet  filter filter-name term term-name
then]
```

Description

Reference a prefix-specific action.

Options

prefix-action-name	Name of a prefix-specific action to use to rate-limit traffic.

Required Privilege Level

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Firewall Filter Nonterminating Actions

Prefix-Specific Counting and Policing Actions

prefix (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 13071
- Hierarchy Level | 13072
- Description | 13072
- Options | 13072
- Required Privilege Level | 13072
- Release Information | 13072

Syntax

```
prefix ipv6-prefix;
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet6]
```

Description

Specify the IPv6 prefix for the IPv6 address-assignment pool. This statement is mandatory for IPv6 address-assignment pools.

Options

ipv6-prefix—The IPv6 prefix.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Address-Assignment Pools Overview

Address-Assignment Pool Configuration Overview

prefix-attribute-flags

IN THIS SECTION

- Syntax | 13073
- Hierarchy Level | 13073

- Description | 13073
- Options | 13073
- Required Privilege Level | 13073
- Release Information | 13074

Syntax

```
prefix-attribute-flags [node-segment-flag | external-prefix-flag]
```

Hierarchy Level

```
[edit policy-options policy-statement <policy-name> then]
```

Description

When the prefix attribute flags CLI is configured, the originating router will include the prefix attributes sub TLV.

Options

Set when the prefix `node-segment-flag` identifies the advertising router,that is when the prefix is a host prefix advertising a globally reachable address typically associated with a loopback address.

Set if the prefix `external-prefix-flag` has been redistributed from another protocol. This includes the case where multiple IS-IS instances are configured and the source of the redistributed prefix is another IS-IS instance. The flag MUST be preserved when leaked between levels.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

prefix (Circuit ID for Option 82)

IN THIS SECTION

- Syntax | 13074
- Hierarchy Level | 13074
- Description | 13075
- Default | 13075
- Options | 13075
- Required Privilege Level | 13076
- Release Information | 13076

Syntax

```
prefix {
    host-name;
    logical-system-name;
    routing-instance-name;
}
```

Hierarchy Level

- For platforms with enhanced Layer 2 software (ELS):

```
[edit vlans forwarding-options dhcp-security option-82 circuit-id]
```

- For platforms without ELS:

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name) dhcp-option82 circuit-id],
[edit forwarding-options helpers bootp dhcp-option82 circuit-id],
[edit forwarding-options helpers bootp interface interface-name dhcp-option82 circuit-id]
```

- For MX Series platforms:

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security option-82circuit-id]
```

Description

Configure an optional prefix for the circuit ID suboption in the DHCP option 82 information that is inserted by the switch or router into the packet header of a DHCP request before it forwards or relays the request to a DHCP server.

Default

If the `prefix` statement is not explicitly specified, no prefix is prepended to the circuit ID.

Options

host-name Add router host name to DHCP option 82 circuit ID.

logical-system-name Add logical system name to DHCP option-82 circuit ID.

This option is not used for the `prefix` statement at any of the above hierarchy levels.

routing-instance-name Add routing instance name to DHCP option-82 circuit ID.

This option is not used for the `prefix` statement occurring at the following hierarchy levels:

- [edit bridge-domains *bridge-domain-name* forwarding-options dhcp-security option-82circuit-id]
- Any of the hierarchy levels for the platforms without ELS

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security option-82 circuit-id]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

> *Example: Setting Up DHCP Option 82 Using the Same VLAN*
>
> *Example: Setting Up DHCP Option 82*
>
> http://tools.ietf.org/html/rfc3046.

prefix (DHCP Client)

IN THIS SECTION

- Syntax | 13076
- Hierarchy Level | 13077
- Description | 13077
- Required Privilege Level | 13077
- Release Information | 13077

Syntax

```
prefix {
    host-name;
```

```
    logical-system-name;
    routing-instance-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcp-client client-
identifier]
```

Description

Specify a prefix as a client identifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

prefix (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 13078
- Hierarchy Level | 13078
- Description | 13078
- Options | 13078
- Required Privilege Level | 13079
- Release Information | 13079

Syntax

```
prefix prefix;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name (relay-agent-interface-id | relay-agent-
remote-id)],
[edit forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)],
[edit forwarding-options dhcp-relay group group-name relay-option-82 (circuit-id | remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay dhcpv6 (relay-agent-
interface-id | relay-agent-remote-id)],
[edit logical-systems logical-system-name ... forwarding-options dhcp-relay ... relay-option-82
(circuit-id | remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6  (relay-agent-
interface-id | relay-agent-remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ... relay-option-82
(circuit-id | remote-id)]
```

Description

Add a prefix to the DHCP base option 82 Agent Circuit ID (suboption 1) or Agent Remote ID (suboption 2) information, or to the DHCPv6 option 18 (Relay Agent Interface-ID) or option 37 (Relay Agent Remote-ID) information in DHCP packets that DHCP relay agent sends to a DHCP server. The prefix can consist of any combination of the hostname, logical system name, and routing instance name.

Options

prefix—Any of the following:

- `host-name`—Prepend the hostname of the router configured with the `host-name` statement at the `[edit system]` hierarchy level to the DHCP option information.
- `logical-system-name`—Prepend the name of the logical system to the option information.
- `routing-instance-name`—Prepend the name of the routing instance to the option information.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... relay-agent-remote-id]` and `[edit ... remote-id]` hierarchy levels introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Including a Prefix in DHCP Options

Using DHCP Relay Agent Option 82 Information

Configuring DHCPv6 Relay Agent Options

prefix (DHCPv6 Options)

IN THIS SECTION

- Syntax | 13080
- Hierarchy Level | 13080
- Description | 13080
- Default | 13080
- Options | 13080
- Required Privilege Level | 13081
- Release Information | 13081

Syntax

```
prefix {
    host-name;
    logical-system-name;
    routing-instance-name;
    vlan-id;
    vlan-name;
}
```

Hierarchy Level

```
[edit vlans forwarding-options dhcp-security dhcpv6-options option-18]
[edit vlans forwarding-options dhcp-security dhcpv6-options option-37]
```

Description

Configure a prefix for DHCPv6 option 18 (Interface ID) or option 37 (Remote ID). When configured, the prefix is inserted into DHCPv6 packets during the DHCPv6 snooping process.

Default

If the `prefix` statement is not explicitly specified, no prefix is inserted in DHCPv6 packets.

Options

host-name	Add the host name of the switch to DHCPv6 options.
logical-system-name	Add the logical system name to the DHCPv6 options.
routing-instance-name	Add the routing instance name to the DHCPv6 options.
vlan-id	Add the VLAN ID to the DHCPv6 options.
vlan-name	Add the VLAN name to the DHCPv6 options.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

- option-37 (DHCPv6 Snooping) | 11872
- option-18 (DHCPv6 Snooping) | 11869

prefix (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 13081
- Hierarchy Level | 13082
- Description | 13082
- Options | 13082
- Required Privilege Level | 13082
- Release Information | 13082

Syntax

```
prefix prefix {
    (autonomous | no-autonomous);
    (on-link | no-on-link);
    preferred-lifetime seconds;
```

```
    valid-lifetime seconds;
}
```

Hierarchy Level

```
[edit dynamic-profiles protocols router-advertisement interface interface-name]
```

Description

Configure the prefix name in router advertisement messages.

Options

prefix—Prefix name. For dynamic configuration, specify the *$junos-ipv6-ndra-prefix* dynamic variable.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *IPv6 WAN Link Addressing with NDRA*
>
> *Dynamic Router Advertisement Configuration Overview*

prefix-export-limit (Protocols IS-IS)

IN THIS SECTION

- Syntax | 13083
- Hierarchy Level | 13083
- Description | 13083
- Options | 13084
- Required Privilege Level | 13084
- Release Information | 13084

Syntax

```
prefix-export-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis level level-number],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis
level level-number],
[edit protocols isis level level-number],
[edit routing-instances routing-instance-name protocols isis level level-number]
```

Description

Configure a limit to the number of prefixes exported into IS-IS.

By default, there is no limit to the number of prefixes that can be exported into IS-IS. To configure a limit to the number of prefixes that can be exported into IS-IS, include the `prefix-export-limit` statement. The `prefix-export-limit` statement protects the rest of the network from a malicious policy by applying a threshold filter for exported routes.

The number of prefixes depends on the size of your network. Good design advice is to set it to double the total number of IS-IS Level 1 and Level 2 routing devices in your network.

If the number of prefixes exported into IS-IS exceeds the configured limit, the overload bit is set and the overload state is reached. When other routers detect that this bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes. The overload state can be cleared by using the `clear isis overload` command.

The `show isis overview` command displays the prefix export limit when it is configured.

You can prevent IS-IS from enterning the overload state even after the prefixes reach the configured limit, by using the `no-overload-on-prefix-export-limit` command.

Options

number—Prefix limit.

- **Range:** 0 through 4,294,967,295 (2^{32} – 1)
- **Default:** None

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Redistributing OSPF Routes into IS-IS

Example: Configuring a Routing Policy to Redistribute BGP Routes with a Specific Community Tag into IS-IS

Understanding BGP Communities, Extended Communities, and Large Communities as Routing Policy Match Conditions

Understanding Routing Policies

prefix-export-limit (Protocols OSPF)

IN THIS SECTION

- Syntax | 13085
- Hierarchy Level | 13085
- Description | 13086
- Options | 13086
- Required Privilege Level | 13086
- Release Information | 13086

Syntax

```
prefix-export-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf topology (default | ipv4-multicast |
name)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
topology (default | ipv4-multicast | name)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf topology (default | ipv4-multicast | name)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf topology (default | ipv4-multicast |
name)],
```

```
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Configure a limit to the number of prefixes exported into OSPF.

Options

number—Prefix limit.

- **Range:** 0 through 4,294,967,295 (2^{32} – 1)
- **Default:** None

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for Multitopology Routing introduced in Junos OS Release 9.0.

Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Example: Limiting the Number of Prefixes Exported to OSPF

Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths

Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Understanding Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Understanding Multitopology Routing in Conjunction with PIM

prefix-limit

IN THIS SECTION

- Syntax | 13087
- Hierarchy Level | 13087
- Description | 13088
- Options | 13088
- Required Privilege Level | 13089
- Release Information | 13089

Syntax

```
prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
    drop-excess <percentage>;
    hide-excess <percentage>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family (inet | inet6) (any | flow |
labeled-unicast | multicast | unicast)],
[edit logical-systems logical-system-name protocols bgp group group-name family (inet | inet6)
(any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
family  (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family (inet | inet6) (any | flow | labeled-unicast |
multicast | unicast)],
```

```
[edit protocols bgp family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit protocols bgp family prefix-limit (inet | inet6)(any | flow | labeled-unicast | multicast
| unicast)],
[edit protocols bgp group group-name family (inet | inet6) (any | labeled-unicast | multicast |
unicast)],
[edit protocols bgp group group-name neighbor address family (inet | inet6) (any | flow |
 labeled-unicast | multicast | unicast)],
[edit routing-instances routing-instance-name protocols bgp family (inet | inet6) (any | flow |
 labeled-unicast | multicast | unicast)],
[edit routing-instances routing-instance-name protocols bgp group group-name family (inet |
inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (inet | inet6) (any | flow |  labeled-unicast | multicast | unicast)]
```

Description

Limit the number of prefixes received on a BGP peer session and a rate-limit logging when injected prefixes exceed a set limit.

This functionality is identical to the `accepted-prefix-limit` functionality except that it operates against received prefixes rather than accepted prefixes.

Options

maximum *number*

When you set the maximum number of prefixes, a message is logged when that number is exceeded.

- **Range:** 1 through 4,294,967,295 ($2^{32} - 1$)

teardown *<percentage>*

(Optional) If you include the `teardown` statement, the session is torn down when the maximum number of prefixes is reached. If you specify a percentage, messages are logged when the number of prefixes exceeds that percentage. After the session is torn down, it is reestablished in a short time unless you include the `idle-timeout` statement. Then the session can be kept down for a specified amount of time, or forever. If you specify `forever`, the session is reestablished only after you issue a `clear bgp neighbor` command.

If the `teardown` statement is not configured, a message is logged when the number of prefixes exceeds the value configured for the `maximum` option.

- **Range:** 1 through 100

idle-timeout (forever | *timeout-in-minutes*)

(Optional) If you include the `idle-timeout` statement, the session is torn down for a specified amount of time, or forever. If you specify a period of time, the session is allowed to reestablish after this timeout period. If you specify `forever`, the session is reestablished only after you intervene with a `clear bgp neighbor` command.

- **Range:** 1 through 2400

drop-excess <*percentage*>

(Optional) If you include the `drop-excess <percentage>` option, the excess routes are dropped when the maximum number of prefixes is reached. If you specify a percentage, the routes are logged when the number of prefixes exceeds that percentage value of the maximum number.

hide-excess <*percentage*>

(Optional) If you include the `hide-excess <percentage>` option, the excess routes are hidden when the maximum number of prefixes is reached. If you specify a percentage, the routes are logged when the number of prefixes exceeds that percentage value of the maximum number. If the percentage is modified, the routes are re-evaluated automatically. If the active routes drop below the specified percentage, those routes are kept as hidden.

To recover the routes from the hidden state, run the following commands to refresh the routes and get a complete re-evaluation of the routes:

`clear bgp table <table-name>` and `clear bgp neighbor <neighbor address> soft-inbound`

NOTE: You can include either the `teardown <percentage>`, `drop-excess <percentage>`, or `hide-excess <percentage>` statement option one at a time.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`drop-excess <percentage>` and `hide-excess <percentage>` options introduced in Junos OS Release 21.2.

`drop-excess <percentage>` and `hide-excess <percentage>` options introduced in Junos OS Evolved Release 21.3R1.

RELATED DOCUMENTATION

accepted-prefix-limit | 1651

Understanding Multiprotocol BGP

show bgp neighbor | 17175

prefix-list

IN THIS SECTION

- Syntax | 13090
- Hierarchy Level | 13090
- Description | 13091
- Options | 13091
- Required Privilege Level | 13091
- Release Information | 13091

Syntax

```
prefix-list name {
    ip-addresses;
    apply-path path;
}
```

Hierarchy Level

```
[edit dynamic policy-options],
[edit logical-systems logical-system-name policy-options],
[edit policy-options]
```

Description

Define a list of IPv4 or IPv6 address prefixes for use in a routing policy statement or firewall filter statement, or a list of IPv6 addresses or address prefixes for use in an IPv6 RA guard policy.

You can configure up to 85,325 prefixes in each prefix list. To configure more than 85,325 prefixes, configure multiple prefix lists and apply them to multiple firewall filter terms.

Options

name—Name that identifies the list of IPv4 or IPv6 addresses or address prefixes.

ip-addresses—These are the IPv4 or IPv6 prefixes specified as prefix/prefix-length. If you omit prefix-length for an IPv4 prefix, the default is /32prefix-length. If you omit prefix-length for an IPv6 prefix, the default is /128.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.

Support for the **vpls** protocol family introduced in Junos OS Release 10.2.

Support for IPv6 RA guard policy lists introduced in Junos OS Release 16.1 for EX Series switches.

RELATED DOCUMENTATION

Understanding Prefix Lists for Use in Routing Policy Match Conditions

dynamic-db | **6105**

Firewall Filter Match Conditions Based on Address Fields

Junos OS Routing Policies, Firewall Filters and Traffic Policers User Guide for Routing Devices

Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List

Junos OS Routing Policies, Firewall Filters and Traffic Policers User Guide for Routing Devices

prefix-list-filter

IN THIS SECTION

- Syntax | 13092
- Hierarchy Level | 13092
- Description | 13092
- Options | 13092
- Required Privilege Level | 13093
- Release Information | 13093

Syntax

```
prefix-list-filter prefix-list-name match-type <actions>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options policy-statement policy-statement-name
term term--name from],
[edit policy-options policy-statement policy-statement-name term term--name from]
```

Description

Evaluate a list of prefixes within a prefix list using specified qualifiers.

Options

prefix-list-name—Name of the prefix list to evaluate.

exact—The prefix-length component of the match prefix is equal to the route's prefix length

longer—The route's prefix length is greater than the prefix-length component of the match prefix.

orlonger—The route's prefix length is equal to or greater than the prefix-length component of the configured match prefix.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding Prefix Lists for Use in Routing Policy Match Conditions

prefix-list (Junos Telemetry Interface)

IN THIS SECTION

- Syntax | 13094
- Hierarchy Level | 13094
- Description | 13094
- Options | 13094
- Required Privilege Level | 13095
- Release Information | 13095

Syntax

```
prefix-list {
    table table-name {
        family {
            inet {
                prefix [ prefix ... ];
            }
            inet6 {
                prefix [ prefix ... ];
            }
        }
    }
}
```

Hierarchy Level

```
[edit system fib-streaming]
```

Description

Enable specific prefix addresses to be periodically streamed from the target (source) to a collector. Only interfaces with those prefixes and their corresponding next-hop and next-hop group containers are exported to the IP OpenConfig Abstract Forwarding Table (oc-aft) collector. Reducing the set of interfaces to only those with prefixes of interest decreases the overall CPU and resource usage on Routing Engines, FPCs, and MPCs.

Export data from your source (target) to a collector using the sensor **/network-instances/network-instance/afts/**. When using prefix-filter mode, we recommend you use streaming subscription mode at a sampling frequency of 5 minutes or more and that all collectors have the same sampling frequency. If you are using non-prefix-filter mode, we recommend you use ON_CHANGE subscription mode. You can also export data using User Datagram Protocol (UDP).

Options

table-name	A forwarding information base (FIB) table.

prefix An IPv4 (`inet`) or IPv6 `inet6` address prefix.

Required Privilege Level

edit

Release Information

Statement introduced in Junos OS 24.2R1 for MX2020 routers.

RELATED DOCUMENTATION

show fib-streaming state | 19927

Configuring IP-AFT Prefix Filtering

prefix-list-name

IN THIS SECTION

- Syntax | 13095
- Hierarchy Level | 13096
- Description | 13096
- Options | 13096
- Required Privilege Level | 13096
- Release Information | 13096

Syntax

```
prefix-list-name prefix-list-name;
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard policy policy-name discard
(forwarding-options)]
[edit forwarding-options access-security router-advertisement-guard policy policy-name accept (Router
Advertisement Guard Policy) match-list]
```

Description

Configure a list of IPv6 address prefixes for an IPv6 Router Advertisement (RA) guard policy. The policy is used to validate the source IPv6 address prefix of an incoming RA message against the IPv6 address prefixes in this list. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.

You can use a list of IPv6 address prefixes for validating RA messages as part of either an accept policy or a discard policy. Before you can include a list in an RA guard policy, you must configure the list name at the [`edit policy-options prefix-list`] hierarchy level. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped. When RA guard is enabled by using a discard policy, any RA messages that match the conditions are dropped, and RA messages that do not match the conditions are forwarded.

Options

prefix-list-name Configure a list of IPv6 address prefixes for an RA guard policy. The policy is used to validate the source of an incoming RA message by comparing the IPv6 address prefix of the RA message to the IPv6 address prefixes contained in the list.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

prefix-list (PIM RPF Selection)

IN THIS SECTION

- Syntax | 13097
- Hierarchy Level | 13097
- Description | 13098
- Options | 13098
- Required Privilege Level | 13098
- Release Information | 13098

Syntax

```
prefix-list prefix-list-addresses {
    source source-address {
        next-hop next-hop-address;
    }
    wildcard-source {
        next-hop next-hop-address;
    }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address
source source-address],
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address
wildcard-source],
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-
```

```
list-addresses source source-address],
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-
list-addresses wildcard-source]
```

Description

(Optional) Configure a list of prefixes (addresses) for multiple PIM groups.

Options

prefix-list-addresses—List of prefixes (addresses) for multiple PIM groups.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring PIM RPF Selection

prefix-name (Security Static NAT)

IN THIS SECTION

- Syntax | 13099
- Hierarchy Level | 13099
- Description | 13099
- Options | 13099

Required Privilege Level | 13099
Release Information | 13100

Syntax

```
prefix-name {
    address-prefix-name;
    mapped-port lower-port-range to upper-port-range;
    routing-instance (routing-instance-name| default);
}
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name then static-nat]
```

Description

Specify an address from the address book.

Options

- *address-prefix-name*—Specify address prefix name from address book.
- `mapped-port` *lower-port-range* `to` *upper-port-range*—Specify a destination port or port range to allow static NAT to map ports.
- `routing-instance` —Specify routing instance type:
  - *routing-instance-name*—Use the user-defined static NAT routing instance to perform static NAT.
  - `default`—Use the default routing-instance.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

prefix (Protocols IPv6 Neighbor Discovery)

IN THIS SECTION

- Syntax | 13100
- Hierarchy Level | 13100
- Description | 13101
- Options | 13101
- Required Privilege Level | 13101
- Release Information | 13101

Syntax

```
prefix prefix {
    (autonomous | no-autonomous);
    (on-link | no-on-link);
    preferred-lifetime seconds;
    valid-lifetime seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Description

Configure prefix properties in router advertisement messages.

Options

prefix—Prefix name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

prefix (Remote ID for Option 82)

IN THIS SECTION

- Syntax | 13102
- Hierarchy Level | 13102
- Description | 13102
- Default | 13102
- Options | 13102
- Required Privilege Level | 13102
- Release Information | 13103

Syntax

```
prefix (hostname | mac | none);
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name) dhcp-option82 remote-id]
[edit forwarding-options helpers bootp dhcp-option82 remote-id]
[edit forwarding-options helpers bootp interface interface-name dhcp-option82 remote-id]
```

Description

Configure an optional prefix for the remote ID suboption in the DHCP option 82 information that is inserted by the switch into the packet header of a DHCP request before it forwards or relays the request to a DHCP server.

Default

If **prefix** is not explicitly specified, no prefix is appended to the remote ID.

Options

hostname—Name of the host system (the switch) that is forwarding or relaying the DHCP request from the DHCP client to the DHCP server.

mac—MAC address of the host system (the switch) that is forwarding or relaying the DHCP request from the DHCP client to the DHCP server.

none—No prefix is applied to the remote ID.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Example: Setting Up DHCP Option 82 Using the Same VLAN

Example: Setting Up DHCP Option 82

http://tools.ietf.org/html/rfc3046.

prefix (Security Static NAT)

IN THIS SECTION

- Syntax | 13103
- Hierarchy Level | 13104
- Description | 13104
- Options | 13104
- Required Privilege Level | 13104
- Release Information | 13104

Syntax

```
prefix {
    address-prefix;
    mapped-port lower-port-range to upper-port-range;
    routing-instance (routing-instance-name| default);
}
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name then static-nat]
```

Description

Specify a static IP address prefix.

> **NOTE**: If you use the *inet* option for translation of IPv6 to IPv4 addresses (and vice versa), you do not need to specify a prefix because the *inet* option automatically strips the first 96 most significant bits from the 128-bit IPv6 address.

Options

- *address-prefix*—Specify address prefix.
- `mapped-port` *lower-port-range* `to` *upper-port-range*—Specify a destination port or port range to allow static NAT to map ports.
- `routing-instance` —Specify routing instance type:
  - *routing-instance-name*—Use the user-defined static NAT routing instance to perform static NAT.
  - `default`—Use the default routing-instance.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

prefix-segment

IN THIS SECTION

- Syntax | 13105
- Hierarchy Level | 13105
- Description | 13105
- Required Privilege Level | 13106
- Release Information | 13106

Syntax

```
prefix-segment {
  algorithm name {
    index index;
    node-segment;
    redistribute;
  }
}
```

Hierarchy Level

```
[edit logical-systems name policy-options policy-statement name from],
[edit logical-systems name policy-options policy-statement name then],
[edit policy-options policy-statement policy-name term term-name from],
[edit policy-options policy-statement polic-name term term-name then],
[edit protocols ospf source-packet-routing]
```

Description

Configure prefix segment attributes such as index and node segment. Prefix segment index is the index assigned to a specific prefix. This is used by all other remote routers in the network to index the prefix into respective segment routing global blocks (SRGB) to derive the segment identifier and to forward the

traffic destined for the current prefix. Prefix segment index is provisioned through policy for each prefix along with the option to mark it as a node segment.

Indicate route which is evaluating against policy must have prefix-segment information associated with it. If route has prefix-segment information associated with it, policy action results in Policy Accept, if not, policy action is Policy Reject. Once policy action result in Policy Accept after evaluating from clause, action specified in the then clause is applied to the matching route.

algorithm *name*

Specify the prefix segment attributes for strict/flex algorithm.

- **Range:** 1 through 255

index *index*

Specify the prefix segment index

- **Range:** 0 through 199,999

node-segment

Set node segment flag for this prefix segment.

reuse-allocated-sid

Specify this option if a large number of prefixes need to be assigned SID value. You can reuse the allocated SID on the IGP route only if they are available during export of IGP routes.

`redistribute`

Enable prefix-sid redistribution between BGP and ISIS via policy configuration without specifying index under prefix-segment.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

Statement introduced under `[edit protocols ospf source-packet-routing]` hierarchy in Junos OS Release 19.1R1 for MX Series.

`reuse-allocated-sid` option introduced in Junos OS Release 20.4R1.

`redistribute` option introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Understanding Adjacency Segments, Anycast Segments, and Configurable SRGB in SPRING

Configuring Anycast and Prefix segments in SPRING for IS-IS Protocol

prefix-segment-range

IN THIS SECTION

- Syntax | 13107
- Hierarchy Level | 13107
- Description | 13108
- Options | 13108
- Required Privilege Level | 13108
- Release Information | 13108

Syntax

```
prefix-segment-range prefix-segment-range-name {
    attached;
    domain-wide-flooding;
    no-node-segment;
    size size;
    start-index start-index;
    start-prefix start-prefix;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options source-packet-routing mapping-
server-entry],
[edit logical-systems name routing-options source-packet-routing mapping-server-entry],
```

```
[edit routing-instances name routing-options source-packet-routing mapping-server-entry],
[edit routing-options source-packet-routing mapping-server-entry]
```

Description

Configure the prefix segment range for the LDP mapping server.

Options

prefix-segment-range-name	Name of the prefix segment range.
attached	(Optional) Set the flag in IS-IS mapping server advertisement to indicate that the prefixes and SIDs advertised in the SID or Label Binding TLV are directly connected to their originators.
domain-wide-flooding	(Optional) Set an S flag in the IS-IS mapping server advertisement to indicate that the SID or Label Binding TLV is flooded across the entire routing domain.
no-node-segment	(Optional) Clear the node segment flag in the mapping server prefix segment to indicate that the prefix has originated from a single node.
size ***size***	Size of prefix segment range. • **Range:** 1 through 1024
start-index ***start-index***	Include start index. • **Range:** 0 through 199999
start-prefix ***start-prefix***	Include start prefix.

Required Privilege Level

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

`attached`, `domain-wide-flooding`, and `no-node-segment` options introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

LDP Mapping Server for Interoperability of Segment Routing with LDP Overview

source-packet-routing (routing-options) | 28264

mapping-server-entry | 10346

prefix-segment (Routing Options)

IN THIS SECTION

- Syntax | 13109
- Hierarchy Level | 13109
- Description | 13110
- Options | 13110
- Required Privilege Level | 13110
- Release Information | 13110

Syntax

```
prefix-segment prefix-segment {
    index index;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options source-packet-routing mapping-
server-entry],
[edit logical-systems name routing-options source-packet-routing mapping-server-entry],
[edit routing-instances name routing-options source-packet-routing mapping-server-entry],
[edit routing-options source-packet-routing mapping-server-entry]
```

Description

Configure the IP address and index number of the prefix segment for the LDP mapping server.

Options

prefix-segment	IP address of the prefix segment.
index *index*	Prefix segment index. • **Range:** 0 through 199999

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

LDP Mapping Server for Interoperability of Segment Routing with LDP Overview

source-packet-routing (routing-options) | **28264**

premium (Hierarchical Policer)

IN THIS SECTION

- Syntax | **13111**
- Hierarchy Level | **13111**
- Description | **13111**
- Options | **13111**
- Required Privilege Level | **13111**
- Release Information | **13111**

Syntax

```
premium {
    if-exceeding {
        bandwidth-limit bandwidth;
        burst-size-limit burst;
    }
    then {
        discard;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer],
[edit firewall hierarchical-policer]
```

Description

On M40e, M120, and M320 edge routers with FPC input as FFPC and FPC output as SFPC, and on MX Series, T320, T640, and T1600 edge routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, specify a premium level for a hierarchical policer.

Options

Options are described separately.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the [edit dynamic-profiles ... hierarchical-policer *name*] hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Applying Policers
- Guidelines for Applying Traffic Policers
- Hierarchical Policer Configuration Overview
- Hierarchical Policers
- aggregate (Hierarchical Policer) | 2064
- bandwidth-limit (Hierarchical Policer)
- burst-size-limit (Hierarchical Policer) | 3084
- hierarchical-policer (aggregate or premium) | 8057
- if-exceeding (Hierarchical Policer) | 8371

premium (Output Priority Map)

IN THIS SECTION

- Syntax | 13112
- Hierarchy Level | 13113
- Description | 13113
- Required Privilege Level | 13113
- Release Information | 13113

Syntax

```
premium {
    forwarding-class class-name {
        loss-priority (high | low);
```

```
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile
output-priority-map classifier]
```

Description

For Gigabit Ethernet IQ interfaces only, define the classifier for egress premium traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

input-priority-map | 8716

premium (Policer)

IN THIS SECTION

Syntax | 13114

- Hierarchy Level | 13114
- Description | 13114
- Required Privilege Level | 13114
- Release Information | 13114

Syntax

```
premium {
    bandwidth-limit bps;
    burst-size-limit bytes;
}
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile
policer cos-policer-name]
```

Description

Define a policer to apply to nonpremium traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers

aggregate (Gigabit Ethernet CoS Policer) | 2063

ieee802.1p | 8366

preprovisioned

IN THIS SECTION

- Syntax | 13115
- Hierarchy Level | 13115
- Description | 13115
- Required Privilege Level | 13116
- Release Information | 13116

Syntax

```
preprovisioned;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Enable the preprovisioned configuration mode for a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.

When the preprovisioned configuration mode is enabled, you cannot use the CLI or the J-Web interface to change the primary-role priority or member ID of member switches.

You must use this statement to configure an EX8200 Virtual Chassis. Nonprovisioned configuration of an EX8200 Virtual Chassis is not supported.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- *Preprovisioning a Virtual Chassis Fabric*
- Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File
- Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines
- Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)
- Configuring an EX8200 Virtual Chassis (CLI Procedure)
- Configuring an EX9200 Virtual Chassis
- Configuring an EX4650 or a QFX Series Virtual Chassis
- *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*

preprovisioned (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 13117
- Hierarchy Level | 13117
- Description | 13117
- Required Privilege Level | 13117
- Release Information | 13117

Syntax

```
preprovisioned;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Enable creation of a Virtual Chassis by means of a preprovisioned configuration.

To configure a Virtual Chassis consisting of MX Series routers, you must create a preprovisioned configuration on the primary router in the Virtual Chassis by specifying the serial number, member ID, and role for each router (member chassis) in the Virtual Chassis. When a new member router joins the Virtual Chassis, its serial number is compared against the values specified in the preprovisioned configuration. If the serial number of a joining router does not match any of the configured serial numbers, the software prevents that router from becoming a member of the Virtual Chassis.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

> *Configuring Preprovisioned Member Information for a Virtual Chassis*
>
> *Example: Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis*

preserve-fpc-poweron-sequence

IN THIS SECTION

- Syntax | 13118
- Hierarchy Level | 13118
- Description | 13118
- Required Privilege Level | 13119
- Release Information | 13119

Syntax

```
preserve-fpc-poweron-sequence;
```

Hierarchy Level

```
[edit chassis]
```

Description

Preserve the sequence in which the line cards—specifically, MPCs on MX series routers—on a router or switch are powered on when the device is restarted. During a system reboot, the line cards are brought online in the sequence specified in the system log file **/var/log/fpc_poweron_seq.log**. When a line card goes offline, its entry is removed from the log file. You can use the `show chassis poweron sequencs` command to view the configured power-on sequence.

NOTE:

- If both `preserve-fpc-poweron-sequence` and `fru-poweron-sequence` statements are configured, then the power-on sequence specified in the `fru-poweron-sequence` statement takes precedence.
- If `preserve-fpc-poweron-sequence` is configured and `fru-poweron-sequence` not configured, then the line cards are powered on in the sequence preserved in the system log file **/var/log/ fpc_poweron_seq.log**.

- If neither of these statements is configured, then the line cards are powered on in the ascending order of their slot numbers. Line cards whose slot numbers are not specified in the log file are powered on in the ascending order of their slot numbers.

Required Privilege Level

chassis—To view this statement in the configuration.

chassis-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

fru-poweron-sequence | 7613

show chassis power sequence | 18688

Understanding How Dynamic Power Management Enables Better Utilization of Power

preserve-incoming-fragment-size

IN THIS SECTION

- Syntax | 13120
- Hierarchy Level | 13120
- Description | 13120
- Required Privilege Level | 13120
- Release Information | 13121

Syntax

```
preserve-incoming-fragment-size;
```

Hierarchy Level

```
[edit security flow] ]
```

Description

Enable the preserve incoming fragment size feature that allows the SRX Series Firewall to preserve the size of incoming fragments to be used in determining the best maximum transmission unit (MTU) size for the egress interface.

When data is sent from one host to another, it is transmitted as a series of packets. Performance is improved and network resources are conserved when packets of the largest size can transit the path from the source node to the destination node without being fragmented at any link in the datapath.

If application-layer services are configured on the SRX Series Firewall, packet fragments at the ingress interface must be reassembled before the services can be applied and the content inspected. These reassembled packet fragments must be broken down again before the data is transmitted out the egress interface.

When a packet must be fragmented into smaller packets to transit a link in the path because the packet is larger than the MTU size established for that link, each of the resulting fragments must contain packet header information, in addition to the payload, or data. The increased overhead can lower throughput and degrade network performance. Also, the packet fragments must be reassembled at the destination node, which consumes additional network resources.

By default, the SRX Series Firewall uses the MTU size configured for the egress interface to determine the size for the packet fragments it transmits. However, if you enable the preserve incoming fragment size feature, the SRX Series Firewall detects and saves the size of incoming packet fragments and takes that into account. To diminish the likelihood of packet fragmentation in the datapath, the SRX Series Firewall sets the egress interface MTU size to the smaller of two values: It identifies the maximum size of all incoming fragments and it compares that size to the existing MTU size of the egress interface. The SRX Series Firewall takes the smaller number and uses it for the egress interface MTU size.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Understanding How Preserving Incoming Fragmentation Characteristics Can Improve Throughput

preserve-interface

IN THIS SECTION

- Syntax | 13121
- Hierarchy Level | 13121
- Description | 13122
- Required Privilege Level | 13122
- Release Information | 13122

Syntax

```
preserve-interface;
```

Hierarchy Level

```
[edit interfaces interface-name sonet-options aps]
```

Description

Provide link PIC replication, providing MLPPP link redundancy at the port level. This feature is supported with SONET APS and the following link PICs:

- Channelized OC3 IQ PIC
- Channelized OC12 IQ PIC
- Channelized STM1 IQ PIC

Link PIC replication provides the ability to add two sets of links, one from the active SONET PIC and the other from the standby SONET PIC, to the same bundle. If the active SONET PIC fails, links from the standby PIC are used without triggering link renegotiation. All the negotiated state is replicated from the active links to the standby links to prevent link renegotiation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Link State Replication for Redundant Link PICs

preserve-nexthop-hierarchy (SRv6 TE)

IN THIS SECTION

- Syntax | 13123
- Hierarchy Level | 13123
- Description | 13123
- Options | 13123

- Required Privilege Level | 13123
- Release Information | 13124

Syntax

```
preserve-nexthop-hierarchy {
                  (sr-mpls | srv6);
}
```

Hierarchy Level

```
[edit protocols source-packet-routing]
```

Description

Enable preserve-nexthop hierarchy support for source packet routing. When you enable `preserve-nexthop-hierarchy`, SRv6 TE transport is populated correctly so that service resolution is performed properly. The `preserve-nexthop-hierarchy` statement at the `[edit protocols source-packet-routing]` hierarchy level is a prerequisite for configuring SR-TE SRv6 tunnels. However, if you enable the global configuration statement `set routing-options resolution preserve-nexthop-hierarchy`, then `preserve-nexthop-hierarchy` statement at the `[edit protocols source-packet-routing]` is not mandatory. Additionally, you have an option to enable the `preserve-nexthop-hierarchy` only for a certain data plane for transports populated by SR-TE. The following options when enabled disable the `preserve-nexthop-hierarchy` for that data plane.

Options

sr-mpls Disable preserve-nexthop-hierarchy for sr-mpls.

srv6 Disable preserve-nexthop-hierarchy for srv6.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

source-packet-routing (Segment Routing Traffic Engineering) | 28257

segment-list | 16521

Understanding SR-TE Policy for SRv6 Tunnel

preserve-parity (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 13124
- Hierarchy Level | 13124
- Description | 13125
- Required Privilege Level | 13125
- Release Information | 13125

Syntax

```
preserve-parity;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Description

Assign a port with the same parity (even or odd) as the incoming source port. This feature is not available if you configure port-block allocation, and is not available for deterministic NAT.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

preserve-range (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 13125
- Hierarchy Level | 13125
- Description | 13126
- Required Privilege Level | 13126
- Release Information | 13126

Syntax

```
preserve-range;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Description

For source NAT with port translation, except for deterministic NAT, assign a port within the same range as the incoming port—either 0 through 1023 or 1024 through 65,535. This feature is not available if you configure port block allocation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

primary

IN THIS SECTION

- Syntax | 13126
- Hierarchy Level | 13127
- Description | 13127
- Options | 13127
- Required Privilege Level | 13127
- Release Information | 13127

Syntax

```
primary path-name;
```

Hierarchy Level

```
[edit protocols mpls                    label-switched-path                    lsp-name]
```

Description

Specify the primary path to use for a label switched path (LSP). You can configure only one primary path.

Options

path-name —Name of the primary path that you created with the `path` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Path Protection in an MPLS Network (CLI Procedure)

primary

IN THIS SECTION

- Syntax | **13128**
- Hierarchy Level | **13128**
- Description | **13128**
- Required Privilege Level | **13128**

Release Information | 13128

Syntax

```
primary interface name;
```

Hierarchy Level

```
[edit protocols ptp client interface interface-name],
[edit protocols ptp master interface interface-name]
```

Description

Configure the primary member link in an aggregated Ethernet bundle to enable Precision Time Protocol (PTP) over a link aggregation group. PTP primary streams are created on the FPC on which the primary interface is present. Announce and sync packets are transmitted on this primary PTP Aggregated Ethernet link.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

primary (Adaptive Services Interfaces)

IN THIS SECTION

- Syntax | 13129
- Hierarchy Level | 13129
- Description | 13129
- Options | 13129
- Required Privilege Level | 13129
- Release Information | 13130

Syntax

```
primary interface-name;
```

Hierarchy Level

```
[edit interfaces (rsp0 | rsp1) redundancy-options]
```

Description

Specify the primary adaptive services interface.

Options

interface-name—The identifier for the AS or Multiservices PIC interface, which must be of the form `sp-fpc/pic/port`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring AS or Multiservices PIC Redundancy

primary (Address on Interface)

IN THIS SECTION

- Syntax | 13130
- Hierarchy Level | 13130
- Description | 13131
- Default | 13131
- Required Privilege Level | 13131
- Release Information | 13131

Syntax

```
primary;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family address address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family  family address address]
```

Description

Configure this address to be the primary address of the protocol on the interface. If the logical unit has more than one address, the primary address is used by default as the source address when packet transfer originates from the interface and the destination address does not indicate the subnet.

> **NOTE**: The `edit logical-systems` hierarchy is not available on QFabric systems.

Default

For unicast traffic, the primary address is the lowest non-127 (in other words, non-loopback) preferred address on the unit.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Interface Address

primary connection (Identity Management Advanced Query)

IN THIS SECTION

- Syntax | 13132
- Hierarchy Level | 13132
- Description | 13132
- Options | 13133

Required Privilege Level | 13133
Release Information | 13134

Syntax

```
primary {
    address address;
    ca-profile ca-profile;
    client-id client-id;
    client-secret client-secret;
    interface interface-name;
    routing-instance routing-instance -name;
    source source-address;
}
```

Hierarchy Level

```
[edit services user-identification identity-management connection]
```

Description

Configure parameters that the SRX Series Firewall uses to connect to the Juniper Identity Management Service (JIMS) primary server and authenticate to it to obtain an access token. JIMS requires that the SRX Series Firewall use OAuth2 to authenticate to it before the SRX Series Firewall can query the JIMS server for user identity information. The SRX Series Firewall must provide the JIMS server with credentials, including a client ID and a client secret. If the client is authenticated-in this case the SRX Series Firewall—it is granted an access token. (See RFC 6749.) Both the client ID and the client secret must be consistent with the API client configured on the JIMS primary server.

In addition to configuring the client ID and the client secret, you configure the filename of the JIMS's ca-certificate. The certificate enables the SRX Series Firewall to verify the identity of JIMS and that it is trusted for the SSL connection.

If the deployment configuration consists of more than one JIMS server, a primary and secondary relationship is established. The SRX Series Firewall always attempts to connect to the primary server. When one or more queries to the primary server fails, the system falls back to the secondary server.

address- Configure the IP address for the primary Juniper Identity Management Service (JIMS) server. The SRX Series Firewall requires the server IP address to connect to the server to obtain an access code that allows it to query the server for user identity information. The IP address is configured as part of a collection of information which includes the SRX Series Firewall's client ID, client secret, and ca-certificate information.

The SRX Series Firewall sends a unique set of identification information to the primary server and the secondary server. The feature supports only IPV4 addresses.

client-id- Client ID that the SRX Series provides to the JIMS primary server as part of its authentication to it. The SRX Series Firewall must authenticate to the server to obtain an access token that allows the SRX Series Firewall to query the server for user identity information The client ID must be consistent with the API client configured on the JIMS primary server.

client-secret- Client secret that the SRX Series provides to the JIMS primary server as part of its authentication to it. The client secret must be consistent with the API client configured on the JIMS primary server.

WARNING: Before you use this feature, you must disable any other actively used options under the [edit services user-identification] hierarchy. You cannot commit this configuration if active directory authentication and the ClearPass query and webapi functions are configured and committed.

Options

address	IP address of the primary server.
ca-profile	CA profile name
client-id	Client ID for OAuth2 grant
client-secret	Client secret for OAuth2 grant

Required Privilege Level

1. services—To view this statement in the configuration.
2. services-control—To add this statement to the configuration.

The remaining statements are explained separately. See CLI Explorer.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

IPv6 address support introduced in Junos OS Release 18.3R1.

Source, interface, and routing-instance options are introduced in Junos OS Release 21.1R1.

Option ca-profile introduced in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

Understanding Advanced Query Feature for Obtaining User Identity Information from JIMS

filter (Identity Management Advanced Query) | **7138**

ip-query (Identity Management Advanced Query) | **9259**

authentication-entry-timeout (Identity Management Advanced Query) | **2597**

batch query | **2886**

primary-data-record-fields

IN THIS SECTION

- Syntax | **13135**
- Hierarchy Level | **13135**
- Description | **13135**
- Default | **13135**
- Options | **13135**
- Additional Information | **13136**
- Required Privilege Level | **13137**
- Release Information | **13137**

Syntax

```
primary-data-record-fields {
   name;
}
```

Hierarchy Level

```
[edit services inline-monitoring template template-name]
```

Description

Configure which IPFIX information elements (IEs) to include in the primary data record, which is exported by the data path and contains the specific properties seen by the data path, either from the sampled packet or the packet notification header or context. With this statement, each inline-monitoring instance can export IPFIX packets using a different, customized data template, resulting in a custom set of IEs in the IPFIX packets for each instance. The fields are encoded using one or more Common Properties IDs. As per RFC 5473, the primary data record is exported using the IPFIX data template.

Default

datalink-frame-section—Datalink Frame Section (IE 315) is always included in the primary data record.

Options

name Configure which IPFIX information elements (IEs) to include in the primary data record. Per RFC 5102, the Common Properties ID (CPID) is an identifier of a set of common properties that is locally unique per observation domain and transport session.

- cpid-egress-interface-index—CPID Egress Interface Index (IE 137); reports zero value when ingress monitoring
- cpid-forwarding-class-drop-priority—CPID Forwarding Class Drop Priority (IE 137)
- cpid-forwarding-exception-code—CPID Forwarding Exception Code (IE 137)
- cpid-forwarding-nexthop-id—(Junos OS only) CPID Forwarding Nexthop ID (IE 137)
- cpid-ingress-interface-index—CPID Ingress Interface Index (IE 137)

- cpid-underlying-ingress-interface-index—(Junos OS only) CPID Underlying Ingress Interface Index (IE 137); reports a zero value when egress monitoring or if this value is not available in the data path.
- datalink-frame-size—(Junos OS only) Datalink Frame Size (IE 312)
- direction—Flow direction (IE 61)
- egress-interface-snmp-id—Egress Interface SNMP ID (IE 14); reports zero value when ingress monitoring
- ingress-interface-snmp-id—Ingress Interface SNMP ID (IE 10)

Additional Information

Table 98: Values Reported for Aggregated Ethernet (AE) Interfaces

CPID	Ingress on AE Interface	Egress on AE Interface
ingressInterfaceIndex	Logical interface index of the AE interface	Logical interface index of the AE interface
underlyingIngressInterfaceIndex (Junos OS only)	Logical interface index of member link	0
egressInterfaceIndex	0	Logical interface index of member link

Table 99: Values Reported for IRB Interfaces

CPID	Ingress on IRB Interface	Egress on IRB interface
ingressInterfaceIndex	Logical interface index of the IRB interface	Logical interface index of ingress interface
underlyingIngressInterfaceIndex (Junos OS only)	Logical interface index of `vlan-bridge` encapsulated interface	0
egressInterfaceIndex	0	Logical interface index of `vlan-bridge` encapsulated interface

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

Statement introduced in Junos Evolved OS Release 22.2R1.

RELATED DOCUMENTATION

Juniper Resiliency Interface

primary-dns

IN THIS SECTION

- Syntax | 13137
- Hierarchy Level | 13138
- Description | 13138
- Options | 13138
- Required Privilege Level | 13138
- Release Information | 13138

Syntax

```
primary-dns primary-dns;
```

Hierarchy Level

```
[edit access group-profile profile-name client client-name ppp],
[edit access profile profile-name ppp]
```

Description

Configure the primary Domain Name System (DNS) server.

Options

primary-dns—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

primary-interface (Aggregated Inline Services)

IN THIS SECTION

- Syntax | 13139
- Hierarchy Level | 13139
- Description | 13139

- Options | 13139
- Required Privilege Level | 13139
- Release Information | 13140

Syntax

```
primary-interface interface-name;
```

Hierarchy Level

```
[edit interfaces asix aggregated-inline-services-options]
```

Description

Specify the primary (active) inline services member link in the asi bundle. You must also configure a secondary (backup) member link on a different MPC with the `secondary-interface` statement. The secondary member provides 1:1 redundancy for subscriber service sessions on the primary member link. The bandwidth configured at the `[edit chassis fpc slot pic number inline-services bandwidth]` hierarchy level must be the same for both member links.

Redundancy is not achievable if you configure the primary and secondary interface on the same MPC, because both member interfaces go down if the card goes down. Consequently, if you configure both interfaces on the same MPC, the subsequent configuration commit fails.

Options

interface-name	Name of an inline services physical interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

primary (Interface for Router)

IN THIS SECTION

- Syntax | 13140
- Hierarchy Level | 13140
- Description | 13141
- Required Privilege Level | 13141
- Release Information | 13141

Syntax

```
primary;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family  family]
```

Description

Configure the primary interface for a device. By default, the multicast-capable interface with the lowest-index address is chosen as the primary interface. If there is no such interface, the point-to-point interface with the lowest-index address is chosen. Otherwise, any interface with an address can be picked. In practice, this means that, on the device, the **fxp0** or **em0** interface is picked by default. To configure a different interface to be the primary interface, you include this statement.

The *primary interface* for the router has the following characteristics:

- It is the interface through which the packets go out when you type a command such as ping 255.255.255.255—that is, a command that does not include an interface name (there is no interface ***type*-0/0/0.0** qualifier) and where the destination address does not imply any particular outgoing interface.
- It is the interface on which multicast applications running locally on the router, such as Session Announcement Protocol (SAP), perform group joins by default.
- It is the interface from which the default local address is derived for packets sourced out of an unnumbered interface if there are no non-127 addresses configured on the loopback interface, lo0.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Default, Primary, and Preferred Addresses and Interfaces

primary (Link Services IQ PIC Interfaces)

IN THIS SECTION

- Syntax | 13142
- Hierarchy Level | 13142
- Description | 13142
- Options | 13142
- Required Privilege Level | 13142
- Release Information | 13143

Syntax

```
primary interface-name;
```

Hierarchy Level

```
[edit interfaces rlsqnumber redundancy-options]
```

Description

Specify the primary Link Services IQ PIC interface.

Options

interface-name—The identifier for the Link Services IQ PIC interface, which must be of the form `lsq-`*fpc*/*pic*/*port*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces

primary (MBGP MVPN)

IN THIS SECTION

- Syntax | 13143
- Hierarchy Level | 13143
- Description | 13144
- Options | 13144
- Required Privilege Level | 13144
- Release Information | 13144

Syntax

```
primary address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvpn
static-umh],
[edit routing-instances routing-instance-name protocols mvpn static-umh]
```

Description

Statically set the primary upstream multicast hop (UMH) for type 7 (S,G) routes.

If the primary UMH is unavailable, the backup UMH is used.

Options

address	Address of the primary UMH.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

Example: Configuring Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

sender-based-rpf (MBGP MVPN) | **16552**

static-umh (MBGP MVPN) | **28585**

unicast-umh-election | **30353**

primary (Protocols MPLS)

IN THIS SECTION

- Syntax | **13145**

- Hierarchy Level | 13145
- Description | 13146
- Options | 13146
- Required Privilege Level | 13146
- Release Information | 13146

Syntax

```
primary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps;
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
    select (manual | unconditional);
    standby;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Specify the primary path to use for an LSP. You can configure only one primary path.

You can optionally specify preference, CoS, and bandwidth values for the primary path, which override any equivalent values that you configure for the LSP (at the `[edit mpls label-switched-path lsp-name]` hierarchy level).

Options

path-name—Name of a path that you created with the `path` statement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Primary and Secondary LSPs

primary-server

IN THIS SECTION

- Syntax | 13147
- Hierarchy Level | 13147
- Description | 13147
- Options | 13147

Required Privilege Level | 13147
Release Information | 13147

Syntax

```
primary-server {
    address ipv4-address;
    login sender-email-address {
        password password;
    }
}
```

Hierarchy Level

```
[edit smtp]
```

Description

Configure Simple Mail Transfer Protocol (SMTP) primary server for access authorization for SMTP requests.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 10.0.

primary (Virtual Tunnel in Routing Instances)

IN THIS SECTION

- Syntax | 13148
- Hierarchy Level | 13148
- Description | 13148
- Default | 13149
- Required Privilege Level | 13149
- Release Information | 13149

Syntax

```
primary;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name interface vt-
fpc/pic/port.unit-number],
[edit routing-instances routing-instance-name interface vt-fpc/pic/port.unit-number]
```

Description

In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure the virtual tunnel (VT) interface to be used as the primary interface for multicast traffic.

Junos OS supports up to eight VT interfaces configured for multicast in a routing instance to provide redundancy for MBGP (next-generation) MVPNs. This support is for RSVP point-to-multipoint provider tunnels as well as multicast Label Distribution Protocol (MLDP) provider tunnels. This feature works for extranets as well.

This statement allows you to configure one of the VT interfaces to be the primary interface, which is always used if it is operational. If a VT interface is configured as the primary, it becomes the nexthop that is used for traffic coming in from the core on the label-switched path (LSP) into the routing instance.

When a VT interface is configured to be primary and the VT interface is used for both unicast and multicast traffic, only the multicast traffic is affected.

If no VT interface is configured to be the primary or if the primary VT interface is unusable, one of the usable configured VT interfaces is chosen to be the nexthop that is used for traffic coming in from the core on the LSP into the routing instance. If the VT interface in use goes down for any reason, another usable configured VT interface in the routing instance is chosen. When the VT interface in use changes, all multicast routes in the instance also switch their reverse-path forwarding (RPF) interface to the new VT interface to allow the traffic to be received.

To realize the full benefit of redundancy, we recommend that when you configure multiple VT interfaces, at least one of the VT interfaces be on a different Tunnel PIC from the other VT interfaces. However, Junos OS does not enforce this.

Default

If you omit this statement, Junos OS chooses a VT interface to be the active interface for multicast traffic.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

> *Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs*
>
> Example: Configuring MBGP MVPN Extranets

primary-vlan

IN THIS SECTION

- Syntax | 13150
- Hierarchy Level | 13150
- Description | 13150
- Required Privilege Level | 13151
- Release Information | 13151

Syntax

```
primary-vlan vlan-name;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Description

For a private VLAN (PVLAN), configure the primary VLAN. The primary VLAN is always tagged.

- If the PVLAN is configured on a single switch, do not assign a tag to the community VLANs.
- If the PVLAN is configured to span multiple switches, you must assign tags to the community VLANs also.

For a community VLAN, configure the primary VLAN. The primary VLAN must be tagged, and the community VLAN must be untagged.

If you want to create a community VLAN, you must configure the primary VLAN to be private using the "pvlan" on page 13882 statement.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type **?** after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Statement updated with enhanced **?** (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.

RELATED DOCUMENTATION

Creating a Private VLAN on a Single QFX Switch without ELS Support

Creating a Private VLAN on a Single EX Series Switch without ELS Support (CLI Procedure)

Creating a Private VLAN Spanning Multiple QFX Series Switches without ELS Support

Example: Configuring a Private VLAN Spanning Multiple EX Series Switches

primary (VPLS Multihoming)

IN THIS SECTION

- Syntax | 13152
- Hierarchy Level | 13152
- Description | 13152

- Default | 13152
- Options | 13153
- Required Privilege Level | 13153
- Release Information | 13153

Syntax

```
primary interface-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-homing active-interface],
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-homing site site-name active-interface],
[edit routing-instances instance-name protocols vpls multi-homing active-interface],
[edit routing-instances instance-name protocols vpls multi-homing site site-name active-interface]
```

Description

Specify the name of the multihomed interface to be used as the primary interface by the VPLS site.

For FEC 128, use the `[edit routing-instances instance-name protocols vpls multi-homing active-interface]` hierarchy level.

Default

If you omit this statement, depending on the order in which interfaces are listed in the PE router's configuration, the first operational interface in the set of configured interfaces is chosen to be the primary interface.

NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Options

interface-name—Name of the interface (for example, `ge-0/1/0.1`).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Specifying an Interface as the Active Interface

any (VPLS Multihoming) | 2282

primary-wins

IN THIS SECTION

- Syntax | 13153
- Hierarchy Level | 13154
- Description | 13154
- Options | 13154
- Required Privilege Level | 13154
- Release Information | 13154

Syntax

```
primary-wins primary-wins;
```

Hierarchy Level

```
[edit access group-profile profile-name client client-name ppp],
[edit access profile profile-name ppp]
```

Description

Configure the primary Windows Internet name server.

Options

primary-wins—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

priority

IN THIS SECTION

- Syntax | 13155
- Hierarchy Level | 13155
- Description | 13155

- Options | 13155
- Required Privilege Level | 13155
- Release Information | 13156

Syntax

```
priority number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group
group-id],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-
group group-id]
```

Description

Configure a switch's priority for becoming the primary default routing platform. The routing platform with the highest priority within the group becomes the primary.

Options

number—Routing platform's priority for being elected to be the primary router in the VRRP group. A larger value indicates a higher priority for being elected.

- **Range:** 1 through 255
- **Default:** 100 (for backup routers)

NOTE: Priority 255 cannot be assigned to routed VLAN interfaces (RVIs).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring VRRP for IPv6 (CLI Procedure)

priority (Application Identification With Next Gen Services)

IN THIS SECTION

- Syntax | 13156
- Hierarchy Level | 13156
- Description | 13157
- Options | 13157
- Required Privilege Level | 13157
- Release Information | 13157

Syntax

```
priority (high | low);
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Description

Define the priority of signatures when both a custom signature and predefined signature apply to a protocol bundle. This statement is valid only if you are running Next Gen Services. If you are not running Next Gen Services, use the `order-priority` statement.

Options

high Custom signatures have priority over predefined signatures.

low Predefined signatures have priority over custom signatures.

- **Default:** high

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 19.3R2 and 19.4R1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

priority (Bootstrap)

IN THIS SECTION

- Syntax | 13158
- Hierarchy Level | 13158
- Description | 13158

- Options | 13158
- Required Privilege Level | 13158
- Release Information | 13159

Syntax

```
priority number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp bootstrap (inet | inet6)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp bootstrap (inet | inet6)],
[edit protocols pim rp bootstrap (inet | inet6)],
[edit routing-instances routing-instance-name protocols pim rp bootstrap (inet | inet6)]
```

Description

Configure the routing device's likelihood to be elected as the bootstrap router.

Options

number—Routing device's priority for becoming the bootstrap router. A higher value corresponds to a higher priority.

- **Range:** 0 through a 32-bit number
- **Default:** 0 (The routing device has the least likelihood of becoming the bootstrap router and sends packets with a priority of 0.)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring PIM Bootstrap Properties for IPv4

Configuring PIM Bootstrap Properties for IPv4 or IPv6

bootstrap-priority | 3019

priority (CFM MEP)

IN THIS SECTION

- Syntax | 13159
- Hierarchy Level | 13159
- Description | 13160
- Options | 13160
- Required Privilege Level | 13160
- Release Information | 13160

Syntax

```
priority value;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
action log-generate-ais]
```

Description

Configure the 802.1p priority of the AIS packet.

Options

value Priority level.

- **Range:** 0 through 7

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Alarm Indication

Ethernet Alarm Indication

priority (Chassis Synchronization Source)

IN THIS SECTION

- Syntax | 13161
- Hierarchy Level | 13161
- Description | 13161
- Options | 13161
- Required Privilege Level | 13161
- Release Information | 13161

Syntax

```
priority number;
```

Hierarchy Level

```
[edit chassis synchronization source (bits | gps)]
[edit chassis synchronization source interfaces]
```

Description

Configure a priority level for the interface specified as a candidate source. When not specified, `gps` has higher default priority than `bits`, and `bits` has higher default priority than other Gigabit Ethernet, 10-Gigabit Ethernet, T1, or E1 clock sources, which have the lowest default priority.

Options

number	Priority level between 1 and 5.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

External Clock Synchronization Overview for ACX Series Routers

Configuring External Clock Synchronization for ACX Series Routers

Synchronous Ethernet Overview

show chassis synchronization

priority (Clock Synchronization)

IN THIS SECTION

- Syntax | 13162
- Hierarchy Level | 13162
- Description | 13162
- Options | 13163
- Required Privilege Level | 13163
- Release Information | 13163

Syntax

```
priority number;
```

Hierarchy Level

```
[edit chassis synchronization source interfaces (external-a | external-b | interface interface-name)]
[edit chassis synchronization source interfaces (external | interface interface-name)]
[edit chassis synchronization source interfaces (external-0/0 | external-1/0 | interface interface-name)]
```

Description

Configure the priority of a clock source in relationship to other clock sources to define a network synchronization flow and to help prevent timing loops.

When the priority is not specified, the external-a interface has higher default priority than the external-b interface, and the external-b interface has higher default priority than other Gigabit Ethernet or 10-Gigabit Ethernet clock sources, which have the lowest default priority. Any priority you configure is higher than any default priority.

Options

number Set the priority level of the clock source.

- **Range:** 1 through 5

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

priority-cost (VRRP)

IN THIS SECTION

- Syntax | 13164
- Hierarchy Level | 13164
- Description | 13164
- Options | 13164
- Required Privilege Level | 13164
- Release Information | 13164

Syntax

```
priority-cost priority;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-
group group-id track interface interface-name],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-
group group-id track interface interface-name],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet address address vrrp-group group-id track interface interface-name],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number
family inet6 address address vrrp-inet6-group group-id track interface interface-name]
```

Description

Configure a Virtual Router Redundancy Protocol (VRRP) router's priority cost for becoming the primary default router. The router with the highest priority within the group becomes the primary.

Options

priority—The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new primary router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.

- **Range:** 1 through 254

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Metro Routers.

RELATED DOCUMENTATION

Configuring a Logical Interface to Be Tracked for a VRRP Group

priority (DDoS)

IN THIS SECTION

- Syntax | 13165
- Hierarchy Level | 13165
- Description | 13166
- Options | 13166
- Required Privilege Level | 13166
- Release Information | 13166

Syntax

```
priority level;
```

Hierarchy Level

- For MX Series routers, T4000 routers, and EX9200 switches:

```
[edit system ddos-protection protocols protocol-group  packet-type]
```

- For ACX Series routers, PTX Series routers, and QFX Series switches:

```
[edit system ddos-protection protocols protocol-group packet-type]
```

Description

Configure the priority for the packet type within the parent protocol group. In the event of downstream traffic congestion, high priority packets are provided bandwidth before medium priority packets. In turn, medium priority packets are provided bandwidth before low priority packets. Packets are dropped when there is insufficient available bandwidth.

> **NOTE**: PTX10003 and PTX10008 routers don't support this option.

Options

level—Priority of the packet type, low, medium, or high.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers

priority (Diameter Network Element)

IN THIS SECTION

- Syntax | 13167
- Hierarchy Level | 13167
- Description | 13167
- Options | 13167

Required Privilege Level | 13167
Release Information | 13167

Syntax

```
priority priority-value;
```

Hierarchy Level

```
[edit access diameter network-element element-name peer peer-name]
```

Description

Set the priority for a peer within a Diameter network element. A peer with a lower number has a higher priority. For load balancing, configure the peers with the same priority.

Options

priority-value Priority for the peer within the network element.

- **Range:** 1 through 65,535

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

priority (Diameter Peer)

IN THIS SECTION

- Syntax | 13168
- Hierarchy Level | 13168
- Description | 13168
- Options | 13168
- Required Privilege Level | 13169
- Release Information | 13169

Syntax

```
priority priority-value;
```

Hierarchy Level

```
[edit diameter network-element element-name peer peer-name]
```

Description

Set the priority for a peer within a Diameter network element. A peer with a higher number has a higher priority.

Options

priority-value—Priority for the peer within the network element.

- **Range:** 1 through 65535

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Network Elements

priority (Dynamic Schedulers)

IN THIS SECTION

- Syntax | 13169
- Hierarchy Level | 13170
- Description | 13170
- Options | 13170
- Required Privilege Level | 13170
- Release Information | 13170

Syntax

```
priority (priority-level | $junos-cos-scheduler-priority);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
```

Description

Specify packet-scheduling priority value in a dynamic profile.

Options

priority-level—one of the following packet-scheduling priority values:

- `low`—Scheduler has low priority.
- `medium-low`—Scheduler has medium-low priority.
- `medium-high`—Scheduler has medium-high priority.
- `high`—Scheduler has high priority. Assigning high priority to a queue prevents the queue from being underserved.
- `strict-high`—Scheduler has strictly high priority. Configure a `high` priority queue with unlimited transmission bandwidth available to it. As long as it has traffic to send, the `strict-high` priority queue receives precedence over `low`, `medium-low`, and `medium-high` priority queues, but not `high` priority queues. You can configure `strict-high` priority on only one queue per interface.

`$junos-cos-scheduler-priority`—Junos predefined variable that is replaced with the packet-scheduling priority value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

The `$junos-cos-scheduler-priority` predefined variable introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Schedulers in a Dynamic Profile for Subscriber Access

Dynamic Variables Overview

scheduler (Dynamic Scheduler Maps) | **16344**

priority (Fabric Priority)

IN THIS SECTION

- Syntax | **13171**
- Hierarchy Level | **13171**
- Description | **13171**
- Options | **13172**
- Required Privilege Level | **13172**
- Release Information | **13172**

Syntax

```
priority (high | low);
```

Hierarchy Level

```
[edit class-of-service forwarding-classes class class-name queue-num queue-number],
[edit class-of-service forwarding-classes queue queue-number class-name]
```

Description

For M320 routers, MX Series routers, T Series routers and EX Series switches only, specify a fabric priority value.

The two hierarchy levels are mutually exclusive. To configure up to eight forwarding classes with one-to-one mapping between forwarding classes and output queues, include this statement at the `[edit class-of-service forwarding-classes queue queue-number class-name]` hierarchy level. To configure up to 16 forwarding classes with multiple forwarding classes mapped to single queues, include this statement at the `[edit class-of-service forwarding-classes class class-name queue-num queue-number]` hierarchy level.

Options

`low`—Forwarding class's fabric queuing has low priority.

`high`—Forwarding class's fabric queuing has high priority.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`[edit class-of-service forwarding-classes class class-name queue-num queue-number]` hierarchy level added in Junos OS Release 8.1.

RELATED DOCUMENTATION

Forwarding Classes and Fabric Priority Queues

Configuring Up to 16 Custom Forwarding Classes

priority (Fabric Queues, Schedulers)

IN THIS SECTION

- Syntax | 13173
- Hierarchy Level | 13173
- Description | 13173

- Options | 13173
- Required Privilege Level | 13173
- Release Information | 13174

Syntax

```
priority (high | low)scheduler scheduler-name;
```

Hierarchy Level

```
[edit class-of-service fabric scheduler-map]
```

Description

Define Fabric traffic priority. For M320, MX Series, T Series routers and EX Series switches only, specify the fabric priority with which a scheduler is associated.

For a scheduler that you associate with a fabric priority, you cannot include the `buffer-size`, `transmit-rate`, or `priority` statements at the `[edit class-of-service schedulers scheduler-name]` hierarchy level.

On EX Series switches, this statement is supported only on EX8200 standalone switches and EX8200 Virtual Chassis.

Options

`high`—Scheduler has high priority.

`low`—Scheduler has low priority.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Associating Schedulers with Fabric Priorities

priority (FIP)

IN THIS SECTION

- Syntax | 13174
- Hierarchy Level | 13174
- Description | 13175
- Options | 13175
- Required Privilege Level | 13175
- Release Information | 13175

Syntax

```
priority priority;
```

Hierarchy Level

```
[edit fc-fabrics fc-fabric-name protocols fip],
[edit fc-fabrics fc-fabric-name protocols fip interface interface-name]
```

Description

Sets the global or interface-specific priority value associated with the switch FCF-MAC. CNAs use the priority value to determine the switch with which they will perform FIP FLOGI. The lower the value, the higher the priority. The switch advertises this value to the server ENodes on the FCoE network. A priority value set at the interface level overrides the global setting.

Options

priority —Value that determines the FCF an ENode selects to perform FIP FLOGI. The lower the priority number, the higher the priority of the FCF.

- **Range:** 0 through 255
- **Default:** 128

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fip | **19938**

show fibre-channel fip interface | **19961**

Overview of FIP

Configuring FIP on an FCoE-FC Gateway

priority-flow-control

IN THIS SECTION

- Syntax | 13176
- Hierarchy Level | 13176
- Description | 13176
- Options | 13176
- Required Privilege Level | 13177
- Release Information | 13177

Syntax

```
priority-flow-control {
    no-auto-negotiation;
}
```

Hierarchy Level

```
[edit protocols dcbx interface (all | interface-name)]
```

Description

Disable autonegotiation of priority-based flow control (PFC) on one or more Ethernet interfaces. Autonegotiation enables PFC on an interface only if the switch and the peer device connected to the switch both support PFC and have the same PFC configuration. Disabling autonegotiation on an interface forces the interface to use the PFC state (enabled or disabled) that is configured on the switch by the configuration and assignment of the congestion notification profile.

Options

`no-auto-negotiation`—Disable automatic negotiation of PFC.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- show dcbx neighbors | **19191**
- Configuring CoS PFC (Congestion Notification Profiles)
- Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)
- *Configuring DCBX Autonegotiation*
- *Example: Configuring CoS PFC for FCoE Traffic*
- Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches
- Understanding Priority-Based Flow Control
- *Understanding DCB Features and Requirements*

priority-hold-time

IN THIS SECTION

- Syntax | **13178**
- Hierarchy Level | **13178**
- Description | **13178**
- Options | **13178**
- Required Privilege Level | **13178**
- Release Information | **13178**

Syntax

```
priority-hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp)],
[edit protocols (mstp | rstp)],
[edit routing-instances routing-instance-name protocols (mstp | rstp)]
```

Description

Specify the number of seconds to hold before switching to the primary priority when the first core domain comes up.

Options

seconds—Number of seconds to hold before switching to primary priority.

- **Range:** 1 through 255
- **Default:** 2 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

priority-hold-time

IN THIS SECTION

- Syntax | 13179
- Hierarchy Level | 13179
- Description | 13179
- Options | 13180
- Required Privilege Level | 13180
- Release Information | 13180

Syntax

```
priority-hold-time seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track]
```

Description

Configure a Virtual Router Redundancy Protocol (VRRP) router's priority hold time to define the minimum length of time that must elapse between dynamic priority changes. If the dynamic priority changes because of a tracking event, the priority hold timer begins running. If another tracking event or manual configuration change occurs while the timer is running, the new dynamic priority update is postponed until the timer expires.

> **NOTE**: When the track feature is configured, and if VRRP should pre-empt due to the tracking interface or route transition, any configured pre-empt hold time will be ignored. VRRP primary will pre-empt according to the configuration of the priority-hold time.

Options

seconds—Minimum length of time that must elapse between dynamic priority changes.

- **Range:** 0through 3600 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring a Logical Interface to Be Tracked for a VRRP Group

priority (Management Traffic)

IN THIS SECTION

- Syntax | 13181
- Hierarchy Level | 13181
- Description | 13181
- Required Privilege Level | 13181
- Release Information | 13181

Syntax

```
priority (low | medium | high)
```

Hierarchy Level

```
[edit system management-traffic protocol protocol-name
```

Description

You can prioritize the internal traffic between the Routing Engine and the Packet Forwarding Engine or other Routing Engine based on the packet classification. This enhancement increases the internal and management traffic performance and provides better packet processing. In addition, this enhancement can introduce more operational commands which provide easily interpretable outputs for debugging packet path issues in the Routing Engine.

Before Junos OS Release 22.4R1, all internal traffic is considered as high priority. Starting in Junos OS Release 22.4R1, the applications in Packet Forwarding Engine and peer Routing Engine can change priority of a traffic using DSCP marking. Different platforms can use this packet prioritization feature based on their deployment scenarios and needs.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

show system statistics netisr qdrops | 27174

show system statistics netisr per-priority-stats | 27173

show system statistics netisr socket-drops | 27176

priority (OAM Connectivity-Fault Management)

IN THIS SECTION

- Syntax | 13182
- Hierarchy Level | 13182
- Description | 13182
- Options | 13182
- Required Privilege Level | 13183
- Release Information | 13183

Syntax

```
priority number;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-namemaintenance-
association ma-name mep mep-id]
```

For EX Series Switches:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domaindomain-namemaintenance-
association ma-name mepmep-id]
```

Description

IEEE 802.1p priority bits used by the continuity check messages.

Options

number—Configure the IEEE 802.1p priority bits to be used in the VLAN header of the CFM packets.

- **Range:** 0 through 7

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring Connectivity Fault Management (CFM)

priority-override (Event Policy)

IN THIS SECTION

- Syntax | 13183
- Hierarchy Level | 13184
- Description | 13184
- Required Privilege Level | 13184
- Release Information | 13184

Syntax

```
priority-override {
    facility facility-type;
    severity severity-level;
}
```

Hierarchy Level

```
[edit event-options policy policy-name then]
```

Description

Configure an event policy to override the default system log priority of the triggering event so that the system logs the event with a different facility type, severity level, or both. If you configure multiple event policies to override the priority of the same event, the event is logged based on the priority set by the last executed event policy to change it.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding the Event System Log Priority in an Event Policy

Example: Configure the Event System Log Priority in an Event Policy

Junos OS System Logging Facilities and Message Severity Levels

facility | **6769**

severity (Event Policy) | **16898**

priority (PIM Interfaces)

IN THIS SECTION

Syntax | **13185**

- Hierarchy Level | 13185
- Description | 13185
- Options | 13185
- Required Privilege Level | 13186
- Release Information | 13186

Syntax

```
priority number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM) interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM) interface-name],
[edit protocols pim interface (Protocols PIM) interface-name],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM) interface-
name]
```

Description

Configure the routing device's likelihood to be elected as the designated router. DR priority is specific to PIM sparse mode; as per RFC 3973, PIM DR priority cannot be configured explicitly in PIM Dense Mode (PIM-DM) in IGMPv2 – PIM-DM only support DRs with IGMPv1.

Options

number—Routing device's priority for becoming the designated router. A higher value corresponds to a higher priority.

- **Range:** 0 through 4294967295
- **Default:** 1 (Each routing device has an equal probability of becoming the DR.)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Interface Priority for PIM Designated Router Selection

priority (PIM RPs)

IN THIS SECTION

- Syntax | 13186
- Hierarchy Level | 13187
- Description | 13187
- Options | 13187
- Required Privilege Level | 13187
- Release Information | 13187

Syntax

```
priority number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp bidirectional address address],
[edit logical-systems logical-system-name routing-instances instance-name protocols pim rp
bidirectional address address],
[edit protocols pim rp bidirectional address address],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances instance-name protocols pim rp bidirectional address address],
[edit routing-instances routing-instance-name protocols pim rp local family (inet | inet6)]
```

Description

For PIM-SM, configure this routing device's priority for becoming an RP.

For bidirectional PIM, configure this RP address' priority for becoming an RP.

The bootstrap router uses this field when selecting the list of candidate rendezvous points to send in the bootstrap message. A smaller number increases the likelihood that the routing device or RP address becomes the RP. A priority value of 0 means that bootstrap router can override the group range being advertised by the candidate RP.

Options

number—Priority for becoming an RP. A lower value corresponds to a higher priority.

- **Range:** 0 through 255
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for bidirectional RP addresses introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Local PIM RPs

Example: Configuring Bidirectional PIM

priority (PoE)

IN THIS SECTION

- Syntax | 13188
- Hierarchy Level | 13188
- Description | 13188
- Default | 13189
- Options | 13189
- Required Privilege Level | 13189
- Release Information | 13189

Syntax

```
priority (high | low);
```

Hierarchy Level

```
[edit poe interface (all | interface-name)]
```

Description

Sets the priority of individual ports. When it is not possible to maintain power to all connected ports, lower-priority ports are powered off before higher priority ports. When a new device is connected on a higher-priority port, a lower-priority port will be powered off automatically if available power is insufficient to power on the higher-priority port. Note that for ports with the same priority configuration, ports on the left are given higher priority than the ports on the right.

Default

low

Options

value—high or low:

- `high`—Specify that this port is to be treated as high priority in terms of power allocation
- `low`—Specify that this port is to be treated as low priority in terms of power allocation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

priority (policy-options)

IN THIS SECTION

- Syntax | 13190
- Hierarchy Level | 13190
- Description | 13190
- Options | 13190
- Required Privilege Level | 13191
- Release Information | 13191

Syntax

```
policy-options {
    policy-statement policy-name {
        term term-name {
            from {
                match-conditions;
                route-filter destination-prefix match-type;
            }
        then {
            priority high | low;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name policy-options policy-statementpolicy-name term term-name then],
[edit logical-systems name policy-options policy-statementpolicy-name then],
[edit policy-options policy-statementpolicy-name term term-name then],
[edit policy-options policy-statementpolicy-name then]
```

Description

Sets the priority for route installation. You can choose a relative priority of `high`, `low` , or `medium` to ensure that high priority IGP and LDP routes are updated in the FIB (forwarding table) before medium and low priority routes. Routes are placed in different priority queues according to the priority. Any routes that are not explicitly assigned a priority are treated as medium priority. Within the same priority level, routes will continue to be updated in lexicographic order.

Options

high—Set priority to high.

low—Set priority to low.

medium—Set priority to medium.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Prefix Prioritization Overview

Example: Configuring the Priority for Route Prefixes in the RPD Infrastructure

priority (Power over Ethernet)

IN THIS SECTION

- Syntax | 13191
- Hierarchy Level | 13192
- Description | 13192
- Default | 13192
- Options | 13192
- Required Privilege Level | 13192
- Release Information | 13192

Syntax

```
priority (low | high);
```

Hierarchy Level

```
[edit poe interface (interface-name | all | all-extended)]
```

Description

Set the power priority for individual interfaces when there is insufficient power for all PoE interfaces. If the switch needs to shut down powered devices because PoE demand exceeds the PoE budget, low-priority devices are shut down before high-priority devices. Among interfaces that have the same assigned priority, priority is determined by port number, with lower-numbered ports having higher priority.

Default

`low`

Options

`high`—Specifies that this interface is to be treated as high-priority in terms of power allocation. If the switch needs to shut down powered devices because PoE demand exceeds the PoE budget, power is not shut down on this interface until it has been shut down on all the low-priority interfaces.

`low`—Specifies that this interface is to be treated as low-priority in terms of power allocation. If the switch needs to shut down powered devices because PoE demand exceeds the PoE budget, power is shut down on this interface before it is shut down on high-priority interfaces.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 17.2R1 for a Junos Fusion Provider Edge.

RELATED DOCUMENTATION

- Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch

Configuring PoE Interfaces on EX Series Switches

Configuring Power over Ethernet in a Junos Fusion

priority (Protocols MPLS)

IN THIS SECTION

- Syntax | 13193
- Hierarchy Level | 13193
- Description | 13193
- Options | 13194
- Required Privilege Level | 13194
- Release Information | 13194

Syntax

```
priority setup-priority reservation-priority;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

Configure the setup priority and reservation priority for an LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established

sessions on the link to determine whether some of them should be preempted to accommodate the new session. Sessions with lower hold priorities are preempted.

Options

reservation-priority—Reservation priority, used to keep a reservation after it has been set up. A smaller number has a higher priority. The priority must be greater than or equal to the setup priority to prevent preemption loops.

- **Range:** 0 through 7, where 0 is the highest and 7 is the lowest priority.
- **Default:** 0 (Once the session is set up, no other session can preempt it.)

setup-priority—Setup priority.

- **Range:** 0 through 7, where 0 is the highest and 7 is the lowest priority.
- **Default:** 7 (The session cannot preempt any existing sessions.)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Priority and Preemption for LSPs

priority (Protocols OAM)

IN THIS SECTION

Syntax | 13195

- Hierarchy Level | 13195
- Description | 13195
- Options | 13195
- Required Privilege Level | 13195
- Release Information | 13196

Syntax

```
priority priority-value;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile profile-namesla-iterator-profile]
```

Description

Configure the priority of the iterator profile, which is the `vlan-pcp` value that is sent in the Y.1731 data frames.

Options

`priority-value`—Priority value, which is the `vlan-pcp` value that is sent in the Y.1731 data frames.

- **Range:** 0 through 7
- **Default:** 0

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

sla-iterator-profile | 28012

Configuring a Remote MEP with an Iterator Profile

priority (Protocols Router Discovery)

IN THIS SECTION

- Syntax | 13196
- Hierarchy Level | 13196
- Description | 13197
- Options | 13197
- Required Privilege Level | 13197
- Release Information | 13197

Syntax

```
priority number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery address address],
[edit protocols router-discovery address address]
```

Description

Specify the preference of the address to become a default router. This preference is set relative to the preferences of other router addresses on the same subnet.

Options

number—Preference of the addresses for becoming the default router. A higher value indicates that the address has a greater preference for becoming the default router.

- **Range:** 0 through 0x80000000
- **Default:** 0 (This address has the least chance of becoming the default router.)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

priority (Protocols RSVP)

IN THIS SECTION

- Syntax | 13198
- Hierarchy Level | 13198
- Description | 13198

- Options | 13198
- Required Privilege Level | 13199
- Release Information | 13199

Syntax

```
priority setup-priority reservation-priority;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection bypass bypass-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Description

Configure the setup priority and reservation priority for a bypass LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. The session with the lower-hold priority is preempted.

Options

reservation-priority—Reservation priority, used to keep a reservation after it has been set up. A smaller number has a higher priority. The priority must be greater than or equal to the setup priority to prevent preemption loops.

- **Range:** 0 through 7, where 0 is the highest and 7 is the lowest priority.
- **Default:** 0 (Once the session is set up, no other session can preempt it.)

setup-priority—Setup priority.

- **Range:** 0 through 7, where 0 is the highest and 7 is the lowest priority.

- **Default:** 7 (The session cannot preempt any existing sessions.)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

priority (Protocols STP)

IN THIS SECTION

- Syntax | 13199
- Hierarchy Level | 13199
- Description | 13200
- Default | 13200
- Options | 13200
- Required Privilege Level | 13201
- Release Information | 13201

Syntax

```
priority interface-priority;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp | vstp) interface interface-name],
[edit logical-systems logical-system-name protocols mstp msti msti-id interface interface-name],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id interface interface-name],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp | vstp) interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name
protocols mstp msti msti-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id interface interface-name],
```

```
[edit protocols (mstp | rstp | vstp) interface interface-name],
[edit protocols mstp msti msti-id interface interface-name],
[edit protocols vstp vlan-id interface interface-name],
[edit protocols mstp ;interface (all | interface-name) arp-on-stp],
[edit protocols mstpmsti msti-id interface interface-name arp-on-stp],
[edit protocols rstp interface (all | interface-name) arp-on-stp],
[edit protocols stp interface (all | interface-name) arp-on-stp],
[edit protocols vstp vlan vlan-id interface (all | interface-name) arp-on-stp]
```

```
[edit routing-instances routing-instance-name protocols (mstp | rstp | vstp) interface interface-
name],
[edit routing-instances routing-instance-name protocols mstp msti msti-id interface interface-
name],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id interface interface-name]
```

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specifies the interface priority to control which interface is elected as the root port. The interface priority must be set in increments of 16.

Default

The default value is 128.

Options

priority—(Optional) Interface priority. The interface priority must be set in increments of 16.

- **Range:** 0 through 240

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding STP

show spanning-tree bridge | 26548

show spanning-tree interface | 26555

priority (Protocols VRRP)

IN THIS SECTION

- Syntax | 13201
- Hierarchy Level | 13202
- Description | 13202
- Options | 13202
- Required Privilege Level | 13202
- Release Information | 13203

Syntax

```
priority priority;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
```

Description

Configure a Virtual Router Redundancy Protocol (VRRP) device's priority for becoming the primary default device. The device with the highest priority within the group becomes the primary. VRRP is designed to eliminate the single point of failure inherent in the static default routed environment. VRRP specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router controlling the IP address(es) associated with a virtual router is called the Primary, and forwards packets sent to these IP addresses. The election process provides dynamic fail-over in the forwarding responsibility when the Primary become unavailable. Any of the virtual router's IP addresses on a LAN can then be used as the default first hop router by end-hosts. The advantage gained from using VRRP is a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end-host.

Options

priority Device's priority for being elected to be the primary device in the VRRP group. A larger value indicates a higher priority for being elected.

- **Range:** 0 through 255
- **Default:** 100. If two or more devices have the highest priority in the VRRP group, the device with the VRRP interface that has the highest IP address becomes the primary, and the others serve as backups.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring Basic VRRP Support
- *Understanding VRRP on SRX Series Devices*
- *Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces*

priority (RADIUS Network Elements)

IN THIS SECTION

- Syntax | 13203
- Hierarchy Level | 13203
- Description | 13204
- Options | 13204
- Required Privilege Level | 13204
- Release Information | 13204

Syntax

```
priority priority;
```

Hierarchy Level

```
[edit access radius network-element name server name]
```

Description

Configure a priority for each RADIUS server in the network element. You can have multiple servers with the same priority in a network element. All access requests are load balanced among the highest priority servers. If all the servers with the highest priority in the network element fail, then requests are load balanced among servers with the next highest priority level.

Options

priority Relative priority for a RADIUS server.

- **Range:** 1 through 16.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Network Elements

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

priority (Redundant Power System)

IN THIS SECTION

Syntax | 13205

Hierarchy Level | 13205

- Description | 13205
- Required Privilege Level | 13206
- Release Information | 13206

Syntax

```
priority (0|1|2|3|4|5|6);
```

Hierarchy Level

```
[edit redundant-power-system member]
[edit redundant-power-system member member-number]
```

Description

Configure the backup of any switch connected to the Redundant Power System (RPS) using the CLI on each switch. The determines the order in which the RPS supplies backup power to the switches connected to the RPS. 6 is the highest priority and 1 is lowest. Zero means off or no RPS backup.

If the switch is not reconfigured from the CLI, the default priority is 1. In this case, priority is determined by connector location with the rightmost connector having the highest priority.

For switches that can only be used as standalone switches, this hierarchy level is used for configuration:

```
[edit redundant-power-system]
```

For switches that can be used either as standalone switches or configured in a Virtual Chassis, this hierarchy level is used for configuration:

```
[edit redundant-power-system member vc-member-number]
```

If two or more connections are assigned the same , then the power of each connection is determined based on its switch connector port location, with the rightmost port receiving power first.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

priority (Schedulers)

IN THIS SECTION

- Syntax | 13206
- Hierarchy Level | 13206
- Description | 13206
- Options | 13207
- Required Privilege Level | 13207
- Release Information | 13207
- Platform-Specific Behavior | 13208

Syntax

```
priority priority-level;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Specify the packet-scheduling priority value.

Options

NOTE: Available options can vary by platform and software release.

priority-level can be one of the following:

- low—Scheduler has low priority.
- low-medium—Scheduler has low-medium priority.
- low-high—Scheduler has low-high priority.
- medium-low—Scheduler has medium-low priority.
- medium-high—Scheduler has medium-high priority.
- high—Scheduler has high priority. Assigning high priority to a queue prevents the queue from being underserved.
- strict-high—Scheduler has strictly high priority. Configure a high priority queue with unlimited transmission bandwidth available to it. As long as it has traffic to send, the strict-high priority queue receives precedence over low, medium-low, and medium-high priority queues, but not high priority queues. You can configure strict-high priority on only one queue per interface.
- low-latency—Scheduler has low latency queuing (LLQ) priority. LLQ enables delay-sensitive data to have preferential treatment over other traffic. A queue configured as a low latency queue has the highest priority over any other priority queues, including strict-high queues.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

low-latency support introduced for Junos OS Evolved Release 23.4R1.

low-high and low-mediumsupport introduced for Junos OS Evolved Release 24.4R1.

Platform-Specific Behavior

Use Feature Explorer to confirm platform and release support for scheduling priority.

Use the following table to review platform-specific behavior for your platforms:

ACX7000	ACX7000 routers support all eight levels of priority for port scheduling and six levels for HCoS scheduling (all except `low-medium` and `low-high`). ACX7000 routers also support multiple queues with the same priority.
EX4400	On EX4400 switches, applying strict-high priority schedulers to queues 0 through 3 also applies strict-high priority to queues 8 through 11. Therefore, Juniper recommends applying strict-high priority schedulers only to queues 4 through 7.
EX4600 and QFX5100	On QFX5100 and EX4600 switches, you can configure only one queue as a strict-high priority queue. We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. A shaping rate (shaper) sets the maximum amount of bandwidth a queue can consume. (Unlike using the transmit rate on a QFX10000 switch to limit traffic that receives strict-high priority treatment, traffic that exceeds the shaping rate is dropped, and is not treated as best-effort traffic that shares in excess bandwidth.) If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

(Continued)

QFX10000	On QFX10000 switches, you can configure as many strict-high priority queues as you want. On QFX10000 switches, we strongly recommend that you apply a transmit rate to strict-high priority queues to prevent them from starving other queues. A transmit rate configured on a strict-high priority queue limits the amount of traffic that receives strict-high priority treatment to the amount or percentage set by the transmit rate.

RELATED DOCUMENTATION

Configuring Schedulers for Priority Scheduling

priority1

IN THIS SECTION

- Syntax | 13209
- Hierarchy Level | 13210
- Description | 13210
- Options | 13210
- Required Privilege Level | 13210
- Release Information | 13210

Syntax

```
priority1 priority1-value;
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

Configure the priority as one of the following:

- In the timeReceiver, the priority value is set to select the best timeTransmitter clock. Note that in order to select a particular timeTransmitter clock, the priority value in the timeTransmitter clock's announce message must be equal to or lower than the configured *priority1-value*.
- In the timeTransmitter, the priority value is set to represent itself in the announce message to other timeReceiver clocks.
- In the boundary node, the timeReceiver uses this value to determine the best timeTransmitter clock, whereas the timeTransmitter uses this value from the announce message of the selected timeTransmitter clock.

Note that the lower value takes precedence.

Options

priority1-value—The priority value of the clock.

- **Range:** 0 through 255
- **Default:** 128

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Precision Time Protocol Overview

IEEE 1588v2 Precision Timing Protocol (PTP)

priority2

IN THIS SECTION

- Syntax | 13211
- Hierarchy Level | 13211
- Description | 13211
- Options | 13211
- Required Privilege Level | 13212
- Release Information | 13212

Syntax

```
priority2 priority2-value;
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

Configure the `priority2` value. This value is used to differentiate and prioritize the timeTransmitter clocks to avoid confusion when the *`priority1-value`* is the same for different timeTransmitter clocks in a network.

Note that the lower value takes precedence.

Options

`priority2-value`—The priority value of the clock.

- **Range:** 0 through 255
- **Default:** 128

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

IEEE 1588v2 Precision Timing Protocol (PTP)

privacy-3des

IN THIS SECTION

- Syntax | 13213
- Hierarchy Level | 13213
- Description | 13213
- Required Privilege Level | 13213
- Release Information | 13213

Syntax

```
privacy-3des {
    (privacy-key privacy-key | privacy-password privacy-password);
}
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]
```

Description

Configure the triple Data Encryption Standard (3DES) as the privacy type for the SNMPv3 user.

For privacy encryption, you can either enter a password in clear text, which is immediately encrypted, or, if you already have the encrypted version (known as a *key*), enter the key directly.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

privacy-aes128

IN THIS SECTION

Syntax | 13214

- Hierarchy Level | 13214
- Description | 13214
- Required Privilege Level | 13214
- Release Information | 13214

Syntax

```
privacy-aes128 {
    (privacy-key privacy-key | privacy-password privacy-password);
}
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]
```

Description

Configure the Advanced Encryption Standard encryption algorithm (CFB128-AES-128 Privacy Protocol) for the SNMPv3 user.

For privacy encryption, you can either enter a password in clear text, which is immediately encrypted, or, if you already have the encrypted version (known as a *key*), enter the key directly.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

privacy-des

IN THIS SECTION

- Syntax | 13215
- Hierarchy Level | 13215
- Description | 13215
- Required Privilege Level | 13215
- Release Information | 13216

Syntax

```
privacy-des {
    (privacy-key privacy-key | privacy-password privacy-password);
}
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]
```

Description

Configure the Data Encryption Standard (DES) as the privacy type for the SNMPv3 user.

For privacy encryption, you can either enter a password in clear text, which is immediately encrypted, or, if you already have the encrypted version (known as a *key*), enter the key directly.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

privacy-key

IN THIS SECTION

- Syntax | 13216
- Hierarchy Level | 13216
- Description | 13216
- Options | 13217
- Required Privilege Level | 13217
- Release Information | 13217

Syntax

```
privacy-key privacy-key;
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username privacy-3des],
[edit snmp v3 usm local-engine user username privacy-aes128],
[edit snmp v3 usm local-engine user username privacy-des],
[edit snmp v3 usm remote-engine engine-id user username privacy-3des],
[edit snmp v3 usm remote-engine engine-id user username privacy-aes128],
[edit snmp v3 usm remote-engine engine-id user username privacy-des]
```

Description

Configure a privacy key for the SNMPv3 user, if you know it. The encrypted version of a password is known as a key.

Options

privacy-key The encrypted version from the cleartext password. It is an alternative to entering a password.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

privacy-none

IN THIS SECTION

- Syntax | 13217
- Hierarchy Level | 13218
- Description | 13218
- Required Privilege Level | 13218
- Release Information | 13218

Syntax

```
privacy-none;
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]
```

Description

Configure that no encryption be used for the SNMPv3 user.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

privacy-password

IN THIS SECTION

- Syntax | 13219
- Hierarchy Level | 13219
- Description | 13219
- Options | 13219
- Required Privilege Level | 13219
- Release Information | 13219

Syntax

```
privacy-password privacy-password;
```

Hierarchy Level

```
[edit snmp v3 usm local-engine user username privacy-3des],
[edit snmp v3 usm local-engine user username privacy-aes128],
[edit snmp v3 usm local-engine user username privacy-des],
[edit snmp v3 usm remote-engine engine-id user username privacy-3des],
[edit snmp v3 usm remote-engine engine-id user username privacy-aes128],
[edit snmp v3 usm remote-engine engine-id user username privacy-des]
```

Description

Configure a privacy password for the SNMPv3 user.

Options

privacy-password—Password that a user enters. The password is then converted into a key that is used for encryption.

SNMPv3 has special requirements when you create plain-text passwords on a router or switch:

- The password must be at least eight characters long.
- The password can include alphabetic, numeric, and special characters, but it cannot include control characters.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

private-vlan

IN THIS SECTION

- Syntax | 13220
- Hierarchy Level | 13220
- Description | 13220
- Options | 13221
- Required Privilege Level | 13221
- Release Information | 13221

Syntax

```
private-vlan (isolated | community) vlan-id number;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Description

Configure a secondary VLAN (either an isolated VLAN or a community VLAN) within a private VLAN (PVLAN) and specify a VLAN ID for that secondary VLAN. This statement essentially converts a VLAN into a PVLAN, by carving out discrete subdomains (secondary VLANs) within the primary VLAN. You must specify a VLAN ID for each secondary PVLAN.

NOTE: After you have configured the secondary VLAN, you must also configure its association with a specific primary VLAN. See "isolated-vlan" on page 9407 and "community-vlan" on page 4477 for additional information.

Options

- `isolated` — The VLAN specified by *vlan-name* is defined as an *isolated* VLAN and a VLAN-ID is assigned to it. An isolated VLAN receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN. The VLAN name is optional. The VLAN ID is required.
- `community` — The VLAN specified by *vlan-name* is defined as community VLAN and a VLAN-ID is assigned to it. A *community* VLAN used to transport frames among members of a community, which is a subset of users within the VLAN, and to forward frames upstream to the primary VLAN. The VLAN name is optional. The VLAN ID is required.

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Creating a Private VLAN on a Single Switch with ELS Support (CLI Procedure)

Creating a Private VLAN Spanning Multiple EX Series Switches with ELS Support (CLI Procedure)

proactive-arp-detection

IN THIS SECTION

- Syntax | 13222
- Hierarchy Level | 13222
- Description | 13222
- Default | 13222
- Required Privilege Level | 13222
- Release Information | 13222

Syntax

```
proactive-arp-detection;
```

Hierarchy Level

```
[edit system arp]
```

Description

Enable proactive ARP detection to check the reachability of connected devices (within an IP subnet range) on a specified interface. After enabling this statement, you can set the ARP configurations at the interface level by the setting the: `host-discovery` *`address-range`*, `ageing-time-out` *`seconds`*, and `discovery-time-interval` *`seconds`* options at the [edit interfaces *interface-name* family inet address *ip-address*] hierarchy level. Likewise, you can delete the ARP configuration settings by using the `delete interfaces` *`interface-name`* `unit` *`unit`* family inet address *ip-address* host-discovery *address-range* command.

Default

This configuration statement is disabled by default.

Required Privilege Level

admin

Release Information

Statement introduced in Junos OS Release 18.3R1-S3, 18.3R2 and 19.3R1.

probe (ICMP Router Discovery)

IN THIS SECTION

Syntax | 13223

- Description | 13223
- Options | 13223
- Required Privilege Level | 13225
- Output Fields | 13225
- Sample Output | 13226
- Release Information | 13228

Syntax

```
probe <proxy-ip-address-or-hostname>
<by-address>
<by-index>
<by-name>
<by-remote-address>
<count>
<detail>
<inet>
<inet6>
<interface>
<interval>
<logical-system>
<routing-instance>
<source>
<tos>
<ttl>
<wait>
```

Description

Query the status of the probed interface.

Options

by-address (Optional) Probe a local interface by address.

by-index	(Optional) Probe a local interface by index.
by-name	(Optional) Probe a local interface by name.
by-remote-address	(Optional) Probe a remote interface by remote address.
count	(Optional) Specify the number of probe requests or packets sent out. • **Range:** 1 through 255
detail	(Optional) Display details of the incoming interface of the received packet.
inet	(Optional) Force probe to IPv4 network family destination.
inet6	(Optional) Force probe to IPv6 network family destination.
interface *interface-name*	(Optional) Specify the source interface to send the probe requests.
interval *seconds*	(Optional) Specify the time interval between the probe requests. • **Default:** 1 second
logical-system *logical-system-name*	(Optional) Specify the name of the logical system used to send the probe requests.
routing-instance *routing-instance-name*	(Optional) Specify the routing instance for the probe attempt.
source *source-address*	(Optional) Specify the source IP address of the probing interface.
tos *type-of-service*	(Optional) Specify the type-of-service (ToS) field in the IP header of the probe packets. • **Range:** 0 through 255
ttl *value*	(Optional) Specify the time-to-live (TTL) value for IPv4 packets and hop-limit value for IPv6 packets for the number of hops/devices to cross. • **Range:** 0 through 255
wait *seconds*	(Optional) Specify the maximum wait time in seconds after the final packet is sent. If this option is not specified, the default delay is 10 seconds. • **Default:** 10 seconds

Required Privilege Level

system-To view this statement in the configuration.

Output Fields

Table 1 lists the output fields for the `probe` command. The output fields are listed in the appropriate order in which they appear.

Table 100: probe Output Fields

Field Name	Field Description	Level of Output
icmp-seq	ICMP sequence number of the proxy IP address.	All levels
ttl	Time to live value (maximum IPv6 hop-limit-value) of the packets sent—0 to 255.	none
code	ICMP code values where—0 indicates No Error, 1 indicates Malformed Query, 2 indicates No Such Interface, 3 indicates No Such Table Entry, and 4 indicates Multiple Interfaces Satisfy Query.	All levels
state	State values are—0 indicates Reserved, 1 indicates Incomplete, 2 indicates Reachable, 3 indicates Stale, 4 indicates Delay, 5 indicates Probe, and 6 indicates Failed.	All levels
active	Inactive/active state (0 indicates inactive state and 1 indicates active state) of the IPv4 or IPv6 address configured.	All levels
IPv4	Unconfigured/configured state (0 indicates unconfigured and 1 indicates configured state) of IPv4 address.	All levels
IPv6	Unconfigured/configured state (0 indicates unconfigured and 1 indicates configured state) of IPv6 address.	All levels
time	Time taken time taken to receive the reply after the request is transmitted in milliseconds (ms).	All levels
hlim	Hop limit to the destination.	detail

Sample Output

Probing Proxy IP Address Using Interface Names of the Probed Interface that Resides in the Proxy Node

```
user@host> probe 2001:db8:1::10 by-name ge-0/0/1.0 count 1
PROBE6(74=40+8+26 bytes) 2001:db8:1::10 --> 2001:db8:1::11
34 bytes from 2001:db8:1::11, icmp_seq=0 hlim=255 code=0 state=0 active=1 IPv4=1 IPv6=1
time=4.750 ms
--- 2001:db8:1::10 probe6 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 0.231/0.231/0.231/0.000 ms
```

Probing Proxy IP Address Using Interface-index of the Probed Interface that Resides in the Proxy Node

```
user@host> probe 10.0.1.10 by-index 333 count 2
PROBE 10.0.1.10 (10.0.1.10):
28 bytes from 10.0.1.10: icmp_seq=0 ttl=255 code=0 state=0 active=1 IPv4=1 IPv6=1 time=6.767 ms
28 bytes from 10.0.1.10: icmp_seq=1 ttl=255 code=0 state=0 active=1 IPv4=1 IPv6=1 time=3.796 ms

--- 10.0.1.10 probe statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.796/5.282/6.767/1.485 ms
```

Probing Proxy IP Address Using IPv4 or IPv6 Address Configured in the Probed Interface on the Proxy Node

```
user@host> probe 2001:db8:1::11 by-address 10.0.12.1 count 1
PROBE6(72=40+8+24 bytes) 2001:db8:1::10 --> 2001:db8:1::11
32 bytes from 2001:db8:1::11, icmp_seq=0 hlim=255 code=0 state=0 active=1 IPv4=1 IPv6=1
time=6.443 ms

--- 2001:db8:1::11 probe6 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 6.443/6.443/6.443/0.000 ms
```

Probing the Proxy IP Address Configured on the Proxy Node for Information About the Probed IP Address Configured on the Probed Node

```
user@host> probe 10.0.1.10 by-remote-address 10.0.12.22 count 1
PROBE 10.0.1.10 (10.0.1.10):
32 bytes from 10.0.1.10: icmp_seq=0 ttl=255 code=3 state=0 active=0 IPv4=0 IPv6=0 time=5.054 ms

--- 10.0.1.10 probe statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 5.054/5.054/5.054/0.000 ms
```

Probing the Proxy IP Address for Information About the Probed Interface

```
user@host> probe 10.0.1.10 by-address 10.0.12.21 count 1
PROBE 10.0.1.10 (10.0.1.10):
32 bytes from 10.0.1.10: icmp_seq=0 ttl=255 code=2 state=0 active=0 IPv4=0 IPv6=0 time=4.569 ms

--- 10.0.1.10 probe statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 4.569/4.569/4.569/0.000 ms
```

Probing the Proxy IP Address for Information About the Probed Interface Details Using the Probed Interface Name That Resides on the Proxy Node

```
user@host> probe 2001:db8:1::11 by-name ge-0/0/1.0 detail
PROBE6(74=40+8+26 bytes) 2001:db8:1::10 --> 2001:db8:1::11
34 bytes from 2001:db8:1::11 via ge-0/0/0.0: icmp_seq=0 hlim=255 code=0 state=0 active=1 IPv4=1
IPv6=1 time=10.021 ms
34 bytes from 2001:db8:1::11 via ge-0/0/0.0: icmp_seq=1 hlim=255 code=0 state=0 active=1 IPv4=1
IPv6=1 time=37.991 ms
34 bytes from 2001:db8:1::11 via ge-0/0/0.0: icmp_seq=2 hlim=255 code=0 state=0 active=1 IPv4=1
IPv6=1 time=3.705 ms
34 bytes from 2001:db8:1::11 via ge-0/0/0.0: icmp_seq=3 hlim=255 code=0 state=0 active=1 IPv4=1
IPv6=1 time=7.961 ms
34 bytes from 2001:db8:1::11 via ge-0/0/0.0: icmp_seq=4 hlim=255 code=0 state=0 active=1 IPv4=1
IPv6=1 time=5.936 ms
^C
--- 2001:db8:1::11 probe6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

```
round-trip min/avg/max/std-dev = 3.705/13.123/37.991/12.610 ms
```

Release Information

Command introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

> *How to Use the Probe Command*
>
> extended-echo | **6697**

probe

IN THIS SECTION

- Syntax | **13228**
- Hierarchy Level | **13230**
- Description | **13230**
- Options | **13230**
- Required Privilege Level | **13230**
- Release Information | **13231**

Syntax

```
probe owner {
    test test-name {
        data-fill data;
        data-size size;
        delegate-probes probes;
        destination-interface interface-name;
        destination-port port;
        dscp-code-point dscp-bits;
```

```
        history-size size;
        inet6-options source-address ipv6-address;
        moving-average-size number;
        next-hop next-hop;
        offload-type {
            none;
            pfe-timestamp;
        }
        probe-count count;
        probe-interval seconds;
        probe-type type;
        routing-instance instance-name;
        rpm-scale {
            destination {
                interface interface-name.logical-unit-number;
                subunit-cnt subunit-cnt;
            }
            source {
                address-base ipv4-address-base;
                count ipv4-count;
                step ipv4-step;
            }
            source-inet6 {
                address-base ipv6-address-base;
                count ipv6-count;
                step ipv6-step;
            }
            target {
                address-base ipv4-address-base;
                count ipv4-count;
                step ipv4-step;
            }
            target-inet6 {
                address-base ipv6-address-base;
                count ipv6-count;
                step ipv6-step;
            }
            tests-count tests-count;
        }
        source-address address;
        target (url url | address ipv4-address | inet6-url url | inet6-address ipv6-address);
        test-interval interval;
        thresholds
```

```
        {
            egress-time microseconds;
            ingress-time microseconds;
            jitter-egress microseconds;
            jitter-ingress microseconds;
            jitter-rtt microseconds;
            rtt microseconds;
            std-dev-egress microseconds;
            std-dev-ingress microseconds;
            std-dev-rtt microseconds;
            successive-loss count;
            total-loss count;
        }
        traps [trap-names];
        ttl [hop-count];
    }
}
```

Hierarchy Level

```
[edit services rpm]
```

Description

Specify an owner name. The owner name combined with the test name represent a single RPM configuration instance.

Options

owner—Owner name up to 32 characters in length.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

probe-count

IN THIS SECTION

- Syntax | 13231
- Junos OS Hierarchy Level | 13231
- Junos OS Evolved Hierarchy Level | 13231
- Description | 13232
- Options | 13232
- Required Privilege Level | 13232
- Release Information | 13232

Syntax

```
probe-count count;
```

Junos OS Hierarchy Level

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name test-session session-
name],
[edit protocols isis interface interface-name], [edit protocols ospf interface interface-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the number of probes within a test.

Options

count 1 through 15 for RPM, for TWAMP 1 through 4,294,967,290.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [`edit services rpm twamp client control-connection` *`control-client-name`*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

Support at the [`edit protocols isis interface` *`interface-name`*] and [`edit protocols ospf interface` *`interface-name`*] hierarchy levels introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Configuring BGP Neighbor Discovery Through RPM

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

probe-interval

IN THIS SECTION

- Syntax | 13233
- Junos OS Hierarchy Level | 13233
- Junos OS Evolved Hierarchy Level | 13233
- Description | 13233
- Options | 13234
- Required Privilege Level | 13234
- Release Information | 13234

Syntax

```
probe-interval interval;
```

Junos OS Hierarchy Level

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name test-session session-
name],
[edit protocols isis interface interface-name], [edit protocols ospf interface interface-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the time to wait between sending packets, in seconds.

Options

interval—Number of seconds, from 1 through 255.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [`edit services rpm twamp client control-connection` *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

Support at the [`edit protocols isis interface` *interface-name*] and [`edit protocols ospf interface` *interface-name*] hierarchy levels introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Configuring BGP Neighbor Discovery Through RPM

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

Understanding Two-Way Active Measurement Protocol on Routers

probe-limit

IN THIS SECTION

- Syntax | 13235
- Hierarchy Level | 13235
- Description | 13235

- Options | 13235
- Required Privilege Level | 13235
- Release Information | 13236

Syntax

```
probe-limit limit;
```

Hierarchy Level

```
[edit services rpm]
```

Description

Configure the maximum number of concurrent probes allowed.

Options

limit—Maximum number of concurrent probes allowed.

- **Default:** 100 concurrent probes
- **Range:** (MX Series routers only) Starting in Junos OS Release 17.2R2 and 17.3R1, 1 through 2000 concurrent probes. In Junos OS releases earlier than 17.2R1, the range is 1 through 500 concurrent probes.
- **Range:** (PTX Series routers only) 1 through 200 concurrent probes
- **Range:** (SRX1500, SRX1600, SRX2300, and SRX4300 firewalls only) Starting in Junos Release 24.2R1, 1 through 2000 concurrent probes. In Junos OS releases earlier than 24.2R1, the range is 1 through 500 concurrent probes.
- **Range:** (Other platforms) 1 through 500 concurrent probes.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Limiting the Number of Concurrent RPM Probes on M, MX, T and PTX Routers and EX Series Switches*

probe-server

IN THIS SECTION

- Junos OS | **13236**
- Junos OS Evolved | **13237**
- Junos OS Hierarchy Level | **13237**
- Junos OS Evolved Hierarchy Level | **13237**
- Description | **13237**
- Required Privilege Level | **13238**
- Release Information | **13238**

Junos OS

```
probe-server {
    icmp {
        destination-interface destination-interface;
    }
    tcp {
        destination-interface interface-name;
        port number;
    }
    udp {
```

```
        destination-interface interface-name;
        port number;
    }
}
```

Junos OS Evolved

```
probe-server {
    icmp;
    tcp {
        port number;
    }
    udp {
        port number;
    }
}
```

Junos OS Hierarchy Level

```
[edit services rpm]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm]
```

Description

Enable the RPM server to act as a receiver for the probes, specified per protocol.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

> **NOTE**: The `destination-interface` statement is not supported on PTX Series routers or for Junos OS Evolved.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Junos OS Evolved Syntax and hierarchy level introduced in Junos OS Evolved Release 20.1R1.

`tcp` option introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Configuring RPM Receiver Servers

probe-type

IN THIS SECTION

- Syntax | 13238
- Junos OS Hierarchy Levels | 13239
- Junos OS Evolved Hierarchy Level | 13239
- Description | 13239
- Options | 13239
- Required Privilege Level | 13239
- Release Information | 13240

Syntax

```
probe-type type;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the packet and protocol contents of a probe.

Options

type—One of the following probe type values:

- `http-get`—(Not available at the `[edit services rpm bgp]` hierarchy level for Junos OS.) Sends a Hypertext Transfer Protocol (HTTP) get request to a target URL. For Junos OS Evolved, you must set the `offload-type` statement to `none` to configure this probe type.
- `http-metadata-get`—(Not available at the `[edit services rpm bgp]` hierarchy level for Junos OS.) Sends an HTTP get request for metadata to a target URL. For Junos OS Evolved, you must set the `offload-type` statement to `none` to configure this probe type.
- `icmp-ping`—Sends ICMP echo requests to a target address.
- `icmp-ping-timestamp`—Sends ICMP timestamp requests to a target address.
- `tcp-ping`—Sends TCP packets to a target. For Junos OS Evolved, you must set the `offload-type` statement to `none` to configure this probe type.
- `udp-ping`—(Junos OS only) Sends UDP packets to a target.
- `udp-ping-timestamp`—Sends UDP timestamp requests to a target address.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

`tcp-ping`, `http-get`, and `http-metadata-get` probe types introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Configuring BGP Neighbor Discovery Through RPM

process-inform

IN THIS SECTION

- Syntax | 13240
- Hierarchy Level | 13241
- Description | 13242
- Default | 13242
- Required Privilege Level | 13242
- Release Information | 13242

Syntax

```
process-inform {
    pool pool-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name interface interface-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name interface interface-name overrides],
[edit logical-systems logical-system-name system services dhcp-local-server overrides],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 overrides],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name overrides],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name interface interface-name overrides],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
overrides],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
interface interface-name overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name interface interface-name overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
overrides],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
interface interface-name overrides],
[edit system services dhcp-local-server overrides],
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name
overrides],
[edit system services dhcp-local-server group group-name overrides],
[edit system services dhcp-local-server group group-name interface interface-name overrides]
```

Description

Enable the processing of DHCP information request messages (DHCPINFORM for DHCPv4 and INFORMATION-REQUEST for DHCPv6) sent from the client to request DHCP options. For DHCP local servers, the messages are also passed to the configured server list.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Default

Information request messages are not processed.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Enabling Processing of Client Information Requests

Overriding the Default DHCP Local Server Configuration Settings

process (Memory)

IN THIS SECTION

- Syntax | 13243
- Hierarchy Level | 13243
- Description | 13243
- Options | 13244

Required Privilege Level | 13244
Release Information | 13244

Syntax

```
process {
    frequency (minutes | hours | days) value;
    minor-event {
        threshold process name memory-limit value ;
        action snmp-trap;
    }
    major-event {
        threshold process name memory-limit value ;
        action [snmp-trap | live-core | (terminate-process| terminate-process-force)];
    }
    critical-event {
        threshold process name memory-limit value ;
        action [snmp-trap | live-core | (terminate-process| terminate-process-force)];
    }
}
```

Hierarchy Level

```
[edit system monitor memory]
```

Description

Configure process level memory monitoring. Because you cannot specify a threshold for all processes, an algorithm (heuristics-based) is used to identify the processes that are likely to cross the threshold. The algorithm can identify the processes based on learning the memory growth pattern of the processes and comparing the memory usage of the processes with the baseline of the process. The algorithm measures the virtual memory footprint, increase in memory foot print since baseline, and memory growth pattern score to identify the processes. You can also specify exceptions. Learning is always on at the configured time interval for process monitoring. Also, prediction is possible only when the module is aware that the system is running low on memory.

The remaining statements are explained separately.

Options

frequency Specify the frequency of system level memory monitoring. You can specify the frequency in days, hours or minutes. Possible values: 1 through 65535.

minor-event Indicates a non-critical event that, if left unchecked, may cause an interruption in service or degradation in performance but does not significantly affect the device.

major-event Indicates an event that, if left unchecked, may result in significant loss of service or interruption of traffic.

critical-event Indicates an event that, if left unchecked, may result in serious loss of service or interruption of traffic. In case a critical event is triggered, you can terminate the process but only if the device is running out of memory.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1.

RELATED DOCUMENTATION

System Log Overview

System Log Explorer

process-monitor

IN THIS SECTION

- Syntax | 13245
- Hierarchy Level | 13245
- Description | 13245
- Options | 13246
- Required Privilege Level | 13246
- Release Information | 13246

Syntax

```
process-monitor {
    disable;
    traceoptions {
        file filename files number match regex size size (world-readable | no-world-readable);
        flag flag;
        level level;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Configure tracing options for the process health monitor process (pmond).

> **NOTE**: Starting with Junos OS Release 15.1R2, the pmond process is enabled by default on the Routing Engines of MX Series routers, even when no service interfaces are configured.

Options

`disable`—Disable the process health monitor process.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

traceoptions (Process Monitor) | 29648

process-non-null-as-null-register

IN THIS SECTION

- Syntax | 13247
- Hierarchy Level | 13247
- Description | 13247
- Required Privilege Level | 13248
- Release Information | 13248

Syntax

```
process-non-null-as-null-register;
```

Hierarchy Level

```
[edit protocols pim rp local]
```

Description

When `process-non-null-as-null-register` is enabled on a PTX device serving as PIM Rendezvous Point (RP) for multicast traffic, it allows the device to treat non-null registers, such as may be sent from any first hop router (FHR), as null registers, and thus to form a register state with the device. This statement is required when RP is enabled on PTX devices running Junos OS Evolved.

> **NOTE**: Currently supported PTX devices include PTX10003, PTX10004, PTX10008, PTX10016, and PTX10001-36MR.

More Information

In typical operation, for PIM any-source multicast (ASM), all *,G PIM joins travel hop-by-hop towards the RP, where they ultimately end. When the FHR receives its first traffic, it forms a register state with the RP in the network for the corresponding S,G. It does this by sending a PIM non-null register to form a multicast route with the downstream encapsulation interface. The RP decapsulates the non-null register and forms a multicast route with the upstream decapsulation device. In this way, multicast data traffic flows across the encapsulation/decapsulation tunnel interface, from the FHR to the RP, to all the downstream receivers until the RP has formed the S,G multicast tree in the direction of the source.

Without `process-non-null-as-null-register` enabled, for PIM ASM, PTX devices can only act as a PIM transit router or last hop router. These devices can receive a PIM join from downstream interfaces and propagate the joins towards the RP, or they can receive an IGMP/MLD join and propagate it towards a PIM RP, but they cannot act as a PIM RP itself. Nor can they form a register state machine with the PIM FHR in the network.

> **NOTE**: When you enable `process-non-null-as-null-register` on a PTX device, an initial drop in traffic is observed while the multicast tree is setup from the PIM RP to the FHR.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 19.3R1.

RELATED DOCUMENTATION

Configuring Local PIM RPs

processes

IN THIS SECTION

- Syntax | 13248
- Hierarchy Level | 13249
- Description | 13249
- Default | 13249
- Options | 13249
- Required Privilege Level | 13250
- Release Information | 13250

Syntax

```
processes {
    process-name (enable | disable) failover (alternate-media | other-routing-engine);
    timeout seconds;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure which Junos OS processes are running on the router or switch.

> **CAUTION**: Never disable any of the software processes unless instructed to do so by a customer support engineer.

Default

All processes are enabled by default.

Options

(enable | disable) (Optional) Enable or disable a specified process.

failover (alternate-media | other-routing-engine) (Optional) For routers or switches with redundant Routing Engines only, switch to backup media if a process fails repeatedly. If a process fails four times within 30 seconds, the router or switch reboots from the alternate media or the other Routing Engine.

process-name One of the valid process names. You can obtain a complete list of process names by using the CLI command completion feature. After specifying a process name, command completion also indicates any additional options for that process.

timeout *seconds* (Optional) How often the system checks the watchdog timer, in seconds. If the watchdog timer has not been checked in the specified number of seconds, the system reloads. If you set the time value too low, it is possible for the system to reboot immediately after it loads.

- **Values:** 15, 60, or 180
- **Default:** 180 seconds (rounded up to 291 seconds by the Junos kernel)

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Disabling Junos OS Processes

product-model (Junos Fusion)

IN THIS SECTION

- Syntax | 13250
- Hierarchy Level | 13251
- Description | 13251
- Options | 13251
- Required Privilege Level | 13251
- Release Information | 13252

Syntax

```
product-model model-name;
```

Hierarchy Level

```
[edit policy-options satellite-policies candidate-uplink-port-policy policy-name term term-name
from]
[edit policy-options satellite-policies candidate-uplink-port-policy policy-name term term-name
from]
[edit policy-options satellite-policies forwarding-policy policy-name term term-name from]
```

Description

Define the satellite device product models that will use the candidate uplink port policy defined in the `from` statement.

The other statements in the same `from` statement are applied to satellite devices that match the `product-model` *`model-name`* definition. Those configuration statements are not applied to satellite devices that do not match the definition.

Options

model-name Defines the satellite device product models that will use the candidate uplink port policy. It can be a complete device model name, to match that device model exactly. You can also use the wildcard character (*) in the match term to match zero or more of any character.

Some examples of using the wildcard character in the match term:

- To apply the satellite policy to all EX4300 switches in the satellite device role, enter `EX4300*` as the *model-name*.
- To apply the satellite policy to all QFX5100 switches in the satellite device role, enter `QFX5100*` as the *model-name*.
- To apply the satellite policy to QFX5100 switches with model names that start with QFX5100-96, enter `QFX5100-96*` as the *model-name*.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

product-name

IN THIS SECTION

- Syntax | 13252
- Hierarchy Level | 13252
- Description | 13253
- Options | 13253
- Required Privilege Level | 13253
- Release Information | 13253

Syntax

```
product-name name;
```

Hierarchy Level

```
[edit access diameter]
```

Description

Configure the product name that is advertised in the Capabilities-Exchange-Request or Capabilities-Exchange-Answer message.

Options

name Name of product that is the advertised value of the Product-Name AVP.

- **Default:** Juniper Diameter Client

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

products

IN THIS SECTION

- Syntax | 13254
- Hierarchy Level | 13254
- Description | 13254
- Options | 13254
- Required Privilege Level | 13254
- Release Information | 13254

Syntax

```
products {
    values [product-value];
}
```

Hierarchy Level

```
[edit security idp dynamic-attack-group dynamic-attack-group-name filters]
```

Description

Specify a products filter to add attack objects based on the application that is vulnerable to the attack.

Options

`values`—Name of the products filter. You can configure multiple filters separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

profile (Application Identification)

IN THIS SECTION

Syntax | 13255

- Hierarchy Level | 13255
- Description | 13255
- Options | 13255
- Required Privilege Level | 13255
- Release Information | 13255

Syntax

```
profile profile-name {
    rule-set rule-set-name;
}
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Define members of the application profile, which consists of one or more rule sets.

Options

profile-name—Identifier for the application profile.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support added in Junos OS release 19.3R2 and 19.4R1 for Next Gen Services on MX240, MX480, and MX960.

RELATED DOCUMENTATION

Configuring Application Profiles

profile

IN THIS SECTION

- Syntax | 13256
- Hierarchy Level | 13256
- Description | 13256
- Options | 13257
- Required Privilege Level | 13257
- Release Information | 13257

Syntax

```
profile profile-name;
```

Hierarchy Level

```
[edit access]
```

Description

Configure an access profile that specifies authentication and accounting parameters. Authentication and accounting parameters run only if you specify the profile name. The access profile contains the entire

authentication, authorization, and accounting (AAA) configuration that aids in handling AAA requests, authentication method and order, AAA server addresses, and AAA accounting.

Options

profile-name	Name of the access profile.

Required Privilege Level

admin—To view this statement in the configuration.admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1 at tenant system level.

RELATED DOCUMENTATION

Firewall Authentication for Tenant Systems

profile

IN THIS SECTION

- Syntax | 13258
- Hierarchy Level | 13258
- Description | 13258
- Default | 13258
- Options | 13258
- Required Privilege Level | 13259
- Release Information | 13259

Syntax

```
profile profile-name {
     accounting {
          accounting-stop-on-access-deny;
          accounting-stop-on-failure;
          order (radius | [ accounting-order-data-list ];
    }
     authentication-order  [authentication-method];
     radius {
          (accounting-server [server-addresses] | accounting-server-name hostname);
          (authentication-server [server-addresses] | authentication-server-name hostname);
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure an access profile. The access profile contains the entire authentication, authorization, and accounting (AAA) configuration that aids in handling AAA requests, including the authentication method and order, AAA server addresses, and AAA accounting.

Default

Not enabled.

Options

profile-name Profile name of up to 32 characters.

The remaining statements are explained separately. See CLI Explorer.

NOTE: The `[edit access]` hierarchy is not available on QFabric systems.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch

Configuring 802.1X RADIUS Accounting (CLI Procedure)

profile (Access)

IN THIS SECTION

- Syntax | 13259
- Hierarchy Level | 13261
- Description | 13261
- Options | 13261
- Required Privilege Level | 13263
- Release Information | 13263

Syntax

```
profile profile-name  {
    accounting {
        accounting-stop-on-access-deny;
        accounting-stop-on-failure;
        duplication;
        duplication-attribute-format;
```

```
        duplication-filter;
        duplication-vrf;
        order [accounting-method];
        statistics (time | volume-time);
    }
    address-assignment {
        inet6-pool inet6-pool-name;
        pool pool-name;
    }
    authentication-order (ldap | none | password | radius | s6a | securid);
    charging-service-list;
    client client-name  {
        chap-secret  chap-secret;
        client-group [ group-names ];
        firewall-user {
            password password;
        }
        no-rfc2486;
        pap-password pap-password;
        x-auth ip-address;
    }
    client-name-filter {
        count number;
        domain-name domain-name;
        separator special-character;
    }
    domain-name-server name;
    domain-name-server-inet name;
    domain-name-server-inet6 name;
    jsrc;
    ldap-options {
        assemble {
            common-name common-name;
        }
        base-distinguished-name base-distinguished-name;
        revert-interval seconds;
        search {
            admin-search {
                distinguished-name distinguished-name;
                password password;
            }
            search-filter search-filter-name;
        }
```

```
    }
    ldap-server server-address {
        port port-number;
        retry attempts;
        routing-instance routing-instance-name;
        source-address source-address;
        timeout seconds;
    }
    provisioning-order (gx-plus | jsrc);
    radius;
    radius-options;
    radius-server;
    session-limit-per-username;
    session-options {
        client-group [group-name];
        client-idle-timeout minutes;
        client-session-timeout minutes;
    }
    subscriber;
    wins-server;
}
```

Hierarchy Level

```
[edit access]
```

Description

Create a profile containing a set of attributes that define device management access.

Options

name	Profile name
accounting	Specifies the accounting options
address-assignment	Specify the address assignment pool
authentication-order	Order in which authentication mechanisms are used

- Values:
  - ldap—Light weight directory access protocol
  - none—No authentication performed
  - password—Locally configured password in access profile
  - radius—Remote authentication dial-in user service
  - s6a—S6a authentication
  - securid—RSA secure ID authentication

charging-service-list List of used 3gpp charging services

- Values:
  - ocs—Online charging service

client Entity requesting access

client-name-filter Restrictions on client names authenticated on this server

domain-name-server Default DNS server's IPv4 address

domain-name-server-inet DNS server's IPv4 address

domain-name-server-inet6 DNS server's IPv6 address

jsrc Set of JSRC configurations

ldap-options Light weight directory access protocol options

ldap-server Light weight directory access protocol server

preauthentication-order Order in which pre authentication mechanisms are used

- Values:
  - radius—Remote Authentication Dial-In User Service

radius Set of RADIUS configurations

radius-options RADIUS options

radius-server RADIUS server configuration

session-limit-per-username	Maximum number of sessions allowed per username • **Range:** 1 through 16
session-options	Options for an authenticated client's session
subscriber	Locally authenticated subscriber configuration
wins-server	Default WINS server's IPv4 address

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`inet6-pool` option is introduced in Junos OS Release 20.3R1.

`none` option is introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Understanding Interfaces

Understanding User Authentication for Security Devices

Ethernet Switching and Layer 2 Transparent Mode Overview

profile (Access)

IN THIS SECTION

- Syntax | **13264**
- Hierarchy Level | **13270**
- Description | **13270**

- Options | 13270
- Required Privilege Level | 13270
- Release Information | 13270

Syntax

```
profile profile-name {
    accounting {
        address-change-immediate-update
        accounting-stop-on-access-deny;
        accounting-stop-on-failure;
        ancp-speed-change-immediate-update;
        coa-immediate-update;
        coa-no-override service-class-attribute;
        duplication;
        duplication-filter;
        duplication-vrf {
            access-profile-name profile-name;
            vrf-name vrf-name;
        }
        immediate-update;
        order [ accounting-method ];
        send-acct-status-on-config-change;
        statistics (time | volume-time);
        update-interval minutes;
        wait-for-acct-on-ack;
    }
    accounting-order (radius | [accounting-order-data-list]);
    authentication-order [ authentication-methods ];
    client client-name {
        chap-secret chap-secret;
        group-profile profile-name;
        ike {
            allowed-proxy-pair {
                remote remote-proxy-address local local-proxy-address;
            }
            pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
            ike-policy policy-name;
```

```
            interface-id string-value;
        }
        l2tp {
            aaa-access-profile profile-name;
            interface-id interface-id;
            lcp-renegotiation;
            local-chap;
            maximum-sessions number;
            maximum-sessions-per-tunnel number;
            multilink {
                drop-timeout milliseconds;
                fragment-threshold bytes;
            }
            override-result-code session-out-of-resource;
            ppp-authentication (chap | pap);
            ppp-profile profile-name;
            service-profile profile-name(parameter)&profile-name;
            sessions-limit-group limit-group-name;
            shared-secret shared-secret;
        }
        pap-password pap-password;
        ppp {
            cell-overhead;
            encapsulation-overhead bytes;
            framed-ip-address ip-address;
            framed-pool framed-pool;
            idle-timeout seconds;
            interface-id interface-id;
            keepalive seconds;
            primary-dns primary-dns;
            primary-wins primary-wins;
            secondary-dns secondary-dns;
            secondary-wins secondary-wins;
        }
        user-group-profile profile-name;
    }
    domain-name-server;
    domain-name-server-inet;
    domain-name-server-inet6;
    local {
        flat-file-profile profile-name;
    }
    preauthentication-order preauthentication-method;
```

```
    provisioning-order  (gx-plus | jsrc | pcrf);
    radius {
        accounting-server [ ip-address ];
        attributes {
            exclude {
                attribute-name packet-type;
                standard-attribute number {
                    packet-type [ access-request | accounting-off | accounting-on | accounting-
start | accounting-stop ];
                }
                vendor-id id-number {
                    vendor-attribute vsa-number {
                        packet-type [ access-request | accounting-off | accounting-on |
accounting-start | accounting-stop ];
                    }
                }
            }
            ignore {
                dynamic-iflset-name;
                framed-ip-netmask;
                idle-timeout;
                input-filter;
                logical-system:routing-instance;
                output-filter;
                session-timeout;
                standard-attribute number;
                vendor-id id-number {
                    vendor-attribute vsa-number;
                }
            }
        }
        authentication-server [ ip-address ];
        options {
            accounting-session-id-format (decimal | description);
            calling-station-id-delimiter delimiter-character;
            calling-station-id-format {
                agent-circuit-id;
                agent-remote-id;
                interface-description;
                interface-text-description;
                mac-address;
                nas-identifier;
                stacked-vlan;
```

```
            vlan;
        }
        chap-challenge-in-request-authenticator;
        client-accounting-algorithm (direct | round-robin);
        client-authentication-algorithm  (direct | round-robin);
        coa-dynamic-variable-validation;
        ethernet-port-type-virtual;
        interface-description-format {
            exclude-adapter;
            exclude-channel;
            exclude-sub-interface;
        }
        juniper-access-line-attributes;
        nas-identifier identifier-value;
        nas-port-extended-format {
            adapter-width width;
            ae-width width;
            port-width width;
            pw-width width;
            slot-width width;
            stacked-vlan-width width;
            vlan-width width;
            atm {
                adapter-width width;
                port-width width;
                slot-width width;
                vci-width width;
                vpi-width width;
            }
        }
        nas-port-id-delimiter delimiter-character;
        nas-port-id-format {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            interface-text-description;
            nas-identifier;
            order {
                agent-circuit-id;
                agent-remote-id;
                interface-description;
                interface-text-description;
                nas-identifier;
```

```
                    postpend-vlan-tags;
                }
                postpend-vlan-tags;
            }
            nas-port-type {
                ethernet {
                    port-type;
                }
            }
            override {
                calling-station-id remote-circuit-id;
                nas-ip-address tunnel-client-gateway-address;
                nas-port tunnel-client-nas-port;
                nas-port-type tunnel-client-nas-port-type;
            }
            remote-circuit-id-delimiter;
            remote-circuit-id-fallback {
            remote-circuit-id-format;
                agent-circuit-id;
                agent-remote-id;
            }
            revert-interval interval;
            service-activation {
                dynamic-profile (optional-at-login | required-at-login);
                extensible-service (optional-at-login | required-at-login);
            }
            vlan-nas-port-stacked-format;
        }
        preauthentication-server ip-address;
    }
    radius-server server-address {
        accounting-port port-number;
        accounting-retry number;
        accounting-timeout seconds;
        dynamic-request-port
        port port-number;
        preauthentication-port port-number;
        preauthentication-secret password;
        retry attempts;
        routing-instance routing-instance-name;
        secret password;
        max-outstanding-requests value;
         source-address source-address;
```

```
        timeout seconds;
    }
    service {
        accounting {
            statistics (time | volume-time);
            update-interval minutes;
        }
        accounting-order (activation-protocol | local | radius);
    }
    session-limit-per-username number;
    session-options {
        client-idle-timeout minutes;
        client-idle-timeout-ingress-only;
        client-session-timeoutminutes;
        pcc-context {
            input-service-filter-name filter-name;
            input-service-set-name service-set-name;
            ipv6-input-service-filter-name filter-name;
            ipv6-input-service-set-name service-set-name;
            ipv6-output-service-filter-name filter-name;
            ipv6-output-service-set-name service-set-name;
            output-service-filter-name filter-name;
            output-service-set-name service-set-name;
            profile-name pcef-profile-name;
        }
        strip-user-name {
            delimiter [ delimiter ];
            parse-direction (left-to-right | right-to-left);
        }
    }
    subscriber username {
        delegated-pool delegated-pool-name;
        framed-ip-address ipv4-address;
        framed-ipv6-pool ipv6-pool-name;
        framed-pool ipv4-pool-name;
        password password;
        target-logical-system logical-system-name <target-routing-instance (default | routing-
instance-name>;
        target-routing-instance (default | routing-instance-name);
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure a subscriber access profile that includes subscriber access, L2TP, or PPP properties.

Options

profile-name—Name of the profile.

For CHAP, the name serves as the mapping between peer identifiers and CHAP secret keys. This entity is queried for the secret key whenever a CHAP challenge or response is received.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Point-to-Point Protocol (PPP)*
- *Layer 2 Tunneling Protocol (L2TP)*
- *L2TP LNS Inline Service Interfaces*
- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Configuring the PPP Password Authentication Protocol*
- *JSRC for Subscriber Provisioning and Accounting*
- *Configuring Service Accounting in Local Flat Files*
- *AAA Service Framework Overview*

Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management

profile (Application Firewall)

IN THIS SECTION

- Syntax | 13271
- Hierarchy Level | 13271
- Description | 13272
- Options | 13272
- Required Privilege Level | 13272
- Release Information | 13272

Syntax

```
profile profile-name {
    block-message {
        type {
            custom-redirect-url {
                content content;
            }
            custom-text {
                content content;
            }
        }
    }
}
```

Hierarchy Level

```
[edit security application-firewall]
```

Description

Define the profile of the response to be issued when an application firewall rule set blocks HTTP or HTTPS traffic with a `deny` or `reject` action.

Although drop and reject actions are logged, application firewall does not notify users when either action is taken. To provide an explanation for the action or to redirect the users to an informative webpage, you can use the `block-message` option with the reject or deny action in an application firewall rule.

You can customize the redirect action by including additional text on the splash screen or by specifying a URL to which the user is redirected. To customize the block message, define the `type` and `content` in a block message profile defined in the rule set.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Options

name Profile name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Application Firewall Overview

profile (Captive Portal Content Delivery)

IN THIS SECTION

- Syntax | 13273
- Hierarchy Level | 13273
- Description | 13273
- Options | 13274
- Required Privilege Level | 13274
- Release Information | 13274

Syntax

```
profile name
    cpcd-rule-sets rule-set-name;
    cpcd-rules rule-name;
    dynamic;
    http-redirect-options url;
    ipda-rewrite-options {
        destination-address destination-address;
        destination-port destination-port;
    }
}
```

Hierarchy Level

```
[edit services captive-portal-content-delivery]
```

Description

Configure the service profile for HTTP redirect services, also known as Captive Portal and Content Delivery (CPCD) services. The profile contains rules or rule sets that specify the details of the service. The CPCD service profile is included in a service set that applies the service to a service interface. This statement is supported only for static CPCD; it is not supported for converged services CPCD.

Options

cpcd-rule-sets *rule-set-name*	Specify a list of sets of service rules.
cpcd-rules *rule-name*	Specify a service rule.
dynamic	Indicate that the service is a dynamic, converged service.
http-redirect-options *url*	Redirect the packets to the specified URL. The URL must begin with `http://` or `https://`.
ipda-rewrite-options *destination*	Rewrite the destination for the packets to send them to a remote HTTP redirect server. • `destination-address` *address*—IP address of the remote HTTP redirect server. • `destination-port` *number*—(Optional) Port number for the remote HTTP redirect server. Requires the destination address to be specified as well.

Required Privilege Level

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

HTTP Redirect Service Overview

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

Configuring Routing Engine-Based, Static HTTP Redirect Services

Configuring Routing Engine-Based, Converged HTTP Redirect Services

profile(dynamic-application)

IN THIS SECTION

- Syntax | 13275
- Hierarchy Level | 13275
- Description | 13276
- Options | 13276
- Required Privilege Level | 13277
- Release Information | 13277

Syntax

```
profile name {
    redirect-message {
        type {
            custom-text {
                content content;
            }
            redirect-url {
                content content;
            }
        }
    }
}
```

Hierarchy Level

```
[edit security dynamic-application]
```

Description

Define a profile to provide an explanation for the policy action or to redirect the client request to an informative webpage when a policy blocks HTTP or HTTPS traffic with a deny or reject action in a security policy.

Although drop and reject actions can be logged, the users might not be notified when either action is taken. To provide a reason for the action or to redirect the users to an informative webpage, use the `redirect-message` option at the `[edit security dynamic-application profile` *name*`]` hierarchy level to display a custom message.

Customize the redirect action by adding your text message on the splash screen or specify the URL to which the user is redirected. When the `redirect message` option is specified, a splash screen and message informs the client that the traffic has been blocked.

Starting in Junos OS Release 18.2R1, a redirect profile can be configured in a unified policy.

Options

name Name of the profile.

- Values:

  redirect-message Defines the profile of the notification sent to clients when HTTP or HTTPS traffic is blocked by a reject or deny action from an application firewall.

  type The message type displayed after a reject or deny action.

  custom-text Text message in HTML added to the default text. If custom-text is specified, the splash screen displays both the default block message and the custom-defined block message.

  The user is redirected when a reject or deny action is taken during one of the following HTTP methods: GET, POST, OPTIONS, HEAD, PUT, DELETE, TRACE, CONNECT, PROPFIND, PROPPATCH, LOCK, UNLOCK, COPY, MOVE, MKCOL, BCOPY, BDELETE, BCOPY, BMOVE, BPROPFIND, BPROPPATCH, POLL, SEARCH, SUBSCRIBE, and UNSUBSCRIBE. If the reject or deny action occurs during a different HTTP method, the traffic is silently dropped.

  content Custom text added to the splash screen. Custom text is inserted below the default message. Add the characters **\n** to insert a line break in the displayed text.

  redirect-url Custom URL redirection.

content The URL of the webpage to which the client is directed. When traffic is rejected or denied, the client is redirected to the specified webpage for further action. The URL can be hosted on either the SRX Series device or an external server.

Enter the redirect URL in quotation marks for an HTTP or HTTPS site, as shown in the following example:

```
"https://custom-redirect-url"
```

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Unified Security Policies

dynamic-application (Security) | 6093

profile (Explicit Proxy)

IN THIS SECTION

- Syntax | 13278
- Hierarchy Level | 13281
- Description | 13281
- Options | 13281
- Required Privilege Level | 13281
- Release Information | 13281

Syntax

```
profile name {
 default-policy(deny-all | permit-all);
 policy name {
  description description;
   match {
   destination-address;
   destination-address-excluded;
   destination-identity-feed [ destination-identity-feed ... ];
   dynamic-application;
   application;
   source-address;
   source-address-excluded;
   source-identity;
   source-identity-feed [ source-identity-feed ... ];

   source-end-user-profile< source-end-user-profile-name>;

   url-category;
  }
  report-skip;
  scheduler-name scheduler-name;
  then {
   deny {
    application-services {
     security-intelligence {
      add-destination-identity-to-feed name-of-feed;
      add-destination-ip-to-feed name-of-feed;
      add-source-identity-to-feed name-of-feed;
      add-source-ip-to-feed name-of-feed;
     }
    }
   }
   log {
    profile profile;
    session-close;
    session-init;
    session-update minutes;
   }
   permit {
```

```
advanced-connection-tracking;
application-services {
 advanced-anti-malware-policy advanced-anti-malware-policy;
 anti-virus-policy anti-virus-policy;
 application-traffic-control {
  rule-set rule-set;
 }
 icap-redirect icap-redirect;
 idp-policy idp-policy;
 packet-capture;
 security-intelligence {
  add-destination-identity-to-feed name-of-feed;
  add-destination-ip-to-feed name-of-feed;
  add-source-identity-to-feed name-of-feed;
  add-source-ip-to-feed name-of-feed;
 }
 security-intelligence-policy security-intelligence-policy;
 ssl-proxy {
  profile-name profile-name;
 }
 utm-policy utm-policy;
destination-address {
 (drop-translated | drop-untranslated);
}
firewall-authentication {
 pass-through {
  access-profile access-profile;
  auth-only-browser;
  auth-user-agent name;

  client-match [ client-match ... ];
  ssl-termination-profile ssl-termination-profile;
  web-authentication-server web-authentication-server;
  web-redirect;
  web-redirect-to-https;
 }
 user-firewall {
  access-profile access-profile;
  auth-only-browser;
  auth-user-agent name;

  domain domain;
  ssl-termination-profile ssl-termination-profile;
```

```
          web-authentication-server web-authentication-server;
          web-redirect;
          web-redirect-to-https;
         }
         web-authentication {
          client-match [ client-match ... ];
         }
         push-to-identity-management;
        }
        no-services-offload;
        tcp-options {
         initial-tcp-mss initial-tcp-mss;
         reverse-tcp-mss reverse-tcp-mss;
         sequence-check-required;
         syn-check-required;
         window-scale;
        }
        tunnel {
         ipsec-vpn ipsec-vpn;
         pair-policy pair-policy;
        }
       tunnel-inspection profile-name;
       }
       reject {
        application-services {
         security-intelligence {
          add-destination-identity-to-feed name-of-feed;
          add-destination-ip-to-feed name-of-feed;
          add-source-identity-to-feed name-of-feed;
          add-source-ip-to-feed name-of-feed;
         }
        }
        profile profile;
        ssl-proxy {
         profile-name profile-name;
        }
       }
       count {

       }
```

```
    }
   }
```

Hierarchy Level

```
[edit logical-systems name tenants name security policies explicit-proxy],
[edit security policies explicit-proxy],
[edit tenants name security policies explicit-proxy]
```

Description

Define a explicit proxy profile policy context.

The SRX Series Firewall applies security enforcement based on the rules created in the explicit web proxy profile policy.

Options

name Explicit proxy profile name

default-policy Default policy action when no policy matches in rulebase.

- deny-all—Deny all traffic if no policy match
- permit-all—Permit all traffic if no policy match

policy Define security policy for explicit proxy profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

profile (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 13282
- Hierarchy Level | 13282
- Description | 13282
- Options | 13282
- Required Privilege Level | 13283
- Release Information | 13283

Syntax

```
profile profile-name {
    tag rule rule-name;
}
```

Hierarchy Level

```
[edit services hcm]
```

Description

Configure an HCM profile, which points to one or more tag rules that Junos OS uses to enrich HTTP headers with the appropriate tags. You can configure a maximum of 100 HCM profiles.

For subscriber traffic under static policy control, a tag rule is used if it is included in the HCM profile specified in a PCC rule that the traffic matches. For subscribers under dynamic policy control, a message from the PCRF identifies the configured HCM profile and tag rules to use for HTTP header enrichment.

Options

profile-name	Name of the HCM profile.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

Configuring HCM Profiles and Assigning Tag Rules

profile (icap-redirect)

IN THIS SECTION

- Syntax | 13283
- Hierarchy Level | 13284
- Description | 13284
- Options | 13285
- Required Privilege Level | 13285
- Release Information | 13285

Syntax

```
profile name {
    fallback-option {
        connectivity (block | log-permit | permit);
        default-action (block | log-permit | permit);
```

```
        timeout (block | log-permit | permit);
    }
    http {
        redirect-request redirect-request;
        redirect-response redirect-response;
    }
    server name {
        authorization {
            authorization-type authorization-type;
            credentials (ascii ascii | base64 base64);
        }
        host host;
        port port;
        reqmod-uri reqmod-uri;
        respmod-uri respmod-uri;
        routing-instance ri-name;
        sockets sockets;
        tls-profile tls-profile;
    }
    timeout timeout;
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the ICAP redirect profile.

The ICAP server profile allows the ICAP server to process request messages, response messages, fallback options, and so on, for the permitted traffic.

When you configure an ICAP redirect service on SRX Series Firewalls, you must configure the ICAP redirect profile. The ICAP redirect profile defines the settings for ICAP server to process request messages, response messages, fallback options incase of a timeout, connectivity issues, too many requests, or other conditions.

This profile is applied to a security policy as an application service when the traffic is permitted by the security policy.

Options

profile *name* ICAP redirect profile name.

fallback-option Fallback options to specify the actions the device applies if the ICAP server is unavailable.

http Redirect request and redirect response for HTTP traffic.

- Values:
  - redirect-request—Enable the redirect service on HTTP request
  - redirect-response—Enable the redirect service on HTTP response

timeout Server response timeout in milliseconds. Timeout is the interval after which the server is considered inactive if there is no response from the server. A new incoming requests can bypass inactive status server.

- **Default:** 500
- **Range:** 100 through 50000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1 R1.

RELATED DOCUMENTATION

Example: Configuring ICAP Redirect Service on SRX Series Firewalls

profile (Juniper Secure Connect)

IN THIS SECTION

- Syntax | 13286
- Hierarchy Level | 13286
- Description | 13286
- Options | 13287
- Required Privilege Level | 13288
- Release Information | 13288

Syntax

```
profile realm-name {
    access-profile access-profile;
    client-config client-config;
    compliance {
        pre-logon compliance-rule;
    }
    description description;
    ipsec-vpn ipsec-vpn;
    options {
        multi-access;
        user-domain user-domain;
    }
}
```

Hierarchy Level

```
[edit security remote-access]
```

Description

Configure remote user connection profiles for the Juniper Secure Connect clients.

The remote access profiles allow you to deploy connection settings for the remote users by pushing the configuration file on the client devices. You can create multiple profiles and set one of the profiles as the default profile.

> **NOTE**: Starting in Junos OS Release 23.1R1, we've hidden the `default-profile` option at the [`edit security remote-access`] hierarchy level. In releases before Junos OS Release 23.1R1, you use this option to specify one of the remote-access profiles as the default profile in Juniper Secure Connect. But with changes to the format of remote-access profile names, we no longer require the `default-profile` option.
>
> We've deprecated `default-profile` option—rather than immediately removing it—to provide backward compatibility and a chance to make your existing configuration conform to the changed configuration. You'll receive a warning message if you continue to use the `default-profile` option in your configuration. However existing deployments are not affected if you modify the current configuration. See default-profile (Juniper Secure Connect).

Each remote access profile includes a *realm-name* mapping to a URL either in *FQDN/RealmName* or *FQDN* format, authentication settings, VPN settings, and client configurations. You can create different remote access profiles for different names or functions.

Example—You can create a configuration profile for the engineering department, and another for the human resource department. You name the profile for engineering department and human resource department as *ra.example.com/engineering* and *ra.example.com/hr* respectively.

When a Juniper Secure connect remote user selects a connection profile such as *ra.example.com/engineering*, the SRX Series Firewall receives the configuration request and selects a remote-access profile with same name —that is—*ra.example.com/engineering* for pushing the configuration on client device.

Options

realm-name

Set realm-name as remote-access profile name. This is the profile identifier in FQDN/RealmName format.

Examples:

ra.example.com/hr, if FQDN is *ra.example.com* and Realm name is *hr*.

ra.example.com/engineering, if FQDN is *ra.example.com* and Realm name is *engineering*.

ra.example.com, if FQDN is *ra.example.com* and Realm name is empty.

Specify an IP address if you do not have an *FQDN* (*192.168.1.10/hr* or *192.168.1.10)*.

access-profile Select the access profile for authentication and accounting for clients.

client-config Select the client configuration object.

compliance Select pre-logon compliance rule object name.

description Text description of the remote access profile.

ipsec-vpn Select the IPsec VPN policy object used for IKE and IPsec proposals.

multi-access Set to allow an user to connect from multiple devices.

user-domain (Optional) Specify a VPN user domain alias name.

When you associate an IPsec VPN object with two different remote access profiles, you must use the same domain alias name in those remote access profiles. You must not configure different domain alias names for the different remote access profiles that use the same IPsec VPN object.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.3R1.

`compliance` option added in Junos OS Release 23.1R1.

Realms format changed to FQDN/RealmName or FQDN in Junos OS Release 23.1R1.

`user-domain` option added in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

Juniper Secure Connect Administrator Guide

profile (LRF)

IN THIS SECTION

- Syntax | 13289
- Hierarchy Level | 13290
- Description | 13290
- Options | 13290
- Required Privilege Level | 13290
- Release Information | 13290

Syntax

```
profile profile-name {
    collector collector-name {
        destination {
            address collector-address;
            port collector-port-number;
        }
        source-address source-address;
    }
    http-log-multiple-transactions;
    policy-based-logging;
    rule lrf-rule-name {
        then {
            report {
                collector collector-name;
                template template-name;
                time-limit time-interval;
                volume-limit volume;
            }
        }
    }
    template template-name {
        format ipfix;
        template-tx-interval tx-time;
        template-type template-type;
```

```
        trigger-type (session-close | time | volume);
    }
    vendor-support ibm;
}
```

Hierarchy Level

```
[edit services lrf]
```

Description

Configure an LRF profile to specify a set of logging and reporting parameters, which includes data templates, collectors, and LRF rules.

For Junos OS Subscriber Aware, you can then assign an LRF profile to a subscriber by assigning the profile to the TDF service set associated with the subscriber's TDF domain.

For Junos OS Broadband Subscriber Management, you can then assign the LRF profile to the service set that is configured for application-aware policy control.

Options

profile-name

Name of the LRF profile.

- **Range:** Up to 63 characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

vendor-support option introduced in Junos OS Release 17.2.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

Logging and Reporting Function for Subscribers

profile (Rule Sets)

IN THIS SECTION

- Syntax | 13291
- Hierarchy Level | 13291
- Description | 13291
- Options | 13292
- Required Privilege Level | 13292
- Release Information | 13292

Syntax

```
profile profile-name;
```

Hierarchy Level

```
[edit security application-firewall rule-sets rule-set-name]
```

Description

Specifies the profile of the block message to be used for any deny or reject action in the rule set that specifies the `block-message` option.

The block-message option enables you to provide an explanation for the action or to redirect the client to an informative webpage. You can configure the block-message in `set security application-firewall profile` hierarchy.

Options

profile-name—Name of the block-message profile to be used for this rule set.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Application Firewall Overview

profile (security)

IN THIS SECTION

- Syntax | 13293
- Hierarchy Level | 13294
- Description | 13294
- Default | 13295
- Options | 13295
- Required Privilege Level | 13310
- Release Information | 13310

Syntax

```
profile name {
category {
aamw {
field-extra-name name;
field-name name;
}
antispam {
field-extra-name name;
field-name name;
}
antivirus {
field-extra-name name;
field-name name;
}
apptrack {
field-extra-name name;
field-name name;
}
content-filter {
field-extra-name name;
field-name name;
}
idp {
field-extra-name name;
field-name name;
}
secintel {
field-extra-name name;
field-name name;
}
session {
field-extra-name name;
field-name name;
}
ssl-proxy {
field-extra-name name;
field-name name;
}
webfilter {
field-extra-name name;
```

```
field-name name;
}
}
stream-name name;
template (sd-wan | traditional-firewall | unified-ngfw);
}
```

Hierarchy Level

```
[edit security log]
```

Description

Configure the security log profile to define the policy related log format including fields selection and fields order for the log. You can apply the log profile to a policy to get the log in a predefined format. Configure the stream and the log server to send the logs to the destination before configuring the profile. You can specify the log profile only for the policy related logs. After configuring the policy category, configure the field-name from the options that are available for the policy log per category. After you have selected the field name, select the field-extra name from the available options. You can apply the profile on the policy related logs mentioned in the Table 101 on page 13294.

Table 101: You can apply log profile on below policy related logs and their categories.

Category	Policy related logs
SESSION	RT_FLOW_SESSION_CREATE RT_FLOW_SESSION_CLOSE RT_FLOW_SESSION_DENY RT_FLOW_SESSION_UPDATE
WEBFILTER	WEBFILTER_URL_BLOCKED WEBFILTER_URL_PERMITTED WEBFILTER_URL_REDIRECTED
ANTIVIRUS	AV_VIRUS_DETECTED_MT

Table 101: You can apply log profile on below policy related logs and their categories. ***(Continued)***

Category	Policy related logs
CONTENT-FILTER	CONTENT_FILTERING_BLOCKED_MT
ANTISPAM	ANTISPAM_SPAM_DETECTED_MT
IDP	IDP_ATTACK_LOG_EVENT
SECINTEL	SECINTEL_ACTION_LOG
AAMW	AAMW_ACTION_LOG AAMW_IMAP_ACTION AAMW_SMTP_ACTION
SSL-PROXY	SSL_PROXY_SSL_SESSION_DROP SSL_PROXY_SSL_SESSION_ALLOW
APPTRACK	APPTRACK_SESSION_VOL_UPDATE APPTRACK_SESSION_CREATE APPTRACK_SESSION_CLOSE

Default

Options

category Selects the category of events for the profile

stream-name Specify the destination to send the log. One log profile can point to multiple streams. Different log profiles can point to the same stream. You need to configure the stream mode for the logging before configuring stream-name.

template Select the template from the available options for the profile to get the log in a predefined format. You can apply the profile template on the policy related logs and categories mentioned in the Table 102 on page 13296.

Table 102: You can apply log profile templates on the below listed policy related logs and categories.

TEMPLATES	CATEGORIES	FIELDS
traditional-firewall	SESSION	reason source-address source-port destination-address destination-port service-name nat-source-address nat-source-port nat-destination-address nat-destination-port protocol-id policy-name source-zone-name destination-zone-name session-id-32 packets-from-client bytes-from-client packets-from-server bytes-from-server elapsed-time username packet-incoming-interface flags
	WEBFILTER	all-fields

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
	ANTIVIRUS	all-fields
	CONTENT-FILTER	all-fields
	ANTISPAM	all-fields
	IDP	all-fields
	SECINTEL	all-fields
	AAMW	all-fields
	SSL-PROXY	all-fields
	APPTRACK	all-fields
unified-ngfw	SESSION	session-id
		reason
		source-address
		source-port
		destination-address
		destination-port
		nat-source-address
		nat-source-port
		nat-destination-address
		nat-destination-port

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		protocol-id
		policy-name
		source-zone-name
		destination-zone-name
		packets-from-client
		bytes-from-client
		packets-from-server
		bytes-from-server
		elapsed-time
		application
		nested-application
		username
		packet-incoming-interface
		application-category
		application-sub-category
		application-risk
		flags
	WEBFILTER	session-id

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		category
		reason
		profile
		url
		username
		urlcategory-risk
	ANTIVIRUS	Session-id
		profile-name
		filename
		temporary-filename
		name
		url
		username
		action
	CONTENT-FILTER	Session-id
		argument
		profile-name
		action

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		reason
		username
		filename
	ANTISPAM	Session-id
		profile-name
		action
		reason
		username
	IDP	Session-id
		Message-type
		rule-name
		rulebase-name
		policy-name
		export-id
		repeat-count
		action
		threat-severity
		attack-name

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		packet-log-id
		alert
		username
		xff-header
		message
	SECINTEL	Session-id
		Category
		sub-category
		action
		action-detail
		http-host
		threat-severity
		feed-name
		policy-name
		profile-name
		username
	AAMW AAMW_ACTION_LOG	session-id
		Hostname

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		file-category
		verdict-number
		malware-info
		action
		list-hit
		file-hash-lookup
		policy-name
		username
		url
	SSL-PROXY	session-id
		profile-name
		message
		username
	APPTRACK APPTRACK_SESSION_CLOSE	session-id
		application
		nested-application
		username
		profile-name

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		rule-name
		routing-instance
		destination-interface-name
		uplink-incoming-interfacename
		uplink-tx-bytes
		uplink-rx-bytes
		apbr-policy-name
		amr-rule-name
sd-wan	SESSION	session-id
		reason
		source-address
		source-port
		destination-address
		destination-port
		nat-source-address
		nat-source-port
		nat-destination-address
		nat-destination-port

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		protocol-id
		policy-name
		source-zone-name
		destination-zone-name
		packets-from-client
		bytes-from-client
		packets-from-server
		bytes-from-server
		elapsed-time
		application
		nested-application
		username
		packet-incoming-interface
		application-category
		application-sub-category
		application-risk
		flags
	WEBFILTER	session-id

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		category
		reason
		profile
		url
		username
		urlcategory-risk
	ANTIVIRUS	Session-id [this does not exist today – needs to be added]
		profile-name
		filename
		temporary-filename
		name
		url
		username
		action [this does not exist today either]
	CONTENT-FILTER	Session-id [this does not exist today]
		argument

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		profile-name
		action
		reason
		username
		filename
	ANTISPAM	Session-id [this does not exist today]
		profile-name
		action
		reason
		username
	IDP	Session-id (does not exist today)
		Message-type
		rule-name
		rulebase-name
		policy-name
		export-id
		repeat-count

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		action
		threat-severity
		attack-name
		packet-log-id
		alert
		username
		xff-header
		message
	SECINTEL	Session-id
		Category
		sub-category
		action
		action-detail
		http-host
		threat-severity
		feed-name
		policy-name
		profile-name

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		username
	AAMW AAMW_ACTION_LOG	session-id
		Hostname
		file-category
		verdict-number
		malware-info
		action
		list-hit
		policy-name
		username
		url
	SSL-PROXY	session-id
		profile-name
		message
		username [this field does not exist today]
	APPTRACK APPTRACK_SESSION_CLOSE	source-address
		source-port
		destination-address

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		destination-port
		service-name
		application
		nested-application
		policy-name
		source-zone-name
		destination-zone-name
		packets-from-client
		bytes-from-client
		packets-from-server
		bytes-from-server
		elapsed-time
		profile-name
		rule-name
		routing-instance
		destination-interface-name
		uplink-incoming-interfacename
		uplink-tx-bytes

Table 102: You can apply log profile templates on the below listed policy related logs and categories. ***(Continued)***

TEMPLATES	CATEGORIES	FIELDS
		uplink-rx-bytes
		multipath-rule-name
		src-vrf-grp
		dst-vrf-grp
		username
		session-id
		reason

Required Privilege Level

Release Information

Configuration statement `profile` is added in Junos OS Release 21.1R1.

profile (Security Antispam SBL)

IN THIS SECTION

- Syntax | 13311
- Hierarchy Level | 13311
- Description | 13311
- Options | 13311
- Required Privilege Level | 13311

Release Information | 13311

Syntax

```
profile profile-name {
    custom-tag-string [string];
    (sbl-default-server | no-sbl-default-server);
    spam-action (block | tag-header | tag-subject);
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-spam sbl]
```

Description

Create a profile for the antispam sbl feature. This profile includes all subsequent configuration options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

profile (Security Antivirus Juniper Express Engine)

IN THIS SECTION

- Syntax | 13312
- Hierarchy Level | 13313
- Description | 13313
- Options | 13313
- Required Privilege Level | 13313
- Release Information | 13314

Syntax

```
profile profile-name {
     fallback-options {
        content-size (block | log-and-permit);
        default (block | log-and-permit);
        engine-not-ready (block | log-and-permit);
        out-of-resources (block | (log-and-permit);
        timeout (block | log-and-permit);
        too-many-requests (block | log-and-permit);
    }
    notification-options {
        fallback-block {
            administrator-email email-address;
            allow-email;
            custom-message message;
            custom-message-subject message-subject;
            display-host;
            (notify-mail-sender | no-notify-mail-sender);
             type (message | protocol-only);
        }
        fallback-non-block {
            custom-message message;
            custom-message-subject message-subject;
            (notify-mail-recipient | no-notify-mail-recipient);
        }
```

```
        virus-detection {
            custom-message message;
            custom-message-subject message-subject;
            (notify-mail-sender | no-notify-mail-sender);
             type (message | protocol-only);
        }
    }
    scan-options {
        content-size-limit value;
        (intelligent-prescreening | no-intelligent-prescreening);
        timeout value;
    }
    trickling {
        timeout value;
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine]
```

Description

Create a profile for the Juniper express engine. This profile includes all subsequent configuration options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The express engine feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Express Antivirus Configuration Overview

profile (Security Antivirus Kaspersky Lab Engine)

IN THIS SECTION

- Syntax | 13314
- Hierarchy Level | 13316
- Description | 13316
- Options | 13316
- Required Privilege Level | 13316
- Release Information | 13316

Syntax

```
profile profile-name {
     fallback-options {
        content-size (block | log-and-permit);
        corrupt-file (block | log-and-permit);
        decompress-layer (block | log-and-permit);
        default (block | log-and-permit);
        engine-not-ready (block | log-and-permit);
        out-of-resources (block | (log-and-permit);
        password-file (block | (log-and-permit);
        timeout (block | log-and-permit);
```

```
        too-many-requests (block | log-and-permit);
    }
    notification-options {
        fallback-block {
            administrator-email email-address;
            allow-email;
            custom-message message;
            custom-message-subject message-subject;
            display-host;
            (notify-mail-sender | no-notify-mail-sender);
             type (message | protocol-only);
        }
        fallback-non-block {
            custom-message message;
            custom-message-subject message-subject;
            (notify-mail-recipient | no-notify-mail-recipient);
        }
        virus-detection {
            custom-message message;
            custom-message-subject message-subject;
            (notify-mail-sender | no-notify-mail-sender);
             type (message | protocol-only);
        }
    }
    scan-options {
        content-size-limit value;
        decompress-layer-limit value;
        (intelligent-prescreening | no-intelligent-prescreening);
        scan-extension filename;
        scan-mode  (all | by-extension);
        timeout value;
    }
    trickling {
        timeout value;
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
```

Description

Create a profile for the Kaspersky Lab engine. This profile includes all subsequent configuration options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Kaspersky feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

kaspersky-lab-engine | **9450**

profile (Security Antivirus Juniper Express Engine) | **13312**

profile (Security Content Filtering)

IN THIS SECTION

- Syntax | 13317
- Hierarchy Level | 13317
- Description | 13318
- Options | 13318
- Required Privilege Level | 13318
- Release Information | 13318

Syntax

```
profile profile-name {
    block-command protocol-command-list;
    block-content-type (activex | exe | http-cookie | java-applet | zip);
    block-extension extension-list;
    block-mime {
        exception list-name;
        list list-name;
    }
    notification-options {
        custom-message message;
        (notify-mail-sender | no-notify-mail-sender);
         type (message | protocol-only);
    }
    permit-command protocol-command-list;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering]
```

Description

Create a profile for the content-filtering feature. This profile includes all subsequent configuration options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

profile (Security GTP)

IN THIS SECTION

- Syntax | 13319
- Hierarchy Level | 13321
- Description | 13322
- Options | 13322
- Required Privilege Level | 13323
- Release Information | 13323

Syntax

```
profile profile_name {
    apn pattern-string {
        imsi-prefix imsi-prefix-digits {
            action {
                drop;
                pass;
                selection {
                    ms;
                    net;
                    vrf;
                }
            }
        }
    }
    drop {
        aa-create-pdp 0;
        aa-delete-pdp 0;
        bearer-resource 2;
        change-notification 2;
        config-transfer 2;
        context 2;
        create-bearer 2;
        create-data-forwarding 2;
        create-pdp (0 | 1 | all);
        create-session 2;
        create-tnl-forwarding 2;
        cs-paging 2;
        data-record (0 | 1 | all);
        delete-bearer 2;
        delete-command 2;
        delete-data-forwarding 2;
        delete-pdn 2;
        delete-pdp (0 | 1 | all);
        delete-session 2;
        detach 2;
        downlink-notification 2;
        echo (0 | 1 | 2 | all);
        error-indication (0 | 1 | all);
        failure-report (0 | 1 | all);
        fwd-access 2;
```

```
        fwd-relocation (1 | 2 | all);
        fwd-srns-context 1;
        g-pdu (0 | 1 | all);
        identification (0 | 1 | 2 | all);
        mbms-session-start (1 | 2 | all);
        mbms-session-stop (1 | 2 | all);
        mbms-session-update (1 | 2 | all);
        modify-bearer 2;
        modify-command 2;
        node-alive (0 | 1 | all);
        note-ms-present (0 | 1 | all);
        pdu-notification (0 | 1 | all);
        ran-info (1 | 2 | all);
        redirection (0 | 1 | all);
        release-access 2;
        relocation-cancel (1 | 2 | all);
        resume 2;
        send-route (0 | 1 | all);
        sgsn-context (0 | 1 | all);
        stop-paging 2;
        supported-extension 1;
        suspend 2;
        trace-session 2;
        update-bearer 2;
        update-pdn 2;
        update-pdp (0 | 1 | all);
        ver-not-supported (0 | 1 | 2 | all);
    }
    end-user-address-validated;
    gtp-in-gtp-denied;
    handover-group group-name;
    handover-on-roaming-intf;
    listening-mode;
    log {
        forwarded (basic | detail);
        gtp-u name;
        prohibited (basic | detail);
        rate-limited (basic | detail);
        state-invalid (basic | detail);
    }
    max-message-length max-message-length;
    min-message-length min-message-length;
    must-ie-v1 msgie-prf-v1-name;
```

```
    must-ie-v2 msgie-prf-v2-name;
    ne-group group-name;
    path-rate-limit {
        message-type (create-req | delete-req | echo-req | other) {
            alarm-threshold {
                forward forward;
                reverse reverse;
            }
            drop-threshold {
                forward forward;
                reverse reverse;
            }
        }
    }
    rate-limit {
        alarm-threshold alarm-threshold;
        drop-threshold drop-threshold;
    }
    remove-ie  {
        version v1 {
            release name;
        }
    }
    remove-ie-v1 ieset-name;
    remove-ie-v2 ieset-name;
    req-timeout seconds;
    restart-path (all | create | echo);
    seq-number-validated;
    timeout hours;
    u-tunnel-validated;
    ue-group group-name;
}
```

Hierarchy Level

```
[edit security gtp]
```

Description

Use the profile option to create a profile for the GPRS tunneling protocol (GTP) feature. This profile includes all subsequent configuration options.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp]` hierarchy level is replaced by `[edit security gtp]`.

Options

- `apn-control`–You can specify APN control profile using APN string. APN string consists of Network Identifier (NI) and Operator Identifier (OI). You can specify APN in the below formats:

  - <NI>: for example internet, internet.public

  - <NI>.<OI>: for example internet.mnc005.mcc244.gprs, internet.public.mnc005.mcc244.gprs

  APN string configuration supports wild card "*" in place of OI. You can specify APN as <NI>."*" and configure the matching APN string. You don't need to specify the OI when you configure APN using wild card "*".

  You can also use this configuration to group all the APNs with the same rate limiting requirements and apply the group to the GTP profile.

- `end-user-address-validated`—During the GTP-U security check procedure, IPv4 and IPv6 addresses for user equipment is checked against the end-user address stored in the user tunnel. Once the GTP-U packet is determined to match the user equipment address, the packet data unit (PDU) is parsed to obtain the user equipment address. Use the end-user-address-validated function to validate the IP address. By default, the UE address check is disabled, it will not check the IP address of GTP-U payload.

- `gtp-in-gtp-denied`—To enable the security device to detect and drop a GTP packet that contains another GTP packet in its message body.

- `handover-on-roaming-intf`—To enable the security device to receive context and forward relocation messages, inspect the packets, and to set up PDP contexts on the device.

- `u-tunnel-validated`—To specify GTP-U tunnel validation. When u-tunnel-validated is enabled, the GTP ALG checks whether the tunnel endpoint identifier (TEID) in the GTP-U packet matches the user tunnel in the tunnel table. If the user tunnel is not found, then the GTP-U packet is dropped. By default, the check is disabled.

- `must-ie-v1`—To check the presence of IEs that should be contained in a GTPv1 message. The device checks the presence of Must-IEs of specific GTP messages and forwards the messages only if Must-IEs are present.

- `must-ie-v2`—To check the presence of IEs that should be contained in a GTPv2 message. The device checks the presence of Must-IEs of specific GTP messages and forwards the messages only if Must-IEs are present.
- `remove-ie-v1`—To remove IEs of specific types from all messages for GTPv1. This helps to remove all instances of specified IEs such as supporting IE, Grouped IE, Embedded IE, or embedded grouped IE.
- `remove-ie-v2`—To remove IEs of specific types from all messages for GTPv2. This helps to remove all instances of specified IEs such as supporting IE, Grouped IE, Embedded IE, or embedded grouped IE.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. The `restart-path` option added in Junos OS Release 11.4. New GPRS tunneling protocol (GTP) message types added in Junos OS Release 11.4. Support for GTPv2 added in Junos OS Release 11.4. Statement modified in Junos OS Release 15.1X49-D40.The `must-ie-v1`, `must-ie-v2`, `remove-ie-v1`, and `remove-ie-v2` options added in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Understanding GTP Path Restart

Understanding GTPv2 Tunnel Cleanup

GTPv1 Message Filtering

profile (Security SCTP)

IN THIS SECTION

Syntax | 13324

Hierarchy Level | 13325

- Description | 13325
- Options | 13325
- Required Privilege Level | 13325
- Release Information | 13325

Syntax

```
profile profile-name {
    association-timeout association-timeout;
    drop (Security SCTP) {
        m3ua-service name;
        payload-protocol name;
    }
    handshake-timeout handshake-timeout;
    limit (Security SCTP) {
        address name {
            payload-protocol (asap | bicc | ddp-segment | ddp-stream | diameter-dtls | diameter-
sctp | dua | enrp | h248 | h323 | id | iua | m2pa | m2ua | m3ua | others | qipc | reserved |
s1ap | simco | sua | tali | v5ua | x2ap) {
                rate rate;
            }
        }
        payload-protocol (asap | bicc | ddp-segment | ddp-stream | diameter-dtls | diameter-sctp
| dua | enrp | h248 | h323 | id | iua | m2pa | m2ua | m3ua | others | qipc | reserved | s1ap |
simco | sua | tali | v5ua | x2ap) {
            rate rate;
        }
        rate {
            address name {
                sccp sccp;
                ssp ssp;
                sst sst;
            }
            sccp sccp;
            ssp ssp;
            sst sst;
        }
    }
```

```
    nat-only;
    permit (Security SCTP) {
        payload-protocol name;
    }
}
```

Hierarchy Level

```
[edit security sctp]
```

Description

Use the profile option to create a profile of the Stream Control Transmission Protocol (SCTP) feature. This profile includes all subsequent configuration options.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs sctp]` hierarchy level is replaced by `[edit security sctp]`.

Options

name Specify SCTP configuration name.

nat-only Specify only do payload IPs translation for SCTP packet.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2. Support for the `nat-only` option added in Junos OS Release 12.1X45-D10. Support for the `permit` option is added in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Understanding Stream Control Transmission Protocol

Understanding SCTP Multichunk Inspection

profile (Security Sophos Engine Antivirus)

IN THIS SECTION

- Syntax | 13326
- Hierarchy Level | 13327
- Description | 13327
- Options | 13327
- Required Privilege Level | 13328
- Release Information | 13328

Syntax

```
profile <name> {
    fallback-options {
        content-size (block | log-and-permit | permit);
        default (block | log-and-permit | permit);
        engine-not-ready (block | log-and-permit | permit);
        out-of-resources (block | log-and-permit | permit);
        timeout (block | log-and-permit | permit);
        too-many-requests (block | log-and-permit | permit);
    }
    notification-options {
        fallback-block {
            administrator-email email-address;
            allow-email;
            custom-message message;
            custom-message-subject message-subject;
            display-host;
            (notify-mail-sender | no-notify-mail-sender);
```

```
             type (message | protocol-only);
        }
        fallback-non-block {
            custom-message message;
            custom-message-subject message-subject;
            (notify-mail-recipient | no-notify-mail-recipient);
        }
        virus-detection {
            custom-message message;
            custom-message-subject message-subject;
            (notify-mail-sender | no-notify-mail-sender);
             type (message | protocol-only);
        }
    }
    scan-options {
        content-size-limit value;
        (no-uri-check | uri-check);
        timeout value;
    }
    trickling {
        timeout value;
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus sophos-engine]
```

Description

Create a profile for the Sophos antivirus engine. This profile includes all subsequent configuration options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Sophos Antivirus Configuration Overview

profile (Security Web Filtering Juniper Enhanced)

IN THIS SECTION

- Syntax | 13328
- Hierarchy Level | 13329
- Description | 13329
- Options | 13329
- Required Privilege Level | 13329
- Release Information | 13330

Syntax

```
profile profile-name {
    category customurl-list name {
        action  (block | log-and-permit | permit | quarantine);
    }
    custom-block-message value;
```

```
    custom-quarantine-message value;
    default  (block | log-and-permit | permit | quarantine);
    fallback-settings {
        default  (block | log-and-permit);
        server-connectivity   (block | log-and-permit);
        timeout    (block | log-and-permit);
        too-many-requests  (block | log-and-permit);
    }
    no-safe-search;
    site-reputation-action {
        fairly-safe   (block | log-and-permit | permit | quarantine);
        harmful   (block | log-and-permit | permit | quarantine);
        moderately-safe  (block | log-and-permit | permit | quarantine);
        suspicious  (block | log-and-permit | permit | quarantine);
        very-safe   (block | log-and-permit | permit | quarantine);
    }
    timeout value;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering juniper-enhanced]
```

Description

Create a profile for the juniper-enhanced feature. This profile includes all subsequent configuration options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Monitoring Web Filtering Configurations

profile (Security Web Filtering Juniper Local)

IN THIS SECTION

- Syntax | 13330
- Hierarchy Level | 13331
- Description | 13331
- Options | 13331
- Required Privilege Level | 13331
- Release Information | 13331

Syntax

```
profile profile-name {
    custom-block-message value;
    default  (block | log-and-permit | permit);
    fallback-settings {
        default (block | log-and-permit);
        server-connectivity (block | log-and-permit);
        timeout  (block | log-and-permit);
        too-many-requests  (block | log-and-permit);
    }
    timeout value;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering juniper-local]
```

Description

Create a profile for the web-filtering juniper–local feature. This profile includes all subsequent configuration options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

> *Monitoring Web Filtering Configurations*
>
> *Example: Configuring Local Web Filtering*

profile (Security Web Filtering Surf Control Integrated)

IN THIS SECTION

- Syntax | 13332
- Hierarchy Level | 13332
- Description | 13333
- Options | 13333
- Required Privilege Level | 13333
- Release Information | 13333

Syntax

```
profile profile-name {
    category customurl-list name {
        action  (block | log-and-permit | permit);
    }
    custom-block-message value;
    default (block | log-and-permit | permit);
    fallback-settings {
        default  (block | log-and-permit);
        server-connectivity   (block | log-and-permit);
        timeout    (block | log-and-permit);
        too-many-requests  (block | log-and-permit);
    }
    timeout value;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated]
```

Description

Create a profile for the web-filtering surf-control-integrated feature. This profile includes all subsequent configuration options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Surf-Control feature is not supported from Junos OS Release 15.1x49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Monitoring Web Filtering Configurations

profile (Security Web Filtering Websense Redirect)

IN THIS SECTION

- Syntax | 13334
- Hierarchy Level | 13334
- Description | 13334
- Options | 13334
- Required Privilege Level | 13334
- Release Information | 13335

Syntax

```
profile profile-name {
   account value;
   custom-block-message value;
   fallback-settings {
       default  (block | log-and-permit);
       server-connectivity  (block | log-and-permit);
       timeout   (block | log-and-permit);
       too-many-requests  (block | log-and-permit);
   }
   server {
       host host-name;
       port number;
   }
   sockets value;
   timeout value;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering websense-redirect]
```

Description

Create a profile for the web-filtering web-sense feature. This profile includes all subsequent configuration options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Monitoring Web Filtering Configurations

profile (Services Application Identification)

IN THIS SECTION

- Syntax | 13335
- Hierarchy Level | 13335
- Description | 13336
- Options | 13336
- Required Privilege Level | 13336
- Release Information | 13336

Syntax

```
profile app-id-profile-name;
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Configure an application identification profile. This profile is a placeholder profile with no configuration options, but it must be created to enable application identification functionality on the services plane.

Options

app-id-profile-name	Name of the application identification profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Applying Services to Subscriber-Aware Traffic with a Service Set

profile (Services PCEF)

IN THIS SECTION

- Syntax | 13337
- Hierarchy Level | 13337
- Description | 13337
- Options | 13337
- Required Privilege Level | 13337
- Release Information | 13337

Syntax

```
profile pcef-profile-name;
```

Hierarchy Level

```
[edit services pcef]
```

Description

Configure a policy and charging enforcement function (PCEF) profile that is a placeholder profile with no configuration options. This profile must be created to enable PCEF functionality on the services plane. You apply this placeholder profile to the subscriber-aware service set.

Options

pcef-profile-name	Name of the PCEF profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> Applying Services to Subscriber-Aware Traffic with a Service Set
>
> *Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control*

profile (Services Proxy)

IN THIS SECTION

- Syntax | 13338
- Hierarchy Level | 13338
- Description | 13338
- Options | 13339
- Required Privilege Level | 13339
- Release Information | 13339

Syntax

```
profile name {
    profile (Services Proxy) {
        http {
            host host;
            port port;
        }
    }
}
```

Hierarchy Level

```
[edit services proxy]
```

Description

Define the proxy profile settings for application signature package download.

You can download the application signature package hosted on an external server, using a proxy sever. To use the proxy server for downloading, you must configure a profile with host and port details of the proxy server, and apply the proxy profile in the `set services application-identification download` command.

This configuration enables you to download the signature package when you have already deployed a web proxy on your device as part of your overall security solution.

Options

name Proxy profile name.

protocol Protocol type for the profile. Support is available for only HTTP connections.

host IP address of the proxy server.

port Port number used by the proxy server.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Predefined Application Signatures for Application Identification

profile (Services SSL Proxy)

IN THIS SECTION

- Syntax | 13340
- Hierarchy Level | 13341
- Description | 13341
- Options | 13341
- Required Privilege Level | 13343
- Release Information | 13343

Syntax

```
profile   name {
    actions {
        allow-strong-certificate;
        crl {
            disable;
            if-not-present (allow | drop);
            ignore-hold-instruction-code;
        }
        disable-session-resumption;
        ignore-server-auth-failure;
        log {
            all;
            errors;
            info;
            sessions-allowed;
            sessions-dropped;
            sessions-ignored;
            sessions-whitelisted;
            warning;
        }
        renegotiation {
            (allow | allow-secure | drop);
        }
    }
    custom-ciphers ;
    disable-deferred-profile-selection;
    enable-flow-tracing    enable-flow-tracing;
    mirror-decrypt-traffic {
        interface interface-name;
        only-after-security-policies-enforcement;
        destination-mac-address mac-address;
    }
    preferred-ciphers (custom | medium | strong | weak);
    ( root-ca     root-ca | server-certificate[ server-certificate ... ]);
    trusted-ca  ;
    whitelist [ whitelist ... ];
    whitelist-url-categories  [ whitelist-url-categories ... ];
}
```

Hierarchy Level

```
[edit services ssl proxy],
[edit logical-systems logical-system-name services ssl proxy]
```

Description

Specify the SSL server profile. An SSL proxy profile defines SSL behavior for the SRX Series Firewall.

The SSL proxy profile will be applied to the security policy as application services.

Options

profile-name Profile identifier.

actions Logging and traffic related actions.

custom-ciphers Custom cipher list.

- Values:
  - `ecdhe-rsa-with-3des-ede-cbc-sha`—ECDHE/RSA, 3DES EDE/CBC, SHA hash
  - `ecdhe-rsa-with-aes-128-cbc-sha`—ECDHE/RSA, 128-bit AES/CBC, SHA hash
  - `ecdhe-rsa-with-aes-128-cbc-sha256`—ECDHE/RSA, 128-bit AES/CBC, SHA256 hash
  - `ecdhe-rsa-with-aes-128-gcm-sha256`—ECDHE/RSA, 128-bit AES/GCM, SHA256 hash
  - `ecdhe-rsa-with-aes-256-cbc-sha`—ECDHE/RSA, 256-bit AES/CBC, SHA hash
  - `ecdhe-rsa-with-aes-256-cbc-sha384`—ECDHE/RSA, 256-bit AES/CBC, SHA384 hash
  - `ecdhe-rsa-with-aes-256-gcm-sha384`—ECDHE/RSA, 256-bit AES/gcm, SHA384 hash
  - `rsa-export-with-des40-cbc-sha`—RSA-export, 40-bit DES/CBC, SHA hash
  - `rsa-export-with-rc4-40-md5`—RSA-export, 40-bit RC4, MD5 hash
  - `rsa-export1024-with-des-cbc-sha`—RSA 1024-bit export, DES/CBC, SHA hash

- `rsa-export1024-with-rc4-56-md5`—RSA 1024-bit export, 56 bit RC4, MD5 hash
- `rsa-export1024-with-rc4-56-sha`—RSA 1024-bit export, 56 bit RC4, SHA hash
- `rsa-with-3des-ede-cbc-sha`—RSA, 3DES EDE/CBC, SHA hash
- `rsa-with-aes-128-cbc-sha`—RSA, 128-bit AES/CBC, SHA hash
- `rsa-with-aes-128-cbc-sha256`—RSA, 128-bit AES/CBC, SHA256 hash
- `rsa-with-aes-128-gcm-sha256`—RSA, 128-bit AES/gcm, SHA256 hash
- `rsa-with-aes-256-cbc-sha`—RSA, 256-bit AES/CBC, SHA hash
- `rsa-with-aes-256-cbc-sha256`—RSA, 256-bit AES/CBC, SHA256 hash
- `rsa-with-aes-256-gcm-sha384`—RSA, 256-bit AES/gcm, SHA384 hash
- `rsa-with-des-cbc-sha`—RSA, DES CBC, SHA hash
- `rsa-with-null-md5`—RSA, no symmetric cipher, MD5 hash
- `rsa-with-null-sha`—RSA, no symmetric cipher, SHA hash
- `rsa-with-rc4-128-md5`—RSA, 128-bit RC4, MD5 hash
- `rsa-with-rc4-128-sha`—RSA, 128-bit RC4, SHA hash

disable-deferred-profile-selection

Disable the deferred profile selection mechanism. In the defered profile selection mechanism, the SSL proxy module defers SSL profile selection until the dynamic application is detected in a client hello message based on the Server Name Indication (SNI). After detecting dynamic application, SSL proxy module does a firewall rule lookup based on the identified application and selects an appropriate SSL proxy profile.

enable-flow-tracing

Enable flow tracing for the profile.

preferred-ciphers

Select preferred ciphers.

- Values:
  - `custom`—Configure custom cipher suite and order of preference.
  - `medium`—Use ciphers with key strength of 128-bits or greater.
  - `strong`—Use ciphers with key strength of 168-bits or greater.
  - `weak`—Use ciphers with key strength of 40-bits or greater.

root-ca	Root certificate for interdicting server certificates in proxy mode.
server-certificate	Local certificate identifier.
trusted-ca	List of trusted certificate authority profiles.
whitelist	Addresses exempted from SSL proxy.
whitelist-url-categories	URL categories exempted from SSL proxy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

The `crl` statement is supported from 15.1X49-D30.

The `logical system` option is introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

SSL Proxy Overview

Configuring SSL Forward Proxy

Enabling Debugging and Tracing for SSL Proxy

profile (SSL Initiation)

IN THIS SECTION

Syntax | 13344

- Hierarchy Level | 13344
- Description | 13345
- Options | 13345
- Required Privilege Level | 13346
- Release Information | 13346

Syntax

```
profile profile-name {
    actions {
        crl {
            disable disable;
            if-not-present (allow | drop);
            ignore-hold-instruction-code ignore-hold-instruction-code;
        }
        ignore-server-auth-failure ignore-server-auth-failure;
    }
    client-certificate client-certificate;
    crypto-hardware-offload {
        tpm;
    }
    custom-ciphers ;
    enable-flow-tracing  enable-flow-tracing;
    enable-session-cache enable-session-cache;
    preferred-ciphers (custom | medium | strong | weak);
    protocol-version (all | ssl3 | tls1 | tls11 | tls12);
    trusted-ca  ;
}
```

Hierarchy Level

```
[edit services ssl initiation]
```

Description

Specify the name of the profile for SSL initiation support service.

SSL initiation is a process where the SRX Series Firewall acts as in SSL proxy client, initiates the SSL sessions with an SSL server. The SRX Series Firewall receives clear text from an HTTP client. It encrypts and transmits the data as ciphertext to the SSL server. On the reverse side, the SRX Series decrypts the ciphertext that it receives from the SSL server and sends the data to the client as clear text.

The profile contains the settings for the SSL-initiated connections. This includes the list of supported ciphers and their priority, the supported versions of SSL/TLS, and a few other options.

Options

actions

Specify the certification revocation checks and traffic related actions for configuring SSL initiation support service.

- **crl**

  Specify certificate revocation actions. The certificate revocation list (CRL) contains the list of digital certificates that have been canceled before their expiration date. When a participating device uses a digital certificate, it checks the certificate signature and validity. It also acquires the most recently issued CRL and checks that the certificate serial number is not on that CRL. By default, CRL verification is enabled on SSL profile.

- **ignore-hold-instruction-code**

  Ignore server authentication failure. By selecting this option, you can choose to ignore certificate validation, root CA expiration dates, and other such issues based on your requirements.

client-certificate

Local certificate. It is a certificate that client connects to server with. It is usually signed by a CA that the SRX Series Firewall trusts.

crypto-hardware-offload

Trusted Platform Module (TPM) mode. Configure tpm based SSL Initiation sessions.

custom-ciphers

Configure custom cipher for an SSL profile.

Custom ciphers allow you to define your own cipher list. If you do not want to use one of the three categories (strong, medium, or week) of preferred ciphers, you can select ciphers from each of the categories to form a custom cipher set.

To configure custom ciphers, you must set preferred-ciphers to custom. See "preferred-ciphers" on page 13055 for more details.

enable-flow-tracing Enable flow tracing to enable debug tracing.

enable-session-cache Enable SSL session cache. You can enable session caching to cache session information, such as the pre-master secret key and agreed-upon ciphers, for both the client and server.

ignore-server-auth-failure Ignore server authentication completely. In this case, SSL forward proxy ignores errors encountered during the server certificate verification process (such as CA signature verification failure, self-signed certificates, and certificate expiry).

preferred-ciphers Select preferred ciphers. Preferred ciphers allow you to define an SSL cipher that can be used with acceptable key strength. Ciphers are divided in three categories depending on their key strength: strong, medium, or weak.

protocol-version Specify the accepted SSL protocol version. You can specify the SSL/TLS protocol version the security device uses to negotiate in SSL connections.

trusted-ca List of trusted certificate authority profiles. SSL forward proxy uses trusted CA certificates for server authentication. Junos OS provides a default list of trusted CA certificates that you can easily load on to your system using a default command option.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. The `protocol-version` statement is updated to include `tls11` and `tls12` from Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

profile (SSL Termination)

IN THIS SECTION

- Syntax | 13347
- Hierarchy Level | 13347
- Description | 13347
- Options | 13348
- Required Privilege Level | 13348
- Release Information | 13349

Syntax

```
profile name {
    custom-ciphers;
    enable-flow-tracing   enable-flow-tracing;
    enable-session-cache enable-session-cache;
    preferred-ciphers (custom | medium | strong | weak);
    protocol-version (all | ssl3 | tls1 | tls11 | tls12);
    server-certificate  server-certificate;
    trusted-ca  ;
}
```

Hierarchy Level

```
[edit services ssl termination]
```

Description

Specify the name of the profile for SSL termination support service.

Traffic from the client to SRX Series is encrypted and terminated at SRX Series, which then re-encrypts traffic to the back-end server.

SSL termination is a process where the SRX Series Firewall acts as an SSL proxy server, terminates the SSL session from the client. The SRX Series Firewall receives encrypted data from the HTTP client. It decrypts and transmits the data as unencrypted request to the other servers (HTTP server).

The profile contains the settings for the SSL-terminated connections. This includes the list of supported ciphers and their priority, the supported versions of SSL/TLS, and a few other options.

Options

custom-ciphers

Configure custom cipher for an SSL profile.

Custom ciphers allow you to define your own cipher list. If you do not want to use one of the three categories (strong, medium, or week) of preferred ciphers, you can select ciphers from each of the categories to form a custom cipher set.

To configure custom ciphers, you must set preferred-ciphers to custom. See "preferred-ciphers" on page 13055 for more details.

enable-flow-tracing

Enable flow tracing to enable debug tracing.

enable-session-cache

Enable SSL session cache. You can enable session caching to cache session information, such as the pre-master secret key and agreed-upon ciphers, for both the client and server.

preferred-ciphers

Select preferred ciphers. Preferred ciphers allow you to define an SSL cipher that can be used with acceptable key strength. Ciphers are divided in three categories depending on their key strength: strong, medium, or weak.

protocol-version

Specify the accepted SSL protocol version. You can specify the SSL/TLS protocol version the security device uses to negotiate in SSL connections.

server-certificate

Local certificate identifier. Server certificates are used to authenticate the identity of a server. A server is required to present a certificate as part of the initial connection setup. SSL proxy generates a new certificate by replacing the original issuer of the certificate with its own identity and signs this new certificate with its own public key (provided as a part of the proxy profile configuration).

trusted-ca

List of trusted certificate authority profiles. SSL forward proxy uses trusted CA certificates for server authentication. Junos OS provides a default list of trusted CA certificates that you can easily load on to your system using a default command option.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. The `protocol-version` statement is updated to include `tls11` and `tls12` from Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

> *Configuring SSL Forward Proxy*
>
> Firewall User Authentication Overview

profile (Traffic Load Balancer)

IN THIS SECTION

- Syntax | 13349
- Hierarchy Level | 13350
- Description | 13350
- Options | 13350
- Required Privilege Level | 13352
- Release Information | 13352

Syntax

```
profile profile-name {
    custom {
        cmd priority {
            default-real-service-status (down | up);
            expect (ascii | binary) receive-string;
            port port;
            real-service-action (down | up);
            send (ascii | binary) send-string;
```

```
        }
        protocol (tcp | udp);
    }
    failure-retries number-of-retries;
    http {
        host hostname;
        method (get | option);
        port http-port-number;
        url url;
    }
    icmp;
    probe-interval interval;
    recovery-retries number-of-recovery-retries;
    ssl-hello {
        port port;
        ssl-version;
    }
    tcp {
        port tcp-port-number;
    }
}
```

Hierarchy Level

```
[edit services network-monitoring]
```

Description

Configure a monitoring profile that can be used for health-checking a group of TLB servers.

Options

custom Use custom probes for server health checking.

cmd *priority* Use the specified command priority to send for a custom probe.

- **Values:** 1 or 2

default-real-service-status (down | up) Assign a server status for when the probe times out. The **up** value is used when the server or the intermediate network nodes are only expected to send a negative response to a probe.

- **Default:** down

expect (ascii | binary) ***receive-string*** Use the specified ascii or binary string as an expected probe response.

- **Range:** 1 through 512 characters

port ***port*** Use the specified port for custom probes.

protocol (tcp | udp) Use the selected protocol for custom probes.

real-service-action (down | up) Assign a server status for when the expected response to the probe is received.

- **Default:** down

send (ascii | binary) ***send-string*** Send the specified ascii or binary string as a probe.

- **Range:** 1 through 512 characters

failure-retries ***number-of-retries*** Use the specified number of probes that are sent after which the real server is tagged as down.

- **Default:** 5

http Use HTTP probes for server health checking.

host ***hostname*** Use the specified hostname for HTTP probes for server health checks.

method (get | option) Use the get or option HTTP method for server health checks.

port ***http-port-number*** Use the specified port number for HTTP probes.

url ***url*** Use the specified URL for HTTP probes. Maximum length is 128 bytes.

icmp Use ICMP probes for server health checking.

probe-interval ***interval*** Use the specified interval of time, in seconds, at which health check probes are sent.

- **Default:** 5

profile-name Identifier for the network monitoring profile.

recovery-retries *number-of-recovery-retries* Use the specified number of successful probe attempts after which the server is declared up.

- **Default:** 5

ssl-hello Use a `Client Hello` for server health checks

port *port* Use the specified port number for `Client Hello` server health checks.

ssl-version SSL version.

- **Default:** 3

tcp Use TCP probes for server health checks.

port *tcp-port-number* Use the specified port number for TCP probes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Traffic Load Balancer Overview

Configuring TLB

profile-type (Dynamic Service Profiles)

IN THIS SECTION

- Syntax | 13353
- Hierarchy Level | 13353
- Description | 13353
- Options | 13354
- Required Privilege Level | 13354
- Release Information | 13354

Syntax

```
profile-type remote-device-service;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Specify that the dynamic service profile containing this statement is not applied locally to the router. Instead, it is applied to an external device by means of the remote device services manager daemon (rdmd). It enables an external authority, such as PCRF to reference the dynamic service profile to provision or deprovision services (charging rules) on the remote device.

The content of the service is limited to a set of variables that are translated to NETCONF XML protocol remote procedure calls (RPCs) by a dictionary provisioned on the MX series router. The dictionary contains entries corresponding to different service profiles. Each entry contains the RPCs that correspond to the service variables and install the service on the remote device.

The `remote-device-service` type effectively separates the external authority from how and where a service is provisioned. The external authority simply references the profile name; if the service profile has this type, then the service is provisioned on a remote device. If the profile does not have this type, then the service is provisioned on the router.

Options

remote-device-service Service profile type that causes the profile to be applied to an external device by means of the remote device services manager daemon (rdmd).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Configuring Remote Device Management for Service Provisioning

Remote Device Services Manager (RDSM) Overview

profile-variable-set (Dynamic Profiles)

IN THIS SECTION

- Syntax | 13355
- Hierarchy Level | 13355
- Description | 13355
- Options | 13355
- Required Privilege Level | 13355
- Release Information | 13355

Syntax

```
profile-variable-set {
    variable-set-name {
        junos-layer2-output-policer policer-name;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Specify the policer used in the variable set.

Options

junos-layer2-output-policer *policer-name*—Layer 2 policer that you want to substitute in the dynamic profile. You define this policer at the `[edit firewall policer]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Applying the Policers to Dynamic Profile Interfaces

Creating a Dynamic Profile for the Complex Configuration

profile-variable-set (Routing Instances)

IN THIS SECTION

- Syntax | 13356
- Hierarchy Level | 13356
- Description | 13356
- Options | 13356
- Required Privilege Level | 13356
- Release Information | 13357

Syntax

```
profile-variable-set variable-set-name
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols vpls associate-profile]
```

Description

Specify the variable set to apply to the dynamic profile for the routing instance.

Options

variable-set-name—Name of the variable set to use when this dynamic profile is applied to the routing instance. You define this variable set at the `[edit dynamic-profiles profile-name]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Attaching Dynamic Profiles to Routing Instances for the Complex Configuration

profile (Web Filter)

IN THIS SECTION

- Syntax | 13357
- Hierarchy Level (starting in Junos OS Release 18.3R1 | 13359
- Hierarchy Level (before Junos OS Release 18.3R1) | 13359
- Description | 13359
- Options | 13359
- Required Privilege Level | 13360
- Release Information | 13360

Syntax

```
profile profile-name {
dns-filter {
    database-file filename;
    dns-resp-ttl seconds;
    dns-server [ ip-address ];
    hash-key key-string;
    hash-method hash-method-name;
    statistics-log-timer minutes;
    wildcarding-level level;
}
dns-filter-template template-name {
    client-interfaces [ client-interface-name ];
```

```
    client-routing-instance client-routing-instance-name;
    dns-filter {
        database-file filename;
        dns-resp-ttl seconds;
        dns-server [ ip-address ];
        hash-key key-string;
        hash-method hash-method-name;
        statistics-log-timer minutes;
        wildcarding-level level;
    }
    server-interfaces  [ server-interface-name ];
    server-routing-instance server-routing-instance-name;
    term term-name {
        from {
            src-ip-prefix [ source-prefix ];
        }
        then {
            accept;
            dns-sinkhole;
        }
    }
}
global-dns-stats-log-timer minutes;
url-filter-database filename;
(url-filter-template | template) template-name {
    client-interfaces [ client-interface-name1 client-interface-name2 ];
    disable-url-filtering;
    dns-resolution-interval minutes;
    dns-resolution-rate seconds;
    dns-retries number;
    dns-routing-instance dns-routing-instance-name;
    dns-server [ ip-address1 ip-address2 ip-address3 ];
    dns-source-interface loopback-interface-name;
    dns-routing-instance dns-routing-instance-name;
    routing-instance routing-instance-name;
    server-interfaces  [ server-interface-name1 server-interface-name2 ];
    term term-name {
        from {
            src-ip-prefix [prefix1 prefix2];
            dest-port   [port1 port2];
        }
        then {
            accept;
```

```
            custom-page custom-page;
            http-status-code http-status-code;
            redirect-url redirect-url;
            tcp-reset;
        }
    }
    url-filter-database filename
}
```

Hierarchy Level (starting in Junos OS Release 18.3R1

```
[edit services web-filter]
```

Hierarchy Level (before Junos OS Release 18.3R1)

```
[edit services url-filter]
```

Description

Define URL filter profile or DNS filter profile.

A URL filter profile is for filtering access to disallowed URLs. A URL filter profile includes a general database setting and templates. The template settings apply to specific interfaces or to access from specific source IP address prefixes, and override the database setting at the profile level.

A DNS filter profile is used to filter DNS requests for disallowed website domains. A DNS filter profile includes general DNS filtering settings and up to 32 templates. The template settings apply to DNS requests on specific interfaces or to DNS requests from specific source IP address prefixes, and override the corresponding settings at the profile level. You can configure up to eight DNS filter profiles.

NOTE: For URL filtering, use the `url-filter-template` option starting in Junos OS Release 18.3R1 and use the `template` option in Junos OS Releases before 18.3R1.

Options

profile-name	Name of the filter profile.

`url-filter-database` *filename* Specify the filename of the URL filter database. This option is mandatory.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

`dns-filter`, `dns-filter-templates`, `global-dns-stats-log-timer`, and `url-filter-template` options introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

Configuring URL Filtering

profilerd

IN THIS SECTION

- Syntax | 13361
- Hierarchy Level | 13361
- Description | 13361
- Options | 13361
- Required Privilege Level | 13361
- Release Information | 13362

Syntax

```
profilerd {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the profiler process.

Options

- `command` *`binary-file-path`*—Path to binary for process.
- `disable`—Disable the profiler process.
- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

profiles

IN THIS SECTION

- Syntax | 13362
- Hierarchy Level | 13362
- Description | 13363
- Options | 13363
- Required Privilege Level | 13363
- Release Information | 13363

Syntax

```
profiles {
    profile-name {
        sip-user-id simple-ip-user-id;
        sip-password simple-ip-password;
        access-point-name apn;
        authentication-method (pap | chap | none);
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options]
```

Description

Configure a profile to establish a data call with a Global System for Mobile Communications (GSM) cellular network. You can configure up to 16 profiles.

Options

profile-name—Name of the profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

profiles (AAA)

IN THIS SECTION

- Syntax | 13364
- Hierarchy Level | 13364
- Description | 13364
- Options | 13364
- Required Privilege Level | 13365
- Release Information | 13365

Syntax

```
profiles aaa-profile-name {
    radius {
        accounting {
            network-element network-element-name;
        }
        authentication {
            network-element network-element-name;
        }
        policy {
            activation-attribute {
                <code numeric-code;>
                <vendor-id vendor-id;>
            }
            deactivation-attribute {
                <code numeric-code;>
                <vendor-id vendor-id;>
            }
            coa-accounting (enable | disable);
        }
    }
}
```

Hierarchy Level

```
[edit unified-edge aaa]
```

Description

Configure a profile of the policy control attributes for RADIUS servers. This profile is used by the policy and charging enforcement function (PCEF) profile.

Options

aaa-profile-name	Name of the AAA profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring an AAA Profile

Understanding AAA Profiles

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

profiles (PCEF)

IN THIS SECTION

- Syntax | 13365
- Hierarchy Level | 13366
- Description | 13366
- Options | 13367
- Required Privilege Level | 13367
- Release Information | 13367

Syntax

```
profiles profile-name {
    aaa-policy-control {
        aaa-profile aaa-profile-name;
        pcc-rulebases [rulebase-name <time-of-day-profile profile-name>];
```

```
        user-password password;
    }
    dynamic-policy-control {
        pcc-rules {
            [rule-name  precedence number <time-of-day-profile profile-name>];
        }
        pcc-rulebases {
            [rulebase-name <time-of-day-profile profile-name>];
        }
        diameter-profile gx-profile-name;
    }
    static-policy-control {
        pcc-rules {
            [rule-name precedence number <time-of-day-profile profile-name>];
        }
        pcc-rulebases {
            [rulebase-name <time-of-day-profile profile-name>];
        }
    }
```

Hierarchy Level

```
[edit unified-edge pcef],
[edit services pcef]
```

Description

Set up the overall policy and charging enforcement function (PCEF) configuration that can be applied to subscribers.

> **NOTE**: You can configure only one of the following statements in a PCEF profile: `aaa-policy-control`, `static-policy-control`, or `dynamic-policy-control`.

You can configure a maximum of 32 policy and charging control (PCC) rules in a PCEF profile. There is no limit to the number of PCC rulebases you can configure in a PCEF profile.

If you are using Junos OS Subscriber Aware, configure the PCEF profile at the `[edit unified-edge pcef]` hierarchy level. You then assign this profile to the subscriber's TDF domain or to the domain selection configuration.

If you are using Junos OS Broadband Subscriber Management, configure the PCEF profile at the `[edit services pcef]` hierarchy level. The `static-policy-control` option is applicable to PCC rule activation through a dynamic profile, and you assign the PCEF profile to the dynamic profile. Starting in Junos OS Release 18.2R1, the `dynamic-policy-control` option is also available and is applicable to direct rule activation by a policy and charging rules function (PCRF) server; you assign the PCEF profile to the access profile. The `aaa-policy-control` option is not applicable to subscriber management.

Options

profile-name Name of the PCEF profile.

- **Range:** 1 through 63 characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

- Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Dynamic Policies
- Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Static Policies

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Policies That a RADIUS Server Controls

Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management

profiles (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 13368
- Hierarchy Level | 13368
- Description | 13369
- Options | 13369
- Required Privilege Level | 13369
- Release Information | 13369

Syntax

```
profiles {
    test-profile profile-name {
        test-type (throughput | latency | frame-loss | back-back-frames);
        packet-size bytes;
        step-percent percent;
        bandwidth-kbps kpbs;
    }
}
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarking]
```

Description

Configure the test profile to specify attributes, such as the period for the test and the type of test to be performed, for the RFC 2544-based benchmarking test. The test profile is referenced in the test interface to perform a specific type of benchmarking test and compute statistics to describe the performance characteristics of a network interconnecting device.

Options

`profiles` Define the test profile for the RFC 2544-based benchmarking test to examine and analyze the performance characteristics of a network interconnecting device.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

programmable-rpd

IN THIS SECTION

- Syntax (Junos OS) | **13370**
- Syntax (Junos OS Evolved) | **13370**

- Hierarchy Level | 13371
- Description | 13371
- Required Privilege Level | 13371
- Release Information | 13371

Syntax (Junos OS)

```
programmable-rpd {
 client name {
  interface-notification name;
 }
 purge-timeout (never | timeout);
 rib-service {
  dynamic-next-hop-interface (disable | enable);
 }
 statistics {
  flex-routes;
   <destination ip-destination-address>;
   table;
 }
 traceoptions {
  file filename <files files> <size size> <(world-readable | no-world-readable)>;
  flag name;
 }
}
```

Syntax (Junos OS Evolved)

```
programmable-rpd {
 client name {
  interface-notification name;
 }
 purge-timeout (never |  timeout);
 rib-service {
  dynamic-next-hop-interface (disable);
 }
```

```
  stats-aggregation-interval seconds;
  stats-sampling-interval seconds;
  traceoptions {
   file filename <files files> <size size> <(world-readable | no-world-readable)>;
   flag name;
  }
 }
```

Hierarchy Level

```
[edit logical-systems name routing-options],
[edit routing-options]
```

Description

Configure the programmable routing protocol process (prpd). The prpd provides public APIs to program routing systems, making it possible for users to directly access the APIs to customize, create, and modify behavior of their network.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

`stats-aggregation-interval` and `stats-sampling-interval` statements introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

show programmable-rpd clients | 23287

promiscuous

IN THIS SECTION

- Syntax | 13372
- Hierarchy Level | 13372
- Description | 13372
- Required Privilege Level | 13372
- Release Information | 13372

Syntax

```
promiscuous;
```

Hierarchy Level

```
[edit vlans vlan-name interfaces interface-name]
```

Description

Configure an access or trunk port to be promiscuous.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Creating a Private VLAN on a Single QFX Switch without ELS Support

Creating a Private VLAN Spanning Multiple QFX Series Switches without ELS Support

promiscuous-mode (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | 13373
- Hierarchy Level | 13373
- Description | 13373
- Required Privilege Level | 13373
- Release Information | 13374

Syntax

```
promiscuous-mode;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name],]
```

Description

Specify that the interface accepts IGMP reports from hosts on any subnetwork. When you enable promiscuous mode, all routing devices on the Ethernet segment must be configured for promiscuous mode. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Accepting IGMP Messages from Remote Subnetworks

promiscuous-mode (Interfaces)

IN THIS SECTION

- Syntax | 13374
- Hierarchy Level | 13374
- Description | 13375
- Required Privilege Level | 13375
- Release Information | 13375

Syntax

```
promiscuous-mode;
```

Hierarchy Level

```
[edit interfaces interface-name ]
```

Description

Enable promiscuous mode on Layer 3 Ethernet interfaces. When promiscuous mode is enabled on an interface, all packets received on the interface are sent to the central point or Services Processing Unit regardless of the destination MAC address of the packet.

You can also enable promiscuous mode on chassis cluster redundant Ethernet interfaces and on aggregated Ethernet interfaces. If you enable promiscuous mode on a redundant Ethernet interface, promiscuous mode is then enabled on any child physical interfaces. If you enable promiscuous mode on an aggregated Ethernet interface, promiscuous mode is then enabled on all member interfaces.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure)

promiscuous-mode (Protocols IGMP)

IN THIS SECTION

- Syntax | 13376
- Hierarchy Level | 13376
- Description | 13376
- Required Privilege Level | 13376
- Release Information | 13376

Syntax

```
promiscuous-mode;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Description

Specify that the interface accepts IGMP reports from hosts on any subnetwork. Note that when enabling promiscuous-mode, all routing devices on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Accepting IGMP Messages from Remote Subnetworks

promote

IN THIS SECTION

- Syntax | 13377
- Hierarchy Level | 13377
- Description | 13378
- Options | 13378
- Required Privilege Level | 13378
- Release Information | 13378

Syntax

Syntax for `promote gre-key`

```
promote gre-key;
```

Syntax for `promote dscp`

```
promote dscp;
```

Syntax for `promote traffic-class`

```
promote traffic-class;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name]
```

Description

Use the `promote gre-key` or `promote dscp` configuration statements to use `gre-key` or `dscp` match conditions in an inet firewall filter. Use the `promote traffic-class` configuration statement to use the `traffic-class` match condition in an inet6 firewall filter.

> **NOTE**: When you apply any of these configuration statements, the entire filter is compiled in a way to optimize the performance of these match conditions for their respective firewall filters. As an example, when you configure `promote gre-key` and use `gre-key` in any of the terms in a filter, the entire filter is compiled in a way that optimizes performance of the filter for gre-key matching.

> **NOTE**: All these configuration statements are supported on PTX Series routers only when network services is set to `enhanced-mode`. For more information, see `enhanced-mode` .

Options

The following are the options for the `promote` configuration command.

- `gre-key`
- `dscp`
- `traffic-class`

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F3 and 16.1R2.

RELATED DOCUMENTATION

Guidelines for Configuring Firewall Filters

Guidelines for Applying Standard Firewall Filters

Firewall Filter Match Conditions for IPv4 Traffic

propagate-settings

IN THIS SECTION

- Syntax | 13379
- Hierarchy Level | 13379
- Description | 13379
- Options | 13379
- Required Privilege Level | 13380
- Release Information | 13380

Syntax

```
propagate-settings interface-name;
```

Hierarchy Level

```
[edit system services dhcp]
[edit system services dhcp pool]
```

Description

Enable or disable the propagation of TCP/IP settings received on the device acting as Dynamic Host Configuration Protocol (DHCP) client. The settings can be propagated to the server pool running on the device. Use the `system services dhcp` to set this feature globally. Use the `system services dhcp pool` to set the feature for the address pool and override the global setting.

Options

`logical-interface-name` —Name of the logical interface to receive TCP/IP settings from the external network for propagation to the DHCP pool running on the device.

Required Privilege Level

system—To view this in the configuration.

system-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

propagate-tc

IN THIS SECTION

- Syntax | 13380
- Hierarchy Level | 13380
- Description | 13381
- Required Privilege Level | 13381
- Release Information | 13381

Syntax

```
propagate-tc;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Enable topology change propagation from a sub-ring to an interconnected major-ring. By default, topology change propagation is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

propagation-delay

IN THIS SECTION

- Syntax | 13382
- Hierarchy Level | 13382
- Description | 13382
- Options | 13382
- Required Privilege Level | 13382
- Release Information | 13382

Syntax

```
propagation-delay milliseconds;
```

Hierarchy Level

```
[edit protocols pim],
[edit protocols pim interface (Protocols PIM)  interface-name],
[edit routing-instances routing-instance-name protocols pim],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name],
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name]
```

Description

Set a delay for implementing a PIM prune message on the upstream routing device on a multicast network for which join suppression has been enabled. The routing device waits for the prune pending period to detect whether a join message is currently being suppressed by another routing device.

Options

milliseconds—Interval for the prune pending timer, which is the sum of the **propagation-delay** value and the **override-interval** value.

- **Range:** 250 through 2000 milliseconds
- **Default:** 500 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Example: Enabling Join Suppression

override-interval | 12058

reset-tracking-bit | 15583

proposal (Security Group VPN Member IKE)

IN THIS SECTION

- Syntax | 13383
- Hierarchy Level | 13383
- Description | 13384
- Options | 13384
- Required Privilege Level | 13385
- Release Information | 13385

Syntax

```
proposal proposal-name {
    authentication-algorithm (sha-256 | sha-384);
    authentication-method pre-shared-keys;
    description description;
    dh-group (group14 | group24);
    encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
    lifetime-seconds seconds;
}
```

Hierarchy Level

```
[edit security group-vpn member ike]
```

Description

Define an IKE proposal. You can configure one or more IKE proposals. Each proposal is a list of IKE attributes to protect the IKE connection between the IKE host and its peer.

Options

`proposal` *`proposal-name`*—Name of the IKE proposal. The proposal name can be up to 32 alphanumeric characters long.

`authentication-algorithm`—Configure the Internet Key Exchange (IKE) authentication algorithm. Hash algorithm that authenticates packet data. It can be one of the following algorithms:

- `sha-256`—Produces a 256-bit digest. This is the default value.
- `sha-384`—Produces a 384-bit digest.

`authentication-method` *`pre-shared-keys`*—Specify the method the device uses to authenticate the source of Internet Key Exchange (IKE) messages. The `pre-shared-keys` option refers to a preshared key, which is a secret key shared between the two peers, is used during authentication to identify the peers to each other. The same key must be configured for each peer. This is the default method.

`description` *`description`*—Specify descriptive text for an IKE proposal.

`dh-group`—Specify the IKE Diffie-Hellman group for key establishment.

- `group14`—2048-bit group. This is the default value.
- `group24`—2048-bit, 256 bit subgroup. Support for the `group24` option added in Junos OS Release 15.1X49-D30 for vSRX Virtual Firewall.

`encryption-algorithm`—Configure an encryption algorithm for an IKE proposal.

- `aes-128-cbc`—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- `aes-192-cbc`—AES 192-bit encryption algorithm.
- `aes-256-cbc`—AES 256-bit encryption algorithm.

`lifetime-seconds` *`seconds`*—Specify the lifetime (in seconds) of an IKE or IPsec security association (SA) for group VPN. When the SA expires, it is replaced by a new SA and security parameter index (SPI) or terminated.

- **Range:** 180 through 86,400 seconds
- **Default:** 3600 seconds

The device does not delete existing IPsec SAs when you update the `authentication-algorithm`, `authentication-method`, `dh-group`, and `encryption-algorithm` configuration in the IKE proposal.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

proposal (Security Group VPN Server IKE)

IN THIS SECTION

- Syntax | 13385
- Hierarchy Level | 13386
- Description | 13386
- Options | 13386
- Required Privilege Level | 13387
- Release Information | 13387

Syntax

```
proposal proposal-name {
    authentication-algorithm (sha-256 | sha-384);
    authentication-method pre-shared-keys;
    description description;
```

```
    dh-group (group14 | group24);
    encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
}
```

Hierarchy Level

```
[edit security group-vpn server ike]
```

Description

Define an IKE proposal for group VPN server. You can configure one or more IKE proposals. Each proposal is a list of IKE attributes to protect the IKE connection between the IKE host and its peer.

Options

`proposal` *`proposal-name`*—Name of the IKE proposal. The proposal name can be up to 32 alphanumeric characters long.

`authentication-algorithm`—Configure the Internet Key Exchange (IKE) authentication algorithm. Hash algorithm that authenticates packet data. It can be one of the following algorithms:

- `sha-256`—Produces a 256-bit digest. This is the default value.
- `sha-384`—Produces a 384-bit digest.

`authentication-method` *`pre-shared-keys`*—Specify the method the device uses to authenticate the source of Internet Key Exchange (IKE) messages. The `pre-shared-keys` option refers to a preshared key, which is a secret key shared between the two peers, is used during authentication to identify the peers to each other. The same key must be configured for each peer. This is the default method.

`description` *`description`*—Specify descriptive text for an IKE proposal.

`dh-group`—Specify the IKE Diffie-Hellman group for key establishment.

- `group14`—2048-bit group. This is the default value.
- `group24`—2048-bit, 256 bit subgroup. Support for the `group24` option added in Junos OS Release 15.1X49-D30 for vSRX Virtual Firewall.

`encryption-algorithm`—Configure an encryption algorithm for an IKE proposal.

- `aes-128-cbc`—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- `aes-192-cbc`—AES 192-bit encryption algorithm.

- `aes-256-cbc`—AES 256-bit encryption algorithm.

The device does not delete existing IPsec SAs when you update the `authentication-algorithm`, `authentication-method`, `dh-group`, and `encryption-algorithm` configuration in the IKE proposal.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

ike (Security Group VPN Server) | 8440

proposal (Security Group VPN Server IPsec)

IN THIS SECTION

- Syntax | 13388
- Hierarchy Level | 13388
- Description | 13388
- Options | 13388
- Required Privilege Level | 13389
- Release Information | 13389

Syntax

```
proposal proposal-name {
    authentication-algorithm (hmac-sha-256-128);
    description description;
    encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
    lifetime-seconds seconds;
}
```

Hierarchy Level

```
[edit security group-vpn server ipsec]
```

Description

Define an IPsec proposal. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 devices and vSRX Virtual Firewall instances.

Options

proposal-name—Name of the IPsec proposal.

`authentication-algorithm` *`hmac-sha-256-128`*—Configure the IPsec authentication algorithm. Produces a 256-bit digest, truncated to 128 bits. This is the default value.

`description` *`description`*—Text the description of IPsec proposal.

`encryption-algorithm`—Configure an encryption algorithm. The device deletes existing IPsec SAs when you update the `encryption-algorithm` configuration in the IPsec proposal.

- `aes-128-cbc`—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- `aes-192-cbc`—AES 192-bit encryption algorithm.
- `aes-256-cbc` —AES 256-bit encryption algorithm. This is the default value.

`lifetime-seconds` *`seconds`*—Specify the lifetime (in seconds) of an IPsec security association (SA) for group VPN. When the SA expires, it is replaced by a new SA and security parameter index (SPI) or terminated. Specify a value from 180 to 86,400 seconds. The default is 3600 seconds.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

ipsec (Security Group VPN Server) | 9317

proposal (Security IKE)

IN THIS SECTION

- Syntax | 13389
- Hierarchy Level | 13390
- Description | 13390
- Options | 13390
- Required Privilege Level | 13392
- Release Information | 13392

Syntax

```
proposal proposal-name {
    authentication-algorithm (md5 | sha-256 | sha-384| sha1 | sha-512);
    authentication-method (certificates | digital-signature | dsa-signatures | ecdsa-
signatures-256 | ecdsa-signatures-384 | ecdsa-signatures-521 | pre-shared-keys | rsa-signatures
| ecdsa-signatures-521);
    description description;
```

```
    dh-group (group1 | group14 | group19 | group2 | group20 | group24 | group5 | group15 |
group16 | group21);
    encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-256-cbc |
aes-256-gcm | chacha20-poly1305 | des-cbc);
    lifetime-seconds seconds;
    Signature Hash Algorithm (Security IKE) (sha1 | sha256 | sha384 | sha512);
}
```

Hierarchy Level

```
[edit security ike]
```

Description

Define an IKE proposal.

Options

proposal-name

Name of the IKE proposal. The proposal name can be up to 32 alphanumeric characters long.

authentication-algorithm

Configure the Internet Key Exchange (IKE) authentication hash algorithm that authenticates packet data. It can be one of the following algorithms:

- `md5`—Produces a 128-bit digest.
- `sha-256`—Produces a 256-bit digest.
- `sha-384`—Produces a 384-bit digest.
- In Power Mode IPSec mode and in normal mode—
  - `sha1`—Produces a 160-bit digest.
  - `sha-512`—Produces a 512-bit digest.

The device deletes existing IPsec SAs when you update the `authentication-algorithm` configuration either in the IKE proposal or IPsec proposal.

authentication-method

Specify the method the device uses to authenticate the source of Internet Key Exchange (IKE) messages. The `pre-shared-keys` option refers to a preshared key, which is a key for encryption and decryption that both participants must have

before beginning tunnel negotiations. The other options refer to types of digital signatures, which are certificates that confirm the identity of the certificate holder. The device deletes existing IPsec SAs when you update the `authentication-method` configuration in the IKE proposal.

- `certificates`—You can establish the IKEv2 and IPsec SA tunnels irrespective of the type of certificate used on initiator and responder. The `authentication-method certificates` option cannot be used with IKEv1.
- `digital-signature`—Specify digital signature as an authentication method for IKEv2. Note that you cannot use the option with IKEv1.
- `dsa-signatures`—Specify that the Digital Signature Algorithm (DSA) is used.
- `ecdsa-signatures-256`—Specify that the Elliptic Curve DSA (ECDSA) using the 256-bit elliptic curve secp256r1, as specified in the *Federal Information Processing Standard (FIPS) Digital Signature Standard (DSS) 186-3*, is used.
- `ecdsa-signatures-384`—Specify that the ECDSA using the 384-bit elliptic curve secp384r1, as specified in the *FIPS DSS 186-3*, is used.
- `pre-shared-keys`—Specify that a preshared key, which is a secret key shared between the two peers, is used during authentication to identify the peers to each other. The same key must be configured for each peer. This is the default method.
- `rsa-signatures`—Specify that a public key algorithm, which supports encryption and digital signatures, is used.
- `ecdsa-signatures-521`—Specify that the ECDSA using the 521-bit elliptic curve secp521r1 is used.

description *description*

Text the description of IKE proposal.

dh-group

Specify the IKE Diffie-Hellman group.

encryption-algorithm

Configure an encryption algorithm for an IKE proposal.

lifetime-seconds *seconds*

Specify the lifetime (in seconds) of an IKE security association (SA). When the SA expires, it is replaced by a new SA and security parameter index (SPI) or terminated.

- Range: 180 through 86,400 seconds
- Default: 28,800 seconds

signature-hash-algorithm (Optional) Specify the digital signature hash algorithms for IKEv2. Note that you cannot use the option with IKEv1.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

Support for `dh-group group 14` and `dsa-signatures` added in Junos OS Release 11.1.

Support for `sha-384`, `ecdsa-signatures-256`, `ecdsa-signatures-384`, `group19`, `group20`, and `group24` options added in Junos OS Release 12.1X45-D10.

Support for `ecdsa-signatures-256` and `ecdsa-signatures-384` options added in Junos OS Release 12.1X45-D10.

Support for `sha-512`, `group15`, `group16`, `group21`, and `ecdsa-signatures-521` options added in Junos OS Release 19.1R1 on SRX5000 line of devices with `junos-ike` package installed.

Support for authentication algorithm (SH1: hmac-sha1-96) added to vSRX Virtual Firewall in Junos OS Release 19.3R1 for Power Mode IPSec mode, along with the existing support in normal mode.

Support for `group15`, `group16`, and `group21` options added in Junos OS Release 20.3R1 on vSRX Virtual Firewall instances with `junos-ike` package installed.

Support for `group15`, `group16`, and `group21` options added in Junos OS Release 21.1R1 on vSRX Virtual Firewall 3.0 instances with `junos-ike` package installed.

Support for `certificates` option added in Junos OS Release 22.4R1 on MX240, MX480, and MX960 in USF mode, SRX1500, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0 running the iked process.

Support for the `chacha20-poly1305` option added to SRX1600, SRX2300, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0 in Junos OS Release 24.2R1.

Support for the `digital-signature` and `signature-hash-algorithm` options added in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

IPsec Overview

ike (Security Services) | 8446

Configuring an IKE Proposal for Dynamic SAs

proposal (Security IKE Dynamic SA)

IN THIS SECTION

- Syntax | 13393
- Hierarchy Level | 13393
- Description | 13393
- Options | 13394
- Required Privilege Level | 13394
- Release Information | 13394

Syntax

```
proposal ike-proposal-name {
    authentication-algorithm (md5 | sha1 | sha-256);
    authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group2 | group 5 | group14);
    encryption-algorithm algorithm;
    lifetime-seconds seconds;
}
```

Hierarchy Level

```
[edit security ike]
```

Description

Define an IKE proposal for a dynamic SA.

Options

ike-proposal-name—Specify an IKE proposal name.

The remaining statements are explained separately.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an IKE Proposal for Dynamic SAs

proposal (Security IPsec)

IN THIS SECTION

- Syntax | 13395
- Hierarchy Level | 13395
- Description | 13395
- Options | 13395
- Required Privilege Level | 13398
- Release Information | 13398

Syntax

```
proposal proposal-name {
    authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha-256-96 | hmac-sha-384 |
hmac-sha-512 | hmac-sha1-96);
    description description;
    encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-192-gcm |
aes-256-cbc | aes-256-gcm | chacha20-poly1305 | des-cbc);
    extended-sequence-number;
    lifetime-kilobytes kilobytes;
    lifetime-seconds seconds;
    protocol (ah | esp);
}
```

Hierarchy Level

```
[edit security ipsec]
```

Description

Define an IPsec proposal. An IPsec proposal lists protocols and algorithms (security services) to be negotiated with the remote IPsec peer.

Options

proposal-name Name of the IPsec proposal.

authentication-algorithm Configure the IPsec authentication algorithm. Authentication algorithm is the hash algorithm that authenticates packet data. It can be one of six algorithms:

- Values:

  The hash algorithm to authenticate data can be one of the following:

  - `hmac-md5-96`—Produces a 128-bit digest.

- `hmac-sha-256-128`—Provides data origin authentication and integrity protection. This version of the hmac-sha-256 authenticator produces a 256-bit digest and specifies truncation to 128 bits.
- `hmac-sha1-96`—Hash algorithm that authenticates packet data. It produces a 160-bit digest. Only 96 bits are used for authentication.
- `hmac-sha-512`—Produces a 512-bit digest.
- `hmac-sha-384`—Produces a 384-bit digest.
- `hmac-sha-256-96`—HMAC-SHA-256-96 authentication algorithm (non-RFC compliant)

description

Text description of IPsec proposal

encryption-algorithm

Define encryption algorithm. The device deletes existing IPsec SAs when you update the `encryption-algorithm` configuration in the IPsec proposal.

- Values:
  - `3des-cbc`—Encryption algorithm with block size of 8 bytes (64 bits) and key size of 192 bits.
  - `aes-128-cbc`—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
  - `aes-128-gcm`—AES Galois/Counter Mode (GCM) 128-bit encryption algorithm.

    For an IKE proposal, AES 128-bit authenticated encryption algorithm is supported with IKEv2 only. When this option is used, `aes-128-gcm` should be configured at the [`edit security ipsec proposal` *`proposal-name`*] hierarchy level, and the `authentication-algorithm` option should not be configured at the [`edit security ike proposal` *`proposal-name`*] hierarchy level.

    When `aes-128-gcm`, `aes-192-gcm`, or `aes-256-gcm` encryption algorithms are configured in the IPsec proposal, it is not mandatory to configure AES-GCM encryption algorithm in the corresponding IKE proposal.
  - `aes-192-cbc`—AES 192-bit encryption algorithm.
  - `aes-192-gcm`—AES GCM 192-bit encryption algorithm.
  - `aes-256-cbc`—AES 256-bit encryption algorithm.

- `aes-256-gcm`—AES GCM 256-bit encryption algorithm.

  For an IKE proposal, AES 256-bit authenticated encryption algorithm is supported with IKEv2 only. When this option is used, `aes-256-gcm` should be configured at the [`edit security ipsec proposal` *`proposal-name`*] hierarchy level, and the `authentication-algorithm` option should not be configured at the [`edit security ike proposal` *`proposal-name`*] hierarchy level.

- `chacha20-poly1305`—ChaCha20-Poly1305 authenticated encryption algorithm.

  You can use the algorithm for data path security with Encapsulating Security Payload (ESP) for IPsec VPN. You can configure the algorithm in:

  - PowerMode IPsec (PMI) mode on SRX1600, SRX2300, SRX4300, SRX4600, SRX5400, SRX5600, and SRX5800 Firewalls.
  - PMI and non-PMI modes on vSRX 3.0.

  You can use the algorithm for IPsec when the VPN monitoring feature is not configured at [`edit security ipsec vpn` *`vpn-name`* `vpn-monitor`].

  Note that you cannot use the algorithm on SRX1500 Firewalls.

- `des-cbc`—Encryption algorithm with block size of 8 bytes (64 bits) and key size 48 bits.

extended-sequence-number

Use the `extended-sequence-number` option to enable ESN support. ESN allows IPsec to use 64-bit sequence numbers for the sequence number. If ESN is not enabled, 32-bit sequence number will be used by default. Ensure ESN is not enabled when anti-replay is disabled.

lifetime-kilobytes

Specify the lifetime (in kilobytes) of an IPsec security association (SA). If this statement is not configured, the number of kilobytes used for the SA lifetime is unlimited.

- **Range:** 64000 through 4,294,967,294 kilobytes

lifetime-seconds

Lifetime in seconds.

- **Range:** 180 through 86400
- **Default:** 3600 seconds

protocol

Define the IPsec protocol for a manual or dynamic security association (SA).

- Values:
  - ah—Authentication header
  - esp—Encapsulated Security Payload header

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`extended-sequence-number` option introduced in Junos OS Release 19.4R1.

Starting in Junos OS Release 20.2R1, we've changed the help text description as `NOT RECOMMENDED` for the CLI options `hmac-md5-96`, `hmac-sha1-96`, `3des-cbc`, and `des-cbc` for devices running iked process with `junos-ike` package installed.

`hmac-sha-512` and `hmac-sha-384` options introduced in Junos OS Release 19.1R1 on SRX5000 line of devices with SRX5K-SPC3 card.

Support for `aes-128-gcm`, `aes-192-gcm`, and `aes-256-gcm` options added in Junos OS Release 15.1X49-D70 for vSRX Virtual Firewall.

Support for `aes-128-gcm`, `aes-192-gcm`, and `aes-256-gcm` options added in Junos OS Release 12.1X45-D10.

Support for the `hmac-sha-256-128` added to SRX5400, SRX5600, and SRX5800 devices in Junos OS Release 12.1X46-D20.

Support for the `lifetime-kilobytes` option with IPsec VPN running iked process added in Junos OS Release 23.4R1.

Support for the `chacha20-poly1305` option added to SRX1600, SRX2300, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0 in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Configuring an IPsec Proposal for an ES PIC

Installing Junos IKE package

proposal (Services IKE)

IN THIS SECTION

- Syntax | 13399
- Hierarchy Level | 13399
- Description | 13399
- Options | 13400
- Required Privilege Level | 13400
- Release Information | 13400

Syntax

```
proposal proposal-name {
    authentication-algorithm (md5 | sha1 | sha-256);
    authentication-method (ecdsa-signatures-256 | ecdsa-signatures-384 | pre-shared-keys | rsa-
signatures);
    description description;
    dh-group  (group1 | group2 | group5 |group14 | group15 | group16 | group19 | group20 |
group24);
    encryption-algorithm algorithm;
    lifetime-seconds seconds;
}
```

Hierarchy Level

```
[edit services ipsec-vpn ike]
```

Description

Define an IKE proposal for a dynamic SA.

> **NOTE**: In Junos FIPS mode, ECDSA options of the `authentication-method` statement are not supported in Junos OS Release 17.3R1. Starting in Junos OS Release 17.4R1, ECDSA is supported in Junos FIPS mode.

Options

`proposal-name`—IKE proposal name.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IKE Proposals

proposal (Services IPsec VPN)

IN THIS SECTION

- Syntax | 13401
- Hierarchy Level | 13401
- Description | 13401
- Options | 13401
- Required Privilege Level | 13401
- Release Information | 13401

Syntax

```
proposal proposal-name {
    authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
    description description;
    encryption-algorithm algorithm;
    lifetime-seconds seconds;
    protocol (ah | esp | bundle);
}
```

Hierarchy Level

```
[edit services ipsec-vpn ipsec]
```

Description

Define an IPsec proposal for a dynamic SA.

Options

proposal-name—IPsec proposal name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Proposals

proposals (Services)

IN THIS SECTION

- Syntax | 13402
- Hierarchy Level | 13402
- Description | 13402
- Options | 13402
- Required Privilege Level | 13402
- Release Information | 13403

Syntax

```
proposals [ proposal-names ];
```

Hierarchy Level

```
[edit services ipsec-vpn ike policy policy-name],
[edit services ipsec-vpn ipsec policy policy-name]
```

Description

Define a list of proposals to include in the IKE or IPsec policy.

Options

proposal-names—List of IKE or IPsec proposal names.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IKE Proposals

Configuring IPsec Proposals

proposals

IN THIS SECTION

- Syntax | 13403
- Hierarchy Level | 13403
- Description | 13404
- Options | 13404
- Required Privilege Level | 13404
- Release Information | 13404

Syntax

```
proposals [ proposal-names ];
```

Hierarchy Level

```
[edit security ike policy ike-peer-address],
[edit security ipsec policy ipsec-policy-name]
```

Description

Associate one or more proposals with an IKE or IPsec policy.

Options

proposal-names—Name of one or more proposals.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an IKE Policy for Preshared Keys

Configuring the IPsec Policy for an ES PIC

protect

IN THIS SECTION

- Syntax | 13405
- Description | 13405
- Options | 13405
- Required Privilege Level | 13405
- Release Information | 13405

Syntax

```
protect (hierarchy | statement | identifier)
```

Description

Protect a hierarchy, statement, or identifier from modification or deletion.

Options

hierarchy	(Optional) Protect a specific hierarchy.
statement	(Optional) Protect a specific statement.
identifier	(Optional) Protect a specific identifier.

none

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Example: Protecting the Configuration from Modification or Deletion

unprotect | **30428**

protect (Security)

IN THIS SECTION

- Syntax | 13406
- Hierarchy Level | 13406
- Description | 13406
- Options | 13407
- Required Privilege Level | 13407
- Release Information | 13407

Syntax

```
protect {
    deny {
        all {
            timeout value-in-seconds;
        }
        destination-ip address;
        timeout value-in-seconds;
    }
}
```

Hierarchy Level

```
[edit security alg],
[edit security sip (Security)],
[edit security application-screen (Security SIP)]
```

Description

Configure options to protect servers against INVITE attacks.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sip (Security) | 27977

protect core

IN THIS SECTION

- Syntax | 13407
- Hierarchy Level | 13408
- Description | 13408
- Required Privilege Level | 13408
- Release Information | 13408

Syntax

```
protect core;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name routing-options]
```

Description

In an MPLS VPN Core network, enable BGP Prefix-Independent Convergence (PIC) Edge to install a VPN route in the forwarding table for use as an alternate path when a provider edge (PE) router fails or connectivity to a PE router is lost. This alternate link is used until global convergence through the interior gateway protocol (IGP) occurs. Both OSPF with LDP or RSVP, and IS-IS with LDP or RSVP are supported.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Support in Layer 3 VPN routers is extended to BGP PIC with multiple routes in the global tables such as inet and inet6 unicast, and inet and inet6 labeled unicast in Junos OS Release 15.1.

Support for Resource Reservation Protocol (RSVP) introduced in Junos OS Release 16.1.

Support for protect core in EVPN-VXLAN introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

Configuring BGP PIC Edge for MPLS Layer 3 VPNs

Example: Configuring Provider Edge Link Protection in Layer 3 VPNs

Introduction to Configuring Layer 3 VPNs

LDP Operation

Understanding Multiprotocol BGP

protect-interface

IN THIS SECTION

- Syntax | 13409
- Hierarchy Level | 13409
- Description | 13410
- Options | 13410
- Required Privilege Level | 13410
- Release Information | 13410

Syntax

```
protect-interface interface-name;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name interface],
[edit logical-systems logical-system-name bridge-domains bridge-domains interface],
[edit logical-systems logical-system-name protocols l2circuit neighbor address
interface interface-name],
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name],
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name end-interface],
[edit logical-systems logical-system-name routing-instances name interface],
[edit protocols l2circuit local-switching interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name],
[edit protocols l2circuit local-switching interface interface-name end-interface],
[edit routing-instances routing-instances-name interface]
```

Description

Provide a backup for the protected interface in case of failure. Network traffic uses the primary interface only, as long as the primary interface functions.

Options

interface-name—Name of the protect interface to configure.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the following hierarchy levels introduced in Junos OS Release 17.4: `[edit bridge domains]`, `[edit logical-systems]`, and `[edit routing-instances]`.

RELATED DOCUMENTATION

Configuring the Protect Interface

Configuring EVPN Active-Standby Multihoming to a Single PE Device

protect-maintenance-association (OAM)

IN THIS SECTION

- Syntax | 13411
- Hierarchy Level | 13411
- Description | 13411
- Options | 13411
- Required Privilege Level | 13411

Release Information | 13411

Syntax

```
protect-maintenance-association protect-ma-name;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name ]
```

Description

Configure the name of the protect transport path for the maintenance-association.

Options

protect-ma-name—The name of the protect transport path.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4

RELATED DOCUMENTATION

Configuring Connectivity Fault Management (CFM)

protected-l2circuit

IN THIS SECTION

- Syntax | 13412
- Hierarchy Level | 13412
- Description | 13412
- Options | 13412
- Required Privilege Level | 13413
- Release Information | 13413

Syntax

```
protected-l2circuit {
    egress-pe address;
    ingress-pe address;
    virtual-circuit-id identifier;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name egress-protection],
[edit protocols l2circuit neighbor address interface interface-name egress-protection]
```

Description

Configures the protected Layer 2 circuit as part of an egress protection virtual circuit (EPVC).

Options

egress-pe *address* Specify the address of the egress PE router for the protected Layer 2 circuit.

ingress-pe ***address***	Specify the address of the ingress PE router for the protected Layer 2 circuit.
virtual-circuit-id ***identifier***	Specify the virtual circuit identifier for the protected Layer 2 circuit.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS release 10.4.

RELATED DOCUMENTATION

Example: Configuring an Egress Protection LSP for a Layer 2 Circuit

protection-group

IN THIS SECTION

- Syntax | 13413
- Hierarchy Level | 13415
- Description | 13416
- Required Privilege Level | 13416
- Release Information | 13416

Syntax

```
protection-group {
        ethernet-ring ring-name {
            data-channel {
```

```
                vlan number
            }
            east-interface {
                control-channel channel-name {
                    vlan number;
                    interface name interface-name
                }
            }
            guard-interval number;
            node-id mac-address;
            restore-interval number;
            ring-protection-link-owner;
            non-revertive;
            wait-to-block-interval number;
            major-ring-name name;
            propagate-tc;
            compatibility-version (1|2);
            ring-id number;
            non-vc-mode;
            dot1p-priority number;
            west-interface {
                control-channel channel-name {
                    vlan number;
                    interface name interface-name
                }
                virtual-control-channel {
                    west-interface name;
                    east-interface name;
                }
            }
        }
    control-vlan (vlan-id | vlan-name);
            east-interface {
                node-id mac-address;
                control-channel channel-name {
                    vlan number;
                    interface name interface-name
                }
                interface-none
                ring-protection-link-end;
            }
        }
            control-channel channel-name {
```

```
                vlan number;
                interface name interface-name
            }
        }
        data-channel {
            vlan number
        }
        guard-interval number;
        node-id mac-address;
        restore-interval number;
        ring-protection-link-owner;
            west-interface {
                node-id mac-address;
                control-channel channel-name {
                    vlan number;
                    interface name interface-name
                }
                interface-none
                ring-protection-link-end;
            }
            control-channel channel-name {
                vlan number;
                interface name interface-name
            }
        }
    }
    guard-interval number;
    restore-interval number;
    traceoptions {
        file filename <no-stamp> <world-readable | no-world-readable> <replace> <size size>;
        flag flag;
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure Ethernet ring protection switching.

The statements are explained separately. All statements apply to MX Series routers. EX Series switches do not assign `node-id` and use `control-vlan` instead of `control-channel`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

- *Ethernet Ring Protection Switching Overview*
- Ethernet Ring Protection Using Ring Instances for Load Balancing
- Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers
- Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*

protection (Protocols BGP)

IN THIS SECTION

- Syntax | 13417
- Hierarchy Level | 13417
- Description | 13417
- Required Privilege Level | 13417

Syntax

```
protection;
```

Hierarchy Level

```
[edit routing-instances instance-name protocols bgp family inet unicast],
[edit routing-instances instance-name protocols bgp family inet6 unicast],
[edit routing-instances instance-name protocols bgp family inet labeled-unicast],
[edit routing-instances instance-name protocols bgp family inet6 labeled-unicast]
```

Description

Configure the backup path to protect the active provider edge path in a Layer 3 VPN or a BGP labeled unicast path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Provider Edge Link Protection in Layer 3 VPNs

Example: Configuring Provider Edge Link Protection for BGP Labeled Unicast Paths

protection (Protocols MPLS)

IN THIS SECTION

- Syntax | 13418
- Hierarchy Level | 13418

- Description | 13418
- Required Privilege Level | 13418
- Release Information | 13418

Syntax

```
protection;
```

Hierarchy Level

```
[edit  logical-systems logical-system-name protocols bgp group group-name family inet labeled-unicast],
[edit protocols bgp group group-name family inet labeled-unicast]
```

Description

Configure protection on a link between two routers in different autonomous systems.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding MPLS Inter-AS Link Protection

protection-revert-time

IN THIS SECTION

- Syntax | 13419
- Hierarchy Level | 13419
- Description | 13419
- Options | 13419
- Required Privilege Level | 13420
- Release Information | 13420

Syntax

```
protection-revert-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls interface interface-name static],
[edit protocols mpls interface interface-name static]
```

Description

Specify the amount of time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path.

If you have configured a value of 0 seconds for the `protection-revert-time` statement and traffic is switched to the bypass path, the traffic remains on that path indefinitely. It is never switched back to the original path unless the bypass path is down or you intervene.

Options

seconds—Time in seconds.

- **Range:** 0 through 65,535 seconds

- **Default:** 5 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring Static LSPs

protector-interface

IN THIS SECTION

- Syntax | 13420
- Hierarchy Level | 13421
- Description | 13421
- Options | 13421
- Required Privilege Level | 13421
- Release Information | 13421

Syntax

```
protector-interface interface-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name egress-protection],
[edit protocols l2circuit neighbor address interface interface-name egress-protection]
```

Description

Configures the protector interface for an egress protection LSP.

Options

interface-name —Name of the interface used to protect traffic for an egress protection LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS release 10.4.

RELATED DOCUMENTATION

Example: Configuring an Egress Protection LSP for a Layer 2 Circuit

protector-pe

IN THIS SECTION

- Syntax | 13422
- Hierarchy Level | 13422

- Description | 13422
- Options | 13422
- Required Privilege Level | 13422
- Release Information | 13423

Syntax

```
protector-pe address {
    context-identifier identifier;
    lsp lsp-name;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name egress-protection],
[edit protocols l2circuit neighbor address interface interface-name egress-protection]
```

Description

Configures the protector PE router for an egress protection LSP. Test.

Options

address —IPv4 address for the protector PE router.

context-identifier *identifier* —Identifies the context for the egress protection LSP.

lsp *lsp-name* —Specifies the LSP for the egress protection LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS release 10.4.

RELATED DOCUMENTATION

Example: Configuring an Egress Protection LSP for a Layer 2 Circuit

protocol (Firewall Filter)

IN THIS SECTION

- Syntax | 13423
- Hierarchy Level | 13423
- Description | 13424
- Options | 13424
- Required Privilege Level | 13424
- Release Information | 13424

Syntax

```
protocol (tcp|udp);
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name from ip-version ip-version]
```

Description

Configure the protocol field of IPv4 or the next-header field of the IPv6 address.

Options

tcp Transmission Control Protocol.

udp User Datagram Protocol.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1R1.

RELATED DOCUMENTATION

Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic

protocol

IN THIS SECTION

- Syntax (ACX Series Routers) | **13425**
- Syntax (MX Series Routers) | **13425**
- Syntax (QFX Series Switches) | **13425**
- Syntax (EX Series Switches) | **13425**
- Hierarchy Level | **13426**
- Description | **13426**
- Options | **13426**

Required Privilege Level | 13427
Release Information | 13427

Syntax (ACX Series Routers)

```
protocol (cdp | elmi | ieee8021x | ieee8023ah | lacp | lldp | mmrp | mvrp | stp | vtp);
```

Syntax (MX Series Routers)

```
protocol (cdp | pvstp | stp | vtp);
```

Syntax (QFX Series Switches)

```
protocol (cdp | elmi | gvrp | ieee8021x | ieee8023ah | lacp | lldp | mmrp | mvrp | stp | udld |
vstp | vtp);
```

Syntax (EX Series Switches)

Syntax (EX2300 Multigigabit Model Switches)

```
protocol (cdp | gvrp | ieee8023ah | lacp | lldp | mvrp | stp | vstp | vtp);
```

Syntax (EX2300, EX3400, EX4300, EX4300 Multigigabit Model, EX4600, and EX4650 Switches)

```
protocol (cdp | elmi | gvrp | ieee8021x | ieee8023ah | lacp | lldp | mmrp | mvrp | stp | udld |
vstp | vtp);
```

Syntax (EX4100, EX4300-MP, and EX4400 Switches)

```
protocol (all | cdp | elmi | gvrp | ieee8021x | ieee8023ah | lacp | lldp | mmrp | mvrp | stp |
udld | vstp | vtp);
```

Syntax (EX9200 Switches)

```
protocol (cdp | elmi | gvrp | ieee8021x | ieee8023ah | lacp | lldp | mmrp | mvrp | pvstp | stp |
udld | vtp);
```

Hierarchy Level

```
[edit logical-systems name protocols layer2-control mac-rewrite              interface interface-
name],
[edit protocols layer2-control mac-rewrite interface interface-name]
```

Description

Configure the protocol to be tunneled on an interface using Layer 2 protocol tunneling (L2PT). To enable tunneling multiple protocols, include multiple `protocol` statements.

You can tunnel different protocols listed in the Options section on different types of devices. The Syntax and Release Information sections list the available options for the protocols that different devices can tunnel as of a particular Junos OS release (for devices that support L2PT).

When a service provider edge (PE) port configured for L2PT receives a control packet for a supported protocol, the device rewrites the multicast destination MAC address with the predefined multicast tunneling MAC address 01:00:0c:cd:cd:d0. The packet travels across the provider network transparently to the other end of the tunnel, and the destination device restores the original multicast destination MAC address to deliver the packet at its destination.

Options

`all`—Tunnel all supported Layer 2 control protocols.

`cdp`—Tunnel the Cisco Discovery Protocol (CDP).

`elmi`—Tunnel Ethernet Local Management Interface (E-LMI) packets.

`gvrp`—Tunnel Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP) packets.

`ieee8021x`—Tunnel IEEE 802.1X authentication packets.

`ieee8023ah`—Tunnel IEEE 802.3AH Operation, Administration, and Maintenance (OAM) link fault management (LFM) packets.

`lacp`—Tunnel Link Aggregation Control Protocol (LACP) packets.

`lldp`—Tunnel Link Layer Discovery Protocol (LLDP) packets.

`mmrp`—Tunnel Multiple MAC Registration Protocol (MMRP) packets.

`mvrp`—Tunnel Multiple VLAN Registration Protocol (MVRP) packets.

`pvstp`—Tunnel VLAN Spanning Tree Protocol (VSTP), Per-VLAN Spanning Tree (PVST), and Per-VLAN Spanning Tree Plus (PVST+) Protocol packets.

`stp`—Tunnel packets for all versions of Spanning-Tree Protocols.

`udld`—Tunnel Unidirectional Link Detection (UDLD) packets.

`vstp`—Tunnel VLAN Spanning Tree Protocol (VSTP) packets.

`vtp`—Tunnel VLAN Trunking Protocol (VTP) packets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support for PVST/PVST+ introduced in Junos OS Release 13.3 for MX Series routers and EX9200 switches.

Support for E-LMI, IEEE 802.1X, MMRP, and UDLD introduced in Junos OS Release 17.3R1 for EX4300 switches.

Support for E-LMI, GVRP, IEEE 802.1x, IEEE 802.3AH, LACP, LLDP, MMRP, MVRP, and UDLD introduced in Junos OS Release 17.3R1 for EX9200 switches.

Support for E-LMI, IEEE 802.1X, MMRP, and UDLD introduced in Junos OS Release 18.2R1 for EX2300 and EX3400 switches.

`all` option added in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Understanding Layer 2 Protocol Tunneling

Configuring Layer 2 Protocol Tunneling

protocol (early-aging)

IN THIS SECTION

- Syntax | 13428
- Hierarchy Level | 13428
- Description | 13429
- Options | 13429
- Required Privilege Level | 13429
- Release Information | 13429

Syntax

```
protocol name {
 cutoff-time {
  (percent percent | seconds seconds);
 }
}
```

Hierarchy Level

```
[edit security flow aging early-aging]
```

Description

Set of attributes that define early-ageout by-protocol.

Options

name Match IP protocol type

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

This statement is introduced in Junos OS Release 23.4R1.

protocol (Application Identification)

IN THIS SECTION

- Syntax | 13429
- Hierarchy Level | 13430
- Description | 13430
- Options | 13430
- Required Privilege Level | 13430
- Release Information | 13430

Syntax

```
protocol (http | ssl | tcp | udp);
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type
signature l4-l7-signature-name]
```

Description

Identify the protocol bundles to be monitored to classify applications. This statement is not available if the MX Series router is running Next Gen Services.

Options

http	Use the HTTP protocol .
ssl	Use the SSL protocol.
tcp	Use the TCP protocol.
udp	Use the UDP protocol.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

protocol (Applications)

IN THIS SECTION

- Syntax | 13431
- Hierarchy Level | 13431
- Description | 13431
- Options | 13431
- Required Privilege Level | 13432
- Release Information | 13432

Syntax

```
protocol protocol-name;
```

Hierarchy Level

```
[edit applications application  application-name  ]
[edit applications application  application-name  term  term-name]
```

Description

Specify the networking protocol name or number.

The networking protocol type combines with **destination-port** to identify an application type. For example, to create an application for iSCSI, use the protocol **tcp** with the destination port number **3260**.

Options

protocol-name— Networking protocol name. The following text values are supported. For a complete list of possible numeric values, see RFC 1700, *Assigned Numbers (for the Internet Protocol Suite)*.

- ah—IP Security Authentication Header
- egp—Exterior gateway protocol

- `esp`—IPsec Encapsulating Security Payload
- `gre`—Generic routing encapsulation
- `icmp`—Internet Control Message Protocol
- `igmp`—Internet Group Management Protocol
- `ipip`—IP over IP
- `node`—Clear each session that uses the specified IP protocol on a specific node.
- `ospf`—Open Shortest Path First
- `pim`—Protocol Independent Multicast
- `rsvp`—Resource Reservation Protocol
- `sctp`—Stream Control Transmission Protocol
- `tcp`—Transmission Control Protocol
- `udp`—User Datagram Protocol

> **NOTE**: Internet Protocol version 6 (IPv6) is not supported as a network protocol in application definitions.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policy Application Sets Overview

protocol-binding

IN THIS SECTION

- Syntax | 13433
- Hierarchy Level | 13434
- Description | 13434
- Options | 13434
- Required Privilege Level | 13434
- Release Information | 13434

Syntax

```
protocol-binding {
    application application-name;
    icmp;
    icmpv6;
    ip (Security IDP Custom Attack) {
        protocol-number transport-layer-protocol-number;
    }
    ipv6 {
        protocol-number transport-layer-protocol-number;
    }
    rpc {
        program-number rpc-program-number;
    }
    tcp (Security IDP Signature Attack) {
        minimum-port port-number <maximum-port port-number>;
    }
    udp (Security IDP Signature Attack) {
        minimum-port port-number <maximum-port port-number>;
    }
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain],

[edit security idp custom-attack attack-name attack-type signature]
```

Description

Select a protocol that the attack uses to enter your network.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

protocol-command

IN THIS SECTION

- Syntax | 13435
- Hierarchy Level | 13435
- Description | 13435
- Options | 13435
- Required Privilege Level | 13435
- Release Information | 13435

Syntax

```
protocol-command object-name {
    value [value];
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm custom-objects]
```

Description

Different protocols use different commands to communicate between servers and clients. By blocking or allowing certain commands, traffic can be controlled on the protocol command level.

Options

- *object-name*—Name of the command-list object.
- value *value*—Value of the command-list object. You can configure multiple values separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The [edit security utm default-configuration] hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

> *Unified Threat Management Overview*

protocol-down

IN THIS SECTION

- Syntax | 13436
- Hierarchy Level | 13436
- Description | 13436
- Required Privilege Level | 13436
- Release Information | 13436

Syntax

```
protocol-down;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management event]
```

Description

Configure the upper layer indication of a protocol down event. When the protocol-down statement is included, the protocol down event triggers the action specified under the `action` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 9.4 .

RELATED DOCUMENTATION

Example: Configure Ethernet OAM Link Fault Management

Configure Ethernet OAM Link Fault Management

protocol (Drop Profile Map)

IN THIS SECTION

- Syntax | 13437
- Hierarchy Level | 13437
- Description | 13437
- Options | 13437
- Required Privilege Level | 13438
- Release Information | 13438

Syntax

```
protocol protocol drop-profile profile-name;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name drop-profile-map loss-priority (low | medium-high |
high)]
```

Description

Configure the protocol type for the specified weighted random early detection (WRED) drop profile.

Options

protocol—Type of protocol. The protocol can be:

- `any`—Accept any protocol type.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

protocol (Drop Profiles)

IN THIS SECTION

- Syntax | 13438
- Hierarchy Level | 13438
- Description | 13439
- Options | 13439
- Required Privilege Level | 13439
- Release Information | 13439

Syntax

```
protocol protocol drop-profile profile-name;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Specify the protocol type for the specified drop profile.

Options

drop-profile *profile-name* —Name of the drop profile.

protocol —Type of protocol. It can be:

- **any**—Accept any protocol type.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches

Configuring CoS Tail Drop Profiles (CLI Procedure)

Understanding CoS Tail Drop Profiles

protocol (Dynamic Schedulers)

IN THIS SECTION

- Syntax | 13440
- Hierarchy Level | 13440
- Description | 13440
- Options | 13440
- Required Privilege Level | 13440

Release Information | 13440

Syntax

```
protocol (any | non-tcp | tcp);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name drop-profile-map]
```

Description

Specify the protocol type for the specified scheduler in a dynamic profile.

Options

`any`—Accept any protocol type.

`non-tcp`—Accept any protocol type other than TCP/IP.

`tcp`—Accept only TCP/IP protocol.

NOTE: Protocol types `non-tcp` and `tcp` are not supported on MX Series routers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Schedulers in a Dynamic Profile for Subscriber Access

protocol (Flow Descriptions)

IN THIS SECTION

- Syntax | 13441
- Hierarchy Level | 13441
- Description | 13441
- Default | 13442
- Options | 13442
- Required Privilege Level | 13442
- Release Information | 13442

Syntax

```
protocol number;
```

Hierarchy Level

```
[edit unified-edge pcef flow-description flow-identifier],
[edit services pcef flow-description flow-identifier]
```

Description

Specify a protocol type to identify the subscriber traffic that you want the service data flow (SDF) filter to detect. If you specify the `protocol` statement, you must specify a protocol number.

If you are using Junos OS Subscriber Aware, specify the protocol type at the `[edit unified-edge pcef flow-description flow-identifier]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the protocol type at the `[edit services pcef flow-description flow-identifier]` hierarchy level.

Default

If you do not specify the `protocol` statement, the default is any protocol.

Options

number Number that specifies the IP protocol type.

- **Range:** 1 through 255

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef flow-description flow-identifier]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

Configuring Service Data Flow Filters

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Understanding Application-Aware Policy Control for Subscriber Management

protocol (Host Outbound Traffic)

IN THIS SECTION

- Syntax | 13443
- Hierarchy Level | 13443
- Description | 13443
- Required Privilege Level | 13443
- Release Information | 13444

Syntax

```
protocol {
    isis-over-gre {
        dscp-code-point dscp-code-point;
    }
}
```

Hierarchy Level

```
[edit class-of-service host-outbound-traffic],
[edit dynamic-profiles name class-of-service host-outbound-traffic]
```

Description

Set the DSCP code point for IS-IS traffic originating on the host and being sent over a GRE tunnel. By default, this DSCP code point is `000000`, which is normally mapped to `best-effort` forwarding class and can lead to packet drops during congestion. Set this DSCP code point to enable traffic prioritization for IS-IS traffic originating on a host and being sent over a GRE tunnel.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 17.3.

protocol (IPsec)

IN THIS SECTION

- Syntax | 13444
- Hierarchy Level | 13444
- Description | 13444
- Options | 13444
- Required Privilege Level | 13445
- Release Information | 13445

Syntax

```
protocol (ah | esp | bundle);
```

Hierarchy Level

```
[edit services ipsec-vpn ipsec proposal proposal-name],
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]
```

Description

Define an IPsec protocol for a dynamic or manual SA.

Options

ah—(Not supported on MS-MPCs and MS-MICs on MX Series routers) Use the Authentication Header protocol.

`esp`—Use the Encapsulating Security Payload protocol.

`bundle`—(Not supported on MS-MPCs and MS-MICs on MX Series routers) Use the AH and ESP protocol.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Security Associations

protocol (Junos-FIPS Software)

IN THIS SECTION

- Syntax | 13445
- Hierarchy Level | 13446
- Description | 13446
- Options | 13446
- Required Privilege Level | 13446
- Release Information | 13446

Syntax

```
protocol esp;
```

Hierarchy Level

```
[edit security ipsec internal security-association manual direction]
```

Description

The protocol used for the internal Routing-Engine–to–Routing-Engine IPsec security association (SA) configuration.

Options

Only `esp` is supported.

Required Privilege Level

Crypto Officer—To add and view this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring IPsec for Enabling Internal Communications Between Routing Engines for Junos OS in FIPS Mode*

protocol (Junos OS)

IN THIS SECTION

- Syntax | 13447
- Hierarchy Level | 13447
- Description | 13447
- Options | 13447

Required Privilege Level | 13448
Release Information | 13448

Syntax

```
protocol (ah | esp | bundle);
```

Hierarchy Level

```
[edit security ipsec proposal ipsec-proposal-name],
[edit security ipsec security-association sa-name manual direction (inbound | outbound |
bidirectional)]
```

Description

Define the IPsec protocol for a manual or dynamic SA.

> **NOTE**: The Junos OS supports only encapsulating security payload (ESP) when you use tunnel mode.
>
> In transport mode, the Junos OS does not support authentication header (AH) and ESP header bundles.
>
> In transport mode, the Junos OS supports only Border Gateway Protocol (BGP).

Options

ah—Authentication Header protocol

bundle—AH and ESP protocols

esp—ESP protocol (the **tunnel** statement must be included at the **[edit security ipsec security-association *sa-name* mode** hierarchy level)

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Example: Using IPsec to Protect BGP Traffic
- *Configuring Manual IPsec Security Associations for an ES PIC*
- *Configuring the Protocol for a Dynamic IPsec SA*

protocol-master

IN THIS SECTION

- Syntax | 13448
- Hierarchy Level | 13449
- Description | 13450
- Options | 13450
- Required Privilege Level | 13450
- Release Information | 13450

Syntax

```
protocol-master (inet | inet6);
```

Hierarchy Level

[edit bridge-domains *name* forwarding-options dhcp-relay dhcpv6 group *group-name* dual-stack-group *dual-stack-group-name*],
[edit bridge-domains *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* bridge-domains *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* bridge-domains *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* forwarding-options dhcp-relay dual-stack-group],
[edit logical-systems *name* routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* system services dhcp-local-server dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* vlans *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* routing-instances *name* vlans *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* system services dhcp-local-server dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* vlans *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit logical-systems *name* vlans *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit routing-instances *name* bridge-domains *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],
[edit routing-instances *name* forwarding-options dhcp-relay dhcpv6 group *name* dual-stack-group *dual-stack-group-name*],
[edit routing-instances *name* forwarding-options dhcp-relay dual-stack-group *dual-stack-group-name*],

```
[edit routing-instances name system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit routing-instances name vlans name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name],
[edit routing-instances name vlans name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit vlans name forwarding-options dhcp-relay dhcpv6 group name dual-stack-group dual-stack-group-name],
[edit vlans name forwarding-options dhcp-relay dual-stack-group dual-stack-group-name]
```

Description

Select family as protocol primary. In some customer use cases, CPE devices have reachability to multiple BNG routers for load sharing. The CPE's DHCP client broadcasts a DHCP protocol handshake to the reachable routers; more than one router may respond. In a DHCP dual-stack environment, the DHCPv4 and DHCPv6 protocol handshakes are independent of each other, meaning that each arm of the subscriber session could connect to a different router. You can avoid this situation by specify the primary protocol. For a given dual-stack subscriber, this configuration causes the rejection of any binding attempt from the secondary address family client when a binding is not currently active for the primary protocol family. If bindings are currently active for both arms when the primary protocol family binding is released or deleted, then the binding for the secondary address family is also torn down.

Options

inet INET family has protocol primary behavior.

inet6 INET6 family has protocol primary behavior.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

Configuring Access Profile Options for Interactions with RADIUS Servers

Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution

Configuring Access Components for the PPPoE Wholesale Network Solution

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

protocol-name

IN THIS SECTION

- Syntax | 13451
- Hierarchy Level | 13451
- Description | 13452
- Options | 13452
- Required Privilege Level | 13452
- Release Information | 13452

Syntax

```
protocol-name protocol-name {
    tunable-name  tunable-name {
        tunable-value protocol-value;
    }
}
```

Hierarchy Level

```
[edit security idp sensor-configuration detector]
```

Description

Specify the name of the protocol to be used to configure each of the protocol detector engines.

Options

protocol-name—Name of the specific protocol.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. Support.

protocol-origin

IN THIS SECTION

- Syntax | 13452
- Hierarchy Level | 13453
- Description | 13453
- Options | 13453
- Required Privilege Level | 13453

Syntax

```
protocol-origin {
                          bgp;
                          cli;
```

```
                                pcep;
                                   }
```

Hierarchy Level

```
[edit policy-options policy-statement name from traffic-engineering spring-te]
```

Description

Specify protocol originating the SPRING-TE policy candidate path.

Options

bgp SPRING-TE policy candidate path originated by BGP

cli SPRING-TE policy candidate path originated by cli

pcep SPRING-TE policy candidate path originated by PCEP

Required Privilege Level

routing

RELATED DOCUMENTATION

Link-State Distribution using BGP

protocol (OVSDB)

IN THIS SECTION

- Syntax | 13454
- Hierarchy Level | 13454
- Description | 13454
- Options | 13455
- Required Privilege Level | 13455
- Release Information | 13455

Syntax

```
protocol protocol {
    port number;
}
```

Hierarchy Level

```
[edit protocols ovsdb controller]
```

Description

Configure the security protocol that protects the connection between a Juniper Networks device that supports the *Open vSwitch Database (OVSDB)* management protocol and a software-defined networking (SDN) controller.

The *Secure Sockets Layer (SSL)* connection requires a private key and certificates, which must be stored in the **/var/db/certs** directory of the Juniper Networks device. See *Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with SDN Controllers*.

Options

protocol Establish a secure connection to the SDN controller, using SSL or TCP.

> **NOTE**: SSL is the only supported connection protocol.

- **Default:** `ssl`

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Understanding How to Set Up OVSDB Connections on a Juniper Networks Device

protocol (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 13456
- Hierarchy Level | 13456
- Description | 13456
- Options | 13456
- Required Privilege Level | 13456
- Release Information | 13456

Syntax

```
protocol tcp {
    port port;
}
```

Hierarchy Level

```
[edit protocols openflow switch switch-name controller]
```

Description

Specify the connection protocol that the OpenFlow virtual switch uses to connect to the OpenFlow controller.

Options

tcp Establish a TCP connection to the controller.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Support for OpenFlow on Devices Running Junos OS

controller (Protocols OpenFlow) | **4654**

port (Protocols OpenFlow) | **12857**

protocol (Rewrite Rules)

IN THIS SECTION

- Syntax | 13457
- Hierarchy Level | 13457
- Description | 13457
- Options | 13458
- Required Privilege Level | 13458
- Release Information | 13458

Syntax

```
protocol protocol-types;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules exp
rewrite-name],
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules dscp
rewrite-name],
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules dscp-
ipv6 rewrite-name],
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules inet-
precedence rewrite-name]
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules inet6-
precedence rewrite-name]
```

Description

Apply a rewrite rule to MPLS packets only, and write the CoS value to MPLS headers only; or apply a rewrite rule to MPLS and IPv4 or IPv6 packets, and write the CoS value to MPLS and IPv4 or IPv6 headers.

Options

protocol-types can be one of the following:

- `mpls-any`—Apply a rewrite rule to MPLS packets and write the CoS value to MPLS headers only.
- `mpls-inet-both`—Apply a rewrite rule to VPN MPLS packets with IPv4 payloads. Write the CoS value to the MPLS and IPv4 headers.
- `mpls-inet-both-non-vpn`—Apply a rewrite rule to non-VPN MPLS packets with IPv4 payloads. Write the CoS value to the MPLS and IPv4 headers.

> Junos OS Evolved supports only the `mpls-any` option. Devices running Junos OS Evolved do not support different rewrite rules for VPN and non-VPN traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Option for `dscp` and `inet-prec` introduced in Junos OS Release 8.4.

Option for `dscp-ipv6` introduced in Junos OS Release 10.4R2.

RELATED DOCUMENTATION

> *Rewriting MPLS and IPv4 Packet Headers*

protocol (Schedulers)

IN THIS SECTION

- Syntax | 13459
- Hierarchy Level | 13459
- Description | 13459
- Options | 13459
- Required Privilege Level | 13460
- Release Information | 13460

Syntax

```
protocol (any | non-tcp | tcp);
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name drop-profile-map]
```

Description

Specify the protocol type for the specified scheduler.

Options

`any`—Accept any protocol type.

`non-tcp`—Accept any protocol type other than TCP/IP.

> **NOTE**: On ACX Series Routers, when you configure the `non-tcp` option, only the `any` option is supported for "loss-priority" on page 10169.

`tcp`—Accept TCP/IP protocol type.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Schedulers

protocol (Security Destination NAT)

IN THIS SECTION

- Syntax | 13460
- Hierarchy Level | 13460
- Description | 13461
- Options | 13461
- Required Privilege Level | 13461
- Release Information | 13461

Syntax

```
protocol [protocol-name-or-number];
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name rule rule-name match]
```

Description

Specify an IP protocol to match the rule. You can configure multiple protocol names or protocol numbers.

Options

protocol-name-or-number—Name or number of the specific protocol.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

protocol (Security Forwarding Options)

IN THIS SECTION

- Syntax | 13461
- Hierarchy Level | 13462
- Description | 13462
- Options | 13462
- Required Privilege Level | 13463
- Release Information | 13463

Syntax

```
protocol protocol;
```

Hierarchy Level

```
[edit security forwarding-options mirror-filter filter-name]
```

Description

Specify the networking protocol name or number to be matched for mirroring.

Options

protocol-name

Networking protocol name or number. The following text values are supported. For a complete list of possible numeric values, see RFC 1700, *Assigned Numbers (for the Internet Protocol Suite)*.

ah	IP Security Authentication header
egp	Exterior gateway protocol
esp	IPsec Encapsulating Security Payload
gre	Generic routing encapsulation
icmp	Internet Control Message Protocol
icmp6	Internet Control Message Protocol version 6
igmp	Internet Group Management Protocol
ipip	IP over IP
ospf	Open Shortest Path First
pim	Protocol Independent Multicast
rsvp	Resource Reservation Protocol
sctp	Stream Control Transmission Protocol
tcp	Transmission Control Protocol
udp	User Datagram Protocol

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

mirror-filter (Security Forwarding Options) | 10768

show security forwarding-options mirror-filter | 24363

protocol (Security IDP Signature Attack)

IN THIS SECTION

- Syntax | 13463
- Hierarchy Level | 13469
- Description | 13469
- Options | 13469
- Required Privilege Level | 13469
- Release Information | 13469

Syntax

```
protocol {
    icmp (Security IDP Signature Attack) {
        checksum-validate {
            match (equal | greater-than | less-than | not-equal);
            value checksum-value;
        }
```

```
        code {
            match (equal | greater-than | less-than | not-equal);
            value code-value;
        }
        data-length {
            match (equal | greater-than | less-than | not-equal);
            value data-length;
        }
        identification (Security ICMP Headers) {
            match (equal | greater-than | less-than | not-equal);
            value identification-value;
        }
        sequence-number (Security IDP ICMP Headers) {
            match (equal | greater-than | less-than | not-equal);
            value sequence-number;
        }
        type (Security IDP ICMP Headers) {
            match (equal | greater-than | less-than | not-equal);
            value type-value;
        }
    }
    icmpv6 {
        checksum-validate {
            match (equal | greater-than | less-than | not-equal);
            value checksum-value;
        }
        code {
            match (equal | greater-than | less-than | not-equal);
            value code-value;
        }
        data-length {
            match (equal | greater-than | less-than | not-equal);
            value data-length;
        }
        identification {
            match (equal | greater-than | less-than | not-equal);
            value identification-value;
        }
        sequence-number {
            match (equal | greater-than | less-than | not-equal);
            value sequence-number;
        }
        type {
```

```
            match (equal | greater-than | less-than | not-equal);
            value type-value;
        }
    }
    ipv4 {
        checksum-validate {
            match (equal | greater-than | less-than | not-equal);
            value checksum-value;
        }
        destination {
            match (equal | greater-than | less-than | not-equal);
            value ip-address-or-hostname;
        }
        identification {
            match (equal | greater-than | less-than | not-equal);
            value identification-value;
        }
        ihl {
            match (equal | greater-than | less-than | not-equal);
            value ihl-value;
        }
        ip-flags {
            (df | no-df);
            (mf | no-mf);
            (rb | no-rb);
        }
        protocol {
            match (equal | greater-than | less-than | not-equal);
            value transport-layer-protocol-id;
        }
        source {
            match (equal | greater-than | less-than | not-equal);
            value ip-address-or-hostname;
        }
        tos {
            match (equal | greater-than | less-than | not-equal);
            value type-of-service-in-decimal;
        }
        total-length {
            match (equal | greater-than | less-than | not-equal);
            value total-length-of-ip-datagram;
        }
        ttl {
```

```
            match (equal | greater-than | less-than | not-equal);
            value time-to-live;
        }
    }
    ipv6 {
        destination {
            match (equal | greater-than | less-than | not-equal);
            value ip-address-or-hostname;
        }
        extension-header {
            destination-option {
                home-address {
                    match (equal | greater-than | less-than | not-equal);
                    value header-value;
                }
                    option-type {
                        match (equal | greater-than | less-than | not-equal);
                        value header-value;
                    }
            }
            routing-header {
                header-type {
                    match (equal | greater-than | less-than | not-equal);
                    value header-value;
                    }
                }
        }
        flow-label {
            match (equal | greater-than | less-than | not-equal);
            value flow-label-value;
        }
        hop-limit {
            match (equal | greater-than | less-than | not-equal);
            value hop-limit-value;
        }
        next-header {
            match (equal | greater-than | less-than | not-equal);
            value next-header-value;
        }
        payload-length {
            match (equal | greater-than | less-than | not-equal);
            value payload-length-value;
        }
```

```
        source {
            match (equal | greater-than | less-than | not-equal);
            value ip-address-or-hostname;
        }
        traffic-class {
            match (equal | greater-than | less-than | not-equal);
            value traffic-class-value;
        }
    tcp (Security IDP Signature Attack) {
        ack-number {
            match (equal | greater-than | less-than | not-equal);
            value acknowledgement-number;
        }
        checksum-validate {
            match (equal | greater-than | less-than | not-equal);
            value checksum-value;
        }
        data-length {
            match (equal | greater-than | less-than | not-equal);
            value tcp-data-length;
        }
        destination-port {
            match (equal | greater-than | less-than | not-equal);
            value destination-port;
        }
        header-length {
            match (equal | greater-than | less-than | not-equal);
            value header-length;
        }
        mss {
            match (equal | greater-than | less-than | not-equal);
            value maximum-segment-size;
        }
        option {
            match (equal | greater-than | less-than | not-equal);
            value tcp-option;
        }
        reserved {
            match (equal | greater-than | less-than | not-equal);
            value reserved-value;
        }
        sequence-number {
            match (equal | greater-than | less-than | not-equal);
```

```
            value sequence-number;
        }
        source-port {
            match (equal | greater-than | less-than | not-equal);
            value source-port;
        }
        tcp-flags {
            (ack | no-ack);
            (fin | no-fin);
            (psh | no-psh);
            (r1 | no-r1);
            (r2 | no-r2);
            (rst | no-rst);
            (syn | no-syn);
            (urg | no-urg);
        }
        urgent-pointer {
            match (equal | greater-than | less-than | not-equal);
            value urgent-pointer;
        }
        window-scale {
            match (equal | greater-than | less-than | not-equal);
            value window-scale-factor;
        }
        window-size {
            match (equal | greater-than | less-than | not-equal);
            value window-size;
        }
    }
    udp (Security IDP Signature Attack) {
        checksum-validate {
            match (equal | greater-than | less-than | not-equal);
            value checksum-value;
        }
        data-length {
            match (equal | greater-than | less-than | not-equal);
            value data-length;
        }
        destination-port {
            match (equal | greater-than | less-than | not-equal);
            value destination-port;
        }
        source-port {
```

```
            match (equal | greater-than | less-than | not-equal);
            value source-port;
        }
    }
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature]
```

Description

Specify a protocol to match the header information for the signature attack.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3. Statement modified in Junos OS Release 12.3X48-D25 to add ICMPv6 protocol support.

protocol (Security Source NAT)

IN THIS SECTION

- Syntax | 13470
- Hierarchy Level | 13470

- Description | 13470
- Options | 13470
- Required Privilege Level | 13470
- Release Information | 13470

Syntax

```
protocol [protocol-name-or-number];
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name match]
```

Description

Specify an IP protocol to match the rule. You can configure multiple protocol names or protocol numbers.

Options

protocol-name-or-number—Name or number of the specific protocol.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

protocol (Server MEP)

IN THIS SECTION

- Syntax | 13471
- Hierarchy Level | 13471
- Description | 13471
- Options | 13471
- Required Privilege Level | 13472
- Release Information | 13472

Syntax

```
protocol (l2circuit | l2vpn | ethernet) {
    interface interface-name;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

Description

Configure the protocol as Layer 2 circuit, Layer 2 VPN, or Ethernet and associate the interface to the protocol that needs to be monitored for ETH-AIS.

Options

l2circuit	Configure the protocol for the server MEP as Layer 2 circuit.
l2vpn	Configure the protocol for the server MEP as Layer 2 VPN.
ethernet	Configure the protocol for the server MEP as Ethernet.

interface *interface-name* — The interface that is to be associated with the protocol that needs monitoring for ETH-AIS.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Alarm Indication

Ethernet Alarm Indication

protocol (Services Proxy)

IN THIS SECTION

- Syntax | 13472
- Hierarchy Level | 13473
- Description | 13473
- Options | 13473
- Required Privilege Level | 13473
- Release Information | 13473

Syntax

```
protocol {
    http {
```

```
        host host;
        port port;
    }
}
```

Hierarchy Level

```
[edit services proxy]
```

Description

Define the proxy profile settings for application signature package download.

You can download the application signature package hosted on an external server, using a proxy sever. To use the proxy server for downloading, you must configure a profile with host and port details of the proxy server, and apply the proxy profile in the `set services application-identification download` command.

This configuration enables you to download the signature package when you have already deployed a web proxy on your device as part of your overall security solution.

Options

name Proxy profile name.

protocol Protocol type for the profile. Support is available for only HTTP connections.

host IP address of the proxy server.

port Port number used by the proxy server.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Predefined Application Signatures for Application Identification

protocol (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 13474
- Hierarchy Level | 13474
- Description | 13474
- Options | 13474
- Required Privilege Level | 13475
- Release Information | 13475

Syntax

```
protocol  protocol;
```

Hierarchy Level

```
[edit services radius-flow-tap policy policy-name inet drop-policy rule-name from],
[edit services radius-flow-tap policy policy-name inet6 drop-policy rule-name from]
```

Description

Specify the match IP protocol type for the radius-flow-tap policy.

Options

protocol—Protocol for the IPv4 or IPv6 address for the radius-flow-tap policy.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Subscriber Secure Policy Overview

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

protocol-version

IN THIS SECTION

- Syntax | 13475
- Hierarchy Level | 13476
- Description | 13476
- Options | 13476
- Required Privilege Level | 13476
- Release Information | 13476

Syntax

```
protocol-version (all | tls1  | tls11 | tls12 | tls12-and-lower | tls13);
```

Hierarchy Level

```
[edit services ssl termination profile profile-name]
[edit services ssl initiation profile profile-name]
```

Description

Specify the accepted SSL protocol version.

You can specify the SSL/TLS protocol version the SRX Series Firewall uses to negotiate in SSL connections.

Options

- `all`—Accept all versions of TLS. This is enabled by default.
- `TLS version 1.0`—Accept TLS version 1.0. It provides secure communication over networks by providing privacy and data integrity between communicating applications
- `TLS version 1.1`—Accept TLS version 1.1. This enhanced version of TLS provides protection against cipher-block chaining (CBC) attacks.
- `TLS version 1.2` —Accept TLS version 1.2. This enhanced version of TLS provides improved flexibility for negotiation of cryptographic algorithms.
- `TLS version 1.2 and lower` —Accept TLS version 1.2 and lower.
- `TLS version 1.3` —Accept TLS version 1.3. This enhanced version of TLS provides improved security and better performance.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. The `tls11` and `tls12` options are introduced in 15.1X49-D30. Starting in Junos OS Release 21.2R1, on SRX Series Firewalls, SSL proxy supports TLS version 1.3.

RELATED DOCUMENTATION

Firewall User Authentication Overview

SSL Proxy Overview

protocols

IN THIS SECTION

- Syntax | 13477
- Hierarchy Level | 13747
- Description | 13747
- Required Privilege Level | 13747
- Release Information | 13748

Syntax

```
protocols {
 all-fiber-channel-enode {
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
   flow-detect-time seconds;
   flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
```

```
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  amtv4 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  amtv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
ancp {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
  }
  flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
  }
  flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ancpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
```

```
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



   }
  }
  arp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;



    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
```

```
 recover-time seconds;




}
arp {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
```

```
 timeout-active-flows;


}
bcast {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
    }
    ucast {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    unclassified {
     bandwidth packets per second;
```

```
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  atm {
   aggregate {
    bandwidth packets per second;
    burst packets;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }


  bfd {
   aggregate {
    bandwidth packets per second;
    burst packets;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   bundle-bfd {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   multihop-bfd {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    unclassified {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
```

```
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  bfdv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
```

```
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
bgp {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
```

```
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  bgpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  custom {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
demux-autosense {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
  }
  flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
  }
  flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  dhcpv4 {
   ack {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
```

```
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
aggregate {
 bandwidth packets per second;
 burst packets;
 disable-fpc;
 disable-logging;
 disable-routing-engine;


 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
```

```
 recover-time seconds;




}
bad-packets {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
    }
    bootp {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;



    }
    decline {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   discover {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   force-renew {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   inform {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   lease-active {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   lease-query {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
```

```
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    lease-unassigned {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
```

```
 timeout-active-flows;


}
lease-unknown {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
nak {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   no-message-type {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   offer {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
```

```
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    rebind {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
```

```
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    release {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
     }
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    renew {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   request {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
```

```
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
unclassified {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
 }
}
dhcpv4v6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
```

```
}
dhcpv6 {
 advertise {
  bandwidth packets per second;
  burst packets;
  bypass-aggregate;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
 aggregate {
```

```
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;

    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }

    recover-time seconds;

   }
   confirm {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
```

```
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    decline {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   information-request {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
```

```
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    leasequery {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
```

```
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    leasequery-data {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   leasequery-done {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   leasequery-reply {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    rebind {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
```

```
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
reconfigure {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
```

```
 recover-time seconds;
 timeout-active-flows;


}
relay-forward {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
}
relay-reply {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
release {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   renew {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   reply {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   request {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    solicit {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
```

```
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  diameter {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  dns {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
```

```
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
dtcp {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
```

```
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  dynamic-vlan {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  egpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  eoam {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
```

```
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    oam-cfm {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
```

```
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  esmc {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
```

```
    timeout-active-flows;



   }
  }
  exceptions {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }


    recover-time seconds;
```

```
   }
   mcast-rpf-err {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   mtu-exceeded {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  fab-probe {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  filter-action {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;




}
filter-v4 {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   filter-v6 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   other {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
 }
 fip-snooping {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  firewall-host {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
```

```
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  frame-relay {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;

    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }

    recover-time seconds;

   }
   frf15 {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   frf16 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ftp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ftpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
```

```
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  garp-reply {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  gre {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;





}
hag {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    hbc {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   punt {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



  }
}
gtp-path-mgmt {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
```

```
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
  }

  recover-time seconds;
```

```
 }
}
icmp {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;
```

```
 }
}
icmpv6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;
```

```
 }
}
igmp {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
 query {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   report {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  igmpv4v6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  igmpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
```

```
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   mld {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
```

```
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    mld-query {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   mld-report {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
```

```
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    unclassified {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
```

```
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
inline-ka {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
  }
  flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
  }
  flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
```

```
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   inline-svcs {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



   }
  }
  ip-fragments {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
```

```
 recover-time seconds;




}
first-fragment {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
```

```
 recover-time seconds;
 timeout-active-flows;


}
trail-fragment {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
 }
}
ipmcast-miss {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
```

```
}
ipmc-reserved {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
}
ip-options {
```

```
aggregate {
 bandwidth packets per second;
 burst packets;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }
 recover-time seconds;
}
non-v4v6 {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
```

```
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    router-alert {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
```

```
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ipsec {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;




}
ike {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    mtu-error {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
```

```
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    on-trigger {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
```

```
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
sn-alarm {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
```

```
 timeout-active-flows;


}
spi-inval {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
    }
    tx-alarm {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    unclassified {
     bandwidth packets per second;
```

```
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
```

```
isis {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
 isis-data {
  bandwidth packets per second;
  burst packets;
```

```
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   isis-hello {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  jfm {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
```

```
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  l2pt {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
```

```
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  l2tp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
```

```
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  l3dest-miss {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
```

```
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  l3mc-sgv-hit-icl {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  l3mtu-fail {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   l3nhop {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  lacp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
```

```
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ldp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    ldp-hello {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
```

```
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    unclassified {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
```

```
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ldpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
```

```
    recover-time seconds;
    timeout-active-flows;


   }
  }
  lldp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;
```

```
 }
}
lmp {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;
```

```
 }
}
lmpv6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;



 }
}
```

```
localnh {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;



 }
}
mac-host {
 aggregate {
```

```
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  martian-address {
   aggregate {
    bandwidth packets per second;
    burst packets;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  mcast-snoop {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    recover-time seconds;
   }
   igmp {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
```

```
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    mld {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   pim {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   mld {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  mlp {
   add {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;



    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
```

```
 recover-time seconds;




}
delete {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
}
lookup {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
```

```
macpin-exception {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
unclassified {
 bandwidth packets per second;
 burst packets;
```

```
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  msdp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
```

```
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  msdpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }


  mvrp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ndpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }

 recover-time seconds;

}
invalid-hop-limit {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   neighbor-advertisement {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    neighbor-solicitation {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   redirect {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    router-advertisement {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
```

```
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;


  }
  router-solicitation {
   bandwidth packets per second;
   burst packets;
   bypass-aggregate;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
   flow-detect-time seconds;
   flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
```

```
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  nonucast-switch {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
```

```
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ntp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;
```

```
 }
}
oam-lfm {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;
```

```
 }
}
ospf {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
 ospf-hello {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ospf-uncls {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ospf-hello {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ospfv3v6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
```

```
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ospfv3v6-hello {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
```

```
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    ospfv3v6-uncls {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  overlay {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    arp {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ndpv6 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
```

```
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



   }
  }
  pfcp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;



    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
```

```
    disable-fpc;
   }



   recover-time seconds;





  }
 }
 pfe-alive {
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
   flow-detect-time seconds;
   flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
```

```
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pim {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
```

```
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pim-ctrl {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
```

```
 timeout-active-flows;


}
hello {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
    }
    join {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   pim-data {
```

```
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;



 }
}
pimv6 {
 aggregate {
  bandwidth packets per second;
```

```
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pkt-inject {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
```

```
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pmvrp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pos {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
```

```
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ppp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }

 recover-time seconds;

}
authentication {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
```

```
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    echo-rep {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   echo-req {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ipcp {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    ipv6cp {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
```

```
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   isis {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
```

```
 recover-time seconds;
 timeout-active-flows;


}
lcp {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
}
mlppp-lcp {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
mplscp {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  pppoe {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
    recover-time seconds;
   }
   padi {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
```

```
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    padm {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   padn {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    pado {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   padr {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    pads {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
```

```
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    padse {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
```

```
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
padt {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
```

```
    timeout-active-flows;


  }
 }
 proto-802-1x {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;
```

```
 }
}
ptp {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
```

```
}
ptpv6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
}
pvstp {
```

```
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;



 }
}
radius {
 accounting {
  bandwidth packets per second;
```

```
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
```

```
    disable-routing-engine;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }
```

```
    recover-time seconds;
```

```
   }
   authorization {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
```

```
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    server {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  redirect {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
```

```
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
reject {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
  }
  flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
  }
  flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
```

```
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   rejectv6 {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  re-services {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   captive-portal {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
re-services-v6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
  captive-portal {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
```

```
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



   }
  }
  resolve {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }

    recover-time seconds;

   }
   mcast-v4 {
    bandwidth packets per second;
```

```
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   mcast-v6 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
```

```
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   other {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
```

```
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   ucast-v4 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
```

```
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    ucast-v6 {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  rip {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
ripv6 {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
  }
  flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
```

```
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  rsvp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
```

```
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    rsvp-hello {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
```

```
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   rsvp-uncls {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  rsvpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
```

```
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  sample {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
    }

    recover-time seconds;
```

```
}
host {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;
```

```
}
pfe {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
sflow {
 bandwidth packets per second;
```

```
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   syslog {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
```

```
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   tap {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
```

```
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  sample-dest {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  sample-source {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
```

```
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  sctp {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  services {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;




}
bsdt {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
```

```
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   packet {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
```

```
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   snmp {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
```

```
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  snmpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
```

```
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ssh {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
```

```
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  sshv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
```

```
    timeout-active-flows;


  }
 }
 stp {
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
   flow-detect-time seconds;
   flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;
```

```
 }
}
syslog {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
```

```
}
tacacs {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
   hostbound-queue hostbound-queue;
  }
  hostbound-queue hostbound-queue;
  no-flow-logging;
  priority(high | low | medium);
  recover-time seconds;
  timeout-active-flows;


 }
}
tcc {
```

```
aggregate {
 bandwidth packets per second;
 burst packets;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;



}
ethernet-tcc {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   iso-tcc {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;


  }
 }
 tcp-flags {
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
```

```
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
 }


 recover-time seconds;





}
established {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
```

```
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    initial {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
```

```
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
```

```
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
telnet {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  flow-detection-mode(automatic | off | on);
  flow-detect-time seconds;
  flow-level-bandwidth {
   logical-interface packets per second;
   physical-interface packets per second;
   subscriber packets per second;
  }
  flow-level-control {
   logical-interface(drop | keep | police);
   physical-interface(drop | keep | police);
   subscriber(drop | keep | police);
  }
  flow-level-detection {
   logical-interface(automatic | off | on);
   physical-interface(automatic | off | on);
   subscriber(automatic | off | on);
  }
  flow-recover-time seconds;
  flow-timeout-time seconds;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
```

```
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
 }
 telnetv6 {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
```

```
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  ttl {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
```

```
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  tunnel-fragment {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
```

```
    recover-time seconds;
    timeout-active-flows;


   }
  }
  tunnel-ka {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;
```

```
 }
}
unclassified {
 aggregate {
  bandwidth packets per second;
  burst packets;
  disable-fpc;
  disable-logging;
  disable-routing-engine;
  fpc name {
   bandwidth-scale percent;
   burst-scale percent;
   disable-fpc;
  }
  recover-time seconds;
 }
 control-layer2 {
  bandwidth packets per second;
  burst packets;
  bypass-aggregate;
  disable-fpc;
```

```
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   control-v4 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
```

```
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   control-v6 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
```

```
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;


  }




  fw-host {
   bandwidth packets per second;
   burst packets;
   bypass-aggregate;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
   flow-detect-time seconds;
   flow-level-bandwidth {
```

```
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    host-route-tcp-v4 {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
```

```
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   host-route-udp-v4 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    host-route-v4 {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   host-route-v6 {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
```

```
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
    mcast-copy {
     bandwidth packets per second;
     burst packets;
     bypass-aggregate;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
```

```
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   other {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
```

```
    burst-scale percent;
    disable-fpc;
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;


  }





 }
 unknown-l2mc {
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
   flow-detection-mode(automatic | off | on);
   flow-detect-time seconds;
   flow-level-bandwidth {
    logical-interface packets per second;
    physical-interface packets per second;
    subscriber packets per second;
   }
   flow-level-control {
    logical-interface(drop | keep | police);
    physical-interface(drop | keep | police);
    subscriber(drop | keep | police);
   }
   flow-level-detection {
    logical-interface(automatic | off | on);
    physical-interface(automatic | off | on);
    subscriber(automatic | off | on);
   }
   flow-recover-time seconds;
   flow-timeout-time seconds;
```

```
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  urpf-fail {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
```

```
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
     }
     hostbound-queue hostbound-queue;
     no-flow-logging;
     priority(high | low | medium);
     recover-time seconds;
     timeout-active-flows;


    }
   }
   vcipc-udp {
    aggregate {
     bandwidth packets per second;
     burst packets;
     disable-fpc;
     disable-logging;
     disable-routing-engine;
     flow-detection-mode(automatic | off | on);
     flow-detect-time seconds;
     flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
     }
     flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
     }
     flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
     }
     flow-recover-time seconds;
     flow-timeout-time seconds;
     fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
```

```
    hostbound-queue hostbound-queue;
   }
   hostbound-queue hostbound-queue;
   no-flow-logging;
   priority(high | low | medium);
   recover-time seconds;
   timeout-active-flows;


  }
 }
 virtual-chassis {
  aggregate {
   bandwidth packets per second;
   burst packets;
   disable-fpc;
   disable-logging;
   disable-routing-engine;
```

```
   fpc name {
    bandwidth-scale percent;
    burst-scale percent;
    disable-fpc;
   }


   recover-time seconds;
```

```
}
control-high {
 bandwidth packets per second;
 burst packets;
 bypass-aggregate;
 disable-fpc;
 disable-logging;
 disable-routing-engine;
 flow-detection-mode(automatic | off | on);
 flow-detect-time seconds;
 flow-level-bandwidth {
  logical-interface packets per second;
  physical-interface packets per second;
  subscriber packets per second;
 }
 flow-level-control {
  logical-interface(drop | keep | police);
  physical-interface(drop | keep | police);
  subscriber(drop | keep | police);
 }
 flow-level-detection {
  logical-interface(automatic | off | on);
  physical-interface(automatic | off | on);
  subscriber(automatic | off | on);
 }
 flow-recover-time seconds;
 flow-timeout-time seconds;
 fpc name {
  bandwidth-scale percent;
  burst-scale percent;
  disable-fpc;
  hostbound-queue hostbound-queue;
 }
 hostbound-queue hostbound-queue;
 no-flow-logging;
 priority(high | low | medium);
 recover-time seconds;
 timeout-active-flows;


}
control-low {
```

```
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;



   }
   unclassified {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
```

```
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   vc-packets {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
```

```
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   vc-ttl-errors {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
      logical-interface packets per second;
      physical-interface packets per second;
      subscriber packets per second;
    }
    flow-level-control {
      logical-interface(drop | keep | police);
      physical-interface(drop | keep | police);
      subscriber(drop | keep | police);
    }
    flow-level-detection {
      logical-interface(automatic | off | on);
      physical-interface(automatic | off | on);
      subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
      bandwidth-scale percent;
      burst-scale percent;
      disable-fpc;
      hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


  }
}
vrrp {
  aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
      logical-interface packets per second;
      physical-interface packets per second;
```

```
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  vrrpv6 {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
```

```
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
  vxlan {
   aggregate {
    bandwidth packets per second;
    burst packets;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
```

```
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
   bfd {
    bandwidth packets per second;
    burst packets;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode(automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
     logical-interface packets per second;
     physical-interface packets per second;
     subscriber packets per second;
    }
    flow-level-control {
     logical-interface(drop | keep | police);
     physical-interface(drop | keep | police);
     subscriber(drop | keep | police);
    }
```

```
    flow-level-detection {
     logical-interface(automatic | off | on);
     physical-interface(automatic | off | on);
     subscriber(automatic | off | on);
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc name {
     bandwidth-scale percent;
     burst-scale percent;
     disable-fpc;
     hostbound-queue hostbound-queue;
    }
    hostbound-queue hostbound-queue;
    no-flow-logging;
    priority(high | low | medium);
    recover-time seconds;
    timeout-active-flows;


   }
  }
 }
```

Hierarchy Level

```
[edit system ddos-protection]
```

Description

DDOS protocol parameters

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

protocols

IN THIS SECTION

- Syntax | 13748
- Hierarchy Level | 13749
- Description | 13749
- Options | 13750
- Required Privilege Level | 13750
- Release Information | 13751

Syntax

```
protocols {
    bgp {
        ... bgp-configuration ...
    }
    isis {
        ... isis-configuration ...
    }
    ldp {
        ... ldp-configuration ...
    }
    mpls {
        ... mpls -configuration ...
    }
    msdp {
        ... msdp-configuration ...
    }
    mstp {
        ... mstp-configuration ...
    }
```

```
    ospf {
        ... ospf-configuration ...
    }
    ospf3 {
        ... ospf3-configuration ...
    }
    pim {
        ... pim-configuration ...
    }
    rip {
        ... rip-configuration ...
    }
    ripng {
        ... ripng-configuration ...
    }
    rstp {
        rstp-configuration;
    }
    rsvp{
        ... rsvp-configuration ...
    }
    vstp {
        vstp configuration;
    }
    vpls {
        vpls configuration;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Specify the protocol for a routing instance. You can configure multiple instances of many protocol types. Not all protocols are supported on the switches. See the switch CLI.

Options

bgp Specify BGP as the protocol for a routing instance.

isis Specify IS-IS as the protocol for a routing instance.

ldp Specify LDP as the protocol for a routing instance or for a virtual router instance.

l2vpn Specify Layer 2 VPN as the protocol for a routing instance.

mpls Specify MPLS as the protocol for a routing instance.

msdp Specify the Multicast Source Discovery Protocol (MSDP) for a routing instance.

mstp Specify the Multiple Spanning Tree Protocol (MSTP) for a virtual switch routing instance.

ospf Specify OSPF as the protocol for a routing instance.

ospf3 Specify OSPF version 3 (OSPFv3) as the protocol for a routing instance.

> **NOTE**: OSPFv3 supports the `no-forwarding`, `virtual-router`, and `vrf` routing instance types only.

pim Specify the Protocol Independent Multicast (PIM) protocol for a routing instance.

rip Specify RIP as the protocol for a routing instance.

ripng Specify RIP next generation (RIPng) as the protocol for a routing instance.

rstp Specify the Rapid Spanning Tree Protocol (RSTP) for a virtual switch routing instance.

rsvp Specify the RSVP for a routing instance.

vstp Specify the VLAN Spanning Tree Protocol (VSTP) for a virtual switch routing instance.

vpls Specify VPLS as the protocol for a routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for RIPng introduced in Junos OS Release 9.0.

`mpls` and `rsvp` options added in Junos OS Release 15.1.

RELATED DOCUMENTATION

Example: Configuring Multiple Routing Instances of OSPF

protocols (DDoS)

IN THIS SECTION

- Syntax | 13751
- Hierarchy Level | 13752
- Description | 13752
- Options | 13753
- Required Privilege Level | 13765
- Release Information | 13765

Syntax

```
protocols protocol-group (aggregate | packet-type) {
    bandwidth packets-per-second;
    burst size;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode (automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
```

```
        logical-interface flow-bandwidth;
        physical-interface flow-bandwidth;
        subscriber flow-bandwidth;
    }
    flow-level-control {
        logical-interface flow-control-mode;
        physical-interface flow-control-mode;
        subscriber flow-control-mode;
    }
    flow-level-detection {
        logical-interface flow-operation-mode;
        physical-interface flow-operation-mode;
        subscriber flow-operation-mode;
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc slot-number {
        bandwidth-scale percentage;
        burst-scale percentage;
        disable-fpc;
    }
    no-flow-logging
    priority level;
    recover-time seconds;
    timeout-active-flows;
}
```

Hierarchy Level

```
[edit system ddos-protection]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure control plane DDoS protection policers for all supported packet types within a protocol group or for a particular supported packet type within a protocol group.

Starting in Junos OS Release 22.2R1, we’ve enabled support for following DDoS protocol statements for MX10008 devices also. In earlier releases, the MX10008 devices did not support these DDoS protocol statements.

- `Filter-action`
- `Virtual-chassis`
- `Ttl`
- `Redirect`
- `Re-services`
- `Re-services-v6`
- `Rejectv6`
- `L2pt`
- `Syslog`
- `Vxlan`

> **NOTE**: For the available control plane DDoS protection policer configuration options on PTX Series routers and QFX Series switches, which are different from the options described here, see protocols (DDoS) (ACX Series, PTX Series, and QFX Series).

> **NOTE**: Although the term bandwidth usually refers to bits per second (bps), this feature's `bandwidth` option represents a packets per second (pps) value, and the `burst` option represents number of packets in a burst. These options are explained separately.

Options

aggregate Configure the policer to monitor all control packets within the protocol group. You can configure an aggregate policer for any protocol group.

packet-type (Optional) Name of the control packet type to be policed. You can configure a specific policer for only the following packet types and protocol groups:

- `arp`—The following ARP packet types are available:
  - `aggregate`—Applies to the combination of all types of control packet traffic for this protocol group.

On MX Series devices, prior to Junos OS Release 23.2R1, the ARP request and reply traffic had a single DDoS protocol. Starting in Junos OS Release 23.2R1, you can configure separate DDoS protocol packet-types, `bcast` and `ucast`, at the [`edit ddos-protection protocols arp`] hierarchy level for ARP request and reply traffic. The separate DDoS policers provide an improved packet rate limiting and priority handling for ARP traffic.

- `bgp`—The following BGP packet types are available:
  - `aggregate`—Applies to the combination of all types of control packet traffic for this protocol group.
- `bgpv6`—The following BGPv6 packet types are available:
  - `aggregate`—Applies to the combination of all types of control packet traffic for this protocol group.
- `dhcpv4`—The following packet types are available for DHCPv4 traffic:
  - `ack`—DHCPACK packets.
  - `bad-packets`—DHCPv4 packets with bad formats.
  - `bootp`—DHCPBOOTP packets.
  - `decline`—DHCPDECLINE packets.
  - `discover`—DHCPDISCOVER packets.
  - `force-renew`—DHCPFORCERENEW packets.
  - `inform`—DHCPINFORM packets.
  - `lease-active`—DHCPLEASEACTIVE packets.
  - `lease-query`—DHCPLEASEQUERYpackets.
  - `lease-unassigned`—DHCPLEASEUNASSIGNED packets.
  - `lease-unknown`—DHCPLEASEUNKNOWN packets.
  - `nak`—DHCPNAK packets.
  - `no-message-type`—DHCP packets that are missing the message type.
  - `offer`—DHCPOFFER packets.

    - `release`—DHCPRELEASE packets.
    - `renew`—DHCPRENEW packets.
    - `request`—DHCPREQUEST packets.
    - `unclassified`—All unclassified packets in the protocol group.
- `dhcpv6`—The following packet types are available for DHCPv6 traffic:
    - `advertise`—ADVERTISE packets.
    - `confirm`—CONFIRM packets.
    - `decline`—DECLINE packets.
    - `information-request`—INFORMATION-REQUEST packets.
    - `leasequery`—LEASEQUERY packets.
    - `leasequery-data`—LEASEQUERY-DATA packets.
    - `leasequery-done`—LEASEQUERY-DONE packets.
    - `leasequery-reply`—LEASEQUERY-REPLY packets.
    - `rebind`—REBIND packets.
    - `reconfigure`—RECONFIGURE packets.
    - `relay-forward`—RELAY-FORWARD packets.
    - `relay-reply`—RELAY-REPLY packets.
    - `release`—RELEASE packets.
    - `renew`—RENEW packets.
    - `reply`—REPLY packets.
    - `request`—REQUEST packets.
    - `solicit`—SOLICIT packets.
    - `unclassified`—All unclassified packets in the protocol group.
- `filter-action`—The following packet types are available for unclassified firewall filter action packets, sent to the host because of reject terms in firewall filters:

    - `filter-v4`—Unclassified IPv4 filter action packets.
    - `filter-v6`—Unclassified IPv6 filter action packets.
    - `other`—All other unclassified filter action packets that are not IPv4 or IPv6.
- `frame-relay`—The following packet types are available for Frame Relay traffic:
    - `frf15`—Multilink frame relay FRF.15 packets.
    - `frf16`—Multilink frame relay FRF.16 packets.
- `ip-fragments`—The following packet types are available for IP fragments:
    - `first-fragment`—First IP fragment.
    - `trail-fragment`—Last IP fragment.
- `ip-options`—The following packet types are available for IP option traffic:
    - `non-v4v6`—Options packets other than IPv4/v6.
    - `router-alert`—Router alert options packets.
    - `unclassified`—All unclassified packets in the protocol group.
- `l2tp`—The following packet types are available for L2TP LNS subscriber management network environments in Junos OS releases 13.3R5 and 14.1X50 (this option has been obsoleted by L2TP ERA in current Enhanced Subscriber Management environments):
    - `cdn`—Call-Disconnect-Notify message packets.
    - `hello`—Hello message packets.
    - `iccn`—Incoming-Call-Connected message packets.
    - `icrq`—Incoming-Call-Request message packets.
    - `scccn`—Start-Control-Connection-Connected message packets.
    - `sccrq`—Start-Control-Connection-Request message packets.
    - `stopccn`—Stop-Control-Connection-Notification message packets.
    - `unclassified`—All unclassified packets in the protocol group.
- `mcast-snoop`—Control traffic for multicast snooping.

    - `igmp`—Snooped IGMP traffic.
    - `mld`—Snooped MLD traffic.
    - `pim`—Snooped PIM control traffic.
- `mlp`—The following MLP packet types are available:
    - `aggregate`—Applies to the combination of all types of control packet traffic for this protocol group.
    - `add`—Add requests; internal MAC address learning request packets sent to the host.
    - `delete`—Delete requests; internal MAC address learning request packets sent to the host.
    - `lookup`—Lookup requests; internal MAC address learning request packets sent to the host.
    - `unclassified`—All unclassified packets in the protocol group.
    - `macpin-exception`—Exceptions to MAC address pinning (wherein dynamically learned MAC addresses are pinned to prevent looping caused by MAC moves from duplicate MAC detection).
- `ndpv6`—The following NDPv6 packet types are available, except where noted, starting in 14.1R8, 14.2R8, 15.1R5, 15.1F7, and 16.1R1:
    - `aggregate`—Applies to the combination of all types of control packet traffic for this protocol group.
    - `invalid-hop-limit`—(Starting in 16.1R2) Invalid hop limit packets. These messages might represent crafted packets in a malicious network-based packet flood.
    - `neighbor-advertisement`—Neighbor advertisement packets. These are messages used for duplicate address detection and to test reachability of neighbors. Neighbor advertisements are sent in response to neighbor solicitation messages.
    - `neighbor-solicitation`—Neighbor solicitation packets. These are messages used for duplicate address detection and to test reachability of neighbors.
    - `redirect`—Redirect packets.

    - `router-advertisement`—Router advertisement packets. These are messages sent to announce the presence of the router, advertise prefixes, assist in address configuration, and share other link information such as MTU size and hop limit. The IPv6 nodes on the link can use this information to configure themselves with an IPv6 address and routing information such as the default gateway.
    - `router-solicitation`—Router solicitation packets. These are messages sent by IPv6 nodes when they come online to solicit immediate router advertisements from the router.
- `ppp`—The following PPP packet types are available:
    - `authentication`—PPP authentication protocol packets.
    - `echo-rep`—LCP echo reply packets.
    - `echo-req`—LCP echo request packets.
    - `ipcp`—IP Control Protocol packets.
    - `ipv6cp`—IPv6 Control Protocol packets.
    - `isis`—IS-IS packets.
    - `lcp`—Link Control Protocol packets.
    - `mlppp-lcp`—MLPPP LCP packets.
    - `mplscp`—MPLS Control Protocol packets.
    - `unclassified`—All unclassified packets in the protocol group.
- `pppoe`—The following PPPoE packet types are available:
    - `padi`—PADI packets.
    - `padm`—PADM packets.
    - `padn`—PADN packets.
    - `pado`—PADO packets.
    - `padr`—PADR packets.
    - `pads`—PADS packets.

  - `padt`—PADT packets.

- `radius`—The following RADIUS packet types are available:
  - `accounting`—RADIUS accounting packets.
  - `authorization`—RADIUS authorization packets.
  - `server`—RADIUS server traffic.
  - `unclassified`—All unclassified packets in the protocol group.

- `re-services`—The following packet type is available for Routing Engine-based HTTP redirect IPv4 traffic:
  - `captive-portal`—Routing Engine-based captive portal content delivery packets.

- `re-services-v6`—The following packet type is available for Routing Engine-based HTTP redirect IPv6 traffic:
  - `captive-portal`—Routing Engine-based captive portal content delivery packets.

- `resolve`—The following packet types are available for unclassified resolve packets, which are sent to the host because of a traffic request resolve action:
  - `mcast-v4`—Unclassified IPv4 multicast resolve packets.
  - `mcast-v6`—Unclassified IPv6 multicast resolve packets.
  - `ucast-v4`—Unclassified IPv4 unicast resolve packets.
  - `ucast-v6`—Unclassified IPv6 unicast resolve packets.
  - `other`—All other unclassified resolve packets.

- `sample`—The following sample packet types are available:
  - `host`—Host packets.
  - `pfe`—Packet Forwarding Engine packets.
  - `syslog`—System log message packets.
  - `tap`—TAP packets.

- `tcp-flags`—The following TCP-flagged packet types are available:

- `established`—TCP packets with ACK or RST flags set.
- `initial`—TCP packets with SYN flag set and ACK flag not set.
- `unclassified`—TCP packets with flags set any other way than the established and initial packets.

- `unclassified`—The following unclassified packet types are available:
  - `control-layer2`—Unclassified layer 2 control packets.
  - `control-v4`—Unclassified IPv4 control packets.
  - `control-v6`—Unclassified IPv6 control packets.
  - `fw-host`—Unclassified send-to-host firewall packets.
  - `host-route-v4`—Unclassified IPv4 routing protocol and host packets in traffic sent to the router local interface address.
  - `host-route-v6`—Unclassified IPv6 routing protocol and host packets in traffic sent to the router local interface address.
  - `other`—All unclassified packets that do not belong to another type.
- `virtual-chassis`—The following packet types are available for virtual chassis packets:
  - `control-low`—Low-priority control packets.
  - `control-high`—High-priority control packets.
  - `unclassified`—All unclassified packets in the protocol group.
  - `vc-packets`—All exception packets on the virtual chassis link.
  - `vc-ttl-errors`—Virtual chassis TTL error packets.

protocol-group

Name of the protocol group for which traffic is policed. You can configure a policer for any of the following protocol groups:

- `amtv4`—IPv4 AMT traffic.
- `amtv6`—IPv6 AMT traffic.
- `ancp`—ANCP traffic.
- `ancpv6`—ANCPv6 traffic.

- `arp`—ARP traffic.
- `atm`—ATM traffic.
- `bfd`—BFD traffic.
- `bfdv6`—BFDv6 traffic.
- `bgp`—BGP traffic.
- `bgpv6`—BGPv6 traffic.
- `control`—Control traffic.
- `demux-autosense`—Demux autosensing traffic.
- `dhcpv4`—DHCPv4 traffic.
- `dhcpv6`—DHCPv6 traffic.
- `diameter`—Diameter and Gx-Plus traffic.
- `dns`—DNS traffic.
- `dtcp`—DTCP traffic.
- `dynamic-vlan`—Dynamic VLAN exception traffic.
- `egpv6`—EGPv6 traffic.
- `eoam`—EOAM traffic.
- `esmc`—ESMC traffic.
- `fab-probe`—Fab out probe packets.
- `filter-action`—IPv4 and IPv6 firewall filter action packets sent to the host because of reject terms in firewall filters
- `frame-relay`—Frame relay traffic.
- `ftp`—FTP traffic.
- `ftpv6`—FTPv6 traffic.
- `gre`—GRE traffic.
- `gtp-path-mgmt`—GTP path management traffic.

- `icmp`—ICMP traffic.
- `igmp`—IGMP traffic
- `igmpv4v6`—IGMP v4/v6 traffic.
- `igmpv6`—IGMPv6 traffic.
- `inline-ka`—Inline service interfaces keepalive traffic.
- `inline-svcs`—Inline services traffic.
- `ip-fragments`—IP fragments traffic.
- `ip-options`–IP traffic with IP packet header options.
- `isis`—IS-IS traffic.
- `jfm`—JFM traffic.
- `l2pt`—Layer 2 protocol tunneling traffic.
- `lacp`—LACP traffic.
- `ldp`—LDP traffic.
- `ldpv6`—LDPv6 traffic.
- `lldp`—LLDP traffic.
- `lmp`—LMP traffic.
- `lmpv6`—LMPv6 traffic.
- `mac-host`—Layer 2 MAC send-to-host traffic.
- `mcast-snoop`—Control traffic for multicast snooping.
- `mlp`—MLP traffic.
- `msdp`—MSDP traffic.
- `msdpv6`—MSDPv6 traffic.
- `multicast-copy`—Host copy traffic due to multicast routing.
- `mvrp`—MVRP traffic.

- `ndpv6`—NDPv6 traffic.
- `ntp`—NTP traffic.
- `oam-lfm`—OAM-LFM traffic.
- `ospf`—OSPF traffic.

  Starting in Junos OS Release 23.2R1, the `ospf` protocol has two sub protocols: `ospf-hello` to classify the DDOS high priority OSPF traffic and `ospf-unclassified` to classify the DDOS low priority OSPF traffic.
- `ospfv3v6`—OSPFv3/IPv6 traffic.

  Starting in Junos OS Release 23.2R1, the `ospfv3v6` protocol has two sub protocols: `ospfv3v6-hello` to classify the DDOS high priority OSPF traffic and `ospfv3v6-unclassified` to classify the DDOS low priority OSPF traffic.
- `pfcp`—Packet Forwarding Control Protocol (PFCP) traffic.
- `pfe-alive`—Packet Forwarding Engine keepalive traffic.
- `pim`—PIM traffic.
- `pimv6`—PIMv6 traffic.
- `pmvrp`—PMVRP traffic.
- `pos`—POS traffic.
- `ppp`—PPP traffic.
- `pppoe`—PPPoE traffic.
- `ptp`—PTP traffic.
- `pvstp`—PVSTP traffic.
- `radius`—RADIUS traffic.
- `re-services`—Captive portal content delivery IPv4 traffic for Routing Engine HTTP redirect.
- `re-services-v6`—Captive portal content delivery IPv6 traffic for Routing Engine HTTP redirect.
- `redirect`—Traffic that triggers ICMP redirects.

- `reject`—Packets rejected by a next-hop forwarding decision.
- `rejectv6`—V6 packets rejected by a next-hop forwarding decision.
- `resolve`—Unclassified IPv4 and IPv6 resolve packets sent to the host because of a traffic request resolve action.
- `rip`—RIP traffic.
- `ripv6`—RIPv6 traffic.
- `rsvp`—RSVP traffic.

  Starting in Junos OS Release 23.2R1, the `rsvp` protocol has two sub protocols: `rsvp-hello` to classify the DDOS high priority RSVP traffic and `ospfv3v6-unclassified` to classify the DDOS low priority RSVP traffic.
- `rsvpv6`—RSVPv6 traffic.
- `services`–Service traffic.
- `snmp`—SNMP traffic.
- `snmpv6`—SNMPv6 traffic.
- `ssh`—SSH traffic.
- `sshv6`—SSHv6 traffic.
- `stp`—STP traffic.
- `syslog`—System log messages UDP traffic on port 6333 for the Routing Engine syslog server.
- `tacacs`—TACACS traffic.
- `tcp-flags`—Traffic with TCP flags.
- `telnet`—TELNET traffic.
- `telnetv6`—TELNETv6 traffic.
- `ttl`—TTL traffic.
- `tunnel-fragment`—Tunnel fragments traffic.
- `tunnel-ka`—Tunnel keepalive traffic.

- `unclassified`—Unclassified traffic.

  Starting in Junos OS Release 23.2R1, the `unclassified` protocol has two sub protocols: `host-route-tcp-v4` and `host-route-udp-v4` to classify the DDOS low priority TCP and UDP traffic.

- `virtual-chassis`—Virtual chassis traffic.
- `vrrp`—VRRP traffic.
- `vrrpv6`—VRRPv6 traffic.
- `vxlan`—VXLAN Layer 2 and Layer 3 traffic.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers

Understanding Distributed Denial-of-Service Protection with IPv6 Neighbor Discovery Protocol

protocols (DDoS) (ACX Series, PTX Series, and QFX Series)

IN THIS SECTION

Syntax (ACX Series Routers) | 13766

- Syntax (PTX Series Routers and QFX Series Switches) | **13766**
- Syntax (PTX Series and ACX7100-48L for Junos OS Evolved) | **13767**
- Hierarchy Level | **13767**
- Description | **13767**
- Options | **13768**
- Required Privilege Level | **13790**
- Release Information | **13790**

Syntax (ACX Series Routers)

```
protocols protocol-group aggregate {
    bandwidth packets-per-second;
    burst size;
    disable-logging;
    disable-routing-engine;
    priority level;
}
```

Syntax (PTX Series Routers and QFX Series Switches)

```
protocols protocol-group (aggregate | packet-type) {
    bandwidth packets-per-second;
    burst size;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    fpc slot-number {
        bandwidth-scale percentage;
        burst-scale percentage;
        disable-fpc;
    }
    priority level;
}
```

Syntax (PTX Series and ACX7100-48L for Junos OS Evolved)

```
protocols{
    arp;
    bfd;
    bfdv6;
    bgp;
    custom;
    sample;
    dhcpv4;
    dhcpv6;
    dhcpv4v6;
    l2tp;
    ppp;
    pppoe;
}
```

Hierarchy Level

```
[edit system ddos-protection]
```

Description

Change default configurable control plane DDoS protection policer parameters for all packet types within a protocol group or for a particular packet type within a protocol group.

> **NOTE**: PTX10003 router doesn't support the `priority` option to change default priority values for aggregate or individual packet type policers.
> QFX10002-60C switches and PTX Series routers do not support the `bypass-aggregate` option.

> **NOTE**: Although the term bandwidth usually refers to bits per second (bps), this feature's `bandwidth` option represents a packets per second (pps) value, and the `burst` option represents number of packets in a burst. These options are explained separately.

Not all protocol groups and packet types listed in Table 103 on page 13769 or Table 104 on page 13778 are supported on all devices. Exceptions include:

- ACX Series routers support only the following protocol group options: `arp`, `bfd`, `bfdv6`, `bgp`, `dhcpv4v6`, `eoam`, `esmc`, `igmp`, `ipmcast-miss` (unknown multicast packets), `isis`, `lacp`, `ldp`, `lldp`, `ndpv6`, `oam-lfm`, `ospf`, `pim`, `pvstp`, `rip`, `rsvp`, `stp`, and `vrrp`.
- PTX10003 and PTX10008 routers do not support the following policer protocol group options: `all-fiber-channel-enode`, `bridge-control`, `diameter`, `garp-reply`, `l2pt`, `ptp`, `radius`, and `tacacs`.
- Other PTX Series routers do not support the following policer protocol group options: `all-fiber-channel-enode`, (however, `arp` protocol group is supported), `bridge-control`, `diameter`, `garp-reply`, `martian-address`, `proto-802-1x`, `ptp`, `pvstp`, `radius`, `stp`, and `tacacs`
- QFX10002-60C switches do not support the following policer protocol group options: `all-fiber-channel-enode`, (however, `arp` protocol group is supported), `bridge-control`, `diameter`, `garp-reply`, `martian-address`, `proto-802-1x`, `ptp`, `radius`, and `tacacs`
- QFX10002, QFX10008, and QFX10016 switches do not support the `ttl` protocol group option.

Options

aggregate Configure parameters for the policer that polices all control packets belonging to the specified protocol as a combined group. An aggregate policer exists for all protocol groups.

> **NOTE**: ACX Series routers support only the aggregate policer for the supported protocol groups.

packet-type Configure policer values for the specified individual control packet type within a protocol group. On some devices, you can configure the packet-type policers in the protocol groups listed in Table 103 on page 13769. For all other protocol groups not listed in Table 103 on page 13769, only aggregate policers are available.

Table 103 on page 13769 lists the protocol groups with packet-type policers available on some devices, and common values for default-configured parameters. Default values can differ among supporting devices and across different Junos OS releases; you can run the `show ddos-protection protocols` or `show ddos-protection protocols parameters` CLI command before modifying any configurable values to see the default policer values for all supported protocol groups and packet types. You can also include a specific protocol group and packet type (or aggregate) with those commands.

Each of the protocol groups in Table 103 on page 13769 also supports the aggregate policer. See Table 104 on page 13778 for the default aggregate policer values for all protocol groups.

Table 103: Packet Types Supported by Control Plane DDoS Protection on PTX Series Routers and QFX Series Switches

Protocol Group	Packet Type	Description	Default Bandwidth (pps)	Default Burst (number of packets)	Default Priority
`arp`	`arp-snoop` or `arp` **NOTE**: Starting in Junos OS Release 20.3R1 and several other maintenance releases on PTX and QFX Series switches, option name `arp-snoop` is renamed simply `arp`.	ARP traffic	10000	1024 or 2048	High
	`unclassified`	Unclassified ARP packets	500	1024	High
`bfd`	`bundle-bfd`	(PTX 10003 only) Link bundle BFD traffic	10000	10000	High
	`multihop-bfd`	Multihop BFD traffic	5000 or 10000	2048 or 10000	High
	`unclassified`	Unclassified BFD packets	10000	2048	High

Table 103: Packet Types Supported by Control Plane DDoS Protection on PTX Series Routers and QFX Series Switches *(Continued)*

Protocol Group	Packet Type	Description	Default Bandwidth (pps)	Default Burst (number of packets)	Default Priority
`dhcpv4` (PTX10003 and PTX10008 routers only; for rate-limiting at line card and RE levels)	`ack`	DHCPACK packets	500	500	Medium
	`bad-packets`	DHCPv4 packets with bad formats	0	0	Low
	`bootp`	DHCPBOOTP packets	300	300	Low
	`decline`	DHCPDECLINE packets	500	500	Low
	`discover`	DHCPDISCOVER packets	500	500	Low
	`force-renew`	DHCPFORCE RE-NEW packets	2000	2000	High
	`inform`	DHCPINFORM packets	500	500	Low
	`lease-active`	DHCPLEASE ACT-IVE packets	2000	2000	High
	`lease-query`	DHCPLEASE-QUERY packets	2000	2000	High

Table 103: Packet Types Supported by Control Plane DDoS Protection on PTX Series Routers and QFX Series Switches *(Continued)*

Protocol Group	Packet Type	Description	Default Bandwidth (pps)	Default Burst (number of packets)	Default Priority
	`lease-unassigned`	DHCPLEASE UN-ASSIGNED packets	2000	2000	High
	`lease-unknown`	DHCPLEASE UN-KNOWN packets	2000	2000	High
	`nak`	DHCPNAK packets	500	500	Low
	`no-message-type`	DHCP packets that are missing the message type	1000	1000	Low
	`offer`	DHCPOFFER packets	1000	1000	Low
	`rebind`	DHCPv4 REBIND packets	2000	2000	High
	`release`	DHCPRELEASE packets	2000	2000	High
	`renew`	DHCPRENEW packets	2000	2000	High

Table 103: Packet Types Supported by Control Plane DDoS Protection on PTX Series Routers and QFX Series Switches *(Continued)*

Protocol Group	Packet Type	Description	Default Bandwidth (pps)	Default Burst (number of packets)	Default Priority
	`request`	DHCPREQUEST packets	1000	1000	Medium
	`unclassified`	All unclassified DHCPv4 packets	300	150	Low
`dhcpv6` (PTX10003 and PTX10008 routers only; for rate-limiting at line card and RE levels)	`advertise`	DHCPv6 ADVERTISE packets	500	500	Low
	`confirm`	DHCPv6 CONFIRM packets	1000	1000	Medium
	`decline`	DHCPv6 DECLINE packets	1000	1000	Low
	`information-request`	DHCPv6 INFORMATION-REQUEST packets	1000	1000	Low
	`leasequery`	DHCPv6 LEASEQUERY packets	1000	1000	Low

Table 103: Packet Types Supported by Control Plane DDoS Protection on PTX Series Routers and QFX Series Switches *(Continued)*

Protocol Group	Packet Type	Description	Default Bandwidth (pps)	Default Burst (number of packets)	Default Priority
	`leasequery-data`	DHCPv6 LEASEQUERY - DATA packets	1000	1000	Low
	`leasequery-done`	LEASEQUERY - DONE packets	1000	1000	Low
	`leasequery-reply`	DHCPv6 LEASEQUERY - REPLY packets	1000	1000	Low
	`rebind`	DHCPv6 REBIND packets	2000	2000	Medium
	`reconfigure`	DHCPv6 RECONFIGURE packets	1000	1000	Low
	`relay-forward`	DHCPv6 RELAY-FORWARD packets	1000	1000	Low
	`relay-reply`	DHCPv6 RELAY-REPLY packets	1000	1000	Low

Table 103: Packet Types Supported by Control Plane DDoS Protection on PTX Series Routers and QFX Series Switches *(Continued)*

Protocol Group	Packet Type	Description	Default Bandwidth (pps)	Default Burst (number of packets)	Default Priority
	`release`	DHCPv6 RELEASE packets	2000	2000	High
	`renew`	DHCPv6 RENEW packets	2000	2000	Medium
	`reply`	DHCPv6 REPLY packets	1000	1000	Medium
	`request`	DHCPv6 REQUEST packets	1000	1000	Medium
	`solicit`	DHCPv6 SOLICIT packets	500	500	Low
	`unclassified`	All unclassified DHCPv6 packets	3000	3000	Low
`eoam`	`oam-cfm`	Ethernet OAM CFM traffic	1000	1024 or 2048	High
	`unclassified`	Unclassified Ethernet OAM traffic	100000	1024 or 2048	High

Table 103: Packet Types Supported by Control Plane DDoS Protection on PTX Series Routers and QFX Series Switches *(Continued)*

Protocol Group	Packet Type	Description	Default Bandwidth (pps)	Default Burst (number of packets)	Default Priority
`igmpv6`	`mld`	MLD traffic	1000 or 5000	1024 or 2048	High
	`unclassified`	Unclassified IGMPv6 packets	1000 or 90000	1024 or 2048	High
`ldp`	`ldp-hello` Some devices have an `ldp-hello` aggregate policer. Only the following devices support this packet type policer: • PTX10003 and PTX10008 routers • QFX10002, QFX10008, and QFX10016 switches	LDP HELLO traffic	1000	1024	High
	`unclassified`	LDP unclassified packets	1000	1024	High

Table 103: Packet Types Supported by Control Plane DDoS Protection on PTX Series Routers and QFX Series Switches *(Continued)*

Protocol Group	Packet Type	Description	Default Bandwidth (pps)	Default Burst (number of packets)	Default Priority
`mcast-snoop`	`igmp`	Control packets for IGMP snooping	500, 5000, or 20000	2048 or 5000	High
	`mld`	Control packets for MLD snooping	500, 2000, or 5000	2048	High
	`pim`	Control packets for PIM snooping	500, 2000, or 5000	2048	High
	`unclassified`	Unclassified multicast snooping control packets	500	2048	High
`radius`	`accounting`	RADIUS accounting packets	200	2048	High
	`authorization`	RADIUS authorization packets	200	2048	High
	`server`	RADIUS server traffic	200	2048	High

Table 103: Packet Types Supported by Control Plane DDoS Protection on PTX Series Routers and QFX Series Switches *(Continued)*

Protocol Group	Packet Type	Description	Default Bandwidth (pps)	Default Burst (number of packets)	Default Priority
	`unclassified`	Unclassified RADIUS traffic	200	2048	High
`tcc`	`ethernet-tcc`	TCC-encapsulated Ethernet traffic	100	100, 1024 or 2048	High
	`iso-tcc`	TCC-encapsulated ISO traffic	100	100, 1024 or 2048	High
	`unclassified`	Unclassified TCC-encapsulated traffic	100	1024 or 2048	High

protocol-group

Configure policer values for the specified protocol group. You can configure the aggregate policer for any of the following protocol groups listed in Table 104 on page 13778. The table shows the aggregate policer default-configured parameters for each protocol group. Default values can differ among supporting devices and across different Junos OS releases; you can run the `show ddos-protection protocols` or `show ddos-protection protocols parameters` CLI command before modifying any configurable values to see the default policer values for all supported protocol groups. You can also include a specific protocol group with or without the `aggregate` option with those commands to see the aggregate policer parameters.

The protocol groups in Table 104 on page 13778 also support any corresponding individual packet-type policers listed in Table 103 on page 13769.

Table 104: Protocol Groups Supported by Control Plane DDoS Protection on ACX Series Routers, PTX Series Routers, and QFX Series Switches

Protocol Group	Description	Default Bandwidth (pps)	Default Burst (number of packets)
		The default bandwidth and burst values varies based on the platform type and underlying DDoS infra settings.	
`all-fiber-channel-enode`	Fiber channel ENode traffic	10	1024 or 2048
`arp` or `arp-snoop` (Different platforms support an aggregate policer protocol group option for ARP traffic named either `arp` or `arp-snoop`.)	ARP traffic **NOTE**: On some platforms, the `arp` protocol group encompasses `arp-snoop` as a packet type. See Table 103 on page 13769. Starting in Junos OS Release 20.3R1 and other maintenance releases on PTX and QFX Series switches, the `arp-snoop` packet type option in the `arp` protocol group is renamed `arp`.	1000, 2000, or 10000	512, 1024, 2000, or 2048
`bfd`	Single-hop BFD traffic	1000, 6200, 10000, 20000, or 250000	512, 2048 or 20000
`bfdv6`	BFDv6 traffic	512, 3000, 10000, 20000, or 250000	512, 2048 or 20000
`bgp`	BGP traffic	1200, 1500, 3000, 5000, 10000, 20000, or 250000	512, 2048, 4096, or 20000

Table 104: Protocol Groups Supported by Control Plane DDoS Protection on ACX Series Routers, PTX Series Routers, and QFX Series Switches *(Continued)*

Protocol Group	Description	Default Bandwidth (pps)	Default Burst (number of packets)
		The default bandwidth and burst values varies based on the platform type and underlying DDoS infra settings.	
`bridge-control`	Bridge Control traffic	10	2048
`dhcpv4` (PTX10003 and PTX10008 routers only)	Aggregate for all DHCPv4 traffic (priority Medium) **NOTE**: On PTX10003 and PTX10008 routers, use this option for rate-limiting at PFE line card and RE levels. Use aggregate option `dhcpv4v6` for rate-limiting at PFE chip level.	5000	5000
`dhcpv6` (PTX10003 and PTX10008 routers only)	Aggregate for all DHCPv6 traffic (priority Low) **NOTE**: On PTX10003 and PTX10008 routers, use this option for rate-limiting at PFE line card and RE levels. Use aggregate option `dhcpv4v6` for rate-limiting at PFE chip level.	5000	5000

Table 104: Protocol Groups Supported by Control Plane DDoS Protection on ACX Series Routers, PTX Series Routers, and QFX Series Switches *(Continued)*

Protocol Group	Description	Default Bandwidth (pps)	Default Burst (number of packets)
		The default bandwidth and burst values varies based on the platform type and underlying DDoS infra settings.	
`dhcpv4v6`	DHCPv4 and DHCPv6 traffic (limits apply to combined traffic) **NOTE**: On PTX10003 and PTX10008 routers, use this aggregate option for rate-limiting at PFE chip level only (priority is Low). Use `dhcpv4` and `dhcpv6` protocol group and individual packet type options for rate-limiting at line card and RE levels.	600 or 5000	512, 2048 or 5000
`diameter`	Diameter and Gx-Plus traffic	200	2048
`dns`	DNS traffic	200	200 or 2048
`dtcp`	DTCP traffic	200	200 or 2048
`egpv6`	EGPv6 traffic	10	10 or 2048

Table 104: Protocol Groups Supported by Control Plane DDoS Protection on ACX Series Routers, PTX Series Routers, and QFX Series Switches *(Continued)*

Protocol Group	Description	Default Bandwidth (pps)	Default Burst (number of packets)
		The default bandwidth and burst values varies based on the platform type and underlying DDoS infra settings.	
`eoam`	Ethernet OAM traffic **NOTE**: On PTX10003 and PTX10008 routers, the aggregate `eoam` protocol group option includes OAM-CFM packets (no `oam-cfm` individual packet type option).	1000, 6200, or 100000	102, 512, 2048, or 10000
`esmc`	ESMC traffic	200	512
`ethernet-tcc`	TCC-encapsulated Ethernet traffic **NOTE**: The `tcc` protocol group option encompasses this as a packet type option on some devices.	100	100 or 2048
`exception`	• MTU traffic • Multicast traffic • TTL traffic (QFX10002, QFX10008, and QFX10016 switches only)	100	2048

Table 104: Protocol Groups Supported by Control Plane DDoS Protection on ACX Series Routers, PTX Series Routers, and QFX Series Switches *(Continued)*

Protocol Group	Description	Default Bandwidth (pps)	Default Burst (number of packets)
		The default bandwidth and burst values varies based on the platform type and underlying DDoS infra settings.	
`ftp`	FTP traffic	500 or 1500	1500 or 2048
`garp-reply`	Gratuitous ARP reply traffic	100	2048
`gre`	GRE traffic	500	500 or 2048
`icmp`	ICMP traffic	500, 1000, or 20000	500, 2048, or 20000
`igmp`	IGMPv4 and IGMPv6 traffic **NOTE**: Use this option on PTX Series and QFX10002-60C devices for IGMPv4 traffic only, and `igmpv6` option for IGMPv6 traffic. On PTX10003 and PTX10008 routers, this option encompasses aggregated IGMP and MLD traffic.	1000, 1600, 5000, or 90000	512, 2048 or 5000
`igmpv6`	IGMPv6 traffic	20000 or 90000	2048 or 5000
`ip-options`	IP traffic with IP packet header options	100	100 or 2048

Table 104: Protocol Groups Supported by Control Plane DDoS Protection on ACX Series Routers, PTX Series Routers, and QFX Series Switches *(Continued)*

Protocol Group	Description	Default Bandwidth (pps)	Default Burst (number of packets)
		The default bandwidth and burst values varies based on the platform type and underlying DDoS infra settings.	
`ipmcast-miss` (ACX Series only)	Unknown IPv4 and IPv6 multicast packets	600	512
`isis`	IS-IS traffic	1000, 1200, 5000, or 20000	512, 2048, 4096, or 20000
`isis-data`	ISIS-Data traffic	5000, 8000, or 10000	4096 or 8000
`isis-hello`	ISIS-Hello traffic	1000, 5000, or 12000	4096 or 12000
`iso-tcc`	TCC-encapsulated ISO traffic **NOTE**: The `tcc` protocol group option encompasses this as a packet type option on some devices.	100	100 or 2048
`l2pt`	Layer 2 protocol tunneling traffic	500	2048
`l2tp`	Layer 2 tunneling protocol traffic	500	500 or 2048
`lacp`	LACP traffic	800, 1000, or 2000	512, 300, 2000, or 2048
`ldp`	LDP traffic	1200, 5000, 10000, or 20000	200, 512, 2048, or 20000

Table 104: Protocol Groups Supported by Control Plane DDoS Protection on ACX Series Routers, PTX Series Routers, and QFX Series Switches *(Continued)*

Protocol Group	Description	Default Bandwidth (pps)	Default Burst (number of packets)
		The default bandwidth and burst values varies based on the platform type and underlying DDoS infra settings.	
`ldp-hello`	LDP hello packets **NOTE**: The following devices have an `ldp-hello` packet type policer and do not use this aggregate policer: • PTX10003 and PTX10008 routers • QFX10002, QFX10008, and QFX10016 switches	1000 or 5000	2048 or 5000
`lldp`	LLDP traffic	100, 800, or 2000	300, 512, 2000, or 2048
`lmp`	LMP traffic	100	100 or 2048
`martian-address`	Martian address	200	20
`mcast-snoop`	Control traffic for multicast snooping	5000, 20000, or 22000	2048, 6000, or 20000

Table 104: Protocol Groups Supported by Control Plane DDoS Protection on ACX Series Routers, PTX Series Routers, and QFX Series Switches *(Continued)*

Protocol Group	Description	Default Bandwidth (pps)	Default Burst (number of packets)
		The default bandwidth and burst values varies based on the platform type and underlying DDoS infra settings.	
`mld`	MLD traffic **NOTE**: The `igmpv6` protocol group option encompasses this as a packet type option on some devices.	1000	2048
`msdp`	MSDP traffic	20000	20000
`multihop-bfd`	Multihop BFD traffic **NOTE**: The `bfd` protocol group option encompasses this as a packet type option on some devices.	1500	2048
`ndpv6`	NDPv6 traffic	100, 1000, or 2000	512, 1024, or 2000
`ntp`	NTP traffic	20000	20000

Table 104: Protocol Groups Supported by Control Plane DDoS Protection on ACX Series Routers, PTX Series Routers, and QFX Series Switches *(Continued)*

Protocol Group	Description	Default Bandwidth (pps)	Default Burst (number of packets)
		The default bandwidth and burst values varies based on the platform type and underlying DDoS infra settings.	
`oam-cfm`	OAM CFM traffic **NOTE**: The `eoam` protocol group option encompasses this as a packet type option on some devices. On PTX10003 and PTX10008 routers, the aggregate `eoam` protocol group option includes OAM-CFM packets (no `oam-cfm` individual packet type option).	200	2048
`oam-lfm`	OAM LFM traffic	200, 800, 1000, or 20000	512, 1000, 2048, or 20000
`ospf`	OSPF traffic	1200, 5000, 10000, or 20000	200, 512, 2048, 4096, or 20000
`ospf-hello`	OSPF hello packets	1000, 1500, or 10000	2048, 4096, or 20000
overlay	Packets such as ARP and NDP coming over VxLAN tunnel with VxLAN header.	500	200
`pim` (ACX Series only)	PIM IPv4 and IPv6 packets	1600	512

Table 104: Protocol Groups Supported by Control Plane DDoS Protection on ACX Series Routers, PTX Series Routers, and QFX Series Switches *(Continued)*

Protocol Group	Description	Default Bandwidth (pps)	Default Burst (number of packets)
		The default bandwidth and burst values varies based on the platform type and underlying DDoS infra settings.	
`pim-ctrl`	PIM control packets	1000 or 1500	200 or 2048
`pim-data`	PIM data	2000 or 3000	1024 or 2048
`proto-802-1x`	802.1X traffic	200	200 or 2048
`ptp`	PTP traffic	100	2048
`pvstp`	PVSTP traffic	800, 2000, or 20000	512, 2048, or 20000
`radius`	RADIUS traffic	200	2048
`reject`	Packets rejected by a next-hop forwarding decision	100, 200, or 2000	200, 2000, or 2048
`resolve`	Unclassified IPv4 and IPv6 resolve packets sent to the host because of a traffic request resolve action	100, 500, or 5000	100, 2048, or 5000
`rip`	RIP traffic	200, 1200, or 20000	200, 512, 2048, or 20000
`rsvp`	RSVP traffic	1200, 5000, 10000, or 20000	512, 2048, 10000, or 20000

Table 104: Protocol Groups Supported by Control Plane DDoS Protection on ACX Series Routers, PTX Series Routers, and QFX Series Switches *(Continued)*

Protocol Group	Description	Default Bandwidth (pps)	Default Burst (number of packets)
		The default bandwidth and burst values varies based on the platform type and underlying DDoS infra settings.	
`snmp`	SNMP traffic	1000 or 20000	1024, 2048, or 20000
`ssh`	SSH traffic	5000 or 20000	500, 2048, or 20000
`stp`	STP traffic	800 or 20000	512, 2048, or 20000
`tacacs`	TACACS+ traffic	200	2048
`tcc`	Transitional Cross-connect encapsulated traffic	100 or 200	200, 1024, or 2048
`telnet`	Telnet traffic	5000 or 20000	500, 2048, or 20000
`ttl`	Time to Live packets	100 or 2000	2048
`unclassified`	Traffic that cannot be classified into one of the other available protocol groups	100 or 10000	2048 or 10000
`vrrp`	VRRP traffic	512, 1000, 2000, or 20000	512, 1000, 2048, or 20000
`vxlan`	VXLAN Layer 2 and Layer 3 packets	300	10

Table 105: Protocol Groups Supported by Control Plane DDoS Protection on PTX Series Routers for Junos OS Evolved

Protocol Group	Description
`arp` or	Configure ARP traffic
`bfd`	Configure BFD traffic
`bfdv6`	BFDv6 traffic
`bgp`	Configure BGP traffic
`custom`	Configure CUSTOM traffic
`sample`	Configure Sampling traffic
`dhcpv4`	Configure DHCPv4 traffic
`dhcpv6`	Configure DHCPv6 traffic
`dhcpv4v6`	Configure DHCPv4/v6 traffic
`l2tp`	Configure L2TP traffic.
`ppp`	Configure PPP traffic.
`pppoe`	Configure PPPoE traffic.

Starting in Junos OS Evolved Release 23.1R1, on QFX5220, QFX5130, and QFX5700 Series devices, the default DDoS `localnh aggregate bandwidth` value is 1500 pps. Earlier to Junos OS Evolved Release 23.1R1, you can configure the bandwidth value using the [`set system ddos-protection protocols localnh aggregate bandwidth 1500 burst 200`] CLI command. This configuration enables you to prevent LAG interface down time during the heavy local IP traffic.

Starting in Junos OS Release 23.3R1, on QFX Series devices, we've added the `ip-option` protocol support at the [`edit ddos-protection protocols`] hierarchy level.

Starting in Junos OS Release 23.4R1, on QFX5000 Series and EX4000 Series devices, the address resolution protocol (ARP) packets over the virtual tunnel endpoints (VTEP) tunnel goes to the overlay ARP. You can check the ARP details using the `show ddos-protection protocols overlay arp` CLI command. We also introduce the `overlay` statement at the [`edit system ddos-protection protocols`] hierarchy level to control the VxLANDDOS packets.

Starting Junos OS Evolved release 23.4R1 supports, DVLAN (single and dual tag) for subscriber services scaling, performance on ACX7100-48L devices that includes:

- DHCP (IP-DEMUX lite) & PPPoE subscribers (IPv4, IPV6, and Dual stack) with CoS Lawful Intercept and filter support.
- DVLAN (Single and dual tag) with L2TP (LAC) DDOS policers configuration for BBE protocols. Individual BBE protocol and DDOS protocol group configuration is enabled.
- Support for protocol groups DHCPv6, L2TP, PPP, and PPPoE
- Only the aggregate DDOS protocol group `dhcpv4v6` is active.
- Subscriber scale qualification with Class of Support (CoS) for Layer 2 Tunneling Protocol (L2TP), L2TP access concentrator (LAC) subscriber interfaces for IPV4, IPV6 and dual stack.

  The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Control Plane Distributed Denial-of-Service (DDoS) Protection Overview

Configuring Control Plane DDoS Protection

protocols (DDoS) | **13751**

pppoe (DDoS Protection) | **12962**

protocols (Dynamic Profiles)

IN THIS SECTION

- Syntax | 13791
- Hierarchy Level | 13793
- Description | 13793
- Default | 13793
- Required Privilege Level | 13793
- Release Information | 13794

Syntax

```
protocols {
    igmp {
        interface interface-name {
            accounting;
            disable;
            group-limit limit;
            group-policy;
            group-threshold value;
            immediate-leave
            log-interval seconds;
            no-accounting;
            oif-map;
            passive;
            promiscuous-mode;
            ssm-map ssm-map-name;
            ssm-map-policy ssm-map-policy-name
            static {
                group group {
                    source source;
                }
            }
            version version;
        }
    }
```

```
mld {
     interfaceinterface-name {
        (accounting | no-accounting);
        disable;
        group-limit limit;
        group-policy;
        group-threshold value;
        immediate-leave;
        log-interval seconds;
        oif-map;
        passive;
        ssm-map ssm-map-name;
        ssm-map-policy ssm-map-policy-name;
        static {
            group multicast-group-address {
                exclude;
                group-count number;
                group-increment increment;
                source ip-address {
                    source-count number;
                    source-increment increment;
                }
            }
        }
        version version;
    }
}
router-advertisement {
    interfaceinterface-name {
         current-hop-limitnumber;
        default-lifetime seconds;
        dns-server-address
        (managed-configuration | no-managed-configuration);
        max-advertisement-interval seconds;
        min-advertisement-interval seconds;
        (other-stateful-configuration | no-other-stateful-configuration);
        prefixprefix {
            (autonomous | no-autonomous);
            (on-link | no-on-link);
            preferred-lifetime seconds;
            valid-lifetime seconds;
        }
        reachable-time milliseconds;
```

```
            retransmit-timer milliseconds;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Enable Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) Protocol on the router and configure interface-specific values on dynamic interfaces for each protocol.

PIM and MLD manage multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use these protocols to learn which groups have members on each of their attached physical networks. Enable IGMP for the router to receive IPv4 or IPv6 multicast traffic. Enable MLD for the router to receive IPv6 multicast traffic. MLD is needed only for IPv6 networks.

You can also use this statement to enable router advertisement for IPv6 Neighbor Discovery protocol and configure interface-specific values on dynamic interfaces.

You can configure these protocols in a dynamic client profile or a dynamic service profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Default

IGMP is disabled on the router. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the `[edit dynamic-profiles profile-name protocols mld]` and `[edit dynamic-profiles profile-name protocols router-advertisement]` hierarchy levels introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

- Configuring IGMP
- Configuring MLD

protocols (Fabric)

IN THIS SECTION

- Syntax | 13794
- Hierarchy Level | 13795
- Description | 13795
- Required Privilege Level | 13795
- Release Information | 13795

Syntax

```
protocols {
    fabric-control {
        graceful-restart {
            restart-time seconds;
            stale-routes-time seconds;
        }
    }
}
```

Hierarchy Level

```
[edit fabric]
```

Description

Specify attributes for the fabric control protocol.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X52-D10.

RELATED DOCUMENTATION

Understanding Routing Engines in the QFabric System

protocols (FIP)

IN THIS SECTION

- Syntax | 13796
- Hierarchy Level | 13796
- Description | 13796
- Options | 13796
- Required Privilege Level | 13796
- Release Information | 13797

Syntax

```
protocols {
    fip {
        fcoe-trusted;
        fc-map fc-map-value;
        fka-adv-period milliseconds;
        interface {
            interface-name {
                fka-adv-period milliseconds;
                priority priority;
            }
        }
        max-sessions-per-enode max-sessions-per-enode;
        priority priority;
        traceoptions {
            file filename <replace> <size size> <files number> <no-stamp>;
            <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
    }
}
```

Hierarchy Level

```
[edit fc-fabrics fc-fabric-name]
```

Description

Configure global or interface-specific FC protocol options. Individual interface settings override global settings.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

storage—To view this statement in the configuration.

storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fip | 19938

Configuring FIP on an FCoE-FC Gateway

Overview of FIP

protocols rsvp expand-flood-reflector-hop

IN THIS SECTION

- Syntax | 13797
- Description | 13797
- Required Privilege Level | 13797
- Release Information | 13798

Syntax

```
protocols rsvp expand-flood-reflector-hop
```

Description

Expands the flood reflector hops in the Explicit Route Objects (EROs) to enable the Level 1-Level 2 (L1-L2) routers to signal RSVP over ISIS non-forwarding adjacency.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2.

protocols (Security Zones Host Inbound Traffic)

IN THIS SECTION

- Syntax | 13798
- Hierarchy Level | 13798
- Description | 13798
- Options | 13799
- Required Privilege Level | 13800
- Release Information | 13800

Syntax

```
protocols {
    (protocol-name | all <protocol-name except>);
}
```

Hierarchy Level

```
[edit security zones security-zone zone-name host-inbound-traffic]
```

Description

Specify the types of protocol traffic that can reach the device for all interfaces in a zone. You can do this in one of several ways:

- You can enable traffic from each protocol individually.

- You can enable traffic from all protocols.
- You can enable traffic from all but some protocols.

Options

protocol-name Protocol for which traffic is allowed. The following protocols are supported:

- `all`—Enable traffic from all possible protocols available. Use the *except* option to disallow specific protocols.
- `bfd`—Enable incoming Bidirectional Forwarding Detection (BFD) protocol traffic.
- `bgp`—Enable incoming BGP traffic.
- `dvmrp`—Enable incoming Distance Vector Multicast Routing Protocol (DVMRP) traffic.
- `igmp`—Enable incoming Internet Group Management Protocol (IGMP) traffic.
- `ldp`—Enable incoming Label Distribution Protocol (LDP) traffic (UDP and TCP port 646).
- `msdp`—Enable incoming Multicast Source Discovery Protocol (MSDP) traffic.
- `nhrp`—Enable incoming Next Hop Resolution Protocol (NHRP) traffic.
- `ospf`—Enable incoming OSPF traffic.
- `ospf3`—Enable incoming OSPF version 3 traffic.
- `pgm`—Enable incoming Pragmatic General Multicast (PGM) protocol traffic (IP protocol number 113).
- `pim`—Enable incoming Protocol Independent Multicast (PIM) traffic.
- `rip`—Enable incoming RIP traffic.
- `ripng`—Enable incoming RIP next generation traffic.
- `router-discovery`—Enable incoming router discovery traffic.
- `rsvp`—Enable incoming Resource Reservation Protocol (RSVP) traffic (IP protocol number 46).
- `sap`— Enable incoming Session Announcement Protocol (SAP) traffic. SAP always listens on `224.2.127.254:9875`. New addresses and ports can be added dynamically. This information must be propagated to the Packet Forwarding Engine (PFE).
- `vrrp`—Enable incoming Virtual Router Redundancy Protocol (VRRP) traffic.

except (Optional) Disable specific incoming protocol traffic, but only when the *all* option has been defined . For example, to enable all but BGP and VRRP protocol traffic:

```
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust host-inbound-traffic protocols bgp except
set security zones security-zone trust host-inbound-traffic protocols vrrp except
```

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Zones Overview

Understanding Functional Zones

protocols (Security Zones Interfaces)

IN THIS SECTION

- Syntax | 13801
- Hierarchy Level | 13801
- Description | 13801
- Options | 13801
- Required Privilege Level | 13802
- Release Information | 13802

Syntax

```
protocols protocol-name {
               except;
}
```

Hierarchy Level

```
[edit security zones security-zone zone-name interfaces interface-name host-inbound-traffic]
```

Description

Specify the types of routing protocol traffic that can reach the device on a per-interface basis.

Options

- *protocol-name* —Protocol for which traffic is allowed. The following protocols are supported:
  - `all`—Enable traffic from all possible protocols available.
  - `bfd`—Enable incoming Bidirectional Forwarding Detection (BFD) Protocol traffic.
  - `bgp`—Enable incoming BGP traffic.
  - `dvmrp`—Enable incoming Distance Vector Multicast Routing Protocol (DVMRP) traffic.
  - `igmp`—Enable incoming Internet Group Management Protocol (IGMP) traffic.
  - `ldp`—Enable incoming Label Distribution Protocol (LDP) traffic (UDP and TCP port 646).
  - `msdp`—Enable incoming Multicast Source Discovery Protocol (MSDP) traffic.
  - `nhrp`—Enable incoming Next Hop Resolution Protocol (NHRP) traffic.
  - `ospf`—Enable incoming OSPF traffic.
  - `ospf3`—Enable incoming OSPF version 3 traffic.
  - `pgm`—Enable incoming Pragmatic General Multicast (PGM) protocol traffic (IP protocol number 113).
  - `pim`—Enable incoming Protocol Independent Multicast (PIM) traffic.

- `rip`—Enable incoming RIP traffic.
- `ripng`—Enable incoming RIP next generation traffic.
- `router-discovery`—Enable incoming router discovery traffic.
- `rsvp`—Enable incoming Resource Resolution Protocol (RSVP) traffic (IP protocol number 46).
- `sap`— Enable incoming Session Announcement Protocol (SAP) traffic. SAP always listens on `224.2.127.254:9875`.
- `vrrp`—Enable incoming Virtual Router Redundancy Protocol (VRRP) traffic.

`except`—(Optional) except can only be used if all has been defined.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Zones Overview

Understanding Functional Zones

protocols (STP Type)

IN THIS SECTION

- Syntax | 13803
- Hierarchy Level | 13803
- Description | 13803
- Options | 13803

Required Privilege Level | 13804
Release Information | 13804

Syntax

```
protocols {
    mstp { ... }
    rstp { ... }
    vstp { ... }
}
```

Hierarchy Level

```
[edit],
```

```
[edit logical-systems logical-system-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name],
```

```
[edit routing-instances routing-instance-name]
```

Description

Configure the Spanning Tree Protocol type as MSTP, RSTP, or VSTP.

Options

mstp—Configure the protocol as Multiple Spanning Tree.

rstp—Configure the protocol as Rapid Spanning Tree.

vstp—Configure the protocol as VLAN Spanning Tree.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

- *Configuring RSTP on EX Series Switches (CLI Procedure)*
- *Configuring MSTP on Switches*
- *Configuring MSTP Instances on a Physical Interface*
- *Configuring Rapid Spanning Tree Protocol*
- *Configuring Multiple Spanning Tree Protocol*
- Configuring VLAN Spanning Tree Protocol
- *Understanding MSTP*

protocols (Tenant Systems)

IN THIS SECTION

- Syntax | 13805
- Hierarchy Level | 13805
- Description | 13805
- Options | 13805
- Required Privilege Level | 13805
- Release Information | 13805

Syntax

```
protocols (bgp | ospf | ospf3);
```

Hierarchy Level

```
edit tenants tenant-name routing-instances instance_name
```

Description

Specify the types of routing protocol traffic that can reach a tenant system supported device.

Options

Tenant system supports the following routing protocols:

bgp Enable incoming BGP traffic.

ospf Enable incoming OSPF traffic.

ospf3 Enable incoming OSPF version 3 traffic.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Tenant Systems Overview

provider-tunnel

IN THIS SECTION

- Syntax | 13806
- Hierarchy Level | 13810
- Description | 13810
- Options | 13811
- Required Privilege Level | 13811
- Release Information | 13811

Syntax

```
provider-tunnel {
    external-controller pccd;
    family {
        inet {
            ingress-replication {
                create-new-ucast-tunnel;
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                }
            ldp-p2mp;
            mdt {
                data-mdt-reuse;
                group-range multicast-prefix;
                threshold {
                    group group-address {
                        source source-address {
                            rate threshold-rate;
                        }
                    }
                    }
                tunnel-limit limit;
            }
            pim-asm {
```

```
                group-address (Routing Instances) address;
            }
            pim-ssm {
                group-address (Routing Instances) address;
            }
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
                static-lsp lsp-name;
            }
        }
    inet6 {
        ingress-replication {
            create-new-ucast-tunnel;
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
            }
        ldp-p2mp;
        mdt {
            data-mdt-reuse;
            group-range multicast-prefix;
            threshold {
                group group-address {
                    source source-address {
                        rate threshold-rate;
                    }
                }
                }
            tunnel-limit limit;
        }
    }
        pim-asm {
            group-address (Routing Instances) address;
        }
        pim-ssm {
            group-address (Routing Instances) address;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
```

```
            static-lsp lsp-name;
        }
    ingress-replication {
        create-new-ucast-tunnel;
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
        }
    inter-as{
        ingress-replication {
            create-new-ucast-tunnel;
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
        }
        inter-region-segmented {
            fan-out| <leaf-AD routes>);
            threshold| <kilobits>);
        }
        ldp-p2mp;
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
        }
    }
    ldp-p2mp;
    pim-asm {
        group-address (Routing Instances) address;
    }
    pim-ssm {
        group-address (Routing Instances) address;
    }
    rsvp-te {
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
        static-lsp lsp-name;
    }
    selective {
        group multicast--prefix/prefix-length {
            source ip--prefix/prefix-length {
                ldp-p2mp;
```

```
                create-new-ucast-tunnel;
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                }
                pim-ssm {
                    group-range multicast-prefix;
                }
                rsvp-te {
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                    static-lsp point-to-multipoint-lsp-name;
                }
                threshold-rate kbps;
            }
            wildcard-source {
                pim-ssm {
                    group-range multicast-prefix;
                }
                rsvp-te {
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                    static-lsp point-to-multipoint-lsp-name;
                }
                threshold-rate kbps;
            }
        }
        tunnel-limit number;
        wildcard-group-inet {
            wildcard-source {
                pim-ssm {
                    group-range multicast-prefix;
                }
                rsvp-te {
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                    static-lsp lsp-name;
                }
                threshold-rate number;
            }
```

```
        }
        wildcard-group-inet6 {
            wildcard-source {
                pim-ssm {
                    group-range multicast-prefix;
                }
                rsvp-te {
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                    static-lsp lsp-name;
                }
                threshold-rate number;
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Configure virtual private LAN service (VPLS) flooding of unknown unicast, broadcast, and multicast traffic using point-to-multipoint LSPs. Also configure point-to-multipoint LSPs for MBGP MVPNs.

Starting in Junos OS Release 21.1R1, following provider tunnel types are supported on QFX10002, QFX10008, and QFX10016 Switches:

- Ingress Replication
- RSVP-TE P2MP LSP
- mLDP P2MP LSP

A point-to-multipoint (P2MP) is a MPLS LSP with a single source and multiple destinations. By taking advantage of MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the ingress router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.

Following are some of the properties of point-to-multipoint LSPs:

- A P2MP LSP enables use of MPLS for point-to-multipoint data distribution. This functionality is similar to that provided by the IP multicast.
- The branch LSPs can be added and removed without disruption traffic.
- A node can be configured as both a transit and egress router for different LPSs of the same point-to-multipoint LSP.
- LSPs can be configured statically or dynamic or as a combination of both static and dynamic LSPs.

P2MP LSPs are used to carry the IP unicast and multicast traffic.

Following tunnel types are not supported on QFX10002, QFX10008, and QFX10016 Switches:

- PIM-SSM tree
- PIM-SM tree
- PIM-Bidir tree
- mLDP MP2MP LSP

Options

external-controller *pccd*

(Optional) Specifies that point-to-multipoint LSP and (S,G) for MVPN can be provided by an external controller.

This option enables an external controller to dynamically configure (S,G) and point-to-multipoint LSP for MVPN. This is for only selective types. When not configured for a particular MVPN routing-instance, the external controller is not allowed to configure (S,G) and map point-to-multipoint LSP to that (S,G).

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.
The `selective` statement and substatements added in Junos OS Release 8.5.

The `ingress-replication` statement and substatements added in Junos OS Release 10.4.

In Junos OS Release 17.3R1, the `mdt` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default MDT in Rosen 7, and data MDT for Rosen 6 and Rosen 7. The `provider-tunnel mdt` hierarchy is now hidden for backward compatibility with existing scripts.

The `inter-as` statement and its substatements were added in Junos OS Release 19.1R1 to support next generation MVPN inter-AS option B.

`external-controller` option introduced in Junos OS Release 19.4R1 on all platforms.

RELATED DOCUMENTATION

Flooding Unknown Traffic Using Point-to-Multipoint LSPs in VPLS

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

providers

IN THIS SECTION

- Syntax | 13812
- Hierarchy Level | 13813
- Description | 13813
- Options | 13813
- Required Privilege Level | 13813
- Release Information | 13813

Syntax

```
providers {
    provider-id {
        license-type license deployment-scope [ deployments ];
```

```
    }
}
```

Hierarchy Level

```
[edit system extensions]
```

Description

Activate the certificate of the provider of the application and enable the PIC for loading of the application.

Options

provider-id Provider ID for the application package. The provider ID identifies the provider of the application to the system. The provider ID must be activated on the device to allow the application to be deployed on the device and run.

license-type Configure the license type and the scope of application deployment.

license Type of license. Obtain correct value from the application's provider.

deployment Scope of application deployment. You can configure a set of deployments. Obtain correct value from the application's provider.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`license-type` option introduced in Junos OS Release 10.2.

provisioning-method (Remote Device Management)

IN THIS SECTION

- Syntax | 13814
- Hierarchy Level | 13814
- Description | 13815
- Required Privilege Level | 13815
- Release Information | 13815

Syntax

```
provisioning-method {
    netconf {
        bulk-interval milliseconds;
        bulk-limit number;
        connection-retry-interval seconds;
        password password;
        port port-number;
        reconfigure-bulk-limit number;
        response-timeout seconds;
        response-timeout-count number;
        user-name name;
    }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name system services remote-device-management service-device device-name],
[edit system services remote-device-management service-device device-name]
```

Description

Configure the method for provisioning and deprovisioning services on the remote device. The NETCONF XML management protocol is currently the only supported method.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Configuring Remote Device Management for Service Provisioning

Remote Device Services Manager (RDSM) Overview

provisioning-order (Diameter Applications)

IN THIS SECTION

- Syntax | 13816
- Hierarchy Level | 13816
- Description | 13816
- Options | 13816
- Required Privilege Level | 13816
- Release Information | 13816

Syntax

```
provisioning-order (gx-plus | jsrc | pcrf);
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure AAA to use the specified application for subscriber service provisioning.

Options

gx-plus—Specify Gx-Plus as the application used to communicate with a PCRF server for subscriber service provisioning. Sets the Subscription-Id-Type Diameter AVP sub-attribute (450) to 4 (END_USER_PRIVATE) and sets the Subscription-Id-Data Diameter AVP sub-attribute (444) to `reserved`. Both of these sub-attributes are conveyed in the Diameter AVP Subscription-ID (443) by a CCR-I message.

`jsrc`—Specify JSRC as the application used to communicate with the SAE for subscriber service provisioning. JSRC is used in an SRC environment to request services from the SAE for an authenticated subscriber. JSRC attempts to activate these services. If successful, JSRC returns an ACK message. If unsuccessful, the subscriber is denied access.

pcrf—Specify Policy Control and Charging Rules Function (PCRF) as the application used to request provisioning from the PCRF server over the Gx protocol. If you change this configuration, any existing subscriber sessions are unaffected.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support for Gx-Plus introduced in Junos OS Release 11.2.

`pcrf` option added in Junos OS Release 16.2.

RELATED DOCUMENTATION

JSRC Configuration Overview

Provisioning Subscribers with JSRC

Configuring Gx-Plus

Provisioning Subscribers with Gx-Plus

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

proxy (interfaces)

IN THIS SECTION

- Syntax | **13817**
- Hierarchy Level | **13818**
- Description | **13818**
- Options | **13818**
- Required Privilege Level | **13818**
- Release Information | **13818**

Syntax

```
proxy inet-address address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family tcc],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family tcc]
```

Description

For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the IP address for which the TCC router is proxying. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only. Ethernet TCC is not supported on the T640 router.

Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms. See Configuring VLAN TCC Encapsulation for more details.

Options

`inet-address`—Configure the IP address of the neighbor to the TCC router.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
20.1R1	Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms.

RELATED DOCUMENTATION

- Configuring TCC Interface Switching
- Configuring TCC Interface Switching
- remote
- Junos OS VPNs Library for Routing Devices

proxy (Bridge Domains)

IN THIS SECTION

- Syntax | 13819
- Hierarchy Level | 13819
- Description | 13820
- Default | 13820
- Required Privilege Level | 13820
- Release Information | 13820

Syntax

```
proxy {
    source-address ip-address;
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
snooping],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols vlan
vlan-id igmp-snooping]
```

Description

Configure proxy mode and options, including source address. All the queries generated by IGMP snooping are sent using 0.0.0.0 as the source address in order to avoid participating in IGMP querier election. Also, all reports generated by IGMP snooping are sent with 0.0.0.0 as the source address unless there is a configured source address to use.

Default

By default, IGMP snooping does not employ proxy mode.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping

proxy-arp

IN THIS SECTION

- Syntax | 13821
- Hierarchy Level | 13821
- Description | 13821
- Default | 13821
- Options | 13821

Required Privilege Level | 13822
Release Information | 13822

Syntax

```
proxy-arp (restricted | unrestricted);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.

NOTE: You must configure the IP address and the `inet` family for the interface when you enable proxy ARP.

Default

Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.

Options

- `none`—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.
- `restricted`—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address.

- `unrestricted`—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.
- **Default:** `unrestricted`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`restricted` added in Junos OS Release 10.0 for EX Series switches.

RELATED DOCUMENTATION

- Configuring Restricted and Unrestricted Proxy ARP
- Configuring Proxy ARP on Switches
- Example: Configuring Proxy ARP on an EX Series Switch
- Configuring Gratuitous ARP

proxy-arp (Dynamic Profiles)

IN THIS SECTION

- Syntax | 13823
- Hierarchy Level | 13823
- Description | 13823
- Required Privilege Level | 13823
- Release Information | 13823

Syntax

```
proxy-arp;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

Description

For Ethernet interfaces only, configure the router to respond to any ARP request, as long as the router has an active route to the target address of the ARP request.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

- Configuring Restricted and Unrestricted Proxy ARP
- Configuring Gratuitous ARP

proxy-arp (Security NAT)

IN THIS SECTION

Syntax | 13824

- Hierarchy Level | 13824
- Description | 13824
- Options | 13824
- Required Privilege Level | 13824
- Release Information | 13825

Syntax

```
proxy-arp {
    interface interface-name {
        address ip-address {
            to ip-address;
        }
    }
}
```

Hierarchy Level

```
[edit security nat]
```

Description

Configure Address Resolution Protocol (ARP) proxy.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

proxy (Fibre Channel)

IN THIS SECTION

- Syntax | 13825
- Hierarchy Level | 13825
- Description | 13826
- Options | 13826
- Required Privilege Level | 13826

Syntax

```
proxy {
    auto-load-rebalance
    load-balance-algorithm (simple | enode-based | flog-based);
    no-fabric-wwn-verify;
    traceoptions {
        file filename <replace> <size size> <files number> <no-stamp>
        <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit fc-fabrics fabric-name]
```

Description

Configure proxy fabric operations.

Options

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding FCoE-FC Gateway Functions

proxy-generate

IN THIS SECTION

- Syntax | 13826
- Hierarchy Level | 13827
- Description | 13827
- Options | 13827
- Required Privilege Level | 13827
- Release Information | 13827

Syntax

```
proxy-generate <route-target-policy route-target-policy-name>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp group group-name family route-target],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
route-target],
[edit protocols bgp group group-name family route-target],
[edit protocols bgp group group-name neighbor address family route-target]
```

Description

Enable proxy BGP route target filtering (also known as proxy route target constrain, or proxy RTC). This feature is useful if you have a network environment where route target filtering is not widely deployed or fully supported. When configured for proxy BGP route target filtering, the device creates route target membership (RT membership) on behalf of its peers that do not have the route target filtering functionality. The device then distributes the RT membership advertisements from incoming BGP VPN routes to other devices in the network that need them.

Options

route-target-policy ***route-target-policy-name***	(Optional) Apply a routing policy that defines a subset of VPN routes to be used in your proxy BGP route target filter. The subset of VPN routes control which proxy BGP route targets are used to create the proxy BGP route target routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring Proxy BGP Route Target Filtering for VPNs

Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs

Configuring BGP Route Target Filtering for VPNs

family route-target | 6911
rtf-prefix-list

proxy (Interfaces)

IN THIS SECTION

- Syntax | 13828
- Hierarchy Level | 13828
- Description | 13828
- Options | 13829
- Required Privilege Level | 13829
- Release Information | 13829

Syntax

```
proxy inet-address address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family tcc],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family tcc]
```

Description

For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the IP address for which the TCC router is proxying. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only. Ethernet TCC is not supported on the T640 routing node.

Options

`inet-address` *address*—IP address for which the TCC router is acting as a proxy.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits

proxy-mac

IN THIS SECTION

- Syntax | 13830
- Hierarchy Level | 13830
- Description | 13830
- Options | 13830
- Required Privilege Level | 13830
- Release Information | 13831

Syntax

```
proxy-mac {
    irb |proxy-mac-address;
}
```

Hierarchy Level

```
[edit routing-instances name protocols evpn],
[edit routing-instances name bridge-domains]
```

Description

Specifies the proxy MAC address.

When CE devices connect directly to gateway devices, configure the MAC proxy on the gateway devices with the `irb` option. Junos OS will use the virtual gateway MAC address or the MAC address of the IRB interface.

In a spine-and-leaf network where the routing is performed on the spine devices, configure the MAC proxy on the leaf devices to use the actual MAC address of the IRB interface.

Options

irb Reply with the manually configured virtual gateway, automatically derived virtual gateway, or IRB MAC address. The virtual gateway MAC address takes precedence over a IRB MAC address. If the MAC address of the IRB interface is not configured, Junos OS uses a MAC address derived from the physical interface.

proxy-mac-address Reply with the configured proxy MAC address for all requests. A configured proxy MAC address should correspond to the virtual or physical gateway MAC address that is deployed on centralized Layer 3 gateway IRB interfaces.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

ARP and NDP Request with a proxy MAC address

proxy-macip-advertisement

IN THIS SECTION

- Syntax | 13831
- Hierarchy Level | 13831
- Description | 13831
- Required Privilege Level | 13832
- Release Information | 13833

Syntax

```
proxy-macip-advertisement;
```

Hierarchy Level

```
[edit interfaces irb unit logical-unit-number ]
```

Description

Enable the proxy advertisement feature on a QFX Series switch that can function as a Layer 3 (L3) gateway. With this feature enabled, the L3 gateway advertises the MAC and IP routes (EVPN Type 2 MAC-IP routes) on behalf of Layer 2 (L2) Virtual Extensible LAN (VXLAN) gateways.

In an Ethernet VPN-VXLAN (EVPN-VXLAN) centrally-routed bridging (CRB) overlay, spine devices typically function as L3 VXLAN gateways, and leaf devices typically function as L2 gateways. In this overlay network, the L2 VXLAN gateways can advertise only the EVPN Type 2 MAC routes for the attached hosts. Because the L2 gateways are unable to resolve the MAC-IP bindings for the hosts, each of the L3 gateways rely on the Address Resolution Protocol (ARP) and the Neighbor Discovery Protocol (NDP) to discover and install the bindings.

For example, an L3 gateway receives a host MAC route advertisement from an L2 gateway, and ARP and NDP resolve the MAC-IP bindings. The L3 gateway in turn advertises the host MAC and IP routes along with the next hop, which is set to the L2 gateway to which the host is attached. Upon receiving this advertisement, L2 and 3 gateways in the topology install the MAC-IP bindings along with the associated next hops. Then when a gateway device receives a packet with a destination MAC that matches an address in its MAC table, the gateway does the following:

- Checks the next hop associated with the MAC address
- Forwards the packet directly to the L2 gateway to which the host is attached.

As a result, the packet doesn't need to be sent first to an L3 gateway that then forwards the packet to the L2 gateway.

Enable this feature in a CRB overlay fabric. With the `proxy-macip-advertisement` setting, the L3 gateway device manages learning and aging for the ARP and NDP entries. By default, the L3 gateway does not refresh the ARP and NDP entries. As a result, the MAC-IP bindings on the spine devices expire after a timeout interval. Upon receiving the next IPv4 or IPv6 packet after the MAC-IP entry expires, the spine devices tries to resolve the MAC-IP binding again using ARP (for IPv4 packets) or NDP (for IPv6 packets). The spine device might drop traffic until it records the new MAC-IP binding.

> **NOTE**: We support the `proxy-macip-advertisement` statement on Junos OS devices only. In an EVPN CRB overlay fabric with a mix of Junos OS and Junos OS Evolved spine devices, if you configure any Junos OS spine devices with the `proxy-macip-advertisement` statement, you must also configure the EVPN leaf devices with the `crb-proxy-mac` option at the `[edit protocols l2-learning]` hierarchy level. See "crb-proxy-mac" on page 4704 for details.

We recommend that you don't use this feature in an EVPN-VXLAN edge-routed bridging overlay (ERB) fabric with anycast IRB addresses. If you set this option in ERB fabrics, the device can have IP reachability issues because remote ARP entries are deleted when the entries age out. With anycast IRB addresses, ARP replies don’t reach the device that initiated the ARP request.

Required Privilege Level

interface—To view this statement in the configuration

interface-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 15.1X53-D60.

proxy-mode

IN THIS SECTION

- Syntax | 13833
- Hierarchy Level | 13833
- Description | 13834
- Required Privilege Level | 13834
- Release Information | 13834

Syntax

```
proxy-mode;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name  overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay  overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay  overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name  overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay  overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
```

```
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Enable DHCP relay proxy mode on the extended DHCP relay. Proxy mode supports all extended DHCP relay functionality.

You cannot configure both the DHCP relay proxy and the extended DHCP local server on the same interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

DHCP Relay Proxy Overview

Extended DHCP Relay Agent Overview

Enabling DHCP Relay Proxy Mode

proxy (Multicast VLAN Registration)

IN THIS SECTION

- Syntax | 13835
- Hierarchy Level | 13835
- Description | 13835

- Default | 13836
- Options | 13836
- Required Privilege Level | 13836
- Release Information | 13836

Syntax

```
proxy source-address ip-address;
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan  (all | vlan-name)]
```

Description

Specify that a VLAN operate in IGMP snooping proxy mode.

On EX Series switches that do not use the Enhanced Layer 2 Software (ELS) configuration style, this statement is used only to set proxy mode for multicast VLAN registration (MVR) on a VLAN acting as a data-forwarding source (an MVLAN).

On ELS EX Series switches, this statement is available to enable IGMP snooping proxy mode either with or without MVR configuration. When you configure this option for a VLAN without MVR, the switch acts as an IGMP proxy to the multicast router for ports in that VLAN. When you configure this option with MVR on an MVLAN, the switch acts as an IGMP proxy between the multicast router and hosts in any MVR receiver VLANs associated with the MVLAN. This mode is configured on the MVLAN only, not on MVR receiver VLANs.

> **NOTE**: ELS switches also support *MVR proxy mode*, which is configured on individual MVR receiver VLANs associated with an MVLAN rather than on an MVLAN (unlike IGMP snooping proxy mode). To enable MVR proxy mode on an MVR receiver VLAN on ELS switches, use the "mode" on page 10809 statement with the proxy option.
>
> See *Understanding Multicast VLAN Registration* for details on MVR modes.

Default

Disabled

Options

`source-address` *`ip-address`*—IP address of the source VLAN to act as proxy.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring Multicast VLAN Registration on EX Series Switches Without ELS

Configuring Multicast VLAN Registration on EX Series Switches

mode (Multicast VLAN Registration) | **10809**

proxy-ndp (Security NAT)

IN THIS SECTION

- Syntax | **13837**
- Hierarchy Level | **13837**
- Description | **13837**
- Options | **13837**
- Required Privilege Level | **13837**
- Release Information | **13837**

Syntax

```
proxy-ndp {
    interface interface-name {
        address ip-address {
            to ip-address;
        }
    }
}
```

Hierarchy Level

```
[edit security nat]
```

Description

Configure Neighbor Discovery Protocol (NDP) proxy.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

proxy-profile

IN THIS SECTION

- Syntax | 13838
- Hierarchy Level | 13838
- Description | 13839
- Required Privilege Level | 13839
- Release Information | 13839

Syntax

```
proxy-profile proxy profile name
```

Hierarchy Level

```
[edit security utm default-configuration anti-virus avira-engine pattern-update],
[edit security utm default-configuration anti-virus sophos-engine pattern-update],
[edit security utm feature-profile anti-virus sophos-engine pattern-update],
[edit logical-systems name security utm default-configuration anti-virus avira-engine pattern-
update],
[edit logical-systems name security utm default-configuration anti-virus sophos-engine pattern-
update],
[edit logical-systems name security utm feature-profile anti-virus sophos-engine pattern-update],
[edit logical-systems name tenants name security utm default-configuration anti-virus avira-
engine pattern-update],
[edit logical-systems name tenants name security utm default-configuration anti-virus sophos-
engine pattern-update],
[edit logical-systems name tenants name security utm feature-profile anti-virus sophos-engine
pattern-update],
[edit tenants name security utm default-configuration anti-virus avira-engine pattern-update],
[edit tenants name security utm default-configuration anti-virus sophos-engine pattern-update],
```

```
[edit tenants name security utm feature-profile anti-virus sophos-engine pattern-update]
```

Description

Specify the proxy profile name and is used for configuring the explicit proxy.

For logical system and tenant you can configure the proxy profile only in root system at device level. You cannot configure the proxy profile in the logical system or tenant level.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

proxy (Security Antivirus)

IN THIS SECTION

- Syntax | 13840
- Hierarchy Level | 13840
- Description | 13840
- Options | 13840
- Required Privilege Level | 13840
- Release Information | 13840

Syntax

```
proxy {
    password password-string;
    port port-number;
    server address-or-url;
    username name;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update]
[edit security utm feature-profile anti-virus sophos-engine pattern-update]
```

Description

Update the pattern file on the proxy server.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 11.2.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

proxy (Services)

IN THIS SECTION

- Syntax | 13841
- Hierarchy Level | 13842
- Description | 13842
- Options | 13842
- Required Privilege Level | 13842
- Release Information | 13842

Syntax

```
proxy {
    global-config {
        session-cache-timeout seconds;
    }
    profile profile-name {
        actions {
            crl {
                disable;
                if-not-present (allow | drop);
                ignore-hold-instruction-code;
            }
            disable-session-resumption;
            ignore-server-auth-failure;
            logs {
                all;
                errors;
                info;
                sessions-allowed;
                sessions-dropped;
                sessions-ignored;
                sessions-whitelisted;
                warning;
            }
            renegotiation {
```

```
                (allow | allow-secure | drop);
            }
        }
        custom-ciphers [cipher];
        enable-flow-tracing;
         preferred-ciphers (custom | medium | strong | weak);
        root-ca root-certificate;
        trusted-ca (all | [ca-profile] );
        whitelist [global-address-book-addresses];
    }
}
```

Hierarchy Level

```
[edit services ssl]
```

Description

Specify the configuration for Secure Socket Layer (SSL) proxy support service.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. The `crl` statement is supported from 15.1X49-D30.

RELATED DOCUMENTATION

SSL Proxy Overview

Configuring SSL Forward Proxy

Enabling Debugging and Tracing for SSL Proxy

proxy (snmp)

IN THIS SECTION

- Syntax | 13843
- Hierarchy Level | 13844
- Description | 13844
- Options | 13844
- Required Privilege Level | 13845
- Release Information | 13845

Syntax

```
proxy proxy-name{
    device-name device-name;
    logical-system logical-system {
        routing-instance routing-instance;
    }
    routing-instance routing-instance;
    (version-v1 | version-v2c) {

        snmp-community community-name;

        no-default-comm-to-v3-config;
    }

    version-v3 {

        security-name security-name;

        context context-name;
```

```
        }
    }
```

Hierarchy Level

```
[edit snmp]
```

Description

Configure a device to act as a proxy SNMP agent, and specify a name for the proxy.

Options

context *context-name*

Specify the SNMPv3 context name as configured on the device specified at `edit snmp proxy` *`proxy-name`* `device-name` *`device-name`*.

For more information about this statement, see "context" on page 4615.

device-name *device-name*

Specify the name of the device to be managed through the proxy SNMP agent.

no-default-comm-to-v3-config

(Optional) Specify whether you have to manually configure the statements at the `[edit snmp v3 snmp-community` *`community-name`*`]` and `[edit snmp v3 vacm]` hierarchy levels.

If this statement is not included in the configuration, the `[edit snmp v3 snmp-community` *`community-name`*`]` and `[edit snmp v3 vacm]` hierarchy level configurations are automatically initialized.

proxy-name

Specify the name of the proxy.

security-name *security-name*

Specify the SNMPv3 security name as configured on the device specified at `edit snmp proxy` *`proxy-name`* `device-name` *`device-name`*.

For more information about this statement, see "security-name" on page 16501.

snmp-community *community-name*

Specify the name of the SNMP community. The community name you configure should match the `snmp-community` configuration on the device specified at `edit snmp proxy` *`proxy-name`* `device-name` *`device-name`*. For more information about this statement, see "snmp-community" on page 28065.

(version-v1 | version-v2c)

Specify the SNMP version, and add the relevant configuration.

version-v3 Add the SNMPv3 configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

proxy (System)

IN THIS SECTION

- Syntax | 13845
- Hierarchy Level | 13846
- Description | 13846
- Options | 13846
- Required Privilege Level | 13846
- Release Information | 13846

Syntax

```
proxy {
    server (hostname | ip-address);
    port port-number;
    username username;
    password password;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the proxy server properties for a device.

Options

server Configure the server by hostname or IP address.

port Set the port number for the proxy server ranging from 0 through 65535.

username Specify the user name configured in the proxy server.

password Specify the password associated with the `username` for the proxy server.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring a Proxy Server for License Updates

ps0 (Pseudowire Subscriber Interfaces)

IN THIS SECTION

- Syntax | 13847
- Hierarchy Level | 13847
- Description | 13847
- Required Privilege Level | 13847
- Release Information | 13848

Syntax

```
ps0 {
    anchor-point lt-device;
    mtu bytes;
    mac mac-address;
    no-gratuitous-arp-request;
    (flexible-vlan-tagging | stacked-vlan-tagging | untagged | vlan-tagging);
}
```

Hierarchy Level

```
[edit logical-systems transport-ls interfaces]
```

Description

Configure the pseudowire logical device.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

pseudowire-service (Pseudowire Subscriber Interfaces)

IN THIS SECTION

- Syntax | 13848
- Hierarchy Level | 13849
- Description | 13849
- Required Privilege Level | 13849
- Release Information | 13849

Syntax

```
pseudowire-service {
    device-count number;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure properties for the pseudowire devices on the router.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

MPLS Pseudowire Subscriber Logical Interfaces

pseudowire-status-tlv

IN THIS SECTION

- Syntax | 13850
- Hierarchy Level | 13850
- Description | 13850

Required Privilege Level | 13850
Release Information | 13850

Syntax

```
pseudowire-status-tlv hot-standby-vc-on;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit protocols l2circuit neighbor address interface interface-name]
```

Description

Enables the pseudowire type length variable (TLV). The pseudowire status TLV is used to communicate the status of a pseudowire back and forth between two PE routers. The pseudowire status TLV is configurable for each pseudowire connection and is disabled by default.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring the Pseudowire Status TLV

Example: Configuring Pseudowire Redundancy in a Mobile Backhaul Scenario

psn-tunnel-endpoint

IN THIS SECTION

- Syntax | 13851
- Hierarchy Level | 13851
- Description | 13851
- Options | 13852
- Required Privilege Level | 13852
- Release Information | 13852

Syntax

```
psn-tunnel-endpoint address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address
interface interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name backup-neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor address backup-neighbor address],
[edit protocols l2circuit neighbor address interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name backup-neighbor address],
[edit routing-instances routing-instance-name protocols vpls neighbor address backup-neighbor
address]
```

Description

Specify the endpoint of the packet switched network (PSN) tunnel on the remote PE router.

Options

address—Address for the tunnel endpoint.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.
Hierarchy levels associated with the `backup-neighbor` statement added in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Layer 2 Circuits over Both RSVP and LDP LSPs

Configuring Pseudowire Redundancy on the PE Router

psu

IN THIS SECTION

- Syntax | 13852
- Hierarchy Level | 13853
- Description | 13853
- Required Privilege Level | 13853
- Release Information | 13853

Syntax

```
psu {
    redundancy {
```

```
        n-plus-n (Power Management);
    }
}
```

Hierarchy Level

```
[edit chassis (EX Series)]
```

Description

Configure *N*+*N* power supply redundancy for power management on an EX6200 or EX8200 switch or SRX380 device.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Power Supply Redundancy (CLI Procedure)

psu (satellite-management)

IN THIS SECTION

- Syntax | 13854
- Hierarchy Level | 13854

- Description | 13854
- Required Privilege Level | 13854
- Release Information | 13854

Syntax

```
psu {
    redundancy {
        n-plus-n;
    }
}
```

Hierarchy Level

```
[edit chassis satellite-management]
```

Description

Configure *N*+*N* power supply redundancy for the satellite devices in a Junos Fusion.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

> *Understanding Power over Ethernet in a Junos Fusion*

ptp

IN THIS SECTION

- Syntax | 13855
- Syntax (MX10008) | 13856
- Syntax (ACX7100-32C, ACX7100-48L and ACX7024) | 13857
- Syntax (QFX Series) | 13857
- Syntax (QFX5120-48T) | 13858
- Syntax (QFX5130-48C) | 13859
- Syntax (PTX10004 and PTX10008) | 13859
- Hierarchy Level | 13860
- Description | 13861
- Options | 13861
- Required Privilege Level | 13864
- Release Information | 13864

Syntax

```
ptp {
    acquiring-state-announce-grandmaster;
    apts (domain domain | apply-groups | apply-groups-except)
    clock-mode(boundary | ordinary);
    disable-lag-revertive-switchover;
    domain domain;
    e2e-transparent;
    frequency-lock-threshold frequency-lock-threshold;
    holdover-in-spec-duration time duration;
    holdover-time-error-budget holdover-time-error-budget;
    ipv4-dscp ipv4-dscp;
    ipv6-dscp ipv6-dscp;
    local-priority local-priority;
    master {...}
    path-trace;
    performance-monitor {...}
```

```
    phase-adjust-threshold phase-adjust-threshold;
    phase-lock-threshold phase-lock-threshold;
    phy-timestamping;
    priority1 priority1;
    priority2 priority2;
    profile-type(aes67 | aes67-smpte | enterprise-profile | g.8275.1 | g.8275.1.enh | g.8275.2 |
g.8275.2.enh | ieee1588v2-custom | ieee-2008 | smpte);
    slave {...}
    stateful {...}
    syntonized-e2e-transparent;
    timescale arbitrary;
    unicast-negotiation;
    utc-leap-seconds utc-leap-seconds;
}
```

Syntax (MX10008)

```
ptp {
    clock-mode(boundary | ordinary);
    disable-lag-revertive-switchover;
    domain domain;
    e2e-transparent;
    frequency-lock-threshold frequency-lock-threshold;
    holdover-time-error-budget holdover-time-error-budget;
    ipv4-dscp ipv4-dscp;
    ipv6-dscp ipv6-dscp;
    local-priority local-priority;
    master {...}
    path-trace;
    performance-monitor {...}
    phase-adjust-threshold phase-adjust-threshold;
    phase-lock-threshold phase-lock-threshold;
    phy-timestamping;
    priority1 priority1;
    priority2 priority2;
    profile-type(aes67 | aes67-smpte | enterprise-profile | g.8275.1 | g.8275.1.enh | g.8275.2 |
g.8275.2.enh | ieee1588v2-custom | ieee-2008 | smpte);
    slave {...}
    stateful {...}
    unicast-negotiation;
```

```
    utc-leap-seconds utc-leap-seconds;
}
```

Syntax (ACX7100-32C, ACX7100-48L and ACX7024)

```
ptp {
    apts (domain domain | apply-groups | apply-groups-except)
    clock-mode(boundary | ordinary);
    disable-lag-revertive-switchover;
    domain domain;
    e2e-transparent;
    frequency-lock-threshold frequency-lock-threshold;
    holdover-in-spec-duration time duration;
    holdover-time-error-budget holdover-time-error-budget;
    ipv4-dscp ipv4-dscp;
    ipv6-dscp ipv6-dscp;
    local-priority local-priority;
    master {...}
    path-trace;
    performance-monitor {...}
    phase-adjust-threshold phase-adjust-threshold;
    phase-lock-threshold phase-lock-threshold;
    priority1 priority1;
    priority2 priority2;
    profile-type(enterprise-profile | g.8275.1 | g.8275.1.enh | g.8275.2.enh );
    slave {...}
    stateful {...}
    syntonized-e2e-transparent;
    timescale arbitrary;
    unicast-negotiation;
    utc-leap-seconds utc-leap-seconds;
}
```

Syntax (QFX Series)

```
ptp {
     clock-mode(boundary | ordinary);
     domain domain;
     ipv4-dscp ipv4-dscp;
```

```
    ipv6-dscp ipv6-dscp;
    local-priority local-priority;
    priority2 priority2;
    profile-type(aes67 | aes67-smpte | enterprise-profile | g.8275.2.enh | smpte);
    unicast-negotiation;
    slave {
        interface interface-name;
        announce-interval announce-interval;
        sync-interval sync-interval;
        delay-request delay-request;
        announce-timeout announce-timeout;
        grant-duration grant-duration;
    }
    master {
         interface interface-name;
         announce-interval announce-interval;
         sync-interval sync-interval;
         delay-req-timeout delay-req-timeout;
    }
}
```

Syntax (QFX5120-48T)

```
ptp {
    clock-mode(boundary | ordinary);
    domain domain;
    ipv4-dscp ipv4-dscp;
    ipv6-dscp ipv6-dscp;
    local-priority local-priority;
    priority2 priority2;
    profile-type(g.8275.2.enh);     unicast-negotiation;
    slave {
         interface interface-name;
         announce-interval announce-interval;
         sync-interval sync-interval;
         delay-request delay-request;
         announce-timeout announce-timeout;
         grant-duration grant-duration;
    }
    master {
         interface interface-name;
```

```
            announce-interval announce-interval;
            sync-interval sync-interval;
            delay-req-timeout delay-req-timeout;
        }
}
```

Syntax (QFX5130-48C)

```
ptp {
        clock-mode(boundary | ordinary);
        domain domain;
        e2e-transparent;
        ipv4-dscp ipv4-dscp;
        local-priority local-priority;
        priority2 priority2;
        profile-type(aes67 | aes67-smpte | enterprise-profile | smpte);
        slave {
            interface interface-name;
            announce-interval announce-interval;
            sync-interval sync-interval;
            delay-request delay-request;
            announce-timeout announce-timeout;
        }
        master {
            interface interface-name;
            announce-interval announce-interval;
            sync-interval sync-interval;
            delay-req-timeout delay-req-timeout;
        }
}
```

Syntax (PTX10004 and PTX10008)

```
ptp {
        acquiring-state-announce-grandmaster;
        clock-mode(boundary | ordinary);
        disable-lag-revertive-switchover;
        domain domain;
        e2e-transparent;
```

```
     holdover-time-error-budget holdover-time-error-budget;
     ipv4-dscp ipv4-dscp;
     ipv6-dscp ipv6-dscp;
     local-priority local-priority;
     master {
          announce-interval announce-interval;
     }
     phy-timestamping;
     priority1 priority1;
     priority2 priority2;
     profile-type(aes67 | aes67-smpte | enterprise-profile | g.8275.1 | g.8275.2 | g.8275.2.enh |
ieee-2008 | ieee1588v2-custom | smpte);
     slave {
          announce-interval announce-interval;
          announce-timeout announce-timeout;
     }
     syntonized-e2e-transparent;
     timescale arbitrary;
     unicast-negotiation;
     utc-leap-seconds utc-leap-seconds;
     performance-monitor passive-port {
          delay-request-rate delay-request-rate;
     }
     stateful interface interface-number {
          passive-port-monitor {
               passive-port-monitor;
               phase-threshold <offset-threshold-value>;
          }
         not-master;
     }
}
```

Hierarchy Level

```
[edit logical-systems name protocols],
[edit protocols]
```

Description

Configure PTP on the device. PTP, also known as IEEE 1588v2, is a packet-based technology that enables the operator to deliver synchronization services on packet-based mobile backhaul networks. IEEE 1588 PTP (Version 2) clock synchronization standard is a highly precise protocol for time synchronization that synchronizes clocks in a distributed system.

The system clocks can be categorized based on the role of the node in the network. They are broadly categorized into ordinary clocks and boundary clocks. The timeTransmitter clock and the timeReceiver clock are known as ordinary clocks. The boundary clock can operate as either a timeTransmitter clock or a timeReceiver clock.

You can configure the timeTransmitter clock and the timeReceiver clock for Precision Time Protocol (PTP) to help synchronize clocks in a distributed system. The time synchronization is achieved through packets that are transmitted and received in a session between the timeTransmitter clock and the timeReceiver clock.

Starting in Junos OS Evolved Release 22.2R1, with PTP timing profiles, such as G.8275.1 and G.8275.2, the accepted clock attribute `priority2` value is 255 for T-TSC clock. Any value other than 255 is rejected by the T-TSC clocks.

Starting in Junos OS Evolved Release 23.2R1, to support PTP timing profiles, such as G.8275.1 and G.8275.2, you can configure timing performance and monitoring for passive ports and stateful interfaces (only on PTX10008 devices).

Options

acquiring-state-announce-grandmaster
Use grandmaster data set for announce messages when in acquiring state.

clock-mode
Clock mode

- boundary—The clock mode of the node is a boundary clock where the clock acts as both timeTransmitter and timeReceiver.
- ordinary—The clock mode of the node is a system clock where the clock acts either as a timeTransmitter or as a timeReceiver. The ordinary option is not supported on ACX710 devices.

disable-lag-revertive-switchover
Disable revertive switchover for LAG.

domain
PTP domain number. This option is applicable to G.8275.1 profile.

- 0 through 127.

e2e-transparent

Enable end-to-end IEEE1588 transparent clock functionality.

frequency-lock-threshold

Control the frequency offset in parts per billion (ppb) that the servo will lock with. Review "frequency-lock-threshold" on page 7580.

> **NOTE**: This option is available only on supported platforms.

holdover-in-spec-duration

In case, both GNSS and the PTP timing reference are lost or inactive, then the internal clock oscillator provides synchronization for the specified duration.

After the holdover time expires, the clock class transition happens based on the mode of the system (T-GM or APTS). For example, in T-GM mode, when the holdover timer is set to 120 minutes, the system sends the clock class 7 for 120 minutes and thereafter clock class 160 is sent. Similarly, in T-BC-A mode, when the holdover timer is set to 120 minutes, the system sends the clock class 135 for 120 minutes and thereafter clock class 165 is sent.

> **NOTE**: This option is available only on supported platforms.

- Default: 240 (minutes)
- Range: 5 through 1440 (minutes)

holdover-time-error-budget

PTP clock holdover in specification time error budget in nano second (ns)

- Default: 400
- Range: 0 through 10000.

ipv4-dscp

IPv4 DiffServ code point (DSCP) value to be used for PTP packets

- 0 through 63.

ipv6-dscp

IPv6 DiffServ code point (DSCP) value to be used for PTP packets

- 0 through 63.

local-priority

Priority assigned to the local clock

- 1 through 255.

phase-adjust-threshold

Set the maximum phase offset that will be adjusted by the G.8275.2 enhanced servo in a phase-aligned state. Review "phase-adjust-threshold" on page 12469.

NOTE: This option is available only on supported platforms.

phase-lock-threshold

Control the maximum phase offset with which G.8275.2 enhanced servo enters the phase-aligned state. Review "phase-lock-threshold" on page 12471.

NOTE: This option is available only on supported platforms.

phy-timestamping

PHY timestamping feature.

priority1

Used in selecting best timeTransmitter clock

- 0 through 255.

priority2

Tie-breaker in selecting best timeTransmitter clock

- 0 through 255.

profile-type

PTP profile type

- aes67—Enable the AES67 PTP profile.
- aes67-smpte—Enable the AES67 and SMPTE ST-2059-2 PTP profile.
- enterprise-profile—Enable the enterprise profile. The enterprise profile supports IEEE 1588 PTPv2 transport over multicast IPv4. When the enterprise profile is enabled, no other profiles can be enabled. Also, unicast negotiation is disabled when you enable the enterprise profile.
- g.8275.1—Enable the G.8275.1 PTP profile.
- g.8275.1.enh—Enable the enhanced G.8275.1 PTP profile. This profile supports PTP over IPv4 and IPv6.
- g.8275.2—Enable the G.8275.2 PTP profile.
- g.8275.2.enh—Enable the enhanced G.8275.2 PTP profile.
- ieee-2008—Enable the IEEE 2008 profile type (default).
- ieee1588v2-custom—Enable the IEEE1588v2 with Multicast.
- smpte—Enable the SMPTE ST-2059-2 PTP Profile.

syntonized-e2e-transparent

Enable end-to-end syntonized IEEE1588 transparent clock functionality.

timescale	The timescale of the timeTransmitter. • arbitrary—TimeTransmitter clock with arbitrary timescale.
unicast-negotiation	Enable unicast negotiation.
utc-leap-seconds	UTC leap seconds offset • 0 through 255.
The following options are available only for PTX10008 devices: **performance-monitor** **Configure performance-monitor PTP packet delay metrics**	passive-port – Enable PTP performance monitoring for passive ports • delay-request-rate – Log mean interval between delay requests. Valid range is -4 to 4. This represents the packet rates as 2^(-4) that is (16 packets per second) to 2^4 (1 packets per 1 second). Default value is -4.
stateful interface	Stateful parameters to set interfaces to PTP stateful role. *<interface name>* – Interface name or port range. For example, et-0/0/4:0.0, et-0/0/4:1.0, et-2/0/4.0 passive-port-monitor – Configure passive port monitoring. • phase-threshold *<offset-threshold-value>* – Configure phase offset threshold value for raising alarm, relative to active timeReceiver port. Valid range is 20 to 100000000 nanoseconds. not-master – Configures port as timeReceiver or passive or listening port. measure-only – Configures port as passive or listening port.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing-To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 12.2R1

`frequency-lock-threshold`, `phase-adjust-threshold`, and `phase-lock-threshold` options introduced in Junos OS Release 23.4R1 for MX10008 routers.

`frequency-lock-threshold`, `phase-adjust-threshold`, and `phase-lock-threshold` options supported on ACX7100-32C, ACX7100-48L, and ACX7024 routers in Junos OS Evolved Release 24.2R1.

PTP profile-type options added to QFX5130-48C devices in Junos OS Evolved Release 24.2R1.

`apts` and `holdover-in-spec-duration` options introduced in Junos OS Evolved Release 24.2R2 for ACX7024.

RELATED DOCUMENTATION

Precision Time Protocol (PTP) Overview

PTP Passive Port Performance Monitoring on PTX10004 and PTX10008 Devices

show ptp passive-port-monitor-status | **23383**

ptp-mode

IN THIS SECTION

- Syntax | **13865**
- Hierarchy Level | **13866**
- Description | **13866**
- Required Privilege Level | **13866**
- Release Information | **13866**

Syntax

```
ptp-mode;
```

Hierarchy Level

```
[edit chassis fpc name pic name]
```

Description

Configure the PTP mode to enable PTP boundary or ordinary clock functionality.

When you enable PTP mode, one of the port that is used for data traffic gets disabled. For more information about valid port configurations on ACX7100-32C, ACX7100-48L, and ACX7024 routers, see Port Speed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 21.1R1.

ptsp-rules

IN THIS SECTION

- Syntax | 13867
- Hierarchy Level | 13867
- Description | 13867
- Options | 13867
- Required Privilege Level | 13867
- Release Information | 13867

Syntax

```
(ptsp-rules rule-name | ptsp-rules-sets rule-set-name);
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the PTSP rules or rule set included in this service set. You can configure multiple rules but only one rule set for each service.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

rule-set-name—Identifier for the set of rules to be included.

Required Privilege Level

interface—To view this statement in the configuration.

interface–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

> *Configuring Service Sets to be Applied to Services Interfaces*

pulse-per-second-enable

IN THIS SECTION

- Syntax | 13868
- Hierarchy Level | 13868
- Description | 13868
- Required Privilege Level | 13868
- Release Information | 13868

Syntax

```
pulse-per-second-enable;
```

Hierarchy Level

```
[edit chassis synchronization interfaces external]
[edit chassis synchronization interfaces (external-0/0 | external-1/0)]
```

Description

Configure the external interface to receive the pulse per second (PPS) signal on the GPS interface of the router.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

purge-flow-timer (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 13869
- Hierarchy Level | 13869
- Description | 13869
- Options | 13870
- Required Privilege Level | 13870
- Release Information | 13870

Syntax

```
purge-flow-timer seconds;
```

Hierarchy Level

```
[edit protocols openflow switch switch-name]
```

Description

For an OpenFlow virtual switch, specify the number of seconds after which an invalid OpenFlow flow entry is deleted from the flow table.

If you do not configure the `purge-flow-timer` statement, the device purges invalid flows from hardware, but indefinitely retains the corresponding flow entries in the flow table. If you configure the `purge-flow-timer` statement, the device purges invalid flows from hardware, and after the specified number of seconds, the device deletes the invalid flow entries from the flow table. Configuring a value of 0 causes the device to immediately delete invalid flow entries from the flow table.

> **NOTE**: By default, if you remove an active OpenFlow interface from an existing OpenFlow configuration, flow entries that match on this interface as the ingress interface and flow entries that include this interface in their action list (for OpenFlow v1.0) or flow instructions (for OpenFlow v1.3.1) are invalid and are automatically purged from the flow table and from the hardware regardless of whether you configure the `purge-flow-timer` statement.

Options

seconds Number of seconds after which an invalid flow entry is deleted from the flow table.

- **Range:** 0 through 300

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Support for OpenFlow on Devices Running Junos OS

Understanding OpenFlow Flow Entry Timers on Devices Running Junos OS

purge-timeout (programmable-rpd)

IN THIS SECTION

- Syntax | **13871**
- Hierarchy Level | **13871**

- Description | 13871
- Options | 13871
- Required Privilege Level | 13872
- Release Information | 13872

Syntax

```
purge-timeout {
    never;
    timeout <seconds>;
}
```

Hierarchy Level

```
[edit routing-options programmable-rpd]
```

Description

Set the time, in seconds, after which a disconnected client times out. Upon disconnect, the client state remains available but no operations occur. If the disconnected client reconnects before the set time has elapsed, the states are restored on the router. If it does not, all client operations are reverted and the programmable routing protocol process (prpd) server notifies any other modules of the disconnect so they can purge any other client operations.

Options

never When this option is configured, the purge timeout is infinite for the BGP route service. Client-added routes never time out. The prpd-client-added routes are not deleted when the client disconnects and does not reconnect back. The routes are deleted only when the client explicitly deletes the routes. If routing is restarted, then the client-added routes are lost.

timeout *seconds* (Optional) After this time, in seconds, disconnected clients time out on the prpd server and the routes added by the client are purged.

- **Default:** 120
- **Range (Junos OS Releases 18.4R1 and later):** 1 through 604,800 (7 days)
- **Range (Before Junos OS Releases 18.4R1):** 1 through 1000

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

`never` option introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

traceoptions (programmable-rpd) | **29651**
show programmable-rpd clients | **23287**
show route | **23512**

purging

IN THIS SECTION

- Syntax | **13873**
- Hierarchy Level | **13873**
- Description | **13873**
- Required Privilege Level | **13873**
- Release Information | **13873**

Syntax

```
purging;
```

Hierarchy Level

```
[edit system arp]
```

Description

Purge obsolete ARP entries from the cache when an interface or link goes offline.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

push (protocols mpls)

IN THIS SECTION

- Syntax | 13874
- Hierarchy Level | 13874
- Description | 13874
- Options | 13874
- Required Privilege Level | 13874
- Release Information | 13874

Syntax

```
push out-label;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
bypass],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit protocols mpls static-label-switched-path lsp-name bypass],
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Description

Add a new label to the top of the label stack. This statement is used to configure static LSPs at ingress routers and to configure bypass LSPs for static LSPs.

Options

out-label—Manually assigned outgoing label value.

- **Range:** 0 through 1,048,575.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

pop (protocols mpls) | **12797**

swap (protocols mpls) | **28738**

push

IN THIS SECTION

- Syntax | 13875
- Hierarchy Level | 13875
- Description | 13875
- Required Privilege Level | 13876
- Release Information | 13876

Syntax

```
push;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

Description

NOTE: On EX4300 switches, push is not supported at the [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map] hierarchy level.

Specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.

You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.

If you include the `push` statement in the configuration, you must also include the "pop" on page 12799 statement at the `[edit interfaces interface-name unit logical-unit-number output-vlan-map]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Stacking a VLAN Tag

Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support

push-backup-to-master (Accounting Options)

IN THIS SECTION

- Syntax | 13877
- Hierarchy Level | 13877
- Description | 13877
- Required Privilege Level | 13877
- Release Information | 13877

Syntax

```
push-backup-to-master;
```

Hierarchy Level

[edit accounting-options file *filename*]

Description

Configure the router to save the accounting files from the new backup Routing Engine to the new primary Routing Engine when a change in primary role occurs. The files are saved to the **/var/log/pfedBackup** directory on the router. The primary Routing Engine includes these accounting files with its own current accounting files when it transfers the files from the backup directory to the archive site at the next transfer interval. Use this statement when the new backup Routing Engine is not able to connect to the archive site; for example, when site is not connected by means of an out-of-band interface or the path to the site is routed through a line card.

NOTE: The backup Routing Engine's files on the primary Routing Engine are sent at each interval even though the files remain the same. If this is more activity than you want, consider using the `backup-on-failure master-and-slave` statement instead.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

- Configuring Accounting-Data Log Files
- *Configuring Flat-File Accounting for Layer 2 Wholesale*
- *Configuring Flat-File Accounting for Extensible Subscriber Services Management*
- *Flat-File Accounting Overview*

push (Dynamic VLANs)

IN THIS SECTION

- Syntax | 13878
- Hierarchy Level | 13878
- Description | 13878
- Required Privilege Level | 13878
- Release Information | 13879

Syntax

```
push;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number input-vlan-map]
```

Description

For dynamic VLAN interfaces, specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag. If you include the `push` statement in the configuration, you must also include the pop statement at the `[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution

push-push

IN THIS SECTION

- Syntax | 13879
- Hierarchy Level | 13879
- Description | 13880
- Required Privilege Level | 13880
- Release Information | 13880

Syntax

```
push-push;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

Description

Specify the VLAN rewrite operation to push two VLAN tags in front of the frame.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Stacking Two VLAN Tags

push-to-identity-management

IN THIS SECTION

- Syntax | 13881
- Hierarchy Level | 13881
- Description | 13881
- Required Privilege Level | 13881
- Release Information | 13881

Syntax

```
push-to-identity-management;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication]
```

Description

Configure firewall authentication to push authentication entries with a successful authentication state to the Juniper Identity Management Service server. You use this statement in conjunction with the query-api/user statement, which sets the path of the URL for querying user identities.

When the SRX Series Firewall does not have authentication information for a user based on the user's IP address, it can force the user to authenticate through captive portal to obtain the user ID information and authenticate the user. If a security policy that specifies firewall authentication is configured with the `push-to-identity-management` statement, the user information is pushed to the Juniper Identity Management Service server.

After you push the entry to the Juniper Identity Management Service server, you can use the batch query function to obtain authentication information for that user from the Juniper Identity Management Service server, including the groups that the user belongs to.

> **NOTE**: The SRX Series Firewall does not update the authentication-entry time-out state to Juniper Identity Management Service.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Understanding Advanced Query Feature for Obtaining User Identity Information from JIMS

filter (Identity Management Advanced Query) | 7138

primary connection (Identity Management Advanced Query) | 13131

secondary connection (Identity Management Advanced Query) | 16409

pvlan

IN THIS SECTION

- Syntax | 13882
- Hierarchy Level | 13882
- Description | 13882
- Options | 13883
- Required Privilege Level | 13883
- Release Information | 13883

Syntax

```
pvlan;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Description

Specify that the VLAN is private and access ports in the VLAN do not forward packets to each other. You use this statement with primary VLANs and isolated secondary VLANs.

Options

none

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Creating a Private VLAN on a Single QFX Switch without ELS Support

Creating a Private VLAN Spanning Multiple QFX Series Switches without ELS Support

pvlan-trunk

IN THIS SECTION

- Syntax | 13883
- Hierarchy Level | 13884
- Description | 13884
- Required Privilege Level | 13884
- Release Information | 13884

Syntax

```
pvlan-trunk;
```

Hierarchy Level

```
[edit vlans  vlan-name vlan-id number interface interface-name]
```

Description

Configure an interface to be the trunk port, connecting switches that are configured with a private VLAN (PVLAN) across these switches.

Required Privilege Level

routing—To view this statement in the configuration.

routing–control—To add this statement to the configuration.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)

Creating a Private VLAN on a Single QFX Switch

Creating a Private VLAN Spanning Multiple QFX Series Switches

python-script-user (Event and SNMP Scripts)

IN THIS SECTION

Syntax | 13885

Hierarchy Level | 13885

- Description | 13885
- Default | 13885
- Options | 13886
- Required Privilege Level | 13886
- Release Information | 13886

Syntax

```
python-script-user username;
```

Hierarchy Level

```
[edit event-options event-script file filename],
[edit system scripts snmp file filename]
```

Description

Define the user under whose access privileges the Python event or SNMP script executes. If you do not associate a user with the script, then by default, Junos OS executes the script with the privileges of the *nix user and group `nobody`, which is the generic, unprivileged system account.

> **NOTE**: You cannot configure Python event and SNMP scripts to execute with root access privileges.

> **NOTE**: Junos OS executes Python commit and op scripts under the access privileges of the user who executes the command or operation that invokes the script.

Default

If you do not include the `python-script-user` statement, Junos OS executes the script under the access privileges of the *nix user and group `nobody`.

Options

username Username that is configured at the `[edit system login]` hierarchy level.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R3.

RELATED DOCUMENTATION

Requirements for Executing Python Automation Scripts on Devices Running Junos OS

Understanding Python Automation Scripts for Devices Running Junos OS

Q

IN THIS SECTION

- qos-adjust | 13888
- qos-adjust-adsl | 13890
- qos-adjust-adsl2 | 13892
- qos-adjust-adsl2-plus | 13894
- qos-adjust-hierarchical | 13896
- qos-adjust-other | 13897
- qos-adjust-sdsl | 13899
- qos-adjust-vdsl | 13901
- qos-adjust-vdsl2 | 13903
- qualified-bum-pruning-mode | 13905
- qualified-next-hop (Access) | 13906

- qualified-next-hop (Access-Internal) | **13907**
- qualified-next-hop (Dynamic Access-Internal Routes) | **13908**
- qualified-next-hop (Static Routes) | **13910**
- qualified-vlan | **13912**
- quality (Interfaces) | **13914**
- quality-level | **13915**
- quality-level (Clock Synchronization) | **13917**
- quality-level (hybrid) | **13919**
- quality-mode-enable (MX) | **13920**
- quality-mode-enable (PTX) | **13922**
- quarantine-message (Security Content Security) | **13924**
- querier (performance-monitoring) | **13925**
- query-interval (Bridge Domains) | **13927**
- query-interval (Protocols IGMP) | **13929**
- query-interval (Protocols IGMP AMT) | **13930**
- query-interval (Protocols MLD) | **13932**
- query-last-member-interval (Bridge Domains) | **13933**
- query-last-member-interval (Protocols IGMP) | **13935**
- query-last-member-interval (Protocols MLD) | **13936**
- query-response-interval (Bridge Domains) | **13938**
- query-response-interval (Protocols IGMP) | **13940**
- query-response-interval (Protocols IGMP AMT) | **13941**
- query-response-interval (Protocols MLD) | **13943**
- queue (Global Queues) | **13944**
- queue-monitoring | **13946**
- queue-num | **13947**
- queue (Restricted Queues) | **13949**
- queue-statistics | **13950**
- queue-threshold | **13953**
- queues | **13955**
- quit | **13956**

qos-adjust

IN THIS SECTION

- For Junos OS Release 17.3 and Earlier Releases | 13888
- For Junos OS Release 17.4 and Later Releases | 13888
- Hierarchy Level | 13889
- Description | 13889
- Required Privilege Level | 13889
- Release Information | 13889

For Junos OS Release 17.3 and Earlier Releases

```
qos-adjust {
    adsl-bytes bytes;
    adsl2-bytes bytes;
    adsl2-plus-bytes bytes;
    other-bytes;
    other-overhead-adjust;
    sdsl-bytes bytes;
    sdsl-overhead-adjust percentage;
    vdsl-bytes bytes;
    vdsl-overhead-adjust percentage;
    vdsl2-bytes bytes;
    vdsl2-overhead-adjust percentage;
}
```

For Junos OS Release 17.4 and Later Releases

```
qos-adjust;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Specify that the ANCP agent reports data rates for downstream traffic to CoS. When this statement is not configured, the ANCP agent does not report traffic rates to CoS.

For Junos OS Release 17.3 and earlier releases, configure the values by which the actual downstream data rates are adjusted. The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Starting in Junos OS Release 17.4, all of the previously supported subordinate statements are deprecated. The ANCP agent ignores them if they are present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustments with the `access-line` statement at the `[edit system]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Subordinate statements deprecated in Junos OS Release 17.4.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS

Traffic Rate Reporting and Adjustment by the ANCP Agent

Configuring the ANCP Agent

qos-adjust-adsl

IN THIS SECTION

- Syntax | 13890
- Hierarchy Level | 13890
- Description | 13890
- Options | 13891
- Required Privilege Level | 13891
- Release Information | 13891

Syntax

```
qos-adjust-adsl adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on an ADSL line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-adsl` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

> *Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates*
>
> *Configuring the ANCP Agent*

qos-adjust-adsl2

IN THIS SECTION

- Syntax | 13892
- Hierarchy Level | 13892
- Description | 13892
- Options | 13893
- Required Privilege Level | 13893
- Release Information | 13893

Syntax

```
qos-adjust-adsl2 adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on an ADSL2 line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-adsl2` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates

Configuring the ANCP Agent

qos-adjust-adsl2-plus

IN THIS SECTION

- Syntax | 13894
- Hierarchy Level | 13894
- Description | 13894
- Options | 13895
- Required Privilege Level | 13895
- Release Information | 13895

Syntax

```
qos-adjust-adsl2-plus adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on an ADSL2+ line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-plus-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-adsl2-plus` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `adsl2-plus-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

> *Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates*
>
> *Configuring the ANCP Agent*

qos-adjust-hierarchical

IN THIS SECTION

- Syntax | 13896
- Hierarchy Level | 13896
- Description | 13896
- Options | 13896
- Required Privilege Level | 13897
- Release Information | 13897

Syntax

```
qos-adjust-hierarchical [interface-set];
```

Hierarchy Level

```
[edit dynamic-profiles dynamic-profile-name access-cac interface $junos-interface-name]
```

Description

Specify that multicast QoS adjustments received on the subscriber contribute to hierarchical QoS adjustments.

Options

`interface-set`—(Optional) Specify that the immediate parent interface-set receive QoS adjustments.

NOTE: For any multicast-enabled interface, only directly stacked interface-set parents are adjusted. The following scenario is not supported: multicast-enabled client interface is stacked over a VLAN interface and the VLAN interface in turn is stacked over and interface-set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers

qos-adjust-other

IN THIS SECTION

- Syntax | 13897
- Hierarchy Level | 13897
- Description | 13898
- Options | 13898
- Required Privilege Level | 13898
- Release Information | 13898

Syntax

```
qos-adjust-other adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on an access line of DSL type OTHER. The ANCP agent reports the adjusted rate only to AAA.

The router reports some access technology types—such as Gigabit passive optical network (GPON) lines—as DSL type OTHER.

> **NOTE**: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-other` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

`adjustment-factor`—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `other-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates

Configuring the ANCP Agent

qos-adjust-sdsl

IN THIS SECTION

- Syntax | 13899
- Hierarchy Level | 13899
- Description | 13899
- Options | 13900
- Required Privilege Level | 13900
- Release Information | 13900

Syntax

```
qos-adjust-sdsl adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on an SDSL line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-sdsl` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

`adjustment-factor`—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates

Configuring the ANCP Agent

qos-adjust-vdsl

IN THIS SECTION

- Syntax | 13901
- Hierarchy Level | 13901
- Description | 13901
- Options | 13902
- Required Privilege Level | 13902
- Release Information | 13902

Syntax

```
qos-adjust-vdsl adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on a VDSL line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-vdsl` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

> *Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates*
>
> *Configuring the ANCP Agent*

qos-adjust-vdsl2

IN THIS SECTION

- Syntax | 13903
- Hierarchy Level | 13903
- Description | 13903
- Options | 13904
- Required Privilege Level | 13904
- Release Information | 13904

Syntax

```
qos-adjust-vdsl2 adjustment-factor;
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Configure an adjustment factor that is applied globally to the downstream and upstream data rates reported by the ANCP agent for all subscribers on a VDSL2 line. The ANCP agent reports the adjusted rate only to AAA.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `qos-adjust-vdsl2` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

adjustment-factor—Adjustment factor applied to upstream and downstream data rates for the DSL type.

- **Range:** 0 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-total-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

> *Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates*
>
> *Configuring the ANCP Agent*

qualified-bum-pruning-mode

IN THIS SECTION

- Syntax | 13905
- Hierarchy Level | 13905
- Description | 13905
- Required Privilege Level | 13905
- Release Information | 13905

Syntax

```
qualified-bum-pruning-mode;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

For Junos OS Layer 2 Wholesale configurations, prune (constrain) distribution of broadcast, unicast, and multicast (BUM) packets of unknown origin to only those interfaces that match the traffic from a specific VLAN pair.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers
- Layer 2 and Layer 3 Wholesale Overview
- *Support of Inner VLAN List and Inner VLAN Range for Qualified BUM Pruning on a Dual-Tagged Interface for a VPLS Routing Instance Overview*

qualified-next-hop (Access)

IN THIS SECTION

- Syntax | 13906
- Hierarchy Level | 13906
- Description | 13906
- Options | 13907
- Required Privilege Level | 13907
- Release Information | 13907

Syntax

```
qualified-next-hop next-hop;
```

Hierarchy Level

```
[edit routing-options access route ip-prefix</prefix-length>]
```

Description

Configure the qualified next-hop address for an access route.

Options

next-hop—Specific qualified next-hop address you want to assign to the access route.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

qualified-next-hop (Access-Internal)

IN THIS SECTION

- Syntax | 13907
- Hierarchy Level | 13907
- Description | 13908
- Options | 13908
- Required Privilege Level | 13908
- Release Information | 13908

Syntax

```
qualified-next-hop next-hop;
```

Hierarchy Level

```
[edit routing-options access-internal route ip-prefix</prefix-length>]
```

Description

Configure the qualified next-hop address for an internal access route.

Options

next-hop—Specific qualified next-hop address you want to assign to the internal access route.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

qualified-next-hop (Dynamic Access-Internal Routes)

IN THIS SECTION

- Syntax | 13908
- Hierarchy Level | 13909
- Description | 13909
- Options | 13909
- Required Privilege Level | 13909
- Release Information | 13909

Syntax

```
qualified-next-hop interface-name {
    mac-address address;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
access-internal route subscriber-ip-address],
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
rib routing-table-name access-internal route subscriber-ip-address],
[edit dynamic-profiles profile-name routing-options access-internal route subscriber-ip-address
qualified-next-hop underlying-interface].
```

Description

Dynamically configure the qualified next-hop and the MAC address for an access-internal route for DHCP and PPP subscriber interfaces.

Options

interface-name—Either the specific interface you want to assign to the access route or the variable, or the `$junos-interface-name` variable. The variable is dynamically replaced with the value supplied by DHCP or PPP when a subscriber logs in.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options route *subscriber-ip-address*] and [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* route *subscriber-ip-address*] hierarchy levels introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Routes for DHCP and PPP Subscriber Access Networks

qualified-next-hop (Static Routes)

IN THIS SECTION

- Syntax | 13910
- Hierarchy Level | 13911
- Description | 13911
- Options | 13911
- Required Privilege Level | 13912
- Release Information | 13912

Syntax

```
qualified-next-hop (address | interface-name) {
    bfd-liveness-detection {
        authentication {
            algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 | meticulous-keyed-sha-1 |
simple-password);
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        holddown-interval milliseconds;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
```

```
        }
        version (1 | automatic);
    }
    interface interface-name;
    metric metric;
    preference preference;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options
static route destination-prefix],
[edit logical-systems logical-system-name routing-options rib inet6.0 static route destination-
prefix],
[edit logical-systems logical-system-name routing-options static route destination-prefix],
[edit routing-instances routing-instance-name routing-options static route destination-prefix],
[edit routing-options rib inet6.0 static route destination-prefix],
[edit routing-options static route destination-prefix]
```

Description

Configure a static route with multiple possible next hops, each of which can have its own preference value, IGP metric that is used when the route is exported into an IGP, and Bidirectional Forwarding Detection (BFD) settings. If multiple links are operational, the one with the most preferred next hop is used. The most preferred next hop is the one with the lowest preference value.

Options

address—IPv4, IPv6, or ISO network address of the next hop.

interface-name—Name of the interface on which to configure an independent metric or preference for a static route. To configure an unnumbered interface as the next-hop interface for a static route, specify **qualified-next-hop *interface-name***, where ***interface-name*** is the name of the IPv4 or IPv6 unnumbered interface.

> **NOTE**: For an Ethernet interface to be configured as the qualified next hop for a static route, it must be an unnumbered interface.

To configure an Ethernet interface as an unnumbered interface, configure the `unnumbered-address <interface-name>` statement at the `[edit interfaces <interface-name> unit <logical-unit-number> family <family-name>]` hierarchy level as described in *Configuring an Unnumbered Interface*.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection

qualified-vlan

IN THIS SECTION

- Syntax | 13913
- Hierarchy Level | 13913
- Description | 13913
- Options | 13913
- Required Privilege Level | 13913
- Release Information | 13913

Syntax

```
qualified-vlan vlan-id;
```

Hierarchy Level

```
[edit protocols mld-snooping vlan vlan-name]
[edit routing-instances instance-name protocols mld-snooping vlan vlan-name]
[edit protocols igmp-snooping vlan vlan-name]
```

Description

Configure VLAN options for qualified learning.

Options

vlan-id—VLAN ID of the learning domain.

- **Range:** 0 through 1023

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping on SRX Series Devices

Configuring MLD Snooping on a Switch VLAN with ELS Support (CLI Procedure)

Configuring IGMP Snooping on Switches

IGMP Snooping Overview

igmp-snooping | **8387**

quality (Interfaces)

IN THIS SECTION

- Syntax | 13914
- Hierarchy Level | 13914
- Description | 13914
- Required Privilege Level | 13914
- Release Information | 13914

Syntax

```
quality <value>;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number radio-router]
```

Description

This option controls relative link quality weight (value 0–100).

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

quality-level

IN THIS SECTION

- Syntax | 13915
- Hierarchy Level | 13915
- Description | 13915
- Options | 13916
- Required Privilege Level | 13916
- Release Information | 13916

Syntax

```
quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
```

Hierarchy Level

```
[edit chassis synchronization source (bits | gps | interfaces interface-name)]
```

Description

For ACX Series routers operating with Synchronous Ethernet, configure the quality level (QL) for a timing source so that the router can know the best available source with which to synchronize. The QL specifies the expected accuracy level of the clock and is transmitted across the network through Synchronization Status Messages (SSMs) over the Ethernet Synchronization Message Channel (ESMC).

The quality level options are valid depending on the configured **network-option**; **option-1** or **option-2** at the [**edit chassis synchronization**] hierarchy level.

Both **option-1** and **option-2** SSM quality levels (QL) are supported:

- For **option-1**, QL must be configured for external clocks (**gps** or **bits**) whether or not QL is enabled.
- For **option-2**, the default QL for external clocks is QL_STU whether or not QL is enabled.

Options

prc—Timing quality of a primary reference clock (option-1 only).

prs—Clock traceable to a primary reference source (option-2 only).

sec—Timing quality of an SDH equipment clock (option-1 only).

smc— Clock traceable to a self-timed SONET (option-2 only).

ssu-a—Timing quality of a type I or IV client clock (option-1 only).

ssu-b—Timing quality of a type VI client clock (option-1 only).

st2—Clock traceable to Stratum 2 (option-2 only).

st3—Clock traceable to Stratum 3 (option-2 only).

st3e—Clock traceable to Stratum 3E (option-2 only).

st4—Clock traceable to Stratum 4 free-run (option-2 only).

stu—Clock traceable to an unknown quality (option-2 only).

tnc—Clock traceable to a transit node clock (option-2 only).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

External Clock Synchronization Overview for ACX Series Routers

Configuring External Clock Synchronization for ACX Series Routers

Synchronous Ethernet Overview

show chassis synchronization

quality-level (Clock Synchronization)

IN THIS SECTION

- Syntax | 13917
- Hierarchy Level | 13917
- Description | 13917
- Options | 13918
- Required Privilege Level | 13918
- Release Information | 13918

Syntax

```
quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
```

Hierarchy Level

```
[edit chassis synchronization source interfaces (external-a | external-b | interface interface-name)]
[edit chassis synchronization source interfaces (external | interface interface-name)]
[edit chassis synchronization source interfaces (external-0/0 | external-1/0 | interface interface-name)]
```

Description

Configure the quality level for a timing source so that the router knows the best available source with which to synchronize. The quality level specifies the accuracy level of the clock and is transmitted across the network through Synchronization Status Messages (SSMs) over the Ethernet Synchronization Messaging Channel (ESMC), or through SSMs contained in SONET/SDH frames.

SONET SSM messages are either Generation 1 or Generation 2. Generation 1 is the first and most widely deployed SSM message set. Generation 2 is a newer version. Quality level options are available for both Generation 1 and Generation 2.

Options

prs—Primary reference source—Stratum 1

st2—Stratum 2

tnc—Transit node clock

st3e—Stratum 3E

st3—Stratum 3

smc—SONET minimum clock

st4—Stratum 4

prc—Primary reference clock

ssu-a—Synchronization supply unit A

ssu-b—Synchronization supply unit B

sec—SDH equipment clock

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

quality-level (hybrid)

IN THIS SECTION

- Syntax | 13919
- Hierarchy Level | 13919
- Description | 13919
- Options | 13919
- Required Privilege Level | 13920
- Release Information | 13920

Syntax

```
quality-level (prs | st2 | tnc | st3e | st3 | smc | st4) | (prc | ssu-a | ssu-b | sec);
```

Hierarchy Level

```
[edit protocols ptp slave clock-class-to-quality-level-mapping]
```

Description

Configure the quality level for the timing source so that the router knows the best available source with which to synchronize. The quality level specifies the accuracy level of the clock.

Options

`prs` —Primary reference source—Stratum 1

`st2`—Stratum 2

`tnc`—Transit node clock

`st3e`—Stratum 3E

`st3`—Stratum 3

`smc`—SONET minimum clock

`st4`—Stratum 4

`prc`—Primary reference clock

`ssu-a`—Synchronization supply unit A

`ssu-b`—Synchronization supply unit B

`sec`—SDH equipment clock

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2R2.

RELATED DOCUMENTATION

Configuring Hybrid Mode and ESMC Quality Level Mapping

Example: Configuring Hybrid Mode and ESMC Quality Level Mapping

Understanding Hybrid Mode

Precision Time Protocol Overview

Synchronous Ethernet Overview

quality-mode-enable (MX)

IN THIS SECTION

- Syntax | 13921
- Hierarchy Level | 13921
- Description | 13921

- Default | 13922
- Required Privilege Level | 13922
- Release Information | 13922

Syntax

```
quality-mode-enable;
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Enable Synchronous Ethernet ESMC quality mode. The quality level parameter for a Synchronous Ethernet interface is optional when the `quality-mode-enable` and the `selection-mode received-quality` statements are included at the `[edit chassis synchronization]` hierarchy level.

The default quality level for a Synchronous Ethernet interface is based on the value of network option: The `option-1` statement, when set, selects the `sec` quality level; and the `option-2` statement, when set, selects the `st3` quality level.

Table 106 on page 13921 shows whether SSM quality level is supported for a given external interface signal type and framing. The default setting is disabled.

- **Table 106: SSM-Quality LevelSupport by Signal Type and Framing**

Signal Type	Framing	SSM Quality Level Supported
E1	G.704	yes
E1	G.704 no CRC4	no
T1	ESF	yes

Table 106: SSM-Quality LevelSupport by Signal Type and Framing *(Continued)*

Signal Type	Framing	SSM Quality Level Supported
T1	SF	no
2048 KHz	Not applicable	no

Default

By default, this statement is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

show chassis synchronization | 18879

quality-mode-enable (PTX)

IN THIS SECTION

- Syntax | 13923
- Hierarchy Level | 13923
- Description | 13923

- Default | 13923
- Required Privilege Level | 13923
- Release Information | 13923

Syntax

```
quality-mode-enable;
```

Hierarchy Level

```
[edit chassis synchronization (PTX Series)]
```

Description

Enable Synchronous Ethernet ESMC quality mode. The quality level parameter for a Synchronous Ethernet interface is optional when the `quality-mode-enable` and the `selection-mode received-quality` statements are included at the `[edit chassis synchronization]` hierarchy level.

The default quality level for a Synchronous Ethernet interface is based on the value of the network option: The `option-1` statement, when included, selects the `sec` quality level; and the `option-2` statement, when included, selects the `st3` quality level.

Default

By default, this statement is not included.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R1.

RELATED DOCUMENTATION

synchronization (PTX Series) | 28802

Synchronizing Internal Stratum 3 Clock to External Clock Sources on PTX Series Routers

show chassis synchronization | 18879

quarantine-message (Security Content Security)

IN THIS SECTION

- Syntax | 13924
- Hierarchy Level | 13924
- Description | 13925
- Options | 13925
- Required Privilege Level | 13925
- Release Information | 13925

Syntax

```
quarantine-message {
    type {
        custom-redirect-url;
    }
    url url;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name]
```

Description

Configure Juniper enhanced quarantine message settings.

Options

- `type`—Specify the following type of the quarantine message:
  - `custom-redirect-url`—Specify Custom redirect URL server.
- `url` *`url`*—Specify an URL of the quarantine message.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

querier (performance-monitoring)

IN THIS SECTION

- Syntax | 13926
- Hierarchy Level | 13926
- Description | 13927
- Required Privilege Level | 13927
- Release Information | 13927

Syntax

```
querier {
    delay {
        traffic-class tc-value {
            average-sample-size sample size;
            padding-size size;
            query-interval milliseconds;
            rtt-delay-threshold rtt threshold value;
            twcd-delay-threshold twcd threshold value;
        }
    }
    loss {
        traffic-class tc-value {
            average-sample-size sample size;
            loss-threshold loss threshold value;
            loss-threshold-window number of samples for loss threshold;
            measurement-quantity bytes|packets;
            query-interval milliseconds;
        }
    }
    loss-delay {
        traffic-class tc-value {
            average-sample-size sample size;
            loss-threshold loss threshold value;
            loss-threshold-window number of samples for loss threshold;
            measurement-quantity bytes|packets;
            padding-size size;
            query-interval milliseconds;
            rtt-delay-threshold rtt threshold value;
            twcd-delay-threshold twcd threshold value;
        }
    }
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring],
```

```
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring]
```

Description

Configure querier options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

query-interval (Bridge Domains)

IN THIS SECTION

- Syntax | 13927
- Hierarchy Level | 13928
- Description | 13928
- Options | 13928
- Required Privilege Level | 13928
- Release Information | 13928

Syntax

```
query-interval seconds;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-name],
[edit bridge-domains bridge-domain-name protocols mld-snooping],
[edit protocols igmp-snooping vlan],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols mld-snooping],
[edit routing-instances routing-instance-name protocols mld-snooping]
```

Description

Configure the interval for host-query message timeouts.

Options

seconds—Time interval. This value must be greater than the interval set for query-response-interval.

- **Range:** 1 through 1024
- **Default:** 125 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Example: Configuring IGMP Snooping*

query-last-member-interval (Bridge Domains) | 13933
query-response-interval (Bridge Domains) | 13938
mld-snooping | 10783
igmp-snooping | 8387
robust-count (IGMP Snooping) | 15815
IGMP Snooping Overview

query-interval (Protocols IGMP)

IN THIS SECTION

- Syntax | 13929
- Hierarchy Level | 13929
- Description | 13929
- Options | 13930
- Required Privilege Level | 13930
- Release Information | 13930

Syntax

```
query-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp],
[edit protocols igmp]
```

Description

Specify how often the querier routing device sends general host-query messages.

Options

seconds—Time interval.

- **Range:** 1 through 1024
- **Default:** 125 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Modifying the IGMP Host-Query Message Interval

query-last-member-interval (Protocols IGMP) | **13935**

query-response-interval (Protocols IGMP) | **13940**

query-interval (Protocols IGMP AMT)

IN THIS SECTION

- Syntax | **13931**
- Hierarchy Level | **13931**
- Description | **13931**
- Options | **13931**
- Required Privilege Level | **13931**
- Release Information | **13931**

Syntax

```
query-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp amt relay defaults],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols igmp
amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances routing-instance-name protocols igmp amt relay defaults]
```

Description

Specify how often the querier router sends IGMP general host-query messages through an Automatic Multicast Tunneling (AMT) interface.

Options

seconds—Number of seconds between sending of general host query messages.

- **Range:** 1 through 1024
- **Default:** 125 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

> *Configuring Default IGMP Parameters for AMT Interfaces*

query-interval (Protocols MLD)

IN THIS SECTION

- Syntax | 13932
- Hierarchy Level | 13932
- Description | 13932
- Options | 13932
- Required Privilege Level | 13933
- Release Information | 13933

Syntax

```
query-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld],
[edit protocols mld]
```

Description

Specify how often the querier router sends general host-query messages.

Options

seconds—Time interval.

- **Range:** 1 through 1024
- **Default:** 125 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Modifying the MLD Host-Query Message Interval

query-last-member-interval (Protocols MLD) | 13936

query-response-interval (Protocols MLD) | 13943

query-last-member-interval (Bridge Domains)

IN THIS SECTION

- Syntax | 13933
- Hierarchy Level | 13934
- Description | 13934
- Options | 13934
- Required Privilege Level | 13934
- Release Information | 13934

Syntax

```
query-last-member-interval seconds;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-name],
[edit bridge-domains bridge-domain-name protocols mld-snooping],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols mld-snooping],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id],
 [edit routing-instances routing-instance-name protocols mld-snooping interface interface-name],
[edit protocols igmp-snooping vlan]
```

Description

Configure the interval for group-specific query timeouts.

Options

seconds—Time interval, in fractions of a second or seconds.

- **Range:** 0.1 through 0.9, then in 1-second intervals 1 through 1024
- **Default:** 1 second

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Example: Configuring IGMP Snooping*

query-interval (Bridge Domains) | 13927
query-response-interval (Bridge Domains) | 13938
mld-snooping | 10783
igmp-snooping | 8387
Example: Configuring IGMP Snooping on SRX Series Devices

query-last-member-interval (Protocols IGMP)

IN THIS SECTION

- Syntax | 13935
- Hierarchy Level | 13935
- Description | 13935
- Options | 13936
- Required Privilege Level | 13936
- Release Information | 13936

Syntax

```
query-last-member-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp],
[edit protocols igmp]
```

Description

Specify how often the querier routing device sends group-specific query messages.

Options

seconds—Time interval, in fractions of a second or seconds.

- **Range:** 0.1 through 0.9, then in 1-second intervals 1 through 999999
- **Default:** 1 second

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Modifying the IGMP Last-Member Query Interval

query-interval (Protocols IGMP) | **13929**

query-response-interval (Protocols IGMP) | **13940**

query-last-member-interval (Protocols MLD)

IN THIS SECTION

- Syntax | **13937**
- Hierarchy Level | **13937**
- Description | **13937**
- Options | **13937**
- Required Privilege Level | **13937**
- Release Information | **13937**

Syntax

```
query-last-member-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld],
[edit protocols mld]
[edit protocols mld-snooping vlan vlan-id]
[edit routing-instances instance-name protocols mld-snooping vlan vlan-id]
```

Description

Specify how often the querier routing device sends group-specific query messages.

Options

seconds—Time interval, in fractions of a second or seconds.

- **Range:** 0.1 through 0.9, then in 1-second intervals from 1 through 1024
- **Default:** 1 second

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit protocols mld-snooping vlan vlan-id]` and the `[edit routing-instances instance-name protocols mld-snooping vlan vlan-id]` hierarchy levels introduced in Junos OS Release 13.3 for EX Series switches.

Support at the `[edit protocols mld-snooping vlan vlan-id]` hierarchy level introduced in Junos OS Release 18.1R1 for the SRX1500 devices.

RELATED DOCUMENTATION

Example: Configuring MLD Snooping on SRX Series Devices

mld-snooping | **10783**

Modifying the MLD Last-Member Query Interval

query-interval (Protocols MLD) | **13932**

query-response-interval (Protocols MLD) | **13943**

Understanding MLD Snooping

query-response-interval (Bridge Domains)

IN THIS SECTION

- Syntax | **13938**
- Hierarchy Level | **13938**
- Description | **13939**
- Options | **13939**
- Required Privilege Level | **13939**
- Release Information | **13939**

Syntax

```
query-response-interval seconds;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-
name],
[edit bridge-domains bridge-domain-name protocols mld-snooping],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
snooping interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
```

```
snoopingvlan vlan-id],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols mld-
snooping interface interface-name],
[edit routing-instances routing-instance-name protocols mld-snooping],
[edit protocols igmp-snooping vlan]
```

Description

Specify how long to wait to receive a response to a specific query message from a host.

Options

seconds—Time interval. This interval should be less than the host-query interval.

- **Range:** 1 through 1024
- **Default:** 10 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping on SRX Series Devices

Example: Configuring IGMP Snooping

query-interval (Bridge Domains) | **13927**

query-last-member-interval (Bridge Domains) | **13933**

mld-snooping | **10783**

igmp-snooping | **8387**

query-response-interval (Protocols IGMP)

IN THIS SECTION

- Syntax | 13940
- Hierarchy Level | 13940
- Description | 13940
- Options | 13940
- Required Privilege Level | 13941
- Release Information | 13941

Syntax

```
query-response-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp],
[edit protocols igmp]
```

Description

Specify how long the querier routing device waits to receive a response to a host-query message from a host.

Options

seconds—The query response interval must be less than the query interval.

- **Range:** 1 through 1024
- **Default:** 10 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Modifying the IGMP Query Response Interval

query-interval (Protocols IGMP) | 13929

query-last-member-interval (Protocols IGMP) | 13935

query-response-interval (Protocols IGMP AMT)

IN THIS SECTION

- Syntax | 13941
- Hierarchy Level | 13942
- Description | 13942
- Options | 13942
- Required Privilege Level | 13942
- Release Information | 13942

Syntax

```
query-response-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp amt relay defaults],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols igmp
amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances routing-instance-name protocols igmp amt relay defaults]
```

Description

Specify how long the IGMP querier router waits to receive a response to a host query message from a host through an Automatic Multicast Tunneling (AMT) interface.

Options

seconds—Time to wait to receive a response to a host query message. The query response interval must be less than the query interval.

- **Range:** 1 through 1024
- **Default:** 10 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

> *Configuring Default IGMP Parameters for AMT Interfaces*

query-response-interval (Protocols MLD)

IN THIS SECTION

- Syntax | 13943
- Hierarchy Level | 13943
- Description | 13943
- Options | 13943
- Required Privilege Level | 13944
- Release Information | 13944

Syntax

```
query-response-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld],
[edit protocols mld]
[edit protocols mld-snooping vlan vlan-id]
[edit routing-instances instance-name protocols mld-snooping vlan vlan-id]
```

Description

Specify how long the querier routing device waits to receive a response to a host-query message from a host.

Options

seconds—Time interval.

- **Range:** 1 through 1024
- **Default:** 10 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit protocols mld-snooping vlan vlan-id]` and the `[edit routing-instances instance-name protocols mld-snooping vlan vlan-id]` hierarchy levels introduced in Junos OS Release 13.3 for EX Series switches.

RELATED DOCUMENTATION

Modifying the MLD Query Response Interval

query-interval (Protocols MLD) | **13932**

query-last-member-interval (Protocols MLD) | **13936**

queue (Global Queues)

IN THIS SECTION

- Syntax | **13945**
- Hierarchy Level | **13945**
- Description | **13945**
- Options | **13945**
- Required Privilege Level | **13945**
- Release Information | **13945**

Syntax

```
queue queue-number class-name;
```

Hierarchy Level

```
[edit class-of-service forwarding-classes]
```

Description

Specify the output transmission queue to which to map all input from an associated forwarding class.

This statement enables you to configure up to eight forwarding classes with one-to-one mapping to output queues. If you want to configure up to 16 forwarding classes with multiple forwarding classes mapped to single output queues, include the `class` statement instead of the `queue` statement at the `[edit class-of-service forwarding-classes]` hierarchy level.

Options

class-name—Name of forwarding class.

queue-number—Output queue number.

- **Range:** 0 through 7.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring a Custom Forwarding Class for Each Queue*
>
> class (Forwarding Classes) | 3364

queue-monitoring

IN THIS SECTION

- Syntax | 13946
- Hierarchy Level | 13946
- Release Information | 13946
- Description | 13946
- Required Privilege Level | 13946

Syntax

```
queue-monitoring;
```

Hierarchy Level

```
[edit class-of-service interface interface-name],
[edit class-of-service traffic-control-profile lcp-name],
[edit class-of-service schedulers scheduler-name]
```

Release Information

Statement introduced in Junos OS Release 21.2R1.

Description

Enable queue depth monitoring.

Required Privilege Level

interface

RELATED DOCUMENTATION

show class-of-service interface | 19056

show class-of-service scheduler-map | 19125

show interfaces voq | 21593

queue-num

IN THIS SECTION

- Syntax | 13947
- Hierarchy Level | 13947
- Description | 13947
- Options | 13949
- Required Privilege Level | 13949
- Release Information | 13949

Syntax

```
queue-num queue-number <no-loss>;
```

Hierarchy Level

```
[edit class-of-service forwarding-classes class class-name]
```

Description

Map a forwarding class to an output queue number. Optionally, configure the forwarding class as a lossless forwarding class. Each switch provides enough output queues so that you can map forwarding classes to queues on a one-to-one basis, so each forwarding class can have a dedicated output queue.

On switches that use different forwarding classes and output queues for unicast and multidestination (multicast, broadcast, destination lookup fail) traffic, the switch supports 12 forwarding classes and 12

output queues, eight of each for unicast traffic and four of each for multidestination traffic. You can map some or all of the eight unicast forwarding classes to a unicast queue (0 through 7) and some or all of the four multidestination forwarding classes to the a multidestination queue (8 through 11). You cannot map a forwarding class to more than one queue (each forwarding class maps to one and only one queue), but you can map multiple forwarding classes to one queue. The queue to which you map a forwarding class determines if the forwarding class is a unicast or multidestination forwarding class.

On switches that use the same forwarding classes and output queues for unicast and multidestination traffic, the switch supports eight forwarding classes and eight output queues. You can map some or all of the eight of the forwarding classes to queues (0 through 7). You cannot map a forwarding class to more than one queue (each forwarding class maps to one and only one queue), but you can map multiple forwarding classes to one queue.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the no-loss packet drop attribute. Do not associate a drop profile with lossless forwarding classes. Instead, use priority-based flow control (PFC) to prevent frame drop on lossless forwarding classes.

NOTE: If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

NOTE: On systems that do not use the ELS CLI, if you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless `fcoe` and `no-loss` forwarding classes. If you explicitly configure lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

NOTE: On systems that do not use the ELS CLI, if you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the `fcoe` and `no-loss` forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the `no-loss` option. If you do not specify the `no-loss` option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the `fcoe` forwarding class and you do not include the `no-loss` option, the `fcoe` forwarding class is lossy, not lossless.

Options

queue-number—(Switches that use different output queues for unicast and multidestination traffic) Number of the CoS unicast queue (0 through 7) or the CoS multidestination queue (8 through 11).

queue-number—(Switches that use the same output queues for unicast and multidestination traffic) Number of the CoS queue (0 through 7).

`no-loss`—Optional packet drop attribute keyword to configure the forwarding class as lossless.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

No-loss option introduced in Junos OS Release 12.3 for the QFX Series.

queue (Restricted Queues)

IN THIS SECTION

- Syntax | **13950**
- Hierarchy Level | **13950**
- Description | **13950**
- Options | **13950**
- Required Privilege Level | **13950**
- Release Information | **13950**

Syntax

```
queue queue-number;
```

Hierarchy Level

```
[edit class-of-service restricted-queues forwarding-class class-name]
```

Description

For M320, MX Series, T Series routers and EX Series switches only, map forwarding classes to restricted queues.

Options

queue-number—Output queue number.

- **Range:** 0 through 3.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

queue-statistics

IN THIS SECTION

- Syntax | 13951
- Hierarchy Level | 13951

- Description | 13951
- Default | 13951
- Options | 13952
- Required Privilege Level | 13952
- Release Information | 13952

Syntax

```
queue-statistics {
    file filename {
        files number-of-files;
        size size;
    }
    interval interval;
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Enable the logging of queue statistics in a local file. This statement does not enable queue statistics monitoring.

To enable queue monitoring, you must specify the `queue-statistics` configuration statement at the `[edit services analytics interfaces]` hierarchy level.

Default

This feature is disabled by default.

Options

interval *interval*

Configure the polling interval in milliseconds.

> **NOTE**: You can configure the polling interval for queue statistics globally for all interfaces only. Due to limitations and variations in the hardware capability of different devices, you might see a difference in value between the actual interval and configured interval.

- Range:
  - Junos OS Release 13.2X50-D15—8 to 1000 milliseconds (8 milliseconds to 1 second)
  - Junos OS Release 13.2X51-D10 or later—10 to 1000 milliseconds (10 milliseconds to 1 second)

    > **NOTE**: In Junos OS Release 13.2X51-D10 or later, if you configured an interval of less than 10 milliseconds, the following warning messages appear during the commit process: `Queue statistics polling interval can not be less than 10 milliseconds` and `Setting Queue statistics polling interval to 10 milliseconds`. These messages do not stop the commit operation, but the interval is automatically set to 10 milliseconds.

- Default:
  - Junos OS Release 13.2X50-D15—8 milliseconds
  - Junos OS Release 13.2X51-D10 or later—10 milliseconds

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

queue-threshold

IN THIS SECTION

- Syntax | 13953
- Hierarchy Level | 13953
- Description | 13954
- Options | 13954
- Required Privilege Level | 13954
- Release Information | 13955

Syntax

```
queue-threshold {
    fabric-queue {
        priority priority {
            threshold threshold-percentage;
        }
    }
    wan-queue {
        priority priority {
            threshold threshold-percentage;
        }
    }
}
```

Hierarchy Level

```
[edit chassis fpc slot-number traffic-manager]
```

Description

Enable monitoring of fabric and WAN queues.

When the `fabric-queue` statement is configured, an SNMP trap is generated when the fabric queue depth exceeds the configured threshold value.

When `wan-queue` is configured, an SNMP trap is generated whenever the WAN queue depth exceeds the configured threshold value.

By default, monitoring of fabric queue and WAN queue is disabled.

Options

fabric-queue Configure threshold values to monitor fabric queue depth. An SNMP trap is generated when fabric queue depth exceeds the configured threshold value. When this option is configured, you can use the `show class-of-service fabric statistics detail` command or the `show pfe statistics traffic detail` command to display detailed statistics at the level of the Packet Forwarding Engine.

- `priority` *`priority`*—Specify the queue priority for fabric queues. The options are `high` and `low`. Traffic with `low` priority is sent after any queue that has a `high` priority. The priorities map to numeric priorities in the underlying hardware.
- `threshold` *`threshold-percentage`*—Configured threshold value for the priority level for fabric queues. An SNMP trap is generated when the actual queue utilization percentage exceeds the configured threshold value.

wan-queue Configure threshold values to monitor the WAN queue depth. An SNMP trap is generated when the WAN queue depth exceeds the configured value.

- `priority` *`priority`*—Specify the queue priority for WAN queues. The priority level can be `high`, `medium-high`, `medium-low`, and `low`.
- `threshold` *`threshold-percentage`*—Configured threshold value for the priority level for WAN queues. An SNMP trap is generated when the actual queue utilization percentage exceeds the configured threshold value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

traffic-manager | 30036

queues

IN THIS SECTION

- Syntax | 13955
- Hierarchy Level | 13955
- Description | 13956
- Options | 13956
- Required Privilege Level | 13956
- Release Information | 13956

Syntax

```
queues [ queue-numbers ];
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number compression rtp],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
compression rtp]
```

Description

For voice services interfaces only, assign queue numbers on which RTP compression takes place.

Options

`queues` *`queue-numbers`*—Assign one or more of the following queues: `q0`, `q1`, `q2`, and `q3`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Compression of Voice Traffic

quit

IN THIS SECTION

- Syntax | 13957
- Description | 13957
- Options | 13957
- Required Privilege Level | 13957
- Release Information | 13957

Syntax

```
quit <configuration-mode>
```

Description

Exit the current level of the statement hierarchy, returning to the level prior to the last `edit` command, or exit from configuration mode. The `quit` and `exit` commands are synonyms.

Options

none Return to the previous edit level. If you are at the top of the statement hierarchy, exit configuration mode.

configuration-mode (Optional) Exit from configuration mode.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Displaying the Current Configuration

R

IN THIS SECTION

- r2cp | 13990
- radio-router (Interfaces) | 13991
- radius (Access Profile) | 13993
- radius-accounting | 13997
- radius-coa (Adjustment Control Profiles) | 13999
- radius-disconnect | 14001
- radius-disconnect (DHCP Local Server) | 14003
- radius-disconnect-port | 14005
- radius-flow-tap | 14006
- radius-options (Access) | 14010
- radius-options (Interfaces) | 14011
- radius-options (System) | 14013
- radius-reachability | 14015
- radius-realm | 14017
- radius (Security Log) | 14018
- radius-server | 14020
- radius-server (Access) | 14026
- radius-server-name | 14028
- radius-server (System) | 14030
- radius (System) | 14034
- radsec | 14036
- radsec-destination | 14040
- raise-threshold | 14041
- raise-trap (Event Policy) | 14042
- random | 14043
- range (Access) | 14045
- range-address | 14047
- range (Address-Assignment Pools) | 14048

- range (Source NAT Next Gen Services) | **14049**
- ranges (Dynamic Stacked VLAN) | **14051**
- ranges (Dynamic VLAN) | **14052**
- rapid-commit (DHCPv6 Client) | **14054**
- rapid-commit (DHCPv6 Local Server) | **14055**
- rate | **14057**
- rate-cap | **14058**
- rate (Forwarding Options) | **14060**
- rate (Interface Services) | **14061**
- rate-limit | **14063**
- rate-limit (Aggregated rate limit) | **14065**
- rate-limit (Security GTP) | **14068**
- rate-limit (Security Log) | **14069**
- rate-limiters | **14071**
- rate (Routing Instances) | **14073**
- rcp-count | **14075**
- rdi | **14077**
- rdma-opcode | **14078**
- re-assembler | **14079**
- re-enroll-trigger-time-percentage | **14081**
- re-generate-keypair | **14082**
- reachability-fault | **14083**
- reachable-time | **14086**
- reachable-time (Dynamic Router Advertisement) | **14087**
- read-view | **14089**
- readvertise | **14090**
- real-service (Traffic Load Balancer) | **14092**
- realm | **14093**
- realm-delimiter (Domain Map) | **14095**
- realm (Diameter Origin) | **14096**
- realm-parse-direction (Domain Map) | **14098**
- realtime-ukern-thread (BFD) | **14099**

- reassemble-packets | **14100**
- reassembly-timeout | **14102**
- reauthenticate (DHCP Local Server) | **14103**
- rebalance (Aggregated Ethernet Interfaces) | **14106**
- rebalance-periodic (Aggregated Ethernet Subscriber Interfaces) | **14108**
- rebalance-periodic (PS Interface) | **14109**
- rebalance-subscriber-granularity (PS interface) | **14111**
- receive-ais | **14112**
- receive-failure-threshold (RFC 2544 Benchmarking) | **14113**
- receive-options-packets (Passive Flow Monitoring) | **14114**
- receive-options-packets | **14116**
- receive (Protocols RIP) | **14117**
- receive (Protocols RIPng) | **14119**
- receive-ttl-exceeded (Passive Flow Monitoring) | **14120**
- receive-ttl-exceeded | **14121**
- receive-window | **14123**
- receiver | **14124**
- receiver (Graceful Restart for BGP Helper) | **14127**
- recommendation-tlv | **14129**
- recommended | **14130**
- recommended-action | **14131**
- reconfigure (DHCP Local Server) | **14133**
- reconnect-time | **14135**
- record | **14136**
- record (Origin Validation for BGP) | **14138**
- record-route-option (IDS Screen Next Gen Services) | **14139**
- recover-time (DDoS) | **14141**
- recovered-clock | **14142**
- recovery | **14143**
- recovery | **14146**
- recovery-time (LDP) | **14148**
- recovery-timeout | **14149**

- recycle-bandwidth-profiles | **14152**
- red-buffer-occupancy | **14154**
- red-differential-delay | **14156**
- redial-delay | **14157**
- redirect (PCC Action Profiles) | **14159**
- redirect-traffic | **14160**
- redirect-url | **14162**
- redirect-wx (Application Services) | **14164**
- redistribute-all-traffic (Aggregated Multiservices) | **14165**
- redundancy | **14166**
- redundancy-event (Services Redundancy Daemon) | **14168**
- redundancy (Graceful Switchover) | **14170**
- redundancy-group | **14171**
- redundancy-group (Chassis Cluster) | **14173**
- redundancy-group (Chassis - MX Series) | **14176**
- redundancy-group-id (Junos Fusion) | **14177**
- redundancy-group (Interfaces) | **14179**
- redundancy-group (Redundant Tunnel) | **14180**
- redundancy-interface-process | **14182**
- redundancy (Junos OS) | **14183**
- redundancy (Junos OS Evolved) | **14185**
- redundancy (M:N Subscriber Redundancy) | **14186**
- redundancy-mode | **14191**
- redundancy-options (Adaptive Services Interfaces) | **14192**
- redundancy-options (Aggregated Multiservices) | **14194**
- redundancy-options (Link Services IQ PIC Interfaces) | **14195**
- redundancy-options (Stateful Synchronization) | **14197**
- redundancy-policy (Interchassis Services Redundancy) | **14199**
- redundancy (Power Management) | **14201**
- redundancy (satellite-management) | **14203**
- redundancy-set | **14204**
- redundancy-set-id (Service Set) | **14206**

- redundancy-mode | **14207**
- redundant-ether-options | **14209**
- redundant-logical-tunnel | **14211**
- redundant-parent (Interfaces) | **14212**
- redundant-parent (Interfaces Fast Ethernet) | **14214**
- redundant-parent (Interfaces Gigabit Ethernet) | **14215**
- redundant-power-system | **14216**
- redundant-pseudo-interface-options | **14218**
- redundant-sources | **14219**
- redundant-trunk-group | **14221**
- redundant-virtual-tunnel | **14222**
- reference-bandwidth (Protocols IS-IS) | **14224**
- reference-bandwidth (Protocols OSPF) | **14226**
- reflect-etype | **14228**
- reflect-mode | **14229**
- reflective-relay | **14232**
- reflexive | revert | reverse | **14233**
- refresh (Commit Scripts) | **14234**
- refresh (Event Scripts) | **14236**
- refresh-from (Commit Scripts) | **14238**
- refresh-from (Event Scripts) | **14240**
- refresh-from (JET) | **14242**
- refresh-from (Op Scripts) | **14244**
- refresh-from (SNMP Scripts) | **14246**
- refresh-interval | **14248**
- refresh (JET) | **14250**
- refresh (Op Scripts) | **14251**
- refresh-rate (Flow Monitoring Logs for NAT) | **14253**
- refresh (SNMP Scripts) | **14255**
- refresh-time | **14257**
- regex-additive-logic | **14258**
- regex (Class Attribute) | **14260**

- regexp | **14261**
- region | **14262**
- register-limit | **14264**
- register-probe-time | **14266**
- registration | **14268**
- reject (Security) | **14270**
- reject-unauthorized-ipv6cp | **14271**
- rejoin-timeout (Aggregated Multiservices) | **14273**
- relay-agent-interface-id (DHCP Local Server) | **14275**
- relay-agent-interface-id (DHCPv6 Relay Agent) | **14276**
- relay-agent-interface-id (DHCPv6 Relay Agent Username) | **14278**
- relay-agent-option | **14279**
- relay-agent-option-79 | **14281**
- relay-agent-remote-id (DHCP Local Server) | **14282**
- relay-agent-remote-id (DHCPv6 Relay Agent) | **14284**
- relay-agent-remote-id (DHCPv6 Relay Agent Username) | **14286**
- relay-agent-subscriber-id (DHCP Local Server) | **14288**
- relay-agent-subscriber-id (DHCPv6 Relay Agent) | **14289**
- relay (AMT Protocol) | **14291**
- relay (IGMP) | **14292**
- relay-option-82 | **14294**
- relay-option (DHCP Relay Agent) | **14297**
- relay-option-vendor-specific (dhcpv6) | **14299**
- relay-server-group (DHCP Relay Agent Option) | **14300**
- relay-source | **14302**
- reliable | **14304**
- remote | **14305**
- remote-access (Juniper Secure Connect) | **14307**
- remote-address | **14311**
- remote-address (for LMP Control Channel) | **14312**
- remote-address (for LMP Traffic Engineering) | **14314**
- remote-debug-permission | **14315**

- remote-device-management | **14317**
- remote-engine | **14318**
- remote-execution (Event Scripts) | **14320**
- remote-gateway | **14322**
- remote-gateway (Tunnel Profile) | **14323**
- remote (Host VPN) | **14324**
- remote-id (Services) | **14326**
- remote-id | **14327**
- remote-id (protocols link-management) | **14330**
- remote-id (DHCP Relay Agent) | **14331**
- remote-id-mismatch (DHCP Local Server and DHCP Relay Agent) | **14334**
- remote-id (MX Series) | **14337**
- remote-id (VLAN Authentication Username) | **14338**
- remote-identity | **14340**
- remote-interface-switch | **14342**
- remote-ip-address | **14344**
- remote-ip-host-routes | **14346**
- remote-loop-enable | **14348**
- remote-loopback | **14350**
- remote-maintenance-association (OAM) | **14351**
- remote-mep | **14352**
- remote-mep (EX Series Switch Only) | **14354**
- remote-port-range | **14355**
- remote-ports | **14357**
- remote-site-id | **14360**
- remote-traffic-selector | **14361**
- remote-vtep-list | **14362**
- remotely-controlled | **14364**
- remove-ie | **14366**
- remove-private | **14368**
- remove-when-no-subscribers | **14370**
- rename | **14371**

- renegotiation (Services) | **14373**
- reopen-time | **14374**
- replace | **14376**
- replace-ip-source-with (DHCP Relay Agent) | **14377**
- replay-attacks | **14379**
- replay-protect | **14380**
- replay-protect (MX Series) | **14382**
- replay-window-size | **14383**
- replay-window-size (MX Series) | **14385**
- replicator | **14387**
- report-ingress-shaping-rate (Dynamic CoS Interfaces) | **14390**
- report-interface-descriptions (Access) | **14392**
- report-local-rule (PCRF Partition) | **14393**
- report (LRF Rule) | **14395**
- report-resource-allocation (PCRF Partition) | **14396**
- report (Security Log) | **14398**
- report-skip | **14403**
- report-successful-resource-allocation (PCRF Partition) | **14404**
- report-zero-oif-gw-on-discard | **14406**
- req-option | **14408**
- req-timeout | **14409**
- request | **14411**
- request access-security router-advertisement-guard-block | **14413**
- request access-security router-advertisement-guard-forward | **14414**
- request access-security router-advertisement-guard-learn interface | **14415**
- request access-security slaac-snooping unblock | **14417**
- request ancp oam interface | **14419**
- request ancp oam neighbor | **14421**
- request ancp oam port-down | **14423**
- request ancp oam port-up | **14425**
- request app-engine cleanup | **14427**
- request app-engine file-copy (crash | log) from-jhost to-vjunos | **14429**

- request auto-configuration reconnect-pending | **14431**
- request-cache-timeout (RADIUS Snoop Segment) | **14432**
- request chassis afeb | **14434**
- request chassis beacon | **14435**
- request chassis cb (Security) | **14440**
- request chassis cb | **14441**
- request chassis ccg | **14447**
- request chassis cfeb | **14449**
- request chassis clock master switch | **14450**
- request chassis cluster configuration-synchronize | **14452**
- request chassis cluster failover node | **14454**
- request chassis cluster failover redundancy-group | **14455**
- request chassis cluster failover reset | **14457**
- request chassis cluster in-service-upgrade abort (ISSU) | **14459**
- request chassis device-mode | **14460**
- request chassis device-mode satellite | **14463**
- request chassis fabric fpc | **14465**
- request chassis fabric pfe | **14467**
- request chassis fabric plane | **14468**
- request chassis feb | **14472**
- request chassis fpc | **14475**
- request chassis fpc-control-port | **14486**
- request chassis fpc slc | **14488**
- request chassis fpm resync | **14489**
- request chassis high-availability failover services-redundancy-group | **14492**
- request chassis mcs | **14494**
- request chassis mic | **14495**
- request chassis optics | **14499**
- request chassis pcg | **14501**
- request chassis pic | **14502**
- request chassis pic-mode | **14509**
- request chassis port-led | **14512**

- request chassis primary-ha-control-port-transition | 14516
- request chassis redundancy feb slot | 14518
- request chassis routing-engine hard-disk-test | 14520
- request chassis routing-engine master | 14529
- request chassis satellite beacon | 14535
- request chassis satellite disable | 14537
- request chassis satellite enable | 14539
- request chassis satellite file-copy | 14541
- request chassis satellite install | 14542
- request chassis satellite interface | 14545
- request chassis satellite login | 14547
- request chassis satellite reboot | 14548
- request chassis satellite restart | 14550
- request chassis satellite shell-command | 14551
- request chassis scg | 14553
- request chassis sfb | 14555
- request chassis sfm | 14558
- request chassis sfm master switch | 14559
- request chassis sib | 14562
- request chassis spmb restart | 14565
- request chassis ssb master switch | 14568
- request chassis synchronization mode | 14570
- request (Chassis Synchronization Source) | 14573
- request chassis synchronization switch | 14574
- request chassis system-mode | 14577
- request (Clock Synchronization) | 14579
- request component login | 14580
- request csds add-vsrx | 14583
- request csds authenticate-host | 14584
- request csds delete-vsrx | 14586
- request csds extract-vsrx-keys | 14587
- request csds jdm | 14589

- request csds sync-controller | **14592**
- request debug information | **14594**
- request dhcp client renew | **14595**
- request dhcp relay bulk-leasequery | **14597**
- request dhcp relay leasequery | **14599**
- request dhcp server reconfigure | **14601**
- request dhcpv6 client renew | **14604**
- request dhcpv6 relay bulk-leasequery | **14605**
- request dhcpv6 relay leasequery | **14607**
- request dhcpv6 server reconfigure | **14609**
- request diagnostics tdr | **14611**
- request extension-service restart-daemonize-app | **14614**
- request extension-service (start | stop) | **14615**
- request fabric administration director-group change-master | **14616**
- request fabric administration power-off interconnect-device | **14618**
- request fabric administration power-off interconnect-device fpc | **14620**
- request fabric administration power-off node-device | **14622**
- request fabric administration remove | **14624**
- request fabric administration system mac-pool add | **14626**
- request fabric administration system mac-pool delete | **14628**
- request fibre-channel proxy load-rebalance | **14630**
- request flight-recorder set high-cpu | **14633**
- request forwarding-options port-mirroring instance family self-mirror-start | **14636**
- request forwarding-options port-mirroring instance family self-mirror-stop | **14638**
- request interface bounce | **14639**
- request interface link-degrade-recover | **14642**
- request interface load-balancing revert (Aggregated Multiservices) | **14646**
- request interface load-balancing switchover (Aggregated Multiservices) | **14647**
- request interface mc-ae switchover (Multichassis Link Aggregation) | **14649**
- request interface optics-reset | **14651**
- request interface rebalance (Aggregated Ethernet for Subscriber Management) | **14653**
- request interface revert | **14657**

- request interface (revert | switchover) (Adaptive Services) | **14658**
- request interface (revert | switchover) (Aggregated Ethernet Link Protection) | **14660**
- request interface (revert | switchover) (Aggregated Inline Service Interfaces) | **14662**
- request interface switchover | **14664**
- request ipsec switch | **14665**
- request jdaf clear statistics service | **14667**
- request jnu role | **14668**
- request jnu satellite sync | **14670**
- request jnu schema add | **14671**
- request jnu schema delete | **14673**
- request jnu schema merge | **14674**
- request l2circuit-switchover | **14675**
- request lacp link-switchover | **14677**
- request-max-tcp-connections (System Process) | **14678**
- request message | **14680**
- request modem wireless create-profile | **14681**
- request modem wireless fota | **14684**
- request modem wireless sim-lock | **14685**
- request modem wireless sim-unlock | **14687**
- request mpls container-lsp | **14689**
- request mpls lsp adjust-autobandwidth | **14691**
- request network-access aaa accounting | **14693**
- request network-access aaa replay pending-accounting-stops | **14695**
- request network-access aaa subscriber add session-id | **14696**
- request network-access aaa subscriber delete session-id | **14699**
- request network-access aaa subscriber modify session-id | **14702**
- request network-access aaa subscriber set session-id | **14704**
- request network-access aaa saml export-sp-metadata | **14706**
- request network-access aaa saml load-idp-metadata | **14707**
- request node halt (Junos OS Evolved) | **14708**
- request node (offline | online) (Junos OS Evolved) | **14711**
- request node power-off (Junos OS Evolved) | **14713**

- request node power-on (Junos OS Evolved) | **14715**
- request node reboot (re0 | re1) (Junos OS Evolved) | **14717**
- request packet-capture start | **14720**
- request packet-capture stop | **14726**
- request path-computation-client active-pce | **14728**
- request pim multicast-tunnel rebalance | **14729**
- request pppoe connect | **14731**
- request pppoe disconnect | **14732**
- request protection-group ethernet-aps | **14734**
- request protection-group ethernet-ring | **14736**
- request-rate (Access) | **14738**
- request redundant-power-system multi-backup | **14739**
- request-response | **14741**
- request restart packet-forwarding-engine | **14742**
- request routing-engine | **14743**
- request security certificate enroll (Signed) | **14745**
- request security certificate enroll (Unsigned) | **14747**
- request security datapath-debug capture start | **14749**
- request security datapath-debug capture stop | **14750**
- request security idp jist-conversion | **14751**
- request security idp security-package download | **14754**
- request security idp security-package install | **14757**
- request security idp security-package offline-download | **14760**
- request security idp security package rollback | **14761**
- request security idp ssl-inspection key add | **14763**
- request security idp ssl-inspection key delete | **14766**
- request security idp storage-cleanup | **14768**
- request security ike debug-disable | **14770**
- request security ike debug-enable | **14771**
- request security integrity | **14773**
- request security internal-security-association refresh | **14774**
- request security key-pair | **14776**

- request security pki ca-certificate ca-profile-group default-trusted-ca-certs | **14778**
- request security pki ca-certificate ca-profile-group load | **14781**
- request security pki ca-certificate enroll | **14784**
- request security pki ca-certificate enroll (Security) | **14786**
- request security pki ca-certificate load | **14787**
- request security pki ca-certificate load (Security) | **14789**
- request security pki ca-certificate verify | **14791**
- request security pki ca-certificate verify (Security) | **14792**
- request security pki ca-profile-group-status | **14795**
- request security pki crl load | **14797**
- request security pki crl load (Security) | **14798**
- request security pki generate-certificate-request | **14800**
- request security pki generate-certificate-request (Security) | **14802**
- request security pki generate-key-pair | **14805**
- request security pki generate-key-pair | **14806**
- request security pki generate-key-pair (Security) | **14808**
- request security pki key-pair export | **14810**
- request security pki local-certificate enroll | **14812**
- request security pki local-certificate enroll cmpv2 | **14814**
- request security pki local-certificate enroll scep | **14817**
- request security pki local-certificate export | **14822**
- request security pki local-certificate generate-self-signed | **14823**
- request security pki local-certificate generate-self-signed | **14825**
- request security pki local-certificate generate-self-signed (Security) | **14827**
- request security pki local-certificate load | **14830**
- request security pki local-certificate re-enroll cmpv2 | **14831**
- request security pki local-certificate re-enroll scep | **14833**
- request security pki local-certificate verify | **14835**
- request security pki node-local generate-certificate-request | **14837**
- request security pki node-local generate-key-pair | **14840**
- request security pki node-local key-pair export | **14842**
- request security pki node-local local-certificate enroll | **14843**

- request security pki node-local local-certificate export | **14847**
- request security pki node-local local-certificate load | **14848**
- request security pki node-local local-certificate re-enroll | **14850**
- request security pki node-local local-certificate verify | **14852**
- request security pki sync-from-peer | **14854**
- request security pki verify-integrity-status | **14856**
- request security policies check | **14857**
- request security policies resync | **14861**
- request security re-distribution ipsec-vpn | **14863**
- request security ssh key-pair-identity generate | **14865**
- request security ssh password-less-authentication (Junos OS Evolved) | **14867**
- request security system-keystore | **14869**
- request security system-keystore unlink | **14870**
- request security tpm master-encryption-password set | **14871**
- request security uefi revoke | **14873**
- request security user-identification local-authentication-table delete | **14875**
- request security user-identification local-authorization-table add | **14877**
- request security utm anti-virus avira-engine | **14879**
- request security utm anti-virus juniper-express-engine | **14881**
- request security utm anti-virus kaspersky-lab-engine | **14883**
- request security utm anti-virus sophos-engine | **14885**
- request security utm custom-objects url-feed update | **14886**
- request security utm custom-objects url-feed update feed-name force | **14888**
- request security utm web-filtering categorize | **14889**
- request security utm web-filtering category download-install [version] | **14891**
- request security utm web-filtering category download [version] | **14892**
- request security utm web-filtering category install | **14893**
- request security utm web-filtering category migrate-to-ng-juniper | **14895**
- request security utm web-filtering category uninstall | **14897**
- request security utm web-filtering custom-page reload | **14898**
- request security utm web-filtering recategorize | **14899**
- request server authenticate-peer-server | **14901**

- request services advanced-anti-malware data-connection | **14903**
- request services advanced-anti-malware diagnostic | **14905**
- request services advanced-anti-malware redirect-file | **14910**
- request services anti-virus update | **14912**
- request services anti-virus machine-learning-engine | **14914**
- request services application-identification application | **14916**
- request services application-identification clear packet-capture all | **14918**
- request services application-identification download | **14919**
- request services application-identification download status | **14923**
- request services application-identification group | **14925**
- request services application identification install ignore duplicate version check | **14928**
- request services application-identification install status | **14929**
- request services application identification new to production | **14931**
- request services application-identification rollback status | **14933**
- request services application-identification install | **14934**
- request services application-identification offline-download package-path | **14936**
- request services application-identification proto-bundle-status | **14937**
- request services application-identification uninstall | **14939**
- request services application-identification uninstall status | **14940**
- request services extensible-subscriber-services reload-dictionary | **14942**
- request services flow-collector change-destination primary interface | **14943**
- request services flow-collector change-destination secondary interface | **14945**
- request services flow-collector test-file-transfer | **14946**
- request services ip-monitoring preempt-restore policy | **14948**
- request services ipsec-vpn ipsec switch tunnel | **14950**
- request services l2tp destination unlock | **14951**
- request services monitoring twamp client | **14953**
- request services paa install | **14954**
- request services paa uninstall | **14958**
- request services redundancy-set trigger | **14960**
- request services remote-device-management reconfigure service-device | **14961**
- request services remote-device-management reload-dictionary | **14962**

- request services rpm twamp | **14964**
- request services static-subscribers login group | **14966**
- request services static-subscribers login interface | **14968**
- request services static-subscribers logout group | **14969**
- request services static-subscribers logout interface | **14971**
- request services subscribers | **14972**
- request services subscribers clear | **14974**
- request services url-filter delete gencfg-data | **14975**
- request services url-filter force dns-resolution | **14976**
- request services url-filter update url-filter-database file | **14979**
- request services url-filter validate | **14980**
- request services user-identification active-directory-access active-directory-authentication-table delete | **14981**
- request services user-identification active-directory-access domain-controller | **14983**
- request services user-identification active-directory-access ip-user-probe | **14985**
- request services user-identification authentication-source aruba-clearpass user-query | **14987**
- request services user-identification authentication-source jims groups domain <domain-name> (force-fetch|status) | **14989**
- request services user-identification authentication-source jims validate (user <user-name>|group <group-name>|device <device-name>) domain <domain-name> | **14991**
- request services user-identification authentication-table delete | **14994**
- request services web-filter delete gencfg-data | **15004**
- request services web-filter force dns-resolution | **15005**
- request services web-filter update dns-filter-database | **15006**
- request services web-filter update url-filter-database file | **15007**
- request services web-filter validate dns-filter-file-name | **15009**
- request services web-filter validate url-filter-file-name | **15010**
- request session member | **15011**
- request snmp spoof-trap | **15012**
- request snmp utility-mib clear instance | **15021**
- request snmp utility-mib set instance | **15023**
- request support information | **15024**
- request system application (Junos OS Evolved) | **15042**

- request system autorecovery state | **15044**
- request system boot console | **15047**
- request system certificate add | **15048**
- request system commit server pause | **15050**
- request system commit server queue cleanup | **15051**
- request system commit server start | **15053**
- request system configuration database resize | **15054**
- request system configuration rescue delete | **15057**
- request system configuration rescue save | **15059**
- request system convert-json-configuration | **15060**
- request system debug-info | **15062**
- request system decrypt password | **15066**
- request system disable unified-services | **15068**
- request system download abort | **15069**
- request system download clear | **15071**
- request system download pause | **15072**
- request system download resume | **15074**
- request system download start | **15076**
- request system enable unified-services | **15078**
- request system filesystem encryption enable | **15079**
- request system filesystem encryption keys delete | **15082**
- request system firmware downgrade optics | **15085**
- request system firmware reload (Junos OS Evolved) | **15087**
- request system firmware upgrade | **15089**
- request system firmware upgrade (Junos OS Evolved) | **15099**
- request system firmware upgrade poe | **15102**
- request system halt | **15105**
- request-system-integrity-attestation | **15112**
- request-system-integrity-log-retrieval | **15114**
- request system license update | **15115**
- request system logout | **15117**
- request system malware-scan | **15119**

- request system partition abort | **15122**
- request system partition compact-flash | **15125**
- request system partition hard-disk | **15127**
- request system power-off | **15131**
- request system power-off (SRX Series) | **15135**
- request system process terminate | **15138**
- request system reboot (Junos OS) | **15140**
- request system reboot (Junos OS Evolved) | **15151**
- request system reboot (Junos OS with Upgraded FreeBSD) | **15154**
- request system reboot (SRX Series) | **15159**
- request system reboot usb (SRX Series) | **15160**
- request system recover | **15161**
- request system restore state path | **15165**
- request system save state path | **15166**
- request system scripts add (Chassis Cluster) | **15167**
- request system scripts add | **15171**
- request system scripts convert | **15172**
- request system scripts delete | **15175**
- request system scripts event-scripts reload | **15176**
- request system scripts refresh-from | **15177**
- request system scripts rollback | **15181**
- request system scripts synchronize | **15182**
- request system services dhcp | **15185**
- request system snapshot (Junos OS Evolved) | **15187**
- request system snapshot (Junos OS with FreeBSD Prior to Release 10) | **15190**
- request system snapshot (Junos OS with Upgraded FreeBSD) | **15203**
- request system snapshot (Maintenance) | **15208**
- request system software abort | **15212**
- request system software abort in-service-upgrade (ICU) | **15215**
- request system software add (Junos OS) | **15216**
- request system software add (Junos OS Evolved) | **15233**
- request system software add (Maintenance) | **15238**

- request system software add optional://auto-snapshot | **15240**
- request system software add restart | **15242**
- request system software configuration-backup | **15245**
- request system software configuration-restore | **15247**
- request system software delete-backup | **15248**
- request system software delete (Junos OS) | **15250**
- request system software delete (Junos OS Evolved) | **15254**
- request system software download | **15258**
- request system software format-qfabric-backup | **15260**
- request system software in-service-upgrade | **15262**
- request system software in-service-upgrade (Maintenance) | **15282**
- request system software in-service-upgrade (MX Series 5G Universal Routing Platforms and EX9200 Switches) | **15288**
- request system software nonstop-upgrade (QFabric Systems) | **15311**
- request system software nonstop-upgrade | **15319**
- request system software recover-from-restore-point | **15332**
- request system software recovery-package | **15334**
- request system software restore-point | **15338**
- request system software rollback (Junos OS) | **15340**
- request system software rollback (Junos OS Evolved) | **15345**
- request system software sync | **15348**
- request system software system-backup | **15353**
- request system software validate in-service-upgrade | **15355**
- request system software validate (Junos OS) | **15360**
- request system software validate (Junos OS Evolved) | **15365**
- request system software validate on (Junos OS with Upgraded FreeBSD) | **15368**
- request system software validate-restart (Junos OS Evolved) | **15371**
- request system storage cleanup (Junos OS) | **15374**
- request system storage cleanup (Junos OS Evolved) | **15388**
- request system storage cleanup (SRX Series) | **15400**
- request system yang add | **15404**
- request system yang delete | **15408**

- request system yang disable | 15410
- request system yang enable | 15413
- request system yang update | 15414
- request system yang validate | 15417
- request system zeroize | 15419
- request system zeroize (Junos OS) | 15421
- request system zeroize (Junos OS Evolved) | 15425
- request system zeroize ztp-option secure | 15430
- request-timeout | 15432
- request-type | 15433
- request unified-edge tdf call-trace clear | 15435
- request unified-edge tdf call-trace show | 15436
- request unified-edge tdf call-trace start | 15440
- request unified-edge tdf call-trace stop | 15443
- request validation policy | 15445
- request virtual-chassis device-reachability | 15447
- request virtual-chassis member-id delete (MX Series Virtual Chassis) | 15451
- request virtual-chassis member-id set | 15452
- request virtual-chassis mode | 15455
- request virtual-chassis reactivate | 15461
- request virtual-chassis recycle | 15462
- request virtual-chassis renumber | 15464
- request virtual-chassis routing-engine master switch | 15465
- request virtual-chassis vc-port | 15468
- request virtual-chassis vc-port (Dedicated VCP) | 15472
- request virtual-chassis vc-port delete (MX Series Virtual Chassis) | 15474
- request virtual-chassis vc-port diagnostics optics | 15476
- request virtual-chassis vc-port set (MX Series Virtual Chassis) | 15478
- request virtual-network-functions | 15481
- request vmhost cleanup | 15482
- request vmhost cleanup (NFX Series) | 15485
- request vmhost copy jnode-to-vjunos | 15486

- request vmhost copy vjunos-to-jnode | **15487**
- request vmhost file-copy | **15488**
- request vmhost file-copy (NFX Series) | **15490**
- request vmhost halt | **15492**
- request vmhost halt (NFX Series) | **15494**
- request vmhost hard-disk-test | **15495**
- request vmhost jdm add (In-Chassis Model) | **15497**
- request vmhost jdm delete (In-Chassis Model) | **15499**
- request vmhost jdm login (In-Chassis Model) | **15500**
- request vmhost jdm start (In-Chassis Model) | **15501**
- request vmhost jdm stop (In-Chassis Model) | **15502**
- request vmhost local-file-copy | **15503**
- request vmhost mode | **15506**
- request-vmhost-power-cycle-cb | **15507**
- request-vmhost-power-cycle-re | **15516**
- request vmhost power-off | **15525**
- request vmhost power-off (NFX Series) | **15526**
- request vmhost power-on | **15528**
- request vmhost reboot | **15529**
- request vmhost reboot (NFX Series) | **15532**
- request vmhost remote-file-copy | **15535**
- request vmhost snapshot | **15539**
- request vmhost software abort in-service-upgrade | **15542**
- request vmhost software add | **15543**
- request vmhost software add (NFX Series) | **15548**
- request vmhost software in-service-upgrade | **15551**
- request vmhost software rollback | **15556**
- request vmhost software validate | **15559**
- request vmhost storage | **15561**
- request vmhost zeroize | **15563**
- request vmhost zeroize ztp-option secure | **15566**
- request vpls-switchover | **15568**

- request wlan access-point packet capture | **15570**
- requested-ip-network-match (DHCP Local Server) | **15572**
- required-depth | **15574**
- reroute-address | **15575**
- reserve ingress-egress-mcast | **15577**
- reserved-bandwidth | **15579**
- reserved (Security IDP Custom Attack) | **15580**
- reset (Security IDP) | **15582**
- reset-tracking-bit | **15583**
- reset-count | **15584**
- resignal-minimum-bandwidth | **15585**
- resiliency | **15586**
- resilient-hash | **15589**
- resize (Ephemeral Database) | **15590**
- resize-ephemeral-default | **15592**
- resolution | **15594**
- resolution-map | **15597**
- resolution-ribs | **15599**
- resolve | **15600**
- resolve-vpn | **15602**
- resource (Analytics) | **15604**
- resource-category jtree (Resource Monitor) | **15605**
- resource-cleanup | **15607**
- resource (Interfaces) | **15608**
- resource-list (Network Management) | **15610**
- resource-list (Telemetry) | **15612**
- resource-manager | **15614**
- resource-manager (Forwarding-options) | **15616**
- resource-monitor | **15618**
- resource-profiles (Analytics) | **15620**
- resource-watch | **15622**
- resources | **15625**

- respond-bad-spi (Services IKE Policy) | **15626**
- responder (performance-monitoring) | **15628**
- response-cache-timeout (RADIUS Client) | **15629**
- rest | **15631**
- restart | **15633**
- restart chassis-control | **15648**
- restart-duration (Protocols) | **15651**
- restart-duration (Multicast Snooping) | **15653**
- restart-duration (Routing Options) | **15655**
- restart extensible-subscriber-services | **15656**
- restart (JDM) | **15657**
- restart (Junos OS Evolved) | **15660**
- restart-path | **15662**
- restart-time (BGP Graceful Restart) | **15663**
- restart-time (Fabric Control) | **15665**
- restarter (Graceful Restart for BGP Restarter) | **15667**
- restore-interval | **15669**
- restricted-queues | **15671**
- retain | **15672**
- retain-hold-resource | **15674**
- reth-count (Chassis Cluster) | **15675**
- retransmission-attempt (DHCP Client) | **15677**
- retransmission-attempt (DHCPv6 Client) | **15678**
- retransmission-count-established (L2TP) | **15680**
- retransmission-count-not-established (L2TP) | **15681**
- retransmission-interval | **15683**
- retransmission-interval (DHCP) | **15684**
- retransmission-interval (DHCP Client) | **15685**
- retransmit-interval (Services) | **15687**
- retransmit-timer | **15689**
- retransmit-timer (Dynamic Router Advertisement) | **15690**
- retry (Access LDAP) | **15691**

- retry (Access RADIUS) | **15693**
- retry (Adaptive Services Interface) | **15694**
- retry-count (Security Log) | **15696**
- retry-count (SNMP) | **15697**
- retry-count (FPC Error Condition) | **15698**
- retry-count (Chassis Cluster) | **15700**
- retry-count (Event Policy) | **15701**
- retry-delay | **15703**
- retry (Event Policy) | **15704**
- retry-interval | **15706**
- retry-interval (Chassis Cluster) | **15707**
- retry-limit | **15709**
- retry-limit (Protocols Link Management) | **15710**
- retry-options | **15712**
- retry (RADIUS Server) | **15714**
- retry (Services Flow Collector) | **15716**
- retry-timer | **15717**
- reverse | **15718**
- reverse-oif-mapping | **15720**
- reverse-route-packet-mode-vr | **15721**
- reverse-tcp-mss | **15722**
- revert-interval (Access) | **15724**
- revert-interval (Access LDAP) | **15725**
- revert-interval (Access RADIUS) | **15727**
- revert-interval (RADIUS Server) | **15728**
- revert-time (Protocols Layer 2 Circuits) | **15729**
- revert-timer | **15732**
- revertive | **15734**
- revision-level | **15735**
- revocation-check | **15737**
- revocation-check (Security PKI) | **15739**
- rewrite-rules (Definition) | **15741**

- rewrite-rules (CoS Host Outbound Traffic) | **15743**
- rewrite-rules (Definition) | **15744**
- rewrite-rules (Definition) | **15746**
- rewrite-rules (Dynamic CoS Interfaces) | **15748**
- rewrite-rules (Interfaces) | **15749**
- rewrite-rules (Interfaces) | **15751**
- rewrite-value (Fibre Channel Interfaces) | **15753**
- rfc-compliant (NETCONF) | **15755**
- rfc2544 | **15757**
- rfc2544-benchmarking | **15759**
- rfc6514-compliant-safi129 (Protocols BGP) | **15763**
- rib (General) | **15764**
- rib-group (Protocols BGP) | **15767**
- rib-group (Protocols DVMRP) | **15769**
- rib-group (Protocols IS-IS) | **15770**
- rib-group (Protocols MSDP) | **15772**
- rib-group (Protocols OSPF) | **15774**
- rib-group (Protocols PIM) | **15775**
- rib-group (Protocols RIP) | **15777**
- rib-group (Routing Options) | **15778**
- rib-group (Source Packet Routing) | **15780**
- rib-groups | **15781**
- rib-list | **15783**
- rib (Multitopology Routing) | **15785**
- rib (Protocols BGP) | **15787**
- rib (Route Resolution) | **15789**
- rib-service (programmable-rpd) | **15790**
- rib-sharding | **15792**
- rift | **15794**
- ring-id | **15797**
- ring-protection-link-end | **15799**
- ring-protection-link-owner | **15800**

- riot-loopback | **15801**
- rip | **15803**
- ripng | **15804**
- rising-event-index | **15805**
- rising-threshold (Health Monitor) | **15807**
- rising-threshold (RMON) | **15808**
- rmon | **15810**
- rng | **15812**
- roaming-mode | **15813**
- robust-count (IGMP Snooping) | **15815**
- robust-count (MLD Snooping) | **15816**
- robust-count (Protocols IGMP) | **15818**
- robust-count (Protocols IGMP AMT) | **15820**
- robust-count (Protocols MLD) | **15821**
- robustness-count | **15823**
- role | **15824**
- role (MX Series Virtual Chassis) | **15827**
- role (Protocols OpenFlow) | **15829**
- rollback | **15830**
- root-authentication | **15832**
- root-ca (Services) | **15834**
- root-logical-system | **15835**
- root-login (JDM) | **15837**
- root-streaming | **15838**
- root-address | **15840**
- route (Access) | **15841**
- route (Access Internal) | **15843**
- route (Access-Internal) | **15845**
- route-accounting | **15846**
- route-active-on | **15848**
- route-attributes | **15850**
- route-change-timeout | **15852**

- route (chassis) | **15853**
- route (Diameter Network Element) | **15854**
- route-distinguisher | **15856**
- route-distinguisher-id | **15860**
- route-filter | **15861**
- route-filter-list | **15863**
- route (Interfaces) | **15868**
- route-localization | **15869**
- route-monitoring | **15871**
- route-record | **15874**
- route-refresh-priority | **15876**
- route-suppression (DHCP Local Server and Relay Agent) | **15877**
- route-target-filter | **15879**
- route-target (Protocols BGP) | **15881**
- route-target (Protocols MVPN) | **15883**
- route-timeout (Protocols RIP) | **15885**
- route-timeout (Protocols RIPng) | **15887**
- route-distinguisher-id-use-router-id | **15889**
- router-advertisement | **15890**
- router-advertisement (Dynamic Profiles) | **15895**
- router-advertisement-guard | **15897**
- router-discovery | **15900**
- router-id | **15902**
- routing-engine | **15904**
- routing-engine (Chassis) | **15905**
- routing-engine (Chassis Redundancy) | **15908**
- routing-engine-profile | **15909**
- routing-engine-services | **15911**
- routing-engine (SNMP Resource Level) | **15912**
- routing-header | **15914**
- routing-instance (Services) | **15915**
- routing-instance (SNMPv3) | **15917**

- routing-instance (SNMP) | **15918**
- routing-instance (Forwarding-options Analyzer) | **15920**
- routing-instance | **15921**
- routing-instance (Interfaces) | **15923**
- routing-instance-access | **15924**
- routing-instance (Access LDAP) | **15925**
- routing-instance (Access RADIUS) | **15927**
- routing-instance (Advanced Policy-Based Routing) | **15928**
- routing-instance (Flow Monitoring) | **15929**
- routing-instance (Diameter Peer) | **15931**
- routing-instance (Diameter Transport) | **15932**
- routing-instance (Event Scripts) | **15934**
- routing-instance (JET Scripts) | **15935**
- routing-instance (Junos Telemetry Interface) | **15937**
- routing-instance (L2TP Destination) | **15938**
- routing-instance (L2TP Tunnel Destination) | **15940**
- routing-instance-list (TWAMP) | **15941**
- routing-instance-name | **15943**
- routing-instance-name (circuit-id) | **15944**
- routing-instance-name (DHCP Local Server) | **15945**
- routing-instance-name (DHCP Relay Agent) | **15947**
- routing-instance-name (Static Subscribers) | **15949**
- routing-instance (PCC Action Profiles) | **15951**
- routing-instance (Port Mirroring) | **15953**
- routing-instance (PPPoE Service Name Tables) | **15954**
- routing-instance (RPM) | **15956**
- routing-instance (Security Content Security) | **15958**
- routing-instance (Security Destination NAT) | **15959**
- routing-instance (Security Source NAT) | **15960**
- routing-instance (Syslog) | **15962**
- routing-instance (System Scripts) | **15963**
- routing-instance (Tunnel Profile) | **15965**

- routing-instances (Logical Systems) | **15966**
- routing-instances (JDAF) | **15967**
- routing-instances | **15969**
- routing-instances (RPM) | **15970**
- routing-instances (CoS) | **15971**
- routing-instances (Dynamic Profiles) | **15973**
- routing-instances (Multiple Routing Entities) | **15976**
- routing-interface | **15978**
- routing-options | **15980**
- routing-options (Dynamic Profiles) | **15981**
- routing-options (Junos Node Slicing) | **15983**
- routing-options-max-ecmp | **15985**
- routing-options (QFabric System) | **15987**
- routing-service (Dynamic Profiles) | **15988**
- routing-services (Enhanced Subscriber Management) | **15991**
- routing (System Processes) | **15993**
- routing-engine-mode (TLB) | **15995**
- rp | **15996**
- rp-register-policy | **16000**
- rp-set | **16001**
- rpc | **16003**
- rpc-program-number | **16004**
- rpc-program-number (Applications) | **16005**
- rpf-check | **16007**
- rpf-check (Dynamic Profiles) | **16009**
- rpf-check-policy (Routing Options) | **16011**
- rpf-check-policy (Routing Options RPF) | **16012**
- rpf-loose-mode-discard | **16014**
- rpf-selection | **16015**
- rpf-vector (PIM) | **16017**
- rpm (Interfaces) | **16019**
- rpm-scale | **16020**

- rpm (Services) | **16023**
- rpm-tracking | **16038**
- rpt-spt | **16040**
- rsh | **16041**
- rst-invalidate-session | **16043**
- rst-sequence-check | **16044**
- rstp | **16046**
- rsvp | **16050**
- rsvp-error-hold-time | **16056**
- rsvp-te (Protocols MVPN) | **16058**
- rsvp-te (Routing Instances Provider Tunnel) | **16059**
- rsvp-te (Routing Instances Provider Tunnel Selective) | **16061**
- rsvp-te (Routing Options) | **16062**
- rtf-prefix-list | **16064**
- rtlog (Next Gen Services Global System Logging) | **16066**
- rtp | **16068**
- rts | **16069**
- rts-polarity | **16071**
- rtsp | **16072**
- rule | **16074**
- rule (AACL Rule Set) | **16076**
- rule-activation-time | **16078**
- rule (Advanced Policy-Based Routing) | **16080**
- rule (Application Firewall) | **16083**
- rule (Captive Portal Content Delivery) | **16085**
- rule (casb-policy) | **16087**
- rule-deactivation-time | **16089**
- rule (Destination NAT Next Gen Services) | **16091**
- rule (HTTP Header Enrichment for Tag Rule Set) | **16093**
- rule (IDS MS-DPC) | **16094**
- rule (IDS MS-MPC) | **16096**
- rule (IP Reassembly) | **16100**

- rule (LRF) | **16101**
- rule (PCP) | **16103**
- rule (Security Destination NAT) | **16105**
- rule (Security Exempt Rulebase) | **16107**
- rule (Security IPS Rulebase) | **16109**
- rule (Security Source NAT) | **16111**
- rule (Security Static NAT) | **16113**
- rule (Services CoS) | **16115**
- rule (Services CoS Next Gen Services) | **16117**
- rule (Services NAT) | **16120**
- rule (Services Stateful Firewall) | **16122**
- rule-session-count-alarm (Security Destination NAT Rule Set) | **16124**
- rule-session-count-alarm (Security Source NAT Rule Set) | **16125**
- rule-session-count-alarm (Security Static NAT Rule Set) | **16126**
- rule-set | **16128**
- rule-set (Captive Portal Content Delivery) | **16129**
- rule-set (Security Destination NAT) | **16131**
- rule-set (Security Source NAT) | **16133**
- rule-set (Security Static NAT) | **16135**
- rule-set (Services AACL) | **16138**
- rule-set (Services CoS) | **16139**
- rule-set (Services CoS Next Gen Services) | **16140**
- rule-set (Services IDS) | **16142**
- rule-set (Services NAT) | **16143**
- rule-set (Services Stateful Firewall) | **16144**
- rule-set (Softwire) | **16146**
- rule-set (Softwires Next Gen Services) | **16147**
- rule-sets (CoS AppQoS) | **16149**
- rule-sets (Security Application Firewall) | **16151**
- rule (Softwire) | **16154**
- rule (Source NAT Next Gen Services) | **16156**
- rulebase-exempt | **16157**

- rulebase-ips | **16159**
- run | **16162**
- run-length | **16163**
- rx-buffers | **16164**
- rx-connect-speed-when-equal (L2TP LAC) | **16166**
- rx-enable | **16168**
- rx-max-duration | **16169**
- rx-window-size (L2TP) | **16171**

r2cp

IN THIS SECTION

- Syntax | **13990**
- Hierarchy Level | **13991**
- Description | **13991**
- Options | **13991**
- Required Privilege Level | **13991**
- Release Information | **13991**

Syntax

```
r2cp {
    command binary-file-path;
    disable;
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the Radio-to-Router Control Protocol (R2CP) used to exchange dynamic metric changes in the network that routers use to update the OSPF topologies.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the Radio-to-Router Control Protocol process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

radio-router (Interfaces)

IN THIS SECTION

Syntax | 13992
Hierarchy Level | 13992

- Description | 13992
- Options | 13992
- Required Privilege Level | 13993
- Release Information | 13993

Syntax

```
radio-router {
    bandwidth number;
    credit {
        interval number;
    }
    data-rate number;
    latency number;
    quality number;
    resource number;
    threshold number;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Point-to-Point Protocol over Ethernet (PPPoE)-based radio-to-router protocols include messages that define how an external system will provide the device with timely information about the quality of a link's connection. They also include a flow control mechanism to indicate how much data the device can forward. The device can then use the information provided in the PPPoE messages to dynamically adjust the interface speed of PPP links.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

radius (Access Profile)

IN THIS SECTION

- Syntax | 13993
- Hierarchy Level | 13996
- Description | 13996
- Options | 13996
- Required Privilege Level | 13997
- Release Information | 13997

Syntax

```
radius {
    accounting-server [ ip-address ];
    attributes {
        exclude
            attribute-name packet-type;
            standard-attribute number {
                packet-type [ access-request | accounting-off | accounting-on | accounting-start
```

```
| accounting-stop ];
            }
            vendor-id id-number {
                vendor-attribute vsa-number {
                    packet-type [ access-request | accounting-off | accounting-on | accounting-
start | accounting-stop ];
                }
            }
        }
        ignore {
            dynamic-iflset-name;
            framed-ip-netmask;
            idle-timeout;
            input-filter;
            logical-system-routing-instance;
            output-filter;
            session-timeout;
            standard-attribute number;
            vendor-id id-number {
                vendor-attribute vsa-number;
            }
        }
    }
    authentication-server [ ip-address ];
    options {
        accounting-session-id-format (decimal | description);
        calling-station-id-delimiter delimiter-character;
        calling-station-id-format {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            nas-identifier;
        }
        chap-challenge-in-request-authenticator;
        client-accounting-algorithm (direct | round-robin);
        client-authentication-algorithm  (direct | round-robin);
        coa-dynamic-variable-validation;
        ethernet-port-type-virtual;
        interface-description-format {
            exclude-adapter;
            exclude-channel;
            exclude-sub-interface;
        }
```

```
        ip-address-change-notify message;
        juniper-access-line-attributes;
        nas-identifier identifier-value;
        nas-port-extended-format {
            adapter-width width;
            ae-width width;
            port-width width;
            slot-width width;
            stacked-vlan-width width;
            vlan-width width;
            atm {
                adapter-width width;
                port-width width:
                slot-width width;
                vci-width width:
                vpi-width width;
            }
        }
        nas-port-id-delimiter delimiter-character;
        nas-port-id-format {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            interface-text-description;
            nas-identifier;
            order {
                agent-circuit-id;
                agent-remote-id;
                interface-description;
                interface-text-description;
                nas-identifier;
                postpend-vlan-tags;
            }
            postpend-vlan-tags;
        }
        nas-port-type {
            ethernet {
                port-type;
            }
        }
        override {
            calling-station-id remote-circuit-id;
            nas-ip-address tunnel-client-gateway-address;
```

```
            nas-port tunnel-client-nas-port;
            nas-port-type tunnel-client-nas-port-type;
        }
        remote-circuit-id-delimiter;
        remote-circuit-id-fallback;
        remote-circuit-id-format {
            agent-circuit-id;
            agent-remote-id;
        }
        revert-interval interval;
        service-activation {
            dynamic-profile (optional-at-login | required-at-login);
            extensible-service (optional-at-login | required-at-login);
        }
        vlan-nas-port-stacked-format;
    }
    preauthentication-server ip-address;
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure the RADIUS parameters that the router uses for AAA authentication and accounting for subscribers.

Options

accounting-server

(MX Series only) Specify a list of the RADIUS accounting servers used for accounting for DHCP, L2TP, and PPP clients.

- **Values:** *ip-address*—IP version 4 (IPv4) address.

authentication-server

(SRX Series only) Specify a list of the RADIUS authentication servers used to authenticate DHCP, L2TP, and PPP clients. The servers in the list are also used as RADIUS dynamic-request servers, from which the router accepts and processes RADIUS disconnect requests, CoA requests, and dynamic service activations and deactivations.

- **Values:** *ip-address*—IPv4 address.

preauthentication-server

(MX Series only) Starting in Junos OS Release 13.3, specify the RADIUS preauthentication server, which is used for the LLID service.

> **NOTE**: You cannot configure this statement if the Calling-Station-ID attribute is excluded from RADIUS Access-Request messages by the `exclude` statement.

- **Values:** *ip-address*—IPv4 address.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Authentication and Accounting Basic Configuration

RADIUS Logical Line Identification

radius-accounting

IN THIS SECTION

- Syntax | 13998
- Hierarchy Level | 13998

- Description | 13998
- Options | 13998
- Required Privilege Level | 13999
- Release Information | 13999

Syntax

```
radius-accounting {
   session-drop;
}
```

Hierarchy Level

```
[edit security nat source]
```

Description

Configure RADIUS accounting options for NAT.

NOTE: If you have configured both syslog and RADIUS server for the same stream, NAT PBA logs are sent to the RADIUS server and other logs are sent to the syslog server.

If you want to receive all logs (including NAT PBA and flow logs), you must use different streams for the RADIUS server and syslog server.

Options

session-drop Drop Source NAT sessions until RADIUS accounting succeeds.

When you enable the `session-drop` statement at the `[edit security nat source radius-accounting]` hierarchy, subscribers cannot use a port block until RADIUS logging succeeds for the port block. The device drops the subscriber's NAT PBA sessions until it receives a positive

acknowledgement for the RADIUS log sent. The port block is released when the RADIUS logging timeout or a negative acknowledgement is received from the RADIUS accounting server.

When you disable this command, subscribers can continue using the allocated port block even if RADIUS logging fails for the port block.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

radius-coa (Adjustment Control Profiles)

IN THIS SECTION

- Syntax | 13999
- Hierarchy Level | 14000
- Description | 14000
- Options | 14000
- Required Privilege Level | 14000
- Release Information | 14001

Syntax

```
radius-coa {
    priority priority;
    algorithm algorithm;
}
```

Hierarchy Level

```
[edit class-of-service adjustment-control-profiles profile-name application]
```

Description

Configure the shaping rate adjustment controls for the RADIUS CoA application.

Options

priority Priority of the RADIUS CoA application in the adjustment control profile.

- **Range:** 1 through 10; 1 being the highest priority.
- **Default:** 1

algorithm Rate adjustment algorithm used by the RADIUS CoA application.

- Values:
  - adjust-never—Do not perform rate adjustments.
  - adjust-always—Adjust the shaping rate unconditionally.
  - adjust-less—Adjust the shaping rate if it is less than the configured value.
  - adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
  - adjust-greater—Adjust the shaping rate if it is greater than the configured value.
  - adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.
- **Default:** adjust-always

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

CoS Adjustment Control Profiles Overview

Configuring CoS Adjustment Control Profiles

Verifying the CoS Adjustment Control Profile Configuration

adjustment-control-profiles | **1961**

application (Adjustment Control Profiles) | **2301**

radius-disconnect

IN THIS SECTION

- Syntax | **14001**
- Hierarchy Level | **14002**
- Description | **14002**
- Options | **14002**
- Required Privilege Level | **14002**
- Release Information | **14002**

Syntax

```
radius-disconnect {
     client-address {
        secret (RADIUS) password;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure a disconnect server that listens on a configured User Datagram Protocol (UDP) port for disconnect messages from a configured client and processes these disconnect messages.

Options

client-address A valid IP address configured on one of the router interfaces.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

RADIUS Authentication and Accounting Basic Configuration

Example: Configuring CHAP Authentication with RADIUS

Configuring RADIUS Authentication for L2TP

radius-disconnect (DHCP Local Server)

IN THIS SECTION

- Syntax | 14003
- Hierarchy Level | 14003
- Description | 14004
- Default | 14004
- Required Privilege Level | 14004
- Release Information | 14004

Syntax

```
radius-disconnect;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure trigger],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure trigger],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure trigger],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure trigger],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure trigger],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure
trigger],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
reconfigure trigger],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure trigger],
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure
trigger],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
```

```
reconfigure trigger],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
reconfigure trigger],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure trigger],
[edit system services dhcp-local-server reconfigure trigger],
[edit system services dhcp-local-server dhcpv6 reconfigure trigger],
[edit system services dhcp-local-server group group-name reconfigure trigger],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure trigger]
```

Description

Configure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces to be reconfigured when a RADIUS-initiated disconnect is received by the DHCP client or group of clients. A group configuration takes precedence over a DHCP local server configuration.

Default

The client is deleted when a RADIUS-initiated disconnect is received.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview

Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect

radius-disconnect-port

IN THIS SECTION

- Syntax | 14005
- Hierarchy Level | 14005
- Description | 14005
- Options | 14005
- Required Privilege Level | 14006
- Release Information | 14006

Syntax

```
radius-disconnect-port port-number;
```

Hierarchy Level

```
[edit access]
```

Description

Specify a port number on which to contact the RADIUS disconnect server. Most RADIUS servers use port number 1700.

Options

port-number—The server port to which disconnect requests from the RADIUS client are sent. The L2TP network server, which accepts these disconnect requests, is the server.

NOTE: The Junos OS accepts disconnect requests only from the client address configured at the `[edit access radius-disconnect client client-address]` hierarchy level.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring RADIUS Authentication for L2TP

radius-flow-tap

IN THIS SECTION

- Syntax | 14006
- Hierarchy Level | 14007
- Description | 14008
- Options | 14008
- Required Privilege Level | 14009
- Release Information | 14009

Syntax

```
radius-flow-tap {
    forwarding-class class-name;
    interfaces interface-name;
    logical-system logical-system-name name <routing-instance routing-instance>;
    multicast-interception;
    policy policy-name {
        inet
```

```
{
            drop-policy rule-name {
                from
{
                    apply-groups group-name;
                    apply-groups-except group-name;
                    destination-address address;
                    destination-port port-number;
                    dscp dscp-value;
                    protocol protocol;
                    source-address address;
                    source-port port-number;
                }
            }
        }
        inet6 {
            drop-policy rule-name {
                from {
                    apply-groups group-name;
                    apply-groups-except group-name;
                    destination-address address;
                    destination-port port-number;
                    dscp dscp-value;
                    protocol protocol;
                    source-address address;
                    source-port port-number;
                }
            }
        }
    }
    snmp (
        notify-targets ip-address;
    }
    routing-instance routing-instance-name;
    source-ipv4-address ipv4-address;
)
```

Hierarchy Level

```
[edit services]
```

Description

Configure the radius-flow-tap service for subscriber secure policy mirroring. Both RADIUS-initiated and Dynamic Tasking Control Protocol (DTCP)-initiated mirroring are supported.

Starting in Junos OS Release 17.3R1, the radius-flow-tap service can run concurrently on the same router with the FlowTapLite service. The FlowTapLite service is a version of the flow-tap service (`[edit services flow-tap]`) that is configured only on tunnel interfaces on MX Series routers.

In earlier releases, the radius-flow-tap and FlowTapLite services cannot run concurrently on an MX Series router, preventing you from running FlowTapLite monitoring and subscriber secure policy mirroring at the same time.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

forwarding-class *class-name*

Specify the forwarding class that is applied to mirrored packets sent to a mediation device.

interfaces *interface-name*

Assign virtual tunnel interfaces to mirror the interfaces created by extensible subscriber services manager (ESSM)

If a currently used tunnel interface is deleted from the pool of interfaces, the active mirroring sessions are redistributed from the deleted interface to other tunnel interfaces in the pool. Also, when a new tunnel interface is added into the pool, the service adds the new interface to the list of interfaces available for new mirroring sessions or for existing sessions transferred from a failed interface.

logical-system *logical-system-name*

Specify the logical system that is used to send mirrored packets to a mediation device for subscriber secure policy traffic mirroring. When you specify a logical system, you must also specify a routing instance.

- **Default:** Logical system `default`

multicast-interception

Enables subscriber secure policy to mirror IPv4 multicast traffic sent to subscribers. It enables the mirroring of multicast traffic for all subscribers on the chassis.

Mirroring of multicast traffic is supported only for subscribers in the default logical system.

routing-instance *routing-instance-name*

Specify the routing instance that is used to send mirrored packets to a mediation device for subscriber secure policy traffic mirroring.

- **Default:** Routing instance `default`

snmp notify-targets *ip-address*

Specify the IP address for a target mediation device (trap target) to receive SNMPv3 encrypted trap notifications subscriber secure policy mirroring trap. Only these configured targets can receive the notifications. This is required for secure SNMPv3 notifications for subscriber secure policy mirroring. If you configure multiple targets, you must configure them one at a time.

source-ipv4-address *ipv4-address*

Specify the source IPv4 address used in the IP header that is prepended to mirrored packets sent to a mediation device.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

`logical-system` option added in Junos OS Release 15.1R3 for enhanced subscriber management.

`multicast-interception` option added in Junos OS Release 11.4.

`snmp notify-targets` option added in Junos OS Release 16.1R1.

`routing-instance` option added in Junos OS Release 15.1R3 for enhanced subscriber management.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, the radius-flow-tap service can run concurrently on the same router with the FlowTapLite service.

RELATED DOCUMENTATION

Configuring Support for Subscriber Secure Policy Mirroring

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

Disabling RADIUS-Initiated Subscriber Secure Policy Mirroring

Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic

Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview

Subscriber Secure Policy Overview

radius-options (Access)

IN THIS SECTION

- Syntax | 14010
- Hierarchy Level | 14010
- Description | 14011
- Required Privilege Level | 14011
- Release Information | 14011

Syntax

```
radius-options {
    interim-rate interim-rate;
    request-rate rate;
    revert-interval seconds;
    timeout-grace seconds;
    unique-nas-port {
        chassis-id chassis-id;
        chassis-id-width chassis-id-width;
    }
}
```

Hierarchy Level

```
[edit access]
[edit access profile profile-name]
```

Description

Configure RADIUS options that apply to all RADIUS servers globally.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable

Configuring Authentication and Accounting Parameters for Subscriber Access

radius-options (Interfaces)

IN THIS SECTION

- Syntax | 14012
- Hierarchy Level | 14012
- Description | 14012
- Required Privilege Level | 14012
- Release Information | 14013

Syntax

```
radius-options {
    nas-port-options nas-port-options-name {
        nas-port-extended-format {
            adapter-width width;
            ae-width width;
            port-width width;
            slot-width width;
            stacked;
            stacked-vlan-width width;
            vci-width width:
            vlan-width width;
            vpi-width width;
        }
        nas-port-type port-type;
        stacked-vlan-ranges (any | low-outer-tag-high-outer-tag),any;
        vlan-ranges (any | low-tag-high-tag);
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure RADIUS options to set the NAS-Port-Type (61) RADIUS IETF attribute, and an extended format for the NAS-Port (5) RADIUS IETF attribute, on a per-physical interface, per-VLAN, or per-stacked VLAN (S-VLAN) basis.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN

Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN

radius-options (System)

IN THIS SECTION

- Syntax | **14013**
- Hierarchy Level | **14014**
- Description | **14014**
- Options | **14014**
- Required Privilege Level | **14014**
- Release Information | **14014**

Syntax

```
radius-options {
    attributes {
        nas-id nas-id
        nas-ip-address ip-address;
    }
    enhanced-accounting;
    password-protocol mschap-v2;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure RADIUS options for the NAS-IP address for outgoing RADIUS packets and password protocol used in RADIUS packets.

Options

enhanced-accounting	Configure audit of TACACS+ or RADIUS authentication events such as access method, remote port, and access privileges.
nas-id *nas-id*	Value of NAS-ID in outgoing RADIUS packets.
nas-ip-address *ip-address*	IP address of the network access server (NAS) that requests user authentication.
password-protocol *mschap-v2*	Protocol MS-CHAPv2, used for password authentication and password changing.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

MS-CHAPv2 password protocol configuration option introduced in Junos OS Release 9.2.

MS-CHAPv2 password protocol configuration option introduced in Junos OS Release 9.2 for EX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Support for network access server (NAS) IPv6 address added in Junos OS Release 12.1X47-D15 for SRX1500, SRX5400, SRX5600, and SRX5800 devices.

Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

NOTE: The `radius-options` statement is not available on QFabric systems.

`enhanced-accounting` statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring MS-CHAPv2 for Password-Change Support

Configuring RADIUS System Accounting

radius-reachability

IN THIS SECTION

- Syntax | **14015**
- Hierarchy Level | **14016**
- Release Information | **14016**
- Description | **14016**
- Options | **14016**
- Required Privilege Level | **14016**

Syntax

```
radius-reachability {
 query-period seconds;
}
```

Hierarchy Level

```
[edit logical-systems name protocols dot1x authenticator],
[edit protocols dot1x authenticator]
```

Release Information

Statement introduced in Junos OS Release 21.4R1.

Description

Enable the RADIUS reachability feature to authenticate sessions that are in server fail state once the authentication server becomes reachable. The authenticator checks for reachability by triggering reauthentication for one of the authentications sessions that previously failed to reach the server.

When an authentication session fails because the server is not reachable, the session is placed in held state or in the server-fail VLAN, if it is configured. If RADIUS reachability is not enabled, these sessions will not reauthenticate until the reauthentication timer expires. The reauthentication interval is typically long, one hour by default, forcing the client to wait even after the server is reachable.

Options

query-period *seconds*

Query period interval.

- Default: 120 seconds.
- Range: 20 through 65535 seconds.

Required Privilege Level

routing

radius-realm

IN THIS SECTION

- Syntax | 14017
- Hierarchy Level | 14017
- Description | 14017
- Options | 14017
- Required Privilege Level | 14017
- Release Information | 14018

Syntax

```
radius-realm radius-realm-string;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication username-include]
```

Description

Specify that the user-defined RADIUS realm string is appended as a last piece to the username and used by RADIUS to direct the authentication request to a profile that does not allocates addresses.

Options

radius-realm-string—A string to describe the RADIUS realm.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

radius (Security Log)

IN THIS SECTION

- Syntax | 14018
- Hierarchy Level | 14019
- Description | 14019
- Options | 14019
- Required Privilege Level | 14020
- Release Information | 14020

Syntax

```
radius {
   ipaddr {
      backup backup;
      primary primary;
   }
   port port;
   secret secret;
}
```

Hierarchy Level

```
[edit logical-systems name security log stream],
[edit logical-systems name tenants name security log stream],
[edit security log stream],
[edit tenants name security log stream]
```

Description

Configure RADIUS accounting for security events by specifying the RADIUS accounting server.

Options

- `ipaddr` — Specify the IP address of the RADIUS server for accounting messages. Use the `primary` and `backup` options in the hierarchy `edit security log stream radius idappr` to configure the primary and backup RADIUS servers for authentication event logging. These options provide redundancy in case the primary RADIUS server is unavailable, ensuring that authentication events can still be logged to the backup RADIUS server.

  RADIUS accounting supports all the eight streams and each stream can have both a primary and backup radius server.

  Both the primary and backup RADIUS server shares same secret key and port number. Without a primary RADIUS server, you cannot configure the backup RADIUS server.

  - `primary` — IP address of the primary RADIUS server. Configure the primary IP address to enable `radius` stream.

  - `backup` — IP address of the backup RADIUS server. You can configure one backup RADIUS server per stream for the primary RADIUS server. The backup server handles log delivery if the primary server goes down.

    The backup server is optional. Backup server becomes active only if message delivery fails to primary server.

    If the backup server becomes unavailable, disable the backup server to switch back to the primary server.

- `port`- Port number of the RADIUS server. Configure the port number that receives RADIUS accounting messages.

  - **Default**: 1813

- `secret`- Configure a secret key on the RADIUS server. This secret key is used to calculate the MD5 hash for verifying the message's validity on the receiver side. The maximum length of the secret key is 256 characters.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

radius-server

IN THIS SECTION

- Syntax | 14020
- Hierarchy Level | 14021
- Description | 14021
- Options | 14021
- Required Privilege Level | 14025
- Release Information | 14025

Syntax

```
radius-server server-address {
    accounting-port port-number;
    accounting-retry number;
    accounting-timeout seconds;
    dynamic-request-port port-number;
    max-outstanding-requests value;
```

```
    port port-number;
    preauthentication-port port-number;
    preauthentication-secret password;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level

```
[edit access],
[edit access profile profile-name]
```

Description

Configure RADIUS for subscriber access management, L2TP, or PPP.

To configure multiple RADIUS servers, include multiple `radius-server` statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options

server-address IPv4 or IPv6 address of the RADIUS server.

accounting-port Configure the port number on which to contact the RADIUS accounting server.

> **NOTE**: Specifying the accounting port is optional, and port 1813 is the default. However, we recommend that you configure it in order to avoid confusion, as some RADIUS servers might refer to an older default.

- **Values:** *port-number*—Port number on which to contact the RADIUS accounting server. Most RADIUS servers use port 1813, as specified in RFC 2866.
- **Default:** 1813

accounting-retry

Configure the number of times the device retransmits RADIUS accounting messages when no response is received from the server. When you do not configure this statement, the number of retry attempts is determined by the `retry` statement.

> **NOTE**: To successfully set a retry limit for the accounting servers different from the authentication servers, you must configure both the `accounting-retry` and `accounting-timeout` statements . If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the `retry` and `timeout` statements.

> **NOTE**: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

- **Values:** *number*—Number of retry attempts.
- **Range:** 0 through 100
- **Default:** 0 (disabled)

accounting-timeout

Configure how long the local device waits to receive a response from a RADIUS accounting server before retransmitting the message. When you do not configure this statement, the length of the timeout is determined by the `timeout` statement.

> **NOTE**: To successfully set a timeout value for the accounting servers different from the authentication servers, you must configure both the `accounting-retry` and `accounting-timeout` statements . If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the `retry` and `timeout` statements.

> **NOTE**: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

- **Values:** *seconds*—Duration of timeout period.

- **Range:** 0 through 1000 seconds
- **Default:** 0 (disabled)

dynamic-request-port

Specify the port that the router monitors for dynamic (CoA) requests from the specified RADIUS servers. You can configure a port globally or for a specific access profile.

You must either use the default port for all RADIUS servers or configure the same nondefault port for all RADIUS servers. This rule applies at both the global access and access profile levels.

> **NOTE**: Any other configuration results in a commit check failure. Multiple port numbers—that is, different port numbers for different servers—are not supported.

- **Values:** *port-number*—Number of the monitored port.
- **Default:** 3799 (as specified in RFC 5176)

max-outstanding-requests

Configure the maximum number of outstanding requests for this RADIUS server. An increase in this value is immediate while a decrease is more gradual if the current number of outstanding requests exceeds the new value.

- **Values:** *requests*—Maximum number of outstanding requests for this RADIUS server.
- **Range:** 0 through 2000 outstanding requests per server
- **Default:** 1000 outstanding requests per server

port

Configure the port number on which to contact the RADIUS server.

- **Values:** *port-number*—Port number on which to contact the RADIUS server.
- **Default:** 1812 (as specified in RFC 2865)

preauthentication-port

Configure the port number on which to contact the RADIUS server for logical line identification (LLID) preauthentication requests. If you do not configure a separate UDP port for preauthentication purposes, the same UDP port that you configure for authentication messages by including the `port port-number` statement is used.

- **Values:** *port-number*—Port number used for preauthentication requests to contact the RADIUS server.

preauthentication-secret

Configure the password to use with the RADIUS server for LLID preauthentication requests. If you do not configure a separate UDP password for preauthentication purposes, the same password that you configure for authentication messages by including the `secret` *password* statement is used. The secret password used by the local router must match that used by the server.

- **Values:** *password*—Password to use. To include spaces enclose the character string in quotation marks.

retry

Specify the number of times that the device is allowed to attempt to contact a RADIUS authentication or accounting server. You can override the retry limit for accounting servers with the `accounting-retry` statement.

> **NOTE**: To successfully set a retry limit for the accounting servers different from the authentication servers, you must configure both the `accounting-retry` and `accounting-timeout` statements . If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the `retry` and `timeout` statements.

> **NOTE**: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

- **Values:** *attempts*—Number of times that the router is allowed to attempt to contact a RADIUS server.
- **Range:** 1 through 100
- **Default:** 3

routing-instance

Configure the routing instance used to send RADIUS packets to the RADIUS server.

- **Values:** *routing-instance-name*—Routing instance name.

source-address

Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. Support for IPv6 *source-address* was introduced in Junos OS Release 16.1.

- **Values:** *source-address*—Valid IPv4 or IPv6 address configured on one of the router or switch interfaces. On M Series routers only, the source address can be an IPv6 address and the UDP source port is 514.

timeout

Configure the amount of time that the local device waits to receive a response from RADIUS authentication and accounting servers. You can override the timeout value for accounting servers with the `accounting-timeout` statement.

> **NOTE**: To successfully set a timeout value for the accounting servers different from the authentication servers, you must configure both the `accounting-retry` and `accounting-timeout` statements . If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the `retry` and `timeout` statements.

> **NOTE**: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

- **Values:** *seconds*—Amount of time to wait.
- **Range:** 1 through 1000 seconds
- **Default:** 3 seconds

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`max-outstanding-requests` introduced in Junos OS Release 11.4.

`accounting-retry` and `accounting-timeout` introduced in Junos OS Release 14.1.

`dynamic-request-port` option added in Junos OS Release 14.2R1 for MX Series routers.

`preauthentication-port` and `preauthentication-secret` options added in Junos OS Release 15.1 for MX Series routers.

`accounting-port` introduced in Junos OS Release 13.2X50-D10 for EX Series switches with support for Enhanced Layer 2 software (ELS). It was introduced in Junos OS without ELS in the following releases: Junos OS Releases 12.3R10, 14.1X53-D25, and 15.1R4 for EX Series switches.

Support for IPv6 *`server-address`* introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

RADIUS Authentication and Accounting Basic Configuration

Configuring the PPP Password Authentication Protocol

Configuring RADIUS Authentication for L2TP

RADIUS Authentication

Configuring RADIUS-Initiated Dynamic Request Support

RADIUS Logical Line Identification

show network-access aaa statistics | **22428**

clear network-access aaa statistics | **3736**

radius-server (Access)

IN THIS SECTION

- Syntax | **14027**
- Hierarchy Level | **14027**
- Description | **14027**
- Options | **14027**
- Required Privilege Level | **14027**
- Release Information | **14027**

Syntax

```
radius-server server-address {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level

```
[edit access],
[edit access profile profile-name]
```

Description

Configure RADIUS for Layer 2 Tunneling Protocol (L2TP) or Point-to-Point Protocol (PPP) authentication.

To configure multiple RADIUS servers, include multiple `radius-server` statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options

server-address—Address of the RADIUS authentication server.

The remaining statements are explained separately.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

Firewall User Authentication Overview

radius-server-name

IN THIS SECTION

- Syntax | 14028
- Hierarchy Level | 14029
- Description | 14029
- Options | 14029
- Required Privilege Level | 14029
- Release Information | 14030

Syntax

```
radius-server-name hostname {
    accounting-port port-number;
    accounting-retry number;
    accounting-timeout seconds;
    dynamic-request-port port-number;
    dns-query-interval minutes;
    max-outstanding-requests value;
    port port-number;
    preauthentication-port port-number;
    preauthentication-secret password;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level

```
[edit access],
[edit access profile profile-name]
```

Description

Configure RADIUS for 802.1x authentication using a fully-qualified domain name (FQDN). Use an FQDN if the server name can translate to more than one address. The FQDN resolves dynamically at regular intervals so there is no need to modify the configuration for address changes.

The FQDN is first resolved at the time of configuration commit, and then at fixed intervals based on the configured value for the `dns-query-interval` statement.

You can combine the FQDN configuration with the round-robin algorithm to achieve load distribution among the available servers. Configure the round-robin algorithm at the [`edit access profile profile-name radius options`] hierarchy level.

To configure a RADIUS server using a fixed IP address, use the `radius-server` statement.

Options

hostname Configure the FQDN to identify the RADIUS server. The *hostname* parameter supports only one FQDN. To configure more than one FQDN, use multiple `radius-server-name` statements.

dns-query-interval Configure the interval for resolving an FQDN as the server address. The FQDN is resolved dynamically at fixed intervals based on the configured value.

- **Values:** *minutes*—Interval of the FQDN resolution in minutes.
- **Range:** 10 through 1440
- **Default:** 60

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

RADIUS Server Configuration for Authentication

radius-server (System)

IN THIS SECTION

- Syntax | 14030
- Hierarchy Level | 14031
- Description | 14031
- Options | 14031
- Required Privilege Level | 14033
- Release Information | 14033

Syntax

```
radius-server {
    server-address {
        accounting-port port-number;
        accounting-retry number;
        accounting-timeout seconds;
        dynamic-request-port number;
        max-outstanding-requests value;
        port number;
        preauthentication-port number;
        preauthentication-secret secret;
        retry number;
        routing-instance routing-instance-name;
        secret password;
```

```
        source-address source-address;
        timeout seconds;
        tls {
            ca-profiles name;
    }
    }
```

Hierarchy Level

```
[edit system]
```

Description

Configure the RADIUS authentication server for subscriber access management, Layer 2 Tunnelling Protocol (L2TP), or Point-to-Point Protocol (PPP).

To configure multiple RADIUS servers, include multiple `radius-server` *`server-address`* statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

> **NOTE**: The accounting-port and source-address options are not available on QFabric systems.

Options

server-address	Specify the IPv4 or IPv6 address of the RADIUS authentication server.
accounting-port ***port-number***	Configure the accounting port number on which to contact the RADIUS server. • **Range:** 1 through 65,335 • **Default:** 1813 (as specified in RFC 2866)
accounting-retry ***number***	Configure the number of accounting retry attempts. • **Range:** 0 through 100 attempts • **Default:** 0

accounting-timeout *seconds*

Configure the accounting request timeout period.

- **Range:** 0 through 1000 seconds
- **Default:** 0

dynamic-request-port *number*

Configure the RADIUS client dynamic request port number

- **Range:** 1 through 65535
- **Default:** 3799

max-outstanding-requests *value*

Configure the maximum number of outstanding requests in flight to the server.

- **Range:** 0 through 2000 requests
- **Default:** 1000 requests

port *port-number*

Configure the port number on which to contact the RADIUS server.

- **Range:** 1 through 65,335
- **Default:** 1812 (as specified in RFC 2865)

preauthentication-port *number*

Configure the RADIUS server preauthentication-port number.

- **Range:** 1 through 65535

preauthentication-secret *secret*

Configure the shared secret with the RADIUS server; it can include spaces if the character string is enclosed in quotation marks. The secret used by the local device must match that used by the RADIUS server.

retry *value*

Configure the number of times that the device is allowed to try to contact a RADIUS authentication server.

- **Range:** 1 through 100
- **Default:** 3

routing-instance *routing-instance-name*

Configure the routing instance name for the management routing instance. In the case of configuring the non-default management instance, use the value `mgmt_junos`. that is `mgmt_junos`. Configuring this option along with the `management-instance` statement enables authentication processes (for example, RADIUS and TACACS+) to use the non-default management routing instance for packet traffic.

> **NOTE**: You must also define the mgmt_junos routing instance under the `[edit routing-instances]` hierarchy level.

If you do not configure the mgmt_junos instance under the `[edit routing-instances]` hierarchy level and configure it only under `tacplus-server` or `radius-server`, the commit will fail.

secret *password*

(Required) Configure the password (shared secret) to use with the RADIUS server; it can include spaces if the character string is enclosed in quotation marks. The secret password used by the local device must match that used by the RADIUS server.

source-address *source-address*

Configure a valid IPv4 or IPv6 address configured on one of the device's interfaces.

timeout *seconds*

Configure the amount of time the local device waits to receive a response from a RADIUS server.

- **Range:** 1 through 1000 seconds
- **Default:** 3 seconds

tls

Configure RADIUS over the Transport Layer Security (TLS) protocol.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`routing-instance` introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Configuring RADIUS Server Authentication

Configuring RADIUS Authentication (QFX Series or OCX Series)

Configuring RADIUS System Accounting

management-instance | **10309**

Management Interface in a Non-Default Instance

radius (System)

IN THIS SECTION

- Syntax | **14034**
- Hierarchy Level | **14035**
- Description | **14035**
- Options | **14035**
- Required Privilege Level | **14035**
- Release Information | **14035**

Syntax

```
radius {
    server {
        server-address {
            accounting-port port-number;
            accounting-retry number;
            accounting-timeout seconds;
            dynamic-request-port number;
            max-outstanding-requests value;
            port number;
            preauthentication-port number;
            preauthentication-secret secret;
            retry number;
            routing-instance routing-instance-name;
            secret password;
            source-address source-address;
            timeout seconds;
```

```
        }
}
```

Hierarchy Level

```
[edit system accounting destination]
```

Description

Configure the RADIUS accounting server.

Options

server-address—Address of the RADIUS accounting server.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

`routing-instance` introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Configuring RADIUS System Accounting

radsec

IN THIS SECTION

- Syntax | 14036
- Hierarchy Level | 14037
- Description | 14037
- Default | 14037
- Options | 14037
- Required Privilege Level | 14039
- Release Information | 14039

Syntax

```
radsec {
    destination id-number {
        address ip-address;
        dynamic-requests {
            (logical-system ls-name <routing-instance ri-name>| routing-instance ri-name);
            source-address source-address;
            source-port source-port;
        }
        id-reuse-timeout seconds;
        (logical-system ls-name <routing-instance ri-name>| routing-instance ri-name);
        max-tx-buffers number;
        port port-number;
        source-address ip-address;
        tls-certificate certificate-name;
        tls-force-ciphers [medium | low];
        tls-min-version [v1.1 | v1.2];
        tls-peer-name tls-peer-name;
        tls-timeout seconds;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure RADIUS over TLS, also known as RADSEC, to redirect regular RADIUS traffic to remote RADIUS servers connected over TLS. The TLS connection provides encryption, authentication, and data integrity for the exchange of RADIUS messages.

To configure RADIUS over TLS, you need to configure the RADSEC server as a destination for RADIUS traffic. Traffic that is destined for a RADIUS server can then be redirected to the RADSEC destination. RADSEC destinations are identified by a unique numeric ID. You can configure multiple RADSEC destinations with different parameters pointing to the same RADSEC server.

TLS relies on certificates and private-public key exchange pairs to secure the transmission of data between the RADSEC client and server. The RADSEC destination uses local certificates that are dynamically acquired from the Junos PKI infrastructure.

To enable RADSEC, you must specify the name of the local certificate. If a certificate is not available, or if the certificate was revoked, the RADSEC destination attempts to retrieve it every 300 seconds.

Default

RADSEC is not enabled by default.

Options

destination ***id-number***

Globally unique ID number for the RADSEC destination.

- **Range:** 1 through 65535

address ***ip-address***

Specify the IP address of the RADSEC server.

id-reuse-timeout ***seconds***

Configure the number of seconds after which the RADIUS ID field value can be reused.

- **Default:** 120 seconds
- **Range:** 60 to 3600 seconds

logical-system *ls-name* **routing-instance** *ri-name*

Specify the logical system or the routing instance for transport.

Default: If you do not explicitly configure the logical system or routing instance, the default is used. You can specify the logical system, the routing instance, or both.

- No configuration specifies default:default.
- Configuring only the routing instance specifies default:*ri-name*.
- Configuring only the logical system specifies *ls-name*:default.
- Configuring both the logical system and the routing instance specifies *ls-name*:*ri-name*.

max-tx-buffers *number*

Configure the maximum number of packets buffered on transmission.

> **NOTE**: The buffer allocation should be able to accommodate the `max-outstanding-requests` for mapped RADIUS servers configured at the [`edit access radius-server`] hierarchy level.

- **Default:** 100
- **Range:** 32 to 3200

port *port-number*

(Optional) Configure the port number of the RADSEC server.

- **Default:** 2083
- **Range:** 1 through 65535

source-address *ip-address*

Configure the source IP address, which is the IP address of the RADSEC server. If the source address is not configured for dynamic requests, dynamic requests will be rejected.

tls-certificate *certificate-name*

Specify the name of the local certificate.

tls-force-ciphers [medium | low]

(Optional) Allow lower-grade ciphers than the default.

- Values:
  - low—Add medium and low grade ciphers.
  - medium—Add medium grade ciphers.

- **NOTE**: The "tls-force-ciphers" option is not applicable. Strong cipher suites are always used by default.

tls-min-version [v1.1 | v1.2]

(Optional) Configure the TLS version to limit the lowest supported versions of TLS that are enabled for SSL connections.

- Values:
  - v1.1—Require TLS 1.1 and 1.2.
  - v1.2—Require TLS 1.2.
- **Default:** v1.2

tls-peer-name *name*

Certified name of the RADSEC server.

tls-timeout *seconds*

Specify a limit on TLS negotiation.

- **Default:** 5 seconds
- **Range:** 3 through 90 seconds

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

`dynamic-requests` introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

RADIUS over TLS (RADSEC)

radsec-destination

IN THIS SECTION

- Syntax | 14040
- Hierarchy Level | 14040
- Description | 14040
- Options | 14040
- Required Privilege Level | 14040
- Release Information | 14041

Syntax

```
radsec-destination id-number;
```

Hierarchy Level

```
[edit access radius-server server-address],
[edit access profile profile-name radius-server server-address]
```

Description

Configure a RADIUS over TLS (RADSEC) server as the destination for RADIUS traffic. The RADIUS traffic is redirected from the RADIUS server to the RADSEC destination. You can redirect more than one RADIUS server to the same RADSEC destination.

Options

id-number—The unique ID number for the RADSEC destination.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

RADIUS over TLS (RADSEC)

raise-threshold

IN THIS SECTION

- Syntax | 14041
- Hierarchy Level | 14041
- Description | 14042
- Options | 14042
- Required Privilege Level | 14042
- Release Information | 14042

Syntax

```
raise-threshold value;
```

Hierarchy Level

```
[edit security nat source pool-utilization-alarm]
```

Description

Configure the upper threshold at which an SNMP trap is triggered when pool utilization for a source pool without Port Address Translation (PAT) rises above the threshold. This feature is disabled by default.

Options

raise-threshold ***value*** Threshold at which an SNMP trap is triggered.

- **Range:** 50 through 100

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

raise-trap (Event Policy)

IN THIS SECTION

- Syntax | 14043
- Hierarchy Level | 14043
- Description | 14043
- Required Privilege Level | 14043
- Release Information | 14043

Syntax

```
raise-trap;
```

Hierarchy Level

```
[edit event-options policy policy-name then]
```

Description

Define an event policy that raises an SNMP trap in response to an event. If one or more of the listed events occur, the system log message for the event is converted into a trap. This enables an agent to notify a trap-based *network management system* (NMS) of significant events.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Overview of Using Event Policies to Raise SNMP Traps

Example: Raise an SNMP Trap in Response to an Event

random

IN THIS SECTION

- Syntax | 14044
- Hierarchy Level | 14044

- Description | 14044
- Default | 14044
- Required Privilege Level | 14045
- Release Information | 14045

Syntax

```
(random | least-fill | most-fill);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Configure the preferred path when several equal-cost candidate paths to a destination exist, and prefer the path with the highest available bandwidth (with the largest minimum available bandwidth ratio). The available bandwidth ratio of a link is the available bandwidth on a link divided by the maximum reservable bandwidth on the link.

- `least-fill`—Prefer the path with the most available bandwidth (with the largest available bandwidth ratio).
- `most-fill`—Prefer the path with the least available bandwidth (with the minimum available bandwidth ratio). The minimum available bandwidth ratio of a path is the smallest available bandwidth ratio belonging to any of the links in the path.
- `random`—Choose the path at random.

Default

`random`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring CSPF Tie Breaking

range (Access)

IN THIS SECTION

- Syntax | 14045
- Hierarchy Level | 14046
- Description | 14046
- Options | 14046
- Required Privilege Level | 14046
- Release Information | 14046

Syntax

```
range range-name {
    high upper-limit;
    low lower-limit;
    prefix-length delegated-prefix-length;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet6]
[edit access address-assignment pool pool-name family inet]
```

Description

Configure an IP name range used within an address-assignment pool. For IPv4, you do not create a prefix-length.

Options

- *range-name*—Name of the range.
- `high` *upper-limit*—Upper limit of IPv6 address range.
- `low` *lower-limit*—Lower limit of IPv6 address range.
- `prefix-length` *delegated-prefix-length*—IPv6 delegated prefix length.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.4 of Junos OS.

RELATED DOCUMENTATION

Firewall User Authentication Overview

range-address

IN THIS SECTION

- Syntax | 14047
- Hierarchy Level | 14047
- Description | 14047
- Options | 14047
- Required Privilege Level | 14047
- Release Information | 14048

Syntax

```
range-address lower-limit to upper-limit;
```

Hierarchy Level

```
[edit security address-book address address-name]
```

Description

Configure the address range for an address book.

Options

- `lower-limit`—Lower limit of an address range.
- `upper-limit`—Upper limit of an address range.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Address Books

Understanding Address Sets

range (Address-Assignment Pools)

IN THIS SECTION

- Syntax | 14048
- Hierarchy Level | 14048
- Description | 14049
- Options | 14049
- Required Privilege Level | 14049
- Release Information | 14049

Syntax

```
range range-name {
    high upper-limit;
    low lower-limit;
    prefix-length prefix-length;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family (inet | inet6)]
```

Description

Configure a named range of IPv4 addresses or IPv6 prefixes, used within an address-assignment pool.

Options

`high` *`upper-limit`*—Upper limit of an address range or IPv6 prefix range.

`low` *`lower-limit`*—Lower limit of an address range or IPv6 prefix range.

`prefix-length` *`prefix-length`*—Assigned length of the IPv6 prefix.

`range-name`—Name assigned to the range of IPv4 addresses or IPv6 prefixes.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

IPv6 support introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Address-Assignment Pools for Subscriber Management

range (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 14050
- Hierarchy Level | 14050
- Description | 14050

- Options | **14050**
- Required Privilege Level | **14051**
- Release Information | **14051**

Syntax

```
range {
    port-low to port-high;
    (random-allocation | round-robin);
}
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Description

To configure a range of ports to assign to a pool, specify the low and high values for the port. If you do not configure automatic port assignment, you must configure a range of ports. This statement applies to source NAT with port translation, but not to deterministic NAT.

If you specify a range, ports are selected a round-robin fashion. If you specify a range of ports to assign, the automatic statement is ignored.

Options

port-low	Lowest port number.
port-high	Highest port number.
random-allocation	Randomly assigns a port from the range 1024 through 65535 for each port translation.
round-robin	First assigns port 1024, and uses the next higher port for each successive port assignment. Round robin allocation is the default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

ranges (Dynamic Stacked VLAN)

IN THIS SECTION

- Syntax | 14051
- Hierarchy Level | 14051
- Description | 14051
- Options | 14052
- Required Privilege Level | 14052
- Release Information | 14052

Syntax

```
ranges (any | low-tag-high-tag),(any | low-tag-high-tag);
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure stacked-vlan-ranges dynamic-profile profile-name]
```

Description

Configure VLAN ranges for dynamic, auto-sensed stacked VLANs.

Options

any—The entire VLAN range.

low-tag—The lower limit of the VLAN range.

high-tag—The upper limit of the VLAN range.

- **Range:** 1 through 4094

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

ranges (Dynamic VLAN)

IN THIS SECTION

- Syntax | **14053**
- Hierarchy Level | **14053**
- Description | **14053**
- Options | **14053**
- Required Privilege Level | **14053**
- Release Information | **14053**

Syntax

```
ranges (any | low-tag)-(any | high-tag);
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges dynamic-profile profile-name]
```

Description

Configure VLAN ranges for dynamic, auto-sensed VLANs.

Options

`any`—The entire VLAN range.

low-tag—The lower limit of the VLAN range.

high-tag—The upper limit of the VLAN range.

- **Range:** 1 through 4094

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

> *Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs*

rapid-commit (DHCPv6 Client)

IN THIS SECTION

- Syntax | 14054
- Hierarchy Level | 14054
- Description | 14054
- Required Privilege Level | 14054
- Release Information | 14055

Syntax

```
rapid-commit;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcpv6-client]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet6 dhcpv6-client]
[edit tenants tenant-name interfaces interface-name unit logical-unit-number family inet6dhcpv6-
client]
```

Description

Used to signal the use of the two-message exchange for address assignment.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

The `logical-systems` and `tenants` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

DHCPv6 Client Overview

Understanding DHCPv6 Client and Server Identification

rapid-commit (DHCPv6 Local Server)

IN THIS SECTION

- Syntax | 14055
- Hierarchy Level | 14055
- Description | 14056
- Default | 14056
- Required Privilege Level | 14056
- Release Information | 14056

Syntax

```
rapid-commit;
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6 overrides],
[edit system services dhcp-local-server dhcpv6 group group-name overrides],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name
overrides],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 ...],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 ...],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 ...]
```

Description

Configure DHCPv6 local server to recognize the Rapid Commit option (DHCPv6 option 14) in DHCPv6 solicit messages sent from the DHCPv6 client. When rapid commit is enabled for both DHCPv6 local server and the DHCPv6 client, a two-message handshake is used instead of the standard four-message handshake. You can enable rapid commit support on DHCPv6 local server globally, for a named group, or for a specific interface.

Default

Rapid commit support is not enabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring DHCPv6 Rapid Commit (MX Series, EX Series)

Overriding the Default DHCP Local Server Configuration Settings

rate

IN THIS SECTION

- Syntax | 14057
- Hierarchy Level | 14057
- Description | 14057
- Options | 14057
- Required Privilege Level | 14058
- Release Information | 14058

Syntax

```
rate {
    (fixed-stuff-bytes | no-fixed-stuff-bytes);
    otu4; oc192;
    (pass-through | no-pass-through);
}
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify the line rate or speed of the OTN signals.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Statement and `otu4` option introduced in Junos OS Release 13.2 for PTX Series routers.

Option `oc192` introduced in Junos OS Release 13.3 for MX Series routers.

RELATED DOCUMENTATION

10-Gigabit Ethernet OTN Options Configuration Overview

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

rate-cap

IN THIS SECTION

- Syntax | 14058
- Hierarchy Level | 14059
- Description | 14059
- Options | 14059
- Required Privilege Level | 14059
- Release Information | 14059

Syntax

```
rate-cap <rate-cap-value>;
```

Hierarchy Level

```
[edit security log]
[edit logical-systems name security log]
[edit tenants tenant-name security log]
```

Description

Limits the rate at which data plane logs will be generated per second.

Options

`rate-cap` *rate-cap-value*—Works with event mode only. Limits the rate at which data plane logs will be generated per second

- **Range:** 0 through 5000 logs per second
- **Default:** 5000 logs per second

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

The [edit `logical-systems` *name* security log] and [edit `tenants` *tenant-name* security log] hierarchy levels introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

syslog (System) | 28836

rate (Forwarding Options)

IN THIS SECTION

- Syntax | 14060
- Hierarchy Level | 14060
- Description | 14060
- Options | 14061
- Required Privilege Level | 14061
- Release Information | 14061

Syntax

```
rate number;
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input],
[edit forwarding-options port-mirroring input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance instance-name input]
```

Description

Set the ratio of the number of packets to be sampled. For example, if you specify a rate of 10, every tenth packet (1 packet out of 10) is sampled.

Native analyzer sessions (that is, the `[edit forwarding-options analyzer analyzer-name input]` hierarchy level for MX Series routers) can be configured without specifying input parameters, which means that the instance uses default input values: rate = 1 and maximum-packet-length = 0.

NOTE: The recommended sampling rate for the MX150 is 1000 or greater. If you configure less than 1000, a warning is issued.

Options

number—Denominator of the ratio.

- **Range:** 1 through 16000000 (16M)

  For QFX Series switches, the maximum sampling rate for inline Junos Traffic Vision (J-Flow) is 65535.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit forwarding-options analyzer analyzer-name input]` hierarchy level for MX Series routers introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring Port Mirroring

Configuring Traffic Sampling

rate (Interface Services)

IN THIS SECTION

- Syntax | 14062
- Hierarchy Level | 14062
- Description | 14062
- Options | 14062
- Required Privilege Level | 14062
- Release Information | 14062

Syntax

```
rate new-sessions-per-second;
```

Hierarchy Level

```
[edit interfaces interface-name services-options session-limit]
```

Description

Specify the maximum number of new sessions allowed per second on services cards.

Options

rate *new-sessions-per-second*

Specify the maximum number of new sessions allowed per second.

- **Range:** 0, which indicates no limit, or greater.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

rate-limit

IN THIS SECTION

- Syntax | 14063
- Hierarchy Level | 14063
- Description | 14063
- Default | 14064
- Options | 14064
- Required Privilege Level | 14065
- Release Information | 14065

Syntax

```
rate-limit limit;
```

Hierarchy Level

```
[edit system services finger],
[edit system services ftp],
[edit system services netconf ssh],
[edit system services ssh],
[edit system services telnet],
[edit system services tftp-server],
[edit system services xnm-clear-text],
[edit system services xnm-ssl]
```

Description

Configure the maximum number of connections attempts per minute, per protocol (either IPv6 or IPv4) on an access service. For example, a rate limit of 10 allows 10 IPv6 ssh session connection attempts per minute and 10 IPv4 ssh session connection attempts per minute.

Default

150 connections

Options

`rate-limit` *`limit`*—(Optional) Maximum number of connection attempts allowed per minute, per IP protocol (either IPv4 or IPv6).

- Junos OS
  - **Range:** 1 through 250
  - **Default:** 150
- Junos OS Evolved
  - **Range:** 1 through 250
  - **Default:** 150

NOTE: For certain Junos OS Evolved releases, the `rate-limit` option was specified as the number of connection attempts allowed per second, and had various default values.

Table 107: Releases with Connection Attempts Allowed per Second

Release	Range	Default
Junos OS Evolved Release 20.4R1	1 through 5	3
Junos OS Evolved Release 20.4R2	1 through 5	(unlimited - no rate-limit enforced)
Junos OS Evolved Release 21.1R1	1 through 5	3
Junos OS Evolved Release 21.1R2	1 through 50	25
Junos OS Evolved Release 21.2R1	1 through 50	25

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring clear-text or SSL Service for Junos XML Protocol Client Applications

rate-limit (Aggregated rate limit)

IN THIS SECTION

- Syntax | 14065
- Hierarchy Level | 14066
- Description | 14066
- Options | 14067
- Required Privilege Level | 14067
- Release Information | 14067

Syntax

```
rate-limit name {
    message {
        v0 msg-list-name;
        v1 msg-list-name;
        v2 msg-list-name;
```

```
    }
}
```

Hierarchy Level

```
[edit security gtp]
```

Description

Use the rate-limit function to configure the message that needs to be rate limited.

Use the message-list function to configure the list of messages for aggregated rate limiting. This helps to specify aggregated rate limit messages of various GTP versions. It is also used to customize aggregated rate limit for different GTP message-list profiles.

The configured message-list replaces the default messages. The configured *msg-list-name* should be a defined message-list.

If the message-lists are not configured, aggregated rate limit uses the default messages.

Table 108: Default aggregated rate limit messages for various GTP versions.

GTP version	Default aggregated rate limit messages
GTPv0	Echo Request
	Create PDP Request
	Create AA PDP Request
	Update PDP Request
	Delete PDP Request
	Delete AA PDP Request
GTPv1	Echo Request
	Create PDP Request

Table 108: Default aggregated rate limit messages for various GTP versions. ***(Continued)***

GTP version	Default aggregated rate limit messages
	Update PDP Request
	Delete PDP Request
GTPv2	Echo Request
	Create Session Request
	Create Bearer Request
	Modify Bearer Request
	Delete Session Request
	Delete Bearer Request

Options

name Name of the rate-limit. Currently, only default option is available.

message Specifies GTP limited message type.

v0 Specifies GTPv0 limited message type.

v1 Specifies GTPv1 limited message type.

v2 Specifies GTPv2 limited message type.

msg-list-name Specifies the GTP message list name. The configured msg-list-name should be a defined message-list.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

show security gtp rate-limit default | 24460

rate-limit (Security GTP)

IN THIS SECTION

- Syntax | 14068
- Hierarchy Level | 14068
- Description | 14068
- Options | 14069
- Required Privilege Level | 14069
- Release Information | 14069

Syntax

```
rate-limit {
    alarm-threshold alarm-threshold;
    drop-threshold drop-threshold;
}
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

Use the rate-limit option to limit the GTP messages per second. The default value of rate limit is 0, which means there is no limit. A drop threshold option and an alarm threshold option are added to the existing rate-limit configuration to reduce the duplicate drop logs for destination GSN. In one second, if number of certain messages received reaches the alarm threshold, one alarm log is generated. If number

of certain messages received reaches the drop threshold, one packet drop log is generated while all other messages received within this second will be dropped.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name]` hierarchy level is replaced by `[edit security gtp profile profile-name]`.

Options

alarm-threshold

Sets alarm threshold for rate-limit.

- **Range:** 1 through 80000

drop-threshold

Sets drop threshold for rate-limit.

- **Range:** 1 through 80000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. Support. Support for `alarm-threshold` option and `drop-threshold` option are added in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Understanding Path Rate Limiting for GTP Control Messages

Example: Limiting the Message Rate and Path Rate for GTP Control Messages

rate-limit (Security Log)

IN THIS SECTION

Syntax | 14070

- Hierarchy Level | 14070
- Description | 14070
- Required Privilege Level | 14071
- Release Information | 14071

Syntax

```
rate-limit rate-limit;
```

Hierarchy Level

```
[edit security log stream stream-name]
[edit logical-systems name security log stream stream-name]
[edit tenants tenant-name security log stream stream-name]
```

Description

The Integrated Authentication and Enforcement feature sends threat and attack logs generated by the SRX Series Firewall security modules to the ClearPass Policy Manager (CPPM) to use in its security policy assessment.

The logs are sent in stream mode. To avoid overburdening the SRX Series Firewall and the log server, you can control the rate at which these logs are sent. By setting a rate-limit value, you can constrain the number of logs that are sent in 1 second. After the limit is reached, no more logs are sent.

Range: 1 through 65,535.

NOTE: For devices with multicore systems that use SPUs, each SPU is programmed with the configured-rate, which results in an aggregate-rate proportional to the number of SPUs.

```
SPU configured-rate = aggregate-rate/number-of-SPUs
```

Rate limiting on SRX5400, SRX5600, and SRX5800 devices is generally not as accurate as it is on SRX100, SRX210, SRX220, SRX240, SRX650, SRX1400, SRX1500, SRX3400, SRX3600, SRX4100, and SRX4200 devices, because the generation of logs is not entirely balanced between SPUs.

Required Privilege Level

1. security—To view this statement in the configuration.
2. security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

The [edit `logical-systems` *name* security log stream] hierarchy level introduced in Junos OS Release 18.2R1.

The [edit `tenants` *tenant-name* security log stream] hierarchy level introduced in Junos OS Release 18.3R1.

rate-limiters

IN THIS SECTION

- Syntax | 14072
- Hierarchy Level | 14072
- Description | 14072
- Options | 14072
- Required Privilege Level | 14073
- Release Information | 14073

Syntax

```
rate-limiters {
    rate-limiter-name {
        bandwidth-limit value-in-kbps;
        burst-size-limit value-in-bytes;
    }
}
```

Hierarchy Level

```
[edit class-of-service application-traffic-control]
[edit logical-systems logical-system-name class-of-service application-traffic-control]
[edit tenants tenant-name class-of-service application-traffic-control]
```

Description

Share the available bandwidth and burst size of a device's PICs by defining rate limiter profiles and applying them in AppQoS rules.

Options

- *rate-limiter-name*—Name of the rate limiter. It is applied in AppQoS rules to share device resources based on quality-of-service requirements.

  The combination of rate limiting parameters, namely bandwidth- limit and burst-size-limit rate limit, make up the rate limiter profile. A maximum of 16 profiles are allowed per device. The same profile can be used by multiple rate limiters. For example, a profile with a bandwidth-limit of 200 Kbps and a burst-limit of 130,000 bytes, could be used in several rate limiters.

  A maximum of 1000 rate limiters can be created. Rate limiters are defined for the device, and are assigned in rules in a rule set. A single rate limiter can be used multiple times within the same rule set. However, the rate limiter cannot be used in another rule set.

- `bandwidth-limit` *value-in-Kbps*—Maximum number of kilobits to be transmitted per second for this rate limiter. Up to 2 GB of bandwidth can be provisioned among multiple rate limiters to share the resource proportionally.

- `burst-size-limit` *`value-in-bytes`*—Maximum number of bytes to be transferred in a single burst or time-slice. This limit ensures that a high-priority transmission does not keep a lower priority transmission from transmitting.

> **NOTE**: The number of `bandwidth-limit` and `burst-size-limit` combinations cannot exceed 16.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the following hierarchy levels introduced in Junos OS Release 19.3R1: `[edit logical-systems` *`logical-system-name`* `class-of-service application-traffic-control]`, and `[edit tenants` *`tenant-name`* `class-of-service application-traffic-control]`.

RELATED DOCUMENTATION

Example: Configuring AppQoS

rate (Routing Instances)

IN THIS SECTION

- Syntax | 14074
- Hierarchy Level | 14074
- Description | 14074
- Options | 14074
- Required Privilege Level | 14074
- Release Information | 14074

Syntax

```
rate threshold-rate;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
mdt threshold group group-address source source-address],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel family inet | inet6mdt threshold group group-address source source-address],
[edit routing-instances routing-instance-name protocols pim mdt threshold group group-address
source source-address],
[edit routing-instances routing-instance-name  provider-tunnel family inet | inet6mdt threshold
group group-address source source-address]
```

Description

Apply a rate threshold to a multicast source to automatically create a data MDT.

Options

threshold-rate—Rate in kilobits per second (Kbps) to apply to source.

- **Range:** 10 Kbps through 1 Gbps (1,000,000 Kbps)
- **Default:** 10 Kbps

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4. `mdt` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default MDT in Rosen 7, and data MDT for Rosen 6 and Rosen 7. The `provider-tunnel mdt` hierarchy is now hidden for backward compatibility with existing scripts.

RELATED DOCUMENTATION

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

rcp-count

IN THIS SECTION

- Syntax | 14075
- Hierarchy Level | 14075
- Description | 14075
- Options | 14076
- Required Privilege Level | 14076
- Release Information | 14076

Syntax

```
rcp-count number;
```

Hierarchy Level

```
[edit chassis nssu]
```

Description

(QFX5100 Virtual Chassis or Virtual Chassis Fabric [VCF] only) Optionally override the default algorithm that defines the number of parallel `rcp` sessions used for copying the new software image to member switches in a Virtual Chassis or VCF for nonstop software upgrade (NSSU).

You can upgrade the software running on all members of a Virtual Chassis or VCF using NSSU. At the beginning of the upgrade process, NSSU uses `rcp` to copy the new software from the primary switch to each of the member switches. Starting in Junos OS Release 14.1X53-D40, to minimize the time to copy

the files to all members, NSSU uses parallel `rcp` sessions to copy the software to multiple members at the same time, rather than copying sequentially to each member in turn.

By default, the number of parallel `rcp` sessions that NSSU launches at a time follows an algorithm that optimizes the transfer time based on the total number of members in the Virtual Chassis or VCF, so that the new software is transferred to about half the members in parallel up to a maximum of 8 members at one time. After the first set of parallel copy operations are complete, NSSU launches the next set of parallel copy operations, and so on as needed until all members are updated.

Although the default value is usually optimal, you can configure the `rcp-count` statement to set a specific value for the number of parallel copy sessions instead. For example, you might want to use a lower number of parallel sessions to avoid potential impact on Virtual Chassis or VCF forwarding activity during NSSU.

After copying the new software to all members using as many parallel `rcp` sessions as needed, the NSSU process continues by rebooting each member with the new software in sequence, starting with the member in the backup Routing Engine role.

> **NOTE**: If copying the new software to any member fails, NSSU terminates the upgrade process for the entire Virtual Chassis or VCF, logs the error condition, and initiates an error recovery measure to remove the new software from the members to which it was already transferred.

Options

number Number of parallel `rcp` sessions for NSSU to launch at a time, which overrides the default algorithm.
Allowable values are 0 through 8. Specifying a value of 0 suppresses parallel copy sessions; the new software is copied to each member sequentially.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D40.

RELATED DOCUMENTATION

Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis

Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric

rdi

IN THIS SECTION

- Syntax | 14077
- Hierarchy Level | 14077
- Description | 14077
- Required Privilege Level | 14078
- Release Information | 14078

Syntax

```
rdi;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile tlv-action event]
```

Description

Define a new event **rdi**. The remote defect indication (rdi) event is triggered whenever CCM packets are received from a remote location with the rdi bit set.

This event is cleared and action is reverted when none of the remote MEPs send the CCM packets with the RDI bit.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

rdma-opcode

IN THIS SECTION

- Syntax | 14078
- Hierarchy Level | 14078
- Description | 14079
- Required Privilege Level | 14079
- Release Information | 14079

Syntax

```
rdma-opcode
```

Hierarchy Level

```
[edit firewall family ethernet-switching filter filter name term term name from]
[edit firewall family inet filter filter name term term name from]
[edit firewall family inet6 filter filter name term term name from]
[edit family ethernet-switching filter filter name term term name from ip-version ipv6]
[edit family ethernet-switching filter filter name term term name from ip-version ipv4]
```

Description

Match on InfiniBand Base Transport header opcode. Input a value or range of values. Apply it as a match condition only in an input firewall filter.

Required Privilege Level

firewall - To view this statement in the configuration

firewall-control - To add this statement to the configuration

Release Information

Statement introduced in Junos 24.2R1

Statement introduced in Junos OS Evolved 22.2X100-D20

re-assembler

IN THIS SECTION

- Syntax | 14079
- Hierarchy Level | 14080
- Description | 14080
- Options | 14080
- Required Privilege Level | 14081
- Release Information | 14081

Syntax

```
re-assembler {
    action-on-reassembly-failure (drop | drop-session | ignore);
    (force-tcp-window-checks | no-force-tcp-window-checks);
    (ignore-memory-overflow | no-ignore-memory-overflow);
    (ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-overflow);
```

```
    ignore-reassembly-overflow;
    max-flow-mem value;
    max-packet-mem-ratio percentage-value;
     max-synacks-queued value;
    (tcp-error-logging | no-tcp-error-logging);
}
```

Hierarchy Level

```
[edit security idp sensor-configuration]
```

Description

Configure TCP reassembler for IDP sensor settings.

Options

max-flow-mem

Define the maximum TCP flow memory that the IDP sensor can handle.

- **Syntax:** *value*—Maximum TCP flow memory in kilobytes.
- **Range:** 64 through 4,294,967,295 kilobytes
- **Default:** 1024 kilobytes

max-synacks-queued

Define the maximum limit for queuing Syn/Ack packets with different SEQ numbers.

- **Syntax:** *value*—Maximum synchronization acknowledgements queued with different SEQ numbers.
- **Range:** 0 through 5

max-packet-mem-ratio

By default, values for IDP reassembler packet memory are established as percentages of all memory. In most cases, these default values are adequate.

If a deployment exhibits an excessive number of dropped TCP packets or retransmissions resulting in high IDP reassembly memory usage, use the `max-packet-mem-ratio` option to reset the percentage of available IDP memory for IDP reassembly packet memory. Acceptable values are between 5 percent and 40 percent.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. Packet memory ratios added in Junos OS Release 12.1X44-D20.

re-enroll-trigger-time-percentage

IN THIS SECTION

- Syntax | 14081
- Hierarchy Level | 14081
- Description | 14082
- Options | 14082
- Required Privilege Level | 14082
- Release Information | 14082

Syntax

```
re-enroll-trigger-time-percentage percentage;
```

Hierarchy Level

```
[edit security pki auto-re-enrollment certificate-id]
```

Description

Percentage of the router certificate "validity-period" on page 30650 statement value, in days, when auto-reenrollment should start before expiration.

Options

percentage—Percentage for the reenroll trigger time.

- **Range:** 1 through 99

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Auto-Reenrollment of a Router Certificate

auto-re-enrollment | 2727

re-generate-keypair

IN THIS SECTION

- Syntax | 14083
- Hierarchy Level | 14083
- Description | 14083
- Required Privilege Level | 14083
- Release Information | 14083

Syntax

```
<re-generate-keypair>;
```

Hierarchy Level

```
[edit security pki auto-re-enrollment certificate-id]
```

Description

(Optional) Automatically generate a new key pair when auto-reenrolling a router certificate. If this statement is not configured, the current key pair is used.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Auto-Reenrollment of a Router Certificate

auto-re-enrollment | 2727

reachability-fault

IN THIS SECTION

Syntax (PTX Series Devices) | 14084

- Hierarchy Level | 14084
- Description | 14085
- Options | 14085
- Required Privilege Level | 14085
- Release Information | 14085

Syntax (PTX Series Devices)

```
reachability-fault {
      actions {
                  fpc-restart-disable;
                  recovery-failure
                  {
                     pfe-restart-disable;
                  }
              }
      degraded {
            apply-groups
            apply-groups-except
            error-threshold {
                  degradation percentage;
            }
      }
}
```

Hierarchy Level

```
[edit chassis fabric event],
[edit chassis member name fabric event]
```

Description

Fabric reachability fault detection. Configure options that apply to detect fabric reachability fault conditions and trigger automatic fabric connectivity restoration.

Options

- `actions`— Reachability fault recovery actions. Depending on the reachability fault conditions, allow users to configure one or more recovery actions for the router to attempt automatic fabric connectivity restoration.
  - `fpc-restart-disable`— Turn off FPC restart recovery action.
  - `recovery failure (pfe-restart-disable)`— Configure PFE-offline to override the default action of disabling interfaces, if fabric connectivity cannot be restored.
- `degraded`— Degraded fabric configuration. Allow users to configure the error threshold percentage value. The percentage value can range from 1 to 99, and it represents the percentage of fabric degradation at which the recovery action is taken.
  - `apply-groups` *`group-name,`* — Inherit configuration data from these groups.
  - `+ apply-groups-except`— Do not inherit configuration data from these groups.
  - `error-threshold`— Fabric error threshold when PFE is degraded.
  - `degradation` *`percentage`*— Degradation percentage values from 1 to 99.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 22.4R1 (PTX10004, PTX10008, PTX10016).

reachable-time

IN THIS SECTION

- Syntax | 14086
- Hierarchy Level | 14086
- Description | 14086
- Options | 14086
- Required Privilege Level | 14087
- Release Information | 14087

Syntax

```
reachable-time milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-
name],
[edit protocols router-advertisement interface interface-name]
```

Description

Set the length of time that a node considers a neighbor reachable until another reachability confirmation is received from that neighbor.

After receiving a reachability confirmation from a neighbor, a node considers that neighbor reachable for a certain amount of time without receiving another confirmation. This mechanism is used for neighbor unreachability detection, a mechanism for finding link failures to a target node.

Options

milliseconds—Reachability time limit.

- **Range:** 0 through 3,600,000 milliseconds
- **Default:** 0 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery
>
> *Secure IPv6 Neighbor Discovery*

reachable-time (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 14087
- Hierarchy Level | 14088
- Description | 14088
- Options | 14088
- Required Privilege Level | 14088
- Release Information | 14088

Syntax

```
reachable-time milliseconds;
```

Hierarchy Level

```
[edit protocols router-advertisement interface interface-name]
```

Description

Set the length of time that a node considers a neighbor reachable until another reachability confirmation is received from that neighbor.

Options

milliseconds—Reachability time limit.

- **Range:** 0 through 3,600,000 milliseconds
- **Default:** 0 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *IPv6 WAN Link Addressing with NDRA*
>
> *Dynamic Router Advertisement Configuration Overview*

read-view

IN THIS SECTION

- Syntax | 14089
- Hierarchy Level | 14089
- Description | 14089
- Options | 14089
- Required Privilege Level | 14089
- Release Information | 14090

Syntax

```
read-view view-name;
```

Hierarchy Level

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix context-
prefix) security-model (any | usm | v1 | v2c) security-level (authentication | none | privacy)]
```

Description

Associate the read-only view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).

Options

view-name—The name of the view to which the SNMP user group has access.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

readvertise

IN THIS SECTION

- Syntax | 14090
- Hierarchy Level | 14090
- Description | 14091
- Default | 14091
- Options | 14091
- Required Privilege Level | 14091
- Release Information | 14091

Syntax

```
(readvertise | no-readvertise);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name static (defaults | route)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options static (defaults | route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
(defaults | route)],

[edit logical-systems logical-system-name routing-options static (defaults | route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
(defaults | route)],
[edit routing-instances routing-instance-name routing-options static (defaults | route)],
```

```
[edit routing-options rib routing-table-name static (defaults | route)],
[edit routing-options static (defaults | route)]
```

Description

Configure whether static routes are eligible to be readvertised by routing protocols:

Default

Static routes are eligible to be readvertised (that is, exported from the routing table into dynamic routing protocols) if a policy to do so is configured. To mark an IPv4 static route as being ineligible for readvertisement, include the `no-readvertise` statement.

Options

readvertise Readvertise static routes. Include the `readvertise` statement when configuring an individual route in the **route** portion of the `static` statement to override a **no-readvertise** option specified in the **defaults** portion of the statement.

no-readvertise Mark a static route as being ineligible for readvertisement. Include the **no-readvertise** option when configuring the route.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Preventing a Static Route from Being Readvertised

Understanding Static Route Control in Routing and Forwarding Tables

static (Routing Options) | **28570**

real-service (Traffic Load Balancer)

IN THIS SECTION

- Syntax | 14092
- Hierarchy Level | 14092
- Description | 14092
- Options | 14092
- Required Privilege Level | 14093
- Release Information | 14093

Syntax

```
real-service real-service-name {
    address server-ip-address;
    admin-down;
}
```

Hierarchy Level

```
[edit services traffic-load-balance instance instance-name]
```

Description

Configure a traffic load balancer server.

Options

admin-down Set a server's status to Down.

real-service-name Identifier for a server to which sessions can be distributed using the server distribution table in conjunction with the session distribution API.

server-ip-address IP address for the server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Traffic Load Balancer Overview

Configuring TLB

realm

IN THIS SECTION

- Syntax | 14094
- Hierarchy Level | 14094
- Description | 14094
- Options | 14094
- Required Privilege Level | 14094
- Release Information | 14094

Syntax

```
realm (ipv4-unicast | ipv4-multicast | ipv6-unicast) {
    area area-id {
        interface interface-name;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf3],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3],
[edit protocols ospf3],
[edit routing-instances routing-instance-name protocols ospf3]
```

Description

Configure OSPFv3 to advertise address families other than unicast IPv6. Junos OS maps each address family you configure to a separate realm with its own set of neighbors and link-state database.

Options

ipv4-unicast—Configure a realm for IPv4 unicast routes.

ipv4-multicast—Configure a realm for IPv4 multicast routes.

ipv6-multicast—Configure a realm for IPv6 multicast routes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Example: Configuring Multiple Address Families for OSPFv3

realm-delimiter (Domain Map)

IN THIS SECTION

- Syntax | 14095
- Hierarchy Level | 14095
- Description | 14095
- Default | 14095
- Options | 14096
- Required Privilege Level | 14096
- Release Information | 14096

Syntax

```
realm-delimiter [delimiter-character];
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the characters that the router uses to separate usernames from realm names.

Default

none

Options

delimiter-character—One or more characters used as delimiters. You can specify a maximum of eight delimiters. You cannot use the semicolon (;) as a delimiter. Do not include spaces between characters.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Specifying Domain and Realm Name Delimiters

Configuring Domain and Realm Name Usage for Domain Maps

realm (Diameter Origin)

IN THIS SECTION

- Syntax | 14097
- Hierarchy Level | 14097
- Description | 14097
- Options | 14097
- Required Privilege Level | 14097
- Release Information | 14097

Syntax

```
realm realm-name;
```

Hierarchy Level

```
[edit access diameter origin]
```

Description

Specify the realm of the host that originates the Diameter message.

Options

realm-name Name of the message origin realm. Supplied as the value of Origin-Realm AVP for all messages sent by the Diameter instance.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

realm-parse-direction (Domain Map)

IN THIS SECTION

- Syntax | 14098
- Hierarchy Level | 14098
- Description | 14098
- Default | 14098
- Options | 14098
- Required Privilege Level | 14099
- Release Information | 14099

Syntax

```
realm-parse-direction (left-to-right | right-to-left);
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the direction in which the router searches for the realm name in a username.

Default

`left-to-right`

Options

`left-to-right`—The router searches starting at the left-most character. When the router reaches a realm delimiter, it uses anything to the left of the delimiter as the realm name.

`right-to-left`—The router searches starting at the right-most character. When the router reaches a realm delimiter, it uses anything to the left of the realm as the domain name.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Specifying the Parsing Direction for Domain and Realm Names

Configuring Domain and Realm Name Usage for Domain Maps

realtime-ukern-thread (BFD)

IN THIS SECTION

- Syntax | 14099
- Hierarchy Level | 14100
- Description | 14100
- Required Privilege Level | 14100
- Release Information | 14100

Syntax

```
realtime-ukern-thread;
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable the real-time Bidirectional Forwarding Detection (BFD) protocol. After real-time BFD is enabled, the priority of the flowd ukernel thread is changed to the highest level and, therefore, the flowd ukernel thread gets more CPU cycles for processing the BFD packets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Enabling Dedicated and Real-Time BFD

show chassis realtime-ukern-thread | **18692**

bfd-liveness-detection (LAG) | **2926**

authentication (LAG) | **2622**

detection-time (LAG) | **5549**

transmit-interval (LAG) | **30090**

reassemble-packets

IN THIS SECTION

Syntax | **14101**

- Hierarchy Level | **14101**
- Description | **14101**
- Default | **14101**
- Required Privilege Level | **14101**
- Release Information | **14102**

Syntax

```
reassemble-packets;
```

Hierarchy Level

```
[edit interfaces gr-fpc/pic/port unit logical-unit-number],
[edit logical-systems logical-system-name interfaces gr-fpc/pic/port unit logical-unit-number]
```

Description

Enable reassembly of fragmented generic routing encapsulation (GRE) encapsulated packets on GRE tunnel interfaces at the endpoint of the GRE tunnel.

GRE-encapsulated packets are fragmented if the "allow-fragmentation" on page 2179 statement is configured for the GRE tunnel and the size of the GRE-encapsulated packet exceeds the maximum transmission unit (MTU) value of a link that the packet passes through.

Default

If you do not include the `reassemble-packets` statement, the GRE tunnel interface does not reassemble fragmented GRE-encapsulated packets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Enabling Fragmentation and Reassembly on Packets After GRE-Encapsulation

reassembly-timeout

IN THIS SECTION

- Syntax | 14102
- Hierarchy Level | 14102
- Description | 14103
- Options | 14103
- Required Privilege Level | 14103
- Release Information | 14103

Syntax

```
reassembly-timeout seconds;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
[edit security flow]
```

Description

The maximum acceptable time, in seconds, from the receipt of the first and latest fragments in a packet. When the number is exceeded, the packet is dropped.

Options

seconds—Maximum seconds allowed.

- **Range:** 1 to 60 seconds.
- **Default:** 4 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

Statement added in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480, and MX960 routers.

RELATED DOCUMENTATION

Configuring Fragmentation Control for MS-DPC and MS-PIC Service Interfaces

reauthenticate (DHCP Local Server)

IN THIS SECTION

- Syntax | 14104
- Hierarchy Level | 14104
- Description | 14104

- Options | 14105
- Required Privilege Level | 14106
- Release Information | 14106

Syntax

```
reauthenticate (<lease-renewal> <remote-id-mismatch > <actual-data-rate-change>);
```

Hierarchy Level

```
[edit logical-systems name routing-instances name system services dhcp-local-server],
[edit logical-systems name routing-instances name system services dhcp-local-server dhcpv6],
[edit logical-systems name routing-instances name system services dhcp-local-server dual-stack-
group name],
[edit logical-systems name system services dhcp-local-server],
[edit logical-systems name system services dhcp-local-server dhcpv6],
[edit logical-systems name system services dhcp-local-server dual-stack-group name],
[edit routing-instances name system services dhcp-local-server ],
[edit routing-instances name system services dhcp-local-server dhcpv6],
[edit routing-instances name system services dhcp-local-server
dual-stack-group name],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dual-stack-group name]
```

Description

Enable DHCP client reauthentication, that is, trigger jdhcpd to request reauthentication from authd, which in turn reissues the RADIUS Access-Request for subscriber authentication. The purpose of the reauthentication is to change characteristics of the subscriber session, such as activating subscriber services or changing attributes. You can use reauthentication as an alternative to a RADIUS CoA request.

Starting in Junos OS Release 18.1R1, reauthentication can be triggered by discover and solicit messages in addition to the previously supported renew and rebind messages. The release also introduces reauthentication support for dual-stack, single-session subscribers.

You can specify that reauthentication occurs in response to all DHCP renew, rebind, discover, or solicit messages or only in response to discover and solicit messages that include a new (different) Agent Remote ID for the DHCP client.

You can use the Juniper Networks VSA, Reauthentication-On-Renew (26-206) as an alternative to the CLI configuration to enable reauthentication. The `reauthenticate` statement overrides the VSA when the VSA is present with a value of disable.

NOTE: Reauthentication for dual-stack, single-session subscribers requires that the `on-demand-address-allocation` statement is configured for the dual-stack group. This is true whether you enable reauthentication with the `reauthenticate` statement or the Reauthenticate-On-Renew VSA (26-206).

NOTE: You cannot configure both the `reauthenticate` statement and the remote-id-mismatch (DHCP Local Server and DHCP Relay Agent) statement at the global level, `[edit system services dhcp-local-server]`. However, DHCP precedence rules do permit you to configure both statements when they are at different levels. For example, you can configure `reauthenticate` at the global level and remote-id-mismatch (DHCP Local Server and DHCP Relay Agent) for DHCPv6 at the `[edit system services dhcp-local-server dhcpv6]` or for a specific group at the `[edit system services dhcp-local-server group name]` hierarchy level, and so on.

NOTE: Reauthentication does not support Extensible Services Subscriber Management (essmd) services. Activation or deactivation of any such service causes the request to fail.

Options

lease-renewal Reauthenticate when a renew, rebind, discover, or solicit message is received from the DHCP client. This re-authentication is an alternative to RADIUS Change of Authorization (CoA) to change subscriber session characteristics based on the actual data rate change without interrupting the subscriber service.

remote-id-mismatch Reauthenticate when a discover or solicit message is received from the DHCP client with a new value for the DHCP client's Agent Remote ID. The change in value corresponds to a change in subscriber service plan. The Agent Remote ID is conveyed in option 82, suboption 2 for DHCPv4 clients and in option 37 for DHCPv6 clients.

actual-data-rate-change Optical line terminal (OLT) adds `option 82` with `sub-option 9` with Broadband Forum (Vendor ID 3561). It contains the sub-attributes `Actual-Data-Rate-Upstream` and `Actual-Data-Down-Stream` encoded. The decoded upstream and downstream data send to RADIUS server as a part of authentication request. RADIUS server analyzes the data received. Based on the `Actual-Data-Rate-Upstream` and `Actual-Data-Down-Stream` values, the RADIUS server selects the service profile for the subscriber interface.

When the actual data rate changes, the DHCP server re-authenticates the subscriber service. This re-authentication is an alternative to RADIUS Change of Authorization (CoA) to change subscriber session characteristics based on the actual data rate change without interrupting the subscriber service.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

Support at the `[edit ... system services dhcp-local-server dual-stack-group]` hierarchy level introduced in Junos OS Release 18.1R1.

`actual-data-rate-change` option introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Configuring RADIUS Reauthentication for DHCP Subscribers

RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers

rebalance (Aggregated Ethernet Interfaces)

IN THIS SECTION

Syntax | 14107

- Hierarchy Level | 14107
- Description | 14107
- Options | 14107
- Required Privilege Level | 14107
- Release Information | 14108

Syntax

```
rebalance interval
```

Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options load-balance-stateful per-flow]
```

Description

Configure periodic rebalancing of traffic flows of an aggregated Ethernet bundle by clearing the load balance state at a specified interval.

Options

interval Number of minutes after which the load-balancing state must be cleared for the specified interface.

- **Range:** 1 through 1000 flows per minute

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces

rebalance-periodic (Aggregated Ethernet Subscriber Interfaces)

IN THIS SECTION

- Syntax | 14108
- Hierarchy Level | 14108
- Description | 14108
- Options | 14109
- Required Privilege Level | 14109
- Release Information | 14109

Syntax

```
rebalance-periodic time hour:minute <interval hours>
```

Hierarchy Level

```
[edit interfaces ae number aggregated-ether-options]
```

Description

Configure periodic rebalancing of distribution of subscribers on an aggregated Ethernet bundle.

Options

hour:minute Time at which the rebalancing occurs, in military time.

hours Interval at which the rebalancing occurs, in hours. Default: 24 hours.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface

rebalance-periodic (PS Interface)

IN THIS SECTION

- Syntax | 14110
- Hierarchy Level | 14110
- Description | 14110
- Options | 14110
- Required Privilege Level | 14110
- Release Information | 14110

Syntax

```
rebalance-periodic {
        interval hours;
        start-time hour:minute;
```

Hierarchy Level

```
[edit interfaces ps number]
```

Description

Configure periodic rebalancing for distribution of subscribers on a redundant logical tunnel (RLT) bundle.

Options

interval (*hours*) Interval at which the rebalancing occurs, in hours. Default: 24 hours.

start-time (*hour:minute*) Time at which the rebalancing occurs, in military time.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

MPLS Pseudowire Subscriber Logical Interfaces

rebalance-subscriber-granularity (PS interface)

IN THIS SECTION

- Syntax | 14111
- Hierarchy Level | 14111
- Description | 14111
- Default | 14112
- Required Privilege Level | 14112
- Release Information | 14112

Syntax

```
rebalance-subscriber-granularity <subscriber-granularity-value>;
```

Hierarchy Level

```
[edit interfaces psnumber targeted-options]
```

Description

Rebalancing takes place when the member links have more subscribers than the configured value on the rebalance-subscriber-granularity option. Changing the value to a value less than the current active value does not force IFLs to be reassigned to a different pseudo IFL.

BEST PRACTICE: We recommend that you do not configure a low granularity value. A low value can have undesirable effects, such as the router running out of pseudo logical interfaces or an increase in the convergence time for rebalancing.

Leaving subscriber granularity at the default value of 500 subscribers is sufficient in most cases. Whenever more than one member link is active, targeted distribution places new subscribers on a link with fewer subscribers than other member links on the interface.

Default

The default subscriber granularity value is 500 subscribers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

MPLS Pseudowire Subscriber Logical Interfaces

receive-ais

IN THIS SECTION

- Syntax | 14112
- Hierarchy Level | 14113
- Description | 14113
- Required Privilege Level | 14113
- Release Information | 14113

Syntax

```
receive-ais;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
event ais-trigger-condition]
```

Description

Send a message to the peer MEPs when an AIS message is received by a peer MEP at its own maintenance level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Ethernet Alarm Indication

Ethernet Alarm Indication

receive-failure-threshold (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 14114
- Hierarchy Level | 14114
- Description | 14114
- Required Privilege Level | 14114
- Release Information | 14114

Syntax

```
receive-failure-threshold;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specifies the failure threshold value of the received test frames.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

receive-options-packets (Passive Flow Monitoring)

IN THIS SECTION

Syntax | **14115**

- Hierarchy Level | 14115
- Description | 14115
- Required Privilege Level | 14115
- Release Information | 14115

Syntax

```
receive-options-packets;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

When you enable passive monitoring, this statement is required for conformity with cflowd records structure.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers

receive-options-packets

IN THIS SECTION

- Syntax | 14116
- Hierarchy Level | 14116
- Description | 14116
- Required Privilege Level | 14116
- Release Information | 14117

Syntax

```
receive-options-packets;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet]
```

Description

For a Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Enabling Passive Monitoring on ATM Interfaces
- Enabling Packet Flow Monitoring on SONET/SDH Interfaces

receive (Protocols RIP)

IN THIS SECTION

- Syntax | 14117
- Hierarchy Level | 14117
- Description | 14118
- Options | 14118
- Required Privilege Level | 14118
- Release Information | 14118

Syntax

```
receive receive-options;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip],
```

```
[edit protocols rip group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

Configure RIP receive options.

Options

receive-options—One of the following:

- **both**—Accept both RIP version 1 and version 2 packets.
- **none**—Do not receive RIP packets.
- **version-1**—Accept only RIP version 1 packets.
- **version-2**—Accept only RIP version 2 packets.
- **Default: both**

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets*
>
> send (Protocols RIP) | **16546**

receive (Protocols RIPng)

IN THIS SECTION

- Syntax | 14119
- Hierarchy Level | 14119
- Description | 14119
- Options | 14120
- Required Privilege Level | 14120
- Release Information | 14120

Syntax

```
receive <none>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng],
[edit logical-systems logical-system-name protocols ripng group group-name neighbor neighbor-
name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng group group-name neighbor neighbor-name],
[edit protocols ripng],
[edit protocols ripng group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols ripng],
[edit routing-instances routing-instance-name protocols ripng group group-name neighbor neighbor-
name]
```

Description

Enable or disable receiving of update messages.

Options

none—(Optional) Disable receiving update messages.

- **Default:** Enabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

send (Protocols RIPng) | **16548**

receive-ttl-exceeded (Passive Flow Monitoring)

IN THIS SECTION

- Syntax | **14121**
- Hierarchy Level | **14121**
- Description | **14121**
- Required Privilege Level | **14121**
- Release Information | **14121**

Syntax

```
receive-ttl-exceeded;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

When you enable passive monitoring, this statement is required for conformity with cflowd records structure.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers

receive-ttl-exceeded

IN THIS SECTION

- Syntax | 14122
- Hierarchy Level | 14122

- Description | 14122
- Required Privilege Level | 14122
- Release Information | 14122

Syntax

```
receive-ttl-exceeded;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet]
```

Description

For Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Enabling Passive Monitoring on ATM Interfaces
- Enabling Packet Flow Monitoring on SONET/SDH Interfaces

receive-window

IN THIS SECTION

- Syntax | 14123
- Hierarchy Level | 14123
- Description | 14123
- Options | 14123
- Required Privilege Level | 14124
- Release Information | 14124

Syntax

```
receive-window packets;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Specify the size of the receive window for L2TP tunnels, which limits the number of packets the server processes concurrently.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options

packets—Maximum number of packets the receive window can hold at one time.

- **Default:** 16

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Window Size for L2TP Tunnels

receiver

IN THIS SECTION

- Syntax | 14124
- Hierarchy Level | 14125
- Description | 14125
- Default | 14125
- Options | 14125
- Required Privilege Level | 14126
- Release Information | 14126

Syntax

```
receiver {
    install;
    mode (proxy | transparent);
    (source-list | source-vlans) vlan-list;
```

```
    translate;
}
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan (vlan-name) data-forwarding]
[edit logical-systems name protocols igmp-snooping vlan vlan-name data-forwarding]
```

Description

Configure a VLAN as a multicast receiver VLAN of a multicast source VLAN (MVLAN) using the multicast VLAN registration (MVR) feature.

You must associate an MVR receiver VLAN with at least one data-forwarding source MVLAN. You can configure an MVR receiver VLAN with multiple source MVLANs using the `source-list` or `source-vlans` statement.

The remaining statements are explained separately.

> **NOTE**: The `mode`, `source-list`, and `translate` statements are only applicable to MVR configuration on EX Series switches that support the Enhanced Layer 2 Software (ELS) configuration style.
>
> The `source-vlans` statement is applicable only to EX Series switches that do not support ELS, and is equivalent to the ELS `source-list` statement.

See CLI Explorer.

Default

MVR not enabled

Options

install Install forwarding table entries (also called bridging entries) on the MVR receiver VLAN when MVR is enabled. By default, MVR only installs bridging entries on the source MVLAN for a group address.

You cannot configure the `install` option for a data-forwarding receiver VLAN that is configured in proxy mode (see the MVR "mode" on page 10809 option). In MVR

transparent mode, by default, the device installs bridging entries only on the MVLAN for a multicast group, so upon receiving MVR receiver VLAN traffic for that group, the switch doesn't forward the traffic to receiver ports on the MVR receiver VLAN that sent the join message for that group. The traffic is only forwarded on the MVLAN to MVR receiver interfaces. Configure this option when in transparent mode to enable MVR receiver VLAN ports to receive traffic forwarded on the MVR receiver VLAN.

mode (proxy | transparent)

(ELS devices only) Set proxy or transparent mode for an MVR receiver VLAN. This statement is explained separately. The mode is `transparent` by default.

source-list *vlan-list*

(ELS devices only) Specify a list of multicast source VLANs (MVLANs) from which a multicast receiver VLAN receives multicast traffic when multicast VLAN registration (MVR) is configured. This option is available only on on-ELS devices. (Use the `source-vlans` option for the same function on non-ELS switches.)

source-vlans *vlan-list*

(Non-ELS switches only) Specify a list of MVLANs for MVR operation from which the MVR receiver VLAN receives multicast traffic when multicast VLAN registration (MVR) is configured. Either all of these MVLANs must be in proxy mode or none of them can be in proxy mode (see "proxy" on page 13834). This option is available only on non-ELS switches. (Use the `source-list` option for the same function on ELS devices.)

translate

(ELS devices only) Translate VLAN tags in multicast VLAN (MVLAN) packets from the MVLAN tag to the multicast receiver VLAN tag on an MVR receiver VLAN. Without this option, tagged traffic has the MVLAN ID by default.

We recommend you set this option for MVR receiver VLANs with trunk ports, so hosts on the trunk interfaces receive multicast traffic tagged with the expected VLAN ID (the MVR receiver VLAN ID).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Statement and `mode`, `source-list`, and `translate` options introduced in Junos OS Release 18.3R1 for EX4300 switches (ELS switches).

Statement and `mode`, `source-list`, and `translate` options added in Junos OS Release 18.4R1 for EX2300 and EX3400 switches (ELS switches).

RELATED DOCUMENTATION

Understanding Multicast VLAN Registration

Configuring Multicast VLAN Registration on EX Series Switches

Example: Configuring Multicast VLAN Registration on EX Series Switches Without ELS

receiver (Graceful Restart for BGP Helper)

IN THIS SECTION

- Syntax | 14127
- Hierarchy Level | 14127
- Description | 14128
- Options | 14128
- Required Privilege Level | 14128
- Release Information | 14128

Syntax

```
receiver {
    enable:
    disable;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp graceful-restart long-lived],
[edit logical-systems logical-system-name protocols bgp group group-name graceful-restart long-
lived],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
graceful-restart long-lived],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp graceful-
restart long-lived],
```

```
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name graceful-restart long-lived],
[edit logical-systems logical-system-name routing-instances instance-name protocols bgp group
group-name neighbor address graceful-restart],
[edit routing-instances instance-name protocols bgp graceful-restart long-lived],
[edit routing-instances instance-name protocols bgp group group-name graceful-restart long-
lived],
[edit routing-instances instance-name protocols bgp group group-name neighbor address graceful-
restart],
[edit protocols bgp graceful-restart long-lived],
[edit protocols bgp group group-name graceful-restart long-lived],
[edit protocols bgp group group-name neighbor address graceful-restart long-lived]
```

Description

Enable the long-lived graceful restart mechanism for a BGP receiver or helper router to preserve BGP routing details for a longer period from a failed BGP peer. Long-lived graceful restart receiver or helper mode is enabled by default, unless ordinary graceful restart receiver or helper mode is disabled.

The `long-lived receiver enable` overrides a disable option inherited from a higher level in the configuration. When the LLGR receiver or helper mode is enabled or disabled, the session is reset. This behavior enables the new capability value to be sent to the neighbor.

Options

enable Enable long-lived BGP graceful restart for a receiver or helper router

disable Disable long-lived BGP graceful restart for a receiver or helper router

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

recommendation-tlv

IN THIS SECTION

- Syntax | 14129
- Hierarchy Level | 14129
- Description | 14129
- Default | 14130
- Options | 14130
- Required Privilege Level | 14130
- Release Information | 14130

Syntax

```
recommendation-tlv {
    no-auto-negotiation;
}
```

Hierarchy Level

```
[edit protocols dcbx interface interface-name enhanced-transmission-selection]
```

Description

Enable DCBX to send the ETS Recommendation TLV (also known as the Information TLV) on egress. This feature is valid only if the interface DCBX mode is IEEE DCBX. If the interface DCBX mode is DCBX

version 1.01, this statement has no effect. (DCBX version 1.01 does not advertise separate TLVs for individual attributes.)

Default

DCBX-enabled interfaces send the ETS recommendation TLV unless it is disabled.

Options

`no-auto-negotiation`—Disable sending of the ETS recommendation TLV.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

show dcbx neighbors | **19191**

Configuring DCBX Autonegotiation

recommended

IN THIS SECTION

- Syntax | **14131**
- Hierarchy Level | **14131**
- Description | **14131**
- Required Privilege Level | **14131**
- Release Information | **14131**

Syntax

```
recommended;
```

Hierarchy Level

```
[edit security idp dynamic-attack-group dynamic-attack-group-name filters]
```

Description

Specify recommended filter to add predefined attacks recommended by Juniper Networks to the dynamic attack group.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

recommended-action

IN THIS SECTION

- Syntax | 14132
- Hierarchy Level | 14132
- Description | 14132
- Options | 14132
- Required Privilege Level | 14132
- Release Information | 14132

Syntax

```
recommended-action (close | close-client | close-server | drop | drop-packet | ignore | none);
```

Hierarchy Level

```
[edit security idp custom-attack attack-name]
```

Description

When the security device detects an attack, it performs the specified action.

Options

The seven actions are as follows, from most to least severe:

- `close`—Reset the client and the server.
- `close-client`—Reset the client.
- `close-server`—Reset the server.
- `drop`—Drop the particular packet and all subsequent packets of the flow.
- `drop-packet`—Drop the particular packet of the flow.
- `ignore`—Do not inspect any further packets.
- `none`—Do not perform any action.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

reconfigure (DHCP Local Server)

IN THIS SECTION

- Syntax | 14133
- Hierarchy Level | 14133
- Description | 14134
- Options | 14134
- Required Privilege Level | 14134
- Release Information | 14134

Syntax

```
reconfigure {
    attempts attempt-count;
    clear-on-terminate;
    strict;
    support-option-pd-exclude;
    timeout timeout-value;
    token token-value;
     trigger {
        radius-disconnect;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name],
```

```
[edit logical-systems logical-system-name system services dhcp-local-server],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name],
[edit routing-instances routing-instance-name system services dhcp-local-server],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-
name],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name]
```

Description

Enable dynamic reconfiguration triggered by the DHCP local server of all DHCP clients or only the DHCP clients serviced by the specified group of interfaces. A group configuration takes precedence over a DHCP local server configuration. The `strict` statement is available only for DHCPv6.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

support-option-pd-exclude Request to exclude prefix option in the reconfigure message.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

`support-option-pd-exclude` statement introduced in Junos OS Release 17.3 for the MX Series.

RELATED DOCUMENTATION

Dynamic Reconfiguration of DHCP Servers and Clients

reconnect-time

IN THIS SECTION

- Syntax | 14135
- Hierarchy Level | 14135
- Description | 14135
- Options | 14135
- Required Privilege Level | 14136
- Release Information | 14136

Syntax

```
reconnect-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp graceful-restart],
[edit protocols ldp graceful-restart],
[edit routing-instances routing-instance-name protocols ldp graceful-restart]
```

Description

Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.

Options

seconds—Time required for reconnection.

- **Range:** 30 through 300
- **Default:** 60 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Configuring LDP Graceful Restart

MPLS Applications User Guide

Configuring Graceful Restart Options for LDP

record

IN THIS SECTION

- Syntax | 14137
- Hierarchy Level | 14137
- Description | 14137
- Default | 14137
- Required Privilege Level | 14137
- Release Information | 14137

Syntax

```
(record | no-record);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path  (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

Specify whether an LSP should actively record the path by sending the record route object (RRO). The RRO is used to record the path that the LSP traverses. It includes the IP address and router/node ID of each RSVP node in the path. Recording the LSP path can be useful for diagnostics and also provides loop detection.

Default

Recording LSP path is enabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Disabling Path Route Recording by LSPs

record (Origin Validation for BGP)

IN THIS SECTION

- Syntax | 14138
- Hierarchy Level | 14138
- Description | 14138
- Options | 14139
- Required Privilege Level | 14139
- Release Information | 14139

Syntax

```
record destination {
    maximum-length prefix-length {
        origin-autonomous-system as-number {
            validation-state (invalid | valid);
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation static],
[edit logical-systems logical-system-name routing-options validation static],
[edit routing-instances instance-name routing-options validation static],
[edit routing-options validation static]
```

Description

Configure the network prefix for the route validation (RV) record.

An RV record matches any route whose prefix matches the RV prefix, whose prefix length does not exceed the "maximum-length" on page 10531 given in the RV record, and whose origin AS equals the "origin-autonomous-system" on page 11932 given in the RV record. RV records are received from the cache server using the protocol defined in Internet draft draft-ietf-sidr-rpki-rtr-19, *The RPKI/Router Protocol*, and can also be configured statically, as shown here.

Options

destination—Network prefix for the RV record.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- *Use Case and Benefit of Origin Validation for BGP*
- *Understanding Origin Validation for BGP*
- *Example: Configuring Origin Validation for BGP*

record-route-option (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 14140
- Hierarchy Level | 14140
- Description | 14140
- Required Privilege Level | 14140

Release Information | **14140**

Syntax

```
record-route-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop IPv4 packets that have the IP option of 7 (Record Route).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

recover-time (DDoS)

IN THIS SECTION

- Syntax | 14141
- Hierarchy Level | 14141
- Description | 14141
- Options | 14141
- Required Privilege Level | 14142
- Release Information | 14142

Syntax

```
recover-time seconds;
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group (aggregate | packet-type)]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure how much time must pass since the last detected DDoS violation before the traffic is considered to have recovered from the attack and returned to normal.

Options

seconds—Period required for the traffic to recover.

- **Range:** 1 through 3600 seconds
- **Default:** 300

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers

recovered-clock

IN THIS SECTION

- Syntax | 14142
- Hierarchy Level | 14143
- Description | 14143
- Options | 14143
- Required Privilege Level | 14143
- Release Information | 14143

Syntax

```
recovered-clock {
    port port-number;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number]
```

Description

Select the port where the synchronous clock may be recovered.

Options

port-number—Port number where the synchronous clock may be recovered.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X48.

RELATED DOCUMENTATION

Clock Sources for PTX Series Packet Transport Routers

Synchronizing Internal Stratum 3 Clock to External Clock Sources on PTX Series Routers

synchronization (PTX Series) | **28802**

recovery

IN THIS SECTION

- Syntax | **14144**
- Hierarchy Level | **14144**

- Description | 14144
- Options | 14144
- Required Privilege Level | 14145
- Release Information | 14145

Syntax

```
recovery {
    automatic;
    configuration-only;
    configuration-preference;
    no-automatic;
    on-upgrade;
}
```

Hierarchy Level

```
[edit system snapshot]
```

Description

Configure system snapshot recovery management options.

Options

automatic

This option automatically recovers the Junos volume when necessary.

When you issue the `request system snapshot` command, the configuration on the device is stored in the OAM volume along with the creation of the recovery snapshot.

configuration-only

This option allows you to create or update a recovery snapshot along with storing the configuration on the device.

When you configure this option, the following scenarios occur:

- If a recovery snapshot does not exist for the current Junos image, issuing the `request system snapshot` command create a recovery snapshot along with storing the configuration on the device.
- If a recovery snapshot exists for the current image and additional packages are added or removed, a new recovery snapshot is created. The configuration changes are stored along with the new snapshot on the device.
- If a recovery snapshot exists for the current image and no changes are made to the current image, only the configuration changes are stored to the existing recovery snapshot. No new recovery snapshot is created.

configuration-preference This option allows you to specify a preferred configuration when booting from the OAM volume.

no-automatic When you configure this option, the system does not recover the Junos volume automatically.

on-upgrade Update the recovery snapshot when the Junos release changes.

This option enables the system to automatically create a recovery snapshot when you load a different version of the Junos image on the device. The system generates a recovery snapshot with the new Junos image version and stores the configuration on the device.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

recovery

IN THIS SECTION

- Syntax | 14146
- Hierarchy Level | 14146
- Description | 14146
- Options | 14146
- Required Privilege Level | 14147
- Release Information | 14147

Syntax

```
recovery {
    (auto | manual);
    timer timer-value;
}
```

Hierarchy Level

```
[edit interfaces interfaces-name link-degrade-monitor]
```

Description

Configure the mechanism to be used to recover a degraded link. The recovery options supported are auto and manual.

Options

auto Recover a degraded link automatically. Use this option with the media-based action when there are no Layer 2 or Layer 3 protocols configured on the interface. If this option is configured, the degraded link is monitored at user-configured intervals; and if the link quality

is found to have improved (if bit error rate hits the clear threshold), the link is automatically recovered. With this configuration, you must configure a timer value.

manual Recover a degraded link manually. Use this option with the media-based action configuration when Layer 2 and Layer 3 protocols are configured on the interface. If this option is configured, you need to use the `request interface link-degrade-recover` *`interface-name`* statement to recover the link.

> **NOTE**: The manual recovery option is recommended for user deployments that have static route configurations causing the remote end of the link to start forwarding packets (as soon as the physical link is up) while autorecovery is in progress.

timer ***timer-value*** Specify the interval value (in seconds) after which autorecovery of the degraded link must be triggered. This option is applicable if you configure the autorecovery option. The interval period starts from the time the link is degraded. The default interval is 1800 seconds. The autorecovery attempt is repeated until the link is recovered or the link monitoring feature is disabled through configuration.

> **NOTE**: During autorecovery, you might notice link flaps at the remote end of the link.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Link Degrade Monitoring Overview

link-degrade-monitor | **9784**

thresholds | **29346**

request interface link-degrade-recover | **14642**

recovery-time (LDP)

IN THIS SECTION

- Syntax | 14148
- Hierarchy Level | 14148
- Description | 14148
- Options | 14148
- Required Privilege Level | 14149
- Release Information | 14149

Syntax

```
recovery-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp
graceful-restart],
[edit protocols ldp graceful-restart],
[edit routing-instances routing-instance-name protocols ldp graceful-restart]
```

Description

Specify the length of time a router or switch waits for Label Distribution Protocol (LDP) neighbors to assist it with a graceful restart.

Options

seconds—Time the router waits for LDP to restart gracefully.

- **Range:** 120 through 1800

- **Default:** 160

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for LDP

Configuring Recovery Time and Maximum Recovery Time

maximum-neighbor-recovery-time | **10538**

recovery-timeout

IN THIS SECTION

- Syntax | **14150**
- Hierarchy Level (EX Series and QFX Series) | **14150**
- Hierarchy Level (MX Series) | **14150**
- Description | **14150**
- Default | **14151**
- Options | **14151**
- Required Privilege Level | **14151**
- Release Information | **14151**

Syntax

```
recovery-timeout seconds;
```

Hierarchy Level (EX Series and QFX Series)

```
[edit interfaces interface-name unit 0 family ethernet-switching]
```

Hierarchy Level (MX Series)

```
[edit interfaces interface-name unit 0 family bridge]
```

Description

Configure an interface to be temporarily disabled when MAC limiting, MAC move limiting, or rate-limiting is in effect with the action `shutdown`. This enables the affected interface to recover automatically from the error condition after the specified period of time:

- If you configure MAC limiting with the `shutdown` option and you enable `recovery-timeout`, the interface is temporarily disabled when the MAC address limit is reached. The interface will recover automatically after the number of seconds specified.
- If you enable MAC move limiting with the `shutdown` option and you enable `recovery-timeout`, the interface is temporarily disabled when the maximum number of MAC address moves is reached. The interface will recover automatically after the number of seconds specified.
- If you enable MAC move limiting with the `vlan-member-shutdown` option and you enable `recovery-timeout`, the interface is temporarily disabled when the maximum number of MAC address moves is reached. The interface will recover automatically after the number of seconds specified. If the recovery timeout is not configured, the interface will recover automatically after 180 seconds.
- If you enable storm control with the `action-shutdown` option and you enable `recovery-timeout`, the interface is temporarily disabled when applicable traffic exceeds the specified levels. Depending upon the configuration, applicable traffic can include broadcast, unknown unicast, and multicast traffic.

NOTE: The `recovery-timeout` configuration does not apply to preexisting error conditions. It impacts only error conditions that are detected after the `recovery-timeout` statement is

configured and committed. To clear a preexisting error condition and restore the interface to service, use the operational mode commands `clear ethernet-switching recovery-timeout` for EX Series and QFX Series and `clear bridge recovery-timeout` for MX Series routers.

Default

The interface does not automatically recover from an error condition.

NOTE: On EX9200 switches, if a MAC move limit is configured with the action `vlan-member-shutdown`, the interface automatically recovers from the disabled condition after 180 seconds by default.

Options

seconds Number of seconds that the interface remains in a disabled state due to a port error prior to automatic recovery.

- **Range:** 10 through 3600

Required Privilege Level

system—To view this statement in the configuration.
system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

action-shutdown | **1787**

clear ethernet-switching recovery-timeout (SRX) | **3583**

Understanding MAC Limiting

Configuring Autorecovery From the Disabled State on Secure Interfaces of a Security Device (CLI Procedure)

Configuring MAC Limiting (ELS)

Configuring MAC Move Limiting (ELS)

Enabling and Disabling Storm Control (ELS)

Example: Configuring MAC Limiting on a Security Device

recycle-bandwidth-profiles

IN THIS SECTION

- Syntax | 14152
- Hierarchy Level | 14152
- Description | 14153
- Options | 14153
- Required Privilege Level | 14153
- Release Information | 14153

Syntax

```
recycle-bandwidth-profiles profile name {
 ingress-egress-mcast percentage;
 srv6 percentage;
 srv6-ext percentage;
 tunnel-services percentage;
 }
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

This statement allows you to enable and configure a recycle bandwidth profile. This profile holds the map of the recycle applications and the percentage of recycle bandwidth value assigned to each application within that profile.

There is no cap on the number of recycle bandwidth profiles.

Options

The table below indicates supported recycle applications for which you can reserve recycle bandwidth. This list is not exhaustive.

> **NOTE**: Since the values are in percentage, the cumulative recycle bandwidth should not exceed 100.

Application	Description
ingress-egress-mcast	Set bandwidth in percentage to be allocated for ingress-egress-multicast replication.
srv6	Set bandwidth in percentage to be allocated for srv6 replication.
srv6-ext	Set bandwidth in percentage to be allocated for srv6-ext replication.
tunnel-services	Set bandwidth in percentage to be allocated for tunnel-services replication.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.4R1.

Statement introduced in Junos OS Evolved Release 24.1R1 for ACX7K Series Routers.

RELATED DOCUMENTATION

Recycle Bandwidth Management

Example: Configuring Recycle Bandwidth

red-buffer-occupancy

IN THIS SECTION

- Syntax | 14154
- Hierarchy Level | 14154
- Description | 14154
- Options | 14155
- Required Privilege Level | 14155
- Release Information | 14155

Syntax

```
red-buffer-occupancy {
    weighted-averaged [ instant-usage-weight-exponent exponent-value ];
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number ],
[edit chassis lcc number fpc slot-number pic pic-number]
```

Description

Configure the IQ PIC to base random early detection (RED) queue management on a *simple moving average* buffer occupancy calculation. If you do not include this statement, the IQ PIC bases RED on an *instantaneous* buffer occupancy value. As an option, you can specify that the IQ PIC bases RED on a *weighted moving average* of buffer occupancy values.

If you configure this feature on a channelized OC12 intelligent queuing (IQ) PIC, the PIC reboots.

Options

`weighted-averaged`—Configure the IQ PIC to base RED processing on a simple moving average of instantaneous buffer occupancy values instead of an instantaneous buffer occupancy.

`instant-usage-weight-exponent` *`exponent-value`*—(Optional) Specify the integer to be used as the negative exponent of 2 to express a weight value. The PIC performs weighted RED (WRED) by based on a calculation of average buffer occupancy that applies the specified weight value to the instantaneous buffer occupancy and then factors the weighted value into the calculation of average buffer occupancy. Valid exponent range is from 1 through 31 (weight values from 2^{-1} through 2^{-31}). If you do not specify this option, the default exponent value is 0, which results in a weight value of $2^0 = 1$. With a weight value of 1, the calculation of weighted average buffer occupancy yields the same value as the simple average buffer occupancy.

> **NOTE**: You can specify an exponent value greater than 31, and the value displays in the output of `show` commands. However, the PIC replaces the out-of-range value with the *operational* value of 31, which results in a weight value of $2^{-31} = 1 / 2^{31} = 0.0000000004656612873077392578125$.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Managing Transient Traffic Bursts by Configuring Weighted RED Buffer Occupancy

Example: Managing Transient Traffic Bursts by Configuring Weighted RED Buffer Occupancy

red-differential-delay

IN THIS SECTION

- Syntax | 14156
- Hierarchy Level | 14156
- Description | 14156
- Options | 14156
- Required Privilege Level | 14157
- Release Information | 14157

Syntax

```
red-differential-delay milliseconds;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

For link services interfaces only, configure the red differential delay among bundle links to give warning when a link has a differential delay that exceeds the configured threshold.

Options

milliseconds—Red differential delay threshold.

- **Range:** 1 through 2000 milliseconds
- **Default:** 120 milliseconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16

action-red-differential-delay | **1782**

yellow-differential-delay | **31168**

redial-delay

IN THIS SECTION

- Syntax | **14157**
- Hierarchy Level | **14158**
- Description | **14158**
- Options | **14158**
- Required Privilege Level | **14158**
- Release Information | **14158**

Syntax

```
redial-delay time;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with interfaces configured for ISDN with dialout, specify the delay (in seconds) between two successive calls made by the dialer. To configure callback mode, include the `callback` statement at the `[edit interfaces dl*n* unit *logical-unit-number* dialer-options]` hierarchy level.

If the `callback` statement is configured, you cannot use the `caller *caller-id*` statement at the `[edit interfaces dl*n* unit *logical-unit-number* dialer-options]` hierarchy level.

Options

time—Delay (in seconds) between two successive calls.

- **Range:** 2 through 255 seconds
- **Default:** 3 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

ISDN Interfaces Overview

redirect (PCC Action Profiles)

IN THIS SECTION

- Syntax | 14159
- Hierarchy Level | 14159
- Description | 14159
- Required Privilege Level | 14160
- Release Information | 14160

Syntax

```
redirect {
    url url-name;
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],
[edit services pcef pcc-action-profiles profile-name]
```

Description

Specify HTTP redirection to a URL. If you configure this, the PCC action profile can only be used in PCC rules that match only HTTP-based applications and all flows.

If you are using Junos OS Subscriber Aware, specify the redirection at the `[edit unified-edge pcef pcc-action-profiles profile-name]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the redirection at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

Configuring Policy and Charging Control Action Profiles for Subscriber Management

redirect-traffic

IN THIS SECTION

- Syntax | **14161**
- Hierarchy Level | **14161**
- Description | **14161**
- Options | **14161**
- Required Privilege Level | **14161**

Release Information | 14161

Syntax

```
redirect-traffic (all | unauthenticated);
```

Hierarchy Level

```
[edit services unified-access-control captive-portal policy]
```

Description

Specify to redirect traffic destined for protected sources to the IC Series device. You can choose to redirect all traffic or only unauthenticated traffic.

Options

- `all`—Redirect all traffic destined for the protected sources to the IC Series device. Specify this option if you want to redirect all traffic (IPsec or source IP) to the currently connected IC Series device or to an IP address or domain name that you specify in a redirect URL.
- `unauthenticated`—Redirect unauthenticated traffic destined for the protected sources to the IC Series device. Select this option if your deployment uses source IP only or a combination of source IP and IPsec. The Junos OS Enforcer redirects clear-text traffic from unauthenticated users to the currently connected IC Series device or to an IP address or domain name that you specify in a redirect URL.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Firewall User Authentication Overview

redirect-url

IN THIS SECTION

- Syntax | 14162
- Hierarchy Level | 14162
- Description | 14162
- Required Privilege Level | 14163
- Release Information | 14163

Syntax

```
redirect-url url;
```

Hierarchy Level

```
[edit services unified-access-control captive-portal policy]
```

Description

Specify to redirect traffic destined for protected sources to a specified URL.

You can configure the following options in the redirect URL string:

- *%dest-url%*—Specifies the protected resource which the user is trying to access.
- *%enforcer-id%*—Specifies the ID assigned to the Junos OS Enforcer by the IC Series device.
- *%policy-id%*—Specifies the encrypted policy ID for the security policy that redirected the traffic.

- *%dest-ip%*—Specifies the IP address or hostname of the protected resource that the user is trying to access.
- *%ic-ip%*—Specifies the IP address or hostname of the IC Series device to which the Junos OS Enforcer is currently connected.

If you do not specify the redirect URL, the Junos OS Enforcer uses the following default configuration:

```
https://%ic-ip%/?target = %dest-url%  &enforcer = %enforcer-id%  &policy = %policy-id%  &dest-ip
= %dest-ip%
```

NOTE: The maximum size of a redirect payload is 1450 bytes. The size of the redirect URL is restricted to 1407 bytes (excluding a few HTTP headers). If a user accesses a destination URL that is larger than 1407 bytes, the Infranet Controller authenticates the payload, calculates the exact length of the redirect URL, and trims the destination URL so that it can fit into the redirect URL. The destination URL can be fewer than 1407 bytes based on what else is present in the redirect URL (for example, policy ID). The destination URL in the default redirect URL is trimmed so that the redirect packet payload size is limited to 1450 bytes. If the length of the payload is larger than 1450 bytes, the excess length is trimmed and the user is directed to the destination URL that has been resized to 1450 bytes.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Firewall User Authentication Overview

redirect-wx (Application Services)

IN THIS SECTION

- Syntax | 14164
- Hierarchy Level | 14164
- Description | 14164
- Required Privilege Level | 14164
- Release Information | 14165

Syntax

```
redirect-wx;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
application-services]
```

Description

ForSRX300, SRX320, SRX340, SRX345, SRX380, or SRX550M devices, define the acceleration zone security policy for WX redirection of packets to the WXC Integrated Service Module (ISM 200) for WAN acceleration. During the redirection process, the direction of the WX packet and its type determine further processing of the packet.

Specify the WX redirection needed for the packets that arrive from the LAN. Use the `reverse-redirect-wx` statement to specify the WX redirection needed for the reverse flow of the packets that arrive from the WAN.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

redistribute-all-traffic (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 14165
- Hierarchy Level | 14165
- Description | 14165
- Required Privilege Level | 14166
- Release Information | 14166

Syntax

```
redistribute-all-traffic {
    enable-rejoin;
}
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options member-failure-options]
```

Description

Enable the option to redistribute traffic of a failed active member to the other active members.

For many-to-one (N:1) high availability support for Network Address Translation (NAT), the traffic for the failed member is automatically redistributed to the other active members.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

member-failure-options (Aggregated Multiservices) | **10608**

redundancy

IN THIS SECTION

- Syntax | **14167**
- Hierarchy Level | **14167**
- Description | **14167**
- Options | **14167**
- Required Privilege Level | **14167**
- Release Information | **14168**

Syntax

```
redundancy {
    cfeb slot (always | preferred);
    failover {
        on-disk-failure;
        on-loss-of-keepalives;
        on-re-to-fpc-stale;;
    }
    feb {
        redundancy-group group-name {
            description description;
            feb slot-number (backup | primary);
            no-auto-failover;
        }
    }
    graceful-switchover;
    keepalive-time seconds;
    routing-engine slot-number (backup | disabled | master);
    sfm slot-number (always | preferred);
    ssb slot-number (always | preferred);
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure redundancy options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routing Engine Redundancy

Configuring CFEB Redundancy on the M10i Router

Configuring FEB Redundancy on the M120 Router

Configuring SFM Redundancy on M40e and M160 Routers

Configuring SSB Redundancy on the M20 Router

redundancy-event (Services Redundancy Daemon)

IN THIS SECTION

- Syntax | 14168
- Hierarchy Level | 14169
- Description | 14169
- Options | 14169
- Required Privilege Level | 14169
- Release Information | 14169

Syntax

```
redundancy-event event-name {
    monitor {
        link-down interface-name;
        peer {
            (mastership-acquire | mastership-release);
        }
```

```
        process routing abort;
        process routing restart;
    }
]
```

Hierarchy Level

```
[edit  event-options]
```

Description

Configure events monitored to trigger change of primary role and routing using inter-chassis redundancy.

Options

event-name	Alphanumeric name for a monitored event.
link-down *interface-name*	Name of an interface, link, or link aggregation, to monitor.
peer mastership-acquire	(Optional) Monitor primary-role acquisition peer events.
peer mastership-release	(Optional) Monitor primary role release peer events.
process routing abort	(Optional, and only applies to Next Gen Services) Monitor process routing daemon (rpd) terminate requests.
process routing restart	(Optional) Monitor process routing daemon (rpd) restart requests.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring Inter-Chassis Services Redundancy for Next Gen Services

Configuring the Service Redundancy Daemon

redundancy (Graceful Switchover)

IN THIS SECTION

- Syntax | 14170
- Hierarchy Level | 14170
- Description | 14171
- Default | 14171
- Required Privilege Level | 14171
- Release Information | 14171

Syntax

```
redundancy {
    failover {
        on-disk-failure;
        on-loss-of-keepalives;
    }
    graceful-switchover;
}
```

Hierarchy Level

```
[edit chassis (EX Series)]
```

Description

Enable redundant Routing Engines on a Virtual Chassis with two or more member switches or on a Virtual Chassis Fabric, on a standalone EX6200 or EX8200 switch with more than one Routing Engine.

The remaining statements are explained separately. See CLI Explorer.

Default

Redundancy is enabled for the Routing Engines.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

- graceful-switchover (Chassis) | **7786**
- *Configuring Graceful Routing Engine Switchover in a Virtual Chassis*
- Configuring Graceful Routing Engine Switchover
- Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)
- High Availability Features for EX Series Switches Overview

redundancy-group

IN THIS SECTION

- Syntax | **14172**
- Hierarchy Level | **14172**
- Description | **14172**

- Options | 14172
- Required Privilege Level | 14172
- Release Information | 14173

Syntax

```
redundancy-group group-name {
    description description;
    feb slot-number (backup | primary);
    no-auto-failover;
}
```

Hierarchy Level

```
[edit chassis redundancy feb]
```

Description

On M120 routers only, configure a Forwarding Engine Board (FEB) redundancy group.

Options

group-name is the unique name for the redundancy group. The maximum length is 39 alphanumeric characters.

Other statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring FEB Redundancy on the M120 Router

redundancy-group (Chassis Cluster)

IN THIS SECTION

- Syntax | 14173
- Hierarchy Level | 14174
- Description | 14174
- Options | 14174
- Required Privilege Level | 14175
- Release Information | 14175

Syntax

```
redundancy-group  group-number {
    gratuitous-arp-count number;
    hold-down-interval number;
    interface-monitor interface-name {
        weight number;
    }
    ip-monitoring{
        family {
            inet {
                ipv4-address {
                    interface {
                        logical-interface-name;
                        secondary-ip-address ip-address;
```

```
                    }
                    weight number;
                }
            }
        }
        global-threshold number;
        global-weight number;
        retry-count number;
        retry-interval seconds;
    }
    node (0 | 1 ) {
        priority number;
    }
    preempt (Chassis Cluster){
        delay seconds;
        limit limit;
        period seconds;
    }
}
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Define a redundancy group. Except for redundancy group 0, a redundancy group is a logical interface consisting of two physical Ethernet interfaces, one on each chassis. One interface is active, and the other is on standby. When the active interface fails, the standby interface becomes active. The logical interface is called a redundant Ethernet interface (`reth`).

Redundancy group 0 consists of the two Routing Engines in the chassis cluster and controls which Routing Engine is primary. You must define redundancy group 0 in the chassis cluster configuration.

Options

group-number Redundancy group identification number.

> **NOTE**: The maximum number of redundancy groups is equal to the number of redundant Ethernet interfaces that you configure.

- **Range:** 0 through 128

interface-monitor Specify a redundancy group interface to be monitored for failover and the relative weight of the interface.

ip-monitor Specify IP address of interface to be monitored for end-to-end connectivity.

gratuitous-arp-count Number of gratuitous ARPs to send on an active interface after failover

- **Range:** `1-16`

hold-down-interval RG failover interval. RG0(300-1800) RG1+(0-1800) (seconds)

- **Range:** `0-1800`

node Identify each cluster node in a redundancy group and set its relative priority for primary role.

preempt Allow preemption of the primary system based on the priority within a redundancy group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

ip-monitoring | 9244

redundancy-group (Chassis - MX Series)

IN THIS SECTION

- Syntax | 14176
- Hierarchy Level | 14176
- Description | 14176
- Required Privilege Level | 14177
- Release Information | 14177

Syntax

```
redundancy-group {
    interface-type {
        redundant-logical-tunnel {
            device count;
        }
        redundant-virtual-tunnel {
            device count;
        }
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure redundant logical tunnels, redundant virtual tunnels, or both on MX Series 5G Universal Routing Platforms.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Redundant Virtual Tunnels Providing Resiliency in Delivering Multicast Traffic Overview

Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic

Example: Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic

redundancy-group (Redundant Tunnel) | **14180**

redundancy-group-id (Junos Fusion)

IN THIS SECTION

- Syntax | **14177**
- Hierarchy Level | **14178**
- Description | **14178**
- Default | **14178**
- Options | **14178**
- Required Privilege Level | **14178**
- Release Information | **14178**

Syntax

```
redundancy-group-id redundancy-group-id-number;
```

Hierarchy Level

```
[edit chassis satellite-management redundancy-groups redundancy-group-name]
```

Description

Defines a redundancy group ID number for a redundancy group in a Junos Fusion using dual aggregation devices.

A Junos Fusion using dual aggregation devices must have one redundancy group that includes both aggregation devices and all satellite devices. The redundancy group must be configured individually on each aggregation device. The redundancy group name and redundancy group ID number must match on the redundancy group configuration on each aggregation device in the Junos Fusion.

Default

A redundancy group does not have a redundancy group ID number by default.

Options

redundancy-group-id-number The user-defined redundancy group ID number.

- **Range:** 1 through 255

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

redundancy-group (Interfaces)

IN THIS SECTION

- Syntax | 14179
- Hierarchy Level | 14179
- Description | 14179
- Options | 14179
- Required Privilege Level | 14179
- Release Information | 14180

Syntax

```
redundancy-group number ;
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options]
```

Description

Specify the redundancy group that a redundant Ethernet interface belongs to.

Options

number —Number of the redundancy group that the redundant interface belongs to. Failover properties of the interface are inherited from the redundancy group.

- **Range:** 1 through 255

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Interfaces User Guide for Security Devices

redundancy-group (Redundant Tunnel)

IN THIS SECTION

- Syntax | 14180
- Hierarchy Level | 14181
- Description | 14181
- Options | 14181
- Required Privilege Level | 14181
- Release Information | 14181

Syntax

```
redundancy-group {
    member-interface interface-name {
        (active | backup);
    minimum-links number-of-links;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure member tunnels of redundant logical or virtual tunnels.

Options

active Set the interface to the active mode.

backup Set the interface to the backup mode.

minimum-links Specify the minimum number of active links required for the interface to remain up.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To view this statement in the configuration.

Release Information

`minimum-links` option introduced in Junos OS Release 24.2.

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Redundant Virtual Tunnels Providing Resiliency in Delivering Multicast Traffic Overview

Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic

Example: Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic

Example: Configuring Redundant Logical Tunnels

Configuring Redundant Logical Tunnels

Redundant Logical Tunnels Overview

redundancy-group (Chassis - MX Series) | **14176**

redundancy-interface-process

IN THIS SECTION

- Syntax | 14182
- Hierarchy Level | 14182
- Description | 14182
- Options | 14182
- Required Privilege Level | 14183
- Release Information | 14183

Syntax

```
redundancy-interface-process {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify as an active or backup process of an application server, configure to process traffic for more than one logical application server.

Options

command *binary-file-path*	Path to the binary process.
disable	Disable the redundancy interface management process.

failover Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.

alternate-media Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.

other-routing-engine Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

redundancy (Junos OS)

IN THIS SECTION

- Syntax | 14184
- Hierarchy Level | 14184
- Description | 14184
- Options | 14184
- Required Privilege Level | 14184
- Release Information | 14184

Syntax

```
redundancy {
    feed-redundancy {
        slot pem slot number;
    }
    source-redundancy;
}
```

Hierarchy Level

```
[edit chassis lcc name pem],
[edit chassis pem]
```

Description

PEM Source/Feed Redundancy configuration. This command is only supported on the MX10004 and MX10008 platforms with JNP10K-PWR-AC3 PSM.

Options

feed-redundancy	Toggle Feed Redundancy for PEM
source-redundancy	Toggle Source Redundancy for PEM

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Managing Power

redundancy (Junos OS Evolved)

IN THIS SECTION

- Syntax | 14185
- Hierarchy Level | 14185
- Description | 14185
- Options | 14186
- Required Privilege Level | 14186
- Release Information | 14186

Syntax

```
redundancy {
    feed-redundancy {
        slot psm slot number;
    }
    source-redundancy;
}
```

Hierarchy Level

```
[edit chassis psm]
```

Description

PSM Source/Feed Redundancy configuration. This command is only supported on the PTX10004, PTX10008, and PTX10016 platforms with JNP10K-PWR-AC3 PSM.

Options

source-redundancy	Toggle Source Redundancy for PSM
feed-redundancy	Toggle Feed Redundancy for PSM

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Managing Power

redundancy (M:N Subscriber Redundancy)

IN THIS SECTION

- Syntax | 14187
- Hierarchy Level | 14187
- Description | 14187
- Options | 14187
- Required Privilege Level | 14190
- Release Information | 14190

Syntax

```
redundancy {
    interface name {
        local-inet-address v4-address;
        local-inet6-address v6-address;
        shared-key string;
        virtual-inet-address virtual-inet-address;
        virtual-inet6-address virtual-inet6-address;
    }
    group group-name {
            interface interface-name {
                 standby-mode [hot-standby | service-activation-on-failover];
            }
        }
    }
    no-advertise-routes-on-backup;
    protocol {
        pseudo-wire;
        vrrp;
    }
}
```

Hierarchy Level

```
[edit system services subscriber-management]
```

Description

Enable M:N subscriber group redundancy. This global configuration applies to all subscribers in redundancy groups across the chassis.

For dual-stack subscribers, you must configure both virtual IPv4 and virtual IPv6 addresses on the interface.

Options

group *group-name*	Name of the redundancy group.

hot-standby Specifies the backup BNG, which is ready to switchover the traffic from the primary BNG immediately without any disruption during primary BNG failure.

interface *interface-name* Identify the interface name for the logical access interface for the subscriber redundancy group covered by M:N redundancy.

Configuring the interfaces inside the group is for logical grouping only. You can group all the interfaces to one interface-group or you can have single interface-group for each redundancy interfaces.

> **NOTE**: You must configure the names for all such interfaces on the chassis for all redundancy groups.
>
> For VRRP redundancy, only ge and xe interfaces are supported for the underlying physical interface. Ethernet connections can be untagged, single-tagged, or double-tagged.
>
> For pseudowire redundancy, only pseudowire interfaces are supported for the underlying physical interface.

> **NOTE**: You cannot configure an interface for redundancy if it already has active non-DHCP subscribers. If the interface has existing DHCP subscribers, then redundancy is enabled for those subscribers.

local-inet-address *v4-address* (Pseudowires only) Local IP address for the associated pseudowire interface. This address must match one of the access-facing GE interface addresses. It is unique per redundancy group identified by the ps*x*.0 interface. The DHCP relay agent uses this address as the giaddr for DHCP messages sent to a DHCP server. For dynamic subscriber Interfaces, this address is the same as the preferred-source address (the unnumbered address derived from the access-facing GE interface).

local-inet6-address *v6-address* (Pseudowires only) Local IPv6 address for the associated pseudowire interface. This address must match one of the access-facing GE interface addresses. It is unique per redundancy group identified by the ps*x*.0 interface. The DHCP relay agent uses this address as the linkaddr for DHCP messages sent to a DHCP server. For dynamic subscriber Interfaces, this address is the same as the preferred-source address (the unnumbered address derived from the access-facing GE interface).

no-advertise-routes-on-backup (Optional) Suppress advertisement of subscriber access routes or framed routes at the backup BNG towards the core or install the routes in the forwarding table. The routes are added to the routing table when the primary BNG fails over to the backup BNG. This option applies to all subscribers that are covered by redundancy on the chassis and that log in after you configure the option. Existing subscribers are not affected.

> **BEST PRACTICE**: We recommend that you always configure `no-advertise-routes-on-backup` when you use the non-aggregated mode of address allocation. In this mode, the DHCP server uses a common pool for all redundant subscribers and possibly all non-redundant subscribers. This mode increases the routing services overhead due to the number of routes being advertised for all subscribers. It increases the size of the routing table on the core side and can increase core-side convergence if a DHCP relay agent fails. It also increases downstream traffic convergence. The `no-advertise-routes-on-backup` option reduces the number of routes advertised and the associated potential issues.

- **Default:** Advertise the access and framed routes to the backup BNG during subscriber login.

protocol

Specify a method for BNG M:N redundancy. You can configure both VRRP and pseudowires when the underlying access interfaces do not overlap between these methods.

- `pseudo-wire`—Use pseudowires for redundancy on pseudowire interfaces for L2VPN and Layer 2 circuit-based networks over IP/MPLS.
- `vrrp`—Use VRRP for redundancy on GE and XE interfaces.
- **Default:** `vrrp`

shared-key ***string***

(Pseudowires only) Shared key that is used by the topology discovery process to identify a backup pseudowire interface on a peer BNG for a subscriber redundancy group. You configure the same shared-key on the local access interfaces on the primary and backup BNGs for a particular redundancy group. A topology discovery match is achieved when both the query and response messages include the pseudowire local access interface name, the same shared key, and the same transaction ID. You define the shared key as a string of up to 64 characters that you choose. You must not use the shared key for any other peer BNGs.

> **NOTE**: If the interface already has non-DHCP subscribers, the router rejects the configuration. The configuration is accepted for existing DHCP relay subscribers.

> **NOTE**: You cannot dynamically change a shared key after you have configured a key for a subscriber redundancy group (represented by the pseudowire interface). This means that if you change the value after committing it, the change has no effect. In order to have the change take effect, you must delete the entire redundancy group configuration and then add it back with the new key.

service-activation-on-failover
Specifies the interface mode on backup BNG, which uses less resources to forward traffic instantly on traffic switchover with basic statistics.

standby-mode
Specifies how to program the Packet Forwarding Engine with subscriber state when the interface is inactive.

virtual-inet-address ***virtual-inet-address***
(VRRP) IPv4 address for the virtual router. This address is used as the gateway address for all redundancy peers associated with a particular subscriber redundancy group. You must specify the same virtual IP address that you also configure on VRRP so that it can create the virtual router for redundancy.

virtual-inet6-address ***virtual-inet6-address***
(VRRP) IPv6 address for the virtual router. This address is used as the gateway address for all redundancy peers associated with a particular subscriber redundancy group. You must specify the same virtual IP address that you also configure on VRRP so that it can create the virtual router for redundancy.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

`local-inet-address`, `local-inet6-address`, `pseudo-wire`, and `shared-key` options added in Junos OS Release 20.1R1.

`service-activation-on-failover` introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

How to Configure M:N Subscriber Redundancy with VRRP and DHCP Binding Synchronization

redundancy-mode

IN THIS SECTION

- Syntax | 14191
- Hierarchy Level | 14191
- Description | 14191
- Options | 14192
- Required Privilege Level | 14192
- Release Information | 14192

Syntax

```
redundancy-mode (increased-bandwidth | redundant)
```

Hierarchy Level

```
[edit chassis fabric]
```

Description

Configure the active control boards to be in redundancy mode or increased fabric bandwidth mode.

In increased fabric bandwidth mode, which is the default behavior for MX Series routers with Switch Control Board (SCB), the maximum number of available fabric planes are used. The MX Series routers that contain the Enhanced SCB—SCBE—and the MPC3E, the control boards operate in redundancy fabric mode (all the FPCs use 4 fabric planes as active planes) by default.

In increased fabric bandwidth mode, which is the default mode on the SRX5600 and SRX5800 Services Gateways with the SRX5K-SCB4 Switch Control Board, six active planes (without any spare planes) are used. With the redundant fabric mode, the SRX5600 and SRX5800 Services Gateways use four active planes and two spare planes.

Options

increased-bandwidth Enable increased fabric bandwidth mode for the control boards, which causes all the available fabric planes to be used.

redundant Enable redundancy mode for the control boards, which causes all the FPCs to use 4 fabric planes as active planes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Detection and Corrective Actions of Line Cards on MX Series Routers

Detection and Recovery of Fabric-Related Failures Caused by Loss of Connectivity on MX Series Routers

MX Series Routers Fabric Resiliency

show chassis fabric redundancy-mode | **18181**

Configuring Fabric Redundancy Mode for Active Control Boards on MX Series Routers

redundancy-options (Adaptive Services Interfaces)

IN THIS SECTION

- Syntax | **14193**
- Hierarchy Level | **14193**
- Description | **14193**
- Options | **14193**

Required Privilege Level | 14193
Release Information | 14193

Syntax

```
redundancy-options {
    primary sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
    hot-standby
}
```

Hierarchy Level

```
[edit interfaces rspnumber]
[edit interfaces rmsnumber]
```

Description

Specify the primary and secondary (backup) adaptive services interfaces.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring AS or Multiservices PIC Redundancy

redundancy-options (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 14194
- Hierarchy Level | 14194
- Description | 14194
- Options | 14195
- Required Privilege Level | 14195
- Release Information | 14195

Syntax

```
redundancy-options {
    primary mams-a/b/0;
    secondary mams-a/b/0;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure warm standby for an aggregated multiservices (AMS) interface. Specify a primary and a secondary (backup) member services interface for the AMS interface. The primary interface is the service interface that you want to back up, and it is the active interface unless it fails. The secondary interface is the backup interface, and does not handle any traffic unless the primary interface fails. You can use the same services interface as the backup in multiple warm standby AMS interfaces.

You cannot use both the `redundancy-options` and the `load-balancing-options` statements in the same AMS interface.

Options

primary mams-*a/b*/0 Name of the primary services interface, where *a* is the FPC slot number and *b* is the PIC slot number.

secondary mams-*a/b*/0 Name of the secondary (backup) services interface, where *a* is the FPC slot number and *b* is the PIC slot number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring Warm Standby for Services Interfaces

redundancy-options (Link Services IQ PIC Interfaces)

IN THIS SECTION

- Syntax | 14196
- Hierarchy Level | 14196
- Description | 14196
- Options | 14196

Required Privilege Level | 14196
Release Information | 14196

Syntax

```
redundancy-options {
    (hot-standby | warm-standby);
    primary lsq-fpc/pic/port;
    secondary lsq-fpc/pic/port;
}
```

Hierarchy Level

```
[edit interfaces rlsqnumber]
```

Description

Specify the primary and secondary (backup) Link Services IQ PIC interfaces.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces

redundancy-options (Stateful Synchronization)

IN THIS SECTION

- Syntax | 14197
- Hierarchy Level | 14198
- Description | 14198
- Options | 14198
- Required Privilege Level | 14199
- Release Information | 14199

Syntax

```
redundancy-options {
    redundancy-local {
        data-address address;
    }
    redundancy-peer {
        ipaddress address;
    }
    replication-threshold seconds;
    routing-instance instance-name;
    apply-groups (apply-groups-except |redundancy-local | redundancy-peer)
    replication-options (apply-groups | apply-groups-except | mtu | replication-threshold |
replication-threshold routing-instance )
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Specify the primary and secondary (backup) adaptive services PIC interfaces.

Options

data-address *address*	Internal IP address of the local redundant PIC.
ipaddress *address*	Internal IP address of the remote redundant PIC.
instance-name	Name of the routing instance to apply to the HA synchronization traffic between the high availability pair.
seconds	Length of time that the flow remains active for replication. • **Default:** 180 seconds
apply-groups *apply-groups-except*	Specify the groups from which NOT to inherit the configuration.
apply-groups *redundancy-local*	Specify information for the local peer.
apply-groups *redundancy-peer*	Specify information for peer.
replication-options *apply-groups*	Specify groups from which to inherit the configuration.
replication-options *apply-groups-except*	Specify the groups from which NOT to inherit the configuration.
replication-options *mtu*	Specify the maximal packet size for the replicated data. • **Range:** 1500 through 8000 bytes
replication-options *replication-threshold*	Specify the duration for which flow should remain active for replication. • **Range:** 60 through 3600 seconds
replication-options *replication-threshold routing-instance*	Specify routing-instance for the HA traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card (interfaces of type `vms-x/y/z`).

RELATED DOCUMENTATION

Configuring Inter-Chassis Stateful Synchronization for Long Lived NAT and Stateful Firewall Flows (MS-MPC, MS-MIC) (Release 16.1 and later)

Inter-Chassis High Availability for MS-MIC and MS-MPC (Release 15.1 and earlier)

redundancy-policy (Interchassis Services Redundancy)

IN THIS SECTION

- Syntax | **14199**
- Hierarchy Level | **14200**
- Description | **14200**
- Options | **14200**
- Required Privilege Level | **14201**
- Release Information | **14201**

Syntax

```
redundancy-policy policy-name {
    redundancy-events [event-list] {
```

```
        then {
            acquire-mastership;
            <add-static-route destination {
                (next-hop next-hop | receive);
                routing-instance routing-instance
            }>
            <broadcast-warning> ;
            <delete-static-route destination  {
                routing-instance routing-instance;
            }>
            <(release-mastership | release-mastership-force);>
        }
    }
}
```

Hierarchy Level

```
[edit policy-options]
```

Description

Specify the actions to be taken for redundancy events. These include acquiring or releasing primary role and adding or deleting static routes.

Options

acquire-mastership	Switch from standby to primary role.
add-static-route ***destination***	(Optional) Use the specified destination IP address and prefix for an added signal route.
broadcast-warning	(Optional) Switch status from Standby to Standby (Warned).
delete-static-route ***destination***	(Optional) Use the specified destination IP address and prefix for a deleted signal route.
event-list	List of names of one or more monitored events that trigger the actions specified in this policy.
next-hop	Interface name for the next hop for an added signal route.

policy-name	Name of the redundancy policy.
receive	Use the added signal route as a receive route.
release-mastership	(Optional) Switch from primary to standby role.
release-mastership-force	(Optional) Force switch from primary to standby role.
routing-instance *routing-instance*	(Optional) Name of the vrf used for the added signal route.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring Inter-Chassis Services Redundancy for Next Gen Services

Configuring the Service Redundancy Daemon

redundancy (Power Management)

IN THIS SECTION

- Syntax | **14202**
- Hierarchy Level | **14202**
- Description | **14202**
- Default | **14202**
- Required Privilege Level | **14202**

Release Information | 14202

Syntax

```
redundancy {
                          n-plus-n (Power Management);
}
```

Hierarchy Level

```
[edit chassis (EX Series)                psu]
```

Description

Configure *N*+*N* power supply redundancy for power management on an EX6200 or EX8200 switch.

The remaining statement is explained separately. See CLI Explorer.

Default

N+1 power supply redundancy is configured by default.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Power Supply Redundancy (CLI Procedure)

redundancy (satellite-management)

IN THIS SECTION

- Syntax | 14203
- Hierarchy Level | 14203
- Description | 14203
- Default | 14203
- Required Privilege Level | 14203
- Release Information | 14204

Syntax

```
redundancy {
    n-plus-n;
}
```

Hierarchy Level

```
[edit chassis satellite-management psu]
```

Description

Configure *N*+*N* power supply redundancy for the satellite devices in a Junos Fusion.

The remaining statement is explained separately. See CLI Explorer.

Default

N+1 power supply redundancy is configured on each satellite device by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Understanding Power over Ethernet in a Junos Fusion

redundancy-set

IN THIS SECTION

- Syntax | 14204
- Hierarchy Level | 14205
- Description | 14205
- Options | 14205
- Required Privilege Level | 14205
- Release Information | 14205

Syntax

```
redundancy-set redundancy-set {
    healthcheck-timer-interval healthcheck-timer-interval;
    hold-time hold-time;
    keepalive keepalive;
    redundancy-group redundancy-group;
    redundancy-policy [redundancy-policy-list]
}
```

Hierarchy Level

```
[edit services]
```

Description

Specify the characteristics of a redundancy set.

Options

healthcheck-timer-interval *healthcheck-timer-interval*

Frequency of health check probes in seconds.

- **Range:** 0 through 3600 seconds

hold-time

Maximum wait time for a health check response. When this time expires, the peer is considered down.

- **Range:** 0 through 3600 seconds

keepalive

Frequency of srd hello messages in seconds.

- **Range:** 1 through 60 seconds

redundancy-group

Redundancy group identifier. This must match a redundancy group ID in the ICCP configuration.

- **Range:** 1 through 100

redundancy-policy-list

Names of one or more redundancy policies applied to the redundancy set.

redundancy-set

Redundancy set identifier.

- **Range:** 1 through 100

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring Inter-Chassis Services Redundancy for Next Gen Services

Configuring the Service Redundancy Daemon

redundancy-set-id (Service Set)

IN THIS SECTION

- Syntax | 14206
- Hierarchy Level | 14206
- Description | 14206
- Options | 14207
- Required Privilege Level | 14207
- Release Information | 14207

Syntax

```
redundancy-set-id redundancy-set;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the identifier of the redundancy set to use in the stateful synchronization of services for a service set.

Options

redundancy-set Identifier for the redundancy set. The identifier can be a number from 1-100.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring Inter-Chassis Services Redundancy for Next Gen Services

Configuring the Service Redundancy Daemon

redundancy-mode

IN THIS SECTION

- Syntax | 14208
- Hierarchy Level | 14208
- Description | 14208
- Options | 14208
- Required Privilege Level | 14208
- Release Information | 14208

Syntax

```
redundancy-mode(active-active | active-backup);
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Redundancy mode for the cluster (Reboot all nodes simultaneously after change)

Options

active-active	Primary and secondary nodes in active mode
active-backup	Primary in active, secondary in backup mode

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

show chassis cluster information | 17364

cluster (Chassis) | 4369

redundant-ether-options

IN THIS SECTION

- Syntax | 14209
- Hierarchy Level | 14209
- Description | 14209
- Options | 14210
- Required Privilege Level | 14210
- Release Information | 14211

Syntax

```
redundant-ether-options {
     (flow-control | no-flow-control);
    lacp {
        (active | passive);
        periodic (fast | slow);
    }
    link-speed speed;
    (loopback | no-loopback);
    minimum-links number;
    redundancy-group number;
    source-address-filter mac-address;
    (source-filtering | no-source-filtering);
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure Ethernet redundancy options for a chassis cluster.

In a chassis cluster setup, a redundant Ethernet interface is a pseudointerface that includes at minimum one physical interface from each node of the cluster.

A reth is a special type of interface that has the characteristics of aggregated Ethernet interface.

Options

flow-control

Enable flow control.

link-speed

Link speed of individual interface that joins the reth interface.

- Values:
  - `100m`—Links are 100 Mbps
  - `10g`—Links are 10 Gbps
  - `10m`—Links are 10 Mbps
  - `1g`—Links are 1Gbps

loopback

Enable loopback.

minimum-links

Minimum number of active links.

- **Default:** `1`
- **Range:** `1-8`

no-flow-control

Do not enable flow control.

no-loopback

Do not enable loopback.

no-source-filtering

Do not enable source address filtering.

redundancy-group

Redundancy group of this interface.

- **Range:** `1-128`

source-filtering

Enable source address filtering.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Example: Enabling Eight-Queue Class of Service on Redundant Ethernet Interfaces on SRX Series Devices in a Chassis Cluster

Example: Configuring Chassis Cluster Redundant Ethernet Interfaces

redundant-logical-tunnel

IN THIS SECTION

- Syntax | 14211
- Hierarchy Level | 14211
- Description | 14212
- Options | 14212
- Required Privilege Level | 14212
- Release Information | 14212

Syntax

```
redundant-logical-tunnel {
    device-count count;
}
```

Hierarchy Level

```
[edit chassis redundancy-group interface-type]
```

Description

Configure redundant logical tunnels on MX Series 5G Universal Routing Platforms.

Options

count Specify the number of the redundant logical tunnels. For Junos OS Release 13.3R1, 13.3R2, and 14.1R1 the valid range is from 1 through 16. For Junos OS Release 13.3R3 and 14.1R2 and later releases the valid range is from 1 through 255.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Support for up to 255 redundant logical tunnels added to Junos OS Release 13.3R3 and 14.1R2.

RELATED DOCUMENTATION

- Example: Configuring Redundant Logical Tunnels
- Configuring Redundant Logical Tunnels
- Redundant Logical Tunnels Overview
- redundancy-group (Redundant Tunnel) | **14180**
- redundancy-group (Chassis - MX Series) | **14176**

redundant-parent (Interfaces)

IN THIS SECTION

- Syntax | **14213**

- Hierarchy Level | 14213
- Description | 14213
- Options | 14213
- Required Privilege Level | 14213
- Release Information | 14213

Syntax

```
redundant-parent redundant-ethernet-interface-name;
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options]
[edit interfaces interface-name fastether-options]
```

Description

Specify redundant Ethernet interfaces (reth) and assign local (child) interfaces to the reth interfaces. A redundant Ethernet interface contains a pair of Fast Ethernet interfaces or a pair of Gigabit Ethernet interfaces that are referred to as child interfaces of the redundant Ethernet interface.

Options

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Example: Configuring Chassis Cluster Redundant Ethernet Interfaces

redundant-parent (Interfaces Fast Ethernet)

IN THIS SECTION

- Syntax | 14214
- Hierarchy Level | 14214
- Description | 14214
- Options | 14214
- Required Privilege Level | 14215
- Release Information | 14215

Syntax

```
redundant-parent interface-name;
```

Hierarchy Level

```
[edit interfaces interface-name fastether-options]
```

Description

Configure Fast Ethernet-specific interface properties for Ethernet redundancy in a chassis cluster.

Options

interface —Parent redundant interface of the Fast Ethernet interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

redundant-parent (Interfaces Gigabit Ethernet)

IN THIS SECTION

- Syntax | 14215
- Hierarchy Level | 14215
- Description | 14216
- Options | 14216
- Required Privilege Level | 14216
- Release Information | 14216

Syntax

```
redundant-parent interface-name;
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options]
```

Description

Configure Gigabit Ethernet-specific interface properties for Ethernet redundancy in a chassis cluster.

Options

interface —Parent redundant interface of the Gigabit Ethernet interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 9.0 of Junos OS.

redundant-power-system

IN THIS SECTION

- Syntax | 14216
- Hierarchy Level | 14217
- Description | 14217
- Required Privilege Level | 14217
- Release Information | 14217

Syntax

EX2200 switch:

```
redundant-power-system {
    priority (0|1|2|3|4|5|6)
```

```
    }
}
```

EX3300 switch:

```
redundant-power-system {
    member vc-member-number  {
        priority (0|1|2|3|4|5|6)
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure Redundant Power System (RPS) member to ensure higher- switches always receive power backup.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Determining and Setting Priority for Switches Connected to an EX Series RPS

redundant-pseudo-interface-options

IN THIS SECTION

- Syntax | 14218
- Hierarchy Level | 14218
- Description | 14218
- Options | 14219
- Required Privilege Level | 14219
- Release Information | 14219

Syntax

```
redundant-pseudo-interface-options {
    redundancy-group redundancy-group-number;
}
```

Hierarchy Level

```
[edit interfaces lo0]
```

Description

Configure the loopback pseudointerface in a redundancy group.

Redundancy groups are used to bundle interfaces into a group for failover purpose in a chassis cluster setup. You can configure a loopback interface as an alternative physical interface to reach the peer gateway. Loopback interfaces can be configured on any redundancy group.

For example: An Internet Key Exchange (IKE) gateway operating in chassis cluster, needs an external interface to communicate with a peer device. When an external interface (a reth interface or a standalone interface) is used for communication; the interface might go down when the physical interfaces are down. Instead, use loopback interfaces as an alternative to physical interfaces.

Options

redundancy-group-number Configure the redundancy group number.

- **Range:** 0 through 255

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Understanding the Loopback Interface for a High Availability VPN

redundant-sources

IN THIS SECTION

- Syntax | 14220
- Hierarchy Level | 14220
- Description | 14220
- Options | 14220
- Required Privilege Level | 14220
- Release Information | 14220

Syntax

```
redundant-sources [ addresses ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast flow-map flow-map-name],
[edit logical-systems logical-system-name routing-options multicast flow-map flow-map-name],
[edit routing-instances routing-instance-name routing-options multicast flow-map flow-map-name],
[edit routing-options multicast flow-map flow-map-name]
```

Description

Configure a list of redundant sources for multicast flows defined by a flow map.

Options

addresses—List of IPv4 or IPv6 addresses for use as redundant (backup) sources for multicast flows defined by a flow map.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

> *Example: Configuring a Multicast Flow Map*

redundant-trunk-group

IN THIS SECTION

- Syntax | 14221
- Hierarchy Level | 14221
- Description | 14222
- Required Privilege Level | 14222
- Release Information | 14222

Syntax

```
redundant-trunk-group {
    group name {
        interface interface-name <primary>;
        interface interface-name;
        preempt-cutover-timer seconds;
    }
}
```

Hierarchy Level

- For platforms with ELS:

```
[edit switch-options]
```

- For platforms without ELS:

```
[edit ethernet-switching-options]
```

Description

Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over without waiting for normal spanning-tree protocol convergence.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Hierarchy level `[edit switch-options]` introduced in Junos OS Release 13.2X50-D10 (ELS). (See *Using the Enhanced Layer 2 Software CLI* for information about ELS.)

RELATED DOCUMENTATION

Example: Configuring Redundant Trunk Links for Faster Recovery on EX Series Switches

Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support

Understanding Redundant Trunk Links (Legacy RTG Configuration)

redundant-virtual-tunnel

IN THIS SECTION

- Syntax | 14223
- Hierarchy Level | 14223
- Description | 14223
- Options | 14223
- Required Privilege Level | 14223
- Release Information | 14223

Syntax

```
redundant-virtual-tunnel {
    device-count count;
}
```

Hierarchy Level

```
[edit chassis redundancy-group interface-type]
```

Description

Configure redundant virtual tunnels on MX Series 5G Universal Routing Platforms.

Options

`device-count` *`count`*—Specify the number of redundant virtual tunnels.

- **Range:** 1 through 16

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.2.

RELATED DOCUMENTATION

Redundant Virtual Tunnels Providing Resiliency in Delivering Multicast Traffic Overview

Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic

Example: Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic

redundancy-group (Redundant Tunnel) | **14180**

redundancy-group (Chassis - MX Series) | **14176**

reference-bandwidth (Protocols IS-IS)

IN THIS SECTION

- Syntax | 14224
- Hierarchy Level | 14224
- Description | 14224
- Options | 14225
- Required Privilege Level | 14225
- Release Information | 14225

Syntax

```
reference-bandwidth reference-bandwidth;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Optimize routing based on bandwidth by setting the reference bandwidth used in calculating the default interface cost.

All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. When there are several equal-cost routes to a destination, traffic is distributed equally among them. Tweak the reference bandwidth to route traffic across the fastest interface.

The cost of a route is described by a single dimensionless metric that is determined using the following formula:

```
cost = reference-bandwidth/bandwidth
```

For example, if you set the reference bandwidth to 1 Gbps (that is, *reference-bandwidth* is set to 1,000,000,000), a 100-Mbps interface has a routing metric of 10.

All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. By default, IS-IS has a metric of 10. Modify the configured reference bandwidth to increase the cost of an interface in order to avoid a slow route or you can modify the reference bandwidth value to decrease the cost of a preferred interface. A route that has a lower cost is selected over a high cost route. Note that fractional metric values are not allowed. Therefore, if you configure a reference bandwidth that is higher than the actual bandwidth, which results in a fractional metric value then IS-IS takes the metric value as 1 because fractional metric values are not allowed.

Options

reference-bandwidth—Reference bandwidth value in bits per second.

- **Range:** 9600 through 4,000,000,000,000 bps
- **Default:** None

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Understanding IS-IS Configuration*
>
> *Example: Configuring IS-IS*
>
> https://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

reference-bandwidth (Protocols OSPF)

IN THIS SECTION

- Syntax | 14226
- Hierarchy Level | 14226
- Description | 14227
- Options | 14227
- Required Privilege Level | 14227
- Release Information | 14227

Syntax

```
reference-bandwidth reference-bandwidth;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula:

```
cost = ref-bandwidth/bandwidth
```

Options

reference-bandwidth—Reference bandwidth, in bits per second.

- **Range:** 9600 through 4,000,000,000,000 bits
- **Default:** 100 Mbps (100,000,000 bits)

> **NOTE**: The default behavior is to use the reference-bandwidth value to calculate the cost of OSPF interfaces. You can override this behavior for any OSPF interface by configuring a specific cost with the `metric` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Example: Controlling the Cost of Individual OSPF Network Segments

reflect-etype

IN THIS SECTION

- Syntax | 14228
- Junos OS Hierarchy Level | 14228
- Junos OS Evolved Hierarchy Level | 14228
- Description | 14228
- Options | 14229
- Required Privilege Level | 14229
- Release Information | 14229

Syntax

```
reflect-etype ethertype-value;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the EtherType to be used for reflection of the test frames. EtherType is a two-octet field in an Ethernet frame that defines the protocol in the frame payload. This statement is valid only if you configure the test mode to be a reflector. If you do not configure this statement, all EtherTypes are reflected.

Options

ethertype-value

Identifier for the EtherType. The EtherType value appears in the Ethernet type field of the packet. It specifies the protocol being transported in the Ethernet frame. For instance, the EtherType for IPv4 is 0x0800. So, if you specify the value as 2048, IPv4 packets are reflected.

- **Range:** 1 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

Supported RFC 2544-Based Benchmarking Statements on MX Series Routers

Configuring an RFC 2544-Based Benchmarking Test

reflect-mode

IN THIS SECTION

- Junos OS Syntax | **14230**
- Junos OS Evolved Syntax | **14230**
- Junos OS Hierarchy Level | **14230**

- Junos OS Evolved Hierarchy Level | 14230
- Description | 14230
- Options | 14231
- Required Privilege Level | 14231
- Release Information | 14231

Junos OS Syntax

```
reflect-mode (mac-rewrite | mac-swap | no-mac-swap );
```

Junos OS Evolved Syntax

```
reflect-mode (mac-swap | no-mac-swap);
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the reflection mode for the benchmarking test.

Options

mac-rewrite (Junos OS ACX Series routers only) Enable rewriting of the MAC address on the reflected frames. You configure which MAC addresses the software should use with the `source-mac-address` and `destination-mac-address` configuration statements.

mac-swap Swap the source and destination MAC addresses in the test frame. This is the default behavior.

> **NOTE**: In `bridge` families, when the service type is `ELAN`, MAC addresses are swapped by default on the reflected frames. When the service type is `ELINE` , MAC addresses are not swapped by default.

no-mac-swap Do not swap the source and destination MAC addresses in the test frame. The frame is returned to the originator without any modification to the MAC addresses.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

Understanding RFC2544-Based Benchmarking Tests for E-LAN and E-Line Services on MX Series Routers

rfc2544-benchmarking | **15759**

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

Configuring an RFC 2544-Based Benchmarking Test

reflective-relay

IN THIS SECTION

- Syntax | 14232
- Hierarchy Level | 14232
- Description | 14232
- Default | 14232
- Required Privilege Level | 14232
- Release Information | 14233

Syntax

```
reflective-relay;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family ethernet-switching]
```

Description

Configure a switch interface to return packets back to a device on the same interface that was used to deliver the packets.

Default

Switch interfaces are not configured for reflective relay.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Example: Configuring Reflective Relay for Use with VEPA Technology on QFX Switches
- Configuring Reflective Relay on Switches

reflexive | revert | reverse

IN THIS SECTION

- Syntax | 14233
- Hierarchy Level | 14233
- Description | 14234
- Required Privilege Level | 14234
- Release Information | 14234

Syntax

```
reflexive; | revert; | reverse {
    application-profile profile-name;
    dscp (alias | bits);
    forwarding-class class-name;
    syslog;
}
```

Hierarchy Level

```
[edit services cos rule rule-name term term-name then]
```

Description

`reflexive`—Applies the CoS rule actions to flows in the reverse direction as well as to flows in the matching direction.

`revert`—Stores the DSCP and forwarding class of a packet that is received in the match direction of the rule and then applies that DSCP and forwarding class to packets that are received in the reverse direction of the same session.

`reverse`—Allows you to define CoS behavior for flows in the reverse direction.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

`revert` option introduced in Junos OS Release 16.1R5 and 17.4R1.

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

Configuring CoS Rules

refresh (Commit Scripts)

IN THIS SECTION

- Syntax | 14235
- Hierarchy Level | 14235
- Description | 14235

Required Privilege Level | 14236
Release Information | 14236

Syntax

```
refresh;
```

Hierarchy Level

```
[edit system scripts commit],
[edit system scripts commit file filename]
```

Description

Overwrite the local copy of all enabled commit scripts or a single enabled commit script with the copy located at the source URL, as specified in the `source` statement at the same hierarchy level. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

The update operation occurs as soon as you issue the `set refresh` configuration mode command. Issuing the `set refresh` command does not add the `refresh` statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

Before you refresh a local script from an HTTPS server, you must first configure the `cert-file` statement for that script in supported releases. Specify the absolute path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate during the refresh operation.

NOTE: On the QFabric system, commit scripts are stored in the **/pbdata/mgd_shared/*partition-ip*/var/db/scripts/commit/** directory on the Director device.

As of Junos OS Release 18.1R1, you can specify which routing instance the update is done through. To specify the routing instance to use for updating commit scripts, configure the routing instance in two places in the CLI:

```
user@host# set system routing-instances routing-instance-name description description
user@host# set system scripts commit file filename routing-instance routing-instance-name
```

If you enable the non-default management instance and use `mgmt_junos` for *routing-instance-name*, you can configure scripts to update using the dedicated management instance mgmt_junos.

Required Privilege Level

maintenance—To execute this statement in configuration mode.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring and Using a Master Source Location for a Script

Example: Configuring and Refreshing from the Master Source for a Script

refresh-from (Commit Scripts) | **14238**

source (Commit Scripts) | **28194**

routing-instance (System Scripts) | **15963**

refresh (Event Scripts)

IN THIS SECTION

- Syntax | **14237**
- Hierarchy Level | **14237**
- Description | **14237**
- Required Privilege Level | **14238**

Release Information | 14238

Syntax

```
refresh;
```

Hierarchy Level

```
[edit event-options event-script],
[edit event-options event-script file filename]
```

Description

Overwrite the local copy of all enabled event scripts or a single enabled event script with the copy located at the source URL, as specified in the `source` statement at the same hierarchy level. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

The update operation occurs as soon as you issue the `set refresh` configuration mode command. Issuing the `set refresh` command does not add the `refresh` statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

Before you refresh a local script from an HTTPS server, you must first configure the `cert-file` statement for that script in supported releases. Specify the absolute path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate during the refresh operation.

As of Junos OS Release 18.1R1, you can specify which routing instance the update is done through. To specify the routing instance to use for updating event scripts, configure the routing instance in two places in the CLI:

```
user@host# set system routing-instances routing-instance-name description description
user@host# set event-options event-script event file filename routing-instance-
name
```

If you enable the non-default management instance and use `mgmt_junos` for *routing-instance-name*, you can configure scripts to update using the dedicated management instance mgmt_junos.

Required Privilege Level

maintenance—To execute this statement in configuration mode.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configure and Use a Master Source Location for a Script

Example: Configure and Refresh from the Master Source for a Script

refresh-from (Event Scripts) | **14240**

source (Event Scripts) | **28206**

routing-instance (Event Scripts) | **15934**

refresh-from (Commit Scripts)

IN THIS SECTION

- Syntax | **14239**
- Hierarchy Level | **14239**
- Description | **14239**
- Options | **14240**
- Required Privilege Level | **14240**
- Release Information | **14240**

Syntax

```
refresh-from url;
```

Hierarchy Level

```
[edit system scripts commit],
[edit system scripts commit file filename]
```

Description

Overwrite the local copy of all enabled commit scripts or a single enabled commit script with the copy located at the specified URL. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

The update operation occurs as soon as you issue the `set refresh-from` *`url`* configuration mode command. Issuing the `set refresh-from` command does not add the `refresh-from` statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

Before you refresh a local script from an HTTPS server, you must first configure the `cert-file` statement for that script in supported releases. Specify the absolute path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate during the refresh operation.

NOTE: This statement is not supported on the QFabric system.

As of Junos OS Release 18.1R1, you can specify which routing instance the update is done through. To specify the routing instance to use for updating op scripts, configure the routing instance in two places in the CLI:

```
user@host# set system routing-instances routing-instance-name description description
user@host# set system scripts op file filename routing-instance routing-instance-name
```

If you enable the non-default management instance and use `mgmt_junos` for *`routing-instance-name`*, you can configure scripts to update using the dedicated management instance mgmt_junos.

Options

url The source specified as an HTTPS URL, HTTP URL, FTP URL, or secure copy (scp)-style remote file specification.

Required Privilege Level

maintenance—To execute this statement in configuration mode.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Using an Alternate Source Location for a Script

refresh (Commit Scripts) | **14234**

source (Commit Scripts) | **28194**

routing-instance (System Scripts) | **15963**

refresh-from (Event Scripts)

IN THIS SECTION

- Syntax | **14241**
- Hierarchy Level | **14241**
- Description | **14241**
- Options | **14241**
- Required Privilege Level | **14242**
- Release Information | **14242**

Syntax

```
refresh-from url;
```

Hierarchy Level

```
[edit event-options event-script],
[edit event-options event-script file filename]
```

Description

Overwrite the local copy of all enabled event scripts or a single enabled event script with the copy located at the specified URL. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

The update operation occurs as soon as you issue the `set refresh-from` *`url`* configuration mode command. Issuing the `set refresh-from` command does not add the `refresh-from` statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

Before you refresh a local script from an HTTPS server, you must first configure the `cert-file` statement for that script in supported releases. Specify the absolute path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate during the refresh operation.

As of Junos OS Release 18.1R1, you can specify which routing instance the update is done through. To specify the routing instance to use for updating event scripts, configure the routing instance in two places in the CLI:

```
user@host# set system routing-instances routing-instance-name description description
user@host# set event-options event-script event file filename routing-instance routing-instance-
name
```

If you enable the non-default management instance and use `mgmt_junos` for *`routing-instance-name`*, you can configure scripts to update using the dedicated management instance mgmt_junos.

Options

url Source specified as an HTTPS URL, HTTP URL, FTP URL, or secure copy (scp)-style remote file specification.

Required Privilege Level

maintenance—To execute this statement in configuration mode.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Use an Alternate Source Location for a Script

refresh (Event Scripts) | **14236**

source (Event Scripts) | **28206**

routing-instance (Event Scripts) | **15934**

refresh-from (JET)

IN THIS SECTION

- Syntax | **14242**
- Hierarchy Level | **14243**
- Description | **14243**
- Options | **14243**
- Required Privilege Level | **14244**
- Release Information | **14244**

Syntax

```
refresh-from url;
```

Hierarchy Level

```
[edit system extensions extension-service application file filename]
```

Description

Overwrite the local copy of all enabled commit scripts or a single enabled commit script with the copy located at the specified URL. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

The update operation occurs as soon as you issue the `set refresh-from` *`url`* configuration mode command. Issuing the `set refresh-from` command does not add the `refresh-from` statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

Before you refresh a local script from an HTTPS server, you must first configure the `cert-file` statement for that script in supported releases. Specify the absolute path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate during the refresh operation.

NOTE: This statement is not supported on the QFabric system.

As of Junos OS Release 18.1R1, you can specify which routing instance the update is done through. To specify the routing instance to use for updating op scripts, configure the routing instance in two places in the CLI:

```
user@host# set system routing-instances routing-instance-name description description
user@host# set system extensions extension-service application file filename routing-instance
routing-instance-name
```

If you enable the non-default managment instance and use `mgmt_junos` for *`routing-instance-name`*, you can configure scripts to update using the dedicated management instance mgmt_junos.

Options

url The source specified as a HTTPS URL, HTTP URL, FTP URL, or secure copy (scp)-style remote file specification.

Required Privilege Level

maintenance—To execute this statement in configuration mode.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

- Using an Alternate Source Location for a Script
- refresh (JET) | 14250
- source (JET Scripts) | 28245
- routing-instance (JET Scripts) | 15935

refresh-from (Op Scripts)

IN THIS SECTION

- Syntax | 14244
- Hierarchy Level | 14245
- Description | 14245
- Options | 14245
- Required Privilege Level | 14246
- Release Information | 14246

Syntax

```
refresh-from url;
```

Hierarchy Level

```
[edit system scripts op],
[edit system scripts op file filename]
```

Description

Overwrite the local copy of all enabled op scripts or a single enabled op script with the copy located at the specified URL. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

The update operation occurs as soon as you issue the `set refresh-from` *`url`* configuration mode command. Issuing the `set refresh-from` command does not add the `refresh-from` statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

Before you refresh a local script from an HTTPS server, you must first configure the `cert-file` statement for that script in supported releases. Specify the absolute path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate during the refresh operation.

NOTE: This statement is not supported on the QFabric system.

As of Junos OS Release 18.1R1, you can specify which routing instance the update is done through. To specify the routing instance to use for updating op scripts, configure the routing instance in two places in the CLI:

```
user@host# set system routing-instances routing-instance-name description description
user@host# set system scripts op file filename routing-instance routing-instance-name
```

If you enable the non-default management instance and use `mgmt_junos` for *`routing-instance-name`*, you can configure scripts to update using the dedicated management instance mgmt_junos.

Options

url Source specified as an HTTPS URL, HTTP URL, FTP URL, or secure copy (scp)-style remote file specification.

Required Privilege Level

maintenance—To execute this statement in configuration mode.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Using an Alternate Source Location for a Script

refresh (Op Scripts) | **14251**

source (Op Scripts) | **28255**

routing-instance (System Scripts) | **15963**

refresh-from (SNMP Scripts)

IN THIS SECTION

- Syntax | **14246**
- Hierarchy Level | **14247**
- Description | **14247**
- Options | **14247**
- Required Privilege Level | **14248**
- Release Information | **14248**

Syntax

```
refresh-from url;
```

Hierarchy Level

```
[edit system scripts snmp],
[edit system scripts snmp file filename]
```

Description

Overwrite the local copy of all enabled SNMP scripts or a single enabled SNMP script with the copy located at the specified URL. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard drive.

The update operation occurs as soon as you issue the `set refresh-from` *`url`* configuration mode command. Issuing the `set refresh-from` command does not add the `refresh-from` statement to the configuration. Thus, the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

Before you refresh a local script from an HTTPS server, you must first configure the `cert-file` statement for that script in supported releases. Specify the absolute path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate during the refresh operation.

NOTE: This statement is not supported on the QFabric system.

As of Junos OS Release 18.1R1, you can specify which routing instance the update is done through. To specify the routing instance to use for updating snmp scripts, configure the routing instance in two places in the CLI:

```
user@host# set system routing-instances routing-instance-name description description
user@host# set system scripts snmp file filename routing-instance routing-instance-name
```

If you enable the non-default management instance and use `mgmt_junos` for *`routing-instance-name`*, you can configure scripts to update using the dedicated management instance mgmt_junos.

Options

url The source specified as an HTTPS URL, HTTP URL, FTP URL, or secure copy (scp)-style remote file specification.

Required Privilege Level

maintenance—To execute this statement in configuration mode.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Use an Alternate Source Location for a Script

snmp (SNMP Scripts) | **28068**

refresh (SNMP Scripts) | **14255**

routing-instance (System Scripts) | **15963**

refresh-interval

IN THIS SECTION

- Syntax | **14248**
- Hierarchy Level | **14249**
- Description | **14249**
- Options | **14249**
- Required Privilege Level | **14249**
- Release Information | **14249**

Syntax

```
refresh-interval number-of-hours;
```

Hierarchy Level

```
[edit security pki ca-profile ca-profile-name revocation-check crl]
```

Description

(Adaptive services interfaces only) Specify the amount of time between certificate revocation list (CRL) updates.

Options

number-of-hours—Time interval, in hours, between CRL updates.

- **Range:** 0 through 8784
- **Default:** 24

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring the Certificate Revocation List

crl (Adaptive Services Interface) | **4718**

refresh (JET)

IN THIS SECTION

- Syntax | 14250
- Hierarchy Level | 14250
- Description | 14250
- Required Privilege Level | 14251
- Release Information | 14251

Syntax

```
refresh;
```

Hierarchy Level

```
[edit system extensions extension-service application file filename]
```

Description

Overwrite the local copy of all enabled commit scripts or a single enabled commit script with the copy located at the source URL, as specified in the `source` statement at the same hierarchy level. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

The update operation occurs as soon as you issue the `set refresh` configuration mode command. Issuing the `set refresh` command does not add the `refresh` statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

Before you refresh a local script from an HTTPS server, you must first configure the `cert-file` statement for that script in supported releases. Specify the absolute path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate during the refresh operation.

> **NOTE**: On the QFabric system, commit scripts are stored in the **/pbdata/mgd_shared/*partition-ip*/var/db/scripts/commit/** directory on the Director device.

As of Junos OS Release 18.1R1, you can specify which routing instance the update is done through. To specify the routing instance to use for updating commit scripts, configure the routing instance in two places in the CLI:

```
user@host# set system routing-instances routing-instance-name description description
user@host# set system extensions extension-service application file filename routing-instance
routing-instance-name
```

If you enable the non-default managment instance and use `mgmt_junos` for *routing-instance-name*, you can configure scripts to update using the dedicated management instance mgmt_junos.

Required Privilege Level

maintenance—To execute this statement in configuration mode.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

- refresh-from (JET) | **14242**
- Configuring and Using a Master Source Location for a Script
- Example: Configuring and Refreshing from the Master Source for a Script
- Management Interface in a Non-Default Instance

refresh (Op Scripts)

IN THIS SECTION

- Syntax | **14252**

- Hierarchy Level | 14252
- Description | 14252
- Required Privilege Level | 14253
- Release Information | 14253

Syntax

```
refresh;
```

Hierarchy Level

```
[edit system scripts op],
[edit system scripts op file filename]
```

Description

Overwrite the local copy of all enabled op scripts or a single enabled op script with the copy located at the source URL, as specified in the `source` statement at the same hierarchy level. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

The update operation occurs as soon as you issue the `set refresh` configuration mode command. Issuing the `set refresh` command does not add the `refresh` statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

Before you refresh a local script from an HTTPS server, you must first configure the `cert-file` statement for that script in supported releases. Specify the absolute path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate during the refresh operation.

NOTE: On the QFabric system, op scripts are stored in the **/pbdata/mgd_shared/*partition-ip*/var/db/scripts/op/** directory on the Director device.

As of Junos OS Release 18.1R1, you can specify which routing instance the update is done through. To specify the routing instance to use for updating op scripts, configure the routing instance in two places in the CLI:

```
user@host# set system routing-instances routing-instance-name description description
user@host# set system scripts op file filename routing-instance routing-instance-name
```

If you enable the non-default management instance and use `mgmt_junos` for *routing-instance-name*, you can configure scripts to update using the dedicated management instance mgmt_junos.

Required Privilege Level

maintenance—To execute this statement in configuration mode.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring and Using a Master Source Location for a Script

Example: Configuring and Refreshing from the Master Source for a Script

refresh-from (Op Scripts) | **14244**

source (Op Scripts) | **28255**

routing-instance (System Scripts) | **15963**

refresh-rate (Flow Monitoring Logs for NAT)

IN THIS SECTION

- Syntax | **14254**
- Hierarchy Level | **14254**
- Description | **14254**
- Options | **14254**

Required Privilege Level | 14254
Release Information | 14255

Syntax

```
refresh-rate packets packets seconds seconds;
```

Hierarchy Level

```
[edit services jflow-log template-profile template-profile-name]
```

Description

Specify the refresh rate for transmitting flow template records with version 9 and IPFIX templates for NAT events to the collector, in either packets or seconds.

Options

packets— Number of packets after which templates are sent to the collector.

- **Range:** 1 through 480,000
- **Default:** 4800

seconds—Number of seconds after which templates are sent to the collector

- **Range:** 10 through 600
- **Default:** 600

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

refresh (SNMP Scripts)

IN THIS SECTION

- Syntax | 14255
- Hierarchy Level | 14255
- Description | 14256
- Required Privilege Level | 14256
- Release Information | 14256

Syntax

```
refresh;
```

Hierarchy Level

```
[edit system scripts snmp],
[edit system scripts snmp file filename]
```

Description

Overwrite the local copy of all enabled SNMP scripts or a single enabled SNMP script with the copy located at the source URL, as specified in the `source` statement at the same hierarchy level. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard drive.

The update operation occurs as soon as you issue the `set refresh` configuration mode command. Issuing the `set refresh` command does not add the `refresh` statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

Before you refresh a local script from an HTTPS server, you must first configure the `cert-file` statement for that script in supported releases. Specify the absolute path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate during the refresh operation.

As of Junos OS Release 18.1R1, you can specify which routing instance the update is done through. To specify the routing instance to use for updating snmp scripts, configure the routing instance in two places in the CLI:

```
user@host# set system routing-instances routing-instance-name description description
user@host# set system scripts snmp file filename routing-instance routing-instance-name
```

If you enable the non-default management instance and use `mgmt_junos` for *`routing-instance-name`*, you can configure scripts to update using the dedicated management instance mgmt_junos.

Required Privilege Level

maintenance—To execute this statement in configuration mode.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring and Using a Master Source Location for a Script

Example: Configuring and Refreshing from the Master Source for a Script

snmp (SNMP Scripts) | **28068**

refresh-from (SNMP Scripts) | **14246**

routing-instance (System Scripts) | 15963

refresh-time

IN THIS SECTION

- Syntax | 14257
- Hierarchy Level | 14257
- Description | 14257
- Options | 14257
- Required Privilege Level | 14258
- Release Information | 14258

Syntax

```
refresh-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Set the refresh time.

Options

seconds—Refresh time.

- **Range:** 1 through 65,535
- **Default:** 1200 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Timers for RSVP Refresh Messages

regex-additive-logic

IN THIS SECTION

- Syntax | 14258
- Hierarchy Level | 14258
- Description | 14259
- Default | 14259
- Required Privilege Level | 14259
- Release Information | 14259

Syntax

```
regex-additive-logic;
```

Hierarchy Level

```
[edit system]
```

Description

Enable additive logic (that is, deny by default / allow some as specified) to be used in regular expressions in the `allow-configuration-regexps` and `deny-configuration-regexps` statements at the [`edit system login class` *`class-name`*] hierarchy level.

When you enable additive logic, the `allow-configuration-regexps` statement takes precedence over the `deny-configuration-regexps` statement. Thus, if you configure the `regex-additive-logic` statement and the same configuration hierarchy is referenced in both the `allow-configuration-regexps` statement and the `deny-configuration-regexps` statement, then the configuration hierarchy is allowed. If you do not configure the `regex-additive-logic`, then the same hierarchy would be denied in this case because the `deny-configuration-regexps` statement takes precedence by default.

Default

By default, this statement is disabled; the `deny-configuration-regexps` statement takes precedence over the `allow-configuration-regexps` statement.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Example: Configuring User Permissions with Access Privileges for Configuration Statements and Hierarchies

Example: Using Additive Logic With Regular Expressions to Specify Access Privileges

Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands, Configuration Statements, and Hierarchies

class (Defining Login Classes) | **3347**

user (Access) | **30547**

regex (Class Attribute)

IN THIS SECTION

- Syntax | 14260
- Hierarchy Level | 14260
- Description | 14260
- Options | 14260
- Required Privilege Level | 14261
- Release Information | 14261

Syntax

```
regex "value";
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name subscription-id use-class]
```

Description

Configure a regular expression to parse the Class attribute contents of the accounting request from the BNG, PGW, or GGSN.

Options

value Regular expression that parses the contents of the Class attribute.

For example, the following configuration generates " 000118191129ALICE:DRAV3:" out of " 000118191129#000118191129#ALICE:DRAV3:#7168#nflat#ADSL##":

```
[edit unified-edge gateways tdf TDF1 domains domain1 subscription-id ]
user@host# set use-class regex "[^#]*#\([^#]*\)\#\([^#]*\)"
```

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

regexp

IN THIS SECTION

- Syntax | 14262
- Hierarchy Level | 14262
- Description | 14262
- Required Privilege Level | 14262
- Release Information | 14262

Syntax

```
regexp regular-expression;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature]
```

Description

Specify a Perl Compatible Regular Expression (PCRE) expression.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

region

IN THIS SECTION

- Syntax | 14263
- Hierarchy Level | 14263
- Description | 14263
- Options | 14263
- Required Privilege Level | 14264
- Release Information | 14264

Syntax

```
region region-name{
    incoming;
    ingress-replication {
        create-new-ucast-tunnel;
        label-switched-path {
            llabel-switched-path-template (Multicast){
                (default-template | lsp-template-name);
            }
        }
    }
    ldp-p2mp;
    rsvp-te {
        label-switched-path-template (Multicast) {
            (default-template | lsp-template-name);
        }
        static-lsp static-lsp;
    }
}
```

Hierarchy Level

```
[edit protocols mvpn inter-region-template template template-name]
```

Description

Specify the regions based on the BGP peer groups. The regions support different tunnel types such as incoming, ingress replication, ldp-p2mp, or rsvp-te. The incoming tunnel indicates that the tunnel type across the area boundary router (ABR) is not changed. A template is used to define the tunnel type for each region.

Options

incoming Specify the tunnel type to be same across area boundary router (ABR).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview

Example: Configuring Segmented Inter-Area P2MP LSP

Configuring Segmented Inter-Area P2MP LSP

all-regions | **2157**

inter-region | **8789**

inter-region-segmented | **8791**

inter-region-template | **8793**

template(Protocols MVPN) | **29092**

register-limit

IN THIS SECTION

- Syntax | **14265**
- Hierarchy Level | **14265**
- Description | **14265**
- Options | **14265**
- Required Privilege Level | **14266**
- Release Information | **14266**

Syntax

```
register-limit {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
    log-interval seconds;
    maximum limit;
    threshold value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Configure a limit for the number of incoming (S,G) PIM registers.

> **NOTE**: The maximum limit settings that you configure with the `maximum` and the `family (inet | inet6) maximum` statements are mutually exclusive. For example, if you configure a global maximum PIM register message limit, you cannot configure a limit at the family level for IPv4 or IPv6. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.

Options

`family (inet | inet6)`—(Optional) Specify either IPv4 or IPv6 messages to be counted towards the configured register message limit.

- **Default:** Both IPv4 and IPv6 messages are counted towards the configured register message limit.

The remaining statements are described separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring PIM State Limits

clear pim join

clear pim register

register-probe-time

IN THIS SECTION

- Syntax | 14266
- Hierarchy Level | 14267
- Description | 14267
- Options | 14267
- Required Privilege Level | 14267
- Release Information | 14267

Syntax

```
register-probe-time register-probe-time;
```

Hierarchy Level

```
[edit protocols pim rp]
```

Description

Specify the amount of time before the register suppression time (RST) expires when a designated switch can send a NULL-Register to the rendezvous point (RP).

Options

register-probe-time

Amount of time before the RST expires.

- **Default: 5** seconds
- **Range: 5** to **60** seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

PIM Overview

Understanding PIM Sparse Mode

registration

IN THIS SECTION

- Syntax | 14268
- Hierarchy Level | 14268
- Description | 14268
- Default | 14269
- Options | 14269
- Required Privilege Level | 14269
- Release Information | 14269

Syntax

```
registration  (forbidden | normal | restricted);
```

Hierarchy Level

```
[edit protocols mvrp interface (all | interface-name)]
```

```
[edit routing-instances routing-instance-name protocols mvrp interface (all | interface-name)]
(for virtual switch instance type)
```

```
[edit logical-systems logical-system-name protocols mvrp interface (all | interface-name)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvrp
interface (all | interface-name)] (for virtual switch instance type)
```

Description

Specifies the Multiple VLAN Registration Protocol (MVRP) registration mode for the interface if MVRP is enabled.

Default

`normal`—The interface or interfaces accept MVRP messages and participate in MVRP.

Options

forbidden—The interface or interfaces do not register and do not participate in MVRP.

normal—The interface or interfaces accept MVRP messages and participate in MVRP.

restricted—The interface or interfaces ignore all MVRP JOIN messages received for VLANs that are not statically configured for MVRP on the interface.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring Multiple VLAN Registration Protocol (MVRP) on Switches

Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers

Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration

Verifying That MVRP Is Working Correctly

join-timer (MVRP) | **9434**

leaveall-timer (MVRP) | **9707**

leave-timer (MVRP) | **9705**

point-to-point (MVRP) | **12620**

reject (Security)

IN THIS SECTION

- Syntax | 14270
- Hierarchy Level | 14270
- Description | 14270
- Options | 14271
- Required Privilege Level | 14271
- Release Information | 14271

Syntax

```
reject {
    profile redirect-profile-name;
    ssl-proxy {
        profile-name ssl-proxy-profile-name;
    }
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then]
```

Description

Block the service at the firewall. The device drops the packet and sends a TCP reset (RST) segment to the source host for TCP traffic and an ICMP “destination unreachable, port unreachable” message (type 3, code 3) for UDP traffic. For types of traffic other than TCP and UDP, the device drops the packet without notifying the source host, which is also what occurs when the action is `deny`.

You can configure reject action with one of the following options for the dynamic-applications:

- profile - You can chose to provide a notification to the clients or redirect client request to an informative Web page when a policy blocks HTTP or HTTPS traffic with a deny or reject action. To apply a profile, you must define the redirect profile for the dynamic applications.
- ssl-proxy - You can apply a redirect SSL proxy profile when a policy blocks HTTPS traffic with a reject action. When you apply am SSL proxy profile, SSL proxy decrypts the traffic and application identification functionality identifies the application. Next, you can take action to redirect or drop the traffic as per the configuration.

Options

profile	Profile for redirect HTTP/HTTPS traffic.
SSL-proxy	SSL proxy profile for decrypt HTTPS traffic.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Starting in Junos OS Release 19.2, options *profile* and *ssl-proxy* are added.

RELATED DOCUMENTATION

Security Policies Overview

reject-unauthorized-ipv6cp

IN THIS SECTION

Syntax | 14272
Hierarchy Level | 14272

- Description | 14272
- Default | 14272
- Required Privilege Level | 14272
- Release Information | 14273

Syntax

```
reject-unauthorized-ipv6cp;
```

Hierarchy Level

```
[edit protocols ppp-service]
```

Description

Configure the router to reject any IPv6 Control Protocol (IPv6CP) negotiation messages on dynamic interfaces when no appropriate IPv6 address or prefix has been received from AAA. IPv6CP negotiation attempts are also rejected when only a Framed-IPv6-Prefix attribute is received but router advertisement is not enabled in the dynamic profile.

NOTE: IPv6CP negotiation messages are not rejected for static interfaces.

Default

IPv6CP negotiation is allowed regardless of the presence of IPv6 attributes received from AAA.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address

rejoin-timeout (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 14273
- Hierarchy Level | 14273
- Description | 14274
- Default | 14274
- Options | 14274
- Required Privilege Level | 14274
- Release Information | 14274

Syntax

```
rejoin-timeout rejoin-timeout;
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options member-failure-options drop-member-
traffic]
```

Description

Configure the time by when failed members (members in the `DISCARD` state) should rejoin the aggregated Multiservices (AMS) interface automatically. All members that do not rejoin by the configured time are moved to the `INACTIVE` state and the traffic meant for each of the members is dropped.

If multiple members fail around the same time, then they are held in the `DISCARD` state using a single timer. When the timer expires, all the failed members move to `INACTIVE` state at the same time.

Default

If you do not configure a value, the default value of 120 seconds is used.

Options

`rejoin-timeout`—Time, in seconds, by which a failed member must rejoin.

- **Default:** 120 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

drop-member-traffic (Aggregated Multiservices) | **5949**

relay-agent-interface-id (DHCP Local Server)

IN THIS SECTION

- Syntax | 14275
- Hierarchy Level | 14275
- Description | 14276
- Required Privilege Level | 14276
- Release Information | 14276

Syntax

```
relay-agent-interface-id;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  group group-
name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Interface-ID option (option 18) in the client PDU name is concatenated with the username during the subscriber authentication or DHCP client authentication process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients

relay-agent-interface-id (DHCPv6 Relay Agent)

IN THIS SECTION

- Syntax | 14276
- Hierarchy Level | 14277
- Description | 14277
- Required Privilege Level | 14277
- Release Information | 14277

Syntax

```
relay-agent-interface-id {
    include-irb-and-l2;
```

```
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 ...],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dual-stack-group dual-
stack-group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dual-stack-group dual-stack-group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dual-stack-group ],
```

Description

Insert the DHCPv6 Relay Agent Interface-ID option (option 18) in DHCPv6 packets destined for the DHCPv6 server.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the [edit ... dual-stack-group *dual-stack-group-name*] hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

dhcp-relay | 5620

Extended DHCP Relay Agent Overview

DHCPv6 Relay Agent Overview

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets

relay-agent-interface-id (DHCPv6 Relay Agent Username)

IN THIS SECTION

- Syntax | 14278
- Hierarchy Level | 14278
- Description | 14279
- Required Privilege Level | 14279
- Release Information | 14279

Syntax

```
relay-agent-interface-id;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Interface-ID option (option 18) in the client PDU name is concatenated with the username during the subscriber authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

DHCPv6 Relay Agent Overview

Creating Unique Usernames for DHCP Clients

relay-agent-option

IN THIS SECTION

- Syntax | 14280
- Hierarchy Level | 14280

- Description | 14280
- Required Privilege Level | 14280
- Release Information | 14280

Syntax

```
relay-agent-option;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit logical-systems routing-instances instance-name forwarding-options helpers bootp],
[edit routing-instances instance-name forwarding-options helpers bootp]
```

Description

Enable the DHCP relay agent information option which allows DHCP to forward information from clients on different VRF routing instances. The functionality is described in RFC 3046, *DHCP Relay Agent Information Option*. For the Junos OS implementation, the DHCP option number is 82, and the suboption ID is 1. The suboption length is the length required to contain an interface name in addition to the terminating null character. The overall option length is the suboption length plus 2 bytes (for the option header). The DHCP relay agent information option is only present on packets sent between the relay and the DHCP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

relay-agent-option-79

IN THIS SECTION

- Syntax | 14281
- Hierarchy Level | 14281
- Description | 14281
- Required Privilege Level | 14282
- Release Information | 14282

Syntax

```
relay-agent-option-79;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6]
```

Description

Configure the DHCPv6 client link layer option (option 79) to include the source MAC address in DHCPv6 requests from clients before the messages are forwarded to a DHCPv6 server. In a dual-stack network environment in which devices act as both DHCPv4 and DHCPv6 clients, the client MAC address can be used to associate DHCPv4 and DHCPv6 messages with the same client interface. The client MAC address also provides network operators with additional information for event debugging and logging related to the client at the relay agent and the server.

You can configure `relay-agent-option-79` to enable DHCPv6 option 79 for a layer 3 DHCPv6 relay agent (LDRA). When DHCPv6 option 79 is enabled, the relay agent reads the source MAC address of DHCPv6 Solicit and DHCPv6 Request messages that it receives from a client. The relay agent encapsulates the

Solicit and Request messages within a DHCPv6 Relay-Forward message, and inserts the client MAC address as option 79 in the Relay-Forward header before relaying the message to the server.

You can also configure DHCPv6 option 79 for a lightweight DHCPv6 relay agent (LDRA). An LDRA resides on the same IPv6 link as the DHCPv6 client and relay agent or server and acts as a layer 2 relay agent. To configure DHCPv6 option 79 for an LDRA, use the `option-79` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
21.2R1	Starting with Junos OS Release 21.2R1, QFX Series switches support the `relay-agent-option-79` CLI statement in a stateless DHCPv6 relay configuration. You can configure stateless DHCPv6 relay using the `forward-only` CLI statement at the [`edit forwarding-options dhcp-relay dhcpv6`] hierarchy level.

RELATED DOCUMENTATION

Inserting the DHCPv6 Client MAC Address Option (Option 79) In DHCPv6 Packets

relay-agent-remote-id (DHCP Local Server)

IN THIS SECTION

- Syntax | 14283
- Hierarchy Level | 14283

- Description | 14283
- Required Privilege Level | 14284
- Release Information | 14284

Syntax

```
relay-agent-remote-id;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  group group-
name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Remote-ID option (option 37) in the client PDU name is concatenated with the username during the subscriber authentication or DHCP client authentication process. In order to generate an ASCII version of the username, the router concatenates only the remote-id portion of option 37 to the username, and ignores the enterprise number.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

For MX Series routers only, `enterprise-id` and `remote-id` options introduced in Junos OS Release 12.3R3.

For MX Series routers only, the `enterprise-id` and `remote-id` options are obsoleted starting in Junos OS Releases 12.3R7, 13.2R4, 13.3R2, and 14.1R1.

RELATED DOCUMENTATION

> *Creating Unique Usernames for DHCP Clients*

relay-agent-remote-id (DHCPv6 Relay Agent)

IN THIS SECTION

- Syntax | 14284
- Hierarchy Level | 14285
- Description | 14285
- Required Privilege Level | 14285
- Release Information | 14285

Syntax

```
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-remote-id ;
    no-vlan-interface-name;
```

```
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 ... ],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 ... ],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 ... ]
```

Description

Specify that the DHCPv6 relay agent include Relay Agent Remote-ID (option 37) in DHCPv6 packets destined for a DHCPv6 server. Optionally specify that the option includes a prefix, the interface textual description, or both. In dual-stack environments, you can also specify that the DHCPv6 relay agent use the DHCPv4 option 82 value for DHCPv6 option 37.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
21.2R1	Starting with Junos OS Release 21.2R1, QFX Series switches support the `relay-agent-remote-id` CLI statement in a stateless DHCPv6 relay configuration. You can configure stateless DHCPv6 relay using the `forward-only` CLI statement at the [`edit forwarding-options dhcp-relay dhcpv6`] hierarchy level.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

DHCPv6 Relay Agent Overview

DHCPv6 Relay Agent Options

Configuring DHCPv6 Relay Agent Options

relay-agent-remote-id (DHCPv6 Relay Agent Username)

IN THIS SECTION

- Syntax | **14286**
- Hierarchy Level | **14287**
- Description | **14287**
- Required Privilege Level | **14287**
- Release Information | **14287**

Syntax

```
relay-agent-remote-id;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Remote-ID option (option 37) in the client PDU name is concatenated with the username during the subscriber authentication or client authentication process. In order to generate an ASCII version of the username, the router concatenates only the remote-id portion of option 37 to the username, and ignores the enterprise number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

For MX Series routers only, `enterprise-id` and `remote-id` options introduced in Junos OS Release 12.3R3.

For MX Series routers only, the `enterprise-id` and `remote-id` options are obsoleted starting in Junos OS Releases 12.3R7, 13.2R4, 13.3R2, and 14.1R1.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

DHCPv6 Relay Agent Overview

Creating Unique Usernames for DHCP Clients

relay-agent-subscriber-id (DHCP Local Server)

IN THIS SECTION

- Syntax | 14288
- Hierarchy Level | 14288
- Description | 14289
- Required Privilege Level | 14289
- Release Information | 14289

Syntax

```
relay-agent-subscriber-id;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  group group-
name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
```

```
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication username-include],
[edit system services dhcp-local-server dhcpv6  authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Subscriber-ID option (option 38) in the client PDU name is concatenated with the username during the subscriber authentication process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients

relay-agent-subscriber-id (DHCPv6 Relay Agent)

IN THIS SECTION

- Syntax | **14290**
- Hierarchy Level | **14290**
- Description | **14290**
- Required Privilege Level | **14290**
- Release Information | **14290**

Syntax

```
relay-agent-subscriber-id;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include]
```

Description

Specify that the DHCPv6 Relay Agent Subscriber-ID option (option 38) in the client PDU name is concatenated with the username during the subscriber authentication process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

DHCPv6 Relay Agent Overview

Creating Unique Usernames for DHCP Clients

relay (AMT Protocol)

IN THIS SECTION

- Syntax | 14291
- Hierarchy Level | 14292
- Description | 14292
- Required Privilege Level | 14292
- Release Information | 14292

Syntax

```
relay {
    accounting;
    family {
        inet {
            anycast-prefix ip-prefix/<prefix-length>;
            local-address ip-address;
        }
    }
    secret-key-timeout minutes;
     tunnel-devicesvalue ;
    tunnel-limit number;
     unicast-stream-limitnumber;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols amt],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols amt],
[edit protocols amt],
[edit routing-instances routing-instance-name protocols amt]
```

Description

Configure the protocol address family, secret key timeout, and tunnel limit for Automatic Multicast Tunneling (AMT) relay functions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

relay (IGMP)

IN THIS SECTION

- Syntax | 14293
- Hierarchy Level | 14293
- Description | 14293

Required Privilege Level | 14293
Release Information | 14294

Syntax

```
relay {
    defaults {
        (accounting | no-accounting);
        group-policy [ policy-names ];
        query-interval seconds;
        query-response-interval seconds;
        robust-count number;
        ssm-map ssm-map-name;
        version version;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name statement-name protocols igmp amt],
[edit logical-systems logical-system-name routing-instances routing-instance-name statement-name
protocols igmp amt],
[edit protocols igmp amt],
[edit routing-instances routing-instance-name statement-name protocols igmp amt]
```

Description

Configure default Automatic Multicast Tunneling (AMT) interface attributes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Default IGMP Parameters for AMT Interfaces

relay-option-82

IN THIS SECTION

- Syntax | 14294
- Syntax (QFX Series) | 14295
- Hierarchy Level | 14296
- Description | 14296
- Required Privilege Level | 14296
- Release Information | 14297

Syntax

```
relay-option-82 {
    circuit-id {
        include-irb-and-l2;
        keep-incoming-circuit-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-vlan-id;
    }
    remote-id {
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
```

```
        prefix prefix;
        use-interface-description (logical | device);
        use-vlan-id;
    }
    server-id-override;
    vendor-specific{
        host-name;
        location;
    }
}
```

Syntax (QFX Series)

```
relay-option-82 {
    circuit-id {
        include-irb-and-l2;
        keep-incoming-circuit-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-vlan-id;
    }
    remote-id {
        include-irb-and-l2;
        keep-incoming-remote-id ;
        no-vlan-interface-name;
        prefix prefix;
        use-interface-description (logical | device);
        use-vlan-id;
    }
    server-id-override;
    link-selection;
    vendor-specific{
        host-name;
        location;
    }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name]
```

Description

Enable or disable the insertion of the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server.

To enable insertion of option 82 information in DHCP packets, you must specify at least one of the `circuit-id` or `remote-id` statements.

You can use the `relay-option-82` statement and its subordinate statements at the `[edit forwarding-options dhcp-relay]` hierarchy level to control insertion of option 82 information globally, or at the `[edit forwarding-options dhcp-relay group group-name]` hierarchy level to control insertion of option 82 information for a named group of interfaces.

To restore the default behavior (option 82 information is not inserted into DHCP packets), use the `delete relay-option-82` statement.

Starting in Junos OS Release 21.2R1, on QFX Series devices, we've introduced `link-selection` statement at the `edit forwarding-options dhcp-relay relay-option-82` hierarchy level, which allows DHCP relay to add suboption 5 to option 82. Suboption 5 allows DHCP proxy clients and relay agents to request an IP address for a specific subnet from a specific IP address range and scope. Earlier to this release, the DHCP relay drops packets during the renewal DHCP process as the DHCP Server uses the leaf's address as a destination to acknowledge DHCP renewal message.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

`link-selection` option added in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information

dhcp-relay | 5620

relay-option (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 14297
- Hierarchy Level | 14298
- Description | 14298
- Required Privilege Level | 14298
- Release Information | 14299

Syntax

```
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
```

```
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Configure the extended DHCP relay agent selective processing that is based on DHCP options in DHCP client packets and specify the action to perform on client traffic. You can configure support globally or for a named group of interfaces, and for either DHCP or DHCPv6 relay agent.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic

relay-option-vendor-specific (dhcpv6)

IN THIS SECTION

- Syntax | 14299
- Hierarchy Level | 14299
- Description | 14300
- Required Privilege Level | 14300
- Release Information | 14300

Syntax

```
relay-option-vendor-specific {
    host-name;
    location;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 ],
```

Description

Add vendor-specific information to the vendor-specific option (17) field of DHCPv6 control messages on server-facing interfaces. The vendor-specific information can be a hostname, a location (such as a unique connection identifier), or both. The hostname can be a string of characters such as **Juniper-AB-1**. The location should be specified as interface, VLAN ID, and if applicable, stacked VLAN ID. For example, **<ifd-name>:<vlan>** (ae0:100) or **<ifd-name>:<svlan> -<vlan>** (ae0:100-10).

Junos OS automatically adds the remaining vendor-specific information as per RFC 4243. The enterprise ID is 2636. The hostname is option-data 1, and the location is option-data 2. The DHCPv6 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query a single entity to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

relay-server-group (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 14301
- Hierarchy Level | 14301
- Description | 14301
- Options | 14301
- Required Privilege Level | 14301
- Release Information | 14301

Syntax

```
relay-server-group relay-server-group;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option (default-action | equals | starts-with),
[edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action | equals | starts-with),
[edit forwarding-options dhcp-relay group group-name relay-option (default-action | equals | starts-
with),
[edit forwarding-options dhcp-relay dhcpv6 group group-name relay-option (default-action | equals
| starts-with),
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Relay DHCP client packets to the specified group of DHCP servers when you use the DHCP relay selective processing feature. You can configure the relay operation globally or for a group of interfaces, and for either DHCP or DHCPv6 relay agent.

Options

relay-server-group	Name of DHCP server group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Using DHCP Option Information to Selectively Process DHCP Client Traffic

relay-source

IN THIS SECTION

- Syntax | 14302
- Hierarchy Level | 14302
- Description | 14303
- Options | 14303
- Required Privilege Level | 14303
- Release Information | 14303

Syntax

```
relay-source interface-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
interface interface-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
```

```
options dhcp-relay group group-name interface interface-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Configure DHCP relay server to use the configured loopback address as the source address in both the IP header and in DHCP messages relayed to the server.

In network configurations where a firewall on the broadband network gateway (BNG) is between the DHCP relay agent and the DHCP server, only the BNG loopback address passes through the BNG firewall. In that case, DHCP unicast packets do not pass through the firewall and are discarded. This configuration statement places the loopback address in IP headers and DHCP messages, which enables DHCP unicast packets to pass through the firewall to the DHCP server.

Options

interface-name Specify the loopback interface. The *interface-name* must be `lo0`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Configuring the DHCP Relay Agent Source Address to Enable DHCP Packets to Pass Through a Firewall*

reliable

IN THIS SECTION

- Syntax | 14304
- Hierarchy Level | 14304
- Description | 14304
- Default | 14305
- Required Privilege Level | 14305
- Release Information | 14305

Syntax

```
(reliable | no-reliable);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit logical-systems logical-system-name protocols rsvp peer-interface peer-interface-name],
[edit protocols rsvp interface interface-name],
[edit protocols rsvp peer-interface peer-interface-name]
```

Description

Enable reliable message delivery on the interface.

To have both refresh reduction and reliable delivery, enable both the `aggregate` and `reliable` statements.

> **NOTE**: For Junos OS Release 16.1 running on MX Series or PTX Series routers, setting `no-reliable` on an interface automatically disables the fast reroute (FRR) scalability enhancements, including refresh reduction, for all label-switched paths (LSPs) traversing the interface.

Default

Starting in Junos OS Release 16.1R1, all refresh reduction extensions are enabled by default.

Prior to Junos OS Release 16.1R1, the `reliable` option is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring RSVP Refresh Reduction

aggregate (Protocols RSVP) | **2068**

no-enhanced-frr-bypass (Protocols RSVP) | **11451**

remote

IN THIS SECTION

- Syntax | **14306**
- Hierarchy Level | **14306**
- Description | **14306**
- Options | **14306**
- Required Privilege Level | **14306**
- Release Information | **14306**

Syntax

```
remote (inet-address | mac-address) address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family tcc],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family tcc]
```

Description

For Layer 2.5 VPNs employing an Ethernet interface as the TCC router, configure the location of the remote router. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only.

Options

`inet-address` *address*—The IP address of the remote site.

`mac-address` *address*—The MAC address of the remote site.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits

remote-access (Juniper Secure Connect)

IN THIS SECTION

- Syntax | 14307
- Hierarchy Level | 14309
- Description | 14309
- Options | 14309
- Required Privilege Level | 14310
- Release Information | 14310

Syntax

```
remote-access {
    client-config name {
        application-bypass {
            term name {
                description description;
                protocol protocol;
                domain-name domain-name
     }
}
        biometric-authentication;
        certificate {
            no-expiry-warning;
            no-pin-request-per-connection;
            warn-before-expiry days;
        }
        connection-mode (always | manual);
        dead-peer-detection {
            interval seconds;
            threshold threshold;
        }
        no-dead-peer-detection;
        no-eap-tls;
        no-tcp-encap;
        windows-logon {
```

```
            auto-dialog-open;
            disconnect-at-logoff;
            domain domain;
            eap-auth;
            flush-credential-at-logoff;
            lead-time-duration seconds;
            mode (automatic | manual);
        }
    }
    compliance
        pre-logon name {
        term term-name {
             match {
                 platform {
                    (android | ios | macos | windows) {
                        (app-version | os-version) {
                            (equal | greater-than | greater-than-or-equal | less-than | less-
than-or-equal) version;
                    }
                }
            }
                hostname value;
                ms-domain value;
                ms-workgroup value;
                deviceid value;
            }
             action (accept | reject);
    }
}
    default-profile default-profile;
    global-options {
        auth-token-valid-time seconds;
    }
    profile realm-name {
        access-profile access-profile;
        client-config client-config;
        compliance {
            pre-logon compliance-rule;
        description description;
        ipsec-vpn ipsec-vpn;
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
```

```
readable)>;
        flag name;
        level (brief | detail | extensive | verbose);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure remote access settings.

You must configure the remote client settings on SRX Series Firewall to facilitate auto configuration for Juniper Secure Connect remote clients.

When a remote client downloads Juniper Secure Connect application, the application establishes an HTTPS connection with the security device. All authenticated clients fetch the configuration file from the security device and establish a VPN tunnel. This step eliminates the need for the remote clients to configure parameters for certificate identifier parameters, remote access client settings, and IKE and IPsec parameters on their device to establish a VPN connection.

Options

client-config Define Juniper Secure Connect remote client configuration parameters.

compliance Configure the compliance rules for the Juniper Secure Connect client's connection request

default-profile Configure default profile. On your security device, you must specify one of the remote-access profiles as the default profile.

> **NOTE**: Starting in Junos OS Release 23.1R1, we've hidden the `default-profile` option at the [`edit security remote-access`] hierarchy level. In releases before Junos OS Release 23.1R1, you use this option to specify one of the remote-access profiles as the default profile in Juniper

Secure Connect. But with changes to the format of remote-access profile names, we no longer require the `default-profile` option.

We've deprecated `default-profile` option—rather than immediately removing it—to provide backward compatibility and a chance to make your existing configuration conform to the changed configuration. You'll receive a warning message if you continue to use the `default-profile` option in your configuration. However existing deployments are not affected if you modify the current configuration. See default-profile (Juniper Secure Connect).

global-options Define global parameters for Juniper Secure Connect remote access configuration.

profile Configure remote user connection profiles for the Juniper Secure Connect clients.

traceoptions Configure remote access tracing operations for Juniper Secure Connect.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.3R1.

Support for prelogon compliance rule is added in Junos OS Release 23.1R1.

Support for application bypass is added in Junos OS Release 23.1R1.

Usage of `default-profile` options is not allowed starting Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Juniper Secure Connect Administrator Guide

remote-address

IN THIS SECTION

- Syntax | 14311
- Hierarchy Level | 14311
- Description | 14311
- Options | 14311
- Required Privilege Level | 14312
- Release Information | 14312

Syntax

```
remote-address (ipv4-address ipv4-address | ipv6-address ipv6-address);
```

Hierarchy Level

```
[edit unified-edge pcef flow-descriptions flow-identifier],
[edit services pcef flow-descriptions flow-identifier]
```

Description

Specify a remote IP address for the service data flow (SDF) filter.

If you are using Junos OS Subscriber Aware, specify the remote IP address at the `[edit unified-edge pcef flow-descriptions flow-identifier]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the remote IP address at the `[edit services pcef flow-descriptions flow-identifier]` hierarchy level.

Options

ipv4-address	IPv4 address.

ipv6-address IPv6 address.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef flow-descriptions flow-identifier]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

Configuring Service Data Flow Filters

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Understanding Application-Aware Policy Control for Subscriber Management

remote-address (for LMP Control Channel)

IN THIS SECTION

- Syntax | **14313**
- Hierarchy Level | **14313**
- Description | **14313**
- Options | **14313**

- Required Privilege Level | 14313
- Release Information | 14313

Syntax

```
remote-address ip-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-control-channel control-channel-interface],
[edit protocols link-management peer peer-name lmp-control-channel control-channel-interface]
```

Description

Specify the remote IP address for the Link Management Protocol (LMP) control channel interface.

Options

ip-address—Remote IP address mapped to the LMP control channel interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring the Remote IP Address for LMP Control Channels

remote-address (for LMP Traffic Engineering)

IN THIS SECTION

- Syntax | 14314
- Hierarchy Level | 14314
- Description | 14314
- Options | 14314
- Required Privilege Level | 14315
- Release Information | 14315

Syntax

```
remote-address ip-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],
[edit logical-systems logical-system-name protocols link-management te-link te-link-name
interface interface-name],
[edit protocols link-management te-link te-link-name],
[edit protocols link-management te-link te-link-name interface interface-name]
```

Description

Specify the remote IP address for the traffic engineering link. If you configure the remote IP address, you must also configure the "local-address" on page 9908 statement.

Options

ip-address—Remote IP address mapped to the traffic engineering link.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Remote IP Address for Traffic Engineering Links

Configuring the Remote IP Address for Forwarding Adjacencies

local-address (Protocols Link Management) | 9908

remote-debug-permission

IN THIS SECTION

- Syntax | 14315
- Hierarchy Level | 14316
- Description | 14316
- Default | 14316
- Options | 14316
- Required Privilege Level | 14316
- Release Information | 14316

Syntax

```
remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);
```

Hierarchy Level

```
[edit system login user username authentication]
[edit system root-authentication]
```

Description

(QFabric systems only) Configure authentication classes that permit or deny user access to individual components of the QFabric system.

Default

qfabric-user

Options

qfabric-admin Permits a user to log in to individual QFabric system components, view operations, and change component configurations.

qfabric-operator Permits a user to log in to individual QFabric system components and view component operations.

qfabric-user Prevents a user from logging in to individual QFabric system components.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D20.

RELATED DOCUMENTATION

Example: Configuring QFabric System Login Classes

request component login | **14580**

remote-device-management

IN THIS SECTION

- Syntax | 14317
- Hierarchy Level | 14318
- Description | 14318
- Required Privilege Level | 14318
- Release Information | 14318

Syntax

```
remote-device-management {
    service-device device-name {
        access-domain {
            vlan-id-list [vlan-id-low-vlan-id-high vlan-id]
        }
        address ip-address;
        dictionary absolute file path;
        provisioning-method {
            netconf {
                bulk-interval milliseconds;
                bulk-limit number;
                connection-retry-interval seconds;
                password password;
                port port-number;
                reconfigure-bulk-limit number;
                response-timeout seconds;
                response-timeout-count number;
                user-name name;
            }
        }
```

```
    }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name system services],
[edit system services]
```

Description

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Configuring Remote Device Management for Service Provisioning

Remote Device Services Manager (RDSM) Overview

remote-engine

IN THIS SECTION

- Syntax | 14319
- Hierarchy Level | 14319
- Description | 14320

- Options | 14320
- Required Privilege Level | 14320
- Release Information | 14320

Syntax

```
remote-engine engine-id {
    user username {
        authentication-md5 {
            (authentication-key authentication-key | authentication-password authentication-
password);
        }
        authentication-none;
        authentication-sha {
            (authentication-key authentication-key | authentication-password authentication-
password);
        }
        privacy-aes128 (privacy-key privacy-key | privacy-password privacy-password);
        }
        privacy-des {
             (privacy-key privacy-key | privacy-password privacy-password);
        }
        privacy-3des {
             (privacy-key privacy-key | privacy-password privacy-password);
        }
        privacy-none {
        }
    }
}
```

Hierarchy Level

```
[edit snmp v3 usm]
```

Description

Configure the remote engine information for the user-based security model (USM). To send inform messages to an SNMPv3 user on a remote device, you must configure the engine identifier for the SNMP agent on the remote device where the user resides.

Options

engine-id—Specify engine identifier in hexadecimal format. Used to compute the security digest for authenticating and encrypting packets sent to a user on the remote host.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

remote-execution (Event Scripts)

IN THIS SECTION

- Syntax | 14321
- Hierarchy Level | 14321
- Description | 14321
- Options | 14321
- Required Privilege Level | 14321
- Release Information | 14322

Syntax

```
remote-execution {
    remote-hostname {
        passphrase user-password;
        username user-login;
    }
}
```

Hierarchy Level

```
[edit event-options event-script file filename]
```

Description

Configure connection information for a local or remote host, which the event script can use to connect to and execute RPCs on the device.

Remote execution details, which include the hostname, username, and passphrase for one or more hosts, enable an event script to invoke remote procedure calls on a host without encoding the connection information directly in the event script. When you include remote execution details in the configuration instead of in individual event scripts, it has the advantage of capturing the information in a single location and the passphrase is encrypted.

When an event policy invokes an event script, it passes any configured remote execution details to the script. An event script references the remote execution details in the argument list of the function used to create the connection to the remote host.

Options

passphrase *user-password*	User's password for the host.
remote-hostname	Name of the host with which the event script will communicate.
username *username*	User's login name for the host.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Use RPCs and Operational Mode Commands in Event Scripts

Use Event and Remote Execution Details in Event Scripts

remote-gateway

IN THIS SECTION

- Syntax | 14322
- Hierarchy Level | 14322
- Description | 14322
- Options | 14323
- Required Privilege Level | 14323
- Release Information | 14323

Syntax

```
remote-gateway address;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Define the remote address to which the IPsec traffic is directed.

Options

address—Remote IPv4 or IPv6 address.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

remote-gateway (Tunnel Profile)

IN THIS SECTION

- Syntax | 14323
- Hierarchy Level | 14324
- Description | 14324
- Required Privilege Level | 14324
- Release Information | 14324

Syntax

```
remote-gateway {
    address server-ip-address;
```

```
    gateway-name server-name;
}
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Description

Specify the IP address and hostname of the remote gateway at the L2TP tunnel endpoint, the LNS.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

remote (Host VPN)

IN THIS SECTION

- Syntax | 14325
- Hierarchy Level | 14325

- Description | 14325
- Options | 14325
- Required Privilege Level | 14325
- Release Information | 14325

Syntax

```
remote {
    id remote-id;
}
```

Hierarchy Level

```
[edit security host-vpn]
```

Description

Configure the identity details for authenticating the remote device during IKE negotiations.

Options

id *remote-id* Specify the remote IKE identity to use when authenticating the host-to-host VPN connection. The identity can be an IP address, a domain name, or an e-mail address. This identity matches the local identity.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 18.3R1.

remote-id (Services)

IN THIS SECTION

- Syntax | 14326
- Hierarchy Level | 14326
- Description | 14326
- Options | 14326
- Required Privilege Level | 14327
- Release Information | 14327

Syntax

```
remote-id {
    any-remote-id;
    ipv4_addr [ values ];
    ipv6_addr [ values ];
    key_id [ values ];
    fqdn fqdn
    distinguished-name  container container-string-values |wildcard wildcard-string-values
}
```

Hierarchy Level

```
[edit services ipsec-vpn ikepolicy policy-name]
```

Description

Define the remote identification values to which the IKE policy applies.

Options

`any-remote-id`–Allow any remote address to connect. This option is supported only in dynamic configurations and cannot be configured with specific values.

`ipv4_addr [ values ]`—Define one or more IPv4 address identification values.

`ipv6_addr [ values ]`—Define one or more IPv6 address identification values.

`key_id [ values ]`—Define one or more key identification values.

`fqdn fqdn`—Fully-qualified domain name.

`distinguished-name container container-string-values |wildcard wildcard-string-values`—One or more distinguished name values.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`ipv6_addr` option added in Junos OS Release 7.6.

`any-remote-id` option added in Junos OS Release 8.2.

`distinguished-name container container-string-values |wildcard wildcard-string-values` option added in Junos OS Release 19.1.

RELATED DOCUMENTATION

Configuring IKE Policies

remote-id

IN THIS SECTION

- Syntax | 14328
- Hierarchy Level | 14328
- Description | 14328

- Default | 14329
- Required Privilege Level | 14329
- Release Information | 14329

Syntax

```
remote-id {
    host-name host-name;
    mac (Option 82);
    prefix ( hostname | mac | none);
    use-interface-description (logical | device);
    use-string string;
}
```

Hierarchy Level

- For platforms with Enhanced Level 2 Software (ELS):

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]
```

- For platforms without ELS:

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name) dhcp-option82],
[edit forwarding-options helpers bootp dhcp-option82],
[edit forwarding-options helpers bootp interface interface-name dhcp-option82]
```

Description

Insert the `remote-id` suboption of DHCP option 82 (also known as the DHCP relay agent information option) in DHCP request packet headers before forwarding or relaying requests to a DHCP server. This suboption provides a trusted identifier for the host system that has forwarded or relayed requests to the server.

The remaining statements are explained separately, and their availability depends on the hierarchy level at which the `remote-id` suboption is specified, as follows:

- The statement `prefix`, is *not* supported at the `[edit vlans vlan-name forwarding-options dhcp-security option-82]` hierarchy level.
- The statement `host-name` is supported *only* at the `[edit vlans vlan-name forwarding-options dhcp-security option-82]` hierarchy level.

Default

If the `remote-id` statement is not explicitly set, no remote ID value is inserted in the DHCP request packet header.

If the `remote-id` statement is explicitly set, but is not qualified by a keyword, the following are true:

- At the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level, the default keyword value is *interface-name*.
- At all other hierarchy levels, the default value of the `remote-id` keyword is the MAC address of the switch.

> **NOTE**: When you configure `remote-id`, `circuit-id` is also enabled, even if you do not explicitly configure `circuit-id`.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security option-82]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

> *Example: Setting Up DHCP Option 82*
>
> http://tools.ietf.org/html/rfc3046

remote-id (protocols link-management)

IN THIS SECTION

- Syntax | 14330
- Hierarchy Level | 14330
- Description | 14330
- Options | 14330
- Required Privilege Level | 14331
- Release Information | 14331

Syntax

```
remote-id id-number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],
[edit logical-systems logical-system-name protocols link-management te-link te-link-name
interface interface-name],
[edit protocols link-management te-link te-link-name],
[edit protocols link-management te-link te-link-name interface interface-name]
```

Description

Specify the ID assigned to a traffic engineering link or an interface (resource) on the peer node.

Options

id-number—ID number for the remote device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Remote ID for Traffic Engineering Links

remote-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 14331
- Hierarchy Level | 14332
- Description | 14332
- Required Privilege Level | 14334
- Release Information | 14334

Syntax

```
remote-id {
    include-irb-and-l2;
    keep-incoming-remote-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-vlan-id;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82],
[edit forwarding-options dhcp-relay group group-name relay-option-82],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ... relay-option-82],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ... relay-option-82],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ... relay-option-82]
```

Description

Specify the Agent Remote ID suboption (suboption 2) of the DHCP relay agent information option (option 82) to include in DHCP packets destined for a DHCP server. Optionally specify that the suboption includes a prefix, textual description, or VLAN tag.

> **NOTE**: For Layer 3 interfaces, when you configure relay-option-82 only, the Agent Remote ID is the default If no VLAN tags are configured, then the default is the logical interface device (IFL) name. For integrated routing and bridging (IRB) interfaces, the default is the Layer 2 IFL name and bridge domain name.
>
> The interface to bridge domain relationship may be implicit (the interface is mapped to the bridge domain by the system based on VLAN tag) or explicit (the interface is mapped to the bridge domain by configuring it in the bridge domain definition). For the explicit case, tagging might not be relevant for the mapping.

The format of the Agent Remote ID information for Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual LANs (VLANs), stacked VLANs (S-VLANs), or bridge domains is as follows:

```
(fe | ge)-fpc/pic/port.subunit
```

> **NOTE**: For remote systems, the *subunit* is required and is used to differentiate an interface.

The format of the Agent Remote ID information for Fast Ethernet or Gigabit Ethernet interfaces that use VLANs is as follows:

```
(fe | ge)-fpc/pic/port:vlan-id
```

The format of the Agent Remote ID information for Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs is as follows:

```
(fe | ge)-fpc/pic/port:svlan-id--vlan-id
```

In the case of an IRB interface, the format displays the Layer 2 interface instead of the IRB interface along with the bridge domain name. For IRB interfaces (or other pseudo devices) the default format is as follows:

- IRB interfaces that use bridge domains but do not use VLANs or S-VLANs:

  ```
  (fe | ge)-fpc/pic/port.subunit:bridge-domain-name
  ```

- IRB interfaces that use VLANs:

  ```
  (fe | ge)-fpc/pic/port.subunit:vlan-name
  ```

To include the IRB interface name with the Layer 2 interface name, configure the `include-irb-and-l2` statement. The format is as follows:

- IRB interfaces that use bridge domains but do not use VLANs or S-VLANs:

  ```
  (fe | ge)-fpc/pic/port:bridge-domain-name+irb.subunit
  ```

- IRB interfaces that use VLANs:

  ```
  (fe | ge)-fpc/pic/port:vlan-name+irb.subunit
  ```

To include only the IRB interface name without the Layer 2 interface and bridge domain or VLAN, configure the `no-vlan-interface-name` statement. The format is as follows:

```
irb.subunit
```

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
21.2R1	Starting with Junos OS Release 21.2R1, QFX Series switches support the `remote-id` CLI statement in a stateless DHCP relay configuration. You can configure stateless DHCP relay using the `forward-only` CLI statement at the [`edit forwarding-options dhcp-relay`] hierarchy level.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information

Configuring Option 82 Information

remote-id-mismatch (DHCP Local Server and DHCP Relay Agent)

IN THIS SECTION

- Syntax | 14335
- Hierarchy Level | 14335
- Description | 14335
- Required Privilege Level | 14336
- Release Information | 14336

Syntax

```
remote-id-mismatch disconnect;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name]
```

Description

Configure the DHCP local server or DHCP relay agent to detect a mismatch in the Agent Remote ID value to trigger a new connection request. Information about a subscriber's service plan is encoded in the Agent Remote ID, which is conveyed in option 82, suboption 2, for DHCPv4 clients and in option 37 for DHCPv6 clients. When a subscriber session is activated, the Agent Remote ID value for the authorized service plan is stored in the session database. When you configure `remote-id-mismatch`, the DHCP local server and relay agent inspect incoming renew and rebind messages and compare the Agent Remote ID in the message against the initial value that DHCP stored in the database. When DHCP local server discovers a mismatch between the stored value and the value in the message, DHCP local server sends a NAK to the client and tears down the client binding. If the client is a DHCPv6 client, because DHCPv6 does not support an explicit NAK message, the local server sends a reply packet with lifetime set to 0 to signify a logical NAK.

When DHCP relay agent discovers the mismatch, it sends a NAK or logical NAK (for DHCPv6) to the DHCP client. The relay agent cannot tear down the binding itself, so it sends a release message to the local server, causing the local server to tear down the binding. For this to happen, you must configure the `send-release-on-delete` statement on the DHCP relay agent; otherwise it will not send the release message to the local server. In that case, the local server retains the client entry in the database until it times out or the IP address is used for a different binding.

> **NOTE**: `remote-id-mismatch` functionality overrides the default DHCP relay agent bind-on-request behavior. By default, when a stray DHCP request is received, that is, one for which there is an entry in the local server database but not in the relay agent database, a complete binding is automatically made with the relay agent and the local server.

The DHCP client initiates renegotiation when it receives the NAK. The changed Agent Remote ID value is conveyed as part of the request, enabling the new service plan to be submitted for authorization.

The `remote-id-mismatch` statement is typically used in an environment that uses local authorization instead of RADIUS authorization.

> **NOTE**: You cannot configure both the `remote-id-mismatch` statement and the `reauthenticate` statement at the global level, `[edit system services dhcp-local-server]`. However, DHCP precedence rules do permit you to configure both statements when they are at different levels. For example, you can configure `reauthenticate` at the global level and `remote-id-mismatch` for DHCPv6 at the `[edit system services dhcp-local-server dhcpv6]` hierarchy level or for a specific group at the `[edit system services dhcp-local-server group name]` hierarchy level, and so on.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring DHCP-Initiated Service Change Based on Remote ID

remote-id (MX Series)

IN THIS SECTION

- Syntax | 14337
- Hierarchy Level | 14337
- Description | 14337
- Default | 14338
- Required Privilege Level | 14338
- Release Information | 14338

Syntax

```
remote-id {
    host-name;
    use-interface-description (logical | device);
    use-string string;
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security option-82]
```

Description

Insert the `remote-id` suboption of DHCP option 82 (also known as the DHCP relay agent information option) in DHCP request packet headers before forwarding or relaying requests to a DHCP server. This suboption provides a trusted identifier for the host system that has forwarded or relayed requests to the server.

The remaining statements are explained separately. See CLI Explorer.

Default

If the `remote-id` statement is not explicitly set, no remote ID value is inserted in the DHCP request packet header.

If the `remote-id` statement is explicitly set, but is not qualified by a keyword, the default value is the device MAC address.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

> *Configuring DHCP Option 82 on a Router with Bridge Domain*
>
> http://tools.ietf.org/html/rfc3046

remote-id (VLAN Authentication Username)

IN THIS SECTION

- Syntax | 14339
- Hierarchy Level | 14339
- Description | 14339
- Required Privilege Level | 14339
- Release Information | 14339

Syntax

```
remote-id;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include]
```

Description

Include the agent remote identifier (ARI) in the username sent to RADIUS for authentication of the dynamic VLAN. The ARI is conveyed by the Access-Loop-Remote-ID TLV in an out-of-band ANCP Port Up message.

NOTE: This statement is not supported for stacked VLANs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs

Configuring VLAN Interface Username Information for AAA Authentication

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

remote-identity

IN THIS SECTION

- Syntax | 14340
- Hierarchy Level | 14340
- Description | 14341
- Options | 14341
- Usage | 14341
- Required Privilege Level | 14342
- Release Information | 14342

Syntax

```
remote-identity {
    distinguished-name {
        container container-string;
        wildcard wildcard-string;
    }
    hostname hostname;
    inet ip-address;
    inet6 ipv6-address;
    key-id;
    user-at-hostname e-mail-address;
}
```

Hierarchy Level

```
[edit security ike gateway gateway-name]
[edit system ntp server server-name nts]
```

Description

Specify the remote IKE identity to exchange with the destination peer to establish communication. If you do not configure a remote-identity, the device uses the IPv4 or IPv6 address corresponding to the remote endpoint by default.

For Network Address Translation Traversal (NAT-T), both remote identity and local identity must be configured. You can use the NTS 'remote-identity' configuration to verify server identity.

Options

- `distinguished-name`—Specify identity as the distinguished name (DN) from the certificate. If there is more than one certificate on the device, use the `security ike gateway` *`gateway-name`* `policy` *`policy-name`* `certificate local-certificate` *`certificate-id`*.

  Optional container and wildcard strings can be specified:

  - `container` *`container-string`*—Specify a string for the container.
  - `wildcard` *`wildcard-string`*—Specify a string for the wildcard.
- `hostname` *`hostname`*—Specify identity as a fully qualified domain name (FQDN).
- `inet` *`ip-address`*—Specify identity as an IPv4 address.
- `inet6` *`ipv6-address`*—Specify identity as an IPv6 address.
- `key-id` *`string-key-id`*—Specify the key ID in ASCII sring.
- `user-at-hostname` *`e-mail-address`*—Specify identity as an e-mail address.

Usage

The command option `hostname` is the FQDN of the server. This hostname must match the FQDN of server in the `Subject Alternative Name` field of the local certificate.

For example:

```
set system ntp server 10.10.10.10 nts remote-identity hostname junos.example.net
```

You can use the optional parameter `distinguished-name`, to verify the distinguished name (DN) in the server certificate. You can specify the DN and values in either container string or wildcard string format. If we use container option, the order of the fields in DN and their values must exactly match with the fields in the server certificate. If we use wildcard option, the order of the fields in the DN is not considered.

For example,

```
set system ntp server 10.10.10.10 nts remote-identity distinguished-name container
                ou=eng, o=example
```

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

IPsec Overview

remote-interface-switch

IN THIS SECTION

- Syntax | 14343
- Hierarchy Level | 14343
- Description | 14343
- Options | 14343
- Required Privilege Level | 14343
- Release Information | 14344

Syntax

```
remote-interface-switch connection-name {
    interface interface-name.unit-number;
    receive-lsp label-switched-path;
    transmit-lsp label-switched-path;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols connections],
[edit protocols connections (MPLS)]
```

Description

Configure MPLS LSP tunnel cross-connects. This makes an association between a CCC interface and two LSPs, one for transmitting MPLS packets from the local provider edge switch to the remote provider edge switch and the other for receiving MPLS packets on the local provider edge switch from the remote provider edge switch.

Options

connection-name—Connection name.

interface ***interface-name.unit-number***—Interface name. Include the logical portion of the name, which corresponds to the logical unit number of the CCC interface.

receive-lsp ***label-switched-path***—Name of the LSP from the connection's source. This LSP name was specified by the `label-switched-path` statement on the remote provider edge switch in the **protocols mpls** stanza.

transmit-lsp ***label-switched-path***—Name of the LSP to the connection's destination. This LSP name was specified by the `label-switched-path` statement on the local provider edge switch in the **protocols mpls** stanza.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Example: Configuring MPLS on EX8200 and EX4500 Switches*
- *Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect*
- *Configuring MPLS on Provider Edge Switches Using IP-Over-MPLS*
- MPLS Applications User Guide

remote-ip-address

IN THIS SECTION

- Syntax | **14344**
- Hierarchy Level | **14345**
- Description | **14345**
- Options | **14345**
- Required Privilege Level | **14346**
- Release Information | **14346**

Syntax

```
remote-ip-address ip-address {
    local-ip-address;
    announce-interval announce-interval;
    asymmetry asymmetry;
    delay-request delay-request;
    dscp dscp;
    grant-duration grant-duration;
```

```
    sync-interval sync-interval;
}
```

Hierarchy Level

```
[edit logical-systems name              protocols ptp stateful interface name unicast-mode],
[edit protocols ptp stateful interface name unicast-mode]
```

Description

Configure remote IP address for PTP stateful parameters

Options

ip-address IP address of remote PTP partner.

local-ip-address IP address on local interface.

announce-interval Log mean interval between announce messages.

- **Range:** -3 through 0

asymmetry Adjust the timeReceiver-to-timeTransmitter delay by value specified in nanoseconds.

- **Range:** -100000000 through 100000000

delay-request Log mean interval between delay requests.

- **Range:** -7 through 0

dscp IP DSCP value for the stream.

- **Default:** 56
- **Range:** 0 through 63

grant-duration Length of grants in seconds requested during unicast-negotiation

- **Default:** 300
- **Range:** 60 through 1000

sync-interval Requested log mean interval between sync messages

- **Range:** -7 through 0

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

remote-ip-host-routes

IN THIS SECTION

- Syntax | 14346
- Hierarchy Level | 14347
- Description | 14347
- Options | 14347
- Required Privilege Level | 14347
- Release Information | 14347

Syntax

```
remote-ip-host-routes {
    import [ import ... ];
```

```
    no-advertise-community;
}
```

Hierarchy Level

```
[edit routing-instances name protocols evpn],
[edit protocols evpn]
```

Description

Allows you to configure virtual machine traffic optimization (VMTO) for EVPN.

Options

import Policy to control the creation of remote IP host routes

no-advertise-community Enable this option to prevent type-2 MAC and IP community advertisements from being included in type-5 route imports, even though the information may be available. This was the default behavior when importing NLRI type-5 routes from remote PEs type prior to Junos 19.4R1. Starting in Junos 19.4R1, the default behavior is to include the *community* information when importing type-5 routes from type-2 route advertisements (unless this option is enabled).

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.4R1.

The option, `no-advertise-community`, was introduced in Junos OS Release 19.4R1.

remote-loop-enable

IN THIS SECTION

- Syntax | 14348
- Hierarchy Level | 14348
- Description | 14348
- Default | 14349
- Options | 14349
- Required Privilege Level | 14349
- Release Information | 14349

Syntax

```
(remote-loop-enable | no-remote-loop-enable);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

Description

Enable the remote interface to loop back the delay measurement pattern to the local interface. Delay is measured by transmitting a known pattern (delay measurement pattern) in a selected bit of the delay measurement (DM) field and measuring the number of frames that are missed when the delay measurement pattern is received at the transmitting end (local interface).

NOTE: Do not enable remote loopback on both ends (local and remote). If you enable remote loopback on both interfaces, the delay measurement pattern is looped back continuously between the two interfaces.

Default

Delay measurement is disabled by default.

Options

remote-loop-enable Enables loopback of the delay measurement pattern at the remote interface.

no-remote-loop-enable Disables loopback of the delay measurement pattern at the remote interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- *100-Gigabit Ethernet OTN Options Configuration Overview*
- 100-Gigabit DWDM OTN MIC with CFP2-ACO
- 100-Gigabit DWDM OTN PIC with CFP2-ACO (PTX Series)
- *Configuring OTN Interfaces on MIC3-100G-DWDM MIC*
- *Configuring OTN Interfaces on PTX-5-100G-WDM PIC*
- *Disabling ODU Path Delay Measurement*
- *Enabling ODU Path Delay Measurement*
- *Understanding ODU Path Delay Measurement*

remote-loopback

IN THIS SECTION

- Syntax | 14350
- Hierarchy Level | 14350
- Description | 14350
- Required Privilege Level | 14350
- Release Information | 14351

Syntax

```
remote-loopback;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management
interface interface-name]
```

Description

Set the data terminal equipment (DTE) in loopback mode. Remove the statement from the configuration to take the DTE out of loopback mode. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support.

For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, set the remote DTE into loopback mode. Remove the statement from the configuration to take the remote DTE out of loopback mode. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Example: Configure Ethernet OAM Link Fault Management

Configure Ethernet OAM Link Fault Management

remote-maintenance-association (OAM)

IN THIS SECTION

- Syntax | 14351
- Hierarchy Level | 14351
- Description | 14352
- Options | 14352
- Required Privilege Level | 14352
- Release Information | 14352

Syntax

```
remote-maintenance-association remote-ma-name;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-
association ma-name]
```

Description

Configure the name of the remote maintenance association.

Options

remote-ma-name—Name of the remote maintenance association.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Connectivity Fault Management (CFM)

remote-mep

IN THIS SECTION

- Syntax | 14353
- Hierarchy Level | 14353
- Description | 14353
- Options | 14353
- Required Privilege Level | 14353
- Release Information | 14353

Syntax

```
remote-mep mep-id {
    action-profile profile-name;
    sla-iterator-profile profile-name {
        data-tlv-size size;
        iteration-count count-value;
        priority priority-value;
    }
    detect-loc;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-
association ma-name mep mep-id]
```

Description

Configure the numeric identifier of the remote maintenance association end point (MEP) within the maintenance association.

Options

mep-id—Numeric identifier of the MEP.

- **Range:** 1 through 8191

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring Connectivity Fault Management (CFM)

detect-loc | 5538

remote-mep (EX Series Switch Only)

IN THIS SECTION

- Syntax | 14354
- Hierarchy Level | 14354
- Description | 14354
- Options | 14355
- Required Privilege Level | 14355
- Release Information | 14355

Syntax

```
remote-mep mep-id {
    action-profile profile-name;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id ]
```

Description

Specify the numeric identifier of the remote maintenance association end point (MEP) within the maintenance association.

Options

mep-id—Specify the numeric identifier of the MEP.

- **Range:** 1 through 8191

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

> *Configure Ethernet OAM Connectivity Fault Management (CLI Procedure)*
>
> Junos OS Network Interfaces Configuration Guide

remote-port-range

IN THIS SECTION

- Syntax | 14356
- Hierarchy Level | 14356
- Description | 14356
- Default | 14356
- Options | 14356
- Required Privilege Level | 14357
- Release Information | 14357

Syntax

```
remote-port-range {
    low low-value;
    high high-value;
}
```

Hierarchy Level

```
[edit unified-edge pcef flow-descriptions flow-identifier],
[edit services pcef flow-descriptions flow-identifier]
```

Description

Specify the remote port range to identify the subscriber traffic that you want the service data flow (SDF) filter to detect.

If you are using Junos OS Subscriber Aware, specify the remote port range at the `[edit unified-edge pcef flow-descriptions flow-identifier]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the remote port range at the `[edit services pcef flow-descriptions flow-identifier]` hierarchy level.

NOTE: You can specify either a remote port range or a list of remote ports, but not both.

Default

If you configure neither the `remote-port-range` nor the `remote-ports` statement, the default is any remote port.

Options

high-value Upper boundary for the remote port range.

- **Range:** 1 through 65,535

low-value Lower boundary for the remote port range.

- **Range:** 1 through 65,535

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef flow-descriptions flow-identifier]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

Configuring Service Data Flow Filters

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Understanding Application-Aware Policy Control for Subscriber Management

remote-ports

IN THIS SECTION

- Syntax | 14358
- Hierarchy Level | 14358
- Description | 14358
- Default | 14358

- Options | 14359
- Required Privilege Level | 14359
- Release Information | 14359

Syntax

```
remote-ports [number];
```

Hierarchy Level

```
[edit unified-edge pcef flow-description flow-identifier],
[edit services pcef flow-description flow-identifier]
```

Description

Specify a remote port or list of remote ports to identify the subscriber traffic that you want the service data flow (SDF) filter to detect.

If you are using Junos OS Subscriber Aware, specify the remote ports at the `[edit unified-edge pcef flow-descriptions flow-identifier]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the remote ports at the `[edit services pcef flow-descriptions flow-identifier]` hierarchy level.

NOTE: You can specify either a list of remote ports or a remote port range, but not both.

Default

If you configure neither the `remote-ports` nor the `remote-port-range` statement, the default is any remote port.

Options

number Port number or list of port numbers. You can specify a maximum of three port numbers in a list.

- **Range:** 1 through 65,535

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef flow-descriptions flow-identifier]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

Configuring Service Data Flow Filters

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Understanding Application-Aware Policy Control for Subscriber Management

remote-site-id

IN THIS SECTION

- Syntax | 14360
- Hierarchy Level | 14360
- Description | 14360
- Options | 14360
- Required Privilege Level | 14361
- Release Information | 14361

Syntax

```
remote-site-id remote-site-ID;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn site site-name interface interface-name],
[edit routing-instances routing-instance-name protocols l2vpn site site-name interface interface-
name]
```

Description

Control the remote interface to which the interface should connect. If you do not explicitly configure the remote site ID, the order of the interfaces configured for the site determines the default value. This statement is optional.

Options

remote-site-ID—Identifier specifying the interface on the remote PE router the Layer 2 VPN routing instance connects to.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring the Remote Site ID*
>
> Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)

remote-traffic-selector

IN THIS SECTION

- Syntax | 14361
- Hierarchy Level | 14362
- Description | 14362
- Options | 14362
- Required Privilege Level | 14362
- Release Information | 14362

Syntax

```
remote-traffic-selector {
    (ipv4-prefix ipv4-prefix | ipv6-prefix ipv6-prefix);
    port port;
    protocol protocol;
}
```

Hierarchy Level

```
[edit security host-vpn connections connection-name children child-name]
```

Description

Configure the remote IPsec traffic to be protected by the child security association. A traffic selector is a traffic filter that defines and identifies the traffic flow permitted between two systems (a specified pair of local and remote addresses) that have IPsec protection.

Options

(ipv4-prefix *ipv4-prefix* | ipv6-prefix *ipv6-prefix*)	Specify traffic to be protected by the child security association using either IPv4 or IPv6 with a prefix. The prefix allows for specifying more general traffic.
port *port*	Specify the port to protect by number or name. For example, port 21 and port ftp refer to the same port.
protocol *protocol*	Specify the protocol to protect by number or name. For example, protocol 6 and protocol tcp refer to the same protocol.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 18.3R1.

remote-vtep-list

IN THIS SECTION

Syntax | 14363

- Hierarchy Level | 14363
- Description | 14363
- Required Privilege Level | 14364
- Release Information | 14364

Syntax

```
remote-vtep-list [remote-vtep-loopback-address ... ];
```

Hierarchy Level

```
[edit routing-instances name],
[edit switch-options]
```

Description

Configure static remote VXLAN tunnel endpoints (VTEPs). This statement configures VTEPs at the global level in a Layer 2 (L2) VXLAN gateway device. Specify one or more loopback addresses of the remote VTEPs.

You don't configure EVPN as part of a static VXLAN configuration. However, on platforms that support EVPN-VXLAN and static VXLAN, you can configure EVPN-VXLAN and static VXLAN on the same device if the two configurations use different sets of remote peer VTEPs.

On some platforms, you can use this statement with the `static-remote-vtep-list` statement to also configure remote VTEPs at the VLAN or bridge domain level. When you use the `static-remote-vtep-list` option, you must configure the `remote-vtep-list` statement to specify the remote VTEPs at the global level as follows:

- In an L2 routing instance of type `virtual-switch` at the `[edit routing-instances name]` hierarchy level.
- In the default switch instance at the `[edit switch-options]` hierarchy level.

Then configure the `static-remote-vtep-list` statement at the VLAN or bridge domain level in the same L2 routing instance or the default switch instance at the following hierarchy levels:

- In a `virtual-switch` L2 routing instance:

- VLAN level: `[edit routing-instances name vlans name vxlan]`
- Bridge domain level: `[edit routing-instances name bridge-domains name vxlan]`

- In the default switch instance:
  - VLAN level: `[edit vlans name vxlan]`
  - Bridge domain level: `[edit bridge-domains name vxlan]`

> **NOTE**: On the PTX10000 line of routers that run Junos OS Evolved, we support static VXLANs as follows:
>
> - Only in routing instances of type `virtual-switch`.
> - At the global and VLAN levels.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 14.1R3 and 14.2R1.

RELATED DOCUMENTATION

Static VXLAN

static-remote-vtep-list | 28566

remotely-controlled

IN THIS SECTION

Syntax | 14365
Hierarchy Level | 14365

- Description | 14365
- Options | 14365
- Required Privilege Level | 14365
- Release Information | 14365

Syntax

```
remotely-controlled;
```

Hierarchy Level

```
[edit services nat pool nat-pool-name pgcp]
```

Description

Configure the addresses and ports in a NAT pool to be remotely controlled by the gateway controller.

Options

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

remove-ie

IN THIS SECTION

- Syntax | 14366
- Hierarchy Level | 14366
- Description | 14366
- Options | 14367
- Required Privilege Level | 14367
- Release Information | 14367

Syntax

```
remove-ie {
    version v1 {
        number ie-number;
        release name;
    }
}
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

Use the remove-ie option to enable the security device to detect and remove 3G-specific attributes from the GPRS tunneling protocol (GTP) packet header when the packet passes into a 2G network. It allows you to retain interoperability when roaming between 2G and 3G networks.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name]` hierarchy level is replaced by `[edit security gtp profile profile-name]`.

Options

- `version v1`—To specify GTP version 1.
- `number` —To specify the user-configured IE number. IE removal by IE number supports all IEs, ranging from 1 to 255.
- `release` —To Specify release number. Available options are:
  - Release R6—To specify R6 IE removal.
  - Release R7—To specify R7 IE removal.
  - Release R8—To specify R8 IE removal.
  - Release R9—To specify R9 IE removal.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Removing R6, R7, R8, and R9 Information Elements from GTP Messages

Understanding GTPv1 Information Element Removal

Example: Removing GTPv1 Information Elements Using IE Number

remove-private

IN THIS SECTION

- Syntax | 14368
- Hierarchy Level | 14368
- Description | 14369
- Options | 14369
- Required Privilege Level | 14370
- Release Information | 14370

Syntax

```
remove-private;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

When advertising AS paths to remote systems, have the local system strip private AS numbers from the AS path. The numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. If the AS path contains the AS number of the external BGP (EBGP) neighbor, BGP does not remove the private AS number.

> **NOTE**: As of Junos OS 10.0R2 and higher, if there is a need to send prefixes to an EBGP peer that has an AS number that matches an AS number in the AS path, consider using the `as-override` statement instead of the `remove-private` statement.

The operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.

Junos OS recognizes the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document.

The set of reserved AS numbers is in the range from 64,512 through 65,535. The 32-bit private ASN scope is in the range from 4,200,000,000 through 4,294,967,294.

Options

all Remove all private AS numbers from the original path. Do not stop the process of removing private AS numbers, even if a public AS number is encountered.

nearest When you use the `all` and `replace` options, choose the last (right-most) public AS number encountered in the original AS path for the replacement value, as the AS path is processed from left to right. If no public AS number is encountered, the default replacement value is used. (See the `replace` option for information about the default replacement value.)

replace When you use the `all` option, instead of a removing private AS numbers, perform a replace operation. The default replacement value for the private AS number is the local AS number at the BGP group level for the BGP peer. If you are unsure about the replacement value, check the local AS value displayed in the output of the `show bgp group` *`group-name`* command.

no-peer-loop-check Peer loop check is removed. By default, the `remove-private` statement has a peer loop check restriction. If a private AS in the AS path has the same value as the configured `peer-as` for the neighbor, `remove-private` does not remove or replace this private AS number. This restriction provides peer-as loop protection. However, you can remove this restriction using the `no-peer-loop-check` option.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`no-peer-loop-check` option added in Junos OS Release 15.1.

RELATED DOCUMENTATION

Example: Removing Private AS Numbers from AS Paths

remove-when-no-subscribers

IN THIS SECTION

- Syntax | 14370
- Hierarchy Level | 14371
- Description | 14371
- Required Privilege Level | 14371
- Release Information | 14371

Syntax

```
remove-when-no-subscribers;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure]
```

Description

Remove subscriber VLANs automatically when no client sessions (for example, DHCP or PPPoE) exist on the VLAN.

Required Privilege Level

routing—To view this statement in the configuration.

routing–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Automatically Removing VLANs with No Subscribers

rename

IN THIS SECTION

- Syntax | 14372
- Description | 14372
- Options | 14372
- Required Privilege Level | 14372
- Release Information | 14373

Syntax

```
rename <statement-path> identifier1 to identifier2
```

Description

Rename an existing configuration statement or identifier.

Options

identifier1	Existing identifier to rename.
identifier2	New name of identifier.
statement-path	(Optional) Path to an existing statement or identifier.

NOTE: For example, to rename interface `ge-0/1/0.0` to `ge-0/1/10.0` at the following hierarchy level:

```
logical-systems {
    logical-system-abc {
    (...)
        protocols {
            ospf {
                area 0.0.0.0 {
                interface ge-0/1/0.0;
```

Issue the following command:

```
rename logical-systems logical-system-abc protocols ospf area 0.0.0.0 interface
ge-0/1/0.0.0 to interface ge-0/1/10.0
```

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Renaming an Identifier in a Configuration

renegotiation (Services)

IN THIS SECTION

- Syntax | 14373
- Hierarchy Level | 14373
- Description | 14373
- Options | 14374
- Required Privilege Level | 14374
- Release Information | 14374

Syntax

```
renegotiation (allow | allow-secure | drop);
```

Hierarchy Level

```
[edit services ssl proxy profile profile-name actions]
```

Description

Specify the renegotiation options.

Options

- `allow`—Allow secure and nonsecure renegotiation.
- `allow-secure`—Allow secure negotiation only.
- `drop`—Drop session on renegotiation request.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

reopen-time

IN THIS SECTION

- Syntax | 14375
- Hierarchy Level | 14375
- Description | 14375
- Default | 14375
- Options | 14375
- Required Privilege Level | 14375
- Release Information | 14375

Syntax

```
reopen-time seconds;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

(MX Series routers only) Configure the value for the reopen timer.

Default

180 seconds

Options

seconds—Time duration after which the port is unblocked.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring MAC Move Parameters

replace

IN THIS SECTION

- Syntax | 14376
- Description | 14376
- Options | 14376
- Required Privilege Level | 14377
- Release Information | 14377

Syntax

```
replace pattern pattern1 with pattern2 <upto n>
```

Description

Replace identifiers or values in a configuration.

Options

pattern1 Text string or regular expression that defines the identifiers or values you want to match.

pattern2 Text string or regular expression that replaces the identifiers and values located with *pattern1*. Juniper Networks uses standard UNIX-style regular expression syntax (as defined in POSIX 1003.2). If the regular expression contains spaces, operators, or wildcard characters, enclose the expression in quotation marks. Greedy qualifiers (match as much as possible) are supported. Lazy qualifiers (match as little as possible) are not.

upto *n* Number of objects replaced. The value of *n* controls the total number of objects that are replaced in the configuration (not the total number of times the pattern occurs). Objects at the same hierarchy level (siblings) are replaced first. Multiple occurrences of a pattern within a given object are considered a single replacement. If you do not specify an upto option, all identifiers and values in the configuration that match *pattern1* are replaced.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

> *Using Global Replace in the Configuration*

replace-ip-source-with (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 14377
- Hierarchy Level | 14377
- Description | 14378
- Required Privilege Level | 14378
- Release Information | 14378

Syntax

```
replace-ip-source-with giaddr;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
```

```
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Configure DHCP Relay Agent to Replace Request and Release Packets with Gateway IP address

replay-attacks

IN THIS SECTION

- Syntax | 14379
- Hierarchy Level | 14379
- Description | 14379
- Default | 14379
- Options | 14380
- Required Privilege Level | 14380
- Release Information | 14380

Syntax

```
replay-attacks {
    threshold value;
}
```

Hierarchy Level

```
[edit security alarms potential-violation]
```

Description

Raise a security alarm when the device detects a replay attack. A replay attack is a form of network attack in which a valid data transmission is maliciously or fraudulently repeated or delayed.

Default

Replay attacks do not raise security alarms.

Options

- `threshold` *value*—Number of reply attacks up to which an alarm is not raised. When the configured number is exceeded, an alarm is raised.
- **Range:** 1 through 100,00,00,000.
- **Default:** 1000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

IPsec Overview

Monitoring VPN Traffic

replay-protect

IN THIS SECTION

- Syntax | 14381
- Hierarchy Level | 14381
- Description | 14381
- Options | 14381
- Required Privilege Level | 14381
- Release Information | 14382

Syntax

```
replay-protect {
    replay-window-size number-of-packets;
}
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Enable replay protection for MACsec.

A replay window size specified using the "replay-window-size" on page 14383 *number-of-packets* statement must be specified to enable replay protection.

When replay protection is enabled, the receiving interface checks the ID number of all packets that have traversed the MACsec-secured link. If a packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window. Replay protection is especially useful for fighting man-in-the-middle attacks.

A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network. Replay protection should not be enabled in cases where packets are expected to arrive out of order. You can require that all packets arrive in order by setting the replay window size to 0.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced on SRX Series Firewalls in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

replay-protect (MX Series)

IN THIS SECTION

- Syntax | **14382**
- Hierarchy Level | **14383**
- Description | **14383**
- Options | **14383**
- Required Privilege Level | **14383**
- Release Information | **14383**

Syntax

```
replay-protect {
    replay-window-size number-of-packets;
}
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Enable replay protection for MACsec.

A replay window size specified using the `replay-window-size` *`number-of-packets`* statement must be specified to enable replay protection.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

replay-window-size

IN THIS SECTION

- Syntax | 14384
- Hierarchy Level | 14384

- Description | 14384
- Default | 14385
- Options | 14385
- Required Privilege Level | 14385
- Release Information | 14385

Syntax

```
replay-window-size number-of-packets;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name replay-protect]
```

Description

Specifies the size of the replay protection window.

This statement has to be configured to enable replay protection.

When MACsec is enabled on an Ethernet link, an ID number is assigned to each packet entering the link. The ID number of the packet is checked by the receiving interface after the packet has traversed the MACsec-enabled link.

When replay protection is enabled, the sequence of the ID number of received packets are checked. If the packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

Default

Replay protection is disabled.

Options

number-of-packets Specifies the size of the replay protection window, in packets.

When this variable is set to 0, all packets that arrive out-of-order are dropped.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

Statement introduced on SRX Series Firewalls in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

Configuring Media Access Control Security (MACsec) on Routers

Configuring MACsec on EX, QFX and SRX Devices

replay-window-size (MX Series)

IN THIS SECTION

Syntax | **14386**

- Hierarchy Level | 14386
- Description | 14386
- Default | 14387
- Options | 14387
- Required Privilege Level | 14387
- Release Information | 14387

Syntax

```
replay-window-size number-of-packets;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name replay-protect]
```

Description

Specifies the size of the replay protection window.

This statement has to be configured to enable replay protection.

When MACsec is enabled on an Ethernet link, an ID number is assigned to each packet entering the link. The ID number of the packet is checked by the receiving interface after the packet has traversed the MACsec-enabled link.

When replay protection is enabled, the sequence of the ID number of received packets are checked. If the packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

Default

Replay protection is disabled.

Options

number-of-packets Specifies the size of the replay protection window, in packets.

When this variable is set to 0, all packets that arrive out-of-order are dropped.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

replicator

IN THIS SECTION

- Syntax | 14388
- Hierarchy Level | 14388
- Description | 14388
- Options | 14388
- Required Privilege Level | 14389

Release Information | 14390

Syntax

```
replicator {
    inet ip;
    vxlan-encapsulation-source-ip (assisted-replicator-ip | ingress-replication-ip | retain-
incoming-source-ip);
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols evpn assisted-replication],
[edit routing-instances name protocols evpn assisted-replication],
[edit protocols evpn assisted-replication]
```

Description

Configure a device in an EVPN network into the assisted replication (AR) replicator role.

AR helps to optimize broadcast, unknown unicast, and multicast (BUM) traffic in EVPN networks. To enable AR, you configure devices in the EVPN network to operate as AR replicator and AR leaf devices.

AR replicator devices take on BUM traffic replication and forwarding tasks for AR leaf devices. When an AR leaf device receives BUM traffic, instead of always performing ingress replication itself to the EVPN core, it forwards the traffic on an AR tunnel to an AR replicator device that can better handle the replication load. The AR replicator device then replicates and forwards the traffic to other overlay tunnels. You configure a loopback interface IP address on each AR replicator device for the AR tunnel.

You can configure multiple AR replicator devices in an EVPN network, and AR leaf devices load-balance among the available AR replicator devices.

Options

inet *ip-address* (Required) IPv4 address for the AR tunnel to the AR replicator.

This must be a loopback interface configured on the AR replicator device that is part of the EVPN instance.

vxlan-encapsulation-source-ip

(Mandatory for EVPN-VXLAN networks with QFX Series AR devices) Specify the type of VXLAN encapsulation source IP address to include in replicated and forwarded traffic.

- Values:
  - `ingress-replication-ip`—(Mandatory for EVPN-VXLAN networks with QFX Series AR devices) Use the ingress replication IP address.

    > **NOTE**: When configuring AR in EVPN-VXLAN networks with QFX Series devices as AR replicators and AR leaf devices, the `vxlan-encapsulation-source-ip` statement must be set to `ingress-replication-ip`.
    >
    > The `ingress-replication-ip` option supports AR on AR replicator devices that are not capable of retaining the source IP address of the originating AR leaf device (sent on the AR overlay tunnel) when replicating the traffic onto other overlay tunnels. In that case, AR must operate in *extended AR mode* (see *Extended AR Mode for Multihomed Ethernet Segments*).You must explicitly set this option to override the default factory configuration setting, `retain-incoming-source-ip`, which is not supported in this case.
    >
    > AR is currently only supported in EVPN-VXLAN network with QFX Series AR devices operating in extended AR mode.

  (Not currently used) The following `vxlan-encapsulation-source-ip` options are reserved for future AR compatibility among AR replicators with different capabilities:

  - `assisted-replicator-ip`—Use the AR replicator IP address.
  - `retain-incoming-source-ip`—(Default) Retain the VXLAN encapsulation source IP address of the incoming packet.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.4R2 and 19.4R1.

RELATED DOCUMENTATION

Assisted Replication Multicast Optimization in EVPN Networks

show evpn multicast-snooping assisted-replication next-hops | **19807**

show evpn multicast-snooping assisted-replication replicators | **19811**

report-ingress-shaping-rate (Dynamic CoS Interfaces)

IN THIS SECTION

- Syntax | **14390**
- Hierarchy Level | **14390**
- Description | **14391**
- Options | **14391**
- Required Privilege Level | **14391**
- Release Information | **14391**

Syntax

```
report-ingress-shaping-rate bps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Report the ingress shaping rate in bits per second that is used as the receive speed (Rx) for the LAC to send to the LNS. The ingress shaping rate is used when the method for deriving the connect speed is configured as `service-profile` with the `tx-connect-speed` statement at the `[edit services l2tp]` hierarchy level.

> **NOTE**: This statement is supported only when the `effective shaping-rate` statement is included at the `[edit chassis]` hierarchy level. If it is not, the subscriber login fails and a system log message is generated. There is no commit check failure because a commit check cannot be performed at run time.

Options

bps Ingress shaping rate in bits per second.

- **Range:** 1000 through 6,400,000,000,000 (6.4Tbps)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

L2TP Subscriber Access Lines and Connection Speeds

L2TP Subscriber Access Lines and Connection Speeds

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying a Classifier to a Subscriber Interface in a Dynamic Profile

report-interface-descriptions (Access)

IN THIS SECTION

- Syntax | 14392
- Hierarchy Level | 14392
- Description | 14392
- Required Privilege Level | 14392
- Release Information | 14392

Syntax

```
report-interface-descriptions;
```

Hierarchy Level

```
[edit access]
```

Description

Enable storing and reporting of interface descriptions through RADIUS. To disable storing and reporting of interface descriptions, configure the `[edit access profile profile-name radius attributes exclude]` statement to exclude the interface description attribute. The description can contain letters, numbers, and hyphens (-), and can be up to 64 characters long.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Storage and Reporting of Interface Descriptions to Uniquely Identify Subscribers

RADIUS Servers and Parameters for Subscriber Access

Configuring Authentication and Accounting Parameters for Subscriber Access

report-local-rule (PCRF Partition)

IN THIS SECTION

- Syntax | 14393
- Hierarchy Level | 14393
- Description | 14393
- Required Privilege Level | 14394
- Release Information | 14394

Syntax

```
report-local-rule;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the Policy and Charging Rule Function (PCRF) partition to send local report downstream messages by default. If you configure this statement, this generates the same result as if you enabled the Resource-Allocation-Notification message trigger with the local rule definition.

The following example shows predefined rule installation commands, including the Resource-Allocation-Notification command, for Credit Control Answer (CCA-GX) and Reauthorization Request (RAR-GX) messages between the PCRF and Gx:

```
{ Charging-Rule-Install
  { Charging-Rule-Name:            “fixed-cos” }
  { Charging-Rule-Definition:
    { Charging-Rule-Name:        “firewall” }
    { Service-Identifier:         10 }
    { Rating-Group:            292 }
    { Juniper-Param-V4-Input-Filter:    “my_input_filter” }
    { Juniper-Param-V4-Output-Filter:    “my_output_filter” }
  }
  [ Resource-Allocation-Notification: ENABLE_NOTIFICATION(0) ]
 }
```

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

Understanding Upstream and Downstream Messages for the PCRF

report (LRF Rule)

IN THIS SECTION

- Syntax | 14395
- Hierarchy Level | 14395
- Description | 14395
- Required Privilege Level | 14395
- Release Information | 14396

Syntax

```
report {
    collector collector-name;
    template template-name;
    time-limit time-interval;
    volume-limit volume;
}
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name then]
```

Description

Configure the actions to take if the LRF rule is matched.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

report-resource-allocation (PCRF Partition)

IN THIS SECTION

- Syntax | 14396
- Hierarchy Level | 14396
- Description | 14397
- Required Privilege Level | 14397
- Release Information | 14397

Syntax

```
report-resource-allocation;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the Policy and Charging Rule Function (PCRF) partition to generate reports for resource allocation failures by default. If you configure this statement, this generates the same result as if you enabled the Resource-Allocation-Notification message trigger with every rule received from the PCRF.

The following example shows predefined rule installation commands, including the Resource-Allocation-Notification command, for Credit Control Answer (CCA-GX) and Reauthorization Request (RAR-GX) messages between the PCRF and Gx:

```
{ Charging-Rule-Install
  { Charging-Rule-Name:            “fixed-cos” }
  { Charging-Rule-Definition:
    { Charging-Rule-Name:        “firewall” }
    { Service-Identifier:         10 }
    { Rating-Group:            292 }
    { Juniper-Param-V4-Input-Filter:    “my_input_filter” }
    { Juniper-Param-V4-Output-Filter:    “my_output_filter” }
  }
  [ Resource-Allocation-Notification: ENABLE_NOTIFICATION(0) ]
 }
```

NOTE: Some PCRFs may be unable to generate a Resource-Allocation-Notification AVP. As a result, the `report-resource-allocation` statement provides generated reports by default.

For login and update interactions between the PCRF and the Online Charging System (OCS), the OCS sends applicable reports to the PCRF using a CCR-GX-U message. If the service-dynamic-profile instantiation fails, and the Resource-Allocation-Notification (ENABLE_NOTIFICATION) is set for the charging rule, then the OCS sends a report. The router acknowledges the subscriber activation, allows subscriber traffic to flow, and continues to repeat the process.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

Understanding Upstream and Downstream Messages for the PCRF

report (Security Log)

IN THIS SECTION

- Syntax | **14398**
- Syntax (SRX4600) | **14399**
- Hierarchy Level | **14400**
- Description | **14400**
- Options | **14400**
- Required Privilege Level | **14402**
- Release Information | **14403**

Syntax

```
report {
    logs-per-table {
        idp log-number;
        ipsec-vpn log-number;
        screen log-number;
        session-all log-number;
        sky log-number;
        utm log-number;
    }
    table-lifetime table-lifetime;
}
    table-mode {
```

```
        dense;
    }
```

Syntax (SRX4600)

```
report {
    database-capacity {
        max-db-size;
    }
    database-filter {
        event-category;
        exclude;
    }
    log-share-per-table {
        atp log-share;
        dns log-share;
        idp log-share;
        ipsec-vpn log-share;
        screen log-share;
        session-all log-share;
        utm log-share;
    }
    logs-per-table {
        idp log-number;
        ipsec-vpn log-number;
        screen log-number;
        session-all log-number;
        sky log-number;
        utm log-number;
    }
    table-lifetime table-lifetime;
    }
    table-mode {
        dense;
        disable-dense;
    }
}
```

Hierarchy Level

```
[edit security log]
[edit logical-systems name security log]
[edit tenants tenant-name security log]
```

Description

Set security log report settings.

On-box reporting offers a comprehensive reporting facility where your security management team can spot a security event when it occurs, immediately access and review pertinent details about the event, and quickly decide appropriate remedial action.

The on-box reporting feature is enabled by default on a SRX Series Firewall with Junos OS Release 15.1X49-D100 or later.

If you are upgrading your SRX Series Firewall from a Junos OS Release prior to Junos OS 15.1X49-D100, then on-box reporting feature is disabled by default. You need to run the `set security log report` command to enable the on-box reporting feature on the device.

Options

report Enable log report.

logs-per-table Log number for each table.

idp Log number of idp.

- **Range:** For SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, and vSRX Virtual Firewall: 40000 through 80000
- **Default:** 40000
- **Range:** For SRX1500, SRX4100, and SRX4200: 500000 through 1000000
- **Default:** 500000
- **Range:** For SRX4600: 1333333 through 2666666
- **Default:** 1333333

ipsec-vpn Log number of IPsec-VPN.

- **Range:** For SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, and vSRX Virtual Firewall: 26666 through 26666
- **Default:** 26666
- **Range:** For SRX1500, SRX4100, and SRX4200: 333333 through 333333
- **Default:** 333333
- **Range:** For SRX4600: 666666 through 666666
- **Default:** 666666

screen Log number of screen.

- **Range:** For SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, and vSRX Virtual Firewall: 40000 through 80000
- **Default:** 40000
- **Range:** For SRX1500, SRX4100, and SRX4200: 500000 through 1000000
- **Default:** 500000
- **Range:** For SRX4600: 1333333 through 2666666
- **Default:** 1333333

session-all Log number of session.

- **Range:** For SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, and vSRX Virtual Firewall: 80000 through 800000
- **Default:** 80000
- **Range:** For SRX1500: 533333 through 5333333
- **Default:** 533333
- **Range:** For SRX4100 and SRX4200: 666666 through 6666666
- **Default:** 666666
- **Range:** For SRX4600: 1000000 through 10000000
- **Default:** 1000000

sky Log number of SKY.

- **Range:** For SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, and vSRX Virtual Firewall: 40000 through 80000
- **Default:** 40000
- **Range:** For SRX1500, SRX4100, and SRX4200: 500000 through 1000000
- **Default:** 500000
- **Range:** For SRX4600: 1333333 through 2666666
- **Default:** 1333333

Content Security

Log number of Content Security.

- **Range:** For SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, and vSRX Virtual Firewall: 40000 through 80000
- **Default:** 40000
- **Range:** For SRX1500, SRX4100, and SRX4200: 500000 through 1000000
- **Default:** 500000
- **Range:** For SRX4600: 1333333 through 2666666
- **Default:** 1333333

table-lifetime

Table lifetime days.

- **Default:** 90
- **Range:** 0 through 365

table-mode

Report table mode.

- `dense`- Enable table dense mode.
- `disable-dense`- Disable table dense mode.

Required Privilege Level

The remaining statements are explained separately. See CLI Explorer.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100

The [edit `logical-systems` *name* security "log" on page 10026] and [edit `tenants` *tenant-name* security "log" on page 10026] hierarchy levels introduced in Junos OS Release 19.1R1.

`table-mode` option added in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

log (Security) | 10026

report-skip

IN THIS SECTION

- Syntax | 14403
- Hierarchy Level | 14403
- Description | 14404
- Required Privilege Level | 14404
- Release Information | 14404

Syntax

```
report-skip
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name],
[edit security policies global policy policy-name],
[edit security policies stateful-firewall-rule [name] policy [name]],
[edit security policies explicit-proxy profile [name] policy [name] report-skip],
```

```
[edit logical-systems logical-system-name security policies from-zone zone-name to-zone zone-name
policy policy-name],
[edit logical-systems logical-system-name security policies global policy policy-name],
[edit tenants tenant-name security policies from-zone zone-name policy policy-name],
[edit tenants tenant-name security policies global policy policy-name]
```

Description

You can configure this command to skip policy reports for a policy.

Adding the policy to the policy allowlist excludes the policy from any policy analysis and excluded it from appearing in any future report.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

show security policy-report | 25089

report-successful-resource-allocation (PCRF Partition)

IN THIS SECTION

- Syntax | 14405
- Hierarchy Level | 14405
- Description | 14405
- Required Privilege Level | 14405
- Release Information | 14406

Syntax

```
report-successful-resource-allocation;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the Policy and Charging Rule Function (PCRF) partition to generate reports for both resource allocation successes and failures by default. If you configure this statement, this generates the same result as if you enabled the Successful-Resource-Allocation-Notification message trigger with every CCA-GX-I, CCA-GX-U, and RAR-GX-A message.

The following example shows a predefined event trigger command for CCA-GX and RAR-GX messages between the PCRF and Gx:

```
{ Event-Trigger: SUCCESSFUL_RESOURCE_ALLOCATION(22) }
```

If the SUCCESSFUL_RESOURCE_ALLOCATION (22) trigger value exists in the downstream data, the Broadband Policy and Charging Enforcement Function (PCEF) reports successful installations of rules marked with Resource-Allocation-Notification AVP in the Charging-Rule-Install AVP.

> **NOTE**: Some PCRFs may be unable to generate this event trigger. As a result, the `report-successful-resource-allocation` statement provides generated reports by default.

For login and update interactions between the PCRF and the Online Charging System (OCS), the OCS sends applicable reports to the PCRF using a CCR-GX-U message. If the service-dynamic-profile instantiation succeeds, and the Resource-Allocation-Notification (ENABLE_NOTIFICATION) is set for the charging rule, and the SUCCESSFUL_RESOURCE_ALLOCATION event trigger is set in the request, then the OCS sends a report.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

Understanding Upstream and Downstream Messages for the PCRF

report-zero-oif-gw-on-discard

IN THIS SECTION

- Syntax | 14406
- Hierarchy Level | 14406
- Description | 14407
- Default | 14407
- Required Privilege Level | 14407
- Release Information | 14407

Syntax

```
report-zero-oif-gw-on-discard;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services]
```

Description

(MX Series only) For inline active flow monitoring, for only the flow records for rejected, discarded, and policed flows, set the flow record value to 0 (zero) for the IPFIX information elements `ipNextHopIPv6Address` (Element ID 62), `egressInterface` (Element ID 14), and `ipNextHopIPv4Address` (Element ID 15). If the `report-zero-oif-gw-on-discard` statement is not configured, the flow records display the available information for these elements, which is the default behavior for not only rejected, discarded, or policed flows, but also for all other flows. The available information for these elements is the result of a route lookup in the forwarding table, and is reported in the OIF/GW fields in the export flow record.

> **NOTE**: This statement does not apply to flows that are policed and sampled by the Packet Forwarding Engine for egress interfaces. After you configure the `report-zero-oif-gw-on-discard` statement, each sampled packet updates the forwarding action of that flow in the flow record. That is, the last sampled packet of the flow just before export determines the forwarding action of that flow. For example, in the case of a rate-limiting policer, the forwarding action taken on a flow is not deterministic. The flow is treated either as forwarded or as policed based on the forwarding action of the last sampled packet of that flow.

You configure this statement to help you track what traffic is being dropped.

Default

The flow records display the available information for the information elements `ipNextHopIPv6Address` (Element ID 62), `egressInterface` (Element ID 14), and `ipNextHopIPv4Address` (Element ID 15).

Required Privilege Level

Release Information

Statement introduced in Junos OS Release 15.1F7.

req-option

IN THIS SECTION

- Syntax | 14408
- Hierarchy Level | 14408
- Description | 14408
- Options | 14408
- Required Privilege Level | 14409
- Release Information | 14409

Syntax

```
req-option (dns-server | domain | fqdn | nis-domain | nis-server | ntp-server | sip-domain | sip-server | time-zone | vendor-spec);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcpv6-client]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet6 dhcpv6-client]
[edit tenants tenant-name interfaces interface-name unit logical-unit-number family inet6 dhcpv6-client]
```

Description

The configuration options requested by the DHCPv6 client.

Options

dns-server Specify a DNS server.

domain Specify a domain name.

fqdn	Specify a fully qualified domain name.
nis-domain	Specify a Network Information Service (NIS) domain.
nis-server	Specify a Network Information Service (NIS) server.
ntp-server	Specify a Network Time Protocol (NTP) server.
sip-domain	Specify a Session Initiation Protocol (SIP) domain.
sip-server	Specify a Session Initiation Protocol (SIP) server.
time-zone	Specify a time zone.
vendor-spec	Specify vendor specification.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

The `logical-systems` and `tenants` options are introduced in Junos OS Release 18.4R1.

req-timeout

IN THIS SECTION

- Syntax | 14410
- Hierarchy Level | 14410
- Description | 14410
- Options | 14410
- Required Privilege Level | 14410
- Release Information | 14410

Syntax

```
req-timeout second;
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

Use the req-timeout option to specify a GTP request message timeout. The default timeout value is 5 seconds.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name]` hierarchy level is replaced by `[edit security gtp profile profile-name]`.

Options

- **Range:** 1 to 30 seconds.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D35.

RELATED DOCUMENTATION

> *Example: Configure the Validity of IP Address in GTP Messages*

request

IN THIS SECTION

- Syntax | 14411
- Description | 14411
- Additional Information | 14412
- Required Privilege Level | 14412
- Release Information | 14412

Syntax

```
request <access-security | ancp | chassis | debug | dhcp | dhcpv6 | extension-service | flight-
recorder | interface | ipsec | jdaf | jnu | l2circuit-switchover | lacp | message | mpls |
network-access | path-computation-client | pim | protection-group | routing-engine | security |
services | session | snmp | support | system | unified-edge |  validation | virtual-chassis |
vpls-switchover>
```

Description

Stop or reboot router components, switch between primary and backup components, display messages, and display system information.

CAUTION: Halt the backup Routing Engine before you remove it or shut off the power to the router; otherwise, you might need to reinstall the Junos OS.

NOTE: If your router contains two Routing Engines and you want to shut the power off to the router or remove a Routing Engine, you must first halt the backup Routing Engine (if it has been upgraded) and then the primary Routing Engine. To halt a Routing Engine, enter the `request system halt` command. You can also halt both Routing Engines at the same time by issuing the request `system halt both-routing-engines` command.

If you want to reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) and then the primary Routing Engine.

NOTE: If you reboot the TX Matrix router, all the T640 primary Routing Engines connected to the TX Matrix router reboot. If you halt both Routing Engines on a TX Matrix router, all the T640 Routing Engines connected to the TX Matrix router are also halted. Likewise, if you reboot the TX Matrix Plus router, all the T1600 or T4000 primary Routing Engines connected to the TX Matrix Plus router reboot. If you halt both Routing Engines on a TX Matrix Plus router, all the T1600 or T4000 Routing Engines connected to the TX Matrix Plus router are also halted.

NOTE: If you insert a Flexible PIC Concentrator (FPC) into your router, you may need to issue the `request chassis fpc` command (or press the **online** button) to bring the FPC online. This applies to FPCs in M20, M40, M40e, M160, M320, and T Series routers. For command usage, see the `request chassis fpc` command description in the CLI Explorer.

Additional Information

Most `request` commands are described in the *Junos System Basics and Services Command Reference*. The following `request` commands are described in the *Junos Interfaces Command Reference*: `request ipsec switch` and `request services`.

Required Privilege Level

maintenance

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Overview of Operational Mode Commands

request access-security router-advertisement-guard-block

IN THIS SECTION

- Syntax | 14413
- Description | 14413
- Options | 14413
- Required Privilege Level | 14413
- Output Fields | 14414
- Release Information | 14414

Syntax

```
request access-security router-advertisement-guard-block interface (interface-name)
```

Description

Initiate the blocking state on an interface or range of interfaces for stateful IPv6 Router Advertisement (RA) guard. Stateful RA guard learns about legitimate senders of RA messages and stores this information in order to validate senders of subsequent RA messages. For example, an interface that is in the learning state and receives RA messages that are not sent from valid IPv6 routers will dynamically transition to the blocking state. While the interface is in blocking state, all RA messages received on that interface are dropped.

You can override the dynamic state transitions by requesting the blocking state on an interface. If you issue the request for the blocking state on an interface, the interface will remain in forwarding state until either the learning or forwarding state is requested on that interface.

Options

`interface` *`interface-name`*	Initiate the blocking state on the specified interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Configuring Stateful IPv6 Router Advertisement Guard

request access-security router-advertisement-guard-forward

IN THIS SECTION

- Syntax | 14414
- Description | 14414
- Options | 14415
- Required Privilege Level | 14415
- Output Fields | 14415
- Release Information | 14415

Syntax

```
request access-security router-advertisement-guard-forward interface (interface-name)
```

Description

Initiate the forwarding state on an interface or range of interfaces for stateful IPv6 Router Advertisement (RA) guard. Stateful RA guard learns about legitimate senders of RA messages and stores

this information in order to validate senders of subsequent RA messages. For example, an interface that is in the learning state and receives RA messages sent from legitimate sources will dynamically transition to the forwarding state. While the interface is in forwarding state, all RA messages received on that interface that can be validated against the configured policy are forwarded.

You can override the dynamic state transitions by requesting the forwarding state on an interface. If you issue the request for the forwarding state on an interface, the interface will remain in forwarding state until either the learning or blocking state is requested on that interface.

Options

`interface` *`interface-name`*	Initiate the forwarding state on the specified interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Configuring Stateful IPv6 Router Advertisement Guard

request access-security router-advertisement-guard-learn interface

IN THIS SECTION

Syntax | 14416

- Description | 14416
- Options | 14416
- Required Privilege Level | 14417
- Output Fields | 14417
- Release Information | 14417

Syntax

```
request access-security router-advertisement-guard-learn interface (interface-name) duration
seconds (forward | block)
```

Description

Request the learning state on an interface or range of interfaces for stateful IPv6 Router Advertisement (RA) guard. Stateful RA guard learns about legitimate senders of RA messages and stores this information in order to validate senders of subsequent RA messages. For example, an interface that is in the learning state and receives RA messages sent from legitimate sources dynamically transitions to the forwarding state after the learning period ends. While the interface is in forwarding state, all RA messages received on that interface that can be validated against the configured policy are forwarded.

Before you can request learning on an interface, you must enable RA guard at the [`edit forwarding-options access-security router-advertisement-guard`] hierarchy level and configure the `stateful` option. When you enable stateful RA guard, the default state is `Off`. An interface in the `Off` state operates as if RA guard is not available. The learning state can be initiated only by configuring the `request access-security router-advertisement-guard-learn` command.

When you request the learning state, you must configure the duration of the learning period in seconds. This is the amount of time the interface will remain in the learning state before it transitions to another state. RA messages that are received during the learning period can be either forwarded or blocked. Configure the `forward` option to forward RA messages during the learning period, or configure the `block` option to block RA messages during the learning period.

Options

interface *interface-name*	Initiate the learning state on the specified interface.

`duration` *`seconds`*	Configure the duration of the learning state in seconds. When the learning period ends, the state dynamically transitions to either the forwarding state or the blocking state.
`forward`	Configure the interface to forward RA messages received during the learning period.
`block`	Configure the interface to block RA messages received during the learning period.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Configuring Stateful IPv6 Router Advertisement Guard

request access-security slaac-snooping unblock

IN THIS SECTION

- Syntax | 14418
- Description | 14418
- Options | 14418
- Required Privilege Level | 14418
- Output Fields | 14418

Sample Output | 14418
Release Information | 14418

Syntax

```
request access-security slaac-snooping unblock
<interface interface-name | all>
```

Description

Clear an interface that has been blocked by SLAAC snooping.

Options

interface (*interface-name* | all)	Clear SLAAC snooping information for the specified interface or all interfaces.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

request access-security slaac-snooping unblock

```
user@switch> request access-security slaac-snooping unblock
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show access-security slaac-snooping binding | 17003

request ancp oam interface

IN THIS SECTION

- Syntax | 14419
- Description | 14419
- Options | 14419
- Required Privilege Level | 14420
- Output Fields | 14420
- Sample Output | 14420
- Release Information | 14420

Syntax

```
request ancp oam interface
(interface-name | interface-set set-name)
<count count>
<timeout duration>
```

Description

Trigger the access node to run a loopback test on the local loop between the access node and the customer premises equipment. You must specify either an ANCP interface or an ANCP interface set. The access node responds to the NAS with the results of the test.

Options

interface-name Name of the ANCP interface on whose local loop the loopback test is run.

interface-set ***set-name***	Name of the ANCP interface set on whose local loop the loopback test is run.
count ***count***	(Optional) Number of times a loopback message is sent on the local loop. Range: `1` through `32`. Default: `1`.
timeout ***duration***	(Optional) Period of time in seconds that the NAS waits for a response to the OAM request. Range: `0` through `255`. Default: `5`.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request, including the result of the test, the response code, and the response string returned with the OAM response in the event of failure, an error code is displayed.

Sample Output

request ancp oam interface

```
user@host> request ancp oam interface ge-1/0/4.12 count 5 timeout 40
request succeeded
 0x503 : DSL line status showtime
 DEFAULT RESPONSE
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

ANCP Monitoring and Management

request ancp oam neighbor

IN THIS SECTION

- Syntax | 14421
- Description | 14421
- Options | 14421
- Required Privilege Level | 14422
- Output Fields | 14422
- Sample Output | 14422
- Release Information | 14422

Syntax

```
request ancp oam neighbor
(ip-address ip-address | system-name neighbor-name)
subscriber identifier-string
<count count>
<timeout duration>
```

Description

Trigger the access node to run a loopback test on the local loop between the access node and the customer premises equipment. You must specify both the access node and the subscriber. The access node responds to the NAS with the results of the test.

Options

ip-address *ip-address*	IP address that specifies the access node on whose local loop the loopback test is run.
system-name *neighbor-name*	System name that specifies the access node on whose local loop the loopback test is run.

subscriber ***identifier-string***	Access identifier that specifies the subscriber on whose local loop the loopback test is run.
count ***count***	(Optional) Number of times a loopback message is sent on the local loop. Range: 1 through 32. Default: 1.
timeout ***duration***	(Optional) Period of time in seconds that the NAS waits for a response to the OAM request. Range: 0 through 255. Default: 5.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request, including the result of the test, the response code, and the response string returned with the OAM response in the event of failure, an error code is displayed.

Sample Output

request ancp oam subscriber

```
user@host> request ancp oam neighbor 203.0.113.21 subscriber “dslam port-1-11”
request succeeded
 0x503 : DSL line status showtime
 DEFAULT RESPONSE
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

ANCP Monitoring and Management

request ancp oam port-down

IN THIS SECTION

- Syntax | 14423
- Description | 14423
- Options | 14423
- Required Privilege Level | 14424
- Output Fields | 14424
- Sample Output | 14424
- Release Information | 14424

Syntax

```
request ancp oam port-down
(neighbor ip-address | subscriber-interface physical-interface-name)
circuit-id aci remote-id ari outer-vlan-id vlan-id
```

Description

Simulate an ANCP Port Down message on the specified access loop for troubleshooting or to mitigate an abnormal condition. Triggers removal of the corresponding out-of-band triggered, autosensed dynamic VLAN session for which no ANCP-sourced information exists. You must specify an ACI, an ARI, and an outer VLAN tag. This command is overridden by a genuine ANCP Port-Up message, meaning that you cannot use this command to initiate a Port Down condition when the access node has already reported a Port Up condition.

Options

aci ANCP Access-Loop-Circuit-ID TLV that corresponds to a subscriber interface on the access node; used to identify the access node from which the message is simulated.

ari	ANCP Access-Loop-Remote-ID TLV that identifies the subscriber associated with an interface on the access node; used to identify the access node from which the message is simulated.
ip-address	IP address that specifies the access node from which the message is simulated.
physical-interface-name	Name of the access-facing subscriber interface that specifies the access node on whose local loop the loopback test is run.
vlan-id	ANCP Access-Aggregation-Circuit-ID-Binary TLV, the outer VLAN tag inserted by the access node on upstream traffic; used to identify the access node from which the message is simulated.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp neighbor detail`, `show subscribers client-type vlan-oob detail`, and the `show subscribers summary` commands before and after initiating the Port Down message to verify the operation.

Sample Output

request ancp oam port-down neighbor circuit-id remote-id outer-vlan-id

```
user@host> request ancp oam port-down neighbor 192.168.25.31 circuit-id line-aci-1 remote-id
line-ari-1 outer-vlan-id 126
```

Release Information

Command introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

request ancp oam port-up

IN THIS SECTION

- Syntax | 14425
- Description | 14425
- Options | 14425
- Required Privilege Level | 14426
- Output Fields | 14426
- Sample Output | 14426
- Release Information | 14426

Syntax

```
request ancp oam port-up
(neighbor ip-address | subscriber-interface physical-interface-name)
circuit-id aci remote-id ari outer-vlan-id vlan-id
```

Description

Simulate an ANCP Port Up message on the specified access loop for troubleshooting or to mitigate an abnormal condition. You must specify an ACI, an ARI, and an outer VLAN tag. This command is overridden by a genuine ANCP Port Down message, meaning that you cannot use this command to initiate a Port Up condition when the access node has already reported a Port Down condition.

Options

aci ANCP Access-Loop-Circuit-ID TLV that corresponds to a subscriber interface on the access node; used to identify the access node from which the message is simulated.

ip-address IP address that specifies the access node from which the message is simulated.

ari ANCP Access-Loop-Remote-ID TLV that identifies the subscriber associated with an interface on the access node; used to identify the access node from which the message is simulated.

physical-interface-name	Name of the access-facing subscriber interface that specifies the access node on whose local loop the loopback test is run.
vlan-id	ANCP Access-Aggregation-Circuit-ID-Binary TLV, the outer VLAN tag inserted by the access node on upstream traffic; used to identify the access node from which the message is simulated.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp neighbor detail`, `show subscribers client-type vlan-oob detail`, and the `show subscribers summary` commands before and after initiating the Port Up message to verify the operation.

Sample Output

request ancp oam port-up neighbor circuit-id remote-id outer-vlan-id

```
user@host> request ancp oam port-up neighbor 192.168.25.31 circuit-id line-aci-1 remote-id line-
ari-1 outer-vlan-id 126
```

Release Information

Command introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

request app-engine cleanup

IN THIS SECTION

- Syntax | 14427
- Description | 14427
- Options | 14427
- Required Privilege Level | 14427
- Output Fields | 14428
- Sample Output | 14428
- Release Information | 14428

Syntax

```
request app-engine cleanup
<compute-cluster compute-cluster-name>
<compute-cluster compute-cluster-name compute-node compute-node-name>
```

Description

Cleans up temporary files on the Junos V App Engine or the host OS of QFX Series switches. This command deletes all files from the directories where temporary files are normally stored, such as /var/tmp and /var/log.

Be sure to collect any temporary files and logs you want to save before running this command.

Options

- `compute-cluster` *`compute-cluster-name`* `compute-node` *`compute-node-name`*—Name of the compute cluster and name of the compute node. On QFX Series switches, the default names of the compute cluster and compute node are default-cluster and default-node, which applies to the host OS.

Required Privilege Level

view

Output Fields

For a description of the output fields, see Table 109 on page 14428. Output fields are listed in the approximate order in which they appear.

Table 109: request app-engine cleanup Output Fields

Field Name	Field Description
`Compute Cluster`	Name of compute cluster.
`Compute Node`	Name of compute node.
`Cleanup`	Lists the directories that were cleaned up.

Sample Output

request app-engine cleanup (QFX 5100 Switches)

```
user@host> request app-engine cleanup
user@host> request app-engine cleanup
Compute cluster: default-cluster

  Compute node: default-node

   Cleanup (/var/tmp)
   =======
```

Release Information

Command introduced in Junos OS Release 13.1X51-D10.

RELATED DOCUMENTATION

show app-engine crash | 17091

show app-engine logs | 17093

request app-engine file-copy (crash | log) from-jhost to-vjunos | 14429

request app-engine file-copy (crash | log) from-jhost to-vjunos

IN THIS SECTION

- Syntax | 14429
- Description | 14429
- Options | 14430
- Additional Information | 14430
- Required Privilege Level | 14430
- Output Fields | 14430
- Sample Output | 14430
- Release Information | 14431

Syntax

```
request app-engine file-copy (crash | log)
from-jhost host-os-filename
to-vjunos vjunos-filename
```

Description

In Junos OS environments with a host OS, copies core files (`crash` option) or system log files (`log` option) from the host OS to a Junos OS filename.

- `crash` option—Core files are copied from the directory where the host OS normally stores core files. For example, when the host OS is Linux, the source directory is **/var/crash**. The `show app-engine crash` command displays the list of core files available to copy from the host OS source directory.
- `log` option—Log files are copied from the directory where the host OS normally stores system log files. For example, when the host OS is Linux, the source directory is **/var/log**. The `show app-engine logs` command displays the list of system log files available to copy from the host OS source directory.

Either the `crash` or `log` option is required. When the destination Junos OS path for the `to-vjunos` argument is a directory, the destination filename is the source filename by default. To rename the file at the destination, specify a full path that includes a destination filename in the path with the `to-vjunos` argument.

Options

- `from-jhost` *host-os-filename*—Source filename to copy from the host OS.
- `to-vjunos` *vjunos-filename*—Junos OS destination path (with or without a filename) to which the file is copied. If this path is a directory, the source filename is used as the destination filename by default.

Additional Information

This command does not apply to accessing files on guest virtual machines (VMs) on QFX Series devices that support guest VMs.

Required Privilege Level

view

Output Fields

This command produces no output. To see the results of the copy operation, you can run Junos OS commands for viewing files, such as `file list` with the destination directory, `show system core-dumps`, or `show log`.

Sample Output

request app-engine file-copy crash

```
user@host> request app-engine file-copy crash from-jhost
localhost.dcfpe.6449.1454328456.core.tgz to-junos /var/tmp

user@host> show system core-dumps
re0:
--------------------------------------------------------------------------
-rw-r--r--  1 user  test    2538949 Feb 2  19:01 /var/tmp/
localhost.dcpfe.6449.1454328456.core.tgz
total files: 1
```

Release Information

Command introduced in Junos OS Release 13.1X51-D10.

RELATED DOCUMENTATION

show app-engine crash | 17091
show app-engine logs | 17093
request app-engine cleanup | 14427

request auto-configuration reconnect-pending

IN THIS SECTION

- Syntax | 14431
- Description | 14431
- Required Privilege Level | 14432
- Output Fields | 14432
- Sample Output | 14432
- Release Information | 14432

Syntax

```
request auto-configuration reconnect-pending
```

Description

Initiate reestablishment of Layer 2 wholesale sessions that correspond to access lines that are in the pending state. Ordinarily, the most likely situations with pending sessions are handled automatically This statement is intended to be used only when an uncommon condition might prevent automatic reestablishment. This command has no effect when no pending sessions are present.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show ancp neighbor detail`, `show subscribers client-type vlan-oob detail`, and the `show subscribers summary` commands before and after initiating the Port Up message to verify the operation.

Sample Output

request auto-configuration reconnect-pending

```
user@host> request auto-configuration reconnect-pending
```

Release Information

Command introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Reestablishing Pending Access Line Sessions for Layer 2 Wholesale

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

request-cache-timeout (RADIUS Snoop Segment)

IN THIS SECTION

- Syntax | 14433
- Hierarchy Level | 14433
- Description | 14433
- Options | 14433

Required Privilege Level | 14433
Release Information | 14433

Syntax

```
request-cache-timeout timeout;
```

Hierarchy Level

```
[edit access radius snoop-segments segment-name]
```

Description

Configure the length of time to cache the accounting request that was snooped. If the same request is received by the MX Series router within this time, the duplicate request is dropped.

Options

timeout Length of time, in seconds.

- **Range:** 1 through 30

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview

IP-Based Subscriber Setup Overview

request chassis afeb

IN THIS SECTION

- Syntax | 14434
- Description | 14434
- Options | 14434
- Required Privilege Level | 14435
- Output Fields | 14435
- Sample Output | 14435
- Release Information | 14435

Syntax

```
request chassis afeb (offline | online | restart)
```

Description

Control the operation of the compact Forwarding Engine Board (FEB).

Options

offline	Take the FEB offline.
online	Bring the FEB online.
restart	Restart the FEB.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis afeb online (MX104 Router)

```
user@host> request chassis afeb online
AFEB is already online
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

show chassis afeb | 17299

request chassis beacon

IN THIS SECTION

- Syntax (EX Series) | 14436
- Syntax (QFX Series) | 14436
- Syntax (PTX Series) | 14436
- Description | 14437
- Options | 14437
- Required Privilege Level | 14438

- Output Fields | 14438
- Sample Output | 14438
- Release Information | 14439

Syntax (EX Series)

```
request chassis beacon
<fpc slot-number (off | on)>
<fpc slot-number (off | on) timer minutes>
<fpc slot-number pic-slot slot-number port port-number (off | on)>
<fpc slot-number pic-slot slot-number port port-number (off | on) timer minutes>
```

Syntax (QFX Series)

This is the general syntax for the QFX Series of switches. The exact syntax and options are different depending on the specific model of QFX switch and Junos OS release.

```
request chassis beacon
<all (off | on)>
<fpc slot-number (off | on)>
<interconnect-device name (cb slot-number | fpc slot-number | (off | on)>
<node-device name (off | on)>
```

Syntax (PTX Series)

```
request chassis beacon
<all (off | on)>
<fpc slot-number (off | on)>
<interconnect-device name (cb slot-number | fpc slot-number | (off | on)>
<node-device name (off | on)>
```

NOTE: `request chasis beacon` is not supported on the PTX10001-36MR.

Description

(QFX Series switches only) Enable or disable the beacon LED on a QFX Series device.

(EX Series switches only) Supports the following for one or all network port status LEDs on a specified FPC:

- Enable the beacon for:
  - 5 minutes (default)
  - A specified number of minutes (1 through 120)
- Disable the beacon:
  - Instantly
  - After a specified number of minutes (1 through 120) is over

> **NOTE**: To turn the beacon on or off on a Virtual Chassis, you must issue the `request chassis beacon` command on the primary switch in the Virtual Chassis. Also, when specifying the FPC slot number, you must use the target Virtual Chassis member number.

Options

all	Turn the beacon LED either **on** or **off** on all QFabric system Interconnect and Node devices.
cb *slot-number*	Turn the beacon LED either **on** or **off** on the Control Board of the QFX3008-I Interconnect device.
fpc *slot-number*	Turn all beacon LEDs on the specified FPC either **on** or **off**.
fpc *slot-number* **pic-slot** *slot-number* **port** *port-number*	Turn the beacon LED for the specified port either **on** or **off**.
interconnect-device *name*	Turn the beacon LED either **on** or **off** on the Interconnect device.
node-device *name*	Turn the beacon LED either **on** or **off** on the Node device.
off	Turn the specified beacon LEDs **off**.
on	Turn the specified beacon LEDs **on**.

timer ***minutes***

Turn the beacon LEDs **on** for the specified number of minutes or **off** after the specified number of minutes elapses.

- **Range:** 1 through 120

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis beacon fpc on (EX Series and QFX Series Switches)

```
user@switch> request chassis beacon fpc 0 on

Beacon set to ON
```

request chassis beacon fpc on timer (EX Series Switches)

```
user@switch> request chassis beacon fpc 0 on
timer 10
Beacon set to ON
```

request chassis beacon fpc pic-slot port on (EX Series Switches)

```
user@switch> request chassis beacon fpc 0 pic-slot
0 port 1 on
Beacon set to ON
```

request chassis beacon fpc pic-slot port on timer (EX Series Switches)

```
user@switch> request chassis beacon fpc 0 pic-slot
0 port 1 on timer 10
Beacon set to ON
```

request chassis beacon node-device (QFabric System)

```
user@switch> request chassis beacon node-device
node1 on

node1                 ON
```

request chassis beacon on interconnect-device fpc (QFabric System)

```
user@switch> request chassis beacon on interconnect-device
fpc 2

FPC 2                ON
```

request chassis beacon fpc 0 on (PTX Router)

```
user@switch> request chassis beacon fpc 0 on

FPC 0                ON
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

> *How to Locate a Device or Port Using the Chassis Beacon*
>
> show chassis beacon | **17334**

request chassis cb (Security)

IN THIS SECTION

- Syntax | 14440
- Description | 14440
- Options | 14440
- Required Privilege Level | 14440
- Output Fields | 14441
- Sample Output | 14441
- Release Information | 14441

Syntax

```
request chassis cb (offline | online) slot slot-number
```

Description

Control the operation (take the CB offline or bring online) of the Control Board (CB).

Options

offline	Take the Control Board offline.
online	Bring the Control Board online.
slot *slot-number*	Control Board slot number.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis cb (SRX Series)

```
user@host> request chassis cb offline slot 2 node local
node0:
--------------------------------------------------------------------------
Offline initiated, use "show chassis environment cb" to verify
```

Release Information

Command introduced in Junos OS Release 9.2.

Starting with Junos OS Release 12.1X47-D15, the SRX5K-SCBE (SCB2) is introduced and starting with Junos OS Release 15.1X49-D10, the SRX5K-SCB3 (SCB3) with enhanced midplane is introduced.

RELATED DOCUMENTATION

show chassis environment cb (Security) | 17569

request chassis cb

IN THIS SECTION

- Syntax | 14442
- Syntax (TX Matrix Router) | 14442
- Syntax (TX Matrix Plus Router) | 14442
- Syntax (QFabric System) | 14442
- Syntax (EX9253 Switches) | 14443

- Syntax (ACX7900) | 14443
- Output | 14443
- Description | 14443
- Options | 14443
- Required Privilege Level | 14445
- Output Fields | 14445
- Sample Output | 14445
- Usage | 14447
- Release Information | 14447

Syntax

```
request chassis cb (offline | online) slot slot-number
```

Syntax (TX Matrix Router)

```
request chassis cb (offline | online) <slot slot-number | lcc number slot cb-slot-number | scc
number slot cb-slot-number>
```

Syntax (TX Matrix Plus Router)

```
request chassis cb (offline | online) <slot slot-number | lcc number slot cb-slot-number | sfc
number slot cb-slot-number>
```

Syntax (QFabric System)

```
request chassis cb  (offline | online) interconnect-device name slot slot-number
<interconnect-device name slot slot-number (offline | online)>
```

Syntax (EX9253 Switches)

```
request chassis cb  (offline | online) name slot slot-number
```

Syntax (ACX7900)

```
request chassis cb
<(offline | online ) <slot cb-slot-number>
<slot slot-number (offline | online)>
```

request chassis cb (PTX10008 Router)

```
user@host> request chassis cb offline slot 1
Offline initiated, use "show chassis environment cb" to verify
```

Output

The following error message is displayed if user tries to offine a slot when the related node is online.

```
root@re0-re0> request chassis cb slot 1 offline

Error: CB offline is not allowed, when node is in online state.
Try to offline the node using "request node offline [node]"
```

Description

(M120, M320, and MX Series routers and T Series routers, QFabric systems, and EX8200 switches only) Control the operation of the Control Board (CB).

Options

offline Take the Control Board offline.

> **NOTE**: On a QFabric system, to bring the backup Control Board on a QFX3008-I Interconnect device offline, issue the `request chassis cb slot backup-slot-number offline` command.

> **NOTE**: Only backup Control Board can be turned offline or online. To turn a Control Board offline or to bring it back online, the Routing Engine should be turned offline first.

online

Bring the Control Board online.

interconnect-device *name*

(QFabric systems only) (Optional) Bring the QFX3008-I Interconnect device Control Board either offline or online:

slot *slot-number*

The following options can be configured for CB slots:

- offline —Take the CB slot offline.
- online — Bring the CB slot online.

Control Board slot number:

- (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, if you specify the number of the T640 router by using the lcc *number* option (the recommended method), replace *cb-slot-number* with a value from 0 through 1.

  Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 or T4000 router by using the lcc *number* option (the recommended method), replace *cb-slot-number* with a value from 0 through 1.
- M320 router—Replace *slot-number* with a value from 0 through 1.
- MX480/MX240 routers—Replace *slot-number* with a value from 0 through 1.
- MX960 router—Replace *slot-number* with a value from 0 through 2.
- MX2020, MX2010, and MX2008 routers—Replace *slot-number* with 0 or 1.
- EX8208 switch—Replace *slot-number* with a value from 0 through 2.
- EX8216 switch—Replace *slot-number* with a value from 0 through 1.
- QFabric System—Replace *slot-number* with a value from 0 through 1.

lcc *number*

(TX Matrix, TX Matrix Plus routers only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

sfc *number* (TX Matrix Plus routers only) (Optional) Change the CB status for the TX Matrix Plus router (switch-fabric chassis). Replace *number* with 0.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis cb

```
user@host> request chassis cb offline slot 1

Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb interconnect-device (QFabric System)

```
user@switch> request chassis cb interconnect-device
interconnect1 offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb (MX2020 Router)

```
user@host> request chassis cb offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb (MX2010 Router)

```
user@host> request chassis cb offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb (MX2008 Router)

```
user@host>request chassis cb offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

request chassis cb (MX10003 Router)

```
user@host>request chassis cb online slot 1
CB 1 appears to be online already
```

request chassis cb (EX9253 Switch)

```
user@switch>request chassis cb offline slot 1
Offline initiated, use "show chassis environment cb" to verify
```

request chassis cb (ACX7509)

```
user@host>request chassis cb offline slot 1
Offline initiated
```

```
user@host>request chassis cb online slot 1
Online initiated
```

Usage

When `request chassis cb slot 1 offline`, is configured before the node is made offline, the following error `Error: CB offline is not allowed, when the node is online state. Try to offline the node using "request node offline` *node* is displayed.

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Switching Control Board Redundancy

request chassis ccg

IN THIS SECTION

- Syntax | 14447
- Description | 14448
- Options | 14448
- Required Privilege Level | 14448
- Output Fields | 14448
- Sample Output | 14448
- Release Information | 14448

Syntax

```
request chassis ccg (offline | online) slot slot-number
```

Description

(PTX5000 Packet Transport Routers) Control the operation of the Centralized Clock Generator (CCG).

Options

offline	Take the CCG offline.
online	Bring the CCG online.
slot *slot-number*	CCG slot number. Replace *slot-number* with a value from 0 through 1.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis ccg

```
user@host> request chassis ccg offline slot 1
CCG 1 is  offline, Backup CCG 0 is now online.
```

Release Information

Command introduced in Junos OS 12.1x48 .

RELATED DOCUMENTATION

Clock Sources for PTX Series Packet Transport Routers

show chassis environment ccg | **17613**

request chassis cfeb

IN THIS SECTION

- Syntax | 14449
- Description | 14449
- Options | 14449
- Required Privilege Level | 14449
- Output Fields | 14450
- Sample Output | 14450
- Release Information | 14450

Syntax

```
request chassis cfeb (offline | online | restart)
```

Description

(M7i and M10i routers only) Control the operation of the Compact Forwarding Engine Board (CFEB).

Options

offline	Take the CFEB offline.
online	Bring the CFEB online.
restart	Restart the CFEB.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis cfeb

```
user@host> request chassis cfeb offline
CFEB Offlined
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show chassis cfeb | 17337

Configuring CFEB Redundancy on the M10i Router

CFEB Overview

request chassis clock master switch

IN THIS SECTION

- Syntax | 14451
- Description | 14451
- Options | 14451
- Additional Information | 14451
- Required Privilege Level | 14451
- Output Fields | 14451
- Sample Output | 14452
- Release Information | 14452

Syntax

```
request chassis clock master switch
```

Description

(PTX Series Packet Transport Routers only) Control which Centralized Clock Generator (CCG) is the primary.

Options

This command has no options.

Additional Information

By default, the CCG in slot 0 (CCG0) is the primary and the CCG in slot 1 (CCG1) is the backup. If you use this command to change the primary, and then restart the chassis software for any reason, the primary reverts to the default setting. To change the default primary CCG, include the `ccg` statement at the [**edit chassis redundancy**] hierarchy level in the configuration. For more information, see the Junos OS Administration Library.

The configurations on the two CCGs do not have to be the same, and they are not automatically synchronized. If you configure both CCGs as primaries, when the chassis software restarts for any reason, the CCG in slot 0 becomes the primary and the one in slot 1 becomes the backup.

The switchover from the primary CCG to the backup CCG is immediate.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis clock primary switch

```
user@host> request chassis clock master switch
CCG master switch initiated, use “show chassis environment ccg” to verify
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

show chassis environment | 17429

request chassis cluster configuration-synchronize

IN THIS SECTION

- Syntax | 14452
- Description | 14453
- Required Privilege Level | 14453
- Output Fields | 14453
- Sample Output | 14453
- Release Information | 14453

Syntax

```
request chassis cluster configuration-synchronize
```

Description

On the secondary node, you can synchronize the configuration from the primary node to the secondary node when the secondary node joins the primary node in a cluster.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis cluster configuration-synchronize

```
{secondary:node1}
user@host> request chassis cluster configuration-synchronize
Performing configuration synchronization from remote node.
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Understanding Automatic Chassis Cluster Synchronization Between Primary and Secondary Nodes

Verifying Chassis Cluster Configuration Synchronization Status

NTP Time Synchronization on SRX Series Devices

request chassis cluster failover node

IN THIS SECTION

- Syntax | 14454
- Description | 14454
- Options | 14454
- Required Privilege Level | 14455
- Output Fields | 14455
- Sample Output | 14455
- Release Information | 14455

Syntax

```
request chassis cluster failover node          node-number
redundancy-group          group-number
```

Description

For chassis cluster configurations, initiate manual failover in a redundancy group from one node to the other, which becomes the primary node, and automatically reset the priority of the group to 255. The failover stays in effect until the new primary node becomes unavailable, the threshold of the redundancy group reaches 0, or you use the `request chassis cluster failover reset` command.

After a manual failover, you must use the `request chassis cluster failover reset` command before initiating another failover.

Options

- `node` *`node-number`*—Number of the chassis cluster node to which the redundancy group fails over.
- **Range:** 0 through 1
- `redundancy-group` *`group-number`*—Number of the redundancy group on which to initiate manual failover. Redundancy group 0 is a special group consisting of the two Routing Engines in the chassis cluster.

Range: 0 through 255

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis cluster failover node

```
user@host> request chassis cluster failover node 0 redundancy-group 1
Initiated manual failover for redundancy group 1
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

clear chassis cluster failover-count | 3458
request chassis cluster failover reset | 14457
show chassis cluster status | 17405

request chassis cluster failover redundancy-group

IN THIS SECTION

- Syntax | 14456
- Description | 14456
- Options | 14456

- Required Privilege Level | 14456
- Output Fields | 14456
- Sample Output | 14457
- Release Information | 14457

Syntax

```
request chassis cluster failover node node-number redundancy-group redundancy-group-number
```

Description

For chassis cluster configurations, initiate manual failover in a redundancy group from one node to the other, which becomes the primary node, and automatically reset the priority of the group to 255. The failover stays in effect until the new primary node becomes unavailable, the threshold of the redundancy group reaches 0, or you use the `request chassis cluster failover reset` command.

After a manual failover, you must use the `request chassis cluster failover reset` command before initiating another failover.

Options

- `node` *`node-number`*—Number of the chassis cluster node to which the redundancy group fails over.
- **Range:** 0 or 1
- `redundancy-group` *`group-number`*—Number of the redundancy group on which to initiate manual failover. Redundancy group 0 is a special group consisting of the two Routing Engines in the chassis cluster.

  Range: 0 through 255

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis cluster failover redundancy-group

```
user@host> request chassis cluster failover redundancy-group 0 node 1
{primary:node0}
user@host> request chassis cluster failover redundancy-group 0 node 1
- - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
Initiated manual failover for redundancy group 0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Initiating a Chassis Cluster Manual Redundancy Group Failover

Verifying Chassis Cluster Failover Status

request chassis cluster failover reset

IN THIS SECTION

- Syntax | 14458
- Description | 14458
- Options | 14458
- Required Privilege Level | 14458
- Output Fields | 14458
- Sample Output | 14458
- Release Information | 14458

Syntax

```
request chassis cluster failover reset
redundancy-group           group-number
```

Description

In chassis cluster configurations, undo the previous manual failover and return the redundancy group to its original settings.

Options

`redundancy-group` *`group-number`*—Number of the redundancy group on which to reset manual failover. Redundancy group 0 is a special group consisting of the two Routing Engines in the chassis cluster.

Range: 0 through 255

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis cluster failover reset

```
user@host> request chassis cluster failover reset redundancy-group 0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

clear chassis cluster failover-count | 3458
request chassis cluster failover node | 14454
show chassis cluster status | 17405

request chassis cluster in-service-upgrade abort (ISSU)

IN THIS SECTION

- Syntax | 14459
- Description | 14459
- Options | 14459
- Required Privilege Level | 14460
- Output Fields | 14460
- Sample Output | 14460
- Release Information | 14460

Syntax

```
request chassis cluster in-service-upgrade abort
```

Description

Terminate an upgrade in a chassis cluster during an in-service software upgrade (ISSU). Use this command to end the ISSU on any nodes in a chassis cluster followed by `reboot` to terminate the ISSU on that device.

Options

This command has no options.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis cluster in-service-upgrade terminate

```
user@host> request chassis cluster in-service-upgrade abort
Exiting in-service-upgrade window
Chassis ISSU Aborted
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Upgrading a Chassis Cluster Using In-Service Software Upgrade

request chassis device-mode

IN THIS SECTION

- Syntax | 14461
- Description | 14461
- Options | 14461
- Required Privilege Level | 14462
- Output Fields | 14462

Sample Output | 14462
Release Information | 14463

Syntax

```
request chassis device-mode (interconnect-device | node-device | standalone)
```

Description

Select the operating mode for the device, which acts either as a device within a QFabric system or as a standalone switch.

NOTE:

- Issue the `request chassis device-mode` command only when your management station is connected directly to the device over a console port connection.
- Changing the device mode erases all configuration data on the device. When you convert a device to a different device mode, we recommend that you back up your device configuration to an external location before issuing the `request chassis device-mode` command.

Options

interconnect-device Set the device to operate as an Interconnect device within a QFabric system. To complete the Interconnect device mode conversion process, you must connect the device to the QFabric system management control plane and reboot the device.

node-device Set the device to operate as a Node device within a QFabric system. To complete the Node device mode conversion process, you must connect the device to the QFabric system management control plane and reboot the device.

standalone Set the device to operate as a standalone switch. If the device starts in Node device or Interconnect device mode, you must reboot the device to return to standalone mode. Standalone mode is the factory default setting.

Required Privilege Level

admin

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis device-mode interconnect-device (Starting in Node Device or Standalone Mode)

```
user@switch> request chassis device-mode interconnect-device
Device mode set to 'interconnect-device' mode.
Please reboot the system to complete the process.
```

request chassis device-mode node-device (Starting in Interconnect Device or Standalone Mode)

```
user@switch> request chassis device-mode node-device
Device mode set to 'node-device' mode.
Please reboot the system to complete the process.
```

request chassis device-mode standalone (Starting in Interconnect Device or Node Device Mode)

```
user@switch> request chassis device-mode standalone
Device mode set to 'standalone' mode.
Please reboot the system to complete the process.
```

request chassis device-mode standalone (Starting in Standalone Mode)

```
user@switch> request chassis device-mode standalone
Device mode set to 'standalone' mode.
No reboot required.
```

Release Information

Command introduced in Junos OS Release 11.2.

interconnect-device option introduced in Junos OS Release 13.1 for the QFX Series.

RELATED DOCUMENTATION

- *Converting the Device Mode for a QFabric System Component*
- show chassis device-mode | **17419**
- *Understanding Interconnect Devices*
- *Understanding Node Devices*
- *Understanding the QFabric System Hardware Architecture*

request chassis device-mode satellite

IN THIS SECTION

- Syntax | **14464**
- Description | **14464**
- Options | **14464**
- Required Privilege Level | **14464**
- Sample Output | **14464**
- Release Information | **14464**

Syntax

```
request chassis device-mode satellite package-name
```

Description

Manually install satellite software onto a switch before interconnecting the switch as a satellite device into a Junos Fusion.

There are other methods of installing satellite software onto a satellite device, and each Junos Fusion has individual requirements for manually installing satellite software. See *Configuring Junos Fusion Provider Edge* or Configuring or Expanding a Junos Fusion Enterprise before manually installing satellite software.

This command is entered from a standalone device before it is configured into a satellite device in a Junos Fusion.

Options

package-name	The URL to the satellite software package.

Required Privilege Level

system-control

Sample Output

request chassis device-mode satellite /var/tmp/satellite-3.0R1.1-signed.tgz

```
user@satellite-device> request chassis device-mode satellite /var/tmp/satellite-3.0R1.1-
signed.tgz
```

Release Information

Command introduced in Junos OS Release 14.1X53-D16.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

request chassis fabric fpc

IN THIS SECTION

- Syntax | 14465
- Description | 14465
- Options | 14465
- Required Privilege Level | 14466
- Output Fields | 14466
- Sample Output | 14466
- Release Information | 14466

Syntax

```
request chassis fabric fpc interconnect-device interconnect-device-name slot slot-number
(offline | online)
```

Description

(QFabric systems only) Set the Interconnect device Flexible PIC Concentrator (FPC) offline or online for the QFabric system. When the FPC is offline, traffic is redirected to other FPCs and is not lost while you remove or install an FPC. After issuing this command, you must issue the `request chassis fpc` command.

Options

interconnect-device *interconnect-device-name*	Set the Interconnect device containing the FPC you want to bring either offline or online.

slot ***slot-number***	Set the specific FPC slot on the Interconnect device.
offline	Set the Interconnect device FPC to offline for removal.
online	Set the Interconnect device FPC to online after installation.

Required Privilege Level

admin

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis fabric fpc online

```
user@qfabric> request chassis fabric fpc interconnect-device IC-YW3781 offline slot 15
Graceful offline of the fabric card has been initiated. Please wait 20 seconds before offlining
or removing the card.
```

request chassis fabric fpc offline

```
user@qfabric> request chassis fabric fpc interconnect-device IC-YW3781 online slot 15
Bring the FPC online by issuing the "request chassis fpc online" command.
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

- request chassis fpc

request chassis fabric pfe

IN THIS SECTION

- Syntax | 14467
- Description | 14467
- Options | 14467
- Required Privilege Level | 14468
- Release Information | 14468

Syntax

```
request chassis fabric pfe pfe-number fpc fpc-number offline
```

Description

Make a specified fabric of the packet forwarding engine (PFE) offline.

This command makes only the specified PFE of the FPC offline and rest of the PFEs of the FPC are not affected. If a PFE of an FPC is affected because of fabric path wedge errors, the affected PFE is disabled and the associated fabric goes offline as part of fabric hardening actions. The output of the `show chassis fabric fpcs` and `show chassis fabric plane` commands show a new state for the PFE as `Fabric Disabled`.

Fabric stream wedge occurs when the ASIC of the FPC is in the stuck state, and the ingress PFE fails to send traffic to the destination PFE. You can use the request `chassis fabric pfe` *`pfe-number`* `fpc` *`fpc-number`* `offline` command to make any PFE offline.

> **NOTE**: This statement does not have an option to bring the PFE back online. You must restart the FPC to bring the PFE back online.

Options

pfe-number [0–3]

fpc-number [0–11] or [0–5] or [0–19] or [0–2], depending on the type of the chassis.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

show chassis fabric plane | 18113

show chassis fabric fpcs | 18079

request chassis fabric plane

IN THIS SECTION

- Syntax | 14468
- Description | 14469
- Options | 14469
- Required Privilege Level | 14470
- Output Fields | 14470
- Sample Output | 14470
- Release Information | 14471

Syntax

```
request chassis fabric plane plane-number  (offline | online)
```

Description

(M120 and MX Series routers and EX8200 switches only) Control the operation of the specified fabric plane.

On an MX480 or MX240 series router, you can configure the active control board for redundancy mode or increased bandwidth mode. When running in increased bandwidth mode, MX series routers with Trio chips and the MPC3E will use all eight active fabric planes.

To take both plane 0 and plane 1 offline on a MX480 and MX240 series routers with one or more MPC4E MICs installed, a X86 Media Service Blade, and/or 100G PFE, and where redundancy-mode is configured for "increased-bandwidth", Juniper recommends taking plane 1 offline before plane 0. Likewise, when the router is configured for increased-bandwidth mode, taking fabric planes 0, 2, 4, and 6 offline can cause the chassis to run in a reduced fabric bandwidth mode. Plane 7 may remain in a "spare" state (as seen in the "show chassis fabric summary" command output) until plane 3 is taken offline and then brought back up.

Options

offline

Take the fabric plane offline. Use the `request chassis fabric plane` *`plane-number`* `offline` command to clear a FAULT state on a fabric plane. To bring the fabric plane back online, use the `request chassis fabric plane` *`plane-number`* `online` command.

online

Bring the fabric plane online.

plane ***plane-number***

Fabric plane number.

- For the M120 router, replace *plane-number* with a value from 0 through 3.
- For the MX480 and MX240 routers, replace *plane-number* with a value from 0 through 7.
- For the MX2020, MX2010, and MX2008 routers, replace *plane-number* with a value from 0 through 7.
- For the MX960, MX10004, and MX10008 routers, replace *plane-number* with a value from 0 through 11.
- For the EX8208 switch, replace *plane-number* with a value from 0 through 11.
- For the EX8216 switch, replace *plane-number* with a value from 0 through 7.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis fabric plane 0 online

```
user@host> request chassis fabric plane 0 online
Online initiated, use “show chassis fabric plane” to verify
```

request chassis fabric plane 0 offline

```
user@host> request chassis fabric plane 0 offline
Offline initiated, use “show chassis fabric plane” to verify
```

request chassis fabric plane 0 online (EX8200 switch)

```
user@host> request chassis fabric plane 0 online

Plane 0 is already active
```

request chassis fabric plane (MX2020 Router)

```
user@host> request chassis fabric plane 2 online
Plane 2 is already active
```

request chassis fabric plane (MX2010 Router)

```
user@host> request chassis fabric plane 4 online
Plane 4 is already active
```

request chassis fabric plane (MX2008 Router)

```
user@host>request chassis fabric plane 4 online
Plane 4 is already active
```

request chassis fabric plane (MX10003 and MX10004 Router)

```
user@host>request chassis fabric plane 4 online
Plane 4 is already active
```

request chassis fabric plane (EX9253 Switch)

```
user@switch>request chassis fabric plane 0 online
Plane 0 is already active
```

Release Information

Command introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

show chassis fabric plane | 18113

show chassis fabric plane-location | 18156

show chassis fabric summary | 18212

request chassis feb

IN THIS SECTION

- Syntax | 14472
- Syntax (ACX7509 Router) | 14472
- Description | 14472
- Options | 14473
- Required Privilege Level | 14473
- Output Fields | 14473
- Sample Output | 14473
- Release Information | 14474

Syntax

```
request chassis feb slot slot-number (offline | online | restart)
```

Syntax (ACX7509 Router)

```
request chassis feb slot slot-number (offline | online | restart) <force>
```

Description

Offline, online, restart the specified FEB. Upon restart, the FEB goes offline and online within the stipulated time. FEB traffic is forwarded before power off. Traffic is recovered after FEB is online.

On ACX7509 routers, when a primary FEB is made offline or restarted, the switchover to backup FEB does not take place. This results in traffic loss as none of the FEB is forwarding the traffic. If you make the primary FEB offline or restart it, a traffic loss warning message is displayed and the request is rejected. You can use the `force` option to forcefully offline or restart the primary FEB.

Options

slot ***slot-number***	FEB slot number. Replace ***slot-number*** with the FEB slot number value of your platform.
offline	Bring the specified FEB offline.
online	Bring the specified FEB online.
restart	Restart the specified FEB.
force	Force execute the offline/restart command. Only available for ACX7509 routers.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis feb (ACX1000 Router)

```
user@host> request chassis feb slot 0 restart
FEB will be restarted NOW.
```

request chassis feb (ACX7509 Router) (Primary FEB)

NOTE: Slot 1 is the primary FEB for the example commands provided below.

```
user@host> request chassis feb restart slot 1
warning: RCB and FEB work in the paired slot mode. FEB 1 offline/restart will result in traffic
loss and does not cause a switchover. Please re-try after initiating a mastership switchover
```

```
using 'request chassis routing-engine master switch' CLI. If offline/restart is still intended,
use 'force' option in addition to this CLI.
```

```
user@host> request chassis feb offline slot 1
warning: RCB and FEB work in the paired slot mode. FEB 1 offline/restart will result in traffic
loss and does not cause a switchover. Please re-try after initiating a mastership switchover
using 'request chassis routing-engine master switch' CLI. If offline/restart is still intended,
use 'force' option in addition to this CLI.
```

If offline/restart is still intended for primary FEB, use `force` option in addition to the command. For example:

```
user@host> request chassis feb offline slot 1 force
Offline initiated, use "show chassis feb" to verify
```

request chassis feb (QFX5700 Switch)

```
user@host> request chassis feb slot 0 online
Online initiated, use "show chassis feb" to verify
```

Release Information

Command introduced in Junos OS Release 8.0.

Command support introduced for ACX7509 devices in Junos OS Evolved Release 21.4R1.

Warning message introduced for ACX7509 devices in Junos OS Evolved Release 23.2R2.

RELATED DOCUMENTATION

- request chassis feb | **14472**
- show chassis feb | **18278**
- show chassis fabric feb | **18078**
- show chassis fpc-feb-connectivity | **18397**
- *Understanding Switching Control Board Redundancy*

request chassis fpc

IN THIS SECTION

- Syntax | 14475
- Syntax (SRX Series Firewalls) | 14475
- Syntax (TX Matrix and TX Matrix Plus Routers) | 14476
- Syntax (MX Series Routers) | 14476
- Syntax (MX10004) | 14476
- Syntax (QFabric Systems) | 14476
- Syntax (PTX Series Devices)) | 14476
- Syntax (ACX7900, ACX7509, ACX7332, and ACX7348) | 14477
- Syntax (ACX7100 and ACX7024) | 14477
- Description | 14477
- Options | 14479
- Required Privilege Level | 14482
- Output Fields | 14482
- Sample Output | 14483
- Release Information | 14486

Syntax

```
request chassis fpc slot slot-number (offline | online | restart)
<pfe-instance pfe-id> (offline | online | restart)
```

Syntax (SRX Series Firewalls)

```
request chassis fpc
<node (node-id | local | primary)>
<offline | online | restart>
<slot slot-number>
```

Syntax (TX Matrix and TX Matrix Plus Routers)

```
request chassis fpc (offline | online | restart)
<slot slot-number>
<lcc number>
```

Syntax (MX Series Routers)

```
request chassis fpc (offline | online | restart)
<slot slot-number (offline | online | restart)>
<pfe-instance pfe-id (offline | online | restart)>
<all-members>
<local>
<member member-id>
<tag>
```

Syntax (MX10004)

```
request chassis fpc (offline | online | restart)
<slot slot-number (offline | online | restart)>
<pfe-instance pfe-id (offline | online | restart)>
<tag>
```

Syntax (QFabric Systems)

```
request chassis fpc (offline | online | restart)
<interconnect-device name slot slot-number (offline | online | restart)>
<(offline | online) interconnect-device name slot slot-number>
<slot slot-number interconnect-device name (offline | online)>
<pfe-instance pfe-id (offline | online | restart)>
```

Syntax (PTX Series Devices))

```
request chassis fpc slot slot-number (offline | online | restart)
<pfe-instance pfe-id> (offline | online | restart)
```

```
<pfe pfe-id> (offline | online | restart)
<force>
<tag>
```

Syntax (ACX7900, ACX7509, ACX7332, and ACX7348)

```
request chassis fpc
< slot slot-number
<pfe-instance pfe-id> (offline | online | restart)
<tag>
```

Syntax (ACX7100 and ACX7024)

```
request chassis fpc
< slot slot-number
<tag>
```

Description

(M20, M40, M40e, M120, M160, M320, MX Series, and TSeries routers, QFabric systems, EX Series switches, and PTX Series Devices only) Control the operation of the Flexible PIC Concentrator (FPC). An error message is displayed when restart, offline, or online is configured using `request chassis fpc slot` command on an FPC with empty slot.

Control the operation of the Flexible PIC Concentrator (FPC).

NOTE: The SRX5K-SPC-2-10-40 (SPC1), SRX5K-SPC-4-15-320 (SPC2), and SRX5K-SPC3 does not support the `request chassis fpc` command.

NOTE:

- Starting with Junos OS Release 12.3, it is possible that FPCs brought offline by using the `request chassis fpc slot` *`fpc-slot`* `offline` operational-mode CLI command can come online during a configuration commit or power-supply replacement procedure. As an

alternative, use the `set fpc fpc-slot power off` configuration-mode command at the `[edit chassis]` hierarchy level to ensure that the FPCs remain offline.

- Starting in Junos OS Evolved Release 19.1R1, PTX10003-80C and PTX10003-160C devices do not support the `request chassis fpc slot <slot-number> online` command. The only way to bring up an FPC (MPC) that is offline is by rebooting the chassis. So, starting in Junos OS Evolved Release 21.1R1, when you take an FPC offline by using the `request chassis fpc slot <slot-number> offline` command, the screen displays the following message.

  ```
  Warning : FPC <slot> cannot be made online using a CLI command. You need to perform router
  reboot using "request system reboot" to online the FPC <slot>. Do you wish to continue ?
  [yes,no] (no).
  ```

  Starting in Junos OS Evolved Release 22.4R1, when you perform FPC restart by using the `request chassis fpc slot <slot-number> offline` command, the screen displays the following message.

On Junos OS Evolved, an FPC does not restart when you enter a commit command that configures an element of that FPC.

NOTE: In releases earlier than Junos OS Release 15.1F3 and Junos OS Release 16.1, offline FPCs in the PTX5000 router might be powered on by the router during a reboot, or when triggered by other power management events on the router, such as when you take another FPC offline.

Starting with Junos OS Release 15.1F3 and Junos OS Release 16.1, offline FPCs do not come online during reboots or other power management events. To bring such an FPC online:

1. Delete the `fpc fpc-slot power off` statement from the `[edit chassis]` hierarchy level, if that statement is configured, and commit the configuration.

2. Either issue the `request chassis fpc online slot fpc-slot` operational-mode CLI command or press and hold the FPC **ONLINE/OFFLINE** button for about 5 seconds until the green **OK** LED next to the button lights steadily.

NOTE: If a CLI-based firmware upgrade is in progress, the specified FPC does not restart. Starting with Junos OS Release 15.1, the following message is displayed when this occurs:

```
user@host> request chassis fpc slot 0 restart
FPC 0 Firmware update in progress. Wait!!!
```

NOTE: The command `request chassis fpc (offline | online) slot slot-number` is not supported on PTX1000 router. Whereas, `request chassis fpc restart slot slot-number` is supported on PTX1000 router.

NOTE: On EX9204, EX9208, EX9214, and EX9253 switches, when a line card is brought online, if the aggregate interface is initialized before the child interface is marked as part of the aggregate interface, there might be a loss of traffic from the aggregate interface for up to 30 seconds and the CPU usage of the line card installed on the switch might go up to 100%.

Options

node Change Flexible PIC Concentrator status of specific node.

offline Take the FPC offline.

online Bring the FPC online.

interconnect-device *name* (QFabric systems only) Bring the FPC on the QFX3008-I Interconnect device either offline or online:

- (QFabric System) On a QFabric system, specify the name of the QFX3008-I Interconnect device containing the FPC you want to bring either offline or online.

restart Restart the FPC.

force Restart an unresponsive FPC on PTX platforms. The `<force>` option can be used before or after the `<(offline | online | restart)>` options

slot *slot-number* The following options can be configured for FPC slots:

- offline —Take the FPC slot offline.
- online — Bring the FPC slot online.
- restart — Restart the FPC slot.

FPC slot number:

- M20 router—0 through 3.
- M120 router—0 through 5.
- MX240 router—0 through 2. On the MX240 router, slot-number corresponds to the (DPC slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX480 router—0 through 5. On the MX480 router, slot-number corresponds to the DPC slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX960 router—0 through 11. On the MX960 router, slot-number corresponds to the DPC slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX2020 router—0 through 19.
- MX2010 router—0 through 9.
- MX2008 router—0 through 9.
- MX10004 router—0 through 3.

  MX10008 router—0 through 7.

  MX10016 router—0 through 15.
- TX Matrix and TX Matrix Plus routers only—On the TX Matrix router, if you specify the number of the T640 router by using the lcc *number* option (the recommended method), replace *slot-number* with a value from 0 through 7. Otherwise, replace *slot-number* with a value from 0 through 31.

  Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 or T4000 router by using the lcc *number* option (the recommended method), replace *slot-number* with a value from 0 through 7. Otherwise, replace *slot-number* with a value from 0 through 31. In case of TX Matrix Plus router with 3D SIBs, replace *slot-number* with a value from `0` through `63`. For example, the following commands have the same result:

```
user@host> request chassis fpc lcc 1 slot 1 offline
user@host> request chassis fpc slot 9 offline
```

- Other routers—0 through 7.
- QFabric System—Replace *slot-number* with a value from 0 through 2.
- EX Series switches:
  - EX4200 switches in a Virtual Chassis configuration—Replace *slot-number* with a value from 0 through 9.
  - EX6210 switches—Replace *slot-number* with a value from 0 through 9.

    NOTE: These commands are not supported for slots 4 and 5 when a Switch Fabric and Routing Engine (SRE) module is installed in those slots. These commands are supported for slots 4 and 5 only if a line card is installed in them.

  - EX8208 switches—Replace *slot-number* with a value from 0 through 7.
  - EX8216 switches—Replace *slot-number* with a value from 0 through 15.
  - EX9204 switches—Replace *slot-number* with a value from 0 through 2.
  - EX9208 switches—Replace *slot-number* with a value from 0 through 5.
  - EX9214 switches—Replace *slot-number* with a value from 0 through 11.
- PTX5000 Packet Transport Router—Replace *slot-number* with a value from 0 through 7.

tag

This option is used for upgrading the latest firmware for components Mezz, I2CS, CPLD or Boot CPLD on the FPC.

pfe-instance *pfe-number*

PFE instance number for the FPC slot. For PTX10008 device, it is 0 through 4.
The following options can be configured for PFEs:

- offline —Take the PFE instance on the FPC offline.
- online — Bring the PFE instance on the FPC online.
- restart — Restart the PFE instance on the FPC.

pfe *pfe-id*

(PTX10002-36QDD routers only) PFE number for the FPC slot. Range: 0 through 3.
The following options can be configured for PFEs:

- offline —Take the PFE instance on the FPC offline.

- online — Bring the PFE instance on the FPC online.
- restart — Restart the PFE instance on the FPC.

all-members (MX Series routers only) (Optional) Change FPC status of all members of the Virtual Chassis configuration.

local (MX Series routers only) (Optional) Change FPC status of the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Change FPC status of the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

lcc *number* (TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis fpc (SRX Series)

```
user@host> request chassis fpc online slot 0
FPC 0 already online
```

request chassis fpc

```
user@host> request chassis fpc online slot 0

Online initiated, use "show chassis fpc" to verify
```

request chassis fpc (MX Series Routers with Media Services Blade [MSB])

```
user@host> request chassis fpc slot 0
Possible completions:
  offline              Take FPC offline
  online               Bring FPC online
  restart              Restart FPC
```

request chassis fpc (MX2020 Router)

```
user@host> request chassis fpc online slot 2
Online initiated, use "show chassis fpc" to verify
```

request chassis fpc (MX2010 Router)

```
user@host> request chassis fpc offline slot 5
Offline initiated, use "show chassis fpc" to verify
```

request chassis fpc (MX2008 Router)

```
user@host> request chassis fpc online slot 5
Online initiated, use "show chassis fpc" to verify
```

request chassis fpc (MX10003, MX10004, and MX10008)

```
user@host> request chassis fpc online slot 1
Online initiated, use "show chassis fpc" to verify
```

request chassis fpc (MX204 Router)

```
user@host> request chassis fpc online slot 0
Online initiated, use "show chassis fpc" to verify
```

request chassis fpc (PTX10008 Router)

```
user@host> request chassis fpc online slot 1
Online initiated, use "show chassis fpc" to verify
```

```
user@host> request chassis fpc slot 0 pfe-instance 0 ?
Possible completions:
  offline              Take FPC offline
  online               Bring FPC online
  restart              Restart FPC
```

request chassis fpc (EX9200 Switch)

```
user@host> request chassis fpc slot 0
Possible completions:
  offline              Take FPC offline
  online               Bring FPC online
  restart              Restart FPC
```

request chassis fpc (EX9251 Switch)

```
user@switch> request chassis fpc online slot
0
Online initiated, use "show chassis fpc" to verify
```

request chassis fpc (EX9253 Switch)

```
user@switch> request chassis online fpc slot
0
Online initiated, use "show chassis fpc" to verify
```

request chassis fpc (PTX10002-36QDD)

```
user@host> request chassis fpc slot 0 pfe 0 offline
Warning: pfe 1 will also be offlined. Do you wish to continue? [yes,no]
```

NOTE: On a PTX10002-36QDD router, you must commit command for both PFEs.

request chassis fpc (ACX7509, ACX7332, and ACX7348 Routers)

```
user@host> request chassis fpc slot 1 offline

Offline initiated, use "show chassis fpc" to verify
```

```
user@host> request chassis fpc slot 1 online

Online initiated, use "show chassis fpc" to verify
```

```
user@host> request chassis fpc slot 1 restart
Fpc 1 restart initiated!
```

Release Information

Command introduced before Junos OS Release 7.4.

Command support added to MX10004 device in Junos OS Release 22.3R1.

Command introduced on SRX Series Firewalls in Junos OS Release 17.2.

RELATED DOCUMENTATION

- MX960 Flexible PIC Concentrator Description
- show chassis fpc (View)

request chassis fpc-control-port

IN THIS SECTION

- Syntax | 14486
- Description | 14486
- Options | 14487
- Required Privilege Level | 14487
- Output Fields | 14487
- Sample Output | 14487
- Release Information | 14488

Syntax

```
request chassis fpc-control-port (disable | enable) fpc FPC slot number port Port number
```

Description

You can disable FPC HA control port 0 on primary node and enable FPC HA control port 0 on secondary node.

Options

disable Disable HA control port on the local node.

enable Enable HA control port on the local node.

fpc Specify the fpc slot number.

Range: 0 through 23

port Specify the port number.

Range: 0 through 1

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis fpc-control-port

```
user@host> request chassis fpc-control-port disable fpc 0 port 0
```

```
fpc 0 HA port 0 disabled
```

```
user@host> request chassis fpc-control-port enable fpc 0 port 0
```

```
fpc 0 HA port 0 enabled
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Chassis Cluster Dual Control Links

request chassis fpc slc

IN THIS SECTION

- Syntax | 14488
- Description | 14488
- Options | 14489
- Required Privilege Level | 14489
- Release Information | 14489

Syntax

`request chassis fpc` *`fpc-slot`* `slc` *`slc-id`* `(offline | restart | online)`

Description

(Applies to the base system or BSYS). Start, stop or restart the sub line cards (SLCs).

NOTE: On BSYS, entering the `request chassis fpc` *`fpc-slot`* `offline/restart/online` command initiates the offline/online/restart process of all the SLCs (sub line cards) of the selected line card (FPC). On GNF, entering the `request chassis fpc` *`fpc-slot`* `offline/restart/online` command initiates the offline/online/restart process of only those SLCs that are assigned to the selected GNF.

Options

fpc-slot Specify the line card slot.

slc-id Specify the SLC you want to take offline, bring online or restart.

Required Privilege Level

View

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

fpc-slice | 7530

request chassis fpm resync

IN THIS SECTION

- Syntax | 14490
- Syntax (TX Matrix Routers) | 14490
- Syntax (TX Matrix Plus Routers) | 14490
- Syntax (MX Series Routers) | 14490
- Description | 14490
- Options | 14490
- Required Privilege Level | 14491
- Output Fields | 14491
- Sample Output | 14491
- Release Information | 14492

Syntax

```
request chassis fpm resync
```

Syntax (TX Matrix Routers)

```
request chassis fpm resync (lcc number / scc)
```

Syntax (TX Matrix Plus Routers)

```
request chassis fpm resync (lcc number /  sfc number)
```

Syntax (MX Series Routers)

```
request chassis fpm resync
<all-members>
<local>
<member member-id>
```

Description

(M40e, M120, M160, M320, MX Series, and T Series routers only) Resynchronize the craft interface status.

Options

all-members (MX Series routers only) (Optional) Resynchronize the craft interface status on all members of the Virtual Chassis configuration.

lcc *number* (TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.

- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Resynchronize the craft interface status on the local Virtual Chassis member.

member ***member-id*** (MX Series routers only) (Optional) Resynchronize the craft interface status on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

scc (TX Matrix routers only) Resynchronize the craft interface status on the TX Matrix router (switch-card chassis).

sfc ***number*** (TX Matrix Plus routers only) Resynchronize the craft interface status on the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with **0**.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis fpm resync

```
user@host> request chassis fpm resync
Front Panel resynced
```

request chassis fpm resync (MX2010 Router)

```
user@host > request chassis fpm resync
Front Panel resynced.
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

show chassis environment fpm | **17705**

Resynchronizing FPC Sequence Numbers with Active FPCs when an FPC Comes Online

request chassis high-availability failover services-redundancy-group

IN THIS SECTION

- Syntax | **14492**
- Description | **14493**
- Options | **14493**
- Required Privilege Level | **14493**
- Output Fields | **14493**
- Sample Output | **14493**
- Release Information | **14493**

Syntax

```
request chassis high-availability failover services-redundancy-group group-id peer-id peer-id
```

Description

Initiate a manual failover on service redundancy group of the peer node. Use the command from the active node of the service redundancy group.

Options

services-redundancy-group

Service redundancy group on which to initiate manual failover.

- **Range:** 1 through 255

peer-id

Node identifier of the SRG. After the failover, this node transitions as the new active node.

- **Range:** 1 through 10

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis high-availability failover services-redundancy-group 1

```
user@host> request chassis high-availability failover services-redundancy-group 1 peer-id 2
Initiated manual failover for redundancy group 1
```

Release Information

Command introduced in Junos OS 20.4R1.

RELATED DOCUMENTATION

Hardware Upgrade for SRX5000-Line SPC3 in a Multinode High Availability Setup

request chassis mcs

IN THIS SECTION

- Syntax | 14494
- Description | 14494
- Options | 14494
- Required Privilege Level | 14494
- Output Fields | 14495
- Sample Output | 14495
- Release Information | 14495

Syntax

```
request chassis mcs (offline | online | restart) slot slot-number
```

Description

(M40e and M160 routers only) Control the operation of the Miscellaneous Control Subsystem (MCS).

Options

offline	Take the MCS offline.
online	Bring the MCS online.
restart	Restart the MCS.
slot *slot-number*	MCS slot number. Replace ***slot-number*** with 0 or 1.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis mcs

```
user@host> request chassis mcs online slot 0
MCS 0 appears to be online already
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show chassis environment mcs | 17714

request chassis mic

IN THIS SECTION

- Syntax | 14496
- Description | 14496
- Options | 14496
- Required Privilege Level | 14497
- Output Fields | 14497
- Sample Output | 14497
- Platform-Specific `request chassis mic` Behavior | 14498
- Release Information | 14498

Syntax

```
request chassis mic (offline | online) fpc-slot slot-number mic-slot slot-number
```

Description

(MX Series routers only) Control the operation of the Modular Interface Cards (MICs) installed on a Modular Port Concentrator (MPC).

> **NOTE**: On MX960 routers, if the MIC is not functioning correctly, you should take the MPC offline, replace it with a new MPC, and reinstall the MIC. On MX104 routers, the `request chassis mic` command is not supported on FPC slot 2 and MIC slot 0.

Options

offline Take the MIC offline.

online Bring the MIC online.

fpc-slot *slot-number* FPC slot number where the MIC is installed:

- ACX4000 router—Replace ***fpc-slot*** with the value `0` or `1`.
- MX80 router—Replace ***fpc-slot*** with the value **1**. This command is not supported on FPC slot 0.
- MX240 router—Replace ***fpc-slot*** with a value from **0** through **2**.
- MX480 router—Replace ***fpc-slot*** with a value from **0** through **5**.
- MX960 router—Replace ***fpc-slot*** with a value from **0** through **11**.
- MX2020 router—Replace ***fpc-slot*** with a value from **0** through **19**.
- MX2010 router—Replace ***fpc-slot*** with a value from **0** through **9**.
- MX2008 router—Replace ***fpc-slot*** with a value from **0** through **9**.

mic-slot *slot-number* MIC slot number.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis mic (MX Routers with Media Services Blade [MSB])

```
user@host> request chassis mic fpc-slot 1 mic-slot 0
Possible completions:
  offline              Take MIC offline
  online               Bring MIC online
```

request chassis mic offline (MX304 Router)

```
user@host > request chassis mic mic-slot 0 fpc-slot 0 offline
fpc 0 mic 0 offline initiated, use "show chassis fpc mic-status" to verify
```

A 120-second delay is required between consecutive requests. If a command is issued within this period, the request fails with an error message. For example:

```
user@host > request chassis mic mic-slot 0 fpc-slot 0 offline
FPC0 MIC0 throttling too frequent offline and online requests, please try after 73 seconds
```

request chassis mic online (MX304 Router)

```
user@host > request chassis mic mic-slot 0 fpc-slot 0 online
fpc 0 mic 0 online initiated, use "show chassis fpc mic-status" to verify
```

request chassis mic online (MX2010 Router)

```
user@host> request chassis mic online fpc-slot 1 mic-slot
0
FPC 1, MIC 0 is already online
```

request chassis mic online (MX2008 Router)

```
user@host>request chassis mic online fpc-slot 0 mic-slot
0
FPC 0 is not online
```

Platform-Specific `request chassis mic` Behavior

Use Feature Explorer to confirm platform, and release support for specific features.

Use the following table to review platform-specific behaviors for your platform:

Table 110:

Platform	Difference
MX304 devices	Once the `request chassis mic` command is executed, you need to verify the MIC status using the `show chassis fpc mic-status` command before any further action. If you try to begin another MIC online/offline process during MIC online/offline progress, the system will display the `MIC <slot> is in transition, try again.` error message. A 120-second delay is required between consecutive requests. If a command is issued within this period, The request fails with the error message: `FPC <fpc-slot-number> MIC <mic-slot-number> throttling too frequent offline and online requests, please try after <number of seconds> seconds.` The command option `request chassis mic fpc-slot <> mic-slot <> <offline | online>` is supported for MX304-LMIC-OIR-PFE on MX304.

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

show chassis hardware | 18430

request chassis optics

IN THIS SECTION

- Syntax | 14499
- Description | 14499
- Options | 14499
- Required Privilege Level | 14500
- Output Fields | 14500
- Sample Output | 14500
- Release Information | 14501

Syntax

```
request chassis optics fpc-slot fpc-slot-number reactivate
```

Description

(MX240, MX480, and MX960 routers) Control the status of the optical transceiver.

Options

fpc-slot *fpc-slot-number* Slot number of the line card that houses the optical transceiver.

- MX240 router—Replace *fpc-slot-number* with a value from 0 through 2.
- MX480 router—Replace *fpc-slot-number* with a value from 0 through 5.
- MX960 router—Replace *fpc-slot-number* with a value from 0 through 11.

reactivate Reactivate the optical transceiver.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis optics (MX480 router)

```
user@host> request chassis optics fpc-slot 5
reactivate
Enable FPC 5 non-nebs optics.
```

request chassis optics (MX10003 router)

```
user@host>request chassis optics fpc-slot 1 reactivate
Enable FPC 1 non-nebs optics.
```

request chassis optics (EX9251 swich)

```
user@switch>request chassis optics fpc-slot 0 reactivate
Enable FPC 0 non-nebs optics.
```

request chassis optics (EX9253 swich)

```
user@switch>request chassis optics fpc-slot 1 reactivate
Enable FPC 1 non-nebs optics.
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Determining Transceiver Support and Specifications for M Series and T Series Routers

request chassis pcg

IN THIS SECTION

- Syntax | 14501
- Description | 14501
- Options | 14501
- Required Privilege Level | 14502
- Output Fields | 14502
- Sample Output | 14502
- Release Information | 14502

Syntax

```
request chassis pcg (offline | online) slot slot-number
```

Description

(M40e and M160 routers) Control the operation of the Packet Forwarding Engine (PFE) clock generator (PCG).

Options

offline	Take the PCG offline.

online	Bring the PCG online.
slot *slot-number*	PCG slot number. Replace ***slot-number*** with **0** or **1**.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis pcg

```
user@host> request chassis pcg online slot 0
PCG 1 appears to be already online
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show chassis environment pcg | 17748

request chassis pic

IN THIS SECTION

- Syntax | 14503
- Syntax (MX Series Routers) | 14503

- Syntax (TX Matrix and TX Matrix Plus Routers) | 14503
- Description | 14503
- Options | 14505
- Required Privilege Level | 14508
- Output Fields | 14508
- Sample Output | 14508
- Release Information | 14509

Syntax

```
request chassis pic (offline | online) fpc-slot slot-number pic-slot slot-number
```

Syntax (MX Series Routers)

```
request chassis pic (offline | online) fpc-slot slot-number pic-slot slot-number
<member member-id>
```

Syntax (TX Matrix and TX Matrix Plus Routers)

```
request chassis pic (offline | online) fpc-slot slot-number pic-slot slot-number
<lcc number>
```

Description

Control the operation of the PIC.

NOTE: Starting with Junos OS Release 12.3, it is possible that PICs brought offline by using the `request chassis fpc slot fpc-slot pic pic-slot offline` operational mode command can come online during a configuration commit or power-supply replacement procedure. (See the following note about difference in behavior on Junos OS Evolved.) As an

alternative, use the `set fpc` *`fpc-slot`* `pic` *`pic-slot`* `power off` configuration mode command at the `[edit chassis]` hierarchy level to ensure that PICs remain offline.

NOTE: On Junos OS Evolved, a PIC does not restart when you enter a `commit` command that configures an element of that PIC.

NOTE: The `request chassis pic (offline | online) fpc-slot` *`slot number`* `pic-slot` *`slot-number`* command is not supported for built-in PICs on MX Series routers.

To view a list of built-in PICs on the router or switch chassis, use the `show chassis hardware` command.

NOTE: This command is not supported on MX960 and MX2020 routers with MPC5EQ.

NOTE: T1600 routers and TX Matrix Plus routers with 100-Gigabit Ethernet PICs require two adjacent PIC slots, 0 and 1, for each PIC. Therefore, only `online` and `offline` command options to PIC slot 0 are allowed. Use of the `online` and `offline` command options for PIC slot 1 with the described router and PIC combination is not allowed.

NOTE: In T Series routers, when the PIC state is set from `offline` to `online` or vice-versa before the processing is complete for the previous command, you are provided feedback on the status of your request. The following sample messages are displayed if you try to set a PIC `offline` or `online`:

```
user@switch> request chassis pic fpc-slot 1
pic-slot 0 online
fpc 1 pic 0 online initiated, use "show chassis fpc pic-status" to verify
user@switch> request chassis pic fpc-slot 1
pic-slot 0 online
FPC 1 PIC 0 already transitioning to online
```

When the same PIC is set to a different state while the transition is in progress, you are provided feedback on the status of your request.

```
user@switch> request chassis pic fpc-slot 1
pic-slot 0 offline
FPC 1, PIC 0 already transitioning to online.  Please retry later.
```

NOTE: If a CLI-based firmware upgrade is in progress, it prevents the specified PIC from restarting. Starting in Junos OS Release 15.1, the following message is displayed:

```
user@host> request chassis pic fpc-slot 0 pic-slot
1 offline
PIC's Firmware update in progress. Wait!!!
```

Options

offline Take the PIC offline.

online Bring the PIC online.

fpc-slot *slot-number* Flexible PIC Concentrator (FPC) slot number. Replace *slot-number* with a value appropriate for your router or switch:

- ACX4000 routers—1 or 2.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—0.
  - EX4200 switches in a Virtual Chassis configuration—0 through 9 (switch's member ID).
  - EX8208 switches—0 through 7 (line card).
  - EX8216 switches—0 through 15 (line card).
- M5, M7i, M10, and M10i routers—0 or 1.
- M20 routers—0 through 3.
- M40 and M40e routers—0 through 7.
- M120 routers—0 through 5.

- M160 routers—0 through 7.
- M320 routers—0 through 7.
- MX 5, MX10, and MX40 routers—0 or 1.
- MX80 routers—0 or 1.
- MX240 routers—0 through 2
- MX480 routers—0 through 5
- MX2020 routers—0 through 19.
- MX2010 routers—0 through 9.
- MX960 routers—0 through 11.
- MX10003 routers—0 or 1.
- MX204 routers—0.
- PTX5000 routers—0 or 1.
- T Series routers—0 through 7.
- TX Matrix and TX Matrix Plus routers only—On a TX Matrix router, if you specify the number of the T640 router by using the `lcc` *`number`* option (the recommended method), replace *`slot-number`* with a value from 0 through 7. Otherwise, replace *`slot-number`* with a value from 0 through 31.

  Likewise, on a TX Matrix Plus router, if you specify the *`number`* of the T1600 or T4000 router by using the lcc number option (the recommended method), replace *`slot-number`* with a value from 0 through 7. Otherwise, for the FPC slot number, replace *`slot-number`* with a value from 0 through 31. On a TX Matrix Plus router with 3D SIBs to assign the FPC slot number, replace *slot-number* with a value from `0` through `63`. For example, the following commands have the same result:

  ```
  user@host> request chassis pic fpc-slot 1 lcc 1 pic-slot 0 offline
  user@host> request chassis pic fpc-slot 9 pic-slot 0 offline
  ```

- QFX5100 standalone switches—0.

lcc *number* (TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

member ***member-id***
(MX Series routers only) (Optional) Change the PIC status on the specified member of the Virtual Chassis configuration. Replace *member-id* with the value that is assigned to the specified member.

offline
Take the PIC offline.

online
Bring the PIC online.

pic-slot ***slot-number***
PIC slot number.

- EX3200 and EX4200 switches—0 for built-in network interfaces and 1 for interfaces on uplink modules.
- EX8208 and EX8216 switches—0.
- M Series routers—0, 1, 2, or 3
- MX960 router—***slot-number*** corresponds to the slot number of the Packet Forwarding Engine.
- MX204 router—0 or 1.
- PTX5000 routers—0 or 1.
- T320 router—0 or 1.
- T640 router—0, 1, 2, or 3.
- T1600 router —0, 1, 2, or 3.
- T4000 router—0, 1, 2, or 3.
- QFX5100 standalone switches—0, 1, or 2. PIC 0 is used for all interfaces that are not configured on expansion modules, and PIC 1 and PIC 2 are used for interfaces configured on expansion modules.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis pic

```
user@host> request chassis pic pic-slot
0 online fpc-slot 0
FPC 0, PIC 0 is already online
```

request chassis pic online member (MX Series Routers)

```
user@host> request chassis pic online member
1 fpc-slot 11 pic-slot 3
fpc 11 pic 3 online initiated
```

request chassis pic offline member (MX Series Routers)

```
user@host> request chassis pic offline member
1 fpc-slot 11 pic-slot 3
fpc 11 pic 3 offline initiated
```

request chassis pic (MX10003 Router)

```
user@host> request chassis pic online pic-slot
1 fpc-slot 0
FPC 0 is not online
```

request chassis pic online member (PTX10008 Router)

```
user@host> request chassis pic online pic-slot
1 fpc-slot 0
FPC 0, PIC 1 is empty
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `member` introduced in Junos OS Release 14.2 for MX Series routers.

RELATED DOCUMENTATION

show chassis hardware | 18430

show chassis pic | 18588

request chassis pic-mode

IN THIS SECTION

- Syntax | 14509
- Description | 14510
- Options | 14510
- Required Privilege Level | 14511
- Sample Output | 14511
- Release Information | 14511

Syntax

```
request chassis pic-mode (intraconnect | virtual-chassis)
<all-members>
```

```
<local>
<member member-id>
```

Description

Set the PIC mode on an EX4500 switch. The PIC mode command determines whether the module that is plugged into the EX4500 switch operates as an intraconnect module or a Virtual Chassis module.

> **NOTE**: Operating an EX4500 switch without the intraconnect module or the Virtual Chassis module is not supported.

> **NOTE**: If you use the EX4500 switch as a standalone switch without changing the PIC mode setting to `intraconnect`, you will not be able to achieve 960 Gbps data rate throughput. Some of the bandwidth is reserved for the dedicated Virtual Chassis ports (VCPs).

The PIC mode must be set to `virtual-chassis` if the EX4500 switch is cabled into a Virtual Chassis using the dedicated VCPs on the Virtual Chassis module. This is the setting in the default factory configuration. The `intraconnect` PIC mode is not supported by the default factory configuration.

If you want to use the EX4500 switch as a standalone switch, you must take several steps to reset the PIC mode as `intraconnect`. You cannot reset the PIC mode `intraconnect` simply by issuing the `request chassis pic-mode intraconnect` command. See *Resetting the EX4500 Switch to Run as a Standalone Switch (CLI Procedure)*.

Use the "show chassis pic-mode" on page 18628 command to verify the current PIC mode setting.

The PIC mode setting is maintained through reboots even though it is set in operational mode.

You do not need to enter this command to use the Virtual Chassis module on an EX4550 switch.

Options

intraconnect	Set the PIC mode to intraconnect.
virtual-chassis	Set the PIC mode to Virtual Chassis.
all-members	(Optional) Set the PIC mode for all members of the Virtual Chassis configuration. This is the default option.
local	(Optional) Set the PIC mode on the member switch where the command is issued.

member *member-id* (Optional) Set the PIC mode on the specified member of the Virtual Chassis configuration.

Required Privilege Level

view

Sample Output

request chassis pic-mode virtual-chassis

```
user@switch> request chassis pic-mode virtual-chassis
```

request chassis pic-mode intraconnect

```
user@switch> request chassis pic-mode intraconnect

fpc0:
--------------------------------------------------------------------------------------------------
virtual-chassis (mode unchanged; intraconnect mode not supported on factory-default config)
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- *Verifying the Setting for the PIC Mode on an EX4500 Switch in a Virtual Chassis*
- *Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)*
- *Adding an EX4500 Switch to a Preprovisioned EX4200 Virtual Chassis (CLI Procedure)*
- *Adding an EX4500 Switch to a Nonprovisioned EX4200 Virtual Chassis (CLI Procedure)*

request chassis port-led

IN THIS SECTION

- Syntax | 14512
- Description | 14512
- Options | 14512
- Required Privilege Level | 14514
- Output Fields | 14514
- Sample Output | 14514
- Release Information | 14516

Syntax

```
request chassis port-led (start | stop) fpc-slot fpc-slot-number pic-slot pic-slot-number port
(port-number| all-10g | all-25g | all-40g | all-50g | all-100g | all-400g | all-port) duration
duration
```

Description

Enable remote port identification of the Modular Interface Card (MIC) or PIC by making the LED of the appropriate active port blink for a duration of time. You can also make the LEDs of all active ports, specific active ports, and specific port types blink. For instance, on MX480 routers with MPC7E-MRATE, you can make the LED of all active ports that support port speeds of 100 Gbps blink. Enabling remote port identification provides cabling assistance and reduces cabling mistakes.

NOTE: You can stop a LED from blinking before the end of duration. After the LED stops blinking, it resumes its normal operation of indicating current link status of the port.

Options

start Start blinking of the LEDs of the specified ports.

stop Stop the LEDs of the specified ports from blinking.

fpc-slot *fpc-slot-number* (MX Series routers) Modular Port Concentrator (MPC) slot number. Replace *fpc-slot-number* with a value appropriate for your router:

- MX240 routers—0 through 2.
- MX480 routers—0 through 5.
- MX960 routers—0 through 11.
- MX2010 routers—0 through 9.
- MX2020 routers—0 through 19.
- MX10003 routers—0 through 1.
- MX10008 routers—0 through 6.

pic-slot *pic-slot-number* (MX Series routers) MIC slot number. Replace *pic-slot-number* with a value from 0 through 3. For MX10008 routers with JNP10K-LC2101 and JNP10K-LC9600 MPC only, replace *pic-slot-number* with a value from 0 through 5.

port *port-number* Port number. Replace *port-number* with a value appropriate for your MPC or line card:

- MIC-MRATE—0 through 11.
- MX10003 MPC—0 through 5 for PIC 0 and 0 through 11 for PIC1
- MX10K-LC2101 and MX10K-LC9600 MPC—0 through 3 per PIC. MX10K-LC2101 and MX10K-LC9600 supports 6 built-in PICs and each PIC supports 4 ports.

all-10g Active ports that support port speed of 10 Gbps.

all-25g Active ports that support port speed of 25 Gbps.

all-40g Active ports that support port speed of 40 Gbps.

all-50g Active ports that support port speed of 50 Gbps.

all-100g Active ports that support port speed of 100 Gbps.

all-400g Active ports that support port speed of 400 Gbps.

all-ports All active ports.

duration *duration*

Duration, in seconds, to perform LED blinking. Replace *duration* with a value from 0 through 65,535. The default duration is 5 minutes (300 seconds).

Required Privilege Level

view

Output Fields

When you enter this command, you are asked to verify the LED status based on your request.

Sample Output

request chassis port-led (MX2020 Routers with MPC8E)

```
user@host> request chassis port-led start fpc-slot
3 pic-slot 0 port all-10g duration 5
Command sent to FPC_3. Check physically about LED status on the PIC_0 ports.
```

request chassis port-led (MX480 Routers with MPC7E-MRATE)

```
user@host> request chassis port-led start fpc-slot
3 pic-slot 0 port 2 duration 5
Command sent to FPC_3. Check physically about LED status on the PIC_0 ports.
```

request chassis port-led (MX10003 Router)

```
user@host> request chassis port-led start fpc-slot
1 pic-slot 1 port 0 duration 5
Command sent to FPC_1. Check physically about LED status on the PIC_1 ports.
```

request chassis port-led (MX10003 Router for all active 100GE ports)

```
user@host> request chassis port-led start fpc-slot
1 pic-slot 1 port all-100g duration 5
Command sent to FPC_1. Check physically about LED status on the PIC_1 ports.
```

request chassis port-led (MX10003 Router for all active 40GE ports)

```
user@host> request chassis port-led start fpc-slot
0 pic-slot 1 port all-40g duration 5
Command sent to FPC_0. Check physically about LED status on the PIC_1 ports.
```

request chassis port-led (MX10003 Router for all the active ports in a PIC)

```
user@host> request chassis port-led start fpc-slot
0 pic-slot 0 port all-ports duration 5
Command sent to FPC_0. Check physically about LED status on the PIC_0 ports.
```

request chassis port-led (MX204 Router for all the active ports in a PIC)

```
user@host> request chassis port-led start fpc-slot
0 pic-slot 0 port all-ports duration 5
Command sent to FPC_0. Check physically about LED status on the PIC_0 ports.
```

request chassis port-led (EX9251 Switch)

```
user@switch> request chassis port-led start
fpc-slot 0 pic-slot 0 port all-ports duration 5
Command sent to FPC_0. Check physically about LED status on the PIC_0 ports.
```

request chassis port-led (MX10008 Router with MX10K-LC2101 and MX10K-LC9600)

```
user@host>request chassis port-led start fpc-slot
0 pic-slot 2 all-ports duration 15
Command sent to FPC_0. Check physically about LED status on the PIC_2 ports.
```

request chassis port-led (ACX Series)

```
user@host>request chassis port-led start fpc-slot 0 pic-slot 0 port 20 duration 1
Command sent to FPC-0. Check the LEDs status on the PIC-0 ports.
```

```
user@host>request chassis port-led stop fpc-slot 0 pic-slot 0 port 20
Command sent to FPC-0. Check the LEDs status on the PIC-0 ports.
```

Release Information

Command introduced in Junos OS Release 15.1F4.

RELATED DOCUMENTATION

Craft Interface

request chassis primary-ha-control-port-transition

IN THIS SECTION

- Syntax | 14517
- Description | 14517
- Options | 14517
- Required Privilege Level | 14517
- Output Fields | 14517

Sample Output | 14518
Release Information | 14518

Syntax

```
request chassis primary-ha-control-port-transition <fpc FPC slot number | from-fpc-to-scb | from-scb-to-fpc>
```

Description

With this command, you can disable the *FPC slot number* HA port 0 on local node, and then enable the SCB0 HA control port on local node or disable the SCB0 HA control port on local node, and then enable the *FPC slot number* HA port 0 on local node.

This command is used during single control link online transition and it can be carried out in cluster mode on SRX5K series devices.

Options

fpc	Specify the FPC slot number. **Range:** 0 through 23
from-fpc-to-scb	Transition from FPC to SCB.
from-scb-to-fpc	Transition from SCB to FPC.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis primary-ha-control-port-transition

```
user@host> request chassis primary-ha-control-port-transition from-fpc-to-scb fpc 0
```

```
fpc 0 HA control port 0 disabled & scb 0 HA control port enabled
```

```
user@host> request chassis primary-ha-control-port-transition from-scb-to-fpc fpc 0
```

```
scb 0 HA control port disabled & fpc 0 HA control port 0 enabled
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Chassis Cluster Dual Control Links

request chassis redundancy feb slot

IN THIS SECTION

- Syntax | 14519
- Description | 14519
- Options | 14519
- Required Privilege Level | 14519
- Output Fields | 14519
- Sample Output | 14519

Release Information | 14520

Syntax

```
request chassis redundancy feb slot slot-number (switch-to-backup | revert-from-backup)
```

Description

(M120 routers only) Control the operation of the specified Forwarding Engine Board (FEB) in a redundancy group.

Options

slot-number	FEB slot number. Replace ***slot-number*** with a value from **0** through **5**.
switch-to-backup	Initiate a switchover from the specified active FEB to the backup FEB for the redundancy group.
revert-from-backup	Initiate a revert to the specified FEB following a switchover from the backup FEB for a redundancy group.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis redundancy feb slot 2 switch-to-backup

```
user@host> request chassis redundancy feb slot 2 switch-to-backup
Switch initiated, use “show chassis redundancy febs” to verify
```

request chassis redundancy feb slot 3 revert-from-backup

```
user@host> request chassis redundancy feb slot 3 revert-from-backup
Revert initiated, use “show chassis redundancy febs” to verify
```

Release Information

Command introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

show chassis redundancy feb | 18694

Configuring FEB Redundancy on the M120 Router

Understanding Switching Control Board Redundancy

request chassis routing-engine hard-disk-test

IN THIS SECTION

- Syntax | 14520
- Description | 14521
- Required Privilege Level | 14521
- Output Fields | 14521
- Sample Output | 14524
- Release Information | 14528

Syntax

```
request chassis routing-engine hard-disk-test
```

Description

Run a Hard Disk Drive (HDD) memory and diagnostics monitoring test for an XRE200 External Routing Engine. The test provides output that is helpful in monitoring the status of the hard disk memory on the XRE200 External Routing Engine.

Required Privilege Level

system-control

Output Fields

Table 111 on page 14521 lists the output fields for the `request chassis routing-engine hard-disk-test` command. Output fields are listed in the approximate order in which they appear.

Table 111: request chassis routing-engine hard-disk-test Output Fields

Field Name	Field Description
RAID device path	The device path to the RAID. This output is generally only useful for customer support personnel.
Firmware Version	The firmware version of the RAID. This output is generally only useful for customer support personnel.
RAID controller s/n	The serial number of the RAID controller. This output is generally only useful for customer support personnel.
RAID Chip ID	The RAID chip identification number. This output is generally only useful for customer support personnel.
RAID policy	The RAID policy status. SAFE is the only RAID policy used by an XRE200 External Routing Engine. This output is generally only useful for customer support personnel.

Table 111: request chassis routing-engine hard-disk-test Output Fields *(Continued)*

Field Name	Field Description
Drive0 model	The model number of the HDD module in HDD slot 0. This output is generally only useful for customer support personnel.
Drive1 model	The model number of the HDD module in HDD slot 1. This output is generally only useful for customer support personnel.
Drive0 s/n	The serial number of the HDD module in HDD slot 0. This output is generally only useful for customer support personnel.
Drive1 s/n	The serial number of the HDD module in HDD slot 1. This output is generally only useful for customer support personnel.
Drive0 capacity	The memory capacity of the HDD module in HDD slot 0, in gigabytes.
Drive1 capacity	The memory capacity of the HDD module in HDD slot 1, in gigabytes.
Drive0 (RAID STATUS)	The RAID status of the HDD module in HDD slot 0. Outputs include: • **On-line**—The HDD module is online. • **Off-line**—The HDD module is offline. If the status is **Off-line**, verify that the HDD module is securely seated in its slot. If the HDD module is securely seated and you cannot identify the reason for this status, contact Juniper Networks customer support (JTAC). • **Unknown**—HDD module RAID status is unknown. If the status is **Unknown**, wait a few minutes then retry the test. If the status remains **Unknown** after the second test, contact Juniper Networks customer support (JTAC).

Table 111: request chassis routing-engine hard-disk-test Output Fields *(Continued)*

Field Name	Field Description
Drive1 (RAID STATUS)	The RAID status of the HDD module in HDD slot 1. Outputs include: • **On-line**—HDD module is online. • **Off-line**—HDD module is offline. If the status is **Off-line**, verify that the HDD module is securely seated in its slot. If the HDD module is securely seated and you cannot identify the reason for this status, contact Juniper Networks customer support (JTAC). • **Unknown**—HDD module RAID status is unknown. If the status is **Unknown**, wait a few minutes then retry the test. If the status remains **Unknown** after the second test, contact Juniper Networks customer support (JTAC).
Number of partitions	The number of RAID memory partitions.
Size of Partitions	The size of each partition, in gigabytes.

Table 111: request chassis routing-engine hard-disk-test Output Fields *(Continued)*

Field Name	Field Description
RAID Status	The RAID status. Outputs include: • **Healthy**—RAID is healthy. • **Degraded**—RAID has detected a problem with at least one HDD module. If **Degraded** appears in the RAID Status field, contact customer support. • **Unknown**—RAID status is unknown. If the status is **Unknown**, wait a few minutes then retry the test. If the status remains **Unknown** after the second test, contact Juniper Networks customer support (JTAC). • **Rebuilding**—The RAID is rebuilding. The output includes a drive number and a percentage complete indication. The RAID rebuilding process can take up to 40 minutes. Wait several minutes after the process has finished and then retry the test. • **Verifying**—The RAID is rebuilt and being verified by the system. The output includes a drive number and a percentage complete indication. The RAID verification process can take up to 40 minutes. Wait several minutes after the process has finished and then retry the test.

Sample Output

request chassis routing-engine hard-disk-test

```
user@switch> request chassis routing-engine hard-disk-test
RAID INFORMATION
RAID device path: /dev/ad4
```

```
Firmware Version: 11594
RAID controller s/n: 12345678
RAID Chip ID: 123
RAID policy: SAFE

Drive0 model: WDC WD123AAJS-4567A0
Drive1 model: WDC WD345JD-18MSA1
Drive0 s/n:      WD-WCAT30214999
Drive1 s/n:      WD-WMAM9DTK4111
Drive0 capacity: 74(GB)
Drive1 capacity: 74(GB)

RAID STATUS
Drive0: On-line
Drive1: On-line
Number of partitions: 1
Size of Partitions:
    Partition 0: 74(GB)
RAID Status: Healthy
```

request chassis routing-engine hard-disk-test (ACX Series)

```
user@host> request chassis routing-engine hard-disk-test long disk /dev/sda
smartctl 7.3 2022-02-28 r5338 [x86_64-linux-5.15.76-yocto-standard-juniper-02183-g46c537862907]
(local build)
Copyright (C) 2002-22, Bruce Allen, Christian Franke, www.smartmontools.org

=== START OF OFFLINE IMMEDIATE AND SELF-TEST SECTION ===
Sending command: "Execute SMART Extended self-test routine immediately in off-line mode".
Drive command "Execute SMART Extended self-test routine immediately in off-line mode" successful.
Testing has begun.
Please wait 15 minutes for test to complete.
Test will complete after Fri Sep 20 08:16:32 2024 PDT
Use smartctl -X to abort test.
```

```
user@host> request chassis routing-engine hard-disk-test short disk /dev/sda
smartctl 7.3 2022-02-28 r5338 [x86_64-linux-5.15.76-yocto-standard-juniper-02183-g46c537862907]
(local build)
smartctl 7.3 2022-02-28 r5338 [x86_64-linux-5.15.76-yocto-standard-juniper-02183-g46c537862907]
(local build)
```

```
Copyright (C) 2002-22, Bruce Allen, Christian Franke, www.smartmontools.org

=== START OF OFFLINE IMMEDIATE AND SELF-TEST SECTION ===
Can't start self-test without aborting current test (90% remaining),
add '-t force' option to override, or run 'smartctl -X' to abort test.
```

```
user@host> request chassis routing-engine hard-disk-test show-status disk /dev/sda
smartctl 7.3 2022-02-28 r5338 [x86_64-linux-5.15.76-yocto-standard-juniper-02183-g46c537862907]
(local build)
Copyright (C) 2002-22, Bruce Allen, Christian Franke, www.smartmontools.org

=== START OF INFORMATION SECTION ===
Model Family:     SwissBit Slim Sata SSD
Device Model:     SFSA100GM2AK4TO-I-6B-66A-JUN
Serial Number:    0000602219249A000046
Firmware Version: SBR10021
User Capacity:    100,030,242,816 bytes [100 GB]
Sector Size:      512 bytes logical/physical
Rotation Rate:    Solid State Device
Form Factor:      2.5 inches
TRIM Command:     Available, deterministic, zeroed
Device is:        In smartctl database 7.3/5319
ATA Version is:   ACS-3 (minor revision not indicated)
SATA Version is:  SATA 3.2, 6.0 Gb/s (current: 6.0 Gb/s)
Local Time is:    Fri Sep 20 08:05:24 2024 PDT
SMART support is: Available - device has SMART capability.
SMART support is: Enabled

=== START OF READ SMART DATA SECTION ===
SMART overall-health self-assessment test result: PASSED

General SMART Values:
Offline data collection status:  (0x03)    Offline data collection activity
                     is in progress.
                     Auto Offline Data Collection: Disabled.
Self-test execution status:      ( 247)    Self-test routine in progress...
                     70% of test remaining.
Total time to complete Offline
data collection:         (    0) seconds.
Offline data collection
capabilities:              (0x53) SMART execute Offline immediate.
```

```
                    Auto Offline data collection on/off support.
                    Suspend Offline collection upon new
                    command.
                    No Offline surface scan supported.
                    Self-test supported.
                    No Conveyance Self-test supported.
                    Selective Self-test supported.
SMART capabilities:            (0x0003)    Saves SMART data before entering
                    power-saving mode.
                    Supports SMART auto save timer.
Error logging capability:        (0x01)    Error logging supported.
                    General Purpose Logging supported.
Short self-test routine
recommended polling time:      (   2) minutes.
Extended self-test routine
recommended polling time:      (  15) minutes.
SCT capabilities:            (0x003d)    SCT Status supported.
                    SCT Error Recovery Control supported.
                    SCT Feature Control supported.
                    SCT Data Table supported.

SMART Attributes Data Structure revision number: 1
Vendor Specific SMART Attributes with Thresholds:
ID# ATTRIBUTE_NAME          FLAG     VALUE WORST THRESH TYPE      UPDATED  WHEN_FAILED RAW_VALUE
  1 Raw_Read_Error_Rate     0x000b   100   100   000    Pre-fail  Always       -       0
  5 Reallocated_Sector_Ct   0x0013   100   100   000    Pre-fail  Always       -       0
  9 Power_On_Hours          0x0012   100   100   000    Old_age   Always       -       7720
 12 Power_Cycle_Count       0x0012   100   100   000    Old_age   Always       -       136
 16 Unknown_Attribute       0x0112   094   094   001    Old_age   Always       -       1955
 17 Unknown_Attribute       0x0012   100   100   000    Old_age   Always       -       30000
160 Uncorrect_Sect_Cnt_RW   0x0002   100   100   000    Old_age   Always       -       0
161 Spare_Block_Count       0x0012   100   100   000    Old_age   Always       -       61539243
163 Initial_Invalid_Blocks  0x0003   100   100   000    Pre-fail  Always       -       10
164 Total_Erase_Count       0x0012   100   100   000    Old_age   Always       -       593053
165 Maximum_Erase_Count     0x0002   100   100   000    Old_age   Always       -       449
166 Minimum_Erase_Count     0x0012   100   100   000    Old_age   Always       -       348
167 Average_Erase_Count     0x0012   087   087   000    Old_age   Always       -       408
168 MaxSpecified_Erase_Count0x0012   100   100   000    Old_age   Always       -       3000
169 PowerOn_UECC_Count      0x0003   100   100   000    Pre-fail  Always       -       102
193 Dynamic_Remaps          0x0012   100   100   000    Old_age   Always       -       0
194 Temperature_Celsius     0x0023   100   100   000    Pre-fail  Always       -       3871788
195 Hardware_ECC_Recovered  0x0012   100   100   000    Old_age   Always       -       12821
196 Reallocated_Event_Count 0x0012   000   000   000    Old_age   Always       -       0
```

```
198 Uncorrectable_Sect_Count0x0012   100   100   000    Old_age   Always       -       0
199 SATA_Phy_CRC_Error_Count0x000b   100   100   000    Pre-fail  Always       -       0
215 Trim_Count              0x0012   100   100   000    Old_age   Always       -       1816
235 TotalFlashLBA_Written   0x0012   100   100   000    Old_age   Always       -
233195432448
237 TotalFlashLBA_Written_E 0x0012   100   100   000    Old_age   Always       -       0
241 TotalHostLBA_Write      0x0012   100   100   000    Old_age   Always       -
167103620097
242 TotalHostLBA_Read       0x0012   100   100   000    Old_age   Always       -
19887123867
243 TotalHostLBA_Written_E  0x0012   100   100   000    Old_age   Always       -       0
244 TotalHostLBA_Read_E     0x0012   100   100   000    Old_age   Always       -       0
248 SSD_Life_Left           0x0112   087   087   001    Old_age   Always       -       87
249 Spare_Block_Life_Left   0x1913   100   100   025    Pre-fail  Always       -       100

SMART Error Log Version: 1
No Errors Logged

SMART Self-test log structure revision number 1
No self-tests have been logged.  [To run self-tests, use: smartctl -t]

SMART Selective self-test log data structure revision number 1
 SPAN  MIN_LBA  MAX_LBA  CURRENT_TEST_STATUS
    1        0        0  Not_testing
    2        0        0  Not_testing
    3        0        0  Not_testing
    4        0        0  Not_testing
    5        0        0  Not_testing
Selective self-test flags (0x0):
  After scanning selected spans, do NOT read-scan remainder of disk.
If Selective self-test is pending on power-up, resume after 0 minute delay.
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Hard Disk Drive (HDD) Module in an XRE200 External Routing Engine
- XRE200 External Routing Engine Hardware Overview

request chassis routing-engine master

IN THIS SECTION

- Syntax | 14529
- Syntax (M Series, MX Series, ACX7509, ACX1000 Devices) | 14529
- Syntax (MX Series Virtual Chassis) | 14529
- Syntax (QFX Series) | 14530
- Description | 14530
- Options | 14530
- Additional Information | 14531
- Required Privilege Level | 14532
- Output Fields | 14532
- Sample Output | 14532
- Release Information | 14535

Syntax

```
request chassis routing-engine master (acquire | release | switch)
<no-confirm>
```

Syntax (M Series, MX Series, ACX7509, ACX1000 Devices)

```
request chassis routing-engine master (acquire | release | switch)
<no-confirm>
<check>
```

Syntax (MX Series Virtual Chassis)

```
request chassis routing-engine master (acquire | release | switch)
<all-members>
<check>
```

```
<local>
<member member-id>
<no-confirm>
```

Syntax (QFX Series)

```
request chassis routing-engine master (release | switch)
<check>
<interconnect-device name>
<node-group name>
<no-confirm>
```

Description

For routers or switches with multiple Routing Engines, control which Routing Engine is the primary.

NOTE: Successive graceful Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

If the router or switch displays a warning message similar to “Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset,” do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines that were not ready for graceful switchover are reset. None of the Flexible PIC concentrators (FPCs) should spontaneously restart. We recommend that you wait until the warning no longer appears and then proceed with the switchover.

You will receive an error message stating “Command terminated. Not ready for primary role switch, try after n seconds” when this command is re-entered before 240 seconds have elapsed on EX Series switches.

NOTE: On a QFabric system, to avoid traffic loss on the network Node group, switch primary role of the routing engine to the backup routing engine, and then reboot.

Options

acquire (Not available for Junos OS Evolved) Attempt to become the primary Routing Engine.

release (Not available for Junos OS Evolved) Request that the other Routing Engine become the primary.

switch Toggle primary role between Routing Engines.

> **NOTE**: The `acquire` option should be used with caution because acquiring a Routing Engine may result in a corrupted database. If possible, use the `switch` option instead.

The `acquire`, `release`, and `switch` options have the following suboptions:

all-members (MX Series routers only) (Optional) Control Routing Engine primary role on the Routing Engines in all member routers of the Virtual Chassis configuration.

check (ACX7509, MX104, MX480, MX960, MX2010, MX2020, and MX2008 routers, QFabric systems, and PTX5000 routers only) (Optional) Available with the `switch`, `release`, and `acquire` options. Check graceful switchover status of the standby Routing Engine before toggling primary role between Routing Engines (RE). Check if the RE is ready for a switchover.

interconnect-device *name* (QFabric systems only) (Optional) Control Routing Engine primary role on the Routing Engines on an Interconnect device.

local (MX Series routers only) (Optional) Control Routing Engine primary role on the Routing Engines in the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Control Routing Engine primary role on the Routing Engines of the specified member in the Virtual Chassis configuration. Replace *`member-id`* with a value of 0 or 1.

no-confirm (Optional) Do not request confirmation for the switch.

node-group *name* (QFabric systems only) (Optional) Control Routing Engine primary role on the Routing Engines on a Node group.

Additional Information

Always use the `show system switchover` command on the backup Routing Engine to determine if the backup is ready to take over as the primary Routing Engine.

By default, the Routing Engine in slot 0 (`RE0`) is the primary and the Routing Engine in slot 1 (`RE1`) is the backup. To change the default primary Routing Engine, include the `routing-engine` statement at the [`edit chassis redundancy`] hierarchy level in the configuration.

To have the backup Routing Engine become the primary Routing Engine, use the `request chassis routing-engine master switch` command. If you use this command to change the primary and then restart the chassis software for any reason, the primary reverts to the default setting.

> **NOTE**: Although the configurations on the two Routing Engines do not have to be the same and are not automatically synchronized, we recommend making both configurations the same.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis routing-engine primary acquire

```
user@host> request chassis routing-engine master acquire

warning: Traffic will be interrupted while the PFE is re-initialized

warning: The other routing engine's file system could be corrupted

Reset other routing engine and become master ? [yes,no] (no)
```

request chassis routing-engine primary switch

```
user@host> request chassis routing-engine master switch

warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between Routing Engines ? [yes,no] (no) yes

Resolving mastership...
Complete. The other Routing Engine becomes the master.
```

Switch primary role back to the local Routing Engine:

```
user@host> request chassis routing-engine master switch

warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between routing engines ? [yes,no] (no) yes

Resolving mastership...
Complete. The local routing engine becomes the master.
```

request chassis routing-engine master switch (Junos OS Evolved)

```
user@host> request chassis routing-engine master switch
Resolving mastership...
Complete. The other Routing Engine becomes the master.
```

Switching back to primary router:

```
user@host> request chassis routing-engine master switch
Resolving mastership...
Complete. The local Routing Engine becomes the master.
```

If you did not switch back and tried to enter configuration mode, you would get the following error message:

```
user@host> configure
error: unknown command: configure
Configuration is allowed only from the master Routing Engine.
```

request chassis routing-engine master switch (ACX7348)

```
user@host> request chassis routing-engine master switch
Toggle mastership between routing engines ? [yes,no] (no) yes

Resolving mastership...
Mastership will switchover within 45 secs
Please check mastership status using 'show chassis routing-engine'
```

request chassis routing-engine primary switch check (ACX7509, M Series, and MX Series Devices)

```
warning: Standby Routing Engine is not ready for graceful switchover.
{master}[edit]

user@host> request chassis routing-engine master switch check
Platform specific components not ready for switchover
```

Output when system is ready for graceful switchover.

```
user@host> request chassis routing-engine master switch check
Switchover Ready
```

You can similarly check the backup Routing Engine for the switchover readiness.

request chassis routing-engine master switch check (DRAM Size Mismatch Between Primary and Standby)

```
user@host> request chassis routing-engine master switch check
error: Standby mirror connection is not up:RE DRAM Size Mismatch

{master}
```

To check switchover readiness, use the `show system switchover` command before changing the mastership. Please note that the switchover is denied until 360 seconds have passed from the previous switchover.

request chassis routing-engine master switch check (Disabled Backup Routing Engine)

```
user@host> request chassis routing-engine master switch check
Backup RE disabled
```

request chassis routing-engine master switch no-confirm (Disabled Backup Routing Engine)

```
user@host> request chassis routing-engine master switch no-confirm
Backup RE disabled
```

Release Information

Command introduced before Junos OS Release 7.4.

`all-chassis` option added in Junos OS Release 8.0.

Command support added to ACX7509 platform in Junos OS Evo Release 22.1.

RELATED DOCUMENTATION

show system switchover | 27246

request chassis satellite beacon

IN THIS SECTION

- Syntax | 14535
- Description | 14536
- Options | 14536
- Required Privilege Level | 14536
- Sample Output | 14536
- Release Information | 14537

Syntax

```
request chassis satellite beacon
    <(off | on)>
    <fpc-slot slot-id (off | on)>
```

```
    <range range (off | on)>
```

Description

This command is used to enable or disable the beacon LED on satellite devices in a Junos Fusion.

To display the status and colors of the beacon LEDs of the satellite devices in a Junos Fusion, use the `show chassis led satellite` command.

Options

off	Turn the beacon LED `off`.
on	Turn the beacon LED `on`.
range *range*	Enable or disable the beacon LED in a range of FPC slot identifiers. For example, you can specify FPC slot identifiers 101, 102, and 103 by entering a range of 101-103.
fpc-slot *slot-id*	Enable or disable beacon LED for the satellite device using the specified FPC slot identifier in the Junos Fusion. The *slot-id* is the FPC slot ID number.

Required Privilege Level

maintenance

Sample Output

request chassis satellite beacon fpc-slot 136 on

```
user@aggregation-device> request chassis satellite beacon fpc-slot 136 on
```

request chassis satellite beacon range 101-103 off

```
user@aggregation-device> request chassis satellite beacon range 101-103 off
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

show chassis led satellite | 18551

Configuring or Expanding a Junos Fusion Enterprise

request chassis satellite disable

IN THIS SECTION

- Syntax | 14537
- Description | 14537
- Options | 14538
- Required Privilege Level | 14538
- Sample Output | 14538
- Sample Output | 14538
- Release Information | 14539

Syntax

```
request chassis satellite disable
    <device-alias alias-name>
    <fpc-slot fpc-slot>
```

Description

Disable the specified satellite device from the Junos Fusion.

When a satellite device is disabled from a Junos Fusion, all extended ports are immediately placed in the down state. The satellite device cannot send or receive traffic for the Junos Fusion until it is reenabled.

This command is useful whenever you need to disable a satellite device from a Junos Fusion, such as for troubleshooting scenarios. If you are removing a satellite device from a Junos Fusion to use the device elsewhere on the network, use the "request chassis satellite install" on page 14542 command to install Junos OS onto your satellite device before removing it from the Junos Fusion. See Removing a Satellite Device from a Junos Fusion.

You can reenable a satellite device that was disabled using this command using the "request chassis satellite enable" on page 14539 command.

Options

device-alias ***alias-name*** Disable the satellite device with the specified alias name from the Junos Fusion.

fpc ***fpc-slot*** Disable the satellite device with the specified FPC slot identifier from the Junos Fusion.

Required Privilege Level

system-control

Sample Output

request chassis satellite disable device-alias satellite-01

```
user@aggregation-device> request chassis satellite disable device-alias satellite-01
```

Sample Output

request chassis satellite disable fpc-slot 101

```
user@aggregation-device> request chassis satellite disable fpc-slot 101
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

request chassis satellite enable

IN THIS SECTION

- Syntax | 14539
- Description | 14539
- Options | 14540
- Required Privilege Level | 14540
- Sample Output | 14540
- Sample Output | 14540
- Release Information | 14540

Syntax

```
request chassis satellite enable
    <device-alias alias-name>
    <fpc-slot fpc-slot>
```

Description

Enable the specified device as a satellite device in a Junos Fusion.

This command is typically not used in any standard Junos Fusion initial configuration procedure. This command is typically needed in cases where the satellite device or cascade port has been disabled and needs to be re-enabled.

Options

device-alias *alias-name*	Enable the satellite device with the specified alias name in the Junos Fusion.
fpc *fpc-slot*	Enable the device with the specified FPC slot ID as a satellite device in the Junos Fusion.

Required Privilege Level

system-control

Sample Output

request chassis satellite enable device-alias satellite-01

```
user@aggregation-device> request chassis satellite enable device-alias satellite-01
```

Sample Output

request chassis satellite enable fpc-slot 101

```
user@aggregation-device> request chassis satellite enable fpc-slot 101
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

request chassis satellite file-copy

IN THIS SECTION

- Syntax | 14541
- Description | 14541
- Options | 14541
- Required Privilege Level | 14542
- Sample Output | 14542
- Release Information | 14542

Syntax

```
request chassis satellite file-copy [remote | local] <source-URL> <destination-URL>
```

Description

Copy a file between a satellite device and an aggregation device in a Junos Fusion.

Options

local Indicate that the file-copy from satellite-device has been initiated by a local user.

remote Indicate that the file-copy from satellite-device has been initiated by a remote user.

source-URL Specify the URL of the file that is copied.

If no device is specified as the *source-URL*, the file is copied from the aggregation device.

To specify a satellite device in the *source-URL*, enter **sd*slot-id-number*** at the beginning of the *source-URL*. For example, enter **sd101:/var/tmp/filename.txt** to specify that filename.txt in the /var/tmp directory on the satellite device using FPC slot ID number 101 is the *source-URL*.

destination-URL Specify the destination URL where the file is copied into.

If no device is specified as the *destination-URL*, the file is copied into the aggregation device.

To specify a satellite device in the *destination-URL*, enter **sd*slot-id-number*** at the beginning of the *destination-URL*. For example, enter **sd101:/var/tmp/** to specify the /var/tmp directory on the satellite device using FPC slot ID number 101 as the *source-URL*.

Required Privilege Level

system-control

Sample Output

request chassis satellite file-copy

```
root@aggregation-device> request chassis satellite file-copy local /var/tmp/file_name
sd101:/var/tmp/
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

request chassis satellite install

IN THIS SECTION

- Syntax | 14543
- Description | 14543

- Options | 14544
- Required Privilege Level | 14544
- Sample Output | 14544
- Release Information | 14545

Syntax

```
request chassis satellite install package-name
    [fpc-slot fpc-slot | device-alias device-alias]
    <no-confirm>
```

Description

Install a version of Junos OS software onto a satellite device in a Junos Fusion.

Any device operating as a satellite device in a Junos Fusion is running satellite software. A device running satellite software cannot operate as a standalone network device until it is running a version of Junos OS software.

You would typically enter this command to install Junos OS onto a satellite device before removing the satellite device from a Junos Fusion. Installing the Junos OS software onto the satellite device before removing it from the Junos Fusion allows you to more easily install the device elsewhere in your network.

If you are using the automatic satellite conversion feature to convert devices into satellite devices in your Junos Fusion, remove the FPC slot ID to the satellite device from the automatic satellite conversion configuration before using this command to install the Junos OS software. You can update the automatic satellite conversion feature using the `set chassis satellite-management auto-satellite-conversion satellite slot-id` configuration statement.

You must install a PXE version of compatible Junos OS to convert the satellite device running satellite software to a standalone device running Junos OS software on QFX5100 switches acting as satellite devices. The PXE version of Junos OS is the software that includes **pxe** in the Junos OS package name when it is downloaded from the Software Center—for example, the PXE image for Junos OS Release 14.1X53-D16 is named **install-media-pxe-qfx-5-14.1X53-D16.2.tgz**.

For Junos Fusion systems running Junos OS Release 17.2R1 and later, you must install a signed PXE version of Junos OS to convert the satellite device running satellite software to a standalone device

running Junos OS software. The signed PXE Junos OS package name uses the format **install-media-pxe-qfx-5-*version*-domestic-signed.tgz**.

The device uses a factory-default configuration after the Junos OS installation is complete. No Junos OS configuration is modified and the previous Junos OS configuration is not restored after the Junos OS software installation.

Options

package-name	Specify the URL to the Junos OS image to install onto the satellite device.
fpc *fpc-slot*	Install the Junos OS software onto the satellite device with the specified FPC slot ID in the Junos Fusion.
device-alias *device-alias*	Install the Junos OS software onto the satellite device with the alias name in the Junos Fusion.
no-confirm	(Optional) Install the Junos OS software onto the satellite device immediately without further confirmation prompting.

Required Privilege Level

system-control

Sample Output

request chassis satellite install /var/tmp/jinstall-ex-4300-14.1X53-D16.1-domestic-signed.tgz fpc-slot 101 (EX4300 switch as satellite device)

```
user@aggregation-device> request chassis satellite install /var/tmp/jinstall-ex-4300-14.1X53-D16.1-domestic-signed.tgz fpc-slot 101
Response from device:
    Conversion Started
```

request chassis satellite install /var/tmp/install-media-pxe-qfx-5-14.1X53-D16.2.tgz fpc-slot 102 (QFX5100 switch as satellite device)

```
user@aggregation-device> request chassis satellite install /var/tmp/install-media-pxe-qfx-5-14.1X53-D16.2.tgz fpc-slot 102
```

```
Response from device:
    Conversion Started
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

- Configuring or Expanding a Junos Fusion Enterprise
- *Configuring Junos Fusion Provider Edge*

request chassis satellite interface

IN THIS SECTION

- Syntax | 14545
- Description | 14545
- Options | 14546
- Required Privilege Level | 14546
- Sample Output | 14546
- Release Information | 14546

Syntax

```
request chassis satellite interface interface-name
    device-mode satellite
```

Description

Change the device mode for a device.

This command is used to change a device into a satellite device for a Junos Fusion. After interconnecting a device to an aggregation device in a Junos Fusion, enter this command from the aggregation device to begin the manual satellite device conversion procedure.

Other configuration steps, such as configuring the cascade port and creating a satellite software upgrade group, must be completed before this command can be used to convert a device into a satellite device. See *Configuring Junos Fusion Provider Edge* or Configuring or Expanding a Junos Fusion Enterprise.

Options

interface-name Specify the name of the cascade port interface on the aggregation device that connects to the device that will be converted into a satellite device.

Required Privilege Level

system-control

Sample Output

request chassis satellite interface xe-0/0/1 device-mode satellite

```
user@aggregation-device> request chassis satellite interface xe-0/0/1 device-mode satellite
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

request chassis satellite login

IN THIS SECTION

- Syntax | 14547
- Description | 14547
- Options | 14547
- Required Privilege Level | 14548
- Sample Output | 14548
- Release Information | 14548

Syntax

```
request chassis satellite login
    <fpc-slot fpc-slot>
    <interface-name interface-name>
```

Description

Log in to the satellite device from the aggregation device.

This command is typically used to log in to the satellite device by expert users for debugging purposes. You can perform all configuration and administration tasks in a Junos Fusion from the aggregation device.

Options

fpc *fpc-slot*	Log in to the satellite device with the specified FPC slot ID.
interface-name *interface-name*	Log in to the satellite device connected to the specified interface. The *interface-name* is the cascade port on the aggregation device.

Required Privilege Level

system-control

Sample Output

request chassis satellite login fpc-slot 101

```
user@aggregation-device> request chassis satellite login fpc-slot 101
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

request chassis satellite reboot

IN THIS SECTION

- Syntax | 14549
- Description | 14549
- Options | 14549
- Required Privilege Level | 14549
- Sample Output | 14549
- Sample Output | 14549
- Release Information | 14550

Syntax

```
request chassis satellite reboot
    <fpc-slot fpc-slot>
    <range range>
```

Description

Reboot the satellite device or devices from the aggregation device in a Junos Fusion.

Options

fpc *fpc-slot* Reboot the satellite device with the specified FPC slot identifier.

range *range* Reboot all satellite devices in a range of FPC slot identifiers.

For instance, you can reboot the satellite devices using FPC slot identifiers 101, 102, and 103 by entering a *range* of `101-103`.

Required Privilege Level

system-control

Sample Output

request chassis satellite reboot fpc 101

```
user@aggregation-device> request chassis satellite reboot fpc 101
```

Sample Output

request chassis satellite reboot range 101-103

```
user@aggregation-device> request chassis satellite reboot range 101-103
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

request chassis satellite restart

IN THIS SECTION

- Syntax | 14550
- Description | 14550
- Options | 14551
- Required Privilege Level | 14551
- Release Information | 14551

Syntax

```
request chassis satellite restart
    [fpc-slot fpc-slot | range range]
    <process-name>
```

Description

Restart a process on a satellite device or devices from the aggregation device in a Junos Fusion.

You would typically restart a process in a Junos Fusion for troubleshooting or debugging purposes.

This command is intended for use by expert users for debugging purposes.

Options

fpc ***fpc-slot*** Restart the specified process on the satellite device in the specified FPC slot ID only.

range Restart the process on the satellite devices in the specified range of FPC slot IDs only.

For instance, if you want to reboot the satellite devices using FPC slot IDs 101, 102, and 103, you can enter a *range* of `101-103`.

process-name Restart the specified process on the specified FPC slot ID or range of FPC slot IDs.

Required Privilege Level

system-control

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

request chassis satellite shell-command

IN THIS SECTION

- Syntax | 14552
- Description | 14552
- Options | 14552
- Required Privilege Level | 14552
- Release Information | 14552

Syntax

```
request chassis satellite shell-command
    [fpc-slot fpc-slot |range ]
    <remote-command>
```

Description

Run a UNIX shell command for a satellite device from the aggregation device in a Junos Fusion.

Options

fpc *fpc-slot*
Run the shell command on the satellite device using the specified FPC slot identifier only.

range
Run the shell command on the satellite devices in the specified range of FPC slot identifiers only.

For instance, you can run the shell command on the satellite devices in FPC slot identifiers 101, 102, and 103 by entering a *range* of `101-103`.

remote-command
Specify the UNIX shell command to run on the satellite device or devices.

Required Privilege Level

system-control

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

request chassis scg

IN THIS SECTION

- Syntax | 14553
- Syntax (TX Matrix and TX Matrix Plus Routers) | 14553
- Description | 14553
- Options | 14553
- Required Privilege Level | 14554
- Output Fields | 14554
- Sample Output | 14554
- Release Information | 14554

Syntax

```
request chassis scg (offline | online) slot slot-number
```

Syntax (TX Matrix and TX Matrix Plus Routers)

```
request chassis scg lcc number  (offline | online) slot slot-number
```

Description

(T Series routers only) Control the operation of the specified SONET Clock Generator (SCG).

Options

lcc *number* (TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

offline Take the SCG offline. When you change the SCG status to offline, the unit is not powered down.

online Bring the SCG online.

slot *slot-number* SCG slot number. Replace ***slot-number*** with 0 or 1.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis scg

```
user@host> request chassis scg online slot 0
Online initiated, use "show chassis environment scg" to verify
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show chassis environment scg | 17811

Configuring the Clock Source

T320 SONET Clock Generator (SCG) Description

request chassis sfb

IN THIS SECTION

- Syntax | 14555
- Syntax (MX10004 and MX10008) | 14555
- Description | 14556
- Options | 14556
- Required Privilege Level | 14556
- Output Fields | 14556
- Sample Output | 14556
- Release Information | 14557

Syntax

```
request chassis sfb (offline | online) slot slot-number
<all-members>
<local>
<member member-id>
```

Syntax (MX10004 and MX10008)

```
request chassis sfb
<(offline | online) slot slot-number>
<slot slot-number (offline | online)>
```

Description

Control the operation of the Switch Fabric Board (SFB). Configure the SFBs on a slot to online or offline.

command.

Options

all-members	(Optional) Control the operation of the SFB in all members of the Virtual Chassis configuration.
local	(Optional) Control the operation of the SFB in the local Virtual Chassis member.
member *member-id*	(Optional) Control the operation of the SFB in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.
offline	Take the Switch Fabric Board offline.
online	Bring the Switch Fabric Board online.
slot *slot-number*	Switch Fabric Board slot number. Replace *slot-number* with a value of 0 through 7.
plane *plane-number*	Switch Fabric Board plane number. For SFB, the *plane-number* value range is 0 through 3. For SFB2, the *plane-number* value range is 0 through 1.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis sfb

```
user@host> request chassis sfb offline slot
1
Backup SFB 1 cannot be set offline, backup RE is online
```

request chassis sfb (MX2010 Routers)

```
user@host> request chassis sfb offline slot 7
Offline initiated, use "show chassis sfb" to verify
```

request chassis sfb (MX2008 Routers)

```
user@host>request chassis sfb offline slot 1
Offline initiated, use "show chassis sfb" to verify
```

request chassis fabric sfb (MX10008 Routers)

```
user@host>request chassis fabric sfb slot 1 plane 0 offline
Offline initiated, use "show chassis fabric plane" to verify
```

request chassis sfb (MX10004 and MX10008 Routers)

```
user@host>request chassis sfb  slot 0 online
Online initiated, use "show chassis sfb" to verify
```

```
user@host>request chassis sfb  slot 0 offline
Offline initiated, use "show chassis sfb" to verify
```

Release Information

Command introduced in Junos OS Release 12.3.

`all-members`, `local`, and `member` *`member-id`* options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.

RELATED DOCUMENTATION

show chassis hardware | **18430**

show chassis sfb | **18823**

request chassis sfm

IN THIS SECTION

- Syntax | 14558
- Description | 14558
- Options | 14558
- Required Privilege Level | 14558
- Output Fields | 14559
- Sample Output | 14559
- Release Information | 14559

Syntax

```
request chassis sfm (offline | online | restart) slot slot-number
```

Description

(M40e and M160 routers only) Control the operation of the specified Switching and Forwarding Module (SFM).

Options

offline	Take the SFM offline.
online	Bring the SFM online.
restart	Restart the SFM.
slot *slot-number*	SFM slot number. Replace ***slot-number*** with a value from **0** through **3**.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis sfm (M40e)

```
user@host> request chassis sfm slot 1 restart
M40e router:
error: SFM 0 is transitioning to online state.
```

request chassis sfm (M160)

```
user@host> request chassis sfm slot 1 restart
M160 router:
Restart initiated, use "show chassis sfm" to verify
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show chassis sfm | **18831**

Configuring SFM Redundancy on M40e and M160 Routers

M40e Switching and Forwarding Module (SFM) Description

request chassis sfm master switch

IN THIS SECTION

Syntax | **14560**

- Description | 14560
- Options | 14560
- Additional Information | 14560
- Required Privilege Level | 14560
- Output Fields | 14561
- Sample Output | 14561
- Release Information | 14561

Syntax

```
request chassis sfm master switch
 <no-confirm>
```

Description

(M40e and M160 routers only) Control which Switching and Forwarding Module (SFM) is primary.

Options

no-confirm (Optional) Do not display a switch warning or query.

Additional Information

By default, the SFM in slot 0 (SFM0) is the primary and the SFM in slot 1 (SFM1) is the backup. If you use this command to change the primary, and then restart the chassis software for any reason, the primary reverts to the default setting. To change the default primary SFM, include the `sfm` statement at the `[edit chassis redundancy]` hierarchy level in the configuration. For more information, see the Junos OS Administration Library for Routing Devices.

All installed SFMs are always working together to forward packets. If an SFM fails, the other SFMs take over and traffic continues to flow uninterrupted.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis sfm primary switch

```
user@host> request chassis sfm master switch
warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between system forwarding module? [yes,no] (no) yes

Switch initiated, use “show chassis sfm” to verify
```

request chassis sfm primary switch no-confirm

```
user@host> request chassis sfm master switch no-confirm
Switch initiated, use “show chassis sfm” to verify
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show chassis sfm | 18831

Switching the Global Primary and Backup Roles in a Virtual Chassis Configuration

request chassis sib

IN THIS SECTION

- Syntax | 14562
- Syntax (TX Matrix Router) | 14562
- Syntax (TX Matrix Plus Router) | 14562
- Description | 14563
- Options | 14563
- Required Privilege Level | 14564
- Output Fields | 14564
- Sample Output | 14564
- Release Information | 14565

Syntax

```
request chassis sib (offline | online) slot slot-number
```

Syntax (TX Matrix Router)

```
request chassis sib (all-chassis | lcc number | scc) (offline | online) slot slot-number
(start-receiver number | stop-receiver number)
```

Syntax (TX Matrix Plus Router)

```
request chassis sib (all-lcc | f13 slot-number | f2s sib-slot/sib-f2s-slot-number |lcc number  |
(offline | online) slot slot-number)
```

Description

(M320 routers and T Series routers only) Control the operation of the specified Switch Interface Board (SIB).

Options

all-chassis (TX Matrix routers only) Control the status of the specified SIB.

all-lcc (TX Matrix Plus router only) On TX Matrix Plus router, control the operation of the SIB on all routers connected to the TX Matrix Plus router.

f13 ***slot-number*** Control the operation of F13 SIBs. Replace *slot-number* with a value `0`, `1`, `3`, `4`, `6`, `7`, `8`, `9`, `11`, or `12`.

f2s ***sib-slot/sib-f2s-slot-number*** (TX Matrix Plus routers only) (Optional) Control the operation of the SIB F2s. Replace *sib-slot* with a value from `0` through `4`, followed by a *sib-f2s-slot-number* value `0`, `2`, `4` or `6`.

lcc ***number*** (TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

scc (TX Matrix router only) TX Matrix router (switch-card chassis) on a routing matrix.

offline Take the SIB offline.

> **NOTE**: In PTX Series (PTX3000 and PTX5000) and T Series (T640 and T1600) Routers with active PFE interfaces, when the last SIB is taken offline, a message displays that if no SIB is brought online within 10 seconds, the system will take action to address the fabric null route condition. Taking all SIBs offline in these PTX Series or T Series Routers

with active PFE interfaces results in null route condition, and the software takes action to rectify this condition if it persists for more than 10 seconds. If these routers do not have active PFE interfaces, taking all SIBs offline does not result in null route condition, and the message is not displayed when the last active SIB is taken offline. For details on null route condition, see *Fabric Resiliency and Degradation*.

online Bring the SIB online.

slot *slot-number* SIB slot number. For the T320 router, replace *slot-number* with a value from 0 through 2. For the T640 router, TX Matrix router, and T1600 router in a routing matrix, replace *slot-number* with a value from 0 through 4.

start-receiver *number* (TX Matrix routers only) Start the SIB optical receiver. Replace *number* with a value from 0 through 3.

stop-receiver *number* (TX Matrix routers only) Stop the SIB optical receiver. Replace *number* with a value from 0 through 3.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis sib

```
user@host> request chassis sib slot 0 online
Online initiated, use "show chassis sibs" to verify
```

request chassis sib

```
user@host> request chassis sib f13 slot 0 offline
Offline initiated, use "show chassis sibs" to verify
```

Release Information

Command introduced before Junos OS Release 7.4.

`f13` and `f2s` options for the TX Matrix Plus router introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Configuring Junos OS to Upgrade and Downgrade Switch Interface Boards on a TX Matrix Router
- M320 SIB Description
- show chassis sibs | **18835**
- show chassis environment sib | **17849**

request chassis spmb restart

IN THIS SECTION

- Syntax | **14565**
- Syntax (MX2020, MX2010, and 2008 Routers) | **14566**
- Syntax (TX Matrix Router) | **14566**
- Syntax (TX Matrix Plus Router) | **14566**
- Description | **14566**
- Options | **14566**
- Required Privilege Level | **14567**
- Output Fields | **14567**
- Sample Output | **14567**
- Release Information | **14568**

Syntax

```
request chassis spmb restart slot slot-number
```

Syntax (MX2020, MX2010, and 2008 Routers)

```
request chassis spmb restart slot slot-number
<all-members>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
request chassis spmb restart (lcc number | scc) slot slot-number
```

Syntax (TX Matrix Plus Router)

```
request chassis spmb restart (lcc number | sfc number) slot slot-number
```

Description

Restart the specified Switch Processor Mezzanine Board (SPMB) on the Control Board (CB).

Options

all-members (MX2010, MX2020, and MX2008 routers only) (Optional) Restart the SPMB on the CB in all members of the Virtual Chassis configuration.

lcc *number* (TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX2010, MX2020, and MX2008 routers only) (Optional) Restart the SPMB on the CB in the local Virtual Chassis member.

member ***member-id*** (MX2010, MX2020, and MX2008 routers only) (Optional) Restart the SPMB on the CB in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.

scc (TX Matrix routers only) TX Matrix router (switch-card chassis) in the routing matrix.

sfc ***number*** (TX Matrix Plus routers only) The switch-fabric chassis number of the TX Matrix Plus router. Replace the *number* variable with a value `0`.

slot ***slot-number*** The SPMB slot number. Replace *slot-number* with `0` or `1`.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis spmb restart

```
user@host> request chassis spmb restart slot
0
```

request chassis spmb restart (MX2010 Router)

```
user@host> request chassis spmb restart slot 0
Restart initiated, use "show chassis spmb" to verify
```

request chassis spmb restart (MX2008 Router)

```
user@host>request chassis spmb restart slot 0
Restart initiated, use "show chassis spmb" to verify
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option for the TX Matrix Plus router introduced in Junos OS Release 9.6.

`all-members`, `local`, and `member` *`member-id`* options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.

RELATED DOCUMENTATION

show chassis spmb | 18853

show chassis spmb sibs | 18867

request chassis ssb master switch

IN THIS SECTION

- Syntax | 14569
- Description | 14569
- Options | 14569
- Additional Information | 14569
- Required Privilege Level | 14569
- Output Fields | 14569
- Sample Output | 14570
- Release Information | 14570

Syntax

```
request chassis ssb master switch
<no-confirm>
```

Description

(M20 router only) Control which System and Switch Board (SSB) is primary.

Options

no-confirm (Optional) Do not request confirmation for the switch.

Additional Information

By default, the SSB in slot 0 (SSB0) is the primary and the SSB in slot 1 (SSB1) is the backup. If you use this command to change the primary, and then restart the chassis software for any reason, the primary reverts to the default setting. To change the default primary SSB, include the `ssb` statement at the [**edit chassis redundancy**] hierarchy level in the configuration. For more information, see the Junos OS Administration Library for Routing Devices.

The configurations on the two SSBs do not have to be the same, and they are not automatically synchronized. If you configure both SSBs as primaries, when the chassis software restarts for any reason, the SSB in slot 0 becomes the primary and the one in slot 1 becomes the backup.

The switchover from the primary SSB to the backup SSB is immediate. The SSB takes several seconds to reinitialize the Flexible PIC Concentrators (FPCs) and restart the PICs. The interior gateway protocol (IGP) and BGP convergence times depend on the specific network environment.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis ssb primary switch

```
user@host> request chassis ssb master switch
warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between system switch boards ? [yes,no] (no) yes

Switch initiated, use “show chassis ssb” to verify
```

request chassis ssb primary switch no-confirm

```
user@host> request chassis ssb master switch no-confirm
Switch initiated, use “show chassis ssb” to verify
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show chassis ssb | 18876

request chassis synchronization mode

IN THIS SECTION

- Syntax | 14571
- Description | 14571
- Options | 14571
- Required Privilege Level | 14571
- Output Fields | 14571

- Sample Output | 14572
- Sample Output | 14572
- Sample Output | 14572
- Release Information | 14572

Syntax

```
request chassis synchronization mode (free-run | holdover | auto-select)
```

Description

(MX5-T, MX10-T, MX40-T, MX80, MX80-T, MX240, MX480, and MX960 routers only) Change the chassis synchronization source used for Synchronous Ethernet configuration.

Options

freerun	Change chassis synchronization to free-run mode.
holdover	Change chassis synchronization to holdover mode.
auto-select	Change chassis synchronization to auto-select mode.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the current status of your request is displayed.

- **Not configured**—Indicates that the source is not configured.
- **Present**—Indicates that the source is configured and present.
- **Qualified**—Indicates that the source is being used for synchronization.

Sample Output

request chassis synchronization mode freerun

```
user@host> request chassis synchronization mode freerun
mode is freerun, status: qualified
```

Sample Output

request chassis synchronization mode holdover

```
user@host> request chassis synchronization mode holdover
mode is holdover, status: qualified
```

Sample Output

request chassis synchronization mode auto-select

```
user@host> request chassis synchronization mode auto-select
mode is auto-select, status: qualified
```

Release Information

Command introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

Synchronous Ethernet Overview

Configuring an External Clock Synchronization Interface for MX Series Routers

request (Chassis Synchronization Source)

IN THIS SECTION

- Syntax | 14573
- Hierarchy Level | 14573
- Description | 14573
- Options | 14573
- Required Privilege Level | 14574
- Release Information | 14574

Syntax

```
request (force-switch | lockout);
```

Hierarchy Level

```
[edit chassis synchronization source (bits | gps)]
[edit chassis synchronization source interfaces]
```

Description

Configure an interface with an upstream clock source to force a switch to another source, provided that the source is enabled and not locked out, or to lockout a source from the clock selection process.

Options

force-switch Force a switch to the source provided that the source is enabled and not locked out. Only one configured source may be force-switched.

lockout A lockout may be configured for any source. When a lockout is configured for a source, that source will not be considered by the selection process.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

External Clock Synchronization Overview for ACX Series Routers

Configuring External Clock Synchronization for ACX Series Routers

Synchronous Ethernet Overview

show chassis synchronization

request chassis synchronization switch

IN THIS SECTION

- Syntax | 14575
- Syntax (M Series, T Series) | 14575
- Syntax (PTX Series) | 14575
- Description | 14575
- Options | 14575
- Required Privilege Level | 14576
- Output Fields | 14576
- Sample Output | 14576
- Release Information | 14576

Syntax

```
request chassis synchronization switch
```

Syntax (M Series, T Series)

```
request chassis synchronization switch (external-a | external-b)
```

Syntax (PTX Series)

```
request chassis synchronization switch (bits-a | bits-b | fpc-slot-number | gps-0 | gps-1 )
```

Description

(M320, M40e, M120, T320, T640, and T1600 routers and PTX Packet Transport Routers only) Change the external clock source used for chassis synchronization.

Options

external-a (Routing matrix only) Change the synchronization source to external source A.

external-b (Routing matrix only) Change the synchronization source to external source B.

bits-a (PTX Series only) Change the synchronization source to the BITS external source A.

bits-b (PTX Series only) Change the synchronization source to the BITS external source B.

fpc-*slot-number* (PTX Series only) Change the synchronization source to an FPC in the slot specified. For the PTX5000 Packet Transport Router, replace *slot-number* with a value from 0 through 7.

gps-0-10mhz (PTX Series only) Change the synchronization source to the 10 MHz GPS source on CCG port 0.

gps-0-5mhz (PTX Series only) Change the synchronization source to the 5 MHz GPS source on CCG port 0.

gps-1-10mhz (PTX Series only) Change the synchronization source to the 10 MHz GPS source on CCG port 1.

gps-1-5mhz (PTX Series only) Change the synchronization source to the 5 MHz GPS source on CCG port 1.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request. **Not configured** indicates that the source is not configured. **Present** indicates that the source is configured and present. **Qualified** indicates that the source is being used for synchronization.

Sample Output

request chassis synchronization switch (M Series, T Series)

```
user@host> request chassis synchronization switch external-a
switching to external-a, status: qualified
```

request chassis synchronization switch (PTX Series)

```
user@host> request chassis synchronization switch fpc-2
switching to fpc-2, status: qualified
```

Release Information

Command introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

show chassis synchronization | 18879

Configuring Clock Synchronization Interface on MX Series Routers

Supported Time Synchronization Standard

request chassis system-mode

IN THIS SECTION

- Syntax | 14577
- Description | 14577
- Options | 14577
- Required Privilege Level | 14578
- Sample Output | 14578
- Release Information | 14578

Syntax

```
request chassis system-mode
<mode-2x100G>
<default-mode>
```

Description

Set the operating mode of the ports to 100-Gigabit Ethernet mode on an EX4300-48MP switch. The `request chassis system-mode` command sets the ports to 100-Gigabit Ethernet mode by overriding the `set chassis fpc` *`fpc-slot`* `pic` *`pic-slot`* `port` *`port-number`* `speed` *`speed value`* command. This command restarts the DC packet forwarding engine (DCPFE), disables the virtual chassis ports (VCP) and brings up both ports of the uplink module to operate at 100Gbps speed. Use the `show chassis system-mode` command to verify configured system mode and the `request chassis system-mode default-mode` command to return to the default operating mode.

Options

`mode-2x100G` PIC mode configuration is 2x100Gbps. This will restart PFE and disable VCP ports.

default-mode Default-mode is the oversubscribed mode. This will restart PFE.

Required Privilege Level

view

Sample Output

request chassis system-mode

```
user@switch> request chassis system-mode mode-2x100G (mode - 2x100G on EX4300-48MP)
System-mode has changed. This will restart PFE and disable VCP ports.
```

request chassis system-mode (EX4300-48MP)

```
user@switch> request chassis system-mode default-mode
System-mode has changed. This will restart PFE.
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Uplink Modules in EX4300 Switches

Setting the Operating Mode on a 2-Port 40-Gigabit Ethernet QSFP+/100-Gigabit Ethernet QSFP28 Uplink Module

request (Clock Synchronization)

IN THIS SECTION

- Syntax | 14579
- Hierarchy Level | 14579
- Description | 14579
- Options | 14579
- Required Privilege Level | 14580
- Release Information | 14580

Syntax

```
request (force-switch lockout);
```

Hierarchy Level

```
[edit chassis synchronization source interfaces (external-a | external-b | interface interface-name)]
[edit chassis synchronization source interfaces (external | interface interface-name)]
[edit chassis synchronization source interfaces (external-0/0 | external-1/0 | interface interface-name)]
```

Description

Specify the clock selection request criterion.

Options

force-switch Force switching to a clock source, provided the clock source is enabled and not locked out. Only one configured source may be force-switched.

lockout Clock source is not considered by the selection process. Lockout may be configured for any source.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2 R4.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

request component login

IN THIS SECTION

- Syntax | 14580
- Description | 14581
- Options | 14581
- Required Privilege Level | 14581
- Sample Output | 14581
- Release Information | 14582

Syntax

```
request component login component-name
```

Description

(QFabric systems only) Log in to a QFabric system component. To gain access to individual components by way of the `request component login` command, you must first provide the `qfabric-admin` or `qfabric-operator` class privilege to your user (for more information, see: remote-debug-permission).

Options

component-name	Specify the QFabric system component to which you wish to log in.

Required Privilege Level

admin

Sample Output

The three sample output displays show the results of attempts to log in to Node device EE3093. The results differ depending on the privilege level assigned to the user.

request component login (with qfabric-admin Privileges)

```
admin@qfabric> request component login EE3093
Warning: Permanently added 'qfabric-node-ee3093,192.0.2.0' (RSA) to the list of known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
qfabric-admin@node-ee3093> ?
Possible completions:
  clear                Clear information in the system
  file                 Perform file operations
  help                 Provide help information
  load                 Load information from file
  monitor              Show real-time debugging information
  mtrace               Trace multicast path from source to receiver
  op                   Invoke an operation script
  ping                 Ping remote target
  quit                 Exit the management session
  request              Make system-level requests
  restart              Restart software process
  save                 Save information to file
  set                  Set CLI properties, date/time, craft interface message
```

```
  show                  Show system information
  ssh                   Start secure shell on another host
  start                 Start shell
  telnet                Telnet to another host
  test                  Perform diagnostic debugging
  traceroute            Trace route to remote host{master}
qfabric-admin@node-ee3093>
```

request component login (with qfabric-operator Privileges)

```
operator@qfabric> request component login EE3093
Warning: Permanently added 'qfabric-node-EE3093,192.0.2.0' (RSA) to the list of known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
qfabric-operator@node-EE3093> ?
Possible completions:
  file                  Perform file operations
  help                  Provide help information
  load                  Load information from file
  op                    Invoke an operation script
  quit                  Exit the management session
  request               Make system-level requests
  save                  Save information to file
  set                   Set CLI properties, date/time, craft interface message
  show                  Show system information
  start                 Start shell
  test                  Perform diagnostic debugging
{master}
qfabric-operator@node-ee3093>
```

request component login (with qfabric-user Privileges)

```
user0@qfabric> request component login EE3093
error: User user0 does not have sufficient permissions to login to device ee3093
```

Release Information

Command introduced in Junos OS Release 14.1X53-D20.

request csds add-vsrx

IN THIS SECTION

- Syntax | 14583
- Description | 14583
- Options | 14583
- Required Privilege Level | 14583
- Sample Output | 14584
- Release Information | 14584

Syntax

```
request csds add-vsrx csds-instance-id csds-instance-id image image-name
```

Description

The command copies the vSRX Virtual Firewall image from Junos Node Unifier (JNU) controller to Junos Device Manager (JDM) in the Connected Security Distributed Services (CSDS) architecture. The command also starts the virtual network functions (VNFs).

Options

csds-instance-id ***csds-instance-id*** Specify CSDS instance ID.

image ***image-name*** Specify vSRX Virtual Firewall image name.

Required Privilege Level

maintenance

Sample Output

request csds add-vsrx

```
user@host> request csds add-vsrx csds-instance-id 1 image /var/tmp/junos-vsrx3-x86-64-24.4.qcow2

Starting to copy the VSRX package from MX to server
Finished Copying the VSRX package from Controller to host
Successfully updated the VSRX image path in platform_attributes ...
Starting VNfs
vnf0 Started Successfully
```

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

No Link Title

request csds authenticate-host

IN THIS SECTION

- Syntax | 14585
- Description | 14585
- Options | 14585
- Required Privilege Level | 14585
- Sample Output | 14585
- Release Information | 14585

Syntax

```
request csds authenticate-host csds-instance-id csds-instance-id
```

Description

The command sets up a one-way key exchange between the Junos Node Unifier (JNU) controller and the host. This process provides the controller with passwordless SSH access to the server. As a result, the controller can operate Junos Device Manager (JDM) remotely in the Connected Security Distributed Services (CSDS) framework.

Options

csds-instance-id *csds-instance-id*	Specify CSDS instance ID.

Required Privilege Level

maintenance

Sample Output

request csds authenticate-host

```
user@host> request csds authenticate-host csds-instance-id 3

The authenticity of host '10.157.79.104 (10.157.79.104)' can't be established.
ED25519 key fingerprint is SHA256:$ABC123.
This key is not known by any other names.
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '10.157.79.104' (ED25519) to the list of known hosts.
root@10.157.79.104's password:
```

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

No Link Title

request csds delete-vsrx

IN THIS SECTION

- Syntax | 14586
- Description | 14586
- Options | 14586
- Required Privilege Level | 14587
- Release Information | 14587

Syntax

```
request csds delete-vsrx csds-instance-id csds-instance-id
```

```
request csds delete-vsrx csds-instance-id csds-instance-id force
```

Description

Remove currently running VNFs. The command deletes the running vSRX Virtual Firewall in Connected Security Distributed Services (CSDS) architecture in JNU topology.

Options

csds-instance-id *csds-instance-id* Specify CSDS instance ID.

force Perform operation by discarding any uncommitted changes.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

No Link Title

request csds extract-vsrx-keys

IN THIS SECTION

- Syntax | 14587
- Description | 14588
- Options | 14588
- Required Privilege Level | 14588
- Sample Output | 14588
- Release Information | 14589

Syntax

```
request csds extract-vsrx-keys csds-instance-id csds-instance-id
```

```
request csds extract-vsrx-keys csds-instance-id csds-instance-id force
```

Description

Within the Connected Security Distributed Services (CSDS) architecture, the Junos Node Unifier (JNU) controller extracts public key information of vSRX Virtual Firewall spawned by the Junos Device Manager (JDM) using this command.

Options

csds-instance-id Specify CSDS instance ID.

force Perform operation by discarding any uncommitted changes.

Required Privilege Level

maintenance

Sample Output

request csds extract-vsrx-keys

```
user@host> request csds extract-vsrx-keys csds-instance-id 5

Perform Extracting the Keys on vnf0
The authenticity of host '10.157.80.183 (10.157.80.183)' can't be established.
ED25519 key fingerprint is SHA256:$ABC123.
This key is not known by any other names.
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '10.157.80.183' (ED25519) to the list of known hosts.
Performing xml-commit-sync to add the keys to jnuadmin user account
xml-commit-sync is successfull added keys of jnuadmin user account
Added known host entry of VSRX to Controller
Successfully appended the controller known host entry..
Successfully appended the backup re known host entry..
Created symlink for vnfs ...
Extracting VSRX keys on vnf0 is Successfull
Restarting JNUD on vnf0...
Restarted JNUD on  vnf0...
```

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

No Link Title

request csds jdm

IN THIS SECTION

- Syntax | 14589
- Description | 14590
- Options | 14590
- Required Privilege Level | 14590
- Sample Output | 14590
- Release Information | 14592

Syntax

```
request csds jdm add csds-instance-id csds-instance-id image vsrx-image-with-path
```

```
request csds jdm add csds-instance-id csds-instance-id image vsrx-image-with-path force
```

```
request csds jdm delete csds-instance-id csds-instance-id
```

```
request csds jdm delete csds-instance-id csds-instance-id force
```

Description

The command either installs a new Junos Device Manager (JDM) package or removes currently installed JDM package in Junos Node Unifier (JNU) for Connected Security Distributed Services (CSDS) architecture.

Options

add Install a JDM package.

delete Remove currently installed JDM package.

csds-instance-id *csds-instance-id* Specify CSDS instance ID.

force Perform operation by discarding any uncommitted changes.

Required Privilege Level

maintenance

Sample Output

request csds jdm add

```
user@host> request csds jdm add csds-instance-id 3 image /var/tmp/csds-jdm-24.4-
I.20241027.0.2259.x86_64.deb

Starting to validate the Package
Finished validating the Package
Starting to copy the DEB package from Controller to host
Finished Copying the DEB package from Controller to host
Starting to install the JDM DEB package
Selecting previously unselected package csds-jdm.
(Reading database ... 248607 files and directories currently installed.)
Preparing to unpack .../csds-jdm-24.4-I.20241027.0.2259.x86_64.deb ...
Detailed log of jdm setup saved in /var/log/csds-jdm-setup.log
Server Profile:IntelGold
Interface selected: eno8303np0
Unpacking csds-jdm (24.4-I.20241027.0.2259) ...
Setting up csds-jdm (24.4-I.20241027.0.2259) ...
```

```
Setup host for jdm...
Interface detected:eno8303np0 carrier:1
Setup kernel cmdline for host
kernel cmdline OK!
Done Setup host for jdm
Installing /juniper/.tmp-jdm-install/juniper_ubuntu_rootfs.tgz...
Configure /juniper/lxc/jdm/jdm1/rootfs...
etc/ssh/ssh_host_rsa_key already exists.
Overwrite (y/n)? etc/ssh/ssh_host_ecdsa_key already exists.
Overwrite (y/n)? Configure /juniper/lxc/jdm/jdm1/rootfs DONE
Setup Junos cgroups...Skipped setting cgroups
Done
Created symlink /etc/systemd/system/multi-user.target.wants/jdm.service → /lib/systemd/system/
jdm.service.
Done Setup jdm
Processing triggers for libc-bin (2.35-0ubuntu3.8) ...
Finished installing the JDM DEB package: csds-jdm-24.4-I.20241027.0.2259.x86_64.deb
JDM Installation Successful !
Installing Mellanox driver !
OneWay Key Exchange from MX to JDM is Successfull...
Successfully updated the MX root and CSDS public key on the host.
Successfully updated the ipprefix_range and controller_ip to host
Successfully appended the JDM known host entry to the controller.
Performing xml-commit-sync to add the keys to jnuadmin user account
xml-commit-sync is successfull added keys of jnuadmin user account
Successfully appended the controller known host entry to the JDM side.
Thu Nov 14 11:46:58 AM UTC 2024
Exisitng Profile:IntelGold

Starting JDM

JDM Starting Successfully !
```

request csds jdm delete

```
user@host> request csds jdm delete csds-instance-id 0

JDM is currently running with version: csds-jdm-24.4-I.20241120.0.0829.x86_64
Warning: Uninstallation of JDM will result into deletion of GNFs
Do you still want to uninstall the JDM package ? Confirm (y/n) [yes,no] (no) yes
```

```
Package Uninstallation successful for package: csds-jdm
```

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

No Link Title

request csds sync-controller

IN THIS SECTION

- Syntax | 14592
- Description | 14593
- Options | 14593
- Required Privilege Level | 14593
- Sample Output | 14593
- Release Information | 14594

Syntax

```
request csds sync-controller other-controller-ip other-controller-ip
```

```
request csds sync-controller other-controller-ip other-controller-ip force
```

Description

The command synchronizes the primary controller with the other controller in dual controller Junos Node Unifier (JNU) topology in Connected Security Distributed Services (CSDS) architecture. Run the command from the second controller.

Options

other-controller-ip ***other-controller-ip*** Specify the other controller IP in dual node JNU controller topology..

force Perform operation by discarding any uncommitted changes.

Required Privilege Level

maintenance

Sample Output

request csds sync-controller

```
user@host> request csds sync-controller other-controller-ip 10.205.163.102

The authenticity of host '10.205.163.102 (10.205.163.102)' can't be established.
ED25519 key fingerprint is SHA256:$ABC123.
This key is not known by any other names.
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '10.205.163.102' (ED25519) to the list of known hosts.
(root@10.205.163.102) Password:
Please confirm authentication from other controller ? Confirm (y/n) [yes,no] (no) yes
Perform Two way controller Key exchange ...
Perform CSDS config sync from other controller ...
Successfully synced csds config from other controller ...
Successfully synced ssh keys between two controllers...
Perform updating the controller keys at satellite side
Perform updating controller keys via xml-commit ...
Perform xml-commit on all the vnfs
Successsfully done the xml-commit on all the vnfs
Performing the xml-commit of satellite conf on vnfs
Successsfully done xml-commit on 10.209.11.87
```

```
Restarting JNUD at vnfs
Successfully done the xml-commit of satellite conf on vnfs
Successfully updated the controller keys at satellite side
Restarting the JDM ..
```

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

No Link Title

request debug information

IN THIS SECTION

- Syntax | 14594
- Description | 14595
- Options | 14595
- Required Privilege Level | 14595
- Sample Output | 14595
- Release Information | 14595

Syntax

```
request debug information
<mgd>
```

Description

Consolidate debugging information for the component specified. You can issue the command without a component specified

Using this command speeds up the initial information-gathering phase of debugging. The output includes the combined output of multiple member commands.

Options

mgd (Optional) Display debug information for the Mangement process (mgd) workflows.

Required Privilege Level

view

Sample Output

request debug information mgd

```
user@host> request debug information mgd
```

Release Information

Command introduced in Junos OS Release 19.4R1 with `mgd` component.

request dhcp client renew

IN THIS SECTION

- Syntax | 14596
- Description | 14596
- Options | 14596
- Required Privilege Level | 14596

Output Fields | 14596
Release Information | 14596

Syntax

```
request dhcp client renew
[all|interface <interface-name>]
routing-instance <routing-instance-name>
```

Description

Initiates a renew request for the specified clients if they are in the bound state.

Options

all	Initiate renew requests for all DHCP clients. If you specify a routing instance, renew requests are initiated for all DHCP clients within that routing instance.
interface <interface-name>	Initiate renew requests for DHCP clients on the specified interface.
routing-instance <routing-instance-name>	Initiate renew requests for DHCP clients in the specified routing instance. If you do not specify a routing instance, renew requests are initiated on the default routing instance.

Required Privilege Level

view

Output Fields

This command produces no output.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

request dhcpv6 client renew | 14604

request dhcp relay bulk-leasequery

IN THIS SECTION

- Syntax | 14597
- Description | 14597
- Options | 14598
- Required Privilege Level | 14599
- Output Fields | 14599
- Sample Output | 14599
- Release Information | 14599

Syntax

```
request dhcp relay bulk-leasequery
<all-configured-ip>
<client-id | ipv4-address | mac-address>
<relay-id relay-id>
<remote-id remote-id>
<server-address server-address | server-group group-name>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Initiate a bulk leasequery operation to update the binding information for multiple subscribers by requesting that the DHCPv4 relay agent send the bulk leasequery message to the configured DHCPv4 servers. The DHCP bulk leasequery feature must be configured on the DHCPv4 relay agent. If you do not specify a `server-address` or `server-group`, then the query is sent to all DHCPv4 servers known to the relay agent.

Options

NOTE: By default, the bulk leasequery operation uses the relay ID of the DHCPv4 relay agent if you do not explicitly specify any of the following options: *client-id*, *ipv4-address*, *mac-address*, *relay-id*, or *remote-id*.

all-configured-ip	Update binding information for all IP addresses configured in the local server. The information is returned regardless of whether the IP addresses are part of a currently active binding and regardless of remote Id or relay ID. This enables the relay agent to update its database with all address changes that occurred after some point in time.
client-id	Update binding information for clients identified by the specified DHCPv4 client option (option 61).
ipv4-address	Update binding information for the most recent client that was assigned that IP address.
logical-system ***logical-system-name***	Specify an optional logical system for the DHCP server being queried. The default logical system is used by default.
mac-address	Update binding information for the most recent client that has that MAC address.
relay-id ***relay-id***	Update binding information for all currently active leases assigned to the client that has the specified Relay Agent Identifier suboption (suboption 12) of the DHCP relay agent information option (option 82). Suboption 12 uniquely identifies a relay agent within an administrative domain, which consists of all DHCP servers and relay agents that communicate with each other.
remote-id ***remote-id***	Update binding information for all currently active leases assigned to clients that use that Agent Remote ID (suboption 2) of the DHCP relay agent information option (option 82).
routing-instance ***routing-instance-name***	Specify an optional routing instance for the DHCP server being queried. The default routing instance is used by default.
server-address ***server-address***	Specify the address of a DHCP local server to query.
server-group ***group-name***	Specify the name of a group of DHCP local servers to query.

Required Privilege Level

view

Output Fields

When you enter this command, DHCP relay agent initiates the bulk leasequery operation.

Sample Output

request dhcp relay bulk-leasequery

```
user@host> request dhcp relay bulk-leasequery 4/1224 server-address 192.168.10.10
```

Release Information

Command introduced in Junos OS Release 16.1.

`all-configured-ip`, `relay-id`, and `remote-id` options added in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

DHCP Leasequery Methods

request dhcp relay leasequery

IN THIS SECTION

- Syntax | 14600
- Description | 14600
- Options | 14600
- Required Privilege Level | 14601
- Output Fields | 14601
- Sample Output | 14601

Release Information | 14601

Syntax

```
request dhcp relay leasequery (client-id | ipv4-address | mac-address)
gateway-address giaddr;
<server-address address | server-group group-name>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Initiate a leasequery operation to update the binding information for a specific subscriber by requesting that DHCPv4 relay agent send the leasequery message to the configured DHCP servers. The DHCP leasequery feature must be configured on the DHCPv4 relay agent. If you do not specify a `server-address` or `server-group`, then the query is sent to all DHCPv4 servers known to the relay agent.

Options

NOTE: You must configure at least one of *client-id*, *ipv4-address*, or *mac-address*.

client-id	Use the DHCPv4 client identifier option (option 61) to identify the client whose binding information is requested.
gateway-address *giaddr*	Gateway address of the relay agent making the request; the value in the giaddr field of the DHCP packet. You must always specify a gateway address in your request.
ipv4-address	Use the IPv4 address to identify the client whose binding information is requested.
logical-system *logical-system-name*	Specify an optional logical system for the DHCP servers being queried. The default logical system is used by default.
mac-address	Use the MAC address to identify the client whose binding information is requested.

routing-instance *routing-instance-name*	Specify an optional routing instance for the DHCPv4 servers being queried. The default routing instance is used by default.
server-address *server-address*	Specify the IP address of the DHCPv4 local server to query.
server-group *group-name*	Specify the name of a group of DHCPv4 local servers to query.

Required Privilege Level

view

Output Fields

When you enter this command, DHCP relay agent initiates the leasequery operation.

Sample Output

request dhcp relay leasequery

```
user@host> request dhcp relay leasequery 192.168.25.25 server-group DHCPgroup-10
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

DHCP Leasequery Methods

request dhcp server reconfigure

IN THIS SECTION

Syntax | 14602

- Description | 14602
- Options | 14602
- Required Privilege Level | 14603
- Output Fields | 14603
- Sample Output | 14603
- Release Information | 14603

Syntax

```
request dhcp server reconfigure (all | address | interface interface-name | logical-system logical-system-name | routing-instance routing-instance-name)
```

Description

Initiate reconfiguration processing for the specified DHCP clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the `clear dhcp server binding` command.

When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfiguring state and the local server sends a forcerenew message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the forcerenew message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the `clear-on-abort` statement to configure the client to be cleared when reconfiguration fails.

Options

all	Initiate reconfiguration for all DHCP clients.
address	Initiate reconfiguration for DHCP client with the specified IP address or MAC address.
interface ***interface-name***	Initiate reconfiguration for all DHCP clients on this logical interface (clients whose initial login requests were received over the specified interface).

> **NOTE**: You cannot use the `interface` *`interface-name`* option with the `request dhcp server reconfigure` command for DHCP passive clients (clients that are added as a result of DHCP snooped packets). For passive clients, the interface is not guaranteed to be the next-hop interface to the client, as is the case for active clients.

logical-system *logical-system-name*	Initiate reconfiguration for all DHCP clients on the specified logical system.
routing-instance *routing-instance-name*	Initiate reconfiguration reconfigured for all DHCP clients in the specified routing instance.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request dhcp server reconfigure

```
user@host> request dhcp server reconfigure interface fe-0/0/0.100
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview

request dhcpv6 client renew

IN THIS SECTION

- Syntax | 14604
- Description | 14604
- Options | 14604
- Required Privilege Level | 14605
- Output Fields | 14605
- Release Information | 14605

Syntax

```
request dhcpv6 client renew
[all | interface interface-name]
routing-instance <routing-instance-name>
```

Description

Initiate a renew request for the specified DHCPv6 clients if they are in the bound state.

Options

all	Initiate renew requests for all DHCPv6 clients. If you specify a routing instance, renew requests are initiated for all DHCPv6 clients within that routing instance.
interface-name *interface-name*	Initiate renew requests for DHCPv6 clients on the specified interface.
routing-instance *routing-instance-name*	Initiate renew requests for DHCPv6 clients in the specified routing instance. If you do not specify a routing instance, renew requests are initiated on the default routing instance.

Required Privilege Level

view

Output Fields

This command produces no output.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

request dhcpv6 relay bulk-leasequery

IN THIS SECTION

- Syntax | 14605
- Description | 14606
- Options | 14606
- Required Privilege Level | 14607
- Output Fields | 14607
- Sample Output | 14607
- Release Information | 14607

Syntax

```
request dhcpv6 relay bulk-leasequery
<client-id | ipv6-prefix| link-address ipv6-link-address | relay-id relay-id | remote-id remote-id>
<server-address server-address | server-group group-name>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Initiate a bulk leasequery operation to update the binding information for multiple subscribers by requesting that the DHCPv6 relay agent send the bulk leasequery message to the configured DHCPv6 servers. The DHCP bulk leasequery feature must be configured on the DHCPv6 relay agent. If you do not specify a `server-address` or `server-group`, then the query is sent to all DHCPv6 servers known to the relay agent.

Options

> **NOTE**: By default, the bulk leasequery operation uses the relay ID of the DHCPv6 relay agent if you do not explicitly specify any of the following options: *client-id*, *ipv6-prefix*, *ipv6-link-address*, *relay-id*, or *remote-id*.

client-id	Update binding information for clients identified by the DHCP Unique Identifier (DUID) specified DHCPv6 Client ID option (option 1).
ipv6-prefix	Update binding information for clients identified by the specified IPv6 prefix.
link-addr *ipv6-link-address*	Update binding information for clients associated with the specified IPv6 network segment. The link address is an address that a relay agent may have used in a Relay-Forward message.
logical-system *logical-system-name*	Specify an optional logical system for the clients. The default logical system is used by default.
relay-id *relay-id*	Update binding information for clients associated with the DHCPv6 relay identified by the DHCP Unique Identifier (DUID) in the specified Relay-ID option (option 53).
remote-id *remote-id*	Update binding information for clients associated with the specified Relay Agent Remote-ID option (option 37).
routing-instance *routing-instance-name*	Specify an optional routing instance for the clients. The default routing instance is used by default.
server-address *server-address*	Specify the IPv6 address of the DHCP local server to query.
server-group *group-name*	Specify the name of a group of DHCP local servers to query.

Required Privilege Level

view

Output Fields

When you enter this command, DHCPv6 relay agent initiates the bulk leasequery operation.

Sample Output

request dhcpv6 relay bulk-leasequery

```
user@host> request dhcpv6 relay bulk-leasequery LL0x1-00:00:65:03:01:02 server-address
2001:db8:3000:0:8001::5/128
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

DHCP Leasequery Methods

request dhcpv6 relay leasequery

IN THIS SECTION

- Syntax | 14608
- Description | 14608
- Options | 14608
- Required Privilege Level | 14609
- Output Fields | 14609
- Sample Output | 14609

Release Information | 14609

Syntax

```
request dhcpv6 relay leasequery (client-id | ipv6-prefix)
<server-address server-address | server-group group-name>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Initiate a leasequery operation to update the binding information for a specific subscriber by requesting that the DHCPv6 relay agent send the leasequery message to the configured DHCP servers. The DHCP leasequery feature must be configured on the DHCPv6 relay agent. If you do not specify a `server-address` or `server-group`, then the query is sent to all DHCP servers known to the relay agent.

Options

NOTE: You must configure at least one of *client-id* or *ipv6-prefix*.

client-id	Use the DHCPv6 Client ID option (option 1) to identify the client whose binding information is requested.
ipv6-prefix	Use the IPv6 prefix to identify the client whose binding information is requested.
logical-system *logical-system-name*	Specify an optional logical system for the DHCPv6 servers being queried. The default logical system is used by default.
routing-instance *routing-instance-name*	Specify an optional routing instance for the DHCPv6 servers being queried. The default routing instance is used by default.
server-address *server-address*	Specify the IPv6 address of the DHCPv6 local server to query.
server-group *group-name*	Specify the name of a group of DHCPv6 local servers to query.

Required Privilege Level

view

Output Fields

When you enter this command, DHCPv6 relay agent initiates the leasequery operation.

Sample Output

request dhcpv6 relay leasequery

```
user@host> request dhcpv6 relay leasequery 2001:db8:3000:0:8001::5/128
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

DHCP Leasequery Methods

request dhcpv6 server reconfigure

IN THIS SECTION

- Syntax | 14610
- Description | 14610
- Options | 14610
- Required Privilege Level | 14610
- Output Fields | 14611
- Sample Output | 14611
- Release Information | 14611

Syntax

```
request dhcpv6 server reconfigure (all | address | client-id | interface interface-name | logical-system
logical-system-name | routing-instance routing-instance-name | session-id)
```

Description

Initiate reconfiguration processing for the specified DHCPv6 clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the `clear dhcpv6 server binding` command.

When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfigure state and the local server sends a reconfigure message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the reconfigure message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the `clear-on-abort` statement to configure the client to be cleared when reconfiguration fails.

Options

all	Initiate reconfiguration for all DHCPv6 clients.
address	Initiate reconfiguration for DHCPv6 client with the specified IPv6 address.
client-id	Initiate reconfiguration for DHCPv6 client with the specified client ID.
interface *interface-name*	Initiate reconfiguration for all DHCPv6 clients on this logical interface (clients whose initial login requests were received over the specified interface).
logical-system *logical-system-name*	Initiate reconfiguration for all DHCPv6 clients on the specified logical system.
routing-instance *routing-instance-name*	Initiate reconfiguration reconfigured for all DHCPv6 clients in the specified routing instance.
session-id	Initiate reconfiguration for DHCPv6 client with the specified session ID.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request dhcpv6 server reconfigure

```
user@host> request dhcpv6 server reconfigure 2001db8::2/16
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview

request diagnostics tdr

IN THIS SECTION

- Syntax | 14612
- Description | 14612
- Options | 14612
- Required Privilege Level | 14612
- Output Fields | 14613
- Sample Output | 14613
- Release Information | 14613

Syntax

```
request diagnostics tdr (abort | start) interface interface-name
```

Description

Start a time domain reflectometry (TDR) diagnostic test on the specified interface. This test characterizes and locates faults on twisted-pair Ethernet cables. For example, it can detect a broken twisted pair and provide the approximate distance to the break. It can also detect polarity swaps, pair swaps, and excessive skew.

The TDR test is supported on the following switches and interfaces:

- EX2200, EX3200, EX3300, and EX4200 switches—RJ45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.
- EX6200 and EX8200 switches—RJ45 interfaces on line cards.

NOTE:

- Run the TDR test when the interface traffic under test is nil.
- TDR is a cable diagnostic feature that must be triggered only when a link is down.

You view the results of the TDR test with the `show diagnostics tdr` command.

Options

abort Stop the TDR test currently in progress on the specified interface. No results are reported, and previous results, if any, are cleared.

interface-name The name of the interface.

start Start a TDR test on the specified interface.

Required Privilege Level

maintenance

Output Fields

Table 112 on page 14613 lists the output fields for the `request diagnostics tdr` command. Output fields are listed in the approximate order in which they appear.

Table 112: request diagnostics tdr Output Fields

Field Name	Field Description
Test Status	Information about the status of the TDR test request: • **Admin Down *interface-name***—The interface is administratively down. The TDR test cannot run on interfaces that are administratively down. • **Interface *interface-name* not found**—The interface does not exist. • **Test successfully executed *interface-name***—The test has successfully started on the interface. You can view the test results with the `show diagnostics tdr` command. • **VCT not supported on *interface-name***—The TDR test is not supported on the interface.

Sample Output

request diagnostics tdr start interface ge-0/0/19

```
user@switch> request diagnostics tdr start interface ge-0/0/19

Interface TDR detail:
Test status                     : Test successfully executed  ge-0/0/19
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show diagnostics tdr | 19465

request extension-service restart-daemonize-app

IN THIS SECTION

- Syntax | 14614
- Description | 14614
- Options | 14614
- Required Privilege Level | 14614
- Output Fields | 14615
- Release Information | 14615

Syntax

```
request extension-service restart-daemonize-app application-name
```

Description

Restart a daemonized application running on a Junos device.

A daemonized application runs as a process in the background. An application configured to run as a daemonized process is automatically triggered upon commit. Restarting the application can assist you with debugging and troubleshooting.

Options

application-name	Name of application to be restarted.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Release Information

Command introduced in Junos OS and Junos OS Evolved Release 23.2R1.

request extension-service (start | stop)

IN THIS SECTION

- Syntax | 14615
- Description | 14615
- Options | 14615
- Required Privilege Level | 14616
- Output Fields | 14616
- Sample Output | 14616
- Release Information | 14616

Syntax

```
request extension-service (start | stop) application-name
<invoke-debugger cli>
```

Description

Start or stop a JET application running on a device running Junos OS.

Options

application-name	Name of application to be started or stopped.

invoke-debugger cli (Optional) Starts the extension service process in debugger mode.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

command-name

```
user@device> request extension-service start cmdline_args.py
Extension-service application 'cmdline_args.py' started with pid: 99418
```

Release Information

Command introduced in Junos OS Release 16.1.

request fabric administration director-group change-master

IN THIS SECTION

- Syntax | 14617
- Description | 14617
- Options | 14617
- Required Privilege Level | 14617
- Sample Output | 14617
- Release Information | 14617

Syntax

```
request fabric administration director-group change-master (director-device director-device-name)
```

Description

(QFabric systems only) Select a Director device to become the new primary device within a Director group. The specified device becomes the new primary Director device, and the previous primary Director device becomes a backup Director device.

Options

none	Change the device that controls the Director group. Assign the current backup Director device as the new primary and the current primary Director device as the backup.
director-device *director-device-name*	Specify which Director device should become the primary device within the Director group.

Required Privilege Level

admin

Sample Output

request fabric administration director-group change-primary

```
user@qfabric> request fabric administration director-group change-master
Do you intend to switchover mastership? [yes,no] (no) yes

Cluster master successfully switched
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

show fabric administration inventory director-group status | 19885

Performing the QFabric System Initial Setup on a QFX3100 Director Group

Understanding the Director Group

request fabric administration power-off interconnect-device

IN THIS SECTION

- Syntax | 14618
- Description | 14618
- Options | 14619
- Additional Information | 14619
- Required Privilege Level | 14619
- Output Fields | 14619
- Sample Output | 14619
- Release Information | 14620

Syntax

```
request fabric administration power-off interconnect--device (serial-id | alias-name)
```

Description

In QFX3000-G QFabric systems, you can use this command to power off QFX3008-I Interconnect devices.

This command systematically takes the device offline and gracefully shuts down the device while preserving system state information. A message appears on the console or console log, confirming that the operating system has stopped on the device.

Options

serial-id | alias-name Provide the serial ID or alias name of the device.

Additional Information

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request fabric administration power-off interconnect-device

```
user@qfabric request fabric administration power-off interconnect-device IC-12345
STEP 1 of 7 (Acquiring lock):
Acquiring lock to perform this operation
Acquired lock to perform this operation
STEP 2 of 7 (Performing pre-checks):
interconnect-device IC-12345 is online
STEP 3 of 7 (Gracefully offlining the interconnect):
Disabling fabric protocol on interconnect-device
Disabled fabric protocol on interconnect-device
STEP 4 of 7 (Waiting for convergence):
Waiting for convergence(this will take few minutes)
Convergence complete
STEP 5 of 7 (Powering-off the interconnect):
Now, powering-off the interconnect-device (this may take some time)
Powered-off the interconnect-device
STEP 6 of 7 (Updating inventory):
Removing the interconnect-device - IC-A0004 from inventory
interconnect-device - IC-12345 is removed from inventory
STEP 7 of 7 (Releasing lock):
Releasing the lock
```

Release Information

Command introduced in Junos OS Release 14.1X53-D15.

RELATED DOCUMENTATION

- request fabric administration power-off interconnect-device fpc | 14620
- request fabric administration power-off node-device | 14622
- show fabric administration inventory | 19878

request fabric administration power-off interconnect-device fpc

IN THIS SECTION

- Syntax | 14620
- Description | 14620
- Options | 14621
- Additional Information | 14621
- Required Privilege Level | 14621
- Output Fields | 14621
- Sample Output | 14621
- Release Information | 14622

Syntax

```
request fabric administration power-off interconnect--device (serial-id | alias-name) fpc slot-number
```

Description

In QFX3000-G QFabric systems, you can use this command to power off a 16-Port QSFP+ front card in a QFX3008-I Interconnect device.

A message appears on the console or console log, confirming that the operating system has stopped on the device.

Options

serial-id | alias-name	Provide the serial ID or alias name of the QFX3008-I Interconnect device.
slot *slot-number*	Provide the slot number of the 16-Port QSFP+ front card in the QFX3008-I Interconnect device. The range of slot numbers is 0 through 15

Additional Information

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request fabric administration power-off interconnect-device fpc 13

```
user@qfabric request fabric administration power-off interconnect-device IC-12345 fpc 13
STEP 1 of 6 (Acquiring lock):
Acquiring lock to perform this operation
Acquired lock to perform this operation
STEP 2 of 6 (Performing pre-checks):
interconnect-device IC-12345 is online
STEP 3 of 6 (Gracefully offlining the interconnect linecard):
Gracefully offlining the fpc
fpc is successfully offlined
STEP 4 of 6 (Waiting for convergence):
Waiting for convergence(this will take few minutes)
request fabric administration power-off interconnect-device IC-12345
Convergence complete
STEP 5 of 6 (Powering-off the interconnect linecard):
Now, powering-off the fpc on the interconnect-device
```

```
Powered-off the fpc on the interconnect-device
STEP 6 of 6 (Releasing lock):
Releasing the lock
```

Release Information

Command introduced in Junos OS Release 14.1X53-D15.

RELATED DOCUMENTATION

- request fabric administration power-off interconnect-device | 14618
- request fabric administration power-off node-device | 14622
- show fabric administration inventory | 19878

request fabric administration power-off node-device

IN THIS SECTION

- Syntax | 14622
- Description | 14623
- Options | 14623
- Additional Information | 14623
- Required Privilege Level | 14623
- Output Fields | 14623
- Sample Output | 14623
- Release Information | 14624

Syntax

```
request fabric administration power-off node-device (serial-id | alias-name)
```

Description

In QFX3000-G QFabric systems, you can use this command to power off QFX3500, QFX3600, QFX5100-48S, QFX5100-48T, and QFX5100-24Q devices in server Node groups, redundant server Node groups, and network Node groups. In QFX3000-M QFabric systems, you can use this command to power off QFX5100-24Q Interconnect devices and QFX3600-I Interconnect devices.

This command systematically takes the device offline and gracefully shuts down the device while preserving system state information. A message appears on the console or console log, confirming that the operating system has stopped on the device.

Options

serial-id | alias-name	Provide the serial ID or alias name of the device.

Additional Information

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request fabric administration power-off node-device

```
user@qfabric request fabric administration power-off node-device ED1234
STEP 1 of 8 (Acquiring lock):
Acquiring lock to perform this operation
Acquired lock to perform this operation
STEP 2 of 8 (Performing pre-checks):
node-device ED1491 is online
STEP 3 of 8 (Mastership switch-over):
node-device ED1491 is Server Node-group
STEP 4 of 8 (Gracefully offlining the node-device):
node-device is being offlined
```

```
node-device is successfully offlined
STEP 5 of 8 (Waiting for convergence):
Waiting for convergence(this will take few minutes)
Convergence complete
STEP 6 of 8 (Powering-off the node-device):
Now, powering-off the node-device
Powered-off the node-device
STEP 7 of 8 (Updating inventory):
Removing the node-device - ED1491 from inventory
node-device - ED1491 is removed from inventory
STEP 8 of 8 (Releasing lock):
Releasing the lock
```

Release Information

Command introduced in Junos OS Release 14.1X53-D15.

RELATED DOCUMENTATION

- request fabric administration power-off interconnect-device | 14618
- request fabric administration power-off interconnect-device fpc | 14620
- show fabric administration inventory | 19878

request fabric administration remove

IN THIS SECTION

- Syntax | 14625
- Description | 14625
- Options | 14625
- Required Privilege Level | 14625
- Sample Output | 14626
- Sample Output | 14626
- Release Information | 14626

Syntax

```
request fabric administration remove (interconnect-device interconnect-device-name | node-device
node-device-name)
```

Description

(QFabric systems only) Remove a disconnected Interconnect or Node device from the QFabric system inventory so that it does not appear in the output of the `show fabric administration inventory` command.

NOTE:

- You cannot remove any devices that appear in the `Connected` state in the output of the `show fabric administration inventory` command.
- For Node devices, you can only remove a device if it belongs to an autogenerated server Node group that contains a single Node device. Node devices contained within redundant server Node groups or network Node groups cannot be removed directly. To remove a Node device that is part of a group, delete the device from the Node group configuration first before attempting to remove the device from the inventory.

Options

interconnect-device ***interconnect-device-name***	Remove a disconnected Interconnect device from the QFabric system inventory.
node-device ***node-device-name***	Remove a disconnected Node device from the QFabric system inventory.

Required Privilege Level

admin

Sample Output

request fabric administration remove interconnect-device

```
user@qfabric> request fabric administration remove interconnect-device IC1
Device successfully removed
```

Sample Output

request fabric administration remove node-device

```
user@qfabric> request fabric administration remove node-device node5
Device successfully removed
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

- show fabric administration inventory | 19878
- show fabric administration inventory interconnect-devices | 19898
- show fabric administration inventory node-devices | 19900
- show fabric administration inventory node-groups | 19903

request fabric administration system mac-pool add

IN THIS SECTION

- Syntax | 14627
- Description | 14627
- Options | 14627

- Required Privilege Level | 14627
- Output Fields | 14627
- Sample Output | 14628
- Release Information | 14628

Syntax

```
request fabric administration system mac-pool add mac-base starting-mac-address count number-of-mac-address
```

Description

(QFabric systems only) Add a MAC address pool to expand the initial set of MAC addresses assigned to the QFabric system.

Options

mac-base *starting-mac-address*	Set the starting MAC address for a pool of addresses assigned to the QFabric system.
count *number-of-mac-address*	Set the total number of MAC addresses in the specified address pool assigned to the QFabric system.

Required Privilege Level

admin

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request fabric administration system mac-pool add mac-base starting-mac-address count

```
user@switch> request fabric administration system mac-pool add mac-base 02:00:00:11:22:00 count
10
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

request fabric administration system mac-pool delete | 14628

show fabric administration system mac-pool | 19906

request fabric administration system mac-pool delete

IN THIS SECTION

- Syntax | 14629
- Description | 14629
- Options | 14629
- Additional Information | 14629
- Required Privilege Level | 14629
- Output Fields | 14629
- Sample Output | 14629
- Release Information | 14630

Syntax

```
request fabric administration system mac-pool delete mac-base starting-mac-address
```

Description

(QFabric systems only) Delete a range of MAC addresses assigned manually to the QFabric system.

> **NOTE**: You cannot delete the MAC address range assigned during the initial setup of the QFabric system. Also, you cannot delete a MAC address range if the MAC address block is still in use.

Options

mac-base ***starting-mac-address***	Specify the starting MAC address for a pool of addresses you wish to remove from the QFabric system.

Additional Information

After you issue the `request fabric administration system mac-pool delete` command, issue the "show fabric administration system mac-pool" on page 19906 command to verify that the MAC address range has been deleted.

Required Privilege Level

admin

Output Fields

Sample Output

request fabric administration system mac-pool delete mac-base

```
user@switch> request fabric administration system mac-pool delete mac-base 02:00:00:11:22:00
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

request fabric administration system mac-pool add | 14626

show fabric administration system mac-pool | 19906

request fibre-channel proxy load-rebalance

IN THIS SECTION

- Syntax | 14630
- Description | 14631
- Options | 14631
- Additional Information | 14631
- Required Privilege Level | 14631
- Output Fields | 14631
- Sample Output | 14632
- Release Information | 14632

Syntax

```
request fibre-channel proxy load-rebalance
<dry-run>
fabric <fabric-name>
<brief | detail>
```

Description

Rebalance the link load on one or more FCoE-FC gateway proxy fabrics (local Fibre Channel fabrics on the gateway) on demand. Load rebalancing is a disruptive action that forces some or all sessions (depending on the configured load-balancing algorithm) to log out and then log in again. When sessions log in again, they are placed on NP_Port interfaces so that the link loads are balanced.

Link load rebalancing occurs 10 seconds after you run the rebalancing command, unless another rebalancing trigger occurs before the 10 seconds elapse. If another rebalancing event occurs before the 10-second timer elapses, the timer is extended. Rebalancing occurs a maximum of 30 seconds after you run the rebalancing command, regardless of whether more rebalancing events occur.

You can also perform a *dry run* to see a list of sessions that might be affected (logged out) if you request a load rebalance. A dry run does not rebalance the link loads; it only lists the sessions that might be affected if you rebalance.

Options

dry-run	(Optional) Simulates performing link load rebalancing and displays a list of sessions that might be affected if you rebalance the link loads.
fabric ***fabric-name***	Name of the fabric on which you want to rebalance the link loads. If you do not specify a fabric name with the fabric keyword, all fabrics on the FCoE-FC gateway rebalance their link loads.
brief | detail	(Optional) Display the specified level of output.

Additional Information

Requesting link load rebalancing is a one-time, on-demand operation. You must explicitly request load rebalancing every time you want to rebalance the link loads. Alternatively, you can configure automated load rebalancing if you want the NP_Port links to be rebalanced automatically whenever a load-rebalancing trigger occurs.

Required Privilege Level

maintenance

Output Fields

Table 113 on page 14632 lists the output fields for the `request fibre-channel proxy load-rebalance dry-run` command. Output fields are listed in the approximate order in which they appear.

Table 113: request fibre-channel proxy load-rebalance dry-run Output Fields

Field Name	Field Description
Fabric	Name of the fabric.
Fabric-id	Fabric ID number.
F-Port	FCoE VLAN interface (VF_Port interface to the FCoE network).
FCID	VN_Port Fibre Channel identifier provided by the Fibre Channel over Ethernet Forwarder (FCoE forwarder) or the Fibre Channel switch.
Port-WWN	Unique worldwide name (WWN) of the VN_Port.
NP-Port	Name of the native Fibre Channel interface.

Sample Output

request fibre-channel proxy load-rebalance dry-run fabric fc_fabric_100

```
user@host> request fibre-channel proxy load-rebalance dry-run fabric fc_fabric_100
Fabric: fc_fabric_100, Fabric-id: 100
F-Port                        FCID     Port-WWN                NP-Port
vlan.100                      0x8a013a 02:01:00:64:00:00:00:2a fc-0/0/1.0
vlan.100                      0x8a013c 02:01:00:64:00:00:00:2b fc-0/0/1.0
vlan.100                      0x8a0146 02:01:00:64:00:00:00:2e fc-0/0/1.0
vlan.100                      0x8a014c 02:01:00:64:00:00:00:2f fc-0/0/1.0
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Monitoring Fibre Channel Interface Load Balancing

Simulating On-Demand Fibre Channel Link Load Rebalancing (Dry Run Test)

Defining the Proxy Load-Balancing Algorithm

Example: Configuring Automated Fibre Channel Interface Load Rebalancing

Understanding Load Balancing in an FCoE-FC Gateway Proxy Fabric

request flight-recorder set high-cpu

IN THIS SECTION

- Syntax | 14633
- Description | 14633
- Options | 14634
- Required Privilege Level | 14635
- Output Fields | 14635
- Sample Output | 14635
- Release Information | 14636

Syntax

```
request flight-recorder set high-cpu
<disable>
<backoff-duration seconds>
<collect-core>
<cpu-threshold percentage>
<logical-system>
<num-snapshots number>
<polling-frequency frequency>
```

Description

Enable flight recorder tool to collect snapshots of historical data on when the CPU utilization for the routing protocol process on a device was high and what processes caused the high utilization. The detection of high CPU usage enables faster resolution of issues.

The recorded snapshots and core files are saved as log files in a folder under the **/var/log/flight_recorder/** directory. The log files are listed in the order of time stamp saved. The folder format is `Flr_MONTH_DD_YYYY_HH:MM:SS`; for example, `Flr_May_09_2018_02:20:50`. Each log file in the directory includes the following information:

- Output from the `show task accounting detail` command (after enabling and waiting for 10 seconds).
- Output from the `show task jobs` command.
- Running core data stored in a separate core log file, if enabled.

Options

none

Enable flight recorder tool to collect snapshots of data used for detecting high CPU utilization. The recorded snapshots and core files are saved as log files in a folder under the **/var/log/flight_recorder/**.

disable

Disable flight recorder tool that has been enabled using the `request flight-recorder set high-cpu` command.

- **Default:** Disabled.

backoff-duration *seconds*

(Optional) Specify the time interval in seconds between two snapshots of data.

- **Default:** 100 seconds.
- **Range:** 10 through 1000.

collect-core

(Optional) Perform snapshot collection of the running core with every snapshot of data taken.

When the `collect-core` option is enabled, the data snapshots are stored in a separate core log file in a folder under the **/var/log/flight_recorder/** directory. The folder format is `Flr_MONTH_DD_YYYY_HH:MM:SS`; for example, `Flr_May_09_2018_02:20:50`.

- **Default:** Disabled.

cpu-threshold *percentage*

(Optional) Specify the maximum value of CPU utilization in percentage, beyond which the collection of data is triggered.

- **Default:** 80
- **Range:** 1 through 400.

logical-system

(Optional) Enable data collection on logical systems.

- **Default:** Disabled.

num-snapshots *number*

(Optional) Specify the number of snapshots of data to be collected before quitting the collection process.

- **Default:** 0
- **Range:** 1 through 20

polling-frequency *seconds*

(Optional) Specify the time in seconds for polling for high CPU utilization.

- **Default:** 10 seconds.
- **Range:** 5 through 100.

Required Privilege Level

root

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request flight-recorder set high-cpu (Enable flight-recorder)

```
user@host> request flight-recorder set high-cpu cpu-threshold 10 polling-frequency 5 backoff-
duration 10 collect-core num-snapshots 1
Please wait....Starting flight-recorder process.
```

request flight-recorder set high-cpu disable (When flight-recorder is enabled)

```
user@host> request flight-recorder set high-cpu disable
Disabling Done
```

request flight-recorder set high-cpu disable (When flight-recorder is disabled)

```
user@host> request flight-recorder set high-cpu disable
Flight Recorder is not running!
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show flight-recorder status | 20083

request forwarding-options port-mirroring instance family self-mirror-start

IN THIS SECTION

- Syntax | 14636
- Description | 14637
- Options | 14637
- Required Privilege Level | 14637
- Output Fields | 14637
- Sample Output | 14637
- Release Information | 14637

Syntax

```
request forwarding-options port-mirroring instance instance-name <family family-name> self-
mirror-start start-duration-seconds
```

Description

Start the on-device packet capture (self-mirroring) of network packets to the CPU and to PCAP files.

Options

family *family-name*	(Optional) Protocol family to which this command is applied.
instance *instance-name*	Port-mirroring instance to which this command is applied.
self-mirror-start *start-duration-seconds*	Duration in seconds before the start of the file recording (1–1800 seconds).

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show forwarding-options port-mirroring self-mirror start` command to verify the operation.

Sample Output

request forwarding-options port-mirroring instance I1 family inet self-mirror-start 900

```
user@host> request forwarding-options port-mirroring instance I1 family inet self-mirror start
900
```

Release Information

Command introduced in Junos OS Evolved Release 24.2R1.

request forwarding-options port-mirroring instance family self-mirror-stop

IN THIS SECTION

- Syntax | 14638
- Description | 14638
- Options | 14638
- Required Privilege Level | 14638
- Output Fields | 14639
- Sample Output | 14639
- Release Information | 14639

Syntax

```
request forwarding-options port-mirroring instance instance-name <family family-name> self-
mirror-stop
```

Description

Stop the on-device packet capture (self-mirroring) of network packets to the CPU and to PCAP files.

Options

family *family-name*	(Optional) Protocol family to which this command is applied.
instance *instance-name*	Port-mirroring instance to which this command is applied.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the `show forwarding-options port-mirroring self-mirror stop` command to verify the operation.

Sample Output

request forwarding-options port-mirroring instance I1 family inet self-mirror-stop

```
user@host> request forwarding-options port-mirroring instance I1 family inet self-mirror-stop
```

Release Information

Command introduced in Junos OS Evolved Release 24.2R1.

request interface bounce

IN THIS SECTION

- Syntax | 14639
- Description | 14640
- Options | 14640
- Additional Information | 14640
- Required Privilege Level | 14640
- Output Fields | 14640
- Sample Outputs | 14641
- Release Information | 14641

Syntax

```
request interface bounce interface_name poe interval seconds
```

Description

You can shut down the interface for a given time interval. The interface powers up at the end of the configured time.

Options

poe

An optional parameter that bounces the poe port. The time interval range is from 1 to 60 seconds. An example of the syntax with poe is shown:

```
request interface bounce interface_name poe interval seconds
```

> **NOTE**: You can only bounce poe on ports that have been configured for poe.

interface *interface-name*

Name of the interface to shut down for the given time interval.

interval *seconds*

The interface shut down for the configured time interval. The time interval range is from 1 to 30 seconds.

Additional Information

- The default time interval is 0 seconds.
- In Junos OS the `poe bounce` command is supported only on Gigabit Ethernet and Multi rate Gigabit Ethernet.
- In Junos OS Evolved the `poe bounce` command is supported only on Ethernet interface.

Required Privilege Level

View

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Outputs

request interface bounce

```
user@host> request interface bounce et-0/0/0 interval 20
```

```
Bounce operation on interface et-0/0/0 started with interval 20 secs.
```

poe port bounce with interval

```
user@host> request interface bounce poe mge-0/0/29 interval 30
```

```
PoE bounce request received for mge-0/0/29 with interval 30s.
```

poe port bounce without interval

```
user@host> request interface bounce poe mge-0/0/29
```

```
PoE bounce request received for mge-0/0/29 with interval 0s.
```

See Sample outputs for information on what the output looks like when poe bounce is enabled and when poe bounce is disabled.

Release Information

Starting in 24.2 R1 release, we've introduced the `poe` optional parameter. This statement was originally introduced in Junos OS release and Junos OS Evolved release 23.4R1.

request interface link-degrade-recover

IN THIS SECTION

- Syntax | 14642
- Description | 14642
- Options | 14642
- Required Privilege Level | 14642
- Output Fields | 14643
- Sample Output | 14643
- Release Information | 14645

Syntax

```
request interface link-degrade-recover interfaces-name
```

Description

Manually recover a degraded physical link. Manual recovery is used when the interface has any Layer 2 and Layer 3 protocols that prevents autorecovery. This command is applicable only if you have configured the manual link recovery option on the interface.

NOTE: Manual recovery option is recommended for user deployments that have static route configurations causing the remote end of the link to start forwarding packets (as soon as the physical link is up) while auto-recovery is in progress.

Options

interfaces-name	Name of the interface.

Required Privilege Level

View

Output Fields

When you enter this command, Junos OS displays the status of your request.

Sample Output

Manual recovery

```
user@host>run request interface link-degrade-recover xe-9/1/11
FPC 9 PIC 1 PORT 11 Link        Degrade Recovery Started
```

Interface status when link degrade is enabled

```
user@host>run show interfaces xe-9/1/11
Physical interface: xe-9/1/11, Enabled, Physical link is Up
Interface index: 181, SNMP ifIndex: 664
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, Loop Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: None, Source filtering:
Disabled,
  Flow control: Enabled, Speed Configuration: Auto
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: 28:8a:1c:c9:0e:32, Hardware address: 28:8a:1c:c9:0e:32
  Last flapped   : 2017-10-25 01:53:17 PDT (00:00:10 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Link Degrade :
    Link Monitoring                   :  Enable
    Link Degrade Set Threshold        :  1E-8
    Link Degrade Clear Threshold      :  1E-11
    Link Degrade War Set Threshold    :  1E-9
```

```
    Link Degrade War Clear Threshold  :  1E-10
    Estimated BER                     :  <= 1E-16
    Link-degrade event                :  Seconds              Count                State
                                         0                    0                    OK
  Interface transmit statistics: Disabled

  Logical interface xe-9/1/11.0 (Index 32368) (SNMP ifIndex 33153)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop
cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

Interface status when the defect is active

```
user@host>run show interfaces xe-9/1/11
Physical interface: xe-9/1/11, Enabled, Physical link is Down
Interface index: 181, SNMP ifIndex: 664
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, Loop Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: None, Source filtering:
Disabled,
  Flow control: Enabled, Speed Configuration: Auto
  Pad to minimum frame size: Disabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: 28:8a:1c:c9:0e:32, Hardware address: 28:8a:1c:c9:0e:32
  Last flapped   : 2017-10-25 01:54:09 PDT (00:00:03 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK, LOCAL-FAULT
  PCS statistics                      Seconds
```

```
    Bit errors                              0
    Errored blocks                          0
  Link Degrade :
    Link Monitoring                   :  Enable
    Link Degrade Set Threshold        :  1E-8
    Link Degrade Clear Threshold      :  1E-11
    Link Degrade War Set Threshold    :  1E-9
    Link Degrade War Clear Threshold  :  1E-10
    Estimated BER                     :  1E-4
    Link-degrade event                :  Seconds              Count                State
                                         4                    1                    Defect Active
  Interface transmit statistics: Disabled

  Logical interface xe-9/1/11.0 (Index 32368) (SNMP ifIndex 33153)
    Flags: Device-Down SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop
cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Link Degrade Monitoring Overview

link-degrade-monitor | **9784**

thresholds | **29346**

recovery | **14146**

request interface load-balancing revert (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 14646
- Description | 14646
- Options | 14646
- Required Privilege Level | 14646
- Output Fields | 14646
- Sample Output | 14647
- Release Information | 14647

Syntax

```
request interface load-balancing revert interface-name
```

Description

Revert the aggregated multiservices member interface (mams-) from the inactive state to the active or backup state based on the configuration and the operational state of the aggregated multiservices interface.

Options

interface-name	Name of the member interface. The member interface format is mams-*a*/*b*/0, where *a* is the FPC slot number and *b* is the PIC slot number.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request interface load-balancing revert

```
user@host> request interface load-balancing revert mams-4/0/0
request succeeded
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

request interface load-balancing switchover (Aggregated Multiservices) | 14647

request interface load-balancing switchover (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 14648
- Description | 14648
- Options | 14648
- Required Privilege Level | 14648
- Output Fields | 14648
- Sample Output | 14649
- Release Information | 14649

Syntax

```
request interface load-balancing switchover interface-name
<force>
```

Description

Switch the active member interface to the backup state.

In the case of mobile control plane redundancy, the behavior depends on the replication state of the member interface:

- If the sync state is `in-sync`, then the active member is rebooted and the backup member becomes the new active member.
- If the sync-state is `in-progress`, then the `force` option must be used to force the switchover.

> **CAUTION**: In this case, there is a risk of losing subscriber information because the synchronization has not yet been completed.

Options

interface-name	Name of the member interface. The member interface format is mams-*a*/*b*/0, where *a* is the FPC slot number and *b* is the PIC slot number.
force	(Optional) Force the switchover from the active member to the backup member.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request interface load-balancing switchover force

```
user@host> request interface load-balancing switchover force mams-4/0/0
Switchover Initiated
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

request interface load-balancing revert (Aggregated Multiservices) | 14646

request interface mc-ae switchover (Multichassis Link Aggregation)

IN THIS SECTION

- Syntax | 14650
- Description | 14650
- Options | 14650
- Required Privilege Level | 14650
- Output Fields | 14650
- Sample Output | 14651
- Sample Output | 14651
- Release Information | 14651

Syntax

```
request interface mc-ae switchover
<immediate> mcae-id mcae-id;
mcae-id mcae-id;
```

Description

Manually revert egress traffic from the active node to the designated preferred node of a multichassis aggregated Ethernet interface. You can use this command to manually switch over traffic to the preferred node when the `switchover-mode` statement for the multichassis aggregated Ethernet interface is configured as `non-revertive` at the [`edit interfaces ae`*X* `mc-ae`] hierarchy level.

> **NOTE**: To run this command successfully, the `status-control` statement should be configured as `active` at the [`edit interfaces ae`*X* `mc-ae`] hierarchy level.

Options

`immediate` (Optional) Trigger immediate switchover to the preferred node. If this option is not configured, Junos OS waits for the timer configured using the `revert-time` statement at the [`edit interfaces ae`*X* `mc-ae`] hierarchy level to expire before it triggers the switchover.

`mcae-id` *`mcae-id`* Triggers switchover for the specified mc-ae interface.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request interface mc-ae switchover immediate mcae-id

```
user@host >request interface mc-ae switchover immediate mcae-id 2
MCAE: Switchover Done
```

Sample Output

request interface mc-ae switchover mcae-id

```
user@host >request interface mc-ae switchover mcae-id 2
Switchover In Progress: Please check after 1 minutes,
Use “show interfaces mc-ae revertive-info” to check for the status
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

- Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up
- Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on MX Series Routers

request interface optics-reset

IN THIS SECTION

- Syntax | 14652
- Description | 14652
- Options | 14652
- Required Privilege Level | 14652

Sample Output | 14652
Release Information | 14652

Syntax

```
request interface optics-reset transceiver-name
```

Description

(PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX100016 routers only) Perform manual optics-reset.

Options

transceiver-name

Transceiver name. For example:

- xcvr-fpc/pic/port

Required Privilege Level

maintenance

Sample Output

request interface optics-reset (PTX10002-36QDD)

```
user@host> request interface optics-reset xcvr-0/0/0
```

Release Information

Command introduced in Junos OS Evolved Release 23.4R2-S1.

request interface rebalance (Aggregated Ethernet for Subscriber Management)

IN THIS SECTION

- Syntax | 14653
- Description | 14653
- Options | 14653
- Additional Information | 14653
- Required Privilege Level | 14654
- Output Fields | 14654
- Sample Output | 14655
- Release Information | 14656

Syntax

```
request interface rebalance interface interface-name
```

Description

Manually rebalance the subscribers on an aggregated Ethernet bundle with targeted distribution enabled. For example, when you add a new member link to an existing aggregated Ethernet interface, you might want to rebalance the subscribers across the new member link.

Options

interface-name	Aggregated Ethernet logical interface number.

Additional Information

The `targeted-distribution` statement in a dynamic profile automatically distributes subscribers to the aggregated Ethernet links associated with the profile. If you apply the profile and then bring up a single

link on an aggregated Ethernet interface, all newly connected subscribers are located on that single, active link.

When you then bring up a second link on the aggregated Ethernet interface, you can use the `request interface rebalance interface` command to rebalance the existing subscribers so that some remain on the first link and others are redistributed to the second link.

In some cases, it might appear that rebalancing does not work. By default, the router rebalances the subscribers only when the difference in subscriber numbers on the member links is greater than 500. This number is the default threshold for subscriber granularity.

Consider the following cases with two member links on the interface:

- Case 1—One link has 300 subscribers and the other link has 10 subscribers. Issuing the command does not rebalance the subscribers because the difference between the links is only 290.
- Case 2—One link has 530 subscribers and the other link has 10 subscribers. The difference between the links is 520 subscribers. Issuing the command rebalances the subscribers so that 500 subscribers are on one link and 40 subscribers are on the other link.

You can modify the subscriber granularity with the `rebalance-subscriber-granularity` option at the `[edit interfaces ae.`*`interface-number`*` aggregated-ether-options aggregated-ether-options]` hierarchy level. For example, if you set the granularity to 1 subscriber, then rebalancing always takes place when the member links differ by 2 or more subscribers.

BEST PRACTICE: We recommend that you do not configure a low granularity value. A low value can have undesirable effects, such as the router running out of pseudo logical interfaces or an increase in the convergence time for rebalancing.

Leaving subscriber granularity at the default value of 500 subscribers is sufficient in most cases. Whenever more than one member link is active, targeted distribution places new subscribers on a link with fewer subscribers than other member links on the interface.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. You can compare the output of the `show interfaces targeting` command before and after the request to view the effect of the rebalancing operation. .

Sample Output

request interface rebalance (Adding a New Link)

```
user@host >show interfaces targeting ae0
Aggregated interface: ae0
Targeting Type: Auto
Redundancy mode: Link Level Redundancy
Total number of distribution groups: 1
Total number of distributed interfaces: 1000

Distribution Group name: default
Number of distributed interfaces: 1000
Physical interface: xe-0/2/0, Link status: Up
Number of primary distributions: 1000     /*All subscribers are distributed on this active link*/
Number of backup distributions: 0
Physical interface: xe-0/2/1, Link status: Down   /*This link is inactive*/
Number of primary distributions: 0     /*No subscribers*/
Number of backup distributions: 1000
user@host >request interface rebalance interface ae0 (After Link xe-0/2/1 is Activated)
Rebalance operation on interface ae0 started
user@host >show interfaces targeting ae0
Aggregated interface: ae0
Targeting Type: Auto
Redundancy mode: Link Level Redundancy
Total number of distribution groups: 1
Total number of distributed interfaces: 1000

Distribution Group name: default
Number of distributed interfaces: 1000
Physical interface: xe-0/2/0, Link status: Up
  Number of primary distributions: 500    /*500 subscribers were redistributed from this link*/
  Number of backup  distributions: 500
Physical interface: xe-0/2/1, Link status: Up    /*member link is now active*/
  Number of primary distributions: 500    /*500 subscribers were redistributed to this newly
active member link*/
  Number of backup  distributions: 500
```

request interface rebalance (Rebalancing Two Existing Links)

```
user@host >show interfaces targeting ae0
Aggregated interface: ae0
Targeting Type: Auto
Redundancy mode: Link Level Redundancy
Total number of distribution groups: 1
Total number of distributed interfaces: 540

Distribution Group name: default
Number of distributed interfaces: 540
Physical interface: xe-0/2/0, Link status: Up
  Number of primary distributions: 530     /*This interface has 520 more subscribers than the
other*/
  Number of backup  distributions: 10
Physical interface: xe-0/2/1, Link status: Up
  Number of primary distributions: 10
  Number of backup  distributions: 530

user@host >request interface rebalance interface ae0
Rebalance operation on interface ae0 started
user@host >show interfaces targeting ae0
Aggregated interface: ae0
Targeting Type: Auto
Redundancy mode: Link Level Redundancy
Total number of distribution groups: 1
Total number of distributed interfaces: 540

Distribution Group name: default
Number of distributed interfaces: 540
Physical interface: xe-0/2/0, Link status: Up
  Number of primary distributions: 500    /*30 subscribers were redistributed from this link*/
  Number of backup  distributions: 40
Physical interface: xe-0/2/1, Link status: Up
  Number of primary distributions: 40    /*30 subscribers were redistributed to the link with
fewer subscribers*/
  Number of backup  distributions: 500
```

Release Information

Command introduced in Junos OS Release 11.2.

request interface revert

IN THIS SECTION

- Syntax | 14657
- Description | 14657
- Options | 14657
- Required Privilege Level | 14657
- Output Fields | 14657
- Sample Output | 14658
- Release Information | 14658

Syntax

```
request interface revert interface interface-name
```

Description

Revert from the secondary to the primary interface in the specified warm standby AMS interface.

Options

interface *interface-name*	Name of the AMS interface in which you want to revert from the secondary to the primary interface.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request interface revert interface

```
user@host> request interface revert interface ams1
```

Release Information

Statement introduced in Junos OS Release 17.2.

request interface (revert | switchover) (Adaptive Services)

IN THIS SECTION

- Syntax | 14658
- Description | 14658
- Options | 14659
- Required Privilege Level | 14659
- Output Fields | 14659
- Sample Output | 14659
- Release Information | 14660

Syntax

```
request interface (revert | switchover) (rspnumber | rlsqnumber)
```

Description

(M Series and T Series routers only) Manually revert to the primary adaptive services interface or link services IQ interface, or to switch from the primary to the secondary interface.

> **NOTE**: All **rlsq** switchover or revert operations are allowed from the **rlsq***number* level only and not for individual channelized interfaces (**rlsq***number.unit*).

On an aggregated Ethernet interface with link protection enabled, use the **request interface (revert | switchover)** (Aggregated Ethernet Link Protection) operational command to manually revert egress traffic from the designated backup link to the designated primary link, or to manually switch egress traffic from the primary link to the backup link. For information about this command, see "request interface (revert | switchover) (Aggregated Ethernet Link Protection)" on page 14660.

Options

(revert | switchover)	The **revert** keyword restores active processing to the primary adaptive services (sp) or link services IQ (lsq) interface. The **switchover** keyword transfers active processing to the secondary (backup) interface.
rsp*number*	Redundant adaptive services interface name.
rlsq*number*	Redundant link services IQ interface name.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request interface revert

```
user@host> request interface revert rlsq0
request succeeded
```

request interface switchover

```
user@host> request interface switchover rlsq0
error: rlsq0: already on secondary
```

Release Information

Command introduced before Junos OS Release 7.4.

Support for **rlsq** interfaces added in Junos OS Release 7.6.

request interface (revert | switchover) (Aggregated Ethernet Link Protection)

IN THIS SECTION

- Syntax | 14660
- Description | 14660
- Options | 14661
- Required Privilege Level | 14661
- Output Fields | 14661
- Sample Output | 14661
- Release Information | 14662

Syntax

```
request interface (revert | switchover) aex
```

Description

Manually revert egress traffic from the designated backup link to the designated primary link of an aggregated Ethernet interface for which link protection is enabled, or manually switch egress traffic from

the primary link to the backup link. This traffic includes transit traffic and local traffic originated on the router itself.

> **NOTE**: When link protection is enabled on an aggregated Ethernet interface, if the primary link fails, the router automatically routes egress traffic to the backup link. However, the router does not automatically route egress traffic back to the primary link when the primary link is subsequently reestablished. Instead, you manually control when to have traffic diverted back to the primary link by issuing the **request interface (revert | switchover)** (Aggregated Ethernet Link Protection) operational command and specifying the **revert** keyword.

On M Series and T Series routers, use the **request interface (revert | switchover)** (Adaptive Services) operational command to manually revert to the primary adaptive services interface or link services interface, or to switch from the primary to the secondary interface. For information about this command, see "request interface (revert | switchover) (Adaptive Services)" on page 14658.

Options

revert	Restores egress traffic processing to the primary link.
switchover	Transfers egress traffic processing to the secondary (backup) link.
ae*x*	Aggregated Ethernet logical interface number: 0 through 15.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request interface revert

```
user@host >request interface revert ae1
```

Release Information

Command introduced in Junos OS Release 8.3.

request interface (revert | switchover) (Aggregated Inline Service Interfaces)

IN THIS SECTION

- Syntax | 14662
- Description | 14662
- Options | 14663
- Required Privilege Level | 14663
- Output Fields | 14663
- Sample Output | 14663
- Sample Output | 14663
- Release Information | 14664

Syntax

```
request interface (revert | switchover) bundle-name
```

Description

Manually revert L2TP data traffic from the designated backup link to the designated primary link of an aggregated inline service interface bundle interface for which 1:1 redundancy is configured, or manually switch data traffic from the primary link to the backup link.

NOTE: When 1:1 redundancy protection is configured for an aggregated inline service interface, if the primary link fails, the router automatically routes data traffic destined for the L2TP session on that link to the backup link. However, the router does not

automatically route data traffic back to the primary link when the primary link is subsequently reestablished. Instead, you manually divert traffic back to the primary link by issuing the `request interface revert` operational command.

Options

revert	Restore data traffic for the LNS session to the primary link.
switchover	Transfer data traffic for the LNS session to the secondary (backup) link.
bundle-name	Name of the aggregated inline service interface bundle.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request interface switchover

```
user@host >request interface switchover asi0
error: requesting cmd SWITCH when primary is not active
```

Sample Output

request interface revert

```
user@host >request interface revert asi0
request succeeded
```

Release Information

Command introduced in Junos OS Release 16.2.

request interface switchover

IN THIS SECTION

- Syntax | 14664
- Description | 14664
- Options | 14664
- Required Privilege Level | 14665
- Output Fields | 14665
- Sample Output | 14665
- Release Information | 14665

Syntax

```
request interface switchover interface interface-name
```

Description

Switch over from the primary to the secondary (backup) interface in the specified warm standby AMS interface. If the secondary interface is already in use, then the operation fails.

Options

interface *interface-name*	Name of the AMS interface in which you want to switchover from the primary to the secondary interface.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request interface switchover interface

```
user@host> request interface switchover interface ams1
```

Release Information

Statement introduced in Junos OS Release 17.2.

request ipsec switch

IN THIS SECTION

- Syntax | 14666
- Description | 14666
- Options | 14666
- Required Privilege Level | 14666
- Output Fields | 14666
- Sample Output | 14666
- Release Information | 14666

Syntax

```
request ipsec switch (interface <es-fpc/pic/port> | security-associations <sa-name>)
```

Description

(Encryption interface on M Series, PTX Series, and T Series routers and EX Series switches only) Manually switch from the primary to the backup encryption services interface, or switch from the primary to the backup IP Security (IPsec) tunnel.

Options

interface <es-*fpc/pic/port*>	Switch to the backup encryption interface.
security-associations <*sa-name*>	Switch to the backup tunnel.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request ipsec switch security-associations

```
user@host> request ipsec switch security-associations sa-private
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ipsec redundancy | 21672

request jdaf clear statistics service

IN THIS SECTION

- Syntax | 14667
- Description | 14667
- Options | 14667
- Required Privilege Level | 14668
- Sample Output | 14668
- Release Information | 14668

Syntax

```
request jdaf clear statistics service <(cmd | config | firewall | interface | route) client
client-id | server server-id>
```

Description

Clear Juniper distributed application framework (JDAF) client or server side statistics for either command, configuration, firewall, interface, or route service.

Options

cmd	Clear JDAF command service message statistics.
config	Clear JDAF config service message statistics.
firewall	Clear JDAF firewall service message statistics.
interface	Clear JDAF interface service message statistics.
route	Clear JDAF route service message statistics.

client Clear JDAF client side statistics.

server Clear JDAF server statistics.

Required Privilege Level

view

Sample Output

request jdaf clear statistics service

```
user@host> request jdaf clear statistics service cmd client
warning: app-engine-management-service subsystem not running - not needed by configuration.
```

Release Information

Command introduced in Junos OS Release 14.2.

request jnu role

IN THIS SECTION

- Syntax | 14669
- Description | 14669
- Options | 14669
- Required Privilege Level | 14669
- Sample Output | 14669
- Release Information | 14670

Syntax

```
request jnu role (controller | invoke-on <all-routing-engines | other-routing-engine> | no-
forwarding | satellite)
```

Description

Perform Junos Node Unifier (JNU) operations such as controller, remote command execution, no forwarding, and satellite operations.

Options

controller	Controller.
invoke-on	Remote command execution either on all routing engines or on other routing engine.
no-forwarding	No forwarding.
satellite	Satellite

Required Privilege Level

view

Sample Output

The following commands produce no output.

request jnu role controller invoke-on all-routing-engines

```
user@host> request jnu role controller invoke-on all-routing-engines
```

request jnu role no-forwarding controller

```
user@host> request jnu role no-forwarding controller
```

Release Information

Command introduced in Junos OS Release 22.1R1.

request jnu satellite sync

IN THIS SECTION

- Syntax | 14670
- Description | 14670
- Required Privilege Level | 14670
- Sample Output | 14671
- Release Information | 14671

Syntax

```
request jnu satellite sync
```

Description

Applicable on JNU satellite in Junos Node Unifier (JNU) topology. The command synchronizes the JNU satellite with the JNU controller manually if the initial automatic synchronization fails. This command also restarts the *jnud* process on the satellite.

Required Privilege Level

view

Sample Output

request jnu satellite sync

```
user@host> request jnu satellite sync
```

```
Junos node unifier process started, pid 96869
```

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

request jnu schema add

IN THIS SECTION

- Syntax | 14671
- Description | 14672
- Options | 14672
- Required Privilege Level | 14672
- Sample Output | 14672
- Release Information | 14672

Syntax

```
request jnu schema add ( filename filename | model model | version version
```

Description

Perform Junos Node Unifier (JNU) schema related operations such as adding satellite schema TLV package files, satellite models, or satellite version.

Options

filename	Satellite schema TLV package filename.
model	Satellite model
version	Satellite version

Required Privilege Level

view

Sample Output

The following command produces no output if the operation is successful. It displays a failure message such as **Satellite Schema creation failed** if the operation is unsuccessful.

request jnu schema add

```
user@host> request jnu schema add model mx240 version 21.1 /root/tunnel_plugin-
x86-64-20220803-080233.tgz
```

Release Information

Command introduced in Junos OS Release 22.1R1.

request jnu schema delete

IN THIS SECTION

- Syntax | 14673
- Description | 14673
- Options | 14673
- Required Privilege Level | 14673
- Sample Output | 14674
- Release Information | 14674

Syntax

```
request jnu schema delete ( model model | version version
```

Description

Delete a JNU satellite schema.

Options

model	Satellite model
version	Satellite version

Required Privilege Level

view

Sample Output

request jnu schema delete

The following command produces no output if the operation is successful. It displays a failure message such as **Satellite Schema deletion failed** if the operation is unsuccessful.

```
user@host> request jnu schema delete model mx240 version 21.1 /root/tunnel_plugin-
x86-64-20220803-080233.tgz
```

Release Information

Command introduced in Junos OS Release 22.1R1.

request jnu schema merge

IN THIS SECTION

- Syntax | 14674
- Description | 14674
- Options | 14675
- Required Privilege Level | 14675
- Sample Output | 14675
- Release Information | 14675

Syntax

```
request jnu schema merge ( conschema | satschema)
```

Description

Merge JNU schema to the controller.

Options

conschema	Controller Schema Path.
satschema	Satellite Schema Path.

Required Privilege Level

view

Sample Output

request jnu schema merge

The following command produces no output if the operation is successful. It displays a failure message such as **Satellite Schema merge failed** if the operation is unsuccessful.

```
user@host> request jnu schema merge
```

Release Information

Command introduced in Junos OS Release 22.1R1.

request l2circuit-switchover

IN THIS SECTION

- Syntax | 14676
- Description | 14676
- Options | 14676
- Required Privilege Level | 14676
- Output Fields | 14676
- Sample Output | 14676
- Release Information | 14677

Syntax

```
request l2circuit-switchover
<logical-system (all | logical-system-name)>
<neighbor address>
<virtual-circuit-id identifier>
```

Description

Manually trigger a switch from the active pseudowire to the redundant pseudowire. This command can be useful when performing network maintenance.

Options

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
neighbor *address*	(Optional) Trigger a switch of all of the active pseudowire connections with the specified neighbor to their respective redundant pseudowires.
virtual-circuit-id *identifier*	(Optional) Trigger a switch from the active pseudowire connection of the specified Layer 2 circuit to its redundant pseudowire.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request l2circuit-switchover virtual-circuit-id

```
user@host>request l2circuit-switchover virtual-circuit-id 12
```

Release Information

Command introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

MPLS Feature Support on QFX Series and EX4600 Switches

request lacp link-switchover

IN THIS SECTION

- Syntax | 14677
- Description | 14677
- Options | 14678
- Required Privilege Level | 14678
- Output Fields | 14678
- Sample Output | 14678
- Release Information | 14678

Syntax

```
request lacp link-switchover aex
```

Description

Manually switch aggregated Ethernet active or standby LACP links.

NOTE: Because this command overrides LACP priority calculations, we strongly recommend that you use this command only when the actor (in this case, the Juniper

Networks router) is controlling the active or standby link and the partner (peer) is following. This scenario occurs when you configure only the actor for link protection.

Options

ae*x* Aggregated Ethernet logical interface number: 0 through 15.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. To view the switchover, use the `show lacp interfaces` command.

Sample Output

request lacp link-switchover aeX

```
user@host >request lacp link-switchover ae0ae0: Request succeeded
```

Release Information

Command introduced in Junos OS Release 9.3.

request-max-tcp-connections (System Process)

IN THIS SECTION

- Syntax | 14679
- Hierarchy Level | 14679
- Description | 14679

- Options | 14679
- Required Privilege Level | 14679
- Release Information | 14679

Syntax

```
request-max-tcp-connections max-tcp-connections;
```

Hierarchy Level

```
[edit system processes dhcp-service]
```

Description

Specify the number of simultaneous TCP connections that the DHCP relay can request when sending bulk leasequery messages to DHCP servers.

Options

max-tcp-connections

Number of simultaneous connections.

- **Range:** 1 through 10
- **Default:** 3

Required Privilege Level

system—to view this statement in the configuration.

system-control—to add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring and Using DHCP Individual Leasequery

Configuring and Using DHCP Bulk Leasequery

request message

IN THIS SECTION

- Syntax | 14680
- Description | 14680
- Options | 14680
- Required Privilege Level | 14681
- Output Fields | 14681
- Sample Output | 14681
- Release Information | 14681

Syntax

```
request message all message "text"
request message message "text" (terminal terminal-name  | user user-name)
```

Description

Display a message on the screens of all users who are logged in to the router or switch or on specific screens.

Options

all	Display a message on the terminal of all users who are currently logged in.
message "*text*"	Message to display.

terminal *terminal-name* Name of the terminal on which to display the message.

user *user-name* Name of the user to whom to direct the message.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request message message

```
user@host> request message message "Maintenance window in 10 minutes" user maria
Message from user@host on ttyp0 at 20:27 ...
Maintenance window in 10 minutes
EOF
```

Release Information

Command introduced before Junos OS Release 7.4.

request modem wireless create-profile

IN THIS SECTION

- Syntax | 14682
- Description | 14682
- Options | 14682
- Required Privilege Level | 14683
- Sample Output | 14683

Release Information | 14683

Syntax

```
request modem wireless create-profile interface-name access-point-name access-point-name
authentication-method authentication-method  profile-id profile-id sip-password sip-password sip-
user-id sip-id slot sim-slot-number
```

Description

Create a profile. The Subscriber Identity Module (SIM) uses a profile to establish a connection with the network. You can configure up to 16 profiles for each SIM card. The LTE Mini-PIM supports two SIM cards and so you can configure a total of 32 profiles, although only one profile can be active at a time.

To create a profile, you must obtain the following information from the service provider:

- Username and password
- Access point name (APN)
- Authentication (Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP))

Options

- *interface-name*—The LTE interface is `cl-x/0/0`, where *x* is the slot number in which the LTE Mini-PIM is installed.
- `access-point-name` *access-point-name*—Access point name (APN). Obtain the APN from the service provider. You can specify only a single APN in a profile.
- `authentication-method`—The authentication protocol that the SIM card uses to authenticate with the wireless network. Obtain the authentication information from the service provider. The authentication protocol used by the SIM card must match the protocol used by the service provider. The `authentication-method` can be one of the following:
  - CHAP
  - PAP

- None

- `profile-id` *`profile-id`*—Profile identification number for the profile. The default value is 1. The range of possible values is from 1 through 16.
- `sip-password` *`sip-password`*—Simple IP password. Obtain the password from the service provider.
- `sip-user-id` *`sip-id`*—Simple IP user identification. Obtain the username from the service provider.
- `slot` *`sim-slot-number`*—The slot in which the SIM card is inserted. The value can be either 1 or 2.

Required Privilege Level

maintenance

Sample Output

request modem wireless create-profile

```
user@host> request modem wireless create-profile cl-1/0/0 access-point-name apn authentication-
method pap profile-id 2 sip-password 123 sip-user-id userid slot 1
Issued create profile request successfully.
Please use 'show modem wireless profiles' to check profile status
```

Release Information

Command introduced in Junos OS 9.5. The `slot` *`sim-slot-number`* option is introduced in Junos OS 15.1X49-D100.

RELATED DOCUMENTATION

show modem wireless profiles | 22120

request modem wireless fota

IN THIS SECTION

- Syntax | 14684
- Description | 14684
- Required Privilege Level | 14684
- Sample Output | 14684
- Release Information | 14685

Syntax

```
request modem wireless fota interface-name (enable | disable)
```

Description

Enable or disable over-the-air (OTA) firmware upgrade for the modem on the LTE Mini-PIM. OTA firmware upgrade enables automatic and timely upgrade of modem firmware when new firmware versions are available. The OTA upgrade can be enabled or disabled on the LTE Mini-PIM. OTA is disabled by default.

Required Privilege Level

maintenance

Sample Output

request modem wireless fota (enable)

```
user@host> request modem wireless fota cl-1/0/0 enable
Set FOTA on modem succeeded
```

request modem wireless fota (disable)

```
user@host> request modem wireless fota cl-1/0/0 disable
Set FOTA on modem succeeded
```

Release Information

Command introduced in Junos OS 15.1X49-D100.

RELATED DOCUMENTATION

show modem wireless firmware | 22111

request modem wireless sim-lock

IN THIS SECTION

- Syntax | 14685
- Description | 14686
- Options | 14686
- Required Privilege Level | 14686
- Sample Output | 14686
- Release Information | 14686

Syntax

```
request modem wireless sim-lock enable interface-name pin pin
```

Description

Lock the Subscriber Identity Module (SIM) on the Mini-PIM. The SIM lock does not take effect until the next reboot of the services gateway. You can verify the locked mode using the `show modem wireless firmware` command.

> **NOTE**: If there are two SIMs installed on the LTE Mini-PIM, then only the active SIM is locked. After the SIM is locked, it cannot connect to the network. The SIM must be unlocked before it is used to connect to the network.

Options

- *interface-name*—The LTE Mini-PIM is denoted as cl-*x*/0/0, where *x* is the slot number in which the LTE Mini-PIM is installed.
- `pin` *pin*—Four-digit personal identification number (PIN). Obtain the PIN from the service provider.

> **NOTE**: If the PIN is entered incorrectly three consecutive times, the SIM card is blocked. Obtain a PIN unblocking key (PUK) from the service provider.

Required Privilege Level

maintenance

Sample Output

request modem wireless sim-lock

```
user@host> request modem wireless sim-lock enable cl-1/0/0 pin 4321
Issued SIM 2 lock state request successfully.
Please use 'show modem wireless firmware' to check SIM status
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

request modem wireless sim-unlock | 14687

request modem wireless sim-unlock

IN THIS SECTION

- Syntax | 14687
- Description | 14687
- Options | 14688
- Required Privilege Level | 14688
- Sample Output | 14688
- Release Information | 14688

Syntax

```
request modem wireless sim-unlock interface-name pin unlock-code
```

Description

Unlock the Subscriber Identity Module (SIM) on the LTE Mini-PIM. Some service providers lock the SIM to prevent unauthorized access to the service provider's network. If this is the case, you will need to unlock the SIM by using an personal identification number (PIN), which is provided by the service provider. You can verify the unlocked mode using the `show modem wireless firmware` command.

NOTE: If there are two SIM cards installed on the Mini-PIM, then only the active SIM card is unlocked.

The SIM must be unlocked before it can be used to connect to the service provider’s network.

Options

- *interface-name*—The LTE interface is denoted as cl-*x*/0/0, where *x* is the slot number in which the LTE Mini-PIM is installed.
- pin *unlock-code*—Four-digit personal identification number (PIN). Obtain the PIN from the service provider.

> **NOTE**: If the PIN is entered incorrectly three consecutive times, the SIM card is blocked. Obtain a PIN unblocking key (PUK) from the service provider.

Required Privilege Level

maintenance

Sample Output

request modem wireless sim-unlock

```
user@host> request modem wireless sim-unlock cl-1/0/0 pin 1234
Issued SIM 2 unlock request successfully.
Please use 'show modem wireless firmware' to check SIM status
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

request modem wireless sim-lock | 14685

request mpls container-lsp

IN THIS SECTION

- Syntax | 14689
- Description | 14689
- Options | 14689
- Required Privilege Level | 14690
- Output Fields | 14690
- Sample Output | 14690
- Release Information | 14690

Syntax

```
request  mpls container-lsp
<logical-system (all | logical-system-name)>
<name lsp-name>
<adjust-autobandwidth>
<normalization>
```

Description

Manually trigger a bandwidth allocation adjustment for the container label-switched path (LSP).

Options

none	Manually trigger a bandwidth allocation adjustment for all active member LSP paths.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
name *lsp-name*	(Optional) Manually trigger a bandwidth allocation adjustment on the specified member LSP only.

adjust-autobandwidth	(Optional) Request LSP autobandwidth adjustment.
normalization	(Optional) Request container LSP normalization.

Required Privilege Level

clear, maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request mpls container-lsp

```
user@host> request  mpls container-lsp lsp-name normalize
```

request mpls container-lsp

```
user@host> request  mpls container-lsp normalize bandwidth bps
```

Release Information

Command introduced in Junos OS Release 14.2.

Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.

RELATED DOCUMENTATION

show mpls container-lsp | 22136

clear mpls container-lsp | 3711

request mpls lsp adjust-autobandwidth

IN THIS SECTION

- Syntax | 14691
- Syntax (EX and QFX Series Switches) | 14691
- Description | 14691
- Options | 14692
- Additional Information | 14692
- Required Privilege Level | 14692
- Output Fields | 14692
- Sample Output | 14692
- Release Information | 14693

Syntax

```
request mpls lsp adjust-autobandwidth
<logical-system (all | logical-system-name)>
<name lsp-name>
```

Syntax (EX and QFX Series Switches)

```
request mpls lsp adjust-autobandwidth
<name lsp-name>
```

Description

Manually trigger a bandwidth allocation adjustment for active label-switched paths (LSPs).

Without running this command, the bandwidth adjustment is recomputed at a configurable interval. The default interval is 5 minutes. If you do not want to wait for the periodic adjustment (for example, during a software demonstration), this command is useful.

During bandwidth allocation adjustment, the LSP stays up to enable the bandwidth to be changed without dropping any traffic. This functionality is often referred to as *make-before-break*.

Options

none	Manually trigger a bandwidth allocation adjustment for all active LSP paths.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
name *lsp-name*	(Optional) Manually trigger a bandwidth allocation adjustment on the specified LSP only.

Additional Information

For this command to work properly, the following conditions must exist:

- Automatic bandwidth allocation must be enabled on the LSP. The parameters for adjustment interval and maximum average bandwidth are not reset after you issue the `request mpls lsp adjust-autobandwidth` command.
- The difference between the adjusted bandwidth and the current LSP path bandwidth must be greater than the threshold limit.

Required Privilege Level

clear, maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request mpls lsp adjust-auto-bandwidth

```
user@host> request mpls lsp adjust-auto-bandwidth
```

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

auto-bandwidth (MPLS Tunnel) | 2686

Configuring Automatic Bandwidth Allocation for LSPs

request network-access aaa accounting

IN THIS SECTION

- Syntax | 14693
- Description | 14693
- Options | 14694
- Required Privilege Level | 14694
- Sample Output | 14694
- Release Information | 14694

Syntax

```
request network-access aaa accounting (baseline | suspend | resume)
```

Description

Suspend accounting processes; determine a baseline of the statistical details while accounting is suspended; and restart accounting operations after baselining is completed. This command is useful in service provider environments when an upgrade of the server infrastructure is critical and needed immediately. RADIUS Acct-Start, Interim-Update, and Acct-Stop messages are not generated while accounting is suspended; the router does not send any accounting messages to the RADIUS server. While accounting is suspended, subscribers can continue to log in and log out.

Options

baseline (Optional) Determine a baseline of accounting statistics for current subscriber sessions. Applies to only those subscribers for which interim accounting is configured. The router implements the baseline by reading and storing the statistics at the time the baseline is set and then subtracting this baseline when you retrieve baseline-relative statistics after accounting resumes.

resume Restart the accounting processes for all logged-in subscriber sessions after baselining of statistics completes.

suspend Temporarily halt accounting processes for all logged-in subscriber sessions.

Required Privilege Level

view

Sample Output

request network-access aaa accounting suspend

```
user@host> request network-access aaa accounting suspend
```

request network-access aaa accounting baseline

```
user@host> request network-access aaa accounting baseline
```

request network-access aaa accounting resume

```
user@host> request network-access aaa accounting resume
```

Release Information

Command introduced in Junos OS Release 14.2R1.

RELATED DOCUMENTATION

RADIUS Accounting for Subscriber Access

RADIUS Accounting for Subscriber Access

request network-access aaa replay pending-accounting-stops

IN THIS SECTION

- Syntax | 14695
- Description | 14695
- Options | 14695
- Required Privilege Level | 14695
- Output Fields | 14696
- Sample Output | 14696
- Release Information | 14696

Syntax

```
request network-access aaa replay pending-accounting-stops
```

Description

Force the router to attempt contact with the accounting sever immediately, rather than allowing it to wait until the periodic interval has expired.

Options

This command has no options.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request network-access aaa replay pending-accounting-stops

```
user@host> request network-access aaa replay pending-accounting-stops
replay started
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

RADIUS Accounting for Subscriber Access

show accounting pending-accounting-stops | **17011**

request network-access aaa subscriber add session-id

IN THIS SECTION

- Syntax | **14697**
- Description | **14697**
- Options | **14697**
- Required Privilege Level | **14697**
- Output Fields | **14697**
- Sample Output | **14698**
- Sample Output | **14698**
- Release Information | **14698**

Syntax

```
request network-access aaa subscriber add session-id subscriber-session-id service-profile
profile-name
```

Description

Locally activate (provision) a dynamic subscriber service for a subscriber who is currently logged in to the network. Starting in Junos OS Release 18.3R1, when the dynamic service profile is configured with the `profile-type remote-device-service` statement, the service is provisioned on a remote device by the remote device services manager daemon (rdmd).

Options

profile-name Name of service-profile to activate.

subscriber-session-id ID of the subscriber session for which the service will be added.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. Table 114 on page 14697 lists possible error messages that might be returned if the service activation fails.

Table 114: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
`Command failed: reason`	–	–
`Error: AUTHD ISSU in progress`	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.

Table 114: Service Activation/Deactivation Error Messages *(Continued)*

Message	Description	Corrective Action
`Provisioning is already active`	Remote provisioning by a JSRC server or Gx-plus server is active.	–
`Service activation/deactivation already in progress`	Another service activation/deactivation operation is currently in progress.	Wait until the active operation completes and then retry the activation/deactivation operation.
`Session identifier is not for a subscriber session`	The session ID is incorrect.	Verify the correct session ID for the subscriber and then retry the activation/deactivation operation.

Sample Output

request network-access aaa subscriber add session-id service-profile

```
user@host> request network-access aaa subscriber add session-id 49 service-profile service-bronze
Successful completion
```

Sample Output

request network-access aaa subscriber add session-id service-profile (Parameters for Profile on Remote Device)

```
user@host> request network-access aaa subscriber add session-id 131 service-profile
“upstreamBandwidth(100,100,100)”
Successful completion
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

CLI-Activated Subscriber Services

Local and Remote Service Activation and Deactivation Using the CLI

request network-access aaa subscriber delete session-id | 14699

request network-access aaa subscriber delete session-id

IN THIS SECTION

- Syntax | 14699
- Description | 14699
- Options | 14700
- Required Privilege Level | 14700
- Output Fields | 14700
- Sample Output | 14701
- Release Information | 14701

Syntax

```
request network-access aaa subscriber delete session-id subscriber-session-id service-profile
profile-name
```

Description

Deactivate (deprovision) a dynamic subscriber service for a subscriber who is currently logged in to the network. Starting in Junos OS Release 18.3R1, when the dynamic service profile is configured with the `profile-type remote-device-service` statement, the service is deprovisioned on a remote device by the remote device services manager daemon (rdmd).

Options

profile-name Name of the service-profile to deactivate. To deactivate a single instance of a subscriber service that has multiple instances, you can specify the service-profile name and its service parameters.

subscriber-session-id ID of the subscriber session for which the service will be deleted.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. Table 115 on page 14700 lists possible error messages that might be returned if the service deactivation fails.

Table 115: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
`Command failed: reason`	Error condition that caused the command to fail.	Correct the error condition.
`Error: AUTHD ISSU in progress`	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.
`Provisioning is already active`	Remote provisioning by a JSRC server or Gx-plus server is active.	Disable provisioning.
`Service activation/deactivation already in progress`	Another service activation/deactivation operation is currently in progress.	Wait until the active operation completes and then retry the activation/deactivation operation.
`Session identifier is not for a subscriber session`	The session ID is incorrect.	Verify the correct session ID for the subscriber and then retry the activation/deactivation operation.

Sample Output

request network-access aaa subscriber delete session-id service-profile

```
user@host> request network-access aaa subscriber delete session-id 49 service-profile service-silver
Successful completion
```

request network-access aaa subscriber delete session-id service-profile (Deactivating a Single Server Instance)

```
user@host> request network-access aaa subscriber delete session-id 6 service-profile economy-service(up-filter,down-filter)
Successful completion
```

request network-access aaa subscriber delete session-id service-profile (Deactivating All Server Instances)

```
user@host> request network-access aaa subscriber delete session-id 6 service-profile economy-service
Successful completion
```

request network-access aaa subscriber add session-id service-profile (Parameters for Profile on Remote Device)

```
user@host> request network-access aaa subscriber delete session-id 131 service-profile
“upstreamBandwidth(100,100,100)”
Successful completion
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

CLI-Activated Subscriber Services

Local and Remote Service Activation and Deactivation Using the CLI

Deactivating a Single Instance of a Subscriber Service

Deactivating All Instances of a Subscriber Service

request network-access aaa subscriber add session-id | **14696**

request network-access aaa subscriber modify session-id

IN THIS SECTION

- Syntax | **14702**
- Description | **14702**
- Options | **14702**
- Required Privilege Level | **14703**
- Output Fields | **14703**
- Sample Output | **14703**
- Release Information | **14703**

Syntax

```
request network-access aaa subscriber modify session-id subscriber-session-id predefined-
variable variable-option
```

Description

Modify a predefined variable that is applied to a subscriber who is currently logged in to the network.

Options

predefined-variable Name of the predefined variable that you want to modify.

subscriber-session-id ID of the subscriber session.

variable-option Name of the variable option that you want to apply to the predefined variable.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. Table 116 on page 14703 lists possible messages that might be returned.

Table 116: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
`Successful completion`	Variable was successfully modified	–
`Error: AUTHD ISSU in progress`	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.

Sample Output

request network-access aaa subscriber modify session-id

```
user@host> request network-access aaa subscriber modify session-id 49 junos-cos-traffic-control-
profile TCP-gold
Successful completion
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers

CLI-Activated Subscriber Services

request network-access aaa subscriber set session-id

IN THIS SECTION

- Syntax | 14704
- Description | 14704
- Options | 14704
- Required Privilege Level | 14705
- Output Fields | 14705
- Sample Output | 14705
- Release Information | 14705

Syntax

```
request network-access aaa subscriber set session-id subscriber-session-id provisioning-state
none
```

Description

Release control of the PCRF over the specified subscriber session. In response, AAA clears the subscriber's provisioning state and sends a terminated request to the PCRF indicating the subscriber is no longer available.

Options

subscriber-session-id	ID of the subscriber session.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. Table 117 on page 14705 lists possible error messages that might be returned if the service activation fails.

Table 117: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
`Error: AUTHD ISSU in progress`	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.
`Service activation/ deactivation already in progress`	Another service activation/ deactivation operation is currently in progress.	Wait until the active operation completes and then retry the activation/deactivation operation.
`Session identifier is not for a subscriber session`	The session ID is incorrect.	Verify the correct session ID for the subscriber and then retry the activation/deactivation operation.

Sample Output

request network-access aaa subscriber set session-id

```
user@host> request network-access aaa subscriber set session-id session-id 49 provisioning-state
none
Successful completion
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Disabling PCRF Control of a Subscriber Session

Local and Remote Service Activation and Deactivation Using the CLI

request network-access aaa saml export-sp-metadata

IN THIS SECTION

- Syntax | 14706
- Description | 14706
- Options | 14706
- Required Privilege Level | 14706
- Release Information | 14707

Syntax

```
request network-access aaa saml export-sp-metadata access-profile-name access-profile-name
```

Description

Export Security Assertion Markup Language (SAML) authentication service provider metadata to a file.

Options

access-profile-name ***access-profile-name***	Exports service provider configuration metadata to /var/tmp/*access-profile-name*.xml

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

saml | 16281

request network-access aaa saml load-idp-metadata | 14707

request network-access aaa saml load-idp-metadata

IN THIS SECTION

- Syntax | 14707
- Description | 14707
- Options | 14708
- Required Privilege Level | 14708
- Release Information | 14708

Syntax

```
request network-access aaa saml load-idp-metadata
<filename filename>
<metadata-id metadata-id>
```

Description

Load Security Assertion Markup Language (SAML) authentication identity provider (IdP) metadata details.

Options

filename *filename* — Provide filename along with the absolute path to load identity provider (IdP) metadata details such as /var/tmp/filename.xml.

metadata-id *metadata-id* — Provide metadata identifier of the identity provider (IdP).

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

saml | 16281

clear network-access aaa saml idp-metadata | 3759

request node halt (Junos OS Evolved)

IN THIS SECTION

- Syntax | 14709
- Description | 14709
- Options | 14709
- Required Privilege Level | 14709
- Sample Output | 14710
- Release Information | 14711

Syntax

```
request node halt node-name
<(at time | in minutes)>
<message message>
```

Description

Use this command to halt a Routing Engine. Halt instructs the hardware to stop all CPU functions but leave the node in a powered-on, standby state. To un-halt the node, do one of the following:

- Issue the `request chassis cb slot slot-number online` operational mode command from the primary Routing Engine.
- Log in to the console port for that node and press any key to reboot the node.

> **NOTE**: We do not recommend leaving a node halted for a long period of time, because the node is not available as a backup in case something happens to the primary Routing Engine.

Options

node-name	Specify the Routing Engine node (re0 | re1) to halt. You can halt the primary Routing Engine.
(at *time* | in *minutes*)	(Optional) Specify when the action should occur, either in time, in *hh:mm* format, or in number of minutes.
message *message*	(Optional) Message to display to all users.

Required Privilege Level

view

Sample Output

request node halt re1

On the primary Routing Engine:

```
user@host-re0> request node halt re1
Halt the node ? [yes,no] (no) yes
*** System shutdown message from user@host-re0***


Shutdown at Thu Apr 29 16:31:01 2021
 .
{master}
user@host-re0>
```

When logged into the console port on the node during the halt:

```
[...output truncated, processes stopped...]
Shutdown target is 'halt'
Checking on root FS daemon ........Done
Cleaning up root FS daemon...done
/sbin/bom.sh: line 1242: /tmp/mp_aid: Read-only file system
/sbin/bom.sh: line 1793: /tmp/sku_mode_log: Read-only file system
/sbin/bom.sh: line 1242: /tmp/mp_aid: Read-only file system
/sbin/bom.sh: line 1793: /tmp/sku_mode_log: Read-only file system
Clearing linux_up bit
warning - 'debugfs' was not mounted
Deactivating swap...
Putting SSD devices in standby mode...

/dev/sdb:
 issuing standby command

/dev/sdc:
 issuing standby command
SG_IO: bad/missing sense data, sb[]:  70 00 05 00 00 00 00 14 00 00 00 00 20 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00

/dev/sda:
 issuing standby command
```

```
The operating system has halted.

Please press any key to reboot...
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- request node (offline | online) (Junos OS Evolved) | 14711
- request node power-off (Junos OS Evolved) | 14713
- request node power-on (Junos OS Evolved) | 14715
- request node reboot (re0 | re1) (Junos OS Evolved) | 14717

request node (offline | online) (Junos OS Evolved)

IN THIS SECTION

- Syntax | 14711
- Description | 14712
- Options | 14712
- Required Privilege Level | 14712
- Sample Output | 14712
- Release Information | 14713

Syntax

```
request node (offline | online) node-name
```

Description

Use this command to change the node status to offline or online.

- To add a node to the system, set the node status to `online`.
- To remove a node from the system, set the node status to `offline`.

  You can use the `offline` option to stop all applications on the node and move them to other nodes if applicable. The node is not allowed to join the system until the node is brought online using the `request node online` command.

  > **NOTE**: We do not recommend leaving the secondary Routing Engine offline for a long period of time, because the secondary Routing Engine is not available as a backup in case something happens to the primary Routing Engine.

  When you use the `request node offline` command for FPC nodes, the node is powered off. When used for a Routing Engine node, the node is powered off.

Options

node-name	Specify the node name for the Routing Engine or the FPC. You can take the primary Routing Engine (re0) offline. The backup or secondary node is re1. You can not specify this command for a device that supports only one Routing Engine.
(offline | online)	Change the node status to online or offline. When you specify the online option, the node reboots, which can take a few minutes.

Required Privilege Level

view

Sample Output

request node offline re1

```
user@host-re0> request node offline re1
This may take a few minutes. Offline the node ? [yes,no] (no) yes

Node re1 is set to be offline
```

request node online re1

```
user@host-re0> request node online re1
This may take a few minutes. Online the node ? [yes,no] (no) yes

Node re1 is set to be online
```

Release Information

Statement introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- request node halt (Junos OS Evolved) | 14708
- request node power-off (Junos OS Evolved) | 14713
- request node power-on (Junos OS Evolved) | 14715
- request node reboot (re0 | re1) (Junos OS Evolved) | 14717
- request system application (Junos OS Evolved) | 15042

request node power-off (Junos OS Evolved)

IN THIS SECTION

- Syntax | 14714
- Description | 14714
- Options | 14714
- Required Privilege Level | 14714
- Sample Output | 14714
- Release Information | 14715

Syntax

```
request node power-off node-name
<(at time | in minutes)>
<message message>
```

Description

Use this command to power off a Routing Engine, stopping the CPU and commencing a complete shutdown.

> **NOTE**: We do not recommend leaving a node powered off for a long period of time, because the node is not available as a backup in case something happens to the primary Routing Engine.

Options

node-name	Specify the Routing Engine node (re0 | re1) to shut down. You can shut down the primary Routing Engine.
(at *time* | in *minutes*)	(Optional) Specify when the action should occur, either in time, in *hh:mm* format, or in number of minutes.
message *message*	(Optional) Message to display to all users.

Required Privilege Level

view

Sample Output

request node power-off re1

On the primary Routing Engine:

```
user@host-re0> request node power-off re1
Power-off the node ? [yes,no] (no) yes
```

```
*** System shutdown message from user@host-re0***


Shutdown at Fri Apr 30 10:47:01 2021
 .

{master}
user@host-re0>
```

When logged in to the console port on the node:

```
[...output truncated, stopping processes...]

Shutdown target is 'poweroff'
Checking on root FS daemon ........Done
Cleaning up root FS daemon...done
Powering off.
[  285.750267] reboot: Power down
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- request node halt (Junos OS Evolved) | **14708**
- request node (offline | online) (Junos OS Evolved) | **14711**
- request node power-on (Junos OS Evolved) | **14715**
- request node reboot (re0 | re1) (Junos OS Evolved) | **14717**

request node power-on (Junos OS Evolved)

IN THIS SECTION

- Syntax | **14716**

- Description | 14716
- Options | 14716
- Required Privilege Level | 14716
- Sample Output | 14716
- Release Information | 14717

Syntax

```
request node power-on node-name
```

Description

Use this command to power on a Routing Engine, re0 or re1. The node reboots, which can take a few minutes. You can not specify this command for a device that supports only one Routing Engine.

Options

node-name	Specify the Routing Engine node to power on.

Required Privilege Level

view

Sample Output

request node power-on re1

```
user@host-re0> request node power-on re1
Power-on the node ? [yes,no] (no) yes
OK
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- request node halt (Junos OS Evolved) | 14708
- request node (offline | online) (Junos OS Evolved) | 14711
- request node power-off (Junos OS Evolved) | 14713
- request node reboot (re0 | re1) (Junos OS Evolved) | 14717

request node reboot (re0 | re1) (Junos OS Evolved)

IN THIS SECTION

- Syntax | 14717
- Description | 14718
- Options | 14718
- Additional Information | 14718
- Required Privilege Level | 14719
- Sample Output | 14719
- Release Information | 14719

Syntax

```
request node reboot (re0 |re1)
<(at time | in minutes)>
<(disk1 | disk2)>
<message message>
<network1>
<usb>
```

Description

Use this command to reboot one of the Routing Engines in a system. You can reboot the primary Routing Engine with this command. To reboot all nodes at once, use the `request system reboot` command.

Options

(at *time* | in *minutes*) (Optional) Specify when the reboot is performed, either at a particular time, in *hh:mm* format, or in number of minutes.

(disk1 | disk2) (Optional) Boot from the primary solid-state drive (SSD) (`disk1`) or the secondary SSD (`disk2`).

Default: `disk1`

message *message* (Optional) Message to display to all users.

(re0 | re1) Specify which Routing Engine to reboot. You can reboot the primary Routing Engine using this command.

network1 (Optional) Network boot using PXE over port1.

usb (Optional) Boot from the USB device.

Additional Information

Starting in Junos OS Evolved 23.4R1, when rebooting a node with this command, the Layer 2 applications `l2ald`, `l2ald-agent`, `l2cpd`, and `l2cpd-agent` are initialized and started only if any of the following configuration hierarchy levels contain any configuration statements:

- `[edit interface interface-name unit number family ethernet-switching]`
- `[edit vlans]`
- `[edit routing-instance instance-name instance-type virtual-switch]`
- `[edit routing-instance instance-name instance-type mac-vrf]`
- `[edit protocols l2-learning]`

Additionally, `l2cpd`, and `l2cpd-agent` are initialized and started if the `[edit protocols lldp]` hierarchy level contains any configuration statements.

As a result of this change, if your configuration already contains these configuration statements and you then delete all of the above configuration statements, these Layer 2 applications stop running.

Required Privilege Level

view

Sample Output

request node reboot re1

On the primary Routing Engine:

```
user@host-re0> request node reboot re1
This may affect traffic in system. Proceed ? [yes,no] (no) yes

*** System shutdown message from user@host-re0***


Shutdown at Fri Apr 30 10:47:01 2021
 .
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- request node halt (Junos OS Evolved) | 14708
- request node (offline | online) (Junos OS Evolved) | 14711
- request node power-off (Junos OS Evolved) | 14713
- request node power-on (Junos OS Evolved) | 14715
- request system application (Junos OS Evolved) | 15042

request packet-capture start

IN THIS SECTION

- Syntax | 14720
- Description | 14720
- Options | 14721
- Required Privilege Level | 14722
- Output Fields | 14722
- Sample Output | 14724
- Release Information | 14726

Syntax

```
request packet-capture start
<bidirectional bidirectional>
<capture-file capture-file>
<count count>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface>
<maximum-capture-size maximum-capture-size>
<protocol protocol-number>
<size size>
<source-port source-port>
<source-prefix source-prefix>
```

Description

Captures packet information from the operational mode. You can execute the packet capture from the operational mode with minimal impact to the production system without committing the configurations. You can define the packet filter to trace the type of traffic based on your requirement.

You can request for only one packet capture at a time, and you need to stop each packet capture request before you give an another packet capture request. For each packet capture request, you need to give the respective show command to view the output.

In chassis cluster mode, you must request operational mode packet capture from the primary node. Operational mode packet capture does not function when executed from the secondary node.

Options

bidirectional Allows to collect bidirectional information.

capture-file Name of the capture file. It generates a pcap file, which can support Wireshark by adding the suffix '.pcap'.

count Number of packets to capture.

- **Range:** 10 through 1000000
- **Default:** 100

destination-port TCP/UDP destination port.

destination-prefix Destination IPv4/IPv6 address prefix to filter the packets.

interface Name of the logical interface.

maximum-capture-size The maximum size of packet capture. The packet truncates if the capture size is more than the specified capture size.

- **Range:** 68 through 10000
- **Default:** 1514

protocol-number Numeric protocol value.

- **Range:** 0 through 255
- **Default:** 0

size The maximum size of packet capture file.

- **Range:** 10240 through 1073741824
- **Default:** 50M

source-port TCP/UDP source port.

source-prefix Source IPv4/IPv6 address prefix to filter the packets.

Required Privilege Level

view

Output Fields

Table 118 on page 14722 lists the output fields for the `request packet-capture start` command. Output fields are listed in the approximate order in which they appear.

Table 118: request packet-capture start Output Fields

Field Name	Field Description
`packet capture is running`	Status of the packet capture session: • `0`: No packet capture is performing. • `1`: Packet capture is performing.
`counter of packet capture`	Number of packets captured.
`packet capture capture-file name`	Name, location, size and capture length of the packet capture file.

Table 118: request packet-capture start Output Fields *(Continued)*

Field Name	Field Description
`Flag`	Decimal value corresponding to the type of filter configured. Convert this value to binary or hexadecimal format such as: • 0x000 (No filters) • 0x001 (Source IP) • 0x002 (Destination IP) • 0x004 (Protocol) • 0x008 (Source port) • 0x010 (Destination port) • 0x040 (Filter is valid) • 0x200 (IPv6 Source IP) • 0x400 (IPv6 Destination IP) • 0x800 (Interface) For example if the flag value is `72`, convert it to binary or hexadecimal format: 72 = 0x48 • 0x040 (Filter is valid) • 0x008 (Source port) This means that the filter is valid and only configured with the filter source port.
`Source`	Source IPv4/IPv6 address.
`Destination`	Destination IPv4/IPv6 address.
`Interface`	Name of the logical interface.

Table 118: request packet-capture start Output Fields *(Continued)*

Field Name	Field Description
bidirectional	Status of the bidirectional keyword configuration: • 0: Not configured. • 1: Configured.
Protocol	Numeric protocol value.
rotate	Number of capture files.
snap_len	The maximum packet capture size.
ifl	If index value of interface.

Sample Output

Packet capture using protocol, destination port, destination prefix, and bidirectional options

```
root@host> request packet-capture start protocol 17 destination-port 53 destination-prefix
203.0.113.0 bidirectional
{primary:node1}
root@host> show packet-capture status
    packet capture is running: 1
    counter of packet capture: 100
    packet capture capture-file name: /var/log/packet-capture.pcap size: 52428800 rotate: 0
snap_len: 1514
  bidirectional: 1
    Flag: 86 Protocol: udp
    Source: 0.0.0.0 0.0.0.0 (port 0~65535)
    Destination: 203.0.113.0 255.255.255.255 (port 53~53)
    Interface: None ifl: 0
     Flag: 77 Protocol: udp
    Source: 203.0.113.0 255.255.255.255 (port 53~53)
```

```
    Destination: 0.0.0.0 0.0.0.0 (port 0~65535)
    Interface: None ifl: 0
```

Packet capture using interface

```
user@host> request packet-capture start capture-file cli-e2e interface xe-1/0/2.0
user@host> show packet-capture status
packet capture is running: 1
    counter of packet capture: 100
    packet capture capture-file name: /var/log/cli-e2e.pcap size: 52428800 rotate: 0 snap_len:
1514
  bidirectional: 0
    Flag: 2112
    Source: 0.0.0.0 0.0.0.0 (port 0~65535)
    Destination: 0.0.0.0 0.0.0.0 (port 0~65535)
    Interface: xe-1/0/2.0 ifl: 68
```

Packet capture using source prefix

```
user@host> request packet-capture start capture-file cli-e2e source-prefix 192.0.2.0
user@host> show packet-capture status
packet capture is running: 1
counter of packet capture: 100
packet capture capture-file name: /var/log/cli-e2e.pcap size: 52428800 rotate: 0 snap_len: 1514
bidirectional: 0
Flag: 65
Source: 192.0.2.0.2 255.255.255.255 (port 0~65535)
Destination: 0.0.0.0 0.0.0.0 (port 0~65535)
Interface: None ifl: 0
```

Packet capture using source port

```
user@host> request packet-capture start capture-file cli-e2e count 100 source-port telnet
user@host> show packet-capture status
packet capture is running: 1
counter of packet capture: 100
packet capture capture-file name: /var/log/cli-e2e.pcap size: 52428800 rotate: 0 snap_len: 1514
bidirectional: 0
```

```
Flag: 72
Source: 0.0.0.0 0.0.0.0 (port 23~23)
Destination: 0.0.0.0 0.0.0.0 (port 0~65535)
Interface: None ifl: 0
```

Packet capture using protocol

```
user@host> request packet-capture start protocol 6
user@host> show packet-capture status
packet capture is running: 1
counter of packet capture: 100 packet
capture capture-file name: /var/log/packet-capture.pcap size: 52428800 rotate: 0 snap_len: 1514
bidirectional: 0
Flag: 68
Protocol: tcp
Source: 0.0.0.0 0.0.0.0 (port 0~65535)
Destination: 0.0.0.0 0.0.0.0 (port 0~65535)
Interface: None ifl: 0
```

Release Information

Statement introduced in Junos OS Release 19.3R1.

request packet-capture stop

IN THIS SECTION

- Syntax | 14727
- Description | 14727
- Required Privilege Level | 14727
- Output Fields | 14727
- Sample Output | 14727
- Release Information | 14727

Syntax

```
request packet-capture stop
```

Description

Stops the packet capture request from the operational mode. You can execute the packet capture from the operational mode with minimal impact to the production system without committing the configurations. You must stop the packet capture request to generate the packet capture report.

Required Privilege Level

view

Output Fields

This command produces no output.

Sample Output

request packet-capture stop

```
user@host> request packet-capture stop
```

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

request packet-capture start | 14720

request path-computation-client active-pce

IN THIS SECTION

- Syntax | 14728
- Description | 14728
- Options | 14728
- Required Privilege Level | 14728
- Output Fields | 14728
- Release Information | 14729

Syntax

```
request path-computation-client active-pce pce-id
```

Description

Request a new active Path Computation Element (PCE).

Options

pce-id	Unique user defined ID for this PCE.
retry-delegation	Retry label-switched path (LSP) delegation

Required Privilege Level

request

Output Fields

This command produces no output. To verify the operation of the command, run the "show path-computation-client active-pce" on page 22922 before and after running the `request path-computation-client active-pce` command.

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

show path-computation-client active-pce | 22922

request pim multicast-tunnel rebalance

IN THIS SECTION

- Syntax | 14729
- Syntax (EX Series Switches) | 14729
- Description | 14730
- Options | 14730
- Required Privilege Level | 14730
- Output Fields | 14730
- Release Information | 14730

Syntax

```
request pim multicast-tunnel rebalance
<instance instance-name>
<logical-system (all | logical-system-name>
```

Syntax (EX Series Switches)

```
request pim multicast-tunnel rebalance
<instance instance-name>
```

Description

Rebalance the assignment of multicast tunnel encapsulation interfaces across available tunnel-capable PICs or across a configured list of tunnel-capable PICs. You can determine whether a rebalance is necessary by running the `show pim interfaces instance` *`instance-name`* command.

Options

none	Re-create and rebalance all tunnel interfaces for all routing instances.
instance ***instance-name***	Re-create and rebalance all tunnel interfaces for a specific instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

maintenance

Output Fields

This command produces no output. To verify the operation of the command, run the `show pim interface instance` *`instance-name`* before and after running the `request pim multicast-tunnel rebalance` command.

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

> show pim interfaces | **23066**
>
> *Load Balancing Multicast Tunnel Interfaces Among Available PICs*

request pppoe connect

IN THIS SECTION

- Syntax | 14731
- Description | 14731
- Options | 14731
- Required Privilege Level | 14731
- Output Fields | 14731
- Sample Output | 14732
- Release Information | 14732

Syntax

```
request pppoe connect
```

Description

Connect all sessions that are down.

Options

`pppoe interface name`— (Optional) Connect to a specified session.

Required Privilege Level

maintenance

Output Fields

When you enter this command, this command returns no output.

Sample Output

request pppoe connect

```
user@host> request pppoe connect
```

Release Information

Statement supported on SRX300, SRX320, SRX340, and SRX345 is introduced in Junos OS Release 15.1X49-D60.

Statement supported on SRX1500 and vSRX Virtual Firewall instances is introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

- Understanding PPPoE Interfaces
- Example: Configuring PPPoE Interfaces

request pppoe disconnect

IN THIS SECTION

- Syntax | 14733
- Description | 14733
- Options | 14733
- Required Privilege Level | 14733
- Output Fields | 14733
- Sample Output | 14733
- Release Information | 14733

Syntax

```
request pppoe disconnect
```

Description

Disconnect all active sessions.

Options

`session id` — (Optional) Disconnect the session for which the session ID is specified.

`pppoe interface name`— (Optional) Disconnect the session for a specific pppoe interface name.

Required Privilege Level

maintenance

Output Fields

When you enter this command, this command returns no output.

Sample Output

request pppoe disconnect

`user@host>` **request pppoe disconnect**

Release Information

Statement supported on SRX300, SRX320, SRX340, and SRX345 is introduced in Junos OS Release 15.1X49-D60.

Statement supported on SRX1500 and vSRX Virtual Firewall instances is introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Understanding PPPoE Interfaces

Example: Configuring PPPoE Interfaces

request protection-group ethernet-aps

IN THIS SECTION

- Syntax | 14734
- Description | 14734
- Options | 14734
- Required Privilege Level | 14735
- Sample Output | 14735
- Release Information | 14735

Syntax

```
request-protection-group ethernet-aps (clear | exercise | force-switch | lockout | logical-
system | maintenance-association | maintenance-domain | manual-switch | switch-to-protect |
switch-to-working)
```

Description

Specify Ethernet Automatic Protection Switching (APS) related operations.

Options

clear	Clears active local command
exercise	Excercise APS protocol
force-switch	Force switch data path to use protection path.
lockout	Lockout protection.
logical-system	Name of logical system, or 'all'.
maintenance-association	Name of maintenance association.
maintenance-domain	Name of maintenance domain.

manual-switch	Manual switch data path to use protection path.
switch-to-protect	Switch data path to use protection path.
switch-to-working	Switch data path to use working path.

Required Privilege Level

view

Sample Output

request-protection-group ethernet-aps manual-switch

```
user@host> request-protection-group ethernet-aps manual-switch maintenance-domain mdut1 logical-
system all maintenance-association maut1
Successfully switched to protect.
```

request-protection-group ethernet-aps clear

```
user@host> request protection-group ethernet-aps md nsn maintenance-association working clear
Manual Switch cleared.
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Ethernet Automatic Protection Switching Overview

request protection-group ethernet-ring

IN THIS SECTION

- Syntax | 14736
- Description | 14736
- Options | 14736
- Additional Information | 14736
- Required Privilege Level | 14736
- Sample Output | 14737
- Release Information | 14737

Syntax

```
request protection-group ethernet-ring (clear | force-switch | manual-switch)
```

Description

Specify Ethernet Automatic Protection Switching (APS) related operations on an Ethernet ring. APS is a linear protection scheme designed to protect VLAN based Ethernet networks.

Options

clear	Clears active local command
force-switch	Force switch data path to use protection path.
manual-switch	Manual switch data path to use protection path.

Additional Information

Required Privilege Level

view

Sample Output

request-protection-group ethernet-ring manual-switch

```
user@host> request-protection-group ethernet-ring manual-switch interface-name ge-0/0/1
Local MS command is completed for pg (erp1) intf (ge-0/0/1.0)
```

request-protection-group ethernet-ring force-switch

```
user@host> request-protection-group ethernet-ring force-switch interface-name ge-0/0/1
```

```
Local FS command is completed for pg (erp1) intf (ge-0/0/1.0)
```

request protection-group ethernet-ring clear

```
user@host> request protection-group ethernet-ring clear
CLEAR request is completed for pg (erp1)
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Understanding Ethernet Ring Protection Switching Functionality

Ethernet Automatic Protection Switching Overview

request-rate (Access)

IN THIS SECTION

- Syntax | 14738
- Hierarchy Level | 14738
- Description | 14738
- Options | 14738
- Required Privilege Level | 14739
- Release Information | 14739

Syntax

```
request-rate rate;
```

Hierarchy Level

```
[edit access radius-options]
```

Description

Configure the number of requests the router can send per second to all configured RADIUS servers collectively. By limiting the flow of requests from the router to the RADIUS servers, you can prevent the RADIUS servers from being flooded with requests.

Options

rate—Number of requests per second.

- **Range:** 100 through 4000 requests per second
- **Default:** 500 requests per second

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

request redundant-power-system multi-backup

IN THIS SECTION

- Syntax | 14739
- Description | 14740
- Required Privilege Level | 14740
- Sample Output | 14740
- Release Information | 14740

Syntax

EX2200 switch:

```
request redundant-power-system multi-backup
request redundant-power-system no-multi-backup
```

EX3300 switch:

```
 request redundant-power-system multi-backup member member-number

request redundant-power-system no-multi-backup member member-number
```

Description

Configure a redundant power system (RPS) to back up six non-Power-over-Ethernet (PoE) powered switches instead of the default which is to back up three PoE-powered switches.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Sample Output

request redundant-power-system multi-backup

```
user@switch> request redundant-power-system multi-backup member 1
Sending multi-backup setting to RPS
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

EX Series Redundant Power System Hardware Overview

request-response

IN THIS SECTION

- Syntax | 14741
- Hierarchy Level | 14741
- Description | 14742
- Required Privilege Level | 14742
- Release Information | 14742

Syntax

```
request-response {
    grpc {
        max-connections max-connections;
        routing-instance routing-instance;
        ssl {
            address ip-address;
            hot-reloading;
            local-certificate (System) local-certificate;
            mutual-authentication {
                certificate-authority certificate-authority-profile-name;
                client-certificate-request (no-certificate | request-certificate | request-
certificate-and-verify | require-certificate | require-certificate-and-verify);
            }
            port port;
            use-pki;
        }
    }
}
```

Hierarchy Level

```
[edit system services extension-service]
```

Description

Allow request-response API execution.

Statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`grpc` option introduced in Junos OS Release 16.2.

request restart packet-forwarding-engine

IN THIS SECTION

- Syntax | 14742
- Description | 14743
- Options | 14743
- Required Privilege Level | 14743
- Output Fields | 14743
- Release Information | 14743

Syntax

```
request restart packet-forwarding-engine
```

Description

Restart the packet forwarding engine (PFE).

Options

This command has no options.

Required Privilege Level

maintenance

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 23.2R2 for EX2300, EX4100, EX4300-MP and EX4400 platforms.

request routing-engine

IN THIS SECTION

- Syntax | 14744
- Description | 14744
- Options | 14744
- Required Privilege Level | 14744
- Sample Output | 14744
- Release Information | 14744

Syntax

```
request routing-engine login (backup | master | other-routing-engine | re0 | re1)
```

Description

On a router with two Routing Engines, specify a tty connection for login.

Options

backup	Log in to the backup Routing Engine.
master	Log in to the master (primary) Routing Engine.
other-routing-engine	Log in to the other Routing Engine.
re0	Log in to the Routing Engine in slot 0.
re1	Log in to the Routing Engine in slot 1.

Required Privilege Level

maintenance

Sample Output

request routing-engine login other-routing-engine

user@host> **request routing-engine login other-routing-engine**

```
--- JUNOS 21.2R1.10 Kernel 64-bit  JNPR-12.1-20210529.2f59a40_build
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

request | 14411

request security certificate enroll (Signed)

IN THIS SECTION

- Syntax | 14745
- Description | 14745
- Options | 14745
- Required Privilege Level | 14746
- Output Fields | 14746
- Sample Output | 14746
- Release Information | 14747

Syntax

```
request security certificate enroll filename filename   subject subject
alternative-subject alternative-subject   certification-authority certification-authority
encoding (binary | pem) key-file key-file   domain-name domain-name
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Obtain a signed certificate from a certificate authority (CA). The signed certificate validates the CA and the owner of the certificate. The results are saved in a specified file to the **/var/etc/ikecert** directory.

NOTE: For FIPS mode, the digital security certificates must be compliant with the National Institute of Standards and Technology (NIST) SP 800-131A standard. The `request security key-pair` command is deprecated and not available with Junos in FIPS mode because it generates RSA and DSA keys with sizes of 512 and 1024 bits that are not compliant with the NIST SP 800-131A standard.

Options

filename *filename* File that stores the certificate.

subject *subject*	Distinguished name (**dn**), which consists of a set of components—for example, an organization (**o**), an organization unit (**ou**), a country (**c**), and a locality (**l**).
alternative-subject *alternative-subject*	Tunnel source address.
certification-authority *certification-authority*	Name of the certificate authority profile in the configuration.
encoding (binary | pem)	File format used for the certificate. The format can be a binary file or privacy-enhanced mail (PEM), an ASCII base64-encoded format. The default format is binary.
key-file *key-file*	File containing a local private key.
domain-name *domain-name*	Fully qualified domain name.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security certificate enroll filename subject alternative-subject certification-authority key-file domain-name (Signed)

```
user@host> request security certificate enroll filename host.crt subject c=uk,o=london
alternative-subject 10.50.1.4 certification-authority verisign                  key-file
host-1.prv domain-name host.example.com
CA name: example.com CA file: ca_verisign
local pub/private key pair: host.prv
subject: c=uk,o=london domain name: host.example.com
alternative subject: 10.50.1.4
Encoding: binary
Certificate enrollment has started. To view the status of your enrollment, check the key
management process (kmd) log file at /var/log/kmd. <--------------
```

Release Information

Command introduced before Junos OS Release 7.4.

request security certificate enroll (Unsigned)

IN THIS SECTION

- Syntax | 14747
- Description | 14747
- Options | 14747
- Required Privilege Level | 14748
- Output Fields | 14748
- Sample Output | 14748
- Release Information | 14748

Syntax

```
request security certificate enroll filename filename   ca-file ca-file   ca-name ca-
name
encoding (binary | perm) url url
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Obtain a certificate from a certificate authority (CA). The results are saved in a specified file to the **/var/etc/ikecert** directory.

Options

filename ***filename***	File that stores the public key certificate.
ca-file ***ca-file***	Name of the certificate authority profile in the configuration.

ca-name *ca-name*	Name of the certificate authority.
encoding (binary | pem)	File format used for the certificate. The format can be a binary file or privacy-enhanced mail (PEM), an ASCII base64-encoded format. The default value is **binary**.
url *url*	Certificate authority URL.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security certificate enroll filename ca-file ca-name url (Unsigned)

```
user@host> request security certificate enroll filename ca_verisign ca-file verisign ca-name
example.com urlxyzcompany URL
http://<verisign ca-name xyzcompany url>/cgi-bin/pkiclient.exe CA name: example.com CA file:
verisign Encoding: binary
Certificate enrollment has started. To view the status of your enrollment, check the key
management process (kmd) log file at /var/log/kmd. <--------------
```

Release Information

Command introduced before Junos OS Release 7.4.

request security datapath-debug capture start

IN THIS SECTION

- Syntax | 14749
- Description | 14749
- Required Privilege Level | 14749
- Output Fields | 14749
- Sample Output | 14750
- Release Information | 14750

Syntax

```
request security datapath-debug capture start
```

Description

Start the data path debugging capture.

NOTE: Data path debugging is supported on SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security datapath-debug capture start

```
user@host> request  security datapath-debug capture start
datapath-debug capture started on file
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Understanding Data Path Debugging for Logical Systems

request security datapath-debug capture stop

IN THIS SECTION

- Syntax | 14750
- Description | 14751
- Required Privilege Level | 14751
- Output Fields | 14751
- Sample Output | 14751
- Release Information | 14751

Syntax

```
request security datapath-debug capture stop
```

Description

Stop the data path debugging capture.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security datapath-debug capture stop

```
user@host> request  security datapath-debug capture stop
datapath-debug capture successfully stopped, use show security datapath-debug
capture to view
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Understanding Data Path Debugging for Logical Systems

request security idp jist-conversion

IN THIS SECTION

- Syntax | 14752
- Description | 14752

- Options | 14752
- Required Privilege Level | 14753
- Output Fields | 14753
- Sample Output | 14753
- Release Information | 14754

Syntax

```
request security idp jist-conversion
<input-file>
<output-file>
```

Description

Use the request security idp jist-conversion command to generate the set commands. These set commands are generated as the CLI output.

Juniper Integration of Snort Tool (JIST) is a tool that converts the Snort IPS rules into Juniper IDP custom attack signatures. JIST is included in Junos OS by default. The tool supports Snort version 2 and version 3 rules. JIST converts the Snort rules with snort-ids into equivalent custom attack signatures on Junos with respective snort-ids as the custom attack names, JIST generates set commands equivalent to the Snort IPS rules. To load the set commands, use the load set terminal statement or copy the commands and paste in the configuration mode, then commit. You can then configure the existing IDP policy with the converted custom attack signatures.

Snort IPS rule files that are not converted are placed at **/tmp/jist-failed.rules**. The error log files generated during the conversion are placed at **/tmp/jist-error.log**.

Options

Input-file (Mandatory)— Specifies the pathname of the snort file that needs to be processed. For example, **/var/tmp/snort_rule.rules**.

Output-file (Optional)— Writes the set commands into a file. You can load the set commands using load set <output-file> configuration statement.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security idp jist-conversion

```
user@host> request security idp jist-conversion input-file /var/preserve/s.rules
```

```
set security idp custom-attack TCP-SNORT-77600859 severity minor

  set security idp custom-attack TCP-SNORT-77600859 attack-description "Backdoor.BEACON"

  set security idp custom-attack TCP-SNORT-77600859 attack-type chain member m01 attack-type
signature context stream

  set security idp custom-attack TCP-SNORT-77600859 attack-type chain member m01 attack-type
signature direction server-to-client

  set security idp custom-attack TCP-SNORT-77600859 attack-type chain member m01 attack-type
signature content pattern "<meta name=\"msvalidate\.01\" content=
\"ECEE9516DDABFC7CCBBF1EACC04CAC20\">"

  set security idp custom-attack TCP-SNORT-77600859 attack-type chain member m02 attack-type
signature context stream

  set security idp custom-attack TCP-SNORT-77600859 attack-type chain member m02 attack-type
signature direction server-to-client

  set security idp custom-attack TCP-SNORT-77600859 attack-type chain member m02 attack-type
signature content pattern "<meta name=\"google-site-verification\" content=
\"CD5EF1FCB54FE29C838ABCBBE0FA57AE\">"
```

```
set security idp custom-attack TCP-SNORT-77600859 attack-type chain protocol-binding tcp
minimum-port 1 maximum-port 65535
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

show security idp jist-package-version | 24599

request security idp security-package download

IN THIS SECTION

- Syntax | 14754
- Description | 14755
- Options | 14755
- Additional Information | 14755
- Required Privilege Level | 14755
- Output Fields | 14755
- Sample Output | 14756
- Sample Output | 14756
- Sample Output | 14756
- Release Information | 14757

Syntax

```
request security idp security-package download
<check-server>
<full-update>
<policy-templates>
```

```
<version          version-number       >
<status>
```

Description

Manually download the individual components of the security package from the Juniper Security Engineering portal. The components are downloaded into a staging folder inside the device.

By default, this command tries to download the delta set attack signature table. It also downloads IDP, IPS, and application package signatures.

Options

- `check-server`—(Optional) Retrieve the version information of the latest security package from the security portal server.
- `full-update`—(Optional) Download the latest security package with the full set of attack signature tables from the portal.
- `policy-templates`—(Optional) Download the latest policy templates from the portal.
- `version` *`version-number`* —(Optional) Download the security package of a specific version from the portal.
- `status`—(Optional) Provide detailed status of security package download operation.

Additional Information

The `request security idp security-package download` command does not download security package files if the installed version on the device is same as the security package version on the server (https://signatures.juniper.net/cgi-bin/index.cgi always). The `request security idp security-package download full-update` command downloads the latest security package files on the device from the server, irrespective of the version on the device and the server.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security idp security-package download

```
user@host> request security idp security-package download
Successfully downloaded from (https://signatures.juniper.net/cgi-bin/index.cgi).
Version info:1152(Thu Apr 24 14:37:44 2008, Detector=9.1.140080400)
```

Sample Output

request security idp security-package download policy-templates

```
user@host> request security idp security-package download policy-templates
Successfully downloaded from (https://signatures.juniper.net/cgi-bin/index.cgi).
Version info:35
```

Sample Output

request security idp security-package download version 1151 full-update

```
user@host> request security idp security-package download version 1151 full-update
Successfully downloaded from (https://signatures.juniper.net/cgi-bin/index.cgi).
Version info:1151(Wed Apr 23 14:39:15 2008, Detector=9.1.140080400)
```

request security idp security-package download status

To request status for a package download:

```
user@host> request security idp security-package download status
Done;Successfully downloaded from (https://signatures.juniper.net/cgi-bin/index.cgi).
Version info:2014(Thu Oct 20 12:07:01 2011, Detector=11.6.140110920)
```

To request status for a template download:

```
user@host> request security idp security-package download status
Done; Successfully downloaded from (https://signatures.juniper.net/cgi-bin/index.cgi).
```

When devices are operating in chassis cluster mode, when you check the security package download status, a message is displayed confirming that the downloaded security package is being synchronized to the primary and secondary nodes.

```
user@host> request security idp security-package download status
node0:
--------------------------------------------------------------------------
Done;Successfully downloaded from(https://signatures.juniper.net/cgi-bin/index.cgi)
and synchronized to backup.
Version info:2011(Mon Oct 17 15:13:06 2011, Detector=11.6.140110920)
```

Release Information

Command introduced in Junos OS Release 9.2. Detailed status added in Junos OS Release 10.1. Description modified in Junos OS Release 11.1. Application package support added in Junos OS Release 11.4.

RELATED DOCUMENTATION

show security idp active-policy | 24478

show security idp security-package-version | 24615

request security idp security-package install

IN THIS SECTION

- Syntax | 14758
- Description | 14758

- Options | 14758
- Required Privilege Level | 14759
- Output Fields | 14759
- Sample Output | 14759
- Sample Output | 14759
- Release Information | 14759

Syntax

```
request security idp security-package install
<policy-templates>
<status>
<update-attack-database-only>
```

Description

Updates the attack database inside the device with the newly downloaded one from the staging folder, recompiles the existing running policy, and pushes the recompiled policy to the data plane.

Also, if there is an existing running policy, and the previously installed detector's version is different from the newly downloaded one, the downloaded components are pushed to the data plane. This command installs IDP, IPS, and application package signatures.

Options

- `policy-templates`—(Optional) Installs the policy template file into /var/db/scripts/commit/templates.
- `status`—(Optional) The command `security-package install` may take a long time depending on the new Security database size. Hence, `security-package install` command returns immediately and a background process performs the task. User can check the status using `security-package install status` command.
- `update-attack-database-only`—(Optional) Loads the security package into IDP database but does not compile/push the active policy or the new detector to the data plane.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security idp security-package install

```
user@host> request security idp security-package install
Will be processed in async mode. Check the status using the status checking CLI
```

Sample Output

request security idp security-package install status

To request status on a package installation:

```
user@host> request security idp security-package install status
Done;Attack DB update : successful - [UpdateNumber=1152,ExportDate=Thu Apr 24 14:37:44 2008]
     Updating data-plane with new attack or detector : not performed
     due to no existing active policy found.
```

To request status on a template installation:

```
user@host> request security idp security-package install status
Done; policy-template has been successfully updated into internal repository
(=>/var/db/scripts/commit/templates.xsl)!
```

Release Information

Command introduced in Junos OS Release 9.2. Description modified in Junos OS Release 11.1. Added application package support in Junos OS Release 11.4.

RELATED DOCUMENTATION

show security idp active-policy | 24478

show security idp security-package-version | 24615

request security idp security-package offline-download

IN THIS SECTION

- Syntax | 14760
- Description | 14760
- Usage | 14760
- Options | 14761
- Required Privilege Level | 14761
- Release Information | 14761

Syntax

```
request security idp security-package offline-download ( package-path package-path|status )
```

Description

Unzip the security package and copy the xml files.

Manually download the security package from the Juniper Security Engineering portal. The package will have both IDP and application package signatures. Copy the files over to the device into a certain folder and then issues the `request security idp security-package offline-download package-path` *`package-path`* command. The command will unzip the security package and copy the xml files to staging directory. Signature package installation should follow an offline-download. There is no change in installation process.

Usage

- Use this command to download and install the IDP Signature Database when the JUNOS SRX Series Firewall does not have an Internet connection.

- You can perform offline IDP update without a direct Internet connection. Download a single compressed file that can be copied to the JUNOS SRX device.

Options

- `package-path`—Package path of the zipped security package.
- `status`—Retrieve the status of offline package download operation.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 12.3X48-D10.

RELATED DOCUMENTATION

show security idp active-policy | 24478
show security idp security-package-version | 24615
request security idp security-package install | 14757

request security idp security package rollback

IN THIS SECTION

Syntax | 14762
Description | 14762
Options | 14762
Required Privilege Level | 14762
Output Fields | 14762
Sample Output | 14762
Sample Output | 14763

Release Information | 14763

Syntax

```
request security idp security-package rollback <status>
```

Description

Manually rollback to the previous version of IDP security package on SRX Series Firewalls and platforms.

Also view the status of IDP security package rollback operation on device.

Options

- `status`—(Optional) Provide status of IDP security package rollback operation.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security idp security-package rollback

```
user@host> request security idp security-package rollback
Will be processed in async mode. Check the status using the status checking CLI
```

Sample Output

request security idp security-package rollback status

```
user@host> request security idp security-package rollback status
Ready to accept a new request
```

Release Information

This command is introduced in Junos OS Release 21.4 and 22.1.

RELATED DOCUMENTATION

request security idp security-package download | 14754

request security idp ssl-inspection key add

IN THIS SECTION

- Syntax | 14764
- Description | 14764
- Options | 14764
- Required Privilege Level | 14764
- Output Fields | 14764
- Sample Output | 14764
- Sample Output | 14765
- Sample Output | 14765
- Sample Output | 14765
- Sample Output | 14765
- Release Information | 14765

Syntax

```
request security idp ssl-inspection key add <key-name> [file <file-name>] [password <password-string>] [server <server-ip>]
```

Description

Install a Privacy-Enhanced Mail (PEM) key that is optionally password protection, and associate a server with an installed key. The length of each key name and password string should not exceed 32 alphanumeric characters.

Options

- *key-name*—Name of the SSL private key.
- `file` *<file-name>*—(Optional) Location of RSA private key (PEM format) file.
- `password` *<password-string>*—(Optional) Password used to encrypt specified key.
- `server` *<server-ip>* —(Optional) Server IP address to be added to the specified key.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security idp ssl-inspection key add key1 file /var/tmp/enc1.key password encrypted

```
user@host> request security idp  ssl-inspection key add key1 file /var/tmp/enc1.key password encrypted
Added key 'key1'
```

Sample Output

request security idp ssl-inspection key add key2 file /var/tmp/enc2.key password encrypted

```
user@host> request security idp ssl-inspection key add key2 file /var/tmp/enc2.key password
encrypted
Added key 'key2', server 2.2.0.1
```

Sample Output

request security idp ssl-inspection key add key3 file /var/tmp/norm.key

```
user@host> request security idp ssl-inspection key add key3 file /var/tmp/norm.key
Added key 'key3'
```

Sample Output

request security idp ssl-inspection key add key1 server 1.1.0.1

```
user@host> request security idp ssl-inspection key add key1 server 1.1.0.1
Added key 'key1', server 1.1.0.1
```

Sample Output

request security idp ssl-inspection key add key1 server 1.1.0.2

```
user@host> request security idp ssl-inspection key add key1 server 1.1.0.2
Added key 'key1', server 1.1.0.2
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show security idp ssl-inspection key | 24618

request security idp ssl-inspection key delete

IN THIS SECTION

- Syntax | 14766
- Description | 14766
- Options | 14767
- Required Privilege Level | 14767
- Output Fields | 14767
- Sample Output | 14767
- Sample Output | 14767
- Sample Output | 14768
- Release Information | 14768

Syntax

```
request security idp ssl-inspection key delete          [<key-name>        [server
<server-ip>]]
```

Description

Delete the specified server IP from the given key if the server is specified. If the server IP is not specified, the given key will be deleted along with all the server addresses associated with it.

NOTE: You will get a delete confirmation question before deleting one or more keys or server.

Options

- *key-name*—(Optional) Name of the SSL private key.
- server *<server-ip>* —(Optional) Server IP address associated with the specified key to be deleted.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security idp ssl-inspection key delete

```
user@host> request security idp ssl-inspection key delete

This command will delete one or more ssl keys.
Continue? [yes,no] (no) yes

Number of keys 4, server 3 deleted
```

Sample Output

request security idp ssl-inspection key delete key1

```
user@host> request security idp ssl-inspection key delete key1

This command will delete one or more ssl keys.
Continue? [yes,no] (no) yes

Number of keys 1, server 2 deleted
```

Sample Output

request security idp ssl-inspection key delete key2 server 2.2.0.1

```
user@host> request security idp ssl-inspection key delete key2 server 2.2.0.1

This command will delete one or more ssl keys.
Continue? [yes,no] (no) yes

Number of keys 0, server 1 deleted
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show security idp ssl-inspection key | 24618

request security idp storage-cleanup

IN THIS SECTION

- Syntax | 14769
- Description | 14769
- Options | 14769
- Required Privilege Level | 14769
- Output Fields | 14769
- Sample Output | 14769
- Release Information | 14769

Syntax

```
request security idp storage-cleanup
```

Description

Delete unused files to free up storage space on a device.

Options

`cache-files`— Delete DFA cache files used for optimizing idp policy compilation.

`downloaded-files`— Delete downloaded security-package files (with out affecting the installed database).

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security idp storage-cleanup

```
user@host> request security idp storage-cleanup downloaded-files
Successfully deleted downloaded secdb files
```

Release Information

Command introduced in Junos OS Release 11.4.

request security ike debug-disable

IN THIS SECTION

- Syntax | 14770
- Description | 14770
- Required Privilege Level | 14770
- Output Fields | 14770
- Release Information | 14770

Syntax

```
request security ike debug-disable
```

Description

Disable IKE debugging.

Required Privilege Level

maintenance

Output Fields

This command produces no output.

Release Information

Command introduced in Release Junos OS 11.4R3.

RELATED DOCUMENTATION

- request security ike debug-enable | 14771
- show security ike debug-status | 24638

request security ike debug-enable

IN THIS SECTION

- Syntax | 14771
- Description | 14771
- Options | 14772
- Required Privilege Level | 14772
- Release Information | 14772

Syntax

```
request security ike debug-enable local local-ip-address remote remote-ip-address
```

Description

Enable IKE tracing on a single VPN tunnel specified by a local and a remote IP address. Use of this command is an alternative to configuring IKE traceoptions; you do not require any configuration to use this command. This command only traces a single tunnel, whereas configuring IKE traceoptions affects all VPN tunnels on the SRX Series Firewalls.

> **NOTE**: SRX Series Firewalls and MX-SPC3 Services Card supports this command. MX Series device with Multiservices Modular Interfaces Card (MS-MIC) or Multiservices Modular PIC Concentrator (MS-MPC) does not support this command.

To use this command:

1. Identify the local and remote IP addresses of the VPN tunnel you want to trace.
2. Enable IKE tracing on the VPN tunnel with this command.
3. Attempt tunnel establishment to capture trace information to the log file:
   - For the SRX Series Firewalls and vSRX Virtual Firewall running kmd process, the trace information is stored in `/var/log/kmd` file.

- For the MX-SPC3 Services Card, SRX Series Firewalls and vSRX Virtual Firewall running iked process (including mixed mode), the trace information is stored in `/var/log/iked` file.

If you've configured to save the trace messages into a specific file under the `[edit security ike traceoptions]` hierarchy level, the trace information is stored in the specified file name.

4. Disable per-tunnel IKE tracing with the **request security ike debug-disable** command.
5. Review the log file with the following command:
   - For the SRX Series Firewalls and vSRX Virtual Firewall running kmd process, execute the `show log kmd` or the file name specified under the `[edit security ike traceoptions]` hierarchy level.
   - For the MX-SPC3 Services Card, SRX Series Firewalls and vSRX Virtual Firewall running iked process (including mixed mode), execute the `show log iked` or the file name specified under the `[edit security ike traceoptions]` hierarchy level.

You can use the **show security ike debug-status** command:

- to view the status of the per-tunnel IKE tracing operation.
- to view the status of the interchassis link tunnel only.

Options

- `local` *`local-ip-address`*—The address of the local VPN peer.
- `remote` *`remote-ip-address`*—The address of the remote VPN peer.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 11.4R3.

RELATED DOCUMENTATION

request security ike debug-disable | **14770**

show security ike debug-status | **24638**

request security integrity

IN THIS SECTION

- Syntax | 14773
- Description | 14773
- Options | 14773
- Required Privilege Level | 14773
- Release Information | 14773

Syntax

```
request security integrity (appraise | measure)
```

Description

Verify the validity of keys added to the system IMA key chain.

Options

appraise (file | key) Verify the IMA signature on a key file or key.

measure (file | key) Compute and add an IMA signature to a key file or a key.

Required Privilege Level

maintenance

Release Information

Statement introduced in Junos OS Evolved Release 20.1R1.

request security internal-security-association refresh

IN THIS SECTION

- Syntax | 14774
- Description | 14774
- Required Privilege Level | 14774
- Output Fields | 14774
- Sample Output | 14775
- Release Information | 14776

Syntax

```
request security internal-security-association refresh node <node-id | all | local | primary>
```

Description

Activate internal IPsec so an attacker cannot gain unauthorized information.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request. Table 119 on page 14774 shows the list of output fields.

Table 119: request security internal-security-association refresh Output Fields

Field Name	Field Description
RE	Specify key for the routing engine.

Table 119: request security internal-security-association refresh Output Fields *(Continued)*

Field Name	Field Description
direction	Direction of the manual security association. The direction must be `bidirectional`.

Sample Output

request security internal-security-association refresh

with `security ipsec internal security-association` enabled

```
user@host> request security internal-security-association refresh
node0:
--------------------------------------------------------------------------
RE: set up internal sa

node1:
--------------------------------------------------------------------------
RE: set up internal sa
```

with `security ipsec internal security-association` disabled

```
user@host> request security internal-security-association refresh
node0:
--------------------------------------------------------------------------
RE: no key
direction 0
delete __juniper_internal_ipsec_sa__, direction 0
direction 1
delete __juniper_internal_ipsec_sa__, direction 1

node1:
--------------------------------------------------------------------------
RE: no key
direction 0
delete __juniper_internal_ipsec_sa__, direction 0
```

```
direction 1
delete __juniper_internal_ipsec_sa__, direction 1
```

Release Information

Command introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

- show security internal-security-association | 24686
- internal (Security IPsec) | 9190

request security key-pair

IN THIS SECTION

- Syntax | 14776
- Description | 14777
- Options | 14777
- Required Privilege Level | 14777
- Output Fields | 14777
- Sample Output | 14777
- Release Information | 14777

Syntax

```
request security key-pair filename
<size key-size>
<type (rsa | dsa)>
```

Description

(Encryption interface on M Series and T Series routers and EX Series switches only) Generate a public and private key pair for a digital certificate.

> **NOTE**: The `request security-certificates` command is deprecated and are not available with Junos in FIPS mode because security certificates are not compliant with the NIST SP 800-131A standard.

Options

filename Name of a file in which to store the key pair.

size *key-size* (Optional) Key size, in bits. The key size can be **512**, **1024**, or **2048**. The default value is **1024**.

type (Optional) Algorithm used to encrypt the key:

- **rsa**—RSA algorithm. This is the default.
- **dsa**—Digital signature algorithm with Secure Hash Algorithm (SHA).

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security key-pair

```
user@host> request security key-pair security-key-file
```

Release Information

Command introduced before Junos OS Release 7.4.

request security pki ca-certificate ca-profile-group default-trusted-ca-certs

IN THIS SECTION

- Syntax | 14778
- Description | 14778
- Options | 14779
- Required Privilege Level | 14779
- Output Fields | 14779
- Sample Output | 14779
- Sample Output | 14780
- Sample Output | 14780
- Release Information | 14781

Syntax

```
request security pki ca-certificate ca-profile-group default-trusted-ca-certs download [check-server | no-forwarding | status]
```

Description

When you setup dynamic update of trusted CA bundle, you use this command to -

- Explicitly instruct the SRX Series Firewall or a Junos OS device to manually download the default trusted CA certificates from CDN server. It checks the available version on the CDN server and downloads it.
- Check connectivity to the CDN server for the Junos OS device.
- Monitor the status of default trusted CA certificates downloaded from CDN server.

Options

download Download default trusted CA certificates from a CDN server.

Use this option when you need to explicitly download default trusted CA certificates in addition to periodic download.

check-server Check connectivity to CDN server to download default trusted CA certificates.

This command downloads the manifest file and displays the trusted-ca-bundle version available in CDN server

no-forwarding No forwarding.

status Check current status of default trusted CA certificates downloaded from CDN server.

This option displays the default trusted CA certificates version number and version date.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki ca-certificate ca-profile-group default-trusted-ca-certs download

- Success scenario

```
user@host> request security pki ca-certificate ca-profile-group default-trusted-ca-certs
download
Connection to CDN server is successful. Default trusted CA certs bundle version available is
   <version> dated <date>
   Downloading the trusted CA certs bundle...
   Download successful
   Updating trusted CA certs to default CA profile group <ca-profile-group>
   Added <num> CAs and removed <num> CAs from default trusted CA profile group test
```

- Failure scenario - 1

```
user@host> request security pki ca-certificate ca-profile-group default-trusted-ca-certs
download
Connection to CDN server is unsuccessful
```

- Failure scenario - 2

```
user@host> request security pki ca-certificate ca-profile-group default-trusted-ca-certs
download
Connection to CDN server is successful. Default trusted CA certs bundle
   version available is <version_no> dated <date>.
   Downloading the trusted CA certs...
   Download Failed.
```

Sample Output

request security pki ca-certificate ca-profile-group default-trusted-ca-certs download check-server

```
user@host> request security pki ca-certificate ca-profile-group default-trusted-ca-certs
download check-server
Connection to CDN server is successful. Default trusted CA certs bundle version available is
<version_no> dated <date>.
```

Sample Output

request security pki ca-certificate ca-profile-group default-trusted-ca-certs download status

- Success scenario

```
user@host> request security pki ca-certificate ca-profile-group default-trusted-ca-certs
download status
Default trusted CA certs bundle available on device is <version_no> dated <date>.
```

- Failure scenario

```
user@host> request security pki ca-certificate ca-profile-group default-trusted-ca-certs
download status
Default trusted CA certs bundle is not available.
```

Release Information

Statement introduced in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

- default-trusted-ca-certs (Security)
- Dynamic Update of Trusted CA Bundle

request security pki ca-certificate ca-profile-group load

IN THIS SECTION

- Syntax | 14782
- Description | 14782
- Options | 14782
- Required Privilege Level | 14782
- Output Fields | 14783
- Sample Output | 14783
- Sample Output | 14783
- Release Information | 14784

Syntax

```
request security pki ca-certificate ca-profile-group load ca-group-name ca-group-name filename
[path/filename | default]
```

Description

For SSL forward proxy, you need to load trusted CA certificates on your system. By default, Junos OS provides a list of trusted CA certificates that include default certificates used by common browsers. Alternatively, you can define your own list of trusted CA certificates and import them on to your system.

Use this command to load the default certificates or to specify a path and filename of trusted CA certificates that you define.

The `default` option is not supported on MX204, MX240, MX304, MX 480, MX960, MX2008, MX2010, MX 2020, MX10004, MX10008, PTX10003-80C, PTX10003-160C, and PTX10008 routers.

Starting in Junos OS Release 21.4R1, you can get the status of CA certificates configured under default CA profile group by executing "request security pki ca-profile-group-status" on page 14795 command . With "request security pki ca-profile-group-status" on page 14795 command, you can verify the number of CA certificates loaded and number of CA certificates missing within a CA profile group.

Starting in Junos OS Release 23.2R1, when you configure dynamic update of trusted CA bundle using the statement default-trusted-ca-certs (Security), the process of loading the default trusted CA certificates happens in the background. During this process, PKID response might slowdown for few minutes.

Options

ca-group-name ***ca-group-name***	Load the specified CA group profile.
filename ***path/filename***	Directory location and filename of the trusted CA certificates defined by you.
filename default	Load the trusted CA certificates available by default.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki ca-certificate ca-profile-group load (default)

```
user@host> request security pki ca-certificate ca-profile-group load ca-group-name ca-default
filename default

Loading of certs started
Loading of <no-of-certs> trusted CA certs started in the background. PKID response might be slow
for next several minutes.
```

Sample Output

request security pki ca-certificate ca-profile-group load (path/filename)

```
user@host> request security pki ca-certificate ca-profile-group load ca-group-name ca-manual
filename /var/tmp/firefox-all.pem

Do you want to load this CA certificate ? [yes,no] (no) yes

Loading 196 certificates for group 'ca-manual'.
ca-manual_1_sysgen: Loading done.
ca-manual_2_sysgen: Loading done.
ca-manual_3_sysgen: Loading done.
ca-manual_4_sysgen: Loading done.
ca-manual_5_sysgen: Loading done.
ca-manual_6_sysgen: Loading done.

...
ca-manual_195_sysgen: Loading done.
ca-manual_196_sysgen: Loading done.
ca-profile-group 'ca-manual’ successfully loaded. Success[193] Skipped[3]
```

Release Information

Command introduced in Junos OS Release 12.1; `default` option added in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

- show security pki ca-certificate | **24957**
- *Basic Elements of PKI in Junos OS*
- request security pki ca-profile-group-status | **14795**

request security pki ca-certificate enroll

IN THIS SECTION

- Syntax | **14784**
- Description | **14784**
- Options | **14785**
- Required Privilege Level | **14785**
- Output Fields | **14785**
- Sample Output | **14785**
- Release Information | **14785**

Syntax

```
request security pki ca-certificate enroll ca-profile ca-profile-name
```

Description

Request a digital certificate from a certificate authority (CA) online by using the Simple Certificate Enrollment Protocol (SCEP).

Options

ca-profile ***ca-profile-name***	CA profile name.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki ca-certificate enroll

```
user@host> request security pki ca-certificate enroll ca-profile entrust
Received following certificates:
  Certificate: C=us, O=juniper, CN=First Officer
    Fingerprint: 46:71:15:34:f0:a6:41:76:65:81:33:4f:68:47:c4:df:78:b8:e3:3f
  Certificate: C=us, O=juniper, CN=First Officer
    Fingerprint: bc:78:87:9b:a7:91:13:20:71:db:ac:b5:56:71:42:ad:1a:b6:46:17
  Certificate: C=us, O=juniper
    Fingerprint: 00:8e:6f:58:dd:68:bf:25:0a:e3:f9:17:70:d6:61:f3:53:a7:79:10
Do you want to load the above CA certificate ? [yes,no] (no) yes
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

clear security pki ca-certificate | 4017

show security pki ca-certificate | 24957

request security pki ca-certificate enroll (Security)

IN THIS SECTION

- Syntax | 14786
- Description | 14786
- Options | 14786
- Required Privilege Level | 14786
- Output Fields | 14786
- Sample Output | 14787
- Release Information | 14787

Syntax

```
request security pki ca-certificate enroll ca-profile ca-profile-name
```

Description

Request a digital certificate from a certificate authority (CA) online by using the Simple Certificate Enrollment Protocol (SCEP).

Options

ca-profile ***ca-profile-name***	CA profile name.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki ca-certificate enroll

```
user@host> request security pki ca-certificate enroll ca-profile entrust
Received following certificates:
  Certificate: C=us, O=example, CN=First Officer
    Fingerprint: 46:71:15:34:f0:a6:41:76:65:81:33:4f:68:47:c4:df:78:b8:e3:3f
  Certificate: C=us, O=example, CN=First Officer
    Fingerprint: bc:78:87:9b:a7:91:13:20:71:db:ac:b5:56:71:42:ad:1a:b6:46:17
  Certificate: C=us, O=example
    Fingerprint: 00:8e:6f:58:dd:68:bf:25:0a:e3:f9:17:70:d6:61:f3:53:a7:79:10
Do you want to load the above CA certificate ? [yes,no] (no) yes
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

show security pki ca-certificate (View) | 24962

Basic Elements of PKI in Junos OS

request security pki ca-certificate load

IN THIS SECTION

- Syntax | 14788
- Description | 14788
- Options | 14788
- Required Privilege Level | 14788
- Output Fields | 14788
- Sample Output | 14788

Release Information | 14789

Syntax

```
request security pki ca-certificate load ca-profile ca-profile-name  filename path/
filename
```

Description

Manually load a certificate authority (CA) digital certificate from a specified location.

Options

ca-profile ***ca-profile-name***	Load the specified CA profile.
filename ***path/filename***	Directory location and filename of the CA digital certificate.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki ca-certificate load

```
user@host> request security pki ca-certificate load ca-profile ca-private filename pki-
file
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

clear security pki ca-certificate | 4017

show security pki ca-certificate | 24957

request security pki ca-certificate load (Security)

IN THIS SECTION

- Syntax | 14789
- Description | 14789
- Options | 14790
- Required Privilege Level | 14790
- Output Fields | 14790
- Sample Output | 14790
- Release Information | 14790

Syntax

```
request security pki ca-certificate load ca-profile ca-profile-name  filename path/
filename
```

Description

Manually load a certificate authority (CA) digital certificate from a specified location.

Options

ca-profile *ca-profile-name*	Load the specified CA profile.
filename *path/filename*	Directory location and filename of the CA digital certificate.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki ca-certificate load

```
user@host> request security pki ca-certificate load ca-profile 2Kkey filename /var/tmp/
2Kkey.pem

Fingerprint:
  a0:08:bb:1f:75:96:76:cd:ee:db:36:10:b6:c6:d8:df:5e:02:05:05 (sha1)
  f5:58:6b:de:7c:d6:cd:90:5a:18:c3:0e:3d:95:da:25 (md5)
Do you want to load this CA certificate ? [yes,no] (no) yes

CA certificate for profile 2Kkey loaded successfully
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

show security pki ca-certificate | 24957

Basic Elements of PKI in Junos OS

request security pki ca-certificate verify

IN THIS SECTION

- Syntax | 14791
- Description | 14791
- Options | 14791
- Required Privilege Level | 14791
- Output Fields | 14791
- Sample Output | 14792
- Release Information | 14792

Syntax

```
request security pki ca-certificate verify ca-profile ca-profile-name
```

Description

Verify the digital certificate installed for the specified certificate authority (CA).

Options

ca-profile *ca-profile-name*	Name of the local digital certificate identifier.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

command-name

You receive the following response before the certificate revocation list (CRL) is downloaded:

```
request security pki ca-certificate verify ca-profile ca1 (CRL not downloaded)
user@host> request security pki ca-certificate verify ca-profile ca1

CA certificate ca1: CRL verification in progress. Please check the PKId debug logs for
completion status
```

Release Information

Command introduced in Junos OS Release 8.5.

request security pki ca-certificate verify (Security)

IN THIS SECTION

- Syntax | 14792
- Description | 14793
- Options | 14793
- Required Privilege Level | 14793
- Output Fields | 14793
- Sample Output | 14793
- Release Information | 14794

Syntax

```
request security pki ca-certificate verify ca-profile ca-profile-name
```

Description

Verify the digital certificate installed for the specified certificate authority (CA).

Options

`ca-profile` *`ca-profile-name`* —Display the specified CA profile.

Required Privilege Level

maintenance and security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki ca-certificate verify ca-profile ca1 (CRL downloaded)

This user has downloaded the certificate revocation list (CRL).

```
user@host> request security pki ca-certificate verify ca-profile ca1
CA certificate ca1 verified successfully
```

request security pki ca-certificate verify ca-profile ca1 (CRL not downloaded)

This user has not downloaded the certificate revocation list (CRL).

```
user@host> request security pki ca-certificate verify ca-profile ca1
CA certificate ca1: CRL verification in progress. Please check the PKId debug logs for
completion status
```

request security pki ca-certificate verify ca-profile Root-CA (Verify enrolled CA certificate validity status on MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall)

You receive the following response when the CA certificate verification is failed. In this sample, the CA certificate verification is failed due to invalid CA certificate:

```
user@host> request security pki ca-certificate verify ca-profile Root-CA
CA certificate Root-CA verification failed. CA cert is not valid untill <05-19-2021 08:05>
```

request security pki ca-certificate verify ca-profile Root-CA (Verify enrolled CA certificate present in MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall)

You receive the following response when the CA certificate is missing:

```
user@host> request security pki ca-certificate verify ca-profile Root-CA
CA cert Root-CA Verification Failed. CA cert is missing
```

request security pki ca-certificate verify ca-profile CSO_37 (Verify local certificate status when the CA is unreachable for MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall)

You receive the following response when a CA is not reachable or CRL download has failed:

```
user@host> request security pki ca-certificate verify ca-profile CSO_37
CA certificate CSO_37 Verification Failed. Unreachable CA or CRL Download Failed
```

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

- ca-profile (Security PKI) | **3157**
- show security pki ca-certificate (View) | **24962**
- request security pki ca-profile-group-status (MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall) | **14796**

request security pki ca-profile-group-status

IN THIS SECTION

- Syntax | 14795
- Description | 14795
- Options | 14795
- Required Privilege Level | 14795
- Output Fields | 14796
- Sample Output | 14796
- Release Information | 14796

Syntax

```
request security pki ca-profile-group-status ca-group-name [ca-group-name|default]
```

Description

Get the status of CA certificates configured under default CA profile group. With this command, you can verify the number of CA certificates loaded and number of CA certificates missing within a CA profile group.

Options

ca-group-name *ca-group-name*	Load the specified CA group profile.
ca-group-name default	Load the trusted CA group profile available by default.

Required Privilege Level

maintenance

Output Fields

Table 120 on page 14796 lists the output fields for the `request security pki ca-profile-group-status` command.

Table 120: request security pki ca-profile-group-status Output Fields

Field Name	Field Description
No-of-Ca-Certs-Loaded	Total number of certificates successfully loaded for the specified CA profile group.
No-of-Ca-Certs-Load-Failure	Total number of certificates failed to load for the specified CA profile group.
Missing-Cert-id	Missing certificate Ids for the specified CA profile group.
Total-Certs	Total number of certificates imported for the specified CA profile group.

Sample Output

request security pki ca-profile-group-status (MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall)

```
user@host> request security pki ca-profile-group-status ca-group-name DEFAULT_CSO
No-of-Ca-Certs-Loaded:135 , No-of-Ca-Certs-Load-Failure: 0 , Missing-Cert-id: None ,Total-Certs:
135
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

show security ipsec statistics (MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall) | **25023**

request security pki crl load

IN THIS SECTION

- Syntax | 14797
- Description | 14797
- Options | 14797
- Required Privilege Level | 14797
- Output Fields | 14797
- Sample Output | 14798
- Release Information | 14798

Syntax

```
request security pki crl load ca-profile ca-profile-name   filename path/filename
```

Description

Manually install a certificate revocation list (CRL) on the router from a specified location.

Options

ca-profile ***ca-profile-name***	Load the specified certificate authority (CA) profile.
filename ***path/filename***	Directory location and filename of the CRL.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki crl load

```
user@host> request security pki crl load ca-profile ca-private filename pki-file
```

Release Information

Command introduced in Junos OS Release 8.1.

request security pki crl load (Security)

IN THIS SECTION

- Syntax | 14798
- Description | 14798
- Options | 14799
- Required Privilege Level | 14799
- Output Fields | 14799
- Sample Output | 14799
- Release Information | 14799

Syntax

```
request security pki crl load ca-profile ca-profile-name   filename path/filename
```

Description

Manually install a certificate revocation list (CRL) on the device from a specified location.

Options

ca-profile ***ca-profile-name***	Load the specified certificate authority (CA) profile.
filename ***path/filename***	Directory location and filename of the CRL.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki crl load

```
user@host> request security pki crl load ca-profile ca-test filename example-inter-
ca.crl
CRL for CA profile ca-test loaded successfully
```

Release Information

Command introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Basic Elements of PKI in Junos OS

request security pki generate-certificate-request

IN THIS SECTION

- Syntax | 14800
- Description | 14800
- Options | 14800
- Required Privilege Level | 14801
- Output Fields | 14801
- Sample Output | 14801
- Release Information | 14802

Syntax

```
request security pki generate-certificate-request certificate-id certificate-id-name  domain-
name domain-name  subject subject-distinguished-name
<email email-address>
<filename (path | terminal)>
<ip-address ip-address>
```

Description

Manually generate a local digital certificate request in the Public-Key Cryptography Standards #10 (PKCS-10) format.

Options

certificate-id ***certificate-id-name***	Name of the local digital certificate and the public/private key pair.
domain-name ***domain-name***	Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.
subject ***subject-distinguished-name***	Distinguished name format that contains the common name, department, company name, state, and country:

- CN—Common name
- OU—Organizational unit name
- O—Organization name
- ST—State
- C—Country

email *email-address* — (Optional) E-mail address of the certificate holder.

filename (*path* | terminal) — (Optional) Location where the local digital certificate request should be placed or the login terminal.

ip-address *ip-address* — (Optional) IP address of the router.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki generate-certificate-request

```
user@host> request security pki generate-certificate-request certificate-id local-entrust2
domain-name router2.example.net filename entrust-req2 subject cn=router2.example.net

Generated certificate request
-----BEGIN CERTIFICATE REQUEST-----
MIIBoTCCAQoCAQAwGjEYMBYGA1UEAxMPdHAxLmp1bmlwZXIubmV0MIGfMA0GCSqG
SIb3DQEBAQUAA4GNADCBiQKBgQCiUFklQws1Ud+AqN5DDxRs2kVyKEhh9qoVFnz+
Hz4c9vsy3B8ElwTJlkmIt2cB3yifB6zePd+6WYpf57Crwre7YqPkiXM31F6z3YjX
H+1BPNbCxNWYvyrnSyVYDbFj8o0Xyqog8ACDfVL2JBWrPNBYy7imq/K9soDBbAs6
5hZqqwIDAQABoEcwRQYJKoZIhvcNAQkOMTgwNjAOBgNVHQ8BAf8EBAMCB4AwJAYD
VR0RAQH/BBowGIIWdHAxLmVuZ2xhYi5qdW5pcGVyLm5ldDANBgkqhkiG9w0BAQQF
AAOBgQBc2rq1v5SOQXH7LCb/FdqAL8ZM6GoaN5d6cGwq4bB6a7UQFgtoH406gQ3G
3iH0Zfz4xMIBpJYuGd1dkqgvcDoH3AgTsLkfn7Wi3x5H2qeQVs9bvL4P5nvEZLND
```

```
EIMUHwteolZCiZ70fO9Fer9cXWHSQs1UtXtgPqQJy2xIeImLgw==
-----END CERTIFICATE REQUEST-----
Fingerprint:
0d:90:b8:d2:56:74:fc:84:59:62:b9:78:71:9c:e4:9c:54:ba:16:97 (sha1)
1b:08:d4:f7:90:f1:c4:39:08:c9:de:76:00:86:62:b8 (md5)
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

clear security pki certificate-request | 4018

show security pki certificate-request | 24968

request security pki generate-certificate-request (Security)

IN THIS SECTION

- Syntax | 14802
- Description | 14803
- Options | 14803
- Required Privilege Level | 14804
- Output Fields | 14804
- Sample Output | 14804
- Release Information | 14805

Syntax

```
request security pki generate-certificate-request certificate-id certificate-id-name domain-name
domain-name   subject subject-distinguished-name
<add-ca-constraint>
<digest (sha1 | sha256)>
```

```
<email email-address>
<filename (path | terminal)>
<ip-address ip-address>
```

Description

Manually generate a local digital certificate request in the Public-Key Cryptography Standards #10 (PKCS-10) format.

Options

certificate-id *certificate-id-name*
Name of the local digital certificate and the public/private key pair.

domain-name *domain-name*
Fully qualified domain name (FQDN) provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.

subject *subject-distinguished-name*
Distinguished name format contains the following information:

- `DC`—Domain component
- `CN`—Common name
- `OU`—Organizational unit name
- `O`—Organization name
- `L`—Locality
- `ST`—State
- `C`—Country

digest
(Optional) Hash algorithm used to sign the certificate request.

- `sha1`—SHA-1 digests (default value for RSA or DSA only).
- `sha256`—SHA-256 digests for RSA or ECDSA only (default value for ECDSA).
- `sha-384`—SHA-384 digests for ECDSA only.

Starting in Junos OS Release 18.1R3, the default encryption algorithm that is used for validating automatically and manually generated self-signed PKI certificates is Secure Hash Algorithm 256 (SHA-256). Prior to Junos OS Release 18.1R3, SHA-1 is used as default encryption algorithm.

email ***email-address***	(Optional) E-mail address of the certificate holder.
filename (*path* | terminal)	(Optional) Location where the local digital certificate request should be placed or the login terminal.
ip-address ***ip-address***	(Optional) IP address of the router.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki generate-certificate-request

```
user@host> request security pki generate-certificate-request certificate-id local-entrust2
domain-name router2.example.net filename entrust-req2 subject cn=router2.example.net

Generated certificate request
-----BEGIN CERTIFICATE REQUEST-----
MIIBoTCCAQoCAQAwGjEYMBYGA1UEAxMPdHAxLmp1bmlwZXIubmV0MIGfMA0GCSqG
SIb3DQEBAQUAA4GNADCBiQKBgQCiUFklQws1Ud+AqN5DDxRs2kVyKEhh9qoVFnz+
Hz4c9vsy3B8ElwTJlkmIt2cB3yifB6zePd+6WYpf57Crwre7YqPkiXM31F6z3YjX
H+1BPNbCxNWYvyrnSyVYDbFj8o0Xyqog8ACDfVL2JBWrPNBYy7imq/K9soDBbAs6
5hZqqwIDAQABoEcwRQYJKoZIhvcNAQkOMTgwNjAOBgNVHQ8BAf8EBAMCB4AwJAYD
VR0RAQH/BBowGIIWdHAxLmVuZ2xhYi5qdW5pcGVyLm5ldDANBgkqhkiG9w0BAQQF
AAOBgQBc2rq1v5SOQXH7LCb/FdqAL8ZM6GoaN5d6cGwq4bB6a7UQFgtoH406gQ3G
3iH0Zfz4xMIBpJYuGd1dkqgvcDoH3AgTsLkfn7Wi3x5H2qeQVs9bvL4P5nvEZLND
EIMUHwteolZCiZ70fO9Fer9cXWHSQs1UtXtgPqQJy2xIeImLgw==
-----END CERTIFICATE REQUEST-----
Fingerprint:
0d:90:b8:d2:56:74:fc:84:59:62:b9:78:71:9c:e4:9c:54:ba:16:97 (sha1)
1b:08:d4:f7:90:f1:c4:39:08:c9:de:76:00:86:62:b8 (md5)
```

Release Information

Command introduced in Junos OS Release 7.5. Support for `digest` option added in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

show security pki certificate-request (View) | 24972

request security pki generate-key-pair

IN THIS SECTION

- Syntax | 14805
- Description | 14805
- Options | 14806
- Required Privilege Level | 14806
- Output Fields | 14806
- Sample Output | 14806
- Release Information | 14806

Syntax

```
request security pki generate-key-pair certificate-id certificate-id-name
<size (512 | 1024 | 2048)>
```

Description

Generate a Public Key Infrastructure (PKI) public and private key pair for a local digital certificate.

Options

certificate-id *certificate-id-name*	Name of the local digital certificate and the public/private key pair.
size	(Optional) Key pair size. The key pair size can be **512**, **1024**, or **2048** bits.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki generate-key-pair

```
user@host> request security pki generate-key-pair certificate-id billy size 2048
Generated key pair billy, key size 2048 bits
```

Release Information

Command introduced in Junos OS Release 7.5.

request security pki generate-key-pair

IN THIS SECTION

- Syntax | 14807
- Description | 14807
- Options | 14807
- Required Privilege Level | 14807
- Output Fields | 14807

Sample Output | 14808
Release Information | 14808

Syntax

```
request security pki generate-key-pair certificate-id certificate-id-name
<size (512 | 1024 | 2048)>
<type (dsa | rsa)>
```

Description

Generate a public key infrastructure (PKI) public/private key pair for a local digital certificate.

Options

certificate-id *certificate-id-name*	Name of the local digital certificate and the public/private key pair.
size	(Optional) Key pair size. The key pair size can be **512**, **1024**, or **2048** bits. If a key pair size is not specified, the default value, **1024** bits, is applied.
type	(Optional) The algorithm to be used for encrypting the public/private key pair. The encryption algorithm can be **dsa** or **rsa** . If an encryption algorithm is not specified, the default value, **rsa**, is applied.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki generate-key-pair

```
user@switch> request security pki generate-key-pair certificate-id billy size 2048
Generated key pair billy, key size 2048 bits
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Manually Generating Self-Signed Certificates on Switches (CLI Procedure)

request security pki generate-key-pair (Security)

IN THIS SECTION

- Syntax | 14809
- Description | 14809
- Options | 14809
- Required Privilege Level | 14809
- Output Fields | 14810
- Sample Output | 14810
- Release Information | 14810

Syntax

```
request security pki generate-key-pair certificate-id certificate-id-name
<size (256 | 384 | 1024 | 2048 | 4096 | 521)>
<type (dsa | ecdsa | rsa)>
```

Description

Generate a public key infrastructure (PKI) public/private key pair for a local digital certificate.

Options

certificate-id *certificate-id-name*

Name of the local digital certificate and the public/private key pair.

size

Key pair size. The key pair size can be 256, 384, 521, 1024, 2048, or 4096 bits. Key pair sizes of 256, 384, and 521 bits are compatible with ECDSA. For Digital Signal Algorithm (DSA) and Rivest Shamir Adleman (RSA), algorithms the size must be 1024, 2048, or 4096. The default key pair size is 1024 for DSA and 2048 for RSA.

The following are supported when ECDSA-521 signatures are used:

- Load a complete certificate, which is generated using an external tool like OpenSSL into PKI.
- Manually generate a Certificate Signing Request (CSR) for a local certificate and sending the CSR to a (Certificate Authority) CA server to enroll.
- Automatic enroll with CA server.

type

The algorithm to be used for encrypting the public/private key pair:

- `ecdsa`—ECDSA encryption
- `dsa`— DSA encryption
- `rsa`—RSA encryption (default)

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki generate-key-pair

```
user@host> request security pki generate-key-pair type [xxx] size [xxx] certificate-id
test
Generated key pair test, key size [xxx] bits
```

Release Information

Command introduced in Junos OS Release 11.1.

Options to support Elliptic Curve Digital Signature Algorithm (ECDSA) added in Junos OS Release 12.1X45-D10.

521 option to support ECDSA introduced in Junos OS Release 19.1R1 on SRX5000 line of devices with SRX5K-SPC3 card.

RELATED DOCUMENTATION

Basic Elements of PKI in Junos OS

request security pki key-pair export

IN THIS SECTION

- Syntax | 14811
- Description | 14811
- Options | 14811
- Required Privilege Level | 14811

- Output Fields | **14811**
- Release Information | **14812**

Syntax

```
request security pki key-pair export certificate-id certificate-id filename filename
<passphrase string>
< type (der | pem)>
```

Description

Export the keypair for an end-entity (EE) certificate. The exported keypair is encrypted and can be imported along with the EE certificate. Using the CLI `request security pki key-pair export` command, you can export the pki key-pairs file as a backup or to check the file for troubleshooting purposes. We recommend denying access to the CLI `request security pki key-pair export` command to all users and restrict this command only to the privileged users.

Options

certificate-id *certificate-id*	Name of the local digital certificate.
filename *filename*	Target directory location and filename of the CA digital certificate.
passphrase *passphrase*	(Optional) Passphrase to protect the keypair data for PEM format. The passphrase can be up to 64 characters. If specified, the passphrase must be used when importing the keypair.
type (der | pem)	(Optional) Type of format, either DER or PEM. PEM is the default.

Required Privilege Level

maintenance

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

request security pki local-certificate export | 14822

request security pki local-certificate enroll

IN THIS SECTION

- Syntax | 14812
- Description | 14812
- Options | 14813
- Required Privilege Level | 14813
- Output Fields | 14813
- Sample Output | 14814
- Release Information | 14814

Syntax

```
request security pki local-certificate enroll ca-profile ca-profile-name  certificate-
id certificate-id-name  challenge-password password  domain-name domain-name  subject subject-
distinguished-name
<email email-address>
<ip-address ip-address>
```

Description

Request that a certificate authority (CA) enroll and install a local digital certificate online by using the Simple Certificate Enrollment Protocol (SCEP).

Options

ca-profile *ca-profile-name*	CA profile name.
certificate-id *certificate-id-name*	Name of the local digital certificate and the public/private key pair.
challenge-password *password*	Password set by the administrator and normally obtained from the SCEP enrollment webpage of the CA. The password is 16 characters in length.
domain-name *domain-name*	Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.
subject *subject-distinguished-name*	Distinguished name format that contains the common name, department, company name, state, and country: • `CN`—Common name • `OU`—Organizational unit name • `O`—Organization name • `ST`—State • `C`—Country
email *email-address*	(Optional) E-mail address of the certificate holder.
ip-address *ip-address*	(Optional) IP address of the router.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

command-name

```
user@host> request security pki local-certificate enroll certificate-id r3-entrust-scep ca-
profile entrust domain-name router3.example.net subject
"CN=router3,OU=Engineering,O=juniper,C=US" challenge-password 123

Certificate enrollment has started. To view the status of your enrollment, check the public key
infrastructure log (pkid) log file at /var/log/pkid. Please save the challenge-password for
revoking this certificate in future.  Note that this password is not stored on the router.
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

show security pki local-certificate | 24982

request security pki local-certificate enroll cmpv2

IN THIS SECTION

- Syntax | 14815
- Description | 14815
- Options | 14815
- Required Privilege Level | 14816
- Output Fields | 14816
- Sample Output | 14816
- Release Information | 14817

Syntax

```
request security pki local-certificate enroll cmpv2
    ca-dn  subject-dn
    ca-profile  ca-profile name
    ca-reference  reference
    ca-secret  shared-secret
    certificate-id certificate-id-name
    domain-name domain-name
    email email-address
     ip-address ip-address
    ipv6-address ipv6-address
    subject subject-distinguished-name
```

Description

Enroll and install a local digital certificate online by using CMPv2. This command loads both end-entity (EE) and CA certificates based on the CA server configuration. Certificate revocation list (CRL) or Online Certificate Status Protocol (OCSP) can be used to check the revocation status of a certificate.

Options

ca-dn *subject-dn*	The distinguished name (DN) of the CA enrolling the EE certificate must be specified during enrollment. This optional parameter is mandatory if the CA certificate is not already enrolled. If the CA certificate is already enrolled, the subject DN is extracted from the CA certificate.
ca-profile *ca-profile-name*	CA profile name.
ca-reference *reference*	Out-of-band reference value received from the CA server.
ca-secret *shared-secret*	Out-of-band secret value received from the CA server.
certificate-id *certificate-id-name*	Name of the local digital certificate and the public/private key pair.
domain-name *domain-name*	Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.
email *email-address*	E-mail address of the certificate holder.

ip-address *ip-address*

IP address of the router.

ipv6-address *ipv6-address*

IPv6 address of the router for the alternate subject.

subject *subject-distinguished-name*

Distinguished Name (DN) format that contains the domain component, common name, department, serial number, company name, state, and country in the following format: DC, CN, OU, O, SN, L, ST, C.

- `DC`—Domain component
- `CN`—Common name
- `OU`—Organizational unit name
- `O`—Organization name
- `SN`—Serial number of the device

  If you define SN in the subject field without the serial number, then the serial number is read directly from the device and added to the certificate signing request (CSR).
- `ST`—State
- `C`—Country

Required Privilege Level

maintenance and security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

command-name

```
user@host> request security pki local-certificate enroll cmpv2 ca-profile root-552 ca-dn
DC=example,CN=root-552 certificate-id tc552 email tc552-root@example.net domain-name example.net
ip-address 192.0.2.22 ca-secret example ca-reference 51892 subject CN=example,OU=SBU,O=552-22
```

```
Certificate enrollment has started. To view the status of your enrollment, check the public key
infrastructure log (pkid) log file at /var/log/pkid.
```

Release Information

Command introduced in Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

- show security pki local-certificate (View) | 24992
- clear security pki local-certificate (Device) | 4027

request security pki local-certificate enroll scep

IN THIS SECTION

- Syntax | 14817
- Release Information | 14818
- Description | 14818
- Options | 14818
- Required Privilege Level | 14820
- Output Fields | 14820
- Sample Output | 14820
- Sample Output | 14821

Syntax

```
request security pki local-certificate enroll scep
    ca-profile  ca-profile name
    certificate-id certificate-id-name
    challenge-password challenge-password
    digest (sha-1 | sha-256)
```

```
domain-name domain-name
email email-address
ip-address ip-address
ipv6-address ipv6-address
logical-system (logical-system-name | all)
scep-digest-algorithm (md5 | sha-1)
scep-encryption-algorithm (des | des3)
subject subject-distinguished-name
```

Release Information

Command introduced in Junos OS Release 9.1. Serial number (SN) option added to the subject string output field in Junos OS Release 12.1X45. `scep` keyword and `ipv6-address` option added in Junos OS Release 15.1X49-D40.

Starting in Junos OS Release 20.1R1 on vSRX Virtual Firewall 3.0, you can safeguard the private keys used by PKID and IKED using Microsoft Azure Key Vault hardware security module (HSM) service. You can establish a PKI based VPN tunnel using the keypairs generated at the HSM. The hub `certificate-id` option under certificate-id is not available for configuration after generating HSM key-pair.

Starting in Junos OS Release 20.4R1 on vSRX Virtual Firewall 3.0, you can safeguard the private keys used by PKID and IKED using AWS Key Management Service (KMS). You can establish a PKI based VPN tunnel using the keypairs generated by the KMS. The hub `certificate-id` option under certificate-id is not available for configuration after generating PKI key-pair.

Starting in Junos OS Release 22.4R2, `logical-system` is introduced in the statement for PKI SCEP certificate enrollment.

Description

Enroll and install a local digital certificate online by using Simple Certificate Enrollment Protocol (SCEP).

If you enter the `request security pki local-certificate enroll` command without specifying the `scep` or `cmpv2` keyword, SCEP is the default method for enrolling a local certificate.

Options

ca-profile *ca-profile-name*	CA profile name.
certificate-id *certificate-id-name*	Name of the local digital certificate and the public/private key pair.

challenge-password ***password***	Password set by the administrator and normally obtained from the SCEP enrollment webpage of the CA. The password is maximum 256 characters in length. You can enforce the limit to the required characters.
digest (sha-1 | sha-256)	Hash algorithm used for signing RSA certificates, either SHA-1 or SHA-256. SHA-1 is the default.
domain-name ***domain-name***	Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.
email ***email-address***	E-mail address of the certificate holder.
ip-address ***ip-address***	IP address of the router.
ipv6-address ***ipv6-address***	IPv6 address of the router for the alternate subject.
logical-system (*logical-system-name* | all)	Name of the logical system or all. This is optional.
scep-digest-algorithm (md5 | sha-1)	Hash algorithm digest, either MD5 or SHA-1; SHA-1 is the default.
scep-encryption-algorithm (des | des3)	Encryption algorithm, either DES or DES3; DES3 is the default.
subject ***subject-distinguished-name***	Distinguished Name (DN) format that contains the domain component, common name, department, serial number, company name, state, and country in the following format: DC, CN, OU, O, SN, L, ST, C. • `DC`—Domain component • `CN`—Common name • `OU`—Organizational unit name • `O`—Organization name • `SN`—Serial number of the device If you define SN in the subject field without the serial number, then the serial number is read directly from the device and added to the certificate signing request (CSR). • `ST`—State • `C`—Country

Required Privilege Level

maintenance and security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

command-name

```
user@host> request security pki local-certificate enroll scep certificate-id r3-entrust-scep ca-
profile entrust domain-name router3.example.net subject
"CN=router3,OU=Engineering,O=example,C=US" challenge-password 123
```

```
Certificate enrollment has started. To view the status of your enrollment, check the public key
infrastructure log (pkid) log file at /var/log/pkid. Please save the challenge-password for
revoking this certificate in future.  Note that this password is not stored on the router.
```

Sample Output

Sample output for vSRX Virtual Firewall 3.0

```
user@host> request security pki generate-key-pair certificate-id example
```

```
Generated key pair example, key size 2048 bits
```

```
user@host> request security pki local-certificate enroll certificate-id ?
```

```
Possible completions:
<certificate-id> Certificate identifier
example
```

```
user@host> request security pki generate-key-pair certificate-id Hub
```

```
error: Failed to generate key pair at HSM. Found a key with the same name at HSM. Use a
different certificate id next time. Refer to PKID logs for more details
```

RELATED DOCUMENTATION

- request security pki local-certificate enroll cmpv2
- show security pki local-certificate (View)
- clear security pki local-certificate (Device)

request security pki local-certificate export

IN THIS SECTION

- Syntax | 14822
- Description | 14822
- Options | 14822
- Required Privilege Level | 14822
- Output Fields | 14823
- Sample Output | 14823
- Release Information | 14823

Syntax

```
request security pki local-certificate export
```

Description

Export a generated self-signed certificate from the default location (var/db/certs/common/local) to a specific location within the device.

Options

certificate id ***certificate-id-name***	Name of the local digital certificate.
filename ***path/filename***	Target directory location and filename of the CA digital certificate.
type (der | pem)	Certificate format: DER (distinguished encoding rules) or PEM (privacy-enhanced mail).

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki local-certificate export

```
user@host> request security pki local-certificate export filename /var/tmp/my-cert.pem
certificate-id nss-cert type pem
certificate exported successfully
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Basic Elements of PKI in Junos OS

request security pki local-certificate generate-self-signed

IN THIS SECTION

- Syntax | 14824
- Description | 14824
- Options | 14824
- Required Privilege Level | 14824
- Output Fields | 14825
- Sample Output | 14825
- Release Information | 14825

Syntax

```
request security pki local-certificate  generate-self-signed certificate-id certificate-id-name
domain-name domain-name ip-address ip-address email email-address subject subject-distinguished-
name
```

Description

Manually generate a self-signed certificate for the given distinguished name.

Options

certificate-id ***certificate-id-name***	Name of the local digital certificate and the public/private key pair.
domain-name ***domain-name***	Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.
email ***email-address***	E-mail address of the certificate holder.
ip-address ***ip-address***	IP address of the router.
subject ***subject-distinguished-name***	Distinguished name format that contains the common name, department, company name, state, and country: • `CN`—Common name • `OU`—Organizational unit name • `O`—Organization name • `ST`—State • `C`—Country

Required Privilege Level

maintenance

security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

command-name

```
user@host> request security pki local-certificate generate-self-signed certificate-id self-cert
subject cn=abc domain-name example.net email user1@example.net
Self-signed certificate generated and loaded successfully
```

Release Information

Command introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Requesting for and Installing a Digital Certificates on Your Router

request security pki local-certificate generate-self-signed

IN THIS SECTION

- Syntax | 14826
- Description | 14826
- Options | 14826
- Required Privilege Level | 14826
- Output Fields | 14827
- Sample Output | 14827
- Release Information | 14827

Syntax

```
request security pki local-certificate  generate-self-signed certificate-id certificate-id-name
domain-name domain-name ip-address ip-address email email-address subject subject-distinguished-
name
```

Description

Manually generate a self-signed certificate for the given distinguished name.

Options

certificate-id *certificate-id-name*	Name of the local digital certificate and the public/private key pair.
domain-name *domain-name*	Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.
email *email-address*	E-mail address of the certificate holder.
ip-address *ip-address*	IP address of the switch.
subject *subject-distinguished-name*	Distinguished name format that contains the common name, department, company name, state, and country: • **CN**—Common name • **OU**—Organizational unit name • **O**—Organization name • **ST**—State • **C**—Country

Required Privilege Level

maintenance

security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki local-certificate generate-self-signed

```
user@switch> request security pki local-certificate generate-self-signed certificate-id self-
cert subject cn=abc domain-name abc.net email jdoe@abc.net
Self-signed certificate generated and loaded successfully
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Manually Generating Self-Signed Certificates on Switches (CLI Procedure)

request security pki local-certificate generate-self-signed (Security)

IN THIS SECTION

- Syntax | 14828
- Description | 14828
- Options | 14828
- Required Privilege Level | 14829
- Output Fields | 14829
- Sample Output | 14829
- Release Information | 14829

Syntax

```
request security pki local-certificate generate-self-signed certificate-id certificate-id-
namedomain-name domain-name subject subject-distinguished-name
<add-ca-constraint>
<digest (sha1 | sha256)>
<email email-address>
<ip-address ipv4-address>
<ipv6-address ipv6-address>
```

Description

Manually generate a self-signed certificate for the given distinguished name.

Options

`certificate-id` *`certificate-id-name`*—Name of the certificate and the public/private key pair.

`domain-name` *`domain-name`*—Fully qualified domain name (FQDN) provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.

`subject` *`subject-distinguished-name`*—Distinguished name format contains the following information:

- `DC`—Domain component
- `CN`—Common name
- `OU`—Organizational unit name
- `O`—Organization name
- `L`—Locality
- `ST`—State
- `C`—Country

`add-ca-constraint`—(Optional) Specifies that the certificate can be used to sign other certificates.

`digest`—(Optional) Hash algorithm used to sign the certificate.

- `sha1`—SHA-1 digest (default)
- `sha256`—SHA-256 digest

Starting in Junos OS Release 18.1R3, the default encryption algorithm that is used for validating automatically and manually generated self-signed PKI certificates is Secure Hash Algorithm 256 (SHA-256). Prior to Junos OS Release 18.1R3, SHA-1 is used as default encryption algorithm.

`email` *`email-address`*—(Optional) E-mail address of the certificate holder.

`ip-address` *`ipv4-address`*—(Optional) Static IPv4 address of the device.

`ipv6-address` *`ipv6-address`*—(Optional) Static IPv6 address of the device.

Required Privilege Level

maintenance and security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki local-certificate generate-self-signed certificate-id self-cert subject cn=abc domain-name example.net email mholmes@example.net

```
user@host> request security pki local-certificate generate-self-signed certificate-id self-cert
subject cn=abc domain-name example.net email mholmes@example.net
Self-signed certificate generated and loaded successfully
```

Release Information

Command introduced in Junos OS Release 9.1.

Support for `digest` option added in Junos OS Release 12.1X45-D10.

Support for `ipv6-address` option added in Junos OS Release 22.1R1.

RELATED DOCUMENTATION

clear security pki local-certificate (Device) | 4027

show security pki local-certificate (View) | 24992

request security pki local-certificate load

IN THIS SECTION

- Syntax | 14830
- Description | 14830
- Options | 14830
- Required Privilege Level | 14831
- Output Fields | 14831
- Sample Output | 14831
- Release Information | 14831

Syntax

```
request security pki local-certificate load filename ssl_proxy_ca.crt key ssl_proxy_ca.key
certificate-id certificate id
```

Description

Manually load a local digital certificate from a specified location.

Options

filename Filename that contains the certificate to load

key File pathname that contains the private key/key-pair to loaded

certificate-id Name of the certificate identifier

Starting in Junos OS Release 19.1R1, a commit check is added to prevent user from adding `.`, `/`, `%`, and space in a certificate identifier while generating a local or remote certificates or a key pair.

Required Privilege Level

maintenance and security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki local-certificate load

```
user@host> request security pki local-certificate load filename cert_name.crt key key_name.key
certificate-id test
Local certificate cert_name.crt loaded successfully
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- show security pki local-certificate (View) | 24992
- clear security pki local-certificate (Device) | 4027
- *request security pki local-certificate verify (Security)*
- *Basic Elements of PKI in Junos OS*

request security pki local-certificate re-enroll cmpv2

IN THIS SECTION

- Syntax | 14832
- Description | 14832

- Options | 14832
- Required Privilege Level | 14832
- Output Fields | 14832
- Release Information | 14833

Syntax

```
request security pki local-certificate re-enroll cmpv2 certificate-id certificate-id
<ca-profile-name ca-profile>
<re-generate-keypair>
```

Description

Manually reenroll an end-entity (EE) certificate with Certificate Management Protocol version 2 (CMPv2). This command allows the administrator to initiate renewal of the EE certificate using CMPv2 and can be used in conjunction with the `set security pki auto-re-enrollment cmpv2` automatic enrollment configuration.

Options

certificate-id *certificate-id-name*	Name of the local digital certificate.
ca-profile-name *ca-profile-name*	(Optional) CA profile name.
re-generate-keypair	(Optional) Generate a PKI public/private key pair for the EE certificate. Key generation might take a few seconds.

Required Privilege Level

maintenance and security

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

request security pki local-certificate enroll cmpv2 | 14814

request security pki local-certificate re-enroll scep

IN THIS SECTION

- Syntax | 14833
- Description | 14833
- Options | 14834
- Required Privilege Level | 14834
- Output Fields | 14834
- Release Information | 14835

Syntax

```
request security pki local-certificate re-enroll scep certificate-id certificate-id
<ca-profile-name ca-profile>
<challenge-password password>
<re-generate-keypair>
<scep-digest-algorithm (md5 | sha-1)>
<scep-encryption-algorithm (des | des3)>
```

Description

Manually reenroll an end-entity (EE) certificate with Simple Certificate Enrollment Protocol (SCEP). This command allows the administrator to initiate renewal of the EE certificate using SCEP and can be used in conjunction with the `set security pki auto-re-enrollment scep` automatic enrollment configuration.

Starting in Junos OS Release 20.1R1 on vSRX Virtual Firewall 3.0, you can safeguard the private keys used by PKID and IKED to establish a PKI based VPN tunnel using the keypairs generated at the Microsoft Azure Key Vault hardware security module (HSM) service and starting in Junos OS Release 20.4R1 on vSRX Virtual Firewall 3.0, the same feature is supported through AWS Key Management Service (KMS).

You cannot manually re-enroll the local certificates with the “re-generate key-pair” option. An error message is displayed.

Warning message upon re-enrollment - sample output:

```
[edit]
root@vsrx-1# ...te-id hsm1 ca-profile azure-ca challenge-password juniper re-generate-keypair
error: HSM Error: Re-enrollment is not allowed with re-generate key-pair option.
```

Options

certificate-id ***certificate-id-name***	Name of the local digital certificate.
ca-profile-name ***ca-profile-name***	(Optional) CA profile name.
challenge-password ***password***	Password set by the administrator and normally obtained from the SCEP enrollment webpage of the CA. The password is 16 characters in length.
re-generate-keypair	(Optional) Generate a PKI public/private key pair for the EE certificate. Key generation might take a few seconds.
scep-digest-algorithm	(Optional) Hash algorithm digest, either MD5 or SHA-1; SHA-1 is the default.
scep-encryption-algorithm	(Optional) Encryption algorithm, either DES or DES3; DES3 is the default.

Required Privilege Level

maintenance and security

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

request security pki local-certificate enroll scep

request security pki local-certificate verify

IN THIS SECTION

- Syntax | 14835
- Description | 14835
- Options | 14835
- Required Privilege Level | 14836
- Output Fields | 14836
- Sample Output | 14836
- Release Information | 14836

Syntax

```
request security pki local-certificate verify certificate-id certificate-id-name
```

Description

Verify the validity of the local digital certificate identifier.

Options

certificate-id ***certificate-id-name***	Display the specified certificate identifier name.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

command-name

You receive the following response before the certificate revocation list (CRL) is downloaded:

```
request security pki local-certificate verify certificate-id bme1 (not downloaded)
user@host> request security pki local-certificate verify certificate-id bme1

Local certificate bme1: CRL verification in progress. Please check the PKId debug logs for
completion status
```

command-name

You receive the following response after the certificate revocation list (CRL) is downloaded:

```
request security pki local-certificate verify certificate bme1 (downloaded)
user@host> request security pki local-certificate verify certificate-id bme1
Local certificate bme1 verification success
```

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

show security pki local-certificate | 24982

request security pki node-local generate-certificate-request

IN THIS SECTION

- Syntax | 14837
- Description | 14837
- Options | 14837
- Required Privilege Level | 14838
- Output Fields | 14838
- Sample Output | 14839
- Release Information | 14839

Syntax

```
request security pki node-local generate-certificate-request certificate-id certificate-id-name
domain-name domain-name subject subject-distinguished-name
<digest (sha1 | sha256)>
<email email-address>
<filename (path | terminal)>
<ip-address ip-address>
```

Description

Manually generate a local digital certificate request in the Public-Key Cryptography Standards #10 (PKCS-10) format on the local device in a Multinode High Availability setup.

Options

certificate-id *certificate-id-name*	Name of the local digital certificate and the public/private key pair.
domain-name *domain-name*	Fully qualified domain name (FQDN) provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.

subject ***subject-distinguished-name***	Distinguished name format contains the following information: • `DC`—Domain component • `CN`—Common name • `OU`—Organizational unit name • `O`—Organization name • `L`—Locality • `ST`—State • `C`—Country
digest	(Optional) Hash algorithm used to sign the certificate request. • `sha-1`—SHA-1 digests (default value for RSA or DSA only). • `sha-256`—SHA-256 digests for RSA or ECDSA only (default value for ECDSA). • `sha-384`—SHA-384 digests for ECDSA only.
email ***email-address***	(Optional) E-mail address of the certificate holder.
filename (*path* | terminal)	(Optional) Location where the local digital certificate request should be placed or the login terminal.
ip-address ***ip-address***	(Optional) IP address of the router.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki generate-certificate-request

```
user@host> request security pki node-local generate-certificate-request certificate-id local-
entrust2 domain-name router2.example.net filename entrust-req2 subject cn=router2.example.net

Generated certificate request
-----BEGIN CERTIFICATE REQUEST-----
MIIBoTCCAQoCAQAwGjEYMBYGA1UEAxMPdHAxLmp1bmlwZXIubmV0MIGfMA0GCSqG
SIb3DQEBAQUAA4GNADCBiQKBgQCiUFklQws1Ud+AqN5DDxRs2kVyKEhh9qoVFnz+
Hz4c9vsy3B8ElwTJlkmIt2cB3yifB6zePd+6WYpf57Crwre7YqPkiXM31F6z3YjX
H+1BPNbCxNWYvyrnSyVYDbFj8o0Xyqog8ACDfVL2JBWrPNBYy7imq/K9soDBbAs6
5hZqqwIDAQABoEcwRQYJKoZIhvcNAQkOMTgwNjAOBgNVHQ8BAf8EBAMCB4AwJAYD
VR0RAQH/BBowGIIWdHAxLmVuZ2xhYi5qdW5pcGVyLm5ldDANBgkqhkiG9w0BAQQF
AAOBgQBc2rq1v5SOQXH7LCb/FdqAL8ZM6GoaN5d6cGwq4bB6a7UQFgtoH406gQ3G
3iH0Zfz4xMIBpJYuGd1dkqgvcDoH3AgTsLkfn7Wi3x5H2qeQVs9bvL4P5nvEZLND
EIMUHwteolZCiZ70fO9Fer9cXWHSQs1UtXtgPqQJy2xIeImLgw==
-----END CERTIFICATE REQUEST-----
Fingerprint:
0d:90:b8:d2:56:74:fc:84:59:62:b9:78:71:9c:e4:9c:54:ba:16:97 (sha1)
1b:08:d4:f7:90:f1:c4:39:08:c9:de:76:00:86:62:b8 (md5)
```

Release Information

Command introduced in Junos OS Release 22.3R1

RELATED DOCUMENTATION

Multinode High Availability

request security pki node-local generate-key-pair

IN THIS SECTION

- Syntax | 14840
- Description | 14840
- Options | 14840
- Required Privilege Level | 14841
- Output Fields | 14841
- Sample Output | 14841
- Release Information | 14841

Syntax

```
request security pki node-local generate-key-pair certificate-id certificate-id-name
<size (256 | 384 | 521 |1024 | 2048 | 4096)>
<type (dsa | ecdsa | rsa)>
```

Description

Generate a public key infrastructure (PKI) public/private key pair for a local digital certificate on the local node in a Multinode High Availability setup.

Options

certificate-id *certificate-id-name*
: Name of the local digital certificate and the public/private key pair.

size
: Key pair size. The key pair size can be 256, 384, 521, 1024, 2048, or 4096 bits.

  Key size compatibility

  - ECDSA-256, 384, and 521

- DSA and RSA - 1024, 2048, or 4096. The default key pair size is 1024 for DSA and 2048 for RSA.

When you use ECDSA-521 signatures, you can:

- Load a complete certificate, which is generated using an external tool like OpenSSL into PKI.
- Manually generate a Certificate Signing Request (CSR) for a local certificate and sending the CSR to a (Certificate Authority) CA server to enroll.
- Automatic enroll with CA server.

type The algorithm to be used for encrypting the public/private key pair:

- `ecdsa`—ECDSA encryption
- `dsa`— DSA encryption
- `rsa`—RSA encryption (default)

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki generate-key-pair

```
user@host> request security pki node-local generate-key-pair certificate-id cert-123 type rsa
size 4096
Generated key pair cert-123, key size [4096] bits
```

Release Information

Command introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Multinode High Availability

clear security pki node-local key-pair

request security pki node-local key-pair export

request security pki node-local key-pair export

IN THIS SECTION

- Syntax | 14842
- Description | 14842
- Options | 14843
- Required Privilege Level | 14843
- Output Fields | 14843
- Release Information | 14843

Syntax

```
request security pki node-local key-pair export certificate-id certificate-id filename filename
<passphrase string>
< type (der | pem)>
```

Description

Export the keypair for an end-entity (EE) certificate. Junos OS encrypts the exported keypair.

You can export the PKI key-pairs file as a backup or to check the file for troubleshooting purposes.

We recommend providing permission to the `request security pki node-local key-pair export` command only to the privileged users.

Options

certificate-id *certificate-id*	Name of the local digital certificate.
filename *filename*	Target directory location and filename of the CA digital certificate.
passphrase *passphrase*	(Optional) Passphrase to protect the keypair data for PEM format. The passphrase can be up to 64 characters. If specified, the passphrase must be used when importing the keypair.
type (der | pem)	(Optional) Type of format, either DER or PEM. PEM is the default.

Required Privilege Level

maintenance

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 22.3R1

RELATED DOCUMENTATION

Multinode High Availability

request security pki node-local generate-key-pair

clear security pki node-local key-pair

request security pki node-local local-certificate enroll

IN THIS SECTION

Syntax | 14844

- Description | **14844**
- Options | **14845**
- Required Privilege Level | **14846**
- Output Fields | **14846**
- Sample Output | **14846**
- Release Information | **14846**

Syntax

```
request security pki node-local local-certificate enroll
    ca-dn  subject-dn
    ca-profile  ca-profile name
    ca-reference  reference
    ca-secret  shared-secret
    certificate-id certificate-id-name
    challenge-password password
    cmpv2
    digist
    domain-name domain-name
    email email-address
    ip-address ip-address
    ipv6-address ipv6-address
    scep
    scep-digest-algorithm
    scep-encryption-algorithm
    subject subject-distinguished-name
```

Description

Enroll and install a local digital certificate online by using CMPv2 or Simple Certificate Enrollment Protocol (SCEP). This command loads both end-entity (EE) and CA certificates based on the CA server configuration. Certificate revocation list (CRL) or Online Certificate Status Protocol (OCSP) can be used to check the revocation status of a certificate.

Options

ca-profile *ca-profile-name*	CA profile name.
certificate-id *certificate-id-name*	Name of the local digital certificate and the public/private key pair.
challenge-password *password*	Password set by the administrator and normally obtained from the SCEP enrollment webpage of the CA. The password is 16 characters in length
cmpv2	Enroll certificate using CMPv2 protocol.
domain-name *domain-name*	Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.
email *email-address*	E-mail address of the certificate holder.
ip-address *ip-address*	IP address of the router.
ipv6-address *ipv6-address*	IPv6 address of the router for the alternate subject.
scep	Enroll certificate using Simple Certificate Enrollment Protocol (SCEP) protocol.
scep-digest-algorithm	Hash algorithm digest, either MD5 or SHA-1; SHA-1 is the default.
scep-encryption-algorithm	Encryption algorithm, either DES or DES3; DES3 is the default.
subject *subject-distinguished-name*	Distinguished Name (DN) format that contains the domain component, common name, department, serial number, company name, state, and country in the following format: DC, CN, OU, O, SN, L, ST, C. • `DC`—Domain component • `CN`—Common name • `OU`—Organizational unit name • `O`—Organization name • `SN`—Serial number of the device

If you define SN in the subject field without the serial number, then the serial number is read directly from the device and added to the certificate signing request (CSR).

- `ST`—State
- `C`—Country

Required Privilege Level

maintenance and security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

command-name

```
user@host> request security pki node-local local-certificate enroll cmpv2 ca-profile root-552 ca-
dn DC=example,CN=root-552 certificate-id tc552 email tc552-root@example.net domain-name
example.net ip-address 10.192.0.22 ca-secret example ca-reference 51892 subject
CN=example,OU=SBU,O=552-22

Certificate enrollment has started. To view the status of your enrollment, check the public key
infrastructure log (pkid) log file at /var/log/pkid.
```

Release Information

Command introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Multinode High Availability

show security pki node-local local-certificate

request security pki node-local local-certificate load

request security pki node-local local-certificate export

request security pki node-local local-certificate re-enroll

request security pki node-local local-certificate export

IN THIS SECTION

- Syntax | 14847
- Description | 14847
- Options | 14847
- Required Privilege Level | 14848
- Output Fields | 14848
- Sample Output | 14848
- Release Information | 14848

Syntax

```
request security pki node-local local-certificate export certificate id certificate-id-name
filename path/filename
<type (der | pem)>
```

Description

Export a generated self-signed certificate from the default location (var/db/certs/common/local) to a specific location within the local device in a Multinode High Availability setup.

Options

certificate id *certificate-id-name*	Name of the local digital certificate.
filename *path/filename*	Target directory location and filename of the CA digital certificate.
type (der | pem)	Certificate format: DER (distinguished encoding rules) or PEM (privacy-enhanced mail).

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki node-local local-certificate export

```
user@host> request security pki node-local local-certificate export filename /var/tmp/my-
cert.pem certificate-id nss-cert type pem
certificate exported successfully
```

Release Information

Command introduced in Junos OS Release 22.3R1

RELATED DOCUMENTATION

Multinode High Availability

show security pki node-local local-certificate

request security pki node-local local-certificate enroll

request security pki node-local local-certificate load

request security pki node-local local-certificate re-enroll

request security pki node-local local-certificate load

IN THIS SECTION

Syntax | 14849

- Description | 14849
- Options | 14849
- Required Privilege Level | 14849
- Output Fields | 14849
- Sample Output | 14850
- Release Information | 14850

Syntax

```
request security pki node-local local-certificate load certificate-id certificate-id
filename file-name key key-string
<passphrase passphrase-string>
```

Description

Manually load a local digital certificate from a specified location on the local device in a Multinode High Availability setup.

Options

certificate-id Name of the certificate identifier

filename Filename that contains the certificate to load

key File pathname that contains the private key/key-pair to loaded

passphrase Passphrase of the private key/key-pair (PEM) file

Required Privilege Level

maintenance and security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki node-local local-certificate load

```
user@host> request security pki node-local local-certificate load filename cert_name.crt key
key_name.key certificate-id test
Local certificate cert_name.crt loaded successfully
```

Release Information

Command introduced in Junos OS Release 22.3R1

RELATED DOCUMENTATION

- *Multinode High Availability*
- *show security pki node-local local-certificate*
- *request security pki node-local local-certificate enroll*
- request security pki node-local local-certificate export
- *request security pki node-local local-certificate re-enroll*

request security pki node-local local-certificate re-enroll

IN THIS SECTION

- Syntax | 14851
- Description | 14851
- Options | 14851
- Required Privilege Level | 14851
- Output Fields | 14851
- Release Information | 14852

Syntax

```
request security pki node-local local-certificate re-enroll (cmpv2 | scep) certificate-id
certificate-id ca-profile profile-name challenge-password password
<re-generate-keypair>
<scep-digest-algorithm>
<scep-encryption-algorithm>
```

Description

Manually reenroll an end-entity (EE) certificate with Certificate Management Protocol version 2 (CMPv2) or with Simple Certificate Enrollment Protocol (SCEP). This command initiates renewal of the EE certificate using the selected protocol and you can use the command in conjunction with the `set security pki auto-re-enrollment` command for automatic enrollment.

Options

scep	Enroll end-entity certificate using SCEP protocol
cmpv2	Enroll certificate using CMPv2 protocol
ca-profile-name *ca-profile-name*	(Optional) CA profile name.
certificate-id *certificate-id-name*	Name of the local digital certificate.
challenge-password	Password used by CA for enrollment and revocation
re-generate-keypair	(Optional) Generate a PKI public/private key pair for the EE certificate. Key generation might take a few seconds.
scep-digest-algorithm	Hash algorithm used for SCEP-PKCS7
scep-encryption-algorithm	Encryption algorithm used for SCEP-PKCS7

Required Privilege Level

maintenance and security

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

- *Multinode High Availability*
- *show security pki node-local local-certificate*
- *request security pki node-local local-certificate enroll*
- *request security pki node-local local-certificate load*
- request security pki node-local local-certificate export

request security pki node-local local-certificate verify

IN THIS SECTION

- Syntax | 14852
- Description | 14852
- Options | 14853
- Required Privilege Level | 14853
- Output Fields | 14853
- Sample Output | 14853
- Release Information | 14853

Syntax

```
request security pki node-local local-certificate verify certificate-id  certificate-id-name
```

Description

Verify the validity of the local digital certificate identifier on the local node in Multinode High Availability setup.

Options

`certificate-id` *`certificate-id-name`* — Name of the local digital certificate identifier.

Required Privilege Level

maintenance and security

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki node-local local-certificate verify certificate-id bme1 (not downloaded)

You receive the following response before the certificate revocation list (CRL) is downloaded:

```
user@host> request security pki node-local local-certificate verify certificate-id
bme1
Local certificate bme1: CRL verification in progress. Please check the PKId debug logs for
completion status
```

request security pki local-certificate verify certificate bme1 (downloaded)

You receive the following response after the certificate revocation list (CRL) is downloaded:

```
user@host> request security pki node-local local-certificate verify certificate-id
bme1
Local certificate bme1 verification success
```

Release Information

Command introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Multinode High Availability

show security pki node-local local-certificate

request security pki node-local local-certificate re-enroll

request security pki node-local local-certificate load

request security pki node-local local-certificate export

request security pki node-local local-certificate enroll

request security pki sync-from-peer

IN THIS SECTION

- Syntax | 14854
- Description | 14854
- Required Privilege Level | 14855
- Output Fields | 14855
- Sample Output | 14855
- Release Information | 14855

Syntax

```
request security pki sync-from-peer
```

Description

Synchronize the PKI file system on the peer node in a Multinode High Availability setup. You can use this command to replicate the PKI directory in the remote node to your local node. Replicating PKI directory is helpful when one of the two nodes or ICL goes down.

Note that you can run this command only you've enabled Multinode High Availabiliy.

Consider a set up with node 0 (local node) and node 1 (remote node). To replicate the PKI directory of the remote node (node 1), run this command in your local node (node 0).

When you run this command on your local node, all the local PKI files are deleted and replaced by the remote node PKI directory. Hence, be sure on which node you are executing this command. After running this command, we recommend you to verify whether the files are synchronized between the two nodes.

Required Privilege Level

maintenance

Output Fields

This command produces no output.

Sample Output

request security pki sync-from-peer

```
user@host> request security pki sync-from-peer
File syncing is in progress... This will take a few seconds. Please confirm that the files are
synched. If not, run this command once again.
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

- ike (High Availability) | 8424
- ipsec (High Availability) | 9301
- High-Availability (Chassis)
- Multinode High Availability

request security pki verify-integrity-status

IN THIS SECTION

- Syntax | 14856
- Description | 14856
- Required Privilege Level | 14856
- Output Fields | 14856
- Sample Output | 14857
- Release Information | 14857

Syntax

```
request security pki verify-integrity-status
```

Description

Verify the integrity of public key infrastructure (PKI) files. This feature is supported only on SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall instances.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki verify-integrity-status

```
user@host> request security pki verify-integrity-status
All PKI objects: verification success
```

Release Information

Command introduced in Junos OS Release 11.2.

Do not use this command for non-FIPS or Common Criteria releases. We recommend that you do not use this command for any Junos OS Release 15.1X49-D40 or later releases.

request security policies check

IN THIS SECTION

- Syntax | 14857
- Description | 14858
- Options | 14858
- Additional Information | 14858
- Required Privilege Level | 14859
- Sample Output | 14859
- Release Information | 14860

Syntax

```
request security policies check
<from-zone zone name>
<global>
<logical-system (logical-system name | all)>
<pfe>
```

```
<root-logical-system>
<to-zone zone-name>
<tenant tenant-name>
```

Description

Displays the security policy sync status between the Routing Engine and the Packet Forwarding Engine. Use the command to display a list of all security polices which are in-sync or out-of-sync on the device.

Use the `show security policies checksum` command to display the security policy checksum value and use the `request security policies resync` command to synchronize the configuration of security policies in the Routing Engine and Packet Forwarding Engine.

Options

<from-zone *zone-name*	Displays security policies sync status from this zone.
global	Displays global policies sync status.
logical-system (*logical-system name* | all)	Displays security policies sync status for the security policies configured on a logical system or on all logical systems.
pfe	Displays security policies sync status for the security policies on the Packet Forwarding Engine.
root-logical-system	Displays security policies sync status for the security policies configured on the root logical system. This is the default outcome.
to-zone *zone-name*	Displays security policies sync status to this zone.
tenant *tenant-name*	Displays security policies sync status for the security policies configured on a tenant.

Additional Information

Security policies are stored in the routing engine and the packet forwarding engine. Security policies are pushed from the Routing Engine to the Packet Forwarding Engine when you commit configurations. If the security policies on the Routing Engine are out of sync with the Packet Forwarding Engine, the commit of a configuration fails. Core dump files may be generated if the commit is tried repeatedly. The out of sync can be due to:

- A policy message from Routing Engine to the Packet Forwarding Engine is lost in transit.

- An error with the routing engine, such as a reused policy UID.

When the policy configurations are modified and the policies are out of sync, the following error message displays - `error: Warning: policy might be out of sync between RE and PFE <SPU-name(s)>. Please request security policies check/resync`.

Required Privilege Level

maintenance

Sample Output

request security policies check

```
user@host> request security policies check
Start sending policies ...
Success
Total sent 1 policy.

Policy Checking Result:
 PFE master            in-sync
```

request security policies check logical-system LSYS1

```
user@host> request security policies check logical-system LSYS1
Start sending policies ...
Success
Total sent 1 policy.

Policy Checking Result:
 PFE fpc0.pic3         in-sync
 PFE fpc0.pic1         in-sync
 PFE fpc0.pic0         in-sync
 PFE fpc0.pic2         in-sync
```

request security policies check logical-system all

```
user@host> request security policies check logical-system all
Start sending policies ...
Success
Total sent 2 policies.

Policy Checking Result:
 PFE fpc0.pic1         in-sync
 PFE fpc0.pic3         in-sync
 PFE fpc0.pic2         in-sync
 PFE fpc0.pic0         in-sync
```

request security policies check from-zone trust to-zone untrust

```
user@host> request security policies check from-zone trust to-zone untrust
Start sending policies ...
Success
Total sent 2 policies.

Policy Checking Result:
 PFE fpc4.pic0         in-sync
 PFE fpc4.pic1         out-of-sync
```

Release Information

Command introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

- request security policies resync | 14861
- show security policies checksum | 25056

request security policies resync

IN THIS SECTION

- Syntax | 14861
- Description | 14861
- Options | 14861
- Additional Information | 14862
- Required Privilege Level | 14862
- Sample Output | 14862
- Release Information | 14863

Syntax

```
request security policies resync
<from-zone zone name>
<global>
<logical-system (logical-system name | all)>
<pfe>
<root-logical-system>
<to-zone zone-name>
<tenant tenant-name>
```

Description

Synchronize the configuration of security policies in the Routing Engine and Packet Forwarding Engine.

This command recovers the security policies in the Packet Forwarding Engine. If policy inconsistencies between the Routing Engine and Packet Forwarding Engine are determined, the security policies resync.

Options

<from-zone *zone-name*	Recover the policies from this zone.
global	Recover global policies.

logical-system (*logical-system name* | all)	Recover the policies on all logical systems or on a particular logical system.
pfe	Recover the policies on the Packet Forwarding Engine.
root-logical-system	Recover the policies on the root logical system. This is the default option.
to-zone *zone-name*	Recover the policies to this zone.
tenant *tenant-name*	Recover the policies of a tenant.

Additional Information

Security policies are stored in the routing engine and the packet forwarding engine. Security policies are pushed from the Routing Engine to the Packet Forwarding Engine when you commit configurations. If the security policies on the Routing Engine are out of sync with the Packet Forwarding Engine, the commit of a configuration fails. Core dump files may be generated if the commit is tried repeatedly. The out of sync can be due to:

- A policy message from Routing Engine to the Packet Forwarding Engine is lost in transit.
- An error with the routing engine, such as a reused policy UID.

When the policy configurations are modified and the policies are out of sync, the following error message displays - `error: Warning: policy might be out of sync between RE and PFE <SPU-name(s)>. Please request security policies check/resync`.

Use the `show security policies checksum` command to display the security policy checksum value and use the `request security policies check` to display the security policy sync status.

Required Privilege Level

maintenance

Sample Output

request security policies resync

```
user@host> request security policies resync
Start sending policies ...
```

```
Success
Total sent 1 policy.
```

request security policies resync logical-system all

```
user@host> request security policies resync logical-system all
Start sending policies ...
Success
Total sent 2 policies.
```

Release Information

Command introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

request security policies check | 14857
show security policies checksum | 25056

request security re-distribution ipsec-vpn

IN THIS SECTION

- Syntax | 14864
- Description | 14864
- Options | 14864
- Required Privilege Level | 14865
- Release Information | 14865

Syntax

```
request security re-distribution ipsec-vpn
gateway-name <gateway-name>
fpc <fpc-number>
pic <pic-number>
[thread-id <tid>]
[remote-id <rid>]
```

Description

Redistribute the tunnels that belongs to a Auto VPN or site-to-site gateway to a new processing unit.

This command migrates the tunnels only once and is valid only for 30 minutes, if the peer does not bring up the tunnel(s) immediately. After execution of the command, subsequent tunnels for the peer is established on the same FPC, PIC, and thread-id (only if specified).

In case of Auto VPN gateways, once the tunnels are brought down, it is expected that peer re-establishes the tunnel.

This command causes traffic disruption when used on an already established tunnel. If the command is used on a tunnel which is already anchored on the destination processing unit, it will not tear down the tunnel and re-establish it.

This feature is supported only on SRX5K-SPC3 (SPC3) card and in mixed-mode (SPC3 or SRX5K-SPC-4-15-320 (SPC2) cards).

When a tunnel goes down, you can use only the syslog to trace why a tunnel is anchored on a different processing unit.

If you want to migrate the tunnel back to the previous FPC or PIC (that is, default profile), you can either redistribute the tunnel again or run the `clear security ike security-associations index` *`SA-index-number`* command.

Options

gateway-name ***gateway-name***	Name of the gateway.
fpc ***fpc-number***	FPC slot number (0..63).
pic ***pic-number***	PIC slot number (0..3).
thread-id ***tid***	(Optional) Thread ID number. Only valid for SPC3. (1..27)

remote-id *rid*

If you provide Auto VPN as a gateway, then it is mandatory to provide the *remote-id*. If you provide site-to-site as a gateway, then you need not provide the *remote-id*.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show security ipsec tunnel-distribution | 24755

show security re-distribution ipsec-vpn | 25093

request security ssh key-pair-identity generate

IN THIS SECTION

- Syntax | 14866
- Description | 14866
- Options | 14866
- Required Privilege Level | 14866
- Output Fields | 14866
- Sample Output | 14867
- Release Information | 14867

Syntax

```
request security ssh key-pair-identity generate <identity-name> passphrase passphrase
```

Description

Generate the SSH private and public key pair for a specified identity. The private and public key files are stored in the **/var/db** directory, which is accessible through root only. Filenames are based on the `identity-name` with extensions. The files are similar to the certificate files that are stored in Junos OS.

Options

- *identity-name*—Identity name.
- `passphrase` *passphrase*— An SSH identity generated with a passphrase. The passphrase is used to protect the private key file stored in the file system. This option does not allow the user to enter a weak passphrase, which ensures stronger security. A private key is used to connect to a remote server and is never displayed or transferred between servers, even if the system is compromised. The private key cannot be used to connect to a remote server if the passphrase is not known.

> **NOTE**: By default, the `passphrase` uses Advanced Encryption Standard (AES) 128 in cipher block chaining (CBC) mode to encrypt a private key. All generated keys are stored in the **/var/db/ssh_key** directory.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security ssh key-pair-identity with passphrase

```
user@host> request security ssh key-pair-identity generate myident passphrase 1q2w3e
Created SSH key myident
```

Release Information

Command introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

show security ssh key-pair-identity | 25144

clear security ssh key-pair-identity | 4044

request security ssh password-less-authentication (Junos OS Evolved)

IN THIS SECTION

- Syntax | 14868
- Description | 14868
- Options | 14868
- Required Privilege Level | 14868
- Sample Output | 14869
- Release Information | 14869

Syntax

```
request security ssh password-less-authentication remote-host host remote-user username
<all-conf-users>
<routing-instance routing-instance>
```

Description

Configure SSH key-based authentication between the local device and a remote server.

When you execute the command, the device generates SSH keys for the current user, provided the user does not already have existing keys, and transfers the user's public key to the **authorized_keys** file of the specified user on the remote host. If you include the `all-conf-users` option, the device performs the operation for all configured users with SSH keys instead of only for the current user.

You can back up your device's configuration to an archive site periodically or after each commit by configuring statements at the `[edit system archival configuration]` hierarchy level. Before you can configure the device to archive the configuration using an SCP-style URL, you must set up passwordless authentication between the local device (as the root user) and the remote server (for a remote user configured on the server).

Options

all-conf-users	(Optional) Configure SSH key-based authentication for all configured users that have an SSH key pair. If you omit this option, the device configures SSH key-based authentication for the user executing the command.
remote-host *host*	Hostname or IP address of the remote server.
remote-user *username*	Username that the local device uses to log in to the remote server.
routing-instance *routing-instance-name*	(Optional) Name of the routing instance through which to reach the remote server.

Required Privilege Level

maintenance

Sample Output

request security ssh password-less-authentication

```
root@host> request security ssh password-less-authentication remote-host 198.51.100.1 remote-
user admin
INFO: Reading public key of current logged in user 'root'.
INFO: Logging in to 198.51.100.1 as 'admin' user to append content of public key file(s), please
authenticate..
The authenticity of host '198.51.100.1 (198.51.100.1)' can't be established.
ECDSA key fingerprint is SHA256:ZZYzAS1mzURvmmxnoUkDuBp3UfS3ryD46CDei1MoTT4.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '198.51.100.1' (ECDSA) to the list of known hosts.
admin@198.51.100.1's password:
INFO: Password-less authentication from host (for user 'root') to 198.51.100.1 (for user
'admin') is setup.
```

Release Information

Command introduced in Junos OS Evolved Release 22.3R1.

RELATED DOCUMENTATION

Backing Up Configurations to an Archive Site

request security system-keystore

IN THIS SECTION

- Syntax | 14870
- Description | 14870
- Options | 14870
- Required Privilege Level | 14870
- Release Information | 14870

Syntax

```
request security system-keystore import keyname key-name (private-key | public-key | x509-cert)
key-file
request security system-keystore unlink key-name
```

Description

Add or remove keys from the system IMA keyring.

Options

import Import a key of the specified type onto the system IMA key chain.

unlink Remove the key with the specified name from the system IMA key chain.

Required Privilege Level

maintenance

Release Information

Statement introduced in Junos OS Evolved Release 20.1R1.

request security system-keystore unlink

IN THIS SECTION

- Syntax | 14871
- Description | 14871
- Options | 14871
- Required Privilege Level | 14871
- Release Information | 14871

Syntax

```
request security system-keystore unlink key-name
<force>
```

Description

Remove the key with the specified name from the system IMA key chain. This command will not remove a key if there are third-party packages that require the key, unless you use the `force` option.

Options

force *(Optional)* Force the removal of a key.

Required Privilege Level

maintenance

Release Information

Statement introduced in Junos OS Evolved Release 20.1R1.

`force` option introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

No Link Title

request security system-keystore | 14869

request security tpm master-encryption-password set

IN THIS SECTION

Syntax | 14872

- Description | 14872
- Options | 14872
- Required Privilege Level | 14872
- Output Fields | 14872
- Sample Output | 14873
- Release Information | 14873

Syntax

```
request security tpm master-encryption-password set plain-text-password
```

Description

Use this command to set or replace the password (in plain text).

> **NOTE**: This command runs on the primary Routing Engine and synchronizes the password to the backup Routing Engine. If you use this command while the backup Routing Engine is unreachable, it will return an error and you will not be able to set the password until the backup Routing Engine is available.

Options

plain-text-password	Set or replace the password (in plain text).

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security tpm master-encryption-password set plain-text-password

```
user@host> request security tpm master-encryption-password set plain-text-password

master encryption password change requested
Enter current master encryption password:
Enter new master encryption password:
Repeat new master encryption password:
Binding password with TPM
Master encryption password is bound to TPM
Encoding master password ..
Successfully encoded master password
Encoding master password ..
Successfully encoded master password
Successfully generated and saved TPM encrypted hash of config file
Encrypted key-pair files
Syncing password to other re: re1
Password should be changed whenever redundant RE is replaced in future
```

Release Information

Command introduced in Junos OS Release 15.1X49-D80.

RELATED DOCUMENTATION

show security tpm status | 25150

request security uefi revoke

IN THIS SECTION

- Syntax | 14874
- Description | 14874

- Options | 14874
- Required Privilege Level | 14874
- Release Information | 14874

Syntax

```
request security uefi revoke
<development-keys>
<load-cert-file>
```

Description

Revoke the platform Unified Extensible Firmware Interface (UEFI) keys.

Options

development-keys (Optional) Revoke the development keys for the operating system.

> **WARNING**: If you revoke the development keys, the device will not boot up. You will have to manually reprogram the BIOS image provided by Juniper Networks.

load-cert-file (Optional) Revoke the certificate. Whenever any key is compromised, the keys and certificates have to be revoked in a controlled environment.

Required Privilege Level

view

Release Information

Statement introduced in Junos OS Evolved Release 18.4R1.

request security user-identification local-authentication-table delete

IN THIS SECTION

- Syntax | 14875
- Description | 14875
- Options | 14875
- Required Privilege Level | 14875
- Output Fields | 14876
- Sample Output | 14876
- Release Information | 14876

Syntax

```
request security user-identification local-authentication-table delete ip-address | user-name
```

Description

This command removes an entry from the local authentication table. You can identify the entry by IP address or user-name. To change the user name of an entry or to remove or change entries in a role list, you must delete the existing entry and create a new one.

Options

ip-address—The IP address of the entry to be deleted.

user-name—The user name of the entry to be deleted. To change the user name of an entry, you must delete the old entry and create a new one.

Required Privilege Level

maintenance

Output Fields

The specified show command verifies the table content before and after an entry has been deleted from the local authentication table.

Sample Output

command-name

```
user@host> show security user-identification local-authentication-table all
Total entries: 2
    Ip-address: 192.0.2.1
    Username: user1
    Roles: role1

    Ip-address: 203.0.113.2
    Username: user2
    Roles: role2, role3, role1

user@host> request security user-identification local-authentication-table delete
203.0.113.2
user@host> show security user-identification local-authentication-table all
Total entries: 1
    Ip-address: 192.0.2.1
    Username: user1
    Roles: role1
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

request security user-identification local-authorization-table add | 14877

Understanding the User Identification Table

request security user-identification local-authorization-table add

IN THIS SECTION

- Syntax | 14877
- Description | 14877
- Options | 14877
- Required Privilege Level | 14878
- Output Fields | 14878
- Sample Output | 14878
- Release Information | 14878

Syntax

```
request security user-identification local-authorization-table add user user-name ip-address ip-address roles [role-name]
```

Description

This command adds user and role information to the local authentication table. The table is used to retrieve user and role information for traffic from the specified IP address to enforce a user role firewall.

To add an entry, specify the user name, IP address, and up to 40 roles to be associated with this user. Subsequent commands for the same user and IP address aggregates any new roles to the existing list. An authentication entry can contain up to 200 roles.

> **NOTE**: To change the user name of an entry or to remove or change entries in a role list, you must delete the existing entry and create a new one.

An IP address can be associated with only one user. If a second request is made to add a different user using the same IP address, the second authentication entry overwrites the existing entry.

Options

user *user-name*—Specify the name of the user to be added to the table.

ip-address *ip-address*—Specify the IP address of the user. Either IPv4 or IPv6 addresses are supported.

roles [*role-name*]—(Optional) Specify the role or list of roles to be associated with the specified user. If the specified user and IP address already exist, any roles specified in the command are added to the existing role list.

Required Privilege Level

maintenance

Output Fields

When you enter this command, either an entry is added to the local authentication table, or the roles of an existing entry are aggregated with additional roles.

Sample Output

request security user-identification local-authentication-table add

```
user@host> request security user-identification local-authentication-table add user user1 ip-address 192.0.2.1 roles role1
user@host> request security user-identification local-authentication-table add user user2 ip-address 203.0.113.2 roles [role2 role3]
user@host> request security user-identification local-authentication-table add user user2 ip-address 203.0.113.2 roles role1
user@host> show security user-identification local-authentication-table all
Total entries: 2
Source IP       Username         Roles
192.0.2.1       user1                  role1
203.0.113.2     user2                  role2, role3, role1
```

Release Information

Command introduced in Junos OS Release 12.1. Command updated in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

request security user-identification local-authentication-table delete | 14875

Understanding the User Identification Table

request security utm anti-virus avira-engine

IN THIS SECTION

- Syntax | 14879
- Description | 14879
- Options | 14879
- Required Privilege Level | 14880
- Output Fields | 14880
- Sample Output | 14880
- Sample Output | 14880
- Sample Output | 14880
- Sample Output | 14881
- Release Information | 14881

Syntax

```
request security utm anti-virus avira-engine
```

Description

Manually update the Avira antivirus pattern database using the command described. To update automatically you use the configuration statement `set security utm default-configuration anti-virus avira-engine pattern-update interval` *seconds*. Avira is an internal scan engine that provides a full file-based anitvirus scanning feature. The full file-based antivirus scanning feature is a separately licensed subscription service. The Avira scan engine is provided as a downloadable Content Security module. You can download and install virus signature database.

Options

- `pattern-delete` — Delete the current Avira antivirus pattern database.
- `pattern-local-update` — Update the Avira antivirus pattern database from local folder.
- `pattern-reload` — Reload the Avira antivirus pattern database.

- `pattern-update` — Update the Avira antivirus pattern database with the latest signatures.

Required Privilege Level

maintenance

Output Fields

request security utm anti-virus avira-engine pattern-delete

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security utm anti-virus avira-engine pattern-delete

```
user@host> request security utm anti-virus avira-engine pattern-delete
Anti-virus update request results:
Anti-virus update request results: Starting to delete avira files.
```

Sample Output

request security utm anti-virus avira-engine pattern-local-update <path>

```
user@host> request security utm anti-virus avira-engine pattern-local-update from </var/tmp/
db_0531>
Anti-virus update request results:
av_mgr: pattern updater 30445 is started, updating from /var/tmp/db_0531
```

Sample Output

request security utm anti-virus avira-engine pattern-reload

```
user@host> request security utm anti-virus avira-engine pattern-reload
Anti-virus update request results:
     Reloading good database starts ...
```

Sample Output

request security utm anti-virus avira-engine pattern-update

```
user@host> request security utm anti-virus avira-engine pattern-update
Anti-virus update request results:
av_mgr: pattern updater 44934 is started, downloading from https://update.juniper-updates.net/
avira.
```

Release Information

Command introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

clear security utm antivirus statistics | 4051

show security utm anti-virus statistics | 25183

show security utm anti-virus status | 25191

request security utm anti-virus juniper-express-engine

IN THIS SECTION

- Syntax | 14882
- Description | 14882
- Options | 14882
- Required Privilege Level | 14882
- Output Fields | 14882
- Sample Output | 14882
- Release Information | 14882

Syntax

```
request security utm anti-virus juniper-express-engine
```

Description

The Express Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, manually update the express antivirus pattern database using the command described. You can update the express antivirus pattern database automatically or manually. With full chassis cluster support for Content Security this command is operational on both the nodes.

Options

- `pattern-delete` — Delete the current express antivirus pattern database.
- `pattern-reload` — Reload the express antivirus pattern database.
- `pattern-update` — Update the express antivirus pattern database with the latest signatures.

Required Privilege Level

maintenance

Output Fields

request security utm anti-virus juniper-express-engine pattern-update

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security utm anti-virus juniper-express-engine pattern-update

```
user@host> request security utm anti-virus juniper-express-engine pattern-update
```

Release Information

The Express Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, command introduced in Junos OS Release 9.5.

Support for Content Security in chassis cluster added in Junos OS Release 11.4 .

RELATED DOCUMENTATION

clear security utm antivirus statistics | 4051
show security utm anti-virus statistics | 25183
show security utm anti-virus status | 25191

request security utm anti-virus kaspersky-lab-engine

IN THIS SECTION

- Syntax | 14883
- Description | 14883
- Options | 14884
- Required Privilege Level | 14884
- Output Fields | 14884
- Sample Output | 14884
- Release Information | 14884

Syntax

```
request security utm anti-virus kaspersky-lab-engine
```

Description

The Kaspersky Antivirus feature is not supported from Junos OS Release 15.1x49-D10 onwards.For previous releases, manually update the full file-based antivirus pattern database using the commands described. You can update the full file-based antivirus pattern database automatically or manually. With full chassis cluster support for Content Security this command is operational on both the nodes.

Options

- `pattern-delete` — Delete the current full file-based antivirus pattern database.
- `pattern-reload` — Reload the full file-based antivirus pattern database.
- `pattern-update` — Update the full file-based antivirus pattern database with the latest signatures.

Required Privilege Level

maintenance

Output Fields

request security utm anti-virus kaspersky-lab-engine pattern-update

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security utm anti-virus kaspersky-lab-engine pattern-update

```
user@host> request security anti-virus kaspersky-lab-engine pattern-update
```

Release Information

Command introduced in Junos OS Release 11.1 .

Support for Content Security in chassis cluster added in Junos OS Release 11.4 .

RELATED DOCUMENTATION

request security utm anti-virus juniper-express-engine | 14881

clear security utm antivirus statistics | 4051

show security utm anti-virus statistics | 25183

show security utm anti-virus status | 25191

request security utm anti-virus sophos-engine

IN THIS SECTION

- Syntax | 14885
- Description | 14885
- Options | 14885
- Required Privilege Level | 14886
- Output Fields | 14886
- Sample Output | 14886
- Release Information | 14886

Syntax

```
request security utm anti-virus sophos-engine
```

Description

Manually update the Sophos antivirus pattern database using the command described. To update automatically you use the configuration statement `set security utm feature-profile anti-virus sophos-engine pattern-update interval` *seconds*. With full chassis cluster support for Content Security this command is operational on both the nodes.

Starting in Junos OS Release 23.1R1, we've deprecated the `request security utm anti-virus sophos-engine` command.

Options

- `pattern-delete` — Delete the current Sophos antivirus pattern database.
- `pattern-reload` — Reload the Sophos antivirus pattern database.
- `pattern-update` — Update the Sophos antivirus pattern database with the latest signatures.

Required Privilege Level

maintenance

Output Fields

request security utm anti-virus sophos-engine pattern-update

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security utm anti-virus sophos-engine pattern-update

```
user@host> request security utm anti-virus sophos-engine pattern-update
```

Release Information

Command introduced in Junos OS Release 11.1.

Support for Content Security in chassis cluster added in Junos OS Release 11.4.

RELATED DOCUMENTATION

clear security utm antivirus statistics | 4051

show security utm anti-virus statistics | 25183

show security utm anti-virus status | 25191

request security utm custom-objects url-feed update

IN THIS SECTION

- Syntax | 14887
- Description | 14887

- Options | 14887
- Required Privilege Level | 14887
- Sample Output | 14887
- Release Information | 14888

Syntax

```
request security utm custom-objects url-feed update feed-name
```

Description

Manually update the URL feed. To update the URL feed automatically you use the configuration statement [`edit security utm custom-objects url-feed` *name* interval *interval*].

Options

url-feed *name*	Define the name of the URL feed object

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Sample Output

request security utm custom-objects url-feed update *feed-name*

```
user@host> request security utm custom-objects url-feed update temp
Update for feed temp initiated
Please use this command to check the status of url-feed :-
show security utm custom-objects url-feed status <feed-name>
```

Release Information

Statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

url-feed | 30489

request security utm custom-objects url-feed update feed-name force

IN THIS SECTION

- Syntax | 14888
- Description | 14888
- Options | 14889
- Required Privilege Level | 14889
- Sample Output | 14889
- Release Information | 14889

Syntax

```
request security utm custom-objects url-feed update feed-name force
```

Description

Forcefully update the URL feed. The forceful update enables the system to download the latest URL feed file and update the URL pattern. This update is irrespective of whether the file has any updates or not. To update the URL feed automatically you use the configuration statement [`edit security utm custom-objects url-feed` *`name`* interval *interval*].

Options

feed-name	Define the name of the URL feed object

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Sample Output

request security utm custom-objects url-feed update *feed-name* force

```
user@host> request security utm custom-objects url-feed update temp force
Update for feed temp initiated
Please use this command to check the status of url-feed :-
show security utm custom-objects url-feed status <feed-name>
```

Release Information

Statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

url-feed | 30489

request security utm web-filtering categorize

IN THIS SECTION

- Syntax | 14890
- Description | 14890
- Required Privilege Level | 14890

Sample Output | 14890
Release Information | 14890

Syntax

```
request security utm web-filtering categorize
<url>
<category-name>
```

NOTE: Specifying category-name is optional.

Description

Juniper NextGen Web filtering enables you to categorize an uncategorized URL. Using this feature, you can submit the categorization request.

Required Privilege Level

maintenance

Sample Output

request security utm web-filtering categorize

```
user@host> request security utm web-filtering categorize http://mail.abc.com category-name
NG_Web_based_Mail
```

```
NG categorize response: Your request is accepted
```

Release Information

Command introduced in Junos OS Release 23.4R1.

request security utm web-filtering category download-install [version]

IN THIS SECTION

- Syntax | 14891
- Description | 14891
- Required Privilege Level | 14891
- Sample Output | 14891
- Release Information | 14891

Syntax

```
request security utm web-filtering category download-install version;
```

Description

Download and install the category file, if no version is specified, the latest version is downloaded and installed during category upgrade.

Required Privilege Level

maintenance

Sample Output

request security utm web-filtering category download-install version 5

```
user@host> request security utm web-filtering category download-install version 5
Category updater result: Download scheduled
```

Release Information

Command introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

category (Security Web Filtering) | 3226

request security utm web-filtering category install | 14893

request security utm web-filtering category download [version] | 14892

request security utm web-filtering category download [version]

IN THIS SECTION

- Syntax | 14892
- Description | 14892
- Required Privilege Level | 14892
- Sample Output | 14893
- Release Information | 14893

Syntax

```
request security utm web-filtering category download version;
```

Description

Download the category file, if no version is specified, the latest version of the category file is downloaded during category upgrade.

Required Privilege Level

maintenance

Sample Output

request security utm web-filtering category download version 3

```
user@host> request security utm web-filtering category download version 3
Category updater result: Download done
```

Release Information

Command introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

- category (Security Web Filtering) | 3226
- request security utm web-filtering category install | 14893
- request security utm web-filtering category uninstall | 14897
- request security utm web-filtering category download-install [version] | 14891

request security utm web-filtering category install

IN THIS SECTION

- Syntax | 14893
- Description | 14894
- Required Privilege Level | 14894
- Sample Output | 14894
- Release Information | 14894

Syntax

```
request security utm web-filtering category install
```

Description

Install the predefined category and predefined filter on the system. Users could check the category or filter using the following command: `show security utm web-filtering category base-filter`.

> **NOTE**: During new category file installation, if the category filename is changed, then the new category file overwrites the old category file in the internal system and all related output information is replaced with the new category name.

Required Privilege Level

maintenance

Sample Output

request security utm web-filtering category install

```
user@host> request security utm web-filtering category install
Category updater result: install done
```

Release Information

Command introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

category (Security Web Filtering) | 3226

request security utm web-filtering category uninstall | 14897

request security utm web-filtering category migrate-to-ng-juniper

IN THIS SECTION

- Syntax | 14895
- Description | 14895
- Required Privilege Level | 14896
- Sample Output | 14896
- Release Information | 14897

Syntax

```
request security utm web-filtering category migrate-to-ng-juniper
```

Description

This command is processed in asynchronous mode. Use following command to check the category migration status:

```
show security utm web-filtering category migrate-to-ng-juniper status
```

Juniper NextGen Web filtering is a Web filtering solution which acts as a gateway for the SRX Series devices seeking URL reputation or category from the Juniper NextGen Web Filtering (NGWF) cloud. You can migrate from the existing Juniper Enhanced Web filtering solutions to the Juniper NextGen Web filtering solution based on your requirements.

> **NOTE**: If you change the policy names, you cannot commit the configuration during migration.

The Enhanced Web Filtering (EWF) and Juniper NextGen Web filtering require separate licenses. Starting in Junos OS Release 23.4R1, the SRX Series Firewall license installs have the NGWF license by default. Customers using the Enhanced Web Filtering (EWF) have an option to manually migrate from the existing EWF to the NGWF. You can also migrate to Juniper NGWF with your current Websense license.

Starting in Junos OS Release 23.4R1, during new installs or in case of an upgrade, the wf_key_ng_juniper key is installed. Juniper NextGen Web Filtering and URL category download and installation work when wf_key_websense_ewf or wf_key_ng_juniper is present.

Starting in Junos OS Release 23.4R1, you can use Juniper NextGen Web Filtering feature with wf_key_ng_juniper key.

> **NOTE**: The time required for the migration to complete from EWF to NGWF depends on the platforms and policies configured for migration. We recommend that you migrate to NGWF during downtime.

Required Privilege Level

view

Sample Output

request security utm web-filtering category migrate-to-ng-juniper

```
user@host> request security utm web-filtering category migrate-to-ng-juniper
```

This command is processed in asynchronous mode. Use following command to check the category migration status:

```
show security utm web-filtering category migrate-to-ng-juniper status
```

.

```
show security utm web-filtering category migrate-to-ng-juniper status Category migration status:
Status: Migration is In-progress

show security utm web-filtering category migrate-to-ng-juniper status Category migration status:
Status: Migration is Successful

Ignored Commands during migration:
Note:You must add the appropriate NG Juniper categories manually.
set security utm feature-profile web-filtering juniper-enhanced profile temp category
Enhanced_File_Download_Servers action block
```

```
set security utm feature-profile web-filtering juniper-enhanced profile temp category
Enhanced_Files_Containing_Passwords action block
set security utm feature-profile web-filtering juniper-enhanced profile temp category
Enhanced_Freeware_and_Software_Download action block
```

Release Information

Command introduced in Junos OS Release 23.4R1.

request security utm web-filtering category uninstall

IN THIS SECTION

- Syntax | 14897
- Description | 14897
- Required Privilege Level | 14897
- Sample Output | 14898
- Release Information | 14898

Syntax

```
request security utm web-filtering category uninstall
```

Description

Reset the predefined category and the base filters to the factory default. This option helps for category rollback.

Required Privilege Level

maintenance

Sample Output

request security utm web-filtering category uninstall

```
user@host> request security utm web-filtering category uninstall
Category updater result: Uninstall done
```

Release Information

Command introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

category (Security Web Filtering) | 3226

request security utm web-filtering category install | 14893

request security utm web-filtering custom-page reload

IN THIS SECTION

- Syntax | 14898
- Description | 14899
- Required Privilege Level | 14899
- Sample Output | 14899
- Release Information | 14899

Syntax

```
request security utm web-filtering custom-page reload
```

Description

You can this command to check and reload all custom page files in the configuration. It displays the similar error or warning messages when you commit the configuration.

Required Privilege Level

View

Sample Output

request security utm web-filtering custom-page reload

```
user@host> request security utm web-filtering custom-page reload
No such file cust_page1.html
The size of this page file (cust_page2.html) exceeds the limit, the maximum size is 2k bytes.
Custom page reload result: Reload Complete
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

custom-page-file | 4778

request security utm web-filtering recategorize

IN THIS SECTION

- Syntax | 14900
- Description | 14900
- Required Privilege Level | 14900
- Sample Output | 14900

Release Information | 14901

Syntax

```
request security utm web-filtering recategorize
<url>
<category-name>
<status>
```

NOTE: Specifying category-name is optional.

Description

The Juniper NextGen Web filtering enables you to re-categorize a URL. Using the re-categorization feature, you can submit the re-categorization request and view the status of your requests.

Required Privilege Level

maintenance

Sample Output

request security utm web-filtering recategorize

```
user@host> request security utm web-filtering recategorize http://mail.abc.com category-name
“Web based Mail"
```

```
Recategorize Accepted.
```

request security utm web-filtering recategorize

```
user@host> request security utm web-filtering recategorize http://mail.abc.com status
```

```
The request is being processed.
```

Release Information

Command introduced in Junos OS Release 23.4R1.

request server authenticate-peer-server

IN THIS SECTION

- Syntax | 14901
- Description | 14901
- Required Privilege Level | 14902
- Sample Output | 14902
- Release Information | 14902

Syntax

```
request server authenticate-peer-server
```

Description

Copy the **ssh** public key to the peer JDM. This command is equivalent to `ssh-copy-id user@jdm-server<0/1>`.

> **NOTE**: If the JDM fails to establish SSH connection with its peer on either of the two CB links, you need to run the JDM CLI command `request server authenticate-peer-server`. You can use the JDM CLI command `show server connections` to view the status of the SSH

connection between the JDM peers. Note that the command `request server authenticate-peer-server` will prompt for user confirmation twice - once per CB link.

Required Privilege Level

View

Sample Output

command-name

```
user@jdm> request server authenticate-peer-server
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter out anythat are
already installed
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted now it is to
install the new keys
user@jdm-server1's password:

Number of key(s) added: 1

Now try logging into the machine, with:   "ssh 'user@jdm-server1'"
and check to make sure that only the key(s) you wanted were added.
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

show virtual-network-functions (Junos Node Slicing) | **27728**

request virtual-network-functions | **15481**

request services advanced-anti-malware data-connection

IN THIS SECTION

- Syntax | 14903
- Description | 14903
- Options | 14903
- Required Privilege Level | 14903
- Output Fields | 14904
- Sample Output | 14904
- Release Information | 14905

Syntax

```
request services advanced-anti-malware data-connection test (start <0-32768> | status)
```

Description

Tests the connection between the SRX Series Firewall and the Juniper ATP Cloud by initiating a websocket connection and then sending data payloads of a given size. The SRX Series Firewall must already be enrolled with Juniper ATP Cloud before running this command.

Run this command when the `show services advanced-anti-malware statistics` CLI command shows that several files failed to be sent to the cloud (see the “File Send to Cloud Failed” result.)

Options

start <0-32768>	Start the data connection test and specify the packet payload size in bytes.
status	Returns the result of the data connection test. See Table 121 on page 14904.

Required Privilege Level

View

Output Fields

This CLI command returns a single line that indicates the data connection results. Table 121 on page 14904 lists the possible results.

Table 121: Data Connection Test Output

Message	Description
Test not started.	You cannot view the status without first running the data connection test. Run the `request services advanced-anti-malware data-connection test start` CLI command and then check the status again.
Test in progress.	The data connection test has not finished. Wait a few seconds and try the command again. Depending on your environment, it can take up to 20 seconds for the test to complete.
Test OK.	The data connection test passed.
Test failed.	The data connection test failed and indicates where it failed. Possible failures are: • Connect error—The websocket connection cannot be established. • Ping pong error—Successfully connected to the cloud server, but the payload delivery is not reliable.

Sample Output

request services advanced-anti-malware data-connection test start

```
user@host> request services advanced-anti-malware data-connection test start
Cloud connectivity test started. Ping payload size: 128 bytes.
```

request services advanced-anti-malware data-connection test status

```
user@host> request services advanced-anti-malware data-connection test status
  fpc0: Test OK. RTT = 38 ms. Test time: 2016-08-11 20:53:02 UTC.
```

request services advanced-anti-malware data-connection test status

```
user@host> request services advanced-anti-malware data-connection test status
  fpc0: Test failed. Reason: Ping pong error. Test time: 2016-08-11 21:13:05 UTC.
```

Release Information

Command introduced in Junos OS Release 15.1X49-D60.

request services advanced-anti-malware diagnostic

IN THIS SECTION

- Syntax | 14905
- Description | 14905
- Options | 14906
- Additional Information | 14907
- Required Privilege Level | 14908
- Sample Output | 14909
- Release Information | 14910

Syntax

```
request services advanced-anti-malware diagnostic url (detail | pre-detection url | routing-
instance instance-name)
```

Description

Use this command before you enroll your SRX Series Firewall with Juniper Advanced Threat Prevention Cloud to verify your Internet connection to the cloud. If you already enrolled your SRX Series Firewall, you can still use this command and the `request services aamw data-connection` CLI command to check and troubleshoot your connection to the cloud.

This CLI command checks the following:

- DNS lookup—Performs a forward DNS lookup of the cloud hostname to verify it returns an IP address. The examining process is terminated if it cannot get an interface name to the cloud. This issue may be caused by a connection error. Please check your network connection.
- Route to cloud—Tests your network connection using telnet.
- Whether server is live—Uses the telnet and ping commands to verify connection with the cloud.
- Outgoing interface—Checks that both the Routing Engine (RE) and the Packet Forwarding Engine (PFE) can connect to the Internet.
- IP path MTU—Determines the maximum transmission unit (MTU) size on the network path between the SRX Series Firewall and the cloud server. The examining process is terminated if the outgoing interface MTU is less than 1414. As a workaround, set the outgoing interface MTU to the default value or to a value greater than 1414.

  A warning message appears if the path MTU is less than the outgoing interface MTU. This is a minor issue and you can ignore the message. A higher path MTU is recommended but a low path MTU will work.
- SSL configuration consistency—Verifies that the SSL profile, client certificate and CA exists in both the RE and the PFE.
- Client and server clock check—When you run this CLI command, it first checks the difference between the server time and the local time. The time difference is expected to be less than one minute. If the time difference is more than one minute, an error message is displayed. See Table 122 on page 14907.

Options

url URL to the Juniper Advanced Threat Prevention Cloud server.

detail (optional) Debug mode that provides more verbose output.

pre-detection *url* (optional) Pre-detection mode where you can test your connection to the cloud server prior to actually enrolling your SRX Series Firewall.

To use this option, in the Web UI, click **Devices** and then click **Enroll**. You will receive an ops script similar to this:

```
op url https://abc.def.junipersecurity.net/bootstrap/enroll/AaBbCc/DdEeFf.slax
```

Use the root URL from the ops script as the url for the pre-detection option. For example, using the above ops script run the command as:

```
request services advanced-anti-malware diagnostic  pre-detection
abc.def.junipersecurity.net
```

routing-instance (optional) Routing instance used during enrollment. Specifying this option lets you diagnose the data plane connection to the Juniper ATP Cloud server with a customized routing instance. If you add `routing-instance ?` to the command line and press Enter, a list of known routing instances is displayed.

Additional Information

Table 122 on page 14907 lists the error conditions detected by this CLI command.

Table 122: aamw-diagnostics Script Error Messages

Error Message	Description
URL unreachable is detected, please make sure URL *url* port *port* is reachable.	Could not access the cloud server.
SSL profile *ssl profile name* is inconsistent between PFE and RE.	The SSL profile exists in the RE but does not exist in the PFE.
SSL profile *ssl profile name* is empty.	The SSL profile has neither trusted CA nor client certificate configured.
SSL local certificate *local certificate* is inconsistent between PFE and RE.	The SSL client certificate does not exist in PFE.
SSL CA *ca name* is inconsistent between PFE and RE.	The SSL CA exists in the RE but does not exist in the PFE.

Table 122: aamw-diagnostics Script Error Messages *(Continued)*

Error Message	Description
DNS lookup failure is detected, please check your DNS configuration.	The IP address of the cloud server could not be found. If this test fails, check to make sure your Internet connection is working properly and your DNS server is configured and has an entry for the cloud URL.
To-SKYATP connection through management interface is detected. Please make sure to-SKYATP connection is through packet forwarding plane.	The test detected that the Internet connection to the cloud server is through the management interface. This may result in your PFE connection to the cloud server failing. To correct this, change the Internet connection to the cloud to be through the PFE and not the management interface.
Unable to get server time.	Could not retrieve the server time.
Time difference is too large between server and this device.	The difference between the server time and the local SRX Series Firewall's time is more than a minute. To correct this, ensure that the clock on the local SRX Series Firewall is set correctly. Also, verify that you are using the correct NTP server.
Unable to perform IP path MTU check since ICMP service is down.	Unable to connect to the Juniper ATP Cloud server.
Required ICMP session not found.	Unable to establish an ICMP session with the specified URL. Check that you have specified a valid URL.

Required Privilege Level

View

Sample Output

request services advanced-anti-malware diagnostic

```
user@host> request services advanced-anti-malware diagnostic abc.def.junipersecurity.net

Time check                                          : [OK]
DNS check                                           : [OK]
SKYATP reachability check                           : [OK]
SKYATP ICMP service check                           : [OK]
Interface configuration check                       : [OK]
Outgoing interface MTU is default value
IP Path MTU check                                   : [OK]
IP Path MTU is 1472
SSL configuration consistent check                  : [OK]
```

request services advanced-anti-malware diagnostic detail

```
user@host> request services advanced-anti-malware diagnostic abc.def.junipersecurity.net detail

Time check                                          : [OK]
    [INFO]    Try to get IP address for hostname abc.def.junipersecurity.net
DNS check                                           : [OK]
    [INFO]    Try to test SKYATP server connectivity
SKYATP reachability check                           : [OK]
    [INFO]    Try ICMP service in SKYATP
SKYATP ICMP service check                           : [OK]
    [INFO]    To-SKYATP connection is using ge-0/0/3.0, according to route
Interface configuration check                       : [OK]
Outgoing interface MTU is default value
    [INFO]    Check IP MTU with length 1472
IP Path MTU check                                   : [OK]
IP Path MTU is 1472
SSL configuration consistent check                  : [OK]
```

request services advanced-anti-malware diagnostic pre-detection

```
user@host> request services advanced-anti-malware diagnostic pre-detection
abc.def.junipersecurity.net
Time check                                          : [OK]
DNS check                                           : [OK]
SKYATP reachability check                           : [OK]
SKYATP ICMP service check                           : [OK]
Interface configuration check                       : [OK]
Outgoing interface MTU is default value
IP Path MTU check                                   : [OK]
IP Path MTU is 1472
```

Release Information

Command introduced in Junos OS Release 15.1X49-D60. The interface name to cloud check, MTU warning, and client and server clock check added in Junos OS Release 15.1X49-D90. `routing-instance` option added in Junos OS Release 15.1X49-D100.

request services advanced-anti-malware redirect-file

IN THIS SECTION

- Syntax | 14911
- Description | 14911
- Options | 14911
- Additional Information | 14911
- Required Privilege Level | 14911
- Release Information | 14911

Syntax

```
request services advanced-anti-malware redirect-file
<add|delete ></var/tmp/custom_page.html>
```

Note that the file path is required when adding a file. When deleting a file, only the file name, no path, is required. For example: `request services advanced-anti-malware redirect-file delete custom_page.html`

Description

Allows you to add a customized file for users to be directed to. The file must be no larger than 8KB.

A default rich html file named “aamw_default_page.html” is included on the SRX Series Firewall for you to use. It says the following:

THIS SITE HAS BEEN BLOCKED

Sorry, the page you are looking for has been blocked due to the current security policy. If you believe this is an error or need to access this link, please contact your security department.

Also use the following command to view added files: `show services advanced-anti-malware redirect-files ?`

Options

add	Add a redirect file
delete	Delete a redirect file

Additional Information

See "advanced-anti-malware policy" on page 1997 for the `set services advanced-anti-malware policy p1 http client-notify file` command.

Required Privilege Level

View

Release Information

Command introduced in Junos OS Release 19.3R1.

request services anti-virus update

IN THIS SECTION

- Syntax | 14912
- Description | 14912
- Options | 14912
- Additional Information | 14912
- Required Privilege Level | 14913
- Output Fields | 14913
- Sample Output | 14913
- Release Information | 14914

Syntax

```
request services anti-virus update
request services anti-virus update status
request services anti-virus update-check-server
```

Description

Run this command to trigger antivirus database update immediately.

Options

status	Check the download status of antivirus database.
update-check-server	Check the connection status with the CDN server, from which you download the antivirus database package.

Additional Information

See "anti-virus (Services)" on page 2275 topic to configure automatic update of antivirus database package.

Required Privilege Level

View

Output Fields

This CLI command returns a single line that indicates the status of antivirus update. Table 123 on page 14913 lists the possible results.

Table 123: Antivirus Update Output Fields

Message	Description
No AV license detected or license is expired	Specifies that either Juniper AV license is not installed or the license has expired.
Anti-virus update is in progress	Specifies that the antivirus update is in progress.
Virus DB version	Specifies the antivirus database version.
Virus DB total count	Total number of antivirus database signatures at any given time.
Virus DB export date	The date and time when the antivirus signature database was last updated.
Server URL	URL of the CDN server.
Server status	Connection status with CDN server.

Sample Output

request services anti-virus update

```
request services anti-virus update

No AV license detected or license is expired
```

request services anti-virus update

```
request services anti-virus update

Anti-virus update is in progress
```

request services anti-virus update-check-server

```
request services anti-virus update-check-server
Anti-virus server:
  Virus DB version:       1671134540
  Virus DB total count:   31256
  Virus DB export date:   Thu Dec 15 20:02:20 2022 UTC
  Server URL:             https://signatures.juniper.net/phase
  Server status:          Connection success
```

Release Information

Command introduced in Junos OS Release 23.4R1.

request services anti-virus machine-learning-engine

IN THIS SECTION

- Syntax | 14915
- Description | 14915
- Options | 14915
- Required Privilege Level | 14915
- Sample Output | 14915
- Release Information | 14916

Syntax

```
request services anti-virus machine-learning-engine start
request services anti-virus machine-learning-engine stop
request services anti-virus machine-learning-engine status
```

Description

Run these commands to enable and disable machine learning (ML)-based zero-day threat detection.

Options

Table 124: Commands

start	Enable ML-based threat detection.
stop	Disable ML-based threat detection.
status	Check the status of ML-based threat detection.

Required Privilege Level

View

Sample Output

`request services anti-virus machine-learning-engine start`

```
request services anti-virus machine-learning-engine start
```

`request services anti-virus machine-learning-engine status`

```
request services anti-virus machine-learning-engine status
Active (enabled)
```

request services anti-virus machine-learning-engine stop

```
request services anti-virus machine-learning-engine stop
```

request services anti-virus machine-learning-engine status

```
request services anti-virus machine-learning-engine status
Inactive (disabled)
```

Release Information

Command introduced in Junos OS Release 24.2R1.

request services application-identification application

IN THIS SECTION

- Syntax | 14916
- Description | 14917
- Options | 14917
- Required Privilege Level | 14917
- Output Fields | 14917
- Sample Output | 14917
- Release Information | 14918

Syntax

request services application-identification application [disable | enable] *predefined-application-name*

Description

Disable, or enable a predefined application signature.

Options

predefined-application-name—Application name; a maximum of up to 31 characters. Predefined applications have the prefix junos- to avoid conflict with user-defined ones. Do not name your custom application signature with the junos prefix; this prefix is reserved for predefined application signatures.

disable—(Optional) Disable a predefined application signature, initiate signature recompilation, and commit all pending uncompiled signatures to the configuration.

The following conditions apply:

- You cannot disable a predefined application signature that is referenced by an active security policy or custom application signature. First modify or deactivate the policy or custom application signature.
- If you disable an application signature, for example, junos:HTTP, that has nested applications, the nested applications are not recognized.

enable—(Optional) Enable a predefined application signature, initiate signature recompilation, and commit all pending uncompiled signatures to the configuration.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

request services application-identification application disable

```
user@host> request services application-identification application disable junos:163

 Please wait while we are updating signatures ...
Please wait while we are updating signatures ...
Please wait while we are updating signatures ...
```

```
Please wait while we are updating signatures ...
Please wait while we are updating signatures ...
Please wait while we are updating signatures ...
 Disable application junos:163 succeed.
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show services application-identification application | 25389

request services application-identification clear packet-capture all

IN THIS SECTION

- Syntax | 14918
- Description | 14918
- Required Privilege Level | 14919
- Output Fields | 14919
- Sample Output | 14919
- Release Information | 14919

Syntax

```
request services application-identification clear packet-capture all
```

Description

Clear all packet capture files generated for unknown application traffic.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

command-name

```
user@host> request services application-identification clear packet-capture all
Packet-capture file(s) deleted
```

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Configure Packet Capture For Unknown Application Traffic

packet-capture | 12137

show services application-identification packet-capture counters | 25443

clear services application-identification packet-capture counters | 4081

request services application-identification download

IN THIS SECTION

- Syntax | 14920
- Description | 14920
- Options | 14920

- Required Privilege Level | 14921
- Output Fields | 14921
- Sample Output | 14921
- Release Information | 14923

Syntax

```
request services application-identification download
<check-server>
<minor-only>
<status>
<version>
```

Description

Manually download the application signature package for Junos OS application identification.

Juniper Networks regularly updates the predefined application signature package database and makes it available to subscribers on the Juniper Networks website.

When you enter the command, the system retrieves the application package from the Juniper Networks security website https://signatures.juniper.net/cgi-bin/index.cgi and downloads on the device.

Options

- `check-server` —Display the details of the download server URL. The output also displays the major and minor versions of application package in the output. Do not run the check-server option after download or install commands are executed.
- `minor-only` —Download the minor type of application package. If you do not specify this option, the firewall downloads the major (regular) package.
- `status`—Check the download status of the application signature package.
- *`version`*—(Optional) Download a specific version of the application package from the Juniper Networks security website. If you do not enter a version, the most recent version is downloaded.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services application-identification download

```
user@host> request services application-identifications download
Please use command "request services application-identification download status"
 to check status
```

request services application-identifications download minor-only

```
user@host> request services application-identifications download minor-only
Please use command "request services application-identification download status"
 to check status
```

request services application-identification download status

```
user@host> request services application-identifications download status
Application package 1608 is downloaded successfully.
```

request services application-identification download status

```
user@host> request services application-identifications download status
Downloading application package 3670 (Minor) succeeded
```

request services application-identification download status

```
user@host> request services application-identifications download status
Downloading application package (latest) failed with error (No suitable version available for
this device,
please re-try the download manually without minor
```

request services application-identification download check-server

```
user@host> request services application-identifications download check-server
Download server URL: https://signatures.juniper.net/cgi-bin/index.cgi
Sigpack Version: 3377
Protobundle version: 1.500.2-31
Build Time: Apr 05 2021 13:57:06
```

request services application-identification check-server

```
user@host> request services application-identifications download check-server
Download server URL: https://signatures.juniper.net/cgi-bin/index.cgi
Sigpack Version: 3666 (Major)
Protobundle version: 1.550.2-43
Build Time: Apr 05 2023 06:28:09
Sigpack Version: 3659 (Minor)
Protobundle version: 1.550.2-43
Build Time: Apr 05 2023 06:28:09
```

NOTE: We recommend not to use the `request services application-identification download check-server` command when running the following commands:

`request services application-identification download`

`request services application-identification install`

Following messages are displayed in such cases:

```
user@host> request services application-identifications download check-
server
Currently "download" command is running, please try again later.

user@host> request services application-identifications download check-
server
Currently "install" command is running, please try again later.
```

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

- request services application-identification download status | 14923
- request services application-identification install | 14934

request services application-identification download status

IN THIS SECTION

- Syntax | 14924
- Description | 14924
- Required Privilege Level | 14924
- Output Fields | 14924
- Sample Output | 14924
- Release Information | 14925

Syntax

```
request services application-identification download status
```

Description

Check the download status of the application signature package.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request. The output displays the application package version.

Sample Output

request services application-identification download status

```
user@host> request services application-identifications download status
Application package 1608 is downloaded successfully.
```

request services application-identification download status

```
user@host> request services application-identifications download status
Application package 1608 is downloaded successfully.
```

request services application-identification download status

```
user@host> request services application-identifications download status
Downloading application package 3670 (Minor) succeeded
```

request services application-identification download status

```
user@host> request services application-identifications download status
Downloading application package (latest) failed with error (No suitable version available for
this device, please re-try the download manually without minor keyword)
```

request services application-identification download status

```
user@host> request services application-identifications download status
Downloading application package (50250) failed with error (Application package version 50250 is
already installed)
```

Release Information

Statement introduced in Junos OS Release 10.2.

request services application-identification group

IN THIS SECTION

- Syntax | 14926
- Description | 14926
- Options | 14926
- Required Privilege Level | 14926
- Output Fields | 14927
- Sample Output | 14927
- Release Information | 14927

Syntax

```
request services application-identification group [copy | disable | enable] predefined-application-group-name
```

Description

Copy, disable, or enable a predefined application signature group.

Options

copy (Optional) Copy a predefined application signature group from the database to the configuration and change the name (for example, my:FTP). The ID and order are generated automatically. Do not name your custom application signature group with the `junos` prefix; this prefix is reserved for predefined application signature groups. You can copy the same predefined application signature group only once; duplicate custom signature groups are not allowed.

> **NOTE**: In configuration mode, if an uncommitted action is pending, the `request services application-identification group copy` command fails.

disable (Optional) Disable a predefined application signature group.

> **NOTE**: You cannot disable a predefined application signature group that is referenced by an active security policy or custom application signature group. First modify or deactivate the policy or custom application signature group.

enable (Optional) Enable a predefined application signature group.

predefined-application-group-name Name of the predefined application signature group.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

request services application-identification group

```
user@host> request services application-identification group disable
junos:infrastructure:networking
Disable application group junos:infrastructure:networking succeed.
```

request services application-identification group

```
user@host> request services application-identification group enable
junos:infrastructure:networking
Enable application group junos:infrastructure:networking succeed.
```

request services application-identification group

```
user@host> request services application-identification group copy
junos:infrastructure:networking
Please wait while we are copying group ...
Copy application group junos:infrastructure:networking succeed.
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show services application-identification group | 25438

request services application identification install ignore duplicate version check

IN THIS SECTION

- Syntax | 14928
- Description | 14928
- Required Privilege Level | 14928
- Sample Output | 14928
- Release Information | 14929

Syntax

```
request services application-identification install ignore-duplicate-version-check
```

Description

Forcefully installs the application signature package over the same version of the signature package.

Required Privilege Level

maintenance

Sample Output

When you enter this command, the system provides feedback on the status of your request.

request services application-identification install ignore-duplicate-version-check

```
user@host> request services application-identification install ignore-duplicate-version-check
```

```
Please use command
                "request services application-identification install status" to check install
status
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Predefined Application Signatures for Application Identification

request services application-identification install status

IN THIS SECTION

- Syntax | 14929
- Description | 14930
- Required Privilege Level | 14930
- Output Fields | 14930
- Sample Output | 14930
- Release Information | 14931

Syntax

```
request services application-identification install status
```

Description

Display the status of the install operation.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

request services application-identification install status

```
user@host> request services application-identification install status
Installed Application package (3702) and Protocol bundle successfully
```

The system displays following output when the installed application signature package is detected as faulty (Junos OS 24.2R1):

```
user@host> request services application-identification install status
Install Application package 3656 and Protocol bundle failed ( Install Application package (3656)
failed in data plane validation, auto rollback triggered)
```

request services application-identification install status (For Chassis Cluster) (Junos OS 24.2R1)

```
user@host> request services application-identification install status
node0:
--------------------------------------------------------------------------
Checking compatibility of application package version 3577 ...
node1:
--------------------------------------------------------------------------
Waiting for primary node to finish installation and validation of the application package ...
```

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

request services application-identification install | 14934

request services application identification new to production

IN THIS SECTION

- Syntax | 14931
- Description | 14931
- Options | 14932
- Required Privilege Level | 14932
- Sample Output | 14932
- Release Information | 14932

Syntax

```
request services application-identification new-to-production [applications-list
application-name | all]]
```

Description

Migrate the applications tagged as new in the installed application signature pack on your security device to normal applications.

The applications tagged as new are part of `junos:all-new-apps` group. Once you migrate the new applications to normal applications, these applications will no longer be tagged as new and will not be part of the `junos:all-new-apps` group.

Options

applications-list	Specify list of applications to migrate to normal applications.
all	Migrate all new applications as normal applications

Required Privilege Level

view

Sample Output

request services application-identification new-to -production application-list

```
user@host> request services application-identification new-to-production applications-list
[junos:RLOGIN junos:LINKEDIN]
```

```
Please wait while we are updating signatures ...
new-to-production command successfully executed
```

request services application-identification new-to-production all

```
user@host> request services application-identification new-to-production applications-list
[junos:RLOGIN junos:LINKEDIN]
```

```
Please wait while we are updating signatures ...
new-to-production command successfully executed
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Configure Packet Capture For Unknown Application Traffic

request services application-identification rollback status

IN THIS SECTION

- Syntax | 14933
- Description | 14933
- Required Privilege Level | 14933
- Output Fields | 14933
- Sample Output | 14934
- Release Information | 14934

Syntax

```
request services application-identification rollback status
```

Description

Displays the status of the application signature package rollback.

You can use this command when you manually rollback the application signature package to the previously installed version using the **request services application-identification rollback** command.

Required Privilege Level

maintenance

Output Fields

This command provides the application signature package rollback status.

Sample Output

request services application-identification rollback status

```
user@host> request services application-identification rollback status
```

Sample: When rollback is successful.

```
Application package (3250) and Protocol bundle successful.
```

Sample: When rollback is failure.

```
Application package (3250) and Protocol bundle failed.
```

Release Information

Command introduced in Junos OS Release 20.3R1.

request services application-identification install

IN THIS SECTION

- Syntax | 14935
- Description | 14935
- Required Privilege Level | 14935
- Output Fields | 14935
- Sample Output | 14935
- Release Information | 14935

Syntax

```
request services application-identification install
```

Description

Install the downloaded predefined application signature package.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

command-name

```
user@host> request services application-identification install
Please use command "request services application-identification install status" to check status
and use command "request services application-identification proto-bundle-status" to check
protocol bundle status
```

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

request services application-identification install status | 14929

request services application-identification download | 14919

request services application-identification offline-download package-path

IN THIS SECTION

- Syntax | 14936
- Description | 14936
- Required Privilege Level | 14936
- Output Fields | 14936
- Sample Output | 14937
- Release Information | 14937

Syntax

```
request services application-identification offline-download package-path <path>
```

Description

Extract the application signature package from a downloaded local tar file.

When you enter this command, the system untars the signature package and places the extracted files in the proper locations on the device.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services application-identification offline-download package-path

```
user@host> request services application-identification offline-download package-path /root/appid-offline-update.tar-2.gz

Please use command
"request services application-identification offline-download status" to check offline download status
```

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

request services application-identification download status

request services application-identification install

request services application-identification proto-bundle-status

IN THIS SECTION

- Syntax | 14938
- Description | 14938
- Required Privilege Level | 14938
- Output Fields | 14938
- Sample Output | 14938
- Release Information | 14938

Syntax

```
request services application-identification proto-bundle-status
```

Description

Display the status of the install operation of the protocol bundle.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

command-name

```
user@host> request services application-identification proto-bundle-status
Protocol Bundle Version (1.30.4-22.005 (build date Jan 17 2014)) and application secpack version
(2345) is loaded and activated.
```

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

request services application-identification install | 14934

request services application-identification uninstall

IN THIS SECTION

- Syntax | 14939
- Description | 14939
- Required Privilege Level | 14939
- Output Fields | 14939
- Sample Output | 14940
- Release Information | 14940

Syntax

```
request services application-identification uninstall
```

Description

Uninstall the predefined application package.

The uninstall operation will fail if any active security policies reference predefined application signatures or predefined application signature groups in the Junos OS configuration.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

command-name

```
user@host> request services application-identification uninstall
Please use command "request services application-identification uninstall status" to check
status and use command "request services application-identification proto-bundle-status" to
check protocol bundle status
```

Release Information

Statement introduced in Junos OS Release 10.2. Statement modified in Junos OS Release 10.4. Statement modified in Junos OS Release 11.4.

RELATED DOCUMENTATION

request services application-identification install | 14934

request services application-identification uninstall status

IN THIS SECTION

- Syntax | 14941
- Description | 14941
- Required Privilege Level | 14941
- Output Fields | 14941
- Sample Output | 14941
- Release Information | 14941

Syntax

```
request services application-identification uninstall status
```

Description

Display the status of the uninstall operation.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

command-name

```
user@host> request services application-identification uninstall status
Uninstall application package version (1776) succeed.
```

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

request services application-identification uninstall | 14939

request services extensible-subscriber-services reload-dictionary

IN THIS SECTION

- Syntax | 14942
- Description | 14942
- Options | 14942
- Required Privilege Level | 14942
- Sample Output | 14942
- Release Information | 14943

Syntax

```
request services extensible-subscriber-services reload-dictionary
```

Description

Reload the configured dictionary to essmd.

Options

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Sample Output

command-name

```
# request services extensible-subscriber-services reload-dictionary
Dictionary reloaded successfully
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

- dictionary | 5703
- show extensible-subscriber-services dictionary | 19853
- show extensible-subscriber-services dictionary attributes | 19859
- show extensible-subscriber-services dictionary services | 19863
- *Understanding the Dictionary File*

request services flow-collector change-destination primary interface

IN THIS SECTION

- Syntax | 14943
- Description | 14944
- Options | 14944
- Required Privilege Level | 14944
- Output Fields | 14944
- Sample Output | 14944
- Release Information | 14944

Syntax

```
request services flow-collector change-destination primary interface cp-fpc/pic/port
<clear-files>
<clear-logs>
<immediately | gracefully>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Switch to the primary File Transfer Protocol (FTP) server that is configured as a flow collector.

Options

none	Switch to the primary FTP server.
cp-*fpc/pic/port*	Use the specified flow collector interface name for the primary destination.
clear-files	(Optional) Request clearing of existing data files in the FTP wait queue when the switch takes place.
clear-logs	(Optional) Request clearing of existing logs when the switch takes place.
immediately | gracefully	(Optional) Specify whether you want the switch to take place immediately, or to affect only newly created files.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services flow-collector change-destination primary interface

```
user@host> request services flow-collector change-destination primary interface
cp-6/0/0
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Destination change successful
```

Release Information

Command introduced before Junos OS Release 7.4.

request services flow-collector change-destination secondary interface

IN THIS SECTION

- Syntax | 14945
- Description | 14945
- Options | 14945
- Required Privilege Level | 14946
- Output Fields | 14946
- Sample Output | 14946
- Release Information | 14946

Syntax

```
request services flow-collector change-destination secondary interface cp-fpc/pic/
port
<clear-files>
<clear-logs>
<immediately | gracefully>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Switch to the secondary File Transfer Protocol (FTP) server that is configured as a flow collector.

Options

none	Switch to the secondary FTP server.
cp-*fpc/pic/port*	Use the specified flow collector interface name (**cp-*fpc/pic/port***) for the secondary destination.
clear-files	(Optional) Request clearing of existing data files in the FTP wait queue when the switch takes place.

clear-logs	(Optional) Request clearing of existing logs when the switch takes place.
immediately | gracefully	(Optional) Specify whether you want the switch to take place immediately, or to affect only newly created files.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services flow-collector change-destination secondary interface

```
user@host> request services flow-collector change-destination secondary interface
cp-6/0/0
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Destination change successful
```

Release Information

Command introduced before Junos OS Release 7.4.

request services flow-collector test-file-transfer

IN THIS SECTION

- Syntax | 14947
- Description | 14947
- Options | 14947

- Required Privilege Level | 14947
- Output Fields | 14947
- Sample Output | 14948
- Release Information | 14948

Syntax

```
request services flow-collector test-file-transfer filename  interface (all | cp-fpc/pic/port)
(channel-zero | channel-one) (primary | secondary)
```

Description

(M40e, M160, and M320 Series routers, PTX Series, and T Series routers only) Transfer a test file to the primary or secondary File Transfer Protocol (FTP) server that is configured as a flow collector. This command verifies that the output side of the flow collector interface is operating properly.

Options

filename	Name of the test file to transfer.
interface (all | cp-*fpc/pic/port*)	Transfer a test file of flows from all configured flow collector interfaces or from only the specified interface.
channel-zero | channel-one	Transfer a file from export channel 0 (unit 0) or channel 1 (unit 1) of the PIC.
primary | secondary	Transfer a file to the primary or secondary server configured as a flow collector.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services flow-collector test-file-transfer interface channel-one primary

```
user@host> request services flow-collector test-file-transfer test_file interface cp-7/1/0
channel-one primary


Flow collector interface: cp-7/1/0
Interface state: Collecting flows
Response: Test file transfer successfully scheduled
```

Release Information

Command introduced before Junos OS Release 7.4.

request services ip-monitoring preempt-restore policy

IN THIS SECTION

- Syntax | 14948
- Description | 14949
- Options | 14949
- Required Privilege Level | 14949
- Output Fields | 14949
- Sample Output | 14949
- Release Information | 14949

Syntax

```
request services ip-monitoring preempt-restore policy
<policy-name>
```

Description

If the no–preempt option is specified, the policy will not perform preemptive failback when it is in a failover state, and when the RPM probe test recovers from failure. To manually revert to the failback state, run the `request services ip-monitoring preempt-restore policy` command.

> **NOTE**: The `request services ip-monitoring preempt-restore policy` command takes effect only when the RPM probe is in the pass state, and when the policy is in a failover state.

Options

policy name	Name of the policy.

Required Privilege Level

maintenance

Output Fields

When you run this command, the policy is restored to the failback state.

Sample Output

run request services ip-monitoring preempt-restore policy <policy name>

```
user@host> run request services ip-monitoring preempt-restore policy policy1
Restore request succeeded: Policy policy1
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show services ip-monitoring status | 25604

request services ipsec-vpn ipsec switch tunnel

IN THIS SECTION

- Syntax | 14950
- Description | 14950
- Options | 14950
- Required Privilege Level | 14950
- Output Fields | 14951
- Sample Output | 14951
- Release Information | 14951

Syntax

```
request services ipsec-vpn ipsec switch tunnel local-gateway address   remote-gateway
address
<routing-instance instance-name>
```

Description

(Adaptive services interface only) Manually switch between primary and backup IP Security (IPsec) tunnels.

Options

local-gateway *address*	Gateway address of the local system.
remote-gateway *address*	Gateway address of the remote system.
routing-instance *instance-name*	(Optional) VRF instance associated with local gateway address.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services ipsec-vpn ipsec switch tunnel

```
user@host> request services ipsec-vpn ipsec switch tunnel local-gateway 10.1.1.1 remote gateway
10.100.10.1
```

Release Information

Command introduced before Junos OS Release 7.4.

routing-instance option added in Release 8.1.

RELATED DOCUMENTATION

show services ipsec-vpn ipsec security-associations | 25622

request services l2tp destination unlock

IN THIS SECTION

- Syntax | 14952
- Description | 14952
- Options | 14952
- Required Privilege Level | 14952
- Output Fields | 14952
- Sample Output | 14952
- Release Information | 14952

Syntax

```
request services l2tp destination unlock destination-name
```

Description

Remove the specified destination from the destination lockout list immediately, before its lockout period expires. As a result, the L2TP process can again consider the destination during the selection of new tunnels.

Options

destination-name	Name of the destination to be unlocked.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services l2tp destination unlock

```
user@host> request services l2tp destination unlock dest-a
Destination dest-a unlocked
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

L2TP Tunnel Switching For Multiple-Domain Networks

show services l2tp destination lockout | **25647**

request services monitoring twamp client

IN THIS SECTION

- Syntax | **14953**
- Description | **14953**
- Options | **14954**
- Required Privilege Level | **14954**
- Release Information | **14954**

Syntax

```
request services monitoring twamp client (start | stop)
<control-connection-name>
repeat <number>
```

Description

Start or stop a Two-Way Active Measurement Protocol (TWAMP) session. You can start or stop all of the sessions for all of the TWAMP clients, or start or stop a session for a specific TWAMP client control connection. When you start all the test sessions configured for a particular TWAMP client, the control-client starts all requested testing with a Start-Sessions message, and the server sends an acknowledgment. If the control connection is not active between the server and the client, the control connection is also established and the test connections are started later. If the control-client name is not specified, all the configured test sessions are begun.

When you stop the test session, the control connection is closed only after the Stop-sessions message is sent from the TWAMP client to the TWAMP server. If the control-client name is not specified, all the configured test sessions are closed.

Options

control-connection-name	(Optional) Start or stop the TWAMP session with the server for only the specified control-connection or TWAMP control-client. If you do not specify this option, all control connections are either started or stopped.

repeat *number*

Required Privilege Level

view

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

control-connection (Junos OS Evolved) | **4627**

Understanding Two-Way Active Measurement Protocol on Routers

request services paa install

IN THIS SECTION

- Syntax | **14955**
- Description | **14955**
- Options | **14955**
- Additional Information | **14956**
- Required Privilege Level | **14957**
- Sample Output | **14957**
- Release Information | **14958**

Syntax

```
request services paa install
cc-account account
cc-host host
cc-user user@domain
cc-password password
ta-version version
ta-name name
cc-port port
<cc-insecure>
<ta-debug>
```

Description

Install the Paragon Active Assurance (PAA) test agent. Junos OS Evolved Release 22.3R1 supports installing test agents for PAA 4.0 and later. This command fetches the PAA test agent software image from the PAA Control Center for you and installs the software into a Docker container.

Before issuing this command, make sure you have an HTTPS connection to the PAA Control Center (port 443 is open) and either port 6800 or a port of your choosing is open to connect the test agent with the PAA Control Center. In the PAA Control Center, make sure your SSL certificate is correct on both ports, and both are either secure or self-signed. Also, configure 127.0.0.1 as the loopback address (lo0.0) on the router. For more information about loopback addresses and how to configure them, see Loopback Interfaces (Junos OS Evolved).

Options

cc-account *account*	(Required) Your PAA Control Center user account.
cc-host *host*	(Required) IP address or DNS host name of the PAA Control Center.
cc-user *user@domain*	(Required) The email address registered with the PAA Control Center.
cc-password *password*	(Required) Your PAA Control Center password. The password is passed as-is from the router to the PAA Control Center. Special characters must be escaped using the standard escaping scheme for the Linux shell.
ta-version *version*	(Optional) The PAA test agent software version you would like to install. Check the PAA Control Center for the versions available. If not specified, the operating system installs the latest PAA test agent software version.

ta-name *name*

(Optional) The name you want to use to refer to this PAA test agent. When upgrading the test agent, make sure to use the existing `ta-name`. The default is the router's hostname.

The `ta-name` is used to identify the test agent in the PAA Control Center, and PAA won't allow two test agents with the same name to run at the same time.

cc-port *port*

(Optional) Port over which the router can communicate with the PAA Control Center. The default is port 6800. Make sure the port you specify here (or port 6800 if taking the default) is open.

cc-insecure

(Optional) Allow an insecure connection between the test agent and the PAA Control Center. In other words, disable SSL certificate checking and use self-signed SSL certificates.

ta-debug

(Optional) Enable debug log files for test agent operation. Specify this option if you want to see some debug-level log messages if something goes wrong with the installation. This option also allows you to collect debug-level messages about test agent operation after a successful install. To see these logs, use the `docker logs` command at the shell prompt:

```
user@host> start shell
user@host:~$ export DOCKER_HOST=unix:///run/docker-vrf0.sock
user@host:~$ docker logs paa_agent
```

Additional Information

You can also use NETCONF to install the PAA test agent:

```
<rpc>
   <install-paa-ta>
     <cc-host>host</cc-host>
     <cc-account>account</cc-account>
     <cc-user>user@domain</cc-user>
     <cc-password>password</cc-password>
     <ta-version>version</ta-version>
     <cc-port>port</cc-port>
     <ta-name>name</ta-name>
     <cc-insecure></cc-insecure>
     <ta-debug></ta-debug>
```

```
    </install-paa-ta>
  </rpc>
  ]]>]]>
```

Required Privilege Level

maintenance

Sample Output

request services paa install (first-time installation)

```
user@host> request services paa install cc-account my_account cc-host host cc-user user@domain
cc-password $1234abc ta-version 4.0.0.29 ta-name my_test_agent cc-port 6800
```

```
PAA installation files copied from 10.83.153.119.
Starting docker daemon.
Created symlink /etc/systemd/system/extensions.target.wants/docker@vrf0.service -> /lib/systemd/
system/docker@.service.
Starting PAA test agent process.
Loaded image: paa/test-agent-application:4.0.0.29
Setting environment.
459d83560855faa6bae16873d3753344f252cb5cd860f790228cf53d5e0ff046
Done. Starting the test agent with environment file /var/opt/paa.env
```

request services paa install (subsequent upgrade)

```
user@host> request services paa install cc-account my_account cc-host host cc-user user@domain
cc-password $1234abc ta-version 4.0.0.36 ta-name my_test_agent cc-port 6800
```

```
PAA installation files copied from 10.83.153.119.
Starting docker daemon.
Starting PAA test agent process.
Loaded image: paa/test-agent-application:4.0.0.36
```

```
Setting environment.
A0c12feaddb312fd2fe3625a659304a448e9eeac4767d2eccd7749bc6f24e8ca
Done. Starting the test agent with environment file /var/opt/paa.env
```

Release Information

Command introduced in Junos OS Evolved Release 22.3R1.

Command deprecated in Junos OS Evolved Release 23.4R1. Replaced by the `test-agent` configuration statement at the [edit services paa] hierarchy level.

RELATED DOCUMENTATION

Install the Paragon Active Assurance (PAA) Test Agent

request services paa uninstall | 14958

show services paa status | 25893

request services paa uninstall

IN THIS SECTION

- Syntax | 14958
- Description | 14959
- Additional Information | 14959
- Required Privilege Level | 14959
- Sample Output | 14959
- Release Information | 14959

Syntax

```
request services paa uninstall
```

Description

Uninstall the Paragon Active Assurance (PAA) test agent. Before installing a different version of the test agent software, you must uninstall the previous test agent software.

Additional Information

You can also use NETCONF to uninstall the PAA test agent:

```
<rpc>
  <uninstall-paa-ta>
  </uninstall-paa-ta>
</rpc>
]]>]]>
```

Required Privilege Level

maintenance

Sample Output

request services paa uninstall

```
user@host> request services paa uninstall
```

```
Stopping PAA test agent.
Done. Un-installation of PAA test agent.
```

Release Information

Command introduced in Junos OS Evolved Release 22.3R1.

Command deprecated in Junos OS Evolved Releases 23.4R1. Replaced by deleting the `test-agent` configuration statement at the [edit services paa] hierarchy level.

RELATED DOCUMENTATION

Install the Paragon Active Assurance (PAA) Test Agent

request services paa install | 14954

request services redundancy-set trigger

IN THIS SECTION

- Syntax | 14960
- Description | 14960
- Required Privilege Level | 14960
- Release Information | 14960

Syntax

```
request services redundancy-set (redundancy-set |all) trigger redundancy-event
```

Description

Manually trigger a redundancy event for the specified redundancy set or all sets on the current gateway.

Required Privilege Level

view

Release Information

Statement introduced in Junos OS Release 16.1.

request services remote-device-management reconfigure service-device

IN THIS SECTION

- Syntax | 14961
- Description | 14961
- Options | 14961
- Required Privilege Level | 14961
- Output Fields | 14961
- Sample Output | 14962
- Release Information | 14962

Syntax

```
request services remote-device-management reconfigure service-device device-name
```

Description

Reconfigure a remote device to provision all active subscriber services matching the access domain (list of VLAN ranges and IDs) configured for this remote device with the `vlan-id-list` option at the [[`edit system services remote-device-management service-device` *`device-name`*`access-domain`] hierarchy level.

Options

device-name System-wide name of the remote device that uniquely identifies the device across routing instances.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services remote-device-management reconfigure service-device (Successful)

```
user@host> request services remote-device-management reconfigure service-device OLT1
succeeded
```

request services remote-device-management reconfigure service-device (Failed)

```
user@host> request services remote-device-management reconfigure service-device OLT1
Service device reconfiguration failed
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Reconfiguring a Remote Device for RDSM

request services remote-device-management reload-dictionary | **14962**

Remote Device Services Manager (RDSM) Overview

request services remote-device-management reload-dictionary

IN THIS SECTION

- Syntax | **14963**
- Description | **14963**
- Options | **14963**
- Required Privilege Level | **14963**
- Output Fields | **14963**
- Sample Output | **14964**

Release Information | 14964

Syntax

```
request services remote-device-management reload-dictionary absolute file path
```

Description

Reload the specified dictionary to the RDSM database. The dictionary is configured with the `dictionary` option at the `[edit system services remote-device-management service-device device-name]` hierarchy level. The reload affects all remote service devices that are configured with this dictionary. When you modify an existing dictionary, this is how you apply the updated file.

NOTE: You cannot modify a dictionary when there is at least one active subscriber service configured on a remote device using that dictionary.

Options

absolute file path Absolute file path for the vendor-specific dictionary that defines the set of NETCONF XML protocol commands required to provision, deprovision, and roll back a subscriber service on the remote device. The dictionary is stored on the BNG. An example absolute path is **/var/home/dict/remote-device.xml**.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services remote-device-management reload-dictionary (Successful)

```
user@host> request services remote-device-management reload-dictionary /var/home/dict/vendor-1-
dictionary.xml
Dictionary reloaded successfully
```

request services remote-device-management reload-dictionary (Failed)

```
user@host> request services remote-device-management reload-dictionary /var/home/dict/vendor-2-
dictionary.xml
Dictionary reload failed
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Reloading a Dictionary File for RDSM

request services remote-device-management reconfigure service-device | **14961**

Remote Device Services Manager (RDSM) Overview

request services rpm twamp

IN THIS SECTION

- Syntax | **14965**
- Description | **14965**
- Options | **14965**
- Required Privilege Level | **14965**

- Output Fields | 14966
- Sample Output | 14966
- Release Information | 14966

Syntax

```
request services rpm twamp  (start | stop) client <control-connection-name>
```

Description

Start or stop a TWAMP session. You can start or stop all of the sessions for all of the TWAMP clients, or start or stop a session for a specific TWAMP client. When you start all the test session configured for a particular TWAMP client, the control-client initiates all requested testing with a Start-Sessions message, and the server sends an acknowledgment. If the control connection is not active between the server and the client, the control connection is also established and the test connections are started later. If the control-client name is not specified, all the configured test sessions are commenced.

When you stop the test session, the control connection is closed only after the Stop-sessions message is sent from the TWAMP client to the TWAMP server. If the control-client name is not specified, all the configured test sessions are closed.

Options

start client	Start the TWAMP session between the TWAMP client and the TWAMP server.
stop client	Terminate the TWAMP session between the TWAMP client and the TWAMP server.
control-connection-name	(Optional) Start or stop the TWAMP session with the server only for the specified control-connection or TWAMP control-client.

Required Privilege Level

view

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

request services rpm twamp start client

```
user@host> request services rpm twamp start client c1
```

Release Information

Command introduced in Junos OS Release 15.1.

request services static-subscribers login group

IN THIS SECTION

- Syntax | 14966
- Description | 14967
- Options | 14967
- Required Privilege Level | 14967
- Output Fields | 14967
- Sample Output | 14967
- Release Information | 14967

Syntax

```
request services static-subscribers login group group-name
```

Description

Resets the state of an interface group on which static subscribers were forcibly logged out by the `request services static-subscribers logout group` command. This action enables static subscriber to login on the interfaces in the group.

Options

group ***group-name*** Group of static subscriber interfaces on which static subscribers have been created.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services static-subscribers login group

```
user@host> request services static-subscribers login group boston
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

> *Monitoring and Managing Static Subscriber Information*
>
> request services static-subscribers logout group | **14969**

request services static-subscribers login interface

IN THIS SECTION

- Syntax | 14968
- Description | 14968
- Options | 14968
- Required Privilege Level | 14968
- Output Fields | 14968
- Sample Output | 14969
- Release Information | 14969

Syntax

```
request services static-subscribers login interface interface-name
```

Description

Resets the state of an interface on which a static subscriber was forcibly logged out by the `request services static-subscribers logout interface` command. This action enables a static subscriber to login on the interface.

Options

interface *interface-name* Static interface on which a static subscriber has been created.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services static-subscribers login interface

```
user@host> request services static-subscribers login interface ge-2/0/1.5
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Monitoring and Managing Static Subscriber Information

request services static-subscribers logout interface | 14971

request services static-subscribers logout group

IN THIS SECTION

- Syntax | 14969
- Description | 14970
- Options | 14970
- Required Privilege Level | 14970
- Output Fields | 14970
- Sample Output | 14970
- Release Information | 14970

Syntax

```
request services static-subscribers logout group igroup-name
```

Description

Force static subscribers on the interfaces in the group to be logged out. No subscriber can subsequently log in on the interface group until the interface state is reset by a router reset or the `request services static-subscribers login group` command.

Options

group ***group-name*** Group of static subscriber interfaces on which static subscribers have been created.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services static-subscribers logout group

```
user@host> request services static-subscribers logout group boston
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Monitoring and Managing Static Subscriber Information

request services static-subscribers login group | **14966**

request services static-subscribers logout interface

IN THIS SECTION

- Syntax | 14971
- Description | 14971
- Options | 14971
- Required Privilege Level | 14971
- Output Fields | 14971
- Sample Output | 14972
- Release Information | 14972

Syntax

```
request services static-subscribers logout interface interface-name
```

Description

Force static subscriber on the interface to be logged out. No subscriber can subsequently log in on the interface until the interface state is reset by a router reset or the `request services static-subscribers login interface` command.

Options

interface *interface-name* Static interface on which a static subscriber has been created.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services static-subscribers logout interface

```
user@host> request services static-subscribers logout interface ge-2/0/1.5
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Monitoring and Managing Static Subscriber Information

request services static-subscribers login interface | 14968

request services subscribers

IN THIS SECTION

- Syntax | 14972
- Description | 14973
- Options | 14973
- Required Privilege Level | 14973
- Sample Output | 14973
- Release Information | 14973

Syntax

```
request services subscribers set subscriber-profile profile client-id client-id
```

Description

Set the subscriber profile associated with the given subscriber.

Options

profile Name of the subscriber profile to create or override the currently active subscriber profile for the given subscriber.

client-id Client session ID assigned to the subscriber.

Required Privilege Level

view

Sample Output

request services subscriber set subscriber-profile tc_act_prof client-id

```
user@host> request services subscriber set subscriber-profile tc_act_prof client-id
2533274790395909 | display xml
rpc-reply xmlns:junos="http://xml.juniper.net/junos/11.1I0/junos"
    packet-triggered-subscribers-information xmlns="http://xml.juniper.net/junos/11.1I0/junos-
packet-triggered-subscribers"
        service-subscribers-request-result junos:style="success"
        /service-subscribers-request-result
    /packet-triggered-subscribers-information
    cli
        banner/banner
    /cli
/rpc-reply
```

Release Information

Command introduced in Junos OS Release 11.4.

request services subscribers clear

IN THIS SECTION

- Syntax | 14974
- Description | 14974
- Options | 14974
- Required Privilege Level | 14974
- Sample Output | 14975
- Release Information | 14975

Syntax

```
request services subscribers clear subscriber-profile profile client-id client-id
```

Description

Clear the subscriber profile associated with the given subscriber.

Options

profile Name of the subscriber profile to clear the active subscriber profile for the given subscriber.

client-id Client session ID assigned to the subscriber.

Required Privilege Level

clear

Sample Output

request services subscriber clear subscriber-profile tc_act_prof client-id

```
user@host>request services subscriber clear subscriber-profile tc_act_prof client-id
2533274790395909 | display xml
rpc-reply xmlns:junos="http://xml.juniper.net/junos/11.1I0/junos"
    packet-triggered-subscribers-information xmlns="http://xml.juniper.net/junos/11.1I0/junos-
packet-triggered-subscribers"
        service-subscribers-request-result junos:style="success"
        /service-subscribers-request-result
    /packet-triggered-subscribers-information
    cli
        banner/banner
    /cli
/rpc-reply
```

Release Information

Command introduced in Junos OS Release 11.4.

request services url-filter delete gencfg-data

IN THIS SECTION

- Syntax | 14975
- Description | 14976
- Required Privilege Level | 14976
- Release Information | 14976

Syntax

```
request services url-filter delete gencfg-data
```

Description

Delete url-filterd–ased objects. This is historical information.

WARNING: Do not use this command unless explicitly instructed to do so.

NOTE: Starting in Junos OS Release 18.3R1, the `request services url-filter delete gencfg-data` command is deprecated and has been replaced by the `request services web-filter delete gencfg-data` command.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

- request services url-filter force dns-resolution | **14976**
- request services url-filter update url-filter-database file | **14979**
- request services url-filter validate | **14980**
- *Configuring URL Filtering*

request services url-filter force dns-resolution

IN THIS SECTION

- Syntax | **14977**
- Description | **14977**
- Options | **14977**
- Required Privilege Level | **14977**

Sample Output | 14978
Release Information | 14978

Syntax

```
request services url-filter force dns-resolution (all | profile profile-name)
<template template-name>
```

Description

Force the domain name system (DNS) resolution request.

NOTE: Starting in Junos OS Release 18.3R1, the `request services url-filter force dns-resolution` command is deprecated and has been replaced by the `request services web-filter force dns-resolution` command.

Options

all	Force a DNS resolution request for all profiles.
profile *profile-name*	Force a DNS resolution request for the profile (all templates).
template *template-name*	Force a DNS resolution request for the template.

Required Privilege Level

view

Sample Output

request services url-filter force dns-resolution profile profile1

```
user@host> request services url-filter force dns-resolution profile profile1
STATUS: DNS resolution is triggered for profile profile1
```

request services url-filter force dns-resolution profile profile1 template template1

```
user@host> request services url-filter force dns-resolution profile profile1 template template1
STATUS: DNS resolution is triggered for template template1 in profile profile1
```

request services url-filter force dns-resolution all

```
user@host> request services url-filter force dns-resolution all
STATUS: DNS resolution has been triggered for all profiles
```

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

- request services url-filter delete gencfg-data | 14975
- request services url-filter update url-filter-database file | 14979
- request services url-filter validate | 14980
- *Configuring URL Filtering*

request services url-filter update url-filter-database file

IN THIS SECTION

- Syntax | 14979
- Description | 14979
- Options | 14979
- Required Privilege Level | 14979
- Release Information | 14980

Syntax

```
request services url-filter update url-filter-database file filename
```

Description

Update the URL filter database. If you change the database file, run this command. It sends a request to the DNS server for only the new hostnames that were not in the previous version of URL database file.

NOTE: Starting in Junos OS Release 18.3R1, the `request services url-filter update url-filter-database file` command is deprecated and has been replaced by the `request services web-filter update url-filter-database file` command.

Options

filename	Use the database filename.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

- request services url-filter delete gencfg-data | 14975
- request services url-filter force dns-resolution | 14976
- request services url-filter validate | 14980
- *Configuring URL Filtering*

request services url-filter validate

IN THIS SECTION

- Syntax | 14980
- Description | 14980
- Options | 14981
- Required Privilege Level | 14981
- Release Information | 14981

Syntax

```
request services url-filter validate (all | file-name filename)
```

Description

Validate the URL filter database to ensure the database is correct.

> **NOTE**: Starting in Junos OS Release 18.3R1, the `request services url-filter validate` command is deprecated and has been replaced by the `request services web-filter validate url-filter-file-name` command.

Options

`all`	Validate the URL filter database for all profiles.
`file-name` *`filename`*	Validate the URL filter database for the database file specified.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

- request services url-filter delete gencfg-data | **14975**
- request services url-filter force dns-resolution | **14976**
- request services url-filter update url-filter-database file | **14979**
- *Configuring URL Filtering*

request services user-identification active-directory-access active-directory-authentication-table delete

IN THIS SECTION

- Syntax | **14982**
- Description | **14982**

- Options | 14982
- Required Privilege Level | 14982
- Output Fields | 14983
- Release Information | 14983

Syntax

```
request services user-identification active-directory-access active-directory-authentication-
table delete
(domain name | ip-address ip-address | group group-name <domain name> | user name <domain name>
```

Description

Delete entries from the active directory authentication table by domain, address, group, or user. This command provides the network administrator with flexibility and control over the table entries beyond what is automatically added to or deleted from the table. For example, if a person leaves the company, the corresponding username can be deleted; after a department reorganization, a group can be deleted.

Options

- `domain` *`name`*—Delete the entries from the authentication table for the specified domain.
- `ip-address` *`ip-address`*—Delete the entry from the authentication table for the specified IP address.
- `group` *`group-name`*—Delete the entries from the authentication table for the specified group.
  - `domain` *`name`*—Delete the group only from the specified domain.
- `user` *`name`*—Delete the entries from the authentication table for the specified username.
  - `domain` *`name`*—Delete the user only from the specified domain.

Required Privilege Level

maintenance

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

show services user-identification active-directory-access active-directory-authentication-table

user-identification (Services) | 30560

Understanding Active Directory Authentication Tables

request services user-identification active-directory-access domain-controller

IN THIS SECTION

- Syntax | 14983
- Description | 14984
- Options | 14984
- Required Privilege Level | 14984
- Output Fields | 14984
- Sample Output | 14984
- Release Information | 14984

Syntax

```
request services user-identification active-directory-access domain-controller discovery
domain name
```

Description

Discover and display the name and address of all domain controllers in the specified domain.

Options

- `domain` *name*—Name of the domain for which to get and display domain controller names and addresses.

Required Privilege Level

maintenance

Output Fields

This command displays the discovered domain controllers.

Sample Output

request services user-identification active-directory-access domain-controller discovery domain <domain-name>

```
user@host> request services user-identification active-directory-access domain-controller
discovery domain example.net
Domain: example.net
 Domain controller: example-dc.example.net
   Address: 192.0.2.2
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

active-directory-access | 1802

show services user-identification active-directory-access domain-controller status | 26385

user-identification (Services) | 30560

request services user-identification active-directory-access ip-user-probe

IN THIS SECTION

- Syntax | 14985
- Description | 14985
- Options | 14985
- Required Privilege Level | 14985
- Output Fields | 14986
- Sample Output | 14986
- Release Information | 14986

Syntax

```
request services user-identification active-directory-access ip-user-probe
address ip-address <domain name>
```

Description

Probe the PC at the specified IP address to get an authentication entry, which is used for the integrated user firewall feature. You can display the authentication table to see the results. If the probe succeeded, there will be a valid authentication entry. If the probe failed, there will be an invalid authentication entry.

Options

- `address` *`ip-address`*—Probe the PC at this IP address.

  `domain` *`name`*—Probe the IP address in the specified domain.

Required Privilege Level

maintenance

Output Fields

The following command displays the results of the IP address probe:

Sample Output

show services user-identification active-directory-access active-directory-authentication-table address <ip-address>

```
user@host> show services user-identification active-directory-access active-directory-
authentication-table address 192.0.2.3
Domain: example.net
 Source-ip: 192.0.2.3
  Username: user1
  Groups:group1
  State: Valid
  Source: wmic
  Access start date: 2014-03-10
  Access start time: 13:59:56
  Age time: 1437
user@host> show services user-identification active-directory-access active-directory-
authentication-table address 2001:db8::1:1
Domain: example.net
 Source-ip: 2001:db8::1:1
  Username: user2
  Groups:group1
  State: Valid
  Source: wmic
  Access start date: 2017-03-10
  Access start time: 13:59:56
  Age time: 1437
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

clear services user-identification active-directory-access | 4194

show services user-identification active-directory-access active-directory-authentication-table
show services user-identification active-directory-access statistics | 26389
user-identification (Services) | 30560

request services user-identification authentication-source aruba-clearpass user-query

IN THIS SECTION

- Syntax | 14987
- Description | 14987
- Options | 14988
- Required Privilege Level | 14988
- Sample Output | 14988
- Release Information | 14988

Syntax

```
request services user-identification authentication-source authentication-source user-query
address  ip-address
```

Description

Manually send to the ClearPass website a request for user authentication and identity information for an individual user. The command specifies the IP address of the user's device to identify the user whose information you want to obtain. If the user query command executes successfully, an entry for the user (IP address) has been created in the ClearPass authentication table, and no output is displayed.

The user query function is part of the SRX Series integrated ClearPass authentication and enforcement feature. The user query function, if configured, allows the SRX Series Firewall to send requests for individual user information. This command also allows you to manually send requests. Normally administrators send query requests manually to troubleshoot issues.

The user query function supplements use of the Web API function. The SRX Series Firewall exposes to ClearPass a Web API that ClearPass uses to send POST request messages to the SRX Series Firewall. These messages contain user authentication and identity information.

Options

ip-address The IP address of the user's device for whom you are manually requesting authentication information.

> **NOTE**: Starting with Junos OS Release 15.1X49-D130, SRX Series Firewall can query ClearPass for IPv6 addresses, in addition to IPv4 addresses for an individual user.

Required Privilege Level

maintenance

Sample Output

request services user-identification authentication-source authentication-source user-query address ip-address

```
user@host> request services user-identification authentication-source aruba-clearpass user-query
address 40.0.0.1
user@host> request services user-identification authentication-source aruba-clearpass user-query
address 2001:db8:4136:e378:8000:63bf:3fff:fdd2
```

Release Information

Command introduced in Junos OS Release 12.3X48-D30.

request services user-identification authentication-source jims groups domain <domain-name> (force-fetch|status)

IN THIS SECTION

- Syntax | 14989
- Description | 14989
- Options | 14989
- Required Privilege Level | 14990
- Sample Output | 14990
- Release Information | 14990

Syntax

```
request services user-identification authentication-source jims groups domain <domain-name>
request services user-identification authentication-source jims groups domain <domain-name>
status
request services user-identification authentication-source jims groups domain <domain-name>
force-fetch
```

Description

You can request JIMS server to retrieve the group list for Active Directory (AD) domain using root logic-system. There are no new requests triggered if the group query for the domain is ongoing.

Options

force-fetch This option is used to send out the group query even if you retrieve the group-list. The force-fetch does not take effect even if the group query is already sent out and that does not retrieve result.

status This option displays the status of the received message. You can verify that the Juniper Identity Management Service server is online and provides primary or secondary server is responding to queries from the SRX Series device.

Required Privilege Level

view

Sample Output

request services user-identification authentication-source jims groups domain <domain-name> (force-fetch|status)

```
user@host> request services user-identification authentication-source jims groups domain <domain-name>

node0:
--------------------------------------------------------------------------

user@host> request services user-identification authentication-source jims groups domain <domain-name> status
node0:
--------------------------------------------------------------------------
            Server-type                 : JIMS
            Logical system name         : root-logical-system
            Domain                      : domain-name
            Groups                      : finished:(48 groups are processed)
            Received status message     : OK(200)
            Update time                 : 2020-04-20 04:57:03

user@host> request services user-identification authentication-source jims groups domain <domain-name> force-fetch
node0:
--------------------------------------------------------------------------
```

Release Information

Command introduced in Junos OS Release 20.2R1.

request services user-identification authentication-source jims validate (user <user-name>|group <group-name>|device <device-name>) domain <domain-name>

IN THIS SECTION

- Syntax | 14991
- Description | 14991
- Options | 14992
- Required Privilege Level | 14992
- Output Fields | 14992
- Sample Output | 14992
- Release Information | 14994

Syntax

```
request services user-identification authentication-source jims validate (user <user-name>|group
<group-name>|device <device-name>) domain <domain-name>
```

Description

Displays the list of status of some user or group or device information.

- When the client-id or client secret is error, the command shows the unauthorized error in the received status message.
- When the server is offline, the command shows couldn't connect to server error in the received status message.
- When the HTTP request timeouts for different reasons, the command shows timeout reached in the received status message.
- If there are any other communication error, then there is no validation carried out on a status message.

Options

user-name	User ID.
group-name	Name of the group.
device-name	Name of the connected device.
domain-name	Name of the domain.

Required Privilege Level

view

Output Fields

The following examples cover how to validate the valid user, device, and group for JIMS.

Sample Output

request services user-identification authentication-source jims validate (user <user-name>| group <group-name>|device <device-name>) domain <domain-name>

The following commands are used to request JIMS validator for validating a specified user, device, and group.

```
user@host> request services user-identification authentication-source jims validate user <user-name> domain <domain-name>

node0:
--------------------------------------------------------------------------

user@host> request services user-identification authentication-source jims validate user <user-name> domain <domain-name> status
node0:
--------------------------------------------------------------------------
            Server-type                 : JIMS
            Logical system name         : root-logical-system
            Domain                      : domain-name
            Nametype                    : user
            Namestring                  : user-name
```

```
            Validate result             : valid
            Received status message     : OK(200)
            Update time                 : 2020-04-19 23:57:37

user@host>  request services user-identification authentication-source jims validate user <user-name> domain <domain-name> force-fetch
node0:
--------------------------------------------------------------------------



user@host> request services user-identification authentication-source jims validate group <group-name> domain <domain-name>
node0:
--------------------------------------------------------------------------



user@host> request services user-identification authentication-source jims validate group <group-name> domain <domain-name> status
node0:
--------------------------------------------------------------------------
            Server-type                 : JIMS
            Logical system name         : root-logical-system
            Domain                      : domain-name
            Nametype                    : group
            Namestring                  : group-name
            Validate result             : valid
            Received status message     : OK(200)
            Update time                 : 2020-04-08 21:19:35

user@host> request services user-identification authentication-source jims validate group <group-name> domain <domain-name> force-fetch
node0:
--------------------------------------------------------------------------

user@host> request services user-identification authentication-source jims validate device <device-name>$ domain <domain-name>
```

You must add an $ after device name or computer name.

```
node0:
--------------------------------------------------------------------------
```

```
user@host> request services user-identification authentication-source jims validate device
<device-name>$ domain <domain-name> status
node0:
--------------------------------------------------------------------------
            Server-type                 : JIMS
            Logical system name         : root-logical-system
            Domain                      : domain-name
            Nametype                    : device
            Namestring                  : device-name$
            Validate result             : failed
            Received status message     : Forbidden(403)
            Update time                 : 2020-04-24 01:34:18

user@host> request services user-identification authentication-source jims validate device
<device-name>$ domain <domain-name> force-fetch
node0:
--------------------------------------------------------------------------
```

Release Information

Command introduced in Junos OS Release 20.2R1.

request services user-identification authentication-table delete

IN THIS SECTION

- Syntax | 14995
- Description | 14995
- Options | 14995
- Required Privilege Level | 14996
- Output Fields | 14996
- Sample Output | 14996
- Release Information | 15004

Syntax

```
request services user-identification authentication-table delete (ip-address  ip-address |
authentication-source (all | active-directory | authentication-source (domain domain-name |
group group-name /user user-name) )
```

Description

Delete entries from the ClearPass authentication table based on the IP address of the user's device, or on the authentication source and the name of a domain, a group, or a user. When only the authentication source is specified, the entire ClearPass authentication table is deleted. For the integrated ClearPass authentication and enforcement feature, the authentication source is always aruba-clearpass.

Options

ip-address Deletes a user authentication entry from the ClearPass authentication table, and the Active Directory (AD) table, based on the IP address of the user's device.

> **NOTE**: Starting with Junos OS Release 15.1X49-D130, SRX Series device supports to delete IPv6 addresses if IPv6 addresses were configured.

authentication-source Deletes user entries from the ClearPass authentication table. In the CLI, ClearPass as the authentication source is referred to by the value aruba-clearpass as is the ClearPass authentication table. To identify the user entries to be deleted, you specify a domain, a group, or a username.

- ***domain-name*** Deletes from the ClearPass authentication table user entries for users who belong to the specified domain.
- **group *group-name*** Deletes the entry entry from the ClearPass authentication table for users who belong to the group, regardless of whether they belong to other groups.
- **user *user-name*** Deletes the entry for the specified user from the ClearPass authentication table.

Required Privilege Level

maintenance

Output Fields

The following examples cover how to delete various user entries from the ClearPass authentication table based on the specified parameter. It also shows how to check to ensure that the user entries were deleted successfully.

Sample Output

request services user-identification authentication-table delete ip-address

```
The following command deletes the entry for the user whose device IP address is specified.
user@host> request services user-identification authentication-table delete ip-address 50.0.0.1
user@host> request services user-identification authentication-table delete ip-address
2001:db8:4136:e378:8000:63bf:3fff:fdd2
```

Before you delete the entry:

To ensure that the entry exists in the ClearPass authentication table, use the following command to display the entry for the user. Note that the ClearPass authentication table includes the user entry with the IP address 50.0.0.1 and 2001:db8:4136:e378:8000:63bf:3fff:fdd2.

```
user@host> show services user-identification authentication-table ip-address 50.0.0.1
Domain: GLOBAL
  Source-ip: 50.0.0.1
    Username: guest1
    Groups:posture-healthy, guest, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2015-12-14
    Access start time: 17:07:23
    Last updated timestamp: 2015-12-22 05:50:47
    Age time: 0

user@host> show services user-identification authentication-table ip-address
2001:db8:4136:e378:8000:63bf:3fff:fdd2
Domain: GLOBAL
  Source-ip: 2001:db8:4136:e378:8000:63bf:3fff:fdd2
```

```
    Username: guest2
    Groups:posture-healthy1, guest, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2015-12-14
    Access start time: 17:07:23
    Last updated timestamp: 2015-12-22 05:50:47
    Age time: 0
```

After you delete the user entry associated with the IP address, enter the command again to verify that the entry has been deleted.

```
user@host> show services user-identification authentication-table ip-address 50.0.0.1
warning: “This IP address isn’t in authentication table.”
user@host> show services user-identification authentication-table ip-address
2001:db8:4136:e378:8000:63bf:3fff:fdd2
warning: “This IP address isn’t in authentication table.”
```

request services user-identification authentication-table delete authentication-source aruba-clearpass domain

```
The following command deletes the specified domain.
user@host> request services user-identification authentication-table delete authentication-
source domain global
```

Before you delete the domain contents from the ClearPass authentication table, use the following command to display the domain information to ensure that it exists. Note that the ClearPass authentication table includes the global domain.

```
user@host> show services user-identification authentication-table authentication-source aruba-
clearpass domain global extensive
Domain: GLOBAL
Total entries: 6
  Source-ip: 10.0.0.1
    Username: viki2
    Groups:posture-healthy, accounting-grp, accounting-grp-and-company-device,
    corporate-limited, [user authenticated]
    Groups referenced by policy:accounting-grp-and-company-device
    State: Valid
    Source: Aruba ClearPass
```

```
    Access start date: 2016-03-08
    Access start time: 17:20:30
    Last updated timestamp: 2015-12-22 04:02:48
    Age time: 0
  Source-ip: 20.0.0.1
    Username: abew1
    Groups:posture-unknown, marketing-access-limited-grp, [user authenticated]
    Groups referenced by policy:marketing-access-limited-grp
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:31:40
    Last updated timestamp: 2015-12-22 04:18:48
    Age time: 0
  Source-ip: 30.0.0.1
    Username: jxchan
    Groups:posture-healthy, marketing-access-for-pcs-limited-group,
    marketing-general, sales-limited, corporate-limited, [user authenticated]
    Groups referenced by policy:marketing-access-for-pcs-limited-group
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:22:48
    Last updated timestamp: 2015-12-22 05:46:21
    Age time: 0
  Source-ip: 40.0.0.1
    Username: lchen1
    Groups:posture-healthy, human-resources-grp, accounting-limited,
    corporate-limited, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:21:37
    Last updated timestamp: 2015-12-22 05:41:18
    Age time: 0
  Source-ip: 50.0.0.1
    Username: guest1
    Groups:posture-healthy, guest, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:23:10
    Last updated timestamp: 2015-12-22 05:50:47
```

```
    Age time: 0
  Source-ip: 50.0.0.2
    Username: guest2
    Groups:posture-healthy, guest-device-byod, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:23:21
    Last updated timestamp: 2015-12-22 05:52:44
    Age time: 0
```

After you delete the domain, use the command again to verify that the domain and its user members was deleted.

```
user@host> show services user-identification authentication-table authentication-source aruba-
clearpass domain global
warning: "There is no related auth entry in authentication-table."
```

request services user-identification authentication-table delete authentication-source aruba-clearpass group

The following command deletes the entries for any users who belong to the group posture-healthy.

```
user@host> request services user-identification authentication-table delete authentication-
source aruba-clearpass group posture-healthy
```

Before you delete the group contents from the ClearPass authentication table, use the following command to display it to ensure that the group is used in some user entries. Notice that the appropriate user entries contain the posture-healthy group.

```
Domain: GLOBAL
Total entries: 6
  Source-ip: 10.0.0.1
    Username: viki2
    Groups:posture-healthy, accounting-grp, accounting-grp-and-company-device,
    corporate-limited, [user authenticated]
    Groups referenced by policy:accounting-grp-and-company-device
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
```

```
    Access start time: 17:20:30
    Last updated timestamp: 2015-12-22 04:02:48
    Age time: 0
  Source-ip: 20.0.0.1
    Username: abew1
    Groups:posture-unknown, marketing-access-limited-grp, [user authenticated]
    Groups referenced by policy:marketing-access-limited-grp
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:31:40
    Last updated timestamp: 2015-12-22 04:18:48
    Age time: 0
  Source-ip: 30.0.0.1
    Username: jxchan
    Groups:posture-healthy, marketing-access-for-pcs-limited-group,
    marketing-general, sales-limited, corporate-limited, [user authenticated]
    Groups referenced by policy:marketing-access-for-pcs-limited-group
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:22:48
    Last updated timestamp: 2015-12-22 05:46:21
    Age time: 0
  Source-ip: 40.0.0.1
    Username: lchen1
    Groups:posture-healthy, human-resources-grp, accounting-limited,
    corporate-limited, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:21:37
    Last updated timestamp: 2015-12-22 05:41:18
    Age time: 0
  Source-ip: 50.0.0.1
    Username: guest1
    Groups:posture-healthy, guest, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:23:10
    Last updated timestamp: 2015-12-22 05:50:47
    Age time: 0
```

```
  Source-ip: 50.0.0.2
    Username: guest2
    Groups:posture-healthy, guest-device-byod, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:23:21
    Last updated timestamp: 2015-12-22 05:52:44
    Age time: 0
```

Enter the **show services user-identification authentication-table authentication-source aruba-clearpass group posture-healthy** to display the entries for the users who belong to the group posture-healthy.

Notice that the group name does not show up in the column for groups referenced by policy because it is not one. Notice, too, that the output contains information for only those users who belong to the group. It does not include an entry for the user abew1, who does not belong to the group.

```
Domain: GLOBAL
Source IP       Username       groups(Ref by policy)          state
10.0.0.1        viki2          accounting-grp-and-company-dev Valid
30.0.0.1        jxchan         marketing-access-for-pcs-limit Valid
40.0.0.1        lchen1         corporate-limited              Valid
50.0.0.1        guest1                                        Valid
50.0.0.2        guest2                                        Valid
```

After you delete the group, use the command again to verify that it has been deleted.

```
user@host> show services user-identification authentication-table authentication-source aruba-
clearpass group posture-healthy
warning: "There is no related auth entry in authentication-table."
```

For further verification, you can use the following command to check the entry for one of the users who belonged to the group:

```
user@host> show services user-identification authentication-table authentication-source aruba-
clearpass user viki2
warning: "There is no related auth entry in authentication-table."
```

request services user-identification authentication-table delete authentication-source aruba-clearpass

The following command deletes the ClearPass authentication table (aruba-clearpass).

```
user@host> request services user-identification authentication-table delete authentication-source aruba-clearpass
```

Before you delete the ClearPass authentication table, use the following command to display it to ensure that the table exists.

```
user@host> show services user-identification authentication-table authentication-source aruba-clearpass
Domain: GLOBAL
Total entries: 6
  Source-ip: 10.0.0.1
    Username: viki2
    Groups:posture-healthy, accounting-grp, accounting-grp-and-company-device,
    corporate-limited, [user authenticated]
    Groups referenced by policy:accounting-grp-and-company-device
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:20:30
    Last updated timestamp: 2015-12-22 04:02:48
    Age time: 0
  Source-ip: 20.0.0.1
    Username: abew1
    Groups:posture-unknown, marketing-access-limited-grp, [user authenticated]
    Groups referenced by policy:marketing-access-limited-grp
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:31:40
    Last updated timestamp: 2015-12-22 04:18:48
    Age time: 0
  Source-ip: 30.0.0.1
    Username: jxchan
    Groups:posture-healthy, marketing-access-for-pcs-limited-group,
    marketing-general, sales-limited, corporate-limited, [user authenticated]
    Groups referenced by policy:marketing-access-for-pcs-limited-group
```

```
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:22:48
    Last updated timestamp: 2015-12-22 05:46:21
    Age time: 0
  Source-ip: 40.0.0.1
    Username: lchen1
    Groups:posture-healthy, human-resources-grp, accounting-limited,
    corporate-limited, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:21:37
    Last updated timestamp: 2015-12-22 05:41:18
    Age time: 0
  Source-ip: 50.0.0.1
    Username: guest1
    Groups:posture-healthy, guest, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:23:10
    Last updated timestamp: 2015-12-22 05:50:47
    Age time: 0
  Source-ip: 50.0.0.2
    Username: guest2
    Groups:posture-healthy, guest-device-byod, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:23:21
    Last updated timestamp: 2015-12-22 05:52:44
    Age time: 0
```

To verify that you deleted the authentication table successfully, enter the command again:

```
user@host> show services user-identification authentication-table authentication-source aruba-
clearpass

warning: "There is no authentication-table entry."
```

Release Information

Command introduced in Junos OS Release 12.3X48-D30.

request services web-filter delete gencfg-data

IN THIS SECTION

- Syntax | 15004
- Description | 15004
- Required Privilege Level | 15004
- Release Information | 15004

Syntax

```
request services web-filter delete gencfg-data
```

Description

Delete url-filterd–ased objects. This is historical information.

WARNING: Do not use this command unless explicitly instructed to do so.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

request services web-filter force dns-resolution | 15005

request services web-filter update url-filter-database file | 15007

request services web-filter validate url-filter-file-name | 15010

Configuring URL Filtering

request services web-filter force dns-resolution

IN THIS SECTION

- Syntax | 15005
- Description | 15005
- Options | 15005
- Required Privilege Level | 15006
- Release Information | 15006

Syntax

```
request services web-filter force dns-resolution (all | profile profile-name)
<template template-name>
```

Description

Force the domain name system (DNS) resolution request.

Options

all	Force a DNS resolution request for all profiles.
profile *profile-name*	Force a DNS resolution request for the profile (all templates).
template *template-name*	Force a DNS resolution request for the template.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

- request services web-filter delete gencfg-data | 15004
- request services web-filter update url-filter-database file | 15007
- request services web-filter validate url-filter-file-name | 15010
- *Configuring URL Filtering*

request services web-filter update dns-filter-database

IN THIS SECTION

- Syntax | 15006
- Description | 15007
- Options | 15007
- Required Privilege Level | 15007
- Release Information | 15007

Syntax

```
request services web-filter update dns-filter-database filename
```

Description

When you make changes to the domain filter database file, which is used in filtering DNS requests for disallowed domains, apply the changes.

Options

filename	File name of the database file.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

request services web-filter update url-filter-database file

IN THIS SECTION

- Syntax | 15008
- Description | 15008
- Options | 15008
- Required Privilege Level | 15008
- Release Information | 15008

Syntax

```
request services web-filter update url-filter-database file filename
```

Description

Update the URL filter database. If you change the database file, run this command. It sends a request to the DNS server for only the new hostnames that were not in the previous version of the URL database file.

Options

filename	Use the database filename.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

- request services web-filter delete gencfg-data | **15004**
- request services web-filter force dns-resolution | **15005**
- request services web-filter validate url-filter-file-name | **15010**
- *Configuring URL Filtering*

request services web-filter validate dns-filter-file-name

IN THIS SECTION

- Syntax | 15009
- Description | 15009
- Options | 15009
- Required Privilege Level | 15009
- Release Information | 15010

Syntax

```
request services web-filter validate dns-filter-file-name filename hash-key key-string hash-method hash-method-name
```

Description

Validate the file format of the domain filter database file, which is used in filtering DNS requests for disallowed domains.

Options

filename	File name of the database file.
hash-method-name	Hash method you used to produce the hashed domain name values in the database file.
key-string	Hash key you used to produce the hashed domain name values in the database file.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

request services web-filter validate url-filter-file-name

IN THIS SECTION

- Syntax | 15010
- Description | 15010
- Options | 15010
- Required Privilege Level | 15011
- Release Information | 15011

Syntax

```
request services web-filter validate (all | url-filter-file-name filename)
```

Description

Validate the URL filter database to ensure the database is correct.

Options

`all`	Validate the URL filter database for all profiles.

`url-filter-file-name` *filename* Validate the URL filter database for the database file specified.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

- request services web-filter delete gencfg-data | 15004
- request services web-filter force dns-resolution | 15005
- request services web-filter update url-filter-database file | 15007
- *Configuring URL Filtering*

request session member

IN THIS SECTION

- Syntax | 15011
- Description | 15012
- Options | 15012
- Required Privilege Level | 15012
- Release Information | 15012

Syntax

```
request session member member-id
```

Description

Start a session with the specified member of a Virtual Chassis or a VCF.

Options

member-id Member ID for the specific member of the Virtual Chassis or VCF.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

member | 10604

Understanding Virtual Chassis Components

request snmp spoof-trap

IN THIS SECTION

- Syntax | 15013
- Description | 15013
- Options | 15013
- Required Privilege Level | 15013
- Sample Output | 15014
- Release Information | 15021

Syntax

```
request snmp spoof-trap
<trap> variable-bindings <object> <instance> <value>
```

Description

Spoof (mimic) the behavior of a Simple Network Management Protocol (SNMP) trap.

Options

<*trap*>	Name of the trap to spoof.
variable-bindings <*object*> <*instance*> <*value*>	(Optional) List of variables and values to include in the trap. Each variable binding is specified as an object name, the object instance, and the value (for example, **ifIndex[14] = 14**). Enclose the list of variable bindings in quotation marks (“ “) and use a comma to separate each object name, instance, and value definition (for example, **variable-bindings “ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2”**). Objects included in the trap definition that do not have instances and values specified as part of the command are included in the trap and spoofed with automatically generated instances and values.
<*dummy name*>	A dummy trap name to display the list of available traps.
Question mark (?)	Question mark**?** to display possible completions.

Required Privilege Level

request

Sample Output

request snmp spoof-trap (with Variable Bindings)

```
user@host> request snmp spoof-trap linkUp variable-bindings “ifIndex[14] = 14, ifAdminStatus[14]
= 1, ifOperStatus[14] = 2”
Spoof trap request result: trap sent successfully
```

```
user@host> request snmp spoof-trap jnxFruOnline variable-bindings "jnxFruContentsIndex[1][1][1]
[1]=10,
jnxFruL1Index[1][1][1][1]=10, jnxFruL2Index[1][1][1][1]=10, jnxFruL3Index[1][1][1][1]=10,
jnxFruName[1][1][1][1]=Fpc, jnxFruType[1][1][1][1]=10, jnxFruSlot[1][1][1][1]=10"

Spoof trap request result: trap sent successfully
```

```
user@host> request snmp spoof-trap jnxFruOffline variable-bindings "jnxFruL2Index[1][1][1]
[1]=10,
jnxFruL3Index[1][1][1][1]=10, jnxFruName[1][1][1][1]=Fpc, jnxFruType[1][1][1][1]=10,
jnxFruSlot[1][1][1][1]=10, jnxFruOfflineReason[1][1][1][1]=10, jnxFruLastPowerOff[1][1][1]
[1]=10, jnxFruLastPowerOn[1][1][1][1]=10"

Spoof trap request result: trap sent successfully
```

request snmp spoof-trap (Illegal Trap Name)

```
user@host> request snmp spoof-trap xx
Spoof trap request result: trap not found

Allowed Traps:
adslAtucInitFailureTrap
  adslAtucPerfESsThreshTrap
  adslAtucPerfLofsThreshTrap
  adslAtucPerfLolsThreshTrap
  adslAtucPerfLossThreshTrap
  adslAtucPerfLprsThreshTrap
  adslAtucRateChangeTrap
  adslAturPerfESsThreshTrap
```

```
adslAturPerfLofsThreshTrap
adslAturPerfLossThreshTrap
adslAturPerfLprsThreshTrap
adslAturRateChangeTrap
apsEventChannelMismatch
apsEventFEPLF
apsEventModeMismatch
apsEventPSBF
apsEventSwitchover
authenticationFailure
bfdSessDown
bfdSessUp
bgpBackwardTransition
bgpEstablished
coldStart
dlswTrapCircuitDown
dlswTrapCircuitUp
dlswTrapTConnDown
dlswTrapTConnPartnerReject
dlswTrapTConnProtViolation
dlswTrapTConnUp
dsx1LineStatusChange
dsx3LineStatusChange
entConfigChange
fallingAlarm
frDLCIStatusChange
ggsnTrapChanged
ggsnTrapCleared
ggsnTrapNew
gmplsTunnelDown
ifMauJabberTrap
ipv6IfStateChange
isisAreaMismatch
isisAttemptToExceedMaxSequence
isisAuthenticationFailure
isisAuthenticationTypeFailure
isisCorruptedLSPDetected
isisDatabaseOverload
isisIDLenMismatch
isisLSPTooLargeToPropagate
isisManualAddressDrops
isisMaxAreaAddressesMismatch
isisOriginatingLSPBufferSizeMismatch
```

```
isisOwnLSPPurge
isisProtocolsSupportedMismatch
isisRejectedAdjacency
isisSequenceNumberSkip
isisVersionSkew
jnxAccessAuthServerDisabled
jnxAccessAuthServerEnabled
jnxAccessAuthServiceDown
jnxAccessAuthServiceUp
jnxBfdSessDetectionTimeHigh
jnxBfdSessTxIntervalHigh
jnxBgpM2BackwardTransition
jnxBgpM2Established
jnxCmCfgChange
jnxCmRescueChange
jnxCollFlowOverload
jnxCollFlowOverloadCleared
jnxCollFtpSwitchover
jnxCollMemoryAvailable
jnxCollMemoryUnavailable
jnxCollUnavailableDest
jnxCollUnavailableDestCleared
jnxCollUnsuccessfulTransfer
jnxDfcHardMemThresholdExceeded
jnxDfcHardMemUnderThreshold
jnxDfcHardPpsThresholdExceeded
jnxDfcHardPpsUnderThreshold
jnxDfcSoftMemThresholdExceeded
jnxDfcSoftMemUnderThreshold
jnxDfcSoftPpsThresholdExceeded
jnxDfcSoftPpsUnderThreshold
jnxEventTrap
jnxExampleStartup
jnxFEBSwitchover
jnxFanFailure
jnxFanOK
jnxFruCheck
jnxFruFailed
jnxFruInsertion
jnxFruOK
jnxFruOffline
jnxFruOnline
jnxFruPowerOff
```

```
jnxFruPowerOn
jnxFruRemoval
jnxHardDiskFailed
jnxHardDiskMissing
jnxJsAvPatternUpdateTrap
jnxJsChassisClusterSwitchover
jnxJsFwAuthCapacityExceeded
jnxJsFwAuthFailure
jnxJsFwAuthServiceDown
jnxJsFwAuthServiceUp
jnxJsNatAddrPoolThresholdStatus
jnxJsScreenAttack
jnxJsScreenCfgChange
jnxLdpLspDown
jnxLdpLspUp
jnxLdpSesDown
jnxLdpSesUp
jnxMIMstCistPortLoopProtectStateChangeTrap
jnxMIMstCistPortRootProtectStateChangeTrap
jnxMIMstErrTrap
jnxMIMstGenTrap
jnxMIMstInvalidBpduRxdTrap
jnxMIMstMstiPortLoopProtectStateChangeTrap
jnxMIMstMstiPortRootProtectStateChangeTrap
jnxMIMstNewRootTrap
jnxMIMstProtocolMigrationTrap
jnxMIMstRegionConfigChangeTrap
jnxMIMstTopologyChgTrap
jnxMacChangedNotification
jnxMplsLdpInitSesThresholdExceeded
jnxMplsLdpPathVectorLimitMismatch
jnxMplsLdpSessionDown
jnxMplsLdpSessionUp
jnxOspfv3IfConfigError
jnxOspfv3IfRxBadPacket
jnxOspfv3IfStateChange
jnxOspfv3LsdbApproachingOverflow
jnxOspfv3LsdbOverflow
jnxOspfv3NbrRestartHelperStatusChange
jnxOspfv3NbrStateChange
jnxOspfv3NssaTranslatorStatusChange
jnxOspfv3RestartStatusChange
jnxOspfv3VirtIfConfigError
```

```
jnxOspfv3VirtIfRxBadPacket
jnxOspfv3VirtIfStateChange
jnxOspfv3VirtNbrRestartHelperStatusChange
jnxOspfv3VirtNbrStateChange
jnxOtnAlarmCleared
jnxOtnAlarmSet
jnxOverTemperature
jnxPMonOverloadCleared
jnxPMonOverloadSet
jnxPingEgressJitterThresholdExceeded
jnxPingEgressStdDevThresholdExceeded
jnxPingEgressThresholdExceeded
jnxPingIngressJitterThresholdExceeded
jnxPingIngressStddevThresholdExceeded
jnxPingIngressThresholdExceeded
jnxPingRttJitterThresholdExceeded
jnxPingRttStdDevThresholdExceeded
jnxPingRttThresholdExceeded
jnxPortBpduErrorStatusChangeTrap
jnxPortLoopProtectStateChangeTrap
jnxPortRootProtectStateChangeTrap
jnxPowerSupplyFailure
jnxPowerSupplyOK
jnxRedundancySwitchover
jnxRmonAlarmGetFailure
jnxRmonGetOk
jnxSecAccessIfMacLimitExceeded
jnxSecAccessdsRateLimitCrossed
jnxSonetAlarmCleared
jnxSonetAlarmSet
jnxSpSvcSetCpuExceeded
jnxSpSvcSetCpuOk
jnxSpSvcSetZoneEntered
jnxSpSvcSetZoneExited
jnxStormEventNotification
jnxSyslogTrap
jnxTemperatureOK
jnxVccpPortDown
jnxVccpPortUp
jnxVpnIfDown
jnxVpnIfUp
jnxVpnPwDown
jnxVpnPwUp
```

```
jnxl2aldGlobalMacLimit
jnxl2aldInterfaceMacLimit
jnxl2aldRoutingInstMacLimit
linkDown
linkUp
lldpRemTablesChange
mfrMibTrapBundleLinkMismatch
mplsLspChange
mplsLspDown
mplsLspInfoChange
mplsLspInfoDown
mplsLspInfoPathDown
mplsLspInfoPathUp
mplsLspInfoUp
mplsLspPathDown
mplsLspPathUp
mplsLspUp
mplsNumVrfRouteMaxThreshExceeded
mplsNumVrfRouteMidThreshExceeded
mplsNumVrfSecIllglLblThrshExcd
mplsTunnelDown
mplsTunnelReoptimized
mplsTunnelRerouted
mplsTunnelUp
mplsVrfIfDown
mplsVrfIfUp
mplsXCDown
mplsXCUp
msdpBackwardTransition
msdpEstablished
newRoot
ospfIfAuthFailure
ospfIfConfigError
ospfIfRxBadPacket
ospfIfStateChange
ospfLsdbApproachingOverflow
ospfLsdbOverflow
ospfMaxAgeLsa
ospfNbrStateChange
ospfOriginateLsa
ospfTxRetransmit
ospfVirtIfAuthFailure
ospfVirtIfConfigError
```

```
  ospfVirtIfRxBadPacket
  ospfVirtIfStateChange
  ospfVirtIfTxRetransmit
  ospfVirtNbrStateChange
  pethMainPowerUsageOffNotification
  pethMainPowerUsageOnNotification
  pethPsePortOnOffNotification
  pingProbeFailed
  pingTestCompleted
  pingTestFailed
  ptopoConfigChange
  risingAlarm
  rpMauJabberTrap
  sdlcLSStatusChange
  sdlcPortStatusChange
  topologyChange
  traceRoutePathChange
  traceRouteTestCompleted
  traceRouteTestFailed
  vrrpTrapAuthFailure
  vrrpTrapNewMaster
  warmStart
```

request snmp spoof-trap (Question Mark ?)

```
user@host> request snmp spoof-trap ?
Possible completions:
  <trap>               The name of the trap to spoof
  adslAtucInitFailureTrap
  adslAtucPerfESsThreshTrap
  adslAtucPerfLofsThreshTrap
  adslAtucPerfLolsThreshTrap
  adslAtucPerfLossThreshTrap
  adslAtucPerfLprsThreshTrap
  adslAtucRateChangeTrap
  adslAturPerfESsThreshTrap
  adslAturPerfLofsThreshTrap
  adslAturPerfLossThreshTrap
  adslAturPerfLprsThreshTrap
  adslAturRateChangeTrap
  apsEventChannelMismatch
```

```
  apsEventFEPLF
  apsEventModeMismatch
  apsEventPSBF
  apsEventSwitchover
  authenticationFailure
  bfdSessDown
  bfdSessUp
  bgpBackwardTransition
  bgpEstablished
  coldStart
  dlswTrapCircuitDown
  dlswTrapCircuitUp
---(more 10%)---
```

Release Information

Command introduced in Junos OS Release 8.2.

request snmp utility-mib clear instance

IN THIS SECTION

- Syntax | 15021
- Description | 15022
- Options | 15022
- Required Privilege Level | 15022
- Release Information | 15022

Syntax

```
request snmp utility-mib clear instance name
object-type type
```

Description

Clear the data stored in the specified container object in the SNMP Utility MIB.

Options

name Name of the SNMP instance that is used to identify the data stored in the container object.

object-type ***type*** Type of container object in which the data is stored. The following container object types are supported:

- **counter**—Stores a 32-bit counter value.
- **counter64**—Stores a 64-bit counter value.
- **integer**—Stores a 32-bit signed integer value.
- **unsigned-integer**—Stores a 32-bit unsigned integer value.
- **string**—Stores an octet string value.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Understanding the Implementation of SNMP on the QFabric System

request snmp utility-mib set instance | **15023**

request snmp utility-mib set instance

IN THIS SECTION

- Syntax | 15023
- Description | 15023
- Options | 15023
- Required Privilege Level | 15024
- Release Information | 15024

Syntax

```
request snmp utility-mib set instance name
object-type type
object-value value
```

Description

Store data in the specified container object in the SNMP Utility MIB. The data may be retrieved by SNMP operations.

Options

name Name of the SNMP instance that is used to identify the data stored in the container object.

object-type ***type*** Type of container object in which to store data. The following container object types are supported:

- **counter**—Stores a 32-bit counter value.
- **counter64**—Stores a 64-bit counter value.
- **integer**—Stores a 32-bit signed integer value.
- **unsigned-integer**—Stores a 32-bit unsigned integer value.

- **string**—Stores an octet string value.

object-value *value* Data that is stored in the container object.

Required Privilege Level

request

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Understanding the Implementation of SNMP on the QFabric System

request snmp utility-mib clear instance | **15021**

request support information

IN THIS SECTION

- Syntax | **15025**
- Syntax (ACX Series Router) | **15025**
- Syntax (EX Series Switch) | **15025**
- Syntax (MX Series Router) | **15025**
- Syntax (PTX Series Router) | **15026**
- Syntax (SRX Series Firewalls) | **15026**
- Description | **15026**
- Options | **15027**
- Additional Information | **15029**
- Required Privilege Level | **15037**
- Output Fields | **15037**
- Sample Output | **15037**

- Release Information | **15041**

Syntax

```
request support information
<brief>
<evpn-mpls | evpn-vxlan>
<ipsec-vpn>
```

Syntax (ACX Series Router)

```
request support information
<archive>
<cos>
<dot1x>
<evpn-mpls>
<evpn-vxlan>
<with-components>
<with-logs>
<with-options>
```

Syntax (EX Series Switch)

```
request support information
<brief>
<all-members>
<local>
<member member-id>
```

Syntax (MX Series Router)

```
request support information
<archive>
<cos>
```

```
<dot1x>
<evpn-mpls>
<evpn-vxlan>
<with-components>
<with-logs>
<with-options>
```

Syntax (PTX Series Router)

```
request support information
<archive>
<cos>
<dot1x>
<evpn-mpls>
<evpn-vxlan>
<with-components>
<with-logs>
<with-options>
```

Syntax (SRX Series Firewalls)

```
request support information
<archive>
<dot1x>
<evpn-mpls>
<evpn-vxlan>
<security-components component>
<with-components>
<with-logs>
<with-options>
```

Description

Display all configuration data for the system, including data hidden with the **apply-flags omit** command. Issue this command before contacting customer support, and then include the command output in your support request. Output from this command varies somewhat, depending on which platform you issue the command from. However, the command always executes a series of `show` commands, with the appropriate information for your device automatically included.

The output from this command could be very long (thousands of lines of output). You might often want to redirect the output to a file (such as /var/log/rsi.log), as follows:

```
request support information <options> | save filename
```

Options

brief (Optional) Display brief information for the command output. Without this option, display of the output can take a long time to complete.

archive Display a compressed archive file (**tar.gz**) in the **/var/tmp/rsi** directory containing all output from the `request support information` command.

dot1x Display detailed dot1x status report.

evpn-mpls (Optional) Display system information including in-depth EVPN control plane and Packet Forwarding Engine (PFE) status for EVPN-MPLS fabric configurations. The output includes:

- EVPN instance, database, and routing information for EVPN route types (including EVPN Type 5 IP prefix routes).
- EVPN multicast and multicast snooping information.

evpn-vxlan (Optional) Display system information including in-depth EVPN control plane and PFE status for EVPN-VXLAN fabric configurations. The output includes:

- EVPN instance, database, routing (including EVPN Type 5 IP prefix routes), multicast and multicast snooping information.
- Ethernet switching, flooding, and VXLAN tunnel endpoint (VTEP) status and next hop information.
- (Starting in Junos OS and Junos OS Evolved Releases 22.4R2 and 23.1R1) Additional debugging information for EVPN-VXLAN multicast and multicast snooping, including command outputs for optimized intersubnet multicast (OISM) and assisted replication (AR) features.

ipsec-vpn (Optional) Display detailed system information for IPsec VPN.

all-members (EX Series switches and MX Series routers) (Optional) Display system information for all members of the Virtual Chassis configuration.

local

(EX Series switches and MX Series routers) (Optional) Display system information for the local Virtual Chassis member.

member *member-id*

(EX Series switches and MX Series routers) (Optional) Display system information for the specified member of the Virtual Chassis configuration. On EX Series switches, replace *member-id* with a value appropriate for that Virtual Chassis configuration. On MX Series routers, replace *member-id* with a value of 0 or 1.

secure-gateway

(SRX5400, SRX5600, and SRX5800 devices) (Optional) Display system information for secure gateway deployment scenarios.

security-components *component*

(SRX Series) (Optional) Display information on security components.

- `alg`— Display detailed security report for Application Layer Gateway (ALG).
- `apbr`— Display detailed security report for Advanced policy-based routing (APBR).
- `appid`— Display detailed security report for application identification (AppID).
- `application-firewall`— Display detailed security report for application firewall (AppFW).
- `appqoe`— Display detailed security report for application quality of experience (AppQoE).
- `appqos`— Display detailed security report for application quality of service (AppQoS).
- `ha`— Display detailed security report for firewall system with chassis cluster.
- `idp`— Display detailed security report for intrusion detection and prevention (IDP).
- `ipsec-vpn`— Display detailed security report for IPsec VPN.
- `macsec`— Display detailed security report for MACsec.
- `packet-flow`— Display detailed security report for packet flow.
- `perf`— Display detailed security report for system performance.
- `skyatp`— Display detailed security report for Juniper Advanced Threat Prevention Cloud (Juniper ATP Cloud).
- `sslfp`— Display detailed security report for ssl proxy.
- `system`— Display detailed security report for firewall system.
- `utm`— Display detailed security report for Content Security.

- `web-proxy`— Display detailed security report for web proxy.

node

(SRX5400, SRX5600, and SRX5800 devices) (Optional) Display detailed system status report for all nodes or a specific node.

- *`node-id`*—Display detailed system status report for a specific node. Replace *`node-id`* with a value of 0 or 1.
- `all`—Display detailed system status report for all nodes.
- `local`—Display detailed system status report for the local node.
- `primary`—Display detailed system status report for the primary node.

with-logs

Display a compressed archive file (**tar.gz**) containing all output from the `request support information` command and log files located in the **/var/log** directory.

> **NOTE**: The compressed archive file (**tar.gz**) contains log files located only in the **/var/log** directory, not in the **/var/log/trace** directory.

with-components

Display a compressed archive file (**tar.gz**) containing output from the component specific `request support information` commands along with output from the generic `request support information` command and log files located in the **/var/log** directory.

- `timing-debug`—Include timing specific information.

with-options

Display a compressed archive file (tar.gz) containing output from the component specific `request support information` command with a specific parameter. This command provides flexibility in gathering system information and allows you to customize the output based on your requirements.

Additional Information

The `show` commands issued as a result of this command vary depending on which platform you issue the command from. Output is always appropriate for the device. For example, Table 125 on page 15029 lists the `show` commands that are called when you issue `request support information` on an MX Series router.

Table 125: Sample show Commands Called by the request support information command on an MX Series Router

`show chassis alarms no-forwarding`	`show interfaces extensive no-forwarding`

show chassis environment no-forwarding	show krt queue
show chassis fabric degradation	show krt state
show chassis fabric destinations	show pfe statistics error
show chassis fabric fpcs	show pfe statistics traffic
show chassis fabric plane	show route summary
show chassis fabric reachability	show system boot-messages no-forwarding
show chassis fabric summary	show system buffer no-forwarding
show chassis fpc	show system commit
show chassis fpc detail	show system core-dumps no-forwarding
show chassis firmware no-forwarding	show system processes extensive no-forwarding
show chassis hardware detail no-forwarding	show system queues no-forwarding
show chassis hardware extensive no-forwarding	show system statistics no-forwarding
show pfe statistics traffic	show system storage no-forwarding
show chassis power	show system uptime no-forwarding
show chassis routing-engine no-forwarding	show system virtual-memory no-forwarding
show configuration | except SECRET-DATA	show version detail no-forwarding

> **NOTE**: The `no-forwarding` option ensures that all `mgd` processes associated with the `show` command are properly halted if you break into the output (Ctrl+C) while the command is still running.
>
> Show command `show interfaces extensive no-forwarding` is not supported for `request support information brief` command.

Table 126 on page 15031 lists the `show` commands that are called when you issue `request support information ipsec-vpn` on a MX Series Router.

Table 126: Sample show commands called by the request support information ipsec-vpn command on an MX Series Router

`show version no-forwarding | no-more`	`show services ipsec-vpn ipsec statistics | no-more`
`show system uptime no-forwarding | no-more`	`show security pki local-certificate detail no-forwarding`
`show system core-dumps no-forwarding | no-more`	`show security pki local-certificate system-generated detail no-forwarding`
`show system processes extensive no-forwarding | grep ike`	`show security pki ca-certificate detail no-forwarding`
`show system processes extensive no-forwarding | grep kmd`	`show security pki crl detail no-forwarding`
`show chassis fpc pic-status`	`show security pki certificate-request detail no-forwarding`
`show chassis cluster status`	`show security pki statistics no-forwarding | no-more`
`show chassis cluster interfaces`	`show ipsec memory-usage | no-more`
`show chassis cluster information`	`show services sessions analysis | no-more`

show configuration | no-more	show services sessions extensive | no-more
show route forwarding-table | no-more	show pfe statistics traffic | no-more
show interfaces extensive | no-more	show pfe statistics exceptions | no-more
show services ipsec-vpn ike security-associations detail | no-more	show pfe statistics error | no-more
show services ipsec-vpn ike statistics | no-more	show extension-provider system processes extensive | match mspmand
show services ipsec-vpn ipsec security-associations detail | no-more	show services service-sets statistics packet-drops | no-more
plugin ipsec show stats all	plugin ipsec show sadb detail
plugin ipsec show rules all	plugin ipsec show rules no-sa
plugin ipsec show sa-trigger-seqno	plugin ipsec show invalid-spi brief
plugin ipsec show invalid-spi history	show msp pkt-cntrs terse
show msp pkt-cntrs	show msp plugins pkt-cntrs
show msp service-sets	show msp service-sets summary
show msp service-sets mem-usage	show msp ip-stats
show msp mem usage	show msp mem ocwm jnx-msp-shm-data mum
show msp mem ocwm jnx-msp-shm-data mum objcache	show msp stats ctrl

show msp cpu-throttle	show svcs-xlp intf counters
show svcs-xlp exceptions	show svcs-xlp exceptions detail
show svcs-xlp cpu stats	show svcs-xlp fifo
show svcs-xlp poe stats	show svcs-xlp jbuf-stats
plugin ha show ckpt	plugin ha show ha-pvt-obj stats
plugin ha show info-ha-plugin	plugin ha show stats
plugin ha show udp	plugin ha show mtu
plugin ha show junos-sync stats	plugin ha show junos-sync stats aggregation
plugin ha show junos-sync stats history	

Starting in Release 20.2R1, Junos OS introduces `ipsec-vpn` option to the existing `request support information` command. The `request support information ipsec-vpn` command displays all the configurations, states, and statistics at Routing Engine and Service Card level. This new option helps in debugging IPsec-VPN related issues. The information collection is streamlined and reduces the output file size.

Table 127 on page 15033 lists the `show` commands that are called when you issue `request support information` on an EX Series 9200 switch. The table does not include the `no-forwarding` option, which is used for purposes of the `request support information`, itself.

Table 127: Sample show Commands Called by the request support information command on an EX Series 9200 Switch

show chassis alarms	show interfaces extensive
show chassis environment	
show chassis firmware	show pfe statistics traffic

show chassis fpcdetail	show spanning-tree bridge detail
show chassis hardware detail	show spanning-tree interface
show chassis routing-engine	
show configuration | except SECRET-DATA | display omit	show system boot-messages
"show dhcp-security binding" on page 19344	show system queues
"show dhcp-security ipv6 binding" on page 19351	show system processes extensive
	show system queues
show ethernet-switching interface detail	show system statistics
show ethernet-switching table	show vlans extensive
	show vrrp summary

Table 128 on page 15034 lists the show commands that are called when you issue request support information on SRX Series Firewalls.

Table 128: Sample show Commands Called by the request support information command on SRX Series Firewalls

show pfe statistics traffic	show interfaces queue
show chassis environment no-forwarding	show security monitoring fpc 0
show chassis fpc detail	show system license
show system storage no-forwarding	show security policies hit-count no-forwarding

show system virtual-memory no-forwarding	show security policies information no-forwarding
show system buffer no-forwarding	show security dns-cache
show system queues no-forwarding	show security flow statistics
show chassis hardware extensive no-forwarding	show security flow status
show krt queue	show security flow session summary no-forwarding
show route summary	show security utm anti-virus status
show version no-forwarding	show system processes extensive no-forwarding
show system uptime no-forwarding	show chassis fpc pic-status
show system core-dumps no-forwarding	show chassis cluster status
show chassis cluster interfaces	show chassis cluster information
show route forwarding-table	show interfaces extensive
show security pki local-certificate detail no-forwarding	show security pki local-certificate system-generated detail no-forwarding
show security pki ca-certificate detail no-forwarding	show security pki crl detail no-forwarding
show security pki certificate-request detail no-forwarding	show security pki statistics no-forwarding
show security internal-security-association no-forwarding	show security ike security-associations detail no-forwarding

show security ike active-peer stats no-forwarding	show security ike stats no-forwarding
show security ipsec security-associations detail no-forwarding	show security ipsec statistics no-forwarding
show security ipsec tunnel-distribution no-forwarding	show security ipsec tunnel-distribution summary no-forwarding
show security ipsec next-hop-tunnels no-forwarding	show security ipsec tunnel-events-statistics no-forwarding
show security ipsec inactive-tunnels no-forwarding	show pfe statistics traffic
show usp ipsec sa	show usp ipsec pmi-sa
show usp ipsec global-stat	show usp ipsec pkt-stat
show usp ipsec stat	show usp ipsec spd table
show usp ipsec spd stats	show usp ipsec traffic-selector stats
show usp ipsec traffic-selector pkt-stats	show usp ipsec tunnel-distribution
show usp ipsec nsp-tunnel-table	show usp ipsec nhtb
show usp ipsec ager stats	show talus *number* stats all
show talus *number* drop-pkt	show talus *number* queue-bp
show jspec talus *number* registers cntr fi_bp *number*	show jspec talus *number* registers cntr fi_bp *number*
show jspec talus *number* registers cntr fi_bp *number*	show jspec talus *number* registers cntr fi_bp *number*

show jspec talus *number* registers cntr fo_free *number*	show jspec talus *number* registers cntr fo_free *number*
show jspec talus *number* registers cntr fo_free *number*	show jspec talus *number* registers cntr fo_free *number*
show jspec talus *number* registers cntr ing_eop *number*	show jspec talus *number* registers cntr ing_eop *number*
show jspec talus *number* registers cntr ing_eop *number*	show jspec talus *number* registers cntr ing_eop *number*

Starting in Junos OS Release 15.1X49-D110, on SRX5400, SRX5600, and SRX5800 devices, a new option `secure-gateway` is added to the existing `request support information` command. This new option displays all the required information that is relevant for secure gateway deployment scenarios. In Junos OS Release 15.1X49-D100 and earlier, request support information displays the information about all features that might not be relevant for secure gateway deployments.

Starting in Junos OS Release 20.2R1, we've introduced the CLI `ipsec-vpn` option to the existing `request support information security-components` command. This new option displays all the configuration, states, and statistics information necessary for debugging IPsec VPN related issues.

Required Privilege Level

maintenance, view, and view-configuration

Output Fields

For information about output fields, see the description for the specific command--listed in the output–in which you are interested.

Sample Output

request support information | save

```
user@host> request support information | save hostA
Wrote 1143 lines of output to 'hostA'

user@host>
```

request support information (SRX Series Firewall)

```
user@host> request support information
root@host> show security dns-cache
DNS entry number: 0

root@host> show security utm web-filtering statistics

 UTM web-filtering statistics:
    Web-filtering sessions in total:    512000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit           block
          Default                       0               0
          Timeout                       0               0
     Connectivity                       0               0
Too-many-requests                       0               0

root@host> show security utm session

 UTM session info:
 Maximum sessions:                256000
 Total allocated sessions:        0
 Total freed sessions:            0
 Active sessions:                 0

root@host> show system uptime no-forwarding

Current time: 2017-07-09 21:38:01 PDT
Time Source:  LOCAL CLOCK
System booted: 2017-07-09 20:58:06 PDT (00:39:55 ago)
Protocols started: 2017-07-09 20:58:07 PDT (00:39:54 ago)
Last configured: 2017-06-30 08:56:45 PDT (1w2d 12:41 ago) by user
 9:38PM  up 40 mins, 1 user, load averages: 0.10, 0.07, 0.04
...
```

request support information ipsec-vpn (MX Series Router)

```
user@host> request support information ipsec-vpn
root@mobst480e> show services ipsec-vpn ike security-associations detail | no-more
IKE peer 10.0.1.2
```

```
  Role: Responder, State: Matured
  Initiator cookie: 9425c926b3614be5, Responder cookie: 1bcf54a314423b3f
  Exchange type: IKEv2, Authentication method: Pre-shared-keys
  Local gateway interface: ge-4/3/3.0
  Local: 10.0.1.1, Remote: 10.0.1.2
  Lifetime: Expires in 1923 seconds
  Peer ike-id: ipv4(any:0,[0..3]=10.0.1.2)
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : 2
  Traffic statistics:
   Input  bytes  :                 8024
   Output bytes  :                 7744
   Input  packets:                   52
   Output packets:                   52
  Flags: IKE SA created
  IPSec security associations: 34 created, 36 deleted

root@mobst480e> show services ipsec-vpn ike statistics | no-more
IKE peer 10.0.1.2
  Initiator cookie: 9425c926b3614be5, Responder cookie: 1bcf54a314423b3f
  Local: 10.0.1.1, Remote: 10.0.1.2
  IPSec security associations: 34 created, 36 deleted
  IPSec tunnel rekeys: 17
  Exchange type: IKEv2
  Traffic statistics:
   Input  bytes  :                                        8024
   Output bytes  :                                        7744
   Input  packets:                                          52
   Output packets:                                          52
  Delete IPsec SA payload statistics:
   Delete IPsec SA payloads received:                       18
   Delete IPsec SA payloads sent    :                       17
  DPD statistics:
   DPD request received           :                         18
   DPD request sent               :                          0
   DPD response received          :                          0
   DPD response sent              :                         18
   DPD response missed            :                          0
   Worst case DPD response delay  :                          0 msec
   DPD sequential response missed :                          0
```

```
  Invalid SPI notification statistics:
   Notifications received:                                  0
   Notifications sent    :                                  0

root@mobst480e> show services ipsec-vpn ipsec security-associations detail | no-more
Service set: ss1, IKE Routing-instance: default

  Rule: ipsec_vpn_rules_1, Term: term1, Tunnel index: 1
  Local gateway: 10.0.1.1, Remote gateway: 10.0.1.2
  IPSec inside interface: ms-2/0/0.1, Tunnel MTU: 1500
  UDP encapsulate: Disabled, UDP Destination port: 0
  Local identity: ipv4_subnet(any:0,[0..7]=192.168.60.0/24)
  Remote identity: ipv4_subnet(any:0,[0..7]=192.168.70.0/24)
  NATT Detection: Not Detected, NATT keepalive interval: 0
  Total uptime:  0 days 2 hrs 45 mins 21 secs

    Direction: inbound, SPI: 2957881284, AUX-SPI: 0
    Mode: tunnel, Type: dynamic, State: Installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Soft lifetime: Expires in 48 seconds
    Hard lifetime: Expires in 88 seconds
    Anti-replay service: Enabled, Replay window size: 4096
    Copy ToS: Enabled
    Copy TTL: Disabled, TTL value: 64
    SA lifetime: 180 seconds

    Direction: outbound, SPI: 3292864325, AUX-SPI: 0
    Mode: tunnel, Type: dynamic, State: Installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Soft lifetime: Expires in 48 seconds
    Hard lifetime: Expires in 88 seconds
    Anti-replay service: Enabled, Replay window size: 4096
    Copy ToS: Enabled
    Copy TTL: Disabled, TTL value: 64
    SA lifetime: 180 seconds

root@mobst480e> show services ipsec-vpn ipsec statistics | no-more

PIC: ms-2/0/0, Service set: ss1

ESP Statistics:
  Encrypted bytes:                0
  Decrypted bytes:                0
```

```
  Encrypted packets:              0
  Decrypted packets:              0
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures:     0
  ESP authentication failures:    0
  ESP decryption failures:        0
  Bad headers: 0, Bad trailers: 0
  Replay before window drops: 0, Replayed pkts: 0
  IP integrity errors: 0, Exceeds tunnel MTU: 0
  Rule lookup failures: 0, No SA errors: 0
  Flow errors: 0, Misc errors: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `brief` introduced in Junos OS Release 13.2.

Option `secure-gateway` introduced in Junos OS Release 15.1X49-D110 for SRX5400, SRX5600, and SRX5800 devices.

Options `security-components` and `ipsec-vpn` introduced in Junos OS Release 20.2R1.

Options `evpn-mpls` and `evpn-vxlan` introduced in Junos OS and Junos OS Evolved Releases 21.2R3, 21.4R2, 21.4R3, and 22.2R1.

Additional EVPN multicast debug information added to `evpn-vxlan` option output in Junos OS and Junos OS Evolved Releases 22.4R2 and 23.1R1.

Options `archive`, `with-logs`, `with-components`, and `with-options` introduced in Junos OS and Junos OS Evolved Release 24.2R1.

request system application (Junos OS Evolved)

IN THIS SECTION

- Syntax | 15042
- Description | 15042
- Options | 15042
- Additional Information | 15043
- Required Privilege Level | 15043
- Sample Output | 15043
- Release Information | 15043

Syntax

```
request system application app application-name node node-name restart
```

Description

Use this command to stop and then start (restart) an application on the specified node. Use the `show system applications` command to verify if an application is started or stopped.

Options

app *application-name*	Specify the application you want started or stopped.
node *node-name*	Specify the name of the node on which to start or stop the application.
restart	Restart the application.

NOTE: Make sure to specify the correct node for the application. The sFlow application (`sflowd`) runs on the Routing Engine node on PTX10001-36MR and PTX10003 routers. It runs on the FPC node on PTX10004 and PTX10008 routers.

Additional Information

(PTX10001-36MR only) Starting in Junos OS Evolved Release 23.2R1, when using this command to restart one of the forwarding applications, such as `evo-cda-bt` or `evo-aftmand-bt`, the restart causes traffic forwarding to stop until all applications are back to normal. You might want to restart these forwarding applications if those applications are leaking memory or file handles. When `evo-cda-bt` or `evo-aftmand-bt` are restarted, other forwarding applications are restarted as well.

(PTX10001-36MR only) Starting in Junos OS Evolved Release 24.2R1, when the forwarding applications `evo-cda-bt` and `evo-aftmand-bt` restart after failure, other forwarding applications also restart (`packetio-bt`, `ppman-aft-bt`, `evoaft-jvisiond-bt`, `aft-sysinfo`, `fabspoked-pfe`, and `securityd`). This behavior makes the system more resilient to abnormal termination of the `evo-cda-bt` and `evo-aftmand-bt` applications. When these forwarding applications restart, traffic forwarding stops until all applications are back to normal. Prior to this release, the entire chassis would restart if either `evo-cda-bt` or `evo-aftmand-bt` failed. To manually restart the `evo-cda-bt`, `evo-aftmand-bt`, or `packetio-bt` applications if they ever leak memory or file handles, use the `request system application` *`application-name`* `node` *`nodename`* `restart` operational mode command.

Required Privilege Level

view

Sample Output

request system application app application-name node node-name restart

```
user@host> request system application app cmdd node fpc0 restart
This may affect traffic in the system. Proceed ? [yes,no] (no) yes

App cmdd on node fpc0 restart request is submitted
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- request node halt (Junos OS Evolved) | **14708**
- request node (offline | online) (Junos OS Evolved) | **14711**
- request node power-off (Junos OS Evolved) | **14713**

request node power-on (Junos OS Evolved) | 14715
request node reboot (re0 | re1) (Junos OS Evolved) | 14717
show system applications (Junos OS Evolved) | 26695

request system autorecovery state

IN THIS SECTION

- Syntax | 15044
- Description | 15044
- Options | 15045
- Required Privilege Level | 15045
- Output Fields | 15045
- Sample Output | 15045
- Sample Output | 15046
- Sample Output | 15046
- Release Information | 15046

Syntax

```
request system autorecovery state  (save | recover | clear)
```

Description

Use this command to prepare the system for autorecovery of configuration, licenses, and disk information.

In devices running FreeBSD Release 12 or later, you cannot back up data with the autorecovery feature. Instead, back up data using snapshots. To learn if your device is running FreeBSD Release 12 or later, issue the `show version` command and look for the `fbsd_builder_stable` string in the module names. If the string includes the number 12 or later, your device is running FreeBSD Release 12 or later.

Options

save Save the current state of the disk partitioning, configuration, and licenses for autorecovery.

The active Junos OS configuration is saved as the Junos rescue configuration, after which the rescue configuration, licenses, and disk partitioning information is saved for autorecovery. Autorecovery information must be initially saved using this command for the autorecovery feature to verify integrity of data on every bootup.

Any recovery performed at a later stage will restore the data to the same state as it was when the save command was executed.

A fresh rescue configuration is generated when the command is executed. Any existing rescue configuration will be overwritten.

recover Recover the disk partitioning, configuration, and licenses.

After autorecovery data has been saved, the integrity of saved items is always checked automatically on every bootup. The recovery command allows you to forcibly re-run the tests at any time if required.

clear Clear all saved autorecovery information.

Only the autorecovery information is deleted; the original copies of the data used by the router are not affected. Clearing the autorecovery information also disables all autorecovery integrity checks performed during bootup.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system autorecovery state save

```
user@host> request system autorecovery state save
Saving config recovery information
```

```
       Saving license recovery information
       Saving bsdlabel recovery information
```

Sample Output

request system autorecovery state recover

```
user@host> request system autorecovery state recover

Configuration:
File             Recovery Information    Integrity Check    Action / Status
rescue.conf.gz   Saved                   Passed             None
Licenses:
File             Recovery Information    Integrity Check    Action / Status
JUNOS282736.lic  Saved                   Passed             None
JUNOS282737.lic  Saved                   Failed             Recovered
BSD Labels:
Slice            Recovery Information    Integrity Check    Action / Status
s1               Saved                   Passed             None
s2               Saved                   Passed             None
s3               Saved                   Passed             None
s4               Saved                   Passed             None
```

Sample Output

request system autorecovery state clear

```
user@host> request system autorecovery state clear
Clearing config recovery information
       Clearing license recovery information
       Clearing bsdlabel recovery information
```

Release Information

Command introduced in Junos OS Release 15.1X49-D35.

RELATED DOCUMENTATION

show system autorecovery state | 26725

request system boot console

IN THIS SECTION

- Syntax | 15047
- Description | 15047
- Options | 15047
- Required Privilege Level | 15048
- Sample Output | 15048
- Release Information | 15048

Syntax

```
request system boot-console <all-members|auxiliary|console|local|member>
```

Description

Use this command to choose the port for early boot logs.

Options

all-members	Switch boot-console port all virtual chassis members
auxiliary	Set Auxiliary port as boot console
console	Set Console port as boot console
local	Switch boot-console port local virtual chassis member
member	Switch boot-console port specific virtual chassis member

Required Privilege Level

maintenance

Sample Output

request system boot-console

```
user@host> request system boot-console
Command options:
console   : RJ-45 console port
auxiliary : USB-C console port

(rest of the options are the usual option related to VC)
```

Release Information

Command introduced in Junos OS Release 21.3R1.

request system certificate add

IN THIS SECTION

- Syntax | 15049
- Description | 15049
- Options | 15049
- Required Privilege Level | 15049
- Output Fields | 15049
- Sample Output | 15049
- Release Information | 15049

Syntax

```
request system certificate add (filename | terminal)
```

Description

(Encryption interface on M Series and T Series routers, PTX Series, and QFX Series switches only) Add a certificate provided by the Juniper Networks certificate authority (CA).

Options

filename	Filename (URL, local, or remote).
terminal	Use login terminal.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system certificate add terminal

```
user@host> request system certificate add terminal
```

Release Information

Command introduced before Junos OS Release 7.4.

request system commit server pause

IN THIS SECTION

- Syntax | 15050
- Description | 15050
- Options | 15050
- Required Privilege Level | 15050
- Sample Output | 15050
- Release Information | 15051

Syntax

```
request system commit server pause
```

Description

Pause the commit server.

> **NOTE**: If you issue this command when a commit job is in process, the batch commit server pauses only after the current commit job is completed.

Options

This command has no options.

Required Privilege Level

view

Sample Output

When you enter the `request system commit server pause` command, you are provided feedback on the status of your request.

request system commit server pause

```
user@host> request system commit server pause

Successfully paused the commit server.
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Batch Commit Server Properties

request system commit server queue cleanup

IN THIS SECTION

- Syntax | 15051
- Description | 15052
- Options | 15052
- Required Privilege Level | 15052
- Sample Output | 15052
- Release Information | 15053

Syntax

```
request system commit server queue cleanup <id commit-id | job-status (error | pending |
success)>
```

Description

Clean up the batch commit queue. Note that the `id` argument cleans up batch commit operation messages for a specific commit ID, whereas `job-status` cleans up more broadly, based on categories of status messages. You can use either option, but not both.

Options

`id` *`commit-id`* (Optional) Clean up batch commit operation status messages for a specific commit ID.

`job-status` (Optional) Clean up batch commit operation status messages for the following:

- `error`—Clean up status messages for batch commit operations that have errors.
- `pending`—Clean up status messages for batch commit operations that are pending.
- `success`—Clean up status messages for batch commit operations that are successful.

Required Privilege Level

view

Sample Output

When you enter the `request system commit server queue cleanup` command, you are provided feedback on the status of your request. The first example demonstrates cleaning up job ID 1008, while the second shows a queue clean up for all jobs marked as successfully completed.

request system commit server queue cleanup id

```
user@host> request system commit server queue cleanup id 1008

Successfully cleaned up jobs.
```

request system commit server queue cleanup job-status success

```
user@host> request system commit server queue cleanup job-status success

Successfully cleaned up jobs.
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Batch Commit Server Properties

request system commit server start

IN THIS SECTION

- Syntax | 15053
- Description | 15053
- Options | 15053
- Required Privilege Level | 15054
- Sample Output | 15054
- Release Information | 15054

Syntax

```
request system commit server start
```

Description

Start the commit server.

Options

This command has no options.

Required Privilege Level

view

Sample Output

When you enter the `request system commit server start` command, you are provided feedback on the status of your request.

request system commit server start

```
user@host> request system commit server start

Successfully started the commit server.
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Batch Commit Server Properties

request system configuration database resize

IN THIS SECTION

- Syntax | 15055
- Description | 15055
- Options | 15055
- Additional Information | 15056
- Required Privilege Level | 15056
- Output Fields | 15056

Sample Output | 15057
Release Information | 15057

Syntax

```
request system configuration database resize
<force>
<re0 | re1>
<routing-engine backup | master | both | local>
```

Description

You might require to perform a resizing of the configuration database as a result of performing frequent load replacement operations which replaces the configuration hierarchy with the new configuration. When this occurs, deleted objects store in the database, and the memory allocated in the database is freed after the commit is completed. However, during this process the database size on disk can grow although the actual configuration size is actually less.

The `request system configuration database resize` command provides you with a manual method to resize the configuration database. The system will display `Database resize completed` when the configuration database has been successfully resized.

Use the `show system configuration database usage` command to display configuration database disk space usage statistics. The Current database size on disk field will display the change in database size related to the use of the `resize` command. See show system configuration database usage.

Options

force (Optional) If there are uncommitted changes and you execute the `request system configuration database resize` command, the command will fail and an error will be generated (`error: Database resize failed, Configuration Database Modified.`). To execute the command when this occurs, use the `force` option to discard the uncommitted changes.

re0 | re1 (Optional) If a router has dual Routing Engines, `re0` and `re1` identify the specific routing engine that is to execute the `request system configuration database resize` command. `re0` is for the Routing Engine in slot 0 and `re1` is for the Routing Engine in slot 1.

routing-engine (Optional) To execute the `request system configuration database resize` command on a specific routing engine, or on both routing engines, in a redundant configuration you can specify one of the following:

- backup—Resizes the configuration database on the backup routing engine in a redundant configuration.
- primary—Resizes the configuration database on the primary routing engine in a redundant configuration.
- both—Resizes the configuration database on both routing engines.
- local—Resizes the configuration database on the local routing engine.

Additional Information

Include below is a set of usage guidelines when using the `request system configuration database resize` command:

- All other configuration sessions will be closed on the routing engine to remove the old mappings to the configuration database.
- This command cannot be run from configuration mode.
- Do not execute the `request system configuration database resize` command when a commit is in progress.
- During the time when the `request system configuration database resize` command is in progress, a new configuration session cannot be opened. A warning message will appear instructing you that the configuration database is being resized (`warning: Database is being resized`).
- If there are uncommitted changes and the command is executed, the command will fail and an error will be generated (`error: Database resize failed, Configuration Database Modified.`). To execute the command when this occurs, use the `force` option to discard the uncommitted changes.
- To execute the command on both routing engines use the `request system configuration database resize routing-engine both` command.
- Do not terminate the command when it is in progress (for example, by specifying **Ctrl+c** or **Ctrl+z**).

Required Privilege Level

maintenance

Output Fields

Sample Output

request system configuration database resize

```
user@host> show system configuration database usage
Maximum size of the database: 1309.99 MB
Current database size on disk: 500.00 MB
Actual database usage: 176.81 MB
Available database space: 1133.18 MB

user@host> request system configuration database resize
Database resize completed
user@host> show system configuration database usage
Maximum size of the database: 1309.99 MB
Current database size on disk: 176.50 MB
Actual database usage: 176.48 MB
Available database space: 1133.52 M
```

Release Information

Command introduced in Junos OS Release 17.2R1.

request system configuration rescue delete

IN THIS SECTION

- Syntax | 15058
- Description | 15058
- Options | 15058
- Required Privilege Level | 15058
- Output Fields | 15058
- Sample Output | 15058
- Release Information | 15058

Syntax

```
request system configuration rescue delete
```

Description

Delete an existing rescue configuration.

NOTE: The `[edit system configuration]` hierarchy is not available on QFabric systems.

Options

This command has no options.

Required Privilege Level

maintenance

Output Fields

This command produces no output.

Sample Output

request system configuration rescue delete

```
user@host> request system configuration rescue delete
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Evolved Release 20.4R2.

request system configuration rescue save

IN THIS SECTION

- Syntax | 15059
- Description | 15059
- Options | 15059
- Required Privilege Level | 15059
- Output Fields | 15060
- Sample Output | 15060
- Release Information | 15060

Syntax

```
request system configuration rescue save
```

Description

Save the most recently committed configuration as the rescue configuration so that you can return to it at any time by using the `rollback` command. If saved on a device with redundant Routing Engines, the rescue configuration file is saved on both Routing Engines.

NOTE: The `[edit system configuration]` hierarchy is not available on QFabric systems.

Options

This command has no options.

Required Privilege Level

maintenance

Output Fields

This command produces no output.

Sample Output

request system configuration rescue save

```
user@host> request system configuration rescue save
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Evolved Release 20.4R2.

request system convert-json-configuration

IN THIS SECTION

- Syntax | 15060
- Description | 15061
- Options | 15061
- Required Privilege Level | 15061
- Sample Output | 15061
- Release Information | 15062

Syntax

```
request system convert-json-configuration filename output-filename output-filename
<origin-schema schema>
```

Description

Convert the JSON configuration data in the specified input file so that list keys precede any other siblings within a list entry as required by the Junos schema. The converted data is saved in the specified output file in the current working directory. The command does not emit output when the conversion is successful.

If the input file includes only OpenConfig configuration statements, or if it includes both OpenConfig and Junos configuration statements, include the `origin-schema openconfig` command option to specify the schema.

Options

filename Path to the input JSON configuration file.

origin-schema ***schema*** (Optional) The schema that defines the input configuration data. Specify one of the following:

- `juniper`—(Default) Input configuration data is defined only by the native Junos schema.
- `openconfig`—Input configuration data is defined by the OpenConfig schema or a combination of the OpenConfig and Junos schemas.

output-filename ***output-filename*** Filename for the output JSON configuration file. The destination path is the current working directory, which default's to the user's home directory.

Required Privilege Level

maintenance

Sample Output

request system convert-json-configuration

```
user@host> request system convert-json-configuration config.json output-filename config-
converted.json origin-schema openconfig

user@host>
```

Release Information

Command introduced in Junos OS Release 22.1R1 and Junos OS Evolved Release 22.1R1.

RELATED DOCUMENTATION

- json | 9436

Load JSON Configuration Data With Unordered List Entries

request system debug-info

IN THIS SECTION

- Syntax | 15062
- Syntax (Junos EVO devices) | 15062
- Description | 15063
- Options | 15063
- Additional Information | 15063
- Required Privilege Level | 15065
- Sample Output | 15065
- Release Information | 15066

Syntax

```
request system debug-info
<node node-name>
```

Syntax (Junos EVO devices)

```
request system debug-info
<mode custom | lite | normal>
```

```
<since duration>
<node node-name>
```

Description

Collect debug information from Junos OS Evolved, such as logs. Logs are stored in the **/var/tmp/** directory. Debug collector collects cores, application traces, configurations, commits and so on. It also collects outputs from CLI and CLI-PFE commands. You need to login to the device as root user to use this command. Enter the root password and confirm it when requested.

Options

node ***node-name*** (Optional) Identifier of the node. If a value for `node` (for example, `re0`, `re1`, `fpc0`) is given, this command collects logs only from the given node. If the node option is not provided, log information is collected from all active nodes.

mode

- *custom*—Support for custom debug collection. The values are normal, lite, l2, dot1x, vrrp, lacp, infra, mcast, vxlan, extensive, logs.
- `lite`—You can use this option to run selected debug commands.
- `normal`—You can use this option to run all debug commands.

since Specifies the log duration. The supported format is yyyy-mm-dd or #d (ex: 2d) or #h (ex: 5h). By default, you can collect the logs for a duration of 90 days.

Additional Information

Debug collector takes maximum time while collecting journal logs and traces. You can use the `request system debug-info` command for quick debug collection from **/var/tmp** in the following sample format:

- `request system debug-info node re0,fpc0,fpc1,fpc2 since 1h`—Normal mode with minimal journal and trace logs collection from master RE and all FPCs.
- `request system debug-info node re1 since 1h`—Normal mode with minimal journal and trace logs collection from backup RE only.
- `request system debug-info mode lite node re0,fpc0,fpc1,fpc2 since 1h` —Lite mode with minimal journal and trace logs collection from master RE and all FPCs.
- `request system debug-info mode lite node re1 since 1h`—Lite mode with minimal journal and trace logs collection from backup RE only.

To view the **debug_collector.log** file:

1. Go to the **/var/tmp/** directory.

```
root@host:~# cd /var/tmp
root@host:/var/tmp# ls
debug_collector_2017-11-20_13_16_45.tar.gz
```

2. Uncompress the **/var/tmp/** file.

```
root@host:/var/tmp# tar -xf debug_collector_2017-11-20_13_16_45.tar.gz
root@host:/var/tmp# ls
debug_collector_2017-11-20_13_16_45
debug_collector_2017-11-20_13_16_45.tar.gz
```

3. Go to the **/var/tmp/debug_collector_*timestamp*** directory.

```
root@host:/var/tmp# cd debug_collector_2017-11-20_13_16_45
root@host:/var/tmp/debug_collector_2017-11-20_13_16_45# ls
debug_collector.log re0  re1  fpc0
```

4. Look at the log file.

```
root@host:/var/tmp/debug_collector_2017-11-20_13_16_45# cat debug_collector.log
2017-11-20 13:16:45 INFO: debug_collector.py:704 <module>(): Running debug_collector as user
'root' with super user privileges...
2017-11-20 13:16:45 INFO: debug_collector.py:705 <module>(): You can monitor detailed
progress using:
2017-11-20 13:16:45 INFO: debug_collector.py:706 <module>():    tail -f /var/tmp/
debug_collector_2017-11-20_13_16_45/debug_collector.log
2017-11-20 13:16:45 INFO: debug_collector.py:323 fetch_root_pswd(): Debug collector needs to
login into all nodes in the system as root user
2017-11-20 13:25:31 DEBUG: debug_collector.py:332 fetch_root_pswd(): Passwords matched,
proceeding with debug collection...
2017-11-20 13:25:31 INFO: debug_collector.py:225 detect_res(): Detecting REs in the system...
2017-11-20 13:25:31 DEBUG: debug_collector.py:239 detect_res(): Into for
2017-11-20 13:25:31 DEBUG: debug_collector.py:242 detect_res(): Into m1 not None
2017-11-20 13:25:31 DEBUG: debug_collector.py:245 detect_res(): Found RE 're0' in /etc/hosts
```

```
2017-11-20 13:25:31 DEBUG: debug_collector.py:248 detect_res(): Into if 1
2017-11-20 13:25:31 DEBUG: debug_collector.py:239 detect_res(): Into for
2017-11-20 13:25:31 DEBUG: debug_collector.py:242 detect_res(): Into m1 not None
2017-11-20 13:25:31 DEBUG: debug_collector.py:245 detect_res(): Found RE 're1' in /etc/hosts
2017-11-20 13:25:31 DEBUG: debug_collector.py:261 detect_res(): Into if 4
2017-11-20 13:25:31 INFO: debug_collector.py:265 detect_res(): This RE is 're0'
2017-11-20 13:25:31 DEBUG: debug_collector.py:266 detect_res(): This RE is 're0' and the
other RE is 're1'
2017-11-20 13:25:32 DEBUG: debug_collector.py:269 detect_res(): Other RE 're1' is reachable
2017-11-20 13:25:32 INFO: debug_collector.py:274 detect_res():  List of reachable REs:
['re0', 're1']
2017-11-20 13:25:32 INFO: debug_collector.py:293 detect_fpcs(): Detecting FPCs in the system,
might take upto a couple of minutes...
2017-11-20 13:25:33 DEBUG: debug_collector.py:309 detect_fpcs(): FPC 'fpc0' is reachable
2017-11-20 13:25:36 DEBUG: debug_collector.py:312 detect_fpcs(): FPC 'fpc1' is not reachable
2017-11-20 13:25:39 DEBUG: debug_collector.py:312 detect_fpcs(): FPC 'fpc2' is not reachable
2017-11-20 13:25:40 INFO: debug_collector.py:977 <module>(): User interrupt detected,
stopping further debug collection...
2017-11-20 13:25:40 INFO: debug_collector.py:372 create_final_archive(): Creating archive of
all logs collected...
```

Required Privilege Level

view

Sample Output

request system debug-info

```
root@host> request system debug-info
2017-11-20 13:16:45 INFO: debug_collector.py:704 <module>(): Running debug_collector as user
'root' with super user privileges...
2017-11-20 13:16:45 INFO: debug_collector.py:705 <module>(): You can monitor detailed progress
using:
2017-11-20 13:16:45 INFO: debug_collector.py:706 <module>():    tail -f /var/tmp/
debug_collector_2017-11-20_13_16_45/debug_collector.log
2017-11-20 13:16:45 INFO: debug_collector.py:323 fetch_root_pswd(): Debug collector needs to
login into all nodes in the system as root user
Please enter root password:
Please re-enter the password:
2017-11-20 13:25:31 INFO: debug_collector.py:225 detect_res(): Detecting REs in the system...
```

```
2017-11-20 13:25:31 INFO: debug_collector.py:265 detect_res(): This RE is 're0'
2017-11-20 13:25:32 INFO: debug_collector.py:274 detect_res():  List of reachable REs: ['re0',
're1']
2017-11-20 13:25:32 INFO: debug_collector.py:293 detect_fpcs(): Detecting FPCs in the system,
might take upto a couple of minutes...
^C2017-11-20 13:25:40 INFO: debug_collector.py:977 <module>(): User interrupt detected, stopping
further debug collection...
2017-11-20 13:25:40 INFO: debug_collector.py:372 create_final_archive(): Creating archive of all
logs collected...
2017-11-20 13:25:40 INFO: debug_collector.py:380 create_final_archive(): Here it is: /var/tmp/
debug_collector_2017-11-20_13_16_45.tar.gz
2017-11-20 13:25:40 INFO: debug_collector.py:389 create_final_archive(): Debug collector run
time: 0h 8m 54s
2017-11-20 13:25:40 INFO: debug_collector.py:401 cleanup_and_exit(): Over n out!
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

request system decrypt password

IN THIS SECTION

- Syntax | 15067
- Description | 15067
- Options | 15067
- Required Privilege Level | 15067
- Output Fields | 15067
- Sample Output | 15067
- Sample Output | 15067
- Release Information | 15068

Syntax

```
request system decrypt password
```

Description

Use to display plain text versions of obfuscated ($9) or encrypted ($8) passwords. If the password was encrypted using the new 8 method, you are prompted for the primary password.

Options

- `decrypt`—Decrypt a 8-encrypted or 9-encrypted password.

Required Privilege Level

system

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
// Decrypting a $9 password
user@host> request system decrypt password $9$ABC123
Plaintext password: mysecret
```

Sample Output

```
// Decrypting a $8 password
user@host> request system decrypt password $8$ABC123
Master password:
Plaintext password: mysecret
(Simple passwords like "mysecret" are discouraged. This is an example only.)
```

Release Information

Statement introduced in Junos OS Release 15.1X49-D50.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

master-password | 10365

Hardening Shared Secrets in Junos OS

request system disable unified-services

IN THIS SECTION

- Syntax | 15068
- Description | 15068
- Required Privilege Level | 15069
- Output Fields | 15069
- Sample Output | 15069
- Release Information | 15069

Syntax

```
request system disable unified-services
```

Description

Disable Next Gen Services services on the MX Series.

Before you disable Next Gen Services, delete any router configuration for services. This includes configuration under the `[edit services]` hierarchy, configuration for services interfaces, and any configuration that refers to services interfaces.

After you enter `request system enable unified-services`, reboot the chassis.

Required Privilege Level

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system disable unified-services

```
user@host> request system disable unified-services
Before disabling unified services, please move to baseline configuration.
Are above conditions satisfied ? [yes,no]
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Enabling and Disabling Next Gen Services

request system download abort

IN THIS SECTION

- Syntax | 15070
- Description | 15070
- Options | 15070
- Required Privilege Level | 15070
- Output Fields | 15070

Sample Output | 15070
Release Information | 15071

Syntax

```
request system download abort <download-id>
```

Description

Use this command to terminate a download. The download instance is stopped and cannot be resumed. Any partially downloaded file is automatically deleted to free disk space. Information regarding the download is retained and can be displayed with the `show system download` command until a `request system download clear` operation is performed.

Downloads in the active, paused, and error states can be terminated.

Options

`download-id`—(Required) The ID number of the download to be terminated.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system download abort

```
user@host> request system download abort 1
Aborted download #1
```

Release Information

Command introduced in Junos OS Release 15.1X49-D35.

RELATED DOCUMENTATION

- request system download start | 15076
- request system download pause | 15072
- request system download resume | 15074
- request system download clear | 15071

request system download clear

IN THIS SECTION

- Syntax | 15071
- Description | 15071
- Required Privilege Level | 15072
- Output Fields | 15072
- Sample Output | 15072
- Release Information | 15072

Syntax

```
request system download clear
```

Description

Use this command to delete the history of completed and aborted downloads.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system download clear

```
user@host> request system download clear
Cleared information on completed and aborted downloads
```

Release Information

Command introduced in Junos OS Release 15.1X49-D35.

RELATED DOCUMENTATION

- request system download start | 15076
- request system download pause | 15072
- request system download resume | 15074
- request system download abort | 15069

request system download pause

IN THIS SECTION

- Syntax | 15073
- Description | 15073
- Options | 15073

- Required Privilege Level | 15073
- Output Fields | 15073
- Sample Output | 15073
- Release Information | 15074

Syntax

```
request system download pause <download-id>
```

Description

Use this command to suspend a particular download instance. Downloads in the active state can be paused.

Options

`download-id`—(Required) The ID number of the download to be paused.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system download pause

```
user@host> request system download pause 1
Paused download #1
```

Release Information

Command introduced in Junos OS Release 15.1X49-D35.

RELATED DOCUMENTATION

- request system download start | 15076
- request system download resume | 15074
- request system download abort | 15069
- request system download clear | 15071

request system download resume

IN THIS SECTION

- Syntax | 15074
- Description | 15074
- Options | 15075
- Required Privilege Level | 15075
- Output Fields | 15075
- Sample Output | 15075
- Release Information | 15075

Syntax

```
request system download resume download-id <max-rate>
```

Description

Use this command to resume a download that has been paused. You can resume the downloaded instances that are not in progress because of an error or that have been explicitly paused. The file will continue downloading from the point where it paused. By default, the download resumes with the same

bandwidth specified with the `request system download start` command. You can specify a new (maximum) bandwidth with the `request system download resume` command.

Downloads in the paused and error states can be resumed.

Options

`download-id`—(Required) The ID number of the download to be resumed.

`max-rate`—(Optional) The maximum bandwidth for the download.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system download resume

```
user@host> request system download resume 1
Resumed download #1
```

Release Information

Command introduced in Junos OS Release 15.1X49-D35.

RELATED DOCUMENTATION

- request system download start | 15076
- request system download pause | 15072
- request system download abort | 15069
- request system download clear | 15071

request system download start

IN THIS SECTION

- Syntax | 15076
- Description | 15076
- Options | 15076
- Required Privilege Level | 15077
- Output Fields | 15077
- Sample Output | 15077
- Release Information | 15077

Syntax

```
request system download start  (sftp-url | delay | identity-file | login | max-rate | passphrase
| save as )
```

Description

Use this command to create a download instance and identify it with a unique integer called the download ID.

Options

`sftp-url`—(Required) The FTP or HTTP URL location of the file to be downloaded securely.

`delay`—(Optional) The number of hours after which the download should start. Ranges from 1 through 48 hours.

`identity-file`—(Required) The name of the file requesting a Secure FTP (SFTP) download. The SFTP in smart download leverages public key authentication to authenticate a download request. You must generate a private or public key pair before starting a download, and then upload a public key to an SFTP server.

`login`—(Optional) The username and password for the server in the format `username:password`.

`max-rate`—(Optional) The maximum average bandwidth for the download. Numbers with the suffix k or K, m or M, and g or G are interpreted as Kbps, Mbps, or Gbps, respectively.

`passphrase`—(Required) The passphrase to protect the private key file stored on the file system. This option does not allow the user to enter a weak passphrase, which ensures stronger security.

`save-as`—(Optional) The filename to be used for saving the file in the **/var/tmp** location.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system download start

```
user@host> request system download start identity-file mytestkey sftp://mysftpserver/homes/kelly/
test.tgz max-rate 200 save as newfile.tgz
Starting download #8
```

Release Information

Command introduced in Junos OS Release 15.1X49-D35.

RELATED DOCUMENTATION

- request system download pause | 15072
- request system download resume | 15074
- request system download abort | 15069
- request system download clear | 15071

request system enable unified-services

IN THIS SECTION

- Syntax | 15078
- Description | 15078
- Required Privilege Level | 15078
- Output Fields | 15078
- Sample Output | 15079
- Release Information | 15079

Syntax

```
request system enable unified-services
```

Description

Enable Next Gen Services services on the MX Series.

Before you enable Next Gen Services, delete any router configuration for services. This includes configuration under the `[edit services]` hierarchy, configuration for services interfaces, and any configuration that refers to services interfaces.

After you enter `request system enable unified-services`, reboot the chassis.

In Junos node slicing, you can enable unified services at guest network function (GNF), by using the CLI `request system enable unified-services` at GNF.

Required Privilege Level

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system enable unified-services

```
user@host> request system enable unified-services
Before enabling unified services, please move to baseline configuration.
Are above conditions satisfied ? [yes,no]
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Enabling and Disabling Next Gen Services

request system filesystem encryption enable

IN THIS SECTION

- Syntax | 15080
- Prerequisites | 15080
- Description | 15080
- Options | 15080
- Required Privilege Level | 15081
- Sample Output | 15081
- Release Information | 15082

Syntax

```
request system filesystem encryption enable
<dry-run | re0 | re1>
```

Prerequisites

Following are the prerequisites to enable file-system encryption:

- System contains a TPM2.0 with IDevID provisioned.
- System having single or redundant disk are supported.
- Take data backup of configurations and log files.

Description

When you enable encryption process on the file-system, the conversion process starts with the backup routing engine followed by the active Routing Engine. In the case of redundant disks, the conversion starts with the primary disk followed by the secondary disk to avoid loss of data.

Once enabled, the encryption cannot be disabled and all the software image versions that does not support file-system encryption are deleted.

Options

none Enable file-system encryption on all Routing Engines.

dry-run (Optional) Display the file-system encryption message without running the encryption process.

re0 (Optional) Enable file-system encryption on RE0.

re1 (Optional) Enable file-system encryption on RE1.

routing-engine (Optional) Enable file-system encryption on the specified Routing Engine. Use one of the following options to specify the Routing Engine:

Backup	Backup Routing Engine
Both	Both Routing Engines
Local	Local Routing Engine
Master	Primary Routing Engine

Other	Other Routing Engine

Required Privilege Level

maintenance

Sample Output

request system filesystem encryption enable (Junos OS)

```
user@host> request system filesystem encryption enable
```

```
Please check the message on enabling filesystem encryption enable using the dry-run option. i.e.
request system filesystem encrypton enable dry-run
Do you want to proceed ? [yes,no] (no) yes

Enable filesystem encryption ? [yes,no] (no) yes



A vmhost reboot is required to start filesystem encryption.
```

During the conversion process, the vmhost reboot using `request vmhost reboot` is required to start file-system encryption and to reflect the changes.

request system filesystem encryption enable (Junos OS Evolved)

```
user@host> request system filesystem encryption enable
```

```



Please check the message on enabling filesystem encryption enable using the dry-run option. i.e.
'request system filesystem encrypton enable dry-run'
Do you want to proceed ? [yes,no] (no) yes

-------------------------------
node: re0
-------------------------------
Removing version junos-evo-install-ptx-x86-64-23.4R2.14-EVO...
```

```
Removing external packages for junos-evo-install-ptx-x86-64-23.4R2.14-EVO...
Done.
Done.

A reboot is required to start the encryption.
Issue 'request system reboot' command when ready.
```

Release Information

Command introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Encryption with Trusted Platform Module

request system filesystem encryption keys delete | 15082

show system filesystem encryption status | 26852

request system filesystem encryption keys delete

IN THIS SECTION

- Syntax | 15082
- Description | 15083
- Options | 15083
- Required Privilege Level | 15083
- Sample Output | 15083
- Release Information | 15085

Syntax

```
request system filesystem encryption keys delete
<dry-run | re0 | re1 | routing-engine>
```

Description

Key slots are the Keys used to encrypt an area or slot on the hard disk. You must enable file-system encryption on the disk before deleting the key slots. After deleting the key slots, reboot the vmhost to reflect the changes. You need to use USB based reinstall or recovery method to restore the system.

Options

none Erase all keyslots on all Routing Engines.

dry-run (Optional) Display the keyslot deletion message without deleting any keyslots.

re0 (Optional) Erase all keyslots on RE0.

re1 (Optional) Erase all keyslots on RE1.

routing-engine (Optional) Erase all keyslots on the specified Routing Engine. Use one of the following options to specify the Routing Engine:

- **Backup** Backup Routing Engine
- **Both** Both Routing Engines
- **Local** Local Routing Engine
- **Master** Primary Routing Engine
- **Other** Other Routing Engine

Required Privilege Level

maintenance

Sample Output

request system filesystem encryption keys delete (Junos OS)

```
user@host> request system filesystem encryption keys delete
Please check the message on deleting filesystem encryption keyslots using the dry-run option.
i.e.

request system filesystem encrypton keys delete dry-run

Do you want to proceed ? [yes,no] (no) yes
```

```
Delete filesystem encryption keyslots ? [yes,no] (no) yes

Erasing keyslots on LUKS2 device /dev/nvme0n1p5

Erasing keyslots on LUKS2 device /dev/nvme1n1p5

All keyslots successfully erased, a vmhost reboot is needed for changes to take effect.
```

```
user@host> request vmhost reboot
```

request system filesystem encryption keys delete (Junos OS Evolved)

```
user@host> request system filesystem encryption keys delete no-forwarding routing-engine local
```

```
Please check the message on deleting filesystem encryption keyslots using the dry-run option.
i.e.
'request system filesystem encrypton keys delete dry-run'
Do you want to proceed ? [yes,no] (no) yes

-------------------------------
node: re0
-------------------------------
Erasing keyslots on LUKS2 device /dev/sda3
Erasing keyslots on LUKS2 device /dev/sda5
Erasing keyslots on LUKS2 device /dev/sda6
Erasing keyslots on LUKS2 device /dev/sda7
Erasing keyslots on LUKS2 device /dev/sdb3
Erasing keyslots on LUKS2 device /dev/sdb5
Erasing keyslots on LUKS2 device /dev/sdb6
Erasing keyslots on LUKS2 device /dev/sdb7

All keyslots successfully erased, a system reboot is needed for changes to take effect.
```

```
user@host> request system reboot
```

Release Information

Command introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Encryption with Trusted Platform Module

request system filesystem encryption enable | 15079

show system filesystem encryption status | 26852

request system firmware downgrade optics

IN THIS SECTION

- Syntax | 15085
- Description | 15086
- Options | 15086
- Required Privilege Level | 15086
- Output Fields | 15086
- Sample Output | 15087

Syntax

```
request system firmware downgrade optics
[
<part-number>
 <all>
fpc-slot <fpc_slot_number>
pic-slot <pic_slot_number>
port <port_number>
]
```

Description

Use this command to downgrade the firmware for a specific QSFP-DD optic module/all QSFP-DD modules plugged in the router on Junos OS Evolved.

Options

<*part-number*> (Junos OS Evolved only, for PTX Series routers) Downgrade optics for a specific QSFP-DD optic module part number .

all (Junos OS Evolved only, for PTX Series routers) Downgrade optics for all optic modules on device.

fpc (Junos OS Evolved only, for PTX Series routers) Downgrade FPC optics module.

- `fpc<fpc_slot_number>`—(Optional) Downgrade optics module for a specific FPC slot.

pem Upgrade PEM firmware.

pic-slot (Junos OS Evolved only, for PTX Series routers) Downgrade optics module for PIC slot.

- `pic-slot <pic_slot_number>`— Downgrade optics module for a specific PIC slot.

port Downgrade optics on ports.

- `port<port_number>`—(Optional) Downgrade optics on specific port.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system firmware downgrade optics

```
user@host> request system firmware downgrade optics
Firmware optics module downgrade initiated. Optics downgrade takes about 10 minutes
Use 'show system firmware optics' to get the status
```

RELATED DOCUMENTATION

request system firmware upgrade | 15089

request system firmware reload (Junos OS Evolved)

IN THIS SECTION

- Syntax | 15087
- Description | 15088
- Options | 15088
- Required Privilege Level | 15088
- Output Fields | 15088
- Release Information | 15088

Syntax

```
request system firmware reload
<fpc slot slot-number>
<sib slot slot-number>
```

Description

> **NOTE**: Use extreme caution when using this command. You can cause significant damage to your configuration if you use this command incorrectly.

Use this command to reload firmware after you issue the `request system firmware upgrade` operational mode command. Requesting a firmware reload avoids having to power cycle the line card or SIB after you perform a firmware upgrade.

The target FPC or SIB must be offline for the reload firmware command to take effect.

Options

`fpc slot` *`slot-number`* Reload FPC firmware.

- `slot` *`slot-number`*—Specify the particular FPC slot (0-7) that should reload its firmware.

`sib slot` *`slot-number`* Reload SIB firmware.

- `slot` *`slot-number`*—Specify the particular SIB slot (0-5) that should reload its firmware.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Release Information

Command introduced in Junos OS Evolved Release 21.4R1 for the PTX Series routers with the LC1201 and LC1202 line cards.

RELATED DOCUMENTATION

request system firmware upgrade | 15089

request system firmware upgrade

IN THIS SECTION

- Syntax | 15089
- Description | 15090
- Options | 15090
- Required Privilege Level | 15094
- Output Fields | 15094
- Sample Output | 15094
- Platform-Specific Behavior | 15097
- Release Information | 15098

Syntax

```
request system firmware upgrade

    <bios <progress>
    <cb <(fancpld | fpga | optics | port-fpga)>>
    <cpld cpu <progress>>
    <cpld sys <progress>>
    <eusb <progress>>
    <fpc <(bcm-pfe | dpll | dpll-cfg | opticscpld<0|1|2|3> | slot slot-number)>>
    <fpm>
    <ftc slot (0 | 1)>
    <mmc <progress>>
    <pem slot slot-number mcu (primary |secondary)>
    <poe fpc-slot slot-number>
    <poe file jfirmware filename fpc-slot (number | all-numbers) <poe-at-firmware|poe-bt-
firmware>
    <re <(bios | fancpld | fpga | i210 | i40nvm | ssd (disk1 | disk2) | xmcfpga)>>
    <sfb slot slot-number tag tag-number>
    <uboot <progress>>
    <vcpu>
```

Description

Use this command to upgrade firmware and optics module on a system running Junos OS.

Options

The following table lists all possible options for the *request system firmware upgrade* command. The fields that are valid may vary depending on the platform and software release.

Use Feature Explorer to confirm platform and release support for specific features.

Review the "Platform-Specific Behavior" on page 15097 section for notes related to your platform.

bios Upgrade BIOS.

- `progress`—(Optional) Check the progress of BIOS upgrade.

cb Upgrade baseboard FPGA.

- `fancpld`—(Optional) Upgrade fanboard CPLD.
- `fpga`—(Optional) Upgrade baseboard FPGA.
- `optics`—(Optional) Upgrade optics CPLD.
- `port-fpga`—(Optional) Upgrade baseboard PORT FPGA.

cpld cpu Upgrade CPU CPLD.

- `progress`—(Optional) Check the progress of CPU CPLD upgrade.

cpld sys Upgrade system CPLD.

- `progress`—(Optional) Check the progress of system CPLD upgrade.

eusb Upgrade eUSB firmware.

- `progress`—(Optional) Check the progress of eUSB firmware upgrade.

fpc Upgrade FPC ROM monitor.

- bcm-pfe—(Optional) Upgrade BCM PFE chip.
- dpll—(Optional) Upgrade FPC DPLL firmware.
- dpll-cfg—(Optional) Upgrade FPC DPLL configuration.

- opticscpld *<0|1|2|3>*—(Optional) Upgrade optics CPLD.
  - opticscpld0—Upgrade optics CPLD 0.
  - opticscpld1—Upgrade optics CPLD 1.
  - opticscpld2—Upgrade optics CPLD 2.
  - opticscpld3—Upgrade optics CPLD 3.

  > **NOTE**: Run the `show system firmware` command to verify the number of optics CPLDs supported on your device.

- slot *slot-number*—(Optional) Upgrade all devices in a particular FPC slot.

  After you upgrade the firmware on the LC4800 or the LC9600 line card, the line card may go offline. If this happens, use the `request chassis fpc slot-number restart` command to restart the line card.

fpm

Upgrade front panel module firmware.

ftc slot (0 | 1)

Upgrade fan tray controller firmware.

mmc

Upgrade the eMMC firmware.

- `progress`—(Optional) Check the progress of eMMC upgrade.

pem slot *slot-number* mcu (primary | secondary)

Upgrade PEM firmware.

Within the JNP10K-PWR-AC3 module, there are four types of MCUs; primary, secondary, communication, and LED. This module contains four primary MCUs; however, you can operate the device with three or fewer input leads attached. Before upgrading the firmware on this module, however, you must reattach all of the unattached input leads to ensure that all of the primary MCUs are upgraded to the same firmware version. If you do not reattach an unattached lead, the primary MCU corresponding to that lead is not upgraded, and the PEM remains offline with an alarm of `firmware mismatch`. This mismatch causes a corresponding PEM output shutdown, causing the PEM to appear offline to avoid issues caused by different firmware versions. If you did not ensure that all of the leads were attached before upgrading and

now have the alarm `firmware mismatch`, attach all of the leads and repeat the firmware upgrade to clear the alarm.

```
root@host-re0> request system firmware upgrade pem slot 0
WARNING: Please ensure all inputs are connected to this PEM; otherwise, PEM
will undergo a partial upgrade and go offline.
Perform indicated firmware upgrade? [yes, no] (no)
```

pic

Upgrade PIC firmware.

poe fpc-slot *slot-number*

Upgrade Power over Ethernet (PoE) firmware.

poe file *<jfirmware filename>* fpc-slot *(number | all-numbers)* <poe-at-firmware | poe-bt-firmware>

Upgrade the PoE controller firmware for the Virtual Chassis member or line card specified by `number`, or for all Virtual Chassis members and line cards, specified by `all-members`.

- poe-at-firmware—Rollback the PoE firmware to IEEE 802.3at (PoE-at).
- poe-bt-firmware—Upgrade the PoE firmware to IEEE 802.3bt (PoE-bt).

re

Upgrade baseboard BIOS/FPGA. There is an active BIOS image and a backup BIOS image.

- bios—(Optional) Upgrade BIOS.
- ctrl-fpga—(Optional) Upgrade Control FPGA.
- fancpld—(Optional) Upgrade fanboard CPLD.
- fpga—(Optional) Upgrade baseboard FPGA.
- i210-ha0—(Optional) Upgrade Intel I210-HA0 GbE NIC.
- i210-ha1—(Optional) Upgrade Intel I210-HA1 GbE NIC.
- i210—(Optional) Upgrade baseboard i210 GbE NIC.
- i40nvm—(Optional) Upgrade baseboard i40.

  Starting in Junos OS Release 19.3R1, you can upgrade the i40e NVM firmware on routers with VM Host support.

- ssd—(Optional) Upgrade Routing Engine solid-state drive (SSD) firmware.

- disk1—Upgrade SSD disk1 firmware.
- disk2—Upgrade SSD disk2 firmware.

  Starting in Junos OS Release 17.2R1, you can upgrade the SSD firmware on routers with the VM Host support.

- ucd—(Optional) Upgrade RE UCD firmware.
- xmcfpga—(Optional) Upgrade XMC FPGA.

> **NOTE**: To upgrade the `re fpga` component from version 0x0108 to 0x010A on QFX5130-32CD, the upgrade must be performed twice since the first upgrade command upgrades the component to version 0x0109.

sfb slot *slot-number* tag *tag-number*

Upgrade the SF2 and SFB2 switch fabric firmware. *slot-number* can be 0 to 5. For `tag` *`tag-number`* option, specify the tag number that indicates you want to update the FPGA. To find out what number you should use for the `tag` option, issue the `show system firmware` command. For example, the `show system firmware` command shows the tag numbers in the third column as follows:

```
user@host> show system firmware
Part           Type             Tag  Current             Available
Status
                                       version             version
[output truncated]
...
SFB 0          FPGA PRIM         0   0.13.0               0.13.0
OK
SFB 1          FPGA PRIM         0   0.13.0               0.13.0
OK
SFB 2          FPGA PRIM         0   0.13.0               0.13.0
OK
SFB 3          FPGA PRIM         0   0.13.0               0.13.0
OK
SFB 4          FPGA PRIM         0   0.13.0               0.13.0
OK
SFB 5          FPGA PRIM         0   0.13.0               0.13.0
OK
```

```
...
[output truncated]
```

After you upgrade the firmware on the SFB, you must take the SFB offline by using the `request chassis sfb slot slot-number offline` command. Once the SFB is offline, bring the SFB back online and make the new firmware take effect by using the `request chassis sfb slot slot-number online-reload` command.

uboot <progress>

Upgrade the U-boot firmware.

- `progress`—(Optional) Check the progress of U-boot upgrade.

vcpu

Upgrade VCPU ROM monitor.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided with the resulting status of the command.

Sample Output

request system firmware upgrade

```
user@host> request system firmware upgrade re bios
Part             Type           Tag Current   Available Status
                                    version   version
Routing Engine 0 RE BIOS        0   1.5       1.9       OK
Routing Engine 0 RE BIOS Backup 1   1.7       1.9       OK
Perform indicated firmware upgrade ? [yes,no] (no) yes
user@host> request system firmware upgrade re bios backup
Part             Type           Tag Current   Available Status
                                    version   version
Routing Engine 0 RE BIOS        0   1.5       1.9       OK
Routing Engine 0 RE BIOS Backup 1   1.7       1.9       OK
Perform indicated firmware upgrade ? [yes,no] (no) yes

user@host> request system firmware upgrade re ssd disk1
Part    Type    Tag                    Current   Available       Status
```

```
                                    version   version
Routing Engine   0 RE SSD1    4      12028     12029        OK
Perform indicated firmware upgrade ? [yes,no] (no) yes

Firmware upgrade initiated, use "show system firmware"  to monitor status.

user@host> request system firmware upgrade pem slot 0
...
...
Firmware upgrade initiated, use "show system firmware"  to monitor status.

user@host> request system firmware upgrade optics fpc-slot 0
...
...
Firmware optics upgrade initiated, use "show system firmware"  to monitor status.
```

request system firmware upgrade (SRX1600, SRX2300, and SRX4300)

```
user@host> request system firmware upgrade re ssd ?
Possible completions:
  disk1                SSD Disk1
  disk2                SSD Disk2
Perform indicated firmware upgrade ? [yes,no] (no) yes
user@host> request system firmware upgrade re ssd disk1
Part             Type              Tag Current               Available             Status
                                       version               version
Routing Engine 0 SSD1              3   40004                 40007                 OK
Perform indicated firmware upgrade ? [yes,no] (no) yes
user@host> request system firmware upgrade re ssd disk2
Part             Type              Tag Current               Available             Status
                                       version               version
Routing Engine 0 SSD2              3   40004                 40007                 OK
Perform indicated firmware upgrade ? [yes,no] (no) yes
user@host> request system firmware upgrade re ctrl-fpga
Part             Type              Tag Current               Available             Status
                                       version               version
Routing Engine 0 CTRL-FPGA         11  1.21                  1.29                  OK
Perform indicated firmware upgrade ? [yes,no] (no) yes
user@host> request system firmware upgrade re ucd
Part             Type              Tag Current               Available             Status
                                       version               version
```

```
Routing Engine 0 UCD              12  0.2                  0.3                  OK
Perform indicated firmware upgrade ? [yes,no] (no) yes
user@host> request system firmware upgrade re i210
Part             Type             Tag Current              Available            Status
                                      version              version
Routing Engine 0 I210             7   3.30.1908            3.30.2032
OK
Routing Engine 0 I210-HA0         8   3.30.2047            3.30.2047
OK
Routing Engine 0 I210-HA1         9   3.30.2047            3.30.2047                 OK
Perform indicated firmware upgrade ? [yes,no] (no) yes
user@host> request system firmware upgrade re bios
Part             Type             Tag Current              Available            Status
                                      version              version
Routing Engine 0 BIOS            1    0.23.4               0.24.1               OK
Perform indicated firmware upgrade ? [yes,no] (no) yes
user@host> request system firmware upgrade re fpga
Part             Type             Tag Current              Available            Status
                                      version              version
Routing Engine 0 FPGA             2   1.71                 1.75                 OK
Perform indicated firmware upgrade ? [yes,no] (no) yes
```

request system firmware upgrade cpld sys (EX4400-48XP and EX4400-48MXP)

```
user@host> request system firmware upgrade jfirmware cpld sys
cpld sys firmware upgrade initiated
Check progress using "request system firmware upgrade jfirmware <type> progress"

user@host> request system firmware upgrade jfirmware cpld sys progress
cpld sys firmware upgrade complete
Use "request system reboot" to initiate remote power cycle to install the firmware
```

request system firmware upgrade fpc slot

```
user@host> request system firmware upgrade fpc slot 0
...
...
FPC may go offline after the upgrade, Please restart FPC post upgrade.
```

```
“request chassis fpc <slot> restart” command can be used for restarting the fpc.
```

request system firmware upgrade sfb

```
user@host> request system firmware upgrade sfb slot 0 tag 0
...
...
"Firmware upgrade initiated, use "show system firmware" to monitor status. After upgrade, do
"request chassis sfb slot <slot> offline" and “request chassis sfb slot <slot> online-reload”
for new firmware to take effect.
```

Platform-Specific Behavior

Use Feature Explorer to confirm platform and release support for specific features.

Use the following table to review platform-specific behaviors for your platform:

Platform	Difference
MX Series	• MX10004 and MX10008 support `cb`, `fpm`, `ftc slot (0 | 1)` and `sfb slot slot-number tag tag-number` • MX10004, MX10008, MX10016, and MX304 support `pem slot slot-number mcu (primary | secondary)` **NOTE**: The `mcu` option upgrades the firmware on one micro controller unit at a time, applies only to the MX304 router, and is required for the MX304 PEM firmware upgrade.

(Continued)

Platform	Difference
EX Series Switches	• EX4100 support `uboot <progress>` and `eusb` • EX4100 and EX4400 support `bios`, `cpld sys`, and `poe file <jfirmware filename>fpc-slot (number | all-numbers) <poe-at-firmware | poe-bt-firmware>` • EX4100 and EX4400 support `mmc`, and `cpld cpu` • EX4300-48MP support `poe-at-firmware`
SRX Series	• SRX1600, SRX2300, and SRX4300 does not support backup BIOS image.
PTX Series (Junos OS)	• PTX10008 and PTX10016 support `pem slot slot-number mcu (primary | secondary)`
QFX Series Switches (Junos OS)	• QFX10008 and QFX10016 support `pem slot slot-number mcu (primary | secondary)`

Release Information

Command introduced in Junos OS Release 10.2.

`cb` option added in Junos OS Evolved Releases 21.1R2 and 21.2R1. Support for the MX10004 router added in Junos OS Release 22.3R1.

`pem` option introduced in Junos OS Release 21.2R1.

`sfb` option introduced in Junos OS Release 21.4R1 for the MX10008 router. Support for the MX10004 router added in Junos OS Release 22.3R1.

`mcu` option introduced in Junos OS Release 22.2R1-S1 and 22.3R1 for the MX304 router. `opticscpld`, and `fancpld` options introduced in Junos OS Evolved Release 23.4R1-S1 for QFX5130-48C devices.

`ctrl-fpga`, `i210-ha0`, `i210-ha1`, and `ucd` options introduced in Junos OS Release 23.4R1 for SRX1600 and SRX2300 devices.

`ctrl-fpga`, `i210-ha0`, `i210-ha1`, and `ucd` options introduced in Junos OS Release 24.1R1 for SRX4300 device.

`dpll`, and `dpll-cfg` options introduced in Junos OS Evolved Release 23.4R1-S1 for QFX5130-48C.

`dpll`, and `dpll-cfg` options introduced in Junos OS Evolved Release 23.4R2 for QFX5130-48CM.

`dpll`, and `dpll-cfg` options introduced in Junos OS Evolved Release 23.4R2 for QFX5230-64CD.

RELATED DOCUMENTATION

No Link Title

request system firmware upgrade (Junos OS Evolved)

IN THIS SECTION

- Syntax | 15099
- Description | 15100
- Options | 15100
- Required Privilege Level | 15101
- Sample Output | 15102
- Release Information | 15102

Syntax

```
request system firmware upgrade
<fpc <(bcm-pfe | dpll | dpll-cfg | opticscpld<0|1|2|3> | slot slot-number)>>
<ftc slot (0 | 1)>
<pic>
<psm <slot slot-number>>
<re <(bios | fancpld | fpga | ssd (disk1 | disk2))>>
<sib <slot slot-number>>
```

Description

After installing the `jfirmware-junos-evo-install*` package with the `request system software add` command, use this command to upgrade firmware on a system running Junos OS Evolved. You can only upgrade one option at a time; you cannot specify multiple options.

> **NOTE**: For firmware upgrades on Junos OS Evolved, we recommend that you work with the Juniper Networks Technical Assistance Center (JTAC).

Options

The fields that are valid may vary depending on the platform and software release.

fpc Upgrade FPC ROM monitor.

- bcm-pfe—(Optional) Upgrade BCM PFE chip.
- dpll—(Optional) Upgrade FPC DPLL firmware.
- dpll-cfg—(Optional) Upgrade FPC DPLL configuration.
- opticscpld *<0|1|2|3>*—(Optional) Upgrade optics CPLD.
  - opticscpld0—Upgrade optics CPLD 0.
  - opticscpld1—Upgrade optics CPLD 1.
  - opticscpld2—Upgrade optics CPLD 2.
  - opticscpld3—Upgrade optics CPLD 3.

  > **NOTE**: Run the `show system firmware` command to verify the number of optics CPLDs supported on your device.

- slot *slot-number*—(Optional) Upgrade all devices in a particular FPC slot.

ftc slot (0 | 1) Upgrade fan tray controller firmware.

pic Upgrade PIC firmware.

psm Upgrade power supply module firmware.

Within the JNP10K-PWR-AC3 module, there are four types of MCUs; primary, secondary, communication, and LED. This module contains four primary MCUs; however, you can operate the device with three or fewer input leads attached. Before upgrading the firmware on this module, however, you must reattach all of the unattached input leads to ensure that all of the primary MCUs are upgraded to the same firmware version. If you do not reattach an unattached lead, the primary MCU corresponding to that lead is not upgraded, and the PSM remains offline with an alarm of `firmware mismatch`. This mismatch causes a corresponding PSM output shutdown, causing the PSM to appear offline to avoid issues caused by different firmware versions. If you did not ensure that all of the leads were attached before upgrading and now have the alarm `firmware mismatch`, attach all of the leads and repeat the firmware upgrade to clear the alarm.

```
root@host-re0> request system firmware upgrade psm slot 0
WARNING: Please ensure all inputs are connected to this PSM; otherwise, PSM will
undergo a partial upgrade and go offline.
Perform indicated firmware upgrade? [yes, no] (no)
```

- `slot` *`slot-number`*—(Optional) Upgrade a particular power supply module.

re (bios | fancpld | fpga | ssd (disk1 | disk2))

Upgrade baseboard BIOS/FPGA. There is an active BIOS image and a backup BIOS image.

- bios—(Optional) Upgrade BIOS.
- fancpld—(Optional) Upgrade fanboard CPLD.
- fpga—(Optional) Upgrade baseboard FPGA.
- ssd—(Optional) Upgrade Routing Engine solid-state drive (SSD) firmware.
  - disk1—Upgrade SSD disk1 firmware.
  - disk2—Upgrade SSD disk2 firmware.

> **NOTE**: To upgrade the `re fpga` component from version 0x0108 to 0x010A on QFX5130-32CD, the upgrade must be performed twice, since the first upgrade command upgrades the component to version 0x0109.

sib

Upgrade SIB firmware.

Required Privilege Level

maintenance

Sample Output

name-of-sample-output-command

```
user@host> command-name option1 option2
```

Release Information

Command introduced in Junos OS Evolved ????.

request system firmware upgrade poe

IN THIS SECTION

- Syntax | 15102
- Description | 15103
- Options | 15104
- Required Privilege Level | 15104
- Output Fields | 15104
- Sample Output | 15105
- Release Information | 15105

Syntax

```
request system firmware upgrade poe fpc-slot (number | all-members) <poe-at-firmware | poe-bt-
firmware>
```

Description

Upgrade the PoE controller software on switches and line cards.

The Junos OS image running on the switch contains a copy of the PoE controller software. This command compares the Junos OS version with the version of the software running on the PoE controller. If the Junos OS version is a more recent version, the command downloads the more recent version to the controller.

For all Virtual Chassis except EX8200 Virtual Chassis, execute this command on the primary. The primary itself need not support PoE for this command to work—for example, you can execute this command on the primary of a mixed EX4200 and EX4500 Virtual Chassis when the primary is an EX4500 switch, which does not support PoE. On an EX8200 Virtual Chassis, you must execute this command on the member switch, not the primary XRE200 External Routing Engine.

We recommend that all members of a Virtual Chassis run the same version of the PoE controller software.

Upgrading the controller software can take up to 10 minutes. Use the `show poe controller` command to monitor the progress of the software download.

> **NOTE**: When you enter the `request system firmware upgrade poe` command, a message advises you that the controller software upgrade has started and that it will take about 10 minutes to complete. This message appears even if the FPC you have specified does not have a PoE controller or if the PoE controller software is up-to-date. To determine whether or not the controller software upgrade has actually started, use the `show poe controller` command.

> **NOTE**: While the upgrade is in progress, power to the powered devices is not guaranteed. We recommend that you upgrade the controller software during a regularly scheduled maintenance window.

If you load a version of the Junos OS software that is not compatible with the PoE firmware version, a major alarm will be raised. See Table 129 on page 15104 for firmware compatibility.

Table 129: PoE Firmware Compatibility

PoE Firmware Version	Junos OS Release	
	Prior to 19.3R1	19.3R1 and later
2.xx	PoE-at is supported. POE-bt is not available.	PoE-at is supported. PoE-bt is not available without firmware upgrade.
3.xx	PoE-at is not supported. PoE-bt is not available. **NOTE**: Firmware downgrade is required for PoE.	PoE-at is supported. PoE-bt is supported.

Options

fpc-slot (*number* | all-members)	Upgrade the PoE controller firmware for the Virtual Chassis member or line card specified by *number*, or for all Virtual Chassis members and line cards, specified by `all-members`.
poe-at-firmware	(EX4300-48MP switches only) Rollback the PoE firmware to IEEE 802.3at (PoE-at).
poe-bt-firmware	(EX4300-48MP switches only) Upgrade the PoE firmware to IEEE 802.3bt (PoE-bt).

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system firmware upgrade poe (Specific FPC Slot)

```
user@switch> request system firmware upgrade poe fpc-slot 8
Firmware upgrade initiated. Poe Upgrade takes about 10 minutes
Use 'show poe controller' to get the download status
```

Release Information

Command introduced in Junos OS Release 12.1.

Support for `poe-at-firmware` and `poe-bt-firmware` options introduced in 19.3 for EX4300-48MP switches.

RELATED DOCUMENTATION

show poe controller | 23171

Upgrading the PoE Controller Software

Upgrading to IEEE 802.3bt PoE

request system halt

IN THIS SECTION

- Syntax | 15106
- Syntax (EX Series Switches) | 15106
- Syntax (PTX Series) | 15107
- Syntax (MX Series Router) | 15107
- Syntax (QFX Series) | 15107
- Syntax (SRX Series) | 15108
- Description | 15108
- Options | 15108
- Additional Information | 15110

- Required Privilege Level | 15110
- Output Fields | 15110
- Sample Output | 15110
- Release Information | 15112

Syntax

```
request system halt
<at time>
<backup-routing-engine>
<both-routing-engines>
<other-routing-engine>
<in minutes>
<media (compact-flash | disk | removable-compact-flash | usb)>
<message "text">
```

Syntax (EX Series Switches)

```
request system halt
<all-members>
<at time>
<backup-routing-engine>
<both-routing-engines>
<in minutes>
<local>
<media (external | internal)>
<member member-id>
<message "text">
<other-routing-engine>
<slice slice>
```

Syntax (PTX Series)

```
request system halt
<at time>
<backup-routing-engine>
<both-routing-engines>
<other-routing-engine>
<in minutes>
<media (compact-flash | disk)>
<message "text">
```

Syntax (MX Series Router)

```
request system halt
<all-members>
<at time>
<backup-routing-engine>
<both-routing-engines>
<in minutes>
<local>
<media (external | internal)>
<member member-id>
<message "text">
<other-routing-engine>
```

Syntax (QFX Series)

```
request system halt
<all-members>
<at time>
<both-routing-engines>
<director-device director-device-id>
<in minutes>
<local>
<media >
<member member-id>
<message "text">
```

```
<other-routing-engine>
<slice slice>
```

Syntax (SRX Series)

```
request system halt
<at time>
<in minutes>
<media (compact-flash | disk | usb)>
<message "text">
```

Description

Stop the device software.

> **NOTE**: When you issue this command on an individual component—for example, a Node device—in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member of a Node group, issue this command with the `member` option on the Node device CLI, because you cannot issue this command from the QFabric CLI. Also, issuing this command might cause traffic loss on an individual component.
>
> When you issue this command on a QFX5100 switch, you are not prompted to reboot. You must power cycle the switch to reboot.

> **NOTE**: For the routers with the Routing Engines RE-S-2x00x6, RE-PTX-2x00x8, and RE-S-2x00x8, this command is deprecated and might be removed completely in a future release.
>
> On these routers, this command is replaced with the `request vmhost halt` command which provides similar functionality.

Options

none Stop the router or switch software immediately.

all-members (Optional) Halt all members of the Virtual Chassis configuration.

at *time* (Optional) Time at which to stop the software, specified in one of the following ways:

- `now`—Stop the software immediately. This is the default.
- `+minutes`—Number of minutes from now to stop the software.
- `yymmddhhmm`—Absolute time at which to stop the software, specified as year, month, day, hour, and minute.
- `hh:mm`—Absolute time on the current day at which to stop the software.

backup-routing-engine
(Optional) Halt the backup Routing Engine. This command halts the backup Routing Engine, regardless from which Routing Engine the command is executed. For example, if you issue the command from the primary Routing Engine, the backup Routing Engine is halted. If you issue the command from the backup Routing Engine, the backup Routing Engine is halted.

both-routing-engines
(Optional) Halt both Routing Engines at the same time.

director-device *director-device-id*
(QFabric systems only) Halt a specific Director device.

local
(Optional) Halt the local Virtual Chassis member.

in *minutes*
(Optional) Number of minutes from now to stop the software. This option is an alias for the at +*minutes* option.

media (compact-flash /disk)
(Optional) Boot medium for the next boot.

media (external | internal)
(EX Series and QFX Series switches and MX Series routers only) (Optional) Halt the boot media:

- `external`—Halt the external mass storage device.
- `internal`—Halt the internal flash device.

media (compact-flash | disk | usb)
(SRX Series only) (Optional) Boot media for the next boot.

- `compact-flash`— Standard boot from a flash device.
- `disk`— Boot from a hard disk.
- `usb`— Boot from a USB device.

member *member-id*
(Optional) Halt the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, `member-id` can only be 0 or 1.

message "*text*"
(Optional) Message to display to all system users before stopping the software.

other-routing-engine (Optional) Halt the other Routing Engine from which the command is issued. For example, if you issue the command from the primary Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the primary Routing Engine is halted.

slice *slice* (EX Series and QFX Series switches only) (Optional) Halt a partition on the boot media. This option has the following suboptions:

- 1—Halt partition 1.
- 2—Halt partition 2.
- `alternate`—Reboot from the alternate partition.

Additional Information

On the M7i router, the `request system halt` command does not immediately power down the Packet Forwarding Engine. The power-down process can take as long as 5 minutes.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system halt

```
user@host> request system halt
Halt the system ? [yes,no] (no) yes

*** FINAL System shutdown message from root@section2 ***
System going down IMMEDIATELY
Terminated
...
syncing disks... 11 8 done
```

```
The operating system has halted.
Please press any key to reboot.
```

request system halt (SRX Series)

```
user@host> request system halt
Halt the system ? [yes,no] (no) yes

*** FINAL System shutdown message from user@host ***

System going down IMMEDIATELY



Shutdown NOW!
[pid 7560]

user@host> Dec  8 08:57:37 Waiting (max 60 seconds) for system process `vnlru' to stop...done
Waiting (max 60 seconds) for system process `vnlru_mem' to stop...done
Waiting (max 60 seconds) for system process `bufdaemon' to stop...done
Waiting (max 60 seconds) for system process `syncer' to stop...
Syncing disks, vnodes remaining...2 2 2 2 2 2 2 2 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 done

syncing disks... All buffers synced.
Uptime: 2d16h25m9s
recorded reboot as normal shutdown

The operating system has halted.
Please press any key to reboot.
```

request system halt (In 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request that the system stop 2 hours from now:

```
user@host> request system halt at +120
user@host> request system halt in 120
user@host> request system halt at 19:00
```

request system halt (Immediately)

```
user@host> request system halt at now
```

request system halt (At 1:20 AM)

To stop the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system halt at yymmdd120
request system halt at 120
Halt the system at 120? [yes,no] (no) yes
```

Release Information

Command introduced before Junos OS Release 7.4.

`other-routing-engine` option introduced in Junos OS Release 8.0.

`director-device` option introduced for QFabric systems in Junos OS Release 12.2.

`backup-routing-engine` option introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

- clear system reboot
- request system power-off
- request vmhost halt

request-system-integrity-attestation

IN THIS SECTION

Syntax | 15113

- Description | 15113
- Options | 15113
- Required Privilege Level | 15113
- Output Fields | 15114
- Sample Output | 15114
- Release Information | 15114

Syntax

```
request system integrity attestation (pcr-index pcr number | nonce-value | tpm-name)
```

Description

The Attester is a device on the network that provides evidence to the Verifier on demand. The first part of that evidence is a signed collection of the TPM's Platform Configuration Registers (PCRs) called a TPM Quote. To ensure the freshness of the quote, the Verifier sends a random nonce to the Attester to include in the TPM Quote.

Options

nonce-value Defines a base-64 encoded 20 byte random hex value. A cryptographically generated random number which should not be predictable prior to its issuance from a random number generation function. The random number MUST be derived from an entropy source external to the Attester.

pcr-index The numbers/indexes of the PCRs. At the moment this is limited to 32 items. (0..31)

tpm-name Unique system generated name for a TPM on a device.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system integrity attestation nonce-value $NONCE tpm-name $SERIAL_NUMBER

```
user@host> request system integrity attestation nonce-value $NONCE tpm-name $SERIAL_NUMBER
```

Release Information

Command introduced in Junos OS Release 22.4R1.

request-system-integrity-log-retrieval

IN THIS SECTION

- Syntax | 15114
- Description | 15114
- Options | 15115
- Required Privilege Level | 15115
- Release Information | 15115

Syntax

```
request system integrity log-retrieval log-type
```

Description

The remaining evidence provided to the Verifier on demand is comprised of the boot event log recorded by the TPM. The Verifier will compare this log against the TPM Quote for validity. If the log is valid, the

Verifier will check each entry in the log against a known-good entry. Note that as of this release, only the bios log option is supported.

Options

bios Specifies an event type associated with BIOS/UEFI.

ima Specifies an event type recorded in IMA.

netequip_boot Specifies an event type associated with Network Equipment Boot.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 22.4R1.

request system license update

IN THIS SECTION

- Syntax | 15116
- Description | 15116
- Options | 15116
- Required Privilege Level | 15116
- Output Fields | 15116
- Sample Output | 15116
- Release Information | 15117

Syntax

```
request system license update
```

Description

Starts autoupdating license keys from the license portal.

- The **request system license update** command always uses the default Juniper license server: **https://ae1.juniper.net/**.
- The **request system license update** command is supported only on SRX, vSRX Virtual Firewall, and QFX Series devices.

The products supported by the Juniper Agile Licensing (JAL) portal includes: QFX series, SRX Series, EX Series, NFX, vBNG, vMX, vSRX Virtual Firewall, and ACX. For other Juniper products (SPACE, JSA, SBR Carrier, Screen OS and so on) access the License Management System (LMS).

Options

`trial`—Immediately updates trial license keys from the license portal.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system license update

```
user@host> request system license update

Trying to update license keys from https://ae1.juniper.net has been sent, use show system
license to check status.
```

request system license update trial

```
user@host> request system license update trial

Request to automatically update trial license keys from https://ae1.juniper.net has been sent,
use show system license to check status.
```

Release Information

Command introduced in Junos OS Release 9.5.

request system logout

IN THIS SECTION

- Syntax | 15117
- Description | 15117
- Options | 15118
- Required Privilege Level | 15118
- Output Fields | 15118
- Sample Output | 15118
- Release Information | 15118

Syntax

```
request system logout (pid pid / terminal terminal / user username)
<all>
```

Description

Log out users from the router or switch and the configuration database. If a user held the `configure exclusive` lock, this command clears the exclusive lock.

Options

all	(Optional) Log out all sessions owned by a particular PID, terminal session, or user. (On a TX Matrix or TX Matrix Plus router, this command is broadcast to all chassis.)
pid *pid*	Log out the user session using the specified management process identifier (PID). The PID type must be management process.
terminal *terminal*	Log out the user for the specified terminal session.
user *username*	Log out the specified user.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system logout

```
user@host> request system logout user test all
Connection closed by foreign host.
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Log a User Out of the Device

request system malware-scan

IN THIS SECTION

- Syntax | 15119
- Description | 15119
- Options | 15119
- Required Privilege Level | 15120
- Sample Output | 15121
- Release Information | 15121

Syntax

```
request system malware-scan
integrity-check
quick-scan
  clean-action (clean | warn)
  pids set-of-pids
  test
run-fake-malware | list-fake-malware
```

Description

Run the Juniper Malware Removal Tool (JMRT), which scans for and removes malware running on Juniper Networks devices. This command can perform multiple types of scans, detailed in the following section.

Options

integrity-check Check whether integrity mechanisms are enabled and working properly. Integrity mechanisms allow only signed binaries to run on the system.

> **NOTE**: Junos OS with Junos Automation Enhancements does not run Veriexec. As such, running the `integrity-check` command on Junos OS with Junos Automation Enhancements always shows that integrity is not enforced.

quick-scan Start a quick scan, which attempts to scan each process' executable for malware. If the executable file does not exist, it scans the process memory instead.

clean-action (clean | warn) Determine the action that JMRT should take when potential malware is detected:

- `clean`—Remove infected files and processes. This is the default action.
- `warn`—Notify the user of files and processes containing malware, but do not remove them.

pids *set-of-pids* Set of process IDs (PIDs) to scan. By default, JMRT scans all processes.

You can access the following options by default in Junos OS Evolved. To access them in Junos OS, install the `jmrt-test` package.

> **NOTE**: You will see the `run-fake-malware` and `list-fake-malware` options in Junos OS only if you have installed `jmrt-test`.

> **NOTE**: Before 23.3, you must enable the `jmrt-test` package to gain access to the `test` option for `quick-scan`, `run-fake-malware`, and `list-fake-malware`.

test Run a test scan that will detect fake malware. Use this option to observe how JMRT works, without needing actual malware on the system.

> **NOTE**: This option is deprecated from Junos 24.2R2
> Now the `quick-scan` option is sufficient to scan for fake-malware.

run-fake-malware Spawn a fake malware process on the system.

list-fake-malware List all fake malware processes (PIDs) that are running on the system.

Required Privilege Level

admin

Sample Output

request system malware-scan

```
user@host> request system malware-scan integrity-check
Integrity is enforced: No
```

```
user@host> request system malware-scan list-fake-malware
```

```
user@host> request system malware-scan quick-scan
Found potential malware: No
```

```
user@host> request system malware-scan run-fake-malware
Fake malware PID: 28423
```

Release Information

Command introduced in Junos OS Release 19.2R1.

Command option `veriexec-check` renamed as `integrity-check` in Junos OS Release 21.4R1.

Command introduced in Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

Juniper Malware Removal Tool

request system software add

request system partition abort

IN THIS SECTION

- Syntax | 15122
- Syntax (TX Matrix Router) | 15122
- Syntax (TX Matrix Plus Router) | 15122
- Syntax (MX Series Router) | 15123
- Description | 15123
- Options | 15123
- Required Privilege Level | 15124
- Output Fields | 15124
- Sample Output | 15124
- Release Information | 15124

Syntax

```
request system partition abort
```

Syntax (TX Matrix Router)

```
request system partition abort
<all-chassis | all-lcc | lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
request system partition abort
<all-chassis | all-lcc | lcc number | sfc number>
```

Syntax (MX Series Router)

```
request system partition abort
<all-members>
<local>
<member member-id>
```

Description

Terminate a previously scheduled storage media partition operation. If the command is issued between the time of a partition request and a reboot, the partition request is terminated and the storage media is not affected.

Options

all-chassis (TX Matrix and TX Matrix Plus routers only) (Optional) Terminate a previously scheduled partition operation for all chassis.

all-lcc (TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, terminate a previously scheduled partition operation on all T640 routers (line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, terminate a previously scheduled partition operation on all routers (line-card chassis) connected to the TX Matrix Plus router.

all-members (MX Series routers only) (Optional) Terminate a previously scheduled partition operation for all members of the Virtual Chassis configuration.

lcc *number* (TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix Plus router, terminate a previously scheduled partition operation on a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, terminate a previously scheduled partition operation on a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Terminate a previously scheduled partition operation for the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Terminate a previously scheduled partition operation for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

scc (TX Matrix routers only) (Optional) Terminate a previously scheduled partition operation on the TX Matrix router (or switch-card chassis).

sfc *number* (TX Matrix Plus routers only) (Optional) Terminate a previously scheduled partition operation on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system partition terminate

```
user@host> request system partition abort
The hard disk is no longer scheduled to be partitioned.
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Command deprecated for Junos OS with Upgraded FreeBSD in Junos OS Release 15.1.

NOTE: To determine which platforms run Junos OS with Upgraded FreeBSD, see the table listing the platforms currently running Junos OS with upgraded FreeBSD in Release Information for Junos OS with Upgraded FreeBSD.

RELATED DOCUMENTATION

request system partition hard-disk | 15127

request system partition compact-flash

IN THIS SECTION

- Syntax | 15125
- Description | 15125
- Required Privilege Level | 15126
- Output Fields | 15126
- Sample Output | 15126
- Sample Output | 15126
- Release Information | 15126

Syntax

```
request system partition compact-flash
```

Description

Use this command to reboot the device and repartition the compact flash. The CompactFlash card is repartitioned only if it is possible to restore all the data on the CompactFlash card. Otherwise, the operation is terminated, and a message is displayed indicating that the current disk usage needs to be reduced.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system partition compact-flash (If Yes)

```
user@host> request system partition compact-flash
Are you sure you want to reboot
and partition the compact-flash ? [yes,no] yes
Initiating repartition operation.
The operation may take several minutes to complete.
System will reboot now...
<System reboots>
<Repartition operation is performed>
<System reboots and starts up normally>
```

Sample Output

request system partition compact-flash (If No)

```
user@host> request system partition compact-flash
Are you sure you want to reboot
and partition the compact-flash ? [yes,no] no
```

Release Information

Command introduced in Junos OS Release 9.2.

Command deprecated for Junos OS with Upgraded FreeBSD in Junos OS Release 15.1.

NOTE: To determine which platforms run Junos OS with Upgraded FreeBSD, see the table listing the platforms currently running Junos OS with upgraded FreeBSD in Release Information for Junos OS with Upgraded FreeBSD.

request system partition hard-disk

IN THIS SECTION

- Syntax | 15127
- Syntax (TX Matrix Router) | 15127
- Syntax (TX Matrix Plus Router) | 15128
- Syntax (MX Series Router) | 15128
- Description | 15128
- Options | 15128
- Additional Information | 15129
- Required Privilege Level | 15129
- Output Fields | 15129
- Sample Output | 15130
- Release Information | 15130

Syntax

```
request system partition hard-disk
```

Syntax (TX Matrix Router)

```
request system partition hard-disk
<all-chassis | all-lcc | lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
request system partition hard-disk
<all-chassis | all-lcc | lcc number | sfc number>
```

Syntax (MX Series Router)

```
request system partition hard-disk
<all-members>
<local>
<member member-id>
```

Description

Set up the hard disk for partitioning. After this command is issued, the hard disk is partitioned the next time the system is rebooted. When the hard disk is partitioned, the contents of **/altroot** and **/altconfig** are saved and restored. All other data on the hard disk is at risk of being lost.

Options

all-chassis (TX Matrix routers and TX Matrix Plus routers only) (Optional) Schedule a partition of the hard disk for all routers in the chassis at its next reboot.

all-lcc (TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, schedule a partition of the hard disk on all T640 routers connected to the TX Matrix router at their next reboot. On a TX Matrix Plus router, schedule a partition of the hard disk on all connected LCCs.

all-members (MX Series routers only) (Optional) Schedule a partition of the hard disk for all members of the Virtual Chassis configuration.

lcc *number* (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix Plus router, schedule a partition of the hard disk on a specific T640 router connected to the TX Matrix router. On a TX Matrix Plus router, schedule a partition of the hard disk on a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Schedule a partition of the hard disk for the local member of the Virtual Chassis.

member *member-id* (MX Series routers only) (Optional) Schedule a partition of the hard disk for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

scc (TX Matrix routers only) (Optional) Schedule a partition of the hard disk on the T640 router connected to the TX Matrix router (or switch-card chassis).

sfc *number* (TX Matrix Plus routersonly) (Optional) Schedule a partition of the hard disk on the connected T1600 or T4000 LCCs connected to the TX Matrix Plus router . Replace *number* with 0.

Additional Information

To immediately partition the hard disk, use the `request system reboot` command. To cancel the partition request, use the `request system partition abort` command.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system partition hard-disk

```
user@host> request system partition hard-disk
WARNING:   The hard disk is about to be partitioned.  The contents
WARNING:   of /altroot and /altconfig will be saved and restored.
WARNING:   All other data is at risk.  This is the setup stage, the
WARNING:   partition happens during the next reboot.

Setting up to partition the hard disk ...

WARNING:   A REBOOT IS REQUIRED TO PARTITION THE HARD DISK.  Use the
WARNING:   'request system reboot' command when you are ready to proceed
WARNING:   with the partitioning.  To abort the partition of the hard disk
WARNING:   use the 'request system partition abort' command.
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Command deprecated for Junos OS with Upgraded FreeBSD in Junos OS Release 15.1.

NOTE: To determine which platforms run Junos OS with Upgraded FreeBSD, see the table listing the platforms currently running Junos OS with upgraded FreeBSD in Release Information for Junos OS with Upgraded FreeBSD.

RELATED DOCUMENTATION

request system partition abort | **15122**

Routing Matrix with a TX Matrix Plus Router Solutions Page

request system power-off

IN THIS SECTION

- Syntax | 15131
- Syntax (EX Series Switches) | 15131
- Syntax (MX Series Router) | 15132
- Syntax (QFX Series) | 15132
- Description | 15132
- Options | 15133
- Required Privilege Level | 15134
- Output Fields | 15135
- Sample Output | 15135
- Release Information | 15135

Syntax

```
request system power-off
<both-routing-engines>
<other-routing-engine>
<at time>
<in minutes>
<media (compact-flash | disk | removable-compact-flash | usb)>
<message "text">
```

Syntax (EX Series Switches)

```
request system power-off
<at time>
<both-routing-engines>
<in minutes>
<junos>
<message "text">
<network>
```

```
<oam>
<other-routing-engine>
<usb>
```

Syntax (MX Series Router)

```
request system power-off
<at time>
<both-routing-engines>
<in minutes>
<junos>
<message "text">
<network>
<no-forwarding>
<oam>
<other-routing-engine>
<usb>
```

Syntax (QFX Series)

```
request system power-off
<all-members>
<at time>
<both-routing-engines>
<in minutes>
<local>
<member>
<message "text">
<other-routing-engine>
<slice slice>
```

Description

Power off the Routing Engines.

NOTE: When you issue this command on an individual component in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos

Routing Engines will reboot." If you want to halt only one member, use the `member` option. You cannot issue this command from the QFabric CLI.

NOTE: For a standalone chassis (such as MX Series and PTX Series), the request to power off the system is applicable only to the Routing Engines. When you request to power off both Routing Engines, all the FPCs in the chassis shut down after approximately 10 minutes and the chassis fans run at full speed. The FPCs shut down because they no longer have communication with the Routing Engines and an Inter-Integrated Circuit (I2C) timeout occurred.

NOTE: For the devices with Routing Engines RE-S-2x00x6, RE-PTX-2x00x8, and RE-S-2x00x8, this command is deprecated and might be removed completely in a future release.

On these devices, this command is replaced with the `request vmhost power-off` command which provides similar functionality.

Options

none — Power off the router or switch software immediately.

all-members — (QFX Series devices only) (Optional) Power off all members of the Virtual Chassis configuration.

at *time* — (Optional) Time at which to power off the software, specified in one of the following ways:

- `now`—Power off the software immediately. This is the default.
- `+`*`minutes`*—Number of minutes from now to power off the software.
- *`yymmddhhmm`*—Absolute time at which to power off the software, specified as year, month, day, hour, and minute.
- *`hh:mm`*—Absolute time on the current day at which to power off the software.

both-routing-engines — (Optional) Power off both Routing Engines at the same time.

in *minutes* — (Optional) Number of minutes from now to power off the software. This option is an alias for the `at +`*`minutes`* option.

junos — (EX Series and MX Series devices only) (Optional) Boot off Junos volume.

local	(EX4200 switches and MX Series devices only) (Optional) Power off the local Virtual Chassis member.
media (compact-flash | disk | removable-compact-flash | usb)	(Optional) Boot medium for the next boot.
media (external | internal)	(EX Series, MX Series, and QFX Series devices only) (Optional) Power off the boot media: • `external`—Power off the external mass storage device. • `internal`—Power off the internal flash device.
member *member-id*	(EX4200 and MX Series devices only) (Optional) Power off the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *`member-id`* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *`member-id`* with a value of 0 or 1.
message "*text*"	(Optional) Message to display to all system users before powering off the software.
no-forwarding	(MX series devices only) (Optional) No forwarding.
network	(EX Series and MX Series devices only) (Optional) Network boot through PXE.
oam	(Optional) Boot off OAM volume.
other-routing-engine	(Optional) Power off the other Routing Engine from which the command is issued. For example, if you issue the command from the primary Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the primary Routing Engine is halted.
slice *slice*	(EX Series and QFX Series switches only) (Optional) Power off a partition on the boot media. This option has the following suboptions: • `1`—Power off partition 1. • `2`—Power off partition 2. • `alternate`—Reboot from the alternate partition.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system power-off

```
user@host> request system power-off message “This router will be powered off in 30 minutes.
Please save your data and log out immediately.”
warning: This command will not halt the other routing-engine.
If planning to switch off power, use the both-routing-engines option.
Power Off the system ? [yes,no] (no) yes

*** FINAL System shutdown message from remote@nutmeg ***
System going down IMMEDIATELY

This router will be powered off in 30 minutes. Please save your data and log out
immediately.

Shutdown NOW!
[pid 5177]
```

Release Information

Command introduced in Junos OS Release 8.0.

request system power-off (SRX Series)

IN THIS SECTION

- Syntax | 15136
- Description | 15136
- Options | 15136
- Required Privilege Level | 15137

- Output Fields | 15137
- Sample Output | 15137
- Release Information | 15138

Syntax

```
request system power-off
at <time>
in <minutes>
junos
message <message>
network
oam
power-off
usb
media (compact-flash | disk | usb | internal)
```

Description

Use this command to power off the system.

Options

at *time*
Time at which to power off the system.

in *minutes*
Number of minutes to delay before powering off the system.

media
Boot media for the next boot.

- `compact-flash`— Standard boot from a flash device.
- `disk`— Boot from a hard disk.
- `usb`— Boot from a USB device.
- `internal`— Boot from internal flash.

message *message*
Message that is displayed to all system users before powering off the system.

junos	(SRX1500, SRX5400, SRX5600, and SRX5800) Boot off Junos volume.
network	(SRX1500, SRX5400, SRX5600, and SRX5800) Network boot through PXE.
oam	(SRX1500, SRX5400, SRX5600, and SRX5800) Boot off OAM volume.
usb	(SRX1500, SRX5400, SRX5600, and SRX5800) Boot off USB device.
power-off	(SRX1500) Power off the software on RE.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system power-off

```
user@host> request system power-off
Power Off the system ? [yes,no] (no) yes

Shutdown NOW!
[pid 3300]


*** FINAL System shutdown message from user@host ***

System going down IMMEDIATELY


user@host> Dec  8 09:37:45 Waiting (max 60 seconds) for system process `vnlru' to stop...done
Waiting (max 60 seconds) for system process `vnlru_mem' to stop...done
Waiting (max 60 seconds) for system process `bufdaemon' to stop...done
Waiting (max 60 seconds) for system process `syncer' to stop...
Syncing disks, vnodes remaining...2 2 2 2 2 2 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 done

syncing disks... All buffers synced.
```

```
Uptime: 38m33s
recorded reboot as normal shutdown

The operating system has halted.
Turning the system power off.
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

request system halt | 15105

request system process terminate

IN THIS SECTION

- Syntax | 15138
- Description | 15139
- Options | 15139
- Required Privilege Level | 15139
- Output Fields | 15139
- Sample Output | 15139
- Release Information | 15139

Syntax

```
request system process terminate process-id
```

Description

Terminate any process that you specify with the process identification number (process ID or pid). The `request system process terminate` command is an alternative to using the `restart` command. The `restart` command terminates and restarts a process that you specify by process name, but limits you to only certain well-known processes.

> **CAUTION**: Caution: Never terminate a software process unless instructed to do so by a customer support engineer. Restarting processes could cause unknown system behavior resulting in partial or complete traffic loss.

Use the `show system processes` command to display a list of processes by process ID.

Options

process-id Identification number for a process.

- **Range:** 1 through 99999

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are not provided feedback on the status of your request. You can use the `show system processes extensive` command to confirm the process that was terminated.

Sample Output

request system process terminate

```
user@host> request system process terminate 1514
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show system processes | 26957

Junos OS Routing Processes

request system reboot (Junos OS)

IN THIS SECTION

- Syntax | 15140
- Syntax (EX Series Switches and EX Series Virtual Chassis) | 15141
- Syntax (MX Series Routers and MX Series Virtual Chassis, EX9200 Switches and EX9200 Virtual Chassis) | 15141
- Syntax (QFabric Systems) | 15141
- Syntax (QFX Series Switches and QFX Series Virtual Chassis, Virtual Chassis Fabric) | 15142
- Syntax (TX Matrix Router) | 15142
- Syntax (TX Matrix Plus Router) | 15142
- Description | 15143
- Options | 15143
- Additional Information | 15147
- Required Privilege Level | 15147
- Output Fields | 15147
- Sample Output | 15148
- Release Information | 15150

Syntax

```
request system reboot
<at time>
<both-routing-engines>
<in minutes>
<media (compact-flash | disk | removable-compact-flash | usb)>
```

```
<message "text">
<other-routing-engine>
```

Syntax (EX Series Switches and EX Series Virtual Chassis)

```
request system reboot
<all-members | local | member member-id>
<at time>
<in minutes>
<media (external | internal)> | <media (compact-flash | disk | removable-compact-flash | usb)>
<message "text">
<slice slice>
```

Syntax (MX Series Routers and MX Series Virtual Chassis, EX9200 Switches and EX9200 Virtual Chassis)

```
request system reboot
<all-members | local | member member-id>
<at time>
<both-routing-engines>
<in minutes>
<media (external | internal)> | <media (compact-flash | disk | usb)> | <junos | network | oam |
usb>
<message "text">
<other-routing-engine>
```

Syntax (QFabric Systems)

```
request system reboot
<all <graceful>>
<at time>
<director-device name>
<director-group <graceful>>
<fabric <graceful>>
<in minutes>
<in-service>
<media>
<message “text”>
```

```
<node-group name>
<slice slice>
```

Syntax (QFX Series Switches and QFX Series Virtual Chassis, Virtual Chassis Fabric)

```
request system reboot
<all-members | local | member member-id>
<at time>
<in minutes>
<in-service>
<hypervisor>
<junos | network | oam | usb>
<message “text”>
<slice slice>
```

Syntax (TX Matrix Router)

```
request system reboot
<all-chassis | all-lcc | lcc number /  scc>
<at time>
<both-routing-engines>
<in minutes>
<media (compact-flash | disk)>
<message "text">
<other-routing-engine>
```

Syntax (TX Matrix Plus Router)

```
request system reboot
<all-chassis | all-lcc | lcc number | sfc number>
<at time>
<both-routing-engines>
<in minutes>
<media (compact-flash | disk)>
<message "text">
<other-routing-engine>
<partition (1 | 2 | alternate)>
```

Description

Use this command to reboot the device software.

This command can be used on standalone devices and on devices supported in a Virtual Chassis, Virtual Chassis Fabric, or QFabric system.

Starting with Junos OS Release 15.1F3, the `request system reboot` command reboots only the guest operating system on the PTX5000 with RE-PTX-X8-64G and, MX240, MX480, and MX960 with RE-S-X6-64G.

Starting with Junos OS Release 15.1F5, the `request system reboot` command reboots only the guest operating system on the MX2010, and MX2020 with REMX2K-X8-64G.

Starting from Junos OS Release 17.2R1, PTX10008 routers do not support the `request system reboot` command. Starting from Junos OS Release 17.4R1, PTX10016 routers do not support the `request system reboot` command. Use the `request vmhost reboot` command instead of the `request system reboot` command on the PTX10008 and PTX10016 routers to reboot the Junos OS software package or bundle on the router. See request vmhost reboot.

Starting from Junos OS Release 19.1R1, the PTX10002-60C router and the QFX10002-60C switch do not support the `request system reboot` command. Use the `request vmhost reboot` command instead of the `request system reboot` command on these devices to reboot the Junos OS software package or bundle on the device. See request vmhost reboot.

On a QFabric system, to avoid traffic loss on the network Node group, switch mastership of the Routing Engine to the backup Routing Engine, and then reboot.

Options

The options described here are not all supported on every platform or release of Junos OS. Refer to the Syntax sections for the options commonly available on each type of platform.

none Reboot the software immediately.

> **NOTE**: On EX Series and QFX Series switches, the device waits for one minute before rebooting. To reboot the device immediately, use the `request system reboot at now` command.

all-chassis (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all routers connected to the TX Matrix or TX Matrix Plus router, respectively.

all-lcc

(Optional) On a TX Matrix router or TX Matrix Plus router, reboot all line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively.

all-members | local | member *member-id*

(Optional) Specify which member of the Virtual Chassis to reboot:

- `all-members`—Reboots each switch that is a member of the Virtual Chassis.
- `local`—Reboots only the local switch (switch where you are logged in).
- `member` *member-id*—Reboots the specified member switch of the Virtual Chassis

at *time*

(Optional) Time at which to reboot the software, specified in one of the following ways:

- `now`—Stop or reboot the software immediately. This is the default.
- *+minutes*—Number of minutes from now to reboot the software.
- *yymmddhhmm*—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.
- *hh:mm*—Absolute time on the current day at which to stop the software, specified in 24-hour time.

both-routing-engines

(Optional) Reboot both Routing Engines at the same time.

fast-boot

(Optional, QFX Series) Minimizes traffic loss and downtime of network ports by leaving the network ports up during the system reboot.

hypervisor

(Optional) Reboot Junos OS, host OS, and any installed guest VMs.

in *minutes*

(Optional) Number of minutes from now to reboot the software. The minimum value is 1. This option is an alias for the `at +`*minutes* option.

in-service

(Optional) Enables you to reset the software state (no software version change) of the system with minimal disruption in data and control traffic.

junos

(Optional) Reboot from the Junos OS (main) volume.

lcc *number*

—(Optional) Line-card chassis (LLC) number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.

- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

media (compact-flash | disk | removable-compact-flash | usb)

(Optional) Use the indicated boot medium for the next boot.

media (external | internal)

(Optional) Use the indicated boot medium for the next boot:

- `external`—Reboot the device using a software package stored on an external boot source, such as a USB flash drive.
- `internal`—Reboot the device using a software package stored in an internal memory source.

message "*text*"

(Optional) Message to display to all system users before stopping or rebooting the software.

network

(Optional) Reboot using the Preboot Execution Environment (PXE) boot method over the network.

oam

(Optional) Reboot from the maintenance volume (OAM volume, usually the compact flash drive).

other-routing-engine

(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the primary Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the primary Routing Engine is rebooted.

partition *partition*

(Optional) Reboot using the specified partition on the boot media. This option is equivalent to the `slice` option that is supported on some devices. Specify one of the following *partition* values:

- `1`—Reboot from partition 1.
- `2`—Reboot from partition 2.
- `alternate`—Reboot from the alternate partition.

scc (Optional) Reboot the Routing Engine on the TX Matrix switch-card chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted.

sfc *number* (Optional) Reboot the Routing Engine on the TX Matrix Plus switch-fabric chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted. Replace *number* with 0.

slice *slice* (Optional) Reboot using the specified partition on the boot media. This option was originally the `partitiion` option but was renamed to `slice` on EX Series and QFX Series switches. Specify one of the following *slice* values:

- `1`—Reboot from partition 1.
- `2`—Reboot from partition 2.
- `alternate`—Reboot from the alternate partition (which did not boot the switch at the last bootup).

> **NOTE**: The `slice` option is not supported on QFX Series switches that have no alternate slice when Junos OS boots as a Virtual Machine (VM). To switch to the previous version of Junos OS, issue the `request system software rollback` command.

usb (Optional) Reboot from a USB device.

The following options are available only on QFabric Systems:

all (Optional) Reboots the software on the Director group, fabric control Routing Engines, fabric manager Routing Engines, Interconnect devices, and network and server Node groups.

director-device *name* (Optional) Reboots the software on the Director device and the default partition (QFabric CLI).

director-group (Optional) Reboots the software on the Director group and the default partition (QFabric CLI).

fabric (Optional) Reboots the fabric control Routing Engines and the Interconnect devices.

node-group *name* (Optional) Reboots the software on a server Node group or a network Node group.

graceful (Optional) Enables the QFabric component to reboot with minimal impact to network traffic. This sub-option is only available for the `all`, `fabric`, and`director-group` options.

Additional Information

Reboot requests are recorded in the system log files, which you can view with the `show log` command (see "show log" on page 21980). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the `show system processes` command (see "show system processes" on page 26957).

On a TX Matrix or TX Matrix Plus router, if you issue the `request system reboot` command on the primary Routing Engine, all the primary Routing Engines connected to the routing matrix are rebooted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are rebooted.

> **NOTE**: Before issuing the `request system reboot` command on a TX Matrix Plus router with no options or the `all-chassis`, `all-lcc`, `lcc` *number*, or `sfc` options, verify that primary Routing Engine for all routers in the routing matrix are in the same slot number. If the primary Routing Engine for a line-card chassis is in a different slot number than the primary Routing Engine for a TX Matrix Plus router, the line-card chassis might become logically disconnected from the routing matrix after the `request system reboot` command.

> **NOTE**: To reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the primary Routing Engine.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system reboot

```
user@host> request system reboot
Reboot the system ? [yes,no] (no)
```

request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes

shutdown: [pid 186]
*** System shutdown message from root@test.example.net ***
System going down at 23:00
```

request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
 user@host> request system reboot in 120
 user@host> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@host> request system reboot at now
```

request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

request system reboot in-service

```
user@switch> request system reboot in-service
Reboot the system ? [yes,no] yes
[Feb 22 02:37:04]:ISSU: Validating Image

 PRE ISSR CHECK:
 ---------------
 PFE Status                           : Online
 Member Id zero                       : Valid
 VC not in mixed or fabric mode       : Valid
 Member is single node vc             : Valid
 BFD minimum-interval check done      : Valid
 GRES enabled                         : Valid
 NSR enabled                          : Valid
 drop-all-tcp not configured          : Valid
 Ready for ISSR                       : Valid

warning: Do NOT use /user during ISSR. Changes to /user during ISSR may get lost!
Current image is jinstall-jcp-i386-flex-18.1.img
[Feb 22 02:37:14]:ISSU: Preparing Backup RE
Prepare for ISSR
[Feb 22 02:37:19]:ISSU: Backup RE Prepare Done
Spawning the backup RE
Spawn backup RE, index 1 successful
Starting secondary dataplane
Second dataplane container started
GRES in progress
Waiting for backup RE switchover ready
```

```
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade for ISSR
Chassis ISSU Started
[Feb 22 02:42:55]:ISSU: Preparing Daemons
[Feb 22 02:43:00]:ISSU: Daemons Ready for ISSU
[Feb 22 02:43:05]:ISSU: Starting Upgrade for FRUs
[Feb 22 02:43:15]:ISSU: FPC Warm Booting
[Feb 22 02:44:16]:ISSU: FPC Warm Booted
[Feb 22 02:44:27]:ISSU: Preparing for Switchover
[Feb 22 02:44:31]:ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online (ISSU)
Send ISSR done to chassisd on backup RE
Chassis ISSU Completed
Removing dcpfe0 eth1 128.168.0.16 IP
Bringing down bme00
Post Chassis ISSU processing done
[Feb 22 02:44:33]:ISSU: IDLE
Stopping primary dataplane
Clearing ISSU states
Console and management sessions will be disconnected. Please login again.
device_handoff successful ret: 0
Shutdown NOW!
[pid 14305]

*** FINAL System shutdown message from root@sw-duckhorn-01 ***

System going down IMMEDIATELY
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `other-routing-engine` introduced in Junos OS Release 8.0.

Option `sfc` introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Option `partition` changed to `slice` in Junos OS Release 10.0 for EX Series switches.

Option `both-routing-engines` introduced in Junos OS Release 12.1.

Option `fast-boot` introduced in Junos OS Release 14.1X53-D10 for QFX Series.

RELATED DOCUMENTATION

- request system halt | 15105
- Routing Matrix with a TX Matrix Plus Router Solutions Page
- request vmhost reboot | 15529
- clear system reboot | 4226

request system reboot (Junos OS Evolved)

IN THIS SECTION

- Syntax | 15151
- Description | 15151
- Options | 15152
- Additional Information | 15152
- Required Privilege Level | 15153
- Output Fields | 15153
- Sample Output | 15153
- Release Information | 15154

Syntax

```
request system reboot
<at time>
<in minutes>
<message "text">
```

Description

Reboot the entire system (all nodes).

To reboot a single node, use the `request node reboot` command instead.

Options

none Reboot the software immediately.

at *time* (Optional) Time at which to reboot the software, specified in one of the following ways:

- `now`—Stop or reboot the software immediately. This is the default.
- `+minutes`—Number of minutes from now to reboot the software.
- `hh:mm`—Absolute time on the current day at which to stop the software, specified in 24-hour time.

in *minutes* (Optional) Number of minutes from now to reboot the software. The minimum value is 1. This option is an alias for the `at +minutes` option.

message "*text*" (Optional) Message to display to all system users before stopping or rebooting the software.

Additional Information

Starting in Junos OS Evolved 23.4R1, when rebooting a system with this command, the Layer 2 applications `l2ald`, `l2ald-agent`, `l2cpd`, and `l2cpd-agent` are initialized and started only if any of the following configuration hierarchy levels contain any configuration statements:

- `[edit interface interface-name unit number family ethernet-switching]`
- `[edit vlans]`
- `[edit routing-instance instance-name instance-type virtual-switch]`
- `[edit routing-instance instance-name instance-type mac-vrf]`
- `[edit protocols l2-learning]`

Additionally, `l2cpd`, and `l2cpd-agent` are initialized and started if the `[edit protocols lldp]` hierarchy level contains any configuration statements.

As a result of this change, if your configuration already contains these configuration statements and you then delete all of the above configuration statements, these Layer 2 applications stop running.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system reboot

```
user@host-re0> request system reboot
The entire system (all nodes) will reboot causing traffic loss, do you wish to continue?
[yes,no] (no) yes

*** System shutdown message from root@host-re0 ***


reboot the system at Tue Dec  6 13:32:24 2022

System going down IMMEDIATELY
```

request system reboot at 23:00

```
user@host-re0> request system reboot at 23:00
The entire system (all nodes) will reboot causing traffic loss, do you wish to continue?
[yes,no] (no) yes

*** System shutdown message from root@host-re0 ***


reboot the system at Mon Dec  5 23:00:00 2022
```

request system reboot in 1

```
user@host-re0> request system reboot in 1
The entire system (all nodes) will reboot causing traffic loss, do you wish to continue?
```

```
[yes,no] (no) yes

*** System shutdown message from root@host-re0 ***


reboot the system at Mon Dec  5 15:08:00 2022
```

request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host-re0> request system reboot at +120
user@host-re0> request system reboot in 120
user@host-re0> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@host-re0> request system reboot at now
The entire system (all nodes) will reboot causing traffic loss, do you wish to continue?
[yes,no] (no) yes
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

request system reboot (Junos OS with Upgraded FreeBSD)

IN THIS SECTION

- Syntax | 15155
- Description | 15155
- Options | 15155
- Additional Information | 15156

- Required Privilege Level | 15157
- Output Fields | 15157
- Sample Output | 15157
- Release Information | 15158

Syntax

```
request system reboot
<at time>
<both-routing-engines>
<in minutes>
<junos>
<message "text">
<network>
<oam>
<other-routing-engine>
<usb>
```

Description

For devices running FreeBSD Release 10 or above, use this command to reboot the device software.

Options

none Reboot the device software immediately.

at *time* (Optional) Time at which to reboot the software, specified in one of the following ways:

- `now`—Stop or reboot the software immediately. This is the default.
- *+minutes*—Number of minutes from now to reboot the software.
- *yymmddhhmm*—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute. Omitting a value will default to the current date for that value.
- *hh:mm*—Absolute time on the current day at which to stop the software, specified in 24-hour time.

both-routing-engines	(Optional) Reboot both Routing Engines at the same time.
in *minutes*	(Optional) Number of minutes from now to reboot the software. This option is an alias for the `at +minutes` option.
junos	(Optional) Reboot from the **junos** volume.
message "*text*"	(Optional) Message to display to all system users before stopping or rebooting the software.
network	(Optional) Reboot from the network.
oam	(Optional) Reboot from the **oam** volume.
other-routing-engine	(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the primary Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the primary Routing Engine is rebooted.

usb	(Optional) Reboot from the USB device. From U-Boot release 3.12 and later and FreeBSD release 10 or later, the USB boot behaves differently than in prior releases.
	
	Therefore, to stop the device from booting from the USB in the future, you can either:
	
	• Unplug the USB stick and powercycle the device. Then from the U-Boot prompt, issue these commands:
	
	`setenv bootdev 0`
	`saveenv`
	`reset`
	
	• Or if the CLI is available, you can issue the `request system reboot junos` command.

Additional Information

Reboot requests are recorded in the system log files, which you can view with the `show log` command (see "show log" on page 21980). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the `show system processes` command (see show system processes).

> **NOTE**: To reboot a device that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the primary Routing Engine.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system reboot

```
user@host> request system reboot
Reboot the system ? [yes,no] (no)
```

request system reboot (at 2300)

```
user@host> request system reboot at 2300 message “Maintenance time!”
Reboot the system ? [yes,no] (no) yes

shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
 user@host> request system reboot in 120
 user@host> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@host> request system reboot at now
```

request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

Release Information

Command introduced in Junos OS Release 15.1.

Command introduced for all platforms using Junos OS with upgraded FreeBSD. To find which platforms in which releases use Junos OS with upgraded FreeBSD, see Feature Explorer, enter **freebsd**, and select **Junos kernel upgrade to FreeBSD 10+**.

RELATED DOCUMENTATION

- request system snapshot (Junos OS with Upgraded FreeBSD) | **15203**
- show system snapshot (Junos OS with Upgraded FreeBSD) | **27075**
- clear system reboot
- request system halt | **15105**

request system reboot (SRX Series)

IN THIS SECTION

- Syntax | 15159
- Description | 15159
- Options | 15159
- Required Privilege Level | 15160
- Release Information | 15160

Syntax

```
request system reboot <at time> <in minutes><media><message "text">
```

Description

Reboot the software.

Options

- at *time* (Optional)— Specify the time at which to reboot the device. You can specify time in one of the following ways:
  - now— Reboot the device immediately. This is the default.
  - *+minutes*— Reboot the device in the number of minutes from now that you specify.
  - *yymmddhhmm*— Reboot the device at the absolute time on the date you specify. Enter the year, month, day, hour (in 24-hour format), and minute.
  - *hh:mm*— Reboot the device at the absolute time you specify, on the current day. Enter the time in 24-hour format, using a colon (:) to separate hours from minutes.
- in *minutes*(Optional)— Specify the number of minutes from now to reboot the device. This option is a synonym for the at +minutes option
- media *type*(Optional)— Specify the boot device to boot the device from:

- disk/internal— Reboot from the internal media. This is the default.
- usb— Reboot from the USB storage device.
- compact flash— Reboot from the external CompactFlash card.

> **NOTE**: The `media` command option is not available on vSRX Virtual Firewall.

- message *"text"*(Optional)— Provide a message to display to all system users before the device reboots.

  Example: **request system reboot at 5 in 50 media internal message stop**

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 10.1.

Command `hypervisor` option introduced in Junos OS Release 15.1X49-D10 for vSRX Virtual Firewall.

RELATED DOCUMENTATION

request system software rollback (SRX Series)

request system reboot usb (SRX Series)

IN THIS SECTION

- Syntax | **15161**
- Description | **15161**
- Options | **15161**
- Required Privilege Level | **15161**
- Release Information | **15161**

Syntax

```
request system reboot usb <yes|no>
```

Description

Use this command to boot from the USB, without power cycling the SRX device. This command helps to detect USB with a soft reboot and provides the user with control over the boot sequence for soft reboot without turning off the device.

Options

- yes – Reboot system
- No – Do not reboot system

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 21.2R1.

request system recover

IN THIS SECTION

- Syntax | 15162
- Syntax (EX Series Switches) | 15162
- Description | 15162
- Options | 15162
- Additional Information | 15163
- Required Privilege Level | 15163

Sample Output | 15163
Release Information | 15165

Syntax

```
request system recover (junos-volume | oam-volume)
```

Syntax (EX Series Switches)

```
request system recover (junos-volume | oam-volume)
<(all-members | local | member member-id)>
```

Description

For devices running FreeBSD 10 or above, use this command to recover a specified volume of the system.

Guest network functions (GNFs) do not support the `recover` option under the `request system` command. See Components of Junos Node Slicing for more details on GNF.

Options

(all-members | local | member *member-id*) (EX Series Virtual Chassis only) (Optional) Specify the members who need to recover archive data and executable areas in a Virtual Chassis:

- `all-members`—Recover archive data and executable areas for all members of the Virtual Chassis.
- `local`—Recover archive data and executable areas for the member of the Virtual Chassis that you are currently logged into.
- `member` *member-id*—Recover archive data and executable areas for the specified member of the Virtual Chassis.

junos-volume Specify the **/junos** volume to be recovered. The **/junos** volume is the main drive and contains all the software and files needed for the day-to-day running of the device,

including configuration information and logs. The**/junos** volume also contains non-recovery snapshots, which are new with Junos OS with upgraded FreeBSD. Non-recovery snapshots cannot be used for recovery of a failed system.

oam-volume Specify the **/oam** volume to be recovered. The compact flash drive is the **/oam** volume and stores recovery snapshot backup information. In case of failure of the **/junos** volume, the **/oam** volume can be used to boot the system. The **/oam** volume has the recovery snapshot, which is created with the `request system snapshot recovery` command. A recovery snapshot is always replaced when a new recovery snapshot is taken.

Additional Information

If you try to recover the **/junos** volume while you are booted on the **/junos** volume, you will get an error message.

To recover the **/junos** volume, do the following:

1. Recover the **/oam** volume.

```
user@host> request system recover oam-volume
```

2. Reboot on the **/oam** volume.

```
user@host> request system reboot oam-volume
```

Required Privilege Level

view

Sample Output

request system recover junos-volume (While booted on the /junos volume)

```
user@host> request system recover junos-volume
ERROR: You are currently running on the Junos volume
ERROR: A recovery of the Junos volume is not possible
```

request system recover junos-volume (While booted on the /oam volume)

```
user@host> request system recover junos-volume
NOTICE: Recovering the Junos volume ...
ada0p3 deleted
ada0 created
ada0p1 added
bootcode written to ada0
ada0p3 added
ada0p2 added
/dev/gpt/junos: 20303.9MB (41582448 sectors) block size 32768, fragment size 4096
        using 33 cylinder groups of 626.22MB, 20039 blks, 80256 inodes.
super-block backups (for fsck_ffs -b #) at:
 192, 1282688, 2565184, 3847680, 5130176, 6412672, 7695168, 8977664, 10260160,
 11542656, 12825152, 14107648, 15390144, 16672640, 17955136, 19237632,
 20520128, 21802624, 23085120, 24367616, 25650112, 26932608, 28215104,
 29497600, 30780096, 32062592, 33345088, 34627584, 35910080, 37192576,
 38475072, 39757568, 41040064
NOTICE: Junos volume recovered
```

request system recover oam-volume

```
user@host> request system recover oam-volume
NOTICE: Recovering the OAM volume ...
ada1p2 deleted
ada1 created
ada1p1 added
bootcode written to ada1
ada1p2 added
/dev/gpt/oam: 3831.6MB (7847136 sectors) block size 32768, fragment size 4096
        using 7 cylinder groups of 626.09MB, 20035 blks, 80256 inodes.
super-block backups (for fsck_ffs -b #) at:
 192, 1282432, 2564672, 3846912, 5129152, 6411392, 7693632
Verified oam signed by PackageProductionEc_2017 method ECDSA256+SHA256
Installing OAM volume contents ...
The OAM volume is now installed
NOTICE: Creating a recovery snapshot on the OAM volume ...
Creating image ...
Compressing image ...
Image size is 1717MB
```

```
Recovery snapshot created successfully
NOTICE: OAM volume recovered
```

Release Information

Command introduced in Junos OS Release 15.1.

Command introduced for all platforms using Junos OS with upgraded FreeBSD. To find which platforms in which releases use Junos OS with upgraded FreeBSD, see Feature Explorer, enter **freebsd**, and select **Junos kernel upgrade to FreeBSD 10+**.

Options (`all-members` | `local` | `member` *`member-id`*) introduced in Junos OS Release 24.4R1.

request system restore state path

IN THIS SECTION

- Syntax | **15165**
- Description | **15165**
- Release Information | **15165**

Syntax

`request system restore state path` *`/host/tmp/ folder`*

Description

Use this command to restore your backup JDM configuration from the `/host/tmp/` folder.

Release Information

Command introduced in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

show virtual-network-functions (Junos Node Slicing) | **27728**

request server authenticate-peer-server | **14901**

request system save state path

IN THIS SECTION

- Syntax | **15166**
- Description | **15166**
- Release Information | **15166**

Syntax

```
request system save state path /host/tmp/ folder
```

Description

Use this command to back up your JDM configuration to the /host/tmp/ folder.

Release Information

Command introduced in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

show virtual-network-functions (Junos Node Slicing) | **27728**

request server authenticate-peer-server | **14901**

request system scripts add (Chassis Cluster)

IN THIS SECTION

- Syntax | 15167
- Description | 15167
- Options | 15167
- Additional Information | 15168
- Required Privilege Level | 15168
- Sample Output | 15168
- Release Information | 15171

Syntax

```
request system scripts add package-name <no-copy | unlink>
<master>
<backup>
```

Description

CLI command to install AI-Script install packages on SRX Series Firewalls in a chassis cluster.

Options

package-name	Specify AI-Script install package name.
no-copy	(Optional) Do not save a copy of the AI script package file.
unlink	(Optional) Remove the AI script package after successful installation.
master	(Optional) Install AI script packages on the primary node.
backup	(Optional) Install AI script packages on the secondary node.

Additional Information

This command eliminates the AI script installation on both primary node and secondary node separately.

Required Privilege Level

maintenance

Sample Output

request system scripts add package-name

```
user@host> request system scripts add jais-5.0R1.0-signed.tgz master



[: -a: unexpected operator
grep: /etc/db/pkg/jais/+COMMENT: No such file or directory
Installing package '/var/tmp/jais-5.0R4.0-signed.tgz' ...
Verified jais-5.0R4.0.tgz signed by PackageProductionRSA_2016
Adding jais...
Available space: 798414 require: 1814
Installing AI-Scripts version: 5.0R4
Saving package file in /var/db/scripts/commit/jais-5.0R4.0-signed.tgz ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
ln: ///etc/rc.d/ais: Read-only file system
Model: srx5600
Model: srx5600
Linking in Junos ES manifest file.
Creating srx5800/srx5600 trend data file.
Creating SRX intelligence attachments file.
Creating SRX events attachments file.
Creating AI-Scripts FIFO
Starting AI-Scripts FIFO handler
77834: old priority 0, new priority 20
77842: old priority 0, new priority 20
RSI parameters are now being set
BIOS validation parameter is now being set
BIOS interval parameter is now being set
JMB cleanup age is now being set
JMB Event file is now being set
JMB User Event file is now being set
```

```
PHDC collect parameter is now being set
PHDC duration parameter is now being set
PHDC commands file is now being set
JMB Progress Logging parameter is now being set
iJMB generation parameters are now being set
AI-Scripts platform support flag is now being set
Interval event commands file is now being set
Interval event enabled parameter is now being set
All node log collect parameter is now being set
Disk Warning Threshold is now being set
Disk Full Threshold is now being set
RSI Lite Enabled is now being set
Removing any old files that need to be updated
Copying updated files
Restarting eventd ...
Event processing process started, pid 78147
Installation completed
Saving package file in /var/sw/pkg/jais-5.0R4.0-signed.tgz ...
Saving state for rollback ...
```

request system scripts add package-name

```
user@host> request system scripts add jais-5.0R1.0-signed.tgz backup
Pushing bundle to node1
[: -a: unexpected operator
grep: /etc/db/pkg/jais/+COMMENT: No such file or directory
Installing package '/var/tmp/jais-5.0R4.0-signed.tgz' ...
Verified jais-5.0R4.0.tgz signed by PackageProductionRSA_2016
Adding jais...
Available space: 2619677 require: 1814
Installing AI-Scripts version: 5.0R4
chmod: /var/db/scripts/event/cron.slax: No such file or directory
chmod: /var/db/scripts/event/bit_event.slax: No such file or directory
chmod: /var/db/scripts/event/bit_event2.slax: No such file or directory
chmod: /var/db/scripts/op/ais_bit.slax: No such file or directory
Saving package file in /var/db/scripts/commit/jais-5.0R4.0-signed.tgz ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
ln: ///etc/rc.d/ais: Read-only file system
Mounted jais package on /dev/md2...
Verified manifest signed by PackageProductionRSA_2016
Verified jais-5.0R4.0 signed by PackageProductionRSA_2016
```

```
Model: srx5600
Model: srx5600
Linking in Junos ES manifest file.
Creating srx5800/srx5600 trend data file.
Creating SRX intelligence attachments file.
Creating SRX events attachments file.
Creating AI-Scripts FIFO
Starting AI-Scripts FIFO handler
99423: old priority 0, new priority 20
99428: old priority 0, new priority 20
99429: old priority 0, new priority 20
99430: old priority 0, new priority 20
RSI parameters are now being set
BIOS validation parameter is now being set
BIOS interval parameter is now being set
JMB cleanup age is now being set
JMB Event file is now being set
JMB User Event file is now being set
PHDC collect parameter is now being set
PHDC duration parameter is now being set
PHDC commands file is now being set
JMB Progress Logging parameter is now being set
iJMB generation parameters are now being set
AI-Scripts platform support flag is now being set
Interval event commands file is now being set
Interval event enabled parameter is now being set
All node log collect parameter is now being set
Disk Warning Threshold is now being set
Disk Full Threshold is now being set
RSI Lite Enabled is now being set
chmod: /var/db/scripts/event/cron.slax: No such file or directory
chmod: /var/db/scripts/event/bit_event.slax: No such file or directory
chmod: /var/db/scripts/event/bit_event2.slax: No such file or directory
chmod: /var/db/scripts/op/ais_bit.slax: No such file or directory
Removing any old files that need to be updated
Copying updated files
Restarting eventd ...
Event processing process started, pid 99730
Installation completed
Saving package file in /var/sw/pkg/jais-5.0R4.0-signed.tgz ...
Saving state for rollback ...
```

Release Information

Command introduced before Junos OS Release 9.0. The options `master` and `backup` are introduced in Junos OS Release 15.1X49-D50.

RELATED DOCUMENTATION

Understanding Automatic Chassis Cluster Synchronization Between Primary and Secondary Nodes

request system scripts add

IN THIS SECTION

- Syntax | 15171
- Description | 15171
- Options | 15172
- Required Privilege Level | 15172
- Release Information | 15172

Syntax

```
request system scripts add <package-name>
<no-copy>
<unlink>
```

Description

Use this command to install AI-Script (jais) packages on Juniper Networks devices.

Options

no-copy Don't save a copy of the jais package file.

```
user@host> request system scripts add no-copy <package-name>
```

If you use the no-copy option during the jais installation, the jais package cannot be rolled back.

unlink Remove the package after successful installation.

```
user@host> request system scripts add unlink <package-name>
```

Required Privilege Level

maintenance

Release Information

Command introduced before Junos OS Release 9.0.

RELATED DOCUMENTATION

- request system scripts delete | 15175
- request system scripts rollback | 15181
- request system scripts event-scripts reload | 15176

request system scripts convert

IN THIS SECTION

- Syntax | 15173
- Description | 15173
- Options | 15173

- Required Privilege Level | 15174
- Output Fields | 15174
- Sample Output | 15174
- Release Information | 15174

Syntax

```
request system scripts convert (slax-to-xslt | xslt-to-slax) source source/filename destination destination/
<filename> <partial> <version (1.0 | 1.1)>
```

Description

Convert an Extensible Stylesheet Language Transformations (XSLT) script to Stylesheet Language Alternative syntaX (SLAX), or convert a SLAX script to XSLT.

Options

destination *destination/<filename>*

Path and filename for the converted file.

Optionally, you can specify a filename for the converted file. If you do not specify a filename, the software assigns one automatically. The default destination filename is **SLAX-Conversion-Temp** or **slax-temp** depending on the Junos OS release, with a randomly generated, five-character, alpha-numeric extension. For example, the software converts a source file called **test.xsl** to **slax-temp.kWwQk**. The software converts a source file called **test1.slax** to **slax-temp.zN61h**.

partial

(Optional) Convert partial script input.

slax-to-xslt

Convert a SLAX script to XSLT.

source *source/filename*

Path and filename of the file to convert.

version

(Optional) SLAX version number that will be included in the version statement of the generated script for XSLT-to-SLAX conversions. Acceptable values are 1.0 and 1.1.

- **Default:** 1.1

xslt-to-slax Convert an XSLT script to SLAX.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system scripts convert slax-to-xslt

```
user@host> request system scripts convert slax-to-xslt source /var/db/scripts/op/script1.slax
destination /var/db/scripts/op
conversion complete
```

request system scripts convert xslt-to-slax

```
user@host> request system scripts convert xslt-to-slax source /var/db/scripts/commit/script1.xsl
destination /var/db/scripts/commit partial version 1.0
conversion complete
```

Release Information

Command introduced in Junos OS Release 8.2.

`partial` and `version` options added in Junos OS Release 12.2.

RELATED DOCUMENTATION

Converting Scripts Between SLAX and XSLT

request system scripts delete

IN THIS SECTION

- Syntax | 15175
- Description | 15175
- Options | 15175
- Required Privilege Level | 15175
- Release Information | 15175

Syntax

```
request system scripts delete <package-name>
```

Description

(Junos OS only) Use this command to delete AI-Script (jais) packages on Juniper Networks devices.

Options

No options are available.

Required Privilege Level

maintenance

Release Information

Command introduced before Junos OS Release 9.0.

NOTE: This command is not supported on Junos OS Evolved. AI-Scripts and Service Now are not supported on Junos OS Evolved.

RELATED DOCUMENTATION

request system scripts add | 15171

request system scripts rollback | 15181

request system scripts event-scripts reload | 15176

request system scripts event-scripts reload

IN THIS SECTION

- Syntax | 15176
- Description | 15176
- Options | 15176
- Required Privilege Level | 15177
- Output Fields | 15177
- Sample Output | 15177
- Release Information | 15177

Syntax

```
request system scripts event-scripts reload
```

Description

Reload all active event scripts into the eventd process' memory. This enables you to replace or update an existing event script on the device and refresh the event policy embedded in the script without changing the device's configuration or disrupting operations.

Options

reload	Reload all active event-scripts.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system scripts event-scripts reload

```
user@host> request system scripts event-scripts reload
Event scripts loaded
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Replace an Event Script

request system scripts refresh-from

IN THIS SECTION

- Syntax | 15178
- Description | 15178
- Options | 15179
- Required Privilege Level | 15179
- Sample Output | 15179

Release Information | 15181

Syntax

```
request system scripts refresh-from (commit | event | extension-service | op | snmp) file
filename url url cert-file path
<routing-instance routing-instance-name>
<sync>
```

Description

Refresh a local script from a remote source at the specified URL. The system uses the script type to determine the directory on the device where the script resides. If you include the `sync` option, the system refreshes the script on the requesting Routing Engine and then synchronizes the script on the other Routing Engine. If the `load-scripts-from-flash` statement is present in the configuration for that Routing Engine, the system refreshes the script in flash memory. Otherwise, the system refreshes the script on the hard disk.

If you specify an HTTPS URL, you must also include the `cert-file` option in supported releases. The `cert-file` option specifies the path to the certificate (Root CA or self-signed) that the device uses to validate the server's certificate.

> **NOTE**: We do not recommend refreshing scripts from an HTTPS server when the device is unable to authenticate the server because malicious users could compromise the integrity of a script.

The Junos XML management protocol equivalent for this operational mode command is:

```
<request-script-refresh-from>
   <type>(commit | op | event | extension-service | snmp)</type>
   <file>filename</file>
   <url>url</url>
   <routing-instance>routing-instance-name</routing-instance>
   <cert-file>path</cert-file>
   <sync/>
</request-script-refresh-from>
```

Options

(commit | event | extension-service | op | snmp)	Specify in which **scripts** subdirectory the script resides: **commit**, **event**, **jet**, **op**, or **snmp**.
cert-file *path*	Absolute path to the certificate (Root CA or self-signed) in Privacy-Enhanced Mail (PEM) format that is used to validate the server's certificate when you specify an HTTPS URL.
file *filename*	Filename of the script to refresh.
routing-instance *routing-instance-name*	(Optional) Name of the routing instance through which you want to refresh the script.
sync	(Optional) Synchronize the refreshed script from the requesting Routing Engine to the other Routing Engine.
url *url*	URL of the file that will replace the script provided as an HTTPS URL, HTTP URL, FTP URL, or secure copy (scp)-style remote file specification.

Required Privilege Level

maintenance

Sample Output

request system scripts refresh-from

The following command refreshes the op script **config.slax** from the remote source at the specified URL.

```
user@host> request system scripts refresh-from op file config.slax url http://host1.juniper.net/
config.slax
config.slax                              100%    9     0.0KB/s   00:00
```

request system scripts refresh-from

The following command refreshes the op script **config1.slax** from the remote source at the specified URL. However, in this case, the script does not exist at the remote site, and the device generates an error message.

```
user@host> request system scripts refresh-from op file config1.slax url http://host1.juniper.net/
config1.slax
refreshing 'config1.slax' from 'http://host1.juniper.net/config1.slax'
fetch-secure: http://host1.juniper.net/config1.slax: Not Found
error: file-fetch failed
error: communication error: rpc failed (file-fetch)
error: error retrieving file http://host1.juniper.net/config1.slax
```

request system scripts refresh-from sync

The following command refreshes the op script **config.slax** on the requesting Routing Engine from the remote source at the specified URL. The system then synchronizes the updated script to the other Routing Engine.

```
user@host> request system scripts refresh-from op file config.slax url http://host1.juniper.net/
config.slax sync
config.slax                               100%    9     0.0KB/s   00:00
```

request system scripts refresh-from cert-file

The following command refreshes the op script **ospf-neighbors.slax** from the remote source at the HTTPS URL. The command includes the `cert-file` option, which references the Root CA certificate file required to validate the server's certificate.

```
user@host> request system scripts refresh-from op file ospf-neighbors.slax url https://
example.juniper.net/scripts/ospf-neighbors.slax cert-file /var/tmp/root-ca-cert-example-juniper-
net.pem
refreshing 'ospf-neighbors.slax' from 'https://example.juniper.net/scripts/ospf-neighbors.slax'
/var/home/user/...transferring.file........100% of
```

Release Information

Command introduced in Junos OS Release 10.1.

`sync` option introduced in Junos OS Release 13.2.

`extension-service` and `routing-instance` options introduced in Junos OS Release 18.1R1.

`cert-file` option introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Refreshing a Script from an Alternate Location

Example: Refreshing a Script from an Alternate Source

Example: Configuring and Refreshing from the Master Source for a Script

Synchronizing Scripts Between Routing Engines

request system scripts synchronize | **15182**

request system scripts rollback

IN THIS SECTION

- Syntax | **15181**
- Description | **15182**
- Options | **15182**
- Required Privilege Level | **15182**
- Release Information | **15182**

Syntax

```
request system scripts rollback
```

Description

(Junos OS only) Use this command to roll back to most recent installation of AI-Scripts (jais) package.

Options

No options are available.

Required Privilege Level

maintenance

Release Information

Command introduced before Junos OS Release 9.0.

> **NOTE**: This command is not supported on Junos OS Evolved. AI-Scripts and Service Now are not supported on Junos OS Evolved.

RELATED DOCUMENTATION

- request system scripts add | 15171
- request system scripts delete | 15175
- request system scripts event-scripts reload | 15176

request system scripts synchronize

IN THIS SECTION

- Syntax | 15183
- Description | 15183
- Options | 15183
- Required Privilege Level | 15184
- Sample Output | 15184

Release Information | 15185

Syntax

```
request system scripts synchronize (all | commit | event | jet | lib | op | snmp)
<file filename>
<newer-than YYYY-MM-DD.HH:MM:SS>
```

Description

This command is for devices with multiple Routing Engines only. Synchronize scripts from the requesting Routing Engine to the other Routing Engine.

Specify `all` to synchronize all scripts present on the requesting Routing Engine to the responding Routing Engine. Specify `commit`, `event`, `jet`, `lib`, `op`, or `snmp` to synchronize all scripts of the given type to the other Routing Engine. Include the `file` option or the `newer-than` option to narrow the scope to only synchronize scripts with the specified filename or date criteria.

If the `load-scripts-from-flash` statement is configured for the requesting Routing Engine, the device synchronizes the scripts that are in flash memory. Otherwise, the device synchronizes the scripts that are on the hard disk. If the `load-scripts-from-flash` statement is configured for the responding Routing Engine, the scripts are synchronized to flash memory. Otherwise, the scripts are synchronized to the hard disk. The device synchronizes a script regardless of whether it is enabled in the configuration.

The Junos XML management protocol equivalent for this operational mode command is:

```
<request-scripts-synchronize>
   <all/> | <commit/> | <event/> | <jet/> | <lib/> | <op/> | <snmp/>
   <file>filename</file>
   <newer-than>YYYY-MM-DD.HH:MM:SS</newer-than>
</request-scripts-synchronize>
```

Options

(all | commit | event | jet | lib | op | snmp) Synchronize all scripts or all scripts of a given type that are present on the requesting Routing Engine to the responding Routing Engine.

file ***filename***	(Optional) Synchronize an individual script with the specified filename from the requesting Routing Engine to the responding Routing Engine. This option can only be used with a specific script type.
newer-than ***YYYY-MM-DD.HH:MM:SS***	(Optional) Synchronize scripts that are more current than the specified date and time from the requesting Routing Engine to the responding Routing Engine.

Required Privilege Level

maintenance

Sample Output

request system scripts synchronize all

The following command synchronizes all scripts from the requesting Routing Engine to the responding Routing Engine.

```
user@host> request system scripts synchronize all
user@host>
```

request system scripts synchronize commit file

The following command synchronizes the commit script **vpn-commit.slax** from the requesting Routing Engine to the responding Routing Engine. In this case, the script does not exist on the requesting Routing Engine, and the device issues an error message.

```
user@host> request system scripts synchronize commit file vpn-commit.slax
error: Invalid directory: No such file or directory
warning: No script will be pushed to other RE
```

request system scripts synchronize event file

The following command synchronizes the event script **ospf-neighbor.slax** from the requesting Routing Engine to the responding Routing Engine. In this case, the responding Routing Engine does not have the resources to perform the operation, and the device issues an error message.

```
user@host> request system scripts synchronize event file ospf-neighbor.slax
error: Unable to copy scripts to re: re1
```

request system scripts synchronize op newer-than

The following command synchronizes any op scripts that have been modified after the specified date. In this case, no op scripts meet this requirement.

```
user@host> request system scripts synchronize op newer-than 2012-05-30
warning: No script will be pushed to other RE
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding Script Synchronization Between Routing Engines

Synchronizing Scripts Between Routing Engines

synchronize (Scripts) | **28806**

request system scripts refresh-from | **15177**

request system services dhcp

IN THIS SECTION

Syntax | **15186**

- Description | 15186
- Options | 15186
- Required Privilege Level | 15186
- Output Fields | 15186
- Release Information | 15187

Syntax

```
request system services dhcp (release interface-name | renew interface-name)
```

Description

Release or renew the acquired IP address for a specific interface.

To view the status of the Dynamic Host Configuration Protocol (DHCP) clients on the specified interfaces, enter the `show system services dhcp client` *`interface-name`* command.

Options

- `release` *`interface-name`* —Clears other resources received earlier from the server, and reinitializes the client state to INIT for the particular interface.
- `renew` *`interface-name`* —Reacquires an IP address from the server for the interface. When you use this option, the command sends a discover message if the client state is INIT and a renew request message if the client state is BOUND. For all other states it performs no action.

Required Privilege Level

maintenance

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

dhcp (DHCP Client) | 5598

show system services dhcp client | 27038

request system snapshot (Junos OS Evolved)

IN THIS SECTION

- Syntax | 15187
- Description | 15187
- Options | 15188
- Additional Information | 15188
- Required Privilege Level | 15189
- Output Fields | 15189
- Sample Output | 15189
- Release Information | 15190

Syntax

```
request system snapshot
<configuration | no-configuration>
<routing-engine (both | local)>
```

Description

On devices running Junos OS Evolved, take a “snapshot” of the files currently used to run the device and copy the files onto the alternate solid-state drive (SSD). The snapshot contains the complete contents of

the **/soft**, **/config**, and **/root** directories, which include the current and all rollback software images, copies of user data, the active configuration, the rescue configuration, and content from the **/var** directory (except the **/var/core**, **/var/external**, **/var/log**, and **/var/tmp** directories). You can then use this snapshot to boot the device at the next boot up or as a backup boot option.

Starting with Junos OS Evolved 21.2, for dual-Routing Engine systems, this command only runs on the local Routing Engine, and not on both Routing Engines in the system. Any data in the **/var/home** directory of the secondary disk is no longer overwritten.

> **CAUTION**: After you run the `request system snapshot` command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Options

none Take a snapshot of the local Routing Engine.

configuration Copy the **/config** directory and the configuration stored for each installed software version to the alternate SSD as part of the snapshot. This is the default.

no-configuration Do not copy the **/config** directory or the configuration stored for each installed software version to the alternate SSD as part of the snapshot.

routing-engine (both | local) (Optional) Routing Engine on which to perform the snapshot. Supported values include:

- `both`—In a dual Routing Engine environment, take a snapshot of each Routing Engine. If the backup Routing Engine is outside of the system, for example, because of a software version mismatch, Junos OS Evolved takes a snapshot of the primary Routing Engine only.
- `local`—Take a snapshot of the local Routing Engine only.

Additional Information

Before you upgrade the software on the router or replace one of the Routing Engines, when you have a known stable system, issue the `request system snapshot` command to back up the software, including the configuration, to the **/soft** directory. After you have upgraded the software or have replaced one of the Routing Engines, and are satisfied that the software packages are successfully installed and running, issue the `request system snapshot` command again to back up the software to the **/soft** directory.

Required Privilege Level

view

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

request system snapshot (Junos OS Evolved)

```
user@host> request system snapshot
```

```
-------------------------------
node: re0
-------------------------------
........
Starting Snapshot in device /dev/sdb
List of software versions getting copied to Snapshot...
[1] junos-evo-install-ptx-x86-64-20.4-202103111254.0-EVO
[2] junos-evo-install-ptx-x86-64-20.4-202103150459.0-EVO
[3] junos-evo-install-ptx-x86-64-20.4-202103141559.0-EVO
........
........
........
........
........
........
........
........
........
........
........
........
........
Software Snapshot completed.
```

Release Information

Command introduced in Junos OS Evolved Release 20.4R2.

`routing-engine both` and `local` options added in Junos OS Evolved Release 22.4R1.

`configuration` and `no-configuration` options added in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

show system snapshot (Junos OS Evolved) | **27073**

Back up and Recover Software with Snapshots

request system snapshot (Junos OS with FreeBSD Prior to Release 10)

IN THIS SECTION

- Syntax (ACX Series Routers) | **15191**
- Syntax (EX Series Switches; for EX4600, see QFX Series Syntax) | **15191**
- Syntax (MX Series Routers) | **15191**
- Syntax (PTX Series) | **15192**
- Syntax (QFX Series, OCX1100, and EX4600) | **15192**
- Syntax (SRX Series Firewalls) | **15192**
- Syntax (TX Matrix Routers) | **15192**
- Syntax (TX Matrix Plus Routers) | **15193**
- Description | **15193**
- Options | **15193**
- Additional Information | **15198**
- Required Privilege Level | **15198**
- Output Fields | **15198**
- Sample Output | **15198**
- Release Information | **15202**

If your device is running an upgraded version of FreeBSD (10+), use this topic instead: request system snapshot (Junos OS with Upgraded FreeBSD).

To determine which platforms support Junos OS with upgraded FreeBSD, see Feature Explorer and enter one of the following:

- For non-virtualized, enter **freebsd** and select **Junos kernel upgrade to FreeBSD 10+**.
- For virtualized, enter **virtualization** and select **Virtualization of the Routing Engine**.

Syntax (ACX Series Routers)

```
request system snapshot
<media type>
<partition>
```

Syntax (EX Series Switches; for EX4600, see QFX Series Syntax)

```
request system snapshot
<all-members | local | member member-id>
<media type>
<partition>
<re0 | re1 | routing-engine routing-engine-id>
<slice alternate>
```

Syntax (MX Series Routers)

```
request system snapshot
<all-members>
<config-partition>
<local>
<member member-id>
<media usb-port-number>
<partition>
<root-partition>
```

Syntax (PTX Series)

```
request system snapshot
<partition>
```

Syntax (QFX Series, OCX1100, and EX4600)

```
request system snapshot
<all-members | local | member member-id>
<config-partition>
<partition>
<root-partition>
<slice alternate>
```

Syntax (SRX Series Firewalls)

```
request system snapshot
<config-partition>
<media (compact-flash | hard-disk | internal | usb)>
<partition>
<root-partition>
<factory>
<node (all | local | node-id | primary)>
<slice (alternate) >
```

Syntax (TX Matrix Routers)

```
request system snapshot
<all-chassis | all-lcc | lcc number /  scc>
<config-partition>
<partition>
<root-partition>
```

Syntax (TX Matrix Plus Routers)

```
request system snapshot
<all-chassis | all-lcc | lcc number /  sfc number>
<config-partition>
<partition>
<root-partition>
```

Description

- On routers and firewalls running Junos OS, back up the currently running and active file system partitions to standby partitions that are not running. Specifically, the root file system (/) is backed up to **/altroot**, and **/config** is backed up to **/altconfig**. The root and **/config** file systems are on the flash drive, and the **/altroot** and **/altconfig** file systems are on the hard drive.
- On switches running Junos OS, take a snapshot of the files currently used to run the switch—the complete contents of the root (/) , **/altroot**, **/config**, **/var**, and **/var-tmp** directories, which include the running version of Junos OS, the active configuration, and log files.

> **NOTE**: System snapshot is not supported on EX4650 switches. System snapshot is not supported on QFX10000,QFX5110, and QFX5200 switches. System snapshot using an external USB drive is supported on EX4600 switches.

Starting with Junos OS Release 15.1F3, the command `request system snapshot` creates a snapshot of the guest OS image only for the PTX5000 with RE-DUO-C2600-16G, and the MX240, MX480, and MX960 routers with RE-S-1800X4-32G-S.

Starting with Junos OS Release 15.1F5, the command `request system snapshot` creates a snapshot of the guest OS image only for the MX2010 and MX2020 routers with REMX2K-1800-32G-S.

On these routers, in order to create snapshot of the host OS image along with Junos OS image, use the `request vmhost snapshot` command.

> **CAUTION**: After you run the `request system snapshot` command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Options

The specific options available depend upon the device:

none

Back up the currently running software as follows:

- On the router or firewall, back up the currently running and active file system partitions to standby partitions that are not running. Specifically, the root file system (/) is backed up to **/altroot**, and **/config** is backed up to **/altconfig**. The root and **/config** file systems are on the flash drive, and the **/altroot** and **/altconfig** file systems are on the hard drive.

- On the switch, take a snapshot of the files currently used to run the switch and copy them to the media that the switch did not boot from. If the switch is booted from internal media, the snapshot is copied to external (USB) media. If the switch is booted from external (USB) media, the snapshot is copied to internal media.

  - If the snapshot destination is external media but a USB flash drive is not connected, an error message is displayed.

  - If the automatic snapshot procedure is already in progress, the command returns the following error: `Snapshot already in progress. Cannot start manual snapshot.` For additional information about the automatic snapshot feature, see *Configuring Dual-Root Partitions*.

all-chassis | all-lcc | lcc *number*

(TX Matrix and TX Matrix Plus router only) (Optional)

- `all-chassis`—On a TX Matrix router, archive data and executable areas for all Routing Engines in the chassis. On a TX Matrix Plus router, archive data and executable areas for all Routing Engines in the chassis.

- `all-lcc`—On a TX Matrix router, archive data and executable areas for all T640 routers (or line-card chassis) connected to a TX Matrix router. On a TX Matrix Plus router, archive data and executable areas for all routers (or line-card chassis) connected to a TX Matrix Plus router.

- `lcc` *`number`*—On a TX Matrix router, archive data and executable areas for a specific T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, archive data and executable areas for a specific router (line-card chassis) that is connected to a TX Matrix Plus router.

  Replace *number* with the following values depending on the LCC configuration:

  - 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.

- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

all-members | local | member *member-id*

(EX Series Virtual Chassis, MX Series routers, QFX Series switches, QFabric System, and OCX1100 only) (Optional) Specify where to place the snapshot (archive data and executable areas) in a Virtual Chassis:

- `all-members`—Create a snapshot (archive data and executable areas) for all members of the Virtual Chassis.
- `local`—Create a snapshot (archive data and executable areas) on the member of the Virtual Chassis that you are currently logged into.
- `member` *member-id*—Create a snapshot (archive data and executable areas) for the specified member of the Virtual Chassis.

config-partition

(EX Series Virtual Chassis, MX Series routers, QFX Series switches, QFabric System, OCX1100, SRX Series Firewalls, and T and TX Series routers only) Create a snapshot of the configuration partition only and store it onto the default **/altconfig** on the hard disk device or an **/altconfig** on a USB device. Option deprecated for Junos OS with Upgraded FreeBSD in Junos OS Release 15.1.

To determine which platforms support Junos OS with upgraded FreeBSD, see Feature Explorer and enter one of the following:

- For non-virtualized, enter **freebsd** and select **Junos kernel upgrade to FreeBSD 10+**.
- For virtualized, enter **virtualization** and select **Virtualization of the Routing Engine**.

factory

(SRX Series Firewalls only) (Optional) Specifies that only the files shipped from the factory are included in the snapshot.

media *type*

(ACX Series, M320, T640, and MX Series routers, SRX Series Firewalls) (Optional) Specify the boot device the software is copied to:

- `compact-flash`—(SRX5800, SRX5600, and SRX5400 devices only) Copy software to the primary compact flash drive.

- `external`—(Switches and SRX650 Services Gateway only) Copy software to an external mass storage device, such as a USB flash drive. If a USB drive is not connected, the switch displays an error message. This option is available for the compact flash on the SRX650 Services Gateway.
- `hard-disk`—(SRX5800, SRX5600, and SRX5400 devices only) Copy software to the hard disk drive.
- `internal`—(SRX300, SRX320, SRX340, SRX345, SRX380 devices only) Copy software to an internal flash drive. This is the default option.
- `removable-compact-flash`—(ACX Series, M320, T640, and MX Series routers only) Copy software to the removable compact flash drive.
- `usb`—(ACX Series, SRX Series Firewalls, M320, T640, and, except for MX104, MX Series routers) Copy software to the device connected to the USB port.
- `usb0`—(MX104 routers only) Copy software to the device connected to the USB0 port.
- `usb1`—(MX104 routers only) Copy software to the device connected to the USB1 port.

node

(SRX Series Firewalls only) (Optional) Specify the archive data and executable areas of a specific node. If you do not specify the `node` option, the device considers the current node as the default option.

- `node-id`—Specify for node (0, 1).
- `all`—Specify for all nodes.
- `local`—Specify for local nodes.
- `primary`— Specify for primary nodes.

partition

(Optional) Repartition the flash drive before a snapshot occurs. If the partition table on the flash drive is corrupted, the `request system snapshot` command fails and reports errors. The partition option is only supported for restoring the software image from the hard drive to the flash drive.

(Routers only) You cannot issue the request system snapshot command when you enable flash disk mirroring. We recommend that you disable flash disk mirroring when you upgrade or downgrade the software.

(EX Series switches only) If the snapshot destination is the media that the switch did not boot from, you must use the `partition` option.

(SRX Series Firewalls only) (Default) The target media is partitioned whether or not it is specified in the command, because this is a mandatory option. For example:

```
user@host> request system snapshot media usb partition
```

```
user@host> request system snapshot media usb partition factory
```

re0 | re1 | routing-engine *routing-engine-id*

(EX6200 and EX8200 switches only) Specify where to place the snapshot in a redundant Routing Engine configuration.

- `re0`—Create a snapshot on Routing Engine 0.
- `re1`—Create a snapshot on Routing Engine 1.
- `routing-engine` *routing-engine-id*—Create a snapshot on the specified Routing Engine.

root-partition

(M, MX, T, and TX Series routers; EX Series Virtual Chassis; QFX Series switches; QFabric System; SRX Series Firewalls; and OCX1100 only) Create a snapshot of the root partition only and store it onto the default **/altroot** on the hard disk device or an **/altroot** on a USB device. Option deprecated for Junos OS with Upgraded FreeBSD in Junos OS Release 15.1.

To determine which platforms run Junos OS with Upgraded FreeBSD, see the information in Release Information for Junos OS with Upgraded FreeBSD.

slice alternate

(EX Series switches, EX Series Virtual Chassis, QFX Series switches, QFabric System, SRX Series Firewalls, and OCX1100 only) (Optional) Take a snapshot of the active root partition and copy it to the alternate slice on the boot media.

(SRX Series Firewalls) The slice option cannot be used along with the other `request system snapshot` options, because the options are mutually exclusive. If you use the `factory`, `media`, or `partition` option, you cannot use the `slice` option; if you use the `slice` option, you cannot use any of the other options.

Option deprecated for Junos OS with Upgraded FreeBSD in Junos OS Release 15.1.

To determine which platforms support Junos OS with upgraded FreeBSD, see Feature Explorer and enter one of the following:

- For non-virtualized, enter **freebsd** and select **Junos kernel upgrade to FreeBSD 10+**.
- For virtualized, enter **virtualization** and select **Virtualization of the Routing Engine**.

scc	(TX Matrix router only) (Optional) Archive data and executable areas for a TX Matrix router (or switch-card chassis).
sfc *number*	(TX Matrix Plus router only) (Optional) Archive data and executable areas for a TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

Additional Information

- (Routers only) Before you upgrade the software on the router, when you have a known stable system, issue the `request system snapshot` command to back up the software, including the configuration, to the **/altroot** and **/altconfig** file systems. After you have upgraded the software on the router and are satisfied that the new packages are successfully installed and running, issue the `request system snapshot` command again to back up the new software to the **/altroot** and **/altconfig** file systems.
- (Routers only) You cannot issue the `request system snapshot` command when you enable flash disk mirroring. We recommend that you disable flash disk mirroring when you upgrade or downgrade the software.
- (TX Matrix and TX Matrix Plus router only) On a routing matrix, if you issue the `request system snapshot` command on the primary Routing Engine, all the primary Routing Engines connected to the routing matrix are backed up. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are backed up.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system snapshot (Routers)

```
 user@host> request system snapshot
umount: /altroot: not currently mounted
Copying / to /altroot.. (this may take a few minutes)
umount: /altconfig: not currently mounted
Copying /config to /altconfig.. (this may take a few minutes)
```

```
The following filesystems were archived: / /config
```

request system snapshot (EX Series Switches)

```
 user@switch> request system snapshot partition
Clearing current label...
Partitioning external media (/dev/da1) ...
Partitions on snapshot:

  Partition  Mountpoint  Size    Snapshot argument
      s1a    /altroot    179M    none
      s2a    /           180M    none
      s3d    /var/tmp    361M    none
      s3e    /var        121M    none
      s4d    /config     60M     none
Copying '/dev/da0s1a' to '/dev/da1s1a' .. (this may take a few minutes)
Copying '/dev/da0s2a' to '/dev/da1s2a' .. (this may take a few minutes)
Copying '/dev/da0s3d' to '/dev/da1s3d' .. (this may take a few minutes)
Copying '/dev/da0s3e' to '/dev/da1s3e' .. (this may take a few minutes)
Copying '/dev/da0s4d' to '/dev/da1s4d' .. (this may take a few minutes)
The following filesystems were archived: /altroot / /var/tmp /var /config
```

request system snapshot partition (EX4600, QFX Series, QFabric System, and OCX1100)

```
user@switch> request system snapshot partition
Clearing current label...
Partitioning external media (da1) ...
Verifying compatibility of destination media partitions...
Running newfs (334MB) on external media  / partition ...
Running newfs (404MB) on external media  /config partition ...
Running newfs (222MB) on external media  /var partition ...
Copying '/dev/da0s2a' to '/dev/da1s1a' .. (this may take a few minutes)
Copying '/dev/da0s3e' to '/dev/da1s3e' .. (this may take a few minutes)
Copying '/dev/da0s2f' to '/dev/da1s1f' .. (this may take a few minutes)
The following filesystems were archived: / /config /var
```

request system snapshot partition (SRX Series Firewalls)

```
user@host> request system snapshot partition
Verifying compatibility of destination media partitions...
Running newfs (439MB) on internal media  / partition (da0s1a)...
Running newfs (46MB) on internal media  /config partition (da0s1e)...
Copying '/dev/da1s1a' to '/dev/da0s1a' .. (this may take a few minutes)
Copying '/dev/da1s1e' to '/dev/da0s1e' .. (this may take a few minutes)
The following filesystems were archived: / /config
```

request system snapshot config-partition (SRX Series Firewalls)

```
user@host> request system snapshot config-partition
Doing the initial labeling...
Verifying compatibility of destination media partitions...
Running newfs (391MB) on hard-disk media  /config partition (ad1s1e)...
Copying '/dev/ad0s1e' to '/dev/ad1s1e' .. (this may take a few minutes)
The following filesystems were archived: /config
```

request system snapshot root-partition (SRX Series Firewalls)

```
user@host> request system snapshot root-partition
Doing the initial labeling...
Verifying compatibility of destination media partitions...
Running newfs (3GB) on hard-disk media  / partition (ad1s1a)...
Copying '/dev/ad0s1a' to '/dev/ad1s1a' .. (this may take a few minutes)
The following filesystems were archived: /
```

request system snapshot (When the Partition Flag Is On)

```
 user@host>  request system snapshot partition
Performing preliminary partition checks ...
Partitioning ad0 ...
umount: /altroot: not currently mounted
Copying / to /altroot.. (this may take a few minutes)

The following filesystems were archived: / /config
```

request system snapshot (MX104 Routers When Media Device is Missing)

```
user@host > request system snapshot media usb0
error: usb0 media missing or invalid
```

request system snapshot (When Mirroring Is Enabled)

```
 user@host>  request system snapshot
Snapshot is not possible since mirror-flash-on-disk is configured.
```

request system snapshot all-lcc (Routing Matrix)

```
 user@host>  request system snapshot all-lcc
lcc0-re0:
--------------------------------------------------------------------------
Copying '/' to '/altroot' .. (this may take a few minutes)
Copying '/config' to '/altconfig' .. (this may take a few minutes)
The following filesystems were archived: / /config

lcc2-re0:
--------------------------------------------------------------------------
Copying '/' to '/altroot' .. (this may take a few minutes)
Copying '/config' to '/altconfig' .. (this may take a few minutes)
The following filesystems were archived: / /config
```

request system snapshot all-members (Virtual Chassis)

```
user@switch>  request system snapshot all-members media internal

fpc0:
--------------------------------------------------------------------------
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc1:
--------------------------------------------------------------------------
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
```

```
The following filesystems were archived: /

fpc2:
--------------------------------------------------------------------------
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc3:
--------------------------------------------------------------------------
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc4:
--------------------------------------------------------------------------
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc5:
--------------------------------------------------------------------------
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /
```

Release Information

Command introduced before Junos OS Release 7.4.

Options `<config-partition>` and `<root-partition>` introduced in Junos OS Release 13.1 for M Series, MX Series, T Series, and TX Series routers.

Option `media` *`usb-port-number`* introduced in Junos OS Release 13.2 for MX104 routers.

Options `<config-partition>`, `<root-partition>`, and `<slice>` deprecated for Junos OS with Upgraded FreeBSD in Junos OS Release 15.1

To determine which platforms support Junos OS with upgraded FreeBSD, see Feature Explorer and enter one of the following:

- For non-virtualized, enter **freebsd** and select **Junos kernel upgrade to FreeBSD 10+**.
- For virtualized, enter **virtualization** and select **Virtualization of the Routing Engine**.

RELATED DOCUMENTATION

request system snapshot (Junos OS with Upgraded FreeBSD) | 15203
show system snapshot (Junos OS) | 27069
show system auto-snapshot | 26720

request system snapshot (Junos OS with Upgraded FreeBSD)

IN THIS SECTION

- Syntax | 15203
- Description | 15203
- Options | 15204
- Additional Information | 15205
- Required Privilege Level | 15206
- Output Fields | 15206
- Sample Output | 15206
- Release Information | 15207

Syntax

```
request system snapshot
<configuration>
<delete snapshot-name>
<load snapshot-name>
<media type>
<name snapshot-name>
<no-configuration>
<recovery>
```

Description

Junos OS with upgraded FreeBSD has two types of snapshots:

- Recovery snapshot
- Non-recovery snapshot

Recovery snapshots—Use the `request system snapshot recovery` command to create a recovery snapshot. The recovery snapshot contains copies of the packages and configuration taken at the time of the snapshot. The recovery snapshot is stored at **recovery.ufs** on the **/oam** volume.

QFX with TVP platforms do not support these commands.

> **CAUTION**: Once the platform performs a Junos OS upgrade to an upgraded FreeBSD version (Junos OS Release 15.1R1 or higher for most devices; or Junos OS Release 24.4R1 or higher for the SRX300, SRX320, SRX340, SRX345, or SRX380 devices)), Junos OS automatically generates a recovery snapshot with the new Junos OS version. Thereafter, you should request a recovery snapshot every time you upgrade to a new software version. To find which platforms in which releases use Junos OS with upgraded FreeBSD, see Feature Explorer, enter **freebsd**, and select **Junos kernel upgrade to FreeBSD 10+**.

Non-recovery snapshots—Use the `request system snapshot` command to create a non-recovery snapshot. The non-recovery snapshots are essentially lists of software components and configuration files, which can be helpful when major software or configuration changes are occurring and establishing a known stable system baseline is required.

On the device, back up the currently running and active file system partitions to standby partitions that are not running. By default, non-recovery snapshots are named **snap.*date.time*** and are stored in the **/packages/sets** directory.

After you run the `request system snapshot` command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Options

none	On the router or switch, back up the currently running and active file system partitions to standby partitions that are not running. Specifically, this creates a non-recovery snapshot named **snap.< date>.< time>** which is stored in **/packages/sets**.
configuration	Save a copy of the configuration files (contents of the **/config** and **/var** directories and SSH files) as part of the snapshot. This is the default.

`delete` *`snapshot-name`*	(Optional) Delete a specific non-recovery snapshot from **/packages/sets**. Wildcards are supported, so `request system snapshot delete snap*` deletes all snapshots that start with "**snap**".
`load` *`snapshot-name`*	(Optional) Load a specific snapshot from **/packages/sets**.
`media` *`type`*	(Optional) Specify the boot device the software is copied to: • `usb`—(MX960 routers and SRX300, SRX320, SRX340, SRX345, or SRX380 firewalls only) Copy software to the device connected to the USB port.
name *snapshot-name*	Specify the name for a non-recovery snapshot.
no-configuration	Do not save a copy of the configuration files (contents of the **/config** and **/var** directories and SSH files) as part of the snapshot.
`recovery`	Create a recovery snapshot and store it in the **/oam** volume.

Additional Information

Before upgrading the software on the device, when you have a known stable system, issue the `request system snapshot` command to back up the software, including the configuration, to the **/packages/sets** file systems. After you have upgraded the software on the router or switch and are satisfied that the new packages are successfully installed and running, issue the `request system snapshot` command again to back up the new software to the **/packages/sets** file systems.

> **CAUTION**: After you have upgraded the software on the device and are satisfied that the new packages are successfully installed and running or if you have replaced one of the Routing Engines, you should also issue the `request system snapshot recovery` command to have a new updated recovery snapshot. Junos OS automatically recovers the Junos OS volume upon reboot, if considered defective, from the recovery snapshot, along with the configuration saved in the recovery snapshot.

The snapshot script (which is the script that generates output for non-recovery snapshots) does not generate XML output. In such cases, the <output> tag is used.

```
user@host> request system snapshot | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/18.1I0/junos">
<output>
NOTICE: Snapshot snap.20180105.165049 created successfully
</output>
```

```
<cli>
<banner></banner>
</cli>
</rpc-reply>
```

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system snapshot recovery

```
user@host> request system snapshot recovery
Creating image ...
Compressing image . . .
Image size is 777MB
Recovery snapshot created successfully
```

request system snapshot

```
user@host> request system snapshot
NOTICE: Snapshot snap.20190412.124526 created successfully
```

request system snapshot name

```
user@host> request system snapshot name pre-maintenance-snapshot
NOTICE: Snapshot pre-maintenance-snapshot created successfully
```

show system snapshot

```
user@host> show system snapshot
Non-recovery snapshots:
Snapshot snap.20190412.124526:
Location: /packages/sets/snap.20190412.124526
Creation date: Apr 12 12:45:26 2019
Junos version: 18.3R1.8

Total non-recovery snapshots: 1

Recovery Snapshots:
Snapshots available on the OAM volume:
recovery.ufs
Date created: Thu Mar 28 07:44:25 PDT 2019
Junos version: 18.3R1.8

Total recovery snapshots: 1
```

request system snapshot delete

```
user@host> request system snapshot delete snap.20150112.122106
NOTICE: Snapshot 'snap.20150112.122106'
deleted successfully
```

Release Information

Command introduced in Junos OS Release 15.1.

Command introduced for all platforms using Junos OS with upgraded FreeBSD. To find which platforms in which releases use Junos OS with upgraded FreeBSD, see Feature Explorer, enter **freebsd**, and select **Junos kernel upgrade to FreeBSD 10+**.

`configuration` and `no-configuration` options added in Junos OS Release 23.2R1.

`name` option added in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

show system snapshot (Junos OS) | **27069**

request system snapshot (Maintenance)

IN THIS SECTION

- Syntax | 15208
- Description | 15208
- Options | 15208
- Required Privilege Level | 15210
- Output Fields | 15210
- Sample Output | 15210
- Release Information | 15211

Syntax

```
request system snapshot
<config-partition>
<media (compact-flash | hard-disk | internal | usb)>
<partition>
<root-partition>
<factory>
<node (all | local | node-id | primary)>
<slice (alternate) >
```

Description

Use this command to back up the currently running and active file system partitions on the device.

Options

- `config-partition`— Creates a snapshot of the configuration partition only and stores it onto the default **/altconfig** on the hard disk device or an **/altconfig** on a USB device.
- `root-partition`— Creates a snapshot of the root partition only and stores it onto the default **/altroot** on the hard disk device or an **/altroot** on a USB device.
- `factory`— (Optional) Specifies that only the files shipped from the factory are included in the snapshot.

- `media`—(Optional) Specify the boot device the software is copied to:
  - `compact-flash`—Copy software to the primary compact flash drive.
  - `hard-disk`— Copy software to the hard disk.
  - `usb`— Copy software to the device connected to the USB port.
  - `internal`— Copy software to an internal flash drive. This is the default option.

    USB option is available on all SRX Series Firewalls; hard disk and compact-flash options are available only on SRX5800, SRX5600, and SRX5400 devices; media internal option is available only on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550M devices.
  - external— Copies software to an external storage device. This option is available for the compact flash on the SRX650 Services Gateway.
- `node`—(Optional) Specify the archive data and executable areas of a specific node. If you do not specify the node option, the device considers the current node as default option.
  - `node-id`—Specify for node (0, 1).
  - `all`—Specify for all nodes.
  - `local`—Specify for local nodes.
  - `primary`— Specify for primary nodes.
- `partition`—(Default) Specify that the target media should be repartitioned before the backup is saved to it.

  The target media is partitioned whether or not it is specified in the command, because this is a mandatory option.

  Example: **request system snapshot media usb partition**

  Example: **request system snapshot media usb partition factory**
- `slice`—(Optional) Take a snapshot of the root partition the system has currently booted from to another slice in the same media.
  - `alternate`—(Optional) Store the snapshot on the other root partition in the system.

  The slice option cannot be used along with the other `request system snapshot` options, because the options are mutually exclusive. If you use the factory, media, or partition option, you cannot use the slice option; if you use the slice option, you cannot use any of the other options.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system snapshot config-partition

```
user@host> request system snapshot config-partition
Doing the initial labeling...
Verifying compatibility of destination media partitions...
Running newfs (391MB) on hard-disk media  /config partition (ad1s1e)...
Copying '/dev/ad0s1e' to '/dev/ad1s1e' .. (this may take a few minutes)
The following filesystems were archived: /config
```

request system snapshot root-partition

```
user@host> request system snapshot root-partition
Doing the initial labeling...
Verifying compatibility of destination media partitions...
Running newfs (3GB) on hard-disk media  / partition (ad1s1a)...
Copying '/dev/ad0s1a' to '/dev/ad1s1a' .. (this may take a few minutes)
The following filesystems were archived: /
```

request system snapshot media hard-disk

```
user@host> request system snapshot media hard-disk
Verifying compatibility of destination media partitions...
Running newfs (880MB) on hard-disk media  / partition (ad2s1a)...
Running newfs (98MB) on hard-disk media  /config partition (ad2s1e)...
Copying '/dev/ad0s1a' to '/dev/ad2s1a' .. (this may take a few minutes)
...
```

request system snapshot media usb (when usb device is missing

```
user@host> request system snapshot media usb
Verifying compatibility of destination media partitions...
Running newfs (254MB) on usb media  / partition (da1s1a)...
Running newfs (47MB) on usb media  /config partition (da1s1e)...
Copying '/dev/da0s2a' to '/dev/da1s1a' .. (this may take a few minutes)
Copying '/dev/da0s2e' to '/dev/da1s1e' .. (this may take a few minutes)
The following filesystems were archived: / /config
```

request system snapshot media compact-flash

```
user@host> request system snapshot media compact-flash
error: cannot snapshot to current boot device
```

request system snapshot partition

```
user@host> request system snapshot partition
Verifying compatibility of destination media partitions...
Running newfs (439MB) on internal media  / partition (da0s1a)...
Running newfs (46MB) on internal media  /config partition (da0s1e)...
Copying '/dev/da1s1a' to '/dev/da0s1a' .. (this may take a few minutes)
Copying '/dev/da1s1e' to '/dev/da0s1e' .. (this may take a few minutes)
The following filesystems were archived: / /config
```

Release Information

Command introduced in Junos OS Release 10.2.

request system software abort

IN THIS SECTION

- Syntax | 15212
- Description | 15212
- Options | 15212
- Required Privilege Level | 15212
- Output Fields | 15213
- Sample Output | 15213
- Release Information | 15214

Syntax

```
request system software abort in-service-upgrade
```

Description

Terminate a unified in-service software upgrade (ISSU). The unified ISSU must be in progress and you must issue this command from a router session other than the one on which you issued the `request system in-service-upgrade` command that launched the unified ISSU.

Options

This command has no options.

Required Privilege Level

view

Output Fields

When you enter the `request system software abort` command on a new router session, you are provided feedback on the status of your request in the router session on which you issued the `request system software in-service-upgrade` command.

Sample Output

request system software abort (New Router Session)

```
user@host> request system software abort
```

request system software in-service-upgrade (Unified ISSU Session)

```
user@host> request system software in-service-upgrade /var/tmp/jinstall-9.0-20080117.0-domestic-
signed.tgz
ISSU: Preparing Backup RE
Pushing bundle to re1
Checking compatibility with configuration Initializing...
Using jbase-9.0-20080116.2
Verified manifest signed by PackageProduction_9_0_0 Using /var/tmp/jinstall-9.0-20080117.0-
domestic-signed.tgz
Verified jinstall-9.0-20080117.0-domestic.tgz signed by PackageProduction_9_0_0 Using
jinstall-9.0-20080117.0-domestic.tgz
Using jbundle-9.0-20080117.0-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.0-20080117.0.tgz
Verified manifest signed by PackageProduction_9_0_0 Using jkernel-9.0-20080117.0.tgz Verified
manifest signed by PackageProduction_9_0_0 Using jcrypto-9.0-20080117.0.tgz Verified manifest
signed by PackageProduction_9_0_0 Using jpfe-9.0-20080117.0.tgz Using jdocs-9.0-20080117.0.tgz
Verified manifest signed by PackageProduction_9_0_0 Using jroute-9.0-20080117.0.tgz Verified
manifest signed by PackageProduction_9_0_0 Hardware Database regeneration succeeded Validating
against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-9.0-20080117.0-domestic-signed.tgz'
...
Verified jinstall-9.0-20080117.0-domestic.tgz signed by PackageProduction_9_0_0 Adding
jinstall...
Verified manifest signed by PackageProduction_9_0_0
```

```
WARNING:     This package will load JUNOS 9.0-20080117.0 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in
/var/sw/pkg/jinstall-9.0-20080117.0-domestic-signed.tgz ...
Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
error: ISSU Aborted! Backup RE maybe in inconsistent state, Please restore backup RE
ISSU: IDLE

{master}
user@host>
```

Release Information

Command introduced in JUNOS Release 9.0.

RELATED DOCUMENTATION

request system software in-service-upgrade | 15262

show chassis in-service-upgrade | 18507

Getting Started with Unified In-Service Software Upgrade

Example: Performing a Unified ISSU

request system software abort in-service-upgrade (ICU)

IN THIS SECTION

- Syntax | 15215
- Description | 15215
- Options | 15216
- Required Privilege Level | 15216
- Output Fields | 15216
- Sample Output | 15216
- Release Information | 15216

Syntax

```
request system software abort in-service-upgrade
```

Description

Use this command to terminate an in-band cluster upgrade (ICU).

This command must be issued from a router session other than the one on which you issued the `request system in-service-upgrade` command that launched the ICU. If an ICU is in progress, this command terminates it. If the node is being upgraded, this command will cancel the upgrade. This command is also helpful in recovering the node in case of a failed ICU.

We recommend that you use the command only when there is an issue with the ongoing session of ISSU. You may need to manually intervene to bring the system to sane state if after issuing the command the system does not recover from the terminate.

Options

This command has no options.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software abort in-service-upgrade

```
user@host> request system software abort in-service-upgrade
In-Service-Upgrade aborted
```

Release Information

Command introduced in Junos OS Release 15.1X49-D35.

RELATED DOCUMENTATION

Upgrading Devices in a Chassis Cluster Using ICU

request system software add (Junos OS)

IN THIS SECTION

- Syntax | 15217
- Syntax (EX Series Switches) | 15217
- Syntax (MX Series Router) | 15218

- Syntax (QFX Series) | 15218
- Syntax (OCX Series) | 15219
- Description | 15219
- Options | 15221
- Additional Information | 15225
- Required Privilege Level | 15225
- Output Fields | 15226
- Sample Output | 15226
- Release Information | 15232

Syntax

```
request system software add package-name
<best-effort-load>
<delay-restart>
<device-alias alias-name>
<force>
<no-copy>
<no-validate>
<re0 | re1>
<reboot>
<satellite slot-id>
<set [package-name1 package-name2]>
<unlink>
<upgrade-group [all |upgrade-group-name]>
<upgrade-with-config>
<satellite slot-id>
<validate>
<version version-string>
```

Syntax (EX Series Switches)

```
request system software add package-name
<best-effort-load>
<delay-restart>
```

```
<force>
<no-copy>
<no-validate>
<re0 | re1>
<reboot>
<set [package-name1 package-name2]>
<upgrade-with-config>
<validate>
<validate-on-host hostname>
<validate-on-routing-engine routing-engine>
```

Syntax (MX Series Router)

```
request system software add package-name
<best-effort-load>
<delay-restart>
<device-alias alias-name>
<force>
<member member-id>
<no-copy>
<no-validate>
<re0 | re1>
<reboot>
<satellite slot-id>
<set [package-name1 package-name2]>
<upgrade-group [all |upgrade-group-name]>
<unlink>
<upgrade-with-config>
<validate>
<version version-string>
<validate-on-host hostname>
<validate-on-routing-engine routing-engine>
```

Syntax (QFX Series)

```
request system software add package-name
<best-effort-load>
<component all>
<delay-restart>
```

```
<force>
<force-host>
<no-copy>
<partition>
<reboot>
<unlink>
<upgrade-with-config>
```

Syntax (OCX Series)

```
request system software add package-name
<best-effort-load>
<delay-restart>
<force>
<force-host>
<no-copy>
<no-validate>
<reboot>
<unlink>
<upgrade-with-config>
<validate>
```

Description

Install a software package or bundle on the device.

We recommend that you always download the software image to **/var/tmp** only. On EX Series and QFX Series switches, you must use the **/var/tmp** directory. Other directories are not supported.

Starting in Junos OS Release 23.4R1 for MX Series routers, you can configure the `dual-phase-bootup` statement at the [edit system] hierarchy level. During device upgrades or downgrades, this feature enables a two-step restart process. If the `dual-phase-bootup` statement is already configured and committed, the device picks up the **rescue.conf** file during any upgrade or downgrade. The image is then validated against both the **rescue.conf** file and the **juniper.conf** file. An internal commit is performed post-reboot, with the latest last-run user configuration being committed. While this configuration is being committed, you cannot make any changes to the configuration.

When you are upgrading to a different release of Junos OS, you usually use the `validate` option on this command. The `validate` option checks the candidate software against the current configuration of the device to ensure they are compatible. (Validate is the default behavior when the software package being added is a different release.) However, there are circumstances under which you cannot validate the

running configuration in this way. One such circumstance is when you are upgrading to Junos OS with upgraded FreeBSD from Junos OS based on FreeBSD 6.1. Another such circumstance is when you are updating between different releases of Junos OS with upgraded FreeBSD, and the newest version of FreeBSD uses system calls that are not available in earlier versions of FreeBSD.

Therefore, for platforms already running an upgraded FreeBSD release, you cannot use the `validate` option when upgrading to Junos OS Release 21.2R1, because this release runs on FreeBSD release 12; previous releases with upgraded FreeBSD run either FreeBSD release 10 or 11.

For SRX300, SRX320, SRX340, SRX345, and SRX380 Firewalls, when you are upgrading to Junos OS Release 24.4R1 from Junos OS 23.4R2-S3 or from Release 24.2R2, you must use the `request system software add` *`package-name`* `no-copy no-validate reboot` command to upgrade the software. Similarly, to downgrade the software, you must use the `request system software add` *`package-name`* `no-validate reboot` command. When downgrading the software, sometimes the non-SRX380 firewalls cannot complete the downgrade because the **/oam** partition does not have enough free space.

If you are upgrading between releases that cannot use direct validation, you need to specify one of the following on the `request system software add` operational mode command when you upgrade:

- The `no-validate` option—this option does not validate the software package against the current configuration. Therefore, the current configuration might fail once you upgrade the system. Choose this option for the first time you upgrade a system to the newer version.
- The `validate-on-host` option—this option validates the software package by comparing it to the running configuration on a remote Junos OS host. Be sure to choose a host that you have already upgraded to the newer version of software.
- The `validate-on-routing-engine` option—(for systems with redundant REs) this option validates the software package by comparing it to the running configuration on a Routing Engine in the same chassis. Use this option when you have already upgraded the other Routing Engine to the newer version.

For information on valid filename and URL formats, see Format for Specifying Filenames and URLs in Junos OS CLI Commands.

Any configuration changes performed after inputting the `request system software add` command remain after the system reboots with an upgraded version of Junos OS.

Starting from Junos OS Release 17.2R1, PTX10008 routers do not support the `request system software add` command. Starting from Junos OS Release 17.4R1, PTX10016 routers do not support the `request system software add` command. Use the `request vmhost software add` command instead of the `request system software add` command on the PTX10008 and PTX10016 routers to install or upgrade the Junos OS software package or bundle on the router. See request vmhost software add.

When graceful Routing Engine switchover (GRES) is enabled on a device, you must perform a unified in-service software upgrade (ISSU) operation to update the software running on the device. With GRES

enabled, if you attempt to perform a software upgrade by entering the `request system software add` *`package-name`* command, an error message is displayed stating that only in-service software upgrades are supported when GRES is configured. In such a case, you must either remove the GRES configuration before you attempt the upgrade or perform a unified ISSU.

Starting with Junos OS Release 15.1F3, the statement `request system software add` installs a software package for the guest OS only for the PTX5000 router with RE-DUO-C2600-16G, and for MX240, MX480, and MX960 routers with RE-S-1800X4-32G-S.

Starting with Junos OS Release 15.1F5, the statement `request system software add` installs a software package for the guest OS only for the MX2010 and MX2020 routers with REMX2K-1800-32G-S.

On these routers, in order to install both Junos software and host software packages, use the `request vmhost software add` command.

Options

package-name Location from which the software package or bundle is to be installed.

In Junos OS, *`package-name`* can be either the URL of a remote location or the pathname of a local package. But Junos OS Evolved does not support a remote iso for upgrade, so “URL” is removed from the help string in the CLI.

For example:

- **/var/tmp/*package-name***—For a software package or bundle that is being installed from a local directory on the router or switch.
- ***protocol://hostname/pathname/package-name***—For a software package or bundle that is to be downloaded and installed from a remote location. Replace ***protocol*** with one of the following:
  - **ftp**—File Transfer Protocol.
    Use **ftp://*hostname/pathname/package-name***. To specify authentication credentials, use **ftp://<*username*>:<*password*>@*hostname/pathname/package-name***. To have the system prompt you for the password, specify `prompt` in place of the password. If a password is required, and you do not specify the password or `prompt`, an error message is displayed.
  - **http**—Hypertext Transfer Protocol.
    Use **http://*hostname/pathname/package-name***. To specify authentication credentials, use **http://<*username*>:<*password*>@*hostname/pathname/***

package-name. If a password is required and you omit it, you are prompted for it.

- **scp**—Secure copy (not available for limited editions).
  Use **scp://*hostname/pathname/package-name***. To specify authentication credentials, use **scp://<*username*>:<*password*>@*hostname/pathname/package-name***.

- The ***pathname*** in the protocol is the relative path to the user's home directory on the remote system and not the root directory.

- Do not use the **scp** protocol in the `request system software add` command to download and install a software package or bundle from a remote location. The previous statement does not apply to the QFabric switch. The software upgrade is handled by the management process (mgd), which does not support scp.
  Use the `file copy` command to copy the software package or bundle from the remote location to the **/var/tmp** directory on the hard disk:
  file copy scp://*source/package-name* /var/tmp
  Then install the software package or bundle using the `request system software add` command:
  request system software add /var/tmp/*package-name*

best-effort-load (Optional) Activate a partial load and treat parsing errors as warnings instead of errors.

component all (QFabric systems only) (Optional) Install the software package on all of the QFabric components.

delay-restart (Optional) Install a software package or bundle, but do not restart software processes.

device-alias *alias-name* (Junos Fusion only) (Optional) Install the satellite software package onto the specified satellite device using the satellite device's alias name.

force (Optional) Force the addition of the software package or bundle (ignore warnings).

force-host (Optional) Force the addition of the host software package or bundle (ignore warnings) on the QFX5100 device.

member *member-id* (MX Series routers only) (Optional) Install a software package on the specified Virtual Chassis member. Replace *`member-id`* with a value of 0 or 1.

partition (QFX3500 switches and SRX3xx firewalls only) (Optional) Format and repartition the media before installation.

satellite *slot-id*

(Junos Fusion only) (Optional) Install the satellite software package onto the specified satellite device using the satellite devices FPC slot identifier.

no-copy

(Optional) Install a software package or bundle, but do not save copies of the package or bundle files.

no-validate

(Optional) When loading a software package or bundle with a different release, suppress the default behavior of the `validate` option.

To upgrade to Junos OS Release 21.2R1, you cannot use the `validate` option. Instead, choose one of the following options:

- `no-validate`
- `validate-on-host`
- `validate-on-routing-engine`

Software packages from unidentified providers cannot be loaded. To authorize providers, include the `provider-id` statement at the `[edit system extensions provider]` hierarchy level.

re0 | re1

(Optional) On routers or switches that support dual or redundant Routing Engines, load a software package or bundle on the Routing Engine in slot 0 (re0) or the Routing Engine in slot 1 (re1).

reboot

(Optional) After adding the software package or bundle, reboot the system. On a QFabric switch, the software installation is not complete until you reboot the component for which you have installed the software.

set [*package-name1 package-name2*]

(Mixed EX4200 and EX4500 Virtual Chassis, and MX Series routers only) (Optional) Install multiple packages at same time:

- In the case of mixed EX4200 and EX4500 Virtual Chassis, install two software packages—a package for an EX4200 switch and the same release of the package for an EX4500 switch—to upgrade all member switches in a mixed EX4200 and EX4500 Virtual Chassis.
- In the case of MX Series routers, install multiple (two or more) software packages and software add-on packages at the same time. The variable ***package-name*** can either be a list of installation packages, each separated by a blank space, or the full URL to the directory or tar file containing the list of installation packages.

In each case, ***installation-package*** can either be a list of installation packages, each separated by a blank space, or the full URL to the directory or tar file containing the list of installation packages.

Use the `request system software add set` command to retain any SDK configuration by installing the SDK add-on packages along with the core Junos OS installation package.

unlink

(Optional) On MX Series routers, use the unlink option to remove the software package from this directory after a successful upgrade is completed.

upgrade-group [all | *upgrade-group-name*]

(Junos Fusion only) (Required to configure a Junos Fusion using autoconversion or manual conversion) Associate a satellite software image with a satellite software upgrade group. The satellite software package is associated with the specified satellite software upgrade group using the *upgrade-group-name*, or for all satellite software upgrade groups in a Junos Fusion when the all keyword is specified.

A satellite software upgrade group is a group of satellite devices in a Junos Fusion that are designated to upgrade to the same satellite software version using the same satellite software package. See , *Understanding Software in a Junos Fusion Provider EdgeUnderstanding Software in a Junos Fusion Enterprise*, and *Managing Satellite Software Upgrade Groups in a Junos Fusion*.

upgrade-with-config

(Optional) Install one or more configuration files.

Configuration files specified with this option must have the extension **.text** or **.xml** and have the extension specified. Using the extension **.txt** will not work.

validate

(Optional) Validate the software package or bundle against the current configuration as a prerequisite to adding the software package or bundle. This is the default behavior when the software package or bundle being added is a different release.

To upgrade to Junos OS Release 21.2R1, you cannot use the `validate` option. Instead, choose one of the following options:

- `no-validate`
- `validate-on-host`
- `validate-on-routing-engine`

The `validate` option only works on systems that do not have `graceful-switchover` (*GRES*) enabled. To use the `validate` option on a system with GRES, either disable GRES for the duration of the installation, or install using the command `request system software in-service-upgrade`, which requires nonstop active routing *(NSR)* to be enabled when using GRES.

validate-on-host *hostname*

(Optional) Validate the software package by comparing it to the running configuration on a remote Junos OS host. Specify a host, replacing *`hostname`* with the remote

hostname. You can optionally provide the username that will be used to log in to the remote host by specifying the hostname in the format `user@hostname`.

validate-on-routing-engine *routing-engine*

(Optional) Validate the software bundle or package by comparing it to the running configuration on a Junos OS Routing Engine on the same chassis. Specify a Routing Engine, replacing *routing-engine* with the routing engine name.

Additional Information

Before upgrading the software on the device, when you have a known stable system, issue the `request system snapshot` command to back up the software, including the configuration, to the **/altroot** and **/altconfig** file systems. After you have upgraded the software on the device and are satisfied that the new package or bundle is successfully installed and running, issue the `request system snapshot` command again to back up the new software to the **/altroot** and **/altconfig** file systems.

The `request system snapshot` command is currently not supported on the QFabric system. Also, you cannot add or install multiple packages on a QFabric system.

After you run the `request system snapshot` command, you cannot return to the previous version of the software because the running and backup copies of the software are identical.

If you are upgrading more than one package at the same time, delete the operating system package, jkernel, last. Add the operating system package, jkernel, first and the routing software package, jroute, last. If you are upgrading all packages at once, delete and add them in the following order:

```
user@host> request system software add /var/tmp/jbase
user@host> request system software add /var/tmp/jkernel
user@host> request system software add /var/tmp/jpfe
user@host> request system software add /var/tmp/jdocs
user@host> request system software add /var/tmp/jroute
user@host> request system software add /var/tmp/jcrypto
```

Before installing software on a device that has one or more custom YANG data models added to it, back up and remove the configuration data corresponding to the custom YANG data models from the active configuration. For more information see Managing YANG Packages and Configurations During a Software Upgrade or Downgrade.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software add validate

```
user@host> request system software add validate /var/tmp/ jinstall-7.2R1.7-domestic-signed.tgz
Checking compatibility with configuration
Initializing...
Using jbase-7.1R2.2
Using /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Using /var/validate/tmp/jinstall-signed/jinstall-7.2R1.7-domestic.tgz
Using /var/validate/tmp/jinstall/jbundle-7.2R1.7-domestic.tgz
Checking jbundle requirements on /
Using /var/validate/tmp/jbundle/jbase-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jkernel-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jcrypto-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jpfe-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jdocs-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jroute-7.2R1.7.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Validating against /config/rescue.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-7.2R1.7-domestic-signed.tgz' ...
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Adding jinstall...

WARNING:     This package will load JUNOS 7.2R1.7 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.
```

```
Saving the config files ...
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-7.2R1.7-domestic-signed.tgz ...
Saving state for rollback ...
```

request system software add /var/tmp/ no-validate

```
user@host> request system software add no-validate /var/tmp/junos-install-mx-x86-32-15.1R1.9.tgz
Installing package '/var/tmp/junos-install-mx-x86-32-15.1R1.9.tgz' ...
Verified manifest signed by PackageProductionEc_2015
Verified manifest signed by PackageProductionRSA_2015
Verified contents.iso
Verified issu-indb.tgz
Verified junos-x86-32.tgz
Verified kernel
Verified metatags
Verified package.xml
Verified pkgtools.tgz
camcontrol: not found
camcontrol: not found
Verified manifest signed by PackageProductionEc_2015
Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Saving package file in /var/sw/pkg/junos-install-x86-32-
domestic-20150618.043753_builder_junos_151_r1.tgz ...
Saving state for rollback ...
```

request system software add no-copy no-validate reboot

```
user@host> request system software add no-copy no-validate junos-install-srx-x86-64-17.3R1.tgz
reboot
Verified junos-install-srx-x86-64-17.3R1 signed by PackageProductionEc_2017 method
ECDSA256+SHA256
```

```
Verified manifest signed by PackageProductionEc_2017 method ECDSA256+SHA256
Checking PIC combinations
Verified fips-mode signed by PackageProductionEc_2017 method ECDSA256+SHA256
Adding fips-mode-x86-32-20170728.153050_builder_junos_173_r1 ...
Verified jail-runtime signed by PackageProductionEc_2017 method ECDSA256+SHA256
Adding jail-runtime-x86-32-20170725.352915_builder_stable_10 ...
Verified jdocs signed by PackageProductionEc_2017 method ECDSA256+SHA256
Adding jdocs-x86-32-20170728.153050_builder_junos_173_r1 ...
Verified jfirmware signed by PackageProductionEc_2017 method ECDSA256+SHA256
Adding jfirmware-x86-32-17.3R1 ...
Verified jpfe-X signed by PackageProductionEc_2017 method ECDSA256+SHA256
Adding jpfe-X-x86-32-20170728.153050_builder_junos_173_r1 ...
Verified jpfe-X960 signed by PackageProductionEc_2017 method ECDSA256+SHA256
Adding jpfe-X960-x86-32-20170728.153050_builder_junos_173_r1 ...
Verified jpfe-common signed by PackageProductionEc_2017 method ECDSA256+SHA256
Adding jpfe-common-x86-32-20170728.153050_builder_junos_173_r1 ...
Verified jpfe-fips signed by PackageProductionEc_2017 method ECDSA256+SHA256
Verified jpfe-wrlinux signed by PackageProductionEc_2017 method ECDSA256+SHA256
Adding jpfe-wrlinux-x86-32-20170728.153050_builder_junos_173_r1 ...
Verified jsd-jet-1 signed by PackageProductionEc_2017 method ECDSA256+SHA256
Adding jsd-x86-32-17.3R1-jet-1 ...
```

request system software add validate-on-host

```
user@host> request system software add validate-on-host user@xyz :/var/tmp/
jinstall-15.1-20150516_ib_15_2_psd.0-domestic-signed.tgz
user@host> request system software add validate-on-host user@xyz :/var/tmp/
jinstall-15.1-20150516_ib_15_2_psd.0-domestic-signed.tgz
Extracting JUNOS version from package...
Connecting to remote host xyz...
Password:
Sending configuration to xyz...
Validating configuration on xyz...
PACKAGETYPE: not found
Checking compatibility with configuration
Initializing...
Using jbase-15.1-20150516_ib_15_2_psd.0
Verified manifest signed by PackageDevelopmentEc_2015
Using jruntime-15.1-20150516_ib_15_2_psd.0
Verified manifest signed by PackageDevelopmentEc_2015
Using jkernel-15.1-20150516_ib_15_2_psd.0
```

```
Verified manifest signed by PackageDevelopmentEc_2015
Using jroute-15.1-20150516_ib_15_2_psd.0
Verified manifest signed by PackageDevelopmentEc_2015
Using jcrypto-15.1-20150516_ib_15_2_psd.0
Verified manifest signed by PackageDevelopmentEc_2015
Using jweb-15.1-20150516_ib_15_2_psd.0
Verified manifest signed by PackageDevelopmentEc_2015
Using /var/packages/jtools-15.1-20150516_ib_15_2_psd.0
Verified manifest signed by PackageDevelopmentEc_2015
Using /var/tmp/config.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: warning: schema: init: 'logical-systems-vlans' contains-node 'juniper-config vlans': not
found
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-15.1-20150516_ib_15_2_psd.0-domestic-signed.tgz' ...
Verified jinstall-15.1-20150516_ib_15_2_psd.0-domestic.tgz signed by PackageDevelopmentEc_2015
Adding jinstall...

WARNING:     The software that is being installed has limited support.
WARNING:     Run 'file show /etc/notices/unsupported.txt' for details.


WARNING:     This package will load JUNOS 15.1-20150516_ib_15_2_psd.0 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.
```

```
Saving package file in /var/sw/pkg/jinstall-15.1-20150516_ib_15_2_psd.0-domestic-signed.tgz ...
Saving state for rollback ...
```

request system software add (Mixed EX4200 and EX4500 Virtual Chassis)

```
user@switch> request system software add set [/var/tmp/jinstall-ex-4200-11.1R1.1-domestic-
signed.tgz /var/tmp/jinstall-ex-4500-11.1R1.1-domestic-signed.tgz]
...
```

request system software add component all (QFabric Systems)

```
user@switch> request system software add /pbdata/packages/jinstall-qfabric-12.2X50-D1.3.rpm
component all
...
```

request system software add upgrade-group (Junos Fusion)

```
user@aggregation-device> request system software add /var/tmp/satellite-3.0R1.1-signed.tgz
upgrade-group group1
```

request system software add no-validate (SRX Series Firewall)

In Junos OS Release 24.4R1, we re-named the software install images. Images named `junos-install-srxsme*` are recognized by the system to install Junos OS Release 24.4R1 or later. Images named `junos-srxsme*` are recognized by the system to install Junos OS Release 23.4R1 or previous releases.

```
user@host> request system software add /var/tmp/junos-
srxsme-20.4I-20200810_dev_common.0.0833.tgz no-copy no-validate
Formatting alternate root (/dev/ad0s2a)...
/dev/ad0s2a: 600.0MB (1228732 sectors) block size 16384, fragment size 2048
        using 4 cylinder groups of 150.00MB, 9600 blks, 19200 inodes.
super-block backups (for fsck -b #) at:
32, 307232, 614432, 921632
Installing package '/altroot/cf/packages/install-tmp/junos-20.4I-20200810_dev_common.0.0833' ...
Verified junos-boot-srxsme.tgz signed by PackageDevelopmentECP256_2020 method ECDSA256+SHA256
Verified junos-srxsme-domestic signed by PackageDevelopmentECP256_2020 method ECDSA256+SHA256
```

```
Verified manifest signed by PackageDevelopmentECP256_2020 method ECDSA256+SHA256

WARNING:     The software that is being installed has limited support.
WARNING:     Run 'file show /etc/notices/unsupported.txt' for details.

JUNOS 20.4I-20200810_dev_common.0.0833 will become active at next reboot
WARNING: A reboot is required to load this software correctly
WARNING:     Use the 'request system reboot' command
WARNING:         when software installation is complete
Saving state for rollback ...

user@host> request system software add /var/tmp/junos-srxsme-19.4R1.3.tgz no-copy no-validate
WARNING: Package junos-19.4R1.3 version 19.4R1.3 is not compatible with current loader
WARNING: Automatic recovering loader, please wait ...
Upgrading Loader...
####################################
Verifying the loader image... OK
WARNING: The new boot firmware will take effect when the system is rebooted.
WARNING: Loader recover finish.
Formatting alternate root (/dev/ad0s1a)...
/dev/ad0s1a: 598.5MB (1225692 sectors) block size 16384, fragment size 2048
     using 4 cylinder groups of 149.62MB, 9576 blks, 19200 inodes.
super-block backups (for fsck -b #) at:
32, 306464, 612896, 919328
Installing package '/altroot/cf/packages/install-tmp/junos-19.4R1.3' ...
Verified junos-boot-srxsme-19.4R1.3.tgz signed by PackageProductionEc_2019 method
ECDSA256+SHA256
Verified junos-srxsme-19.4R1.3-domestic signed by PackageProductionEc_2019 method
ECDSA256+SHA256
Verified junos-boot-srxsme-19.4R1.3.tgz signed by PackageProductionEc_2019 method
ECDSA256+SHA256 V
erified junos-srxsme-19.4R1.3-domestic signed by PackageProductionEc_2019 method ECDSA256+SHA256
JUNOS 19.4R1.3 will become active at next reboot
WARNING: A reboot is required to load this software correctly
WARNING: Use the 'request system reboot' command
WARNING: when software installation is complete Saving state for rollback ...
```

request system software add (SRX Series Firewall)

In Junos OS Release 24.4R1, we re-named the software install images. Images named `junos-install-srsxme*` are recognized by the system to install Junos OS Release 24.4R1 or later. Images named `junos-srxsme*` are recognized by the system to install Junos OS Release 23.4R1 or previous releases.

```
user@host> request system software add /var/tmp/junos-srxsme-19.4R2.3.tgz
WARNING: Package junos-19.4R2.3 version 19.4R2.3 is not compatible with this system.
WARNING: Please install a package with veloadr support, 20.3 or higher.
```

Release Information

Command introduced before Junos OS Release 7.4.

`best-effort-load` and `unlink` options added in Junos OS Release 7.4.

`set [package-name1 package-name2]` option added in Junos OS Release 11.1 for EX Series switches. Added in Junos OS Release 12.2 for MX Series routers.

On EX Series switches, the `set [package-name1 package-name2]` option allows you to install only two software packages on a mixed EX4200 and EX4500 Virtual Chassis. Whereas, on MX Series routers, the `set [package-name1 package-name2 package-name3]` option allows you to install multiple software packages and software add-on packages at the same time.

`upgrade-with-config` and `upgrade-with-config-format format` options added in Junos OS Release 12.3 for MX Series routers routers, EX Series Ethernet switches, and QFX Series devices.

`device-alias`, `satellite`, `upgrade-group`, and `version` options introduced in Junos OS Release 14.2R3 for Junos Fusion.

`validate-on-host` and `validate-on-routing-engine` options added in Junos OS Release 15.1F3 for PTX5000 routers and MX240, MX480, and MX960 routers.

`upgrade-with-config-format format` option deleted in Junos OS Release 16.1 for MX Series routers routers, EX Series Ethernet switches, and QFX Series devices.

RELATED DOCUMENTATION

Format for Specifying Filenames and URLs in Junos OS CLI Commands

request system software rollback (Junos OS) | **15340**

request system storage cleanup (Junos OS) | **15374**

Installing Software Packages on QFX Series Devices (Junos OS)

Upgrading Software on a QFabric System

Managing Satellite Software Upgrade Groups in a Junos Fusion

request system software add (Maintenance) | **15238**

request system software add (Junos OS Evolved)

IN THIS SECTION

- Syntax | **15233**
- Description | **15233**
- Options | **15234**
- Additional Information | **15236**
- Required Privilege Level | **15237**
- Output Fields | **15237**
- Sample Output | **15237**
- Release Information | **15238**

Syntax

```
request system software add package-name
<force>
<no-confirm>
<validate | no-validate>
<reboot | restart>
```

Description

Install a software package on all Routing Engines in a cluster, as seen in the output of the `show system nodes` operational mode command. The default option is `validate`. We recommend that you always download the software image to **/var/tmp** only. For another way to validate the configuration before trying to install the software package (rather than at the same time), see "request system software validate (Junos OS Evolved)" on page 15365.

Jfirmware packages: Starting in Junos OS Evolved Release 24.4R1, for ACX and PTX routers that support this feature, you can also use this command to install a `jfirmware` software package. When the package is installed, the system makes the latest firmware binaries available on the filesystem, ready for you to upgrade the firmware of individual components using the `request system firmware upgrade` command. The older firmware binaries keep running if you do not update the firmware on the individual components. At any given time, only one `jfirmware` package can be present in the system. Before you install another `jfirmware` package, the older `jfirmware` package needs to be either deleted first or you need to specify the `force` option when you add the package to the system. Installed `jfirmware` packages are not carried over to the next Evo version either during upgrade or during rollback.

For Junos OS Evolved, the `request system software add` command has a built-in feature to not start the upgrade if a reboot is pending after an upgrade or rollback.

After you add a new software image, you must reboot the system to run the new software. If you have added the software image but have not yet rebooted, you can issue the `request system software delete` *`package-name`* command to remove the newly added package and cancel the installation.

Any configuration changes you perform after inputting the `request system software add` command are lost when the system reboots with an upgraded version of Junos OS Evolved.

> **NOTE**: Software packages from unidentified providers cannot be loaded. To authorize providers, include the `provider-id` statement at the `[edit system extensions provider]` hierarchy level.

You can use the `request system software add restart` command to perform a unified in-service software upgrade (ISSU), which upgrades to a more recent version of Junos OS Evolved with minimal disruption of the control plane and data plane traffic.

For information on the valid filename and URL formats, see Format for Specifying Filenames and URLs in Junos OS CLI Commands.

Options

package-name

Location from which the software package or bundle is to be installed. Junos OS Evolved does not support a remote `.iso` file for upgrade, so specify the pathname of a package to be installed from a local directory on the router or switch (for example, **/var/tmp/*package-name***).

Use the `file copy` command to copy the software package or bundle from the remote location to the **/var/tmp** directory on the hard disk:
file copy scp:// *package-name* /var/tmp
Then install the software package or bundle using the `request system software add` command:
request system software add /var/tmp/*package-name*

force

(Optional) Force the addition of the software package or bundle (ignore warnings). The `force` option automatically removes software versions until there is enough space for the new software install.

For Junos OS Evolved, if you are trying to reinstall an already installed application, use the `force` option. The `force` option will cause the program to remove the existing application before reinstalling it.

no-confirm

(Optional) Starting from Junos OS Evolved Release 22.4R3, you can use this option to suppress the prompt that asks you to enter `yes` or `no` before starting a unified ISSU. The upgrade will start without confirmation, but if there is any error that can impact traffic, then the upgrade is automatically cancelled.

no-validate

(Optional) When loading a software package or bundle with a different release, suppress the default behavior of the `validate` option and skip the validation of the configuration. A subsequent reboot can cause the system to lose its configuration if the configuration is not compatible with the new software package. The `no-validate` option should only be used if you have previously issued the `request system software validate` operational mode command on the same target version and target configuration.

reboot

(Optional) After installing the software package, reboot the system.

The `reboot` command is not needed to install third-party applications on devices running Junos OS Evolved.

Starting in Junos OS Evolved 23.4R1, when using the `reboot` option, the Layer 2 applications `l2ald`, `l2ald-agent`, `l2cpd`, and `l2cpd-agent` are initialized and started only if any of the following configuration hierarchy levels contain any configuration statements:

- `[edit interface` *`interface-name`* `unit` *`number`* `family ethernet-switching]`
- `[edit vlans]`
- `[edit routing-instance` *`instance-name`* `instance-type virtual-switch]`
- `[edit routing-instance` *`instance-name`* `instance-type mac-vrf]`
- `[edit protocols l2-learning]`

Additionally, `l2cpd`, and `l2cpd-agent` are initialized and started if the `[edit protocols lldp]` hierarchy level contains any configuration statements.

As a result of this change, if your configuration already contains these configuration statements and you then delete all of the above configuration statements, these Layer 2 applications stop running.

restart (Optional) Starting from Junos OS Evolved Release 22.1R1, you can use this option to perform a unified ISSU. For more details, see "request system software add restart" on page 15242,

> **NOTE**: This is applicable only for QFX5220-CD switches.

validate (Default) When loading a software package or bundle with a different release, validate the candidate software against the current configuration of the node.

Additional Information

Back Up the Software First

Before you upgrade the software on the router or replace one of the Routing Engines, when you have a known stable system, issue the `request system snapshot` command to back up the software, including the configuration, to the **/soft** directory. After you have upgraded the software or have replaced one of the Routing Engines, and are satisfied that the software packages are successfully installed and running, issue the `request system snapshot` command again to back up the software to the **/soft** directory.

After you run the `request system snapshot` command, you cannot return to the previous version of the software because the running and backup copies of the software are identical.

Migrate to GPT Disk Partitioning

Starting in Junos OS Evolved Release 24.2R1, we support migrating to GUID Partition Table (GPT) disk partitioning. GPT is the native disk partitioning scheme used by UEFI BIOSes. GPT is similar to the Master Boot Record (MBR) disk partitioning scheme used by traditional BIOSes. All Junos OS Evolved platforms support GPT natively. However, we default to MBR disk partitioning because Junos OS Evolved was originally ported to systems that used traditional BIOSes.

When you use this command to install a release that supports GPT disk partitioning, you can:

- For new installations, change the default partition scheme for both the primary and secondary disks to GPT immediately (for example, scratch installations to empty disks).
- For existing installations, migrate to GPT disk partitioning for both the primary and secondary disks after a reboot of the system.

For more details about this feature, see *Migrate to GPT Disk Partitioning*.

Unsupported Configurations

Starting in Junos OS Evolved Release 23.4R1, the system accepts unsupported configurations silently, and ignores unsupported configurations during commit. Downgrading to a Junos OS Evolved release prior to 23.2R1 will fail if an unsupported configuration is present on the device. Remove any unsupported configurations to downgrade successfully.

YANG Data Models

Before installing software on a device that has one or more custom YANG data models added to it, back up and remove the configuration data corresponding to the custom YANG data models from the active configuration.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

request system software add

```
user@host-re0> request system software add /var/tmp/ptxinstall.iso
Adding software images. This process can take several minutes. Please be patient...
Download and Validate in Progress
re0: Starting upgrade : /var/tmp/ptxinstall.iso
re0: Upgrade version : junos-evo-install-ptx-x86-64-20.4-202103131143.0-EVO
re0: Running pre-checks for 'junos-evo-install-ptx-x86-64-20.4-202103131143.0-EVO'
re0: Pre-checks pass successfully, copying files to software area
re0: Running post install commands...
re0: Post install sequence was successful.
re0: Validating existing configs. See /var/log/validation_config.log for config validation logs.
re0: Validation Passed
re0: Going ahead with Installation
re0: Boot version is now 'junos-evo-install-ptx-x86-64-20.4-202103131143.0-EVO'
re0: Updating all nodes...
re1: Running pre-checks for 'junos-evo-install-ptx-x86-64-20.4-202103131143.0-EVO'
re1: Pre-checks pass successfully, copying files to software area
```

```
re1: Running post install commands...
re1: Post install sequence was successful.
re1: Boot version is now 'junos-evo-install-ptx-x86-64-20.4-202103131143.0-EVO'
re0: Other nodes have been updated successfully
re0: Cluster wide installation was successful
Image validation and installation succeeded.
WARNING: NOTE: A reboot is required to start using the new software.
WARNING: Use the 'request system reboot' command when ready.
```

Release Information

The following options are deprecated in Junos OS Evolved Release 18.3R1: `best-effort-load`, `delay-restart`, `no-copy`, `on-primary`, (`re0` | `re1`), `set`, `unlink`, `validate`, `validate-on-host`, and `validate-on-routing-engine`.

Support for jfirmware software packages introduced in Junos OS Evolved 24.4R1.

RELATED DOCUMENTATION

- request system snapshot (Junos OS Evolved) | **15187**
- request system software delete (Junos OS Evolved) | **15254**
- request system software rollback (Junos OS Evolved) | **15345**
- request system software sync | **15348**
- request system software validate (Junos OS Evolved) | **15365**
- request system software add restart | **15242**

request system software add (Maintenance)

IN THIS SECTION

- Syntax | **15239**
- Description | **15239**
- Options | **15239**
- Required Privilege Level | **15239**
- Release Information | **15239**

Syntax

```
request system software add package-name
```

Description

Use this command to install new software package on the device, for example: **request system software add junos-srxsme-10.0R2-domestic.tgz no-copy no-validate partition reboot**.

Options

- `delay-restart`—Install the software package but does not restart the software process.
- `best-effort-load`—Activate a partial load and treat parsing errors as warnings instead of errors.
- `no-copy`—Install the software package but does not saves the copies of package files.
- `no-validate`—Do not check the compatibility with current configuration before installation starts.
- `partition`—Format and re-partition the media before installation.
- `reboot`—Reboot the device after installation is completed.
- `unlink`—Remove the software package after successful installation.
- `validate`—Check the compatibility with current configuration before installation starts.

Required Privilege Level

maintenance

Release Information

Partition option introduced in the command in Junos OS Release 10.1.

request system software add optional://auto-snapshot

IN THIS SECTION

- Syntax | 15240
- Description | 15240
- Additional Information | 15240
- Required Privilege Level | 15241
- Output Fields | 15241
- Sample Output | 15241
- Release Information | 15241

Syntax

```
request system software add optional://auto-snapshot
```

Description

Activates the optional auto-snapshot package to generate a recovery snapshot after the next reboot. Issue this command before you upgrade the software. Then, when the device reboots after the software upgrade, the system automatically generates a recovery snapshot. Enabling auto-snapshot ensures a recovery image for the latest installed image.

Additional Information

The `request system software add optional://auto-snapshot` command activates the auto-snapshot package. Auto-snapshot performs a recovery snapshot during the next reboot, and stores the recovery snapshot on the OAM volume.

Enabling auto-snapshot on an active Routing Engine might extend the time for the next reboot. Auto-snapshot uses system resources only during this reboot to generate the recovery snapshot. In this way, it avoids resource contention with other applications, such as the routing protocol process (rpd).

You can verify the results of auto-snapshot with the `show system snapshot` command after the software installation completes.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software add optional://auto-snapshot

```
user@host> request system software add optional://auto-snapshot
Verified auto-snapshot signed by PackageDevelopmentECP256_2022 method ECDSA256+SHA256
```

show system snapshot

```
user@host> show system snapshot

No non-recovery snapshots available on the Junos volume

Recovery Snapshots:
Snapshots available on the OAM volume:
recovery.ufs
Date created: Sat Sep  3 15:33:10 CEST 2022
Junos version: 22.1-20220902.421

Total recovery snapshots: 1
```

Release Information

Command introduced in Junos OS Release 22.1R2.

RELATED DOCUMENTATION

request system software add (Junos OS) | 15216

request system software add restart

IN THIS SECTION

- Syntax | 15242
- Description | 15242
- Options | 15242
- Required Privilege Level | 15242
- Sample Output | 15243
- Release Information | 15245

Syntax

```
request system software add package-name restart
```

Description

Perform a unified in-service software upgrade (unified ISSU) to a more recent version of Junos OS Evolved. A unified ISSU involves minimal disruption of the control plane and data plane traffic. The software is upgraded by using an application-level restart or warm restart instead of a reboot, when possible. Use the `request system software validate-restart` command before using the `request system software add restart` command to determine whether you need an application-level restart or a reboot.

NOTE: You can currently use this command only on QFX5220-32CD switches.

Options

package-name Path and filename of the software package to be installed.

Required Privilege Level

maintenance

Sample Output

request system software add restart

```
user@host> request system software add /var/home/regress/junos-evo-install-qfx-ms-
x86-64-22.1-202203012330.0-EVO.iso restart
Adding software images. This process can take several minutes. Please be patient...
Download and Validate in Progress
re0: Starting upgrade : /var/home/regress/junos-evo-install-qfx-ms-x86-64-22.1-202203012330.0-
EVO.iso
re0: Upgrade version : junos-evo-install-qfx-ms-x86-64-22.1-202203012330.0-EVO.iso
re0: Running pre-checks for 'junos-evo-install-qfx-ms-x86-64-22.1-202203012330.0-EVO.iso'
re0: Pre-checks pass successfully, copying files to software area
re0: Running post install commands...
re0: Post install sequence was successful.
re0: Validating existing configs. See /var/log/validation_config.log for config validation logs.
re0: Validation Passed
re0: Validating in-service-upgrade application configs. See /var/log/validation_appconfig.log
for in-service-upgrade application configs validation logs.
re0: In-service-upgrade application configs validation Passed
re0: Validating in-service-upgrade SDK compatibility. See /var/log/validation_sdk.log for in-
service-upgrade SDK compatibility validation logs.
re0: In-service-upgrade SDK Validation Passed
re0: Going ahead with Installation
re0: Boot version is now 'junos-evo-install-qfx-ms-x86-64-22.1-202203012330.0-EVO.iso'
re0: Generating local impact report...
re0: Installation was successful
Image validation and installation succeeded. Restarting Applications.

*** Restart Apps list ***
distributord

*** Applications that do not support restart upgrade ***
distributord

This platform supports in-service kernel warm restart upgrade.

Enter yes to proceed with in-service kernel warm restart or no to proceed with the reboot
upgrade.
Proceed with in-service kernel warm restart upgrade ? [yes,no] (yes) yes
```

```
----- Impact report for kernel warm restart upgrade -----

Actions prior to warm restart:

*** Applications that need prep to upgrade ***
rpdagent

*** Applications that need prep to upgrade final ***
evo-pfemand
l2ald-agent
l2cpd-agent
ndp
picd
rpdagent

Actions post warm restart:

*** Applications that need sw sync ***
evo-pfemand

*** Applications that need hw sync ***
evo-pfemand

*** Applications that need unprep to upgrade ***
evo-pfemand
l2cpd-agent
ndp
picd
rpdagent

Sending prepare notification to app rpdagent on node re0
Prepare to upgrade succeeded for app rpdagent on node re0
Sending prepare final notification to app evo-pfemand on node re0
Sending prepare final notification to app l2ald-agent on node re0
Sending prepare final notification to app l2cpd-agent on node re0
Sending prepare final notification to app ndp on node re0
Sending prepare final notification to app picd on node re0
Sending prepare final notification to app rpdagent on node re0
Prepare to upgrade succeeded for app l2ald-agent on node re0
Prepare to upgrade succeeded for app picd on node re0
Prepare to upgrade succeeded for app evo-pfemand on node re0
Prepare to upgrade succeeded for app l2cpd-agent on node re0
Prepare to upgrade succeeded for app rpdagent on node re0
```

```
Prepare to upgrade succeeded for app ndp on node re0
Saving system snapshot and rebooting. See /var/log/issu.log for ISSU logs
```

Release Information

Command introduced in Junos OS Evolved Release 22.1R1.

RELATED DOCUMENTATION

request system software validate-restart (Junos OS Evolved) | 15371

request system software configuration-backup

IN THIS SECTION

- Syntax | 15245
- Description | 15245
- Options | 15246
- Required Privilege Level | 15246
- Output Fields | 15246
- Sample Output | 15246
- Release Information | 15246

Syntax

```
request system software configuration-backup (path)
```

Description

Use this command to save the currently active configuration and any installation-specific parameters such as a configuration that you have entered outside of the CLI, director group IP addresses, and the default partition IP address.

Options

path—(QFabric System) Provide the path to the location of the backup configuration files. You can save the backup configuration files to either a URL, local directory, remote server, or removable drive.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software configuration-backup

```
user@switch request system software configuration-backup ftp://ftp.test.net/test
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100        4035     0    0    100 4035    0      138k --:--:-- --:--:-- --:--:--  0
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

request system software configuration-restore | 15247

request system software configuration-restore

IN THIS SECTION

- Syntax | 15247
- Description | 15247
- Options | 15247
- Required Privilege Level | 15247
- Output Fields | 15248
- Sample Output | 15248
- Release Information | 15248

Syntax

```
request system software configuration-restore (path)
```

Description

Use this command to restore a previously saved configuration and any installation-specific parameters, such as a configuration that you have entered outside of the CLI, director group IP addresses, and the default partition IP address.

Options

`path`—(QFabric System) Provide the path to the location of the backup configuration files. The path can be to a local file, a file on an external flash drive, or an SCP or FTP destination.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software configuration-restore

```
user@switch request system software configuration-restore ftp://ftp.test.net/test
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100  4035  100  4035    0     0   153k      0 --:--:-- --:--:-- --:--:-- 3803k
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

request system software configuration-backup | 15245

Performing a QFabric System Recovery Installation on the Director Group

request system software delete-backup

IN THIS SECTION

- Syntax | 15249
- Description | 15249
- Required Privilege Level | 15249
- Output Fields | 15249
- Sample Output | 15249
- Release Information | 15249

Syntax

```
request system software delete-backup
```

Description

(Junos OS only) Use this command to delete the backup image previously created. Junos OS keeps a backup image of the software that was previously installed so that you can downgrade to that version of the software if necessary. If you delete this image, you cannot downgrade to the particular version of the software. Note that this command does not perform any function on SRX Series Firewalls with the dual-root partitioning scheme (units that are shipped with Junos OS 10.0 version or later).

In devices running FreeBSD Release 12 or later, you cannot issue this command. To learn if your device is running FreeBSD Release 12 or later, issue the `show version` command and look for the `fbsd_builder_stable` string in the module names. If the string includes the number 12 or later, your device is running FreeBSD Release 12 or later.

Required Privilege Level

maintenance

Output Fields

Sample Output

request system software system-backup

```
user@host> request system software delete-backup
Delete backup system software package [yes,no]
```

Release Information

Command introduced in Junos OS Release 9.6R1.

RELATED DOCUMENTATION

request system software delete (Junos OS) | 15250

request system software delete (Junos OS)

IN THIS SECTION

- Syntax | 15250
- Syntax (Junos Fusion) | 15250
- Description | 15250
- Options | 15251
- Additional Information | 15252
- Required Privilege Level | 15252
- Output Fields | 15252
- Sample Output | 15252
- Release Information | 15254

Syntax

```
request system software delete  software-package
<force>
<re0 | re1>
<reboot>
<set [package-name package-name]>
```

Syntax (Junos Fusion)

```
request system software delete
<upgrade-group [all |upgrade-group-name]>
<version version-string>
```

Description

Use this command to remove a software package or bundle from the device.

> **CAUTION**: Before removing a software package or bundle, make sure that you have already placed the new software package or bundle that you intend to load onto the device.

Options

software-package

Software package or bundle name. You can see this software package name by using the `show system software` command. Type the software package name explicitly and do not use the tab key to auto-complete the command.

You can delete any or all of the following software bundles or packages:

- `jbase`—(Optional) Junos base software suite
- `jcrypto`—(Optional, in domestic version only) Junos security software
- `jdocs`—(Optional) Junos online documentation file
- `jkernel`—(Optional) Junos kernel software suite
- `jpfe`—(Optional) Junos Packet Forwarding Engine support
- `jroute`—(Optional) Junos routing software suite
- `junos`—(Optional) Junos base software

On EX Series switches, some of the package names are different than those listed. To see the list of packages that you can delete on an EX Series switch, enter the command **show system software**.

force

(Optional) Ignore warnings and force removal of the software.

re0 | re1

(Optional) On routers or switches that support dual or redundant Routing Engines, delete a software package or bundle on the Routing Engine in slot 0 (re0) or the Routing Engine in slot 1 (re1).

reboot

Automatically reboot upon completing the `request system software delete` command.

`set` [*`package-name package-name`*]

(Optional) Install multiple software packages or software add-on packages at the same time.

upgrade-group [**all** |*upgrade-group-name*]

(Junos Fusion only) Delete the satellite software image association with the specified satellite software upgrade group.

A satellite software upgrade group is a group of satellite devices in the same Junos Fusion that are designated to upgrade to the same satellite software version using the same satellite software package.

version *version-string*

(Junos Fusion only) (Optional) Delete a satellite software package association with a satellite software upgrade group by selecting the satellite software package's version.

Additional Information

Before upgrading the software on the router or switch, when you have a known stable system, issue the `request system snapshot` command to back up the software, including the configuration, to the /altroot and /altconfig file systems (on routers) or the /, /altroot, /config, /var, and /var/tmp file systems (on switches). After you have upgraded the software on the router or switch and are satisfied that the new packages are successfully installed and running, issue the `request system snapshot` command again to back up the new software to the /altroot and /altconfig file systems (on routers) or the /, /altroot, /config, /var, and /var/tmp file systems (on switches). After you run the `request system snapshot` command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Devices that support Junos Fusion must be configured as part of a Junos Fusion deployment before you can use Junos Fusion commands. For platforms that support Junos Fusion, see Understanding Junos Fusion Enterprise Components.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software delete jdocs

The following example displays the system software packages before and after the `jdocs` package is deleted through the `request system software delete` command:

```
user@host> show system software
Information for jbase:
```

```
Comment:
JUNOS Base OS Software Suite [7.2R1.7]


Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]


Information for jdocs:

Comment:
JUNOS Online Documentation [7.2R1.7]


Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

...


user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]


Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]


Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]
```

```
...
```

Release Information

Command introduced before Junos OS Release 7.4.

`set [package-name package-name]` option added in Junos OS Release 12.2.

`reboot` option introduced in Junos OS Release 12.3.

`upgrade-group` and `version` options introduced in Junos OS Release 14.2R3 for Junos Fusion.

RELATED DOCUMENTATION

- request system software add (Junos OS) | **15216**
- request system software rollback (Junos OS) | **15340**
- Managing Satellite Software Upgrade Groups in a Junos Fusion

request system software delete (Junos OS Evolved)

IN THIS SECTION

- Syntax | **15255**
- Description | **15255**
- Options | **15255**
- Additional Information | **15256**
- Required Privilege Level | **15256**
- Output Fields | **15256**
- Sample Output | **15256**
- Release Information | **15258**

Syntax

```
request system software delete
<force>
<package-name>
<archived>
<all-third-party-packages>
```

Description

Use this command to remove a software package from the device, as long as it is not the software version currently running on the system. The `force` option is required if the requested version is the rollback version.

> **CAUTION**: Before removing a software package, make sure that you have already placed the new software package that you intend to load onto the device, in the **/var/tmp** directory.

If you have added a new software package, but you have not yet rebooted the system to complete the installation of the package, use this command to remove the newly added package and cancel the installation.

Starting in Junos OS Evolved Release 24.4R1, for ACX and PTX routers that support this feature, you can use this command to delete a `jfirmware` software package. To install a newer `jfirmware` package, you must first delete the `jfirmware` package already on the device.

Options

package-name Name of the Junos OS Evolved package running on the device. Type the package-name explicitly and do not use the tab key to auto-complete the command. You can see this package name by issuing the `show system software list` operational mode command.

all-third-party-packages (Optional) Delete all third-party software on the device.

archived (Optional) Delete all archived software versions except the current and the rollback versions. When there is a pending next-boot software version, you must reboot the system to finish installing that version, or delete just that version, before you can delete any versions with this option. If the other Routing Engine in the system has more images than the one from which you are issuing the command, the older versions are deleted

from the current Routing Engine and the current and the rollback versions are synchronized to the other Routing Engine.

force (Optional) Ignore warnings and force removal of the software. The `force` option is required if the requested version is the rollback version.

Additional Information

Before you upgrade the software on the router or replace one of the Routing Engines, when you have a known stable system, issue the `request system snapshot` command to back up the software, including the configuration, to the **/soft** directory. After you have upgraded the software or have replaced one of the Routing Engines, and are satisfied that the software packages are successfully installed and running, issue the `request system snapshot` command again to back up the software to the **/soft** directory.

After you run the `request system snapshot` command, you cannot return to the previous version of the software because the running and backup copies of the software are identical.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

request system software delete (Junos OS Evolved)

```
user@host> request system software delete junos-evo-install-qfx-fixed-
x86-64-18.3I20180911102422
Removing version 'junos-evo-install-qfx-fixed-x86-64-18.3I20180911102422'.
Software ... done.
Data ... done.
Version 'junos-evo-evo-qfx-fixed-x86-64-18.3I20180911102422' removed successfully.
```

request system software delete archived

```
user@host-re0> request system software delete archived
ALERT:  This command will delete all archived SW versions except current and rollback.
        Do you want to proceed? [yes,no] (no) yes

Software delete in progress...
re0: Executing Software delete...
re0: Cannot delete junos-evo-install-ptx-x86-64-20.4-202103141559.0-EVO - It is the current
version
re0: Rollback or scratch install
re0: Removing version junos-evo-install-ptx-x86-64-20.4-202103150459.0-EVO...
re0: Removing version junos-evo-install-ptx-x86-64-20.4-202103111254.0-EVO...
re1: Removing version junos-evo-install-ptx-x86-64-20.4-202103150459.0-EVO...
re1: Removing version junos-evo-install-ptx-x86-64-20.4-202103111254.0-EVO...
Image deletion succeeded.
```

request system software delete archived (with Next-boot Software Version in System)

```
user@host-re0> request system software delete archived
ALERT: This command will delete all archived SW versions except current and rollback.
Do you want to proceed? [yes,no] (no) yes

re0: Software delete cannot proceed as reboot is pending after upgrade/rollback to junos-evo-
install-ptx-x86-64-20.4I20210212000536-EVO.
re0: Please reboot before doing delete operation.
re0: Or Delete junos-evo-install-ptx-x86-64-20.4I20210212000536-EVO using 'request system
software delete'.
re0: Run 'show system software list' to get all installed software versions
Image deletion failed.
```

request system software delete archived (with Only a Current and Rollback Version Available)

```
user@host-re0> request system software delete archived
ALERT: This command will delete all archived SW versions except current and rollback.
Do you want to proceed? [yes,no] (no) yes
```

```
re0: Only minimal set of software versions exists. Cannot delete Current or Rollback versions.
Image deletion failed.
```

Release Information

`all-third-party-packages` option introduced in Junos OS Evolved Release 19.4R2.

`archived` option added in Junos OS Evolved Release 20.4R2.

Support for deleting `jfirmware` packages introduced in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

- request system software add (Junos OS Evolved) | **15233**
- request system software rollback (Junos OS Evolved) | **15345**
- request system software sync | **15348**
- request system software validate (Junos OS Evolved) | **15365**
- show system software list | **27088**

request system software download

IN THIS SECTION

- Syntax (QFabric System) | **15259**
- Description | **15259**
- Options | **15259**
- Additional Information | **15259**
- Required Privilege Level | **15259**
- Output Fields | **15259**
- Sample Output | **15260**
- Release Information | **15260**

Syntax (QFabric System)

```
request system  software download path/package-name
```

Description

Use this command to download a software package from a location on the director device, mounted external USB flash drive, remote FTP or SCP location, or other location.

Options

path Location where the software package is located. For example:

- `/pbdata/packages/package-name`—For a software package that is being installed from a local directory on the switch.
- ***protocol://hostname/pathname/package-name***—For a software package or bundle that is to be downloaded and installed from a remote location. Replace *protocol* with one of the following:
  - `ftp`—File Transfer Protocol.
    Use **ftp://*hostname/pathname/package-name***. To specify authentication credentials, use **ftp://<*username*>:<*password*>@*hostname/pathname/package-name***. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
  - `scp`—Secure copy (available only for Canada and U.S. version).
    Use **scp://*hostname/pathname/package-name***. To specify authentication credentials, use **scp://<*username*>:<*password*>@*hostname/pathname/package-name***.

Additional Information

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software download

```
user@switch> request system software download ftp://ftp.install-directory/jinstall-
qfabric-11.3X30.6.rpm
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100  186M  100  186M    0     0  18.4M      0  0:00:10  0:00:10 --:--:-- 18.6M
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

- request system software add (Junos OS) | **15216**
- request system software delete (Junos OS) | **15250**
- request system software rollback (Junos OS) | **15340**
- request system storage cleanup (Junos OS) | **15374**
- *Installing Software Packages on QFX Series Devices*

request system software format-qfabric-backup

IN THIS SECTION

- Syntax | **15261**
- Description | **15261**
- Options | **15261**
- Required Privilege Level | **15261**
- Output Fields | **15261**
- Sample Output | **15261**
- Release Information | **15261**

Syntax

```
request system software format-qfabric-backup
```

Description

(QFabric systems only) Copy the install media files from a USB flash drive to your QFabric system recovery directory on the Director group. You must issue this command before you can use the `request system software system-backup` and `request system software system-backup copy-to-usb` commands.

Options

none Copy the install media files from a USB flash drive to a Director group recovery directory.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software format-qfabric-backup

```
user@qfabric> request system software format-qfabric-backup
Copying QFabric USB template image from /dev/sdb(Unigen,PQS4000,4009 MB)......
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Performing System Backup and Recovery for a QFabric System

request system software system-backup | 15353

Performing a QFabric System Recovery Installation on the Director Group

request system software in-service-upgrade

IN THIS SECTION

- Syntax | 15262
- Syntax | 15262
- Description | 15263
- Options | 15263
- Additional Information | 15264
- Required Privilege Level | 15265
- Output Fields | 15265
- Sample Output | 15265
- Release Information | 15282

Syntax

```
request system software in-service-upgrade package-name
<no-old-master-upgrade>
<reboot>
<status>
<enhanced-mode>
<no-validate>
<handle-incompatible-config>
```

Syntax

Syntax (QFX Series)

```
request system software in-service-upgrade package-name
```

Description

Perform a unified in-service software upgrade (ISSU). A unified ISSU enables you to upgrade from one Junos OS release to another with no disruption on the control plane and with minimal disruption of traffic.

On QFX5100 and QFX5200 switches, enable nonstop active routing (NSR) and nonstop bridging (NSB).

Options

package-name Location from which the software package or bundle is to be installed. For example:

- `/var/tmp/`*`package-name`*— For a software package or bundle that is being installed from a local directory on the router.
- *`protocol`*`://`*`hostname`*`/`*`pathname`*`/`*`package-name`*—For a software package or bundle that is to be downloaded and installed from a remote location. Replace *`protocol`* with one of the following:
  - `ftp`—File Transfer Protocol
  - `http`—Hypertext Transfer Protocol
  - `scp`—Secure copy (available only for Canada and U.S. version)

no-old-master-upgrade (Optional) When the `no-old-master-upgrade` option is included, after the backup Routing Engine is rebooted with the new software package and a switchover occurs to make it the new primary Routing Engine, the former primary (new backup) Routing Engine will not be upgraded to the new software. In this case, you must manually upgrade the former primary (new backup) Routing Engine. If you do not include the `no-old-master-upgrade` option, the system will automatically upgrade the former primary Routing Engine.

> **NOTE**: This option is not available on QFX5100 and QFX5200 switches.

reboot (Optional) When the `reboot` option is included, the former primary (new backup) Routing Engine is automatically rebooted after being upgraded to the new software. When the `reboot` option is not included, you must manually reboot the former primary (new backup) Routing Engine using the `request system reboot` command.

> **NOTE**: This option is not available on the QFX Series switches.

status

(Optional) Starting in Junos OS Release 19.4R1, use this option to display the status of a unified ISSU during the upgrade. You will need to run this command on the Routing Engine where the ISSU was triggered to display the correct ISSU log file.

> **NOTE**: This option is only available on MX240, MX480, MX960, MX2010, MX2020, PTX3000, and PTX5000 routers.

enhanced-mode

(Optional) Starting in Junos OS Release 20.1R1, runs unified ISSU in enhanced mode, a mode of ISSU that eliminates packet loss during the unified ISSU process.

> **NOTE**: This option is only available on MPC8E, and MPC9E line cards. Support for MPC11E line cards is provided from Junos OS Release 21.4R1 onwards.

no-validate

(Optional) Starting in Junos OS Release 20.4R2, skip the validation step of the unified ISSU process.

handle-incompatible-config

(Optional) Starting in Junos OS Release 22.1R1, automatically deactivate the incompatible configurations at the start of an ISSU and reactivate them once the ISSU is completed. If the ISSU is aborted, the deactivated configurations are automatically reactivated. This is currently applicable only to clock synchronization configurations for Precision Time Protocol (PTP) and Synchronous Ethernet.

> **NOTE**: This option is available only for the MX960, MX10003, MX10008, MX10016, MX2010, and MX2020 routers.

Additional Information

The following conditions apply to unified ISSUs:

- Unified ISSU is not supported on every platform. For a list of supported platforms, see *Unified ISSU System Requirements*.
- Unsupported PICs are restarted during a unified ISSU on certain routing devices. For information about supported PICs, see the Junos OS High Availability User Guide.

- Unsupported protocols will experience packet loss during a unified ISSU. For information about supported protocols, see the Junos OS High Availability User Guide.
- During a unified ISSU, you cannot bring any PICs online or offline on certain routing devices.
- The `no-validate` option is required when upgrading to Junos OS Releases 21.2R1 or 22.1R1. If you are upgrading from a release prior to Junos OS Release 20.4R2, then you will need to upgrade to a release that supports the `no-validate` option before using unified ISSU to upgrade to Junos OS Releases 21.2R1 or 22.1R1.

For more information, see the Junos OS High Availability User Guide.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software in-service-upgrade reboot

```
{master}
user@host> request system software in-service-upgrade /var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz
reboot
ISSU: Validating Image
PIC 0/3 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

ISSU: Preparing Backup RE
Pushing bundle to re1
Checking compatibility with configuration
Initializing...
Using jbase-9.0-20080114.2
Verified manifest signed by PackageProduction_9_0_0
Using /var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Using jinstall-9.0-20080114.2-domestic.tgz
Using jbundle-9.0-20080114.2-domestic.tgz
Checking jbundle requirements on /
```

```
Using jbase-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jkernel-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jcrypto-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jpfe-9.0-20080114.2.tgz
Using jdocs-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jroute-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz' ...
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

WARNING:     This package will load JUNOS 9.0-20080114.2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-9.0-20080114.2-domestic-signed.tgz ...
Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE
```

```
Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU started
ISSU: Backup RE Prepare Done
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
Installing package '/var/tmp/paKEuy' ...
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

WARNING:     This package will load JUNOS 9.0-20080114.2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
```

```
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-9.0-20080114.2-domestic-signed.tgz ...
cp: /var/tmp/paKEuy is a directory (not copied).
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
Reboot consistency check bypassed - jinstall 9.0-20080114.2 will complete installation upon
reboot
[pid 30227]

*** FINAL System shutdown message from root@host ***

System going down IMMEDIATELY

Connection to host closed.
```

request system software in-service-upgrade reboot (TX Matrix Plus Router)

```
{master}
user@host> request system software in-service-upgrade /var/tmp/jinstall-12.3R2-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
PIC 8/1 will be offlined (In-Service-Upgrade not supported)
PIC 19/2 will be offlined (In-Service-Upgrade not supported)
PIC 15/3 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

Checking compatibility with configuration
Initializing...
Using jbase-12.3R2
Verified manifest signed by PackageProduction_12_3_0
Using /var/tmp/jinstall-12.3R2-domestic-signed.tgz
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Using jinstall-12.3R2-domestic.tgz
Using jbundle-12.3R2-domestic.tgz
Checking jbundle requirements on /
Using jbase-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
```

```
Verified jbase-12.3R2 signed by PackageProduction_12_3_0
Using /var/validate/chroot/tmp/jbundle/jboot-12.3R2.tgz
Using jcrypto-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jcrypto-12.3R2 signed by PackageProduction_12_3_0
Using jdocs-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jdocs-12.3R2 signed by PackageProduction_12_3_0
Using jkernel-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jkernel-12.3R2 signed by PackageProduction_12_3_0
Using jpfe-12.3R2.tgz
WARNING: jpfe-12.3R2.tgz: not a signed package
WARNING: jpfe-common-12.3R2.tgz: not a signed package
Verified jpfe-common-12.3R2 signed by PackageProduction_12_3_0
WARNING: jpfe-T-12.3R2.tgz: not a signed package
Verified jpfe-T-12.3R2 signed by PackageProduction_12_3_0
Using jplatform-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jplatform-12.3R2 signed by PackageProduction_12_3_0
Using jroute-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jroute-12.3R2 signed by PackageProduction_12_3_0
Using jruntime-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jruntime-12.3R2 signed by PackageProduction_12_3_0
Using jservices-12.3R2.tgz
Using jservices-crypto-12.3R2.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
ISSU: Preparing LCC Backup REs
Pushing bundle to lcc0-re1
Pushing bundle to lcc1-re1
Pushing bundle to lcc2-re1
Pushing bundle to lcc3-re1
Pushing bundle to sfc0-re1
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0
```

```
WARNING:     This package will load JUNOS 12.3R2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:     This package will load JUNOS 12.3R2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.
```

```
Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:     This package will load JUNOS 12.3R2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:     This package will load JUNOS 12.3R2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...
```

```
WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
ISSU: Preparing SFC Backup RE
NOTICE: Validating configuration against jinstall-12.3R2-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Checking compatibility with configuration
Initializing...
Using jbase-12.3R2
Verified manifest signed by PackageProduction_12_3_0
Using /var/tmp/jinstall-12.3R2-domestic-signed.tgz
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Using jinstall-12.3R2-domestic.tgz
Using jbundle-12.3R2-domestic.tgz
Checking jbundle requirements on /
Using jbase-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jbase-12.3R2 signed by PackageProduction_12_3_0
Using /var/validate/chroot/tmp/jbundle/jboot-12.3R2.tgz
Using jcrypto-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jcrypto-12.3R2 signed by PackageProduction_12_3_0
Using jdocs-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jdocs-12.3R2 signed by PackageProduction_12_3_0
Using jkernel-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jkernel-12.3R2 signed by PackageProduction_12_3_0
Using jpfe-12.3R2.tgz
WARNING: jpfe-12.3R2.tgz: not a signed package
WARNING: jpfe-common-12.3R2.tgz: not a signed package
Verified jpfe-common-12.3R2 signed by PackageProduction_12_3_0
WARNING: jpfe-T-12.3R2.tgz: not a signed package
Verified jpfe-T-12.3R2 signed by PackageProduction_12_3_0
Using jplatform-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jplatform-12.3R2 signed by PackageProduction_12_3_0
```

```
Using jroute-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jroute-12.3R2 signed by PackageProduction_12_3_0
Using jruntime-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jruntime-12.3R2 signed by PackageProduction_12_3_0
Using jservices-12.3R2.tgz
Using jservices-crypto-12.3R2.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:     This package will load JUNOS 12.3R2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
SFC Backup upgrade done
Rebooting SFC Backup RE

Rebooting sfc0-re1
ISSU: SFC Backup RE Prepare Done
Waiting for SFC Backup RE reboot
```

```
Rebooting lcc0-re1
Rebooting LCC [lcc0-re1]

Rebooting lcc1-re1
Rebooting LCC [lcc1-re1]

Rebooting lcc2-re1
Rebooting LCC [lcc2-re1]

Rebooting lcc3-re1
Rebooting LCC [lcc3-re1]
LCC Backup REs have rebooted
Waiting for LCC Backup REs come back online
ISSU: LCC Backup REs Prepare Done
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
lcc0-re0:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 1          Online (ISSU)
    PIC 0        Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
    PIC 1        Online (ISSU)
  FPC 4          Online (ISSU)
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)

lcc1-re0:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online (ISSU)
    PIC 3        Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
```

```
  FPC 4          Online (ISSU)
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)

lcc2-re0:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
    PIC 0        Online (ISSU)
  FPC 4          Online (ISSU)
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)
    PIC 1        Online (ISSU)

lcc3-re0:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online (ISSU)
    PIC 0        Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
    PIC 2        Online (ISSU)
  FPC 4          Online (ISSU)
  FPC 5          Online (ISSU)
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)
    PIC 1        Online (ISSU)

lcc0-re0:
--------------------------------------------------------------------------
Resolving mastership...
Complete. The other routing engine becomes the master.

lcc1-re0:
--------------------------------------------------------------------------
Resolving mastership...
Complete. The other routing engine becomes the master.

lcc2-re0:
--------------------------------------------------------------------------
```

```
Resolving mastership...
Complete. The other routing engine becomes the master.

lcc3-re0:
--------------------------------------------------------------------------
Resolving mastership...
Complete. The other routing engine becomes the master.
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading SFC Old Master RE

lcc0-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:     This package will load JUNOS 12.3R2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...

lcc1-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
```

```
Verified manifest signed by PackageProduction_12_3_0

WARNING:     This package will load JUNOS 12.3R2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...

lcc2-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:     This package will load JUNOS 12.3R2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
```

```
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...

lcc3-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:     This package will load JUNOS 12.3R2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/paBWTg' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:     This package will load JUNOS 12.3R2 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
```

```
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed ...
cp: /var/tmp/paBWTg is a directory (not copied).
Saving state for rollback ...
ISSU: SFC Old Master Upgrade Done
ISSU: IDLE
```

request system software in-service-upgrade (QFX5100 Switch)

```
{master}
user@switch> request system software in-service-upgrade /var/tmp/jinstall-qfx-132_x51_vjunos.0-domestic.tgz
ISSU: Validating Image
Prepare for ISSU
spawn the backup VM
ISSU: Preparing Backup RE
Backup upgrade done
ISSU: Backup RE Prepare Done
waiting for backup RE switchover ready
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
```

```
  Item           Status                  Reason
  FPC 0          Online (ISSU)
send ISSU done to chassisd on backup VM
Chassis ISSU Completed
ISSU: IDLE
mgd_package_opus_issu: Initiate em0 device handoff
```

request system software in-service-upgrade status

```
{master}
user@host> request system software in-service-upgrade /var/tmp/jinstall-12.3R2-domestic-signed.tgz
[Apr 29 01:31:11]:ISSU: Validating Image
[Apr 29 01:43:13]:ISSU: Validating Image Done
[Apr 29 01:43:13]:ISSU: Preparing Backup RE
[Apr 29 01:43:13]:ISSU: Pushing /var/tmp/jinstall-12.3R2-domestic-signed.tgz to re1:/var/tmp/
jinstall-12.3R2-domestic-signed.tgz
[Apr 29 01:44:48]:ISSU: Pushing package /var/tmp/jinstall-12.3R2-domestic-signed.tgz to re1 done
[Apr 29 01:44:48]:ISSU: Installing package /var/tmp/jinstall-12.3R2-domestic-signed.tgz on re1
[Apr 29 01:52:35]:ISSU: Installing package /var/tmp/jinstall-12.3R2-domestic-signed.tgz on re1
done
[Apr 29 01:52:35]:ISSU: Rebooting Backup RE
[Apr 29 01:52:36]:ISSU: Backup RE Prepare Done
[Apr 29 01:52:36]:ISSU: Waiting for Backup RE reboot
[Apr 29 01:56:45]:ISSU: Backup RE reboot done. Backup RE is up
[Apr 29 01:56:45]:ISSU: Waiting for Backup RE state synchronization
[Apr 29 01:57:10]:ISSU: Backup RE state synchronization done
[Apr 29 01:57:10]:ISSU: GRES operational
[Apr 29 01:58:16]:ISSU: Preparing Daemons
[Apr 29 01:58:40]:ISSU: Daemons Ready for ISSU
[Apr 29 01:58:46]:ISSU: Offline Incompatible FRUs
[Apr 29 01:58:51]:ISSU: Starting Upgrade for FRUs
[Apr 29 02:03:32]:ISSU: Preparing for Switchover
[Apr 29 02:03:57]:ISSU: Ready for Switchover
[Apr 29 02:03:59]:ISSU: RE switchover Done
[Apr 29 02:03:59]:ISSU: Upgrading Old Master RE
[Apr 29 02:12:51]:ISSU: Old Master Upgrade Done
[Apr 29 02:12:51]:ISSU: IDLE
```

request system software in-service-upgrade enhanced-mode

```
{master}
user@host> request system software in-service-upgrade /var/tmp/junos-install-mx-x86-32-20.1.tgz reboot
enhanced-mode
Chassis ISSU enhanced-mode
ISSU: set chassis enhanced-mode
Chassis ISSU Check Done
ISSU: Validating Image
..
mgd: commit complete
Validation succeeded
 Validating Image Done
 Preparing Backup RE
 Pushing /var/tmp/junos-install-mx-x86-32-20.1.tgz to re1:/var/tmp/junos-install-mx-
x86-32-20.1.tgz
 Pushing package /var/tmp/junos-install-mx-x86-32-20.1.tgz to re1 done
 Installing package /var/tmp/junos-install-mx-x86-32-20.1.tgz on re1
…
Verified sflow-mx signed by PackageDevelopmentEc_2019 method ECDSA256+SHA256
NOTICE: 'pending' set will be activated at next reboot...
ISSU: Installing package /var/tmp/junos-install-mx-x86-32-20.1.tgz on re1 done
ISSU: Rebooting Backup RE

Rebooting re1
 Backup RE Prepare Done
 Waiting for Backup RE reboot
 Backup RE reboot done. Backup RE is up
 Waiting for Backup RE state synchronization
 Backup RE state synchronization done
 GRES operational
 "Initiating Chassis In-Service-Upgrade"
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Offline Incompatible FRUs
ISSU: Starting Upgrade for FRUs
…

ISSU: Preparing for Switchover
ISSU: Ready for Switchover
 Checking In-Service-Upgrade status
```

```
  Item           Status                  Reason
  FPC 1          Online (ISSU)
  FPC 2          Offline                 Configured power off
Resolving mastership...
Complete. The other routing engine becomes the master.
```

Release Information

Command introduced in Junos OS Release 9.0.

`status` option introduced in Junos OS Release 19.4R1 for MX240, MX480, MX960, MX2010, MX2020, PTX3000, and PTX5000 routers.

`handle-incompatible-config` option introduced in Junos OS Release 22.1R1 for MX960, MX10003, MX10008, MX10016, MX2010, and MX2020 routers.

RELATED DOCUMENTATION

- request system software abort | **15212**
- show chassis in-service-upgrade | **18507**
- *Getting Started with Unified In-Service Software Upgrade*
- *Performing an In-Service Software Upgrade (ISSU) with Non-Stop Routing*
- *Example: Performing a Unified ISSU*

request system software in-service-upgrade (Maintenance)

IN THIS SECTION

- Syntax | **15283**
- Description | **15283**
- Options | **15283**
- Required Privilege Level | **15284**
- Output Fields | **15284**
- Sample Output | **15284**

Sample Output | 15285
Release Information | 15287

Syntax

```
request system software in-service-upgrade image_name
<no-copy>
<no-old-master-upgrade>
<no-sync>
<no-tcp-syn-check>
<no-validate>
<status>
<unlink>
```

Description

The in-service software upgrade (ISSU) feature allows a chassis cluster pair to be upgraded from supported Junos OS versions with a traffic impact similar to that of redundancy group failovers. Before upgrading, you must perform failovers so that all redundancy groups are active on only one device. We recommend that graceful restart for routing protocols be enabled before you initiate an ISSU.

For SRX300, SRX320, SRX340, SRX345, and SRX380 devices, you must use the `no-sync` parameter to perform an in-band cluster upgrade (ICU). This allows a chassis cluster pair to be upgraded with a minimal service disruption of approximately 30 seconds.

For SRX1500, SRX4100, and SRX4200 devices, the `no-sync` parameter is not supported when using ISSU to upgrade. The `no-sync` option specifies that the state is not synchronized from the primary node to the secondary node.

For SRX1500 devices, the `no-tcp-syn-check` parameter is not supported when using ISSU to upgrade.

Options

- *image_name*—Specify the location and name of the software upgrade package to be installed.
- `no-copy`—(Optional) Install the software upgrade package but do not save the copies of package files.
- `no-old-master-upgrade`—(Optional) Do not upgrade the old primary after switchover.

This parameter applies to SRX5400, SRX5600, and SRX5800 devices only.

- `no-sync`—(Optional) Stop the flow state from synchronizing when the old secondary node has booted with a new Junos OS image.

  This parameter applies to SRX300, SRX320, SRX340, SRX345, and SRX380 devices only. It is required for an ICU.

- `no-tcp-syn-check`—(Optional) Create a window wherein the TCP SYN check for the incoming packets is disabled. The default value for the window is 7200 seconds (2 hours).

  This parameter applies to SRX300, SRX320, SRX340, SRX345, and SRX380 devices only.

- `no-validate`—(Optional) Disable the configuration validation step at installation. The system behavior is similar to that of the `request system software add` command.

  This parameter applies to SRX300, SRX320, SRX340, SRX345, and SRX380 devices only.

- `status`—(Optional) Display the status of a unified ISSU during the upgrade. You will need to run this command on the Routing Engine where the ISSU was triggered to display the correct ISSU log file.

  This parameter applies to SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800 devices only.

- `unlink`—(Optional) Remove the software package after successful installation.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software in-service-upgrade status (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800)

```
user@host> request system software in-service-upgrade status
[Apr 29 01:31:11]:ISSU: Validating Image
[Apr 29 01:43:13]:ISSU: Validating Image Done
[Apr 29 01:43:13]:ISSU: Preparing Backup RE
[Apr 29 01:43:13]:ISSU: Pushing /var/tmp/junos-install-mx-x86-32-19.3I20190425_1
```

```
100_divyansh.tgz to re1:/var/tmp/junos-install-mx-x86-32-19.3I20190425_1100_divy
ansh.tgz
[Apr 29 01:44:48]:ISSU: Pushing package /var/tmp/junos-install-mx-x86-32-19.3I20
190425_1100_divyansh.tgz to re1 done
[Apr 29 01:44:48]:ISSU: Installing package /var/tmp/junos-install-mx-x86-32-19.3
I20190425_1100_divyansh.tgz on re1
[Apr 29 01:52:35]:ISSU: Installing package /var/tmp/junos-install-mx-x86-32-19.3
I20190425_1100_divyansh.tgz on re1 done
[Apr 29 01:52:35]:ISSU: Rebooting Backup RE
[Apr 29 01:52:36]:ISSU: Backup RE Prepare Done
[Apr 29 01:52:36]:ISSU: Waiting for Backup RE reboot
[Apr 29 01:56:45]:ISSU: Backup RE reboot done. Backup RE is up
[Apr 29 01:56:45]:ISSU: Waiting for Backup RE state synchronization
[Apr 29 01:57:10]:ISSU: Backup RE state synchronization done
[Apr 29 01:57:10]:ISSU: GRES operational
[Apr 29 01:58:16]:ISSU: Preparing Daemons
[Apr 29 01:58:40]:ISSU: Daemons Ready for ISSU
[Apr 29 01:58:46]:ISSU: Offline Incompatible FRUs
[Apr 29 01:58:51]:ISSU: Starting Upgrade for FRUs
[Apr 29 02:03:32]:ISSU: Preparing for Switchover
[Apr 29 02:03:57]:ISSU: Ready for Switchover
[Apr 29 02:03:59]:ISSU: RE switchover Done
[Apr 29 02:03:59]:ISSU: Upgrading Old Master RE
[Apr 29 02:12:51]:ISSU: Old Master Upgrade Done
[Apr 29 02:12:51]:ISSU: IDLE
```

Sample Output

request system software in-service-upgrade (SRX300, SRX320, SRX340, SRX345, and SRX380 Devices)

```
user@host> request system software in-service-upgrade /var/tmp/junos-srxsme-15.1I20160520_0757-
domestic.tgz no-sync

ISSU: Validating package
WARNING: in-service-upgrade shall reboot both the nodes
         in your cluster. Please ignore any subsequent
         reboot request message
ISSU: start downloading software package on secondary node
Pushing /var/tmp/junos-srxsme-15.1I20160520_0757-domestic.tgz to node0:/var/tmp/junos-
srxsme-15.1I20160520_0757-domestic.tgz
```

```
Formatting alternate root (/dev/da0s1a)...
/dev/da0s1a: 2510.1MB (5140780 sectors) block size 16384, fragment size 2048
        using 14 cylinder groups of 183.62MB, 11752 blks, 23552 inodes.
super-block backups (for fsck -b #) at:
32, 376096, 752160, 1128224, 1504288, 1880352, 2256416, 2632480, 3008544,
3384608, 3760672, 4136736, 4512800, 4888864
Installing package '/altroot/cf/packages/install-tmp/junos-15.1I20160520_0757-domestic' ...
Verified junos-boot-srxsme-15.1I20160520_0757.tgz signed by PackageDevelopmentEc_2016 method
ECDSA256+SHA256
Verified junos-srxsme-15.1I20160520_0757-domestic signed by PackageDevelopmentEc_2016 method
ECDSA256+SHA256

WARNING:     The software that is being installed has limited support.
WARNING:     Run 'file show /etc/notices/unsupported.txt' for details.

Verified junos-boot-srxsme-15.1I20160520_0757.tgz signed by PackageDevelopmentEc_2016 method
ECDSA256+SHA256
Verified junos-srxsme-15.1I20160520_0757-domestic signed by PackageDevelopmentEc_2016 method
ECDSA256+SHA256
JUNOS 15.1I20160520_0757 will become active at next reboot
WARNING: A reboot is required to load this software correctly
WARNING:     Use the 'request system reboot' command
WARNING:         when software installation is complete
cp: cannot overwrite directory /altroot/cf/etc/ssh with non-directory /cf/etc/ssh
Saving state for rollback ...
ISSU: finished upgrading on secondary node node0
ISSU: start upgrading software package on primary node
Formatting alternate root (/dev/da0s1a)...
/dev/da0s1a: 2510.1MB (5140780 sectors) block size 16384, fragment size 2048
        using 14 cylinder groups of 183.62MB, 11752 blks, 23552 inodes.
super-block backups (for fsck -b #) at:
32, 376096, 752160, 1128224, 1504288, 1880352, 2256416, 2632480, 3008544,
3384608, 3760672, 4136736, 4512800, 4888864
Installing package '/altroot/cf/packages/install-tmp/junos-15.1I20160520_0757-domestic' ...
Verified junos-boot-srxsme-15.1I20160520_0757.tgz signed by PackageDevelopmentEc_2016 method
ECDSA256+SHA256
Verified junos-srxsme-15.1I20160520_0757-domestic signed by PackageDevelopmentEc_2016 method
ECDSA256+SHA256

WARNING:     The software that is being installed has limited support.
WARNING:     Run 'file show /etc/notices/unsupported.txt' for details.

Verified junos-boot-srxsme-15.1I20160520_0757.tgz signed by PackageDevelopmentEc_2016 method
```

```
ECDSA256+SHA256
Verified junos-srxsme-15.1I20160520_0757-domestic signed by PackageDevelopmentEc_2016 method
ECDSA256+SHA256
JUNOS 15.1I20160520_0757 will become active at next reboot
WARNING: A reboot is required to load this software correctly
WARNING:     Use the 'request system reboot' command
WARNING:         when software installation is complete
cp: cannot overwrite directory /altroot/cf/etc/ssh with non-directory /cf/etc/ssh
Saving state for rollback ...
ISSU: failover all redundancy-groups 1...n to primary node
Successfully reset all redundancy-groups priority back to configured priority.
Successfully reset all redundancy-groups priority back to configured priority.
error: Command failed. None of the redundancy-groups has been failed over.
        Some redundancy-groups' priority on node1 are 0.
        e.g.: priority of redundancy-groups-1 on node1 is 0.
Use 'force' option at the end to ignore this check.
WARNING: Using force option may cause traffic loss.
ISSU: rebooting Secondary Node
Shutdown NOW!
ISSU: Waiting for secondary node node0 to reboot.
ISSU: node 0 went down
ISSU: Waiting for node 0 to come up
ISSU: node 0 came up
ISSU: secondary node node0 booted up.
ISSU: failover all redundancy-groups 1...n to remote node, before reboot.
Successfully reset all redundancy-groups priority back to configured priority.

Shutdown NOW!

{primary:node1}
user@host>
*** FINAL System shutdown message from user@host ***

System going down IMMEDIATELY
```

Release Information

For SRX5400, SRX5600, and SRX5800 devices, command introduced in Junos OS Release 9.6. For SRX5400 devices, the command is introduced in Junos OS Release 12.1X46-D20. For SRX300, SRX320, SRX340, and SRX345 devices, command introduced in Junos OS Release 15.1X49-D40. For SRX1500

devices, command introduced in Junos OS Release 15.1X49-D50. For SRX380 devices, command introduced in Junos OS Release 20.1R1.

Starting with Junos OS Release 15.1X49-D80, SRX4100 and SRX4200 devices support ISSU.

Starting with Junos OS Release 17.4R1, SRX4600 devices support ISSU.

SRX300 Series devices and vSRX Virtual Firewall do not support ISSU.

For SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800, the `status` option is introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

request system software rollback (SRX Series)

request system software in-service-upgrade (MX Series 5G Universal Routing Platforms and EX9200 Switches)

IN THIS SECTION

- Syntax | 15288
- Description | 15289
- Options | 15289
- Additional Information | 15290
- Required Privilege Level | 15290
- Output Fields | 15290
- Sample Output | 15291
- Release Information | 15311

Syntax

```
request system software in-service-upgrade package-name
<no-copy>
<no-old-master-upgrade>
```

```
<reboot>
<unlink>
```

Description

Perform a unified in-service software upgrade (unified ISSU). Unified ISSU enables you to upgrade from one Junos OS release to another with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is supported only by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled.

Options

package-name

Location from which the software package or bundle is to be installed. For example:

- **/var/tmp/*package-name***— For a software package or bundle that is being installed from a local directory on the router.
- ***protocol://hostname/pathname/package-name***—For a software package or bundle that is to be downloaded and installed from a remote location. Replace ***protocol*** with one of the following:
  - **ftp**—File Transfer Protocol
  - **http**—Hypertext Transfer Protocol
  - **scp**—Secure copy (available only for Canada and U.S. version)

no-copy

(Optional) When the `no-copy` option is included, copies of package files are not saved on the Packet Forwarding Engine.

The `no-copy` option is not available for an MX Series Virtual Chassis or an EX9200 Virtual Chassis.

no-old-master-upgrade

(Optional) When the `no-old-master-upgrade` option is included, after the backup Routing Engine is rebooted with the new software package and a switchover occurs to make it the new primary Routing Engine, the former primary (new backup) Routing Engine is not upgraded to the new software. In this case, you must manually upgrade the former primary (new backup) Routing Engine. If you do not include the `no-old-master-upgrade` option, the system automatically upgrades the former primary Routing Engine.

The `no-old-master-upgrade` option is not available for an MX Series Virtual Chassis or an EX9200 Virtual Chassis.

reboot (Optional) When the `reboot` option is included, the former primary (new backup) Routing Engine is automatically rebooted after being upgraded to the new software. When the `reboot` option is not included, you must manually reboot the former primary (new backup) Routing Engine using the `request system reboot` command.

The `reboot` option is accepted but ignored for an MX Series Virtual Chassis or an EX9200 Virtual Chassis. A unified ISSU in an MX Series Virtual Chassis or EX9200 Virtual Chassis always reboots all Routing Engines in the member routers or switches.

unlink (Optional) When the `unlink` option is included, the package is removed from **/var/home** whether the installation is successful or unsuccessful.

The `unlink` option is not available for an MX Series Virtual Chassis or an EX9200 Virtual Chassis.

Additional Information

The following conditions apply to unified ISSUs:

- Unified ISSUs are supported on MX Series 5G Universal Routing Platforms and EX9200 switches.
- Unsupported PICs (on EX9200, PICs are known as “line cards”) are restarted during a unified ISSU. For information about supported PICs, see the Junos OS High Availability User Guide. For information about supported EX9200 line cards, see *Unified ISSU System Requirements*.
- Unsupported protocols will experience packet loss during a unified ISSU. For information about supported protocols, see the Junos OS High Availability User Guide or, for EX9200, see *Unified ISSU System Requirements*.
- During a unified ISSU, you cannot bring any PICs online or offline.

For more information, see the Junos OS High Availability User Guide.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback about the status of your request.

Sample Output

request system software in-service-upgrade reboot

```
{master}
user@host> request system software in-service-upgrade /var/tmp/jinstall-11.2B2.1-domestic-
signed.tgz reboot
Chassis ISSU Check Done
ISSU: Validating Image
Checking compatibility with configuration
Initializing...
Using jbase-11.2B1.5
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B1.5 signed by PackageProduction_11_2_0
Using /var/tmp/jinstall-11.2B2.1-domestic-signed.tgz
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Using jinstall-11.2B2.1-domestic.tgz
Using jbundle-11.2B2.1-domestic.tgz
Checking jbundle requirements on /
Using jbase-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B2.1 signed by PackageProduction_11_2_0
Using /var/validate/chroot/tmp/jbundle/jboot-11.2B2.1.tgz
Using jcrypto-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jcrypto-11.2B2.1 signed by PackageProduction_11_2_0
Using jdocs-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jdocs-11.2B2.1 signed by PackageProduction_11_2_0
Using jkernel-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jkernel-11.2B2.1 signed by PackageProduction_11_2_0
Using jpfe-11.2B2.1.tgz
Using jroute-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jroute-11.2B2.1 signed by PackageProduction_11_2_0
Using jruntime-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jruntime-11.2B2.1 signed by PackageProduction_11_2_0
Using jservices-11.2B2.1.tgz
Auto-deleting old jservices-voice ...
Removing /opt/sdk/service-packages/jservices-voice ...
```

```
Removing jservices-voice-bsg-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /var/sw/pkg ...
Creating /opt/sdk/service-packages/jservices-voice ...
Storing jservices-voice-bsg-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-voice/jservices-voice-bsg -> /var/sw/pkg/jservices-
voice-bsg-11.2B2.1.tgz...
Auto-deleting old jservices-bgf ...
Removing /opt/sdk/service-packages/jservices-bgf ...
Removing jservices-bgf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-bgf ...
Storing jservices-bgf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-bgf/jservices-bgf-pic -> /var/sw/pkg/jservices-bgf-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-aacl ...
Removing /opt/sdk/service-packages/jservices-aacl ...
Removing jservices-aacl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aacl ...
Verified jservices-aacl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-aacl ...
Storing jservices-aacl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-aacl/jservices-aacl-pic -> /var/sw/pkg/jservices-aacl-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/service-packages/jservices-llpdf ...
Removing jservices-llpdf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-llpdf ...
Storing jservices-llpdf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-llpdf/jservices-llpdf-pic -> /var/sw/pkg/jservices-
llpdf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ptsp ...
Removing /opt/sdk/service-packages/jservices-ptsp ...
Removing jservices-ptsp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
```

```
Installing new jservices-ptsp ...
Verified jservices-ptsp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ptsp ...
Storing jservices-ptsp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ptsp/jservices-ptsp-pic -> /var/sw/pkg/jservices-ptsp-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/service-packages/jservices-sfw ...
Removing jservices-sfw-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-sfw ...
Storing jservices-sfw-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-sfw/jservices-sfw-pic -> /var/sw/pkg/jservices-sfw-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-nat ...
Removing /opt/sdk/service-packages/jservices-nat ...
Removing jservices-nat-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-nat ...
Verified jservices-nat-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-nat ...
Storing jservices-nat-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-nat/jservices-nat-pic -> /var/sw/pkg/jservices-nat-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-alg ...
Removing /opt/sdk/service-packages/jservices-alg ...
Removing jservices-alg-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-alg ...
Verified jservices-alg-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-alg ...
Storing jservices-alg-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-alg/jservices-alg-pic -> /var/sw/pkg/jservices-alg-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-cpcd ...
Removing /opt/sdk/service-packages/jservices-cpcd ...
Removing jservices-cpcd-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-cpcd ...
Verified jservices-cpcd-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-cpcd ...
```

```
Storing jservices-cpcd-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-cpcd/jservices-cpcd-pic -> /var/sw/pkg/jservices-cpcd-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-rpm ...
Removing /opt/sdk/service-packages/jservices-rpm ...
Removing jservices-rpm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-rpm ...
Verified jservices-rpm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-rpm ...
Storing jservices-rpm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-rpm/jservices-rpm-pic -> /var/sw/pkg/jservices-rpm-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-hcm ...
Removing /opt/sdk/service-packages/jservices-hcm ...
Removing jservices-hcm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-hcm ...
Verified jservices-hcm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-hcm ...
Storing jservices-hcm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-hcm/jservices-hcm-pic -> /var/sw/pkg/jservices-hcm-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-appid ...
Removing /opt/sdk/service-packages/jservices-appid ...
Removing jservices-appid-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-appid ...
Storing jservices-appid-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-appid/jservices-appid-pic -> /var/sw/pkg/jservices-
appid-pic-11.2B2.1.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/service-packages/jservices-idp ...
Removing jservices-idp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-idp ...
Storing jservices-idp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-idp/jservices-idp-pic -> /var/sw/pkg/jservices-idp-
pic-11.2B2.1.tgz...
```

```
Using jservices-crypto-11.2B2.1.tgz
Auto-deleting old jservices-crypto-base ...
Removing /opt/sdk/service-packages/jservices-crypto-base ...
Removing jservices-crypto-base-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-crypto-base ...
Verified jservices-crypto-base-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-crypto-base ...
Storing jservices-crypto-base-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-crypto-base/jservices-crypto-base-pic -> /var/sw/pkg/
jservices-crypto-base-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ssl ...
Removing /opt/sdk/service-packages/jservices-ssl ...
Removing jservices-ssl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ssl ...
Verified jservices-ssl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ssl ...
Storing jservices-ssl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ssl/jservices-ssl-pic -> /var/sw/pkg/jservices-ssl-
pic-11.2B2.1.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
ISSU: Preparing Backup RE
Pushing bundle to re1
NOTICE: Validating configuration against jinstall-11.2B2.1-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Checking compatibility with configuration
Initializing...
Using jbase-11.2B1.5
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B1.5 signed by PackageProduction_11_2_0
Using /var/tmp/jinstall-11.2B2.1-domestic-signed.tgz
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Using jinstall-11.2B2.1-domestic.tgz
Using jbundle-11.2B2.1-domestic.tgz
Checking jbundle requirements on /
Using jbase-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B2.1 signed by PackageProduction_11_2_0
Using /var/validate/chroot/tmp/jbundle/jboot-11.2B2.1.tgz
```

```
Using jcrypto-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jcrypto-11.2B2.1 signed by PackageProduction_11_2_0
Using jdocs-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jdocs-11.2B2.1 signed by PackageProduction_11_2_0
Using jkernel-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jkernel-11.2B2.1 signed by PackageProduction_11_2_0
Using jpfe-11.2B2.1.tgz
Using jroute-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jroute-11.2B2.1 signed by PackageProduction_11_2_0
Using jruntime-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jruntime-11.2B2.1 signed by PackageProduction_11_2_0
Using jservices-11.2B2.1.tgz
Auto-deleting old jservices-voice ...
Removing /opt/sdk/service-packages/jservices-voice ...
Removing jservices-voice-bsg-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /var/sw/pkg ...
Creating /opt/sdk/service-packages/jservices-voice ...
Storing jservices-voice-bsg-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-voice/jservices-voice-bsg -> /var/sw/pkg/jservices-
voice-bsg-11.2B2.1.tgz...
Auto-deleting old jservices-bgf ...
Removing /opt/sdk/service-packages/jservices-bgf ...
Removing jservices-bgf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-bgf ...
Storing jservices-bgf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-bgf/jservices-bgf-pic -> /var/sw/pkg/jservices-bgf-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-aacl ...
Removing /opt/sdk/service-packages/jservices-aacl ...
Removing jservices-aacl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aacl ...
```

```
Verified jservices-aacl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-aacl ...
Storing jservices-aacl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-aacl/jservices-aacl-pic -> /var/sw/pkg/jservices-aacl-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/service-packages/jservices-llpdf ...
Removing jservices-llpdf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-llpdf ...
Storing jservices-llpdf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-llpdf/jservices-llpdf-pic -> /var/sw/pkg/jservices-
llpdf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ptsp ...
Removing /opt/sdk/service-packages/jservices-ptsp ...
Removing jservices-ptsp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ptsp ...
Verified jservices-ptsp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ptsp ...
Storing jservices-ptsp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ptsp/jservices-ptsp-pic -> /var/sw/pkg/jservices-ptsp-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/service-packages/jservices-sfw ...
Removing jservices-sfw-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-sfw ...
Storing jservices-sfw-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-sfw/jservices-sfw-pic -> /var/sw/pkg/jservices-sfw-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-nat ...
Removing /opt/sdk/service-packages/jservices-nat ...
Removing jservices-nat-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-nat ...
Verified jservices-nat-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-nat ...
Storing jservices-nat-pic-11.2B2.1.tgz in /var/sw/pkg ...
```

```
Link: /opt/sdk/service-packages/jservices-nat/jservices-nat-pic -> /var/sw/pkg/jservices-nat-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-alg ...
Removing /opt/sdk/service-packages/jservices-alg ...
Removing jservices-alg-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-alg ...
Verified jservices-alg-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-alg ...
Storing jservices-alg-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-alg/jservices-alg-pic -> /var/sw/pkg/jservices-alg-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-cpcd ...
Removing /opt/sdk/service-packages/jservices-cpcd ...
Removing jservices-cpcd-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-cpcd ...
Verified jservices-cpcd-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-cpcd ...
Storing jservices-cpcd-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-cpcd/jservices-cpcd-pic -> /var/sw/pkg/jservices-cpcd-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-rpm ...
Removing /opt/sdk/service-packages/jservices-rpm ...
Removing jservices-rpm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-rpm ...
Verified jservices-rpm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-rpm ...
Storing jservices-rpm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-rpm/jservices-rpm-pic -> /var/sw/pkg/jservices-rpm-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-hcm ...
Removing /opt/sdk/service-packages/jservices-hcm ...
Removing jservices-hcm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-hcm ...
Verified jservices-hcm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-hcm ...
Storing jservices-hcm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-hcm/jservices-hcm-pic -> /var/sw/pkg/jservices-hcm-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-appid ...
```

```
Removing /opt/sdk/service-packages/jservices-appid ...
Removing jservices-appid-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-appid ...
Storing jservices-appid-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-appid/jservices-appid-pic -> /var/sw/pkg/jservices-
appid-pic-11.2B2.1.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/service-packages/jservices-idp ...
Removing jservices-idp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-idp ...
Storing jservices-idp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-idp/jservices-idp-pic -> /var/sw/pkg/jservices-idp-
pic-11.2B2.1.tgz...
Using jservices-crypto-11.2B2.1.tgz
Auto-deleting old jservices-crypto-base ...
Removing /opt/sdk/service-packages/jservices-crypto-base ...
Removing jservices-crypto-base-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-crypto-base ...
Verified jservices-crypto-base-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-crypto-base ...
Storing jservices-crypto-base-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-crypto-base/jservices-crypto-base-pic -> /var/sw/pkg/
jservices-crypto-base-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ssl ...
Removing /opt/sdk/service-packages/jservices-ssl ...
Removing jservices-ssl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ssl ...
Verified jservices-ssl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ssl ...
Storing jservices-ssl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ssl/jservices-ssl-pic -> /var/sw/pkg/jservices-ssl-
pic-11.2B2.1.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
```

```
Validation succeeded
Installing package '/var/tmp/jinstall-11.2B2.1-domestic-signed.tgz' ...
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Adding jinstall...
Verified manifest signed by PackageProduction_11_2_0

WARNING:     This package will load JUNOS 11.2B2.1 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-11.2B2.1-domestic-signed.tgz ...
Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 1          Online (ISSU)
```

```
  FPC 4          Online (ISSU)
  FPC 8          Online (ISSU)
  FPC 10         Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
NOTICE: Validating configuration against jinstall-11.2B2.1-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Checking compatibility with configuration
Initializing...
Using jbase-11.2B1.5
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B1.5 signed by PackageProduction_11_2_0
Using /var/tmp/jinstall-11.2B2.1-domestic-signed.tgz
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Using jinstall-11.2B2.1-domestic.tgz
Using jbundle-11.2B2.1-domestic.tgz
Checking jbundle requirements on /
Using jbase-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B2.1 signed by PackageProduction_11_2_0
Using /var/validate/chroot/tmp/jbundle/jboot-11.2B2.1.tgz
Using jcrypto-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jcrypto-11.2B2.1 signed by PackageProduction_11_2_0
Using jdocs-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jdocs-11.2B2.1 signed by PackageProduction_11_2_0
Using jkernel-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jkernel-11.2B2.1 signed by PackageProduction_11_2_0
Using jpfe-11.2B2.1.tgz
Using jroute-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jroute-11.2B2.1 signed by PackageProduction_11_2_0
Using jruntime-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jruntime-11.2B2.1 signed by PackageProduction_11_2_0
Using jservices-11.2B2.1.tgz
Auto-deleting old jservices-voice ...
Removing /opt/sdk/service-packages/jservices-voice ...
Removing jservices-voice-bsg-11.2B1.5.tgz from /var/sw/pkg ...
```

```
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /var/sw/pkg ...
Creating /opt/sdk/service-packages/jservices-voice ...
Storing jservices-voice-bsg-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-voice/jservices-voice-bsg -> /var/sw/pkg/jservices-
voice-bsg-11.2B2.1.tgz...
Auto-deleting old jservices-bgf ...
Removing /opt/sdk/service-packages/jservices-bgf ...
Removing jservices-bgf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-bgf ...
Storing jservices-bgf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-bgf/jservices-bgf-pic -> /var/sw/pkg/jservices-bgf-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-aacl ...
Removing /opt/sdk/service-packages/jservices-aacl ...
Removing jservices-aacl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aacl ...
Verified jservices-aacl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-aacl ...
Storing jservices-aacl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-aacl/jservices-aacl-pic -> /var/sw/pkg/jservices-aacl-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/service-packages/jservices-llpdf ...
Removing jservices-llpdf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-llpdf ...
Storing jservices-llpdf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-llpdf/jservices-llpdf-pic -> /var/sw/pkg/jservices-
llpdf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ptsp ...
Removing /opt/sdk/service-packages/jservices-ptsp ...
Removing jservices-ptsp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ptsp ...
```

```
Verified jservices-ptsp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ptsp ...
Storing jservices-ptsp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ptsp/jservices-ptsp-pic -> /var/sw/pkg/jservices-ptsp-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/service-packages/jservices-sfw ...
Removing jservices-sfw-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-sfw ...
Storing jservices-sfw-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-sfw/jservices-sfw-pic -> /var/sw/pkg/jservices-sfw-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-nat ...
Removing /opt/sdk/service-packages/jservices-nat ...
Removing jservices-nat-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-nat ...
Verified jservices-nat-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-nat ...
Storing jservices-nat-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-nat/jservices-nat-pic -> /var/sw/pkg/jservices-nat-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-alg ...
Removing /opt/sdk/service-packages/jservices-alg ...
Removing jservices-alg-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-alg ...
Verified jservices-alg-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-alg ...
Storing jservices-alg-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-alg/jservices-alg-pic -> /var/sw/pkg/jservices-alg-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-cpcd ...
Removing /opt/sdk/service-packages/jservices-cpcd ...
Removing jservices-cpcd-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-cpcd ...
Verified jservices-cpcd-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-cpcd ...
Storing jservices-cpcd-pic-11.2B2.1.tgz in /var/sw/pkg ...
```

```
Link: /opt/sdk/service-packages/jservices-cpcd/jservices-cpcd-pic -> /var/sw/pkg/jservices-cpcd-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-rpm ...
Removing /opt/sdk/service-packages/jservices-rpm ...
Removing jservices-rpm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-rpm ...
Verified jservices-rpm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-rpm ...
Storing jservices-rpm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-rpm/jservices-rpm-pic -> /var/sw/pkg/jservices-rpm-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-hcm ...
Removing /opt/sdk/service-packages/jservices-hcm ...
Removing jservices-hcm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-hcm ...
Verified jservices-hcm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-hcm ...
Storing jservices-hcm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-hcm/jservices-hcm-pic -> /var/sw/pkg/jservices-hcm-
pic-11.2B2.1.tgz...
Auto-deleting old jservices-appid ...
Removing /opt/sdk/service-packages/jservices-appid ...
Removing jservices-appid-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-appid ...
Storing jservices-appid-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-appid/jservices-appid-pic -> /var/sw/pkg/jservices-
appid-pic-11.2B2.1.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/service-packages/jservices-idp ...
Removing jservices-idp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-idp ...
Storing jservices-idp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-idp/jservices-idp-pic -> /var/sw/pkg/jservices-idp-
pic-11.2B2.1.tgz...
Using jservices-crypto-11.2B2.1.tgz
```

```
Auto-deleting old jservices-crypto-base ...
Removing /opt/sdk/service-packages/jservices-crypto-base ...
Removing jservices-crypto-base-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-crypto-base ...
Verified jservices-crypto-base-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-crypto-base ...
Storing jservices-crypto-base-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-crypto-base/jservices-crypto-base-pic -> /var/sw/pkg/
jservices-crypto-base-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ssl ...
Removing /opt/sdk/service-packages/jservices-ssl ...
Removing jservices-ssl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ssl ...
Verified jservices-ssl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ssl ...
Storing jservices-ssl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ssl/jservices-ssl-pic -> /var/sw/pkg/jservices-ssl-
pic-11.2B2.1.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-11.2B2.1-domestic-signed.tgz' ...
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Adding jinstall...
Verified manifest signed by PackageProduction_11_2_0

WARNING:     This package will load JUNOS 11.2B2.1 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
```

```
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-11.2B2.1-domestic-signed.tgz ...
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
Reboot consistency check bypassed - jinstall 11.2B2.1 will complete installation upon reboot
[pid 66780]

*** FINAL System shutdown message from user@host> ***
System going down IMMEDIATELY
```

request system software in-service-upgrade (MX Series Virtual Chassis)

```
{master:member0-re0}
user@host> request system software in-service-upgrade jinstall-14.1-20140114.2-domestic-
signed.tgz
[Jan 30 10:45:32]:ISSU: IDLE

Beginning in-service-upgrade at Jan 30, 2014; 10:45:34
[Jan 30 10:45:34]:ISSU: Validating Image
Validating VC readiness...
Validating required configuration...
Validating release compatibility...
Validation successful
Initiating chassis in-service-upgrade
[Jan 30 10:46:56]:ISSU: Preparing LCC Backup REs
Copying new release to all RE's
Pushing bundle to member0-re0
Pushing bundle to member1-re0
Pushing bundle to member1-re1
[Jan 30 10:51:11]:ISSU: Preparing Backup RE
Arming new release on all RE's
member0-re0:
--------------------------------------------------------------------------
Installing package '/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by PackageDevelopmentEc_2014
Adding jinstall...
```

```
WARNING:     The software that is being installed has limited support.
WARNING:     Run 'file show /etc/notices/unsupported.txt' for details.

veriexec: accepting signer: PackageDevelopmentEc_2014
Verified manifest signed by PackageDevelopmentEc_2014

WARNING:     This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...

member1-re0:
--------------------------------------------------------------------------
Installing package '/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by PackageDevelopmentEc_2014
Adding jinstall...

WARNING:     The software that is being installed has limited support.
WARNING:     Run 'file show /etc/notices/unsupported.txt' for details.

veriexec: accepting signer: PackageDevelopmentEc_2014
Verified manifest signed by PackageDevelopmentEc_2014

WARNING:     This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
```

```
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...

member1-re1:
--------------------------------------------------------------------------
Installing package '/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by PackageDevelopmentEc_2014
Adding jinstall...

WARNING:     The software that is being installed has limited support.
WARNING:     Run 'file show /etc/notices/unsupported.txt' for details.

veriexec: accepting signer: PackageDevelopmentEc_2014
Verified manifest signed by PackageDevelopmentEc_2014

WARNING:     This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
```

```
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by PackageDevelopmentEc_2014
Adding jinstall...

WARNING:     The software that is being installed has limited support.
WARNING:     Run 'file show /etc/notices/unsupported.txt' for details.

veriexec: accepting signer: PackageDevelopmentEc_2014
Verified manifest signed by PackageDevelopmentEc_2014

WARNING:     This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...
[Jan 30 11:03:12]:ISSU: Backup RE Prepare Done
Rebooting standby RE's
Sending Reboot Command to member0-re0
Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will complete
installation upon reboot
```

```
[pid 2757]
Sending Reboot Command to member1-re1
Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will complete
installation upon reboot
[pid 2670]
Waiting for standby RE's to boot
[Jan 30 11:18:26]:ISSU: LCC Backup REs Prepare Done
Waiting for standby RE's to have the correct ISSU state
Waiting for protocol backup to be ready to switch mastership
Switching mastership on the protocol backup chassis to slot 1
Waiting for protocol backup chassis master switch to complete
Globally updating ISSU state
Waiting for protocol backup chassis to become GRES ready
[Jan 30 11:19:18]:ISSU: VC Protocol Backup has Switched
Passing ISSU control to chassisd
Chassis ISSU Started
[Jan 30 11:21:01]:ISSU: Preparing Daemons
[Jan 30 11:22:02]:ISSU: Daemons Ready for ISSU
[Jan 30 11:22:06]:ISSU: Starting Upgrade for FRUs
[Jan 30 11:25:42]:ISSU: Preparing for Switchover
[Jan 30 11:26:06]:ISSU: Ready for Switchover
[Jan 30 11:26:20]:ISSU: All VC Members Ready for Switchover
Waiting for master chassis to be switch ready
Switching mastership locally
Resolving mastership...
Complete. The other routing engine becomes the master.
Waiting for virtual chassis roles to switch
Globally updating ISSU state to IDLE
[Jan 30 11:26:33]:ISSU: IDLE
Rebooting protocol backup standby RE.
Sending Reboot Command to member1-re0

member1-re0:
--------------------------------------------------------------------------
Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will complete
installation upon reboot
[pid 10462]
Rebooting locally to complete the in service upgrade.
Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will complete
installation upon reboot
```

```
[pid 13458]

{local:member0-re1}
user@host>
*** FINAL System shutdown message from user@host ***

System going down IMMEDIATELY



Connection closed by foreign host.
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

request system software abort | 15212

show chassis in-service-upgrade | 18507

request system software nonstop-upgrade (QFabric Systems)

IN THIS SECTION

- Syntax | 15312
- Description | 15312
- Options | 15312
- Required Privilege Level | 15313
- Output Fields | 15313
- Sample Output | 15313
- Release Information | 15319

Syntax

```
request system  software nonstop-upgrade package-name
<fabric >
<director-group>
<node-group name>
```

Description

Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. You should upgrade the devices in the following order: Director group, fabric controls and Interconnect devices, and network and server Node groups.

> **NOTE**: Before you perform a nonstop software upgrade, contact JTAC to perform a pre-upgrade health check on the QFabric system.

Options

package-name Location from which the software is to be installed. For example:

- ***protocol://hostname/pathname/package-name***—For a software package or bundle that is to be downloaded and installed from a remote location. Replace ***protocol*** with one of the following:
  - **ftp**—File Transfer Protocol.
    Use **ftp://*hostname/pathname/package-name***. To specify authentication credentials, use **ftp://<*username*>:<*password*>@*hostname/pathname/package-name***. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
  - **http**—Hypertext Transfer Protocol.
    Use **http://*hostname/pathname/package-name***. To specify authentication credentials, use **http://<*username*>:<*password*>@*hostname/pathname/package-name***. If a password is required and you omit it, you are prompted for it.
  - **scp**—Secure copy (available only for Canada and U.S. version).

Use **scp://** ***hostname/ pathname/ package-name***. To specify authentication credentials, use **scp://<*username*>:<*password*>@*hostname/ pathname/ package-name***.

> **NOTE:**
>
> - The ***pathname*** in the protocol is the relative path to the user's home directory on the remote system and not the root directory.

director-group Install software package on the Director group and Fabric managers.

fabric Install software package on the Interconnect devices and Fabric controls.

node-group *name* Install software package on the redundant server Node group, server Node group, or network Node group.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software nonstop-upgrade director-group

```
user@qfabric> request system software nonstop-upgrade director-group jinstall-qfabric-12.2X50-D10.3.rpm
Validating update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing fabric images version 12.2X50-D10.3
Performing cleanup
Package install complete
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm on peer
Triggering Initial Stage of Fabric Manager Upgrade
Updating CCIF default image to 12.2X50-D10.3
Updating FM-0 to Junos version 12.2X50-D10.3
[Status   2012-06-05 15:25:29]: Fabric Manager: Upgrade Initial Stage started
```

```
[FM-0     2012-06-05 15:25:38]: FM-0 Master already running on LOCAL DG
[NW-NG-0  2012-06-05 15:25:45]: NW-NG-0 Master already running on LOCAL DG
[FM-0     2012-06-05 15:26:12]: Retrieving package
[FM-0     2012-06-05 15:27:11]: Pushing bundle to re0
[Status   2012-06-05 15:29:06]: Load completed with 0 errors...
[Status   2012-06-05 15:29:06]: Reboot is required to complete upgrade ...
[Status   2012-06-05 15:29:07]: Trying to Connect to Node: FM-0
[Status   2012-06-05 15:29:13]: Rebooting FM-0
[FM-0     2012-06-05 15:29:13]: Waiting for FM-0 to terminate ...
Starting Peer upgrade

Initiating rolling upgrade of Director peer:  version 12.2X50-D10.3

Inform CCIF regarding rolling upgrade
[Peer Update Status]: Validating install package jinstall-qfabric-12.2X50-D10.3.rpm
[Peer Update Status]: Cleaning up node for rolling phase one upgrade
[Peer Update Status]: Director group upgrade complete
[Peer Update Status]: COMPLETED
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling
```

```
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to complete phase one of rolling upgrade
[Peer Update Status]: Peer completed phase one of rolling upgrade
Setting peer DG node as the master SFC

Delaying start of local upgrade to allow peer services time to initialize [15 minutes]
Delaying start of local upgrade to allow peer services time to initialize [15 minutes]
Delaying start of local upgrade to allow peer services time to initialize [12 minutes]
Delaying start of local upgrade to allow peer services time to initialize [9 minutes]
Delaying start of local upgrade to allow peer services time to initialize [6 minutes]
Delaying start of local upgrade to allow peer services time to initialize [3 minutes]
[Peer Update Status]: Check for VMs on dg0
Triggering Final Stage of Fabric Manager Upgrade:
Updating FM-0 to Junos version 12.2X50-D10.3
[Status   2012-06-05 16:10:12]: Fabric Manager: Upgrade Final Stage started
[NW-NG-0  2012-06-05 16:10:22]: Transferring NW-NG-0 Mastership to REMOTE DG
[NW-NG-0  2012-06-05 16:11:44]: Finished NW-NG-0 Mastership switch
[Status   2012-06-05 16:11:45]: Upgrading FM-0 VM on worker DG to 12.2X50-D10.3
[DRE-0    2012-06-05 16:12:43]: Retrieving package
[DRE-0    2012-06-05 16:13:46]: ------- re0: -------
[Status   2012-06-05 16:15:17]: Load completed with 0 errors...
[Status   2012-06-05 16:15:17]: Reboot is required to complete upgrade ...
[DRE-0    2012-06-05 16:15:22]: Waiting for DRE-0 to terminate ...
[DRE-0    2012-06-05 16:15:34]: Waiting for DRE-0 to come back ...
[DRE-0    2012-06-05 16:18:44]: Running Uptime Test for DRE-0
[DRE-0    2012-06-05 16:18:51]: Uptime Test for DRE-0 Passed ...
[Status   2012-06-05 16:18:51]: DRE-0 booted successfully ...
Performing post install shutdown and cleanup

Broadcast message from root (Tue Jun  5 16:18:51 2012):

The system is going down for reboot NOW!
Director group upgrade complete

root@qfabric> Read from remote host qfabric-partition0: Connection reset by peer
Connection to qfabric-partition0 closed.
```

request system software nonstop-upgrade fabric

```
user@qfabric> request system software nonstop-upgrade fabric jinstall-qfabric-12.2X50-D10.3.rpm
[FC-0     2012-06-05 16:48:53]: Retrieving package
[FC-1     2012-06-05 16:48:53]: Retrieving package
[IC-F4912 2012-06-05 16:48:59]: Retrieving package
[FC-0     2012-06-05 16:49:51]: ------- re0: -------
[FC-1     2012-06-05 16:49:52]: ------- re0: -------
[IC-F4912 2012-06-05 16:49:54]: ------- re0: -------
[IC-F4912 2012-06-05 16:50:42]: Step 1 of 20 Creating temporary file system
[IC-F4912 2012-06-05 16:50:42]: Step 2 of 20 Determining installation source
[IC-F4912 2012-06-05 16:50:43]: Step 3 of 20 Processing format options
[IC-F4912 2012-06-05 16:50:43]: Step 4 of 20 Determining installation slice
[IC-F4912 2012-06-05 16:50:43]: Step 5 of 20 Creating and labeling new slices
[IC-F4912 2012-06-05 16:50:44]: Step 6 of 20 Create and mount new file system
[IC-F4912 2012-06-05 16:50:53]: Step 7 of 20 Getting OS bundles
[IC-F4912 2012-06-05 16:50:53]: Step 8 of 20 Updating recovery media
[IC-F4912 2012-06-05 16:51:17]: Step 9 of 20 Extracting incoming image
[IC-F4912 2012-06-05 16:52:56]: Step 10 of 20 Unpacking OS packages
[IC-F4912 2012-06-05 16:52:59]: Step 11 of 20 Mounting jbase package
[IC-F4912 2012-06-05 16:53:28]: Step 12 of 20 Creating base OS symbolic links
[IC-F4912 2012-06-05 16:54:45]: Step 13 of 20 Creating fstab
[IC-F4912 2012-06-05 16:54:45]: Step 14 of 20 Creating new system files
[IC-F4912 2012-06-05 16:54:46]: Step 15 of 20 Adding jbundle package
[IC-F4912 2012-06-05 16:58:15]: Step 16 of 20 Backing up system data
[IC-F4912 2012-06-05 16:58:18]: Step 17 of 20 Setting up shared partition data
[IC-F4912 2012-06-05 16:58:18]: Step 18 of 20 Checking package sanity in installation
[IC-F4912 2012-06-05 16:58:18]: Step 19 of 20 Unmounting and cleaning up temporary file systems
[IC-F4912 2012-06-05 16:58:22]: Step 20 of 20 Setting da0s1 as new active partition
[Status   2012-06-05 16:58:34]: Load completed with 0 errors...
[Status   2012-06-05 16:58:34]: Reboot is required to complete upgrade ...
[Status   2012-06-05 16:58:34]: Trying to Connect to Node: FC-0
[Status   2012-06-05 16:58:39]: Rebooting FC-0
[Status   2012-06-05 16:58:39]: Trying to Connect to Node: FC-1
[Status   2012-06-05 16:58:44]: Rebooting FC-1
[Status   2012-06-05 16:58:44]: Trying to Connect to Node: IC-F4912
[Status   2012-06-05 16:58:50]: Rebooting IC-F4912
Success
```

request system software nonstop-upgrade node-group (Redundant Server Node Group)

```
user@qfabric> request system software nonstop-upgrade node-group RSNG jinstall-qfabric-12.2X50-
D10.3.rpm
Upgrading target(s): RSNG

[RSNG     2012-06-05 17:26:44]: Starting with package ftp://169.254.0.3/pub/images/12.2X50-D10.3/
jinstall-qfx.tgz
[RSNG     2012-06-05 17:26:44]: Retrieving package
[RSNG     2012-06-05 17:28:56]: Pushing bundle to fpc1
[RSNG     2012-06-05 17:29:26]: fpc1: Validate package...
[RSNG     2012-06-05 17:35:22]: fpc0: Validate package...
[RSNG     2012-06-05 17:35:49]: ------- fpc1 -------
[RSNG     2012-06-05 17:36:25]: Step 1 of 20 Creating temporary file system
[RSNG     2012-06-05 17:36:26]: Step 2 of 20 Determining installation source
[RSNG     2012-06-05 17:36:26]: Step 3 of 20 Processing format options
[RSNG     2012-06-05 17:36:26]: Step 4 of 20 Determining installation slice
[RSNG     2012-06-05 17:36:27]: Step 5 of 20 Creating and labeling new slices
[RSNG     2012-06-05 17:36:27]: Step 6 of 20 Create and mount new file system
[RSNG     2012-06-05 17:36:35]: Step 7 of 20 Getting OS bundles
[RSNG     2012-06-05 17:36:35]: Step 8 of 20 Updating recovery media
[RSNG     2012-06-05 17:36:56]: Step 9 of 20 Extracting incoming image
[RSNG     2012-06-05 17:38:07]: Step 10 of 20 Unpacking OS packages
[RSNG     2012-06-05 17:38:16]: Step 11 of 20 Mounting jbase package
[RSNG     2012-06-05 17:38:41]: Step 12 of 20 Creating base OS symbolic links
[RSNG     2012-06-05 17:39:41]: Step 13 of 20 Creating fstab
[RSNG     2012-06-05 17:39:42]: Step 14 of 20 Creating new system files
[RSNG     2012-06-05 17:39:42]: Step 15 of 20 Adding jbundle package
[RSNG     2012-06-05 17:42:16]: Step 16 of 20 Backing up system data
[RSNG     2012-06-05 17:42:32]: Step 17 of 20 Setting up shared partition data
[RSNG     2012-06-05 17:42:33]: Step 18 of 20 Checking package sanity in installation
[RSNG     2012-06-05 17:42:33]: Step 19 of 20 Unmounting and cleaning up temporary file systems
[RSNG     2012-06-05 17:42:36]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG     2012-06-05 17:42:51]: ------- fpc0 - master -------
[RSNG     2012-06-05 17:42:51]: Step 1 of 20 Creating temporary file system
[RSNG     2012-06-05 17:42:51]: Step 2 of 20 Determining installation source
[RSNG     2012-06-05 17:42:51]: Step 3 of 20 Processing format options
[RSNG     2012-06-05 17:42:51]: Step 4 of 20 Determining installation slice
[RSNG     2012-06-05 17:42:51]: Step 5 of 20 Creating and labeling new slices
[RSNG     2012-06-05 17:42:51]: Step 6 of 20 Create and mount new file system
[RSNG     2012-06-05 17:42:51]: Step 7 of 20 Getting OS bundles
[RSNG     2012-06-05 17:42:51]: Step 8 of 20 Updating recovery media
```

```
[RSNG     2012-06-05 17:42:51]: Step 9 of 20 Extracting incoming image
[RSNG     2012-06-05 17:42:51]: Step 10 of 20 Unpacking OS packages
[RSNG     2012-06-05 17:42:51]: Step 11 of 20 Mounting jbase package
[RSNG     2012-06-05 17:42:51]: Step 12 of 20 Creating base OS symbolic links
[RSNG     2012-06-05 17:42:51]: Step 13 of 20 Creating fstab
[RSNG     2012-06-05 17:42:51]: Step 14 of 20 Creating new system files
[RSNG     2012-06-05 17:42:51]: Step 15 of 20 Adding jbundle package
[RSNG     2012-06-05 17:42:51]: Step 16 of 20 Backing up system data
[RSNG     2012-06-05 17:42:51]: Step 17 of 20 Setting up shared partition data
[RSNG     2012-06-05 17:42:51]: Step 18 of 20 Checking package sanity in installation
[RSNG     2012-06-05 17:42:51]: Step 19 of 20 Unmounting and cleaning up temporary file systems
[RSNG     2012-06-05 17:42:51]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG     2012-06-05 17:43:36]: Rebooting Backup RE
[RSNG     2012-06-05 17:43:36]: ------- Rebooting fpc1 -------
[RSNG     2012-06-05 17:50:12]: Initiating Chassis In-Service-Upgrade
[RSNG     2012-06-05 17:50:33]: Upgrading group: 0 fpc: 0
[RSNG     2012-06-05 17:52:38]: Upgrade complete for group:0
[RSNG     2012-06-05 17:52:38]: Upgrading group: 1 fpc: 1
[RSNG     2012-06-05 17:54:42]: Upgrade complete for group:1
[RSNG     2012-06-05 17:54:42]: Finished processing all upgrade groups, last group :1
[RSNG     2012-06-05 17:54:48]: Preparing for Switchover
[RSNG     2012-06-05 17:55:38]: Switchover Completed
[Status   2012-06-05 17:55:41]: Upgrade completed with 0 errors
Success
```

request system software nonstop-upgrade node-group (Server Node Group)

```
user@qfabric> request system software nonstop-upgrade node-group P1507-C jinstall-
qfabric-12.2X50-D10.3.rpm
Upgrading target(s): P1507-C

[P1507-C  2012-06-26 14:02:44]: Retrieving package
[P1507-C  2012-06-26 14:03:21]: ------- P1507-C: -------
[P1507-C  2012-06-26 14:03:59]: Step 1 of 20 Creating temporary file system
[P1507-C  2012-06-26 14:03:59]: Step 2 of 20 Determining installation source
[P1507-C  2012-06-26 14:03:59]: Step 3 of 20 Processing format options
[P1507-C  2012-06-26 14:03:59]: Step 4 of 20 Determining installation slice
[P1507-C  2012-06-26 14:04:00]: Step 5 of 20 Creating and labeling new slices
[P1507-C  2012-06-26 14:04:00]: Step 6 of 20 Create and mount new file system
[P1507-C  2012-06-26 14:04:08]: Step 7 of 20 Getting OS bundles
[P1507-C  2012-06-26 14:04:09]: Step 8 of 20 Updating recovery media
```

```
[P1507-C  2012-06-26 14:04:29]: Step 9 of 20 Extracting incoming image
[P1507-C  2012-06-26 14:05:42]: Step 10 of 20 Unpacking OS packages
[P1507-C  2012-06-26 14:05:49]: Step 11 of 20 Mounting jbase package
[P1507-C  2012-06-26 14:06:14]: Step 12 of 20 Creating base OS symbolic links
[P1507-C  2012-06-26 14:07:15]: Step 13 of 20 Creating fstab
[P1507-C  2012-06-26 14:07:15]: Step 14 of 20 Creating new system files
[P1507-C  2012-06-26 14:07:16]: Step 15 of 20 Adding jbundle package
[P1507-C  2012-06-26 14:09:52]: Step 16 of 20 Backing up system data
[P1507-C  2012-06-26 14:10:07]: Step 17 of 20 Setting up shared partition data
[P1507-C  2012-06-26 14:10:07]: Step 18 of 20 Checking package sanity in installation
[P1507-C  2012-06-26 14:10:08]: Step 19 of 20 Unmounting and cleaning up temporary file systems
[P1507-C  2012-06-26 14:10:11]: Step 20 of 20 Setting da0s2 as new active partition
[Status   2012-06-26 14:10:25]: Trying to Connect to Node: P1507-C
[Status   2012-06-26 14:10:32]: Rebooting P1507-C
[Status   2012-06-26 14:10:32]: Upgrade completed with 0 errors
Success
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- *Nonstop Software Upgrade Checklist for QFabric Systems*
- *Performing a Nonstop Software Upgrade on the QFabric System*
- *Verifying Nonstop Software Upgrade for QFabric Systems*
- show chassis nonstop-upgrade node-group | **18585**

request system software nonstop-upgrade

IN THIS SECTION

- Syntax | **15320**
- Description | **15320**
- Options | **15322**

- Required Privilege Level | 15324
- Output Fields | 15324
- Sample Output | 15324
- Release Information | 15332

Syntax

```
request system software nonstop-upgrade (package-name | set [package-name package-name])
<force-host>
<no-copy>
<no-old-master-upgrade>
<reboot >
<unlink>
```

Description

Perform a nonstop software upgrade (NSSU) on a switch with redundant Routing Engines or on a Virtual Chassis or Virtual Chassis Fabric (VCF). The behavior of this command depends on the type of switch, Virtual Chassis, or VCF where you run it, as follows:

- When you run this command on any of the following Virtual Chassis or VCF configurations, NSSU upgrades all members of the Virtual Chassis:
  - EX3300, EX3400, EX4200, EX4300,EX4400, EX4500, EX4550, EX4600, or EX4650-48Y Virtual Chassis
  - Mixed Virtual Chassis composed of any combination of EX4200, EX4500, and EX4550 switches, or EX4300 and EX4600 switches
  - QFX3500 and QFX3600 Virtual Chassis
  - QFX5100 Virtual Chassis
  - QFX5120-48Y, QFX5120-48T or QFX5120-32C Virtual Chassis
  - Fixed configuration of switches in a VCF (QFX3500/QFX3600 and QFX5100 switches)
  - Mixed VCF composed of any combination of QFX3500/QFX3600, QFX5100, and EX4300 switches

The original Virtual Chassis or VCF backup becomes the primary. The new primary automatically upgrades and reboots the original primary, which then rejoins the Virtual Chassis or VCF as the backup.

- When you run this command on an EX6200 or EX8200 switch, NSSU upgrades both the backup and primary Routing Engines. The original backup Routing Engine becomes the new primary at the end of the upgrade.
  - On an EX6200 switch, NSSU automatically reboots the original primary Routing Engine.
  - On an EX8200 switch, NSSU does not automatically reboot the original primary Routing Engine unless you specify the `reboot` option.
- When you run this command on an EX8200 Virtual Chassis, NSSU upgrades all primary and backup Routing Engines in the Virtual Chassis, including the external Routing Engines. The original backup Routing Engines become the new primary Routing Engines. NSSU does not automatically reboot the original primary Routing Engines unless you specify the `reboot` option.

This command has the following requirements:

- All Virtual Chassis members, VCF members, and all Routing Engines must be running the same Junos OS release.
- You must enable Graceful Routing Engine switchover (GRES)..
- You must enable Nonstop active routing (NSR).

> **NOTE**: Although not required, we recommend you also enable nonstop bridging (NSB). NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover during NSSU. See Configuring Nonstop Bridging on EX Series Switches (CLI Procedure).

- You must run the command from the primary Routing Engine on a standalone switch or from the primary on a Virtual Chassis.
- For minimal traffic disruption, you must define link aggregation groups (LAGs) such that the member links reside on different Virtual Chassis or VCF member switches (or on different line cards for EX6200 and EX8200 switches and EX8200 Virtual Chassis).
- For all Virtual Chassis (except EX8200 Virtual Chassis):
  - The Virtual Chassis members must be connected in a ring topology. A ring topology prevents the Virtual Chassis from splitting during an NSSU.

- The Virtual Chassis primary and backup must be adjacent to each other in the ring topology. With adjacent placement, the primary and backup are always in sync while the switches in line-card roles are rebooting.
- The Virtual Chassis must be preprovisioned so the line-card role is explicitly assigned to member switches acting in a line-card role. During an NSSU, the primary and backup member switches must maintain their Routing Engine roles (although the primary role switches to the backup), and the remaining switches must maintain their line-card roles.
- In a two-member Virtual Chassis, you must configure `no-split-detection` so the Virtual Chassis doesn't split during NSSU.

- For Virtual Chassis Fabric:
  - You can only have two members preprovisioned in the Routing Engine role. If more than two Routing Engines are configured, NSSU issues a warning message and the NSSU process stops.
  - The VCF members should be connected in a spine and leaf topology. A spine and leaf topology prevents the VCF from splitting during NSSU. Each leaf device must be connected to both spine devices.
  - The VCF must be preprovisioned so that the line-card role has been explicitly assigned to member switches acting in a line-card role, and likewise the Routing Engine role has been explicitly assigned to the member switches acting in a Routing Engine role. During an NSSU, the primary and backup member switches must maintain their Routing Engine roles (although the primary role switches to the backup), and the remaining switches must maintain their line-card roles.
  - You must configure `no-split-detection` in a two-member VCF so the VCF does not split during NSSU.

Options

package-name

Location of the software package or bundle to be installed. For example:

- `/var/tmp/`*`package-name`*—For a software package or bundle installed from a local directory on the switch.
- *`protocol`*`://`*`hostname`*`/`*`pathname`*`/`*`package-name`*—For a software package or bundle downloaded and installed from a remote location. Replace *`protocol`* with one of the following:
  - `ftp`—File Transfer Protocol.
    Use `ftp://`*`hostname`*`/`*`pathname`*`/`*`package-name`*.

To specify authentication credentials, use `ftp://<username>:<password>@hostname/pathname/package-name`.
To have the system prompt you for the password, specify `prompt` in place of the password.
The command displays an error message If a password is required and you do not specify the password or `prompt`.

- `http`—Hypertext Transfer Protocol.
  Use `http://hostname/pathname/package-name`.
  To specify authentication credentials, use `http://<username>:<password>@hostname/pathname/package-name`.
  The command prompts you for a password if one is required and you didn't include it.

- `scp`—Secure copy (available only for Canada and U.S. version).
  Use `scp://hostname/pathname/package-name`.
  To specify authentication credentials, use `scp://<username>:<password>@hostname/pathname/package-name`.

> **NOTE**: The *pathname* in the protocol is the relative path to the user home directory on the remote system and not the root directory.

set [*package-name package-name*]
(Mixed Virtual Chassis only) Locations of the different installation packages required by the different types of member switches. These packages must be for the same Junos OS release. See this command's *package-name* option for information about how to specify the installation packages.

force-host
(Optional) Force adding the host software package or bundle (and ignore warnings) on EX4650, QFX5100, or QFX5120 devices.

no-copy
(Optional) Install a software package or bundle, but do not save copies of the package or bundle files.

no-old-master-upgrade
(Optional) (EX8200 switches only) Upgrade the backup Routing Engine only. After the upgrade completes, the original primary Routing Engine becomes the backup Routing Engine and continues running the previous software version.

reboot
(Optional) (EX8200 switches and EX8200 Virtual Chassis only) When you include the `reboot` option, NSSU automatically reboots the original primary (new backup) Routing Engine after being upgraded to the new software. When you omit the `reboot` option, you must manually reboot the original primary (new backup) Routing Engine using the `request system reboot` command.

> **NOTE**: If you do not use the `reboot` option on an EX8200 Virtual Chassis, you must establish a connection to the console port on the Switch Fabric and Routing Engine (SRE) module or Routing Engine (RE) module to manually reboot the backup Routing Engines.

unlink (Optional) Remove the software package after a successful upgrade.

Required Privilege Level

maintenance

Output Fields

This command reports feedback on the status of the request. Some functions are shared between NSSU and the in-service software upgrade (ISSU) feature, so you might see what appear to be ISSU messages as well as NSSU messages in the output from this command.

Sample Output

request system software nonstop-upgrade (EX4200 Virtual Chassis)

```
user@switch> request system software nonstop-upgrade
/var/tmp/jinstall-ex-4200-12.1R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Installing image on other FPC's along with the backup

Checking pending install on fpc1
Pushing bundle to fpc1
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc1

Checking pending install on fpc2
Pushing bundle to fpc2
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
```

```
Completed install on fpc2

Checking pending install on fpc3
Pushing bundle to fpc3
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc3

Checking pending install on fpc4
Pushing bundle to fpc4
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc4

Checking pending install on fpc5
Pushing bundle to fpc5
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc5

Checking pending install on fpc6
Pushing bundle to fpc6
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc6

Checking pending install on fpc7
Pushing bundle to fpc7
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc7
Backup upgrade done
Rebooting Backup RE

Rebooting fpc1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
```

```
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online
  FPC 1          Online
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
  FPC 4          Online (ISSU)
  FPC 5          Online (ISSU)
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)
Going to install image on master
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE

*** FINAL System shutdown message from root@switch ***

System going down IMMEDIATELY



Shutdown NOW!
[pid 9336]
```

request system software nonstop-upgrade (EX6200 Switch)

```
{master}
user@switch> request system software nonstop-upgrade
/var/tmp/jinstall-ex-6200-12.2R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re0
NOTICE: Validating configuration against jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Backup upgrade done
```

```
Rebooting Backup RE

Rebooting re0
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
  FPC 4          Online
  FPC 5          Online
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)
  FPC 8          Online (ISSU)
  FPC 9          Online (ISSU)
Going to install image on master
NOTICE: Validating configuration against jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE
Trying to relinquish mastership before rebooting...
Resolving mastership...
Complete. The other routing engine becomes the master.

*** FINAL System shutdown message from user@switch ***

System going down IMMEDIATELY
```

request system software nonstop-upgrade reboot (EX8200 Switch)

```
{master}
user@switch> request system software nonstop-upgrade reboot
/var/tmp/jinstall-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 2          Offline                 Offlined by CLI command
  FPC 3          Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
[pid 2635]
```

```
*** FINAL System shutdown message from user@switch ***
System going down IMMEDIATELY
```

request system software nonstop-upgrade no-old-master-upgrade (EX8200 Switch)

```
{master}
user@switch> request system software nonstop-upgrade no-old-master-upgrade
/var/tmp/jinstall-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 3          Offline                 Offlined by CLI command
  FPC 4          Online (ISSU)
  FPC 5          Online (ISSU)
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
```

```
Skipping Old Master Upgrade
ISSU: IDLE
```

request system software nonstop-upgrade reboot (EX8200 Virtual Chassis)

```
{master:9}
user@external-routing-engine> request system software nonstop-upgrade reboot
/var/tmp/jinstall-ex-xre200-11.1-20101130.0-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing LCC Backup REs
ISSU: Preparing Backup RE
Pushing bundle /var/tmp/jinstall-ex-xre200-11.1-20101130.0-domestic-signed.tgz to member8
----------------------------------------------------------------------------
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
VC Backup upgrade done
Rebooting VC Backup RE

Rebooting member8
ISSU: Backup RE Prepare Done
Waiting for VC Backup RE reboot
Pushing bundle to member0-backup
Pushing bundle to member1-backup
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately

Rebooting member0-backup
Rebooting LCC [member0-backup]

Rebooting member1-backup
Rebooting LCC [member1-backup]
ISSU: LCC Backup REs Prepare Done
GRES operational
Initiating Chassis Nonstop-Software-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
```

```
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking Nonstop-Upgrade status
member0:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 5          Online (ISSU)

member1:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Offline                 Offlined due to config
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
  FPC 4          Online (ISSU)
  FPC 5          Online (ISSU)
  FPC 7          Online (ISSU)

member0:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 5          Online (ISSU)

member1:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Offline                 Offlined due to config
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
  FPC 4          Online (ISSU)
  FPC 5          Online (ISSU)
  FPC 7          Online (ISSU)
ISSU: Upgrading Old Master RE
Pushing bundle /var/tmp/incoming-package-8200.tgz to member0-master
Pushing bundle /var/tmp/incoming-package-8200.tgz to member1-master
```

```
ISSU: RE switchover Done
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Rebooting ...
shutdown: [pid 2188]
Shutdown NOW!
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!

*** FINAL System shutdown message from root@ ***
System going down IMMEDIATELY
```

Release Information

Command introduced in Junos OS Release 10.4.

Option `set [package-name package-name]` added in Junos OS Release 12.1 for EX Series switches.

RELATED DOCUMENTATION

- show chassis nonstop-upgrade | **18581**
- Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)
- Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)
- Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade
- *Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade*

request system software recover-from-restore-point

IN THIS SECTION

- Syntax | **15333**
- Description | **15333**

- Required Privilege Level | 15333
- Sample Output | 15333
- Release Information | 15334

Syntax

```
request system software recover-from-restore-point
```

Description

Use to command to rollback to a previously created restore-point.

Rolling back to a previously created restore-point might disrupt traffic, as both Director devices reboot from the restore-point partition.

Required Privilege Level

configure

Sample Output

request system software recover-from-restore-point

```
root@qfabric> request system software recover-from-restore-point
Start Restore
Checking if the restore-point exists
    LogVol00 has the root filesystem
    Found Restore-Point:  Fri Aug 15 07:42:39 UTC 2014 /dev/VolGroup00/LogVol03
    Mounting restore-volume LogVol03
Checking the sanity of restore-point
    Checking if the restore DB is present
    Checking if the restore grub is present
Checking the current state of the system
Checking the state of cluster services
Checking the inventory
Checking if the peer is reachable
```

```
    Checking if peer is reachable via Compute Node Monitor
    Successfully communicated with peer over 169.254.0.2
Intimating the peer to do stage INITIATE_PEER_INITIAL_STAGE of downgrade
Preparing the system to downgrade
Prepping all Junos devices
Checking status at Peer
    Downgrade first stage at peer concluded successfully
Initiating final stage of downgrade in peer
Intimating the peer to do stage INITIATE_PEER_FINAL_STAGE of downgrade
Modify loader to boot from restore-point
Move mount points to new filesystem
Force Reboot
Rebooting....
```

Release Information

Command introduced in Junos OS Release 14.1X53-D15.

RELATED DOCUMENTATION

request system software restore-point | 15338

request system software recovery-package

IN THIS SECTION

- Syntax | 15335
- Description | 15335
- Options | 15335
- Required Privilege Level | 15336
- Output Fields | 15336
- Sample Output | 15336
- Release Information | 15337

Syntax

```
request system software recovery-package ( add package-name | delete package-name | extract (ex-
xre200-package | ex-8200-package) package-name)
<backup>
<master>
```

Description

Add, remove, or extract a software recovery package for the EX8200 switch or the XRE200 External Routing Engine in an EX8200 Virtual Chassis.

Use this command to load a software recovery package onto the USB flash drive of the external Routing Engine and then extract it to the recovery partition of the external Routing Engine—so that it is available for use during a software recovery installation. During a software recovery installation procedure, you are prompted to select which software recovery package to use for the recovery.

> **NOTE**: This command is available only on the external Routing Engine.

If you do not have the desired software recovery package already installed in the recovery partition of the external Routing Engine, download the software package from the Internet.

Options

An EX8200 Virtual Chassis supports up to two external Routing Engines for the purpose of external Routing Engine redundancy. If you have configured two external Routing Engines, one functions as the primary external Routing Engine and the other functions as the backup. If you have not configured two external Routing Engines, the `backup` and `master` options are not relevant.

add *package-name* backup | master

Load the EX8200 software recovery package or the XRE200 software recovery package onto the USB flash drive of the external Routing Engine.

> **NOTE**: You can have up to five software recovery packages loaded in the USB flash drive.

- ***package-name***—path name of the software recovery package.
- **backup**—(Optional) Load the specified package onto the USB flash drive of the backup external Routing Engine.

- **master**—(Optional) Load the specified package onto the USB flash drive of the primary external Routing Engine.

delete *package-name* backup | master

Delete the EX8200 software recovery package or the XRE200 software recovery package from the USB drive of the external Routing Engine.

- ***package-name***—Path name of the software package.
- **backup**—(Optional) Delete the specified package from the USB flash drive of the backup external Routing Engine.
- **master**—(Optional) Delete the specified package from the USB flash drive of the primary external Routing Engine.

extract (ex-xre200-package | ex-8200-package) *package-name* backup | master

Extract either the XRE200 software recovery package or the EX8200 software recovery package from the USB flash drive of the external Routing Engine and put it into the recovery partition of the external Routing Engine.

- ***package-name***—Path name of the software recovery package.
- **backup**—(Optional) Extract the specified package from the USB flash drive to the recovery partition of the backup external Routing Engine.
- **master**—(Optional) Extract the specified package from the USB flash drive to the recovery partition of the primary external Routing Engine.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software recovery-package add

```
user@external-routing-engine> request system software recovery-package add /tftpboot/jinstall-ex-
xre200-11.1I20110608_1356_pravasp-domestic.tgz

...-ex-xre200-11.1I20110608_1356_pravasp-domestic.tgz
```

```
Member 1 :
This feature is not supported on this device
Inorder to activate this feature, please reinstall the device from USB with a package not lesser
than 11.4
Member 0 :
Validating package... Done.
Please wait while the package is being uploaded.
jinstall-ex-xre200-11.1I20110608_1356_pravasp-domestic.tgz was added successfully.
```

request system software recovery-package delete

```
user@external-routing-engine> request system software recovery-package delete jinstall-ex-
xre200-11.1I20110608_1356_pravasp-domestic.tgz

Member 1 : jinstall-ex-xre200-11.1I20110608_1356_pravasp-domestic.tgz was removed successfully.
Member 0 : Requested package doesn't exist in the recovery media.
```

request system software recovery-package extract

```
user@external-routing-engine> request system software recovery-package extract ex-xre200-package
jinstall-ex-xre200-11.1I20110608_1356_pravasp-domestic.tgz

...kage jinstall-ex-xre200-11.1I20110608_1356_pravasp-domestic.tgz
Package extracted to member1:/tftpboot/jinstall-ex-xre200-11.1I20110608_1356_pravasp-domestic.tgz
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

show system software recovery-package | 27092

request system software restore-point

IN THIS SECTION

- Syntax | 15338
- Description | 15338
- Required Privilege Level | 15338
- Sample Output | 15338
- Release Information | 15340

Syntax

```
request system software restore-point
```

Description

Use this command to create a restore-point. A restore-point is a snapshot of the software on the QFabric system as well as the configuration that can be rolled back to in cases where a software upgrade or configuration changes have made the QFabric system unstable or inoperable.

Required Privilege Level

configure

Sample Output

request system software restore-point

```
root@qfabric> request system software restore-point
Checking if director-device upgrade is currently in progress.
Checking VM status.
Checking for communication between director devices.
Checking inventory status of all components.
Checking Server INE passwords.
```

```
Checking FC passwords.
Checking CCPC passwords.
Checking FM-0 passwords.
Checking DRE-0 passwords.
Checking NW-NG-0 passwords.
Checking chassis alarms.
0
sent command to peer to start operation
sanity checks passed
Performing fdisk
restore partition created
creating restore partition on physical disk
device /dev/sda: start 0 size -388718592
gpt: 0 slices
dos: 4 slices
# 1:        63-   208844 (   208782 sectors,    106 MB)
# 2:    208845-1048771394 (1048562550 sectors, 536864 MB)
# 3: 1048771395-1146446594 ( 97675200 sectors,  50009 MB)
# 4: 1146446595-2146460714 (1000014120 sectors, 512007 MB)
performing physical volume creation
  Physical volume "/dev/sda4" successfully created
  "/dev/sda4" is a new physical volume of "476.84 GB"
  PV Name               /dev/sda4
extending volume group 00
  Volume group "VolGroup00" successfully extended
Creating Logical Volume
  Logical volume "LogVol03" created
  LV Name                /dev/VolGroup00/LogVol03
Restore volume selected is /dev/VolGroup00/LogVol03
Formatting restore volume
mke2fs 1.39 (29-May-2006)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
62504960 inodes, 124993536 blocks
6249676 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
3815 block groups
32768 blocks per group, 32768 fragments per group
16384 inodes per group
Superblock backups stored on blocks:
```

```
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000

Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 22 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
/dev/VolGroup00/LogVol03: UUID="a9fafbaf-da3c-417f-bd53-def01fbf3936" SEC_TYPE="ext2"
TYPE="ext3"
Restore Volume mounted
backing up root filesystem..this will take a few minutes
Copying files from tmp..this may take a few minutes
Dumping databases...this may take a few minutes
backing up shared filesystem..this may take a few minutes
Restore point creation finished for dg0 on /dev/VolGroup00/LogVol03
waiting 10 mins for for peer dg to finish
Restore point creation success on both DGs
```

Release Information

Command introduced in Junos OS Release 14.1X53-D15.

RELATED DOCUMENTATION

request system software recover-from-restore-point | 15332

request system software rollback (Junos OS)

IN THIS SECTION

- Syntax | 15341
- Syntax (EX Series Switches) | 15341
- Syntax (MX Series Router) | 15341
- Syntax (SRX Series Firewalls) | 15342

- Description | 15342
- Options | 15342
- Required Privilege Level | 15343
- Output Fields | 15343
- Sample Output | 15343
- Release Information | 15344

Syntax

```
request system software rollback
```

Syntax (EX Series Switches)

```
request system software rollback
<all-members>
<local>
<member member-id>
<reboot>
```

Syntax (MX Series Router)

```
request system software rollback
<all-members>
<device-alias alias-name>
<local>
<member member-id
<reboot>
<satellite slot-id>
<upgrade-group [all |upgrade-group-name]>
```

Syntax (SRX Series Firewalls)

```
request system software rollback <node-id>
```

Description

Use this command to revert to the last successfully installed package before the `request system software (add | delete)` command. It uses the copy stored in the **/var/sw/pkg** directory.

Additional Information

- On Junos Fusion, the `request system software rollback` command can be used to roll back the version of satellite software associated with a satellite software upgrade group. Rolling back the version of satellite software associated with a satellite software upgrade group triggers a satellite software upgrade.
- On SRX Series Firewalls with dual root systems, when `request system software rollback` is run, the system switches to the alternate root. Each root can have a different version of Junos OS. Roll back takes each root back to the previously installed image.
- For SRX300, SRX320, SRX340, SRX345, and SRX380 firewalls, you cannot use the `request system rollback` command to roll back from Junos OS Release 24.4R1 to either Junos OS Release 23.4R2-S3 or to Release 24.2R2. Instead, you must treat the rollback as a downgrade, and use the `request system software add` *`package-name`* `no-validate reboot` command instead.
- On QFX3500 and QFX3600 devices in a mixed Virtual Chassis, when the `request system software rollback` command is issued, the system does not rollback to the image stored in the alternate partition.
- On QFX5100 switches, the `reboot` option has been removed. To reboot the switch after a software rollback, issue the `request system reboot` command as a separate, secondary command.

Options

all-members (EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on all members of the Virtual Chassis configuration.

device-alias *alias-name* (Junos Fusion only) (Optional) Rollback the satellite software package onto the specified satellite device using the satellite devices FPC slot identifier.

local (EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on the local Virtual Chassis member.

member *member-id*	(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.
node-id	(SRX Series Firewalls only) Identification number of the chassis cluster node. It can be 0 or 1.
none	Revert to the last known good state before the most recent `request system software (add | delete)` command.
reboot	(Optional) The system reboots automatically to complete the rollback.
satellite *slot-id*	(Junos Fusion only) (Optional) Roll back the satellite software package onto the specified satellite device using the satellite devices FPC slot identifier.
upgrade-group [all |*upgrade-group-name***]**	(Junos Fusion only) Roll back the satellite software image associated with the specified satellite software upgrade group, or for all satellite software upgrade groups in the Junos Fusion when `all` is entered.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software rollback

```
user@host> request system software rollback
Verified SHA1 checksum of ./jbase-7.2R1.7.tgz
Verified SHA1 checksum of ./jdocs-7.2R1.7.tgz
Verified SHA1 checksum of ./jroute-7.2R1.7.tgz
Installing package './jbase-7.2R1.7.tgz' ...
Available space: 35495 require: 7335
Installing package './jdocs-7.2R1.7.tgz' ...
Available space: 35339 require: 3497
Installing package './jroute-7.2R1.7.tgz' ...
Available space: 35238 require: 6976
```

```
NOTICE: uncommitted changes have been saved in /var/db/config/juniper.conf.pre-install
Reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...
Restarting aprobed ...
Restarting apsd ...
Restarting cosd ...
Restarting fsad ...
Restarting fud ...
Restarting gcdrd ...
Restarting ilmid ...
Restarting irsd ...
Restarting l2tpd ...
Restarting mib2d ...
Restarting nasd ...
Restarting pppoed ...
Restarting rdd ...
Restarting rmopd ...
Restarting rtspd ...
Restarting sampled ...
Restarting serviced ...
Restarting snmpd ...
Restarting spd ...
Restarting vrrpd ...

WARNING: cli has been replaced by an updated version:
CLI release 7.2R1.7 built by builder on 2005-04-22 02:03:44 UTC
Restart cli using the new version ? [yes,no] (yes) yes

Restarting cli ...
user@host
```

Release Information

Command introduced before Junos OS Release 7.4.

`reboot` option introduced in Junos OS Release 12.3.

`device-alias`, `satellite`, and `upgrade-group` options introduced in Junos OS Release 14.2R3 for Junos Fusion.

`force` option deprecated in Junos OS Release 15.1 for Junos OS with Upgraded FreeBSD.

To find which platforms in which releases use Junos OS with upgraded FreeBSD, see Feature Explorer, enter **freebsd**, and select **Junos kernel upgrade to FreeBSD 10+**.

RELATED DOCUMENTATION

- request system software delete (Junos OS) | 15250
- request system configuration rescue delete | 15057
- request system configuration rescue save | 15059

request system software rollback (Junos OS Evolved)

IN THIS SECTION

- Syntax | 15345
- Description | 15346
- Options | 15346
- Required Privilege Level | 15346
- Output Fields | 15346
- Sample Output | 15347
- Release Information | 15348

Syntax

```
request system software rollback reboot
<no-validate>
<package-name version>
<with-old-snapshot-config>
```

Description

Use this command to revert to the last successfully installed package before the last-issued `request system software (add | delete)` command. By default, once the software is rolled back, the device uses the current configuration file. You can use this command on either Routing Engine in a dual-Routing Engine system.

On Junos OS Evolved, the `reboot` option is required in order to complete the rollback.

Options

package-name *version*
Select any installed version for the rollback. The `request system software rollback` operational mode command uses the version instead of the package-name. You can see the available versions by using the `show system software list` operational mode command. If you do not specify the version, the system rolls back to the default rollback version (the one with the '<' before it on the `show system software list` command output). You can specify any previous Junos OS Evolved release as long as it is neither the one that is currently running nor the rollback version.

no-validate
Do not check compatibility with the current configuration.

Default: validate

reboot
(Optional) Reboot to complete the rollback. If you do not specify the `reboot` option, then when this command completes, you need to issue the `request system reboot` operational mode command to reboot the system to finish the rollback process.

with-old-snapshot-config
(Optional) Rolls back the system to the specified version with the old snapshot of the configuration used in that version. Otherwise, the rollback, by default, takes the current configuration.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

request system software rollback with-old-snapshot-config

```
user@host> request system software rollback junos-evo-install-ptx-x86-64-20.4-202103141559.0
with-old-snapshot-config
```

```
Starting software rollback to version junos-evo-install-ptx-x86-64-20.4-202103141559.0-EVO
Software rollback in progress...
re0: Boot version is now 'junos-evo-install-ptx-x86-64-20.4-202103141559.0-EVO'
re0: Rollback Done. Next Boot version is now 'junos-evo-install-ptx-x86-64-20.4-202103141559.0-
EVO'. Must reboot now to activate.
re0: Syncing nodes
re1: Boot version is now 'junos-evo-install-ptx-x86-64-20.4-202103141559.0-EVO'
re0: All nodes synced
Software rollback succeeded.
NOTICE: 'pending' rollback version will be activated at next reboot...
```

request system software rollback reboot

```
user@host> request system software rollback reboot
Starting software rollback to default rollback version
Software rollback in progress...
re0: Validating current config for rollback version junos-evo-install-ptx-
x86-64-20.4-202103151929.0-EVO
re0: Validating existing configs. See /var/log/validation_config.log for config validation logs.
re0: Validation Passed
re0: Validation passed for version junos-evo-install-ptx-x86-64-20.4-202103151929.0-EVO
re0: Copying current config to rollback version
re0: Boot version is now 'junos-evo-install-ptx-x86-64-20.4-202103151929.0-EVO'
re0: Rollback Done. Next Boot version is now 'junos-evo-install-ptx-x86-64-20.4-202103151929.0-
EVO'. Must reboot now to activate.
re0: Syncing nodes
re1: Config fetch successful
re1: Res Config fetch successful
re1: Boot version is now 'junos-evo-install-ptx-x86-64-20.4-202103151929.0-EVO'
re0: All nodes synced
Software rollback succeeded.
Rebooting...
```

```
{master}
user@host-re0>
System going down IMMEDIATELY

Software rollback in reboot mode succeeded. Rebooting

Connection to host closed by remote host.
Connection to host closed.
```

Release Information

`validate` and `no-validate` options introduced for Junos OS Evolved Release 18.3R1.

`package-name` *version* option introduced for Junos OS Evolved Release 18.3R1.

`with-old-snapshot-config` option introduced for Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- request system snapshot (Junos OS Evolved) | 15187
- request system software add (Junos OS Evolved) | 15233
- request system software delete (Junos OS Evolved) | 15254
- request system software sync | 15348
- request system software validate (Junos OS Evolved) | 15365
- show system software list | 27088

request system software sync

IN THIS SECTION

- Syntax | 15349
- Description | 15349
- Options | 15349
- Additional Information | 15350

- Required Privilege Level | 15351
- Sample Output | 15351
- Release Information | 15352

Syntax

```
request system software sync (current | rollback | all-versions)
```

Description

Use this command on the primary Routing Engine of a system to synchronize the software and configurations from the primary Routing Engine to the other nodes and reboot the other nodes. The configurations are synchronized even if the images are identical. If specified on the backup Routing Engine, the command fails.

NOTE: The `request system software sync` command is only supported on PTX Series and QFX Series devices that have two Routing Engines.

Options

`current | rollback | all-versions`

Specify which software version (current, rollback, or all versions) to sync to the other node:

- For the `current` option, the system syncs the current version from the primary Routing Engine to the other node and reboots the other node with that version. If the current version on the primary Routing Engine of the system matches the version on the other node of the system, the command fails.
- For the `rollback` option, the system synchronizes the rollback version to the other node. If the rollback version on the primary Routing Engine of the system matches the rollback version on the other node, the command fails.
- For the `all-versions` option, the system synchronizes all software versions and configurations from the primary Routing Engine to the other node and reboots the other node if there's a mismatch between current versions.

Additional Information

For the `all-versions` option, the synchronization proceeds as follows:

1. All versions are deleted on the other node, except for the current version.
2. The current and rollback images and configurations are copied from the primary Routing Engine to the other node. Even if the software versions match, the configuration is copied and the software proceeds to the next image.
3. Any other versions and configurations are copied from the primary Routing Engine to the other node.
4. If the current version on the other node is not the same as current on the primary Routing Engine, then the other node is rebooted after warning the user.

The `request system software sync all-versions` command is successful if the first two steps of the synchronization are successful. If step 3 fails, a warning message is displayed. To make the versions match, you can delete the extra versions using the `request system software delete` operational mode command.

To see what the software versions are available on the device, use the `show system software list` command.

If an Routing Engine that has a different software version is inserted into the system, the Routing Engine is kept outside the system and a software mismatch alarm is generated, which specifies the Routing Engine name and the version of software on that Routing Engine, similar to the following: `Software Version Mismatch on re1:junos-evo-install-ptx-x86-64-20.4R2.6-EVO.` To clear this alarm, use the `request system software sync all-versions` command to synchronize the software. Once the new Routing Engine comes back up, it joins the system.

```
user@host-re0> show system alarms
2 alarms currently active
Alarm time               Class  Description
2021-04-19 16:02:26 PDT  Major  Re1 Node unreachable
2021-04-19 16:04:46 PDT  Major  Software Version Mismatch on re1:junos-evo-install-ptx-
x86-64-20.4R2.6-EVO
```

For the `current` option, before you switch control to a newly-inserted Routing Engine, ensure all images are synchronized to the newly-inserted Routing Engine by using the output from the `show system software list` operational mode command to compare the images installed on both Routing Engines and make sure they are the same. You must make sure that the system has finished synchronizing all of the images in the background before you switch control to the newly- inserted Routing Engine to ensure that the newly-inserted Routing Engine does not remove any images from the existing Routing Engine.

Required Privilege Level

view

Sample Output

request system software sync current

```
user@host-re0> request system software sync current
warning: Erase software versions present on the other RE node and sync software versions from
Master RE node
Erase software versions on the other RE and sync from Master RE? [yes,no] (no) yes

The current version on master RE - junos-evo-install-ptx-x86-64-20.4R2.14-EVO
The current version on other RE - junos-evo-install-ptx-x86-64-20.4R2.13-EVO
Transfer software version files for junos-evo-install-ptx-x86-64-20.4R2.14-EVO to node re1...
Sync in progress for /soft/junos-evo-install-ptx-x86-64-20.4R2.14-EVO...
/data/var/home/root
Sync in progress for /data/junos-evo-install-ptx-x86-64-20.4R2.14-EVO...
Sync in progress for /data/config/junos-evo-install-ptx-x86-64-20.4R2.14-EVO...
Sync in progress for /data/etc/junos-evo-install-ptx-x86-64-20.4R2.14-EVO...
Sync in progress for /data/var_db/junos-evo-install-ptx-x86-64-20.4R2.14-EVO...
Sync in progress for /data/usr_conf/junos-evo-install-ptx-x86-64-20.4R2.14-EVO...
Sync in progress for /data/usr_evo_share/junos-evo-install-ptx-x86-64-20.4R2.14-EVO...
Sync in progress for /data/var_pfe/junos-evo-install-ptx-x86-64-20.4R2.14-EVO...
Sync in progress for /data/var_etc/junos-evo-install-ptx-x86-64-20.4R2.14-EVO...
Setting up software version files for junos-evo-install-ptx-x86-64-20.4R2.14-EVO on re1
Sync in progress /data/var/home...
Software sync completed for junos-evo-install-ptx-x86-64-20.4R2.14-EVO
Warning: Rebooting re1
Please run 'show system software list' to see SW versions installed in all nodes
```

request system software sync rollback (Versions Are Already The Same)

```
user@host> request system software sync rollback
warning: Erase software versions present on the other RE node and sync software versions from
Master RE node
Erase software versions on the other RE and sync from Master RE? [yes,no] (no) yes

The rollback version on master RE - junos-evo-install-ptx-x86-64-20.4R2.14-EVO
```

```
The rollback version on other RE - junos-evo-install-ptx-x86-64-20.4R2.14-EVO
Warning: The rollback version junos-evo-install-ptx-x86-64-20.4R2.14-EVO matches to rollback
version on other-RE. Skipping software sync
Software sync completed for junos-evo-install-ptx-x86-64-20.4R2.14-EVO

Please run 'show system software list' to see SW versions installed in all nodes
```

request system software sync all-versions

```
user@host-re0 request system software sync all-versions
warning: Erase software versions present on the other RE node and sync software versions from
Master RE node
Erase software versions on the other RE and sync from Master RE? [yes,no] (no) ...yes

Cleanup old software versions on re1
The current version on master RE - junos-evo-install-ptx-x86-64-20.4R2.14-EVO
The current version on other RE - junos-evo-install-ptx-x86-64-20.4R2.13-EVO
Sync in progress for /soft/junos-evo-install-ptx-x86-64-20.4R2.14-EVO...
The rollback version on master RE - junos-evo-install-ptx-x86-64-20.4R2.13-EVO
The rollback version on other RE - junos-evo-install-ptx-x86-64-20.4R2.14-EVO
Sync in progress for /soft/junos-evo-install-ptx-x86-64-20.4R2.13-EVO...
Software sync completed for all versions
Warning: Rebooting re1
Please run 'show system software list' to see SW versions installed in all nodes
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

show system software list | 27088

request system software rollback (Junos OS Evolved) | 15345

request system software system-backup

IN THIS SECTION

- Syntax | 15353
- Description | 15353
- Options | 15353
- Required Privilege Level | 15354
- Output Fields | 15354
- Sample Output | 15354
- Release Information | 15355

Syntax

```
request system software system-backup
<usb-create>
```

Description

(QFabric systems only) Save a copy of the current QFabric system configuration file and the current software package for recovery purposes. You can use these saved files to restore your QFabric system to full operation after a system failure or shutdown.

Options

none Copy the QFabric system software package and system configuration file to a Director group recovery directory.

> **NOTE**: If this command fails, insert a Juniper Networks software installation USB flash drive into the primary Director device and issue the `request system software format-qfabric-backup` command. For more details about this prerequisite procedure that is required before you can use the QFabric system backup and

recovery feature, see *Performing System Backup and Recovery for a QFabric System*.

usb-create Copy the QFabric system software package and system configuration file from the Director group recovery directory to a USB flash drive. When the files have been copied, you can use the USB flash drive to help your QFabric system recover from a failure condition.

> **NOTE**: You must issue the `request system software system-backup` command (which saves the files to the Director group) before you can issue the `request system software system-backup usb-create` command.

Required Privilege Level

maintenance

Output Fields

When you enter these commands, you are provided feedback on the status of your request.

Sample Output

request system software system-backup

```
user@qfabric> request system software system-backup

user@qfabric>
```

request system software system-backup usb-create

```
user@qfabric> request system software system-backup usb-create /dev/sdb
Issuing this command will overwrite the contents of the USB drive.
Continue? [yes,no] (no) yes

This operation will access the USB drive on 0281042010000013.
Are you sure you want to continue? [yes,no] (no) yes
```

```
Copying QFabric recovery media to /dev/sdb...
Successfully copied QFabric recovery media to /dev/sdb
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Performing System Backup and Recovery for a QFabric System

request system software format-qfabric-backup | **15260**

save | **16315**

request system software configuration-backup

Performing a QFabric System Recovery Installation on the Director Group

request system software validate in-service-upgrade

IN THIS SECTION

- Syntax | **15356**
- Description | **15356**
- Options | **15356**
- Additional Information | **15356**
- Required Privilege Level | **15357**
- Output Fields | **15357**
- Sample Output | **15357**
- Release Information | **15360**

Syntax

```
request system software validate in-service-upgrade package-name
<enhanced-mode>
```

Description

Perform a compatibility check to ensure that the software and hardware components and the configuration on the device support unified ISSU. The `request system software validate in-service-upgrade` command enables you to detect any compatibility issues before actually issuing the `request system software in-service-upgrade` command to initiate unified ISSU.

Options

package-name

Location from which the software package or bundle is to be installed. For example:

- `/var/tmp/`*`package-name`*—For a software package or bundle that is being installed from a local directory on the router.
- *`protocol`*`://`*`hostname`*`/`*`pathname`*`/`*`package-name`*—For a software package or bundle that is to be downloaded and installed from a remote location. Replace *`protocol`* with one of the following:
  - `ftp`—File Transfer Protocol
  - `http`—Hypertext Transfer Protocol
  - `scp`—Secure copy (available only for Canada and U.S. version)

enhanced-mode

(Optional) Checks for compatibility with enhanced mode—a mode of ISSU that eliminates packet loss during the unified ISSU process.

> **NOTE**: The enhanced mode option is available on MPC7E, MPC8E, and MPC9E line cards starting from Junos OS Release 20.1R1. Support for MPC11E line cards is provided from Junos OS Release 21.4R1 onwards.

Additional Information

Unified ISSU is not supported on every platform. For a list of supported platforms, see *Unified ISSU System Requirements*.

Required Privilege Level

view

Output Fields

When you enter this command, Junos OS displays the status of your request.

Sample Output

request system software validate in-service-upgrade

```
{master}
user@host> request system software validate in-service-upgrade /var/tmp/jinstall-9.0-20080114.2-domestic-
signed.tgz reboot
Checking compatibility with configuration
Initializing...
Using jbase-9.5-20090127.0
Verified manifest signed by PackageProduction_9_5_0
Using /var/tmp/jinstall-9.6-daily-domestic-signed.tgz
Verified jinstall-9.6-20090706.0-domestic.tgz signed by PackageProduction_9_6_0
Using jinstall-9.6-20090706.0-domestic.tgz
Using jbundle-9.6-20090706.0-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jkernel-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jcrypto-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jpfe-9.6-20090706.0.tgz
Using jdocs-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jroute-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jservices-9.6-20090706.0.tgz
[: /var/validate/chroot/tmp/jservices/packages/jservices-voice-9.6-20090706.0.tgz: unexpected
operator
Auto-deleting old jservices-voice ...
Removing /opt/sdk/jservices-voice ...
Removing jservices-voice-bsg-9.5-20090127.0.tgz from /var/sw/pkg ...
```

```
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /var/sw/pkg ...
Creating /opt/sdk/jservices-voice ...
Storing jservices-voice-bsg-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-voice/jservices-voice-bsg -> /var/sw/pkg/jservices-voice-
bsg-9.6-20090706.0.tgz...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-bgf ...
Storing jservices-bgf-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-bgf/jservices-bgf-pic -> /var/sw/pkg/jservices-bgf-
pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-aacl ...
Removing /opt/sdk/jservices-aacl ...
Removing jservices-aacl-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aacl ...
Verified jservices-aacl-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-aacl ...
Storing jservices-aacl-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-aacl/jservices-aacl-pic -> /var/sw/pkg/jservices-aacl-
pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/jservices-llpdf ...
Removing jservices-llpdf-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-llpdf ...
Storing jservices-llpdf-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-llpdf/jservices-llpdf-pic -> /var/sw/pkg/jservices-llpdf-
pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/jservices-sfw ...
Removing jservices-sfw-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-sfw ...
Storing jservices-sfw-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-sfw/jservices-sfw-pic -> /var/sw/pkg/jservices-sfw-
```

```
pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-appid ...
Removing /opt/sdk/jservices-appid ...
Removing jservices-appid-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-appid ...
Storing jservices-appid-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-appid/jservices-appid-pic -> /var/sw/pkg/jservices-appid-
pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/jservices-idp ...
Removing jservices-idp-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-idp ...
Storing jservices-idp-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-idp/jservices-idp-pic -> /var/sw/pkg/jservices-idp-
pic-9.6-20090706.0.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
PIC 7/0 will be offlined (In-Service-Upgrade not supported)
PIC 7/1 will be offlined (In-Service-Upgrade not supported)
PIC 4/2 will be offlined (In-Service-Upgrade not supported)
PIC 4/3 will be offlined (In-Service-Upgrade not supported)
```

request system software validate in-service-upgrade enhanced-mode

```
{master}
user@host> request system software validate in-service-upgrade /var/tmp/junos-install-mx-x86-32-20.1.tgz
enhanced-mode
ISSU: enhanced-mode check passed
Verified junos-install-mx-x86-32-20.1 signed by PackageDevelopmentEc_2019 method ECDSA256+SHA256
Verified manifest signed by PackageDevelopmentEc_2019 method ECDSA256+SHA256
…
Checking PIC combinations
Adding junos-mx-x86-32-20.1 ...
```

```
Verified fips-mode signed by PackageDevelopmentEc_2019 method ECDSA256+SHA256
…
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
…
Validation succeeded
ISSU: Validating Image Done
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- request system software validate (Junos OS) | **15360**
- request system software in-service-upgrade | **15262**
- request system software abort | **15212**
- show chassis in-service-upgrade | **18507**
- *Getting Started with Unified In-Service Software Upgrade*
- *Example: Performing a Unified ISSU*

request system software validate (Junos OS)

IN THIS SECTION

- Syntax | **15361**
- Syntax (TX Matrix Router) | **15361**
- Syntax (TX Matrix Plus Router) | **15361**
- Syntax (MX Series Router) | **15362**
- Description | **15362**
- Options | **15362**
- Additional Information | **15363**
- Required Privilege Level | **15363**

- Output Fields | 15364
- Sample Output | 15364
- Release Information | 15365

Syntax

```
request system software validate package-name
<on (host host <username username> | routing-engine routing-engine)>
<set [package-name package-name]>
<upgrade-with-config>
<upgrade-with-config-format format>
```

Syntax (TX Matrix Router)

```
request system software validate package-name
<lcc number | scc>
<on (host host <username username> | routing-engine routing-engine)>
<set [package-name package-name]>
<upgrade-with-config>
<upgrade-with-config-format format>
```

Syntax (TX Matrix Plus Router)

```
request system software validate package-name
<lcc number | sfc number>
<on (host host <username username> | routing-engine routing-engine)>
<set [package-name package-name]>
<upgrade-with-config>
<upgrade-with-config-format format>
```

Syntax (MX Series Router)

```
request system software validate <package-name>
<member member-id>
<on (host host <username username> | routing-engine routing-engine)>
<set [package-name package-name]>
<upgrade-with-config>
<upgrade-with-config-format format>
```

Description

Use this command to validate candidate software against the current configuration of the router, the switch, or a remote host.

Options

lcc *number*

(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, validate the software bundle or package on a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, validate the software bundle or package for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

on (host *host*<username *username*> | routing-engine *routing-engine*)

(Optional) Validate the software bundle or package by comparing it to the running configuration on a remote host or Routing Engine. Specify either a host, replacing *host* with the remote hostname, or a Routing Engine, replacing *routing-engine* with the Routing Engine name. If you specify a remote host, you can optionally provide the username to be used to log in to the remote host.

member *member-id*	(MX Series routers only) (Optional) Validate the software bundle or package on the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.
package-name	Name of the software bundle or package to test.
scc	(TX Matrix routers only) (Optional) Validate the software bundle or package for the TX Matrix router (or switch-card chassis).
set [*package-namepackage-name*]	(M Series, MX Series, T Series routers) (Optional) Install multiple software packages or software add-on packages at the same time.
sfc *number*	(TX Matrix Plus routers only) (Optional) Validate the software bundle or package for the TX Matrix Plus router.
upgrade-with-config	(Optional) Install one or more configuration files.
upgrade-with-config-format *format*	(Optional) Specify the configuration file fomat, `text` or `xml`. The default format is `text`. The `upgrade-with-config` and `upgrade-with-config-format` options are only available locally on the router or switch. In a routing matrix, the configuration is applied only to the local router and is not propagated to other routers. The options are validated during the validation process and applied to the router or switch during the upgrade process. If the upgrade process is successful, the options are removed from the configuration. If the upgrade process fails, the configuration file is renamed with the .failed suffix.

Additional Information

By default, when you issue the `request system software validate` command on a TX Matrix primary Routing Engine, all the T640 primary Routing Engines that are connected to it are validated. If you issue the same command on the TX Matrix backup Routing Engine, all the T640 backup Routing Engines that are connected to it are upgraded to the same version of software.

Likewise, if you issue the `request system software validate` command on a TX Matrix Plus primary Routing Engine, all the T1600 or T4000 primary Routing Engines that are connected to it are validated. If you issue the same command on a TX Matrix Plus backup Routing Engine, all the T1600 or T4000 backup Routing Engines that are connected to it are upgraded to the same version of software.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software validate (Successful Case)

```
user@host> request system software validate /var/sw/pkg/
jbundle-5.3I20020124_0520_sjg.tgz
Checking compatibility with configuration
Initializing...
Using /packages/jbase-5.3I20020122_1901_sjg
Using /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jbase-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jkernel-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jcrypto-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jpfe-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jdocs-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jroute-5.3I20020124_0520_sjg.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete

WARNING: cli has been replaced by an updated version:
CLI release 5.3I0 built by sjg on 2002-01-24 05:23:53 UTC
Restart cli using the new version ? [yes,no] (yes)
```

request system software validate (Failure Case)

```
user@host> request system software validate 6.3/
Pushing bundle to lcc0-re0
error: Failed to transfer package to lcc0-re0



user@host> request system software validate test
Pushing bundle to lcc0-re0
Pushing bundle to lcc2-re0

lcc0-re0:
gzip: stdin: not in gzip format
```

```
tar: child returned status 1
ERROR: Not a valid package: /var/tmp/test
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

`set [package-namepackage-name]` option added in Junos OS Release 12.2 for M Series, MX Series, T Series routers.

`upgrade-with-config` and `upgrade-with-config-format format` options added in Junos OS Release 12.3 for M Series routers, MX Series routers, and T Series routers.

`on (host host<username username> | routing-engine routing-engine)` option introduced in Junos OS Release 13.3, Junos OS Release 14.1, and Junos OS Release 15.1.

RELATED DOCUMENTATION

- request system software delete (Junos OS) | 15250
- request system software rollback (Junos OS) | 15340
- Routing Matrix with a TX Matrix Plus Router Solutions Page
- request system software validate in-service-upgrade | 15355

request system software validate (Junos OS Evolved)

IN THIS SECTION

- Syntax | 15366
- Description | 15366
- Options | 15366
- Required Privilege Level | 15366
- Output Fields | 15366
- Sample Output | 15367

Release Information | 15367

Syntax

```
request system software validate package-name
<restart>
```

Description

Use this command to validate the candidate software package against the current configuration of the node. The configuration check does not change the current software or the file system.

You can use the `request system software validate` command before using the `request system software add restart` command to determine if you can upgrade to the new image with an application restart or with a reboot.

Options

package-name	Name of the software bundle or package to test.
restart	(Optional) Verify the new software configuration compatibility. When you issue the command with this option, the output lists those services that might be restarted.

Required Privilege Level

maintenance

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

request system software validate **/var/tmp/*package-name***

```
user@host-re0> request system software validate /var/tmp/ptxinstall.iso
Validating software image /var/tmp/ptxinstall.iso...
Download and Validate in Progress
re0: Starting validation : /var/tmp/ptxinstall.iso
re0: Validating version : junos-evo-install-ptx-x86-64-20.4-202103131143.0-EVO
re0: Running pre-checks for 'junos-evo-install-ptx-x86-64-20.4-202103131143.0-EVO'
re0: Pre-checks pass successfully, copying files to software area
re0: Running post install commands...
re0: Post install sequence was successful.
re0: Validating existing configs. See /var/log/validation_config.log for config validation logs.
re0: Validation Passed
re1: Running pre-checks for 'junos-evo-install-ptx-x86-64-20.4-202103131143.0-EVO'
re1: Pre-checks pass successfully, copying files to software area
re1: Running post install commands...
re1: Post install sequence was successful.
re0: Validation successful - cleaning up
re0: Removing version junos-evo-install-ptx-x86-64-20.4-202103131143.0-EVO...
re1: Removing version junos-evo-install-ptx-x86-64-20.4-202103131143.0-EVO...
Image validation succeeded.
```

Release Information

`restart` option added in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- request system snapshot (Junos OS Evolved) | **15187**
- request system software add (Junos OS Evolved) | **15233**
- request system software delete (Junos OS Evolved) | **15254**
- request system software rollback (Junos OS Evolved) | **15345**
- request system software sync | **15348**

request system software validate on (Junos OS with Upgraded FreeBSD)

IN THIS SECTION

- Syntax | 15368
- Description | 15368
- Options | 15368
- Additional Information | 15369
- Required Privilege Level | 15369
- Output Fields | 15369
- Sample Output | 15369
- Release Information | 15371

Syntax

```
request system software validate on
<host host-name [ username user-name ]>
<routing-engine (re0 | re1)>
```

Description

Use this command to validate the current configuration on a Routing Engine that is not running Junos OS with upgraded FreeBSD or a remote host.

Direct validation of a running configuration is not possible on a device running Junos OS with upgraded FreeBSD. Nevertheless, validation is an important step in the installation of an upgraded operating system. This command allows validation on a device that is not running Junos OS with upgraded FreeBSD.

Options

The specific options available are:

host *host-name* **[username** *user-name* **]**	Validate the current configuration on a remote host. The host-name is resolved through DNS. Optionally, you can supply a username to employ on the remote

host. If you omit the username option, the currently logged-in user-name is sent to the remote host.

routing-engine (re0 | re1) Validate the current configuration on another Routing Engine on the same device. The other Routing Engine cannot be running Junos OS with upgraded FreeBSD or the validation does not succeed.

Additional Information

If the authenticity of the remote host cannot be established, you are prompted to continue the validation or not. If you choose not to continue, the validation process does not take place.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software validate on host

```
user@host> request system software validate on host remote-validator
The authenticity of host 'remote-validator (192.168.164.174)' can't be established.
ECDSA key fingerprint is 73:d0:78:ce:8d:09:aa:92:4c:ce:45:52:1d:76:86:b5.
Are you sure you want to continue connecting (yes/no)? yes
Password:
**********

Sending /var/tmp/config.tgz to remote-validator...
Validating /var/tmp/config.tgz on remote-validator...
Checking compatibility with configuration
Initializing...
Using jbase-15.1-20150416.2
Verified manifest signed by PackageDevelopmentEc_2015
Using jruntime-15.1-20150416.2
Verified manifest signed by PackageDevelopmentEc_2015
Using jkernel-15.1-20150416.2
```

```
Verified manifest signed by PackageDevelopmentEc_2015
Using jroute-15.1-20150416.2
Verified manifest signed by PackageDevelopmentEc_2015
Using jcrypto-15.1-20150416.2
Verified manifest signed by PackageDevelopmentEc_2015
Using jweb-15.1-20150416.2
Verified manifest signed by PackageDevelopmentEc_2015
Using /var/tmp/config.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: warning: schema: init: 'logical-systems-vlans' contains-node 'juniper-config vlans': not
found
mgd: commit complete
Validation succeeded
```

request system software validate on routing-engine

```
user@host> request system software validate on routing-engine re1

Sending /var/tmp/config.tgz to re1...
Validating /var/tmp/config.tgz on re1...
Checking compatibility with configuration
Initializing...
Using jbase-15.1-20150416.2
Verified manifest signed by PackageDevelopmentEc_2015
Using jruntime-15.1-20150416.2
Verified manifest signed by PackageDevelopmentEc_2015
Using jkernel-15.1-20150416.2
Verified manifest signed by PackageDevelopmentEc_2015
Using jroute-15.1-20150416.2
Verified manifest signed by PackageDevelopmentEc_2015
Using jcrypto-15.1-20150416.2
Verified manifest signed by PackageDevelopmentEc_2015
Using jweb-15.1-20150416.2
Verified manifest signed by PackageDevelopmentEc_2015
Using /var/tmp/config.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: warning: schema: init: 'logical-systems-vlans' contains-node 'juniper-config vlans': not
found
```

```
mgd: commit complete
Validation succeeded
```

Release Information

Command introduced in Junos OS Release 15.1 for supported platforms. See Feature Explorer.

RELATED DOCUMENTATION

- request system reboot (Junos OS with Upgraded FreeBSD) | 15154
- show system snapshot (Junos OS with Upgraded FreeBSD) | 27075
- Release Information for Junos OS with Upgraded FreeBSD

request system software validate-restart (Junos OS Evolved)

IN THIS SECTION

- Syntax | 15371
- Description | 15372
- Options | 15372
- Additional Information | 15372
- Required Privilege Level | 15372
- Sample Output | 15372
- Release Information | 15373

Syntax

```
request system software validate-restart package-name
```

Description

Verify the new package against the current configuration to ensure that it supports unified in-service software upgrade (unified ISSU). The command output also indicates which applications can be upgraded by using an application-level restart or an in-service kernel warm restart. Use this command before you run `request system software add restart` to detect any compatibility issues before actually upgrading the software.

NOTE: You can currently use this command only on QFX5220-32CD switches.

Options

package-name Path and filename of the software package to be installed.

Additional Information

Starting in Junos OS Evolved Release 23.4R1, the `request system software validate-restart` command output summarizes the method required to perform the indicated upgrade. For example:

- Image validation succeeded. Application restart or Reboot is needed for this software image upgrade.
- Image validation succeeded. In-service kernel warm restart or Reboot is needed for this software image upgrade.
- Image validation succeeded. Reboot is needed for this software image upgrade.

Required Privilege Level

maintenance

Sample Output

request system software validate-restart

```
user@host> request system software validate-restart /var/home/regress/junos-evo-install-qfx-ms-
x86-64-22.2-202203012330.0-EVO.iso
Validating software image and getting ISSU services impact /var/home/regress/junos-evo-install-
qfx-ms-x86-64-22.1-202203012330.0-EVO.iso...
Download and Validate in Progress
re0: Starting validation : /var/home/regress/junos-evo-install-qfx-ms-x86-64-22.1-202203012330.0-
```

```
EVO.iso
re0: Validating version : junos-evo-install-qfx-ms-x86-64-22.1-202203012330.0-EVO.iso
re0: Running pre-checks for 'junos-evo-install-qfx-ms-x86-64-22.1-202203012330.0-EVO.iso'
re0: Pre-checks pass successfully, copying files to software area
re0: Running post install commands...
re0: Post install sequence was successful.
re0: Validating existing configs. See /var/log/validation_config.log for config validation logs.
re0: Validation Passed
re0: Validating in-service-upgrade application configs. See /var/log/validation_appconfig.log
for in-service-upgrade application configs validation logs.
re0: In-service-upgrade application configs validation Passed
re0: Validating in-service-upgrade SDK compatibility. See /var/log/validation_sdk.log for in-
service-upgrade SDK compatibility validation logs.
re0: In-service-upgrade SDK Validation Passed
re0: Generating local impact report...
re0: Installation was successful
Image validation succeeded. ISSU impact report:

*** Restart Apps list ***
distributord

*** Applications that do not support restart upgrade ***
distributord

This platform supports in-service kernel warm restart upgrade.
Validate cleanup succeeded.
Image validation succeeded. Reboot is needed for this software image upgrade.
```

Release Information

Command introduced in Junos OS Evolved Release 22.1R1.

RELATED DOCUMENTATION

request system software add restart | 15242

request system storage cleanup (Junos OS)

IN THIS SECTION

- Syntax | 15374
- Syntax (EX Series Switches) | 15374
- Syntax (MX Series Router) | 15375
- Syntax (QFX Series) | 15375
- Syntax (SRX Series) | 15375
- Description | 15376
- Options | 15376
- Additional Information | 15377
- Required Privilege Level | 15377
- Output Fields | 15377
- Sample Output | 15378
- Release Information | 15388

Syntax

```
request system storage cleanup
<dry-run>
<no-confirm>
<re0  | re1 | routing-engine (backup | both | local | master | other)>
```

Syntax (EX Series Switches)

```
request system storage cleanup
<all-members>
<dry-run>
<local>
<member member-id>
<no-confirm>
```

```
<re0  | re1 | routing-engine (backup | both | local | master | other)>
<satellite [slot-id slot-id |device-alias  alias-name]>
```

Syntax (MX Series Router)

```
request system storage cleanup
<all-members>
<dry-run>
<local>
<member member-id>
<no-confirm>
<re0  | re1 | routing-engine (backup | both | local | master | other)>
<satellite [slot-id slot-id |device-alias  alias-name]>
```

Syntax (QFX Series)

```
request system storage cleanup
<component (serial number | UUID | all)>
<director-group name>
<dry-run>
<infrastructure name>
<interconnect-device name>
<name-tag name-tag>
<no-confirm>
<node-group name>
<prune>
<qfabric (component name) | dry-run | name-tag | repository)>
<repository (core | log)>
<re0  | re1 | routing-engine (backup | both | local | master | other)>
```

Syntax (SRX Series)

```
request system storage cleanup
<dry-run>
<no-confirm>
<re0  | re1 | routing-engine (backup | both | local | master | other)>
```

Description

Use this command to free storage space on the router or switch by rotating log files and proposing a list of files for deletion. On a QFabric system, you can delete debug files located on individual devices or on the entire QFabric system.

Options

all-members	(EX4200 switches and MX Series routers only) (Optional) Delete files on the Virtual Chassis primary Routing Engine only. To delete files on the other members of the Virtual Chassis configuration, log in to each backup Routing Engine and delete the files using the `request system storage cleanup local` command.
component (*UUID* | *serial number* | all)	(QFabric systems only) (Optional) Delete files located on individual QFabric system devices or on the entire QFabric system.
director-group *name*	(QFabric systems only) (Optional) Delete files on the director group.
dry-run	(Optional) List files proposed for deletion (without deleting them).
infrastructure *name*	(QFabric systems only) (Optional) Delete files on the fabric control Routing Engine and fabric manager Routing Engine.
interconnect-device *name*	(QFabric systems only) (Optional) Delete files on the Interconnect device.
local	(EX4200 switches and MX Series routers only) (Optional) Delete files on the local Virtual Chassis member.
member *member-id*	(EX4200 switches and MX Series routers only) (Optional) Delete files on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.
name-tag *name-tag*	(QFabric systems only) (Optional) Delete debug files that match a specific regular expression.
node-group *name*	(QFabric systems only) (Optional) Delete files on the Node group.
no-confirm	(Optional) Do not ask for confirmation before doing the cleanup.
prune	(QFabric systems only) (Optional) Delete debug files located in either the core or log debug repositories of a QFabric system device.

qfabric component *name*

(QFabric systems only) (Optional) Delete debug files located in the debug repositories of a QFabric system device.

(re0 | re1 | routing-engine (backup | both | local | master | other))

(Optional) Request operation on system storage on RE0, RE1, or on specified Routing Engine by these classifications: backup, both, local, primary, or other.

When Routing Engine is specified, the below message is shown before listing the files and deleting them.

```
Please check the list of files to be deleted using the dry-run option.
i.e.
request system storage cleanup dry-run
Do you want to proceed ? [yes,no] (no)
```

repository (core | log)

(QFabric systems only) (Optional) Specify the repository on the QFabric system device for which you want to delete debug files.

satellite [slot-id *slot-id* | device-alias *alias-name*]

(Junos Fusion only) (Optional) Specify the satellite device in the Junos Fusion by FPC ID or device alias name for which you want to delete debug files.

Additional Information

If logging is configured and being used, the `dry-run` option rotates the log files. In that case, the output displays the message “Currently rotating log files, please wait.” If no logging is currently under way, the output displays only a list of files to delete.

Required Privilege Level

maintenance

Output Fields

Table 130 on page 15378 describes the output fields for the `request system storage cleanup` command. Output fields are listed in the approximate order in which they appear.

Table 130: request system storage cleanup Output Fields

Field Name	Field Description
`List of files to delete:`	Shows list of files available for deletion.
`Size`	Size of the core-dump file.
`Date`	Last core-dump file modification date and time.
`Name`	Name of the core-dump file.
`Directory to delete:`	Shows list of directories available for deletion.
`Repository scope:`	Repository where core-dump files and log files are stored. The core-dump files are located in the `core` repository, and the log files are located in the `log` repository. The default `Repository scope` is shared since both the `core` and `log` repositories are shared by all of the QFabric system devices.
`Repository head:`	Name of the top-level repository location.
`Repository name:`	Name of the repository: `core` or `log`.
`Creating list of debug artifacts to be removed under:`	Shows location of files available for deletion.
`List of debug artifacts to be removed under:`	Shows list of files available for deletion.

Sample Output

request system storage cleanup dry-run

```
user@host> request system storage cleanup dry-run
Currently rotating log files, please wait.
```

```
This operation can take up to a minute.

List of files to delete:

         Size Date         Name
  11.4K Mar  8 15:00 /var/log/messages.1.gz
  7245B Feb  5 15:00 /var/log/messages.3.gz
  11.8K Feb 22 13:00 /var/log/messages.2.gz
  3926B Mar 16 13:57 /var/log/messages.0.gz
  3962B Feb 22 12:47 /var/log/sampled.1.gz
  4146B Mar  8 12:20 /var/log/sampled.0.gz
  4708B Dec 21 11:39 /var/log/sampled.2.gz
  7068B Jan 16 18:00 /var/log/messages.4.gz
  13.7K Dec 27 22:00 /var/log/messages.5.gz
   890B Feb 22 17:22 /var/tmp/sampled.pkts
  65.8M Oct 26 09:10 /var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz
  63.1M Oct 26 09:13 /var/sw/pkg/jbundle-7.4R1.7.tgz
```

request system storage cleanup

```
user@host> request system storage cleanup
Currently rotating log files, please wait.
This operation can take up to a minute.

List of files to delete:

         Size Date         Name
  11.4K Mar  8 15:00 /var/log/messages.1.gz
  7245B Feb  5 15:00 /var/log/messages.3.gz
  11.8K Feb 22 13:00 /var/log/messages.2.gz
  3926B Mar 16 13:57 /var/log/messages.0.gz
  11.6K Mar  8 15:00 /var/log/messages.5.gz
  7254B Feb  5 15:00 /var/log/messages.6.gz
  12.9K Feb 22 13:00 /var/log/messages.8.gz
  3726B Mar 16 13:57 /var/log/messages.7.gz
  3962B Feb 22 12:47 /var/log/sampled.1.gz
  4146B Mar  8 12:20 /var/log/sampled.0.gz
  4708B Dec 21 11:39 /var/log/sampled.2.gz
  7068B Jan 16 18:00 /var/log/messages.4.gz
  13.7K Dec 27 22:00 /var/log/messages.5.gz
   890B Feb 22 17:22 /var/tmp/sampled.pkts
```

```
  65.8M Oct 26 09:10 /var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz
  63.1M Oct 26 09:13 /var/sw/pkg/jbundle-7.4R1.7.tgz

Delete these files ? [yes,no] (yes)
```

request system storage cleanup director-group (QFabric Systems)

```
user@switch> request system storage cleanup director-group
List of files to delete:

         Size Date         Name
4.0K   2011-11-07 05:16:29  /tmp/2064.sfcauth
4.0K   2011-11-07 05:07:34  /tmp/30804.sfcauth
4.0K   2011-11-07 04:13:41  /tmp/26792.sfcauth
4.0K   2011-11-07 04:13:39  /tmp/26432.sfcauth
0      2011-11-07 07:45:40  /tmp/cluster_cleanup.log
1.3M   2011-11-07 07:39:11  /tmp/cn_monitor.20111107-052401.log
4.0K   2011-11-07 07:36:29  /tmp/clustat.28019.log
4.0K   2011-11-07 07:36:29  /tmp/clustat_x.28019.log
9.6M   2011-11-07 05:30:24  /tmp/sfc.2.log
4.0K   2011-11-07 05:28:11  /tmp/mgd-init.1320672491.log
248K   2011-11-07 05:19:24  /tmp/cn_monitor.20111107-045111.log
4.0K   2011-11-07 05:17:18  /tmp/clustat.3401.log
4.0K   2011-11-07 05:17:18  /tmp/clustat_x.3401.log
8.0K   2011-11-07 04:58:25  /tmp/mgd-init.1320670633.log
0      2011-11-07 04:54:01  /tmp/mysql_db_install_5.1.37.log
4.0K   2011-11-07 04:52:08  /tmp/cn_send.log
0      2011-11-07 04:52:00  /tmp/init_eth0.log
4.0K   2011-11-07 04:49:35  /tmp/install_interfaces.sh.log
4.0K   2011-11-07 04:48:15  /tmp/bootstrap.sh.log
160K   2011-11-07 04:47:43  /tmp/bootstrap_cleanup.log
38M    2011-11-07 04:42:42  /tmp/cn_monitor.20111104-110308.log
4.0K   2011-11-07 04:38:47  /tmp/clustat.30913.log
4.0K   2011-11-07 04:38:47  /tmp/clustat_x.30913.log
4.0K   2011-11-07 04:38:03  /tmp/dcf_upgrade.sh.remove.log
4.0K   2011-11-07 04:38:03  /tmp/peer_update.log
4.0K   2011-11-07 04:38:02  /tmp/dcf_upgrade.log
4.0K   2011-11-07 04:38:02  /tmp/perl_mark_upgrade.log
8.0K   2011-11-07 04:13:42  /tmp/install_dcf_rpm.log
4.0K   2011-11-07 04:13:06  /tmp/00_cleanup.sh.1320667986.log
0      2011-11-07 04:13:06  /tmp/ccif_patch_4410_4450.sh.1320667986.log
```

```
4.0K   2011-11-07 04:13:06  /tmp/dcf-tools.sh.1320667986.log
0      2011-11-07 04:13:06  /tmp/initial.sh.1320667986.log
0      2011-11-07 04:13:06  /tmp/inventory.sh.1320667986.log
4.0K   2011-11-07 04:13:06  /tmp/qf-db.sh.1320667986.log
4.0K   2011-11-07 04:13:06  /tmp/sfc.sh.1320667986.log
8.0K   2011-11-07 04:13:05  /tmp/jinstall-qfabric.log
8.0K   2011-11-04 11:10:24  /tmp/mgd-init.1320430192.log
4.0K   2011-11-04 11:07:03  /tmp/mysql_dcf_db_install.log
8.0K   2011-11-04 10:55:07  /tmp/ccif_patch_4410_4450.sh.1320429307.log
8.0K   2011-11-04 10:55:07  /tmp/initial.sh.1320429307.log
4.0K   2011-11-04 10:55:07  /tmp/inventory.sh.1320429307.log
8.0K   2011-11-04 10:55:07  /tmp/sfc.sh.1320429307.log
4.0K   2011-11-04 10:54:09  /tmp/ks-script-Ax0tz5.log
4.0K   2011-11-07 04:13:06  /tmp//sfc.sh.1320667986.log
8.0K   2011-11-04 10:55:07  /tmp//sfc.sh.1320429307.log

Directory to delete:
45M    2011-11-08 10:57:43  /tmp/sfc-captures

List of files to delete:

         Size Date         Name
4.0K   2011-11-08 05:47:47  /tmp/5713.sfcauth
4.0K   2011-11-08 05:14:32  /tmp/14494.sfcauth
4.0K   2011-11-08 05:11:47  /tmp/9978.sfcauth
4.0K   2011-11-08 05:09:37  /tmp/6128.sfcauth
4.0K   2011-11-08 05:04:28  /tmp/29703.sfcauth
4.0K   2011-11-07 11:59:10  /tmp/7811.sfcauth
4.0K   2011-11-07 11:36:08  /tmp/32415.sfcauth
4.0K   2011-11-07 11:30:30  /tmp/22406.sfcauth
4.0K   2011-11-07 11:24:37  /tmp/12131.sfcauth
4.0K   2011-11-07 10:48:42  /tmp/12687.sfcauth
4.0K   2011-11-07 09:27:20  /tmp/31082.sfcauth
4.0K   2011-11-07 07:33:58  /tmp/14633.sfcauth
4.0K   2011-11-07 05:08:25  /tmp/15447.sfcauth
4.0K   2011-11-07 04:12:29  /tmp/26874.sfcauth
4.0K   2011-11-07 04:12:27  /tmp/26713.sfcauth
4.0K   2011-11-07 03:49:17  /tmp/17691.sfcauth
4.0K   2011-11-05 01:32:23  /tmp/5716.sfcauth
4.0K   2011-11-07 08:00:17  /tmp/sfcsnmpd.log
4.0K   2011-11-07 07:57:50  /tmp/cluster_cleanup.log
824K   2011-11-07 07:38:37  /tmp/cn_monitor.20111107-053643.log
4.0K   2011-11-07 07:36:30  /tmp/clustat.18399.log
```

```
4.0K   2011-11-07 07:36:30  /tmp/clustat_x.18399.log
4.0K   2011-11-07 07:35:47  /tmp/command_lock.log
4.0K   2011-11-07 05:39:54  /tmp/mgd-init.1320673194.log
92K    2011-11-07 05:19:25  /tmp/cn_monitor.20111107-050412.log
4.0K   2011-11-07 05:17:20  /tmp/clustat.30115.log
4.0K   2011-11-07 05:17:20  /tmp/clustat_x.30115.log
8.0K   2011-11-07 05:08:07  /tmp/mgd-init.1320671241.log
4.0K   2011-11-07 05:04:57  /tmp/cn_send.log
0      2011-11-07 05:04:52  /tmp/init_eth0.log
4.0K   2011-11-07 05:02:38  /tmp/install_interfaces.sh.log
4.0K   2011-11-07 05:01:19  /tmp/bootstrap.sh.log
160K   2011-11-07 05:00:47  /tmp/bootstrap_cleanup.log
28M    2011-11-07 04:42:27  /tmp/cn_monitor.20111104-112954.log
4.0K   2011-11-07 04:38:49  /tmp/clustat.6780.log
4.0K   2011-11-07 04:38:49  /tmp/clustat_x.6780.log
4.0K   2011-11-07 04:38:05  /tmp/issue_event.log
4.0K   2011-11-07 04:38:05  /tmp/peer_upgrade_reboot.log
12K    2011-11-07 04:38:05  /tmp/primary_update.log
4.0K   2011-11-07 04:38:04  /tmp/dcf_upgrade.sh.remove.log
4.0K   2011-11-07 04:38:04  /tmp/peer_rexec_upgrade.log
4.0K   2011-11-07 04:13:42  /tmp/peer_install_dcf_rpm.log
4.0K   2011-11-07 04:11:57  /tmp/dcf-tools.sh.1320667917.log
0      2011-11-07 04:11:57  /tmp/initial.sh.1320667917.log
0      2011-11-07 04:11:57  /tmp/inventory.sh.1320667917.log
4.0K   2011-11-07 04:11:57  /tmp/qf-db.sh.1320667917.log
4.0K   2011-11-07 04:11:57  /tmp/sfc.sh.1320667917.log
4.0K   2011-11-07 04:11:56  /tmp/00_cleanup.sh.1320667916.log
0      2011-11-07 04:11:56  /tmp/ccif_patch_4410_4450.sh.1320667916.log
8.0K   2011-11-07 04:11:56  /tmp/jinstall-qfabric.log
4.0K   2011-11-07 04:11:33  /tmp/dcf_upgrade.log
8.0K   2011-11-04 11:53:12  /tmp/mgd-init.1320432782.log
8.0K   2011-11-04 11:06:17  /tmp/ccif_patch_4410_4450.sh.1320429977.log
8.0K   2011-11-04 11:06:17  /tmp/initial.sh.1320429977.log
4.0K   2011-11-04 11:06:17  /tmp/inventory.sh.1320429977.log
8.0K   2011-11-04 11:06:17  /tmp/sfc.sh.1320429977.log
4.0K   2011-11-04 11:05:19  /tmp/ks-script-_tnWeb.log
4.0K   2011-11-07 04:11:57  /tmp//sfc.sh.1320667917.log
8.0K   2011-11-04 11:06:17  /tmp//sfc.sh.1320429977.log

Directory to delete:
49M    2011-11-08 10:45:20  /tmp/sfc-captures
```

request system storage cleanup infrastructure device-name (QFabric Systems)

```
user@switch> request system storage cleanup infrastructure FC
re0:
--------------------------------------------------------------------------

List of files to delete:

         Size Date         Name
   139B Nov  8 19:03 /var/log/default-log-messages.0.gz
  5602B Nov  8 19:03 /var/log/messages.0.gz
  28.4K Nov  8 10:15 /var/log/messages.1.gz
  35.2K Nov  7 13:45 /var/log/messages.2.gz
   207B Nov  7 16:02 /var/log/wtmp.0.gz
    27B Nov  7 12:14 /var/log/wtmp.1.gz
 184.4M Nov  7 12:16 /var/sw/pkg/jinstall-dc-re-11.3I20111104_1216_dc-builder-domestic-signed.tgz
 124.0K Nov  7 15:59 /var/tmp/gres-tp/env.dat
     0B Nov  7 12:57 /var/tmp/gres-tp/lock
   155B Nov  7 16:02 /var/tmp/krt_gencfg_filter.txt
     0B Nov  7 12:35 /var/tmp/last_ccif_update
  1217B Nov  7 12:15 /var/tmp/loader.conf.preinstall
 184.4M Nov  6 07:11 /var/tmp/mchassis-install.tgz
  10.8M Nov  7 12:16 /var/tmp/preinstall/bootstrap-install-11.3I20111104_1216_dc-builder.tar
  57.4K Nov  7 12:16 /var/tmp/preinstall/configs-11.3I20111104_1216_dc-builder.tgz
   259B Nov  7 12:16 /var/tmp/preinstall/install.conf
 734.3K Nov  4 13:46 /var/tmp/preinstall/jboot-dc-re-11.3I20111104_1216_dc-builder.tgz
 177.8M Nov  7 12:16 /var/tmp/preinstall/jbundle-dc-re-11.3I20111104_1216_dc-builder-domestic.tgz
   124B Nov  7 12:15 /var/tmp/preinstall/metatags
  1217B Nov  7 12:16 /var/tmp/preinstall_boot_loader.conf
     0B Nov  7 16:02 /var/tmp/rtsdb/if-rtsdb
```

request system storage cleanup interconnect-device device-name (QFabric Systems)

```
user@switch> request system storage cleanup interconnect IC
re1:
--------------------------------------------------------------------------

List of files to delete:

         Size Date         Name
```

```
    11B Nov  7 15:55 /var/jail/tmp/alarmd.ts
   128B Nov  8 19:06 /var/log/default-log-messages.0.gz
  9965B Nov  8 19:06 /var/log/messages.0.gz
  15.8K Nov  8 12:30 /var/log/messages.1.gz
  15.8K Nov  8 11:00 /var/log/messages.2.gz
  15.7K Nov  8 07:30 /var/log/messages.3.gz
  15.8K Nov  8 04:00 /var/log/messages.4.gz
  15.7K Nov  8 00:30 /var/log/messages.5.gz
  18.7K Nov  7 21:00 /var/log/messages.6.gz
  17.6K Nov  7 19:00 /var/log/messages.7.gz
  58.3K Nov  7 16:00 /var/log/messages.8.gz
  20.3K Nov  7 15:15 /var/log/messages.9.gz
    90B Nov  7 15:41 /var/log/wtmp.0.gz
    57B Nov  7 12:41 /var/log/wtmp.1.gz
 124.0K Nov  7 15:42 /var/tmp/gres-tp/env.dat
     0B Nov  7 12:40 /var/tmp/gres-tp/lock
     0B Nov  7 12:41 /var/tmp/if-rtsdb/env.lck
  12.0K Nov  7 15:41 /var/tmp/if-rtsdb/env.mem
 132.0K Nov  7 15:55 /var/tmp/if-rtsdb/shm_usr1.mem
2688.0K Nov  7 15:41 /var/tmp/if-rtsdb/shm_usr2.mem
2048.0K Nov  7 15:41 /var/tmp/if-rtsdb/trace.mem
   730B Nov  7 19:57 /var/tmp/juniper.conf+.gz
   155B Nov  7 15:53 /var/tmp/krt_gencfg_filter.txt
     0B Nov  7 15:41 /var/tmp/rtsdb/if-rtsdb

re0:
--------------------------------------------------------------------------

List of files to delete:

         Size Date         Name
    11B Nov  7 15:55 /var/jail/tmp/alarmd.ts
   121B Nov  8 19:06 /var/log/default-log-messages.0.gz
  16.7K Nov  8 19:06 /var/log/messages.0.gz
  22.2K Nov  8 17:45 /var/log/messages.1.gz
  K Nov  8 17:00 /var/log/messages.2.gz
  21.6K Nov  8 16:00 /var/log/messages.3.gz
  17.9K Nov  8 14:30 /var/log/messages.4.gz
  19.4K Nov  8 13:30 /var/log/messages.5.gz
  18.2K Nov  8 12:30 /var/log/messages.6.gz
  20.4K Nov  8 11:30 /var/log/messages.7.gz
  21.4K Nov  8 10:15 /var/log/messages.8.gz
  21.0K Nov  8 09:00 /var/log/messages.9.gz
```

```
  19.9K Nov  8 08:13 /var/log/snmp-traps.0.gz
   203B Nov  8 15:36 /var/log/wtmp.0.gz
    57B Nov  7 12:41 /var/log/wtmp.1.gz
 124.0K Nov  7 15:42 /var/tmp/gres-tp/env.dat
     0B Nov  7 12:40 /var/tmp/gres-tp/lock
     0B Nov  7 12:41 /var/tmp/if-rtsdb/env.lck
  12.0K Nov  7 15:41 /var/tmp/if-rtsdb/env.mem
 132.0K Nov  7 15:55 /var/tmp/if-rtsdb/shm_usr1.mem
2688.0K Nov  7 15:41 /var/tmp/if-rtsdb/shm_usr2.mem
2048.0K Nov  7 15:41 /var/tmp/if-rtsdb/trace.mem
   727B Nov  7 15:54 /var/tmp/juniper.conf+.gz
   155B Nov  7 15:55 /var/tmp/krt_gencfg_filter.txt
     0B Nov  7 15:41 /var/tmp/rtsdb/if-rtsdb
```

request system storage cleanup node-group group-name (QFabric Systems)

```
user@switch> request system storage cleanup node-group NW-NG
BBAK0372:
--------------------------------------------------------------------------

List of files to delete:

         Size Date         Name
   126B Nov  8 19:07 /var/log/default-log-messages.0.gz
   179B Nov  7 13:32 /var/log/install.0.gz
  22.9K Nov  8 19:07 /var/log/messages.0.gz
  26.5K Nov  8 17:30 /var/log/messages.1.gz
  20.5K Nov  8 13:15 /var/log/messages.2.gz
  33.2K Nov  7 17:45 /var/log/messages.3.gz
  35.5K Nov  7 15:45 /var/log/messages.4.gz
   339B Nov  8 17:10 /var/log/wtmp.0.gz
    58B Nov  7 12:40 /var/log/wtmp.1.gz
 124.0K Nov  8 17:08 /var/tmp/gres-tp/env.dat
     0B Nov  7 12:39 /var/tmp/gres-tp/lock
     0B Nov  7 12:59 /var/tmp/if-rtsdb/env.lck
  12.0K Nov  8 17:09 /var/tmp/if-rtsdb/env.mem
2688.0K Nov  8 17:09 /var/tmp/if-rtsdb/shm_usr1.mem
 132.0K Nov  8 17:09 /var/tmp/if-rtsdb/shm_usr2.mem
2048.0K Nov  8 17:09 /var/tmp/if-rtsdb/trace.mem
  1082B Nov  8 17:09 /var/tmp/juniper.conf+.gz
   155B Nov  7 17:39 /var/tmp/krt_gencfg_filter.txt
```

```
     0B Nov  8 17:09 /var/tmp/rtsdb/if-rtsdb

EE3093:
---------------------------------------------------------------------------

List of files to delete:

         Size Date         Name
    11B Nov  8 17:33 /var/jail/tmp/alarmd.ts
   119B Nov  8 19:08 /var/log/default-log-messages.0.gz
   180B Nov  7 17:41 /var/log/install.0.gz
   178B Nov  7 13:32 /var/log/install.1.gz
  2739B Nov  8 19:08 /var/log/messages.0.gz
  29.8K Nov  8 18:45 /var/log/messages.1.gz
  31.8K Nov  8 17:15 /var/log/messages.2.gz
  20.6K Nov  8 16:00 /var/log/messages.3.gz
  15.4K Nov  8 10:15 /var/log/messages.4.gz
  15.4K Nov  8 02:15 /var/log/messages.5.gz
  25.5K Nov  7 20:45 /var/log/messages.6.gz
  48.0K Nov  7 17:45 /var/log/messages.7.gz
  32.8K Nov  7 13:45 /var/log/messages.8.gz
   684B Nov  8 17:02 /var/log/wtmp.0.gz
    58B Nov  7 12:40 /var/log/wtmp.1.gz
 124.0K Nov  7 17:34 /var/tmp/gres-tp/env.dat
     0B Nov  7 12:40 /var/tmp/gres-tp/lock
     0B Nov  7 12:59 /var/tmp/if-rtsdb/env.lck
  12.0K Nov  7 17:39 /var/tmp/if-rtsdb/env.mem
2688.0K Nov  7 17:39 /var/tmp/if-rtsdb/shm_usr1.mem
 132.0K Nov  7 17:40 /var/tmp/if-rtsdb/shm_usr2.mem
2048.0K Nov  7 17:39 /var/tmp/if-rtsdb/trace.mem
   155B Nov  7 17:40 /var/tmp/krt_gencfg_filter.txt
     0B Nov  7 17:39 /var/tmp/rtsdb/if-rtsdb
```

request system storage cleanup qfabric component device-name (QFabric Systems)

```
user@switch> request system storage cleanup qfabric component Test
Repository type: regular
Repository head: /pbstorage
Creating list of debug artifacts to be removed under: /pbstorage/rdumps/Test
Removing debug artifacts ... (press control C to abort)
Removing /pbstorage/rdumps/Test/cosd.core.0.0.05162011123308.gz ... done
```

```
Removing /pbstorage/rdumps/Test/cosd.core.1.0.05162011123614.gz ... done
Removing /pbstorage/rdumps/Test/cosd.core.2.0.05162011123920.gz ... done
Removing /pbstorage/rdumps/Test/livekcore.05132011163930.gz ... done
Removing /pbstorage/rdumps/Test/tnetd.core.0.1057.05162011124500.gz ... done
Removing /pbstorage/rdumps/Test/vmcore.05132011120528.gz ... done
Removing /pbstorage/rdumps/Test/vmcore.kz ... done
Creating list of debug artifacts to be removed under: /pbstorage/rlogs/Test
Removing debug artifacts ... (press control C to abort)
Removing /pbstorage/rlogs/Test/kdumpinfo.05132011120528 ... done
Removing /pbstorage/rlogs/Test/kernel.tarball.0.1039.05122011234415.tgz ... done
Removing /pbstorage/rlogs/Test/kernel.tarball.1.1039.05132011175544.tgz ... done
Removing /pbstorage/rlogs/Test/tnetd.tarball.0.1057.05162011175453.tgz ... done
```

request system storage cleanup qfabric component device-name repository core (QFabric Systems)

```
user@switch> request system storage cleanup qfabric component Test repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
Creating list of debug artifacts to be removed under: /pbdata/export/rdumps/Test
NOTE: core repository under /pbdata/export/rdumps/Test empty
```

request system storage cleanup qfabric component all (QFabric Systems)

```
user@switch> request system storage cleanup qfabric component all
Repository scope: shared
Repository head: /pbdata/export
Creating list of debug artifacts to be removed under: /pbdata/export/rdumps
NOTE: core repository under /pbdata/export/rdumps/all empty
Creating list of debug artifacts to be removed under: /pbdata/export/rlogs
List of debug artifacts to clean up ... (press control C to abort)
/pbdata/export/rlogs/73747cd8-0710-11e1-b6a4-00e081c5297e/install-11072011125819.log
/pbdata/export/rlogs/77116f18-0710-11e1-a2a0-00e081c5297e/install-11072011125819.log
/pbdata/export/rlogs/BBAK0372/install-11072011121538.log
/pbdata/export/rlogs/BBAK0394/install-11072011121532.log
/pbdata/export/rlogs/EE3093/install-11072011121536.log
/pbdata/export/rlogs/WS001/YN5999/install-11072011121644.log
/pbdata/export/rlogs/WS001/YW3803/install-11072011122429.log
```

```
/pbdata/export/rlogs/cd78871a-0710-11e1-878e-00e081c5297e/install-11072011125932.log
/pbdata/export/rlogs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011125930.log
/pbdata/export/rlogs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011133211.log
/pbdata/export/rlogs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011155302.log
/pbdata/export/rlogs/d31ab7a6-0710-11e1-ad1b-00e081c5297e/install-11072011125931.log
/pbdata/export/rlogs/d4d0f254-0710-11e1-90c3-00e081c5297e/install-11072011125932.log
```

Release Information

Command introduced in Junos OS Release 7.4.

`dry-run` option introduced in Junos OS Release 7.6.

`satellite` option introduced in Junos OS Release 14.2R3.

`no-confirm` and `(re0 |re1 | routing-engine (backup | both | local | master | other))` options introduced in Junos OS 17.3R1.

request system storage cleanup (Junos OS Evolved)

IN THIS SECTION

- Syntax | 15388
- Description | 15389
- Options | 15389
- Additional Information | 15390
- Required Privilege Level | 15390
- Sample Output | 15390
- Release Information | 15399

Syntax

```
request system storage cleanup (dry-run | force-deep | no-confirm)
```

Description

Use this command to free storage space on the router or switch by rotating log files and proposing a list of files for deletion.

The Junos OS Evolved implementation of the `request system storage cleanup` command is slightly different from the implementation on Junos OS:

- The system prompts you to specify the `dry-run` option:

```
Please check the list of files to be deleted using the dry-run option.
Continue anyway without checking? [yes,no] (yes)
```

- When you issue the `request system storage cleanup` command, Junos OS Evolved displays the types of files to be deleted. See the Sample Output section below for an example.
- Prior to Junos OS Evolved Release 20.1R1, the command cleans up any ISO files on the system, rotates system log files, and clears trace files. It does not remove user-created files. Starting in Junos OS Evolved Release 20.1R1, this command does not remove ISO images from the system. It removes all core files, log files from **/var/log/**, and all **/var/log/*** files. To remove old images from the device, use the `request system software delete` command.
- In Junos OS Evolved, the system computes the available space and emits o/p on console for reference.

In Junos OS Evolved, the `request system storage cleanup | display xml rpc` command displays different XML tags for different file types. In Junos OS, the command displays only the `file` tag for all types of files. For more information about the differences between Junos OS and Junos OS Evolved, see How Junos OS Evolved Differs from Junos OS.

Options

dry-run (Prompted if not specified) List the files proposed for deletion (without deleting the files).

force-deep (Optional) Deep clean all temporary files and rotate logs. This option cleans up all the user-created files under the **/tmp** and **/var/tmp** directories.

no-confirm (Optional) Do not ask for confirmation before doing the cleanup.

Additional Information

If logging is configured and being used, the `dry-run` option rotates the log files. In that case, the output displays the message “Currently rotating log files, please wait”. If no logging is currently under way, the output displays only a list of files to delete.

Required Privilege Level

maintenance

Sample Output

request system storage cleanup (Junos OS Evolved)

```
user@host-re0> request system storage cleanup
Please check the list of files to be deleted using the dry-run option.
Continue anyway without checking? [yes,no] (no) yes
-------------------------------
  node:  fpc0
-------------------------------
  Clearing all core files


  Clearing all local host core files and files from /var/log/watchdog


  Clearing node specific core files


  Clearing FPC log  files


  Clearing logical-systems log files


  Clearing journal logs


  Clearing all /var/log/* files

  Size      Date              Name
```

```
  4.0K  Tue Feb  2 13:32  /var/log/beacon_getconfid.log
  4.0K  Tue Feb  2 13:32  /var/log/beacon_set_dhcpd.log
  4.0K  Tue Feb  2 13:31  /var/log/boot_init.log
  556K  Tue Feb  2 13:37  /var/log/evo-cda-bt.log
  8.0K  Tue Feb  2 13:32  /var/log/evoinit.log
     0  Tue Feb  2 13:31  /var/log/ftp_fail
  4.0K  Tue Feb  2 13:32  /var/log/imgd.log
     0  Tue Feb  2 13:31  /var/log/interactive-commands
     0  Tue Feb  2 13:33  /var/log/mcelog.log
  4.0K  Tue Feb  2 13:33  /var/log/mem_mgmt.log
     0  Tue Feb  2 13:32  /var/log/mem_monitor.dat
     0  Tue Feb  2 13:31  /var/log/messages
  4.0K  Tue Feb  2 13:32  /var/log/mstr.log
  4.0K  Tue Feb  2 13:32  /var/log/mstr_change.log
  292K  Tue Feb  2 13:37  /var/log/ofp-packet.log
  688K  Tue Feb  2 13:37  /var/log/ofp.log
  680K  Tue Feb  2 13:37  /var/log/picd.log
   40K  Tue Feb  2 13:32  /var/log/sinet.log
  4.0K  Tue Feb  2 13:32  /var/log/ss.log
  4.0K  Tue Feb  2 13:37  /var/log/storageCleanup.log
  4.0K  Tue Feb  2 13:32  /var/log/sysconfig.log
   60K  Tue Feb  2 13:32  /var/log/sysepochman.log
  4.0K  Tue Feb  2 13:32  /var/log/sysstart.log
   72K  Tue Feb  2 13:32  /var/log/uswitch.log
   68K  Tue Feb  2 13:31  /var/log/uswitch.log.prev
  4.0K  Tue Feb  2 13:32  /var/log/wtmp

  Clearing all JSON files in /var/log/objmon


  busy with previous request.


  Clearing SI traces


  Removing any ISO files in /data


  Current space available in /soft: 4233300 K and /data: 69765 K

-------------------------------
  node:  re0
```

```
-------------------------------
  Clearing all core files


  Clearing all local host core files and files from /var/log/watchdog


  Clearing node specific core files


  Clearing FPC log  files


  Clearing logical-systems log files


  Clearing journal logs


  Clearing all /var/log/* files

  Size      Date              Name
  4.0K  Tue Feb  2 13:28  /var/log/__policy_names_rpdc__
  4.0K  Tue Feb  2 13:32  /var/log/__policy_names_rpdn__
  8.0K  Tue Feb  2 13:32  /var/log/alarm-mgmtd
  4.0K  Tue Feb  2 13:28  /var/log/beacon_getconfid.log
   48K  Tue Feb  2 13:27  /var/log/beacon_mountVersion.log
  4.0K  Tue Feb  2 13:29  /var/log/beacon_set_dhcpd.log
  4.0K  Tue Feb  2 13:28  /var/log/boot_init.log
  4.0K  Tue Feb  2 13:31  /var/log/charonctl_trace.log
  4.0K  Tue Feb  2 13:27  /var/log/check_restore_recovery_bios.log
  4.0K  Tue Feb  2 13:32  /var/log/clksynced.log
   84K  Tue Feb  2 13:32  /var/log/configd-streamer.log
  4.0K  Tue Feb  2 13:31  /var/log/copy_other_re_keys.log
  4.0K  Tue Feb  2 13:37  /var/log/core_mgr.log
     0  Tue Feb  2 13:28  /var/log/cscript.log
  4.0K  Tue Feb  2 13:32  /var/log/ddosd.log
  4.0K  Tue Feb  2 13:31  /var/log/disk_mgmt
  4.0K  Tue Feb  2 13:31  /var/log/evo_dns_relay.log
  8.0K  Tue Feb  2 13:28  /var/log/evoinit.log
   12K  Tue Feb  2 13:36  /var/log/fibd-proxy.log
   12K  Tue Feb  2 13:31  /var/log/imgd.log
   32K  Tue Feb  2 13:33  /var/log/interactive-commands
```

```
 28K  Tue Feb  2 13:37  /var/log/kfirewall-agent.log
   0  Tue Feb  2 13:32  /var/log/mcelog.log
8.0K  Tue Feb  2 13:32  /var/log/mem_mgmt.log
   0  Tue Feb  2 13:31  /var/log/mem_monitor.dat
128K  Tue Feb  2 13:37  /var/log/messages
   0  Tue Feb  2 13:31  /var/log/mgd-api
4.0K  Tue Feb  2 13:28  /var/log/mgmt-ethd-helper.log
 84K  Tue Feb  2 13:33  /var/log/mib2d
 12K  Tue Feb  2 13:33  /var/log/mirrord.log
4.0K  Tue Feb  2 13:28  /var/log/mstr.log
4.0K  Tue Feb  2 13:28  /var/log/mstr_change.log
464K  Tue Feb  2 13:37  /var/log/ofp-packet.log
984K  Tue Feb  2 13:37  /var/log/ofp.log
 48K  Tue Feb  2 13:27  /var/log/postinstall.log
4.0K  Tue Feb  2 13:37  /var/log/security
4.0K  Tue Feb  2 13:28  /var/log/set_mgmt_mac.log
 68K  Tue Feb  2 13:28  /var/log/sinet.log
 16K  Tue Feb  2 13:33  /var/log/snmpd
4.0K  Tue Feb  2 13:31  /var/log/ss.log
4.0K  Tue Feb  2 13:31  /var/log/ssh-key-utils.log
4.0K  Tue Feb  2 13:37  /var/log/storageCleanup.log
4.0K  Tue Feb  2 13:31  /var/log/sync_other_re.log
4.0K  Tue Feb  2 13:28  /var/log/sysconfig.log
552K  Tue Feb  2 13:32  /var/log/sysepochman.log
4.0K  Tue Feb  2 13:28  /var/log/sysstart.log
4.0K  Tue Feb  2 13:33  /var/log/system-events
 84K  Tue Feb  2 13:28  /var/log/uswitch.log
 88K  Tue Feb  2 13:26  /var/log/uswitch.log.prev
8.0K  Tue Feb  2 13:31  /var/log/validator_debug.log
4.0K  Tue Feb  2 13:33  /var/log/vrf.log
8.0K  Tue Feb  2 13:36  /var/log/wtmp
4.0K  Tue Feb  2 13:32  /var/log/xferlog
364K  Tue Feb  2 13:37  /var/log/zookeeper--server-host-re0.log
4.0K  Tue Feb  2 13:29  /var/log/zookeeper--server-host-re0.out
4.0K  Tue Feb  2 13:33  /var/log/ztp.log

Clearing all JSON files in /var/log/objmon


Cleared traces for application all node all pid all


Clearing SI traces
```

```
  Removing any ISO files in /data


  Current space available in /soft: 14158432 K and /data: 2857732 K

-------------------------------
  node:  re1
-------------------------------
  Clearing all core files


  Clearing all local host core files and files from /var/log/watchdog


  Clearing node specific core files


  Clearing FPC log  files


  Clearing logical-systems log files


  Clearing journal logs


  Clearing all /var/log/* files

  Size      Date              Name
  4.0K  Tue Feb  2 13:32  /var/log/__policy_names_rpdc__
  4.0K  Tue Feb  2 13:31  /var/log/alarm-mgmtd
  4.0K  Tue Feb  2 13:29  /var/log/beacon_getconfid.log
   48K  Tue Feb  2 13:27  /var/log/beacon_mountVersion.log
  4.0K  Tue Feb  2 13:29  /var/log/beacon_set_dhcpd.log
  4.0K  Tue Feb  2 13:28  /var/log/boot_init.log
  4.0K  Tue Feb  2 13:31  /var/log/charonctl_trace.log
  4.0K  Tue Feb  2 13:28  /var/log/check_restore_recovery_bios.log
   12K  Tue Feb  2 13:32  /var/log/copy_other_re_keys.log
  4.0K  Tue Feb  2 13:32  /var/log/core_mgr.log
     0  Tue Feb  2 13:28  /var/log/cscript.log
  4.0K  Tue Feb  2 13:31  /var/log/disk_mgmt
```

```
  4.0K  Tue Feb  2 13:31  /var/log/evo_dns_relay.log
  8.0K  Tue Feb  2 13:31  /var/log/evoinit.log
   24K  Tue Feb  2 13:37  /var/log/fibd-proxy.log
  4.0K  Tue Feb  2 13:29  /var/log/imgd.log
     0  Tue Feb  2 13:28  /var/log/interactive-commands
   24K  Tue Feb  2 13:37  /var/log/kfirewall-agent.log
     0  Tue Feb  2 13:32  /var/log/mcelog.log
  4.0K  Tue Feb  2 13:31  /var/log/mem_mgmt.log
     0  Tue Feb  2 13:31  /var/log/mem_monitor.dat
     0  Tue Feb  2 13:32  /var/log/messages
  4.0K  Tue Feb  2 13:31  /var/log/mgmt-ethd-helper.log
  4.0K  Tue Feb  2 13:31  /var/log/mib2d
  4.0K  Tue Feb  2 13:28  /var/log/mstr.log
  4.0K  Tue Feb  2 13:28  /var/log/mstr_change.log
  444K  Tue Feb  2 13:37  /var/log/ofp-packet.log
 1016K  Tue Feb  2 13:37  /var/log/ofp.log
   48K  Tue Feb  2 13:28  /var/log/postinstall.log
     0  Tue Feb  2 13:32  /var/log/security
  4.0K  Tue Feb  2 13:28  /var/log/set_mgmt_mac.log
   68K  Tue Feb  2 13:28  /var/log/sinet.log
  4.0K  Tue Feb  2 13:31  /var/log/snmpd
  4.0K  Tue Feb  2 13:31  /var/log/ss.log
  8.0K  Tue Feb  2 13:32  /var/log/ssh-key-utils.log
  4.0K  Tue Feb  2 13:37  /var/log/storageCleanup.log
  4.0K  Tue Feb  2 13:32  /var/log/sync_other_re.log
  4.0K  Tue Feb  2 13:28  /var/log/sysconfig.log
  204K  Tue Feb  2 13:32  /var/log/sysepochman.log
  4.0K  Tue Feb  2 13:28  /var/log/sysstart.log
     0  Tue Feb  2 13:32  /var/log/system-events
   84K  Tue Feb  2 13:29  /var/log/uswitch.log
   88K  Tue Feb  2 13:26  /var/log/uswitch.log.prev
   12K  Tue Feb  2 13:32  /var/log/validator_debug.log
  4.0K  Tue Feb  2 13:33  /var/log/vrf.log
  8.0K  Tue Feb  2 13:33  /var/log/wtmp
  4.0K  Tue Feb  2 13:31  /var/log/xferlog
   88K  Tue Feb  2 13:37  /var/log/zookeeper--server-host-re1.log
  4.0K  Tue Feb  2 13:29  /var/log/zookeeper--server-host-re1.out

  Clearing all JSON files in /var/log/objmon



  Cleared traces for application all node all pid all
```

```
  Clearing SI traces


  Removing any ISO files in /data


  Current space available in /soft: 14158432 K and /data: 2775780 K


{master}
user@host-re0>
```

request system storage cleanup dry-run (Junos OS Evolved)

```
user@host> request system storage cleanup dry-run
  -------------------------------
  node:  re0
  -------------------------------
  List of all core files to be cleared:


  List of local_host core files to be cleared:


  List of core sub directory files to be cleared:


  List of log files from FPCs to be cleared:


  List of log files from logical systems to be cleared:


  Clears all App logs, App traces, App SI traces and App core files from /var/log/*, /var/log/
traces/*, /var/log/si_traces/* and /var/core/*


  Clears all JSON files in /var/log/objmon/ during cleanup
```

```
  List of ISO files from /data partition to be cleared:

  Current list of software versions installed
  Removes older software versions - Minimum two versions would be left around

Active boot device is primary : /dev/vda
List of installed version(s) :

    '-' running version
    '>' next boot version after upgrade/downgrade
    '<' rollback boot version

 -   junos-evo-install-ptx-fixed-x86-64-20.1-201911201458.0-EVO - [2019-11-21 11:23:21]
```

request system storage cleanup force-deep (Junos OS Evolved)

```
user@host> request system storage cleanup force-deep
Please check the list of files to be deleted using the dry-run option.
 Continue anyway without checking? [yes,no] (no) yes

-------------------------------
node: re0
-------------------------------
........
========= Start cleanup now ==========
=== Start removing other logs, traces, core files  ===
Clearing core files
Clearing FPC logs
Clearing logical-systems logs
=== Clearing journal logs ===
Clearing log: /var/log/RE_journal.log
Clearing log: /var/log/RE_journal_boot.log
Clearing log: /var/log/alarm-mgmtd
Clearing log: /var/log/appDemo_stdout
Clearing log: /var/log/charonctl_trace.log
Clearing log: /var/log/configd-streamer.log
Clearing log: /var/log/core_mgr.log
Clearing log: /var/log/cscript.log
Clearing log: /var/log/eth_linkmon.log
```

```
Clearing log: /var/log/evo-cda-zx.log
Clearing log: /var/log/evoinit.log
Clearing log: /var/log/fibd-proxy.log
Clearing log: /var/log/i2ctrace.log
Clearing log: /var/log/i2ctrace_spmb0.log
Clearing log: /var/log/i2ctrace_spmb1.log
Clearing log: /var/log/icmpd.log
Clearing log: /var/log/ifinfo.log
Clearing log: /var/log/imgd_svr.log
Clearing log: /var/log/install
Clearing log: /var/log/interactive-commands
Clearing log: /var/log/jsd
Clearing log: /var/log/lastlog
Clearing log: /var/log/mcelog.log
Clearing log: /var/log/messages
Clearing log: /var/log/mgd-api
Clearing log: /var/log/mgmt-ethd-helper.log
Clearing log: /var/log/mib2d
Clearing log: /var/log/na-grpcd
Clearing log: /var/log/objmon_sync.json
Clearing log: /var/log/packetio-cout.log
Clearing log: /var/log/picd.log
Clearing log: /var/log/platform_mon.log
Clearing log: /var/log/policerd.log
Clearing log: /var/log/postinstall.log
Clearing log: /var/log/ptp_fpga.log
Clearing log: /var/log/reboot_node.log
Clearing log: /var/log/rollback.log
Clearing log: /var/log/security
Clearing log: /var/log/semctl.log
Clearing log: /var/log/set_mgmt_mac.log
Clearing log: /var/log/shutdown_complete.log
Clearing log: /var/log/sinet.log
Clearing log: /var/log/smartd-attr-SFSA200GM3AA4TO_C_HC_636_JUN-000060139624B1000020.log
Clearing log: /var/log/smartd-attr-SFSA200GM3AA4TO_C_HC_636_JUN-000060139624B1000022.log
Clearing log: /var/log/snmpd
Clearing log: /var/log/ss.log
Clearing log: /var/log/ssh-key-utils.log
Clearing log: /var/log/sshd_lua.log
Clearing log: /var/log/sysconfig.log
Clearing log: /var/log/sysman.conf
Clearing log: /var/log/system-events
Clearing log: /var/log/upgrade_master.log
```

```
Clearing log: /var/log/uswitch.log
Clearing log: /var/log/uswitch.log.prev
Clearing log: /var/log/validator_debug.log
Clearing log: /var/log/wtmp
Clearing log: /var/log/zookeeper--server-re.log
Clearing log: /var/log/zookeeper--server-re.out
Clearing log: /var/log/ztp.log
=== Clearing all traces ===
=== Clearing SI traces ===
=== Removing other logs, traces, core files completed ===
=== Started removing any ISO files in /data
=== Removing any ISO files in /data completed
=== Start Software versions cleanup ===
Removing older software versions except current and rollback
=== Software versions cleanup completed ===
========= Cleanup done ==========
Current space available in /soft: 12372572 K
Current space available in /data: 2638752 K
Cannot delete junos-evo-install-qfx-fixed-x86-64-18.3I20180906130134_mkamil - It is the rollback
version
Cannot delete junos-evo-install-qfx-fixed-x86-64-18.3-20180906.3 - It is the current version
Removing version junos-evo-install-qfx-x86-64-16.2I20180516093649...
Done.
```

Release Information

Command introduced in Junos OS Evolved Release 20.4R2.

RELATED DOCUMENTATION

request system software delete (Junos OS Evolved) | **15254**

request system storage cleanup (SRX Series)

IN THIS SECTION

- Syntax | 15400
- Description | 15400
- Options | 15400
- Additional Information | 15400
- Required Privilege Level | 15401
- Output Fields | 15401
- Sample Output | 15401
- Release Information | 15404

Syntax

```
request system storage cleanup <dry-run>
```

Description

Use this command to free storage space on the device by rotating log files and proposing a list of files for deletion.

Options

dry-run (Optional) List files proposed for deletion (without deleting them).

Additional Information

If logging is configured and being used, the `dry-run` option rotates the log files. In that case, the output displays the message “Currently rotating log files, please wait.” If no logging is currently under way, the output displays only a list of files to delete.

Required Privilege Level

maintenance

Output Fields

Table 131 on page 15401 describes the output fields for the `request system storage cleanup` command. Output fields are listed in the approximate order in which they appear.

Table 131: request system storage cleanup Output Fields

Field Name	Field Description
`List of files to delete:`	Shows list of files available for deletion.
`Size`	Size of the core-dump file.
`Date`	Last core-dump file modification date and time.
`Name`	Name of the core-dump file.

Sample Output

request system storage cleanup dry-run

```
user@host> request system storage cleanup dry-run
List of files to delete:

         Size Date         Name
    11B Jul 14 22:51 /var/jail/tmp/alarmd.ts
  84.3K Jul 20 22:09 /var/log/chassisd.0.gz
  83.0K Jul 20 04:35 /var/log/chassisd.1.gz
  84.0K Jul 19 10:52 /var/log/chassisd.2.gz
  90.4K Jul 18 17:16 /var/log/chassisd.3.gz
  91.8K Jul 20 04:30 /var/log/hostlogs/auth.log.1.gz
  93.1K Jul 17 05:45 /var/log/hostlogs/auth.log.2.gz
  97.6K Jun  7 01:30 /var/log/hostlogs/auth.log.3.gz
  92.0K Apr 25 15:15 /var/log/hostlogs/auth.log.4.gz
```

```
 78.0K Jul 21 05:44 /var/log/hostlogs/daemon.log.1.gz
 78.6K Jul 21 02:59 /var/log/hostlogs/daemon.log.2.gz
 78.5K Jul 21 00:14 /var/log/hostlogs/daemon.log.3.gz
 78.8K Jul 20 21:30 /var/log/hostlogs/daemon.log.4.gz
 58.7K Jul 21 05:14 /var/log/hostlogs/debug.1.gz
 58.5K Jul 21 00:59 /var/log/hostlogs/debug.2.gz
 58.7K Jul 20 20:44 /var/log/hostlogs/debug.3.gz
 58.7K Jul 20 16:29 /var/log/hostlogs/debug.4.gz
166.9K Jul 13 00:33 /var/log/hostlogs/kern.log.1.gz
166.5K Jun  1 02:32 /var/log/hostlogs/kern.log.2.gz
163.5K May  5 00:03 /var/log/hostlogs/kern.log.3.gz
152.3K Mar  2 23:23 /var/log/hostlogs/kern.log.4.gz
260.0K Apr 13 10:28 /var/log/hostlogs/lcmd.log.1.gz
257.3K Mar  7 00:38 /var/log/hostlogs/lcmd.log.2.gz
240.8K Feb  7 19:45 /var/log/hostlogs/lcmd.log.3.gz
241.1K Feb  7 14:00 /var/log/hostlogs/lcmd.log.4.gz
370.6K Jul 21 00:45 /var/log/hostlogs/syslog.1.gz
370.9K Jul 20 12:30 /var/log/hostlogs/syslog.2.gz
370.4K Jul 20 00:15 /var/log/hostlogs/syslog.3.gz
370.2K Jul 19 12:00 /var/log/hostlogs/syslog.4.gz
 55.0K Jul 14 22:50 /var/log/hostlogs/vjunos0.log.1.gz
 1467B Oct 28  2015 /var/log/install.0.gz
119.9K Jul 21 07:37 /var/log/messages.0.gz
147.4K May 27 01:30 /var/log/messages.1.gz
 71.4K Apr 14 11:19 /var/log/messages.2.gz
 90.7K Feb 28 14:15 /var/log/messages.3.gz
 10.1K Jan 12  2016 /var/log/messages.4.gz
 55.1K Jan  6  2016 /var/log/messages.5.gz
 81.5K Dec  1  2015 /var/log/messages.6.gz
 43.3K Oct 28  2015 /var/log/messages.7.gz
 54.8K Oct 20  2015 /var/log/messages.8.gz
 35.8K Oct 19  2015 /var/log/messages.9.gz
 12.4K Jul 21 07:37 /var/log/security.0.gz
 59.4K Jul 19 01:30 /var/log/security.1.gz
 51.8K Apr 25 10:00 /var/log/security.2.gz
 43.6K Apr 14 11:19 /var/log/security.3.gz
 52.7K Apr  5 02:15 /var/log/security.4.gz
 54.4K Mar 25 17:15 /var/log/security.5.gz
 51.9K Mar 16 05:15 /var/log/security.6.gz
 52.0K Mar  5 02:15 /var/log/security.7.gz
 53.4K Feb 22 22:15 /var/log/security.8.gz
 55.6K Feb 13 13:00 /var/log/security.9.gz
 4063B Jul 14 22:51 /var/tmp/cleanup-pkgs.log
```

```
     0B Jul 14 22:51 /var/tmp/eedebug_bin_file
  50.9K Feb  8 20:33 /var/tmp/event_tags.php
    34B Jul 14 22:51 /var/tmp/gksdchk.log
 124.0K Apr 26 06:12 /var/tmp/gres-tp/env.dat
     0B Oct  9  2015 /var/tmp/gres-tp/lock
     4B Jul 14 22:52 /var/tmp/idp_license_info
    46B Jul 14 22:51 /var/tmp/kmdchk.log
    57B Jul 14 22:51 /var/tmp/krt_rpf_filter.txt
    30B Jul 14 22:53 /var/tmp/policy_status
     0B Jul 14 22:51 /var/tmp/rtsdb/if-rtsdb
   349B Jul 14 22:51 /var/tmp/sd-upgrade/debug_log
     0B Oct  9  2015 /var/tmp/spu_kmd_init
    53B Feb  7 23:11 /var/tmp/vjunos-install.log
     0B Jul 14 22:51 /var/tmp/vpn_tunnel_orig.id
```

request system storage cleanup

```
user@host> request system storage cleanup
List of files to delete:

         Size Date         Name
    11B Oct 28 23:40 /var/jail/tmp/alarmd.ts
  92.4K Jan 11 17:12 /var/log/chassisd.0.gz
  92.4K Jan 11 06:06 /var/log/chassisd.1.gz
  92.5K Jan 10 19:00 /var/log/chassisd.2.gz
  92.5K Jan 10 07:53 /var/log/chassisd.3.gz
  92.2K Jan 10 15:00 /var/log/hostlogs/auth.log.1.gz
  92.2K Jan  1 18:45 /var/log/hostlogs/auth.log.2.gz
  92.1K Jan  4 17:30 /var/log/hostlogs/auth.log.3.gz
  92.2K Jan  1 18:45 /var/log/hostlogs/auth.log.4.gz
  79.0K Jan 12 01:59 /var/log/hostlogs/daemon.log.1.gz
  78.8K Jan 11 23:15 /var/log/hostlogs/daemon.log.2.gz
  78.7K Jan 11 20:30 /var/log/hostlogs/daemon.log.3.gz
  79.1K Jan 11 17:44 /var/log/hostlogs/daemon.log.4.gz
  59.1K Jan 11 21:59 /var/log/hostlogs/debug.1.gz
  59.2K Jan 11 17:44 /var/log/hostlogs/debug.2.gz
  59.2K Jan 11 13:29 /var/log/hostlogs/debug.3.gz
  59.3K Jan 11 09:14 /var/log/hostlogs/debug.4.gz
 186.6K Oct 20 16:31 /var/log/hostlogs/kern.log.1.gz
 238.3K Jan 11 23:15 /var/log/hostlogs/lcmd.log.1.gz
 238.4K Jan 11 17:30 /var/log/hostlogs/lcmd.log.2.gz
```

```
  238.6K Jan 11 11:45 /var/log/hostlogs/lcmd.log.3.gz
  238.5K Jan 11 06:00 /var/log/hostlogs/lcmd.log.4.gz
  372.5K Jan 11 17:00 /var/log/hostlogs/syslog.1.gz
  372.5K Jan 11 04:45 /var/log/hostlogs/syslog.2.gz
  371.9K Jan 10 16:30 /var/log/hostlogs/syslog.3.gz
  372.7K Jan 10 04:15 /var/log/hostlogs/syslog.4.gz
   10.1K Jan 12 02:03 /var/log/messages.0.gz
   55.1K Jan  6 21:25 /var/log/messages.1.gz
   81.5K Dec  1 21:30 /var/log/messages.2.gz


Delete these files ? [yes,no] (no)
```

Release Information

Command introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

- Cleaning Up Files with the CLI

request system yang add

IN THIS SECTION

- Syntax | 15405
- Description | 15405
- Options | 15406
- Required Privilege Level | 15406
- Sample Output | 15407
- Release Information | 15407

Syntax

```
request system yang add package package-name module [modules]
<action-script [scripts]>
<translation-script [scripts]>
<deviation-module [modules]>
<proxy-xml>
<snmp>
```

Description

Define a new YANG package with the modules, deviation modules, and scripts that are added to the device as part of the package, and merge the data models defined in the modules with the Junos OS schema. When you add a custom YANG data model to the device, you must also add at least one translation script or one action script, which provides the mapping between the new data model and Junos OS. To add multiple modules or scripts, include a space-delimited list of absolute or relative file paths enclosed in brackets.

> **NOTE**: To install OpenConfig modules that are packaged as a compressed tar file, use the `request system software add` command. OpenConfig modules and scripts that are installed using the `request system software add` command are always associated with the package identifier `openconfig`.

When you create a new package, the device stores copies of the module and script files in a new location. The device also stores copies of the action script and translation script files under the **/var/db/scripts/action** and **/var/db/scripts/translation** directories, respectively. Junos OS validates the syntax of the modules and scripts, rebuilds its schema to include the new data models, and then validates the active configuration against this schema. Newly added RPCs and configuration hierarchies are immediately available for use.

> **NOTE**: Devices that use the ephemeral configuration database will delete all ephemeral configuration data in the process of rebuilding the schema.

> **NOTE**: To prevent CLI-related or configuration database errors, we recommend that you do not perform any CLI operations, change the configuration, or terminate the operation while a device is in the process of adding, updating, or deleting a YANG package and modifying the schema.

> **NOTE**: Starting in Junos OS Release 18.3R1, adding, deleting, or updating YANG packages in configuration mode with the `run` command is not supported.

Options

action-script [*scripts*]	List of paths for one or more action scripts to add to the device as part of the package.
module [*modules*]	List of paths for one or more YANG modules to add to the device as part of the package. The device merges the data models defined in the modules with the Junos OS schema.
deviation-module [*modules*]	(Optional) List of paths for one or more modules that define deviation statements that should be applied to modules in the package.
package *package-name*	User-defined identifier that represents the collection of YANG modules and scripts.
proxy-xml	(Optional) Specify that `module` is a list of paths for one or more modules that provide user-defined OpenConfig mappings for the XML Proxy process to translate Junos Telemetry Interface statistics exported through gRPC into key-value pairs.
snmp	(Optional) Specify that `module` is a list of paths for one or more YANG modules that define custom MIBs. The system converts the modules to JSON format, and the snmpd process parses the JSON data and builds its internal database.
translation-script [*scripts*]	List of paths for one or more translation scripts to add to the device as part of a package. YANG modules that define configuration data models require one or more translation scripts to map the nonnative configuration syntax to the corresponding Junos OS syntax.

Required Privilege Level

maintenance

Sample Output

request system yang add

```
user@host> request system yang add package p1 module [yang/if.yang yang/if-aggregate.yang
yang/if-show.yang] deviation-module yang/deviation/if-devs.yang translation-script translation/
if.slax action-script action/if-show.py

YANG modules validation : START
YANG modules validation : SUCCESS
Scripts syntax validation : START
script check succeeds
Scripts syntax validation : SUCCESS
Scripts syntax validation : START
Scripts syntax validation : SUCCESS
TLV generation: START
TLV generation: SUCCESS
Building schema and reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...

WARNING: cli has been replaced by an updated version:
CLI release 16.1R1 built by builder on 2016-03-30 13:46:11 UTC
Restart cli using the new version ? [yes,no] (yes) yes

Restarting cli ...
user@host>
```

Release Information

Command introduced in Junos OS Release 16.1R1.

`proxy-xml` option introduced in Junos OS Release 17.3R1 on MX Series and PTX Series routers.

`snmp` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Managing YANG Packages, Modules, and Scripts on Devices Running Junos OS

Understanding the Management of Nonnative YANG Modules on Devices Running Junos OS

Configure a NETCONF Proxy Telemetry Sensor in Junos

request system yang update | **15414**

show system yang package | **27323**

Customized SNMP MIBs for Syslog Traps

request system yang delete

IN THIS SECTION

- Syntax | **15408**
- Description | **15408**
- Options | **15409**
- Required Privilege Level | **15409**
- Sample Output | **15410**
- Release Information | **15410**

Syntax

```
request system yang delete package-name
```

Description

Remove the given YANG package and all of its modules and scripts from the device, and remove the data models associated with that package from the Junos OS schema.

CAUTION: Before you delete a YANG package, ensure that the active configuration does not contain configuration data that has dependencies on the data models added by that package.

> **NOTE**: You must use the `request system software delete` command to remove OpenConfig packages that were installed from a compressed tar file using the `request system software add` command.

When you delete a package, Junos OS rebuilds its schema to remove the data models associated with that package and then validates the active configuration against the newly updated schema. The device removes the copies of the module and script files that were generated when the package was created. The device also removes the copies of the package's action script and translation script files that are stored under the **/var/db/scripts/action** and **/var/db/scripts/translation** directories. If you downloaded the original module and script files to a different location, the original files remain unchanged.

> **NOTE**: Devices that use the ephemeral configuration database will delete all ephemeral configuration data in the process of rebuilding the schema.

> **NOTE**: To prevent CLI-related or configuration database errors, we recommend that you do not perform any CLI operations, change the configuration, or terminate the operation while a device is in the process of adding, updating, or deleting a YANG package and modifying the schema.

> **NOTE**: Starting in Junos OS Release 18.3R1, adding, deleting, or updating YANG packages in configuration mode with the `run` command is not supported.

Options

package-name	Name of the YANG package to remove.

Required Privilege Level

maintenance

Sample Output

request system yang delete

```
user@host> request system yang delete p1
Building schema and reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...

WARNING: cli has been replaced by an updated version:
CLI release 16.1R1 built by builder on 2016-03-30 13:46:11 UTC

Restart cli using the new version ? [yes,no] (yes) yes

Restarting cli ...
```

Release Information

Command introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Managing YANG Packages, Modules, and Scripts on Devices Running Junos OS

Understanding the Management of Nonnative YANG Modules on Devices Running Junos OS

request system yang add | **15404**

show system yang package | **27323**

request system yang disable

IN THIS SECTION

- Syntax | **15411**
- Description | **15411**

- Options | 15412
- Required Privilege Level | 15412
- Sample Output | 15412
- Release Information | 15412

Syntax

```
request system yang disable package-name
```

Description

Disable the translation scripts associated with the given YANG package.

Translation scripts convert configuration data corresponding to YANG data models into Junos OS syntax and add the translated configuration data as a transient change in the checkout configuration during the commit operation. Translation scripts are enabled by default as soon as you add the scripts and related YANG modules to the device using the appropriate operational command.

Use this command to temporarily disable translation scripts for a package to help troubleshoot translation issues instead of deleting the entire package, which would remove the associated data models from the Junos OS schema as well as remove the package and related files from the device. After you disable translation for a package and commit the configuration, the configuration data associated with the YANG data models in that package can be present in the active configuration, but the configuration has no impact on the functioning of the device.

When translation is disabled, you can still configure and commit the statements and hierarchies in the data models added by that package. However, the device does not commit the corresponding Junos OS configuration statements as transient changes during the commit operation for any statements in the data models added by that package, even for those statements that were committed prior to disabling translation.

NOTE: When you disable translation for a package, the device retains any transient configuration changes that were committed prior to disabling translation until the next commit operation.

> **TIP**: Use the `show system yang package` *`package-name`* command to verify the translation status of a package.

Options

package-name Name of the YANG package for which to disable translation.

Required Privilege Level

maintenance

Sample Output

request system yang disable

```
user@host> request system yang disable p1
user@host>
```

Release Information

Command introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Disabling and Enabling YANG Translation Scripts on Devices Running Junos OS

Creating Translation Scripts for YANG Configuration Models

request system yang enable | **15413**

show system yang package | **27323**

request system yang enable

IN THIS SECTION

- Syntax | 15413
- Description | 15413
- Options | 15414
- Required Privilege Level | 15414
- Sample Output | 15414
- Release Information | 15414

Syntax

```
request system yang enable package-name
```

Description

Enable the translation scripts associated with the given YANG package.

Translation scripts convert configuration data corresponding to YANG data models into Junos OS syntax and add the translated configuration data as a transient change in the checkout configuration during the commit operation. Translation scripts are enabled by default as soon as you add the scripts and related YANG modules to the device using the appropriate operational command. Use this command to enable translation scripts that were previously disabled using the `request system yang disable` command.

NOTE: When you enable translation for a package, configuration data that is associated with the YANG data models in that package and that is present in the active configuration does not impact the functioning of the device until the next commit operation.

TIP: Use the `show system yang package` *`package-name`* command to verify the translation status of a package.

Options

package-name Name of the YANG package for which to enable translation.

Required Privilege Level

maintenance

Sample Output

request system yang enable

```
user@host> request system yang enable p1
user@host>
```

Release Information

Command introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Disabling and Enabling YANG Translation Scripts on Devices Running Junos OS

Creating Translation Scripts for YANG Configuration Models

request system yang disable | **15410**

show system yang package | **27323**

request system yang update

IN THIS SECTION

- Syntax | **15415**
- Description | **15415**
- Options | **15415**

- Required Privilege Level | 15416
- Sample Output | 15416
- Release Information | 15417

Syntax

```
request system yang update package-name action-script [scripts] deviation-module [modules]
module [modules] proxy-xml [file-path-names] translation-script [scripts]
```

Description

Update an existing YANG package to include new or modified YANG modules or scripts, and merge the updated data models in that package with the Junos OS schema.

When you update a package, the device stores copies of the new and modified module and script files. Junos OS then rebuilds its schema to include the changes to the data models and validates the active configuration against this schema.

NOTE: Devices that use the ephemeral configuration database will delete all ephemeral configuration data in the process of rebuilding the schema.

NOTE: To prevent CLI-related or configuration database errors, we recommend that you do not perform any CLI operations, change the configuration, or terminate the operation while a device is in the process of adding, updating, or deleting a YANG package and modifying the schema.

NOTE: Starting in Junos OS Release 18.3R1, adding, deleting, or updating YANG packages in configuration mode with the `run` command is not supported.

Options

package-name	Name of the YANG package to update.
action-script [*scripts*]	List of paths for one or more action scripts to add to or update in the package.

deviation-module [*modules*]	List of paths for one or more deviation modules to add to or update in the package.
module [*modules*]	List of paths for one or more YANG modules to add to or update in the package.
proxy-xml [*file-path-names*]	List of paths for one or more YANG modules to add to or update in the package that provide user-defined OpenConfig mappings for the XML Proxy process to translate Junos Telemetry Interface statistics exported through gRPC into key-value pairs.
translation-script [*scripts*]	List of paths for one or more translation scripts to add to or update in the package.

Required Privilege Level

maintenance

Sample Output

request system yang update

```
user@host> request system yang update p1 module yang/if.yang

YANG modules validation : START
YANG modules validation : SUCCESS
TLV generation: START
TLV generation: SUCCESS
Building schema and reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...

WARNING: cli has been replaced by an updated version:
CLI release 16.1R1 built by builder on 2016-03-30 13:46:11 UTC
Restart cli using the new version ? [yes,no] (yes) yes

Restarting cli ...
```

Release Information

Command introduced in Junos OS Release 16.1R1.

`proxy-xml` option introduced in Junos OS Release 17.3R1 on MX Series and PTX Series routers.

RELATED DOCUMENTATION

Managing YANG Packages, Modules, and Scripts on Devices Running Junos OS

Configure a NETCONF Proxy Telemetry Sensor in Junos

request system yang add | **15404**

show system yang package | **27323**

request system yang validate

IN THIS SECTION

- Syntax | **15417**
- Description | **15417**
- Options | **15418**
- Required Privilege Level | **15418**
- Sample Output | **15418**
- Release Information | **15418**

Syntax

```
request system yang validate action-script [scripts] module [modules] proxy-xml module [modules]
translation-script [scripts]
```

Description

Validate the syntax of one or more YANG modules, translation scripts, or action scripts.

Options

action-script *scripts*	List of paths for one or more action scripts to validate.
module *modules*	List of paths for one or more YANG modules to validate.
proxy-xml module *modules*	List of paths for one or more YANG modules to validate that provide user-defined OpenConfig mappings for the XML Proxy process to translate Junos Telemetry Interface statistics exported through gRPC into key-value pairs.
translation-script *scripts*	List of paths for one or more translation scripts to validate.

Required Privilege Level

maintenance

Sample Output

request system yang validate

```
user@host> request system yang validate module [yang/if.yang yang/if-aggregate.yang] translation-
script translation/if.slax
YANG modules validation : START
YANG modules validation : SUCCESS
Scripts syntax validation : START
script check succeeds
Scripts syntax validation : SUCCESS
```

Release Information

Command introduced in Junos OS Release 16.1R1.

`proxy-xml` option introduced in Junos OS Release 17.3R1 on MX Series and PTX Series routers.

RELATED DOCUMENTATION

Managing YANG Packages, Modules, and Scripts on Devices Running Junos OS

Understanding the Management of Nonnative YANG Modules on Devices Running Junos OS

Configure a NETCONF Proxy Telemetry Sensor in Junos

request system zeroize

IN THIS SECTION

- Syntax | 15419
- Description | 15419
- Options | 15419
- Required Privilege Level | 15420
- Sample Output | 15420

Syntax

```
request system zeroize <media>
```

Description

Use this command to remove all configuration information and resets all key values. The command removes all data files, including customized configuration and log files, by unlinking the files from their directories. If the device has two Routing Engines, the command is broadcast to both Routing Engines on the device. As a result, both Routing Engines are zeroized.

The command removes all user-created files from the system including all plain-text passwords, secrets, and private keys for SSH, local encryption, local authentication, IPsec, RADIUS, TACACS+, and SNMP.

This command reboots the device and sets it to the factory default configuration. After the reboot, you cannot access the device through the management Ethernet interface. Log in through the console as root and start the Junos OS CLI by typing `cli` at the prompt.

Options

media (Optional) In addition to removing all configuration and log files, the media option causes memory and the media to be scrubbed, removing all traces of any user-created files. Every storage device attached to the system is scrubbed, including disks, flash drives, removable USBs, and the like. The duration of the scrubbing process is dependent on the size of the media being erased. As a result, the request system zeroize media operation can take considerably more time

than the request system zeroize operation. However, the critical security parameters are all removed at the beginning of the process.

> **NOTE**: The `media` option is not supported on SRX5000 line devices.

Required Privilege Level

Not applicable.

Sample Output

request system zeroize

```
user@host> request system zeroize
warning: System will be rebooted and may not boot without configuration


Erase all data, including configuration and log files?. In case of Dual RE system, both Routing
Engines will be zeroized [yes,no] (no) yes

warning: zeroizing re0

Loading /boot/loader            Consoles: serial port
BIOS driver C: is disk0
BIOS 607kB/2087552kB available memory

FreeBSD/i386 bootstrap loader, Revision 1.1
(builder@youcompany.com, Mon Mar 28 20:49:26 UTC 2011)
Loading /boot/defaults/loader.confg
/kernel text-0x837a60 data=0x46a78+0x9d44c syms=[0x4+0x8f38+0x4+0xca1ee]

Hit [Enter[ to boot immediately, or space bar for command prompt.
Booting [/kernel]...
platform_early_bootinit: MAG Series Early Boot Initilaization
GDB: debug ports: sio
GDB: current port: sio
KDB: debugger backends: ddb gdb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
```

```
All rights resrved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 18\989, 1991, 1992, 1993,1994
    The Regents of the University of California. All rights reserved.
...
output truncated
```

RELATED DOCUMENTATION

request system software rollback (SRX Series)

request system zeroize (Junos OS)

IN THIS SECTION

- Syntax | 15421
- Description | 15421
- Options | 15423
- Required Privilege Level | 15423
- Sample Output | 15423
- Release Information | 15425

Syntax

```
request system zeroize
<media>
<local>
```

Description

Use this command to remove all configuration information on the Routing Engines and reset all key values on the device where you run the command.

- If the device has dual Routing Engines, the command is broadcast to all Routing Engines on the device. As a result, both Routing Engines are zeroized.
- In a Virtual Chassis or Virtual Chassis Fabric (VCF) composed of EX Series switches (except EX8200 Virtual Chassis) or QFX Series switches, this command operates only on the member switch where you run the command, even if that switch is in the primary Routing Engine role. The command is not forwarded to the backup Routing Engine member or to member switches in the line-card role. To apply this command to more than one member of an EX Series or QFX Series Virtual Chassis or VCF, we recommend you remove and disconnect each of those members from the Virtual Chassis or VCF, and then run the command on each isolated switch individually.

This command removes all data files, including customized configuration and log files, by unlinking the files from their directories. The command removes all user-created files from the system, including all plain-text passwords, secrets, and private keys for SSH, local encryption, local authentication, IPsec, RADIUS, TACACS+, and SNMP.

This command reboots the device and sets it to the factory default configuration. After the reboot, you cannot access the device through the management Ethernet interface. Log in through the console as `root` and start the Junos OS CLI by typing `cli` at the prompt.

If the configuration contains the `commit synchronize` statement at the `[edit system]` hierarchy level, and you issue a `commit` in the primary Routing Engine, the primary configuration is automatically synchronized with the backup. If the backup Routing Engine is down when you issue the `commit`, the Junos OS displays a warning and commits the candidate configuration in the primary Routing Engine. When the backup Routing Engine comes up, its configuration will automatically be synchronized with the primary. A newly inserted backup Routing Engine or a Routing Engine that comes up after running the `request system zeroize` command also automatically synchronizes its configuration with the primary Routing Engine configuration.

Starting with Junos OS Release 15.1F3, the `request system zeroize` command removes all configuration information on the guest OS for the PTX5000 router with RE-DUO-C2600-16G, and MX240, MX480, and MX960 with RE-S-1800X4-32G-S.

Starting with Junos OS Release 15.1F5, the `request system zeroize` command removes all configuration information on the guest OS for the MX2010 and MX2020 with REMX2K-1800-32G-S.

On these routers, in order to remove all configuration information on both guest OS and host OS, use the `request vmhost zeroize` command.

To completely erase user-created data so that it is unrecoverable, use the `media` option.

Options

media (Optional) In addition to removing all configuration and log files, causes memory and the media to be scrubbed, removing all traces of any user-created files. Every storage device attached to the system is scrubbed, including disks, flash drives, removable USBs, and so on. The duration of the scrubbing process is dependent on the size of the media being erased. As a result, the `request system zeroize media` operation can take considerably more time than the `request system zeroize` operation. However, the critical security parameters are all removed at the beginning of the process.

On QFX Series platforms running Junos OS Release 14.1X53 or earlier, the `media` option is not available. On QFX Series platforms running releases later than Junos OS Release 14.1X53 that do not have the upgraded FreeBSD kernel (10+), the `media` option is available, but if you use it, the system will issue a warning that the `media` option is not supported and will continue with the zeroize operation. On platforms that are not QFX Series platforms, the `media` option is not available in Junos OS Release 17.2 or later with Junos with upgraded FreeBSD.

local (Optional) Remove all the configuration information and restore all the key values on the active Routing Engine.

> **CAUTION**: The `local` option is not available on devices that run Junos OS Evolved, so if you execute `request system zeroize` on those devices and they have dual Routing Engines, all Routing Engines on the local chassis are rebooted. See the Junos® OS Evolved Software Installation and Upgrade Guide.

Specifying this option has no effect on switches in a Virtual Chassis or VCF composed of EX Series switches (except EX8200 Virtual Chassis) or QFX switches, because in these configurations, the `request system zeroize` command only operates locally by default.

Required Privilege Level

maintenance

Sample Output

request system zeroize

```
user@host> request system zeroize
warning: System will be rebooted and may not boot without configuration
```

```
Erase all data, including configuration and log files?. In case of Dual RE system, both Routing
Engines will be zeroized [yes,no] (no) yes
0 1 1 0 0 0 done

syncing disks... All buffers synced.
Uptime: 5d19h20m26s
recorded reboot as normal shutdown
Rebooting...

U-Boot 1.1.6 (Mar 11 2011 - 04:39:06)

Board: EX4200-24T 2.11
EPLD:  Version 6.0 (0x85)
DRAM:  Initializing (1024 MB)
FLASH: 8 MB

Firmware Version: --- 01.00.00 ---
USB:   scanning bus for devices... 2 USB Device(s) found
       scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.4
(user@device.example.net, Fri Mar 11 03:03:36 UTC 2011)
Memory: 1024MB
bootsequencing is enabled
bootsuccess is set
new boot device = disk0s1:
Loading /boot/defaults/loader.conf
/kernel data=0x915c84+0xa1260 syms=[0x4+0x7cbd0+0x4+0xb1c19]



Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
```

```
        The Regents of the University of California. All rights reserved.
JUNOS 11.1R1.8 #0: 2011-03-09 20:14:25 UTC
    user@device.example.net:/volume/build/junos/11.1/release/11.1R1.8/obj-powerpc/bsd/kernels/
        JUNIPER-EX/kernel
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080
...
```

Release Information

Command introduced before Junos OS Release 9.0.

Option `media` added in Junos OS Release 11.4 for EX Series switches.

Option `local` added in Junos OS Release 14.1.

RELATED DOCUMENTATION

- request system snapshot (Junos OS with FreeBSD Prior to Release 10) | **15190**
- *Reverting to the Default Factory Configuration for the EX Series Switch*
- Reverting to the Rescue Configuration for the EX Series Switch
- *Reverting to the Default Factory Configuration*
- *Reverting to the Rescue Configuration*
- Reverting to the Default Factory Configuration by Using the request system zeroize Command

request system zeroize (Junos OS Evolved)

IN THIS SECTION

- Syntax | **15426**
- Description | **15426**
- Options | **15427**
- Determine whether both Routing Engines are Booted from the Same Disk | **15428**
- Required Privilege Level | **15428**

Sample Output | 15428
Release Information | 15430

Syntax

```
request system zeroize <(disk1 | disk2)>
```

Description

Use this command to sanitize the disks on a Routing Engine, to make it difficult to access data on the disks, using various levels of sanitization corresponding to degrees of difficulty, as defined in the NIST 800-88 standard. If the device has two Routing Engines, the command is broadcast to both Routing Engines on the device. As a result, both Routing Engines are zeroized. When you issue this command, the console logs show which sanitization level was actually used on the disks. The console logs also provide a list of the erase procedures supported by the disk. The same list order is used by this command to erase the disk. If any erase procedure fails, the next erase mechanism is tried, and so on, until the last erase mechanism is what is used to sanitize the disk.

Starting in Junos OS Evolved Release 21.3R1, for devices that support this feature, if the disks in the Routing Engine support the ATA standard, this command sanitizes the disks on the Routing Engine using the ATA `secure erase` command to overwrite the data. The ATA `secure erase` command overwrites the contents of LBA 0 through the greater of either READ NATIVE MAX or READ NATIVE MAX EXT, and replaces the contents with 0s or 1s. If the disks do not support the ATA standard (for example, they support the SCSI standard), they are sanitized using the older method described below. The secure erase capability is classified under the CLEAR NIST media sanitization level, according to the NIST 800-88 standard. When the secure erase is complete, the system copies the current running OS from RAMDISK to the ATA disk. Once the current running OS is installed, the system reboots and comes back to the factory default configuration.

Starting in Junos OS Evolved 24.4R1, for devices that support this feature, if the disks in the Routing Engine support the SATA standard, this command sanitizes the disks using the PURGE NIST media sanitization level, according to the NIST 800-88 standard. The PURGE level is comprised of both the CRYPTO_SCRAMBLE (if supported by the SATA SSD controller) and the BLOCK_ERASE mechanisms. The CRYPTO_ERASE mechanism (if supported by the SATA SSD controller) is followed by the BLOCK_ERASE mechanism in all cases of sanitization. Whenever the CRYPTO_SCRAMBLE mechanism is not supported, only the BLOCK_ERASE mechanism is run. When the disk sanitization is complete, the system copies the current running OS from RAMDISK to the SATA disk. Once the current running OS is installed, the system reboots and comes back to the factory default configuration.

Prior to Junos OS Evolved Release 21.3R1, and in any release for disks that do not support the NIST 800-88 sanitization standard, this command removes all data files, including any customized configuration and log files, by unlinking the files from their directories. The command removes all user-created files from the system, including all plain-text passwords, secrets, and private keys for SSH, local encryption, local authentication, IPSec, RADIUS, TACACS+, and SNMP.

This command reboots the device and sets it to the factory default configuration. After the reboot, you cannot access the device through the management Ethernet interface. Log in through the console as `root` and start the CLI by typing `cli` at the prompt.

Options

(disk1 | disk2) (For Routing Engines that contain two disks only. For Routing Engines that contain only one disk, the (disk1|disk2) option does not apply.) The `disk1` and `disk2` options refer to the absolute disk; that is, to the primary disk and to the snapshot disk, respectively. One use case for this option is when you need to replace a disk on the Routing Engine, and you want to securely erase the information before removing the disk from the Routing Engine.

The `request system zeroize disk1` command secure erases the primary disk and installs the current running Junos OS Evolved version on the Routing Engine (or on the primary disk on both Routing Engines in a chassis-based system. The contents on disk 2 (that is, the snapshot disk) remain **unchanged**.

The `request system zeroize disk2` command secure erases the snapshot disk and installs the current running Junos OS Evolved version on the Routing Engine (or on the secondary disk on both Routing Engines in a chassis-based system). The contents on disk 1 (that is, the primary disk) remain **unchanged**.

When you use any zeroize option that cleans the active disk (disk 1 in the case of primary boot or disk 2 in the case of snapshot boot), the system automatically reboots to finish the zeroize of the active disk. When you zeroize the backup disk (that is, disk 2 in the case of primary boot or disk 1 in the case of snapshot boot), the system does not reboot.

Note that the behavior of the `request system zeroize (disk1 | disk2)` command is non-deterministic in a dual Routing Engine system where one Routing Engine is in snapshot boot and other Routing Engine is in primary boot. For example, if the primary Routing Engine is in primary boot and the backup Routing Engine is in snapshot boot, then the `request system zeroize disk1` command acts on the active disk in the primary Routing Engine and on the backup disk in the backup Routing Engine. Because the primary Routing Engine reboots the chassis after zeroizing its active disk, the backup Routing Engine, which is zeroizing its snapshot disk, is interrupted by the reboot and the data on its snapshot disk will be inconsistent. Therefore, the final behavior is non-deterministic in the case of the master and backup Routing Engine, because they are running from different boot disks.

Default: Zeroize all disks and both Routing Engines, if applicable. The default order is to first create a snapshot on secondary disk and then to install the OS on the primary disk. There is no difference in behavior between single Routing Engine or dual Routing Engine systems, except that for dual Routing Engine systems, the command is run on both Routing Engines.

Determine whether both Routing Engines are Booted from the Same Disk

Log in into both Routing Engines. The `show system software list` command displays a banner if the Routing Engine was booted from the snapshot disk (disk2). The same command doesn't display a banner if the Routing Engine was booted from the primary disk (disk1). In this way, you can check if both Routing Engines are booted from the same disk. For example, if the command `show system software list` on RE0 does not display a banner for booting from snapshot and RE1 shows a banner, then RE0 was booted from its primary disk, but RE1 was booted from its snapshot disk instead.

In case of issues, you can boot the affected Routing Engine from either disk1 or disk2, whichever disk did not boot from the required disk. Both Routing Engines may boot from disk1 or disk2, based on where this command was issued.

Required Privilege Level

maintenance

Sample Output

request system zeroize

```
user@host> request system zeroize
warning: System will be rebooted and may not boot without configuration

Erase all data, including configuration and log files?. In case of Dual RE system, both Routing
Engines will be zeroized [yes,no] (yes)
```

Example: Console logs for SATA drive sanitization

```
user@host> request system zeroize
NIST secure erase needed..

=========================================
```

```
NIST Secure Erase START...

========================================

SATA drive.. Checking capabilities.

******************************************

Supported Secure Erase features of SATA HDD

******************************************

Device supports SANITIZE feature

Device supports CRYPTO_SCRAMPLE feature

Device supports Block Erase feature

Device supports Enhanced Secure Erase feature

Device supports Secure Erase feature
*********************************************************************************

Secure Erase mechanism: Crypto Scramble

Secure erase command: hdparm --sanitize-crypto-scramble /dev/sda

*****************************************

    Last Sanitize Operation Completed Without Error

Crypto erase succeeded..

Executing Block Erase to finish disk scrubbing

*****************************************

Secure Erase mechanism: Block EraseSecure erase

command: hdparm --sanitize-block-erase /dev/sda

*****************************************
```

```
####

   Last Sanitize Operation Completed Without Error

Secure erase succeeded

=========================================

NIST Secure Erase END...

=========================================
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

Secure erase capability according to the NIST 800-88 standard introduced in Junos OS Evolved Release 21.3R1.

`disk1` and `disk2` options introduced in Junos OS Evolved Release 23.4R1.

Support for NIST PURGE capability for SATA drives introduced in Junos OS Evolved 24.4R1.

RELATED DOCUMENTATION

NIST Special Publication 800-88, Guidelines for Media Sanitization

request system zeroize ztp-option secure

IN THIS SECTION

- Syntax | **15431**
- Description | **15431**
- Options | **15431**
- Required Privilege Level | **15431**
- Sample Output | **15431**

Release Information | 15432

Syntax

```
request system zeroize ztp-option secure-(enable | disable)
```

Description

On secure devices, you can switch between using secure zero-touch provisioning (SZTP) and zero-touch provisioning (ZTP).

Options

enable The CLI checks to see if the device is a secure device. If the device is secure, the process ends. The next time the device boots, the device uses secure ZTP as the provisioning solution. If the device is not a secure device, you will receive an error message that says the device is not secure, and the process ends.

disable The CLI checks to see if the device is a secure device. If the device is secure, the next time the device boots, the device uses ZTP as the provisioning solution. If the device is not secure, the process ends.

Required Privilege Level

maintenance

Sample Output

request system zeroize ztp-option secure-enable

```
user@device> request system zeroize ztp-option secure-enable
```

request system zeroize ztp-option secure-disable

```
user@device> request system zeroize ztp-option secure-disable
```

Release Information

Command introduced in Junos OS Evolved Release 23.2.

RELATED DOCUMENTATION

Switching between Secure Zero Touch Provisioning and Zero Touch Provisioning

request-timeout

IN THIS SECTION

- Syntax | 15432
- Hierarchy Level | 15432
- Description | 15433
- Options | 15433
- Required Privilege Level | 15433
- Release Information | 15433

Syntax

```
request-timeout seconds;
```

Hierarchy Level

```
[edit unified-edge diameter-profiles gx-profile profile-name]
```

Description

Configure the time to wait for a response from the server.

Options

seconds Length of timeout interval.

- **Range:** 0 through 65,535 seconds.

> **NOTE**: 0 seconds indicates that the request timeout is not be enabled.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

gx-profile | 7954

request-type

IN THIS SECTION

- Syntax | 15434
- Hierarchy Level | 15434
- Description | 15434
- Options | 15434

Required Privilege Level | 15434
Release Information | 15434

Syntax

```
request-type (get-next-request | get-request | walk-request);
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Extend monitoring to a specific SNMP object instance (`get-request`), or extend monitoring to all object instances belonging to a MIB branch (`walk-request`), or extend monitoring to the next object instance after the instance specified in the configuration (`get-next-request`).

Options

`get-next-request`—Performs an SNMP get next request.

`get-request`—Performs an SNMP get request.

`walk-request`—Performs an SNMP walk request.

- **Default:** `walk-request`

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Remote Nework Monitoring (RMON)

variable | 30653

request unified-edge tdf call-trace clear

IN THIS SECTION

- Syntax | 15435
- Description | 15435
- Options | 15435
- Required Privilege Level | 15435
- Output Fields | 15436
- Sample Output | 15436
- Release Information | 15436

Syntax

```
request unified-edge tdf call-trace clear
```

Description

Clear the completed or duplicate subscriber call traces on one or more TDF gateways.

Options

This command has no options.

Required Privilege Level

unified-edge

Output Fields

No message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

request unified-edge tdf call-trace clear

```
user@host> request unified-edge tdf call-trace clear
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- request unified-edge tdf call-trace show | 15436
- request unified-edge tdf call-trace start | 15440
- request unified-edge tdf call-trace stop | 15443

request unified-edge tdf call-trace show

IN THIS SECTION

- Syntax | 15437
- Description | 15437
- Options | 15437
- Required Privilege Level | 15437
- Output Fields | 15437
- Sample Output | 15439
- Release Information | 15440

Syntax

```
request unified-edge tdf call-trace show
<all | completed | current>
<brief | detail>
```

Description

Display the information related to subscriber call tracing on one or more TDF gateways.

Options

none (Same as brief) Display information related to subscriber call tracing in brief.

all | completed | current (Optional) Display call trace information for the following:

- `all`—All calls.
- `completed`—Completed calls only.
- `current`—Call traces that are currently active.

brief | detail (Optional) Display the specified level of output.

Required Privilege Level

unified-edge

Output Fields

Table 132 on page 15437 lists the output fields for the `request unified-edge tdf call-trace show` command. Output fields are listed in the approximate order in which they appear.

Table 132: request unified-edge tdf call-trace show Output Fields

Field Name	Field Description	Level of Output
`Identifier`	Identifier for the call trace.	All levels

Table 132: request unified-edge tdf call-trace show Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`File name`	Name of the call trace file.	none `brief`
`Trace file`	Name of the call trace file.	`detail`
`Status`	Status of the call trace: • `done`—Call trace complete. • `not-done`—Call trace in progress. • `duplicate`—Another call trace record is present that has the same attributes.	All levels
`SPIC Mask create`	Internal mask of the services PIC where this call trace was enabled.	none `brief`
`Create Mask`	Internal mask of the services PIC where this call trace was enabled.	`detail`
`SPIC Mask complete`	Internal mask of the services PIC where this call trace was completed.	none `brief`
`Complete Mask`	Internal mask of the services PIC where this call trace was completed.	`detail`
`IMSI`	International Mobile Subscriber Identity (IMSI) of the subscriber's user equipment (UE).	`detail`
`MSISDN`	Mobile station ISDN of the subscriber's user equipment.	
`Calls Traced`	Number of calls traced.	`detail`

Table 132: request unified-edge tdf call-trace show Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Next Call	Number of next calls to be traced. For example, a value of 10 indicates that the next 10 calls are traced.	detail
TDF domain	TDF domain pertaining to the subscriber's call.	detail
FPC	FPC slot on which the call trace was enabled. This field is displayed only if the call trace is enabled on the FPC slot.	detail
PIC	PIC slot on which the call trace was enabled. This field is displayed only if the call trace is enabled on the PIC slot.	detail

Sample Output

request unified-edge tdf call-trace show brief

```
user@host> request unified-edge tdf call-trace show brief
  Identifier           File name           Status       SPIC Mask   SPIC Mask
                                                         create      complete
 call_trace_id_2 call_trace_id_2_02112012_060450        done 0x10   0x10
 call_trace_id_3 call_trace_id_3_02112012_070614        done 0x10   0x10
 call_trace_id_4 call_trace_id_4_02112012_071342   duplicate 0x0    0x0
 call_trace_id_5 call_trace_id_5_02112012_201317   duplicate 0x0    0x0
 call_trace_id_6 call_trace_id_6_02112012_201649   duplicate 0x0    0x0
 call_trace_id_7 call_trace_id_7_02112012_202501        done 0x0    0x0
 call_trace_id_8 call_trace_id_8_02112012_204718   duplicate 0x0    0x0
 call_trace_id_9 call_trace_id_9_02112012_204759    not-done 0x10   0x0
```

request unified-edge tdf call-trace show detail

```
user@host> request unified-edge tdf call-trace show detail
 Call trace information :
```

```
Identifier : call_trace_id_13        Trace file : call_trace_id_13_02292012_001343
Status : not-done    Create Mask : 0x200      Complete Mask : 0x0
IMSI : 29299
Calls Traced : 0
Identifier : call_trace_id_14        Trace file : call_trace_id_14_02292012_001348
Status : not-done    Create Mask : 0x200      Complete Mask : 0x0
                                                   MS-ISDN: 2929910000000000
Calls Traced : 0
Identifier : call_trace_id_15        Trace file : call_trace_id_15_02292012_001408
Status : not-done    Create Mask : 0x200      Complete Mask : 0x0
Next Call : 1                        TDF domain : jnpr-sunnyvale
Calls Traced : 0
Identifier : call_trace_id_16        Trace file : call_trace_id_16_02292012_001416
Status : not-done    Create Mask : 0x200      Complete Mask : 0x0
Calls Traced : 0                     FPC : 3   PIC : 1
Identifier : call_trace_id_17        Trace file : call_trace_id_17_02292012_001424
Status : done        Create Mask : 0x200      Complete Mask : 0x200
Next Call : 2
Calls Traced : 2
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

request unified-edge tdf call-trace clear | 15435

request unified-edge tdf call-trace start | 15440

request unified-edge tdf call-trace stop | 15443

request unified-edge tdf call-trace start

IN THIS SECTION

- Syntax | 15441
- Description | 15441

- Options | 15441
- Required Privilege Level | 15442
- Output Fields | 15442
- Sample Output | 15443
- Release Information | 15443

Syntax

```
request unified-edge tdf call-trace start
<imsi imsi>
<msisdn msisdn>
<next-call next-call>
<routing-instance routing-instance>
<subscriber-name string>
<user-name string>
<v4-addr v4-addr>
<v6-addr v6-addr>
```

Description

Start TDF subscriber call tracing.

Options

none	Start call tracing for all TDF subscribers.
imsi *imsi*	(Optional) Start the call tracing for subscribers with the specified International Mobile Subscriber Identity (IMSI) number.
msisdn *msisdn*	(Optional) Start call tracing for subscribers with the specified Mobile station ISDN (MSIDSN) number.
next-call *next-call*	(Optional) Start call tracing for the specified number of next call events (1 through 50). For example, if you specify 10, then the next 10 calls are traced.

> **NOTE**: If you do not include the `next-call` keyword while tracing subscribers on a domain, the default value of 1 is used.

routing-instance *routing-instance* — (Optional) Start call tracing for subscribers for the specified routing instance.

subscriber-name *string* — (Optional) Start call tracing for the specified IFL-based subscriber.

user-name *string* — (Optional) Start call tracing for the specified IP-based subscriber.

v4-addr *v4-addr* — (Optional) Start call tracing for subscribers for the specified IPv4 address of the subscriber's user equipment (UE).

v6-addr *v6-addr* — (Optional) Start call tracing for subscribers for the specified IPv6 address of the subscriber's user equipment.

Required Privilege Level

unified-edge

Output Fields

Table 133 on page 15442 lists the output fields for the `request unified-edge tdf call-trace start` command. Output fields are listed in the approximate order in which they appear.

Table 133: request unified-edge tdf call-trace start Output Fields

Field Name	Field Description
`Session PIC`	Identifier of the session PIC for which the call trace status is displayed.
`Status`	Status of the call trace: • `duplicate`—Another call trace record is present that has the same attributes. • `success`—Call trace started successfully. • `fail`—Call tracing cannot be started.

Sample Output

request unified-edge tdf call-trace start next-call

```
user@host> request unified-edge tdf call-trace start next-call 10
       Session PIC            Status
       ms-0/1/0             success
       ms-1/1/0             success
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- request unified-edge tdf call-trace clear | 15435
- request unified-edge tdf call-trace show | 15436
- request unified-edge tdf call-trace stop | 15443

request unified-edge tdf call-trace stop

IN THIS SECTION

- Syntax | 15444
- Description | 15444
- Options | 15444
- Required Privilege Level | 15444
- Output Fields | 15444
- Sample Output | 15445
- Release Information | 15445

Syntax

```
request unified-edge tdf call-trace stop
<all>
<identifier call-trace-identifier>
```

Description

Stop the previously configured subscriber call tracing on one or more TDF gateways.

Options

none	(Same as all) Stop all subscriber call tracing.
all	(Optional) Stop all subscriber call tracing.
identifier *call-trace-identifier*	(Optional) Stop call tracing for the specified call trace identifier.

Required Privilege Level

unified-edge

Output Fields

Table 134 on page 15444 lists the output fields for the `request unified-edge tdf call-trace stop` command. Output fields are listed in the approximate order in which they appear.

Table 134: request unified-edge tdf call-trace stop Output Fields

Field Name	Field Description
`Session PIC`	Identifier of session PIC for which the call trace status is displayed.
`Status`	Status of the call trace: • `success`—Call trace stopped successfully. • `fail`—Call tracing cannot be stopped.

Sample Output

request unified-edge tdf call-trace stop

```
user@host> request unified-edge tdf call-trace stop
       Session PIC            Status
       ms-0/1/0             success
       ms-1/1/0             success
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- request unified-edge tdf call-trace clear | 15435
- request unified-edge tdf call-trace show | 15436
- request unified-edge tdf call-trace start | 15440

request validation policy

IN THIS SECTION

- Syntax | 15446
- Description | 15446
- Options | 15446
- Required Privilege Level | 15446
- Output Fields | 15446
- Sample Output | 15447
- Release Information | 15447

Syntax

```
request validation policy
<instance instance-name>
<logical-system logical-system-name>
<record ip-prefix>
```

Description

When BGP origin validation is configured, manually request a route validation record policy to be reevaluated. This command causes dependent route validation records to be reevaluated. Dependent route validation records are exactly matching and more specific records.

Options

none	Request a policy reevaluation of all dependent route validation records.
instance ***instance-name***	(Optional) Request a policy reevaluation of all dependent route validation records for the specified routing instance. The instance name can be primary for the main instance, or any valid configured instance name or its prefix.
logical-system ***logical-system-name***	(Optional) Perform this operation on a particular logical system.
record ***ip-prefix***	(Optional) Request a policy reevaluation of all route validation records that match a given prefix.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request validation policy

```
user@host> request validation policy
  Enqueued 1 IPv4 records
  Enqueued 0 IPv6 records
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

request virtual-chassis device-reachability

IN THIS SECTION

- Syntax | 15448
- Description | 15448
- Options | 15448
- Required Privilege Level | 15449
- Output Fields | 15449
- Sample Output | 15450
- Release Information | 15450

Syntax

```
request virtual-chassis device-reachability test-name name  (source-fpc source-fpc-id | source-
ip-address source-ip-address) (destination-device device-id | destination-fpc destination-fpc-id
| destination-ip-address destination-ip-address)
<probe-count count>
<probe-interval interval>
<probe-pattern pattern>
<test-count count>
<test-interval interval>
```

Description

Run a Virtual Chassis device reachability test. A Virtual Chassis device reachability test sends ping packets from one member of a Virtual Chassis to another member of a Virtual Chassis directly through the Virtual Chassis connections. The receiving Virtual Chassis member sends reply packets to confirm receipt of the ping packets from the sending device. The results of the test immediately provide information that is helpful in verifying connectivity between Virtual Chassis member devices.

Options

test-name *name*	Specify the name of the test. The name can be any single-word character string. Spaces are not allowed.
source-fpc *source-fpc-id*	Specify the FPC that sends the first ping message.
source-ip-address *source-ip-address*	Specify the source IP address that receives the ping message.
destination-device *device-id*	Specify the device in the Virtual Chassis that receives the ping message.
destination-fpc *destination-fpc-id*	Specify the FPC that receives the first ping message.
destination-ip-address *destination-ip-address*	Specify the destination IP address that receives the ping message.
probe-count *count*	(Optional) Specify the number of ping messages to send for the test. The count can be any number from 1 through 10. The default count is 5.
probe-interval *interval*	(Optional) Specify the time in seconds between the ping messages that are sent during the test. The interval can be any number from 1 through 3. The default interval is 1 second.

probe-pattern *pattern* (Optional) Specify a payload pattern in the ping message. Enter the pattern in hexadecimal format.

test-count *count* (Optional) Specify the number of times to run the test. The count can be any number from 1 through 3. The default count is 1.

test-interval *interval* (Optional) Specify the time in seconds between tests when multiple tests are specified. The interval can be any number from 1 through 3. The default interval is 1.

Required Privilege Level

system-control

Output Fields

Table 135 on page 15449 lists the output fields for the `request virtual chassis device reachability` command. Output fields are listed in the approximate order in which they appear.

Table 135: request virtual chassis device reachability Output Fields

Field Name	Field Description
Test Name	The name of the test.
Performing Test	The test number.
Session ID	The session ID. This output is always **0** and not useful for Virtual Chassis device reachability tests.
Packets Count/Sent/Received/Sendfail	The number of total ping packets sent, including the total number of ping packets counted (**Count**), sent (**Sent**), received (**Received**), and packets that could not be sent (**Sendfail**).
Unknown/Timedout/Duplicate packets received	The failed pings, including the number of ping packets that failed for an unknown reason (**Unknown**), timed out (**Timed Out**), and the number of duplicate received ping packets (**Duplicate packets received**).

Table 135: request virtual chassis device reachability Output Fields *(Continued)*

Field Name	Field Description
Round-trip Min/Avg/Max	The average round-trip ping time, including the shortest ping time (**Min**), the average ping time (**Avg**), and the maximum ping time (**Max**).

Sample Output

request virtual-chassis device-reachability test-name member0-to-member2 source-fpc 0 destination-fpc 2

```
user@switch> request virtual-chassis device-reachability test-name member0-to-member2 source-fpc
0 destination-fpc 2
Device Reachability Statistics:
Test Name                   : member0-to-member2

Performing Test: 1

56 bytes from 0: session-id 0 seq-id 0
56 bytes from 0: session-id 0 seq-id 1
56 bytes from 0: session-id 0 seq-id 2
56 bytes from 0: session-id 0 seq-id 3
56 bytes from 0: session-id 0 seq-id 4
--- sping statistics ---
Session ID                                  : 0
Packets Count/Sent/Received/Sendfail        : 5/5/5/0
Unknown/Timedout/Duplicate packets received : 0/0/0
Round-trip Min/Avg/Max                      : 181/807/3010 usec
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Verifying Connectivity Between Virtual Chassis Member Devices

request virtual-chassis member-id delete (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 15451
- Description | 15451
- Required Privilege Level | 15451
- Sample Output | 15452
- Release Information | 15452

Syntax

```
request virtual-chassis member-id delete
<force>
```

Description

Remove (**delete**) the member ID from a router or switch that you want to remove from a Virtual Chassis configuration.

> **NOTE**: Issuing the command to remove the member ID causes the device to reboot, and requires you to confirm that you want to proceed with this operation. If you do not confirm the operation, the software cancels the command. You can alternatively use the `force` option to reboot without requiring confirmation.

Required Privilege Level

system-control

Sample Output

request virtual-chassis member-id delete

```
user@host> request virtual-chassis member-id delete
This command will disable virtual-chassis mode and reboot the system.
 Continue? [yes,no] (no)
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Deleting Member IDs in a Virtual Chassis Configuration

Example: Deleting a Virtual Chassis Configuration for MX Series 5G Universal Routing Platforms

request virtual-chassis member-id set

IN THIS SECTION

- Syntax | 15453
- Syntax (MX960, MX2010, and MX2020 Routers) | 15453
- Description | 15453
- Options | 15453
- Required Privilege Level | 15454
- Sample Output | 15454
- Release Information | 15454

Syntax

```
request virtual-chassis member-id set member member-id
```

Syntax (MX960, MX2010, and MX2020 Routers)

```
request virtual-chassis member-id set member member-id
<slots-per-chassis slot-count>
```

Description

Assign (`set`) a member ID and, optionally, a slot count to a router or switch that you want to add as a member of a Virtual Chassis configuration.

> **NOTE**: Issuing the `request virtual-chassis member-id set` command causes the device to reboot, and requires you to confirm that you want to proceed with this operation. If you do not confirm the operation, the software cancels the command. After the reboot all MPCs remain powered off until the Virtual Chassis port connection is configured.

Options

member *member-id*

Assign the numeric value that identifies a member router or switch in a Virtual Chassis configuration. When you assign a member ID to a router or switch, assign the same member ID defined for this device in the preprovisioned configuration. Replace *member-id* with the value 0 or 1.

slots-per-chassis *slot-count*

(MX960, MX2010, and MX2020 routers) (Optional) Identify the number of chassis slots in the Virtual Chassis member router. To ensure that a Virtual Chassis consisting of an MX2020 member router and either an MX960 or MX2010 member router forms properly, you must explicitly set the *slot-count* value for the MX960 router or MX2010 router to 20 to match the slot count of the MX2020 router.

- Values:

  The valid values for *slot-count* are as follows:

- MX960 router: 12 or 20
- MX2010 router: 12 or 20
- MX2020 router: 20

- Default:

  The default values for *slot-count* are as follows:

  - MX960 router: 12
  - MX2010 router: 12
  - MX2020 router: 20

Required Privilege Level

system-control

Sample Output

request virtual-chassis member-id set (Assigning a Member ID)

```
user@host> request virtual-chassis member-id set member 0
This command will enable virtual-chassis mode and reboot the system.
 Continue? [yes,no] (no)
```

request virtual-chassis member-id set (Assigning a Member ID and Slot Count)

```
user@host> request virtual-chassis member-id set member 1 slots-per-chassis 20
This command will enable virtual-chassis mode and reboot the system.
 Continue? [yes,no] (no)
```

Release Information

Command introduced in Junos OS Release 11.2.

`slots-per-chassis` option added in Junos OS Release 15.1 for MX960 routers, MX2010 routers, and MX2020 routers.

RELATED DOCUMENTATION

Configuring Member IDs for a Virtual Chassis

Configuring an EX9200 Virtual Chassis

Example: Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis

request virtual-chassis mode

IN THIS SECTION

- Syntax | 15455
- Description | 15456
- Options | 15459
- Required Privilege Level | 15459
- Sample Output | 15460
- Sample Output | 15460
- Sample Output | 15460
- Release Information | 15460

Syntax

```
request virtual-chassis mode
    fabric
    mixed <ieee-clause-82>
    hgoe
    network-port
    <disable>
    <reboot>
    <all-members> | <local> | <member member-id>
```

Description

Enable special device modes when operating as part of a Virtual Chassis or Virtual Chassis Fabric (VCF). Supported device modes include:

- Fabric mode—Enable the device to operate as part of a VCF. You must set `fabric` mode on all of the devices in a VCF.
- Mixed mode—Enable the device to operate in a mixed Virtual Chassis or VCF.
- (EX4100 and EX4400 switches only) Network port mode—Enable the default Virtual Chassis ports (VCPs) to operate as network ports instead of as VCPs.

You must reboot the standalone device, Virtual Chassis or VCF for this command to take effect, and the setting persists across subsequent reboot operations. You can use this command's `reboot` option to automatically reboot the device when setting or disabling a device mode.

> **NOTE**: For devices in a mixed VCF, you can set both mixed and fabric modes in the same command and then reboot the device once for both settings to take effect.

Mixed Mode Requirements and Best Practices

You must set `mixed` mode on all devices in a Virtual Chassis or a VCF when the devices have differences in how they interoperate. See Understanding Mixed EX Series and QFX Series Virtual Chassis or *Understanding Mixed Virtual Chassis Fabric* for details on which devices can be combined to form a mixed Virtual Chassis or VCF.

Some combinations of different types of switches form non-mixed Virtual Chassis or VCFs because the devices can run the same software images. You *don't* need to configure mixed mode if the only devices in your Virtual Chassis are the following combinations of switches:

- Only EX4500 and EX4550 switches in an EX Series Virtual Chassis.
- EX2300 multigigabit switches (EX2300-24MP and EX2300-48MP models) with any other EX2300 model switches in an EX2300 Virtual Chassis.
- EX4100 multigigabit switches (EX4100-24MP and EX4100-48MP models) with any other EX4100 model switches in an EX4100 Virtual Chassis.
- EX4400 multigigabit switches (EX4400-24MP and EX4400-48MP models) with any other EX4400 model switches in an EX4400 Virtual Chassis.
- Only QFX3500 and QFX3600 switches in a QFX Series Virtual Chassis.
- QFX5110 and QFX5100 switches in a QFX5110 Virtual Chassis or QFX5110 VCF.

However, EX4300 switches that are not multigigabit models combined with EX4300 multigigabit model (EX4300-48MP) switches form a mixed EX4300 Virtual Chassis. In this case, you must set mixed mode on all member switches. In addition, when you set mixed mode on the EX4300 non-multigigabit model member switches, you must include the `ieee-clause-82` option with the mode command. This option sets a special port mode (IEEE Clause 82) on the Virtual Chassis ports (VCPs) that enables them to communicate when interconnected with VCPs on EX4300 multigigabit switches.

> **NOTE**: If you remove an EX4300 member switch from a mixed EX4300 Virtual Chassis with multigigabit model members, you must disable `ieee-clause-82` port mode on the switch if you want to reconfigure it as a standalone switch or use it in any other type of mixed Virtual Chassis or any non-mixed Virtual Chassis. Otherwise, the VCPs will not connect with other members in the new Virtual Chassis. To do this, use the `request virtual-chassis mode mixed ieee-clause-82 disable` command and then reboot the switch. To add the switch to any other type of mixed Virtual Chassis, enable mixed mode again without this option and reboot the switch again for the change to take effect.

Best practices when setting `mixed` mode include:

- Don't enable mixed mode on a standalone device or for a member switch that is in a non-mixed Virtual Chassis or VCF. Mixed mode reduces the maximum scaling numbers for some features on the standalone switch, Virtual Chassis, or VCF to accommodate different types of devices that can be interconnected.

- To avoid potential traffic disruptions and configuration issues, on each device we recommend setting mixed mode for a mixed Virtual Chassis or VCF before cabling them into a mixed Virtual Chassis or VCF. You must then reboot the device, Virtual Chassis or VCF for the mode change to take effect.

  > **NOTE**: The same best practice applies to setting fabric mode for a VCF—set this mode on the member devices before interconnecting them.

- You can change mode settings manually after a device has been added to a Virtual Chassis or VCF if you determine the settings are not the same on all the devices.

  However, if you set only some of the devices in a mixed Virtual Chassis or VCF to mixed mode, the mixed Virtual Chassis or VCF might not form. If you experience this issue, enter the `request virtual-chassis mode mixed all-members` command to force setting mixed mode on all devices in the Virtual Chassis or VCF. (Include the `fabric` mode option as well for a VCF.) You can then reboot the entire Virtual Chassis or VCF, or only the devices where the mode changed. The Virtual Chassis or VCF forms when the rebooted devices are online again.

> **NOTE**: On switches running Junos OS Release 11.4 or later, if you don't set any of the switches in a mixed EX4200 and EX4500 Virtual Chassis to mixed mode, a mixed EX4200 and EX4500 Virtual Chassis forms where one of the switches assumes the primary role, and all other switches take on the linecard role. If you see this behavior, enter the `request virtual-chassis mode mixed all-members` command to set all members to mixed mode and reboot the Virtual Chassis.

Network Port Mode (EX4100 and EX4400 Switches Only)

If you are not using an EX4100 or EX4400 switch as a member switch in a Virtual Chassis, you can use this command with the `network-port` option to enable network port mode, which converts the default VCPs on the switch into network ports. By default, the two 100-Gigabit Ethernet QSFP28 ports on the rear panel are set as two logical 50-Gbps VCP interfaces each, forming a total of four logical 50-Gbps VCPs. (The default VCPs are in PIC slot 1, so the VCP interfaces are named vcp-255/1/0 through vcp-255/1/3.) If you convert these ports into network port mode, they become two 100 Gigabit Ethernet network ports, or you can channelize each of them into four 25-Gigabit Ethernet or four 10-Gigabit Ethernet network ports.

You can't use one of the rear panel 100-Gigabit Ethernet ports as a VCP and the other as a network port. When you use this command to enable network port mode, the mode affects both 100-Gbps ports together.

You can use the `show virtual-chassis mode` command to see the current port mode and what the port mode will be the next time the switch is rebooted. The two values might be different if you changed the mode but have not rebooted the switch yet.

To disable network port mode and return these ports to their default settings as VCPs, use the `network-port` and `disable` options with the `request virtual-chassis mode` command. You must reboot the switch for network port mode changes to take effect, so you can include the reboot option in the same command. For example:

```
request virtual-chassis mode network-port disable reboot
```

You can use the `request virtual-chassis vc-port delete` command to disable one of these ports as a VCP, but that action doesn't change the port mode to network port mode. You must enter the `request virtual-chassis mode network-port` command (and reboot the switch) to use the port as a network port. Also, if you run the `request virtual-chassis vc-port delete interface` *`interface-name`* command to disable vcp-255/1/0, that command disables both logical ports 0 and 1 (vcp-255/1/0 and vcp-255/1/1). Similarly, specifying to disable vcp-255/1/2 disables both logical ports 2 and 3 (vcp-255/1/2 and vcp-255/1/3).

Options

all-members (Optional) Set the Virtual Chassis mode for all members of the Virtual Chassis or VCF.

disable Disable the specified Virtual Chassis mode setting if it was previously enabled. You must include at least one of the supported modes that you want to disable with this option: `fabric`, `mixed`, or `network-port`.

The `request virtual-chassis mode mixed disable` command does not also automatically disable the mixed mode `ieee-clause-82` port mode option. To disable `ieee-clause-82` port mode on the device, you must include that option in the disable command as follows:

```
request virtual-chassis mode mixed ieee-clause-82 disable
```

fabric Set the device into fabric mode so it can participate as a member device in a VCF.

hgoe (EX4400-24X) Configure the QSFP28 ports as VCPs in HGoE mode.

local (Optional) Set the Virtual Chassis mode only on the member device where the command is issued.

member *member-id* (Optional) Set the Virtual Chassis mode on the specified member of the Virtual Chassis or VCF.

mixed <ieee-clause-82> Set the device into mixed mode so that the device can participate in a mixed Virtual Chassis or mixed VCF.

You must set mixed mode with the `ieee-clause-82` option on EX4300 switches that are not multigigabit models when mixing them with EX4300 multigigabit (EX4300-48MP) model switches in an EX4300 Virtual Chassis.

network-port (EX4400 switches only) Set an EX4400 switch into network port mode, which converts the default VCPs on the switch to network port mode so you can use them as network ports instead of VCPs.

reboot After applying the mode change specified by the other options, reboot the device automatically. You must reboot the device for any mode change to take effect.

Required Privilege Level

system-control

Sample Output

request virtual-chassis mode mixed

```
user@switch> request virtual-chassis mode mixed
```

Sample Output

request virtual-chassis mode fabric mixed reboot

```
user@switch> request virtual-chassis mode fabric mixed reboot
```

Sample Output

request virtual-chassis mode mixed (with IEEE Clause 82 mode for EX4300 switches mixed with EX4300 multigigabit switches in a Virtual Chassis)

```
user@switch> request virtual-chassis mode ieee-clause-82 mixed
```

Release Information

Command introduced in Junos OS Release 11.1.

`fabric` mode introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices in a Virtual Chassis Fabric (VCF).

`ieee-clause-82` mixed-mode option introduced for EX4300 switches mixed with EX4300 multigigabit switches in a Virtual Chassis in Junos OS Release 18.2R1.

`network-port` mode introduced in Junos OS Release 21.1R1 for EX4400 switches in a Virtual Chassis.

RELATED DOCUMENTATION

- Understanding Mixed EX Series and QFX Series Virtual Chassis
- *Understanding Mixed Virtual Chassis Fabric*
- Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)

Configuring an EX2300, EX3400, or EX4300 Virtual Chassis
Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis
Configuring an EX4650 or a QFX Series Virtual Chassis
Adding a New Switch to an Existing EX2300, EX3400, or EX4300 Virtual Chassis
Removing or Replacing a Member Switch of a Virtual Chassis Configuration
Verifying the Virtual Chassis Fabric Mode Settings
Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member

request virtual-chassis reactivate

IN THIS SECTION

- Syntax | **15461**
- Description | **15461**
- Required Privilege Level | **15461**
- Sample Output | **15462**
- Release Information | **15462**

Syntax

```
request virtual-chassis reactivate
```

Description

Reactivate a device that has been assigned a member ID but is not currently connected to the Virtual Chassis or VCF.

You can use this command to reactivate a device that was previously part of the Virtual Chassis or VCF but whose status is no longer **Prsnt**.

Required Privilege Level

system-control

Sample Output

request virtual-chassis reactivate

```
user@switch> request virtual-chassis reactivate
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

- Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member
- Verifying the Member ID, Role, and Neighbor Member Connections of an EX8200 Virtual Chassis Member

request virtual-chassis recycle

IN THIS SECTION

- Syntax | 15462
- Description | 15463
- Options | 15463
- Required Privilege Level | 15463
- Sample Output | 15463
- Sample Output | 15463
- Release Information | 15463

Syntax

```
request virtual-chassis recycle member-id member-id
```

Description

Make a previously used member ID available for reassignment.

When you remove a member switch from the Virtual Chassis configuration, the primary reserves that member ID. To make the member ID available for reassignment, you must use this command.

NOTE: You must run this command from the Virtual Chassis member in the primary role.

Options

member-id *member-id* Specify the member ID that you want to make available for reassignment to a different member.

Required Privilege Level

system-control

Sample Output

request virtual-chassis recycle member-id 3

```
user@switch> request virtual-chassis recycle member-id 3
```

Sample Output

request virtual-chassis recycle member-id 1

```
user@external-routing-engine> request virtual-chassis recycle member-id 1
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

request virtual-chassis renumber | 15464

Removing or Replacing a Member Switch of a Virtual Chassis Configuration

request virtual-chassis renumber

IN THIS SECTION

- Syntax | 15464
- Description | 15464
- Options | 15464
- Required Privilege Level | 15465
- Sample Output | 15465
- Release Information | 15465

Syntax

```
request virtual-chassis renumber member-id old-member-id new-member-id
```

Description

Renumber a member of a Virtual Chassis configuration.

NOTE: You must run this command from the Virtual Chassis member in the primary role.

Options

member-id *old-member-id*	Specify the ID of the member that you wish to renumber.
new-member-id *new-member-id*	Specify an unassigned member ID.

Required Privilege Level

system-control

Sample Output

request virtual-chassis renumber member-id 5 new-member-id 4

```
user@switch> request virtual-chassis renumber member-id 5 new-member-id 4
```

request virtual-chassis renumber member-id 1 new-member-id 0

```
user@external-routing-engine> request virtual-chassis renumber member-id 1 new-member-id 0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

request virtual-chassis recycle | 15462

Removing or Replacing a Member Switch of a Virtual Chassis Configuration

request virtual-chassis routing-engine master switch

IN THIS SECTION

- Syntax | 15466
- Description | 15466
- Options | 15467
- Required Privilege Level | 15467
- Sample Output | 15467

Release Information | 15468

Syntax

```
request virtual-chassis routing-engine master switch
<check>
```

Description

Change the primary role in an MX Series Virtual Chassis or EX9200 Virtual Chassis by switching the global roles of the primary router or switch and backup router or switch in the Virtual Chassis configuration. The `request virtual-chassis routing-engine master switch` command must be issued from the primary router or switch (VC-Pp).

(MX Series routers only) The local roles (`master` and `standby`, or `m` and `s`) of the Routing Engines in the Virtual Chassis primary router change after a global switchover, but the local roles of the Routing Engines in the Virtual Chassis backup router do not change. For example, the primary Routing Engine in the Virtual Chassis primary router (VC-Pp) becomes the standby Routing Engine in the Virtual Chassis backup router (VC-Bs) after the global switchover. By contrast, the primary Routing Engine in the Virtual Chassis backup router (VC-Bp) remains the primary Routing Engine in the Virtual Chassis primary router (VC-Pp) after the global switchover.

NOTE: Before you issue the `request virtual-chassis routing-engine master switch` command from the primary router or switch in the Virtual Chassis, make sure that the system configuration is synchronized between the primary and backup router or switch. If the configuration is not synchronized, or if you attempt to issue the `request virtual-chassis routing-engine master switch` command from the backup router or switch instead of from the primary router or switch, the device displays an error message and rejects the command.

If you issue the `request virtual-chassis routing-engine master switch` command when the Virtual Chassis is in a transition state (for example, the backup router or switch is disconnecting from the Virtual Chassis), the device does not process the command.

Options

check (Optional) Perform a check from the primary router or switch to determine whether the member routers or switches are ready for GRES from a database synchronization perspective, without initiating the GRES operation itself.

Required Privilege Level

system-control

Sample Output

request virtual-chassis routing-engine master switch (From Primary Router)

```
{master:member0-re0}
user@host> request virtual-chassis routing-engine master switch
Do you want to continue ? [yes,no] (no)
```

request virtual-chassis routing-engine master switch (Error When Configuration Not Synchronized)

```
{master:member0-re0}
user@host> request virtual-chassis routing-engine master switch
Error: mastership switch request NOT honored, backup not ready
```

request virtual-chassis routing-engine master switch (Error When Run from Backup Router)

```
{backup:member1-re0}
user@host> request virtual-chassis routing-engine master switch
error: Virtual Chassis member is not the protocol master
```

request virtual-chassis routing-engine master switch check (Ready for GRES)

```
{master:member0-re0}
user@host> request virtual-chassis routing-engine master switch check
Switchover Ready
```

request virtual-chassis routing-engine master switch check (Not Ready for GRES)

```
{master:member0-re0}
user@host> request virtual-chassis routing-engine master switch check
error: chassisd Not ready for mastership switch, try after 217 secs.
mastership switch request NOT honored, backup not ready
```

Release Information

Command introduced in Junos OS Release 11.2.

Option `check` introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Switching the Global Primary and Backup Roles in a Virtual Chassis Configuration

Determining GRES Readiness in a Virtual Chassis Configuration

Switchover Behavior in an MX Series Virtual Chassis

Primary-role Election in a Virtual Chassis

request virtual-chassis vc-port

IN THIS SECTION

- Syntax | 15469
- Description | 15469
- Options | 15470

- Required Privilege Level | 15471
- Sample Output | 15471
- Release Information | 15471

Syntax

```
request virtual-chassis vc-port (set | delete)
<fpc-slot fpc-slot> pic-slot pic-slot port port-number
interface interface-name
<member  | local>
```

member-id

Description

Set a port to operate as a Virtual Chassis port (VCP), or delete the VCP setting on a port. See Virtual Chassis Port Options for details on which ports you can set as VCPs on different switches. After setting a port as a VCP, you can't use the port for any other purpose unless you remove the VCP setting.

You can specify the port using the `interface` option or the `pic-slot` and `port` options.

If you don't include the `member` *`member-id`* option, this command defaults to setting the port as a VCP or deleting the VCP setting on the switch where you run the command. You can alternatively set the `local` option to ensure the command applies to the specified port on the local switch where you run it.

You might experience a temporary traffic disruption immediately after creating or deleting a user-configured VCP in an EX8200 Virtual Chassis.

Some switches have ports that are dedicated VCPs (you can only use them as VCPs and for no other purpose) or that are configured as VCPs by default (set as VCPs in the default factory configuration). You do not need to explicitly set those ports as VCPs to use them to interconnect the switch into a Virtual Chassis. See Virtual Chassis Port Options for details.

If you don't need a default-configured VCP to interconnect Virtual Chassis member switches, you can run this command with the `delete` option to remove the VCP setting and use the port as a regular uplink or network port. If you want to use the port as a VCP again later, then you need to run this command with the `set` option to reapply the VCP setting.

Special Behavior for VCPs on EX4400 Switches

On EX4400 switches, you don't use this command to convert ports on the switch into VCPs. EX4400 switches have default VCPs which are the only ports you can use as VCPs on the switch. These are the two 100-Gbps QSFP28 ports on the rear panel in PIC slot 1, which are set by default as four logical 50-Gbps VCP interfaces (vcp-255/1/0 through vcp-255/1/3). You can convert them into network ports by enabling network port mode on the switch (`request virtual-chassis mode network-port <reboot>`). In that case, to subsequently use the switch in a Virtual Chassis, you must disable network port mode (`request virtual-chassis mode network-port disable <reboot>`), which converts those ports back into VCPs. (Port mode changes apply to all VCPs on the switch,)

You can enter the `request virtual-chassis vc-port delete` command to disable these ports as VCPs, but, as mentioned above, that action doesn't convert them into network ports.

> **NOTE**: On EX4400 switches, if you use the `request virtual-chassis vc-port delete interface` *`interface-name`* command to disable logical VCP interface vcp-255/1/0, the switch disables both logical ports 0 and 1 (vcp-255/1/0 and vcp-255/1/1). Similarly, if you disable logical VCP interface vcp-255/1/2, that action disables both logical ports 2 and 3 (vcp-255/1/2 and vcp-255/1/3).

Options

set	Set a port as a VCP to convert an uplink or network port into a VCP.
delete	Delete the VCP setting on a port to convert a VCP back into an uplink or network port.
pic-slot *pic-slot*	Number of the PIC slot for the port on the switch.
port *port-number*	Number of the port that you want to enable or disable as a VCP.
interface *interface-name*	Interface name of the port that you want to enable or disable as a VCP. You can use this option to specify the port instead of using the `pic-slot` and `port` options.
member *member-id*	(Optional) Enable or disable the specified VCP on the specified member of the Virtual Chassis or VCF.
local	(Optional) Enable or disable the specified VCP on the local switch where you run the command.

Required Privilege Level

system-control

Sample Output

request virtual-chassis vc-port set pic-slot 1 port 0

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port 0
```

To check the results of this command, use the "show virtual-chassis vc-port" on page 27681 command.

request virtual-chassis vc-port set pic-slot 1 port 1 member 3

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port 1 member 3
```

To check the results of this command, use the "show virtual-chassis vc-port" on page 27681 command.

request virtual-chassis vc-port delete pic-slot 1 port 1 member 3

```
user@switch> request virtual-chassis vc-port delete pic-slot 1 port 1 member 3
```

To check the results of this command, use the "show virtual-chassis vc-port" on page 27681 command.

Release Information

Command introduced in Junos OS Release 9.0.

Option `fpc-slot` introduced in Junos OS Release 10.4 for EX Series switches.

RELATED DOCUMENTATION

- request virtual-chassis vc-port (Dedicated VCP)
- show virtual-chassis vc-port | **27681**
- show virtual-chassis vc-port statistics | **27713**
- clear virtual-chassis vc-port statistics | **4283**
- Virtual Chassis Port (VCP) Interface Names in an EX8200 Virtual Chassis

Understanding Virtual Chassis Components

request virtual-chassis vc-port (Dedicated VCP)

IN THIS SECTION

- Syntax | 15472
- Description | 15472
- Options | 15472
- Additional Information | 15473
- Required Privilege Level | 15473
- Sample Output | 15473
- Release Information | 15474

Syntax

```
request virtual-chassis vc-port set interface vcp-interface-name
<member member-id> <disable>
```

Description

Disable or enable a Virtual Chassis port (VCP) on a dedicated VCP.

Configure a Gigabit Ethernet link on an EX8200 switch as a VCP in an EX8200 Virtual Chassis. You can use this command only to configure a link between two EX8200 member switches as a VCP link; all other links in an EX8200 Virtual Chassis are automatically VCP links.

Options

interface ***vcp-interface-name***

Name of the interface to enable or disable.

If you omit **member** ***member-id*** in an EX4200 Virtual Chassis, EX4500 Virtual Chassis, EX4550 Virtual Chassis, mixed Virtual Chassis, this command defaults to disabling or

enabling the dedicated VCP on the switch where the command is issued. The dedicated VCPs are enabled in the factory default configuration.

member *member-id*
(Optional on EX4200, EX4500, and EX4550 switches; required on EX8200 switches) Enable or disable the specified VCP on the specified member of the Virtual Chassis configuration.

This option must be specified when using this command to configure a VCP link between two EX8200 switches.

disable
(Optional) Disable the specified VCP. If you omit this keyword, the command enables the dedicated VCP.

Additional Information

Required Privilege Level

system-control

Sample Output

request virtual-chassis vc-port set interface vcp-0 disable (EX4200 Virtual Chassis)

```
user@switch> request virtual-chassis vc-port set interface vcp-0 disable
```

To check the results of this command, use the "show virtual-chassis vc-port" on page 27681 command.

request virtual-chassis vc-port set fpc-slot 4 pic-slot 0 port 2 member 1 (EX8200 Virtual Chassis)

```
user@external-routing-engine>
 request virtual-chassis vc-port set fpc-slot 4 pic-slot 0 port 2 member 1
```

To check the results of this command, use the "show virtual-chassis vc-port" on page 27681 command.

request virtual-chassis vc-port set interface vcp-0 member 3 disable (EX8200 Virtual Chassis)

```
user@external-routing-engine> request virtual-chassis vc-port set interface vcp-0 member 3
disable
```

To check the results of this command, use the "show virtual-chassis vc-port" on page 27681 command.

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- request virtual-chassis vc-port | 15468
- show virtual-chassis vc-port | 27681
- show virtual-chassis vc-port statistics | 27713
- clear virtual-chassis vc-port statistics | 4283
- *Understanding Virtual Chassis Components*

request virtual-chassis vc-port delete (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 15475
- Description | 15475
- Options | 15475
- Required Privilege Level | 15476
- Sample Output | 15476
- Release Information | 15476

Syntax

```
request virtual-chassis vc-port delete fpc-slot fpc-slot-number pic-slot pic-slot-number port
port-number
<(local | member member-id)>
```

Description

Remove (**delete**) a Virtual Chassis port from a member router in an MX Series Virtual Chassis configuration. After a Virtual Chassis port is created, it is renamed **vcp-*slot/pic/port***, and is no longer available for configuration as a standard network port. After you remove a Virtual Chassis port, it becomes available to the global configuration and can again function as a standard network port.

> **NOTE**: If the member ID has not been set on the router where you issue the `request virtual-chassis vc-port delete` command, the software prevents the removal of the Virtual Chassis port on the router. To set the member ID, use the `request virtual-chassis member-id set` command.

Options

fpc-slot *fpc-slot-number*

Number of the Flexible PIC Concentrator (FPC) slot on which the Virtual Chassis port resides. The slot number corresponds to the Modular Port Concentrator (MPC) slot number. Replace ***fpc-slot-number*** with a value appropriate for your router:

- MX960 router—0 through 11.
- MX480 router—0 through 5.
- MX240 router—0 through 2.

pic-slot *pic-slot-number*

Number of the PIC slot on which the Virtual Chassis port resides. Replace ***pic-slot-number*** with a value in the range 0 through 3.

port *port-number*

Number of the port on the PIC on which the Virtual Chassis port resides. Replace ***port-number*** with a value appropriate for your PIC.

local

(Optional) Delete the Virtual Chassis port on the member router on which you are issuing the command. This is the default behavior if you do not specify the **local** or **member** options.

member *member-id* (Optional) Numeric value that identifies the remote Virtual Chassis member on which you want to delete the Virtual Chassis port. Replace ***member-id*** with the value 0 or 1.

Required Privilege Level

system-control

Sample Output

request virtual-chassis vc-port delete (Remove vcp-3/2/1)

```
user@host> request virtual-chassis vc-port delete fpc-slot 3 pic-slot 2 port 1
vc-port successfully deleted
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Deleting Virtual Chassis Ports in a Virtual Chassis Configuration

Example: Deleting a Virtual Chassis Configuration for MX Series 5G Universal Routing Platforms

request virtual-chassis vc-port diagnostics optics

IN THIS SECTION

- Syntax | 15477
- Description | 15477
- Required Privilege Level | 15477
- Sample Output | 15477
- Release Information | 15477

Syntax

```
request virtual-chassis vc-port diagnostics optics
```

Description

Run a digital optical monitoring (DOM) scan on the optical ports configured as Virtual Chassis ports (VCPs).

Enter the `show virtual-chassis vc-port diagnostics optics` command to view the results of the diagnostic scan.

On certain EX Series switches, you must enter the `request virtual-chassis vc-port diagnostics optics` command to run a diagnostic scan before you can gather the `show virtual-chassis vc-port diagnostics optics` output.

Required Privilege Level

system-control

Sample Output

request virtual-chassis vc-port diagnostics optics

```
user@switch> request virtual-chassis vc-port diagnostics optics
fpc0:
--------------------------------------------------------------------------
vc-port Diagnostics Optics Done
```

Release Information

Command introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

show virtual-chassis vc-port diagnostics optics | 27688

request virtual-chassis vc-port set (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 15478
- Description | 15478
- Options | 15479
- Required Privilege Level | 15479
- Sample Output | 15480
- Release Information | 15480

Syntax

```
request virtual-chassis vc-port set fpc-slot fpc-slot-number pic-slot pic-slot-number port port-number
<(local | member member-id)>
```

Description

Create (**set**) a Virtual Chassis port on an MX Series router through which the router connects to other member routers in the Virtual Chassis. You can create Virtual Chassis ports only on Modular Port Concentrator/Modular Interface Card (MPC/MIC) network ports on MX Series routers.

After a Virtual Chassis port is created, it is renamed **vcp-*slot/pic/port***, and is no longer available for configuration as a standard network port. Virtual Chassis ports can be used only to interconnect the MX Series routers in the Virtual Chassis.

NOTE: If the member ID has not been set on the router where you issue the `request virtual-chassis vc-port set` command, the software prevents the creation of the Virtual Chassis port on the router. To set the member ID, use the `request virtual-chassis member-id set` command.

Options

fpc-slot ***fpc-slot-number***

Number of the Flexible PIC Concentrator (FPC) slot on which the Virtual Chassis port resides. The slot number corresponds to the Modular Port Concentrator (MPC) slot number. Replace ***fpc-slot-number*** with a value appropriate for your router:

- MX960 router—0 through 11.
- MX480 router—0 through 5.
- MX240 router—0 through 2.

When you issue the `show interfaces` command on a member router in an MX Series Virtual Chassis, the FPC slot number displayed in the command output reflects the FPC slot numbering and offset used in the Virtual Chassis instead of the physical slot number where the FPC is actually installed. The router with member ID 0 in the Virtual Chassis uses FPC slot numbers 0 through 11 with no offset, and the router with member ID 1 uses FPC slot numbers 12 through 23, with an offset of 12. For example, a 10-Gigabit Ethernet interface that appears as **xe-14/2/2** (FPC slot 14, PIC slot 2, port 2) in the `show interfaces` command is actually interface **xe-2/2/2** (FPC slot 2, PIC slot 2, port 2) on member ID 1 after deducting the FPC slot numbering offset of 12 for member ID 1.

pic-slot ***pic-slot-number***

Number of the PIC slot on which the Virtual Chassis port resides. Replace ***pic-slot-number*** with a value in the range 0 through 3.

port ***port-number***

Number of the port on the PIC on which the Virtual Chassis port resides. Replace ***port-number*** with a value appropriate for your PIC.

local

(Optional) Set the Virtual Chassis port on the member router on which you are issuing the command. This is the default behavior if you do not specify the **local** or **member** options.

member ***member-id***

(Optional) Numeric value that identifies the remote Virtual Chassis member on which you want to create the Virtual Chassis port. Replace ***member-id*** with the value 0 or 1.

Required Privilege Level

system-control

Sample Output

request virtual-chassis vc-port set (No Existing Network Port)

```
user@host> request virtual-chassis vc-port set fpc-slot 1 pic-slot 1 port 0
vc-port successfully set
```

request virtual-chassis vc-port set (Existing Network Port Converted)

```
user@host> request virtual-chassis vc-port set fpc-slot 2 pic-slot 1 port 1
Port conversion initiated, use “show virtual chassis vc-port” to verify
```

request virtual-chassis vc-port set (On Local Router)

```
user@host> request virtual-chassis vc-port set fpc-slot 2 pic-slot 1 port 3 local
vc-port successfully set
```

request virtual-chassis vc-port set (On Remote Member Router 1)

```
user@host> request virtual-chassis vc-port set fpc-slot 5 pic-slot 3 port 10 member 1
vc-port successfully set
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Virtual Chassis Ports to Interconnect Member Routers or Switches

Example: Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis

Guidelines for Configuring Virtual Chassis Ports

request virtual-network-functions

IN THIS SECTION

- Syntax | 15481
- Description | 15481
- Options | 15481
- Required Privilege Level | 15482
- Release Information | 15482

Syntax

`request virtual-network-functions` *`vnf-name`* `(add-image | all-servers | delete-image | force | restart | server | start | stop)`

Description

Start, stop or restart the VNFs. Also, you can add or remove the base image.

> **NOTE**: You can issue these commands either on both the servers (server0 and server1) or on one specific server.

Options

vnf-name	Name of the VNF.
add-image	Add the base image to the VNF repository.
all-servers	Issue the command on both the servers.
delete-image	Delete the base image from the VNF repository.
force	Overwrite the image that was added earlier (example: `request virtual-network-functions` *`vnf-name`* `add-image` *`new-image-name`* `force`). The `force` option is also used with `delete-image` command to perform a cleanup in case, for example, you abruptly stopped an earlier `add-image` process by pressing Ctrl-C.

restart Restart the VNF specified.

server Issue the command on a specific server. Applicable value is `0` or `1`.

start Start the VNF.

stop Stop the VNF.

Required Privilege Level

View

Release Information

Command introduced in Junos OS Release 17.2R1.

The `force` option is available from Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

show virtual-network-functions (Junos Node Slicing) | **27728**

request server authenticate-peer-server | **14901**

request vmhost cleanup

IN THIS SECTION

- Syntax | **15483**
- Description | **15483**
- Options | **15483**
- Required Privilege Level | **15484**
- Output Fields | **15484**
- Release Information | **15484**

Syntax

```
request vmhost cleanup
invoke-on;
(re0 | re1);
routing engine;
```

Description

Clean up temporary files, crash generated files, and log files located in the **/var/tmp**, **/var/crash**, and **/var/log directories respectively** on the host OS.

> **NOTE**: You can clean up files of any format in the **/var/tmp** and **/var/crash** location. However, you can cleanup only the .gz format files in the **/var/log/** location.

Options

none—Clean up temporary files, crash generated files, and log files located in the **/var/tmp**, **/var/crash**, and **/var/log directories** on the host OS running on the Routing Engine on the local Virtual Chassis member.

invoke-on

(Optional) Clean up temporary files, crash generated files, and log files on all the Routing Engines or the other Routing Engine.

Clean up files in **/var/tmp**, **/var/crash**, and **/var/log** on the host OS running on a router that has dual Routing Engines. You can use the `all-routing-engine` option to clean up the files in these directories on the host OS running on all the Routing Engines or the `other-routing-engine` option to clean up the files in these directories on the host OS running on the other Routing Engine. If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

re0 | re1

(Optional) On devices that support dual or redundant Routing Engines, clean up files in **/var/tmp**, **/var/crash**, and **/var/log** on the host OS running on the Routing Engine in slot 0 (`re0`) or the Routing Engine in slot 1 (`re1`).

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

routing-engine (Optional) Specify the Routing Engine for which the files in **/var/tmp**, **/var/crash**, and **/var/log** on the host OS are to be cleaned up. The following options are available:

- `backup`—Backup Routing Engine.
- `both`—Primary and backup Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

request system storage cleanup (Junos OS) | **15374**

request vmhost cleanup (NFX Series)

IN THIS SECTION

- Syntax | 15485
- Description | 15485
- Required Privilege Level | 15485
- Output Fields | 15485
- Release Information | 15485

Syntax

```
request vmhost cleanup
```

Description

Clean up temporary files, crash generated files, and log files located in the **/var/tmp**, **/var/crash**, and **/var/log** directories respectively on the host OS.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Release Information

Command introduced in Junos OS Release 18.1R1.

request vmhost copy jnode-to-vjunos

IN THIS SECTION

- Syntax | 15486
- Description | 15486
- Options | 15486
- Required Privilege Level | 15486
- Release Information | 15486

Syntax

```
request vmhost copy jnode-to-vjunos from-jnode host-filename to-vjunos junos-filename
```

Description

Copy files from Linux host to Junos OS.

Options

from-jnode *host-filename*	Name of the host file to be copied.
to-vjunos *junos-filename*	Name of the Junos OS file to which the host file is copied.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 18.4R1 on Enhanced Automation variants of Junos OS. For more information, see *Overview of Junos Automation Enhancements on Devices Running Junos OS with Enhanced Automation*.

RELATED DOCUMENTATION

request vmhost copy vjunos-to-jnode | 15487

request vmhost copy vjunos-to-jnode

IN THIS SECTION

- Syntax | 15487
- Description | 15487
- Options | 15487
- Required Privilege Level | 15487
- Release Information | 15488

Syntax

```
request vmhost copy vjunos-to-jnode from-vjunos junos-filename to-jnode host-filename
```

Description

Copy files from Junos OS to Linux host.

Options

from-vjunos *junos-filename*	Name of the Junos OS file to be copied.
to-jnode *host-filename*	Name of the host file to which to copy the file.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 18.4R1 on Enhanced Automation variants of Junos OS. For more information, see *Overview of Junos Automation Enhancements on Devices Running Junos OS with Enhanced Automation*.

RELATED DOCUMENTATION

request vmhost copy jnode-to-vjunos | 15486

request vmhost file-copy

IN THIS SECTION

- Syntax | 15488
- Description | 15489
- Options | 15489
- Additional Information | 15489
- Required Privilege Level | 15489
- Sample Output | 15489
- Release Information | 15490

Syntax

```
request vmhost file-copy
           from-jnode filename{
                     (crash | log);
                     to-vjunos directory filename;
                     }
                    invoke-on;
```

Description

Copy crash files or log files from the host OS to Junos OS. You can use these files for analysis and debugging purposes.

Options

crash	Files in **/var/crash** on the host.
from-jnode *filename*	Name of the host file to be copied.
log	Files in **/var/log** on the host.
to-vjunos *directory/<filename>*	Name of the JUNOS directory/<filename> to which the host file is copied.
invoke-on	Remote command execution.

Additional Information

You can use the `show vmhost crash` and `show vmhost logs` commands to list or identify the files in the host OS to be copied to Junos OS.

Required Privilege Level

maintenance

Sample Output

request vmhost file-copy

```
user@host> request vmhost file-copy from-jnode messages log to-vjunos /var/tmp/linux_log_file

user@host> file list /var/tmp/linux_log_file detail
-rw-r-----  1 root  wheel    7253996 Mar 2  23:08 /var/tmp/linux_log_file
total files: 1
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

Command introduced in Junos OS Release 18.1R1 for NFX150 devices.

Command introduced in Junos OS Release 18.4R1 for NFX250 (NG) devices.

Command introduced in Junos OS Release 19.1R1 for NFX350 devices.

request vmhost file-copy (NFX Series)

IN THIS SECTION

- Syntax | 15490
- Description | 15490
- Options | 15491
- Additional Information | 15491
- Required Privilege Level | 15491
- Sample Output | 15491
- Release Information | 15491

Syntax

```
request vmhost file-copy (crash|log) from-jnode host file-name to-vjunos host file-name
```

Description

Copy crash files or log files from the host OS to Junos OS. You can use these files for analysis and debugging purposes.

Options

- crash—Files in /var/crash on the host.
- from-jnode *filename*—Name of the host file to be copied.
- log—Files in **/var/log** on the host.
- to-vjunos *filename*—Name of the Junos OS file to which the host file is copied.

Additional Information

You can use the `show vmhost crash` and `show vmhost logs` commands to list or identify the files in the host OS to be copied to Junos OS.

Required Privilege Level

maintenance

Sample Output

request vmhost file-copy

```
user@host> request vmhost file-copy log from-jnode daemon.log to-vjunos /var/tmp
:/var/tmp # ls -lrt daemon.log
-rw-r--r--  1 root  wheel  1035126 Mar  4 20:33 daemon.log
```

Release Information

Command introduced in Junos OS Release 18.1R1.

request vmhost halt

IN THIS SECTION

- Syntax | 15492
- Description | 15492
- Options | 15492
- Required Privilege Level | 15493
- Sample Output | 15493
- Release Information | 15493

Syntax

```
request vmhost halt
<re0 | re1>
<routing-engine>
```

Description

Stop the host OS and Junos OS running on the device.

Options

none—Stop the host OS and Junos OS on the device.

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, stop the host OS and Junos OS running on the Routing Engine in slot 0 (`re0`) or the Routing Engine in slot 1 (`re1`).

routing-engine (Optional) Specify the Routing Engine on which the host OS and Junos OS needs to be stopped. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Primary and backup Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

maintenance

Sample Output

request vmhost halt

```
request vmhost halt
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

request system halt | 15105

vmhost | 30895

request vmhost halt (NFX Series)

IN THIS SECTION

- Syntax | 15494
- Description | 15494
- Required Privilege Level | 15494
- Sample Output | 15494
- Release Information | 15495

Syntax

```
request vmhost halt
```

Description

Stop the host OS and Junos OS running on the device.

Required Privilege Level

maintenance

Sample Output

request vmhost halt

```
user@host> request vmhost halt
Halt the vmhost ? [yes,no] (no) yes

Initiating vmhost halt... ok
Initiating Junos shutdown...  shutdown: [pid 8782]
Shutdown NOW!
ok
Junos shutdown is in progress...
```

```
*** FINAL System shutdown message from root@ ***

System going down IMMEDIATELY




...
...

Operating System halted
Please press any key to reboot
```

Release Information

Command introduced in Junos OS Release 18.1R1.

request vmhost hard-disk-test

IN THIS SECTION

- Syntax | 15495
- Description | 15496
- Options | 15496
- Required Privilege Level | 15496
- Sample Output | 15496
- Release Information | 15497

Syntax

```
request vmhost hard-disk-test {disk disk-name|long|short|show-status}
```

Description

Run memory and diagnostics monitoring test on the solid-state drive (SSD). The test output provides information about the various attributes of the SSD that is help monitor the status of the hard disk memory.

Options

disk *disk-name* Name of the SSD.

long Run extended Self-Monitoring, Analysis and Reporting Technology (SMART) tests of the SSD.

short Run short SMART tests of the SSD.

show-status Display the status of the test.

Required Privilege Level

maintenance

Sample Output

request vmhost hard-disk-test

```
user@host> request vmhost hard-disk-test show-status disk /dev/sda
smartctl 5.42 2014-07-28 r3460 [x86_64-linux-3.10.79-ovp-rt74-WR6.0.0.20_preempt-rt] (local
build)Copyright (C) 2002-11 by Bruce Allen, http://smartmontools.sourceforge.net

=== START OF INFORMATION SECTION ===
Model Family:     StorFly Slim Sata SSD
Device Model:     StorFly VSF202CC050G-JUN
Serial Number:    P1T13003443810130041
Firmware Version: 0729-000
User Capacity:    50,020,540,416 bytes [50.0 GB]
Sector Size:      512 bytes logical/physical
Device is:        In smartctl database [for details use: -P show]
ATA Version is:   8
ATA Standard is:  ACS-2 (revision not indicated)
Local Time is:    Fri Jun 17 17:30:57 2016 IST
SMART support is: Available - device has SMART capability.
```

```
SMART support is: Enabled

=== START OF READ SMART DATA SECTION ===
SMART overall-health self-assessment test result: PASSED

General SMART Values:
Offline data collection status:  (0x02) Offline data collection activity
                                        was completed without error.
                                        Auto Offline Data Collection: Disabled.
Self-test execution status:      (   0) The previous self-test routine completed
                                        without error or no self-test has ever
                                        been run.

...
...
...
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

vmhost | 30895

request vmhost jdm add (In-Chassis Model)

IN THIS SECTION

- Syntax | 15498
- Description | 15498
- Options | 15498

Required Privilege Level | 15498
Release Information | 15498

Syntax

```
request vmhost jdm add package-name
```

Description

Adds the Juniper Device Manager (JDM) software package. This command is applicable to in-chassis Junos node slicing.

Options

add *package-name* Add the Juniper Device Manager (JDM) software package. You need to type the package name after the command `add`.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

show virtual-network-functions (Junos Node Slicing) | 27728

request server authenticate-peer-server | 14901

request vmhost jdm delete (In-Chassis Model)

IN THIS SECTION

- Syntax | 15499
- Description | 15499
- Required Privilege Level | 15499
- Release Information | 15499

Syntax

```
request vmhost jdm delete
```

Description

Uninstalls the Juniper Device Manager (JDM). This command is applicable to in-chassis Junos node slicing.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

show virtual-network-functions (Junos Node Slicing) | 27728

request server authenticate-peer-server | 14901

request vmhost jdm login (In-Chassis Model)

IN THIS SECTION

- Syntax | 15500
- Description | 15500
- Required Privilege Level | 15500
- Sample Output | 15500
- Release Information | 15501

Syntax

```
request vmhost jdm login
```

Description

Logs in to the Juniper Device Manager (JDM). This command is applicable to in-chassis Junos node slicing.

NOTE: To log in to JDM, you must be a root user.

Required Privilege Level

maintenance

Sample Output

request vmhost jdm login

```
root@router> request vmhost jdm login

***************************************************************************
* The Juniper Device Manager (JDM) must only be used for orchestrating the *
* Virtual Machines for Junos Node Slicing *
* *
```

```
* Host Linux Distro: Wind River Linux *
* JDM Version: jns-jdm-vmhost-19.1-20181003.dev.common.0.x86_64 *
* Free Disk Space on JDM's root-fs ("/"): 125081(MiB) *
***************************************************************************
Last login: Thu Oct 4 15:26:30 2018 from 192.168.1.1
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

show virtual-network-functions (Junos Node Slicing) | **27728**

request server authenticate-peer-server | **14901**

request vmhost jdm start (In-Chassis Model)

IN THIS SECTION

- Syntax | **15501**
- Description | **15501**
- Required Privilege Level | **15502**
- Sample Output | **15502**
- Release Information | **15502**

Syntax

```
request vmhost jdm start
```

Description

Starts the Juniper Device Manager (JDM). This command is applicable to in-chassis Junos node slicing.

Required Privilege Level

maintenance

Sample Output

request vmhost jdm start

```
user@router> request vmhost jdm start
Starting JDM
Starting jdm: Domain jdm defined from
/vm/vm/iapps/jdm//install/juniper/lxc/jdm/current/config/jdm.xml
Domain jdm started
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

show virtual-network-functions (Junos Node Slicing) | **27728**

request server authenticate-peer-server | **14901**

request vmhost jdm stop (In-Chassis Model)

IN THIS SECTION

- Syntax | **15503**
- Description | **15503**
- Required Privilege Level | **15503**
- Release Information | **15503**

Syntax

```
request vmhost jdm stop
```

Description

Stops the Juniper Device Manager (JDM). This command is applicable to in-chassis Junos node slicing.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

show virtual-network-functions (Junos Node Slicing) | **27728**

request server authenticate-peer-server | **14901**

request vmhost local-file-copy

IN THIS SECTION

- Syntax | **15504**
- Description | **15504**
- Options | **15504**
- Additional Information | **15504**
- Required Privilege Level | **15504**
- Sample Output | **15505**
- Release Information | **15505**

Syntax

```
request vmhost local-file-copy source source_path destination destination_path;
```

Description

Copy a file or a directory present under **/var/public/** directory structure on the hypervisor to a different name under **/var/public/** directory structure.

Options

- source *source-path*—Name of the source path from which the file is to be copied.
- destination*destination-path*—Name of the destination path to which the file is copied.

Additional Information

Consider the following guidelines while creating a local path.

- Both paths should be under **/var/public/** directory
- Avoid special characters. You can include alphanumeric characters, hyphen, underscore, dot, and asterisk.
- Avoid relative path components **..** and **/./**
- Avoid ending the path with **/.**

Following are the examples of a path:

- **/var/public/sanity_packages**
- **/var/public/sanity_packages/readme.txt**
- **/var/public/12-hours.data**

Required Privilege Level

Maintenance

Sample Output

request vmhost local-file-copy (To copy a directory)

```
user@host> request vmhost local-file-copy source /var/public/sanity_packages destination /var/
public/new_sanity_packages
'/var/public/sanity_packages' -> '/var/public/new_sanity_packages'
'/var/public/sanity_packages/centos-linux-1.img.gz' -> '/var/public/new_sanity_packages/centos-
linux-1.img.gz'
'/var/public/sanity_packages/vnf_test_pkg.tar.gz' -> '/var/public/new_sanity_packages/
vnf_test_pkg.tar.gz'
'/var/public/sanity_packages/readme.txt' -> '/var/public/new_sanity_packages/readme.txt'
```

request vmhost local-file-copy (To copy a file specifying the destination file name)

```
user@host> request vmhost local-file-copy source
/var/public/sanity_packages/readme.txt destination
/var/public/new_sanity_packages/new_readme.txt
'/var/public/sanity_packages/readme.txt' ->
'/var/public/new_sanity_packages/new_readme.txt'
```

request vmhost local-file-copy (To copy a file specifying the destination directory)

```
user@host> request vmhost local-file-copy source
 /var/public/sanity_packages/readme.txt destination
/var/public/new_sanity_packages/
'/var/public/sanity_packages/readme.txt' ->
'/var/public/new_sanity_packages/readme.txt'
```

Release Information

Command introduced in Junos OS Release 24.2R1 for NFX150, NFX250, and NFX350 devices.

request vmhost mode

IN THIS SECTION

- Syntax | 15506
- Description | 15506
- Required Privilege Level | 15507
- Sample Output | 15507
- Release Information | 15507

Syntax

```
request vmhost mode [compute | hybrid | throughput | mode-name]
```

Description

Select the operational mode of the device from a pre-defined list of modes (compute, hybrid, or throughput) or specify a *mode-name* for a custom mode.

NOTE:

- Starting from Junos OS Release 19.3R1, if the same physical CPU is used for both VNFs and the Junos OS or device components, the request to change the mode fails and an error message is displayed. For example:

  ```
  root> request vmhost mode throughput
  error: Mode cannot be changed; Reason: Reserved CPUs conflict with VNF cpu
  pinnings: 3
  ```

- When you upgrade the software image that has a VNF CPU conflict to Junos OS Release 19.3R1 by using the CLI upgrade option, the upgrade succeeds and the VNF

configuration is applied. The VNF CPU conflict is reported by JDM only if you issue a `commit` command. You must modify the VNF configurations accordingly.

Required Privilege Level

maintenance

Sample Output

request vmhost mode compute

```
user@host> request vmhost mode compute
warning: Device will be rebooted to change the mode from hybrid to compute
Do you want to continue? [yes,no] (no)
```

Release Information

The `request vmhost mode` command is introduced in Junos OS Release 19.1R1.

The `request vmhost mode` *`mode-name`* command is introduced in Junos OS Release 21.1R1 for NFX250 NextGen and NFX350 devices.

request-vmhost-power-cycle-cb

IN THIS SECTION

- Syntax | 15508
- Description | 15508
- Required Privilege Level | 15508
- Sample Output | 15508
- Release Information | 15516

Syntax

```
request vmhost power-cycle-cb
```

Description

Initiates the power cycle for the Chassis.

Required Privilege Level

maintenance

Sample Output

request vmhost power-cycle-cb

```
user@host> request vmhost power-cycle-cb
Power cycle the Chassis/CB ? [yes,no] (no) yes

Initiating vmhost power-cycle-cb...
warning: Initiating Junos shutdown...
shutdown: [pid 14832]

*** FINAL System shutdown message from root@sw-bugatti-pfe-b ***

System going down IMMEDIATELY



Stopping cron.
Waiting for PIDS: 13380.
 .
 .
uma_zalloc_arg: zone "2048" with the following non-sleepable locks held:
shared rm hhook_head rm lock (hhook_head rm lock) r = 0 (0xfffff8010592c310) locked @ /.amd/svl-
engdata5vs2/occamdev/build/freebsd/stable_12/sandbox-stable-12-202301241026/freebsd/
stable_12/20230123.172357__ci_fbsd_builder_stable_12.23fb5ff/src/sys/kern/kern_hhook.c:113
stack backtrace:
Feb  1 03:08:09 jlaunchd 13356 - - ethernet-link-fault-management (PID 13585) terminate signal
15 sent
#0 0xffffffff80548bcd at witness_debugger+0x6d
```

```
#1 0xffffffff80Feb  1 03:08:09 jlaunchd 13356 -549aeb at witness_warn+0x41b
#2 0xffffffff8064092d at uma_zall - wireless-wan-service (PID 135oc_arg+0x9d
#3 0xffffffff804b4788) terminate signal 15 sent
7b at malloc+0x7b
#4 0xffffffff82c853f2 at junos_malloc+0xe2
#5 0xffffffff82c89abe at net_jlock_osd_get+0x5e
#6 0xffffffff82c89a2d at net_jlock_rw_is_locked+0x2d
#7 0xffffffff82c92dd3 at sock_close_hook+0x73
#8 0xffffffff8049a55a at hhook_run_hooks+0x7a
#9 0xffffffff80579ea4 at sodealloc+0xc4
#10 0xffffffff80Feb  1 03:08:09 57b67c at soclose+0x31c
#11 0xfjlaunchd 13356 -ffffff80482865 at _fdrop+0x25
#12 0xffffffff80485bf0 at close - shm-rtsdbd (PID 13595) terminate signal 15 sent
f+0x1f0
#13 0xffffffff8048566c Feb  1 03:08:09 at fdescfree_fds+0x4c
#14 0xffffffff804851db at fdescfree+0x3djlaunchd 13356 -b
 - gstatd (PID 13598) terminate signal 15 sent
#15 0xffffffff80494b24 at exit1+0x564
#16 0xffffffff804dff4f at sigexit+0x35f
#17 0xffffffff804e1347 at postsig+0x347
Feb  1 03:08:09 jlaunchd 13356 - - rpcbind-service (PID 13607) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - ftp-inet-process (PID 13608) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - smid (PID 13609) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - mountd-service (PID 13612) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - process-monitor (PID 13613) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - application-identification (PID 13620) terminate signal 15
sent
Feb  1 03:08:09 jlaunchd 13356 - - ifstate-tracing-process (PID 13622) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - charged (PID 13624) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - mobiled (PID 13626) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - system-monitor (PID 13629) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - bbe-up-pfcp-proxy (PID 13637) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - license-service (PID 13639) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - pki-service (PID 13892) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - alarm-control (PID 13895) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - interface-control (PID 13896) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - snmp (PID 13897) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - sntp (PID 13898) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - internal-routing-service (PID 13899) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - firewall (PID 13900) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - general-authentication-service (PID 13901) terminate signal
15 sent
Feb  1 03:08:09 jlaunchd 13356 - - l2tp-universal-edge (PID 13902) terminate signal 15 sent
```

```
Feb  1 03:08:09 jlaunchd 13356 - - ppp-service (PID 13904) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - alarm-management (PID 13907) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - transport-control (PID 13912) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - up-smg-service (PID 13917) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - SDN-Telemetry (PID 13921) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - smg-service (PID 13924) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - replication-client-process (PID 13926) terminate signal 15
sent
Feb  1 03:08:09 jlaunchd 13356 - - bbe-stats-daemon (PID 13928) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - replication-serverd-process (PID 13929) terminate signal 15
sent
Feb  1 03:08:09 jlaunchd 13356 - - kernel-health-monitoring (PID 13931) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - watchdog (PID 13995) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - dot1x-protocol (PID 14023) terminate signal 15 sent
Feb  1 03:08:09 jlaunchd 13356 - - cron (PID 14861) terminate signal 15 sent
Feb  1 03:08:13 jlaunchd 13356 - - interface-control (PID 13896) exited with status=0 Normal Exit
Feb  1 03:08:13 jlaunchd 13356 - - mpls-traceroute (PID 13859) exited with status=0 Normal Exit
Feb  1 03:08:13 jlaunchd 13356 - - l2cpd-service (PID 13858) exited with status=0 Normal Exit
Feb  1 03:08:13 jlaunchd 13356 - - dynamic-flow-capture (PID 13857) exited with status=0 Normal
Exit
Feb  1 03:08:13 jlaunchd 13356 - - license-service (PID 13639) exited with status=0 Normal Exit
Feb  1 03:08:13 jlaunchd 13356 - - bbe-up-pfcp-proxy (PID 13637) exited with status=0 Normal Exit
Feb  1 03:08:13 jlaunchd 13356 - - ethernet-link-fault-management (PID 13585) exited with
status=0 Normal Exit
Waiting (max 60 seconds) for system process `vnlru' to stop... done
Waiting (max 60 seconds) for system process `syncer' to stop...
Syncing disks, vnodes remaining... 0 0 0 0 0 done
Waiting (max 60 seconds) for system thread `bufdaemon' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-0' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-1' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-2' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-3' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-4' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-5' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-6' to stop... done
All buffers synced.
jnpr-ctrl: port 1(eth2) entered disabled state
Uptime: 3m20s
pci 0000:17:00.0: Removing from iommu group 37
pcieport 0000:00:09.0: bridge window [io  0x1000-0x0fff] to [bus 02] add_size 1000
pcieport 0000:00:0a.0: bridge window [io  0x1000-0x0fff] to [bus 03] add_size 1000
pcieport 0000:00:09.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: failed to assign [io  size 0x1000]
```

```
pcieport 0000:00:0a.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: failed to assign [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: failed to assign [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: failed to assign [io  size 0x1000]
pci 0000:17:00.0: [10b5:8749] type 01 class 0x060400
pci 0000:17:00.0: reg 0x10: [mem 0xb4c00000-0xb4c3ffff]
pci 0000:17:00.0: PME# supported from D0 D3hot D3cold
pci 0000:17:00.0: Adding to iommu group 37
pci 0000:18:10.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:10.0: PME# supported from D0 D3hot D3cold
pci 0000:18:10.0: Adding to iommu group 37
pci 0000:18:11.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:11.0: PME# supported from D0 D3hot D3cold
pci 0000:18:11.0: Adding to iommu group 37
pci 0000:18:12.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:12.0: PME# supported from D0 D3hot D3cold
pci 0000:18:12.0: Adding to iommu group 37
pci 0000:17:00.0: PCI bridge to [bus 18-1b]
pci 0000:17:00.0:   bridge window [mem 0xb4000000-0xb4bfffff]
pci 0000:19:00.0: [1304:008d] type 00 class 0x030000
pci 0000:19:00.0: reg 0x10: [mem 0xb4000000-0xb47fffff 64bit]
pci 0000:19:00.0: PME# supported from D0
pci 0000:19:00.0: 4.000 Gb/s available PCIe bandwidth, limited by 5 GT/s x1 link at 0000:18:10.0
(capable of 16.000 Gb/s with 5 GT/s x4 link)
pci 0000:19:00.0: vgaarb: VGA device added: decodes=io+mem,owns=none,locks=none
pci 0000:19:00.0: Adding to iommu group 37
pci 0000:18:10.0: PCI bridge to [bus 19]
pci 0000:18:10.0:   bridge window [mem 0xb4000000-0xb47fffff]
pci 0000:1a:00.0: [1304:00e6] type 00 class 0x058000
pci 0000:1a:00.0: reg 0x10: [mem 0xb4800000-0xb48fffff]
pci 0000:1a:00.0: Adding to iommu group 37
pci 0000:18:11.0: PCI bridge to [bus 1a]
pci 0000:18:11.0:   bridge window [mem 0xb4800000-0xb48fffff]
pci 0000:1b:00.0: [10b5:87b0] type 00 class 0x068000
pci 0000:1b:00.0: reg 0x10: [mem 0xb4900000-0xb493ffff]
pci 0000:1b:00.0: Setting PLX NTB proxy ID aliases
pci 0000:1b:00.0: Enabling fixed DMA alias for devfn range from 00.0 to 1f.7
pci 0000:1b:00.0: Adding to iommu group 37
pci 0000:18:12.0: PCI bridge to [bus 1b]
pci 0000:18:12.0:   bridge window [mem 0xb4900000-0xb49fffff]
pci 0000:17:00.0: BAR 14: assigned [mem 0xb4000000-0xb4bfffff]
```

```
pci 0000:17:00.0: BAR 0: assigned [mem 0xb4c00000-0xb4c3ffff]
pci 0000:18:10.0: BAR 14: assigned [mem 0xb4000000-0xb47fffff]
pci 0000:18:11.0: BAR 14: assigned [mem 0xb4800000-0xb48fffff]
pci 0000:18:12.0: BAR 14: assigned [mem 0xb4900000-0xb49fffff]
pci 0000:19:00.0: BAR 0: assigned [mem 0xb4000000-0xb47fffff 64bit]
pci 0000:18:10.0: PCI bridge to [bus 19]
pci 0000:18:10.0:   bridge window [mem 0xb4000000-0xb47fffff]
pci 0000:1a:00.0: BAR 0: assigned [mem 0xb4800000-0xb48fffff]
pci 0000:18:11.0: PCI bridge to [bus 1a]
pci 0000:18:11.0:   bridge window [mem 0xb4800000-0xb48fffff]
pci 0000:1b:00.0: BAR 0: assigned [mem 0xb4900000-0xb493ffff]
pci 0000:18:12.0: PCI bridge to [bus 1b]
pci 0000:18:12.0:   bridge window [mem 0xb4900000-0xb49fffff]
pci 0000:17:00.0: PCI bridge to [bus 18-1b]
pci 0000:17:00.0:   bridge window [mem 0xb4000000-0xb4bfffff]
pci 0000:19:00.0: Removing from iommu group 37
pci_bus 0000:19: busn_res: [bus 19] is released
pci 0000:18:10.0: Removing from iommu group 37
pci 0000:1a:00.0: Removing from iommu group 37
pci_bus 0000:1a: busn_res: [bus 1a] is released
pci 0000:18:11.0: Removing from iommu group 37
pci 0000:1b:00.0: Removing from iommu group 37
pci_bus 0000:1b: busn_res: [bus 1b] is released
pci 0000:18:12.0: Removing from iommu group 37
pci_bus 0000:18: busn_res: [bus 18-1b] is released
pci 0000:17:00.0: Removing from iommu group 37
pcieport 0000:00:09.0: bridge window [io  0x1000-0x0fff] to [bus 02] add_size 1000
pcieport 0000:00:0a.0: bridge window [io  0x1000-0x0fff] to [bus 03] add_size 1000
pcieport 0000:00:09.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: failed to assign [io  size 0x1000]
Khelp module "jsocket" can't unload until its refcount drops fropcieport 0000:00:0a.0: BAR 13:
no space for [io  size 0x1000]
m 4 to 0.
acpi0: Powering system off
pcieport 0000:00:0a.0: BAR 13: failed to assign [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: failed to assign [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: failed to assign [io  size 0x1000]
pci 0000:17:00.0: [10b5:8749] type 01 class 0x060400
pci 0000:17:00.0: reg 0x10: [mem 0xb4c00000-0xb4c3ffff]
pci 0000:17:00.0: PME# supported from D0 D3hot D3cold
pci 0000:17:00.0: Adding to iommu group 37
```

```
pci 0000:18:10.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:10.0: PME# supported from D0 D3hot D3cold
pci 0000:18:10.0: Adding to iommu group 37
pci 0000:18:11.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:11.0: PME# supported from D0 D3hot D3cold
pci 0000:18:11.0: Adding to iommu group 37
pci 0000:18:12.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:12.0: PME# supported from D0 D3hot D3cold
pci 0000:18:12.0: Adding to iommu group 37
pci 0000:17:00.0: PCI bridge to [bus 18-1b]
pci 0000:17:00.0:   bridge window [mem 0xb4000000-0xb4bfffff]
pci 0000:19:00.0: [1304:008d] type 00 class 0x030000
pci 0000:19:00.0: reg 0x10: [mem 0xb4000000-0xb47fffff 64bit]
pci 0000:19:00.0: PME# supported from D0
pci 0000:19:00.0: 4.000 Gb/s available PCIe bandwidth, limited by 5 GT/s x1 link at 0000:18:10.0
(capable of 16.000 Gb/s with 5 GT/s x4 link)
pci 0000:19:00.0: vgaarb: VGA device added: decodes=io+mem,owns=none,locks=none
pci 0000:19:00.0: Adding to iommu group 37
pci 0000:18:10.0: PCI bridge to [bus 19]
pci 0000:18:10.0:   bridge window [mem 0xb4000000-0xb47fffff]
pci 0000:1a:00.0: [1304:00e6] type 00 class 0x058000
pci 0000:1a:00.0: reg 0x10: [mem 0xb4800000-0xb48fffff]
pci 0000:1a:00.0: Adding to iommu group 37
pci 0000:18:11.0: PCI bridge to [bus 1a]
pci 0000:18:11.0:   bridge window [mem 0xb4800000-0xb48fffff]
pci 0000:1b:00.0: [10b5:87b0] type 00 class 0x068000
pci 0000:1b:00.0: reg 0x10: [mem 0xb4900000-0xb493ffff]
pci 0000:1b:00.0: Setting PLX NTB proxy ID aliases
pci 0000:1b:00.0: Enabling fixed DMA alias for devfn range from 00.0 to 1f.7
pci 0000:1b:00.0: Adding to iommu group 37
pci 0000:18:12.0: PCI bridge to [bus 1b]
pci 0000:18:12.0:   bridge window [mem 0xb4900000-0xb49fffff]
pci 0000:17:00.0: BAR 14: assigned [mem 0xb4000000-0xb4bfffff]
pci 0000:17:00.0: BAR 0: assigned [mem 0xb4c00000-0xb4c3ffff]
pci 0000:18:10.0: BAR 14: assigned [mem 0xb4000000-0xb47fffff]
pci 0000:18:11.0: BAR 14: assigned [mem 0xb4800000-0xb48fffff]
pci 0000:18:12.0: BAR 14: assigned [mem 0xb4900000-0xb49fffff]
pci 0000:19:00.0: BAR 0: assigned [mem 0xb4000000-0xb47fffff 64bit]
pci 0000:18:10.0: PCI bridge to [bus 19]
pci 0000:18:10.0:   bridge window [mem 0xb4000000-0xb47fffff]
pci 0000:1a:00.0: BAR 0: assigned [mem 0xb4800000-0xb48fffff]
pci 0000:18:11.0: PCI bridge to [bus 1a]
pci 0000:18:11.0:   bridge window [mem 0xb4800000-0xb48fffff]
```

```
pci 0000:1b:00.0: BAR 0: assigned [mem 0xb4900000-0xb493ffff]
pci 0000:18:12.0: PCI bridge to [bus 1b]
pci 0000:18:12.0:   bridge window [mem 0xb4900000-0xb49fffff]
pci 0000:17:00.0: PCI bridge to [bus 18-1b]
pci 0000:17:00.0:   bridge window [mem 0xb4000000-0xb4bfffff]
High host-mode time vcpu 0 delay 91333 us exit reason IO_INSTRUCTION rip 0xffffffff84f1b1bc info
c3920000 0
virbr0: port 2(tap1) entered disabled state
device tap1 left promiscuous mode
virbr0: port 2(tap1) entered disabled state
jnpr-ctrl: port 2(tap3) entered disabled state
device tap3 left promiscuous mode
jnpr-ctrl: port 2(tap3) entered disabled state
jnpr-re-re: port 2(tap2) entered disabled state
device tap2 left promiscuous mode
jnpr-re-re: port 2(tap2) entered disabled state
pci 0000:19:00.0: Removing from iommu group 37
pci_bus 0000:19: busn_res: [bus 19] is released
pci 0000:18:10.0: Removing from iommu group 37
pci 0000:1a:00.0: Removing from iommu group 37
pci_bus 0000:1a: busn_res: [bus 1a] is released
pci 0000:18:11.0: Removing from iommu group 37
pci 0000:1b:00.0: Removing from iommu group 37
pci_bus 0000:1b: busn_res: [bus 1b] is released
pci 0000:18:12.0: Removing from iommu group 37
pci_bus 0000:18: busn_res: [bus 18-1b] is released
ctrlbr0: port 1(vnet100) entered disabled state
device vnet100 left promiscuous mode
ctrlbr0: port 1(vnet100) entered disabled state
[FAILED] Failed unmounting /.old/tdev/.uniprintk: systemd-shutdow: 54 output lines suppressed
due to ratelimiting
on/.s.
      [  OK  ] Unmounted /.old/tmp.systemd-shutdown[1]: Syncing filesystems and block devices.

[  OK  ] Unmounted /.old/vm.
                            [  systemd-shutdown[1]: Sending SIGTERM to remaining processes...
OK  ] Unmounted /spare.
                       [  OK  ] Unmsystemd-journald[399]: Received SIGTERM from PID 1 (systemd-
shutdow).
ounted /tmp/junos-7ZtJR.
                        [  OK  ] Unmounted /var/volatile.
                                                          [  OK  ] Unmounted /vm.
```

```
Unmounting /.old...

                 pci 0000:17:00.0: Removing from iommu group 37
   Unmounting Temporary Directory (/tmp)...
                                          [  OK  ] Unmounted /junos.
                                                                      [FAILED] Failed
unmounting /.old.

     [  OK  ] Unmounted Temporary Directory (/tmp).

                                                   [  OK  ] Stopped target Local File Systems
(Pre).
           [  OK  ] Stopped target Swap.
                                       [  OK  ] Reached target Unmount All Filesystems.
                                                                                          [  OK  ]
Stopped Remount Root and Kernel File Systems.

                                             [  OK  ] Stopped Create Static Device Nodes in /dev.
        [  OK  ] Reached target Shutdown.
                                        [  OK  ] Reached target Final Step.
                                                                            [  OK  ] Started
Power-Off.

      [  OK  ] Reached target Power-Off.

                                       systemd-shutdown[1]: Sending SIGKILL to remaining
processes...
systemd-shutdown[1]: Unmounting file systems.
[40069]: Remounting '/' read-only in with options 'si=cb688aeaeeec04f4'.
[40070]: Remounting '/.old/tdev/.union/.s' read-only in with options '(null)'.
[40070]: Failed to remount '/.old/tdev/.union/.s' read-only: Device or resource busy

Booting from Flash A

FPGA Reset Reason = 0x87

Primary BIOS version CICL_P_AVG0_00.11.01

Total Memory Size = 128 GB

Checking Primary BIOS code integrity...Passed!
```

Release Information

request-vmhost-power-cycle-re

IN THIS SECTION

- Syntax | 15516
- Description | 15516
- Required Privilege Level | 15516
- Sample Output | 15516
- Release Information | 15524

Syntax

```
request vmhost power-cycle-cb
```

Description

Initiates the power cycle for the Routing Engine.

Required Privilege Level

maintenance

Sample Output

request vmhost power-cycle-re

```
user@host> request vmhost power-cycle-re
Power cycle the RE ? [yes,no] (no) yes

Initiating vmhost power-cycle...
```

```
warning: Initiating Junos shutdown...
shutdown: [pid 20669]

*** FINAL System shutdown message from root@sw-bugatti-pfe-b ***

System going down IMMEDIATELY



Stopping cron.
Waiting for PIDS: 18970.
 .
 .
uma_zalloc_arg: zone "2048" with the following non-sleepable locks held:
shared rm hhook_head rm lock (hhook_head rm lock) r = 0 (0xfffff80105a3bc10) locked @ /.amd/svl-
engdata5vs2/occamdev/build/freebsd/stable_12/sandbox-stable-12-202301241026/freebsd/
stable_12/20230123.172357__ci_fbsd_builder_stable_12.23fb5ff/src/sys/kern/kern_hhook.c:113
stack backtrace:
#0 0xffffffff80548bcd at witness_debugger+0x6d
#1 0xffffffff80549aeb at witness_warn+0x41b
#2 0xffffffff8064092d at uma_zalloc_arg+0x9d
#3 0xffffffff804b47Feb  1 02:27:13 7b at malloc+0x7b
#4 0xffffffffjlaunchd 18899 - - disk-monitoring (PID 19183) t82c853f2 at junos_malloc+0xe2
#5 0xffffffff82c89abe at net_jloerminate signal 15 sent
ck_osd_get+0x5e
#6 0xffffffff82c89a2d at net_jlock_rw_is_lockeFeb  1 02:27:13 jlaunchd 18899 -d+0x2d
#7 0xffffffff82c92dd3 at - neighbor-liveness (PID 19196) terminate signal 15 sent
 sock_close_hook+0x73
#8 0xffffffff8049a55a at hhook_run_hooks+0x7a
#9 0xffffffff80579ea4 at sodealloc+0xc4
#10 0xffffffff8057b67c at soclose+0x31c
#11 0xffffffff80482865 at _fdrop+0x25
#12 0xffffffff80485bf0 at closef+0x1f0
#13 0xffffffff8048566c at fdescfree_fds+0x4c
#14 0xffffffff804851db at fdescfree+0x3db
#15 0xffffffff80494b24 at exit1+0x564
#16 0xffffffff804dff4f at sigexit+0x35f
#17 0xffffffff804e1347 at postsig+0x347
Feb  1 02:27:13 jlaunchd 18899 - - event-processing (PID 18877) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - clksyncd-service (PID 19201) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - ethernet-link-fault-management (PID 19203) terminate signal
15 sent
Feb  1 02:27:13 jlaunchd 18899 - - wireless-wan-service (PID 19208) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - shm-rtsdbd (PID 19216) terminate signal 15 sent
```

```
Feb  1 02:27:13 jlaunchd 18899 - - gstatd (PID 19219) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - rpcbind-service (PID 19232) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - ftp-inet-process (PID 19234) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - jsqlsync (PID 19292) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - watchdog (PID 20030) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - ntp (PID 20176) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - routing (PID 20177) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - l2-learning (PID 20178) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - remote-operations (PID 20180) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - class-of-service (PID 20182) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - link-management (PID 20183) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - pppoe (PID 20185) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - redundancy-device (PID 20187) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - ppp (PID 20190) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - dynamic-flow-capture (PID 20192) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - l2cpd-service (PID 20193) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - mpls-traceroute (PID 20196) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - secure-neighbor-discovery (PID 20198) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - iccp-service (PID 20200) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - ddos-service (PID 20201) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - nfsd-service (PID 20205) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - commit-syncd (PID 20207) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - mspd (PID 20211) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - kernel-debug-streaming (PID 20213) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - datapath-trace-service (PID 20215) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - captive-portal-content-delivery for dBNG (PID 20216)
terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - subscriber-management-helper (PID 20217) terminate signal 15
sent
Feb  1 02:27:13 jlaunchd 18899 - - kernel-offload-service (PID 20218) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - health-mon (PID 20220) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - jsd (PID 20221) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - grpc-tunnel (PID 20222) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - xmlproxyd (PID 20223) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - analytics-agent (PID 20224) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - na-grpc-server (PID 20225) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - rpd-telemetry-daemon (PID 20226) terminate signal 15 sent
Feb  1 02:27:13 jlaunchd 18899 - - pki-service (PID 20227) terminate signal 15 sent
Feb  1 02:27:17 jlaunchd 18899 - - alarm-control (PID 20230) exited with status=0 Normal Exit
Feb  1 02:27:17 jlaunchd 18899 - - ddos-service (PID 20201) exited with status=0 Normal Exit
Feb  1 02:27:17 jlaunchd 18899 - - iccp-service (PID 20200) exited with status=0 Normal Exit
Feb  1 02:27:17 jlaunchd 18899 - - secure-neighbor-discovery (PID 20198) exited with status=0
Normal Exit
```

```
Feb  1 02:27:17 jlaunchd 18899 - - aaasd (PID 19282) exited with status=0 Normal Exit
Feb  1 02:27:17 jlaunchd 18899 - - license-service (PID 19280) exited with status=0 Normal Exit
Feb  1 02:27:17 jlaunchd 18899 - - bbe-up-pfcp-proxy (PID 19278) exited with status=0 Normal Exit
Feb  1 02:27:17 jlaunchd 18899 - - ethernet-link-fault-management (PID 19203) exited with
status=0 Normal Exit
Feb  1 02:27:17 jlaunchd 18899 - - clksyncd-service (PID 19201) exited with status=0 Normal Exit
Waiting (max 60 seconds) for system process `vnlru' to stop... done
Waiting (max 60 seconds) for system process `syncer' to stop...
Syncing disks, vnodes remaining... 0 0 0 0 0 done
Waiting (max 60 seconds) for system thread `bufdaemon' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-0' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-1' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-2' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-3' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-5' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-4' to stop... done
Waiting (max 60 seconds) for system thread `bufspacedaemon-6' to stop... done
All buffers synced.
pci 0000:19:00.0: Removing from iommu group 37
pci_bus 0000:19: busn_res: [bus 19] is released
pci 0000:18:10.0: Removing from iommu group 37
pci 0000:1a:00.0: Removing from iommu group 37
pci_bus 0000:1a: busn_res: [bus 1a] is released
pci 0000:18:11.0: Removing from iommu group 37
pci 0000:1b:00.0: Removing from iommu group 37
pci_bus 0000:1b: busn_res: [bus 1b] is released
pci 0000:18:12.0: Removing from iommu group 37
pci_bus 0000:18: busn_res: [bus 18-1b] is released
Uptime: 2m52s
pci 0000:17:00.0: Removing from iommu group 37
pcieport 0000:00:09.0: bridge window [io  0x1000-0x0fff] to [bus 02] add_size 1000
pcieport 0000:00:0a.0: bridge window [io  0x1000-0x0fff] to [bus 03] add_size 1000
pcieport 0000:00:09.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: failed to assign [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: failed to assign [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: failed to assign [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: failed to assign [io  size 0x1000]
pci 0000:17:00.0: [10b5:8749] type 01 class 0x060400
pci 0000:17:00.0: reg 0x10: [mem 0xb4c00000-0xb4c3ffff]
pci 0000:17:00.0: PME# supported from D0 D3hot D3cold
```

```
pci 0000:17:00.0: Adding to iommu group 37
pci 0000:18:10.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:10.0: PME# supported from D0 D3hot D3cold
pci 0000:18:10.0: Adding to iommu group 37
pci 0000:18:11.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:11.0: PME# supported from D0 D3hot D3cold
pci 0000:18:11.0: Adding to iommu group 37
pci 0000:18:12.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:12.0: PME# supported from D0 D3hot D3cold
pci 0000:18:12.0: Adding to iommu group 37
pci 0000:17:00.0: PCI bridge to [bus 18-1b]
pci 0000:17:00.0:   bridge window [mem 0xb4000000-0xb4bfffff]
pci 0000:19:00.0: [1304:008d] type 00 class 0x030000
pci 0000:19:00.0: reg 0x10: [mem 0xb4000000-0xb47fffff 64bit]
pci 0000:19:00.0: PME# supported from D0
pci 0000:19:00.0: 4.000 Gb/s available PCIe bandwidth, limited by 5 GT/s x1 link at 0000:18:10.0
(capable of 16.000 Gb/s with 5 GT/s x4 link)
pci 0000:19:00.0: vgaarb: VGA device added: decodes=io+mem,owns=none,locks=none
pci 0000:19:00.0: Adding to iommu group 37
pci 0000:18:10.0: PCI bridge to [bus 19]
pci 0000:18:10.0:   bridge window [mem 0xb4000000-0xb47fffff]
pci 0000:1a:00.0: [1304:00e6] type 00 class 0x058000
pci 0000:1a:00.0: reg 0x10: [mem 0xb4800000-0xb48fffff]
pci 0000:1a:00.0: Adding to iommu group 37
pci 0000:18:11.0: PCI bridge to [bus 1a]
pci 0000:18:11.0:   bridge window [mem 0xb4800000-0xb48fffff]
pci 0000:1b:00.0: [10b5:87b0] type 00 class 0x068000
pci 0000:1b:00.0: reg 0x10: [mem 0xb4900000-0xb493ffff]
pci 0000:1b:00.0: Setting PLX NTB proxy ID aliases
pci 0000:1b:00.0: Enabling fixed DMA alias for devfn range from 00.0 to 1f.7
pci 0000:1b:00.0: Adding to iommu group 37
pci 0000:18:12.0: PCI bridge to [bus 1b]
pci 0000:18:12.0:   bridge window [mem 0xb4900000-0xb49fffff]
pci 0000:17:00.0: BAR 14: assigned [mem 0xb4000000-0xb4bfffff]
pci 0000:17:00.0: BAR 0: assigned [mem 0xb4c00000-0xb4c3ffff]
pci 0000:18:10.0: BAR 14: assigned [mem 0xb4000000-0xb47fffff]
pci 0000:18:11.0: BAR 14: assigned [mem 0xb4800000-0xb48fffff]
pci 0000:18:12.0: BAR 14: assigned [mem 0xb4900000-0xb49fffff]
pci 0000:19:00.0: BAR 0: assigned [mem 0xb4000000-0xb47fffff 64bit]
pci 0000:18:10.0: PCI bridge to [bus 19]
pci 0000:18:10.0:   bridge window [mem 0xb4000000-0xb47fffff]
pci 0000:1a:00.0: BAR 0: assigned [mem 0xb4800000-0xb48fffff]
pci 0000:18:11.0: PCI bridge to [bus 1a]
```

```
pci 0000:18:11.0:   bridge window [mem 0xb4800000-0xb48fffff]
pci 0000:1b:00.0: BAR 0: assigned [mem 0xb4900000-0xb493ffff]
pci 0000:18:12.0: PCI bridge to [bus 1b]
pci 0000:18:12.0:   bridge window [mem 0xb4900000-0xb49fffff]
pci 0000:17:00.0: PCI bridge to [bus 18-1b]
pci 0000:17:00.0:   bridge window [mem 0xb4000000-0xb4bfffff]
Khelp module "jsocket" can't unload until its refcount drops fropci 0000:19:00.0: Removing from
iommu group 37
m 4 to 0.
acpi0: Powering system off
pci_bus 0000:19: busn_res: [bus 19] is released
pci 0000:18:10.0: Removing from iommu group 37
pci 0000:1a:00.0: Removing from iommu group 37
pci_bus 0000:1a: busn_res: [bus 1a] is released
pci 0000:18:11.0: Removing from iommu group 37
pci 0000:1b:00.0: Removing from iommu group 37
pci_bus 0000:1b: busn_res: [bus 1b] is released
pci 0000:18:12.0: Removing from iommu group 37
pci_bus 0000:18: busn_res: [bus 18-1b] is released
virbr0: port 2(tap1) entered disabled state
device tap1 left promiscuous mode
virbr0: port 2(tap1) entered disabled state
jnpr-ctrl: port 2(tap3) entered disabled state
device tap3 left promiscuous mode
jnpr-ctrl: port 2(tap3) entered disabled state
jnpr-re-re: port 2(tap2) entered disabled state
pci 0000:17:00.0: Removing from iommu group 37
pcieport 0000:00:09.0: bridge window [io  0x1000-0x0fff] to [bus 02] add_size 1000
pcieport 0000:00:0a.0: bridge window [io  0x1000-0x0fff] to [bus 03] add_size 1000
pcieport 0000:00:09.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: failed to assign [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: failed to assign [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:0a.0: BAR 13: failed to assign [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: no space for [io  size 0x1000]
pcieport 0000:00:09.0: BAR 13: failed to assign [io  size 0x1000]
pci 0000:17:00.0: [10b5:8749] type 01 class 0x060400
pci 0000:17:00.0: reg 0x10: [mem 0xb4c00000-0xb4c3ffff]
pci 0000:17:00.0: PME# supported from D0 D3hot D3cold
pci 0000:17:00.0: Adding to iommu group 37
pci 0000:18:10.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:10.0: PME# supported from D0 D3hot D3cold
```

```
pci 0000:18:10.0: Adding to iommu group 37
pci 0000:18:11.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:11.0: PME# supported from D0 D3hot D3cold
pci 0000:18:11.0: Adding to iommu group 37
pci 0000:18:12.0: [10b5:8749] type 01 class 0x060400
pci 0000:18:12.0: PME# supported from D0 D3hot D3cold
pci 0000:18:12.0: Adding to iommu group 37
pci 0000:17:00.0: PCI bridge to [bus 18-1b]
pci 0000:17:00.0:   bridge window [mem 0xb4000000-0xb4bfffff]
pci 0000:19:00.0: [1304:008d] type 00 class 0x030000
pci 0000:19:00.0: reg 0x10: [mem 0xb4000000-0xb47fffff 64bit]
pci 0000:19:00.0: PME# supported from D0
pci 0000:19:00.0: 4.000 Gb/s available PCIe bandwidth, limited by 5 GT/s x1 link at 0000:18:10.0
(capable of 16.000 Gb/s with 5 GT/s x4 link)
pci 0000:19:00.0: vgaarb: VGA device added: decodes=io+mem,owns=none,locks=none
pci 0000:19:00.0: Adding to iommu group 37
pci 0000:18:10.0: PCI bridge to [bus 19]
pci 0000:18:10.0:   bridge window [mem 0xb4000000-0xb47fffff]
pci 0000:1a:00.0: [1304:00e6] type 00 class 0x058000
pci 0000:1a:00.0: reg 0x10: [mem 0xb4800000-0xb48fffff]
pci 0000:1a:00.0: Adding to iommu group 37
pci 0000:18:11.0: PCI bridge to [bus 1a]
pci 0000:18:11.0:   bridge window [mem 0xb4800000-0xb48fffff]
pci 0000:1b:00.0: [10b5:87b0] type 00 class 0x068000
pci 0000:1b:00.0: reg 0x10: [mem 0xb4900000-0xb493ffff]
pci 0000:1b:00.0: Setting PLX NTB proxy ID aliases
pci 0000:1b:00.0: Enabling fixed DMA alias for devfn range from 00.0 to 1f.7
pci 0000:1b:00.0: Adding to iommu group 37
pci 0000:18:12.0: PCI bridge to [bus 1b]
pci 0000:18:12.0:   bridge window [mem 0xb4900000-0xb49fffff]
pci 0000:17:00.0: BAR 14: assigned [mem 0xb4000000-0xb4bfffff]
pci 0000:17:00.0: BAR 0: assigned [mem 0xb4c00000-0xb4c3ffff]
pci 0000:18:10.0: BAR 14: assigned [mem 0xb4000000-0xb47fffff]
pci 0000:18:11.0: BAR 14: assigned [mem 0xb4800000-0xb48fffff]
pci 0000:18:12.0: BAR 14: assigned [mem 0xb4900000-0xb49fffff]
pci 0000:19:00.0: BAR 0: assigned [mem 0xb4000000-0xb47fffff 64bit]
pci 0000:18:10.0: PCI bridge to [bus 19]
pci 0000:18:10.0:   bridge window [mem 0xb4000000-0xb47fffff]
pci 0000:1a:00.0: BAR 0: assigned [mem 0xb4800000-0xb48fffff]
pci 0000:18:11.0: PCI bridge to [bus 1a]
pci 0000:18:11.0:   bridge window [mem 0xb4800000-0xb48fffff]
pci 0000:1b:00.0: BAR 0: assigned [mem 0xb4900000-0xb493ffff]
pci 0000:18:12.0: PCI bridge to [bus 1b]
```

```
pci 0000:18:12.0:   bridge window [mem 0xb4900000-0xb49fffff]
pci 0000:17:00.0: PCI bridge to [bus 18-1b]
pci 0000:17:00.0:   bridge window [mem 0xb4000000-0xb4bfffff]
device tap2 left promiscuous mode
jnpr-re-re: port 2(tap2) entered disabled state
ctrlbr0: port 1(vnet100) entered disabled state
device vnet100 left promiscuous mode
ctrlbr0: port 1(vnet100) entered disabled state
[FAILED] Failed unmounting /.old/tdev/.union/.s.
                                                [ printk: systemd-shutdow: 48 output lines
suppressed due to ratelimiting
 OK  ] Unmounted /.old/tmp.
                          [  OK  ]systemd-shutdown[1]: Syncing filesystems and block devices.
 Unmounted /.old/vm.
                    [FAILED] Failed usystemd-shutdown[1]: Sending SIGTERM to remaining
processes...
nmounting /junos.
                 [  OK  ] Unmounted ystemd-journald[400]: Received SIGTERM from PID 1 (systemd-
shutdow).
;1;39m/spare.
             [  OK  ] Unmounted /tmp/junos-vfmBX.
                                                 [  OK  ] Unmounted /var/volatile.
                                                                                      [FAILED]
Failed unmounting /vm.

                         Unmounting /.old...

                                                      Unmounting Temporary Directory (/tmp)...
    [FAILED] Failed unmounting /.old.
                                     [  OK  ] Stopped target Local File Systems (Pre).
                                                                                          [  OK  ]
Stopped Create Static Device Nodes in /dev.

                                         [  OK  ] Stopped Create System Users.

                                                                                 [  OK  ] Stopped
Remount Root and Kernel File Systems.
                                           [  OK  ] Unmounted Temporary Directory (/tmp).

[  OK  ] Stopped target Swap.

                      [  OK  ] Reached target Shutdown.
```

```
                                                        [  OK  ] Reached target Unmount All
Filesystems.
             [  OK  ] Reached target Final Step.
                                                [  OK  ] Started Power-Off.
                                                                                [  OK  ] Reached
target Power-Off.

            systemd-shutdown[1]: Sending SIGKILL to remaining processes...
systemd-shutdown[1]: Unmounting file systems.
[55498]: Remounting '/vm' read-only in with options '(null)'.
EXT4-fs (dm-2): re-mounted. Opts: (null)
[55499]: Unmounting '/vm'.
[55500]: Remounting '/junos' read-only in with options '(null)'.
EXT4-fs (dm-1): re-mounted. Opts: (null)
[55501]: Unmounting '/junos'.
[55502]: Remounting '/' read-only in with options 'si=e65be7872f990b8b'.
[55503]: Remounting '/.old/tdev/.union/.s' read-only in with options '(null)'.
[55503]: Failed to remount '/.old/tdev/.union/.s' read-only: Device or resource busy
[55504]: Unmounting '/.old/tdev/.union/.s'.
[55504]: Failed to unmount /.old/tdev/.union/.s: Device or resource busy
[55505]: Remounting '/.old' read-only in with options '(null)'.
EXT4-fs (sda3): re-mounted. Opts: (null)
[55506]: Unmounting '/.old'.
[55506]: Failed to unmount /.old: Device or resource busy

Booting from Flash A

FPGA Reset Reason = 0x83

Primary BIOS version CICL_P_AVG0_00.11.01

Total Memory Size = 128 GB

Checking Primary BIOS code integrity...Passed!
```

Release Information

request vmhost power-off

IN THIS SECTION

- Syntax | 15525
- Description | 15525
- Options | 15525
- Required Privilege Level | 15526
- Sample Output | 15526
- Release Information | 15526

Syntax

```
request vmhost power-off
<other routing-engine>
```

Description

Power off the Routing Engine on which Junos OS and the host OS are running. In a PTX3000, the Routing and Control Board is powered off.

NOTE: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

Options

none—Power off the Routing Engine immediately.

other-routing-engine (Optional) Power off the other Routing Engine on which the Junos OS and the host OS are running. For example, if you issue the command from the primary Routing Engine, the backup Routing Engine is powered off. Similarly, if you issue the command from the backup Routing Engine, the primary Routing Engine is powered off.

Required Privilege Level

maintenance

Sample Output

request vmhost power-off

```
request vmhost power-off
```

Release Information

Command introduced in Junos OS Release 15.1F3.

> **NOTE**: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

request system power-off

vmhost | 30895

request vmhost power-off (NFX Series)

IN THIS SECTION

- Syntax | 15527
- Description | 15527
- Required Privilege Level | 15527
- Sample Output | 15527
- Release Information | 15527

Syntax

```
request vmhost power-off
```

Description

Shut down the Junos OS software and the host OS.

Required Privilege Level

maintenance

Sample Output

request vmhost power-off

```
user@host> request vmhost power-off
Power-off the vmhost ? [yes,no] (no) yes

Initiating vmhost shutdown...  ok
Initiating Junos shutdown...  shutdown: [pid 3884]
Shutdown NOW!
ok

*** FINAL System shutdown message from root@host ***

System going down IMMEDIATELY
...
...
```

Release Information

Command introduced in Junos OS Release 18.1R1.

NOTE: `request vmhost power-on` is not supported on NFX150 and NFX250 (NG) devices.

request vmhost power-on

IN THIS SECTION

- Syntax | 15528
- Description | 15528
- Options | 15528
- Required Privilege Level | 15528
- Sample Output | 15529
- Release Information | 15529

Syntax

```
request vmhost power-on other-routing-engine
```

Description

Power on the Routing Engine on which Junos OS and the host OS are running. In a PTX3000, the Routing and Control Board is powered on.

Options

other-routing-engine Power on the other Routing Engine on which the Junos OS and the host OS are running. For example, if you issue the command from the primary Routing Engine, the backup Routing Engine is powered on. Similarly, if you issue the command from the backup Routing Engine, the primary Routing Engine is powered on.

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

Required Privilege Level

maintenance

Sample Output

request vmhost power-on

```
user@host> request vmhost power-on other-routing-engine
Routing Engine 1 power-on initiated, use "show chassis routing-engine" to verify
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

request vmhost power-off | 15525

request vmhost reboot | 15529

request vmhost reboot

IN THIS SECTION

- Syntax | 15530
- Description | 15530
- Options | 15530
- Required Privilege Level | 15531
- Sample Output | 15531
- Release Information | 15531

Syntax

```
request vmhost reboot
<disk1>
<disk2>
<network>
<re0 | re1>
<routing engine>
<usb>
```

Description

Reboot both the Junos OS and the host OS running on the device.

Options

none—Reboot the device software immediately.

disk1 (Optional) Reboot both Junos OS and the host OS on the Routing Engine and boot the Routing Engine from the primary disk.

disk2 (Optional) Reboot both Junos OS and the host OS on the Routing Engine and boot the Routing Engine from backup disk.

network (Optional) Reboot both Junos OS and the host OS on the Routing Engine and boot the Routing Engine from network by using the PXE boot method.

re0 | re1 (Optional) On routers that support dual or redundant Routing Engines, reboot both Junos OS and the host OS on the Routing Engine in slot 0 (`re0`) or on the the Routing Engine in slot 1 (`re1`).

routing-engine (Optional) Specify the Routing Engine on which Junos OS and the host OS are to be rebooted. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Both Routing Engines.

- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

usb (Optional) Reboot both Junos OS and the host OS on the Routing Engine using the USB installation media.

Required Privilege Level

maintenance

Sample Output

request vmhost reboot

```
request vmhost reboot
<disk1>
<disk2>
<in>
<network>
<usb>
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

request system reboot (Junos OS) | **15140**

request vmhost power-on | **15528**

request vmhost reboot (NFX Series)

IN THIS SECTION

- Syntax | 15532
- Description | 15532
- Required Privilege Level | 15532
- Sample Output | 15532
- Release Information | 15534

Syntax

```
request vmhost reboot [disk1 | disk2] [primary | alternate]
```

Description

Reboot the Junos OS software and the host OS from the specified disk and the partition within the disk.

Required Privilege Level

maintenance

Sample Output

request vmhost reboot (NFX150)

```
user@host> request vmhost reboot disk1 primary
Reboot the vmhost ? [yes,no] (no) yes

Switching boot to disk1 primary
Initiating vmhost reboot...  ok
Stopping jrestartd: [  OK  ]
/etc/init.d/functions: line 286: usleep: command not found
Initiating Junos shutdown...  shutdown: [pid 12151]
```

```
Shutdown NOW!
user@host> request vmhost reboot disk1 alternate
Reboot the vmhost ? [yes,no] (no) yes

Switching boot to disk1 alternate
Initiating vmhost reboot...  ok
Stopping jrestartd: [  OK  ]
/etc/init.d/functions: line 286: usleep: command not found
Initiating Junos shutdown...  shutdown: [pid 16368]
Shutdown NOW!
```

request vmhost reboot (NFX250 NextGen)

```
user@host> request vmhost reboot disk1 primary
Reboot the vmhost ? [yes,no] (no) yes

Switching boot to disk1 primary
Initiating vmhost reboot...  ok
Stopping jrestartd: [  OK  ]
/etc/init.d/functions: line 286: usleep: command not found
Initiating Junos shutdown...  shutdown: [pid 52663]
Shutdown NOW!

user@host> request vmhost reboot disk1 alternate
Reboot the vmhost ? [yes,no] (no) yes

Switching boot to disk1 alternate
Initiating vmhost reboot...  ok
Stopping jrestartd: [  OK  ]
/etc/init.d/functions: line 286: usleep: command not found
Initiating Junos shutdown...  shutdown: [pid 18763]
Shutdown NOW!
```

request vmhost reboot (NFX350)

```
user@host> request vmhost reboot disk1  primary
Reboot the vmhost ? [yes,no] (no) yes

Switching boot to disk1 primary
Initiating vmhost reboot...  ok
```

```
Stopping jrestartd: [  OK  ]
/etc/init.d/functions: line 286: usleep: command not found
Initiating Junos shutdown...  shutdown: [pid 15575]
Shutdown NOW!
user@host> request vmhost reboot disk1  alternate
Reboot the vmhost ? [yes,no] (no) yes

Switching boot to disk1 alternate
Initiating vmhost reboot...  ok
Stopping jrestartd: [  OK  ]
/etc/init.d/functions: line 286: usleep: command not found
Initiating Junos shutdown...  shutdown: [pid 14189]
Shutdown NOW!

user@host> request vmhost reboot disk2  primary
Reboot the vmhost ? [yes,no] (no) yes

Switching boot to disk2 primary
Initiating vmhost reboot...  ok
Stopping jrestartd: [  OK  ]
/etc/init.d/functions: line 286: usleep: command not found
Initiating Junos shutdown...  shutdown: [pid 12956]
Shutdown NOW!
user@host> request vmhost reboot disk2  alternate
Reboot the vmhost ? [yes,no] (no) yes

Switching boot to disk2 alternate
Initiating vmhost reboot...  ok
Stopping jrestartd: [  OK  ]
/etc/init.d/functions: line 286: usleep: command not found
Initiating Junos shutdown...  shutdown: [pid 13025]
Shutdown NOW!
```

Release Information

Command introduced in Junos OS Release 18.1R1.

request vmhost remote-file-copy

IN THIS SECTION

- Syntax | 15535
- Description | 15535
- Options | 15535
- Additional Information | 15536
- Required Privilege Level | 15537
- Sample Output | 15537
- Release Information | 15539

Syntax

```
request vmhost remote-file-copy
                 interface {
                       (in-band | out-of-band);
                       }
                   source source-path;
                   destination destination-path;
```

Description

Enable remote file copy between the hypervisor and a remote fileserver. The configuration allows creating an external network interface directly on the hypervisor and associating an IPv4 address with it.

The remote file copy feature allows upload or download operations from or into the NFX device. These operations are initiated from the NFX device, not from the remote fileserver.

Options

- in-band—Enable external network access through one of the front panel revenue ports.
- out-of-band—Enable external network access through the front panel management interface.

- source *source-path*—Name of the source path from which the file is to be copied. The source path should refer to a specific filename, not a directory, and without any wild cards. If the *source-path* is local, the file-copy is an upload operation. If the *source-path* is remote, the file-copy is a download operation.
- destination*destination-path*—Name of the destination path to which the host file is copied. The destination path must refer to a specific filename.

Additional Information

Either *source-path* or *destination-path* must be remote while the other should be local. Consider the following guidelines while creating a local path:

- Local path should refer to a file under **/var/public/** directory.
- Avoid special characters. You can include alphanumeric characters, hyphen, underscore, and dot.
- Avoid relative path components **..** and **/./**
- Avoid ending the path with **/.**

Following are the examples of a local path:

- **/var/public/myfile**
- **/var/public/mydir/my_file**
- **/var/public/12-hours.data**

Consider the following guidelines while creating a local path:

- A remote path must be in the following format:
  [<remote_user>[:<remote_user_pwd>]@]<protocol>:// <remote_ip>[:<remote_port>]:<remote_path>

  The protocols supported are, SCP, HTTP, HTTPS, and FTP.

- `remote_user` is optional, depending on protocol and the remote fileserver configuration.

  For example, `remote_user` is required when scp protocol is used whereas `remote_user` is unnecessary when HTTP protocol is used.

- Depending on the protocol and the remote fileserver configuration, when `remote_user` is specified, you might have to specify `remote_user_pwd`.

  For example, when ftp is used, username and password are required.

- `remote_port` is optional and could be skipped if the remote fileserver is reachable on a standard port for the specified protocol.

Following are the examples of a remote path:

```
user@scp://10.10.10.2:22:/home/user/myfile
user@scp://10.10.10.2:/home/user/myfile
http://10.10.10.2:80:/myfile
https://10.10.10.2:80:/myfile
user:user@ftp://10.10.10.2:21:/home/user/myfile
user@ftp://10.10.10.2:21:/home/user/myfile
```

Required Privilege Level

Maintenance

Sample Output

request vmhost remote-file-copy (To upload a file from an NFX device to remote fileserver using an in-band interface and scp protocol, with user interaction)

```
user@host> request vmhost remote-file-copy interface in-band source /var/public/1G.dat
destination root@scp://10.10.10.11:22:/var/www/html/1G.dat
The authenticity of host '10.10.10.11 (10.10.10.11)' can't be established.
ECDSA key fingerprint is SHA256:1ldwJbrMb47DBsf+wwG4W06a00yfXt9hAS6ozIVAPVc.
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '10.10.10.11' (ECDSA) to the list of known hosts.
user@host@10.10.10.11's password:
1G.dat
100% 1024MB  42.8MB/s   00:23
```

request vmhost remote-file-copy (To upload a file from an NFX device to remote fileserver using an out-of-band interface and scp protocol, with user interaction)

```
user@host> request vmhost remote-file-copy interface out-of-band source /var/public/1G.dat
destination root@scp://10.204.103.43:22:/var/www/html/1G.dat
user@host@10.204.103.43's password:
```

```
1G.dat  100% 1024MB  34.9MB/s   00:29
```

request vmhost remote-file-copy (To download a file from the remote fileserver to an NFX device using an in-band interface and http protocol)

```
user@host> request vmhost remote-file-copy interface in-band source http://
10.10.10.11:80:/1G.dat destination /var/public/1G.dat
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100 1024M  100 1024M    0     0  77.5M      0  0:00:13  0:00:13 --:--:-- 75.0M
```

request vmhost remote-file-copy (To download a file from the remote fileserver to an NFX device using an in-band interface and https protocol)

```
user@nhost> request vmhost remote-file-copy interface in-band source https://10.10.10.11:/1G.dat
destination /var/public/1G.dat
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100 1024M  100 1024M    0     0  22.9M      0  0:00:44  0:00:44 --:--:-- 23.4M
```

request vmhost remote-file-copy (To upload a file from an NFX device to the remote fileserver using an in-of-band interface and ftp protocol, with user interaction)

```
user@host> request vmhost remote-file-copy interface in-band source /var/public/1G.dat
destination user1@ftp://10.10.10.11:21:/home/user1/1G.dat
Enter host password for user 'user1':
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100 1024M    0     0  100 1024M      0   105M  0:00:09  0:00:09 --:--:--  110M
```

request vmhost remote-file-copy (To upload a file from an NFX device to the remote fileserver using an in-band interface and ftp protocol, without user interaction)

```
user@host> request vmhost remote-file-copy interface in-band source /var/public/1G.dat
destination user1:user1@ftp://10.10.10.11:21:/home/user1/1G.dat
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100 1024M    0     0  100 1024M      0   106M  0:00:09  0:00:09 --:--:--  111M
```

Release Information

Command introduced in Junos OS Release 24.2R1 for NFX150, NFX250, and NFX350 devices.

request vmhost snapshot

IN THIS SECTION

- Syntax | 15539
- Description | 15540
- Options | 15540
- Required Privilege Level | 15541
- Sample Output | 15541
- Release Information | 15542

Syntax

```
request vmhost snapshot
<partition>
<re0 | re1>
<recovery>
<routing-engine>
```

Description

Create a recovery snapshot of the currently running and active file system partitions on the backup disk to recover the primary disk in case of failure.

On the device, back up the snapshot of the host OS image along with the Junos OS image. In case of failure of the primary disk, you can boot from the image available in the backup disk and then recover the primary disk with the snapshot created using the `recovery` option.

Options

none—Create a snapshot from the current disk to the target disk without partitioning the target disk. Contents on target disk is lost.

`partition` (Optional) Create a snapshot from the current disk to target disk and partition the target disk. Contents on the target disk are lost

re0 | re1 (Optional) Create a snapshot from the current disk to target disk and partition the target disk on Routing Engine in slot 0 (`re0`) or on the Routing Engine in slot 1 (`re1`). The snapshot is taken without partitioning the target disk on corresponding Routing Engines. Contents on the target disk on the Routing Engines are lost.

For PTX1000 routers, since there is only one RE, re0|re1 is not supported.

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

recovery (Optional) Recover the primary disk from the snapshot content stored in the backup disk. This option is applicable only when the Routing engine is booted from backup disk. Contents in the primary disk are lost.

routing-engine (Optional) Specify the Routing Engine on which the snapshot is to be created. The following options are available:

- `backup`—Backup Routing Engine.
- `both`—Both Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—Other Routing Engine.

Required Privilege Level

maintenance

Sample Output

request vmhost snapshot

```
user@host> request vmhost snapshot
warning: Existing data on the target may be lost
Proceed ? [yes,no] (no) yes

warning: Proceeding with vmhost snapshot
Current root details,           Device sda, Label: jrootb_P, Partition: sda4
Snapshot admin context from current boot disk to target disk ...
Proceeding with snapshot on secondary disk
Mounting device in preparation for snapshot...
Cleaning up target disk for snapshot ...
Creating snapshot on target disk from current boot disk ...
Snapshot created on secondary disk.
Software snapshot done
```

request vmhost snapshot recovery

```
user@host> request vmhost snapshot recovery
warning: Existing data on the target may be lost
Proceed ? [yes,no] (no) yes

warning: Proceeding with vmhost snapshot
Current root details,           Device sdb, Label: jrootb_S, Partition: sdb4
Snapshot admin context from current boot disk to target disk ...
Proceeding with snapshot on primary disk
Mounting device in preparation for snapshot...
Cleaning up target disk for snapshot ...
Creating snapshot on target disk from current boot disk ...
Primary disk is recovered now. Please issue "request vmhost reboot" to boot from the primary
disk.
Software snapshot done
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

show vmhost snapshot | 27821

request vmhost software abort in-service-upgrade

IN THIS SECTION

- Syntax | 15542
- Description | 15542
- Options | 15543
- Required Privilege Level | 15543
- Release Information | 15543

Syntax

```
request vmhost software abort in-service-upgrade
```

Description

Terminate unified in-service software upgrade (unified ISSU). The unified ISSU must be in progress and you must issue this command from a router session other than the one on which you issued the `request vmhost software in-service-upgrade` command to launch the unified ISSU.

Options

in-service-upgrade	Terminate unified ISSU.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 15.1F3.

> **NOTE**: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

- request system software abort
- request vmhost software in-service-upgrade | 15551

request vmhost software add

IN THIS SECTION

- MX Series | 15544
- PTX Series | 15544
- Description | 15544
- Options | 15545
- Additional Information | 15546
- Required Privilege Level | 15547
- Sample Output | 15547
- Release Information | 15547

MX Series

```
request vmhost software add package-name
<no-validate>
<re0 | re1>
<reboot>
<set [package-names]>
```

PTX Series

```
request vmhost software add package-name
<in>
<log_collector.sh>
<no-validate>
<reboot>
<set [package-names]>
<unlink>
<upgrade-to-model>
```

Description

Install Junos OS and host software packages on the device.

For installing the host software as well as Junos OS, specify the package name **junos-vmhost-install-*x*.tgz** in the `request vmhost software add` command. Junos OS installation alone can be achieved by specifying the regular package name **junos-install-*x*.tgz** in the `request system software add` command. However, installation using the vmhost package is recommended as it upgrades both the host software and Junos OS.

To upgrade to Junos OS Release 21.2R1, you must include the `no-validate` option when using this command to install the package.

> **NOTE**: You must load the PTX1000, PTX10008, PTX10016, PTX10002-60C, and QFX10002-60C devices only with **junos-vmhost-install-*x*.tgz** package using the `request vmhost software add` command. The **junos-vmhost-install-*x*.tgz** package upgrades both the host software and Junos OS. The PTX1000, PTX10008, PTX10016, PTX10002-60C, and QFX10002-60C devices do not support Junos only upgrade. If you try to load Junos only

image, then these devices go down or vmhost commands do not work or the device state is unpredictable.

Options

none—Install Junos OS and host software packages on the Routing Engine on the local Virtual Chassis member.

package-name Location from which the software package or bundle is to be installed. For example:

- **/var/tmp/*package-name***—For a software package or bundle that is being installed from a local directory on the device.
- ***protocol://hostname/pathname/package-name***—For a software package or bundle that is to be downloaded and installed from a remote location. Replace ***protocol*** with one of the following:
  - **ftp**—File Transfer Protocol (FTP).
    Use **ftp://*hostname/pathname/package-name***. To specify authentication credentials, use **ftp://<*username*>:<*password*>@*hostname/pathname/package-name***. To have the system prompt you for the password, specify `prompt` in place of the password. If a password is required, and you do not specify the password or `prompt`, an error message is displayed.
  - **http**—Hypertext Transfer Protocol (HTTP).
    Use **http://*hostname/pathname/package-name***. To specify authentication credentials, use **http://<*username*>:<*password*>@*hostname/pathname/package-name***. If a password is required and you omit it, you are prompted for it.
  - **scp**—Secure Copy Protcol (SCP)(not available for limited editions).
    Use **scp://*hostname/pathname/package-name***. To specify authentication credentials, use **scp://<*username*>:<*password*>@*hostname/pathname/package-name***.

NOTE:

- The ***pathname*** in the protocol is the relative path to the user's home directory on the remote system and not the root directory.

- Do not use the **scp** protocol in the `request vmhost software add` command to download and install a software package or bundle from a remote location. The software upgrade is handled by the mgd process that does not support SCP.

  To install a software package or bundle from a remote location:

  1. Use the `file copy` command to copy the software package or bundle from the remote location to the **/var/tmp** directory on the hard disk:
     file copy scp:// *source/ package-name* /var/tmp
  2. Install the software package or bundle by using the `request vmhost software add` command:
     request vmhost software add /var/tmp/*package-name*

no-validate (Optional) When loading a software package or bundle with a different release, suppress the default behavior of the `validate` option.

To upgrade to Junos OS Release 21.2R1, you must include the `no-validate` option when using this command to install the package.

re0 | re1 (Optional) Load a software package or bundle on the Routing Engine in slot 0 (`re0`) or the Routing Engine in slot 1 (`re1`).

> **NOTE**: The option `re1` is not supported on the PTX1000 Packet Transport Router.

reboot (Optional) After adding the software package or bundle, reboot the system.

set [*package-names*] (Optional)

Additional Information

Before upgrading the software on the device, when you have a known stable system, issue the `request vmhost snapshot` command to back up the software. After you have upgraded the software on the device and are satisfied that the new package or bundle is successfully installed and running, issue the `request vmhost snapshot` command again to back up the new software to the backup disk.

After you run the `request vmhost snapshot` command, you cannot return to the previous version of the snapshot, because the previous snapshot is replaced by the new snapshot.

Before installing software on a device that has one or more custom YANG data models added to it, back up and remove the configuration data corresponding to the custom YANG data models from the active configuration. For more information see *Managing YANG Packages and Configurations During a Software Upgrade or Downgrade*.

Required Privilege Level

maintenance

Sample Output

request vmhost software add (Multiple Packages)

```
user@host> request vmhost software add set [/var/tmp/junos-vmhost-install-ptx-
x86-64-15.1F-20160518.0.tgz /var/tmp/junos-vmhost-jdiag-15.1F-20160518.0.tgz no-validate
Verified /var/tmp/junos-vmhost-install-ptx-x86-64-15.1F-20160518.0.tgz signed by
PackageDevelopmentEc_2016
Copied the config and other data to the aux disk.
Transfer junos-host-upgrade.sh
Transfer Done
Transfer /packages/db/pkginst.7286/junos-vmhost-install*.tgz
Transfer Done
Starting upgrade...
Preparing for upgrade...
/tmp/pkg-1dX/unpack/install/
...
...
..
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

request system software add (Junos OS) | **15216**

request vmhost software rollback | 15556

request vmhost software add (NFX Series)

IN THIS SECTION

- Syntax | 15548
- Description | 15548
- Options | 15548
- Required Privilege Level | 15549
- Sample Output | 15549
- Release Information | 15551

Syntax

```
request vmhost software add package-name <in>| <no-validate>| <reboot>| <set>| <unlink>|
<upgrade-to-model model-number>
```

Description

Install or upgrade the Junos OS and host software packages on the device.

Options

- in—(Optional) Number of minutes to delay before the reboot operation.
- no-validate—(Optional) When loading a software package or bundle with a different release, suppress the default behavior of the `validate` option.
- reboot—(Optional) After adding the software package or bundle, reboot the system.
- set—(Optional) List of URLs or pathnames corresponding to the software packages.
- unlink—(Optional) Removes the software package after successful installation.

- upgrade-to-model—(Optional) *model number*—(Optional) Name of the model to upgrade to.

Required Privilege Level

maintenance

Sample Output

request vmhost software add (NFX150)

```
user@host> request vmhost software add /var/public/jinstall-host-nfx-3-x86-64-18.1R1.8-secure-
signed.tgz no-validate reboot
Verified jinstall-host-nfx-3-x86-64-18.1R1.8-secure-signed signed by PackageProductionEc_2018
method ECDSA256+SHA256
Pushing Junos image package to the host...
File already present in Host. Skipping pushing the image
Mounting primary partitions to stage upgrade operation
Installing /mnt/.share/lshare/public/pkginst.7565/install-media-nfx-3-junos-18.1R1.8-secure.tgz
Extracting the package ...
..
..
```

request vmhost software add (NFX250 (NG))

```
user@host> request vmhost software add /var/public/jinstall-host-nfx-3-x86-64-18.4R1.8-secure-
signed.tgz
Verified jinstall-host-nfx-3-x86-64-18.4R1.8-secure-signed signed by PackageProductionEc_2018
method ECDSA256+SHA256
Pushing Junos image package to the host...
File already present in Host. Skipping pushing the image
Mounting alternate partitions to stage upgrade operation
Installing /mnt/.share/lshare/public/pkginst.39634/install-media-nfx-3-junos-18.4R1.8-secure.tgz
Extracting the package ...
============================================
Host OS upgrade is FORCED
Current Host kernel version : 4.1.27-rt30-WR8.0.0.25_ovp
Package Host kernel version : 4.1.27-rt30-WR8.0.0.25_ovp
Current Host version        : 3.0.3
Package Host version        : 3.0.3
Min host version required for applications: 3.0.2
```

```
============================================
Validate linux image...
upgrade_platform: -------------------
upgrade_platform: Parameters passed:
upgrade_platform: silent=0
upgrade_platform: package=/var/tmp/tmp.rV7S1sxWedjunos_cli_upg/jinstall-nfx-3-junos-18.4R1.8-
secure-linux.tgz
upgrade_platform: clean install=0
upgrade_platform: on primary   =0
upgrade_platform: clean upgrade=0
upgrade_platform: Need reboot after staging=1
upgrade_platform: -------------------
upgrade_platform:
upgrade_platform: Checking input /var/tmp/tmp.rV7S1sxWedjunos_cli_upg/jinstall-nfx-3-
junos-18.4R1.8-secure-linux.tgz ...
upgrade_platform: Input package /var/tmp/tmp.rV7S1sxWedjunos_cli_upg/jinstall-nfx-3-
junos-18.4R1.8-secure-linux.tgz is valid.
Secure Boot is enforced.
ALLOW:usr/secureboot/grub/BOOTX64.EFI
ALLOW:boot/bzImage-intel-x86-64.bin
ALLOW:boot/initramfs.cpio.gz
Setting up Junos host applications for installation ...
Current junos instance is 0
Installing Host OS ...
upgrade_platform: -------------------
upgrade_platform: Parameters passed:
upgrade_platform: silent=0
upgrade_platform: package=/var/tmp/jinstall-nfx-3-junos-18.4R1.8-secure-linux.tgz
upgrade_platform: clean install=0
upgrade_platform: on primary   =0
upgrade_platform: clean upgrade=0
upgrade_platform: Need reboot after staging=0
upgrade_platform: -------------------
upgrade_platform:
upgrade_platform: Checking input /var/tmp/jinstall-nfx-3-junos-18.4R1.8-secure-linux.tgz ...
upgrade_platform: Input package /var/tmp/jinstall-nfx-3-junos-18.4R1.8-secure-linux.tgz is valid.
Secure Boot is enforced.
ALLOW:usr/secureboot/grub/BOOTX64.EFI
ALLOW:boot/bzImage-intel-x86-64.bin
ALLOW:boot/initramfs.cpio.gz
upgrade_platform: Backing up boot assets..
upgrade_platform: Staging the upgrade package - /var/tmp/jinstall-nfx-3-junos-18.4R1.8-secure-
linux.tgz..
```

```
upgrade_platform: Checksum verified and OK...
upgrade_platform: Staging of /var/tmp/jinstall-nfx-3-junos-18.4R1.8-secure-linux.tgz completed
upgrade_platform: System needs *REBOOT* to complete the upgrade
Host OS upgrade staged. Reboot the system to complete installation!
```

Release Information

Command introduced in Junos OS Release 18.1R1.

request vmhost software in-service-upgrade

IN THIS SECTION

- Syntax | 15551
- Description | 15551
- Options | 15552
- Additional Information | 15553
- Required Privilege Level | 15553
- Sample Output | 15553
- Sample Output | 15554
- Release Information | 15555

Syntax

```
request vmhost software in-service-upgrade package-name
< no-old-master-upgrade>
< reboot>
< verbose>
```

Description

Perform a unified in-service software upgrade (ISSU). A unified ISSU enables you to upgrade from one Junos OS release and host OS release to another with no disruption on the control plane and with

minimal disruption of traffic. For an unified ISSU, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled.

Starting inJunos OS Release 18.2R1, MX10003 routers support unified in-service software upgrade (ISSU) using `request vmhost software in-service-upgrade` command. MX10003 does not support upgrading Junos OS only image using request system software in-service-upgrade command.

NOTE: On MX10003 routers:

- ISSU is not supported on MACsec MIC (JNP-MIC1-MACSEC).
- ISSU is not supported for the interfaces that are configured with 1-Gigabit Ethernet mode. If ISSU upgrade is carried out in 1-Gigabit Ethernet mode, then the behavior is unexpected and traffic loss can be expected.
- ISSU is not supported on timing protocols (like, Precision Time Protocol and Synchronous Ethernet), MACsec protocols, and BBE protocols. If these protocols are already enabled, then it will not work after ISSU is enabled.
- The MAC statistics (retrieved using "show interfaces extensive" on page 20676 command) are reset during ISSU which means that the MAC statistics does not provide the correct statistics after ISSU.

Options

package-name

Location from which the software package or bundle is to be installed. For example:

- `/var/tmp/`*`package-name`*— For a software package or bundle that is being installed from a local directory on the router.
- *`protocol`*`://`*`hostname`*`/`*`pathname`*`/`*`package-name`*—For a software package or bundle that is to be downloaded and installed from a remote location. Replace *`protocol`* with one of the following:
  - `ftp`—File Transfer Protocol (FTP).
  - `http`—Hypertext Transfer Protocol (HTTP).
  - `scp`—Secure Copy Protocol (SCP) (not available for limited editions).

no-old-master-upgrade

(Optional) When the `no-old-master-upgrade` option is included, after the backup Routing Engine is rebooted with the new software package and a switchover occurs to make it the new primary Routing Engine, the former primary (new backup) Routing Engine is not upgraded to the new software. In this case, you must manually upgrade the former

primary (new backup) Routing Engine. If you do not include the `no-old-master-upgrade` option, the system automatically upgrades the former primary Routing Engine.

reboot (Optional) Automatically reboot the former primary (new backup) Routing Engine after the ISSU. If you don not include the `reboot` option in the command, you must manually reboot the former primary (new backup) Routing Engine by using the `request vmhost reboot` command.

verbose (Optional) (MX Series) Use this option to display the daemon related information during the upgrade.

Additional Information

The following conditions apply to unified ISSU:

- Unified ISSU is not supported on every platform. For a list of supported platforms, see Unified ISSU System Requirements.
- Unsupported PICs are restarted during a unified ISSU on certain routing devices.
- During a unified ISSU, any unsupported protocols running on the device causes packet loss.
- During a unified ISSU, you cannot bring any PICs online or take them offline on certain routing devices.

Required Privilege Level

maintenance

Sample Output

request vmhost software in-service-upgrade

```
user@host> request vmhost software in-service-upgrade
/var/tmp/junos-vmhost-install-ptx-x86-64-15.1F5.6.tgz reboot
Chassis ISSU Check Done
[Feb 24 01:12:09]: Starting VMHOST ISSU
[Feb 24 01:12:09]:ISSU: Validating Image
FPC 2 will be offlined (In-Service-Upgrade not supported)
FPC 11 will be offlined (In-Service-Upgrade not supported)
MIC 11/0 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes
```

```
Junos Validation begin. Procedure will take few minutes.
  ...
  ...
```

Sample Output

request vmhost software in-service-upgrade verbose

```
user@host> request vmhost software in-service-upgrade verbose


  ...
  ...
Oct 26 15:20:02 [INFO ] Verified py-extensions signed by PackageDevelopmentEc_2017 method
ECDSA256+SHA256
Oct 26 15:20:02 [INFO ] Adding py-extensions-
x86-32-20171024.002108_builder_release_174_throttle ...
Oct 26 15:20:02 [INFO ] Verified vrr-mx signed by PackageDevelopmentEc_2017 method
ECDSA256+SHA256
Oct 26 15:20:02 [INFO ] NOTICE: 'pending' set will be activated at next reboot...
Oct 26 15:20:02 [INFO ] [Oct 26 02:36:10]:ISSU: Installing package /var/tmp/junos-install-mx-
x86-64-17.4-20171024.0.tgz on re1 done
Oct 26 15:20:02 [INFO ] [Oct 26 02:36:10]:ISSU: Rebooting Backup RE
Oct 26 15:20:02 [INFO ]
Oct 26 15:20:02 [INFO ] Rebooting re1
Oct 26 15:20:02 [INFO ] [Oct 26 02:36:11]:ISSU: Backup RE Prepare Done
Oct 26 15:20:02 [INFO ] [Oct 26 02:36:11]:ISSU: Waiting for Backup RE reboot
Oct 26 15:20:02 [INFO ] [Oct 26 02:39:26]:ISSU: Backup RE reboot done. Backup RE is up
Oct 26 15:20:02 [INFO ] [Oct 26 02:39:26]:ISSU: Waiting for Backup RE state synchronization
Oct 26 15:20:02 [INFO ] [Oct 26 02:39:51]:ISSU: Backup RE state synchronization done
Oct 26 15:20:02 [INFO ] [Oct 26 02:39:51]:ISSU: GRES operational
Oct 26 15:20:02 [INFO ] [Oct 26 02:40:52]: "Initiating Chassis In-Service-Upgrade"
Oct 26 15:20:02 [INFO ] Chassis ISSU Started
Oct 26 15:20:02 [INFO ] [Oct 26 02:40:57]:ISSU: Preparing Daemons
Oct 26 15:20:02 [INFO ] [Oct 26 02:40:57]: Daemon [rpd] transitioned to READY state
Oct 26 15:20:02 [INFO ] [Oct 26 02:40:57]: Daemon [lfmd] transitioned to READY state
Oct 26 15:20:02 [INFO ] [Oct 26 02:40:57]: Daemon [l2cpd] transitioned to READY state
Oct 26 15:20:02 [INFO ] [Oct 26 02:40:57]: Daemon [smid] transitioned to READY state
Oct 26 15:20:02 [INFO ] [Oct 26 02:40:57]: Daemon [bfdd] transitioned to READY state
Oct 26 15:20:02 [INFO ] [Oct 26 02:41:17]:ISSU: Daemons Ready for ISSU
Oct 26 15:20:02 [INFO ] [Oct 26 02:41:17]: Daemon [apsd] transitioned to READY state
```

```
Oct 26 15:20:02 [INFO ] [Oct 26 02:41:22]:ISSU: Offline Incompatible FRUs
Oct 26 15:20:02 [INFO ] [Oct 26 02:41:27]:ISSU: Starting Upgrade for FRUs
Oct 26 15:20:02 [INFO ] [Oct 26 02:41:27]: [FPC 1] None -> Prepare
Oct 26 15:20:02 [INFO ] [Oct 26 02:41:48]: [FPC 1] Prepare -> Ready for Reboot
Oct 26 15:20:02 [INFO ] [Oct 26 02:41:52]: [FPC 1] Ready for Reboot -> Reboot
Oct 26 15:20:02 [INFO ] [Oct 26 02:42:01]: [FPC 1] Reboot -> Blob Resync
Oct 26 15:20:02 [INFO ] [Oct 26 02:42:28]: [FPC 1] Blob Resync -> Ready Software State Sync
Oct 26 15:20:02 [INFO ] [Oct 26 02:42:32]: [FPC 1] Ready Software State Sync -> Software State
Sync
Oct 26 15:20:02 [INFO ] [Oct 26 02:44:02]: [FPC 1] Software State Sync -> Ready Hardware State
Sync
Oct 26 15:20:02 [INFO ] [Oct 26 02:44:02]: [FPC 1] Ready Hardware State Sync -> Hardware State
Sync
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:07]: [FPC 1] Hardware State Sync -> Reconnected
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:07]:ISSU: FRU Upgrade Done
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:07]:ISSU: Preparing for Switchover
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:07]: Daemon [lfmd] transitioned to SWITCHOVER_READY state
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:07]: Daemon [rpd] transitioned to SWITCHOVER_READY state
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:07]: Daemon [l2cpd] transitioned to SWITCHOVER_READY state
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:07]: Daemon [smid] transitioned to SWITCHOVER_READY state
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:07]: Daemon [bfdd] transitioned to SWITCHOVER_READY state
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:27]: Daemon [apsd] transitioned to SWITCHOVER_READY state
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:29]: Checking In-Service-Upgrade status
Oct 26 15:20:02 [INFO ]   Item           Status                  Reason
Oct 26 15:20:02 [INFO ]   FPC 1          Online (ISSU)
Oct 26 15:20:02 [INFO ] Resolving mastership...
Oct 26 15:20:02 [INFO ] Complete. The other routing engine becomes the master.
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:30]:ISSU: RE switchover Done
Oct 26 15:20:02 [INFO ] [Oct 26 02:47:30]:ISSU: Upgrading Old Master RE
Oct 26 15:20:02 [INFO ] Verified junos-install-mx-x86-64-17.4-20171024.0 signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
Oct 26 15:20:02 [INFO ] Verified manifest signed by PackageDevelopmentEc_2017 method
ECDSA256+SHA256
  ...
    ...
```

Release Information

Command introduced in Junos OS Release 15.1F3.

Option introduced in Junos OS Release 18.2R1 for the MX Series routers.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

request system software in-service-upgrade

request vmhost software rollback

IN THIS SECTION

- Syntax | 15556
- Description | 15556
- Options | 15557
- Required Privilege Level | 15557
- Output Fields | 15557
- Sample Output | 15557
- Release Information | 15559

Syntax

```
request vmhost  software rollback
<re0 | re1>
<routing-engine>
```

Description

Roll back the Junos OS and the host OS software packages to the previous versions. You can revert to the previous versions of software packages that were loaded at the last successful `request vmhost software add` command.

Options

none—Roll back the software packages of the Routing Engine on the local Virtual Chassis member.

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, roll back the software packages in Routing Engine in slot 0 (`re0`) or software packages in the Routing Engine in slot 1 (`re1`).

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

routing-engine (Optional) Specify the Routing Engine on which the software packages needs to be rolled back to the previous set of software packages. The following options are available:

- `backup`—Backup Routing Engine.
- `both`—Both Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—Other Routing Engine.

Required Privilege Level

maintenance

Output Fields

Sample Output

request vmhost software rollback

```
request vmhost  software rollback
<reboot>
```

command-name

```
user@host> request vmhost software rollback
Current root details,         Device sda, Label: jrootp_P, Partition: sda3
Finding alternate root for rollback
Rollback to software on jrootb_P ...
sh /etc/install/mk-mtre-rollback.sh jrootb_P b
Mounting device in preparation for rollback...
Updating boot partition for rollback...
Rollback complete, please reboot the node for it to take effect.
Cmos Write successfull
Cmos Write successfull for Boot_retry
Cmos Write successfull for Boot_retry
```

command-name

```
user@host> show vmhost version
Current root details,         Device sda, Label: jrootp_P, Partition: sda3
Current boot disk: Primary
Current root set: p
UEFI     Version: NGRE_v00.53.00.01

Primary Disk, Upgrade Time: Wed Feb 24 17:51:53 UTC 2016
    Pending reboot.

Version: set p
VMHost Version: 2.951
VMHost Root: vmhost-x86_64-15.1I20160210_2212_builder
VMHost Core: vmhost-core-x86_64-15.1I20160210_2212_builder
kernel: 3.10.79-ovp-rt74-WR6.0.0.20_preempt-rt
Junos Disk: junos-install-x86-64-15.1F5.5

Version: set b
VMHost Version: 2.953
VMHost Root: vmhost-x86_64-15.1F520160222_1052_builder
VMHost Core: vmhost-core-x86_64-15.1F520160222_1052_builder
kernel: 3.10.79-ovp-rt74-WR6.0.0.20_preempt-rt
Junos Disk: junos-install-x86-64-15.1F5.6
```

command-name

```
user@host> request vmhost reboot
Reboot the vmhost ? [yes,no] (no) yes

warning: Rebooting re1
Initiating vmhost reboot...  ok
Initiating Junos shutdown...  shutdown: [pid 9733]
Shutdown NOW!
ok
Junos shutdown is in progress...

*** FINAL System shutdown message from root@nikon1 ***

System going down IMMEDIATELY
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

request vmhost software add | 15543

request vmhost software abort in-service-upgrade | 15542

request vmhost software validate

IN THIS SECTION

- Syntax | 15560
- Description | 15560
- Required Privilege Level | 15560

Sample Output | 15560
Release Information | 15560

Syntax

```
request vmhost software validate package-name
```

Description

Verify and validate the software package compatibility with the current configuration of the router.

Required Privilege Level

maintenance

Sample Output

request vmhost software validate

```
request vmhost software validate
<package-name>
<log_collector.sh>
```

Release Information

Statement introduced in Junos OS Release 18.4R1

NOTE: The command is supported on the routers with RE-MX-X6, RE-MX-X8, and RE-PTX-X8 Routing Engines with VM host support only.

RELATED DOCUMENTATION

request system software abort

request vmhost software in-service-upgrade | 15551

request vmhost storage

IN THIS SECTION

- Syntax | 15561
- Description | 15561
- Options | 15562
- Required Privilege Level | 15562
- Sample Output | 15562
- Release Information | 15563

Syntax

```
request vmhost storage
request vmhost storage external-ssd initialize slot [0 | 1] public-dir-name [public-disk0 |
public-disk1] force
request vmhost storage external-ssd [add | remove] slot [0 | 1]
```

Description

Initializes an SSD in the specified external SSD slot. It prompts you to confirm and then formats the external SSD so that you can use it for NFX350 device.

Adds or removes an external SSD from its slot. This command also checks the configuration for any VNF path that requires the external disk to be present.

NOTE: External SSDs are not supported on NFX150 and NFX250 devices.

Options

- initialize—Initializes an SSD in the specified external SSD slot 0 or slot 1.
- public-dir-name—Shows the same public-directory path for an SSD even if you move the SSD from one slot to another.
- add—Adds an external SSD to slot 0 or slot 1.
- remove—Removes an external SSD from slot 0 or slot 1.

Required Privilege Level

maintenance

Sample Output

request vmhost storage (NFX150)

```
user@host> request vmhost storage ?
Possible completions:
  <storage-name>       Storage  name
  self-test-long       Long Storage Test
  self-test-messages   Storage Self Test messages
  self-test-short      Long Storage Test
```

request vmhost storage (NFX250 NextGen)

```
user@host> request vmhost storage
Possible completions:
  <storage-name>       Storage  name
  self-test-long       Long Storage Test
  self-test-messages   Storage Self Test messages
  self-test-short      Long Storage Test
```

request vmhost storage (NFX350)

```
user@host> request vmhost storage external-ssd initialize slot 0 public-dir-name public-disk0
force
```

```
Destroy all files on this external SSD and initialize? [yes,no] (no) yes

External SSD in slot 0  initialized, public directory name public-disk0

user@host> request vmhost storage external-ssd add slot 0
External SSD in slot 0 successfully added, accessible at /var/public-disk0
user@host> request vmhost storage external-ssd remove slot 0
Remove SSD paths from device? [yes,no] (no) yes

External SSD in slot 0 successfully removed
user@host> request vmhost storage external-ssd initialize slot 1 public-dir-name public-disk1
force
Destroy all files on this external SSD and initialize? [yes,no] (no) yes

External SSD in slot 1  initialized, public directory name public-disk1
user@host> request vmhost storage external-ssd add slot 1
External SSD in slot 1 successfully added, accessible at /var/public-disk1
user@host> request vmhost storage external-ssd remove slot 1
Remove SSD paths from device? [yes,no] (no) yes

External SSD in slot 1 successfully removed
```

Release Information

Command introduced in Junos OS Release 18.1R1.

request vmhost zeroize

IN THIS SECTION

- Syntax | 15564
- Description | 15564
- Options | 15564
- Required Privilege Level | 15565
- Sample Output | 15565
- Release Information | 15566

Syntax

```
request vmhost zeroize
<re0 | re1>
<routing-engine>
```

Description

Use this command to remove all configuration information on the Routing Engines and reset all key values. If the device has dual Routing Engines, the command is broadcast to both Routing Engines on the device. As a result, both Routing Engines are zeroized. The command removes all data files, including customized configuration and log files from their directories. The command removes all user-created files from the system including all plain-text passwords, secrets, and private keys for SSH, local encryption, local authentication, IPsec, RADIUS, TACACS+, and SNMP.

This command reboots the device and sets it to the factory-default configuration. After the reboot, you cannot access the device through the management Ethernet interface. Log in through the console as the root user and start the Junos OS CLI by typing **cli** at the prompt.

Options

none—Remove all configuration information on all the Routing Engines and reset all key values.

re0 | re1 (Optional) Remove all configuration information on the Routing Engine in slot 0 (`re0`) or the Routing Engine in slot 1 (`re1`).

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

routing-engine (Optional) Remove all configuration information on the specified Routing Engine. The following options are available:

- `backup`—Backup Routing Engine.
- `both`—Both Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Master Routing Engine.
- `other`—Other Routing Engine.

Required Privilege Level

maintenance

Sample Output

request vmhost zeroize

```
user@host> request vmhost zeroize
VMHost Zeroization :

Erase all data, including configuration and log files?. In case of Dual RE system, both Routing
Engines will be zeroized [yes,no] (no) yes

re0:
--------------------------------------------------------------------------
warning: Vmhost will reboot and may not boot without configuration
warning: Proceeding with vmhost zeroize
Zeroise secondary internal disk ...
Proceeding with zeroize on secondary disk
Mounting device in preparation for zeroize...
Cleaning up target disk for zeroize ...
Zeroize done on target disk.
Zeroize of secondary disk completed
Zeroize primary internal disk ...
Proceeding with zeroize on primary disk
Mounting device in preparation for zeroize...
Cleaning up target disk for zeroize ...
Zeroize done on target disk.
Zeroize of primary disk completed
Zeroize done
mv: cannot stat '/tmp/zero-UytUWY/tgt_jlvmrootfs/etc/fstab': No such file or directory mv:
cannot stat '/tmp/zero-UytUWY/fstab': No such file or directory mv: cannot stat '/tmp/
zero-6gvrWj/tgt_jlvmrootfs/etc/fstab': No such file or directory mv: cannot stat '/tmp/
zero-6gvrWj/fstab': No such file or directory
warning: Proceeding with vmhost reboot

*** FINAL System shutdown message from root@user ***

System going down IMMEDIATELY
```

```
Initiating vmhost reboot...  ok
Initiating Junos shutdown...  shutdown: [pid 8565]
Shutdown NOW!
ok
Junos shutdown is in progress...
Shutdown NOW!

System shutdown time has arrived\x07\x07

re1:
--------------------------------------------------------------------------
warning: Vmhost will reboot and may not boot without configuration
warning: Proceeding with vmhost zeroize
Zeroise secondary internal disk ...
Proceeding with zeroize on secondary disk
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

request vmhost snapshot | 15539

request vmhost zeroize ztp-option secure

IN THIS SECTION

- Syntax | 15567
- Description | 15567
- Options | 15567

- Required Privilege Level | 15567
- Sample Output | 15568
- Release Information | 15568

Syntax

```
request vmhost zeroize ztp-option secure-(enable | disable)
```

Description

On secure platforms that run secure ZTP (SZTP), you can switch between using SZTP to using zero-touch provisioning (ZTP).

> **NOTE**: On MX304 devices without a backup Routing Engine, when you issue the `request vmhost zeroize ztp-option secure-(enable | disable)` command, you will see the following warning on the console: `Backup RE is not present. Zeroize backup RE when it is inserted.`

Options

enable The CLI checks to see if the device is a secure device. If the device is secure, the process ends. The next time the device boots, the device uses secure ZTP as the provisioning solution. If the device is not a secure device, you will receive an error message that says the device is not secure, and the process ends.

disable The CLI checks to see if the device is a secure device. If the device is a secure device, the next time the device boots, the device uses ZTP as the provisioning solution. If the device is not secure, the process ends.

Required Privilege Level

maintenance

Sample Output

request vmhost zeroize ztp-option-enable

```
user@device> request vmhost zeroize ztp-option secure-enable
```

request vmhost zeroize ztp-option secure-disable

```
user@device> request vmhost zeroize ztp-option secure-disable
```

Release Information

Command introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Switching between Secure Zero Touch Provisioning and Zero Touch Provisioning

request vpls-switchover

IN THIS SECTION

- Syntax | 15569
- Description | 15569
- Options | 15569
- Required Privilege Level | 15569
- Output Fields | 15569
- Sample Output | 15569
- Release Information | 15569

Syntax

```
request vpls-switchover (instance | logical-system | neighbor | vpls-id)
```

Description

Perform LDP VPLS pseudowire switchover. This provides a way to gracefully recover from pseudowire failures in such a way that Layer 2 circuit and VPLS services can be restarted when an LDP session with the remote PE router goes down.

Options

instance	Name of the VPLS instance.
logical-system	Name of the logical system, or 'all'.
neighbor	IP address of the neighbor.
vpls-id	Layer 2 circuit ID.

Required Privilege Level

view

Output Fields

When you enter this command, you are not provided feedback on the status of your request. You can use the `show vpls connections instance` command to confirm the switchover.

Sample Output

request vpls-switchover vpls-id

```
user@host> request vpls-switchover instance VPLS_16040 vpls-id 16040 neighbor 101.0.0.1
```

Release Information

Command introduced in Junos OS Release 9.2R1.

RELATED DOCUMENTATION

Redundant Pseudowires for Layer 2 Circuits and VPLS

request wlan access-point packet capture

IN THIS SECTION

- Syntax | 15570
- Description | 15570
- Options | 15570
- Required Privilege Level | 15571
- What's Next | 15571
- Sample Output | 15571
- Release Information | 15571

Syntax

```
request wlan access-point packet-capture start ap-name interface interface-name duration
duration filename capture-file sizefile-size-max promiscuous filter-mac mac-address
```

Description

Capture the packets for the interfaces on wireless LAN interface Mini-PIM and to check whether the packets are received from Wi-Fi interface and sent to JUNOS. In chassis cluster mode, packet capture command is supported only on active wireless LAN interface.

Options

- `start` *`access-point name`*—Access point name to start the packet capture.
  - `duration`—Specify the capture duration. Range: 10 through 3600 seconds
  - `filename`—Specify the captured filename.

- `filter-mac`—Specify the MAC address of the interface.
- `interface`—Specify the name of the interface.
- `promiscuous`—Enable the promiscuous mode.
- `size`—Specify the maximum file size. Range: 64 through 4096 kilobytes.

- `stop` *`access-point name`*—Stop the packet capture.

Required Privilege Level

view

WHAT'S NEXT

Sample Output

request wlan access-point packet capture

```
user@host> request wlan access-point packet-capture start filename wlan.pcap e09-22-ha-ap
interface Radio1VAP0 duration 10
Starting packet capture

Capture interface:   Radio1VAP0

File             :   wlan.pcap

Duration         :   10 seconds

File size max    :   1024 kilobytes
```

Release Information

Command introduced in Junos OS release 20.3R1.

RELATED DOCUMENTATION

> *Wi-Fi Mini-Physical Interface Module Overview*

Configure Wi-Fi Mini-PIM

wlan | 31137

requested-ip-network-match (DHCP Local Server)

IN THIS SECTION

- Syntax | 15572
- Hierarchy Level | 15572
- Description | 15572
- Options | 15573
- Required Privilege Level | 15573
- Release Information | 15573

Syntax

```
requested-ip-network-match subnet-mask;
```

Hierarchy Level

```
[edit system services dhcp-local-server]
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Configure the subnet to which the DHCP local server matches the requested IP address (IPv6 address for DHCPv6 local server). The server accepts and uses the active client's requested address for address assignment only when the requested address and the IP address of the DHCP server interface (or IPv6

address of the DHCPv6 local server) are in the same subnet. The server accepts and uses the passive client's requested address for address assignment only when the requested address and the address of the relay interface are in the same subnet.

Options

subnet-mask

Length of the subnet mask.

- Range:
  - DHCP: 0 through 31
  - DHCPv6: 0 through 127
- Default:
  - DHCP: 8
  - DHCPv6: 16

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Specifying the Subnet for DHCP Client Address Assignment

required-depth

IN THIS SECTION

- Syntax | 15574
- Hierarchy Level | 15574
- Description | 15574
- Options | 15575
- Required Privilege Level | 15575
- Release Information | 15575

Syntax

```
required-depth number;
```

Hierarchy Level

```
[edit interfaces interface-name atm-options mpls pop-all-labels],
[edit interfaces interface-name fastether-options mpls pop-all-labels],
[edit interfaces interface-name gigether-options mpls pop-all-labels],
[edit interfaces interface-name sonet-options mpls pop-all-labels]
```

Description

For passive monitoring on ATM, SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, specify the number of MPLS labels an incoming packet must have for the `pop-all-labels` statement to take effect. For passive monitoring on the MX304 router, MX Series routers with MPCs, and MX series routers with the LC9600 line card, all labels are popped by default and the `required-depth` statement is ignored.

If you include the `required-depth 1` statement, the `pop-all-labels` statement takes effect for incoming packets with one label only. If you include the `required-depth 2` statement, the `pop-all-labels` statement takes effect for incoming packets with two labels only.

Options

number—Number of MPLS labels on incoming IP packets.

- **Range:** 1 through 2 labels.
- **Default:** If you omit this statement, the `pop-all-labels` statement takes effect for incoming packets with one or two labels. The default is equivalent to including the `required-depth [ 1 2 ]` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Passive Flow Monitoring for MPLS Encapsulated Packets

Interfaces Fundamentals for Junos OS

reroute-address

IN THIS SECTION

- Syntax | 15576
- Hierarchy Level | 15576
- Description | 15576
- Default | 15576
- Options | 15576
- Required Privilege Level | 15577
- Release Information | 15577

Syntax

```
reroute-address {
    inet address | inet6 address);
}
```

Hierarchy Level

```
[edit forwarding-options evpn-vxlan]
```

Description

Enable fast reroute egress link protection (ELP) on provider edge (PE) devices for access links to multihomed customer edge (CE) devices in an EVPN-VXLAN fabric. This feature helps to minimize load-balanced traffic loss when the link from a PE device to a multihomed CE device goes down.

Remote sources load balance traffic toward a multihomed CE device by distributing portions of the traffic among the multihoming peer PE devices. When you enable this feature on a multihoming PE device, the device creates a backup VXLAN tunnel called an egress link protection (ELP) tunnel to each of its multihoming peer PE devices. If the PE device link to the CE device goes down, the PE device reroutes its portion of the load-balanced traffic to its peer PE devices on the ELP tunnels. The peer PE devices then send the excess load-balanced traffic to the CE device.

The PE device reroutes traffic on the ELP tunnels only until the source device routing table converges. Then the source device stops sending any traffic to the PE device with the bad link, and load balances the traffic only to the other peer PE devices.

To enable this feature, configure this statement and specify an IP address for the ELP tunnel.

Default

Fast reroute egress link protection is not enabled.

Options

inet *address*	Specify the IPv4 address for an ELP reroute tunnel.
inet6 *address*	Specify the IPv6 address for an ELP reroute tunnel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

Fast Reroute for Egress Link Protection with EVPN-VXLAN Multihoming

reserve ingress-egress-mcast

IN THIS SECTION

- Syntax | 15577
- Hierarchy Level | 15577
- Description | 15578
- Options | 15578
- Required Privilege Level | 15578
- Release Information | 15578

Syntax

```
reserve ingress-egress-mcast <percentage>;
```

Hierarchy Level

```
[edit system packet-forwarding-options internal-forwarding-bandwidth]
```

Description

You can configure the bandwidth of the recycle interface to enable a percentage of the recycle interface bandwidth for ingress-egress multicast replication.

The following is an example of multicast replication bandwidth calculation:

internal-forwarding-bandwidth = 100Gbps

Reserved ingress-egress-multicast = 40% means

For multicast replication, you can obtain 100Gbps x 40/100 = 40 Gbps per PFE and receiver.

In the above example, 40Gbps is the reserved bandwidth on the recycle port. When you send 100Gbps multicast traffic, then that multicast traffic is forwarded without any drop till 40Gbps, because the remaining 60Gpbs bandwidth was planned for other than multicast applications.

Options

percentage

Recycle interface bandwidth percentage.

- **Range:** 1 through 100 %
- **Default:** 100%

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS EVO Release 21.4R1.

RELATED DOCUMENTATION

Ingress and Egress Multicast Replication using Recycle Replication Model

reserved-bandwidth

IN THIS SECTION

- Syntax | 15579
- Hierarchy Level | 15579
- Description | 15579
- Options | 15580
- Required Privilege Level | 15580
- Release Information | 15580

Syntax

```
reserved-bandwidth {
      express percentage;
      gbr percentage;
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw system class-of-service]
```

Description

Reserve a fraction of the system's bandwidth for the express and guaranteed bit rate forwarding queues on the virtual routing and forwarding loopback interface on the anchor Packet Forwarding Engine. The system reserves the same fraction of bandwidth across all anchor Packet Forwarding Engines.

This configuration does not affect any bandwidth reservations on any of the WAN ports on any of the Packet Forwarding Engines (including the anchor Packet Forwarding Engines) in the system.

Options

express *percentage*

Specify the percentage of bandwidth to be reserved for the express queue.

- Range: 1 through 100
- Default: 2

gbr *percentage*

Specify the percentage of bandwidth to be reserved for each guaranteed bit rate queue.

- Range: 1 through 100
- Default: 5

Required Privilege Level

system

Release Information

Juniper Networks introduced this statement in Junos OS Release 22.2R1.

reserved (Security IDP Custom Attack)

IN THIS SECTION

- Syntax | 15581
- Hierarchy Level | 15581
- Description | 15581
- Options | 15581
- Required Privilege Level | 15581
- Release Information | 15581

Syntax

```
reserved {
    match (equal | greater-than | less-than | not-equal);
    value reserved-value;
}
```

Hierarchy Level

```
[edit security idp custom-attack ipv4_custom attack-type signature protocol tcp]
```

Description

Specify the three reserved bits in the TCP header field.

Options

match (equal | greater-than | less-than | not-equal)	Match an operand.
value	Match a decimal value. • **Range:** 0 through 7

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

reset (Security IDP)

IN THIS SECTION

- Syntax | 15582
- Hierarchy Level | 15582
- Description | 15582
- Required Privilege Level | 15582
- Release Information | 15582

Syntax

```
reset;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain]
```

Description

Select `reset` if the compound attack should be matched more than once within a single session or transaction.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

reset-tracking-bit

IN THIS SECTION

- Syntax | 15583
- Hierarchy Level | 15583
- Description | 15583
- Required Privilege Level | 15584
- Release Information | 15584

Syntax

```
reset-tracking-bit;
```

Hierarchy Level

```
[edit protocols pim],
[edit protocols pim interface (Protocols PIM)  interface-name],
[edit routing-instances routing-instance-name protocols pim],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name],
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name]
```

Description

Change the value of a tracking bit (T-bit) field in the LAN prune delay hello option from the default of 1 to 0, which enables join suppression for a multicast interface. When the network starts receiving multiple identical join messages, join suppression triggers a random timer with a value of 66 through 84 milliseconds (1.1 × periodic through 1.4 × periodic, where periodic is 60 seconds). This creates an interval during which no identical join messages are sent. Eventually, only one of the identical messages is sent. Join suppression is triggered each time identical messages are sent for the same join.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Example: Enabling Join Suppression

override-interval | **12058**

propagation-delay | **13381**

reset-count

IN THIS SECTION

- Syntax | **15584**
- Hierarchy Level | **15585**
- Description | **15585**
- Options | **15585**
- Release Information | **15585**

Syntax

```
reset-count reset-count;
```

Hierarchy Level

```
[edit chassis fabric-card]
```

Description

Number of times sib can be reset by cmerror

Options

reset-count

- Range: 1 to 10

Release Information

Statement introduced in Junos OS Release 24.3R1.

resignal-minimum-bandwidth

IN THIS SECTION

- Syntax | 15585
- Hierarchy Level | 15586
- Description | 15586
- Required Privilege Level | 15586
- Release Information | 15586

Syntax

```
resignal-minimum-bandwidth;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Description

Resignal the LSP using the configured minimum bandwidth when an LSP comes back up after going down.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Automatic Bandwidth Allocation for LSPs

resiliency

IN THIS SECTION

- Syntax | 15587
- Hierarchy Level | 15588
- Description | 15588
- Default | 15588
- Options | 15588

- Additional information | 15588
- Required Privilege Level | 15588
- Release Information | 15589

Syntax

```
resiliency {
   exceptions {
      forwarding;
      os;
      routing;
   }
   store {
      database;
      file name {
         files number;
         (no-world-readable | world-readable);
         size size;
      }
      fwding-file name{
         files number;
         (no-world-readable | world-readable);
         size size;
      }
   }
   traceoptions {
      file name {
         files number;
         match;
         (no-world-readable | world-readable);
         size size;
      }
      flag flag;
      no-remote-trace;
   }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the Juniper Resiliency Interface (JRI) on-box collector to detect exceptions from multiple modules and collect those exceptions in a file or a database. If you want to store exceptions data in a file, you need to configure a filename for the file and the size for the file.

Default

The on-box collector is disabled.

Options

The statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Additional information

To fully configure JRI for forwarding exceptions, you must also:

- Configure inline-monitoring services as you normally would, and configure the Juniper-specific IEs with the `primary-data-record-fields` statement at the `[edit services inline-monitoring templates` *`template-name`*`]` hierarchy level.
- Configure the Observation Cloud identifier with the `observation-cloud-id` statement at the `[edit services inline-monitoring]` hierarchy level.
- Configure exception reporting for a particular inline-monitoring instance with the `exception-reporting` statement at the `[edit chassis fpc` *`name`* `pfe` *`name`*`]` hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 for MX Series routers.

Statement introduced in Junos Evolved OS Release 22.2R1 for PTX Series routers.

RELATED DOCUMENTATION

Juniper Resiliency Interface

inline-monitoring | 8672

primary-data-record-fields | 13134

exception-reporting | 6572

resilient-hash

IN THIS SECTION

- Syntax | 15589
- Hierarchy Level | 15589
- Description | 15590
- Required Privilege Level | 15590
- Release Information | 15590

Syntax

```
resilient-hash;
```

Hierarchy Level

```
[edit interfaces aex  aggregated-ether-options]]
```

Description

Enable resilient hashing for a LAG to minimize remapping of destination paths.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Configuring Resilient Hashing for LAGs/ECMP Groups

resize (Ephemeral Database)

IN THIS SECTION

- Syntax | 15590
- Hierarchy Level | 15591
- Description | 15591
- Default | 15591
- Options | 15591
- Required Privilege Level | 15592
- Release Information | 15592

Syntax

```
resize {
    database-size-diff size MB;
```

```
    database-size-on-disk size MB;
}
```

Hierarchy Level

```
[edit system configuration-database ephemeral instance instance-name]
```

Description

Resize the ephemeral instance's configuration database. The system resizes the database during the commit operation when it meets the criteria in both the `database-size-diff` and `database-size-on-disk` statements. The system displays `Database resize completed` after it successfully resizes the configuration database.

A resize operation involves copying the active configuration objects to a new database and setting the version for each object to the latest version. Resizing the database effectively removes all deleted configuration objects and defragments the database.

Default

If you do not include this statement, the device does not resize the ephemeral database.

Options

none	Enable database resizing and use the default values for the `database-size-diff` and `database-size-on-disk` statements.
`database-size-diff` ***size***	Specify the minimum difference between the database size and the actual usage. • **Default:** 100 MB • **Range:** 1 MB through 1024 MB
`database-size-on-disk` ***size***	Specify the minimum configuration database size on disk. • **Default:** 450 MB • **Range:** 1 MB through 1024 MB

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Managing Ephemeral Configuration Database Space

resize-ephemeral-default

IN THIS SECTION

- Syntax | 15592
- Hierarchy Level | 15593
- Description | 15593
- Default | 15593
- Options | 15593
- Required Privilege Level | 15593
- Release Information | 15594

Syntax

```
resize-ephemeral-default {
    database-size-diff size MB;
    database-size-on-disk size MB;
}
```

Hierarchy Level

```
[edit system configuration-database ephemeral]
```

Description

Resize the default ephemeral configuration database. The system resizes the database during the commit operation when it meets the criteria in both the `database-size-diff` and `database-size-on-disk` statements. The system displays `Database resize completed` after it successfully resizes the configuration database.

A resize operation involves copying the active configuration objects to a new database and setting the version for each object to the latest version. Resizing the database effectively removes all deleted configuration objects and defragments the database.

Default

If you do not include this statement, the device does not resize the ephemeral database.

Options

none	Enable database resizing and use the default values for the `database-size-diff` and `database-size-on-disk` statements.
`database-size-diff` ***size***	Specify the minimum difference between the database size and the actual usage. • **Default:** 100 MB • **Range:** 1 MB through 1024 MB
`database-size-on-disk` ***size***	Specify the minimum configuration database size on disk. • **Default:** 450 MB • **Range:** 1 MB through 1024 MB

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Managing Ephemeral Configuration Database Space

resolution

IN THIS SECTION

- Syntax | 15594
- Hierarchy Level | 15595
- Description | 15595
- Options | 15596
- Required Privilege Level | 15596
- Release Information | 15596

Syntax

```
resolution {
    (preserve-nexthop-hierarchy | no-preserve-nexthop-hierarchy);
    rib routing-table-name {
        import [ policy-names ];
        inet-import [ policy-names ];
        inet-resolution-ribs [ routing-table-names ];
        inet6-import [ policy-names ];
        inet6-resolution-ribs [ routing-table-names ];
        iso-import [ policy-names ];
        iso-resolution-ribs [ routing-table-names ];
        mpls-import [ policy-names ];
        resolution-family resolution-family;
        resolution-ribs [ routing-table-names ];
```

```
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Configure the router to perform custom route resolution on protocol next hops of routes in a certain routing table. The protocol next hop is used to determine the forwarding next hop.

For example, you might want to direct **inet.2** route resolution to use topology routing tables **:red.inet.0** and **:blue.inet.0** for protocol next-hop IP address lookups. Or you might want to direct bgp.l3vpn.0 to use the information in inet.0 to resolve routes, thus overriding the default behavior, which is to use inet.3.

You can specify up to two routing tables in the "resolution-ribs" on page 15599 statement. The route resolution scheme first checks the first-listed routing table for the protocol next-hop address. If the address is found, it uses this entry. If it is not found, the resolution scheme checks the second-listed routing table. Hence, only one routing table is used for each protocol next-hop address. For example, if you configure `resolution rib bgp.l3vpn.0 resolution-ribs [inet.0 inet.3]`, inet.0 is checked first and then inet.3 is checked.

> **NOTE**: Customizing route resolution might cause the routing protocol process (rpd) to consume more memory resources than it ordinarily would. When you customize route resolution, we recommend that you check the memory resources by running the `show system processes` and the `show task memory` commands. For more information, see *Routing Protocol Process Overview for EX Series Switches*.

The remaining statements are explained separately. See CLI Explorer.

Options

disable-auto-resolve-on-scheme-changes	(Optional) Service routes are not re-resolved automatically when resolution schemes are created, deleted, or changed. However, you can manually trigger route re-resolution on resolution scheme changes.
inet-import [*policy-names*]	(Optional) Import policy for IPv4 family resolution tree.
inet-resolution-ribs [*routing-table-names*	(Optional) Specify routing tables to use for IPv4 family protocol-next-hop resolution.
inet6-import [*policy-names*]	(Optional) Import policy for IPv6 family resolution tree.
inet6-resolution-ribs [*routing-table-names*]	(Optional) Specify routing tables to use for IPv6 family protocol-next-hop resolution. Enabling the `inet6-resolution-ribs` option causes the static LSP route resolution to happen over the more preferred resolving route (lowest protocol preference) among the longest-matching-prefix routes in both the inet6.0 and inet6.3 routing tables.
iso-import [*policy-names*]	(Optional) Import policy for ISO family resolution tree.
iso-resolution-ribs [*routing-table-names*]	(Optional) Specify routing tables to use for ISO family protocol-next-hop resolution.
resolution-family *resolution-family*	(Optional) Specify a family of resolution tree.
preserve-nexthop-hierarchy	Preserve nexthop hierarchy when installing nexthop to forwarding plane.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

inet6-import and inet6-resolution-ribs options added in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Example: Configuring Route Resolution on PE Routers

Example: Configuring Route Resolution on Route Reflectors

Understanding Multitopology Routing in Conjunction with PIM

Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths

resolution-map

IN THIS SECTION

- Syntax | 15597
- Hierarchy Level | 15597
- Description | 15598
- Options | 15598
- Required Privilege Level | 15598
- Release Information | 15598

Syntax

```
resolution-map name {
    mode (color-only | ip-color);
}
```

Hierarchy Level

```
[edit logical-systems name policy-options],
[edit policy-options]
```

Description

Define a set of protocol next hop resolution modes.

A resolution-map can be referred by a new resolution-map action of a policy statement, which is in turn applied to a VPN service through the routing-instance `import-vrf`. A given resolution-map may be shared by multiple policy-statements.

Options

name Resolution Map name.

mode List of resolution modes in order that defines fallback mechanism.

- Values:
  - color-only—Color-only protocol next hop resolution mode.
  - ip-color—Colored-IP protocol next hop resolution mode.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.2R1-S1.

RELATED DOCUMENTATION

Color-Based Mapping of VPN Services Overview

Static Segment Routing Label Switched Path

resolution-ribs

IN THIS SECTION

- Syntax | 15599
- Hierarchy Level | 15599
- Description | 15599
- Options | 15600
- Required Privilege Level | 15600
- Release Information | 15600

Syntax

```
resolution-ribs [ routing-table-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options resolution rib],
[edit logical-systems logical-system-name routing-options resolution rib],
[edit routing-instances routing-instance-name routing-options resolution rib],
[edit routing-options resolution rib]
```

Description

Specify one or more routing tables to use for route resolution.

This statement enables you to override the default routing tables that Junos OS uses for route resolution. For example, suppose that the resolution routing table is **inet.3**, but you want to allow fallback resolution through **inet.0**. One example use case is overriding the **bgp.rtarget.0** (**family route-target**) routing table resolution from using only **inet.3** to using both **inet.3** and **inet.0**.

Options

routing-table-names—Name of one or more routing tables.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring Route Resolution on PE Routers

Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths

Understanding Multitopology Routing in Conjunction with PIM

resolve

IN THIS SECTION

- Syntax | 15601
- Hierarchy Level | 15601
- Description | 15601
- Default | 15601
- Required Privilege Level | 15601
- Release Information | 15602

Syntax

```
resolve;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name static (defaults | route)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options static (defaults | route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
(defaults | route)],
[edit logical-systems logical-system-name routing-options static (defaults | route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
(defaults | route)],
[edit routing-instances routing-instance-name routing-options static (defaults | route)],
[edit routing-options rib routing-table-name static (defaults | route)],
[edit routing-options static (defaults | route)]
```

Description

Statically configure routes to be resolved to a next hop that is not directly connected. The route is resolved through the `inet.0` and `inet.3` routing tables.

NOTE: You cannot configure both `resolve` and `retain` options for a statically configured route because resolved next hops cannot be retained.

Default

Static routes can point only to a directly connected next hop.

TIP: We recommend configuring the `no-resolve` option for individual routes to override default configuration.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

static (Routing Options) | 28570

resolve-vpn

IN THIS SECTION

- Syntax | 15602
- Hierarchy Level | 15602
- Description | 15603
- Required Privilege Level | 15603
- Release Information | 15603

Syntax

```
resolve-vpn;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family inet labeled-unicast],
[edit logical-systems logical-system-name protocols bgp group group-name family inet labeled-
unicast],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
family inet labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
```

```
family inet labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family inet labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family inet labeled-unicast],
[edit protocols bgp family inet labeled-unicast],
[edit protocols bgp group group-name family inet labeled-unicast],
[edit protocols bgp group group-name neighbor address family inet labeled-unicast],
[edit routing-instances routing-instance-name protocols bgp family inet labeled-unicast],
[edit routing-instances routing-instance-name protocols bgp group group-name family inet labeled-
unicast],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family inet labeled-unicast]
```

Description

Allow labeled routes to be placed in the inet.3 routing table for route resolution. These routes are then resolved for PE router connections where the remote PE is located across another AS. For a PE router to install a route in the VRF, the next hop must resolve to a route stored within the inet.3 table.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding Multiprotocol BGP

resource (Analytics)

IN THIS SECTION

- Syntax | 15604
- Hierarchy Level | 15604
- Description | 15605
- Required Privilege Level | 15605
- Release Information | 15605

Syntax

```
resource {
    interfaces {
        interface-name {
            resource-profile profile-name;
        }
    }
    system {
        polling-interval {
            queue-monitoring interval;
            traffic-monitoring interval;
        }
        resource-profile profile-name;
    }
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Configure network analytics resources such as resource profiles (for interfaces and system), and polling intervals (for queue and traffic monitoring).

> **NOTE**: The `resource` statement is not available in Junos OS Releases prior to 13.2X51-D15.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

resource-category jtree (Resource Monitor)

IN THIS SECTION

- Syntax | 15606
- Hierarchy Level | 15606
- Description | 15606
- Options | 15606
- Required Privilege Level | 15606
- Release Information | 15607

Syntax

```
resource-category jtree {
resource-type (contiguous-pages | free-dwords | free-pages) {
    low-watermark number;
    high-watermark number;
}
}
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Define the resource category that you want to monitor and analyze for ensuring system stability, especially the health and operating efficiency of I-chip-based line cards and Trio-based FPCs on MX Series routers. The resource category includes detailed CPU utilization, session rate, and session count statistics. You use the resource category statistics to understand the extent to which new attack objects or applications affect performance. The jtree memory on all MX Series router Packet Forwarding Engines has two segments: one segment primarily stores routing tables and related information, and the other segment primarily stores firewall-filter-related information. The Junos OS provides the memory-enhanced statement to reallocate the jtree memory for routes, firewall filters, and Layer 3 VPNs.

Options

jtree Specify the Jtree resource category for which you want to monitor the health (working condition), operating efficiency, traffic-handling capacity, and performance status.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

> *Diagnosing and Debugging System Performance by Configuring Memory Resource Usage Monitoring on MX Series Routers*

resource-cleanup

IN THIS SECTION

- Syntax | 15607
- Hierarchy Level | 15608
- Description | 15608
- Options | 15608
- Required Privilege Level | 15608
- Release Information | 15608

Syntax

```
resource-cleanup {
    disable;
    traceoptions {
        file filename files number match regex size size (world-readable | no-world-readable);
        flag flag;
        level level;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Selectively turn on or off the debugging of trace messages for the resource cleanup process.

Options

`disable`—Disable the resource cleanup process.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

traceoptions (Resource Cleanup) | 29755

resource (Interfaces)

IN THIS SECTION

- Syntax | 15609
- Hierarchy Level | 15609
- Description | 15609

Required Privilege Level | 15609
Release Information | 15609

Syntax

```
resource number;
```

Hierarchy Level

```
[edit interfaces interface-name radio-router]
```

Description

This option controls the resource weight (value 1–100).

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

resource-list (Network Management)

IN THIS SECTION

- Syntax | 15610
- Hierarchy Level | 15610
- Description | 15610
- Options | 15611
- Required Privilege Level | 15611
- Release Information | 15612

Syntax

```
        resource-list resource-list-name resource-names <all-resources> / <ecmp-group> <ecmp-
member> <efp> <egress-l3-interface> <host-ipv4> <host-ipv6> <ifp> <l3-next-hop> <lpm-ipv4> <lpm-
ipv6-128> <lpm-ipv6-64> <mac> <mpls-ingress> <mpls-swap> <tunnel> <vfp> <vport>;
        }
```

Hierarchy Level

```
[edit system packet-forwarding-options hw-resource-monitor]
```

Description

Create a resource list to monitor hardware resource utilization. You can configure multiple resource lists, but the same resource can only be used once and cannot be duplicated on multiple resource lists.

Configure a resource list with multiple resource names, each resource name separated by a space. To add all available resources to the resource list, specify `all-resources` option.

You can map a resource list with a monitor profile (optional) and configure the upper, lower thresholds, and notification type of the monitor profile. You can also configure the polling interval (optional).

No default values.

Options

resource-list-name The name of the resource list.

resource-names The list of resources to monitor.

- `all-resources`—All applicable hardware resources.
- `ecmp-group`—Equal Cost Multi Path (ECMP) group hardware table.
- `ecmp-member`—ECMP member hardware table.
- `efp`—TCAM: Egress Field Processor hardware table.
- `egress-l3-interface`—Egress layer 3 Interface hardware table.
- `host-ipv4`—IPv4 host routes hardware table.
- `host-ipv6`—IPv6 host routes hardware table.
- `ifp`—TCAM: Ingress field processor hardware table.
- `l3-next-hop`—Layer 3 next hop hardware table.
- `lpm-ipv4`—Longest prefix match IPv4 32 bits wide hardware table.
- `lpm-ipv6-128`—Longest prefix match IPv6 128 bits wide hardware table.
- `lpm-ipv6-64`—Longest prefix match IPv6 64 bits wide hardware table.
- `mac`—Layer 2 entry hardware table.
- `mpls-ingress`—Multiprotocol Label Switching (MPLS) ingress hardware table.
- `mpls-swap`—MPLS swap hardware table.
- `tunnel`—IP/Layer 3 tunnel hardware table.
- `vfp`—TCAM: VLAN field processor hardware table.
- `vport`—Virtual port hardware table.

Required Privilege Level

edit

Release Information

Statement introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

No Link Title

polling-interval (Network Management) | 12752

show system packet-forwarding-options hw-resource-utilization-info | 26952

monitor-profile | 10912

resource-list (Telemetry)

IN THIS SECTION

- Syntax | 15612
- Hierarchy Level | 15613
- Description | 15613
- Options | 15613
- Required Privilege Level | 15614
- Release Information | 15614

Syntax

```
resource-list {ecmp-group-table | ecmp-member-table | efp | egress-l3-interface | host-ipv4 |
host-ipv6 | ifp | l3-next-hop | lpm-table-ipv4 | lpm-table-ipv6-128 | lpm-table-ipv6-64 | mpls-
ingress-table | mpls-table-swap | tunnel-table | vfp | vport}
```

Hierarchy Level

```
[edit forwarding-options pfe-sensor npu-memory],
[edit logical-systems name routing-instances name forwarding-options pfe-sensor npu-memory],
[edit logical-systems name tenants name routing-instances name forwarding-options pfe-sensor npu-
memory],
[edit routing-instances name forwarding-options pfe-sensor npu-memory],
[edit tenants name routing-instances name forwarding-options pfe-sensor npu-memory]
```

Description

Specify a list of resources to monitor and export for Packet Forwarding Engine sensors. If you do not include the `resource-list` statement, information for all active groups in the system is exported. To see a list of groups, use the `show pfe filter hw summary` command.

Export group information to a collector by including the resource path **/junos/system/linecard/npu/memory** in a subscription.

Options

ecmp-group-table	Retrieve equal-cost multipath (ECMP) group table information.
ecmp-member-table	Retrieve ECMP member table information.
efp	Retrieve egress filter information.
egress-l3-interface	Retrieve egress Layer 3 (L3 interface information.
host-ipv4	Retrieve IPv4 host information.
host-ipv6	Retrieve IPv6 host information.
ifp	Retrieve ingress filter information.
l3-next-hop	Retrieve L3 next hop information.
lpm-table-ipv4	Retrieve longest prefix match (LPM) IPv4 table information.
lpm-table-ipv6-128	Retrive LPM IPv6 table information for /128 prefix routes.

lpm-table-ipv6-64	Retrieve LPM IPv6 table information for /64 prefix routes.
mpls-ingress-table	Retrieve MPLS ingress table information
mpls-swap-table	Retrieve MPLS swap table information.
tunnel-table	Retrieve MPLS/LP tunnel table information.
vfp	Retrieve VLAN filter information.
vport	Retrieve Virtual Chassis port (VCP) information.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

show pfe filter hw summary | 22983

resource-manager

IN THIS SECTION

- Syntax | 15615
- Hierarchy Level | 15615
- Description | 15615
- Options | 15615
- Required Privilege Level | 15615
- Release Information | 15615

Syntax

```
resource-manager {
    traceoptions (Security) {
        flag flag;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure resource manager security options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced before Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

resource-manager (Forwarding-options)

IN THIS SECTION

- Syntax | 15616
- Syntax (SRX4600) | 15616
- Hierarchy Level | 15617
- Description | 15617
- Options | 15617
- Required Privilege Level | 15617
- Release Information | 15617

Syntax

```
resource-manager {
 cpu {
  re re;
}
```

Syntax (SRX4600)

```
resource-manager {
 cpu {
  re re;
 }
 enhanced-logging;
}
```

Hierarchy Level

```
[edit security forwarding-options]
```

Description

Manage the resources for special usage.

Options

cpu Configure CPU resources.

enhanced-logging Configure a dedicated CPU resource for on-box logging on SRX4600.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

`enhanced-login` option introduced to assign dedicated CPU resource for on-box logging on SRX4600 in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

show security forward-options resource-manager | 24360

resource-monitor

IN THIS SECTION

- Syntax | 15618
- Hierarchy Level | 15619
- Description | 15619
- Required Privilege Level | 15619
- Release Information | 15620

Syntax

```
resource-monitor {
    free-fw-memory-watermark number;
    free-heap-memory-watermark number;
    free-nh-memory-watermark number;
    high-cos-queue-threshold number;
    high-threshold number;
    no-logging;
    no-throttle;
    resource-category jtree {
        resource-type (contiguous-pages | free-dwords | free-pages) {
            low-watermark number;
            high-watermark number;
        }
    }
    subscribers-limit {
        client-type (any | dhcp | l2tp | pppoe) {
            chassis {
                limit limit;
            }
            fpc slot-number {
                limit limit;
                pic number {
                    limit limit;
                    port number {
                        limit limit;
```

```
                    }
                }
            }
        }
    }
    traceoptions {
        file filename <files number> <match regular-expression> <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Enable the resource monitoring capability to provision sufficient headroom (memory space limits that are set for the application or virtual router) for monitoring the health and operating efficiency of DPCs and MPCs. This feature also enables the memory resource monitoring mechanism to avoid the system operations from compromising on the health and traffic-handling stability of the line cards by generating error logs when a specified watermark value for memory regions and threshold value for the jtree memory region are exceeded. A trade-off on the system performance can be detrimental for supporting live traffic and protocols.

The variable *number* in the Syntax section represents a percentage.

You can configure the resource-monitoring capability on MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

`high-cos-queue-threshold` option introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

- Diagnosing and Debugging System Performance by Configuring Memory Resource Usage Monitoring on MX Series Routers
- *Resource Monitoring for Subscriber Management and Services*
- Resource Monitoring Usage Computation Overview
- *Limiting Subscribers by Client Type and Hardware Element with Resource Monitor*

resource-profiles (Analytics)

IN THIS SECTION

- Syntax | 15620
- Hierarchy Level | 15621
- Description | 15621
- Options | 15621
- Required Privilege Level | 15622
- Release Information | 15622

Syntax

```
resource-profiles {
    profile-name {
        depth-threshold {
            high number;
            low number;
        }
        latency-threshold {
```

```
            high number;
            low number;
        }
        no-queue-monitoring;
        no-traffic-monitoring;
        queue-monitoring;
        traffic-monitoring;
    }
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Configure resource profiles that are used as templates for specifying network analytics parameters. You use resource profiles to enable and disable queue and traffic monitoring, and specify depth and latency thresholds as applicable. Once you have defined a resource profile, you can apply it specifically to individual interfaces, or globally to a system.

> **NOTE**: The `resource-profiles` statement is not available in Junos OS Releases prior to 13.2X51-D15.

The remaining statements are explained separately. See CLI Explorer.

Options

profile-name	Specify a name for the resource profile.
no-queue-monitoring	Disable queue monitoring.
no-traffic-monitoring	Disable traffic monitoring.
queue-monitoring	Enable queue monitoring.
traffic-monitoring	Enable traffic monitoring.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

resource-watch

IN THIS SECTION

- Syntax | 15622
- Hierarchy Level | 15623
- Description | 15623
- Options | 15623
- Required Privilege Level | 15624
- Release Information | 15625

Syntax

```
resource-watch {
    cpu-statistics;
        junos-statistics;
        orange-zone-enable;
        orange-zone-threshold;
        red-zone-cpu-threshold;
        red-zone-enable;
```

```
        red-zone-jkl-threshold;
        red-zone-snmp-enable;
    resource-watch-enable;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

In SRX chassis cluster mode, split brain caused by redundant failover is triggered due to Routing Engine busy related to CPU overhead causing SRX HA cluster resiliency. To analyze CPU load, Reswatchd is a process running on Linux, Routing Engine3, and Routing Engine2 to collect and monitor the resource usage data. You can generate CPU load statistics report to analyze the system load (resource assignment and Junos kernel usage). You can enable statistics function through CLI to log CPU usage every second and generate a statistics report. This feature is supported on chassis standalone mode to monitor Routing Engine CPU usage and Junos kernel usage and is disabled by default.

Resource availability is classified to three zones:

For RE2, the threshold values are:

- **Green: CPU idle > 20%**
- **Orange: CPU idle <= 15%**
- **Red: CPU idle <= 5% or Junos Kernel load >= 50%**

For RE3, the threshold values are:

- **Green: CPU idle > 10%**
- **Orange: CPU idle <= 10%**
- **Red: CPU idle <= 3% or Junos Kernel load >= 50%**

Options

cpu-statistics Display CPU load statistics report. When enable CPU statistics, reswatchd logs CPU idle values **/var/log/reswatch_CPU_statistics_[1-5]** in .csv format.

- **Range:** `0-604800` seconds.
- Values:
  - `0`—Disable

junos-statistics

Display Junos kernel statistics report. When enable Junos Kernel statistics, reswatchd logs Junos Kernel usage values at **/var/log/reswatch_JK_statistics_[1-5]** in .csv format.

- **Range:** `0-604800` seconds.
- Values:
  - `0`—Disable

orange-zone-enable

Enable orange zone.

orange-zone-threshold

Specify orange zone value (CPU idle %)

- **Default:** `20`
- **Range:** `1-99`

red-zone-cpu-threshold

Specify red zone value (CPU idle %).

- **Default:** `10`
- **Range:** `1-99`

red-zone-enable

Enable red zone.

red-zone-jkl-threshold

Specify red zone value (Junos load %)

- **Default:** `50`
- **Range:** `1-99`

red-zone-snmp-enable

Enable red zone snmp.

resource-watch-enable

Enable resource watch.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS release 20.4R1.

RELATED DOCUMENTATION

show chassis reswatch | 18701

resources

IN THIS SECTION

- Syntax | 15625
- Hierarchy Level | 15625
- Description | 15626
- Required Privilege Level | 15626
- Release Information | 15626

Syntax

```
resources {
    node-group node-group-name {
        network-domain;
        node-device node-device-name;
    }
}
```

Hierarchy Level

```
[edit fabric]
```

Description

(QFabric systems only) Define resources within a QFabric system and handle a variety of component assignments.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Node Groups for the QFabric System

Understanding Node Groups

respond-bad-spi (Services IKE Policy)

IN THIS SECTION

- Syntax | 15627
- Hierarchy Level | 15627
- Description | 15627
- Options | 15627
- Required Privilege Level | 15627
- Release Information | 15627

Syntax

```
respond-bad-spi max-responses;
```

Hierarchy Level

```
[edit services ipsec-vpn ike policy]
```

Description

Enable response to invalid IPsec Security Parameter Index (SPI) values. If the security associations (SAs) between two peers of an IPsec VPN become unsynchronized, the device resets the state of a peer so that the two peers are synchronized.

Options

max-responses Number of times to respond to invalid SPI values per gateway.

- **Range:** 1 through 30
- **Default:** 5

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring IKE Policies

responder (performance-monitoring)

IN THIS SECTION

- Syntax | 15628
- Hierarchy Level | 15628
- Description | 15628
- Required Privilege Level | 15629
- Release Information | 15629

Syntax

```
responder {
    delay {
        min-query-interval milliseconds;
    }
    loss {
        min-query-interval milliseconds;
    }
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring]
```

Description

Configure responder options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Pro-Active Loss and Delay Measurements

On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview

performance-monitoring (Protocols MPLS) | **12434**

response-cache-timeout (RADIUS Client)

IN THIS SECTION

- Syntax | **15629**
- Hierarchy Level | **15630**
- Description | **15630**
- Options | **15630**
- Required Privilege Level | **15630**
- Release Information | **15630**

Syntax

```
response-cache-timeout seconds;
```

Hierarchy Level

```
[edit access radius clients client-name accounting]
```

Description

Configure the timeout for the RADIUS response cache. This timeout indicates how long to store the RADIUS response messages (sent for request messages) in the MX Series router response cache.

Options

seconds Length of timeout interval.

- **Range:** 5 through 20 seconds
- **Default:** 15 seconds

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers

rest

IN THIS SECTION

- Syntax | 15631
- Hierarchy Level | 15632
- Description | 15632
- Options | 15632
- Required Privilege Level | 15632
- Release Information | 15632

Syntax

```
rest {
    control {
        allowed-sources [ value-list ];
        connection-limit limit;
    }
    enable-explorer;
    http {
        addresses [ address-list ];
        port port-number;
    }
    https {
        addresses [ address-list ];
        cipher-list [cipher-1 cipher-2 cipher-3 ... ];
        mutual-authentication {
            certificate-authority certificate-authority-profile-name;
        }
        port port-number;
        server-certificate local-certificate-identifier;
    }
    routing-instance routing-instance;
    traceoptions {
        flag flag;
```

```
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Execute Junos OS commands over HTTP or HTTPS using REST. Optionally, specify JSON output for operational and configuration commands.

Options

routing-instance *routing-instance*

Specify the routing instance through which the remote server is reachable. This can be the non-default management instance mgmt_junos or a user-defined non-default routing instance. If you do not specify a routing instance, REST API requests are sent over the default routing instance.

> **NOTE**: If you want to use the non-default management instance mgmt_junos, you must also define the mgmt_junos routing instance under the `[edit routing-instances]` hierarchy level and configure the `management-instance` statement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

`routing-instance` option introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Management Interface in a Non-Default Instance

restart

IN THIS SECTION

- Syntax | 15633
- Syntax (ACX Series Routers) | 15634
- Syntax (EX Series Switches) | 15634
- Syntax (MX Series Routers) | 15635
- Syntax (QFX Series) | 15635
- Syntax (Routing Matrix) | 15636
- Syntax (SRX Series) | 15636
- Syntax (TX Matrix Routers) | 15636
- Syntax (TX Matrix Plus Routers) | 15637
- Syntax (QFX Series) | 15637
- Description | 15637
- Options | 15638
- Required Privilege Level | 15647
- Output Fields | 15647
- Sample Output | 15647
- Release Information | 15647

Syntax

```
restart
<adaptive-services |ancpd-service | application-identification |audit-process | auto-
configuration |captive-portal-content-delivery |ce-l2tp-service | chassis-control  | class-of-
service |clksyncd-service |database-replication|datapath-trace-service |dhcp-service | diameter-
service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | ethernet-connectivity-
fault-management |ethernet-link-fault-management |event-processing | firewall | general-
```

```
authentication-service | gracefully | iccp-service |idp-policy | immediately |interface-control
| ipsec-key-management | kernel-health-monitoring | kernel-replication | l2-learning | l2cpd-
service | l2tp-service | l2tp-universal-edge | lacp | license-service |link-management |local-
policy-decision-function |mac-validation |mib-process | mountd-service |mpls-traceroute |mspd |
multicast-snooping |named-service | nfsd-service | packet-triggered-subscribers |peer-selection-
service |pgm | pic-services-logging | pki-service |ppp | ppp-service | pppoe | protected-system-
domain-service | redundancy-interface-process | remote-operations | root-system-domain-service |
routing <logical-system logical-system-name>  | sampling | sbc-configuration-process | sdk-
service |service-deployment | services | snmp |soft |static-subscribers |statistics-service|
subscriber-management | subscriber-management-helper | tunnel-oamd |usb-control| vrrp |web-
management>
<gracefully | immediately | soft>
```

Syntax (ACX Series Routers)

```
restart
<adaptive-services |audit-process | auto-configuration | autoinstallation |chassis-control |
class-of-service |clksyncd-service |database-replication| dhcp-service | diameter-service | disk-
monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management | ethernet-link-fault-
management |event-processing | firewall | general-authentication-service | gracefully |
immediately |interface-control | ipsec-key-management | l2-learning | lacp |link-management |mib-
process | mountd-service |mpls-traceroute |mspd | named-service | nfsd-service | pgm | pki-
service | ppp | pppoe | redundancy-interface-process | remote-operations | routing | sampling |
sdk-service |secure-neighbor-discovery | service-deployment | services | snmp |soft | statistics-
service| subscriber-management | subscriber-management-helper | tunnel-oamd | vrrp>
```

Syntax (EX Series Switches)

```
restart
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp | dhcp-
service | diameter-service | dot1x-protocol | ethernet-link-fault-management | ethernet-
switching | event-processing | firewall | general-authentication-service | interface-control |
kernel-health-monitoring | kernel-replication | l2-learning | lacp | license-service | link-
management | lldpd-service | mib-process | mountd-service | multicast-snooping | pgm |
redundancy-interface-process | remote-operations | routing | secure-neighbor-discovery | service-
deployment | sflow-service | snmp | vrrp | web-management>
```

Syntax (MX Series Routers)

```
restart
<adaptive-services | ancpd-service | application-identification | audit-process | auto-
configuration | bbe-stats-service | captive-portal-content-delivery | ce-l2tp-service | chassis-
control | class-of-service | clksyncd-service | database-replication | datapath-trace-service |
dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging |
ethernet-connectivity-fault-management | ethernet-link-fault-management | event-processing |
firewall | general-authentication-service | gracefully | iccp-service | idp-policy | immediately
|interface-control | ipsec-key-management |kernel-health-monitoring |  kernel-replication | l2-
learning | l2cpd-service | l2tp-service | l2tp-universal-edge | lacp | license-service | link-
management | local-policy-decision-function | mac-validation | mib-process | mountd-service |
mpls-traceroute | mspd | multicast-snooping |named-service | nfsd-service | packet-triggered-
subscribers |peer-selection-service | pgm | pic-services-logging | pki-service | ppp | ppp-
service | pppoe | protected-system-domain-service | redundancy-interface-process | remote-
operations | root-system-domain-service | routing | routing <logical-system logical-system-
name>  | sampling | sbc-configuration-process | sdk-service | service-deployment | services |
snmp |soft |static-subscribers |statistics-service| subscriber-management | subscriber-
management-helper | tunnel-oamd | usb-control | vrrp | web-management>
<all-members>
<gracefully | immediately | soft>
<local>
<member member-id>
```

Syntax (QFX Series)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
diameter-service | dlsw | ethernet-connectivity | event-processing | fibre-channel | firewall |
general-authentication-service | igmp-host-services | interface-control | ipsec-key-management |
isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process | named-service | network-
access-service | nstrace-process | pgm | ppp | pppoe | redundancy-interface-process | remote-
operations |logical-system-name>  | routing | sampling |secure-neighbor-discovery | service-
deployment | snmp | usb-control | web-management>
<gracefully | immediately | soft>
```

Syntax (Routing Matrix)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | disk-monitoring |
dynamic-flow-capture | ecc-error-logging | event-processing | firewall | interface-control |
ipsec-key-management | kernel-replication | l2-learning | l2tp-service | lacp | link-management
| mib-process | pgm | pic-services-logging | ppp | pppoe | redundancy-interface-process | remote-
operations | routing <logical-system logical-system-name>  | sampling | service-deployment |
snmp>
<all | all-lcc | lcc number>
<gracefully | immediately | soft>
```

Syntax (SRX Series)

```
restart
<application-identification |application-security |audit-process |commitd-service |chassis-
control | class-of-service |database-replication |datapath-trace-service |ddns |dhcp |dhcp-
service |dynamic-flow-capture |disk-monitoring |event-processing | ethernet-connectivity-fault-
management |ethernet-link-fault-management |extensible-subscriber-services |fipsd |firewall |
firewall-authentication-service |general-authentication-service |gracefully |gprs-process |idp-
policy |immediately |interface-control | ipmi |ipsec-key-management |jflow-service |jnu-
management |jnx-wmicd-service |jsrp-service |kernel-replication |l2-learning |l2cpd-service |
lacp |license-service |logical-system-service |mib-process |mountd-service |named-service |
network-security |network-security-trace |nfsd-service |ntpd-service |pgm |pic-services-logging |
profilerd |pki-service |remote-operations |rest-api |routing |sampling |sampling-route-record |
scc-chassisd |secure-neighbor-discovery |security-intelligence |security-log |services |service-
deployment |simple-mail-client-service |soft |snmp |static-routed |statistics-service |
subscriber-management |subscriber-management-helper |system-log-vital |tunnel-oamd |uac-service |
user-ad-authentication |vrrp |web-management >
```

Syntax (TX Matrix Routers)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service | diameter-
service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | event-processing |
firewall | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2tp-
service | lacp | link-management | mib-process |pgm | pic-services-logging | ppp | pppoe |
redundancy-interface-process | remote-operations | routing <logical-system logical-system-name>
| sampling | service-deployment | snmp| statistics-service>
```

```
<all-chassis | all-lcc | lcc number | scc>
<gracefully | immediately | soft>
```

Syntax (TX Matrix Plus Routers)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service | diameter-
service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | event-processing |
firewall | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2tp-
service | lacp | link-management | mib-process | pgm | pic-services-logging | ppp | pppoe |
redundancy-interface-process | remote-operations | routing <logical-system logical-system-name>
| sampling | service-deployment | snmp| statistics-service>
<all-chassis | all-lcc | all-sfc | lcc number | sfc number>
<gracefully | immediately | soft>
```

Syntax (QFX Series)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
diameter-service | dlsw | ethernet-connectivity | event-processing | fibre-channel | firewall |
general-authentication-service | igmp-host-services | interface-control | ipsec-key-management |
isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process | named-service | network-
access-service | nstrace-process | pgm | ppp | pppoe | redundancy-interface-process | remote-
operations |logical-system-name>  | routing | sampling |secure-neighbor-discovery | service-
deployment | snmp | usb-control | web-management>
<gracefully | immediately | soft>
```

Description

Restart a Junos OS process.

> **CAUTION**: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

The `restart` command expands all applications names including applications that are not required for the current platform. Therefore, a user could try to do a restart for an application that is not running for the

current platform. This error message communicates that the restart failed because the application was not running on the system.

Options

none	Same as `gracefully`.
adaptive-services	(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.
all-chassis	(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.
all-lcc	(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.
all-members	(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.
all-sfc	(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).
ancpd-service	(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.
application-identification	(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.
application-security	(Optional) Restart the application security process.
audit-process	(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.
auto-configuration	(Optional) Restart the Interface Auto-Configuration process.
autoinstallation	(EX Series switches only) (Optional) Restart the autoinstallation process.

bbe-stats-service	(MX Series routers only) (Optional) Restart bbe-statsd, the BBE statistics collection and management process.
captive-portal-content-delivery	(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.
ce-l2tp-service	(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.
chassis-control	(Optional) Restart the chassis management process.
class-of-service	(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch’s CoS configuration.
clksyncd-service	(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).
commitd-service	(Optional) Restart the committed services.
database-replication	(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.
datapath-trace-service	(Optional) Restart the packet path tracing process.
dhcp	(EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.
dhcp-service	(Optional) Restart the Dynamic Host Configuration Protocol process.
dialer-services	(EX Series switches only) (Optional) Restart the ISDN dial-out process.
diameter-service	(Optional) Restart the diameter process.
disk-monitoring	(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.
dlsw	(QFX Series only) (Optional) Restart the data link switching (DLSw) service.
dot1x-protocol	(EX Series switches only) (Optional) Restart the port-based network access control process.

dynamic-flow-capture	(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.
ecc-error-logging	(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.
ethernet-connectivity-fault-management	(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.
ethernet-link-fault-management	(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.
ethernet-switching	(EX Series switches only) (Optional) Restart the Ethernet switching process.
event-processing	(Optional) Restart the event process (eventd).
extensible-subscriber-services	(Optional) Restart the extensible subscriber services process.
fibre-channel	(QFX Series only) (Optional) Restart the Fibre Channel process.
fipsd	(Optional) Restart the fipsd services.
firewall	(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.
general-authentication-service	(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.
gprs-process	(Optional) Restart the General Packet Radio Service (GPRS) process.
gracefully	(Optional) Restart the software process.
iccp-service	(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.
idp-policy	(Optional) Restart the intrusion detection and prevention (IDP) protocol process.
immediately	(Optional) Immediately restart the software process.
interface-control	(Optional) Restart the interface process, which controls the router's or switch’s physical interface devices and logical interfaces.

ipmi	(Optional) Restart the intelligent platform management interface process.
ipsec-key-management	(Optional) Restart the IPsec key management process.
isdn-signaling	(QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.
jflow-service	(Optional) Restart jflow service process.
jnu-management	(Optional) Restart jnu management process.
jnx-wmicd-service	(Optional) Restart jnx wmicd service process.
jsrp-service	(Optional) Restart the Juniper Services Redundancy Protocol (jsrdp) process, which controls chassis clustering.
kernel-health-monitoring	(Optional) Restart the Routing Engine kernel health monitoring process, which enables health parameter data to be sent from kernel components to data collection applications. When you change the polling interval through `sysctl kern.jkhmd_polling_time_secs`, you must restart the kernel health monitoring process for the new polling interval to take effect.
kernel-replication	(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.
l2-learning	(Optional) Restart the Layer 2 address flooding and learning process.
l2cpd-service	(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.
l2tp-service	(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.
l2tp-universal-edge	(MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.
lacp	(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to

move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc *number*

(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

license-service

(EX Series switches only) (Optional) Restart the feature license management process.

link-management

(TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

lldpd-service

(EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

local

(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

local-policy-decision-function

(Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

logical-system-service

(Optional) Restart the logical system service process.

mac-validation

(Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member ***member-id***	(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace *member-id* with a value of `0` or `1`.
mib-process	(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.
mobile-ip	(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.
mountd-service	(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.
mpls-traceroute	(Optional) Restart the MPLS Periodic Traceroute process.
mspd	(Optional) Restart the Multiservice process.
multicast-snooping	(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.
named-service	(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.
network-access-service	(QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.
network-security	(Optional) Restart the network security process.
network-security-trace	(Optional) Restart the network security trace process.
nfsd-service	(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.
ntpd-service	(Optional) Restart the Network Time Protocol (NTP) process.
packet-triggered-subscribers	(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.
peer-selection-service	(Optional) Restart the Peer Selection Service process.

pgcp-service	(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the `services pgcp gateway` option.
pgm	(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.
pic-services-logging	(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.
pki-service	(Optional) Restart the PKI Service process.
ppp	(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.
ppp-service	(Optional) Restart the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.
pppoe	(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.
profilerd	(Optional) Restart the profiler process.
protected-system-domain-service	(Optional) Restart the Protected System Domain (PSD) process.
redundancy-interface-process	(Optional) Restart the ASP redundancy process.
remote-operations	(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.
rest-api	(Optional) Restart the rest api process.
root-system-domain-service	(Optional) Restart the Root System Domain (RSD) service.
routing	(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.
routing <logical-system *logical-system-name*>	(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables.

	Optionally, restart the routing protocol process for the specified logical system only.
sampling	(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.
sampling-route-record	(Optional) Restart the sampling route record process.
sbc-configuration-process	(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).
scc	(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).
scc-chassisd	(Optional) Restart the scc chassisd process.
sdk-service	(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.
secure-neighbor-discovery	(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.
security-intelligence	(Optional) Restart security intelligence process.
security-log	(Optional) Restart the security log process.
sfc *number*	(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with `0`.
service-deployment	(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.
services	(Optional) Restart a service.
services pgcp gateway *gateway-name*	(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the `pgcp-service` option.
sflow-service	(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

simple-mail-client-service	(Optional) Restart the simple mail client service process.
snmp	(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.
soft	(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.
static-routed	(Optional) Restart the static routed process.
static-subscribers	(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.
statistics-service	(Optional) Restart the process that manages the Packet Forwarding Engine statistics.
subscriber-management	(Optional) Restart the Subscriber Management process.
subscriber-management-helper	(Optional) Restart the Subscriber Management Helper process.
system-log-vital	(Optional) Restart system log vital process.
tunnel-oamd	(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.
uac-service	(Optional) Restart the Unified Access Control (UAC) process.
usb-control	(MX Series routers) (Optional) Restart the USB control process.
user-ad-authentication	(Optional) Restart User ad Authentication process
vrrp	(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

web-management (QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

Required Privilege Level

reset

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```

restart interface-control gracefully

```
user@host> restart interface-control gracefully
Interface control process started, pid 41129
```

restart interface-control (Junos OS Evolved)

```
user@host> restart interface-control
interface-control restart requested
Restarted aggd on re0
Restarted ifmand on re0
```

Release Information

Command introduced before Junos OS Release 7.4.

Options added:

- `dynamic-flow-capture` in Junos OS Release 7.4.
- `dlsw` in Junos OS Release 7.5.
- `event-processing` in Junos OS Release 7.5.
- `ppp` in Junos OS Release 7.5.
- `l2ald` in Junos OS Release 8.0.
- `link-management` in Junos Release 8.0.
- `pgcp-service` in Junos OS Release 8.4.
- `sbc-configuration-process` in Junos OS Release 9.5.
- `services pgcp gateway` in Junos OS Release 9.6.
- `sfc` and `all-sfc` for the TX Matrix Router in Junos OS Release 9.6.
- Command introduced before Junos OS Release 9.2 on SRX Series Firewalls.
- `bbe-stats-service` in Junos OS Release 18.4R1 on MX Series routers.
- `kernel-health-monitoring` in Junos OS Release 19.1R1.
- Introduced in Junos OS Evolved Release 19.1R1.

RELATED DOCUMENTATION

Overview of Operational Mode Commands

restart chassis-control

IN THIS SECTION

- Syntax | **15649**
- Syntax (TX Matrix Routers) | **15649**
- Syntax (TX Matrix Plus Routers) | **15649**
- Description | **15649**

- Options | 15650
- Required Privilege Level | 15650
- Output Fields | 15650
- Sample Output | 15651
- Release Information | 15651

Syntax

```
restart chassis-control
<gracefully | immediately | soft>
```

Syntax (TX Matrix Routers)

```
restart chassis-control
<all-chassis | all-lcc | lcc number | scc>
<gracefully | immediately | soft>
```

Syntax (TX Matrix Plus Routers)

```
restart chassis-control
<all-chassis | all-lcc | all-sfc | lcc number | sfc number>
<gracefully | immediately | soft>
```

Description

Restart the chassis management process.

NOTE: When GRES is configured and the `restart chassis-control` command is executed on a TX Matrix Plus router with 3D SIBs, we cannot ascertain which Routing Engine becomes a primary. This is due to the chassisd restart. The chassis process or chassisd is responsible for maintaining and retaining primary role and when it is restarted, the new

chassisd is processed based on the router load. This results in one of the Routing Engines being made primary.

Options

lcc *number* (Routing matrix only) (Optional) Restart the software process for a specific routing node that is connected to a PTX Matrix platform. Replace ***number*** with a value from **0** through **3**.

sfc *number* (PTX routers only) (Optional) Restart the software process on the PTX series router (or switch-fabric chassis). Replace ***number*** with **0**.

gracefully (Optional) Restart the software process.

immediately (Optional) Immediately restart the software process.

> **NOTE**: For PTX Series routers, when the `restart chassis-control` command is executed with the `immediately` option set, the PFE's state will be reset to enabled regardless of whether or not the PFE's state is set to disabled when the action occurs.

soft (Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, Border Gateway Protocol (BGP) peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

Required Privilege Level

reset

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

restart chassis-control gracefully

```
user@host> restart chassis-control gracefully
Chassis control process started, pid 1631
```

restart chassis-control soft

```
user@host> restart chassis-control soft
Chassis control process started, pid 1653
```

Release Information

Command introduced before JUNOS Release 7.4.

Command added in Junos Evolved Release 21.4R1.

RELATED DOCUMENTATION

restart | 15633

restart-duration (Protocols)

IN THIS SECTION

- Syntax | 15652
- Hierarchy Level | 15652
- Description | 15652
- Options | 15652
- Required Privilege Level | 15653
- Release Information | 15653

Syntax

```
restart-duration seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (isis | ospf | ospf3 | pim) graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3 | pim) graceful-restart],
[edit protocols (esis | isis | ospf | ospf3 | pim) graceful-restart],
[edit routing-instances routing-instance-name protocols (ospf | ospf3 | pim) graceful-restart],
[edit routing-options graceful-restart]
```

Description

Configure the grace period for graceful restart globally.

Additionally, you can individually configure the duration of the graceful restart period for the End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), and OSPFv3 protocols and for Protocol Independent Multicast (PIM) sparse mode.

Options

seconds—Time for the graceful restart period.

Range:

- The range of values varies according to whether the graceful restart period is being set globally or for a particular protocol:
  - `[edit routing-options graceful-restart]` (global setting)—120 through 900
  - ES-IS—30 through 300
  - IS-IS—30 through 300
  - OSPF/OSPFv3—1 through 3600
  - PIM—30 through 300

Default:

- The default value varies according to whether the graceful restart period is being set globally or for a particular protocol:
  - `[edit routing-options graceful-restart]` (global setting)—300
  - ES-IS—180
  - IS-IS—210
  - OSPF/OSPFv3—180
  - PIM—60

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Graceful Restart

Configuring Graceful Restart for MPLS-Related Protocols

Configuring VPN Graceful Restart

Configuring Graceful Restart for VPNs

Configuring Logical System Graceful Restart

Configuring PIM Sparse Mode Graceful Restart

restart-duration (Multicast Snooping)

IN THIS SECTION

Syntax | 15654

- Hierarchy Level | 15654
- Description | 15654
- Options | 15654
- Required Privilege Level | 15654
- Release Information | 15654

Syntax

```
restart-duration seconds;
```

Hierarchy Level

```
[edit multicast-snooping-options graceful-restart]
```

Description

Configure the duration of the graceful restart interval.

Options

seconds— Graceful restart duration for multicast snooping.

- **Range:** 0 through 300
- **Default:** 180

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Example: Configuring Multicast Snooping

restart-duration (Routing Options)

IN THIS SECTION

- Syntax | 15655
- Hierarchy Level | 15655
- Description | 15655
- Options | 15656
- Required Privilege Level | 15656
- Release Information | 15656

Syntax

```
restart-duration seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options
graceful-restart],
[edit logical-systems logical-system-name routing-options graceful-restart],
[edit routing-instances routing-instance-name routing-options graceful-restart],
[edit routing-options graceful-restart]
```

Description

Configure the restart timer for graceful restart.

Options

seconds—Configure the time period for the restart to last.

- **Range:** 120 through 900 seconds
- **Default:** 300 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS High Availability User Guide

restart extensible-subscriber-services

IN THIS SECTION

- Syntax | 15657
- Description | 15657
- Required Privilege Level | 15657
- Sample Output | 15657
- Release Information | 15657

Syntax

```
restart extensible-subscriber-services
```

Description

Restart essmd.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Sample Output

command-name

```
# restart extensible-subscriber-services
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

disable (Extensible Subscriber Services Manager) | 5753

restart (JDM)

IN THIS SECTION

Syntax | 15658

- Description | 15658
- Options | 15658
- Required Privilege Level | 15659
- Output Fields | 15659
- Sample Output | 15659
- Release Information | 15659

Syntax

```
restart (gracefully | immediately | soft)
<jdmd>
<jdmmon>
<jinventoryd>
<jlinkmon>
```

Description

Restart daemons belonging to Juniper Device Manager (JDM).

CAUTION: Never restart a software process unless instructed to do so by a customer support engineer.

NOTE: Starting in Junos OS Release 19.2R1, JDM does not support the `jinventoryd` option in the `restart` command.

Options

gracefully (Optional) Restart a specified module or component by sending the equivalent of a UNIX SIGTERM signal.

immediately (Optional) Immediately restart a module or component by sending the equivalent of a UNIX SIGKILL signal.

soft (Optional) Reread and reactivate the configuration without completely restarting a module or component. This option is the equivalent of a UNIX SIGHUP signal.

jdmd Restart the JDM service process.

jdmmon Restart the JDM link monitoring daemon.

jinventoryd Restart the JDM inventory management daemon.

jlinkmon Restart the JDM link monitor daemon, which runs on the Linux host.

> **NOTE**: The options `gracefully`, `immediately`, and `soft` are not available for restarting the Juniper link monitor daemon.

Required Privilege Level

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

restart jdmd gracefully

```
user@jdm> restart jdmd gracefully
Juniper Device Manager service process started, pid 12058
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | **15481**

request server authenticate-peer-server | **14901**

restart (Junos OS Evolved)

IN THIS SECTION

- Syntax | 15660
- Description | 15660
- Options | 15661
- Required Privilege Level | 15661
- Output Fields | 15661
- Sample Output | 15661
- Release Information | 15662

Syntax

```
restart application-name
<gracefully | immediately | soft>
```

Description

Restart a Junos OS Evolved process.

> **CAUTION**: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

For Junos OS Evolved, the `restart` command also triggers a restart of the dependent applications (apps). To inform you which dependent apps restarted, the following message appears in the log whenever you use the `restart` command:

`App restarting <app name>. Related apps that may be impacted - <related-app name>` . For example: `Jan 14 11:42:08 RE0 sysman[5100]: SYSTEM_APP_RESTARTING_WITH_RELAPPS_EVENT: App restarting re0-ifmand. Related apps that may be impacted - aggd`

Starting in Junos OS Evolved Release 20.1R1, if you specify the `restart` *`app-name`* command and the application is not supposed to run on the platform, the error message is as follows:

```
user@device> restart fabspoked-pfe
Restart failed for fabspoked-pfe on node re0. Application is not running.
```

The `restart` command expands all application names, including applications not required for the current platform. Therefore, you could try to restart an application that is not running for the current platform. This error message communicates that the restart failed because the application was not running on the system.

Options

application-name	Specify the name of the application you want to restart. Use the `show system applications` operational mode command for information about what applications are running.
none	Same as `gracefully`.
`gracefully`	(Optional) Gracefully restart the software process.
`immediately`	(Optional) Immediately restart the software process.
`soft`	(Optional) Re-read and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

Required Privilege Level

reset

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

restart interface-control (Junos OS Evolved)

```
user@host> restart interface-control
interface-control restart requested
```

```
Restarted aggd on re0
Restarted ifmand on re0
```

Release Information

Introduced in Junos OS Evolved Release 19.1R1.

restart-path

IN THIS SECTION

- Syntax | 15662
- Hierarchy Level | 15662
- Description | 15662
- Options | 15663
- Required Privilege Level | 15663
- Release Information | 15663

Syntax

```
restart-path (all | create | echo);
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

Use the restart-path function to restart a GTP path. By default, the restart path is disabled and then related function does not work. Tunnels are not deleted when restart counter changes.

From Junos OS Release 20.4R1 onwards, the [`edit security gprs gtp profile` *`profile-name`*] hierarchy level is replaced by [`edit security gtp profile` *`profile-name`*].

Options

- **all**—Restart GTP paths by detecting the changed restart number obtained from the Recovery information element (IE) in all GTP messages.
- **create**—Restart GTP paths by detecting the changed restart number obtained from the Recovery IE in create-session messages.
- **echo**—Restart GTP paths by detecting the changed restart number obtained from the Recovery IE in echo messages.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4. Support.

RELATED DOCUMENTATION

Understanding GTP Path Restart

Example: Restarting a GTPv2 Path

restart-time (BGP Graceful Restart)

IN THIS SECTION

- Syntax | 15664
- Hierarchy Level | 15664
- Description | 15664

- Options | 15664
- Required Privilege Level | 15665
- Release Information | 15665

Syntax

```
restart-time seconds;
```

Hierarchy Level

```
[edit protocols (bgp | rip | ripng) graceful-restart],
[edit logical-systems logical-system-name protocols (bgp | rip | ripng) graceful-restart (Enabling Globally)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp graceful-restart],
[edit routing-instances routing-instance-name protocols bgp graceful-restart]
```

Description

Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.

Options

seconds—Length of time for the graceful restart period.

- **Range:** Varies by protocol
  - BGP— 1 through 1800 seconds
  - RIP — 1 through 600 seconds
- **Default:** Varies by protocol:
  - BGP—120 seconds
  - RIP and RIPng—60 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Configuring Graceful Restart Options for RIP and RIPng

Configuring Graceful Restart for QFabric Systems

stale-routes-time | 28478

restart-time (Fabric Control)

IN THIS SECTION

- Syntax | 15665
- Hierarchy Level | 15666
- Description | 15666
- Options | 15666
- Required Privilege Level | 15666
- Release Information | 15666

Syntax

```
restart-time seconds;
```

Hierarchy Level

```
[edit fabric protocols fabric-control graceful-restart]
```

Description

Configure the duration of the graceful restart period for the fabric control Routing Engine.

The graceful restart resynchronization process takes longer when the QFabric system contains node groups that have a large number of VLANs. The graceful-restart duration should, therefore, be set higher when the QFabric system contains at least one node group with a large number of VLANs.

Configure a restart time of 600 seconds if the number of VLAN members (vmembers) exceeds 32k.

> **CAUTION**: Configuring the restart time restarts the session between the fabric control Routing Engine and the Node groups. Traffic is dropped as a result of this restart. Normal QFabric system operations should resume once the session has restarted without any further user actions.

Options

seconds Duration of the graceful restart period.

- **Default:** 300 seconds
- **Range:** 300 to 900 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X52-D10.

RELATED DOCUMENTATION

Understanding Bridging and VLANs on Switches

Understanding Routing Engines in the QFabric System

restarter (Graceful Restart for BGP Restarter)

IN THIS SECTION

- Syntax | 15667
- Hierarchy Level | 15667
- Description | 15668
- Options | 15669
- Required Privilege Level | 15669
- Release Information | 15669

Syntax

```
restarter {
    disable;
    stale-time interval;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family (l2vpn | route-target | inet)
(labeled-unicast | flow) graceful-restart long-lived],
[edit logical-systems logical-system-name protocols bgp group group-name family (l2vpn | route-
target | inet) (labeled-unicast | flow) graceful-restart long-lived],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
(l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
```

```
group group-name family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart
long-lived],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family (l2vpn | route-target | inet) (labeled-unicast | flow)
graceful-restart long-lived],
[edit routing-instances routing-instance-name protocols bgp family (l2vpn | route-target | inet)
(labeled-unicast | flow) graceful-restart long-lived],
[edit routing-instances routing-instance-name protocols bgp group group-name family (l2vpn |
route-target | inet) (labeled-unicast | flow) graceful-restart long-lived],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived],
[edit protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-
restart long-lived],
[edit protocols bgp group group-name family (l2vpn | route-target | inet) (labeled-unicast |
flow) graceful-restart long-lived],
[edit protocols bgp group group-name neighbor address family (l2vpn | route-target | inet)
(labeled-unicast | flow) graceful-restart long-lived]
```

NOTE: Each routing table is identified by the protocol family or address family indicator (AFI) and a subsequent address family identifier (SAFI). The AFI parameter can be one of the `(l2vpn | inet | route-target)` protocols and the SAFI parameter can be either of the `(flow | labeled-unicast)` protocols for inet family and one of the `(auto-discovery-mspw | auto-discovery-only | signaling)` protcols for L2VPN family..

Configuring LLGR does not require that BGP graceful restart also be configured. The long-lived-graceful-restart section is visible only for families l2vpn, inet labeled-unicast, inet flow and route-target. It is prohibited for inet-mvpn, inet6-mvpn and inet-mdt. It is hidden for other families.

Description

Configure the long-lived graceful restart mechanism for a BGP restarter router to preserve BGP routing details for a longer period from a failed BGP peer. You can also configure the BGP long-lived graceful restarter mode negotiation mechanism for a particular address family instead of configuring this capability for all address families in a system, logical system, or routing instance.

The stanzas in the per-family graceful-restart long-lived restarter configuration section enables LLGR restarter mode negotiation for BGP globally, or for a group or neighbor. The values are inherited by groups from the global configuration, and by neighbors from the group configuration. The disable attribute is used to override configuration inherited from a higher level. It does not disable LLGR receiver mode; you must disable LLGR receiver mode explicitly for all families as necessary. A hidden `enable` attribute can be used to override an inherited disable attribute. Configuring graceful-restart long-lived

restarter at the neighbor level (when it is not configured at the containing group level or globally) causes an internal group to be split. When LLGR restarter is enabled or disabled for a family or the stale- time is changed, the session is reset so that the new capability can be sent to the neighbor.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

restore-interval

IN THIS SECTION

- Syntax | 15670
- Hierarchy Level | 15670
- Description | 15670
- Options | 15670
- Required Privilege Level | 15670
- Release Information | 15670

Syntax

```
restore-interval number;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Configures the number of minutes that the node does not process any Ethernet ring protection (ERP) protocol data units (PDUs).. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

Options

number—Specify the restore interval.

- **Range:** 1 through 12 minutes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Example: Configuring Ethernet Ring Protection Switching on EX Series Switches

Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS

Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

restricted-queues

IN THIS SECTION

- Syntax | 15671
- Hierarchy Level | 15671
- Description | 15671
- Required Privilege Level | 15672
- Release Information | 15672

Syntax

```
restricted-queues {
    forwarding-class class-name queue queue-number;
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

For M320, MX Series, T Series routers and EX Series switches only, map forwarding classes to restricted queues.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

retain

IN THIS SECTION

- Syntax | 15672
- Hierarchy Level | 15672
- Description | 15673
- Default | 15673
- Options | 15673
- Required Privilege Level | 15674
- Release Information | 15674

Syntax

```
(no-retain | retain);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name static (defaults | route)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options static (defaults | route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
```

```
(defaults | route)],
[edit logical-systems logical-system-name routing-options static (defaults | route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
(defaults | route)],
[edit routing-instances routing-instance-name routing-options static (defaults | route)],
[edit routing-options rib routing-table-name static (defaults | route)],
[edit routing-options static (defaults | route)]
```

Description

Configure statically configured routes to be deleted from or retained in the forwarding table when the routing protocol process shuts down normally.

> **NOTE**: You cannot configure both `retain` and `resolve` options for a statically configured route because resolved next hops cannot be retained.

Default

Statically configured routes are deleted from the forwarding table when the routing protocol process shuts down normally.

The retention policy applied to defaults and route is as follows:

- defaults— The default route retention policy. By default, the default route retention policy is no-retain for the configured scope using the `defaults` configuration statement. This default retention policy is overridden on a per route basis using the routes option.
- route—The retention policy for this route. The default retention policy for a route is the policy of its routing table.

Options

no-retain Delete statically configured routes from the forwarding table when the routing protocol process shuts down normally. To explicitly specify that routes be deleted from the forwarding table, include the `no-retain` statement. Include this statement when configuring an individual route in the `route` portion of the `static` statement to override a `retain` option specified in the `defaults` portion of the statement.

retain Have a static route remain in the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

static (Routing Options) | 28570

retain-hold-resource

IN THIS SECTION

- Syntax | 15674
- Hierarchy Level | 15674
- Description | 15675
- Required Privilege Level | 15675
- Release Information | 15675

Syntax

```
retain-hold-resource;
```

Hierarchy Level

```
[edit logical-systems name security alg sip (Security)],
[edit logical-systems name tenants name security alg sip (Security)],
[edit security alg sip (Security)],
```

```
[edit services alg sip (Security)],
[edit tenants name security alg sip (Security)]
```

Description

Configure retain-hold-resource to enable the device to not to free media resources for a Session Initiation Protocol (SIP) Application Layer Gateway (ALG), even when a media stream is placed on hold. By default, media stream resources are released when the media stream is held.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sip (Security) | 27977

reth-count (Chassis Cluster)

IN THIS SECTION

- Syntax | 15676
- Hierarchy Level | 15676
- Description | 15676
- Options | 15676
- Required Privilege Level | 15676
- Release Information | 15676

Syntax

```
reth-count number;
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Specify the number of redundant Ethernet (`reth`) interfaces allowed in the chassis cluster. You need to specify the total number of interfaces in the chassis cluster before redundant Ethernet interfaces are created. For example, the `set chassis cluster reth-count 2` allow you to create two reth interfaces (example: reth0 and reth1)

Note that the number of `reth` interfaces configured determines the number of redundancy groups that can be configured and each SRX Series Firewall has a maximum number of reths that it can support.

Options

number —Number of redundant Ethernet interfaces allowed.

- **Range:** 1 through 128
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

retransmission-attempt (DHCP Client)

IN THIS SECTION

- Syntax | 15677
- Hierarchy Level | 15677
- Description | 15677
- Options | 15677
- Required Privilege Level | 15678
- Release Information | 15678

Syntax

```
retransmission-attempts number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcp-client]
```

Description

Specify the number of times the device attempts to retransmit a Dynamic Host Control Protocol (DHCP) packet fallback.

Options

number Number of retransmit attempts.

Range:

- - For IPv4 — 0 through 50000 from Junos OS Release 17.3R1 onwards and 0 through 6 on Junos OS earlier releases.
  - For IPv6 — 0 through 9.

- **Default:** 4

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

Starting with Junos OS Release 17.3R1, on all SRX Series Firewalls and vSRX Virtual Firewall instances, the CLI option `dhcp-clinet` at `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy is changed to `dhcp` to keep consistent with other Junos platforms. There is no change in the functionality with the changed CLI option `dhcp`.

RELATED DOCUMENTATION

Understanding DHCP Client Operation

Minimum DHCP Client Configuration

retransmission-attempt (DHCPv6 Client)

IN THIS SECTION

- Syntax | 15679
- Hierarchy Level | 15679
- Description | 15679
- Options | 15679
- Required Privilege Level | 15679
- Release Information | 15679

Syntax

```
retransmission-attempt number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcpv6-client]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet6 dhcpv6-client]
[edit tenants tenant-name interfaces interface-name unit logical-unit-number family inet6 dhcpv6-
client]
```

Description

Specify the number of times the device retransmits a DHCPv6 client packet if a DHCPv6 server fails to respond. After the specified number of attempts, no further attempts at reaching a server are made.

Options

number Number of retransmit attempts.

- **Range:** 0 through 50000 (from Junos OS Release 20.3R1 onwards)
- **Range:** 0 through 9

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

The `logical-systems` and `tenants` options are introduced in Junos OS Release 18.4R1.

retransmission-count-established (L2TP)

IN THIS SECTION

- Syntax | 15680
- Hierarchy Level | 15680
- Description | 15680
- Options | 15680
- Required Privilege Level | 15681
- Release Information | 15681

Syntax

```
retransmission-count-established count;
```

Hierarchy Level

```
[edit services l2tp tunnel]
```

Description

Set the maximum number of times control messages are retransmitted for established tunnels.

BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing `no services l2tp tunnel retransmission-count-established`.

Options

count—Number of retransmissions.

- **Range:** 2 through 30

- **Default:** 7

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

L2TP for Subscriber Access Overview

L2TP LAC Subscriber Configuration

L2TP LNS Inline Service Interfaces

retransmission-count-not-established (L2TP)

IN THIS SECTION

- Syntax | 15681
- Hierarchy Level | 15682
- Description | 15682
- Options | 15682
- Required Privilege Level | 15682
- Release Information | 15682

Syntax

```
retransmission-count-not-established count;
```

Hierarchy Level

```
[edit services l2tp tunnel]
```

Description

Set the maximum number of times control messages are retransmitted for tunnels that are not established.

> **BEST PRACTICE**: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing `no services l2tp tunnel retransmission-count-not-established`.

Options

count—Number or retransmissions.

- **Range:** 2 through 30
- **Default:** 5

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

> *L2TP for Subscriber Access Overview*
>
> *L2TP LAC Subscriber Configuration*
>
> *L2TP LNS Inline Service Interfaces*

retransmission-interval

IN THIS SECTION

- Syntax | 15683
- Hierarchy Level | 15683
- Description | 15683
- Options | 15683
- Required Privilege Level | 15684
- Release Information | 15684

Syntax

```
retransmission-interval milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-protocol],
[edit protocols link-management peer peer-name lmp-protocol]
```

Description

Specify how often Link Management Protocol (LMP) sends Config and LinkSummary messages on the LMP control channel.

Options

milliseconds—Length of time between Config messages.

- **Range:** 500 through 300,000
- **Default:** 500 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

retry-limit (Protocols Link Management) | 15710

Controlling Message Exchange for LMP Control Channels

retransmission-interval (DHCP)

IN THIS SECTION

- Syntax | 15684
- Hierarchy Level | 15685
- Description | 15685
- Options | 15685
- Required Privilege Level | 15685
- Release Information | 15685

Syntax

```
retransmission-interval  seconds;
```

Hierarchy Level

```
[edit interfaces interface-name  unit                logical-unit-number  family inet dhcp]
```

Description

Specify the time between successive retransmissions of the client DHCP request if a DHCP server fails to respond.

Options

seconds Number of seconds between successive retransmissions.

- **Range:** 4 through 64 seconds
- **Default:** 4 seconds

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring a DHCP Client

retransmission-interval (DHCP Client)

IN THIS SECTION

Syntax | 15686

- Hierarchy Level | 15686
- Description | 15686
- Options | 15686
- Required Privilege Level | 15687
- Release Information | 15687

Syntax

```
retransmission-interval seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcp-client]
```

Description

Specify the initial retransmission interval. Successive retransmission intervals are doubled as per RFC2131.

> **NOTE**: Though the SRX Series Firewalls implement the exponential backoff, as described in RFC 2131, the retransmit attempt does not stop when the retransmission interval reaches 64 seconds. The packet is transmitted till the retransmission attempt is reached. For example, if you configure the `retransmission-attempt` to 5 and the `retransmission-interval` to 20, the sequence of retransmission-interval is 20, 40, 80, 160, 320.

Options

seconds Number of seconds before initial retransmission.

Range: The range is 4 through 64. The default is 4 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

Starting with Junos OS Release 17.3R1, on all SRX Series Firewalls and vSRX Virtual Firewall instances, the CLI option `dhcp-clinet` at `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy is changed to `dhcp` to keep consistent with other Junos platforms. There is no change in the functionality with the changed CLI option `dhcp`.

RELATED DOCUMENTATION

Understanding DHCPv6 Client and Server Identification

retransmit-interval (Services)

IN THIS SECTION

- Syntax | 15687
- Hierarchy Level | 15688
- Description | 15688
- Options | 15688
- Required Privilege Level | 15688
- Release Information | 15688

Syntax

```
retransmit-interval seconds;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Specify the maximum retransmit interval for L2TP tunnels.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options

seconds—Interval, in seconds, after which the server retransmits data if no acknowledgment is received.

- **Default:** 30 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Timers for L2TP Tunnels

retransmit-timer

IN THIS SECTION

- Syntax | 15689
- Hierarchy Level | 15689
- Description | 15689
- Options | 15689
- Required Privilege Level | 15690
- Release Information | 15690

Syntax

```
retransmit-timer milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Description

Set the retransmission frequency of neighbor solicitation messages. This timer is used to detect when a neighbor has become unreachable and to resolve addresses.

Options

milliseconds—Retransmission frequency.

- **Default:** 0 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

Secure IPv6 Neighbor Discovery

retransmit-timer (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 15690
- Hierarchy Level | 15691
- Description | 15691
- Options | 15691
- Required Privilege Level | 15691
- Release Information | 15691

Syntax

```
retransmit-timer milliseconds;
```

Hierarchy Level

```
[edit protocols router-advertisement interface interface-name]
```

Description

Set the retransmission frequency of neighbor solicitation messages.

Options

milliseconds—Retransmission frequency.

- **Default:** 0 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *IPv6 WAN Link Addressing with NDRA*
>
> *Dynamic Router Advertisement Configuration Overview*

retry (Access LDAP)

IN THIS SECTION

- Syntax | 15692
- Hierarchy Level | 15692

- Description | 15692
- Options | 15692
- Required Privilege Level | 15692
- Release Information | 15692

Syntax

```
retry attempts;
```

Hierarchy Level

```
[edit access ldap-server server-address],
[edit access profile profile-name ldap-server server-address]
```

Description

Specify the number of retries that a device can attempt to contact an LDAP server.

Options

attempts—Number of retries that the device is allowed to attempt to contact an LDAP server.

- **Range:** 1 through 10
- **Default:** 3

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring RADIUS and LDAP User Authentication

retry (Access RADIUS)

IN THIS SECTION

- Syntax | 15693
- Hierarchy Level | 15693
- Description | 15693
- Options | 15693
- Required Privilege Level | 15694
- Release Information | 15694

Syntax

```
retry attempts;
```

Hierarchy Level

```
[edit access radius-server server-address],
[edit access profile profile-name radius-server server-address]
```

Description

Specify the number of retries that a device can attempt to contact a RADIUS authentication server.

Options

attempts—Number of retries that the device is allowed to attempt to contact a RADIUS server.

- **Range:** 1 through 10

- **Default:** 3

Required Privilege Level

secret—To view this statement in the configuration.

secret-control—To add this statement to the configuration.

Release Information

Statement modified in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Firewall User Authentication Overview

retry (Adaptive Services Interface)

IN THIS SECTION

- Syntax | 15694
- Hierarchy Level | 15695
- Description | 15695
- Options | 15695
- Required Privilege Level | 15695
- Release Information | 15695

Syntax

```
retry number-of-attempts;
```

Hierarchy Level

```
[edit security pki ca-profile ca-profile-name enrollment]
```

Description

(Adaptive services interfaces only) Specify how many times a router can resend a digital certificate request.

Options

number-of-attempts—Number of enrollment retries.

- **Range:** 0 through 100
- **Default:** 0

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Specifying the Enrollment Properties

enrollment | 6408

retry-count (Security Log)

IN THIS SECTION

- Syntax | 15696
- Hierarchy Level | 15696
- Description | 15696
- Options | 15696
- Required Privilege Level | 15697
- Release Information | 15697

Syntax

```
retry-count retry-count;
```

Hierarchy Level

```
[edit logical-systems name security log],
[edit logical-systems name tenants name security log],
[edit security log],
[edit tenants name security log]
```

Description

Specify the maximum number attempts to send accounting message to the RADIUS accounting server in case of a transmission failure.

Options

- `retry-count` — The number of times you attempt to send accounting messages to the RADIUS accounting server. The interval between each retry message is 1 second.

- **Default**: 5
- **Range**: 1 through 5

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

retry-count (SNMP)

IN THIS SECTION

- Syntax | 15697
- Hierarchy Level | 15697
- Description | 15698
- Options | 15698
- Required Privilege Level | 15698
- Release Information | 15698

Syntax

```
retry-count number;
```

Hierarchy Level

```
[edit snmp v3 target-address target-address-name]
```

Description

Configure the retry count for SNMP informs.

Options

number—Maximum number of times the inform is transmitted if no acknowledgment is received. If no acknowledgment is received after the inform is transmitted the maximum number of times, the inform message is discarded.

- **Default:** 3 times

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

timeout (SNMP) | 29365

retry-count (FPC Error Condition)

IN THIS SECTION

- Syntax | 15699
- Hierarchy Level | 15699
- Description | 15699
- Options | 15699
- Required Privilege Level | 15699
- Release Information | 15699

Syntax

```
retry-count number;
```

Hierarchy Level

```
[edit chassis cfeb slot-number],
[edit chassis cluster redundancy-group group-number ip-monitoring],
[edit chassis feb slot-number],
[edit chassis fpc slot-number sanity-poll],
[edit chassis lcc number fpc number sanity-poll]
```

Description

Number of times sanity polling periodically checks for an error condition in the FPC.

Options

number—Number of times sanity polling is allowed to check for an error condition.

- **Range:** 1 through 30
- **Default:** 10

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Sanity Polling

sanity-poll | **16302**

on-error | 11819

retry-count (Chassis Cluster)

IN THIS SECTION

- Syntax | 15700
- Hierarchy Level | 15700
- Description | 15700
- Options | 15701
- Required Privilege Level | 15701
- Release Information | 15701

Syntax

```
retry-count number;
```

Hierarchy Level

```
[edit chassis cluster redundancy-group group-number ip-monitoring]
```

Description

Specify the number of consecutive ping attempts that must fail before an IP address monitored by the redundancy group is declared unreachable in IP address monitoring.

In the IP address monitoring, you can configure redundancy groups to monitor upstream resources by pinging specific IP addresses that are reachable through redundant Ethernet interfaces on either node in a cluster. In this configuration, the retry interval determines the ping interval for each IP address monitored by the redundancy group and the retry count sets the number of allowed consecutive ping failures for each IP address monitored by the redundancy group.

(See "retry-interval (Chassis Cluster)" on page 15707 for a related redundancy group IP address monitoring variable.)

Options

number—Number of consecutive ping attempt failures before a monitored IP address is declared unreachable.

- **Default:** 5
- **Range:** 5-15

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

retry-count (Event Policy)

IN THIS SECTION

- Syntax | 15702
- Hierarchy Level | 15702
- Description | 15702
- Default | 15702
- Options | 15702
- Required Privilege Level | 15702

Release Information | 15702

Syntax

```
retry-count number retry-interval seconds;
```

Hierarchy Level

```
[edit event-options policy policy-name then event-script filename destination destination-name],
[edit event-options policy policy-name then execute-commands destination destination-name],
[edit event-options policy policy-name then upload filename (filename | committed)
destination destination-name]
```

Description

Configure an event policy to retry a file upload operation if the first attempt fails.

Default

If you do not include this statement, the file upload operation is attempted one time only.

Options

number	Number of retries.
retry-interval *seconds*	Length of time to wait between retries.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configure an Event Policy to Retry the File Upload Action

retry-delay

IN THIS SECTION

- Syntax | 15703
- Hierarchy Level | 15703
- Description | 15703
- Options | 15703
- Required Privilege Level | 15704
- Release Information | 15704

Syntax

```
retry-delay seconds;
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure the amount of time the flow collector interface waits between retry attempts.

Options

seconds—Amount of time between transfer retry attempts.

- **Range:** 0 through 60

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Retry Attempts

retry (Event Policy)

IN THIS SECTION

- Syntax | 15704
- Hierarchy Level | 15705
- Description | 15705
- Default | 15705
- Options | 15705
- Required Privilege Level | 15705
- Release Information | 15705

Syntax

```
retry count number interval seconds;
```

Hierarchy Level

```
[edit event-options policy policy-name then change-configuration]
```

Description

Specify the number of times that Junos OS attempts the `change-configuration` event policy action if the initial attempt fails while acquiring a lock on the configuration database. If you include the `retry` statement, you must configure both the `count` and `interval` statements.

Default

If you do not include the `retry` statement, and the `change-configuration` event policy action fails, the configuration changes specified in the event policy are not implemented or committed.

Options

count ***number*** — The number of attempts to retry the `change-configuration` event policy action upon failure of the initial attempt.

- **Range:** 0 through 10
- **Default:** 0

interval ***seconds*** — The time interval specified in seconds between retry attempts.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configure an Event Policy to Change the Configuration

Example: Change the Configuration Using an Event Policy

Example: Change the Interface Configuration in Response to an Event

change-configuration (Event Policy) | **3284**

commands (Event Policy Change Configuration) | **4432**

retry-interval

IN THIS SECTION

- Syntax | **15706**
- Hierarchy Level | **15706**
- Description | **15706**
- Options | **15706**
- Required Privilege Level | **15707**
- Release Information | **15707**

Syntax

```
retry-interval seconds;
```

Hierarchy Level

```
[edit security pki ca-profile ca-profile-name enrollment]
```

Description

(Adaptive services interfaces only) Specify the amount of time the router should wait between enrollment retries.

Options

seconds—Time interval, in seconds, between enrollment retries.

- **Range:** 0 through 3600
- **Default:** 0

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Specifying the Enrollment Properties

enrollment | 6408

retry-interval (Chassis Cluster)

IN THIS SECTION

- Syntax | 15707
- Hierarchy Level | 15708
- Description | 15708
- Options | 15708
- Required Privilege Level | 15708
- Release Information | 15708

Syntax

```
retry-interval interval;
```

Hierarchy Level

```
[edit chassis cluster redundancy-group group-number  ip-monitoring ]
```

Description

Specify the ping packet send frequency (in seconds) for each IP address monitored by the redundancy group.

In the IP address monitoring, you can configure redundancy groups to monitor upstream resources by pinging specific IP addresses that are reachable through redundant Ethernet interfaces on either node in a cluster. In this configuration, the retry interval determines the ping interval for each IP address monitored by the redundancy group and the retry count sets the number of allowed consecutive ping failures for each IP address monitored by the redundancy group.

(See "retry-count (Chassis Cluster)" on page 15700 for a related IP address monitoring configuration variable.)

Options

interval—Pause time (in seconds) between each ping sent to each IP address monitored by the redundancy group.

- **Default:** `1`
- **Range:** `1-30`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

ip-monitoring | 9244

retry-limit

IN THIS SECTION

- Syntax | 15709
- Hierarchy Level | 15709
- Description | 15709
- Options | 15709
- Required Privilege Level | 15710
- Release Information | 15710

Syntax

```
retry-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Maximum number of times the ingress router tries to establish the primary path. This counter is reset each time a primary path is created successfully. When the limit is exceeded, no more connection attempts are made. Intervention is then required to restart the connection.

Options

number—Maximum number of tries to establish the primary path.

- **Range:** 0 through 10,000
- **Default:** 0 (The ingress node never stops trying to establish the primary path.)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Connection Between Ingress and Egress Routers

retry-limit (Protocols Link Management)

IN THIS SECTION

- Syntax | 15710
- Hierarchy Level | 15711
- Description | 15711
- Options | 15711
- Required Privilege Level | 15711
- Release Information | 15711

Syntax

```
retry-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-protocol],
[edit protocols link-management peer peer-name lmp-protocol]
```

Description

Specify how many times the Link Management Protocol (LMP) sends Config and LinkSummary messages on the LMP control channel without receiving an appropriate acknowledgment before it logs a message and restarts the LMP control channel configuration process.

Options

number—Maximum number of times messages are sent without receiving an acknowledgment.

- **Range:** 3 through 1000
- **Default:** 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

retransmission-interval | **15683**

Controlling Message Exchange for LMP Control Channels

retry-options

IN THIS SECTION

- Syntax (Junos OS) | 15712
- Syntax (Junos OS Evolved) | 15712
- Hierarchy Level | 15712
- Description | 15713
- Options | 15713
- Required Privilege Level | 15714
- Release Information | 15714

Syntax (Junos OS)

```
retry-options {
    backoff-factor  seconds;
    backoff-threshold number;
    lockout-period minutes
    maximum-time seconds;
    minimum-time seconds;
    tries-before-disconnect number;
}
```

Syntax (Junos OS Evolved)

```
retry-options {
    lockout-period minutes;
}
```

Hierarchy Level

```
[edit system login]
```

Description

Limit the number of times a user can attempt to log in through SSH or Telnet before being disconnected.

Options

backoff-factor *seconds*

Length of delay in seconds after each failed login attempt. The length of delay increases by this value for each subsequent failed login attempt after the value specified in the `backoff-threshold` option.

- **Default:** 5
- **Range:** 5 through 10

backoff-threshold *number*

Threshold for the number of failed login attempts before the user experiences a delay when attempting to log in again. Use the `backoff-factor` option to specify the length of delay, in seconds.

- **Default:** 2
- **Range:** 1 through 3

lockout-period *minutes*

Amount of time before the user can attempt to log in to the device after being locked out. The user is locked out when the number of failed login attempts specified in the `tries-before-disconnect` option is reached.

- **Range:** 1 through 43200

maximum-time *seconds*

Maximum length of time that the connection remains open for the user to enter a username and password to log in. If the user remains idle and does not enter a username and password within the time period configured with this option, the connection is closed.

- **Default:** 120
- **Range:** 20 through 300

minimum-time *seconds*

Minimum length of time that the connection remains open while the user is attempting to enter a username and password to log in.

- **Default:** 20
- **Range:** 20 through 60

tries-before-disconnect

Maximum number of times a user is allowed to attempt to log in through SSH or Telnet before closing the connection.

- **Default:** 3
- **Range:** 2 through 10

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

`maximum-time` option introduced in Junos OS Release 9.6.

`lockout-period` option introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Limiting the Number of User Login Attempts for SSH and Telnet Sessions

rate-limit | **14063**

retry (RADIUS Server)

IN THIS SECTION

- Syntax | **15715**
- Hierarchy Level | **15715**
- Description | **15715**
- Options | **15715**
- Required Privilege Level | **15715**
- Release Information | **15715**

Syntax

```
retry attempts;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure a limit to the number of times the MX Series router can resend a request to the RADIUS server when no response from the RADIUS server is received. If the number of retries reaches this limit, the RADIUS server is marked as dead, and the MX Series router begins to send requests to other RADIUS servers in the network element.

Options

attempts Number of attempts allowed.

- **Range:** 1 through 10
- **Default:** 3

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

retry (Services Flow Collector)

IN THIS SECTION

- Syntax | 15716
- Hierarchy Level | 15716
- Description | 15716
- Options | 15716
- Required Privilege Level | 15717
- Release Information | 15717

Syntax

```
retry number;
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure the maximum number of attempts the flow collector interface make to transfer log files to the FTP server.

Options

number—Maximum number of transfer retry attempts.

- **Range:** 0 through 10

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Retry Attempts

retry-timer

IN THIS SECTION

- Syntax | 15717
- Hierarchy Level | 15718
- Description | 15718
- Options | 15718
- Required Privilege Level | 15718
- Release Information | 15718

Syntax

```
retry-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Amount of time the ingress router waits between attempts to establish the primary path.

Options

seconds—Amount of time between attempts to connect to the primary path.

- **Range:** 1 through 600 seconds
- **Default:** 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Connection Between Ingress and Egress Routers

reverse

IN THIS SECTION

Syntax | 15719

- Hierarchy Level | 15719
- Description | 15719
- Required Privilege Level | 15719
- Release Information | 15719

Syntax

```
reverse {
    ssh {
        port port;
    }
    telnet {
        port port;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Allow connections to device connected to the AUX port. Available for SRX series devices only.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

reverse-oif-mapping

IN THIS SECTION

- Syntax | 15720
- Hierarchy Level | 15720
- Description | 15720
- Required Privilege Level | 15720
- Release Information | 15721

Syntax

```
reverse-oif-mapping {
    no-qos-adjust;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast interface interface-name],
[edit logical-systems logical-system-name routing-options multicast interface interface-name],
[edit routing-instances routing-instance-name routing-options multicast interface interface-name],
[edit routing-options multicast interface interface-name]
```

Description

Enable the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD request it receives over the multicast VLAN.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

The **no-qos-adjust** statement added in Junos OS Release 9.5.

The **no-qos-adjust** statement introduced in Junos OS Release 9.5 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring Multicast with Subscriber VLANs

reverse-route-packet-mode-vr

IN THIS SECTION

- Syntax | 15721
- Hierarchy Level | 15721
- Description | 15722
- Required Privilege Level | 15722
- Release Information | 15722

Syntax

```
reverse-route-packet-mode-vr;
```

Hierarchy Level

```
[edit security flow advanced-options]
```

Description

When processing a server to client traffic, if the route is changed, the server to client traffic is rerouted. The traffic is rerouted using the virtual router from the first incoming session interface. The command allows reverse route lookup with packet mode using virtual router.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Understanding Packet-Based Processing

reverse-tcp-mss

IN THIS SECTION

- Syntax | 15722
- Hierarchy Level | 15723
- Description | 15723
- Required Privilege Level | 15723
- Release Information | 15723

Syntax

```
reverse-tcp-mss mss-value;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name
policy policy-name then permit tcp-options]
```

Description

Configure the TCP maximum segment size (MSS) for packets that match a specific policy and travel in the reverse direction of a session. The value you configure replaces the TCP MSS value when the value in the packet is higher than the one you specify.

The `reverse-tcp-mss` value per policy takes precedence over a global `tcp-mss` value (`all-tcp`, `ipsec-vpn`, `gre-in`, `gre-out`), if one is configured. However, when the `syn-flood-protection-mode syn-proxy` statement at the `[edit security flow] hierarchy level` is used to enable SYN proxy defenses against SYN attacks, the TCP MSS value is not overridden.

Because each policy has two directions, you can configure a value for both directions or for just one direction. To configure the TCP MSS value for the initial session, use the `initial-tcp-mss` option.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D20.

RELATED DOCUMENTATION

initial-tcp-mss | 8659

tcp-mss (Security Flow)

syn-flood-protection-mode

revert-interval (Access)

IN THIS SECTION

- Syntax | 15724
- Hierarchy Level | 15724
- Description | 15724
- Options | 15724
- Required Privilege Level | 15725
- Release Information | 15725

Syntax

```
revert-interval interval;
```

Hierarchy Level

```
[edit access profile profile-name radius options],
[edit access radius-options]
```

Description

Configure the amount of time the router or switch waits after a server has become unreachable. The router or switch rechecks the connection to the server when the specified interval expires. If the server is then reachable, it is used in accordance with the order of the server list.

Options

interval—Amount of time to wait.

- **Range:** 0 through 604,800 seconds
- **Default:** 60 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

RADIUS Authentication and Accounting Basic Configuration

revert-interval (Access LDAP)

IN THIS SECTION

- Syntax | 15725
- Hierarchy Level | 15726
- Description | 15726
- Options | 15726
- Required Privilege Level | 15726
- Release Information | 15726

Syntax

```
revert-interval seconds;
```

Hierarchy Level

```
[edit access ldap-options],
[edit access profile profile-name ldap-options]
```

Description

Specify the amount of time that elapses before the primary server is contacted if a backup server is being used.

Options

seconds—Number of seconds that elapse before the primary server is contacted.

- **Range:** 60 through 4,294,967,295 seconds
- **Default:** 600 seconds

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring RADIUS and LDAP User Authentication

revert-interval (Access RADIUS)

IN THIS SECTION

- Syntax | 15727
- Hierarchy Level | 15727
- Description | 15727
- Options | 15727
- Required Privilege Level | 15728
- Release Information | 15728

Syntax

```
revert-interval seconds;
```

Hierarchy Level

```
[edit access radius-options]
```

Description

Specify the amount of time that elapses before the primary server is contacted if a backup server is being used.

Options

seconds—Number of seconds that elapse before the primary server is contacted.

- **Range:** 60 through 4,294,967,295 seconds
- **Default:** 600 seconds

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

revert-interval (RADIUS Server)

IN THIS SECTION

- Syntax | 15728
- Hierarchy Level | 15728
- Description | 15729
- Options | 15729
- Required Privilege Level | 15729
- Release Information | 15729

Syntax

```
revert-interval seconds;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure the amount of time that must pass after a RADIUS server is first marked dead until it is marked as alive by Junos OS. When Junos OS marks the RADIUS server as alive, it can again send requests to the RADIUS server.

Options

seconds Number of seconds after which a dead server is marked active.

- **Range:** 0 through 4,294,967,295
- **Default:** 300 seconds

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

revert-time (Protocols Layer 2 Circuits)

IN THIS SECTION

Syntax | 15730

Hierarchy Level | 15730

- Description | 15730
- Default | 15731
- Options | 15731
- Required Privilege Level | 15731
- Release Information | 15732

Syntax

```
revert-time seconds maximum seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor address],
[edit protocols l2circuit neighbor address interface interface-name],
[edit routing-instances routing-instance-name protocols vpls neighbor address]
```

Description

Specify a revert time for redundant Layer 2 circuits and VPLS pseudowires. When you have configured redundant pseudowires for Layer 2 circuits or VPLS, traffic is switched to the backup connection in the event that the primary connection fails. If you configure a revert time, when the configured time expires traffic is reverted to the primary path, assuming the primary path has been restored.

With the `maximum` option, specify a maximum reversion interval to add after the `revert-time` delay. If a revert-time delay is defined but a maximum timer is not defined, VCs are restored upon the revert-timer's expiration.

To reduce as much as possible the amount of traffic discarded, and potential data-path asymmetries observed during primary-to-backup transition periods, you can use this restoration timer. This restoration timer is activated when the backup path is performing as active, and then the primary path is restored. The goal is to avoid moving traffic back to the primary path right away, to make sure that the control plane's related tasks (such as IGP, LDP, RSVP, and internal BGP) have enough time to complete their updating cycle.

By enabling a gradual return of traffic to the primary path, you can ensure that the relatively-slow control-plane processing and updating does not have a negative impact on the restoration process.

The `maximum` option extends the revert timer's functionality to provide a jittered interval over which a certain number of circuits can be transitioned back to the primary path. By making use of this maximum value, you can define a time interval during which circuits are expected to switch over. As a consequence, circuits' effective transitions are scattered during restoration periods.

When making use of `revert-time` *`x`* `maximum` *`y`* statement, you can ensure that the corresponding circuit that is active is moved to the primary path within a time-slot (t1) such as that: x <= t1 <= y. In other words, by activating this statement, you can ensure the following:

- VCs stay in the backup path for at least x seconds after the primary path comes back up.
- VCs are moved back to the primary path before y seconds have elapsed.
- y maximum value = x maximum value * 2 = 7200 seconds.

The ideal values for x and y will are conditioned to internal aspects of your network. For this reason, there are no default values for these settings. If no revert-time is set, the default behavior is non-revertive. That is, circuits are not returned to the primary path upon restoration. They are kept on the backup path.

Default

Without the `revert-time` statement, virtual circuit (VC) traffic is not transitioned to the primary path upon restoration of the primary path.

Options

seconds—Revert time in seconds.

- **Range:** 0 through 3600 seconds

maximum *seconds*—Number of seconds to delay path restoration after the `revert-time` delay.

- **Range:** 0 through 7200 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

`maximum` option introduced in Junos OS Release 13.2.

The range and maximum range for revert-time increased from 600 to 3600 and 1200 to 7200 seconds in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

> *Configuring Redundant Pseudowires for Layer 2 Circuits and VPLS*

revert-timer

IN THIS SECTION

- Syntax | **15732**
- Hierarchy Level | **15732**
- Description | **15733**
- Options | **15733**
- Required Privilege Level | **15733**
- Release Information | **15733**

Syntax

```
revert-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
```

```
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Specify the amount of time (in seconds) that an LSP must wait before traffic reverts to a primary path. If during this time the primary path experiences any connectivity problem or stability problem, the timer is restarted.

If you have configured BFD on the LSP, the Junos OS waits until the BFD session is restored before starting the revert timer counter.

If you have configured a value of 0 seconds for the `revert-timer` statement and traffic is switched to the secondary path, the traffic remains on that path indefinitely. It is never switched back to the primary path unless you intervene.

Options

seconds—Time in seconds.

- **Range:** 0 through 65,535 seconds
- **Default:** 60 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

BFD behavior modified in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring the Revert Timer for LSPs

revertive

IN THIS SECTION

- Syntax | 15734
- Hierarchy Level | 15734
- Description | 15734
- Required Privilege Level | 15734
- Release Information | 15735

Syntax

```
revertive;
```

Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
```

Description

Enable the ability to switch to a better priority link (if one is available).

NOTE: By default, LACP link protection is revertive. However, you can use this statement to define a specific aggregated Ethernet interface as revertive to override a global `non-revertive` statement specified at the `[edit chassis]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

non-revertive (Chassis) | 11661

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

revision-level

IN THIS SECTION

- Syntax | 15735
- Hierarchy Level | 15736
- Description | 15736
- Options | 15736
- Required Privilege Level | 15736
- Release Information | 15736

Syntax

```
revision-level revision-level;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mstp],
[edit logical-systems logical-system-name routing-instances routing-instance-name
protocols mstp],
```

```
[edit protocols mstp],
[edit routing-instances routing-instance-name protocols mstp]
```

Description

For Multiple Spanning Tree Protocol (MSTP), set the revision number of the MSTP configuration.

Options

revision-level—Configure the revision number of the MSTP region configuration.

- **Range:** 0 through 65,535

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding MSTP

Configuring Multiple Spanning Tree Protocol

Example: Configuring Network Regions for VLANs with MSTP on Switches

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

revocation-check

IN THIS SECTION

- Syntax | 15737
- Hierarchy Level | 15737
- Description | 15737
- Options | 15738
- Required Privilege Level | 15738
- Release Information | 15738

Syntax

```
revocation-check {
    disable;
    crl {
        refresh-interval number-of-hours;
        url {
            url-name;
        }
    }
}
```

Hierarchy Level

```
[edit security pki ca-profile ca-profile-name]
```

Description

Specify the method to verify revocation status of digital certificates for J Series Services Routers and Adaptive Services (AS) and MultiServices PICs installed in M Series and T Series routers.

Options

`disable`—Disable verification of status of digital certificates. Use `disable` temporarily in cases where a certificate authority (CA) server is unreachable and certificate cannot be renewed or if the certificate download fails.

`crl`—Only certificate revocation list (CRL) is supported. A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis. By default, `crl` is enabled.

The PKID process might fail after RG0 failover on the new node causing all the IPsec VPNs using the public key infrastructure (PKI) to go down when:

- A local certificate used for IPsec VPN is revoked by the Certificate Authority (CA).
- Certificate revocation list (CRL) check is disabled.
- CRL is not cleared.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring the Certificate Revocation List

revocation-check (Security PKI)

IN THIS SECTION

- Syntax | 15739
- Hierarchy Level | 15739
- Description | 15739
- Options | 15740
- Required Privilege Level | 15740
- Release Information | 15741

Syntax

```
revocation-check {
    crl:
    disable;
    ocsp:
    use-crl;
    use-ocsp;
}
```

Hierarchy Level

```
[edit security pki ca-profile ca-profile-name]
```

Description

Specify the method the device uses to verify the revocation status of digital certificates.

Options

crl Only certificate revocation list (CRL) is supported. A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis.

You should also specify the location (URL) to retrieve the CRL (HTTP or LDAP). By default, the URL is empty and uses CDP information embedded in the CA certificate.

For Example: `set security pki ca-profile ms-ca revocation-check crl url http://labsrv1.labdomain.com/CertEnroll/LABDOMAIN.crl`

The URL can include the server-name or port information such as, ldap://<ip-or-fqdn>:<port>). If the port number is missing, HTTP uses port 80, or LDAP uses port 443. Currently, you can configure only one URL. We do not support for configuring backup URL.

By default, `crl` is enabled. Local certificates are being validated against certificate revocation list (CRL) even when CRL check is disabled. This can be stopped by disabling the CRL check through the Public Key Infrastructure (PKI) configuration. When CRL check is disabled, PKI will not validate local certificate against CRL.

disable Disable verification of status of digital certificates.

ocsp Configure Online Certificate Status Protocol (OCSP) to check the revocation status of a certificate.

use-crl Specify the CRL as the method to check the revocation status of a certificate. CRL is the default method.

When you enable this option, you choose CRL as a method to verify the revocation status of digital certificates.

use-ocsp Specify the Online Certificate Status Protocol (OCSP) as the method to check the revocation status of a certificate. CRL is the default method.

When you enable this option, you choose OCSP as a method to verify the revocation status of digital certificates.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5. Support for `ocsp`, `use-crl`, and `use-ocsp` options added in Junos OS Release 12.1X46-D20.

RELATED DOCUMENTATION

Basic Elements of PKI in Junos OS

ca-profile (Security PKI) | **3157**

rewrite-rules (Definition)

IN THIS SECTION

- Syntax | **15741**
- Hierarchy Level | **15742**
- Description | **15742**
- Options | **15742**
- Required Privilege Level | **15742**
- Release Information | **15742**

Syntax

```
rewrite-rules {
    type rewrite-name{
        import (rewrite-name | default);
        forwarding-class class-name {
            loss-priority level code-point [ aliases ] [ 6-bit-patterns ];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.

Options

- ***rewrite-name***—Name of a **rewrite-rules** mapping.
- ***type***—Traffic type.
- **Values: dscp**, **dscp-ipv6**, **exp**, **frame-relay-de** (J Series only), **ieee-802.1**, **ieee-802.1ad**, **inet-precedence**

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

`ieee-802.1ad` option introduced in 9.2 of Junos OS.

RELATED DOCUMENTATION

rewrite-rules (Interfaces) | **15749**

rewrite-rules (CoS Host Outbound Traffic)

IN THIS SECTION

- Syntax | 15743
- Hierarchy Level | 15743
- Description | 15743
- Required Privilege Level | 15743
- Release Information | 15744

Syntax

```
rewrite-rules;
```

Hierarchy Level

```
[edit class-of-service host-outbound-traffic ieee-802.1]
```

Description

Apply the IEEE 802.1p rewrite rules associated with the egress logical interface to the IEEE 802.1p PCP field for all host outbound traffic on that interface.

NOTE: Enabling IEEE 802.1p rewrite rules for host outbound traffic on a DPC without creating any corresponding IEEE 802.1p rewrite rules on a logical interface on the DPC causes the IEEE 802.1p code point to be automatically set to 000 for all host generated traffic that exits that logical interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface

Rewriting Packet Headers to Ensure Forwarding Behavior

Configuring Rewrite Rules

rewrite-rules (Definition)

IN THIS SECTION

- Syntax (Rewrite Rule Configuration) | **15744**
- Syntax (Rewrite Rule Association with Interface) | **15745**
- Hierarchy Level (Rewrite Rule Configuration) | **15745**
- Hierarchy Level (Rewrite Rule Association with Interface) | **15745**
- Description | **15745**
- Options | **15745**
- Required Privilege Level | **15745**
- Release Information | **15746**

Syntax (Rewrite Rule Configuration)

```
rewrite-rules {
    (dscp | dscp-ipv6 | ieee-802.1 | exp) rewrite-name {
        import (rewrite-name | default);
        forwarding-class class-name {
            loss-priority priority code-point (alias | bits);
        }
```

```
    }
}
```

Syntax (Rewrite Rule Association with Interface)

```
rewrite-rules {
    (dscp | dscp-ipv6 | ieee-802.1 | exp) rewrite-name;
}
```

Hierarchy Level (Rewrite Rule Configuration)

```
[edit class-of-service],
```

Hierarchy Level (Rewrite Rule Association with Interface)

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Configure rewrite rules that map traffic to code points when traffic exits the system, and apply the rewrite rules to a specific interface.

MPLS EXP rewrite rules can only be bound to logical interfaces, not to physical interfaces. You can configure up to 64 EXP rewrite rules, but you can use only 16 EXP rewrite rules on switch interfaces at any given time.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Defining CoS Rewrite Rules

Configuring Rewrite Rules for MPLS EXP Classifiers

Understanding CoS Rewrite Rules

Understanding CoS MPLS EXP Classifiers and Rewrite Rules

rewrite-rules (Definition)

IN THIS SECTION

- Syntax | 15746
- Hierarchy Level | 15747
- Description | 15747
- Options | 15747
- Required Privilege Level | 15747
- Release Information | 15747

Syntax

```
rewrite-rules {
    type rewrite-name {
        import (rewrite-name | default);
        forwarding-class class-name {
            loss-priority level code-point [ aliases ] [ bit-patterns ];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.

Options

rewrite-name—Name of a `rewrite-rules` mapping.

type—Traffic type.

- **Values:** `dscp`, `dscp-ipv6`, `exp`, `ieee-802.1`, `ieee-802.1ad`, `inet-precedence`, `inet6-precedence`

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`ieee-802.1ad` option introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Rewrite Rules

Understanding CoS MPLS EXP Classifiers and Rewrite Rules

Configuring Rewrite Rules for MPLS EXP Classifiers

rewrite-rules (Dynamic CoS Interfaces)

IN THIS SECTION

- Syntax | 15748
- Hierarchy Level | 15748
- Description | 15748
- Options | 15748
- Required Privilege Level | 15749
- Release Information | 15749

Syntax

```
rewrite-rules {
    dscp (rewrite-name | default);
    dscp-ipv6 (rewrite-name | default);
    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
    inet-precedence (rewrite-name | default);
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Associate a rewrite-rules configuration or default mapping with a specific interface in a dynamic profile.

Options

rewrite-name—Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules]` hierarchy level.

`default`—The default mapping.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *Guidelines for Configuring Dynamic CoS for Subscriber Access*
>
> rewrite-rules

rewrite-rules (Interfaces)

IN THIS SECTION

- Syntax | **15750**
- Hierarchy Level | **15750**
- Description | **15750**
- Options | **15750**
- Required Privilege Level | **15750**
- Release Information | **15750**

Syntax

```
rewrite-rules {
    dscp (rewrite-name | default);
    dscp-ipv6 (rewrite-name | default);
    exp (rewrite-name | default) protocol protocol-types;
    exp-push-push-push default;
    exp-swap-push-push default;
    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
    inet-precedence (rewrite-name | default);
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Associate a rewrite-rules configuration or default mapping with a specific interface.

Options

- ***rewrite-name***—Name of a **rewrite-rules** mapping configured at the `[edit class-of-service rewrite-rules]` hierarchy level.
- **default**—The default mapping.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

rewrite-rules (Definition) | 15741

rewrite-rules (Interfaces)

IN THIS SECTION

- Syntax | 15751
- Hierarchy Level | 15751
- Description | 15752
- Options | 15752
- Required Privilege Level | 15752
- Release Information | 15752

Syntax

```
rewrite-rules {
    dscp (rewrite-name | default) protocol (inet-both | inet-outer | mpls);
    dscp-ipv6 (rewrite-name | default) protocol mpls;
    exp (rewrite-name | default) protocol protocol-types;
    exp-push-push-push default;
    exp-swap-push-push default;
    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
    ieee-802.1ad (rewrite-name | default) vlan-tag (outer | outer-and-inner);
    inet-precedence (rewrite-name | default) protocol (inet-both | inet-outer | mpls);
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Associate a rewrite-rules configuration or default mapping with a specific interface.

Integrated Bridging and Routing (IRB) interfaces are used to tie together Layer 2 switched and Layer 3 routed domains on MX routers. MX routers support classifiers and rewrite rules on the IRB interface at the `[edit class-of-service interfaces irb unit logical-unit-number]` level of the hierarchy. All types of classifiers and rewrite rules are allowed, including IEEE 802.1p.

> **NOTE**: The IRB classifiers and rewrite rules are used only for *routed* packets; in other words, it is for traffic that originated in the Layer 2 domain and is then routed through IRB into the Layer 3 domain, or vice versa. Only IEEE classifiers and IEEE rewrite rules are allowed for pure Layer 2 interfaces within a bridge domain.

On an MX Series router and on an EX Series switch, `exp-push-push-push`, `exp-swap-push-push`, and `frame-relay-de` are not supported on an integrated routing and bridging (IRB) interface.

The `[edit class-of-service interfaces interface-name unit logical-unit-number]` hierarchy level is not supported on ACX Series routers.

On an ACX Series router, only the outer tag is supported for `dscp`, `inet-precedence`, and `ieee802.1`.

For IQ PICs, you can configure only one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.

Options

`rewrite-name`—Name of a `rewrite-rules` mapping configured at the `[edit class-of-service rewrite-rules]` hierarchy level.

`default`—The default mapping.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Rewrite Rules

rewrite-rules (Definition) | **15746**

Applying Rewrite Rules to Output Logical Interfaces

rewrite-value (Fibre Channel Interfaces)

IN THIS SECTION

- Syntax | **15753**
- Hierarchy Level | **15753**
- Description | **15753**
- Required Privilege Level | **15754**

Syntax

```
rewrite-value {
    input {
        ieee-802.1{
            code-point code-point-bits;
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service interfaces fibre-channel-interface-name]
```

Description

Configure the IEEE 802.1p code point value (priority) for all traffic received from the Fibre Channel (FC) network on the specified FC interface. Instead of using the default priority 3 (011) for FCoE traffic, the

priority is rewritten to the specified priority before being forwarded. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.

FCoE traffic typically uses priority 3 (IEEE code point 011). The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011) for FCoE traffic.

The system supports only one IEEE 802.1p code point value per FC interface, so you cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.

NOTE: In order to avoid fate sharing (separate flows that affect each others' throughput), the code point used for the rewrite value should be the only code point used for the lossless FCoE forwarding class (the forwarding class used for the fixed classifier on the Fibre Channel interface). When you configure classifiers for ingress Ethernet interfaces, map only the rewrite value code point to the forwarding class.

For example, if the rewrite value uses code point 101 for lossless FCoE forwarding class fcoe_fc1, then in the classifier configuration attached to ingress Ethernet interfaces, code point 101 is the only code point that should be classified to the fcoe_fc1 forwarding class. Now if you also attach a classifier to an interface that maps code point 110 to forwarding class fcoe_fc1, then congestion on priority 110 unfairly (and unintentionally) affects the FCoE traffic that uses priority 101. Both priorities 101 and 110 are classified into forwarding class fcoe_fc1, so the traffic from both priorities shares the same fate.

The remaining statements are described separately.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

forwarding-class (Fibre Channel Interfaces) | 7421

Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

Understanding CoS Classifiers

rfc-compliant (NETCONF)

IN THIS SECTION

- Syntax | 15755
- Hierarchy Level | 15755
- Description | 15755
- Default | 15756
- Required Privilege Level | 15757
- Release Information | 15757

Syntax

```
rfc-compliant;
```

Hierarchy Level

```
[edit system services netconf]
```

Description

Require that the NETCONF server enforce certain behaviors that are compliant with RFC 4741, *NETCONF Configuration Protocol*, during NETCONF sessions.

When you configure the `rfc-compliant` statement:

- The NETCONF server explicitly declares the NETCONF namespace in its replies and qualifies all NETCONF tags with the `nc` prefix.

- <get> and <get-config> operations that return no configuration data do not include an empty <configuration> element in RPC replies.
- On devices running Junos OS Release 17.2R1 or later, the NETCONF server sets the default namespace for the <configuration> element to the same namespace as in the corresponding YANG model.
- On devices running Junos OS Release 17.4R3, 18.2R2, 18.3R2, and 18.4R1 or later, the NETCONF server omits <rpc-error> elements with a severity level of warning in its replies when the operation is successful and returns an <ok/> element.
- On devices running Junos OS Release 21.2R1 or later, the NETCONF server's response to <commit> operations includes the following changes:
  - If a successful <commit> operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
  - The NETCONF server response emits the <source-daemon> element as a child of the <error-info> element instead of the <rpc-error> element.
  - If you also configure the flatten-commit-results statement at the [edit system services netconf] hierarchy level, the NETCONF server suppresses any <commit-results> XML subtree and only emits an <ok/> or <rpc-error> element in its response.
- On devices running Junos OS Release 23.2R1 or later and on devices running Junos OS Evolved Release 23.4R1 or later, the NETCONF server emits only an <ok/> or <rpc-error> element in response to <validate> operations. In earlier releases, the RPC reply also includes the <commit-results> element.

Default

If you do not include the rfc-compliant statement:

- The NETCONF server sets the default namespace to the NETCONF namespace in RPC replies.
- <get> and <get-config> operations that return no configuration data include an empty <configuration> element in RPC replies.
- The NETCONF server does not set the default namespace for the <configuration> element to the same namespace as in the corresponding YANG model.
- The NETCONF server might issue an RPC reply that includes both an <rpc-error> element with a severity level of warning and an <ok/> element.
- The NETCONF server's response to <commit> and <validate> operations might not be RFC compliant.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring RFC-Compliant NETCONF Sessions

rfc2544

IN THIS SECTION

- Syntax | 15757
- Hierarchy Level | 15758
- Description | 15758
- Default | 15759
- Required Privilege Level | 15759
- Release Information | 15759

Syntax

```
rfc2544 {
    test-profile profile-name {
        bandwidth bandwidth;
        packet-size size;
        step-percent percent;
        test-type (throughput | latency | frame-loss | back-back-frames);
    }
```

```
    tests{
        test-name test-name {
            destination-ipv4-address address;
            destination-mac-address mac-address;
            destination-udp-port port-number;
            direction (egress | ingress);
            disable-signature-check;
            dscp-code-points codepoint;
            family (ccc | ethernet-switching |inet);
            halt-on-prefix-down;
            in-service;
            ip-swap;
            mode (initiate-and-terminate | reflect);
            reflect-etype;
            reflect-mode (mac-swap | no-mac-swap);
            service-type service-type;
            skip-arp-iteration;
            source-ipv4-address address;
            source-mac-address mac-address;
            source-udp-port port-number;
            test-interface interface-name;
            test-iterator-duration seconds;
            test-profile profile-name;
            udp-tcp-port-swap;
        }
    }
}
```

Hierarchy Level

```
[edit services monitoring]
```

Description

Configure the parameters for the RFC 2544-based benchmarking test to examine and analyze the performance characteristics of a network interconnecting device. You must configure a test profile, which specifies the type of test and the manner in which it must be performed, and associate the test profile with a test name. The test name that you configure contains details, such as the address family and the test mode, for the test. You can associate the same test profile with multiple test names.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Default

Disabled

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 21.1R1.

Support for the `destination-mac-address`, `direction`, `disable-signature-check`, `in-service`, `ip-swap`, `reflect-etype`, `reflect-mode`, `service-type`, `source-mac-address`, and `udp-tcp-port-swap` configuration statements introduced in Junos OS Evolved 22.4R1.

Support for `bandwidth`, `dscp-code-points`, `halt-on-prefix-down`, `mode initiate-and-terminate`, `packet-size`, `step-percent`, `test-iterator-duration`, `test-profile`, and `test-type` configuration statements introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

rfc2544-benchmarking

IN THIS SECTION

- Junos OS Syntax (except for SRX Series Firewalls) | **15760**
- Junos OS Syntax (SRX300 and SRX550HM) | **15761**
- Hierarchy Level | **15762**
- Description | **15762**

- Required Privilege Level | **15762**
- Release Information | **15762**

Junos OS Syntax (except for SRX Series Firewalls)

```
rfc2544-benchmarking {
        profiles {
            test-profile profile-name {
                test-type (throughput | latency | frame-loss | back-back-frames);
                packet-size bytes;
                step-percent percent;
                bandwidth-kbps kpbs;
            }
        }
        tests {
            test-name test-name {
                destination-ipv4-address address;
                destination-mac-address destination-mac-address;
                destination-udp-port  port-number;
                direction (egress | ingress);
                disable-signature-check;
                dscp-code-points dscp-code-points;
                family (bridge| inet | ccc | vpls);
                forwarding-class forwarding-class;
                halt-on-prefix-down;
                in-service;
                ip-swap;
                ivlan-cfi ivlan-cfi;
                ivlan-id ivlan-id;
                ivlan-priority ivlan-priority;
                mode (initiate-and-terminate | reflect);
                outer-tag-protocol-id outer-tag-protocol-id;
                ovlan-cfi ovlan-cfi;
                ovlan-id ovlan-id;
                ovlan-priority  ovlan-priority;
                packet-loss-priority (high | low | medium-high);
                receive-failure-threshold receive-failure-threshold;
                reflect-etype reflect-etype;
```

```
                reflect-mode (mac-rewrite | mac-swap | no-mac-swap);
                service-type (elan | eline);
                skip-arp-iteration;
                source-ipv4-address address;
                source-mac-address source-mac-address;
                source-udp-port port-number;
                test-finish-wait-duration test-finish-wait-duration;
                test-interface interface-name;
                test-iterator-duration test-iterator-duration;
                test-iterator-pass-threshold test-iterator-pass-threshold;
                test-profile test-profile;
                timestamp-format (microseconds | nanoseconds);
                transmit-failure-threshold transmit-failure-threshold;
                udp-tcp-port-swap;
                vlan-cfi vlan-cfi;
                vlan-id vlan-id;
                vlan-priority vlan-priority;
            }
        }
    }
```

Junos OS Syntax (SRX300 and SRX550HM)

```
rpm {
    rfc2544-benchmarking {
        tests {
            test-name test-name {
                destination-ipv4-address address;
                destination-udp-port  port-number;
                disable-signature-check;
                family inet
                mode reflect;
                source-ipv4-address address;
                source-udp-port port-number;
                test-interface interface-name;
            }
        }
    }
}
```

Hierarchy Level

```
[edit services rpm]
```

Description

Configure the parameters for the RFC 2544-based benchmarking test to examine and analyze the performance characteristics of a network interconnecting device. You must configure a test profile for the initiator (ACX routers only), which specifies the type of test and the manner in which it must be performed, and associate the test profile with a test name. The test name that you configure contains details, such as the address family and the test mode (`initiate-and-terminate`), for the test. You can associate the same test profile with multiple test names. For a reflector, you need only configure a test name and its associated details, and set the test mode to `reflect`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

RELATED DOCUMENTATION

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

show services rpm rfc2544-benchmarking | **25960**

show services rpm rfc2544-benchmarking test-id | **25969**

rfc6514-compliant-safi129 (Protocols BGP)

IN THIS SECTION

- Syntax | 15763
- Hierarchy Level | 15763
- Description | 15764
- Required Privilege Level | 15764
- Release Information | 15764

Syntax

```
rfc6514-compliant-safi129
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor neighbor-name],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor neighbor-
name]
```

Description

Parse and send BGP VPN multicast traffic according to Subsequent Address Family Identifier (SAFI) 129, as defined in RFC 6514 (that is, *length, prefix*). The Network Layer Reachability Information (NLRI) format used for BGP VPN multicast in previous releases of Junos OS was SAFI 128, which was *length, label, prefix*.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring BGP Neighbor Discovery Through RPM

rib (General)

IN THIS SECTION

- Syntax | 15765
- Hierarchy Level | 15766
- Description | 15766
- Default | 15766
- Options | 15766
- Required Privilege Level | 15767
- Release Information | 15767

Syntax

```
rib routing-table-name {
    aggregate {
        defaults {
            ... aggregate-options ...
        }
        route destination-prefix {
            policy policy-name;
            ... aggregate-options ...
        }
        generate {
            defaults {
                generate-options;
            }
            route destination-prefix {
                policy policy-name;
                generate-options;
            }
        }
        martians {
            destination-prefix match-type <allow>;
        }
    }
    policy-multipath;
    static {
        defaults {
            static-options;
        }
        rib-group group-name;
        route destination-prefix {
            next-hop;
            static-options;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Create a routing table.

Explicitly creating a routing table with ***routing-table-name*** is optional if you are not adding any static, martian, aggregate, or generated routes to the routing table and if you also are creating a routing table group.

> **NOTE**: The IPv4 multicast routing table (**inet.1**) and the IPv6 multicast routing table (**inet6.1**) are not supported for this statement.

Default

If you do not specify a routing table name with the ***routing-table-name*** option, the software uses the default routing tables, which are **inet.0** for unicast routes and **inet.1** for the multicast cache.

Options

routing-table-name—Name of the routing table, in the following format: ***protocol***[.***identifier***].

In a routing instance, the routing table name must include the routing instance name. For example, if the routing instance name is **link0**, the routing table name might be **link0.inet6.0**.

- ***protocol*** is the protocol family. It can be **inet6** for the IPv6 family, **inet** for the IPv4 family, **iso** for the ISO protocol family, or ***instance-name*.iso.0** for an ISO routing instance.
- ***identifier*** is a positive integer that specifies the instance of the routing table.
- **Default: inet.0**

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`policy-multipath` statement introduced in Junos Release 19.1R1 for all platfroms.

RELATED DOCUMENTATION

Example: Creating Routing Tables

passive (Routing Options) | **12225**

policy-multipath | **12694**

Policy-Based Multipath Routes Overview

rib-group (Protocols BGP)

IN THIS SECTION

- Syntax | **15767**
- Hierarchy Level | **15768**
- Description | **15768**
- Options | **15768**
- Required Privilege Level | **15768**
- Release Information | **15769**

Syntax

```
rib-group group-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family inet (labeled-unicast | unicast |
multicast)],
[edit logical-systems logical-system-name protocols bgp group group-name family inet (labeled-
unicast | unicast | multicast)],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
inet (labeled-unicast | unicast | multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family inet (labeled-unicast | unicast | multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family inet (labeled-unicast | unicast | multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family inet (labeled-unicast | unicast | multicast)],
[edit protocols bgp family inet (labeled-unicast | unicast | multicast)],
[edit protocols bgp group group-name family inet (labeled-unicast | unicast | multicast)],
[edit protocols bgp group group-name neighbor address family inet (labeled-unicast | unicast |
multicast)],
[edit routing-instances routing-instance-name protocols bgp family inet (labeled-unicast |
unicast | multicast)],
[edit routing-instances routing-instance-name protocols bgp group group-name family inet
(labeled-unicast | unicast | multicast)],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family inet (labeled-unicast | unicast | multicast)]
```

Description

Add unicast prefixes to unicast and multicast tables.

Options

group-name—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. You generally specify only one routing table group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table

Example: Populating a Routing Table Created by Virtual Router Configuration

Understanding Multiprotocol BGP

rib-group (Protocols DVMRP)

IN THIS SECTION

- Syntax | 15769
- Hierarchy Level | 15769
- Description | 15770
- Options | 15770
- Required Privilege Level | 15770
- Release Information | 15770

Syntax

```
rib-group group-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols dvmrp],
[edit protocols dvmrp]
```

Description

Associate a routing table group with DVMRP.

Options

group-name—Name of the routing table group. The name must be one that you defined with the `rib-groups` statement at the `[edit routing-options]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

> **NOTE**: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring DVMRP

rib-group (Protocols IS-IS)

IN THIS SECTION

- Syntax | **15771**
- Hierarchy Level | **15771**
- Description | **15771**

- Options | 15771
- Required Privilege Level | 15772
- Release Information | 15772

Syntax

```
rib-group {
    inet group-name;
    inet6 group-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Install routes learned from IS-IS routing instances into routing tables in the IS-IS routing table group. You can install IPv4 routes or IPv6 routes.

Support for IPv6 routing table groups in IS-IS enables IPv6 routes that are learned from IS-IS routing instances to be installed into other routing tables defined in an IS-IS routing table group.

Options

group-name—Name of the routing table group.

inet—Install IPv4 IS-IS routes.

inet6—Install IPv6 IS-IS routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table

Example: Populating a Routing Table Created by Virtual Router Configuration

Understanding Multiprotocol BGP

rib-group (Protocols MSDP)

IN THIS SECTION

- Syntax | 15772
- Hierarchy Level | 15773
- Description | 15773
- Options | 15773
- Required Privilege Level | 15773
- Release Information | 15773

Syntax

```
rib-group group-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
msdp],
[edit protocols msdp],
[edit routing-instances routing-instance-name protocols msdp]
```

Description

Associate a routing table group with MSDP.

Options

group-name—Name of the routing table group. The name must be one that you defined with the `rib-groups` statement at the `[edit routing-options]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Example: Configuring MSDP in a Routing Instance*

rib-group (Protocols OSPF)

IN THIS SECTION

- Syntax | 15774
- Hierarchy Level | 15774
- Description | 15774
- Options | 15775
- Required Privilege Level | 15775
- Release Information | 15775

Syntax

```
rib-group group-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Install routes learned from OSPF routing instances into routing tables in the OSPF routing table group.

Options

group-name—Name of the routing table group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

- Example: Exporting Specific Routes from One Routing Table Into Another Routing Table
- Example: Populating a Routing Table Created by Virtual Router Configuration
- Understanding Multiprotocol BGP
- interface-routes
- rib-group

rib-group (Protocols PIM)

IN THIS SECTION

- Syntax | 15776
- Hierarchy Level | 15776
- Description | 15776
- Options | 15776
- Required Privilege Level | 15776

Release Information | 15776

Syntax

```
rib-group {
    inet group-name;
    inet6 group-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Associate a routing table group with PIM.

Options

table-name—Name of the routing table. The name must be one that you defined with the `rib-groups` statement at the `[edit routing-options]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring a Dedicated PIM RPF Routing Table

rib-group (Protocols RIP)

IN THIS SECTION

- Syntax | 15777
- Hierarchy Level | 15777
- Description | 15777
- Options | 15777
- Required Privilege Level | 15778
- Release Information | 15778

Syntax

```
rib-group group-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit protocols rip],
[edit routing-instances routing-instance-name protocols rip]
```

Description

Install RIP routes into multiple routing tables by configuring a routing table group.

Options

group-name—Name of the routing table group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Redistributing Routes Between Two RIP Instances

rib-group (Routing Options)

IN THIS SECTION

- Syntax | 15778
- Hierarchy Level | 15779
- Description | 15779
- Options | 15779
- Required Privilege Level | 15779
- Release Information | 15779

Syntax

```
rib-group group-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options interface-routes],
[edit logical-systems logical-system-name routing-options interface-routes],
[edit logical-systems logical-system-name routing-options rib routing-table-name static],
[edit logical-systems logical-system-name routing-options static],
[edit routing-instances routing-instance-name routing-options interface-routes],
[edit routing-options interface-routes],
[edit routing-options rib routing-table-name static],
[edit routing-options static]
```

Description

Configure which routing table groups interface routes are imported into.

Options

group-name—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. It generally does not make sense to specify more than a single routing table group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Populating a Routing Table Created by Virtual Router Configuration

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table

interface-routes | **9010**

rib-groups | **15781**

rib-group (Source Packet Routing)

IN THIS SECTION

- Syntax | 15780
- Hierarchy Level | 15780
- Description | 15780
- Options | 15780
- Required Privilege Level | 15781
- Release Information | 15781

Syntax

```
rib-group {
    ipv4 rib-group-name;
    ipv4-color rib-group-name;
    ipv6 rib-group-name;
    ipv6-color rib-group-name;
    tag rib-group-name;
}
```

Hierarchy Level

```
[edit logical-systems name protocols source-packet-routing],
[edit protocols source-packet-routing]
```

Description

Enable rib-group import policies on segment routing traffic-engineering (SR-TE).

Options

rib-group-name Rib-group import policy.

`ipv4 | ipv6`	Import policy to be applied on IPv4 or IPv6 uncolored route, respectively.
`ipv4-color | ipv6-color`	Import policy to be applied on IPv4 or IPv6 colored route, respectively.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Static Segment Routing Label Switched Path

rib-groups

IN THIS SECTION

- Syntax | 15781
- Hierarchy Level | 15782
- Description | 15782
- Options | 15783
- Required Privilege Level | 15783
- Release Information | 15783

Syntax

```
rib-groups {
    group-name {
        export-rib group-name;
        import-policy [ policy-names ];
```

```
        import-rib [ group-names ];
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit routing-options]
```

Description

Group one or more routing tables to form a routing table group. A routing protocol can import routes into all the routing tables in the group and can export routes from a single routing table.

Each routing table group must contain one or more routing tables that Junos OS uses when importing routes (specified in the `import-rib` statement) and optionally can contain one routing table group that Junos OS uses when exporting routes to the routing protocols (specified in the `export-rib` statement).

The first routing table you specify is the *primary routing table*, and any additional routing tables are the *secondary routing tables*.

The primary routing table determines the address family of the routing table group. To configure an IP version 4 (IPv4) routing table group, specify **inet.0** as the primary routing table. To configure an IP version 6 (IPv6) routing table group, specify **inet6.0** as the primary routing table. If you configure an IPv6 routing table group, the primary and all secondary routing tables must be IPv6 routing tables (**inet6.*x***).

In Junos OS Release 9.5 and later, you can include both IPv4 and IPv6 routing tables in an IPv4 import routing table group using the `import-rib` statement. In releases prior to Junos OS Release 9.5, you can only include either IPv4 or IPv6 routing tables in the same `import-rib` statement. The ability to configure an import routing table group with both IPv4 and IPv6 routing tables enables you, for example, to populate the **inet6.3** routing table with IPv6 addresses that are compatible with IPv4. Specify **inet.0** as the primary routing table, and specify **inet6.3** as a secondary routing table.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

NOTE: If you configure an import routing table group that includes both IPv4 and IPv6 routing tables, any corresponding export routing table group must include only IPv4 routing tables.

If you have configured a routing table, configure the OSPF primary instance at the `[edit protocols ospf]` hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group. For more information, see *Example: Configuring Multiple Routing Instances of OSPF*.

After specifying the routing table from which to import routes, you can apply one or more policies to control which routes are installed in the routing table group. To apply a policy to routes being imported into the routing table group, include the `import-policy` statement.

Options

group-name—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table

rib-group (Routing Options) | **15778**

rib-list

IN THIS SECTION

- Syntax | **15784**
- Hierarchy Level | **15784**

- Description | 15784
- Options | 15784
- Required Privilege Level | 15784
- Release Information | 15785

Syntax

```
rib-list name {
rib-list-item name;

}
```

Hierarchy Level

```
[edit logical-systems name policy-options],
[edit policy-options]
```

Description

Define a set of RIBs and assign a name to the list. You can even use wildcards as a list item.

Options

name Assign a name to the RIB list.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4 R1

RELATED DOCUMENTATION

loc-rib | 9897

rib (Multitopology Routing)

IN THIS SECTION

- Syntax | 15785
- Hierarchy Level | 15786
- Description | 15786
- Options | 15786
- Required Privilege Level | 15786
- Release Information | 15786

Syntax

```
rib routing-table-name  {
    static {
        route destination-prefix {
            next-hop;
        }
        static-options;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Configure a static route to install routes in the routing table for a specific topology.

Options

routing-table-name—Name of the routing table for a topology. Use the following format: ***logical-system-name/ routing-instance-name:topology-name.protocol.identifier***. Include the routing instance string only if the instance is not the primary. The logical system string is included only if the logical system identifier has a value other than 0 (zero). Each routing table for a topology includes a colon (:) before the topology name. ***protocol*** is the protocol family, which can be **inet** or **inet6**. ***identifier*** is the positive integer that specifies the instance of the routing table. For example, to install IPv6 routes to the routing table for a topology named voice in the primary instance, include **:voice.inet6.0**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement support for multitopology routing introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Removing the Class E Prefix on Martian Addresses

Example: Configuring IPv6 BGP Routes over IPv4 Transport

Enabling Layer 2 VPN and VPLS Signaling

Understanding Martian Addresses

static (Routing Options) | **28570**

rib (Protocols BGP)

IN THIS SECTION

- Syntax | **15787**
- Hierarchy Level | **15787**
- Description | **15788**
- Options | **15788**
- Required Privilege Level | **15788**
- Release Information | **15788**

Syntax

```
rib (inet.3 | inet6.3) ;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family (inet | inet6) labeled-unicast],
[edit logical-systems logical-system-name protocols bgp group group-name family (inet | inet6)
labeled-unicast],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address
family (inet | inet6) labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family (inet | inet6) labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family (inet | inet6) labeled-unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family (inet | inet6) labeled-unicast],
[edit protocols bgp family (inet | inet6) labeled-unicast],
```

```
[edit protocols bgp group group-name family (inet | inet6) labeled-unicast],
[edit protocols bgp group group-name neighbor address family (inet | inet6) labeled-unicast],
[edit routing-instances routing-instance-name protocols bgp family (inet | inet6) labeled-
unicast],
[edit routing-instances routing-instance-name protocols bgp group group-name family (inet |
inet6) labeled-unicast],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
(inet | inet6) labeled-unicast]
```

Description

You can allow both labeled and unlabeled routes to be exchanged in a single session. The labeled routes are placed in the inet.3 or inet6.3 routing table, and both labeled and unlabeled unicast routes can be sent or received by the router.

Options

inet.3—Name of the routing table for IPv4.

inet6.3—Name of the routing table for IPv6.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Example: Configuring IPv6 BGP Routes over IPv4 Transport*
>
> *Enabling Layer 2 VPN and VPLS Signaling*
>
> *Understanding Multiprotocol BGP*

rib (Route Resolution)

IN THIS SECTION

- Syntax | 15789
- Hierarchy Level | 15789
- Description | 15789
- Required Privilege Level | 15789
- Release Information | 15790

Syntax

```
rib routing-table-name {
    import [ policy-names ];
    resolution-ribs [ routing-table-names ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options resolution],
[edit logical-systems logical-system-name routing-options resolution],
[edit routing-instances routing-instance-name routing-options resolution],
[edit routing-options resolution]
```

Description

Specify a routing table name for route resolution.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring Route Resolution on PE Routers

rib-service (programmable-rpd)

IN THIS SECTION

- Syntax | 15790
- Hierarchy Level | 15790
- Description | 15791
- Options | 15791
- Required Privilege Level | 15791
- Release Information | 15791

Syntax

```
rib-service {
    dynamic-next-hop-interface (enable | disable);
}
```

Hierarchy Level

```
[edit routing-options programmable-rpd]
```

Description

Configure programmable routing protocol process (prpd) options that apply only to the RIB Service API.

Options

dynamic-next-hop-interface (enable | disable)

Disable or enable dynamic next-hop interface binding. When enabled, programmed RIB routes react to Up, Down, Add, and Delete events for direct next-hop interfaces. When all next-hop interfaces are unconfigured or down, the route is hidden and becomes inactive. When a next-hop interface is configured or comes up, the route becomes visible and active. This prevents dropped traffic and keeps inactive routes from being propagated through the network.

This feature applies to all routes programmed using the rib_service JET API where an interface is configured as a direct next-hop, including interfaces that are part of a flexible tunnel. It also applies to tunnels configured with the flexible_tunnel_service JET API. Indirect next-hops are resolved by the RPD resolver normally.

Changes to the configuration of this statement only affect routes programmed using Junos OS Release 20.2R1 or later.

- **Default:** This option is enabled by default starting in Junos OS Release 20.2R1. In earlier releases, dynamic next-hop interface binding is disabled by default.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Understanding Programmable Flexible VXLAN Tunnels

show programmable-rpd clients | **23287**

traceoptions (programmable-rpd) | **29651**

rib-sharding

IN THIS SECTION

- Syntax | 15792
- Hierarchy Level | 15792
- Description | 15792
- Options | 15793
- Required Privilege Level | 15793
- Release Information | 15793

Syntax

```
rib-sharding {
    number-of-shards;
}
```

Hierarchy Level

```
[edit system processes routing bgp]
```

Description

Enable BGP RIB sharding. BGP RIB sharding is to split the BGP process across routes. Different routes are hashed into different threads to achieve concurrency. BGP RIB sharding splits a unified BGP RIB into several sub RIBs and each sub RIB handles a subset of BGP routes. Each sub RIB is served by a separate RPD thread to achieve concurrency. BGP RIB sharding is disabled by default. To enable, you need to explicitly configure it on a multicore routing engine. This feature is supported only on 64-bit routing protocol process (rpd) where the Routing Engine has more than one core. We recommend configuring this feature on a device with atleast 4CPU cores and 16GB of memory.

NOTE:

- RPD would restart automatically when rib-sharding or update-threading configuration is changed.
- Sharding requires Update/IO thread. If `rib-sharding` is configured and `update-threading` is not configured, the commit check fails.
- BGP RIB sharding and update IO features are supported for the following IPv4 and IPv6 address families:
  - IPv4 VPN Unicast
  - IPv4 VPN Multicast (vrf.inet.2)
  - IPv6 VPN Unicast
  - IPv6 VPN Multicast (vrf.inet.2)
  - IPv4 Labeled Unicast
  - IPv6 Labeled Unicast
- All the other RIBs are still processed without sharding.

Options

`number of shards`—the number of sharding threads created. If you configure `rib-sharding` on a routing engine, RPD creates sharding threads. By default, it is the same as the number of CPU cores on the routing engine. Optionally, you can specify the number-of-shards you want to create.

- **Range:** 1-31 shards

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

Support for NSR with sharding introduced in Junos OS Release 22.2.

BGP sharding supports IPv4, IPv6, L3VPN and BGP-LU from Junos OS Release 20.4R1.

rift

IN THIS SECTION

- Syntax | 15794
- Hierarchy Level | 15795
- Description | 15795
- Options | 15795
- Required Privilege Level | 15797
- Release Information | 15797

Syntax

```
rift {
    capabilities;
    default-prefixes;
    default-prefixes-advertisement;
    export;
    external-preference external-preference;
    interface interface-name;
    level;
    lie-receive-address;
    lie-receive-port port-number;
    name name;
    node-id (node-id | auto);
    overload timeout seconds;
    preference route-preference;
    proxy-process traceoptions (file | flag | no-remote-trace);
    tie-authentication (loose | none | permissive | strict);
    tie-origination-key tie-origination-key;
    traceoptions (file | flag | level | peer-prefixes);
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure the Routing in Fat Trees (RIFT) protocol for IP abrics that have CLOS-based and fat tree model topologies.

The RIFT protocol enables:

- Automatic disaggregation of prefixes on link and node failures.
- Minimal storing of routing state information at every level.
- Zero-configuration capabilities with automatic pruning.
- Load balancing of traffic towards the spine based on available bandwidth.
- Synchronization of a limited key-value data-store that can be used after protocol convergence.

Options

capabilities

Enable the RIFT capability of reducing flooding globally. See "capabilities" on page 3197 for more information.

deafult-prefixes

Explicitly specify default prefixes generated in southbound direction. See "default-prefixes" on page 5206 for more information.

default-prefixes-advertisement (always | never)

(Optional) Enable default route generation strategy.

- always—Always originate default prefixes southbound.
- never—Never originate default prefix southbound.

export

Enable exporting of routes from other protocols into the RIFT protocol. See "export" on page 6685 for more information.

external-preference ***external-preference***

—(Optional) External route preference for the RIFT protocol. The external preference configured is analogous to the OSPF configuration. This value is shown when RIFT installs the preference as static route values.

- **Range:** 1 through 256

interface

(Optional) Configure the interfaces for the RIFT protocol. See "interface" on page 8994 for more information.

level

(Optional) Configure the level to identify the role of node in the RIFT topology. See "level" on page 9734 for more information.

lie-receive-address

(Optional) Configure the address on which the link information elements (LIEs) should be received. See "lie-receive-address" on page 9742 for more information.

lie-receive-port ***port-number***

(Optional) Port on which link information elements (LIES) should be received.

- **Range:** 512 through 65535

name ***name***

(Optional) Name of the node for identification.

node-id (*node-id* | auto)

(Optional) ID of the configured node, or enable zero touch provisioning where the node ID is configured automatically.

overload

(Optional) Configure the overload bit that takes the node out of traffic bearing paths.

- Values:

  timeout seconds (Optional) Specify in seconds the time after which the overload bit is reset. The overload timeout is similar to IS-IS timeout.

  - **Range:** 10 through 1800 seconds

preference ***route-preference***

(Optional) Route preference for the RIFT protocol.

- **Range:** 15 through 256

proxy-process

(Optional) Configure the proxy process options for the RIFT protocol.

tie-authentication (loose | none | permissive | strict)

(Optional) Configure the method to authenticate received TIEs (inner fingerprint).

- Values:

  loose Verify authentication only if present, that is, when the key ID is not 0.

  none Disable authentication checking completely.

permissive Accept authentication if key identifier is unknown.

strict Accept authentication only if a key is present and it is valid.

tie-origination-key *tie-origination-key* (Optional) Configure the key ID used to protect self-originated TIEs (inner key).

- **Range:** 1 through 16777215

traceoptions (Optional) Enable traceoptions for the RIFT protocol.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

ring-id

IN THIS SECTION

- Syntax | 15798
- Hierarchy Level | 15798
- Description | 15798
- Options | 15798
- Required Privilege Level | 15798
- Release Information | 15798

Syntax

```
ring-id number;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Specify the ring ID.

Options

number—Ring ID number.

- **Range:** 1 through 239
- **Default:** 1

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

> *Ethernet Ring Protection Switching Overview*
>
> Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

ring-protection-link-end

IN THIS SECTION

- Syntax | 15799
- Hierarchy Level | 15799
- Description | 15799
- Required Privilege Level | 15799
- Release Information | 15799

Syntax

```
ring-protection-link-end;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name (east-interface | west-interface)]
```

Description

Specify that the port is one side of a ring protection link (RPL) by setting the RPL end flag.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Example: Configuring Ethernet Ring Protection Switching on EX Series Switches

Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS

Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

ring-protection-link-owner

IN THIS SECTION

- Syntax | 15800
- Hierarchy Level | 15800
- Description | 15800
- Required Privilege Level | 15801
- Release Information | 15801

Syntax

```
ring-protection-link-owner;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Specify the ring protection link (RPL) owner flag in the Ethernet protection ring. Include this statement only once for each ring (only one node can function as the RPL owner).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Example: Configuring Ethernet Ring Protection Switching on EX Series Switches

Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS

riot-loopback

IN THIS SECTION

- Syntax | 15801
- Hierarchy Level | 15802
- Description | 15802
- Required Privilege Level | 15802
- Release Information | 15802

Syntax

```
riot-loopback;
```

Hierarchy Level

```
[edit vlans name vxlan]
```

Description

Assign a VLAN that enables routing in and out of VXLAN tunnels (RIOT) for symmetric EVPN Type 2 and Type 5 routing.

Some Juniper switching platforms don't have native VXLAN routing hardware support for Layer 3 VXLAN gateway functions. Instead, supporting switches can use an intermediary loopback port in a two-pass process for routing in and out of VXLAN tunnels (RIOT). You configure this intermediary port as a loopback link aggregation group (LAG) that is a member of all VLANs with IRB interfaces used for VXLAN routing.

To support symmetric EVPN Type 2 and Type 5 routes with the RIOT loopback LAG port, for each Layer 3 virtual routing and forwarding (VRF) instance, you:

- Configure an extra VLAN with an IRB interface.
- Map that VLAN to the same VXLAN network identifier (VNI) that you configure for a corresponding Layer 3 VRF instance.
- Configure this `riot-loopback` statement with the extra VLAN so the Layer 3 VRF instance uses the RIOT loopback LAG port for Type 2 and Type 5 VXLAN routing.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced for EVPN Type 5 routing in Junos OS Release 21.3R1.

Support added for symmetric EVPN Type 2 routing in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Using a RIOT Loopback Port to Route Traffic in an EVPN Network

Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics

rip

IN THIS SECTION

- Syntax | 15803
- Hierarchy Level | 15803
- Description | 15803
- Default | 15803
- Required Privilege Level | 15803
- Release Information | 15804

Syntax

```
rip {...}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable RIP routing on the routing device.

Default

RIP is disabled on the routing device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring RIP

ripng

IN THIS SECTION

- Syntax | 15804
- Hierarchy Level | 15804
- Description | 15805
- Default | 15805
- Required Privilege Level | 15805
- Release Information | 15805

Syntax

```
ripng {...}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Enable RIPng routing on the routing device.

Default

RIPng is disabled on the routing device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

rising-event-index

IN THIS SECTION

- Syntax | 15805
- Hierarchy Level | 15806
- Description | 15806
- Options | 15806
- Required Privilege Level | 15806
- Release Information | 15806

Syntax

```
rising-event-index index;
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Set the index of the event entry that is used when a rising alarm threshold is exceeded. The rising-event index is specified when you configure an SNMP RMON alarm. If this value is zero, no event is triggered.

Options

index—Index of the event entry that is used when a rising threshold is crossed.

- **Range:** 0 through 65,535
- **Default:** 0

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Remote Nework Monitoring (RMON)

falling-event-index | 6808

rising-threshold (Health Monitor)

IN THIS SECTION

- Syntax | 15807
- Hierarchy Level | 15807
- Description | 15807
- Options | 15807
- Required Privilege Level | 15808
- Release Information | 15808

Syntax

```
rising-threshold percentage;
```

Hierarchy Level

```
[edit snmp health-monitor],
[edit snmp health-monitor idp]
```

Description

Set the upper threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored object, you can be alerted whenever the value of the variable falls outside the allowable operational range.

Options

percentage—The lower threshold for the alarm entry.

- **Range:** `1` through `100`
- **Default:** `80` percent of the maximum possible value

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

falling-threshold (percentage) | 6815

Configure Health Monitoring on Devices Running Junos OS

rising-threshold (RMON)

IN THIS SECTION

- Syntax | 15808
- Hierarchy Level | 15809
- Description | 15809
- Options | 15809
- Required Privilege Level | 15809
- Release Information | 15809

Syntax

```
rising-threshold integer;
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Set the upper threshold for the sampled variable (monitored object) when you configure an SNMP RMON alarm. By setting a rising and a falling threshold for a variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.

Options

integer—The lower threshold for the alarm entry.

- **Range:** –2,147,483,648 through 2,147,483,647

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Remote Nework Monitoring (RMON)

falling-threshold (integer) | **6812**

rmon

IN THIS SECTION

- Syntax | 15810
- Hierarchy Level | 15811
- Description | 15811
- Default | 15811
- Required Privilege Level | 15811
- Release Information | 15811

Syntax

```
rmon {
    alarm index {
        description description;
        falling-event-index index;
        falling-threshold integer;
        No Link Title seconds;
        interval seconds;
        No Link Title;
        No Link Title index;
        rising-threshold integer;
        No Link Title (absolute-value | delta-value);
        No Link Title (falling-alarm | rising-alarm | rising-or-falling alarm);
        syslog-subtag syslog-subtag;
        variable oid-variable;
    }
    No Link Title index {
        community community-name;
        description description;
        type (RMON Notification) type;
    }
    history history-index {
        bucket-size number;
        interface interface-name;
        interval seconds;
```

```
        owner (RMON) owner-name;
    }
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Provide comprehensive network fault diagnosis, planning, and performance tuning information. RMON delivers this information in nine groups of monitoring elements, each providing specific sets of data to meet common network monitoring requirements. Each group is optional, so that vendors do not need to support all the groups within the MIB.

Junos OS supports the RMON statistics, history, alarm, and event groups.

The remaining statements are explained separately. See CLI Explorer.

Default

Disabled.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

RMON MIB Event, Alarm, Log, and History Control Tables

RMON Overview

Junos OS Network Management Configuration Guide

rng

IN THIS SECTION

- Syntax | 15812
- Hierarchy Level | 15812
- Description | 15812
- Default | 15813
- Options | 15813
- Required Privilege Level | 15813

Syntax

```
rng [dyce | fortuna | hmac-drbg];
```

Hierarchy Level

```
[edit system]
```

Description

Configure a system cryptographically secure pseudo-random number generator (CSPRNG). By default, rng is not configured.

NOTE:

- Only root users can configure rng options.
- A device reboot is required for the change to take effect.
- In FIPS mode, HMAC-DRBG will be the only active RNG available.

Default

Federal Information Processing Standards (FIPs) and Non-FIPs mode behavior:

- (SRX1500, SRX4100, SRX4200, and SRX4600 Devices) FIPS Mode uses the hash-based message authentication code-deterministic random bit generator (HMAC-DRBG) , while Non-FIPS mode uses the Junos default Fortuna Random Number Generator.
- (General Routing) Starting with Junos OS Release 21.4R1, in non-FIPS mode, the RNG in FreeBSD 12 based Junos OS versions has been changed from the default FreeBSD Fortuna RNG to the FIPS/SP800-90A&B HMAC-DRBG CSPRNG.

Options

fortuna Configure the fortuna CSPRNG.

hmac-drbg Configure the HMAC DRBG CSPRNG.

dyce Configure the Da Yan Concentrator and Extender of Random Data (DYCE) CSPRNG.

Required Privilege Level

system-To view this statement in the configuration.

roaming-mode

IN THIS SECTION

- Syntax | 15814
- Hierarchy Level | 15814
- Description | 15814
- Options | 15814
- Required Privilege Level | 15814
- Release Information | 15814

Syntax

```
roaming-mode (home-only | automatic)
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options]
```

Description

Specify whether the 3G wireless modem interface can access networks other than the home network.

Options

- `home-only`—No roaming is allowed.
- `automatic`—Allows access to networks other than the home network. This is the default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

robust-count (IGMP Snooping)

IN THIS SECTION

- Syntax | 15815
- Hierarchy Level | 15815
- Description | 15815
- Options | 15816
- Required Privilege Level | 15816
- Release Information | 15816

Syntax

```
robust-count number;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-name],
[edit protocols igmp-snooping vlan vlan-name]
```

Description

Configure the number of queries a device sends before removing a multicast group from the multicast forwarding table. We recommend that the robust count be set to the same value on all multicast routers and switches in the VLAN.

This option provides fine-tuning to allow for expected packet loss on a subnet. You can wait more intervals if subnet packet loss is high and IGMP report messages might be lost.

Use the `query-interval`, `query-last-member-interval`, or `query-response-interval` statements in the same hierarchy to configure interval lengths.

Options

number—Number of intervals the switch waits before timing out a multicast group.

- **Range:** 2 through 10
- **Default:** 2

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Example: Configuring IGMP Snooping*
>
> *Example: Configuring IGMP Snooping on SRX Series Devices*

robust-count (MLD Snooping)

IN THIS SECTION

- Syntax | **15817**
- Hierarchy Level | **15817**
- Description | **15817**
- Default | **15817**
- Options | **15817**
- Required Privilege Level | **15817**

Release Information | 15818

Syntax

```
robust-count number;
```

Hierarchy Level

```
[edit protocols mld-snooping vlan (all | vlan-name)]
```

```
[edit routing-instances instance-name protocols mld-snooping vlan vlan-name]
```

Description

Configure the number of queries the switch sends before removing a multicast group from the multicast forwarding table. We recommend that the robust count be set to the same value on all multicast routers and switches in the VLAN.

Default

The default is the value of the `robust-count` statement configured for MLD. The default for the MLD `robust-count` statement is 2.

Options

number—Number of queries the switch sends before timing out a multicast group.

- **Range:** 2 through 10

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

Support at the `[edit routing-instances instance-name protocols mld-snooping vlan vlan-name]` hierarchy level introduced in Junos OS Release 13.3 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring MLD Snooping on SRX Series Devices

mld-snooping | **10783**

Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)

Understanding MLD Snooping

robust-count (Protocols IGMP)

IN THIS SECTION

- Syntax | **15818**
- Hierarchy Level | **15819**
- Description | **15819**
- Options | **15819**
- Required Privilege Level | **15819**
- Release Information | **15819**

Syntax

```
robust-count number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp],
[edit protocols igmp]
```

Description

Tune the expected packet loss on a subnet. This factor is used to calculate the group member interval, other querier present interval, and last-member query count.

Options

number—Robustness variable.

- **Range:** 2 through 10
- **Default:** 2

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Modifying the IGMP Robustness Variable

robust-count (Protocols IGMP AMT)

IN THIS SECTION

- Syntax | 15820
- Hierarchy Level | 15820
- Description | 15820
- Options | 15820
- Required Privilege Level | 15821
- Release Information | 15821

Syntax

```
robust-count number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp amt relay defaults],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols igmp
amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances routing-instance-name protocols igmp amt relay defaults]
```

Description

Configure the expected IGMP packet loss on an Automatic Multicast Tunneling (AMT) tunnel. If a tunnel is expected to have packet loss, increase the robust count.

Options

number—Number of packets that can be lost before the AMT protocol deletes the multicast state.

- **Range:** 2 through 10

- **Default:** 2

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Default IGMP Parameters for AMT Interfaces

robust-count (Protocols MLD)

IN THIS SECTION

- Syntax | 15821
- Hierarchy Level | 15822
- Description | 15822
- Options | 15822
- Required Privilege Level | 15822
- Release Information | 15822

Syntax

```
robust-count number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld],
[edit protocols mld]
```

Description

Tune for the expected packet loss on a subnet.

Options

number—Time interval. This interval must be less than the interval between general host-query messages.

- **Range:** 2 through 10
- **Default:** 2 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Example: Modifying the MLD Robustness Variable*

robustness-count

IN THIS SECTION

- Syntax | 15823
- Hierarchy Level | 15823
- Description | 15823
- Options | 15824
- Required Privilege Level | 15824
- Release Information | 15824

Syntax

```
robustness-count number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM)  interface-name
bidirectional df-election],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM)  interface-name bidirectional df-election],
[edit protocols pim interface (Protocols PIM)  interface-name bidirectional df-election],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-
name bidirectional df-election]
```

Description

Configure the designated forwarder (DF) election robustness count for bidirectional PIM. When a DF election Offer or Winner message fails to be received, the message is retransmitted. The **robustness-count** statement sets the minimum number of DF election messages that must fail to be received for DF election to fail. To prevent routing loops, all routers on the link must have a consistent view of the DF. When the DF election fails because DF election messages are not received, forwarding on bidirectional PIM routes is suspended.

If a router receives from a neighbor a better offer than its own, the router stops participating in the election for a period of **robustness-count** * **offer-period**. Eventually, all routers except the best candidate stop sending Offer messages.

Options

number—Number of transmission attempts for DF election messages.

- **Range:** 1 through 10
- **Default:** 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Bidirectional PIM

Example: Configuring Bidirectional PIM

role

IN THIS SECTION

- Syntax | **15825**
- Hierarchy Level | **15825**
- Description | **15825**
- Options | **15826**
- Required Privilege Level | **15827**

Release Information | 15827

Syntax

```
role (line-card | routing-engine);
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

Specify the roles of the members of the Virtual Chassis or a Virtual Chassis Fabric (VCF) in a preprovisioned Virtual Chassis. For a mixed Virtual Chassis or VCF, see Understanding Mixed EX Series and QFX Series Virtual Chassis or *Understanding Mixed Virtual Chassis Fabric* for any recommendations or requirements for assigning the Routing Engine role based on the types of switches comprising the Virtual Chassis or VCF.

Virtual Chassis Fabric

Specify the role to be performed by each switch. In a VCF, two spine devices are configured into the Routing Engine role and the remaining spine devices and all leaf devices are configured into the `line-card` role. The role must also be associated with the member's serial number (see "serial-number" on page 16599).

EX Series (except EX8200 Virtual Chassis) and QFX Series Virtual Chassis

Specify the role to be performed by each member switch. Associate the role with the member's serial number (see "serial-number" on page 16599).

When you use a preprovisioned configuration, you cannot modify the primary-role priority or member ID of member switches through the user interfaces. The primary-role priority value is generated by the software, based on the assigned role:

- A member configured as `routing-engine` is assigned the primary-role priority 129.
- A member configured as `line-card` is assigned the primary-role priority 0.

- A member listed in the preprovisioned configuration without an explicitly specified role is assigned the `line-card` role and primary-role priority 0 by default.

The configured role specifications are permanent. If both `routing-engine` members fail, a `line-card` member cannot take over as primary of the Virtual Chassis configuration. You must delete the preprovisioned configuration to change the specified roles in a Virtual Chassis.

Explicitly configure two members as `routing-engine` and configure additional switches as members of the preprovisioned Virtual Chassis by specifying only their serial numbers. If you do not explicitly configure the role of the additional members, they function in a `line-card` role by default. In that case, a member that is functioning in a `line-card` role can take over primary role if the members functioning as primary and backup (`routing-engine` role) both fail.

EX8200 Virtual Chassis

Specify the role to be performed by each XRE200 External Routing Engine and each EX8200 member switch. Associate the role with the member's serial number (see "serial-number" on page 16599). An EX8200 Virtual Chassis cannot function when both external Routing Engines, which must be configured in the `routing-engine` role, have failed.

Options

- `line-card`—Enables the member to be eligible to function only in the linecard role. Any member of the Virtual Chassis or VCF configuration other than the primary or backup functions in the linecard role and runs only a subset of Junos OS for EX Series switches. A member functioning in the linecard role does not run the control protocols or the chassis management processes.

  A Virtual Chassis must have at least three members for one member to function in the linecard role.

  In an EX8200 Virtual Chassis configuration, all member switches must be in the linecard role.

- `routing-engine`—Enables the member to function as a primary or backup Routing Engine of the Virtual Chassis or VCF configuration. The primary manages all members and runs the chassis management processes and control protocols. The backup synchronizes with the primary in terms of protocol states, forwarding tables, and so forth, so that it is prepared to preserve routing information and maintain network connectivity without disruption in case the primary is unavailable.

  In a Virtual Chassis composed of EX Series switches (except EX8200 switches) or QFX Series switches, specify two and only two members in the `routing-engine` role. The software determines which of the two members assigned the `routing-engine` role functions as primary, based on the primary-role election algorithm. See Understanding How the Primary in a Virtual Chassis Is Elected. In these Virtual Chassis, the `routing-engine` role is associated with a switch. For a mixed Virtual Chassis or VCF, see Understanding Mixed EX Series and QFX Series Virtual Chassis or *Understanding Mixed Virtual Chassis Fabric* for specific recommendations or requirements for assigning the Routing Engine

role based on the types of switches comprising the Virtual Chassis or VCF. The remaining switches are configured into the linecard role.

In an EX8200 Virtual Chassis, all XRE200 External Routing Engines must be in the `routing-engine` role.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis
- Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis
- Configuring an EX9200 Virtual Chassis
- Configuring an EX4650 or a QFX Series Virtual Chassis
- *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*

role (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | **15828**
- Hierarchy Level | **15828**
- Description | **15828**
- Options | **15828**
- Required Privilege Level | **15828**
- Release Information | **15829**

Syntax

```
role (routing-engine | line-card);
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

As part of the preprovisioned configuration for an MX Series Virtual Chassis, assign the role to be performed by each member router in the Virtual Chassis. The preprovisioned configuration permanently associates the member ID and role with the member router's chassis serial number.

Options

routing-engine—Enable the member router to function as the primary router or backup router of the Virtual Chassis configuration. The primary router maintains the global configuration and state information for both members of the Virtual Chassis, and runs the chassis management processes. The backup router synchronizes with the primary router and relays chassis control information, such as line-card presence and alarms, to the primary router. If the primary router is unavailable, the backup router takes primary role of the Virtual Chassis to preserve routing information and maintain network connectivity without disruption. You must assign the **routing-engine** role to both members of the Virtual Chassis. When the Virtual Chassis is formed, the software runs a primary-role election algorithm to determine which of the two member routers functions as the primary router and which functions as the backup router of the Virtual Chassis.

line-card—Explicitly configuring a member router with the **line-card** role is *not supported* in the current release. However, when split detection is enabled (the default behavior for a Virtual Chassis) and either the Virtual Chassis ports go down or the backup router fails, the primary router takes a **line-card** role. The **line-card** role effectively removes the former primary router from the Virtual Chassis configuration until connectivity is restored.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Preprovisioned Member Information for a Virtual Chassis

Example: Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis

Virtual Chassis Components Overview

Global Roles and Local Roles in a Virtual Chassis

Split Detection Behavior in a Virtual Chassis

role (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 15829
- Hierarchy Level | 15829
- Description | 15830
- Options | 15830
- Required Privilege Level | 15830
- Release Information | 15830

Syntax

```
role equal;
```

Hierarchy Level

```
[edit protocols openflow switch switch-name controller]
```

Description

Specify the role of each OpenFlow controller when configuring more than one controller for a virtual switch. A single controller configuration automatically puts the controller in active mode. In active mode, the virtual switch automatically initiates a connection to the controller.

Options

equal Configure the controller as the active controller in a single controller configuration.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Support for OpenFlow on Devices Running Junos OS

controller (Protocols OpenFlow) | 4654

rollback

IN THIS SECTION

- Syntax | 15831
- Description | 15831
- Options | 15831
- Required Privilege Level | 15831
- Release Information | 15831

Syntax

```
rollback <number | rescue | revision revision-string>
```

Description

Return to a previously committed configuration. The software saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the `commit` configuration command.

The currently operational configuration is stored in the file **juniper.conf**, and the last three committed configurations are stored in the files **juniper.conf.1**, **juniper.conf.2**, and **juniper.conf.3**. These four files are located in the directory **/config**, which is on the router's flash drive. The remaining 46 previous committed configurations, the files **juniper.conf.4** through **juniper.conf.49**, are stored in the directory **/var/db/config**, which is on the router's hard disk.

During rollback, the configuration you specify is loaded from the associated file. Only objects in the rollback configuration that differ from the previously loaded configuration are marked as changed (equivalent to `load update`).

Options

none	(Optional) Return to the most recently saved configuration.
number	(Optional) Configuration to return to. The range of values is from 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49. The default is 0.
rescue	(Optional) Return to the rescue configuration.
revision ***revision-string***	(Option) Use a configuration revision identifier to rollback to a specific configuration. Use the "show system commit include-configuration-revision" on page 26754 command to view the configuration revision identifier for each revision.

Required Privilege Level

rollback—To roll back to configurations other than the one most recently committed.

Release Information

Command introduced before Junos OS Release 7.4.

Option `revision` introduced in Junos OS Release 20.4R1 and Junos OS Evolved Release 20.4R1.

root-authentication

IN THIS SECTION

- Syntax | 15832
- Hierarchy Level | 15832
- Description | 15833
- Options | 15833
- Required Privilege Level | 15833
- Release Information | 15834

Syntax

```
root-authentication {
    encrypted-password "password";
    no-public-keys
    ssh-ecdsa name {
        from from;
    }
    ssh-ed25519 name {
        from from;
    }
    ssh-rsa name {
        from from;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the authentication methods for the root-level user, whose username is `root`.

You can use the `ssh-ecdsa`, `ssh-ed25519`, or `ssh-rsa` statements to directly configure SSH ECDSA, ED25519, or RSA keys to authenticate root logins. You can configure more than one public key for SSH authentication of root logins as well as for user accounts. When a user logs in as root, the public keys are referenced to determine whether the private key matches any of them.

> **NOTE**: There is a difference in the way a plain-text password is hashed and stored as encrypted in Junos OS release 15.1 and newer vs. that used in release 12.3 and earlier.
>
> In Junos OS release 12.3 and earlier images, encrypting and decrypting of passwords using SHA-256 is not supported. This is a limitation in Junos OS release 12.3. Therefore, if a password that is encrypted by using SHA-256 in Junos OS release 15.1 is provided to Junos OS releases 12.3 and earlier, the password decryption will fail.
>
> See this KB article for instructions on how to first configure a compatible MD5 hashed root password before you downgrade. If you downgrade without first configuring a compatible root password you will not be able to login as root after the downgrade.

Options

encrypted-password "*password*"	Specify the MD5 or other password. You can specify only one encrypted password. You cannot configure a blank password using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.
no-public-keys	Disable SSH public key-based authentication.
ssh-ecdsa *name* from *from*	Use an SSH ECDSA public key. You can specify one or more public keys.
ssh-ed25519 *name* from *from*	Use an SSH ED25519 public key. You can specify one or more public keys.
ssh-rsa *name* from *from*	Use an SSH RSA public key. You can specify one or more public keys.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

class (Defining Login Classes) | 3347

user (Access) | 30547

root-ca (Services)

IN THIS SECTION

- Syntax | 15834
- Hierarchy Level | 15834
- Description | 15835
- Options | 15835
- Required Privilege Level | 15835
- Release Information | 15835

Syntax

```
root-ca root-certificate;
```

Hierarchy Level

```
[edit services ssl proxy profile profile-name]
[edit services ssl termination profile profile-name]
```

Description

Root certificate for interdicting server certificates in proxy mode. This statement is supported on the SRX1500, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall.

Options

root-ca-name—Specify root certificate for interdicting server certificates in proxy mode.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

Firewall User Authentication Overview

root-logical-system

IN THIS SECTION

- Syntax | 15836
- Hierarchy Level | 15836
- Description | 15836
- Options | 15836
- Required Privilege Level | 15836
- Release Information | 15836

Syntax

```
root-logical-system;
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify root-logical-system to bind the security profile to the primary logical system.

The primary administrator uses security profiles to provision logical systems with resources. The security profile containing this statement must be bound to root-logical-system only.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

none

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways

root-login (JDM)

IN THIS SECTION

- Syntax | 15837
- Hierarchy Level | 15837
- Description | 15837
- Default | 15837
- Options | 15837
- Required Privilege Level | 15838
- Release Information | 15838

Syntax

```
root-login (allow | deny );
```

Hierarchy Level

```
[edit system services ssh]
```

Description

Control user access to the JDM through SSH.

Default

`root-login allow` is the default.

Options

`allow`—Allow users to log in to the JDM as root through SSH.

`deny`—Disable users from logging in to the JDM as root through SSH. This configuration option is applicable only to the JDM management interface (jmgmt0). Setting this configuration option does not block the internal JDM to JDM communication, which uses root account with password-less authentication method.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3.

RELATED DOCUMENTATION

Setting Up Junos Node Slicing

root-streaming

IN THIS SECTION

- Syntax | 15838
- Hierarchy Level | 15839
- Description | 15839
- Required Privilege Level | 15840
- Release Information | 15840

Syntax

```
root-streaming;
```

Hierarchy Level

```
[edit security log],
[edit logical-systems name security log],
[edit tenants name security log]
```

Description

The `root-streaming` option allows the user logical systems and tenant systems to generate the logs using the root system's stream configuration. If you configure `root-streaming` option under user logical systems and tenant systems in stream mode, all the logs of this user logical system and tenant systems will be sent out according to root logical and tenant system's configuration.

The `root-streaming` option works together with the stream mode. If user logical system is in event mode, the `root-streaming` option will not take effect. If user logical system is in stream mode and root logical systems is in event mode, the `root-streaming` option will try to get the streams in root but the logs will not be sent out. Table 1 provides a detailed behavior of user logical system logs.

Table 136: root-streaming option for root and logical systems in different mode

Logical systems mode	Root logical systems mode	Status of the root-streaming option for logical systems	Behavior of logical systems logs
Event mode		Enabled	Send to Routing Engine
Stream mode	Event mode	Enabled	Cannot send out
	Stream mode	Enabled	Send to root log server
	Stream-event mode	Enabled	Send to root log server
Stream-event mode	Event mode	Enabled	Send to Routing Engine
	Stream mode	Enabled	Send to Routing Engine and root log server

Table 136: root-streaming option for root and logical systems in different mode *(Continued)*

Logical systems mode	Root logical systems mode	Status of the root-streaming option for logical systems	Behavior of logical systems logs
	Stream-event mode	Enabled	Send to Routing Engine and root log server
Event mode		Not Enabled	Send to Routing Engine
Stream mode		Not Enabled	Send to logical systems log server
Stream-event mode		Not Enabled	Send to Routing Engine and logical systems log server

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

root-address

IN THIS SECTION

- Syntax | 15841
- Hierarchy Level | 15841
- Description | 15841

- Release Information | 15841

Syntax

```
root-address root-address;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name provider-tunnel ldp-p2mp],
[edit logical-systems name tenants name routing-instances name provider-tunnel ldp-p2mp],
[edit routing-instances name provider-tunnel ldp-p2mp],
[edit tenants name routing-instances name provider-tunnel ldp-p2mp]
```

Description

Configure the non-default loopback address for LDP point to multi-point provider tunnels.

NOTE: A maximum of 9000 ingress replication provider tunnels are supported with 3 root addresses.

A maximum of 3 VRFs can be mapped to 3 root addresses.

Release Information

Statement introduced in Junos OS Release 24.4R1.

route (Access)

IN THIS SECTION

- Syntax | 15842

- Hierarchy Level | 15842
- Description | 15842
- Options | 15842
- Required Privilege Level | 15843
- Release Information | 15843

Syntax

```
route prefix {
    metric route-cost;
    next-hop next-hop;
    preference route-distance;
    qualified-next-hop
    tag route-tag;
    tag2 route-tag2;
}
```

Hierarchy Level

```
[edit routing-options access]
[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options
access],
[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options
rib routing-table-name access],
[edit dynamic-profiles profile-name routing-options access],
[edit dynamic-profiles profile-name routing-options rib routing-table-name access]
```

Description

Dynamically configure the parameters for access routes.

Options

prefix—Either the specific route prefix that you want to assign to the access route or one of the following route prefix variables.

- For IPv4 access routes, use the variable, `$junos-framed-route-ip-address-prefix`. The route prefix variable is dynamically replaced with the value in Framed-Route RADIUS attribute [22].
- For IPv6 access routes, use the variable, `$junos-framed-route-ipv6-address-prefix`. The variable is dynamically replaced with the value in Framed-IPv6-Route RADIUS attribute [99]. When you use this variable, you must configure it at either of the following hierarchy levels:
  - `[edit dynamic-profiles` *`profile-name`* `routing-instances $junos-routing-instance routing-options rib $junos-ipv6-rib access]`
  - `[edit dynamic-profiles` *`profile-name`* `routing-options rib $junos-ipv6-rib access]`

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles` *`profile-name`* `routing-instances $junos-routing-instance routing-options access]` and `[edit dynamic-profiles` *`profile-name`* `routing-instances $junos-routing-instance routing-options rib` *`routing-table-name`* `access]` hierarchy levels introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Routes for DHCP and PPP Subscriber Access Networks

route (Access Internal)

IN THIS SECTION

- Syntax | 15844
- Hierarchy Level | 15844

- Description | 15844
- Options | 15844
- Required Privilege Level | 15845
- Release Information | 15845

Syntax

```
route subscriber-ip-address {
    next-hop next-hop;
    qualified-next-hop underlying-interface {
        mac-address address;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
access-internal],
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
rib routing-table-name access-internal],
[edit dynamic-profiles profile-name routing-options access-internal]
```

Description

Dynamically configure parameters for an access-internal route.

Options

subscriber-ip-address—Either the specific IP address you want to assign to the access-internal route or the subscriber IP address variable ($junos-subscriber-ip-address). The subscriber IP address variable is dynamically replaced with the value supplied by DHCP or PPP when a subscriber logs in.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options access-internal]` and `[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options rib routing-table-name access-internal]` hierarchy levels introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Routes for DHCP and PPP Subscriber Access Networks

route (Access-Internal)

IN THIS SECTION

- Syntax | 15845
- Hierarchy Level | 15846
- Description | 15846
- Options | 15846
- Required Privilege Level | 15846
- Release Information | 15846

Syntax

```
route  ip-prefix</prefix-length> {
    next-hop next-hop;
```

```
    qualified-next-hop next-hop;
}
```

Hierarchy Level

```
[edit routing-options access-internal]
```

Description

Configure the parameters for internal access routes.

Options

ip-prefix</prefix-length>—Specific route prefix that you want to assign to the internal access route.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

route-accounting

IN THIS SECTION

- Syntax | 15847
- Hierarchy Level | 15847
- Description | 15847
- Default | 15848

Required Privilege Level | 15848
Release Information | 15848

Syntax

```
route-accounting;
```

Hierarchy Level

```
[edit forwarding-options family inet6]
```

Description

Enables the collection of transit IPv6 statistics at both the physical interface and logical interface levels on third-generation FPCs (FPC3-PTX-U2 and FPC3-PTX-U3 on PTX5000 and FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1 on PTX3000), PTX1000, and PTX10008. Enable this statement prior to the route being created.

Egress accounting for IPv6 traffic is not performed for cases where MPLS packets arrive on a TCC interface and then egress the router as IPv6 packets.

PTX Series routers with first or second-generation FPCs do not display IPv6statistics for physical interfaces or logical interfaces. Transit statistics on child links in aggregated Ethernet interfaces are not included.

To view the IPv6 statistics for a given physical or logical interface, run the operational command, `show interfaces extensive interface | find IPv6` or `show interfaces statistics` *`interface-name`* `detail`. Note that the output displays packet and byte counts for transit traffic only. Locally generated packets are not included in the metrics.

NOTE:

- Enabling this option soon after disabling it on a third-generation FPCs (FPC3-PTX-U2 and FPC3-PTX-U3 on PTX5000 and FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1 on

PTX3000), PTX1000, and PTX10008, restarts the accounting from zero and does not retain the statistics that was previously accounted.

- On third-generation FPCs in PTX series, egress accounting for IPV6 traffic is not performed for cases where MPLS packets arrives on TCC interface and egress out of the router as IPV6 packets.

Default

Disabled

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring IPv6 Accounting

route-active-on

IN THIS SECTION

- Syntax | 15849
- Hierarchy Level | 15849
- Description | 15849
- Options | 15849
- Required Privilege Level | 15849
- Release Information | 15849

Syntax

```
route-active-on (node0 | node1);
```

Hierarchy Level

```
[edit policy-options condition condition-name]
```

Description

In chassis cluster configurations, the incoming traffic from the core network is sent to the interface that exists on the current active node . For a route advertised by BGP, a policy is applied on BGP configuration before exporting routes. An additional term in the policy match condition determines the current active device (node) before making the routing decision. In this way, the traffic is processed by the active node.

Options

node0	Route active on node 0
node1	Route active on node 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

Example: Configuring Conditional Route Advertising in a Chassis Cluster

route-attributes

IN THIS SECTION

- Syntax | 15850
- Hierarchy Level | 15850
- Description | 15851
- Options | 15851
- Required Privilege Level | 15851

Syntax

```
route-attributes {
 as-path {
  export-action (allow | skip);
  import-action (allow | skip);
 }
 community {
  export-action (allow | skip);
  import-action (allow | skip);
 }
 preference {
  export-action (allow | skip);
  import-action (allow | skip);
 }
}
```

Hierarchy Level

```
[edit routing-instances name protocols evpn ip-prefix-routes]
```

Description

You can use the route-attributes import and export options to control which attributes to carry between the IP and EVPN routes. Use the export option when generating EVPN routes from IP routes and the import option when generating IP routes from EVPN routes.

Options

as-path

export-action

- allow — Inherit as-path from IP routes to EVPN routes.
- skip — Do not inherit as-path from IP routes to EVPN routes.

import-action

- allow — Inherit as-path from EVPN routes to IP routes.
- skip — Do not inherit as-path from EVPN routes to IP routes.

community

export-action

- allow — Inherit community from EVPN routes to IP routes.
- skip — Do not inherit community from EVPN routes to IP routes.

import-action

- allow — Inherit community from EVPN routes to IP routes.
- skip — Do not inherit community from EVPN routes to IP routes.

preference

export-action

- allow — Inherit preference from EVPN routes to IP routes.
- skip — Do not inherit preference from EVPN routes to IP routes.

import-action

- allow — Inherit preference from EVPN routes to IP routes.
- skip — Do not inherit preference from EVPN routes to IP routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

ip-prefix-routes | 9250

route-change-timeout

IN THIS SECTION

- Syntax | 15852
- Hierarchy Level | 15852
- Description | 15852
- Options | 15853
- Required Privilege Level | 15853
- Release Information | 15853

Syntax

```
route-change-timeout seconds;
```

Hierarchy Level

```
[edit security flow]
```

Description

Specify the session timeout when a session is rerouted but there is a reroute failure (for example, the new route uses a different egress zone from the previous route).

Options

seconds —Amount of time before sessions are timed out.

- **Range:** 6 through 1800 seconds
- **Default:** 30 seconds

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

route (chassis)

IN THIS SECTION

- Syntax | 15854
- Hierarchy Level | 15854
- Description | 15854
- Required Privilege Level | 15854
- Release Information | 15854

Syntax

```
route;
```

Hierarchy Level

```
[edit chassis memory-enhanced]
```

Description

Allocate more jtree memory for routing tables over firewall filters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Allocating More Memory for Routing Tables, Firewall Filters, and Layer 3 VPN Labels

route (Diameter Network Element)

IN THIS SECTION

- Syntax | 15855
- Hierarchy Level | 15855
- Description | 15855

- Options | 15855
- Required Privilege Level | 15855
- Release Information | 15856

Syntax

```
route dne-route-name {
    destination realm realm-name  <host hostname>;
    function function-name <partition partition-name>;
    metric route-metric;
}
```

Hierarchy Level

```
[edit diameter network-element element-name forwarding]
```

Description

Define a route reachable through the Diameter network element by associating a metric with a combination of destination and function partition.

Options

dne-route-name—Route name defined for the Diameter network element.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Network Elements

route-distinguisher

IN THIS SECTION

- Syntax | 15856
- Hierarchy Level | 15856
- Description | 15857
- Options | 15858
- Required Privilege Level | 15859
- Release Information | 15859

Syntax

```
route-distinguisher (as-number:id | ip-address:id);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn mesh-group mesh-group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name],
[edit protocols evpn interconnect]
```

```
[edit routing-instances routing-instance-name],
[edit routing-instances routing-instance-name protocols evpn interconnect]
[edit routing-instances routing-instance-name protocols l2vpn mesh-group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]
```

Description

Specify an identifier attached to a route that distinguishes to which VPN or virtual private LAN service (VPLS) the route belongs. Each routing instance must have a unique route distinguisher (RD) associated with it. The RD places bounds around a VPN so the device can use the same IP address prefixes in different VPNs without having the addresses overlap. You must configure the route-distinguisher statement for instances with instance type `vrf`.

Use the following guidelines when you assign RDs:

- For Layer 2 (L2) VPNs and VPLS, if you configure the `l2vpn-use-bgp-rules` statement, you must configure a unique RD for each PE router participating in the routing instance.

  If you configure mesh groups, the RD in each mesh group must also be unique.

- For Ethernet VPNs (EVPNs), you must configure a unique RD for each provider edge (PE) device participating in the routing instance to ensure that the prefixes generated by different PEs are unique.

- For other VPNs besides L2 VPNs, VPLS, and EVPNs, we recommend that you use a unique RD for each PE device participating in a particular routing instance. You can alternatively use the same RD on all PE devices for the same VPN routing instance, but if you use a unique RD, you can determine the customer edge (CE) router from which a route originated within the VPN.

- On EVPN data center interconnect (DCI) gateway devices, if you configure an interconnect RD at the `[edit routing-instances name protocols evpn interconnect]` hierarchy, the interconnect RD must be different from the local RD in the routing instance.

> **NOTE**: When you configure DCI with seamless stitching for EVPN Type 2 routes, the device throws a commit error if you try to configure the same value for the interconnect RD and the local RD.
>
> To enforce this condition for DCI seamless stitching with EVPN Type 5 routes as well, you also see a commit error with Junos OS and Junos OS Evolved Releases starting in 22.4R2 and 23.1R1.

CAUTION: We strongly recommend that if you change an RD that you configured and committed previously, or change the routing instance type from `virtual-router` to `vrf`, make either of those changes during a maintenance window, as follows:

1. Deactivate the routing instance.
2. Change the RD.
3. Activate the routing instance.

Options

as-number:*number*—*as-number* is an assigned AS number, and *number* is any 2-byte or 4-byte value. The AS number can be from 1 through 4,294,967,295. If the AS number is a 2-byte value, the administrative number is a 4-byte value. If the AS number is 4-byte value, the administrative number is a 2-byte value. An RD consisting of a 4-byte AS number and a 2-byte administrative number is defined as a type 2 RD in RFC 4364 *BGP/MPLS IP VPNs*.

NOTE: In Junos OS Release 9.1 and later, the numeric range for AS numbers is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. All releases of Junos OS support 2-byte AS numbers. To configure an RD that includes a 4-byte AS number, append the letter "L" to the end of the AS number. For example, an RD with the 4-byte AS number 7,765,000 and an administrative number of 1,000 is represented as `77765000L:1000`.

In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS dot notation format of two integer values joined by a period: <*16-bit high-order value in decimal*>.<*16-bit low-order value in decimal*>. For example, the 4-byte AS number of 65,546 in the plain-number format is represented as 1.10 in AS dot notation format.

number:*id*—Number and identifer expressed in one of these formats: *16-bit number*.*32-bit identifier* or *32-bit number*.*16-bit identifier*.

ip-address:*id*—IP address (*ip-address* is a 4-byte value) within your assigned prefix range and a 2-byte value for the *id*. The IP address can be any globally unique unicast address.

- **Range:** 0 through 4,294,967,295 (2^{32} – 1). If the router you are configuring is a BGP peer of a router that does not support 4-byte AS numbers, you need to configure a local AS number. For more information, see *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.

> **NOTE**: For Ethernet VPNs (EVPNs), an RD that includes zero as the *id* value is reserved for the default EVPN routing instance by default. Because you can't assign the same RD for two routing instances, the device throws a commit error if you use an RD of the form *ip-address*:*id* with *id* value zero for another routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at `[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]` hierarchy level introduced in Junos OS Release 11.2.

Support at `[edit routing-instances routing-instance-name protocols l2vpn mesh-group mesh-group-name]` hierarchy level introduced in Junos OS Release 13.2.

Support at the following hierarchy levels introduced in Junos OS Release 20.3R1 on QFX Series switches: `[edit protocols evpn interconnect]` and `[edit routing-instances routing-instance-name protocols evpn interconnect]`.

RELATED DOCUMENTATION

- *Example: Configuring BGP Route Target Filtering for VPNs*
- *Example: Configuring FEC 129 BGP Autodiscovery for VPWS*
- *Configuring EVPN Routing Instances*
- *Configuring Routing Instances on PE Routers in VPNs*
- Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)
- Configuring an MPLS-Based Layer 3 VPN (CLI Procedure)
- path-selection | **12266**

route-distinguisher-id

IN THIS SECTION

- Syntax | 15860
- Hierarchy Level | 15860
- Description | 15860
- Options | 15861
- Required Privilege Level | 15861
- Release Information | 15861

Syntax

```
route-distinguisher-id ip-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit routing-options]
```

Description

The automatic Route-Distinguisher feature uses the `route-distinguisher-id` command to create a Type 1 Route-Distinguisher to the routing instance, based on RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs).

The 4-octet `Administrator` subfield of the Route-Distinguisher uses the IP address that is configured as the `route-distinguisher-id`, which is common for all routing instances in the local system. Additionally, the system assigns an autogenerated 2-octet number as the `Assigned Number` subfield that is unique per the routing instances.

If you configure the `route-distinguisher` statement in addition to the `route-distinguisher-id` statement, the value configured for `route-distinguisher` supersedes the value generated from `route-distinguisher-id`.

> **NOTE**: To avoid a conflict between the configured and auto-generated route distinguisher values, you must ensure that the first half (IP address) of the `route-distinguisher` obtained by configuring the `route-distinguisher` statement is different from the IP-address configured in the `route-distinguisher-id` statement.

Options

`ip-address`—Address for routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring BGP Route Target Filtering for VPNs

Configuring Routing Instances on PE Routers in VPNs

route-filter

IN THIS SECTION

- Syntax | 15862
- Hierarchy Level | 15862
- Description | 15862
- Default | 15862
- Required Privilege Level | 15862

Release Information | 15862

Syntax

```
route-filter (no-walkup | walkup);
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options defaults],
[edit logical-systems logical-system-name policy-options policy-statement policy-statement-name
defaults],
[edit policy-options defaults],
[edit policy-options policy-statement policy-statement-name defaults]
```

Description

Enable or disable walkup globally or locally for route filters in a particular policy statement or globally. The walkup feature examines more than the longest match route filters in a policy statement term with more than one route filter, allowing consolidation of terms and a potential performance enhancement.

Default

By default, no route filter walkup is performed.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

no-walkup | 11623

walkup | 31019

Walkup for Route Filters Overview

Configuring Walkup for Route Filters to Improve Operational Efficiency

Example: Configuring Walkup for Route Filters Globally to Improve Operational Efficiency

Example: Configuring Walkup for Route Filters Locally to Improve Operational Efficiency

route-filter-list

IN THIS SECTION

- route-filter-list | 15863
- Hierarchy Level | 15864
- route-filter-list (usage) | 15864
- Hierarchy Level | 15864
- Description | 15864
- Options | 15867
- Required Privilege Level | 15867
- Release Information | 15867

route-filter-list

```
route-filter-list route-filter-list-name {
ip-addresses <exact | longer | orlonger | prefix-length-range | through | upto> label range
start : end;
ip-addresses <exact | longer | orlonger | prefix-length-range | through | upto> label-allocation-
fallback-reject;
ip-addresses exact label value;
ip-addresses exact label-allocation-fallback-reject;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options],
[edit policy-options]
```

route-filter-list (usage)

```
route-filter-list route-filter-list-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-statement policy-statement-name term term-name
from],
[edit policy-options policy-statement policy-statement-name term term-name from]
```

Description

The route filter list is a user-configured list of individual route filters that you create at the `[edit policy-options]` hierarchy level. Each item in the list consists of a complete route filter statement, made up of a destination prefix, a match type, and an optional action. You can re-use the list in different policies, adding whatever qualifiers you need, rather than having to re-create it for every different case, as shown in the examples below.

```
[edit]
user@router# show policy-options route-filter-list rf-test-list
203.0.113.0/24 address-mask 255.255.255.0;
192.0.2.0/26 orlonger reject;
198.51.100.8/29 exact accept;
```

Once configured, the `route-filter-list` is used by referencing its *`route-filter-list-name`* in a `policy-statement` at the `[edit policy-options policy-statement` *`policy-statement-name`* `term` *`term-name`* `from]` hierarchy level. Route filter lists can be used in conjunction with other `route-filter` statements.

For example:

```
[edit]
user@router# show policy-options policy-statement test-route-filter-list-statement
from {
    route-filter 198.51.100.32/29 exact accept;
    route-filter 192.0.2.1/32 exact;
    route-filter-list rf-test-list;
}
then reject;
```

If one action is associated with the `route-filter-list` entry and another action is associated with the list-level entry, the action associated with the list-level entry takes precedence.

```
policy-options {
    route-filter-list RFL-1 {
        198.1.1.0/24 exact;
        198.2.1.0/24 orlonger;
    }
    route-filter-list RFL-2 {
        198.1.1.0/24 exact reject;
        198.2.1.0/24 orlonger accept;
    }
    policy-statement pol {
        term 1 {
            from {
                route-filter-list RFL-1 longer accept;
            }
        }
        term 2 {
            from {
                route-filter-list RFL-2 longer reject;
            }
        }
    }
}
```

Here is an example on how to specify policy-based control for BGP-LU labels being allocated for a given prefix:

```
[edit]
user@host# show policy-options route-filter-list rfl-1{
198.51.1.1/32 exact {
    label 1000101;
    next-hop 198.52.1.2;
    accept;
    label-allocation-fallback-reject;
}
198.51.1.0/24 prefix-length-range /24-/32 {
    label range 1000000:1000200;
    next-hop 198.53.2.2;
    accept;
}
}
```

```
[edit]
user@host# show policy-options policy-statement p-1{
from {
    route-filter-list rfl-1;
}
then accept;
}
```

```
[edit]
user@host# show protocols bgp group ibgp{
type internal;
local-address 198.0.23.3;
family inet {
    labeled-unicast {
        per-prefix-label;
        }
    }
neighbor 198.0.12.1 {
    export p-1;
}
}
```

Options

When specifying a match prefix, you can specify an exact match with a particular route or a less precise match. You can configure either a common action that applies to the entire list or an action associated with each prefix.

- *list-name*—Name of route-filter-list of routes to match.
- *exact*—Exactly match the prefix length.
- *longer*—Mask is greater than the prefix length.
- *orlonger*—Mask is greater than or equal to the prefix length.
- *prefix-length-range*—Mask falls between two prefix lengths.
- *upto*—Mask falls between two prefix length.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support on PTX1000, PTX10000, and PTX10003 Packet Transport Routers introduced in Junos OS Release 20.2R1 for the following list-level qualifiers: `exact`, `longer`, `orlonger`, `prefix-length-range`, and `upto`.

RELATED DOCUMENTATION

Understanding Route Filters for Use in Routing Policy Match Conditions

Example: Configuring Routing Policy Prefix Lists

show policy | **23199**

route (Interfaces)

IN THIS SECTION

- Syntax | 15868
- Hierarchy Level | 15868
- Description | 15868
- Options | 15868
- Required Privilege Level | 15869
- Release Information | 15869

Syntax

```
route prefix routing-instance instance-name priority-cost priority;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track]
```

Description

Enable route tracking for a Virtual Router Redundancy Protocol (VRRP) group.

Options

prefix—Route to be tracked for this VRRP group.

priority-cost ***priority***—The value subtracted from the configured VRRP priority when the tracked interface or route is down, forcing a new primary router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.

routing-instance ***instance-name***—Routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, the value for ***instance-name*** must be **default**.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring a Route to Be Tracked for a VRRP Group

route-localization

IN THIS SECTION

- Syntax | **15870**
- Hierarchy Level | **15870**
- Description | **15870**
- Options | **15870**
- Required Privilege Level | **15870**
- Release Information | **15870**

Syntax

```
route-localization {
    inet;
    inet6;
}
```

Hierarchy Level

[edit chassis]

Description

Configure FIB localization for IPv4 and IPv6 routes.

Options

inet Configure FIB localization for IPv4 routes.

inet6 Configure FIB localization for IPv6 routes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring Packet Forwarding Engine FIB Localization

route-monitoring

IN THIS SECTION

- Syntax | 15871
- Hierarchy Level | 15871
- Description | 15872
- Options | 15872
- Required Privilege Level | 15874
- Release Information | 15874

Syntax

```
route-monitoring {
    none;
    loc-rib;
    post-policy {
        exclude-non-eligible;
    }
    pre-policy {
        exclude-non-feasible;
    }
    rib-out {
        post-policy;
        pre-policy;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp bmp],
[edit logical-systems logical-system-name protocols bgp group group-name bmp],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address bmp],
[edit logical-systems logical-system-name routing-options bmp],
[edit logical-systems logical-system-name routing-options bmp station station-name],
```

```
[edit protocols bgp bmp],
[edit protocols bgp group group-name bmp],
[edit protocols bgp group group-name neighbor address bmp],
[edit routing-options bmp],
[edit routing-options bmp station station-name]
```

Description

Specify whether BMP should send pre-policy route monitoring messages, post-policy route monitoring messages, both types of messages, or none at all. The pre-policy can be configured to exclude routes that are non-feasible for the decision process (for example, a route loop). The post-policy can be configured to exclude routes that are not eligible for the decision process (for example, protocol next-hop not resolved).

You can also selectively enable or disable BMP route monitoring at various hierarchy levels (for example, `[edit protocols bgp group group-name]` or `[edit protocols bgp group group-name neighbor address]`).

> **NOTE**: If you have initially configured pre-policy route monitoring, but later delete the initial configuration and configure post-policy route monitoring, then the previously advertised pre-policy routes are withdrawn. Conversely, if you have initially configured post-policy route monitoring and later modify it to pre-policy route monitoring, then the previously advertised post-policy routes are withdrawn.

Options

none Explicitly disables BMP route monitoring.

- **Default:** If you configure the `route-monitoring` statement at the `[edit routing-options bmp]` hierarchy level, the default option is `pre-policy`.

> **NOTE**: If `post-policy` or `pre-policy` is not configured explicitly, then the default mode `pre-policy` would be applicable.
>
> The `route-monitoring` options `pre-policy`, `post-policy`, `rib-out pre-policy`, and `rib-out post-policy` are not mutually exclusive. You can configure any combination, or none, or all. If none of the options are configured, `pre-policy` remains as the default.

If you configure the `route-monitoring` statement at any of the `[edit protocols bgp]` hierarchy levels, the default option is to inherit the configuration from the `route-monitoring` statement configured at the `[edit routing-options bmp]` hierarchy level.

loc-rib Send local-rib route monitoring messages.

post-policy For BMP route monitoring, allows you to exclude routes that are non-eligible for the decision process (for example, protocol next-hop not resolved). This represents the view of the BGP routes after running the import policy. If the import policy has rejected the BGP route, the route does not exist in the post policy view.

> **NOTE**: If post-policy is configured, the device applies the import policies on the routes received so that BMP can display the routes as follows:
>
> - If an import policy modifies any parameters in the route, the route is displayed by BMP with the new values.
> - If an import policy modifies the next-hop and the next-hop is unresolved by the local device, the route is displayed by BMP as reachable.
> - If an import policy rejects a route, it is displayed as unreachable in the BMP update.
> - If a route is received with an unreachable next-hop, but the import policy doesn't reject the route, then BMP displays the route as reachable.

- **Values:** exclude-non-eligible—Exclude routes that are non-eligible from the decision process.

  Using post-policy with the exclude-non-eligible option causes an unreachable BMP update in the following cases:

  - If an import policy modifies the next-hop and the next-hop is unresolved by the local router, the route is displayed by BMP as unreachable.
  - if a route is received with a next-hop that is unreachable and the import policy does not reject this route, it is displayed by BMP as unreachable.

pre-policy Excludes routes that are non-feasible from the BMP route monitoring decision process (for example, a route loop). This represents the view of the BGP routes before running the import policy.

- **Values:** exclude-non-feasible—Exclude routes that are non-feasible from the decision process.

rib-out Send adj-ribs-out route monitoring messages.

- Values:
  - post-policy—Send post-policy adj-ribs-out route monitoring messages.
  - pre-policy—Send pre-policy adj-ribs-out route monitoring messages.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D15.

`rib-out` option introduced in Junos OS Release 19.1R1.

`loc-rib` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Configuring BGP Monitoring Protocol Version 3

route-record

IN THIS SECTION

- Syntax | 15875
- Hierarchy Level | 15875
- Description | 15875
- Required Privilege Level | 15875
- Release Information | 15875

Syntax

```
route-record;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit routing-options]
```

Description

Export the AS path and routing information to the traffic sampling process.

Before you can perform flow aggregation, the routing protocol process must export the AS path and routing information to the sampling process.

> **NOTE**: Starting with Junos OS Release 15.1, when you commit a minor configuration change, the routing protocol process sends only AS paths that are active routes to the FPCs. Not all known AS paths are sent to the FPC, thereby considerably reducing the memory and CPU usage, resulting in a faster route record database update.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Enabling Flow Aggregation*
>
> Junos OS Services Interfaces Library for Routing Devices

route-refresh-priority

IN THIS SECTION

- Syntax | 15876
- Hierarchy Level | 15876
- Description | 15876
- Required Privilege Level | 15877
- Release Information | 15877

Syntax

```
route-refresh-priority (expedited | priority priority-queue-number (1-16));
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family family-name sub-family],
[edit logical-systems logical-system-name protocols bgp group group-name family family-name sub-
family],
[edit protocols bgp family family-name sub-family],
[edit protocols bgp group group-name family family-name sub-family],
[edit protocols bgp group group-name neighbor neighbor-id family family-name]
```

Description

Within BGP route prioritization, the `route-refresh-priority` statement allows you to set specific priority levels for BGP routes that are to be refreshed. The `route-refresh-priority` statement can be configured for BGP neighbors during BGP configuration, or for sub-families within the following address families:

- `evpn`
- `inet`
- `inet-mdt`

- inet-mvpn
- inet-vpn
- inet6
- inet6-mvpn
- inet6-vpn
- iso-vpn
- l2vpn
- route-target
- traffic-engineering

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

output-queue-priority | 12008

withdraw-priority | 31134

Understanding BGP Route Prioritization

route-suppression (DHCP Local Server and Relay Agent)

IN THIS SECTION

Syntax | 15878

- Hierarchy Level | 15878
- Description | 15878
- Options | 15879
- Required Privilege Level | 15879
- Release Information | 15879

Syntax

```
route-suppression (access | access-internal | destination);
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name]
```

Description

Configure the jdhcpd process to suppress the installation of access, access-internal, or destination routes during client binding.

NOTE: You cannot suppress access-internal routes when the subscriber is configured with both IA_NA and IA_PD addresses over IP demux interfaces—the IA_PD route relies on the IA_NA route for next hop connectivity.

Options

access (DHCPv6 only) Suppress installation of access routes. You can use the `access` and `access-internal` options in the same statement for DHCPv6.

access-internal In a DHCPv4 hierarchy, suppress installation of both access-internal and destination routes. In a DHCPv6 hierarchy, suppress access-internal routes only. Can be configured in the same statement with the `access` option.

destination (DHCPv4 only) Suppress installation of destination routes. This option and the `access-internal` option are mutually exclusive; however, the `access-internal` option also suppresses destination routes.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default

route-target-filter

IN THIS SECTION

- Syntax | 15880
- Hierarchy Level | 15880
- Description | 15880
- Options | 15880

Required Privilege Level | 15881
Release Information | 15881

Syntax

```
route-target-filter destination {
    group group-name;
    local;
    neighbor address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options rib bgp.rtarget.0 static],
[edit routing-options rib bgp.rtarget.0 static]
```

Description

Statically configure route target filtering. Route target filtering allows you to distribute VPN routes to only the routers that need them. In VPN networks without route target filtering configured, BGP distributes all static VPN routes to all of the VPN peer routers. You can add static routes to the bgp.rtarget.0 routing table with specific NLRI-imposed constraints.

Options

destination Allows you to specify the static route destination. This value must be in the format x:y/len. The x value is either an IP address or an AS number followed by an optional L to indicate a 4 byte AS number, and y is a number (for example, 123456L:100/64).

group *group-name(s)* Installs an RT-Constrain filter for the destination for all peers in the specified BGP group. The route and corresponding BGP group are displayed in the output of the `show bgp group rtf detail` command.

local Causes the router to originate the route target constrain NLRI, but does not install any filtering state for the prefix. This behavior can be useful when the router should always

receive VPN routes with this route-target regardless of the state of a given BGP peering session or group status. The route is not displayed in the output of the `show bgp group rtf detail` command unless it is also included in either the BGP neighbor or BGP group configuration.

neighbor *address*

Installs an RT-Constrain filter for the destination for this BGP neighbor. The route is displayed in the output of the `show bgp group rtf detail` command. Specify the BGP neighbor using the router's IP address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Static Route Target Filtering for VPNs

Reducing Network Resource Use with Static Route Target Filtering for VPNs

route-target (Protocols BGP)

IN THIS SECTION

- Syntax | **15882**
- Hierarchy Level | **15882**
- Description | **15883**
- Options | **15883**
- Required Privilege Level | **15883**
- Release Information | **15883**

Syntax

```
route-target {
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | time-in-minutes)>;
        drop-excess <percentage>;
        hide-excess <percentage>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | time-in-minutes)>;
        drop-excess <percentage>;
        hide-excess <percentage>;
   }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family],
[edit logical-systems logical-system-name protocols bgp group group-name family],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family],
[edit protocols bgp family],
[edit protocols bgp group group-name family],
[edit protocols bgp family prefix-limit (inet | inet6)(any | flow | labeled-unicast | multicast |
unicast)],
[edit protocols bgp group group-name neighbor address family],
[edit routing-instances routing-instance-name protocols bgp group group-name family],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family]
```

Description

Limit the number of prefixes advertised on BGP peers specifically to the peers that need the updates.

Options

advertise-default—Advertise default routes and suppress more specific routes.

external-paths *number*—Number of external paths accepted for route filtering.

- **Range:** 1 through 256 paths
- **Default:** 1 path

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs

Example: Configuring Proxy BGP Route Target Filtering for VPNs

route-target (Protocols MVPN)

IN THIS SECTION

Syntax | 15884

Hierarchy Level | 15884

- Description | 15885
- Default | 15885
- Options | 15885
- Required Privilege Level | 15885
- Release Information | 15885

Syntax

```
route-target {
    export-target {
        target target-community;
        unicast;
    }
    import-target {
        target {
            target-value;
            receiver target-value;
            sender target-value;
        }
        unicast {
            receiver;
            sender;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
mvpn],
[edit routing-instances routing-instance-name protocols mvpn]
```

Description

Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN NLRI.

Default

The multicast VPN routing instance uses the import and export route targets configured for the Layer 3 VPN.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring VRF Route Targets for Routing Instances for an MBGP MVPN

route-timeout (Protocols RIP)

IN THIS SECTION

- Syntax | 15886
- Hierarchy Level | 15886
- Description | 15886
- Options | 15886

Required Privilege Level | 15887
Release Information | 15887

Syntax

```
route-timeout seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit logical-systems logical-system-name protocols rip group group-name],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip],
[edit protocols rip group group-name],
[edit protocols rip group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip],
[edit routing-instances routing-instance-name protocols rip group group-name],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

Configure the route timeout interval for RIP. If a route is not refreshed after being installed in the routing table by the specified timeout interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.

Options

seconds—Estimated time to wait before making updates to the routing table.

- **Range:** 30 through 360 seconds

- **Default:** 180 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Example: Configuring RIP Timers

RIP Demand Circuits Overview

route-timeout (Protocols RIPng)

IN THIS SECTION

- Syntax | **15887**
- Hierarchy Level | **15888**
- Description | **15888**
- Options | **15888**
- Required Privilege Level | **15888**
- Release Information | **15888**

Syntax

```
route-timeout seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng],
[edit protocols ripng],
[edit routing-instances routing-instance-name protocols ripng]
```

Description

Configure the route timeout interval for RIPng.

Options

seconds—Estimated time to wait before making updates to the routing table.

- **Range:** 30 through 360 seconds
- **Default:** 180 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Support for routing instances introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring RIPng Update Interval

route-distinguisher-id-use-router-id

IN THIS SECTION

- Syntax | 15889
- Hierarchy Level | 15889
- Description | 15889
- Required Privilege Level | 15889
- Release Information | 15890

Syntax

```
route-distinguisher-id-use-router-id;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options],
[edit logical-systems name routing-options],
[edit logical-systems name tenants name routing-instances name routing-options],
[edit routing-instances name routing-options],
[edit routing-options]
```

Description

Configure automatic route distinguisher to use the master instance router-ID as the route-distinguisher-id so that you do not need to allocate and manage another IP address for that purpose.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

vpn-global-import | 30953

router-advertisement

IN THIS SECTION

- Syntax | 15890
- Hierarchy Level | 15892
- Description | 15892
- Options | 15892
- Required Privilege Level | 15895
- Release Information | 15895

Syntax

```
router-advertisement {
    interface interface-name {
        current-hop-limit current-hop-limit;
        default-lifetime seconds;
        dns-search-list name {
            lifetime seconds;
        }
        dns-server-address name {
            lifetime seconds;
        }
```

```
        downstream name;
        downstream-mode;
        (link-mtu | no-link-mtu);
        (managed-configuration | no-managed-configuration);
        max-advertisement-interval seconds;
        min-advertisement-interval seconds;
        nat-prefix name {
            lifetime seconds;
        }
        (other-stateful-configuration | no-other-stateful-configuration);
        parameter-preference (configured | proxied);
        passive-mode;
        preference (high | low | medium);
        prefix name {
            (autonomous | no-autonomous);
            (on-link | no-on-link);
            preferred-lifetime seconds;
            valid-lifetime seconds;
        }
        reachable-time milliseconds;
        retransmit-timer milliseconds;
        rio-prefix name {
            rio-lifetime seconds;
            rio-preference (high | low | medium);
        }
        solicit-router-advertisement-unicast;
        upstream-mode;
        virtual-router-only;
    }
    ra-secure {
        accept-current-hop-limit-max unit;
        accept-current-hop-limit-min unit;
        accept-reachable-time-max milliseconds;
        accept-reachable-time-min milliseconds;
        accept-retransmit-time-max milliseconds;
        accept-retransmit-time-min milliseconds;
    }
    traceoptions {
        file filename <files files> <size size> <world-readable | no-world-readable>;
        flag name;
    }
}
```

Hierarchy Level

```
[edit logical-systems name protocols],
[edit protocols]
```

Description

Configure router advertisement to multicast router's availability. Router advertisement announces the address autoconfiguration and discovers the link local addresses of the neighboring routers. The router discovery messages enables the host to discover the existence of neighboring routers. Router advertisment messages contain prefixes that are used for address autoconfiguration.

Options

accept-current-hop-limit-max

(Optional) Current hop limit acceptable maximum for incoming RA.

- **Range:** 1 through 255 units.

accept-current-hop-limit-min

(Optional) Current hop limit acceptable minimum for incoming RA.

- **Range:** 1 through 255 units.

accept-reachable-time-max

(Optional) Reachable time acceptable maximum for incoming RA.

- **Range:** 1000 through 3600000 milliseconds.

accept-reachable-time-min

(Optional) Reachable time acceptable minimum for incoming RA.

- **Range:** 1000 through 3600000 milliseconds.

accept-retransmit-time-max

(Optional) Retransmit Time acceptable maximum for incoming RA.

- **Range:** 1000 through 3600000 milliseconds.

accept-retransmit-time-min

(Optional) Retransmit Time acceptable min for incoming RA (1000..3600000 milliseconds).

- **Range:** 1000 through 3600000 milliseconds.

autonomus

(Optional) Indicates the value to be placed in the autonomous flag field in the prefix information option. The "autonomous address-configuration flag", indicates if the option applies to the stateless autoconfiguration.

no-autonomus	(Optional) Indicates to silently ignore the prefix information option as stateless autoconfiguration is not applicable.
current-hop-limit	(Optional) Current hop limit is the default value placed in the hop count field of the IP header for outgoing packets. • **Range:** 0 through 255.
default-lifetime	(Optional) Router lifetime is the lifetime associated with a default router. • **Range:** 0 through 9000 seconds.
dns-search-list	(Optional) DNS search list configuration.
dns-server-address	(Optional) Recursive DNS address configuration.
downstream	(Optional) Used to configure the list of downstream interfaces a given upstream interface.
downstream-mode	(Optional) Used for the interface as downstream interface for RA proxy.
file	(Optional) Trace file options.
flag	(Optional) Tracing parameters.
link-mtu	(Optional) Link MTU. Include the maximum transmission unit (MTU).
no-link-mtu	(Optional) Don't link MTU. Does not include the maximum transmission unit (MTU).
managed-configuration	(Optional) Set managed address configuration. Enable host to use stateful autoconfiguration.
no-managed-configuration	(Optional) Don't set managed address configuration. Disable host from using stateful autoconfiguration.
max-advertisement-interval	(Optional) Maximum advertisement interval. Set the maximum interval between each router advertisement message. • **Range:** 4 through 1800 seconds.
min-advertisement-interval*seconds*	(Optional) Minimum advertisement interval. Set the minimum interval between each router advertisement message. • **Range:** 3 through 1350 seconds.

other-stateful-configuration	(Optional) Set other stateful configuration. Enable autoconfiguration of other nonaddress-related information.
no-other-stateful-configuration	(Optional) Don't set other stateful configuration. Disable autoconfiguration of other nonaddress-related information.
parameter-preference	(Optional) Used for RA proxy. The TLVs/options from RA packet will be proxied if parameter-preference is set as “proxied”. If parameter-preference is set as “configured”, configured options will have preference and configured options will be sent in the RA packet.
passive-mode	(Optional) RA receive only mode.
preference	(Optional) Set the preference for router selection.
prefix	(Optional) Specifies the prefixes that are on-link and/or are advertised for stateless-address autoconfiguration.
nat-prefix	(Optional) Used to configure NAT prefix to be advertised in the RA packet.
preferred-lifetime	(Optional) Indicates the length of time in seconds (relative to the time the packet is sent) for which addresses are generated from the prefix via stateless address autoconfiguration remain preferred.
reachable-time	(Optional) Reachable time. Set the length of time that a node considers a neighbor reachable until another reachability confirmation is received from that neighbor. • **Range:** 0 through 3600000 milliseconds
retransmit-timer*seconds*	(Optional) Retransmit timer. Set the retransmission frequency of neighbor solicitation messages. • 0 through 4294967295 milliseconds.
rio-prefix	(Optional) Route information option configuration.
rio-lifetime	(Optional) Time in seconds for which the advertised route information in the route information option is valid.
rio-preference	(Optional) High, Low, or Medium preference value indicates whether to prefer the advertised route information in the route information option over other default routes.
solicit-router-advertisement-unicast	(Optional) Enable solicited router advertisement as unicast.

upstream-mode (Optional) Used for RA proxy, when enabled for upstream interface. The packets received on the upstream interface will be proxied to one or more downstream interfaces.

valid-lifetime (Optional) Indicates the length of time in seconds (relative to the time the packet is sent) for which the prefix is valid for the purpose of on-link determination.

virtual-router-only (Optional) Send advertisements only for vrrp-inet6-group.

- **Range:** 0 through 3600000 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

Secure IPv6 Neighbor Discovery

router-advertisement (Dynamic Profiles)

IN THIS SECTION

- Syntax | 15896
- Hierarchy Level | 15896
- Description | 15896
- Required Privilege Level | 15896
- Release Information | 15897

Syntax

```
router-advertisement {
    interface interface-name {
         current-hop-limit number;
        default-lifetime seconds;
        (managed-configuration | no-managed-configuration);
        max-advertisement-interval seconds;
        min-advertisement-interval seconds;
        (other-stateful-configuration | no-other-stateful-configuration);
        prefixprefix {
            (autonomous | no-autonomous);
            (on-link | no-on-link);
            preferred-lifetime seconds;
            valid-lifetime seconds;
        }
        reachable-time milliseconds;
        retransmit-timer milliseconds;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles protocols]
```

Description

Enable router advertisement in a dynamic client profile or a dynamic service profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

IPv6 WAN Link Addressing with NDRA

Dynamic Router Advertisement Configuration Overview

router-advertisement-guard

IN THIS SECTION

- Syntax | 15897
- Hierarchy Level | 15898
- Description | 15898
- Required Privilege Level | 15899
- Release Information | 15899

Syntax

```
router-advertisement-guard {
    interface  interface-name {
        mark-interface (trusted | block);
        policy policy-name (stateful | stateless);
    }
    vlans (vlan-name | all) {
        policy policy-name (stateful | stateless);
    }
    policy policy-name {
        accept {
            match-list {
                match-criteria {
                    (match-all | match-any);
```

```
                }
                prefix-list-name prefix-list-name;
                source-ip-address-list address-list-name;
                source-mac-address-list address-list-name;
            }
            match-option {
                hop-limit {
                    (maximum | minimum) value;
                }
                managed-config-flag;
                other-config-flag;
                router-preference (high | low | medium);
            }
        }
        discard {
            prefix-list-name prefix-list-name;
            source-ip-address-list address-list-name;
            source-mac-address-list address-list-name;
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options access-security]
```

Description

Configure IPv6 Router Advertisement (RA) guard. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in a policy. The policy can be either an accept policy or a discard policy. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped. When RA guard is enabled by using a discard policy, any RA messages that match the conditions are dropped, and RA messages that do not match the conditions are forwarded.

You can enable RA guard on an interface or on a VLAN. You must first configure a policy at the [`edit forwarding-options access-security router-advertisement-guard`] hierarchy level. The policy is then applied to an

interface at the [`edit forwarding-options access-security router-advertisement-guard interface` *`interface-name`*] hierarchy level, or to a VLAN at the [`edit forwarding-options access-security router-advertisement-guard vlan` *`vlan-name`*] hierarchy level.

> **NOTE**: If you apply an RA guard policy on an interface, you must enable RA guard on the VLAN that is associated with that interface using the `vlan` statement at the [`edit forwarding-options access-security router-advertisement-guard`] hierarchy level.

You can configure RA guard to be stateless or stateful. Stateless RA guard enables a switch to examine incoming RA messages and filter each message on the basis of whether it matches the conditions configured in the policy. For example, an interface can be statically configured to forward RA messages only from predefined sources. Stateful RA guard enables a switch to learn about legitimate senders of RA messages and store this information, which is used to validate senders of subsequent RA messages. For example, an interface that is in the learning state and receives RA messages from legitimate senders dynamically transitions to the forwarding state, in which RA messages from valid senders are forwarded to their destination.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

router-discovery

IN THIS SECTION

- Syntax | 15900
- Hierarchy Level | 15901
- Description | 15901
- Default | 15901
- Required Privilege Level | 15901
- Release Information | 15901

Syntax

```
router-discovery {
    disable;
    address address {
        (advertise | ignore);
        (broadcast | multicast);
        (ineligible | priority number);
    }
    interface interface-name {
        lifetime seconds;
        max-advertisement-interval seconds;
        min-advertisement-interval seconds;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit protocols]
```

Description

Enable ICMP router discovery (server mode) on the router.

The remaining statements are explained separately. See CLI Explorer.

Default

Router discovery is disabled on the router. When router discovery is enabled, the default behavior is to advertise all interfaces. If the router supports multicast, all the IPv4 Layer 3 interfaces are advertised through multicast. Otherwise, all the IPv4 Layer 3 interfaces are advertised through broadcast.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

router-id

IN THIS SECTION

- Syntax | 15902
- Hierarchy Level | 15902
- Description | 15902
- Options | 15903
- Required Privilege Level | 15904
- Release Information | 15904

Syntax

```
router-id address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Specify the routing device's IP address.

The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier usually is the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.

> **NOTE**: We strongly recommend that you configure the router identifier under the `[edit routing-options]` hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

You must configure a router-id in order for BGP and OSPF to function in a routing instance. Use the **show route instance detail** command to display the router-id value for a routing instance. If the router-id is `0.0.0.0`, then the routing instance has no router-id.

For more information about the router identifier in OSPF, see *Example: Configuring an OSPF Router Identifier*.

> **NOTE**: If you run OSPF3 or IPv6 based BGP peering in a routing instance, you must configure a router identifier (`router-id`) in the routing instance itself. In other words, the `router-id` used for the main routing instance is **not** inherited by other routing instances. Even if you run *only* OSPF3 or IPv6 BGP peering in a routing instance, a 32-bit `router-id` must be configured in the instance. This is because IPv6 routing protocols use the `router-id` for handshaking. **The router ID must be configured as a 4 octet (32 bit) unsigned non-zero integer value.**
>
> **It's often convenient to use an IPv4 address as the router ID. However, a valid IPv4 address is not required. The RID does not have to be a routable IPv4 address. You can configure any 32-bit value that is unique within the routing domain.** If you do not configure the `router-id` in an IPv6 OSPF or BGP routing instance the IPv6 protocols will use an invalid value for the router ID (`0.0.0.0`) and the adjacency and connections will fail

When you configure a router ID with an IP address that differs from the BGP local address for EVPN routing instances, you must also configure a policy that exports the local IP address for BGP routes. This ensures that EVPN traffic will be properly routed. For more information on exporting BGP routes, see *Basic BGP Routing Policies*.

Options

address—IP address of the routing device.

- **Default:** Address of the first interface encountered by Junos OS

Starting in Junos OS Release 20.4R1, you can establish a BGP connection with the BGP identifier that is a 4-octet, unsigned, and non-zero integer and it needs to be unique only within the Autonomous System (AS) per *RFC6286, Autonomous-System-Wide Unique BGP Identifier for BGP-4*

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Examples: Configuring External BGP Peering
- Examples: Configuring Internal BGP Peering

routing-engine

IN THIS SECTION

- Syntax | 15904
- Hierarchy Level | 15905
- Description | 15905
- Required Privilege Level | 15905
- Release Information | 15905

Syntax

```
routing-engine {
    on-disk-failure {
        disk-failure-action (halt | reboot);
    }
}
```

Hierarchy Level

```
[edit chassis]
[edit chassis interconnect-device name],
[edit chassis node-group name]
```

Description

Configure a Routing Engine to halt or reboot automatically when a hard disk error occurs. A hard disk error may cause a Routing Engine to enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding. Rebooting or halting prevents this.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- Enabling a Routing Engine to Reboot on Hard Disk Errors
- Junos OS High Availability User Guide

routing-engine (Chassis)

IN THIS SECTION

- Syntax | 15906
- Hierarchy Level | 15906
- Description | 15906

- Options | 15907
- Required Privilege Level | 15907
- Release Information | 15907

Syntax

```
routing-engine {
    on-disk-failure (Chassis Routing Engine)
    {
        disk-failure-action (halt | reboot);
    }
    disk {
        ssd-series ssd-series-name {
            id id-num {
                id-threshold id-threshold-value;
                id-value (norm | raw);
                id-flag (high | low);
            }
        }
        smart-check{
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure a Routing Engine to halt or reboot automatically when a hard disk error occurs. A hard disk error may cause a Routing Engine to enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding. Rebooting or halting prevents this. You can also configure the device to perform certain health checks on the Routing Engine solid-state drive (SSD) and log a health event or raise an alarm in case a predefined health attribute threshold is breached.

Options

disk Configure health check parameters for the SSD.

ssd-series *ssd-series-name* Specify the SSD series name.

id *id-num* Specify the ID of the SSD smart attribute for which you are configuring the health check. You can view the smart attributes and their IDs of an SSD by using the command `show vmhost hard-disk-test status disk /dev/sda`.

- **Range:** 1 - 255

id-threshold Specify the SSD smart attribute threshold, breaching which will result in the device logging a health event or raising an alarm.

id-value Choose the attribute value type from the following options:

- norm—Use this option to instruct the device to consider normalized value for the specified smart attribute.
- raw—Use this option to instruct the device to consider raw value for the specified smart attribute.

id-flag Choose the attribute value flag from the following options:

- high—The software considers the threshold to be breached if the value is greater than smart threshold.
- low—The software considers the threshold to be breached if the value is less than smart threshold.

smart-check Configure the device to raise alarm when an error occurs. See "smart-check" on page 28045.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The `disk-failure-action` statement added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Enabling a Routing Engine to Reboot on Hard Disk Errors

routing-engine (Chassis Redundancy)

IN THIS SECTION

- Syntax | 15908
- Hierarchy Level | 15908
- Description | 15908
- Default | 15908
- Options | 15909
- Required Privilege Level | 15909
- Release Information | 15909

Syntax

```
routing-engine slot-number (backup | disabled | master);
```

Hierarchy Level

```
[edit chassis redundancy]
```

Description

Configure Routing Engine redundancy.

Default

By default, the Routing Engine in slot 0 is the primary Routing Engine and the Routing Engine in slot 1 is the backup Routing Engine.

Options

slot-number—Specify the slot number (0 or 1).

Set the function of the Routing Engine for the specified slot:

- **master**—Routing Engine in the specified slot is the primary.
- **backup**—Routing Engine in the specified slot is the backup.
- **disabled**—Routing Engine in the specified slot is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routing Engine Redundancy

routing-engine-profile

IN THIS SECTION

- Syntax | 15910
- Hierarchy Level | 15910
- Description | 15910
- Options | 15910
- Required Privilege Level | 15910
- Release Information | 15911

Syntax

```
routing-engine-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
    start-time time;
}
```

Hierarchy Level

```
[edit accounting-options]
```

Description

Create a Routing Engine profile to collect selected Routing Engine statistics and write them to a file in the `/var/log` directory.

Options

profile-name Name of the Routing Engine statistics profile.

start-time *time* (Optional) Specify the start time in *yyyy-mm-dd.hh:mm* local time format for the accounting profile. Sometimes, if the time at which the records are written to the accounting file is beyond the transfer interval window of the file, drifting of accounting records occurs. To prevent drifting and to have predictable timestamps of when accounting data is written to an accounting file, we can specify the start time option for the accounting profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configure the Routing Engine Profile

routing-engine-services

IN THIS SECTION

- Syntax | 15911
- Hierarchy Level | 15911
- Description | 15911
- Required Privilege Level | 15912
- Release Information | 15912

Syntax

```
routing-engine-services;
```

Hierarchy Level

```
[edit services service-set service-set service-set-options]
```

Description

When configuring a Routing Engine-based captive portal service, specify the service set options to apply to a service set. The services interfaces on the Routing Engine are identified with an si- prefix (for example, si-1/1/0). The si- interface contains all redirect and rewrite traffic and services for the Routing Engine.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

HTTP Redirect Service Overview

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

Configuring Routing Engine-Based, Static HTTP Redirect Services

Configuring Routing Engine-Based, Converged HTTP Redirect Services

routing-engine (SNMP Resource Level)

IN THIS SECTION

- Syntax | 15913
- Hierarchy Level | 15913
- Description | 15913
- Options | 15913
- Required Privilege Level | 15914
- Release Information | 15914

Syntax

```
routing-engine  {
    resource <cpu | memory | open-files-count | process-count  | storage | temperature |
traceoptions > ;
    {
        interval <interval in secs>;
        moderate-threshold <percentage level>;
        high-threshold <percentage level>;
        critical-threshold <percentage level>;
        action <monitor | prevent | recover>;
    }
}
```

Hierarchy Level

```
[edit snmp health-monitor routing-engine]
```

Description

Override the global configuration for a resource.

Options

- `interval`—Monitoring interval in seconds.

  Range: 1 through 604800 seconds

  Default: 300 seconds

- `moderate-threshold`—Percentage of moderate threshold level resource utilization.

  Range: 30 through 99 percent

  Default: 70 percent

- `high-threshold` —Percentage of high-threshold level resource utilization.

  Range: 30 through 99 percent

  Default: 80 percent

- `critical-threshold` —Percentage of critical threshold level resource utilization.

  Range: 30 through 99 percent

  Default: 90 percent

- `action`–Enable action for all resources.

  Default: If action is not enabled, the default action is prevent.

> **WARNING**: If the system health management action for an affected resource is configured to recover, then certain instrusive operations necessary for preventing system breakdown are taken. Instrusive operations can include restarting or terminating processes, deleting files, and so on. Such action information is logged in the system health management history and system log.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. Statement modified in Junos OS Release 15.1X49-D10.

routing-header

IN THIS SECTION

- Syntax | 15915
- Hierarchy Level | 15915
- Description | 15915
- Required Privilege Level | 15915
- Release Information | 15915

Syntax

```
routing-header {
    header-type {
        match (equal | greater-than | less-than | not-equal);
        value header-value;
    }
}
```

Hierarchy Level

```
[edit  set security idp custom-attack attack-name attack-type signature protocol ipv6 extension-
header]
```

Description

Specify the IPv6 routing header type. The `routing-header` option inspects the routing-header type field and reports a custom attack if a match with the specified value is found. The `routing-header` option supports the following routing header types: `routing-header-type0`, `routing-header-type1`, and so on.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

routing-instance (Services)

IN THIS SECTION

Syntax | 15916

- Hierarchy Level | 15916
- Description | 15916
- Options | 15917
- Required Privilege Level | 15917
- Release Information | 15917

Syntax

```
routing-instance routing-instance-name{
    family inet {
        address ip-address;
    }
}
```

Hierarchy Level

```
[edit services app-engine compute-cluster compute-cluster-name local-management],
[edit services app-engine virtual-machines instance instance-name local-management]
```

Description

(Optional) Configure a routing instance for the JunosV App Engine infrastructure. Packets are restricted to the specified routing instance. By specifying the routing instance, the JunosV App Engine infrastructure can be made to operate in a non-default routing instance. This feature ensures that JunosV App Engine traffic can be segregated from other traffic.

If the `routing-instance` statement is absent from the `[edit services app-engine virtual-machines instance instance-name local-management]` hierarchy level, the virtual machine (VM) will inherit this routing instance from its `[edit services app-engine compute-cluster compute-cluster-name local-management]` hierarchy level.

NOTE: Additional configuration under the `[edit routing-instances]` hierarchy level is required to ensure that the Routing Engine interface address belongs to that particular routing instance on the routing engine. Each routing instance consists of sets of the

following: Routing tables, interfaces that belong to these routing tables (optional, depending upon the routing instance type), and routing option configurations.

Options

routing-instance-name	Name of the routing instance.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

routing-instance (SNMPv3)

IN THIS SECTION

- Syntax | 15917
- Hierarchy Level | 15918
- Description | 15918
- Options | 15918
- Required Privilege Level | 15918
- Release Information | 15918

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit snmp v3 target-address target-address-name]
```

Description

Specify a routing instance for an SNMPv3 trap target.

Options

routing-instance-name—Name of the routing instance.

To configure a routing instance within a logical system, specify the logical system name followed by the routing instance name. Use a slash (/) to separate the two names (for example, `test-ls/test-ri`). To configure the default routing instance on a logical system, specify the logical system name followed by `default` (for example, `test-ls/default`).

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

routing-instance (SNMP)

IN THIS SECTION

- Syntax | 15919
- Hierarchy Level | 15919
- Description | 15919
- Options | 15919

Required Privilege Level | 15919
Release Information | 15920

Syntax

```
routing-instance routing-instance-name {
    source-address address;
}
```

Hierarchy Level

```
[edit snmp community community-name],
[edit snmp community community-name logical-system logical-system-name],
[edit snmp trap-group group],
[edit snmp trap-options]
```

Description

Specify a routing instance for SNMPv1 and SNMPv2 trap targets. All targets configured in the trap group use this routing instance.

If the routing instance is defined within a logical system, include the `logical-system logical-system-name` statement at the `[edit snmp community community-name]` hierarchy level and specify the `routing-instance` statement under the `[edit snmp community community-name logical-system logical system-name]` hierarchy level.

Options

routing-instance-name—Name of the routing instance.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Added to the [edit snmp community *community-name*] hierarchy level in Junos OS Release 8.4.

Added to the [edit snmp community *community-name* logical-system *logical-system-name*] hierarchy level in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specify a Routing Instance in an SNMPv1 or SNMPv2c Community

routing-instance (Forwarding-options Analyzer)

IN THIS SECTION

- Syntax | 15920
- Hierarchy Level | 15921
- Description | 15921
- Options | 15921
- Required Privilege Level | 15921
- Release Information | 15921

Syntax

```
routing-instance {
    instance-name {
        bridge-domain bridge-domain-name;
    }
    vlan (vlan-id | vlan-name);
}
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input egress],
[edit forwarding-options analyzer analyzer-name input ingress],
[edit forwarding-options analyzer analyzer-name output]
```

Description

Configure routing instance.

Options

instance-name	Name of the routing instance.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

routing-instance

IN THIS SECTION

- Syntax | 15922
- Hierarchy Level | 15922
- Description | 15922
- Options | 15922
- Required Privilege Level | 15922
- Release Information | 15922

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit firewall family inet filter filter-name term
term-name then]
```

Description

Specify a specific virtual routing instance to which the switch sends matched packets.

Options

routing-instance-name —Name of a virtual routing instance.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

> *Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device*
>
> Configuring Virtual Routing Instances on EX Series Switches
>
> *Understanding Filter-Based Forwarding for EX Series Switches*

routing-instance (Interfaces)

IN THIS SECTION

- Syntax | 15923
- Hierarchy Level | 15923
- Description | 15923
- Default | 15923
- Required Privilege Level | 15924
- Release Information | 15924

Syntax

```
routing-instance {
    destination routing-instance-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number tunnel],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
tunnel]
```

Description

Specify the destination routing instance that points to the routing table containing the tunnel destination address.

Default

The default Internet routing table `inet.0`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Tunnel Interfaces for Routing Table Lookup

routing-instance-access

IN THIS SECTION

- Syntax | 15924
- Hierarchy Level | 15925
- Description | 15925
- Required Privilege Level | 15925
- Release Information | 15925

Syntax

```
[edit snmp]
    routing-instance-access {
        access-list {
            routing-instance;
            routing-instance restrict;
        }
    }
```

Hierarchy Level

```
[edit snmp]
```

Description

Enable SNMP managers in routing instances other than the default routing instance to access SNMP information. For information about the **access-list** option, see "access-list" on page 1686.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Enable SNMP Access over Routing Instances

routing-instance (Access LDAP)

IN THIS SECTION

- Syntax | **15926**
- Hierarchy Level | **15926**
- Description | **15926**
- Options | **15926**
- Required Privilege Level | **15926**
- Release Information | **15926**

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit access ldap-server server-address],
[edit access profile profile-name ldap-server server-address]
```

Description

Configure the routing instance used to send LDAP packets to the LDAP server. A routing instance is a collection of routing tables, the interfaces contained in the routing tables, and the routing protocol parameters that control the information in the routing tables.

Options

routing-instance-name—Name of the routing instance.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring RADIUS and LDAP User Authentication

routing-instance (Access RADIUS)

IN THIS SECTION

- Syntax | 15927
- Hierarchy Level | 15927
- Description | 15927
- Options | 15927
- Required Privilege Level | 15927
- Release Information | 15928

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit access radius-server server-address],
[edit access profile profile-name radius-server server-address]
```

Description

Configure the routing instance used to send RADIUS packets to the RADIUS server. A routing instance is a collection of routing tables, the interfaces contained in the routing tables, and the routing protocol parameters that control the information in the routing tables.

Options

routing-instance-name —Name of the routing instance.

Required Privilege Level

secret—To view this statement in the configuration.

secret-control—To add this statement to the configuration.

Release Information

Statement modified in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Firewall User Authentication Overview

routing-instance (Advanced Policy-Based Routing)

IN THIS SECTION

- Syntax | 15928
- Hierarchy Level | 15928
- Description | 15928
- Options | 15929
- Required Privilege Level | 15929

Syntax

```
routing-instance name ;
```

Hierarchy Level

[edit security advance-policy-based-routing profile *profile-name* rule *rule-name* then]

Description

Specify a specific routing instance to which the device sends the matched packets.

When traffic arrives at the specified zone or interface, it is matched by the advanced policy-based routing (APBR) profile (application profile). The application profile matches applications and application groups and if the matching rule is found, the packets are routed to the routing instance that sends the traffic to a different interface as specified in the next-hop IP address.

The routing instances specify the routing table and the destination to which a packet is forwarded. The following types of routing instances are supported:

- Forwarding—Use this routing instance type for filter-based forwarding applications.
- Virtual router—Similar to the forwarding instance type, but used for non-VPN-related applications.

Options

name Specify the name of the routing instance.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Advanced Policy-Based Routing for Application-Aware Traffic Management Solution

Understanding Advanced Policy-Based Routing

routing-instance (Flow Monitoring)

IN THIS SECTION

- Syntax | 15930
- Hierarchy Level | 15930
- Description | 15930
- Options | 15930

Required Privilege Level | 15930
Release Information | 15931

Syntax

```
routing-instance instance-name;
```

Hierarchy Level

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server hostname]
[edit forwarding-options sampling instance sampling-instance-name family (inet | inet6 | mpls)
output flow-server hostname]
```

Description

Configure a non-default VPN routing and forwarding (VRF) instance through which flow collectors can be reachable for inline flow monitoring. You cannot configure a flow collector to be reachable through non-default VRF instances for version 5 and version 8 flows. You can configure a flow collector to be reachable through WAN ports belonging to non-default VRF instances for version 9 and IPFIX flows. You must also configure the routing instance to be a VRF instance by including the `instance-type vrf` statement at the `[edit routing-instances instance-name]` hierarchy level.

Options

instance-name—Name of a routing instance that has been configured at the `[edit routing-instance]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Evolved Release 24.2R1 at the [edit forwarding-options sampling instance *sampling-instance-name* family (inet | inet6 | mpls) output flow-server *hostname*] hierarchy level.

RELATED DOCUMENTATION

Directing Traffic Sampling Output to a Server Running the cflowd Application

Understand Inline Active Flow Monitoring

routing-instance (Diameter Peer)

IN THIS SECTION

- Syntax | 15931
- Hierarchy Level | 15931
- Description | 15932
- Options | 15932
- Required Privilege Level | 15932
- Release Information | 15932

Syntax

```
routing-instance routing-instance-name ;
```

Hierarchy Level

```
[edit diameter peer peer-name]
```

Description

Specify a routing instance for a Diameter peer. Alternatively, you can include the `logical-system` statement at the `[edit diameter peer peer-name]` hierarchy level to configure a logical and routing instance.

Options

routing-instance-name—Name of the routing instance.

- **Default:** Master routing instance

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Diameter

Configuring Diameter Peers

routing-instance (Diameter Transport)

IN THIS SECTION

- Syntax | 15933
- Hierarchy Level | 15933
- Description | 15933
- Options | 15933
- Required Privilege Level | 15933
- Release Information | 15933

Syntax

```
routing-instance routing-instance-name ;
```

Hierarchy Level

```
[edit diameter transport transport-name]
```

Description

Specify a routing instance for the Diameter transport layer connection.

Options

routing-instance-name—Name of the routing instance.

- **Default:** Master routing instance

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

> *Configuring the Diameter Transport*

routing-instance (Event Scripts)

IN THIS SECTION

- Syntax | 15934
- Hierarchy Level | 15934
- Description | 15934
- Options | 15934
- Required Privilege Level | 15935
- Release Information | 15935

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit event-options event-script file filename]
```

Description

Configure the routing instance you want to use to update automation scripts. To use a management instance, configure the `management-instance` statement along with the `routing-instance` statement, thus enabling event scripts to use the non-default management routing instance mgmt_junos when refreshing the scripts.

Options

routing-instance-name	Name of the routing instance. For the management instance, use `mgmt_junos`. Otherwise, you can specify any routing instance name.

> **NOTE**: You must also define the routing instance under the `[edit routing-instances]` hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

- management-instance
- Management Interface in a Non-Default Instance

routing-instance (JET Scripts)

IN THIS SECTION

- Syntax | 15936
- Hierarchy Level | 15936
- Description | 15936
- Options | 15936
- Required Privilege Level | 15936
- Release Information | 15936

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit system extensions extension-service application file filename]
```

Description

Configure the routing instance you want to use to update Automation scripts. To use a management instance, configure the `management-instance` statement along with the `routing-instance` statement, thus enabling JET scripts to use the non-default management routing instance mgmt_junos when refreshing the scripts.

Options

routing-instance-name Name of the routing instance. For the management instance, use `mgmt_junos`. Otherwise, you can specify any routing instance name.

> **NOTE**: You must also define the routing instance under the `[edit routing-instances]` hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

- management-instance

Management Interface in a Non-Default Instance

routing-instance (Junos Telemetry Interface)

IN THIS SECTION

- Syntax | 15937
- Hierarchy Level | 15937
- Description | 15937
- Options | 15938
- Required Privilege Level | 15938
- Release Information | 15938

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit services analytics export-profile export-profile-name]
```

Description

For legacy gRPC dialout connections, configure a routing instance in an export profile. You can only configure the routing instance in an export profile for legacy remote procedure call (gRPC) dial-out connections. An error is returned if the routing instance configured is used with User Datagram Protocol (UDP) transport.

Configuring the routing instance in an export profile is not mandatory. If you choose not to configure a routing instance, the default routing instance is applied.

Options

routing-instance A routing instance to be used for gRPC transport.

Required Privilege Level

Edit

Release Information

Statement introduced in Junos OS Evolved 24.2R1 (for IPv6) on

Statement introduced in Junos OS Evolved 24.2R1 (for IPv6) on ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7900, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, and QFX5700E.

RELATED DOCUMENTATION

Using gRPC Dial-Out for Secure Telemetry Collection

export-profile (Junos Telemetry Interface) | **6669**

local-address (Junos Telemetry Interface) | **9902**

routing-instance (L2TP Destination)

IN THIS SECTION

- Syntax | **15939**
- Hierarchy Level | **15939**
- Description | **15939**
- Options | **15939**
- Required Privilege Level | **15939**
- Release Information | **15939**

Syntax

```
routing-instance routing-instance-name {
    drain;
}
```

Hierarchy Level

```
[edit services l2tp destination address ip-address]
```

Description

Specify the routing instance in which the destination is created.

Options

routing-instance-name— Name of the routing instance.

- **Default:** Routing instance *default*

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

routing-instance (L2TP Tunnel Destination)

IN THIS SECTION

- Syntax | 15940
- Hierarchy Level | 15940
- Description | 15940
- Options | 15940
- Required Privilege Level | 15941
- Release Information | 15941

Syntax

```
routing-instance routing-instance-name {
    drain;
}
```

Hierarchy Level

```
[edit services l2tp tunnel name tunnel-name address ip-address]
```

Description

Specify the routing instance in which the tunnel to the destination address is created.

Options

routing-instance-name— Name of the routing instance.

- **Default:** Routing instance *default*

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

routing-instance-list (TWAMP)

IN THIS SECTION

- Syntax | 15941
- Hierarchy Level | 15942
- Description | 15942
- Options | 15942
- Required Privilege Level | 15942
- Release Information | 15942

Syntax

```
routing-instance-list {
    instance-name {
        port number;
    }
}
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Configure the Two-Way Active Measurement Protocol (TWAMP) servers on specific routing instances, instead of associating the TWAMP server at the system-level to apply to all routing instances configured on a router. The default routing instance is Internet routing table `inet.0`. If you do not specify a routing instance, the TWAMP probe applies to all routing instances. To apply the TWAMP probe to only the default routing instance, you must explicitly set the value of instance-name to default. If an interface is not part of any routing instance, the default port is used for TWAMP probes. You can configure up to 100 routing instances for a TWAMP server.

Options

instance-name—Name of the routing instance, a maximum of 31 characters.

number—Port number.

- **Range:** 1 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

> *Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches*

routing-instance-name

IN THIS SECTION

- Syntax | 15943
- Hierarchy Level (EX Series) | 15943
- Hierarchy Level (MX Series) | 15943
- Description | 15943
- Required Privilege Level | 15944
- Release Information | 15944

Syntax

```
routing-instance-name;
```

Hierarchy Level (EX Series)

```
[edit vlans forwarding-options dhcp-security option-82 circuit-id prefix]
```

Hierarchy Level (MX Series)

```
[edit bridge-domains bridge domain name forwarding-options dhcp-security option-82 circuit-id
prefix]
```

Description

Specify that the routing instance name be included within the optional prefix for the circuit ID suboption in the DHCP option 82 information that is inserted by the switch into the packet header of a DHCP request before it forwards or relays the request to a DHCP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Setting Up DHCP Option 82 on the Switch with No Relay (ELS)

Configuring DHCP Option 82 on a Router with Bridge Domain

routing-instance-name (circuit-id)

IN THIS SECTION

- Syntax | 15944
- Hierarchy Level | 15944
- Description | 15945
- Required Privilege Level | 15945
- Release Information | 15945

Syntax

```
routing-instance--name;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security option-82 circuit-id prefix]
```

Description

Specify that the routing instance name used by the VLAN is included with the circuit ID suboption in the DHCP option 82 information that is inserted by the switch into the packet header of a DHCP request before it forwards or relays the request to a DHCP server

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Setting Up DHCP Option 82 on the Switch with No Relay (ELS)

Understanding DHCP Option 82

routing-instance-name (DHCP Local Server)

IN THIS SECTION

- Syntax | 15946
- Hierarchy Level | 15946
- Description | 15947
- Required Privilege Level | 15947
- Release Information | 15947

Syntax

```
routing-instance-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication username-
include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  authentication
username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication
username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
```

```
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify that the routing instance name be concatenated with the username during the subscriber authentication or DHCP client authentication process. No routing instance name is concatenated if the configuration is in the default routing instance.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

routing-instance-name (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 15948
- Hierarchy Level | 15948
- Description | 15949
- Required Privilege Level | 15949
- Release Information | 15949

Syntax

```
routing-instance-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication username-
include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-
include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name authentication username-include],
[edit tenants tenant-name routing-instances routing-instance-name forwarding-options dhcp-relay
dhcpv6 group group-name authentication username-include]
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication
username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
authentication username-include]
```

Description

Specify that the routing instance name is concatenated with the username during the subscriber authentication or client authentication process. No routing instance name is concatenated if the configuration is in the default routing instance. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group` *`dual-stack-group-name`*`]` hierarchy level introduced in Junos OS Release 15.1.

The `tenants` option is introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Specifying Authentication Support

Creating Unique Usernames for DHCP Clients

routing-instance-name (Static Subscribers)

IN THIS SECTION

- Syntax | 15950
- Hierarchy Level | 15950
- Description | 15950
- Required Privilege Level | 15950

Release Information | 15951

Syntax

```
routing-instance-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication username-include],
[edit logical-systems logical-system-name system services static-subscribers  authentication
username-include],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],
[edit system services static-subscribers authentication username-include],
[edit system services static-subscribers group group-name authentication username-include]
```

Description

Specify that the name of the routing instance is included as part of the username created for all static subscribers or for the static subscribers in the specified group. The group version of the statement takes precedence over the global version. The username is also sent to RADIUS in the Access-Request message.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Configuring the Static Subscriber Global Username

Configuring the Static Subscriber Group Username

routing-instance (PCC Action Profiles)

IN THIS SECTION

- Syntax | 15951
- Hierarchy Level | 15952
- Description | 15952
- Options | 15952
- Required Privilege Level | 15952
- Release Information | 15952

Syntax

```
routing-instance {
    downlink downlink-vrf-name;
    uplink uplink-vrf-name;
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name steering],
[edit services pcef pcc-action-profiles profile-name steering]
```

Description

Specify the routing instance that a PCC action profile uses for steering traffic.

Options

downlink *downlink-vrf-name*
Use the specified name of the routing instance for downlink traffic (to the access side) or the predefined dynamic interface variable .

uplink *uplink-vrf-name*
Use the specified name of the routing instance for uplink traffic (from the access side).

NOTE: The routing instances must have been previously configured.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support at the `[edit services pcef pcc-action-profiles profile-name steering]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

- Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware
- Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment
- *Configuring Policy and Charging Control Action Profiles for Subscriber Management*
- *Understanding Application-Aware Policy Control for Subscriber Management*

routing-instance (Port Mirroring)

IN THIS SECTION

- Syntax | **15953**
- Hierarchy Level | **15953**
- Description | **15954**
- Required Privilege Level | **15954**
- Release Information | **15954**

Syntax

```
routing-instance instance-name;
```

Hierarchy Level

```
[edit forwarding-options] analyzer name output]
[edit forwarding-options port-mirroring [instance name] family inet output interface name]
```

Description

Configure a port mirroring instance. You do not specify an input for this instance. Instead, you, create a firewall filter that specifies the required traffic and directs it to the mirror. This instance type is useful for controlling which types of traffic should be mirrored.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

routing-instance (PPPoE Service Name Tables)

IN THIS SECTION

- Syntax | 15954
- Hierarchy Level | 15955
- Description | 15955
- Options | 15955
- Required Privilege Level | 15955
- Release Information | 15955

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name],
[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci
circuit-id-string ari remote-id-string]
```

Description

Use in conjunction with the `dynamic-profile` statement at the same hierarchy levels to specify the routing instance in which to instantiate a dynamic PPPoE interface. You can associate a routing instance with a named service entry, `empty` service entry, or `any` service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The routing instance associated with a service entry in a PPPoE service name table overrides the routing instance associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the `routing-instance` statement at the `[edit protocols pppoe service-name-tables` *`table-name`* `service` *`service-name`* `agent-specifier aci` *`circuit-id-string`* `ari` *`remote-id-string`*`]` hierarchy level, you cannot also include the `static-interface` statement at this level. The `routing-instance` and `static-interface` statements are mutually exclusive for ACI/ARI pair configurations.

Options

`routing-instance-name`—Name of the routing instance in which the router instantiates the dynamic PPPoE interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation

routing-instance (RPM)

IN THIS SECTION

- Syntax | 15956
- Junos OS Hierarchy Level | 15956
- Junos OS Evolved Hierarchy Level | 15956
- Description | 15957
- Options | 15957
- Required Privilege Level | 15957
- Release Information | 15957

Syntax

```
routing-instance instance-name;
```

Junos OS Hierarchy Level

```
[edit services rpm probe owner test test-name]
[edit services rpm twamp client control-connection control-client-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the routing instance used by the probes. The routing instance is also applicable for control connection.

> **NOTE**: The media interface from where the TWAMP control and test or data packets arrive and exit the `si- logical interface` must be a part of the same routing instance.

Options

`instance-name`—Routing instance configured at the `[edit routing-instance]` hierarchy level.

- **Default:** Internet (IPv4) routing table `inet.0`.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.

Support at the [`edit services rpm twamp client control-connection` *`control-client-name`*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

> *Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches*
>
> *Understanding Two-Way Active Measurement Protocol on Routers*

routing-instance (Security Content Security)

IN THIS SECTION

- Syntax | 15958
- Hierarchy Level | 15958
- Description | 15958
- Options | 15958
- Required Privilege Level | 15959
- Release Information | 15959

Syntax

```
routing-instance name;
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus sophos-engine pattern-update]
[edit security utm feature-profile web-filtering juniper-enhanced server]
[edit security utm feature-profile web-filtering websense-redirect profile wr server]
[edit security utm feature-profile anti-virus sophos-engine server]
[edit security utm default-configuration anti-virus avira-engine pattern-update]
```

Description

Configure the routing instance name. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables. Each routing instance has a unique name.

Options

name Specify the name of the routing instance.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D90.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for Avira engine added in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

admin-email | 1963

url (Security Antivirus) | 30504

routing-instance (Security Destination NAT)

IN THIS SECTION

- Syntax | 15959
- Hierarchy Level | 15960
- Description | 15960
- Options | 15960
- Required Privilege Level | 15960
- Release Information | 15960

Syntax

```
routing-instance (routing-instance-name | default);
```

Hierarchy Level

```
[edit security nat destination pool pool-name]
```

Description

Specify the routing instance on which to perform the route lookup for the address in the pool. It is not a mandatory flag.

A destination NAT pool that does not specify a specific routing instance will default to the routing instance of the ingress zone. You can configure a NAT pool to exist in the default routing instance. As a result, the NAT pool is reachable from zones in the default routing instance and from zones in other routing instances.

Options

routing-instance-name	Name of the routing instance.
default	Use the default routing instance.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

routing-instance (Security Source NAT)

IN THIS SECTION

- Syntax | 15961
- Hierarchy Level | 15961

- Description | 15961
- Options | 15961
- Required Privilege Level | 15961
- Release Information | 15961

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit security nat source pool pool-name]
```

Description

Specify the routing instance to which the pool is bound. It is not a mandatory flag. If the user does not configure the routing instance, by default the pool belongs to routing-instance `inet.0`.

Options

routing-instance-name—Name of the routing instance.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

routing-instance (Syslog)

IN THIS SECTION

- Syntax | 15962
- Hierarchy Level | 15962
- Description | 15962
- Required Privilege Level | 15963
- Release Information | 15963

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit system syslog],
[edit system syslog host ip-address]
```

Description

Configure the routing instance name for the routing instance that you want the syslog client to use. If you want to use the reserved non-default management routing instance mgmt_junos, make sure you configure the `management-instance` statement.

NOTE: You must also define the routing instance you want to use under the `[edit routing-instances]` hierarchy level.

Starting in Junos OS Evolved Release 20.1R1 and 19.4R2, if you are sending syslog messages to a remote host that is identified by its IP address at the `[edit system syslog host ip-address]` hierarchy, you only need to configure the `management-instance` statement to use the mgmt_junos routing instance. You do not need to configure the mgmt_junos routing instance at the `[edit system syslog host ip-address routing-instance]` hierarchy.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Management Interface in a Non-Default Instance

routing-instance (System Scripts)

IN THIS SECTION

- Syntax | 15963
- Hierarchy Level | 15964
- Description | 15964
- Options | 15964
- Required Privilege Level | 15964
- Release Information | 15964

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit system scripts commit file filename],
[edit system scripts op file filename],
[edit system scripts snmp file filename]
```

Description

Configure the routing instance you want to use to update automation scripts. To use a management instance, configure the `management-instance` statement along with the `routing-instance` statement, thus enabling commit, op, and SNMP scripts to use the non-default management routing instance mgmt_junos when refreshing the scripts.

Options

routing-instance-name Name of the routing instance. For the management instance, use `mgmt_junos`. Otherwise, you can specify any routing instance name.

> **NOTE**: You must also define the routing instance under the `[edit routing-instances]` hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

- management-instance
- Management Interface in a Non-Default Instance

routing-instance (Tunnel Profile)

IN THIS SECTION

- Syntax | 15965
- Hierarchy Level | 15965
- Description | 15965
- Options | 15965
- Required Privilege Level | 15965
- Release Information | 15966

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Description

Specify a routing instance for a tunnel.

Options

routing-instance-name—Name of the routing instance.

- **Default:** Routing instance *default*

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

routing-instances (Logical Systems)

IN THIS SECTION

- Syntax | 15966
- Hierarchy Level | 15966
- Description | 15967
- Default | 15967
- Options | 15967
- Required Privilege Level | 15967
- Release Information | 15967

Syntax

```
routing-instances routing-instance-name { ... }
```

Hierarchy Level

```
[edit],
[edit logical-systems logical-system-name]
```

Description

Configure an additional routing entity for a router or switch. You can create multiple instances of BGP, IS-IS, OSPF, OSPF version 3 (OSPFv3), and RIP for a router or switch.

Default

Routing instances are disabled for the router or switch.

Options

routing-instance-name—Name of the routing instance, a maximum of 31 characters. The remaining statements are explained separately.

non-forwarding-vrf—Enable this option to not create a routing and forwarding (VRF) table for local or transit routes belonging to the given VPN.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Configuring Routing Instances on PE Routers in VPNs

routing-instances (JDAF)

IN THIS SECTION

Syntax | 15968

- Hierarchy Level | 15968
- Description | 15968
- Options | 15968
- Required Privilege Level | 15968
- Release Information | 15968

Syntax

```
routing-instances [ routing-instance-names ];
```

Hierarchy Level

```
[edit services jdaf]
```

Description

Configure the routing instances on which Juniper Distributed Application Framework (JDAF) is enabled. If the `jdaf` statement is not configured, then JDAF is disabled.

Options

`routing-instances` **[*`routing-instance-names`*]**	Name or names of routing instances for JDAF clients. If multiple names are being configured, these can be set as an open set of values.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

jdaf | 9422

routing-instances

IN THIS SECTION

- Syntax | 15969
- Hierarchy Level | 15969
- Description | 15969
- Options | 15969
- Required Privilege Level | 15970
- Release Information | 15970

Syntax

```
routing-instances routing-instance-name {
    instance-type virtual-router;
    interface interface-name;
}
```

Hierarchy Level

```
[edit]
```

Description

Configure a virtual routing entity.

Options

routing-instance-name—Name for this routing instance.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches

Configuring Virtual Routing Instances on EX Series Switches

routing-instances (RPM)

IN THIS SECTION

- Syntax | 15970
- Hierarchy Level | 15971
- Description | 15971
- Options | 15971
- Required Privilege Level | 15971
- Release Information | 15971

Syntax

```
routing-instances instance-name;
```

Hierarchy Level

```
[edit services rpm bgp],
[edit services rpm bgp logical-system logical-system-name]
```

Description

Specify the routing instance used by the probes.

Options

instance-name—A routing instance configured at the `[edit routing-instances]` hierarchy level.

- **Default:** Internet routing table `inet.0`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring BGP Neighbor Discovery Through RPM

routing-instances (CoS)

IN THIS SECTION

- Syntax | 15972
- Hierarchy Level | 15972

- Description | 15972
- Default | 15973
- Options | 15973
- Required Privilege Level | 15973
- Release Information | 15973

Syntax

```
routing-instances routing-instance-name {
    classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        exp (classifier-name | default);
        ieee-802.1 (classifier-name | default);
        no-default;
    }
    policy-map policy-map-name {
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

For routing instances with VRF table labels enabled, apply a DSCP, MPLS EXP, or IEEE 802.1 classifier to the routing instance. You can apply the default classifier or one that is previously defined.

Apply the `no-default` option to disable the application of any default classifier to the routing instance.

You can also apply a policy map that defines the rewrite rules for a routing instance.

Default

If you do not include this statement with the `classifiers` option, the default MPLS EXP classifier is applied to the routing instance. When no DSCP classifier is configured, the default MPLS EXP classifier is applied.

Options

`routing-instance-name`—Name of the routing instance.

`classifier-name`—Name of the behavior aggregate DSCP, MPLS EXP, or IEEE 802.1 classifier.

`policy-map-name`—Name of the policy map that defines the rewrite rules for a routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`no-default` and `policy-map` options added in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Custom Forwarding Class for Each Queue

Applying Custom MPLS EXP Classifiers to Routing Instances in Layer 3 VPNs

routing-instances (Dynamic Profiles)

IN THIS SECTION

- Syntax | **15974**
- Hierarchy Level | **15975**

- Description | 15975
- Options | 15975
- Required Privilege Level | 15975
- Release Information | 15976

Syntax

```
routing-instances routing-instance-name {
    interface interface-name;
    multicast-snooping-options {
    }
    routing-options {
        access {
            route prefix {
                metric route-cost;
                next-hop next-hop;
                preference route-distance;
                tag route-tag;
                tag2 route-tag2;
            }
        }
        access-internal {
            route subscriber-ip-address {
                qualified-next-hop underlying-interface {
                    mac-address address;
                }
            }
        }
        multicast {
            interface interface-name {
                no-qos-adjust;
            }
        }
        rib routing-table-name {
            access {
                route prefix {
                    metric route-cost;
                    next-hop next-hop;
```

```
                    preference route-distance;
                    tag route-tag;
                    tag2 route-tag2;
                }
            }
            access-internal {
                route subscriber-ip-address {
                    qualified-next-hop underlying-interface {
                        mac-address address;
                    }
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles]
[edit logical-systems logical-system-name ]
```

Description

Dynamically configure an additional routing entity for a router in a dynamic client profile or a dynamic service profile.

Options

routing-instance-name—The routing instance variable (*$junos-routing-instance*). The routing instance variable is dynamically replaced with the routing instance the accessing client uses when connecting to the router.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support at the `logical-systems` hierarchy level was introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution

routing-instances (Multiple Routing Entities)

IN THIS SECTION

- Syntax | 15976
- Hierarchy Level | 15976
- Description | 15977
- Default | 15977
- Options | 15977
- Required Privilege Level | 15978
- Release Information | 15978

Syntax

```
routing-instances routing-instance-name { ... }
```

Hierarchy Level

```
[edit],
[edit logical-systems logical-system-name]
```

Description

Configure an additional routing entity for a router. You can create multiple instances of BGP, IS-IS, OSPF, OSPFv3, and RIP for a router. You can also create multiple routing instances for separating routing tables, routing policies, and interfaces for individual wholesale subscribers (retailers) in a Layer 3 wholesale network.

Each routing instance consist of the following:

- A set of routing tables
- A set of interfaces that belong to these routing tables
- A set of routing option configurations

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name `my-instance`, its corresponding IP unicast table is my-instance.inet.0. All routes for `my-instance` are installed into my-instance.inet.0.

Routes are installed into the default routing instance inet.0 by default, unless a routing instance is specified.

In Junos OS Release 9.0 and later, you can no longer specify a routing-instance name of *primary*, *default*, or *bgp* or include special characters within the name of a routing instance.

In Junos OS Release 9.6 and later, you can include a slash (/) in a routing-instance name only if a logical system is not configured. That is, you cannot include the slash character in a routing-instance name if a logical system other than the default is explicitly configured. Routing-instance names, further, are restricted from having the form `__.*__` (beginning and ending with underscores). The colon : character cannot be used when multitopology routing (MTR) is enabled.

Default

Routing instances are disabled for the router.

Options

routing-instance-name	—Name of the routing instance. This must be a non-reserved string of not more than 128 characters.
remote-vtep-list	Configure static remote VXLAN tunnel endpoints.
remote-vtep-v6-list	Configure static IPv6 remote VXLAN tunnel endpoints.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`remote-vtep-v6-list` statement introduced in Junos OS Release 17.3 for MX Series routers with MPC and MIC interfaces.

RELATED DOCUMENTATION

Example: Configuring Interprovider Layer 3 VPN Option A

Example: Configuring Interprovider Layer 3 VPN Option B

Example: Configuring Interprovider Layer 3 VPN Option C

routing-interface

IN THIS SECTION

- Syntax | 15978
- Hierarchy Level | 15979
- Description | 15979
- Options | 15979
- Required Privilege Level | 15979
- Release Information | 15979

Syntax

```
routing-interface routing-interface-name;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]
```

Description

(MX Series routers only) Specify a routing interface to include in a bridge domain or a VPLS routing instance.

When you configure `routing-interface irb.x`, the VPLS connection comes up, even if no customer edge (CE) interfaces are configured. This works with one site configured, but not when multiple sites (multisite) are configured.

Options

routing-interface-name—Name of the routing interface to include in the bridge domain or the VPLS routing instance. The format of the routing interface name is **irb.*x***, where ***x*** is the unit number of the routing interface you configured at the `[edit interfaces irb]` hierarchy level. For more information about how to configure a routing interface, see the Junos OS Network Interfaces Library for Routing Devices.

> **NOTE**: You can specify only one routing interface for each bridge domain or VPLS instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring a Bridge Domain

Configuring a Layer 2 Virtual Switch

routing-options

IN THIS SECTION

- Syntax | 15980
- Hierarchy Level | 15980
- Description | 15980
- Required Privilege Level | 15981
- Release Information | 15981

Syntax

```
routing-options { ... }
```

For information on the complete list of `routing-options`, see the Protocol-Independent Routing Properties User Guide .

Hierarchy Level

```
[edit],
[edit logical-systems logical-system-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit tenants tenant-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Configure protocol-independent routing properties.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The [edit tenants *tenant-name* routing-instances *routing-instance-name*] hierarchy level introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Protocol-Independent Routing Properties User Guide

routing-options (Dynamic Profiles)

IN THIS SECTION

- Syntax | 15981
- Hierarchy Level | 15982
- Description | 15983
- Required Privilege Level | 15983
- Release Information | 15983

Syntax

```
routing-options {
    access {
        route prefix {
            metric route-cost;
            next-hop next-hop;
            preference route-distance;
```

```
            tag route-tag;
            tag2 route-tag2;
        }
    }
    access-internal {
        route subscriber-ip-address {
            qualified-next-hop underlying-interface {
                mac-address address;
            }
        }
    }
    multicast {
        interface interface-name {
            no-qos-adjust;
        }
    }
    rib routing-table-name {
        access {
            route prefix {
                metric route-cost;
                next-hop next-hop;
                preference route-distance;
                tag route-tag;
                tag2 route-tag2;
            }
        }
        access-internal {
            route subscriber-ip-address {
                qualified-next-hop underlying-interface {
                    mac-address address;
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name],
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance]
```

Description

Configure protocol-independent routing properties in a dynamic client profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support at the `[edit dynamic-profiles profile-name routing-instances $junos-routing-instance]` hierarchy level introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring Dynamic Access Routes for Subscriber Management

Configuring Dynamic Access-Internal Routes for DHCP and PPP Subscribers

routing-options (Junos Node Slicing)

IN THIS SECTION

- Syntax | 15984
- Hierarchy Level | 15984
- Description | 15984
- Options | 15984
- Required Privilege Level | 15984
- Release Information | 15984

Syntax

```
routing-options {
    static {
        route route {
            next-hop next-hop;
        }
    }
}
```

Hierarchy Level

[edit]

Description

Configure static routes for JDM.

Options

route *route*	Static route destination.
next-hop *next-hop*	Specify the address of the next hop to destination.

Required Privilege Level

root

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

server (Junos node slicing) | 16606

virtual-network-functions (Junos node slicing) | 30773

routing-options-max-ecmp

IN THIS SECTION

- Syntax | 15985
- Hierarchy Level | 15985
- Description | 15986
- Default | 15986
- Options | 15986
- Required Privilege Level | 15987
- Release Information | 15987

Syntax

```
maximum-ecmp (16 | 32 | 64 | 96 | 128 | 160 | 192 | 224 | 256 | 288 | 320 | 352 | 384 | 416 |
448 | 480 | 512);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options]
[edit logical-systems logical-system-name routing-options],

[edit logical-systems logical-system-name tenants tenant-name routing-instances routing-instance-
name routing-options],
[edit routing-instances routing-instance-name routing-options]
[edit routing-options]
[edit tenants tenant-name routing-instances routing-instance-name routing-options].
```

Description

On cRPD, you can control the maximum ECMP limit for nexthops using **routing-options maximum-ecmp**. You can use **routing-options maximum-ecmp** on other platforms but the ECMP limit specified using **routing-options maximum-ecmp** must be less than **chassis maximum-ecmp**.

Default

The default value of ECMP next-hop limit on cRPD is 16.

Options

16	Maximum 16 ECMP routes.
32	Maximum 32 ECMP routes
64	Maximum 64 ECMP routes
96	Maximum 96 ECMP routes
128	Maximum 128 ECMP routes
160	Maximum 160 ECMP routes
192	Maximum 192 ECMP routes
224	Maximum 224 ECMP routes
256	Maximum 256 ECMP routes
288	Maximum 288 ECMP routes
320	Maximum 320 ECMP routes
352	Maximum 352 ECMP routes
384	Maximum 384 ECMP routes
416	Maximum 416 ECMP routes
448	Maximum 448 ECMP routes
480	Maximum 480 ECMP routes

512	Maximum 512 ECMP routes

Required Privilege Level

routing—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Protocol-Independent Routing Properties User Guide

routing-options (QFabric System)

IN THIS SECTION

- Syntax | 15987
- Hierarchy Level | 15988
- Description | 15988
- Required Privilege Level | 15988
- Release Information | 15988

Syntax

```
routing-options {
    multicast {
        no-make-before-break;
    }
}
```

Hierarchy Level

```
[edit fabric]
```

Description

(QFabric systems only) Set routing options in a QFabric system.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Understanding the QFabric System Data Plane

routing-service (Dynamic Profiles)

IN THIS SECTION

- Syntax | 15989
- Hierarchy Level | 15989
- Description | 15989
- Options | 15990
- Required Privilege Level | 15990
- Release Information | 15990

Syntax

```
routing-service (disable <:$junos-routing-services> | enable<:$junos-routing-services>);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-name]
```

Description

Enable configuration of additional routing protocols required by client connections on dynamically created enhanced subscriber management interfaces.

NOTE: When clients use only the standard access-internal routes, access routes, and framed routes, you do not need to configure routing services. In other words, the routing service configuration is not required for simple client reachability purposes.

NOTE: You do not need to configure routing services for Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocol because these protocols are natively supported on enhanced subscriber management interfaces.

NOTE: `routing-service` is supported for the Layer 2 Tunneling Protocol (L2TP) network server subscribers and Point-to-Point Protocol over Ethernet (PPPoE) subscribers. Any other access models, otherwise mentioned, for example aggregated inline service interface, Multilink Point-to-Point Protocol (MLPPP) over L2TP network server (LNS) are not supported.

`routing-service` is not supported for pseudowire subscriber interfaces.

Enable configuration of all routing protocols except Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocol on dynamically created enhanced subscriber management interfaces. IGMP and MLD are natively supported on enhanced subscriber management interfaces, and therefore do not require you to specify the `routing-service` statement in the dynamic profile.

When a dynamic profile containing the `routing-service` statement is instantiated, the router creates an enhanced subscriber management logical interface, also referred to as a pseudo logical interface, in the form demux0.*nnnnnnnnnn* (for example, demux0.3221225472). Any associated subscriber routes or

routes learned from a routing protocol running on the enhanced subscriber management interface use this pseudo interface as the next-hop interface.

> **NOTE**: The `routing-service` statement replaces the deprecated "routing-services" on page 15991 statement. If you upgrade to Junos OS Release 18.4R1 or higher with a configuration that has active subscribers with "routing-services" on page 15991 enabled, then routing services remain enabled for those subscribers. For all new subscribers after the upgrade, you must use the "routing-service" on page 15988 statement.

You can specify the $junos-routing-services predefined variable to determine, per subscriber, whether routing services can be enabled or disabled by the value of the Routing-Services VSA (26-212) returned in the RADIUS Access-Accept message during subscriber authentication. A value of 0x0000 disables installation of routing services. A value of 0x0001 enables the installation of routing services. Any other value is rejected.

> **NOTE**: The `routing-service` statement is not supported for DHCP dynamic profiles. Do not enable it for DHCP dynamic profiles. However, if the `routing-service` statement is enabled for a dynamic VLAN profile, that VLAN can stack DHCP subscribers.

Options

disable	Disable routing services for all subscribers instantiated on the interface by the profile.
disable:$junos-routing-services	Prevent routing services from being enabled or disabled by the value of the Routing-Services VSA (26-212) for the subscriber being authenticated.
enable	Enable routing services for all subscribers instantiated on the interface by the profile.
enable:$junos-routing-services	Allow routing services to be enabled or disabled by the value of the Routing-Services VSA (26-212) for the subscriber being authenticated.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Junos OS Enhanced Subscriber Management

Juniper Networks VSAs Supported by the AAA Service Framework

AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS

Predefined Variables in Dynamic Profiles

routing-services (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 15991
- Hierarchy Level | 15991
- Description | 15991
- Default | 15992
- Required Privilege Level | 15993
- Release Information | 15993

Syntax

```
routing-services;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-name]
```

Description

Enable configuration of additional routing protocols required by client connections on dynamically created enhanced subscriber management interfaces.

NOTE: When clients use only the standard access-internal routes, access routes, and framed routes, you do not need to configure routing services. In other words, the routing service configuration is not required for simple client reachability purposes.

NOTE: You do not need to configure routing services for Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocol because these protocols are natively supported on enhanced subscriber management interfaces.

NOTE: `routing-services` is supported only for the Layer 2 Tunneling Protocol (L2TP) network server subscribers. Any other access models, otherwise mentioned, for example aggregated inline service interface, Multilink Point-to-Point Protocol (MLPPP) over L2TP network server (LNS) are not supported.

When a dynamic profile containing the `routing-services` statement is instantiated, the router creates an enhanced subscriber management logical interface, also referred to as a pseudo logical interface, in the form demux0.*nnnnnnnnnn* (for example, demux0.3221225472). Any associated subscriber routes or routes learned from a routing protocol running on the enhanced subscriber management interface use this pseudo interface as the next-hop interface.

NOTE: Starting in Junos OS Release 18.4R1, the `routing-services` statement is deprecated. It is replaced by the `routing-service` statement. If you upgrade to Junos OS Release 18.4R1 or higher with a configuration that has active subscribers with "routing-services" on page 15991 enabled, then routing services remain enabled for those subscribers. For all new subscribers after the upgrade, you must use the "routing-service" on page 15988 statement.

NOTE: The `routing-services` statement is not supported for DHCP dynamic profiles. Do not enable it for DHCP dynamic profiles. However, if the `routing-services` statement is enabled for a dynamic VLAN profile, that VLAN can stack DHCP subscribers.

Default

Configuration of additional routing protocols associated with the client connection is disabled by default on enhanced subscriber management dynamic interfaces when you do not include this statement in the dynamic profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

Statement deprecated in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Junos OS Enhanced Subscriber Management

routing (System Processes)

IN THIS SECTION

- Syntax | 15993
- Hierarchy Level | 15994
- Description | 15994
- Default | 15994
- Options | 15994
- Required Privilege Level | 15995
- Release Information | 15995

Syntax

```
routing {
    force-32-bit | force-64-bit | auto-64-bit;
}
```

Hierarchy Level

```
[edit system processes],
[edit logical-systems logical-system name system processes]
```

Description

Configure routing protocols process (rpd) mode.

Default

`force-32-bit` mode is used in Junos 15.1F2 and prior releases

`auto-64-bit` mode is used starting in Junos 15.1F3

Options

`auto-64-bit` (Optional) Enable to use 64-bit mode. If the system is 64-bit capable and has at least 16 GB of RAM, then `auto-64-bit` will cause the Routing Engine to run in 64-bit mode. Otherwise, it will run in 32-bit mode.

> **NOTE**: This option is not applicable for logical systems.
> Starting in Junos 15.1F3, 15.1R2, 15.1R3, and 15.2R1, 64-bit mode is enabled by default on systems that support it and which have at least 16 GB of RAM.

`force-32-bit` (Optional) Enable to always use 32-bit mode.

> **NOTE**: For MX Series routers, virtual private LAN service (VPLS) dynamic profiles are not supported with the 64-bit mode routing protocol process (rpd). To enable VPLS dynamic profiles configuration, configure the routing process to use 32-bit mode.

`force-64-bit` (Optional) Enable to always use 64-bit mode.

> **TIP**: You need not restart the routing protocol process (rpd) to use the 64-bit mode. However, forcing rpd from 32-bit to 64-bit or 64-bit-to 32-bit will restart the rpd process, which can impact the routing protocols. For this reason, it is recommended to perform these changes in a maintenance window.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3 R4.

RELATED DOCUMENTATION

Disabling Junos OS Processes

processes | 13248

routing-engine-mode (TLB)

IN THIS SECTION

- Syntax | 15996
- Hierarchy Level | 15996
- Description | 15996
- Required Privilege Level | 15996
- Release Information | 15996

Syntax

```
routing-engine-mode;
```

Hierarchy Level

```
[edit services traffic-load-balance]
```

Description

Enable TLB heath-check on the RE.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

rp

IN THIS SECTION

- Syntax | 15997
- Hierarchy Level | 15999
- Description | 15999
- Default | 15999
- Required Privilege Level | 15999

Release Information | 15999

Syntax

```
rp {
    auto-rp {
        (announce | discovery | mapping);
        (mapping-agent-election | no-mapping-agent-election);
    }
    bidirectional {
        address address {
            group-ranges {
                destination-ip-prefix</prefix-length>;
            }
            hold-time seconds;
            priority number;
            }
        }
    bootstrap {
        family (inet | inet6) {
            export [ policy-names ];
            import [ policy-names ];
            priority number;
        }
    }
    bootstrap-export [ policy-names ];
    bootstrap-import [ policy-names ];
    bootstrap-priority number;
    dr-register-policy [ policy-names ];
    embedded-rp {
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
        maximum-rps limit;
    }
    group-rp-mapping {
        family (inet | inet6) {
            log-interval seconds;
            maximum limit;
```

```
            threshold value;
            }
        }
        log-interval seconds;
        maximum limit;
        threshold value;
        }
    }
    local {
        family (inet | inet6) {
            disable;
            address address;
            anycast-pim {
                local-address address;
                    address address <forward-msdp-sa>;
                    rp-set {
                }
            }
            group-ranges {
                destination-ip-prefix</prefix-length>;
            }
            hold-time seconds;
            override;
            priority number;
        }
    }
    register-limit {
        family (inet | inet6) {
            log-interval seconds;
            maximum limit;
            threshold value;
            }
        }
        log-interval seconds;
        maximum limit;
        threshold value;
        }
    }
    register-probe-time register-probe-time;
    }
    rp-register-policy [ policy-names ];
    static {
        address address {
```

```
            override;
            version version;
            group-ranges {
                destination-ip-prefix</prefix-length>;
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Configure the routing device as an actual or potential RP. A routing device can be an RP for more than one group.

The remaining statements are explained separately. See CLI Explorer.

Default

If you do not include the `rp` statement, the routing device can never become the RP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding PIM Sparse Mode

rp-register-policy

IN THIS SECTION

- Syntax | 16000
- Hierarchy Level | 16000
- Description | 16000
- Options | 16001
- Required Privilege Level | 16001
- Release Information | 16001

Syntax

```
rp-register-policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Apply one or more policies to control incoming PIM register messages.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Register Message Filters on a PIM RP and DR

dr-register-policy | 5938

rp-set

IN THIS SECTION

- Syntax | 16002
- Hierarchy Level | 16002
- Description | 16002
- Required Privilege Level | 16002
- Release Information | 16002

Syntax

```
rp-set {
    address address <forward-msdp-sa>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim local family (inet | inet6) anycast-pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
local family (inet | inet6) anycast-pim],
[edit protocols pim local family (inet | inet6) anycast-pim],
[edit routing-instances routing-instance-name protocols pim local family (inet | inet6) anycast-
pim]
```

Description

Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring PIM Anycast With or Without MSDP

rpc

IN THIS SECTION

- Syntax | 16003
- Hierarchy Level | 16003
- Description | 16003
- Options | 16003
- Required Privilege Level | 16004
- Release Information | 16004

Syntax

```
rpc {
    program-number rpc-program-number;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain protocol-binding],

[edit security idp custom-attack attack-name attack-type signature protocol-binding]
```

Description

Allow IDP to match the attack for a specified remote procedure call (RPC) program number.

Options

`program-number` *`rpc-program-number`*—RPC program number.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

rpc-program-number

IN THIS SECTION

- Syntax | 16004
- Hierarchy Level | 16004
- Description | 16004
- Options | 16005
- Required Privilege Level | 16005
- Release Information | 16005

Syntax

```
rpc-program-number number;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Remote procedure call (RPC) or Distributed Computing Environment (DCE) value.

Options

number—RPC or DCE program value.

- **Range:** 100,000 through 400,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

ALG Descriptions

Configuring an RPC Program Number

Examples: Configuring Application Protocols

Verifying the Output of ALG Sessions

rpc-program-number (Applications)

IN THIS SECTION

- Syntax | 16006
- Hierarchy Level | 16006
- Description | 16006
- Options | 16006
- Required Privilege Level | 16006
- Release Information | 16006

Syntax

```
rpc-program-number number;
```

Hierarchy Level

```
[edit applications application  application-name  ]
[edit applications application  application-name  term  term-name]
```

Description

Specify the remote procedure call (RPC) or Distributed Computing Environment (DCE) value.

Options

number—RPC or DCE program value.

- **Range:** 0 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policy Application Sets Overview

rpf-check

IN THIS SECTION

- Syntax (MX Series, SRX Series, M Series, T Series, PTX Series) | 16007
- Syntax (ACX7k Series) | 16007
- Syntax (EX Series and QFX Series) | 16008
- Hierarchy Level (MX Series, SRX Series, M Series, T Series, PTX Series) | 16008
- Hierarchy Level (EX Series and QFX Series) | 16008
- Description | 16008
- Default | 16009
- Options | 16009
- Required Privilege Level | 16009
- Release Information | 16009

Syntax (MX Series, SRX Series, M Series, T Series, PTX Series)

```
rpf-check {
    fail-filter filter-name;
    mode (strict | loose);
}
```

Syntax (ACX7k Series)

```
rpf-check {
    fail-filter filter-name;
    mode (loose);
}
```

Syntax (EX Series and QFX Series)

```
rpf-check;
```

Hierarchy Level (MX Series, SRX Series, M Series, T Series, PTX Series)

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit interfaces interface-name unit logical-unit-number family inet6],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet6]
```

Hierarchy Level (EX Series and QFX Series)

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit interfaces interface-name unit logical-unit-number family inet6]
```

Description

Enable a reverse-path forwarding (RPF) check on unicast traffic.

On EX3200 and EX4200 switches, enable a reverse-path forwarding (RPF) check on unicast traffic (except ECMP packets) on all ingress interfaces.

On EX4300 switches, enable a reverse-path forwarding (RPF) check on unicast traffic, including ECMP packets, on all ingress interfaces.

On EX8200 and EX6200 switches, enable an RPF check on unicast traffic, including ECMP packets, on the selected ingress interfaces.

On MX204 and ACX7k devices, `strict` is the default mode unless you over-ride it by setting the option to `loose`.

On QFX Series switches, enable an RPF check on unicast traffic on the selected ingress interfaces. ECMP packets are checked by QFX5000 Series switches only.

The mode statement is explained separately.

Default

Unicast RPF is disabled on all interfaces.

Options

`fail-filter`—A filter to evaluate when packets are received on the interface. If the RPF check fails, this optional filter is evaluated. If the fail filter is not configured, the default action is to silently discard the packet.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for interface `ps0` (pseudowire subscriber logical interface device) added in Junos OS Release 15.1.

RELATED DOCUMENTATION

- Configuring Unicast RPF Strict Mode
- Configuring Unicast RPF Loose Mode
- *Configuring a Pseudowire Subscriber Logical Interface Device*
- Example: Configuring Unicast RPF (On a Switch)

rpf-check (Dynamic Profiles)

IN THIS SECTION

- Syntax | **16010**
- Hierarchy Level | **16010**

- Description | 16010
- Required Privilege Level | 16011
- Release Information | 16011

Syntax

```
rpf-check {
    fail-filter filter-name;
    mode loose;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family]
```

Description

Reduce forwarding of IP packets that might be spoofing and address by checking whether traffic is arriving on an expected path that the sender would use to reach the destination. You can include this statement with the `inet` protocol family only. When the traffic passes the check, it is forwarded to the destination address; otherwise it is discarded. When you configure `rpf-check` alone, then unicast RPF is in strict mode, meaning that the check passes only when the packet's source address is in the FIB and the interface matches the routes RPF.

Starting in Junos OS Release 19.1, the `show interfaces statistics` *`logical-interface-name`* `detail` command displays unicast RPF statistics for dynamic logical interfaces when either `rpf-check` or `rpf-check mode loose` is enabled on the interface. No additional statistics are displayed when `rpf-check fail-filter` *`filter-name`* is configured on the interface. The `clear interfaces statistics` *`logical-interface-name`* command clears RPF statistics.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

> *Unicast RPF in Dynamic Profiles for Subscriber Interfaces*
>
> Configuring Unicast RPF Strict Mode

rpf-check-policy (Routing Options)

IN THIS SECTION

- Syntax | **16011**
- Hierarchy Level | **16012**
- Description | **16012**
- Options | **16012**
- Required Privilege Level | **16012**
- Release Information | **16012**

Syntax

```
rpf-check-policy policy;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-options multicast]
```

Description

Enable you to control whether a reverse path forwarding (RPF) check is performed for a source and group entry before installing a route in the multicast forwarding cache. This makes it possible to use point-to-multipoint LSPs to distribute multicast traffic to Protocol Independent Multicast (PIM) islands situated downstream from the egress routers of the point-to-multipoint LSPs.

Options

policy—Name of the RPF check routing policy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs

rpf-check-policy (Routing Options RPF)

IN THIS SECTION

Syntax | 16013

- Hierarchy Level | **16013**
- Description | **16013**
- Options | **16013**
- Required Privilege Level | **16013**
- Release Information | **16014**

Syntax

```
rpf-check-policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Description

Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.

Options

policy-names—Name of one or more multicast RPF check policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Example: Configuring RPF Policies

rpf-loose-mode-discard

IN THIS SECTION

- Syntax | 16014
- Hierarchy Level | 16014
- Description | 16015
- Options | 16015
- Required Privilege Level | 16015
- Release Information | 16015

Syntax

```
rpf-loose-mode-discard {
    family {
        inet;
        inet6;
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure unicast reverse path forwarding (unicast RPF) loose mode with the ability to discard packets with the source address pointing to the discard next hop.

Options

`inet`—IPv4 address family.

`inet6`—IPv6 address family.

Required Privilege Level

interface-control—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Unicast RPF (Routers)

rpf-selection

IN THIS SECTION

- Syntax | 16016
- Hierarchy Level | 16016
- Description | 16016
- Default | 16016
- Options | 16017
- Required Privilege Level | 16017
- Release Information | 16017

Syntax

```
rpf-selection {
    group group-address {
        sourcesource-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
    prefix-list prefix-list-addresses {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols pim]
[edit protocols pim]
```

Description

Configure the PIM RPF next-hop neighbor for a specific group and source for a VRF routing instance.

> **NOTE**: Starting in Junos OS 17.4R1, you can configure `rpf-selection` statement at the `[edit protocols pim]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Default

If you omit the `rpf-selection` statement, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows.

Options

source-address—Specific source address for the PIM group.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring PIM RPF Selection

rpf-vector (PIM)

IN THIS SECTION

- Syntax | 16018
- Hierarchy Level | 16018
- Description | 16018
- Options | 16018
- Required Privilege Level | 16018
- Release Information | 16019

Syntax

```
rpf-vector {
    policy (rpf-vector)[ policy-name];
}
```

Hierarchy Level

```
[edit dynamic-profiles name protocols pim],
[edit logical-systems name protocols pim],
[edit logical-systems name routing-instances name protocols pim],
[edit protocols pim],
[edit routing-instances name protocols pim]
```

Description

This feature provides a way for PIM source-specific multicast (SSM) to resolve Vector Type Length (TLV) for multicast in a seamless Multiprotocol Label Switching (MPLS) networks. In other words, it enables PIM to build multicast trees through an MPLS core. `rpf-vector` implements RFC 5496, Reverse Path Forwarding (RPF) Vector TLV .

When `rpf-vector` is enabled on an edge router that sends PIM join messages into the core, the join message includes a vector specifying the IP address of the next edge router along the path to the root of the multicast distribution tree (MDT). The core routers can then process the join message by sending it towards the specified edge router (i.e., toward the Vector). The address of the edge router serves as the RPF vector in the PIM join message so routers in the core can resolve the next-hop towards the source without the need for BGP in the core.

Only the IPv4 address family is supported.

Options

policy Create a filter policy to determine whether or not to apply `rpf-vector`.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

- show pim join | 23070
- show pim neighbors | 23106
- policy (rpf-vector)

rpm (Interfaces)

IN THIS SECTION

- Syntax | 16019
- Hierarchy Level | 16019
- Description | 16020
- Options | 16020
- Required Privilege Level | 16020
- Release Information | 16020

Syntax

```
rpm (client client | server server | twamp-client twamp-client | twamp-server twamp-server);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Associate an RPM or TWAMP client (router or switch that originates RPM or TWAMP probes) or RPM or TWAMP server with a specified interface.

> **NOTE**: The TWAMP client is applicable only for `si-` interfaces.

Options

client—Identifier for RPM client router or switch.

server—Identifier for RPM server.

twamp-client—Identifier for RPM TWAMP client router.

twamp-server—Identifier for RPM TWAMP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches

rpm-scale

IN THIS SECTION

Syntax | 16021

- Hierarchy Level | 16022
- Description | 16022
- Options | 16022
- Required Privilege Level | 16023
- Release Information | 16023

Syntax

```
rpm-scale {
    destination {
        interface interface-name.logical-unit-number;
        subunit-cnt subunit-cnt;
    }
    source {
        address-base ipv4-address-base;
        count ipv4-count;
        step ipv4-step;
    }
    source-inet6 {
        address-base ipv6-address-base;
        count ipv6-count;
        step ipv6-step;
    }
    target {
        address-base ipv4-address-base;
        count ipv4-count;
        step ipv4-step;
    }
    target-inet6 {
        address-base ipv6-address-base;
        count ipv6-count;
        step ipv6-step;
    }
    tests-count tests-count;
}
```

Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Description

Configure the generation of multiple IPv4 RPM tests for a probe owner. Starting in Junos OS Release 18.2R1, you can also configure the generation of multiple IPv6 RPM tests for a probe owner. Tests are generated for multiple combinations of source and target addresses, which are incremented based on your configuration. Tests are first generated for all the source addresses with the initial target address, then tests are generated for all the source addresses with the next available target address, and so on.

Options

interface-name.logical-unit-number	The services interface that is generating RPM probes and the logical unit number that is used for the first test that is generated.
ipv4-address-base	The IPv4 source or target address that is incremented to generate the addresses used in the RPM tests.
ipv4-count	The maximum number of IPv4 source or target addresses to use for the generated RPM tests.
ipv4-step	The amount to increment the IPv4 source or target address for each generated RPM test.
ipv6-address-base	The IPv6 source or target address that is incremented to generate the addresses used in the RPM tests.
ipv6-count	The maximum number of IPv6 source or target addresses to use for the generated RPM tests.
ipv6-step	The amount to increment the IPv6 source or target address for each generated RPM test.
subunit-cnt	The maximum number of logical units used by the services interface in the generated tests. The first generated test uses the logical unit specified in the *interface-name.logical-unit-number* option, and each successive test increments the logical unit number by one. Once the maximum number of logical units has been used, the next generated test cycles back to the logical unit that was used in the first test.

tests-count

The maximum number of RPM tests to generate. This number must be less than or equal to the number of generated source addresses multiplied by the number of generated target addresses.

- **Range:** 1 through 500,000 for probes generated on an MS-MPC or MS-MIC.
  1 through 2,000 for probes generated on the Routing Engine.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

rpm (Services)

IN THIS SECTION

- Syntax (Junos OS, except for QFX10000 switches, EX9200 switches, and SRX Series Firewalls) | **16024**
- Syntax (Junos OS, for EX9200 and QFX10000 switches) | **16029**
- Syntax (Junos OS, for SRX300 and SRX550HM devices) | **16033**
- Syntax (Junos OS Evolved) | **16034**
- Junos OS Hierarchy Level | **16035**
- Junos OS Evolved Hierarchy Level | **16035**
- Description | **16035**
- Required Privilege Level | **16036**
- Release Information | **16036**

Syntax (Junos OS, except for QFX10000 switches, EX9200 switches, and SRX Series Firewalls)

```
rpm {
    bgp {
        data-fill data;
        data-size size;
        destination-port port;
        history-size size;
        logical-system logical-system-name [routing-instances routing-instance-name];
        moving-average-size number;
        probe-count count;
        probe-interval seconds;
        probe-type type;
        routing-instances instance-name;
        test-interval interval;
        ttl ttl;
    }
    probe owner {
        delegate-probes;
        test test-name {
            data-fill data;
            data-size size;
            destination-interface interface-name;
            destination-port port;
            dscp-code-points dscp-bits;
            hardware-timestamp;
            history-size size;
            inet6-options {
                source-address source-address;
            }
            moving-average-size number;
            next-hop next-hop;
            one-way-hardware-timestamp;
            probe-count count;
            probe-interval seconds;
            probe-type type;
            routing-instance instance-name;
            rpm-scale {
                target {
                    address-base address-base;
                    count count;
```

```
                    step step;
                }
                target-inet6 {
                    address-base address-base;
                    count count;
                    step step;
                }
                source {
                    address-base address-base;
                    count count;
                    step step;
                }
                source-inet6 {
                    address-base address-base;
                    count count;
                    step step;
                }
                destination {
                    interface interface;
                    subunit-cnt subunit-cnt;
                }
                tests-count tests-count;
            }
            source-address address;
            target (url url | address address);
            test-interval interval;
            thresholds thresholds;
            traps traps;
            ttl ttl;
        }
    }
    probe-server {
        icmp {
            destination-interface destination-interface;
        }
        tcp {
            destination-interface interface-name;
            port number;
        }
        udp {
            destination-interface interface-name;
            port number;
        }
```

```
    }
    probe-limit limit;
    rfc2544-benchmarking {
        profiles {
            test-profile profile-name {
                test-type (throughput | latency | frame-loss | back-back-frames);
                packet-size bytes;
                step-percent percent;
                bandwidth-kbps kpbs;
            }
        }
        tests {
            test-name test-name {
                destination-ipv4-address address;
                destination-mac-address destination-mac-address;
                destination-udp-port  port-number;
                direction (egress | ingress);
                disable-signature-check;
                dscp-code-points dscp-code-points;
                family (bridge| inet | ccc);
                forwarding-class forwarding-class;
                halt-on-prefix-down;
                in-service;
                ip-swap;
                ivlan-cfi ivlan-cfi;
                ivlan-id ivlan-id;
                ivlan-priority ivlan-priority;
                mode reflect;
                outer-tag-protocol-id outer-tag-protocol-id;
                ovlan-cfi ovlan-cfi;
                ovlan-id ovlan-id;
                ovlan-priority  ovlan-priority;
                packet-loss-priority (high | low | medium-high);
                receive-failure-threshold receive-failure-threshold;
                reflect-etype reflect-etype;
                reflect-mode (mac-rewrite | mac-swap | no-mac-swap);
                reflector-port reflector-port;
                service-type (elan | eline);
                skip-arp-iteration;
                source-ipv4-address address;
                source-mac-address source-mac-address;
                source-udp-port port-number;
                test-finish-wait-duration test-finish-wait-duration;
```

```
                test-interface interface-name;
                test-iterator-duration test-iterator-duration;
                test-iterator-pass-threshold test-iterator-pass-threshold;
                test-profile test-profile;
                timestamp-format (microseconds | nanoseconds);
                transmit-failure-threshold transmit-failure-threshold;
                udp-tcp-port-swap;
                vlan-cfi vlan-cfi;
                vlan-id vlan-id;
                vlan-priority vlan-priority;
            }
        }
    }
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
    twamp {
        client {
            control-connection  name {
                authentication-mode none;
                control-type (managed | light);
                destination-interface  destination-interface;
                destination-port    destination-port;
                history-size    history-size;
                moving-average-size   moving-average-size;
                persistent-results;
                routing-instance   routing-instance;
                target-address target-address;
                tcp-keepcnt   count;
                 tcp-keepidle    seconds;
                 tcp-keepintvl    seconds;
                 test-count   test-count;
                 test-interval   seconds;
                 traps {
                    control-connection-closed;
                    test-iteration-done;
                }
                test-session name {
                    data-fill-with zeros;
```

```
                    data-size   data-size;
                    dscp-code-points dscp-code-points;
                    probe-count  probe-count;
                    probe-interval  seconds;
                    source-address source-address;
                    target-address target-address local-link IPv6-link-local-interface-
name;
                    thresholds {
                        egress-time microseconds;
                        ingress-time microseconds;
                        jitter-egress microseconds;
                        jitter-ingress microseconds;
                        jitter-rtt microseconds;
                        max-rtt microseconds;
                        rtt microseconds;
                        std-dev-egress microseconds;
                        std-dev-ingress microseconds;
                        std-dev-rtt microseconds;
                        successive-loss successive-loss;
                        total-loss total-loss;
                    }
                    traps {
                        egress-jitter-exceeded;
                        egress-std-dev-exceeded;
                        egress-time-exceeded;
                        ingress-jitter-exceeded;
                        ingress-std-dev-exceeded;
                        ingress-time-exceeded;
                        jitter-exceeded;
                        max-rtt-exceeded;
                        probe-failure;
                        rtt-exceeded;
                        std-dev-exceeded;
                        test-completion;
                        test-failure;
                    }
                    ttl hop-count;
                }
            }
        }
        post-cli-implicit-firewall;
        server {
            authentication-key-chain name {
```

```
                key-id name {
                    secret secret;
                }
            }
            authentication-mode <authenticated> <control-only-encrypted> <encrypted>
<none>;
            client-list  {
                address address <routing-instance [instance-name...]>;
            }
            max-connection-duration  hours;
            maximum-connections  maximum-connections;
            maximum-connections-per-client  maximum-connections-per-client;
            maximum-sessions  maximum-sessions;
            maximum-sessions-per-connection  maximum-sessions-per-connection;
            port  port;
            routing-instance-list  name {
                port port;
            }
            server-inactivity-timeout  minutes;
            tcp-keepcnt  count;
            tcp-keepidle  seconds;
            tcp-keepintvl seconds;
        }
    }
}
```

Syntax (Junos OS, for EX9200 and QFX10000 switches)

```
rpm {
    bgp {
        data-fill data;
        data-size size;
        destination-port port;
        history-size size;
        logical-system logical-system-name [routing-instances routing-instance-name];
        moving-average-size number;
        probe-count count;
        probe-interval seconds;
        probe-type type;
        routing-instances instance-name;
        test-interval interval;
```

```
        ttl ttl;
    }
    probe owner {
        delegate-probes;
        test test-name {
            data-fill data;
            data-size size;
            destination-interface interface-name;
            destination-port port;
            dscp-code-points dscp-bits;
            hardware-timestamp;
            history-size size;
            inet6-options {
                source-address source-address;
            }
            moving-average-size number;
            next-hop next-hop;
            one-way-hardware-timestamp;
            probe-count count;
            probe-interval seconds;
            probe-type type;
            routing-instance instance-name;
            rpm-scale {
                target {
                    address-base address-base;
                    count count;
                    step step;
                }
                target-inet6 {
                    address-base address-base;
                    count count;
                    step step;
                }
                source {
                    address-base address-base;
                    count count;
                    step step;
                }
                source-inet6 {
                    address-base address-base;
                    count count;
                    step step;
                }
```

```
                destination {
                    interface interface;
                    subunit-cnt subunit-cnt;
                }
                tests-count tests-count;
            }
            source-address address;
            target (url url | address address);
            test-interval interval;
            thresholds thresholds;
            traps traps;
            ttl ttl;
        }
    }
    probe-server {
        icmp {
            destination-interface destination-interface;
        }
        tcp {
            destination-interface interface-name;
            port number;
        }
        udp {
            destination-interface interface-name;
            port number;
        }
    }
    probe-limit limit;
    twamp {
        client {
            control-connection  name {
                authentication-mode none;
                destination-interface  destination-interface;
                destination-port    destination-port;
                history-size    history-size;
                moving-average-size   moving-average-size;
                persistent-results;
                routing-instance   routing-instance;
                target-address target-address;
                tcp-keepcnt   count;
                tcp-keepidle    seconds;
                tcp-keepintvl    seconds;
                test-count   test-count;
```

```
test-interval   seconds;
traps {
        control-connection-closed;
        test-iteration-done;
}
test-session name {
    data-fill-with zeros;
    data-size   data-size;
    dscp-code-points dscp-code-points;
    probe-count  probe-count;
    probe-interval  seconds;
    target-address target-address;
    thresholds {
        egress-time microseconds;
        ingress-time microseconds;
        jitter-egress microseconds;
        jitter-ingress microseconds;
        jitter-rtt microseconds;
        max-rtt microseconds;
        rtt microseconds;
        std-dev-egress microseconds;
        std-dev-ingress microseconds;
        std-dev-rtt microseconds;
        successive-loss successive-loss;
        total-loss total-loss;
    }
    traps {
        egress-jitter-exceeded;
        egress-std-dev-exceeded;
        egress-time-exceeded;
        ingress-jitter-exceeded;
        ingress-std-dev-exceeded;
        ingress-time-exceeded;
        jitter-exceeded;
        max-rtt-exceeded;
        probe-failure;
        rtt-exceeded;
        std-dev-exceeded;
        test-completion;
        test-failure;
    }
    ttl hop-count;
}
```

```
            }
        }
        post-cli-implicit-firewall;
        server {
            authentication-key-chain name {
                key-id name {
                    secret secret;
                }
            }
            authentication-mode <authenticated   > <control-only-encrypted> <encrypted   >
<none>;
            client-list  {
                address address <routing-instance [instance-name...]>;
            }
            max-connection-duration  hours;
            maximum-connections  maximum-connections;
            maximum-connections-per-client  maximum-connections-per-client;
            maximum-sessions  maximum-sessions;
            maximum-sessions-per-connection  maximum-sessions-per-connection;
            port  port;
            routing-instance-list  name {
                port port;
            }
            server-inactivity-timeout  minutes;
            tcp-keepcnt  count;
            tcp-keepidle  seconds;
            tcp-keepintvl seconds;
        }
    }
}
```

Syntax (Junos OS, for SRX300 and SRX550HM devices)

```
rpm {
    rfc2544-benchmarking {
        tests {
            test-name test-name {
                destination-ipv4-address address;
                destination-udp-port  port-number;
                disable-signature-check;
                family inet
```

```
                mode reflect;
                source-ipv4-address address;
                source-udp-port port-number;
                test-interface interface-name;
                test-iterator-duration test-iterator-duration;
            }
        }
    }
}
```

Syntax (Junos OS Evolved)

```
rpm {
    owner name {
        test test-name {
            data-fill data;
            data-size size;
            destination-port port;
            dscp-code-points dscp-bits;
            history-size size;
            moving-average-size number;
            offload-type {
                none;
                pfe-timestamp;
            }
            probe-count count;
            probe-interval seconds;
            probe-server {
                icmp;
                tcp {
                    port number;
                }
                udp {
                    port number;
                }
            }
            probe-type type;
            routing-instance instance-name;
            source-address address;
            target (url url | address address);
            test-interval interval;
```

```
            thresholds thresholds;
            traps traps;
            ttl hop-count
        }
    }
}
```

Junos OS Hierarchy Level

```
[edit services]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring]
```

Description

Configure real-time performance monitoring (RPM). RPM enables you to configure active probes to track and monitor traffic. Probes collect packets per destination and per application, including PING Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets. RPM provides Management Information Base (MIB) support with extensions for RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*.

Starting in Junos OS Evolved Release 20.1R1, you can configure RPM probes. For Junos OS Evolved, RPM is configured at the `[edit services monitoring rpm]` hierarchy level. The scope of support is limited to:

- Probe generation and reception (client) as well as reflection (server) for the following RPM probe types:
  - icmp-ping
  - icmp-timestamp
  - udp-ping
  - udp-timestamp
- Probe history management

- Reporting through syslog only

Starting in Junos OS Evolved 20.3R1, you can configure TWAMP probes. Starting in Junos OS Evolved 21.1R1, you can configure RFC 2544 benchmarking tests. For Junos OS Evolved, TWAMP is configured at the `[edit services monitoring twamp]` hierarchy level and RFC 2544 benchmarking tests are configured at the `[edit services monitoring rfc2544]` hierarchy level.

Starting in Junos OS Evolved Release 21.2R1, reporting through SNMP MIB objects is supported for RPM.

Starting in Junos OS Evolved Release 23.4R1, you can configure `tcp-ping`, `http-get`, and `http-metadata-get` probe types.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

`traps` option introduced in Junos OS Evolved Release 21.2R1.

`source-address` option and the `local-link` sub-option of the `target-address` option for TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

`tcp` sub-option for the `probe-server` option introduced in Junos OS Evolved Release 23.4R1.

`tcp-ping`, `http-get`, and `http-metadata-get` probe types introduced in Junos OS Evolved Release 23.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
23.4R1 Evo	Starting in Junos OS Evolved Release 23.4R1 for the ACX7024, ACX7100, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, and QFX5700, you can configure `tcp-ping`, `http-get`, and `http-metadata-get` probe types for RPM.
21.4R1	Starting in Junos OS Release 21.4R1, you can configure RPM and TWAMP probes for EX9200 switches.
21.4R1	Starting in Junos OS Release 21.4R1, for TWAMP Light test-sessions on MX and PTX routers, you can configure IPv6 link-local target addresses and their corresponding source addresses.
21.3R1	Starting in Junos OS 21.3R1, you can configure RPM and TWAMP probes for QFX10000 Series switches.
21.2R1 Evo	Starting in Junos OS Evolved Release 21.2R1, reporting through SNMP MIB objects is supported for RPM.
21.1R1 Evo	Starting in Junos OS Evolved 21.1R1, you can configure RFC 2544 benchmarking tests.
20.3R1 Evo	Starting in Junos OS Evolved 20.3R1, you can configure TWAMP probes.
20.1R1 Evo	Starting in Junos OS Evolved Release 20.1R1, you can configure RPM probes. For Junos OS Evolved, RPM is configured at the [edit services monitoring rpm] hierarchy level.

RELATED DOCUMENTATION

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Configuring BGP Neighbor Discovery Through RPM

rpm-tracking

IN THIS SECTION

- Syntax | 16038
- Hierarchy Level | 16038
- Description | 16038
- Options | 16039
- Required Privilege Level | 16040
- Release Information | 16040

Syntax

```
rpm-tracking {
    route destination-prefix {
        metric metric;
        next-hop next-hop {;
           rpm-probe name rpm-test test-name;
        }
        preference preference;
        tag tag;
    }
 }
```

Hierarchy Level

```
[edit routing-instances name routing-options],
[edit routing-options]
```

Description

RPM-tracked static routes are coupled with a given RPM test instance. Routes can be installed or removed according to the results of the given RPM test. When installed, these routes are automatically

given a preference of 1, and so are preferred over static routes that may already exist with the same prefix.

If the RPM test result is “success,” then all the RPM-tracked static routes that match the probe owner and test name of the successful test are added to the routing table. If the test result is “failure,” then all the RPM-tracked static routes that match the probe owner and test name of the failed test are removed, if present, from the routing table.

RPM route tracking supports both IPv4 and IPv6 routes. RPM-tracked static routes are configured individually; wildcards, ranges, and regular expressions are not supported.

Options

route ***destination-prefix***
(Required) Configure an RPM-tracked static route. Must be a IPv4 or IPv6 destination prefix.

next-hop ***next-hop***
(Required) Configure a next-hop address. Must be a IPv4 or IPv6 address. You can configure up to 16 multiple paths (next-hops) for any given RPM-tracked static route (RPM-tracked static routes with multiple next-hops can also be configured inside a routing instance).

metric ***metric***
(Optional) Set the route metric for the destination prefix. The route with the lowest metric is active in the routing table.

- **Default:** 1
- **Range:** 1 through 16.

preference ***preference***
(Optional) Set the route preference value for the destination prefix. The route with the lowest preference value is active in the routing table. Qualified next hops are not supported with RPM-tracked static routes. Hence, the setting of the `preference`, `metric`, and `tag` options apply only to the RPM-tracked static route and not to the related next hops.

- **Default:** 1

tag ***tag***
(Optional) Set the route tag value for the destination prefix. The route with the lowest tag value is active in the routing table.

- **Default:** 0

rpm-probe ***name***
(Required) Must be a valid RPM probe owner from the `[edit services rpm]` configuration hierarchy.

rpm-test *test-name* (Required) Must be a valid RPM test name from the `[edit services rpm]` configuration hierarchy.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Support introduced in Junos OS Release 18.4 R1.

`next-hop` option added in Junos OS Release 19.1R1.

`preference` and `tag` options added in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

show route rpm-tracking | **23759**

rpt-spt

IN THIS SECTION

- Syntax | **16041**
- Hierarchy Level | **16041**
- Description | **16041**
- Required Privilege Level | **16041**
- Release Information | **16041**

Syntax

```
rpt-spt;
```

Hierarchy Level

```
[edit logical-systems profile-name routing-instances instance-name protocols mvpn mvpn-mode],
[edit routing-instances instance-name protocols mvpn mvpn-mode]
```

Description

Use rendezvous-point trees for customer PIM (C-PIM) join messages, and switch to the shortest-path tree after the source is known.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

rsh

IN THIS SECTION

- Syntax | 16042
- Hierarchy Level | 16042
- Description | 16042
- Options | 16042
- Required Privilege Level | 16043

Release Information | 16043

Syntax

```
rsh {
    disable;
    traceoptions {
        flag {
            all {
                extensive
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Configure rsh to specify the remote shell (RSH) ALG on the device. The Remote Shell (RSH) provides a conduit to execute commands on a remote host.

Options

- `disable`—Disable the RSH ALG. By default, the RSH ALG is disabled.
- `traceoptions`—Configure RSH ALG tracing options.
  - `flag`—Trace operation to perform.
    - `all`—Trace all events.

- *extensive*—Display extensive amount of data.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding the RSH ALG

rst-invalidate-session

IN THIS SECTION

- Syntax | 16043
- Hierarchy Level | 16044
- Description | 16044
- Required Privilege Level | 16044
- Release Information | 16044

Syntax

```
rst-invalidate-session;
```

Hierarchy Level

```
[edit security flow tcp-session]
```

Description

Enable the device to mark a session for immediate termination when it receives a TCP reset (RST) message. By default, this feature is disabled.

> **NOTE**: On SRX Series Firewalls with `rst-invalidate-session` configured the invalidation of the session occurs immediately whereas without `rst-invalidate-session` configured the session is set to time out 2 seconds after a TCP reset (RST) message has been received.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

rst-sequence-check

IN THIS SECTION

- Syntax | 16045
- Hierarchy Level | 16045
- Description | 16045

Required Privilege Level | 16045
Release Information | 16045

Syntax

```
rst-sequence-check;
```

Hierarchy Level

```
[edit security flow tcp-session]
```

Description

Verify that the TCP sequence number in a TCP segment with the RST bit enabled matches the previous sequence number for a packet in that session or is the next higher number incrementally. If the sequence number does not match either of these expected numbers, the device drops the packet and sends the host a TCP ACK message with the correct sequence number. By default, this check is disabled.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

rstp

IN THIS SECTION

- MX Series | 16046
- EX Series | 16047
- ELS Versions: EX Series, QFX Series, NFX Series | 16047
- Hierarchy Level | 16048
- Description | 16048
- Default | 16049
- Required Privilege Level | 16049
- Release Information | 16049

MX Series

```
rstp {
    bpdu-block-on-edge;
    bpdu-destination-mac-address provider-bridge-group;
    bridge-priority priority;
    extended-system-id;
    force-version stp;
    forward-delay seconds;
    hello-time seconds;
    max-age seconds;
    interface interface-name {
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode (p2p | shared);
        no-root-port;
        priority interface-priority;
    }
    priority-hold-time seconds;
```

```
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

EX Series

```
rstp   {
                            bpdu-block-on-edge;
                            bridge-priority priority;                    priority;
                            disable;
                            forward-delay                      seconds;
                            hello-time                       seconds;
                                interface (all | interface-name) {
        arp-on-stp;
        bpdu-timeout-action {

            block;
            log;
        }
                                    cost                          cost;
                                    edge;
                                    mode                          mode;
                                    no-root-port;
                                    priority                          priority;
    }
                            max-age                      seconds;
                                traceoptions {
        file filename <files                          number > <size size> <no-stamp | no-world-
readable | world-readable>;
        flag flag;
    }
}
```

ELS Versions: EX Series, QFX Series, NFX Series

```
rstp {
    bpdu-block-on-edge;
```

```
    bpdu-destination-mac-address provider-bridge-group;
    bridge-priority priority;
    disable;
    extended-system-id;
    force-version stp;
    forward-delay seconds;
    hello-time seconds;
    max-age seconds;
    priority-hold-time seconds;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
```

```
[edit protocols],
```

```
[edit routing-instances routing-instance-name protocols]
```

Description

Configure Rapid Spanning Tree Protocol (RSTP). RSTP is defined in the IEEE 802.1D-2004 specification and is used to prevent loops in Layer 2 networks, which results in shorter convergence times than those provided by basic Spanning Tree Protocol (STP).

VSTP and RSTP can be configured concurrently. You can selectively configure up to 253 VLANs using VSTP; the remaining VLANs will be configured using RSTP. VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on the switch.

BEST PRACTICE: Configure RSTP when you configure VSTP. RSTP overhead is minimal and this configuration ensures that a spanning-tree protocol is running on all VLANs on your switch, even when your switch is supporting more than 253 VLANs.

The remaining statements are explained separately. See CLI Explorer.

Default

RSTP is enabled on all Ethernet switching interfaces.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

`bpdu-block-on-edge` statement added in Junos OS Release 9.4.

`bpdu-timeout-action` statement added in Junos OS Release 9.4.

Support for logic systems added in Junos OS Release 9.6.

Statement updated in Junos OS Release 15.1 for EX Series and QFX Series switches to support configuration of spanning tree parameters globally on all interfaces.

NOTE: You cannot disable spanning tree parameters globally on all interfaces.

RELATED DOCUMENTATION

Configuring Rapid Spanning Tree Protocol

Understanding RSTP

Configuring RSTP on EX Series Switches (CLI Procedure)

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

rsvp

IN THIS SECTION

- Syntax | 16050
- Hierarchy Level | 16053
- Description | 16053
- Options | 16053
- Default | 16055
- Required Privilege Level | 16055
- Release Information | 16056

Syntax

```
rsvp {
    associated-bidirectional-lspsingle-sided-provisioning;
    authentication-key (Protocols RSVP) authentication-key;
    cross-credibility-cspf;
    (disable | enable);
    expand-flood-reflector-hop;
    fast-rerouteoptimize-timer seconds;
    graceful-deletion-timeout seconds;
    graceful-restart (Protocols RSVP) <(disable | enable)> <helper-disable> <maximum-helper-
recovery-time seconds> <maximum-helper-restart-time seconds>;
    hello-acknowledgements;
    interface (Protocols RSVP) name {
        authentication-key authentication-key;
        bandwidth (Protocols RSVP) bandwidth;
        (disable | enable);
        hello-interval (Protocols RSVP) seconds;
        link-protection (RSVP) {
            admin-groupinclude-any [include-any ...] include-all [include-all ...] exclude
[exclude ...];
            bandwidth <ct0 bps> <ct1 bps> <ct2 bps> <ct3 bps> <per-traffic-class-bandwidth>;
            bypass (Signaled LSP) bypass name {
```

```
                admin-groupinclude-any [include-any ...] include-all [include-all ...] exclude
[exclude ...];
                bandwidth <ct0 bps> <ct1 bps> <ct2 bps> <ct3 bps> <per-traffic-class-bandwidth>;
                class-of-service (Protocols RSVP) class-of-service;
                description description;
                exclude-srlg;
                hop-limit hop-limit;
                no-cspf;
                path (Protocols RSVP) name;
                priority (Protocols RSVP) {
                   reservation-priority;
                   setup-priority;
                }
                subscription <subscription>;
                to to;
             }
             class-of-service class-of-service;
             (disable | enable);
             exclude-srlg;
             hop-limit hop-limit;
             max-bypasses max-bypasses;
             no-cspf (Protocols RSVP);
             no-node-protection;
             optimize-timer (Protocols RSVP) seconds;
             path (Protocols RSVP) name;
             priority (Protocols RSVP) {
             reservation-priority;
             setup-priority;
             }
             subscription subscription;
          }
          non-rsvp-bandwdith {
             local-bw-override-threshold;
          }
          (reliable | no-reliable);
          subscription <ct0 ct0> <ct1 ct1> <ct2 ct2> <ct3 ct3> <link-subscription> priority-value
percent percent;
          update-threshold <adaptive name threshold-percent percent threshold-value threshold-
value> <threshold-percent> <threshold-value threshold-value>;
          update-threshold-max-reservable {
             bandwidth;
             percent percent;
          }
```

```
    }
    keep-multiplier keep-multiplier;
    load-balancebandwidth;
    local-reversion;
    lsp-set name {
       (disable | enable);
       match-criteria(egress | ingress | transit);
       traceoptions {
           file filename <files files> <size size> <(world-readable | no-world-readable)>;
           flag (all | enhanced-frr | error | event | io-event | io-packets | lmp | lsp-prefix
| nsr-synchronization | packets | path | pathtear | resv | resvtear | route | state) {
               detail;
               disable;
               receive;
               send;
           }
       }
    }
    no-authentication-check (Protocols RSVP);
    node-hello <hello-interval seconds>;
    no-enhanced-frr-bypass (Protocols RSVP);
    no-hello-acknowledgements;
    no-interface-hello;
    no-local-reversion;
    no-node-hello;
    no-node-hello-on-bypass;
    no-node-id-subobject;
    no-p2mp-re-merge;
    no-p2mp-sublsp;
    peer-interface (Protocols RSVP) name {
       authentication-key authentication-key;
       dynamic-bidirectional-transport <template template>;
       (disable | enable);
       hello-interval (Protocols RSVP) seconds;
       (reliable | no-reliable);
    }
    pop-and-forward (Protocols RSVP) <application-labeldepth depth>;
    preemption (aggressive | disabled | normal);
    refresh-time refresh-time;
    rfc6510-lsp-attributes;
    setup-protection;
    traceoptions (Protocols RSVP) {
       file filename <files files> <size size> <(world-readable | no-world-readable)>;
```

```
        flag (all | enhanced-frr | error | event | io-event | io-packets | lmp | lsp-prefix |
nsr-synchronization | packets | path | pathtear | resv | resvtear | route | state) {
            detail;
            disable;
            receive;
            send;
        }
     }
     tunnel-services (RSVP) {
        devices [devices ...];
     }
 }
```

Hierarchy Level

```
[edit logical-systems name protocols],
[edit logical-systems name routing-instances name protocols],
[edit logical-systems name tenants name routing-instances name protocols],
[edit protocols],
[edit routing-instances name protocols],
[edit tenants name routing-instances name protocols]
```

Description

Enable Resource Reservation Protocol (RSVP) signaling.

You must include the `rsvp` statement in the configuration to enable RSVP on the router.

The primary purpose of RSVP in Junos OS for EX Series switches is to support dynamic signaling within label switched paths (LSPs).

Options

`cross-credibility-cspf`—Enable path computation across credibility levels. The constraint path computation is run across multi-protocol links and nodes, instead of a credibility-by-credibility basis.

`authentication key`—Authentication key. Neighboring routers use the password to verify the authenticity of packets sent from this interface or peer interface. It is recommended to use the `authentication-key` configuration at the `[edit protocols rsvp]` hierarchy level for the RSVP node-neighbor.

`graceful deletion timeout`—Specify the time, in seconds, before completing graceful deletion of signaling.

`graceful restart`—Configure `graceful restart` on the router. You must configure the `graceful-restart` statement at the `[edit routing-options]` hierarchy level to enable graceful restart on the router.

`hello acknowledgements`—Enable hello messages from nonsession neighbors to be acknowledged with a hello acknowledgment message. Once `hello acknowledgments` are enabled, the router continues to acknowledge hello messages from any nonsession RSVP neighbors unless the interface itself goes down or the configuration is changed by an administrator.

`interface`—Enable RSVP on one or more router interfaces.

`bandwidth`—Enables you to specify the actual available bandwidth.

`hello interval`—Enable the sending of hello packets on the interface.

`link protection`—Using link protection, you can configure a network to reroute traffic quickly around broken links.

`priority`—Configure the setup priority and reservation priority for a bypass LSP.

`exclude-srlg`—Exclude Shared Risk Link Group (SRLG) links for the secondary path for critical links where it is imperative to keep the secondary and primary label-switched paths completely disjoint from any common SRLG.

`hop-limit`—Specify the maximum number of routers that an LSP can traverse. Maximum number of hops.

`bypass`—Enables you to configure specific bandwidth and path constraints for a bypass LSP.

`max bypasses`—Specify the maximum number of dynamic bypass LSPs permitted for protecting this interface.

`no cspf`—Disable CSPF computation on all bypass LSPs or on a specific bypass LSP. You need to disable CSPF for link protection to function properly on interarea paths.

`optimize-timer`—The optimize timer initiates a periodic optimization process that reshuffles data LSPs among bypass LSPs to achieve the most efficient use of network resources.

`path`—Configure an explicit path (a sequence of strict or loose routes) to control where and how a bypass LSP is established.

`update threshold`—Adjust the threshold at which a change in bandwidth triggers an interior gateway protocol (IGP) update.

`keep multiplier`—Indicates the number of messages that can be lost before a particular state is declared stale and must be deleted. The keep multiplier directly affects the lifetime of an RSVP state.

`no authentication check`—Skip authentication check for received messages.

`node hello`—Enable node-ID based RSVP hellos globally on all of the RSVP interfaces on the router to allow Juniper Networks routers to interoperate with the equipment of other vendors.

`peer interface`—Configure the name of the LMP peer device.

`no-enhanced-frr-bypass`—Enable `no-enhanced-frr-bypass` to turn off all Fast reroute (FRR) facility protection enhancements, which includes improved LSP scaling and enhanced RSVP message handling, and reduce the default refresh time to 30 seconds.

`no-interface-hello`—Explicitly disable RSVP interface hellos globally on the router.

`no-local-reversion`—Disable RSVP local revertive mode as specified in RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels.

`no-node-hello-on-bypass`—Starting with Junos release 16.1, Junos routers send RSVP hello packets over a bypass LSP (when one is present), instead of the IGP next hop. To return to the original behavior specify the `no-node-hello-on-bypass` option.

`no-node-id-subobject`—Disable the record route object (RRO) node-ID subobject for compatibility with earlier versions of Junos OS.

`pop-and-forward`—Specify RSVP pop-and-forward LSP tunnel-specific global parameters.

`refresh time`—Periodical refresh timer expressed in seconds.

`setup-protection`—The facility-backup fast reroute mechanism can provide setup protection for LSPs which are in the process of being signaled.

`traceoptions`—The default RSVP-level trace options are those inherited from the routing protocols `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

`tunnel services`—The Junos OS selects one of the available virtual tunnel (VT) interfaces to de-encapsulate the egress traffic.

Default

RSVP is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Minimum RSVP Configuration

Example: Configuring MPLS on EX8200 and EX4500 Switches

Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect

Configuring MPLS on EX8200 and EX4500 Provider Switches

rsvp-error-hold-time

IN THIS SECTION

- Syntax | 16056
- Hierarchy Level | 16056
- Description | 16057
- Options | 16057
- Required Privilege Level | 16057
- Release Information | 16057

Syntax

```
rsvp-error-hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations. The more time you configure, the more time a source node (ingress of an RSVP LSP) can have to learn about the failures of its LSP by monitoring PathErr messages transmitted from downstream nodes.

Information from the PathErr messages is incorporated into subsequent LSP computations, which can improve the accuracy and speed of LSP setup. Some PathErr messages are also used to update traffic engineering database bandwidth information, reducing inconsistencies between the database and the network.

Options

seconds—Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations.

- **Range:** 0 through 240 seconds
- **Default:** 25 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages

rsvp-te (Protocols MVPN)

IN THIS SECTION

- Syntax | 16058
- Hierarchy Level | 16058
- Description | 16058
- Options | 16058
- Required Privilege Level | 16059
- Release Information | 16059

Syntax

```
rsvp-te{
    label-switched-path-template (Multicast){
        (default-template | lsp-template-name);
    }
    static-lsp static-lsp;
}
```

Hierarchy Level

```
[edit protocols mvpn inter-region-template template template-name all-regions],
[edit protocols mvpn inter-region-template template template-name region region-name]
```

Description

Configure to unicast, broadcast, and multicast traffic flooding using point-to-multipoint LSPs.

Options

static-lsp *static-lsp* Create a static point-to-multipoint LSP and automatically include all of the neighbors in the routing instance.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview

all-regions | **2157**

inter-region | **8789**

inter-region-template | **8793**

inter-region-segmented | **8791**

region | **14262**

template(Protocols MVPN) | **29092**

Example: Configuring Segmented Inter-Area P2MP LSP

Configuring Segmented Inter-Area P2MP LSP

rsvp-te (Routing Instances Provider Tunnel)

IN THIS SECTION

- Syntax | **16060**
- Hierarchy Level | **16060**
- Release Information | **16060**

Syntax

```
rsvp-te {
    label-switched-path-template {
        (default-template | lsp-template-name);
    }
    static-lsp lsp-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel],
[edit routing-instances routing-instance-name provider-tunnel]
[edit routing-instances instance-name provider-tunnel]
```

Release Information

Statement introduced in Junos OS Release 8.3.

Options

`static-lsp` *`lsp-name`*—Create a static point-to-multipoint LSP and automatically include all of the neighbors in the VPLS routing instance.

The remaining option is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Flooding Unknown Traffic Using Point-to-Multipoint LSPs in VPLS

rsvp-te (Routing Instances Provider Tunnel Selective)

IN THIS SECTION

- Syntax | 16061
- Hierarchy Level | 16061
- Description | 16062
- Required Privilege Level | 16062
- Release Information | 16062

Syntax

```
rsvp-te {
    label-switched-path-template {
        (default-template | lsp-template-name);
    }
    static-lsp lsp-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group address source source-address],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group address wildcard-source],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet wildcard-source],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet6 wildcard-source],
[edit routing-instances routing-instance-name provider-tunnel],
[edit routing-instances routing-instance-name provider-tunnel selective group address source
source-address],
[edit routing-instances routing-instance-name provider-tunnel selective group address wildcard-
```

```
source],
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet
wildcard-source],
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet6
wildcard-source]
```

Description

Configure the properties of the RSVP traffic-engineered point-to-multipoint LSP for MBGP MVPNs.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: Junos OS Release 11.2 and earlier do not support point-to-multipoint LSPs with next-generation multicast VPNs on MX80 routers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

rsvp-te (Routing Options)

IN THIS SECTION

- Syntax | 16063
- Hierarchy Level | 16063

- Description | 16063
- Options | 16063
- Required Privilege Level | 16063
- Release Information | 16064

Syntax

```
rsvp-te entry-name {
    destination-networks network-prefix;
    label-switched-path-template (Multicast) {
        default-template;
        template-name;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name],
[edit routing-options dynamic-tunnels tunnel-name]
```

Description

Enable RSVP to automatically establish LSPs for any new PE router added to a full mesh of LSPs. To enable this feature, you must configure the `rsvp-te` statement on all of the PE routers in the full mesh.

Options

`entry-name`—Specify the entry for the RSVP tunnel.

The other options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring RSVP Automatic Mesh

Configuring Dynamic Point-to-Multipoint Flooding LSPs

rtf-prefix-list

IN THIS SECTION

- Syntax | 16064
- Hierarchy Level | 16064
- Description | 16065
- Options | 16065
- Required Privilege Level | 16065
- Release Information | 16066

Syntax

```
rtf-prefix-list name route-targets
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options],
[edit logical-systems logical-system-name policy-options policy-statement policy-name term term-
name],
```

```
[edit policy-options],
[edit policy-options policy-statement policy-name term term-name]
```

Description

Define a list of route target prefixes for use in a routing policy statement. These prefixes are only useful for filtering routes in the bgp.rtarget.0 table.

The route target filtering prefix is in the format: *AS number:route target extended community/length*. The first number represents the autonomous system (AS) of the device that sent the advertisement. The second group of numbers represent the route target extended community. The format of the extended community is the same as the extended community type `target:`. For more information about extended communities, see *Understanding How to Define BGP Communities and Extended Communities* .

In this route target prefix example 64500:200:101/96, 64500 is the AS number, 200:101 is the BGP extended community used for the route target, and 96 is the prefix length.

For more information about the route target community, see RFC 4360, *BGP Extended Communities Attribute*.

For more information about the route target filtering prefix format, see RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*.

Options

name—Name that identifies the list of route target filtering prefixes. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in quotation marks (“ “).

route-targets—List of route target filtering prefixes, one route target filter per line in the configuration. When you use the `rtf-prefix-list` statement as a match condition, you do not have the option of configuring the list of route target filtering prefixes. You must first define and configure the route target filtering prefixes with the `policy-options` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- *Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs*
- *Configuring BGP Route Target Filtering for VPNs*
- *Understanding Proxy BGP Route Target Filtering for VPNs*
- *Understanding How to Define BGP Communities and Extended Communities*
- Routing Policies, Firewall Filters, and Traffic Policers User Guide
- family route-target | **6911**

rtlog (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | **16066**
- Hierarchy Level | **16067**
- Description | **16067**
- Required Privilege Level | **16067**
- Release Information | **16067**

Syntax

```
rtlog {
    name {
        apply-groups group-names;
        apply-groups-except group-names;
        flag name;
        file filename,
        no-remote-trace;
```

```
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Enable global system logging for Next Gen Services.

traceoptions Specify the options to include in the trace.

All other options are explained separtely.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

- *Understanding Next Gen Services CGNAT Global System Logging*
- *Enabling Global System Logging for Next Gen Services*
- *Configuring System Logging to One or More Remote Servers for Next Gen Services*
- *Configuring Local System Logging for Next Gen Services*
- traceoptions (Next Gen Services Global System Logging) | **29619**

rtp

IN THIS SECTION

- Syntax | 16068
- Hierarchy Level | 16068
- Description | 16068
- Required Privilege Level | 16069
- Release Information | 16069

Syntax

```
rtp {
    f-max-period number;
    maximum-contexts number <force>;
    port {
        minimum port-number;
        maximum port-number;
    }
    queues [ queue-numbers ];
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number compression],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
compression]
```

Description

Configure the RTP properties for voice services traffic.

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Compression of Voice Traffic

rts

IN THIS SECTION

- Syntax | 16069
- Hierarchy Level | 16070
- Description | 16070
- Options | 16070
- Required Privilege Level | 16070
- Release Information | 16070

Syntax

```
rts (assert | de-assert | normal);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options dce-options],
[edit interfaces interface-name serial-options dte-options]
```

Description

For EIA-530 and V.35 interfaces only, configure the to-DCE signal, request to send (RTS).

Options

`assert`—The to-DCE signal must be asserted.

`de-assert`—The to-DCE signal must be deasserted.

`normal`—Normal RTS signal handling, as defined by the TIA/EIA Standard 530.

- **Default:** `normal`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Signal Handling

rts-polarity

IN THIS SECTION

- Syntax | 16071
- Hierarchy Level | 16071
- Description | 16071
- Options | 16071
- Required Privilege Level | 16071
- Release Information | 16072

Syntax

```
rts-polarity (negative | positive);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

Configure RTS signal polarity.

Options

`negative`—Negative signal polarity.

`positive`—Positive signal polarity.

- **Default:** `positive`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Serial Signal Polarities

rtsp

IN THIS SECTION

- Syntax | 16072
- Hierarchy Level | 16073
- Description | 16073
- Options | 16073
- Required Privilege Level | 16073
- Release Information | 16073

Syntax

```
rtsp {
    disable;
    traceoptions {
        flag {
            all {
                extensive
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Configure Real-Time Streaming Protocol (RTSP) ALG on the device to control he delivery of data with real-time properties such as audio and video.

Options

- `disable`—Disable the RTSP ALG. By default, the RTSP ALG is enabled.
- `traceoptions`—Configure RTSP ALG tracing options.
  - `flag`—Trace operation to perform.
    - `all`—Trace all events.
    - *`extensive`*—Display extensive amount of data.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding the RTSP ALG

rule

IN THIS SECTION

- Syntax | 16074
- Hierarchy Level | 16075
- Description | 16075
- Options | 16075
- Required Privilege Level | 16075
- Release Information | 16075

Syntax

```
rule rule-name {
    match-direction (input | output);
    term term-name {
        from {
            destination-address address;
            ipsec-inside-interface interface-name;
            source-address address;
        }
        then {
            anti-replay-window-size bits;
            backup-remote-gateway address;
            clear-dont-fragment-bit;
            dynamic {
                ike-policy policy-name;
                ipsec-policy policy-name;
            }
            initiate-dead-peer-detection;
            manual {
                direction (inbound | outbound | bidirectional) {
                    authentication {
                        algorithm (hmac-md5-96 | hmac-sha1-96);
                        key (ascii-text key | hexadecimal key);
                    }
                    auxiliary-spi spi-value;
```

```
                    encryption {
                        algorithm algorithm;
                        key (ascii-text key | hexadecimal key);
                    }
                    protocol (ah | bundle | esp);
                    spi spi-value;
                }
            }
            no-anti-replay;
            remote-gateway address;
            syslog;
            tunnel-mtu bytes;
        }
    }
}
```

Hierarchy Level

```
[edit services ipsec-vpn],
[edit services ipsec-vpn rule-set rule-set-name]
```

Description

Specify the rule the router uses when applying this service.

Options

rule-name—Identifier for the collection of terms that comprise this rule.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

Configuring IPsec Rule Sets

Configuring Security Associations

rule (AACL Rule Set)

IN THIS SECTION

- Syntax | **16076**
- Hierarchy Level | **16077**
- Description | **16077**
- Options | **16077**
- Required Privilege Level | **16077**
- Release Information | **16077**

Syntax

```
rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
        from {
            application-group-any;
            application-groups [ application-group-names ];
            application-unknown;
            applications [ application-names ];
            destination-address address <any-unicast>;
            destination-address-range low minimum-value high maximum-value;
            destination-prefix-list list-name;
            nested-application-unknown;
            source-address address <any-unicast>;
            source-address-range low minimum-value high maximum-value;
            source-prefix-list list-name;
        }
```

```
        then {
            (accept | discard);
            count (application | application-group | application-group-any | nested-application
| none);
            forwarding-class class-name;
            policer policer-name;
        }
    }
}
```

Hierarchy Level

```
[edit services aacl],
[edit services aacl rule-set rule-set-name]
```

Description

Specify the rule the router uses when applying this service.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring AACL Rules

rule-activation-time

IN THIS SECTION

- Syntax | 16078
- Hierarchy Level | 16078
- Description | 16078
- Options | 16079
- Required Privilege Level | 16079
- Release Information | 16079

Syntax

```
rule-activation-time {
    <day-of-week | day-of-month month>;
    <hour:min>;
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-time-of-day-profiles profile-name]
```

Description

Specify the time of day, day of the week or day of the month, or month of the year to activate a PCC rule or rulebase. You can specify the time of day, the day, or both. If you specify the day of the month, you can also specify the month of the year, which results in the yearly activation of the rule or rulebase. Use the same combination of options in the `rule-deactivation-time` statement. If a time zone is configured on the router, the time-of-day settings apply to the configured time zone.

If a day is not specified and the activation time of day setting is later than the deactivation time of day setting, then a rule is deactivated the day after it is activated.

Options

day-of-week (Optional) Day of the week on which to activate a PCC rule or rulebase.

day-of-month (Optional) Day of the month on which to activate a PCC rule or rulebase.

- **Syntax: DAY*n***, where *n* can be from 1 through 31, or **Last-day-of-month**, which depends on the current month.

month (Optional) Month of the year in which to activate a PCC rule or rulebase.

hour (Optional) Hour at which to activate a PCC rule or rulebase as a two-digit number from 00 through 23.

min (Optional) Minute at which to activate a PCC rule or rulebase as a two-digit number from 00 through 59.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Static Time-of-Day PCC Rule Activation and Deactivation in a Junos OS Subscriber Aware PCEF Profile

Understanding Static Time-of-Day PCC Rule Activation and Deactivation

rule (Advanced Policy-Based Routing)

IN THIS SECTION

- Syntax | 16080
- Hierarchy Level | 16080
- Description | 16081
- Options | 16081
- Required Privilege Level | 16082
- Release Information | 16082

Syntax

```
rule rule-name {
    disable-midstream-routing;
    match {
        category (juniper-enhanced-category | custom-category);
        dynamic-application [system-application | any];
        dynamic-application-group [system-application-group];
        dscp dscp-value;
    }
    then {
        application-services-bypass;
        routing-instance name;
        sla-rule sla-rule-name;
    }
}
```

Hierarchy Level

[edit security advance-policy-based-routing profile *profile-name*]

Description

Configure rules for the advanced policy-based routing (APBR) profile (application profile). Associate the rule with one or more than one dynamic applications or application groups or URL categories as follows:

- For matching the dynamic applications, APBR consults the application identification (AppID) and application system cache (ASC) to get the application type. If the application matches any of the application or application groups of a rule in a profile, the application profile rule is considered to be a match, and the traffic is redirected to the defined routing instance for the route lookup.
- You can use a DSCP value in an APBR rule as a matching criteria to perform advanced policy-based routing on the traffic with DSCP markings. You can use the DSCP value in addition to the dynamic applications in an APBR rule.
- For matching the URL categories, APBR leverages category identification from the Enhanced Web Filtering (EWF) and local Web filtering results obtained from the Content Security module. Web filtering classifies websites into categories. If the traffic matches the URL categories specified in the rule of the APBR profile, it is redirected to the defined routing instance.

Options

disable-midstream-routing Selectively disable APBR in the middle of a session for a specific APBR rule.

match Define a match criteria for matching the traffic in APBR profile rule.

- **category (*juniper-enhanced-category* | *custom-category*)** Define the category type as the Juniper Enhanced Web Filtering (EWF) or a custom category if you are using local Web filtering.
  - ***juniper-enhanced-category*** Define URL categories such as Enhanced_Social_Web_Facebook, Enhanced_Social_Web_Linkedin, Enhanced_Social_Web_Twitter or Enhanced_Social_Web_Youtube as match criteria in APBR profile rule.
  - ***custom-category*** Define either custom URL or IP address of a site as match criteria in APBR profile rule.
- **dynamic-application [*system-application* | any]** Specify the dynamic application names for match criteria in APBR rule.

dynamic-application-group [*system-application-group*]	Dynamic application groups for match criteria in APBR rule.
dscp [*dscp-value*]	Specify DSCP value as match criteria in APBR rule. • **Range:** 0-63

then Define the action for the match condition by specifying the routing instance name.

application-services-bypass	Bypass applying the application services on the traffic matching the APBR rule. As URL category-based routing enables you to identify and selectively route Web traffic (HTTP and HTTPS) to a specified destination or to another device where further inspection, you can select not to apply or bypass application services on the same session. You can select to exclude traffic from security services when additional throughput is required, or traffic is going from trusted device to another trusted device.
routing-instance *name*	Name of the routing instance for redirecting traffic.
sla-rule *name*	Name of the SLA rule to apply on the traffic.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D60. The option `category` is introduced in Junos OS Release 18.3R1. Junos OS Release 19.3R1 supports the option `dscp`. Junos OS Release 20.1R1 supports the option `any` for `dynamic-application`.

RELATED DOCUMENTATION

Example: Configuring Advanced Policy-Based Routing for Application-Aware Traffic Management Solution

Understanding Advanced Policy-Based Routing

rule (Application Firewall)

IN THIS SECTION

- Syntax | 16083
- Hierarchy Level | 16083
- Description | 16084
- Options | 16084
- Required Privilege Level | 16085
- Release Information | 16085

Syntax

```
rule rule-name {
    match {
        dynamic-application [system-application];
        dynamic-application-groups [system-application-group];
        ssl-encryption (any | yes | no);
    }
    then {
        deny {
            block-message block-message;
        }
        permit permit;
        reject {
            block-message block-message;
        }
    }
}
```

Hierarchy Level

```
[edit security application-firewall  rule-sets name ]
```

Description

Specify rules for application firewall.

You need to create rules to permit, reject, or deny traffic for dynamic applications to configure application firewall rule sets within the security policy. The application firewall support in the policies provides additional security control for dynamic applications.

Starting in Junos OS Release 18.2R1 application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration.

Options

match Specify security rule match-criteria

dynamic-application Select dynamic applications as match criteria.

dynamic-application-group Select dynamic applications group as match criteria.

ssl-encryption Select SSL encryption rules as match criteria.

- Values:
  - `any`—Encrypted and non-encrypted rule.
  - `no`—Non-encrypted rule.
  - `yes`—Encrypted rule.

then Specify the action to be performed when traffic matches the associated match criteria.

deny Block the traffic at the firewall. The device drops the packet. By default, no message is returned to the sender.

block-message ***block-message*** (Optional) In application firewall rules, provide information to the user regarding blocked traffic. Depending on the content of the `profile` option for this rule set, including the `block-message` option displays a default message or customized message, or redirects the user for denied HTTP or HTTPS traffic. All other traffic is dropped silently.

reject Block the traffic at the firewall. For TCP traffic, by default the device drops the packet and returns a TCP reset (RST) message to the source host. For UDP and other protocol

traffic, by default the device drops the packet and returns an ICMP “destination unreachable, port unreachable” message to both the client and the server.

block-message *block-message* (Optional) In application firewall rules, provide information to the user regarding blocked traffic. Depending on the content of the `profile` option for this rule set, including the `block-message` option displays a default message or customized message, or redirects the user for denied HTTP or HTTPS traffic. All other traffic is dropped silently.

permit Permit traffic at the firewall.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 11.1. Statement updated in Junos OS Release 12.1X44-D10 to include the `ssl-encryption` and `reject` options. The `block-message` options added in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Application Firewall Overview

rule-sets (Security Application Firewall) | **16151**

application-firewall (Application Services) | **2320**

rule (Captive Portal Content Delivery)

IN THIS SECTION

- Syntax | **16086**
- Hierarchy Level | **16086**
- Description | **16086**

- Options | 16087
- Required Privilege Level | 16087
- Release Information | 16087

Syntax

```
rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
        from {
            destination-address address <except>;
        }
        then {
            accept;
            insert tag-name tag-name tag-value tag-value;
            redirect url;
            rewrite destination-address address <destination-port port-number>;
            syslog;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery],
[edit services captive-portal-content-delivery]
```

Description

Specify the rule the router uses when applying this service. Use the statement at the `[edit services...]` hierarchy level for static CPCD. Use the statement at the `[edit dynamic-profiles profile-name services...]` hierarchy level for converged services CPCD.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the `[edit dynamic-profiles profile-name services captive-portal-content-delivery rule rule-name term term-name]` hierarchy level added in Junos OS Release 17.2R1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

- *HTTP Redirect Service Overview*
- *Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services*
- *Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services*
- *Configuring Routing Engine-Based, Static HTTP Redirect Services*
- *Configuring Routing Engine-Based, Converged HTTP Redirect Services*

rule (casb-policy)

IN THIS SECTION

- Syntax | 16088
- Hierarchy Level | 16088
- Description | 16089

- Options | 16089
- Required Privilege Level | 16089
- Release Information | 16089

Syntax

```
rules rule-name {
    match {
          application application-name {
        activity activity-name {
        [param-name param-value param-name param-value]
        }
    instance instance-name;
       }
      application-group application-group-name {
             activity activity-name
         }
    }
    then {
    allow;
       deny;
       log-action;
    }
}
```

Hierarchy Level

```
[edit security casb casb-policy policy-name]
```

Description

Configure Cloud Access Security Broker (CASB) policy rule. CASB policy rule includes set of match criteria defining the conditions that must be satisfied to apply the policy and a set of actions to be performed in case of a match—allow, deny, or log actions.

Options

match *rule-name*

Define a match criteria for matching the traffic in CASB policy rule. You can specify following:

- application *application-name*—Enter the supported application name.
- application-group *application-group-name*—Enter the supported application-group name.
- activity—Configure activity-parameters for shared-activity option. You can configure this optional statement to have even more granular control over traffic.
- instance *instance-name*—Enter an application instance.

then

Define the action for the match condition by specifying—allow/deny and log-actions.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

rule-deactivation-time

IN THIS SECTION

- Syntax | 16090
- Hierarchy Level | 16090

- Description | 16090
- Options | 16090
- Required Privilege Level | 16091
- Release Information | 16091

Syntax

```
rule-deactivation-time {
    <day-of-week | day-of-month month>;
    <hour:min>;
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-time-of-day-profiles profile-name]
```

Description

Specify the time of day, day of the week or day of the month, or month of the year to deactivate a PCC rule or rulebase. You can specify the time of day, the day, or both. If you specify the day of the month, you can also specify the month of the year, which results in the yearly deactivation of the rule or rulebase. Use the same combination of options as in the `rule-activation-time` statement. If a time zone is configured on the router, the time-of-day settings apply to the configured time zone.

If a day is not specified and the deactivation time of day setting is earlier than the activation time of day setting, then a rule is deactivated the day after it is activated.

Options

day-of-week (Optional) Day of the week on which to deactivate a PCC rule or rulebase.

day-of-month (Optional) Day of the month on which to deactivate a PCC rule or rulebase.

- **Syntax: DAY*n***, where *n* can be from 1 through 31, or **Last-day-of-month**, which depends on the current month.

month (Optional) Month of the year in which to deactivate a PCC rule or rulebase.

hour (Optional) Hour at which to deactivate a PCC rule or rulebase as a two-digit number from 00 through 23.

min (Optional) Minute at which to deactivate a PCC rule or rulebase as a two-digit number from 00 through 59.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Static Time-of-Day PCC Rule Activation and Deactivation in a Junos OS Subscriber Aware PCEF Profile

Understanding Static Time-of-Day PCC Rule Activation and Deactivation

rule (Destination NAT Next Gen Services)

IN THIS SECTION

- Syntax | 16092
- Hierarchy Level | 16092
- Description | 16092
- Required Privilege Level | 16092
- Release Information | 16092

Syntax

```
rule rule-name {
    match {
        application [application-name]
        destination-address (NAT Next Gen Services) (address | any-unicast);
        destination-address-name address-name;
        source-address (address | any-unicast);
        source-address-name address-name;
    }
}
    then {
        destination-nat {
            destination-prefix destination-prefix;
            off;
            pool nat-pool-name;
        }
        port-forwarding-mappings map-name;
        }
    syslog;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set]
```

Description

Configure a destination NAT rule, which translates the destination address of IP packets.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

rule (HTTP Header Enrichment for Tag Rule Set)

IN THIS SECTION

- Syntax | 16093
- Hierarchy Level | 16093
- Description | 16093
- Options | 16093
- Required Privilege Level | 16094
- Release Information | 16094

Syntax

```
rule rule-name;
```

Hierarchy Level

```
[edit services hcm tag-rule-set]
```

Description

Specify the tag rule that you want to be a part of the tag rule set.

NOTE: The tag rule must already be defined at the `[edit services hcm]` hierarchy level.

Options

rule-name	Name of the tag rule.

To specify multiple tag rules, include the `rule` statement multiple times.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm (HTTP Header Enrichment) | 8005

rule (IDS MS-DPC)

IN THIS SECTION

- Syntax | 16094
- Hierarchy Level | 16096
- Description | 16096
- Options | 16096
- Required Privilege Level | 16096
- Release Information | 16096

Syntax

```
rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
        from {
            application-sets set-name;
            applications [ application-names ];
```

```
            destination-address (address | any-unicast) <except>;
            destination-address-range low minimum-value high maximum-value <except>;
            source-address (address | any-unicast) <except>;
            source-address-range low minimum-value high maximum-value <except>;
        }
        then {
            aggregation (IDS) {
                destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
                source-prefix prefix-value | source-prefix-ipv6 prefix-value;
            }
            (force-entry | ignore-entry);
            logging {
                syslog;
                threshold rate;
            }
            session-limit {
                by-destination (IDS MS-DPC) {
                    hold-time seconds;
                    maximum number;
                    packets number;
                    rate number;
                }
                by-pair (IDS MS-DPC) {
                    hold-time seconds;
                    maximum number;
                    packets number;
                    rate number;
                }
                by-source (IDS MS-DPC) {
                    hold-time seconds;
                    maximum number;
                    packets number;
                    rate number;
                }
            }
            syn-cookie {
                mss value;
                threshold rate;
            }
        }
    }
}
```

Hierarchy Level

```
[edit services ids],
[edit services ids rule-set rule-set-name]
```

Description

Specify the rule the router uses when applying this service on the MS-DPC.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IDS Rules on an MS-DPC

rule (IDS MS-MPC)

IN THIS SECTION

- Syntax | 16097
- Hierarchy Level | 16099
- Description | 16099

- Required Privilege Level | **16099**
- Release Information | **16099**

Syntax

```
rule {
    match-direction (input | output | input-output);
    term {
        then {
            aggregation (IDS) {
                destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
                source-prefix prefix-value | source-prefix-ipv6 prefix-value;
            }
            allow-ip-options {
                any;
                loose-source-route;
                route-record;
                route-alert;
                security;
                stream-id;
                strict-source-route;
                timestamp;
            }
            allow-ipv6-extension-header {
                any;
                ah;
                dstopts;
                esp;
                fragment;
                hop-by-hop;
                mobility;
                routing;
            }
            icmp-fragment-check;
            icmp-large-packet-check;
            land-attack-check (ip-only | ip-port);
            session-limit {
                by-destination {
```

```
                by-protocol {
                    icmp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                    tcp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                    udp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                }
                maximum number;
                packets number;
                rate number;
            }
            by-source {
                by-protocol {
                    icmp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                    tcp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                    udp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                }
                maximum number;
                packets number;
                rate number;
            }
```

```
            }
            tcp-syn-defense;
            tcp-syn-fragment-check;
            tcp-winnuke-check;
        }
    }
}
```

Hierarchy Level

```
[edit services ids ]
```

Description

Configure network attack protection for a service set on an MS-MPC.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

rule (IP Reassembly)

IN THIS SECTION

- Syntax | 16100
- Hierarchy Level | 16100
- Description | 16100
- Options | 16101
- Required Privilege Level | 16101
- Release Information | 16101

Syntax

```
rule rule-name {
    match-direction direction;
}
```

Hierarchy Level

```
[edit services ip-reassembly]
```

Description

Configure an IP reassembly rule, which is used for inline IP reassembly on the inline services interface. The IP reassembly rule can be attached to a service set to indicate that the service set is of type IP reassembly. For inline IP reassembly, each rule must include the `match-direction` statement, which specifies the direction in which the match is applied.

NOTE: If you configure an IP reassembly rule, then you must configure the `match-direction` statement.

> **NOTE**: To create more than one IP reassembly rule, include the `rule` statement multiple times.

Options

rule-name Name of the IP reassembly rule.

- **Range:** Up to 63 characters

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring IP Inline Reassembly for L2TP

IP Packet Fragment Reassembly for L2TP Overview

rule (LRF)

IN THIS SECTION

Syntax | 16102

- Hierarchy Level | 16102
- Description | 16102
- Options | 16102
- Required Privilege Level | 16103
- Release Information | 16103

Syntax

```
rule lrf-rule-name {
    then {
        report {
            collector collector-name;
            template template-name;
            time-limit time-interval;
            volume-limit volume;
        }
    }
}
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Description

Configure an LRF rule, which controls how data sessions are logged and reported. In this release, the matching conditions for an LRF rule are identified in a static PCC rule, not in the LRF rule.

Options

lrf-rule-name

Name of the LRF rule.

- **Range:** Up to 63 characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

rule (PCP)

IN THIS SECTION

- Syntax | 16104
- Hierarchy Level | 16104
- Description | 16104
- Options | 16105
- Required Privilege Level | 16105
- Release Information | 16105

Syntax

```
rule rule-name {
    match-direction (input | output);
    term term-name {
        from {
            application-sets set-name;
            applications [ application-name ];
            destination-address address <except>;
            destination-address-range high maximum-value low minimum-value <except>;
            destination-port  high maximum-value low minimum-value;
            destination-prefix-list list-name <except>;
            source-address address <except>;
            source-address-range high maximum-value low minimum-value <except>;
            source-prefix-list list-name <except>;
        }
        then {
            pcp-server server-name;
        }
    }
    }
}
```

Hierarchy Level

```
[edit services pcp]
```

Description

Configure a rule to assign the port control protocol (PCP) server that handles selected traffic. PCP enables hosts to operate servers for a long time (as in the case of a webcam) or a short time (for example, while playing a game or on a phone call) when behind a NAT device, including when behind a carrier-grade NAT operated by their ISP. PCP enables applications to create mappings from an external IP address and port to an internal IP address and port.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

rule-name	Rule name

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

rule (Security Destination NAT)

IN THIS SECTION

- Syntax | 16106
- Hierarchy Level | 16106
- Description | 16106
- Options | 16106
- Required Privilege Level | 16107
- Release Information | 16107

Syntax

```
rule rule-name {
    description text;
    match {
        application {
            [application];
            any;
        }
        (destination-address ip-address| destination-address-name address-name);
        destination-port (port-or-low <to high>);
        protocol [protocol-name-or-number];
        source-address [ip-address];
        source-address-name [address-name];
    }
    then {
        destination-nat (off | pool pool-name | rule-session-count-alarm (clear threshold value
| raise-threshold value));
    }
}
```

Hierarchy Level

```
[edit security nat destination rule-set          rule-set-name]
```

Description

Define a destination NAT rule.

Options

- *rule-name*—Name of the destination NAT rule.
- `description`—Description of the destination NAT rule.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. The `description` option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

rule (Security Exempt Rulebase)

IN THIS SECTION

- Syntax | 16107
- Hierarchy Level | 16108
- Description | 16108
- Options | 16108
- Required Privilege Level | 16108
- Release Information | 16108

Syntax

```
rule rule-name {
    description text;
    match (Security IDP Policy) {
        attacks (Security IPS Rulebase) {
            custom-attack-groups [attack-group-name];
            custom-attacks [attack-name];
            dynamic-attack-groups [attack-group-name];
            predefined-attack-groups [attack-group-name];
            predefined-attacks [attack-name];
        }
```

```
        destination-address (Security IDP Policy) ([address-name] | any | any-ipv4 | any-ipv6);
        destination-except [address-name];
        from-zone (zone-name | any );
        source-address ([address-name] | any | any-ipv4 | any-ipv6);
        source-except [address-name];
        to-zone (zone-name | any);
    }
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt]
```

Description

Specify exempt rule to create, modify, delete, and reorder the rules in a rulebase.

Options

rule-name—Name of the exempt rulebase rule.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

rule (Security IPS Rulebase)

IN THIS SECTION

- Syntax | 16109
- Hierarchy Level | 16110
- Description | 16110
- Options | 16110
- Required Privilege Level | 16111
- Release Information | 16111

Syntax

```
rule rule-name {
    description text;
    match (Security IDP Policy) {
        application (application-name | any | default);
        attacks (Security IPS Rulebase) {
            custom-attack-groups [attack-group-name];
            custom-attacks [attack-name];
            dynamic-attack-groups [attack-group-name];
            predefined-attack-groups [attack-group-name];
            predefined-attacks [attack-name];
        }
        destination-address (Security IDP Policy) ([address-name] | any | any-ipv4 | any-ipv6);
        destination-except [address-name];
        from-zone (zone-name | any );
        source-address ([address-name] | any | any-ipv4 | any-ipv6);
        source-except [address-name];
        to-zone (zone-name | any);
    }
    terminal;
    then {
        action {
            class-of-service {
                dscp-code-point number;
                 forwarding-class forwarding-class;
```

```
            }
            (close-client | close-client-and-server | close-server |drop-connection | drop-
packet | ignore-connection | mark-diffserv value | no-action | recommended);
        }
        ip-action (Security IDP Rulebase IPS) {
            (ip-block | ip-close | ip-notify);
            log;
            log-create;
            refresh-timeout;
            target (destination-address | service | source-address | source-zone | source-zone-
address | zone-service);
            timeout seconds;
        }
        notification {
            log-attacks {
                alert;
            }
            packet-log (Security IDP Policy) {
                post-attack number;
                post-attack-timeout seconds;
                pre-attack number;
            }
        }
        severity (critical | info | major | minor | warning);
    }
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips]
```

Description

Specify IPS rule to create, modify, delete, and reorder the rules in a rulebase.

Options

rule-name—Name of the IPS rulebase rule.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Starting in Junos OS Release 18.2R1, IDP policy is available within unified security policy. When IDP policy is available within the unified security policy then the IDP polciy configurations are simplified. Configuring source or destination address, source and destination-except, from and to zone, or application is not required with unified policy, as the match happens in the security policy itself.

Additional tags under filters of dynamic attack groups are added in Junos OS Release 18.2R1 for dynamic attacks grouping of IDP signatures.

rule (Security Source NAT)

IN THIS SECTION

- Syntax | 16111
- Hierarchy Level | 16112
- Description | 16112
- Options | 16113
- Required Privilege Level | 16113
- Release Information | 16113

Syntax

```
rule rule-name {
    description text;
    match {
        application {
            [application];
```

```
            any;
        }
        (destination-address <ip-address> | destination-address-name <address-name>);
        destination-port (port-or-low <to high>);
        protocol [protocol-name-or-number];
        source-address [ip-address];
        source-address-name [address-name];
        source-port (port-or-low <to high>)
    }
    then {
        source-nat {
            interface {
                persistent-nat {
                    address-mapping;
                    inactivity-timeout seconds;
                    max-session-number value;
                    permit (any-remote-host | target-host | target-host-port);
                }
            }
            off;
            pool <pool-name>
                 persistent-nat {
                     address-mapping;
                    inactivity-timeout seconds;
                    max-session-number number;
                    permit (any-remote-host | target-host | target-host-port);
                 }
            }
        }
    }
}
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name]
```

Description

Define a source NAT rule.

Options

- *rule-name*—Name of the source NAT rule.
- description—Description of the source NAT rule.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. The description option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

rule (Security Static NAT)

IN THIS SECTION

- Syntax | 16113
- Hierarchy Level | 16114
- Description | 16115
- Options | 16115
- Required Privilege Level | 16115
- Release Information | 16115

Syntax

```
rule rule-name {
    description text;
```

```
    match {
        (destination-address <ip-address> | destination-address-name <address-name>);
        destination-port (port-or-low | <to high>);
        source-address [ip-address];
        source-address-name [ip-address-name];
        source-port (port-or-low <to high>);
    }
    then {
        static-nat {
            inet {
                routing-instance (routing-instance-name| default);
            }
            nptv6-prefix {
                address-prefix;
                    routing-instance routing-instance-name;
            nptv6-prefix-name {
                address-prefix-name;
                    routing-instance routing-instance-name;
            [
            prefix {
                address-prefix;
                mapped-port lower-port-range to upper-port-range;
                routing-instance (routing-instance-name| default);
            }
            prefix-name {
                address-prefix-name;
                mapped-port lower-port-range to upper-port-range;
                routing-instance (routing-instance-name| default);
                }
            rule-session-count-alarm (clear-threshold value | raise-threshold value);
        }
    }
}
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name]
```

Description

Define a static NAT rule.

Options

- *rule-name*—Name of the static NAT rule.
- `Description`—Description of the static NAT rule.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3. The `description` option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.3X48-D25.

rule (Services CoS)

IN THIS SECTION

- Syntax | 16116
- Hierarchy Level | 16116
- Description | 16116
- Options | 16117
- Required Privilege Level | 16117
- Release Information | 16117

Syntax

```
rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
        from {
            application-sets set-name;
            applications [ application-names ];
            destination-address address;
            destination-prefix-list list-name <except>;
            source-address address;
            source-prefix-list list-name <except>;
        }
        then {
            application-profile profile-name;
            dscp (alias | bits);
            forwarding-class class-name;
            syslog;
            reflexive; | revert; | reverse {
                application-profile profile-name;
                dscp (alias | bits);
                forwarding-class class-name;
                syslog;
            }
        }
    }
}
```

Hierarchy Level

```
[edit services cos],
[edit services cos rule-set rule-set-name]
```

Description

Specify the rule the router uses when applying this service.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring CoS Rules

Configuring CoS Rules on Services PICs

rule (Services CoS Next Gen Services)

IN THIS SECTION

- Syntax | **16118**
- Hierarchy Level | **16118**
- Description | **16118**
- Options | **16119**
- Required Privilege Level | **16119**
- Release Information | **16119**

Syntax

```
rule rule-name {
    match-direction (input | input-output | output);
    policy policy-name {
        match {
            application [ application-names ];
            destination-address address;
            destination-address-range low minimum-value high maximum-value;
            destination-port port-number;
            destination-prefix-list list-name;
            source-address address;
            source-address-range low minimum-value high maximum-value;
            source-prefix-list list-name;
        }
        then {
            application-profile profile-name;
            dscp (alias | bits);
            forwarding-class class-name;
            reflexive; | revert; | reverse {
                application-profile profile-name;
                dscp (alias | bits);
                forwarding-class class-name;
            }
        }
    }
}
```

Hierarchy Level

```
[edit services cos]
```

Description

Configure a services CoS rule, which specifies Differentiated Services (DiffServ) code point (DSCP) marking and forwarding-class assignment for packets that are processed by a service set. The CoS rule identifies the matching conditions for packet source and destination addresses and for packet applications, and the actions to take on those packets.

The service set that the CoS rule is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

match-direction (input | input-output | output) The direction in which the rule is matched.

- **input** Apply the rule match on input. If the CoS rule is assigned to an interface service set, input means traffic entering the interface. If the CoS rule is assigned to a next-hop service set, input means traffic routed with the inside interface.
- **input-output** Apply the rule match in both directions.
- **output** Apply the rule match on output. If the CoS rule is assigned to an interface service set, input means traffic leaving the interface. If the CoS rule is assigned to a next-hop service set, output means traffic routed with the outside interface.

rule-name Name of the CoS rule.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Class of Service for Services PICs (Next Gen Services)

rule (Services NAT)

IN THIS SECTION

- Syntax | 16120
- Hierarchy Level | 16121
- Description | 16121
- Options | 16121
- Required Privilege Level | 16122
- Release Information | 16122

Syntax

```
rule rule-name {
    match-direction (input | output);
    term term-name {
        from {
            application-sets set-name;
            applications [ application-names ];
            destination-address (address | any-unicast) <except>;
            destination-address-range low minimum-value high maximum-value <except>;
            destination-port range high maximum-value low minimum-value;
            source-address (address | any-unicast) <except>;
            source-address-range low minimum-value high maximum-value <except>;
        }
        then {
            no-translation;
            port-forwarding-mappings map-name;
            translated {
                address-pooling paired;
                clat-prefix clat-prefix;
                destination-pool nat-pool-name;
                  destination-prefix destination-prefix; destination-prefix;
                dns-alg-pool dns-alg-pool;
                dns-alg-prefix dns-alg-prefix;
                filtering-type endpoint-independent;
                mapping-type endpoint-independent;
```

```
                overload-pool overload-pool;
                overload-prefix overload-prefix;
                source-pool nat-pool-name;
                source-prefix source-prefix;
                translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | deterministic-napt44
| deterministic-napt64 |dnat-44 | dynamic-nat44 | napt-44 | napt-66 | napt-pt | stateful-nat464
| stateful-nat64 | twice-basic-nat-44 | twice-dynamic-nat-44 | twice-napt-44);
                }
            }
            syslog;
        }
    }
}
```

Hierarchy Level

```
[edit services nat],
[edit services nat rule-set rule-set-name]
```

Description

Specify the rule the router uses when applying this service.

NOTE: You are limited to a maximum of 200 terms for a NAT rule that is applied to an inline services (type si) interface. If you specify more than 200 terms, you will receive following error when you commit the configuration:

```
[edit]
  'service-set service-set-name'
    NAT rule rule-name  with more than 200 terms is disallowed for si-n/n/n.n
error: configuration check-out failed
```

Options

rule-name Identifier for the collection of terms that make up this rule.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

rule (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 16122
- Hierarchy Level | 16123
- Description | 16123
- Options | 16123
- Required Privilege Level | 16123
- Release Information | 16123

Syntax

```
rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
        from {
            application-sets set-name;
            applications [ application-names ];
            destination-address (address | any-ipv4 | any-ipv6 | any-unicast) <except>;
```

```
            destination-address-range low minimum-value high maximum-value <except>;
            destination-prefix-list list-name <except>;
            source-address (address | any-ipv4 | any-ipv6 | any-unicast) <except>;
            source-address-range low minimum-value high maximum-value <except>;
            source-prefix-list list-name <except>;
        }
        then {
            (accept <skip-ids>| discard | reject);
            syslog;
        }
    }
}
```

Hierarchy Level

```
[edit services stateful-firewall],
[edit services stateful-firewall rule-set rule-set-name]
```

Description

Specify the rule the router uses when applying this service.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Stateful Firewall Rules

rule-session-count-alarm (Security Destination NAT Rule Set)

IN THIS SECTION

- Syntax | 16124
- Hierarchy Level | 16124
- Description | 16124
- Options | 16124
- Required Privilege Level | 16125
- Release Information | 16125

Syntax

```
rule-session-count-alarm (clear-threshold value | raise-threshold value):
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name rule rule-name then destination-nat]
```

Description

Define session count alarm thresholds for a specific Network Address Translation (NAT) destination rule. When the session count exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered.

Options

clear-threshold *value* Lower threshold at which an SNMP trap is triggered.

raise-threshold *value* Upper threshold at which an SNMP trap is triggered.

> **NOTE**: If you enter a value for **raise-threshold** but not for **clear-threshold**, **clear-threshold** is automatically set to 80 percent of **raise-threshold**.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

rule-session-count-alarm (Security Source NAT Rule Set)

IN THIS SECTION

- Syntax | 16125
- Hierarchy Level | 16126
- Description | 16126
- Options | 16126
- Required Privilege Level | 16126
- Release Information | 16126

Syntax

```
rule-session-count-alarm (clear-threshold value | raise-threshold value):
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name then source-nat ]
```

Description

Define session count alarm thresholds for a specific Network Address Translation (NAT) source rule. When the session count exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered.

Options

clear-threshold *value* Lower threshold at which an SNMP trap is triggered.

raise-threshold *value* Upper threshold at which an SNMP trap is triggered.

> **NOTE**: If you enter a value for `raise-threshold` but not for `clear-threshold`, `clear-threshold` is automatically set to 80 percent of `raise-threshold`.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

rule-session-count-alarm (Security Static NAT Rule Set)

IN THIS SECTION

Syntax | 16127

- Hierarchy Level | 16127
- Description | 16127
- Options | 16127
- Required Privilege Level | 16127
- Release Information | 16128

Syntax

```
rule-session-count-alarm (clear-threshold value | raise-threshold value);
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name then static-nat ]
```

Description

Define session count alarm thresholds for a specific static Network Address Translation (NAT) rule. When the session count exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered.

Options

clear-threshold *value* Lower threshold at which an SNMP trap is triggered.

raise-threshold *value* Upper threshold at which an SNMP trap is triggered.

NOTE: If you enter a value for `raise-threshold` but not for `clear-threshold`, `clear-threshold` is automatically set to 80 percent of `raise-threshold`.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

rule-set

IN THIS SECTION

- Syntax | 16128
- Hierarchy Level | 16128
- Description | 16128
- Options | 16129
- Required Privilege Level | 16129
- Release Information | 16129

Syntax

```
rule-set rule-set-name {
    [ rule rule-names ];
}
```

Hierarchy Level

```
[edit services ipsec-vpn]
```

Description

Specify the rule set the router uses when applying this service.

Options

rule-set-name—Identifier for the collection of rules that constitute this rule set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rule Sets

Configuring IPsec Rules

rule-set (Captive Portal Content Delivery)

IN THIS SECTION

- Syntax | 16130
- Hierarchy Level | 16130
- Description | 16130
- Options | 16130
- Required Privilege Level | 16130
- Release Information | 16130

Syntax

```
rule-set rule-set-name {
    [rule rule-name];
}
```

Hierarchy Level

```
[edit services captive-portal-content-delivery]
```

Description

Define a set of captive portal content delivery rules that the router uses when applying this service. This statement is supported only for static CPCD; it is not supported for converged services CPCD.

Options

rule-set-name—Identifier for the collection of rules that constitute this rule set.

rule *rule-name*–Name of a rule defined at the [edit services captive-portal-content-delivery] hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

HTTP Redirect Service Overview

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

Configuring Routing Engine-Based, Static HTTP Redirect Services

Configuring Routing Engine-Based, Converged HTTP Redirect Services

rule-set (Security Destination NAT)

IN THIS SECTION

- Syntax | 16131
- Hierarchy Level | 16132
- Description | 16132
- Options | 16132
- Required Privilege Level | 16132
- Release Information | 16132

Syntax

```
rule-set rule-set-name {
    description text;
    from {
        interface [interface-name];
        routing-group [routing-group-name]];
        routing-instance [routing-instance-name];
        zone [zone-name];
    }
    rule rule-name {
        description text;
        match {
            application {
                [application];
                any;
            }
            destination-address ip-address| destination-address-name address-name);
            destination-port (port-or-low <to high>);
            protocol [protocol-name-or-number];
            source-address [ip-address];
```

```
            source-address-name [address-name];
        }
        then {
            destination-nat (off | pool pool-name | rule-session-count-alarm (clear-threshold
value | raise-threshold value));
            }
        {
    }
}
```

Hierarchy Level

```
[edit security nat destination]
```

Description

Configure a set of rules for destination NAT.

Options

rule-set-name	Name of the rule set.
description	Description of the rule set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. The `description` option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

The `routing-group` option is added in Junos OS Release 22.2R1.

rule-set (Security Source NAT)

IN THIS SECTION

- Syntax | 16133
- Hierarchy Level | 16134
- Description | 16134
- Options | 16135
- Required Privilege Level | 16135
- Release Information | 16135

Syntax

```
rule-set rule-set-name {
    description text;
    from {
        interface [interface-name];
        routing-group [routing-group-name]];
        routing-instance [routing-instance-name];
        zone [zone-name];
    }
    rule rule-name {
        description text;
        match {
            application {
                [application];
                any;
            }
            (destination-address <ip-address> | destination-address-name <address-name>);
            destination-port (port-or-low <to high>);
            protocol [protocol-name-or-number];
            source-address [ip-address];
            source-address-name [address-name];
            source-port (port-or-low <to high>);
        }
        then {
            source-nat {
```

```
                interface {
                    persistent-nat {
                        address-mapping;
                        inactivity-timeout seconds;
                        max-session-number value;
                        permit (any-remote-host | target-host | target-host-port);
                    }
                }
                off;
                pool <pool-name>
                     persistent-nat {
                         address-mapping;
                        inactivity-timeout seconds;
                        max-session-number number;
                        permit (any-remote-host | target-host | target-host-port);
                     }
                }
                rule-session-count-alarm (raise-threshold value | clear-threshold value);
            }
        }
    }
    to {
        interface [interface-name];
        routing-group [routing-group-name]];
        routing-instance [routing-instance-name];
        zone [zone-name];
    }
}
```

Hierarchy Level

```
[edit security nat source]
```

Description

Configure a set of rules for source NAT.

> **NOTE**: When zones are not configured under rule-set and when active source NAT is configured with missing mandatory statement “from” then, the following message is displayed when performing commit **“Missing mandatory statement: 'from' error: configuration check-out failed”** and the configuration check-out fails.

Starting from 19.3R3, the external node connection for one persistent NAT binding is updated from 8 to 32 .

Options

rule-set-name	Name of the rule set.
description	Description of the rule set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. The `description` option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

The `routing-group` option is added in Junos OS Release 22.2R1.

rule-set (Security Static NAT)

IN THIS SECTION

- Syntax | **16136**
- Hierarchy Level | **16137**

- Description | 16137
- Options | 16137
- Required Privilege Level | 16137
- Release Information | 16137

Syntax

```
rule-set rule-set-name {
    description text;
    from {
        interface [interface-name];
        routing-group [routing-group-name]];
        routing-instance [routing-instance-name];
        zone [zone-name];
    }
    rule rule-name {
        description text;
        match {
            (destination-address ip-address | destination-address-name address-name);
            destination-port (port | low to high);
            source-address ip-address;
            source-address-name address-name;
            source-port (port or low <to high>);
        }
        then {
            static-nat {
                inet {
                    routing-instance (default | routing-instance-name);
                }
                prefix {
                    address-prefix;
                    mapped-port lower-port-range to upper-port-range;
                    routing-instance (default | routing-instance-name);
                }
                prefix-name {
                    address-prefix-name;
                    mapped-port lower-port-range to upper-port-range;
                    routing-instance (default | routing-instance-name);
```

```
                    }
                rule-session-count-alarm (raise-threshold value | clear-threshold value);
            }
        }
    }
}
```

Hierarchy Level

```
[edit security nat static]
```

Description

Configure a set of rules for static NAT.

Options

rule-set-name	Name of the rule set.
description	Description of the rule set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1. The `rule-session-count-alarm`, `source-address`, `source-address-name`, and `source-port` options added in Junos OS Release 12.1X45-D10.

The `routing-group` option is added in Junos OS Release 22.2R1.

rule-set (Services AACL)

IN THIS SECTION

- Syntax | 16138
- Hierarchy Level | 16138
- Description | 16138
- Options | 16138
- Required Privilege Level | 16138
- Release Information | 16139

Syntax

```
rule-set rule-set-name {
    [rule rule-names ];
}
```

Hierarchy Level

```
[edit services aacl]
```

Description

Specify the rule set the router uses when applying this service.

Options

rule-set-name—Identifier for the collection of rules that constitute this rule set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring AACL Rule Sets

rule-set (Services CoS)

IN THIS SECTION

- Syntax | 16139
- Hierarchy Level | 16139
- Description | 16140
- Options | 16140
- Required Privilege Level | 16140
- Release Information | 16140

Syntax

```
rule-set rule-set-name {
    [ rule rule-name ];
}
```

Hierarchy Level

```
[edit services cos]
```

Description

Specify the rule set the router uses when applying this service.

Options

rule-set-name—Identifier for the collection of rules that constitute this rule set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring CoS Rule Sets on Services PICs

rule-set (Services CoS Next Gen Services)

IN THIS SECTION

- Syntax | **16141**
- Hierarchy Level | **16141**
- Description | **16141**
- Options | **16141**
- Required Privilege Level | **16141**
- Release Information | **16141**

Syntax

```
rule-set rule-set-name {
    [ rule rule-name ];
}
```

Hierarchy Level

```
[edit services cos]
```

Description

Configure a set of services CoS rules. You can then assign the rule set to a service set, which processes the rules in the order they appear. Once a rule matches the packet, the router performs the corresponding action, and no further rules are applied.

Options

rule ***rule-name***	The name of each rule in the rule set.
rule-set-name	The name for the set of rules.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Class of Service for Services PICs (Next Gen Services)

rule-set (Services IDS)

IN THIS SECTION

- Syntax | 16142
- Hierarchy Level | 16142
- Description | 16142
- Options | 16142
- Required Privilege Level | 16142
- Release Information | 16143

Syntax

```
rule-set rule-set-name {
    [ rule rule-names ];
}
```

Hierarchy Level

```
[edit services ids]
```

Description

Specify the rule set the router uses when applying this service.

Options

rule-set-name—Identifier for the collection of rules that constitute this rule set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IDS Rule Sets on an MS-DPC

rule-set (Services NAT)

IN THIS SECTION

- Syntax | 16143
- Hierarchy Level | 16143
- Description | 16144
- Options | 16144
- Required Privilege Level | 16144
- Release Information | 16144

Syntax

```
rule-set rule-set-name {
    [ rule rule-names ];
}
```

Hierarchy Level

```
[edit services nat]
```

Description

Specify the rule set the router uses when applying this service.

Options

rule-set-name—Identifier for the collection of rules that constitute this rule set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

rule-set (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 16145
- Hierarchy Level | 16145
- Description | 16145
- Options | 16145
- Required Privilege Level | 16145
- Release Information | 16145

Syntax

```
rule-set rule-set-name {
    [ rule rule-names ];
}
```

Hierarchy Level

```
[edit services stateful-firewall]
```

Description

Specify the rule set the router uses when applying this service.

Options

rule-set-name—Identifier for the collection of rules that constitute this rule set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Stateful Firewall Rule Sets

rule-set (Softwire)

IN THIS SECTION

- Syntax | 16146
- Hierarchy Level | 16146
- Description | 16146
- Options | 16146
- Required Privilege Level | 16146
- Release Information | 16147

Syntax

```
rule-set rule-set-name {
    rule rule-name;
}
```

Hierarchy Level

```
[edit  services softwire]
```

Description

Specify the rule set the router uses when applying this service.

Options

rule-set-name—Identifier for the collection of rules that constitute this rule set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Softwire Rules

rule-set (Softwires Next Gen Services)

IN THIS SECTION

- Syntax | 16147
- Hierarchy Level | 16148
- Description | 16148
- Options | 16148
- Required Privilege Level | 16148
- Release Information | 16148

Syntax

```
rule-set rule-set-name {
    match-direction (input | output);
    rule rule-name {
        then {
            ds-lite ds-lite-concentrator-name
            map-e map-e-concentrator-name
            v6rd v6rd-softwire-concentrator;
        }

    }
}
```

Hierarchy Level

```
[edit services softwires]
```

Description

Configure a rule to apply a DS-Lite, MAP-E, or v6rd softwire concentrator to a flow.

Options

input	Apply the rule on the input side of the interface.
output	Apply the rule on the output side of the interface.
rule *rule-name*	Name of the rule.
rule-set *rule-set-name*	Name of the rule set that contains the rule.
ds-lite *ds-lite-softwire-concentrator*	Name of the softwire concentrator that the rule assigns to a flow.
map-e *map-e-softwire-concentrator*	Name of the softwire concentrator that the rule assigns to a flow.
v6rd *v6rd-softwire-concentrator*	Name of the softwire concentrator that the rule assigns to a flow.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

6rd Softwires in Next Gen Services

rule-sets (CoS AppQoS)

IN THIS SECTION

- Syntax | 16149
- Hierarchy Level | 16150
- Description | 16150
- Options | 16150
- Required Privilege Level | 16151
- Release Information | 16151

Syntax

```
rule-sets {
    rule-set-name {
        rule rule-name {
            match {
                application application-name;
                application-any;
                application-group application-group-name;
                application-known;
                application-unknown;
            }
            then {
                dscp-code-point dscp-value ;
                forwarding-class forwarding-class-name;
                log;
                loss-priority [ high | medium-high | medium-low | low ];
                rate-limit {
                    loss-priority-high;
                    client-to-server rate-limiter-name;
                    server-to-client rate-limiter-name;
                }
            }
        }
```

```
    }
}
```

Hierarchy Level

```
[edit class-of-service application-traffic-control]
[edit logical-systems logical-system-name class-of-service application-traffic-control]
[edit tenants tenant-name class-of-service application-traffic-control]
```

Description

Defines AppQoS rule sets and the rules that establish priorities based on quality-of-service requirements for the associated applications. AppQoS rules can be included in policy statements to implement application-aware quality of service control.

Options

- *rule-set-name*—Name used to refer to a collection of AppQoS rules.
- rule *rule-name*—Name applied to the match criteria and resulting actions that control the quality-of-service provided to any matching applications.
- application *application-name*—Name of the application to be used as match criteria for the rule.
- application-any —Any application encountering this rule. Note that when you use this specification, all application matching ends. Any application rule following this one will never be encountered.
- application-group *application-group-name*—Group of applications to be used as match criteria for the rule. Both applications and application groups can be match criteria for a single rule.
- application-known—Match criteria specifying any session that is identified, but its corresponding application is not specified.
- application-unknown—Match criteria specifying any session that is not identified.
- forwarding-class *forwarding-class-name*—The AppQoS class with which matching applications will be marked. This field identifies the rewriter that has marked the DSCP value . Therefore, the AppQoS forwarding class must be different from those used by IDP or firewall filters. With this class specified, firewall filter class will not overwrite the existing DSCP value.
- dscp-code-point—DSCP alias or bit map with which matching applications will be marked to establish the output queue. This value can be marked by rewriters from IDP, AppQoS, or a firewall filter. The

forwarding-class value identifies which rewriter has re-marked the packet with the current DSCP value. If a packet triggers all three rewriters, IDP takes precedence over AppQoS, which takes precedence over a firewall filter.

- `loss-priority`—Loss priority with which matching applications will be marked. This value is used to determine the likelihood that a packet would be dropped when encountering congestion. A high loss priority means that there is an 80% chance of packet loss in congestion. Possible values are high, medium-high, medium-low and low.
- `rate-limit`—Rate limiters to be associated with client-to-server and with server-to-client traffic for this application. The rate limiter profile defines maximum speed and volume limits for matching applications.
- `log`—AppQoS event logging.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the following hierarchy levels introduced in Junos OS Release 19.3R1: `[edit logical-systems` *`logical-system-name`* `class-of-service application-traffic-control]`, and `[edit tenants` *`tenant-name`* `class-of-service application-traffic-control]`.

RELATED DOCUMENTATION

Example: Configuring AppQoS

rule-sets (Security Application Firewall)

IN THIS SECTION

Syntax | 16152

- Hierarchy Level | 16152
- Description | 16153
- Options | 16153
- Required Privilege Level | 16153
- Release Information | 16154

Syntax

```
rule-sets rule-set-name {
    default-rule {
        (deny [block-message] | permit | reject [block-message]);
    }
    profile profile-name;
    rule rule-name {
        match {
            dynamic-application [system-application];
            dynamic-application-groups [system-application-group];
            ssl-encryption (any | yes | no);
        }
        then {
            deny {
                block-message block-message;
            }
            permit permit;
            reject {
                block-message block-message;
            }
        }
    }
}
```

Hierarchy Level

```
[edit security application-firewall]
```

Description

Configure the set of rules for the application firewall.

The application firewall is defined by a collection of rule sets. These rule sets can be defined independently and shared across network security policies. A rule set defines the rules that specify match criteria, including dynamic applications, and the action to be taken for matching traffic.

To implement an application firewall, you need to:

- Define one or more application firewall rule sets.
- Create rules for each rule set that permit, reject, or deny traffic based on the application ID.
- Configure a security policy to invoke the application firewall service and specify the rule set to be applied to permitted traffic.

The application firewall support in the policies provides additional security control for dynamic applications.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Options

rule-set-name	Name of the rule set.
profile *profile-name*	Profile for block message.
default-rule	Specify default rule.
rule	Specify security rule match-criteria

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1. Statement updated in Junos OS Release 12.1X44-D10 to include the `ssl-encryption` and `reject` options. The `block-message` options added in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Example: Configuring an Application Group for Application Firewall

rule (Softwire)

IN THIS SECTION

- Syntax | **16154**
- Hierarchy Level | **16155**
- Description | **16155**
- Options | **16155**
- Required Privilege Level | **16155**
- Release Information | **16155**

Syntax

```
rule rule-name {
    match-direction (input | output);
    term term-name {
        then {
            (ds-lite ds-lite-softwire-concentrator | map-e name | v6rd v6rd-softwire-
concentrator);
        }
    }
}
```

Hierarchy Level

```
[edit  services softwire],
[edit  services softwire rule-set rule-set-name]
```

Description

Configure a rule to apply a softwire concentrator for a flow.

Softwire rules are supported on the MS-DPC, MS-100, MS-400, and MS-500 line cards. Starting in Junos OS release 17.4R1, softwire rules for DS-Lite are supported on MX Series routers with MS-MPCs and MS-MICs.

Options

rule-name	Identifier for the collection of terms that constitute this rule.
input	Apply the rule match on the input side of the interface.
output	Apply the rule match on the output side of the interface.
ds-lite-softwire-concentrator	Specify the DS-Lite softwire concentrator to use.
map-e *name*	Specify the Mapping of Address and Port with Encapsulation (MAP-E) softwire concentrator to use.
v6rd-softwire-concentrator	Specify the 6rd softwire concentrator to use.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`map-e` *name* option introduced in Junos OS Release 18.2R1 for MX Series Routers with MPC and MIC interfaces.

RELATED DOCUMENTATION

Configuring Softwire Rules

rule (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 16156
- Hierarchy Level | 16157
- Description | 16157
- Required Privilege Level | 16157
- Release Information | 16157

Syntax

```
rule rule-name {
    match {
        application [application-name]
        destination-address (NAT Next Gen Services) address;
        destination-address-name address-name;
        source-address (address | any-unicast);
        source-address-name address-name;
    }
    then {
        source-nat {
            clat-prefix clat-prefix;
            filtering-type {
                endpoint-independent {
                    prefix-list [allowed-host] except [denied-host];
                }

            }
            mapping-type {
                endpoint-independent;
            }
```

```
            pool nat-pool-name;
            secure-nat-mapping {
                eif-flow-limit number-of-flows;
                mapping-refresh (inbound | inbound-outbound | outbound);
            }
        }
        syslog;
    }
}
```

Hierarchy Level

```
[edit services nat source rule-set rule-set]
```

Description

Configure a source NAT rule, which translates the source address of IP packets.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

rulebase-exempt

IN THIS SECTION

- Syntax | 16158
- Hierarchy Level | 16158

- Description | 16158
- Options | 16159
- Required Privilege Level | 16159
- Release Information | 16159

Syntax

```
rulebase-exempt {
    rule rule-name {
        description text;
        match (Security IDP Policy) {
            attacks (Security IPS Rulebase) {
                custom-attack-groups [attack-group-name];
                custom-attacks [attack-name];
                dynamic-attack-groups [attack-group-name];
                predefined-attack-groups [attack-group-name];
                predefined-attacks [attack-name];
            }
            destination-address (Security IDP Policy)([address-name] | any | any-ipv4 | any-ipv6);
            destination-except [address-name];
            from-zone (zone-name | any );
            source-address ([address-name] | any | any-ipv4 | any-ipv6);
            source-except [address-name];
            to-zone (zone-name | any);
        }
    }
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name]
```

Description

Configure the exempt rulebase to skip detection of a set of attacks in certain traffic.

NOTE: You must configure the IPS rulebase before configuring the exempt rulebase.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Starting in Junos OS Release 18.2R1, IDP policy is available within unified security policy. IDP policy configurations are simplified and made available under the unified policy as one of the policy. Configuring source or destination address, source and destination-except, from and to zone, or application is not required with unified policy, as the match happens in the security policy itself.

rulebase-ips

IN THIS SECTION

- Syntax | 16160
- Hierarchy Level | 16161
- Description | 16161
- Options | 16161
- Required Privilege Level | 16161
- Release Information | 16161

Syntax

```
rulebase-ips {
    rule (Security IPS Rulebase) rule-name {
        description text;
        match (Security IDP Policy) {
            application (application-name | any | default);
            attacks (Security IPS Rulebase) {
                custom-attack-groups [attack-group-name];
                custom-attacks [attack-name];
                dynamic-attack-groups [attack-group-name];
                predefined-attack-groups [attack-group-name];
                predefined-attacks [attack-name];
            }
            destination-address (Security IDP Policy) ([address-name] | any | any-ipv4 | any-ipv6);
            destination-except [address-name];
            from-zone (zone-name | any );
            source-address ([address-name] | any | any-ipv4 | any-ipv6);
            source-except [address-name];
            to-zone (zone-name | any);
        }
        terminal;
        then {
            action {
                class-of-service {
                    dscp-code-point number;
                     forwarding-class forwarding-class;
                }
                (close-client | close-client-and-server | close-server |drop-connection | drop-
packet | ignore-connection | mark-diffserv value | no-action | recommended);
            }
            ip-action (Security IDP Rulebase IPS) {
                (ip-block | ip-close | ip-notify);
                log;
                log-create;
                refresh-timeout;
                target (destination-address | service | source-address | source-zone | source-
zone-address | zone-service);
                timeout seconds;
            }
            notification {
                log-attacks {
```

```
                    alert;
                }
                packet-log (Security IDP Sensor Configuration) {
                    post-attack number;
                    post-attack-timeout seconds;
                    pre-attack number;
                }
            }
            severity (critical | info | major | minor | warning);
        }
    }
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name]
```

Description

Configure the IPS rulebase to detect attacks based on stateful signature and protocol anomalies.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

run

IN THIS SECTION

- Syntax | 16162
- Description | 16162
- Options | 16162
- Required Privilege Level | 16162
- Release Information | 16162

Syntax

```
run command
```

Description

Run a top-level CLI command without exiting from configuration mode.

Options

command	CLI top-level command.

Required Privilege Level

configure—To enter configuration mode.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding Configuration Mode

run-length

IN THIS SECTION

- Syntax | 16163
- Hierarchy Level | 16163
- Description | 16163
- Options | 16164
- Required Privilege Level | 16164
- Release Information | 16164

Syntax

```
run-length number;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring input],
[edit forwarding-options port-mirroring instance port-mirroring-instance-name input],
[edit forwarding-options  port-mirroring family (inet|inet6) input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance instance-name input]
```

Description

Set the number of samples following the initial trigger event. The configuration enables you to sample packets following those already being sampled.

NOTE: The `run-length` statement is not supported when you configure inline flow monitoring (by including the `inline-jflow` statement at the `[edit forwarding-options sampling instance instance-name family (inet | inet6) output]` hierarchy level).

Options

number—Number of samples.

- **Range:** 0 through 20
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Table Filters

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

Configuring Traffic Sampling on MX, M and T Series Routers

rx-buffers

IN THIS SECTION

- Syntax | **16165**
- Hierarchy Level | **16165**
- Description | **16165**
- Default | **16165**
- Options | **16166**
- Required Privilege Level | **16166**
- Release Information | **16166**

Syntax

```
rx-buffers (on | off);
```

Hierarchy Level

```
[edit interfaces interface-name ether-options configured-flow-control]
```

Description

Enable or disable an interface to generate and send Ethernet PAUSE messages. If you enable the receive buffers to generate and send PAUSE messages, when the receive buffers reach a certain level of fullness, the interface sends a PAUSE message to the connected peer. If the connected peer is properly configured, it stops transmitting frames to the interface on the entire link. When the interface receive buffer empties below a certain threshold, the interface sends a message to the connected peer to resume sending frames.

Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the `rx-buffers` statement with the `tx-buffers` statement to configure asymmetric Ethernet PAUSE on an interface. (Use the `flow-control` statement to enable symmetric PAUSE and the `no-flow-control` statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)

> **NOTE**: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).
>
> Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

Default

Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.

Options

`on | off`—Enable or disable an interface to generate and send Ethernet PAUSE messages.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

flow-control | 7267

tx-buffers | 30263

Configuring CoS Asymmetric Ethernet PAUSE Flow Control

Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

rx-connect-speed-when-equal (L2TP LAC)

IN THIS SECTION

- Syntax | 16167
- Hierarchy Level | 16167
- Description | 16167
- Required Privilege Level | 16167
- Release Information | 16167

Syntax

```
rx-connect-speed-when-equal
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Enable sending the receive connect speed, which is represented by AVP 38, even when its value is equal to that of the transmit connect speed, which is represented by AVP 24. By default, AVP 38 is sent from the LAC to the LNS during the establishment of an L2TP tunnel session, only when its value is different from AVP 24. You can override the default setting with this configuration statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

> *L2TP Subscriber Access Lines and Connection Speeds*
>
> *L2TP Subscriber Access Lines and Connection Speeds*

rx-enable

IN THIS SECTION

- Syntax | 16168
- Hierarchy Level | 16168
- Description | 16168
- Default | 16168
- Required Privilege Level | 16168
- Release Information | 16169

Syntax

```
expected-defect {
rx-enable ;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

Description

Enable the ethernet expected defect (ETH-ED) function to process the received EDM PDUs.

The remaining statements are explained separately. See CLI Explorer.

Default

The MEP does not process EDM PDUs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

connectivity-fault-management | 4556

show oam ethernet connectivity-fault-management mep-database | 22620

rx-max-duration

IN THIS SECTION

- Syntax | 16169
- Hierarchy Level | 16169
- Description | 16170
- Options | 16170
- Required Privilege Level | 16170
- Release Information | 16170

Syntax

```
expected-defect {
rx-max-duration ;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

Description

Duration to indicate the maximum acceptable value at which the loss of continuity alarms are suppressed. If the duration in the received EDM PDU exceeds this configured value then the duration value will be truncated to this configured value and loss of continuity (LoC) alarms shall be suppressed for this duration.

Options

Minimum value The minimum value at which the loss of continuity alarms will be suppressed is 120 *seconds*.

Minimum value The maximum acceptable value at which the loss of continuity alarms will be suppressed is 3600 *seconds*.

Default 900 *seconds*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

connectivity-fault-management | 4556

show oam ethernet connectivity-fault-management mep-database | 22620

rx-window-size (L2TP)

IN THIS SECTION

- Syntax | 16171
- Hierarchy Level | 16171
- Description | 16171
- Options | 16171
- Required Privilege Level | 16172
- Release Information | 16172

Syntax

```
rx-window-size packets;
```

Hierarchy Level

```
[edit services l2tp tunnel]
```

Description

Configure the L2TP receive window size for an L2TP tunnel.

Options

packets—Number of packets that a peer can transmit without receiving an acknowledgment from the router. By default, this value is set to 4 packets. If the receive window size is configured to its default value, the router does not send the Receive Window Size AVP (AVP 10) in the first tunnel negotiation packet that is sent to its peer.

- **Range:** 4 through 128
- **Default:** 4

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

L2TP LAC Subscriber Configuration

L2TP LNS Inline Service Interfaces

S

IN THIS SECTION

- s-nssai | **16259**
- s6a | **16260**
- sa-hold-time (Protocols MSDP) | **16262**
- sa-multicast (100-Gigabit Ethernet) | **16264**
- sa-multicast (PTX Series Packet Transport Routers) | **16266**
- sabit | **16267**
- saegw | **16269**
- saegw access-network-peers | **16270**
- saegw control-plane-peers | **16272**
- saegw-core-network-peers | **16276**
- saegw system | **16278**
- sak-rekey-interval | **16280**
- saml | **16281**

- saml-options | **16284**
- sample | **16285**
- sample-once | **16287**
- sample-rate | **16288**
- sample-type | **16290**
- sampling | **16291**
- sampling (Forwarding Options) | **16293**
- sampling-instance | **16297**
- sampling (Interfaces) | **16298**
- sampling (Protocols MPLS) | **16300**
- sanity-poll | **16302**
- sap | **16303**
- satellite (Junos Fusion Automatic Satellite Conversion) | **16305**
- satellite (Junos Fusion Satellite Device Homing) | **16307**
- satellite (Junos Fusion Satellite Software Upgrade Groups) | **16308**
- satellite-management (Junos Fusion) | **16310**
- satellite-policies | **16313**
- save | **16315**
- save (Configuration Mode) | **16317**
- saved-core-context (System) | **16320**
- saved-core-files (System) | **16321**
- sbfd | **16323**
- sbfd (protocols bfd) | **16324**
- sbl | **16326**
- sbl-default-server | **16327**
- scale-mode | **16328**
- scale-optimized | **16329**
- scale-mode | **16331**
- scan-extension | **16332**
- scan-mode | **16333**
- scan-options (Security Antivirus Avira Engine) | **16335**
- scan-options (Security Antivirus Juniper Express Engine) | **16336**

- scan-options (Security Antivirus Kaspersky Lab Engine) | **16338**
- scan-options (Security Antivirus Sophos Engine) | **16340**
- sccp | **16341**
- scheduler | **16343**
- scheduler (Dynamic Scheduler Maps) | **16344**
- scheduler (Fabric Queues) | **16345**
- scheduler (Interconnect Device) | **16347**
- scheduler-map-chassis | **16348**
- scheduler-map (CoS Virtual Channels) | **16350**
- scheduler-map (Dynamic Traffic Shaping) | **16351**
- scheduler-map (Fabric Queues) | **16353**
- scheduler-map-forwarding-class-sets | **16354**
- scheduler-map (Interfaces and Traffic-Control Profiles) | **16355**
- scheduler-map (Interfaces and Traffic-Control Profiles) | **16357**
- scheduler-maps (Definition) | **16358**
- scheduler-maps (Dynamic CoS Definition) | **16360**
- scheduler-name | **16361**
- scheduler (Scheduler Map) | **16362**
- scheduler (Security Policies) | **16364**
- scheduler (System Security Profile) | **16366**
- schedulers (CoS) | **16368**
- schedulers (CoS) | **16369**
- schedulers (Dynamic CoS Definition) | **16371**
- schedulers (Interfaces) | **16373**
- schedulers (Security Policies) | **16375**
- schema | **16376**
- schema-version (Flat-File Accounting Options) | **16378**
- scope | **16379**
- scope-policy | **16380**
- scope (Security IDP Chain Attack) | **16382**
- scope (Security IDP Custom Attack) | **16383**
- screen | **16384**

- screen (Security Zones) | **16390**
- scripts | **16392**
- sctp | **16396**
- sdsl-bytes | **16398**
- sdsl-overhead-adjust | **16400**
- search | **16402**
- search-filter | **16404**
- secondary (Label Switched Path) | **16405**
- secondary (Precision Time Protocol) | **16406**
- secondary (Adaptive Services Interfaces) | **16408**
- secondary connection (Identity Management Advanced Query) | **16409**
- secondary-dns | **16412**
- secondary-interface (Aggregated Inline Services) | **16413**
- secondary (Link Services IQ PIC Interfaces) | **16414**
- secondary (Protocols MPLS) | **16416**
- secondary-server | **16418**
- secondary-vtep-address | **16419**
- secondary-wins | **16420**
- secret (Access Profile) | **16422**
- secret-key-timeout | **16423**
- secret (RADIUS) | **16424**
- secret (RADIUS Client) | **16426**
- secret (RADIUS Server) | **16427**
- secret (Tunnel Profile) | **16429**
- secure | **16430**
- secure-access-port | **16432**
- secure-channel | **16435**
- secure-nat-mapping | **16436**
- secure-nat-mapping (Source NAT Next Gen Services) | **16438**
- secure-neighbor-discovery | **16439**
- secure-proxy | **16441**
- secure-wire (System Security Profile) | **16443**

- secured-port-block-allocation | **16444**
- securid-server | **16447**
- security | **16448**
- security-association (MACSec) | **16450**
- security-association | **16452**
- security-association (IPSec) | **16454**
- security-association (Junos-FIPS Software) | **16458**
- security-association (Junos OS) | **16460**
- security-intelligence (Services) | **16462**
- security-intelligence | **16464**
- security-intelligence (IDP) | **16465**
- security-intelligence-policy | **16467**
- security-intelligence(services) | **16469**
- security-level | **16480**
- security-level (Defining Access Privileges) | **16481**
- security-level (Generating SNMP Notifications) | **16482**
- security-log | **16484**
- security-log-percent-full | **16485**
- security-metadata-streaming | **16487**
- security-mode (SRX) | **16492**
- security-mode | **16494**
- security-model (Access Privileges) | **16496**
- security-model (Group) | **16497**
- security-model (SNMP Notifications) | **16499**
- security-name (Community String) | **16500**
- security-name (Security Group) | **16501**
- security-name (SNMP Notifications) | **16503**
- security-option (IDS Screen Next Gen Services) | **16504**
- security-package | **16505**
- security-profile | **16507**
- security-profile-resources | **16512**
- security-service (Security Forwarding Options) | **16513**

- security-to-group | **16515**
- security-zone | **16516**
- segment | **16519**
- segment-list | **16521**
- select | **16525**
- select-profile | **16527**
- selection-mode | **16528**
- selective | **16529**
- selective-vlan-switching | **16532**
- send-acct-status-on-config-change (Access Profile) | **16534**
- send (add-path) | **16535**
- send-critical-event | **16539**
- send-dyn-subscription-indicator (PCRF Partition) | **16540**
- send-network-family-indicator (PCRF Partition) | **16542**
- send-oam | **16544**
- send-origin-state-id (Diameter Applications) | **16545**
- send (Protocols RIP) | **16546**
- send (Protocols RIPng) | **16548**
- send-release-on-delete (DHCP Relay Agent) | **16550**
- sender-based-rpf (MBGP MVPN) | **16552**
- sendid-tlv | **16554**
- sensor-based-stats | **16555**
- sensor-based-stats (Junos Telemetry Interface) | **16558**
- sensor-configuration | **16560**
- sensor (Junos Telemetry Interface) | **16564**
- separator | **16593**
- seq-number-validated | **16594**
- sequence-check-required | **16595**
- sequence-number (Security IDP ICMP Headers) | **16597**
- sequence-number (Security IDP TCP Headers) | **16598**
- serial-number | **16599**
- serial-number (Junos Fusion) | **16601**

- serial-number (MX Series Virtual Chassis) | **16603**
- serial-options | **16604**
- server (Junos node slicing) | **16606**
- server | **16608**
- server-address | **16609**
- server-address (dhcp-client) | **16611**
- server-address (Hosted Services) | **16612**
- server (Batch Commits) | **16613**
- server-certificate (REST API) | **16615**
- server-certificate (Services) | **16617**
- server-certificate-subject | **16618**
- server-cluster (Security Group VPN Server) | **16619**
- server-connectivity | **16621**
- server (DHCP and BOOTP Relay Agent) | **16623**
- server (DNS, Port, and TFTP Service) | **16624**
- server-duid-type (DHCP Local Server) | **16626**
- server (Dynamic PPPoE) | **16627**
- server (Easy EVPN LAG Configuration) | **16628**
- server-group | **16633**
- server (icap-redirect profile) | **16635**
- server-id-override | **16637**
- server-identifier | **16639**
- server-inactivity-timeout | **16641**
- server (Junos OS) | **16642**
- server (Junos OS Evolved) | **16644**
- server-member-communication (Security Group VPN Server) | **16646**
- server (NTP) | **16648**
- server (pcp) | **16651**
- server-profile | **16654**
- server-profile (Active Flow Monitoring) | **16655**
- server (RADIUS Accounting) | **16656**
- server (RADIUS Network Elements) | **16660**

- server-reject-bridge-domain | server-reject-vlan | **16661**
- server-response-time (DHCP Relay Agent) | **16664**
- server (Security Antivirus) | **16665**
- server (Security Group VPN) | **16667**
- server (Security Sophos Engine Antivirus) | **16670**
- server (Security Web Filtering) | **16672**
- server (TACACS+ Accounting) | **16674**
- servers | **16677**
- servers (RADIUS) | **16678**
- service | **16680**
- service-agents (Analytics) | **16681**
- service-context-id (OCS) | **16684**
- service-deployment | **16686**
- service-device-pool (L2TP) | **16687**
- service-device-pools (Service Interfaces) | **16688**
- service-device (Remote Device Management) | **16690**
- service-domain | **16693**
- service (Dynamic Profiles) | **16694**
- service (Dynamic Service Sets) | **16695**
- service-filter (Dynamic Service Sets) | **16697**
- service-filter (Firewall) | **16699**
- service-filter (Interfaces) | **16701**
- service-groups | **16702**
- service-id (Bridge Domains) | **16704**
- service-id (Switch Options) | **16705**
- service-interface (L2TP Processing) | **16706**
- service-interface-pools | **16708**
- service-interface (Services Interfaces) | **16709**
- service-mode | **16711**
- service-mode (mobile-edge SAEGW) | **16713**
- service-name | **16714**
- service-name-table | **16716**

- service-name-tables | **16717**
- service-order | **16719**
- service-package | **16721**
- service-pics | **16723**
- service-port (UDP, Dynamic Flow Capture) | **16724**
- service-port (Inline Services) | **16725**
- service (PPPoE) | **16727**
- service-profile (DHCP Local Server) | **16729**
- service-profile (DHCP Relay Agent) | **16731**
- service-profile (L2TP) | **16733**
- service-profile (Static Subscribers) | **16735**
- service (Security IDP Anomaly Attack) | **16736**
- service (Security IDP Dynamic Attack Group) | **16737**
- service (Service Accounting) | **16739**
- service-set (Application-Aware Control Policy) | **16740**
- service-set (Dynamic Service Sets) | **16742**
- service-set (Interfaces) | **16744**
- service-set-options | **16746**
- service-set-options (Next Gen Services Services) | **16748**
- service-set (Services) | **16750**
- service-set (Subscriber-Aware) | **16754**
- service-set (TDF Interface) | **16756**
- service-tracking-action (Network Isolation Groups) | **16758**
- service-tracking (Network Isolation Groups) | **16760**
- service-type | **16761**
- service-type (RFC 2544 Benchmarking Tests) | **16763**
- service-type (MAC-VRF Routing Instances) | **16764**
- service-provider (SAML) | **16767**
- services (Captive Portal Content Delivery) | **16769**
- services (CoS) | **16772**
- services (Hierarchy) | **16773**
- services (IDS) | **16774**

- services (Interfaces) | **16776**
- services (IPsec VPN) | **16777**
- services (L2TP) | **16778**
- services (L2TP System Logging) | **16780**
- services-memory (resource-manager) | **16781**
- services (NAT) | **16783**
- services-offload | **16784**
- services-offload (Security) | **16786**
- services-options | **16787**
- services-options (Next Gen Services Interfaces) | **16789**
- services-redundancy-group | **16793**
- services (Stateful Firewall) | **16799**
- services (System Logging) | **16800**
- services (system services) | **16802**
- session-affinity | **16804**
- session-close | **16806**
- session-drop-hold-down | **16807**
- session-establishment-hold-time | **16808**
- session-expiry (MX Series in Enhanced LAN Mode) | **16809**
- session-group | **16810**
- session-id-change-limiter-indirect | **16813**
- session-init | **16815**
- session-limit | **16816**
- session-limit (IDS MS-DPC) | **16817**
- session-limit (IDS MS-MPC) | **16819**
- session-limit-per-username (Access Profile) | **16821**
- session-limit (Service Set Next Gen Services) | **16823**
- session-limit (System Services) | **16824**
- session-memory (resource-manager) | **16826**
- session-mode | **16828**
- session-offload | **16829**
- session-options | **16831**

- session-options (Access Profile) | **16835**
- session (Origin Validation for BGP) | **16836**
- session-persistence-scan | **16839**
- session-pics | **16840**
- session-pics (Diameter) | **16842**
- session-protection | **16843**
- session (Protocols LDP) | **16845**
- session-scan | **16847**
- session (Security IKE) | **16848**
- session (System Services) | **16853**
- session-timeout (Service Set Next Gen Services) | **16854**
- session-update-interval | **16856**
- session (Web Management) | **16857**
- sessions-limit-group (L2TP) | **16858**
- set | **16860**
- set chassis cluster cluster-id node node-number reboot | **16861**
- set chassis cluster disable reboot | **16863**
- set chassis display message | **16865**
- set chassis virtual-chassis | **16869**
- set cli complete-on-space | **16870**
- set cli directory | **16872**
- set cli idle-timeout | **16874**
- set cli logical-system | **16875**
- set cli prompt | **16877**
- set cli restart-on-upgrade | **16878**
- set cli screen-length | **16880**
- set cli screen-width | **16881**
- set cli terminal | **16882**
- set cli timestamp | **16884**
- set date | **16885**
- set date ntp | **16888**
- set-dont-fragment-bit (Services IPsec VPN) | **16890**

- set-dont-fragment-bit (Services Set) | **16891**
- set services jinsightd subscribe health-monitor | **16892**
- set services jinsightd traceoptions | **16894**
- identity-provider settings (SAML) | **16895**
- setup-protection | **16897**
- severity (Event Policy) | **16898**
- severity (Next Gen Services Service-Set Remote System Logging) | **16900**
- severity (Security IDP Custom Attack) | **16901**
- severity (Security IDP Dynamic Attack Group) | **16903**
- severity (Security IDP IPS Rulebase) | **16904**
- severity (Security Log) | **16906**
- sflow | **16907**
- sfm (Chassis) | **16912**
- sfm (Chassis Redundancy) | **16913**
- sfpplus | **16915**
- sftp-backup (OCS Partition) | **16916**
- sglimit | **16919**
- sham-link | **16921**
- sham-link-remote | **16923**
- shaping | **16925**
- shaping-rate (Traffic Control Profiles) | **16926**
- shaping-rate (Applying to an Interface) | **16929**
- shaping-rate (CoS Adaptive Shapers) | **16932**
- shaping-rate (CoS Interfaces) | **16933**
- shaping-rate (CoS Virtual Channels) | **16936**
- shaping-rate (Dynamic Traffic Shaping and Scheduling) | **16937**
- shaping-rate-excess-high | **16939**
- shaping-rate-excess-low | **16941**
- shaping-rate-excess-medium-high | **16943**
- shaping-rate-excess-medium-low | **16945**
- shaping-rate (Oversubscribing an Interface) | **16947**
- shaping-rate-priority-high | **16950**

- shaping-rate-priority-low | **16952**
- shaping-rate-priority-medium | **16954**
- shaping-rate-priority-medium-low | **16956**
- shaping-rate-priority-strict-high | **16958**
- shaping-rate (Schedulers) | **16960**
- shared-bandwidth-policer (Configuring) | **16962**
- shared-buffer | **16964**
- shared-instance | **16967**
- shared-key | **16968**
- shared-name | **16970**
- shared-scheduler | **16971**
- shared-secret (RADIUS Snoop Segment) | **16973**
- shdsl-options | **16974**
- shell | **16975**
- shellcode | **16977**
- shmlog (Shared Memory Log) | **16978**
- short-cycle-protection (DHCP Local Server and Relay Agent) | **16982**
- short-cycle-protection (Static and Dynamic Subscribers) | **16984**
- short-name-format | **16986**
- short-reach-mode | **16988**
- short-sequence | **16989**
- shortcuts (Protocols OSPF) | **16991**
- should-secure | **16992**
- show | **16994**
- show access-cac interface-set | **16995**
- show access-security router-advertisement state | **16998**
- show access-security router-advertisement statistics | **17001**
- show access-security slaac-snooping binding | **17003**
- show access-security slaac-snooping state | **17005**
- show access-security slaac-snooping statistics | **17007**
- show accounting pending-accounting-stops | **17011**
- show accounting profile | **17016**

- show accounting server statistics archival-transfer | **17023**
- show agent sensors | **17024**
- show amt statistics | **17033**
- show amt summary | **17037**
- show amt tunnel | **17039**
- show analyzer | **17044**
- show ancp cos | **17046**
- show ancp neighbor | **17054**
- show ancp statistics | **17066**
- show ancp subscriber | **17072**
- show ancp summary | **17083**
- show ancp summary neighbor | **17086**
- show ancp summary subscriber | **17089**
- show app-engine crash | **17091**
- show app-engine logs | **17093**
- show arp | **17096**
- show arp inspection statistics | **17103**
- show as-path | **17105**
- show as-path domain | **17111**
- show as-path summary | **17113**
- show auto-bandwidth | **17116**
- show auto-configuration out-of-band | **17120**
- show backup-selection | **17126**
- show bfd seamless session | **17132**
- show bfd session | **17134**
- show bfd subscriber session | **17144**
- show bgp bmp | **17151**
- show bgp group | **17153**
- show bgp group output-queues | **17166**
- show bgp group traffic-statistics | **17171**
- show bgp neighbor | **17175**
- show bgp output-scheduler | **17216**

- show bgp replication | 17218
- show bgp replication logical-system | 17223
- show bgp summary | 17227
- show bridge domain | 17237
- show bridge flood | 17240
- show bridge flood nexthops satellite | 17249
- show bridge flood satellite | 17253
- show bridge mac-ip-table | 17257
- show bridge mac-table | 17264
- show bridge satellite device | 17272
- show bridge statistics | 17276
- show captive-portal authentication-failed-users | 17281
- show captive-portal firewall | 17283
- show captive-portal interface | 17286
- show chassis adc | 17291
- show chassis adc (GNF) | 17295
- show chassis adc (Node Slicing) | 17297
- show chassis afeb | 17299
- show chassis alarms | 17302
- show chassis alarms | 17304
- show chassis beacon | 17334
- show chassis cfeb | 17337
- show chassis cluster control-interface status | 17341
- show chassis cluster control-plane statistics | 17351
- show chassis cluster data-plane interfaces | 17354
- show chassis cluster data-plane statistics | 17356
- show chassis cluster ethernet-switching interfaces | 17359
- show chassis cluster ethernet-switching status | 17361
- show chassis cluster information | 17364
- show chassis cluster information configuration-synchronization | 17378
- show chassis cluster information issu | 17381
- show chassis cluster interfaces | 17383

- show chassis cluster ip-monitoring status redundancy-group | **17392**
- show chassis cluster port-peering | **17396**
- show chassis cluster statistics | **17399**
- show chassis cluster status | **17405**
- show chassis craft-interface | **17410**
- show chassis dedicated-ukern-cpu | **17418**
- show chassis device-mode | **17419**
- show chassis enhanced-temperature-thresholds | **17423**
- show chassis environment | **17429**
- show chassis environment adc | **17553**
- show chassis environment cb (Security) | **17569**
- show chassis environment cb | **17573**
- show chassis environment ccg | **17613**
- show chassis environment fan | **17616**
- show chassis environment fpc | **17621**
- show chassis environment fpm | **17705**
- show chassis environment mcs | **17714**
- show chassis environment monitored | **17718**
- show chassis environment pcg | **17748**
- show chassis environment pdu | **17752**
- show chassis environment pem | **17760**
- show chassis environment power-supply-unit | **17782**
- show chassis environment psm | **17784**
- show chassis environment psu | **17798**
- show chassis environment routing-engine | **17800**
- show chassis environment scg | **17811**
- show chassis environment (Security) | **17818**
- show chassis environment sfb | **17825**
- show chassis environment sfm | **17844**
- show chassis environment sib | **17849**
- show chassis errors active | **17895**
- show chassis ethernet-switch (Security) | **17901**

- show chassis ethernet-switch | **17908**
- show chassis ethernet-switch interconnect-device cb | **17960**
- show chassis ethernet-switch interconnect-device fpc | **17985**
- show chassis fabric connectivity | **18022**
- show chassis fabric degradation | **18031**
- show chassis fabric degradation actions | **18033**
- show chassis fabric degraded-fabric-reachability | **18035**
- show chassis fabric destinations | **18037**
- show chassis fabric device | **18066**
- show chassis fabric errors | **18068**
- show chassis fabric faults recovery-actions | **18075**
- show chassis fabric feb | **18078**
- show chassis fabric fpcs | **18079**
- show chassis fabric map | **18091**
- show chassis fabric mesh | **18101**
- show chassis fabric plane (Security) | **18104**
- show chassis fabric plane | **18113**
- show chassis fabric plane-location (Security) | **18154**
- show chassis fabric plane-location | **18156**
- show chassis fabric reachability | **18168**
- show chassis fabric redundancy-mode | **18181**
- show chassis fabric sibs | **18183**
- show chassis fabric stream-info | **18203**
- show chassis fabric summary (Security) | **18209**
- show chassis fabric summary | **18212**
- show chassis fabric topology | **18220**
- show chassis fabric unreachable-destinations | **18257**
- show chassis fan | **18260**
- show chassis feb | **18278**
- show chassis firmware | **18283**
- show chassis forwarding | **18304**
- show chassis forwarding-options | **18306**

- show chassis fpc | **18311**
- show chassis fpc errors | **18387**
- show chassis fpc ether-types | **18392**
- show chassis fpc-feb-connectivity | **18397**
- show chassis fpc global pfe-id | **18400**
- show chassis fpc (GNF) | **18404**
- show chassis fpc (Node Slicing) | **18408**
- show chassis fpc pic-status (GNF) | **18411**
- show chassis fpc pic-status (Node Slicing) | **18413**
- show chassis fpc (View) | **18415**
- show chassis gnf | **18427**
- show chassis hardware | **18430**
- show chassis hardware (GNF) | **18453**
- show chassis hardware (View) | **18457**
- show chassis high-availability data-plane statistics | **18476**
- show chassis high-availability information | **18482**
- show chassis high-availability peer-info | **18492**
- show chassis high-availability prefix-srgid-table | **18495**
- show chassis high-availability services-redundancy-group | **18497**
- show chassis in-service-upgrade | **18507**
- show chassis jnu satellite | **18513**
- show chassis lcd | **18515**
- show chassis led | **18534**
- show chassis led satellite | **18551**
- show chassis location | **18557**
- show chassis mac-addresses | **18563**
- show chassis network-services | **18571**
- show chassis network-slices | **18575**
- show chassis network-slices fpcs | **18579**
- show chassis nonstop-upgrade | **18581**
- show chassis nonstop-upgrade node-group | **18585**
- show chassis oss-map | **18587**

- show chassis pic | **18588**
- show chassis pic-mode | **18628**
- show chassis pic (Security) | **18630**
- show chassis power | **18633**
- show chassis power (Security) | **18676**
- show chassis power-budget-statistics | **18681**
- show chassis power sequence | **18688**
- show chassis psd | **18690**
- show chassis realtime-ukern-thread | **18692**
- show chassis redundancy feb | **18694**
- show chassis redundant-power-system | **18698**
- show chassis reswatch | **18701**
- show chassis routing-engine | **18703**
- show chassis routing-engine (View) | **18720**
- show chassis satellite | **18727**
- show chassis satellite-cluster | **18743**
- show chassis satellite-cluster route | **18751**
- show chassis satellite-cluster statistics | **18754**
- show chassis satellite extended-port | **18759**
- show chassis satellite interface | **18765**
- show chassis satellite neighbor | **18772**
- show chassis satellite power-budget-statistics | **18779**
- show chassis satellite redundancy-group | **18783**
- show chassis satellite redundancy-group devices | **18787**
- show chassis satellite software | **18790**
- show chassis satellite statistics | **18794**
- show chassis satellite unprovision | **18801**
- show chassis satellite upgrade-group | **18809**
- show chassis scb | **18813**
- show chassis scb-ha-port | **18816**
- show chassis security uefi variables | **18821**
- show chassis sfb | **18823**

- show chassis sfb errors | **18828**
- show chassis sfm | **18831**
- show chassis sibs | **18835**
- show chassis spmb | **18853**
- show chassis spmb sibs | **18867**
- show chassis ssb | **18876**
- show chassis synchronization | **18879**
- show chassis synchronization (MX Series Routers) | **18894**
- show chassis temperature-thresholds (Security) | **18911**
- show chassis temperature-thresholds | **18914**
- show chassis usb storage | **18945**
- show chassis zones | **18947**
- show class-of-service | **18959**
- show class-of-service adjustment-control-profile | **18965**
- show class-of-service application-traffic-control counter | **18968**
- show class-of-service application-traffic-control rate-limiters | **18974**
- show class-of-service application-traffic-control statistics rate-limiter | **18986**
- show class-of-service application-traffic-control statistics rule | **18990**
- show class-of-service classifier | **18993**
- show class-of-service code-point-aliases | **18999**
- show class-of-service congestion-notification | **19001**
- show class-of-service dedicated-buffer-profile | **19007**
- show class-of-service drop-profile | **19009**
- show class-of-service fabric scheduler-map | **19011**
- show class-of-service fabric statistics | **19014**
- show class-of-service forwarding-class | **19019**
- show class-of-service forwarding-class-set | **19024**
- show class-of-service forwarding-table | **19026**
- show class-of-service forwarding-table classifier | **19031**
- show class-of-service forwarding-table classifier mapping | **19033**
- show class-of-service forwarding-table drop-profile | **19035**
- show class-of-service forwarding-table fabric scheduler-map | **19038**

- show class-of-service forwarding-table rewrite-rule | **19040**
- show class-of-service forwarding-table rewrite-rule mapping | **19042**
- show class-of-service forwarding-table scheduler-map | **19044**
- show class-of-service forwarding-table slice | **19047**
- show class-of-service forwarding-table slice mapping | **19049**
- show class-of-service forwarding-table traffic-class-map | **19051**
- show class-of-service fragmentation-map | **19054**
- show class-of-service interface | **19056**
- show class-of-service interface-set | **19100**
- show class-of-service l2tp-session | **19103**
- show class-of-service loss-priority-rewrite | **19106**
- show class-of-service monitoring-profile | **19108**
- show class-of-service multi-destination | **19110**
- show class-of-service policy-map | **19112**
- show class-of-service rewrite-rule | **19114**
- show class-of-service routing-instance | **19117**
- show class-of-service scheduler-hierarchy interface | **19119**
- show class-of-service scheduler-hierarchy interface-set | **19122**
- show class-of-service scheduler-map | **19125**
- show class-of-service scheduler-map-forwarding-class-sets | **19131**
- show class-of-service scheduler-resources fpc | **19134**
- show class-of-service shared-buffer | **19137**
- show class-of-service slice <slice_name> interface <interface_name> | **19140**
- show class-of-service traffic-class-map | **19142**
- show class-of-service traffic-control-profile | **19145**
- show class-of-service translation-table | **19151**
- show cli | **19157**
- show cli authorization | **19159**
- show cli directory | **19164**
- show cli history | **19166**
- show configuration | **19167**
- show configuration chassis cluster traceoptions | **19171**

- show configuration services analytics sensor ddos | **19173**
- show configuration smtp | **19174**
- show connections | **19176**
- show cos halp flow queue-stats | **19181**
- show database-replication statistics | **19183**
- show database-replication summary | **19185**
- show dcbx | **19189**
- show dcbx neighbors | **19191**
- show ddos-protection protocols | **19222**
- show ddos-protection protocols culprit-flows | **19239**
- show ddos-protection protocols flow-detection | **19249**
- show ddos-protection protocols isis | **19254**
- show ddos-protection protocols parameters | **19265**
- show ddos-protection protocols statistics | **19276**
- show ddos-protection protocols violations | **19292**
- show ddos-protection statistics | **19295**
- show ddos-protection version | **19298**
- show dhcp client binding | **19300**
- show dhcp client statistics | **19305**
- show dhcp relay active-leasequery | **19309**
- show dhcp relay binding (Security) | **19316**
- show dhcp relay binding | **19319**
- show dhcp relay lockout-entries | **19328**
- show dhcp relay statistics (Security) | **19331**
- show dhcp relay statistics | **19334**
- show dhcp-security arp inspection statistics | **19342**
- show dhcp-security binding | **19344**
- show dhcp-security binding ip-source-guard | **19348**
- show dhcp-security ipv6 binding | **19351**
- show dhcp-security ipv6 statistics | **19355**
- show dhcp-security neighbor-discovery-inspection statistics | **19359**
- show dhcp server active-leasequery | **19362**

- show dhcp server active-leasequery statistics | 19364
- show dhcp server binding | 19367
- show dhcp server lockout-entries | 19377
- show dhcp server statistics | 19381
- show dhcp snooping binding | 19387
- show dhcp snooping statistics | 19390
- show dhcpv6 client binding | 19392
- show dhcpv6 client statistics | 19395
- show dhcpv6 relay active-leasequery | 19398
- show dhcpv6 relay binding | 19404
- show dhcpv6 relay lockout-entries | 19417
- show dhcpv6 relay statistics | 19420
- show dhcpv6 server active-leasequery | 19426
- show dhcpv6 server active-leasequery statistics | 19428
- show dhcpv6 server binding | 19431
- show dhcpv6 server binding (View) | 19440
- show dhcpv6 server lockout-entries | 19446
- show dhcpv6 server statistics | 19450
- show dhcpv6 server statistics (View) | 19456
- show dhcpv6 snooping binding | 19460
- show dhcpv6 snooping statistics | 19463
- show diagnostics tdr | 19465
- show diameter | 19472
- show diameter function | 19482
- show diameter function statistics | 19488
- show diameter instance | 19493
- show diameter network-element | 19495
- show diameter network-element map | 19500
- show diameter peer | 19504
- show diameter peer map | 19510
- show diameter peer statistics | 19514
- show diameter route | 19519

- show | display | 19522
- show | display commit-scripts | 19527
- show | display inheritance | 19529
- show | display json | 19532
- show | display omit | 19535
- show | display rfc5952 | 19537
- show | display set | 19538
- show | display set relative | 19540
- show | display translation-scripts | 19542
- show | display xml | 19546
- show dot1x | 19548
- show dot1x accounting attribute | 19557
- show dot1x authentication-bypassed-users | 19561
- show dot1x authentication-failed-users | 19562
- show dot1x firewall | 19564
- show dot1x interface | 19567
- show dot1x static-mac-address | 19577
- show dot1x statistics | 19580
- show dual-phase-bootup-status | 19583
- show dvmrp interfaces | 19584
- show dvmrp neighbors | 19587
- show dvmrp prefix | 19589
- show dvmrp prunes | 19592
- show dynamic-profile session | 19595
- show dynamic-tunnels database | 19601
- show dynamic-tunnels pfe-tunnel-localization | 19609
- show edge-virtual-bridging | 19616
- show ephemeral-configuration | 19620
- show ethernet-switching evpn nd-statistics | 19623
- show ethernet-switching evpn nd-table | 19626
- show ethernet-switching flood | 19630
- show ethernet-switching interface | 19636

- show ethernet-switching interfaces | **19641**
- show ethernet-switching layer2-protocol-tunneling interface | **19652**
- show ethernet-switching layer2-protocol-tunneling statistics | **19654**
- show ethernet-switching layer2-protocol-tunneling vlan | **19658**
- show ethernet-switching mac-ip-table | **19660**
- show ethernet-switching mac-learning-log | **19670**
- show ethernet-switching mgrp-policy | **19676**
- show ethernet-switching redundancy-groups | **19682**
- show ethernet-switching statistics | **19688**
- show ethernet-switching statistics aging | **19692**
- show ethernet-switching statistics mac-learning | **19694**
- show ethernet-switching table | **19700**
- show ethernet-switching vxlan-tunnel-end-point esi | **19724**
- show ethernet-switching vxlan-tunnel-end-point remote | **19728**
- show ethernet-switching vxlan-tunnel-end-point source | **19730**
- show ethernet-switching vxlan-tunnel-end-point svlbnh | **19733**
- show evpn arp-table | **19736**
- show evpn database | **19740**
- show evpn flood | **19748**
- show evpn gbp-src-tag | **19749**
- show evpn igmp-snooping database | **19752**
- show evpn igmp-snooping proxy | **19756**
- show evpn instance | **19761**
- show evpn ip-prefix-database | **19779**
- show evpn l3-context | **19790**
- show evpn mac-ip-table | **19794**
- show evpn mac-table | **19798**
- show evpn mld-snooping database | **19800**
- show evpn multicast-snooping assisted-replication multihomed-peers | **19804**
- show evpn multicast-snooping assisted-replication next-hops | **19807**
- show evpn multicast-snooping assisted-replication replicators | **19811**
- show evpn multicast-snooping status | **19815**

- show evpn nd-table | **19818**
- show evpn oism | **19822**
- show evpn oism peg-df-status | **19826**
- show evpn oism spmsi-ad | **19829**
- show evpn p2mp | **19833**
- show evpn peer-gateway-macs | **19836**
- show evpn prefix | **19838**
- show evpn vpws-instance | **19840**
- show extensible-subscriber-services counters | **19849**
- show extensible-subscriber-services debug-information | **19851**
- show extensible-subscriber-services dictionary | **19853**
- show extensible-subscriber-services dictionary attributes | **19859**
- show extensible-subscriber-services dictionary services | **19863**
- show extensible-subscriber-services service | **19867**
- show extensible-subscriber-services sessions | **19869**
- show extension-service request-response clients | **19871**
- show extension-service request-response servers | **19874**
- show extension-service status | **19876**
- show fabric administration inventory | **19878**
- show fabric administration inventory director-group status | **19885**
- show fabric administration inventory infrastructure | **19893**
- show fabric administration inventory interconnect-devices | **19898**
- show fabric administration inventory node-devices | **19900**
- show fabric administration inventory node-groups | **19903**
- show fabric administration system mac-pool | **19906**
- show fabric inventory | **19908**
- show fabric session-host | **19912**
- show fib-local-accounting ip | **19914**
- show fib-streaming | **19915**
- show fib-streaming route-tables | **19921**
- show fib-streaming routes | **19924**
- show fib-streaming state | **19927**

- show fibre-channel fabric | 19929
- show fibre-channel fc2 sessions | 19932
- show fibre-channel fc2 statistics | 19936
- show fibre-channel fip | 19938
- show fibre-channel fip enode | 19946
- show fibre-channel fip fabric | 19951
- show fibre-channel fip fcf | 19956
- show fibre-channel fip interface | 19961
- show fibre-channel fip statistics | 19966
- show fibre-channel flogi fport | 19971
- show fibre-channel flogi nport | 19973
- show fibre-channel flogi statistics | 19977
- show fibre-channel interfaces | 19980
- show fibre-channel next-hops | 19985
- show fibre-channel proxy fabric-state | 19987
- show fibre-channel proxy login-table | 19993
- show fibre-channel proxy np-port | 19997
- show fibre-channel proxy statistics | 20003
- show fibre-channel routes | 20006
- show fip snooping | 20008
- show fip snooping enode | 20015
- show fip snooping fcf | 20021
- show fip snooping interface | 20024
- show fip snooping statistics | 20029
- show fip snooping vlan | 20034
- show fip vlan-discovery | 20040
- show firewall (QFX) | 20043
- show firewall | 20048
- show firewall filter version | 20060
- show firewall log | 20062
- show firewall policer | 20066
- show firewall prefix-action-stats | 20068

- show firewall templates-in-use | **20071**
- show firewall (View) | **20073**
- show flexible-tunnels profiles | **20077**
- show flight-recorder status | **20083**
- show forwarding-options enhanced-hash-key | **20086**
- show forwarding-options hash-key | **20095**
- show forwarding-options hyper-mode | **20100**
- show forwarding-options load-balance | **20102**
- show forwarding-options load-balance ecmp|trunk | **20107**
- show forwarding-options next-hop-group | **20111**
- show forwarding-options port-mirroring | **20116**
- show groups junos-defaults | **20120**
- show gtp tunnels | **20122**
- show hfrr profiles | **20128**
- show host | **20131**
- show iccp | **20133**
- show igmp group | **20135**
- show igmp interface | **20142**
- show igmp snooping data-forwarding | **20148**
- show igmp snooping evpn status | **20153**
- show igmp snooping interface | **20157**
- show igmp snooping membership | **20165**
- show igmp-snooping membership | **20173**
- show igmp snooping options | **20179**
- show igmp-snooping route | **20181**
- show igmp-snooping route (View) | **20185**
- show igmp snooping statistics | **20187**
- show igmp-snooping statistics | **20196**
- show igmp-snooping vlans | **20198**
- show igmp statistics | **20203**
- show ike security-associations | **20210**
- show ilmi statistics | **20217**

- show ingress-replication mvpn | **20221**
- show interfaces (NFX150) | **20223**
- show interfaces | **20227**
- show interfaces (Abstracted Fabric) | **20371**
- show interfaces (Adaptive Services) | **20401**
- show interfaces (Aggregated Ethernet) | **20411**
- show interfaces anchor-group (Aggregated Packet Forwarding Engine) | **20437**
- show interfaces (ATM) | **20442**
- show interfaces demux0 (Demux Interfaces) | **20497**
- show interfaces destination-class | **20517**
- show interfaces detail | **20521**
- show interfaces diagnostics optics | **20526**
- show interfaces diagnostics optics | **20535**
- show interfaces diagnostics optics (Security) | **20569**
- show interfaces diagnostics optics-applications | **20576**
- show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port) | **20580**
- show interfaces diagnostics optics-profile | **20645**
- show interfaces (Discard) | **20651**
- show interfaces (Dynamic Flow Capture) | **20660**
- show interfaces (Encryption) | **20667**
- show interfaces extensive | **20676**
- show interfaces extensive satellite-device | **20747**
- show interfaces fabric | **20774**
- show interfaces (far-end-interval) | **20803**
- show interfaces (Fast Ethernet) | **20805**
- show interfaces filters | **20832**
- show interfaces (Flow Collector) | **20835**
- show interfaces (Flow Monitoring) | **20844**
- show interfaces flow-statistics | **20852**
- show interfaces forwarding-class-counters | **20859**
- show interfaces fti | **20866**

- show interfaces ge | **20881**
- show interfaces (GRE) | **20900**
- show interfaces interface-set (Ethernet Interface Set) | **20913**
- show interfaces interface-set (PS Interface Set) | **20920**
- show interfaces interface-set queue | **20922**
- show interfaces interval | **20934**
- show interfaces (IP-over-IP) | **20938**
- show interfaces irb | **20945**
- show interfaces l2-routing-instance | **20957**
- show interfaces lb-stats af | **20959**
- show interfaces (Link Services) | **20965**
- show interfaces (Link Services IQ) | **20987**
- show interfaces load-balancing (Aggregated Multiservices) | **21024**
- show interfaces (Logical Tunnel) | **21029**
- show interfaces (Loopback) | **21040**
- show interfaces lsi (Label-Switched Interface) | **21051**
- show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet) | **21056**
- show interfaces mac-database | **21080**
- show interfaces mc-ae (MX Series) | **21090**
- show interfaces mc-ae | **21098**
- show interfaces me0 | **21101**
- show interfaces media | **21113**
- show interfaces (Multicast Tunnel) | **21117**
- show interfaces (Multilink Services) | **21124**
- show interfaces (PIM) | **21137**
- show interfaces policers | **21143**
- show interfaces port-profile | **21146**
- show interfaces (PPPoE) | **21154**
- show interfaces prbs-stats | **21170**
- show interfaces priority-group | **21172**
- show interfaces ps0 (Pseudowire Subscriber Interfaces) | **21175**

- show interfaces (PTX Series Packet Transport Routers) | **21182**
- show interfaces queue | **21208**
- show interfaces queue fabric | **21257**
- show interfaces redundancy | **21288**
- show interfaces (Redundant Adaptive Services) | **21292**
- show interfaces (Redundant Link Services IQ) | **21295**
- show interfaces routing | **21316**
- show interfaces routing-instance | **21325**
- show interfaces routing summary | **21328**
- show interfaces satellite-device | **21332**
- show interfaces (Serial) | **21339**
- show interfaces smart-sfp-defects | **21359**
- show interfaces smart-sfp-statistics | **21372**
- show interfaces snmp-index | **21389**
- show interfaces source-class | **21391**
- show interfaces (SRX Series) | **21395**
- show interfaces statistics | **21442**
- show interfaces statistics fabric | **21461**
- show interfaces statistics st0 | **21488**
- show interfaces summary | **21490**
- show interfaces swfabx | **21492**
- show interfaces targeting (Aggregated Ethernet for Subscriber Management) | **21494**
- show interfaces terse | **21497**
- show interfaces terse satellite-device | **21504**
- show interfaces terse zone | **21508**
- show interfaces transport pm | **21509**
- show interfaces vcp | **21524**
- show interfaces (View Aggregated Ethernet) | **21552**
- show interfaces (Virtual Loopback Tunnel) | **21567**
- show interfaces vlan | **21575**
- show interfaces voq | **21593**
- show interfaces xe | **21612**

- show interfaces xle | 21639
- show ip-source-guard | 21666
- show ipsec certificates | 21668
- show ipsec redundancy | 21672
- show ipsec security-associations | 21675
- show ipv6 neighbors | 21680
- show ipv6 neighbors (SRX Series) | 21683
- show ipv6 router-advertisement | 21685
- show ipv6-source-guard | 21692
- show isis adjacency | 21694
- show isis adjacency holddown | 21702
- show isis authentication | 21706
- show isis backup coverage | 21709
- show isis backup label-switched-path | 21711
- show isis backup spf results | 21714
- show isis context-identifier | 21720
- show isis database | 21722
- show isis flex-algorithm | 21732
- show isis hostname | 21736
- show isis interface | 21738
- show isis interface-group | 21746
- show isis layer2-map | 21748
- show isis overview | 21751
- show isis purge log | 21759
- show isis route | 21762
- show isis route download priority | 21768
- show isis spf | 21772
- show isis spring interface traffic-statistics | 21779
- show isis spring sensor info | 21780
- show isis statistics | 21783
- show l2-learning global-information | 21787
- show l2-learning global-mac-count | 21790

- show l2-learning instance | 21791
- show l2-learning interface | 21794
- show l2-learning mac-move-buffer | 21796
- show l2-learning redundancy-groups | 21798
- show l2circuit connections | 21805
- show l2pt interface | 21816
- show l2vpn connections | 21819
- show lacp interfaces | 21829
- show lacp interfaces (View) | 21836
- show lacp statistics interfaces (View) | 21843
- show lacp timeouts | 21845
- show ldp database | 21848
- show ldp database-label-requests | 21861
- show ldp fec-filters | 21864
- show ldp interface | 21866
- show ldp neighbor | 21868
- show ldp overview | 21871
- show ldp p2mp fec | 21879
- show ldp p2mp path | 21881
- show ldp p2mp tunnel | 21884
- show ldp path | 21886
- show ldp rib-groups | 21888
- show ldp route | 21890
- show ldp session | 21901
- show ldp statistics | 21912
- show ldp traffic-statistics | 21918
- show link-management | 21923
- show link-management peer | 21928
- show link-management routing | 21931
- show link-management statistics | 21935
- show link-management te-link | 21939
- show lldp | 21942

- show lldp local-information | **21952**
- show lldp neighbors | **21956**
- show lldp neighbors-vlan-name-tlv-list interface | **21973**
- show lldp remote-global-statistics | **21974**
- show lldp statistics | **21977**
- show log | **21980**
- show loop-detect enhanced interface | **21987**
- show mac-rewrite interface | **21993**
- show mac-vrf forwarding flood | **21996**
- show mac-vrf forwarding flood-group | **22000**
- show mac-vrf forwarding global-information | **22005**
- show mac-vrf forwarding global-mac-count | **22008**
- show mac-vrf forwarding global-mac-ip-count | **22010**
- show mac-vrf forwarding instance | **22011**
- show mac-vrf forwarding instance-mapping | **22018**
- show mac-vrf forwarding interface | **22020**
- show mac-vrf forwarding mac-ip-table | **22025**
- show mac-vrf forwarding mac-table | **22031**
- show mac-vrf forwarding mgrp-policy | **22035**
- show mac-vrf forwarding statistics | **22042**
- show mac-vrf forwarding vlans | **22046**
- show mac-vrf forwarding vxlan-tunnel-end-point esi | **22050**
- show mac-vrf forwarding vxlan-tunnel-end-point remote | **22054**
- show mac-vrf forwarding vxlan-tunnel-end-point svlbnh | **22063**
- show mld group | **22066**
- show mld interface | **22074**
- show mld snooping evpn database | **22081**
- show mld snooping interface | **22085**
- show mld snooping membership | **22090**
- show mld-snooping route | **22094**
- show mld snooping statistics | **22098**
- show mld-snooping vlans | **22100**

- show mld statistics | **22104**
- show modem wireless firmware | **22111**
- show modem wireless network | **22115**
- show modem wireless profiles | **22120**
- show monitor security flow | **22122**
- show mpls abstract-hop-membership | **22125**
- show mpls admin-groups | **22128**
- show mpls association | **22130**
- show mpls call-admission-control | **22132**
- show mpls container-lsp | **22136**
- show mpls context-identifer | **22147**
- show mpls correlation label | **22150**
- show mpls correlation nexthop-id | **22151**
- show mpls cspf | **22153**
- show mpls diffserv-te | **22156**
- show mpls egress-protection | **22159**
- show mpls interface | **22162**
- show mpls label usage | **22165**
- show mpls label usage label-range | **22169**
- show mpls lsp | **22172**
- show mpls lsp abstract-computation | **22203**
- show mpls lsp autobandwidth | **22206**
- show mpls path | **22210**
- show mpls srlg | **22212**
- show mpls static-lsp | **22214**
- show mpls tunnel manager statistics | **22219**
- show msdp | **22222**
- show msdp source | **22225**
- show msdp source-active | **22228**
- show msdp statistics | **22232**
- show multi-chassis mc-lag configuration-consistency | **22237**
- show multi-chassis mc-lag configuration-consistency global-config | **22243**

- show multi-chassis mc-lag configuration-consistency icl-config | 22246
- show multi-chassis mc-lag configuration-consistency list-of-parameters | 22249
- show multi-chassis mc-lag configuration-consistency mcae-config | 22262
- show multi-chassis mc-lag configuration-consistency vlan-config | 22267
- show multi-chassis mc-lag configuration-consistency vrrp-config | 22270
- show multicast backup-pe-groups | 22273
- show multicast ecid-mapping satellite | 22276
- show multicast flow-map | 22280
- show multicast forwarding-cache statistics | 22283
- show multicast interface | 22286
- show multicast mrinfo | 22290
- show multicast next-hops | 22292
- show multicast next-hops satellite | 22298
- show multicast pim-to-igmp-proxy | 22301
- show multicast pim-to-mld-proxy | 22304
- show multicast route | 22306
- show multicast rpf | 22323
- show multicast scope | 22328
- show multicast sessions | 22331
- show multicast snooping next-hops | 22335
- show multicast snooping next-hops satellite | 22338
- show multicast snooping route | 22344
- show multicast snooping route satellite | 22350
- show multicast statistics | 22356
- show multicast statistics satellite | 22363
- show multicast summary satellite | 22365
- show multicast usage | 22366
- show mvpn c-multicast | 22370
- show mvpn instance | 22375
- show mvpn neighbor | 22381
- show mvpn suppressed | 22388
- show mvrp | 22390

- show mvrp applicant-state | 22394
- show mvrp dynamic-vlan-memberships | 22398
- show mvrp interface | 22401
- show mvrp registration-state | 22403
- show mvrp statistics | 22406
- show neighbor-discovery-inspection statistics | 22413
- show network-access aaa accounting | 22415
- show network-access aaa radius-servers | 22417
- show network-access aaa statistics | 22428
- show network-access aaa statistics accounting | 22447
- show network-access aaa statistics authentication | 22449
- show network-access aaa statistics dynamic-requests | 22453
- show network-access aaa statistics pending-accounting-stops | 22455
- show network-access aaa statistics preauthentication | 22456
- show network-access aaa statistics re-authentication | 22459
- show network-access aaa subscribers | 22461
- show network-access aaa subscribers session-id | 22468
- show network-access aaa terminate-code | 22479
- show network-access address-assignment pool | 22486
- show network-access address-assignment pool (View) | 22488
- show network-access domain-map | 22490
- show network-access gx-plus | 22492
- show network access nas-5g | 22509
- show network-access nasreq statistics | 22514
- show network-access ocs | 22520
- show network-access pcrf | 22523
- show network-access radsec local-certificate | 22528
- show network-access radsec state | 22531
- show network-access radsec statistics | 22535
- show network-access requests pending | 22538
- show network-access requests statistics | 22542
- show network-access s6a | 22544

- show network-access securid-node-secret-file | 22547
- show network-agent statistics | 22549
- show network-access aaa saml assertion-cache | 22560
- show nlsd mode | 22562
- show node reboot | 22564
- show node statistics | 22565
- show nonstop-routing | 22571
- show nsagentd | 22575
- show ntp associations | 22577
- show ntp status | 22581
- show oam ethernet connectivity-fault-management delay-statistics | 22585
- show oam ethernet connectivity-fault-management forwarding-state | 22590
- show oam ethernet connectivity-fault-management interfaces | 22596
- show oam ethernet connectivity-fault-management linktrace path-database | 22611
- show oam ethernet connectivity-fault-management loss-statistics | 22615
- show oam ethernet connectivity-fault-management mep-database | 22620
- show oam ethernet connectivity-fault-management mep-statistics | 22637
- show oam ethernet connectivity-fault-management mip | 22654
- show oam ethernet connectivity-fault-management path-database | 22656
- show oam ethernet connectivity-fault-management policer | 22660
- show oam ethernet connectivity-fault-management sla-iterator-statistics | 22664
- show oam ethernet connectivity-fault-management synthetic-loss-statistics | 22677
- show oam ethernet evc | 22681
- show oam ethernet fnp interface | 22684
- show oam ethernet fnp messages | 22686
- show oam ethernet fnp status | 22688
- show oam ethernet link-fault-management | 22691
- show oam ethernet lmi | 22701
- show oam ethernet lmi statistics | 22704
- show oam fabric flow-specification | 22707
- show oam fabric interfaces | 22711
- show openflow capability | 22713

- show openflow controller | 22722
- show openflow filters | 22726
- show openflow flows | 22730
- show openflow groups | 22736
- show openflow interfaces | 22740
- show openflow statistics flows | 22745
- show openflow statistics groups | 22750
- show openflow statistics interfaces | 22753
- show openflow statistics packet | 22758
- show openflow statistics queue | 22760
- show openflow statistics summary | 22764
- show openflow statistics tables | 22766
- show openflow summary | 22771
- show openflow switch | 22773
- show ospf context-identifier | 22776
- show ospf database | 22780
- show ospf flex-algorithm | 22791
- show (ospf | ospf3) backup coverage | 22800
- show (ospf | ospf3) backup lsp | 22804
- show (ospf | ospf3) backup neighbor | 22807
- show (ospf | ospf3) backup spf | 22809
- show (ospf | ospf3) bgp-orr | 22813
- show (ospf | ospf3) interface | 22817
- show (ospf | ospf3) io-statistics | 22829
- show (ospf | ospf3) log | 22831
- show (ospf | ospf3) neighbor | 22835
- show (ospf | ospf3) overview | 22843
- show (ospf | ospf3) route | 22851
- show (ospf | ospf3) statistics | 22859
- show ospf spring prefix sid map | 22864
- show ospf spring sid database | 22866
- show ospf3 database | 22869

- show ovsdb commit failures | 22881
- show ovsdb controller | 22884
- show ovsdb interface | 22887
- show ovsdb logical-switch | 22890
- show ovsdb mac | 22893
- show ovsdb statistics interface | 22899
- show ovsdb tunnels | 22901
- show ovsdb virtual-tunnel-end-point | 22905
- show passive-monitoring error | 22908
- show passive-monitoring flow | 22911
- show passive-monitoring memory | 22915
- show passive-monitoring status | 22917
- show passive-monitoring usage | 22919
- show path-computation-client active-pce | 22922
- show path-computation-client lsp | 22927
- show path-computation-client statistics | 22936
- show path-computation-client status | 22944
- show performance-monitoring mpls lsp | 22947
- show pfe cfeb | 22955
- show pfe data | 22961
- show pfe feb | 22974
- show pfe filter hw summary | 22983
- show pfe fpc | 22986
- show pfe next-hop | 23013
- show pfe route | 23019
- show pfe sFlow egress-mcast-stats | 23033
- show pfe ssb | 23036
- show pfe statistics bridge | 23045
- show pfe terse | 23051
- show pfe version | 23054
- show pim bidirectional df-election | 23056
- show pim bidirectional df-election interface | 23060

- show pim bootstrap | 23064
- show pim interfaces | 23066
- show pim join | 23070
- show pim mdt | 23093
- show pim mdt data-mdt-joins | 23099
- show pim mdt data-mdt-limit | 23102
- show pim mvpn | 23104
- show pim neighbors | 23106
- show pim rps | 23113
- show pim snooping interfaces | 23124
- show pim snooping join | 23128
- show pim snooping neighbors | 23133
- show pim snooping statistics | 23141
- show pim source | 23147
- show pim statistics | 23152
- show poe controller | 23171
- show poe controller (View) | 23175
- show poe interface | 23177
- show poe notification-control | 23187
- show poe telemetries | 23189
- show policer (EX) | 23192
- show policer | 23194
- show policy | 23199
- show policy conditions | 23203
- show policy damping | 23206
- show ppp address-pool | 23209
- show ppp interface | 23211
- show ppp statistics | 23229
- show ppp summary | 23239
- show pppoe interfaces (Security) | 23241
- show pppoe interfaces | 23246
- show pppoe lockout | 23252

- show pppoe lockout atm-identifier | 23256
- show pppoe lockout vlan-identifier | 23260
- show pppoe service-name-tables | 23265
- show pppoe sessions | 23269
- show pppoe statistics | 23271
- show pppoe underlying-interfaces | 23275
- show pppoe version | 23285
- show programmable-rpd clients | 23287
- show programmable-rpd clients route-summary | 23289
- show programmable-rpd statistics | 23292
- show protection-group ethernet-ring aps | 23294
- show protection-group ethernet-ring configuration | 23300
- show protection-group ethernet-ring data-channel | 23309
- show protection-group ethernet-ring flush-info | 23312
- show protection-group ethernet-ring interface | 23315
- show protection-group ethernet-ring node-state | 23320
- show protection-group ethernet-ring statistics | 23328
- show protection-group ethernet-ring vlan | 23335
- show ptp aggregated-ethernet interfaces | 23342
- show ptp all master clock | 23344
- show ptp clock | 23348
- show ptp global-information | 23354
- show ptp hybrid | 23369
- show ptp lock-status | 23372
- show ptp master | 23377
- show ptp passive-port-monitor-status | 23383
- show ptp path-trace detail | 23389
- show ptp phy-timestamping-interfaces | 23390
- show ptp port | 23392
- show ptp quality-level-mapping | 23400
- show ptp slave | 23401
- show ptp stateful detail | 23406

- show ptp statistics | 23408
- show redundant-power-system led | 23414
- show redundant-power-system multi-backup | 23418
- show redundant-power-system network | 23419
- show redundant-power-system power-supply | 23420
- show redundant-power-system status | 23423
- show redundant-power-system upgrade | 23426
- show redundant-power-system version | 23427
- show redundant-trunk-group | 23430
- show remote-device-management service-devices | 23432
- show remote-device-management services | 23441
- show remote-device-management statistics | 23445
- show remote-device-management subscribers | 23451
- show remote-device-management summary | 23455
- show resource-manager | 23458
- show resource-manager cpu | 23460
- show resource-manager memory | 23461
- show rift database | 23463
- show rift flood-reduction | 23467
- show rift interface | 23469
- show rift node | 23476
- show rift path-computation | 23480
- show rift routes | 23482
- show rift tie | 23486
- show rift topology | 23488
- show rift versions | 23491
- show rip general-statistics | 23493
- show rip neighbor | 23496
- show rip statistics | 23499
- show ripng general-statistics | 23504
- show ripng neighbor | 23506
- show ripng statistics | 23508

- show route | 23512
- show route active-path | 23526
- show route advertising-protocol | 23533
- show route all | 23542
- show route aspath-regex | 23545
- show route best | 23548
- show route brief | 23553
- show route ccc | 23555
- show route community | 23557
- show route community-name | 23560
- show route cumulative | 23562
- show route damping | 23564
- show route detail | 23573
- show route exact | 23608
- show route export | 23612
- show route export vrf-target | 23615
- show route extensive | 23619
- show route extensive <route> expanded-nh | 23643
- show route flow validation | 23649
- show route forwarding-table | 23652
- show route forwarding-table family fibre-channel | 23673
- show route forwarding-table interface-name | 23676
- show route hidden | 23680
- show route inactive-path | 23685
- show route inactive-prefix | 23690
- show route instance | 23692
- show route label | 23698
- show route label-switched-path | 23705
- show route localization | 23707
- show route martians | 23709
- show route match-prefix | 23713
- show route next-hop | 23715

- show route no-community | 23719
- show route output | 23723
- show route pfe-install-error | 23727
- show route protocol | 23730
- show route range | 23740
- show route receive-protocol | 23746
- show route resolution | 23753
- show route rpm-tracking | 23759
- show route sla-tracking | 23764
- show route snooping | 23767
- show route source-gateway | 23771
- show route summary | 23779
- show route table | 23789
- show route tenant | 23821
- show route terse | 23824
- show route vpn-localization | 23829
- show routing nhdb pfe-install-error | 23835
- show rsvp interface | 23837
- show rsvp neighbor | 23847
- show rsvp pop-and-forward | 23854
- show rsvp route-session-id | 23857
- show rsvp session | 23860
- show rsvp statistics | 23878
- show rsvp version | 23887
- show sap listen | 23892
- show security advance-policy-based-routing detail | 23894
- show security advance policy based routing multipath policy statistics | 23898
- show security advance-policy-based-routing multipath rule | 23901
- show security advance-policy-based-routing profile | 23913
- show security advance-policy-based-routing sla active-probe-statistics | 23915
- show security advance-policy-based-routing sla profile (Application Name) | 23919
- show security advance-policy-based-routing sla profile (Next-Hop) | 23921

- show security advance-policy-based-routing sla profile (Status) | **23928**
- show security advance-policy-based-routing sla statistics | **23931**
- show security advance-policy-based-routing sla status | **23934**
- show security advance-policy-based-routing sla version | **23935**
- show security advance-policy-based-routing statistics | **23936**
- show security advance-policy-based-routing status | **23943**
- show security advanced-policy-based-routing policy-name | **23945**
- show security alarms | **23952**
- show security alg h323 counters | **23958**
- show security alg ike-esp-nat summary | **23963**
- show security alg msrpc | **23964**
- show security alg sccp calls | **23966**
- show security alg sccp counters | **23969**
- show security alg sip calls | **23973**
- show security alg sip counters | **23975**
- show security alg sip rate | **23996**
- show security alg status | **23999**
- show security alg status logical-system | **24007**
- show security application-firewall rule-set | **24012**
- show security application-firewall rule-set logical-system | **24017**
- show security application-tracking counters | **24021**
- show security casb policies | **24023**
- show security datapath-debug capture | **24025**
- show security datapath-debug counter | **24027**
- show security dns-cache | **24028**
- show security dynamic-address | **24031**
- show security dynamic-policies | **24045**
- show security dynamic-vpn users | **24053**
- show security dynamic-vpn users terse | **24055**
- show security firewall-authentication history | **24058**
- show security firewall-authentication history address | **24061**
- show security firewall-authentication history identifier | **24065**

- show security firewall-authentication jims | 24069
- show security firewall-authentication users | 24071
- show security firewall-authentication users address | 24075
- show security firewall-authentication users auth-type | 24080
- show security firewall-authentication users identifier | 24082
- show security flow cp-session | 24086
- show security flow cp-session destination-port | 24091
- show security flow cp-session destination-prefix | 24095
- show security flow cp-session family | 24098
- show security flow cp-session protocol | 24102
- show security flow cp-session source-port | 24105
- show security flow cp-session source-prefix | 24109
- show security flow gate | 24112
- show security flow gate brief node | 24118
- show security flow gate destination-port | 24127
- show security flow gate destination-prefix | 24131
- show security flow gate protocol | 24135
- show security flow gate summary node | 24140
- show security flow ip-action | 24147
- show security flow pmi statistics | 24159
- show security flow session | 24162
- show security flow session advanced-anti-malware | 24177
- show security flow session application | 24184
- show security flow session application-firewall | 24190
- show security flow session brief node | 24199
- show security flow session destination-port | 24205
- show security flow session destination-prefix | 24211
- show security flow session extensive node | 24217
- show security flow session family | 24227
- show security flow session idp family | 24235
- show security flow session idp summary | 24237
- show security flow session interface | 24240

- show security flow session nat | 24246
- show security flow session plugins | 24251
- show security flow session policy-id | 24254
- show security flow session pretty | 24259
- show security flow session protocol | 24263
- show security flow session resource-manager | 24270
- show security flow session services-offload | 24276
- show security flow session session-identifier | 24283
- show security flow session session-state | 24291
- show security flow session source-port | 24295
- show security flow session source-prefix | 24300
- show security flow session ssl | 24306
- show security flow session summary family | 24310
- show security flow session summary node | 24313
- show security flow session summary services-offload | 24322
- show security flow session tenant | 24327
- show security flow session tunnel | 24330
- show security flow statistics | 24341
- show security flow status | 24348
- show security flow tunnel inspection statistics | 24358
- show security forward-options resource-manager | 24360
- show security forwarding-options mirror-filter | 24363
- show security group-vpn member ike security-associations | 24366
- show security group-vpn member ipsec inactive-tunnels | 24372
- show security group-vpn member ipsec security-associations | 24375
- show security group-vpn member ipsec statistics | 24381
- show security group-vpn member kek security-associations | 24385
- show security group-vpn member policy | 24391
- show security group-vpn server ike security-associations | 24394
- show security group-vpn server ipsec security-associations | 24401
- show security group-vpn server kek security-associations | 24404
- show security group-vpn server registered-members | 24409

- show security group-vpn server server-cluster | 24412
- show security group-vpn server statistics | 24416
- show security gtp | 24419
- show security gtp counters | 24420
- show security gtp counters path-rate-limit | 24432
- show security gtp grouped-ie-profile | 24434
- show security gtp gsn statistics | 24436
- show security gtp handover-group | 24437
- show security gtp ie-set | 24439
- show security gtp ip-group | 24441
- show security gtp message-ie-profile-v1 | 24443
- show security gtp message-ie-profile-v2 | 24445
- show security gtp message-list | 24447
- show security gtp profile | 24448
- show security gtp rate-limit default | 24460
- show security gtp sctp association | 24461
- show security gtp sctp counters | 24465
- show security host-vpn security-associations | 24473
- show security host-vpn version | 24477
- show security idp active-policy | 24478
- show security idp attack attack-list | 24480
- show security idp attack attack-list policy | 24485
- show security idp attack context | 24490
- show security idp attack deprecated-list | 24491
- show security idp attack description | 24492
- show security idp attack detail | 24494
- show security idp attack group-list | 24498
- show security idp attack table | 24500
- show security idp attacks deprecated-attacks policy policy_name | 24502
- show security idp counters action | 24504
- show security idp counters application-identification | 24507
- show security idp counters dfa | 24514

- show security idp counters flow | 24517
- show security idp counters http-decoder | 24534
- show security idp counters ips | 24537
- show security idp counters log | 24545
- show security idp counters memory | 24552
- show security idp counters packet | 24556
- show security idp counters packet-log | 24567
- show security idp counters pdf-decoder | 24581
- show security idp counters policy-manager | 24585
- show security idp counters ssl-inspection | 24588
- show security idp counters tcp-reassembler | 24592
- show security idp jist-package-version | 24599
- show security idp logical system | 24600
- show security idp logical-system policy-association | 24602
- show security idp memory | 24603
- show security idp policies | 24605
- show security idp policy-commit-status | 24607
- show security idp policy-commit-status clear | 24610
- show security idp policy-templates-list | 24611
- show security idp predefined-attacks | 24613
- show security idp security-package-version | 24615
- show security idp ssl-inspection key | 24618
- show security idp ssl-inspection session-id-cache | 24621
- show security idp status | 24622
- show security idp status detail | 24625
- show security ike active-peer | 24630
- show security ike debug-status | 24638
- show security ike peers | 24640
- show security ike peers statistics | 24644
- show security ike pre-shared-key | 24647
- show security ike security-associations | 24648
- show security ike stats | 24671

- show security ike tunnel-map | 24681
- show security integrity | 24685
- show security internal-security-association | 24686
- show security ipsec control-plane-security-associations | 24688
- show security ipsec inactive-tunnels | 24691
- show security ipsec next-hop-tunnels | 24699
- show security ipsec security-associations | 24701
- show security ipsec statistics | 24744
- show security ipsec traffic-selector | 24753
- show security ipsec tunnel-distribution | 24755
- show security ipsec tunnel-events-statistics | 24761
- show security ipsec vpn vpnname ike idle time | 24763
- show security keychain | 24765
- show security log | 24769
- show security log file | 24773
- show security log profile | 24781
- show security log query | 24791
- show security log report | 24794
- show security log report in-detail | 24796
- show security log report in-interval | 24802
- show security log report summary | 24804
- show security log severity | 24807
- show security macsec connections (SRX) | 24809
- show security macsec connections | 24812
- show security macsec connections (MX Series) | 24815
- show security macsec statistics | 24819
- show security macsec statistics (SRX Series Firewalls) | 24825
- show security match-policies | 24831
- show security mka sessions | 24840
- show security mka sessions (MX Series) | 24844
- show security mka sessions (SRX Series Device | 24849
- show security mka sessions summary | 24852

- show security mka statistics | 24854
- show security mka statistics | 24857
- show security mka statistics (MX Series) | 24860
- show security monitoring | 24864
- show security monitoring fpc fpc-number | 24867
- show security monitoring performance session | 24871
- show security monitoring performance spu | 24873
- show security nat destination pool | 24877
- show security nat destination rule | 24881
- show security nat destination rule-application | 24886
- show security nat destination summary | 24889
- show security nat incoming-table | 24893
- show security nat interface-nat-ports | 24895
- show security nat resource-usage source-pool | 24900
- show security nat source deterministic | 24904
- show security nat source paired-address | 24907
- show security nat source persistent-nat-table | 24910
- show security nat source pool | 24914
- show security nat source port-block | 24927
- show security nat source rule | 24931
- show security nat source rule-application | 24937
- show security nat source summary | 24940
- show security nat static rule | 24948
- show security np-cache summary | 24955
- show security pki ca-certificate | 24957
- show security pki ca-certificate (View) | 24962
- show security pki certificate-request | 24968
- show security pki certificate-request (View) | 24972
- show security pki crl | 24975
- show security pki crl (View) | 24979
- show security pki local-certificate | 24982
- show security pki local-certificate | 24987

- show security pki local-certificate (View) | **24992**
- show security pki node-local certificate-request | **24999**
- show security pki node-local local-certificate | **25003**
- show security pki statistics | **25009**
- show security policies | **25028**
- show security policies checksum | **25056**
- show security policies file-serialization counters | **25059**
- show security policies hit-count | **25061**
- show security policies information | **25068**
- show security policies lookup-intact-on-commit | **25077**
- show security policies policy set | **25080**
- show security policies unknown-source-identity | **25083**
- show security policies zone-context | **25085**
- show security policy-report | **25089**
- show security re-distribution ipsec-vpn | **25093**
- show security resource-manager group active | **25095**
- show security resource-manager resource active | **25099**
- show security resource-manager settings | **25103**
- show security resource-manager summary | **25106**
- show security screen ids-option | **25108**
- show security screen statistics | **25116**
- show security screen status | **25132**
- show security screen white-list | **25133**
- show security shadow-policies | **25135**
- show security softwires | **25138**
- show security softwires map-e confidentiality status | **25140**
- show security softwires map-e domain | **25142**
- show security ssh key-pair-identity | **25144**
- show security tcp-encap connection | **25146**
- show security tcp-encap statistics | **25149**
- show security tpm status | **25150**
- show security tunnel inspection | **25153**

- show security user-identification device-provision authentication-source active-directory start 1 count 9 (match-string|prefix) | **25156**
- show security user-identification device-provision authentication-source jims start 1 count 9 (match-string|prefix) | **25158**
- show security user-identification local-authentication-table | **25159**
- show security user-identification role-provision all | **25163**
- show security user-identification role-provision authentication-source active-directory start 1 count 9 (match-string|prefix) | **25164**
- show security user-identification role-provision authentication-source jims start 1 count 9 (match-string|prefix) | **25166**
- show security user-identification source-identity-provision all | **25169**
- show security user-identification user-provision all | **25171**
- show security user-identification user-provision authentication-source active-directory start 1 count 9 (match-string|prefix) | **25172**
- show security user-identification user-provision authentication-source jims start 1 count 9 (match-string|prefix) | **25174**
- show security utm anti-spam statistics | **25176**
- show security utm anti-spam status | **25182**
- show security utm anti-virus statistics | **25183**
- show security utm anti-virus status | **25191**
- show security utm content-filtering statistics | **25195**
- show security utm custom-objects url-feed status | **25207**
- show security utm session | **25212**
- show security utm status | **25213**
- show security utm web-filtering category base-filter | **25215**
- show security utm web-filtering category category | **25217**
- show security utm web-filtering category status | **25219**
- show security utm web-filtering statistics | **25220**
- show security utm web-filtering status | **25228**
- show security zones | **25232**
- show security zones type | **25242**
- show service l2tp session routing-instance | **25246**
- show service user-identification authentication-source aruba-clearpass user-query counters | **25248**
- show service user-identification authentication-source aruba-clearpass user-query status | **25250**

- show service user-identification identity-management | 25252
- show services accounting | 25256
- show services accounting aggregation | 25260
- show services accounting aggregation template | 25266
- show services accounting aggregation template (View) | 25268
- show services accounting aggregation (View) | 25269
- show services accounting errors | 25270
- show services accounting flow | 25277
- show services accounting flow-detail | 25286
- show services accounting flow-detail (View) | 25293
- show services accounting memory | 25295
- show services accounting packet-size-distribution | 25297
- show services accounting status | 25300
- show services accounting usage | 25305
- show services advanced-anti-malware dynamic-filter status | 25308
- show services advanced-anti-malware machine-learning-scan | 25311
- show services advanced-anti-malware policy | 25313
- show services advanced-anti-malware profile | 25317
- show services advanced-anti-malware statistics | 25321
- show services advanced-anti-malware status | 25328
- show services agf fiveqi remarking | 25331
- show services agf nas | 25334
- show services agf ngap | 25338
- show services alg conversations | 25348
- show services alg statistics | 25356
- show services analytics agent | 25374
- show services anti-virus statistics | 25378
- show services application-aware-access-list flows | 25381
- show services application-aware-access-list statistics | 25386
- show services application-identification application | 25389
- show services application-identification application micro-applications | 25402
- show services application-identification application non-configurable | 25404

- show services application identification application obsolete applications | **25406**
- show services application-identification application-system-cache | **25408**
- show services application-identification application-system-cache (View) | **25413**
- show services application-identification commit-status | **25419**
- show services application-identification commit-status (Next Gen Services) | **25420**
- show services application-identification counter | **25422**
- show services application-identification counter (AppSecure) | **25426**
- show services application-identification entries | **25435**
- show services application-identification group | **25438**
- show services application-identification packet-capture counters | **25443**
- show services application-identification statistics application-groups | **25446**
- show services application-identification statistics applications | **25451**
- show services application-identification status | **25456**
- show services application-identification version | **25469**
- show services application-identification group obsolete-groups | **25472**
- show services application-identification recent-appid-sigpack-versions | **25473**
- show services application-identification recent-appid-sigpack-versions | **25474**
- show services captive-portal-content-delivery | **25476**
- show services cos statistics | **25483**
- show services cos statistics (Next Gen Services) | **25487**
- show services crtp | **25492**
- show services crtp flows | **25496**
- show services dns-filtering cache | **25498**
- show services dynamic-flow-capture content-destination | **25502**
- show services dynamic-flow-capture control-source | **25504**
- show services dynamic-flow-capture statistics | **25508**
- show services fixed-wireless-access statistics | **25512**
- show services flow-collector file interface | **25515**
- show services flow-collector input interface | **25519**
- show services flow-collector interface | **25521**
- show services ha detail | **25532**
- show services ha statistics | **25535**

- show services hcm pic-statistics | 25542
- show services hcm statistics | 25550
- show services icap-redirect server status | 25553
- show services icap-redirect statistic | 25554
- show services icap-redirect status | 25559
- show services ids | 25562
- show services inband-flow-telemetry | 25574
- show services inline ip-reassembly statistics | 25577
- show services inline-monitoring feature-profile-mapping fpc-slot | 25586
- show services inline-monitoring statistics fpc-slot | 25588
- show services inline nat pool | 25591
- show services inline nat statistics | 25594
- show services inline softwire statistics | 25598
- show services ip-monitoring status | 25604
- show services ipsec-vpn certificates | 25610
- show services ipsec-vpn ike security-associations | 25615
- show services ipsec-vpn ipsec security-associations | 25622
- show services ipsec-vpn ipsec statistics | 25632
- show services l2tp client | 25638
- show services l2tp destination | 25641
- show services l2tp destination lockout | 25647
- show services l2tp multilink | 25649
- show services l2tp radius | 25656
- show services l2tp session | 25661
- show services l2tp session-limit-group | 25674
- show services l2tp summary | 25676
- show services l2tp tunnel | 25685
- show services l2tp tunnel-group | 25694
- show services l2tp tunnel-switch destination | 25696
- show services l2tp tunnel-switch session | 25702
- show services l2tp tunnel-switch summary | 25709
- show services l2tp tunnel-switch tunnel | 25711

- show services l2tp user | 25718
- show services link-services cpu-usage | 25724
- show services lrf collector statistics | 25730
- show services lrf rule statistics | 25732
- show services lrf statistics | 25735
- show services lrf template | 25737
- show services mobile-edge anchor class-of-service | 25740
- show services mobile-edge peers | 25742
- show services mobile-edge resources call-admission-control statistics | 25747
- show services mobile-edge sessions | 25749
- show services mobile-edge summary | 25772
- show services monitoring rfc2544 | 25775
- show services monitoring rfc2544 | 25780
- show services monitoring rpm history-results | 25787
- show services monitoring rpm probe-results | 25794
- show services monitoring twamp client control-info | 25804
- show services monitoring twamp client history-results | 25807
- show services monitoring twamp client probe-results | 25814
- show services monitoring twamp client test-info | 25825
- show services monitoring twamp server control-info | 25828
- show services monitoring twamp server test-info | 25831
- show services nat destination pool | 25834
- show services nat destination rule | 25836
- show services nat destination summary | 25839
- show services nat deterministic-nat internal-host | 25842
- show services nat deterministic-nat nat-port-block | 25844
- show services nat ipv6-multicast-interfaces | 25847
- show services nat pool | 25850
- show services nat resource-usage source-pool | 25858
- show services nat source deterministic | 25861
- show services nat source mappings address-pooling-paired | 25864
- show services nat source mappings endpoint-independent | 25868

- show services nat source mappings pcp | 25871
- show services nat source mappings summary | 25873
- show services nat source pool | 25875
- show services nat source port-block | 25881
- show services nat source rule | 25884
- show services nat source rule-application | 25888
- show services nat source summary | 25891
- show services paa status | 25893
- show services pcef pic | 25899
- show services pcef subscribers | 25901
- show services pcp statistics | 25910
- show services policies | 25915
- show services policies detail | 25917
- show services policies hit-count | 25921
- show services policies interface | 25922
- show services policies service-set | 25924
- show services redundancy-group | 25925
- show services rpm active-servers | 25937
- show services rpm history-results | 25939
- show services rpm probe-results | 25944
- show services rpm rfc2544-benchmarking | 25960
- show services rpm rfc2544-benchmarking test-id | 25969
- show services rpm twamp client connection | 25995
- show services rpm twamp client history-results | 25997
- show services rpm twamp client probe-results | 26002
- show services rpm twamp client session | 26013
- show services rpm twamp server connection | 26016
- show services rpm twamp server session | 26018
- show services screen ids-option | 26021
- show services screen-statistics service-set | 26024
- show services security-intelligence | 26032
- show services security-intelligence category summary | 26034

- show services security-intelligence dns-statistics | **26037**
- show services security-intelligence update status | **26041**
- show services security-metadata-streaming | **26042**
- show services service-redirect statistic | **26048**
- show services service-set statistics ids drops | **26049**
- show services service-sets cpu-usage | **26058**
- show services service-sets memory-usage | **26061**
- show services service-sets plug-ins | **26063**
- show services service-sets statistic screen-drops (Next Gen Services) | **26065**
- show services service-sets statistic screen-session-limit-counters (Next Gen Services) | **26073**
- show services service-sets statistics ids session-limits counters | **26084**
- show services service-sets statistics integrity-drops | **26095**
- show services service-sets statistics jflow-log | **26101**
- show services service-sets statistics packet-drops | **26111**
- show services service-sets statistics syslog | **26114**
- show services service-sets statistics tcp | **26123**
- show services service-sets statistics tcp-mss | **26125**
- show services service-sets summary | **26127**
- show services sessions | **26129**
- show services sessions (Aggregated Multiservices) | **26141**
- show services sessions analysis | **26151**
- show services sessions analysis (USF) | **26156**
- show services sessions count | **26162**
- show services sessions (Next Gen Services) | **26163**
- show services sessions service-set | **26176**
- show services sessions softwire | **26178**
- show services sessions tcp-log | **26183**
- show services sessions utilization | **26185**
- show services soft-gre tunnel | **26186**
- show services softwire | **26190**
- show services softwire flows | **26192**
- show services softwire statistics | **26197**

- show services ssl certificate | 26209
- show services ssl droplogs | 26215
- show services ssl initiation counters | 26217
- show services ssl initiation profile | 26222
- show services ssl proxy certificate-cache entries | 26228
- show services ssl proxy certificate–cache statistics | 26231
- show services ssl proxy counters | 26233
- show services ssl proxy profile | 26239
- show services ssl proxy session-cache entries | 26243
- show services ssl proxy session-cache statistics | 26249
- show services ssl proxy statistics | 26252
- show services ssl proxy status | 26255
- show services ssl session | 26258
- show services ssl termination counters | 26262
- show services ssl termination profile | 26267
- show services stateful-firewall conversations | 26273
- show services stateful-firewall flow-analysis | 26279
- show services stateful-firewall flows | 26286
- show services stateful-firewall sip-call | 26294
- show services stateful-firewall sip-register | 26300
- show services stateful-firewall statistics | 26305
- show services stateful-firewall statistics application-protocol sip | 26318
- show services stateful-firewall subscriber-analysis | 26322
- show services subscriber analysis | 26327
- show services tcp-log | 26330
- show services tcp-log connections | 26331
- show services traffic-detection-function hcm statistics | 26333
- show services traffic-detection-function sessions | 26338
- show services traffic-load-balance statistics | 26341
- show services unified-access-control authentication-table | 26358
- show services unified-access-control counters | 26362
- show services unified-access-control policies | 26365

- show services unified-access-control roles | **26367**
- show services unified-access-control status | **26369**
- show services url-filter dns-resolution profile | **26370**
- show services url-filter dns-resolution-statistics profile template | **26374**
- show services url-filter statistics profile template | **26381**
- show services user-identification active-directory-access domain-controller status | **26385**
- show services user-identification active-directory-access statistics | **26389**
- show services user-identification active-directory-access user-group-mapping | **26393**
- show services user-identification authentication-table | **26397**
- show services user-identification device-information table | **26421**
- show services user-identification logical-domain-identity-management | **26426**
- show services user-identification validate-statistics | **26430**
- show services video-monitoring mdi errors fpc-slot | **26431**
- show services video-monitoring mdi flows fpc-slot | **26434**
- show services video-monitoring mdi stats fpc-slot | **26441**
- show services web-filter dns-resolution profile | **26445**
- show services web-filter dns-resolution-statistics profile template | **26449**
- show services web-filter secintel-policy status | **26455**
- show services web-filter statistics dns-filter-template | **26461**
- show services web-filter statistics profile | **26464**
- show services web-proxy dns forwarding-cache | **26470**
- show services web-proxy dns global-cache statistics | **26473**
- show services web-proxy session | **26476**
- show sflow | **26478**
- show sflow collector | **26482**
- show sflow interface | **26485**
- show snmp health-monitor | **26489**
- show snmp inform-statistics | **26499**
- show snmp mib | **26502**
- show snmp rmon | **26516**
- show snmp rmon history | **26523**
- show snmp statistics | **26529**

- show snmp stats-response-statistics | **26539**
- show snmp v3 | **26544**
- show spanning-tree bridge | **26548**
- show spanning-tree interface | **26555**
- show spanning-tree mstp configuration | **26568**
- show spanning-tree statistics | **26572**
- show spanning-tree statistics bridge | **26575**
- show spanning-tree statistics interface | **26576**
- show spanning-tree statistics message-queues | **26578**
- show spanning-tree stp-buffer see-all | **26580**
- show spring-traffic-engineering | **26587**
- show-spring-traffic-statistics | **26600**
- show ssl-certificates | **26602**
- show static-subscribers sessions | **26605**
- show subscribers | **26609**
- show subscribers summary | **26663**
- show synchronous-ethernet esmc statistics | **26674**
- show synchronous-ethernet esmc transmit | **26676**
- show synchronous-ethernet global-information | **26679**
- show system alarms | **26682**
- show system alarms | **26683**
- show system anomalies (GNF) | **26689**
- show system anomalies gnf-id | **26692**
- show system applications (Junos OS Evolved) | **26695**
- show system audit | **26709**
- show system auto-snapshot | **26720**
- show system autoinstallation status | **26723**
- show system autorecovery state | **26725**
- show system boot console | **26727**
- show system boot-messages | **26729**
- show system buffers | **26741**
- show system certificate | **26751**

- show system commit | **26754**
- show system commit server queue | **26758**
- show system commit server status | **26763**
- show system configuration archival | **26765**
- show system configuration database status | **26766**
- show system configuration database usage | **26768**
- show system configuration rescue | **26770**
- show system configuration revision | **26772**
- show system connections | **26775**
- show system core-dumps | **26783**
- show system core dumps (Junos OS Evolved) | **26802**
- show system cpu (JDM) | **26804**
- show system data-models | **26807**
- show system directory-usage | **26808**
- show system download | **26815**
- show system errors | **26817**
- show system errors active | **26822**
- show system errors count | **26829**
- show system errors error-id | **26831**
- show system errors fru | **26834**
- show system errors history | **26841**
- show system errors inactive | **26844**
- show system filesystem encryption status | **26852**
- show system firmware | **26855**
- show system firmware (Junos OS Evolved) | **26865**
- show system health-monitor | **26868**
- show system information | **26873**
- show system inventory hardware cpu | **26874**
- show-system-inventory-hardware-memory | **26881**
- show-system-inventory-hardware-optics | **26884**
- show-system-inventory-hardware-storage | **26886**
- show system khms-stats | **26889**

- show system license (View) | **26896**
- show system login lockout | **26902**
- show system mac-addresses (JDM) | **26903**
- show system max-smbios-version (JDM) | **26907**
- show system memory | **26908**
- show system memory (JDM) | **26918**
- show system monitor memory actions | **26921**
- show system monitor memory events | **26923**
- show system monitor memory status | **26926**
- show system name-resolution | **26929**
- show system network (JDM) | **26931**
- show system node-attributes | **26936**
- show system nodes | **26939**
- show system ntp threshold | **26941**
- show system packet-forwarding-options hw-resource-monitor monitor-profile | **26943**
- show system packet-forwarding-options hw-resource-monitor resource-list | **26946**
- show system packet-forwarding-options hw-resource-monitor utilization-info | **26948**
- show system packet-forwarding-options hw-resource-utilization-info | **26952**
- show system processes | **26957**
- show system queues | **26975**
- show system random-mac-prefix (JDM) | **26981**
- show system reboot | **26982**
- show system resource-monitor | **26987**
- show system resource-monitor ifd-cos-queue-mapping fpc | **26992**
- show system resource-monitor subscribers-limit | **26995**
- show system rollback | **27001**
- show system schema | **27004**
- show system security-profile | **27008**
- show system security-profile scheduler | **27016**
- show system security-profile secure-wire | **27020**
- show system security-profile security-log-stream-number | **27025**
- show system security-profile security-log-stream-number detail | **27028**

- show system security-profile security-log-stream-number summary | **27032**
- show system services dhcp binding | **27034**
- show system services dhcp client | **27038**
- show system services dhcp conflict | **27043**
- show system services dhcp global | **27045**
- show system services dhcp pool | **27047**
- show system services dhcp relay-statistics | **27051**
- show system services dhcp statistics | **27053**
- show system services dns dns-proxy | **27057**
- show system services dynamic-dns | **27061**
- show system services outbound-ssh jefflee | **27063**
- show system services service-deployment | **27067**
- show system snapshot (Junos OS) | **27069**
- show system snapshot (Junos OS Evolved) | **27073**
- show system snapshot (Junos OS with Upgraded FreeBSD) | **27075**
- show system snapshot media | **27078**
- show system software | **27081**
- show system software add-restart (Junos OS Evolved) | **27086**
- show system software list | **27088**
- show system software recovery-package | **27092**
- show system software restore-point-status | **27094**
- show system software upgrade status | **27096**
- show system software usb-software-version | **27098**
- show system statistics | **27101**
- show system statistics arp | **27119**
- show system statistics icmp | **27129**
- show system statistics icmp6 | **27136**
- show system statistics igmp | **27146**
- show system statistics ip | **27151**
- show system statistics ip6 | **27162**
- show system statistics netisr per-priority-stats | **27173**
- show system statistics netisr qdrops | **27174**

- show system statistics netisr socket-drops | 27176
- show system statistics tcp | 27178
- show system storage | 27190
- show system storage partitions | 27198
- show system storage satellite | 27202
- show system subscriber-management redundancy-state dhcp active-leasequery interface | 27205
- show system subscriber-management redundancy-state interface | 27208
- show system subscriber-management resiliency | 27211
- show system subscriber-management route | 27222
- show system subscriber-management statistics | 27229
- show system subscriber-management summary | 27242
- show system switchover | 27246
- show system unified-services status | 27251
- show system uptime | 27253
- show system users | 27260
- show system virtual-memory | 27266
- show system visibility cpu | 27279
- show system visibility host | 27286
- show system visibility jcp | 27296
- show system visibility jdm | 27300
- show system visibility memory | 27304
- show system visibility network | 27308
- show system visibility vnf | 27316
- show system yang package | 27323
- show system ztp | 27325
- show task | 27331
- show task io | 27336
- show task logical-system-mux | 27339
- show task memory | 27342
- show task replication | 27348
- show tcp-forwarding status | 27351
- show ted database | 27357

- show-ted-ipv4-prefix | 27372
- show-ted-ipv6-prefix | 27374
- show ted link | 27377
- show ted protocol | 27384
- show ted spring-te-policy | 27386
- show trace | 27392
- show unified-edge tdf aaa radius client statistics | 27396
- show unified-edge tdf aaa radius client status | 27404
- show unified-edge tdf aaa radius network-element statistics | 27406
- show unified-edge tdf aaa radius server statistics | 27410
- show unified-edge tdf aaa radius server status | 27416
- show unified-edge tdf aaa radius snoop-segment statistics | 27420
- show unified-edge tdf aaa statistics | 27425
- show unified-edge tdf address-assignment pool | 27437
- show unified-edge tdf address-assignment service-mode | 27442
- show unified-edge tdf address-assignment statistics | 27445
- show unified-edge tdf call-admission-control statistics | 27448
- show unified-edge tdf call-rate statistics | 27452
- show unified-edge tdf diameter network-element statistics | 27456
- show unified-edge tdf diameter network-element status | 27459
- show unified-edge tdf diameter pcc-gx statistics | 27462
- show unified-edge tdf diameter peer statistics | 27469
- show unified-edge tdf diameter peer status | 27476
- show unified-edge tdf domain service-mode | 27481
- show unified-edge tdf domain statistics | 27484
- show unified-edge tdf resource-manager clients | 27490
- show unified-edge tdf service-mode | 27493
- show unified-edge tdf statistics | 27496
- show unified-edge tdf status | 27507
- show unified-edge tdf subscribers | 27513
- show unified-edge tdf system interfaces | 27532
- show unified-edge tdf system interfaces service-mode | 27534

- show uplink-failure-detection | **27538**
- show uplink-failure-detection | **27541**
- show v4ov6-tunnels | **27544**
- show validation database | **27547**
- show validation group | **27551**
- show validation replication database | **27553**
- show validation session | **27556**
- show validation statistics | **27563**
- show version (JDM) | **27567**
- show version (Junos OS) | **27569**
- show version (Junos OS Evolved) | **27572**
- show version vnf | **27575**
- show virtual-chassis | **27578**
- show virtual-chassis active-topology | **27587**
- show virtual-chassis active-topology (MX Series Virtual Chassis) | **27593**
- show virtual-chassis device-topology | **27596**
- show virtual-chassis device-topology (MX Series Virtual Chassis) | **27605**
- show virtual-chassis fast-failover | **27609**
- show virtual-chassis heartbeat (MX Series Virtual Chassis) | **27611**
- show virtual-chassis login | **27616**
- show virtual-chassis mode | **27617**
- show virtual-chassis protocol adjacency | **27622**
- show virtual-chassis protocol adjacency (MX Series Virtual Chassis) | **27627**
- show virtual-chassis protocol database | **27633**
- show virtual-chassis protocol database (MX Series Virtual Chassis) | **27638**
- show virtual-chassis protocol interface | **27647**
- show virtual-chassis protocol interface (MX Series Virtual Chassis) | **27650**
- show virtual-chassis protocol route | **27655**
- show virtual-chassis protocol route (MX Series Virtual Chassis) | **27659**
- show virtual-chassis protocol statistics | **27663**
- show virtual-chassis protocol statistics (MX Series Virtual Chassis) | **27668**
- show virtual-chassis status (MX Series Virtual Chassis) | **27672**

- show virtual-chassis vc-path | 27676
- show virtual-chassis vc-port | 27681
- show virtual-chassis vc-port diagnostics optics | 27688
- show virtual-chassis vc-port (MX Series Virtual Chassis) | 27709
- show virtual-chassis vc-port statistics | 27713
- show virtual-network-functions | 27723
- show virtual-network-functions (Junos Node Slicing) | 27728
- show vlans | 27738
- show vlans (JDM) | 27767
- show vlans evpn nd-table | 27769
- show vlans vlan-groups | 27774
- show vmhost bridge | 27775
- show vmhost connections | 27778
- show vmhost control-plane | 27780
- show vmhost crash | 27782
- show vmhost crash (NFX Series) | 27784
- show vmhost forwarding-options analyzer | 27785
- show vmhost hard-disk-test | 27787
- show vmhost hardware | 27789
- show vmhost information | 27791
- show vmhost jdm status (In-Chassis Model) | 27793
- show vmhost logs | 27795
- show vmhost management-if | 27797
- show vmhost memory | 27798
- show vmhost mode | 27799
- show vmhost netstat | 27811
- show vmhost network external-network | 27813
- show vmhost processes | 27815
- show vmhost resource-usage | 27818
- show vmhost snapshot | 27821
- show vmhost snmp mib walk | 27823
- show vmhost status | 27825

- show vmhost status (NFX Series) | **27827**
- show vmhost status (In-Chassis Model) | **27829**
- show vmhost storage | **27831**
- show vmhost uptime | **27838**
- show vmhost uptime (NFX Series) | **27840**
- show vmhost version | **27841**
- show vmhost version (NFX Series) | **27844**
- show vmhost vlans | **27847**
- show vpls connections | **27849**
- show vpls flood event-queue | **27865**
- show vpls flood instance | **27867**
- show vpls flood route | **27871**
- show vpls mac-move-action | **27874**
- show vpls mac-table | **27876**
- show vpls statistics | **27884**
- show vrrp | **27888**
- show vrrp | **27905**
- show vrrp track | **27920**
- show wlan access-points | **27926**
- shutdown (BPDU Block) | **27935**
- shutdown (Protocols BGP) | **27936**
- shutdown-threshold | **27937**
- sib | **27939**
- signal-bandwidth | **27941**
- signal-degrade | **27942**
- signal-degrade-monitor-enable | **27944**
- signal-type | **27946**
- signal-type (BITS Interfaces) | **27947**
- signaling | **27949**
- signature | **27950**
- signature (Application Identification) | **27952**
- Signature Hash Algorithm (Security IKE) | **27953**

- signature (Security IDP) | **27955**
- simple-filter | **27962**
- simple-filter (Applying to an Interface) | **27964**
- simple-filter (Interfaces) | **27965**
- simple-mail-client-service | **27967**
- single-home (Junos Fusion) | **27968**
- single-rate | **27969**
- sip | **27971**
- sip (Application Profile) | **27973**
- sip-call-hold-timeout | **27974**
- sip-password | **27976**
- sip (Security) | **27977**
- sip-server | **27979**
- sip (Services CoS Next Gen Services) | **27981**
- sip-user-id | **27983**
- site-identifier (Layer 2 Circuits) | **27984**
- site-identifier (VPLS) | **27985**
- site (Layer 2 Circuits) | **27987**
- site-preference | **27989**
- site-range | **27990**
- site-reputation-action | **27992**
- site (VPLS Multihoming for FEC 128) | **27994**
- site (VPLS Multihoming for FEC 129) | **27995**
- size | **27997**
- size (VNF) | **27999**
- size (Interface) | **28000**
- size (Next Gen Services Global System Logging) | **28001**
- size (Sampling and Traceoptions) | **28003**
- size (Security Web Filtering Cache) | **28004**
- size (Services) | **28006**
- size (System) | **28007**
- skew-timer-disable | **28009**

- skip-arp-iteration (RFC 2544 Benchmarking) | **28010**
- sla-iterator-profile | **28012**
- sla-iterator-profiles | **28013**
- sla-iterator-profiles (OAM CFM) | **28015**
- sla-options | **28017**
- sla-rule | **28019**
- sla-tracking | **28027**
- slaac-snooping | **28030**
- slamon-services | **28032**
- slave | **28034**
- slice (CoS Interfaces) | **28039**
- slice (export routing policy action) | **28041**
- slice (firewall filter action) | **28042**
- slice (firewall filter match condition) | **28043**
- slow-pfe-alarm | **28044**
- smart-check | **28045**
- smart-optimize-timer | **28048**
- smg-service (Enhanced Subscriber Management) | **28050**
- smtp-profile (Security Content Security Policy Antispam) | **28051**
- smtp-profile (Security Content Security Policy Antivirus) | **28052**
- smtp-profile (Security Content Security Policy Content Filtering) | **28053**
- snapshot | **28055**
- snext | **28056**
- snmp | **28058**
- snmp-command | **28064**
- snmp-community | **28065**
- snmp-options | **28067**
- snmp (SNMP Scripts) | **28068**
- snmp-trap-thresholds | **28070**
- snmp-trap-thresholds (Next Gen Services) | **28072**
- snmp-value-match-msmic (Services NAT Options) | **28074**
- snoop-pseudowires | **28075**

- snoop-segment (TDF Domain Selection) | **28076**
- snoop-segments (RADIUS) | **28078**
- snoop-segments (TDF Gateway) | **28080**
- snr-margin | **28081**
- sockets | **28082**
- soft-gre | **28083**
- soft-limit | **28085**
- soft-limit-clear | **28087**
- soft-preemption (Protocols MPLS) | **28088**
- soft-preemption (Protocols RSVP) | **28089**
- software-upgrade | **28090**
- softwire-concentrator | **28092**
- softwire-name (Next Gen Services) | **28093**
- softwire-options | **28095**
- softwire-rules | **28097**
- softwire-types (Next Gen Services) | **28098**
- softwires | **28102**
- softwires (Next Gen Services) | **28104**
- softwires-rule-set (Service Set Next Gen Services) | **28105**
- solicit-router-advertisement-unicast | **28106**
- sonet | **28108**
- sophos-engine | **28109**
- source (GRE Encapsulation) | **28111**
- source (Interfaces) | **28112**
- source-active-advertisement | **28114**
- source-address (Services IPSec VPN) | **28115**
- source-address (Dynamic Tunnel) | **28117**
- source-address (IGMP and MLD Snooping) | **28118**
- source-address | **28120**
- source-address (Content Security) | **28121**
- source-address (AACL) | **28123**
- source-address (Access LDAP) | **28124**

- source-address (Access RADIUS) | **28125**
- source-address-excluded | **28127**
- source-address-filter | **28128**
- source-address-filter (Interfaces) | **28130**
- source-address-filter-list | **28131**
- source-address (Forwarding Options) | **28133**
- source-address-giaddr | **28135**
- source-address (IDS MS-DPC) | **28136**
- source-address (LRF Profile) | **28138**
- source-address-name (NAT Next Gen Services) | **28139**
- source-address-name (Security Destination NAT) | **28140**
- source-address-name (Security Source NAT) | **28141**
- source-address-name (Security Static NAT Rule Set) | **28142**
- source-address (NAT Next Gen Services) | **28144**
- source-address (Next Gen Services Service-Set Remote System Logging) | **28145**
- source-address (PCP) | **28146**
- source-address (RADIUS, System Logging, or TACACS+) | **28148**
- source-address-range | **28149**
- source-address-range (IDS MS-DPC) | **28150**
- source-address-range (PCP) | **28152**
- source-address-range (Services NAT) | **28153**
- source-address-range (Services Stateful Firewall) | **28155**
- source-address (Routing Options) | **28156**
- source-address (RPM) | **28157**
- source-address (Security Destination NAT) | **28160**
- source-address (Security IDP) | **28161**
- source-address (Security IDP Policy) | **28162**
- source-address (Security IDP Sensor Configuration) | **28163**
- source-address (Security Policies) | **28164**
- source-address (Security Source NAT) | **28166**
- source-address (Security Static NAT Rule Set) | **28167**
- source-address (Service Sets) | **28168**

- source-address (Services CoS) | **28169**
- source-address (Services NAT) | **28171**
- source-address (Services Stateful Firewall) | **28172**
- source-address (SNMP) | **28174**
- source-address (SRC Software) | **28175**
- source-address (Subscriber Secure Policy) | **28176**
- source-address (System Logging) | **28177**
- source-address (TWAMP) | **28179**
- source-addresses | **28181**
- source (Application Identification) | **28182**
- source-attachment-identifier (Protocols VPWS) | **28183**
- source-bmac | **28185**
- source (Bridge Domains) | **28187**
- source (Chassis Synchronization) | **28188**
- source-checking | **28190**
- source-class-usage | **28191**
- source-classes | **28193**
- source (Commit Scripts) | **28194**
- source-count (Dynamic MLD Interface) | **28196**
- source-count (Protocols IGMP) | **28197**
- source-count (Protocols MLD) | **28199**
- source (Distributed IGMP) | **28200**
- source (Dynamic IGMP Interface) | **28201**
- source (Dynamic MLD Interface) | **28203**
- source-end-user-profile | **28204**
- source (Event Scripts) | **28206**
- source-except | **28207**
- source-filtering | **28208**
- source-filtering (Interfaces) | **28210**
- source-gateway (Tunnel Profile) | **28211**
- source-hostname | **28213**
- source-id | **28214**

- source-identity | **28216**
- source-identity-log (Security) | **28218**
- source-increment (Dynamic MLD Interface) | **28220**
- source-increment (Protocols IGMP) | **28221**
- source-increment (Protocols MLD) | **28223**
- source-interface (RADIUS Client) | **28224**
- source-interface | **28225**
- source-interface (RADIUS Server) | **28227**
- source-interface (RADIUS Snoop Segment) | **28228**
- source-interface-set-at-login | **28229**
- source interfaces | **28231**
- source-ip | **28233**
- source-ip-address-list | **28235**
- source-ip-address (RADIUS Snoop Segment) | **28237**
- source-ip-based | **28238**
- source-ip-change (Forwarding Options) | **28239**
- source-ip (Flow Monitoring Logs for NAT) | **28240**
- source-ip (Security Alarms) | **28242**
- source-ipv4-address (RFC 2544 Benchmarking) | **28244**
- source (JET Scripts) | **28245**
- source-mac-address | **28247**
- source-mac-address-list | **28249**
- source-mode | **28250**
- source (Multicast VLAN Registration) | **28252**
- source-nat | **28254**
- source (Op Scripts) | **28255**
- source-packet-routing (Segment Routing Traffic Engineering) | **28257**
- source-packet-routing (Junos Telemetry Interface) | **28261**
- source-packet-routing (routing-options) | **28264**
- source-packet-routing (Protocols IS-IS) | **28266**
- source-packet-routing (Protocols OSPF) | **28270**
- source-packet-routing (Routing Instances Protocols EVPN) | **28273**

- source (PIM RPF Selection) | **28275**
- source-pool | **28276**
- source-port | **28277**
- source-port (firewall) | **28278**
- source-port (Applications) | **28280**
- source-port (Security Forwarding Options) | **28281**
- source-port (Security IDP) | **28286**
- source-port (Security Source NAT Rule Set) | **28288**
- source-port (Security Static NAT Rule Set) | **28289**
- source-port (Subscriber Secure Policy) | **28290**
- source-prefix (IDS) | **28292**
- source-prefix-ipv6 (IDS) | **28293**
- source-prefix-list (PCP) | **28295**
- source-prefix-list (Services AACL) | **28296**
- source-prefix-list (Services CoS) | **28297**
- source-prefix-list (Services IDS) | **28299**
- source-prefix-list (Services NAT) | **28300**
- source-prefix-list (Services Stateful Firewall) | **28301**
- source-prefix (Security Forwarding Options) | **28303**
- source-prefix (Services NAT) | **28304**
- source (Protocols IGMP) | **28305**
- source (Protocols MLD) | **28307**
- source (Protocols MSDP) | **28308**
- source-route-option (IDS Screen Next Gen Services) | **28310**
- source-routing | **28311**
- source (Routing Instances) | **28312**
- source (Routing Instances Provider Tunnel Selective) | **28314**
- source-routing-path | **28316**
- source-routing-path-template | **28319**
- source (Security Source NAT) | **28322**
- source (SNMP Scripts) | **28326**
- source (Source-Specific Multicast) | **28327**

- source-threshold | **28329**
- source-tree (MBGP MVPN) | **28330**
- source-udp-port (RFC 2544 Benchmarking) | **28332**
- source-based-fpolicy-statementorwarding | **28333**
- source-lookup | **28338**
- source-routing | **28339**
- spam-action | **28341**
- speed (24-port and 12-port 10 Gigabit Ethernet PIC) | **28342**
- speed | **28344**
- speed (Chassis Cluster) | **28347**
- speed (Ethernet) | **28349**
- speed (Fibre Channel Interfaces) | **28370**
- speed (Gigabit Ethernet interface) | **28372**
- speed (Gigabit Ethernet interface - MX Series) | **28373**
- speed (Interfaces) | **28376**
- speed (MX Series DPC) | **28377**
- spf-options (Protocols IS-IS) | **28379**
- spf-options (Protocols OSPF) | **28381**
- spi | **28384**
- spi (Junos-FIPS Software) | **28386**
- spi (Junos OS) | **28387**
- spid1 | **28388**
- spid2 | **28389**
- splitting-merging | **28391**
- spring-te | **28393**
- spring-te (Dynamic Tunnels) | **28395**
- spring-te-lsp-next-hop | **28397**
- spring-traffic-engineering-edit-protocols | **28398**
- spt-only | **28400**
- spt-threshold | **28401**
- sql | **28402**
- squelch | **28404**

- srgb (Protocols BGP) | **28407**
- srgb (Protocols IS-IS) | **28408**
- srgb-label-range | **28410**
- srlg | **28411**
- srlg-cost | **28413**
- srlg-value | **28414**
- srv6 (Routing Options) | **28415**
- srv6 (Protocols BGP) | **28417**
- srv6-sid | **28419**
- ssb | **28421**
- ssh | **28422**
- ssh-known-hosts | **28426**
- ssh (NETCONF) | **28428**
- ssh (System Services) | **28430**
- ssl | **28439**
- ssl-encryption | **28442**
- ssl-inspection | **28443**
- ssl-proxy (Application Services) | **28446**
- ssl (Services) | **28447**
- ssl-termination-profile | **28450**
- ssm-groups | **28452**
- ssm-map (Dynamic IGMP Interface) | **28453**
- ssm-map (Dynamic MLD Interface) | **28455**
- ssm-map-policy (Dynamic IGMP Interface) | **28458**
- ssm-map-policy (Dynamic MLD Interface) | **28459**
- ssm-map-policy (IGMP) | **28461**
- ssm-map-policy (MLD) | **28462**
- ssm-map (Protocols IGMP) | **28464**
- ssm-map (Protocols IGMP AMT) | **28465**
- ssm-map (Protocols MLD) | **28466**
- ssm-map (Routing Options Multicast) | **28468**
- stacked-interface-set (Dynamic Profiles) | **28469**

- stacked-vlan-ranges | **28471**
- stacked-vlan-ranges (RADIUS Options) | **28473**
- stacked-vlan-tagging | **28475**
- stale-labels-holddown-period | **28476**
- stale-routes-time | **28478**
- stale-routes-time (Fabric Control) | **28479**
- stale-time (Long-Lived Graceful Restart for BGP Restarter) | **28481**
- stamp (Sampling) | **28484**
- standby | **28485**
- standby | **28486**
- standby-path-creation-delay | **28488**
- standby (Protocols Layer 2 Circuit) | **28489**
- standby (Routing Policy Condition) | **28491**
- start-date | **28492**
- start-log | **28493**
- start-measurement | **28494**
- start (Scripts) | **28496**
- start shell | **28498**
- start-time | **28500**
- start-time (Accounting) | **28502**
- start-time (Schedulers) | **28503**
- start-time (Security IDP) | **28504**
- starts-with (DHCP Relay Agent Option) | **28506**
- starts-with (Event Policy) | **28510**
- startup-alarm | **28511**
- startup-silent-period | **28513**
- stateful (Router Advertisement Guard) | **28514**
- stateful (Precision Time Protocol) | **28516**
- stateful-firewall-rule-set (Next Gen Services) | **28518**
- stateful-firewall-rule-sets (Service Set Next Gen Services) | **28519**
- stateful-firewall-rules | **28520**
- stateful-firewall-rules (Service Set Next Gen Services) | **28521**

- stateful-nat64 | **28523**
- stateless | **28524**
- static (802.1X) | **28526**
- static-binding | **28528**
- static (Bridge Domains) | **28529**
- static (Distributed IGMP) | **28531**
- static (Dynamic IGMP Interface) | **28532**
- static (Dynamic MLD Interface) | **28534**
- static-host-mapping | **28535**
- static (IGMP Snooping) | **28537**
- static-interface | **28539**
- static-ip | **28540**
- static-ip (MX Series) | **28542**
- static-ipv6 | **28544**
- static-label-switched-path | **28545**
- static-lsp | **28547**
- static-mac | **28549**
- static-nat | **28551**
- static (Origin Validation for BGP) | **28553**
- static-policy-control | **28555**
- static (Protocols IGMP) | **28557**
- static (Protocols Layer 2 Circuit) | **28559**
- static (Protocols MLD) | **28560**
- static (Protocols PIM) | **28562**
- static (Protocols VPLS) | **28564**
- static-remote-vtep-list | **28566**
- static (Routing Options) | **28570**
- static (Security NAT) | **28577**
- static-subscriber-application | **28579**
- static-subscribers | **28581**
- static-subscribers (Dynamic Service Provisioning) | **28582**
- static-tei-val | **28584**

- static-umh (MBGP MVPN) | **28585**
- static-mac-ip | **28587**
- static-multihoming-peer | **28589**
- statistics | **28591**
- statistics (Access Profile) | **28592**
- statistics (Protocols MPLS) | **28594**
- statistics-service | **28596**
- statistics (Service Accounting) | **28598**
- statistics (Services) | **28599**
- stats-aggregation-interval (programmable-rpd) | **28601**
- stats-sampling-interval (programmable-rpd) | **28602**
- status | **28603**
- status-control | **28604**
- steering | **28605**
- step-percent (RFC 2544 Benchmarking) | **28607**
- stickydr | **28609**
- stop-date | **28610**
- stop-time | **28612**
- storage | **28613**
- store | **28615**
- storm-control (Video Monitoring) | **28619**
- storm-control | **28620**
- storm-control (switch) | **28622**
- storm-control-profiles | **28624**
- stp | **28625**
- stream (Logical Systems and Tenant Systems) | **28628**
- stream (Next Gen Services Service-Set Remote System Logging) | **28631**
- stream-option (IDS Screen Next Gen Services) | **28633**
- stream-protection (Multicast-Only Fast Reroute) | **28634**
- stream (Security Log) | **28635**
- streaming-server (Junos Telemetry Interface) | **28640**
- streaming-servers | **28642**

- strict-asla-based-flex-algorithm | **28645**
- strict (DHCP Local Server) | **28646**
- strict-dual-isis (Protocols IS-IS) | **28648**
- strict-priority-scheduler | **28650**
- strict-source-route-option (IDS Screen Next Gen Services) | **28652**
- strict-syn-check | **28653**
- strict-targeted-hellos | **28654**
- string | **28655**
- strip-domain (Domain Map) | **28657**
- strip-nexthop | **28658**
- strip-user-name (Access Profile) | **28659**
- strip-username (Domain Map) | **28661**
- strip-as-path | **28662**
- structured-data | **28664**
- stub | **28665**
- stub-network | **28667**
- sub-domain | **28668**
- sub-prefix-length | **28674**
- subnet-address (Easy EVPN LAG Configuration) | **28675**
- subscriber (Access Profile) | **28680**
- subscriber-address | **28683**
- subscriber-aware-services | **28684**
- subscriber-awareness (Service Set Options) | **28685**
- subscriber-context (AAA Options) | **28687**
- subscriber (DDoS Flow Detection) | **28688**
- subscriber-exclude-prefix | **28691**
- subscriber-extension | **28692**
- subscriber-leave-timer | **28694**
- subscriber-management-helper | **28695**
- subscriber-management | **28697**
- subscriber-packet-idle-timeout | **28699**
- subscriber-profile | **28700**

- subscriber-type (TDF Domain) | **28702**
- subscriber-info | **28704**
- subscriber-pool-utilization-alarm | **28705**
- subscribers-limit (Resource Monitor) | **28707**
- subscription | **28709**
- subscription-id | **28711**
- subscription-id-data-include (PCRF Partition) | **28713**
- subscription-id-options | **28715**
- subscription-id-type (Class Attribute) | **28717**
- subscription-id-type (PCRF Partition) | **28719**
- success | **28720**
- sunrpc | **28722**
- supplementary-blackout-timer (Host Fast Reroute) | **28724**
- supplicant | **28726**
- suppress-host-routes-from-dci-to-dc | **28728**
- suppress-interface-leaf | **28730**
- suppression | **28731**
- surf-control-integrated | **28733**
- suspend-for | **28735**
- suspend-on-request | **28737**
- swap (protocols mpls) | **28738**
- swap | **28740**
- swap-by-poppush | **28741**
- swap (Dynamic VLANs) | **28742**
- swap-push | **28744**
- swap-swap | **28745**
- switch-away-lsps | **28746**
- switch-options (Interfaces) | **28748**
- switch-options | **28749**
- switch-options (VLANs) | **28751**
- switch-port | **28753**
- switch (Protocols OpenFlow) | **28755**

- switch-type | **28757**
- switching-type | **28758**
- switchover-delay | **28759**
- switchover-mode | **28761**
- switchover-on-routing-crash | **28763**
- sxl-retry | **28764**
- sxl-timeout | **28765**
- symbol-period | **28767**
- symmetric-hash | **28768**
- syn-ack-ack-proxy | **28770**
- syn-ack-ack-proxy (IDS Screen Next Gen Services) | **28771**
- syn-check-required | **28773**
- syn-cookie (IDS MS-DPC) | **28774**
- syn-fin | **28776**
- syn-fin (IDS Screen Next Gen Services) | **28777**
- syn-flood | **28778**
- syn-flood-protection-mode | **28779**
- syn-frag | **28781**
- syn-frag (IDS Screen Next Gen Services) | **28782**
- sync-active-path-bandwidth | **28783**
- sync-reset | **28785**
- synchronization (ACX Series) | **28787**
- synchronization (M Series and T Series) | **28792**
- synchronization (MX Series) | **28795**
- synchronization (PTX Series) | **28802**
- synchronize | **28804**
- synchronize (Scripts) | **28806**
- synchronous-ethernet-mapping | **28808**
- profile-type (PTP) | **28809**
- syntonized-e2e-transparent | **28813**
- syslog | **28815**
- syslog (Services IPsec VPN) | **28816**

- syslog (Chassis) | **28817**
- syslog (IDS MS-DPC) | **28819**
- syslog (Interfaces) | **28820**
- syslog-local-system-timestamp | **28821**
- syslog (Next Gen Services Service-Set System Logging) | **28823**
- syslog (OAM Action) | **28824**
- syslog (QFabric System) | **28825**
- syslog (Services CoS) | **28827**
- syslog (Services L2TP) | **28829**
- syslog (Services NAT) | **28830**
- syslog (Services Service Set) | **28831**
- syslog (Services Stateful Firewall) | **28833**
- syslog-subtag | **28834**
- syslog (System) | **28836**
- system (Analytics Resource) | **28840**
- system-defaults | **28842**
- system (Export Profiles) | **28844**
- system-id (Interfaces) | **28845**
- system-id | **28847**
- system-id (Junos Fusion) | **28848**
- system-id (Junos Fusion Satellite Device Cluster) | **28850**
- system login (Junos Node Slicing) | **28852**
- system (Memory) | **28853**
- system-priority (chassis) | **28855**
- system-priority (interfaces) | **28857**
- system-services (Security Zones Host Inbound Traffic) | **28859**
- system-services (Security Zones Interfaces) | **28862**

s-nssai

IN THIS SECTION

- Syntax | 16259
- Hierarchy Level | 16259
- Description | 16259
- Options | 16260
- Required Privilege Level | 16260
- Release Information | 16260

Syntax

```
s-nssai name {
 sd sd;
 sst(embb | miot | urllc | v2x);
}
```

Hierarchy Level

```
[edit services agf tracking-area name plmn]
```

Description

Configure the single network slice selection assistance information (S-NSSAI). The S-NSSAI is used to uniquely identify a network slice. It consists of two components. For a fixed network residential gateway (FN-RG), the Access Gateway Function (AGF) sends a proxy S-NSSAI message to the Access and Mobility Management Function (AMF) during the initial setup.

- SST—Slice service type
- SD—Selection differentiator

Options

name Single network slice selection assistance information identifier

- **Range:**0 through 128

sd Selection differentiator

- **Default:** 0
- **Range:** 0 through 16777215

sst Slice service type

- embb—Enhanced mobile broadband slice
- miot—Massive Internet of Things (IoT)
- urllc—Ultra-reliable, low-latency communications
- v2x—V2X services

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3R1.

s6a

IN THIS SECTION

- Syntax | **16261**
- Hierarchy Level | **16261**
- Description | **16261**

- Options | 16261
- Required Privilege Level | 16262
- Release Information | 16262

Syntax

```
s6a {
    partition name {
        destination-host destination-host;
        destination-realm destination-realm;
        diameter-instance diameter-instance;
        max-outstanding-requests max-outstanding-requests;
        response-timeout seconds;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

S6a is a Diameter-based authentication application used by the Mobility Management Entity (MME) nodes to retrieve authentication information from Home Subscriber Server (HSS). Specify the destination and transmission parameters for the S6a protocol.

Options

destination-host Name of the host where the S6A server application resides; generally, this value is not set unless needed.

destination-realm Name of the realm where the S6a server application resides.

diameter-instance Specifies that the master Diameter instance is used. `master` is the only supported value.

max-outstanding-requests

Configure the maximum number of outstanding requests for the S6a application.

- **Default:** 40
- **Range:** 2 through 1024

response-timeout

Configure the amount of time (in seconds) that the MME waits to receive a response from the HSS.

- **Default:** 15
- **Range:** 1 through 30

Required Privilege Level

access

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Example: Configure S6a Application

authentication-order | **2633**

partition (s6a) | **12190**

Diameter Base Protocol Overview

sa-hold-time (Protocols MSDP)

IN THIS SECTION

- Syntax | **16263**
- Hierarchy Level | **16263**
- Description | **16263**
- Options | **16264**

Required Privilege Level | 16264
Release Information | 16264

Syntax

```
sa-hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name routing-instances instance-name protocols msdp],
[edit logical-systems logical-system-name routing-instances instance-name protocols msdp group
group-name peer address],
[edit logical-systems logical-system-name routing-instances instance-name protocols msdp peer
address],
[edit protocols msdp],
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit routing-instances instance-name protocols msdp],
[edit routing-instances instance-name protocols msdp group group-name peer address]
[edit routing-instances instance-name protocols msdp peer address]
```

Description

Specify the source address (SA) message hold time to use when maintaining a connection with the MSDP peer. Each entry in an SA cache has an associated hold time. The hold timer is started when an SA message is received by an MSDP peer. The timer is reset when another SA message is received before the timer expires. If another SA message is not received during the SA message hold-time period, the SA message is removed from the cache.

You might want to change the SA message hold time for consistency in a multi-vendor environment.

Options

seconds—Source address message hold time.

- **Range:** 75 through 300 seconds
- **Default:** 75 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Examples: Configuring MSDP

hold-time (Protocols MSDP) | **8112**

keep-alive (Protocols MSDP) | **9457**

sa-multicast (100-Gigabit Ethernet)

IN THIS SECTION

- Syntax | **16265**
- Hierarchy Level | **16265**
- Description | **16265**
- Required Privilege Level | **16265**
- Release Information | **16265**

Syntax

```
sa-multicast;
```

Hierarchy Level

```
[edit chassis fpc slot pic slot forwarding-mode]
```

Description

Configure the 100-Gigabit Ethernet PIC or MIC to interoperate with other Juniper Networks 100-Gigabit Ethernet PICs.

> **NOTE**: The default packet steering mode for PD-1CE-CFP-FPC4 is SA multicast bit mode. No SA multicast configuration is required to enable this mode.

sa-multicast supports interoperability between the following PICs and MICs:

- 100-Gigabit Ethernet Type 5 PIC with CFP (PF-1CGE-CFP) and the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) .
- 100-Gigabit Ethernet MICs and the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

> *Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP*
>
> *Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4*

Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode

Interoperability Between MPC4E (MPC4E-3D-2CGE-8XGE) and 100-Gigabit Ethernet PICs on Type 4 FPC

Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode

Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP

Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4

forwarding-mode (100-Gigabit Ethernet) | **7459**

sa-multicast (PTX Series Packet Transport Routers) | **16266**

vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) | **30855**

Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP

sa-multicast (PTX Series Packet Transport Routers)

IN THIS SECTION

- Syntax | **16266**
- Hierarchy Level | **16266**
- Description | **16267**
- Required Privilege Level | **16267**
- Release Information | **16267**

Syntax

```
sa-multicast;
```

Hierarchy Level

```
[edit chassis fpc slot pic slot port port-number forwarding-mode]
```

Description

Configure source address (SA) multicast bit mode on the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP to enable interoperability with 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4.

> **NOTE**: When SA multicast bit steering mode is configured on a PTX Series Packet Transport Router 100-Gigabit Ethernet port, VLANs are not supported for that port.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X48R4.

RELATED DOCUMENTATION

Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP

Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4

sabit

IN THIS SECTION

- Syntax | 16268
- Hierarchy Level | 16268
- Description | 16268
- Options | 16268
- Required Privilege Level | 16268
- Release Information | 16268

Syntax

```
sabit bit;
```

Hierarchy Level

```
[edit chassis synchronization interfaces external e1-options]
[edit chassis synchronization interfaces (external-0/0 | external-1/0) e1-options]
[edit chassis synchronization interfaces bits e1-options]
```

Description

Configure the SA bit for exchanging the SSM quality on the E1 interface.

Options

bit SA bit value.

- **Range:** 4 through 8.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | **28795**

saegw

IN THIS SECTION

- Syntax | 16269
- Hierarchy Level | 16269
- Description | 16269
- Options | 16270
- Required Privilege Level | 16270
- Release Information | 16270

Syntax

```
saegw name {
    access-network-peers {
    ...}
    control-plane-peers {
    ...}
    core-network-peers {
    ...}
    system {
    ...}
}
```

Hierarchy Level

```
[edit services mobile-edge gateways]
```

Description

Define the gateway name.

NOTE: Only a single instance of the gateway can be defined on the MX Series router.

Options

name Gateway name

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.4R1.

`core-network-peers` option introduced in Junos OS Release 20.4R1.

saegw access-network-peers

IN THIS SECTION

- Syntax | 16271
- Hierarchy Level | 16271
- Description | 16271
- Options | 16271
- Required Privilege Level | 16272
- Release Information | 16272

Syntax

```
access-network-peers {
    local-address [ local-address ... ];
    n3-requests n3-requests;
    peer-groups name {
        peer {
            address [ address ... ];
            hostname hostname;
        }
        routing-instance routing-instance;
    }
    routing-instance routing-instance;
    t3-response seconds;
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw]
```

Description

Use this section to define the connection to and characteristics of data plane peers.

Options

local-address

Required. IPv4 address(es) of the local end of the GTP-U connection.

> **NOTE**: There is no limit to the number of data plane peers that can connect to the SAEGW-U.

n3-requests

Number of retries of peer management request messages upon `t3-response` timeout.

- **Range:** 1 through 5
- **Default:** 3

routing-instance

Local routing instance of the GTP-U connection.

> **NOTE**: If at least one peer group is used, eNodeBs matching only the peer group address/prefix are accepted by the SAEGW-U/SGW-U during session establishment. The eNodeB is bound to the routing instance within the "peer-groups" on page 12361 stanza, if available. Otherwise, the SAEGW-U/SGW-U uses the routing instance provided here.

t3-response Waiting time before retrying peer management request upon response timeout (seconds).

- **Range:** 1 through 5
- **Default:** 5

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

saegw control-plane-peers | 16272
saegw-core-network-peers | 16276

saegw control-plane-peers

IN THIS SECTION

- Syntax | 16273
- Hierarchy Level | 16273
- Description | 16274
- Options | 16274

Required Privilege Level | 16275
Release Information | 16275

Syntax

```
control-plane-peers {
    apn-services {
        apns name {
            mobile-interface mobile-interface;
            source-validation {
        }
    }
    heartbeat-interval seconds;
    local-address [ local-address ... ];
    n3-requests n3-requests;
    path-management (disable | enable);
    peer-groups name {
        heartbeat-interval seconds;
        initiate-association;
        n3-requests n3-requests;
        path-management (disable | enable);
        peer {
            address [ address ... ];
            hostname hostname;
        }
        routing-instance routing-instance;
        t3-response seconds;
    }
    response-cache-timeout seconds;
    routing-instance routing-instance;
    t3-response seconds;
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw]
```

Description

Use this section to define the connection to and characteristics of control plane peers (SAEGW-C, SGW-C, or PGW-C).

Options

heartbeat-interval

Time between two successive heartbeat requests (seconds).

- **Range:** 60 through 255
- **Default:** 60

local-address

Required. IPv6 or IPv4 or both addresses of the local end of the PFCP connection.

> **NOTE**: If you are connecting to multiple control plane peers, define the local address under the `control-plane-peers` stanza for each control plane peer rather than define a single loopback address.

n3-requests

Maximim number of retries of PFCP request messages upon t3-response timeout.

- **Range:** 1 through 5
- **Default:** 3

path-management

Enable/disable origination of heartbeat message requests to control peers.

- **Default:** Disabled

response-cache-timeout

Configure the timeout for the PFCP response cache (seconds).

- **Range:** 0 through 255
- **Default:** 0

routing-instance

Local routing instance of the PFCP. This is used to determine in which routing instance the responses are to be sent for incoming PFCP messages. If a routing instance is not configured:

- Only PFCP messages coming from control plane peers over the default routing instance are handled.
- PFCP responses are only sent in the default routing instance.

- PFCP messages coming over any other routing instance are dropped even if the destination address matches the defined `local-address`.

If a single routing instance is configured:

- Only PFCP messages coming from control plane peers over the configured routing instance are handled.
- PFCP responses are only sent through this routing instance.
- PFCP messages coming over any other routing instance are dropped.

t3-response

Waiting time of gateway before retrying a PFCP signaling-request upon response timeout (seconds).

- **Range:** 1 through 5
- **Default:** 5

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

saegw access-network-peers | **16270**

saegw-core-network-peers | **16276**

saegw-core-network-peers

IN THIS SECTION

- Syntax | 16276
- Hierarchy Level | 16276
- Description | 16277
- Options | 16277
- Required Privilege Level | 16277
- Release Information | 16277

Syntax

```
core-network-peers {
    echo-interval seconds;
    local-address [ local-address ... ];
    n3-requests n3-requests;
    path-management (disable | enable);
    peer-groups name {
        peer {
            address [ address ... ];
            hostname hostname;
        }
        routing-instance routing-instance;
    }
    routing-instance routing-instance;
    t3-response seconds;
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw]
```

Description

Define core-facing GTP-U peers (SGW-U to PGW-U and vice versa).

Options

echo-interval Time between origination of two successive echo requests (seconds). You must enable `path-management` for echo requests to occur.

- **Range:** 60 through 255
- **Default:** 60

local-address Required. IPv4 address(es) of the local end of the GTP-U connection.

n3-requests Number of retries of peer management request messages upon `t3-response` timeout.

- **Range:** 1 through 5
- **Default:** 3

path-management Enable/disable origination of echo requests to core peers.

- **Default:** Disabled

routing-instance Local routing instance of the GTP-U connection.

t3-response Waiting time before retrying peer management request upon response timeout (seconds).

- **Range:** 1 through 5
- **Default:** 5

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

saegw access-network-peers | 16270

saegw control-plane-peers | 16272

saegw system

IN THIS SECTION

- Syntax | 16278
- Hierarchy Level | 16279
- Description | 16279
- Options | 16279
- Required Privilege Level | 16279
- Release Information | 16279

Syntax

```
system {
    anchor-pfes {
        interface interface-name {
            colocated-4g-pgw;
            colocated-4g-sgw;
        }
    }
    class-of-service {
    ... }
    destination {
    ... }
    service-mode (mobile-edge SAEGW) {
    ... }
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw]
```

Description

Define the anchor PFEs for the SAEGW-U. The anchor PFEs are the line cards where the GTP-U data packets from `access-network-peers` are decapsulated and sent towards the SGi interface to the core network. Similarly, packets from the SGi interface are encapsulated in GTP-U by the anchor PFEs and sent towards the RAN.

Options

anchor-pfes Define the anchor PFEs by providing their interface names.

- **Syntax:** interface *interface-name*

colocated-4g-pgw (Optional) Attach the PGW-U gateway type to the anchor PFE to optimize throughput of 4G traffic.

colocated-4g-sgw (Optional) Attach the SGW-U gateway type to the anchor PFE to optimize throughput of 4G traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.4R1.

`colocated-4g-pgw` and `colocated-4g-sgw` options introduced in Junos OS Release 21.3R1.

`class-of-service` and `destination` options introduced in Junos OS Release 22.2R1.

`service-mode maintenance` options introduced in Junos OS Release 23.1R1.

sak-rekey-interval

IN THIS SECTION

- Syntax | 16280
- Hierarchy Level | 16280
- Description | 16280
- Default | 16281
- Options | 16281
- Required Privilege Level | 16281
- Release Information | 16281

Syntax

```
sak-rekey-interval seconds;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name mka]
```

Description

Configure a timer-based refresh of the secure association key (SAK) on a MACsec-secured link. This ensures that the SAK is frequently updated, making it less vulnerable to attack.

In static CAK mode, the SAK is generated by the key server and is periodically refreshed. The refresh interval is based on packet counter movement by default. Depending on the amount of traffic and the speed of the interface, it might take a long time for the new SAK to be generated. This can provide enough time for a successful attack on the key. You can enhance security of the SAK by configuring a shorter timer-based refresh interval.

When the MACsec session is live with a primary, preceding, or fallback PSK, or with a key from a key-chain, the SAK refresh will occur based on the configured interval independent of key type in a hitless

way. No traffic drop will occur at the time of the SAK rollover. Refresh of the PSK will occur periodically after every configured SAK refresh interval.

When an XPN cipher suite is configured, the refresh interval configured in the key server takes precedence over the refresh interval configured in the non-key server, even if the interval configured in the non-key server is lower.

Default

The SAK refresh interval is not enabled by default.

Options

seconds The length of the SAK refresh interval in seconds.

- **Range:** 60 through 86,400 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Configuring Advanced MACsec Features

saml

IN THIS SECTION

Syntax | 16282

Hierarchy Level | 16283

- Description | 16283
- Options | 16283
- Required Privilege Level | 16283
- Release Information | 16283

Syntax

```
saml {
    identity-provider identity-provider {
        attribute-mapping {
            attribute-name assertion-attribute <mandatory | optional>;
            address-pool assertion-attribute <mandatory | optional>;
            firstname assertion-attribute <mandatory | optional>;
            group assertion-attribute <mandatory | optional>;
            lastname assertion-attribute <mandatory | optional>;
            username assertion-attribute <mandatory | optional>;
        }
        metadata-id metadata-id;
        settings {
        entity-id entity-id;
        idp-certificate idp-certificate;
        single-logout-url single-logout-url;
        single-signon-url single-signon-url;
        }
    }
    service-provider service-provider {
        assertion-waittime assertion-waittime;
        encrypted-assertions;
        entity-id entity-id;
        hash-algorithm <sha-256 | sha-384 | sha-512>;
        sign-auth-request;
        sp-certificate sp-certificate;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Describe Security Assertion Markup Language (SAML) authentication parameters.

Options

identity-provider *identity-provider* Specify the identity provider (IdP) configuration details.

service-provider *service-provider* Specify the service provider configuration details.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

show network-access aaa saml assertion-cache | 22560

clear network-access aaa saml assertion-cache | 3757

saml-options

IN THIS SECTION

- Syntax | 16284
- Hierarchy Level | 16284
- Description | 16284
- Options | 16284
- Required Privilege Level | 16285
- Release Information | 16285

Syntax

```
saml-options {

    no-assertion-cache;
}
```

Hierarchy Level

```
[edit access]
```

Description

Specify the option to avoid caching Security Assertion Markup Language (SAML) assertions from the identity provider (IdP).

Options

no-assertion-cache	Do not cache assertions from IdP.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

saml | 16281

sample

IN THIS SECTION

- Syntax | 16285
- Hierarchy Level | 16286
- Description | 16286
- Options (PTX Series for Junos OS Evolved) | 16286
- Required Privilege Level | 16286
- Release Information | 16287

Syntax

```
sample
 {
   aggregate
     {
       range range;
```

```
        }
    }
```

Hierarchy Level

```
[edit system ddos-protection protocols]
```

Description

Configure aggregate for all sampling traffic.

Options (PTX Series for Junos OS Evolved)

Table 137: Protocol Groups Supported by Control Plane DDoS Protection on PTX Series Routers for Junos OS Evolved

Protocol Group	Description
`aggregate`	Configure aggregate for all Sampling traffic
`range` (PTX10008, PTX10001-36MR, PTX10001, and PTX10004 routers only)	range <0 to 150000>
`range` (PTX10003-80C and PTX10003-160C routers only)	range <0 to 50000>

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

sample-once

IN THIS SECTION

- Syntax | 16287
- Hierarchy Level | 16287
- Description | 16287
- Required Privilege Level | 16288
- Release Information | 16288

Syntax

```
sample-once;
```

Hierarchy Level

```
[edit forwarding-options sampling]
```

Description

Explicitly sample a packet for active monitoring only once. When you configure the `sample-once` statement, the ingress interface sets a flag for any packets that it selects for sampling, and the egress interface does not sample any flagged packets. Setting this option avoids duplication of packets in cases where sampling is enabled at both the ingress and egress interfaces and simplifies analysis of the sampled traffic.

NOTE: The sampling works independently on the ingress and egress interfaces. If the sample `rate` is 1, the ingress interface samples and flags all packets arriving at the

interface. However, if the sample `rate` is higher than 1, then the ingress interface may sample and flag only a subset of packets, as determined by the configured sampling parameters. The ingress and egress interfaces may contain different sets of packets in their sampling windows. As a result, if the ingress interface does not sample (and flag) a packet in a given sampling window, the egress interface may still sample that same packet as part of its sampling window.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

sample-rate

IN THIS SECTION

- Syntax | 16289
- Hierarchy Level | 16289
- Description | 16289
- Default | 16289
- Options | 16289
- Required Privilege Level | 16289
- Release Information | 16290

Syntax

```
sample-rate {
    egress number;
    ingress number;
}
```

Hierarchy Level

```
[edit protocols sflow],
[edit protocols sflow interfaces interface-name]
```

Description

Specify the number of egress or ingress packets out of which one packet is sampled. If no interface sampling rates are configured, the global sampling rates take effect. If neither is configured, by default both ingress and egress packet sampling are disabled.

Default

By default, both egress and ingress sampling rates are disabled.

Options

`egress` *number*—Value for egress sampling rate.

- **Range:** 1 through 1,073,741,823

`ingress` *number*—Value for ingress sampling rate.

- **Range:** 1 through 1,073,741,823

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

The option *number* (which immediately followed `sample-rate`) is no longer available and options `egress` *`number`* and `ingress` *`number`* added in Junos OS Release 10.4 for EX Series switches.

RELATED DOCUMENTATION

> *Configuring sFlow Technology for Network Monitoring (CLI Procedure)*
>
> *Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches*

sample-type

IN THIS SECTION

- Syntax | **16290**
- Hierarchy Level | **16290**
- Description | **16291**
- Options | **16291**
- Required Privilege Level | **16291**
- Release Information | **16291**

Syntax

```
sample-type (absolute-value | delta-value);
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Configure the method of sampling the selected variable (monitored object). When you configure an SNMP RMON alarm, you can specify the sample type.

Options

`absolute-value`—Actual value of the selected variable is used when comparing against the thresholds.

`delta-value`—Difference between samples of the selected variable is used when comparing against the thresholds.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Remote Nework Monitoring (RMON)

sampling

IN THIS SECTION

- Syntax | 16292
- Hierarchy Level | 16292
- Description | 16292
- Options | 16292
- Required Privilege Level | 16292
- Release Information | 16293

Syntax

```
sampling {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Perform packet sampling based on particular input interfaces and various fields in the packet header.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the traffic sampling control process.
- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this in the configuration.

system-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

sampling (Forwarding Options)

IN THIS SECTION

- Syntax | 16293
- Hierarchy Level | 16296
- Description | 16296
- Required Privilege Level | 16296
- Release Information | 16296

Syntax

```
sampling {
    disable;
    family (inet | inet6 | mpls | vpls) {
        disable;
        output {
            aggregate-export-interval seconds;
            file {
                disable;
                filename filename;
                files number;
                size bytes;
                (stamp | no-stamp);
                (world-readable | no-world-readable);
            }
```

```
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            flow-server hostname {
                aggregation {
                    autonomous-system;
                    destination-prefix;
                    protocol-port;
                    source-destination-prefix {
                        caida-compliant;
                    }
                    source-prefix;
                }
                autonomous-system-type (origin | peer);
                (local-dump | no-local-dump);
                port port-number;
                source-address address;
                version format;
                version9 {
                    template template-name;
                }
            }
            interface interface-name {
                engine-id number;
                engine-type number;
                source-address address;
            }
        }
    }
    input {
        max-packets-per-second number;
        maximum-packet-length bytes;
        rate number;
        run-length number;
    }
    instance instance-name {
        disable;
        family (bridge | inet | inet6 | mpls | vpls) {
            disable;
            output {
                aggregate-export-interval seconds;
                extension-service service-name;
                flow-server hostname {
                    aggregation {
```

```
                    autonomous-system;
                    destination-prefix;
                    protocol-port;
                    source-destination-prefix {
                        caida-compliant;
                    }
                    source-prefix;
                }
                autonomous-system-type (origin | peer);
                dscp dscp-value;
                forwarding-class class-name;
                (local-dump | no-local-dump);
                port port-number;
                source-address address;
                version format;
                version9 {
                    template template-name;
                }
                version-ipfix {
                    template template-name;
                }
            }
            inline-jflow {
                source-address address;
                flow-export-rate rate;
            }
            interface interface-name {
                engine-id number;
                engine-type number;
                source-address address;
            }
        }
    }
    input {
        max-packets-per-second number;
        maximum-packet-length bytes;
        rate number;
        run-length number;
    }
}
pre-rewrite-tos;
sample-once;
traceoptions {
```

```
        no-remote-trace;
        file filename <files number> <size bytes> <match expression> <world-readable | no-world-
readable>;
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure traffic sampling.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for IPv6 addresses for the `flow-server` *address* and `source-address` *address* statements within sampling instances introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

Applying Forwarding Table Filters

Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format

Directing Traffic Sampling Output to a Server Running the cflowd Application

Configuring Port Mirroring

Tracing Traffic-Sampling Operations

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

sampling-instance

IN THIS SECTION

- Syntax | 16297
- Hierarchy Level | 16297
- Description | 16297
- Required Privilege Level | 16297
- Release Information | 16298

Syntax

```
sampling-instance instance-name;
```

Hierarchy Level

```
[edit chassis fpc slot-number],
[edit chassis lcc number fpc routing-matrix-slot-number],
[edit chassis member member-number fpc slot slot-number]
```

Description

Associate a defined sampling instance with a specific FPC, MPC, or DPC for active sampling instances configured at the `[edit forwarding-options sampling]` hierarchy level.

For M120 routers with FEB, this statement must also be configured under `[edit chassis feb slot-number]`, in addition to the `[edit forwarding-options sampling]` hierarchy level.

In a two-member MX Series Virtual Chassis, the primary router (member 0) uses FPC slot numbers 0 through 11 with no offset; the backup router (member 1) uses FPC slot numbers 12 through 23, with an offset of 12.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support at the `[edit chassis member member-number fpc slot slot-number]` hierarchy level introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

> *Associating Sampling Instances for Active Flow Monitoring with a Specific FPC, MPC, or DPC*
>
> Inline Flow Monitoring for Virtual Chassis Overview

sampling (Interfaces)

IN THIS SECTION

- Syntax | **16298**
- Hierarchy Level | **16299**
- Description | **16299**
- Options | **16299**
- Required Privilege Level | **16299**
- Release Information | **16299**

Syntax

```
sampling direction;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet]
```

Description

Configure the direction of traffic to be sampled.

Options

direction can be one of the following:

`input`—Configure at least one expected ingress point.

`output`—Configure at least one expected egress point.

`input output`—On a single interface, configure at least one expected ingress point and one expect egress point.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Services Interfaces Library for Routing Devices

Configuring Flow Monitoring

sampling (Protocols MPLS)

IN THIS SECTION

- Syntax | 16300
- Hierarchy Level | 16300
- Description | 16300
- Options | 16301
- Required Privilege Level | 16301
- Release Information | 16301

Syntax

```
sampling {
    cut-off-threshold percentile;
    use-average-aggregate;
    use-percentile percentile;
}
```

Hierarchy Level

```
[edit protocols mpls container-label-switched-path lsp-name splitting-merging]
```

Description

Configure traffic sampling. By default, sampling mode is set to `Max Aggregate` which means the system will compare each new sample with the previous sample. If the new sample is higher than the old sample, then the newer sample is considered `Sampled Aggregate bandwidth`

Max Aggregate Example

If normalization happens every 20min (T0, T20, T40..) then if at time T0 the traffic rate is 185Gbps, and subsequently drops to 7.5Gbps at time T3, the max-aggregate sample for the current normalization window T0-T20 will be 185Gbps. When the current normalization window expires at time T20, we use the previous sampled max-aggregate of 185Gbps to calculate the split/merge activities of the next (or

now current) normalization window between T20 – T40. If traffic remains at 7.5Gbps for this second normalization period, then at time T40 the max-aggregate sample of 7.5Gbps would then be used for split/merge activities. Even though traffic volumes dropped at time T3, LSP split/merge activities would not occur until time T40 which might be unexpected. This default behavior can be modified with use-average-aggregate or use-percentile to achieve alternative normalization behavior if desired.

Options

`cut-off-threshold` *`percentile`*

Specify the percentile value to be used as a cut-off threshold in removing outlier bandwidth samples. All the aggregate bandwidth samples determined as outliers are used for computing aggregate bandwidth used at the time of normalization.

- **Default:** 0 percentile (the ingress considers all aggregate bandwidth samples for normlaization.)
- **Range:** 0 through 100

`use-average-aggregate`

Specify the ingress router to take average of the aggregate samples for normalization.

This option is mutually exclusive with the `use-percentile` configuration option.

`use-percentile` *`percentile`*

Specify the ingress router to compute and use the `pth` percentile from all the bandwidth samples, and use that for normalization.

This option is mutually exclusive with the `use-average-aggregate` configuration option.

- **Range:** 0 through 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

splitting-merging | 28391

sanity-poll

IN THIS SECTION

- Syntax | 16302
- Hierarchy Level | 16302
- Description | 16303
- Options | 16303
- Required Privilege Level | 16303
- Release Information | 16303

Syntax

```
sanity-poll {
    retry-count number;
    on-error {
        raise-alarm;
        power (cycle | off);
        write-coredump;
    }
}
```

Hierarchy Level

```
[edit chassis cfeb slot-number],
[edit chassis feb slot-number],
[edit chassis fpc slot-number],
[edit chassis lcc number fpc number]
```

Description

Enable sanity polling and start periodic sanity checking for a particular FPC. The periodic sanity check includes checking for error conditions such as “register sanity issues,” “high temperature,” “hardware failure,” and so on in the FPC.

NOTE: Currently, periodic sanity check is performed only on the routing chip register.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Sanity Polling

retry-count (FPC Error Condition) | **15698**

on-error | **11819**

sap

IN THIS SECTION

- Syntax | **16304**
- Hierarchy Level | **16304**
- Description | **16304**

- Options | 16304
- Required Privilege Level | 16304
- Release Information | 16305

Syntax

```
sap {
    disable;
    listen address <port port>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit protocols]
```

Description

Enable the router to listen to session directory announcements for multimedia and other multicast sessions.

SAP and SDP always listen on the default SAP address and port, 224.2.127.254:9875. To have SAP listen on additional addresses or pairs of address and port, include a **listen** statement for each address or pair.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Session Announcement Protocol

listen | 9842

satellite (Junos Fusion Automatic Satellite Conversion)

IN THIS SECTION

- Syntax | 16305
- Hierarchy Level | 16305
- Description | 16305
- Options | 16306
- Required Privilege Level | 16306
- Release Information | 16306

Syntax

```
satellite [slot-id | range | all];
```

Hierarchy Level

```
[edit chassis satellite-management auto-satellite-conversion]
```

Description

Specify the interface to enable automatic software conversion in a Junos Fusion.

The device that is cabled to the slot specified in this command is automatically converted into a satellite device.

Additional configuration steps are required to add satellite devices to a Junos Fusion using automatic satellite conversion. See Configuring or Expanding a Junos Fusion Enterprise or *Configuring Junos Fusion Provider Edge*.

Options

slot-id Specifies the FPC slot identifier of the device that will be automatically converted into a satellite device.

The FPC identifier must be mapped to a cascade port interface before this command is operational. See Configuring or Expanding a Junos Fusion Enterprise or *Configuring Junos Fusion Provider Edge*.

range Specifies a range of FPC slot identifiers that will automatically be converted into satellite devices. For instance, to specify that FPC IDs 103, 104, and 105 should be automatically converted into satellite devices, enter a *range* of **103-105**.

all Specifies that all FPC slot identifiers in the Junos Fusion will automatically be converted into satellite devices.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

satellite (Junos Fusion Satellite Device Homing)

IN THIS SECTION

- Syntax | 16307
- Hierarchy Level | 16307
- Description | 16307
- Options | 16307
- Required Privilege Level | 16308
- Release Information | 16308

Syntax

```
satellite [slot-id |slot-id-range | all];
```

Hierarchy Level

```
[edit chassis satellite-management single-home]
```

Description

Specify which satellite device links are single-homed to the aggregation device.

You must enter this statement to configure a Junos Fusion when the aggregation device is running Junos OS Release 14.2R3 or 14.2R4. You are not required to enter this command when the aggregation device is running Junos OS Release 14.2R5 or later.

Options

slot-id Specifies that a link from a specified satellite device is single-homed to the aggregation device. The *slot-id* is the satellite device member number.

slot-id-range	Specifies that the links from a range of specified satellite devices are single-homed to the aggregation device. The *slot-id-range* includes the satellite device member numbers.
all	Specifies that all links from satellite devices are single-homed to the aggregation device.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

Statement introdcued in Junos OS Release 16.1R1 for Junos Fusion Enterprise.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

satellite (Junos Fusion Satellite Software Upgrade Groups)

IN THIS SECTION

- Syntax | 16309
- Hierarchy Level | 16309
- Description | 16309
- Options | 16309
- Required Privilege Level | 16309
- Release Information | 16309

Syntax

```
satellite [slot-id | range | all];
```

Hierarchy Level

```
[edit chassis satellite-management upgrade-groups upgrade-group-name]
```

Description

Specify the satellite device to add to the satellite software upgrade group.

This statement is entered on an aggregation device in a Junos Fusion. Software upgrade groups are configured and managed using the aggregation device.

Options

slot-id Specifies the FPC slot identification number of the satellite device that is being added to the satellite software upgrade group.

range Specifies a range of FPC slot identifiers to add to the satellite software upgrade group. For instance, to specify that FPC IDs 103, 104, and 105 should be automatically converted into satellite devices, enter a *range* of **103-105**.

all Specifies that all FPC slot identifiers in the Junos Fusion are added to the satellite software upgrade group.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

satellite-management (Junos Fusion)

IN THIS SECTION

- Syntax | 16310
- Hierarchy Level | 16312
- Description | 16312
- Options | 16312
- Required Privilege Level | 16312
- Release Information | 16312

Syntax

```
satellite-management {
    aging-timer aging-timer;
    auto-satellite-conversion {
        satellite [slot-id | range | all];
    }
    cluster cluster-name{
        cascade-ports interface-name;
        cluster-id cluster-id-number;
        fpc slot-id{
            alias alias;
            description description;
            member-id member-id-number;
            system-id mac-address;
        }
    }
    environment-monitoring-policy policy;
    firewall
```

```
        family family-name {
            filter filter-name {
                term term-name {
                    from {
                        match-conditions;
                    }
                    then {
                        action;
                        action-modifiers;
                    }
                }
            }
        }
    fpc slot-id{
        alias alias;
        cascade-ports interface-name;
        description description;
        environment-monitoring-policy policy;
        serial-number serial-number;
        system-id mac-address;
        uplink-failure-detection {
            candidate-uplink-policy policy;
        }
    }
    psu {
        redundancy {
            n-plus-n;
        }
    }
    redundancy-groups {
        chassis-id number;
        redundancy-group-name {
            redundancy-group-id redundancy-group-id-number;
            peer-chassis-id peer-chassis-id-number {
            inter-chassis-link interface-name;
            no-auto-iccp-provisioning;
            satellite satellite-device-fpc-IDs;
            }
        }
    }
    single-home {
        satellite [slot-id |slot-id-range | all];
    }
```

```
    upgrade-groups upgrade-group-name {
        satellite [slot-id | range | all];
    }
    uplink-failure-detection {
        candidate-uplink-policy policy;
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure and manage a Junos Fusion.

If you enter the **delete chassis satellite-management** command to delete a Junos Fusion configuration, we recommend also rebooting the Routing Engines on your device to maximize device performance.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

satellite-policies

IN THIS SECTION

- Syntax | 16313
- Hierarchy Level | 16314
- Description | 16314
- Options | 16314
- Required Privilege Level | 16315
- Release Information | 16315

Syntax

```
satellite-policies {
    <candidate-uplink-port-profile policy-name> {
        <holddown holddown-time>;
        <minimum-links number-of-links>;
        <uplink-port-group uplink-port-group-name>;
        term term-name {
            from {
                product-model model-name;
                <holddown holddown-time>;
                <minimum-links number-of-links>;
                <uplink-port-group uplink-port-group-name>;
            }
        }
    }
    environment-monitoring-policy policy-name {
        <alarm <linkdown>
        term term-name {
            from {
                product-model model-name;
            }
        }
    }
    forwarding-policy {
         policy-name {
```

```
            port-group-extended name;
                filter filter-name;
                mirror-egress port-group-mirror port-group-mirror;
            mirror-ingress port-group-mirror port-group-mirror,
                port-group-uplink port-group-uplnk-name
                    holddown time;
            term term-name {
                from {
                    port-group-extended name;
                        filter filter-name;
                        mirror-egress port-group-mirror port-group-mirror;
                    mirror-ingress port-group-mirror port-group-mirror,
                        port-group-uplink port-group-uplnk-name
                            holddown time;
                    product-model model-name;
                    port-group-extended port-group-alias-name {
                        port-group-uplink port-group-alias-name;
                    }
                }
            }
        }
    }
    port-group-alias port-group-alias-name {
        pic pic-number {
        port [port-number | port-number-range | all];
        }
    }
```

Hierarchy Level

```
[edit policy-options]
```

Description

Configure satellite policies for a Junos Fusion.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

save

IN THIS SECTION

- Syntax | 16315
- QFX and VMX Series | 16316
- Description | 16316
- Options | 16316
- Required Privilege Level | 16317
- Release Information | 16317

Syntax

```
save (<dhcp-security-snoop>filename)
```

QFX and VMX Series

```
save (<dhcp-security-snoop | dhcpv6-security-snoop | slaac-security-snoop>filename)
```

Description

Save the configuration to an ASCII file. The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.

For information on valid filename and URL formats, see Format for Specifying Filenames and URLs in Junos OS CLI Commands.

When saving a file to a remote system, the software uses the `scp/ssh` protocol.

Options

filename—Name of the saved file. You can specify a filename in one of the following ways:

- *filename*—File in the user's home directory (the current directory) on the local flash drive.
- *path/filename*—File on the local flash drive.
- `/var/`*filename* or `/var/`*path/filename*—File on the local hard disk.
- `a:`*filename* or `a:`*path/filename*—File on the local drive. The default path is `/` (the root-level directory). The removable media can be in MS-DOS or UNIX (UFS) format.
- *hostname:/path/filename*, *hostname:filename*, *hostname:path/filename*, or `scp://` *hostname/path/filename*—File on an `scp/ssh` client. This form is not available in the worldwide version of Junos OS. The default path is the user's home directory on the remote system. You can also specify *hostname* as *username@hostname*.
- `ftp://`*hostname/path/filename*—File on an FTP server. You can also specify *hostname* as *username @hostname* or *username*:*password @hostname*. The default path is the user's home directory. To specify an absolute path, the path must start with the string `%2F`; for example, `ftp://`*hostname/%2Fpath/filename*. To have the system prompt you for the password, specify `prompt` in place of the password. If a password is required, and you do not specify the password or `prompt`, an error message is displayed:

```
user@host> file copy ftp://username@ftp.hostname.net//filename
file copy ftp.hostname.net: Not logged in.
user@host> file copy ftp://username:prompt@ftphostname.net//filename
```

Password for *username*@ftp.*hostname*.net:

- `http://hostname/path/filename`—File on a Hypertext Transfer Protocol (HTTP) server. You can also specify *hostname* as *username@hostname* or *username:password@hostname*. If a password is required and you omit it, you are prompted for it.
- `re0:/path/filename` or `re1:/path/filename`—File on a local Routing Engine.

Options for QFX and VMX Series

- `dhcp-security-snoop`—Save DHCP snooping entries
- `dhcpv6-security-snoop`—Save DHCPv6 snooping entries
- `slaac-security-snoop`—Slaac-Snooping security entries

Required Privilege Level

> **NOTE**: The `save` commands for `dhcp-security-snoop`, `dhcpv6-security-snoop`, and `slaac-security-snoop` are restricted to admin only.

view

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Deactivating and Reactivating Statements and Identifiers in a Device Configuration

save (Configuration Mode)

IN THIS SECTION

- Syntax | **16318**
- Description | **16318**

- Options | 16318
- Required Privilege Level | 16319
- Sample Output | 16319
- Release Information | 16319
- Hierarchy Level | 16320

Syntax

```
save filename
```

Description

Save the configuration to an ASCII file. The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.

For information on valid filename and URL formats, see Format for Specifying Filenames and URLs in Junos OS CLI Commands.

When saving a file to a remote system, the software uses the scp/ssh protocol.

Options

filename—Name of the saved file. You can specify a filename in one of the following ways:

- *filename*—File in the user's home directory (the current directory) on the local flash drive.
- *path/filename*—File on the local flash drive.
- /var/*filename* or /var/*path/filename*—File on the local hard disk.
- a:*filename* or a:*path/filename*—File on the local drive. The default path is / (the root-level directory). The removable media can be in MS-DOS or UNIX (UFS) format.
- *hostname:/path/filename*, *hostname:filename*, *hostname:path/filename*, or scp:// *hostname/path/filename*—File on an scp/ssh client. This form is not available in the worldwide version of Junos OS. The default path is the user's home directory on the remote system. You can also specify *hostname* as *username@hostname*.

- `ftp://hostname/path/filename`—File on an FTP server. You can also specify *hostname* as *username* `@`*hostname* or *username*:*password*`@`*hostname*. The default path is the user's home directory. To specify an absolute path, the path must start with the string `%2F`; for example, `ftp://hostname/%2F`*path*/*filename*. To have the system prompt you for the password, specify `prompt` in place of the password. If a password is required, and you do not specify the password or `prompt`, an error message is displayed:

```
user@host> file copy ftp://username@ftp.hostname.net//filename
file copy ftp.hostname.net: Not logged in.
user@host> file copy ftp://username:prompt@ftphostname.net//filename
```

Password for *username*@ftp.*hostname*.net: `re0:/path/filename` or `re1:/path/filename`—File on a local Routing Engine.

`routing-instance`—Name of the routing_instance.

`source-address`—Local address to be used in originating the connection.

`terminal`—Use login terminal.

Required Privilege Level

configure—To enter configuration mode.

Sample Output

The following is an example of a save scenario:

Save a File Using Secure Copy Protocol (scp) with 'source-address' and 'routing-instance' options

To save the configuration to a remote system, enter the following save command using `scp` with the `source-address` and `routing-instance` options:

```
root@host# save scp://user@hostname/path/filename routing-instance instance-name source-address
address
```

Release Information

Command introduced before Junos OS Release 7.4.

Hierarchy Level

All hierarchy levels

RELATED DOCUMENTATION

Deactivating and Reactivating Statements and Identifiers in a Device Configuration

saved-core-context (System)

IN THIS SECTION

- Syntax | 16320
- Hierarchy Level | 16320
- Description | 16320
- Required Privilege Level | 16321
- Release Information | 16321

Syntax

```
(saved-core-context | no-saved-core-context);
```

Hierarchy Level

```
[edit system]
```

Description

Configure whether the device saves core files generated by internal Junos processes, along with contextual information (system log files and a copy of the current configuration):

- `saved-core-context`—The device saves each core file and its associated context in a compressed tar file named **/var/tmp/*process-name*.core.*core-number*.tgz**.
- `no-saved-core-context`—The router does not save cores files and their associated context.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Saving Core Files from Junos OS Processes

saved-core-files (System)

IN THIS SECTION

- Syntax | 16321
- Hierarchy Level | 16322
- Description | 16322
- Options | 16322
- Required Privilege Level | 16322
- Release Information | 16322

Syntax

```
saved-core-files number;
```

Hierarchy Level

```
[edit system]
```

Description

Save core files generated by internal Junos processes, but not the associated contextual information (configuration and system log files).

Options

number Maximum number of core files to save per executable.

- **Default:** 5
- **Range:** 1 through 10

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Saving Core Files from Junos OS Processes

sbfd

IN THIS SECTION

- Syntax | 16323
- Hierarchy Level | 16323
- Description | 16323
- Options | 16323
- Required Privilege Level | 16324
- Release Information | 16324

Syntax

```
sbfd {
    remote-discriminator remote-discriminator;
    destination-ipv6-local-host;
}
```

Hierarchy Level

```
[edit protocols source-packet-routing source-routing-path name primary name bfd-liveness-detection
(LSP)],
[edit protocols source-packet-routing source-routing-path name secondary name bfd-liveness-detection
(LSP)]
```

Description

Configure seamless BFD (S-BFD) parameters in the source routing path.

Options

remote-discriminator Remote discriminator of reflector

- **Range:** 1 through 4294967295

destination-ipv6-local-host Enforce the use of an IPv6 local host address, for S-BFD responders that only supports IPv6 local host address.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Routing Engine-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution

bfd-liveness-detection (LSP) | **2931**

sbfd (protocols bfd)

IN THIS SECTION

- Syntax | **16325**
- Hierarchy Level | **16325**
- Description | **16325**
- Options | **16325**
- Required Privilege Level | **16325**
- Release Information | **16325**

Syntax

```
sbfd {
    local-discriminator  local-discriminator;
    local-discriminator-ip local-discriminator-ip;
}
```

Hierarchy Level

```
[edit protocols bfd sbfd]
```

Description

Configure seamless BFD (S-BFD) parameters in the BFD protocol.

Options

local-discriminator	Local discriminator for seamless BFD responder
local-discriminator-ip	Local discriminator IP address for sBFD responder

Required Privilege Level

routing

Release Information

Statement `local-discriminator-ip` introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

> *Routing Engine-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution*

sbl

IN THIS SECTION

- Syntax | 16326
- Hierarchy Level | 16326
- Description | 16326
- Options | 16326
- Required Privilege Level | 16327
- Release Information | 16327

Syntax

```
sbl {
    profile profile-name {
        custom-tag-string [string];
        (sbl-default-server | no-sbl-default-server);
        spam-action (block | tag-header | tag-subject);
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-spam]
```

Description

Configure Content Security server-based antispam features.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

sbl-default-server

IN THIS SECTION

- Syntax | 16327
- Hierarchy Level | 16327
- Description | 16328
- Required Privilege Level | 16328
- Release Information | 16328

Syntax

```
sbl-default-server;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-spam sbl profile profile-name]
```

Description

Enable the default SBL server lookup. You should enable this feature if you are using server-based spam filtering. (The SBL server is predefined on the device. It ships with the name and address of the SBL server.)

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

scale-mode

IN THIS SECTION

- Syntax | 16328
- Hierarchy Level | 16329
- Description | 16329
- Required Privilege Level | 16329
- Release Information | 16329

Syntax

```
scale-mode
```

Hierarchy Level

```
[edit firewall family filter]
```

Description

Programs the filter on a single replication instance instead of the default behavior of programming the filter on two replication instances. Use `scale-mode` to accommodate more firewall filter terms, when the need is to provide more scale than performance.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 24.1

scale-optimized

IN THIS SECTION

- Syntax | 16329
- Hierarchy Level | 16330
- Description | 16330
- Required Privilege Level | 16330
- Release Information | 16331

Syntax

```
scale-optimized;
```

Hierarchy Level

```
[edit firewall family protocol-family-name filter filter-name],
[edit logical-systems logical-system-name firewall family protocol-family-name filter filter-
name]
```

Description

Specify to optimize the interface specific firewall filters in the PFE itself.

When an interface specific firewall filter is configured with multiple Interface bind point instances, the PTX halp software allocates resources for each interface instance separately, and the resources consumption is directly proportional to the number of bind points. This happens because, RE Junos dfw software creates independent instances for each bind point of an interface specific filter.

For example, if an interface specific firewall filter with 'x' number of prefix matches is bound to 'y' number of interfaces (bind points), the Junos software sends 'x' number of independent firewall instances to the PFE software each with 'y' numbers of prefix matches. This automatically consumes 'x * y' numbers of prefixes in the Alpha match block of FLT, and this leads to Alpha block prefix scale issue in FLT. When you add the `scale-optimized` flag under the filter hierarchy, the interface specific firewall filters are optimized in the PFE itself.

The `scale-optimized` flag has the following limitations:

- Can only be configured along with `interface-specific` flag.
- Applicable only to IPv6 and IPv4 families.
- Does not support the `next term` action.
- Does not work with filter lists. Filter lists create unique interface-specific filter for each interface they are applied on. Filters lists do not have a template filter, from which they are copied. It does not add any advantage having scale-optimized flag.
- Cannot be applied to both input and output directions for an interface.
- The `scale-optimized` flag should be configured along with filter configuration. If you add the `scale-optimized` flag to an existing filter configuration, the filter counters do not increase. To resolve the filter counter issue, create a new filter with the `scale-optimized` flag, replace the filter on the interface and commit it.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1.

RELATED DOCUMENTATION

Guidelines for Configuring Firewall Filters

scale-mode

IN THIS SECTION

- Syntax | 16331
- Hierarchy Level | 16331
- Description | 16332
- Required Privilege Level | 16332
- Release Information | 16332

Syntax

```
scale-mode
```

Hierarchy Level

```
[edit firewall family filter]
```

Description

Programs the filter on a single replication instance instead of the default behavior of programming the filter on two replication instances. Use `scale-mode` to accommodate more firewall filter terms, when the need is to provide more scale than performance.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 24.1

scan-extension

IN THIS SECTION

- Syntax | 16332
- Hierarchy Level | 16333
- Description | 16333
- Required Privilege Level | 16333
- Release Information | 16333

Syntax

```
scan-extension filename;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name scan-options]
```

Description

For antivirus file extension scanning, configure the scan extension setting by specifying the name of the defined file extension list.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Kaspersky feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

scan-mode

IN THIS SECTION

- Syntax | 16334
- Hierarchy Level | 16334
- Description | 16334
- Options | 16334
- Required Privilege Level | 16334
- Release Information | 16334

Syntax

```
scan-mode  (all | by-extension);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name scan-
options]
```

Description

You can scan all content or scan content with specific file extensions. You can use a file extension list to define a set of file extensions that are used in file extension scan mode. The antivirus module can then only scan files with extensions on the scan-extension list.

Options

- `all`—Scan all files.
- `by-extension`—Scan only files with extensions specified in a file extension list custom object.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The scan-mode is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, the statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The scan-mode is not supported from Junos OS Release 15.1X49-D10 onwards.

scan-options (Security Antivirus Avira Engine)

IN THIS SECTION

- Syntax | 16335
- Hierarchy Level | 16335
- Description | 16336
- Options | 16336
- Required Privilege Level | 16336
- Release Information | 16336

Syntax

```
scan-options {
    content-size-limit value;
    decompress-layer-limit value;
    (no-pre-detection | pre-detection);
    timeout value;
    (uri-check|no-uri-check);
}
```

Hierarchy Level

```
[edit security utm default-configuration anti-virus]
```

Description

Configure the antivirus feature to scan specific types of traffic based on various scanning configuration parameters. The scan engine scans the data by accessing the virus pattern database. You can download and uninstall the Avira scan engine. The antivirus module is the software subsystem on the gateway device that scans specific application layer traffic to protect the user from virus attacks and to prevent virus from spreading. The antivirus module software subsystem consists of a virus signature database, an application proxy, the scan manager, and the scan engine. The scan engine requires a valid license

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

The `[edit antivirus scan-options]` hierarchy level is deprecated in Junos OS Release 18.2R1.

scan-options (Security Antivirus Juniper Express Engine)

IN THIS SECTION

- Syntax | 16337
- Hierarchy Level | 16337
- Description | 16337
- Options | 16337
- Required Privilege Level | 16337
- Release Information | 16337

Syntax

```
scan-options {
    content-size-limit value;
    (intelligent-prescreening | no-intelligent-prescreening);
    timeout value;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name]
```

Description

Configure the antivirus feature to scan specific types of traffic based on various scanning configuration parameters.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The scan-options (Security Antivirus Juniper Express Engine) is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The scan-options (Security Antivirus Juniper Express Engine) is not supported from Junos OS Release 15.1X49-D10 onwards.

scan-options (Security Antivirus Kaspersky Lab Engine)

IN THIS SECTION

- Syntax | 16338
- Hierarchy Level | 16339
- Description | 16339
- Options | 16339
- Required Privilege Level | 16339
- Release Information | 16339

Syntax

```
scan-options {
    content-size-limit value;
    decompress-layer-limit value;
    (intelligent-prescreening | no-intelligent-prescreening);
    scan-extension filename;
    scan-mode  (all | by-extension);
    timeout value;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name]
```

Description

Configure the antivirus feature to scan specific types of traffic based on various scanning configuration parameters.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

scan-options (Security Antivirus Sophos Engine)

IN THIS SECTION

- Syntax | 16340
- Hierarchy Level | 16340
- Description | 16340
- Options | 16340
- Required Privilege Level | 16341
- Release Information | 16341

Syntax

```
scan-options {
    content-size-limit value;
    (no-uri-check | uri-check);
    timeout value;
}
```

Hierarchy Level

```
[edit security utm default-configuration antivirus scan-options]
```

Description

Configure the antivirus feature to scan specific types of traffic based on various scanning configuration parameters.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

sccp

IN THIS SECTION

- Syntax | 16341
- Hierarchy Level | 16342
- Description | 16342
- Options | 16342
- Required Privilege Level | 16342
- Release Information | 16343

Syntax

```
sccp {
    application-screen {
        call-flood {
            threshold rate;
        }
        unknown-message {
            permit-nat-applied;
            permit-routed;
        }
    }
```

```
    disable;
    dscp-rewrite {
        code-point string;
    }
    inactive-media-timeout value-in-seconds;
    traceoptions {
        flag {
            all {
                extensive
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Configure sccp to specify the Skinny Client Control Protocol (SCCP) ALG on the device. Skinny Client Control Protocol (SCCP) is a Cisco proprietary protocol for call signaling.

Options

`disable`—Disable the SCCP ALG. By default, the SCCP ALG is enabled.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding SCCP ALGs

scheduler

IN THIS SECTION

- Syntax | 16343
- Hierarchy Level | 16343
- Description | 16343
- Options | 16344
- Required Privilege Level | 16344
- Release Information | 16344

Syntax

```
scheduler scheduler-name;
```

Hierarchy Level

```
[edit class-of-service scheduler-maps map-name forwarding-class class-name]
```

Description

Map a scheduler to a forwarding class using a scheduler map.

> **NOTE**: On QFX5200 only, absolute CoS rate limits for transmit rate and shaping rate do not reflect 50g and 100g interfaces. Therefore this statement does not affect those interfaces for QFX5200 in release 15.1X53-D30.

Options

scheduler-name—Name of the scheduler to map to the forwarding class.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

scheduler (Dynamic Scheduler Maps)

IN THIS SECTION

- Syntax | 16344
- Hierarchy Level | 16345
- Description | 16345
- Options | 16345
- Required Privilege Level | 16345
- Release Information | 16345

Syntax

```
scheduler scheduler-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service scheduler-maps map-name forwarding-class class-name]
```

Description

Associate a scheduler with a scheduler map in a dynamic profile.

Options

scheduler-name—Either the specific name of the scheduler configuration block or the scheduler variable (`$junos-cos-scheduler`).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Schedulers in a Dynamic Profile for Subscriber Access

Dynamic Variables Overview

scheduler (Fabric Queues)

IN THIS SECTION

Syntax | 16346

- Hierarchy Level | 16346
- Description | 16346
- Options | 16346
- Required Privilege Level | 16346
- Release Information | 16346

Syntax

```
scheduler scheduler-name;
```

Hierarchy Level

```
[edit class-of-service fabric scheduler-map priority (high | low)]
```

Description

Define scheduler name. For M320, MX Series, T Series routers and for EX Series switches only, specify a scheduler to associate with a fabric queue. For fabric CoS configuration, schedulers are restricted to transmit rates and drop profiles.

Options

scheduler-name—Name of the scheduler configuration block.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Associating Schedulers with Fabric Priorities

scheduler (Interconnect Device)

IN THIS SECTION

- Syntax | 16347
- Hierarchy Level | 16347
- Description | 16347
- Options | 16347
- Required Privilege Level | 16348
- Release Information | 16348

Syntax

```
scheduler scheduler-name;
```

Hierarchy Level

```
[edit class-of-service scheduler-map-forwarding-class-sets scheduler-map-fcset-name forwarding-class-set fabric-forwarding-class-set-name]
```

Description

Map a scheduler to a fabric forwarding class set using a fabric scheduler map. All of the forwarding classes in the fabric forwarding class set use the mapped scheduler.

Options

scheduler-name—Name of the scheduler to map to the fabric forwarding class set.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1R2.

RELATED DOCUMENTATION

Example: Configuring Queue Schedulers

Example: Configuring CoS Scheduling Across the QFabric System

Understanding CoS Scheduling Across the QFabric System

scheduler-map-chassis

IN THIS SECTION

- Syntax | 16348
- Hierarchy Level | 16349
- Description | 16349
- Default | 16349
- Options | 16349
- Required Privilege Level | 16349
- Release Information | 16349

Syntax

```
scheduler-map-chassis (derived | map-name);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-type-fpc/pic/*]
```

Description

For IQ and IQ2 interfaces, as well as on the 10x10GE MIC with SFP+, assign a custom scheduler to the packet forwarding component queues that control the aggregated traffic transmitted into the entire PIC.

Default

On Intelligent Queuing (IQ) and Intelligent Queuing 2 (IQ2) interfaces, as well as on the 10x10GE MIC with SFP+, the traffic that is fed from the packet forwarding components into the PIC uses low packet loss priority (PLP) by default and is distributed evenly across the four chassis queues (not PIC queues), regardless of the scheduling configuration for each logical interface. This default behavior can cause traffic congestion.

Options

`derived`—Sets the chassis queues to derive their scheduling configuration from the associated logical interface scheduling configuration.

`map-name`—Name of the scheduler map configured at the `[edit class-of-service scheduler-maps]` hierarchy level.

> **BEST PRACTICE**: It is always preferred that you apply a custom scheduler map instead of using the `derived` option. We strongly recommend that you do *not* use the `derived` option if the PIC is oversubscribed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Scheduler Maps to Chassis-Level Queues

scheduler-map (Fabric Queues) | 16353

scheduler-map (CoS Virtual Channels)

IN THIS SECTION

- Syntax | 16350
- Hierarchy Level | 16350
- Description | 16350
- Options | 16350
- Required Privilege Level | 16351
- Release Information | 16351

Syntax

```
scheduler-map map-name;
```

Hierarchy Level

```
[edit class-of-service virtual-channel-groups group-name virtual-channel-name]
```

Description

Apply a scheduler map to this virtual channel.

Options

map-name—Name of the scheduler map.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

default (CoS) | 5187

shaping-rate (CoS Virtual Channels) | 16936

virtual-channel-group (CoS Interfaces) | 30740

virtual-channel-groups | 30742

virtual-channels | 30743

scheduler-map (Dynamic Traffic Shaping)

IN THIS SECTION

- Syntax | 16351
- Hierarchy Level | 16352
- Description | 16352
- Options | 16352
- Required Privilege Level | 16352
- Release Information | 16352

Syntax

```
scheduler-map (map-name);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

Associate a scheduler map name with a traffic-control profile in a dynamic profile.

The scheduler map can be defined dynamically (at the `[edit dynamic-profiles profile-name class-of-service scheduler-maps]` hierarchy level) or statically (at the `[edit class-of-service scheduler-maps]` hierarchy level).

Options

map-name—Name of the scheduler map or the Junos predefined variable (`$junos-cos-scheduler-map`). When you specify the variable, the scheduler-map name is obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

The `$junos-cos-scheduler-map` variable introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Traffic Scheduling and Shaping for Subscriber Access

output-traffic-control-profile (Dynamic CoS Definition) | **12018**

scheduler-map (Fabric Queues)

IN THIS SECTION

- Syntax | 16353
- Hierarchy Level | 16353
- Description | 16353
- Required Privilege Level | 16353
- Release Information | 16354

Syntax

```
scheduler-map priority (high | low) scheduler scheduler-name;
```

Hierarchy Level

```
[edit class-of-service fabric]
```

Description

Mapping of fabric traffic to packet schedulers. For M320, MX Series, T Series routers, and for EX Series switches only, associate a scheduler with a fabric priority.

On EX Series switches, this statement is supported only on EX8200 standalone switches and EX8200 Virtual Chassis.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Associating Schedulers with Fabric Priorities

scheduler-map-forwarding-class-sets

IN THIS SECTION

- Syntax | 16354
- Hierarchy Level | 16354
- Description | 16355
- Options | 16355
- Required Privilege Level | 16355
- Release Information | 16355

Syntax

```
scheduler-map-forwarding-class-sets fabric-scheduler-map-name {
    forwarding-class-set fabric-forwarding-class-set-name scheduler scheduler-name;
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify a fabric scheduler map name to map a scheduler configuration to a fabric forwarding class set. All of the forwarding classes in the fabric forwarding class set use the mapped scheduler.

Options

fabric-scheduler-map-name—Name of the fabric scheduler map.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1R2.

RELATED DOCUMENTATION

- *Configuring CoS Fabric Forwarding Class Set Scheduler Maps (Fabric Scheduler to Fabric FC-Set Mapping)*
- *Example: Configuring Queue Schedulers*
- *Example: Configuring CoS Scheduling Across the QFabric System*
- *Understanding CoS Scheduling Across the QFabric System*

scheduler-map (Interfaces and Traffic-Control Profiles)

IN THIS SECTION

- Syntax | 16356
- Enhanced Transmission Selection (ETS) Hierarchical Scheduling | 16356
- Port Scheduling | 16356

- Description | 16356
- Options | 16356
- Required Privilege Level | 16356
- Release Information | 16357

Syntax

```
scheduler-map map-name;
```

Enhanced Transmission Selection (ETS) Hierarchical Scheduling

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
```

Port Scheduling

```
[edit class-of-service interfaces interface-name]
```

Description

Associate a scheduler map with a traffic control profile.

Options

map-name—Name of the scheduler map.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

scheduler-map (Interfaces and Traffic-Control Profiles)

IN THIS SECTION

- Syntax | 16357
- Hierarchy Level | 16357
- Description | 16357
- Options | 16358
- Required Privilege Level | 16358
- Release Information | 16358

Syntax

```
scheduler-map map-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit class-of-service interfaces interface-name unit logical-unit-number],
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
[edit class-of-service multi-destination]
```

Description

Associate a scheduler map with an interface, a logical interface, a traffic-control profile, or with a multidestination traffic configuration. .

Options

map-name—Name of the scheduler map.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Schedulers

Oversubscribing Interface Bandwidth

output-traffic-control-profile | 12016

scheduler-maps (Definition)

IN THIS SECTION

- Syntax | 16359
- Hierarchy Level | 16359
- Description | 16359
- Options | 16359
- Required Privilege Level | 16359
- Release Information | 16359

Syntax

```
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Create a map of schedulers to forwarding classes (which you associate with queues). You then apply the scheduler map to an interface, a logical interface, or a traffic-control profile, thus creating an association of schedulers to queues.

Options

map-name—Name of the scheduler map.

The remaining statements are explained separately. See CLI Explorer.

See *Configuring Schedulers* .

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

scheduler-maps (Dynamic CoS Definition)

IN THIS SECTION

- Syntax | 16360
- Hierarchy Level | 16360
- Description | 16360
- Options | 16360
- Required Privilege Level | 16361
- Release Information | 16361

Syntax

```
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service]
```

Description

Specify a scheduler map name in a dynamic client profile or a dynamic service profile and associate it with the scheduler configuration and forwarding class.

Options

map-name—Name of the scheduler map.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Support at the [`edit dynamic-profiles` *`profile-name`*] hierarchy level introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Schedulers in a Dynamic Profile for Subscriber Access

scheduler-name

IN THIS SECTION

- Syntax | 16361
- Hierarchy Level | 16362
- Description | 16362
- Options | 16362
- Required Privilege Level | 16362
- Release Information | 16362

Syntax

```
scheduler-name scheduler-name;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name]
```

Description

Specify the schedule (as defined by the `scheduler` *`scheduler-name`* statement) for which the policy is in effect.

Options

`scheduler-name` —Name of the scheduler.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

scheduler (Scheduler Map)

IN THIS SECTION

- Syntax | 16363
- Hierarchy Level | 16363
- Description | 16363

- Options | 16363
- Required Privilege Level | 16363
- Release Information | 16363

Syntax

```
scheduler scheduler-name;
```

Hierarchy Level

```
[edit class-of-service scheduler-maps map-name]
```

Description

Associate a scheduler with a scheduler map.

Options

scheduler-name—Name of the scheduler configuration block.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Schedulers

scheduler (Security Policies)

IN THIS SECTION

- Syntax | 16364
- Hierarchy Level | 16365
- Description | 16365
- Options | 16365
- Required Privilege Level | 16365
- Release Information | 16365

Syntax

```
scheduler scheduler-name {
   daily {
       (all-day | exclude | start-time hh:mm stop-time hh:mm);
   }
   description text;
   friday {
       (all-day | exclude | start-time hh:mm stop-time hh:mm);
   }
    monday {
        (all-day | exclude | start-time hh:mm stop-time hh:mm);
   }
   saturday {
        (all-day | exclude | start-time hh:mm stop-time hh:mm);
   }
   start-date date-time stop-date date-time;
   sunday {
       (all-day | exclude | start-time hh:mm stop-time hh:mm);
   }
    thursday {
        (all-day | exclude | start-time hh:mm stop-time hh:mm);
   }
    tuesday {
       (all-day | exclude | start-time hh:mm stop-time hh:mm);
   }
```

```
     wednesday {
        (all-day | exclude | start-time hh:mm stop-time hh:mm);
    }
}
```

Hierarchy Level

```
[edit schedulers]
[edit tenants <tenant-name> schedulers]
```

Description

Create or modify a scheduler that defines when security policies are in effect.

You configure a scheduler to start at a specific date and time or start on a recurrent basis.

Options

scheduler-name —Name of the scheduler. The scheduler name must consist of 1 to 63 characters that can be letters, numbers, dashes, and underscores and can begin with a number or letter.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. The `description` option added in Junos OS Release 12.1. The `scheduler` option added under the `set tenants tenant-name schedulers` hierarchy in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Security Policies Overview

scheduler (System Security Profile)

IN THIS SECTION

- Syntax | 16366
- Hierarchy Level | 16366
- Description | 16366
- Options | 16367
- Required Privilege Level | 16367
- Release Information | 16367

Syntax

```
scheduler {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of schedulers that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 64

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways

schedulers (CoS)

IN THIS SECTION

- Syntax | 16368
- Hierarchy Level | 16368
- Description | 16369
- Options | 16369
- Required Privilege Level | 16369
- Release Information | 16369

Syntax

```
schedulers {
    scheduler-name  {
        buffer-dynamic-threshold value;
        buffer-size (percent percentage | remainder);
        drop-profile-map loss-priority (low | medium-high | high) protocol protocol drop-profile drop-profile-name;
        excess-rate;
        explicit-congestion-notification;
        priority priority;
        shaping-rate (rate | percent percentage);
        transmit-rate (percent percentage) <exact>;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify scheduler name and parameter values such as minimum bandwidth (`transmit-rate`), maximum bandwidth (`shaping-rate`), and priority (`priority`).

Options

> **NOTE**: The syntax above shows all possible options. The options that appear for you can vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

scheduler-name —Name of the scheduler.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

`buffer-dynamic-threshold` option introduced in Junos OS Evolved Release 22.3R1.

schedulers (CoS)

IN THIS SECTION

- Syntax | 16370
- Hierarchy Level | 16370
- Description | 16370
- Options | 16371
- Required Privilege Level | 16371

Release Information | 16371

Syntax

```
schedulers {
    scheduler-name {
        adjust-minimum rate;
        adjust-percent percentage;
        buffer-dynamic-threshold value;
        buffer-rate (percent percent |  rate);
        buffer-size (seconds | percent percentage | remainder | temporal microseconds);
        drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any |
non-tcp | tcp) drop-profile profile-name;
        excess-priority [ low | medium-low | medium-high | high | none];
        excess-rate (percent percentage | proportion value);
        explicit-congestion-notification;
        priority priority-level;
        shaping-rate (percent percentage | rate);
        transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify scheduler name and parameter values such as minimum bandwidth (`transmit-rate`), maximum bandwidth (`shaping-rate`), and priority (`priority`).

Options

> **NOTE**: The syntax above shows all possible options. The options that appear for you can vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

scheduler-name—Name of the scheduler to be configured.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`buffer-dynamic-threshold` option introduced in Junos OS Evolved Release 22.3R1.

RELATED DOCUMENTATION

- How Schedulers Define Output Queue Properties
- Default Schedulers Overview
- Configuring Schedulers
- Configuring a Scheduler

schedulers (Dynamic CoS Definition)

IN THIS SECTION

- Syntax | 16372
- Hierarchy Level | 16372

- Description | 16372
- Options | 16373
- Required Privilege Level | 16373
- Release Information | 16373

Syntax

```
schedulers {
    scheduler-name{
        adjust-minimum rate;
        adjust-percent percentage;
        buffer-size (percent percentage | remainder | temporal microseconds | $junos-cos-
scheduler-bs);
        drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any |
non-tcp | tcp) drop-profile (profile-name | predefined-variable);
        excess-priority (low | high | $junos-cos-scheduler-excess-priority | none);
        excess-rate (percent percentage  | percent $junos-cos-scheduler-excess-rate);
        priority (priority-level | $junos-cos-scheduler-priority);
        shaping-rate (rate | predefined-variable) <burst-size  bytes>;
        transmit-rate  (rate | percent percentage | remainder | percent percentage $junos-cos-
scheduler-tx) <exact | rate-limit>;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service]
```

Description

Specify scheduler name and parameter values in a dynamic client profile or a dynamic service profile.

Options

scheduler-name—Name of the scheduler to be configured or the Junos OS predefined variable (`$junos-cos-scheduler`). The predefined variable is replaced with the scheduler name obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

The `$junos-cos-scheduler` predefined variable introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Schedulers in a Dynamic Profile for Subscriber Access

scheduler (Dynamic Scheduler Maps) | **16344**

schedulers (Interfaces)

IN THIS SECTION

- Syntax | **16374**
- Hierarchy Level | **16374**
- Description | **16374**
- Default | **16374**
- Options | **16374**

Required Privilege Level | 16374
Release Information | 16374

Syntax

```
schedulers number;
```

Hierarchy Level

```
[edit interfaces]
```

Description

Specify number of schedulers for Ethernet IQ2 PIC port interfaces.

Default

If you omit this statement, the 1024 schedulers are distributed equally over all ports in multiples of 4.

Options

number—Number of schedulers to configure on the port.

- **Range:** 1 through 1024

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring the Number of Schedulers per Port for Ethernet IQ2 PICs

schedulers (Security Policies)

IN THIS SECTION

- Syntax | 16375
- Hierarchy Level | 16375
- Description | 16375
- Options | 16375
- Required Privilege Level | 16376
- Release Information | 16376

Syntax

```
schedulers { ... }
```

Hierarchy Level

```
[edit]
```

Description

Configure schedules for security policies that allow you to control network traffic flow and enforce network security.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security— To view this statement in the configuration.

security-control— To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. The `schedulers` option added under the `tenants` hierarchy in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Security Policies Overview

schema

IN THIS SECTION

- Syntax | 16376
- Hierarchy Level | 16377
- Description | 16377
- Options | 16377
- Required Privilege Level | 16377
- Release Information | 16377

Syntax

```
schema {
    openconfig {
        unhide;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Specify whether OpenConfig statements are available and viewable in the CLI.

Options

openconfig unhide

Unhide the OpenConfig statements in the CLI. By default, the OpenConfig schema is not available through the CLI. To rehide the OpenConfig statements, use the following command:

```
delete system schema openconfig unhide
```

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Installing the OpenConfig Package

schema-version (Flat-File Accounting Options)

IN THIS SECTION

- Syntax | 16378
- Hierarchy Level | 16378
- Description | 16378
- Options | 16378
- Required Privilege Level | 16378
- Release Information | 16379

Syntax

```
schema-version schema-name;
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name]
```

Description

Specify the name of the XML schema that defines the contents and format of the accounting file, and appears in the accounting record header.

Options

schema-name	Name of the schema.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Layer 2 Wholesale

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Flat-File Accounting Overview

scope

IN THIS SECTION

- Syntax | 16379
- Hierarchy Level | 16379
- Description | 16380
- Options | 16380
- Required Privilege Level | 16380
- Release Information | 16380

Syntax

```
scope scope-name {
    interface [ interface-names ];
    prefix destination-prefix;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
```

```
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Description

Configure multicast scoping.

Options

scope-name—Name of the multicast scope.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Multicast Snooping

scope-policy

IN THIS SECTION

- Syntax | 16381
- Hierarchy Level | 16381
- Description | 16381
- Options | 16381

Required Privilege Level | 16381
Release Information | 16382

Syntax

```
scope-policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-options multicast]
```

NOTE: You can configure a scope policy at these two hierarchy levels only. You cannot apply a scope policy to a specific routing instance, because all scoping policies are applied to all routing instances. However, you can apply the **scope** statement to a specific routing instance at the **[edit routing-instances *routing-instance-name* routing-options multicast]** or **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast]** hierarchy level.

Description

Apply policies for scoping. The policy must be correctly configured at the **edit policy-options policy-statement** hierarchy level.

Options

policy-names—Name of one or more multicast scope policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

scope | 16379

scope (Security IDP Chain Attack)

IN THIS SECTION

- Syntax | 16382
- Hierarchy Level | 16382
- Description | 16382
- Options | 16383
- Required Privilege Level | 16383
- Release Information | 16383

Syntax

```
scope (session | transaction);
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain]
```

Description

Specify whether the match should occur over a single session or can be made across multiple transactions within a session.

Options

- `session`—Allow multiple matches for the object within the same session.
- `transaction`—Match the object across multiple transactions that occur within the same session.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

scope (Security IDP Custom Attack)

IN THIS SECTION

- Syntax | 16383
- Hierarchy Level | 16384
- Description | 16384
- Options | 16384
- Required Privilege Level | 16384
- Release Information | 16384

Syntax

```
scope (destination | peer | source);
```

Hierarchy Level

```
[edit security idp custom-attack attack-name time-binding]
```

Description

Specify whether the counting of the attack is from the same source IP address, the same destination IP address, or a peer.

Options

- `destination`—IDP detects attacks to a given destination IP address for the specified number of times, regardless of the source IP address.
- `peer`—IDP detects attacks between source and destination IP addresses of the sessions for the specified number of times.
- `source`—IDP detects attacks from a given source IP address for the specified number of times, regardless of the destination IP address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

screen

IN THIS SECTION

Syntax | 16385
Hierarchy Level | 16389

- Description | 16389
- Options | 16389
- Required Privilege Level | 16390
- Release Information | 16390

Syntax

```
screen {
    ids-option name {
        aggregation {
            destination-prefix-mask destination-prefix-mask;
            destination-prefix-v6-mask destination-prefix-v6-mask;
            source-prefix-mask source-prefix-mask;
            source-prefix-v6-mask source-prefix-v6-mask;
        }
        alarm-without-drop;
        description (Security Screen) description;
        icmp (Security Screen) {
            flood (Security ICMP) <threshold ICMP packets per second>;
            fragment;
            icmpv6-malformed;
            ip-sweep <threshold microseconds in which 10 ICMP packets are detected>;
            large;
            ping-death;
        }
        ip (Security Screen) {
            bad-option;
            block-frag;
            ipv6-extension-header {
                AH-header;
                destination-header {
                    home-address-option;
                    ILNP-nonce-option;
                    line-identification-option;
                    tunnel-encapsulation-limit-option;
                    user-defined-option-type name {
                        to type-high;
                    }
```

```
                }
                ESP-header;
                fragment-header;
                HIP-header;
                hop-by-hop-header {
                    CALIPSO-option;
                    jumbo-payload-option;
                    quick-start-option;
                    router-alert-option;
                    RPL-option;
                    SMF-DPD-option;
                    user-defined-option-type name {
                        to type-high;
                    }
                }
                mobility-header;
                no-next-header;
                routing-header;
                shim6-header;
                user-defined-header-type name {
                    to type-high;
                }
            }
            ipv6-extension-header-limit ipv6-extension-header-limit;
            ipv6-malformed-header;
            loose-source-route-option;
            record-route-option;
            security-option;
            source-route-option;
            spoofing;
            stream-option;
            strict-source-route-option;
            tear-drop;
            timestamp-option;
            tunnel (Security Screen) {
                bad-inner-header;
                gre {
                    gre-4in4;
                    gre-4in6;
                    gre-6in4;
                    gre-6in6;
                }
                ip-in-udp {
```

```
                teredo;
            }
            ipip {
                dslite;
                ipip-4in4;
                ipip-4in6;
                ipip-6in4;
                ipip-6in6;
                ipip-6over4;
                ipip-6to4relay;
                isatap;
            }
        }
        unknown-protocol;
    }
    limit-session {
        by-destination {
            by-protocol {
                icmp {
                    maximum-sessions maximum-sessions;
                    packet-rate packet-rate;
                    session-rate session-rate;
                }
                tcp {
                    maximum-sessions maximum-sessions;
                    packet-rate packet-rate;
                    session-rate session-rate;
                }
                udp {
                    maximum-sessions maximum-sessions;
                    packet-rate packet-rate;
                    session-rate session-rate;
                }
            }
            maximum-sessions maximum-sessions;
            packet-rate packet-rate;
            session-rate session-rate;
        }
        by-source {
            by-protocol {
                icmp {
                    maximum-sessions maximum-sessions;
                    packet-rate packet-rate;
```

```
                    session-rate session-rate;
                }
                tcp {
                    maximum-sessions maximum-sessions;
                    packet-rate packet-rate;
                    session-rate session-rate;
                }
                udp {
                    maximum-sessions maximum-sessions;
                    packet-rate packet-rate;
                    session-rate session-rate;
                }
            }
            maximum-sessions maximum-sessions;
            packet-rate packet-rate;
            session-rate session-rate;
        }
        destination-ip-based destination-ip-based;
        source-ip-based source-ip-based;
    }
    match-direction (input | input-output | output);
    tcp (Security Screen) {
        fin-no-ack;
        land;
        port-scan <threshold microseconds in which 10 attack packets are detected>;
        syn-ack-ack-proxy <threshold un-authenticated connections>;
        syn-fin;
        syn-flood {
            alarm-threshold requests per second;
            attack-threshold proxied requests per second;
            destination-threshold SYN pps;
            source-threshold SYN pps;
            timeout (Security Screen) seconds;
            white-list name {
                destination-address [ destination-address ... ];
                source-address [ source-address ... ];
            }
        }
        syn-frag;
        tcp-no-flag;
        tcp-sweep <threshold microseconds in which 10 TCP packets are detected>;
        winnuke;
    }
```

```
        udp (Security Screen) {
            flood (Security UDP) {
                threshold UDP packets per second;
                white-list [ white-list ... ];
            }
            port-scan <threshold microseconds in which 10 attack packets are detected>;
            udp-sweep <threshold microseconds in which 10 UDP packets are detected>;
        }
    }
    traceoptions (Security Screen) {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace;
    }
    trap <interval seconds>;
    white-list name {
        address [ address ... ];
    }
}
```

Hierarchy Level

```
[edit security],
[edit tenants tenant-name security]
```

Description

Configure the security screen options. For every security zone, you can enable a set of predefined screen options that detect and block various kinds of traffic that the device determines as potentially harmful.

Options

ids-options *screen-name*
: Name of the screen configured at the `security screen ids-options` level. Define screens for the intrusion detection service (IDS).

trap
: Configure trap interval. Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes. Traps are

unsolicited messages sent from an SNMP agent to remote network management systems or trap receivers.

`white-list` Set of IP addresses for allowlist. Configure an allowlist of IP addresses that are to be exempt from the SYN cookie and SYN proxy mechanisms that occur during the SYN flood screen protection process. An allowlist contains known trusted IP addresses and URLs. Content downloaded from locations on the allowlist does not have to be inspected for malware.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `description` option added in Junos OS Release 12.1.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

screen (Security Zones)

IN THIS SECTION

- Syntax | 16391
- Hierarchy Level | 16391
- Description | 16391
- Options | 16391

Required Privilege Level | 16391
Release Information | 16391

Syntax

```
screen screen-name;
```

Hierarchy Level

```
[edit security zones functional-zone management],
[edit security zones security-zone zone-name]
```

Description

Specify a security screen for a security zone.

Options

screen-name —Name of the screen.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

- Attack Detection and Prevention Overview
- Example: Configuring Multiple Screening Options

scripts

IN THIS SECTION

- Syntax | 16392
- Hierarchy Level | 16395
- Description | 16395
- Options | 16395
- Required Privilege Level | 16395
- Release Information | 16395

Syntax

```
scripts {
    commit {
        allow-transients;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        direct-access;
        file filename {
            allow-transients;
            cert-file path;
            checksum hash-algorithm hash-value;
            optional;
            refresh;
            refresh-from url;
            routing-instance routing-instance-name;
            source url;
        }
        max-datasize size;
        refresh;
        refresh-from url;
```

```
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
    language (python | python3);
    load-scripts-from-flash;
    op {
        allow-url-for-python;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        file filename {
            allow-commands "regular-expression";
            arguments {
                argument-name {
                    description descriptive-text;
                }
            }
            cert-file path;
            checksum hash-algorithm hash-value;
            command filename-alias;
            dampen {
                dampen-options {
                    cpu-factor cpu-factor;
                    line-interval line-interval;
                    time-interval time-interval;
                }
            }
            description descriptive-text;
            refresh;
            refresh-from url;
            routing-instance routing-instance-name;
            source url;
            start {
                start-options {
                    mem-factor percentage;
                }
```

```
            }
        }
        max-datasize size;
        no-allow-url;
        refresh;
        refresh-from url;
        start {
            start-options {
                mem-factor percentage;
            }
        }
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
    snmp {
        file filename {
            cert-file path;
            checksum hash-algorithm hash-value;
            oid id {
                priority;
            }
            python-script-user username;
            refresh;
            refresh-from;
            routing-instance routing-instance-name;
            source;
            start {
                start-options {
                    mem-factor percentage;
                }
            }
        }
        max-datasize size;
        refresh;
        refresh-from;
        start {
            start-options {
                mem-factor percentage;
            }
        }
```

```
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
    synchronize;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure scripting mechanisms for Junos OS commit, op, or SNMP scripts.

NOTE: The `traceoptions` statement is not supported on QFabric systems.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Commit Script Overview

Op Script Overview

SNMP Script Overview

sctp

IN THIS SECTION

- Syntax | 16396
- Hierarchy Level | 16397
- Description | 16397
- Options | 16398
- Required Privilege Level | 16398
- Release Information | 16398

Syntax

```
sctp {
    log (Security SCTP) name;
    multichunk-inspection disable;
    nullpdu {
        protocol (ID-0x0000 | ID-0xFFFF);
    }
    profile (Security SCTP) name {
        association-timeout association-timeout;
        drop (Security SCTP) {
            m3ua-service name;
            payload-protocol name;
        }
        handshake-timeout handshake-timeout;
        limit (Security SCTP) {
            address name {
                payload-protocol (asap | bicc | ddp-segment | ddp-stream | diameter-dtls |
diameter-sctp | dua | enrp | h248 | h323 | id | iua | m2pa | m2ua | m3ua | others | qipc |
reserved | s1ap | simco | sua | tali | v5ua | x2ap) {
                    rate rate;
                }
```

```
            }
            payload-protocol (asap | bicc | ddp-segment | ddp-stream | diameter-dtls | diameter-
sctp | dua | enrp | h248 | h323 | id | iua | m2pa | m2ua | m3ua | others | qipc | reserved |
s1ap | simco | sua | tali | v5ua | x2ap) {
                rate rate;
            }
            rate {
                address name {
                    sccp sccp;
                    ssp ssp;
                    sst sst;
                }
                sccp sccp;
                ssp ssp;
                sst sst;
            }
        }
        nat-only;
        permit (Security SCTP) {
            payload-protocol name;
        }
    }
    traceoptions (Security SCTP) {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Use the Stream Control Transmission Protocol (SCTP) commands to configure SCTP objects, configure SCTP logs, set trace options, and set address rate limit.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs]` hierarchy level is replaced by `[edit security gtp]`.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2. Support for the `nat-only` option added in Junos OS Release 12.1X45-D10. Support for the `profile` statement added in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Understanding Stream Control Transmission Protocol

SCTP Configuration

sdsl-bytes

IN THIS SECTION

- Syntax | 16399
- Hierarchy Level | 16399
- Description | 16399
- Options | 16399
- Required Privilege Level | 16399
- Release Information | 16399

Syntax

```
sdsl-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of frame overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for an SDSL access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

> **NOTE**: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `sdsl-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream frame overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS

Traffic Rate Reporting and Adjustment by the ANCP Agent

Configuring the ANCP Agent

sdsl-overhead-adjust

IN THIS SECTION

- Syntax | 16400
- Hierarchy Level | 16401
- Description | 16401
- Options | 16401
- Required Privilege Level | 16401
- Release Information | 16401

Syntax

```
sdsl-overhead-adjust percentage;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the actual downstream rate for an SDSL access line reported in the ANCP Port Up message by multiplying the rate by the specified percentage. The ANCP agent reports the adjusted rate to CoS.

> **NOTE**: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `sdsl-overhead-adjust` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

percentage Percentage by which to multiply the rate.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `sdsl-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

Configuring the ANCP Agent to Report Traffic Rates to CoS

Traffic Rate Reporting and Adjustment by the ANCP Agent

Configuring the ANCP Agent

search

IN THIS SECTION

- Syntax | **16402**
- Hierarchy Level | **16403**
- Description | **16403**
- Options | **16403**
- Required Privilege Level | **16403**
- Release Information | **16403**

Syntax

```
search {
    admin-search {
        distinguished-name distinguished-name;
        password password;
    }
```

```
    search-filter filter-name;
}
```

Hierarchy Level

```
[edit access ldap-options],
[edit access profile profile-name ldap-options]
```

Description

Specify that a search is used to get a user's LDAP distinguished name (DN). The search is performed based on the search filter and the part typed in by the user during authentication.

Options

The remaining statements are explained separately.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring Integrated User Firewall on SRX Series

search-filter

IN THIS SECTION

- Syntax | 16404
- Hierarchy Level | 16404
- Description | 16404
- Options | 16404
- Required Privilege Level | 16404
- Release Information | 16405

Syntax

```
search-filter filter-name;
```

Hierarchy Level

```
[edit access ldap-options search],
[edit access profile profile-name ldap-options search]
```

Description

Specify that a search filter is used to find the user's LDAP distinguished name (DN). For example, a filter of **cn** specifies that the search matches a user whose common name is the username.

Options

filter-name—Name of the filter used to find the user's distinguished name.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring Integrated User Firewall on SRX Series

secondary (Label Switched Path)

IN THIS SECTION

- Syntax | 16405
- Hierarchy Level | 16405
- Description | 16406
- Options | 16406
- Required Privilege Level | 16406
- Release Information | 16406

Syntax

```
secondary path-name  {
                        standby;
 }
```

Hierarchy Level

```
[edit protocols mpls                label-switched-path                lsp-name]
```

Description

Specify one or more secondary paths to use for the label switched path (LSP). You can configure more than one secondary path. All secondary paths are equal, and the first one that is available is chosen.

Options

path-name —Name of a secondary path that you created with the `path` statement.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Path Protection in an MPLS Network (CLI Procedure)

secondary (Precision Time Protocol)

IN THIS SECTION

- Syntax | 16407
- Hierarchy Level | 16407
- Description | 16407
- Required Privilege Level | 16407
- Release Information | 16407

Syntax

```
secondary interface name;
```

Hierarchy Level

```
[edit protocols ptp client interface interface-name],
[edit protocols ptp master interface interface-name]
```

Description

Configure the secondary member link in an aggregated Ethernet bundle to enable Precision Time Protocol (PTP) over a link aggregation group. PTP switches over to the secondary member in the aggregated Ethernet bundle when the primary aggregated Ethernet link is down.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

secondary (Adaptive Services Interfaces)

IN THIS SECTION

- Syntax | 16408
- Hierarchy Level | 16408
- Description | 16408
- Options | 16408
- Required Privilege Level | 16408
- Release Information | 16409

Syntax

```
secondary interface-name;
```

Hierarchy Level

```
[edit interfaces (rsp0 | rsp1) redundancy-options]
```

Description

Specify the secondary (backup) adaptive services interface.

Options

interface-name—The identifier for the adaptive services interface, which must be of the form `sp-`*fpc*/*pic*/*port*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring AS or Multiservices PIC Redundancy

secondary connection (Identity Management Advanced Query)

IN THIS SECTION

- Syntax | 16409
- Hierarchy Level | 16410
- Description | 16410
- Options | 16411
- Required Privilege Level | 16411
- Release Information | 16411

Syntax

```
secondary {
    address address;
    ca-profile ca-profile;
    client-id client-id;
    client-secret client-secret;
    interface interface-name;
    routing-instance routing-instance -name;
    source source-address;
}
```

Hierarchy Level

```
[edit services user-identification identity-management connection]
```

Description

Configure parameters that the SRX Series Firewall uses to connect to the Juniper Identity Management Service (JIMS) secondary server and authenticate to it to obtain an access token. JIMS requires that the SRX Series Firewall use OAuth2 to authenticate to it before the SRX Series Firewall can query the JIMS server for user identity information. The SRX Series Firewall must provide the JIMS server with credentials, including a client ID and a client secret. If the client is authenticated-in this case the SRX Series Firewall—it is granted an access token. (See RFC 6749.) Both the client ID and the client secret must be consistent with the API client configured on the JIMS Service primary server.

In addition to configuring the client ID and the client secret, you configure a ca-certificate for the secondary server, if one exists. You configure the file name of the JIMS's ca-certificate. The certificate enables the SRX Series Firewall to verify the identity of JIMS and that it is trusted for the SSL connection.

The SRX Series Firewall always attempts to connect to the primary server first. When one or more queries to the primary server fails, the system falls back to the secondary server.

address- Configure the IP address for the secondary JIMS server. The SRX Series Firewall requires the server IP address to connect to the server to obtain an access code that allows it to query the server for user identity information. The IP address is configured as part of a collection of information which includes the SRX Series Firewall's client ID, client secret, and ca-certificate information.

The SRX Series Firewall uses the secondary server when the primary one fails. You configure the SRX Series Firewall to connect to the primary server separately. This feature supports only IPv4 addresses.

client-id- Client ID that the SRX Series provides to the JIMS Service secondary server as part of its authentication to it. The SRX Series Firewall must authenticate to the server to obtain an access token that allows the SRX Series Firewall to query the server for user identity information The client ID must be consistent with the API client configured on the JIMS primary server.

client-secret- Client secret that the SRX Series provides to the JIMS secondary server as part of its authentication to it. The client secret must be consistent with the API client configured on the JIMS secondary server.

Interface- Client interface name to connect with JIMS server.

routing-instance Client routing instance name to connect with JIMS server. When the client interface connects to JIMS server, routing-instance is auto selected based on the location of interface.

sourceSource address of the request depends on the JIMS server status. If the status is online, then SRX Series Firewall gets source address otherwise the source address is auto selected.

> **WARNING**: Before you use this feature, you must disable active-directory-access and authentication-source options under the user-identification hierarchy. You cannot commit this configuration if active directory authentication or the ClearPass query and webapi functions are configured and committed.

Options

address	IP address of the secondary server.
ca-profile	CA profile name
client-id	Client ID for OAuth2 grant
client-secret	Client secret for OAuth2 grant

Required Privilege Level

1. services—To view this statement in the configuration.
2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D100.

IPv6 address support introduced in Junos OS Release 18.3R1.

Source, interface, and routing-instance options are introduced in Junos OS Release 21.1R1.

Option ca-profile introduced in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

Understanding Advanced Query Feature for Obtaining User Identity Information from JIMS

filter (Identity Management Advanced Query) | **7138**

ip-query (Identity Management Advanced Query) | **9259**

authentication-entry-timeout (Identity Management Advanced Query) | **2597**

batch query | **2886**

secondary-dns

IN THIS SECTION

- Syntax | 16412
- Hierarchy Level | 16412
- Description | 16412
- Options | 16412
- Required Privilege Level | 16412
- Release Information | 16413

Syntax

```
secondary-dns secondary-dns;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp],
[edit access profile profile-name client client-name ppp]
```

Description

Configure the secondary DNS server.

Options

secondary-dns—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

secondary-interface (Aggregated Inline Services)

IN THIS SECTION

- Syntax | 16413
- Hierarchy Level | 16413
- Description | 16413
- Options | 16414
- Required Privilege Level | 16414
- Release Information | 16414

Syntax

```
secondary-interface interface-name;
```

Hierarchy Level

```
[edit interfaces asix aggregated-inline-services-options]
```

Description

Specify the secondary (backup) inline services member link in the asi bundle. You must also configure a primary (active) member link on a different MPC with the `primary-interface` statement. The secondary

member provides 1:1 redundancy for subscriber service sessions on the primary member link. The bandwidth configured at the `[edit chassis fpc slot pic number inline-services bandwidth]` hierarchy level must be the same for both member links.

Redundancy is not achievable if you configure the primary and secondary interface on the same MPC, because both member interfaces go down if the card goes down. Consequently, if you configure both interfaces on the same MPC, the subsequent configuration commit fails.

Options

interface-name	Name of an inline services physical interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

secondary (Link Services IQ PIC Interfaces)

IN THIS SECTION

- Syntax | 16415
- Hierarchy Level | 16415
- Description | 16415
- Options | 16415

Required Privilege Level | 16415
Release Information | 16415

Syntax

```
secondary interface-name;
```

Hierarchy Level

```
[edit interfaces rlsqnumber redundancy-options]
```

Description

Specify the secondary (backup) Link Services IQ PIC interface.

Options

interface-name—The identifier for the Link Services IQ PIC interface, which must be of the form `lsq-fpc/pic/port`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces

secondary (Protocols MPLS)

IN THIS SECTION

- Syntax | 16416
- Hierarchy Level | 16417
- Description | 16417
- Options | 16417
- Required Privilege Level | 16417
- Release Information | 16417

Syntax

```
secondary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps;
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
    retry-limit number;
    retry-timer seconds;
    select (manual | unconditional);
    standby;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Specify one or more secondary paths to use for the LSP. You can configure more than one secondary path. All secondary paths are equal, and the first one that is available is chosen.

You can specify secondary paths even if you have not specified any primary paths.

Optionally, you can specify preference, CoS, and bandwidth values for the secondary path, which override any equivalent values that you configure for the LSP (at the `[edit mpls label-switched-path]` hierarchy level).

Options

path-name—Name of a path that you created with the `path` statement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Primary and Secondary LSPs

secondary-server

IN THIS SECTION

- Syntax | 16418
- Hierarchy Level | 16418
- Description | 16418
- Options | 16418
- Required Privilege Level | 16419
- Release Information | 16419

Syntax

```
secondary-server {
    address ipv4-address;
    login sender-email-address {
        password password;
    }
}
```

Hierarchy Level

```
[edit smtp]
```

Description

Configure Simple Mail Transfer Protocol (SMTP) secondary server for access authorization for SMTP requests.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 10.0.

secondary-vtep-address

IN THIS SECTION

- Syntax | 16419
- Hierarchy Level | 16419
- Description | 16420
- Required Privilege Level | 16420
- Release Information | 16420

Syntax

```
secondary-vtep-address anycast IP address;
```

Hierarchy Level

```
[edit dynamic-profiles name protocols pim],
[edit logical-systems name protocols pim],
[edit logical-systems name routing-instances name protocols pim],
[edit logical-systems name tenants name routing-instances name protocols pim],
[edit protocols pim],
[edit routing-instances name protocols pim],
```

```
[edit tenants name routing-instances name protocols pim]
```

Description

The secondary Virtual Tunnel End Point (VTEP) address.

The *anycast IP address* must match the address under `set routing-instances instance name vtep-source-interface inet evpn-mpls-encap vxlan-source-ip` ***address***.

In earlier versions of Junos, the *anycast IP address* was used by PIM as the primary address to build the VTEP. Starting in Junos OS Release 22.2R1, the *anycast IP address* is not considered the primary address, therefore the `secondary-vtep-address` is necessary for the formation of the VTEP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

No Link Title

secondary-wins

IN THIS SECTION

- Syntax | 16421
- Hierarchy Level | 16421
- Description | 16421
- Options | 16421

Required Privilege Level | 16421
Release Information | 16421

Syntax

```
secondary-wins secondary-wins;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp],
[edit access profile profile-name client client-name ppp]
```

Description

Configure the secondary Windows Internet name server.

Options

secondary-wins—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Point-to-Point Protocol (PPP)

Layer 2 Tunneling Protocol (L2TP)

secret (Access Profile)

IN THIS SECTION

- Syntax | 16422
- Hierarchy Level | 16422
- Description | 16422
- Options | 16422
- Required Privilege Level | 16422
- Release Information | 16423

Syntax

```
secret password;
```

Hierarchy Level

```
[edit access profile profile-name radius-server server-address]
```

Description

Specify the RADIUS secret password, which is shared between the router and the RADIUS server. The device uses this secret to encrypt the user's password that is sent to the RADIUS server.

Options

password—RADIUS secret. Maximum length is 256 characters.

Required Privilege Level

secret—To view this statement in the configuration.

secret-control—To add this statement to the configuration.

Release Information

Statement modified in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring the Device as a Junos OS Enforcer Using IPsec (CLI)

secret-key-timeout

IN THIS SECTION

- Syntax | 16423
- Hierarchy Level | 16423
- Description | 16424
- Default | 16424
- Options | 16424
- Required Privilege Level | 16424
- Release Information | 16424

Syntax

```
secret-key-timeout minutes;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols amt relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols amt
relay],
[edit protocols amt relay],
[edit routing-instances routing-instance-name protocols amt relay]
```

Description

Specify the period in minutes after which the local opaque secret key used in the Automatic Multicast Tunneling (AMT) Message Authentication Code (MAC) times out and is regenerated.

Default

60 minutes

Options

minutes—Number of minutes to wait before generating a new MAC opaque secret key.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

secret (RADIUS)

IN THIS SECTION

- Syntax | **16425**
- Hierarchy Level | **16425**
- Description | **16425**
- Options | **16425**

Required Privilege Level | 16425
Release Information | 16425

Syntax

```
secret password;
```

Hierarchy Level

```
[edit access profile profile-name radius-server server-address],
[edit access radius-disconnect client-address],
[edit access radius-server server-address]
```

Description

Configure the password to use with the RADIUS server. The secret password used by the local router or switch must match that used by the server.

Options

password—Password to use; it can include spaces if the character string is enclosed in quotation marks.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

Example: Configuring CHAP Authentication with RADIUS

Configuring RADIUS Authentication for L2TP

secret (RADIUS Client)

IN THIS SECTION

- Syntax | 16426
- Hierarchy Level | 16426
- Description | 16426
- Options | 16427
- Required Privilege Level | 16427
- Release Information | 16427

Syntax

```
secret password;
```

Hierarchy Level

```
[edit access radius clients client-name accounting]
```

Description

Specify a shared secret to be used by the MX Series router and the RADIUS client for accounting.

Options

password Shared secret to use ; it can include spaces if the character string is enclosed in quotation marks. Maximum length is 256 characters.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers

secret (RADIUS Server)

IN THIS SECTION

- Syntax | 16428
- Hierarchy Level | 16428
- Description | 16428
- Options | 16428
- Required Privilege Level | 16428
- Release Information | 16428

Syntax

```
secret password;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure a shared secret to be used by the MX Series router and the RADIUS server.

Options

password Shared secret to use.

- **Range:** 1 through 64 characters

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

secret (Tunnel Profile)

IN THIS SECTION

- Syntax | 16429
- Hierarchy Level | 16429
- Description | 16429
- Options | 16429
- Required Privilege Level | 16429
- Release Information | 16430

Syntax

```
secret password;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Description

Specify the tunnel password sent to the LNS for authentication.

Options

password—Cleartext password.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

secure

IN THIS SECTION

- Syntax | 16430
- Hierarchy Level | 16431
- Description | 16431
- Required Privilege Level | 16431
- Release Information | 16431

Syntax

```
secure {
    security-level {
        (default | secure-messages-only);
    }
    cryptographic-address {
        key-length number;
        key-pair pathname;
    }
    timestamp {
        clock-drift number;
        known-peer-window seconds;
        new-peer-window seconds;
    }
    traceoptions {
```

```
        file filename <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit protocols neighbor-discovery]
```

Description

Configure parameters for Secure Neighbor Discovery.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Secure IPv6 Neighbor Discovery

secure-access-port

IN THIS SECTION

- Syntax | 16432
- Hierarchy Level | 16434
- Description | 16434
- Required Privilege Level | 16434
- Release Information | 16434

Syntax

```
secure-access-port {
    dhcp-snooping-file {
        location local_pathname | remote_URL;
        timeout seconds;
        write-interval seconds;
    }
    dhcpv6-snooping-file {
        location local_pathname | remote_URL;
        timeout seconds;
        write-interval seconds;
    }
    interface (all | interface-name) {
        allowed-mac {
            mac-address-list;
        }
        (dhcp-trusted | no-dhcp-trusted);
        fcoe-trusted;
        mac-limit   limit action (drop | log | none | shutdown);
        no-allowed-mac-log;
        persistent-learning;
        static-ipip-address {
            vlan  vlan-name;
            mac mac-address;
        }
        static-ipv6ip-address {
```

```
            vlan  vlan-name;
            mac mac-address;
        }
        voip-mac-exclusive;
        (dhcp-trusted | no-dhcp-trusted);
    }
    vlan (all | vlan-name) {
        (arp-inspection | no-arp-inspection) [
            forwarding-class class-name;
        }
        dhcp-option82 {
            circuit-id {
                prefix hostname;
                use-interface-description;
                use-vlan-id;
            }
             remote-id {
                prefix hostname | mac | none;
                use-interface-description;
                use-string                                  string;
            }
            vendor-id <string>;
        }
        (examine-dhcp | no-examine-dhcp) {
            forwarding-class class-name;
        }
        (examine-dhcpv6 | no-examine-dhcpv6) {
            forwarding-class class-name;
        }
        examine-fip {
            fc-map fc-map-value;
        }
        (ip-source-guard | no-ip-source-guard);
        (ipv6-source-guard | no-ipv6-source-guard);
        mac-move-limit limit action (drop | log | none | shutdown);
        }
        (neighbor-discovery-inspection | no-neighbor-discovery-inspection);
        no-option37;
    }
}
```

Hierarchy Level

```
[edit ethernet-switching-options]
```

Description

Configure port security features, including MAC limiting, dynamic ARP inspection, whether interfaces can receive DHCP responses, DHCP snooping, IP source guard, DHCP option 82, MAC move limiting, and FIP snooping.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for IPv6 introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

RELATED DOCUMENTATION

- Example: Configuring Port Security (non-ELS)
- Example: Protecting Against Address Spoofing and Layer 2 DoS Attacks
- Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN
- Example: Setting Up DHCP Option 82 Using the Same VLAN
- *Example: Configuring an FCoE Transit Switch*

secure-channel

IN THIS SECTION

- Syntax | 16435
- Hierarchy Level | 16435
- Description | 16435
- Options | 16436
- Required Privilege Level | 16436
- Release Information | 16436

Syntax

```
secure-channel secure-channel-name {
    direction (inbound | outbound);
    encryption;
    id {
        mac-address mac-address;
        port-id port-id-number;
    }
    offset (0|30|50);
    security-association security-association-number {
        key key-string;
    }
}
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Create and configure a secure channel to enable and configure MACsec when MACsec is enabled using static secure association key (SAK) security mode.

You do not need to use this option to enable MACsec using static connectivity association key (CAK) security mode. All configuration for MACsec using static CAK security mode is done inside of the connectivity association but outside of the secure channel. When MACsec is enabled using static CAK security mode, an inbound and an outbound secure channel—neither of which is user-configurable—is automatically created within the connectivity association.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

secure-nat-mapping

IN THIS SECTION

- Syntax | 16437
- Hierarchy Level | 16437
- Description | 16437
- Options | 16437
- Required Privilege Level | 16437
- Release Information | 16437

Syntax

```
secure-nat-mapping {
    mapping-refresh (inbound | outbound | inbound-outbound);
    eif-flow-limit number-of-flows'
}
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated]
```

Description

Specify configuration options that help prevent or minimize the effect of attempted denial of service (DOS) attacks for NAT operations.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Protecting CGN Devices Against Denial of Service (DOS) Attacks

secure-nat-mapping (Source NAT Next Gen Services)

IN THIS SECTION

- Syntax | 16438
- Hierarchy Level | 16438
- Description | 16438
- Options | 16438
- Required Privilege Level | 16439
- Release Information | 16439

Syntax

```
secure-nat-mapping {
    eif-flow-limit number-of-flows;
    mapping-refresh (inbound | inbound-outbound | outbound);
}
```

Hierarchy Level

```
[edit services nat source rule-set rule-set rule rule-name then source-nat]
```

Description

For endpoint-independent mapping, configure the maximum number of simultaneous inbound flows and the direction in which mappings are refreshed.

Options

eif-flow-limit *number-of-flows*

Maximum number of simultaneous inbound flows.

- **Range:** 0 through 655334

mapping-refresh (inbound | inbound-outbound | outbound)	Direction in which mappings are refreshed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

secure-neighbor-discovery

IN THIS SECTION

- Syntax | 16439
- Hierarchy Level | 16440
- Description | 16440
- Options | 16440
- Required Privilege Level | 16440
- Release Information | 16440

Syntax

```
secure-neighbor-discovery {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Provide support for protecting Secure Neighbor Discovery Protocol (SEND) messages.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the Secure Neighbor Discovery (SEND) protocol process.
- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take the primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

secure-proxy

IN THIS SECTION

- Syntax | 16441
- Hierarchy Level | 16441
- Description | 16441
- Options | 16442
- Required Privilege Level | 16442
- Release Information | 16442

Syntax

```
secure-proxy {
    profile name {
        drop-on-dns-error;
        dynamic-web-application;
        dynamic-web-application-group;
        proxy-address name {
            ip ip-address;
            port port-number;
        }
    }
}
```

Hierarchy Level

```
[edit services web-proxy]
```

Description

Configure secure Web proxy profile.

When you configure secure Web proxy on SRX Series Firewall, it intercepts the session and allows sessions that are interested in specific application and are destined to a configured external web-proxy. The device connects sessions directly to the Web server bypassing the external proxy server. Connections that do not match the applications are routed to the external proxy server.

Since the Secure Web proxy forwards traffic based on applications either to the external proxy server or to the Web server, you can define the routing behavior based on applications. For example, you can specify Office 365 application group in secure Web proxy profile to bypass the external proxy server for connections to Office 365.

To configure secure Web proxy on the SRX Series Firewall, you must define a Web proxy profile by specifying external proxy server details and dynamic application. You can associate this secure Web proxy profile with security policy. The secure Web profile is applied on the traffic matching the application and security policy. The session is now allowed to bypass the external proxy server and connect to the Web server directly.

Options

profile *name*	Name of the secure Web proxy profile.
drop-on-dns-error	Drop the Web proxy session on DNS error.
dynamic-web-application	Dynamic web application.
dynamic-web-application-group	Dynamic web application group.
proxy-address *name*	Name of the external proxy server.
ip *ip-address*	IP address of the external proxy server.
port *port-number*	Port number of the external proxy server.

Required Privilege Level

flow-tap

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Secure Web Proxy

secure-wire (System Security Profile)

IN THIS SECTION

- Syntax | 16443
- Hierarchy Level | 16443
- Description | 16443
- Options | 16444
- Required Privilege Level | 16444
- Release Information | 16444

Syntax

```
secure-wire {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify maximum and reserved quota of a user logical system. If the user logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if the resources are available—that is, if they are not allocated to other user logical systems.

Options

maximum A maximum allowed quota of a user logical system.

- **Range:** 0 through 255

reserved A reserved quota of a user logical system.

- **Default:** 0

Required Privilege Level

admin

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Secure Wire for Logical Systems

secured-port-block-allocation

IN THIS SECTION

- Syntax | 16445
- Hierarchy Level | 16445
- Description | 16445
- Options | 16446
- Required Privilege Level | 16446
- Release Information | 16446

Syntax

```
secured-port-block-allocation {
    active-block-timeout timeout-seconds;
    block-size block-size;
    max-blocks-per-address max-blocks;
}
```

Hierarchy Level

```
[edit services nat pool pool-name port]
```

Description

When you use block allocation, one or more blocks of ports in a NAT pool address range are available for assignment to a subscriber.

Port block allocation is supported on MX series routers with MS-DPCs and on M Series routers with MS-100, MS-400, and MS-500 MultiServices PICS. Port block allocation is supported on MX series routers with MS-MPCs and MS-MICs starting in Junos OS release 14.2R2.

NOTE: If you define the session lifetime globally for a Multiservices (`ms`-) interface (by using the `session-timeout seconds` statement at the `[edit interfaces interface-name services-options]` hierarchy level), the session is terminated even if traffic continues to flow beyond that time period. When continuous traffic transmission occurs, the session is reset immediately after the timeout period. When you configure the same value for the session timeout and the active port block allocation timeout, the system might not determine that the active port block timeout period has elapsed. As a result, when the active port block timeout elapses, the system might use the same block for the initial port allocation that was used previously. However, for the subsequent allocation of a port block, the system identifies the active block timeout value correctly and allocates a port from a new block. This behavior is expected when the session timeout and port block timeout values are identical. To avoid this problem, we recommend that you configure different values for session timeout and port block timeout so that the `JSERVICES_NAT_PORT_BLOCK_ALLOC` system logging message is generated at correct intervals of the active port block timeout value.

> **NOTE**: If you make any configuration changes to a NAT pool that has secured port block allocation configured, you must delete the existing NAT address pool, wait at least 5 seconds, and then configure a new NAT address pool. We also strongly recommend that you perform this procedure if you make any changes to the NAT pool configuration, even when secured port block allocation is not configured.

Options

active-block-timeout *timeout-seconds*

Interval, in seconds, during which a block is active. After the timeout elapses, a new block is allocated, even if ports are available in the active block.

- **Range:** 0 through 86400. When you specify 0, the active block transitions to inactive only when it runs out of ports and a new block is allocated. Any inactive block without any ports in use will be freed to the NAT pool, unless it is active block. The range is 1 through 512.
- **Default:** 120

block-size *block-size*

Number of ports included in a block.

- **Range:** For the Multiservices DPC only, 1 through 32,000
- **Range:** For the Multiservices MPC and Multiservices MIC only, 1 through the total number of configured ports. For example, for a port range of 1024 through 61,024, the block-size range is 1 through 60,000.
- **Default:** 128

max-blocks-per-address *max-blocks*

Maximum number of blocks that can be allocated to a user address. The range is 1 through 512.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Address Pools for Network Address Port Translation (NAPT) Overview

securid-server

IN THIS SECTION

- Syntax | 16447
- Hierarchy Level | 16447
- Description | 16447
- Options | 16447
- Required Privilege Level | 16448
- Release Information | 16448

Syntax

```
            securid-server {
    server-name configuration-file filepath;
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure SecurID server for SecurID authentication type.

Options

The remaining statement is explained separately.

NOTE: You can configure only one SecurID server. SecurID challenges are not yet supported.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 9.1 of Junos OS.

RELATED DOCUMENTATION

Understanding External Authentication Servers

security

IN THIS SECTION

- Syntax | 16448
- Hierarchy Level | 16450
- Description | 16450
- Options | 16450
- Required Privilege Level | 16450
- Release Information | 16450

Syntax

```
security {
    authentication-key-chains {
```

```
        key-chain key-chain-name {
            key key {
                secret secret-data;
                start-time yyyy-mm-dd.hh:mm:ss;
            }
        }
    certificates {
        cache-size bytes;
        cache-timeout-negative seconds;
        certification-authority ca-profile-name {
            ca-name ca-identity;
            crl file-name;
            encoding (binary | pem);
            enrollment-url url-name;
            file certificate-filename;
            ldap-url url-name;
        }
        enrollment-retry attempts;
        local certificate-filename {
            certificate-key-string;
            load-key-file key-file-name;
        }
        maximum-certificates number;
        path-length certificate-path-length;
    }
    ssh-known-hosts {
        host {
            fetch-from-server host-name;
            load-key-file file-name;
        }
    }
    traceoptions {
        file filename <files number> <size size>;
        flag flag;
        level level;
        no-remote-trace
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure security services. Most of the configuration statements do not have default values. If you do not specify an identifier for a statement that does not have a default value, you cannot commit the configuration.

Options

Required Privilege Level

Release Information

Statement introduced in Junos OS Release 11.1.

security-association (MACSec)

IN THIS SECTION

- Syntax | 16451
- Hierarchy Level | 16451
- Description | 16451
- Default | 16451
- Options | 16451
- Required Privilege Level | 16452
- Release Information | 16452

Syntax

```
security-association security-association-number {
    key key-string;
}
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel
secure-channel-name]
```

Description

Specifies the number of one of the security associations in the secure channel when MACsec is enabled using static secure association key (SAK) security mode. Because SAKs are created by the key server when MACsec is enabled using static connectivity association key (CAK) security mode, the `security-association` statement is not used when enabling MACsec using static CAK security mode.

You must configure at least two security associations to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.

Default

No security keys are configured, by default.

Options

security-association-number

Specifies the security association number and creates the SAK.

The security association number is a whole number between 0 and 3. You can configure two security associations in a secure channel when enabling MACsec using static security keys.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

security-association

IN THIS SECTION

- Syntax | 16452
- Hierarchy Level | 16453
- Description | 16453
- Default | 16453
- Options | 16453
- Required Privilege Level | 16453
- Release Information | 16454

Syntax

```
security-association security-association-number {
    key key-string;
}
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name secure-channel secure-channel-name]
```

Description

Specifies the number of one of the security associations in the secure channel when MACsec is enabled using static secure association key (SAK) security mode. Because SAKs are created by the key server when MACsec is enabled using static connectivity association key (CAK) security mode, the `security-association` statement is not used when enabling MACsec using static CAK security mode.

You must configure at least two security associations to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.

Default

No security keys are configured, by default.

Options

security-association-number

Specifies the security association number and creates the SAK.

The security association number is a whole number between 0 and 3. You can configure two security associations in a secure channel when enabling MACsec using static security keys.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

security-association (IPSec)

IN THIS SECTION

- Syntax | 16454
- Hierarchy Level | 16455
- Description | 16455
- Options | 16455
- Required Privilege Level | 16457
- Release Information | 16457

Syntax

```
security-association sa-name {
    manual {
        direction bidirectional {
            authentication {
                algorithm (hmac-md5-96 | hmac-sha1-96 | hmac-sha-256 | hmac-sha-384 | hmac-
sha-512);
                key {
                    ascii-text key;
                    hexadecimal key;
                }
            }
            encryption {
                algorithm (3des-cbc | des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc);
```

```
                key {
                    ascii-text key;
                    hexadecimal key;
                }
            }
            protocol (ah | esp);
            spi spi-value;
        }
    }
    mode transport;
}
```

Hierarchy Level

```
[edit security ipsec]
```

Description

Configure a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link. IPsec can provide authentication and confidentiality to OSPF or OSPFv3 routing packets.

Options

sa-name Name of the SA.

description Specify a text description for the SA.

direction Direction of the manual SA. For this feature, the direction must be `bidirectional`. Decrypt and authenticate the incoming and outgoing traffic using the same algorithm, keys, or SPI in both directions, unlike inbound and outbound SAs that use different attributes in both directions.

- **Values:** `algorithm`—Hash algorithm that authenticates packet data. It can be one of the following:
  - `hmac-md5-96`—Produces a 128-bit digest. This is the default.
  - `hmac-sha1-96`—Produces a 160-bit digest.

- `hmac-sha-256`—Produces a 256-bit digest.
- `hmac-sha-384`—Produces a 384-bit digest.
- `hmac-sha-512`—Produces a 512-bit digest.

Starting in Junos OS Release 22.2R1, MX240, MX480, and MX960 with MX-SPC3, SRX Series Firewalls and vSRX Virtual Firewall running iked process supports all the listed authentication algorithms.

`key`—Type of authentication key. It can be one of the following:

- `ascii-text` *key*—ASCII text key. For `hmac-md5-96`, the key is 16 ASCII characters; for `hmac-sha1-96`, the key is 20 ASCII characters.
- `hexadecimal` *key*—Hexadecimal key. For `hmac-md5-96`, the key is 32 hexadecimal characters; for `hmac-sha1-96`, the key is 40 hexadecimal characters.

- **Values:** `encryption`—Configure an encryption algorithm and key for a manual Security Association (SA). It can be one of the following:
  - `algorithm`—Select the encryption algorithm for the internal Routing-Engine-to-Routing-Engine IPsec security association (SA) configuration.
    - `des-cbc`—Encryption algorithm with block size of 8 bytes (64 bits) and key size 48 bits.
    - `3des-cbc`—Encryption algorithm with block size of 8 bytes (64 bits) and key size of 192 bits.

      For `3des-cbc`, we recommend that the first 8 bytes be different from the second 8 bytes, and the second 8 bytes be the same as the third 8 bytes.
    - `aes-128-cbc`—Advanced encryption algorithm that has a key size of 16 bytes; its key size is 128 bits long.
    - `aes-192-cbc`—Advanced encryption algorithm that has a key size of 24 bytes; its key size is 192 bits long.
    - `aes-256-cbc`—Advanced encryption algorithm that has a key size of 32 bytes; its key size is 256 bits long.

    Starting in Junos OS Release 22.2R1, MX240, MX480, and MX960 with MX-SPC3, SRX Series Firewalls and vSRX Virtual Firewall running iked process supports all the listed encryption algorithms.

- `key`—Type of encryption key. It can be one of the following:
  - `ascii-text key`—ASCII text key. For the `des-cbc` option, the key contains 8 ASCII characters; for `3des-cbc`, the key contains 24 ASCII characters.
  - `hexadecimal key`—Hexadecimal key. For the `des-cbc` option, the key contains 16 hexadecimal characters; for the `3des-cbc` option, the key contains 48 hexadecimal characters.

protocol Define the IPsec protocol for a manual security association (SA). The protocol can be one of the following:

- `ah`—Authentication Header protocol. If you configure AH protocol, it is mandatory to configure the authentication algorithm and the key.
- `esp`—Encapsulating Security Payload (ESP) protocol. This is the default.

  If you configure ESP protocol, it is mandatory to configure either authentication algorithm or encryption algorithm or both. If you did not configure ESP protocol and did not configure either authentication or encryption algorithm, then we do not provide authentication or encryption support.

spi *spi-value* Configure the security parameter index (SPI) for a security association (SA). An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).

- **Range:** 256 through 16,639

mode SA mode. For this feature, the mode must be `transport`.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D20.

Authentication algorithm configuration options, hmac-md5-96, hmac-sha1-96, hmac-sha-256, hmac-sha-384, and hmac-sha-512 are added in Junos OS Release 22.2R1 for MX240, MX480, and MX960 with MX-SPC3, SRX Series Firewalls and vSRX Virtual Firewall running iked process.

Encryption algorithm configuration options, des-cbc, 3des-cbc, aes-128-cbc, aes-192-cbc, and aes-256-cbc are added in Junos OS Release 22.2R1 for MX240, MX480, and MX960 with MX-SPC3, SRX Series Firewalls and vSRX Virtual Firewall running iked process.

RELATED DOCUMENTATION

Understanding OSPF and OSPFv3 Authentication on SRX Series Devices

security-association (Junos-FIPS Software)

IN THIS SECTION

- Syntax | 16458
- Hierarchy Level | 16459
- Description | 16459
- Options | 16459
- Required Privilege Level | 16459
- Release Information | 16459

Syntax

```
security-association sa-name {
    dynamic {
        ipsec-policy policy-name;
        replay-window-size (32 | 64);
    }
    manual {
        direction (inbound | outbound | bi-directional) {
            authentication {
                algorithm (hmac-sha1-96 | hmac-sha2-256);
                key (ascii-text key | hexadecimal key);
            }
            auxiliary-spi auxiliary-spi-value;
            encryption {
```

```
                algorithm 3des-cbc;
                key (ascii-text key | hexadecimal key);
            }
            protocol ( ah | esp | bundle);
            spi spi-value;
        }
        mode (tunnel | transport);
    }
}
```

Hierarchy Level

```
[edit security ipsec]
```

Description

Configure an IPsec security association.

Options

sa-name—Name of the security association.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Crypto Officer—To view and add this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

NOTE: We recommend that you configure the IPsec keys as hexadecimal keys for maximum key strength with Junos OS in FIPS mode.

security-association (Junos OS)

IN THIS SECTION

- Syntax | 16460
- Hierarchy Level | 16461
- Description | 16461
- Options | 16461
- Required Privilege Level | 16461
- Release Information | 16461

Syntax

```
security-association sa-name {
    dynamic {
        ipsec-policy policy-name;
        replay-window-size (32 | 64);
    }
    manual {
        direction (inbound | outbound | bi-directional) {
            authentication {
                algorithm (hmac-sha1-96 | hmac-sha2-256);
                key (ascii-text key | hexadecimal key);
            }
            auxiliary-spi auxiliary-spi-value;
            encryption {
                algorithm (des-cbc | 3des-cbc);
                key (ascii-text key | hexadecimal key);
            }
            protocol ( ah | esp | bundle);
            spi spi-value;
        }
        mode (tunnel | transport);
    }
}
```

Hierarchy Level

```
[edit security ipsec]
```

Description

Configure an IPsec security association.

Options

sa-name—Name of the security association.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

> **NOTE**: You must configure the IPsec keys as hexadecimal keys for maximum key strength with Junos OS in FIPS mode.

RELATED DOCUMENTATION

Configuring Security Associations for IPsec on an ES PIC

security-intelligence (Services)

IN THIS SECTION

- Syntax | 16462
- Hierarchy Level | 16462
- Description | 16462
- Options | 16463
- Required Privilege Level | 16463
- Release Information | 16464

Syntax

```
    authentication {
        auth-token auth-token;
    tls-profile tls-profile;
    traceoptions {
        no-remote-trace;
        file [ filename <files number> <size  bytes> <match expression> <world-readable | no-
world-readable>];
        flag [all | feed | ipc];
        level [all| error | info | notice | verbose | warning];
    no-remote-trace;
url url;
```

Hierarchy Level

```
[edit services]
```

Description

You can configure security intelligence profiles and policies to work with security intelligence feeds, such as infected hosts and C&C. You then configure a firewall policy to include the security intelligence policy, for example, block outgoing requests to a C&C host.

Options

authentication Configure authentication, such as an auth token or TLS profile, to commute with the feed server. This operation is performed by the ops script used to enroll your devices and is typically not required afterwards. If you have problems establishing a connection with the Juniper ATP Cloud cloud server, we recommend that you rerun the ops script instead of manually entering all the CLI commands.

traceoptions Set security intelligence trace options.

- file—Name of the file to receive the output of the tracing operation.
  - `files` *`number`* —Maximum number of trace files

    Range: 2 through 1000
  - match— Regular expression for lines to be logged
  - no-world-readable—Prevent any user from reading the log file
  - size—Maximum size of each trace file

    Range: 10240 through 1073741824
  - world-readable—Allow any user to read the log file
- flag—Tracing operation to perform
  - all—All interface tracing operation
  - feed—Trace feed operation
  - ipc—Trace interface interprocess communication (IPC) module messages
- level—Level of debugging output
- no-remote-trace—Disable the remote trace

url *url-address* Configure the URL of the feed server. This operation is performed by the ops script used to enroll your devices and is typically not required afterwards. If you have problems establishing a connection with the Juniper ATP Cloud cloud server, we recommend that you rerun the ops script instead of manually entering all the CLI commands.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers with Juniper Juniper Advanced Threat Prevention Cloud (ATP).

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480, and MX960. This support runs inline on the MPC card.

security-intelligence

IN THIS SECTION

- Syntax | 16464
- Hierarchy Level | 16464
- Description | 16465
- Options | 16465
- Required Privilege Level | 16465
- Release Information | 16465

Syntax

```
security-intelligence {
(add-source-ip-to-feed feed-name | add-destination-ip-to-feed feed-name)
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
application-services]
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then deny
application-services ]
```

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then reject
application-services]
```

Description

Starting in Junos OS Release 20.2R1, you can add source and destination addresses to the security intelligence (SecIntel) profiles to generate security feeds in a security policy by configuring the `security-intelligence` command. After the feeds are generated, you can configure other security policies to use the feeds as a `dynamic-address` to match designated traffic and perform policy actions.

Options

add-destination-ip-to-feed Add the destination IP address to thesecurity feed.

add-source-ip-to-feed Add the source IP address to the security feed.

The feed-name must be predefined and installed by SecIntel with the feed type and a specific ID.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

security-intelligence (IDP)

IN THIS SECTION

- Syntax | 16466
- Hierarchy Level | 16466
- Description | 16466
- Options | 16466

Required Privilege Level | 16466
Release Information | 16467

Syntax

```
security-intelligence {
    add-attacker-ip-to-feed feed-name;
    add-target-ip-to-feed feed-name;
}
```

Hierarchy Level

```
[edit security idp idp-policy (Security)  name rulebase-ips name rule (Security IPS Rulebase) then
(Security IDP Policy) application-services security-intelligence]
```

Description

Generate security intelligence feeds, you can configure the IDP rule with threat profiles to define the different types of feeds. You can configure attacker IP feed and target IP feed.

To allow the SRX to generate, leverage, and propagate its own threat-intelligence feeds (custom as well as infected-hosts) based off of detection events from IDP.

Options

add-attacker-ip-to-feed	Specify the desired feed-name (attacker-ip-to feed).
add-target-ip-to-feed	Specify the desired feed-name (target-ip-to-feed)

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

show security idp counters ips | 24537

security-intelligence-policy

IN THIS SECTION

- Syntax | 16467
- Hierarchy Level | 16468
- Description | 16468
- Options | 16468
- Required Privilege Level | 16469
- Release Information | 16469

Syntax

```
security-intelligence-policy {
    geo-ip
    threat-level threat-level {
        policy-map policy-name
            };
    threat-action {
        drop;
        drop-and-log;
        drop-and-sample;
        drop-log-and-sample;
        log;
        log-and-sample;
        sample;
```

```
    }
    white-list;
    black-list;
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name]
[edit services web-filter profile profile-name url-filter-template template-name]
```

Description

Define the threat level and action for the Web filter profile. The packets are redirected at the Packet Forwarding Engine based on the configured threat-level action associated with the threat-level of the destination IP address.

Options

threat-level Define the Web filtering threat level. The value ranges from 1 through 10

policy-map Map the threat level to a user-attribute defined by the url-filterd.

threat-action Define the way the Packet Forwarding Engine processes packets in response to a threat. Only one action can be configured for each threat level that is defined. The default threat-action is `accept`.

- `drop`—Drop the packets and do not generate a log message.
- `drop-and-log`—Drop the packets and generate a log message.
- `drop-and-sample`—Drop and sample the packets.
- `drop-log-and-sample`—Drop and sample the packets, and generate a log message.
- `log`—Allow the packets and generate a log message.
- `log-and-sample`—Allow, sample the packets, and generate a log message.
- `sample`—Sample the packets.

white-list Allow the IP addresses configured either as a file or as an IP address-list .

black-list Block the IP addresses configured either as a file or as an IP address-list .

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1 on MX Series routers with Juniper Advanced Threat Prevention Cloud (Juniper ATP Cloud) .

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480, and MX960 with the MX-SPC3 services card

RELATED DOCUMENTATION

web-filter | 31062

security-intelligence(services)

IN THIS SECTION

- Syntax | 16470
- Hierarchy Level | 16473
- Description | 16473
- Options | 16474
- Required Privilege Level | 16479
- Release Information | 16479

Syntax

```
security-intelligence {
    authentication {
        auth-token auth-token;
        tls-profile tls-profile;
    }
    category {
        all;
        category-name {
            Blacklist {
                disable;
            }
            CC {
                disable;
            }
            DNS {
                disable;
            }
            GeoIP {
                disable;
            }
            IPFilter {
                disable;
            }
            Infected-Hosts {
                disable;
            }
            SecProfiling {
                disable;
            }
            Reverse Shell {
                disable;
            }
            Whitelist {
                disable;
            }
        }
    }
    default-policy {
        CC {
            feed-cc-log-only;
```

```
        }
        DNS {
            feed-cc-log-only;
        }
        GeoIP {
            feed-cc-log-only;
        }
        IPFilter {
            feed-cc-log-only;
        }
        Infected-Hosts {
            feed-cc-log-only;
        }
    }
    global-disable-feed {
        all;
        feed-name (cc_ip_dat | cc_ipv6_dat | cc_url_dat))
    }
    policy policy-name category-profile-name{
        CC {
            feed-cc-log-only;
        }
        DNS {
            feed-cc-log-only;
        }
        GeoIP {
            feed-cc-log-only;
        }
        IPFilter {
            feed-cc-log-only;
        }
        Infected-Hosts {
            feed-cc-log-only;
        }
    }
    profile profile-name {
        category {
            Blacklist;
            CC;
            DNS;
            GeoIP;
            IPFilter;
            Infected-Hosts;
```

```
            Whitelist;
        }
        default rule {
            then {
                action {
                    block {
                        close {
                            http {
                                file file-name;
                                message block-message;
                                redirect-url redirect-url;
                            }
                        drop;
                    }
                    recommended;
                }
                log;
                no-log;
            }
        }
        description text;
        rule rule-name {
            match
                feed-name feed-name;
                threat-level (1| 2| 3| 4 |5 | 6 |7 |8 | 9 |10);
            }
            then
                action {
                    block {
                        close {
                            http {
                                file file-name;
                                message block-message;
                                redirect-url redirect-url;
                            }
                        drop;
                    }
                    permit;
                    recommended;
                    sinkhole;
                }
                log;
            }
```

```
        }
    }
    proxy-profile proxy-profile-name;
    traceoptions {
        file {
            filename;
            files number-of-trace-files;
            match match-expression;
            no-world-readable
            size maximum-trace-file-size;
            world-readable
        }
        flag (all | config | connect | dns | filter | gencfg | normal | operational commands |
parse | routing | snmp | statistics | system | timer);
        level (all | error | info | notice | verbose | warning);
         no-remote-trace;
    }
    url url-address;
    url-parameter url-parameter;
}
```

Hierarchy Level

```
[edit services]
```

Description

Using this command, you can configure security intelligence profiles and policies to work with security intelligence feeds, such as infected hosts and C&C. You then configure a firewall policy to include the security intelligence policy, for example, block outgoing requests to a C&C host.

A security intelligence rule can have multiple feed names (`feed-name`) with multiple threat levels. Specifying the threat level is required, but `feed-name` is optional. Juniper ATP Cloud makes sure there is no duplicate feed-name associated with threat levels configured in the same profile. Juniper ATP Cloud uses the following approach:

- If `feed-name` is configured, it looks up the feed-name first.
- If no `feed-name` configured or the `feed-name` is not match, it uses the threat level rules.
- If no rules are present or match, the profile's default rule is used.

Options

authentication

Configure authentication, such as an auth token or TLS profile, to commute with the feed server. This operation is performed by the ops script used to enroll your devices and is typically not required afterwards. If you have problems establishing a connection with the Juniper ATP cloud server, it is recommended that you rerun the ops script instead of manually entering all the CLI commands.

category (all | *category-name*)

Category to be disabled. You can disable a specific category or all. This option is used for temporarily disabling a category during debugging phases.

disable-global-feed (all | feed-name (CC_IP |CC_URL))

Disable the Juniper C&C and URL feed to free the resources on SRX Series Firewalls. The resources are then available for loading custom feeds. The available options are all, CC_IP or CC_URL.

- all—Applies to all feeds.
- CC_IP—Applies to global CC IP feed.
- CC_URL—Applies to global CC URL feed.

policy*policy-name category-profile-name*

Configure the security intelligence policy. You specify the category (such as `CC`) and the security intelligence profile to associate with this policy.

profile *profile-name category* rule *rule-name* (match | then)

Configure security intelligence profile. You specify the profile name, the category (such as CC), and any rules and actions. The actions that you can perform are listed below:

- block—Block with close or drop action
- permit—Permit action
- recommended— Recommended action from feed server
- sinkhole—DNS sinkhole action for malicious DNS queries

> **NOTE**: The action `sinkhole` must be configured only for DNS profile category.

traceoptions

Set security intelligence trace options.

url *url-address*

Configure the URL of the feed server. This operation is performed by the ops script used to enroll your devices and is typically not required afterwards. If you have problems establishing a connection with the Juniper ATP Cloud server, it is

recommended that you rerun the ops script instead of manually entering all the CLI commands.

url-parameter *url-parameter*

This is an internal option. Do not use this option unless instructed to by Juniper Networks Technical Support.

block close infected host file message| redirect-URL

Provides HTTP URL redirection based on infected hosts with the block action.

This allows for administrator notification of Infected Hosts. During the processing of a session IP address, if the IP address is on the infected hosts list and HTTP traffic is using ports 80 or 8080, infected hosts HTTP redirection can be done. If HTTP traffic is using dynamic ports, HTTP traffic redirection cannot be done.

set services security-intelligence profile secintel_profile rule

This example performs feed name-based URL redirection.

```
user@host# set services security-intelligence profile secintel_profile rule secintel_rule1 match
feed-name custom_feed1
uuser@host# set services security-intelligence profile secintel_profile rule secintel_rule1
match threat-level 7
user@host# set services security-intelligence profile secintel_profile rule secintel_rule1 match
threat-level 8
user@host# set services security-intelligence profile secintel_profile rule secintel_rule1 match
threat-level 9
user@host# set services security-intelligence profile secintel_profile rule secintel_rule1 match
threat-level 10
user@host# set services security-intelligence profile secintel_profile rule secintel_rule1 then
action block close http redirect-url http://www.test.com/url1.html

user@host# set services security-intelligence profile secintel_profile rule secintel_rule2 match
feed-name custom_feed2
uuser@host# set services security-intelligence profile secintel_profile rule secintel_rule2
match threat-level 7
user@host# set services security-intelligence profile secintel_profile rule secintel_rule2 match
threat-level 8
user@host# set services security-intelligence profile secintel_profile rule secintel_rule2 match
threat-level 9
user@host# set services security-intelligence profile secintel_profile rule secintel_rule2 match
threat-level 10
user@host# set services security-intelligence profile secintel_profile rule secintel_rule2 then
```

```
action block close http redirect-url http://www.test.com/url2.html
```

set services security-intelligence profile

```
user@host# set services security-intelligence profile cc_profile category CC
```

set services security-intelligence profile

This example configures a profile name, a profile rule and the threat level scores. Anything that matches these scores is considered malware or an infected host.

```
user@host# set services security-intelligence profile cc_profile rule CC_rule match threat-level
[8 9 10]
```

set services advanced-anti-malware connection authentication

This example defines the TLS profile, typically done by the ops script when enrolling devices.

```
user@host# set services advanced-anti-malware connection authentication tls-profile aamw-ssl
```

set services security-intelligence url

This example defines the feed server URL, typically done by the ops script when enrolling devices.

```
user@host# set services security-intelligence url https://cloudfeeds.argon.junipersecurity.net/
api/manifest.xml
```

set services security-intelligence profile secintel_profile rule secintel_rule2 then action block close infected host|http redirect-url http://www.test.com/url2.html

User notification of infected hosts—(Starting in Junos 18.1R1) This command allows you to configure HTTP URL redirection based on infected hosts with the block action. During the processing of a session IP address, if the IP address is on the infected hosts list and HTTP traffic is using ports 80 or 8080, infected hosts HTTP redirection to a specified URL can be used in conjunction with the block action. This allows administrators to receive a notification of the block action. Note that if HTTP traffic is using dynamic ports, HTTP traffic redirection cannot be done

The syntax for the command is as follows:

Syntax: set services security-intelligence profile <name> then action block close <file message|redirect-URL>

For example:

```
user@host# set services security-intelligence profile secintel_profile rule secintel_rule2 then
action block close infected host|http redirect-url http://www.test.com/url2.html
```

To view the HTTP URL redirection counter, type **show services security-intelligence statistics**

set services security-intelligence disable-global-feed all

This example disables the global feed in all SecIntel configurations.

```
user@host# set services security-intelligence disable-global-feed all
```

set services security-intelligence disable-global-feed feed-name CC_IP

This example disables the global CC IP feeds in SecIntel configurations.

```
user@host# set services security-intelligence disable-global-feed feed-name CC_IP
```

set services security-intelligence disable-global-feed feed-name CC_URL

This example disables the global CC URL feeds in SecIntel configurations.

```
user@host# set services security-intelligence disable-global-feed feed-name CC_URL
```

All the security intelligence services are enabled by default. To disable the services, you must use the configuration commands:

- **set services security-intelligence category category-name GeoIP disable**

  This example disables the GeoIP feeds in SecIntel configurations.

  ```
  user@host# set services security-intelligence category category-name GeoIP disable
  ```

- **set services security-intelligence category category-name Whitelist disable**

This example disables allowlist category in SecIntel configurations.

```
user@host# set services security-intelligence category category-name Whitelist disable
```

- **set services security-intelligence category category-name Blacklist disable**

  This example disables blocklist category in SecIntel configurations.

```
user@host# set services security-intelligence category category-name Blacklist disable
```

- **set services security-intelligence category category-name IPFilter disable**

  This example disables IPFilter category in SecIntel configurations.

```
user@host# set services security-intelligence category category-name IPFilter disable
```

Use the `show services security-intelligence` CLI command to verify your settings. Your output should look similar to the following:

```
[edit]
user@host# show services security-intelligence
category {
    category-name {
        Whitelist {
            disable;
        }
        IPFilter {
            disable;
        }
        GeoIP {
            disable;
        }
        Blacklist {
            disable;
        }
    }
}
policy p1 {
```

```
    IPFilter; ## Warning: missing mandatory statement(s): <profile-name>
}
```

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS 12.1X46. Starting with Junos OS 15.1X49-D110, this command adds the `feed-name` option which can be used in security intelligence rules. Prior to Junos OS 15.1X49-D100 you could perform HTTP URL redirect based on threat levels. With `feed-name`, you can now perform HTTP URL redirection based on a feed name.

User notification of infected hosts—As of Junos OS 18.1R1, there is support HTTP URL redirection based on infected hosts with the block action. This allows for administrator notification of Infected Hosts. During the processing of a session IP address, if the IP address is on the infected hosts list and HTTP traffic is using ports 80 or 8080, infected hosts HTTP redirection can be done. If HTTP traffic is using dynamic ports, HTTP traffic redirection cannot be done. See command at bottom of this page.

The `disable-global-feed` option is introduced in Junos OS Release 20.1R1.

The `DNS` option is introduced in Junos OS Release 20.4R1.

The `Reverse-Shell` category is introduced in Junos OS Release 23.3 R1.

RELATED DOCUMENTATION

- Configure DNS Sinkhole
- Configure and Deploy Adaptive Threat Profiling
- Configure the SRX Series Firewall to Block Outbound Requests to a C&C Host
- Configure the SRX Series Firewall to Block Infected Hosts
- Configure Juniper Advanced Threat Prevention Cloud With Geolocation IP
- Configure Reverse Shell Detection
- Configure IPFilter Category
- Example: Configure Juniper Advanced Threat Prevention Cloud Policy

security-level

IN THIS SECTION

- Syntax | 16480
- Hierarchy Level | 16480
- Description | 16480
- Options | 16480
- Required Privilege Level | 16480
- Release Information | 16481

Syntax

```
security-level {
    (default | secure-messages-only);
}
```

Hierarchy Level

```
[edit protocols neighbor-discovery secure]
```

Description

Configure the type of security mode for Secure Neighbor Discovery.

Options

default—Accept and transmit both secure and unsecured messages.

secure-messages-only—Accept secure messages only. Discard unsecured messages.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Secure IPv6 Neighbor Discovery

security-level (Defining Access Privileges)

IN THIS SECTION

- Syntax | 16481
- Hierarchy Level | 16482
- Description | 16482
- Default | 16482
- Options | 16482
- Required Privilege Level | 16482
- Release Information | 16482

Syntax

```
security-level (authentication | none | privacy) {
    notify-view view-name;
    read-view view-name;
    write-view view-name;
}
```

Hierarchy Level

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix context-prefix) security-model (any | usm | v1 | v2c)]
```

Description

Define the security level used for access privileges.

Default

`none`

Options

`authentication`—Provide authentication but no encryption.

`none`—No authentication and no encryption.

`privacy`—Provide authentication and encryption.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

security-level (Generating SNMP Notifications)

IN THIS SECTION

Syntax | 16483

- Hierarchy Level | 16483
- Description | 16483
- Default | 16483
- Options | 16483
- Required Privilege Level | 16483
- Release Information | 16484

Syntax

```
security-level (authentication | none | privacy);
```

Hierarchy Level

```
[edit snmp v3 target-parameters target-parameters-name parameters]
```

Description

Configure the security level to use when generating SNMP notifications.

Default

`none`

Options

`authentication`—Provide authentication but no encryption.

`none`—No authentication and no encryption.

`privacy`—Provide authentication and encryption.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

security-log

IN THIS SECTION

- Syntax | 16484
- Hierarchy Level | 16484
- Description | 16484
- Options | 16485
- Required Privilege Level | 16485
- Release Information | 16485

Syntax

```
security-log {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the security log process.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the security log process.
- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

syslog (System) | 28836

security-log-percent-full

IN THIS SECTION

Syntax | 16486
Hierarchy Level | 16486
Description | 16486
Options | 16486

Required Privilege Level | 16486
Release Information | 16486

Syntax

```
security-log-percent-full percentage;
```

Hierarchy Level

```
[edit security alarms potential-violation]
```

Description

Raise a security alarm when security log exceeds a specified percent of total capacity.

Options

percentage—Percentage of security log capacity at which a security alarm is raised.

- **Range:** 0 through 100 percent

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

syslog (System) | 28836

security-metadata-streaming

IN THIS SECTION

- Syntax | 16487
- Hierarchy Level | 16489
- Description | 16489
- Options | 16489
- Required Privilege Level | 16492
- Release Information | 16492

Syntax

```
security-metadata-streaming {
   dns-cache {
              custom-list [benign <domin-name> | c2 <domain-name>];
              }
    policy policy-name {
        dns {
            cache {
                ttl {
                    benign value;
                    c2 value;
                }
            }
            detections {
                all {
                    action [deny | permit | sinkhole];
                    notification [log |log-detections];
                    fallback-options {
                        notification {
                            log;
                        }
                    }
                }
                dga {
                    action [deny | permit | sinkhole];
```

```
                    verdict-timeout value;
                    notification [log | log-detections];
                    fallback-options {
                        notification {
                            log;
                        }
                    }
                }
                tunneling {
                    action [deny | permit | sinkhole];
                    notification [log | log-detections];
                    inspection-depth value;
                    fallback-options {
                        notification {
                            log;
                        }
                    }
                }
            }
        dynamic-filter;
        http {
        detections {
            encryptedc2 {
                action permit;
                notification log/log detections;
            }
            all {
                action permit;
                notification log/log detections;
                fallback-options {
                    notification {
                        log;
                    }
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure security metadata streaming policy on SRX Series Firewalls to send the metadata and connection patterns of a network traffic to Juniper Networks ATP Cloud for encrypted traffic insights. After configuring the security metadata streaming policy, attach it to the security policy at zone-level.

`set security policies from-zone` *`from-zone`* `to-zone` *`to-zone`* `application-services security-metadata-streaming-policy` *`dns-policy`*

Options

dns-cache Configure a list of static benign and command-and-control (C2) domains in the Domain Name System (DNS) cache to take immediate action on configured domains. Only wildcard domains are allowed. The domain format must be `*.domain_name.domain_ending` . The entries configured in DNS Cache via CLI will remain in the DNS Cache until that configuration is deleted from the device. You can configure a maximum of 500 domains each in benign list and c2 list.

- By default, the action for traffic originating from a benign (allowlisted) domain is `permit`.
- The action for traffic originating from a C2 (blocklisted) domain is based on the action configured under DNS detections.

policy ***policy-name*** Configure the security-metadata-streaming policy.

dns Configure DNS options.

cache Store DNS in cache till time-to-live (TTL). The TTL provided by SRX Series Firewall overrides Juniper ATP Cloud provided TTL.

> **NOTE**: You must configure at least one DNS detection method to configure DNS cache.

- benign—(Optional) Set benign TTL value. The range is 60 to 172800 seconds. Default value is 86400.
- c2—(Optional) Set C2 TTL value. The range is 60 to 172800 seconds. Default value is 86400.

detections Configure the detection type for DNS requests. The available options are all, dga, and tunneling. You can configure any of the following detections.

- all detections
- both dga and tunneling detection
- either dga or tunneling detection

You cannot configure all detections and custom detections(dga and/or tunneling) together. The detections are mutually exclusive.

> **NOTE**: Each detection method has a fallback option which is used in case nothing is detected within a certain number of packets (in case of tunneling) or within a certain time period (in case of DGA).

all Configure all detections.

- action—Specify the action the SRX device will take when a detection is made. The available options are `deny`, `permit`, or `sinkhole`.
- fallback-options—Fallback options for DNS detections. The fallback action is triggered when DNS-based attacks are not detected (DGA verdict is not received within 100ms (default value of verdict-timeout) and DNS tunnel is not detected within 4 packets (default value of inspection-depth)). The available option is to `Log` the DNS requests.
- notification—Global notification action taken for DNS detection methods. The available options are:
  - log—Generates log for DNS requests and DNS detections.
  - log-detections—(recommended) Generates log only for malicious DNS detections.
- verdict-timeout—(Not configurable) Time to wait for a DGA verdict on DNS packet (milliseconds). Default timeout is 100ms for all detections.
- inspection-depth—(Not configurable) Number of packets to be inspected for tunnel detection. Default is 4 packets for all detections.

dga Configure to detect DGA-based attacks on DNS packets.

- action—Specify the action the SRX device will take when a detection is made. The available options are `deny`, `permit`, or `sinkhole`.
- fallback-options—Fallback options for DNS DGA detection. The fallback options are triggered if DGA verdict is not received from Juniper ATP Cloud within the verdict-timeout configured value. The available option is to log the DNS request.
- notification— Notification action taken for DNS DGA detection. The available options are:

- log—Generates log per DNS request and DNS detections.
- log-detections—(recommended) Generates log only for malicious DNS detections.

- verdict-timeout—(Optional) Time to wait for a verdict on DNS Packet (milliseconds). The range is 50 to 500. Default timeout is 100ms.

tunneling Configure to detect DNS tunneling.

- action—Specify the action the SRX device will take when a detection is made. The available options are `deny` (drops tunnel session), `permit` (permits tunnel session), or `sinkhole` (drops the tunnel session and sinkholes the domain).
- fallback-options—Fallback options for DNS tunneling detection. The fallback options are triggered if a tunnel is not detected within the specified number of packets (inspections-depth). The available option is to log the DNS request.
- inspection-depth—(Optional) Number of packets to be inspected for tunnel detection. The range is 0 to 10. Default is 4 packets. 0 indicates forever.
- notification—Notification action taken for DNS tunneling detection. The available options are:
  - log—Generates log per DNS request and DNS detections.
  - log-detections—(recommended) Generates log only for malicious DNS detections.

dynamic-filter Configure dynamic filtering options for security metadata streaming policy on SRX Series Firewalls.

http Configure HTTP options.

- detections—Configure the detection type for HTTP requests. The available options are all and encryptedc2. You can configure any of the following detections:
  - all detections—All detections can only be configured if no other detection is configured.
  - encryptedc2—Encryptedc2 can be configured to detect Command and Control (C&C) communication.
- action—Defines the action taken on the traffic. The default action is permit.
- notification—Defines the notification action taken for the traffic. The available options are:
  - log—Generates log per HTTP request and detections.

- log-detections—(recommended) Generates log only for malicious detections.

- fallback-options—Fallback options for HTTP traffic. The fallback options are triggered if HTTP traffic is not detected. The available option is to log the security metadata streaming actions.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1 on SRX Series Firewalls with Juniper Advanced Threat Prevention Cloud (Juniper ATP Cloud).

RELATED DOCUMENTATION

application-services (security-metadata-streaming)

show services security-metadata-streaming

security-mode (SRX)

IN THIS SECTION

- Syntax | 16493
- Hierarchy Level | 16493
- Description | 16493
- Options | 16493
- Required Privilege Level | 16494
- Release Information | 16494

Syntax

```
security-mode (dynamic | static-cak | static-sak);
```

Hierarchy Level

```
[edit security macsec connectivity-association]
```

Description

Configure the MACsec security mode for the connectivity association.

We recommend enabling MACsec on switch-to-switch Ethernet links using static connectivity association key (CAK) security mode. Static CAK security mode ensures security by frequently refreshing to a new random secure association key (SAK) and by only sharing the SAK between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available when you enable MACsec using static CAK security mode.

Options

security-mode Specifies the MACsec security mode. Options include:

- `dynamic`—Dynamic mode.

  Dynamic security mode is used to enable MACsec on switch-to-host Ethernet links. In dynamic mode, a master key is retrieved from a RADIUS server by a switch and a host as part of the AAA handshake in separate transactions. The MKA protocol is enabled when the master key is exchanged between the switch and the host.

- `static-cak` —Static connectivity association key (CAK) mode.

  Static CAK security mode is used to enable MACsec on switch-to-switch Ethernet links. In `static-cak` mode, the switch at one end of the point-to-point link acts as the key server and regularly transmits a randomized key using a process that does not transmit any traffic outside of the MACsec-secured point-to-point link.

- `static-sak` —Static secure association key (SAK) mode.

Static SAK security mode is used to enable MACsec on switch-to-switch Ethernet links. In static-sak mode, one of two user-configured security keys is used to secure the point-to-point link. The two security keys are regularly rotated.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

security-mode

IN THIS SECTION

- Syntax | **16495**
- Hierarchy Level | **16495**
- Description | **16495**
- Options | **16495**
- Required Privilege Level | **16495**
- Release Information | **16496**

Syntax

```
security-mode security-mode;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name]
```

Description

Configure the MACsec security mode for the connectivity association.

To enable MACsec on a switch-to-host link, you must use dynamic connectivity association key (CAK) mode. For links that connect switches or routers—switch-to-switch, switch-to-router, or router-to-router—you can use static CAK mode or dynamic CAK mode.

NOTE: Dynamic CAK mode is not supported on logical interfaces.

Options

security-mode Specifies the MACsec security mode. Options include:

- `dynamic`—Dynamic CAK mode.

  You can use dynamic CAK mode on switch-to-host links as well as links that connect switches or routers. Dynamic CAK mode relies on 802.1X authentication with EAP-TLS. Public key infrastructure is also required for certificate validation.

- `static-cak`—Static CAK mode.

  You can use static CAK mode on links that connect switching devices. In `static-cak` mode, the CAK must be manually configured on both peer nodes of the point-to-point MACsec link.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D15.

The `dynamic` security mode option was introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Configuring MACsec on EX, QFX and SRX Devices

security-model (Access Privileges)

IN THIS SECTION

- Syntax | 16496
- Hierarchy Level | 16496
- Description | 16497
- Options | 16497
- Required Privilege Level | 16497
- Release Information | 16497

Syntax

```
security-model (usm | v1 | v2c);
```

Hierarchy Level

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix context-prefix)]
```

Description

Configure the security model for an SNMPv3 group. The security model is used to determine access privileges for the group.

Options

`usm`—SNMPv3 security model.

`v1`—SNMPv1 security model.

`v2c`—SNMPv2c security model.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

security-model (Group)

IN THIS SECTION

- Syntax | 16498
- Hierarchy Level | 16498
- Description | 16498
- Options | 16498
- Required Privilege Level | 16498
- Release Information | 16498

Syntax

```
security-model (usm | v1 | v2c) {
    security-name security-name {
        group group-name;
    }
}
```

Hierarchy Level

```
[edit snmp v3 vacm security-to-group]
```

Description

Define a security model for a group.

Options

`usm`—SNMPv3 security model.

`v1`—SNMPv1 security model.

`v2c`—SNMPv2c security model.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

security-model (SNMP Notifications)

IN THIS SECTION

- Syntax | 16499
- Hierarchy Level | 16499
- Description | 16499
- Options | 16499
- Required Privilege Level | 16499
- Release Information | 16500

Syntax

```
security-model (usm | v1 | v2c);
```

Hierarchy Level

```
[edit snmp v3 target-parameters target-parameters-name parameters]
```

Description

Configure the security model for an SNMPv3 group. The security model is used for SNMP notifications.

Options

`usm`—SNMPv3 security model.

`v1`—SNMPv1 security model.

`v2c`—SNMPv2c security model.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

security-name (Community String)

IN THIS SECTION

- Syntax | 16500
- Hierarchy Level | 16500
- Description | 16500
- Options | 16501
- Required Privilege Level | 16501
- Release Information | 16501

Syntax

```
security-name security-name;
```

Hierarchy Level

```
[edit snmp v3 snmp-community community-index]
```

Description

Associate a community string with the security name of a user. The community string, which is used for SNMPv1 and SNMPv2c clients in an SNMPv3 system, is configured at the `[edit snmp v3 snmp-community community-index]` hierarchy level.

Options

security-name—Name that is used for messaging security and user access control.

> **NOTE**: The security name must match the configured security name at the `[edit snmp v3 target-parameters target-parameters-name parameters]` hierarchy level when you configure traps or informs.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

security-name (Security Group)

IN THIS SECTION

- Syntax | 16502
- Hierarchy Level | 16502
- Description | 16502
- Options | 16502
- Required Privilege Level | 16502
- Release Information | 16502

Syntax

```
security-name security-name {
    group group-name;
}
```

Hierarchy Level

```
[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c)]
```

Description

Associate a group or a community string with a configured security group.

Options

security-name—Username configured at the `[edit snmp v3 usm local-engine user username]` hierarchy level. For SNMPv1 and SNMPv2c, the security name is the community string configured at the `[edit snmp v3 snmp-community community-index]` hierarchy level.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

security-name (SNMP Notifications)

IN THIS SECTION

- Syntax | 16503
- Hierarchy Level | 16503
- Description | 16503
- Options | 16503
- Required Privilege Level | 16504
- Release Information | 16504

Syntax

```
security-name security-name;
```

Hierarchy Level

```
[edit snmp v3 target-parameters target-parameters-name parameters]
```

Description

Configure the security name used when generating SNMP notifications.

Options

security-name—If the SNMPv3 USM security model is used, identify the user when generating the SNMP notification. If the v1 or v2c security models are used, identify the SNMP community used when generating the notification.

NOTE: The access privileges for the group associated with this security name must allow this notification to be sent.

If you are using the v1 or v2 security models, the security name at the [edit snmp v3 vacm security-to-group] hierarchy level must match the security name at the [edit snmp v3 snmp-community *community-index*] hierarchy level.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

security-option (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 16504
- Hierarchy Level | 16504
- Description | 16505
- Required Privilege Level | 16505
- Release Information | 16505

Syntax

```
security-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop IPv4 packets that have the IP option of 2 (Security).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

security-package

IN THIS SECTION

- Syntax | 16505
- Hierarchy Level | 16506
- Description | 16506
- Options | 16507
- Required Privilege Level | 16507
- Release Information | 16507

Syntax

```
security-package {
    automatic (Security) {
        download-timeout minutes;
```

```
        enable;
        interval hours;
        start-time start-time;
    }
    install {
        ignore-version-check;
        ignore-appid-failure;
    }
    proxy-profile proxy-profile;
    source-address address;
    url url-name;
}
```

Hierarchy Level

```
[edit security idp]
```

Description

Configure the device to automatically download the updated signature database from the specified URL.

When you configure signature installation to enable the `ignore-appid-failure` option, IDP signature download/installation does not fail even if application identification download/installation fails during IDP signature download/installation. This option is not enabled by default. You have to enable this option.

IDP signature package on an external server can be downloaded and installed on the SRX Series Firewall. Configure the `proxy profile` option of security package download to connect to the external server through a specified proxy server.

IDP uses proxy profile configured at the system level. The proxy profile being used in the security package must be configured at the `[edit services proxy]` hierarchy.

You can configure multiple proxy profiles under `[edit services proxy]` hierarchy. IDP can utilize only one proxy profile. Multiple proxy profiles are not supported for use under IDP simultaneously. When a proxy profile is configured under `[security idp security-package]` hierarchy, then the idpd process connects to the proxy host instead of the signature pack download server. The proxy host then communicates with the download server and provides the response back to the idpd process. The idpd process is notified every time there is a change made at the `[edit services proxy]` hierarchy.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Option `ignore-appid-failure` is introduced in Junos OS Release 18.3R1.

Option `proxy-profile` is introduced in Junos OS Release 18.3R1.

security-profile

IN THIS SECTION

- Syntax | 16507
- Hierarchy Level | 16511
- Description | 16511
- Options | 16511
- Required Privilege Level | 16511
- Release Information | 16511

Syntax

```
security-profile {
    profile name {
        address-book (System) {
            maximum maximum;
```

```
            reserved reserved;
        }
        advanced-anti-malware-policy {
            maximum maximum;
            reserved reserved;
        }
        auth-entry {
            maximum maximum;
            reserved reserved;
        }
        cpu {
            reserved percent;
        }
        dslite-softwire-initiator {
            maximum maximum;
            reserved reserved;
        }
        flow-gate {
            maximum maximum;
            reserved reserved;
        }
        flow-session {
            maximum maximum;
            reserved reserved;
        }
        icap-redirect-profile {
            maximum maximum;
            reserved reserved;
        }
        idp-policy idp-policy;
        logical-system (System Security Profile) logical-system;
        nat-cone-binding {
            maximum maximum;
            reserved reserved;
        }
        nat-destination-pool {
            maximum maximum;
            reserved reserved;
        }
        nat-destination-rule {
            maximum maximum;
            reserved reserved;
        }
```

```
        nat-interface-port-ol (System) {
            maximum maximum;
            reserved reserved;
        }
        nat-nopat-address {
            maximum maximum;
            reserved reserved;
        }
        nat-pat-address {
            maximum maximum;
            reserved reserved;
        }
        nat-pat-portnum {
            maximum maximum;
            reserved reserved;
        }
        nat-port-ol-ipnumber {
            maximum maximum;
            reserved reserved;
        }
        nat-rule-referenced-prefix (System) {
            maximum maximum;
            reserved reserved;
        }
        nat-source-pool {
            maximum maximum;
            reserved reserved;
        }
        nat-source-rule {
            maximum maximum;
            reserved reserved;
        }
        nat-static-rule {
            maximum maximum;
            reserved reserved;
        }
        policy (System Security Profile) {
            maximum maximum;
            reserved reserved;
        }
        policy-with-count {
            maximum maximum;
            reserved reserved;
```

```
        }
        root-logical-system;
        scheduler (System Security Profile) {
            maximum maximum;
            reserved reserved;
        }
        secintel-policy {
            maximum maximum;
            reserved reserved;
        }
        secure-wire {
            maximum maximum;
            reserved reserved;
        }
        security-log-stream-number {
            maximum maximum;
            reserved reserved;
        }
        tenant tenant;
        user-auth-entry {
            maximum maximum;
            reserved reserved;
        }
        vrf-group {
            maximum maximum;
            reserved reserved;
        }
        zone (System Security Profile) {
            maximum maximum;
            reserved reserved;
        }
    }
    resources {
        cpu-control;
        cpu-control-target percent;
    }
}
```

Hierarchy Level

```
[edit system],
[edit tenants tenant-name]
```

Description

Create a security profile and specify the kinds and amounts of resources to allocate to a logical system to which the security profile is bound.

As a primary administrator, you can create a security profile and bind it to more than one logical system if you want to allocate the same kinds and amounts of resources to them. For details on how many security profiles you can create, see *Understanding Logical Systems Security Profiles (Primary Administrators Only)*. When you reach the limit, you must delete a security profile and commit the configuration before you can create and commit the configuration for another security profile.

Only the primary administrator can create security profiles.

Options

profile-name	Name of the security profile.
resources	Name of the resources.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The `dslite-softwire-initiator` option introduced in Junos OS Release 12.1.

The `security-profile` option added under the `tenants` hierarchy in Junos OS Release 18.3R1.

The icap redirect profile option is introduced in Junos OS Release 18.3R1.

`secure-wire` option introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only)

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only)

security-profile-resources

IN THIS SECTION

- Syntax | 16512
- Hierarchy Level | 16512
- Description | 16512
- Options | 16513
- Required Privilege Level | 16513
- Release Information | 16513

Syntax

```
security-profile-resources {
    cpu-control;
    cpu-control-target percent;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure global settings that apply to all logical systems in the device.

Options

`cpu-control`—Enable CPU utilization control.

`cpu-control-target` *percent*—Specify the upper limit for CPU utilization on the device under normal operating conditions.

- **Range:** 0 through 100 percent (decimal point allowed).
- **Default:** 80 percent.

> **NOTE**: The `cpu-control` option must be specified for the `cpu-control-target` value to take effect.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only)

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only)

security-service (Security Forwarding Options)

IN THIS SECTION

- Syntax | 16514
- Hierarchy Level | 16514
- Description | 16514

- Options | 16514
- Required Privilege Level | 16515
- Release Information | 16515

Syntax

```
security-service (fail-open);
```

Hierarchy Level

```
[edit security forwarding-options]
```

Description

The system resource management guarantees the resources are used according to priorities. the fail-open/fail-close impacts the new session creation only when the system resource is busy.

- If system resource is not busy, fail-open/fail-close won't take effect. No impact to traffic.
- If system resource is busy and the new session need not be created with Layer 7 services that require the system resource, then no impact to traffic.
- . If system resource is busy, and the new session needs be created with Layer 7 services that require the system resource:
  - Fail-close (default), drops the packet and won't create the session
  - Fail-open, creates a flow session without the Layer 7 services that require the resource, and forward the packet out.

Options

- `fail-open`—Ignores Layer 7 services with resource requirements, creates a flow session without Layer 7 services, and forward the packet out.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Support on SRX Series Firewalls for flow based security-service in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Traffic Processing on SRX Series Devices Overview

security-to-group

IN THIS SECTION

- Syntax | 16515
- Hierarchy Level | 16516
- Description | 16516
- Required Privilege Level | 16516
- Release Information | 16516

Syntax

```
security-to-group {
    security-model (usm | v1 | v2c) {
        group group-name;
        security-name security-name;
    }
}
```

Hierarchy Level

```
[edit snmp v3 vacm]
```

Description

Configure the group to which a specific SNMPv3 security name belongs. The security name is used for messaging security.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

security-zone

IN THIS SECTION

- Syntax | 16517
- Hierarchy Level | 16518
- Description | 16518
- Options | 16518
- Required Privilege Level | 16518
- Release Information | 16518

Syntax

```
security-zone zone-name {
    address-book {
        address address-name {
            ip-prefix {
                description text;
            }
            description text;
            dns-name domain-name {
                ipv4-only;
                ipv6-only;
            }
            range-address lower-limit to upper-limit;
            wildcard-address ipv4-address/wildcard-mask;
        }
        address-set address-set-name {
            address address-name;
            address-set address-set-name;
            description text;
        }
    }
    advance-policy-based-routing;
    application-tracking;
    description text;
    enable-reverse-reroute;
    host-inbound-traffic {
        protocols protocol-name {
                            except;
        }
        system-services service-name {
                            except;
        }
    }
    interfaces interface-name {
        host-inbound-traffic {
            protocols protocol-name {
                                except;
            }
            system-services service-name {
                                except;
            }
```

```
        }
    }
    screen screen-name;
    source-identity-log;
    tcp-rst;
    unidirectional-session-refreshing;
    white-list white-list name;
}
```

Hierarchy Level

```
[edit security zones]
```

Description

Define a security zone, which allows you to divide the network into different segments and apply different security options to each segment.

Options

zone-name —Name of the security zone.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Support for wildcard addresses added in Junos OS Release 11.1.

The `description` option added in Junos OS Release 12.1.

The `unidirectional-seesion-refreshing` option added in Junos OS Release 20.4R1.

The `white-list` option added in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Security Zones Overview

Example: Configuring Application Firewall Rule Sets Within a Security Policy

segment

IN THIS SECTION

- Syntax | 16519
- Hierarchy Level | 16519
- Description | 16520
- Options | 16520
- Required Privilege Level | 16520
- Release Information | 16520

Syntax

```
segment {
    (pop | swap swap);
    description description;
    next-hop next-hop;
    sid-label;
}
```

Hierarchy Level

```
[edit logical-systems name protocols mpls static-label-switched-path],
[edit logical-systems name routing-instances name protocols mpls static-label-switched-path],
[edit protocols mpls static-label-switched-path],
[edit routing-instances name protocols mpls static-label-switched-path]
```

Description

Static segment for segment routing. A static segment is identified by a unique name. This segment type is assigned a segment identifier (SID) which falls under a default range of 100000 through 1048575. The segment has label operation such as pop-and-forward for adjacency segment and swap-and-forward for prefix or node segment. For both types of label operation, the segment is assigned a next hop that specifies the remote IP address if the outgoing interface is a multi-access interface, or the name of the outgoing interface if the interface is a point-to-point interface. Static segment configuration is used to statically configure or provision the adjacency SIDs, node SIDs, and prefix SIDs on transit routers.

Options

pop	Pop the SID label
swap	Swap the SID label to this label
description	Text description of label-switched path
next-hop	IPv4 address or interface of next-hop router
sid-label	Segment identifier (SID) label

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release OS 18.1R1 for MX Series, PTX Series, and QFX Series.

RELATED DOCUMENTATION

Static Segment Routing Label Switched Path

segment-list | **16521**

source-routing-path | **28316**

segment-list

IN THIS SECTION

- Syntax | 16521
- Hierarchy Level | 16521
- Description | 16522
- Options | 16522
- Required Privilege Level | 16524
- Release Information | 16524

Syntax

```
segment-list name {
    hop-name {
        (loose | strict);
        ip_address IP address;
        label number ;
        label-type node
        srv6-sid;    }
    auto-translate {
         protected mandatory;
         unprotected mandatory;
    }
    dynamic;
    compute;
    inherit-label-nexthops;
    srv6;
}
```

Hierarchy Level

```
[edit logical-systems name protocols source-packet-routing],
[edit protocols source-packet-routing]
```

Description

Specify an name to identify the segment routing list (used in traffic engineering policy) and the explicit path for source routing label switched path (LSPs) to traverse through traffic engineering segments. The segment list is essentially a stack of segment identifiers.

Starting in Junos OS release 19.1R1 for MX and PTX Series routers, you can enable a translation service to translate next-hop IP addresses into the corresponding segment identifier (SID) labels. The translation service keeps track of the node reached at each hop.

When configured, the `segment-list` of a segment routing traffic engineering (SR-TE) LSP accepts IP addresses for all the hops along the path. These IP addresses can be either the loopback address of a node, or the IP address of a link, as identified by the `node-type`. When `auto-translation` is enabled, next hop IP addresses are automatically translated to corresponding SIDs using the translation service. A retry rate can be set for the retry timer at the `source-packet-routing` hierarchy level.

> **NOTE**: The segment list enables BGP and static segment routing LSP to steer traffic based on segment routing policies. When a segment list is used by the protocol BGP, the BGP protocol validates these segment identifiers and selects valid segments for traffic engineering.

Options

name Specify a name to identify a SR-TE segment-list.

> **NOTE**: The combined length of tunnel-name and LSP name must not exceed 53 characters for per-path telemetry to work.

<hop-name> Indicates the next hop in the segment routing traffic engineering policy (SR-TE).

- **ip-address**—Specify the IP address of the hop. For a segment-list to be used by a non-colored segment routing LSP, the first hop must specify an IP address.
- **label**—Specify the SID label of the hop in a segment routing traffic engineering segment list. In static segment routing LSPs, the source routing path uses the segment list only if the second to Nth hop specifies segment identifiers (SID) labels.

  > **NOTE**: The range is from 0 to 1,048,576 and is applies to BGP and static segment routing LSPs.

- **label-type**—Use with the option below to indicate that the specified address is the IP address of the node, for example, its loopback address, as opposed to that of a link.
  - **node**—Hops that have been specified as `node` are translated to a prefix SID, which can be either a node SID or an anycast SID depending on the type of hop IP address. IP addresses not identified as `node` are consider to be a link.

    NOTE: If the first hop is a `node`, for LSP resolution to work correctly, `inherit-label-nexthops` must be enabled at either `source-packet-routing` hierarchy level, or at the relevant `segment-list` hierarchy level.

- **loose | strict**—IP hops specified using router IDs in the sequence can be strict or loose hops. A strict hop must be directly connected to the previous node in the sequence. A loose hop is not necessarily directly connected to the previous node.

  NOTE: You can specify only router IDs as loose or strict hop constraints. Labels and other IP addresses are not supported as loose or strict hop constraints in Junos OS Release 19.2R1-S1.

- **srv6-sid**—Indicates the SID for SRv6 segment-list in the form of IPv6 address.

  NOTE: An SRv6 segment-list can only have SRv6 SID which is a node SID (either the end-sid or end-x sid of the transit and end-sid of the egress nodes). Therefore, you need to correctly configure the last SRv6 SID in the segment-list as the end-sid of the egress PE.

auto-translate

This option must be enabled before a given segment list can use IP addresses instead of SIDs for any hop other than the first hop. In addition, all hops in the segment list must have IP addresses. If any hops on the list have both an IP address and a label configured, the label will be retained. Link addresses are only translated into labels if the preceding node advertises an adjacency SID for the address (otherwise translation fails).

NOTE: In Junos OS Release 19.1R1, for auto-translate to work for OSPF, RSVP for segment routing must be enabled on all participating interfaces.

- **protected**—(Optional) Enable this option to ensure the adjacency SID is eligible to have a backup path, and that a B-flag is set in adjacency SID advertisements. Note that unless `mandatory` is also selected, the choice succeeds regardless.
  - **mandatory**—(Optional) Enable this option to have translation fail if any *unprotected* links are found in the hop-list.
- **unprotected**—(Optional) Enable this option to ensure that no backup path is calculated for a specific adjacency SID, and that a B-flag is not set in adjacency SID advertisements. Note that unless `mandatory` is also selected, the choice succeeds regardless.
- **mandatory**—(Optional) Enable this option to have translation fail if any *protected* links are found the hop-list.

compute (Optional) Enable use of explicit paths specified in segment list for path computation.

inherit-label-nexthops Inherit label next hops for first hop in this segment list that have both IP address and label configured in the first hop.

You can configure the `inherit-label-nexthops` statement globally or individually for each segment list.

The `inherit-label-nexthops` statement takes effect only when the segment list first hop has both IP address and SID label present.

If the `inherit-label-nexthops` is not configured at the `[edit protocols source-packet-routing segment-list]` hierarchy, and the first hop in the segment list has both IP address and label specified, the default behavior is to use the IP address.

srv6 Specify SRv6 TE segment list.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

`ip-address` statement introduced in Junos OS Release 18.1R1 on MX Series routers.

`inherit-label-nexthops`, `node-type`, and `auto-translate` statements introduced in Junos OS Release 19.1R1 on MX Series routers.

`dynamic` statement introduced in Junos OS Release 19.2R1 on all platforms.

`compute`, `loose`, and `strict` statements introduced in Junos OS Release 19.1R1-S1 on MX Series routers.

`srv6-sid` statement introduced in Junos OS Release 21.3R1.

`srv6` statement introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

Static Adjacency Segment Identifier for ISIS

Static Segment Routing Label Switched Path

Segment Routing Traffic Engineering at BGP Ingress Peer Overview

Understanding Static Segment Routing LSP in MPLS Networks

show spring-traffic-engineering | **26587**

extended-nexthop-color

source-routing-path | **28316**

sr-preference-override

select

IN THIS SECTION

- Syntax | **16526**
- Hierarchy Level | **16526**
- Description | **16526**
- Options | **16526**
- Required Privilege Level | **16526**
- Release Information | **16526**

Syntax

```
select (manual | unconditional);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

Specify the conditions under which the path is selected to carry traffic. The `manual` and `unconditional` options are mutually exclusive.

Options

`manual`—The path is selected for carrying traffic if it is up and stable for at least the revert timer window (potentially before the revert timer has elapsed). Traffic is sent to other working paths if the current path is down or degraded (receiving errors).

`unconditional`—The path is always selected for carrying traffic, even if it is currently down or degraded (receiving errors).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Specifying the Conditions for Path Selection

select-profile

IN THIS SECTION

- Syntax | 16527
- Hierarchy Level | 16527
- Description | 16527
- Options | 16527
- Required Privilege Level | 16527
- Release Information | 16528

Syntax

```
select-profile profile-name
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options]
```

Description

Select the active profile to establish a data call with a Global System for Mobile Communications (GSM) cellular network.

Options

profile-name—Name of a configured profile that is to be used to establish a data call.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

selection-mode

IN THIS SECTION

- Syntax | 16528
- Hierarchy Level | 16528
- Description | 16528
- Options | 16529
- Required Privilege Level | 16529
- Release Information | 16529

Syntax

```
selection-mode (configured-quality | received-quality);
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Specify whether the clock source selection must use the configured or the received ESMC or SSM quality level for a qualifying interface. In both the selection modes, the interface qualifies for clock

source selection only when the received ESMC or SSM quality level on the interface is equal to or greater than the configured ESMC or SSM quality level for the interface.

Options

`configured-quality`	Set this option to let the clock source selection algorithm to use the ESMC or SSM quality level that is configured for a qualifying interface.
`received-quality`	Set this option to let the clock source selection algorithm to use the ESMC or SSM quality level that is received on the qualifying interface.

NOTE: For the `selection-mode` statement configuration to take effect, you must set the `quality-mode-enable` statement at the [`edit chassis synchronization`] hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

synchronization (MX Series) | **28795**

selective

IN THIS SECTION

- Syntax | **16530**
- Hierarchy Level | **16532**
- Description | **16532**

- Required Privilege Level | **16532**
- Release Information | **16532**

Syntax

```
selective {
    group multicast-prefix/prefix-length {
        source ip-prefix/prefix-length {
            ingress-replication {
                create-new-ucast-tunnel;
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                }
            ldp-p2mp;
            pim-ssm {
                group-range multicast-prefix;
            }
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
                static-lsp point-to-multipoint-lsp-name;
            }
            threshold-rate kbps;
        }
        wildcard-source {
            ldp-p2mp;
                pim-ssm {
                    group-range multicast-prefix;
                }
                rsvp-te {
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                }
            static-lsp point-to-multipoint-lsp-name;
            }
```

```
            threshold-rate kbps;
        }
    }
    tunnel-limit number;
    wildcard-group-inet {
        wildcard-source {
            ldp-p2mp;
            pim-ssm {
                group-range multicast-prefix;
            }
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
                static-lsp lsp-name;
            }
            threshold-rate number;
        }
    }
    wildcard-group-inet6 {
        wildcard-source {
            ldp-p2mp;
            pim-ssm {
                group-range multicast-prefix;
            }
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
                static-lsp lsp-name;
            }
            threshold-rate number;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-tunnel],
[edit routing-instances routing-instance-name provider-tunnel]
```

Description

Configure selective point-to-multipoint LSPs for an MBGP MVPN. Selective point-to-multipoint LSPs send traffic only to the receivers configured for the MBGP MVPNs, helping to minimize flooding in the service provider's network.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `ingress-replication` statement and substatements added in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

Configuring PIM-SSM GRE Selective Provider Tunnels

selective-vlan-switching

IN THIS SECTION

Syntax | 16533

- Hierarchy Level | 16533
- Description | 16533
- Options | 16533
- Required Privilege Level | 16533
- Release Information | 16533

Syntax

```
selective-vlan-switching {
    routing-instance routing-instance;
}
```

Hierarchy Level

```
[edit forwarding-options satellite fpc]
```

Description

Enables local switching for only a select number of VLANs on the satellite device.

Options

routing-instance	Configure routing instance.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

send-acct-status-on-config-change (Access Profile)

IN THIS SECTION

- Syntax | 16534
- Hierarchy Level | 16534
- Description | 16534
- Required Privilege Level | 16534
- Release Information | 16535

Syntax

```
send-acct-status-on-config-change;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router's authd process to send accounting messages when the RADIUS server status changes for an access profile. When you include this statement, authd sends an Acct-On message when the first RADIUS server is added to an access profile, and to send an Acct-Off message when the last RADIUS server is removed from an access profile.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Configuring Per-Subscriber Session Accounting

RADIUS Servers and Parameters for Subscriber Access

send (add-path)

IN THIS SECTION

- Syntax | 16535
- Hierarchy Level | 16536
- Description | 16536
- Options | 16536
- Required Privilege Level | 16538
- Release Information | 16538

Syntax

```
send {
    include-backup-path include-backup-path;
    multipath;
    path-count path-count;
    path-selection-mode {
        (all-paths | equal-cost-paths);
    }
    prefix-policy [ policy-names ... ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp group group-name family family family-
modifier add-path],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
family-modifier add-path],
[edit logical-systems logical-system-name routing-instances name protocols bgp group group-name
family family family-modifier add-path],
[edit logical-systems logical-system-name routing-instances name protocols bgp group group-name
neighbor address family family-modifier add-path],
[edit logical-systems logical-system-name tenants name routing-instances name protocols bgp
group group-name family family family-modifier add-path],
[edit logical-systems logical-system-name tenants name routing-instances name protocols bgp
group group-name neighbor address family family-modifier add-path],
[edit protocols bgp group group-name family family-modifier add-path],
[edit protocols bgp group group-name neighbor address family family-modifier add-path]
[edit routing-instances name protocols bgp group group-name family family-modifier add-path],
[edit routing-instances name protocols bgp group group-name neighbor address family family-
modifier add-path]
[edit tenants name routing-instances name protocols bgp group group-name family family-modifier
add-path],
[edit tenants name routing-instances name protocols bgp group group-name neighbor address family
family-modifier add-path]
```

Description

Enable advertisement of multiple paths to a destination, instead of advertising only the active path.

NOTE: On PTX1000 devices, the `add-path` configuration statement is not supported under routing instances.

Options

include-backup-path

Configure the number of backup paths that BGP must advertise. Do not configure this option if you have configured `all-paths` at the `[edit protocols bgp group name family name add-path send]` hierarchy level.

- **Range:** 1 through 2

multipath

Restrict BGP `add-path` to advertise contributor multiple paths only. Advertising all available multiple paths might result in a large overhead of processing on device memory. Selective advertising of multiple paths facilitates Internet service providers and data centers that use route reflectors to build in-path diversity in IBGP. You can limit and configure up to six prefixes that the BGP `multipath` algorithm selects. You cannot configure both `multipath` and `path-selection-mode` at the same time.

For example, if a routing device has four paths to a destination in its routing table and is configured to advertise up to two paths, only contributor paths for load balancing are chosen. The best contributor path is the active path and BGP advertises this path by default. The second best contributor path is selected and this process is repeated until the specified number of paths is reached; in this case, two additional paths to the same destination are selected for load balancing.

path-count ***number***

Specify the number of paths to a destination to advertise.

Suppose a routing device has in its routing table four paths to a destination and is configured to advertise up to three paths (`add-path send path-count 3`). The three paths are chosen based on path selection criteria. That is, the three best paths are chosen in path-selection order. The best path is the active path. This path is removed from consideration and a new best path is chosen. This process is repeated until the specified number of paths is reached.

- **Range:** 2 through 64
- **Default:** 1

path-selection-mode

Specify whether BGP can advertise all the available paths or only the equal-cost paths. BGP can advertise up to 64 add-path routes and a second best ECMP path as a backup in addition to the multiple ECMP paths. You cannot configure both `multipath` and `path-selection-mode` at the same time.

- Values:
  - all-paths—Specify the total number of add-path routes to advertise. BGP advertises all the paths that are allowed by the configured prefix policy and the maximum paths configured using the `path-count` option.
  - equal-cost-paths—Specify the equal-cost paths to advertise. You can configure either this option or the `multipath` option; they are mutually exclusive.

prefix-policy ***policy-names***

Specify the name of a policy (or a set of policies) configured at the `[edit policy-options]` hierarchy level to filter the paths to a destination to advertise.

Prefix policies enable you to filter routes on a router that is configured to advertise multiple paths to a destination. Prefix policies can only match prefixes. They cannot match route attributes, and they cannot change the attributes of routes.

The `add-path prefix-policy` allows up to 64 BGP add-paths to be advertised for a subset of prefixes that match the add-path prefix-policy. To enable this feature for a prefix, the `add-path prefix-policy` term matching the prefix should have a new *then* action to set `add-path send-count` *<2...64>*. This new action is not applicable if the policy-statement containing it is used in any place other than `add-path prefix-policy`.

- **Range:** 2 through 64 (for BGP)

> **NOTE**: This range is applicable only under `prefix-policy add-path`.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

Support for range from 2 through 20 (for BGP) introduced in Junos OS Release 14.1.

`multipath` option introduced in Junos OS Release 16.1R2.

Support for 64 BGP add-paths introduced in Junos OS Release 18.4R1 for the MX Series.

`include-backup-path` and `path-selection-mode` options introduced in Junos OS Release 18.4R1 for the MX Series and PTX Series.

RELATED DOCUMENTATION

Understanding the Advertisement of Multiple Paths to a Single Destination in BGP

Example: Advertising Multiple BGP Paths to a Destination

multipath (Protocols BGP) | **11110**

send-critical-event

IN THIS SECTION

- Syntax | 16539
- Hierarchy Level | 16539
- Description | 16539
- Required Privilege Level | 16539
- Release Information | 16539

Syntax

```
send-critical-event;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile action]
```

Description

Send OAM PDUs with the critical event bit set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Specifying the Actions to Be Taken for Link-Fault Management Events

send-dyn-subscription-indicator (PCRF Partition)

IN THIS SECTION

- Syntax | 16540
- Hierarchy Level | 16540
- Description | 16540
- Required Privilege Level | 16541
- Release Information | 16541

Syntax

```
send-dyn-subscription-indicator;
```

Hierarchy Level

```
[edit access pcrf partition partition-name],
```

Description

Include the Juniper-Network-Family-Indicator AVP into its message requests. To enable the Policy Control and Charging Rules Function (PCRF) to dynamically pass subscription-ID parameters, and support a variety of authentication, authorization, and provisioning configuration, the Juniper attribute-value pairs (AVPs) in Table 138 on page 16541 have been allocated from the Juniper Vendor-ID space (2636) vendor-specific attribute (VSA).

NOTE: If no dynamic-subscription ID is received, then neither Online Charging System (OCS) or Offline Charging System (OFCS) communications is initiated.

Table 138: Allocated Juniper AVPs

AVP Name	Vendor-ID	AVP Type	Diameter Type	Diameter Flag
Juniper-Dyn-Subscription-Indicator	2636	10001	Enum	V
Juniper-Dyn-Subscription-Id	2636	10002	Grouped	VM
Juniper-Dyn-Subscription-Id-Type	2636	10003	Integer32	VM
Juniper-Dyn-Subscription-Id-Data	2636	10004	UTF8String	VM

The client system (router) sends the Juniper-Dyn-Subscription-Id-Indicator AVP to indicate support of the dynamic assignment of the subscription ID. The Juniper-Dyn-Subscription-Id-Indicator attribute has two values:

- DYN_SUBSCRIPTION_NOT_SUPPORTED (0)
- DYN_SUBSCRIPTION_SUPPORTED (1)

The server then sends the Juniper-Dyn-Subscription-Id AVP to the client that indicated support. This is a grouped AVP that contains the values to be sent as Subscription-Id-Type and Subscription-Id-Data.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

Understanding Upstream and Downstream Messages for the PCRF

send-network-family-indicator (PCRF Partition)

IN THIS SECTION

- Syntax | 16542
- Hierarchy Level | 16542
- Description | 16542
- Required Privilege Level | 16543
- Release Information | 16543

Syntax

```
send-network-family-indicator;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Include the Juniper-Network-Family-Indicator AVP into its message requests. In many cases, wireline subscribers support only one IP family, which is required information for both AAA service and Policy Control and Charging Rules Function (PCRF). To indicate this information, the Juniper-Network-Family-Indicator AVP has been allocated from the Juniper Vendor-ID space (2636) VSA in Table 139 on page 16543.

Table 139: Network Family Indicator AVP

AVP Name	Vendor-ID	AVP Type	Diameter Type	Diameter Flag
Juniper-Network-Family-Indicator	2636	10010	Enum	V

The client system (router) sends the Juniper-Network-Family-Indicator AVP to indicate which network families are associated with the service request and supported by the subscriber. When you configure the Juniper-Network-Family-Indicator AVP to indicate the associated network family, the system sends the information to the PCRF. The Juniper-Network-Family-Indicator attribute has four values:

- UNSPECIFIED (0)
- IPV4_FAMILY (1)
- IPV6_FAMILY (2)
- IPV4_IPV6_FAMILY (3)

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the PCRF Partition

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gx Interactions Between the Router and the PCRF

Understanding Interactions Between the PCRF, PCEF, and OCS

Understanding Upstream and Downstream Messages for the PCRF

send-oam

IN THIS SECTION

- Syntax | 16544
- Hierarchy Level | 16544
- Description | 16544
- Required Privilege Level | 16544
- Release Information | 16544

Syntax

```
send-oam;
```

Hierarchy Level

```
[edit protocols l2circuit neighbor address interface interface-name static]
```

Description

Enable the ability to ping a static pseudowire. If you configure the `send-oam` statement, it applies to the backup neighbor as well. Once you have configured this statement, you can ping the static pseudowire by issuing the `ping mpls l2circuit` command.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Static Layer 2 Circuits

ping mpls l2circuit | 12550

send-origin-state-id (Diameter Applications)

IN THIS SECTION

- Syntax | 16545
- Hierarchy Level | 16545
- Description | 16545
- Required Privilege Level | 16546
- Release Information | 16546

Syntax

```
send-origin-state-id;
```

Hierarchy Level

```
[edit access pcrf partition partition-name],
[edit access ocs partition partition-name],
[edit diameter peer peer-name]
```

Description

Include the Origin-State AVP for the Diameter peer, OCS partition, or PCRF partition. The AVP is conveyed in Diameter-level or partition-level messages, respectively. The value of the AVP increases every time it changes, such as when the system reboots. Sending the Origin-State AVP in messages enables Diameter entities that receive the message to infer from a changed AVP value that sessions associated with a lower value are no longer active.

If you include the `send-origin-state-id` statement for the OCS or PCRF partition, a new value is sent for the partition only when a cold boot occurs. The system daemon detects the cold boot and changes the Origin-State-Id when required.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- *Configuring the OCS Partition*
- *Configuring the PCRF Partition*
- *Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers*
- *Understanding Gx Interactions Between the Router and the PCRF*
- *Understanding Gy Interactions Between the Router and the OCS*
- *Understanding Interactions Between the PCRF, PCEF, and OCS*

send (Protocols RIP)

IN THIS SECTION

- Syntax | 16547
- Hierarchy Level | 16547
- Description | 16547
- Options | 16547
- Required Privilege Level | 16547
- Release Information | 16548

Syntax

```
send send-options;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip],
[edit protocols rip group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip],
[edit routing-instances routing-instance-name protocols rip group group-name neighbor neighbor-
name]
```

Description

Configure RIP send options.

Options

send-options—One of the following:

- **broadcast**—Broadcast RIP version 2 packets (RIP version 1 compatible).
- **multicast**—Multicast RIP version 2 packets. This is the default.
- **none**—Do not send RIP updates.
- **version-1**—Broadcast RIP version 1 packets.
- **Default: multicast**

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets

receive (Protocols RIP) | **14117**

send (Protocols RIPng)

IN THIS SECTION

- Syntax | **16548**
- Hierarchy Level | **16548**
- Description | **16549**
- Options | **16549**
- Required Privilege Level | **16549**
- Release Information | **16549**

Syntax

```
send <none>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng],
[edit logical-systems logical-system-name protocols ripng group group-name neighbor neighbor-
name],
[edit logical-systems logical-system-name routing-instances routing-instances-name protocols
ripng],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
```

```
ripng group group-name neighbor neighbor-name],
[edit protocols ripng],
[edit protocols ripng group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols ripng],
[edit routing-instances routing-instance-name protocols ripng group group-name neighbor neighbor-
name]
```

Description

Enable or disable sending of update messages.

Options

none—(Optional) Disable sending of update messages.

- **Default:** Enabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

receive (Protocols RIPng) | **14119**

send-release-on-delete (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 16550
- Hierarchy Level | 16550
- Description | 16551
- Required Privilege Level | 16551
- Release Information | 16551

Syntax

```
send-release-on-delete;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 overrides],
```

```
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Send a release message to the DHCP (or DHCPv6) server whenever DHCP relay or relay proxy deletes a client. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Overriding the Default DHCP Relay Configuration Settings

Sending Release Messages When Clients Are Deleted

sender-based-rpf (MBGP MVPN)

IN THIS SECTION

- Syntax | 16552
- Hierarchy Level | 16552
- Description | 16552
- Required Privilege Level | 16553
- Release Information | 16553

Syntax

```
sender-based-rpf;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
mvpn],
[edit routing-instances routing-instance-name protocols mvpn]
```

Description

In a BGP multicast VPN (MVPN) with either RSVP-TE point-to-multipoint or MLDP point-to-multipoint provider tunnels, configure a downstream provider edge (PE) router to forward multicast traffic only from a selected upstream sender PE router.

Starting in Junos OS Release 21.1R1, you can configure MLDP point-to-multipoint provider tunnel on MX Series router.

BGP MVPNs use an alternative to data-driven-event solutions and bidirectional mode DF election because, for one thing, the core network is not exactly a LAN. Because, in an MVPN scenario, it is possible to determine which PE router has sent the traffic, Junos OS uses this information to only forward the traffic if it is sent from the correct PE router. With sender-based RPF, the RPF check is enhanced to check whether data arrived on the correct incoming virtual tunnel (vt-) interface and that the data was sent from the correct upstream PE router.

More specifically, the data must arrive with the correct MPLS label in the outer header used to encapsulate data through the core. The label identifies the tunnel and, if the tunnel is point-to-multipoint, the upstream PE router.

Sender-based RPF is not a replacement for single-forwarder election, but is a complementary feature. Configuring a higher primary loopback address (or router ID) on one PE device (PE1) than on another (PE2) ensures that PE1 is the single-forwarder election winner. The `unicast-umh-election` statement causes the unicast route preference to determine the single-forwarder election. If single-forwarder election is not used or if it is not sufficient to prevent duplicates in the core, sender-based RPF is recommended.

For RSVP point-to-multipoint provider tunnels, the transport label identifies the sending PE router because it is a requirement that penultimate hop popping (PHP) is disabled when using point-to-multipoint provider tunnels with MVPNs. PHP is disabled by default when you configure the MVPN protocol in a routing instance. The label identifies the tunnel, and (because the RSVP-TE tunnel is point-to-multipoint) the sending PE router.

The sender-based RPF mechanism is described in RFC 6513, *Multicast in MPLS/BGP IP VPNs* in section 9.1.1.

Sender-based RPF prevents duplicates from being sent to the customer even if there is duplication in the provider network. Duplication could exist in the provider because of a hot-root standby configuration or if the single-forwarder election is not sufficient to prevent duplicates. Single-forwarder election is used to prevent duplicates to the core network, while sender-based RPF prevents duplicates to the customer even if there are duplicates in the core. There are cases in which single-forwarder election cannot prevent duplicate traffic from arriving at the egress PE router. One example of this (outlined in section 9.3.1 of RFC 6513) is when PIM sparse mode is configured in the customer network and the MVPN is in RPT-SPT mode with an I-PMSI.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Support for MLDP point-to-multipoint provider tunnel is introduced in Junos OS Release 21.1R1 for MX Series router.

RELATED DOCUMENTATION

Understanding Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

Example: Configuring Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

unicast-umh-election | 30353

Example: Configuring Sender-Based RPF in a BGP MVPN with MLDP Point-to-Multipoint Provider Tunnels

sendid-tlv

IN THIS SECTION

- Syntax | 16554
- Hierarchy Level | 16554
- Description | 16555
- Required Privilege Level | 16555
- Release Information | 16555

Syntax

```
sendid-tlv {
    send-chassis-tlv;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management][edit protocols oam ethernet connectivity-fault-management maintenance-domain maintenance-domain-name maintenance-association maintenance-association-name continuity-check]
```

Description

Configures Junos OS to send the sender ID TLV along with the packets. The sender ID TLV is an optional TLV that is sent in continuity check messages (CCMs), loopback messages, and Link Trace Messages (LTMs), as specified in the IEEE 802.1ag standard. TLVs (type, length, and value) are described in the IEEE 802.1ag standard for Connectivity Fault Management (CFM) as a method of encoding variable-length and optional information in a protocol data unit (PDU).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R2.

RELATED DOCUMENTATION

CFM Monitoring between CE and PE Devices

sensor-based-stats

IN THIS SECTION

- Syntax | 16556
- Hierarchy Level | 16556
- Description | 16556
- Options | 16557
- Required Privilege Level | 16558
- Release Information | 16558

Syntax

```
sensor-based stats {
    per-interface-per-member-link {
        egress;
        ingress;
    }
    per-sid {
        egress;
        ingress;
    }
    subscribe {
     interval interval;
    }

}
```

Hierarchy Level

```
[edit protocols isis source-packet-routing],
```

```
[edit protocols (ospf | ospf3) source-packet-routing]
```

Description

Enable the collection of aggregate segment routing statistics for the Source Packet Routing in Networking (SPRING) traffic only, excluding RSVP and LDP-signaled traffic. Family MPLS statistics per interface are accounted for separately. The SR statistics also include SPRING traffic statistics per link aggregation group (LAG) member and per segment identifier (SID). The statistics are collected using the Juniper telemetry interface in an OpenConfig compliant format for Layer 3 interfaces.

> **NOTE**: This statement is also available for LDP and MPLS statistics. For more information see "sensor-based-stats (Junos Telemetry Interface)" on page 16558.
>
> For additional information about configuring these statistics sensors to stream data through gRPC, see *Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)*.

Enables the collection of on box telemetry spring sensor-based traffic statistics using the `subscribe` statement. As a result, the configuration triggers and records the on box telemetry for spring sensor-based statistics. The configuration creates a connection with network-agent and subscribes to the configured sensor path. The sensor path is installed on the Packet Forwarding Engine.

Options

per-interface-per-member-link

Configure sensor-based statistics per nexthop. Sensor-based statistics is the traffic statistics in a segment routing (SR) network that can be recorded in an OpenConfig compliant format for Layer 3 interfaces. The statistics is recorded for the Source Packet Routing in Networking (SPRING) traffic only, excluding RSVP and LDP-signaled traffic, and the family MPLS statistics per interface is accounted for separately. The SR statistics also includes SPRING traffic statistics per link aggregation group (LAG) member, and per segment identifier (SID).

egress Enable sensor based statistics on the egress interface.

ingress Enable sensor based statistics on the ingress interface.

> **NOTE**: On PTX Series Routers, the sensor-based statistics for SPRING traffic is recorded at the ingress interface only.

per-sid

Configure sensor based statistics per Source Packet Routing in Networking (SPRING) route. Sensor-based statistics is the traffic statistics in a segment routing (SR) network that can be recorded in an OpenConfig compliant format for Layer 3 interfaces. The statistics is recorded for SPRING traffic only, excluding RSVP and LDP-signaled traffic, and the family MPLS statistics per interface is accounted for separately. The SR statistics also includes SPRING traffic statistics per link aggregation group (LAG) member, and per segment identifier (SID).

egress Enable sensor based statistics for IP-MPLS egress accounting. This is supported only for segment routing label egress routes at the ingress provider edge (PE) device.

ingress Enable sensor based statistics for per-sid ingress accounting.

subscribe

Enables the collection of on box telemetry spring sensor- based traffic statistics with the interval range for interval configurationbetween 30 to 86400 seconds.

interval "interval" configuration indicates the time interval for sampling the telemetry updates. The configured range for interval can vary from 30 to 86400 seconds.

Required Privilege Level

routing—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

Support at the `[edit protocols (ospf | ospf3) source-packet-routing]` hierarchy level introduced in Junos OS Release 23.1R1 and Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Understanding Source Packet Routing in Networking (SPRING)

Understanding the Junos Telemetry Interface Export Format of Collected Data

sensor (Junos Telemetry Interface) | **16564**

sensor-based-stats (Junos Telemetry Interface)

IN THIS SECTION

- Syntax | **16558**
- Hierarchy Level | **16559**
- Description | **16559**
- Required Privilege Level | **16559**
- Release Information | **16559**

Syntax

```
sensor-based-stats;
```

Hierarchy Level

```
[edit protocols ldp traffic-statistics],
[edit protocols mpls]
```

Description

Enable the collection of certain statistics for the Junos telemetry interface. You must configure this statement when you configure a sensor to monitor and stream data for the following statistics:

- LDP—Enable the collection of LDP traffic statistics. Use the `resource` option at the `[edit services analytics sensor sensor-name]` hierarchy level to configure one of the LDP statistics sensors to stream data through UDP.
- MPLS—Enable the collection of LSP statistics. To enable a sensor to stream data for LSP statistics through UDP, include the `resource /junos/services/label-switched-path/usage/` statement at the `[edit services analytics sensor sensor-name]` hierarchy level.

> **NOTE**: This statement is also available for IS-IS statistics on MX Series, PTX3000, and PTX5000 routers. For more information, see *Understanding Source Packet Routing in Networking (SPRING)* and "sensor-based-stats (Protocols IS-IS)" on page 16555.

For additional information about configuring these statistics sensors to stream data through gRPC, see *Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)*.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F6.

Support at the `[edit protocols mpls]` hierarchy level introduced in Junos OS Release 19.2R1 for ACX6360 routers.

Statement supported in Junos OS Evolved Release 19.1R1 on PTX10003 routers and QFX 10003 switches.

Support at the [edit protocols ldp traffic-statistics] hierarchy level introduced in Junos OS Release 20.2R1 for MX Series routers and PTX Series routers.

RELATED DOCUMENTATION

Understanding the Junos Telemetry Interface Export Format of Collected Data

sensor-configuration

IN THIS SECTION

- Syntax | 16560
- Hierarchy Level | 16563
- Description | 16563
- Options | 16563
- Required Privilege Level | 16563
- Release Information | 16564

Syntax

```
sensor-configuration {
    application-identification                  {
        max-packet-memory-ratio percentage-value;
    }
    detector {
        protocol-name protocol-name {
            tunable-name  tunable-name {
                tunable-value protocol-value;
            }
        }
    }
    flow (Security IDP) {
        (allow-icmp-without-
flow                                                                | no-allow-icmp-
```

```
without-flow);
        fifo-max-size value;
        drop-if-no-policy-loaded;
        drop-on-failover;
        drop-on-limit;
        hash-table-size value;
        idp-bypass-cpu-threshold idp-bypass-cpu-threshold;
        idp-bypass-cpu-tolerance idp-bypass-cpu-tolerance;
        idp-bypass-cpu-usg-overload;
        intel-inspect-cpu-usg-threshold intel-inspect-cpu-usg-threshold;
        intel-inspect-cpu-usg-tolerance intel-inspect-cpu-usg-tolerance;
        intel-inspect-disable-content-decompress;
        intel-inspect-enable;
        intel-inspect-free-mem-threshold intel-inspect-free-mem-threshold;
        intel-inspect-mem-tolerance intel-inspect-mem-tolerance;
        intel-inspect-protocols [ intel-inspect-protocols ... ];
        intel-inspect-session-bytes-depth intel-inspect-session-bytes-depth;
        intel-inspect-signature-severity (critical | major | minor);
        (log-errors | no-log-errors);
        max-sessions-offset value;
        max-timers-poll-ticks value;
        min-objcache-limit-lt lower-threshold-value;
         min-objcache-limit-ut upper-threshold-value;
        reject-timeout value;
        (reset-on-policy | no-reset-on-policy);
        udp-anticipated-timeout value;
    }
    global {
        (enable-all-qmodules | no-enable-all-qmodules);
        (enable-packet-pool | no-enable-packet-pool);
        memory-limit-percent value;
        (policy-lookup-cache | no-policy-lookup-cache);
    }
    high-availability {
        no-policy-cold-synchronization;
    }
    ips {
        content-decompression-max-memory-kb value;
        content-decompression-max-ratio value;
        (detect-shellcode | no-detect-shellcode);
        fifo-max-size value;
        (ignore-regular-
expression
```

```
| no-ignore-regular-expression);
        log-supercede-min minimum-value;
        pre-filter-shellcode;
        (process-ignore-s2c | no-process-ignore-s2c);
        (process-override | no-process-override);
        process-port port-number;
    }
    log (Security IDP Sensor Configuration) {
        cache-size size;
        suppression {
            disable;
            (include-destination-
address
  | no-include-destination-address);
            max-logs-operate value;
            max-time-report value;
            start-log value;
        }
    }
    packet-log {
        host ip-address <                                                    port number>;
        max-sessions percentage;
        source-address ip-address;
        total-memory percentage;
     ssl-profile-name <profile-name>;}
    re-assembler {
        action-on-reassembly-failure (drop | drop-session |
ignore);
        (force-tcp-window-checks | no-force-tcp-window-
checks);
        (ignore-memory-
overflow
           | no-ignore-memory-overflow);
        (ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-
overflow);
        ignore-reassembly-overflow;
        max-flow-mem value;
        max-packet-mem-ratio percentage-
value;
         max-synacks-queued value;
        (tcp-error-logging | no-tcp-error-
logging);
    }
```

```
    ssl-inspection {
        cache-prune-chunk-size number;
        key-protection;
        maximum-cache-size number;
        session-id-cache-timeout seconds;
        sessions number;
    }
}
```

Hierarchy Level

```
[edit security idp]
```

Description

Configure various IDP parameters to match the properties of transiting network traffic.

Options

- `ssl-profile-name` *`profile-name`* —SSL initiation profile name to be used for encrypted packet log transmission.

  Starting in Junos OS Release 22.1R1, you can enable a secure SSL connection and send encrypted IDP packet capture log to the packet capture receiver. To establish the SSL connection you must specify the SSL initiation profile you want to use in the IDP packet log configuration.

  The SSL profile name must be configured in the SSL initiation profile configuration. Configuration and installation of SSL certificates and SSL handshake related configurations that are required to establish secure connection are performed as part of SSL initiation profile configuration. SSL versions are chosen based on the SSL initiation configuration. SSL profile needs to be configured in each logical system separately .

  If SSL profile name is not configured in SSL initiation profile configuration, then the following message is displayed **Referenced SSL initiation profile is not defined**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. Packet memory ratios added in Junos OS Release 12.1X44-D20.

intel-inspect-cpu-usg-threshold, intel-inspect-cpu-usg-tolerance, intel-inspect-disable-content-decompress, intel-inspect-enable, intel-inspect-free-mem-threshold, intel-inspect-mem-tolerance, intel-inspect-protocols, intel-inspect-session-bytes-depth, and intel-inspect-signature-severity options added in Junos OS Release 19.2R1.

sensor (Junos Telemetry Interface)

IN THIS SECTION

- Syntax | **16564**
- Hierarchy Level | **16565**
- Description | **16565**
- Options | **16565**
- Required Privilege Level | **16592**
- Release Information | **16592**

Syntax

```
sensor sensor-name {
    export-name export-profile-name;
    polling-interval seconds;
    resource resource-string;
    <resource-filter regular expression>;
    server-name [ streaming-server-names ];
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Configure a Junos telemetry interface sensor, which defines the parameters of a system resource to monitor and stream data. You can use regular expressions to filter the data collected. Examples include filters for logical and physical interfaces and LSP messages. To apply different filters to the same system resource, you configure multiple sensors. For example, you can configure multiple logical interface sensors and apply a different interface filter to each one.

Options

Each sensor configuration requires you to specify the following: sensor name, an export profile name, a resource identifier string that enables monitoring and streaming of data for the specified system resource, and a server name to collect data. A regular expression to filter data for the specified resource is optional.

sensor-name

Specify a name that defines the sensor configuration. For example, for a sensor configuration that monitors all LSP events, you might choose the name `lsp-mon-global`. For a sensor configuration that monitors events only for an LSP named A2B, you might choose the name `lsp-mon-A2B`.

export-name ***export-profile-name***

Specify the name of an export profile that you configured at the [`edit services analytics export-profile` *name*] hierarchy level to associate with the sensor. This export profile defines the parameters for exporting telemetry data, such as a format for exported data and the rate at which data is generated for export.

> **NOTE**: You can apply only one export profile to each sensor configuration. The only supported transport protocol when you configure a sensor through the CLI is UDP.

polling-interval ***seconds***

Specify the interval at which the Junos Telemetry Interface sensor generates data to export to the collector.

As the configured interval expires, the most recent sample collected by the sensor is gathered and forwarded to the server configured to collect data.

> **NOTE**: For Packet Forwarding Engine sensors, the minimum reporting rate is 2 seconds.

- **Range:** 1 through 3600 (1 hour)

resource *resource-string*

Enable the system resource to monitor and stream data. Each string corresponds to a specific system resource. The format is a file path and must be entered exactly. You can associate only one *resource-string* with a *sensor-name*. Configure a separate sensor for each system resource you want to monitor. The resource string to enable LSP monitoring can be modified to specify a specific LSP.

> **NOTE**: You can configure more than one sensor to monitor the same system resource. Configuring different sensors for the same system resource allows you configure different parameters for monitoring that resource.

Table 140 on page 16566 lists each supported *resource-identifier-string*, a description of the system resource monitored, and additional configuration information.

You can also use the Telemetry Explorer tool to search for and view information about *resource-identifier-string*.

Table 140: resource statement Options

resource string	Description	Release Information
/interfaces/interface/ethernet/state/counters/jnx-pfc-statistics	Transport statistics providing pause frame receive and transmit counter support.	Junos OS Evolved Release 21.4R1 and later on PTX10001-36MR, PTX10004, PTX10008, and PTX10016 routers.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/events	System events sensor. Starting with Junos OS Release 18.1R1, this sensor corresponds to system log messages (syslog). The sensor must be used with an `export-profile` that has a `reporting-rate` of 0, To subscribe for specific events, you can subscribe for /junos/events/event[id='EVENT_NAME'] where event EVENT_NAME is the event id that you are interested in. Alternatively, you can subscribe to any XPATH Many event names can be found in the messages log file.	Junos OS 18.1R1 and later on all JTI platforms.
/junos/services/ip-tunnel/usage/	Packet Forwarding Engine packet statistics sensor. The statistics are used to report various network element performance metrics in a scalable and efficient way, providing visibility into Packet Forwarding Engine errors and drops. A timestamp indicating when the counters were last reset is included with all the exported data to allow collectors to determine if and when a reset event happened; for example, if the Packet Forwarding Engine hardware restarted. Exported statistics are similar to the output of the operational mode command `show nhdb hw dynamic-ip-tunnels`.	Junos OS 17.4R1 and later on MX Series devices.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/services/label-switched-path/usage/	Packet Forwarding Engine sensor for LSP statistics. Starting with Junos OS Release 17.4R1 on MX Series and PTX Series routers only, statistics for bypass LSPs are also exported. Previously, only statistics for ingress LSPs were exported. For bypass LSPs, the following are exported: • Bypass LSP originating at the ingress router of the protected LSP • Bypass LSP originating at the transit router of the protected LSP • Bypass LSP protecting the transit LSP as well as the locally originated LSP When the bypass LSP is active, traffic is exported both on the bypass LSP and the ingress (protected) LSP. On MX Series routers only, bidirectional LSPs for ultimate-hop popping (UHP) are also supported. **NOTE**: You can modify `/junos/services/label-switched-path/usage/` to specify a specific LSP. Add `__instance__/`*`lsp-name`* to the end of the resource string identifier. For example, to monitor and stream data for LSP statistics for an LSP named mirror-to-murano-1, enter the following: `/junos/services/label-switched-path/usage/ __instance__/mirror-to-murano-1`. If you do not specify a specific LSP name, the system resource monitors and streams data for all LSPs. When you enable a sensor for LSP statistics, you must also configure the `sensor-based-stats` statement at the `[edit protocols mpls]` hierarchy level. MX Series routers must also operate in enhanced mode. If not enabled by default, configure either the `enhanced-ip` statement or the `enhanced-ethernet` statement at the `[edit chassis network-services]` hierarchy level.	Junos OS Release 15.1F6 and later. Junos OS Release 17.2R1 and later on QFX10000 switches and PTX1000 routers. Junos OS Release 17.3 and later on EX9200 and QFX5110 switches. Junos OS Release 18.2R1 and later on QFX5100, QFX5110, and QFX5200 switches Junos OS Release 18.3R1 and later on QFX5120-48Y and EX4650 switches Junos OS Release 18.4R1 and later on EX4600 switches Junos OS Release 19.1R1 and later on QFX10002 switches and PTX10002 routers

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/services/ldp/label-switched-path/ingress/usage/	Packet Forwarding Engine sensor for LDP LSP ingress statistics. When you enable a sensor for LDP statistics, you must also configure the `sensor-based-stats` statement at the `[edit protocols ldp traffic-statistics]` hierarchy level. On PTX Series routers, this feature is not supported for the following variants: • PTX3000 and PTX5000 with the RE-DUO-C2600-16G Routing Engine • PTX10008 with the PTX10K-LC1201-36CD line card	Junos OS Release 20.2R1 for MX series and PTX series routers. Junos OS Release 20.3R1 on MPC10E-10C-MRATE, MPC10E-15C-MRATE, and MX2K-MPC11E line cards.
/junos/services/ldp/label-switched-path/transit/usage/	Packet Forwarding Engine sensor for LDP LSP transit statistics. On PTX Series routers, this feature is not supported for the following variants: • PTX3000 and PTX5000 with the RE-DUO-C2600-16G Routing Engine • PTX10008 with the PTX10K-LC1201-36CD line card	Junos OS Release 20.2R1 for MX series and PTX series routers. Junos OS Release 20.3R1 on MPC10E-10C-MRATE, MPC10E-15C-MRATE, and MX2K-MPC11E line cards.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/services/ldp/p2mp/interface/receive/usage/	Packet Forwarding Engine sensor for LDP P2MP receive statistics. On PTX Series routers, this feature is not supported for the following variants: • PTX3000 and PTX5000 with the RE-DUO-C2600-16G Routing Engine • PTX10008 with the PTX10K-LC1201-36CD line card	Junos OS Release 20.2R1 for MX series and PTX series routers. Junos OS Release 20.3R1 on MPC10E-10C-MRATE, MPC10E-15C-MRATE, and MX2K-MPC11E line cards.
/junos/services/ldp/p2mp/interface/transmit/usage/	Packet Forwarding Engine sensor for LDP P2MP transmit statistics. On PTX Series routers, this feature is not supported for the following variants: • PTX3000 and PTX5000 with the RE-DUO-C2600-16G Routing Engine • PTX10008 with the PTX10K-LC1201-36CD line card	Junos OS Release 20.2R1 for MX series and PTX series routers. Junos OS Release 20.3R1 on MPC10E-10C-MRATE, MPC10E-15C-MRATE, and MX2K-MPC11E line cards.
/junos/services/ldp/p2mp/label-switched-path/usage/	Packet Forwarding Engine sensor for LDP P2MP LSP statistics. On PTX Series routers, this feature is not supported for the following variants: • PTX3000 and PTX5000 with the RE-DUO-C2600-16G Routing Engine • PTX10008 with the PTX10K-LC1201-36CD line card	Junos OS Release 20.2R1 for MX series and PTX series routers. Junos OS Release 20.3R1 on MPC10E-10C-MRATE, MPC10E-15C-MRATE, and MX2K-MPC11E line cards.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/services/segment-routing/sid/usage/	Source Packet Routing in Networking (SPRING) sensorSfor transit statistics. SPRING is also known as segment routing. Before statistics can be exported, you must first enable them by including the `sensor-based-stats` statement at the `[edit protocols isis source-packet-routing]` hierarchy level. For more information, see *Enabling Export of Transit SPRING Statistics*.	Junos OS Release 19.1R1 and later on PTX3000 routers and PTX5000 with FPC2
/junos/services/spu/delegated-rpm/	Delegated Realtime Performance Monitoring (RPM) service sensor. Delegated RPM is a mode where RPM probe generation and measurement calculation are done by MS-MIC and MS-MPC cards. This hardware assistance allows a very high scale of concurrent RPM probes. You can use the resulting data from this sensor to improve network design and optimize traffic engineering. Data can also be used to detect problems in individual devices as well as in the overall network and the traffic carried by it. JTI sensor support for other RPM modes was added in Junos OS Release 18.3R1. This sensor has the following limitations: • Configuring multiple export profiles for the same resource for delegated RPM may not provide expected results. • Multiple sensors for single resource-path (such as delegated RPM) is not supported. • Due to an egress packet-size limitation, history outputs are limited to 5 per RPM test. • The sensor exports one RPM test record per export packet.	Junos OS Release 19.1R1 and later on MX Series routers operating with MS-MIC and MS-MPC

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/services/spu/ipsec-vpn	UDP-based PIC sensors. Starting with Junos OS Release 18.1R1, this sensor provides visibility for IPSec services on different service complexes and nodes. Exported data is defined using an IP address and a UDP port. When an export interval expires, the most recent statistics collected by the sensors are gathered, placed in the payload of a UDP packet, and forwarded to a collector. A timestamp indicating when counters are read is included with the exported data to allow collectors to collate data. The timestamp also can determine if and when an event happened, such as a PIC hardware restart or if counters were cleared by means of the CLI.	Junos OS 18.1R1 on MX Series with MS-MICs and MS-MPCs
/junos/services/spu/servicesets	Sensor to export service set statistics. These sensors provide visibility for services on different service complexes and nodes (for example, IPSec services). Exported data is defined using an IP address and a UDP port. When an export interval expires, the most recent statistics collected by the sensors are gathered, placed in the payload of a UDP packet, and forwarded to a collector. A timestamp indicating when counters are read is included with the exported data to allow collectors to collate data. The timestamp also can determine if and when an event happened, such as a PIC hardware restart or if counters were cleared by means of the CLI.	Junos OS 18.2R1 on MX Series with MS-MICs and MS-MPCs

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/services/spu/sessions	Sensor to export session statistics. These sensors provide visibility for services on different service complexes and nodes (for example, IPSec services). Exported data is defined using an IP address and a UDP port. When an export interval expires, the most recent statistics collected by the sensors are gathered, placed in the payload of a UDP packet, and forwarded to a collector. A timestamp indicating when counters are read is included with the exported data to allow collectors to collate data. The timestamp also can determine if and when an event happened, such as a PIC hardware restart or if counters were cleared by means of the CLI.	Junos OS 18.2R1 on MX Series with MS-MICs and MS-MPCs
/junos/system/linecard/node-slicing/af-fab-stats/	Sensor to export abstracted fabric (AF) interface specific load-balancing and fabric queue statistics. This sensor is only supported for a node virtualization configuration on MX series routers with an AF Interface as the connecting link between guest network functions (GNFs). The sensor also reports aggregated statistics across all AF interfaces hosted on a source packet forwarding engine of local GNFs along with the fabric statistics for all traffic ingressing from and egressing to the fabric from that the packet forwarding engine.	Junos OS Release 18.4R1 and later on MX480, MX960, MX2008, MX2010, MX2020, and MX-ELM routers

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/system/linecard/cos/interface/interface-set/output/queue/	Sensor for logical interface (IFL)-set. The sensor streams queue statistics using Juniper proprietary gRPC and gNMI or by means of UDP. Zero suppression (suppressing zero values in statistics from streamed data) is also supported.	Junos OS Release 22.3R1and later on MX204, MX240, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016 routers with Trio chipset EA, ZT, and YT-based fixed systems and modular systems linecards.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/system/linecard/cpu/memory/	Packet Forwarding Engine sensor for CPU memory.	Junos OS Release 16.1R3 and later. Junos OS Release 17.2R1 and later on QFX10000 switches and PTX1000 routers. Junos OS Release 17.3R1 and later on EX9200 and QFX5110 switches. Junos OS Release 18.2R1 and later on QFX5100, QFX5110, and QFX5200 witches Junos OS Release 18.3R1 and later on QFX5120-48Y and EX4650 switches Junos OS Release 18.4R1 and later on EX4600 switches Junos OS Release 19.1R1 and later on QFX10002 switches and PTX10002 Routers Junos OS Release 20.2R1 and later on EX2300, EX2300-MP, and EX3400 switches

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/system/linecard/ddos/	Distributed denial of service (DDoS) sensor. This sensor supports the Openconfig data model junos/ui/openconfig/yang/ and junos-ddos.yang. You can stream information using Juniper proprietary gRPC or UDP (native) export. There are 45 packet types for DDoS. To maintain a reasonably sized data stream, data is exported for all protocols that have seen traffic using the zero-suppression model. On QFX5000 platforms, multiple protocols can share the same CPU queue. DDoS configurations are applied at the CPU queue level. Consequently, DDoS statistics fetched from the CPU queue will return the aggregate value of all protocols using that queue. For example, if BGP, LDP, and RSVP protocols are using a particular CPU queue, but the DDoS limit is violated only by the BGP protocol, the DDoS violation reported will include all three protocols: BGP, LDP, and RSVP. This information will be exported to the collector with the DDoS sensor.	Junos OS Release 22.1R1 and later on EX4650, QFX5110, QFX5120-48Y, QFX5200 and QFX5210 switches Junos OS Evolved Release 22.3R1 and later on PTX10001-36MR, PTX10003, PTX10004, PTX10008 and PTX10016 routers.
/junos/system/linecard/environment/	FPC environment sensor `/junos/system/linecard[name=FPC0]/environment/` is supported to stream voltage, temperature, and power information. You can stream the following endpoints: • `/power-record/max-fpc-power` `/power-record/fpc-power` • `voltage-record/voltage-sensor-name` • `voltage-record/voltage-value` • `temp-sensor-name` • `temp-record/temp-value`	Junos OS Release 22.4R1 and later on MX10004 routers.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/system/linecard/firewall/	Packet Forwarding Engine sensor for firewall filter counters and policer counters. Each line card reports counters separately. **NOTE**: Hierarchical policer statistics are collected for MX Series routers only. Traffic-class counter statistics are collected for PTX Series routers and QFX10000 switches only. Firewall counters are exported even if the interface to which the firewall filer is attached is down.	Junos OS Release 15.1F5 and later. Junos OS Release 17.2R1 and later on QFX10000 switches. Junos OS Release 17.3R1 and later on PTX1000 routers and EX9200 switches and QFX5110 switches. Junos OS Release 18.2R1 and later on QFX5100, QFX5110, and QFX5200 switches Junos OS Release 18.3R1 and later on QFX5120-48Y and EX4650 switches Junos OS Release 18.4R1 and later on EX4600 switches Junos OS Release 20.2R1 and later on EX2300, EX2300-MP, and EX3400 switches.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/system/linecard/interface/	Packet Forwarding Engine sensor for physical interface traffic. **NOTE**: For PTX Series routers, for a specific interface, queue statistics are exported for each line card. For MX series routers, interface queue statistics are exported only from the slot on which an interface is configured. For Aggregated Ethernet interfaces, statistics are exported for the member physical interfaces. You must aggregate the counters at the destination server, or collector. If a physical interface is administratively down or operationally down, interface counters are not exported. Issuing an operational `clear` command, such as `clear interfaces statistics all`, does not reset statistics exported by the line card.	Junos OS Release 15.1F3 and later on PTX Series routers only. Support introduced for MX Series routers in Junos OS Release 15.1F5. Junos OS Release 17.2R1 and later on QFX10000 switches and PTX1000 routers. Junos OS Release 17.3R1 and later on EX9200 switches, QFX5110 switches and MX150 routers. Junos OS Release 18.2R1 and later on QFX5100, QFX5110, and QFX5200 switches Junos OS Release 18.3R1 and later on QFX5120-48Y and EX4650 switches Junos OS Release 18.4R1 and later on EX4600 switches Junos OS Release 19.1R1 and later on QFX10002 Switches and PTX10002 routers

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
		Junos OS Release 20.2R1 and later on EX2300, EX2300-MP, and EX3400 switches Junos OS Evolved Release 20.2R1 and later on PTX10008 routers. UDP streaming is not supported over the management interface.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/system/linecard/interface/logical/usage/	Packet Forwarding Engine sensor for logical interface statistics. **NOTE**: If a logical interface is operationally down, interface statistics continue to be exported. Issuing an operational `clear` command, such as `clear interfaces statistics all`, does not reset statistics exported by the line card. **NOTE**: Locally injected packets from the Routing Engine are not exported.	Junos OS Release 15.1F5 and later. Junos OS Release 17.2R1 and later on QFX10000 switches. Junos OS Release 17.3R1 and later on EX9200 and QFX5110 switches Junos OS Release 18.2R1 and later on QFX5100, QFX5110, and QFX5200 switches Junos OS Release 18.3R1 and later on QFX5120-48Y and EX4650 switches Junos OS Release 18.4R1 and later on EX4600 switches Junos OS Release 20.2R1 and later on EX3400 switches Junos OS Evolved Release 20.2R1 and later on PTX10008 routers. UDP streaming is not supported over the management interface.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/system/linecard/interface/queue/	Packet Forwarding Engine sensor for physical interface traffic. Exports all queue fields from /junos/system/linecard/interface/. To export traffic and queue data for physical interfaces, use /junos/system/linecard/interface/. To export traffic fields only, use /junos/system/linecard/interface/traffic/.	Junos OS Release 18.3R1 and later on PTX Series and ACX Series routers and EX Series, MX Series, and QFX Series switches.
/junos/system/linecard/interfaces/logical/usage/	Packet Forwarding Engine sensor for logical interface statistics.	Junos OS Evolved Release 21.4R1 and later on PTX10001-36MR, PTX10004, PTX10008, and PTX10016 routers.
/junos/system/linecard/interfaces/traffic/	Packet Forwarding Engine sensor for physical interface traffic. Exports traffic fields only	Junos OS Evolved Release 21.4R1 and later on PTX10001-36MR, PTX10004, PTX10008, and PTX10016 routers.
/junos/system/linecard/intf-exp/	Interface express sensor. This sensor leverages statistics out of the physical interface sensor, providing faster and more frequent operational status statistics. Only the physical interfaces' operational status from the Flexible PIC Concentrator (FPC) is collected and reported. Statistics from the Routing Engine interface are not reported.	Junos OS Release 18.1R1 and later on PTX1000, PTX3000, PTX5000, and PTX10000 routers. Junos OS Release 19.3R1 and later on MX960, MX2010, and MX2020 routers.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/system/linecard/npu/memory/	Packet Forwarding Engine sensor for network processing unit (NPU) memory.	Junos OS Release 16.1R3 and later. Junos OS Release 17.2R1 and later on QFX10000 switches. Junos OS Release 17.3R1 and later on EX9200 switches. Junos OS Release 19.1R1 and later on PTX10002 routers. Junos OS Evolved Release 20.2R1 and later on PTX10008 routers. UDP streaming is not supported over the management interface. Junos OS Release 21.4R1 and later on EX4650, QFX5110, QFX5120-48Y, QFX5200, and QFX5210 switches. Supported on QFX5130, QFX5700, QFX5220-32CD, QFX5220-128C, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5230-64CD, QFX5240-64QD,

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
		and QFX5240-64QD switches starting with Junos OS Release 24.4R1 for exporting telemetry information using gRPC or gNMI services.
/junos/system/linecard/npu/utilization/	Packet Forwarding Engine sensor for NPU processor utilization.	Junos OS Release 16.1R3 and later. Junos OS Release 17.2R1 and later on QFX10000 switches. Junos OS Release 17.3R1 and later on EX9200 switches. Junos OS Release 19.1R1 and later on PTX10002 routers. Junos OS Evolved Release 20.2R1 and later on PTX10008 routers. UDP streaming is not supported over the management interface.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/npu-memory/	Sensor that exports both NPU memory statistics from the Packet Forwarding Engine and flow-label statistics from the Routing Engine. To export only flow-label statistics, include the `junos/npu-memory/flabel-memory/` resource string.	Junos OS Release 16.1R3 and later on PTX Series routers only. **NOTE**: Junos OS Release 17.2R1 and later on PTX1000 routers.
/junos/system/linecard/services/inline-jflow/	Packet Forwarding Engine sensor for performance metrics of the inline flow sampling process, such as the number of active flows and the number of exported flows.	Junos OS Release 16.1R3 and later on MX series and PTX series routers only. Junos OS Release and later on EX9200 switches, PTX1000 routers, and MX150 routers.
/junos/system/linecard/optics/	Packet Forwarding Engine sensor for various optical performance metrics, such as transmit and receive power levels.	Junos OS Release 17.1R1 and later. Junos OS Release and later 17.2R1 on QFX10000 switches. Junos OS Release 17.3R1 and later on EX9200 switches and PTX1000 routers.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
`/junos/system/linecard/packet/usage/`	Sensor for Packet Forwarding Engine error and drop statistics. Use these statistics to optimize traffic engineering and improve your network design. When you include the resource path `/junos/system/linecard/packet/usage/` in a subscription, statistics are streamed in the format: `/components/component[name='FPC0:NPU3']/properties/property[name='hwds-dlu-not-routable']/state/value`	Junos OS Release 22.1R1 and later on PTX1000 and PTX5000 routers and QFX10002-60C switches using Juniper proprietary gRPC.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information

/junos/system/linecard/qmon/	Sensor for queue depth statistics for ingress and egress queue traffic. Statistics are exported directly from the line card.
	
	This sensor only supports single -streaming. Configuring this sensor to stream to multiple servers is not supported. If multiple servers are configured, no data is sent to any of the configured servers.
	
	The following example shows a configuration for single-streaming that will send data:
	
	`sensor qmon {`
	`server-name TEMP;`
	`export-name export-common;`
	`resource /junos/system/linecard/qmon/;`
	`}`
	
	The following example shows a multiple-server configuration that will not send data:
	
	`sensor qmon {`
	`server-name TEMP;`
	`server-name digi1;`
	`server-name digi2;`
	`export-name export-common;`
	`resource /junos/system/linecard/qmon/;`
	`}`
	
	NOTE: Issuing an operational `clear` command, such as `clear interfaces statistics all`, does not reset the statistics exported by the line card.
		Junos OS Release 17.1R1 and later on MX Series routers on MPC7E, MPC8E, and MPC9E only.
		
		Junos OS 17.3R1 and later on EX9200 switches.
		
		NOTE: virtual MX Series (vMX) routers are not supported.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/system/linecard/qmon-sw/	Sensor for congestion and latency monitoring statistics.	Junos OS Release 18.2R1 and later on QFX5100, QFX5110, and QFX5200 Switches Junos OS Release 18.3R1 and later on QFX5120-48Y and EX4650 Switches Junos OS Release 18.4R1 and later on EX4600 switches Junos OS Release 20.2R1 and later on EX3400 switches Junos OS Release 21.2R1 and later on PTX5000, PTX1000, PTX10002, PTX10008, and PTX10016 routers and QFX10002, QFX10008, and QFX10016 switches

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/system/linecard/fabric/	Sensor for fabric statistics. The following types of statistics can be exported: • Fabric statistics for Packet Forwarding Engine pairs (`resource-filter` option is not supported) • FPC fabric statistics • Control Board and Switch Fabric Board fabric statistics.	Junos OS Release 17.2R1 and later on MX Series routers only. Junos OS Evolved Release 20.2R1 and later on PTX10008 routers. UDP streaming is not supported over the management interface. **NOTE**: virtual MX Series (vMX) routers are not supported.
/junos/system/linecard/packet/usage/	Sensor for Packet Forwarding Engine Statistics. This sensor exports statistics for counters and provides visibility into Packet Forwarding Engine error and drop statistics.	Junos OS Release 17.4R1 and later on MX Series and PTX Series routers

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
/junos/services/segment-routing/interface/ingress/usage/ /junos/services/segment-routing/interface/egress/usage/ /junos/services/segment-routing/sid/usage/	Sensors for aggregate segment routing traffic with IS-IS. The first path exports inbound traffic. The second path exports outbound traffic. The third path exports inbound segment routing traffic for each segment identifier. **NOTE**: When you enable a sensor for segment routing statistics, you must also configure the `sensor-based-stats` statement at the `[edit protocols isis source-packet-routing]` hierarchy level. MX Series and PTX Series routers must also operate in enhanced mode. On MX Series routers, If not enabled by default, configure either the `enhanced-ip` statement or the `enhanced-ethernet` statement at the `[edit chassis network-services]` hierarchy level. On PTX Series routers, configure the `enhanced-mode` statement at the `[edit chassis network-services]` hierarchy level.	Junos OS Release 17.4 and later on MX Series and PTX5000 routers.
`/junos/system/subscriber-management/aaa/diameter/clients/gx`	Sensor for policy and charging rules function (PCRF) statistics for subscribers.	Junos OS Release 19.3R1 on MX5, MX10, MX40, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX100016 routers.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
`/junos/system/subscriber-management/aaa/diameter/clients/gy`	Sensor for Online Charging System (OCS) statistics for subscribers.	Junos OS Release 19.3R1 on MX5, MX10, MX40, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX100016 routers.
`/junos/system/subscriber-management/aaa/diameter/clients/nasreq`	Sensor for Network Access Server Application (NASREQ) statistics for subscribers.	Junos OS Release 19.3R1 on MX5, MX10, MX40, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX100016 routers.
`/junos/system/subscriber-management/aaa/diameter/peers/peer[peer_address='`*peer-address*`']/gx/response-time`	Diameter peer sensor that provides response time measurements for messages exchanged between an MX router and the peer for PCRF statistics. This sensor includes response-time and delay measurements in milliseconds. **NOTE**: The delay measurements are made over a 60-second measurement interval. As the reporting interval may be as much as 59 seconds out of phase with the measurement interval, the response time values may not be aligned with the reporting interval.	Junos OS Release 19.3R1 on MX5, MX10, MX40, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX100016 routers.

Table 140: resource statement Options *(Continued)*

resource string	Description	Release Information
`/junos/system/subscriber-management/aaa/diameter/peers/peer[peer_address='peer-address']/gy/response-time`	Diameter peer sensor that provides response time measurements for messages exchanged between an MX router and the peer for OCS statistics. This sensor includes response-time and delay measurements in milliseconds. **NOTE**: The delay measurements are made over a 60-second measurement interval. As the reporting interval may be as much as 59 seconds out of phase with the measurement interval, the response time values may not be aligned with the reporting interval.	Junos OS Release 19.3R1 on MX5, MX10, MX40, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX100016 routers.
`/junos/system/subscriber-management/aaa/diameter/peers/peer[peer_address='peer-address']/nasreq/response-time`	Diameter peer sensor that provides response time measurements for messages exchanged between an MX router and the peer for NASREQ statistics. This sensor includes response-time and delay measurements in milliseconds. **NOTE**: The delay measurements are made over a 60-second measurement interval. As the reporting interval may be as much as 59 seconds out of phase with the measurement interval, the response time values may not be aligned with the reporting interval.	Junos OS Release 19.3R1 on MX5, MX10, MX40, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX100016 routers.
`/junos/system/linecard/ddos/`	This PFE sensor exports the statistics of DDOS from MPC1, MPC2, MPC3, MPC5, MPC6, MPC7, MPC8, and MPC9 line cards.	Junos OS Release 21.1R1 on MX Series routers.
`/junos/system/linecard/page-drops/page-drop/`	Page drop statistics include page drop counter, interface and queue details.	Junos OS Evolved Release 22.3R1 on PTX10003 routers.

resource-filter *regular-expression*

(Optional) Specify a regular expression to filter data for a specific resource. For example, you can filter for a specific set of logical or physical interfaces, firewall filters, or LSP messages. When you configure a system resource to monitor and stream data globally—that is, systemwide—you do not need to include a regular expression.

Examples of regular expressions to filter data exported through sensor configuration:

- Logical interface statistics sensor—et-2/0/7:1*
- LSP events sensor—lsp-from-A-to-B*
- Firewall filter counters sensor—f_testl*

server-name [*streaming-server-names*]

Specify one or more servers to transport data for collection. Include at least one server-name configured at the [`edit services analytics streaming-server` *`server-name`*] hierarchy level.

> **NOTE**: Starting in Junos OS Release 15.1F6, you can configure as many as four streaming servers for a single sensor configuration. In previous releases, you can specify only one streaming server for each configured sensor. To specify more than one streaming server for a sensor, you must enclose the names in brackets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F3.

Support for MPC7E, MPC8E, and MPC9E on MX Series routers added in Junos OS Release 15.1F5.

Support for FPC1 and FPC2 on PTX Series routers added in Junos OS Release 16.1R3.

Statement introduced in Junos OS Release 19.1R1 for PTX3000 routers and PTX5000 routers with FPC2.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
19.1R1	Statement introduced in Junos OS Release 19.1R1 for PTX3000 routers and PTX5000 routers with FPC2.

RELATED DOCUMENTATION

export-profile (Junos Telemetry Interface) | 6669

separator

IN THIS SECTION

- Syntax | 16593
- Hierarchy Level | 16593
- Description | 16593
- Options | 16594
- Required Privilege Level | 16594
- Release Information | 16594

Syntax

```
separator special-character;
```

Hierarchy Level

```
[edit access profile profile-name client-name-filter client-name]
```

Description

Specify a character to identify where stripping of characters occurs in a client name. Stripping removes characters to the right of each instance of the specified character, plus the character itself. The stripping begins with the rightmost separator character.

Use the `separator` statement with the `count` statement to determine which characters in a client name are stripped. If the specified number of separator characters (count) exceeds the actual number of separator characters in the client name, stripping stops at the last available separator character.

Options

special-character—Character used to identify where to start the stripping of characters in a client name.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Obtaining Username and Role Information Through Firewall Authentication

seq-number-validated

IN THIS SECTION

- Syntax | 16594
- Hierarchy Level | 16595
- Description | 16595
- Required Privilege Level | 16595
- Release Information | 16595

Syntax

```
seq-number-validated;
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

When the command seq-number-validated is configured, GTP ALG check the sequence number in the GTP-U packet. If it is not configured, the sequence check is disabled by default.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp profile profile-name]` hierarchy level is replaced by `[edit security gtp profile profile-name]`.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

sequence-check-required

IN THIS SECTION

- Syntax | 16596
- Hierarchy Level | 16596
- Description | 16596
- Required Privilege Level | 16596
- Release Information | 16596

Syntax

```
sequence-check-required;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit tcp-
options]
```

Description

Enable sequence check per policy. The sequence-check-required value overrides the global value no-sequence-check.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

sequence-number (Security IDP ICMP Headers)

IN THIS SECTION

- Syntax | 16597
- Hierarchy Level | 16597
- Description | 16597
- Options | 16597
- Required Privilege Level | 16598
- Release Information | 16598

Syntax

```
sequence-number {
    match (equal | greater-than | less-than | not-equal);
    value sequence-number;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol icmp],

[edit security idp custom-attack attack-name attack-type signature protocol icmpv6]
```

Description

Specify the sequence number of the packet. This number identifies the location of the request/reply in relation to the entire sequence.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`sequence-number`*—Match a decimal value.

- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3. Statement modified in Junos OS Release 12.3X48-D25 to add ICMPv6 protocol support.

sequence-number (Security IDP TCP Headers)

IN THIS SECTION

- Syntax | 16598
- Hierarchy Level | 16599
- Description | 16599
- Options | 16599
- Required Privilege Level | 16599
- Release Information | 16599

Syntax

```
sequence-number {
    match (equal | greater-than | less-than | not-equal);
    value sequence-number;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify the sequence number of the packet. This number identifies the location of the data in relation to the entire data sequence.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`sequence-number`*—Match a decimal value.
- **Range:** 0 through 4,294,967,295

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

serial-number

IN THIS SECTION

- Syntax | 16600
- Hierarchy Level | 16600
- Description | 16600
- Options | 16600
- Required Privilege Level | 16600

Release Information | 16600

Syntax

```
serial-number serial-number;
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

In a preprovisioned Virtual Chassis or Virtual Chassis Fabric (VCF), specify the serial number of each member switch to be included in the configuration. If you do not include the serial number within the configuration, the switch cannot be recognized as a member of a preprovisioned configuration. Serial number values are case-sensitive.

In an EX8200 Virtual Chassis configuration, specify the serial number of each XRE200 External Routing Engine and each EX8200 member switch to be included in the Virtual Chassis configuration. If you do not include the serial number within the Virtual Chassis configuration, the external Routing Engine or switch cannot be recognized as a member of the configuration.

Options

serial-number—Permanent serial number for the external Routing Engine or for the member switch.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric

Configuring an EX2300, EX3400, or EX4300 Virtual Chassis

Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis

Configuring an EX9200 Virtual Chassis

Configuring an EX4650 or a QFX Series Virtual Chassis

serial-number (Junos Fusion)

IN THIS SECTION

- Syntax | **16601**
- Hierarchy Level | **16601**
- Description | **16602**
- Default | **16602**
- Options | **16602**
- Required Privilege Level | **16602**
- Release Information | **16602**

Syntax

```
serial-number serial-number;
```

Hierarchy Level

```
[edit chassis satellite-management fpc slot-id]
```

Description

Bind the specified FPC slot ID to a satellite device based on the serial number.

A Junos Fusion provides two methods of assigning an FPC identifier: Unique ID-based FPC identification and connectivity-based FPC identification. Unique ID-based FPC identification maps an FPC slot ID to a satellite device's MAC address or serial number, while unique ID-based FPC identification maps an FPC slot ID to a cascade port. This statement is used to assign an FPC ID using unique ID-based FPC identification by mapping the FPC slot ID to the satellite device's serial number.

In a Junos Fusion, each satellite device must be mapped to an FPC identifier (FPC ID). The FPC ID is used for Junos Fusion configuration, monitoring, and maintenance. Interface names—which are identified using the *type-fpc* / *pic* / *port* format—use the FPC ID as the *fpc* variable when the satellite device is participating in a Junos Fusion. For instance, built-in port 2—a Gigabit Ethernet interface on a satellite device that is using 101 as its FPC ID—uses ge-101/0/2 as its interface name.

For additional information on the role of FPC slot IDs in a Junos Fusion, see Understanding Junos Fusion Enterprise Components or *Understanding Junos Fusion Provider Edge Components*.

If the serial number that is configured using this statement does not match the serial number of the satellite device, the device is not converted into a satellite device.

Default

No FPC slot IDs are associated with satellite devices, by default.

Options

serial-number Specifies the serial number of the satellite device.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

Statement introdcued in Junos OS Release 16.1R1 for Junos Fusion Enterprise.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

serial-number (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 16603
- Hierarchy Level | 16603
- Description | 16603
- Options | 16604
- Required Privilege Level | 16604
- Release Information | 16604

Syntax

```
serial-number serial-number;
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

As part of the preprovisioned configuration for an MX Series Virtual Chassis, specify the chassis serial number of each MX Series member router in the Virtual Chassis configuration. If you do not correctly specify a router's serial number in the preprovisioned configuration, the software does not recognize that router as a member of the Virtual Chassis.

Options

serial-number—Alphanumeric string that represents the chassis serial number of each member router in the Virtual Chassis configuration. The chassis serial number is located on a label affixed to the side of the MX Series chassis. You can also obtain the router's chassis serial number by issuing the `show chassis hardware` command.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Preprovisioned Member Information for a Virtual Chassis

Example: Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis

serial-options

IN THIS SECTION

- Syntax | **16605**
- Hierarchy Level | **16606**
- Description | **16606**
- Required Privilege Level | **16606**
- Release Information | **16606**

Syntax

```
serial-options {
    clock-rate rate;
    clocking-mode (dce | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dsr-polarity (negative | positive);
    dte-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dtr-circuit (balanced | unbalanced);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    indication-polarity (negative | positive);
    line-protocol protocol;
    loopback (dce-local | dce-remote | local | remote);
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);
    transmit-clock invert;
}
```

Hierarchy Level

```
[edit interfaces se-pim/0/port]
```

Description

Configure serial-specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced prior to Junos OS Release 7.4.

RELATED DOCUMENTATION

Serial Interfaces Overview

no-concatenate

server (Junos node slicing)

IN THIS SECTION

- Syntax | 16607
- Hierarchy Level | 16607
- Description | 16607
- Options | 16607
- Required Privilege Level | 16608

Release Information | 16608

Syntax

```
server {
    interfaces {
        cb0 cb0-interface;
        cb1 cb1-interface;
        jdm-management jdm-management-interface;
        vnf-management gnf-management-interface;
    }
}
```

Hierarchy Level

[edit]

Description

Applicable only to Junos node slicing set up using external server model.

Configure the server interfaces for the JDM and GNFs. These include a JDM management interface, a GNF management interface, and two server interfaces that are connected to the MX Series router.

Options

cb0 *cb0-interface*	The server interface that is connected to the control board 0 of the MX Series router.
cb1 *cb1-interface*	The server interface that is connected to the control board 1 of the MX Series router.
jdm-management *jdm-management-interface*	The server interface to be used as the JDM management interface.
vnf-management *gnf-management-interface*	The server interface to be used as GNF management interface.

Required Privilege Level

root

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

virtual-network-functions (Junos node slicing) | 30773

server

IN THIS SECTION

- Syntax | 16608
- Hierarchy Level | 16608
- Description | 16609
- Required Privilege Level | 16609
- Release Information | 16609

Syntax

```
server;
```

Hierarchy Level

```
[edit interfaces pp0 unit logical-unit-number pppoe-options],
[edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number pppoe-options]
```

Description

Configure the router to operate in the PPPoE server mode. Supported on M120 and M320 Multiservice Edge Routers and MX Series 5G Universal Routing Platforms operating as access concentrators.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring the PPPoE Server Mode

server-address

IN THIS SECTION

- Syntax | 16610
- Hierarchy Level | 16610
- Description | 16610
- Default | 16610
- Options | 16610
- Required Privilege Level | 16610
- Release Information | 16610

Syntax

```
server-address  ip-address;
```

Hierarchy Level

```
[edit interfaces interface-nameunit logical-unit-numberfamily inet dhcp]
```

Description

Specify the address of the DHCP server that the client should accept DHCP offers from. If this option is included in the DHCP configuration, the client accepts offers only from this server and ignores all other offers.

Default

The client accepts the first offer it receives from any DHCP server.

Options

ip-address	DHCP server address.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring a DHCP Client

interfaces (EX Series switches) | **9135**

unit

family

server-address (dhcp-client)

IN THIS SECTION

- Syntax | 16611
- Hierarchy Level | 16611
- Description | 16611
- Options | 16611
- Required Privilege Level | 16612
- Release Information | 16612

Syntax

```
server address ip-address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcp-client]
```

Description

Specify the preferred DHCP server address that is sent to DHCP clients.

Options

ip-address	DHCP server address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

Starting with Junos OS Release 17.3R1, on all SRX Series Firewalls and vSRX Virtual Firewall instances, the CLI option `dhcp-clinet` at `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy is changed to `dhcp` to keep consistent with other Junos platforms. There is no change in the functionality with the changed CLI option `dhcp`.

server-address (Hosted Services)

IN THIS SECTION

- Syntax | 16612
- Hierarchy Level | 16612
- Description | 16613
- Options | 16613
- Required Privilege Level | 16613
- Release Information | 16613

Syntax

```
server-address ipv4-address;
```

Hierarchy Level

```
[edit services hosted-services server-profile server-profile-name]
```

Description

Configure the server address where sampled packets are sent.

Options

ipv4-address—IPv4 address of the server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

server (Batch Commits)

IN THIS SECTION

- Syntax | 16614
- Hierarchy Level | 16614
- Description | 16614
- Options | 16614
- Required Privilege Level | 16615
- Release Information | 16615

Syntax

```
server {
    commit-interval number-of-seconds-between-commits;
    commit-schedule-profile;
    days-to-keep-error-logs days-to-keep-error-log-entries;
    maximum-aggregate-pool maximum-number-of-commits-to-aggregate;
    maximum-entries number-of-entries;
    redirect-completion-status;
    retry-attempts;
    retry-interval;
    traceoptions {
        file filename;
        files number;
        flag (all | batch | commit-server | configuration);
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
}
```

Hierarchy Level

```
[edit system commit]
```

Description

Configure the system commit to occur in batches. Configure parameters for aggregating and saving batch commits.

Options

commit-interval	Configure the interval in seconds between commits.
days-to-keep-error-logs	Configure the number of days to keep log entries. Valid range is from 1 to 366 days.
maximum-aggregate-pool	Configure the maximum number of commits to aggregate together. The valid range is 1 through 4294967295.

maximum-entries	Configure the maximum number of commit entries.
redirect-completion-status	Configure the redirect asynchronous commit status to server configured here.
retry-attempts	Configure the retry attempts for commit failure due to db lock error. The default is 5 retries.
retry-interval	Configure the retry interval in seconds for commit failure. The default is 20 seconds.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Batch Commit Server Properties

traceoptions (Batch Commits) | **29499**

server-certificate (REST API)

IN THIS SECTION

- Syntax | **16616**
- Hierarchy Level | **16616**
- Description | **16616**
- Options | **16616**
- Required Privilege Level | **16616**

Release Information | 16616

Syntax

```
server-certificate local-certificate-identifier;
```

Hierarchy Level

```
[edit system services rest https]
```

Description

Set the server certificate when configuring SSL/TLS mutual authentication.

Options

local-certificate-identifier	The server certificate. This must be set when configuring SSL/TLS mutual authentication.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

server-certificate (Services)

IN THIS SECTION

- Syntax | 16617
- Hierarchy Level | 16617
- Description | 16617
- Options | 16617
- Required Privilege Level | 16617
- Release Information | 16618

Syntax

```
server-certificate server-certificate;
```

Hierarchy Level

```
[edit services ssl termination profile profile-name]
```

Description

Specify the local certificate identifier.

Options

`server-certificate`—Specify the name of the local certificate identifier.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. This statement is supported.

server-certificate-subject

IN THIS SECTION

- Syntax | 16618
- Hierarchy Level | 16618
- Description | 16618
- Required Privilege Level | 16619
- Release Information | 16619

Syntax

```
server-certificate-subject subject;
```

Hierarchy Level

```
[edit services unified-access-control infranet-controller hostname]
```

Description

Optionally specify the full subject name of the certificate that the SRX Series Firewall should use to validate the IC Series device's server certificate.

Use this statement when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

ca-profile (Services) | 3160

password (Services) | 12246

server-cluster (Security Group VPN Server)

IN THIS SECTION

- Syntax | 16619
- Hierarchy Level | 16620
- Description | 16620
- Options | 16620
- Required Privilege Level | 16620
- Release Information | 16621

Syntax

```
server-cluster {
    ike-gateway gateway-name;
    retransmission-period seconds;
    server-role (root-server | sub-server);
}
```

Hierarchy Level

```
[edit security group-vpn server group name]
```

Description

Configure the Group Domain of Interpretation (GDOI) group controller/key server (GCKS) cluster for the specified group. All servers in a group VPN server cluster must be SRX Series Firewalls.

Options

ike-gateway ***gateway-name***

(Required) Specify the name of the IKE gateway for the local device in the group server cluster. IKE gateways are configured at the [`edit security group-vpn server ike`] hierarchy level.

If the local device is a root-server, the IKE gateway name must be a sub-server in the cluster; up to four sub-server IKE gateways can be specified.

If the local device is a sub-server, the IKE gateway name must be the root-server.

retransmission-period ***seconds***

(Optional) Specify the time after which the root-server retransmits a `cluster-update` message if it has not received an acknowledgement from a sub-server.

- **Range:** 2 to 60 seconds.
- **Default:** 10 seconds.

server-role

(Required) Assign the role of the local device in the group server cluster, either `root-server` or `sub-server`. Only one device in the cluster can be configured as the root-server. You can configure up to four other devices as a sub-server in a group server cluster.

You must ensure that there is only one root-server at any time for a group VPN server cluster.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Group VPNv2 Overview

Understanding Group VPNv2 Server Clusters

group (Security Group VPN) | **7894**

server-connectivity

IN THIS SECTION

- Syntax | **16621**
- Hierarchy Level | **16621**
- Description | **16622**
- Options | **16622**
- Required Privilege Level | **16622**
- Release Information | **16622**

Syntax

```
server-connectivity  (block | log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated profile profile-name
fallback-settings]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name fallback-
settings]
```

```
[edit security utm feature-profile web-filtering  juniper-enhanced profile profile-name fallback-
settings]
```

Description

Fallback settings tell the system how to handle errors. This is the action that occurs when a request fails for this reason.

Options

- block—Log the error and deny the traffic
- log-and-permit—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The surf-control feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5 .

Statement introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The surf-control feature is not supported from Junos OS Release 15.1X49-D10 onwards.

server (DHCP and BOOTP Relay Agent)

IN THIS SECTION

- Syntax | 16623
- Hierarchy Level | 16623
- Description | 16623
- Options | 16623
- Required Privilege Level | 16624
- Release Information | 16624

Syntax

```
server address {
    logical-system logical-system-name <routing-instance [ <default> routing-instance-names ]>;
    routing-instance [ <default> routing-instance-names ];
}
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface (interface-name | interface-group)]
```

Description

Configure the router or switch to act as a DHCP and BOOTP relay agent. The device forwards all broacast requests within the configured subnet to all configured servers in parallel. To support clients on different VRFs, see the "relay-agent-option" on page 14279 statement.

Options

- *address*—One or more addresses of the server.
- `logical-system` *logical-system-name*—(Optional) Logical system of the server.

- `routing-instance` *routing-instance-names*—(Optional) Routing instance name that belong to the DHCP or BOOTP relay agent.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents*
>
> relay-agent-option | **14279**

server (DNS, Port, and TFTP Service)

IN THIS SECTION

- Syntax | **16624**
- Hierarchy Level | **16625**
- Description | **16625**
- Options | **16625**
- Required Privilege Level | **16625**
- Release Information | **16625**

Syntax

```
server address <logical-system logical-system-name> <routing-instance routing-instance-name>;
```

Hierarchy Level

```
[edit forwarding-options helpers domain],
[edit forwarding-options helpers domain interface interface-name],
[edit forwarding-options helpers port port-number],
[edit forwarding-options helpers port port-number interface interface-name],
[edit forwarding-options helpers tftp],
[edit forwarding-options helpers tftp interface interface-name]
```

Description

Specify the DNS or TFTP server for forwarding DNS or TFTP requests, or specify a destination server address for forwarding LAN broadcast packets as unicast traffic for a custom-configured UDP port.

When configuring port helpers, in releases prior to Junos OS Release 17.2, only one server can be specified for a given port. For Junos OS Release 17.2 and later, multiple servers can be specified for a given port at the global or interface-specific level. When multiple servers are specified, the same packet, with the originator IP address and port requests, is forwarded to the different configured servers; the payload of the UDP packet is not modified.

Options

address—IP address of the server.

logical-system ***logical-system-name***—(Optional) Logical system name of the server.

routing-instance [***routing-instance-names*** **]**—(Optional) Set the routing instance name or names that belong to the DNS server or TFTP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for port helpers in Junos OS Release 17.2R1 for EX4300 switches.

Support for multiple server instances for a given port introduced in Junos OS Release 17.2 for MX Series routers.

Support for multiple server instances for a given port introduced in Junos OS Release 17.3R1 for EX9200 switches.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding

Configuring Port-based LAN Broadcast Packet Forwarding

server-duid-type (DHCP Local Server)

IN THIS SECTION

- Syntax | 16626
- Hierarchy Level | 16626
- Description | 16626
- Options | 16627
- Required Privilege Level | 16627
- Release Information | 16627

Syntax

```
server-duid-type type;
```

Hierarchy Level

```
[edit system services dhcp-local-server dhcpv6]
```

Description

Configure the DHCP Unique Identifier (DUID) type that the DHCPv6 server supports. The server uses the DUID to identify clients for configuration and to group clients into identity associations.

Use this statement to explicitly configure the DUID-LL type. DUID-LL uses a link-layer network interface address on the device as part of the unique identifier. Remove this configuration to return to the default type, DUID-EN. DUID-EN has a vendor-assigned value based on the vendor's enterprise number and an IANA-recognized hardware type code. DUID-EN is not explicitly configurable.

NOTE: The DUID-LLT type is not supported.

Options

type DUID type to support. The only available option is `duid-ll`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

DHCPv6 Local Server Overview

server (Dynamic PPPoE)

IN THIS SECTION

- Syntax | 16628
- Hierarchy Level | 16628
- Description | 16628
- Required Privilege Level | 16628
- Release Information | 16628

Syntax

```
server;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" pppoe-options]
```

Description

In a dynamic profile, configure the router to act as a PPPoE server, also known as a remote access concentrator, when a PPPoE logical interface is dynamically created.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring a PPPoE Dynamic Profile

Subscriber Interfaces and PPPoE Overview

server (Easy EVPN LAG Configuration)

IN THIS SECTION

Syntax | 16629

- Hierarchy Level | 16629
- Description | 16629
- Options | 16630
- Required Privilege Level | 16633
- Release Information | 16633

Syntax

```
server name {
  apply-config-groups config-groups;
  enable-pxe-boot;
  esi-lag-id esi-lag-id;
  interface [ interface ... ];
  mac-vrf-instance instance-id {
      vlan-id-list [ vlan-id-list ... ]
  };
  single-home-id single-home-id;
  vlan-id-list [ vlan-id-list ... ]);
};
```

Hierarchy Level

```
[edit services evpn evpn-vxlan]
```

Description

Specify parameters to generate the configuration statements for servers connected to a provider edge (PE) device in an Ethernet VPN (EVPN) fabric with Virtual Extensible LAN (VXLAN) encapsulation.

The server might be a customer edge (CE) device or end server connected to the EVPN PE device, and can be:

- Single-homed to the PE device.
- Multihomed on an Ethernet segment (identified by an Ethernet segment identifier [ESI]) in a link aggregation group (LAG) to the PE device and a peer PE device in the EVPN fabric.

> **NOTE**: We refer to this type of multihomed Ethernet segment connection as an *ESI LAG*.

You use statements at the `[edit services evpn]` hierarchy level to simplify configuring an EVPN fabric with ESI LAGs in the topology. The feature you enable with statements in this hierarchy is an easy EVPN LAG configuration feature, also called EZ-LAG.

You only need to configure a few statements to provide the parameters required for an EVPN core fabric with multihomed or single-homed customer edge (CE) devices (or end server devices). When you commit the `[edit services evpn]` configuration statements, the device uses a commit script to generate the corresponding standard EVPN fabric configuration statements.

The statements at the `[edit services evpn evpn-vxlan server]` hierarchy level provide the parameters the commit script needs to generate a default configuration for the peer PE to server connections, including configuration for:

- Multihomed server Etehrnet segments, ESI LAG interfaces and member VLANs
- Single-homed server interfaces and member VLANs
- Lightweight leaf to server loop detection on interfaces connecting PE leaf devices to servers

  > **NOTE**: To disable generation of the loop detection configuration statements, you can configure the `no-loop-detect-config` option at the `[edit service evpn defaults-override]` hierarchy level.

- Storm control on PE to server interfaces

  > **NOTE**: To disable generation of the storm control configuration statements, you can configure the `no-storm-control-config` option at the `[edit service evpn defaults-override]` hierarchy level.

Options

(Required) ***name***	Specify a name to identify a server connected to the peer PE device. For example, SERVER_1 or hostA.
apply-config-groups ***config-groups***	Apply the specified configuration groups to the configuration for the server aggregated Ethernet interface bundles that the commit script generates from the statements for this server.

enable-pxe-boot

Enable this connected server to use a preboot execution environment (PXE) boot process.

When you set this option set, the commit script configures LACP `force-up` status on the link to the server. That status ensures the link is UP when the server boots up, which the PXE boot process requires. See *Aggregated Ethernet Interfaces* for more on the LACP `force-up` option and forcing LAG links or interfaces with limited LACP capacity to be up.

(Required) esi-lag-id *esi-lag-id*

Assign an ESI LAG connection ID for this server. With this option, the commit script generates the configuration statements for an ESI LAG connection from the PE device to this multihomed server.

The commit script assigns aggregated Ethernet interfaces (ae*X*) on the peer PE device for ESI LAG connected servers starting with the base number 0. The commit script adds the *esi-lag-id* value to that base. For example, if you configure *esi-lag-id* = 2 for this server, the generated configuration uses interface name ae2.

> **NOTE**: The commit script uses ae0 for the peer to peer PE devices connections, so ae0 is never assigned as a PE-device-to-server ESI LAG interface.

You must configure this option or the `single-home-id` option for each connected server.

- **Range for *esi-lag-id***: 1 through 1024 (which generates interface names ae1 through ae1024)

(Required) interface [*interface-name* ...];

Specify the interface name for the physical link or links from the peer PE device to this multihomed or single-homed server.

Use this option with the `esi-lag-id` option or the `single-home-id` option that identifies this connected server.

mac-vrf-instance *instance-id*

Specify a MAC-VRF instance identifier that hosts this server.

By default, if you don't configure this statement, the commit script generates the configuration for one MAC-VRF instance with 1 as the default MAC-VRF *instance-id*.

If you include the `vlan-id-list` *`vlan-id-list`* option with this statement, the commit script generates configuration statements that associate this MAC-VRF instance with the listed server VLANs.

> **NOTE**: If you don't configure the `mac-vrf-instance` statement, you must configure the `vlan-id-list` *`vlan-id-list`* statement at the `[edit services evpn evpn-vxlan server` *`name`*`]` hierarchy level with at least one member VLAN. The commit script generates configuration statements that associate that list of VLANs with the default MAC-VRF instance (which has *instance-id* 1 and instance name __SERVICES_EVPN_EVPN_VXLAN_MAC_VRF_1).

See the description for the `vlan-id-list` option below for details on how to specify the list of VLAN IDs.

- **Range for *instance-id*:** 1 through 255

(Required) single-home-id *single-home-id*

Assign a single-homed connection ID for this server. With this option, the commit script generates the configuration statements for a connection from the PE device to this single-homed server.

The commit script assigns aggregated Ethernet interface names on the peer PE device for single-homed server connections. The highest `esi-lag-id` index for interfaces to multihomed servers is 1024, which corresponds to interface name ae1024. So to assign interface names for links to single-homed servers, the commit script adds the *single-home-id* value to the base number 1024. As a result, single-homed server interface names start with ae1025 for *single-home-id* = 1 . If you configure *single-home-id* = 2, the generated configuration uses interface name ae1026, and so on.

You must configure this option or the `esi-lag-id` option for each connected server.

- **Range for *single-home-id*:** 1 through 1024 (which generates interface names ae1025 through ae2048)

vlan-id-list [*vlan-id-number* ...]

(Required) List the VLANs this server hosts. Specify one or more VLANs by VLAN ID number. Use spaces to separate multiple individual VLAN IDs (such as [100 200], or use a dash to indicate a contiguous range of VLAN IDs (such as [100-300]). The commit script generates the statements that configure the specified VLANs.

You must configure at least one VLAN ID list for each server. You can configure the server's `vlan-id-list` statement at either of the following hierarchy levels:

- `[edit services evpn evpn-vxlan server` *`name`* `mac-vrf-instance` *`instance-id`*`]`—Associate the listed VLANs with the MAC-VRF instance for the specified instance ID in the generated configuration .

- [edit services evpn evpn-vxlan server *name*]—Associate the listed VLANs with:
  - The default MAC-VRF instance ID, which is 1. The commit script uses MAC-VRF instance ID 1 when you don't specify any mac-vrf-instance statements for this server.
  - All MAC-VRF instance IDs that you configure using the mac-vrf-instance statement for this server.
- **Range for *vlan-id-number* values:** 1 through 4094

See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Easy EVPN LAG (EZ-LAG) Configuration

server-group

IN THIS SECTION

- Syntax | 16634
- Hierarchy Level | 16634
- Description | 16634
- Options | 16634
- Required Privilege Level | 16635
- Release Information | 16635

Syntax

```
server-group {
    server-group-name {
        server-ip-address;
    }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6]
```

Description

Specify the name of a group of DHCP server addresses for use by the extended DHCP relay agent. Apply the group with the `active-server-group` statement globally for all interfaces or for a named group of interfaces configured with the `group` statement. This mechanism enables you to apply different DHCP relay configurations for different groups of servers, with a common configuration for the servers within a server group.

Options

server-group-name Name of the group of DHCP or DHCPv6 server addresses.

server-ip-address IP address of the DHCP server belonging to this named server group. Use IPv6 addresses when configuring DHCPv6 support. Starting in Junos OS Release 18.4R1, you can configure up to 32 server IP addresses per group for DHCPv4 servers. In earlier releases, you can configure only up to 5 server IP addresses for DHCPv4 servers. For DHCPv6

servers, you can configure only up to 32 addresses in all releases. The configuration fails commit check if you configure more than the maximum number of server addresses.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups

server (icap-redirect profile)

IN THIS SECTION

- Syntax | 16636
- Hierarchy Level | 16636
- Description | 16636
- Options | 16636
- Required Privilege Level | 16637
- Release Information | 16637

Syntax

```
server name {
    authorization {
        authorization-type authorization-type;
        credentials (ascii ascii | base64 base64);
    }
    host host;
    port port;
    reqmod-uri reqmod-uri;
    respmod-uri respmod-uri;
    routing-instance ri-name;
    sockets sockets;
    tls-profile tls-profile;
}
```

Hierarchy Level

```
[edit services icap-redirect profile name],
[edit logical-systems logical-system-name services icap-redirect profile name]
```

Description

Configure the ICAP server details.

When you configure the ICAP redirect service on SRX Series Firewalls, you must configure the ICAP server details. ICAP server configuration allows you to define the settings required to process request messages, response messages, authorization, and so on. You can also specify an SSL profile in the ICAP server configuration that enables you to secure the connection to the ICAP server.

You can configure up to two ICAP servers.

Options

name ICAP server name.

host ICAP server hostname or IP address.

port ICAP server listening port, default port is reached according to the protocol defined.

- **Default:** 1344
- **Range:** 1025 through 65535

route-instance Virtual router that is used for launching the service.

reqmod-uri Path to the service that handles Request Modification (REQMOD) requests.

respmod-uri Path to the service that handles Response Modification (RESPMOD) requests.

sockets Number of connections to create the ICAP service.

- **Default:** 8
- **Range:** 1 through 64

tls-profile SSL profile configured to provide a secure connection to the ICAP server.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.1R1.

The logical system option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Example: Configuring ICAP Redirect Service on SRX Series Firewalls

server-id-override

IN THIS SECTION

Syntax | 16638

- Hierarchy Level | **16638**
- Description | **16638**
- Required Privilege Level | **16639**
- Release Information | **16639**

Syntax

```
server-id-override;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82],
[edit forwarding-options dhcp-relay group group-name relay-option-82],
[edit logical-systems logical-system-name forwarding-options dhcp-relay relay-option-82],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name relay-
option-82],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay relay-option-82],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name relay-option-82],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
relay-option-82]
```

Description

Add the suboptions `link-selection` and `server-id override` to option-82 information for DHCP packets relayed to the server.

In network configurations where a firewall on the broadband network gateway (BNG) is between the DHCPv4 relay agent and the DHCP server, only the BNG loopback address passes through the BNG firewall. In that case, DHCP unicast packets do not pass through the firewall and and discarded. To enable DHCP unicast packets to pass through the BNG firewall, use the `relay-source` configuration statement to configure the DHCP relay agent to use the loopback address as the source address in IP headers and DHCP messages. In this case, this configuration statement adds the `link-selection` suboption

to provide the DHCP server information to locate the correct address pool for the DHCP client and adds the `server-id override` suboption to set the server ID option.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Configuring the DHCP Relay Agent Source Address to Enable DHCP Packets to Pass Through a Firewall*

server-identifier

IN THIS SECTION

- Syntax | 16639
- Hierarchy Level | 16640
- Description | 16640
- Default | 16640
- Options | 16640
- Required Privilege Level | 16640
- Release Information | 16640

Syntax

```
server-identifier address;
```

Hierarchy Level

```
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]
```

Description

Configure a server identifier. The identifier can be used to identify a DHCP server in a DHCP message. It can also be used as a destination address from clients to servers (for example, when the boot file is set, but not the boot server).

Servers include the server identifier in **DHCPOFFER** messages so that clients can distinguish between multiple lease offers. Clients include the server identifier in **DHCPREQUEST** messages to select a lease and indicate which offer is accepted from multiple lease offers. Also, clients can use the server identifier to send unicast request messages to specific DHCP servers to renew a current lease.

This address must be a manually assigned, static IP address. The server cannot send a request and receive an IP address from itself or another DHCP server.

Default

If no server identifier is set, the DHCP server sets the server identifier based on the primary interface address used by the server to receive a client request. For example, if the client sends a DHCP request and the server receives it on `fe-0/0/0` and the primary interface address is `10.1.1.1`, then the server identifier is set to `10.1.1.1`.

Options

address—IPv4 address of the server. This address must be accessible by all clients served within a specified range of addresses (based on an address pool or static binding).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

server-inactivity-timeout

IN THIS SECTION

- Syntax | 16641
- Hierarchy Level | 16641
- Description | 16641
- Options | 16641
- Required Privilege Level | 16642
- Release Information | 16642

Syntax

```
server-inactivity-timeout minutes;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

The maximum time the Two-Way Active Measurement Protocol (TWAMP) server has to finish the TWAMP control protocol negotiation.

Options

minutes Number of minutes the TWAMP server has to finish the TWAMP control protocol negotiation. Configure 0 to disable.

- **Default:** 15 minutes
- **Range:** 1 through 30 minutes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

server (Junos OS)

IN THIS SECTION

- Syntax | 16642
- Hierarchy Level | 16643
- Description | 16643
- Options | 16643
- Required Privilege Level | 16643
- Release Information | 16643

Syntax

```
server {
    authentication-mode none;
    client-list list-name {
        address address <routing-instance [instance-name...]>;
    }
    light {
        port number;
```

```
    }
    max-connection-duration hours;
    maximum-connections count;
    maximum-connections-per-client count;
    maximum-sessions count;
    maximum-sessions-per-connection count;
    port number;
    server-inactivity-timeout seconds;
    tcp-keepcnt count;
    tcp-keepidle seconds;
    tcp-intvl seconds;
}
```

Hierarchy Level

```
[edit services rpm twamp]
```

Description

Defines the Two-Way Active Measurement Protocol (TWAMP) server configurations. This configuration describes a TWAMP server on the router that enables a TWAMP client to connect to the server using any media interface IP address.

Options

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

`light` option added in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

server (Junos OS Evolved)

IN THIS SECTION

- Syntax | 16644
- Hierarchy Level | 16645
- Description | 16645
- Options | 16645
- Required Privilege Level | 16645
- Release Information | 16646

Syntax

```
server {
    light {
        offload-type (none | inline-timestamp | pfe-timestamp);
        port number;
        stateful sequence cache-size number-of-entries;
    }
    managed {
        client-limit limit;
        client-list list-name {
            address address <routing-instance [instance-name...]>;
        }
        control-inactivity-timeout seconds;
        control-maximum-duration seconds;
        control-per-client-limit number;
        offload-type (none | inline-timestamp | pfe-timestamp);
        port number;
        stateful-sequence;
        test-per-client-limit limit;
```

```
    }
}
```

Hierarchy Level

```
[edit services monitoring twamp]
```

Description

Configure the Two-Way Active Measurement Protocol (TWAMP) server. The server listens in on all routing instances.

Options

client-limit *limit*	Maximum number of TWAMP clients Range: 0 to 1000 clients
control-inactivity-timeout *seconds*	Inactivity timeout on control connection Range: 0 to 86400 seconds
control-maximum-duration *seconds*	Hard limit on control connection duration Range: 0 to 86400 seconds
control-per-client-limit *number*	Maximum number of TWAMP control connections per client Range: 0 to 1000 control connections per client
test-per-client-limit *limit*	Maximum number of TWAMP test sessions per client Range: 0 to 1000 test sessions per client

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 20.3R1.

`offload-type` and `stateful-sequence` options introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

server-member-communication (Security Group VPN Server)

IN THIS SECTION

- Syntax | 16646
- Hierarchy Level | 16647
- Description | 16647
- Options | 16647
- Required Privilege Level | 16647
- Release Information | 16648

Syntax

```
server-member-communication {
    certificate certificate-id;
    communication-type (unicast);
    encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
    lifetime-seconds seconds;
    number-of-retransmission number;
    retransmission-period seconds;
    sig-hash-algorithm (sha-256 | sha-384);
}
```

Hierarchy Level

```
[edit security group-vpn server group name]
```

Description

Enable and configure server to member communication. When these options are configured, group members receive new keys before current keys expire. Starting with Junos OS Release 15.1X49-D80, the minimum value that you can configure for the `lifetime-seconds` option is 300 seconds instead of 180 seconds.

Options

- `certificate` *`certificate-id`*—Specify the certificate identification. Only RSA keys are supported.
- `communication-type`—Configure `unicast` (the default).
- `encryption-algorithm`—Encryption used for communications between the group server and group member. Specify `aes-128-cbc`, `aes-192-cbc`, or `aes-256-cbc`.
- `lifetime-seconds` *`seconds`*—Lifetime, in seconds, of the key encryption key (KEK). Specify a value from 300 to 86,400. The default is 3600 seconds.
- `number-of-retransmission` *`number`*—For unicast communications, the number of times the group server retransmits messages to a group member when there is no reply. Specify a value from 0 to 60. The default is 2.
- `retransmission-period` *`seconds`*—The time period between a transmission and the first retransmission when there is no reply from the group member. Specify a value from 2 to 60. The default is 10 seconds.
- `sig-hash-algorithm`—Authentication algorithm used to authenticate the group member to the group server. Specify `sha-256` or `sha-384`.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2

RELATED DOCUMENTATION

Group VPNv2 Overview

group (Security Group VPN) | 7894

server (NTP)

IN THIS SECTION

- Syntax | 16648
- Hierarchy Level | 16649
- Description | 16649
- Options | 16649
- Required Privilege Level | 16650
- Release Information | 16650

Syntax

```
server address {
    key key;
    nts remote-identity {
      hostname hostname;
      distinguished-name(container container | wildcard wildcard);
       }
    prefer;
    routing-instance routing-instance;
    version version;
}
```

Hierarchy Level

```
[edit system ntp]
```

Description

For NTP, configure the local router or switch to operate in client mode with the remote system at the specified *address*. In this mode, the local router or switch can be synchronized with the remote system, but the remote system can never be synchronized with the local router or switch.

If the NTP client time drifts so that the difference in time from the NTP server exceeds 128 milliseconds, the client is automatically stepped back into synchronization. If the offset between the NTP client and server exceeds the 1000-second threshold, the client still synchronizes with the server, but it also generates a system log message noting that the threshold was exceeded.

Options

address Address of the remote system. You must specify an address, not a hostname.

key *key-number* (Optional) Use the specified key number to encrypt authentication fields in all packets sent to the specified address.

- **Range:** Any unsigned 32-bit integer

nts Enables Network Time Security (NTS), which uses Transport Layer Security (TLS) protocol and Authenticated Encryption with Associated Data (AEAD) to obtain network time in an authenticated manner to the users. Specified server must also support the NTS feature when you enable NTS on a client device.

remote-identity (Optional) Configure the options to verify the NTS enabled server's identity.

hostname *hostname* (Optional) Specify the fully qualified domain name (FDQN) of the server. This hostname should match with the FQDN of the server in its local certificate's **Subject Alternative Name** field.

distinguished-name (container *container* | wildcard *wildcard*) (Optional) You can verify the identity of NTS-enabled server by its distinguished name in the server certificate.

container *container*
: If you use the *container* option, the order of the fields in the `distinguished-name` and their values must exactly match with the values in the server certificate.

wildcard *wildcard*
: If you use the *wildcard* option, the order of the fields in the `distinguished-name` does not matter.

prefer
: (Optional) Mark the remote system as preferred host, which means that if all other things are equal, this remote system is chosen for synchronization among a set of correctly operating systems.

routing-instance *routing-instance*
: (Optional) Routing instance through which the server is reachable.

> **NOTE**: You can use the routing instance name `mgmt_junos` to enable access to NTP servers through the management interface.

When configuring the NTP service in the management VRF (`mgmt_junos`), you must configure at least one IP address on a physical or logical interface within the default routing instance and ensure that this interface is up in order for the NTP service to work with the mgmt_junos VRF.

version *value*
: (Optional) Specify the version number to be used in outgoing NTP packets.

- **Range:** 1 through 4
- **Default:** 4

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the NTP Time Server and Time Services

server (pcp)

IN THIS SECTION

- Syntax | 16651
- Hierarchy Level | 16651
- Description | 16652
- Options | 16652
- Required Privilege Level | 16653
- Release Information | 16653

Syntax

```
server server-name {
    ipv4-address ipv4-address;
    ipv6-address ipv6-address;
    long-lifetime-error long-lifetime-error;
    mapping-lifetime-max   mapping-lifetime-max;
    mapping-lifetime-min   mapping-lifetime-min;
    max-mappings-per-client max-mappings-per-client;
    nat-options {
        pool pool-name ;
    }
    pcp-options {
        prefer-failure;
        third-party;
    }
    short-lifetime-error   short-lifetime-error;
    softwire-concentrator softwire-concentrator-name;
}
```

Hierarchy Level

```
[edit services pcp]
```

Description

Configure PCP server options. PCP enables hosts to operate servers for a long time (as in the case of a webcam) or a short time (for example, while playing a game or on a phone call) when behind a NAT device, including when behind a carrier-grade NAT operated by their ISP. PCP enables applications to create mappings from an external IP address and port to an internal IP address and port.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC. Starting in Junos OS Release 20.1R1, PCP is also supported for Next Gen Services.

Options

ipv4-address IPv4 address of the PCP server.

ipv6-address IPv6 address of the PCP server.

long-lifetime-error Time limit for generating long lifetime errors.

- **Default:** 1800 seconds
- **Range:** 900 through 18,000 seconds

mapping-lifetime-max Maximum lifetime, in seconds, for PCP mapping. If the PCP client requests a lifetime less than the maximum configured, the server will assign the maximum lifetime and respond accordingly.

- **Default:** 86,400 seconds
- **Range:** 3600 through 4294667 seconds

mapping-lifetime-min Minimum lifetime, in seconds, for PCP mapping. If a PCP client requests a lifetime less than the minimum configured, the server will assign a minimum lifetime and respond accordingly.

- **Default:** 300 seconds
- **Range:** 120 through 3600 seconds

max-mappings-per-client Maximum number of PCP mappings that the PCP client can request.

- **Default:** 32
- **Range:** 1 through 32

pool-name	Name of the NAT pool to use for PCP mapping. You can identify multiple pools. If you do not specify a NAT pool for mapping, the Junos OS performs a partial rule match based on the source IP, source port, and protocol, and the Junos OS uses the NAT pool configured for the first matching rule to allocate mappings for PCP.
prefer-failure	Generate an error message when the PCP client requests a specific IP address or port that is not available, rather than assigning another available address from the NAT pool.
short-lifetime-error	Time limit for generating short lifetime errors. • **Default:** 30 seconds • **Range:** 15 through 300 seconds
softwire-concentrator-name	Softwire concentrator name whose softwire-address is used in creating PCP mappings. The PCP server address must be the same as the softwire-concentrator address.
third-party	Enable third-party requests by the PCP client.

The other statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

server-profile

IN THIS SECTION

- Syntax | 16654
- Hierarchy Level | 16654
- Description | 16654
- Options | 16654
- Required Privilege Level | 16655
- Release Information | 16655

Syntax

```
server-profile server-profile-name {
    client-address ipv4-address;
    server-address ipv4-address;
}
```

Hierarchy Level

```
[edit services hosted-services]
```

Description

Configure the server profile.

Options

server-profile-name—Name to apply to this server profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

server-profile (Active Flow Monitoring)

IN THIS SECTION

- Syntax | 16655
- Hierarchy Level | 16655
- Description | 16656
- Options | 16656
- Required Privilege Level | 16656

Syntax

```
server-profile server-profile-name;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring instance instance-name family (inet | inet6 | mpls)
output]
```

Description

Specify the name of a server profile. This profile specifies a host where sampled traffic is sent.

Options

server-profile-name	Specify the name of a server profile configured at the [`edit services hosted-services server-profile` *`server-profile-name`*] hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

hosted-services | 8178

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

server (RADIUS Accounting)

IN THIS SECTION

- Syntax | 16657
- Hierarchy Level | 16657
- Description | 16657
- Options | 16657
- Required Privilege Level | 16659
- Release Information | 16659

Syntax

```
server {
    server-address {
        accounting-port port-number;
        accounting-retry number;
        accounting-timeout seconds;
        dynamic-request-port number;
        max-outstanding-requests value;
        port number;
        preauthentication-port number;
        preauthentication-secret secret;
        retry number;
        routing-instance routing-instance-name;
        secret password;
        source-addresssource-address;
        timeout seconds;
    }
```

Hierarchy Level

```
[edit system accounting destination radius]
```

Description

Configure RADIUS accounting.

To configure multiple RADIUS accounting servers, include multiple `server` *server-address* statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

NOTE: The `[edit system accounting]` hierarchy is not available on QFabric systems.

Options

server-address Specify the IPv4 or IPv6 address of the RADIUS authentication server.

accounting-port *port-number*

Configure the accounting port number on which to contact the RADIUS server.

- **Range:** 1 through 65,335
- **Default:** 1813

accounting-retry *number*

Configure the number of accounting retry attempts.

- **Range:** 0 through 100 attempts
- **Default:** 0

accounting-timeout *seconds*

Configure the accounting request timeout period.

- **Range:** 0 through 1000 seconds
- **Default:** 0

dynamic-request-port *number*

Configure the RADIUS client dynamic request port number

- **Range:** 1 through 65535
- **Default:** 3799

max-outstanding-requests *value*

Configure the maximum number of outstanding requests in flight to the server.

- **Range:** 0 through 2000 requests
- **Default:** 1000 requests

port *port-number*

Configure the port number on which to contact the RADIUS server.

- **Range:** 1 through 65,335
- **Default:** 1812 (as specified in RFC 2865)

preauthentication-port *number*

Configure the RADIUS server preauthentication-port number.

- **Range:** 1 through 65535

preauthentication-secret *secret*

Configure the shared secret with the RADIUS server; it can include spaces if the character string is enclosed in quotation marks. The secret used by the local device must match that used by the RADIUS server.

retry *value*

Configure the number of times that the device is allowed to try to contact a RADIUS authentication server.

- **Range:** 1 through 100
- **Default:** 3

routing-instance *routing-instance-name*

Configure the routing instance name for the management routing instance. In the case of configuring the non-default management instance, use the value `mgmt_junos`. that is `mgmt_junos`. Configuring this option along with the `management-instance` statement enables authentication processes (for example, RADIUS and TACACS+) to use the non-default management routing instance for packet traffic.

> **NOTE**: You must also define the mgmt_junos routing instance under the `[edit routing-instances]` hierarchy level.

If you do not configure the mgmt_junos instance under the `[edit routing-instances]` hierarchy level and configure it only under `tacplus-server` or `radius-server`, the commit will fail.

secret *password*

(Required) Configure the password (shared secret) to use with the RADIUS server; it can include spaces if the character string is enclosed in quotation marks. The secret password used by the local device must match that used by the RADIUS server.

source-address *source-address*

Configure a valid IPv4 or IPv6 address configured on one of the device's interfaces.

timeout *seconds*

Configure the amount of time the local device waits to receive a response from a RADIUS server.

- **Range:** 1 through 1000 seconds
- **Default:** 3 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

`routing-instance` introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Configuring RADIUS System Accounting

management-instance | **10309**

Management Interface in a Non-Default Instance

Configuring RADIUS Server Authentication

Configuring RADIUS Authentication (QFX Series or OCX Series)

server (RADIUS Network Elements)

IN THIS SECTION

- Syntax | **16660**
- Hierarchy Level | **16660**
- Description | **16661**
- Options | **16661**
- Required Privilege Level | **16661**
- Release Information | **16661**

Syntax

```
server name {
    priority priority;
}
```

Hierarchy Level

```
[edit access radius network-element name]
```

Description

Configure a RADIUS server for the network element, which is a load-balanced group of RADIUS servers providing policy management for TDF subscribers. The RADIUS server must already be defined at the `[edit access radius]` hierarchy level. You can configure multiple RADIUS servers under a network element.

Options

name Name of the RADIUS server.

The remaining statement is described separately.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- *Configuring RADIUS Network Elements*
- *Understanding Network Elements*
- *Understanding How a RADIUS Server Controls Policy and Charging Control Rules*
- *Configuring RADIUS Servers*

server-reject-bridge-domain | server-reject-vlan

IN THIS SECTION

- Syntax | **16662**
- Hierarchy Level | **16662**

- Description | 16662
- Default | 16662
- Options | 16663
- Required Privilege Level | 16663
- Release Information | 16663

Syntax

```
(server-reject-bridge-domain bridge-domain | server-reject-vlan (vlan-id | vlan-name))  {
    block-interval block-interval;
    eapol-block;
    gbp-tag gbp-tag
}
```

Hierarchy Level

```
[edit logical-systems name protocols              dot1x              authenticator
interface (all | [interface-names])],
[edit protocols dot1x                 authenticator interface (all | [interface-names])]
```

Description

For a device configured for 802.1X authentication, specify that when the device receives an Extensible Authentication Protocol Over LAN (EAPoL) Access-Reject message during the authentication process between the device and the RADIUS authentication server, supplicants attempting to access the LAN are granted access and moved to a specific bridge domain or VLAN. Any bridge domain, VLAN name or VLAN ID sent by a RADIUS server as part of the EAPoL Access-Reject message is ignored.

When you specify the bridge domain, VLAN ID, VLAN name, or GBP-tag, bridge domain or VLAN must already be configured on the device.

Default

None

Options

server-reject-bridge-domain *bridge-domain*
(MX Series only) Move the supplicant on the interface to the bridge domain specified by this name or numeric identifier.

server-reject-vlan (*vlan-id* **|** *vlan-name*
(MX Series in enhanced LAN mode, EX, QFX, and SRX Series only) Move the supplicant on the interface to the VLAN specified by this name or numeric identifier.

block-interval *seconds*
Specify the number of seconds that the 802.1X interface ignores Extensible Authentication Protocol (EAP) start messages from the client when an EAPoL block has been enabled on the 802.1X interface.

- **Range:** 120 through 65,535 seconds

gbp-tag *gbp-tag*
Specify the GBP tag to apply when RADIUS rejects the client authentication. If you configure the `gbp-tag` *gbp-tag* and the client authenticates in `server-reject vlan`, then the configured `gbp-tag` filter is also installed for the client. You can only configure the `server-reject gbp-tag` *gbp-tag* when the `server-reject-vlan` *vlan-id* option is configured.

- **Range:** 1 through 65,535

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

`gbp-tag` *gbp-tag* `introduced in Junos OS Release 23.4R1 for EX and QFX Series switches.`

`server-reject-vlan` introduced in Junos OS Release 9.3 for EX Series.

`block-interval` introduced in Junos OS Release 11.2 for EX Series.

`server-reject-vlan` introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

RELATED DOCUMENTATION

show dot1x | **19548**

Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch

Example: Configuring Fallback Options on EX Series Switches for EAP-TTLS Authentication and Odyssey Access Clients

Configuring RADIUS Server Fail Fallback (CLI Procedure)

Understanding Server Fail Fallback and Authentication on Switches

server-response-time (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 16664
- Hierarchy Level | 16664
- Description | 16665
- Options | 16665
- Required Privilege Level | 16665
- Release Information | 16665

Syntax

```
server-response-time seconds;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Configure the timeframe during which the router monitors DHCP server responsiveness within the routing instance. The router generates a system log message when the DHCP server does not respond to relayed packets during the specified timeframe.

Options

seconds Number of seconds the DHCP server is monitored.

- **Range:** 30 through 4,294,967,295 seconds
- **Default:** 0 (no limit)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Monitoring DHCP Relay Server Responsiveness

server (Security Antivirus)

IN THIS SECTION

- Syntax | 16666
- Hierarchy Level | 16666
- Description | 16666
- Required Privilege Level | 16666

Release Information | 16666

Syntax

```
server address-or-url;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update proxy]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update proxy]
[edit security utm feature-profile anti-virus sophos-engine pattern-update proxy]
```

Description

Set the IP address or URL for the proxy server.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 11.2.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

server (Security Group VPN)

IN THIS SECTION

- Syntax | 16667
- Hierarchy Level | 16669
- Description | 16669
- Options | 16670
- Required Privilege Level | 16670
- Release Information | 16670

Syntax

```
server {
    group name {
        anti-replay-time-window milliseconds;
        description description;
        group-id number;
        ike-gateway [gateway-name];
        ipsec-sa name {
            match-policy policy-name {
                destination ip-address/netmask;
                destination-port number;
                protocol number;
                source ip-address/netmask;
                source-port number;
            }
            proposal proposal-name;
        }
        member-threshold number;
        server-cluster {
            ike-gateway gateway-name;
            retransmission-period seconds;
            server-role (root-server | sub-server);
        }
        server-member-communication {
            certificate certificate-id;
```

```
            communication-type unicast;
            encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
            lifetime-seconds seconds;
            number-of-retransmission number;
            retransmission-period seconds;
            sig-hash-algorithm (sha-256 | sha-384);
        }
    }
    ike {
        gateway  gateway-name {
            address ip-address ;
            dead-peer-detection {
                always-send;
                interval seconds;
                threshold number;
            }
            dynamic {
                (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
            }
            ike-policy policy-name;
            local-address ip-address;
            local-identity {
                (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
            }
            remote-identity {
                (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
            }
            routing-instance routing-instance;
        }
        policy policy-name {
            description text;
            mode (aggressive | main);
            pre-shared-key (ascii-text key | hexadecimal key);
            proposals [proposal-name];
        }
        proposal proposal-name {
            authentication-algorithm (sha-256 | sha-384);
            authentication-method pre-shared-keys;
            description description;
            dh-group (group14 | group24);
            encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
        }
    }
```

```
    ipsec {
        proposal proposal-name {
            authentication-algorithm hmac-sha-256-128;
            description description;
            encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
            lifetime-seconds seconds;
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        gateway-filter {
            local-address ip-address;
            remote-address ip-address;
        }
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security group-vpn]
```

Description

Configure group VPN server. You can configure the following on the group server:

- Phase 1 IKE SA for group members
- Phase 2 IPsec proposal
- Group identifier, group members, server-member communications, and group policies to be downloaded to members
- Group VPN trace options

Options

gateway ***gateway-name***	Configure IKE gateway for group VPN server.
ike	Configure Phase 1 security association (SA) with a member on the group server.
ipsec	Configure an IPsec proposal for Phase 2 exchange on the group server.
traceoptions	Configure group VPN tracing options to aid in troubleshooting the IKE issues.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

server (Security Sophos Engine Antivirus)

IN THIS SECTION

- Syntax | 16671
- Hierarchy Level | 16671
- Description | 16671
- Options | 16671
- Required Privilege Level | 16672
- Release Information | 16672

Syntax

```
server {
 host host;
 port port;
 routing-instance routing-instance;
 source-address source-address;
 ssl-profile ssl-profile;
 proxy-profile proxy-profile;
}
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus sophos-engine],
[edit logical-systems name security utm feature-profile anti-virus sophos-engine],
[edit logical-systems name tenants name security utm feature-profile anti-virus sophos-engine],
[edit tenants name security utm feature-profile anti-virus sophos-engine]
```

Description

Enable you to configure the server parameters for the device-to-server communication.

Starting in Junos OS Release 23.1R1, The Content Security supports the new antivirus Sophos Live Protection version 2.0. The new version of Sophos antivirus uses an HTTPS connection for the device-to-server communication. For the HTTPS connection, you must create an SSL initiation profile and add the profile to the default configuration of the Sophos engine. See profile (SSL Initiation) to create an SSL initiation profile.

Options

host

Specify the server host IP address or string host name.

- **Default:** 4.sophosxl.net

port

Specify the server port number.

- **Range:** 1 through 65535

- **Default:** 443

proxy-profile	Specify the proxy profile name.
`routing-instance` *name*	Specify the name of the routing instance.
source-address	Specify the source IP address used to connect Sophos server.
ssl-profile	Specify the SSL initiation profile name.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D90.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

The `host`, `port`, and `ssl-profile` options introduced in Junos OS Release 23.1R1.

The `ip` option is deprecated in Junos OS Release 23.1R1.

The `proxy-profile` option introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

Content Security | **4589**

profile (SSL Initiation) | **13343**

server (Security Web Filtering)

IN THIS SECTION

Syntax | **16673**

- Hierarchy Level | 16673
- Description | 16673
- Options | 16673
- Required Privilege Level | 16673
- Release Information | 16674

Syntax

```
server {
    host host-name;
    port number;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name]
[edit security utm feature-profile web-filtering juniper-enhanced]
```

Description

Set server parameters by entering the server name or IP address.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The surf-control feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5 .

Statement introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The surf-control feature is not supported from Junos OS Release 15.1X49-D10 onwards.

server (TACACS+ Accounting)

IN THIS SECTION

- Syntax | 16674
- Hierarchy Level | 16675
- Description | 16675
- Options | 16675
- Required Privilege Level | 16676
- Release Information | 16676

Syntax

```
server {
    server-address {
        port port-number;
        routing-instance routing-instance;
        secret password;
        single-connection;
```

```
        source-address address
        timeout seconds;
    }
}
```

Hierarchy Level

```
[edit system accounting destination tacplus]
```

Description

Configure TACACS+ logging.

Options

server-address

Address of the IPv4 or IPv6 TACACS+ authentication server.

> **NOTE**: Wildcard characters cannot be used in the TACACS+ server address or source address. This is because the TACACS+ server and source can accept both IPv4 and IPv6 addresses and, if you use wildcard characters for these addresses, Junos OS cannot validate mismatching server and source address families.

port *port-number*

Configure the port number on which to contact the TACACS+ authentication server.

- **Default:** 49

routing-instance *routing-instance*

Configure the routing instance name for the management routing instance, that is `mgmt_junos`. Configuring this parameter along with the `management-instance` statement enables authentication processes (for example, RADIUS and TACACS+) to use the non-default management routing instance for packet traffic.

> **NOTE**: You must also define the mgmt_junos routing instance under the `[edit routing-instances]` hierarchy level.

If you no not configure the mgmt_junos instance under the `[edit routing-instances]` hierarchy level and configure it only under `tacplus-server` or `radius-server`, the commit will fail.

secret *password*

Configure the password to use with the RADIUS or TACACS+ server. The secret password used by the local router or switch must match that used by the server. The password can include spaces included in quotation marks.

> **NOTE**: To ensure better security, we recommend you configure the TACACS+ secret password with a minimum of 14 characters.

single-connection

Optimize attempts to connect to a TACACS+ server. The software maintains one open TCP connection to the server for multiple requests rather than opening a connection for each connection attempt.

source-address *source-address*

Specify a source address for each configured TACACS+ server to record in system log messages that are directed to a remote machine. Configure a valid IP address on one of the device interfaces. For system logging, the address is recorded as the message source in messages sent to the remote machines specified in all `host` *`hostname`* statements at the `[edit system syslog]` hierarchy level.

- **Default:** The primary address of the interface.

timeout *seconds*

The amount of time that the local device waits to receive a response from a TACACS+ server.

- **Default:** 3 seconds
- **Range:** 1 through 90 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`routing-instance` option introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Configuring TACACS+ System Accounting

Configuring TACACS+ Authentication

management-instance | 10309

Management Interface in a Non-Default Instance

servers

IN THIS SECTION

- Syntax | 16677
- Hierarchy Level | 16677
- Description | 16677
- Options | 16678
- Required Privilege Level | 16678
- Release Information | 16678

Syntax

```
servers server-address {
    port port-number;
}
```

Hierarchy Level

```
[edit system services service-deployment]
```

Description

Configure an IPv4 address for the Session and Resource Control (SRC) server.

Options

server-address—The TCP port number.

- **Default:** 3333

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Junos OS to Work with SRC Software

servers (RADIUS)

IN THIS SECTION

- Syntax | 16679
- Hierarchy Level | 16679
- Description | 16679
- Options | 16679
- Required Privilege Level | 16679
- Release Information | 16680

Syntax

```
servers name {
    accounting-port port-number;
    accounting-secret password;
    address server-address;
    allow-dynamic-requests;
    dead-criteria-retries retry-number interval seconds;
    dynamic-requests-secret password;
    port port-number;
    retry attempts;
    revert-interval seconds;
    secret password;
    source-interface interface [ipv4-address address];
    timeout seconds;
}
```

Hierarchy Level

```
[edit access radius]
```

Description

Configure a RADIUS server that provides policy management for TDF subscribers.

Options

name Name of the RADIUS server.

- **Range:** 1 through 32 characters

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

service

IN THIS SECTION

- Syntax | 16680
- Hierarchy Level | 16681
- Description | 16681
- Options | 16681
- Required Privilege Level | 16681
- Release Information | 16681

Syntax

```
service {
    input {
        [ service-set service-set-name <service-filter filter-name> ];
        post-service-filter filter-name;
    }
    output {
        [ service-set service-set-name <service-filter filter-name> ];
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet]
```

Description

Define the service sets and filters to be applied to an interface.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Filters and Services to Interfaces

service-agents (Analytics)

IN THIS SECTION

- Syntax | 16682
- Hierarchy Level | 16683
- Description | 16683

- Required Privilege Level | **16683**
- Release Information | **16683**

Syntax

```
service-agents {
    agent-name {
        inputs {
            analytics {
                parameters {
                    generate-tags value;
                    sample-frequency value;
                    sensors file-path;
                }
            }
            input-ipfix {
                parameters {
                    maximum-connections number;
                    tcp-port port-number;
                    vrf-name name;
                }
            }
            input-jti-ipfix {
                parameters {
                    record-group group-name {
                        record ipfix-record-name;
                        reporting-interval seconds;
                    }
                }
            }
        }
        outputs {
            file {
                parameters {
                    path file-path;
                }
            }
            kafka {
```

```
                parameters {
                    server ip-address;
                    topic topic-name;
                    encoding encoding-type;
                }
            }
            output-ipfix {
                parameters {
                    collector-address ip-address;
                    collector-ca-certificate file-path;
                    collector-certificate file-path;
                    collector-certificate-key file-path;
                    collector-connection-retry-interval seconds;
                    collector-tcp-port port-number;
                    collector-vrf-name vrf-name;
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit services analytics agent]
```

Description

Configure a network analytics service agent that uses input and output plug-ins to collect, transform, and forward network telemetry data.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data

Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator

Configuring NTF Agent

IPFIX Mediation on the BNG

Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector

service-context-id (OCS)

IN THIS SECTION

- Syntax | 16684
- Hierarchy Level | 16684
- Description | 16684
- Options | 16685
- Required Privilege Level | 16685
- Release Information | 16685

Syntax

```
service-context-id service-context ;
```

Hierarchy Level

```
[edit access ocs global]
```

Description

Configure the service-context-id to globally and uniquely identify the Service-Context-Id AVP (cc-service-context) and the Service-Identifier AVP (cc-service-identifier), which is part of the Diameter Credit Control Service charging system. The service provider or operator allocates this identifier.

Options

service-context

3rd Generation Partnership Project (3GPP)-specific values.

- *service-context*—32251
- *customer-domain*—3gpp.org

> **NOTE**: For every service used to support the packet-switching charging infrastructure, you must prepare the necessary documentation and define associated cc-service-context-id values.

For mobile subscribers, the user equipment requests services; whereas for broadband wireline subscribers, the Policy Control and Charging Rules Function (PCRF) requests services. In the wireline environment, PCRF functions as the service requester, and the Policy and Charging Enforcement Function (PCEF) functions as the service receiver and enforcer.

The PCRF controls the PCEF by pushing charging rules. Charging rules are reused as service (policy) rules to push policies. Charging rules may also have an associated rating group, or charging key. As a result, you must configure the PCEF to define a charging rules mapping between credit control services (cc-services) and rating groups.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring OCS Global Parameters

3GPP Policy and Charging Control Overview for Wireline Provisioning and Accounting

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

service-deployment

IN THIS SECTION

- Syntax | 16686
- Hierarchy Level | 16686
- Description | 16686
- Required Privilege Level | 16686
- Release Information | 16687

Syntax

```
service-deployment {
    servers server-address {
        port port-number;
    }
    source-address source-address;
}
```

Hierarchy Level

```
[edit system services]
```

Description

Enable Junos OS to work with the Session and Resource Control (SRC) software.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Junos OS to Work with SRC Software

service-device-pool (L2TP)

IN THIS SECTION

- Syntax | 16687
- Hierarchy Level | 16687
- Description | 16687
- Options | 16688
- Required Privilege Level | 16688
- Release Information | 16688

Syntax

```
service-device-pool pool-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Assign a pool of service interfaces to the tunnel group to balance traffic across.

> **NOTE**: The service interface configuration is required for static LNS sessions. Either the service interface configuration or the service device pool configuration can be used for dynamic LNS sessions.

Options

pool-name	Name of the service device pool.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

service-device-pools (Service Interfaces)

IN THIS SECTION

- Syntax | 16689
- Hierarchy Level | 16689
- Description | 16689
- Options | 16689
- Required Privilege Level | 16689
- Release Information | 16689

Syntax

```
service-device-pools {
    pool pool-name {
        interface service-interface-name;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure one or more pools of service interfaces that can be assigned to an L2TP tunnel group for traffic load-balancing. The interfaces in the pool must be inline service interfaces (si). The service device pool is required for dynamic LNS sessions.

Options

pool-name	Name of the service interface pool.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support for ps service interfaces added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

service-device (Remote Device Management)

IN THIS SECTION

- Syntax | 16690
- Hierarchy Level | 16691
- Description | 16691
- Options | 16691
- Required Privilege Level | 16692
- Release Information | 16692

Syntax

```
service-device device-name {
    access-domain {
        vlan-id-list [vlan-id-low-vlan-id-high vlan-id]
    }
    address ip-address;
    dictionary absolute file path;
    provisioning-method {
        netconf {
            bulk-interval milliseconds;
            bulk-limit number;
            connection-retry-interval seconds;
            password password;
            port port-number;
            reconfigure-bulk-limit number;
            response-timeout seconds;
            response-timeout-count number;
            user-name name;
        }
```

```
    }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name system services remote-device-management],
[edit system services remote-device-management]
```

Description

Configure a remote device to provision and deprovision services for subscribers. This configuration is used when `profile-type remote-device-service` has been configured in the dynamic service profile.

Provides basic configuration for the remote service device and enables NETCONF TCP session to be established with the service-device. Subsequently, dynamic configuration of the access-domain occurs based on allocation and assignment of individual VLAN IDs or VLAN ranges for one or more subscriber sessions.

> **NOTE**: With the exception of the `access-domain` statement, all statements are required to be configured for the service device and are subject to a commit check. This behavior enables basic configuration for the remote device and the NETCONF TCP session to be completed followed later by dynamic configuration of the access-domain based on allocation and assignment of individual VLAN IDs or VLAN ranges for one or more subscriber sessions.

Options

address *ip-address*

Specify the IP address of the remote device used by the BNG to configure the subscriber service. The address must be unique; it is used for all actions performed by the BNG, including service provisioning and deprovisioning. The address must also be unique across all routing instances.

> **NOTE**: You cannot change the IP address for the device when any active subscriber services are mapped to it.

device-name

System-wide name that uniquely identifies the device across routing instances.

dictionary *absolute file path*

Specify the absolute file path for the vendor-specific dictionary that defines the set of NETCONF XML protocol commands required to provision, deprovision, and roll back a subscriber service on the remote device. The dictionary is stored on the BNG. An example absolute path is **/var/home/dict/remote-device.xml**.

Because the dictionary is assigned to a service device, you can use service devices from different vendors at the same time. This assignment method also enables you to use multiple devices from the same vendor that are running different software releases. This is useful because the NETCONF XML protocol commands to configure the devices can differ across releases.

> **NOTE**: The maximum total length of the path is 127 characters. The filename must end in `xml`.

> **NOTE**: You cannot change the dictionary path for the device when any active subscriber services are mapped to it.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Configuring Remote Device Management for Service Provisioning

Remote Device Services Manager (RDSM) Overview

service-domain

IN THIS SECTION

- Syntax | 16693
- Hierarchy Level | 16693
- Description | 16693
- Options | 16693
- Required Privilege Level | 16694
- Release Information | 16694

Syntax

```
service-domain (inside | outside);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet]
```

Description

Specify the service interface domain. If you specify this interface using the `next-hop-service` statement at the `[edit services service-set service-set-name]` hierarchy level, the interface domain must match that specified with the `inside-service-interface` and `outside-service-interface` statements.

Options

`inside`—Interface used within the network.

`outside`—Interface used outside the network.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring the Address and Domain for Services Interfaces

service (Dynamic Profiles)

IN THIS SECTION

- Syntax | 16694
- Hierarchy Level | 16695
- Description | 16695
- Required Privilege Level | 16695
- Release Information | 16695

Syntax

```
service {
    pcef  pcef-profile-name {
        activate rule-name | activate-all;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

Description

Assign a policy and charging enforcement function (PCEF profile) to the dynamic profile. The PCEF profile specifies a set of PCC rules and rulebases to assign to a subscriber for application-aware policy control, and assigns a precedence value to each predefined rule. Also specify which of the PCC rules to activate for the dynamic profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Configuring Application-Aware Policy Control for Subscriber Management

service (Dynamic Service Sets)

IN THIS SECTION

- Syntax | 16696
- Hierarchy Level | 16696
- Description | 16696

Required Privilege Level | 16696
Release Information | 16697

Syntax

```
service {
     input {
         service-set service-set-name {
            service-filter filter-name;
        }
        post-service-filter filter-name;
    }
    output {
         service-set service-set-name {
            service-filter filter-name;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family]
```

Description

Define the service sets and filters to be applied to an interface. This statement is not supported for `family inet6`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family]` hierarchy level introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic Service Sets Overview

Associating Service Sets with Interfaces in a Dynamic Profile

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

Configuring Routing Engine-Based, Static HTTP Redirect Services

Configuring Routing Engine-Based, Converged HTTP Redirect Services

service-filter (Dynamic Service Sets)

IN THIS SECTION

- Syntax | 16697
- Hierarchy Level | 16698
- Description | 16698
- Options | 16698
- Required Privilege Level | 16698
- Release Information | 16698

Syntax

```
service-filter filter-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family service input service-set service-set-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family service output service-set service-set-name],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family
service input service-set service-set-name],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family
service output service-set service-set-name]
```

Description

Define the filter to be applied to traffic before it is accepted for service processing. You can use the predefined dynamic interface variables `$junos-input-service-filter`, `$junos-output-service-filter`, `$junos-input-ipv6-service-filter`, and `$junos-output-ipv6-service-filter`. Configuration of a service filter is optional; if you include the `service-set` statement without a `service-filter` definition, the router software assumes that the match condition is true and selects the service set for processing automatically.

Options

`filter-name`—Identifies the filter to be applied in service processing.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the [`edit dynamic-profiles` *`profile-name`* `interfaces pp0 unit "$junos-interface-unit" family` *`family`* `service input service-set` *`service-set-name`*] and [`edit dynamic-profiles` *`profile-name`* `interfaces pp0 unit "$junos-interface-unit" family` *`family`* `service output service-set` *`service-set-name`*] hierarchy levels introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic Service Sets Overview

Associating Service Sets with Interfaces in a Dynamic Profile

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

Configuring Routing Engine-Based, Static HTTP Redirect Services

Configuring Routing Engine-Based, Converged HTTP Redirect Services

service-filter (Firewall)

IN THIS SECTION

- Syntax | **16699**
- Hierarchy Level | **16700**
- Description | **16700**
- Options | **16700**
- Required Privilege Level | **16700**
- Release Information | **16700**

Syntax

```
service-filter filter-name {
    term term-name {
        from {
            match-conditions;
        }
        then {
            actions;
        }
    }
}
```

Hierarchy Level

```
[edit firewall family (inet | inet6],
[edit logical-systems logical-system-name firewall family (inet | inet6)]
```

Description

Configure service filters.

NOTE: ACX Series routers do not support `family` *`inet6`*.

Options

filter-name—Name that identifies the service filter. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Guidelines for Configuring Service Filters

Guidelines for Applying Service Filters

service-filter (Interfaces)

IN THIS SECTION

- Syntax | 16701
- Hierarchy Level | 16701
- Description | 16701
- Options | 16701
- Required Privilege Level | 16702
- Release Information | 16702

Syntax

```
service-filter filter-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet service (input | output)
service-set service-set-name],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet service (input | output) service-set service-set-name]
```

Description

Define the filter to be applied to traffic before it is accepted for service processing. Configuration of a service filter is optional; if you include the `service-set` statement without a `service-filter` definition, Junos OS assumes the match condition is true and selects the service set for processing automatically.

Options

`filter-name`—Identifies the filter to be applied in service processing. You can include special characters, such as a forward slash (/), colon (:), or a period (.).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Applying Filters and Services to Interfaces*
>
> Junos OS Services Interfaces Library for Routing Devices

service-groups

IN THIS SECTION

- Syntax | 16702
- Hierarchy Level | 16703
- Description | 16703
- Options | 16703
- Required Privilege Level | 16703
- Release Information | 16703

Syntax

```
service-groups {
    service-group-name {
        pbb-service-options {
            default-isid isid-number;
            isid isid-number  vlan-id-list [ vlan-ids ];
            mac-address mac-address;
```

```
        }
        service-type  (elan | eline);
        source-bmac <mac-address> <length>
    }
}
```

Hierarchy Level

```
[edit routing-instances (Multiple Routing Entities) routing-instance-name]
```

Description

For IEEE 802.1ah provider backbone bridge (PBB) configurations, configure the service groups to be supported in the customer routing instance (I-component).

NOTE: The `mac-address` statement under `pbb-service-options` is not supported for PBB-EVPN.

NOTE: The `eline` option for the `service-type` statement is not supported for PBB-EVPN.

The remaining statements are explained separately. See CLI Explorer.

Options

service-group-name—Name of a service group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.0.

service-id (Bridge Domains)

IN THIS SECTION

- Syntax | 16704
- Hierarchy Level | 16704
- Description | 16704
- Options | 16704
- Required Privilege Level | 16705
- Release Information | 16705

Syntax

```
service-id number;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name]
```

Description

Specify a service identifier to include in the packets sent to and from the multichassis link aggregation (MC-LAG) bridge domain when the VLAN identifier is set to `none`. This configuration facilitates media access control (MAC) and Address Resolution Protocol (ARP) synchronization among MC-LAG peers.

NOTE: The VLAN identifier none is supported only for IPv4 traffic.

Options

number—A valid service identifier. You must configure the same service identifier within the bridge domains of MC-LAG peers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2

RELATED DOCUMENTATION

Configuring a Bridge Domain

Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances

Understanding Layer 2 Learning and Forwarding for Bridge Domains

bridge-domains | 3032

service-id (Switch Options)

IN THIS SECTION

- Syntax | 16705
- Hierarchy Level | 16706
- Description | 16706
- Options | 16706
- Required Privilege Level | 16706
- Release Information | 16706

Syntax

```
service-id number;
```

Hierarchy Level

```
[edit switch-options]
[edit vlans vlan-name]
```

Description

Specify a service identifier for each multichassis aggregated Ethernet interface that belongs to a link aggregation group (LAG).

Options

number A number that identifies a particular service.

- **Range:** 1 through 65535

Required Privilege Level

system—To view this statement in the configuration.

system control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R2.

service-interface (L2TP Processing)

IN THIS SECTION

- Syntax | 16707
- Hierarchy Level | 16707
- Description | 16707
- Options | 16707
- Required Privilege Level | 16707

Release Information | 16708

Syntax

```
service-interface interface-name;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Specify the service interface responsible for handling L2TP processing.

NOTE: On MX Series routers, the service interface configuration is required for static LNS sessions. Either the service interface configuration or the service device pool configuration can be used for dynamic LNS sessions.

Options

interface-name Name of the service interface. The ae, si, and sp interface types are supported as follows:

- `asix`—(MPCs on MX Series routers) Aggregated inline services interface.
- `sp-fpc/pic/port`—On AS or Multiservices PICs on M7i, M10i, and M120 routers.
- `si-fpc/pic/port`—On MPCs on MX Series routers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`si-fpc/pic/port` option added in Junos OS Release 11.4.

Option `asifpc` added in Junos OS Release 16.2.

RELATED DOCUMENTATION

Configuring the Local Gateway Address and PIC

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

service-interface-pools

IN THIS SECTION

- Syntax | 16708
- Hierarchy Level | 16709
- Description | 16709
- Options | 16709
- Required Privilege Level | 16709
- Release Information | 16709

Syntax

```
service-interface-pools {
    pool pool-name {
        interface interface-name.unit-number;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure service interface pools used for VPN aggregation.

Options

The options are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Service Interface Pools

service-interface (Services Interfaces)

IN THIS SECTION

- Syntax | 16710
- Hierarchy Level | 16710
- Description | 16710
- Options | 16710

Required Privilege Level | 16710
Release Information | 16710

Syntax

```
service-interface interface-name;
```

Hierarchy Level

```
[edit services service-set service-set-name interface-service]
```

Description

Specify the name for the services interface associated with an interface-wide service set.

Options

interface-name	Identifier of the service interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

- Configuring Service Sets to be Applied to Services Interfaces
- Applying Services to Subscriber-Aware Traffic with a Service Set

service-mode

IN THIS SECTION

- Syntax | 16711
- Hierarchy Level | 16711
- Description | 16711
- Options | 16712
- Required Privilege Level | 16712
- Release Information | 16712

Syntax

```
service-mode service-mode-options;
```

Hierarchy Level

```
[edit routing-instances vrf-name access address-assignment address-pools juniper-pool],
[edit unified-edge gateways tdf gateway-name],
[edit unified-edge gateways tdf gateway-name domains domain-name],
[edit unified-edge tdf gateway-name system interface interface-name],
[edit unified-edge gateways tdf gateway-name system session-pics interface interface-name]
```

Description

Set maintenance mode for a network element so that you can carry out maintenance tasks such as deleting or modifying the element, for example, an address pool.

When in the maintenance mode active phase, you can modify all the valid attributes on the network element. In other cases, you can modify only the non-maintenance mode attributes.

The following network elements must be in maintenance mode before you can modify or delete them:

- Address pools
- AMS interfaces
- PCEF profiles
- Session PICs
- Service PICs
- Static time-of-day settings
- TDF domains
- TDF interfaces
- TDF gateways

Options

service-mode-options Type of the service mode. Currently, only **maintenance** mode is supported.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Maintenance Mode Overview for Subscriber Aware Policy Enforcement

service-mode (mobile-edge SAEGW)

IN THIS SECTION

- Syntax | 16713
- Hierarchy Level | 16713
- Description | 16713
- Required Privilege Level | 16714
- Release Information | 16714

Syntax

```
service-mode {
      maintenance
      }
}
```

Hierarchy Level

```
[edit services mobile-edge gateways saegw gateway-name system]
```

Description

Enter maintenance mode for SAEGW gateways. This makes your gateway inactive for new sessions, but traffic flow and existing sessions will continue uninterrupted. Critical session modify and session delete messages are also allowed in maintenance mode.

When you enter maintenance mode, a duration is set for the maintenance window. The maintenance window can be ended early with a new load control information (LCI) message. If maintenance mode is entered in response to an overload control information (OCI) message, the maintenance window must run for the specified duration in the OCI message.

If the load/overload feature is not enabled, a warning message will be displayed and the command will be rejected.

To enter maintenance mode, enter the following command:

```
user@host# set services mobile-edge gateways saegw gateway-name system service-mode maintenance
```

To leave maintenance mode, enter this command:

```
user@host# delete services mobile-edge gateways saegw gateway-name system service-mode
maintenance
```

To end maintenance mode, use the `delete` command instead of the `set` command.

Required Privilege Level

system

Release Information

Juniper Networks released this statement in Junos OS Release 23.1R1.

service-name

IN THIS SECTION

- Syntax | 16715
- Hierarchy Level | 16715
- Description | 16715
- Options | 16715
- Required Privilege Level | 16715
- Release Information | 16715

Syntax

```
service-name name;
```

Hierarchy Level

```
[edit interfaces pp0 unit logical-unit-number pppoe-options],
[edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number pppoe-options]
```

Description

PPP over Ethernet interfaces, configure the service to be requested from the PPP over Ethernet server; that is, the access concentrator. For example, you can use this statement to indicate an Internet service provider (ISP) name or a class of service.

Options

name—Service to be requested from the PPP over Ethernet server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the PPPoE Service Name

service-name-table

IN THIS SECTION

- Syntax | 16716
- Hierarchy Level | 16716
- Description | 16716
- Options | 16717
- Required Privilege Level | 16717
- Release Information | 16717

Syntax

```
service-name-table table-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
pppoe],
[edit interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
pppoe-underlying-options]
```

Description

Specify the PPPoE service name table assigned to a PPPoE underlying interface. This underlying interface is configured with either the `encapsulation ppp-over-ether` statement or the `family pppoe` statement; the two statements are mutually exclusive.

> **NOTE**: The `[edit ... family pppoe]` hierarchies are supported only on MX Series routers with MPCs.

Options

`table-name`—Name of the PPPoE service name table, a string of up to 32 alphanumeric characters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... family pppoe]` hierarchies introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

Assigning a Service Name Table to a PPPoE Underlying Interface

Configuring the PPPoE Family for an Underlying Interface

service-name-tables

IN THIS SECTION

- Syntax | **16718**
- Hierarchy Level | **16718**
- Description | **16718**
- Options | **16719**
- Required Privilege Level | **16719**

Release Information | 16719

Syntax

```
service-name-tables table-name {
    service service-name {
        drop;
        delay seconds;
        terminate;
        dynamic-profile profile-name;
        routing-instance routing-instance-name;
        max-sessions number;
        agent-specifier {
            aci circuit-id-string ari remote-id-string {
                drop;
                delay seconds;
                terminate;
                dynamic-profile profile-name;
                routing-instance routing-instance-name;
                static-interface interface-name;
            }
        }
    }
}
```

Hierarchy Level

```
[edit protocols pppoe]
```

Description

Create and configure a PPPoE service name table. Specify the action taken for each service and remote access concentrator on receipt of a PPPoE Active Discovery Initiation (PADI) packet. You can also specify the dynamic profile and routing instance that the router uses to instantiate a dynamic PPPoE interface, and the maximum number of active PPPoE sessions that the router can establish with the specified service. A maximum of 32 PPPoE service name tables is supported per router.

Options

table-name—Name of the PPPoE service name table, a string of up to 32 alphanumeric characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

`dynamic-profile`, `routing-instance`, `max-sessions`, and `static-interface` options introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

Creating a Service Name Table

service-order

IN THIS SECTION

- Syntax | 16720
- Hierarchy Level | 16720
- Description | 16720
- Options | 16720
- Required Privilege Level | 16720
- Release Information | 16720

Syntax

```
service-order {
    forward-flow [ service-name1 service-name2 ];
    reverse-flow [ service-name1 service-name2 ];
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Define the order of services in service set to be applied to traffic coming to the PIC.

NOTE: If the `extension-service` statement is specified, the `service-order` statement is mandatory.

Options

`forward-flow`—Order of services in service set to be applied in forward flow.

`reverse-flow`—Order of services in service set to be applied in reverse flow. If you want the order to be the reverse of that specified for forward flow, this is optional. However, if, for example, you want the order to be the same regardless of direction of flow, you must include this statement. (The exception to this is for the sampling service set type. If a service set is a sampling service set and the reverse-flow service order is not configured, all sampled traffic is considered to be forward traffic.)

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

extension-service

service-package

IN THIS SECTION

- Syntax | 16721
- Hierarchy Level | 16722
- Description | 16722
- Options | 16722
- Required Privilege Level | 16722
- Release Information | 16722

Syntax

```
service-package {
    extension-provider {
        control-cores control-number;
        data-cores data-number;
        data-flow-affinity {
            hash-key (layer-3 | layer-4);
        }
        forwarding-db-size size;
        object-cache-size size;
        package package-name;
        policy-db-size size;
        syslog {
            facility {
                severity;
                destination destination;
            }
        }
        wired-max-processes num-procs;
        wired-process-mem-size mem-size;
```

```
    }
    layer-2;
    layer-3;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services]
```

Description

For adaptive services and multi-services interfaces, enable a service package on the specified Physical Interface Card (PIC).

> **NOTE**: By default, the MS-MPC and MS-MIC cards come up with `extension-provider` as `service-package`, irrespective of the configuration. These cards do not support any other `service-package` configuration.

Options

layer-2 Enable a Layer 2 service package on the specified PIC.

layer-3 Enable a Layer 3 service package on the specified PIC.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced on MX Series 5G Universal Routing Platforms with MS-DPCs in Junos OS Release 9.6.

RELATED DOCUMENTATION

Enabling Service Packages

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

service-pics

IN THIS SECTION

- Syntax | 16723
- Hierarchy Level | 16723
- Description | 16723
- Required Privilege Level | 16724
- Release Information | 16724

Syntax

```
service-pics {
    [interface interface-name];
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name system]
```

Description

Specify the service interfaces that represent the service PICs used for anchoring subscriber-aware services in the TDF Gateway.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Service PICs

TDF Gateway Service PICs and Session PICs for Subscriber-Aware Traffic Treatment

service-port (UDP, Dynamic Flow Capture)

IN THIS SECTION

- Syntax | 16724
- Hierarchy Level | 16725
- Description | 16725
- Options | 16725
- Required Privilege Level | 16725
- Release Information | 16725

Syntax

```
service-port port-number;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

Identify the User Datagram Protocol (UDP) port number for control protocol requests.

Options

port-number—Port number for control protocol request messages.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Control Source

service-port (Inline Services)

IN THIS SECTION

- Syntax | 16726
- Hierarchy Level | 16726
- Description | 16726
- Required Privilege Level | 16726

Release Information | 16726

Syntax

```
service-port port-number
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number inline-services]
```

Description

For MX Series 5G Universal Routing Platforms, configure the port number of the inline service interface. The range is 0 through 3.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1R1.

RELATED DOCUMENTATION

bandwidth (Tunnel Services) | 2880

inline-services (PIC level) | 8676

service (PPPoE)

IN THIS SECTION

- Syntax | 16727
- Hierarchy Level | 16728
- Description | 16728
- Default | 16728
- Options | 16728
- Required Privilege Level | 16728
- Release Information | 16729

Syntax

```
service service-name {
    drop;
    delay seconds;
    terminate;
    dynamic-profile profile-name;
    routing-instance routing-instance-name;
    max-sessions number;
    agent-specifier {
        aci circuit-id-string ari remote-id-string {
            drop;
            delay seconds;
            terminate;
            dynamic-profile profile-name;
            routing-instance routing-instance-name;
            static-interface interface-name;
        }
    }
}
```

Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name]
```

Description

Specify the action taken by the interface on receipt of a PPPoE Active Discovery Initiation (PADI) control packet for the specified named service, `empty` service, or `any` service in a PPPoE service name table. You can also specify the dynamic profile and routing instance that the router uses to instantiate a dynamic PPPoE interface, and the maximum number of active PPPoE sessions that the router can establish with the specified service.

Default

The default action is terminate.

Options

service-name—Service entry in the PPPoE service name table:

- *service-name*—Named service entry of up to 32 characters; for example, `premiumService`. You can configure a maximum of 512 named service entries across all PPPoE service name tables on the router.
- `empty`—Service entry of zero length that represents an unspecified service. Each PPPoE service name table includes one `empty` service entry by default.
- `any`—Default service for non-empty service entries that do not match the named or `empty` service entries configured in the PPPoE service name table. Each PPPoE service name table includes one `any` service entry by default.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

`any`, `dynamic-profile`, `routing-instance`, `max-sessions`, and `static-interface` options introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag

Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag

Configuring the Action Taken for the Any Service

service-profile (DHCP Local Server)

IN THIS SECTION

- Syntax | **16729**
- Hierarchy Level | **16730**
- Description | **16730**
- Options | **16730**
- Required Privilege Level | **16730**
- Release Information | **16731**

Syntax

```
service-profile dynamic-profile-name;
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server group group-name interface interface-name],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Specify the default subscriber service or DHCP client management service, which is activated when the subscriber or client logs in and no other service is activated by a RADIUS server or a provisioning server.

- To specify the default service for all DHCP local server clients, include the `service-profile` statement at the `[edit system services dhcp-local-server]` hierarchy level.
- To specify the default service for a named group of interfaces, include the `service-profile` statement at the `[edit system services dhcp-local-server group group-name]` hierarchy level.
- To specify the default service for a particular interface within a named group of interfaces, include the `service-profile` statement at the `[edit system services dhcp-local-server group group-name interface interface-name]` hierarchy level.
- For DHCPv6 clients, use the `service-profile` statement at the `[edit system services dhcp-local-server dhcpv6]` hierarchy level.

Options

dynamic-profile-name—Name of the dynamic profile that defines the service.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Differences Between Legacy DHCP and Extended DHCP

Default Subscriber Service Overview

Configuring a Default Subscriber Service

service-profile (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 16731
- Hierarchy Level | 16731
- Description | 16732
- Options | 16732
- Required Privilege Level | 16732
- Release Information | 16732

Syntax

```
service-profile dynamic-profile-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name],
```

```
[edit forwarding-options dhcp-relay dhcpv6 groupgroup-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify the default subscriber service (or the default DHCP client management service), which is activated when the subscriber (or client) logs in and no other service is activated by a RADIUS server or a provisioning server.

- To specify the default service for all DHCP relay agent clients, include the `service-profile` statement at the `[edit forwarding-options dhcp relay]` hierarchy level.
- To specify the default service for a named group of interfaces, include the `service-profile` statement at the `[edit forwarding-options dhcp relay group group-name]` hierarchy level.
- To specify the default service for a particular interface within a named group of interfaces, include the `service-profile` statement at the `[edit forwarding-options dhcp relay group group-name interface interface-name]` hierarchy level.

Options

dynamic-profile-name—Name of the dynamic service profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

dhcp-relay | 5620

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

Grouping Interfaces with Common DHCP Configurations

Default Subscriber Service Overview

service-profile (L2TP)

IN THIS SECTION

- Syntax | 16733
- Hierarchy Level | 16733
- Description | 16733
- Options | 16734
- Required Privilege Level | 16734
- Release Information | 16734

Syntax

```
service-profile profile-name(parameter)&profile-name;
```

Hierarchy Level

```
[edit access profile profile-name client client-name l2tp]
[edit services l2tp tunnel-group group-name]
```

Description

Configure one or more dynamic service profiles to be applied to subscriber sessions at activation for all subscribers in the specified tunnel group or on the specified LAC. Services are typically applied to L2TP sessions with RADIUS VSAs or CoA requests. In multivendor environments, you might use only standard

attributes to simplify management of multiple vendor VSAs. This statement enables you to apply services without using an external authority such as RADIUS. The locally configured list of services (service profiles) serves as local authorization that is applied by authd during client session activation. This list of services is subject to the same validation and processing as services originating from an external authority, such as RADIUS.

You can optionally specify parameters that are passed to the corresponding service when it is activated for the session. The parameter might override values configured in the profile itself, such as a downstream shaping rate for a CoS service. This enables you to use the same service profile for multiple situations with different requirements, or to modify a previously applied value for a service.

You can still use RADIUS VSAs or CoA requests together with the service profiles. If services are sourced from an external authority as authorization during authentication or during subscriber session provisioning (activation), the services from the external authority take strict priority over those in the local configuration. If a service applied with RADIUS is the same as a service applied with a service profile in the CLI, but with different parameters, the RADIUS service is applied with a new session ID and takes precedence over the earlier service profile.

When service profiles are configured on a LAC client and on a tunnel group that uses that LAC client, the LAC configuration overrides the tunnel group configuration. Only the service profile configured on the LAC client is applied to subscribers in the tunnel group.

Options

profile-name Name of a dynamic service profile that defines a service to be applied to L2TP subscriber sessions. You can specify one or more service profiles, separated by an ampersand (&).

parameter (Optional) Value to be passed to the service when it is activated on the subscriber session.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1 on MX Series routers.

RELATED DOCUMENTATION

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

service-profile (Static Subscribers)

IN THIS SECTION

- Syntax | 16735
- Hierarchy Level | 16735
- Description | 16735
- Options | 16736
- Required Privilege Level | 16736
- Release Information | 16736

Syntax

```
service-profile service-profile-name;
```

Hierarchy Level

```
[edit system services static-subscribers],
[edit system services static-subscribers group group-name]
```

Description

Specify the service profile to apply services for all static subscribers at the global level and at the group level. You can use `service-profile` to assign a default service profile to be applied to a `static-subscriber` session if the policy server is unavailable.

Options

service-profile-name—Name of the service profile.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

JSRC and Subscribers on Static Interfaces

JSRC and Subscribers on Static Interfaces

JSRC and Subscribers on Static Interfaces

service (Security IDP Anomaly Attack)

IN THIS SECTION

- Syntax | 16737
- Hierarchy Level | 16737
- Description | 16737
- Options | 16737
- Required Privilege Level | 16737
- Release Information | 16737

Syntax

```
service service-name;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type anomaly]
```

Description

Service is the protocol whose anomaly is defined in the attack. IP, TCP, UDP, and ICMP are also valid as services. (Protocol names must be entered in lowercase.)

Options

service-name—Name of the protocol in lowercase.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

service (Security IDP Dynamic Attack Group)

IN THIS SECTION

- Syntax | 16738
- Hierarchy Level | 16738
- Description | 16738

- Options | 16738
- Required Privilege Level | 16738
- Release Information | 16738

Syntax

```
service {
    values [service-value];
}
```

Hierarchy Level

```
[edit security idp dynamic-attack-group dynamic-attack-group-name filters]
```

Description

Specify a service filter to add attack objects based on the attack service, such as FTP, HTTP, NetBios, and so on.

Options

`values`—Name of the service filter. You can configure multiple filters separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

service (Service Accounting)

IN THIS SECTION

- Syntax | 16739
- Hierarchy Level | 16739
- Description | 16739
- Required Privilege Level | 16739
- Release Information | 16740

Syntax

```
service {
    accounting-order (activation-protocol | local | radius);
    accounting {
        statistics (time | volume-time);
        update-interval minutes;
    }
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Define the subscriber service accounting configuration.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

`accounting`, `update-interval`, and `statistics` options added in Junos OS Release 14.2R1 for MX Series routers.

RELATED DOCUMENTATION

Configuring Service Accounting with JSRC

Service Accounting with JSRC

Configuring Service Accounting in Local Flat Files

Configuring Service Accounting

Configuring Per-Subscriber Session Accounting

service-set (Application-Aware Control Policy)

IN THIS SECTION

- Syntax | 16740
- Hierarchy Level | 16741
- Description | 16741
- Options | 16741
- Required Privilege Level | 16742
- Release Information | 16742

Syntax

```
service-set service-set-name {
    application-identification-profile app-id-profile-name;
    interface-service {
```

```
        service-interface interface-name;
    }
    hcm-profile profile-name;
    lrf-profile profile-name;
    pcef-profile pcef-profile-name;
    service-set-options {
        subscriber-awareness;
        static-subscriber-application;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure a service set to identify the service interface that handles application-aware policy control.

Options

application-identification–profile *app-id-profile-name*

Specify a dummy application identification profile that you configured at the `[edit services application-identification profile]` hierarchy level. This profile is a placeholder profile with no configuration options, but it must be specified to enable application identification functionality on the services plane.

hcm-profile *profile-name*

Identify the HTTP Content Manager profile that was configured at the `[edit services hcm]` hierarchy level. This placeholder profile has no configuration options, but it must be specified to enable HTTP header enrichment functionality on the services plane.

lrf-profile *profile-name*

Identify the LRF profile that defines the logging of subscriber application-aware data sessions. The LRF profile must first be configured at the `[edit services lrf]` hiearchy level.

pcef-profile *pcef-profile-name*

Specify a dummy PCEF profile that you configured at the `[edit services pcef]` hierarchy level. This profile is a placeholder profile with no configuration options, but it must be specified to enable application-aware policy control functionality on the services plane.

service-interface ***interface-name***	Specify the services PIC interface on which the services are performed.
service-set-name	Name of the service set.
static-subscriber-application	Apply advanced services, such as deep packet inspection, to static subscribers configured under a dynamic profile.
subscriber-awareness	Enable subscriber-aware services on the service set.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement options for application-aware policy control introduced in Junos OS Release 17.2R1 on MX Series for Broadband Subscriber Management.

USF support for statement options for application-aware policy control introduced in Junos OS Release 19.3R2 on MX Series for Broadband Subscriber Management.

`static-subscriber-application` option added in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control

Configuring Logging and Reporting for Subscriber Management

service-set (Dynamic Service Sets)

IN THIS SECTION

- Syntax | **16743**
- Hierarchy Level | **16743**

- Description | 16743
- Options | 16744
- Required Privilege Level | 16744
- Release Information | 16744

Syntax

```
service-set service-set-name {
    service-filter filter-name;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family service input],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family service output],
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” family family
service input],
[edit dynamic-profiles profile-name interfaces pp0 unit  “$junos-interface-unit” family family
service output]
```

Description

Define one or more service sets in a dynamic profile. Service sets are applied to an interface. If you define multiple service sets, the router software evaluates the filters in the order in which they appear in the configuration. You can use the predefined dynamic interface variables `$junos-input-service-set`, `$junos-output-service-set`, `$junos-input-ipv6-service-set`, and `$junos-output-ipv6-service-set`.

NOTE: Starting in Junos OS Release 17.2R1, you can configure converged services at the `edit dynamic-profiles http-redirect-converged` hierarchy level. CPCD rules can also be configured under the dynamic profiles stanza to achieve parameterization of the rules.

This mechanism provides additional flexibility to customize the different rules on a per subscriber basis through service attachment.

Options

service-set-name—Name of the service set.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family service input]` and `[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family service output]` hierarchy levels introduced in Junos OS Release 10.1.

From 17.2R1 onwards, you can configure converged services at the `edit dynamic-profiles http-redirect-converged` hierarchy level.

RELATED DOCUMENTATION

Dynamic Service Sets Overview

Associating Service Sets with Interfaces in a Dynamic Profile

service-set (Interfaces)

IN THIS SECTION

Syntax | 16745

- Hierarchy Level | 16745
- Description | 16745
- Options | 16745
- Required Privilege Level | 16745
- Release Information | 16745

Syntax

```
service-set service-set-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet service (input | output)],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet service (input | output)]
```

Description

Define one or more service sets to be applied to an interface. If you define multiple service sets, the router software evaluates the filters in the order in which they appear in the configuration.

Options

service-set-name—Name of the service set.

Required Privilege Level

System—To view this statement in the configuration.

System-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Guidelines for Configuring Service Filters

service-set-options

IN THIS SECTION

- Syntax | **16746**
- Hierarchy Level | **16747**
- Description | **16747**
- Required Privilege Level | **16747**
- Release Information | **16747**

Syntax

```
service-set-options {
    bypass-traffic-on-exceeding-flow-limits;
    bypass-traffic-on-pic-failure;
    enable-asymmetric-traffic-processing;
    enable-descriptive-session-syslog;
    header-integrity-check;
    routing-engine-services;
    static-subscriber-application;
    subscriber-awareness;
    support-uni-directional-traffic;
    tcp-fast-open  {
        disabled;
        drop;
    }
    tcp-non-syn  {
```

```
        drop-flow;
        drop-flow-send-rst;
    }
unidirectional-session-refreshing {
        input;
        output;
    }
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the service set options to apply to a service set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`enable-asymmetric-traffic-processing` and `support-uni-directional-traffic` options added in Junos OS Release 11.2.

`routing-engine-services` option added in Junos OS Release 15.1.

`enable-change-on-ams-redistribution` option added in Junos OS Release 15.1.

`subscriber-awareness` option added in Junos OS Release 17.1.

`tcp-fast-open` option added in Junos OS Release 17.2.

`enable-descriptive-session-syslog` option added in Junos OS Release 20.3.

`static-subscriber-application` option added in Junos OS Release 21.2.

RELATED DOCUMENTATION

Configuring Service Sets to be Applied to Services Interfaces

Configuring APPID Support for Unidirectional Traffic

service-set-options (Next Gen Services Services)

IN THIS SECTION

- Syntax | 16748
- Hierarchy Level | 16749
- Description | 16749
- Required Privilege Level | 16749
- Release Information | 16749

Syntax

```
service-set-options {
    bypass-traffic-on-exceeding-flow-limits;
    disable-global-timeout-override;
    disable-session-open-syslog ;
    enable-asymmetric-traffic-processing;
    inactivity-non-tcp-timeout ;
    max-sessions-per-subscriber
    session-limit;
    session-timeout;
    tcp-session {
        inactivity-asymm-tcp-timeout ;
        inactivity-tcp-timeout ;
        open-timeout ;
        tcp-fast-open ;
        tcp-mss ;
        tcp-non-syn ;
        tcp-tickles ;
```

```
    }
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the service set options to apply to a service set.

disable-session-open-syslog Disable session open information from being collected in system logs.

inactivity-non-tcp-timeout Specify the inactivity timeout period for non-TCP established sessions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Service Sets to be Applied to Services Interfaces

Configuring APPID Support for Unidirectional Traffic

service-set (Services)

IN THIS SECTION

- Syntax | 16750
- Hierarchy Level | 16753
- Description | 16753
- Options | 16753
- Required Privilege Level | 16753
- Release Information | 16753

Syntax

```
service-set service-set-name {
    allow-multicast;
    captive-portal-content-delivery-profile;
    cos-options {
        match-rules-on-reverse-flow;
    }
    cos-rules [cos-rule-name];
    cos-rule-sets;
    extension-service service-name {
        provider-specific-rules-configuration;
    }
    flow;
    hcm-profile;
    ids-option;
    (ids-rules rule-name | ids-rule-sets rule-set-name);
    interface-service {
        load-balancing-options {
            hash-keys {
                egress-key (destination-ip | source-ip);
                ingress-key (destination-ip | source-ip);
            }
        }
        service-interface interface-name;
    }    ipsec-vpn-options {
```

```
        anti-replay-window-size bits;
        clear-dont-fragment-bit;
        ike-access-profile profile-name;
        local-gateway address;
        no-anti-replay;
        no-certificate-chain-in-ike;
        passive-mode-tunneling;
        trusted-ca [ ca-profile-names ];
        tunnel-mtu bytes;
        udp-encapsulation {
            <udp-dest-port destination-port>;
        }
    }
    ip-reassembly-rules rule-name};
    (ipsec-vpn-rules rule-name | ipsec-vpn-rule-sets rule-set-name);
        ipsec-vpn vpn1;
        jflow-log;
        lrf-profile;
    max-flows number;
    max-drop-flows {
        ingress ingress-flows;
        egress egress-flows;
    }
    max-session-setup-rate max-setup-rate;
    nat-options {
        land-attack-check (ip-only | ip-port);
        max-sessions-per-subscriber session-number;

        stateful-nat64 {
            clear-dont-fragment-bit;
        }
    }
    (nat-rules rule-name | nat-rule-sets rule-set-name);
    next-hop-service {
        inside-service-interface interface-name.unit-number;
        outside-service-interface interface-name.unit-number;
        outside-service-interface-type local;
        service-interface-pool name;
    }
        pcef-profile;pcp-rules rule-name;
    (pgcp-rules rule-name | pgcp-rule-sets rule-set-name);
    (ptsp-rules rule-name | ptsp-rule-sets rule-set-name);
        redundancy-set-id;
```

```
        replicate-services;
    service-set-options {
        bypass-traffic-on-exceeding-flow-limits;
        bypass-traffic-on-pic-failure;
        disable-session-open-syslog;
        enable-asymmetric-traffic-processing;
        header-integrity-check;
        routing-engine-services;
        static-subscriber-application;
        subscriber-awareness;
        support-uni-directional-traffic;
    }
    snmp-trap-thresholds {
        flows high high-threshold | low low-threshold;
        nat-address-port high-threshold | low low-threshold;
        }
    }
    softwire-options {
        dslite-ipv6-prefix-length dslite-ipv6-prefix-length;
    }
    (softwire-rules rule-name | softwire-rule-sets rule-set-name);
    (stateful-firewall-rules rule-name | stateful-firewall-rule-sets rule-set-name);
    syslog {
        host hostname {
            class {
                alg-logs;
                deterministic-nat-configuration-log;
                ids-logs;
                nat-logs;
                packet-logs;
                pcp-logs;
                session-logs <open | close>;
                stateful-firewall-logs ;
            }
            services severity-level;
            facility-override facility-name;
            interface-service prefix-value;
            port port-number;
            services severity-level;
        }
    }
```

```
    (web-filter-profile | url-filter-profile) profile-name;
}
```

Hierarchy Level

```
[edit services]
```

Description

Define the service set.

> **NOTE**: Use the `web-filter-profile` option starting in Junos OS Release 18.3R1 and use the `url-filter-profile` option in Junos OS Releases before 18.3R1.

Options

`service-set-name`—Name of the service set. You can include special characters, such as a forward slash (/), colon (:), or a period (.).

- **Range:** Up to 64 alphanumeric characters.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`pgcp-rules` and `pgcp-rule-sets` options added in Junos OS Release 8.4.

`server-set-options` option added in Junos OS Release 10.1.

`ptsp-rules` and `ptsp-rule-sets` options added in Junos OS Release 10.2.

`softwire-rules` and `clear-rule-sets` options added in Junos OS Release 10.4.

`ip-reassembly-rules` and `outside-service-interface-type` option added in Junos OS Release 13.1R1.

`pcp-rules` option added in Junos OS Release 13.2R1.

`softwire-options` option added in Junos OS Release 14.1.

`subscriber-awareness` option added in Junos OS Release 17.1R1.

`url-filter-profile` option added in Junos OS Release 17.2R1.

`match-rules-on-reverse-flow` option added in Junos OS Release 16.1R5 and 17.4R1.

`no-certificate-chain-in-ike` option added in Junos OS Release 18.2R1.

`web-filter-profile` option added in Junos OS Release 18.3R1, replacing the deprecated `url-filter-profile` option.

`max-session-setup-rate` option added in Junos OS Release 19.1R1, replacing the deprecated option `max-session-creation rate`, which was added in Junos OS Release 17.1R1.

Support added in Junos 20.2R1 for Next Gen Services NAT PT feature.

`static-subscriber-application` option added in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Understanding Service Sets

service-set (Subscriber-Aware)

IN THIS SECTION

- Syntax | **16755**
- Hierarchy Level | **16755**
- Description | **16755**
- Options | **16755**
- Required Privilege Level | **16755**
- Release Information | **16756**

Syntax

```
service-set service-set-name {
    service-set-options {
        subscriber-awareness;
        static-subscriber-application;
    }
    hcm profile hcm-profile-name;
    lrf-profile profile-name;
    pcef-profile pcef-profile-name;
    application-identification-profile app-id-profile-name;
    nat-rules rule-name;
    nat-rule-sets rule-set-name;
    disable-replication-capability;
    }
    interface-service {
        service-interface interface-name;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure subscriber-aware services by creating a subscriber-aware service set to be applied to a TDF interface.

Options

service-set-name	Name of the service set.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

`static-subscriber-application` option added in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Applying Services to Subscriber-Aware Traffic with a Service Set

Configuring Logging and Reporting for Junos OS Subscriber Aware

service-set (TDF Interface)

IN THIS SECTION

- Syntax | 16756
- Hierarchy Level | 16757
- Description | 16757
- Options | 16757
- Required Privilege Level | 16757
- Release Information | 16757

Syntax

```
service-set service-set-name;
```

Hierarchy Level

```
[edit interfaces mif unit number family inet service input],
[edit interfaces mif unit number family inet service output]
```

Description

Apply the service set to the service input and output of the TDF interface (mif) that is part of a TDF domain.

The output service set for the mif is not used by the MX Series router, but it must be configured so that the configuration commit does not fail.

Options

service-set-name Name of the service set that is being applied to the TDF interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Applying Services to Subscriber-Aware Traffic with a Service Set

service-tracking-action (Network Isolation Groups)

IN THIS SECTION

- Syntax | 16758
- Hierarchy Level | 16758
- Description | 16758
- Options | 16759
- Required Privilege Level | 16759
- Release Information | 16759

Syntax

```
service-tracking-action (lacp-out-of-sync | link-down);
```

Hierarchy Level

```
[edit protocols network-isolation group name]
```

Description

Set the action the device takes for the interface or interfaces associated with a configured network isolation group profile, when the device detects a network isolation condition change.

A network isolation group encompasses the parameters that the device uses to:

- Detect when an EVPN instance. a bridge domain, or a VLAN is in a network isolation state.
- Take actions on the associated interfaces after the device detects a network isolation status change.

You set this option as part of configuring parameters that define a network isolation group for network isolation service tracking. When you configure this option in a network isolation group, you also must:

- Enable service tracking in the network isolation group to detect network isolation condition changes.

To do this, include the "service-tracking" on page 16760 statement at the `[edit protocols network-isolation group group-name detection]` hierarchy. With the `service-tracking` statement, you set a supported network isolation condition to track (such as `core-isolation`).

- Assign the network isolation group to an interface where you want the device to perform the action when the configured network isolation condition happens.

  To do this, set the network isolation group name for the desired interface or interfaces using the `network-isolation-profile group group-name` statement at the `[edit interfaces interface-name]` hierarchy.

Options

Actions the device can perform on the interfaces associated with the network isolation group upon detecting the configured network isolation condition:

lacp-out-of-sync Set LACP out of sync state for the associated ESI LAG interfaces.

link-down Bring down the interface or interfaces.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Layer 2 Interface Status Tracking and Shutdown Actions for EVPN Core Isolation Conditions

service-tracking (Network Isolation Groups)

IN THIS SECTION

- Syntax | 16760
- Hierarchy Level | 16760
- Description | 16760
- Options | 16761
- Required Privilege Level | 16761
- Release Information | 16761

Syntax

```
service-tracking {
    core-isolation;
}
```

Hierarchy Level

```
[edit protocols network-isolation group name detection]
```

Description

Enable service tracking in a network isolation group to detect core isolation condition changes.

When you set this option, you must also do the following:

- In the same network isolation group definition, configure the actions the device should take upon detecting a core isolation condition.

  To do this, configure the `service-tracking-action` statement at the `[edit protocols network-isolation group name detection]` hierarchy.

- Assign the network isolation group to the interface or interfaces on which the device should take action upon detecting a core isolation condition.

To do this, configure the `network-isolation-profile group group-name` statement at the `[edit interfaces interface-name]` hierarchy.

The interface to which you assign these parameters can be:

- A logical single link or aggregated Ethernet (AE) link bundle (LAG) to a single-homed customer edge device in the network.
- An Ethernet segment identifier (ESI) LAG interface to a multihomed customer edge device in the network.
- A physical interface.

Options

core-isolation Enable service tracking for EVPN core isolation conditions.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

service-tracking-action (Network Isolation Groups) | **16758**

Layer 2 Interface Status Tracking and Shutdown Actions for EVPN Core Isolation Conditions

service-type

IN THIS SECTION

Syntax | **16762**

- Hierarchy Level | 16762
- Description | 16762
- Options | 16762
- Required Privilege Level | 16762
- Release Information | 16763

Syntax

```
service-type (elan | eline);
```

Hierarchy Level

```
[edit routing-instances (Multiple Routing Entities) routing-instance-name service-groups]
```

Description

For IEEE 802.1ah provider backbone bridge (PBB) configurations, configure the service type for the service group in the customer routing instance (I-component).

Options

`elan`—Connects a set of customer endpoints (like a bridged Ethernet network). E-LAN service is also known as point-to-multipoint service.

`eline`—Connects two customer Ethernet ports over a WAN. E-LINE service is also known as point-to-point service.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.0.

service-type (RFC 2544 Benchmarking Tests)

IN THIS SECTION

- Syntax | 16763
- Hierarchy Level | 16763
- Description | 16763
- Options | 16764
- Required Privilege Level | 16764
- Release Information | 16764

Syntax

```
service-type (elan | eline) ;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Configure the service under test. Possible values are `elan` and `eline`. This statement is applicable only for the `bridge` family, the `ethernet-switching` family, or when the `mode` is configured as reflect. When the service type is `elan`, MAC addresses are swapped by default on the reflected frames. The `no-mac-swap` option is not supported for this service type. When the service type is `eline`, MAC addresses are not swapped by default in the reflected frames. Use the `mac-swap` option to swap the addresses.

> **NOTE**: When you configure Layer 2 reflection, you can specify the service type under test as `eline` if you want to simulate an E-Line service using bridge encapsulation.

Options

elan Specify the ELAN service type.

eline Specify the E-Line service type.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

rfc2544-benchmarking | **15759**

rfc2544 | **15757**

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

service-type (MAC-VRF Routing Instances)

IN THIS SECTION

Syntax | **16765**

- Hierarchy Level | 16765
- Description | 16765
- Options | 16765
- Required Privilege Level | 16766
- Release Information | 16766

Syntax

```
service-type (vlan-aware | vlan-based | vlan-bundle);
```

Hierarchy Level

```
[edit logical-systems name routing-instances mac-vrf routing-instance name],
[edit logical-systems name tenants name routing-instances mac-vrf routing-instance name],
[edit routing-instances mac-vrf routing-instance name],
[edit tenants name routing-instances mac-vrf routing-instance name]
```

Description

Select the type of VLAN service that the MAC-VRF routing instance provides. The selected option controls how the routing instance maps VLANs to the forwarding instances on the device.

Options

vlan-aware Map one or more routing instances of type Virtual Switch to many VLAN IDs (VIDs) and multiple bridge tables. Each bridge table corresponds to a different VLAN.

vlan-based Map one routing instance of type EVPN to one VLAN. Only one bridge table corresponds to one VLAN. If the VLAN consists of multiple VLAN IDs (VIDs), then you must configure VLAN translation for packets that are traveling to the Ethernet segment. You do this when you use a different VID per Ethernet segment on a provider edge device.

vlan-bundle Map multiple broadcast domains to a single bridge domain. Multiple VLANs map to a single EVPN instance (EVI) and share the same bridge table in the MAC-VRF table. This configuration reduces the number of routes and labels stored in the table.

CAUTION: When you want to change a traffic impacting option under a routing instance, use the following procedure.

1. Deactivate the routing instance configuration.
2. Change the traffic impacting option.
3. Reactivate the updated routing instance configuration.

For example, follow this procedure if you need to change settings such as:

- The `service-type` in a MAC-VRF routing instance—When you change the service type of a running instance, the device might incorrectly change the VLAN ID if it is not deactivated prior to making the change.
- The `vlan-id` in an EVPN routing instance—Changing the `vlan-id` on the fly, without first deactivating the associated EVPN routing instance would be catastrophic.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

MAC-VRF Routing Instance Type Overview

service-provider (SAML)

IN THIS SECTION

- Syntax | 16767
- Hierarchy Level | 16767
- Description | 16767
- Options | 16768
- Required Privilege Level | 16768
- Release Information | 16768

Syntax

```
service-provider service-provider {
    assertion-waittime assertion-waittime;
    encrypted-assertions;
    entity-id entity-id;
    hash-algorithm <sha-256 | sha-384 | sha-512>;
    sign-auth-request;
    sp-certificate sp-certificate;
}
```

Hierarchy Level

```
[edit access saml]
```

Description

Describe Security Assertion Markup Language (SAML) service provider configuration details.

Options

assertion-waittime ***assertion-waittime***	Set the maximum waiting time for an assertion from the Identity Provider (IdP). • Value: Within a range of 1 to 600 seconds • Default: 30 seconds
encrypted-assertions	Specify that the IdP must provide encrypted assertions.
entity-id ***entity-id***	Specify service provider entity ID.
hash-algorithm <sha-256 | sha-384 | sha-512>	Define the signature hash algorithm. • Value: Supports sha-256, sha-384, or sha-512 • Default: sha-256
sign-auth-request	Sign the SAML authentication request.
sp-certificate ***sp-certificate***	Specify the service provider certificate.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

saml | **16281**

services (Captive Portal Content Delivery)

IN THIS SECTION

- Syntax | 16769
- Hierarchy Level | 16770
- Description | 16770
- Options | 16771
- Required Privilege Level | 16771
- Release Information | 16771

Syntax

```
services {
    ...
    captive-portal-content-delivery {
        auto-deactivate value;
        profile name
            cpcd-rule-sets rule-set-name;
            cpcd-rules rule-name;
            dynamic;
            http-redirect-options url;
            ipda-rewrite-options {
                destination-address destination-address;
                destination-port destination-port;
            }
        }
        rule rule-name {
            match-direction (input | output | input-output);
            from {
                destination-address address <except>;
            }
            term term-name {
                then {
                    accept;
                    insert tag tag-name tag-value tag-value;
                    redirect url;
```

```
                    rewrite destination-address address <destination-port port-number>;
                    syslog;
                }
            }
        }
        rule-set rule-set-name {
            [rule rule-name];
        }
        traceoptions {
            file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
            flag name;
            no-remote-trace no-remote-trace;
        }
    }
}
```

Hierarchy Level

```
[edit],
[edit dynamic-profiles profile-name]
```

Description

Define the captive portal content delivery set of the rules statements to be applied to traffic. Use the statement at the `[edit services...]` hierarchy level for static CPCD. Use the statement at the `[edit dynamic-profiles profile-name services...]` hierarchy level for converged services CPCD.

The `profile`, `rule-set`, and `traceoptions` stanzas are not supported at the `[edit dynamic-profiles profile-name hierarchy level]`.

Configure the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber. Use the statement at the `[edit services...]` hierarchy level for static CPCD. Use the statement at the `[edit dynamic-profiles profile-name services...]` hierarchy level for converged services CPCD.

The `profile`, `rule-set`, and `traceoptions` stanzas are not supported at the `[edit dynamic-profiles profile-name hierarchy level]`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

auto-deactivate *value*

(Optional) Specify one of the following values to determine whether the redirect service is removed automatically when the router receives the subscriber's initial HTTP-GET message:

- *user-defined-variable*—To use this option in a dynamic profile, you must create a user-defined variable with a name of your choice. The value of the variable can be supplied by the RADIUS server or PCRF. You can also define a default value that is used when the external servers do not supply it. Use the "variables" on page 30654 statement in the dynamic profile to define the default value. Whether supplied by the external server or by the CLI, the value must be either `initial-get` or `never`.
- `initial-get`—Receipt of the subscriber's initial HTTP-GET message triggers removal of the redirect service.
- `never`—Removal of the redirect service is triggered only by the external server, such as by a CoA message from a RADIUS server.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the `[edit dynamic-profiles` *`profile-name`* `services]` hierarchy level added in Junos OS Release 17.2R1.

Statement introduced in Junos OS Release 10.4.

Support at the `[edit dynamic-profiles` *`profile-name`* `services]` hierarchy level added in Junos OS Release 17.2R1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

`auto-deactivate` option added in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

HTTP Redirect Service Overview

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

Configuring Routing Engine-Based, Static HTTP Redirect Services

Configuring Routing Engine-Based, Converged HTTP Redirect Services

Adding Subscriber Information to HTTP Redirect URLs

services (CoS)

IN THIS SECTION

- Syntax | 16772
- Hierarchy Level | 16772
- Description | 16773
- Options | 16773
- Required Privilege Level | 16773
- Release Information | 16773

Syntax

```
services cos { ... }
```

Hierarchy Level

```
[edit]
```

Description

Define the service rules to be applied to traffic.

Options

`cos`—Identifies the class-of-service set of rules statements.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring CoS Rules

services (Hierarchy)

IN THIS SECTION

- Syntax | 16774
- Hierarchy Level | 16774
- Description | 16774
- Required Privilege Level | 16774
- Release Information | 16774

Syntax

```
services { ... }
```

Hierarchy Level

```
[edit]
```

Description

Define the service rules to be applied to traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding Service Sets

services (IDS)

IN THIS SECTION

- Syntax | 16775
- Hierarchy Level | 16775
- Description | 16775

- Options | 16775
- Required Privilege Level | 16775
- Release Information | 16775

Syntax

```
services ids { ... }
```

Hierarchy Level

```
[edit]
```

Description

Define the service rules to be applied to traffic.

Options

`ids`—Identifies the IDS set of rules statements.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring IDS Rules on an MS-DPC*

services (Interfaces)

IN THIS SECTION

- Syntax | 16776
- Hierarchy Level | 16776
- Description | 16776
- Options | 16776
- Required Privilege Level | 16777
- Release Information | 16777

Syntax

```
services severity-level;
```

Hierarchy Level

```
[edit interfaces interface-name services-options syslog host hostname]
```

Description

Specify the system logging severity level.

Starting with Junos OS release 17.4R1, you can configure up to a maximum of four system log servers (combination of local system log hosts and remote system log collectors) for each service set for ms interface under `[edit interfaces interface-name services-options]` hierarchy.

Options

severity-level—Assigns a severity level to the facility. Valid entries include:

- `alert`—Conditions that should be corrected immediately.
- `any`—Matches any level.

- `critical`—Critical conditions.
- `emergency`—Panic conditions.
- `error`—Error conditions.
- `info`—Informational messages.
- `notice`—Conditions that require special handling.
- `warning`—Warning messages.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

You can configure multiple system log hosts from Junos OS Release 17.4R1 onwards.

RELATED DOCUMENTATION

Configuring System Logging for Services Interfaces

services (IPsec VPN)

IN THIS SECTION

- Syntax | 16778
- Hierarchy Level | 16778
- Description | 16778
- Options | 16778
- Required Privilege Level | 16778
- Release Information | 16778

Syntax

```
services ipsec-vpn { ... }
```

Hierarchy Level

```
[edit]
```

Description

Define the service rules to be applied to traffic.

Options

ipsec-vpn—Identifies the IPsec set of rules statements.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

services (L2TP)

IN THIS SECTION

- Syntax | 16779
- Hierarchy Level | 16779
- Description | 16779
- Options | 16779

Required Privilege Level | 16779
Release Information | 16779

Syntax

```
services l2tp { ... }
```

Hierarchy Level

```
[edit]
```

Description

Define the service properties to be applied to traffic.

Options

`l2tp`—Identifies the L2TP set of services statements.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

L2TP Services Configuration Overview

services (L2TP System Logging)

IN THIS SECTION

- Syntax | 16780
- Hierarchy Level | 16780
- Description | 16780
- Options | 16780
- Required Privilege Level | 16781
- Release Information | 16781

Syntax

```
services severity-level;
```

Hierarchy Level

```
[edit services l2tp tunnel-group group-name syslog host hostname]
```

Description

Specify the system logging severity level.

Options

severity-level—Assigns a severity level to the facility. Valid entries include:

- `alert`—Conditions that should be corrected immediately.
- `any`—Matches any level.
- `critical`—Critical conditions.
- `emergency`—Panic conditions.

- `error`—Error conditions.
- `info`—Informational messages.
- `notice`—Conditions that require special handling.
- `warning`—Warning messages.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring L2TP Tunnel Groups

services-memory (resource-manager)

IN THIS SECTION

- Syntax | 16782
- Hierarchy Level | 16782
- Description | 16782
- Options | 16782
- Required Privilege Level | 16783
- Release Information | 16783

Syntax

```
services-memory {
    match {
        threshold threshold;
    }
    then {
        alarm;
        drop-flows;
        fail-close;
        fail-open;
        log;
    }
}
```

Hierarchy Level

```
[edit security resource-manager]
```

Description

Configure resource manager services-memory options to control the behavior of the device based on memory load.

Options

- `threshold` *`threshold`*-Configure the threshold limit to let `resource-manager` take actions when the service memory load exceeds the threshold. The range is 1 to 100 percent.
- `fail-close`-This configuration closes a session when there is not enough memory for a service during session setup. This configuration drops a packet if there is not memory during fast path. This is the default configuration.
- `fail-open`-This configuration opens a infra session even if there is not enough memory for a service to allocate memory.
- `drop-flows`-When you mark a session as a drop-flow, it bypasses the first path for each new packet, which matches the original drop session for next 4 seconds only. This configuration drops the packets that match the `drop-flows`.

- `log`-This configuration logs messages when the resource threshold is exceeded and falls back under the threshold, and what is the protective action taken.
- `alarm`-This configuration sends an SNMP trap when the memory load exceeds the threshold.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Traffic Processing on SRX Series Devices Overview

services (NAT)

IN THIS SECTION

- Syntax | 16783
- Hierarchy Level | 16784
- Description | 16784
- Options | 16784
- Required Privilege Level | 16784
- Release Information | 16784

Syntax

```
services nat { ... }
```

Hierarchy Level

```
[edit]
```

Description

Define the service rules to be applied to traffic.

Options

nat—Identifies the NAT set of rules statements.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

services-offload

IN THIS SECTION

- Syntax | 16785
- Hierarchy Level | 16785
- Description | 16785
- Options | 16785
- Required Privilege Level | 16785
- Release Information | 16785

Syntax

```
services-offload {
    low-latency;
    per-session-statistics;
}
```

Hierarchy Level

```
[edit chassis fpc fpc-slot-number pic pic-slot-number]
```

Description

Enables the Express Path mode (formerly known as *services offloading*) mode on the selected network processor. Services-offload is not enabled by default.

When services-offload is enabled, only the first packets of a session goes to the Services Processing Unit (SPU), rest of packets in services-offload mode does not go to SPU, therefore some security features such as stateful screen are not supported. Only TCP and UDP packets can be services offloaded.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

services-offload (Security)

IN THIS SECTION

- Syntax | 16786
- Hierarchy Level | 16786
- Description | 16786
- Required Privilege Level | 16786
- Release Information | 16786

Syntax

```
services-offload;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit]
```

Description

Enable services offloading within a security policy for SRX4600, SRX5400, SRX5600, and SRX5800 devices.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Elements

services-options

IN THIS SECTION

- Syntax | 16787
- Hierarchy Level | 16788
- Description | 16788
- Required Privilege Level | 16788
- Release Information | 16788

Syntax

```
services-options {
    cgn-pic;
    close-timeout seconds;
    fragment-limit number-of-fragments;
    disable-global-timeout-override;
    ignore-errors <alg> <tcp>;
    inactivity-non-tcp-timeout seconds;
    inactivity-tcp-timeout seconds;
    inactivity-timeout seconds;
    open-timeout seconds;
    pba-interim-logging-interval seconds;
    reassembly-timeout seconds;
    session-limit {
        maximum number;
        rate (Interface Services) new-sessions-per-second;
        cpu-load-threshold percentage;
    }
    session-timeout seconds;
    jflow-log {
```

```
        message-rate-limit messages-per-second;
    }
    syslog {
        host hostname {
            facility-override facility-name;
            log-prefix prefix-value;
            port port-number;
            services severity-level;
        }
        message-rate-limit messages-per-second;
    }
    tcp-tickles tcp-tickles;
    trio-flow-offload minimum-bytes minimum-bytes;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Define the service options to be applied on an interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Default Timeout Settings for Services Interfaces

Configuring System Logging for Services Interfaces

services-options (Next Gen Services Interfaces)

IN THIS SECTION

- Syntax | 16789
- Hierarchy Level | 16790
- Description | 16790
- Options | 16790
- Required Privilege Level | 16793
- Release Information | 16793

Syntax

```
services-options {
    enable-subscriber-analysis
fragment-limit;
    jflow-log {
        message-rate-limit messages-per-second;
    }
    session-limit {
        maximum number;
        rate new-sessions-per-second;
        cpu-load-threshold percentage;
    }
        flow
            traceoptions {
                file {
                    filename;
                    files number;
                    match regular-expression;
                    size maximum-file-size;
                    (world-readable | no-world-readable);
                }
                flag flag;
                no-remote-trace;
                packet-filter filter-name {
                    conn-tag session-conn
```

```
                destination-port port-identifier;
                destination-prefix address;
                interface interface-name;
                protocol protocol-identifier;
                source-port port-identifier;
                source-prefix address;
            }
            rate-limit messages-per-second;
            trace-level (brief | detail | error);
        }
```

Hierarchy Level

```
[edit interfaces interfaces-name]
```

Description

Define the service options to be applied on the virtual multi-service (VMS) interface.

This statement is supported only on the MX-SPC3 Services Card.

The remaining statements are explained separately. See CLI Explorer.

Options

file Configure the trace file options.

filename Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.

files *number* Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000 files
- **Default:** 10 files

match *regular-expression*

Refine the output to include lines that contain the regular expression.

size *maximum-file-size*

Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file*.0. When the trace-file again reaches its maximum size, *trace-file*.0 is renamed *trace-file*.1 and *trace-file* is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the files option and a filename.

Syntax: *x* K to specify KB, *x* m to specify MB, or *x* g to specify GB

- **Range:** 0 KB through 1 GB
- **Default:** 128 KB

world-readable | no-world-readable

By default, log files can be accessed only by the user who configures the tracing operation. The world-readable option enables any user to read the file. To explicitly set the default behavior, use the no-world-readable option.

flag

Trace operation to perform. To specify more than one trace operation, include multiple flag statements.

all

Trace with all flags enabled

basic-datapath

Trace basic packet flow activity

fragmentation

Trace IP fragmentation and reassembly events

high-availability

Trace flow high-availability information

host-traffic

Trace flow host traffic information

multicast

Trace multicast flow information

route

Trace route lookup information

session	Trace session creation and deletion events
session-scan	Trace session scan information
tcp-basic	Trace TCP packet flow information
tunnel	Trace tunnel information

no-remote-trace
Set remote tracing as disabled.

packet-filter *filter-name*
Packet filter to enable during the tracing operation. Configure the filtering options.

destination-port *port-identifier*	Match TCP/UDP destination port
destination-prefix *address*	Destination IP address prefix
interface *interface-name*	Logical interface
protocol *protocol-identifier*	Match IP protocol type
source-port *port-identifier*	Match TCP/UDP source port
source-prefix *address*	Source IP address prefix

rate-limit *messages-per-second*
Limit the incoming rate of trace messages.

trace-level
Set the level for trace logging. This option is available only when the flag is set.

brief	Trace key flow information, such as message types sent between SPU and central point, policy match, and packet drop reasons.
detail	Trace extensive flow information, such as detailed information about sessions and fragments. Detail is the default level.
error	Trace error information, such as system failure, unknown message type, and packet drop.

fragment-limit
Specify the maximum number of fragments to be supported for the PIC. This overrides the value specified, if any, in the `set security flow fragment-limit` statement.

reassembly-timeout
Specify the reassembly timeout value for all fragmentation packets for the PIC. This overrides the value specified, if any, in the `set security flow reassembly-timeout` statement

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

Support introduced in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480 and MX960 routers for the `flow` configuration statement.

services-redundancy-group

IN THIS SECTION

- Syntax | 16793
- Hierarchy Level | 16795
- Description | 16796
- Options | 16796
- Required Privilege Level | 16798
- Release Information | 16798

Syntax

```
services-redundancy-group name {
 activeness-priority activeness-priority;
 activeness-probe {
  bfd-liveliness {
   destination-ip;
   interface;
   minimum-interval;
   multiplier;
```

```
  routing-instance
  source-ip;
 }
 dest-ip {
  ip-address;
  routing-instance routing-instance;
  src-ip src-ip;
 }
 minimum-interval minimum-interval;
 multiplier multiplier;
}
active-signal-route {
  ip-address;mandatory
 routing-instance routing-instance;
}
backup-signal-route {
  ip-address;mandatory
 routing-instance routing-instance;
}
deployment-type(cloud | hybrid | routing | switching);
install-on-failure-route {
 ip-address;
 routing-instance routing-instance;
}
managed-services name;

monitor {
 bfd-liveliness name {
  interface interface;
  routing-instance routing-instance;
  session-type(multihop | singlehop);
  src-ip src-ip;
 }
 interface name;

 ip name {
  routing-instance routing-instance;
 }
 monitor-object name {
  bfd-liveliness {
   destination-ip name {
    interface interface;
    routing-instance routing-instance;
```

```
      session-type(multihop | singlehop);
      src-ip src-ip;
      weight weight;
     }
     threshold threshold;
    }
    interface {
     interface-name name {
      weight weight;
     }
     threshold threshold;
    }
    ip {
     destination-ip name {
      routing-instance routing-instance;
      weight weight;
     }
     threshold threshold;
    }
    object-threshold object-threshold;
   }
  }
  peer-id id;
  preemption;
  prefix-list name {
   routing-instance routing-instance;
  }
  process-packet-on-backup;
  shutdown-on-failure name;

  virtual-ip name {
   interface interface;
   ip ip;mandatory
   use-virtual-mac;
  }
 }
```

Hierarchy Level

```
[edit chassis high-availability]
```

Description

Configure a service redundancy group (SRG). An SRG includes and manages a collection of resources on both nodes of a Multinode High Availability and it fails over between the two devices. You can configure upto 20 SRGs in a Multinode Highavailability setup.

Options

name

Services redundancy group identifier.

- **Range:** 0 through 20

active-signal-route

IP address used for route preference advertisement. You must specify the active signal route along with the route-exists policy in the policy-options statement.

Signal route required for active role enforcement. When you configure the `active-signal-route` with `if-route-exists` condition, the HA module adds this route to the routing table.

- ip-address—IP address for active signal route
- routing-instance—Routing instance of the active signal route.

activeness-priority

Specify priority for the SRG1 in a node to take up the active role in a case where both nodes initialize at the same time. The node where SRG1 is in active state is considered as active node.

If you prefer a certain node to take over as the active node on boot, you can do one of the followings:

- Configure the upstream routers to include preferences for the path where the node is located.
- Configure the activeness priority for SRG1 on the SRX Series device (higher activeness priority). You can configure a priority for each node. As long as the nodes can communicate with each other through the ICL, the priority is honored.
- Allow the node with higher node ID (in case above two options not configured) to take the active role.
- **Range:** 1 through 254

"activeness-probe" on page 1821

Specify the probe destination IP address for activeness determination.

backup-signal-route

Specify the backup signal route to advertise a route with a medium priority. When the HA link is down or the current active node relinquishes active role after any failure, the active signal route is removed from the routing table. The backup overwrites the default routing preference toward the old active node with the medium priority.

- ip-address—IP address for backup signal route
- routing-instance—Routing instance of the backup signal route.

deployment-type

Deployment type of the Services Redundancy Group.

- cloud—Cloud deployment
- hybrid—Hybrid deployment
- routing—Routing deployment
- switching—Switching/default gateway deployment

install-on-failure-route

Divert the traffic by changing the route in a Multinode High Availability setup during a software upgrade. In this case, traffic still traverses through the node and interface remains up.

- *ip-address*—IP address of the route. Mutinode High Availability installs this route to divert the traffic during the upgrade.
- routing-instance *routing-instance*—Routing instance. You must create a dedicated custom virtual router for the route used for diverting traffic during the upgrade.

managed-services

Enable set of services for the specified SRG.

"monitor" on page 10903

Specify to configure the BFD and IP monitoring options.

peer-id

Allows you to choose a specific peer when multiple HA peers are configured globally to the service redundancy group.

preemption

Allow preemption of activeness based on priority. When you configure the activeness priority (1-254) for the SRG1 and enable the preemptive behavior on both nodes, the preempt option ensures that the node with higher activeness priority always remains active after a failover.

prefix-list

Define a named set of address prefixes.

Associate the IP prefix list to the SRG. A prefix list is a listing of IP prefixes that include the local address of IKE gateway. A prefix list is given a name and is configured within the [edit policy-options] configuration hierarchy.

process-packet-on-backup	Enable packet forward engine to forward packets on backup node for the corresponding service redundancy group. When you configure the process packet on back up option, the Packet Forwarding Engine forwards packets on backup node for the corresponding SRG. This configuration processes VPN packets on the backup node even when the node is not active.
services	Enable IPsec VPN service on a particular SRG.
shutdown-on-failure	Configure one or multiple Interfaces which are required to be shut down to isolate the node in case of internal failures or during software upgrades. During software upgrades, you can divert the traffic by closing down interfaces on the node.
"virtual-ip" on page 30761	IP address used for activeness determination and enforcement on the switching side. Required for hybrid and default gateway deployments.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 20.4R1.

Multi SRG1s (SRG1+) support is added in Junos OS Release 22.4R1.

install-on-failure-route option is added in Junos OS Release 22.4R2.

RELATED DOCUMENTATION

Example: Configure Multinode High Availability in a Layer 3 Network

Example: Configure Multinode High Availability in a Default Gateway Deployment

Example: Configure Multinode High Availability in a Default Gateway Deployment

services (Stateful Firewall)

IN THIS SECTION

- Syntax | 16799
- Hierarchy Level | 16799
- Description | 16799
- Options | 16799
- Required Privilege Level | 16799
- Release Information | 16800

Syntax

```
services stateful-firewall { ... }
```

Hierarchy Level

```
[edit]
```

Description

Define the service rules to be applied to traffic.

Options

`stateful-firewall`—Identifies the stateful firewall set of rules statements.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos Network Secure Overview

services (System Logging)

IN THIS SECTION

- Syntax | 16800
- Hierarchy Level | 16800
- Description | 16800
- Options | 16801
- Required Privilege Level | 16801
- Release Information | 16801

Syntax

```
services severity-level;
```

Hierarchy Level

```
[edit services service-set service-set-name syslog host hostname]
```

Description

Specify the severity level for system logging messages.

Starting in Junos OS Release 17.4R1, you can configure up to a maximum of four system log servers (combination of local system log hosts and remote system log collectors) for each service set under [edit services service-set *service-set-name*] hierarchy level.

Options

severity-level—Assigns a severity level to the facility. Valid entries are:

- alert—Conditions that should be corrected immediately.
- any—Matches any level.
- critical—Critical conditions.
- emergency—Panic conditions.
- error—Error conditions.
- info—Informational messages.
- notice—Conditions that require special handling.
- warning—Warning messages.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

You can configure multiple system log hosts from Junos OS Release 17.4R1 onwards.

RELATED DOCUMENTATION

Configuring System Logging for Service Sets

services (system services)

IN THIS SECTION

- Syntax | 16802
- Hierarchy Level | 16803
- Description | 16803
- Options | 16803
- Required Privilege Level | 16804
- Release Information | 16804

Syntax

```
services {
    bbe-stats-service {...}
    database-replication {...}
    dhcp {...}
    dhcp-local-server {...}
    dns {...}
    dtcp-only;
    dynamic-dns {...}
    extensible-subscriber-services {...}
    extension-service {...}
    finger {...}
    flow-tap-dtcp {...}
    ftp {...}
    grpc-tunnel {...}
    jeb {...}
    netconf {...}
    netproxy;
    outbound-https {...}
    resource-monitor {...}
    rest {...}
    reverse {...}
    ssh {...}
    static-subscribers {...}
    subscriber-management {...}
```

```
    subscriber-management-helper {...}
    telnet {...}
    tftp-server {...}
    webapi {...}
    web-management {...}
    xnm-clear-text {...}
    xnm-ssl {...}
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the router or switch so that users on remote systems can access the local router or switch through the DHCP server, DTCP over SSH, finger, outbound HTTPS, rlogin, SSH, telnet, Web management, Junos XML protocol SSL, and network utilities. Also, enable configuration of third-party applications developed using the Juniper Extension Toolkit (JET) to run on Junos OS.

The options `dynamic-dns`, `reverse` and `webapi` option is available only for SRX series devices. For further information, see Related Information section.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

bbe-stats-service Change the default values of BBE subscriber accounting statistics. The accounting statistics are polled or pushed from the PFE to RE depending on the FPC type, to send to RADIUS server

dtcp-only Allow subscriber DTCP based lawful intercept only

netproxy Enable netproxy service for host OS applications. Netproxy proxies TCP/UDP connections of host OS applications to the netproxy-server running in Junos OS for network access via data and management ports

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`extension-service` option added in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.

`grpc` option added in Junos OS Release 16.2 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.

`allow-tcp-forwarding` option added in Junos OS Release 22.2R1.

session-affinity

IN THIS SECTION

- Syntax | 16804
- Hierarchy Level | 16805
- Description | 16805
- Required Privilege Level | 16805
- Release Information | 16805

Syntax

```
session-affinity ipsec
```

Hierarchy Level

```
[edit security flow load-distribution]
```

Description

Enable VPN session affinity. In session affinity feature, we've optimized tunnel redistribution. After tunnel redistribution, the data path might not be optimal, hence we recommend that you enable VPN session affinity to ensure that the data path is optimized. During optimization, the current data path experiences a higher packet delay until it is fully optimized.

This feature is supported on SRX5400, SRX5600, and SRX5800 devices. By default, VPN session affinity is disabled.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4R5.

Starting with Junos OS Release 15.1X49-D10, IPsec session affinity is supported for IPsec tunnel-based traffic by the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) for SRX5400, SRX5600, and SRX5800 devices through improved flow module and session cache.

RELATED DOCUMENTATION

IPsec Overview

session-close

IN THIS SECTION

- Syntax | 16806
- Hierarchy Level | 16806
- Description | 16806
- Required Privilege Level | 16806
- Release Information | 16806

Syntax

```
session-close;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then log]
```

Description

Enable traffic to which the policy applies to be logged at the end of a session.

Session-ager is a process that runs at regular intervals and cleans up inactive sessions. The timestamp of the session-close log is the time that session-ager closes the session.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

session-drop-hold-down

IN THIS SECTION

- Syntax | 16807
- Hierarchy Level | 16807
- Description | 16807
- Required Privilege Level | 16807
- Release Information | 16808

Syntax

```
session-drop-hold-down time;
```

Hierarchy Level

[edit security nat source]

Description

Specify the session hold time value to hold the NAT source session without expiring. The session hold time value must be 30 seconds through 28,800 seconds (eight hours).

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

session-establishment-hold-time

IN THIS SECTION

- Syntax | 16808
- Hierarchy Level | 16808
- Description | 16808
- Options | 16809
- Required Privilege Level | 16809
- Release Information | 16809

Syntax

```
session-establishment-hold-time seconds;
```

Hierarchy Level

```
[edit protocols iccp],
[edit protocols iccp peer]
```

Description

Specify the time during which an Inter-Chassis Control Protocol (ICCP) connection must be established after IP route reachability between MC-LAG peers is up. When an MC-LAG peer detects IP route reachability to the MC-LAG peer, it tries to connect to it during the session-establishment-hold-time.

> **NOTE**: On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

Options

seconds Time (in seconds) within which a successful ICCP connection must be established.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

session-expiry (MX Series in Enhanced LAN Mode)

IN THIS SECTION

- Syntax | 16810
- Hierarchy Level | 16810
- Description | 16810
- Options | 16810
- Required Privilege Level | 16810
- Release Information | 16810

Syntax

```
session-expiry seconds;
```

Hierarchy Level

```
[edit protocols authentication-access-control interface (all | [interface-names])]
```

Description

Configure the maximum duration in seconds of a session.

Options

seconds—Duration of session.

- **Range:** 1 through 65535
- **Default:** 3600

Required Privilege Level

routing—To view this statement in the configuration.

routing–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

session-group

IN THIS SECTION

Syntax | 16811

- Hierarchy Level | 16811
- Description | 16811
- Options | 16812
- Required Privilege Level | 16813
- Release Information | 16813

Syntax

```
session-group name {
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key authentication-key;
    authentication-key-chain authentication-key-chain;
    downstream-on-demand;
    (mtu-discovery | no-mtu-discovery);
    transport-address transport-address;
}
```

Hierarchy Level

```
[edit logical-systems name protocols ldp],
[edit logical-systems name routing-instances name protocols ldp],
[edit protocols ldp],
[edit routing-instances name protocols ldp]
```

Description

Specify the prefix address of the aggregated group of LDP neighbors for the remote end of the LDP session.

The `session-group` statement is useful when an LDP neighbor is dynamic; for instance, in the case of remote loop free alternate (LFA), a targeted or indirect LDP neighbor is automatically picked from any of the nodes in the network.

The `session-group` statement can also be used for configuring authentication for a prefix group, as LDP authentication for all sessions or authentication at the interface level is not supported.

The remaining statements are explained separately. See CLI Explorer.

Options

name

Session destination address or prefix length.

authentication-algorithm

Authentication algorithm name.

- Values:
  - *aes-128-cmac-96*—Cipher-based Message Authentication Code (AES128) (96 bits)
  - *hmac-sha-1-96*—Hash-based Message Authentication Code (SHA1) (96 bits)
  - *md5*—Message Digest 5

> **NOTE**: MD5 authentication algorithm is supported in both Junos OS and Junos OS Evolved Releases. Whereas, aes-128-cmac-96 and hmac-sha-1-96 algorithms are not supported in Junos OS Evolved Releases

authentication-key

MD5 authentication key.

authentication-key-chain

Authentication key chain name. Authentication key chain option is not supported in Junos OS Evolved Releases.

downstream-on-demand

Configure downstream on demand label distribution mode.

mtu-discovery | no-mtu-discovery

Enable and disable TCP path MTU discovery, respectively.

transport-address *transport-address*

IP address used for TCP sessions to the LDP neighbors that have targeted-LDP adjacencies, and fall under the same IP subnet.

You cannot configure this statement for LDP sessions that have discovered adjacencies.

The transport-address configuration can be rejected at the time of validation, if there is no interface with the configured IP address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`transport-address` option introduced in Junos OS Releasse 19.1R1 for all platforms.

Starting in Junos OS Evolved 22.1R1 Release, MD5 authentication algorithm is supported for all platforms.

RELATED DOCUMENTATION

Configuring the TCP MD5 Signature for LDP Sessions

session-id-change-limiter-indirect

IN THIS SECTION

- Syntax | 16814
- Hierarchy Level | 16814
- Description | 16814
- Default | 16814
- Required Privilege Level | 16814
- Release Information | 16814

Syntax

```
forwarding-table {
    session-id-change-limiter-indirect;
}
```

Hierarchy Level

```
[edit routing-options]
```

Description

Creates additional indirect nexthop in the nexthop chain to store session-id changes. When a session-id is changed in any of the indirect nexthop, the Packet Forwarding Engine must reprogram all the parent nodes of the indirect-nexthop. To limit the number of parent nodes that will be re-programmed, an additional indirect nexthop is added in the nexthop chain. This enables the Packet Forwarding Engine to reprogram only the additional indirect-nexthop.

Default

By default, the `session-id-change-limiter-indirect` configuration statement is enabled on the platforms that has Packet Forwarding Engine. To disable the configuration statement, use the `no-session-id-change-limiter` configuration statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4.

Statement introduced in Junos OS Evolved Release 22.2.

session-init

IN THIS SECTION

- Syntax | 16815
- Hierarchy Level | 16815
- Description | 16815
- Required Privilege Level | 16815
- Release Information | 16815

Syntax

```
session-init;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then log]
```

Description

Enable traffic to which the policy applies to be logged at the beginning of a session.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

session-limit

IN THIS SECTION

- Syntax | 16816
- Hierarchy Level | 16816
- Description | 16816
- Required Privilege Level | 16817
- Release Information | 16817

Syntax

```
session-limit {
    maximum number;
    rate (Interface Services) new-sessions-per-second;
    cpu-load-threshold percentage;
}
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Description

Restrict the maximum number of sessions and the session rate on services cards.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

session-limit (IDS MS-DPC)

IN THIS SECTION

- Syntax | 16817
- Hierarchy Level | 16818
- Description | 16818
- Options | 16818
- Required Privilege Level | 16818
- Release Information | 16818

Syntax

```
session-limit {
    by-destination (IDS MS-DPC) {
        hold-time seconds;
        maximum number;
        packets number;
        rate number;
    }
    by-pair (IDS MS-DPC) {
        hold-time seconds;
```

```
        maximum number;
        packets number;
        rate number;
    }
    by-source (IDS MS-DPC) {
        hold-time seconds;
        maximum number;
        packets number;
        rate number;
    }
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Enable flow limitation by configuring thresholds on source, destination, or stateful firewall and network address translation (NAT) paired traffic flows when using the MS-DPC.

Options

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in IDS Rules

session-limit (IDS MS-MPC)

IN THIS SECTION

- Syntax | 16819
- Hierarchy Level | 16820
- Description | 16820
- Required Privilege Level | 16821
- Release Information | 16821

Syntax

```
session-limit {
    by-destination {
        by-protocol {
            icmp {
                maximum number;
                packets number;
                rate number;
            }
            tcp {
                maximum number;
                packets number;
                rate number;
            }
            udp {
                maximum number;
                packets number;
                rate number;
            }
        }
        maximum number;
        packets number;
        rate number;
    }
    by-source {
        by-protocol {
```

```
            icmp {
                maximum number;
                packets number;
                rate number;
            }
            tcp {
                maximum number;
                packets number;
                rate number;
            }
            udp {
                maximum number;
                packets number;
                rate number;
            }
        }
        maximum number;
        packets number;
        rate number;
    }
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Configure the IDS rule session limits for an individual destination or source address or subnet. This protects against network probing attacks and network flooding attacks. This IDS rule can only be assigned to a service set on an MS-MPC.

You can specify limits for specific protocols (ICMP, TCP, and UDP), or specify limits independent of a protocol.

When a session limit is exceeded for a source or destination, packets from the source or to the destination are dropped until the session limit is no longer exceeded.

To specify limits for destination or source subnets rather than individual addresses, include the `aggregation` statement at the `[edit services ids rule rule-name term term-name then]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

session-limit-per-username (Access Profile)

IN THIS SECTION

- Syntax | 16821
- Hierarchy Level | 16822
- Description | 16822
- Options | 16822
- Required Privilege Level | 16822
- Release Information | 16822

Syntax

```
session-limit-per-username number;
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Limit the number of active sessions allowed for a username associated with a given access profile. This enables you to control credential sharing by legitimate subscribers. After the limit is configured, existing subscriber sessions for the access profile are tracked as active sessions, as are any new sessions that do not exceed the limit. Requests for sessions that exceed the configured limit are blocked.

Use the `show network-access aaa statistics session-limit-per-username` command to view statistics for active sessions and blocked requests.

You can use the `clear network-access aaa statistics session-limit-per-username username` command as an aid to debugging by clearing the blocked request statistics for any of the following cases:

- For all usernames across all access profiles.
- For a specific username across all access profiles.
- For a specific username in a specific access profile.
- For all usernames in a specific access profile.

Options

number Maximum number of subscriber sessions allowed per username in the access profile.

- **Range:** 1 through 16

Required Privilege Level

access

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Session Options for Subscriber Access

Session Options for Subscriber Access

session-limit (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 16823
- Hierarchy Level | 16823
- Description | 16823
- Options | 16824
- Required Privilege Level | 16824
- Release Information | 16824

Syntax

```
session-limit {
    maximum number;
}
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Description

Specify the maximum number of sessions allowed simultaneously on the service set. If you specify the maximum number of sessions to be zero, it indicates that the configuration is not effective. You must specify a value higher than zero for the maximum number of sessions.

Options

number Maximum number of sessions.

- **Range:** 1 through 4,294,967,295

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

session-limit (System Services)

IN THIS SECTION

- Syntax | 16824
- Hierarchy Level | 16825
- Description | 16825
- Options | 16825
- Required Privilege Level | 16825
- Release Information | 16825

Syntax

```
session-limit;
```

Hierarchy Level

```
[edit system services web-management session]
```

Description

Configure the maximum allowed number of concurrent web management sessions. Using the CLI, you can limit the number of concurrent sessions from 1 to 1024. Each session remains open for 24 hours (1440 minutes).

You can also set an idle time-out parameter to override the default to specify the length of time a session can be idle before it is terminated.

Options

session-limit *number*

- **Default:** unlimited sessions
- **Range:** 1 to 1024 sessions

Required Privilege Level

system—To view this in the configuration.

system-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Dynamic VPN Overview

session-memory (resource-manager)

IN THIS SECTION

- Syntax | 16826
- Hierarchy Level | 16826
- Description | 16827
- Options | 16827
- Required Privilege Level | 16827
- Release Information | 16827

Syntax

```
session-memory {
    match {
        threshold threshold;
    }
    then {
        alarm;
        drop-flows;
        log;
        session-aggressive-ageout {
            early-ageout early-ageout;
        }
    }
}
```

Hierarchy Level

```
[edit security resource-manager]
```

Description

Configure resource manager session-memory options to control the behavior of the device based on the sessions that the device is handling.

Options

- `threshold` *`threshold`*-Configure the threshold limit to let `resource-manager` take actions. The range is 1 to 100 percent.
- `session-aggressive-ageout early-ageout` *`seconds`*-When you enable aggressive ageout, sessions with timeout less than the early value, become invalid. Range is 1 through 65,535 seconds.
- `drop-flows`-When you mark a session as a drop-flow, it bypasses the first path for each new packet, which matches the original drop session, for next 4 seconds only. This configuration drops the packets that match the `drop-flows`.
- `log`-This configuration logs messages when the resource threshold is exceeded and falls back under the threshold.
- `alarm`-This configuration sends an SNMP trap when the session usage exceeds the threshold.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

> *Traffic Processing on SRX Series Devices Overview*

session-mode

IN THIS SECTION

- Syntax | 16828
- Hierarchy Level | 16828
- Description | 16828
- Options | 16828
- Required Privilege Level | 16829
- Release Information | 16829

Syntax

```
session-mode (automatic | multihop | singlehop);
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay
dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group group-name liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd]
```

Description

Configure the session mode.

Options

- **Default:** automatic

automatic Configure single-hop BFD sessions if the peer is directly connected to the router interface and multihop BFD sessions if the peer is not directly connected to the router interface.

multihop Configure multihop BFD sessions and passive DHCP clients.

single-hop Configure single hop BFD sessions and non-passive DHCP clients.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients
- Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

session-offload

IN THIS SECTION

- Syntax | 16830
- Hierarchy Level | 16830
- Description | 16830
- Default | 16830
- Required Privilege Level | 16830
- Release Information | 16830

Syntax

```
session-offload;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number adaptive-services service-package extension-provider]
```

Description

Enable session offloading on a per-PIC basis for a Multiservices PIC.

Default

Session offloading is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced on MX Series 5G Universal Routing Platforms with MS-DPCs in Junos OS Release 9.6.

RELATED DOCUMENTATION

Enabling Session Offloading for Multiservices DPCs

session-options

IN THIS SECTION

- Syntax | 16831
- Hierarchy Level | 16832
- Description | 16832
- Options | 16832
- Required Privilege Level | 16834
- Release Information | 16834

Syntax

```
session-options {
    client-group [ group-names ];
    client-idle-timeout minutes;
    client-idle-timeout-ingress-only;
    client-session-timeoutminutes;
    pcc-context {
        input-service-filter-name filter-name;
        input-service-set-name service-set-name;
        ipv6-input-service-filter-name filter-name;
        ipv6-input-service-set-name service-set-name;
        ipv6-output-service-filter-name filter-name;
        ipv6-output-service-set-name service-set-name;
        output-service-filter-name filter-name;
        output-service-set-name service-set-name;
        profile-name pcef-profile-name;
    }
    strip-user-name {
        delimiter [ delimiter ];
        parse-direction (left-to-right | right-to-left);
    }
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

(MX Series and SRX Series Firewalls) Define options to place limits on subscriber access based on how long the session has been up, how long the user has been inactive, or both.

(MX Series) Define options to modify a subscriber username at login based on the subscriber's access profile.

(MX Series) Specify characteristics related to policy and charging control (PCC) rules, such as the PCEF profile that contains the rules, service sets to process the rules, and service filters for the service sets.

Options

client-idle-timeout

Specify the grace period that begins after an authenticated user terminates all sessions and connections. Authentication is not required if a new connection is initiated during the grace period by the same user.

During this period, the router determines whether the subscriber is inactive by monitoring data traffic, both upstream from the user (ingress) and downstream to the user (egress). Control traffic is ignored. The subscriber is not considered idle as long as data traffic is detected in either direction. When no traffic is detected for the duration of the idle time out, non-DHCP subscribers (such as L2TP or PPP) are gracefully logged out, similarly to a RADIUS-initiated disconnect or a CLI-initiated logout; DHCP subscribers are disconnected.

When you additionally configure the related `client-idle-timeout-ingress-only` statement (MX Series only), the router monitors only ingress traffic to determine whether the subscriber is inactive; it does not monitor any egress traffic. The related `client-session-timeout` statement terminates the subscriber session when the session timeout expires regardless of user activity.

Client idle timeouts are most often used for residential services rather than business services. The most practical use case for this timeout is in a PPP access model. It is not practical for DHCP or DHCPv6 subscribers.

Although you can use the `client-idle-timeout` statement for dynamically configured subscriber VLANs, this configuration is useful only in limited circumstances (such as IP over Ethernet without DHCP and with fixed addresses) and is not typically used. If you do use the idle timeout for VLANs, the timeout period starts when the VLAN is instantiated. It

resets when a client session is created or an existing session is reactivated. When no traffic is detected on an authenticated VLAN for the duration of the timeout, the VLAN is considered inactive and is deleted. If no client sessions are ever created on the VLAN, then the VLAN is removed when the timeout expires.

- **Default:** The timeout is not configured.
- **Values:** *minutes*—Number of minutes of idle time that elapse before the session is terminated. The value that you specify must be determined locally with consideration of the services and policies that you offer.
- **Range:** 10 through 1440 minutes

client-idle-timeout-ingress-only

Specify that only ingress traffic is monitored for subscriber idle timeout processing for the duration of the idle timeout period that you specify with the `client-idle-timeout` statement. If no ingress traffic is received for the duration of the timeout, then the subscriber is gracefully logged out (non-DHCP subscribers) or disconnected (DHCP subscribers).

If you configure `client-idle-timeout` alone, then both ingress and egress traffic are monitored during the idle timeout. Monitoring only ingress traffic is useful in cases where the LNS sends traffic to the remote peer even when the peer is not up, such as when the LNS does not have PPP keepalives enabled and therefore does not detect that the peer is not up. Because the LAC monitors both ingress and egress traffic by default, in this situation it receives the egress traffic from the LNS and either does not log out the subscriber or delays detection of inactivity until the egress traffic ceases. When you specify that only ingress traffic is monitored in this case, the LAC can detect that the peer is inactive and then initiate logout.

client-session-timeout

Specify the amount of time after which user sessions are terminated, regardless of user activity (also known as a forced or hard authentication timeout).

Alternatively, when you want subscribers to be identified as inactive before they are terminated, use the related statements, `client-idle-timeout` and `client-idle-timeout-ingress-only`. Use `client-idle-timeout` alone to specify a period of time during which both ingress and egress subscriber data traffic is monitored; if no traffic is detected for the duration of the period, the subscriber is considered inactive and is terminated. Add the `client-idle-timeout-ingress-only` statement to monitor only ingress traffic for the duration of the timeout set with the `client-idle-timeout statement`.

BEST PRACTICE: We recommend that you do not configure a session timeout for subscribers receiving voice services. Because the session timeout is a simple time-based timeout, it is likely to interrupt subscribers actively using a voice service and terminate their calls unexpectedly (from

the subscriber viewpoint). This result is a particular concern for emergency services calls.

Client session timeouts are most often used for residential services rather than business services. The most practical use case for this timeout is in a PPP access model when no voice services are offered. For DHCP or DHCPv6 subscribers, the session timeout is used as the DHCP lease timer if no other lease time configuration is present.

Although you can use the `client-session-timeout` statement for dynamically configured subscriber VLANs, this configuration is useful only in limited circumstances (such as IP over Ethernet without DHCP and with fixed addresses) and is not typically used. If you do use the session timeout for VLANs, the timeout period starts when the VLAN is instantiated.

- **Default:** The timeout is not configured.
- **Values:** *minutes*—Number of minutes after which user sessions are terminated. The value that you specify must be determined locally with consideration of the services and policies that you offer.
- **Range:** 1 through 527040 minutes

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Session Options for Subscriber Access

Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management

session-options (Access Profile)

IN THIS SECTION

- Syntax | 16835
- Hierarchy Level | 16835
- Description | 16835
- Options | 16835
- Required Privilege Level | 16836
- Release Information | 16836

Syntax

```
session-options {
    client-group [group-names];
    client-idle-timeout minutes;
    client-session-timeout minutes;
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Define options that control a user's session after successful authentication.

Options

The remaining statements are explained separately.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring RADIUS and LDAP User Authentication

session (Origin Validation for BGP)

IN THIS SECTION

- Syntax | 16836
- Hierarchy Level | 16837
- Description | 16837
- Options | 16837
- Required Privilege Level | 16839
- Release Information | 16839

Syntax

```
session server-ip-address {
    database <database-name>;
    hold-time seconds;
    local-address local-ip-address;
    port port-number;
    preference number;
    record-lifetime seconds;
```

```
    refresh-time seconds;
    traceoptions {
    file filename <files number>     <size size> <(world-readable | no-world-readable)>;
        flag flag {
            disable;
            flag-modifier;
        }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation group group-name],
[edit logical-systems logical-system-name routing-options validation group group-name],
[edit routing-instances instance-name routing-options validation group group-name],
[edit routing-options validation group group-name]
```

Description

Configure a TCP session with a resource public key infrastructure (RPKI) cache server. The router-to-cache transport protocol is carried using a TCP session to a configurable port. Caches are organized in groups. The Junos OS implementation supports up to 63 sessions per group and both IPv4 and IPv6 address families.

The maximum number of sessions in a group is two, by default, and is configurable. If the number of sessions in a group exceeds the `max-sessions` value, the connections are established in order by `preference` value. A numerically higher preference results in a higher probability for session establishment. The order of session establishment is random among sessions with equal preferences.

Options

server-ip-address	Specify the IP address of the RPKI cache server.
database *database-name*	Specify a named validation database.
hold-time *seconds*	Specify the length of time in seconds that the session between the routing device and the RPKI cache server is to be considered operational without any activity. After the hold time expires, the session is dropped.

Receiving any protocol data unit (PDU) from the cache server resets the hold timer. The hold time must be configured to be at least twice the value configured on the `refresh-time` statement. If the hold time expires, the session is considered to be down. This, in turn, triggers a session restart event. During a session restart, the routing device attempts to start a session with the cache server that has the numerically highest `preference` value.

- **Range:** 10 through 3600 seconds. We recommend that you configure the hold time to atleast twice the refresh time.
- **Default:** 600 seconds

local-address *local-ip-address*

Configure a local IP address of the session to be used for the outgoing connection to the RPKI cache server. If the local cache server has inbound firewall filtering, it might be necessary to specify a local IP address to use for this session.

port *port-number*

Configure an alternative TCP port number to be used for the outgoing connection to the cache server. The well-known resource public key infrastructure (RPKI) port is TCP port 2222. For a given deployment, an RPKI cache server might listen on some other TCP port number. If so, configure the alternative port number with this statement.

- **Default:** 2222

preference *number*

Configure the preference number for the RPKI cache server. Each cache server has a static preference. Higher preferences are preferred. During a session start or restart, the device attempts to start a session with the cache server that has the numerically highest preference. The device connects to multiple cache servers in preference order.

- **Range:** 1 through 255
- **Default:** 100

record-lifetime *seconds*

Configure the amount of time that route validation (RV) records learned from an RPKI cache server remain valid after the session to the cache server goes down. RV records expire if the session to the cache server goes down and remains down for the time configured.

- **Range:** 60 (one minute) through 604800 (one week). We recommend that you configure this record lifetime atleast twice the refresh time or greater. If you set this value to 60 seconds or lower, it might trigger very frequent refreshes causing higher CPU usage.
- **Default:** 3600 seconds (one hour)

refresh-time *seconds*

Configure a liveliness check interval for a configured resource public key infrastructure (RPKI) cache server. After every period of time configured on the `refresh-time` statement (in seconds), a serial query protocol data unit (PDU) with the last known serial number is transmitted. The value configured on the `refresh-time` statement cannot be longer than half of the value configured on the `hold-time` statement.

- **Range:** 1 through 1800 seconds
- **Default:** 300 seconds

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

session-persistence-scan

IN THIS SECTION

- Syntax | 16840
- Hierarchy Level | 16840
- Description | 16840
- Required Privilege Level | 16840

- Release Information | 16840

Syntax

```
session-persistence-scan;
```

Hierarchy Level

[edit security nat source]

Description

Specify the sessions to be retained if there is a change in NAT configuration. The existing sessions are retained, if the new session and existing sessions are in the same security zone.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

session-pics

IN THIS SECTION

- Syntax | 16841
- Hierarchy Level | 16841
- Description | 16841

Required Privilege Level | 16841
Release Information | 16841

Syntax

```
session-pics {
    [interface interface-name];
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name system]
```

Description

Specify the service interfaces that represent the session PICs used for the control plane in the TDF gateway.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Session PICs

TDF Gateway Service PICs and Session PICs for Subscriber-Aware Traffic Treatment

session-pics (Diameter)

IN THIS SECTION

- Syntax | 16842
- Hierarchy Level | 16842
- Description | 16842
- Options | 16843
- Required Privilege Level | 16843
- Release Information | 16843

Syntax

```
session-pics {
    group {
        group-name {
            [session-pic interface-name];
        }
    }
}
```

Hierarchy Level

```
[edit unified-edge tdf gateway-name diameter network-element element-name]
```

Description

Configure the session PICs that are serving this Diameter network element for Diameter bindings on this TDF gateway.

NOTE: If you want to set up Diameter bindings for session PICs on the TDF gateway, contact Juniper Networks Professional Services for assistance.

Options

group-name Name of the session PIC group that is serving the Diameter network element.

interface-name Name of interface representing session PIC.

- **Syntax:** The interface must be a valid multiservices interface (ams or ms-*a*/*b*/0, where *a* is the Flexible PIC Concentrator [FPC] slot number and *b* is the PIC slot number); for example, ams0, ams1, or ms-1/0/0.

> **NOTE**: The specified interface for the session PIC must already be configured for this TDF gateway.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (TDF Gateway) | 5702

session-protection

IN THIS SECTION

- Syntax | 16844
- Hierarchy Level | 16844
- Description | 16844
- Options | 16844

Required Privilege Level | 16844

Syntax

```
session-protection {
    timeout seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Configure when an LDP session is torn down and resignaled after the router stops receiving hello messages from a neighboring router. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished. The LDP session remains up for the duration specified as long as the routers maintain IP network connectivity.

Options

`timeout` *`seconds`*—Time in seconds before the LDP session is torn down and resignaled.

- **Range:** 1 through 65,535 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring LDP Session Protection

session (Protocols LDP)

IN THIS SECTION

- Syntax | 16845
- Hierarchy Level | 16845
- Description | 16846
- Options | 16846
- Required Privilege Level | 16846
- Release Information | 16847

Syntax

```
session session-address {
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key authentication-key;
    authentication-key-chain authentication-key-chain;
    downstream-on-demand downstream-on-demand;
    (mtu-discovery | no-mtu-discovery);
    transport-address transport-address;
}
```

Hierarchy Level

```
[edit logical-systems name protocols ldp],
[edit logical-systems name routing-instances name protocols ldp],
[edit protocols ldp],
[edit routing-instances name protocols ldp]
```

Description

Configure LDP session parameters by specifying the session destination address.

Options

session-address Session destination address.

authentication-algorithm Authentication algorithm name.

- Values:
  - `aes-128-cmac-96`—Cipher-based Message Authentication Code (AES128) (96 bits).
  - `hmac-sha-1-96`—Hash-based Message Authentication Code (SHA1) (96 bits).
  - `md5`—Message Digest 5.

authentication-key MD5 authentication key.

authentication-key-chain Key chain name.

downstream-on-demand Configure downstream on demand label distribution mode.

mtu-discovery Enable TCP path MTU discovery.

no-mtu-discovery Disable TCP path MTU discovery.

transport-address *transport-address* IP address used for TCP sessions to the LDP neighbors that have targeted-LDP adjacencies.

You cannot configure this statement for LDP sessions that have discovered adjacencies.

The transport-address configuration can be rejected at the time of validation, if there is no interface with the configured IP address.

Required Privilege Level

routing

Release Information

Statement introduced before Junos OS Release 7.4.

`authentication-algorithm` statement introduced in Junos OS Release 7.6.

`transport-address` option introduced in Junos OS Releasse 19.1R1 for all platforms.

RELATED DOCUMENTATION

Configuring the TCP MD5 Signature for LDP Sessions

session-scan

IN THIS SECTION

- Syntax | 16847
- Hierarchy Level | 16847
- Description | 16848
- Required Privilege Level | 16848
- Release Information | 16848

Syntax

```
session-scan;
```

Hierarchy Level

```
[edit security dynamic-address address-name name]
```

Description

Session scan provides an efficient method to traverse the session table to check and update each session. You can configure the `session-scan` option at two levels, entry level and module level. The entry level `session-scan` requests session scan only when the configured entries have the new addresses, while the module level takes all the entries belong to dynamic address into the scope. By default, the `session-scan` for dynamic address are disabled due to the session scan is CPU intensive.

The session scan checks all sessions source pool ID and ensures all session associated with modified pool can be removed. Session scan consumes more CPU cycle and is better not to trigger the session scan frequently. When one pool gets modified, session scan is not immediately triggered. When the modified pool number is greater than or equal to 1024 or becomes more than ten seconds, the session scan is triggered. The 10 second delay is designed to merge the session scan. For example, the first pool gets modified and after five seconds the second pool is modified, then after five seconds the session scan gets triggered, and thus for first pool and second pool the session scan gets merged together.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

hold-interval | 8089

session (Security IKE)

IN THIS SECTION

- Syntax | 16849
- Hierarchy Level | 16849
- Description | 16849
- Options | 16850

Required Privilege Level | 16852
Release Information | 16852

Syntax

```
session {
    full-open {
        incoming-exchange-max-rates {
            ike-rekey value;
            ipsec-rekey value;
            keepalive value;
        }
    }
    half-open {
        timeout seconds;
        backoff-timeouts {
           init-phase-failure value;
           auth-phase-failure value;
        }
        discard-duplicate;
        max-count value;
        thresholds {
            send-cookie count;
            reduce-timeout count timeout seconds;
        }
     }
}
```

Hierarchy Level

```
[edit security ike]
```

Description

Defines IKE session configuration.

In the configuration hierarchy, you can set the system level parameters of the IKE session to manage the behavior of the negotiations with the remote peer. You configure these settings to control high load scenarios.

Options

full-open

Define the full open IKE session parameters.

half-open

Define the half open IKE session parameters.

incoming-exchange-max-rates

Define the full open IKE session incoming exchange maximum rates.

You can use this parameter to set the maximum rates for various exchanges initiated by the remote peer after establishing an IKE SA.

ike-rekey *value*

Specify the value of the incoming peer initiated IKE rekey maximum rate.

The rate is applicable on a per peer basis.

- Format: Specify the value in [*count*]/[*duration in seconds*] format.
- Value: Range is [1 – 86400] / [1 – 86400].

ipsec-rekey *value*

Specify the value of the incoming peer initiated IPsec SA rekey maximum rate.

The rate is applicable on a per tunnel basis.

- Format: Specify the value in [*count*]/[*duration in seconds*] format.
- Value: Range is [1 – 86400] / [1 – 86400].

keepalive *value*

Specify the value of the incoming peer initiated keepalive, also knows as DPD, maximum rate.

The rate is applicable on a per peer basis.

- Format: Specify the value in [*count*]/[*duration in seconds*] format.
- Value: Range is [1 – 86400] / [1 – 86400].

backoff-timeouts

Define the half-open IKE session backoff timeouts.

You set these timeouts to allow some duration for the remote peer to back off in the event of a session initiation failure, ensuring that the same peer cannot initiate a new session initiation request immediately during that period. After the backoff timeout, the peer can initiate a new session. The scope is applicable at global level and not per peer level.

auth-phase-failure *value*

Specify the backoff timeout when there's a failure during the IKE_AUTH phase.

- Value: Range is 1 - 180 seconds.
- Default: Disabled

init-phase-failure *value*

Specify the backoff timeout when there's a failure during the SA_INIT phase.

- Value: Range is 1 - 180 seconds.
- Default: Disabled

discard-duplicate

Discard duplicate IKE session initiation requests from the peer.

Discard the IKE initiation requests without sending any response, when the half open IKE SA is already present for the same remote peer. The scope is applicable at global level and not per peer level.

- Default: Disabled

max-count *value*

Maximum numbers of half open IKE sessions where the local end is the responder.

The scope is applicable at global level and not per peer level.

- Value: Range is 1-10000
- Default: 300

thresholds

Define the half open IKE session thresholds.

You can set the limits on half open IKE SA count for actions against new a connection. The values denote percentage of the total half open IKE SAs. The scope is applicable at global level and not per peer level.

If you set the `max-count` explicitly, it automatically disables all thresholds unless you configure explicitly.

reduce-timeout *count* timeout *seconds*

Specify the minimum number of half open IKE sessions for enforcing reduce-timeout action at `reduce-timeout` *`count`*.

Set a limit from which you can reduce the lifetime of new half open IKE SAs.

- Value: Range is 1 - 10000.
- Default: Disabled

Specify the reduced timeout value at `timeout` *`seconds`*.

- Value: Range is 1 - 180 seconds.
- Default: Disabled

send-cookie *count*

Specify the minimum number of half open IKE sessions for enforcing cookie action.

Specify the threshold limit from which the responder requests remote peers to retry session initiation with a cookie sent back to the peer in the initial response.

- Value: Range is 1 - 10000.
- Default: 250.

timeout *seconds*

Specify the half open IKE session timeout.

This is the lifetime value of a half-open IKE SA that is applicable on the responder for the new sessions. For the existing sessions where there is no explicit configuration, the value is set to default. Initiator continues to use 60 seconds timeout value. The scope is applicable at global level and not per peer level.

- Value: Range is 1 - 180 seconds.
- Default: 60 seconds

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

`session` statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

IKE Protection from DDoS Attacks

blocklists (Security IKE) | **2997**

policy (Security IKE) | **12713**

ike (Security) | **8432**

session (System Services)

IN THIS SECTION

- Syntax | 16853
- Hierarchy Level | 16853
- Description | 16853
- Options | 16853
- Required Privilege Level | 16854
- Release Information | 16854

Syntax

```
session {
    idle-timout minutes;
    session-limit number;
}
```

Hierarchy Level

```
[edit system services web-management]
```

Description

Configure parameters for web-management sessions. Web management allows for configuration of the SRX Series Firewall using the J-Web interface. You can configure the idle-timeout parameter for web-management sessions and the maximum number of concurrent sessions.

Options

idle-timeout ***minutes***

Configure in minutes the time-out parameter for all web-management sessions. The idle-timeout parameter specifies the length of time that a session can be idle before it is terminated.

- **Default:** 1440
- **Range:** 1 to 1440

session-limit *number*

Configure the maximum allowed number of concurrent web management sessions. By default, an unlimited number of users can log in to the J-Web interface on a Juniper Networks device, and each session remains open for 24 hours (1440 minutes).

- **Default:** unlimited
- **Range:** 1 to 1024

Required Privilege Level

system—To view this in the configuration.

system-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Dynamic VPN Overview

session-timeout (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | **16855**
- Hierarchy Level | **16855**
- Description | **16855**
- Options | **16855**
- Required Privilege Level | **16855**

Release Information | 16855

Syntax

```
session-timeout seconds;
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Description

Define session lifetime for the service set in seconds. The session is closed after this amount of time, even if traffic is running on the session.

Options

seconds—Duration of session.

- **Range:** 4 through 86,400

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

session-update-interval

IN THIS SECTION

- Syntax | 16856
- Hierarchy Level | 16856
- Description | 16856
- Options | 16856
- Required Privilege Level | 16856
- Release Information | 16857

Syntax

```
session-update-interval session-update-interval;
```

Hierarchy Level

```
[edit security application-tracking]
```

Description

Configure the interval between session update messages for long-lived sessions being monitored by AppTrack. Byte count, packet count, and start and end times are updated and logged when the amount of time between session start or the previous update and the current time exceeds the interval.

Options

session-update-interval—Minutes between updates.

- **Default:** 5

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Example: Configuring AppTrack

session (Web Management)

IN THIS SECTION

- Syntax | 16857
- Hierarchy Level | 16857
- Description | 16858
- Options | 16858
- Required Privilege Level | 16858
- Release Information | 16858

Syntax

```
session {
    idle-timeout minutes;
    session-limit number of sessions;
}
```

Hierarchy Level

```
[edit system services web-management]
```

Description

Configure limits for the number of minutes a session can be idle before it times out, and configure the number of simultaneous J-Web user login sessions.

Options

idle-timeout *minutes*

Configure the number of minutes a session can be idle before it times out.

- **Range:** 1 through 1440 minutes
- **Default:** 30 minutes

session-limit *number of sessions*

Configure the maximum number of simultaneous J-Web user login sessions.

- **Range:** 1 through 1024 sessions
- **Default:** Unlimited

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

sessions-limit-group (L2TP)

IN THIS SECTION

- Syntax | 16859
- Hierarchy Level | 16859
- Description | 16859
- Options | 16859

Required Privilege Level | 16859
Release Information | 16859

Syntax

```
sessions-limit-group limit-group-name {
    maximum-sessions number;
}
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Create a group of clients so that you can limit the number of L2TP sessions allowed for the client group. You can create up to 200 groups.

Options

limit-group-name Identifier of the session-limit group for which the session limit is configured.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

L2TP LAC Subscriber Configuration

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

L2TP for Subscriber Access Overview

set

IN THIS SECTION

- Syntax | 16860
- Description | 16860
- Options | 16860
- Required Privilege Level | 16861
- Release Information | 16861

Syntax

```
set <statement-path> identifier
```

Description

Create a statement hierarchy and set identifier values. This is similar to `edit` except that your current level in the hierarchy does not change.

Options

identifier Name of the statement or identifier to set.

statement-path (Optional) Path to an existing statement hierarchy level. If that hierarchy level does not exist, it is created.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

edit | 6201

Displaying the Current Configuration

set chassis cluster cluster-id node node-number reboot

IN THIS SECTION

- Syntax | 16861
- Description | 16862
- Options | 16862
- Required Privilege Level | 16862
- Output Fields | 16862
- Release Information | 16863

Syntax

```
set chassis cluster cluster-id cluster-id node  node-number reboot
```

Description

Sets the chassis cluster identifier (ID) and node ID on each device, and reboots the devices to enable clustering. The system uses the chassis cluster ID and chassis cluster node ID to apply the correct configuration for each node (for example, when you use the `apply-groups` command to configure the chassis cluster management interface). The chassis cluster ID and node ID statements are written to the EPROM, and the statements take effect when the system is rebooted.

Setting a cluster ID to 0 is equivalent to disabling a cluster. A cluster ID greater than 15 can only be set when the fabric and control link interfaces are connected back-to-back.

> **NOTE**: If you have a cluster set up and running with an earlier release of Junos OS, you can upgrade to Junos OS Release 12.1X45-D10 or later and re-create a cluster with cluster IDs greater than 16. If for any reason you decide to revert to the previous version of Junos OS that did not support extended cluster IDs, the system comes up with standalone devices after you reboot. If the cluster ID set is less than 16 and you roll back to a previous release, the system comes back with the previous setup.

Options

cluster-id *cluster-id*

Identifies the cluster within the Layer 2 domain.

- **Range:** 0 through 255

node *node*

Identifies a node within a cluster.

- **Range:** 0 through 1

reboot

reboot the specified devices to configure the chassis cluster.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Release Information

Support for extended cluster identifiers (more than 15 identifiers) added in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

- Example: Setting the Chassis Cluster Node ID and Cluster ID
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Primary Administrators Only)
- Disabling a Chassis Cluster
- set chassis cluster disable reboot

set chassis cluster disable reboot

IN THIS SECTION

- Syntax | 16863
- Description | 16864
- Options | 16864
- Required Privilege Level | 16864
- Output Fields | 16864
- Sample Output | 16864
- Release Information | 16864

Syntax

```
set chassis cluster disable <reboot>
```

Description

After defining chassis cluster configuration, you can disable or remove a chassis cluster and change the SRX Series Firewalls to standalone devices by running the `set chassis cluster disable reboot` command from the operational mode.

After the chassis cluster is disabled using this CLI command, you do not have a similar CLI option to enable it.

Options

reboot You can disable the chassis cluster and run the reboot command. When the reboot is completed, you can view the SRX device functioning in standalone mode.

Required Privilege Level

maintenance

Output Fields

Sample Output

set chassis cluster disable reboot

```
user@host> set chassis cluster disable reboot
Successfully disabled chassis cluster. Going to reboot now.
```

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

set chassis cluster cluster-id node node-number reboot | **16861**

Disabling a Chassis Cluster

set chassis display message

IN THIS SECTION

- Syntax | 16865
- Syntax (TX Matrix Router) | 16865
- Syntax (TX Matrix Plus Router) | 16865
- Description | 16866
- Options | 16866
- Required Privilege Level | 16867
- Output Fields | 16867
- Sample Output | 16867
- Release Information | 16868

Syntax

```
set chassis display message "message"
<permanent>
```

Syntax (TX Matrix Router)

```
set chassis display message "message" (lcc number / scc)
 <permanent>
```

Syntax (TX Matrix Plus Router)

```
set chassis display message “message ” (fpc-slot slot-number | lcc number /  sfc number)
 <permanent>
```

Description

Display or stop a text message on the craft interface display, which is on the front of the router, or on the LCD panel display on the switch. The craft interface alternates the display of text messages with standard craft interface messages three times, switching between messages every 60 seconds.

> **NOTE**: On T Series routers, when this command is executed with the **permanent** option, the display of the text message alternates with that of the standard craft interface message continuously every 60 seconds.

By default, on both the router and the switch, the text message is displayed for 5 minutes. The craft interface display has four 20-character lines. The LCD panel display has two 16-character lines, and text messages appear only on the second line.

Options

"*message*"
Message to display. On the craft interface display, if the message is longer than 20 characters, it wraps onto the next line. If a word does not fit on one line, the entire word moves down to the next line. Any portion of the message that does not fit on the display is truncated. An empty pair of quotation marks (“ ”) deletes the text message from the craft interface display. On the LCD panel display, the message is limited to 16 characters.

fpc-slot *slot-number*
(TX Matrix Plus routers and EX4200 and QFX Series only) On the router or switch, display the text message on the craft interface for a specific Flexible PIC Concentrator (FPC). Replace *slot-number* with a value from **0** through **31**. On the switch, display the text message for a specific member of a Virtual Chassis, where **fpc-slot *slot-number*** corresponds to the member ID. Replace ***slot-number*** with a value from **0** through **9**. On the QFX Series, the *slot-number* is always **0**. On a TX Matrix Plus router with 3D SIBs replace *slot-number* with a value from `0` through `63`.

lcc *number*
(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

permanent (Optional) Display a text message on the craft interface display or LCD panel display permanently.

scc (TX Matrix routers only) Display the text message on the craft interface display of the TX Matrix router (switch-card chassis).

sfc *number* (TX Matrix Plus routers only) Display the text message on the craft interface display of the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level

clear

Output Fields

See show chassis craft-interface for an explanation of output fields.

Sample Output

set chassis display message (Creating)

The following example shows how to set the display message and verify the result:

```
user@host> set chassis display message "NOC
contact Dusty (888) 555-1234"
message sent
```

```
user@host> show chassis craft-interface
Red alarm:     LED off, relay off
Yellow alarm:  LED off, relay off
Host OK LED:   On
Host fail LED: Off
FPCs     0  1  2  3  4  5  6  7
-------------------------------
```

```
Green  ..  *..  *  *.
Red    ........
LCD screen:
     +--------------------+
     |NOC contact Dusty   |
     |(888) 555-1234      |
     +--------------------+
```

set chassis display message (Deleting)

The following example shows how to delete the display message and verify that the message is removed:

```
user@host> set chassis display message ""
message sent


user@host> show chassis craft-interface
Red alarm:     LED off, relay off
Yellow alarm:  LED off, relay off
Host OK LED:   On
Host fail LED: Off
FPCs     0  1  2  3  4  5  6  7
-------------------------------
Green  ..  *..  *  *.
Red    ........
LCD screen:
     +--------------------+
     |host                |
     |Up: 0+17:05:47      |
     |                    |
     |Temperature OK      |
     +--------------------+
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option for TX Matrix Plus router introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring the LCD Panel on EX Series Switches (CLI Procedure)

clear chassis display message

show chassis craft-interface

set chassis virtual-chassis

IN THIS SECTION

- Syntax | 16869
- Description | 16869
- Options | 16870
- Required Privilege Level | 16870
- Sample Output | 16870
- Release Information | 16870

Syntax

```
set chassis virtual-chassis
<disable>
<local>
```

Description

Set the EX8200 switch into Virtual Chassis mode.

An EX8200 switch cannot be a member switch in an EX8200 Virtual Chassis until the switch is in Virtual Chassis mode.

This command is not needed when configuring an EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, or a mixed EX4200 and EX4500 Virtual Chassis.

Options

none Set the internal Routing Engines into Virtual Chassis mode.

disable Disable Virtual Chassis mode for all internal REs on the switch.

local Set the Routing Engine that you are currently logged in to into Virtual Chassis mode.

Required Privilege Level

view

Sample Output

set chassis virtual-chassis

```
user@switch> set chassis virtual-chassis
This will reboot the RE(s)
Do you want to continue ? [yes,no] (no) yes
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

show virtual-chassis | 27578

Configuring an EX8200 Virtual Chassis (CLI Procedure)

set cli complete-on-space

IN THIS SECTION

Syntax | 16871

- Description | 16871
- Options | 16871
- Required Privilege Level | 16871
- Output Fields | 16871
- Sample Output | 16872
- Release Information | 16872

Syntax

```
set cli complete-on-space (off | on)
```

Description

Set the command-line interface (CLI) to complete a partial command entry when you type a space or a tab. This is the default behavior of the CLI.

Options

off Turn off command completion.

on Allow either a space or a tab to be used for command completion.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli complete-on-space

In the following example, pressing the Spacebar changes the partial command entry from **com** to **complete-on-space**. The example shows how adding the keyword **off** at the end of the command disables command completion.

```
user@host> set cli com<Space>
user@host>set cli complete-on-space off
Disabling complete-on-space
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

CLI User Interface Overview

show cli | 19157

set cli directory

IN THIS SECTION

- Syntax | 16873
- Description | 16873
- Options | 16873
- Required Privilege Level | 16873
- Output Fields | 16873
- Sample Output | 16873
- Release Information | 16873

Syntax

```
set cli directory directory
```

Description

Set the current working directory.

Options

directory	Pathname of the working directory.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli directory

```
user@host> set cli directory /var/tmp
Current directory: /var/tmp
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- CLI User Interface Overview
- show cli directory | 19164

set cli idle-timeout

IN THIS SECTION

- Syntax | 16874
- Description | 16874
- Options | 16874
- Required Privilege Level | 16874
- Output Fields | 16875
- Sample Output | 16875
- Release Information | 16875

Syntax

```
set cli idle-timeout <minutes>
```

Description

Set the maximum time that an individual session can be idle before the user is logged off the router or switch. `set cli idle-timeout` holds good only for the session in use when you enter it. If you need to configure the idle timeout permanently for all the CLI sessions, then configure the `idle-timeout` statement at the `[edit system login]` hierarchy level.

Options

minutes (Optional) Maximum idle time. The range of values, in minutes, is 0 through 100,000. If you do not issue this command, and the user's login class does not specify this value, the user is never forced off the system after extended idle times. Setting the value to 0 disables the timeout.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli idle-timeout

```
user@host> set cli idle-timeout 60
Idle timeout set to 60 minutes
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

CLI User Interface Overview

show cli | 19157

set cli logical-system

IN THIS SECTION

- Syntax | 16876
- Description | 16876
- Options | 16876
- Required Privilege Level | 16876
- Output Fields | 16876
- Sample Output | 16876
- Release Information | 16876

Syntax

```
set cli logical-system logical-system
```

Description

Set the CLI to the specified logical system view.

Options

logical-system	logical system name.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli logical-system

```
user@host> set cli logical-system log-router-A
logical system: log-router-A
user@host:log-router-A>
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear cli logical-system | 3481

Logical Systems User Guide for Routers and Switches

set cli prompt

IN THIS SECTION

- Syntax | 16877
- Description | 16877
- Options | 16877
- Required Privilege Level | 16877
- Output Fields | 16877
- Sample Output | 16878
- Release Information | 16878

Syntax

```
set cli prompt string
```

Description

Set the prompt so that it is displayed within the CLI.

Options

string CLI prompt string. To include spaces in the prompt, enclose the string in quotation marks. By default, the string is *username@hostname*.

Required Privilege Level

view

Output Fields

When you enter this command, the new CLI prompt is displayed.

Sample Output

set cli prompt

```
user@host> set cli prompt lab1-router>
lab1-router>
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

CLI User Interface Overview

show cli | 19157

set cli restart-on-upgrade

IN THIS SECTION

- Syntax | 16878
- Description | 16879
- Options | 16879
- Required Privilege Level | 16879
- Output Fields | 16879
- Sample Output | 16879
- Release Information | 16879

Syntax

```
set cli restart-on-upgrade string (off | on)
```

Description

For an individual session, set the CLI to prompt you to restart the router or switch after upgrading the software.

Options

off	Disables the prompt.
on	Enables the prompt.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli restart-on-upgrade

```
user@host> set cli restart-on-upgrade on
Enabling restart-on-upgrade
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

CLI User Interface Overview

show cli | **19157**

set cli screen-length

IN THIS SECTION

- Syntax | 16880
- Description | 16880
- Options | 16880
- Required Privilege Level | 16880
- Release Information | 16881

Syntax

```
set cli screen-length length
```

Description

Set terminal screen length.

```
user@host> set cli screen-length 75
Screen length set to 75
```

Options

length—Number of lines of text that the terminal screen displays. The range of values, in an integer number of lines, 0 through 10,000 (except 1 in range). The value 0 sets no screen length. It also sets the `restart-on-upgrade` command off. The default value is 24.

The point at which the `---(more)---` prompt appears on the screen is a function of this setting and the settings for the `set cli screen-width` and `set cli terminal` commands.

Required Privilege Level

view

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Setting the Junos OS CLI Screen Length and Width

set cli screen-width | **16881**

set cli terminal | **16882**

show cli | **19157**

set cli screen-width

IN THIS SECTION

- Syntax | **16881**
- Description | **16881**
- Options | **16882**
- Required Privilege Level | **16882**
- Release Information | **16882**

Syntax

```
set cli screen-width width
```

Description

Set the terminal screen width.

```
user@host> set cli screen-width
Screen width set to 132
```

Options

width—Number of characters in a line. The value is `0` or in the range of `40` through `1024`. The default value is `80`.

> **NOTE**: In Junos OS Release 13.2 and earlier, the value of *width* is in the range of `0` through `1024`.

Required Privilege Level

view

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Setting the Junos OS CLI Screen Length and Width

set cli screen-length | **16880**

set cli terminal | **16882**

show cli | **19157**

set cli terminal

IN THIS SECTION

- Syntax | **16883**
- Description | **16883**
- Options | **16883**
- Required Privilege Level | **16883**
- Release Information | **16883**

Syntax

```
set cli terminal terminal-type
```

Description

Set the terminal type.

```
user@host> set cli terminal xterm
```

Options

terminal-type—Type of terminal that is connected to the Ethernet management port:

- `ansi`—ANSI-compatible terminal
- `pc`—PC screen command-prompt window
- `small-xterm`—Small xterm window (24 lines long)
- `vt100`—VT100-compatible terminal
- `xterm`—Large xterm window (65 lines long)

Required Privilege Level

view

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

CLI Environment Settings

set cli timestamp

IN THIS SECTION

- Syntax | 16884
- Description | 16884
- Options | 16884
- Required Privilege Level | 16885
- Output Fields | 16885
- Sample Output | 16885
- Release Information | 16885

Syntax

```
set cli timestamp (format timestamp-format  | disable)
```

Description

Set a timestamp for CLI output.

Options

format *timestamp-format*

Set the date and time format for the timestamp. The timestamp format you specify can include the following placeholders in any order:

- `%m`—Two-digit month
- `%d`—Two-digit date
- `%T`—Six-digit hour, minute, and seconds

disable

Remove the timestamp from the CLI.

NOTE: A timestamp is displayed by default when no command output is generated.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli timestamp

```
user@host> set cli timestamp format '%m-%d-%T'
'04-21-17:39:13'
CLI timestamp set to: '%m-%d-%T'
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

CLI User Interface Overview

show cli | 19157

set date

IN THIS SECTION

- Syntax | 16886
- Description | 16886
- Options | 16886
- Required Privilege Level | 16887

Sample Output | 16887
Release Information | 16887

Syntax

```
set date (date-time | ntp <ntp-server> <key authentication-key number> <servers> <source-address
source-address>)
```

Description

Set the date and time.

```
user@host> set date ntp
21 Apr 17:22:02 ntpdate[3867]: step time server 172.17.27.46 offset 8.759252 sec
```

Options

- *date-time*—Specify date and time in one of the following formats:
  - *YYYYMMDDHHMM.SS*
  - "*month DD, YYYY HH:MM*(am | pm)"

  NOTE: In Junos OS Evolved, the date and time will be set only on the node where the command is issued.

- ntp—Configure the router to synchronize the current date and time setting with a Network Time Protocol (NTP) server.

  NOTE: In Junos OS Evolved, if the ntpd server is running, the `set date ntp` command fails with the following error message: `error: ntpd is already running`. To use this command, you must first stop the ntpd server

- *ntp-server*—(Optional) Specify the IP address of one or more NTP servers.

- `key` *`authentication-key number`*—(Optional) Specify a key number to authenticate the NTP server used to synchronize the date and time. You must specify the same key number used to authenticate the server configured at the [`edit system ntp authentication-key` *`number`*] hierarchy level.
- `servers`—(Optional) Specify the IP address of one or more NTP servers.
- `source-address` *`source-address`*—(Optional) Specify the source address that is used by the router to contact the remote NTP server.

Required Privilege Level

view

Sample Output

set date ntp (Junos OS)

```
user@host> set date ntp
22 Jun 10:07:48 ntpdate[51123]: step time server 66.129.255.62 offset -0.013200 sec
```

set date ntp (Junos OS Evolved)

```
user@host> set date ntp
------------------------------
node: re0
------------------------------
error: ntpd is already running
```

Release Information

Command introduced before Junos OS Release 7.4.

`key` option introduced in Junos OS Release 12.1R2

set date ntp

IN THIS SECTION

- Syntax | 16888
- Description | 16888
- Options | 16888
- Required Privilege Level | 16889
- Output Fields | 16889
- Sample Output | 16889
- Release Information | 16889

Syntax

```
set date ntp <servers | force | key key | node node-number | routing-instance routing-instance-name | source-address source-address>
```

Description

Set the date and local time. If reject mode is enabled and the system rejected the update from the NTP server because it exceeds the configured threshold value, an administrator has two options to overrule the reject mode action: manually set the date and time in *YYYYMMDDhhmm.ss* format, or force synchronization of device time with the NTP server update by specifying the **force** option.

Options

servers	Specify the IP address of one or more NTP servers.
force	Force system date and time to update to NTP server values. The device date and time are synchronized with the NTP proposed date and time even if reject is set as the action and the difference between the device time and NTP proposed time exceeds the default or the configured threshold value.

key *key*	Specify a key number to authenticate the NTP server used to synchronize the date and time. You must specify the same key number used to authenticate the server, configured at the [`edit system ntp authentication-key` *number*] hierarchy level.
node *node-name*	Specify system date and time using NTP servers on specific node.
routing-instance *routing-instance-name*	Specify the routing instance through which server is available.
source-address *source-address*	Specify the source address that the SRX Series devices use to contact the remote NTP server.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

set date ntp force

```
user@host> set date ntp force
18 Jul 16:52:43 ntpdate[3319]: NTP update request has been accepted, The time offset is
147605840.624994 sec from the time server 66.129.255.62 which is larger than the maximum
threshold of 400 sec allowed.
```

Release Information

Command introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

show system ntp threshold | 26941

ntp | 11721

ntp threshold | 11731

set-dont-fragment-bit (Services IPsec VPN)

IN THIS SECTION

- Syntax | 16890
- Hierarchy Level | 16890
- Description | 16890
- Required Privilege Level | 16891
- Release Information | 16891

Syntax

```
set-dont-fragment-bit;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Configure the do not fragment (DF) bit in only the outer header of the IPsec packet and leave the inner header unmodified. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. These settings apply for static endpoint tunnels and not for dynamic tunnels, for which you need to include the `set-dont-fragment-bit` statement at the `[edit services service-set service-set-name ipsec-vpn-options]` hierarchy level to set the DF bit in the outer header of the IPv4 packets that enter the dynamic IPsec tunnel. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs.

By default, this statement is disabled on MS-MICs and MS-MPCs (the DF bit value is not configured in the outer header by default).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring IPsec Rules

set-dont-fragment-bit (Services Set)

IN THIS SECTION

- Syntax | 16891
- Hierarchy Level | 16891
- Description | 16892
- Required Privilege Level | 16892
- Release Information | 16892

Syntax

```
set-dont-fragment-bit;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Configure the do not fragment (DF) bit in only the outer header of the IPsec packet and leave the inner header unmodified for dynamic endpoint tunnels. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. These settings apply for dynamic endpoint tunnels and not for static tunnels, for which you need to include the `set-dont-fragment-bit` statement at the `[edit services ipsec-vpn rule` *`rule-name`* `term` *`term-name`* `then]` hierarchy level to set the DF bit in the outer header of the IPv4 packets that enter the static IPsec tunnel. This functionality is supported on MX Series routers with MS-MICs and MS-MPCs.

By default, this statement is disabled on MS-MICs and MS-MPCs (the DF bit value is not configured in the outer header by default).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring IPsec Service Sets

Configuring IPsec Rules

set services jinsightd subscribe health-monitor

IN THIS SECTION

- Syntax | **16893**
- Description | **16893**
- Options | **16893**
- Required Privilege Level | **16893**

Release Information | 16893

Syntax

```
set services jinsightd subscribe health-monitor
```

Description

Enables the J-Insight health monitor. Starting in Junos OS Release 18.2R1, J-Insight provides health monitoring capabilities for FPC FRUs on the MX series routers. The health monitor is disabled by default.

Options

This command has no options.

Required Privilege Level

system

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

J-Insight Device Monitor Basic Configuration

delete services jinsightd subscribe health-monitor | 5284

set services jinsightd traceoptions

IN THIS SECTION

- Syntax | 16894
- Description | 16894
- Options | 16894
- Required Privilege Level | 16895
- Release Information | 16895

Syntax

```
set services jinsightd traceoptions flag
<all>
<core>
<database>
<rule-engine>
<timer>
```

Description

Define tracing operations that track J-Insight functionality. To specify more than one tracing operation, include multiple `flag` statements.

Options

all	All tracing operations.
core	J-Insight core events.
database	Database events.
rule-engine	Rule engine events.
timer	Timer events.

Required Privilege Level

system

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

J-Insight Device Monitor Basic Configuration

identity-provider settings (SAML)

IN THIS SECTION

- Syntax | 16895
- Hierarchy Level | 16896
- Description | 16896
- Options | 16896
- Required Privilege Level | 16896
- Release Information | 16896

Syntax

```
settings {
    entity-id entity-id;
    idp-certificate idp-certificate;
    single-logout-url single-logout-url;
    single-signon-url single-signon-url;
}
```

Hierarchy Level

```
[edit access saml identity-provider name]
```

Description

Describe Security Assertion Markup Language (SAML) identity provider (IdP) settings. If you import IdP metadata, you do not need to manually configure IdP settings.

Options

entity-id ***entity-id***	Specify IdP entity id. Example: http://www.example.com/1123zweqios.
idp-certificate ***idp-certificate***	Specify the IdP's certificate.
single-logout-url ***single-logout-url***	Specify the IdP's single logout service URL. Example: http://dev-6799012.example.com/app/dev-6799012_acc1/1123zweqios/sso/saml.
single-signon-url ***single-signon-url***	Specify the IdP's single sign-on (SSO) service URL. Example: http://dev-6799012.example.com/app/dev-6799012_acc1/1123zweqios/slo/saml.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

saml | 16281

setup-protection

IN THIS SECTION

- Syntax | 16897
- Hierarchy Level | 16897
- Description | 16897
- Required Privilege Level | 16897

Syntax

```
setup-protection;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

The facility-backup fast reroute mechanism can provide setup protection for LSPs which are in the process of being signaled. Both point-to-point LSPs and point-to-multipoint LSPs are supported. You should configure the `setup-protection` statement on each of the routers along the LSP path on which you want to enable LSP setup protection. You should also configure IGP traffic engineering on all of the routers on the LSP path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring RSVP Setup Protection

severity (Event Policy)

IN THIS SECTION

- Syntax | 16898
- Hierarchy Level | 16898
- Description | 16898
- Options | 16898
- Required Privilege Level | 16899
- Release Information | 16899

Syntax

```
severity severity-level;
```

Hierarchy Level

```
[edit event-options policy policy-name then priority-override]
```

Description

Configure an event policy to override the preassigned severity level of a triggering event so that the event is logged based on the configured severity level.

Options

severity-level Severity level logged for the triggering event. Table 141 on page 16899 lists the possible severity levels.

Table 141: System Log Message Severity Levels

Severity Level	Description
emergency	System panic or other conditions that cause the routing platform to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard disk errors
error	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
warning	Conditions that warrant monitoring
notice	Conditions that are not errors but might warrant special handling
info	Events or non-error conditions of interest

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding the Event System Log Priority in an Event Policy

Example: Configure the Event System Log Priority in an Event Policy

Junos OS System Logging Facilities and Message Severity Levels

facility | **6769**

priority-override (Event Policy) | **13183**

severity (Next Gen Services Service-Set Remote System Logging)

IN THIS SECTION

- Syntax | 16900
- Hierarchy Level | 16900
- Description | 16900
- Required Privilege Level | 16901
- Release Information | 16901

Syntax

```
severity severity;
```

Hierarchy Level

```
edit services service-set name syslog stream stream-name
```

Description

Specify the level of severity for the stream.

You can set the following severity levels:

- ANY — Includes all severity levels
- ALERT — Action must be taken immediately
- CRITICAL — Critical conditions
- EMERGENCY — System is unusable
- ERROR — Error conditions
- WARNING — Warning conditions
- NOTICE — Normal but significant condition

- INFO — Informational
- DEBUG — Debug-level messages

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

stream (Next Gen Services Service-Set Remote System Logging) | **28631**

severity (Security IDP Custom Attack)

IN THIS SECTION

- Syntax | **16902**
- Hierarchy Level | **16902**
- Description | **16902**
- Options | **16902**
- Required Privilege Level | **16902**
- Release Information | **16903**

Syntax

```
severity (critical | info | major | minor | warning);
```

Hierarchy Level

```
[edit security idp custom-attack attack-name]
```

Description

Select the severity that matches the lethality of the attack object on your network.

Options

You can set the severity level to the following levels:

- `critical`—Contains attack objects matching exploits that attempt to evade detection, cause a network device to crash, or gain system-level privileges.
- `warning`—Contains attack objects matching exploits that attempt to obtain noncritical information or scan a network with a scanning tool.
- `major`—Contains attack objects matching exploits that attempt to disrupt a service, gain user-level access to a network device, or activate a Trojan horse previously loaded on a device.
- `minor`—Contains attack objects matching exploits that detect reconnaissance efforts attempting to access vital information through directory traversal or information leaks.
- `info`—Contains attack objects matching normal, harmless traffic containing URLs, DNS lookup failures, SNMP public community strings, and Peer-to-Peer (P2P) parameters. You can use informational attack objects to obtain information about your network.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

severity (Security IDP Dynamic Attack Group)

IN THIS SECTION

- Syntax | 16903
- Hierarchy Level | 16903
- Description | 16903
- Options | 16903
- Required Privilege Level | 16904
- Release Information | 16904

Syntax

```
severity {
    values [critical info major minor warning];
}
```

Hierarchy Level

```
[edit security idp dynamic-attack-group dynamic-attack-group-name filters]
```

Description

Specify a severity filter to add attack objects based on the attack severity levels.

Options

`values`—Name of the severity filter. You can select from the following severity:

- `critical`—The attack is a critical one.
- `info`—Provide information of attack when it matches.
- `major`—The attack is a major one.
- `minor`—The attack is a minor one.
- `warning`—Issue a warning when attack matches.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

severity (Security IDP IPS Rulebase)

IN THIS SECTION

- Syntax | 16904
- Hierarchy Level | 16905
- Description | 16905
- Options | 16905
- Required Privilege Level | 16905
- Release Information | 16905

Syntax

```
severity (critical | info | major | minor | warning);
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name then]
```

Description

Set the rule severity levels in logging to support better organization and presentation of log records on the log server. You can use the default severity settings of the selected attack object, or choose a specific severity for your rule. The severity you configure in the rules overrides the inherited attack severity.

Options

You can set the severity level to the following levels:

- `critical`—2
- `info`—3
- `major`—4
- `minor`—5
- `warning`—7

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

severity (Security Log)

IN THIS SECTION

- Syntax | 16906
- Hierarchy Level | 16906
- Description | 16906
- Options | 16906
- Required Privilege Level | 16907
- Release Information | 16907

Syntax

```
severity (alert | critical | debug | emergency | error | info | notice | warning)
```

Hierarchy Level

```
[edit security log stream stream-name]
[edit logical-systems name security log stream stream-name]
[edit tenants tenant-name security log stream stream-name]
```

Description

Set severity threshold for security logs.

Options

- `alert`— Conditions that require immediate attention.
- `critical`—Critical conditions.
- `debug`—Information normally used in debugging.
- `emergency`—Conditions that cause security functions to stop.

- `error`—General error conditions.
- `info`—Information about normal security operations.
- `notice`—Nonerror conditions that are of interest.
- `warning`—General warning conditions.
- **Default:** debug.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

The [edit `logical-systems` *name* security log stream *stream-name*] and [edit `tenants` *tenant-name* security log stream *stream-name*] hierarchy levels introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

syslog (System) | **28836**

sflow

IN THIS SECTION

- Syntax (EX Series) | **16908**
- Syntax (MX Series and EX92xx) | **16909**
- Syntax (PTX Series, QFX Series, OCX Series, ACX Series, and EX4600) | **16910**
- Hierarchy Level | **16911**
- Description | **16911**
- Default | **16911**

- Required Privilege Level | **16911**
- Release Information | **16911**

Syntax (EX Series)

```
sflow {
    adaptive-sample-rate rate {
        fallback;
        sample-limit-threshold sample-limit-threshold;
    }
    agent-id {
        ipv4-address;
        inet6 ipv6-address;
    }
    collector ip-v4/v6-address {
        udp-port port-number;
    }
    ;
    interfaces interface-name {
        polling-interval seconds;
        sample-rate {
            egress rate;
            ingress rate;
        }
    }
    polling-interval seconds;
    sample-rate {
        egress rate;
        ingress rate;
    }
    source-ip {
        ipv4-address;
        inet6 ipv6-address;
    }
    traceoptions {
        file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
        flag flag;
```

```
    }
}
```

Syntax (MX Series and EX92xx)

```
sflow {
    adaptive-sample-rate rate {
        fallback;
        sample-limit-threshold sample-limit-threshold;
    }
    agent-id {
        ipv4-address;
        inet6 ipv6-address;
    }
    collector ip-v4/v6-address {
        dscp dscp;
        forwarding-class forwarding-class;
        udp-port port-number;
    }
    interfaces interface-name {
        polling-interval seconds;
        sample-rate {
            egress rate;
            ingress rate;
        }
    }
    polling-interval seconds;
    sample-rate {
        egress rate;
        ingress rate;
    }
    source-ip {
        ipv4-address;
        inet6 ipv6-address;
    }
    traceoptions {
        file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
        flag flag;
```

```
    }
}
```

Syntax (PTX Series, QFX Series, OCX Series, ACX Series, and EX4600)

```
sflow {
    adaptive-sample-rate rate {
        fallback;
        sample-limit-threshold sample-limit-threshold;
    }
    agent-id {
        ipv4-address;
        inet6 ipv6-address;
    }
    collector ip-v4/v6-address {
        dscp dscp;
        forwarding-class forwarding-class;
        udp-port port-number;
    }
    disable-sw-rate-limiter;

    interfaces interface-name {
        polling-interval seconds;
        sample-rate {
            egress number;
            ingress number;
        }
    }
    polling-interval seconds;
    sample-rate {
        egress number;
        ingress number;
    }
    sample-size Sample-Size;
    source-ip {
        ipv4-address;
        inet6 ipv6-address;
    }
    traceoptions {
        file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
```

```
        flag flag;
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure sFlow technology to continuously monitor traffic at wire speed on specified interfaces simultaneously. sFlow data can be used to provide network traffic visibility information. You can configure sample size of the raw packets in the sFlow record. The range varies from 128 bytes to 512 bytes.

Default

sFlow technology is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

`dscp` and `forwarding-class` options introduced in Junos OS Release 19.1R1, for PTX routers and QFX switches (previously introduced for MX Series).

`sample-size` option introduced in Junos OS Evolved Release 23.1R1, for PTX routers.

RELATED DOCUMENTATION

Understanding How to Use sFlow Technology for Network Monitoring

Understanding Inline Active Flow Monitoring

sfm (Chassis)

IN THIS SECTION

- Syntax | 16912
- Hierarchy Level | 16912
- Description | 16912
- Options | 16912
- Required Privilege Level | 16913
- Release Information | 16913

Syntax

```
sfm slot-number power off;
```

Hierarchy Level

```
[edit chassis]
```

Description

For routers with SFMs, configure an SFM to stay offline.

By default, if you use the **request chassis sfm** CLI command to take an SFM offline, the SFM will attempt to restart when you enter a **commit** CLI command. To prevent a restart, configure an SFM to stay offline. This feature is useful for repair situations. The SFM remains offline until you delete this statement.

Options

slot-number—Slot number in which the SFM is installed.

power off—Take the SFM offline and configure it to remain offline.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring an SFM to Stay Offline

High Availability User Guide

sfm (Chassis Redundancy)

IN THIS SECTION

- Syntax | 16913
- Hierarchy Level | 16914
- Description | 16914
- Default | 16914
- Options | 16914
- Required Privilege Level | 16914
- Release Information | 16914

Syntax

```
sfm slot-number (always | preferred);
```

Hierarchy Level

```
[edit chassis redundancy]
```

Description

On M40e and M160 routers, configure which Switching and Forwarding Module (SFM) is the primary and which is the backup.

Default

By default, the SFM in slot 0 is the primary and the SFM in slot 1 is the backup.

Options

slot-number—Specify which slot is the primary and which is the backup. On the M40e router, ***slot-number*** can be 0 or 1. On the M160 router, ***slot-number*** can be 0 through 3.

always—Define this SFM as the sole device.

preferred—Define this SFM as the preferred device of at least two.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring SFM Redundancy on M40e and M160 Routers*

sfpplus

IN THIS SECTION

- Syntax | 16915
- Hierarchy Level | 16915
- Description | 16915
- Default | 16915
- Required Privilege Level | 16916
- Release Information | 16916

Syntax

```
sfpplus {
    pic-mode mode;
}
```

Hierarchy Level

```
[edit chassis fpc slot pic pic-number]
```

Description

Configure the operating mode for the specified port on the SFP+ uplink module on the EX3200 or EX4200 switch.

The remaining statement is explained separately. See CLI Explorer.

Default

By default, the SFP+ uplink module operates in the 10-gigabit mode and supports SFP+ transceivers.

> **NOTE**: The SFP+ uplink module provides two ports for 10-gigabit small form-factor pluggable (SFP+) transceivers when configured to operate in 10-gigabit mode or four ports for 1-gigabit small form-factor pluggable (SFP) transceivers when configured to operate in 1-gigabit mode.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module

sftp-backup (OCS Partition)

IN THIS SECTION

- Syntax | 16917
- Hierarchy Level | 16917
- Description | 16917
- Options | 16917
- Required Privilege Level | 16918
- Release Information | 16918

Syntax

```
sftp-backup {
    accumulation-timeout seconds;
    accumulation-count messages;
    accumulation-size bytes;
    logical-system;
    retry-interval seconds;
    response-timeout seconds;
    routing-instance;
    address ipv4 | ipv6;
    port number;
    directory filename;
    file-name-prefix filename-prefix;
    node-ipv4-address;
    node-ipv6-address;
    ssh-login;
    ssh-connection-linger;
    ssh-log-verbose;
    ssh-passphrase;
}
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

The data backup for OCS using the SFTP protocol to store data on a remote server.

Options

accumulation-timeout *seconds*

The file accumulation time since the first CCR-GY-T was submitted. The maximum time is 30 seconds.

accumulation-count *messages*

The number of requests that are past file account count. The maximum accumulation count is 2000 messages.

accumulation-size *bytes*	The file once its size has reached the accumulation size limit. The maximum accumulation size is 3000000 bytes.
logical-system	The logical system to be used by sftp.
retry-interval *seconds*	Every failed write operation is retried after this interval until backup timeout is accumulated. The maximum retry interval is 1200 seconds.
response-timeout *seconds*	The timeout on individual SFTP command response. The maximum response timeout is 60 seconds.
routing-instance	The SFTP service connection context for the routing instance.
address	The address of SFTP server
port	The port number. The default port number is 22.
directory	The directory name.
file-prefix *prefix*	The prefix for the file.
node-ipv4-address	The IPv4 address of the sending node.
node-ipv6-address	The IPv6 address of the sending node.
ssh-login	The login name on SFTP server, may include %h and/or %r similar to directory.
ssh-connection-linger	The time to keep existing ssh connection idle state.
ssh-log-verbose	The verbose logs to be collected from ssh session.
ssh-passphrase	The password or later passphrase.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1.

RELATED DOCUMENTATION

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

sglimit

IN THIS SECTION

- Syntax | **16919**
- Hierarchy Level | **16920**
- Description | **16920**
- Options | **16920**
- Required Privilege Level | **16920**
- Release Information | **16920**

Syntax

```
sglimit {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
    log-interval seconds;
    maximum limit;
    threshold value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim ],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Configure a limit for the number of accepted (*,G) and (S,G) PIM join states.

> **NOTE**: The maximum limit settings that you configure with the `maximum` and the `family (inet | inet6) maximum` statements are mutually exclusive. For example, if you configure a global maximum PIM join state limit, you cannot configure a limit at the family level for IPv4 or IPv6 joins. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.

Options

`family (inet | inet6)`—(Optional) Specify either IPv4 or IPv6 join states to be counted towards the configured join state limit.

- **Default:** Both IPv4 and IPv6 join states are counted towards the configured join state limit.

The remaining statements are described separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring PIM State Limits

clear pim join

sham-link

IN THIS SECTION

- Syntax | 16921
- Hierarchy Level | 16921
- Description | 16921
- Options | 16922
- Required Privilege Level | 16922
- Release Information | 16922

Syntax

```
sham-link {
    local address
    {
        keychain keychain-name;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf],
[edit routing-instances routing-instance-name protocols ospf]
```

Description

Configure the local endpoint of a sham link.

You can create an intra-area link or sham link between two provider edge (PE) routing devices so that the VPN backbone is preferred over the back-door link. A back-door link is a backup link that connects customer edge (CE) devices in case the VPN backbone is unavailable. When such a backup link is available and the CE devices are in the same OSPF area, the default behavior is to prefer this backup link over the VPN backbone. This is because the backup link is considered an intra-area link, while the VPN backbone is always considered an inter-area link. Intra-area links are always preferred over inter-area links.

The sham link is an unnumbered point-to-point intra-area link between PE devices. When the VPN backbone has a sham intra-area link, this sham link can be preferred over the backup link if the sham link has a lower OSPF metric than the backup link.

The sham link is advertised using Type 1 link-state advertisements (LSAs). Sham links are valid only for routing instances and OSPFv2.

Each sham link is identified by the combination of a local endpoint address and a remote endpoint address.

Options

`local` *`address`*—The address for the local endpoint of the sham link.

`keychain` *`keychain-name`*—For each keychain, you must specify a name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration

Release Information

Statement introduced before Junos OS Release 7.4.

The option `keychain` *`keychain-name`* introduced in Junos OS and Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

Example: Configuring OSPFv2 Sham Links

sham-link-remote | **16923**

sham-link-remote

IN THIS SECTION

- Syntax | 16923
- Hierarchy Level | 16923
- Description | 16923
- Options | 16924
- Required Privilege Level | 16924
- Release Information | 16924

Syntax

```
sham-link-remote address {
    demand-circuit;
    ipsec-sa name;
    metric metric;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area area-id],
[edit routing-instances routing-instance-name protocols ospf area area-id]
```

Description

Configure the remote endpoint of a sham link.

You can create an intra-area link or sham link between two provider edge (PE) routing devices so that the VPN backbone is preferred over the back-door link. A back-door link is a backup link that connects customer edge (CE) devices in case the VPN backbone is unavailable. When such a backup link is available and the CE devices are in the same OSPF area, the default behavior is to prefer this backup link over the VPN backbone. This is ecause the backup link is considered an intra-area link, while the VPN

backbone is always considered an inter-area link. Intra-area links are always preferred over inter-area links.

The sham link is an unnumbered point-to-point intra-area link between PE devices. When the VPN backbone has a sham intra-area link, this sham link can be preferred over the backup link if the sham link has a lower OSPF metric than the backup link.

The sham link is advertised using Type 1 link-state advertisements (LSAs). Sham links are valid only for routing instances and OSPFv2.

Each sham link is identified by the combination of a local endpoint address and a remote endpoint address.

Options

address	Address for the remote end point of the sham link.
demand-circuit	Configure an interface as a demand circuit.
ipsec-sa ***name***	Apply the named IPsec authentication to the OSPF interface or virtual link or to an OSPFv2 remote sham link.
metric ***metric***	Specify the cost of an OSPF interface. The cost is a routing metric that is used in the link-state calculation. To set the cost of routes exported into OSPF, configure the appropriate routing policy. Range is 1 through 65,535. By default, the cost of an OSPF route is calculated by dividing the reference-bandwidth value by the bandwidth of the physical interface. Any specific value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the cost of the route for that interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for `ipsec-sa` statement added in Junos OS Release 8.3.

RELATED DOCUMENTATION

Example: Configuring OSPFv2 Sham Links

sham-link | 16921

shaping

IN THIS SECTION

- Syntax | 16925
- Hierarchy Level | 16925
- Description | 16926
- Required Privilege Level | 16926
- Release Information | 16926

Syntax

```
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
length);
    queue-length number;
}
```

Hierarchy Level

```
[edit interfaces interface-name atm-options vpi vpi-identifier],
[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-name unit logical-unit-number address address family family
multipoint-destination address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
address address family family multipoint-destination  address]
```

Description

For Circuit Emulation PICs, specify traffic shaping in the ingress and egress directions.

VP tunnels are not supported on multipoint interfaces.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Defining Virtual Path Tunnels

shaping-rate (Traffic Control Profiles)

IN THIS SECTION

- Syntax | 16927
- Hierarchy Level | 16927
- Description | 16927
- Default | 16928
- Options | 16928
- Required Privilege Level | 16928
- Release Information | 16928

Syntax

```
shaping-rate (rate | percent percentage);
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name]
```

NOTE: Only devices that support enhanced transmission selection (ETS) or hierarchical scheduling support the `traffic-control-profiles` hierarchy.

Description

Configure the shaping rate. The shaping rate throttles the rate of packet transmission by setting a maximum bandwidth (rate in bits per second) or a maximum percentage of bandwidth for a queue or a forwarding class set. You specify the maximum bandwidth for a queue by using a scheduler map to associate a forwarding class (queue) with a scheduler that has a configured shaping rate.

For ETS or hierarchical scheduling configuration, you specify the maximum bandwidth for a forwarding class set by setting the shaping rate for a traffic control profile. Then you associate the scheduler map with the traffic control profile, and then you apply the traffic control profile and a forwarding class set to an interface.

The shaping rate as an absolute maximum usage and not as additional usage beyond the configured transmit rate (the minimum guaranteed bandwidth for a queue) or the configured guaranteed rate (the minimum guaranteed bandwidth for a forwarding class set).

NOTE: When you set the maximum bandwidth (`shaping-rate` value) for a queue or for a priority group at 100 Kbps or less, the traffic shaping behavior is accurate only within +/– 20 percent of the configured `shaping-rate` value.

NOTE: On QFX5200, QFX5100, and EX4600 devices, we recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

NOTE: On QFX5200 Series switches, a granularity of 64kbps is supported for the shaping rate. Therefore, the shaping rate on queues for 100g interfaces might not be applied correctly.

NOTE: QFX10000 Series switches do not support the shaping-rate statement. However, you can configure the transmit-rate `exact` option to prevent a queue from consuming more bandwidth than you want the queue to consume.

On QFX10000 Series switches, we recommend that you use the transmit rate to set a limit on the amount of bandwidth that receives strict-high priority treatment on a strict-high priority queue. Traffic up to the transmit rate receives strict-high priority treatment. Traffic in excess of the transmit rate is treated as best-effort traffic that receives the strict-high priority queue excess rate weight of “1”. Do not use a shaping rate to set a maximum bandwidth limit on strict-high priority queues on QFX10000 Series switches.

Default

If you do not configure a shaping rate, the default shaping rate is 100 percent (all of the available bandwidth), which is the equivalent of no rate shaping.

Options

`percent` *`percentage`*—Shaping rate as a percentage of the available interface bandwidth.

- **Range:** 1 through 100 percent

`rate`—Peak (maximum) rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

- **Range:** 1000 through 10,000,000,000 bps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Queue Schedulers

Example: Configuring Queue Schedulers for Port Scheduling

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Understanding CoS Output Queue Schedulers

Understanding CoS Port Schedulers on QFX Switches

Understanding CoS Traffic Control Profiles

shaping-rate (Applying to an Interface)

IN THIS SECTION

- Syntax | **16929**
- Hierarchy Level | **16929**
- Description | **16930**
- Default | **16931**
- Options | **16931**
- Required Privilege Level | **16931**
- Release Information | **16931**

Syntax

```
shaping-rate rate;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface. Applying a shaping rate can help ensure that higher-priority services do not starve lower-priority services.

For physical interfaces, by default, shaping is not configured and traffic can be up to the line rate for that interface. Port shaping enables you to control the amount of traffic passing through a physical interface. Port shaping enables you to shape the aggregate traffic through an interface to a rate that is less than the line rate for that interface. This can be useful to reduce downstream congestion.

NOTE: For platforms that support both logical and physical interface shaping, logical and physical interface traffic shaping rates are mutually exclusive. This means you can include the `shaping-rate` statement at the `[edit class-of-service interfaces interface-name]` hierarchy level or at the `[edit class-of-service interfaces interface-name unit logical-unit-number]` hierarchy level, but not at both.

NOTE: For MX Series routers and for EX Series switches, the shaping rate value for the physical interface at the `[edit class-of-service interfaces interface-name]` hierarchy level must be a minimum of 160 Kbps. If the value is less than the sum of the logical interface guaranteed rates, you cannot apply the shaping rate to a physical interface.

For PTX Series routers, the shaping rate value for the physical interface at the `[edit class-of-service interfaces interface-name]` hierarchy level must be a minimum of 1 Gbps and an incremental granularity of 0.1 percent of the physical interface speed after that (for example, 10 Mbps increments on a 10 Gbps interface).

On EX4650, QFX5110, QFX5120, QFX5200, QFX5210 Series switches, when you configure a shaping rate on an aggregated Ethernet (ae) interface, all members of the ae interface are shaped at the configured shaping rate. For example, consider an interface ae0 that consists of three interfaces: xe-0/0/0, xe-0/0/1, and xe-0/0/2. If you configure a shaping rate of X Mpbs on ae0, traffic up to the rate of X Mpbs flows through each of the three interfaces. Therefore, the total traffic flowing through ae0 can be at the rate of 3X Mbps.

NOTE: Starting in Junos OS Release 24.2R1, we support `shaping-rate` on aggregated Ethernet (AE) interfaces and logical interface of aggregated Ethernet (AE) interfaces on Junos OS PTX Series platforms.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the `shaping-rate` statement at the `[edit class-of-service traffic-control-profiles]`

hierarchy level. With this configuration approach, you can independently control the delay-buffer rate, as described in *Oversubscribing Interface Bandwidth*.

Default

If you do not include this statement at the `[edit class-of-service interfaces interface-name unit logical-unit-number]` hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the `[edit class-of-service interfaces interface-name]` hierarchy level, the default physical interface bandwidth is the average of unused bandwidth for the number of physical interfaces that require default bandwidth treatment.

Options

rate Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000 bps.

> **NOTE**: For all MX Series and EX Series interfaces, the rate can be from 65,535 to 6,400,000,000,000 bps.
>
> For all PTX Series interfaces, the rate can be from 1,000,000,000 to 160,000,000,000 bps in increments of 0.1 percent of the interface speed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`[edit class-of-service interfaces interface-name]` hierarchy level added in Junos OS Release 7.5.

RELATED DOCUMENTATION

Applying Scheduler Maps Overview

shaping-rate (CoS Adaptive Shapers)

IN THIS SECTION

- Syntax | 16932
- Hierarchy Level | 16932
- Description | 16932
- Options | 16932
- Required Privilege Level | 16933
- Release Information | 16933

Syntax

```
shaping-rate (percent percentage | rate);
```

Hierarchy Level

```
[edit class-of-service adaptive-shapers adaptive-shaper-name trigger type]
```

Description

Define the list of trigger types and associated rates.

Options

- **percent *percentage***—Shaping rate as a percentage of the available interface bandwidth.
- **Range:** 0 through 100 percent

- *rate*—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).
- **Range:** 3200 through 32,000,000,000 bps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

trigger (CoS) | **30159**

shaping-rate (CoS Interfaces)

IN THIS SECTION

- Syntax | **16934**
- Hierarchy Level | **16934**
- Description | **16934**
- Default | **16934**
- Options | **16935**
- Required Privilege Level | **16935**
- Release Information | **16935**
- Platform-Specific Behavior | **16935**

Syntax

```
shaping-rate rate <overhead bytes> ;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface.

Logical and physical interface traffic shaping can be configured together. This means you can include the `shaping-rate` statement at the `[edit class-of-service interfaces interface interface-name]` hierarchy level *and* the `[edit class-of-service interfaces interface interface-name unit logical-unit-number]` hierarchy level. If you configure traffic shaping at both the logical and physical interface levels, the logical interface shaping credit is checked and updated before the physical interface shaping credit.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the `shaping-rate` statement at the `[edit class-of-service traffic-control-profiles]` hierarchy level. With this configuration approach, you can independently control the delay-buffer rate.

On the physical interface, you can set the Layer 2 overhead adjustment to the shaping rate calculation at egress.

Default

If you do not include this statement at the `[edit class-of-service interfaces interface interface-name unit logical-unit-number]` hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the `[edit class-of-service interfaces interface interface-name]` hierarchy level, the default physical interface bandwidth is the average of unused bandwidth for the number of physical interfaces that require default bandwidth treatment.

Options

rate Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000 bps

overhead Layer 2 shaping overhead adjustment to be applied at egress (bytes).

- **Range:** -62 through 192

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

`overhead` option introduced in Junos OS Release 18.1.

Platform-Specific Behavior

Use Feature Explorer to confirm platform and release support for specific features.

Use the following table to review platform-specific behavior for your platforms:

Platform	Difference
MX Series	*Only* MX Series routers support the `overhead` *bytes* option.

RELATED DOCUMENTATION

policer-overhead | **12643**

shaping-rate (CoS Virtual Channels)

IN THIS SECTION

- Syntax | 16936
- Hierarchy Level | 16936
- Description | 16936
- Options | 16936
- Required Privilege Level | 16937
- Release Information | 16937

Syntax

```
shaping-rate (percent percentage | rate);
```

Hierarchy Level

```
[edit class-of-service virtual-channel-groups group-name virtual-channel-name]
```

Description

Define the shaping rates to be associated with the virtual channel.

Options

- **percent** ***percentage***—Shaping rate as a percentage of the available interface bandwidth.
- **Range:** 0 through 100 percent
- ***rate***—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).
- **Range:** 3200 through 32,000,000,000 bps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

default (CoS) | 5187

scheduler-map (CoS Virtual Channels) | 16350

virtual-channel-group (CoS Interfaces) | 30740

virtual-channel-groups | 30742

virtual-channels | 30743

shaping-rate (Dynamic Traffic Shaping and Scheduling)

IN THIS SECTION

- Syntax | 16937
- Hierarchy Level | 16938
- Description | 16938
- Options | 16938
- Required Privilege Level | 16938
- Release Information | 16939

Syntax

```
shaping-rate (rate | predefined-variable) <burst-size  bytes | $junos-cos-shaping-rate-burst>;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name],
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
```

Description

Configure a shaping rate for a logical interface or a scheduler. The sum of the shaping rates for all logical interfaces on the physical interface can exceed the physical interface bandwidth. This practice is known as oversubscription of the peak information rate (PIR).

Options

rate—Peak rate in bits per second (bps). You can specify the value as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 160,000,000,000 bps

predefined-variable—One of the following Junos predefined variables. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

- `$junos-cos-shaping-rate`—Variable for the shaping rate that is specified for the logical interface. Use this variable at the [`edit dynamic-profiles` *`profile-name`* `class-of-service traffic-control-profiles` *`profile-name`*] hierarchy level.

- `$junos-cos-scheduler-shaping-rate`—Variable for the shaping rate that is specified for a scheduler. Use this variable at the [`edit dynamic-profiles` *`profile-name`* `class-of-service schedulers` *`scheduler-name`*] hierarchy level.

`burst-size` *`bytes`*—(Optional) Maximum burst size, in bytes.

- **Range:** 0 through 1,000,000,000

`$junos-cos-shaping-rate-burst`—(Optional) Variable for the burst-size that is specified for the shaping rate. Use this variable at the [`edit dynamic-profiles` *`profile-name`* `class-of-service traffic-control-profile`] hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

The `$junos-cos-shaping-rate` variable for traffic-control profiles introduced in Junos OS Release 9.4.

The `$junos-cos-scheduler-shaping-rate` variable for schedulers introduced in Junos OS Release 10.2.

Option `burst-size` introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Traffic Scheduling and Shaping for Subscriber Access

output-traffic-control-profile (Dynamic CoS Definition) | **12018**

shaping-rate-excess-high

IN THIS SECTION

- Syntax | **16939**
- Hierarchy Level | **16940**
- Description | **16940**
- Default | **16940**
- Options | **16940**
- Required Privilege Level | **16941**
- Release Information | **16941**

Syntax

```
shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name],
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For MIC and MPC interfaces on MX Series routers, configure a shaping rate and optional burst size for high-priority excess traffic. This can help to make sure higher priority services do not starve lower priority services.

> **NOTE**: Option `burst-size` is not supported on MPC5 interfaces. The `burst-size` configuration is allowed on these interfaces, but does not take effect.

Default

If you do not include this statement, the default shaping rate for this priority is determined by the `shaping-rate` statement in the traffic control profile.

Options

`percent` *`percentage`*—Specify the shaping rate as a percentage of the overall shaping rate.

- **Range:** 1 through 100

`rate`—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000

> **NOTE**: Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.

- **Default:** None

`burst-size` *`bytes`*—Maximum burst size, in bytes.

- **Range:** 0 through 1,000,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`percent` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Per-Priority Shaping on MIC and MPC Interfaces Overview

Oversubscribing Interface Bandwidth

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

shaping-rate-excess-low | **16941**

shaping-rate-priority-high | **16950**

shaping-rate-priority-low | **16952**

shaping-rate-priority-medium | **16954**

shaping-rate-excess-low

IN THIS SECTION

- Syntax | **16942**
- Hierarchy Level | **16942**
- Description | **16942**
- Default | **16942**
- Options | **16942**
- Required Privilege Level | **16943**
- Release Information | **16943**

Syntax

```
shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name],
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For MIC and MPC interfaces on MX Series routers, configure a shaping rate and optional burst size for low-priority excess traffic. This can help to make sure higher priority services do not starve lower priority services.

> **NOTE**: Option `burst-size` is not supported on MPC5 interfaces. The `burst-size` configuration is allowed on these interfaces, but does not take effect.

Default

If you do not include this statement, the default shaping rate for this priority is determined by the `shaping-rate` statement in the traffic control profile.

Options

`percent` *percentage*—Specify the shaping rate as a percentage of the overall shaping rate.

- **Range:** 1 through 100

rate—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000

> **NOTE**: Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.

- **Default:** None

`burst-size` *`bytes`*—Maximum burst size, in bytes.

- **Range:** 0 through 1,000,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`percent` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Per-Priority Shaping on MIC and MPC Interfaces Overview

Oversubscribing Interface Bandwidth

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

shaping-rate-excess-high | **16939**

shaping-rate-priority-high | **16950**

shaping-rate-priority-low | **16952**

shaping-rate-priority-medium | **16954**

shaping-rate-excess-medium-high

IN THIS SECTION

- Syntax | **16944**
- Hierarchy Level | **16944**
- Description | **16944**
- Default | **16944**

- Options | 16944
- Required Privilege Level | 16945
- Release Information | 16945

Syntax

```
shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name],
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For MPCs (MPC7E-MRATE, MPC7E-10G, MPC8E, and MPC9E) interfaces on MX Series routers, configure a shaping rate and optional burst size for medium-high-priority excess traffic. This can help to make sure higher priority services do not starve lower priority services.

Default

If you do not include this statement, the default shaping rate for this priority is determined by the `shaping-rate` statement in the traffic control profile.

Options

`percent` *`percentage`*—Specify the shaping rate as a percentage of the overall shaping rate.

- **Range:** 1 through 100

`rate`—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000

- **Default:** None

`burst-size` *bytes*—Maximum burst size, in bytes.

- **Range:** 1 through 1,000,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`percent` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Per-Priority Shaping on MIC and MPC Interfaces Overview

Oversubscribing Interface Bandwidth

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

shaping-rate-excess-low | **16941**

shaping-rate-priority-high | **16950**

shaping-rate-priority-low | **16952**

shaping-rate-priority-medium | **16954**

shaping-rate-excess-medium-low

IN THIS SECTION

- Syntax | **16946**
- Hierarchy Level | **16946**
- Description | **16946**
- Default | **16946**

- Options | 16946
- Required Privilege Level | 16947
- Release Information | 16947

Syntax

```
shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name],
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For MPC7E-MRATE, MPC7E-10G, MPC8E, and MPC9E interfaces on MX Series routers, configure a shaping rate and optional burst size for medium-low-priority excess traffic. This can help to make sure higher priority services do not starve lower priority services.

Default

If you do not include this statement, the default shaping rate for this priority is determined by the `shaping-rate` statement in the traffic control profile.

Options

`percent` *`percentage`*—Specify the shaping rate as a percentage of the overall shaping rate.

- **Range:** 1 through 100

`rate`—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000

- **Default:** None

`burst-size` *bytes*—Maximum burst size, in bytes.

- **Range:** 1 through 1,000,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`percent` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

enhanced-priority-mode | **6404**

Per-Priority Shaping on MIC and MPC Interfaces Overview

shaping-rate-excess-medium-high | **16943**

shaping-rate-priority-medium-low | **16956**

shaping-rate-priority-strict-high | **16958**

shaping-rate (Oversubscribing an Interface)

IN THIS SECTION

- Syntax | **16948**
- Hierarchy Level | **16948**
- Description | **16948**
- Default | **16948**
- Options | **16948**
- Required Privilege Level | **16949**

Release Information | 16949

Syntax

```
shaping-rate (percent percentage / rate) <burst-size bytes>;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name]
```

Description

FConfigure a shaping rate for a logical interface. On some interfaces, you can also configure an optional burst size for a logical interface. This can help to ensure that higher-priority services do not starve lower-priority services.

For physical interfaces on T4000 router interfaces on Type 5 FPCs and on PTX Series routers, configure traffic shaping rate.

The sum of the shaping rates for all logical interfaces on the physical interface can exceed the physical interface bandwidth. This practice is known as oversubscription of the peak information rate (PIR).

Default

The default behavior depends on various factors. For more information, see *Bandwidth and Delay Buffer Allocations by Configuration Scenario*.

Options

percent *percentage*

Shaping rate as a percentage of the available interface bandwidth.

- **Range:** 1 through 100 percent

rate

Peak shaping rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

Range:

- 1000 through 6,400,000,000,000 bps
- T4000 routers with Type 5 FPCs—The shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of `shaping-rate` is limited by the maximum transmission rate of the interface.

`burst-size` *bytes*

(Optional) Maximum burst size, in bytes.

- **Range:** 0 through 1,000,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Option `burst-size` introduced in Junos OS Release 9.4.

> **NOTE**: Option `burst-size` is not supported on MPC5 interfaces. The `burst-size` configuration is allowed on these interfaces, but does not take effect.

Statement introduced for PTX Series Packet Transport Routers in Junos OS Release 16.1. PTX Series Packet Transport Routers do not support the `burst-size` option or defining the `shaping-rate` as a percentage.

RELATED DOCUMENTATION

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

Oversubscribing Interface Bandwidth

output-traffic-control-profile | **12016**

shaping-rate-priority-high

IN THIS SECTION

- Syntax | 16950
- Hierarchy Level | 16950
- Description | 16950
- Default | 16951
- Options | 16951
- Required Privilege Level | 16951
- Release Information | 16951

Syntax

```
shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name],
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For MIC and MPC interfaces on MX Series routers, configure a shaping rate and optional burst size for high priority traffic. This can help to make sure higher priority services do not starve lower priority services.

NOTE: Option `burst-size` is not supported on MPC5 interfaces. The `burst-size` configuration is allowed on these interfaces, but does not take effect.

Default

If you do not include this statement, the default shaping rate for this priority is determined by the `shaping-rate` statement in the traffic control profile.

Options

`percent` *`percentage`*—Specify the shaping rate as a percentage of the overall shaping rate.

- **Range:** 1 through 100

`rate`—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000

> **NOTE**: Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.

- **Default:** None

`burst-size` *`bytes`*—Maximum burst size, in bytes.

- **Range:** 0 through 1,000,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`percent` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Per-Priority Shaping on MIC and MPC Interfaces Overview

Oversubscribing Interface Bandwidth

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

shaping-rate-excess-high | 16939

shaping-rate-excess-low | 16941

shaping-rate-priority-low | 16952

shaping-rate-priority-medium | 16954

shaping-rate-priority-low

IN THIS SECTION

- Syntax | 16952
- Hierarchy Level | 16952
- Description | 16953
- Default | 16953
- Options | 16953
- Required Privilege Level | 16953
- Release Information | 16954

Syntax

```
shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name],
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For MIC and MPC interfaces on MX Series routers, configure a shaping rate and optional burst size for low priority traffic. This can help to make sure higher priority services do not starve lower priority services.

> **NOTE**: Option `burst-size` is not supported on MPC5 interfaces. The `burst-size` configuration is allowed on these interfaces, but does not take effect.

Default

If you do not include this statement, the default shaping rate for this priority is determined by the `shaping-rate` statement in the traffic control profile.

Options

`percent` *percentage*—Specify the shaping rate as a percentage of the overall shaping rate.

- **Range:** 1 through 100

rate—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000

> **NOTE**: Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.

- **Default:** None

`burst-size` *bytes*—Maximum burst size, in bytes.

- **Range:** 0 through 1,000,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`percent` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Per-Priority Shaping on MIC and MPC Interfaces Overview

Oversubscribing Interface Bandwidth

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

shaping-rate-excess-high | **16939**

shaping-rate-excess-low | **16941**

shaping-rate-priority-high | **16950**

shaping-rate-priority-medium | **16954**

shaping-rate-priority-medium

IN THIS SECTION

- Syntax | **16954**
- Hierarchy Level | **16955**
- Description | **16955**
- Default | **16955**
- Options | **16955**
- Required Privilege Level | **16956**
- Release Information | **16956**

Syntax

```
shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name],
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For MIC and MPC interfaces on MX Series routers, configure a shaping rate and optional burst size for medium priority traffic. This can help to make sure higher priority services do not starve lower priority services.

> **NOTE**: Option `burst-size` is not supported on MPC5 interfaces. The `burst-size` configuration is allowed on these interfaces, but does not take effect.

Default

If you do not include this statement, the default shaping rate for this priority is determined by the `shaping-rate` statement in the traffic control profile.

Options

`percent` *`percentage`*—Specify the shaping rate as a percentage of the overall shaping rate.

- **Range:** 1 through 100

`rate`—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000

> **NOTE**: Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.

- **Default:** None

`burst-size` *`bytes`*—Maximum burst size, in bytes.

- **Range:** 0 through 1,000,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`percent` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Per-Priority Shaping on MIC and MPC Interfaces Overview

Oversubscribing Interface Bandwidth

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

shaping-rate-excess-high | **16939**

shaping-rate-excess-low | **16941**

shaping-rate-priority-high | **16950**

shaping-rate-priority-low | **16952**

shaping-rate-priority-medium-low

IN THIS SECTION

- Syntax | **16957**
- Hierarchy Level | **16957**
- Description | **16957**
- Default | **16957**
- Options | **16957**
- Required Privilege Level | **16957**
- Release Information | **16958**

Syntax

```
shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name],
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For MPC7E-MRATE, MPC7E-10G, MPC8E, and MPC9E interfaces on MX Series routers, configure a shaping rate and optional burst size for medium-low priority traffic. This can help to make sure higher priority services do not starve lower priority services.

Default

If you do not include this statement, the default shaping rate for this priority is determined by the `shaping-rate` statement in the traffic control profile.

Options

`percent` *`percentage`*—Specify the shaping rate as a percentage of the overall shaping rate.

- **Range:** 1 through 100

`rate`—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000
- **Default:** None

`burst-size` *`bytes`*—Maximum burst size, in bytes.

- **Range:** 1 through 1,000,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`percent` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

- enhanced-priority-mode | 6404
- *Per-Priority Shaping on MIC and MPC Interfaces Overview*
- shaping-rate-excess-medium-high | 16943
- shaping-rate-excess-medium-low | 16945
- shaping-rate-priority-strict-high | 16958

shaping-rate-priority-strict-high

IN THIS SECTION

- Syntax | 16958
- Hierarchy Level | 16959
- Description | 16959
- Default | 16959
- Options | 16959
- Required Privilege Level | 16959
- Release Information | 16960

Syntax

```
shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name],
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
```

Description

For MPC7E-MRATE, MPC7E-10G, MPC8E, and MPC9E interfaces on MX Series routers, configure a shaping rate and optional burst size for strictly high priority traffic. This can help to make sure higher priority services do not starve lower priority services.

Default

If you do not include this statement, the default shaping rate for this priority is determined by the `shaping-rate` statement in the traffic control profile.

Options

`percent` *`percentage`*—Specify the shaping rate as a percentage of the overall shaping rate.

- **Range:** 1 through 100

`rate`—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 6,400,000,000,000
- **Default:** None

`burst-size` *`bytes`*—Maximum burst size, in bytes.

- **Range:** 1 through 1,000,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`percent` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

- enhanced-priority-mode | 6404
- *Per-Priority Shaping on MIC and MPC Interfaces Overview*
- shaping-rate-excess-medium-high | 16943
- shaping-rate-excess-medium-low | 16945
- shaping-rate-priority-medium-low | 16956

shaping-rate (Schedulers)

IN THIS SECTION

- Syntax | 16960
- Hierarchy Level | 16961
- Description | 16961
- Default | 16961
- Options | 16961
- Required Privilege Level | 16962
- Release Information | 16962

Syntax

```
shaping-rate (percent percentage | rate) <burst-size  bytes>;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Define a limit on bandwidth usage for a forwarding class/queue.

The `transmit-rate` statement at the `[edit class-of-service schedulers scheduler-name]` hierarchy level configures the minimum bandwidth allocated to a queue. The transmission bandwidth can be configured as an exact value or allowed to exceed the configured rate if additional bandwidth is available from other queues.

Configure the shaping rate as an absolute maximum usage and not the additional usage beyond the configured transmit rate.

Default

If you do not include this statement, the default shaping rate is 100 percent, which is the same as no shaping at all.

Options

percent *percentage*

Shaping rate as a percentage of the available interface bandwidth.

- **Range:** 0 through 100 percent

> **NOTE**: If you configure a shaping rate as a percent in a scheduler, the effective shaping rate is calculated based on the following hierarchy:
>
> 1. Logical interface shaping rate, if configured
> 2. Physical interface shaping rate, if configured
> 3. Physical interface bandwidth

rate

Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 3200 through 6,400,000,000,000 bps

burst-size *bytes*

Maximum burst size, in bytes. The burst value determines the number of rate credits that can accrue when the queue or scheduler node is held in the inactive round robin.

- **Range:** 0 through 1,000,000,000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The `burst-size` option added in Junos OS Release 11.4.

RELATED DOCUMENTATION

Applying Scheduler Maps Overview

shared-bandwidth-policer (Configuring)

IN THIS SECTION

- Syntax | **16963**
- Hierarchy Level | **16963**
- Description | **16963**
- Required Privilege Level | **16963**
- Release Information | **16963**

Syntax

```
shared-bandwidth-policer;
```

Hierarchy Level

```
[edit firewall policer policer-name],
[edit firewall three-color-policer policer-name],
[edit firewall hierarchical-policer policer-name]
```

Description

Policer instances share bandwidth. This enables configuration of interface-specific policers applied on an aggregated Ethernet bundle or an aggregated SONET bundle to match the effective bandwidth and burst-size to user-configured values. This feature is supported on the following platforms: T Series routers, M120, M10i, M7i (CFEB-E only), M320 (SFPC only), MX240, MX480, and MX960 with DPC, MIC, and MPC interfaces and EX Series switches.

NOTE: This statement is not supported on T4000 Type 5 FPCs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support for MX Series MPC and MIC interfaces added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Policer Support for Aggregated Ethernet Interfaces Overview

shared-buffer

IN THIS SECTION

- Syntax | 16964
- Hierarchy Level | 16964
- Description | 16965
- Options | 16966
- Required Privilege Level | 16966
- Release Information | 16966

Syntax

```
shared-buffer {
    egress {
        buffer-partition (lossless | lossy | multicast) {
            percent percent;
            dynamic-threshold threshold-value;
        }
        percent percent;
    }
    ingress {
        percent percent;
        buffer-partition (lossless | lossless-headroom | lossy) {
            percent percent;
            dynamic-threshold threshold-value;
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Configure the global shared buffer pool allocation to ports. Shared buffers are a pool of buffer space that the system can allocate dynamically across all of its ports as memory space is needed. Some buffer space is reserved for dedicated buffers (buffers allocated permanently to ports), headroom buffers (buffers that help prevent packet loss on lossless flows), and other buffers.

The device uses the shared-buffer pool to absorb traffic bursts after the dedicated-buffer-pool is exhausted. The shared pool threshold is dynamically calculated based on a factor called "*alpha*".

Configure the way the system uses the available (user-configurable) buffer space by setting the `shared-buffer` percentage for the ingress buffer pool and for the egress buffer pool.

The percentage you specify is the percentage of available buffer space allocated to the global shared ingress buffer pool or to the global shared egress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports.)

You can adjust the maximum size of the shared-buffer pool by configuring the dynamic-threshold values:

- By adjusting the value for the egress partition (the calculation includes the alpha value and the number of competing queues).

```
[edit class-of-service shared-buffer egress buffer-partition dynamic-threshold threshold-value]
```

- By adjusting the value for the ingress partition (the calculation includes the alpha value and the number of competing queues).

```
[edit class-of-service shared-buffer ingress buffer-partition dynamic-threshold threshold-value
```

CAUTION: Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until the buffer reprogramming is complete.

You can also partition the ingress shared buffer pool and the egress shared buffer pool to adjust the buffer allocations for different mixes of network traffic (best-effort, lossless, multicast) using the `buffer-partition` statement.

NOTE: If you commit a buffer configuration for which the device does not have sufficient resources, the device might log an error instead of returning a commit error. In that case, a syslog message is displayed on the console. For example:

```
user@host# commit
configuration check succeeds

Message from syslogd@host at Jun 13 11:11:10  ...
host dc-pfe: Not enough Ingress Lossless headroom.(Already allocated more).
Dedicated : 14340 Lossy : 47100 Lossless 4239 Headroom 21195 Avail : 20781
commit complete
```

If the buffer configuration commits but you receive a syslog message that indicates the configuration cannot be implemented, you can:

- Reconfigure the buffers or reconfigure other parameters (for example, the PFC configuration, which affects the need for lossless headroom buffers and lossless buffers—the more priorities you pause, the more lossless and lossless headroom buffer space you need), then attempt the commit operation again.
- Roll back the device to the last successful configuration.

If you receive a syslog message that says the buffer configuration cannot be implemented, you must take corrective action. If you do not fix the configuration or roll back to a previous successful configuration, the system behavior is unpredictable.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic

Configuring Global Ingress and Egress Shared Buffers

Understanding CoS Buffer Configuration

shared-instance

IN THIS SECTION

- Syntax | 16967
- Hierarchy Level | 16967
- Description | 16968
- Options | 16968
- Required Privilege Level | 16968
- Release Information | 16968

Syntax

```
shared-instance instance-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number input-traffic-control-profile],
[edit class-of-service interfaces interface-name unit logical-unit-number output-traffic-control-profile]
```

Description

For Gigabit Ethernet IQ2 and IQ2E PICs only, apply a shared traffic scheduling and shaping profile to the logical interface.

Options

instance-name—Name of the shared scheduler and shaper to be applied to this interface

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Shaping on 10-Gigabit Ethernet IQ2 PICs

traffic-control-profiles | 30009

shared-key

IN THIS SECTION

- Syntax | 16969
- Hierarchy Level | 16969
- Description | 16969
- Options | 16969
- Required Privilege Level | 16969
- Release Information | 16969

Syntax

```
shared-key value;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

Configure the authentication key value.

Options

value—Secret authentication value shared between a control source and destination.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Control Source

shared-name

IN THIS SECTION

- Syntax | 16970
- Hierarchy Level | 16970
- Description | 16970
- Options | 16970
- Required Privilege Level | 16971
- Release Information | 16971

Syntax

```
shared-name filter-shared-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family filter input filter-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family filter output filter-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number filter input
filter-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number filter
output filter-name]
```

Description

Apply a filter shared name to a dynamic filter.

Options

filter-shared-name— Name of the specific shared filter or $junos-interface-set-name.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- Guidelines for Configuring Firewall Filters
- *Understanding Dynamic Firewall Filters*
- *Classic Filters Overview*
- *Basic Classic Filter Syntax*

shared-scheduler

IN THIS SECTION

- Syntax | **16971**
- Hierarchy Level | **16972**
- Description | **16972**
- Required Privilege Level | **16972**
- Release Information | **16972**

Syntax

```
shared-scheduler;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

For Gigabit Ethernet IQ2 PICs only, enable shared schedulers and shapers on this interface. This statement and the `per-unit-scheduler` statement are mutually exclusive. Even so, you can configure one logical interface for each shared instance. This effectively provides the functionality of per-unit scheduling.

For Gigabit Ethernet IQ2 and Ethernet Enhanced IQ2 (IQ2E) PICs on M320 routers, enable shared schedulers on aggregated Ethernet interfaces in link protection mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

Configuring Shaping on 10-Gigabit Ethernet IQ2 PICs

Configuring Shared Scheduling on Aggregated Ethernet Interfaces

traffic-control-profiles | **30009**

shared-secret (RADIUS Snoop Segment)

IN THIS SECTION

- Syntax | 16973
- Hierarchy Level | 16973
- Description | 16973
- Options | 16973
- Required Privilege Level | 16973
- Release Information | 16974

Syntax

```
shared-secret secret;
```

Hierarchy Level

```
[edit access radius snoop-segments segment-name]
```

Description

Configure a shared secret to be used by the MX Series router and the RADIUS client. If the shared secrets do not match, the subscriber session is not set up.

Options

secret Shared secret. The maximum length is 64 characters.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview

IP-Based Subscriber Setup Overview

shdsl-options

IN THIS SECTION

- Syntax | 16974
- Hierarchy Level | 16975
- Description | 16975
- Required Privilege Level | 16975
- Release Information | 16975

Syntax

```
shdsl-options {
    annex (annex-a | annex-b);
    line-rate line-rate;
    loopback (local | remote | payload);
    snr-margin {
        snext margin;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit logical-systems logical-system-name interfaces interface-name
```

Description

For J Series Services Routers only, configure symmetric DSL (SHDSL) options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

shell

IN THIS SECTION

- Syntax | 16976
- Hierarchy Level | 16976
- Description | 16976
- Additional Information | 16976
- Options | 16976
- Required Privilege Level | 16977
- Release Information | 16977

Syntax

```
shell
```

Hierarchy Level

```
[edit system syslog]
```

Description

Configure the logging of commands executed from the shell.

Additional Information

To log command executed from the shell configure shell logging by including the `system syslog shell` statement at the [edit] hierarchy.

```
[edit]
user@host# set system syslog shell
```

By default log entries are written to the messages log file.

> **NOTE**: In Junos OS Evolved, `audit-sh` prompt gets displayed when you enable the feature and perform `start shell` from CLI environment. In Junos OS, `audit-sh` doesn't get displayed.

> **NOTE**: In Junos OS Evolved, `audit-sh` gets displayed in the audit shell prompt when you enable the feature.

Options

A variety of standard syslog options are supported with the `shell` logging option. See "syslog (System)" on page 16975 for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

start shell | 28498

shellcode

IN THIS SECTION

- Syntax | 16977
- Hierarchy Level | 16978
- Description | 16978
- Options | 16978
- Required Privilege Level | 16978
- Release Information | 16978

Syntax

```
shellcode (all | intel | no-shellcode | sparc);
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type anomaly],

[edit security idp custom-attack attack-name attack-type signature]
```

Description

Shellcode signifies that the attack is a shellcode attack and is capable of creating its own shell.

Options

- `all`—All shellcode checks will be performed if this attack matches.
- `intel`—Basic shellcode checks and Intel-specific shellcode checks will be performed.
- `no-shellcode`—No shellcode checks will be performed.
- `sparc`—Basic shellcode checks and Sparc-specific shellcode checks will be performed.
- **Default:** Basic shellcode checks will be performed when this field is not configured.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

shmlog (Shared Memory Log)

IN THIS SECTION

Syntax | 16979

- Hierarchy Level | 16979
- Description | 16979
- Options | 16980
- Required Privilege Level | 16981
- Release Information | 16981

Syntax

```
shmlog {
    disable;
    file filename <files maximum-no-files> <size maximum-file-size>;
    filtering enable;
    log-name {
        all;
        logname {
            <brief | detail | extensive | none | terse>;
            <file-logging |no-file-logging>;
        }
        }
    log-type (debug | info | notice);
|
```

Hierarchy Level

```
[edit system services subscriber-management overrides]
```

Description

Junos OS uses a shared memory space to store log entries for subscriber service daemons including jpppd, jdhcpd, jl2tpd, autoconfd, bbe-smgd, authd, cosd, and dfwd. Shared memory logging is enabled by default and occurs at the client level. You can view the shmlogs on a per subscriber basis, or use filters to retrieve logs according to a variety of different parameters such as interface name, IP address, session

ID, subnet, and VLAN in addition to the Client Identifier or Client DUID. Filtering is disabled by default. To see a complete list of supported filters, use this command:

```
user@rdevice> show shmlog entries logname all ?
```

When viewing logs you can limit results on the basis of event flags that include interface events, routing process interaction events, l2tp tunneling events, and ldap authentication events. To see a complete list of supported flags, use this command:

```
user@rdevice> show shmlog entries logname all flag-name ?
```

NOTE: Some platforms other than MX Series routers use shared memory logs for internal processes. These logs are not intended for customer use.

Options

disable Name of the command to override the default behavior. Use this option to disable shared memory logging; it is always enabled otherwise.

file Name of the file containing the shmlogs. Use this option to redirect shmlogs to a file for file-based logging. Specify the file name, define the number of files (from 2 to 1000), and set the maximum file size (from 10240 to 1073741824 bytes). Data will be written to the **/var/log/shmlog/** directory. Files follow this naming convention: **<cfg-file-name>-<daemon>-<severity>.log**. The shmlog files are not human-readable, so to access the logs you must first run the following command to generate a file in the **/var/log/<file-name>/** directory with logs from all daemons:

```
user@rdevice> show shmlog entries filename /var/log/shmlog/<file-name>* logname all
```

If you then want to view logs from a specific daemon, you need to run the following command to generate a file under the **/var/log/<file-name>/** directory with complete logs:

```
user@rdevice> show shmlog entries filename /var/log/shmlog/<filename>  logname authd*
```

filtering Command to enable filtering. Filtering is subscriber centric and is useful for debugging and troubleshooting. It is disabled by default so you must use this option to enable it.

For example, if you want to quickly view the transmit packet logs for subscribers with interface-name pp0.100, you could use the following command to display only the relevant results:

```
user@rdevice> show shmlog entries logname jpppd* interface-name pp0.100 flag transmit-
packets
```

To debug sessions according to the interface name, use this command:

```
user@rdevice> show shmlog entries logname all interface-name pp0.100
```

To debug sessions that are logging in via VLAN 7 on physical-interface ge-0/0/0, use this command:

```
user@rdevice> show shmlog entries logname all vlan 7 physical-interface ge-0/0/0
```

log-name Name of the file containing the log output. Use this option to override all logs or a specified log, and to set the verbosity level (brief, detail, extensive, none, or terse). For example, to configure **bbe-autoconf-info** for detailed file logging, you would use the following command:

```
user@rdevice> [edit system services subscriber-management overrides shmlog]
user@rdevice> set log-name bbe-autoconf-info detail file-logging
```

log-type Severity level of the collected logs. Use this option to configure the severity level for captured logs (notice, info, or debug).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

short-cycle-protection (DHCP Local Server and Relay Agent)

IN THIS SECTION

- Syntax | 16982
- Hierarchy Level | 16982
- Description | 16983
- Options | 16983
- Required Privilege Level | 16983
- Release Information | 16983

Syntax

```
short-cycle-protection <lockout-max-time seconds> <lockout-min-time seconds>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name]
[edit logical-systems name forwarding-options dhcp-relay ...],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay ...],
[edit routing-instances name forwarding-options dhcp-relay ...],
[edit logical-systems name system services dhcp-local-server ...],
[edit logical-systems name routing-instances name system services dhcp-local-server dhcp-local-
server...],
[edit routing-instances name system services dhcp-local-server ...],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],
```

```
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server group group-name interface interface-name]
```

Description

Enable DHCP short-cycle protection to reduce resource usage associated with connection and authentication processing in highly scaled networks. You must configure both the minimum duration and the maximum duration for the lockout period.

The router detects short-lived client sessions and clients that repeatedly fail session negotiation, then locks them out from access by dropping subsequent DHCP discover or solicit messages from the client. The clients are tracked by the client identifier (client key), which can be a MAC address or some other unique value for DHCPv4 clients or the DUID for DHCPv6 clients. Locked-out clients are entered in the lockout database. If a locked-out client attempts another session before the grace time threshold is reached, it is locked out again. Each successive lockout period is increased exponentially up to the maximum lockout period. The grace time threshold is automatically set at whichever value is larger, 900 seconds or the configured maximum value.

Options

lockout-max-time *seconds*

Maximum length of any lockout period; the upper bound of the lockout range.

- **Range:** 1 through 86400

lockout-min-time *seconds*

Minimum length of any lockout period; the lower bound of the lockout period. The minimum value is the length of the first lockout period for a client. It cannot be greater than the maximum value. If you set it to the same value as the maximum, then the lockout period is fixed and does not increase for a client's subsequent lockouts.

- **Range:** 1 through 86400

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Configuring DHCP Short-Cycle Protection

DHCP Short Cycle Protection Against Frequent Brief or Failed Client Sessions

short-cycle-protection (Static and Dynamic Subscribers)

IN THIS SECTION

- Syntax | 16984
- Hierarchy Level | 16984
- Description | 16985
- Options | 16985
- Required Privilege Level | 16985
- Release Information | 16986

Syntax

```
short-cycle-protection <lockout-time-min minimum-seconds> <lockout-time-max maximum-seconds>
<filter [aci]> ;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number pppoe-
underlying-options],
[edit interfaces demux0 unit logical-unit-number family pppoe]
[edit interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family pppoe],
```

```
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
pppoe-underlying-options]
```

Description

Configure the router to temporarily prevent (lock out) a failed or short-lived (also known as short-cycle) PPPoE subscriber session from reconnecting for a default or configurable period of time. You can optionally override the default lockout time, 1 through 300 seconds (5 minutes), by specifying the minimum lockout time and maximum lockout time as part of the `short-cycle-protection` statement. You can optionally specify the lockout based on the ACI, which locks out all PPPoE subscriber sessions that come from the same household and share the same ACI string.

You can configure PPPoE subscriber session lockout, also known as short-cycle protection, for VLAN, VLAN demux, and PPPoE-over-ATM dynamic subscriber interfaces. Enabling PPPoE subscriber session lockout reduces excessive loading on the router, prevents failed or short-lived sessions from disrupting other sessions on the same underlying interface, and preserves valuable system resources.

Options

`filter aci`—(Optional) Use the agent circuit identifier (ACI) lockout for all subscriber sessions.

`lockout-time-min` *`minimum-seconds`*—(Optional) Use the specified minimum lockout time for failed or short-lived PPPoE subscriber sessions. The *minimum-seconds* value must be less than or equal to the *maximum-seconds* value. Setting *minimum-seconds* and *maximum-seconds* to the same value causes the lockout time to become fixed at that value.

- **Range:** 1 through 86400 (24 hours)
- **Default:** 1

`lockout-time-max` *`maximum-seconds`*—(Optional) Use the specified maximum lockout time for failed or short-lived PPPoE subscriber sessions. The *maximum-seconds* value must be equal to or greater than the *minimum-seconds* value. Setting *maximum-seconds* and *minimum-seconds* to the same value causes the lockout time to become fixed at that value.

- **Range:** 1 through 86400 (24 hours)
- **Default:** 300 (5 minutes)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Lockout of PPPoE Subscriber Sessions

PPPoE Subscriber Session Lockout Overview

Understanding the Lockout Period for PPPoE Subscriber Session Lockout

Configuring Dynamic PPPoE Subscriber Interfaces

Example: Configuring a Static PPPoE Subscriber Interface on a Static Underlying VLAN Demux Interface over Aggregated Ethernet

short-name-format

IN THIS SECTION

- Syntax | **16986**
- Hierarchy Level | **16987**
- Description | **16987**
- Options | **16987**
- Required Privilege Level | **16987**
- Release Information | **16987**

Syntax

```
short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name]
```

Description

Specify the name format of the maintenance association name.

Options

`character-string`—The name is an ASCII character string.

`vlan`—The primary VLAN identifier.

`2octet`—A number in the range 0 through 65,535.

`rfc-2685-vpn-id`—A VPN identifier that complies with RFC 2685.

- **Default:** `character-string`

NOTE: The PTX Series Packet Transport Routers support the `vlan` and `2octet` options only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring Connectivity Fault Management (CFM)

short-reach-mode

IN THIS SECTION

- Syntax | 16988
- Hierarchy Level | 16988
- Description | 16988
- Default | 16988
- Options | 16989
- Required Privilege Level | 16989
- Release Information | 16989

Syntax

```
short-reach-mode (enable | disable);
```

Hierarchy Level

```
[edit chassis fpc fpc-slot pic pic-slot],
[edit chassis fpc fpc-slot pic pic-slot port-range port-range-low port-range-high]
```

Description

Configure short-reach mode for individual as well as a range of copper-based 10-Gigabit Ethernet interfaces using short cable lengths (less than 10m) on the QFX5100-48T switch. Short-reach mode reduces power consumption up to 5W on these interfaces.

Default

This feature is disabled by default.

Options

The following options are available:

- `enable`
- `disable`

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D30.

RELATED DOCUMENTATION

Physical Interface Properties

short-sequence

IN THIS SECTION

- Syntax | 16990
- Hierarchy Level | 16990
- Description | 16990
- Default | 16990
- Required Privilege Level | 16990
- Release Information | 16990

Syntax

```
short-sequence;
```

Hierarchy Level

```
[edit interfaces (ls-fpc/pic/port | ml-fpc/pic/port) unit logical-unit-number],
[edit logical-systems logical-system-name interfaces (ls-fpc/pic/port | ml-fpc/pic/port) unit
logical-unit-number]
```

Description

For multilink interfaces only, set the length of the packet sequence identification number to 12 bits.

Default

If not included in the configuration, the length is set to 24 bits.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces

shortcuts (Protocols OSPF)

IN THIS SECTION

- Syntax | 16991
- Hierarchy Level | 16991
- Description | 16991
- Options | 16992
- Required Privilege Level | 16992
- Release Information | 16992

Syntax

```
shortcuts {
    lsp-metric-into-summary;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) traffic-engineering],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) traffic-engineering],
[edit protocols (ospf | ospf3) traffic-engineering],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)traffic-engineering]
```

Description

Configure OSPF to use MPLS label-switched paths (LSPs) as shortcut next hops. By default, shortcut routes calculated through OSPFv2 are installed in the inet.3 routing table, and shortcut routes calculated through OSPFv3 are installed in the inet6.3 routing table.

Options

lsp-metric-into-summary Advertise the LSP metric in summary LSAs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4.

Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4 for EX Series switches.

RELATED DOCUMENTATION

Example: Enabling OSPF Traffic Engineering Support

should-secure

IN THIS SECTION

- Syntax | 16993
- Hierarchy Level | 16993
- Description | 16993
- Default | 16993
- Required Privilege Level | 16993
- Release Information | 16994

Syntax

```
should-secure;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name mka]
```

Description

Configure the `should-secure` statement to enable fail open mode for Media Access Control Security (MACsec). Fail open mode ensures that traffic continues to flow if the MACsec Key Agreement (MKA) session is not established.

When `should-secure` is configured:

- If the MKA session is not established, traffic is still allowed in clear text without the MACsec header.
- If the MKA session is established successfully, traffic is allowed with MACsec headers.

Fail open mode is recommended for service providers that prioritize network availability. Before you enable fail open mode, consider the sensitivity of the traffic you are attempting to secure with MACsec.

> **NOTE**: Setting or unsetting this configuration statement will cause the tunnel to be re-established.

Default

Fail open mode is not enabled by default. You must configure `should-secure` to enable this feature. When `should-secure` is not configured:

- If the MKA session is not established, all traffic is discarded except Extensible Authentication Protocol over LAN (EAPoL) packets.
- If the MKA session is established, traffic is transmitted with MACsec headers.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4.

show

IN THIS SECTION

- Syntax | 16994
- Description | 16994
- Options | 16994
- Required Privilege Level | 16995
- Release Information | 16995

Syntax

```
show <statement-path> <identifier>
```

Description

Display the current configuration.

Options

none—Display the entire configuration at the current hierarchy level.

identifier—(Optional) Display the configuration for the specified identifier.

statement-path—(Optional) Display the configuration for the specified statement hierarchy path.

Required Privilege Level

configure—To enter configuration mode. In addition to configure, to display the contents under a configuration hierarchy, user should have the required privilege level to view the configuration under that configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show | display inheritance | 19529
- show | display omit | 19535
- show | display set | 19538
- show | display set relative | 19540
- show groups junos-defaults | 20120
- *Displaying the Current Configuration*

show access-cac interface-set

IN THIS SECTION

- Syntax | 16996
- Description | 16996
- Options | 16996
- Required Privilege Level | 16996
- Output Fields | 16996
- Sample Output | 16997
- Release Information | 16998

Syntax

```
show access-cac interface-set
<interface-set-name | detail>
```

Description

Display interface-set adjustment information.

Options

none	List all interface-set entries with the total aggregate adjustment value for each interface-set.
interface-set-name	Restrict the output to a specific interface-set.
detail	List all multicast groups, with adjustment rates, that provide an update on the specified interface-set.

Required Privilege Level

view

Output Fields

Table 142 on page 16996 describes the output fields for the `show show access-cac interface-set` command. Output fields are listed in the approximate order in which they appear.

Table 142: show class-of-service interface-set Output Fields

Field Name	Field Description
`Access cac interface-set`	Then name of the interface set.

Table 142: show class-of-service interface-set Output Fields *(Continued)*

Field Name	Field Description
`Total current adjustment on interface-set`	The total current adjustment on the interface set, in bps, equal to the sum of the measured rates on all configured multicast channels.
`Mcast channels`	The list of all configured multicast channels.
`Configured rate`	The shaping rate configured for each multicast channel.
`Measured rate`	The current data rate for each multicast channel. This rate for a multicast channel adjusts the shaping rate of the interface set by an equal amount.

Sample Output

show access-cac interface-set

```
user@host> show access-cac interface-set
Access cac interface-sets
  Access cac interface-set: ge-5/0/0-1
  Access cac interface-set: ge-5/0/1-1
```

show access-cac interface-set (detail)

```
user@host> show access-cac interface-set detail
Access cac interface-sets
  Access cac interface-set: ge-5/0/0-1
  Total current adjustment on interface-set: 4080908
  Mcast channels:
  0.0.0.0     , 225.0.0.1    Configured rate: 4000000 Measured rate: 560268
  0.0.0.0     , 225.0.0.2    Configured rate: 4000000 Measured rate: 560280
  0.0.0.0     , 225.0.0.3    Configured rate: 4000000 Measured rate: 560280
  ::          , ff8e::e100:1 Configured rate: 4000000 Measured rate: 2400080

  Access cac interface-set: ge-5/0/1-1
```

```
Total current adjustment on interface-set: 500222379
Mcast channels:
0.0.0.0     , 225.0.0.1    Configured rate: 4000000 Measured rate: 166739966
0.0.0.0     , 225.0.0.2    Configured rate: 4000000 Measured rate: 166741131
0.0.0.0     , 225.0.0.3    Configured rate: 4000000 Measured rate: 166741282
::          , ff8e::e100:1 Configured rate: 4000000 Measured rate: 0
```

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Using Hierarchical CoS to Adjust Shaping Rates Based on Multicast Traffic

show access-security router-advertisement state

IN THIS SECTION

- Syntax | **16998**
- Description | **16999**
- Options | **16999**
- Required Privilege Level | **16999**
- Output Fields | **16999**
- Sample Output | **17000**
- Release Information | **17000**

Syntax

```
show access-security router-advertisement state
<interface interface-name>
```

Description

Display the IPv6 Router Advertisement (RA) guard state information. Stateful RA guard enables the switch to learn about the sources of RA messages for a certain period of time. When the learning period ends, the switch has a record of which interfaces are attached to links with valid IPv6 routers. If there is no valid IPv6 router attached to the interface, the switch dynamically transitions the interface from the learning state into the blocking state. Subsequent RA messages received after the transition to blocking state are dropped. If there is a valid IPv6 router attached to the interface, the interface transitions to the forwarding state, and subsequent RA messages that can be validated against the configured policy are forwarded.

Options

interface *interface-name* (Optional) Display the RA guard entries for the specified interface.

Required Privilege Level

view

Output Fields

Table 143 on page 16999 lists the output fields for the `show access-security router-advertisement state` command. Output fields are listed in the approximate order in which they appear.

Table 143: show access-security router-advertisement state Output Fields

Field Name	Field Description
`Interface`	Displays the interface on which stateful IPv6 RA guard is enabled.

Table 143: show access-security router-advertisement state Output Fields *(Continued)*

Field Name	Field Description
State	Displays one of the following states: • OFF—The interface operates as if RA guard is not available. • BLOCKED—The interface blocks ingress RA messages. • FORWARDING—The interface forwards ingress RA messages that can be validated against the configured policy. • LEARNING—The switch is actively acquiring information about the IPv6 routing device connected to the interface. • TRUSTED—The interface forwards all ingress RA messages without performing policy checks.

Sample Output

show access-security router-advertisement state

```
user@device> show access-security router-advertisement state
Interface        state
ge-0/0/0.0     LEARNING
ge-1/0/0.0     FORWARDING
ge-1/0/0.0     BLOCKED
```

Release Information

Command introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

show access-security router-advertisement statistics | 17001

show access-security router-advertisement statistics

IN THIS SECTION

- Syntax | 17001
- Description | 17001
- Options | 17001
- Required Privilege Level | 17001
- Output Fields | 17002
- Sample Output | 17002
- Release Information | 17002

Syntax

```
show access-security router-advertisement statistics
<interface interface-name>
```

Description

Display the IPv6 Router Advertisement (RA) guard entries for received RA messages. RA guard enables a switch to examine incoming RA messages and filter them on the basis of predefined set of criteria. Once the switch has validated that the sender of the RA message is a legitimate IPv6 router, it forwards the RA message to its destination; otherwise, the RA message is dropped.

Options

interface *interface-name* (Optional) Display the RA guard entries for the specified interface.

Required Privilege Level

view

Output Fields

Table 144 on page 17002 lists the output fields for the `show access-security router-advertisement statistics` command. Output fields are listed in the approximate order in which they appear.

Table 144: show access-security router-advertisement statistics Output Fields

Field Name	Field Description	Level of Output
`Interface`	Interface on which the RA packet was received.	All levels
`RA Packets`	Total number of RA packets that were received.	All levels
`RA inspection pass`	Total number of RA packets that passed RA guard inspection.	All levels
`RA inspection fail`	Total number of RA packets that failed RA guard inspection.	All levels

Sample Output

show access-security router-advertisement statistics

```
user@device> show access-security router-advertisement statistics
Interface      RA Packets received    RA inspection pass       RA inspection fail
ge-0/0/7.0     3                      2                        1
ge-0/0/15.0    8                      5                        3
```

Release Information

Command introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

clear access-security router-advertisement statistics | 3403

show access-security slaac-snooping binding

IN THIS SECTION

- Syntax | 17003
- Description | 17003
- Options | 17003
- Required Privilege Level | 17004
- Output Fields | 17004
- Sample Output | 17005
- Release Information | 17005

Syntax

```
show access-security slaac-snooping binding
<interface (interface-name | all>
<vlan vlan-name>
<vlan vlan-name routing-instance routing-instance-name>
<vlan vlan-name logical-system logical-system-name>
```

Description

Display the snooping table entries for IPv6 clients using stateless address auto-configuration (SLAAC). SLAAC snooping extracts address information from Duplicate Address Detection (DAD) packets exchanged during the SLAAC process to build the SLAAC snooping table. The IPv6-MAC address bindings in this table are used to inspect and validate NDP/IP packets sent by IPv6 clients using SLAAC.

Options

`interface (interface-name | all)`	Display the SLAAC snooping entries for the specified interface or for all interfaces.
`vlan vlan-name`	Display the SLAAC snooping entries for the specified VLAN.
`logical-system logical-system-name`	Display the SLAAC snooping entries for the specified logical system.

routing-instance *routing-instance-name*	Display the SLAAC snooping entries for the specified routing instance.

Required Privilege Level

view

Output Fields

Table 145 on page 17004 lists the output fields for the `show access-security slaac-snooping binding` command. Output fields are listed in the approximate order in which they appear.

Table 145: show access-security slaac-snooping binding Output Fields

Field Name	Field Description
`IPv6 address`	IPv6 address of the network device; bound to the MAC address.
`MAC address`	MAC address of the network device; bound to the IPv6 address.
`Vlan`	VLAN name of the network device whose MAC address is shown.
`Expires`	Length of time remaining until the address binding entry expires.
`State`	State of the address binding: • `INIT` • `WAIT` • `BOUND` • `RENEWING` • `AUTO DAD`
`Interface`	Interface address (port).

Sample Output

show access-security slaac-snooping binding

```
user@device> show access-security slaac-snooping binding
IPv6 address                MAC address          Vlan     Expires     State     Interface
2001:db8:fe10::             00:00:01:00:00:03    vlan1    66011       BOUND     xe-0/2/2.0
2001:db8:fe12::             00:00:01:00:00:04    vlan1    78938       BOUND     xe-0/2/0.0
2001:db8:fe14::             00:00:01:00:00:05    vlan1    77946       BOUND     xe-2/2/0.0
2001:db8:fe16::             00:00:01:00:00:06    vlan1    2584539     BOUND     xe-0/2/0.0
2001:db8:fe18::             00:00:01:00:00:07    vlan1    2583254     BOUND     xe-0/2/2.0
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear access-security slaac-snooping binding | 3404

show access-security slaac-snooping state

IN THIS SECTION

- Syntax | 17006
- Description | 17006
- Options | 17006
- Required Privilege Level | 17006
- Output Fields | 17006
- Sample Output | 17007
- Release Information | 17007

Syntax

```
show access-security slaac-snooping state
<interface interface-name I all>
```

Description

Display the SLAAC snooping state information to show interfaces that are blocked. SLAAC snooping extracts address information from Duplicate Address Detection (DAD) packets exchanged during the SLAAC process to build the SLAAC snooping table. The IPv6-MAC address bindings in this table are used to inspect and validate NDP/IP packets sent by IPv6 clients using SLAAC.

To restore an interface that has been blocked, you must issue the `request access-security slaac-snooping unblock` command.

Options

interface (*interface-name* | all)	Display the SLAAC snooping state for the specified interface.

Required Privilege Level

view

Output Fields

Table 146 on page 17006 lists the output fields for the `show access-security slaac-snooping state` command. Output fields are listed in the approximate order in which they appear.

Table 146: show access-security slaac-snooping state Output Fields

Field Name	Field Description
`Interface`	Displays the interface on which SLAAC snooping is enabled.
`State`	Shows that the interface is blocked, meaning incoming DAD and ND requests on the interface will be dropped.

Sample Output

show access-security slaac-snooping state

```
user@device> show access-security slaac-snooping state
Interface         State
ge-0/0/40.0     BLOCKED
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

request access-security slaac-snooping unblock | 14417

show access-security slaac-snooping statistics

IN THIS SECTION

- Syntax | 17008
- Description | 17008
- Options | 17008
- Required Privilege Level | 17008
- Output Fields | 17008
- Sample Output | 17010
- Release Information | 17010

Syntax

```
show access-security slaac-snooping statistics
<interface (interface-name | all>
```

Description

Display the statistics for IPv6 clients using stateless address auto-configuration (SLAAC). SLAAC snooping extracts address information from Duplicate Address Detection (DAD) packets exchanged during the SLAAC process to build the SLAAC snooping table. The IPv6-MAC address bindings in this table are used to inspect and validate NDP/IP packets sent by IPv6 clients using SLAAC.

Options

interface (*interface-name* | all)	Display the SLAAC snooping statistics for the specified interface or for all interfaces.
logical-system *logical-system-name*	Display the SLAAC snooping statistics for the specified logical system.
routing-instance *routing-instance-name*	Display the SLAAC snooping statistics for the specified routing instance.

Required Privilege Level

view

Output Fields

Table 147 on page 17009 lists the output fields for the `show access-security slaac-snooping statistics` command. Output fields are listed in the approximate order in which they appear.

Table 147: show access-security slaac-snooping statistics Output Fields

Field Name	Field Description
`Statistics`	Statistics for SLAAC snooping. • `DAD request queued`—Number of Duplicate Address Detection (DAD) requests queued. • `Confirmed bindings`—Number of SLAAC snooping address bindings in confirmed state. • `Decline bindings`—Number of SLAAC snooping address bindings declined. • `Conflicting DAD entries`—Number of address conflicts found by DAD. • `NS request queued`—Number of Neighbor Solicitation requests queued. • `AUTO DAD packet tx`—Number of auto-DAD packets transmitted.
`Packets dropped`	Number of packets not considered for SLAAC snooping because of errors. • `Total`—Total number of packets discarded by SLAAC snooping. • `Bad read`—Number of packets discarded because they could not be read. • `No configuration`—Number of packets discarded because they did not have a valid configuration. • `No VLAN`—Number of packets discarded because they did not belong to a valid VLAN. • `No interface`—Number of packets discarded because they did not belong to a valid interface. • `Request rx on trusted port`—Number of packets discarded because a Request message was received on a trusted port.
`Interface statistics (interface detail)`	Statistics for the specified interface. • `Interface name`—Name of the interface. • `NA-PKT-RX`—Number of Neighbor Advertisement requests received. • `NA-PKT-TX`—Number of Neighbor Advertisement requests transmitted.

Sample Output

show access-security slaac-snooping statistics

```
user@device> show access-security slaac-snooping statistics
Statistics:
---------------
DAD request queued                      4
Confirmed bindings                      4
Decline bindings                        0
Conflicting DAD entries                 0
NS request queued                       2
AUTO DAD packet tx                      2

Packets dropped:
---------------
Total                                   4
Bad Read                                0
No configuration                        0
No VLAN                                 0
No interface                            0
Request rx on Trusted port              0
```

show access-security slaac-snooping statistics interface ge-0/0/40

```
user@device> show access-security slaac-snooping statistics interface ge-0/0/40
    Interface Statistics:
    ---------------------
    Interface Name             NA-PKT-RX          NA-PKT-TX
    ge-0/0/35.0                0                  0
    ge-0/0/41.0                0                  0
    ge-0/0/40.0                6                  0
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show access-security slaac-snooping binding | 17003

show accounting pending-accounting-stops

IN THIS SECTION

- Syntax | 17011
- Description | 17011
- Options | 17011
- Required Privilege Level | 17012
- Output Fields | 17012
- Sample Output | 17014
- Release Information | 17016

Syntax

```
show accounting pending-accounting-stops
<detail | terse>
<profile-name>
```

Description

Display all statistics for all pending accounting stop requests, including both service and session requests.

Options

none Display information for all access profiles.

detail | terse (Optional) Display the specified level of output.

profile-name (Optional) Particular access profile for which you want to display accounting stop statistics.

Required Privilege Level

view

Output Fields

Table 148 on page 17012 lists the output fields for the `show accounting pending-accounting-stops` command. Output fields are listed in the approximate order in which they appear.

Table 148: show accounting pending-accounting-stops Output Fields

Field Name	Field Description	Level of Output
`Type`	Type of client.	All levels
`Username`	Name of the user logged in to the session.	All levels
`Logical system/Routing instance`	Logical system and routing instance used for the session.	`detail` none
`Access-profile`	Access profile used for AAA services for the session.	`detail` none
`Session ID`	ID of the subscriber session; generated when the subscriber logs in. In the `Service name` block, this is the ID of the service session.	All levels
`Accounting Session ID`	ID of the accounting session (RADIUS attribute 44). The ID appears in decimal or description format, as specified by the `accounting-session-id-format` statement.	`detail` none
`IP Address`	IP address of the subscriber.	`detail` none
`IPv6 Prefix`	IPv6 address of the subscriber.	`detail` none

Table 148: show accounting pending-accounting-stops Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Authentication State`	State of the subscriber authentication session: `AuthInit`, `AuthStart`, `AuthChallenge`, `AuthRedirect`, `AuthClntRespWait`, `AuthAcctVolStatsAckWait`, `AuthAcctStopAckWait`, `AuthServCreateRespWait`, `AuthLogoutStart`, `AuthStateActive`, `AuthClntLogoutRespWait`, `AuthProfileUpdateWait`, `AuthProvisionRespWait`, `AuthProvisionServiceCreationWait`	`detail` none
`Accounting State`	State of the subscriber accounting session: `Acc-Init`, `Acc-Start-Sent`, `Imm-Update-Stats-Pending`, `Acc-Interim-Sent`, `Acc-Stop-Stats-Pending`, `Acc-Stop-Sent`, `Acc-Stop-On-Fail-Deny-Sent`, `Acc-Stop-Ackd`	`detail` none
`Service name`	Name of the attached service or policy.	`detail` none
`Service State`	State of the service provided in the subscriber session.	`detail` none
`Session uptime`	How long the session has been up, in *HH:MM:SS*.	`detail` none
`Accounting status`	Status of the accounting configuration for the service, `on` or `off`, and the type of accounting, `time` or `volume+time`. Configured in RADIUS Service-Statistics VSA [26-69].	`detail` none
`Service accounting session ID`	ID of the service accounting session; RADIUS Acct-Session-Id attribute (44). The ID appears in decimal or description format, as specified by the `accounting-session-id-format` statement.	`detail` none
`Service accounting state`	State of the service accounting session: `Acc-Init`, `Acc-Start-Sent`, `Imm-Update-Stats-Pending`, `Acc-Interim-Sent`, `Acc-Stop-Stats-Pending`, `Acc-Stop-Sent`, `Acc-Stop-On-Fail-Deny-Sent`, `Acc-Stop-Ackd`	`detail` none

Table 148: show accounting pending-accounting-stops Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting interim interval`	Amount of time between interim accounting updates for this service, in seconds; RADIUS Service-Interim-Acct-Interval VSA [26-140] or Diameter Acct-Interim-Interval AVP (85).	`detail` none
`Subscriber ID`	ID of the subscriber; generated when the subscriber logs in.	`detail` none
`Service ID`	ID of the subscriber service.	All levels
`Service`	Name of the attached service or policy.	`terse`

Sample Output

show accounting pending-accounting-stops detail

```
user@host> show accounting pending-accounting-stops detail
Type: pppoe
Username: vjshah29@example.com
AAA Logical system/Routing instance: default:default
Access-profile: ce-ppp-profile
Session ID: 84
Accounting Session ID: 84
IP Address: 192.168.0.25
IPv6 Prefix: 2010:db8:9999:18::/48
Authentication State: AuthAcctStopAckWait
Accounting State: Acc-Stop-Stats-Pending
Service name: cos-service
  Service State: SvcInactive
  Session ID: 94
  Session uptime: 00:08:02
  Accounting status: on/time
  Service accounting session ID: 84:94-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600
Service name: filter-service
```

```
  Service State: SvcInactive
  Session ID: 93
  Session uptime: 00:08:02
  Accounting status: on/volume+time
  Service accounting session ID: 84:93-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600
Service name: filter-service6
  Service State: SvcInactive
  Session ID: 95
  Session uptime: 00:08:02
  Accounting status: on/volume+time
  Service accounting session ID: 84:95-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600
```

show accounting pending-accounting-stops (Specific Profile)

```
user@host> show accounting pending-accounting-stops ce-ppp-profile
Type:         Username:                Session ID:     Service ID:       Service
  pppoe        vjshah29@example.com       84
  pppoe        vjshah29@example.com       84              94                cos-service
  pppoe        vjshah29@example.com       84              93                filter-service
  pppoe        vjshah29@example.com       84              95                filter-service6
```

show accounting pending-accounting-stops terse

```
user@host> show accounting pending-accounting-stops terse
 Type:         Username:                   Session ID:     Service ID:       Service
  pppoe        vjshah29@example.com           84
  pppoe        vjshah29@example.com           84              94                cos-service
  pppoe        vjshah29@example.com           84              93                filter-service
  pppoe        vjshah29@example.com           84              95                filter-service6
  pppoe        larry@example.com              85
  pppoe        larry@example.com              85              94                cos-service
  pppoe        larry@example.com              85              93                filter-service
  pppoe        larry@example.com              85              95                filter-service6
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

request network-access aaa replay pending-accounting-stops | 14695

show network-access aaa statistics pending-accounting-stops | 22455

show accounting profile

IN THIS SECTION

- Syntax | 17016
- Description | 17016
- Options | 17017
- Required Privilege Level | 17017
- Output Fields | 17017
- Sample Output | 17020
- Release Information | 17023

Syntax

```
show accounting profile profile-name
```

Description

Display accounting profile information.

Options

profile-name	Name of the accounting profile.

Required Privilege Level

view

Output Fields

Table 149 on page 17017 lists the output fields for the `show accounting profile` command. Output fields are listed in the approximate order in which they appear.

Table 149: show accounting profile Output Fields

Field Name	Field Description
Profile	Name of the accounting profile.
Sampling interval	Configured interval, in minutes, for statistic collection.
Profile Usage Count	Number of items configured for collecting accounting statistics.
file information	Information about the accounting profile log, including: • **File**—Name of accounting profile log. If no name is explicitly provided, the name of the accounting profile is used. All statistics files are placed in the **/var/log** directory. • **maximum size**—Configured size. When the size is exceeded, the log file closes and a new log file opens. • **maximum number**—Configured maximum number of log files. • **bytes written**—Number of bytes written to the log file.
Transfer Interval	Length of time (in minutes) the file remains open, receiving statistics before it is closed, transferred, and rotated. When either the time or the file size is exceeded, the file is closed and a new one opened, whether or not a transfer site is specified.

Table 149: show accounting profile Output Fields *(Continued)*

Field Name	Field Description
Next Scheduled Transfer	Time at which the next transfer occurs.

Table 149: show accounting profile Output Fields *(Continued)*

Field Name	Field Description
Column Labels	Names of sampled statistics. This list varies depending on the configuration: • **profile-layout**—List of data fields reported, in the order they appear in the output. • **epoch-timestamp**—Number of seconds since the epoch. • **interfaces**—(For interface, filter, and destination class profiles) Name of the interfaces on which the filter is applied. • **filter-name**—(For filter profiles) Name of the filter. • **counter-name**—(For filter profiles) Name of the counter. • **packet-count**—(For filter and destination class profiles) Number of packets for the counter. • **byte-count**—(For filter and destination class profiles) Number of bytes for the counter. • **input-bytes**—(For interface profiles) Input bytes. • **input-errors**—(For interface profiles) Generic input error packets. • **input-multicast**—(For interface profiles) Input packets arriving by multicast. • **input-packets**—(For interface profiles) Input packets. • **input-unicast**—(For interface profiles) Input unicast packets. • **output-bytes**—(For interface profiles) Output bytes. • **output-errors**—(For interface profiles) Generic output error packets. • **output-multicast**—(For interface profiles) Output packets sent by multicast. • **output-packets**—(For interface profiles) Output packets. • **output-unicast**—(For interface profiles) Output unicast packets. • **no-proto**—(For interface profiles) Packets for unsupported protocol. • **snmp-index**—(For interface profiles) SNMP index. • **destination-class-name**—(For destination class profiles) Configured destination class name.

Table 149: show accounting profile Output Fields *(Continued)*

Field Name	Field Description
	• **host name**—(For Routing Engine profiles) Hostname for the router. • **date-*yyyymmdd***—(For Routing Engine profiles) Date. • **timeofday-*hhmmss***—(For Routing Engine profiles) Time of day. • **uptime**—(For Routing Engine profiles) Time since the last reboot, in seconds. • **cpu1min**—(For Routing Engine profiles) Average system load over the last 1 minute. • **cpu5min**—(For Routing Engine profiles) Average system load over the last 5 minutes. • **cpu15min**—(For Routing Engine profiles) Average system load over the last 15 minutes.
Interface name	Name of the interface configured for this accounting profile.
Filter name	Name of the filter configured for this accounting profile.
routing-engine-stats	Routing Engine accounting profile.
Next Scheduled Collection	Time for next collection of statistics for the named interface.

Sample Output

show accounting profile (Interface)

```
user@host> show accounting profile if_prof
Profile if_prof
Sampling interval: 1 minute(s),  Profile Usage Count: 2
File accounting_profile_stats: maximum size 1048576, maximum number 5, bytes written 2196
Transfer Interval: 15 minute(s),  Next Scheduled Transfer: 2001-06-17-18:00:45
Column Labels:
  profile-layout
  epoch-timestamp
  interface-name
```

```
  snmp-index
  input-bytes
  output-bytes
  input-packets
  output-packets
  input-unicast
  output-unicast
  input-multicast
  output-multicast
  no-proto
  input-errors
  output-errors


Interface Name              Next Scheduled Collection
fxp0.0                      2001-06-18-18:00:30
fxp0                        2001-06-18-18:01:00
```

show accounting profile (Filter)

```
user@host> show accounting profile filter_profile
Profile filter_profile
Sampling interval: 1 minute(s),  Profile Usage Count: 0
File accounting_profile_stats: maximum size 1048576, maximum number 5, bytes written 822
Transfer Interval: 15 minute(s),  Next Scheduled Transfer: 2001-06-17-18:00:46
Column Labels:
  profile-layout
  epoch-timestamp
  interfaces
  filter-name
  counter-name
  packet-count
  byte-count


Filter Name                 Next Scheduled Collection
myfiltero                         2001-06-03-04:32:59
```

show accounting profile (Destination Class)

```
user@host> show accounting profile dcu1
Profile dcu1
Sampling interval: 1 minute(s),  Profile Usage Count: 0
File accounting_profile_stats: maximum size 1048576, maximum number 5, bytes written 901
Transfer Interval: 15 minute(s),  Next Scheduled Transfer: 2001-06-17-18:00:46
Column Labels:
  profile-layout
  epoch-timestamp
  interface-name
  destination-class-name
  packet-count
  byte-count


Interface Name              Next Scheduled Collection
so-0/3/3                    2001-06-03-04:34:00
```

show accounting profile (Routing Engine)

```
user@host> show accounting profile rep1
Profile rep1
Sampling interval: 1 minute(s),  Profile Usage Count: 1
File accounting_profile_stats: maximum size 1048576, maximum number 5, bytes written 901
Transfer Interval: 15 minute(s),  Next Scheduled Transfer: 2001-06-17-18:00:46
Column Labels:
  profile-layout
  epoch-timestamp
  hostname
  date-yyyymmdd
  timeofday-hhmmss
  uptime
  cpu1min
  cpu5min
  cpu15min


Interface Name              Next Scheduled Collection
routing-engine-stats        2001-06-18-18:02:31
```

Release Information

Command introduced before Junos OS Release 7.4.

show accounting server statistics archival-transfer

IN THIS SECTION

- Syntax | 17023
- Description | 17023
- Options | 17023
- Required Privilege Level | 17023
- Sample Output | 17024
- Release Information | 17024

Syntax

```
show accounting server statistics archival-transfer
```

Description

Display the statistics of transfer attempted, succeeded, and failed for accounting statistics files and router configuration archives.

Options

This command has no options.

Required Privilege Level

view

Sample Output

show accounting server statistics archival-transfer

```
user@host> show accounting server statistics archival-transfer
                File Name : /var/transfer/config/*
URL                                     :scp://root@ce-bras-nvsrv-a:/var/www/html
Last transfer attempted timestamp       :20190603_143642
Last successful transfer timestamp      :20190603_143642
Success Count                           : 5
Failure Count                           : 0
```

Release Information

Command introduced in Junos OS Release 19.2.

show agent sensors

IN THIS SECTION

- Syntax | 17024
- Description | 17025
- Required Privilege Level | 17025
- Output Fields | 17025
- Sample Output | 17027
- Release Information | 17032

Syntax

```
show agent sensors
```

Description

Display information about sensors configured for Junos Telemetry Interface.

> **NOTE**: Junos telemetry interface was introduced in Junos OS Release 15.1F3 on MX Series routers with interfaces configured on MPC1 through MPC6E and on PTX Series routers with interfaces configured on FPC3. Starting in Junos OS Release 15.1F5, Junos Telemetry Interface is also supported on MPC7E, MPC8E, and MPC9E on MX Series routers.
>
> Starting with Junos OS Release 16.1R3, FPC1, FPC2, and dual Routing Engines on PTX Series routers are also supported.

Required Privilege Level

view

Output Fields

Table 150 on page 17025 lists the output fields for the `show agent sensors` command. Output fields are listed in the approximate order in which they appear.

Table 150: show agent sensors Output Fields

Field Name	Field Description
`Sensor Information`	Information about sensors configured to monitor system resources and stream data.
`Name`	Name of configured sensor. **NOTE**: Junos OS Evolved Release 19.1R1 and later does not show output for generated child sensors.
`Resource`	Resource string used to configure and identify the system resource enabled to monitor and stream data.
`Sensor-id`	Numerical identifier of the sensor.
`Server Information`	Information about servers configured to collect sensor data.

Table 150: show agent sensors Output Fields *(Continued)*

Field Name	Field Description
`Name`	Name of server.
`Scope-id`	Numerical identifier of a scope.
`Remote-Address`	Destination IP address for exported packets.
`Remote-port`	Destination port for exported packets.
`Profile information`	Information about export profiles for sensors.
`Name`	Name of export profile.
`Rep-interval`	Interval, in seconds, at which the sensor generates data to export.
`Address`	Source address of exported packets.
`Port`	Source port of exported packets.
`Format`	Format of exported data message: `GPB`
`DSCP`	Configured DSCP value for exported packets. **NOTE**: The default value is 0. This value is displayed if you do not configure a DSCP value. **NOTE**: This field is not included with gRPC-based streaming.
`Forwarding-class`	Configured forwarding class for exported packets. **NOTE**: The default value is 0. This value is displayed if you do not configure a forwarding class. **NOTE**: This field is not included with gRPC-based streaming.

Table 150: show agent sensors Output Fields *(Continued)*

Field Name	Field Description
Loss-Priority	Configured loss priority for packets streamed through UDP (MX Series only): high, low, medium-high, medium-low **NOTE**: This field is not included with gRPC-based streaming.

Sample Output

show agent sensors (DDOS sensor)

```
user@host> show agent sensors
Sensor Information Sensor name     :__default_fabric_sensor__
        Resource         :/junos/system/linecard/fabric
        Object ID       :2277069355
        Sensor ID      :9288676508521003
        Polling Interval    :2000000000
Server Information Server name     :__default__snmp_server__
        Server ID :33
        Scope ID :0
        Remote-Address :0.0.0.0
        Remote-port :0
Profile Information: Profile
name     :__default_snmp_export_profile__
        profile ID :1
        Rep-interval :2000
        Local-Address :0.0.0.0
        Local-port :0
        Timestamp :0
        Format :1
        Transport :0
Sensor Information Sensor name :ddos
        Resource :/junos/system/linecard/ddos/
        Object ID :3440566888
        Sensor ID :21110626693866088
        Polling Interval :1000
        Server Information Server
            name :sserver
```

```
    Server ID :75
        Scope ID :0 Remote-
Address :10.0.0.1
Remote-port :2000
Local-Address         :10.0.0.10
        Local-port     :21111
        Timestamp         :1
        Format         :1
        Transport         :0
```

show agent sensors (firewall filter sensor)

```
user@host> show agent sensors
Sensor Information :

  Name                                    :firewall-stats
  Resource                                :/junos/system/linecard/firewall/
  Sensor ID                               :93390914

Server Information :

   Name                                   :jvision-server
   Scope ID                               :0
   Remote-Address                         :160.1.1.1
   Remote-port                            :2001

Profile Information :

  Name                                    :export-common
  Rep-interval                            :2
  Address                                 :160.1.1.2
  Port                                    :1000
  Timestamp                               :1
  Format                                  :GPB
  Transport                               :UDP
  DSCP                                    :0
  Forwarding-class                        :0
  Loss-priority                           :high
```

show agent sensors (CPU memory sensor)

```
user@host> show agent sensors
Sensor Information :

       Name                                : se1
       Resource                            :/junos/system/cpu/memory/
       Version                             : 1.0
       Sensor-id                           : 114833
       Subscription-ID                     : 562949953536145
       Parent-Sensor-Name                  : Not applicable
       Component(s)                        : PFE

 Server Information :

        Name                               : ser1
        Scope-id                           : 0
        Remote-Address                     : 10.3.3.3
        Remote-port                        : 6000
        Transport-protocol                 : UDP

 Profile Information :

        Name                               : ex1
        Reporting-interval                 : 1
        Payload-size                       : 5000
        Address                            : 0.0.0.0
        Port                               : 1000
        Timestamp                          : 1
        Format                             : GPB
        DSCP                               : 0
        Forwarding-class                   : assured-forwarding
        Loss-priority                      : high
```

show agent sensors (packet forwarding engine statistics)

```
user@host> show agent sensors
Sensor Information :

    Name                                   : packet_stats
```

```
    Resource                                : /junos/system/linecard/packet/usage/
    Version                                 : 1.0
    Sensor-id                               : 3699
    Subscription-ID                         : 562949953425011
    Parent-Sensor-Name                      : Not applicable
    Component(s)                            : PFE

    Server Information :
        Name                                : s1
        Scope-id                            : 0
        Remote-Address                      : 10.1.1.2
        Remote-port                         : 1000
        Transport-protocol                  : UDP

    Profile Information :
        Name                                : ep1
        Reporting-interval                  : 1
        Payload-size                        : 5000
        Address                             : 10.1.1.1
        Port                                : 1000
        Timestamp                           : 1
        Format                              : GPB
        DSCP                                : 255
        Forwarding-class                    : 255
```

show agent sensors (QFX10008 or QFX10016 switches with Junos OS Release 17.3R1 and later)

```
user@host> show agent sensors
 Sensor Information :

     Name                                    : sensor_1000
     Resource                                : /interfaces/interface/subinterfaces/
     Version                                 : 1.0
     Sensor-id                               : 539528115
     Subscription-ID                         : 1000
     Parent-Sensor-Name                      : Not applicable
     Component(s)                            : PFE,mib2d,xmlproxyd

     Profile Information :
```

```
        Name                                : export_1000
        Reporting-interval                  : 6
        Payload-size                        : 5000
        Format                              : GPB

Sensor Information :

    Name                                    : sensor_1000_1_1
    Resource                                : /junos/system/linecard/interface/logical/usage/
    Version                                 : 1.1
    Sensor-id                               : 3139259737
    Subscription-ID                         : 1000
    Parent-Sensor-Name                      : sensor_1000
    Component(s)                            : PFE

    Profile Information :

        Name                                : export_1000
        Reporting-interval                  : 6
        Payload-size                        : 5000
        Format                              : GPB

Sensor Information :

    Name                                    : sensor_1000_2_1
    Resource                                : /interfaces/interface/subinterfaces/
    Version                                 : 1.0
    Sensor-id                               : 3139256665
    Subscription-ID                         : 1000
    Parent-Sensor-Name                      : sensor_1000
    Component(s)                            : mib2d

    Profile Information :

        Name                                : export_1000
        Reporting-interval                  : 6
        Payload-size                        : 5000
        Format                              : GPB

Sensor Information :

    Name                                    : sensor_1000_4_1
    Resource                                : /interfaces/interface/subinterfaces/
```

```
     Version                                 : 1.0
     Sensor-id                               : 3139262809
     Subscription-ID                         : 1000
     Parent-Sensor-Name                      : sensor_1000
     Component(s)                            : xmlproxyd

     Profile Information :

         Name                                : export_1000
         Reporting-interval                  : 6
         Payload-size                        : 5000
         Format                              : GPB
```

show agent sensors (Junos OS Evolved Release 19.1R1 and later)

```
user@host> show agent sensors
Sensor Information :

    Name                                    : sensor_1000
    Resource                                : /interfaces/interface[name='re0:mgmt-0']/
    Version                                 : 1.0
    Sensor-id                               : 562949953421313
    Subscription-ID                         : 1000
    Component(s)                            : mib2d, mgmt-ethd

    Profile Information :

        Name                                : export_1000
        Reporting-interval                  : 2
        Payload-size                        : 5000
        Address                             : 0.0.0.0
        Port                                : 1000
        Timestamp                           : ntp
        Format                              : GPB
        DSCP                                : 0
        Forwarding-class                    : 0
```

Release Information

Statement introduced in Junos OS Release 15.1F3

RELATED DOCUMENTATION

export-profile (Junos Telemetry Interface) | 6669
sensor (Junos Telemetry Interface) | 16564
streaming-server (Junos Telemetry Interface) | 28640

show amt statistics

IN THIS SECTION

- Syntax | 17033
- Description | 17033
- Options | 17033
- Required Privilege Level | 17034
- Output Fields | 17034
- Sample Output | 17036
- Release Information | 17036

Syntax

```
show amt statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display information about the Automatic Multicast Tunneling (AMT) protocol tunnel statistics.

Options

none	Display summary information about all AMT Protocol tunnels.
instance *instance-name*	(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 151 on page 17034 describes the output fields for the `show amt statistics` command. Output fields are listed in the approximate order in which they appear.

Table 151: show amt statistics Output Fields

Field Name	Field Description
AMT receive message count	Summary of AMT statistics for messages received on all interfaces. • **AMT relay discovery**—Number of AMT relay discovery messages received. • **AMT membership request**—Number of AMT membership request messages received. • **AMT membership update**—Number of AMT membership update messages received.
AMT send message count	Summary of AMT statistics for messages sent on all interfaces. • **AMT relay advertisement**—Number of AMT relay advertisement messages sent. • **AMT membership query**—Number of AMT membership query messages sent.

Table 151: show amt statistics Output Fields *(Continued)*

Field Name	Field Description
AMT error message count	Summary of AMT statistics for error messages received on all interfaces. • **AMT incomplete packet**—Number of messages received with length errors so severe that further classification could not occur. • **AMT invalid mac**—Number of messages received with an invalid message authentication code (MAC). • **AMT unexpected type**—Number of messages received with an unknown message type specified. • **AMT invalid relay discovery address**—Number of AMT relay discovery messages received with an address other than the configured anycast address. • **AMT invalid membership request address**—Number of AMT membership request messages received with an address other than the configured AMT local address. • **AMT invalid membership update address**—Number of AMT membership update messages received with an address other than the configured AMT local address. • **AMT incomplete relay discovery messages**—Number of AMT relay discovery messages received that are not fully formed. • **AMT incomplete membership request messages**—Number of AMT membership request messages received that are not fully formed. • **AMT incomplete membership update messages**—Number of AMT membership update messages received that are not fully formed. • **AMT no active gateway**—Number of AMT membership update messages received for a tunnel that does not exist for the gateway that sent the message. • **AMT invalid inner header checksum**—Number of AMT membership update messages received with an invalid IP checksum. • **AMT gateways timed out**—Number of gateways that timed out because of inactivity.

Sample Output

show amt statistics

```
user@host> show amt statistics

AMT receive message count
AMT relay advertisement                      :           2
AMT membership request                       :           5
AMT membership update                        :           5

AMT send message count
AMT relay advertisement                      :           2
AMT membership query                         :           5

AMT error message count
AMT incomplete packet                        :           0
AMT invalid mac                              :           0
AMT unexpected type                          :           0
AMT invalid relay discovery address          :           0
AMT invalid membership request address       :           0
AMT invalid membership update address        :           0
AMT incomplete relay discovery messages      :           0
AMT incomplete membership request messages   :           0
AMT incomplete membership update messages    :           0
AMT no active gateway                        :           0
AMT invalid inner header checksum            :           0
AMT gateways timed out                       :           0
```

Release Information

Command introduced in JUNOS Release 10.2.

RELATED DOCUMENTATION

clear amt statistics

show amt summary | 17037

show amt tunnel | 17039

show amt summary

IN THIS SECTION

- Syntax | 17037
- Description | 17037
- Options | 17037
- Required Privilege Level | 17037
- Output Fields | 17038
- Sample Output | 17038
- Release Information | 17038

Syntax

```
show amt summary
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display summary information about the Automatic Multicast Tunneling (AMT) protocol.

Options

none	Display summary information about all AMT protocol instances.
instance *instance-name*	(Optional) Display information for the specified instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 152 on page 17038 describes the output fields for the `show amt summary` command. Output fields are listed in the approximate order in which they appear.

Table 152: show amt summary Output Fields

Field Name	Field Description	Level of Output
AMT anycast prefix	Prefix advertised by unicast routing protocols to route AMT discovery messages to the router from nearby AMT gateways.	All levels
AMT anycast address	Anycast address configured from which the anycast prefix is derived.	All levels
AMT local address	Local unique AMT relay IP address configured. Used to send AMT relay advertisement messages, it is the IP source address of AMT control messages and the source address of the data tunnel encapsulation.	All levels
AMT tunnel limit	Maximum number of AMT tunnels that can be created.	All levels
active tunnels	Number of active AMT tunnel interfaces.	All levels

Sample Output

show amt summary

```
user@host> show amt summary
 AMT anycast prefix : 20.0.0.4/32
 AMT anycast address : 20.0.0.4
 AMT local address : 20.0.0.4
 AMT tunnel limit : 1000, active tunnels : 2
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

clear amt tunnel | 3412

show amt statistics | 17033

show amt tunnel | 17039

show amt tunnel

IN THIS SECTION

- Syntax | 17039
- Description | 17039
- Options | 17040
- Required Privilege Level | 17040
- Output Fields | 17040
- Sample Output | 17042
- Release Information | 17044

Syntax

```
show amt tunnel
<brief | detail>
<gateway-address gateway-ip-address> <port port-number>
<instance instance-name>
<logical-system (all | logical-system-name)>
<tunnel-interface interface-name>
```

Description

Display information about the Automatic Multicast Tunneling (AMT) dynamic tunnels.

Options

none	Display summary information about all AMT protocol instances.
brief | detail	(Optional) Display the specified level of detail.
gateway-address ***gateway-ip-address*** **port** ***port-number***	(Optional) Display information for the specified AMT gateway only. If no port is specified, display information for all AMT gateways with the given IP address.
instance ***instance-name***	(Optional) Display information for the specified instance only.
logical-system (all | ***logical-system-name*****)**	(Optional) Perform this operation on all logical systems or on a particular logical system.
tunnel-interface ***interface-name***	(Optional) Display information for the specified AMT tunnel interface only.

Required Privilege Level

view

Output Fields

Table 153 on page 17040 describes the output fields for the `show amt tunnel` command. Output fields are listed in the approximate order in which they appear.

Table 153: show amt tunnel Output Fields

Field Name	Field Description	Level of Output
AMT gateway address	Address of the AMT gateway that is being connected by the AMT tunnel.	All levels
port	Client port used by the AMT tunnel.	All levels
AMT tunnel interface	Dynamically created AMT logical interfaces used by the AMT tunnel in the format **ud-FPC/PIC/Port.unit**.	All levels

Table 153: show amt tunnel Output Fields *(Continued)*

Field Name	Field Description	Level of Output
AMT tunnel state	State of the AMT tunnel. The state is normally **Active**. • **Active**—The tunnel is active. • **Pending**—The tunnel creation is pending. This is a transient state. • **Down**—The tunnel is in the down state. • **Graceful restart pending**—Graceful restart is in progress. • **Reviving**—The routing protocol daemon or Routing Engine was restarted (not gracefully). The tunnel remains in the reviving state until the AMT gateway sends a control message. When the message is received the tunnel is moved to the **Active** state. If no message is received before the AMT tunnel inactivity timer expires, the tunnel is deleted.	All levels
AMT tunnel inactivity timeout	Number of seconds since the most recent control message was received from an AMT gateway. If no message is received before the AMT tunnel inactivity timer expires, the tunnel is deleted.	All levels
Number of groups	Number of multicast groups using the tunnel.	All levels
Group	Multicast group address or addresses using the tunnel.	**detail**
Include Source	Multicast source address for each IGMPv3 group using the tunnel.	**detail**
AMT message count	Statistics for AMT messages: • **AMT Request**—Number of AMT relay tunnel request messages received. • **AMT membership update**—Number of AMT membership update messages received.	All levels

Sample Output

show amt tunnel

```
user@host> show amt tunnel
AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/1/10.1120256
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 15
 Number of groups : 1

AMT message count:
AMT Request    AMT membership update
2              2
```

show amt tunnel detail

```
user@host> show amt tunnel detail
AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/3/10.1120512
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 62
 Number of groups : 1
  Group: 226.2.3.2

AMT message count:
AMT Request    AMT membership update
2              2

 AMT gateway address : 11.11.11.3, port : 2268
 AMT tunnel interface : ud-5/2/10.1120513
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 214
 Number of groups : 1
  Group: 226.2.3.3

AMT message count:
AMT Request    AMT membership update
2              2
```

show amt tunnel tunnel-interface

```
user@host> show amt tunnel tunnel-interface ud-5/3/10.1120512
 AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/3/10.1120512
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 145
 Number of groups : 1

AMT message count:
AMT Request    AMT membership update
2              2
```

show amt tunnel gateway-address

```
user@host> show amt tunnel gateway-address 11.11.11.3 port 2268
AMT gateway address : 11.11.11.3, port : 2268
 AMT tunnel interface : ud-5/2/10.1120513
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 214
 Number of groups : 1
  Group: 226.2.3.3

AMT message count:
AMT Request    AMT membership update
2              2
```

show amt tunnel gateway-address detail

```
user@host> show amt tunnel gateway-address 11.11.11.2 detail
AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/3/10.1120512
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 234
 Number of groups : 1
  Group: 226.2.3.2

AMT message count:
```

```
AMT Request     AMT membership update
2               2
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

clear amt tunnel | 3412
show amt statistics | 17033
show amt summary | 17037

show analyzer

IN THIS SECTION

- Syntax | 17044
- Description | 17045
- Options | 17045
- Required Privilege Level | 17045
- Output Fields | 17045
- Sample Output | 17046
- Release Information | 17046

Syntax

```
show analyzer analyzer-name
```

Description

Display information about analyzers configured for mirroring.

Options

analyzer-name (Optional) Displays the status of a specific analyzer on the switch.

Required Privilege Level

view

Output Fields

Table 154 on page 17045 lists the output fields for the `command-name` command. Output fields are listed in the approximate order in which they appear.

Table 154: show analyzer Output Fields

Field Name	Field Description
Analyzer name	Displays the name of the analyzer.
Output interface	Specifies a local interface to which mirrored packets are sent. An analyzer can have output to either an interface or a VLAN, not both.
Output VLAN	Specifies a VLAN to which mirrored packets are sent. An analyzer can have output to either an interface or a VLAN, not both.
Mirror ratio	Displays the ratio of packets to be mirrored.
Egress monitored interfaces	Displays interfaces for which traffic exiting the interfaces is mirrored.
Ingress monitored interfaces	Displays interfaces for which traffic entering the interfaces is mirrored.
Ingress monitored VLANs	Displays VLANs for which traffic entering the VLAN is mirrored.

Sample Output

show analyzer

```
user@host> show analyzer
Analyzer name           : employee-monitor
  Output interface             : ge-0/0/10.0
  Output VLAN                  : remote-analyzer
  Mirror ratio                 : 1
  Loss priority                : High
  Egress monitored interfaces  : ge-0/0/3.0
  Ingress monitored interfaces : ge-0/0/0.0
  Ingress monitored interfaces : ge-0/0/1.0
```

Release Information

Command introduced in Junos OS Release 9.0.

show ancp cos

IN THIS SECTION

- Syntax | 17047
- Description | 17047
- Options | 17047
- Required Privilege Level | 17047
- Output Fields | 17047
- Sample Output | 17051
- Release Information | 17053

Syntax

```
show ancp cos
<identifier identifier>
<last-update>
<pending-update>
```

Description

Display information about the CoS state for subscriber traffic.

Options

identifier *identifier* (Optional) Display information about the local loops for the specified access identifier.

last-update (Optional) Display the most recently updated CoS information.

pending-update (Optional) Display the pending update of CoS information.

Required Privilege Level

view

Output Fields

Table 155 on page 17047 lists the output fields for the `show ancp cos` command. Output fields are listed in the approximate order in which they appear.

Table 155: show ancp cos Output Fields

Field Name	Field Description
`Per-DSL CoS adjustment`	Adjustment values applied by the ANCP agent to the actual downstream rates and frame overhead for frame-mode DSL types. The agent then reports the adjusted rates to CoS to establish a shaping rate for the CoS node that corresponds to the subscriber access line.

Table 155: show ancp cos Output Fields *(Continued)*

Field Name	Field Description
`QoS Adjust Flag`	State of QoS adjust: • `TRUE`—The ANCP agent is enabled to adjust the actual downstream data rates and frame overhead and report the adjusted values to CoS. • `FALSE`—The ANCP agent is not enabled to adjust and report values to CoS.
`ADSL bytes`	Number of bytes by which the actual ADSL downstream cell overhead is adjusted before reporting it to CoS.
`ADSL2 bytes`	Number of bytes by which the actual ADSL2 downstream cell overhead is adjusted before reporting it to CoS.
`ADSL2-PLUS bytes`	Number of bytes by which the actual ADSL2+ downstream cell overhead is adjusted before reporting it to CoS.
`SDSL overhead adjusted`	Percentage by which the actual SDSL downstream rate is adjusted before reporting it to CoS.
`SDSL bytes`	Number of bytes by which the actual SDSL downstream frame overhead is adjusted before reporting it to CoS.
`OTHER overhead adjusted`	Percentage by which the actual OTHER downstream rate is adjusted before reporting it to CoS.
`OTHER bytes`	Number of bytes by which the actual OTHER downstream frame overhead is adjusted before reporting it to CoS.
`VDSL overhead adjusted`	Percentage by which the actual VDSL downstream rate is adjusted before reporting it to CoS.
`VDSL bytes`	Number of bytes by which the actual VDSL downstream frame overhead is adjusted before reporting it to CoS.

Table 155: show ancp cos Output Fields *(Continued)*

Field Name	Field Description
VDSL2 overhead adjusted	Percentage by which the actual VDSL2 downstream rate is adjusted before reporting it to CoS.
VDSL2 bytes	Number of bytes by which the actual VDSL2 downstream frame overhead is adjusted before reporting it to CoS.
Per-DSL adjustment for reporting	Adjustment values applied by the ANCP agent to the actual downstream rates for individual DSL types to account for traffic overhead. The agent then reports the adjusted rates to AAA.
ADSL adjustment factor	Percentage by which the actual ADSL downstream rate is adjusted before reporting it to AAA.
ADSL2 adjustment factor	Percentage by which the actual ADSL2 downstream rate is adjusted before reporting it to AAA.
ADSL2+ adjustment factor	Percentage by which the actual ADSL2+ downstream rate is adjusted before reporting it to AAA.
VDSL adjustment factor	Percentage by which the actual VDSL downstream rate is adjusted before reporting it to AAA.
VDSL2 adjustment factor	Percentage by which the actual VDSL2 downstream rate is adjusted before reporting it to AAA.
SDSL adjustment factor	Percentage by which the actual SDSL downstream rate is adjusted before reporting it to AAA.
OTHER adjustment factor	Percentage by which the actual OTHER downstream rate is adjusted before reporting it to AAA.
Keepalive Timer	Interval between the keepalive messages that the ANCP agent sends to CoS.

Table 155: show ancp cos Output Fields *(Continued)*

Field Name	Field Description
`Cos State`	State of the interaction between the ANCP agent and CoS: • ANCPD_COS_CONNECT_NEEDED • ANCPD_COS_CONNECT_PENDING • ANCPD_COS_CONNECT_DONE • ANCPD_COS_SESSION_SENT • ANCPD_COS_WRITE_READY
`Connect Time`	Time at which the ANCP agent connected to CoS; useful for debugging.
`Session Time`	Time at which the ANCP agent sent a session connect message to CoS; useful for debugging.
`Routing Instance Time`	Time at which the ANCP agent sent the routing instance to CoS; useful for debugging.
`Keepalive Time`	Time at which the last keepalive message was sent.
`Update Time`	Time at which the shaping rate was last updated.
`Type`	Subscriber access type: `ifl` indicates that a single VLAN carries subscriber traffic and `iflset` indicates that a set of VLANs carries subscriber traffic.
`Name`	System-wide name of the particular subscriber access.
`Index`	Access identifier.
`Pending Update`	Actual downstream data rate to be applied next to this local loop, in Kbps.
`Last Update`	Adjusted downstream data rate last reported to CoS by the ANCP agent for this local loop, in Kbps.

Sample Output

show ancp cos

```
user@host> show ancp cos

Per-DSL CoS adjustment:
  Qos Adjust Flag:         TRUE
  ADSL bytes:              20
  ADSL2 bytes:              20
  ADSL2-PLUS bytes:              20
  VDSL overhead adjusted:  90
  VDSL bytes:              20
  VDSL2 overhead adjusted: 95
  VDSL2 bytes:             -20
  SDSL overhead adjusted:  85
  SDSL bytes:              30
  OTHER overhead adjusted:  85
  OTHER bytes:              30

Per-DSL adjustment for reporting:
  ADSL adjustment factor:  100
  ADSL2 adjustment factor: 100
  ADSL2+ adjustment factor:100
  VDSL adjustment factor:  100
  VDSL2 adjustment factor: 100
  SDSL adjustment factor:  100
  OTHER adjustment factor:  100


Keepalive Timer:         45 secs
State:                   WRITE_READY
Connect Time:            Fri May  2 12:08:49 2016
Session Time:            Fri May  2 12:18:52 2016
Routing Instance Time:   Fri May  2 12:18:53 2016
Keepalive Time:          Fri May  2 13:44:14 2016
Update Time:             Fri May  2 13:02:55 2016


Type      Name             Index         Pending Update      Last Update
iflset    aci-1004-ge-2/0/0.1073741834 4 None                 36000 Kbps
```

show ancp cos last-update

```
Per-DSL CoS adjustment:
  Qos Adjust Flag:         TRUE
  ADSL bytes:              20
  ADSL2 bytes:              20
  ADSL2-PLUS bytes:              20
  VDSL overhead adjusted:  90
  VDSL bytes:              20
  VDSL2 overhead adjusted: 95
  VDSL2 bytes:             -20
  SDSL overhead adjusted:  85
  SDSL bytes:              30
  OTHER overhead adjusted:  85
  OTHER bytes:              30

Per-DSL adjustment for reporting:
  ADSL adjustment factor:  100
  ADSL2 adjustment factor: 100
  ADSL2+ adjustment factor:100
  VDSL adjustment factor:  100
  VDSL2 adjustment factor: 100
  SDSL adjustment factor:  100
  OTHER adjustment factor:  100

Keepalive Timer:         45 secs
State:                   WRITE_READY
Connect Time:            Fri May  2 12:08:49 2016
Session Time:            Fri May  2 12:18:52 2016
Routing Instance Time:   Fri May  2 12:18:53 2016
Keepalive Time:          Fri May  2 13:44:34 2016
Update Time:             Fri May  2 13:02:55 2016


Type      Name             Index         Pending Update      Last Update
iflset    aci-1004-ge-2/0/0.1073741834 4 None                 36000 Kbps
```

show ancp cos pending-update

```
user@host> show ancp cos pending-update
Per-DSL CoS adjustment:
  Qos Adjust Flag:         TRUE
  VDSL overhead adjusted:  90
  VDSL bytes:              20
  VDSL2 overhead adjusted: 95
  VDSL2 bytes:             -20
  SDSL overhead adjusted:  85
  SDSL bytes:              30
  OTHER overhead adjusted:  85
  OTHER bytes:              30

Per-DSL adjustment for reporting:
  ADSL adjustment factor:  100
  ADSL2 adjustment factor: 100
  ADSL2+ adjustment factor:100
  VDSL adjustment factor:  100
  VDSL2 adjustment factor: 100
  SDSL adjustment factor:  100
  OTHER adjustment factor:  100

Keepalive Timer:         45 secs
State:                   WRITE_READY
Connect Time:            Fri May  2 12:08:49 2016
Session Time:            Fri May  2 12:18:52 2016
Routing Instance Time:   Fri May  2 12:18:53 2016
Keepalive Time:          Fri May  2 13:44:34 2016
Update Time:             Fri May  2 13:02:55 2016
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

show ancp neighbor | 17054

show ancp statistics | 17066

show ancp subscriber | 17072

show ancp neighbor

IN THIS SECTION

- Syntax | 17054
- Description | 17054
- Options | 17054
- Required Privilege Level | 17055
- Output Fields | 17055
- Sample Output | 17061
- Release Information | 17065

Syntax

```
show ancp neighbor
<brief | detail>
<ip-address ip-address
<system-name mac-address>
```

Description

Display information about all ANCP neighbors or the specified ANCP neighbor, regardless of operational state.

Options

brief | detail	(Optional) Display the specified level of detail.
ip-address *ip-address*	(Optional) Display information about the neighbor (access node) specified by the IP address.

system-name *mac-address* (Optional) Display information about the neighbor (access node) specified by the MAC address.

Required Privilege Level

view

Output Fields

Table 156 on page 17055 lists the output fields for the `show ancp neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 156: show ancp neighbor Output Fields

Field Name	Field Description	Level of Output
Version	Version of the ANCP implementation: • `0x31`—General Switch Management Protocol (GSMP) version 3, sub-version 1; ANCP version before *RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks*. • `0x32`—ANCP version 1, defined in *RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks*.	`brief` `detail` none
`IP Address`	IP address of the ANCP neighbor.	`brief` `detail` none
`PartId`	Number that associates the ANCP message with a specific partition.	`brief` none

Table 156: show ancp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Operational state of the ANCP adjacency: • Configured—The neighbor has been configured, but has never been in the Established state. An asterisk (*) is prefixed to the neighbor entry for this state. • Establishing—Adjacency negotiations are in progress for the neighbor. An asterisk (*) is prefixed to the neighbor entry for this state. This state is rarely seen because the adjacency is established so quickly. • Established—Adjacency negotiations have succeeded for the neighbor and an ANCP session has been established. • Not Estblshed—Not Established; adjacency negotiations are ready to begin. Indicates that this neighbor previously had been in the Established state; that is, it has lost a previously established adjacency. An asterisk (*) is prefixed to the neighbor entry for this state.	All levels
Time	How long the adjacency has been up in one of the following formats: • *nwndnh*—number of weeks, days, and hours • *nd hh:mm:ss*—number of days, hours, minutes, and seconds	brief detail none
Subscriber Count	Number of subscribers associated with the ANCP neighbor (access local loop).	brief none
Capabilities	Negotiated ANCP capability: • Topo—Topology discovery. • OAM—Performance of local Operations Administration Maintenance (OAM) procedures on an access loop controlled by the router.	All levels
System Name	MAC address of the ANCP neighbor.	detail

Table 156: show ancp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`TCP Port`	TCP port on which ANCP messages are exchanged.	`detail`
`System Instance`	Number identifying the ANCP link instance from the edge device's perspective.	`detail`
`Peer Instance`	Number identifying the ANCP instance from the access node's perspective. This number is unique and changes when the node or link comes back up after going down.	`detail`
`Timer`	Adjacency timer value advertised by the ANCP peer in 100 ms increments; the interval between ANCP ACK messages. This value remains constant for the duration of an ANCP session.	`detail`
`Partition Type`	Number that identifies whether partitions are used and how the ID is negotiated: • `0`—No partition. • `1`—Fixed partition requested. • `2`—Fixed partition assigned.	`detail`
`Partition Flag`	Number that specifies the type of partition requested: `1` (new adjacency) or `2` (recovered adjacency).	`detail`
`Partition Identifier`	Number that identifies a logical partition of an access node with which the ANCP agent has formed an adjacency. A value of zero indicates that the agent supports each neighbor on an IP address over a single TCP session with a partition ID of zero. This is the default support case. A nonzero value indicates that the agent supports each neighbor on an IP address over a single TCP session with a nonzero partition ID.	`detail`
`Partition Adjacencies`	Number of adjacencies that share the partition.	`detail`

Table 156: show ancp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Dead Timer`	Remaining period that the edge device waits for adjacency packets from a neighbor before declaring the neighbor to be down. The maximum dead time value is three times the configured adjacency timer value. This field displays the current value based on the time that the last adjacency packet was received.	`detail`
`Received Syn Count`	Number of synchronization messages received from neighbors to maintain adjacencies.	`detail`
`Received Synack Count`	Number of synchronization acknowledgment messages received from neighbors in response to the node's synchronization messages.	`detail`
`Received Rstack Count`	Number of messages received from neighbors indicating that the link to the neighbor needs to be reset.	`detail`
`Received Ack Count`	Number of acknowledgment messages periodically received from neighbors after an adjacency has been established.	`detail`
`Received Port Up Count`	Number of status messages received from neighbors indicating that a port has transitioned to the up state.	`detail`
`Received Port Down Count`	Number of status messages received from neighbors indicating that a port has transitioned to the down state.	`detail`
`Received Generic Resp Count`	Number of generic response messages received from neighbors.	`detail`
`Received Adjacency Update Count`	Number of adjacency update messages received from neighbors.	`detail`
`Received OAM Count`	Number of OAM responses received from neighbors in reply to `request` commands.	`detail`

Table 156: show ancp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Received Other Count`	Number of all other ANCP message packets received from neighbors that do not fit into one of the other categories.	`detail`
`Sent Syn Count`	Number of synchronization messages sent to neighbors to maintain adjacencies.	`detail`
`Sent Synack Count`	Number of synchronization acknowledgment messages sent to neighbors in response to the their synchronization messages.	`detail`
`Sent Rstack Count`	Number of messages sent to neighbors indicating that the link to the neighbor needs to be reset.	`detail`
`Sent Ack Count`	Number of acknowledgment messages periodically sent to neighbors after an adjacency has been established.	`detail`
`Sent Generic Resp Count`	Number of generic response messages sent to neighbors.	`detail`
`Sent OAM Count`	Number of OAM `request` commands sent to neighbors.	`detail`
`Max Discovery Limit Exceed Count`	Number of times that the maximum number of discovery table entries accepted from the neighbor has been exceeded.	`detail`

Table 156: show ancp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Additional Counters • `Internal Port Up Sent Count` • `Internal Port Down Sent Count` • `Internal Port Reset Count` • `Internal Adjacency Loss Port Down Count` • `Internal Access Change Port Down Count` • `Adjacency Loss Timeout Count`	• Port-up count sent due to protocol triggered port-ups. • Port-down count sent due to protocol triggered port-downs. • Port reset count (which is port down and up) on issue of 'clear ancp access-loop' command. • Port-down count due to adjacency loss. • Count of adjacency losses occurred for neighbor. • Timestamp of sent and received at GSMP stack. • Timestamp of adjacency loss and EOF (end of file / TCP FIN) received from the neighbor. • Count of EOFs received from neighbor.	`detail`

Table 156: show ancp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Result Codes`	Number of generic response messages sent to neighbors that include each of the following result codes: • `Invalid Request Message Count`—A properly formed request message violated the protocol because of timing (such as a race condition) or direction of transmission. • `Specified Port(s) Down Count`—One or more of the specified ports are down because of a state mismatch between the router and an ANCP control application. • `Out of Resources Count`—ANCP is out of resources, probably not related to the access lines. This result code is sent only by an access node. • `Request Msg Not Implemented Count`— • `Malformed Msg Count`—Message is malformed because it was corrupted in transit or there was an implementation error at either end of the connection. • `TLV Missing Count`—One or more mandatory TLVs was missing from a request. • `Invalid TLV Contents Count`—The contents of one or more TLVs in the request do not match its required specification. • `Non-Existent Port(s) Count`—One or more of the ports specified in a request do not exist, possibly because of a configuration mismatch between the access node and the router or AAA.	`detail`

Sample Output

show ancp neighbor

```
user@host> show ancp neighbor
  Version IP Address       PartID  State         Time        Subscriber    Capabilities
                                                             Count
  0x31    203.0.113.13         0       Established   11:24       2             Topo
```

```
  0x31    203.0.113.15         0       Not Estblshd  2:45        2             Topo
* 0x0     198.51.100.102       0       Establishing  0           0
* 0x0     192.0.2.0            0       Configured    0           0
* 0x0     192.0.2.1            0       Configured    0           0
```

show ancp neighbor detail

```
user@host> show ancp neighbor detail
Neighbor
Information

  Version                 : 0x32
  IP Address              : 10.150.240.13
  System Name             : 00:10:94:02:00:01
  Auto-Configure Interface: xe-2/0/0
    Time                              : 1:45:55
    TCP Port                          : 61412
    State                             : Established
    Subscriber Count                  : 100
    Capabilities                      : Topo, OAM
    System Instance                   : 7
    Peer Instance                     : 26
    Adjacency Timer (in 100ms)        : 250
    Peer Adjacency Timer (in 100ms)   : 250
    Adjacency Loss Hold Time (secs)   : 10
    Partition Type                    : 0
    Partition Flag                    : 1
    Partition Identifier              : 0
    Partition Adjacencies             : 0
    Dead Timer                        : 69
    Received Syn Count                : 5
    Received Synack Count             : 5
    Received Rstack Count             : 0
    Received EOF Count                : 1 <<<<<<
    Received Ack Count                : 354
    Received Port Up Count            : 600
    Received Port Down Count          : 100
    Received Generic Resp Count       : 0
    Received Adjacency Update Count   : 0
    Received OAM Count                : 0
    Received Other Count              : 0
```

```
    Sent Syn Count                   : 5
    Sent Synack Count                : 5
    Sent Rstack Count                : 3
    Sent Ack Count                   : 364
    Sent Generic Resp Count          : 0
    Sent OAM Count                   : 0
    Max Discovery Limit Exceed Count : 0
    Internal Port Up Sent Count      : 600 <<<<<<
    Internal Port Down Sent Count    : 100 <<<<<<
    Internal Port Reset Count        : 2 <<<<<<
    Internal Adjacency Loss Port Down Count : 3 <<<<<<98
    Internal Access Change Port Down Count  : 0 <<<<<<
    Adjacency Loss Timeout Count     : 8 <<<<<<
    Sent Rstack Time                 : Wed, 21 Aug 2024 13:09:20 <<<<<<
    Received Rstack Time             : None <<<<<<
    Received EOF Time                : Wed, 21 Aug 2024 12:47:49 <<<<<<
    Adjacency Loss Timeout Time      : Wed, 21 Aug 2024 13:09:31 <<<<<<
```

show ancp neighbor ip-address

```
user@host> show ancp neighbor ip-address 192.0.2.85

Neighbor Information
      Version                 : 0x32
      IP Address   : 192.0.2.85
      System Name  : 00:00:5e:00:53:ba
          Up Time                          : 26
          TCP Port                         : 32666
          State                            : Established
          Subscriber Count                 : 4
          Capabilities                     : Topo
          System Instance                  : 2
          Peer Instance                    : 20
          Adjacency Timer (in 100ms)       : 100
          Peer Adjacency Timer (in 100ms)  : 100
          Partition Type                   : 0
          Partition Flag                   : 1
          Partition Identifier             : 0
          Partition Adjacencies            : 0
          Dead Timer                       : 23
          Received Syn Count               : 1
```

```
            Received Synack Count            : 1
            Received Rstack Count            : 0
            Received Ack Count               : 4
            Received Port Up Count           : 10
            Received Port Down Count         : 0
            Received Generic Resp Count      : 0
            Received Adjacency Update Count  : 0
            Received OAM Count               : 0
            Received Other Count             : 0
            Sent Syn Count                   : 1
            Sent Synack Count                : 2
            Sent Rstack Count                : 0
            Sent Ack Count                   : 3
            Sent Generic Resp Count          : 0
            Sent OAM Count                   : 0
            Max Discovery Limit Exceed Count : 0
        Result Codes:                          Received              Sent
            Invalid Request Message Count    : 0                     0
            Specified Port(s) Down Count     : 0                     0
            Out of Resources Count           : 0                     0
            Request Msg Not Implemented Count: 0                     0
            Malformed Msg Count              : 0                     0
            TLV Missing Count                : 0                     0
            Invalid TLV Contents Count       : 0                     0
            Non-Existent Port(s) Count       : 0                     0
```

show ancp neighbor system-name

```
user@host> show ancp neighbor 00:00:5e:00:53:ba detail

Neighbor Information
   Version                 : 0x31
   IP Address              : 203.0.113.101
   System Name             : 00:00:5e:00:53:ba
     Up Time                              : 19
     TCP Port                             : 1028
     State                                : Established
     Subscriber Count                     : 2
     Capabilities                         : Topology Discovery, OAM
     System Instance                      : 1
     Peer Instance                        : 10
```

```
    Adjacency Timer (in 100ms)          : 100
    Peer Adjacency Timer (in 100ms)     : 250
    Partition Type                      : 0
    Partition Flag                      : 1
    Partition Identifier                : 0
    Partition Adjacencies               : 0
    Dead Timer                          : 55
    Received Syn Count                  : 1
    Received Synack Count               : 1
    Received Rstack Count               : 0
    Received Ack Count                  : 1
    Received Port Up Count              : 34
    Received Port Down Count            : 0
    Received Generic Resp Count         : 0
    Received Adjacency Update Count     : 0
    Received OAM Responses Count        : 2
    Received Other Count                : 0
    Sent Syn Count                      : 1
    Sent Synack Count                   : 1
    Sent Rstack Count                   : 0
    Sent Ack Count                      : 3
    Sent Generic Resp Count             : 0
    Sent OAM Requests Count             : 4
    Max Discovery Limit Exceed Count    : 3
  Result Codes:                         Received              Sent
    Invalid Request Message Count       : 0                     0
    Specified Port(s) Down Count        : 0                     0
    Out of Resources Count              : 0                     0
    Request Msg Not Implemented Count   : 0                     0
    Malformed Msg Count                 : 0                     0
    TLV Missing Count                   : 0                     0
    Invalid TLV Contents Count          : 0                     0
    Non-Existent Port(s) Count          : 0                     0
```

Release Information

Command introduced in Junos OS Release 9.4.

Additional counters added starting Junos OS Release 21.2R1.

RELATED DOCUMENTATION

show ancp cos | 17046

show ancp subscriber | 17072

show ancp statistics

IN THIS SECTION

- Syntax | 17066
- Description | 17066
- Options | 17066
- Required Privilege Level | 17067
- Output Fields | 17067
- Sample Output | 17069
- Release Information | 17072

Syntax

```
show ancp statistics
<ip-address ip-address>
<system-name mac-address>
```

Description

Display statistics for all ANCP neighbors (access nodes) or the specified ANCP neighbor.

Options

none	Display statistics for all ANCP neighbors, including global statistics not show for individual neighbors.
ip-address *ip-address*	(Optional) Display statistics for only the neighbor with the specified IP address.

system-name ***mac-address***	(Optional) Display statistics for only the neighbor with the specified MAC address.

Required Privilege Level

view

Output Fields

Table 157 on page 17067 lists the output fields for the `show ancp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 157: show ancp statistics Output Fields

Field Name	Field Description
`Number of neighbors`	Total count of ANCP neighbors.
`Number of subscribers`	Total count of ANCP subscribers.
`Accept Count`	Number of neighbor TCP/IP sessions accepted on listener socket.
`Accept Fail Count`	Number of neighbor TCP/IP sessions that failed due to one of the following causes: session already exists, maximum number of ANCP connections exceeded, creation of session or neighbor failed, or protocol start failed.
`No Config Accept Deny Count`	Number of neighbor TCP/IP sessions that failed because the neighbor was not configured.
`Received Syn Count`	Number of synchronization messages received from neighbors to maintain adjacencies.
`Received Synack Count`	Number of synchronization acknowledgment messages received from neighbors in response to the node's synchronization messages.
`Received Rstack Count`	Number of messages received from neighbors indicating that the link to the neighbor needs to be reset.

Table 157: show ancp statistics Output Fields *(Continued)*

Field Name	Field Description
`Received Ack Count`	Number of acknowledgment messages periodically received from neighbors after an adjacency has been established.
`Received Port Up Count`	Number of status messages received from neighbors indicating that a port has transitioned to the up state.
`Received Port Down Count`	Number of status messages received from neighbors indicating that a port has transitioned to the down state.
`Received Generic Resp Count`	Number of generic response messages received from neighbors.
`Received Adjacency Update Count`	Number of adjacency update messages received from neighbors.
`Received OAM Count`	Number of OAM responses received from neighbors in reply to `request` commands.
`Received Other Count`	Number of all other ANCP message packets received from neighbors that do not fit into one of the other categories.
`Sent Syn Count`	Number of synchronization messages sent to neighbors to maintain adjacencies.
`Sent Synack Count`	Number of synchronization acknowledgment messages sent to neighbors in response to the their synchronization messages.
`Sent Rstack Count`	Number of messages sent to neighbors indicating that the link to the neighbor needs to be reset.
`Sent Ack Count`	Number of acknowledgment messages periodically sent to neighbors after an adjacency has been established.

Table 157: show ancp statistics Output Fields *(Continued)*

Field Name	Field Description
`Sent Generic Resp Count`	Number of generic response messages sent to neighbors.
`Sent OAM Count`	Number of OAM `request` commands sent to neighbors.
`Result Codes`	Number of generic response messages sent to neighbors that include each of the following result codes: • `Invalid Request Message Count`—A properly formed request messages violated the protocol because of timing (such as a race condition) or direction of transmission. • `Specified Port(s) Down Count`—One or more of the specified ports are down because of a state mismatch between the router and an ANCP control application. • `Out of Resources Count`—the ANCP agent is out of resources, probably not related to the access lines. This result code is sent only by an access node. • `Request Msg Not Implemented Count`— • `Malformed Msg Count`—Message is malformed because it was corrupted in transit or there was an implementation error at either end of the connection. • `TLV Missing Count`—One or more mandatory TLVs was missing from a request. • `Invalid TLV Contents Count`—The contents of one or more TLVs in the request do not match its required specification. • `Non-Existent Port(s) Count`—One or more of the ports specified in a request do not exist, possibly because of a configuration mismatch between the access node and the router or AAA.

Sample Output

show ancp statistics

```
user@root> show ancp statistics
  Statistics
    Number of neighbors             : 0
```

```
    Number of subscribers           : 0
    Accept Count                    : 0
    Accept Fail Count               : 0
    No Config Accept Deny Count     : 0
    Received Syn Count              : 0
    Received Synack Count           : 0
    Received Rstack Count           : 0
    Received Ack Count              : 0
    Received Port Up Count          : 0
    Received Port Down Count        : 0
    Received Generic Resp Count     : 0
    Received Adjacency Update Count : 0
    Received OAM Count              : 0
    Received Other Count            : 0
    Sent Syn Count                  : 0
    Sent Synack Count               : 0
    Sent Rstack Count               : 0
    Sent Ack Count                  : 0
    Sent Generic Resp Count         : 0
    Sent OAM Count                  : 0
  Result Codes:                      Received              Sent
    Invalid Request Message Count    : 0                    0
    Specified Port(s) Down Count     : 0                    0
    Out of Resources Count           : 0                    0
    Request Msg Not Implemented Count: 0                    0
    Malformed Msg Count              : 0                    0
    TLV Missing Count                : 0                    0
    Invalid TLV Contents Count       : 0                    0
    Non-Existent Port(s) Count       : 0                    0
  Orphan Neighbor Statistics:
    IP Address              Orphan Timeout Time             Port Down Count.  <<<<<<<
    11.22.214.2             Wed, 14 Aug 2024 18:19:38       2857
    11.22.215.2             Wed, 14 Aug 2024 18:20:06       2857
```

show ancp statistics ip-address

```
user@host> show ancp statistics ip-address 203.0.113.1
  Statistics
          Received Syn Count              : 2
          Received Synack Count           : 1
          Received Rstack Count           : 0
```

```
            Received Ack Count              : 8
            Received Port Up Count          : 7
            Received Port Down Count        : 0
            Received Generic Resp Count     : 0
            Received Adjacency Update Count : 0
            Received OAM Count              : 0
            Received Other Count            : 0
            Sent Syn Count                  : 1
            Sent Synack Count               : 1
            Sent Rstack Count               : 0
            Sent Ack Count                  : 17
            Sent Generic Resp Count         : 0
            Sent OAM Count                  : 4
        Result Codes:                           Received              Sent
            Invalid Request Message Count    : 0                      0
            Specified Port(s) Down Count     : 0                      0
            Out of Resources Count           : 0                      0
            Request Msg Not Implemented Count: 0                      0
            Malformed Msg Count              : 0                      0
            TLV Missing Count                : 0                      0
            Invalid TLV Contents Count       : 0                      0
            Non-Existent Port(s) Count       : 0                      0
```

show ancp statistics system-name

```
user@host> show ancp statistics system-name 00:00:5E:00:53:02
  Statistics
            Received Syn Count              : 2
            Received Synack Count           : 1
            Received Rstack Count           : 0
            Received Ack Count              : 8
            Received Port Up Count          : 7
            Received Port Down Count        : 0
            Received Generic Resp Count     : 0
            Received Adjacency Update Count : 0
            Received OAM Count              : 0
            Received Other Count            : 0
            Sent Syn Count                  : 1
            Sent Synack Count               : 1
            Sent Rstack Count               : 0
```

```
        Sent Ack Count                  : 17
        Sent Generic Resp Count         : 0
        Sent OAM Count                  : 4
    Result Codes:                          Received              Sent
        Invalid Request Message Count    : 0                     0
        Specified Port(s) Down Count     : 0                     0
        Out of Resources Count           : 0                     0
        Request Msg Not Implemented Count: 0                     0
        Malformed Msg Count              : 0                     0
        TLV Missing Count                : 0                     0
        Invalid TLV Contents Count       : 0                     0
        Non-Existent Port(s) Count       : 0                     0
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

show ancp cos | 17046

show ancp neighbor | 17054

show ancp subscriber | 17072

show ancp subscriber

IN THIS SECTION

- Syntax | 17073
- Description | 17073
- Options | 17073
- Required Privilege Level | 17074
- Output Fields | 17074
- Sample Output | 17078

- Release Information | **17082**

Syntax

```
show ancp subscriber
<brief | detail>
<access-aggregation-circuit-id circuit-identifier>
<identifier identifier>
<ip-address ip-address>
<system-name mac-address>
```

Description

Display information about active subscribers regardless of the subscriber's operational state, for all subscribers (local access loops), the subscriber associated with the access line specified by an ACI, or the subscriber associated with the specified ANCP neighbor (access node).

After an ancpd restart, this command displays orphaned entries (marked with an o) for subscriber sessions that were established before the restart but which have not yet been reestablished. As sessions are reestablished, the number of orphaned entries displayed by the command decreases. The number reaches zero when all sessions are reestablished or when the orphaned-interface timer expires.

Options

none	Display information about all subscribers.
brief | detail	(Optional) Display the specified level of detail.
access-aggregation-circuit-id *circuit-identifier*	(Optional) Display information about ANCP subscribers whose Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003) matches the specified value. A *circuit-identifier* that begins with the# character indicates a backhaul line identifier. You can specify a wildcard (*) anywhere in the string.
identifier *identifier*	(Optional) Display information about the subscriber associated with the access line (ACI) specified by the access identifier.

ip-address *ip-address* (Optional) Display information about the subscribers connected to the access node specified by the IP address.

system-name *mac-address* (Optional) Display information about the subscribers connected to the access node specified by the MAC address.

Required Privilege Level

view

Output Fields

Table 158 on page 17074 lists the output fields for the `show ancp subscriber` command. Output fields are listed in the approximate order in which they appear.

Table 158: show ancp subscriber Output Fields

Field Name	Field Description	Level of Output
`Loop Identifier`	Access loop identifier as sent by the access node and configured to map the subscriber to an interface. An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle. Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established. An `o` indicates that the entry is for an orphaned interface and represents a previously established subscriber session that has not been reestablished after an ancpd restart. The number of orphaned entries decreases as the ANCP neighbors reestablish adjacencies and the protocol subscriber sessions are reestablished. The command output indicates this by removing the `o` marker. Eventually the number of orphaned entries reaches zero, because either all the adjacencies and subscriber sessions have been reestablished or any remaining orphaned entries are removed when the orphaned-interface timer expires.	`brief` none

Table 158: show ancp subscriber Output Fields *(Continued)*

Field Name	Field Description	Level of Output
DSL Line State	State of the DSL line: Idle, Showtime, or Silent.	brief detail
Access Type	Type of access line employed by the access node: ADSL1, ADSL2, ADSL2+, VDSL1, VDSL2, SDSL, G.fast, VDSL2 Annex Q, SDSL bonded, VDSL2 bonded, G.fast bonded VDSL2 Annex Q bonded or OTHER.	brief detail none
Interface	Name of the interface set or logical interface.	brief detail none
Rate Kbps	Actual downstream data rate for this local loop.	brief none
Neighbor	IP address of ANCP neighbor (access node).	brief none
Access Loop Circuit Identifier	Access loop circuit identifier as sent by the access node and configured to map the subscriber to an interface. An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle. Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established.	detail
Neighbor IP Address	IP address of the ANCP neighbor (access node).	detail
Aggregate Circuit Identifier	ASCII identifier for the subscriber access loop; value of the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003).	detail
Aggregate Circuit Identifier Binary	Binary identifier for the VLAN circuit ID.	detail
Tech Type	Type of technology employed by the subscriber. Currently Junos OS supports DSL technology type only.	detail

Table 158: show ancp subscriber Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`DSL Line Data Link`	Data link protocol employed on the access loop: `AAL5` or `Ethernet`.	`detail`
`DSL Line Encapsulation`	Encapsulation type on the access loop, for Ethernet only: • `0`—NA, type not conveyed • `1`—Untagged Ethernet • `2`—Single-tagged Ethernet	`detail`
`DSL Line Encapsulation Payload`	Payload carried across the access loop: • `0`—NA, type not conveyed • `1`—PPPoA LLC • `2`—PPPoA null • `3`—IPoA LLC • `4`—IPoA null • `5`—Ethernet over AAL5 LLC with FCS • `6`—Ethernet over AAL5 LLC without FCS • `7`—Ethernet over AAL5 null with FCS • `8`—Ethernet over AAL5 null without FCS	`detail`
`Interface Type`	Type of interface employed for subscriber traffic: `ifl` for a single VLAN or `interface-set` for a configured group of VLANs.	`detail`
`Actual Net Data Upstream`	Actual upstream data rate for this local loop, in Kbps.	`detail`

Table 158: show ancp subscriber Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Actual Net Data Downstream`	Actual downstream data rate for this local loop, in Kbps.	`detail`
`Minimum Net Data Upstream`	Minimum upstream data rate desired by the operator for this local loop, in Kbps.	`detail`
`Minimum Net Data Downstream`	Minimum downstream data rate desired by the operator for this local loop, in Kbps.	`detail`
`Maximum Net Data Upstream`	Maximum upstream data rate desired by the operator for this local loop, in Kbps.	`detail`
`Maximum Net Data Downstream`	Maximum downstream data rate desired by the operator for this local loop, in Kbps.	`detail`
`Attainable Net Data Upstream`	Maximum attainable upstream data rate for this local loop, in Kbps.	`detail`
`Attainable Net Data Downstream`	Maximum attainable downstream data rate for this local loop, in Kbps.	`detail`
`Minimum Low Power Data Downstream`	Minimum downstream data rate desired by the operator for this local loop in low power state, in Kbps.	`detail`
`Minimum Low Power Data Upstream`	Minimum upstream data rate desired by the operator for this local loop in low power state, in Kbps.	`detail`
`Maximum Interleave Delay Downstream`	Maximum interleaving delay for downstream data, in milliseconds.	`detail`
`Maximum Interleave Delay Upstream`	Maximum interleaving delay for upstream data, in milliseconds.	`detail`

Table 158: show ancp subscriber Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Actual Interleave Delay Downstream	Actual interleaving delay for downstream data, in milliseconds.	detail
Actual Interleave Delay Upstream	Actual interleaving delay for upstream data, in milliseconds.	detail
Port-Up Timestamp	Timestamp when port comes up or port enabled.	detail
Port Down Timestamp	Timestamp when port goes down or disabled.	detail
Port Down Event	Reason for port to be disabled.	detail

Sample Output

show ancp subscriber

```
user@host> show ancp subscriber
  Loop Identifier        DSL Line  Tech Type     Access Type   Interface       Rate   Neighbor
                         State                                 Kbps
**circuit 101            Idle      DSL       ADSL1    ----              32       203.0.113.13
**circuit 102            Idle      DSL       ADSL1    ----              32       203.0.113.13
  circuit 301            Showtime  DSL       ADSL1    ----              32       203.0.113.15
  circuit 302            Showtime  DSL       ADSL1    ----              32       203.0.113.15
```

show ancp subscriber (After ancpd Restart)

```
user@host> show ancp subscriber
  Loop Identifier        DSL Line  Tech Type     Access Type  Interface         Rate     Neighbor
                         State                                 Kbps
o circuit 201            Showtime  DSL        ADSL1    ----              222222
o circuit 202            Showtime  DSL        ADSL1    ----              222222
```

show ancp subscriber brief

```
user@host> show ancp subscriber brief

  Loop Identifier          Type      Interface        Rate          Neighbor
                                                      Kbps
  port-1-10                VDSL2    set-ge-10410       64       203.0.113.102
  port-1-11                VDSL2    set-ge-10411       64       203.0.113.111
  port-2-10                VDSL2    ge-1/0/4.12        64       203.0.113.112
  port-2-11                VDSL2    ge-1/0/4.13        64       203.0.113.113
```

show ancp subscriber detail (Access type DSL)

```
user@host> show ancp subscriber detail
 Subscriber Information
*  Access Loop Circuit Identifier : circuit 101
     Neighbor IP Address                  : 203.0.113.13
     Aggregate Circuit Identifier Binary  : 0/0
     Tech Type                                                                : DSL
     Access Type                          : ADSL1
     DSL Line State                       : Idle
     DSL Line Data Link                   : Data link 2
     DSL Line Encapsulation               : N/A
     DSL Line Encapsulation Payload       : N/A
     Interface Type                       : N/A
     Interface                            : ----
     Actual Net Data Upstream             : 32
     Actual Net Data Downstream           : 32
     Minimum Net Data Upstream            : 0
     Minimum Net Data Downstream          : 0
     Maximum Net Data Upstream            : 0
     Maximum Net Data Downstream          : 0
     Attainable Net Data Upstream         : 1024
     Attainable Net Data Downstream       : 8192
     Minimum Low Power Data Downstream    : 32
     Minimum Low Power Data Upstream      : 32
     Maximum Interleave Delay Downstream  : 20
     Maximum Interleave Delay Upstream    : 20
     Actual Interleave Delay Downstream   : 20
     Actual Interleave Delay Upstream     : 20
```

```
*  Access Loop Circuit Identifier: circuit 102
     Neighbor IP Address                  : 213.0.113.13
     Aggregate Circuit Identifier Binary  : 0/0
     Tech Type                                                                  : DSL
     Access Type                          : ADSL1
     DSL Line State                       : Idle
     DSL Line Data Link                   : Data link 2
     DSL Line Encapsulation               : N/A
     DSL Line Encapsulation Payload       : N/A
     Interface Type                       : N/A
     Interface                            : ----
     Actual Net Data Upstream             : 32
     Actual Net Data Downstream           : 32
     Minimum Net Data Upstream            : 0
     Minimum Net Data Downstream          : 0
     Maximum Net Data Upstream            : 0
     Maximum Net Data Downstream          : 0
     Attainable Net Data Upstream         : 1024
     Attainable Net Data Downstream       : 8192
     Minimum Low Power Data Downstream    : 32
     Minimum Low Power Data Upstream      : 32
     Maximum Interleave Delay Downstream  : 20
     Maximum Interleave Delay Upstream    : 20
     Actual Interleave Delay Downstream   : 20
     Actual Interleave Delay Upstream     : 20
...
     Port-Up Timestamp                    : Wed, 12 Sep 2024 18:19:38
     Port-Down Timestamp                  : Wed, 14 Sep 2024 20:19:00
     Port-Down Event                      : Orphaned subscriber timeout
```

show ancp subscriber detail (Access type PON)

```
user@root> show ancp subscriber detail
Subscriber Information
   Access Loop Circuit Identifier : circuit1
     Neighbor IP Address                  : 192.85.1.4
     Access Loop Remote Identifier        : remote
     Aggregate Circuit Identifier Binary  : 2:1
     Tech Type                            : PON
     Access Type                          : GPON
     Interface Type                       : N/A
```

```
     Interface                           : ----
     Line Attributes:
       ONT/ONU Average Data Rate Downstream      :     100000 Kbps
       ONT/ONU Peak Data Rate Downstream         :     200000 Kbps
       ONT/ONU Maximum Data Rate Upstream        :     300000 Kbps
       ONT/ONU Assured Data Rate Upstream        :     300000 Kbps
     Port-Up Timestamp                   : Wed, 14 Aug 2024 18:19:38
     Port-Down Timestamp                 : Wed, 15 Aug 2024 20:19:00
     Port-Down Event                     : Orphaned subscriber timeout
```

show ancp subscriber access-aggregation-circuit-id detail

```
user@host> show ancp subscriber access-aggregation-circuit-id "#TEST-DPU-C-100" detail
Subscriber Information
 * Access Loop Circuit Identifier : circuit 201
     Neighbor IP Address                 : 192.0.2.1
     Access Loop Remote Identifier       : remote 123
     Aggregate Circuit Identifier        : #TEST-DPU-C-100
     Aggregate Circuit Identifier Binary  : 50
     Tech Type:                          : DSL
     Interface Type                      : interface
     Interface                           : ge-1/0/0.3221225475
     Actual Net Data Upstream            : 1024
     Actual Net Data Downstream          : 2048
     Maximum Net Data Upstream           : 0
     Maximum Net Data Downstream         : 0
* Access Loop Circuit Identifier : circuit 202
     Neighbor IP Address                 : 192.0.2.1
     Access Loop Remote Identifier       : remote 185
     Aggregate Circuit Identifier        : #TEST-DPU-C-100
     Aggregate Circuit Identifier Binary  : 50
     Tech Type:                          : DSL
     Interface Type                      : interface
     Interface                           : ge-1/0/0.3221225476
     Actual Net Data Upstream            : 1024
     Actual Net Data Downstream          : 2048
     Maximum Net Data Upstream           : 0
     Maximum Net Data Downstream         : 0
     Port-Up Timestamp                   : Wed, 14 Aug 2024 18:19:38
```

```
    Port-Down Timestamp                  : Wed, 15 Aug 2024 20:19:00
    Port-Down Event                      : Orphaned subscriber timeout
```

show ancp subscriber identifier identifier-string detail

```
user@host> show ancp subscriber identifier port-1-11 detail
          Access Loop Identifier : port-1-11
          Neighbor IP Address                  : 203.0.113.112
          Aggregate Circuit Identifier Binary  : 0/0
          DSL Type                             :  DSL  0
          Interface Type                       : interface-set
          Interface                            : set-ge-10411
          DSL Line State                       : Show Time
          Actual Net Data Upstream             : 64
          Actual Net Data Downstream           : 64
          DSL Line Data Link                   : AAL5
          DSL Line Encapsulation               : N/A
          DSL Line Encapsulation Payload       : N/A
          Minimum Net Data Upstream            : 64
          Minimum Net Data Downstream          : 64
          Maximum Net Data Upstream            : 64
          Maximum Net Data Downstream          : 64
          Attainable Net Data Upstream         : 64
          Attainable Net Data Downstream       : 64
          Minimum Low Power Data Downstream    : 64
          Minimum Low Power Data Upstream      : 64
          Maximum Interleave Delay Downstream  : 50
          Maximum Interleave Delay Upstream    : 50
          Actual Interleave Delay Downstream   : 50
          Actual Interleave Delay Upstream     : 50
          Port-Up Timestamp                    : Wed, 14 Aug 2024 18:19:38
          Port-Down Timestamp                  : Wed, 15 Aug 2024 20:19:00
          Port-Down Event                      : Orphaned subscriber timeout
```

Release Information

Command introduced in Junos OS Release 9.4.

`neighbor` option replaced with `ip-address` in Junos OS Release 16.1.

`system-name` option introduced in Junos OS Release 16.1.

`access-aggregation-circuit-id` option introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

- clear ancp subscriber | 3421
- show ancp cos | 17046
- show ancp neighbor | 17054
- show ancp statistics | 17066

show ancp summary

IN THIS SECTION

- Syntax | 17083
- Description | 17083
- Options | 17083
- Required Privilege Level | 17084
- Output Fields | 17084
- Sample Output | 17085
- Release Information | 17085

Syntax

```
show ancp summary
```

Description

Display a summary of the counts and states for all ANCP neighbors and subscribers.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 159 on page 17084 lists the output fields for the `show ancp summary` command. Output fields are listed in the approximate order in which they appear.

Table 159: show ancp summary Output Fields

Field Name	Field Description
`Configured`	Number of ANCP neighbors in the `Configured` state; that is, that have been configured but never established.
`Establishing`	Number of ANCP neighbors in the `Establishing` state; that is, where negotiations are in progress.
`Established`	Number of ANCP neighbors in the `Established` state; that is, where negotiations have succeeded and the ANCP session has been established.
`Not Estblshd`	Number of ANCP neighbors in the `Not Estblshd` state; that is, that have lost a previously established adjacency and are ready to begin negotiations.
`Total`	Total number of ANCP neighbors; sum of neighbors in the `Configured`, `Establishing`, `Established`, and `Not Estblshd` states.
`Showtime`	Number of DSL lines in `Showtime` state.
`Idle`	Number of DSL lines in `Idle` state.
`Silent`	Number of DSL lines in `Silent` state.
`Unknown`	Number of DSL lines where the state is not `Showtime`, `Idle`, or `Silent`.

Table 159: show ancp summary Output Fields *(Continued)*

Field Name	Field Description
Total	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the Showtime, Idle, Silent, and Unknown states.

Sample Output

show ancp summary

```
user@host> show ancp summary

Neighbors Summary:
  Configured  Establishing  Established  Not Established     Total
  ----------  ------------  -----------  ---------------  ----------
          22             0            2                0          24

Subscribers Summary:
  Showtime      Idle      Silent     Unknown     Total
  ---------- ---------- ----------- ---------- ----------
           4          0           0          0          4
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

- show ancp neighbor | 17054
- show ancp summary neighbor | 17086
- show ancp subscriber | 17072
- show ancp summary subscriber | 17089

show ancp summary neighbor

IN THIS SECTION

- Syntax | 17086
- Description | 17086
- Options | 17086
- Required Privilege Level | 17086
- Output Fields | 17087
- Sample Output | 17088
- Release Information | 17088

Syntax

```
show ancp summary neighbor
<ip-address ip-address | system-name mac-address>
```

Description

Display a summary of the counts and states for all ANCP neighbors and of the neighbor's subscribers when you specify a particular neighbor.

Options

ip-address *ip-address*	(Optional) IP address of the ANCP neighbor (access node).
system-name *mac-address*	(Optional) MAC address of the ANCP neighbor (access node).

Required Privilege Level

view

Output Fields

Table 160 on page 17087 lists the output fields for the `show ancp summary` command. Output fields are listed in the approximate order in which they appear.

Table 160: show ancp summary neighbor Output Fields

Field Name	Field Description
`Configured`	Number of ANCP neighbors in the `Configured` state; that is, that have been configured but never established.
`Establishing`	Number of ANCP neighbors in the `Establishing` state; that is, where negotiations are in progress.
`Established`	Number of ANCP neighbors in the `Established` state; that is, where negotiations have succeeded and the ANCP session has been established.
`Not Estblshd`	Number of ANCP neighbors in the `Not Estblshd` state; that is, that have lost a previously established adjacency and are ready to begin negotiations.
`Total`	Total number of ANCP neighbors; sum of neighbors in the `Configured`, `Establishing`, `Established`, and `Not Estblshd` states.
`Showtime`	Number of DSL lines for the neighbor in `Showtime` state.
`Idle`	Number of DSL lines for the neighbor in `Idle` state.
`Silent`	Number of DSL lines for the neighbor in `Silent` state.
`Unknown`	Number of DSL lines for the neighbor where the state is not `Showtime`, `Idle`, or `Silent`.
`Total`	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the `Showtime`, `Idle`, `Silent`, and `Unknown` states.

Sample Output

show ancp summary neighbor

```
user@host> show ancp summary neighbor

Neighbors Summary:
  Configured  Establishing  Established  Not Established     Total
  ----------  ------------  -----------  ---------------  ----------
          22             0            2                0          24
```

show ancp summary neighbor (IP Address)

```
user@host> show ancp summary neighbor ip-address 192.168.10.1

Neighbor Summary:192.168.10.1 status Established

Subscribers Summary:
  Show Time     Idle      Silent     Unknown     Total
  ---------- ---------- ----------- ---------- ----------
           6          0           0          0          6
```

show ancp summary neighbor (MAC Address)

```
user@host> show ancp summary neighbor system-name 00:00:5E:00:53:02

Neighbor Summary:00:00:5E:00:53:02 status Established

Subscribers Summary:
  Show Time     Idle      Silent     Unknown     Total
  ---------- ---------- ----------- ---------- ----------
           5          1           2          0          8
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

show ancp summary | 17083
show ancp subscriber | 17072
show ancp summary subscriber | 17089

show ancp summary subscriber

IN THIS SECTION

- Syntax | 17089
- Description | 17089
- Options | 17089
- Required Privilege Level | 17089
- Output Fields | 17090
- Sample Output | 17090
- Release Information | 17090

Syntax

```
show ancp summary subscriber
```

Description

Display a summary of the counts and states for all ANCP subscribers.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 161 on page 17090 lists the output fields for the `show ancp summary subscriber` command. Output fields are listed in the approximate order in which they appear.

Table 161: show ancp summary subscriber Output Fields

Field Name	Field Description
`Showtime`	Number of DSL lines in `Showtime` state.
`Idle`	Number of DSL lines in `Idle` state.
`Silent`	Number of DSL lines in `Silent` state.
`Unknown`	Number of DSL lines where the state is not `Showtime`, `Idle`, or `Silent`.
`Total`	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the `Showtime`, `Idle`, `Silent`, and `Unknown` states.

Sample Output

show ancp summary subscriber

```
user@host> show ancp summary subscriber

Subscribers Summary:
  Show Time     Idle      Silent     Unknown     Total
  ---------- ---------- ----------- ---------- ----------
           8          1           0          1         10
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

show ancp summary | 17083
show ancp neighbor | 17054
show ancp summary neighbor | 17086

show app-engine crash

IN THIS SECTION

- Syntax | 17091
- Description | 17091
- Options | 17092
- Additional Information | 17092
- Required Privilege Level | 17092
- Output Fields | 17092
- Sample Output | 17093
- Release Information | 17093

Syntax

```
show app-engine crash
<compute-cluster compute-cluster-name>
<compute-cluster compute-cluster-name compute-node compute-node-name>
```

Description

Displays information regarding crashes on the Junos V App Engine or host OS. The command displays the list of files present in the **/var/crash/** directory, as well as some core files located in the **/var/tmp/** directory.

Options

- `compute-cluster` *`compute-cluster-name`* `compute-node` *`compute-node-name`*—Name of the compute cluster and compute node. On QFX switches with a host OS, the default compute cluster and compute node names are default-cluster and default-node, and the command applies to the host OS by default.

Additional Information

In the operational mode of the CLI when you type **?** for a name, for example a compute-node name, you would expect to get a list of available compute nodes plus the option to type in a name not listed. This is the auto-complete feature in the CLI. However, in JunosV App Engine, ,if you specify compute cluster and compute node in the operational command, the auto-complete works only if the compute cluster is put before the compute node.

For commands with an optional `compute-cluster` *`compute-cluster-name`* option, if that option is omitted, the command will be executed on all compute nodes of all compute clusters. For commands with an optional `compute-node` *`compute-node-name`* option, if that option is omitted, the command will be executed on all compute nodes of the specified compute cluster.

Required Privilege Level

view

Output Fields

For a description of the output fields, see Table 162 on page 17092. Output fields are listed in the approximate order in which they appear.

Table 162: show app-engine crash Output Fields

Field Name	Field Description
`Compute Cluster`	Name of compute cluster.
`Compute Node`	Name of compute node.
`Crash Info`	List of core files.

Sample Output

show app-engine crash (QFX5100 Switches)

```
user@host> show app-engine crash

Compute cluster: default-cluster

  Compute node: default-node

   Crash Info
   ==========
   total 0
```

Release Information

Command introduced in Junos OS Release 13.2X51-D15.

RELATED DOCUMENTATION

request app-engine file-copy (crash | log) from-jhost to-vjunos | 14429

show app-engine logs

IN THIS SECTION

- Syntax | 17094
- Description | 17094
- Options | 17094
- Additional Information | 17094
- Required Privilege Level | 17094
- Output Fields | 17095
- Sample Output | 17095

Release Information | **17096**

Syntax

```
show app-engine logs
<compute-cluster compute-cluster-name>
<compute-cluster compute-cluster-name compute-node compute-node-name>
```

Description

Displays log files regarding the state of the specified Guest VM or host OS, located in the **/var/log/** directory.

Options

- `compute-cluster` *`compute-cluster-name`* `compute-node` *`compute-node-name`*—Name of the compute cluster and name of the compute node. On QFX switches with a host OS, the default compute cluster and compute node names are default-cluster and default-node, and the command applies to the host OS by default.

Additional Information

In the operational mode of the CLI when you type **?** for a name, for example a compute-node name, you would expect to get a list of available compute nodes plus the option to type in a name not listed. This is the auto-complete feature in the CLI. However, in JunosV App Engine, if you specify compute cluster and compute node in the operational command, the auto-complete works only if the compute cluster is put before the compute node.

For commands with an optional `compute-cluster` *`compute-cluster-name`* option, if that option is omitted, the command will be executed on all compute nodes of all compute clusters. For commands with an optional `compute-node` *`compute-node-name`* option, if that option is omitted, the command will be executed on all compute nodes of the specified compute cluster.

Required Privilege Level

view

Output Fields

For a description of the output fields, see Table 163 on page 17095. Output fields are listed in the approximate order in which they appear.

Table 163: show app-engine logs Output Fields

Field Name	Field Description
`Compute Cluster`	Name of compute cluster.
`Compute Node`	Name of compute node.
`Logs Info`	List of log files.

Sample Output

show app-engine logs (QFX5100 Switches)

```
user@host> show app-engine logs
Compute cluster: default-cluster

  Compute node: default-node

   Logs Info
   =========
   total 1012
   drwxr-xr-x 2 root root   4096 Feb  2 13:53 audit
   -rw-r--r-- 1 root root      0 Feb  2 13:59 named.log
   -rw-r--r-- 1 root root      0 Feb  2 15:30 btmp
   drwxr-xr-x 2 root root   4096 Feb  2 17:11 ntpstats
   -rw------- 1 root root      0 Feb  2 17:27 tallylog
   drwxr-xr-x 2 root root   4096 Feb  7 10:10 stap-server
   -rw------- 1 root root      0 Feb  7 10:16 spooler
   -rw------- 1 root root      0 Feb  7 10:16 maillog
   -rw------- 1 root root      0 Feb  7 10:16 boot.log
   drwxr-xr-x 5 root root   4096 Feb  7 10:16 libvirt
   drwxr-x--- 2 root root   4096 Feb  7 10:16 watchdog
   -rw-r--r-- 1 root root  41693 Feb 10 04:27 dmesg.old
```

```
drwxr-xr-x 2 root root   4096 Feb 10 04:30 sa
-rw-r--r-- 1 root root  41693 Feb 10 09:35 dmesg
-rw-rw-r-- 1 root root   1193 Feb 10 09:47 vjunos_install_log
-rw-rw-r-- 1 root utmp 141312 Feb 10 10:08 wtmp
-rw------- 1 root root  25514 Feb 10 10:08 secure
-rw-r--r-- 1 root root    292 Feb 10 10:08 lastlog
-rw------- 1 root root 338593 Feb 10 10:17
cron
-rw-r--r-- 1 root root 383229 Feb 10 10:17 messages
```

Release Information

Command introduced in Junos OS Release 13.2X51-D15.

RELATED DOCUMENTATION

request app-engine file-copy (crash | log) from-jhost to-vjunos | 14429

show arp

IN THIS SECTION

- Syntax | 17097
- Description | 17097
- Options | 17098
- Required Privilege Level | 17099
- Output Fields | 17099
- Sample Output | 17100
- Release Information | 17102

Syntax

```
show arp
<expiration-time>
<hostname host-name>
<interface interface-name>
<logical-system logical-system-name>
<no-resolve>
<reference-count count>
<tenant name>
<state state>
<vpn vpn-name>
```

Description

Display all entries in the Address Resolution Protocol (ARP) table. To display entries for a particular logical system only, first enter the `set cli logical-system` *`logical-system-name`* command, and then enter the `show arp` command.

To force the output of this command to show the IP address by the hostname of each device, also configure the `force-show-arp-resolve` statement at the `[edit system services subscriber-management overrides]` hierarchy.

NOTE: Starting with Junos OS Release 14.2, the following enhancements have been made to the output of the `show arp interfaces` command:

- For integrated routing and bridging (IRB) interfaces, in the output of the `show arp` command, the IRB interface name is displayed under the Interface field of the output and the Layer 2 interface identifier is specified in square brackets following the IRB name. Until Release 14.1 and earlier, only the layer 2 interface name and not the IRB name was displayed.

- Starting with release 14.2, if you do not specify a subinterface or a logical unit of the interface with the `show arp interface` *`interface-name`* command, an error message is shown. Until Release 14.1 and earlier, if you did not specify the subinterface for a physical interface, the system considered the supplied command to be for subinterface 0 and displayed the output. For example, if you entered ge-2/2/5, it was processed by the system as ge-2/2/5.0.

- When IRB interfaces are configured and if you attempt to specify an interface name that is not configured on the system, an error message is displayed stating the particular interface is not defined on the system. Until release 14.1 and earlier, unrelated and incorrect entries were displayed even for interface names that did not exist.
- Starting with Release 14.2, you can enter the `show arp interface` command with the IRB name and retrieve the statistical details for the IRB interface. This functionality was not available previously. However, you could previously obtain the ARP details of an IRB interface that had a Layer 2 interface configured.

NOTE: Starting with Junos OS Release 16.1, `show arp no-resolve` command does not display the underlying ifl information if `enhanced-convergence` statement at `[edit irb unit unit-number]` hierarchy level and `enhanced-ip` statement at [edit chassis network-services] hierarchy level is configured for the destination interface IRB.

Options

none	Display the entries in the ARP table.
expiration-time	(Optional) Display the amount of time, in seconds, until each ARP entry is set to expire.
hostname *host-name*	(Optional) Display the hostname.
interface *interface-name*	(Optional) Display information about ARP for the specified logical interface
logical-system *logical-system-name*	(Optional) Display ARP entries for the specified logical system; only available on the main router context.
no-resolve	(Optional) Do not attempt to determine the hostname that corresponds to the IP address.
reference-count *count*	(Optional) Display the next-hop reference count.
state *state*	(Optional) Display the next-hop current state.
tenant *name*	(Optional) Display the ARP entries for the specified tenant. Only available in the main router context.

vpn *vpn-name* (Optional) Display entries in the ARP table for the specified virtual private network's (VPN) routing table.

> **NOTE:**
>
> - The `vpn` option displays only entries for non-subscriber interfaces. You must omit this option if you want to view ARP entries for subscriber interfaces.
> - The `show arp vpn` *`vpn-name`* `no-resolve` command does not display entries for subscriber interfaces in the non-default routing-instance, that is when the subscriber interface is coming up in the non-default routing-instance.

Required Privilege Level

view

Output Fields

Table 164 on page 17099 describes the output fields for the `show arp` command. Output fields are listed in the approximate order in which they appear.

Table 164: show arp Output Fields

Field Name	Field Description
MAC Address	Media access control (MAC) address that corresponds to the IP address.
Address	IP address that corresponds to the hostname.
Name	Hostname.
Interface	Interface name.

Table 164: show arp Output Fields *(Continued)*

Field Name	Field Description
Flags	(`no-resolve` option only) Indicates how mappings between IP and MAC addresses are defined: • `permanent`—Static mapping. The ARP entry never times out. • `published`—The next hop is published. This flag tells the host to respond to ARP request and ARP response packets. When you see this flag, the device acts as an ARP server and responds to host name requests even if the host address is not its own. • `gateway`—The device rewrites the source MAC address with the interface MAC address. You will see this flag when the device uses the interface MAC address for the Layer 2 header. • `remote`—The ARP entry is installed by a process (for example, the l2ald process), not the kernel. • `None`—Dynamic mapping.
TTE	(`expiration-time` option only) Amount of time, in seconds, until ARP entry is set to expire.

Sample Output

show arp

```
user@host> show arp
MAC Address       Address         Name                      Interface
00:e0:81:22:fd:74 192.168.64.10   firewall.my.net           fxp0.0
00:04:5a:65:78:e1 192.168.65.13   lab.my net                fxp0.0
```

show arp no-resolve

```
user@host> show arp no-resolve
MAC Address       Address         Interface     Flags
00:90:69:96:00:01 10.10.45.5      fe-0/0/1.0    none
00:00:00:00:00:01 200.200.200.1   fe-0/0/0.0    permanent published
00:00:00:00:00:02 200.200.200.2   fe-0/0/0.0    permanent
```

```
00:90:69:91:b0:00 200.200.200.3   fe-0/0/0.0    none
Total entries: 4
```

show arp no-resolve

```
user@host> show arp no-resolve (QFX Series)
```

The command is displaying the MAC and MAC-IP routes are programmed over the .local interface.

```
MAC Address       Address         Interface                Flags
00:00:10:00:11:00 10.0.1.1        irb.100 [.local..5]      permanent remote
00:00:10:00:22:00 10.0.1.2        irb.100 [.local..5]      permanent remote
00:00:10:00:33:00 10.0.1.3        irb.100 [.local..5]      permanent remote
00:00:10:00:44:00 10.0.1.4        irb.100 [.local..5]      permanent remote
00:00:10:11:22:00 10.0.1.12       irb.100 [.local..5]      permanent remote
00:00:10:11:33:00 10.0.1.13       irb.100 [.local..5]      permanent remote
00:00:10:11:44:00 10.0.1.14       irb.100 [.local..5]      permanent remote
00:00:10:11:66:00 10.0.1.16       irb.100 [.local..5]      permanent remote
40:00:10:11:11:00 10.0.1.254      irb.100                  permanent published gateway
00:00:10:00:11:01 10.0.2.1        irb.200 [.local..5]      permanent remote
00:00:10:00:22:01 10.0.2.2        irb.200 [.local..5]      permanent remote
00:00:10:00:33:01 10.0.2.3        irb.200 [.local..5]      permanent remote
00:00:10:00:44:01 10.0.2.4        irb.200 [.local..5]      permanent remote
00:00:10:11:22:01 10.0.2.12       irb.200 [.local..5]      permanent remote
00:00:10:11:33:01 10.0.2.13       irb.200 [.local..5]      permanent remote
00:00:10:11:44:01 10.0.2.14       irb.200 [.local..5]      permanent remote
00:00:10:11:66:01 10.0.2.16       irb.200 [.local..5]      permanent remote
Total entries: 26
```

command-name

```
user@host> show arp no-resolve
```

The command displaying the underlying l2 ifl information when `enhanced-convergence` statement and `enhanced-ip` statement is not configured.

```
show arp no-resolve
MAC Address       Address         Interface         Flags
```

```
02:01:00:00:00:05 10.0.0.5        em1.0                    none
00:00:5e:00:01:1b 91.91.91.50     irb.0[xe-2/1/0.0]        none  >>> underlying l2 ifl associated
02:01:00:00:00:05 128.0.0.5       em1.0                    none
02:01:00:00:00:05 128.0.0.6       em1.0                    none
02:00:00:00:00:12 128.0.0.18      em0.0                    none
00:26:88:6a:c6:80 192.168.237.126 fxp0.0                   none
Total entries: 6
```

The command not displaying the underlying l2 ifl information when `enhanced-convergence` statement and `enhanced-ip` statement is configured.

```


MAC Address       Address         Interface         Flags
02:01:00:00:00:05 10.0.0.5        em1.0                    none
00:00:5e:00:01:1b 91.91.91.50     irb.0                    none  >>> underlying l2 ifl
association is removed.
02:01:00:00:00:05 128.0.0.5       em1.0                    none
02:01:00:00:00:05 128.0.0.6       em1.0                    none
02:00:00:00:00:12 128.0.0.18      em0.0                    none
00:26:88:6a:c6:80 192.168.237.126 fxp0.0                   none
Total entries: 6
```

show arp expiration-time

```
user@host> show arp expiration-time
MAC Address       Address         Name                 Interface     Flags TTE
00:a0:a5:12:3e:d4 10.0.0.5        10.0.0.5                  fxp1.0 none
00:e0:81:22:fd:74 192.168.64.10   supernova.englab.juniper. fxp0.0 none 1491
00:30:48:84:03:56 192.168.65.46   kgb.englab.juniper.net    fxp0.0 none 1279
00:03:ba:12:f7:5e 192.168.65.226  nmssun1-eri0.englab.junip fxp0.0 none 452
00:90:69:8e:b0:fc 192.168.71.254  stonewall-ge-200.englab.j fxp0.0 none 1421
Total entries: 5
```

Release Information

Command introduced before Junos OS Release 7.4.

`expiration-time` option added in Junos OS Release 8.1.

`logical-system` and `vpn` options added in Junos OS Release 10.1.

`reference-count`, `tenant`, and `state` options added in Junos OS Release 18.3.

RELATED DOCUMENTATION

clear arp | 3424
set cli logical-system
overrides (Enhanced Subscriber Management) | 12074

show arp inspection statistics

IN THIS SECTION

- Syntax | 17103
- Description | 17103
- Required Privilege Level | 17103
- Output Fields | 17104
- Sample Output | 17104
- Release Information | 17105

Syntax

```
show arp inspection statistics
```

Description

Display ARP inspection statistics.

Required Privilege Level

view

Output Fields

Table 165 on page 17104 lists the output fields for the `show arp inspection statistics` command. Output fields are listed in the approximate order in which they appear.

Table 165: show arp inspection statistics Output Fields

Field Name	Field Description	Level of Output
Interface	Interface on which ARP inspection has been applied.	All levels
Packets received	Total number of packets total that underwent ARP inspection.	All levels
ARP inspection pass	Total number of packets that passed ARP inspection.	All levels
ARP inspection failed	Total number of packets that failed ARP inspection.	All levels

Sample Output

show arp inspection statistics

```
user@switch> show arp inspection statistics

Interface     Packets received     ARP inspection pass  ARP inspection failed
---------     -----------------    -------------------  ---------------------
  ge-0/0/0                    0                      0                      0
  ge-0/0/1                    0                      0                      0
  ge-0/0/2                    0                      0                      0
  ge-0/0/3                    0                      0                      0
  ge-0/0/4                    0                      0                      0
  ge-0/0/5                    0                      0                      0
  ge-0/0/6                    0                      0                      0
  ge-0/0/7                  703                    701                      2
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

clear arp inspection statistics | 3426

Example: Configuring Port Security (non-ELS)

Verifying That DAI Is Working Correctly

show as-path

IN THIS SECTION

- Syntax | 17105
- Syntax (EX Series Switches) | 17106
- Description | 17106
- Options | 17106
- Required Privilege Level | 17106
- Output Fields | 17106
- Sample Output | 17108
- Release Information | 17110

Syntax

```
show as-path
<brief | detail>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show as-path
<brief | detail>
```

Description

Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the `as-path-prepend` action) or through aggregation.

AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.

Options

none	Display basic information about AS paths that the local routing device is using (same as brief).
brief | detail	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 166 on page 17107 lists the output fields for the `show as-path` command. Output fields are listed in the approximate order in which they appear.

Table 166: show as-path Output Fields

Field Name	Field Description	Level of Output
`Total AS paths`	Total number of AS paths.	`brief` none
`Bucket`	Bucket number.	All levels
`Count`	Number of AS path entries in this bucket.	All levels
`AS path`	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • `I`—IGP. • `E`—EGP. • `?`—Incomplete; typically, the AS path was aggregated. • `Atomic`—Route is an aggregate of several route prefixes. • `Aggregator`—Routing device has summarized a range of prefixes.	All levels
`domain`	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.	`detail`
`neighbor as`	AS peer address.	`detail`
`length`	Length of the AS path.	`detail`
`segments`	Length of the AS segment descriptor.	`detail`
`unique-count`	Number of unique autonomous systems (ASs) present in the AS path	`detail`
`references`	Path reference count.	`detail`

Sample Output

show as-path

```
user@host> show as-path
Total AS paths: 30382
  Bucket 0    Count: 36
    I
    14203 2914 174 31752 I
    14203 2914 701 21512 I
    14203 2914 1239 26632 I
    14203 2914 1239 29704 I
    14203 2914 4323 10248 I
    14203 2914 4766 23560 I
    14203 2914 6395 32776 I
    14203 2914 7911 11272 I
    14203 2914 12180 18440 I
    14203 2914 17408 17416 I
    14203 2914 701 702 24586 I
    14203 2914 1239 4657 9226 I
    14203 2914 1239 7132 16394 I
    14203 2914 1299 8308 34826 I
    14203 2914 3320 5603 28682 I
    14203 2914 3491 1680 33802 I
    14203 2914 3549 7908 27658 I
    14203 2914 3549 20804 30730 I
    14203 2914 7018 2687 9226 I
    14203 2914 174 9318 9318 23564 I
    14203 2914 701 3786 3786 23564 I
    14203 2914 701 4761 4795 9228 I
    14203 2914 1239 7132 5673 18444 I
    14203 2914 3491 20485 24588 24588 I
    14203 2914 5511 2200 1945 2060 I
    14203 2914 7911 14325 14325 14348 I
    14203 2914 701 4637 9230 9230 9230 I
    14203 2914 6395 14 14 14 14 I
    14203 2914 9299 6163 6163 6163 6163 9232 I
    14203 2914 3356 3356 3356 3356 3356 11955 21522 I
    14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
    14203 2914 174 30209 30222 30222 30222 ?
    14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
    14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
```

```
    14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
  Bucket 1    Count: 28
    14203 2914 35847 I
    14203 2914 174 19465 I
    14203 2914 174 35849 I
    14203 2914 2828 32777 I
    14203 2914 4323 14345 I
    14203 2914 4323 29705 I
    14203 2914 6395 32777 I

...
```

show as-path detail

```
user@host> show as-path detail
Total AS paths: 30410
  Bucket 0    Count: 36
    AS path: I
      domain 0, length 0, segments 0, unique-count 0, references 54
    AS path: 14203 2914 174 31752 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4, references 2
    AS path: 14203 2914 701 21512 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4, references 2
    AS path: 14203 2914 1239 26632 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 5, references 2
    AS path: 14203 2914 1239 29704 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4, references 2
    AS path: 14203 2914 4323 10248 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 6, references 2
    AS path: 14203 2914 4766 23560 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4, references 2
    AS path: 14203 2914 6395 32776 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 5, references 3
    AS path: 14203 2914 7911 11272 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 6, references 2
    AS path: 14203 2914 12180 18440 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 3, references 3
    AS path: 14203 2914 17408 17416 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 8, references 3
    AS path: 14203 2914 701 702 24586 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 4, references 3
```

```
    AS path: 14203 2914 1239 4657 9226 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 5, references 7
    AS path: 14203 2914 1239 7132 16394 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 7, references 2
    AS path: 14203 2914 1299 8308 34826 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 8, references 2
    AS path: 14203 2914 3320 5603 28682 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 4, references 2
    AS path: 14203 2914 3491 1680 33802 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 14, references 2
    AS path: 14203 2914 3549 7908 27658 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 6, references 2
    AS path: 14203 2914 3549 20804 30730 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 24, references 2
    AS path: 14203 2914 7018 2687 9226 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 4, references 3
    AS path: 14203 2914 174 9318 9318 23564 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references 2
    AS path: 14203 2914 701 3786 3786 23564 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references 2
    AS path: 14203 2914 701 4761 4795 9228 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references 14
    AS path: 14203 2914 1239 7132 5673 18444 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references 2
    AS path: 14203 2914 3491 20485 24588 24588 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references 4
    AS path: 14203 2914 5511 2200 1945 2060 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references 2
    AS path: 14203 2914 7911 14325 14325 14348 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references 2
    AS path: 14203 2914 701 4637 9230 9230 9230 I
      domain 1, neighbor as: 14203, length 7, segments 1, unique-count 4, references 3
    AS path: 14203 2914 6395 14 14 14 14 I
      domain 1, neighbor as: 14203, length 7, segments 1, unique-count 4, references 10
...
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show as-path summary | 17113

show as-path domain

IN THIS SECTION

- Syntax | 17111
- Syntax (EX Series Switches) | 17111
- Description | 17111
- Options | 17112
- Required Privilege Level | 17112
- Output Fields | 17112
- Sample Output | 17113
- Release Information | 17113

Syntax

```
show as-path domain
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show as-path domain
```

Description

Display autonomous system (AS) path domain information.

Options

none	(Optional) Display AS path domain information for all routing instances.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 167 on page 17112 lists the output fields for the `show as-path domain` command. Output fields are listed in the approximate order in which they appear

Table 167: show as-path domain Output Fields

Field Name	Field Description
Domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.
Primary	Primary AS number.
References	Path reference count.
***Number* Paths**	Number of known AS paths.
Flags	Information about the AS path: • **ASLoop**—Path contains an AS loop. • **Atomic**—Path includes the ATOMIC_AGGREGATE path attribute. • **Local**—Path was created by local aggregation. • **Master**—Path was created by the master routing instance.

Table 167: show as-path domain Output Fields *(Continued)*

Field Name	Field Description
Local AS	AS number of the local routing device.
Loops	How many times this AS number can appear in an AS path.

Sample Output

show as-path domain

```
user@host> show as-path domain
Domain: 1        Primary: 10458
    References:          3 Paths:      30383
    Flags: Master
    Local AS: 10458  Loops: 1
```

Release Information

Command introduced before Junos OS Release 7.4.

show as-path summary

IN THIS SECTION

- Syntax | 17114
- Syntax (EX Series Switches) | 17114
- Description | 17114
- Options | 17114
- Required Privilege Level | 17114
- Output Fields | 17115

Sample Output | 17115
Release Information | 17115

Syntax

```
show as-path summary
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show as-path summary
```

Description

Display autonomous system (AS) path summary information.

AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.

Options

none	(Optional) Display AS path summary information for all routing instances.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 168 on page 17115 lists the output fields for the `show as-path summary` command. Output fields are listed in the approximate order in which they appear.

Table 168: show as-path summary Output Fields

Field Name	Field Description
AS Paths	Number of AS paths.
Buckets	Number of hash buckets in use.
Max	Maximum number of AS path entries per bucket.
Min	Minimum number of AS path entries per bucket.
Avg	Average number of AS path entries per bucket.
Std deviation	Standard deviation of AS path entries per bucket.

Sample Output

show as-path summary

```
user@host> show as-path summary
AS Paths  Buckets     Max    Min    Avg   Std deviation
   30425     1024      95     12     29        6.481419
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show as-path | 17105

show auto-bandwidth

IN THIS SECTION

- Syntax | 17116
- Description | 17116
- Options | 17116
- Required Privilege Level | 17116
- Output Fields | 17117
- Sample Output | 17118
- Release Information | 17119

Syntax

```
show auto-bandwidth
history-log
traffic (sensor-name brief|detail)
```

Description

Display the auto-bandwidth information. You can view the traffic or the history log information in the output.

Options

history-log	Display the log of the sensor in the threshold activation period.
traffic (*sensor-name* brief | detail)	Display the traffic details and the configured auto-bandwidth parameters.

Required Privilege Level

view

Output Fields

Field Name	Field Description
`Collection Interval`	Time in seconds between two traffic sample collections.
`Adjust Interval`	Time interval after which the average bandwidth utilization is measured.
`Adjust threshold`	Percentage change threshold after which an action is triggered.
`Adjust Subscription`	Threshold for RSVP subscription in percentage.
`Pkt Recv`	Number of packets received.
`Byte Recv`	Number of bytes received.
`Poll Count`	Number of times the data was collected.
`Average`	Average bandwidth utilization across the SPRING traffic samples collected.
`Last Report`	Traffic statistics that was last reported indicating the number of packets received in the last report.
`Last Report time`	Time stamp of the last reporting time.
`Last Query Time`	Time stamp of the last query.
`Last Resp Time`	Time stamp of the last response time.
`Byte Bucket`	Count of bytes sent on the link measured periodically as per the configured collection interval.
`Packet Bucket`	Count of packets sent on the link measured periodically as per the configured collection interval.

Sample Output

show auto-bandwidth (traffic)

```
user@host> show auto-bandwidth traffic
Name                Stats-Id      packets     bytes
ge-0/0/2.0          0xc0000002          0         0
ge-0/0/0.0          0xc0000003          0         0
lo0.0               0xc0000004          0         0
user@host> show auto-bandwidth traffic detail
ge-0/0/0.0:
Collection Interval: 10, Adjust Interval : 200, Adjust threshold : 10
Adjust Subscription : 100
Pkt Recv: 1536   Byte Recv: 135168   Poll Count:29 Average:0
Last Report: 2829 Last Report time:Mon Jan  2 11:07:59 2017
Byte Bucket:
         0          0          0          0          0
         0          0          0          0          0
         0          0          0          0          0
         0          0          0          0          0
Packet Bucket:
         0          0          0          0          0
         0          0          0          0          0
         0          0          0          0          0
         0          0          0          0          0
```

show auto-bandwidth (history-log)

```
user@host> show auto-bandwidth history-log
Tue Jan  3 18:39:36
  ge-0/0/0.0:
   Threshold crossed firsttime, waiting for  5sec (adjust-interval - collection period)
  Byte Recv:4M Poll Count:1 Average:5Mbps
  Last Report:0
  Last Report time:Tue Jan  3 18:39:26 2017
  Chronological order: First bucket is the most recent

     2000000          0          0          0
Tue Jan  3 18:39:46
  ge-0/0/0.0:
```

```
   Threshold crossed early, waiting for  5sec
  Byte Recv:6M Poll Count:2 Average:1Mbps
  Last Report:0
  Last Report time:Tue Jan  3 18:39:26 2017
Chronological order: First bucket is the most recent

     2000000    2000000          0          0
Tue Jan  3 18:39:56
  ge-0/0/0.0:
   Threshold crossed early, waiting for  5sec
  Byte Recv:8M Poll Count:3 Average:1.5Mbps
  Last Report:0
  Last Report time:Tue Jan  3 18:39:26 2017
Chronological order: First bucket is the most recent

     2000000    2000000    2000000          0
Tue Jan  3 18:40:06
  ge-0/0/0.0:
  Threshold crossed, reporting 0 MBps to RSVP
  Byte Recv:10M Poll Count:4 Average:2 Mbps
  Last Report:0
  Last Report time:Tue Jan  3 18:39:26 2017
Chronological order: First bucket is the most recent

     2000000    2000000    2000000    2000000
```

Release Information

Command introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

- auto-bandwidth | 2682
- traffic-statistics (Protocols IS-IS) | 30048
- update-threshold-max-reservable | 30458
- clear isis spring traffic-statistics | 3673
- show isis spring interface traffic-statistics | 21779

show auto-configuration out-of-band

IN THIS SECTION

- Syntax | 17120
- Description | 17120
- Options | 17120
- Required Privilege Level | 17121
- Output Fields | 17121
- Sample Output | 17123
- Release Information | 17126

Syntax

```
show auto-configuration out-of-band
< brief | count | detail>
< all | pending>
< circuit-id>
< interface>
```

Description

Display information about ANCP-triggered out-of-band subscriber sessions. You can display all such subscribers or subscribers associated with a particular ACI or interface. You can also display information for sessions that are in the pending state. Sessions transition to the pending state when the Layer 2 wholesale session is authorized and assigned to an existing, nondefault routing instance, but subsequently profile instantiation to create the dynamic VLAN logical interface fails.

Options

brief | count | detail (Optional) Display the specified level of output. The default output level is the same as `brief`. The `count` level displays the number of out-of-band subscribers and sessions.

all (Optional) Display information for all out-of-band subscriber sessions, including pending sessions.

circuit-id (Optional) Limit the display to subscriber sessions on the specified access loop circuit identifier.

interface (Optional) Limit the display to out-of-band subscriber sessions on the specified interface.

pending (Optional) Display access lines in the pending state per routing instance.

Required Privilege Level

view

Output Fields

Table 169 on page 17121 lists the output fields for the `show auto-configuration out-of-band` command.

Table 169: show auto-configuration out-of-band Output Fields

Field Name	Field Description	Level of Output
`Agent Circuit ID`	Agent circuit ID (ACI, access loop circuit identifier) that maps the subscriber to an interface.	`detail`
`Agent Remote ID`	Agent remote ID (ARI, access loop remote identifier) that identifies the subscriber associated with an interface and therefore the access node.	`detail`
`Circuit Id`	Agent circuit ID (ACI, access loop circuit identifier) that maps the subscriber to an interface.	`brief` none
`Core IFL name`	Name of the core-facing logical interface.	`detail`
`Inner Vlan Map Id`	Inner VLAN tag that replaces the outer VLAN tag when the subscriber traffic is tunneled to the network service provider router.	`detail`

Table 169: show auto-configuration out-of-band Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Login Time`	Timestamp when the subscriber logged in.	`detail`
`Out-of-band subscribers`	Number of active out-of-band subscribers.	`brief` none
`Pending count`	Number of subscriber sessions in the routing instance that are in the pending state.	`brief` none
`Remote Id`	Agent remote ID (ARI, access loop remote identifier) for the subscriber.	`brief` none
`Routing-Instance`	Name of the routing instance for which the pending count is displayed.	`brief` none
`Session ID`	Subscriber session identifier.	All levels
`Subscriber IFL`	Name of the subscriber logical interface.	`brief` none
`Subscriber Interface`	Name of the subscriber logical interface.	`detail`
`total entries`	Total number of active plus pending subscriber sessions.	`brief` none
`VLAN`	VLAN ID of the subscriber' s incoming interface.	`brief` none
`VLAN ID`	VLAN ID of the subscriber' s incoming interface.	`detail`
`Vlan Map Id`	Trunk VLAN ID assigned to the core-facing physical interface within its input VLAN map.	`detail`

Sample Output

show auto-configuration out-of-band

```
user@host> show auto-configuration out-of-band
Circuit Id                  Remote Id               Subscriber IFL        Session ID  VLAN
203.0.13.3/0.0.0.0 eth 1/0  arn.1211.00000000002    ge-2/1/1.1689443721   1301         200
```

show auto-configuration out-of-band (All Sessions)

```
user@host> show auto-configuration out-of-band all
Circuit Id                  Remote Id              Subscriber IFL         Session ID  VLAN
203.0.13.2/0.0.0.0 eth 1/0  arn.4312.00000000002    ge-2/0/4.2547532212   1287        100
203.0.13.3/0.0.0.0 eth 1/0  arn.1211.00000000002    ge-2/1/1.1689443721   1301        200
203.0.13.4/0.0.0.0 eth 1/0  arn.5628.00000000002    ge-2/1/2.7697428544   1628        300
```

show auto-configuration out-of-band (All Sessions, Detail)

```
user@host> show auto-configuration out-of-band all detail
Agent Circuit ID: 203.0.13.2/0.0.0.0 eth 1/0
Agent Remote ID: arn.4312.00000000002
Subscriber interface: ge-2/0/4.2547532212
Session ID: 1287
VLAN ID: 100
Core IFL name: ae1.0
Vlan Map Id: 2
Inner Vlan Map Id: 1010
Login Time: 2018-06-25 09:05:14 EST
Agent Circuit ID: 203.0.13.3/0.0.0.0 eth 1/0
Agent Remote ID: arn.1211.00000000002
Subscriber interface: ge-2/1/1.1689443721
Session ID: 1301
VLAN ID: 200
Core IFL name: ae1.1
Vlan Map Id: 3
Inner Vlan Map Id: 1020
Login Time: 2018-06-25 09:11:21 EST
Agent Circuit ID: 203.0.13.4/0.0.0.0 eth 1/0
```

```
Agent Remote ID: arn.5628.00000000002
Subscriber interface: ge-2/1/2.7697428544
Session ID: 1628
VLAN ID: 300
Core IFL name: ae1.2
Vlan Map Id: 4
Inner Vlan Map Id: 1030
Login Time: 2018-06-25 09:11:58 EST
```

show auto-configuration out-of-band (Brief)

```
user@host> show auto-configuration out-of-band brief
Circuit Id                  Remote Id               Subscriber IFL        Session ID  VLAN
203.0.13.3/0.0.0.0 eth 1/0  arn.1211.00000000002    ge-2/1/1.1689443721   1301         200
```

show auto-configuration out-of-band (Circuit ID)

```
user@host> show auto-configuration out-of-band circuit-id "203.0.13.2/0.0.0.0 eth 1/0"
Circuit Id                  Remote Id               Subscriber IFL        Session ID  VLAN
203.0.13.2/0.0.0.0 eth 1/0  arn.4312.00000000002    ge-2/0/4.2547532212   1287         100
```

show auto-configuration out-of-band (Circuit ID, Detail)

```
user@host> show auto-configuration out-of-band circuit-id "203.0.13.2/0.0.0.0 eth 1/0" detail
Agent Circuit ID: 203.0.13.2/0.0.0.0 eth 1/0
Agent Remote ID: arn.4312.00000000002
Subscriber interface: ge-2/0/4.2547532212
Session ID: 1287
VLAN ID: 100
Core IFL name: ae1.0
Vlan Map Id: 2
Inner Vlan Map Id: 1010
Login Time: 2018-06-25 09:05:14 EST
```

show auto-configuration out-of-band (Count)

```
user@host> show auto-configuration out-of-band count
Out-of-band subscribers: 1, total entries: 1
```

show auto-configuration out-of-band (Detail)

```
user@host> show auto-configuration out-of-band detail
Agent Circuit ID: ACI.0001.aa1000000111
Agent Remote ID: remote-id
Subscriber interface: ge-1/0/0.3221225472
Session ID: 1
VLAN ID: 51
Core IFL name: ge-1/0/4.0
Vlan Map Id: 20
Inner Vlan Map Id: 1
Login Time: 2018-07-10 10:27:55 EDT
```

show auto-configuration out-of-band (Interface)

```
user@host> show auto-configuration out-of-band interface ge-2/0/4.2547532212
Circuit Id                  Remote Id               Subscriber IFL        Session ID  VLAN
203.0.13.2/0.0.0.0 eth 1/0  arn.4312.00000000002    ge-2/0/4.2547532212   1287         100
```

show auto-configuration out-of-band (Interface, Detail)

```
user@host> show auto-configuration out-of-band interface ge-2/0/4.2547532212 detail
Agent Circuit ID: 203.0.13.2/0.0.0.0 eth 1/0
Agent Remote ID: arn.4312.00000000002
Subscriber interface: ge-2/0/4.2547532212
Session ID: 1287
VLAN ID: 100
Core IFL name: ae1.0
Vlan Map Id: 2
Inner Vlan Map Id: 1010
Login Time: 2018-06-25 09:05:14 EST
```

show auto-configuration out-of-band (Pending Sessions)

```
user@host> show auto-configuration out-of-band pending
Circuit Id                  Remote Id               Subscriber IFL        Session ID  VLAN
203.0.13.2/0.0.0.0 eth 1/0  arn.4312.00000000002    ge-2/0/4.2547532212                100
```

Release Information

Command introduced in Junos OS Release 16.1R4.

`all`, `circuit-id`, `count`, and `interface` options introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

show backup-selection

IN THIS SECTION

- Syntax | 17126
- Description | 17127
- Options | 17127
- Required Privilege Level | 17127
- Output Fields | 17127
- Sample Output | 17129
- Release Information | 17132

Syntax

```
show backup-selection
<instance instance-name>
```

```
<interface interface-name>
<logical-system  logical-system-name>
<prefix prefix>
```

Description

Display the configured policies for each destination (IPv4 and IPv6) and a primary next-hop interface.

Options

`instance` *instance-name*	(Optional) Display configured policy for the routing instance.
`interface` *interface-name*	(Optional) Display configured policy for the interface.
`logical-system` *logical-system-name*	(Optional) Display configured policy for a particular logical system or for all logical systems.
`prefix` *prefix*	(Optional) Display configured policy for the destination.

Required Privilege Level

view

Output Fields

Table 170 on page 17127 describes the output fields for the **show backup-selection** command. Output fields are listed in the approximate order in which they appear.

Table 170: show backup-selection Output Fields

Field Name	Field Description
`Prefix`	Destination prefix .
`Interface`	Primary interface to reach the destination. • all — All the interfaces on the router which requires backup path. • <interface-name> — Specific primary interface in the backup path.

Table 170: show backup-selection Output Fields *(Continued)*

Field Name	Field Description
`admin-group exclude`	Specifies the administrative groups to be excluded. The backup path is not selected as the loop free alternative or backup next hop if any of the links in the path have any one of the listed administrative groups.
`admin-group include-all`	Requires each link in the backup path to have all the listed administrative groups in order to accept the path.
`admin-group include-any`	Requires each link in the backup path to have at least one of the listed administrative groups in order to select the path.
`admin-group preference`	Defines an ordered set of administrative groups that specifies the preference of the backup path. The leftmost element in the set is given the highest preference.
`nodes excluded`	Specifies the list of nodes to be excluded. The backup path that has a router from the list is not selected as the loop free alternative or backup next hop.
`nodes preference`	Defines an ordered set of nodes to be preferred. The backup path having the leftmost node is selected.
`node-tags excluded`	Specifies the backup selection to exclude the set of route tags in the backup path selection.
`node-tags preference`	Specifies the set of route tags in descending order of preference.
`protection-type`	Specifies the required protection type of the backup path. • node — Selects a backup path that provides node protection. • link — Selects the backup path that provides link protection. • node-link — Allows either node or link protection LFA where node-protection LFA is preferred over link-protection LFA.
`downstream paths only`	Selects the backup path that is a downstream path to the destination.

Table 170: show backup-selection Output Fields *(Continued)*

Field Name	Field Description
`srlg`	Evaluates common srlgs between the primary link and each link in the backup path. • strict — Rejects the backup path that has common srlgs between the primary link and any link in the backup path. • loose — Allows the backup path that has common srlgs between the primary link and any link in the backup path. The backup path with the fewer number of srlg collisions is preferred.
`B/W >= primary`	Uses backup next hop only if the bandwidth is greater than or equal to the primary next hop.
`root-metric`	Metric to one-hop node or remote router such as an RSVP backup label switched path (LSP) tail-end router.
`dest-metric`	Metric from one-hop node or remote router such as an RSVP backup label switched path (LSP) tail-end router to the final destination.
`metric evaluation order`	Defines the evaluation order of the metric (root and dest metrics) results.
`policy evaluation order`	Defines the evaluation order of the backup policy.

Sample Output

show backup-selection

```
user@host> show backup-selection

Prefix: 0.0.0.0/0
  Interface: all
    Admin-group exclude: c6
    Admin-group include-all: c1 c2
    Admin-group include-any: c3 c4
```

```
    Admin-group preference: c8
    Nodes excluded: 100.0.7.2
    Node preference: 100.2.6.2
    Node-tags excluded: 1004
    Node-tag preference: 1007
    Protection Type: Link, Downstream Paths Only: Enabled, SRLG: Strict, B/w >= Primary:
Enabled, Root-metric: lowest, Dest-metric: lowest
    Metric Evaluation Order: Root-metric, Dest-metric
    Policy Evaluation Order: Admin-group, SRLG, Bandwidth, Protection, Node, Node-Tag, Metric
  Interface: ge-1/2/5.0
    Admin-group include-all: c1
    Admin-group include-any: c2
    Admin-group preference: c4
    Nodes excluded: 10.218.32.0
    Node preference: 10.92.8.0
    Protection Type: Link, Downstream Paths Only: Enabled, SRLG: Strict, B/w >= Primary:
Enabled, Root-metric: lowest, Dest-metric: lowest
    Metric Evaluation Order: Root-metric, Dest-metric
    Policy Evaluation Order: Admin-group, SRLG, Bandwidth, Protection
Prefix: 10.150.0.0/16
  Interface: all
    Admin-group include-all: c1
    Admin-group include-any: c2
    Admin-group preference: c5
    Nodes excluded: 10.218.32.0
    Node preference: 10.92.8.0
    Node-tags excluded: 1004
    Node-tag preference: 1007
    Protection Type: Link, Downstream Paths Only: Enabled, SRLG: Loose, B/w >= Primary: Enabled,
Root-metric: lowest, Dest-metric: lowest
    Metric Evaluation Order: Dest-metric, Root-metric
    Policy Evaluation Order: Admin-group, SRLG, Bandwidth
Prefix: ::/0
  Interface: all
    Admin-group exclude: c2
    Admin-group include-all: c1 c3
    Admin-group include-any: c4 c5
    Admin-group preference: c6
    Node preference: 100.0.1.2
    Protection Type: Link, Downstream Paths Only: Disabled, SRLG: Strict, B/w >= Primary:
Enabled, Root-metric: lowest, Dest-metric: lowest
```

```
    Metric Evaluation Order: Dest-metric, Root-metric
    Policy Evaluation Order: Admin-group, SRLG, Bandwidth
```

show backup-selection instance instance-name

```
user@host> show backup-selection instance r

Prefix: 0.0.0.0/0
  Interface: all
    Admin-group include-any: c4
    Admin-group preference: c6 c0
    Nodes excluded: 100.0.4.2
    Node preference: 100.4.5.1
    Node-tags excluded: 1007
    Node-tag preference: 1004
    Protection Type: Link, Downstream Paths Only: Enabled, SRLG: Loose, B/w >= Primary: Enabled,
Root-metric: lowest, Dest-metric: lowest
    Metric Evaluation Order: Dest-metric, Root-metric
    Policy Evaluation Order: Admin-group, SRLG, Bandwidth
```

show backup-selection interface interface-name

```
user@host> show backup-selection interface ge-1/2/5.0

Prefix: 0.0.0.0/0
  Interface: ge-1/2/5.0
    Admin-group include-all: c1
    Admin-group include-any: c2
    Admin-group preference: c4
    Nodes excluded: 10.218.32.0
    Node preference: 10.92.8.0
    Protection Type: Link, Downstream Paths Only: Enabled, SRLG: Strict, B/w >= Primary:
Enabled, Root-metric: lowest, Dest-metric: lowest
    Metric Evaluation Order: Root-metric, Dest-metric
    Policy Evaluation Order: Admin-group, SRLG, Bandwidth, Protection
Prefix: 10.150.0.0/16
Prefix: ::/0
```

Release Information

Command introduced in Junos OS Release 14.1.

show bfd seamless session

IN THIS SECTION

- Syntax | 17132
- Description | 17132
- Options | 17132
- Required Privilege Level | 17133
- Output Fields | 17133
- Sample Output | 17133
- Release Information | 17134

Syntax

```
show bfd seamless session
 <brief | detail | extensive | summary>
 <logical-system (all | logical-system-name)>
```

Description

Display information about active seamless Bidirectional Forwarding Detection (SBFD) sessions.

Options

none	(Same as `brief`) Display information about active SBFD sessions.
brief | detail | extensive | summary	(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*) (Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

"show bfd seamless session" on page 17132 describes the output fields for the `show bfd seamless session` command. Output fields are listed in the approximate order in which they appear.

Table 171: show bfd seamless session Output Fields

Field Name	Field Description	Level of Output
Type	Type of the SBFD connection.	All levels
Address	Address on which the SBFD session is active.	All levels
State	State of the SBFD session: `Up`, `Down`, `Init` (initializing), or `Failing`.	All levels
Interval	Time interval used by the system to send the SBFD packets.	All levels
Table	Displays the type of routing table.	All levels
Discriminator	Authentication code used by the system to identify that SBFD session.	All levels

Sample Output

show bfd seamless session

```
user@host> show bfd seamless session
Type                 Descriminator  Table      Address              State    Interval
Local                     7007     default  0.0.0.0              Up       0.050
Local                      500     default  0.0.0.0              Up       0.050
Local   184549479(10.1.1.1)        default  0.0.0.0              Up       0.050
```

```
3 local sessions, 0 remote sessions
```

Release Information

Command introduced before Junos OS Release 22.4R1

The output field for Local Discriminator IP address introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

- clear bfd session | 3434
- Understanding BFD for Static Routes for Faster Network Failure Detection
- *Understanding BFD for OSPF*
- Understanding BFD for BGP
- Understanding Bidirectional Forwarding Detection Authentication for PIM
- Configuring BFD for PIM
- Understanding BFD for IS-IS

show bfd session

IN THIS SECTION

- Syntax | 17135
- Syntax (EX Series Switch and QFX Series) | 17135
- Description | 17135
- Options | 17135
- Required Privilege Level | 17136
- Output Fields | 17136
- Sample Output | 17142
- Release Information | 17144

Syntax

```
show bfd session
 <brief | detail | extensive | summary>
<address address>
<client rsvp-oam (brief | detail | extensive | summary) | vpls-oam  (brief | detail | extensive
| instance instance-name | summary)>
<discriminator discriminator>
<logical-system (all | logical-system-name)>
<prefix address>
<subscriber (address destination-address | discriminator discriminator | extensive)>
```

Syntax (EX Series Switch and QFX Series)

```
show bfd session
 <brief | detail | extensive | summary>
<address address>
<client rsvp-oam (brief | detail | extensive | summary) | vpls-oam  (brief | detail | extensive
| instance instance-name | summary)>
<discriminator discriminator>
<prefix address>
```

Description

Display information about active Bidirectional Forwarding Detection (BFD) sessions.

Options

none	(Same as `brief`) Display information about active BFD sessions.
brief | detail | extensive | summary	(Optional) Display the specified level of output.
address *address*	(Optional) Display information about the BFD session for the specified neighbor address.
client rsvp-oam (brief | detail | extensive | summary) | vpls-oam (brief | detail | extensive | instance *instance-name* | summary)	(Optional) Display information about RSVP-OAM or VPLS-OAM BFD sessions in the specified level of output. For VPLS-OAM, display the specified level of output or display information about all of the BFD sessions for the specified VPLS routing instance.

discriminator ***discriminator***	(Optional) Display information about the BFD session using the specified local discriminator.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
<subscriber (address *destination-address* | discriminator *discriminator* | extensive)>	(Optional) Display information about all BFD sessions for subscribers, or for a single BFD subscriber session with a particular destination address, or with a particular denominator.

Required Privilege Level

view

Output Fields

Table 172 on page 17136 describes the output fields for the `show bfd session` command. Output fields are listed in the approximate order in which they appear.

Table 172: show bfd session Output Fields

Field Name	Field Description	Level of Output
Address	Address on which the BFD session is active.	`brief` `detail` `extensive` none
State	State of the BFD session: `Up`, `Down`, `Init` (initializing), or `Failing`.	`brief` `detail` `extensive` none
Interface	Interface on which the BFD session is active.	`brief` `detail` `extensive` none
Detect Time	Negotiated time interval, in seconds, used to detect BFD control packets.	`brief` `detail` `extensive` none
Transmit Interval	Time interval, in seconds, used by the transmitting system to send BFD control packets.	`brief` `detail` `extensive` none

Table 172: show bfd session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Multiplier	Negotiated multiplier by which the time interval is multiplied to determine the detection time for the transmitting system.	`detail` `extensive`
Session up time	How long a BFD session has been established.	`detail` `extensive`
Client	Protocol or process for which the BFD session is active: `ISIS`, `OSPF`, `DHCP`, `Static`, or `VGD`.	`detail` `extensive`
TX interval	Time interval, in seconds, used by the host system to transmit BFD control packets.	`brief` `detail` `extensive` none
RX interval	Time interval, in seconds, used by the host system to receive BFD control packets.	`brief` `detail` `extensive` none
Authenticate	Indicates that BFD authentication is configured.	`detail` `extensive`
keychain	Name of the security authentication keychain being used by a specific client. BFD authentication information for a client is provided in a single line and includes the `keychain`, `algo`, and `mode` parameters. Multiple clients can be configured on a BFD session.	`extensive`
algo	BFD authentication algorithm being used for a specific client: `keyed-md5`, `keyed-sha-1`, `meticulous-keyed-md5`, `meticulous-keyed-sha-1`, or `simple-password`. BFD authentication information for a client is provided in a single line and includes the `keychain`, `algo`, and `mode` parameters. Multiple clients can be configured on a BFD session.	`extensive`

Table 172: show bfd session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
mode	Level of BFD authentication enforcement being used by a specific client: `strict` or `loose`. Strict enforcement indicates that authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated. BFD authentication information for a client is provided in a single line and includes the `keychain`, `algo`, and `mode` parameters. Multiple clients can be configured on a BFD session.	`extensive`
Local diagnostic	Local diagnostic information about failing BFD sessions. Following are the expected values for Local Diagnostic output field: • `None`—No diagnostic • `CtlExpire`—Control detection time expired • `EchoExpire`—Echo detection time expired • `NbrSignal`—Neighbor signalled session down • `FwdPlaneReset`—Forwarding plane reset • `PathDown`—Path down • `ConcatPathDown`—Concatenated path down • `AdminDown`—Administratively down	`detail extensive`

Table 172: show bfd session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Remote diagnostic	Remote diagnostic information about failing BFD sessions. Following are the expected values for Remote Diagnostic output field: • `None`—No diagnostic • `CtlExpire`—Control detection time expired • `EchoExpire`—Echo detection time expired • `NbrSignal`—Neighbor signalled session down • `FwdPlaneReset`—Forwarding plane reset • `PathDown`—Path down • `ConcatPathDown`—Concatenated path down • `AdminDown`—Administratively down	`detail extensive`
Remote state	Reports whether the remote system's BFD packets have been received and whether the remote system is receiving transmitted control packets.	`detail extensive`
Version	BFD version: `0` or `1`.	`extensive`
Replicated	The `replicated` flag appears when the device has a backup Routing Engine installed and the BFD session has been replicated in the backup Routing Engine.	`detail extensive`
Min async interval	Minimum amount of time, in seconds, between asynchronous control packet transmissions across the BFD session.	`extensive`
Min slow interval	Minimum amount of time, in seconds, between synchronous control packet transmissions across the BFD session.	`extensive`
Adaptive async TX interval	Transmission interval being used because of adaptation.	`extensive`

Table 172: show bfd session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
RX interval	Minimum required receive interval.	`extensive`
Local min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the local system.	`extensive`
Local min RX interval	Minimum amount of time, in seconds, between control packet detections on the local system.	`extensive`
Remote min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the remote system.	`extensive`
Remote min TX interval	Minimum amount of time, in seconds, between control packet detections on the remote system.	`extensive`
Threshold transmission interval	Threshold for notification if the transmission interval increases.	`extensive`
Threshold for detection time	Threshold for notification if the detection time increases.	`extensive`
Local discriminator	Authentication code used by the local system to identify that BFD session.	`extensive`
Remote discriminator	Authentication code used by the remote system to identify that BFD session.	`extensive`
Echo mode	Information about the state of echo transmissions on the BFD session.	`extensive`
Prefix	LDP FEC address associated with the BFD session.	All levels

Table 172: show bfd session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Egress, Destination	Displays the LDP FEC destination address. This field is displayed only on a router at the egress of an LDP FEC, where the BFD session has an LDP Operation, Administration, and Maintenance (OAM) client.	All levels
Remote is control-plane independent	The BFD session on the remote peer is running on its Packet Forwarding Engine. In this case, when the remote node undergoes a graceful restart, the local peer can help the remote peer with the graceful restart. The following BFD sessions are not distributed to the Packet Forwarding Engine: tunnel-encapsulated sessions, and sessions over integrated routing and bridging (IRB) interfaces.	`extensive`
Authentication	Summary status of BFD authentication: • `status`—`enabled/active` indicates authentication is configured and active. `enabled/inactive` indicates authentication is configured but not active. This only occurs when the remote end of the session does not support authentication and loose checking is configured. • `keychain`—Name of the security authentication keychain associated with the specified BFD session. • `algo`—BFD authentication algorithm being used: `keyed-md5`, `keyed-sha-1`, `meticulous-keyed-md5`, `meticulous-keyed-sha-1`, or `simple-password`. • `mode`—Level of BFD authentication enforcement: `strict` or `loose`. Strict enforcement indicates authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated. This information is only shown if BFD authentication is configured.	`extensive`
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).	`detail extensive`
sessions	Total number of active BFD sessions.	All levels

Table 172: show bfd session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
clients	Total number of clients that are hosting active BFD sessions.	All levels
Cumulative transmit rate	Total number of BFD control packets transmitted per second on all active sessions.	All levels
Cumulative receive rate	Total number of BFD control packets received per second on all active sessions.	All levels
Multi-hop, min-recv-TTL	Minimum time to live (TTL) accepted if the session is configured for multihop.	`extensive`
route table	Route table used if the session is configured for multihop.	`extensive`
local address	Local address of the source used if the session is configured for multihop. The source IP address for outgoing BFD packets from the egress side of an MPLS BFD session is based on the outgoing interface IP address.	`extensive`

Sample Output

show bfd session

```
user@host> show bfd session
                                                      Transmit
Address              State     Interface     Detect Time  Interval  Multiplier
10.9.1.33            Up        so-7/1/0.0          0.600     0.200      3
10.9.1.29            Up        ge-4/0/0.0          0.600     0.200      3

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

show bfd session brief

The output for the `show bfd session brief` command is identical to that for the `show bfd session` command.

show bfd session detail

```
user@host> show bfd session detail
                                                    Transmit
Address              State     Interface     Detect Time  Interval  Multiplier
10.9.1.33            Up        so-7/1/0.0          0.600     0.200      3
  Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3
  Session up time 3d 00:34:02
  Local diagnostic None, remote diagnostic None
  Remote state Up, version 1
  Replicated
10.9.1.29            Up        ge-4/0/0.0          0.600     0.200      3
  Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3
  Session up time 3d 00:29:04, previous down time 00:00:01
  Local diagnostic NbrSignal, remote diagnostic AdminDown
  Remote state Up, version 1

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

show bfd session extensive

```
user@host> show bfd session extensive                                                Detect
Transmit
Address                  State     Interface      Time     Interval  Multiplier
10.9.1.33                 Up        et-1/0/1.0     6.000     2.000        3
 Client PRPD:6, TX interval 1.000, RX interval 1.000
 Session up time 00:00:01
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated
 Session type: Single hop BFD (Echo Lite)
 Local discriminator 16
 Echo TX interval 2.000, echo detection interval 6.000
  Session ID: 0
```

```
1 sessions, 1 clients
Cumulative transmit rate 0.5 pps, cumulative receive rate 0.5 pps
```

Release Information

Command introduced before Junos OS Release 7.4.

Options `discriminator` and `address` introduced in Junos OS Release 8.2.

Option `prefix` introduced in Junos OS Release 9.0.

Option `client` introduced in Junos OS Release 12.3R3.

Option `subscriber` introduced in Junos OS Release 15.1 for the MX Series.

RELATED DOCUMENTATION

- clear bfd session | 3434
- Understanding BFD for Static Routes for Faster Network Failure Detection
- *Understanding BFD for OSPF*
- Understanding BFD for BGP
- Understanding Bidirectional Forwarding Detection Authentication for PIM
- Configuring BFD for PIM
- Understanding BFD for IS-IS

show bfd subscriber session

IN THIS SECTION

- Syntax | 17145
- Description | 17145
- Options | 17145
- Required Privilege Level | 17145
- Output Fields | 17145
- Sample Output | 17148

- Release Information | **17151**

Syntax

```
show bfd subscriber session
 <brief | detail | extensive | summary>
```

Description

Display information about active Bidirectional Forwarding Detection (BFD) subscriber sessions.

Options

none	(Same as `brief`) Display information about active BFD subscriber sessions.
brief | detail | extensive | summary	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 173 on page 17145 describes the output fields for the `show bfd subscriber session` command. Output fields are listed in the approximate order in which they appear.

Table 173: show bfd subscriber session Output Fields

Field Name	Field Description	Level of Output
Address	IP Address on which the BFD subscriber session is active.	All levels

Table 173: show bfd subscriber session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the BFD subscriber session: `Up`, `Down`, `Init` (initializing), or `Failing`.	All levels
Interface	Interface on which the BFD subscriber session is active.	All levels
Detect Time	Negotiated time interval, in seconds, used to detect BFD control packets.	All levels
Transmit Interval	Time interval, in seconds, used by the transmitting system to send BFD control packets.	All levels
Multiplier	Negotiated multiplier by which the time interval is multiplied to determine the detection time for the transmitting system.	All levels
TX interval	Time interval, in seconds, used by the host system to transmit BFD control packets.	`detail extensive`
RX interval	Time interval, in seconds, used by the host system to receive BFD control packets.	`detail extensive`
Local diagnostic	Local diagnostic information about failing BFD subscriber sessions.	`detail extensive`
Remote diagnostic	Remote diagnostic information about failing BFD subscriber sessions.	`detail extensive`
Remote state	Indication that the remote system's BFD packets have been received and whether the remote system is receiving transmitted control packets.	`detail extensive`
Version	BFD version: `0` or `1`.	`detail extensive`
routing table index	Value of the routing table index. A value of 0 (zero) denotes a route in the default routing table managed by enhanced subscriber management.	`detail extensive`

Table 173: show bfd subscriber session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Min async interval	Minimum amount of time, in seconds, between asynchronous control packet transmissions across the BFD subscriber session.	`extensive`
min slow interval	Minimum amount of time, in seconds, between synchronous control packet transmissions across the BFD subscriber session.	`extensive`
Adaptive async TX interval	Transmission interval being used because of adaptation.	`extensive`
RX interval	Receive interval being used because of adaptation.	extensive
Local min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the local system.	`extensive`
minimum RX interval	Minimum amount of time, in seconds, between control packet detections on the local system.	`extensive`
Remote min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the remote system.	`extensive`
min RX interval	Minimum amount of time, in seconds, between control packet detections on the remote system.	`extensive`
Local discriminator	Authentication code used by the local system to identify that BFD subscriber session.	`extensive`
Remote discriminator	Authentication code used by the remote system to identify that BFD subscriber session.	`extensive`

Sample Output

show bfd subscriber session

```
user@host> show bfd subscriber session
                                                Detect  Transmit
Address              State     Interface        Time    Interval  Multiplier
203.0.113.2          Up        demux0.3221225503  90.000  30.000        3
203.0.113.6          Up        demux0.3221225504  90.000  30.000        3
203.0.113.10         Up        demux0.3221225505  90.000  30.000        3
203.0.113.14         Up        demux0.3221225506  90.000  30.000        3
203.0.113.18         Up        demux0.3221225507  90.000  30.000        3
```

show bfd subscriber session brief

The output for the `show bfd subscriber session brief` command is identical to that for the `show bfd subscriber session` command.

show bfd subscriber session detail

```
user@host> show bfd subscriber session detail
                                               Detect  Transmit
Address              State   Interface         Time    Interval  Multiplier
203.0.113.2          Up      demux0.3221225503  90.000  30.000       3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
203.0.113.6          Up      demux0.3221225504  90.000  30.000       3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
203.0.113.10         Up      demux0.3221225505  90.000  30.000       3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
203.0.113.14         Up      demux0.3221225506  90.000  30.000       3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
203.0.113.18         Up      demux0.3221225507  90.000  30.000       3
```

```
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
```

command-name

show bfd subscriber session extensive

```
user@host> show bfd subscriber session extensive
                                                        Detect   Transmit
Address                  State     Interface            Time     Interval  Multiplier
203.0.113.2              Up        demux0.3221225503    90.000    30.000        3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
 Min async interval 30.000, min slow interval 30.000
 Adaptive async TX interval 30.000, RX interval 30.000
 Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3
 Remote min TX interval 30.000, min RX interval 30.000, multiplier 3
 Local discriminator 28, remote discriminator 1073741825

                                                        Detect   Transmit
Address                  State     Interface            Time     Interval  Multiplier
203.0.113.6              Up        demux0.3221225504    90.000    30.000        3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
 Min async interval 30.000, min slow interval 30.000
 Adaptive async TX interval 30.000, RX interval 30.000
 Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3
 Remote min TX interval 30.000, min RX interval 30.000, multiplier 3
 Local discriminator 29, remote discriminator 1073741826

                                                        Detect   Transmit
Address                  State     Interface            Time     Interval  Multiplier
203.0.113.10             Up        demux0.3221225505    90.000    30.000        3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
 Min async interval 30.000, min slow interval 30.000
```

```
 Adaptive async TX interval 30.000, RX interval 30.000
 Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3
 Remote min TX interval 30.000, min RX interval 30.000, multiplier 3
 Local discriminator 30, remote discriminator 1073741827

                                                        Detect   Transmit
Address                  State     Interface            Time     Interval  Multiplier
203.0.113.14             Up        demux0.3221225506    90.000    30.000        3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
 Min async interval 30.000, min slow interval 30.000
 Adaptive async TX interval 30.000, RX interval 30.000
 Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3
 Remote min TX interval 30.000, min RX interval 30.000, multiplier 3
 Local discriminator 31, remote discriminator 1073741828

                                                        Detect   Transmit
Address                  State     Interface            Time     Interval  Multiplier
203.0.113.18             Up        demux0.3221225507    90.000    30.000        3
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 routing table index 0
 Min async interval 30.000, min slow interval 30.000
 Adaptive async TX interval 30.000, RX interval 30.000
 Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3
 Remote min TX interval 30.000, min RX interval 30.000, multiplier 3
 Local discriminator 32, remote discriminator 1073741829
```

show bfd subscriber session summary

```
user@host> show bfd subscriber session summary
                                                Detect  Transmit
Address              State     Interface        Time    Interval  Multiplier
203.0.113.2          Up        demux0.3221225503  90.000  30.000        3
203.0.113.6          Up        demux0.3221225504  90.000  30.000        3
203.0.113.10         Up        demux0.3221225505  90.000  30.000        3
203.0.113.14         Up        demux0.3221225506  90.000  30.000        3
203.0.113.18         Up        demux0.3221225507  90.000  30.000        3
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

clear bfd session | 3434

Understanding BFD for Static Routes for Faster Network Failure Detection

show bgp bmp

IN THIS SECTION

- Syntax | 17151
- Description | 17151
- Options | 17152
- Required Privilege Level | 17152
- Output Fields | 17152
- Sample Output | 17153
- Release Information | 17153

Syntax

```
show bgp bmp
```

Description

Display information about the BGP Monitoring Protocol (BMP).

Options

statistics (Optional) Show BMP statistics

- **socket-write**
  Show BMP socket write statistics
- **route-monitoring**
  Show BMP route monitoring statistics
- **end-of-rib**
  Show BMP end of RIB statistics

Required Privilege Level

view

Output Fields

Table 174 on page 17152 lists the output fields for the `show bgp bmp` command. Output fields are listed in the approximate order in which they appear.

Table 174: show bgp bmp Output Fields

Field Name	Field Description
`BMP station address/port`	IP address and port number of the monitoring station to which BGP Monitoring Protocol (BMP) statistics are sent.
`BMP session state`	Status of the BMP session: `UP` or `DOWN`.
`Statistics timeout`	Amount of time, in seconds, between transmissions of BMP data to the monitoring station.

Sample Output

show bgp bmp

```
user@host> show bgp bmp
  BMP station address/port: 172.24.24.157+5454
  BMP session state: DOWN
    Statistics timeout: 15
```

Release Information

Command introduced in Junos OS Release 9.5.

show bgp group

IN THIS SECTION

- Syntax | 17154
- Syntax (EX Series Switch and QFX Series) | 17154
- Description | 17154
- Options | 17154
- Required Privilege Level | 17155
- Output Fields | 17155
- Sample Output | 17161
- Release Information | 17166

Syntax

```
show bgp group
<brief | detail | summary>
<group-name>
<exact-instance instance-name>
<instance instance-name>
<logical-system (all | logical-system-name)>
<rtf>
```

Syntax (EX Series Switch and QFX Series)

```
show bgp group
<brief | detail | summary>
<group-name>
<exact-instance instance-name>
<instance instance-name>
```

Description

Display information about the configured BGP groups.

Options

none	Display group information about all BGP groups.
brief | detail | summary	(Optional) Display the specified level of output.
group-name	(Optional) Display group information for the specified group.
exact-instance *instance-name*	(Optional) Display information for the specified instance only.
instance *instance-name*	(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, **cust1**, **cust11**, and **cust111** are all displayed when you run the `show bgp group instance cust1` command). The instance name can be primary for the main instance, or any valid configured instance name or its prefix.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
rtf	(Optional) Display BGP group route targeting information.

Required Privilege Level

view

Output Fields

Table 175 on page 17155 describes the output fields for the `show bgp group` command. Output fields are listed in the approximate order in which they appear.

Table 175: show bgp group Output Fields

Field Name	Field Description	Level of Output
Group Type or **Group**	Type of BGP group: **Internal** or **External**.	All levels
group-index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.	**rtf detail**
AS	AS number of the peer. For internal BGP (IBGP), this number is the same as **Local AS**.	**brief detail** none
Local AS	AS number of the local routing device.	**brief detail** none
Name	Name of a specific BGP group.	**brief detail** none
Options	The Network Layer Reachability Information (NLRI) format used for BGP VPN multicast.	**none** none

Table 175: show bgp group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Index	Unique index number of a BGP group.	**brief detail** none
Flags	Flags associated with the BGP group. This field is used by Juniper Networks customer support.	**brief detail** none
`BGP-Static Advertisement Policy`	Policies configured for the BGP group with the `advertise-bgp-static policy` statement.	`brief` none
`Remove-private options`	Options associated with the `remove-private` statement.	**brief detail** none
Holdtime	Maximum number of seconds allowed to elapse between successive keepalive or update messages that BGP receives from a peer in the BGP group, after which the connection to the peer is closed and routing devices through that peer become unavailable.	**brief detail** none
Export	Export policies configured for the BGP group with the `export` statement.	**brief detail** none
Optimal Route Reflection	Client nodes (primary and backup) configured in the BGP group.	**brief detail** none
MED tracks IGP metric update delay	Time, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire	All levels
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.	**brief detail** none
Total peers	Total number of peers in the group.	**brief detail** none

Table 175: show bgp group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Established	Number of peers in the group that are in the established state.	All levels

Active/Received/Accepted/ Damped	Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established in the main routing device or in a routing instance.
	• If a peer is not established, the field shows the state of the peer session: **Active**, **Connect**, or **Idle**.
	• If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the **inet.0** (main) and **inet.2** (multicast) routing tables. For example, **8/10/10/2 and 2/4/4/0** indicate the following:
	  • 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the **inet.0** routing table.
	  • 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the **inet.2** routing table.
		summary

ip-addresses	List of peers who are members of the group. The address is followed by the peer's port number.	All levels
Allowed Unconfigured Peer Addresses	Configured `peer-limit` and the current `peer-limit` for a dynamic neighbor.	**detail**
Route Queue Timer	Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed.	**detail**

Table 175: show bgp group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Route Queue	Number of prefixes that are queued up for sending to the peers in the group.	**detail**
inet.*number*	Number of active, received, accepted, and damped routes in the routing table. For example, **inet.0: 7/10/9/0** indicates the following: • 7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the **inet.0** routing table.	none

Table 175: show bgp group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Table inet.*number*	Information about the routing table. • **Received prefixes**—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • **Active prefixes**—Number of prefixes received from the peer that are active in the routing table. • **Suppressed due to damping**—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. • **Advertised prefixes**—Number of prefixes advertised to a peer. • **Received external prefixes**—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table. • **Active external prefixes**—Number of prefixes received from the EBGP peers that are active in the routing table. • **Externals suppressed**—Number of routes received from EBGP peers currently inactive because of damping or other reasons. • **Received internal prefixes**—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table. • **Active internal prefixes**—Number of prefixes received from the IBGP peers that are active in the routing table. • **Internals suppressed**—Number of routes received from IBGP peers currently inactive because of damping or other reasons. • **RIB State**—Status of the graceful restart process for this routing table: **BGP restart is complete**, **BGP restart in**	**detail**

Table 175: show bgp group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	progress, **VPN restart in progress**, or **VPN restart is complete**.	
Groups	Total number of groups.	All levels
Peers	Total number of peers.	All levels
External	Total number of external peers.	All levels
Internal	Total number of internal peers.	All levels
Down peers	Total number of unavailable peers.	All levels
Flaps	Total number of flaps that occurred.	All levels
Table	Name of a routing table.	**brief**, none
Tot Paths	Total number of routes.	**brief**, none
Act Paths	Number of active routes.	**brief**, none
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	**brief**, none
History	Number of withdrawn routes stored locally to keep track of damping history.	**brief**, none
Damp State	Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs.	**brief**, none

Table 175: show bgp group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Pending	Routes being processed by the BGP import policy.	**brief**, none
Group	Group the peer belongs to in the BGP configuration.	**detail**
Receive mask	Mask of the received target included in the advertised route.	**detail**
Entries	Number of route entries received.	**detail**
Target	Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer.	**detail**
Mask	Mask which specifies that the peer receive routes with the given route target.	**detail**

Sample Output

show bgp group

```
user@host> show bgp group
Group Type: Internal AS: 65200 Local AS: 65200
Name: ibgp Index: 0 Flags: <>
Options: Preference LocalAddress Cluster AddressFamily Refresh
```

show bgp group

```
user@host> show bgp group
Group Type: Internal    AS: 65101                   Local AS: 65101
  Name: ibgp            Index: 2                   Flags: Export Eval
  Total assigned update threads: 5
```

```
    bgp-updio-4:    50
    bgp-updio-5:    50
    bgp-updio-6:    50
    bgp-updio-3:    50
    bgp-updio-1:    50
  Holdtime: 0
  Optimal Route Reflection: igp-primary 10.1.1.1, igp-backup 10.1.2.1
  Total peers: 1        Established: 1
  10.1.1.2+179
  Trace options:  all
  Trace file: /var/log/bgp-log size 10485760 files 10
  bgp.l3vpn.2: 0/0/0/0
  vpn-1.inet.2: 0/0/0/0

Group Type: Internal    AS: 65101                   Local AS: 65101
  Name: ibgp            Index: 3                   Flags: Export Eval
  Options: RFC6514CompliantSafi129
  Holdtime: 0
  Optimal Route Reflection: igp-primary 10.1.1.1, igp-backup 10.1.2.1
  Total peers: 1        Established: 1
  10.1.1.5+61698
  Trace options:  all
  Trace file: /var/log/bgp-log size 10485760 files 10
  bgp.l3vpn.2: 2/2/2/0
  vpn-1.inet.2: 2/2/2/0

Groups: 2  Peers: 2    External: 0    Internal: 2    Down peers: 0   Flaps: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.2
                       2          2          0          0          0          0
vpn-1.inet.0
                       0          0          0          0          0          0
vpn-1.inet.2
                       2          2          0          0          0          0
vpn-1.inet6.0
                       0          0          0          0          0          0
vpn-1.mdt.0
                       0          0          0          0          0          0
```

show bgp group brief

```
user@host> show bgp group brief
 Groups: 2  Peers: 2    External: 0    Internal: 2    Down peers: 1   Flaps: 0
 Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
 inet.0
                        0          0          0          0          0          0
 bgp.l3vpn.0
                        0          0          0          0          0          0
 bgp.rtarget.0
                        2          0          0          0          0          0
```

show bgp group detail

```
user@host> show bgp group detail
Group Type: Internal    AS: 65100                      Local AS: 65100
  Name: ibgp            Index: 0                   Flags: <Export Eval>
  Holdtime: 0
  Optimal Route Reflection: igp-primary 10.1.1.1, igp-backup 10.1.2.1
  Total peers: 3        Established: 0
  10.22.0.2
  10.22.0.8
  10.22.0.5

Groups: 1  Peers: 3    External: 0    Internal: 3    Down peers: 3   Flaps: 3
  Table bgp.l3vpn.0
    Received prefixes:            0
    Accepted prefixes:            0
    Active prefixes:              0
    Suppressed due to damping:    0
    Received external prefixes:   0
    Active external prefixes:     0
    Externals suppressed:         0
    Received internal prefixes:   0
    Active internal prefixes:     0
    Internals suppressed:         0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
  Table bgp.mdt.0
    Received prefixes:            0
```

```
    Accepted prefixes:            0
    Active prefixes:              0
    Suppressed due to damping:    0
    Received external prefixes:   0
    Active external prefixes:     0
    Externals suppressed:         0
    Received internal prefixes:   0
    Active internal prefixes:     0
    Internals suppressed:         0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
  Table VPN-A.inet.0
    Received prefixes:            0
    Accepted prefixes:            0
    Active prefixes:              0
    Suppressed due to damping:    0
    Received external prefixes:   0
    Active external prefixes:     0
    Externals suppressed:         0
    Received internal prefixes:   0
    Active internal prefixes:     0
    Internals suppressed:         0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
  Table VPN-A.mdt.0
    Received prefixes:            0
    Accepted prefixes:            0
    Active prefixes:              0
    Suppressed due to damping:    0
    Received external prefixes:   0
    Active external prefixes:     0
    Externals suppressed:         0
    Received internal prefixes:   0
    Active internal prefixes:     0
    Internals suppressed:         0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
```

show bgp group rtf detail

```
user@host> show bgp group rtf detail
 Group: internal (group-index: 0)
     Receive mask: 00000002
     Table: bgp.rtarget.0                                    Entries: 2
         Target                           Mask
         10:100/64                       00000002
         10:201/64                       (Group)
 Group: internal (group-index: 1)
     Table: bgp.rtarget.0                                    Entries: 1
         Target                           Mask
         200:201/64                       (Group)
```

show bgp group summary

```
user@host> show bgp group summary
Group        Type       Peers     Established    Active/Received/Accepted/Damped
ibgp         Internal   3         0

Groups: 1  Peers: 3    External: 0    Internal: 3    Down peers: 3   Flaps: 3
  bgp.l3vpn.0      : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  bgp.mdt.0        : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  VPN-A.inet.0     : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  VPN-A.mdt.0      : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
```

show bgp group

```
user@host> show bgp group
Group Type: Internal    AS: 2                      Local AS: 2
  Name: toR1            Index: 0                   Flags: <Export Eval>
  Options: <Cluster>
  Holdtime: 0
+ Optimal Route Reflection: export orr-policy
  Total peers: 1        Established: 1
  1.1.1.1+179
  inet.0: 0/0/0/0
```

Release Information

Command introduced before Junos OS Release 7.4.

exact-instance option introduced in Junos OS Release 11.4.

From Junos OS release 18.4 onwards, `show bgp group` *`group-name`* does an exact match and displays groups with names matching exactly with that of the specified group-name. For all Junos OS releases preceding 18.4, the implementation was performed using the prefix matches (example: if there are two groups grp1, grp2 and the CLI command `show bgp group grp` was issued, then both grp1, grp2 were displayed).

show bgp group output-queues

IN THIS SECTION

- Syntax | **17166**
- Description | **17166**
- Options | **17167**
- Required Privilege Level | **17167**
- Output Fields | **17167**
- Sample Output | **17168**
- Release Information | **17170**

Syntax

```
show bgp group output-queues
<group-name>
<fabric>
<logical-system>
```

Description

Show per group summaries of BGP prioritized output queues. The output includes the number of tokens assigned per queue and the number of routes currently queued within each prioritized queue.

Options

none Display output queue summaries for all BGP groups defined in the system.

group-name Limit the display of queue summaries to the specified group.

fabric Display output queue summaries for the specified fabric.

logical-system Display output queue information within a specified logical system or for all logical systems.

Required Privilege Level

view

Output Fields

Table 176 on page 17167 describes the output fields for the `show bgp group output-queues` command. Output fields are listed in the approximate order in which they appear. Some output fields are self-explanatory and so are not shown in the table.

Table 176: show bgp group output-queues Output Fields

Field Name	Field Description
Index	Group's index number.
Options	Options set within the BGP group definition.
NLRI	Address family for which BGP route prioritization has been implemented within the BGP group.
OutQ	Output priority queue designated for this address family within this group.
RRQ	Route refresh priority queue designated for this address family within this group.
WDQ	Withdraw priority queue designated for this address family within this group.

Table 176: show bgp group output-queues Output Fields *(Continued)*

Field Name	Field Description
Class	Shows the name of the priority queues. There are always 16 numbered priority queues and the expedited queue for a total of 17 priority queues.
Tokens	Shows the number of tokens assigned to each priority queue.
Total Routes	Shows the number of routes currently in each priority queue (class).

Sample Output

show bgp group output-queues

```
user@host> show bgp group output-queues
Group Type: Internal    AS: 64512                  Local AS: 64512
  Name: bgp-group-1             Index: 0                   Flags: <Export Eval>
  Export: [ match-all ]
  Options: <LocalAS>
  Holdtime: 0
  NLRI inet-unicast:
    OutQ: priority 1 RRQ: priority 1 WDQ: priority 1
    Class         Tokens  Total Routes
    ------------  ------  ------------
    Priority   1       1             0
    Priority   2      10             0
    Priority   3      15             0
    Priority   4      20             0
    Priority   5      25             0
    Priority   6      30             0
    Priority   7      35             0
    Priority   8      40             0
    Priority   9      45             0
    Priority  10      50             0
    Priority  11      55             0
    Priority  12      60             0
    Priority  13      65             0
    Priority  14      70             0
```

```
    Priority  15      75             0
    Priority  16      80             0
    Expedited        100             0
  Total peers: 1        Established: 1
  192.0.2.2+179
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                                  0          0
Group Type: External                                Local AS: 69
  Name: reflector       Index: 1                    Flags: <Export Eval>
  Options: <Multihop LocalAS>
  Holdtime: 0
  NLRI inet-unicast:
    OutQ: priority 1 RRQ: priority 1 WDQ: priority 1
    Class         Tokens  Total Routes
    ------------  ------  ------------
    Priority   1       1             0
    Priority   2      10             0
    Priority   3      15             0
    Priority   4      20             0
    Priority   5      25             0
    Priority   6      30             0
    Priority   7      35             0
    Priority   8      40             0
    Priority   9      45             0
    Priority  10      50             0
    Priority  11      55             0
    Priority  12      60             0
    Priority  13      65             0
    Priority  14      70             0
    Priority  15      75             0
    Priority  16      80             0
    Expedited        100             0
  Total peers: 1        Established: 1
  192.0.2.71+179
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                              584195          0

Groups: 2  Peers: 2    External: 1    Internal: 1    Down peers: 0   Flaps: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                   584198     584195          0          0          0          0
```

```
inet.3
                        0          0          0          0          0          0
bgp.l3vpn.0
                        0          0          0          0          0          0
```

show bgp group output-queues <group-name>

```
user@host> show bgp group output-queues bgp-group-2
Group Type: External                                Local AS: 102
  Name: bgp-group-2             Index: 2                   Flags: <>
  Options: <LocalAS>
  Holdtime: 0
  NLRI inet-unicast:
    OutQ: priority 6 RRQ: priority 3 WDQ: priority 3
    Class         Tokens  Total Routes
    ------------  ------  ------------
    Priority   1       1             0
    Priority   2       1             0
    Priority   3      20             0
    Priority   4       1             0
    Priority   5       1             0
    Priority   6      30             0
    Priority   7       1             0
    Priority   8       1             0
    Priority   9      50             0
    Priority  10       1             0
    Priority  11       1             0
    Priority  12       1             0
    Priority  13       1             0
    Priority  14       1             0
    Priority  15       1             0
    Priority  16       1             0
    Expedited          1             0
  Total peers: 1        Established: 1
  192.0.2.3+179
```

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

show bgp neighbor | 17175

show bgp group traffic-statistics

IN THIS SECTION

- Syntax | 17171
- Description | 17171
- Options | 17171
- Required Privilege Level | 17172
- Output Fields | 17172
- Sample Output | 17173
- Release Information | 17174

Syntax

```
show bgp group traffic-statistics
<brief | detail>
<group-name>
<labeled-path label label>
<logical-system (all | logical-system-name)>
```

Description

Display the traffic statistics for configured Border Gateway Protocol (BGP) groups.

Options

none	Display traffic statistics for all BGP groups.
brief | detail	(Optional) Display the specified level of output.

group-name	(Optional) Display BGP traffic statistics for only the specified group.
label-path	(Optional) Display labeled unicast traffic statistics at the ingress.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 177 on page 17172 describes the output fields for the `show bgp group traffic-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 177: show bgp group traffic-statistics Output Fields

Field Name	Field Description
`Group name`	Name of a specific BGP group.
`Group Index`	Index number for the BGP group.
`NLRI`	Network layer reachability information (NLRI) indicating the source of the traffic statistics for the BGP group.
`FEC`	Forwarding equivalence classes (FECs) associated with the BGP group.
`Packets`	Number of packets sent through each FEC.
`Bytes`	Number of bytes transmitted through each FEC.
`EgressAS`	Autonomous system (AS) number of the egress router.
`AdvLabel`	Label associated with each FEC.

Sample Output

show bgp group traffic-statistics (Per-Group-Label Not Configured)

```
user@host> show bgp group traffic-statistics
Group Name: ext1       Group Index: 0          NLRI: inet-labeled-unicast
FEC                   Packets              Bytes     EgressAS    AdvLabel
10.255.245.55               0                  0            I      100224
10.255.245.57               0                  0            I      100240
100.101.0.0               550              48400           25      100256
100.102.0.0               550              48400           25      100256
100.103.0.0               550              48400           25      100272
100.104.0.0               550              48400           25      100272
192.168.25.0                0                  0            I      100288

Group Name: ext2       Group Index: 1          NLRI: inet-labeled-unicast
FEC                   Packets              Bytes     EgressAS    AdvLabel
10.255.245.55               0                  0            I      100224
10.255.245.57               0                  0            I      100240
100.101.0.0               550              48400           25      100256
100.102.0.0               550              48400           25      100256
100.103.0.0               550              48400           25      100272
100.104.0.0               550              48400           25      100272
192.168.25.0                0                  0            I      100288
```

show bgp group traffic-statistics (Per-Group-Label Configured)

```
user@host> show bgp group traffic-statistics
Group Name: ext1       Group Index: 0          NLRI: inet-labeled-unicast
FEC                   Packets              Bytes     EgressAS    AdvLabel
10.255.245.55               0                  0            I      100384
10.255.245.57               0                  0            I      100400
100.101.0.0               101               8888           25      100416
100.102.0.0               101               8888           25      100416
100.103.0.0                 0                  0           25      100432
100.104.0.0                 0                  0           25      100432
192.168.25.0                0                  0            I      100448

Group Name: ext2       Group Index: 1          NLRI: inet-labeled-unicast
FEC                   Packets              Bytes     EgressAS    AdvLabel
```

```
10.255.245.55              0                 0           I      100304
10.255.245.57              0                 0           I      100320
100.101.0.0                0                 0          25      100336
100.102.0.0                0                 0          25      100336
100.103.0.0              101              8888          25      100352
100.104.0.0              101              8888          25      100352
192.168.25.0               0                 0           I      100368
```

show bgp group traffic-statistics labeled-path (Labeled Unicast)

```
user@host> show bgp group traffic-statistics labeled-path
Labels                NextHop            Packets              Bytes
3(top)               10.1.1.1                  0                  0
299840(top)          40.1.1.1                  0                  0
110001(top)          40.1.1.1                  2                168
110002
110003
110001(top)          40.1.1.1                  0                  0
110072
110073
110071(top)          40.1.1.1                  0                  0
110072
110073
120001(top)          40.1.1.1                  0                  0
120002
120003
1000002(top)         40.1.1.1                  2                168
1000003
1000004
```

Release Information

Command introduced before Junos OS Release 7.4.

`labeled-path` option introduced in Junos OS Release 18.1R1 for the MX Series.

show bgp neighbor

IN THIS SECTION

- Syntax | 17175
- Syntax (EX Series Switch, QFX Series, OCX Series, and cRPD) | 17175
- Syntax (SRX Series) | 17176
- Description | 17176
- Options | 17176
- Additional Information | 17177
- Required Privilege Level | 17177
- Output Fields | 17177
- Sample Output | 17194
- Release Information | 17215

Syntax

```
show bgp neighbor
<exact-instance instance-name>
<instance instance-name>
<logical-system (all | logical-system-name)>
<neighbor-address>
<output-queue>
<orf (detail | neighbor-address)
<rib-sharding (main | rib-shard-name)>
```

Syntax (EX Series Switch, QFX Series, OCX Series, and cRPD)

```
show bgp neighbor
<instance instance-name>
<exact-instance instance-name>
<neighbor-address>
<orf (neighbor-address | detail)>
```

```
<rib-sharding neighbor-address>
<validation statistics peer>
```

Syntax (SRX Series)

```
show bgp neighbor
<neighbor-address>
<instance instance-name>
```

Description

Display information about BGP peers.

Options

none	Display information about all BGP peers.
exact-instance *instance-name*	(Optional) Display information for the specified instance only.
instance *instance-name*	(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, `cust1`, `cust11`, and `cust111` are all displayed when you run the `show bgp neighbor instance cust1` command).
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
neighbor-address	(Optional) Display information for only the BGP peer at the specified IP address.
orf (detail | *neighbor-address*)	(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the `detail` option to display detailed output.
output-queue	(Optional) Display information regarding the number of routes currently queued in the 17 prioritized BGP output queues.
rib-sharding (main | junos-bgpshard*shard-number*)	(Optional) Display information for specific shard only. When NSR is configured, display information in the backup Routing Engine. For

example, junos-bgpshard0. If omitted, displays aggregated data from all shards including main shard.

validation statistics *peer* (Optional) Display BGP peer-RIB validation statistics.

Additional Information

For information about the `local-address`, `nlri`, `hold-time`, and `preference` statements, see the Junos OS Routing Protocols Library for Routing Devices.

Required Privilege Level

view

Output Fields

Table 178 on page 17177 describes the output fields for the `show bgp neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 178: show bgp neighbor Output Fields

Field Name	Field Description
`Peer`	Address of the BGP neighbor. The address is followed by the neighbor port number.
`AS`	AS number of the peer.
`Local`	Address of the local routing device. The address is followed by the peer port number.
`Type`	Type of peer: `Internal` or `External`.
`dynamic-neighbor-template-name`	Name of the neighbor for which a peer-limit has been set.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`State`	Current state of the BGP session: • `Active`—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • `Connect`—BGP is waiting for the transport protocol connection to be completed. • `Established`—The BGP session has been established, and the peers are exchanging update messages. • `Idle`—This is the first stage of a connection. BGP is waiting for a Start event. • `OpenConfirm`—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • `OpenSent`—BGP has sent an open message and is waiting to receive an open message from the peer. • `route reflector client`—The BGP session is established with a route reflector client.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Flags`	Internal BGP flags: • `Aggregate Label`—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • `CleanUp`—The peer session is being shut down. • `Delete`—This peer has been deleted. • `Idled`—This peer has been permanently idled. • `ImportEval`—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • `Initializing`—The peer session is initializing. • `PrefixLimitIdle`—The session moves into an idle state when the number of prefixes received from the neighbor exceeds its prefix-limit value. • `PurgePending`—This flag marks one or more routing table (also known as routing information base [RIB]) entries for deletion. The purge job to delete these entries begins after the peer is closed. A purge job keeps running if new routing table entries are marked for deletion. • `PurgeInProgress`—The purge job has started and is not yet complete. • `PurgeImpatient`—The purge begins as a low priority background job. The Adj-RIB-Out can be cleaned up and a new peering can be established in the background before all routes are deleted. After the peer goes down and the group has closed, the purge becomes a normal priority job. • `SendRtn`—Messages are being sent to the peer. • `Sync`—This peer is synchronized with the rest of the peer group. • `RSync`—This peer in the backup Routing Engine is synchronized with the BGP peer in the primary Routing Engine for nonstop active routing. • `TryConnect`—Another attempt is being made to connect to the peer. • `Unconfigured`—This peer is not configured. • `WriteFailed`—An attempt to write to this peer failed.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Last state`	Previous state of the BGP session: • `Active`—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • `Connect`—BGP is waiting for the transport protocol connection to be completed. • `Established`—The BGP session has been established, and the peers are exchanging update messages. • `Idle`—This is the first stage of a connection. BGP is waiting for a Start event. • `OpenConfirm`—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • `OpenSent`—BGP has sent an open message and is waiting to receive an open message from the peer.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Last event`	Last activity that occurred in the BGP session: • `Closed`—The BGP session closed. • `ConnectRetry`—The transport protocol connection failed, and BGP is trying again to connect. • `HoldTime`—The session ended because the hold timer expired. • `KeepAlive`—The local routing device sent a BGP keepalive message to the peer. • `Open`—The local routing device sent a BGP open message to the peer. • `OpenFail`—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • `RecvKeepAlive`—The local routing device received a BGP keepalive message from the peer. • `RecvNotify`—The local routing device received a BGP notification message from the peer. • `RecvOpen`—The local routing device received a BGP open message from the peer. • `RecvUpdate`—The local routing device received a BGP update message from the peer. • `Start`—The peering session started. • `Stop`—The peering session stopped. • `TransportError`—A TCP error occurred.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Last error`	Last error that occurred in the BGP session: • `Cease`—An error occurred, such as a version mismatch, that caused the session to close. • `Finite State Machine Error`—In setting up the session, BGP received a message that it did not understand. • `Hold Time Expired`—The session's hold time expired. • `Message Header Error`—The header of a BGP message was malformed. • `Open Message Error`—A BGP open message contained an error. • `None`—No errors occurred in the BGP session. • `Update Message Error`—A BGP update message contained an error.
`Export`	Name of the export policy that is configured on the peer.
`Import`	Name of the import policy that is configured on the peer.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Options`	Configured BGP options: • `AddressFamily`—Configured address family: inet or inet-vpn. • `AdvertiseBGPStatic`—Configured BGP static routes are advertised. • `AutheKeyChain`—Authentication key change is enabled. • `BfdEnabled`—Status of BFD. • `DontGRHelpFateSharingBfdDown`—Status of the `dont-help-shared-fate-bfd-down` option. If this option is configured the device does not go into graceful restart helper mode. • `DropPathAttributes`—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • `GracefulRestart`—Graceful restart is configured. • `HoldTime`—Hold time configured with the `hold-time` statement. The hold time is three times the interval at which keepalive messages are sent. • `IgnorePathAttributes`—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • `Local Address`—Address configured with the `local-address` statement. • `LLGR`—BGP long-lived graceful restart capability is configured. • `LLGRHelperDisabled`—BGP long-lived graceful restart is completely disabled for a neighbor. • `Multihop`—Allow BGP connections to external peers that are not on a directly connected network. • `NLRI`—Configured MBGP state for the BGP group: `multicast`, `unicast`, or both if you have configured `nlri any`. • `Peer AS`—Configured peer autonomous system (AS). • `Preference`—Preference value configured with the `preference` statement. • `Refresh`—Configured to refresh automatically when the policy changes. • `Rib-group`—Configured routing table group.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
	• RFC6514CompliantSafi129—Configured SAFI 129 according to RFC 6514 (BGP VPN multicast used to use SAFI 128).
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Peer does not support LLGR Restarter or Receiver functionality	BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode completely.
Peer does not support LLGR Restarter functionality	BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode for any family.
Authentication key change	(Appears only if the authentication-keychain statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(Appears only if the authentication-algorithm statement has been configured) Type of authentication algorithm enabled: hmac or md5.
Prefixlimit configured for NLRI	(Appears only if the drop-excess <percentage> or hide-excess <percentage> option in the prefix-limit statement is configured) NLRI for which the prefix-limit statement is configured.
Acceptedprefixlimit configured for NLRI	(Appears only if the drop-excess <percentage> or hide-excess <percentage> option in the accepted-prefix-limit statement is configured) NLRI for which the accepted-prefix-limit statement is configured.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
Action	(Appears only if the drop-excess <percentage> or hide-excess <percentage> option in the prefix-limit and accepted-prefix-limit statements are configured) Performs the configured action: • teardown <percentage>— Tears down the session when the maximum number of prefixes is reached. • drop-excess <percentage>— Drops excess routes when the maximum number of prefixes is exceeded. If you specify a percentage, the routes are logged when the number of prefixes exceeds that percentage value of the maximum number. • hide-excess <percentage>— Hides excess routes when the maximum number of prefixes is exceeded. If you specify a percentage, the routes are logged when the number of prefixes exceeds that percentage value of the maximum number. If the percentage is modified, the routes are re-evaluated automatically. If the active routes drop below the specified percentage, those routes are kept as hidden.
Limit	(Appears only if the drop-excess <percentage> or hide-excess <percentage> option in the prefix-limit and accepted-prefix-limit statements are configured) Number of the maximum prefixes if exceeded, the configured action in the Action field takes place.
Warning percentage	(Appears only if the drop-excess <percentage> or hide-excess <percentage> option in the prefix-limitand accepted-prefix-limit statements are configured) Percentage of the maximum dropped or hidden routes if exceeded, displays a warning message in the logs.
Count	(Appears only if the drop-excess <percentage> or hide-excess <percentage> option in the prefix-limit and accepted-prefix-limit statements are configured) Displays the number of routes that are dropped or hidden after exceeding the limit configured in the Limit field. The counter resets only after the peer resets.
Address families configured	Names of configured address families for the VPN.
BGP-Static Advertisement Policy	Name of the BGP static policy that is configured on the peer.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Local Address`	Address of the local routing device.
`Remove-private options`	Options associated with the `remove-private` statement.
`Holdtime`	Hold time configured with the `hold-time` statement. The hold time is three times the interval at which keepalive messages are sent.
`Flags for NLRI inet-label-unicast`	Flags related to labeled-unicast: • `TrafficStatistics`—Collection of statistics for labeled-unicast traffic is enabled.
`Traffic statistics`	Information about labeled-unicast traffic statistics: • `Options`—Options configured for collecting statistics about labeled-unicast traffic. • `File`—Name and location of statistics log files. • `size`—Size of all the log files, in bytes. • `files`—Number of log files.
`Traffic Statistics Interval`	Time between sample periods for labeled-unicast traffic statistics, in seconds.
`Preference`	Preference value configured with the `preference` statement.
`Outbound Timer`	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the `out-delay` parameter is configured to a non-zero value.
`Number of flaps`	Number of times the BGP session has gone down and then come back up.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Threads related state`	Displays thread related state if update threading is enabled: • `Thread sync pending`—Thread sync is yet to begin. • `Update thread sync`—Syncing peer up with update threads. • `Shard sync`—Syncincg peer up with shards. If the peer is in shard sync state, it also displays a hex value indicating which shards are yet to send peer up acknowledgement. • `Thread sync complete`—Peer has been synced in update threads and shards. • `Peer UP acknowledgement received from Update Thread`—Display peer up acknowledgement received from update threads.
`Peer ID`	Router identifier of the peer.
`Group index`	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.
`Peer index`	Index that is unique within the BGP group to which the peer belongs.
`Local ID`	Router identifier of the local routing device.
`Local Interface`	Name of the interface on the local routing device.
`Active holdtime`	Hold time that the local routing device negotiated with the peer.
`Keepalive Interval`	Keepalive interval, in seconds.
`I/O Session Thread`	Displays the BGP I/O session thread and its state if update threading is enabled.
`I/O Session Thread`	Displays the BGP I/O session thread and its state if update threading is enabled.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGP peering is established.
NLRI and times for LLGR configured on peer	Names of address families and stale time for BGP long-lived graceful restart configured on the BGP peer or neighbor. Times are displayed using the routing protocol daemon (rpd) %#0T format: <weeks>w<days>d <hours>:<minutes>:<seconds> Zero leading elements are omitted, for example, a value less than one week do not include the weeks.
NLRI and times that peer supports LLGR Restarter for	Names of address families and stale time that the BGP peer supports for restarter mode for BGP long-lived graceful restart. Times are displayed using the routing protocol daemon (rpd) %#0T format: <weeks>w<days>d <hours>:<minutes>:<seconds> Zero leading elements are omitted, for example, a value less than one week do not include the weeks.
NLRI that peer saved LLGR forwarding for	Name of the address family for which the BGP peer saved BGP long-lived graceful restart forwarding.
Graceful Restart Details	Amount of time that is remaining until LLGR expires and the time remaining on the GR stale timer, along with RIB details, are displayed while LLGR receiver mode is active (a peer that negotiated LLGR has disconnected and not yet reconnected).
NLRI we are holding stale routes for	Names of address families (NLRIs) for which that stale routes are held or preserved when BGP graceful restart receiver mode is active for a neighbor.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Time until end-of-rib is assumed for stale routes`	Amount of time remaining on the stale timer until which end-of-RIB (EoR) markers are assumed when BGP graceful restart receiver mode is active for a neighbor. Time is displayed in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). Note that the stale timer display ('Time until end-of-rib is assumed') is also present when a session is active, but the neighbor as not yet sent all of the end-of-rib indications.
`Time until stale routes are deleted or become long-lived stale`	Amount of time up to which stale routes are deleted or become long-lived stale routes when BGP graceful restart receiver mode is active for a neighbor.
`NLRI for restart configured on peer`	Names of address families configured for restart.
`NLRI advertised by peer`	Address families supported by the peer: `unicast` or `multicast`.
`NLRI for this session`	Address families being used for this session.
`Peer supports Refresh capability`	Remote peer's ability to send and request full routing table readvertisement (route refresh capability). For more information, see RFC 2918, *Route Refresh Capability for BGP-4*.
`Restart time configured on peer`	Configured time allowed for restart on the neighbor.
`Stale routes from peer are kept for`	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
`Peer does not support Restarter functionality`	Graceful restart restarter-mode is disabled on the peer.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Peer does not support Receiver functionality`	Graceful restart helper-mode is disabled on the peer.
`Restart time requested by this peer`	Restart time requested by this neighbor during capability negotiation.
`Restart flag received from the peer`	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the `end-of-rib` marker from the speaker before advertising routing information to the speaker.
`NLRI that peer supports restart for`	Neighbor supports graceful restart for this address family.
`NLRI peer can save forwarding state`	Neighbor supporting this address family saves all forwarding states.
`NLRI that peer saved forwarding for`	Neighbor saves all forwarding states for this address family.
`NLRI that restart is negotiated for`	Router supports graceful restart for this address family.
`NLRI of received end-of-rib markers`	Address families for which end-of-routing-table markers are received from the neighbor.
`NLRI of all end-of-rib markers sent`	Address families for which end-of-routing-table markers are sent to the neighbor.
`Peer supports 4 byte AS extension (peer-as 1)`	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`NLRIs for which peer can receive multiple paths`	Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route. Possible value is inet-unicast.
`NLRIs for which peer can send multiple paths: inet-unicast`	Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route. Possible value is inet-unicast.
`Table inet.`*`number`*	Information about the routing table: • `RIB State`—BGP is in the graceful restart process for this routing table: `restart is complete` or `restart in progress`. • `Bit`—Number that represents the entry in the routing table for this peer. • `Send state`—State of the BGP group: `in sync`, `not in sync`, or `not advertising`. • `Active prefixes`—Number of prefixes received from the peer that are active in the routing table. • `Received prefixes`—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • `Accepted prefixes`—Total number of prefixes from the peer that have been accepted by a routing policy. • `Suppressed due to damping`—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
`Last traffic (seconds)`	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
`Input messages`	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Output messages`	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
`Input dropped path attributes`	Information about dropped path attributes: • Code—Path attribute code. • Count—Path attribute count.
`Input ignored path attributes`	Information about ignored path attributes: • Code—Path attribute code. • Count—Path attribute count.
`Output queue`	Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue `0` is for unicast NLRIs, and queue `1` is for multicast NLRIs. It also specifies the routing table name and the NLRI that the table was advertised through, in the format (*routing table name*, *NLRI*). If update threading is enabled, the Output Queue field will display the Output Queue count from update threads with an additional field that displays the Output Queue count per RIB as fetched from main or shards. **NOTE**: The output queue of routing tables that are not advertised, will only show up at `extensive` output level.
`Trace options`	Configured tracing of BGP protocol packets and operations.
`Trace file`	Name of the file to receive the output of the tracing operation.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Filter Updates recv`	(`orf` option only) Number of outbound-route filters received for each configured address family. **NOTE**: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
`Immediate`	(`orf` option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes. **NOTE**: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
`Filter`	(`orf` option only) Type of prefix filter received: `prefix-based` or `extended-community`.
`Received filter entries`	(`orf` option only) List of received filters displayed.
`seq`	(`orf` option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
`prefix`	(`orf` option only) Address for the prefix entry that matches the filter.
`minlength`	(`orf` option only) Minimum prefix length, in bits, required to match this prefix.
`maxlength`	(`orf` option only) Maximum prefix length, in bits, required to match this prefix.
`match`	(`orf` option only) For this prefix match, whether to `permit` or `deny` route updates.
`Valid Routes`	(Appears only if the `validation statistics` *peer* option is selected. Count of valid routes received per routing information base (RIB) per BGP neighbor.
`Routes invalid due to max-length`	(Appears only if the `validation statistics` *peer* option is selected. Count of invalid routes due to maximum-length mismatch received per BGP peer-RIB.

Table 178: show bgp neighbor Output Fields *(Continued)*

Field Name	Field Description
`Routes invalid due to origin-as`	(Appears only if the `validation statistics` *peer* option is selected. Count of invalid routes due to `origin-as` mismatch received per BGP peer-RIB.
`Unknown routes`	(Appears only if the `validation statistics` *peer* option is selected. Count of unknown routes received per RIB per BGP neighbor.
`Validation State`	In the case of invalid routes, an invalid reason is displayed: [`max-length` | `origin-as`] mismatch.

Sample Output

show bgp neighbor

```
user@host > show bgp neighbor
```

For M Series, MX Series, and T Series routers running Junos OS Release 16.1 or later, the `show bgp neighbor` output includes the BGP group the peer belongs to, the routing instance (if any) that the peer is configured in, and the routing instance that the peer is using for the forwarding context (if applicable). An example follows.

```
Peer: 10.255.7.250+179 AS 65010   Local: 10.255.7.248+63740 AS 65010
  Group: toAsbr2               Routing-Instance: master
  Forwarding routing-instance: toAsbr2
    Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Options: <AdvertiseBGPStatic>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250    Local ID: 10.255.7.248      Active Holdtime: 90
  Keepalive Interval: 30         Group index: 0    Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
```

```
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 65010)
  Peer does not support Addpath
  NLRI that we support extended nexthop encoding for: inet-unicast
  NLRI that peer supports extended nexthop encoding for: inet-unicast

  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              1
    Received prefixes:            1
    Accepted prefixes:            1
    Suppressed due to damping:    0
    Advertised prefixes:          1
  Last traffic (seconds): Received 9    Sent 5    Checked 5
  Input messages:  Total 36     Updates 2       Refreshes 0     Octets 718
  Output messages: Total 37     Updates 1       Refreshes 0     Octets 796
  Output Queue[0]: 0        (inet.0, inet-unicast)

Peer: 10.255.162.214+52193 AS 65100 Local: 10.255.167.205+179 AS 65100
   Type: Internal    State: Established  (route reflector client)Flags: <Sync>
   Last State: OpenConfirm   Last Event: RecvKeepAlive
   Last Error: None
   Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
   Address families configured: inet-unicast inet-vpn-unicast route-target
   Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
   Number of flaps: 0
   Peer ID: 10.255.162.214  Local ID: 10.255.167.205    Active Holdtime: 90
   Keepalive Interval: 30         Group index: 0    Peer index: 1
```

show bgp neighbor (dont-help-shared-fate-bfd-down is configured)

```
user@host> show bgp neighbor
```

```
Peer: 10.1.1.1 AS 200          Local: unspecified AS 65017
  Group: one                   Routing-Instance: master
  Forwarding routing-instance: master
  Type: External    State: Idle           Flags: <PeerInterfaceError>
  Last State: NoState       Last Event: NoEvent
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Options: <BfdEnabled>
  Options: <DontGRHelpFateSharingBfdDown>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Trace options: bridge
  Trace file: /var/log/bgp-log size 131072 files 10

Peer: 10.20.1.1 AS 200          Local: unspecified AS 65017
  Group: one                   Routing-Instance: master
  Forwarding routing-instance: master
  Type: External    State: Idle           Flags: <PeerInterfaceError>
  Last State: NoState       Last Event: NoEvent
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Options: <BfdEnabled>
  Options: <DontGRHelpFateSharingBfdDown>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.30.1.1 AS 200          Local: unspecified AS 65017
  Group: two                   Routing-Instance: master
  Forwarding routing-instance: master
  Type: External    State: Idle           Flags: <PeerInterfaceError>
  Last State: NoState       Last Event: NoEvent
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Options: <BfdEnabled>
  Options: <DontGRHelpFateSharingBfdDown>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
```

show bgp neighbor (CLNS)

```
user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 65200  Local: 10.245.245.3+3770 AS 65100
```

```
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS     Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1     Local ID: 10.245.245.3     Active Holdtime: 90
  Keepalive Interval: 30         Peer index: 0
  NLRI advertised by peer: iso-vpn-unicast
  NLRI for this session: iso-vpn-unicast
  Peer supports Refresh capability (2)
  Table bgp.isovpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:              3
    Received prefixes:            3
    Suppressed due to damping:    0
    Advertised prefixes:          3
  Table aaaa.iso.0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: not advertising
    Active prefixes:              3
    Received prefixes:            3
    Suppressed due to damping:    0
  Last traffic (seconds): Received 6    Sent 5    Checked 5
  Input messages:  Total 1736   Updates 4       Refreshes 0     Octets 33385
  Output messages: Total 1738   Updates 3       Refreshes 0     Octets 33305
  Output Queue[0]: 0        (bgp.isovpn.0, iso-vpn-unicast)
  Output Queue[1]: 0        (aaaa.iso.0, iso-vpn-unicast)
```

show bgp neighbor (Layer 2 VPN)

```
user@host> show bgp neighbor
Peer: 10.69.103.2     AS 65536 Local: 10.69.103.1     AS 65539
  Type: External    State: Active         Flags: <ImportEval>
  Last State: Idle          Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
```

```
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily   PeerAS Refresh>
  Address families configured: inet-unicast
  Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.69.104.2     AS 65539 Local: 10.69.104.1     AS 65539
  Type: External    State: Active         Flags: <ImportEval>
  Last State: Idle          Last Event: Start
  Last Error: None
  Export: [ BGP-L-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS Refresh>
  Address families configured: inet-labeled-unicast
  Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.255.14.182+179 AS 65069    Local: 10.255.14.176+2131 AS 65069
  Type: Internal    State: Established    Flags: <ImportEval>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily   Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.14.182    Local ID: 10.255.14.176    Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast l2vpn
  NLRI advertised by peer: inet-vpn-unicast l2vpn
  NLRI for this session: inet-vpn-unicast l2vpn
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast l2vpn
  NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
  NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
  NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
  NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
  Table bgp.l3vpn.0 Bit: 10000
    RIB State: BGP restart in progress
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            10
    Received prefixes:          10
    Suppressed due to damping:  0
  Table bgp.l2vpn.0 Bit: 20000
```

```
    RIB State: BGP restart in progress
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            1
    Received prefixes:          1
    Suppressed due to damping:  0
  Table BGP-INET.inet.0 Bit: 30000
    RIB State: BGP restart in progress
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            2
    Received prefixes:          2
    Suppressed due to damping:  0
  Table BGP-L.inet.0 Bit: 40000
    RIB State: BGP restart in progress
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            2
    Received prefixes:          2
    Suppressed due to damping:  0
  Table LDP.inet.0 Bit: 50000
    RIB State: BGP restart is complete
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            1
    Received prefixes:          1
    Suppressed due to damping:  0
  Table OSPF.inet.0 Bit: 60000
    RIB State: BGP restart is complete
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            2
    Received prefixes:          2
    Suppressed due to damping:  0
  Table RIP.inet.0 Bit: 70000
    RIB State: BGP restart is complete
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            2
    Received prefixes:          2
    Suppressed due to damping:  0
  Table STATIC.inet.0 Bit: 80000
    RIB State: BGP restart is complete
```

```
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            1
    Received prefixes:          1
    Suppressed due to damping:  0
  Table L2VPN.l2vpn.0 Bit: 90000
    RIB State: BGP restart is complete
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            1
    Received prefixes:          1
    Suppressed due to damping:  0
  Last traffic (seconds): Received 0    Sent 0    Checked 0
  Input messages:  Total 14     Updates 13      Refreshes 0     Octets 1053
  Output messages: Total 3      Updates 0       Refreshes 0     Octets 105
  Output Queue[0]: 0        (bgp.l3vpn.0, inet-vpn-unicast)
  Output Queue[1]: 0        (bgp.l2vpn.0, inet-vpn-unicast)
  Output Queue[2]: 0        (BGP-INET.inet.0, inet-vpn-unicast)
  Output Queue[3]: 0        (BGP-L.inet.0, inet-vpn-unicast)
  Output Queue[4]: 0        (LDP.inet.0, inet-vpn-unicast)
  Output Queue[5]: 0        (OSPF.inet.0, inet-vpn-unicast)
  Output Queue[6]: 0        (RIP.inet.0, inet-vpn-unicast)
  Output Queue[7]: 0        (STATIC.inet.0, inet-vpn-unicast)
  Output Queue[8]: 0        (L2VPN.l2vpn.0, inet-vpn-unicast)
```

show bgp neighbor (Layer 3 VPN) (Not supported on the OCX Series.)

```
user@host> show bgp neighbor
Peer: 192.0.2.0.179     AS 65045 Local: 192.0.2.1+1214    AS 65045
  Type: Internal    State: Established    Flags: <ImportEval>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ match-all ] Import: [ match-all ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily    Rib-group Refresh>
  Address families configured: inet-vpn-unicast
  Local Address: 192.0.2.1 Holdtime: 90 Preference: 170
  Flags for NLRI inet-labeled-unicast: TrafficStatistics
  Traffic Statistics: Options: all File: /var/log/bstat.log
                                         size 131072 files 10
  Traffic Statistics Interval: 60
  Number of flaps: 0
```

```
  Peer ID: 192.168.1.110    Local ID: 192.168.1.111    Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast
  NLRI advertised by peer: inet-vpn-unicast
  NLRI for this session: inet-vpn-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast
  NLRI peer can save forwarding state: inet-vpn-unicast
  NLRI that peer saved forwarding for: inet-vpn-unicast
  NLRI that restart is negotiated for: inet-vpn-unicast
  NLRI of received end-of-rib markers: inet-vpn-unicast
  NLRI of all end-of-rib markers sent: inet-vpn-unicast
  Table bgp.l3vpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:            2
    Received prefixes:          2
    Suppressed due to damping:  0
  Table vpn-green.inet.0 Bit: 20001
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:            2
    Received prefixes:          2
    Suppressed due to damping:  0
  Last traffic (seconds): Received 15   Sent 20   Checked 20
  Input messages:  Total 40     Updates 2       Refreshes 0     Octets 856
  Output messages: Total 44     Updates 2       Refreshes 0     Octets 1066
  Output Queue[0]: 0        (bgp.l3vpn.0, inet-vpn-unicast)
  Output Queue[1]: 0        (vpn-green.inet.0, inet-vpn-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpgr.log size 131072 files 10
```

show bgp neighbor neighbor-address

```
user@host> show bgp neighbor 10.255.245.12
Peer: 10.255.245.12+179 AS 65035  Local: 10.255.245.13+2884 AS 65035
```

```
  Type: Internal    State: Established  (route reflector client)Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group   Refresh>
  Options: RFC6514CompliantSafi129
  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12    Local ID: 10.255.245.13    Active Holdtime: 90
  Keepalive Interval: 30
BFD: disabled
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 4
    Received prefixes: 6
    Suppressed due to damping: 0
  Table inet6.0 Bit: 20000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 3    Sent 3    Checked 3
  Input messages:  Total 9      Updates 6       Refreshes 0     Octets 403
  Output messages: Total 7      Updates 3       Refreshes 0     Octets 365
  Output Queue[0]: 0        (inet.0, inet-unicast)
  Output Queue[1]: 0        (inet6.0, inet6-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpgr size 131072 files 10
```

show bgp neighbor neighbor-address

```
user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: Cease
  Export: [ export-policy ] Import: [ import-policy ]
  Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
  Address families configured: inet-unicast inet-multicast
  Holdtime: 60000 Preference: 170
  Number of flaps: 4
  Last flap event: RecvUpdate
  Error: 'Cease' Sent: 5 Recv: 0
  Peer ID: 10.255.245.6     Local ID: 10.255.245.5     Active Holdtime: 60000
  Keepalive Interval: 20000      Peer index: 0
  BFD: disabled, down
  Local Interface: fxp0.0
  NLRI advertised by peer: inet-unicast inet-multicast
  NLRI for this session: inet-unicast inet-multicast
  Peer supports Refresh capability (2)
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              8
    Received prefixes:            10
    Accepted prefixes:            10
    Suppressed due to damping:    0
    Advertised prefixes:          3
  Table inet.2 Bit: 20000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              0
    Received prefixes:            0
    Accepted prefixes:            0
    Suppressed due to damping:    0
    Advertised prefixes:          0
  Last traffic (seconds): Received 357  Sent 357  Checked 357
  Input messages:  Total 4    Updates 2     Refreshes 0     Octets 211
  Output messages: Total 4    Updates 1     Refreshes 0     Octets 147
  Output Queue[0]: 0        (inet.0, inet-unicast)
  Output Queue[1]: 0        (inet.2, inet-multiicast)
```

```
  Trace options:  all
  Trace file: /var/log/bgp size 10485760 files 10
```

show bgp neighbor neighbor-address (BGP Graceful Restart Enabled)

```
user@router> show bgp neighbor 10.255.255.16

  Peer: 10.255.255.16 AS 65100     Local: 10.255.255.12 AS 65100
    Type: Internal    State: Active         Flags: <>
    Last State: Idle          Last Event: Start
    Last Error: None
    Options: <Preference LocalAddress AddressFamily Rib-group Refresh>
    Options: <LLGR>
    Address families configured: l2vpn
    Local Address: 10.255.255.12 Holdtime: 90 Preference: 170
    NLRI l2vpn:
    Number of flaps: 6
    Last flap event: Restart
    NLRI we are holding stale routes for: inet-vpn-unicast
    Time until stale routes are deleted or become long-lived stale: 00:01:57
    Time until end-of-rib is assumed for stale routes: 00:04:43
    Table bgp.l3vpn.0
      RIB State: BGP restart is complete
      RIB State: VPN restart is complete
      Send state: not advertising
      Active prefixes:              0
      Received prefixes:            7
      Accepted prefixes:            7
      Suppressed due to damping:    0
    Table foo.inet.0 Bit: 30000
      RIB State: BGP restart is complete
      RIB State: VPN restart is complete
      Send state: not in sync
      Active prefixes:              0
      Received prefixes:            7
      Accepted prefixes:            7
      Suppressed due to damping:    0
```

show bgp neighbor neighbor-address (BGP Long-Lived Graceful Restart)

```
user@router> show bgp neighbor 10.4.12.11

  Peer: 10.4.12.11 AS 65100        Local: 10.6.128.225 AS 65100
    Type: Internal    State: Active         Flags: <>
    Last State: Idle          Last Event: Start
    Last Error: None
    Export: [ foo ]
    Options: <Preference LocalAddress Refresh GracefulRestart>
    Options: <LLGR>
    Local Address: 10.6.128.225 Holdtime: 90 Preference: 170
    Number of flaps: 3
    Last flap event: Restart
    Error: 'Cease' Sent: 0 Recv: 1
    Time until long-lived stale routes deleted: inet-vpn-unicast 10:00:22 route-target 10:00:22
    Table bgp.l3vpn.0
      RIB State: BGP restart is complete
      RIB State: VPN restart is complete
      Send state: not advertising
      Active prefixes:              0
      Received prefixes:            7
      Accepted prefixes:            7
      Suppressed due to damping:    0
    Table foo.inet.0 Bit: 30000
      RIB State: BGP restart is complete
      RIB State: VPN restart is complete
      Send state: not in sync
      Active prefixes:              0
      Received prefixes:            7
      Accepted prefixes:            7
      Suppressed due to damping:    0
```

show bgp neighbor orf neighbor-address detail

```
user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
  Group: ext1

  inet-unicast
```

```
    Filter updates recv:          1 Immediate:          1
    Filter: prefix-based  receive
      Received filter entries:
        seq 1: prefix 10.2.2.2/32: minlen 32: maxlen 32: match deny:

  inet6-unicast
    Filter updates recv:          0 Immediate:          1
    Filter: prefix-based  receive
      Received filter entries:
        *:*
```

show bgp neighbor logical-system

```
user@host > show bgp neighbor logical-system ITR1
Peer: 10.79.8.2+179 AS 65536   Local: 10.79.8.1+50891 AS 65500
  Description: MX1
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
 ....
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              1
    Received prefixes:            1
    Accepted prefixes:            1
    Suppressed due to damping:    0
    Advertised prefixes:          10
    Stale prefixes:               4: <=new, line only appears if count is non-0
It is the Number of prefixes marked as stale;
    LLGR-stale prefixes:          5: <=new, line only appears if count is non-0
It is the Number of prefixes marked as LLGR-stale
```

show bgp neighbor output-queue

```
user@host > show bgp neighbor output-queue
Peer: 192.0.2.2+179 AS 65103       Local: 192.0.2.1+50799 AS 65102
  Output Queue[0]: 0            (inet.0, inet-unicast)
    Priority 1 :   0
    Priority 2 :   0
```

```
    Priority 3 :   0
    Priority 4 :   0
    Priority 5 :   0
    Priority 6 :   0
    Priority 7 :   0
    Priority 8 :   0
    Priority 9 :   0
    Priority 10:   0
    Priority 11:   0
    Priority 12:   0
    Priority 13:   0
    Priority 14:   0
    Priority 15:   0
    Priority 16:   0
    Expedited  :   0
```

show bgp neighbor validation-statistics *peer*

```
user@host > show bgp neighbor validation-statistics 10.1.1.1
        Valid Routes: 432
        Routes invalid due to max-length: 646
        Routes invalid due to origin-as: 21
        Unknown routes: 5353
```

show bgp neighbor (Segment Routing Traffic Engineering)

```
user@host > show bgp neighbor
run show bgp neighbor 10.1.1.254
   Peer: 10.1.1.254+60180 AS 65100   Local: 10.1.1.1+179 AS 65100
  Group: toB                   Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress>
  Address families configured: inet-segment-routing-te
  Local Address: 10.1.1.1 Holdtime: 90 Preference: 170 Local AS: 65100 Local System AS: 0
  Number of flaps: 0
  Peer ID: 10.128.150.15    Local ID: 10.128.150.110     Active Holdtime: 90
  Keepalive Interval: 30         Group index: 0    Peer index: 0
```

```
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: disabled, down
  NLRI for restart configured on peer: inet-segment-routing-te
  NLRI advertised by peer: inet-segment-routing-te
  NLRI for this session: inet-segment-routing-te
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-segment-routing-te
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65100)
  Peer does not support Addpath
  Last traffic (seconds): Received 17628 Sent 25  Checked 17628
  Input messages:  Total 2      Updates 0       Refreshes 0     Octets 82
  Output messages: Total 1      Updates 0       Refreshes 0     Octets 19
  Trace options:  all
  Trace file: /var/log/bgp.log size 10485760 files 10
```

show bgp neighbor (with rib-sharding configured)

```
user@host > show bgp neighbor rib-sharding main
Peer: 10.1.1.1+179 AS 65001         Local: 10.2.2.1+60231 AS 65001
  Group: toFeeder              Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Refresh>
  Options: <ConnectRetryInterval>
  Options: <GracefulShutdownRcv>
  Local Address: 10.2.2.1 Holdtime: 90 Preference: 170
  Graceful Shutdown Receiver local-preference: 0
  Number of flaps: 0
  Threads related state:
    Internal State: Thread sync complete
    Peer UP acknowledgement received from Update Thread
  Peer ID: 10.1.1.1         Local ID: 10.2.2.1           Active Holdtime: 90
  Keepalive Interval: 30         Group index: 0    Peer index: 0    SNMP index: 0
  I/O Session Thread: bgp-updio-2 State: Enabled
  BFD: disabled, down
```

```
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65001)
  Peer does not support Addpath
  NLRI(s) enabled for color nexthop resolution: inet-unicast
  Table inet.0 Bit: 20002
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              0
    Received prefixes:            0
    Accepted prefixes:            0
    Suppressed due to damping:    0
    Advertised prefixes:          0
  Last traffic (seconds): Received 7    Sent 11   Checked 3910
  Input messages:  Total 145    Updates 1       Refreshes 0     Octets 2759
  Output messages: Total 135    Updates 0       Refreshes 0     Octets 2569
  Output Queue[1]: 0            (inet.0, inet-unicast)
  Output Queue[1]: 0            (inet.0, inet-unicast) (Main/Shards)
```

show bgp neighbor (with rib-sharding configured on crpd)

```
user@host > show bgp neighbor rib-sharding junos-bgpshard14


Peer: 10.2.2.1 AS 65100           Local: 10.20.255.10 AS 65100
  Description: To_Adolf
  Group: G101_V4               Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Idle         (route reflector client)Flags: <>
  Last State: Established   Last Event: Stop
  Last Error: None
  Import: [ Block_bgp ]
```

```
  Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group Refresh>
  Options: <GracefulShutdownRcv>
  Address families configured: inet-unicast inet-vpn-unicast inet6-vpn-unicast route-target
  Local Address: 10.20.255.10 Holdtime: 10 Preference: 170
  Graceful Shutdown Receiver local-preference: 0
  Number of flaps: 1
  Last flap event: Stop
Peer: 10.5.1.1 AS 65100          Local: 10.20.255.10 AS 65100
  Description: To_stonepark
  Group: G201_V4              Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Idle         (route reflector client)Flags: <>
  Last State: Established   Last Event: Stop
  Last Error: None
  Import: [ Block_bgp ]
  Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group Refresh>
  Options: <GracefulShutdownRcv>
  Address families configured: inet-vpn-unicast inet6-vpn-unicast route-target
  Local Address: 10.20.255.10 Holdtime: 10 Preference: 170
  Graceful Shutdown Receiver local-preference: 0
  Number of flaps: 2
  Last flap event: Stop
  Trace options:  all
  Trace file: /var/log/aaaaaa size 1073741824 files 10
```

show bgp neighbor (with drop-excess <percentage> option configured and not exceeding the maximum configured percentage.)

```
user@host > show bgp neighbor
Peer: 10.128.139.6+179 AS 65100  Local: 10.128.139.63+55782 AS 65100
  Group: ibgp                 Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established  (route reflector client)Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Import: [ list_based_on_as ]
  Options: <LocalAddress KeepNone Cluster AddressFamily PrefixLimit Rib-group Refresh>
  Options: <GracefulShutdownRcv>
  Address families configured: inet-unicast
  Local Address: 10.128.139.63 Holdtime: 90 Preference: 170
```

```
Graceful Shutdown Receiver local-preference: 0
Prefixlimit configured for NLRI: inet-unicast Limit: 800000 Action: Drop excess
Number of flaps: 0
Peer ID: 10.128.139.6    Local ID: 10.128.139.63     Active Holdtime: 90
Keepalive Interval: 30         Group index: 0    Peer index: 0    SNMP index: 0
I/O Session Thread: bgpio-0 State: Enabled
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast inet-vpn-unicast inet-vpn-multicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
Restart flag received from the peer: Notification
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer does not support LLGR Restarter functionality
Peer supports 4 byte AS extension (peer-as 65100)
Peer does not support Addpath
NLRI(s) enabled for color nexthop resolution: inet-unicast
Table inet.0 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:              800000
  Received prefixes:            800000
  Accepted prefixes:            800000
  Suppressed due to damping:    0
  Advertised prefixes:          29788
Last traffic (seconds): Received 1    Sent 0    Checked 538
Input messages:  Total 137090 Updates 137070  Refreshes 0 Octets 17371127
Output messages: Total 23022  Updates 23001 Refreshes 0 Octets 3539841
Output Queue[1]: 0            (inet.0, inet-unicast)
Trace options: send refresh
Trace file: /var/log/bgp_refresh size 5242880 files 10
```

NOTE: Such similar output is displayed when you configure the `hide-excess <percentage>` option and does not exceed the maximum configured percentage. This sample output is applicable for both the `prefix-limit` and `accepted-prefix-limit` configuration statements.

show bgp neighbor (with drop-excess <percentage> option configured and exceeding the maximum configured percentage.)

```
user@host > show bgp neighbor
Peer: 10.128.139.6+179 AS 65100  Local: 10.128.139.63+55782 AS 65100
  Group: ibgp                  Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established  (route reflector client)Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Import: [ list_based_on_as ]
  Options: <LocalAddress KeepNone Cluster AddressFamily PrefixLimit Rib-group Refresh>
  Options: <GracefulShutdownRcv>
  Address families configured: inet-unicast
  Local Address: 10.128.139.63 Holdtime: 90 Preference: 170
  Graceful Shutdown Receiver local-preference: 0
  Prefixlimit configured for NLRI: inet-unicast Limit: 800000 Action: Drop excess
  Number of flaps: 0
  Dropped prefixes - Exceeded configured prefix-limits
  Peer ID: 10.128.139.6    Local ID: 10.128.139.63     Active Holdtime: 90
  Keepalive Interval: 30         Group index: 0    Peer index: 0    SNMP index: 0
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast inet-vpn-unicast inet-vpn-multicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65100)
  Peer does not support Addpath
  NLRI(s) enabled for color nexthop resolution: inet-unicast
  Table inet.0 Bit: 20000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              800000
    Received prefixes:            800000
```

```
    Accepted prefixes:            800000
    Suppressed due to damping:    0
    Advertised prefixes:          29788
  Prefix-limit dropped routes NLRI: inet-unicast count: 27850
  Last traffic (seconds): Received 1    Sent 0    Checked 538
  Input messages:  Total 137090    Updates 137070    Refreshes 0    Octets 17371127
  Output messages: Total 23022    Updates 23001    Refreshes 0    Octets 3539841
  Output Queue[1]: 0            (inet.0, inet-unicast)
  Trace options: send refresh
  Trace file: /var/log/bgp_refresh size 5242880 files 10
```

```
user@host > show bgp neighbor
Threading mode: BGP I/O
Default eBGP mode: advertise - accept, receive - accept
Groups: 1 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                 1629784     829784          0          0          0          0
bgp.l3vpn.0
                       0          0          0          0          0          0
bgp.l3vpn.2
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
10.128.139.6            65100     137070      23013       0       0        8:49 Establ
  inet.0: 800000/800000/800000/0
```

NOTE: This sample output is applicable for both the `prefix-limit` and `accepted-prefix-limit` configuration statements.

show bgp neighbor (with hide-excess <percentage> option configured and exceeding the maximum configured percentage.)

```
user@host > show bgp neighbor
Peer: 10.128.139.6+50420 AS 65100 Local: 10.128.139.63+179 AS 65100
  Group: ibgp                  Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established  (route reflector client)Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
```

```
Last Error: Cease
Import: [ list_based_on_as ]
Options: <LocalAddress KeepNone Cluster AddressFamily PrefixLimit Rib-group Refresh>
Options: <GracefulShutdownRcv>
Address families configured: inet-unicast
Local Address: 10.128.139.63 Holdtime: 90 Preference: 170
Graceful Shutdown Receiver local-preference: 0
Prefixlimit configured for NLRI: inet-unicast Limit: 800000 Action: Hide excess
Number of flaps: 1
Last flap event: Stop
Error: 'Cease' Sent: 1 Recv: 0
Hidden prefixes - Exceeded configured prefix-limits
Peer ID: 10.128.139.6    Local ID: 10.128.139.63     Active Holdtime: 90
Keepalive Interval: 30         Group index: 0    Peer index: 0    SNMP index: 0
I/O Session Thread: bgpio-0 State: Enabled
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast inet-vpn-unicast inet-vpn-multicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
Restart flag received from the peer: Notification
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer does not support LLGR Restarter functionality
Peer supports 4 byte AS extension (peer-as 65100)
Peer does not support Addpath
NLRI(s) enabled for color nexthop resolution: inet-unicast
Table inet.0 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:              799908
  Received prefixes:            827581
  Accepted prefixes:            799908
  Suppressed due to damping:    0
  Advertised prefixes:          29826
Prefix-limit hidden routes NLRI: inet-unicast count: 28484
Last traffic (seconds): Received 0    Sent 0    Checked 1028
Input messages:  Total 140272 Updates 140232  Refreshes 0 Octets 17794222
Output messages: Total 46362  Updates 46324 Refreshes 0 Octets 6473287
Output Queue[1]: 0            (inet.0, inet-unicast)
```

```
  Trace options: send refresh
  Trace file: /var/log/bgp_refresh size 5242880 files 10
```

```
user@host > show bgp neighbor
Threading mode: BGP I/O
Default eBGP mode: advertise - accept, receive - accept
Groups: 1 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                 1657319     829736          0          0          0          0
bgp.l3vpn.0
                       0          0          0          0          0          0
bgp.l3vpn.2
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
10.128.139.6            65100     140669      46482       0       1       18:46 Establ
  inet.0: 799916/827583/799916/0
```

NOTE: This sample output is applicable for both the `prefix-limit` and `accepted-prefix-limit` configuration statements.

Release Information

Command introduced before Junos OS Release 7.4.

`orf` option introduced in Junos OS Release 9.2.

`exact-instance` option introduced in Junos OS Release 11.4.

`output-queue` option introduced in Junos OS Release 16.1.

`DontGRHelpFateSharingBfdDown` is added to the `options` field of the command output in Junos OS Release 18.3R1.

`PurgePending`, `PurgeInProgress`, and `PurgeImpatient` are added to the `Flags` field of the command output in Junos OS Release 19.4R1.

`rib-sharding` option introduced in cRPD Release 20.1R1.

`Prefixlimit configured for NLRI`, `Acceptedprefixlimit configured for NLRI`, `Action`, `Limit`, `Warning percentage`, and `Count` fields are introduced to the `show bgp neighbor` output in Junos OS Release 21.2R1.

`Prefixlimit configured for NLRI`, `Acceptedprefixlimit configured for NLRI`, `Action`, `Limit`, `Warning percentage`, and `Count` fields are introduced to the `show bgp neighbor` output in Junos OS Evolved Release 21.3R1.

RELATED DOCUMENTATION

clear bgp neighbor | 3437

show bgp output-scheduler

IN THIS SECTION

- Syntax | 17216
- Description | 17216
- Options | 17217
- Required Privilege Level | 17217
- Output Fields | 17217
- Sample Output | 17217
- Release Information | 17218

Syntax

```
show bgp output-scheduler
<exact-instance instance-name>
<fabric <exact-instance instance-name | instance (instance-name | prefix)>>
<instance (instance-name | prefix)>
<logical-system (all | logical-system-name)>
```

Description

Display output scheduler information including the number of tokens assigned to each priority output queue. Output queues are shown as classes.

Options

none	Display the number of tokens assigned to each of the 17 BGP priority output queues for the master routing instance.
exact-instance *instance-name*	(Optional) Display the number of tokens assigned to each of the 17 BGP priority output queues for the specified routing instance name.
fabric	(Optional) Display the internal fabric state. The exact-instance and instance options can be used along with the fabric option.
instance *instance-name*	(Optional) Display information about BGP peers for all routing instances whose name begins with the *instance-name* string. (for example, `cust1`, `cust11`, and `cust111` are all displayed when you run the `show bgp output-scheduler instance cust1` command).
logical-system *logical-system-name.*	(Optional) Display the number of tokens assigned to each of the 17 BGP priority queues within the specified logical-system. The instance and exact-instance options can be used along with the logical-system option.

Required Privilege Level

routing

Output Fields

Sample Output

command-name (optional-text)

```
user@host> show bgp output-scheduler
Instance: master
  Class         Tokens
  ------------  ------
  Priority  1        1
  Priority  2       10
  Priority  3       15
  Priority  4       20
  Priority  5       25
  Priority  6       30
  Priority  7       35
```

```
Priority  8       40
Priority  9       45
Priority 10       50
Priority 11       55
Priority 12       60
Priority 13       65
Priority 14       70
Priority 15       75
Priority 16       80
Expedited        100

Priority  Class
--------  ------------
low       Priority  1
medium    Priority 10
high      Expedited
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

show bgp group output-queues | 17166

show bgp replication

IN THIS SECTION

- Syntax | 17219
- Description | 17219
- Required Privilege Level | 17219
- Output Fields | 17219
- Sample Output | 17221
- Release Information | 17223

Syntax

```
show bgp replication
rib-sharding (rib-shard-name)
```

Description

Displays the status of BGP state replication between the primary and backup Routing Engines on devices that have nonstop active routing configured on them.

> **CAUTION**: Before attempting nonstop active routing switchover, check the output of `show bgp replication` to confirm that BGP routing table synchronization has completed on the backup Routing Engine. The `complete` status in the output of `show task replication` only indicates that the socket replication has completed and the BGP synchronization is in progress.
>
> To determine whether BGP synchronization is complete, you must check the `Protocol state` and `Synchronization state` fields in the output of `show bgp replication` on the primary Routing Engine. The `Protocol state` must be `idle` and the `Synchronization state` must be `complete`. If you perform NSR switchover before the BGP synchronization has completed, the BGP session might flap.

Options

rib-sharding (*rib-shard-name*) (Optional) Display replicated information for specific shard only.

Required Privilege Level

view

Output Fields

Table 179 on page 17220 lists the output fields for the `show bgp replication` command. Output fields are listed in the approximate order in which they appear.

Table 179: show bgp replication Output Fields

Field Name	Field Description
Precision timer registration	State of BGP precision timer feature in the kernel. • **Registered** BGP registers with the precision-timer feature in the kernel for auto keepalive generation after switchover. • **NotRegistered** Keepalive format of BGP is not registered.
session state	State of the current internal BGP state replication session, Up or Down, and the duration for which the session has been in the indicated state.
flaps	Total number of flaps that occurred.
protocol state	Current state of the protocol operation, Active, Connect, Idle, and the duration for which the protocol has been in the indicated state.
synchronization state	Synchronization state at the time of executing the command. The states can be: • **Idle** • **Neighbor**—Indicates that the neighbor state synchronization is in progress. • **AckWait**—Indicates that the request processing is over. • **ORF**—Indicates that the outbound routing filter synchronization is in progress. • **RIB**—Indicates that the routing table synchronization is in progress. • **Complete**
number of peers waiting	Total number of peers waiting for various messages: • **AckWait**—Number of peers waiting for a connection establishment or completed acknowledgment messages. • **SoWait**—Number of peers waiting for TCP socket-related operations. • **Scheduled**—Number of peers being synchronized.

Table 179: show bgp replication Output Fields *(Continued)*

Field Name	Field Description
messages sent	Number of various types of messages that have been sent since internal replication session became active: • **Open**—Number of Open messages sent. • **Establish**—Number of connection establishment acknowledgment messages sent. • **Update**—Number of update messages sent. • **Error**—Number of error messages sent. • **Complete**—Number of connection complete acknowledgment messages sent.

messages received	Total number of messages received:
	• **Open**—Number of Open messages received.
	• **Request**—Number of request messages received:
	  • **Wildcard**—Number of requests received that used wildcards in the target address.
	  • **Targeted**—Number of requests received that used a specific address.
	• **EstablishAck**—Number of connection establishment acknowledgement messages received.
	• **CompleteAck**—Number of connection completed acknowledgement messages received.

Sample Output

show bgp replication (for Primary)

```
user@host> show bgp replication
Synchronization master:
  Precision timer registration: Registered
  Session state: Up, Since: 10:14
  Flaps: 1, Last flap reason: Backup closed connection
  Protocol state: Idle, Since: 10:14
  Synchronization state: Complete
```

```
  Number of peers waiting: AckWait: 0, SoWait: 0, Scheduled: 0
  Messages sent: Open 1, Establish 11, GrHelper 0, Update 0, GrStaleLabel 0 Error 0, Complete 1
  Messages received: Open 1, Request 1 wildcard 0 targeted, EstablishAck 11, GrHelperAck 0,
CompleteAck 1
```

show bgp replication (for Primary with rib-sharding configured)

```
user@host> show bgp replication rib-sharding junos-bgpshard0
Synchronization master:
  Precision timer registration: NotRegistered
  Session state: Up, Since: 15:12
  Flaps: 0
  Protocol state: Idle, Since: 15:12
  Synchronization state: Complete
  Number of peers waiting: AckWait: 0, SoWait: 0, Scheduled: 0
  Messages sent: Open 1, Establish 0, GrHelper 0, Update 0, GrStaleLabel 0 Error 0, Complete 1
  Messages received: Open 1, Request 0 wildcard 0 targeted, EstablishAck 0, GrHelperAck 0,
CompleteAck 1
```

show bgp replication (for Backup)

```
user@host> show bgp replication
Synchronization backup:
  State: Established 13 ago
  , Unsync timer: 2

    Unsync entry queue:
        Instance: 0 Neighbor: 30.30.30.1 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.3 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.4 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.5 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.6 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.1 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.2 elapsed: 7
```

show bgp replication (for Backup with rib-sharding configured)

```
user@host> show bgp replication rib-sharding junos-bgpshard0
Synchronization backup:
State: Established 17:48 ago
```

Release Information

Command introduced in Junos OS Release 8.5.

Support for `logical-system` option introduced in Junos OS Release 13.3.

Support for `rib-sharding` (*`rib-shard-name`*) option introduced in Junos OS and Junos Evolved Release 22.2.

RELATED DOCUMENTATION

show bgp replication logical-system | 17223

show bgp replication logical-system

IN THIS SECTION

- Syntax | 17224
- Description | 17224
- Options | 17224
- Required Privilege Level | 17224
- Output Fields | 17224
- Sample Output | 17226
- Release Information | 17226

Syntax

```
show bgp replication logical-system
<logical-system-name>
```

Description

Display logical system-specific BGP state replication between the primary and backup logical system on Routing Engines that have nonstop active routing configured on them.

Options

This command has no options.

Required Privilege Level

View

Output Fields

Table 180 on page 17224 lists the output fields for the `show bgp replication logical-system` command. Output fields are listed in the approximate order in which they appear.

Table 180: show bgp replication logical-system Output Fields

Field Name	Field Description
session state	State of the current internal BGP state replication session, Up or Down, and the duration for which the session has been in the indicated state.
flaps	Total number of flaps that occurred.
protocol state	Current state of the protocol operation (Active, Connect, Idle) and the duration for which the protocol has been in the indicated state.

Table 180: show bgp replication logical-system Output Fields *(Continued)*

Field Name	Field Description
synchronization state	Synchronization state at the time of executing the command. The states can be: • **Idle** • **Neighbor**—Indicates that the neighbor state synchronization is in progress. • **AckWait**—Indicates that the request processing is over. • **ORF**—Indicates that the outbound routing filter synchronization is in progress. • **RIB**—Indicates that the routing table synchronization is in progress. • **Complete**
number of peers waiting	Total number of peers waiting for various messages: • **AckWait**—Number of peers waiting for connection establishment or completed acknowledgment messages. • **SoWait**—Number of peers waiting for TCP socket-related operations. • **Scheduled**—Number of peers being synchronized.
messages sent	Number of various types of messages that have been sent since internal replication session became active: • **Open**—Number of Open messages sent. • **Establish**—Number of connection establishment acknowledgment messages sent. • **Update**—Number of update messages sent. • **Error**—Number of error messages sent. • **Complete**—Number of connection complete acknowledgment messages sent.

Table 180: show bgp replication logical-system Output Fields *(Continued)*

Field Name	Field Description

messages received	Total number of messages received:
	• **Open**—Number of Open messages received.
	• **Request**—Number of request messages received:
	  • **Wildcard**—Number of requests received that used wildcards in the target address.
	  • **Targeted**—Number of requests received that used a specific address.
	• **EstablishAck**—Number of connection establishment acknowledged messages received.
	• **CompleteAck**—Number of connection completed acknowledged messages received.

Sample Output

show bgp replication logical-system

```
user@host> show bgp replication logical-system lr2
Synchronization master:
  Session state: Up, Since: 24:53
  Flaps: 0
  Protocol state: Idle, Since: 2
  Synchronization state: Complete
  Number of peers waiting: AckWait: 0, SoWait: 0, Scheduled: 0
  Messages sent: Open 1, Establish 145, Update 0, Error 1, Complete 145
  Messages received: Open 1, Request 1 wildcard 144 targeted, EstablishAck 0, CompleteAck 145
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

show bgp replication | 17218

show bgp summary

IN THIS SECTION

- Syntax | 17227
- Syntax (EX Series Switch and QFX Series) | 17227
- Description | 17227
- Options | 17228
- Required Privilege Level | 17228
- Output Fields | 17228
- Sample Output | 17232
- Release Information | 17237

Syntax

```
show bgp summary
<exact-instance instance-name>
<group group-name>
<instance instance-name>
<logical-system (all | logical-system-name)>
<rib-sharding (main | rib-shard-name)>
```

Syntax (EX Series Switch and QFX Series)

```
show bgp summary
<exact-instance instance-name>
<instance instance-name>
```

Description

Display BGP summary information.

Options

none	Display BGP summary information for all routing instances.
exact-instance ***instance-name***	(Optional) Display information for the specified instance only.
group	Display overview of bgp information for a particular group
instance ***instance-name***	(Optional) Display information for all routing instances whose name begins with this string (for example, `cust1`, `cust11`, and `cust111` are all displayed when you run the `show bgp summary instance cust1` command). The instance name can be primary for the main instance, or any valid configured instance name or its prefix.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
rib-sharding (main | *rib-shard-name*)	(Optional) Display name of rib shard. When NSR is configured, display name of the rib shard in the backup Routing Engine.

Required Privilege Level

view

Output Fields

Table 181 on page 17228 describes the output fields for the `show bgp summary` command. Output fields are listed in the approximate order in which they appear.

Table 181: show bgp summary Output Fields

Field Name	Field Description
Default eBGP mode	Default EBGP mode for receive and advertise.
`Groups`	Number of BGP groups.
`Peers`	Number of BGP peers.

Table 181: show bgp summary Output Fields *(Continued)*

Field Name	Field Description
`Down peers`	Number of down BGP peers.
`Table`	Name of routing table.
`Tot Paths`	Total number of paths.
`Act Paths`	Number of active routes.
`Suppressed`	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
`History`	Number of withdrawn routes stored locally to keep track of damping history.
`Damp State`	Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs.
`Pending`	Routes in process by BGP import policy.
`Peer`	Address of each BGP peer. Each peer has one line of output.
`AS`	Peer's AS number.
`InPkt`	Number of packets received from the peer.
`OutPkt`	Number of packets sent to the peer.
`OutQ`	Number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is `0` because the queue usually is emptied quickly.
`Flaps`	Number of times the BGP session has gone down and then come back up.

Table 181: show bgp summary Output Fields *(Continued)*

Field Name	Field Description
`Last Up/Down`	Last time since the neighbor transitioned to or from the established state.

Table 181: show bgp summary Output Fields *(Continued)*

Field Name	Field Description

`State|#Active / Received/ Accepted /Damped`	Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established on the main routing device or in a routing instance.
	• If a peer is not established, the field shows the state of the peer session: `Active`, `Connect`, or `Idle`.
	In general, the Idle state is the first stage of a connection. BGP is waiting for a Start event. A session can be idle for other reasons as well. The reason that a session is idle is sometimes displayed. For example: `Idle (Removal in progress)` or `Idle (LicenseFailure)`.
	• If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the `inet.0` (main) and `inet.2` (multicast) routing tables. For example, `8/10/10/2 and 2/4/4/0` indicate the following:
	  • 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the `inet.0` routing table.
	  • 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the `inet.2` routing table.
	• If a BGP session is established in a routing instance, the field indicates the established `(Establ)` state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, `Establ VPN-AB.inet.0: 2/4/0` indicates the following:
	  • The BGP session is established.
	  • Routes are received in the `VPN-AB.inet.0` routing table.
	  • The local routing device has two active routes, four received routes, and no damped routes from a BGP peer.
	When a BGP session is established, the peers are exchanging update messages.
	NOTE: When graceful restart or LLGR helper mode is active, the RIB information is now displayed by the `show bgp summary` command. If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following:

Table 181: show bgp summary Output Fields *(Continued)*

Field Name	Field Description
	• 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. • 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table.

Sample Output

show bgp summary (When a Peer Is Not Established)

```
user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 6          4          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/Received/
Damped...
10.0.0.3        65002         86         90       0       2       42:54 0/0/0
0/0/0
10.0.0.4        65002         90         91       0       1       42:54 0/2/0
0/0/0
10.0.0.6        65002         87         90       0       3           3 Active
10.1.12.1       65001         89         89       0       1       42:54 4/4/0
0/0/0
```

show bgp summary (When a Peer Is Established)

```
user@host> show bgp summary

Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
10.12.78.2            64531         27         26       0       0       10:49 Establ
```

```
  inet.0: 0/0/0/0

user@host> show bgp summary logical-system R3
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.0
                       2          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
1.1.1.2                   2        204        206       0       0     1:30:59 Establ
  bgp.l3vpn.0: 2/2/2/0
  red.inet.0: 2/2/2/0
10.1.1.10                 3        206        207       0       0     1:31:36 Establ
  red.inet.0: 2/2/2/0
```

show bgp summary (CLNS)

```
user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/Received/
Damped...
10.245.245.1      200       1735       1737       0       0    14:26:12 Establ
  bgp.isovpn.0: 3/3/0
  aaaa.iso.0: 3/3/0
```

show bgp summary (Layer 2 VPN)

```
user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l2vpn.0            1          1          0          0          0          0
inet.0                 0          0          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/Received/
Damped...
10.255.245.35   65299         72         74       0       1       19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0
10.255.245.36   65299       2164       2423       0       4       19:50 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
```

```
10.255.245.37   65299         36         37       0       4       17:07 Establ
  inet.0: 0/0/0
10.255.245.39   65299        138        168       0       6       53:48 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.69   65299        134        140       0       6       53:42 Establ
  inet.0: 0/0/0
```

show bgp summary (Layer 3 VPN)

```
user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State Pending
bgp.l3vpn.0            2          2          0          0          0       0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/Received/
Damped...
10.39.1.5           2         21         22       0       0        6:26 Establ
  VPN-AB.inet.0: 1/1/0
10.255.71.15        1         19         21       0       0        6:17 Establ
  bgp.l3vpn.0: 2/2/0
  VPN-A.inet.0: 1/1/0
  VPN-AB.inet.0: 2/2/0
  VPN-B.inet.0: 1/1/0
```

show bgp summary (with rib-sharding configured)

```
user@host> show bgp summary rib-sharding main
Threading mode: BGP sharding
Thread counts: Update-io: 11 Shards: 8
Groups: 11 Peers: 1010 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
1.30.1.1                  1        147     231346       0       0     4:02:10 Establ
  inet.0: 0/0/0/0
```

show bgp summary group

```
user@host> show bgp summary group Group2
Groups: 3 Peers: 3 Down peers: 3
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
10.0.0.1                 56          0          0       0       0          51 Idle
user@host> show bgp summary logical-system R3 group toR4
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.0
                       2          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
10.1.1.10                 3        207        207       0       0     1:31:40 Establ
  red.inet.0: 2/2/2/0
```

show bgp summary (BGP Graceful Restart or Long-Lived Graceful Restart)

```
user@router> show route receive-protocol bgp 10.4.12.11 detail
  Groups: 2 Peers: 9 Down peers: 1
  ...
  Peer                      AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
  10.255.255.16            100          7          6       0       4           4 Idle
    bgp.l3vpn.0: 0/7/7/0
    foo.inet.0: 0/7/7/0
```

show bgp summary

```
user@router> show bgp summary rib-sharding junos-bgpshard3

Threading mode: BGP sharding
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.rtarget.0
```

```
                      17         10          0          0          0          0
inet.0
                     500        250          0          0          0          0
inet6.0
                     500        250          0          0          0          0
bgp.l3vpn.0
                    2500       2500          0          0          0          0
bgp.l3vpn-inet6.0
                    2500       2500          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
10.1.1.3                100       1669          0       0       0  4d 1:09:05 Establ
  bgp.l3vpn-inet6.0: 2500/2500/2500/0
  bgp.l3vpn.0: 2500/2500/2500/0
  bgp.rtarget.0: 3/10/10/0
  inet.0: 0/250/250/0
  inet6.0: 0/250/250/0
10.1.1.4                100         13          0       0       0  4d 1:09:01 Establ
  bgp.l3vpn-inet6.0: 0/0/0/0
  bgp.l3vpn.0: 0/0/0/0
  bgp.rtarget.0: 7/7/7/0
  inet.0: 250/250/250/0
  inet6.0: 250/250/250/0
```

show bgp summary (with the Default EBGP mode)

```
user@router> show bgp summary

Threading mode: BGP I/O
Default eBGP mode: advertise - accept, receive - accept
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...

192.0.2.2            65551          2          2       0       0           9 Establ
  inet.0: 0/0/0/0
```

Release Information

Command introduced before Junos OS Release 7.4.

`exact-instance` option introduced in Junos OS Release 11.4.

`group` option introduced in Junos OS Release 13.3.

`rib-sharding` option introduced in cRPD Release 20.1R1.

show bridge domain

IN THIS SECTION

- Syntax | 17237
- Description | 17237
- Options | 17238
- Required Privilege Level | 17238
- Sample Output | 17238
- Release Information | 17239

Syntax

```
show bridge domain
<brief | detail | extensive>
<bridge-domain (all | domain-name)>
<instance instance-name>
<operational>
```

Description

(MX Series routers only) Display bridge domain information.

Options

none	Display information for all bridge domains.
brief | detail | extensive	(Optional) Display the specified level of output.
bridge-domain (all | *domain-name*)	(Optional) Display information about all bridge domains or the specified bridge domain.
instance *instance-name*	(Optional) Display information for the specified routing instance.
operational	(Optional) Display information for the operational routing instances.

Required Privilege Level

view

Sample Output

show bridge domain

```
user@host> show bridge domain
Instance                    Bridging Domain          Type
            Primary Table                                          Active
vs1                         vlan100                  bridge
            bridge.0                                               2
vs1                         vlan200                  bridge
            bridge.0                                               0
```

show bridge domain brief

```
user@host> show bridge domain brief
Instance                    Bridging Domain          Type
            Primary Table                                          Active
vs1                         vlan100                  bridge
            bridge.0                                               2
vs1                         vlan200                  bridge
            bridge.0                                               0
```

show bridge domain detail

```
user@host> show bridge domain detail
Routing Instance:vs1
  Bridging Domain:vlan100
  Router ID: 0.0.0.0
  Type: bridge              State: Active
  Interfaces:
    ge-11/0/3.0
    ge-11/1/4.100
    ge-11/1/1.100
    ge-11/1/0.100
    xe-10/2/0.100
    xe-10/0/0.100
  Tables:
     bridge.0               : 2 macs (2 active)
Routing Instance:vs1
  Bridging Domain:vlan200
  Router ID: 0.0.0.0
  Type: bridge              State: Active
  Interfaces:
    ge-11/1/0.200
    ge-11/1/1.200
    ge-11/1/4.200
    xe-10/0/0.200
    xe-10/2/0.200
  Tables:
     bridge.0               : 0 macs (0 active)
```

Release Information

Command introduced in Junos OS Release 8.4.

show bridge flood

IN THIS SECTION

- Syntax | 17240
- Description | 17240
- Options | 17240
- Required Privilege Level | 17241
- Output Fields | 17241
- Sample Output | 17241
- Release Information | 17249

Syntax

```
show bridge flood
<brief | detail | extensive>
<bridge-domain domain-name>
<event-queue>
<instance instance-name>
<route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-flood)>
```

Description

(MX Series routers only) Display bridging flooding information.

Options

none	Display all bridging flooding information for all bridging domains.
brief | detail | extensive	(Optional) Display the specified level of output.
bridge-domain *domain-name*	(Optional) Display bridging flooding information for the specified bridge domain.
event-queue	(Optional) Display the queue of pending bridge flood events.

instance *instance-name*

(Optional) Display bridging flooding information for the specified routing instance.

route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-flood)

(Optional) Display the following:

- **all-ce-flood**—Display the route for flooding traffic to all customer edge routers if **no-local-switching** is enabled.
- **all-ve-flood**—Display the route for flooding traffic to all VPLS edge routers if **no-local-switching** is enabled.
- **alt-root-flood**—Display the Spanning Tree Protocol (STP) alt-root flooding route used for the interface.
- **bd-flood**—Display the route for flooding traffic of a bridge domain if **no-local-switching** is not enabled.
- **mlp-flood**—Display the route for flooding traffic to MAC learning chips.
- **re-flood**—Display the route for Routing Engine flooding to all interfaces.

Required Privilege Level

view

Output Fields

to be provided

Sample Output

show bridge flood

```
user@host> show bridge flood
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood Routes:
  Prefix    Type          Owner                 NhType          NhIndex
  0x36/16   MLP_FLOOD     __vs1+vlan100__       flood           426
  0x3a/16   MLP_FLOOD     __vs1+vlan200__       flood           428
Name: vs1::vlan100
```

```
CEs: 6
VEs: 0
Flood Routes:
  Prefix    Type          Owner                 NhType          NhIndex
  0x35/16   ALL_FLOOD     __vs1+vlan100__       flood           425
  0x35/16   RE_FLOOD      __vs1+vlan100__       flood           425
  0x3780/17 ALT_ROOT_RT   ge-11/0/3.0           flood           425
  0x3b80/17 ALT_ROOT_RT   ge-11/1/4.100         flood           425
  0x3c80/17 ALT_ROOT_RT   ge-11/1/1.100         flood           425
  0x3d80/17 ALT_ROOT_RT   ge-11/1/0.100         flood           425
  0x3e80/17 ALT_ROOT_RT   xe-10/2/0.100         flood           425
  0x3f80/17 ALT_ROOT_RT   xe-10/0/0.100         flood           425
Name: vs1::vlan200
CEs: 5
VEs: 0
Flood Routes:
  Prefix    Type          Owner                 NhType          NhIndex
  0x39/16   ALL_FLOOD     __vs1+vlan200__       flood           427
  0x39/16   RE_FLOOD      __vs1+vlan200__       flood           427
  0x4180/17 ALT_ROOT_RT   ge-11/1/0.200         flood           427
  0x4080/17 ALT_ROOT_RT   ge-11/1/1.200         flood           427
  0x4280/17 ALT_ROOT_RT   ge-11/1/4.200         flood           427
  0x4480/17 ALT_ROOT_RT   xe-10/0/0.200         flood           427
  0x4380/17 ALT_ROOT_RT   xe-10/2/0.200         flood           427
```

show bridge flood brief

```
user@host> show bridge flood brief
Name                Active CEs         Active VEs
__juniper_private1__ 0                 0
vs1::vlan100        6                  0
vs1::vlan200        5                  0
```

show bridge flood detail

```
user@host> show bridge flood detail
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood Routes:
```

```
  Prefix    Type          Owner                 NhType          NhIndex
  0x36/16   MLP_FLOOD     __vs1+vlan100__       flood           426
  0x3a/16   MLP_FLOOD     __vs1+vlan200__       flood           428
Name: vs1::vlan100
CEs: 6
VEs: 0
Flood Routes:
  Prefix    Type          Owner                 NhType          NhIndex
  0x35/16   ALL_FLOOD     __vs1+vlan100__       flood           425
  0x35/16   RE_FLOOD      __vs1+vlan100__       flood           425
  0x3780/17 ALT_ROOT_RT   ge-11/0/3.0           flood           425
  0x3b80/17 ALT_ROOT_RT   ge-11/1/4.100         flood           425
  0x3c80/17 ALT_ROOT_RT   ge-11/1/1.100         flood           425
  0x3d80/17 ALT_ROOT_RT   ge-11/1/0.100         flood           425
  0x3e80/17 ALT_ROOT_RT   xe-10/2/0.100         flood           425
  0x3f80/17 ALT_ROOT_RT   xe-10/0/0.100         flood           425
Name: vs1::vlan200
CEs: 5
VEs: 0
Flood Routes:
  Prefix    Type          Owner                 NhType          NhIndex
  0x39/16   ALL_FLOOD     __vs1+vlan200__       flood           427
  0x39/16   RE_FLOOD      __vs1+vlan200__       flood           427
  0x4180/17 ALT_ROOT_RT   ge-11/1/0.200         flood           427
  0x4080/17 ALT_ROOT_RT   ge-11/1/1.200         flood           427
  0x4280/17 ALT_ROOT_RT   ge-11/1/4.200         flood           427
  0x4480/17 ALT_ROOT_RT   xe-10/0/0.200         flood           427
  0x4380/17 ALT_ROOT_RT   xe-10/2/0.200         flood           427
```

show bridge flood extensive

```
user@host> show bridge flood extensive
Name: __juniper_private1__
CEs: 0
VEs: 0
  Flood route prefix: 0x36/16
  Flood route type: MLP_FLOOD
  Flood route owner: __vs1+vlan100__
  Nexthop type: flood
  Nexthop index: 426
    Interfaces Flooding to:
```

```
    Name             Type          NhType          Index
    lc-11/0/0.32769  LC
    lc-10/2/0.32769  LC
    lc-10/0/0.32769  LC
    lc-11/1/0.32769  LC

  Flood route prefix: 0x3a/16
  Flood route type: MLP_FLOOD
  Flood route owner: __vs1+vlan200__
  Nexthop type: flood
  Nexthop index: 428
    Interfaces Flooding to:
    Name             Type          NhType          Index
    lc-10/0/0.32769  LC
    lc-10/2/0.32769  LC
    lc-11/1/0.32769  LC
Name: vs1::vlan100
CEs: 6
VEs: 0

  Flood route prefix: 0x35/16
  Flood route type: ALL_FLOOD
  Flood route owner: __vs1+vlan100__
  Nexthop type: flood
  Nexthop index: 425
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/0/3.0      CE
    ge-11/1/4.100    CE
    ge-11/1/1.100    CE
    ge-11/1/0.100    CE
    xe-10/2/0.100    CE
    xe-10/0/0.100    CE

  Flood route prefix: 0x35/16
  Flood route type: RE_FLOOD
  Flood route owner: __vs1+vlan100__
  Nexthop type: flood
  Nexthop index: 425
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/0/3.0      CE
    ge-11/1/4.100    CE
```

```
    ge-11/1/1.100    CE
    ge-11/1/0.100    CE
    xe-10/2/0.100    CE
    xe-10/0/0.100    CE

  Flood route prefix: 0x3780/17
  Flood route type: ALT_ROOT_RT
  Flood route owner: ge-11/0/3.0
  Nexthop type: flood
  Nexthop index: 425
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/0/3.0      CE
    ge-11/1/4.100    CE
    ge-11/1/1.100    CE
    ge-11/1/0.100    CE
    xe-10/2/0.100    CE
    xe-10/0/0.100    CE

  Flood route prefix: 0x3b80/17
  Flood route type: ALT_ROOT_RT
  Flood route owner: ge-11/1/4.100
  Nexthop type: flood
  Nexthop index: 425
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/0/3.0      CE
    ge-11/1/4.100    CE
    ge-11/1/1.100    CE
    ge-11/1/0.100    CE
    xe-10/2/0.100    CE
    xe-10/0/0.100    CE

  Flood route prefix: 0x3c80/17
  Flood route type: ALT_ROOT_RT
  Flood route owner: ge-11/1/1.100
  Nexthop type: flood
  Nexthop index: 425
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/0/3.0      CE
    ge-11/1/4.100    CE
    ge-11/1/1.100    CE
```

```
    ge-11/1/0.100    CE
    xe-10/2/0.100    CE
    xe-10/0/0.100    CE

  Flood route prefix: 0x3d80/17
  Flood route type: ALT_ROOT_RT
  Flood route owner: ge-11/1/0.100
  Nexthop type: flood
  Nexthop index: 425
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/0/3.0      CE
    ge-11/1/4.100    CE
    ge-11/1/1.100    CE
    ge-11/1/0.100    CE
    xe-10/2/0.100    CE
    xe-10/0/0.100    CE

  Flood route prefix: 0x3e80/17
  Flood route type: ALT_ROOT_RT
  Flood route owner: xe-10/2/0.100
  Nexthop type: flood
  Nexthop index: 425
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/0/3.0      CE
    ge-11/1/4.100    CE
    ge-11/1/1.100    CE
    ge-11/1/0.100    CE
    xe-10/2/0.100    CE
    xe-10/0/0.100    CE

  Flood route prefix: 0x3f80/17
  Flood route type: ALT_ROOT_RT
  Flood route owner: xe-10/0/0.100
  Nexthop type: flood
  Nexthop index: 425
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/0/3.0      CE
    ge-11/1/4.100    CE
    ge-11/1/1.100    CE
    ge-11/1/0.100    CE
```

```
    xe-10/2/0.100    CE
    xe-10/0/0.100    CE
Name: vs1::vlan200
CEs: 5
VEs: 0

  Flood route prefix: 0x39/16
  Flood route type: ALL_FLOOD
  Flood route owner: __vs1+vlan200__
  Nexthop type: flood
  Nexthop index: 427
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/1/0.200    CE
    ge-11/1/1.200    CE
    ge-11/1/4.200    CE
    xe-10/0/0.200    CE
    xe-10/2/0.200    CE

  Flood route prefix: 0x39/16
  Flood route type: RE_FLOOD
  Flood route owner: __vs1+vlan200__
  Nexthop type: flood
  Nexthop index: 427
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/1/0.200    CE
    ge-11/1/1.200    CE
    ge-11/1/4.200    CE
    xe-10/0/0.200    CE
    xe-10/2/0.200    CE

  Flood route prefix: 0x4180/17
  Flood route type: ALT_ROOT_RT
  Flood route owner: ge-11/1/0.200
  Nexthop type: flood
  Nexthop index: 427
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/1/0.200    CE
    ge-11/1/1.200    CE
    ge-11/1/4.200    CE
    xe-10/0/0.200    CE
```

```
    xe-10/2/0.200    CE

  Flood route prefix: 0x4080/17
  Flood route type: ALT_ROOT_RT
  Flood route owner: ge-11/1/1.200
  Nexthop type: flood
  Nexthop index: 427
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/1/0.200    CE
    ge-11/1/1.200    CE
    ge-11/1/4.200    CE
    xe-10/0/0.200    CE
    xe-10/2/0.200    CE

  Flood route prefix: 0x4280/17
  Flood route type: ALT_ROOT_RT
  Flood route owner: ge-11/1/4.200
  Nexthop type: flood
  Nexthop index: 427
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/1/0.200    CE
    ge-11/1/1.200    CE
    ge-11/1/4.200    CE
    xe-10/0/0.200    CE
    xe-10/2/0.200    CE

  Flood route prefix: 0x4480/17
  Flood route type: ALT_ROOT_RT
  Flood route owner: xe-10/0/0.200
  Nexthop type: flood
  Nexthop index: 427
    Interfaces Flooding to:
    Name             Type          NhType          Index
    ge-11/1/0.200    CE
    ge-11/1/1.200    CE
    ge-11/1/4.200    CE
    xe-10/0/0.200    CE
    xe-10/2/0.200    CE

  Flood route prefix: 0x4380/17
  Flood route type: ALT_ROOT_RT
```

```
  Flood route owner: xe-10/2/0.200
  Nexthop type: flood
  Nexthop index: 427
    Interfaces Flooding to:
    Name              Type          NhType          Index
    ge-11/1/0.200     CE
    ge-11/1/1.200     CE
    ge-11/1/4.200     CE
    xe-10/0/0.200     CE
    xe-10/2/0.200     CE
```

Release Information

Command introduced in Junos OS Release 8.4.

show bridge flood nexthops satellite

IN THIS SECTION

- Syntax | 17249
- Description | 17250
- Options | 17250
- Required Privilege Level | 17250
- Output Fields | 17250
- Sample Output | 17252
- Release Information | 17253

Syntax

```
show bridge flood nexthops satellite
<brief | detail | extensive>
<nexthop-id nexthop-id>
```

Description

Display bridge domain flood next hop information for satellite device destinations.

You can use this command to:

- View the current list of all flood traffic composite next hops.
- See details about a specified composite next hop.
- Follow aggregation device composite next-hop processing as the aggregation device resolves and updates composite next-hop table entries for extended port destinations.

The aggregation device allocates ECID tags that represent multicast or broadcast destinations behind satellite device extended ports, associates them with the corresponding satellite device virtual interfaces (`sd-`*`fpc-id`*`/0/0`), and updates flood next-hop table entries accordingly. More detailed output from this command shows events that result in next-hop table updates.

Options

brief | detail | extensive (Optional) Display the specified level of output. The default output level is `brief`.

nexthop-id *nexthop-id* Display more detailed bridge flooding next hop information only for the specified next hop.

Required Privilege Level

view

Output Fields

Table 182 on page 17250 lists the output fields for the `show bridge flood next-hops satellite` command. Output fields are listed in the approximate order in which they appear.

Table 182: show bridge flood next-hops satellite Command Output Fields

Field Name	Field Description	Level of Output
`Next-hop ID`	Next-hop ID for each next- hop entry displayed.	All

Table 182: show bridge flood next-hops satellite Command Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Composite function`	Purpose of a composite next-hop entry. Values include: • FLOOD_ALL—Flood next hop. • FLOOD_ALL_SPLIT_HZ—Flood next hop including split-horizon information.	All
`Table`	Name of next-hop routing or forwarding table for the listed entry.	All
`Flags`	Flags giving additional information about a next-hop entry. Values include: • SAT—The entry is a satellite destination composite next hop for an extended port destination that has been resolved and updated with the corresponding satellite device interface, `sd-`*`fpc-id`*`/0/0`, and the associated ECID shown in the `label` field. • ST—The entry is stale and waiting to be refreshed. • RU—The entry is stale, but is marked to be reused.	All
`aggregation-device`	Next-hop IDs and corresponding interfaces to reach local flood destination ports on the aggregation device.	detail extensive
`satellite-device-id` *`id`*	Satellite device ID with the next-hop IDs and interface names to reach extended ports that are flood destinations on the listed satellite device. When an ECID has been assigned to destination extended ports on a satellite device, this field lists the satellite next-hop ID, the corresponding virtual satellite device interface (`sd-`*`fpc-id`*`/0/0`), and the allocated ECID (`label` field).	detail extensive

Table 182: show bridge flood next-hops satellite Command Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`label`	ECID associated with a group of flood destination extended ports on the satellite device specified in the `satellite-device-id` field.	detail extensive
`When`	Elapsed time since an event related to a flood next-hop entry change.	extensive
`Event`	Brief description of the event related to a flood next-hop entry change.	extensive
`Action`	Brief description of actions that resulted from the event.	extensive

Sample Output

show bridge flood next hops

```
user@host> > show bridge flood next-hops satellite nexthop-id 962 detail
Next-hop    Composite           Table                   Flags
ID          function
962         FLOOD_ALL_SPLIT_HZ  default-switch
    satellite-device-id 100:
        ->610 xe-100/0/6.0
        ->611 xe-100/0/7.0
    satellite-device-id 103:
        ->682 xe-103/0/30.0
        ->683 xe-103/0/31.0
        ->684 xe-103/0/32.0
        ->685 xe-103/0/33.0
        ->686 xe-103/0/34.0
        ->687 xe-103/0/35.0
```

command-name

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding Multicast Replication in a Junos Fusion

Egress (Local) Replication on the Satellite Devices

show bridge flood satellite

IN THIS SECTION

- Syntax | 17253
- Description | 17254
- Options | 17254
- Required Privilege Level | 17254
- Output Fields | 17254
- Sample Output | 17256
- Release Information | 17257

Syntax

```
show bridge flood satellite
<brief | detail | extensive>
<bridge-domain-name bridge-domain-name>
<vswitch-name vswitch-name>
```

Description

Display bridge flood routing information for the satellite devices in Junos Fusion Provider Edge.

This command lists flood routes by route prefix for each bridge domain. Each flood route prefix entry lists the ingress replication next-hop ID (`NhIndex`). When egress (local) replication is enabled and the bridge domain has multiple destination extended ports on a satellite device, the aggregation device:

- Creates a satellite device next-hop chain to reach those destinations through their corresponding satellite devices.
- Updates the flood route entry with a satellite next-hop chain ID (`Satellite-Nh`).

When a flood route does not have a satellite next-hop chain, the value `0` is displayed in the `Satellite-Nh` column. When the `Satellite-Nh` value is non-zero, the aggregation device uses the satellite next-hop chain instead of the original ingress replication next-hop (`NhIndex`). You can see satellite device flood next-hop chain details, including the ECIDs assigned to satellite device flood destination extended ports, using the `detail` option and the `bridge-domain-name` option for a specific bridge domain.

Options

brief | detail | extensive	(Optional) Display the specified level of output.
bridge-domain-name ***bridge-domain-name***	Display bridge flooding information for the specified bridge domain.
vswitch-name ***vswitch-name***	Display bridge flood information for the specified virtual satellite device.

Required Privilege Level

view

Output Fields

Table 183 on page 17255 lists the output fields for the `show bridge flood satellite` command. Output fields are listed in the approximate order in which they appear, although the display order varies between the different levels of output.

Table 183: show bridge flood satellite Command Output Fields

Field Name	Field Description	Level of Output
`Bridging domain`	Bridge domain name.	All
`Flood Routes`	Flood route information listed by route prefix.	All
`Prefix`	Flood route prefix.	All
`Token`	Internal ID for flood route prefix.	All
`NhIndex`	The ingress replication next-hop for the flood route.	All
`Satellite-Nh`	Satellite device next-hop ID when local replication is enabled and there are multiple flood destination extended ports.	All
`Next-hop information`	Details for each next hop, listed by next-hop or satellite next-hop iID.	detail extensive
`aggregation-device`	Next-hop interfaces for ports that are flood destinations on the aggregation device, listed by next-hop ID.	detail extensive
`satellite-device-id` *`id`*	Next-hop interfaces for extended ports that are flood destinations on satellite devices, listed by satellite next-hop ID.	detail extensive
`label`	ECID assigned to a satellite device interface.	detail extensive
`When`	Elaspsed time since an event related to a flood route entry change.	extensive

Table 183: show bridge flood satellite Command Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Event	Brief description of the event related to a flood route entry change.	extensive
Action	Brief description of actions that resulted from the event.	extensive

Sample Output

show bridge flood satellite

```
user@host> show bridge flood satellite
Paste router command output hereVSwitch instance: default-switch
  ------------ Bridging domain: bd100 ------------
    Flood Routes:
    Prefix                   Token         NhIndex      Satellite-Nh
    00.02.00.01.60.03.e0/51  0x3001F       797          1186
    00.02.ff.fe.60.00.20/51  0x30001       794          1183
  ------------ Bridging domain: bd1000 ------------
    Flood Routes:
    Prefix                   Token         NhIndex      Satellite-Nh
    00.03.00.01.60.04.00/51  0x30020       806          0
    00.03.ff.fe.60.00.40/51  0x30002       803          0
  ------------ Bridging domain: bd1500 ------------
    Flood Routes:
    Prefix                   Token         NhIndex      Satellite-Nh
    00.08.00.01.60.04.a0/51  0x30025       851          1475
    00.08.ff.fe.60.00.e0/51  0x30007       848          1472
  ------------ Bridging domain: bd1600 ------------
    Flood Routes:
    Prefix                   Token         NhIndex      Satellite-Nh
    00.09.00.01.60.04.c0/51  0x30026       860          1481
    00.09.ff.fe.60.01.00/51  0x30008       857          1478
  ------------ Bridging domain: bd1700 ------------
    Flood Routes:
    Prefix                   Token         NhIndex      Satellite-Nh
```

```
    00.0a.00.01.60.04.e0/51  0x30027       869          1487
    00.0a.ff.fe.60.01.20/51  0x30009       866          1484
  ------------ Bridging domain: bd1800 ------------
    Flood Routes:
    Prefix                   Token         NhIndex      Satellite-Nh
    00.0b.00.01.60.05.00/51  0x30028       878          1493
    00.0b.ff.fe.60.01.40/51  0x3000A       875          1490
  ------------ Bridging domain: bd1900 ------------
    Flood Routes:
    Prefix                   Token         NhIndex      Satellite-Nh
    00.0c.00.01.60.05.20/51  0x30029       887          1499
    00.0c.ff.fe.60.01.60/51  0x3000B       884          1496
  ------------ Bridging domain: bd2000 ------------
    Flood Routes:
    Prefix                   Token         NhIndex      Satellite-Nh
    00.0e.00.01.60.05.60/51  0x3002B       905          1511
    00.0e.ff.fe.60.01.a0/51  0x3000D       902          1508
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding Multicast Replication in a Junos Fusion

Egress (Local) Replication on the Satellite Devices

show bridge mac-ip-table

IN THIS SECTION

- Syntax | **17258**
- Description | **17258**
- Options | **17258**
- Required Privilege Level | **17258**

- Output Fields | 17259
- Sample Output | 17262
- Release Information | 17263

Syntax

```
show bridge mac-ip-table
<bridge-domain (all |bridge-domain-name)> <instance instance-name>
<brief | count | detail | detail | extensive>
<instance>
<kernel >
<logical-system (all | logical-system-name)>
```

Description

Displays the MAC-IP address for all IPv4 (ARP) and IPv6 (ND) bindings for bridge domains on a virtual switch.

Options

bridge-domain (all | *bridge-domain-name*)	(Optional) Display the MAC-IP address entries for all bridge domains or for a specified bridge domain.
brief | detail |extensive | summary	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display the MAC-IP address information for the specified routing instance.
kernel	(Optional) Display the MAC-IP address entries in the kernel
logical-system <*logical-system-name*>	(Optional) Display the MAC-IP address information for the specified logical system or all logical systems.

Required Privilege Level

maintenance

Output Fields

Table 184 on page 17259 lists the output fields for the `show bridge mac-ip-table` command. Output fields are listed in the approximate order in which they appear.

Table 184: show bridge mac-ip-table Output Fields

Field Name	Field Description	Level of Output
IP address	IP address associated with the EVPN routing instance.	All levels
MAC address	MAC address	All levels

Table 184: show bridge mac-ip-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	MAC IP flags: Identifies : statically installed MAC-IP entries. • S—Statically installed MAC-IP entries. • D—Dynamic installed MAC-IP address entries. • L—Locally learned MAC-IP address entries • R— MAC-IP address entries that are learnt from a remote device via the control plane. • Lp—Local proxy MAC IP address entries. In multihoming networks, a device that learns of a MAC address as a DLR (local and remote) and also learns the MAC-IP address binding remotely will send a MAC-IP address notification message (as if it was generated by traffic). This entry is identified as a local proxy address. The remote multihomed peer identifies the corresponding MAC-IP entry as remote proxy address. • Rp—Remote proxy MAC IP address entries. See Lp flag for more information • K—MAC-IP address entries that are installed in the kernel.	Brief, detail

Table 184: show bridge mac-ip-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• RT—Destination Route that corresponds to an installed entry. • AD—Advertised to remote device. • NAD—MAC-IP address entries on disabled IRB interfaces. These entries are not advertised. • RE—Specifies temporary ARP/ND entries while the device is waiting for a response from an arp or network solicitation request. The entries are removed after three unsuccessful request attempts. • RO—Router. This applies only to IPv6 MAC-IP addresses. Refer to RFC 4861 for more information on the Router flag. • OV—Override.This applies only to IPv6 MAC-IP addresses. Refer to RFC 4861 for more information on the Override flag. • UR—Unresolved MAC-IP address entries. This flag is specific to the Junos EVO environment • RTS—Device skipped adding an EVPN Type 2 route in favor of an EVPN Type 5 route with a matching prefix (extensive output displays dest_route_skipped).	

Table 184: show bridge mac-ip-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Logical Interface	The logical interface associated with the routing instance.	Brief, detail
Active Source	The source of the learned MAC-IP address entry. Displays the ESI or router ID.	All levels
Routing Instance	Information on the Routing instance where the MAC-IP address entry was learned. • Bridging domain • MAC-IP local mask • MAC-IP remote mask • MAC-IP RPD flags • MAC-IP flags	Extensive
BD ID	Bridge Domain ID	Brief

Sample Output

show bridge mac-ip-table

```
user@host> show bridge mac-ip-table
MAC IP flags  (S - Static, D - Dynamic, L - Local , R - Remote, K - Kernel, RT - Dest Route,
               AD - Advt to remote, RE - Re-ARP/ND, RO - Router, OV - Override)
 Routing instance : evpn
 Bridging domain : bd1
   IP                          MAC                 Flags            Logical            Active
   address                     address                              Interface          source
   10.1.1.1                    00:11:00:00:00:01   DR,K
81.1.1.1
```

```
   10.1.1.2                    00:22:00:00:00:01    DR,K
81.2.2.2
   10.1.1.3                    00:33:00:00:00:01    DL,K,AD           ge-0/0/2.1
   10.1.1.4                    00:44:00:00:00:01    DR,K
81.4.4.4
```

show bridge mac-ip-table extensive

```
user@host> show bridge mac-ip-table extensive
IP address: 10.1.1.1
  MAC address: 00:11:00:00:00:01
  Routing instance: evpn
   Bridging domain: bd1
   MAC-IP local mask: 0x0000000000000000
   MAC-IP remote mask: 0x8000000000000000
   MAC-IP RPD flags: 0x00000081
   MAC-IP flags: remote,kernel
```

show bridge mac-ip-table kernel

```
user@host> show bridge mac-ip-table kernel
    BD       IP                          MAC                 Logical           Flags
    id       address                     address             Interface
    2        192.168.1.1                 00:21:59:aa:77:f0   irb.1             0x00000209
    2        192.168.1.100               00:00:5e:00:01:01   irb.1             0x00000609
    3        192.168.2.1                 00:21:59:aa:77:f0   irb.2             0x00000209
    3        192.168.2.100               00:00:5e:00:01:01   irb.2             0x00000609
    4        192.168.3.1                 00:21:59:aa:77:f0   irb.3             0x00000209
    4        192.168.3.100               00:00:5e:00:01:01   irb.3             0x00000609
    5        192.168.4.1                 00:21:59:aa:77:f0   irb.4             0x00000209
    5        192.168.4.100               00:00:5e:00:01:01   irb.4             0x00000609
```

Release Information

Command introduced in Junos OS Release 17.4R2.

show bridge mac-table

IN THIS SECTION

- Syntax | 17264
- Description | 17264
- Options | 17264
- Additional Information | 17265
- Required Privilege Level | 17265
- Output Fields | 17265
- Sample Output | 17267
- Release Information | 17271

Syntax

```
show bridge mac-table
<age>
<brief | count | detail | extensive>
<bridge-domain (all | bridge-domain-name)>
<global-count>
<instance instance-name>
<interface interface-name>
<mac-address>
<instance instance-name>
<vlan-id (all-vlan | vlan-id)>
```

Description

(MX Series routers only) Display Layer 2 MAC address information.

Options

none	Display all learned Layer 2 MAC address information.

age	(Optional) Display age of a single mac-address.
brief | count | detail | extensive	(Optional) Display the specified level of output.
bridge-domain (all | *bridge-domain-name*)	(Optional) Display learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.
global-count	(Optional) Display the total number of learned Layer 2 MAC addresses on the system.
instance *instance-name*	(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.
interface *interface-name*	(Optional) Display learned Layer 2 MAC addresses for the specified interface.
mac-address	(Optional) Display the specified learned Layer 2 MAC address information.
vlan-id (all-vlan | *vlan-id*)	(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

Additional Information

When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.

Required Privilege Level

view

Output Fields

Table 185 on page 17266 describes the output fields for the `show bridge mac-table` command. Output fields are listed in the approximate order in which they appear.

Table 185: show bridge mac-table Output Fields

Field Name	Field Description
Age	Age of a single mac-address.
`Routing instance`	Name of the routing instance.
Bridging domain	Name of the bridging domain.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: • `S`—Static MAC address is configured. • `D`—Dynamic MAC address is configured. • `L`—Locally learned MAC address is configured. • `C`—Control MAC address is configured. • `SE`—MAC accounting is enabled. • `NM`—Non-configured MAC. • `R`—Remote PE MAC address is configured. • `P`—MAC Pinned interface is configured
Logical interface	Name of the logical interface.
MAC count	Number of MAC addresses learned on the specific routing instance or interface.
Learning interface	Name of the logical interface on which the MAC address was learned.
Learning VLAN	VLAN ID of the routing instance or bridge domain in which the MAC address was learned.
VXLAN ID/VXLAN	VXLAN Network Identifier (VNI).

Table 185: show bridge mac-table Output Fields *(Continued)*

Field Name	Field Description
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

Sample Output

show bridge mac-table

```
user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : test1, VLAN : 1
   MAC                 MAC      Logical          NH     RTR
   address             flags    interface        Index  ID
   01:00:0c:cc:cc:cc   S,NM     NULL
   01:00:0c:cc:cc:cd   S,NM     NULL
   01:00:0c:cd:cd:d0   S,NM     NULL
   64:87:88:6a:17:d0   D        ae0.1
   64:87:88:6a:17:f0   D        ae0.1
```

show bridge mac-table (with Layer 2 Services over GRE Interfaces)

```
user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : vlan-1, VLAN : 1
   MAC                 MAC      Logical
   address             flags    interface
   00:01:01:00:01:f7   D,SE     gr-1/2/10.0
   00:03:00:32:01:f7   D,SE     gr-1/2/10.0
   00:00:21:11:11:10   DL       ge-1/0/0.0
   00:00:21:11:11:11   DL       ge-1/1/0.0


Routing instance : default-switch
 Bridging domain : vlan-2, VLAN : 2
   MAC                 MAC      Logical
   address             flags    interface
   00:02:01:33:01:f7   D,SE     gr-1/2/10.1
   00:00:21:11:21:10   DL       ge-1/0/0.1
   00:00:21:11:21:11   DL       ge-1/1/0.1
```

show bridge mac-table (with VXLAN enabled)

```
user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : vlan-1, VLAN : 1
  VXLAN: Id : 100, Multicast group: 233.252.0.1
   MAC                 MAC      Logical
   address             flags    interface
   00:01:01:00:01:f7   D,SE     vtep.1052010
   00:03:00:32:01:f7   D,SE     vtep.1052011
   00:00:21:11:11:10   DL       ge-1/0/0.0
   00:00:21:11:11:11   DL       ge-1/1/0.0
```

```
Routing instance : default-switch
 Bridging domain : vlan-2, VLAN : 2, VXLAN : 200
  VXLAN: Id : 200, Multicast group: 233.252.0.2
   MAC                 MAC      Logical
   address             flags    interface
   00:02:01:33:01:f7   D,SE     vtep.1052010
   00:04:00:14:01:f7   D,SE     vtep.1052011
   00:00:21:11:21:10   DL       ge-1/0/0.1
   00:00:21:11:21:11   DL       ge-1/1/0.1
```

show bridge mac-table age (for GE interface)

```
user@host> show vpls mac-table age 00:02:03:aa:bb:1a instance vpls_instance_1
MAC Entry Age information
Current Age: 4 seconds
```

show bridge mac-table age (for AE interface)

```
user@host> show vpls mac-table age 00:02:03:aa:bb:1a instance vpls_instance_1
MAC Entry Age information
Current Age on FPC1: 102 seconds
Current Age on FPC2: 94 seconds
```

show bridge mac-table count

```
user@host> show bridge mac-table count
2 MAC address learned in routing instance vs1 bridge domain vlan100

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    ge-11/0/3.0                      1
    ge-11/1/4.100                    0
    ge-11/1/1.100                    0
    ge-11/1/0.100                    0
    xe-10/2/0.100                    1
    xe-10/0/0.100                    0
```

```
  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID            MAC count
                0                    2

0 MAC address learned in routing instance vs1 bridge domain vlan200

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    ge-11/1/0.200                    0
    ge-11/1/1.200                    0
    ge-11/1/4.200                    0
    xe-10/0/0.200                    0
    xe-10/2/0.200                    0

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID            MAC count
                0                    0
```

show bridge mac-table detail

```
user@host> show bridge mac-table detail
MAC address: 00:00:00:19:1c:db
  Routing instance: vs1
   Bridging domain: vlan100
   Learning interface: ge-11/0/3.0     Learning VLAN: 0
   Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
   Epoch: 4                            Sequence number: 0
   Learning mask: 0x800                IPC generation: 0

MAC address: 00:00:00:59:3a:2f
  Routing instance: vs1
   Bridging domain: vlan100
   Learning interface: xe-10/2/0.100   Learning VLAN: 0
   Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
   Epoch: 7                            Sequence number: 0
   Learning mask: 0x400                IPC generation: 0
```

show bridge mac-table instance pbb-evpn

```
user@host> show bridge mac-table instance pbb-evpn
Routing instance : pbb-evpn
 Bridging domain : isid-bd10000, ISID : 10000
   MAC                      MAC      Logical          NH     RTR
   address                 flags      interface        Index  ID
   00:19:e2:b0:76:eb   D         cbp.1000
   aa:bb:cc:dd:ee:f2   DC                        1048576 1048576
   aa:bb:cc:dd:ee:f3   DC                        1048575 1048575
```

show bridge mac-table

```
user@host>run show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC, P -Pinned MAC)

Routing instance : VS-541
Bridging domain : 541, VLAN : 541
MAC MAC Logical NH RTR
address flags interface Index ID
00:00:01:00:00:01 DPRC xe-0/0/3.0
00:00:02:00:00:01 DP xe-0/0/3.0
```

Release Information

Command introduced in Junos OS Release 8.4.

Support for PBB-EVPN instance added in Junos OS Release 16.1

MAC Flag P to indicate a MAC Pinned interface introduced in Junos OS 16.2

show bridge satellite device

IN THIS SECTION

- Syntax | 17272
- Description | 17272
- Options | 17272
- Required Privilege Level | 17272
- Output Fields | 17273
- Sample Output | 17275
- Release Information | 17275

Syntax

```
show bridge satellite device
<brief | detail>
< device-id device-id>
```

Description

Display status and control information for all satellite devices or a specified satellite device.

Options

brief | detail (Optional) Display the specified level of output. The default output level is `brief`.

`device-id` *`device-id`* (Optional) Display information for the specified satellite device.

Required Privilege Level

view

Output Fields

Table 186 on page 17273 lists the output fields for the `show bridge satellite device` command. Output fields are listed in the approximate order in which they appear, although the display order varies between the different levels of output.

Table 186: show bridge satellite device Command Output Fields

Field Name	Field Description	Level of Output
`Device ID` `Device` (detail view)	Satellite device ID.	All
`Interface Index` `Device Interface Index` (detail view)	Internal ID for the satellite device virtual interface.	All
`Interface Name` `Device Interface Name` (detail view)	Satellite device virtual interface name (`sd-`*`fpc-id`*`/0/0`, where *fpc-id* is the satellite device ID).	All
`State` `Device Interface State` (detail view)	State of the satellite device virtual interface. If the interface is created and active, possible values include `Up` or `Down`.	All

Table 186: show bridge satellite device Command Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Connection State` `Device connection Status` (detail view)	State of the satellite device interface connection to the aggregation device. The connection state is monitored using keep-alive messages between satellite and aggregation device control processes. Possible values include `Up` or `Down`.	All
`Requests`	Number of request messages sent to the listed satellite device aggregation device to allocate or update ECID mappings.	All
`Responses`	Number of response messages returned from the listed satellite device to the aggregation device for granting ECID requests.	All
`Device connection uptime`	Duration of the connection between the aggregation device and the specified satellite device.	detail
`Device heartbeat status`	Status of keep-alive message exchange between satellite device and aggregation device control processes.	detail
`Echo packets sent`	Number of keep-alive packets sent to the satellite device from the aggregation device.	detail
`Echo packets received`	Number of keep-alive response packets sent to the aggregation device from the satellite device.	detail
`Multicast IPC stats`	Number of inter-process control (IPC) messages sent from the aggregation device to the satellite device related to multicast functions on the satellite device. This value is displayed for active satellite device connections.	detail

Table 186: show bridge satellite device Command Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Bridge IPC stats`	Number of IPC messages sent from the aggregation device to the satellite device related to Layer 2 bridging functions on the satellite device. This value is displayed for active satellite device connections.	detail

Sample Output

show bridge satellite device

```
user@host> show bridge satellite device
Device    Interface  Interface                  Connection
ID        index      Name                State  State      Requests   Responses
100       370        sd-100/0/0.32770    Up     Up         5          5
101       342        sd-101/0/0.32770    Up     Up         4          4
102       365        sd-102/0/0.32770    Up     Up         3          3
105       364        sd-105/0/0.32770    Up     Up         2          2
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding Multicast Replication in a Junos Fusion

show bridge statistics

IN THIS SECTION

- Syntax | 17276
- Description | 17276
- Options | 17276
- Required Privilege Level | 17276
- Sample Output | 17277
- Release Information | 17281

Syntax

```
show bridge statistics
<bridge-domain domain-name>
<instance instance-name>
```

Description

(MX Series routers only) Display bridge statistics.

Options

none	Display bridge statistics for all bridge domains in all routing instances.
bridge-domain *domain-name*	(Optional) Display statistics for the specified bridge domain.
instance *instance-name*	(Optional) Display statistics for the specified routing instance.

Required Privilege Level

view

Sample Output

show bridge statistics

```
user@host> show bridge statistics
Information for routing instance:

Routing instance : __juniper_private1__
   Index: 1                        Sequence number: 0
   MAC limit: 5000                 MACs learned: 0
   Static MACs learned: 0          Non config Static MACs learned: 0
   Handle: 0x829e800

Information for routing instance:

Routing instance : vs1
 Bridging domain : vlan100
   Index: 3                        Sequence number: 0
   MAC limit: 5120                 MACs learned: 2
   Static MACs learned: 0          Non config Static MACs learned: 0
   Handle: 0x829e400
   Flags: Bridge instance, Config defined, VLAN : 100
   Local interface: ge-11/0/3.0, Index: 79
     Broadcast packets:                     1
     Broadcast bytes  :                    65
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :             358624489
     Unicast bytes    :           23310592305
     Current MAC count:                     1 (Limit 1024)
   Local interface: ge-11/1/4.100, Index: 84
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :                     0
     Unicast bytes    :                     0
     Current MAC count:                     0 (Limit 1024)
```

```
    Local interface: ge-11/1/1.100, Index: 86
      Broadcast packets:                     0
      Broadcast bytes  :                     0
      Multicast packets:                     0
      Multicast bytes  :                     0
      Flooded packets  :                     0
      Flooded bytes    :                     0
      Unicast packets  :                     0
      Unicast bytes    :                     0
      Current MAC count:                     0 (Limit 1024)
    Local interface: ge-11/1/0.100, Index: 87
      Broadcast packets:                     0
      Broadcast bytes  :                     0
      Multicast packets:                     0
      Multicast bytes  :                     0
      Flooded packets  :                     0
      Flooded bytes    :                     0
      Unicast packets  :                     0
      Unicast bytes    :                     0
      Current MAC count:                     0 (Limit 1024)
    Local interface: xe-10/2/0.100, Index: 88
      Broadcast packets:                     0
      Broadcast bytes  :                     0
      Multicast packets:                     0
      Multicast bytes  :                     0
      Flooded packets  :                     0
      Flooded bytes    :                     0
      Unicast packets  :             358627393
      Unicast bytes    :           23310781065
      Current MAC count:                     1 (Limit 1024)
    Local interface: xe-10/0/0.100, Index: 89
      Broadcast packets:                     0
      Broadcast bytes  :                     0
      Multicast packets:                     0
      Multicast bytes  :                     0
      Flooded packets  :                     0
      Flooded bytes    :                     0
      Unicast packets  :                     0
      Unicast bytes    :                     0
      Current MAC count:                     0 (Limit 1024)

Information for routing instance:
```

```
Routing instance : vs1
 Bridging domain : vlan200
   Index: 4                        Sequence number: 0
   MAC limit: 5120                 MACs learned: 0
   Static MACs learned: 0          Non config Static MACs learned: 0
   Handle: 0x829e600
   Flags: Bridge instance, Config defined, VLAN : 200
   Local interface: ge-11/1/0.200, Index: 90
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :                     0
     Unicast bytes    :                     0
     Current MAC count:                     0 (Limit 1024)
   Local interface: ge-11/1/1.200, Index: 91
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :                     0
     Unicast bytes    :                     0
     Current MAC count:                     0 (Limit 1024)
   Local interface: ge-11/1/4.200, Index: 92
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :                     0
     Unicast bytes    :                     0
     Current MAC count:                     0 (Limit 1024)
   Local interface: xe-10/0/0.200, Index: 93
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
```

```
    Flooded bytes    :                    0
    Unicast packets  :                    0
    Unicast bytes    :                    0
    Current MAC count:                    0 (Limit 1024)
  Local interface: xe-10/2/0.200, Index: 94
    Broadcast packets:                    4
    Broadcast bytes  :                  260
    Multicast packets:                    0
    Multicast bytes  :                    0
    Flooded packets  :                    0
    Flooded bytes    :                    0
    Unicast packets  :                    0
    Unicast bytes    :                    0
    Current MAC count:                    0 (Limit 1024)
```

show bridge statistics (vlan, vni)

```
user@host> show bridge statistics
  Local interface: vtep.32777, Index: 371, VLAN ID: 200, VNI: 200
    Broadcast packets:                    0
    Broadcast bytes  :                    0
    Multicast packets:                   16
    Multicast bytes  :                 1280
    Flooded packets  :                    0
    Flooded bytes    :                    0
    Unicast packets  :                    0
    Unicast bytes    :                    0
    Current MAC count:                    1
  Local interface: vtep.32777, Index: 371, VLAN ID: 100, VNI: 100
    Broadcast packets:                    0
    Broadcast bytes  :                    0
    Multicast packets:                   16
    Multicast bytes  :                 1280
    Flooded packets  :                    0
    Flooded bytes    :                    0
    Unicast packets  :                    0
    Unicast bytes    :                    0
    Current MAC count:                    1
```

Release Information

Command introduced in Junos OS Release 8.4.

show captive-portal authentication-failed-users

IN THIS SECTION

- Syntax | 17281
- Description | 17281
- Required Privilege Level | 17281
- Output Fields | 17281
- Sample Output | 17282
- Release Information | 17282

Syntax

```
show captive-portal authentication-failed-users
```

Description

Display the users that have failed captive portal authentication.

Required Privilege Level

view

Output Fields

Table 187 on page 17282 lists the output fields for the `show captive-portal authentication-failed-users` command. Output fields are listed in the approximate order in which they appear.

Table 187: show captive-portal authentication-failed-users Output Fields

Field Name	Field Description	Level of Output
Interface	The MAC address configured to bypass captive portal authentication.	**all**
MAC address	The MAC address configured statically on the interface.	**all**
User	Name of the user that has failed captive portal authentication.	**all**
Failure Count	The number of times that 802.1X authentication has failed on the interface.	**all**

Sample Output

show captive-portal authentication-failed-users

```
user@host> show captive-portal authentication-failed-users

Interface     MAC address          User            Failure Count
ge-0/0/17.0   00:37:00:00:00:00    003700000000        28
ge-0/0/20.0   00:04:10:00:00:00    000410000000        32
ge-0/0/18.0   00:00:03:00:0a:00    000003000a00        4
ge-0/0/19.0   00:00:03:00:0b:00    000003000b00        18
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

show captive-portal interface | **17286**

show captive-portal firewall | **17283**

clear captive-portal | **3450**

Example: Setting Up Captive Portal Authentication on an EX Series Switch

show captive-portal firewall

IN THIS SECTION

- Syntax | 17283
- Description | 17283
- Options | 17283
- Required Privilege Level | 17284
- Output Fields | 17284
- Sample Output | 17284
- Release Information | 17285

Syntax

```
show captive-portal firewall
<brief | detail>
<interface-name>
<interface-name detail>
```

Description

Display information about the firewall filters for each user that is authenticated on each captive portal interface.

Options

none	Display all the firewall filters on all captive portal interfaces.
brief | detail	(Optional) Display the specified level of output.
interface-name	(Optional) Display all the terms of the firewall filters for the specified interface.

***interface-name* detail** (Optional) Display all of the terms of the firewall filters for the specified interface.

Required Privilege Level

view

Output Fields

Output fields for the `show captive-portal firewall` command include any action modifier specified in firewall filters except policers. Policers are not supported in the terms of the internally generated dynamic firewall filters that are created when multiple supplicants authenticate on 802.1X-enabled interfaces.

Sample Output

show captive-portal firewall brief

```
user@switch> show captive-portal firewall brief
Captive Portal Information:
Interface      State           MAC address          User
ge-0/0/1.0     Connecting
ge-0/0/10.0    Connecting      00:30:48:8c:66:bd    No User
```

show captive-portal firewall (Specific Interface)

```
user@switch> show captive-portal firewall ge-0/0/10.0
Filter name: dot1x_ge-0/0/10
Counters:
Name                           Bytes              Packets
dot1x_ge-0/0/10_CP_arp          7616                  119
dot1x_ge-0/0/10_CP_dhcp            0                    0
dot1x_ge-0/0/10_CP_http            0                    0
dot1x_ge-0/0/10_CP_https           0                    0
dot1x_ge-0/0/10_CP_t_dns           0                    0
dot1x_ge-0/0/10_CP_u_dns           0                    0
```

show captive-portal firewall

```
user@switch> show captive-portal firewall
Filter name: dot1x_ge-0/0/0
Counters:
Name                                Bytes              Packets
dot1x_ge-0/0/0_CP_arp                   0                    0
dot1x_ge-0/0/0_CP_dhcp                  0                    0
dot1x_ge-0/0/0_CP_http                  0                    0
dot1x_ge-0/0/0_CP_https                 0                    0
dot1x_ge-0/0/0_CP_t_dns                 0                    0
dot1x_ge-0/0/0_CP_u_dns                 0                    0
Filter name: dot1x_ge-0/0/1
Counters:
Name                                Bytes              Packets
dot1x_ge-0/0/1_CP_arp                   0                    0
dot1x_ge-0/0/1_CP_dhcp                  0                    0
dot1x_ge-0/0/1_CP_http                  0                    0
dot1x_ge-0/0/1_CP_https                 0                    0
dot1x_ge-0/0/1_CP_t_dns                 0                    0
dot1x_ge-0/0/1_CP_u_dns                 0                    0
Filter name: dot1x_ge-0/0/10
Counters:
Name                                Bytes              Packets
dot1x_ge-0/0/10_CP_arp               7616                  119
dot1x_ge-0/0/10_CP_dhcp                 0                    0
dot1x_ge-0/0/10_CP_http                 0                    0
dot1x_ge-0/0/10_CP_https                0                    0
dot1x_ge-0/0/10_CP_t_dns                0                    0
dot1x_ge-0/0/10_CP_u_dns                0                    0
Filter name: dot1x_ge-0/0/11
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

- show captive-portal authentication-failed-users
- show captive-portal interface

clear captive-portal

Example: Setting Up Captive Portal Authentication on an EX Series Switch

Configuring Captive Portal Authentication (CLI Procedure)

show captive-portal interface

IN THIS SECTION

- Syntax | 17286
- Description | 17286
- Options | 17286
- Required Privilege Level | 17287
- Output Fields | 17287
- Sample Output | 17289
- Release Information | 17291

Syntax

```
show captive-portal interface
<interface-name>
 detail
```

Description

Display the current operational state of all captive portal interfaces with the list of connected users and the configured values of captive portal attributes on the interfaces.

Options

none	Display all captive portal interfaces.

interface-name (Optional) Display the state for the specified captive portal interface and lists the MAC address and user names of any clients authenticated on the interface.

interface-name **detail** (Optional) Display the configured values of captive portal attributes on the specified captive portal interface.

Required Privilege Level

view

Output Fields

Table 188 on page 17287 lists the output fields for the `show captive-portal interface` command. Output fields are listed in the approximate order in which they appear.

Table 188: show captive-portal interface Output Fields

Field Name	Field Description	Level of Output
`Interface`	Interface on which captive portal has been configured.	All levels
`State`	The state of the interface: • **Authenticated**—The client has been authenticated through the RADIUS server or has been permitted access through server fail fallback. • **Authenticating**—The client is authenticating through the RADIUS server. • **Connecting**—Switch is attempting to contact the RADIUS server. • **Initialize**—The interface link is down. • **Held**—An action has been triggered through server fail fallback during a RADIUS server timeout. A supplicant is denied access, permitted access through a specified VLAN, or maintains the authenticated state granted to it before the RADIUS server timeout occurred.	All levels
`MAC address`	The MAC address of the connected client on the interface..	brief

Table 188: show captive-portal interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`User`	Users connected to the captive portal interface.	brief
`Fallen back`	Indicates when 802.1X authentication and captive portal are both enabled on an interface: • If 802.1X authentication and captive portal are both enabled, **CP fallen back** status is **Yes**. • If 802.1X authentication and captive portal are not both enabled, **CP fallen back** status is **No**.	
`Supplicant mode`	Mode used to authenticate clients—multiple, single, or single-supplicant.	detail
`Number of retries`	Number of times the user can attempt to submit authentication information.	detail
`Quiet period`	Time, in seconds, after a user exceeds the maximum number of retries before they can attempt to authenticate.	detail
`Configured CP session timeout`	Time, in seconds, that a client can be idle before the session expires.	detail
`Server timeout`	Time, in seconds, that an interface will wait for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action.	detail
`Configured CP User-keepalive timeout`	Time, in minutes, that a captive portal authentication session is extended after the MAC aging timer expires.	detail

Table 188: show captive-portal interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Number of connected supplicants`	Number of users connecting through the captive portal interface. Information for each user includes: • `Supplicant`—User name and MAC address. • `Operational state`—See State (above). • `Dynamic CP session timeout`—Timeout value dynamically downloaded from the RADIUS server for this user, if any. • `CP Session expiration due in`—Time remaining in session. • `Eapol-Block`—Shows whether EAPOL block is in effect or not. • `CP Session User-keepalive Expiration due in`—Time, in seconds, remaining in the keep-alive period.	detail

Sample Output

show captive-portal interface (Only Captive Portal Enabled)

```
user@switch> show captive-portal interface
Captive Portal Information:
Interface      State           MAC address          User          Fallen back
ge-0/0/1.0     Connecting
ge-0/0/10.0    Connecting      00:30:48:8c:66:bd    No User
ge-6/0/5.0     Authenticated   00:30:48:8d:7a:9b    abcdeX      No
```

show captive-portal interface (802.1X Authentication and Captive Portal Enabled)

```
user@switch> show captive-portal interface
Captive Portal Information:
Interface      State           MAC address          User          Fallen back
ge-0/0/1.0     Connecting
ge-0/0/10.0    Connecting      00:30:48:8c:66:bd    No User
```

```
ge-6/0/5.0     Authenticated   00:30:48:8d:7a:9b    abcdeX      Yes
```

show captive-portal interface detail (Only Captive Portal Enabled)

```
user@switch> show captive-portal interface detail ge-6/0/5.0

  Supplicant mode: Multiple
  Number of retries: 3
  Quiet period: 60 seconds
  Configured CP session timeout: 3600 seconds
  Server timeout: 15 seconds
  Configured CP User-keepalive timeout: 7 minutes
  CP fallen back: No
  Number of connected supplicants: 1
      Supplicant: abcdeX, 00:30:48:8d:7a:9b
          Operational state: Authenticated
          Dynamic CP Session Timeout: 3600 seconds
          CP Session Expiration due in: 3583 seconds
                 Eapol-Block: In Effect
          CP session User-keepalive Expiration due in: 420 seconds
```

show captive-portal interface detail (802.1X Authentication and Captive Portal Enabled)

```
user@switch> show captive-portal interface detail ge-6/0/5.0

 Supplicant mode: Multiple
  Number of retries: 3
  Quiet period: 60 seconds
  Configured CP session timeout: 3600 seconds
  Server timeout: 15 seconds
  CP fallen back: Yes
  Number of connected supplicants: 1
      Supplicant: abcdeX, 00:30:48:8d:7a:9b
          Operational state: Authenticated
          Dynamic CP Session Timeout: 3600 seconds
          CP Session Expiration due in: 3583 seconds
              Eapol-Block: In Effect
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

- show captive-portal authentication-failed-users | 17281
- show captive-portal firewall | 17283
- captive-portal | 3199
- clear captive-portal | 3450
- *Example: Setting Up Captive Portal Authentication on an EX Series Switch*
- *Configuring Captive Portal Authentication (CLI Procedure)*

show chassis adc

IN THIS SECTION

- Syntax | 17291
- Description | 17292
- Options | 17292
- Required Privilege Level | 17292
- Output Fields | 17292
- Sample Output | 17293
- Release Information | 17295

Syntax

```
show chassis adc
<all-members>
<local>
<member member-id>
```

Description

Display chassis information about the adapter cards (ADCs).

Options

none Display information about all adapter cards.

all-members (Optional) Display information about the adapter cards (ADCs) in all members of the Virtual Chassis configuration.

local (Optional) Display information about the ADCs in the local member of the Virtual Chassis.

member ***member-id*** (Optional) Display information about the ADCs in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 189 on page 17292 lists the output fields for the `show chassis adc` command. Output fields are listed in the approximate order in which they appear.

Table 189: show chassis adc Output Fields

Field Name	Field Description
`Slot`	Slot number.
`State`	Status of the adapter card. • `Online`—The adapter card is online and running. • `Offline`—Adapter card is powered down.
`Uptime`	How long the Routing Engine has been connected to the adapter card and, therefore, how long the adapter card has been up and running.

Table 189: show chassis adc Output Fields *(Continued)*

Field Name	Field Description
GNF (Node slicing)	GNF identifier for each ADC.

Sample Output

show chassis adc (MX2020 Router)

```
user@host> show chassis adc
Slot  State                           Uptime
 0    Online 1 hour, 21 minutes, 7 seconds
 1    Online 1 hour, 21 minutes, 3 seconds
 2    Online 1 hour, 20 minutes, 59 seconds
 3    Online 1 hour, 20 minutes, 54 seconds
 4    Online 1 hour, 20 minutes, 50 seconds
 5    Online 1 hour, 20 minutes, 46 seconds
 6    Online 1 hour, 20 minutes, 42 seconds
 7    Online 1 hour, 20 minutes, 37 seconds
 8    Online 1 hour, 20 minutes, 33 seconds
 9    Online 1 hour, 20 minutes, 28 seconds
10    Online 1 hour, 20 minutes, 24 seconds
11    Online 1 hour, 20 minutes, 19 seconds
12    Online 1 hour, 20 minutes, 15 seconds
13    Online 1 hour, 20 minutes, 8 seconds
14    Online 1 hour, 20 minutes, 4 seconds
15    Online 1 hour, 19 minutes, 59 seconds
16    Online 1 hour, 19 minutes, 55 seconds
17    Online 1 hour, 19 minutes, 50 seconds
18    Online 1 hour, 19 minutes, 45 seconds
19    Online 1 hour, 19 minutes, 39 seconds
```

show chassis adc (MX2010 Router)

```
user@host > show chassis adc
Slot  State                           Uptime
 0    Online 12 hours, 17 minutes, 38 seconds
```

```
 1    Online 12 hours, 17 minutes, 30 seconds
 2    Online 12 hours, 17 minutes, 22 seconds
 3    Online 12 hours, 17 minutes, 14 seconds
 4    Online 12 hours, 17 minutes, 6 seconds
 5    Online 12 hours, 16 minutes, 58 seconds
 6    Online 12 hours, 16 minutes, 49 seconds
 7    Online 12 hours, 16 minutes, 41 seconds
 8    Online 12 hours, 16 minutes, 33 seconds
 9    Online 12 hours, 16 minutes, 25 seconds
```

show chassis adc (MX2008 Router)

```
user@host > show chassis adc
Slot  State                           Uptime
 0    Empty  --- Native line card ---
 1    Empty  --- Native line card ---
 2    Empty
 3    Empty
 4    Empty
 5    Empty
 6    Empty
 7    Online 1 hour, 14 minutes, 32 seconds
 8    Empty
 9    Empty
```

show chassis adc (Node Slicing)

```
user@router> show chassis adc

Slot  State           Uptime                                        GNF
 0    Online           12 hours, 57 minutes, 46 seconds               3
 1    Empty       --- Native line card ---                            2
 2    Online           12 hours, 57 minutes, 18 seconds               3
 3    Online           11 minutes, 23 seconds                         6
 4    Empty       --- Native line card ---                            6
 5    Empty       --- Native line card ---                            4
 6    Online           13 hours, 38 minutes, 58 seconds               1
 7    Online           13 hours, 3 minutes, 40 seconds                5
```

```
8    Empty       --- Native line card ---                              5
9    Empty       --- Native line card ---                              5
```

Release Information

Command introduced in Junos OS Release 12.3 for MX2020 and MX2010 Universal Routing Platforms.

`all-members`, `local`, and `member` *`member-id`* options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.

RELATED DOCUMENTATION

show chassis environment adc | 17553

show chassis adc (GNF)

IN THIS SECTION

- Syntax | 17295
- Description | 17296
- Required Privilege Level | 17296
- Output Fields | 17296
- Sample Output | 17296
- Release Information | 17297

Syntax

```
show chassis adc
```

Description

Display chassis information about the adapter cards (ADCs) assigned to the guest network function (GNF).

Required Privilege Level

view

Output Fields

Table 190 on page 17296 lists the output fields for the `show chassis adc` command. The output fields are listed in the approximate order in which they appear.

Table 190: show chassis adc Output Fields

Field Name	Field Description
`Slot`	Slot number.
`State`	Status of the adapter card. • `Online`—The adapter card is online and running. • `Offline`—Adapter card is powered down.
`Uptime`	How long the Routing Engine has been connected to the adapter card and, therefore, how long the adapter card has been up and running.
`GNF` (Node slicing)	GNF identifier for each ADC.

Sample Output

command-name

```
user@router> show chassis adc

Slot  State          Uptime                                        GNF
 4    Empty       --- Native line card ---                           2
```

```
6    Online           50 minutes, 2 seconds                          2
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

- show chassis hardware (GNF) | 18453
- show chassis fpc (GNF) | 18404
- show chassis fpc pic-status (GNF) | 18411

show chassis adc (Node Slicing)

IN THIS SECTION

- Syntax | 17297
- Description | 17297
- Required Privilege Level | 17298
- Output Fields | 17298
- Sample Output | 17298
- Release Information | 17299

Syntax

```
show chassis adc
```

Description

Display chassis information about the adapter cards (ADCs).

Required Privilege Level

view

Output Fields

Table 191 on page 17298 lists the output fields for the `show chassis adc` command. Output fields are listed in the approximate order in which they appear.

Table 191: show chassis adc Output Fields

Field Name	Field Description
`Slot`	Slot number.
`State`	Status of the adapter card. • `Online`—The adapter card is online and running. • `Offline`—Adapter card is powered down.
`Uptime`	How long the Routing Engine has been connected to the adapter card and, therefore, how long the adapter card has been up and running.
`GNF` (Node slicing)	GNF identifier for each ADC.

Sample Output

show chassis adc (Node Slicing)

```
user@router> show chassis adc

Slot  State          Uptime                                        GNF
 0    Online          12 hours, 57 minutes, 46 seconds               3
 1    Empty       --- Native line card ---                           2
 2    Online          12 hours, 57 minutes, 18 seconds               3
 3    Online          11 minutes, 23 seconds                         6
 4    Empty       --- Native line card ---                           6
 5    Empty       --- Native line card ---                           4
```

```
6    Online           13 hours, 38 minutes, 58 seconds                1
7    Online           13 hours, 3 minutes, 40 seconds                 5
8    Empty       --- Native line card ---                             5
9    Empty       --- Native line card ---                             5
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

- show chassis network-slices | 18575
- show chassis fpc pic-status (Node Slicing) | 18413
- show chassis fpc (Node Slicing) | 18408
- show chassis network-slices fpcs | 18579

show chassis afeb

IN THIS SECTION

- Syntax | 17299
- Description | 17300
- Options | 17300
- Required Privilege Level | 17300
- Output Fields | 17300
- Sample Output | 17301
- Release Information | 17301

Syntax

```
show chassis afeb
```

Description

Display compact Forwarding Engine Board status.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 192 on page 17300 lists the output fields for the `show chassis afeb` command. Output fields are listed in the approximate order in which they appear.

Table 192: show chassis afeb

Field Name	Field Description

`State`	State of the compact Forwarding Engine Board:
	• `Offline`—FEB is powered down.
	• `Online`—FEB is operational and running.
	• `Check`—FEB is in alarmed state because of the following reasons:
	  • Hardware error.
	  • PFE is unable to boot.

`Temperature`	Temperature of the air passing by the FEB, in degrees Celsius or in both degrees Celsius and degrees Fahrenheit.
`CPU Utilization`	Total percentage of CPU being used.
`Interrupt Utilization`	Of the total CPU being used by the FEB processor, the percentage being used for interrupts.

Table 192: show chassis afeb *(Continued)*

Field Name	Field Description
Heap Utilization	Percentage of heap space (dynamic memory) being used by the FEB processor. If this number exceeds 80 percent, you might experience a software problem (memory leak).
Buffer Utilization	Percentage of buffer space being used by the FEB processor for buffering internal messages.
Total CPU DRAM	Total DRAM, in megabytes, available to the FEB processor.
Start time	Time when the Routing Engine detected that the FEB was running.
Uptime	How long the Routing Engine has been connected to the FEB and, therefore, how long the compact Forwarding Engine Board has been up and running.

Sample Output

show chassis afeb (MX104 Router)

```
user@host> show chassis afeb
FEB status:
Slot 0 information:
  State                                 Online
  Temperature                        31 degrees C / 87 degrees F
  CPU utilization                     3 percent
  Interrupt utilization               0 percent
  Heap utilization                   11 percent
  Buffer utilization                 13 percent
  Total CPU DRAM                   2048 MB
  Start time:                           2013-05-27 08:50:03 IST
  Uptime:                               3 hours, 29 minutes, 34 seconds
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

request chassis afeb | 14434

show chassis alarms

IN THIS SECTION

- Syntax | 17302
- Description | 17302
- Options | 17302
- Additional Information | 17302
- Required Privilege Level | 17303
- Output Fields | 17303
- Sample Output | 17303
- Release Information | 17304

Syntax

```
show chassis alarms
```

Description

Display information about the conditions that have been configured to trigger alarms.

Options

This command has no options.

Additional Information

Chassis alarms are preset. You cannot modify them.

You cannot clear the alarms for chassis components. Instead, you must remedy the cause of the alarm. When a chassis alarm is lit, it indicates that you are running the device in a manner that we do not recommend.

On routers, you can manually silence external devices connected to the alarm relay contacts by pressing the alarm cutoff button, located on the craft interface. Silencing the device does not remove the alarm messages from the display (if present on the router) or extinguish the alarm LEDs. In addition, new alarms that occur after you silence an external device reactivate the external device.

In Junos OS Release 11.1 and later, alarms for fans also show the slot number of the fans in the CLI output.

Required Privilege Level

view

Output Fields

Table 193 on page 17303 lists the output fields for the `show chassis alarms` command. Output fields are listed in the approximate order in which they appear.

Table 193: show chassis alarms Output Fields

Field Name	Field Description
`Alarm time`	Date and time the alarm was first recorded.
`Class`	Severity class for this alarm: Minor or Major.
`Description`	Information about the alarm.

Sample Output

show chassis alarms

```
user@host> show chassis alarms
4 alarms currently active
Alarm time               Class  Description
2012-05-29 16:47:18 UTC  Major  /var partition usage crossed critical threshold
```

```
2012-05-29 16:47:18 UTC  Minor  /var partition usage crossed high threshold
2012-05-29 16:47:18 UTC  Major  /root partition usage crossed critical threshold
2012-05-29 16:47:18 UTC  Minor  /root partition usage crossed high threshold
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show system alarms | 26682

show chassis alarms

IN THIS SECTION

- Syntax | 17305
- Syntax (MX Series Routers) | 17305
- Syntax (SRX1500, SRX4100, and SRX4200) | 17305
- Syntax (SRX4600) | 17305
- Syntax (SRX5400) | 17306
- Syntax (TX Matrix Routers) | 17306
- Syntax (TX Matrix Plus Routers) | 17306
- Syntax (MX104, MX2010, MX2020, and MX2008 Universal Routing Platforms) | 17306
- Syntax (QFX Series) | 17306
- Description | 17306
- Options | 17307
- Additional Information | 17308
- Required Privilege Level | 17313
- Output Fields | 17314
- Sample Output | 17314
- show chassis alarms (ACX7024) | 17333

- Release Information | **17333**

Syntax

```
show chassis alarms
```

Syntax (MX Series Routers)

```
show chassis alarms
<all-members>
<local>
<member member-id>
```

Syntax (SRX1500, SRX4100, and SRX4200)

```
show chassis alarms
1 alarms currently active
Alarm time               Class  Description
2020-02-28 10:07:16 CST  Major  FPC0: PEM 0 Not Present
```

Syntax (SRX4600)

```
show chassis alarms
node0:
--------------------------------------------------------------------------
2 alarms currently active
Alarm time Class Description
2020-10-08 19:42:06 UTC Major FPC 0 BITS CPLD Version Mismatch
2020-10-08 19:42:06 UTC Minor PEM 1 Not Present
```

Syntax (SRX5400)

```
show chassis alarms
1 alarms currently active
Alarm time               Class  Description
2021-04-16 17:21:02 PDT  Major  Too Few AC PEMs
```

Syntax (TX Matrix Routers)

```
show chassis alarms
<lcc number | scc>
```

Syntax (TX Matrix Plus Routers)

```
show chassis alarms
<lcc number | sfc number>
```

Syntax (MX104, MX2010, MX2020, and MX2008 Universal Routing Platforms)

```
show chassis alarms
<satellite [slot-id slot-id]>
```

Syntax (QFX Series)

```
show chassis alarms
<interconnect-device name>
<node-device name>
```

Description

Display information about the conditions that have been configured to trigger alarms. In Junos, the chassis alarms are different from the system alarms (viewed by using the `show system alarms` command). The system alarms indicate a missing rescue configuration or software license, where valid. For more information, see Alarm Overview.

Options

none

Display information about the conditions that have been configured to trigger alarms.

all-members

(MX Series routers only) (Optional) Display information about alarm conditions for all the member routers of the Virtual Chassis configuration.

interconnect-device *name*

(QFabric systems only) (Optional) Display information about alarm conditions for the Interconnect device.

lcc *number*

(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(MX Series routers only) (Optional) Display information about alarm conditions for the local Virtual Chassis member.

member *member-id*

(MX Series routers only) (Optional) Display information about alarm conditions for the specified member of the Virtual Chassis configuration. Replace *member-id* variable with a value of 0 or 1.

node-device *name*

(QFabric systems only) (Optional) Display information about alarm conditions for the Node device.

satellite [slot-id *slot-id*]

(Junos Fusion only)(Optional) Display information about alarm conditions for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

scc

(TX Matrix router only) (Optional) Show information about the TX Matrix router (switch-card chassis).

sfc *number* (TX Matrix Plus router only) (Optional) Show information about the respective TX Matrix Plus router, which is the switch-fabric chassis. Replace *number* variable with `0`.

Additional Information

Chassis alarms are preset. You cannot modify them.

You cannot clear the alarms for chassis components. Instead, you must remedy the cause of the alarm. When a chassis alarm LED is lit, it indicates that you are running the router or switch in a manner that we do not recommend.

On routers, you can manually silence external devices connected to the alarm relay contacts by pressing the alarm cutoff button, located on the craft interface. Silencing the device does not remove the alarm messages from the display (if present on the router) or extinguish the alarm LEDs. In addition, new alarms that occur after you silence an external device reactivate the external device.

NOTE: MX10003 routers do not support craft interface.

In Junos OS release 11.1 and later, alarms for fans also show the slot number of the fans in the CLI output.

In Junos OS Release 11.2 and later, the command output on EX8200 switches shows the detailed location (`Plane/FPC/PFE`) for link errors in the chassis.

In Junos OS Release 10.2 and later, an alarm is shown on T Series routers for a standby SONET Clock Generator (SCG) that is offline or absent.

You may often see the following error messages, in which only the error code is shown and no other information is provided:

```
Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code: 257
Apr 12 08:04:19 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code: 559
```

To understand what CM_ALARM error codes mean, you need to first identify the structure of the CM Alarm codes. A CM_ALARM code has the following structure:

Table 194: CM_ALARM Code Structure

Bits:	Error type:

1-31	Major (1)
0	Minor (0)

According to the table above, the LSB (bit 0) identifies the **Error Type** (major alarm, if the bit is set and minor alarm if the bit is unset). The rest of the bits (1 - 31) identify the actual error code.

Take an example of the following error code, which was logged on a T1600:

```
Apr 12 08:04:10 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code: 559
```

First, you have to convert 559 to binary; that is **1000101111**. The LSB in this case is **1**, which means that this is a major alarm. After removing the LSB, you are left with **100010111**, which is equal to 279 in decimal. This is the actual error code, its meaning can be found from the following list:

Table 195: L Chip Error Code Meanings

Chip Type: L Chip	Code
CMALARM_LCHIP_LOUT_DESRD_PARITY_ERR	1
CMALARM_LCHIP_LOUT_DESRD_UNINIT_ERR	2
CMALARM_LCHIP_LOUT_DESRD_ILLEGALLINK_ERR	3
CMALARM_LCHIP_LOUT_DESRD_ILLEGALSIZE_ERR	4
CMALARM_LCHIP_LOUT_HDRF_TOERR_ERR	5
CMALARM_LCHIP_LOUT_HDRF_PARITY_ERR	6
CMALARM_LCHIP_LOUT_HDRF_UCERR_ERR	7
CMALARM_LCHIP_LOUT_NLIF_CRCDROP_ERR	8
CMALARM_LCHIP_LOUT_NLIF_CRCERR_ERR	9

CMALARM_LCHIP_UCODE_TIMEOUT_ERR	10
CMALARM_LCHIP_LIN_SRCTL_ACCT_DROP_ERR	11
CMALARM_LCHIP_LIN_SRCTL_ACCT_ADDR_SIZE_ERR	12
CMALARM_LCHIP_SRAM_PARITY_ERR	13
CMALARM_LCHIP_UCODE_OVFLW_ERR	14
CMALARM_LCHIP_LOUT_HDRF_MTU_ERR	15

Table 196: M Chip Error Code Meanings

Chip Type: M Chip	Code
CMALARM_MCHIP_ECC_UNCORRECT_ERR	128

Table 197: N Chip Error Code Meanings

Chip Type: N Chip	Code
CMALARM_NCHIP_RDDMA_JBUS_TIMEOUT_ERR	256
CMALARM_NCHIP_RDDMA_FIFO_OVFLW_ERR	257
CMALARM_NCHIP_RDDMA_FIFO_UNFLW_ERR	258
CMALARM_NCHIP_RDDMA_SIZE_ERR	259
CMALARM_NCHIP_RDDMA_JBUS_CRC_ERR	260
CMALARM_NCHIP_WRDMA_PKTR_ERR	261

CMALARM_NCHIP_WRDMA_PKT_CRC_ERR	262
CMALARM_NCHIP_WRDMA_JBUS_TIMEOUT_ERR	263
CMALARM_NCHIP_WRDMA_FIFO_OVFLW_ERR	264
CMALARM_NCHIP_WRDMA_FIFO_UNFLW_ERR	265
CMALARM_NCHIP_WRDMA_PKT_LEN_ERR	266
CMALARM_NCHIP_WRDMA_JBUS_CRC_ERR	267
CMALARM_NCHIP_PKTR_DMA_AGE_ERR	268
CMALARM_NCHIP_PKTR_ICELLSIG_ERR	269
CMALARM_NCHIP_PKTR_FTTL_ERR	270
CMALARM_NCHIP_RODR_OFFSET_OVFLW_ERR	271
CMALARM_NCHIP_PKTR_TMO_CELL_ERR	272
CMALARM_NCHIP_PKTR_TMO_OUTRANGE_ERR	273
CMALARM_NCHIP_PKTR_MD_REQUEST_Q_OVFLW_ERR	274
CMALARM_NCHIP_PKTR_DMA_BUFFER_OVFLW_ERR	275
CMALARM_NCHIP_PKTR_GRT_OVFLW_ERR	276
CMALARM_NCHIP_FRQ_ERR	277
CMALARM_NCHIP_RODR_IN_Q_OVFLW_ERR	278

CMALARM_NCHIP_DBUF_CRC_ERR	279

Table 198: R Chip Error Code Meanings

Chip Type: R Chip	Code
CMALARM_RCHIP_SRAM_PARITY_ERR	512

Table 199: R Chip Error Code Meanings

Chip Type: R Chip	Code
CMALARM_ICHIP_WO_DESRD_ID_ERR	601
CMALARM_ICHIP_WO_DESRD_DATA_ERR	602
CMALARM_ICHIP_WO_DESRD_OFLOW_ERR	603
CMALARM_ICHIP_WO_HDRF_UCERR_ERR	604
CMALARM_ICHIP_WO_HDRF_MTUERR_ERR	605
CMALARM_ICHIP_WO_HDRF_PARITY_ERR	606
CMALARM_ICHIP_WO_HDRF_TOERR_ERR	607
CMALARM_ICHIP_WO_IP_CRC_ERR	608
CMALARM_ICHIP_WO_IP_INTER_ERR	609
CMALARM_ICHIP_WI_WAN_TIMEOUT_ERR	625
CMALARM_ICHIP_WI_FAB_TIMEOUT_ERR	626

CMALARM_ICHIP_RLDRAM_BIST_ERR	630
CMALARM_ICHIP_SDRAM_BIST_ERR	631
CMALARM_ICHIP_RLDRAM_PARITY_ERR	632
CMALARM_ICHIP_SDRAM_UNCORRECT_ERR	633
CMALARM_ICHIP_SDRAM_CORRECT_ERR	634
CMALARM_ICHIP_FUSE_DONE_ERR	635

According to the table above, the **279** error code corresponds to **CMALARM_NCHIP_DBUF_CRC_ERR**; this means that new CRC errors were seen on the NCHIP of this particular FPC, which is FPC as per the logs.

If you do not want to convert decimal to binary and vice versa, you may use the following shortcut:

For major alarms, the **Actual Error Code = (Error Code - 1)/2**, where **Error Code** is the code that you get in the log message. For example, if you get the following log:

```
    Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code:
257
```

Actual Error Code = (257-1)/2 = 128. Similarly, for minor alarms, Actual Error Code = (Error Code)/2

NOTE: Starting in Junos OS Release 18.2R1, on MX Series routers, the `show chassis alarms` output does not display error codes for PFE-related errors. You can use the following commands to view more details of the errors that caused the alarms:

- `show chassis errors active`
- `show chassis errors active detail`

Required Privilege Level

view

Output Fields

Table 200 on page 17314 lists the output fields for the `show chassis alarms` command. Output fields are listed in the approximate order in which they appear.

Table 200: show chassis alarms Output Fields

Field Name	Field Description
`Alarm time`	Date and time the alarm was first recorded.
`Class`	Severity class for this alarm: `Minor` or `Major`.
`Description`	Information about the alarm. When optics are connected to the device, the output displays the transceiver name instead of interface name.

Sample Output

show chassis alarms (Alarms Active)

```
user@host> show chassis alarms
3 alarms are currently active
Alarm time               Class  Description
2000-02-07 10:12:22 UTC Major fxp0: ethernet link down
2000-02-07 10:11:54 UTC Minor YELLOW ALARM - PEM 1 Removed
2000-02-07 10:11:03 UTC Minor YELLOW ALARM - Lower Fan Tray Removed
```

show chassis alarms (No Alarms Active)

```
user@host> show chassis alarms
No alarms are currently active
```

show chassis alarms (Fan Tray)

```
user@host> show chassis alarms
4 alarms currently active
```

```
Alarm time               Class  Description
2010-11-11 20:27:38 UTC  Major  Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC  Minor  Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC  Major  Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC  Major  Side Fan Tray 0 Failure
```

show chassis alarms (MX150)

```
user@host > show chassis alarms
1 alarms currently active
Alarm time               Class  Description
2016-06-04 01:49:43 PDT  Major  Fan Tray 1 Fan 0 failed
```

show chassis alarms (MX104 Router)

```
user@host >show chassis alarms
1 alarms currently active
Alarm time               Class  Description
2013-06-05 14:43:31 IST  Minor  Backup RE Active
```

show chassis alarms (MX2010 Router)

```
user@host> show chassis alarms
7 alarms currently active
Alarm time               Class  Description
2012-08-07 00:46:06 PDT  Major  Fan Tray 2 Failure
2012-08-06 18:24:36 PDT  Minor  Redundant feed missing for PSM 6
2012-08-06 07:41:04 PDT  Minor  Redundant feed missing for PSM 8
2012-08-04 02:42:06 PDT  Minor  Redundant feed missing for PSM 5
2012-08-03 21:14:24 PDT  Minor  Loss of communication with Backup RE
2012-08-03 12:26:03 PDT  Minor  Redundant feed missing for PSM 4
2012-08-03 10:40:18 PDT  Minor  Redundant feed missing for PSM 7
```

show chassis alarms (MX2020 Router)

```
user@host> show chassis alarms
1 alarms currently active
Alarm time Class Description
2012-10-03 12:14:59 PDT Minor Plane 0 not online
```

show chassis alarms (MX10003 Router)

```
user@host> show chassis alarms

9 alarms currently active
Alarm time               Class  Description
2017-07-13 21:50:31 PDT  Major  FPC 1 Temperature Hot
2017-07-13 21:50:04 PDT  Minor  FPC 1 PIC 1 Invalid port profile configuration
2017-07-13 21:49:13 PDT  Minor  FPC 1 PIC 0 Invalid port profile configuration
2017-07-13 21:48:54 PDT  Major  FPC 0 Temperature Hot
2017-07-13 21:43:54 PDT  Minor  CB 1 Voltage Sensor ADS7830_0x4B Sensor Failed
2017-07-13 21:43:54 PDT  Minor  CB 0 Voltage Sensor ADS7830_0x4B Sensor Failed
2017-07-13 21:43:31 PDT  Minor  Loss of communication with Backup RE
```

Starting in Junos OS Release 19.2R1, the MX10003 routers do not raise an alarm if a Power Entry Module (PEM) slot is empty. However, when the number of operational PEMs goes below 2, the router raises a major alarm. This alarm is cleared when the required number of PEMs are made available.

show chassis alarms (MX10004 Router)

```
user@host> show chassis alarms
2 alarms currently active
Alarm time               Class  Description
```

```
2022-07-05 20:34:40 PDT  Major  Fan Tray 0 EEPROM Read Error
2021-17-05 11:06:47 PST  Minor  Check plane 0 Fabric Chip
```

```
user@host> show chassis alarms | match PEM
2021-11-12 13:19:35 UTC Major PEM 2 Not Supported
```

```
user@host> show system alarms | match Fan
2021-09-27 19:13:15 PDT Major Fan Tray 0 Model Not Supported
```

-->

show chassis alarms (MX204 Router)

```
user@host> show chassis alarms

1 alarms currently active
Alarm time               Class  Description
2017-11-05 22:13:03 PST  Major  PEM 0 Not Present
```

show chassis alarms (MX2008 Router)

```
user@host>show chassis alarms
No alarms currently active
```

show chassis alarms (MX960, MX480, and MX240 Routers showing Major CB Failure)

A major CB 0 failure alarm occurs in the event of a bad CB (unknown or mismatched CBs do not trigger this alarm in Junos Release OS 12.3R9 and later). Following GRES or recovery, if the hardware issue persists, the traffic moves to the good CB and continues. If the alarm was triggered by something transient like a power zone budget on GRES, bringing the CB back online can clear the alarm. Otherwise, replace the bad CB. Note that fabric link speed is not impacted by an offline SCB. The alarm migh be raised on CB0, CB1, and CB2.

```
user@host> show chassis alarms
6 alarms currently active
Alarm time               Class  Description
```

```
2014-10-31 16:49:41 EDT  Major  PEM 3 Not OK
2014-10-31 16:49:41 EDT  Major  PEM 2 Not OK
2014-10-31 16:49:31 EDT  Major  CB 0 Failure
2014-10-31 16:49:31 EDT  Minor  CB 0 Fabric Chip 0 Not Online
2014-10-31 16:49:31 EDT  Minor  CB 0 Fabric Chip 1 Not Online
2014-10-31 16:49:31 EDT  Minor  Backup RE Active
```

show chassis alarms (PTX10008 Router)

```
user@host>show chassis alarms
12 alarms currently active
Alarm time               Class  Description
2017-05-09 01:38:55 PDT  Minor  Loss of communication with Backup RE
2017-05-05 06:49:57 PDT  Major  FPC 5 LCPU Temp Sensor Access Failed
2017-05-05 06:49:57 PDT  Major  FPC 5 PE2 Temp Sensor Hot
2017-05-05 06:49:57 PDT  Major  FPC 5 PE1 Temp Sensor Hot
2017-05-05 06:49:57 PDT  Major  FPC 5 PE0 Temp Sensor Hot
2017-05-05 06:49:57 PDT  Major  FPC 5 Exhaust-C Temp Sensor Hot
2017-05-05 06:49:57 PDT  Major  FPC 5 Exhaust-B Temp Sensor Hot
2017-05-05 06:49:57 PDT  Major  FPC 5 Exhaust-A Temp Sensor Hot
2017-05-05 06:49:57 PDT  Major  FPC 5 Intake-B Temp Sensor Access Failed
2017-05-05 06:49:57 PDT  Major  FPC 5 Intake-A Temp Sensor Access Failed
2017-05-05 06:49:57 PDT  Major  Fan Tray 0 Fan 5 running at lower speed
2017-05-05 06:49:57 PDT  Major  Fan Tray 0 Fan 4 running at lower speed
```

show chassis alarms (T4000 Router)

```
user@host> show chassis alarms
9 alarms currently active
Alarm time               Class  Description
2007-06-02 01:41:10 UTC  Minor  RE 0 Not Supported
2007-06-02 01:41:10 UTC  Minor  CB 0 Not Supported
2007-06-02 01:41:10 UTC  Minor  Mixed Master and Backup RE types
2007-05-30 19:37:33 UTC  Major  SPMB 1 not online
2007-05-30 19:37:29 UTC  Minor  Front Bottom Fan Tray Absent
2007-05-30 19:37:13 UTC  Major  PEM 1 Input Failure
2007-05-30 19:37:13 UTC  Major  PEM 0 Not OK
2007-05-30 19:37:03 UTC  Major  PEM 0 Improper for Platform
2007-05-30 19:37:03 UTC  Minor  Backup RE Active
```

show chassis alarms (Unreachable Destinations Present on a T Series Router)

```
user@host> show chassis alarms
10 alarms currently active
Alarm time               Class  Description
2011-08-30 18:43:53 PDT  Major  FPC 7 has unreachable destinations
2011-08-30 18:43:53 PDT  Major  FPC 5 has unreachable destinations
2011-08-30 18:43:52 PDT  Major  FPC 3 has unreachable destinations
2011-08-30 18:43:52 PDT  Major  FPC 2 has unreachable destinations
2011-08-30 18:43:52 PDT  Minor  SIB 0 Not Online
2011-08-30 18:43:33 PDT  Minor  SIB 4 Not Online
2011-08-30 18:43:28 PDT  Minor  SIB 3 Not Online
2011-08-30 18:43:05 PDT  Minor  SIB 2 Not Online
2011-08-30 18:43:28 PDT  Minor  SIB 1 Not Online
2011-08-30 18:43:05 PDT  Major  PEM 1 Not Ok
```

show chassis alarms (FPC Offline Due to Unreachable Destinations on a T Series Router)

```
user@host> show chassis alarms
10 alarms currently active
Alarm time               Class  Description
2011-08-30 18:43:53 PDT  Major  FPC 7 offline due to unreachable destinations
2011-08-30 18:43:53 PDT  Major  FPC 5 offline due to unreachable destinations
2011-08-30 18:43:52 PDT  Major  FPC 3 offline due to unreachable destinations
2011-08-30 18:43:52 PDT  Major  FPC 2 offline due to unreachable destinations
2011-08-30 18:43:52 PDT  Minor  SIB 0 Not Online
2011-08-30 18:43:33 PDT  Minor  SIB 4 Not Online
2011-08-30 18:43:28 PDT  Minor  SIB 3 Not Online
2011-08-30 18:43:05 PDT  Minor  SIB 2 Not Online
2011-08-30 18:43:28 PDT  Minor  SIB 1 Not Online
2011-08-30 18:43:05 PDT  Major  PEM 1 Not Ok
```

show chassis alarms (SCG Absent on a T Series Router)

```
user@host> show chassis alarms
4 alarms currently active
Alarm time               Class  Description
2011-01-23 21:42:46 PST  Major  SCG 0 NO EXT CLK MEAS-BKUP SCG ABS
```

show chassis alarms (Alarms Active on a TX Matrix Router)

```
user@host> show chassis alarms
scc-re0:
--------------------------------------------------------------------------
8 alarms currently active
Alarm time               Class  Description
2004-08-05 18:43:53 PDT  Minor  LCC 0 Minor Errors
2004-08-05 18:43:53 PDT  Minor  SIB 3 Not Online
2004-08-05 18:43:52 PDT  Major  SIB 2 Absent
2004-08-05 18:43:52 PDT  Major  SIB 1 Absent
2004-08-05 18:43:52 PDT  Major  SIB 0 Absent
2004-08-05 18:43:33 PDT  Major  LCC 2 Major Errors
2004-08-05 18:43:28 PDT  Major  LCC 0 Major Errors
2004-08-05 18:43:05 PDT  Minor  LCC 2 Minor Errors
lcc0-re0:
--------------------------------------------------------------------------
5 alarms currently active
Alarm time               Class  Description
2004-08-05 18:43:53 PDT  Minor  SIB 3 Not Online
2004-08-05 18:43:49 PDT  Major  SIB 2 Absent
2004-08-05 18:43:49 PDT  Major  SIB 1 Absent
2004-08-05 18:43:49 PDT  Major  SIB 0 Absent
2004-08-05 18:43:28 PDT  Major  PEM 0 Not OK
lcc2-re0:
--------------------------------------------------------------------------
5 alarms currently active
Alarm time               Class  Description
2004-08-05 18:43:35 PDT  Minor  SIB 3 Not Online
2004-08-05 18:43:33 PDT  Major  SIB 2 Absent
2004-08-05 18:43:33 PDT  Major  SIB 1 Absent
2004-08-05 18:43:33 PDT  Major  SIB 0 Absent
2004-08-05 18:43:05 PDT  Minor  PEM 1 Absent
```

show chassis alarms (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis alarms
sfc0-re0:
--------------------------------------------------------------------------
Alarm time               Class  Description
```

```
2014-04-08 14:35:13 IST  Minor  FPM 0 SFC Config Size Changed
2014-04-08 14:32:58 IST  Major  Fan Tray Failure
2014-04-08 14:31:53 IST  Major  SIB F13 6 Fault
2014-04-08 14:31:43 IST  Major  SIB F13 11 Fault
2014-04-08 14:31:08 IST  Minor  Check SIB F13 12 CXP 14 Fbr Cbl
2014-04-08 14:31:08 IST  Minor  Check SIB F13 12 CXP 8 Fbr Cbl
2014-04-08 14:31:08 IST  Minor  Check SIB F13 12 CXP 3 Fbr Cbl
2014-04-08 14:31:08 IST  Major  SIB F13 12 CXP 15 fault
2014-04-08 14:31:08 IST  Minor  SIB F13 12 CXP 14 LOL
2014-04-08 14:31:08 IST  Minor  Check SIB F13 12 CXP 14
2014-04-08 14:31:08 IST  Major  SIB F13 12 CXP 10 fault
2014-04-08 14:31:08 IST  Minor  SIB F13 12 CXP 8 LOL
2014-04-08 14:31:08 IST  Minor  Check SIB F13 12 CXP 8
2014-04-08 14:31:08 IST  Major  SIB F13 12 CXP 7 fault
2014-04-08 14:31:08 IST  Major  SIB F13 12 CXP 4 fault
2014-04-08 14:31:08 IST  Minor  SIB F13 12 CXP 3 LOL
2014-04-08 14:31:08 IST  Minor  Check SIB F13 12 CXP 3
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 14 Fbr Cbl
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 12 Fbr Cbl
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 8 Fbr Cbl
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 6 Fbr Cbl
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 4 Fbr Cbl
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 2 Fbr Cbl
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 0 Fbr Cbl
2014-04-08 14:31:08 IST  Minor  SIB F13 6 CXP 14 LOL
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 14
2014-04-08 14:31:08 IST  Minor  SIB F13 6 CXP 12 LOL
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 12
2014-04-08 14:31:08 IST  Major  SIB F13 6 CXP 10 fault
2014-04-08 14:31:08 IST  Minor  SIB F13 6 CXP 8 LOL
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 8
2014-04-08 14:31:08 IST  Minor  SIB F13 6 CXP 6 LOL
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 6
2014-04-08 14:31:08 IST  Minor  SIB F13 6 CXP 4 LOL
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 4
2014-04-08 14:31:08 IST  Minor  SIB F13 6 CXP 2 LOL
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 2
2014-04-08 14:31:08 IST  Minor  SIB F13 6 CXP 0 LOL
2014-04-08 14:31:08 IST  Minor  Check SIB F13 6 CXP 0
2014-04-08 14:31:08 IST  Minor  SIB F13 12 CXP 14 XC HSL Link Error
2014-04-08 14:29:27 IST  Minor  LCC 0 Minor Errors
2014-04-08 14:28:37 IST  Major  LCC 0 Major Errors
```

```
2014-04-08 14:28:37 IST  Major  LCC 2 Major Errors
2014-04-08 14:28:37 IST  Minor  LCC 2 Minor Errors
2014-04-08 14:28:24 IST  Major  SIB F2S 4/6 Absent
2014-04-08 14:28:24 IST  Major  SIB F2S 4/4 Absent
2014-04-08 14:28:24 IST  Major  SIB F2S 4/2 Absent
2014-04-08 14:28:24 IST  Major  SIB F2S 4/0 Absent
2014-04-08 14:28:24 IST  Major  SIB F2S 3/6 Absent
2014-04-08 14:28:24 IST  Major  SIB F2S 3/4 Absent
2014-04-08 14:28:24 IST  Major  SIB F2S 3/2 Absent
2014-04-08 14:28:24 IST  Major  SIB F2S 3/0 Absent
2014-04-08 14:28:24 IST  Major  SIB F13 9 Absent
2014-04-08 14:28:24 IST  Major  SIB F13 8 Absent
2014-04-08 14:28:24 IST  Major  SIB F13 7 Absent
2014-04-08 14:28:24 IST  Major  SIB F13 4 Absent
2014-04-08 14:28:24 IST  Major  SIB F13 1 Absent
2014-04-08 14:28:22 IST  Major  PEM 0 Input Failure
2014-04-08 14:28:22 IST  Major  PEM 0 Not OK

lcc0-re0:
--------------------------------------------------------------------------
12 alarms currently active
Alarm time               Class  Description
2014-04-08 14:36:08 IST  Minor  CB 1  M/S Switch Changed
2014-04-08 14:36:08 IST  Minor  CB 1  CHASSIS ID Changed
2014-04-08 14:35:43 IST  Minor  CB 0  M/S Switch Changed
2014-04-08 14:35:43 IST  Minor  CB 0  CHASSIS ID Changed
2014-04-08 14:29:30 IST  Minor  SIB 4 Not Online
2014-04-08 14:29:30 IST  Minor  SIB 3 Not Online
2014-04-08 14:29:30 IST  Minor  SIB 2 Not Online
2014-04-08 14:29:24 IST  Major  Rear Fan Tray Failure
2014-04-08 14:29:24 IST  Major  Front Bottom Fan Tray Improper for Platform
2014-04-08 14:29:24 IST  Major  Front Top Fan Tray Improper for Platform
2014-04-08 14:28:37 IST  Major  SIB 4 Absent
2014-04-08 14:28:37 IST  Major  SIB 3 Absent

lcc2-re0:
--------------------------------------------------------------------------
12 alarms currently active
Alarm time               Class  Description
2014-04-08 14:36:02 IST  Minor  CB 1  M/S Switch Changed
2014-04-08 14:36:02 IST  Minor  CB 1  CHASSIS ID Changed
2014-04-08 14:35:42 IST  Minor  CB 0  M/S Switch Changed
2014-04-08 14:34:42 IST  Minor  CB 0  CHASSIS ID Changed
```

```
2014-04-08 14:29:29 IST  Minor  SIB 0 CXP 7 Unsupported Optics
2014-04-08 14:29:27 IST  Major  Front Bottom Fan Tray Improper for Platform
2014-04-08 14:29:27 IST  Major  Front Top Fan Tray Improper for Platform
2014-04-08 14:29:25 IST  Minor  SIB 4 Not Online
2014-04-08 14:29:25 IST  Minor  SIB 3 Not Online
2014-04-08 14:28:47 IST  Major  PEM 0 Not OK
2014-04-08 14:28:36 IST  Major  SIB 2 Absent
2014-04-08 14:28:36 IST  Minor  Host 0 Boot from alternate media

lcc6-re0:
--------------------------------------------------------------------------
2 alarms currently active
Alarm time               Class  Description
2013-11-06 04:03:56 PST  Minor  SIB 1 CXP 0 XC HSL Link Error
2013-11-06 03:49:32 PST  Major  PEM 1 Not OK
```

show chassis alarms (Alarms on a T4000 Router After the enhanced-mode Statement is Enabled)

To enable improved virtual private LAN service (VPLS) MAC address learning on T4000 routers, you must include the `enhanced-mode` statement at the `[edit chassis network-services]` hierarchy level and reboot the router. When router reboots, only the T4000 Type 5 FPCs are required to be present on the router. If there are any other FPCs (apart from T4000 Type 5 FPCs) on the T4000 router, such FPCs become offline, and FPC misconfiguration alarms are generated. The `show chassis alarm` command output displays FPC misconfiguration (`FPC` *`fpc-slot`* `misconfig`) as the reason for the generation of the alarms.

```
user@host> show chassis alarms
2 alarms currently active
        Alarm time               Class  Description
        2011-10-22 10:10:47 PDT  Major  FPC 1 misconfig
        2011-10-22 10:10:46 PDT  Major  FPC 0 misconfig
```

show chassis alarms (Backup Routing Engine)

```
user@host> show chassis alarms
2 alarms are currently active
Alarm time               Class  Description
2005-04-07 10:12:22 PDT  Minor  Host 1 Boot from alternate media
2005-04-07 10:11:54 PDT  Major  Host 1 compact-flash missing in Boot List
```

show chassis alarms (EX Series Switch)

```
user@switch> show chassis alarms

 4 alarms currently active
Alarm time               Class  Description
2014-03-12 15:36:09 UTC  Minor  Require a Fan Tray upgrade
2014-03-12 15:00:02 UTC  Major  PEM 0 Input Failure
2014-03-12 15:00:02 UTC  Major  PEM 0 Not OK
2014-03-12 14:59:51 UTC  Minor  Host 1 Boot from alternate media
```

show chassis alarms (Alarms Active on the QFX Series and OCX Series Switches)

```
user@switch> show chassis alarms
1 alarms currently active
Alarm time               Class  Description
2012-03-05    2:10:24 UTC  Major  FPC 0 PEM 0 Airflow not matching Chassis Airflow
```

show chassis alarms node-device (Alarms Active on the QFabric System)

```
user@switch> show chassis alarms node-device
Test
node-device ED3694
3 alarms currently active
Alarm time               Class  Description
2011-08-24 16:04:15 UTC  Major  Test:fte-0/1/2: Link down
2011-08-24 16:04:14 UTC  Major  Test:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC  Major  Test PEM 0 is not supported/powered
```

show chassis alarms (Alarms Active on the QFabric System)

```
user@switch> show chassis alarms
IC-1:
--------------------------------------------------------------------------
1 alarms currently active
Alarm time               Class  Description
2011-08-24 16:04:15 UTC  Minor  Backup RE Active
```

```
Test:
---------------------------------------------------------------------------
3 alarms currently active
Alarm time               Class  Description
2011-08-24 16:04:15 UTC  Major  Test:fte-0/1/2: Link down
2011-08-24 16:04:14 UTC  Major  Test:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC  Major  Test PEM 0 is not supported/powered

SNG-0:
---------------------------------------------------------------------------

NW-NG-0:
---------------------------------------------------------------------------
1 alarms currently active
Alarm time               Class  Description
2011-08-24 15:49:27 UTC  Major  Test PEM 0 is not supported/powered
```

show chassis alarms (Alarms Active on an EX8200 Switch)

```
user@switch> show chassis alarms

6 alarms currently active
Alarm time               Class  Description
2010-12-02 19:15:22 UTC  Major  Fan Tray Failure
2010-12-02 19:15:22 UTC  Major  Fan Tray Failure
2010-12-02 19:15:14 UTC  Minor  Check CB 0 Fabric Chip 1 on Plane/FPC/PFE: 1/5/0, 1/5/1, 1/5/2,
1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:15:14 UTC  Minor  Check CB 0 Fabric Chip 0 on Plane/FPC/PFE: 1/5/0, 1/5/1, 1/5/2,
1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:14:18 UTC  Major  PSU 1 Output Failure
2010-12-02 19:14:18 UTC  Minor  Loss of communication with Backup RE
```

show chassis alarms (EX9251 Switch)

```
user@switch> show chassis alarms

 2 alarms currently active
Alarm time               Class  Description
```

```
2018-03-08 05:13:10 PST  Major  PEM 0 Not Powered
2018-03-08 05:13:10 PST  Major  Fan Tray 2 is not present
```

show chassis alarms (EX9253 Switch)

```
user@switch> show chassis alarms

6 alarms currently active
Alarm time               Class  Description
2018-03-07 01:09:01 PST  Major  Power Budget:Insufficient Power
2018-03-06 23:56:34 PST  Minor  Loss of communication with Backup RE
2018-02-15 00:48:10 PST  Minor  PEM 3 Not Present
2018-02-15 00:48:10 PST  Minor  PEM 2 Not Present
2018-02-15 00:48:07 PST  Major  PEM 4 Not Powered
2018-02-15 00:48:07 PST  Major  PEM 1 Not Powered
```

show chassis alarms (MX Series for PFE reset)

```
user@host> show chassis alarms
1 alarm currently active
Alarm time               Class  Description
2021-01-04 11:20:42 PST  Major  FPC 0 PFE 0  reset initiated
```

show chassis alarms (MX Series for PFE reset error)

```
user@host> show chassis alarms
1 alarm currently active
Alarm time               Class  Description
2021-01-04 11:20:42 PST  Major  FPC 0 PFE 0  reset error
```

show chassis alarms (MX2010, MX10008, MX10016 and MX10004)

An optics configuration mismatch alarm may be triggered when there is a discrepancy between the configured speed of an interface and the supported speed of the optic. This alarm indicates that the

optic installed in the specified FPC is incompatible with the speed configured on the interface. This can be observed through the following output:

```
root@user> run show chassis alarms
2 alarms currently active
Alarm time               Class  Description
2024-10-14 21:54:38 PDT  Minor  FPC 5 Optics does not support configured speed
2024-10-14 17:07:23 PDT  Minor  FPC 0 Optics does not support configured speed
```

show chassis alarms (Alarms Active on a PTX5000 Packet Transport Router)

```
user@host> show chassis alarms

23 alarms currently active
Alarm time               Class  Description
2011-07-12 16:22:05 PDT  Minor  No Redundant Power for Rear Chassis
2011-07-12 16:22:05 PDT  Major  PDU 0 PSM 1 Not OK
2011-07-12 16:21:57 PDT  Minor  No Redundant Power for Fan 0-2
2011-07-12 16:21:57 PDT  Major  PDU 0 PSM 0 Not OK
2011-07-12 15:56:06 PDT  Major  PDU 1 PSM 2 Not OK
2011-07-12 15:56:06 PDT  Minor  No Redundant Power for FPC 0-7
2011-07-12 15:56:06 PDT  Major  PDU 0 PSM 3 Not OK
2011-07-12 15:28:20 PDT  Major  PDU 0 PSM 2 Not OK
2011-07-12 15:19:14 PDT  Minor  Backup RE Active
```

show chassis alarms (Mix of PDUs Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A)

All PDUs installed on a PTX5000 router must be of the same type. The `Mix of PDUs` or `Power Manager Non Operational` alarm is raised when different types of PDUs are installed on a PTX5000 router.

```
user@host> show chassis alarms
15 alarms currently active
Alarm time               Class  Description
2013-03-19 23:03:53 PDT  Minor  No Redundant Power
2013-03-19 23:03:48 PDT  Minor  Mix of PDUs
2013-03-19 23:03:47 PDT  Minor  PDU 1 PSM 3 Absent
2013-03-19 23:03:47 PDT  Minor  PDU 1 PSM 2 Absent
2013-03-19 23:03:47 PDT  Minor  PDU 1 PSM 1 Absent
```

```
2013-03-19 23:03:47 PDT  Minor  PDU 1 PSM 0 Absent
2013-03-19 23:03:46 PDT  Major  No CG Online
```

show chassis alarms (PDU Converter Failed Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A)

The `PDU Converter Failed` alarm is raised when one or more 36 V booster converter of a DC PDU fails. If two or more 36 V booster converter fails, fan trays fail and the router might get over heated. Therefore, when this alarm is raised, check the PDU and replace it, if required.

```
user@host> show chassis alarms
11 alarms currently active
Alarm time               Class  Description
2013-12-11 22:14:13 PST  Minor  No Redundant Power for System
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 7 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 6 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 5 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 4 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 3 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 2 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 1 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 PSM 0 Not OK
2013-12-11 22:14:10 PST  Major  PDU 0 Not OK
2013-12-11 22:14:01 PST  Major  PDU 0 Converter Failed
```

show chassis alarms (No Power for System Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis alarms
8 alarms currently active
Alarm time               Class  Description
2013-11-19 01:58:41 PST  Major  No Power for System
2013-11-19 01:58:37 PST  Major  PDU 0 PSM 1 Not OK
2013-11-19 01:56:46 PST  Major  PDU 0 PSM 2 Not OK
2013-11-19 01:54:26 PST  Major  PDU 0 PSM 3 Not OK
2013-11-19 01:53:30 PST  Major  PDU 1 PSM 3 Not OK
2013-11-19 01:53:29 PST  Major  PDU 1 PSM 2 Not OK
2013-11-19 01:53:29 PST  Major  PDU 1 PSM 1 Not OK
2013-11-19 01:53:29 PST  Major  PDU 1 PSM 0 Not OK
```

show chassis alarms (Alarms Active on an ACX2000 Universal Metro Router)

```
user@host> show chassis alarms
7 alarms currently active
Alarm time               Class  Description
2012-05-22 11:19:09 UTC  Major  xe-0/3/1: Link down
2012-05-22 11:19:09 UTC  Major  xe-0/3/0: Link down
2012-05-22 11:19:09 UTC  Major  ge-0/1/7: Link down
2012-05-22 11:19:09 UTC  Major  ge-0/1/6: Link down
2012-05-22 11:19:09 UTC  Major  ge-0/1/3: Link down
2012-05-22 11:19:09 UTC  Major  ge-0/1/2: Link down
2012-05-22 11:19:09 UTC  Major  ge-0/1/1: Link down
```

show chassis alarms (Active Alarm to Indicate Status of the Bad SCB Clock on MX Series)

```
user@host> show chassis alarms
1 alarm currently active
Alarm time               Class  Description
2013-08-06 07:48:35 PDT  Major  CB 0 19.44 MHz clock failure
```

show chassis alarms (Alarms active on a PTX1000 Packet Transport Router)

```
user@host> show chassis alarms
2 alarms currently active
Alarm time               Class  Description
2004-08-10 00:55:49 UTC  Major  PEM 1 Not Present
2004-08-10 00:55:49 UTC  Major  PEM 0 Not Present
```

show chassis alarms (MX10003 Router)

If LCMD is down on the backup RE, then the following alarm is seen on the primary.

```
user@host> show chassis alarms
1 alarm currently active
Alarm time               Class  Description
2017-05-09 13:26:27 PDT  Major  VMHost RE 1 host application failed
```

If LCMD is down on the primary, then following alarms are displayed.

```
user@host> show chassis alarms
3 alarms currently active
Alarm time               Class  Description
2017-05-10 14:12:21 PDT  Major  VMHost RE 0 host application failed
2017-05-10 14:12:16 PDT  Minor  LCM Peer Absent
2017-05-09 13:26:27 PDT  Major  VMHost RE 1 host application failed
```

If the LCMD process is crashing on the primary, the system will switchover after one minute provided the backup RE LCMD connection is stable. The system will not switchover under the following conditions: if the backup RE LCMD connection is unstable or if the current primary just gained primary role. When the primary has just gained primary role, the switchover happens only after four minutes.

The LCM peer connection un-stable alarm is raised when the LCMD-CHASD IPC communication flaps three times within a small interval of two to three minutes. Once LCM peer connection un-stable alarm is raised, the connection status is monitored for two minutes.

```
user@host> show chassis alarms
7 alarms currently active
Alarm time               Class  Description
2017-05-29 10:12:17 PDT  Minor  LCM Peer Connection un-stable
2017-05-29 09:04:17 PDT  Minor  PEM 8 Not Powered
2017-05-29 09:04:17 PDT  Minor  PEM 9 Not Powered
2017-05-29 09:04:17 PDT  Minor  PEM 7 Not Powered
2017-05-29 09:04:17 PDT  Minor  PEM 3 Not Powered
2017-05-29 09:04:17 PDT  Minor  PEM 0 Not Powered
2017-05-29 09:04:08 PDT  Minor  Loss of communication with Backup RE
```

If there are no more connection flaps within this two minutes time interval, the LCM peer connection un-stable alarm is cleared.

```
6 alarms currently active
Alarm time               Class  Description
2017-05-29 09:04:17 PDT  Minor  PEM 8 Not Powered
2017-05-29 09:04:17 PDT  Minor  PEM 9 Not Powered
2017-05-29 09:04:17 PDT  Minor  PEM 7 Not Powered
2017-05-29 09:04:17 PDT  Minor  PEM 3 Not Powered
2017-05-29 09:04:17 PDT  Minor  PEM 0 Not Powered
2017-05-29 09:04:08 PDT  Minor  Loss of communication with Backup RE
```

A major alarm is raised even if there is on one PLL lock error, and this alarm can be cleared only through an FPC restart.

```
user@host> show chassis alarms
4 alarms currently active
Alarm time               Class  Description
2017-02-16 09:06:06 PDT  Major  FPC 0 Major Errors
2017-02-16 09:08:40 PDT  Major  FPC 1 Major Errors
2017-02-16 09:11:47 PST  Minor  Fan Tray 3 Pair 1 Outer Fan running at over speed
2017-02-16 09:11:47 PST  Minor  Fan Tray 3 Pair 1 Inner Fan running at over speed
```

show chassis alarms (Alarms active on a MX10008 Router)

```
user@host> show chassis alarms
13 alarms currently active
Alarm time               Class  Description
2018-07-17 05:48:08 PDT  Major  FPC 2 I2C Failure
2018-07-17 05:47:02 PDT  Minor  Mixed Master and Backup RE types
2018-07-17 05:47:01 PDT  Major  Fan Tray 0 Fan 5 Failed
2018-07-17 05:47:01 PDT  Major  Fan Tray 0 Fan 4 Failed
2018-07-17 05:47:01 PDT  Minor  PEM 5 Not Powered
2018-07-17 05:47:01 PDT  Minor  PEM 5 Feed 2 has no input source
2018-07-17 05:47:01 PDT  Minor  PEM 5 Feed 1 has no input source
2018-07-17 05:47:01 PDT  Minor  PEM 4 Not Powered
2018-07-17 05:47:01 PDT  Minor  PEM 4 Feed 2 has no input source
2018-07-17 05:47:01 PDT  Minor  PEM 4 Feed 1 has no input source
2018-07-17 05:47:01 PDT  Minor  PEM 3 Not Powered
2018-07-17 05:47:01 PDT  Minor  PEM 3 Feed 2 has no input source
2018-07-17 05:47:01 PDT  Minor  PEM 3 Feed 1 has no input source
2018-07-17 05:47:01 PDT  Major  CB 0 external-1 LOS
```

show chassis alarms (ACX Series)

```
user@host> show chassis alarms
Alarm time               Class  Description
2011-01-23 21:42:46 PST  Major  PTP Local Clock OOS
2011-01-23 21:42:46 PST  Major  PTP No Foreign Master
2011-01-23 21:42:46 PST  Major  Chassis Loss of all Equipment Clock Synch References
2011-01-23 21:42:46 PST  Major  Chassis Loss of Equipment Clock Synch Reference 1
```

```
2011-01-23 21:42:46 PST  Major  Chassis Loss of Equipment Clock Synch Reference 2
2011-01-23 21:42:46 PST  Major  Equipment Clock QL Below Threshold
2011-01-23 21:42:46 PST  Major  TOD Input A Signal Fail
2011-01-23 21:42:46 PST  Major  1PPS lost
2011-01-23 21:42:46 PST  Major  SyncE Port incompatible Media Type
```

show chassis alarms (MX10008, MX10016, PTX10008, PTX10016, QFX10008,QFX10016) (Junos OS Release)

Starting Junos OS Evolved Release 21.2R1, if PEM or FET Failure detected, a major alarm is raised, and the identified PSM will shutdown or raise alarms as per predefined configuration in the `set chassis thermal-events fet-failure-check` command.

For FET failure detection and action, the `show chassis alarm` output displays a major alarm and information if PSM is shutdown.

For example, for the `show chassis` configuration:

```
user@root> show chassis thermal-health-check
                   {
                      fet-failure-check;
                      action-onfail auto-shutdown;
                   }
```

The output is displayed as follows:

```
user@root> show chassis alarms

Alarm time                Class  Description
2007-04-10 02:33:11 PDT  Minor  No Redundant Power for System
2007-04-10 02:33:08 PDT  Minor  PDU 1 PSM 4 Thermal Check brought it Down
2007-04-10 02:33:08 PDT  Major  PDU 1 PSM 4 Not OK
2007-04-10 02:33:08 PDT  Minor  PDU 0 PSM 4 Thermal Check brought it Down
2007-04-10 02:33:08 PDT  Major  PDU 0 PSM 4 Not OK
2007-04-10 02:32:03 PDT  Minor  PDU 1 PSM 3 Thermal Check brought it Down
2007-04-10 02:32:03 PDT  Major  PDU 1 PSM 3 Not OK
...
...
```

show chassis alarms (ACX7024)

On ACX7024 device, when `edit chassis pem minimum` is set to 1, The following output is displayed. The PSM absent alarm is removed only after a reboot:

```
user@root> show chassis alarms
4 alarms currently active
Alarm time                Class  Description
2022-05-15 09:12:31   PDT Major  No Redundant Power for System
2022-05-15 09:12:31   PDT Major  PSM 1 Input1 Failed
2022-05-15 09:12:31   PDT Major  PSM 1 Input Failed
2022-05-15 09:12:34   PDT Major PSM 0 Absent
2022-06-04 05:30:11   PDT Minor RE 0 Secure boot disabled or not enforced
```

When the default PSU 1+1 redundancy mode is active, the following output is displayed:

```
user@root> show chassis alarms
6 alarms currently active
Alarm time                Class  Description
2022-05-15 09:12:31   PDT Major  No Redundant Power for System
2022-05-15 09:12:31   PDT Major  PSM 1 Input1 Failed
2022-05-15 09:12:31   PDT Major  PSM 1 Input Failed
2022-06-04 05:30:19   PDT Major  PSM 0 Absent
2022-06-04 05:30:11   PDT Minor  RE 0 Secure boot disabled or not enforced
2022-06-04 05:31:54   PDT Major  chassis No Redundant Power
```

show chassis alarms (MX304 with Optics)

```
user@host> show chassis alarms
4 alarms currently active
Alarm time               Class  Description
2024-12-13 05:51:00 UTC  Minor  FPC 0 xcvr-0/0/0 Temp Sensor Warm
2024-12-12 22:37:43 UTC  Minor  FPC 0 Optics does not support configured speed
2024-12-11 16:07:10 UTC  Major  SFB 0 I2CS CPLD image outdated
2024-12-11 16:04:52 UTC  Major  PEM 0 is Absent
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced in Junos OS Release 9.6 for the TX Matrix Plus router.

`satellite` option introduced in Junos OS Release 14.2R3 for Junos Fusion.

Command introduced in Junos OS Release 18.2R1 for EX9253 Switches and MX10008 Universal Routing Platforms.

RELATED DOCUMENTATION

- Configuring an RMON Alarm Entry and Its Attributes
- Chassis Conditions That Trigger Alarms

show chassis beacon

IN THIS SECTION

- show chassis beacon (QFX Series) | 17334
- Description | 17335
- Options | 17335
- Required Privilege Level | 17335
- Output Fields | 17335
- Sample Output | 17336
- Release Information | 17337

show chassis beacon (QFX Series)

```
show chassis beacon
<cb slot-number>
<fpc slot-number>
<interconnect-device name (cb slot-number | fpc slot-number)>
<node-device name>
```

Description

Display the beacon LED status on a QFX3500, QFX3600, QFX5100, EX4600, OCX Series standalone switch, Node device, and an Interconnect device. You can also display the beacon LED status of the Control Boards and Flexible PIC Concentrators on the Interconnect device.

Options

cb ***slot-number*** (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Control Board on the Interconnect device.

fpc ***slot-number*** (QFabric systems only) (Optional) Display the status of the beacon LEDs for the FPC on the Interconnect device. (Optional) Display the status of the beacon LEDs for the FPC on the standalone switch.

interconnect-device ***name*** (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Interconnect device.

node-device ***name*** (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Node device.

Required Privilege Level

view

Output Fields

Table 201 on page 17335 lists the output fields for the `show chassis beacon` command. Output fields are listed in the approximate order in which they appear.

Table 201: show chassis led Output Fields

Field Name	Field Description
Slot	FPC slot number of the device whose content is being displayed. On QFX3500 standalone switches, the number is always **0**.

Table 201: show chassis led Output Fields *(Continued)*

Field Name	Field Description
Beacon State	Status of the beacon state: • Off—The beacon is **OFF**. • On—The beacon is **ON**.

Sample Output

show chassis beacon (OCX Series and QFX Series Switches)

```
user@switch> show chassis beacon
Slot            Beacon State
            FPC     0           OFF
```

show chassis beacon interconnect-device (QFabric System)

```
user@switch> show chassis beacon interconnect-device
interconnect1
Chassis              OFF
CB 0                 OFF
CB 1                 OFF
FC 0 FPC 0           OFF
FC 1 FPC 1           OFF
RC 0 FPC 8           OFF
RC 1 FPC 9           OFF
```

show chassis beacon interconnect-device fpc (QFabric System)

```
user@switch> show chassis beacon interconnect-device
interconnect1 fpc 0
FPC 0                ON
```

show chassis beacon node-device (QFabric System)

```
user@switch> show chassis beacon node-device
node1
node1                ON
```

show chassis beacon node-device fpc (QFabric System)

```
user@switch> show chassis beacon node-device
node1 fpc 0
FPC 0               ON
```

show chassis beacon (ACX7100 and ACX7024)

```
user@host> show chassis beacon
              ON
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

request chassis beacon | 14435

show chassis cfeb

IN THIS SECTION

- Syntax | 17338
- Description | 17338
- Options | 17338

- Required Privilege Level | 17338
- Output Fields | 17338
- Sample Output | 17340
- Release Information | 17341

Syntax

```
show chassis cfeb
```

Description

(M7i and M10i routers only) Display status information about the Compact Forwarding Engine Board (CFEB).

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 202 on page 17339 lists the output fields for the `show chassis cfeb` command. Output fields are listed in the approximate order in which they appear.

Table 202: show chassis cfeb Output Fields

Field Name	Field Description
State	Status of the CFEB: • **Online**—CFEB is online and running. • **Offline**—CFEB is powered down.
Intake Temperature	Temperature of the air before flowing past the CFEB.
Exhaust Temperature	Temperature of the air after flowing past the CFEB.
CPU utilization	Percentage of CPU being used by the CFEB processor.
Interrupt utilization	Of the total CPU being used by the CFEB processor, the percentage being used for interrupts
Heap Utilization	Percentage of heap space (dynamic memory) being used by the CFEB processor. If this number exceeds 80 percent, there may be a software problem (memory leak).
Buffer Utilization	Percentage of buffer space being used by the CFEB processor for buffering internal messages
Total CPU DRAM	Amount of DRAM available to the CFEB CPU.
Internet Processor II	Information about the CFEB processor.
Start time	Time when the Routing Engine detected that the CFEB was running.
Uptime	How long the Routing Engine has been connected to the CFEB and, therefore, how long the Flexible PIC Concentrator (FPC) has been up and running.

Sample Output

show chassis cfeb (M7i)

```
user@host> show chassis cfeb
CFEB status:
  State                          Online
  Intake Temperature          27 degrees C / 80 degrees F
  Exhaust Temperature         33 degrees C / 91 degrees F
  CPU utilization              3 percent
  Interrupt utilization        0 percent
  Heap utilization             8 percent
  Buffer utilization          21 percent
  Total CPU DRAM             128 MB
  Internet Processor II          Version 1, Foundry IBM, Part number 164
  Start time:                    2003-06-11 11:41:22 PDT
  Uptime:                        1 hour, 39 minutes, 31 seconds
```

show chassis cfeb (M10i)

```
user@host> show chassis cfeb
CFEB status:
Slot 0 information:
  StateMaster
  Intake temperature                 35 degrees C / 95 degrees F
  Exhaust temperature                43 degrees C / 109 degrees F
  CPU utilization                     3 percent
  Interrupt utilization               0 percent
  Heap utilization                   10 percent
  Buffer utilization                 22 percent
  Total CPU DRAM                    128 MB
  Internet Processor II                Version 1, Foundry IBM, Part number 164
  Start time:                          2004-11-01 03:24:15 PST
  Uptime:                             12 hours, 56 minutes, 18 seconds
Slot 1 information:
  State                                Backup
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

request chassis cfeb | 14449

Configuring CFEB Redundancy on the M10i Router

CFEB Overview

show chassis cluster control-interface status

IN THIS SECTION

- Syntax | 17341
- Description | 17341
- Required Privilege Level | 17342
- Output Fields | 17342
- Sample Output | 17348
- Release Information | 17350

Syntax

```
show chassis cluster control-interface status
```

Description

Display information about 1G optical fiber HA control port SFP basic, diagnostic data, HA control port line-side and system-side link status information in both standalone and cluster mode.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show chassis cluster control-interface status` command.

Table 203: show chassis cluster control-interface status Output Fields

Field Name	Field Parameters	Description
Link Status:	Interface—Specifies the interface number.	
	Line-side—Status for line side.	
	System-side—Status for system side.	
Optical Fiber SFP Status:	Port—Specifies the port number.	
	Cable type—Specifies the cable type.	
	Fiber type—Specifies the fiber type model.	
	Xcvr vendor—Specifies the vendor name.	
	Xcvr vendor part number—Specifies the vendor part number.	
	Wave-length—Specifies the wave length number.	
	Xcvr Firmware—Specifies the firmware version.	
	JNPR Rev—Specifies the revision number.	
	MSA Version—Specifies the SFF version number.	

Table 203: show chassis cluster control-interface status Output Fields *(Continued)*

Field Name	Field Parameters	Description
Control Ports Optics Diagnostics Information:	Control port 0:	Laser bias current—Specifies the power information. Laser output power—Specifies the power information. Module temperature—Specifies the temperature information. Module voltage—Specifies the voltage information. Laser receiver power—Specifies the power information. Alarm/warning flags implemented—Specifies the status of the alarm. Laser bias current high alarm / low alarm / high warning / low warning—Specifies the status for the laser bias current alarm and warning.

Table 203: show chassis cluster control-interface status Output Fields *(Continued)*

Field Name	Field Parameters	Description
	Control port 1:	Laser output power high alarm / low alarm / high warning / low warning—Specifies the status for the laser output power alarm and warning. Module temperature high alarm / low alarm / high warning / low warning—Specifies the status for the module temperature alarm and warning. Module voltage high alarm / low alarm / high warning / low warning—Specifies the status for the module voltage alarm and warning. Laser rx power high alarm / low alarm / high warning / low warning—Specifies the status for the

Table 203: show chassis cluster control-interface status Output Fields *(Continued)*

Field Name	Field Parameters	Description
		laser rx power alarm and warning. Laser bias current high alarm threshold—Specifies the laser bias current threshold value. Laser bias current low alarm threshold—Specifies the laser bias current threshold value. Laser bias current high warning threshold—Specifies the laser bias current threshold value. Laser bias current low warning threshold—Specifies the laser bias current threshold value. Laser output power high alarm threshold—Specifies the output power threshold value. Laser output power low alarm threshold—Specifies the

Table 203: show chassis cluster control-interface status Output Fields *(Continued)*

Field Name	Field Parameters	Description
		output power threshold value. Laser output power high warning threshold—Specifies the output power threshold value. Laser output power low warning threshold—Specifies the output power threshold value. Module temperature high alarm threshold—Specifies the temperature threshold value. Module temperature low alarm threshold—Specifies the temperature threshold value. Module temperature high warning threshold—Specifies the temperature threshold value. Module temperature low warning threshold—Specifies the temperature threshold value.

Table 203: show chassis cluster control-interface status Output Fields *(Continued)*

Field Name	Field Parameters	Description
		Module voltage high alarm threshold—Specifies the voltage threshold value. Module voltage low alarm threshold—Specifies the voltage threshold value. Module voltage high warning threshold—Specifies the voltage threshold value. Module voltage low warning threshold—Specifies the voltage threshold value. Laser rx power high alarm threshold—Specifies the laser rx power threshold value. Laser rx power low alarm threshold—Specifies the laser rx power threshold value. Laser rx power high warning

Table 203: show chassis cluster control-interface status Output Fields *(Continued)*

Field Name	Field Parameters	Description
		threshold—Specifies the laser rx power threshold value. Laser rx power low warning threshold—Specifies the laser rx power threshold value.

Sample Output

show chassis cluster control-interface status

```
user@host> show chassis cluster control-interface status
```

```
Link Status:
Interface        Line-side        System-side
0                Up               Up
1                Up               Down

Optical Fiber SFP Status:
                        Fiber                    Xcvr vendor       Wave-
Xcvr          JNPR     MSA
Port   Cable type        type  Xcvr vendor        part number       length
Firmware      Rev      Version
0      1000BASE-SX       MM    FINISAR CORP.      FTLF8519P3BNL-J1  850nm
0.0           REV 02   SFF-8472 ver 9.3
1      1000BASE-SX       MM    AVAGO              AFBR-5715PZ-JU1   850nm
0.0           REV 02   SFF-8472 ver 9.3

Control Ports Optics Diagnostics Information:
Control port 0:
Laser bias current                                                    :  6.336 mA
```

```
Laser output power                                                  :  0.3225 mW / -4.91 dBm
Module temperature                                                  :  20.85 degrees C /
69.53 degrees F
Module voltage                                                      :  3.3004 V
Laser receiver power                                                :  0.3342 mW / -4.76 dBm
Alarm/warning flags implemented                                     :  On
Laser bias current high alarm / low alarm / high warning / low warning  :  Off / Off / Off / Off
Laser output power high alarm / low alarm / high warning / low warning  :  Off / Off / Off / Off
Module temperature high alarm / low alarm / high warning / low warning  :  Off / Off / Off / Off
Module voltage high alarm / low alarm / high warning / low warning      :  Off / Off / Off / Off
Laser rx power high alarm / low alarm / high warning / low warning      :  Off / Off / Off / Off
Laser bias current high alarm threshold                             :  13.000 mA
Laser bias current low alarm threshold                              :  1.000 mA
Laser bias current high warning threshold                           :  12.500 mA
Laser bias current low warning threshold                            :  2.000 mA
Laser output power high alarm threshold                             :  1.0000 mW / 0.00 dBm
Laser output power low alarm threshold                              :  0.0447 mW / -13.50 dBm
Laser output power high warning threshold                           :  0.5012 mW / -3.00 dBm
Laser output power low warning threshold                            :  0.1122 mW / -9.50 dBm
Module temperature high alarm threshold                             :  110.00 degrees C /
230.00 degrees F
Module temperature low alarm threshold                              :  -40.00 degrees C /
-40.00 degrees F
Module temperature high warning threshold                           :  93.00 degrees C /
199.40 degrees F
Module temperature low warning threshold                            :  -30.00 degrees C /
-22.00 degrees F
Module voltage high alarm threshold                                 :  3.6000 V
Module voltage low alarm threshold                                  :  3.0000 V
Module voltage high warning threshold                               :  3.5000 V
Module voltage low warning threshold                                :  3.1000 V
Laser rx power high alarm threshold                                 :  1.1220 mW / 0.50 dBm
Laser rx power low alarm threshold                                  :  0.0079 mW / -21.02 dBm
Laser rx power high warning threshold                               :  0.7943 mW / -1.00 dBm
Laser rx power low warning threshold                                :  0.0200 mW / -16.99 dBm
Control port 1:
Laser bias current                                                  :  6.446 mA
Laser output power                                                  :  0.2983 mW / -5.25 dBm
Module temperature                                                  :  25.81 degrees C /
78.46 degrees F
Module voltage                                                      :  3.2643 V
Laser receiver power                                                :  0.3009 mW / -5.22 dBm
Alarm/warning flags implemented                                     :  On
```

```
Laser bias current high alarm / low alarm / high warning / low warning  :  Off / Off / Off / Off
Laser output power high alarm / low alarm / high warning / low warning  :  Off / Off / Off / Off
Module temperature high alarm / low alarm / high warning / low warning  :  Off / Off / Off / Off
Module voltage high alarm / low alarm / high warning / low warning      :  Off / Off / Off / Off
Laser rx power high alarm / low alarm / high warning / low warning      :  Off / Off / Off / Off
Laser bias current high alarm threshold                                 :  12.000 mA
Laser bias current low alarm threshold                                  :  2.000 mA
Laser bias current high warning threshold                               :  10.000 mA
Laser bias current low warning threshold                                :  2.000 mA
Laser output power high alarm threshold                                 :  1.1000 mW / 0.41 dBm
Laser output power low alarm threshold                                  :  0.0600 mW / -12.22 dBm
Laser output power high warning threshold                               :  1.0000 mW / 0.00 dBm
Laser output power low warning threshold                                :  0.0850 mW / -10.71 dBm
Module temperature high alarm threshold                                 :  100.00 degrees C /
212.00 degrees F
Module temperature low alarm threshold                                  :  -40.00 degrees C /
-40.00 degrees F
Module temperature high warning threshold                               :  85.00 degrees C /
185.00 degrees F
Module temperature low warning threshold                                :  -10.00 degrees C /
14.00 degrees F
Module voltage high alarm threshold                                     :  3.6300 V
Module voltage low alarm threshold                                      :  2.9700 V
Module voltage high warning threshold                                   :  3.4650 V
Module voltage low warning threshold                                    :  3.1349 V
Laser rx power high alarm threshold                                     :  1.8000 mW / 2.55 dBm
Laser rx power low alarm threshold                                      :  0.0000 mW / -inf dBm
Laser rx power high warning threshold                                   :  1.0000 mW / 0.00 dBm
Laser rx power low warning threshold                                    :  0.0200 mW / -16.99 dBm
```

Release Information

Command introduced in Junos OS Release 23.4R1.

show chassis cluster control-plane statistics

IN THIS SECTION

- Syntax | 17351
- Description | 17351
- Required Privilege Level | 17351
- Output Fields | 17351
- Sample Output | 17352
- Sample Output | 17353
- Sample Output | 17353
- Release Information | 17354

Syntax

```
show chassis cluster control-plane statistics
```

Description

Display information about chassis cluster control plane statistics.

Required Privilege Level

view

Output Fields

Table 204 on page 17352 lists the output fields for the `show chassis cluster control-plane statistics` command. Output fields are listed in the approximate order in which they appear.

Table 204: show chassis cluster control-plane statistics Output Fields

Field Name	Field Description
`Control link statistics`	Statistics of the control link used by chassis cluster traffic. Statistics for `Control link 1` are displayed when you use dual control links (SRX5600 and SRX5800 devices only). • `Heartbeat packets sent`—Number of heartbeat messages sent on the control link. • `Heartbeat packets received`—Number of heartbeat messages received on the control link. • `Heartbeat packet errors`—Number of heartbeat packets received with errors on the control link.
`Fabric link statistics`	Statistics of the fabric link used by chassis cluster traffic. Statistics for `Child Link 1` are displayed when you use dual fabric links. • `Probes sent`—Number of probes sent on the fabric link. • `Probes received`—Number of probes received on the fabric link.
`Switch fabric link statistics`	Statistics of the switch fabric link used by chassis cluster traffic. • `Probe state`—State of the probe, `UP` or `DOWN`. • `Probes sent`—Number of probes sent. • `Probes received`—Number of probes received. • `Probe recv error` —Error in receiving probe. • `Probe send error`—Error in sending probe.

Sample Output

show chassis cluster control-plane statistics

```
user@host> show chassis cluster control-plane statistics
Control link statistics:
```

```
    Control link 0:
        Heartbeat packets sent: 11646
        Heartbeat packets received: 8343
        Heartbeat packet errors: 0
Fabric link statistics:
    Child link 0
        Probes sent: 11644
        Probes received: 8266
   Switch fabric link statistics:
    Probe state : DOWN
    Probes sent: 8145
    Probes received: 8013
    Probe recv errors: 0
    Probe send errors: 0
```

Sample Output

Sample Output

show chassis cluster control-plane statistics (SRX5000 Line Devices)

```
user@host> show chassis cluster control-plane statistics
Control link statistics:
    Control link 0:
        Heartbeat packets sent: 2061982
        Heartbeat packets received: 2060367
        Heartbeat packet errors: 0
      Control link 1:
        Heartbeat packets sent: 2061982
        Heartbeat packets received: 0
        Heartbeat packet errors: 0
Fabric link statistics:
    Child link 0
        Probes sent: 3844342
        Probes received: 3843841
    Child link 1
        Probes sent: 0
        Probes received: 0
```

Release Information

Command introduced in Junos OS Release 9.3. Output changed to support dual control ports in Junos OS Release 10.0.

RELATED DOCUMENTATION

clear chassis cluster control-plane statistics | 3455

show chassis cluster data-plane interfaces

IN THIS SECTION

- Syntax | 17354
- Description | 17354
- Required Privilege Level | 17354
- Output Fields | 17355
- Sample Output | 17355
- Release Information | 17355

Syntax

```
show chassis cluster data-plane interfaces
```

Description

Display the status of the data plane interface (also known as a fabric interface) in a chassis cluster configuration.

Required Privilege Level

view

Output Fields

Table 205 on page 17355 lists the output fields for the `show chassis cluster data-plane interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 205: show chassis cluster data-plane interfaces Output Fields

Field Name	Field Description
`fab0/fab1`	Name of the logical fabric interface. • `Name`—Name of the physical Ethernet interface. • `Status`—State of the fabric interface: `up` or `down`.

Sample Output

show chassis cluster data-plane interfaces

```
user@host> show chassis cluster data-plane interfaces
fab0:
    Name                Status
    ge-2/1/9            up
    ge-2/2/5            up
fab1:
    Name                Status
    ge-8/1/9            up
    ge-8/2/5            up
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

show chassis cluster data-plane statistics

IN THIS SECTION

- Syntax | 17356
- Description | 17356
- Required Privilege Level | 17356
- Output Fields | 17356
- Sample Output | 17358
- Release Information | 17358

Syntax

```
show chassis cluster data-plane statistics
```

Description

Display information about chassis cluster data plane statistics.

Required Privilege Level

view

Output Fields

Table 206 on page 17357 lists the output fields for the `show chassis cluster data-plane statistics` command. Output fields are listed in the approximate order in which they appear.

Table 206: show chassis cluster data-plane statistics Output Fields

Field Name	Field Description
`Services Synchronized`	• `Service name`—Name of the service. • `Rtos sent`—Number of runtime objects (RTOs) sent. • `Rtos received`—Number of RTOs received. • `Translation context`—Messages synchronizing Network Address Translation (NAT) translation context. • `Incoming NAT`—Messages synchronizing incoming Network Address Translation (NAT) service. • `Resource manager`—Messages synchronizing resource manager groups and resources. • `Session create`—Messages synchronizing session creation. • `Session close`—Messages synchronizing session close. • `Session change`—Messages synchronizing session change. • `Gate create`—Messages synchronizing creation of pinholes (temporary openings in the firewall). • `Session ageout refresh request`—Messages synchronizing request session after age-out. • `Session ageout refresh reply`—Messages synchronizing reply session after age-out. • `IPsec VPN`—Messages synchronizing VPN session. • `Firewall user authentication`—Messages synchronizing firewall user authentication session. • `MGCP ALG`—Messages synchronizing MGCP ALG sessions. • `H323 ALG`—Messages synchronizing H.323 ALG sessions. • `SIP ALG`—Messages synchronizing SIP ALG sessions. • `SCCP ALG`—Messages synchronizing SCCP ALG sessions.

Table 206: show chassis cluster data-plane statistics Output Fields *(Continued)*

Field Name	Field Description
	• PPTP ALG—Messages synchronizing PPTP ALG sessions. • RTSP ALG—Messages synchronizing RTSP ALG sessions.

Sample Output

show chassis cluster data-plane statistics

```
user@host> show chassis cluster data-plane statistics
Services Synchronized:
    Service name                            RTOs sent    RTOs received
    Translation context                     0            0
    Incoming NAT                            0            0
    Resource manager                        0            0
    Session create                          0            0
    Session close                           0            0
    Session change                          0            0
    Gate create                             0            0
    Session ageout refresh requests         0            0
    Session ageout refresh replies          0            0
    IPsec VPN                               0            0
    Firewall user authentication            0            0
    MGCP ALG                                0            0
    H323 ALG                                0            0
    SIP ALG                                 0            0
    SCCP ALG                                0            0
    PPTP ALG                                0            0
    RTSP ALG                                0            0
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

clear chassis cluster data-plane statistics | 3457

show chassis cluster ethernet-switching interfaces

IN THIS SECTION

- Syntax | 17359
- Description | 17359
- Required Privilege Level | 17359
- Output Fields | 17359
- Sample Output | 17360
- Release Information | 17360

Syntax

```
show chassis cluster ethernet-switching interfaces
```

Description

Display the status of the switch fabric interfaces (swfab interfaces) in a chassis cluster.

Required Privilege Level

view

Output Fields

Table 207 on page 17360 lists the output fields for the `show chassis cluster ethernet-switching interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 207: show chassis cluster ethernet-switching interfaces Output Fields

Field Name	Field Description
swfab *switch fabric interface-name*	Name of the switch fabric interface. • Name—Name of the physical interface. • Status—State of the switch fabric interface: up or down.

Sample Output

show chassis cluster ethernet-switching interfaces

```
user@host> show chassis cluster ethernet-switching interfaces
swfab0:
    Name                Status
    ge-0/0/9            up
    ge-0/0/10           up
swfab1:
    Name                Status
    ge-7/0/9            up
    ge-7/0/10           up
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

Ethernet Switching User Guide

show chassis cluster ethernet-switching status

IN THIS SECTION

- Syntax | 17361
- Description | 17361
- Required Privilege Level | 17361
- Output Fields | 17361
- Sample Output | 17363
- Release Information | 17363

Syntax

```
show chassis cluster ethernet-switching status
```

Description

Display the Ethernet switching status of the chassis cluster.

Required Privilege Level

view

Output Fields

Table 208 on page 17362 lists the output fields for the `show chassis cluster ethernet-switching status` command. Output fields are listed in the approximate order in which they appear.

Table 208: show chassis cluster ethernet-switching status Output Fields

Field Name	Field Description
`Cluster ID`	ID number (1-255) of a cluster. Setting a cluster ID to 0 is equivalent to disabling a cluster. A cluster ID greater than 15 can only be set when the fabric and control link interfaces are connected back-to-back.
`Redundancy-Group`	You can create up to 128 redundancy groups in the chassis cluster.
`Node name`	Node (device) in the chassis cluster (`node0` or `node1`).
`Priority`	Assigned priority for the redundancy group on that node.
`Status`	State of the redundancy group (`Primary`, `Secondary`, `Lost`, or `Unavailable`). • `Primary`—Redundancy group is active and passing traffic. • `Secondary`—Redundancy group is passive and not passing traffic. • `Lost`—Node loses contact with the other node through the control link. Most likely to occur when both nodes are in a cluster and there is a control link failure, one node cannot exchange heartbeats, or when the other node is rebooted. • `Unavailable`—Node has not received a single heartbeat over the control link from the other node since the other node booted up. Most likely to occur when one node boots up before the other node, or if only one node is present in the cluster.
`Preempt`	• `Yes`: Primary Role can be preempted based on priority. • `No`: Change in priority will not preempt primary role.
`Manual failover`	• `Yes`: Primary Role is set manually through the CLI. • `No`: Primary Role is not set manually through the CLI.

Sample Output

show chassis cluster ethernet-switching status

```
user@host> show chassis cluster ethernet-switching status

Monitor Failure codes:
       CS  Cold Sync monitoring        FL  Fabric Connection monitoring
    GR  GRES monitoring             HW  Hardware monitoring
    IF  Interface monitoring        IP  IP monitoring
    LB  Loopback monitoring         MB  Mbuf monitoring
    NH  Nexthop monitoring          NP  NPC monitoring
    SP  SPU monitoring              SM  Schedule monitoring
    CF  Config Sync monitoring

Cluster ID: 1
Node   Priority Status         Preempt Manual   Monitor-failures

Redundancy group: 0 , Failover count: 0
node0  1      primary        no      no       None
node1  1      secondary      no      no       None

Ethernet switching status:
    Probe state is UP. Both nodes are in single ethernet switching domain(s).
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

Ethernet Switching User Guide

show chassis cluster information

IN THIS SECTION

- Syntax | 17364
- Description | 17364
- Required Privilege Level | 17364
- Output Fields | 17364
- Sample Output | 17365
- Sample Output | 17375
- Sample Output | 17376
- Release Information | 17378

Syntax

```
show chassis cluster information
```

Description

Display chassis cluster messages. The messages indicate each node's health condition and details of the monitored failure.

Required Privilege Level

view

Output Fields

Table 209 on page 17365 lists the output fields for the `show chassis cluster information` command. Output fields are listed in the approximate order in which they appear.

Table 209: show chassis cluster information Output Fields

Field Name	Field Description
`Node`	Node (device) in the chassis cluster (`node0` or `node1`).
`Redundancy mode`	Displays configured data plane mode for the cluster. • active-active—Primary and secondary nodes in active mode. • active-backup—Primary in active, secondary in backup mode.
`Redundancy Group Information`	• Redundancy Group—ID number (0 - 255) of a redundancy group in the cluster. • Current State—State of the redundancy group: `primary`, `secondary`, `hold`, or `secondary-hold`. • Weight—Relative importance of the redundancy group. • Time—Time when the redundancy group changed the state. • From—State of the redundancy group before the change. • To—State of the redundancy group after the change. • Reason—Reason for the change of state of the redundancy group.
`Chassis cluster LED information`	• Current LED color—Current color state of the LED. • Last LED change reason—Reason for change of state of the LED.

Sample Output

show chassis cluster information detail

```
user@host> show chassis cluster information detail
node0:
--------------------------------------------------------------------------
Redundancy mode:
    Configured mode: active-active
```

```
    Operational mode: active-active
Cluster configuration:
    Heartbeat interval: 1000 ms
    Heartbeat threshold: 3
    Control link recovery: Disabled
    Fabric link down timeout: 66 sec
Node health information:
    Local node health: Healthy
    Remote node health: Not healthy

Redundancy group: 0, Threshold: 255, Monitoring failures: none
    Events:
        Oct 16 21:44:50.783 : hold->secondary, reason: Hold timer expired
        Oct 16 21:45:41.587 : secondary->primary, reason: Better priority (200/1)

Redundancy group: 1, Threshold: 255, Monitoring failures: none
    Events:
        Oct 16 21:44:50.784 : hold->secondary, reason: Hold timer expired
        Oct 16 21:45:41.599 : secondary->primary, reason: Remote yield (0/0)
Control link statistics:
    Control link 0:
        Heartbeat packets sent: 35917
        Heartbeat packets received: 35372
        Heartbeat packet errors: 0
        Duplicate heartbeat packets received: 0
    Control link 1:
        Heartbeat packets sent: 35917
        Heartbeat packets received: 0
        Heartbeat packet errors: 0
        Duplicate heartbeat packets received: 0
    Control recovery packet count: 3
    Sequence number of last heartbeat packet sent: 35917
    Sequence number of last heartbeat packet received: 32962
Fabric link statistics:
    Child link 0
        Probes sent: 68081
        Probes received: 3385
    Child link 1
        Probes sent: 0
        Probes received: 0

Chassis cluster LED information:
    Current LED color: Green
```

```
    Last LED change reason: No failures

Cold Synchronization:
    Status:
        Cold synchronization completed for: FPC 0 PIC 0, FPC 0 PIC 1
        Cold synchronization failed for: N/A
        Cold synchronization not known for: N/A
        Current Monitoring Weight: 0

    Progress:
        CS Prereq               2 of 2 SPUs completed
           1. if_state sync          2 SPUs completed
           2. fabric link            2 SPUs completed
           3. policy data sync       2 SPUs completed
           4. cp ready               2 SPUs completed
           5. VPN data sync          2 SPUs completed
           6. IPID data sync         2 SPUs completed
           7. All SPU ready          2 SPUs completed
           8. AppID ready            2 SPUs completed
           9. Tunnel Sess ready      2 SPUs completed
        CS RTO sync             2 of 2 SPUs completed
        CS Postreq              2 of 2 SPUs completed

    Statistics:
        Number of cold synchronization completed: 2
        Number of cold synchronization failed: 0

            PFE             Cold-Sync complete           Cold-Sync failed
            ----------------------------------------------------------------
         FPC 0 PIC 0             1                              0
         FPC 0 PIC 1             1                              1

    Events:
        Oct 16 21:53:24.430 : Cold sync for PFE FPC 0 PIC 0 is Pre-req check in process
        Oct 16 21:53:35.272 : Cold sync for PFE FPC 0 PIC 1 is Pre-req check in process
        Oct 16 21:55:46.983 : Cold sync for PFE FPC 0 PIC 0 is RTO sync in process
        Oct 16 21:56:06.077 : Cold sync for PFE FPC 0 PIC 0 is Completed
        Oct 16 22:01:27.414 : Cold sync for PFE FPC 0 PIC 1 is Not complete
        Oct 16 22:03:38.712 : Cold sync for PFE FPC 0 PIC 1 is RTO sync in process
        Oct 16 22:05:48.502 : Cold sync for PFE FPC 0 PIC 1 is Completed

Loopback Information:
```

```
    PIC Name        Loopback        Nexthop     Mbuf
    ------------------------------------------------
    FPC 0 PIC 0     Success         Success     Success
    FPC 0 PIC 1     Success         Success     Success

SPU monitoring:
    Status: Active
        Current monitoring weight: 0

        PIC Name              State
        -----------------------------
            FPC 0 PIC 0      SPC3 cp-flow combo PIC Up
            FPC 0 PIC 1      SPC3 flow PIC Up

    Statistics:
        SPU up count: 2
        NPC up count: 0
        SPU down count: 0
        NPC down count: 0
        Chassis info processing error count: 0

    Events:
        Oct 16 21:50:33.960 : ioc4 PIC Up in slot FPC 3 PIC 0
        Oct 16 21:50:46.888 : ioc4 PIC Up in slot FPC 3 PIC 1
        Oct 16 21:51:00.356 : ioc3 PIC Up in slot FPC 2 PIC 0
        Oct 16 21:51:00.427 : ioc3 PIC Up in slot FPC 2 PIC 1
        Oct 16 21:53:24.430 : SPC3 cp-flow combo PIC  in slot FPC 0 PIC 0
        Oct 16 21:53:35.272 : SPC3 flow PIC Up in slot FPC 0 PIC 1

Interface monitoring:
    Statistics:
        Monitored interface failure count: 0

Fabric monitoring:
    Status:
        Fabric Monitoring: Enabled
        Activation status: Suspended by local node and other node
        Fabric Status reported by data plane: Down
        JSRPD internal fabric status: Up

Fabric link events:
        Oct 16 22:58:05.026 : Child xe-15/0/7 of fab1 is down
        Oct 16 22:58:05.027 : Child xe-15/0/7 of fab1 is up
```

```
        Oct 16 22:58:18.817 : Child xe-15/0/7 of fab1 is down
        Oct 16 22:58:25.266 : Fabric link fab0 is up
        Oct 16 22:58:25.267 : Child xe-3/0/7 of fab0 is up
        Oct 16 22:58:25.470 : Fabric link fab1 is up
        Oct 16 22:58:25.471 : Child xe-15/0/7 of fab1 is up
        Oct 16 22:59:12.120 : Fabric link down, link status timer
        Oct 17 00:35:22.682 : Child xe-3/0/7 of fab0 is up
        Oct 17 00:35:24.348 : Child xe-15/0/7 of fab1 is up

Control link status: Up
    Server information:
        Server status : Connected
        Server connected to 130.32.0.1/48579
    Client information:
        Client status : Inactive
        Client connected to None

Hardware monitoring:
    Status:
        Activation status: Enabled
        Redundancy group 0 failover for hardware faults: Enabled
        Hardware redundancy group 0 errors: 0
        Hardware redundancy group 1 errors: 0

    Hardware monitoring events:
        Oct 17 00:35:06.509 : hw-mon errors notified  RG-0  errors:0 RG-1  errors:0

Schedule monitoring:
    Status:
        Activation status: Disabled
        Schedule slip detected: None
        Timer ignored: No

    Statistics:
        Total slip detected count: 0
        Longest slip duration: 0(s)

Configuration Synchronization:
    Status:
        Activation status: Enabled
        Last sync operation: Auto-Sync
        Last sync result: Not needed
        Last sync mgd messages:
```

```
    Events:
        Oct 16 21:45:41.599 : Auto-Sync: Not needed.

Cold Synchronization Progress:
    CS Prereq               2 of 2 SPUs completed
       1. if_state sync          2 SPUs completed
       2. fabric link            2 SPUs completed
       3. policy data sync       2 SPUs completed
       4. cp ready               2 SPUs completed
       5. VPN data sync          2 SPUs completed
       6. IPID data sync         2 SPUs completed
       7. All SPU ready          2 SPUs completed
       8. AppID ready            2 SPUs completed
       9. Tunnel Sess ready      2 SPUs completed
    CS RTO sync             2 of 2 SPUs completed
    CS Postreq              2 of 2 SPUs completed

 ISSU events:
        System currently in ISSU Window and ISSU state is 80
        Oct 16 21:50:21.708 : node0: Exiting  ISSU window, issu_state: 32
        Oct 16 21:50:21.709 : node0: Exited ISSU, issu_state: 0
        Oct 16 22:08:29.645 : node0: Entered ISSU window, issu_state: 32
        Oct 16 22:08:29.645 : node0: Initiating ISSU, issu_state: 48

 Command history:
        Oct 16 21:50:21.709 : Invoked ISSU abort on node0
        Oct 16 22:08:29.645 : Invoked ISSU on node0

node1:
-------------------------------------------------------------------------
Redundancy mode:
    Configured mode: active-active
    Operational mode: active-active
Cluster configuration:
    Heartbeat interval: 1000 ms
    Heartbeat threshold: 3
    Control link recovery: Disabled
    Fabric link down timeout: 66 sec
Node health information:
    Local node health: Not healthy
    Remote node health: Healthy
```

```
Redundancy group: 0, Threshold: 255, Monitoring failures: none
    Events:
        Oct 16 22:34:12.282 : hold->secondary, reason: Hold timer expired

Redundancy group: 1, Threshold: -765, Monitoring failures: spu-monitoring, cold-sync-monitoring,
hardware-monitoring
    Events:
        Oct 16 22:34:12.282 : hold->secondary, reason: Hold timer expired
Control link statistics:
    Control link 0:
        Heartbeat packets sent: 32962
        Heartbeat packets received: 32987
        Heartbeat packet errors: 0
        Duplicate heartbeat packets received: 0
    Control link 1:
        Heartbeat packets sent: 32962
        Heartbeat packets received: 0
        Heartbeat packet errors: 0
        Duplicate heartbeat packets received: 0
    Control recovery packet count: 3
    Sequence number of last heartbeat packet sent: 32962
    Sequence number of last heartbeat packet received: 35917
Fabric link statistics:
    Child link 0
        Probes sent: 0
        Probes received: 0
    Child link 1
        Probes sent: 0
        Probes received: 0

Chassis cluster LED information:
    Current LED color: Amber
    Last LED change reason: Monitored objects are down

Cold Synchronization:
    Status:
        Cold synchronization completed for: N/A
        Cold synchronization failed for: N/A
        Cold synchronization not known for: N/A
        Current Monitoring Weight: 255

    Progress:
        CS Prereq               0 of 2 SPUs completed
```

```
            1. if_state sync          0 SPUs completed
            2. fabric link            0 SPUs completed
            3. policy data sync       0 SPUs completed
            4. cp ready               0 SPUs completed
            5. VPN data sync          0 SPUs completed
            6. IPID data sync         0 SPUs completed
            7. All SPU ready          0 SPUs completed
            8. AppID ready            0 SPUs completed
            9. Tunnel Sess ready      0 SPUs completed
        CS RTO sync             0 of 2 SPUs completed
        CS Postreq              0 of 2 SPUs completed

    Statistics:
        Number of cold synchronization completed: 0
        Number of cold synchronization failed: 0

SPU monitoring:
    Status: Active
        Current monitoring weight: 510



    Statistics:
        SPU up count: 0
        NPC up count: 0
        SPU down count: 2
        NPC down count: 0
        Chassis info processing error count: 0

    Events:
        Oct 16 22:41:06.095 : ioc4 PIC Up in slot FPC 3 PIC 0
        Oct 16 22:41:18.374 : ioc4 PIC Up in slot FPC 3 PIC 1
        Oct 16 22:53:50.425 : PIC empty in slot FPC 3 PIC 0
        Oct 16 22:53:50.426 : PIC empty in slot FPC 3 PIC 1
        Oct 16 22:54:05.600 : PIC empty in slot FPC 2 PIC 0
        Oct 16 22:54:05.600 : PIC empty in slot FPC 2 PIC 1
        Oct 16 22:58:04.973 : ioc4 PIC Up in slot FPC 3 PIC 0
        Oct 16 22:58:17.586 : ioc4 PIC Up in slot FPC 3 PIC 1
        Oct 16 23:01:18.754 : ioc3 PIC Up in slot FPC 2 PIC 0
        Oct 16 23:01:21.937 : ioc3 PIC Up in slot FPC 2 PIC 1

Interface monitoring:
    Statistics:
        Monitored interface failure count: 0
```

```
Fabric monitoring:
    Status:
        Fabric Monitoring: Enabled
        Activation status: Suspended by local node and other node
        Fabric Status reported by data plane: Down
        JSRPD internal fabric status: Down

Fabric link events:
        Oct 16 22:58:25.270 : Child xe-3/0/7 of fab0 is up
        Oct 16 22:58:25.474 : Fabric link fab1 is up
        Oct 16 22:58:25.474 : Fabric link fab1 is up
        Oct 16 22:58:25.475 : Child xe-15/0/7 of fab1 is up
        Oct 16 23:10:46.812 : Fabric monitoring is suspended due to HWMON failure
        Oct 16 23:10:46.813 : Fabric monitoring is suspended due to HWMON failure
        Oct 17 00:35:22.685 : Child xe-3/0/7 of fab0 is up
        Oct 17 00:35:24.352 : Child xe-15/0/7 of fab1 is up
        Oct 17 01:15:06.691 : Fabric monitoring is suspended due to HWMON failure
        Oct 17 01:15:06.692 : Fabric monitoring is suspended due to HWMON failure

Control link status: Up
    Server information:
        Server status : Inactive
        Server connected to None
    Client information:
        Client status : Connected
        Client connected to 129.32.0.1/32245

Control link heartbeat and jitter events:
        Oct 16 22:32:12.278 : Missed heartbeats, between 0-2912
        Oct 16 22:32:12.661 : Missed heartbeats, between 2937-2957

Hardware monitoring:
    Status:
        Activation status: Enabled
        Redundancy group 0 failover for hardware faults: Enabled
        Hardware redundancy group 0 errors: 0
        Hardware redundancy group 1 errors: 1

    Hardware monitoring events:
        Oct 16 22:53:26.686 : hw-mon errors notified  RG-0  errors:0 RG-1  errors:1
        Oct 16 22:53:26.852 : hw-mon errors notified  RG-0  errors:0 RG-1  errors:1
        Oct 16 22:53:31.774 : hw-mon errors notified  RG-0  errors:0 RG-1  errors:1
```

```
        Oct 16 23:10:46.812 : hw-mon errors notified  RG-0  errors:0 RG-1  errors:1
        Oct 16 23:10:46.813 : hw-mon errors notified  RG-0  errors:0 RG-1  errors:1
        Oct 17 00:35:06.690 : hw-mon errors notified  RG-0  errors:0 RG-1  errors:0
        Oct 17 01:15:06.691 : hw-mon errors notified  RG-0  errors:0 RG-1  errors:1
        Oct 17 01:15:06.692 : hw-mon errors notified  RG-0  errors:0 RG-1  errors:1

Schedule monitoring:
    Status:
        Activation status: Disabled
        Schedule slip detected: None
        Timer ignored: No

    Statistics:
        Total slip detected count: 0
        Longest slip duration: 0(s)

Configuration Synchronization:
    Status:
        Activation status: Enabled
        Last sync operation: NO OP
        Last sync result: Not performed
        Last sync mgd messages:

Cold Synchronization Progress:
    CS Prereq               0 of 2 SPUs completed
       1. if_state sync          0 SPUs completed
       2. fabric link            0 SPUs completed
       3. policy data sync       0 SPUs completed
       4. cp ready               0 SPUs completed
       5. VPN data sync          0 SPUs completed
       6. IPID data sync         0 SPUs completed
       7. All SPU ready          0 SPUs completed
       8. AppID ready            0 SPUs completed
       9. Tunnel Sess ready      0 SPUs completed
    CS RTO sync             0 of 2 SPUs completed
    CS Postreq              0 of 2 SPUs completed

 ISSU events:
        System currently in ISSU Window and ISSU state is 64
```

Sample Output

show chassis cluster information (Monitoring Abnormal Case)

```
user@host> show chassis cluster information
```

The following output is specific to monitoring abnormal (unhealthy) case.

```


node0:
--------------------------------------------------------------------------
Redundancy Group Information:

    Redundancy Group 0 , Current State: secondary, Weight: 255

        Time            From           To             Reason
        Apr  1 11:07:38 hold           secondary      Hold timer expired
        Apr  1 11:07:41 secondary      primary        Only node present
        Apr  1 11:29:20 primary        secondary-hold Manual failover
        Apr  1 11:34:20 secondary-hold secondary      Ready to become secondary

    Redundancy Group 1 , Current State: primary, Weight: 0

        Time            From           To             Reason
        Apr  1 11:07:38 hold           secondary      Hold timer expired
        Apr  1 11:07:41 secondary      primary        Only node present

    Redundancy Group 2 , Current State: primary, Weight: 255

        Time            From           To             Reason
        Apr  1 11:07:38 hold           secondary      Hold timer expired
        Apr  1 11:07:41 secondary      primary        Only node present

Chassis cluster LED information:
    Current LED color: Amber
    Last LED change reason: Monitored objects are down

Failure Information:

    IP Monitoring Failure Information:
    Redundancy Group 1, Monitoring Status: Failed
```

```
    IP Address        Status       Reason
    10.1.1.1           Unreachable  redundancy-group state unknown

node1:
--------------------------------------------------------------------------
Redundancy Group Information:

    Redundancy Group 0 , Current State: primary, Weight: 255

        Time            From           To             Reason
        Apr  1 11:08:40 hold           secondary      Hold timer expired
        Apr  1 11:29:20 secondary      primary        Remote is in secondary hold

    Redundancy Group 1 , Current State: secondary, Weight: 0

        Time            From           To             Reason
        Apr  1 11:08:40 hold           secondary      Hold timer expired

    Redundancy Group 2 , Current State: secondary, Weight: 255

        Time            From           To             Reason
        Apr  1 11:08:40 hold           secondary      Hold timer expired

Chassis cluster LED information:
    Current LED color: Amber
    Last LED change reason: Monitored objects are down

Failure Information:

    IP Monitoring Failure Information:
    Redundancy Group 1, Monitoring Status: Failed
    IP Address        Status       Reason
    10.1.1.1           Unreachable  redundancy-group state unknown
```

Sample Output

show chassis cluster information (Preempt Delay Timer)

```
user@host> show chassis cluster information

node0:
```

```
---------------------------------------------------------------------------
Redundancy Group Information:

    Redundancy Group 0 , Current State: secondary, Weight: 255

        Time            From           To             Reason
        Aug  4 12:30:02 hold           secondary      Hold timer expired
        Aug  4 12:30:05 secondary      primary        Only node present
        Aug  4 14:19:58 primary        secondary-hold Manual failover
        Aug  4 14:24:58 secondary-hold secondary      Ready to become secondary

    Redundancy Group 1 , Current State: secondary, Weight: 255

        Time            From           To             Reason
        Aug  4 14:07:57 secondary      primary        Remote is in secondary hold
        Aug  4 14:20:23 primary        secondary-hold Monitor failed: IF
        Aug  4 14:20:24 secondary-hold secondary      Ready to become secondary
        Aug  4 14:20:54 secondary      primary        Remote is in secondary hold
        Aug  4 14:21:30 primary        secondary-hold Monitor failed: IF
        Aug  4 14:21:31 secondary-hold secondary      Ready to become secondary

Chassis cluster LED information:
    Current LED color: Green
    Last LED change reason: No failures

node1:
---------------------------------------------------------------------------
Redundancy Group Information:

    Redundancy Group 0 , Current State: primary, Weight: 255

        Time            From           To             Reason
        Aug  4 12:33:47 hold           secondary      Hold timer expired
        Aug  4 14:19:57 secondary      primary        Remote is in secondary hold

    Redundancy Group 1 , Current State: primary, Weight: 255

        Time            From           To             Reason
        Aug  4 14:07:56 secondary-hold secondary      Ready to become secondary
        Aug  4 14:20:22 secondary      primary        Remote is in secondary hold
        Aug  4 14:20:37 primary        primary-preempt-hold Preempt (99/101)
        Aug  4 14:20:52 primary-preempt-hold secondary-hold Primary preempt hold timer e
        Aug  4 14:20:53 secondary-hold secondary      Ready to become secondary
```

```
        Aug  4 14:21:28 secondary      primary        Remote yield (99/0)

Chassis cluster LED information:
    Current LED color: Green
    Last LED change reason: No failures
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

The `Redundancy mode` option is introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

show chassis cluster status | 17405

show chassis cluster information configuration-synchronization

IN THIS SECTION

- Syntax | 17378
- Description | 17379
- Required Privilege Level | 17379
- Output Fields | 17379
- Sample Output | 17379
- Release Information | 17380

Syntax

```
show chassis cluster information configuration-synchronization
```

Description

Display chassis cluster messages. The messages indicate the redundancy mode, automatic synchronization status, and if automatic synchronization is enabled on the device.

Required Privilege Level

view

Output Fields

Table 210 on page 17379 lists the output fields for the `show chassis cluster information configuration-synchronization` command. Output fields are listed in the approximate order in which they appear.

Table 210: show chassis cluster information configuration-synchronization Output Fields

Field Name	Field Description
`Node name`	Node (device) in the chassis cluster (`node0` or `node1`).
`Status`	• Activation status—State of automatic configuration synchronization: `Enabled` or `Disabled`. • Last sync operation—Status of the last synchronization. • Last sync result—Result of the last synchronization. • Last sync mgd messages—Management daemon messages of the last synchronization.
`Events`	The timestamp of the event, the automatic configuration synchronization status, and the number of synchronization attempts.

Sample Output

show chassis cluster information configuration-synchronization

```
user@host> show chassis cluster information configuration-synchronization
```

```
node0:
--------------------------------------------------------------------------
Configuration Synchronization:
            Status:
                    Activation status: Enabled
                    Last sync operation: Auto-Sync
                    Last sync result: Not needed
                    Last sync mgd messages:
            Events:
                    Feb 25 22:21:49.174 : Auto-Sync: Not needed
node1:
--------------------------------------------------------------------------
Configuration Synchronization:
            Status:
                    Activation status: Enabled
                    Last sync operation: Auto-Sync
                    Last sync result: Succeeded
                    Last sync mgd messages:
                            mgd: rcp: /config/juniper.conf: No such file or directory
                            Network security daemon: warning: You have enabled/disabled inet6
flow.
                            Network security daemon: You must reboot the system for your change
to take effect.
                            Network security daemon: If you have deployed a cluster, be sure to
reboot all nodes.
                            mgd: commit complete
            Events:
                    Feb 25 23:02:33.467 : Auto-Sync: In progress. Attempt: 1
                    Feb 25 23:03:13.200 : Auto-Sync: Succeeded. Attempt: 1
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Understanding Automatic Chassis Cluster Synchronization Between Primary and Secondary Nodes

NTP Time Synchronization on SRX Series Devices

Example: Simplifying Network Management by Synchronizing the Primary and Backup Nodes with NTP

request chassis cluster configuration-synchronize | 14452

show chassis cluster information issu

IN THIS SECTION

- Syntax | 17381
- Description | 17381
- Required Privilege Level | 17381
- Output Fields | 17381
- Sample Output | 17382
- Release Information | 17383

Syntax

```
show chassis cluster information issu
```

Description

Display chassis cluster messages. The messages indicate the progress of the in-service software upgrade (ISSU).

Required Privilege Level

view

Output Fields

Table 211 on page 17382 lists the output fields for the `show chassis cluster information issu` command. Output fields are listed in the approximate order in which they appear.

Table 211: show chassis cluster information issu Output Fields

Field Name	Field Description
`Node name`	Node (device) in the chassis cluster (`node0` or `node1`).
`CS Prereq`	Status of all cold synchronization prerequisites: • if_state sync—Status of if_state synchronization. • fabric link—Status of fabric link synchronization. • policy data sync—Status of policy data synchronization. • cp ready—Status of the central point. • VPN data sync—Status of the VPN data synchronization.
`CS RTO sync`	Status of cold synchronization runtime objects.
`CS postreq`	Status of cold synchronization postrequirements.

Sample Output

show chassis cluster information issu

```
user@host> show chassis cluster information issu

node0:
--------------------------------------------------------------------------
Cold Synchronization Progress:
    CS Prereq              10 of 10 SPUs completed
       1. if_state sync         10 SPUs completed
       2. fabric link           10 SPUs completed
       3. policy data sync      10 SPUs completed
       4. cp ready              10 SPUs completed
       5. VPN data sync         10 SPUs completed
    CS RTO sync            10 of 10 SPUs completed
    CS Postreq             10 of 10 SPUs completed
```

```
node1:
--------------------------------------------------------------------------
Cold Synchronization Progress:
    CS Prereq              10 of 10 SPUs completed
       1. if_state sync         10 SPUs completed
       2. fabric link           10 SPUs completed
       3. policy data sync      10 SPUs completed
       4. cp ready              10 SPUs completed
       5. VPN data sync         10 SPUs completed
    CS RTO sync            10 of 10 SPUs completed
    CS Postreq             10 of 10 SPUs completed
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

show chassis cluster status | 17405

show chassis cluster interfaces

IN THIS SECTION

- Syntax | 17384
- Description | 17384
- Required Privilege Level | 17384
- Output Fields | 17384
- Sample Output | 17386
- Sample Output | 17387
- Sample Output | 17388
- Sample Output | 17389
- Sample Output | 17390

- Sample Output | 17390
- Sample Output | 17391
- Release Information | 17392

Syntax

```
show chassis cluster interfaces
```

Description

Display the status of the control interface in a chassis cluster configuration.

Required Privilege Level

view

Output Fields

Table 212 on page 17384 lists the output fields for the `show chassis cluster interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 212: show chassis cluster interfaces Output Fields

Field Name	Field Description
`Control link status`	State of the chassis cluster control interface: `up` or `down`.

Table 212: show chassis cluster interfaces Output Fields *(Continued)*

Field Name	Field Description
`Control interfaces`	• `Index`—Index number of the chassis cluster control interface. • `Interface`—Name of the chassis cluster control interface. The control interface names differ based on the routing engine. For RE2, the control interfaces are displayed as em0 and em1 and for RE3, the control interfaces are displayed as ixlv0 and igb0. • `Monitored-Status`—Monitored state of the interface: `up` or `down`. • `Internal SA`—State of the internal SA option on the chassis cluster control link: `enabled` or `disabled`. **NOTE**: This field is available only on SRX5000 line devices. • `Security`—State of MACsec on chassis cluster control interfaces.
`Fabric link status`	State of the fabric interface: `up` or `down`.
`Fabric interfaces`	• `Name`—Name of the fabric interface. • `Child-interface`—Name of the child fabric interface. • `Status`—State of the interface: `up` or `down`. • `Security`—State of MACsec on chassis cluster fabric interfaces.
`Redundant-ethernet Information`	• `Name`—Name of the redundant Ethernet interface. • `Status`—State of the interface: `up` or `down`. • `Redundancy-group`—Identification number (1–255) of the redundancy group associated with the redundant Ethernet interface.
`Redundant-pseudo-interface Information`	• `Name`—Name of the redundant pseudointerface. • `Status`—State of the redundant pseudointerface: `up` or `down`. • `Redundancy-group`—Identification number (1–255) of the redundancy group associated with the redundant pseudointerface.

Table 212: show chassis cluster interfaces Output Fields *(Continued)*

Field Name	Field Description
`Interface Monitoring`	• `Interface`—Name of the interface to be monitored. • `Weight`—Relative importance of the interface to redundancy group operation. • `Status`—State of the interface: `up` or `down`. • `Redundancy-group`—Identification number of the redundancy group associated with the interface.

Sample Output

show chassis cluster interfaces (SRX5000 line devices with RE3)

```
user@host> show chassis cluster interfaces
Control link status: Down

Control interfaces:
    Index   Interface   Monitored-Status   Internal-SA   Security
    0       ixlv0       Down               Enabled       Disabled
    1       igb0        Down               Enabled       Disabled

Fabric link status: Down

Fabric interfaces:
    Name    Child-interface    Status                    Security
                               (Physical/Monitored)
    fab0
    fab0

Redundant-ethernet Information:
    Name         Status      Redundancy-group
    reth0        Down        Not configured
    reth1        Down        Not configured
    reth2        Down        Not configured
    reth3        Down        Not configured
    reth4        Down        Not configured
```

```
Redundant-pseudo-interface Information:
    Name         Status      Redundancy-group
    lo0          Up          0
```

Sample Output

show chassis cluster interfaces (SRX5000 line devices with RE2)

```
user@host> show chassis cluster interfaces
Control link status: Up



 Control interfaces:
    Index   Interface   Monitored-Status   Internal-SA   Security
    0       em0         Up                 Disabled      Disabled
    1       em1         Down               Disabled      Disabled
Fabric link status: Up

Fabric interfaces:
    Name    Child-interface    Status                    Security
                               (Physical/Monitored)
    fab0    xe-1/0/3           Up   / Down               Disabled
    fab0
    fab1    xe-7/0/3           Up   / Down               Disabled
    fab1

Redundant-ethernet Information:
    Name         Status      Redundancy-group
    reth0        Up          1
    reth1        Up          2
    reth2        Down        Not configured
    reth3        Down        Not configured
    reth4        Down        Not configured
    reth5        Down        Not configured
    reth6        Down        Not configured
    reth7        Down        Not configured
    reth8        Down        Not configured
    reth9        Down        Not configured
    reth10       Down        Not configured
    reth11       Down        Not configured
```

```
Redundant-pseudo-interface Information:
    Name        Status       Redundancy-group
     lo0          Up               1

Interface Monitoring:
    Interface         Weight    Status    Redundancy-group
    ge-0/1/9          100       Up        0
    ge-0/1/9          100       Up
```

Sample Output

show chassis cluster interfaces

```
user@host> show chassis cluster interfaces
```

The following output is specific to fabric monitoring failure:

```
Control link status: Up

Control interfaces:
    Index   Interface    Monitored-Status   Internal-SA  Security
    0       em0         Up                 Disabled     Disabled

Fabric link status: Down

Fabric interfaces:
    Name    Child-interface    Status                           Security
                               (Physical/Monitored)
    fab0    ge-0/0/2           Down / Down                    Disabled
    fab0
    fab1    ge-9/0/2           Up   / Up                      Disabled
    fab1

Redundant-pseudo-interface Information:
    Name         Status      Redundancy-group
    lo0          Up          0
```

Sample Output

show chassis cluster interfaces (SRX5400, SRX5600, and SRX5800 Devices with SRX5000 line SRX5K-SCB3 [SCB3] with Enhanced Midplanes and SRX5K-MPC3-100G10G [IOC3] or SRX5K-MPC3-40G10G [IOC3])

```
user@host> show chassis cluster interfaces
```

The following output is specific to SRX5400, SRX5600, and SRX5800 devices in a chassis cluster cluster, when the PICs containing fabric links on the SRX5K-MPC3-40G10G (IOC3) are powered off to turn on alternate PICs. If no alternate fabric links are configured on the PICs that are turned on, RTO synchronous communication between the two nodes stops and the chassis cluster session state will not back up, because the fabric link is missing.

```


Control link status: Up

Control interfaces:
    Index   Interface   Monitored-Status   Internal-SA   Security
    0       em0         Up                 Disabled        Disabled
    1       em1         Down               Disabled        Disabled

Fabric link status: Down

Fabric interfaces:
    Name    Child-interface    Status                            Security
                               (Physical/Monitored)
    fab0     <<< fab child missing once PIC off lined    Disabled
    fab0
    fab1    xe-10/2/7          Up   / Down               Disabled
    fab1

Redundant-ethernet Information:
    Name         Status      Redundancy-group
    reth0        Up          Not configured
    reth1        Down        1

Redundant-pseudo-interface Information:
    Name         Status      Redundancy-group
```

```
    lo0          Up          0
```

Sample Output

show chassis cluster interfaces (vSRX Virtual Firewall)

```
user@host> show chassis cluster interfaces
```

The following output is specific to view control link status with internal SAs.

```


Control link status: Up

Control interfaces:
 Index     Interface           Status         Internal SA
 305
      0       em0              Up              enabled
306
      1       em1              Down            enabled
```

Sample Output

show chassis cluster interfaces (SRX1500 Devices)

```
user@host> show chassis cluster interfaces
```

Check chassis cluster control link and fabric link status.

```
Control link status: Up

Control interfaces:
    Index   Interface   Monitored-Status   Internal-SA   Security
    0       em0         Up                 Disabled      Disabled

Fabric link status: Up
```

```
Fabric interfaces:
    Name    Child-interface    Status                    Security
                               (Physical/Monitored)
    fab0    ge-0/0/11          Up   / Up                 Disabled
    fab0
    fab1    ge-7/0/11          Up   / Up                 Disabled
    fab1

Redundant-ethernet Information:
    Name         Status      Redundancy-group
    reth0        Up          1
    reth1        Up          1
    reth2        Down        Not configured
    reth3        Down        Not configured
    reth4        Down        Not configured

Redundant-pseudo-interface Information:
    Name         Status      Redundancy-group
    lo0          Up          1

Interface Monitoring:
    Interface         Weight    Status                    Redundancy-group
                                (Physical/Monitored)
    ge-0/0/4          255       Up  /  Up                 1
    ge-7/0/4          255       Up  /  Up                 1
    ge-0/0/2          255       Up  /  Up                 1
    ge-7/0/2          255       Up  /  Up                 1
```

command-name

After you disable control link:

```
Control link status: Admin Down

Control interfaces:
    Index   Interface   Monitored-Status   Internal-SA   Security
0       em0         Down                 Disabled      Disabled
```

Sample Output

Release Information

Command modified in Junos OS Release 9.0. Output changed to support dual control ports in Junos OS Release 10.0. Output changed to support control interfaces in Junos OS Release 11.2. Output changed to support redundant pseudo interfaces in Junos OS Release 12.1X44-D10. For SRX5000 line devices, output changed to support the internal security association (SA) option in Junos OS Release 12.1X45-D10. Output changed to support MACsec status on control and fabric interfaces in Junos OS Release 15.1X49-D60. For vSRX Virtual Firewall, output changed to support the internal security association (SA) option in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

show chassis cluster ip-monitoring status redundancy-group

IN THIS SECTION

- Syntax | 17392
- Description | 17393
- Options | 17393
- Required Privilege Level | 17393
- Output Fields | 17393
- Sample Output | 17395
- Sample Output | 17395
- Release Information | 17396

Syntax

```
show chassis cluster ip-monitoring status
  <redundancy-group  group-number>
```

Description

Display the status of all monitored IP addresses for a redundancy group.

Options

- none— Display the status of monitored IP addresses for all redundancy groups on the node.
- `redundancy-group` *`group-number`*— Display the status of monitored IP addresses under the specified redundancy group.

Required Privilege Level

view

Output Fields

Table 213 on page 17393 lists the output fields for the `show chassis cluster ip-monitoring status` command.

Table 213: show chassis cluster ip-monitoring status Output Fields

Field Name	Field Description
`Redundancy-group`	ID number (0 - 255) of a redundancy group in the cluster.
`Global threshold`	Failover value for all IP addresses monitored by the redundancy group.
`Current threshold`	Value equal to the global threshold minus the total weight of the unreachable IP address.
`IP Address`	Monitored IP address in the redundancy group.
`Status`	Current reachability state of the monitored IP address. Values for this field are: `reachable`, `unreachable`, and `unknown`. The status is “unknown” if Packet Forwarding Engines (PFEs) are not yet up and running.
`Failure count`	Number of attempts to reach an IP address.

Table 213: show chassis cluster ip-monitoring status Output Fields *(Continued)*

Field Name	Field Description
`Reason`	Explanation for the reported status. See Table 214 on page 17394.
`Weight`	Combined weight (0 - 255) assigned to all monitored IP addresses. A higher weight value indicates greater importance.

Expanded reason output fields for unreachable IP addresses added in Junos OS Release 10.1. You might see any of the following reasons displayed.

Table 214: show chassis cluster ip-monitoring status redundancy group Reason Fields

Reason	Reason Description
`No route to host`	The router could not resolve the ARP, which is needed to send the ICMP packet to the host with the monitored IP address.
`No auxiliary IP found`	The redundant Ethernet interface does not have an auxiliary IP address configured.
`Reth child not up`	A child interface of a redundant Ethernet interface is down.
`redundancy-group state unknown`	Unable to obtain the state (primary, secondary, secondary-hold, disable) of a redundancy-group.
`No reth child MAC address`	Could not extract the MAC address of the redundant Ethernet child interface.
`Secondary link not monitored`	The secondary link might be down (the secondary child interface of a redundant Ethernet interface is either down or non-functional).
`Unknown`	The IP address has just been configured and the router still does not know the status of this IP. or Do not know the exact reason for the failure.

Sample Output

show chassis cluster ip-monitoring status

```
user@host> show chassis cluster ip-monitoring status
node0:
--------------------------------------------------------------------------

Redundancy group: 1
Global threshold: 200
Current threshold: -120

IP address                Status        Failure count  Reason    Weight
10.254.5.44               reachable       0            n/a       220
10.2.2.1                  reachable       0            n/a       100

node1:
--------------------------------------------------------------------------

Redundancy group: 1
Global threshold: 200
Current threshold: -120

IP address                Status        Failure count  Reason    Weight
10.254.5.44               reachable       0            n/a       220
10.2.2.1                  reachable       0            n/a       100
```

Sample Output

show chassis cluster ip-monitoring status redundancy-group

```
user@host> show chassis cluster ip-monitoring status redundancy-group 1
node0:
--------------------------------------------------------------------------

Redundancy group: 1

IP address                Status        Failure count  Reason
10.254.5.44               reachable       0            n/a
10.2.2.1                  reachable       0            n/a
```

```
10.1.1.5                  reachable       0            n/a
10.1.1.4                  reachable       0            n/a
10.1.1.1                  reachable       0            n/a

node1:
--------------------------------------------------------------------------

Redundancy group: 1

IP address                Status        Failure count  Reason
10.254.5.44               reachable       0            n/a
10.2.2.1                  reachable       0            n/a
10.1.1.5                  reachable       0            n/a
10.1.1.4                  reachable       0            n/a
10.1.1.1                  reachable       0            n/a
```

Release Information

Command introduced in Junos OS Release 9.6. Support.

RELATED DOCUMENTATION

clear chassis cluster failover-count

show chassis cluster port-peering

IN THIS SECTION

- Syntax | 17397
- Description | 17397
- Required Privilege Level | 17397
- Output Fields | 17397
- Sample Output | 17398
- Release Information | 17398

Syntax

```
show chassis cluster port-peering
```

Description

Display the status of the L3 back-end interface and the corresponding L2 mapped interface.

Before running this command, you must deploy the chassis-cluster and map the L2 and L3 interfaces:

```
user@host# set groups $node chassis cluster redundant-interface L3 interface mapping-interface
L2 interface
```

Required Privilege Level

view

Output Fields

Table 215 on page 17397 lists the output fields for the `show chassis cluster port-peering` command. Output fields are listed in the approximate order in which they appear.

Table 215: show chassis cluster port-peering Output Fields

Field Name	Field Description
`Backend L3 Interface`	L3 interface of the peer map.
`Backend L3 Interface status`	Status of the L3 interface.
`Mapped L2 Interface`	L2 interface of the peer map.
`Mapped L2 Interface status`	Status of the L2 interface.

Sample Output

show chassis cluster port-peering

```
user@host> show chassis cluster port-peering
{primary:node0}
user@host> show chassis cluster port-peering
node0:
--------------------------------------------------------------------------

Port peering interfaces:
    Backend L3                Mapping L2
    Interface  Status        Interface  Status
    ge-1/0/1   Up              ge-0/0/1   Up

node1:
--------------------------------------------------------------------------

Port peering interfaces:
    Backend L3                Mapping L2
    Interface  Status        Interface  Status
    ge-8/0/1   Up             ge-7/0/1   Up
```

Release Information

Command introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

show chassis cluster statistics

IN THIS SECTION

- Syntax | 17399
- Description | 17399
- Required Privilege Level | 17399
- Output Fields | 17399
- Sample Output | 17402
- Sample Output | 17403
- Sample Output | 17404
- Release Information | 17405

Syntax

```
show chassis cluster statistics
```

Description

This command displays information about chassis cluster services and interfaces.

Required Privilege Level

view

Output Fields

Table 216 on page 17400 lists the output fields for the `show chassis cluster statistics` command. Output fields are listed in the approximate order in which they appear.

Table 216: show chassis cluster statistics Output Fields

Field Name	Field Description
`Control link statistics`	Statistics of the control link used by chassis cluster traffic. Statistics for `Control link 1` are displayed when you use dual control links (SRX5000 lines only). Note that the output for the SRX5000 lines will always show `Control link 0` and `Control link 1` statistics, even though only one control link is active or working. • `Heartbeat packets sent`—Number of heartbeat messages sent on the control link. • `Heartbeat packets received`—Number of heartbeat messages received on the control link. • `Heartbeat packet errors`—Number of heartbeat packets received with errors on the control link.
`Fabric link statistics`	Statistics of the fabric link used by chassis cluster traffic. Statistics for `Child Link 1` are displayed when you use dual fabric links. • `Probes sent`—Number of probes sent on the fabric link. • `Probes received`—Number of probes received on the fabric link.

Table 216: show chassis cluster statistics Output Fields *(Continued)*

Field Name	Field Description
`Services Synchronized`	• `Service name`—Name of the service. • `Rtos sent`—Number of runtime objects (RTOs) sent. • `Rtos received`—Number of RTOs received. • `Translation context`—Messages synchronizing Network Address Translation (NAT) translation context. • `Incoming NAT`—Messages synchronizing incoming Network Address Translation (NAT) service. • `Resource manager`—Messages synchronizing resource manager groups and resources. • `Session create`—Messages synchronizing session creation. • `Session close`—Messages synchronizing session close. • `Session change`—Messages synchronizing session change. • `Gate create`—Messages synchronizing creation of pinholes (temporary openings in the firewall). • `Session ageout refresh request`—Messages synchronizing request session after age-out. • `Session ageout refresh reply`—Messages synchronizing reply session after age-out. • `IPsec VPN`—Messages synchronizing VPN session. • `Firewall user authentication`—Messages synchronizing firewall user authentication session. • `MGCP ALG`—Messages synchronizing MGCP ALG sessions. • `H323 ALG`—Messages synchronizing H.323 ALG sessions. • `SIP ALG`—Messages synchronizing SIP ALG sessions. • `SCCP ALG`—Messages synchronizing SCCP ALG sessions.

Table 216: show chassis cluster statistics Output Fields *(Continued)*

Field Name	Field Description
	• `PPTP ALG`—Messages synchronizing PPTP ALG sessions. • `RTSP ALG`—Messages synchronizing RTSP ALG sessions. • `MAC address learning`—Messages synchronizing MAC address learning.

Sample Output

show chassis cluster statistics

```
user@host> show chassis cluster statistics
Control link statistics:
    Control link 0:
        Heartbeat packets sent: 798
        Heartbeat packets received: 784
        Heartbeat packets errors: 0
Fabric link statistics:
    Child link 0
        Probes sent: 793
        Probes received: 0
Services Synchronized:
    Service name                            RTOs sent    RTOs received
    Translation context                     0            0
    Incoming NAT                            0            0
    Resource manager                        0            0
    Session create                          0            0
    Session close                           0            0
    Session change                          0            0
    Gate create                             0            0
    Session ageout refresh requests         0            0
    Session ageout refresh replies          0            0
    IPsec VPN                               0            0
    Firewall user authentication            0            0
    MGCP ALG                                0            0
    H323 ALG                                0            0
    SIP ALG                                 0            0
```

```
    SCCP ALG                                0            0
    PPTP ALG                                0            0
    RTSP ALG                                0            0
    MAC address learning                    0            0
```

Sample Output

show chassis cluster statistics (SRX5000 Line Devices)

```
user@host> show chassis cluster statistics
Control link statistics:
    Control link 0:
        Heartbeat packets sent: 258689
        Heartbeat packets received: 258684
        Heartbeat packets errors: 0
    Control link 1:
        Heartbeat packets sent: 258689
        Heartbeat packets received: 258684
        Heartbeat packets errors: 0
Fabric link statistics:
    Child link 0
        Probes sent: 258681
        Probes received: 258681
    Child link 1
        Probes sent: 258501
        Probes received: 258501
Services Synchronized:
    Service name                            RTOs sent    RTOs received
    Translation context                     0            0
    Incoming NAT                            0            0
    Resource manager                        0            0
    Session create                          1            0
    Session close                           1            0
    Session change                          0            0
    Gate create                             0            0
    Session ageout refresh requests         0            0
    Session ageout refresh replies          0            0
    IPSec VPN                               0            0
    Firewall user authentication            0            0
    MGCP ALG                                0            0
    H323 ALG                                0            0
```

```
    SIP ALG                                 0            0
    SCCP ALG                                0            0
    PPTP ALG                                0            0
    RPC ALG                                 0            0
    RTSP ALG                                0            0
    RAS ALG                                 0            0
    MAC address learning                    0            0
    GPRS GTP                                0            0
```

Sample Output

show chassis cluster statistics (SRX5000 Line Devices)

```
user@host> show chassis cluster statistics
Control link statistics:
    Control link 0:
        Heartbeat packets sent: 82371
        Heartbeat packets received: 82321
        Heartbeat packets errors: 0
    Control link 1:
        Heartbeat packets sent: 0
        Heartbeat packets received: 0
        Heartbeat packets errors: 0
Fabric link statistics:
    Child link 0
        Probes sent: 258681
        Probes received: 258681
    Child link 1
        Probes sent: 258501
        Probes received: 258501
Services Synchronized:
    Service name                            RTOs sent    RTOs received
    Translation context                     0            0
    Incoming NAT                            0            0
    Resource manager                        0            0
    Session create                          1            0
    Session close                           1            0
    Session change                          0            0
    Gate create                             0            0
    Session ageout refresh requests         0            0
    Session ageout refresh replies          0            0
```

```
    IPSec VPN                                0            0
    Firewall user authentication             0            0
    MGCP ALG                                 0            0
    H323 ALG                                 0            0
    SIP ALG                                  0            0
    SCCP ALG                                 0            0
    PPTP ALG                                 0            0
    RPC ALG                                  0            0
    RTSP ALG                                 0            0
    RAS ALG                                  0            0
    MAC address learning                     0            0
    GPRS GTP                                 0            0
```

Release Information

Command modified in Junos OS Release 9.0.

Output changed to support dual control ports in Junos OS Release 10.0.

RELATED DOCUMENTATION

clear chassis cluster statistics | 3463

show chassis cluster status

IN THIS SECTION

- Syntax | 17406
- Description | 17406
- Options | 17406
- Required Privilege Level | 17406
- Output Fields | 17406
- Sample Output | 17408
- Sample Output | 17408
- Sample Output | 17409

Release Information | 17409

Syntax

```
show chassis cluster status
<redundancy-group  group-number >
```

Description

Display the current status of the Chassis Cluster. You can use this command to check the status of chassis cluster nodes, redundancy groups, and failover status.

Options

- none—Display the status of all redundancy groups in the chassis cluster.
- `redundancy-group` *`group-number`*—(Optional) Display the status of the specified redundancy group.

Required Privilege Level

view

Output Fields

Table 217 on page 17406 lists the output fields for the `show chassis cluster status` command. Output fields are listed in the approximate order in which they appear.

Table 217: show chassis cluster status Output Fields

Field Name	Field Description
`Cluster ID`	ID number (1-15) of a cluster is applicable for releases upto Junos OS Release 12.1X45-D10. ID number (1-255) is applicable for Releases 12.1X45-D10 and later. Setting a cluster ID to 0 is equivalent to disabling a cluster.

Table 217: show chassis cluster status Output Fields *(Continued)*

Field Name	Field Description
`Redundancy-Group`	You can create up to 128 redundancy groups in the chassis cluster.
`Node name`	Node (device) in the chassis cluster (`node0` or `node1`).
`Priority`	Assigned priority for the redundancy group on that node.
`Status`	State of the redundancy group (`Primary`, `Secondary`, `Lost`, or `Unavailable`). • `Primary`—Redundancy group is active and passing traffic. • `Secondary`—Redundancy group is passive and not passing traffic. • `Lost`—Node loses contact with the other node through the control link. Most likely to occur when both nodes are in a cluster and there is a control link failure, one node cannot exchange heartbeats, or when the other node is rebooted. • `Unavailable`—Node has not received a single heartbeat over the control link from the other node since the other node booted up. Most likely to occur when one node boots up before the other node, or if only one node is present in the cluster.
`Preempt`	• `Yes`: Primary state can be preempted based on priority. • `No`: Change in priority will not preempt the primary state.
`Manual failover`	• `Yes`: Primary state is set manually through the CLI with the `request chassis cluster failover node` or `request chassis cluster failover redundancy-group` command. This overrides `Priority` and `Preempt`. • `No`: Primary state is not set manually through the CLI.
Monitor-failures	• `None`: Cluster working properly. • `Monitor Failure code`: Cluster is not working properly and the respective failure code is displayed.

Sample Output

show chassis cluster status

```
user@host> show chassis cluster status

Monitor Failure codes:
        CS  Cold Sync monitoring        FL  Fabric Connection monitoring
    GR  GRES monitoring             HW  Hardware monitoring
    IF  Interface monitoring        IP  IP monitoring
    LB  Loopback monitoring         MB  Mbuf monitoring
    NH  Nexthop monitoring          NP  NPC monitoring
    SP  SPU monitoring              SM  Schedule monitoring
    CF  Config Sync monitoring

Cluster ID: 1
Node   Priority Status         Preempt Manual   Monitor-failures

Redundancy group: 0 , Failover count: 1
node0  200      primary        no      no       None
node1  1        secondary      no      no       None

Redundancy group: 1 , Failover count: 1
node0  101      primary        no      no       None
node1  1        secondary      no      no       None
```

Sample Output

show chassis cluster status with preemptive delay

```
user@host> show chassis cluster status

Cluster ID: 1
Node   Priority Status         Preempt Manual   Monitor-failures
Redundancy group: 0, Failover count: 1
node0  200      primary        no      no       None
node1  100      secondary      no      no       None
Redundancy group: 1, Failover count: 3
```

```
node0  200      primary-preempt-hold        yes     no       None node1  100
secondary                     yes      no        None
```

Sample Output

show chassis cluster status redundancy-group 1

```
user@host> show chassis cluster status redundancy-group 1

Monitor Failure codes:
    CS  Cold Sync monitoring        FL  Fabric Connection monitoring
    GR  GRES monitoring             HW  Hardware monitoring
    IF  Interface monitoring        IP  IP monitoring
    LB  Loopback monitoring         MB  Mbuf monitoring
    NH  Nexthop monitoring          NP  NPC monitoring
    SP  SPU monitoring              SM  Schedule monitoring
    CF  Config Sync monitoring

Cluster ID: 1
Node   Priority Status         Preempt Manual   Monitor-failures

Redundancy group: 1 , Failover count: 1
node0  101      primary        no      no       None
node1  1        secondary      no      no       None
```

Release Information

Support for monitoring failures added in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

- redundancy-group (Chassis Cluster)
- clear chassis cluster failover-count
- request chassis cluster failover node
- request chassis cluster failover reset

show chassis craft-interface

IN THIS SECTION

- Syntax | 17410
- Syntax (MX Series Routers) | 17410
- Description | 17410
- Options | 17410
- Required Privilege Level | 17411
- Output Fields | 17411
- Sample Output | 17413
- Release Information | 17418

Syntax

```
show chassis craft-interface
```

Syntax (MX Series Routers)

```
show chassis craft-interface
<all-members>
<local>
<member member-id>
```

Description

For routers or switches that have a display on the craft interface, show the messages that are currently displayed. On all routers, you must enter this command on the primary Routing Engine.

Options

none Show messages that are currently displayed on the craft interface.

all-members (MX Series routers only) (Optional) Display information currently on the craft interface for all members of the Virtual Chassis configuration.

local (MX Series routers only) (Optional) Display information currently on the craft interface for the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display information currently on the craft interface for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

Required Privilege Level

view

Output Fields

Table 218 on page 17411 lists the output fields for the `show chassis craft-interface` command. Output fields are listed in the approximate order in which they appear.

Table 218: show chassis craft-interface Output Fields

Field Name	Field Description
LCD screen or **FPM Display Contents**	Contents of the Front Panel Module display: • *router-name*—Name of the router. • **Up**—How long the router has been operational, in days, hours, minutes, and seconds. • *message*—Information about the router traffic load, the power supply status, the fan status, and the temperature status. The display of this information changes every 2 seconds. If a text message has been created with the `set chassis display` command, this message appears on all four lines of the craft interface display. The display alternates between the text message and the standard system status messages every 2 seconds.
`Front Panel System LEDs`	(MX2010, MX2020, and MX2008 Routers) Status of the Front Panel System LEDs. A dot (.) indicates the LED is not lit. An asterisk (*) indicates the LED is lit.
`Front Panel Alarm Indicators`	(MX2010, MX2020, and MX2008 Routers) Status of the Front Panel Alarm indicators. A dot (.) indicates the relay is off. An asterisk (*) indicates the relay is active.

Table 218: show chassis craft-interface Output Fields *(Continued)*

Field Name	Field Description
Input Relay	Status of the configured input relay ports—0 through 3. The mode is normally open or closed. The status is clear or raised.
Output Relay	Status of the configured output ports—0 or 1. The mode is normally open or closed. The status is clear or raised.
Front Panel FPC LEDs	(MX2010, MX2020, and MX2008 Routers) Status of the Front Panel Flexible PIC Concentrator (FPC) LEDs. A dot (**.**) indicates the LED is not lit. An asterisk (*****) indicates the LED is lit. On MX2010 and MX2008 routers, there are 10 (0-9) FPCs LEDs. On MX2020 routers, there are 20 (0-9 and 10-19) FPCs LEDs.
CB LEDs	Status of the Control Board (CB) LEDs. A dot (**.**) indicates the LED is not lit. An asterisk (*****) indicates the LED is lit.
`PS LEDs`	(MX2010, MX2020, and MX2008 Routers) Status of the Power Supply (PS) LEDs. A dot (**.**) indicates the LED is not lit. An asterisk (*****) indicates the LED is lit. On MX2010 and MX2008 routers, there are 9 (0-8) PS LEDs. On MX2020 routers, there are 18 (0-8 and 9-17) PS LEDs.
`FAN Tray LEDs`	(MX2010, MX2020, and MX2008 Routers) Status of the Fan Tray LEDs. A dot (**.**) indicates the LED is not lit. An asterisk (*****) indicates the LED is lit.
`Front Panel SFB LEDs`	(MX2010, MX2020, and MX2008 Routers) Status of the Front Panel Switch Fabric Boards (SFB) LEDs. A dot (**.**) indicates the LED is not lit. An asterisk (*****) indicates the LED is lit.
`Front Panel Chassis Info`	(MX2010, MX2020, and MX2008 Routers) Information about the chassis such as the chassis number and role. User can set the chassis number in multi-chassis configurations.
SIB LEDs	Status of the Switch Interface Board (SIB) LEDs. A dot (**.**) indicates the LED is not lit. An asterisk (*****) indicates the LED is lit.
Front Panel CB LEDs	(PTX Series Routers) Status of the Front Panel Control Board (CB) LEDs. A dot (**.**) indicates the LED is not lit. An asterisk (*****) indicates the LED is lit.

Sample Output

show chassis craft-interface (MX Series Routers)

```
user@host> show chassis craft-interface
Front Panel System LEDs:
Routing Engine    0    1
--------------------------
OK                *    *
Fail              .    .
Master            *    .

Front Panel Alarm Indicators:
-----------------------------
Red LED      .
Yellow LED   .
Major relay  .
Minor relay  .

Front Panel FPC LEDs:
FPC    0    1    2    3    4    5    6    7    8    9
-----------------------------------------------------
Red    .    .    .    .    .    .    .    .    .    .
Green  *    *    .    .    .    .    .    *    .    .

CB LEDs:
 CB    0   1
------------
Amber  .   .
Green  *   *

PS LEDs:
  PS    0    1    2    3    4    5    6    7    8
-------------------------------------------------
Red     .    .    .    .    .    .    .    .    .
Green   .    *    *    *    *    *    *    *    .

Fan Tray LEDs:
  FT   0   1
------------
Red    .   .
Green  *   *
```

```
Front Panel SFB LEDs:
SFB    0   1   2   3   4   5   6   7
------------------------------------
Red    .   .   .   .   .   .   .   .
Green  *   *   *   *   *   *   *   *

Front Panel Chassis Info:
Chassis Number    0x36
Chassis Role      M
```

show chassis craft-interface (PTX Series Routers)

```
user@host> show chassis craft-interface
Front Panel System LEDs:
Routing Engine    0    1
------------------------
OK                *    *
Fail              .    .
Master            *    .

Front Panel Alarm Indicators:
-----------------------------
Red LED      *
Yellow LED   *

Front Panel FPC LEDs:
FPC      0   1   2   3   4   5   6   7
--------------------------------------
Fail     .   .   .   .   .   .   .   .
OK       *   .   .   .   .   .   .   .
Active   .   .   .   .   .   .   .   .

Front Panel SIB LEDs:
SIB      0   1   2   3   4   5
------------------------------
Fail     .   .   .   .   .   .
OK       *   .   .   *   *   .
Active   .   .   .   .   .   .

Front Panel Fan Tray LEDs:
```

```
FAN     0   1
-------------
Red     .   .
Green   *   *

Front Panel PSM LEDs:
PSM     0   1   2   3   4   5
-----------------------------
Red     .   .   .   .   .   .
Green   *   *   *   *   .   .

Front Panel CB LEDs:
CB             0   1
----------------------
Green          *   *
Yellow Blink   .   .
```

show chassis craft-interface (ACX7100)

```
user@host> show chassis craft-interfaceFront Panel System LEDs:
Routing Engine    0
-------------------
OK                *
Fail              .
Master            *

Front Panel Alarm Indicators:
-----------------------------
Red LED      *
Yellow LED   *

Front Panel FPC LEDs:
FPC      0
----------
Fail     .
OK       .
Active   .

Front Panel Fan Tray LEDs:
FAN     0   1   2   3   4   5
-----------------------------
```

```
Red     .   .   .   .   .   .
Green   *   *   *   *   *   *

Front Panel CB LEDs:
CB      0
---------
Amber   .
Green   *
Blue    *
```

show chassis craft-interface (ACX7100 Router)

```
user@host> show chassis craft-interfaceFront Panel System LEDs:
Routing Engine    0
-------------------
OK                *
Fail              .
Master            *

Front Panel Alarm Indicators:
-----------------------------
Red LED      *
Yellow LED   *

Front Panel FPC LEDs:
FPC      0
----------
Fail     .
OK       .
Active   .

Front Panel Fan Tray LEDs:
FAN     0   1   2   3   4   5
-----------------------------
Red     .   .   .   .   .   .
Green   *   *   *   *   *   *

Front Panel CB LEDs:
CB      0
---------
Amber   .
```

```
Green   *
Blue    *
```

show chassis craft-interface (ACX7024 Router)

```
user@host> show chassis craft-interfaceFront Panel System LEDs:Front Panel System LEDs:
Routing Engine    0
-------------------
OK                *
Fail              .
Master            *

Front Panel Alarm Indicators:
-----------------------------
Red LED      *
Yellow LED   *

Front Panel FPC LEDs:
FPC      0
----------
Fail     .
OK       .
Active   .

Front Panel Fan Tray LEDs:
FAN     0
---------
Red     .
Green   *


Front Panel CB LEDs:
CB      0
---------
Amber   .
Green   *
Blue    *
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear chassis display message | 3465

set chassis display message | 16865

show chassis dedicated-ukern-cpu

IN THIS SECTION

- Syntax | 17418
- Description | 17418
- Options | 17419
- Required Privilege Level | 17419
- Output Fields | 17419
- Sample Output | 17419
- Release Information | 17419

Syntax

```
show chassis dedicated-ukern-cpu
```

Description

Display whether dedicated Bidirectional Forwarding Detection (BFD) is enabled or disabled. If dedicated BFD is enabled, the output of the show command displays the value of the `Dedicated Ukern CPU Status` field as `Enabled`.

Options

This command has no options.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

show chassis dedicated-ukern-cpu

```
user@host> show chassis dedicated-ukern-cpu
Dedicated Ukern CPU Status: Enabled
```

Release Information

Command introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Enabling Dedicated and Real-Time BFD

dedicated-ukern-cpu (BFD) | **5174**

Understanding BFD for BGP

show chassis device-mode

IN THIS SECTION

Syntax | **17420**

- Description | 17420
- Options | 17420
- Required Privilege Level | 17420
- Output Fields | 17420
- Sample Output | 17421
- Release Information | 17423

Syntax

```
show chassis device-mode
```

Description

Display information about the operating mode of the device. For example, QFX3500 devices operate either as a single switch in standalone mode or as a QFabric system Node device in node-device mode.

> **NOTE**: Issue the `show chassis device-mode` command only when your management station is connected directly to the device over a console port connection.

Options

There are no options for this command.

Required Privilege Level

admin

Output Fields

Table 219 on page 17421 lists the output fields for the `show chassis device-mode` command. Output fields are listed in the approximate order in which they appear.

Table 219: show chassis device-mode Output Fields

Field Name	Field Description
Current device-mode	Existing operational mode for the device. The device can be in Interconnect device mode, Node device mode, or standalone mode.
Future device-mode after reboot	Future operational mode for the device after you reboot it. The device can be set to enter Interconnect device mode, Node device mode, or standalone mode. **NOTE**: To set the future mode of the device, issue the `request chassis device-mode` command.

Sample Output

show chassis device-mode (Interconnect Device Mode)

```
user@switch> show chassis device-mode
Current device-mode : Interconnect-device
Future device-mode after reboot : Interconnect-device
```

show chassis device-mode (Interconnect Device Mode, but Node Device-Ready)

```
user@switch> show chassis device-mode
Current device-mode : Interconnect-device
Future device-mode after reboot : Node-device
```

show chassis device-mode (Interconnect Device Mode, but Standalone-Ready)

```
user@switch> show chassis device-mode
Current device-mode : Interconnect-device
Future device-mode after reboot : Standalone
```

show chassis device-mode (Node Device Mode)

```
user@switch> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Node-device
```

show chassis device-mode (Node Device Mode, but Interconnect Device-Ready)

```
user@switch> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Interconnect-device
```

show chassis device-mode (Node Device Mode, but Standalone-Ready)

```
user@switch> show chassis device-mode
Current device-mode : Node-device
Future device-mode after reboot : Standalone
```

show chassis device-mode (Standalone Mode)

```
user@switch> show chassis device-mode
Current device-mode : Standalone
Future device-mode after reboot : Standalone
```

show chassis device-mode (Standalone Mode, but Interconnect Device-Ready)

```
user@switch> show chassis device-mode
Current device-mode : Standalone
Future device-mode after reboot : Interconnect-device
```

show chassis device-mode (Standalone Mode, but Node Device-Ready)

```
user@switch> show chassis device-mode
Current device-mode : Standalone
Future device-mode after reboot : Node-device
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Converting the Device Mode for a QFabric System Component

request chassis device-mode | **14460**

Understanding Interconnect Devices

Understanding Node Devices

show chassis enhanced-temperature-thresholds

IN THIS SECTION

- Syntax | **17423**
- Description | **17424**
- Required Privilege Level | **17424**
- Output Fields | **17424**
- Sample Output | **17425**
- Release Information | **17429**

Syntax

```
show chassis enhanced-temperature-thresholds
```

Description

Display chassis temperature threshold settings, in degrees Celsius.

Required Privilege Level

view

Output Fields

The following table lists the output fields for the `show chassis enhanced-temperature-thresholds` command. Output fields are listed in the approximate order in which they appear.

Table 220: show chassis enhanced-temperature-thresholds Output Fields

Field name	Field Description
Item	Chassis component. If per FRU per slot thresholds are configured, the components about which information is displayed include the chassis, the Routing Engines, CPUs, FPCs, and FEBs. If per FRU per slot thresholds are not configured, the components about which information is displayed include the chassis and the Routing Engines.
Fan speed	Temperature threshold settings, in degrees Celsius, for the fans to operate at normal and high speeds. • Normal—The fans operate at normal speed if the component is at or below this temperature and all the fans are present and functioning normally. • High—The fans operate at high speed if the component has exceeded this temperature or a fan has failed or is missing. • Full—The fans operate at it's full speed. • Bad fan—The temperature that must be exceeded on the component when one or more fans have failed or are missing. An alarm is not triggered until the temperature exceeds the threshold settings for a yellow alarm or a red alarm.

Table 220: show chassis enhanced-temperature-thresholds Output Fields *(Continued)*

Field name	Field Description
Yellow alarm	Temperature threshold settings, in degrees Celsius, that trigger a yellow alarm. • Normal—The temperature that must be exceeded on the component to trigger a yellow alarm when the fans are running at full speed. • Bad fan—The temperature that must be exceeded on the component to trigger a yellow alarm when one or more fans have failed or are missing.
Red alarm	Temperature threshold settings, in degrees Celsius, that trigger a red alarm. • Normal—The temperature that must be exceeded on the component to trigger a red alarm when the fans are running at full speed. • Bad fan—The temperature that must be exceeded on the component to trigger a red alarm when one or more fans have failed or are missing.
Fire Shutdown	Temperature threshold settings, in degrees Celsius, for the network device to shut down.
State	State of the temperature sensors. Displays whether sensor is enabled for EM Policy.

Sample Output

show chassis enhanced-temperature-thresholds (SRX4700)

```
user@host> show chassis enhanced-temperature-thresholds
                                        Fan speed     Fan speed      Yellow alarm          Red
alarm      Fire Shutdown    State
                                       (degrees C)   (degrees C)      (degrees C)
(degrees C)      (degrees C)
Item                                  Normal  High   Full  Bad fan    Normal  Bad fan
Normal  Bad fan     Normal
CB 0 TMP435 Exhaust Left                 38    43      53     53        60       57
65       62         70        Enabled
CB 0 TMP435 Intake Left                  38    43      53     53        60       57
65       62         70        Enabled
CB 0 TMP435 Exhaust Right                50    55      60     57        65       62
70       67         75        Enabled
```

```
CB 0 TMP435 Intake Right                  50    55      60     57        65       62
70       67         75        Enabled
CB 0 CPU0 DTS                             75    85      92     89        94       91
95       92         96        Enabled
CB 0 CPU1 DTS                             75    85      92     89        94       91
95       92         96        Enabled
CB 0 DDR5 DIMM CPU0_AB                    60    65      80     77        85       82
95       92        100        Enabled
CB 0 DDR5 DIMM CPU0_CD                    60    65      80     77        85       82
95       92        100        Enabled
CB 0 DDR5 DIMM CPU0_EF                    60    65      80     77        85       82
95       92        100        Enabled
CB 0 DDR5 DIMM CPU0_GH                    60    65      80     77        85       82
95       92        100        Enabled
CB 0 DDR5 DIMM CPU1_AB                    60    65      80     77        85       82
95       92        100        Enabled
CB 0 DDR5 DIMM CPU1_CD                    60    65      80     77        85       82
95       92        100        Enabled
CB 0 DDR5 DIMM CPU1_EF                    60    65      80     77        85       82
95       92        100        Enabled
CB 0 DDR5 DIMM CPU1_GH                    60    65      80     77        85       82
95       92        100        Enabled
CB 0 NVME SSD 0                           55    60      70     67        75       72
85       82         95        Enabled
CB 0 NVME SSD 1                           55    60      70     67        75       72
85       82         95        Enabled
CB 0 PCH DTS 0                            65    70      85     82        90       87
93       90         96        Enabled
CB 0 PCH DTS 1                            65    70      85     82        90       87
93       90         96        Enabled
CB 0 VDD12V0_Zone0                        90   100     120    117       125      122
130      127        140        Enabled
CB 0 VDD12V0_Zone1                        90   100     120    117       125      122
130      127        140        Enabled
CB 0 VDD12V0_Zone2                        90   100     120    117       125      122
130      127        140        Enabled
CB 0 VDD12V0_Zone3                        90   100     120    117       125      122
130      127        140        Enabled
CB 0 VDD12V0_Zone4                        90   100     120    117       125      122
130      127        140        Enabled
CB 0 VDD12V0_Zone5                        90   100     120    117       125      122
130      127        140        Enabled
CB 0 VDD12V0_A                            90   100     120    117       125      122
```

```
135      132        140        Enabled
CB 0 VDD12V0_B                            90   100     120    117       125      122
135      132        140        Enabled
CB 0 VDD12V0_C                            90   100     120    117       125      122
135      132        140        Enabled
CB 0 VDD12V0_D                            90   100     120    117       125      122
135      132        140        Enabled
CB 0 P1V05_PCH                            90   100     120    117       125      122
135      132        140        Enabled
CB 0 VDD0V8_E810                          90   100     120    117       125      122
135      132        140        Enabled
CB 0 VCCD_HV_CPU0                         38    43     113    110       125      122
135      132        140        Enabled
CB 0 VCCFA_EHV_CPU0                       90   100     120    117       125      122
130      127        140        Enabled
CB 0 VCCINFAON_CPU0                       90   100     120    117       125      122
130      127        140        Enabled
CB 0 VCCIN_CPU0                           90   100     120    117       125      122
130      127        140        Enabled
CB 0 VCCFA_EHV_FIVRA_CPU0                 90   100     120    117       125      122
130      127        140        Enabled
CB 0 VCCD_HV_CPU1                         90   100     120    117       125      122
130      127        140        Enabled
CB 0 VCCFA_EHV_CPU1                       90   100     120    117       125      122
130      127        140        Enabled
CB 0 VCCFA_EHV_FIVRA_CPU1                 90   100     120    117       125      122
130      127        140        Enabled
CB 0 VCCINFAON_CPU1                       90   100     120    117       125      122
130      127        140        Enabled
CB 0 VCCIN_CPU1                           45    50     120    117       125      122
135      132        140        Enabled
CB 0 MPM54522_VDD3V3                     100   110     120    117       125      122
145      142        150        Enabled
CB 0 MPM3695_VDD1V1_E810                 100   110     120    117       125      122
145      142        150        Enabled
CB 0 E810 DTS 0                           80    85      95     92       105      102
115      112        125        Enabled
CB 0 E810 DTS 1                           80    85      95     92       105      102
115      112        125        Enabled
FPC 1 Exhaust Left                        50    55      60     57        65       62
70       67         75        Enabled
FPC 1 Exhaust Right                       50    55      60     57        65       62
70       67         75        Enabled
```

```
FPC 1 YT Junction                          75    85     100     97       105      102
115      112        120        Disabled
FPC 1 Inlet Middle                         50    55      60     57        65       62
70       67         75        Enabled
FPC 1 Inlet Left                           40    46      52     49        56       53
60       57         65        Enabled
FPC 1 Inlet Right                          40    46      52     49        56       53
60       57         65        Enabled
FPC 1 YT_therm_sensor                      75    85     100     97       105      102
115      112        120        Enabled
FPC 1 YT_test_macro                        75    85     100     97       105      102
115      112        120        Disabled
FPC 1 YT_pgio_slice0                       75    85     100     97       105      102
115      112        120        Disabled
FPC 1 YT_core_pll                          75    85     100     97       105      102
115      112        120        Disabled
FPC 1 YT_fabio_slice0                      75    85     100     97       105      102
115      112        120        Disabled
FPC 1 YT_pgio_slice1                       75    85     100     97       105      102
115      112        120        Disabled
FPC 1 YT_fabio_slice1                      75    85     100     97       105      102
115      112        120        Disabled
FPC 1 YT_hbmio_slice1                      75    85     100     97       105      102
115      112        120        Enabled
FPC 1 YT_wanio_misc                        75    85     100     97       105      102
115      112        120        Disabled
FPC 1 YT_HBM0                              65    75      90     87        95       92
105      102        110        Enabled
FPC 1 YT_HBM1                              65    75      90     87        95       92
105      102        110        Enabled
FPC 1 WAN RT0.0                            75    85     100     97       105      102
110      107        115        Enabled
FPC 1 WAN RT1.0                            75    85     100     97       105      102
110      107        115        Enabled
FPC 1 WAN RT2.0                            75    85     100     97       105      102
110      107        115        Enabled
FPC 1 WAN RT3.0                            75    85     100     97       105      102
110      107        115        Enabled
FPC 1 WAN RT4.0                            75    85     100     97       105      102
110      107        115        Enabled
FPC 1 WAN RT5.0                            75    85     100     97       105      102
110      107        115        Enabled
FPC 1 PIC0-P00_VDD3V3_T                    90   100     120    117       125      122
```

```
130      127        140        Enabled
FPC 1 PIC1-P00_VDD3V3_T                    90   100     120    117       125      122
130      127        140        Enabled
FPC 1 YT_VDD0V75_T                         90   100     120    117       125      122
130      127        140        Enabled
FPC 1 YT_HBM_VDDQ1V2_T                     90   100     120    117       125      122
130      127        140        Enabled
FPC 1 YT_HBM_VDDC1V2_T                     90   100     120    117       125      122
130      127        140        Enabled
FPC 1 YT_RTVDD0V75_T                       90   100     120    117       125      122
130      127        140        Enabled
FPC 1 WAN_VDD0V8_A_T                       90   100     120    117       125      122
130      127        140        Enabled
FPC 1 WAN_AVDD1V2_A_T                      90   100     120    117       125      122
130      127        140        Enabled
FPC 1 WAN_VDD0V8_B_T                       90   100     120    117       125      122
130      127        140        Enabled
FPC 1 WAN_AVDD1V2_B_T                      90   100     120    117       125      122
130      127        140        Enabled
FPC 1 VDD12V0_A_T                         100   110     120    117       125      122
145      142        150        Enabled
FPC 1 VDD12V0_B_T                         100   110     120    117       125      122
145      142        150        Enabled
```

Release Information

Command introduced in Junos OS Release 24.4R1.

show chassis environment

IN THIS SECTION

- Syntax (T320, T640, T1600, and T4000 Routers) | **17430**
- Syntax (TX Matrix Routers) | **17431**
- Syntax (TX Matrix Plus Routers) | **17431**
- Syntax (MX Series Routers) | **17431**

- Syntax (MX104 Universal Routing Platforms) | **17431**
- Syntax (MX150 Router Appliance) | **17432**
- Syntax (MX2010, MX2020, and MX2008 Universal Routing Platforms) | **17432**
- Syntax (MX10003 and MX204 Universal Routing Platforms) | **17432**
- Syntax (EX8200 Switches) | **17432**
- Syntax (EX Series Switches except EX8200) | **17433**
- Syntax (QFX Series) | **17433**
- Syntax (OCX Series) | **17433**
- Syntax (PTX Series Packet Transport Routers) | **17433**
- Syntax (ACX Series Universal Metro Routers, ACX7509) | **17434**
- Syntax (ACX5048 and ACX5096 Routers) | **17434**
- Syntax (ACX500 Routers) | **17434**
- Syntax (Junos OS Evolved) | **17434**
- Description | **17435**
- Options | **17435**
- Required Privilege Level | **17438**
- Output Fields | **17438**
- Sample Output | **17441**
- show chassis environment (ACX7509) | **17546**
- Release Information | **17552**

Syntax (T320, T640, T1600, and T4000 Routers)

```
show chassis environment
<cb cb-slot-number>
<feb feb-slot-number>
<fpc fpc-slot-number>
<fpm>
<pem pem-slot-number>
<routing-engine re-slot-number>
<scg scg-slot-number>
<sib sib-slot-number>
```

Syntax (TX Matrix Routers)

```
show chassis environment
<lcc number   | scc>
```

Syntax (TX Matrix Plus Routers)

```
show chassis environment
<cb cb-slot-number>
<cip cip-slot-number>
<fpc fpc-slot-number>
<fpm>
<lcc number>
<pem pem-slot-number>
<routing-engine re-slot-number>
<scg scg-slot-number>
< sfc number>
<sib sib-slot-number>
```

Syntax (MX Series Routers)

```
show chassis environment
<all-members>
<local>
<member member-id>
```

Syntax (MX104 Universal Routing Platforms)

```
show chassis environment
<cb>
<pem pem-slot-number>
<routing-engine re-slot-number>
```

Syntax (MX150 Router Appliance)

```
show chassis environment
<pem pem-slot-number>
<routing-engine re-slot-number>
```

Syntax (MX2010, MX2020, and MX2008 Universal Routing Platforms)

```
show chassis environment
<adc adc-slot-number>
<all-members>
<cb cb-slot-number>
<fan fantray-slot-number>
<fpc fpc-slot-number>
<fpm>
<local>
<member member-id>
<monitored>
<psm psm-slot-number>
<routing-engine re-slot-number>
<sfb sfb-slot-number>
<satellite [fpc-slot slot-id |device-alias  alias-name]>
```

Syntax (MX10003 and MX204 Universal Routing Platforms)

```
show chassis environment
<cb cb-slot-number>
<fpc fpc-slot-number>
<pem pem-slot-number>
<routing-engine re-slot-number>
```

Syntax (EX8200 Switches)

```
show chassis environment
<all-members>
<cb cb-slot-number>
<fpc fpc-slot-number>
```

```
<local>
<member member-id>
<psu psu-slot-number>
<routing-engine re-slot-number>
```

Syntax (EX Series Switches except EX8200)

```
show chassis environment
<all-members>
<fpc fpc-slot-number>
<local>
<member member-id>
<power-supply-unit>
<routing-engine>
<satellite [fpc-slot slot-id |device-alias  alias-name]>
```

Syntax (QFX Series)

```
show chassis environment
<cb slot-number <interconnect-device name>>
<fpc slot-number <interconnect-device name>>
<interconnect-device name <slot-number>
<node-device name>
<pem slot-number (interconnect-device name slot-number) | (node-device name)>
<routing-engine name <interconnect-device name slot-number>>
<feb feb-slot-number>
```

Syntax (OCX Series)

```
show chassis environment
```

Syntax (PTX Series Packet Transport Routers)

```
show chassis environment
<cb cb-slot-number>
<ccg ccg-slot-number >
```

```
<fpc fpc-slot-number>
<fpm>
<monitored>
<pdu pdu-slot-number>
<routing-engine re-slot-number>
<sib sib-slot-number>
```

Syntax (ACX Series Universal Metro Routers, ACX7509)

```
show chassis environment
<cb cb-slot-number>
<feb feb-slot-number>
<pem pem-slot-number>
<routing-engine re-slot-number>
```

Syntax (ACX5048 and ACX5096 Routers)

```
show chassis environment
<fpc slot-number>
<pem>
<routing-engine>
```

Syntax (ACX500 Routers)

```
show chassis environment
<cb cb-slot-number>
<routing-engine re-slot-number>
```

Syntax (Junos OS Evolved)

```
show chassis environment
<cb cb-slot-number>
<fan fan-slot-number>
<fpc fpc-slot-number>
<fpm>
<monitored>
```

```
<psm psm-slot-number>
<routing-engine routing-engine-slot-number>
<sib sib-slot-number>
```

Description

Display environmental information about the router or switch chassis, including the temperature and information about the fans, power supplies, and Routing Engine.

In addition, on ACX4000 routers, display temperature information about the different channels of a Modular Interface Card (MIC). The number of channels displayed depends on the type of MIC installed.

Starting with Junos OS Release 14.1, the `show chassis environment cb` *`cb-slot-number`* `| ccg` *`ccg-slot-number`* `| fpc` *`fpc-slot-number`* `| fpm | monitored | pdu` *`pdu-slot-number`* `| routing-engine` *`re-slot-number`* `| sib` *`sib-slot-number`* operational mode command output displays environmental information for the new DC power supply module (PSM) and power distribution unit (PDU) that are added to provide power to the high-density FPC (FPC2-PTX-P1A) and other components in a PTX5000 Packet Transport Router.

Options

none — Display environmental information about the router or switch chassis. On a TX Matrix router, display environmental information about the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about the TX Matrix Plus router and its attached routers.

all-members — (MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for all the members of the Virtual Chassis configuration.

adc *adc-slot-number* — (MX2010, MX2020, and MX2008 routers only) (Optional) Display chassis environmental information for the adapter cards. For MX2020 routers, replace *`adc-slot-number`* with a value from `0` through `19`. For MX2010 and MX2008 routers, replace *`adc-slot-number`* with a value from `0` through `9`.

cb *cb-slot-number* — (ACX Series Universal Metro Routers, EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2020 routers, MX2010 routers, MX2008 routers, PTX Series Packet Transport Routers, QFX Series, and T Series routers, and TX Matrix Plus routers only) (Optional) Display chassis environmental information for the Control Board. On devices other than EX Series switches, replace *`cb-slot`* with `0` or `1`.

feb *feb-slot-number*

(ACX Series Universal Metro Routers, ACX7509 routers, and QFX5700 switches only) (Optional) Display chassis environmental information for the Forwarding Engine Board (FEB).

cip *cip-slot-number*

(TX Matrix Plus routers only) (Optional) Display chassis environmental information for the Connection Interface Panel (CIP). Replace the *cip-slot-number* variable with a value of 0 or 1.

cb interconnect-device *name*

(QFabric systems only) (Optional) Display chassis environmental information for the Control Board on an Interconnect device.

ccg *ccg-slot-number*

(PTX Series only) (Optional) Display chassis environmental information for the Centralized Clock Generator. Replace *cb-slot* with a value of 0 or 1.

fan *fantray-slot-number*

(MX2010, MX2020, and MX2008 routers only) (Optional) Display chassis environmental information for the fan trays. Replace *fantray-slot-number* with a value from 0 through 3.

fpc *fpc-slot*

(EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2010 routers, MX2020 routers, MX2008 routers, PTX Series Packet Transport Routers, QFX Series, QFX3500 switches, QFabric systems, T Series routers, and TX Matrix Plus routers) (Optional) Display chassis environmental information for a specified Flexible PIC Concentrator. For MX2010 and MX2008 routers, replace *fpc-slot* with a value from 0 through 9. For MX2020 routers, replace *fpc-slot* with a value from 0 through 19. For information about FPC numbering, see "show chassis environment fpc" on page 17621. On a QFabric system, display chassis environmental information for a specified Flexible PIC Concentrator on an Interconnect device. On an EX Series switch, display chassis environmental information for a specified Flexible PIC Concentrator; see the hardware documentation for your switch for information on FPC numbering. On a TX Matrix Plus router with 3D SIBs replace *fpc-slot* with a value from 0 through 63.

fpm

(M120, M320, and M40e routers, MX2010 routers, MX2020 routers, MX2008 routers, PTX Series, Packet Transport Routers, T Series routers, and TX Matrix Plus routers only) (Optional) Display chassis environmental information for the craft interface (FPM).

interconnect-device *name*

(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.

lcc *number*

(TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for the local Virtual Chassis member.

member *member-id* (MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for the specified member of the Virtual Chassis configuration. On MX Series routers, replace *member-id* with a value of `0` or `1`. For EX Series switches, see "member" on page 10604 for member ID values.

monitored (MX2020 routers and PTX Series Packet Transport Routers only) (Optional) Display chassis environmental information for monitored temperatures only. Temperatures that are not included in temperature alarm computations are not displayed.

node-device *name* (QFabric systems only) (Optional) Display chassis environmental information for the Node device.

pdu *pdu-slot-number* (PTX Series only) (Optional) Display chassis environmental information for the specified power distribution unit.

pem (QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Interconnect device or Node device.

pem *pem-slot-number* (ACX Series Universal Metro Routers, M120, M320, and M40e routers, MX Series routers, MX104 routers, QFX Series, and T Series routers only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Power Entry Module. For information about the options, see "show chassis environment pem" on page 17760.

psm *psm-slot-number* (MX2010, MX2020, and MX2008 routers only) (Optional) Display chassis environmental information for the power supply module. For MX2020 routers, replace *psm-slot-number* with a value from `0` through `17`. For MX2010 and MX2008 routers, replace *psm-slot-number* with a value from `0` through `8`.

psu ***psu-slot-number***	(EX Series switches only) (Optional) Display chassis environmental information for a specified power supply.
routing-engine	(QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Routing Engine on the specified Interconnect device. For QFX5700 with Junos OS Evolved, details about CPU and memory utilization are also displayed for RCB slot 0 and slot 1.
routing-engine ***re-slot-number***	(Optional) Display chassis environmental information for the specified Routing Engine. For information about the options, see "show chassis environment routing-engine" on page 17800.
satellite [fpc-slot ***slot-id*** **| device-alias** ***alias-name*****]**	(Junos Fusion only)(Optional) Display chassis environmental information for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.
scg	(T Series routers only) (Optional) Display chassis environmental information about the SONET Clock Generator.
scc	(TX Matrix routers only) (Optional) Display chassis environmental information about the TX Matrix router (switch-card chassis).
sfb ***sfb-slot-number***	(MX2010, MX2020, and MX2008 routers only) (Optional) Display chassis environmental information for the switch fabric board. Replace *`sfb-slot-number`* with a value from `0` through `7`.
sfc ***number***	(TX Matrix Plus routers only) (Optional) Display chassis environmental information about the respective TX Matrix Plus router (switch-fabric chassis). Replace *`number`* variable with `0`.
sib ***sib-slot-number***	(M320 routers, PTX Series Packet Transport Routers, and T Series routers only) (Optional) Display chassis environmental information about the specified switch interface board. For information about the options, see show chassis environment sib.

Required Privilege Level

view

Output Fields

Table 221 on page 17439 lists the output fields for the `show chassis environment` command. Output fields are listed in the approximate order in which they appear.

Table 221: show chassis environment Output Fields

Field Name	Field Description

`Class`	Information about the category or class of chassis component:
	• `Power`: Power information:
	  • (M5, M10, M20, and M40 routers and EX Series switches only) Power supply status: `OK`, `Testing`, (during initial power-on), `Failed`, or `Absent`.
	  • (M7i, M10i, M40e, M120, M160, M320, and T Series routers and EX Series switches only) Power Entry Modules status: `OK`, `Testing`, (during initial power-on), `Check`, `Failed`, or `Absent`.
	  • (PTX Series only) Power information is reported in PDU or PSM combinations. The status is: `OK`, `Testing`, (during initial power-on), `Check`, `Failed`, or `Absent`.
	• `Temp`: Temperature of air flowing through the chassis in degrees Celsius (C) and Fahrenheit (F).
	  • On PTX Series Packet Transport Routers and MX2010, MX2020, and MX2008 Routers, multiple cooling zones are supported. FRU temperatures in each zone are coordinated with the fan speed of fan trays in those zones.
	  • EX2200 switches have a side-to-rear cooling system. The `Local Intake` temperature is measured by the sensor on the right side of the chassis, and the `Remote Intake` temperature is measured by the sensor on the left side of the chassis.
	• `Pic`: On ACX4000 routers, multiple temperature channels on a MIC. The status is: `OK` and the `Measurement` is in degrees Celsius (C) and Fahrenheit (F).
	• `Fan`: Fan status: `OK`, `Testing` (during initial power-on), `Failed`, or `Absent`.
	  On PTX Series Packet Transport Routers and MX2010, MX2020, and MX2008 Routers, multiple fan trays are supported. Fan status is reported in Fan Tray or Fan combinations. `Measurement` indicates actual fan RPM (PTX and MX2010, MX2020, and MX2008 Routers only).
	• `Misc`: Information about other components of the chassis.
	  • On some routers, this field indicates the status of one or more additional components.

Table 221: show chassis environment Output Fields *(Continued)*

Field Name	Field Description
	• On the M40e, M160, and M320 router, `Misc` includes `CIP` (Connector Interface Panel). `OK` indicates that the CIP is present. `Absent` indicates that the CIP is not present. • On T Series routers, `Misc` includes `CIP` and `SPMB` (Switch Processor Mezzanine Board). `OK` indicates that the `CIP` or `SPMB` is present. `Absent` indicates that the `CIP` or `SPMB` is not present. • On PTX Series Packet Transport Routers, `Misc` includes the `SPMB` (Switch Processor Mezzanine Board). The SPMB is located on the control boards. `OK` indicates that the control board is present. `Absent` indicates that the control board is not present.
`Item`	(MX2010, MX2020, and MX2008 Routers) Information about the chassis component: Routing Engines, Controls Boards (CBs), Switch Fabric Boards (SFBs), PICs, Flexible PIC Concentrators (FPCs), and Adapter Cards (ADCs). (MX104 Routers) Information about the chassis components: Routing Engines, Control Board (CB), Power Entry Module (PEM), and Compact Forwarding Engine Board (AFEB). (QFabric Systems) Information about the chassis component: Control Boards, Routing Engines, Flexible PIC Concentrators (FPCs), and Power Entry Modules (PEMs), Node Devices, and Interconnect Devices. (QFX Series) Information about the chassis component: Flexible PIC Concentrators (FPCs), and Power Entry Modules (PEMs).

Table 221: show chassis environment Output Fields *(Continued)*

Field Name	Field Description
`Status`	Status of the specified chassis component. For example, if the Class is Fan, the fan status can be: • `OK`: The fans are operational. • `Testing`: The fans are being tested during initial power-on. • `Failed`: The fans have failed or the fans are not spinning. • `Absent`: The fan tray is not installed. If the Class is Power, the power supply status can be: • `OK`: The power component is operational. • `Testing`: The power component is being tested during initial power-on. • `Check`: There is insufficient power---that is, fewer than the minimum required feeds are connected. • `Failed`: The inputs leads have failed. • `Absent`: The power component is not installed.
`Measurement`	(MX104, MX2010, MX2020, and MX2008 Routers) Dependant on the Class. For example, if the Class is Temp, indicates the temperature in degree Celsius and degrees Fahrenheit. If the Class is Fan, indicates actual fan RPM.

Sample Output

show chassis environment (M5 Router)

```
user@host> show chassis environment
Class Item                  Status     Measurement
Power Power Supply A        OK
      Power Supply B        Absent
Temp  FPC 0                 OK         30 degrees C / 86 degrees F
      FEB                   OK         33 degrees C / 91 degrees F
      PS Intake             OK         27 degrees C / 80 degrees F
```

```
      PS Exhaust            OK         27 degrees C / 80 degrees F
      Routing Engine        OK         34 degrees C / 93 degrees F
Fans  Left Fan 1            OK         Spinning at normal speed
      Left Fan 2            OK         Spinning at normal speed
      Left Fan 3            OK         Spinning at normal speed
      Left Fan 4            OK         Spinning at normal speed
Misc  Craft Interface       OK
```

show chassis environment (M7i Router)

```
user@host> show chassis environment
Class Item                  Status     Measurement
Power Power Supply 0        OK
      Power Supply 1        Absent
Temp  Intake                OK         22 degrees C / 71 degrees F
      FPC 0                 OK         23 degrees C / 73 degrees F
      Power Supplies        OK         23 degrees C / 73 degrees F
      CFEB Intake           OK         24 degrees C / 75 degrees F
      CFEB Exhaust          OK         29 degrees C / 84 degrees F
      Routing Engine        OK         26 degrees C / 78 degrees F
Fans  Fan 1                 OK         Spinning at normal speed
      Fan 2                 OK         Spinning at normal speed
      Fan 3                 OK         Spinning at normal speed
      Fan 4                 OK         Spinning at normal speed
```

show chassis environment (M10 Router)

```
user@host> show chassis environment
Class Item                  Status     Measurement
Power Power Supply A        OK
      Power Supply B        Failed
Temp  FPC 0                 OK         36 degrees C / 96 degrees F
      FPC 1                 OK         35 degrees C / 95 degrees F
      FEB                   OK         34 degrees C / 93 degrees F
      PS Intake             OK         31 degrees C / 87 degrees F
      PS Exhaust            OK         34 degrees C / 93 degrees F
      Routing Engine        OK         35 degrees C / 95 degrees F
Fans  Left Fan 1            OK         Spinning at normal speed
```

```
      Left Fan 2            OK         Spinning at normal speed
      Left Fan 3            OK         Spinning at normal speed
      Left Fan 4            OK         Spinning at normal speed
Misc  Craft Interface       OK
```

show chassis environment (M10i Router)

```
user@host> show chassis environment
Class Item                  Status     Measurement
Power Power Supply 0        OK
      Power Supply 1        OK
      Power Supply 2        Absent
      Power Supply 3        Absent
Temp  Intake                OK         26 degrees C / 78 degrees F
      FPC 0                 OK         27 degrees C / 80 degrees F
      FPC 1                 OK         28 degrees C / 82 degrees F
      Lower Power Supplies  OK         29 degrees C / 84 degrees F
      Upper Power Supplies  OK         28 degrees C / 82 degrees F
      CFEB Intake           OK         27 degrees C / 80 degrees F
      CFEB Exhaust          OK         36 degrees C / 96 degrees F
      Routing Engine 0      OK         31 degrees C / 87 degrees F
      Routing Engine 1      OK         27 degrees C / 80 degrees F
Fans  Fan Tray 0 Fan 1      OK         Spinning at normal speed
      Fan Tray 0 Fan 2      OK         Spinning at normal speed
      Fan Tray 0 Fan 3      OK         Spinning at normal speed
      Fan Tray 0 Fan 4      OK         Spinning at normal speed
      Fan Tray 0 Fan 5      OK         Spinning at normal speed
      Fan Tray 0 Fan 6      OK         Spinning at normal speed
      Fan Tray 0 Fan 7      OK         Spinning at normal speed
      Fan Tray 0 Fan 8      OK         Spinning at normal speed
      Fan Tray 1 Fan 1      Absent
      Fan Tray 1 Fan 2      Absent
      Fan Tray 1 Fan 3      Absent
      Fan Tray 1 Fan 4      Absent
      Fan Tray 1 Fan 5      Absent
      Fan Tray 1 Fan 6      Absent
      Fan Tray 1 Fan 7      Absent
      Fan Tray 1 Fan 8      Absent
```

show chassis environment (M20 Router)

```
user@host> show chassis environment
Class Item                  Status     Measurement
Power Power Supply A        OK
      Power Supply B        Absent
Temp  FPC 0                 OK         28 degrees C / 82 degrees F
      FPC 1                 OK         27 degrees C / 80 degrees F
      Power Supply A        OK         22 degrees C / 71 degrees F
      Power Supply B        Absent
      SSB 0                 OK         30 degrees C / 86 degrees F
      Backplane             OK         22 degrees C / 71 degrees F
      Routing Engine 0      OK         26 degrees C / 78 degrees F
      Routing Engine 1      Testing
Fans  Rear Fan              OK         Spinning at normal speed
      Front Upper Fan       OK         Spinning at normal speed
      Front Middle Fan      OK         Spinning at normal speed
      Front Bottom Fan      OK         Spinning at normal speed
Misc  Craft Interface       OK
```

show chassis environment (M40 Router)

```
user@host> show chassis environment
Class Item                  Status     Measurement
Power Power Supply A        OK
      Power Supply B        Absent
Temp  FPC 3                 OK         24 degrees C / 75 degrees F
      FPC 6                 OK         26 degrees C / 78 degrees F
      SCB                   OK         26 degrees C / 78 degrees F
      Backplane @ A1        OK         28 degrees C / 82 degrees F
      Backplane @ A2        OK         23 degrees C / 73 degrees F
      Routing Engine        OK         26 degrees C / 78 degrees F
Fans  Top Impeller          OK         Spinning at normal speed
      Bottom impeller       OK         Spinning at normal speed
      Rear Left Fan         OK         Spinning at normal speed
      Rear Center Fan       OK         Spinning at normal speed
      Rear Right Fan        OK         Spinning at normal speed
Misc  Craft Interface       OK
```

show chassis environment (M40e Router)

```
user@host> show chassis environment
Class Item                    Status     Measurement
Power PEM 0                   OK
      PEM 1                   Absent
Temp  PCG 0                   OK         44 degrees C / 111 degrees F
      PCG 1                   OK         47 degrees C / 116 degrees F
      Routing Engine 0        OK         40 degrees C / 104 degrees F
      Routing Engine 1        OK         37 degrees C / 98 degrees F
      MCS 0                   OK         45 degrees C / 113 degrees F
      MCS 1                   OK         42 degrees C / 107 degrees F
      SFM 0 SPP               OK         40 degrees C / 104 degrees F
      SFM 0 SPR               OK         44 degrees C / 111 degrees F
      SFM 1 SPP               OK         43 degrees C / 109 degrees F
      SFM 1 SPR               OK         45 degrees C / 113 degrees F
      FPC 0                   OK         38 degrees C / 100 degrees F
      FPC 1                   OK         40 degrees C / 104 degrees F
      FPC 2                   OK         38 degrees C / 100 degrees F
      FPC 4                   OK         34 degrees C / 93 degrees F
      FPC 5                   OK         43 degrees C / 109 degrees F
      FPC 6                   OK         41 degrees C / 105 degrees F
      FPC 7                   OK         43 degrees C / 109 degrees F
      FPM CMB                 OK         28 degrees C / 82 degrees F
      FPM Display             OK         28 degrees C / 82 degrees F
Fans  Rear Bottom Blower      OK         Spinning at normal speed
      Rear Top Blower         OK         Spinning at normal speed
      Front Top Blower        OK         Spinning at normal speed
      Fan Tray Rear Left      OK         Spinning at normal speed
      Fan Tray Rear Right     OK         Spinning at normal speed
      Fan Tray Front Left     OK         Spinning at normal speed
      Fan Tray Front Right    OK         Spinning at normal speed
Misc  CIP                     OK
```

show chassis environment (M120 Router)

```
user@host> show chassis environment
Class Item                            Status     Measurement
Temp  PEM 0                           OK
```

```
       PEM 1                        OK
       Routing Engine 0             OK         43 degrees C / 109 degrees F
       Routing Engine 1             OK         44 degrees C / 111 degrees F
       CB 0 Intake                  OK         33 degrees C / 91 degrees F
       CB 0 Exhaust A               OK         36 degrees C / 96 degrees F
       CB 0 Exhaust B               OK         35 degrees C / 95 degrees F
       CB 1 Intake                  OK         34 degrees C / 93 degrees F
       CB 1 Exhaust A               OK         38 degrees C / 100 degrees F
       CB 1 Exhaust B               OK         35 degrees C / 95 degrees F
       FEB 3 Intake                 OK         35 degrees C / 95 degrees F
       FEB 3 Exhaust A              OK         37 degrees C / 98 degrees F
       FEB 3 Exhaust B              OK         39 degrees C / 102 degrees F
       FEB 4 Intake                 OK         33 degrees C / 91 degrees F
       FEB 4 Exhaust A              OK         39 degrees C / 102 degrees F
       FEB 4 Exhaust B              OK         36 degrees C / 96 degrees F
       FPC 2 Exhaust A              OK         32 degrees C / 89 degrees F
       FPC 2 Exhaust B              OK         31 degrees C / 87 degrees F
       FPC 3 Exhaust A              OK         32 degrees C / 89 degrees F
       FPC 3 Exhaust B              OK         33 degrees C / 91 degrees F
       FPC 4 Exhaust A              OK         32 degrees C / 89 degrees F
       FPC 4 Exhaust B              OK         30 degrees C / 86 degrees F
Fans   Front Top Tray Fan 1         OK         Spinning at normal speed
       Front Top Tray Fan 2         OK         Spinning at normal speed
       Front Top Tray Fan 3         OK         Spinning at normal speed
       Front Top Tray Fan 4         OK         Spinning at normal speed
       Front Top Tray Fan 5         OK         Spinning at normal speed
       Front Top Tray Fan 6         OK         Spinning at normal speed
       Front Top Tray Fan 7         OK         Spinning at normal speed
       Front Top Tray Fan 8         OK         Spinning at normal speed
       Front Bottom Tray Fan 1      OK         Spinning at normal speed
       Front Bottom Tray Fan 2      OK         Spinning at normal speed
       Front Bottom Tray Fan 3      OK         Spinning at normal speed
       Front Bottom Tray Fan 4      OK         Spinning at normal speed
       Front Bottom Tray Fan 5      OK         Spinning at normal speed
       Front Bottom Tray Fan 6      OK         Spinning at normal speed
       Front Bottom Tray Fan 7      OK         Spinning at normal speed
       Front Bottom Tray Fan 8      OK         Spinning at normal speed
       Rear Top Tray Fan 1          OK         Spinning at normal speed
       Rear Top Tray Fan 2          OK         Spinning at normal speed
       Rear Top Tray Fan 3          OK         Spinning at normal speed
       Rear Top Tray Fan 4          OK         Spinning at normal speed
       Rear Top Tray Fan 5          OK         Spinning at normal speed
       Rear Top Tray Fan 6          OK         Spinning at normal speed
```

```
      Rear Top Tray Fan 7            OK         Spinning at normal speed
      Rear Top Tray Fan 8            OK         Spinning at normal speed
      Rear Bottom Tray Fan 1         OK         Spinning at normal speed
      Rear Bottom Tray Fan 2         OK         Spinning at normal speed
      Rear Bottom Tray Fan 3         OK         Spinning at normal speed
      Rear Bottom Tray Fan 4         OK         Spinning at normal speed
      Rear Bottom Tray Fan 5         OK         Spinning at normal speed
      Rear Bottom Tray Fan 6         OK         Spinning at normal speed
      Rear Bottom Tray Fan 7         OK         Spinning at normal speed
      Rear Bottom Tray Fan 8         OK         Spinning at normal speed
```

show chassis environment (M160 Router)

```
user@host> show chassis environment
Class Item                    Status     Measurement
Power PEM 0                   OK           PEM 1                  Absent
Temp  PCG 0                   OK         45 degrees C / 113 degrees F
      PCG 1                   Absent
      Routing Engine 0        OK         35 degrees C / 95 degrees F
      Routing Engine 1        Absent
      MCS 0                   OK         50 degrees C / 122 degrees F
      SFM 0 SPP               OK         47 degrees C / 116 degrees F
      SFM 0 SPR               OK         49 degrees C / 120 degrees F
      SFM 1 SPP               OK         50 degrees C / 122 degrees F
      SFM 1 SPR               OK         50 degrees C / 122 degrees F
      SFM 2 SPP               OK         51 degrees C / 123 degrees F
      SFM 2 SPR               OK         52 degrees C / 125 degrees F
      SFM 3 SPP               OK         52 degrees C / 125 degrees F
      SFM 3 SPR               OK         48 degrees C / 118 degrees F
      FPC 0                   OK         45 degrees C / 113 degrees F
      FPC 6                   OK         43 degrees C / 109 degrees F
      FPM CMB                 OK         31 degrees C / 87 degrees F
      FPM Display             OK         33 degrees C / 91 degrees F
Fans  Rear Bottom Blower      OK         Spinning at normal speed
      Rear Top Blower         OK         Spinning at normal speed
      Front Top Blower        OK         Spinning at normal speed
      Fan Tray Rear Left      OK         Spinning at normal speed
      Fan Tray Rear Right     OK         Spinning at normal speed
      Fan Tray Front Left     OK         Spinning at normal speed
      Fan Tray Front Right    OK         Spinning at normal speed
```

```
Misc  CIP                   OK
```

show chassis environment (M320 Router)

```
user@host> show chassis environment
Class Item                  Status     Measurement
Temp  PEM 0                 Absent
      PEM 1                 Absent
      PEM 2                 OK
      PEM 3                 OK
      Routing Engine 0      OK         33 degrees C / 91 degrees F
      Routing Engine 1      OK         32 degrees C / 89 degrees F
      CB 0                  OK         36 degrees C / 96 degrees F
      CB 1                  OK         36 degrees C / 96 degrees F
      SIB 0                 OK         38 degrees C / 100 degrees F
      SIB 1                 OK         29 degrees C / 84 degrees F
      SIB 2                 OK         38 degrees C / 100 degrees F
      SIB 3                 OK         41 degrees C / 105 degrees F
      FPC 0 Intake          OK         28 degrees C / 82 degrees F
      FPC 0 Exhaust         OK         40 degrees C / 104 degrees F
      FPC 1 Intake          OK         29 degrees C / 84 degrees F
      FPC 1 Exhaust         OK         39 degrees C / 102 degrees F
      FPC 2 Intake          OK         28 degrees C / 82 degrees F
      FPC 2 Exhaust         OK         38 degrees C / 100 degrees F
      FPC 3 Intake          OK         28 degrees C / 82 degrees F
      FPC 3 Exhaust         OK         39 degrees C / 102 degrees F
      FPC 6 Intake          OK         27 degrees C / 80 degrees F
      FPC 6 Exhaust         OK         39 degrees C / 102 degrees F
      FPC 7 Intake          OK         27 degrees C / 80 degrees F
      FPC 7 Exhaust         OK         42 degrees C / 107 degrees F
      FPM GBUS              OK         30 degrees C / 86 degrees F
Fan   Top Left Front fan    OK         Spinning at normal speed
      Top Right Rear fan    OK         Spinning at normal speed
      Top Right Front fan   OK         Spinning at normal speed
      Top Left Rear fan     OK         Spinning at normal speed
      Bottom Left Front fan OK         Spinning at normal speed
      Bottom Right Rear fan OK         Spinning at normal speed
      Bottom Right Front fan OK        Spinning at normal speed
      Bottom Left Rear fan  OK         Spinning at normal speed
      Rear Fan 1 (TOP)      OK         Spinning at normal speed
```

```
      Rear Fan 2             OK         Spinning at normal speed
      Rear Fan 3             OK         Spinning at normal speed
      Rear Fan 4             OK         Spinning at normal speed
      Rear Fan 5             OK         Spinning at normal speed
      Rear Fan 6             OK         Spinning at normal speed
      Rear Fan 7 (Bottom)    OK         Spinning at normal speed
Misc  CIP                    OK
```

show chassis environment (MX150)

```
user@host> show chassis environment
Class Item                           Status     Measurement
Power FPC 0 Power Supply 0           OK
Temp  FPC 0 Sensor 1                 OK         42 degrees C / 107 degrees F
      FPC 0 Sensor 2                 OK         39 degrees C / 102 degrees F
      FPC 0 Coretemp                 OK         75 degrees C / 167 degrees F
Fans  FPC 0 Fan Tray 0               OK         Spinning at normal speed
      FPC 0 Fan Tray 1               OK         Spinning at normal speed
```

show chassis environment (MX104 Router)

```
user@host> show chassis environment
Class Item                           Status     Measurement
Temp  PEM 0                          OK         34 degrees C / 93 degrees F
      PEM 1                          Absent
      ABB 0 Intake                   OK         33 degrees C / 91 degrees F
      ABB 0 Exhaust A                OK         42 degrees C / 107 degrees F
      ABB 0 Exhaust B                OK         43 degrees C / 109 degrees F
      ABB 1 Intake                   Absent
      ABB 1 Exhaust A                Absent
      ABB 1 Exhaust B                Absent
      Routing Engine 0               OK         34 degrees C / 93 degrees F
      Routing Engine 0 CPU           OK         46 degrees C / 114 degrees F
      Routing Engine 1               Absent
      Routing Engine 1 CPU           Absent
      AFEB 0 AFEB Processor          OK         33 degrees C / 91 degrees F
Fans  Fan 1                          OK         Spinning at normal speed
      Fan 2                          OK         Spinning at normal speed
      Fan 3                          OK         Spinning at normal speed
```

```
      Fan 4                          OK         Spinning at normal speed
      Fan 5                          OK         Spinning at normal speed
```

show chassis environment (MX240 Router)

```
user@host> show chassis environment
Class Item                           Status     Measurement
Temp  PEM 0                          OK         40 degrees C / 104 degrees F
      PEM 1                          OK         45 degrees C / 113 degrees F
      PEM 2                          Absent
      PEM 3                          Absent
      Routing Engine 0               OK         39 degrees C / 102 degrees F
      Routing Engine 1               OK         37 degrees C / 98 degrees F
      CB 0 Intake                    OK         36 degrees C / 96 degrees F
      CB 0 Exhaust A                 OK         34 degrees C / 93 degrees F
      CB 0 Exhaust B                 OK         38 degrees C / 100 degrees F
      CB 0 ACBC                      OK         37 degrees C / 98 degrees F
      CB 0 SF A                      OK         49 degrees C / 120 degrees F
      CB 0 SF B                      OK         41 degrees C / 105 degrees F
      CB 1 Intake                    OK         37 degrees C / 98 degrees F
      CB 1 Exhaust A                 OK         34 degrees C / 93 degrees F
      CB 1 Exhaust B                 OK         39 degrees C / 102 degrees F
      CB 1 ACBC                      OK         38 degrees C / 100 degrees F
      CB 1 SF A                      OK         47 degrees C / 116 degrees F
      CB 1 SF B                      OK         41 degrees C / 105 degrees F
      FPC 1 Intake                   OK         33 degrees C / 91 degrees F
      FPC 1 Exhaust A                OK         38 degrees C / 100 degrees F
      FPC 1 Exhaust B                OK         53 degrees C / 127 degrees F
      FPC 1 I3 0 TSensor             OK         50 degrees C / 122 degrees F
      FPC 1 I3 0 Chip                OK         53 degrees C / 127 degrees F
      FPC 1 I3 1 TSensor             OK         49 degrees C / 120 degrees F
      FPC 1 I3 1 Chip                OK         52 degrees C / 125 degrees F
      FPC 1 I3 2 TSensor             OK         47 degrees C / 116 degrees F
      FPC 1 I3 2 Chip                OK         49 degrees C / 120 degrees F
      FPC 1 I3 3 TSensor             OK         44 degrees C / 111 degrees F
      FPC 1 I3 3 Chip                OK         46 degrees C / 114 degrees F
      FPC 1 IA 0 TSensor             OK         45 degrees C / 113 degrees F
      FPC 1 IA 0 Chip                OK         44 degrees C / 111 degrees F
      FPC 1 IA 1 TSensor             OK         44 degrees C / 111 degrees F
      FPC 1 IA 1 Chip                OK         48 degrees C / 118 degrees F
      FPC 2 Intake                   OK         32 degrees C / 89 degrees F
```

```
      FPC 2 Exhaust A                OK         40 degrees C / 104 degrees F
      FPC 2 Exhaust B                OK         52 degrees C / 125 degrees F
      FPC 2 I3 0 TSensor             OK         52 degrees C / 125 degrees F
      FPC 2 I3 0 Chip                OK         56 degrees C / 132 degrees F
      FPC 2 I3 1 TSensor             OK         52 degrees C / 125 degrees F
      FPC 2 I3 1 Chip                OK         55 degrees C / 131 degrees F
      FPC 2 I3 2 TSensor             OK         49 degrees C / 120 degrees F
      FPC 2 I3 2 Chip                OK         52 degrees C / 125 degrees F
      FPC 2 I3 3 TSensor             OK         44 degrees C / 111 degrees F
      FPC 2 I3 3 Chip                OK         48 degrees C / 118 degrees F
      FPC 2 IA 0 TSensor             OK         50 degrees C / 122 degrees F
      FPC 2 IA 0 Chip                OK         48 degrees C / 118 degrees F
      FPC 2 IA 1 TSensor             OK         47 degrees C / 116 degrees F
      FPC 2 IA 1 Chip                OK         53 degrees C / 127 degrees F
Fans  Front Fan                      OK         Spinning at normal speed
      Middle Fan                     OK         Spinning at normal speed
      Rear Fan                       OK         Spinning at normal speed
```

show chassis environment (MX240 Router with SCBE)

```
user@host> show chassis environment
Class Item                           Status     Measurement
Temp  PEM 0                          OK         40 degrees C / 104 degrees F
      PEM 1                          OK         45 degrees C / 113 degrees F
      PEM 2                          Absent
      PEM 3                          Absent
      Routing Engine 0               OK         39 degrees C / 102 degrees F
      Routing Engine 1               OK         37 degrees C / 98 degrees F
      CB 0 Intake                    OK         36 degrees C / 96 degrees F
      CB 0 Exhaust A                 OK         34 degrees C / 93 degrees F
      CB 0 Exhaust B                 OK         38 degrees C / 100 degrees F
      CB 0 ACBC                      OK         37 degrees C / 98 degrees F
      CB 0 XF A                      OK         49 degrees C / 120 degrees F
      CB 0 XF B                      OK         41 degrees C / 105 degrees F
      CB 1 Intake                    OK         37 degrees C / 98 degrees F
      CB 1 Exhaust A                 OK         34 degrees C / 93 degrees F
      CB 1 Exhaust B                 OK         39 degrees C / 102 degrees F
      CB 1 ACBC                      OK         38 degrees C / 100 degrees F
      CB 1 XF A                      OK         47 degrees C / 116 degrees F
      CB 1 XF B                      OK         41 degrees C / 105 degrees F
      FPC 1 Intake                   OK         33 degrees C / 91 degrees F
```

```
      FPC 1 Exhaust A                OK         38 degrees C / 100 degrees F
      FPC 1 Exhaust B                OK         53 degrees C / 127 degrees F
      FPC 1 I3 0 TSensor             OK         50 degrees C / 122 degrees F
      FPC 1 I3 0 Chip                OK         53 degrees C / 127 degrees F
      FPC 1 I3 1 TSensor             OK         49 degrees C / 120 degrees F
      FPC 1 I3 1 Chip                OK         52 degrees C / 125 degrees F
      FPC 1 I3 2 TSensor             OK         47 degrees C / 116 degrees F
      FPC 1 I3 2 Chip                OK         49 degrees C / 120 degrees F
      FPC 1 I3 3 TSensor             OK         44 degrees C / 111 degrees F
      FPC 1 I3 3 Chip                OK         46 degrees C / 114 degrees F
      FPC 1 IA 0 TSensor             OK         45 degrees C / 113 degrees F
      FPC 1 IA 0 Chip                OK         44 degrees C / 111 degrees F
      FPC 1 IA 1 TSensor             OK         44 degrees C / 111 degrees F
      FPC 1 IA 1 Chip                OK         48 degrees C / 118 degrees F
      FPC 2 Intake                   OK         32 degrees C / 89 degrees F
      FPC 2 Exhaust A                OK         40 degrees C / 104 degrees F
      FPC 2 Exhaust B                OK         52 degrees C / 125 degrees F
      FPC 2 I3 0 TSensor             OK         52 degrees C / 125 degrees F
      FPC 2 I3 0 Chip                OK         56 degrees C / 132 degrees F
      FPC 2 I3 1 TSensor             OK         52 degrees C / 125 degrees F
      FPC 2 I3 1 Chip                OK         55 degrees C / 131 degrees F
      FPC 2 I3 2 TSensor             OK         49 degrees C / 120 degrees F
      FPC 2 I3 2 Chip                OK         52 degrees C / 125 degrees F
      FPC 2 I3 3 TSensor             OK         44 degrees C / 111 degrees F
      FPC 2 I3 3 Chip                OK         48 degrees C / 118 degrees F
      FPC 2 IA 0 TSensor             OK         50 degrees C / 122 degrees F
      FPC 2 IA 0 Chip                OK         48 degrees C / 118 degrees F
      FPC 2 IA 1 TSensor             OK         47 degrees C / 116 degrees F
      FPC 2 IA 1 Chip                OK         53 degrees C / 127 degrees F
Fans  Front Fan                      OK         Spinning at normal speed
      Middle Fan                     OK         Spinning at normal speed
      Rear Fan                       OK         Spinning at normal speed
```

show chassis environment (MX480 Router)

```
user@host> show chassis environment
Class Item                           Status     Measurement
Temp  PEM 0                          OK         35 degrees C / 95 degrees F
      PEM 1                          OK         40 degrees C / 104 degrees F
      PEM 2                          Absent
      PEM 3                          Absent
```

```
    Routing Engine 0              OK        44 degrees C / 111 degrees F
    Routing Engine 1              OK        45 degrees C / 113 degrees F
    CB 0 Intake                   OK        36 degrees C / 96 degrees F
    CB 0 Exhaust A                OK        38 degrees C / 100 degrees F
    CB 0 Exhaust B                OK        39 degrees C / 102 degrees F
    CB 0 ACBC                     OK        37 degrees C / 98 degrees F
    CB 0 SF A                     OK        51 degrees C / 123 degrees F
    CB 0 SF B                     OK        44 degrees C / 111 degrees F
    CB 1 Intake                   OK        36 degrees C / 96 degrees F
    CB 1 Exhaust A                OK        39 degrees C / 102 degrees F
    CB 1 Exhaust B                OK        40 degrees C / 104 degrees F
    CB 1 ACBC                     OK        37 degrees C / 98 degrees F
    CB 1 SF A                     OK        50 degrees C / 122 degrees F
    CB 1 SF B                     OK        43 degrees C / 109 degrees F
    FPC 0 Intake                  OK        36 degrees C / 96 degrees F
    FPC 0 Exhaust A               OK        39 degrees C / 102 degrees F
    FPC 0 Exhaust B               OK        51 degrees C / 123 degrees F
    FPC 0 I3 0 TSensor            OK        49 degrees C / 120 degrees F
    FPC 0 I3 0 Chip               OK        56 degrees C / 132 degrees F
    FPC 0 I3 1 TSensor            OK        47 degrees C / 116 degrees F
    FPC 0 I3 1 Chip               OK        52 degrees C / 125 degrees F
    FPC 0 I3 2 TSensor            OK        46 degrees C / 114 degrees F
    FPC 0 I3 2 Chip               OK        48 degrees C / 118 degrees F
    FPC 0 I3 3 TSensor            OK        42 degrees C / 107 degrees F
    FPC 0 I3 3 Chip               OK        45 degrees C / 113 degrees F
    FPC 0 IA 0 TSensor            OK        45 degrees C / 113 degrees F
    FPC 0 IA 0 Chip               OK        45 degrees C / 113 degrees F
    FPC 0 IA 1 TSensor            OK        44 degrees C / 111 degrees F
    FPC 0 IA 1 Chip               OK        48 degrees C / 118 degrees F
    FPC 1 Intake                  OK        37 degrees C / 98 degrees F
    FPC 1 Exhaust A               OK        41 degrees C / 105 degrees F
    FPC 1 Exhaust B               OK        52 degrees C / 125 degrees F
    FPC 1 I3 0 TSensor            OK        51 degrees C / 123 degrees F
    FPC 1 I3 0 Chip               OK        57 degrees C / 134 degrees F
    FPC 1 I3 1 TSensor            OK        48 degrees C / 118 degrees F
    FPC 1 I3 1 Chip               OK        52 degrees C / 125 degrees F
    FPC 1 I3 2 TSensor            OK        46 degrees C / 114 degrees F
    FPC 1 I3 2 Chip               OK        50 degrees C / 122 degrees F
    FPC 1 I3 3 TSensor            OK        42 degrees C / 107 degrees F
    FPC 1 I3 3 Chip               OK        46 degrees C / 114 degrees F
    FPC 1 IA 0 TSensor            OK        49 degrees C / 120 degrees F
    FPC 1 IA 0 Chip               OK        48 degrees C / 118 degrees F
    FPC 1 IA 1 TSensor            OK        46 degrees C / 114 degrees F
```

```
      FPC 1 IA 1 Chip                OK         50 degrees C / 122 degrees F
Fans  Top Rear Fan                   OK         Spinning at normal speed
      Bottom Rear Fan                OK         Spinning at normal speed
      Top Middle Fan                 OK         Spinning at normal speed
      Bottom Middle Fan              OK         Spinning at normal speed
      Top Front Fan                  OK         Spinning at normal speed
      Bottom Front Fan               OK         Spinning at normal speed
```

show chassis environment (MX480 Router with SCBE)

```
user@host> show chassis environment
Class Item                           Status     Measurement
Temp  PEM 0                          OK         35 degrees C / 95 degrees F
      PEM 1                          OK         40 degrees C / 104 degrees F
      PEM 2                          Absent
      PEM 3                          Absent
      Routing Engine 0               OK         44 degrees C / 111 degrees F
      Routing Engine 1               OK         45 degrees C / 113 degrees F
      CB 0 Intake                    OK         36 degrees C / 96 degrees F
      CB 0 Exhaust A                 OK         38 degrees C / 100 degrees F
      CB 0 Exhaust B                 OK         39 degrees C / 102 degrees F
      CB 0 ACBC                      OK         37 degrees C / 98 degrees F
      CB 0 XF A                      OK         51 degrees C / 123 degrees F
      CB 0 XF B                      OK         44 degrees C / 111 degrees F
      CB 1 Intake                    OK         36 degrees C / 96 degrees F
      CB 1 Exhaust A                 OK         39 degrees C / 102 degrees F
      CB 1 Exhaust B                 OK         40 degrees C / 104 degrees F
      CB 1 ACBC                      OK         37 degrees C / 98 degrees F
      CB 1 XF A                      OK         50 degrees C / 122 degrees F
      CB 1 XF B                      OK         43 degrees C / 109 degrees F
      FPC 0 Intake                   OK         36 degrees C / 96 degrees F
      FPC 0 Exhaust A                OK         39 degrees C / 102 degrees F
      FPC 0 Exhaust B                OK         51 degrees C / 123 degrees F
      FPC 0 I3 0 TSensor             OK         49 degrees C / 120 degrees F
      FPC 0 I3 0 Chip                OK         56 degrees C / 132 degrees F
      FPC 0 I3 1 TSensor             OK         47 degrees C / 116 degrees F
      FPC 0 I3 1 Chip                OK         52 degrees C / 125 degrees F
      FPC 0 I3 2 TSensor             OK         46 degrees C / 114 degrees F
      FPC 0 I3 2 Chip                OK         48 degrees C / 118 degrees F
      FPC 0 I3 3 TSensor             OK         42 degrees C / 107 degrees F
      FPC 0 I3 3 Chip                OK         45 degrees C / 113 degrees F
```

```
      FPC 0 IA 0 TSensor            OK         45 degrees C / 113 degrees F
      FPC 0 IA 0 Chip               OK         45 degrees C / 113 degrees F
      FPC 0 IA 1 TSensor            OK         44 degrees C / 111 degrees F
      FPC 0 IA 1 Chip               OK         48 degrees C / 118 degrees F
      FPC 1 Intake                  OK         37 degrees C / 98 degrees F
      FPC 1 Exhaust A               OK         41 degrees C / 105 degrees F
      FPC 1 Exhaust B               OK         52 degrees C / 125 degrees F
      FPC 1 I3 0 TSensor            OK         51 degrees C / 123 degrees F
      FPC 1 I3 0 Chip               OK         57 degrees C / 134 degrees F
      FPC 1 I3 1 TSensor            OK         48 degrees C / 118 degrees F
      FPC 1 I3 1 Chip               OK         52 degrees C / 125 degrees F
      FPC 1 I3 2 TSensor            OK         46 degrees C / 114 degrees F
      FPC 1 I3 2 Chip               OK         50 degrees C / 122 degrees F
      FPC 1 I3 3 TSensor            OK         42 degrees C / 107 degrees F
      FPC 1 I3 3 Chip               OK         46 degrees C / 114 degrees F
      FPC 1 IA 0 TSensor            OK         49 degrees C / 120 degrees F
      FPC 1 IA 0 Chip               OK         48 degrees C / 118 degrees F
      FPC 1 IA 1 TSensor            OK         46 degrees C / 114 degrees F
      FPC 1 IA 1 Chip               OK         50 degrees C / 122 degrees F
Fans  Top Rear Fan                  OK         Spinning at normal speed
      Bottom Rear Fan               OK         Spinning at normal speed
      Top Middle Fan                OK         Spinning at normal speed
      Bottom Middle Fan             OK         Spinning at normal speed
      Top Front Fan                 OK         Spinning at normal speed
      Bottom Front Fan              OK         Spinning at normal speed
```

show chassis environment (MX960 Router)

```
user@host> show chassis environment
Class Item                          Status     Measurement
Temp  PEM 0                         Absent
      PEM 1                         Absent
      PEM 2                         Check
      PEM 3                         OK         35 degrees C / 95 degrees F
      Routing Engine 0              OK         37 degrees C / 98 degrees F
      Routing Engine 1              Absent
      CB 0 Intake                   OK         24 degrees C / 75 degrees F
      CB 0 Exhaust A                OK         30 degrees C / 86 degrees F
      CB 0 Exhaust B                OK         27 degrees C / 80 degrees F
      CB 1 Intake                   Absent
      CB 1 Exhaust A                Absent
```

```
      CB 1 Exhaust B                Absent
      CB 1 ACBC                     Absent
      CB 1 SF A                     Absent
      CB 1 SF B                     Absent
      CB 2 Intake                   Absent
      CB 2 Exhaust A                Absent
      CB 2 Exhaust B                Absent
      CB 2 ACBC                     Absent
      CB 2 SF A                     Absent
      CB 2 SF B                     Absent
      FPC 4 Intake                  OK         24 degrees C / 75 degrees F
      FPC 4 Exhaust A               OK         36 degrees C / 96 degrees F
      FPC 4 Exhaust B               OK         38 degrees C / 100 degrees F
      FPC 7 Intake                  OK         24 degrees C / 75 degrees F
      FPC 7 Exhaust A               OK         36 degrees C / 96 degrees F
      FPC 7 Exhaust B               OK         42 degrees C / 107 degrees F
Fans  Top Fan Tray Temp             Failed
      Top Tray Fan 1                OK         Spinning at normal speed
      Top Tray Fan 2                OK         Spinning at normal speed
      Top Tray Fan 3                OK         Spinning at normal speed
      Top Tray Fan 4                OK         Spinning at normal speed
      Top Tray Fan 5                OK         Spinning at normal speed
      Top Tray Fan 6                OK         Spinning at normal speed
      Bottom Fan Tray Temp          Failed
      Bottom Tray Fan 1             OK         Spinning at normal speed
      Bottom Tray Fan 2             OK         Spinning at normal speed
      Bottom Tray Fan 3             OK         Spinning at normal speed
      Bottom Tray Fan 4             OK         Spinning at normal speed
      Bottom Tray Fan 5             OK         Spinning at normal speed
      Bottom Tray Fan 6             OK         Spinning at normal speed
```

show chassis environment (MX960 Router with SCBE)

```
user@host> show chassis environment
Class Item                          Status     Measurement
Temp  PEM 0                         Absent
      PEM 1                         OK         50 degrees C / 122 degrees F
      PEM 2                         OK         50 degrees C / 122 degrees F
      PEM 3                         OK         50 degrees C / 122 degrees F
      Routing Engine 0              OK         42 degrees C / 107 degrees F
      Routing Engine 0 CPU          OK         51 degrees C / 123 degrees F
```

```
    Routing Engine 1              OK        39 degrees C / 102 degrees F
    Routing Engine 1 CPU          OK        44 degrees C / 111 degrees F
    CB 0 Intake                   OK        35 degrees C / 95 degrees F
    CB 0 Exhaust A                OK        36 degrees C / 96 degrees F
    CB 0 Exhaust B                OK        43 degrees C / 109 degrees F
    CB 0 ACBC                     OK        38 degrees C / 100 degrees F
    CB 0 XF A                     OK        53 degrees C / 127 degrees F
    CB 0 XF B                     OK        47 degrees C / 116 degrees F
    CB 1 Intake                   OK        35 degrees C / 95 degrees F
    CB 1 Exhaust A                OK        35 degrees C / 95 degrees F
    CB 1 Exhaust B                OK        41 degrees C / 105 degrees F
    CB 1 ACBC                     OK        38 degrees C / 100 degrees F
    CB 1 XF A                     OK        52 degrees C / 125 degrees F
    CB 1 XF B                     OK        47 degrees C / 116 degrees F
    CB 2 Intake                   OK        32 degrees C / 89 degrees F
    CB 2 Exhaust A                OK        30 degrees C / 86 degrees F
    CB 2 Exhaust B                OK        35 degrees C / 95 degrees F
    CB 2 ACBC                     OK        33 degrees C / 91 degrees F
    CB 2 XF A                     OK        51 degrees C / 123 degrees F
    CB 2 XF B                     OK        50 degrees C / 122 degrees F
    FPC 0 Intake                  OK        35 degrees C / 95 degrees F
    FPC 0 Exhaust A               OK        39 degrees C / 102 degrees F
    FPC 0 Exhaust B               OK        50 degrees C / 122 degrees F
    FPC 0 I3 0 TSensor            OK        50 degrees C / 122 degrees F
    FPC 0 I3 0 Chip               OK        56 degrees C / 132 degrees F
    FPC 0 I3 1 TSensor            OK        47 degrees C / 116 degrees F
    FPC 0 I3 1 Chip               OK        50 degrees C / 122 degrees F
    FPC 0 I3 2 TSensor            OK        45 degrees C / 113 degrees F
    FPC 0 I3 2 Chip               OK        48 degrees C / 118 degrees F
    FPC 0 I3 3 TSensor            OK        41 degrees C / 105 degrees F
    FPC 0 I3 3 Chip               OK        44 degrees C / 111 degrees F
    FPC 0 IA 0 TSensor            OK        45 degrees C / 113 degrees F
    FPC 0 IA 0 Chip               OK        45 degrees C / 113 degrees F
    FPC 0 IA 1 TSensor            OK        44 degrees C / 111 degrees F
    FPC 0 IA 1 Chip               OK        48 degrees C / 118 degrees F
    FPC 1 Intake                  OK        36 degrees C / 96 degrees F
    FPC 1 Exhaust A               OK        47 degrees C / 116 degrees F
    FPC 1 Exhaust B               OK        43 degrees C / 109 degrees F
    FPC 1 LU 0 TCAM TSensor       OK        53 degrees C / 127 degrees F
    FPC 1 LU 0 TCAM Chip          OK        57 degrees C / 134 degrees F
    FPC 1 LU 0 TSensor            OK        53 degrees C / 127 degrees F
    FPC 1 LU 0 Chip               OK        60 degrees C / 140 degrees F
    FPC 1 MQ 0 TSensor            OK        53 degrees C / 127 degrees F
```

```
    FPC 1 MQ 0 Chip              OK        56 degrees C / 132 degrees F
    FPC 1 LU 1 TCAM TSensor      OK        51 degrees C / 123 degrees F
    FPC 1 LU 1 TCAM Chip         OK        52 degrees C / 125 degrees F
    FPC 1 LU 1 TSensor           OK        51 degrees C / 123 degrees F
    FPC 1 LU 1 Chip              OK        53 degrees C / 127 degrees F
    FPC 1 MQ 1 TSensor           OK        51 degrees C / 123 degrees F
    FPC 1 MQ 1 Chip              OK        58 degrees C / 136 degrees F
    FPC 2 Intake                 OK        35 degrees C / 95 degrees F
    FPC 2 Exhaust A              OK        39 degrees C / 102 degrees F
    FPC 2 Exhaust B              OK        54 degrees C / 129 degrees F
    FPC 2 I3 0 TSensor           OK        52 degrees C / 125 degrees F
    FPC 2 I3 0 Chip              OK        59 degrees C / 138 degrees F
    FPC 2 I3 1 TSensor           OK        48 degrees C / 118 degrees F
    FPC 2 I3 1 Chip              OK        52 degrees C / 125 degrees F
    FPC 2 I3 2 TSensor           OK        47 degrees C / 116 degrees F
    FPC 2 I3 2 Chip              OK        49 degrees C / 120 degrees F
    FPC 2 I3 3 TSensor           OK        41 degrees C / 105 degrees F
    FPC 2 I3 3 Chip              OK        44 degrees C / 111 degrees F
    FPC 2 IA 0 TSensor           OK        47 degrees C / 116 degrees F
    FPC 2 IA 0 Chip              OK        46 degrees C / 114 degrees F
    FPC 2 IA 1 TSensor           OK        45 degrees C / 113 degrees F
    FPC 2 IA 1 Chip              OK        49 degrees C / 120 degrees F
    FPC 3 Intake                 OK        34 degrees C / 93 degrees F
    FPC 3 Exhaust A              OK        34 degrees C / 93 degrees F
    FPC 3 Exhaust B              OK        47 degrees C / 116 degrees F
    FPC 3 I3 0 TSensor           OK        48 degrees C / 118 degrees F
    FPC 3 I3 0 Chip              OK        52 degrees C / 125 degrees F
    FPC 3 I3 1 TSensor           OK        46 degrees C / 114 degrees F
    FPC 3 I3 1 Chip              OK        48 degrees C / 118 degrees F
    FPC 3 IA 0 TSensor           OK        41 degrees C / 105 degrees F
    FPC 3 IA 0 Chip              OK        40 degrees C / 104 degrees F
    FPC 5 Intake                 OK        42 degrees C / 107 degrees F
    FPC 5 Exhaust A              OK        42 degrees C / 107 degrees F
    FPC 5 Exhaust B              OK        53 degrees C / 127 degrees F
    FPC 5 LU 0 TSensor           OK        53 degrees C / 127 degrees F
    FPC 5 LU 0 Chip              OK        54 degrees C / 129 degrees F
    FPC 5 LU 1 TSensor           OK        53 degrees C / 127 degrees F
    FPC 5 LU 1 Chip              OK        61 degrees C / 141 degrees F
    FPC 5 LU 2 TSensor           OK        53 degrees C / 127 degrees F
    FPC 5 LU 2 Chip              OK        51 degrees C / 123 degrees F
    FPC 5 LU 3 TSensor           OK        53 degrees C / 127 degrees F
    FPC 5 LU 3 Chip              OK        53 degrees C / 127 degrees F
    FPC 5 MQ 0 TSensor           OK        47 degrees C / 116 degrees F
```

```
       FPC 5 MQ 0 Chip                OK          52 degrees C / 125 degrees F
       FPC 5 MQ 1 TSensor             OK          47 degrees C / 116 degrees F
       FPC 5 MQ 1 Chip                OK          52 degrees C / 125 degrees F
       FPC 5 MQ 2 TSensor             OK          47 degrees C / 116 degrees F
       FPC 5 MQ 2 Chip                OK          46 degrees C / 114 degrees F
       FPC 5 MQ 3 TSensor             OK          47 degrees C / 116 degrees F
       FPC 5 MQ 3 Chip                OK          45 degrees C / 113 degrees F
       FPC 7 Intake                   OK          36 degrees C / 96 degrees F
       FPC 7 Exhaust A                OK          35 degrees C / 95 degrees F
       FPC 7 Exhaust B                OK          33 degrees C / 91 degrees F
       FPC 7 QX 0 TSensor             OK          42 degrees C / 107 degrees F
       FPC 7 QX 0 Chip                OK          47 degrees C / 116 degrees F
       FPC 7 LU 0 TCAM TSensor        OK          42 degrees C / 107 degrees F
       FPC 7 LU 0 TCAM Chip           OK          44 degrees C / 111 degrees F
       FPC 7 LU 0 TSensor             OK          42 degrees C / 107 degrees F
       FPC 7 LU 0 Chip                OK          46 degrees C / 114 degrees F
       FPC 7 MQ 0 TSensor             OK          42 degrees C / 107 degrees F
       FPC 7 MQ 0 Chip                OK          45 degrees C / 113 degrees F
       FPC 8 Intake                   OK          33 degrees C / 91 degrees F
       FPC 8 Exhaust A                OK          33 degrees C / 91 degrees F
       FPC 8 Exhaust B                OK          36 degrees C / 96 degrees F
       FPC 8 I3 0 TSensor             OK          38 degrees C / 100 degrees F
       FPC 8 I3 0 Chip                OK          43 degrees C / 109 degrees F
       FPC 8 BDS 0 TSensor            OK          37 degrees C / 98 degrees F
       FPC 8 BDS 0 Chip               OK          36 degrees C / 96 degrees F
       FPC 8 IA 0 TSensor             OK          37 degrees C / 98 degrees F
       FPC 8 IA 0 Chip                OK          37 degrees C / 98 degrees F
       FPC 10 Intake                  OK          38 degrees C / 100 degrees F
       FPC 10 Exhaust A               OK          36 degrees C / 96 degrees F
       FPC 10 Exhaust B               OK          41 degrees C / 105 degrees F
       FPC 10 I3 0 TSensor            OK          40 degrees C / 104 degrees F
       FPC 10 I3 0 Chip               OK          42 degrees C / 107 degrees F
       FPC 10 I3 1 TSensor            OK          40 degrees C / 104 degrees F
       FPC 10 I3 1 Chip               OK          44 degrees C / 111 degrees F
       FPC 10 I3 2 TSensor            OK          42 degrees C / 107 degrees F
       FPC 10 I3 2 Chip               OK          43 degrees C / 109 degrees F
       FPC 10 I3 3 TSensor            OK          39 degrees C / 102 degrees F
       FPC 10 I3 3 Chip               OK          44 degrees C / 111 degrees F
       FPC 10 IA 0 TSensor            OK          36 degrees C / 96 degrees F
       FPC 10 IA 0 Chip               OK          36 degrees C / 96 degrees F
       FPC 10 IA 1 TSensor            OK          43 degrees C / 109 degrees F
       FPC 10 IA 1 Chip               OK          42 degrees C / 107 degrees F
 Fans  Top Fan Tray Temp              OK          37 degrees C / 98 degrees F
```

```
      Top Tray Fan 1                OK         Spinning at normal speed
      Top Tray Fan 2                OK         Spinning at normal speed
      Top Tray Fan 3                OK         Spinning at normal speed
      Top Tray Fan 4                OK         Spinning at normal speed
      Top Tray Fan 5                OK         Spinning at normal speed
      Top Tray Fan 6                OK         Spinning at normal speed
      Bottom Fan Tray Temp          OK         28 degrees C / 82 degrees F
      Bottom Tray Fan 1             OK         Spinning at normal speed
      Bottom Tray Fan 2             OK         Spinning at normal speed
      Bottom Tray Fan 3             OK         Spinning at normal speed
      Bottom Tray Fan 4             OK         Spinning at normal speed
      Bottom Tray Fan 5             OK         Spinning at normal speed
      Bottom Tray Fan 6             OK         Spinning at normal speed
```

show chassis environment (MX960 Router with MPC5EQ)

```
user@host> show chassis environment
Class Item                          Status     Measurement
Temp  PEM 0                         OK         50 degrees C / 122 degrees F
      PEM 1                         OK         45 degrees C / 113 degrees F
      PEM 2                         OK         45 degrees C / 113 degrees F
      PEM 3                         Absent
      Routing Engine 0              OK         31 degrees C / 87 degrees F
      Routing Engine 0 CPU          OK         30 degrees C / 86 degrees F
      Routing Engine 1              Present
      Routing Engine 1 CPU          Present
      CB 0 Intake                   OK         29 degrees C / 84 degrees F
      CB 0 Exhaust A                OK         29 degrees C / 84 degrees F
      CB 0 Exhaust B                OK         34 degrees C / 93 degrees F
      CB 0 ACBC                     OK         32 degrees C / 89 degrees F
      CB 0 XF A                     OK         49 degrees C / 120 degrees F
      CB 0 XF B                     OK         45 degrees C / 113 degrees F
      CB 1 Intake                   OK         26 degrees C / 78 degrees F
      CB 1 Exhaust A                OK         26 degrees C / 78 degrees F
      CB 1 Exhaust B                OK         27 degrees C / 80 degrees F
      CB 1 ACBC                     OK         26 degrees C / 78 degrees F
      CB 1 XF A                     OK         32 degrees C / 89 degrees F
      CB 1 XF B                     OK         32 degrees C / 89 degrees F
      CB 2 Intake                   OK         28 degrees C / 82 degrees F
      CB 2 Exhaust A                OK         27 degrees C / 80 degrees F
      CB 2 Exhaust B                OK         33 degrees C / 91 degrees F
```

```
     CB 2 ACBC                    OK         30 degrees C / 86 degrees F
     CB 2 XF A                    OK         48 degrees C / 118 degrees F
     CB 2 XF B                    OK         46 degrees C / 114 degrees F
     FPC 0 Intake                 OK         38 degrees C / 100 degrees F
     FPC 0 Exhaust A              OK         48 degrees C / 118 degrees F
     FPC 0 Exhaust B              OK         49 degrees C / 120 degrees F
     FPC 0 XL TSen                OK         48 degrees C / 118 degrees F
     FPC 0 XL Chip                OK         50 degrees C / 122 degrees F
     FPC 0 XL_XR0 TSen            OK         48 degrees C / 118 degrees F
     FPC 0 XL_XR0 Chip            OK         53 degrees C / 127 degrees F
     FPC 0 XL_XR1 TSen            OK         48 degrees C / 118 degrees F
     FPC 0 XL_XR1 Chip            OK         54 degrees C / 129 degrees F
     FPC 0 XQ TSen                OK         48 degrees C / 118 degrees F
     FPC 0 XQ Chip                OK         52 degrees C / 125 degrees F
     FPC 0 XQ_XR0 TSen            OK         48 degrees C / 118 degrees F
     FPC 0 XQ_XR0 Chip            OK         62 degrees C / 143 degrees F
     FPC 0 XQ_XR1 TSen            OK         48 degrees C / 118 degrees F
     FPC 0 XQ_XR1 Chip            OK         62 degrees C / 143 degrees F
     FPC 0 XM 0 TSen              OK         53 degrees C / 127 degrees F
     FPC 0 XM 0 Chip              OK         63 degrees C / 145 degrees F
     FPC 0 XM 1 TSen              OK         53 degrees C / 127 degrees F
     FPC 0 XM 1 Chip              OK         46 degrees C / 114 degrees F
     FPC 0 PLX PCIe Switch TSe    OK         53 degrees C / 127 degrees F
     FPC 0 PLX PCIe Switch Chi    OK         66 degrees C / 150 degrees F
     FPC 1 Intake                 OK         31 degrees C / 87 degrees F
     FPC 1 Exhaust A              OK         38 degrees C / 100 degrees F
     FPC 1 Exhaust B              OK         49 degrees C / 120 degrees F
     FPC 1 LU 0 TSen              OK         41 degrees C / 105 degrees F
     FPC 1 LU 0 Chip              OK         47 degrees C / 116 degrees F
     FPC 1 LU 1 TSen              OK         41 degrees C / 105 degrees F
FPC 1 LU 1 Chip               OK         42 degrees C / 107 degrees F
     FPC 1 LU 2 TSen              OK         41 degrees C / 105 degrees F
     FPC 1 LU 2 Chip              OK         46 degrees C / 114 degrees F
     FPC 1 LU 3 TSen              OK         41 degrees C / 105 degrees F
     FPC 1 LU 3 Chip              OK         51 degrees C / 123 degrees F
     FPC 1 XM 0 TSen              OK         41 degrees C / 105 degrees F
     FPC 1 XM 0 Chip              OK         49 degrees C / 120 degrees F
     FPC 1 XF 0 TSen              OK         41 degrees C / 105 degrees F
     FPC 1 XF 0 Chip              OK         63 degrees C / 145 degrees F
     FPC 1 PLX Switch TSen        OK         41 degrees C / 105 degrees F
     FPC 1 PLX Switch Chip        OK         43 degrees C / 109 degrees F
     FPC 3 Intake                 OK         31 degrees C / 87 degrees F
     FPC 3 Exhaust A              OK         37 degrees C / 98 degrees F
```

```
    FPC 3 Exhaust B                OK          43 degrees C / 109 degrees F
    FPC 3 LU 0 TSen                OK          42 degrees C / 107 degrees F
    FPC 3 LU 0 Chip                OK          43 degrees C / 109 degrees F
    FPC 3 LU 1 TSen                OK          42 degrees C / 107 degrees F
    FPC 3 LU 1 Chip                OK          46 degrees C / 114 degrees F
    FPC 3 LU 2 TSen                OK          42 degrees C / 107 degrees F
    FPC 3 LU 2 Chip                OK          40 degrees C / 104 degrees F
    FPC 3 LU 3 TSen                OK          42 degrees C / 107 degrees F
    FPC 3 LU 3 Chip                OK          41 degrees C / 105 degrees F
    FPC 3 MQ 0 TSen                OK          37 degrees C / 98 degrees F
    FPC 3 MQ 0 Chip                OK          37 degrees C / 98 degrees F
    FPC 3 MQ 1 TSen                OK          37 degrees C / 98 degrees F
    FPC 3 MQ 1 Chip                OK          40 degrees C / 104 degrees F
    FPC 3 MQ 2 TSen                OK          37 degrees C / 98 degrees F
    FPC 3 MQ 2 Chip                OK          36 degrees C / 96 degrees F
    FPC 3 MQ 3 TSen                OK          37 degrees C / 98 degrees F
    FPC 3 MQ 3 Chip                OK          38 degrees C / 100 degrees F
    FPC 4 Intake                   OK          34 degrees C / 93 degrees F
    FPC 4 Exhaust A                OK          45 degrees C / 113 degrees F
    FPC 4 Exhaust B                OK          47 degrees C / 116 degrees F
    FPC 4 XL TSen                  OK          44 degrees C / 111 degrees F
    FPC 4 XL Chip                  OK          47 degrees C / 116 degrees F
    FPC 4 XL_XR0 TSen              OK          44 degrees C / 111 degrees F
    FPC 4 XL_XR0 Chip              OK          48 degrees C / 118 degrees F
    FPC 4 XL_XR1 TSen              OK          44 degrees C / 111 degrees F
    FPC 4 XL_XR1 Chip              OK          47 degrees C / 116 degrees F
    FPC 4 XQ TSen                  OK          44 degrees C / 111 degrees F
    FPC 4 XQ Chip                  OK          47 degrees C / 116 degrees F
    FPC 4 XQ_XR0 TSen              OK          44 degrees C / 111 degrees F
    FPC 4 XQ_XR0 Chip              OK          57 degrees C / 134 degrees F
    FPC 4 XQ_XR1 TSen              OK          44 degrees C / 111 degrees F
    FPC 4 XQ_XR1 Chip              OK          58 degrees C / 136 degrees F
    FPC 4 XM 0 TSen                OK          51 degrees C / 123 degrees F
    FPC 4 XM 0 Chip                OK          61 degrees C / 141 degrees F
    FPC 4 XM 1 TSen                OK          51 degrees C / 123 degrees F
    FPC 4 XM 1 Chip                OK          47 degrees C / 116 degrees F
    FPC 4 PLX PCIe Switch TSe      OK          51 degrees C / 123 degrees F
    FPC 4 PLX PCIe Switch Chi      OK          60 degrees C / 140 degrees F
    FPC 5 Intake                   OK          34 degrees C / 93 degrees F
    FPC 5 Exhaust A                OK          45 degrees C / 113 degrees F
    FPC 5 Exhaust B                OK          47 degrees C / 116 degrees F
    FPC 5 XL TSen                  OK          45 degrees C / 113 degrees F
    FPC 5 XL Chip                  OK          47 degrees C / 116 degrees F
```

```
     FPC 5 XL_XR0 TSen              OK         45 degrees C / 113 degrees F
FPC 5 XL_XR0 Chip              OK         49 degrees C / 120 degrees F
     FPC 5 XL_XR1 TSen              OK         45 degrees C / 113 degrees F
     FPC 5 XL_XR1 Chip              OK         49 degrees C / 120 degrees F
     FPC 5 XQ TSen                  OK         45 degrees C / 113 degrees F
     FPC 5 XQ Chip                  OK         48 degrees C / 118 degrees F
     FPC 5 XQ_XR0 TSen              OK         45 degrees C / 113 degrees F
     FPC 5 XQ_XR0 Chip              OK         60 degrees C / 140 degrees F
     FPC 5 XQ_XR1 TSen              OK         45 degrees C / 113 degrees F
     FPC 5 XQ_XR1 Chip              OK         58 degrees C / 136 degrees F
     FPC 5 XM 0 TSen                OK         50 degrees C / 122 degrees F
     FPC 5 XM 0 Chip                OK         48 degrees C / 118 degrees F
     FPC 5 XM 1 TSen                OK         50 degrees C / 122 degrees F
     FPC 5 XM 1 Chip                OK         47 degrees C / 116 degrees F
     FPC 5 PLX PCIe Switch TSe      OK         50 degrees C / 122 degrees F
     FPC 5 PLX PCIe Switch Chi      OK         59 degrees C / 138 degrees F
     FPC 7 Intake                   OK         32 degrees C / 89 degrees F
     FPC 7 Exhaust A                OK         32 degrees C / 89 degrees F
     FPC 7 Exhaust B                OK         33 degrees C / 91 degrees F
     FPC 7 LU 0 TSen                OK         49 degrees C / 120 degrees F
     FPC 7 LU 0 Chip                OK         44 degrees C / 111 degrees F
     FPC 7 LU 1 TSen                OK         49 degrees C / 120 degrees F
     FPC 7 LU 1 Chip                OK         47 degrees C / 116 degrees F
     FPC 7 LU 2 TSen                OK         49 degrees C / 120 degrees F
     FPC 7 LU 2 Chip                OK         39 degrees C / 102 degrees F
     FPC 7 LU 3 TSen                OK         49 degrees C / 120 degrees F
     FPC 7 LU 3 Chip                OK         43 degrees C / 109 degrees F
     FPC 7 XM 0 TSen                OK         49 degrees C / 120 degrees F
     FPC 7 XM 0 Chip                OK         57 degrees C / 134 degrees F
     FPC 7 XM 1 TSen                OK         49 degrees C / 120 degrees F
     FPC 7 XM 1 Chip                OK         48 degrees C / 118 degrees F
     FPC 7 PLX Switch TSen          OK         49 degrees C / 120 degrees F
     FPC 7 PLX Switch Chip          OK         45 degrees C / 113 degrees F
     FPC 8 Intake                   OK         36 degrees C / 96 degrees F
     FPC 8 Exhaust A                OK         51 degrees C / 123 degrees F
     FPC 8 Exhaust B                OK         46 degrees C / 114 degrees F
     FPC 8 XL TSen                  OK         46 degrees C / 114 degrees F
     FPC 8 XL Chip                  OK         47 degrees C / 116 degrees F
     FPC 8 XL_XR0 TSen              OK         46 degrees C / 114 degrees F
     FPC 8 XL_XR0 Chip              OK         53 degrees C / 127 degrees F
     FPC 8 XL_XR1 TSen              OK         46 degrees C / 114 degrees F
     FPC 8 XL_XR1 Chip              OK         52 degrees C / 125 degrees F
     FPC 8 XQ TSen                  OK         46 degrees C / 114 degrees F
```

```
    FPC 8 XQ Chip                  OK         46 degrees C / 114 degrees F
    FPC 8 XQ_XR0 TSen              OK         46 degrees C / 114 degrees F
    FPC 8 XQ_XR0 Chip              OK         59 degrees C / 138 degrees F
    FPC 8 XQ_XR1 TSen              OK         46 degrees C / 114 degrees F
    FPC 8 XQ_XR1 Chip              OK         57 degrees C / 134 degrees F
    FPC 8 XM 0 TSen                OK         52 degrees C / 125 degrees F
    FPC 8 XM 0 Chip                OK         61 degrees C / 141 degrees F
    FPC 8 XM 1 TSen                OK         52 degrees C / 125 degrees F
    FPC 8 XM 1 Chip                OK         47 degrees C / 116 degrees F
    FPC 8 PLX PCIe Switch TSe      OK         52 degrees C / 125 degrees F
    FPC 8 PLX PCIe Switch Chi      OK         63 degrees C / 145 degrees F
    FPC 9 Intake                   OK         31 degrees C / 87 degrees F
    FPC 9 Exhaust A                OK         34 degrees C / 93 degrees F
    FPC 9 Exhaust B                OK         35 degrees C / 95 degrees F
FPC 9 QX 0 TSen                OK         42 degrees C / 107 degrees F
    FPC 9 QX 0 Chip                OK         45 degrees C / 113 degrees F
    FPC 9 LU 0 TCAM TSen           OK         42 degrees C / 107 degrees F
    FPC 9 LU 0 TCAM Chip           OK         41 degrees C / 105 degrees F
    FPC 9 LU 0 TSen                OK         42 degrees C / 107 degrees F
    FPC 9 LU 0 Chip                OK         43 degrees C / 109 degrees F
    FPC 9 MQ 0 TSen                OK         42 degrees C / 107 degrees F
    FPC 9 MQ 0 Chip                OK         43 degrees C / 109 degrees F
    FPC 9 QX 1 TSen                OK         38 degrees C / 100 degrees F
    FPC 9 QX 1 Chip                OK         40 degrees C / 104 degrees F
    FPC 9 LU 1 TCAM TSen           OK         38 degrees C / 100 degrees F
    FPC 9 LU 1 TCAM Chip           OK         38 degrees C / 100 degrees F
    FPC 9 LU 1 TSen                OK         38 degrees C / 100 degrees F
    FPC 9 LU 1 Chip                OK         41 degrees C / 105 degrees F
    FPC 9 MQ 1 TSen                OK         38 degrees C / 100 degrees F
    FPC 9 MQ 1 Chip                OK         41 degrees C / 105 degrees F
    FPC 10 Intake                  OK         35 degrees C / 95 degrees F
    FPC 10 Exhaust A               OK         51 degrees C / 123 degrees F
    FPC 10 Exhaust B               OK         46 degrees C / 114 degrees F
    FPC 10 XL TSen                 OK         42 degrees C / 107 degrees F
    FPC 10 XL Chip                 OK         44 degrees C / 111 degrees F
    FPC 10 XL_XR0 TSen             OK         42 degrees C / 107 degrees F
    FPC 10 XL_XR0 Chip             OK         47 degrees C / 116 degrees F
    FPC 10 XL_XR1 TSen             OK         42 degrees C / 107 degrees F
    FPC 10 XL_XR1 Chip             OK         48 degrees C / 118 degrees F
    FPC 10 XQ TSen                 OK         42 degrees C / 107 degrees F
    FPC 10 XQ Chip                 OK         46 degrees C / 114 degrees F
    FPC 10 XQ_XR0 TSen             OK         42 degrees C / 107 degrees F
    FPC 10 XQ_XR0 Chip             OK         57 degrees C / 134 degrees F
```

```
      FPC 10 XQ_XR1 TSen            OK         42 degrees C / 107 degrees F
      FPC 10 XQ_XR1 Chip            OK         53 degrees C / 127 degrees F
      FPC 10 XM 0 TSen              OK         51 degrees C / 123 degrees F
      FPC 10 XM 0 Chip              OK         61 degrees C / 141 degrees F
      FPC 10 XM 1 TSen              OK         51 degrees C / 123 degrees F
      FPC 10 XM 1 Chip              OK         49 degrees C / 120 degrees F
      FPC 10 PLX PCIe Switch TSe    OK         51 degrees C / 123 degrees F
      FPC 10 PLX PCIe Switch Chi    OK         61 degrees C / 141 degrees F
      FPC 11 Intake                 OK         33 degrees C / 91 degrees F
      FPC 11 Exhaust A              OK         33 degrees C / 91 degrees F
      FPC 11 Exhaust B              OK         34 degrees C / 93 degrees F
      FPC 11 LU 0 TSen              OK         50 degrees C / 122 degrees F
      FPC 11 LU 0 Chip              OK         48 degrees C / 118 degrees F
      FPC 11 LU 1 TSen              OK         50 degrees C / 122 degrees F
      FPC 11 LU 1 Chip              OK         50 degrees C / 122 degrees F
      FPC 11 LU 2 TSen              OK         50 degrees C / 122 degrees F
      FPC 11 LU 2 Chip              OK         41 degrees C / 105 degrees F
      FPC 11 LU 3 TSen              OK         50 degrees C / 122 degrees F
      FPC 11 LU 3 Chip              OK         48 degrees C / 118 degrees F
      FPC 11 XM 0 TSen              OK         50 degrees C / 122 degrees F
      FPC 11 XM 0 Chip              OK         57 degrees C / 134 degrees F
      FPC 11 XM 1 TSen              OK         50 degrees C / 122 degrees F
      FPC 11 XM 1 Chip              OK         52 degrees C / 125 degrees F
      FPC 11 PLX Switch TSen        OK         50 degrees C / 122 degrees F
      FPC 11 PLX Switch Chip        OK         45 degrees C / 113 degrees F
Fans  Top Fan Tray Temp             OK         42 degrees C / 107 degrees F
      Top Tray Fan 1                OK         Spinning at high speed
 Top Tray Fan 2                 OK         Spinning at high speed
      Top Tray Fan 3                OK         Spinning at high speed
      Top Tray Fan 4                OK         Spinning at high speed
      Top Tray Fan 5                OK         Spinning at high speed
      Top Tray Fan 6                OK         Spinning at high speed
      Top Tray Fan 7                OK         Spinning at high speed
      Top Tray Fan 8                OK         Spinning at high speed
      Top Tray Fan 9                OK         Spinning at high speed
      Top Tray Fan 10               OK         Spinning at high speed
      Top Tray Fan 11               OK         Spinning at high speed
      Top Tray Fan 12               OK         Spinning at high speed
      Bottom Fan Tray Temp          OK         33 degrees C / 91 degrees F
      Bottom Tray Fan 1             OK         Spinning at high speed
      Bottom Tray Fan 2             OK         Spinning at high speed
      Bottom Tray Fan 3             OK         Spinning at high speed
      Bottom Tray Fan 4             OK         Spinning at high speed
```

```
      Bottom Tray Fan 5              OK         Spinning at high speed
      Bottom Tray Fan 6              OK         Spinning at high speed
      Bottom Tray Fan 7              OK         Spinning at high speed
      Bottom Tray Fan 8              OK         Spinning at high speed
      Bottom Tray Fan 9              OK         Spinning at high speed
      Bottom Tray Fan 10             OK         Spinning at high speed
      Bottom Tray Fan 11             OK         Spinning at high speed
      Bottom Tray Fan 12             OK         Spinning at high speed
```

show chassis environment (MX2020 Router)

```
user@host> show chassis environment
Class Item                           Status     Measurement
Temp  PSM 0                          Absent
      PSM 1                          Absent
      PSM 2                          OK         41 degrees C / 105 degrees F
      PSM 3                          OK         39 degrees C / 102 degrees F
      PSM 4                          OK         39 degrees C / 102 degrees F
      PSM 5                          OK         38 degrees C / 100 degrees F
      PSM 6                          OK         38 degrees C / 100 degrees F
      PSM 7                          OK         38 degrees C / 100 degrees F
      PSM 8                          OK         37 degrees C / 98 degrees F
      PSM 9                          Absent
      PSM 10                         Absent
      PSM 11                         OK         47 degrees C / 116 degrees F
      PSM 12                         OK         45 degrees C / 113 degrees F
      PSM 13                         OK         44 degrees C / 111 degrees F
      PSM 14                         OK         44 degrees C / 111 degrees F
      PSM 15                         OK         43 degrees C / 109 degrees F
      PSM 16                         OK         42 degrees C / 107 degrees F
      PSM 17                         OK         41 degrees C / 105 degrees F
      PDM 0                          OK
      PDM 1                          Absent
      PDM 2                          Absent
      PDM 3                          OK
      CB 0 IntakeA-Zone0             OK         45 degrees C / 113 degrees F
      CB 0 IntakeB-Zone1             OK         34 degrees C / 93 degrees F
      CB 0 IntakeC-Zone0             OK         48 degrees C / 118 degrees F
      CB 0 ExhaustA-Zone0            OK         45 degrees C / 113 degrees F
      CB 0 ExhaustB-Zone1            OK         37 degrees C / 98 degrees F
      CB 0 TCBC-Zone0                OK         41 degrees C / 105 degrees F
```

```
    CB 1 IntakeA-Zone0          OK        46 degrees C / 114 degrees F
    CB 1 IntakeB-Zone1          OK        42 degrees C / 107 degrees F
    CB 1 IntakeC-Zone0          OK        49 degrees C / 120 degrees F
    CB 1 ExhaustA-Zone0         OK        46 degrees C / 114 degrees F
    CB 1 ExhaustB-Zone1         OK        41 degrees C / 105 degrees F
    CB 1 TCBC-Zone0             OK        46 degrees C / 114 degrees F
    SPMB 0 Intake               OK        33 degrees C / 91 degrees F
    SPMB 1 Intake               OK        42 degrees C / 107 degrees F
    Routing Engine 0            OK        35 degrees C / 95 degrees F
    Routing Engine 0 CPU        OK        34 degrees C / 93 degrees F
    Routing Engine 1            OK        44 degrees C / 111 degrees F
    Routing Engine 1 CPU        OK        42 degrees C / 107 degrees F
    SFB 0 Intake-Zone0          OK        55 degrees C / 131 degrees F
    SFB 0 Exhaust-Zone1         OK        48 degrees C / 118 degrees F
    SFB 0 IntakeA-Zone0         OK        50 degrees C / 122 degrees F
    SFB 0 IntakeB-Zone1         OK        40 degrees C / 104 degrees F
    SFB 0 Exhaust-Zone0         OK        52 degrees C / 125 degrees F
    SFB 0 SFB-XF2-Zone1         OK        61 degrees C / 141 degrees F
    SFB 0 SFB-XF1-Zone0         OK        69 degrees C / 156 degrees F
    SFB 0 SFB-XF0-Zone0         OK        68 degrees C / 154 degrees F
    SFB 1 Intake-Zone0          OK        56 degrees C / 132 degrees F
    SFB 1 Exhaust-Zone1         OK        47 degrees C / 116 degrees F
    SFB 1 IntakeA-Zone0         OK        51 degrees C / 123 degrees F
    SFB 1 IntakeB-Zone1         OK        40 degrees C / 104 degrees F
    SFB 1 Exhaust-Zone0         OK        51 degrees C / 123 degrees F
    SFB 1 SFB-XF2-Zone1         OK        62 degrees C / 143 degrees F
    SFB 1 SFB-XF1-Zone0         OK        67 degrees C / 152 degrees F
    SFB 1 SFB-XF0-Zone0         OK        69 degrees C / 156 degrees F
    SFB 2 Intake-Zone0          OK        56 degrees C / 132 degrees F
    SFB 2 Exhaust-Zone1         OK        47 degrees C / 116 degrees F
    SFB 2 IntakeA-Zone0         OK        51 degrees C / 123 degrees F
    SFB 2 IntakeB-Zone1         OK        40 degrees C / 104 degrees F
    SFB 2 Exhaust-Zone0         OK        53 degrees C / 127 degrees F
    SFB 2 SFB-XF2-Zone1         OK        65 degrees C / 149 degrees F
    SFB 2 SFB-XF1-Zone0         OK        69 degrees C / 156 degrees F
    SFB 2 SFB-XF0-Zone0         OK        70 degrees C / 158 degrees F
    SFB 3 Intake-Zone0          OK        57 degrees C / 134 degrees F
    SFB 3 Exhaust-Zone1         OK        48 degrees C / 118 degrees F
    SFB 3 IntakeA-Zone0         OK        52 degrees C / 125 degrees F
    SFB 3 IntakeB-Zone1         OK        41 degrees C / 105 degrees F
    SFB 3 Exhaust-Zone0         OK        53 degrees C / 127 degrees F
    SFB 3 SFB-XF2-Zone1         OK        66 degrees C / 150 degrees F
    SFB 3 SFB-XF1-Zone0         OK        69 degrees C / 156 degrees F
```

```
    SFB 3 SFB-XF0-Zone0          OK        71 degrees C / 159 degrees F
    SFB 4 Intake-Zone0           OK        58 degrees C / 136 degrees F
    SFB 4 Exhaust-Zone1          OK        49 degrees C / 120 degrees F
    SFB 4 IntakeA-Zone0          OK        54 degrees C / 129 degrees F
    SFB 4 IntakeB-Zone1          OK        42 degrees C / 107 degrees F
    SFB 4 Exhaust-Zone0          OK        53 degrees C / 127 degrees F
    SFB 4 SFB-XF2-Zone1          OK        64 degrees C / 147 degrees F
    SFB 4 SFB-XF1-Zone0          OK        68 degrees C / 154 degrees F
    SFB 4 SFB-XF0-Zone0          OK        71 degrees C / 159 degrees F
    SFB 5 Intake-Zone0           OK        58 degrees C / 136 degrees F
    SFB 5 Exhaust-Zone1          OK        50 degrees C / 122 degrees F
    SFB 5 IntakeA-Zone0          OK        53 degrees C / 127 degrees F
    SFB 5 IntakeB-Zone1          OK        43 degrees C / 109 degrees F
    SFB 5 Exhaust-Zone0          OK        54 degrees C / 129 degrees F
    SFB 5 SFB-XF2-Zone1          OK        66 degrees C / 150 degrees F
    SFB 5 SFB-XF1-Zone0          OK        69 degrees C / 156 degrees F
    SFB 5 SFB-XF0-Zone0          OK        74 degrees C / 165 degrees F
    SFB 6 Intake-Zone0           OK        58 degrees C / 136 degrees F
    SFB 6 Exhaust-Zone1          OK        49 degrees C / 120 degrees F
    SFB 6 IntakeA-Zone0          OK        53 degrees C / 127 degrees F
    SFB 6 IntakeB-Zone1          OK        43 degrees C / 109 degrees F
    SFB 6 Exhaust-Zone0          OK        53 degrees C / 127 degrees F
    SFB 6 SFB-XF2-Zone1          OK        65 degrees C / 149 degrees F
    SFB 6 SFB-XF1-Zone0          OK        68 degrees C / 154 degrees F
    SFB 6 SFB-XF0-Zone0          OK        72 degrees C / 161 degrees F
    SFB 7 Intake-Zone0           OK        57 degrees C / 134 degrees F
    SFB 7 Exhaust-Zone1          OK        50 degrees C / 122 degrees F
    SFB 7 IntakeA-Zone0          OK        53 degrees C / 127 degrees F
    SFB 7 IntakeB-Zone1          OK        43 degrees C / 109 degrees F
    SFB 7 Exhaust-Zone0          OK        54 degrees C / 129 degrees F
    SFB 7 SFB-XF2-Zone1          OK        68 degrees C / 154 degrees F
    SFB 7 SFB-XF1-Zone0          OK        69 degrees C / 156 degrees F
    SFB 7 SFB-XF0-Zone0          OK        73 degrees C / 163 degrees F
    FPC 0 Intake                 OK        41 degrees C / 105 degrees F
    FPC 0 Exhaust A              OK        48 degrees C / 118 degrees F
    FPC 0 Exhaust B              OK        62 degrees C / 143 degrees F
    FPC 0 LU 0 TSen              OK        59 degrees C / 138 degrees F
    FPC 0 LU 0 Chip              OK        62 degrees C / 143 degrees F
    FPC 0 LU 1 TSen              OK        59 degrees C / 138 degrees F
    FPC 0 LU 1 Chip              OK        64 degrees C / 147 degrees F
    FPC 0 LU 2 TSen              OK        59 degrees C / 138 degrees F
    FPC 0 LU 2 Chip              OK        53 degrees C / 127 degrees F
    FPC 0 LU 3 TSen              OK        59 degrees C / 138 degrees F
```

```
    FPC 0 LU 3 Chip                OK         53 degrees C / 127 degrees F
    FPC 0 MQ 0 TSen                OK         47 degrees C / 116 degrees F
    FPC 0 MQ 0 Chip                OK         49 degrees C / 120 degrees F
    FPC 0 MQ 1 TSen                OK         47 degrees C / 116 degrees F
    FPC 0 MQ 1 Chip                OK         51 degrees C / 123 degrees F
    FPC 0 MQ 2 TSen                OK         47 degrees C / 116 degrees F
    FPC 0 MQ 2 Chip                OK         44 degrees C / 111 degrees F
    FPC 0 MQ 3 TSen                OK         47 degrees C / 116 degrees F
    FPC 0 MQ 3 Chip                OK         45 degrees C / 113 degrees F
    FPC 1 Intake                   OK         40 degrees C / 104 degrees F
    FPC 1 Exhaust A                OK         49 degrees C / 120 degrees F
    FPC 1 Exhaust B                OK         58 degrees C / 136 degrees F
    FPC 1 LU 0 TSen                OK         55 degrees C / 131 degrees F
    FPC 1 LU 0 Chip                OK         56 degrees C / 132 degrees F
    FPC 1 LU 1 TSen                OK         55 degrees C / 131 degrees F
    FPC 1 LU 1 Chip                OK         58 degrees C / 136 degrees F
    FPC 1 LU 2 TSen                OK         55 degrees C / 131 degrees F
    FPC 1 LU 2 Chip                OK         49 degrees C / 120 degrees F
    FPC 1 LU 3 TSen                OK         55 degrees C / 131 degrees F
    FPC 1 LU 3 Chip                OK         51 degrees C / 123 degrees F
    FPC 1 MQ 0 TSen                OK         47 degrees C / 116 degrees F
    FPC 1 MQ 0 Chip                OK         48 degrees C / 118 degrees F
    FPC 1 MQ 1 TSen                OK         47 degrees C / 116 degrees F
    FPC 1 MQ 1 Chip                OK         50 degrees C / 122 degrees F
    FPC 1 MQ 2 TSen                OK         47 degrees C / 116 degrees F
    FPC 1 MQ 2 Chip                OK         44 degrees C / 111 degrees F
    FPC 1 MQ 3 TSen                OK         47 degrees C / 116 degrees F
    FPC 1 MQ 3 Chip                OK         44 degrees C / 111 degrees F
    FPC 2 Intake                   OK         39 degrees C / 102 degrees F
    FPC 2 Exhaust A                OK         49 degrees C / 120 degrees F
    FPC 2 Exhaust B                OK         61 degrees C / 141 degrees F
    FPC 2 LU 0 TSen                OK         58 degrees C / 136 degrees F
    FPC 2 LU 0 Chip                OK         60 degrees C / 140 degrees F
    FPC 2 LU 1 TSen                OK         58 degrees C / 136 degrees F
    FPC 2 LU 1 Chip                OK         65 degrees C / 149 degrees F
    FPC 2 LU 2 TSen                OK         58 degrees C / 136 degrees F
    FPC 2 LU 2 Chip                OK         51 degrees C / 123 degrees F
    FPC 2 LU 3 TSen                OK         58 degrees C / 136 degrees F
    FPC 2 LU 3 Chip                OK         53 degrees C / 127 degrees F
    FPC 2 MQ 0 TSen                OK         47 degrees C / 116 degrees F
    FPC 2 MQ 0 Chip                OK         50 degrees C / 122 degrees F
    FPC 2 MQ 1 TSen                OK         47 degrees C / 116 degrees F
    FPC 2 MQ 1 Chip                OK         52 degrees C / 125 degrees F
```

```
    FPC 2 MQ 2 TSen              OK        47 degrees C / 116 degrees F
    FPC 2 MQ 2 Chip              OK        45 degrees C / 113 degrees F
    FPC 2 MQ 3 TSen              OK        47 degrees C / 116 degrees F
    FPC 2 MQ 3 Chip              OK        46 degrees C / 114 degrees F
    FPC 3 Intake                 OK        40 degrees C / 104 degrees F
    FPC 3 Exhaust A              OK        49 degrees C / 120 degrees F
    FPC 3 Exhaust B              OK        61 degrees C / 141 degrees F
    FPC 3 LU 0 TSen              OK        58 degrees C / 136 degrees F
    FPC 3 LU 0 Chip              OK        61 degrees C / 141 degrees F
    FPC 3 LU 1 TSen              OK        58 degrees C / 136 degrees F
    FPC 3 LU 1 Chip              OK        62 degrees C / 143 degrees F
    FPC 3 LU 2 TSen              OK        58 degrees C / 136 degrees F
    FPC 3 LU 2 Chip              OK        51 degrees C / 123 degrees F
    FPC 3 LU 3 TSen              OK        58 degrees C / 136 degrees F
    FPC 3 LU 3 Chip              OK        53 degrees C / 127 degrees F
    FPC 3 MQ 0 TSen              OK        48 degrees C / 118 degrees F
    FPC 3 MQ 0 Chip              OK        50 degrees C / 122 degrees F
    FPC 3 MQ 1 TSen              OK        48 degrees C / 118 degrees F
    FPC 3 MQ 1 Chip              OK        54 degrees C / 129 degrees F
    FPC 3 MQ 2 TSen              OK        48 degrees C / 118 degrees F
    FPC 3 MQ 2 Chip              OK        45 degrees C / 113 degrees F
    FPC 3 MQ 3 TSen              OK        48 degrees C / 118 degrees F
    FPC 3 MQ 3 Chip              OK        48 degrees C / 118 degrees F
    FPC 4 Intake                 OK        40 degrees C / 104 degrees F
    FPC 4 Exhaust A              OK        49 degrees C / 120 degrees F
    FPC 4 Exhaust B              OK        62 degrees C / 143 degrees F
    FPC 4 LU 0 TSen              OK        59 degrees C / 138 degrees F
    FPC 4 LU 0 Chip              OK        62 degrees C / 143 degrees F
    FPC 4 LU 1 TSen              OK        59 degrees C / 138 degrees F
    FPC 4 LU 1 Chip              OK        65 degrees C / 149 degrees F
    FPC 4 LU 2 TSen              OK        59 degrees C / 138 degrees F
    FPC 4 LU 2 Chip              OK        51 degrees C / 123 degrees F
    FPC 4 LU 3 TSen              OK        59 degrees C / 138 degrees F
    FPC 4 LU 3 Chip              OK        53 degrees C / 127 degrees F
    FPC 4 MQ 0 TSen              OK        48 degrees C / 118 degrees F
    FPC 4 MQ 0 Chip              OK        52 degrees C / 125 degrees F
    FPC 4 MQ 1 TSen              OK        48 degrees C / 118 degrees F
    FPC 4 MQ 1 Chip              OK        53 degrees C / 127 degrees F
    FPC 4 MQ 2 TSen              OK        48 degrees C / 118 degrees F
    FPC 4 MQ 2 Chip              OK        46 degrees C / 114 degrees F
    FPC 4 MQ 3 TSen              OK        48 degrees C / 118 degrees F
    FPC 4 MQ 3 Chip              OK        47 degrees C / 116 degrees F
    FPC 5 Intake                 OK        41 degrees C / 105 degrees F
```

```
    FPC 5 Exhaust A                OK          50 degrees C / 122 degrees F
    FPC 5 Exhaust B                OK          63 degrees C / 145 degrees F
    FPC 5 LU 0 TSen                OK          60 degrees C / 140 degrees F
    FPC 5 LU 0 Chip                OK          63 degrees C / 145 degrees F
    FPC 5 LU 1 TSen                OK          60 degrees C / 140 degrees F
    FPC 5 LU 1 Chip                OK          66 degrees C / 150 degrees F
    FPC 5 LU 2 TSen                OK          60 degrees C / 140 degrees F
    FPC 5 LU 2 Chip                OK          56 degrees C / 132 degrees F
    FPC 5 LU 3 TSen                OK          60 degrees C / 140 degrees F
    FPC 5 LU 3 Chip                OK          54 degrees C / 129 degrees F
    FPC 5 MQ 0 TSen                OK          49 degrees C / 120 degrees F
    FPC 5 MQ 0 Chip                OK          52 degrees C / 125 degrees F
    FPC 5 MQ 1 TSen                OK          49 degrees C / 120 degrees F
    FPC 5 MQ 1 Chip                OK          53 degrees C / 127 degrees F
    FPC 5 MQ 2 TSen                OK          49 degrees C / 120 degrees F
    FPC 5 MQ 2 Chip                OK          48 degrees C / 118 degrees F
    FPC 5 MQ 3 TSen                OK          49 degrees C / 120 degrees F
    FPC 5 MQ 3 Chip                OK          47 degrees C / 116 degrees F
    FPC 6 Intake                   OK          42 degrees C / 107 degrees F
    FPC 6 Exhaust A                OK          51 degrees C / 123 degrees F
    FPC 6 Exhaust B                OK          63 degrees C / 145 degrees F
    FPC 6 LU 0 TSen                OK          61 degrees C / 141 degrees F
    FPC 6 LU 0 Chip                OK          64 degrees C / 147 degrees F
    FPC 6 LU 1 TSen                OK          61 degrees C / 141 degrees F
    FPC 6 LU 1 Chip                OK          66 degrees C / 150 degrees F
    FPC 6 LU 2 TSen                OK          61 degrees C / 141 degrees F
    FPC 6 LU 2 Chip                OK          56 degrees C / 132 degrees F
    FPC 6 LU 3 TSen                OK          61 degrees C / 141 degrees F
    FPC 6 LU 3 Chip                OK          56 degrees C / 132 degrees F
    FPC 6 MQ 0 TSen                OK          50 degrees C / 122 degrees F
    FPC 6 MQ 0 Chip                OK          56 degrees C / 132 degrees F
    FPC 6 MQ 1 TSen                OK          50 degrees C / 122 degrees F
    FPC 6 MQ 1 Chip                OK          59 degrees C / 138 degrees F
    FPC 6 MQ 2 TSen                OK          50 degrees C / 122 degrees F
    FPC 6 MQ 2 Chip                OK          49 degrees C / 120 degrees F
    FPC 6 MQ 3 TSen                OK          50 degrees C / 122 degrees F
    FPC 6 MQ 3 Chip                OK          49 degrees C / 120 degrees F
    FPC 7 Intake                   OK          41 degrees C / 105 degrees F
    FPC 7 Exhaust A                OK          51 degrees C / 123 degrees F
    FPC 7 Exhaust B                OK          63 degrees C / 145 degrees F
    FPC 7 LU 0 TSen                OK          60 degrees C / 140 degrees F
    FPC 7 LU 0 Chip                OK          61 degrees C / 141 degrees F
    FPC 7 LU 1 TSen                OK          60 degrees C / 140 degrees F
```

```
        FPC 7 LU 1 Chip                 OK          65 degrees C / 149 degrees F
        FPC 7 LU 2 TSen                 OK          60 degrees C / 140 degrees F
        FPC 7 LU 2 Chip                 OK          54 degrees C / 129 degrees F
        FPC 7 LU 3 TSen                 OK          60 degrees C / 140 degrees F
        FPC 7 LU 3 Chip                 OK          53 degrees C / 127 degrees F
        FPC 7 MQ 0 TSen                 OK          50 degrees C / 122 degrees F
        FPC 7 MQ 0 Chip                 OK          53 degrees C / 127 degrees F
        FPC 7 MQ 1 TSen                 OK          50 degrees C / 122 degrees F
        FPC 7 MQ 1 Chip                 OK          54 degrees C / 129 degrees F
        FPC 7 MQ 2 TSen                 OK          50 degrees C / 122 degrees F
        FPC 7 MQ 2 Chip                 OK          47 degrees C / 116 degrees F
        FPC 7 MQ 3 TSen                 OK          50 degrees C / 122 degrees F
        FPC 7 MQ 3 Chip                 OK          47 degrees C / 116 degrees F
        FPC 8 Intake                    OK          41 degrees C / 105 degrees F
        FPC 8 Exhaust A                 OK          50 degrees C / 122 degrees F
        FPC 8 Exhaust B                 OK          62 degrees C / 143 degrees F
        FPC 8 LU 0 TSen                 OK          59 degrees C / 138 degrees F
        FPC 8 LU 0 Chip                 OK          62 degrees C / 143 degrees F
        FPC 8 LU 1 TSen                 OK          59 degrees C / 138 degrees F
        FPC 8 LU 1 Chip                 OK          64 degrees C / 147 degrees F
        FPC 8 LU 2 TSen                 OK          59 degrees C / 138 degrees F
        FPC 8 LU 2 Chip                 OK          55 degrees C / 131 degrees F
        FPC 8 LU 3 TSen                 OK          59 degrees C / 138 degrees F
        FPC 8 LU 3 Chip                 OK          54 degrees C / 129 degrees F
        FPC 8 MQ 0 TSen                 OK          49 degrees C / 120 degrees F
        FPC 8 MQ 0 Chip                 OK          51 degrees C / 123 degrees F
        FPC 8 MQ 1 TSen                 OK          49 degrees C / 120 degrees F
        FPC 8 MQ 1 Chip                 OK          52 degrees C / 125 degrees F
        FPC 8 MQ 2 TSen                 OK          49 degrees C / 120 degrees F
        FPC 8 MQ 2 Chip                 OK          46 degrees C / 114 degrees F
        FPC 8 MQ 3 TSen                 OK          49 degrees C / 120 degrees F
        FPC 8 MQ 3 Chip                 OK          47 degrees C / 116 degrees F
        FPC 9 Intake                    OK          42 degrees C / 107 degrees F
        FPC 9 Exhaust A                 OK          51 degrees C / 123 degrees F
        FPC 9 Exhaust B                 OK          63 degrees C / 145 degrees F
        FPC 9 LU 0 TSen                 OK          60 degrees C / 140 degrees F
        FPC 9 LU 0 Chip                 OK          65 degrees C / 149 degrees F
        FPC 9 LU 1 TSen                 OK          60 degrees C / 140 degrees F
        FPC 9 LU 1 Chip                 OK          67 degrees C / 152 degrees F
        FPC 9 LU 2 TSen                 OK          60 degrees C / 140 degrees F
        FPC 9 LU 2 Chip                 OK          54 degrees C / 129 degrees F
        FPC 9 LU 3 TSen                 OK          60 degrees C / 140 degrees F
        FPC 9 LU 3 Chip                 OK          54 degrees C / 129 degrees F
```

```
FPC 9 MQ 0 TSen                OK          51 degrees C / 123 degrees F
FPC 9 MQ 0 Chip                OK          55 degrees C / 131 degrees F
FPC 9 MQ 1 TSen                OK          51 degrees C / 123 degrees F
FPC 9 MQ 1 Chip                OK          59 degrees C / 138 degrees F
FPC 9 MQ 2 TSen                OK          51 degrees C / 123 degrees F
FPC 9 MQ 2 Chip                OK          49 degrees C / 120 degrees F
FPC 9 MQ 3 TSen                OK          51 degrees C / 123 degrees F
FPC 9 MQ 3 Chip                OK          49 degrees C / 120 degrees F
FPC 10 Intake                  OK          44 degrees C / 111 degrees F
FPC 10 Exhaust A               OK          49 degrees C / 120 degrees F
FPC 10 Exhaust B               OK          55 degrees C / 131 degrees F
FPC 10 LU 0 TSen               OK          54 degrees C / 129 degrees F
FPC 10 LU 0 Chip               OK          55 degrees C / 131 degrees F
FPC 10 LU 1 TSen               OK          54 degrees C / 129 degrees F
FPC 10 LU 1 Chip               OK          59 degrees C / 138 degrees F
FPC 10 LU 2 TSen               OK          54 degrees C / 129 degrees F
FPC 10 LU 2 Chip               OK          52 degrees C / 125 degrees F
FPC 10 LU 3 TSen               OK          54 degrees C / 129 degrees F
FPC 10 LU 3 Chip               OK          51 degrees C / 123 degrees F
FPC 10 MQ 0 TSen               OK          48 degrees C / 118 degrees F
FPC 10 MQ 0 Chip               OK          49 degrees C / 120 degrees F
FPC 10 MQ 1 TSen               OK          48 degrees C / 118 degrees F
FPC 10 MQ 1 Chip               OK          52 degrees C / 125 degrees F
FPC 10 MQ 2 TSen               OK          48 degrees C / 118 degrees F
FPC 10 MQ 2 Chip               OK          47 degrees C / 116 degrees F
FPC 10 MQ 3 TSen               OK          48 degrees C / 118 degrees F
FPC 10 MQ 3 Chip               OK          47 degrees C / 116 degrees F
FPC 11 Intake                  OK          30 degrees C / 86 degrees F
FPC 11 Exhaust A               OK          35 degrees C / 95 degrees F
FPC 11 Exhaust B               OK          30 degrees C / 86 degrees F
FPC 11 LU 0 TSen               OK          57 degrees C / 134 degrees F
FPC 11 LU 0 Chip               OK          58 degrees C / 136 degrees F
FPC 11 LU 1 TSen               OK          57 degrees C / 134 degrees F
FPC 11 LU 1 Chip               OK          62 degrees C / 143 degrees F
FPC 11 LU 2 TSen               OK          57 degrees C / 134 degrees F
FPC 11 LU 2 Chip               OK          53 degrees C / 127 degrees F
FPC 11 LU 3 TSen               OK          57 degrees C / 134 degrees F
FPC 11 LU 3 Chip               OK          54 degrees C / 129 degrees F
FPC 11 MQ 0 TSen               OK          52 degrees C / 125 degrees F
FPC 11 MQ 0 Chip               OK          52 degrees C / 125 degrees F
FPC 11 MQ 1 TSen               OK          52 degrees C / 125 degrees F
FPC 11 MQ 1 Chip               OK          57 degrees C / 134 degrees F
FPC 11 MQ 2 TSen               OK          52 degrees C / 125 degrees F
```

```
    FPC 11 MQ 2 Chip              OK        48 degrees C / 118 degrees F
    FPC 11 MQ 3 TSen              OK        52 degrees C / 125 degrees F
    FPC 11 MQ 3 Chip              OK        52 degrees C / 125 degrees F
    FPC 12 Intake                 OK        40 degrees C / 104 degrees F
    FPC 12 Exhaust A              OK        47 degrees C / 116 degrees F
    FPC 12 Exhaust B              OK        52 degrees C / 125 degrees F
    FPC 12 LU 0 TSen              OK        51 degrees C / 123 degrees F
    FPC 12 LU 0 Chip              OK        52 degrees C / 125 degrees F
    FPC 12 LU 1 TSen              OK        51 degrees C / 123 degrees F
    FPC 12 LU 1 Chip              OK        55 degrees C / 131 degrees F
    FPC 12 LU 2 TSen              OK        51 degrees C / 123 degrees F
    FPC 12 LU 2 Chip              OK        47 degrees C / 116 degrees F
    FPC 12 LU 3 TSen              OK        51 degrees C / 123 degrees F
    FPC 12 LU 3 Chip              OK        50 degrees C / 122 degrees F
    FPC 12 MQ 0 TSen              OK        46 degrees C / 114 degrees F
    FPC 12 MQ 0 Chip              OK        46 degrees C / 114 degrees F
    FPC 12 MQ 1 TSen              OK        46 degrees C / 114 degrees F
    FPC 12 MQ 1 Chip              OK        50 degrees C / 122 degrees F
    FPC 12 MQ 2 TSen              OK        46 degrees C / 114 degrees F
    FPC 12 MQ 2 Chip              OK        44 degrees C / 111 degrees F
    FPC 12 MQ 3 TSen              OK        46 degrees C / 114 degrees F
    FPC 12 MQ 3 Chip              OK        46 degrees C / 114 degrees F
    FPC 13 Intake                 OK        40 degrees C / 104 degrees F
    FPC 13 Exhaust A              OK        48 degrees C / 118 degrees F
    FPC 13 Exhaust B              OK        52 degrees C / 125 degrees F
    FPC 13 LU 0 TSen              OK        51 degrees C / 123 degrees F
    FPC 13 LU 0 Chip              OK        52 degrees C / 125 degrees F
    FPC 13 LU 1 TSen              OK        51 degrees C / 123 degrees F
    FPC 13 LU 1 Chip              OK        55 degrees C / 131 degrees F
    FPC 13 LU 2 TSen              OK        51 degrees C / 123 degrees F
    FPC 13 LU 2 Chip              OK        48 degrees C / 118 degrees F
    FPC 13 LU 3 TSen              OK        51 degrees C / 123 degrees F
    FPC 13 LU 3 Chip              OK        48 degrees C / 118 degrees F
    FPC 13 MQ 0 TSen              OK        46 degrees C / 114 degrees F
    FPC 13 MQ 0 Chip              OK        46 degrees C / 114 degrees F
    FPC 13 MQ 1 TSen              OK        46 degrees C / 114 degrees F
    FPC 13 MQ 1 Chip              OK        50 degrees C / 122 degrees F
    FPC 13 MQ 2 TSen              OK        46 degrees C / 114 degrees F
    FPC 13 MQ 2 Chip              OK        44 degrees C / 111 degrees F
    FPC 13 MQ 3 TSen              OK        46 degrees C / 114 degrees F
    FPC 13 MQ 3 Chip              OK        46 degrees C / 114 degrees F
    FPC 14 Intake                 OK        40 degrees C / 104 degrees F
    FPC 14 Exhaust A              OK        50 degrees C / 122 degrees F
```

```
    FPC 14 Exhaust B              OK         51 degrees C / 123 degrees F
    FPC 14 LU 0 TSen              OK         50 degrees C / 122 degrees F
    FPC 14 LU 0 Chip              OK         50 degrees C / 122 degrees F
    FPC 14 LU 1 TSen              OK         50 degrees C / 122 degrees F
    FPC 14 LU 1 Chip              OK         54 degrees C / 129 degrees F
    FPC 14 LU 2 TSen              OK         50 degrees C / 122 degrees F
    FPC 14 LU 2 Chip              OK         47 degrees C / 116 degrees F
    FPC 14 LU 3 TSen              OK         50 degrees C / 122 degrees F
    FPC 14 LU 3 Chip              OK         49 degrees C / 120 degrees F
    FPC 14 MQ 0 TSen              OK         47 degrees C / 116 degrees F
    FPC 14 MQ 0 Chip              OK         46 degrees C / 114 degrees F
    FPC 14 MQ 1 TSen              OK         47 degrees C / 116 degrees F
    FPC 14 MQ 1 Chip              OK         51 degrees C / 123 degrees F
    FPC 14 MQ 2 TSen              OK         47 degrees C / 116 degrees F
    FPC 14 MQ 2 Chip              OK         45 degrees C / 113 degrees F
    FPC 14 MQ 3 TSen              OK         47 degrees C / 116 degrees F
    FPC 14 MQ 3 Chip              OK         48 degrees C / 118 degrees F
    FPC 15 Intake                 OK         44 degrees C / 111 degrees F
    FPC 15 Exhaust A              OK         49 degrees C / 120 degrees F
    FPC 15 Exhaust B              OK         60 degrees C / 140 degrees F
    FPC 15 LU 0 TSen              OK         50 degrees C / 122 degrees F
    FPC 15 LU 0 Chip              OK         56 degrees C / 132 degrees F
    FPC 15 LU 1 TSen              OK         50 degrees C / 122 degrees F
    FPC 15 LU 1 Chip              OK         50 degrees C / 122 degrees F
    FPC 15 LU 2 TSen              OK         50 degrees C / 122 degrees F
    FPC 15 LU 2 Chip              OK         58 degrees C / 136 degrees F
    FPC 15 LU 3 TSen              OK         50 degrees C / 122 degrees F
    FPC 15 LU 3 Chip              OK         63 degrees C / 145 degrees F
    FPC 15 XM 0 TSen              OK         50 degrees C / 122 degrees F
    FPC 15 XM 0 Chip              OK         56 degrees C / 132 degrees F
    FPC 15 XF 0 TSen              OK         50 degrees C / 122 degrees F
    FPC 15 XF 0 Chip              OK         68 degrees C / 154 degrees F
    FPC 15 PLX Switch TSen        OK         50 degrees C / 122 degrees F
    FPC 15 PLX Switch Chip        OK         56 degrees C / 132 degrees F
    FPC 16 Intake                 OK         42 degrees C / 107 degrees F
    FPC 16 Exhaust A              OK         51 degrees C / 123 degrees F
    FPC 16 Exhaust B              OK         53 degrees C / 127 degrees F
    FPC 16 LU 0 TSen              OK         51 degrees C / 123 degrees F
    FPC 16 LU 0 Chip              OK         52 degrees C / 125 degrees F
    FPC 16 LU 1 TSen              OK         51 degrees C / 123 degrees F
    FPC 16 LU 1 Chip              OK         55 degrees C / 131 degrees F
    FPC 16 LU 2 TSen              OK         51 degrees C / 123 degrees F
    FPC 16 LU 2 Chip              OK         48 degrees C / 118 degrees F
```

```
    FPC 16 LU 3 TSen              OK        51 degrees C / 123 degrees F
    FPC 16 LU 3 Chip              OK        49 degrees C / 120 degrees F
    FPC 16 MQ 0 TSen              OK        49 degrees C / 120 degrees F
    FPC 16 MQ 0 Chip              OK        48 degrees C / 118 degrees F
    FPC 16 MQ 1 TSen              OK        49 degrees C / 120 degrees F
    FPC 16 MQ 1 Chip              OK        53 degrees C / 127 degrees F
    FPC 16 MQ 2 TSen              OK        49 degrees C / 120 degrees F
    FPC 16 MQ 2 Chip              OK        46 degrees C / 114 degrees F
    FPC 16 MQ 3 TSen              OK        49 degrees C / 120 degrees F
    FPC 16 MQ 3 Chip              OK        49 degrees C / 120 degrees F
    FPC 17 Intake                 OK        43 degrees C / 109 degrees F
    FPC 17 Exhaust A              OK        51 degrees C / 123 degrees F
    FPC 17 Exhaust B              OK        55 degrees C / 131 degrees F
    FPC 17 LU 0 TSen              OK        54 degrees C / 129 degrees F
    FPC 17 LU 0 Chip              OK        57 degrees C / 134 degrees F
    FPC 17 LU 1 TSen              OK        54 degrees C / 129 degrees F
    FPC 17 LU 1 Chip              OK        60 degrees C / 140 degrees F
    FPC 17 LU 2 TSen              OK        54 degrees C / 129 degrees F
    FPC 17 LU 2 Chip              OK        53 degrees C / 127 degrees F
    FPC 17 LU 3 TSen              OK        54 degrees C / 129 degrees F
    FPC 17 LU 3 Chip              OK        53 degrees C / 127 degrees F
    FPC 17 MQ 0 TSen              OK        49 degrees C / 120 degrees F
    FPC 17 MQ 0 Chip              OK        50 degrees C / 122 degrees F
    FPC 17 MQ 1 TSen              OK        49 degrees C / 120 degrees F
    FPC 17 MQ 1 Chip              OK        54 degrees C / 129 degrees F
    FPC 17 MQ 2 TSen              OK        49 degrees C / 120 degrees F
    FPC 17 MQ 2 Chip              OK        47 degrees C / 116 degrees F
    FPC 17 MQ 3 TSen              OK        49 degrees C / 120 degrees F
    FPC 17 MQ 3 Chip              OK        51 degrees C / 123 degrees F
    FPC 18 Intake                 OK        44 degrees C / 111 degrees F
    FPC 18 Exhaust A              OK        53 degrees C / 127 degrees F
    FPC 18 Exhaust B              OK        57 degrees C / 134 degrees F
    FPC 18 LU 0 TSen              OK        56 degrees C / 132 degrees F
    FPC 18 LU 0 Chip              OK        57 degrees C / 134 degrees F
    FPC 18 LU 1 TSen              OK        56 degrees C / 132 degrees F
    FPC 18 LU 1 Chip              OK        62 degrees C / 143 degrees F
    FPC 18 LU 2 TSen              OK        56 degrees C / 132 degrees F
    FPC 18 LU 2 Chip              OK        53 degrees C / 127 degrees F
    FPC 18 LU 3 TSen              OK        56 degrees C / 132 degrees F
    FPC 18 LU 3 Chip              OK        55 degrees C / 131 degrees F
    FPC 18 MQ 0 TSen              OK        51 degrees C / 123 degrees F
    FPC 18 MQ 0 Chip              OK        54 degrees C / 129 degrees F
    FPC 18 MQ 1 TSen              OK        51 degrees C / 123 degrees F
```

```
FPC 18 MQ 1 Chip          OK        58 degrees C / 136 degrees F
FPC 18 MQ 2 TSen          OK        51 degrees C / 123 degrees F
FPC 18 MQ 2 Chip          OK        50 degrees C / 122 degrees F
FPC 18 MQ 3 TSen          OK        51 degrees C / 123 degrees F
FPC 18 MQ 3 Chip          OK        53 degrees C / 127 degrees F
FPC 19 Intake             OK        48 degrees C / 118 degrees F
FPC 19 Exhaust A          OK        56 degrees C / 132 degrees F
FPC 19 Exhaust B          OK        64 degrees C / 147 degrees F
FPC 19 LU 0 TSen          OK        63 degrees C / 145 degrees F
FPC 19 LU 0 Chip          OK        64 degrees C / 147 degrees F
FPC 19 LU 1 TSen          OK        63 degrees C / 145 degrees F
FPC 19 LU 1 Chip          OK        70 degrees C / 158 degrees F
FPC 19 LU 2 TSen          OK        63 degrees C / 145 degrees F
FPC 19 LU 2 Chip          OK        61 degrees C / 141 degrees F
FPC 19 LU 3 TSen          OK        63 degrees C / 145 degrees F
FPC 19 LU 3 Chip          OK        62 degrees C / 143 degrees F
FPC 19 MQ 0 TSen          OK        56 degrees C / 132 degrees F
FPC 19 MQ 0 Chip          OK        60 degrees C / 140 degrees F
FPC 19 MQ 1 TSen          OK        56 degrees C / 132 degrees F
FPC 19 MQ 1 Chip          OK        62 degrees C / 143 degrees F
FPC 19 MQ 2 TSen          OK        56 degrees C / 132 degrees F
FPC 19 MQ 2 Chip          OK        56 degrees C / 132 degrees F
FPC 19 MQ 3 TSen          OK        56 degrees C / 132 degrees F
FPC 19 MQ 3 Chip          OK        57 degrees C / 134 degrees F
ADC 0 Intake              OK        40 degrees C / 104 degrees F
ADC 0 Exhaust             OK        52 degrees C / 125 degrees F
ADC 0 ADC-XF1             OK        59 degrees C / 138 degrees F
ADC 0 ADC-XF0             OK        66 degrees C / 150 degrees F
ADC 1 Intake              OK        38 degrees C / 100 degrees F
ADC 1 Exhaust             OK        50 degrees C / 122 degrees F
ADC 1 ADC-XF1             OK        59 degrees C / 138 degrees F
ADC 1 ADC-XF0             OK        63 degrees C / 145 degrees F
ADC 2 Intake              OK        37 degrees C / 98 degrees F
ADC 2 Exhaust             OK        52 degrees C / 125 degrees F
ADC 2 ADC-XF1             OK        53 degrees C / 127 degrees F
ADC 2 ADC-XF0             OK        61 degrees C / 141 degrees F
ADC 3 Intake              OK        40 degrees C / 104 degrees F
ADC 3 Exhaust             OK        51 degrees C / 123 degrees F
ADC 3 ADC-XF1             OK        61 degrees C / 141 degrees F
ADC 3 ADC-XF0             OK        64 degrees C / 147 degrees F
ADC 4 Intake              OK        39 degrees C / 102 degrees F
ADC 4 Exhaust             OK        51 degrees C / 123 degrees F
ADC 4 ADC-XF1             OK        60 degrees C / 140 degrees F
```

```
ADC 4 ADC-XF0                    OK          63 degrees C / 145 degrees F
ADC 5 Intake                     OK          38 degrees C / 100 degrees F
ADC 5 Exhaust                    OK          54 degrees C / 129 degrees F
ADC 5 ADC-XF1                    OK          56 degrees C / 132 degrees F
ADC 5 ADC-XF0                    OK          67 degrees C / 152 degrees F
ADC 6 Intake                     OK          39 degrees C / 102 degrees F
ADC 6 Exhaust                    OK          52 degrees C / 125 degrees F
ADC 6 ADC-XF1                    OK          59 degrees C / 138 degrees F
ADC 6 ADC-XF0                    OK          66 degrees C / 150 degrees F
ADC 7 Intake                     OK          39 degrees C / 102 degrees F
ADC 7 Exhaust                    OK          54 degrees C / 129 degrees F
ADC 7 ADC-XF1                    OK          62 degrees C / 143 degrees F
ADC 7 ADC-XF0                    OK          70 degrees C / 158 degrees F
ADC 8 Intake                     OK          39 degrees C / 102 degrees F
ADC 8 Exhaust                    OK          52 degrees C / 125 degrees F
ADC 8 ADC-XF1                    OK          61 degrees C / 141 degrees F
ADC 8 ADC-XF0                    OK          65 degrees C / 149 degrees F
ADC 9 Intake                     OK          41 degrees C / 105 degrees F
ADC 9 Exhaust                    OK          51 degrees C / 123 degrees F
ADC 9 ADC-XF1                    OK          63 degrees C / 145 degrees F
ADC 9 ADC-XF0                    OK          63 degrees C / 145 degrees F
ADC 10 Intake                    OK          48 degrees C / 118 degrees F
ADC 10 Exhaust                   OK          53 degrees C / 127 degrees F
ADC 10 ADC-XF1                   OK          67 degrees C / 152 degrees F
ADC 10 ADC-XF0                   OK          66 degrees C / 150 degrees F
ADC 12 Intake                    OK          49 degrees C / 120 degrees F
ADC 12 Exhaust                   OK          54 degrees C / 129 degrees F
ADC 12 ADC-XF1                   OK          67 degrees C / 152 degrees F
ADC 12 ADC-XF0                   OK          67 degrees C / 152 degrees F
ADC 13 Intake                    OK          49 degrees C / 120 degrees F
ADC 13 Exhaust                   OK          57 degrees C / 134 degrees F
ADC 13 ADC-XF1                   OK          66 degrees C / 150 degrees F
ADC 13 ADC-XF0                   OK          69 degrees C / 156 degrees F
ADC 14 Intake                    OK          51 degrees C / 123 degrees F
ADC 14 Exhaust                   OK          59 degrees C / 138 degrees F
ADC 14 ADC-XF1                   OK          69 degrees C / 156 degrees F
ADC 14 ADC-XF0                   OK          74 degrees C / 165 degrees F
ADC 15 Intake                    OK          50 degrees C / 122 degrees F
ADC 15 Exhaust                   OK          59 degrees C / 138 degrees F
ADC 15 ADC-XF1                   OK          68 degrees C / 154 degrees F
ADC 15 ADC-XF0                   OK          69 degrees C / 156 degrees F
ADC 16 Intake                    OK          52 degrees C / 125 degrees F
ADC 16 Exhaust                   OK          58 degrees C / 136 degrees F
```

```
       ADC 16 ADC-XF1                OK         68 degrees C / 154 degrees F
       ADC 16 ADC-XF0                OK         70 degrees C / 158 degrees F
       ADC 17 Intake                 OK         52 degrees C / 125 degrees F
       ADC 17 Exhaust                OK         59 degrees C / 138 degrees F
       ADC 17 ADC-XF1                OK         69 degrees C / 156 degrees F
       ADC 17 ADC-XF0                OK         71 degrees C / 159 degrees F
       ADC 18 Intake                 OK         53 degrees C / 127 degrees F
       ADC 18 Exhaust                OK         59 degrees C / 138 degrees F
       ADC 18 ADC-XF1                OK         68 degrees C / 154 degrees F
       ADC 18 ADC-XF0                OK         73 degrees C / 163 degrees F
       ADC 19 Intake                 OK         50 degrees C / 122 degrees F
       ADC 19 Exhaust                OK         59 degrees C / 138 degrees F
       ADC 19 ADC-XF1                OK         68 degrees C / 154 degrees F
       ADC 19 ADC-XF0                OK         72 degrees C / 161 degrees F
Fans   Fan Tray 0 Fan 1              OK         7440 RPM
       Fan Tray 0 Fan 2              OK         7200 RPM
       Fan Tray 0 Fan 3              OK         6960 RPM
       Fan Tray 0 Fan 4              OK         7200 RPM
       Fan Tray 0 Fan 5              OK         7080 RPM
       Fan Tray 0 Fan 6              OK         6840 RPM
       Fan Tray 1 Fan 1              OK         6840 RPM
       Fan Tray 1 Fan 2              OK         6960 RPM
       Fan Tray 1 Fan 3              OK         6960 RPM
       Fan Tray 1 Fan 4              OK         7080 RPM
       Fan Tray 1 Fan 5              OK         6960 RPM
       Fan Tray 1 Fan 6              OK         6960 RPM
       Fan Tray 2 Fan 1              OK         8640 RPM
       Fan Tray 2 Fan 2              OK         8640 RPM
       Fan Tray 2 Fan 3              OK         8760 RPM
       Fan Tray 2 Fan 4              OK         8760 RPM
       Fan Tray 2 Fan 5              OK         8640 RPM
       Fan Tray 2 Fan 6              OK         8640 RPM
       Fan Tray 3 Fan 1              OK         8520 RPM
       Fan Tray 3 Fan 2              OK         8520 RPM
       Fan Tray 3 Fan 3              OK         8640 RPM
       Fan Tray 3 Fan 4              OK         8640 RPM
       Fan Tray 3 Fan 5              OK         8520 RPM
       Fan Tray 3 Fan 6              OK         8520 RPM
```

show chassis environment (MX2020 Router with MPC5EQ and MPC6E)

```
Class Item                          Status     Measurement
Temp  PSM 0                         OK         32 degrees C / 89 degrees F
      PSM 1                         OK         32 degrees C / 89 degrees F
      PSM 2                         OK         32 degrees C / 89 degrees F
      PSM 3                         OK         32 degrees C / 89 degrees F
      PSM 4                         OK         32 degrees C / 89 degrees F
      PSM 5                         OK         33 degrees C / 91 degrees F
      PSM 6                         OK         32 degrees C / 89 degrees F
      PSM 7                         OK         32 degrees C / 89 degrees F
      PSM 8                         OK         32 degrees C / 89 degrees F
      PSM 9                         Absent
      PSM 10                        Absent
      PSM 11                        Absent
      PSM 12                        OK         33 degrees C / 91 degrees F
      PSM 13                        OK         33 degrees C / 91 degrees F
      PSM 14                        OK         34 degrees C / 93 degrees F
      PSM 15                        OK         34 degrees C / 93 degrees F
      PSM 16                        OK         33 degrees C / 91 degrees F
      PSM 17                        OK         33 degrees C / 91 degrees F
      PDM 0                         OK
      PDM 1                         OK
      PDM 2                         OK
      PDM 3                         OK
      CB 0 IntakeA-Zone0            OK         34 degrees C / 93 degrees F
      CB 0 IntakeB-Zone1            OK         26 degrees C / 78 degrees F
      CB 0 IntakeC-Zone0            OK         38 degrees C / 100 degrees F
      CB 0 ExhaustA-Zone0           OK         34 degrees C / 93 degrees F
      CB 0 ExhaustB-Zone1           OK         27 degrees C / 80 degrees F
      CB 0 TCBC-Zone0               OK         32 degrees C / 89 degrees F
      CB 1 IntakeA-Zone0            OK         24 degrees C / 75 degrees F
      CB 1 IntakeB-Zone1            OK         22 degrees C / 71 degrees F
      CB 1 IntakeC-Zone0            OK         34 degrees C / 93 degrees F
      CB 1 ExhaustA-Zone0           OK         31 degrees C / 87 degrees F
      CB 1 ExhaustB-Zone1           OK         24 degrees C / 75 degrees F
      CB 1 TCBC-Zone0               OK         27 degrees C / 80 degrees F
      SPMB 0 Intake                 OK         25 degrees C / 77 degrees F
      SPMB 1 Intake                 OK         23 degrees C / 73 degrees F
      Routing Engine 0              OK         28 degrees C / 82 degrees F
      Routing Engine 0 CPU          OK         25 degrees C / 77 degrees F
      Routing Engine 1              OK         25 degrees C / 77 degrees F
```

```
     Routing Engine 1 CPU          OK         24 degrees C / 75 degrees F
     SFB 0 Intake-Zone0            OK         45 degrees C / 113 degrees F
     SFB 0 Exhaust-Zone1           OK         34 degrees C / 93 degrees F
     SFB 0 IntakeA-Zone0           OK         32 degrees C / 89 degrees F
     SFB 0 IntakeB-Zone1           OK         28 degrees C / 82 degrees F
     SFB 0 Exhaust-Zone0           OK         36 degrees C / 96 degrees F
     SFB 0 SFB-XF2-Zone1           OK         46 degrees C / 114 degrees F
     SFB 0 SFB-XF1-Zone0           OK         48 degrees C / 118 degrees F
     SFB 0 SFB-XF0-Zone0           OK         60 degrees C / 140 degrees F
     SFB 1 Intake-Zone0            OK         44 degrees C / 111 degrees F
     SFB 1 Exhaust-Zone1           OK         34 degrees C / 93 degrees F
     SFB 1 IntakeA-Zone0           OK         35 degrees C / 95 degrees F
     SFB 1 IntakeB-Zone1           OK         27 degrees C / 80 degrees F
     SFB 1 Exhaust-Zone0           OK         37 degrees C / 98 degrees F
     SFB 1 SFB-XF2-Zone1           OK         47 degrees C / 116 degrees F
     SFB 1 SFB-XF1-Zone0           OK         49 degrees C / 120 degrees F
SFB 1 SFB-XF0-Zone0            OK         56 degrees C / 132 degrees F
     SFB 2 Intake-Zone0            OK         41 degrees C / 105 degrees F
     SFB 2 Exhaust-Zone1           OK         34 degrees C / 93 degrees F
     SFB 2 IntakeA-Zone0           OK         35 degrees C / 95 degrees F
     SFB 2 IntakeB-Zone1           OK         28 degrees C / 82 degrees F
     SFB 2 Exhaust-Zone0           OK         37 degrees C / 98 degrees F
     SFB 2 SFB-XF2-Zone1           OK         47 degrees C / 116 degrees F
     SFB 2 SFB-XF1-Zone0           OK         55 degrees C / 131 degrees F
     SFB 2 SFB-XF0-Zone0           OK         55 degrees C / 131 degrees F
     SFB 3 Intake-Zone0            OK         43 degrees C / 109 degrees F
     SFB 3 Exhaust-Zone1           OK         33 degrees C / 91 degrees F
     SFB 3 IntakeA-Zone0           OK         35 degrees C / 95 degrees F
     SFB 3 IntakeB-Zone1           OK         27 degrees C / 80 degrees F
     SFB 3 Exhaust-Zone0           OK         36 degrees C / 96 degrees F
     SFB 3 SFB-XF2-Zone1           OK         46 degrees C / 114 degrees F
     SFB 3 SFB-XF1-Zone0           OK         46 degrees C / 114 degrees F
     SFB 3 SFB-XF0-Zone0           OK         57 degrees C / 134 degrees F
     SFB 4 Intake-Zone0            OK         36 degrees C / 96 degrees F
     SFB 4 Exhaust-Zone1           OK         32 degrees C / 89 degrees F
     SFB 4 IntakeA-Zone0           OK         31 degrees C / 87 degrees F
     SFB 4 IntakeB-Zone1           OK         26 degrees C / 78 degrees F
     SFB 4 Exhaust-Zone0           OK         32 degrees C / 89 degrees F
     SFB 4 SFB-XF2-Zone1           OK         44 degrees C / 111 degrees F
     SFB 4 SFB-XF1-Zone0           OK         45 degrees C / 113 degrees F
     SFB 4 SFB-XF0-Zone0           OK         52 degrees C / 125 degrees F
     SFB 5 Intake-Zone0            OK         31 degrees C / 87 degrees F
     SFB 5 Exhaust-Zone1           OK         30 degrees C / 86 degrees F
```

```
     SFB 5 IntakeA-Zone0            OK         26 degrees C / 78 degrees F
     SFB 5 IntakeB-Zone1            OK         24 degrees C / 75 degrees F
     SFB 5 Exhaust-Zone0            OK         29 degrees C / 84 degrees F
     SFB 5 SFB-XF2-Zone1            OK         43 degrees C / 109 degrees F
     SFB 5 SFB-XF1-Zone0            OK         47 degrees C / 116 degrees F
     SFB 5 SFB-XF0-Zone0            OK         49 degrees C / 120 degrees F
     SFB 6 Intake-Zone0             OK         30 degrees C / 86 degrees F
     SFB 6 Exhaust-Zone1            OK         29 degrees C / 84 degrees F
     SFB 6 IntakeA-Zone0            OK         25 degrees C / 77 degrees F
     SFB 6 IntakeB-Zone1            OK         24 degrees C / 75 degrees F
     SFB 6 Exhaust-Zone0            OK         29 degrees C / 84 degrees F
     SFB 6 SFB-XF2-Zone1            OK         43 degrees C / 109 degrees F
     SFB 6 SFB-XF1-Zone0            OK         44 degrees C / 111 degrees F
     SFB 6 SFB-XF0-Zone0            OK         45 degrees C / 113 degrees F
     SFB 7 Intake-Zone0             OK         31 degrees C / 87 degrees F
     SFB 7 Exhaust-Zone1            OK         30 degrees C / 86 degrees F
     SFB 7 IntakeA-Zone0            OK         26 degrees C / 78 degrees F
     SFB 7 IntakeB-Zone1            OK         24 degrees C / 75 degrees F
     SFB 7 Exhaust-Zone0            OK         28 degrees C / 82 degrees F
     SFB 7 SFB-XF2-Zone1            OK         50 degrees C / 122 degrees F
     SFB 7 SFB-XF1-Zone0            OK         43 degrees C / 109 degrees F
     SFB 7 SFB-XF0-Zone0            OK         47 degrees C / 116 degrees F
     FPC 0 Intake                   OK         31 degrees C / 87 degrees F
     FPC 0 Exhaust A                OK         49 degrees C / 120 degrees F
     FPC 0 Exhaust B                OK         43 degrees C / 109 degrees F
     FPC 0 XL TSen                  OK         42 degrees C / 107 degrees F
     FPC 0 XL Chip                  OK         46 degrees C / 114 degrees F
     FPC 0 XL_XR0 TSen              OK         42 degrees C / 107 degrees F
     FPC 0 XL_XR0 Chip              OK         48 degrees C / 118 degrees F
FPC 0 XL_XR1 TSen              OK         42 degrees C / 107 degrees F
     FPC 0 XL_XR1 Chip              OK         48 degrees C / 118 degrees F
     FPC 0 XQ TSen                  OK         42 degrees C / 107 degrees F
     FPC 0 XQ Chip                  OK         44 degrees C / 111 degrees F
     FPC 0 XQ_XR0 TSen              OK         42 degrees C / 107 degrees F
     FPC 0 XQ_XR0 Chip              OK         57 degrees C / 134 degrees F
     FPC 0 XQ_XR1 TSen              OK         42 degrees C / 107 degrees F
     FPC 0 XQ_XR1 Chip              OK         55 degrees C / 131 degrees F
     FPC 0 XM 0 TSen                OK         48 degrees C / 118 degrees F
     FPC 0 XM 0 Chip                OK         62 degrees C / 143 degrees F
     FPC 0 XM 1 TSen                OK         48 degrees C / 118 degrees F
     FPC 0 XM 1 Chip                OK         44 degrees C / 111 degrees F
     FPC 0 PLX PCIe Switch TSe      OK         48 degrees C / 118 degrees F
     FPC 0 PLX PCIe Switch Chi      OK         57 degrees C / 134 degrees F
```

```
    FPC 1 Intake                  OK         29 degrees C / 84 degrees F
    FPC 1 Exhaust A               OK         36 degrees C / 96 degrees F
    FPC 1 Exhaust B               OK         44 degrees C / 111 degrees F
    FPC 1 LU 0 TSen               OK         38 degrees C / 100 degrees F
    FPC 1 LU 0 Chip               OK         45 degrees C / 113 degrees F
    FPC 1 LU 1 TSen               OK         38 degrees C / 100 degrees F
    FPC 1 LU 1 Chip               OK         38 degrees C / 100 degrees F
    FPC 1 LU 2 TSen               OK         38 degrees C / 100 degrees F
    FPC 1 LU 2 Chip               OK         42 degrees C / 107 degrees F
    FPC 1 LU 3 TSen               OK         38 degrees C / 100 degrees F
    FPC 1 LU 3 Chip               OK         47 degrees C / 116 degrees F
    FPC 1 XM 0 TSen               OK         38 degrees C / 100 degrees F
    FPC 1 XM 0 Chip               OK         44 degrees C / 111 degrees F
    FPC 1 XF 0 TSen               OK         38 degrees C / 100 degrees F
    FPC 1 XF 0 Chip               OK         54 degrees C / 129 degrees F
    FPC 1 PLX Switch TSen         OK         38 degrees C / 100 degrees F
    FPC 1 PLX Switch Chip         OK         41 degrees C / 105 degrees F
    FPC 2 Intake                  OK         28 degrees C / 82 degrees F
    FPC 2 Exhaust A               OK         28 degrees C / 82 degrees F
    FPC 2 Exhaust B               OK         28 degrees C / 82 degrees F
    FPC 2 LU 0 TSen               OK         40 degrees C / 104 degrees F
    FPC 2 LU 0 Chip               OK         40 degrees C / 104 degrees F
    FPC 2 LU 1 TSen               OK         40 degrees C / 104 degrees F
    FPC 2 LU 1 Chip               OK         41 degrees C / 105 degrees F
    FPC 2 LU 2 TSen               OK         40 degrees C / 104 degrees F
    FPC 2 LU 2 Chip               OK         34 degrees C / 93 degrees F
    FPC 2 LU 3 TSen               OK         40 degrees C / 104 degrees F
    FPC 2 LU 3 Chip               OK         38 degrees C / 100 degrees F
    FPC 2 XM 0 TSen               OK         40 degrees C / 104 degrees F
    FPC 2 XM 0 Chip               OK         47 degrees C / 116 degrees F
    FPC 2 XM 1 TSen               OK         40 degrees C / 104 degrees F
    FPC 2 XM 1 Chip               OK         42 degrees C / 107 degrees F
    FPC 2 PLX Switch TSen         OK         40 degrees C / 104 degrees F
    FPC 2 PLX Switch Chip         OK         39 degrees C / 102 degrees F
    FPC 3 Intake                  OK         27 degrees C / 80 degrees F
    FPC 3 Exhaust A               OK         38 degrees C / 100 degrees F
    FPC 3 Exhaust B               OK         31 degrees C / 87 degrees F
    FPC 3 QX 0 TSen               OK         38 degrees C / 100 degrees F
    FPC 3 QX 0 Chip               OK         42 degrees C / 107 degrees F
    FPC 3 LU 0 TCAM TSen          OK         38 degrees C / 100 degrees F
    FPC 3 LU 0 TCAM Chip          OK         43 degrees C / 109 degrees F
    FPC 3 LU 0 TSen               OK         38 degrees C / 100 degrees F
FPC 3 LU 0 Chip               OK         42 degrees C / 107 degrees F
```

```
    FPC 3 MQ 0 TSen                OK          38 degrees C / 100 degrees F
    FPC 3 MQ 0 Chip                OK          39 degrees C / 102 degrees F
    FPC 3 QX 1 TSen                OK          32 degrees C / 89 degrees F
    FPC 3 QX 1 Chip                OK          36 degrees C / 96 degrees F
    FPC 3 LU 1 TCAM TSen           OK          32 degrees C / 89 degrees F
    FPC 3 LU 1 TCAM Chip           OK          35 degrees C / 95 degrees F
    FPC 3 LU 1 TSen                OK          32 degrees C / 89 degrees F
    FPC 3 LU 1 Chip                OK          37 degrees C / 98 degrees F
    FPC 3 MQ 1 TSen                OK          32 degrees C / 89 degrees F
    FPC 3 MQ 1 Chip                OK          36 degrees C / 96 degrees F
    FPC 4 Intake                   OK          29 degrees C / 84 degrees F
    FPC 4 Exhaust A                OK          36 degrees C / 96 degrees F
    FPC 4 Exhaust B                OK          40 degrees C / 104 degrees F
    FPC 4 XL TSen                  OK          39 degrees C / 102 degrees F
    FPC 4 XL Chip                  OK          42 degrees C / 107 degrees F
    FPC 4 XL_XR0 TSen              OK          39 degrees C / 102 degrees F
    FPC 4 XL_XR0 Chip              OK          45 degrees C / 113 degrees F
    FPC 4 XL_XR1 TSen              OK          39 degrees C / 102 degrees F
    FPC 4 XL_XR1 Chip              OK          46 degrees C / 114 degrees F
    FPC 4 XQ TSen                  OK          39 degrees C / 102 degrees F
    FPC 4 XQ Chip                  OK          42 degrees C / 107 degrees F
    FPC 4 XQ_XR0 TSen              OK          39 degrees C / 102 degrees F
    FPC 4 XQ_XR0 Chip              OK          54 degrees C / 129 degrees F
    FPC 4 XQ_XR1 TSen              OK          39 degrees C / 102 degrees F
    FPC 4 XQ_XR1 Chip              OK          53 degrees C / 127 degrees F
    FPC 4 XM 0 TSen                OK          45 degrees C / 113 degrees F
    FPC 4 XM 0 Chip                OK          59 degrees C / 138 degrees F
    FPC 4 XM 1 TSen                OK          45 degrees C / 113 degrees F
    FPC 4 XM 1 Chip                OK          41 degrees C / 105 degrees F
    FPC 4 PLX PCIe Switch TSe      OK          45 degrees C / 113 degrees F
    FPC 4 PLX PCIe Switch Chi      OK          58 degrees C / 136 degrees F
    FPC 5 Intake                   OK          29 degrees C / 84 degrees F
    FPC 5 Exhaust A                OK          33 degrees C / 91 degrees F
    FPC 5 Exhaust B                OK          39 degrees C / 102 degrees F
    FPC 5 LU 0 TSen                OK          40 degrees C / 104 degrees F
    FPC 5 LU 0 Chip                OK          40 degrees C / 104 degrees F
    FPC 5 LU 1 TSen                OK          40 degrees C / 104 degrees F
    FPC 5 LU 1 Chip                OK          45 degrees C / 113 degrees F
    FPC 5 LU 2 TSen                OK          40 degrees C / 104 degrees F
    FPC 5 LU 2 Chip                OK          40 degrees C / 104 degrees F
    FPC 5 LU 3 TSen                OK          40 degrees C / 104 degrees F
    FPC 5 LU 3 Chip                OK          46 degrees C / 114 degrees F
    FPC 5 MQ 0 TSen                OK          32 degrees C / 89 degrees F
```

```
      FPC 5 MQ 0 Chip                OK         33 degrees C / 91 degrees F
      FPC 5 MQ 1 TSen                OK         32 degrees C / 89 degrees F
      FPC 5 MQ 1 Chip                OK         35 degrees C / 95 degrees F
      FPC 5 MQ 2 TSen                OK         32 degrees C / 89 degrees F
      FPC 5 MQ 2 Chip                OK         32 degrees C / 89 degrees F
      FPC 5 MQ 3 TSen                OK         32 degrees C / 89 degrees F
      FPC 5 MQ 3 Chip                OK         32 degrees C / 89 degrees F
      FPC 9 Intake                   OK         25 degrees C / 77 degrees F
      FPC 9 Exhaust A                OK         37 degrees C / 98 degrees F
      FPC 9 Exhaust B                OK         40 degrees C / 104 degrees F
      FPC 9 XL 0 TSen                OK         40 degrees C / 104 degrees F
...
```

show chassis environment (MX2010 Router)

```
user@host> show chassis environment
Class Item                           Status     Measurement
Temp  PSM 0                          OK         7 degrees C / 44 degrees F
      PSM 1                          OK         7 degrees C / 44 degrees F
      PSM 2                          OK         7 degrees C / 44 degrees F
      PSM 3                          OK         6 degrees C / 42 degrees F
      PSM 4                          OK         6 degrees C / 42 degrees F
      PSM 5                          OK         6 degrees C / 42 degrees F
      PSM 6                          OK         6 degrees C / 42 degrees F
      PSM 7                          OK         7 degrees C / 44 degrees F
      PSM 8                          OK         7 degrees C / 44 degrees F
      PDM 0                          OK
      PDM 1                          Absent
      CB 0 IntakeA-Zone0             OK         14 degrees C / 57 degrees F
      CB 0 IntakeB-Zone1             OK         7 degrees C / 44 degrees F
      CB 0 IntakeC-Zone0             OK         22 degrees C / 71 degrees F
      CB 0 ExhaustA-Zone0            OK         14 degrees C / 57 degrees F
      CB 0 ExhaustB-Zone1            OK         9 degrees C / 48 degrees F
      CB 0 TCBC-Zone0                OK         11 degrees C / 51 degrees F
      CB 1 IntakeA-Zone0             OK         9 degrees C / 48 degrees F
      CB 1 IntakeB-Zone1             OK         5 degrees C / 41 degrees F
      CB 1 IntakeC-Zone0             OK         20 degrees C / 68 degrees F
      CB 1 ExhaustA-Zone0            OK         12 degrees C / 53 degrees F
      CB 1 ExhaustB-Zone1            OK         7 degrees C / 44 degrees F
      CB 1 TCBC-Zone0                OK         10 degrees C / 50 degrees F
      SPMB 0 Intake                  OK         5 degrees C / 41 degrees F
```

```
      SPMB 1 Intake                  OK         4 degrees C / 39 degrees F
      Routing Engine 0               OK         9 degrees C / 48 degrees F
      Routing Engine 0 CPU           OK         9 degrees C / 48 degrees F
      Routing Engine 1               OK         6 degrees C / 42 degrees F
      Routing Engine 1 CPU           OK         6 degrees C / 42 degrees F
      SFB 0 Intake-Zone0             OK         26 degrees C / 78 degrees F
      SFB 0 Exhaust-Zone1            OK         17 degrees C / 62 degrees F
      SFB 0 IntakeA-Zone0            OK         16 degrees C / 60 degrees F
      SFB 0 IntakeB-Zone1            OK         11 degrees C / 51 degrees F
      SFB 0 Exhaust-Zone0            OK         18 degrees C / 64 degrees F
      SFB 0 SFB-XF2-Zone1            OK         25 degrees C / 77 degrees F
      SFB 0 SFB-XF1-Zone0            OK         23 degrees C / 73 degrees F
      SFB 0 SFB-XF0-Zone0            OK         33 degrees C / 91 degrees F
      SFB 1 Intake-Zone0             OK         27 degrees C / 80 degrees F
      SFB 1 Exhaust-Zone1            OK         15 degrees C / 59 degrees F
      SFB 1 IntakeA-Zone0            OK         20 degrees C / 68 degrees F
      SFB 1 IntakeB-Zone1            OK         10 degrees C / 50 degrees F
      SFB 1 Exhaust-Zone0            OK         19 degrees C / 66 degrees F
      SFB 1 SFB-XF2-Zone1            OK         26 degrees C / 78 degrees F
      SFB 1 SFB-XF1-Zone0            OK         27 degrees C / 80 degrees F
      SFB 1 SFB-XF0-Zone0            OK         32 degrees C / 89 degrees F
      SFB 2 Intake-Zone0             OK         21 degrees C / 69 degrees F
      SFB 2 Exhaust-Zone1            OK         13 degrees C / 55 degrees F
      SFB 2 IntakeA-Zone0            OK         18 degrees C / 64 degrees F
      SFB 2 IntakeB-Zone1            OK         9 degrees C / 48 degrees F
      SFB 2 Exhaust-Zone0            OK         16 degrees C / 60 degrees F
      SFB 2 SFB-XF2-Zone1            OK         24 degrees C / 75 degrees F
      SFB 2 SFB-XF1-Zone0            OK         21 degrees C / 69 degrees F
      SFB 2 SFB-XF0-Zone0            OK         26 degrees C / 78 degrees F
      SFB 4 Intake-Zone0             OK         28 degrees C / 82 degrees F
      SFB 4 Exhaust-Zone1            OK         16 degrees C / 60 degrees F
SFB 4 IntakeA-Zone0            OK         18 degrees C / 64 degrees F
      SFB 4 IntakeB-Zone1            OK         11 degrees C / 51 degrees F
      SFB 4 Exhaust-Zone0            OK         19 degrees C / 66 degrees F
      SFB 4 SFB-XF2-Zone1            OK         27 degrees C / 80 degrees F
      SFB 4 SFB-XF1-Zone0            OK         27 degrees C / 80 degrees F
      SFB 4 SFB-XF0-Zone0            OK         32 degrees C / 89 degrees F
      SFB 5 Intake-Zone0             OK         22 degrees C / 71 degrees F
      SFB 5 Exhaust-Zone1            OK         14 degrees C / 57 degrees F
      SFB 5 IntakeA-Zone0            OK         18 degrees C / 64 degrees F
      SFB 5 IntakeB-Zone1            OK         10 degrees C / 50 degrees F
      SFB 5 Exhaust-Zone0            OK         17 degrees C / 62 degrees F
      SFB 5 SFB-XF2-Zone1            OK         22 degrees C / 71 degrees F
```

```
    SFB 5 SFB-XF1-Zone0          OK        29 degrees C / 84 degrees F
    SFB 5 SFB-XF0-Zone0          OK        27 degrees C / 80 degrees F
    SFB 6 Intake-Zone0           OK        27 degrees C / 80 degrees F
    SFB 6 Exhaust-Zone1          OK        13 degrees C / 55 degrees F
    SFB 6 IntakeA-Zone0          OK        19 degrees C / 66 degrees F
    SFB 6 IntakeB-Zone1          OK        10 degrees C / 50 degrees F
    SFB 6 Exhaust-Zone0          OK        20 degrees C / 68 degrees F
    SFB 6 SFB-XF2-Zone1          OK        24 degrees C / 75 degrees F
    SFB 6 SFB-XF1-Zone0          OK        32 degrees C / 89 degrees F
    SFB 6 SFB-XF0-Zone0          OK        33 degrees C / 91 degrees F
    SFB 7 Intake-Zone0           OK        25 degrees C / 77 degrees F
    SFB 7 Exhaust-Zone1          OK        13 degrees C / 55 degrees F
    SFB 7 IntakeA-Zone0          OK        14 degrees C / 57 degrees F
    SFB 7 IntakeB-Zone1          OK        8 degrees C / 46 degrees F
    SFB 7 Exhaust-Zone0          OK        17 degrees C / 62 degrees F
    SFB 7 SFB-XF2-Zone1          OK        21 degrees C / 69 degrees F
    SFB 7 SFB-XF1-Zone0          OK        21 degrees C / 69 degrees F
    SFB 7 SFB-XF0-Zone0          OK        33 degrees C / 91 degrees F
    FPC 0 Intake                 OK        13 degrees C / 55 degrees F
    FPC 0 Exhaust A              OK        13 degrees C / 55 degrees F
    FPC 0 Exhaust B              OK        14 degrees C / 57 degrees F
    FPC 0 LU 0 TSen              OK        28 degrees C / 82 degrees F
    FPC 0 LU 0 Chip              OK        25 degrees C / 77 degrees F
    FPC 0 LU 1 TSen              OK        28 degrees C / 82 degrees F
    FPC 0 LU 1 Chip              OK        27 degrees C / 80 degrees F
    FPC 0 LU 2 TSen              OK        28 degrees C / 82 degrees F
    FPC 0 LU 2 Chip              OK        19 degrees C / 66 degrees F
    FPC 0 LU 3 TSen              OK        28 degrees C / 82 degrees F
    FPC 0 LU 3 Chip              OK        23 degrees C / 73 degrees F
    FPC 0 XM 0 TSen              OK        28 degrees C / 82 degrees F
    FPC 0 XM 0 Chip              OK        33 degrees C / 91 degrees F
    FPC 0 XM 1 TSen              OK        28 degrees C / 82 degrees F
    FPC 0 XM 1 Chip              OK        26 degrees C / 78 degrees F
    FPC 0 PLX Switch TSen        OK        28 degrees C / 82 degrees F
    FPC 0 PLX Switch Chip        OK        26 degrees C / 78 degrees F
    FPC 1 Intake                 OK        10 degrees C / 50 degrees F
    FPC 1 Exhaust A              OK        24 degrees C / 75 degrees F
    FPC 1 Exhaust B              OK        28 degrees C / 82 degrees F
    FPC 1 LU 0 TSen              OK        22 degrees C / 71 degrees F
    FPC 1 LU 0 Chip              OK        31 degrees C / 87 degrees F
    FPC 1 LU 1 TSen              OK        22 degrees C / 71 degrees F
    FPC 1 LU 1 Chip              OK        21 degrees C / 69 degrees F
    FPC 1 LU 2 TSen              OK        22 degrees C / 71 degrees F
```

```
      FPC 1 LU 2 Chip                OK         25 degrees C / 77 degrees F
      FPC 1 LU 3 TSen                OK         22 degrees C / 71 degrees F
      FPC 1 LU 3 Chip                OK         33 degrees C / 91 degrees F
      FPC 1 XM 0 TSen                OK         22 degrees C / 71 degrees F
      FPC 1 XM 0 Chip                OK         30 degrees C / 86 degrees F
      FPC 1 XF 0 TSen                OK         22 degrees C / 71 degrees F
      FPC 1 XF 0 Chip                OK         37 degrees C / 98 degrees F
      FPC 1 PLX Switch TSen          OK         22 degrees C / 71 degrees F
FPC 1 PLX Switch Chip          OK         22 degrees C / 71 degrees F
      FPC 2 Intake                   OK         9 degrees C / 48 degrees F
      FPC 2 Exhaust A                OK         10 degrees C / 50 degrees F
      FPC 2 Exhaust B                OK         10 degrees C / 50 degrees F
      FPC 2 LU 0 TSen                OK         26 degrees C / 78 degrees F
      FPC 2 LU 0 Chip                OK         25 degrees C / 77 degrees F
      FPC 2 LU 1 TSen                OK         26 degrees C / 78 degrees F
      FPC 2 LU 1 Chip                OK         26 degrees C / 78 degrees F
      FPC 2 LU 2 TSen                OK         26 degrees C / 78 degrees F
      FPC 2 LU 2 Chip                OK         17 degrees C / 62 degrees F
      FPC 2 LU 3 TSen                OK         26 degrees C / 78 degrees F
      FPC 2 LU 3 Chip                OK         22 degrees C / 71 degrees F
      FPC 2 XM 0 TSen                OK         26 degrees C / 78 degrees F
      FPC 2 XM 0 Chip                OK         34 degrees C / 93 degrees F
      FPC 2 XM 1 TSen                OK         26 degrees C / 78 degrees F
      FPC 2 XM 1 Chip                OK         26 degrees C / 78 degrees F
      FPC 2 PLX Switch TSen          OK         26 degrees C / 78 degrees F
      FPC 2 PLX Switch Chip          OK         20 degrees C / 68 degrees F
      FPC 3 Intake                   OK         12 degrees C / 53 degrees F
      FPC 3 Exhaust A                OK         16 degrees C / 60 degrees F
      FPC 3 Exhaust B                OK         26 degrees C / 78 degrees F
      FPC 3 LU 0 TSen                OK         23 degrees C / 73 degrees F
      FPC 3 LU 0 Chip                OK         26 degrees C / 78 degrees F
      FPC 3 LU 1 TSen                OK         23 degrees C / 73 degrees F
      FPC 3 LU 1 Chip                OK         27 degrees C / 80 degrees F
      FPC 3 LU 2 TSen                OK         23 degrees C / 73 degrees F
      FPC 3 LU 2 Chip                OK         22 degrees C / 71 degrees F
      FPC 3 LU 3 TSen                OK         23 degrees C / 73 degrees F
      FPC 3 LU 3 Chip                OK         21 degrees C / 69 degrees F
      FPC 3 MQ 0 TSen                OK         15 degrees C / 59 degrees F
      FPC 3 MQ 0 Chip                OK         18 degrees C / 64 degrees F
      FPC 3 MQ 1 TSen                OK         15 degrees C / 59 degrees F
      FPC 3 MQ 1 Chip                OK         20 degrees C / 68 degrees F
      FPC 3 MQ 2 TSen                OK         15 degrees C / 59 degrees F
      FPC 3 MQ 2 Chip                OK         17 degrees C / 62 degrees F
```

```
     FPC 3 MQ 3 TSen                OK         15 degrees C / 59 degrees F
     FPC 3 MQ 3 Chip                OK         16 degrees C / 60 degrees F
     FPC 4 Intake                   OK         11 degrees C / 51 degrees F
     FPC 4 Exhaust A                OK         22 degrees C / 71 degrees F
     FPC 4 Exhaust B                OK         28 degrees C / 82 degrees F
     FPC 4 LU 0 TSen                OK         22 degrees C / 71 degrees F
     FPC 4 LU 0 Chip                OK         33 degrees C / 91 degrees F
     FPC 4 LU 1 TSen                OK         22 degrees C / 71 degrees F
     FPC 4 LU 1 Chip                OK         21 degrees C / 69 degrees F
     FPC 4 LU 2 TSen                OK         22 degrees C / 71 degrees F
     FPC 4 LU 2 Chip                OK         26 degrees C / 78 degrees F
     FPC 4 LU 3 TSen                OK         22 degrees C / 71 degrees F
     FPC 4 LU 3 Chip                OK         33 degrees C / 91 degrees F
     FPC 4 XM 0 TSen                OK         22 degrees C / 71 degrees F
     FPC 4 XM 0 Chip                OK         30 degrees C / 86 degrees F
     FPC 4 XF 0 TSen                OK         22 degrees C / 71 degrees F
     FPC 4 XF 0 Chip                OK         37 degrees C / 98 degrees F
     FPC 4 PLX Switch TSen          OK         22 degrees C / 71 degrees F
     FPC 4 PLX Switch Chip          OK         23 degrees C / 73 degrees F
     FPC 5 Intake                   OK         12 degrees C / 53 degrees F
     FPC 5 Exhaust A                OK         12 degrees C / 53 degrees F
FPC 5 Exhaust B                OK         12 degrees C / 53 degrees F
     FPC 5 LU 0 TSen                OK         27 degrees C / 80 degrees F
     FPC 5 LU 0 Chip                OK         28 degrees C / 82 degrees F
     FPC 5 LU 1 TSen                OK         27 degrees C / 80 degrees F
     FPC 5 LU 1 Chip                OK         27 degrees C / 80 degrees F
     FPC 5 LU 2 TSen                OK         27 degrees C / 80 degrees F
     FPC 5 LU 2 Chip                OK         19 degrees C / 66 degrees F
     FPC 5 LU 3 TSen                OK         27 degrees C / 80 degrees F
     FPC 5 LU 3 Chip                OK         22 degrees C / 71 degrees F
     FPC 5 XM 0 TSen                OK         27 degrees C / 80 degrees F
     FPC 5 XM 0 Chip                OK         36 degrees C / 96 degrees F
     FPC 5 XM 1 TSen                OK         27 degrees C / 80 degrees F
     FPC 5 XM 1 Chip                OK         26 degrees C / 78 degrees F
     FPC 5 PLX Switch TSen          OK         27 degrees C / 80 degrees F
     FPC 5 PLX Switch Chip          OK         24 degrees C / 75 degrees F
     FPC 6 Intake                   OK         12 degrees C / 53 degrees F
     FPC 6 Exhaust A                OK         17 degrees C / 62 degrees F
     FPC 6 Exhaust B                OK         28 degrees C / 82 degrees F
     FPC 6 LU 0 TSen                OK         24 degrees C / 75 degrees F
     FPC 6 LU 0 Chip                OK         29 degrees C / 84 degrees F
     FPC 6 LU 1 TSen                OK         24 degrees C / 75 degrees F
     FPC 6 LU 1 Chip                OK         30 degrees C / 86 degrees F
```

```
    FPC 6 LU 2 TSen                OK         24 degrees C / 75 degrees F
    FPC 6 LU 2 Chip                OK         24 degrees C / 75 degrees F
    FPC 6 LU 3 TSen                OK         24 degrees C / 75 degrees F
    FPC 6 LU 3 Chip                OK         22 degrees C / 71 degrees F
    FPC 6 MQ 0 TSen                OK         16 degrees C / 60 degrees F
    FPC 6 MQ 0 Chip                OK         19 degrees C / 66 degrees F
    FPC 6 MQ 1 TSen                OK         16 degrees C / 60 degrees F
    FPC 6 MQ 1 Chip                OK         20 degrees C / 68 degrees F
    FPC 6 MQ 2 TSen                OK         16 degrees C / 60 degrees F
    FPC 6 MQ 2 Chip                OK         17 degrees C / 62 degrees F
    FPC 6 MQ 3 TSen                OK         16 degrees C / 60 degrees F
    FPC 6 MQ 3 Chip                OK         16 degrees C / 60 degrees F
    FPC 7 Intake                   OK         10 degrees C / 50 degrees F
    FPC 7 Exhaust A                OK         10 degrees C / 50 degrees F
    FPC 7 Exhaust B                OK         11 degrees C / 51 degrees F
    FPC 7 LU 0 TSen                OK         26 degrees C / 78 degrees F
    FPC 7 LU 0 Chip                OK         26 degrees C / 78 degrees F
    FPC 7 LU 1 TSen                OK         26 degrees C / 78 degrees F
    FPC 7 LU 1 Chip                OK         29 degrees C / 84 degrees F
    FPC 7 LU 2 TSen                OK         26 degrees C / 78 degrees F
    FPC 7 LU 2 Chip                OK         19 degrees C / 66 degrees F
    FPC 7 LU 3 TSen                OK         26 degrees C / 78 degrees F
    FPC 7 LU 3 Chip                OK         24 degrees C / 75 degrees F
    FPC 7 XM 0 TSen                OK         26 degrees C / 78 degrees F
    FPC 7 XM 0 Chip                OK         34 degrees C / 93 degrees F
    FPC 7 XM 1 TSen                OK         26 degrees C / 78 degrees F
    FPC 7 XM 1 Chip                OK         32 degrees C / 89 degrees F
    FPC 7 PLX Switch TSen          OK         26 degrees C / 78 degrees F
    FPC 7 PLX Switch Chip          OK         22 degrees C / 71 degrees F
    FPC 8 Intake                   OK         10 degrees C / 50 degrees F
    FPC 8 Exhaust A                OK         22 degrees C / 71 degrees F
    FPC 8 Exhaust B                OK         28 degrees C / 82 degrees F
    FPC 8 LU 0 TSen                OK         20 degrees C / 68 degrees F
    FPC 8 LU 0 Chip                OK         33 degrees C / 91 degrees F
FPC 8 LU 1 TSen                OK         20 degrees C / 68 degrees F
    FPC 8 LU 1 Chip                OK         23 degrees C / 73 degrees F
    FPC 8 LU 2 TSen                OK         20 degrees C / 68 degrees F
    FPC 8 LU 2 Chip                OK         26 degrees C / 78 degrees F
    FPC 8 LU 3 TSen                OK         20 degrees C / 68 degrees F
    FPC 8 LU 3 Chip                OK         33 degrees C / 91 degrees F
    FPC 8 XM 0 TSen                OK         20 degrees C / 68 degrees F
    FPC 8 XM 0 Chip                OK         29 degrees C / 84 degrees F
    FPC 8 XF 0 TSen                OK         20 degrees C / 68 degrees F
```

```
FPC 8 XF 0 Chip              OK        38 degrees C / 100 degrees F
FPC 8 PLX Switch TSen        OK        20 degrees C / 68 degrees F
FPC 8 PLX Switch Chip        OK        24 degrees C / 75 degrees F
FPC 9 Intake                 OK        11 degrees C / 51 degrees F
FPC 9 Exhaust A              OK        11 degrees C / 51 degrees F
FPC 9 Exhaust B              OK        11 degrees C / 51 degrees F
FPC 9 LU 0 TSen              OK        25 degrees C / 77 degrees F
FPC 9 LU 0 Chip              OK        24 degrees C / 75 degrees F
FPC 9 LU 1 TSen              OK        25 degrees C / 77 degrees F
FPC 9 LU 1 Chip              OK        26 degrees C / 78 degrees F
FPC 9 LU 2 TSen              OK        25 degrees C / 77 degrees F
FPC 9 LU 2 Chip              OK        16 degrees C / 60 degrees F
FPC 9 LU 3 TSen              OK        25 degrees C / 77 degrees F
FPC 9 LU 3 Chip              OK        21 degrees C / 69 degrees F
FPC 9 XM 0 TSen              OK        25 degrees C / 77 degrees F
FPC 9 XM 0 Chip              OK        32 degrees C / 89 degrees F
FPC 9 XM 1 TSen              OK        25 degrees C / 77 degrees F
FPC 9 XM 1 Chip              OK        25 degrees C / 77 degrees F
FPC 9 PLX Switch TSen        OK        25 degrees C / 77 degrees F
FPC 9 PLX Switch Chip        OK        21 degrees C / 69 degrees F
ADC 0 Intake                 OK        12 degrees C / 53 degrees F
ADC 0 Exhaust                OK        20 degrees C / 68 degrees F
ADC 0 ADC-XF1                OK        26 degrees C / 78 degrees F
ADC 0 ADC-XF0                OK        32 degrees C / 89 degrees F
ADC 1 Intake                 OK        11 degrees C / 51 degrees F
ADC 1 Exhaust                OK        21 degrees C / 69 degrees F
ADC 1 ADC-XF1                OK        24 degrees C / 75 degrees F
ADC 1 ADC-XF0                OK        31 degrees C / 87 degrees F
ADC 2 Intake                 OK        14 degrees C / 57 degrees F
ADC 2 Exhaust                OK        21 degrees C / 69 degrees F
ADC 2 ADC-XF1                OK        28 degrees C / 82 degrees F
ADC 2 ADC-XF0                OK        34 degrees C / 93 degrees F
ADC 3 Intake                 OK        13 degrees C / 55 degrees F
ADC 3 Exhaust                OK        19 degrees C / 66 degrees F
ADC 3 ADC-XF1                OK        24 degrees C / 75 degrees F
ADC 3 ADC-XF0                OK        31 degrees C / 87 degrees F
ADC 4 Intake                 OK        9 degrees C / 48 degrees F
ADC 4 Exhaust                OK        22 degrees C / 71 degrees F
ADC 4 ADC-XF1                OK        28 degrees C / 82 degrees F
ADC 4 ADC-XF0                OK        35 degrees C / 95 degrees F
ADC 5 Intake                 OK        12 degrees C / 53 degrees F
ADC 5 Exhaust                OK        22 degrees C / 71 degrees F
ADC 5 ADC-XF1                OK        28 degrees C / 82 degrees F
```

```
      ADC 5 ADC-XF0                  OK         34 degrees C / 93 degrees F
      ADC 6 Intake                   OK         11 degrees C / 51 degrees F
      ADC 6 Exhaust                  OK         21 degrees C / 69 degrees F
      ADC 6 ADC-XF1                  OK         26 degrees C / 78 degrees F
 ADC 6 ADC-XF0                  OK         35 degrees C / 95 degrees F
      ADC 7 Intake                   OK         14 degrees C / 57 degrees F
      ADC 7 Exhaust                  OK         22 degrees C / 71 degrees F
      ADC 7 ADC-XF1                  OK         26 degrees C / 78 degrees F
      ADC 7 ADC-XF0                  OK         34 degrees C / 93 degrees F
      ADC 8 Intake                   OK         14 degrees C / 57 degrees F
      ADC 8 Exhaust                  OK         21 degrees C / 69 degrees F
      ADC 8 ADC-XF1                  OK         24 degrees C / 75 degrees F
      ADC 8 ADC-XF0                  OK         31 degrees C / 87 degrees F
      ADC 9 Intake                   OK         10 degrees C / 50 degrees F
      ADC 9 Exhaust                  OK         22 degrees C / 71 degrees F
      ADC 9 ADC-XF1                  OK         28 degrees C / 82 degrees F
      ADC 9 ADC-XF0                  OK         36 degrees C / 96 degrees F
Fans  Fan Tray 0 Fan 1               OK         3480 RPM
      Fan Tray 0 Fan 2               OK         3480 RPM
      Fan Tray 0 Fan 3               OK         3480 RPM
      Fan Tray 0 Fan 4               OK         3360 RPM
      Fan Tray 0 Fan 5               OK         3360 RPM
      Fan Tray 0 Fan 6               OK         3480 RPM
      Fan Tray 1 Fan 1               OK         3360 RPM
      Fan Tray 1 Fan 2               OK         3360 RPM
      Fan Tray 1 Fan 3               OK         3360 RPM
      Fan Tray 1 Fan 4               OK         3480 RPM
      Fan Tray 1 Fan 5               OK         3480 RPM
      Fan Tray 1 Fan 6               OK         3480 RPM
      Fan Tray 2 Fan 1               OK         3360 RPM
      Fan Tray 2 Fan 2               OK         3360 RPM
      Fan Tray 2 Fan 3               OK         3480 RPM
      Fan Tray 2 Fan 4               OK         3480 RPM
      Fan Tray 2 Fan 5               OK         3360 RPM
      Fan Tray 2 Fan 6               OK         3480 RPM
      Fan Tray 3 Fan 1               OK         3360 RPM
      Fan Tray 3 Fan 2               OK         3360 RPM
      Fan Tray 3 Fan 3               OK         3480 RPM
      Fan Tray 3 Fan 4               OK         3480 RPM
      Fan Tray 3 Fan 5               OK         3480 RPM
      Fan Tray 3 Fan 6               OK         3360 RPM
```

show chassis environment (MX2008 Router)

```
user@host>show chassis environment
Class Item                          Status     Measurement
Temp  PSM 0                         Absent
      PSM 1                         OK         29 degrees C / 84 degrees F
      PSM 2                         OK         30 degrees C / 86 degrees F
      PSM 3                         OK         29 degrees C / 84 degrees F
      PSM 4                         OK         29 degrees C / 84 degrees F
      PSM 5                         OK         30 degrees C / 86 degrees F
      PSM 6                         OK         29 degrees C / 84 degrees F
      PSM 7                         OK         31 degrees C / 87 degrees F
      PSM 8                         Absent
      PDM 0                         OK
      PDM 1                         OK
      CB 0 Inlet1                   OK         37 degrees C / 98 degrees F
      CB 0 Inlet2                   OK         45 degrees C / 113 degrees F
      CB 0 Inlet3                   OK         44 degrees C / 111 degrees F
      CB 0 Inlet4                   OK         41 degrees C / 105 degrees F
      CB 0 Exhaust1                 OK         30 degrees C / 86 degrees F
      CB 0 Exhaust2                 OK         40 degrees C / 104 degrees F
      CB 0 Exhaust3                 OK         48 degrees C / 118 degrees F
      CB 0 Exhaust4                 OK         46 degrees C / 114 degrees F
      CB 1 Inlet1                   OK         30 degrees C / 86 degrees F
      CB 1 Inlet2                   OK         31 degrees C / 87 degrees F
      CB 1 Inlet3                   OK         29 degrees C / 84 degrees F
      CB 1 Inlet4                   OK         32 degrees C / 89 degrees F
      CB 1 Exhaust1                 OK         30 degrees C / 86 degrees F
      CB 1 Exhaust2                 OK         33 degrees C / 91 degrees F
      CB 1 Exhaust3                 OK         34 degrees C / 93 degrees F
      CB 1 Exhaust4                 OK         34 degrees C / 93 degrees F
      Routing Engine 0              OK
      Routing Engine 0 CPU          OK         75 degrees C / 167 degrees F
      Routing Engine 1              OK
      Routing Engine 1 CPU          OK         46 degrees C / 114 degrees F
      SFB 0 Inlet2                  OK         44 degrees C / 111 degrees F
      SFB 0 Exhaust1                OK         39 degrees C / 102 degrees F
      SFB 0 Inlet1                  OK         41 degrees C / 105 degrees F
      SFB 0 Exhaust2                OK         45 degrees C / 113 degrees F
      SFB 0 SFB2-PF-local           OK         45 degrees C / 113 degrees F
      SFB 0 SFB2-PF-die             OK         51 degrees C / 123 degrees F
      SFB 1 Inlet2                  OK         30 degrees C / 86 degrees F
```

```
    SFB 1 Exhaust1                OK          27 degrees C / 80 degrees F
    SFB 1 Inlet1                  OK          28 degrees C / 82 degrees F
    SFB 1 Exhaust2                OK          31 degrees C / 87 degrees F
    SFB 1 SFB2-PF-local           OK          30 degrees C / 86 degrees F
    SFB 1 SFB2-PF-die             OK          37 degrees C / 98 degrees F
    SFB 2 Inlet2                  OK          28 degrees C / 82 degrees F
    SFB 2 Exhaust1                OK          26 degrees C / 78 degrees F
    SFB 2 Inlet1                  OK          27 degrees C / 80 degrees F
    SFB 2 Exhaust2                OK          28 degrees C / 82 degrees F
    SFB 2 SFB2-PF-local           OK          27 degrees C / 80 degrees F
    SFB 2 SFB2-PF-die             OK          33 degrees C / 91 degrees F
    SFB 3 Inlet2                  OK          28 degrees C / 82 degrees F
    SFB 3 Exhaust1                OK          26 degrees C / 78 degrees F
    SFB 3 Inlet1                  OK          26 degrees C / 78 degrees F
    SFB 3 Exhaust2                OK          28 degrees C / 82 degrees F
    SFB 3 SFB2-PF-local           OK          27 degrees C / 80 degrees F
    SFB 3 SFB2-PF-die             OK          33 degrees C / 91 degrees F
    SFB 4 Inlet2                  OK          28 degrees C / 82 degrees F
    SFB 4 Exhaust1                OK          26 degrees C / 78 degrees F
    SFB 4 Inlet1                  OK          26 degrees C / 78 degrees F
    SFB 4 Exhaust2                OK          28 degrees C / 82 degrees F
    SFB 4 SFB2-PF-local           OK          27 degrees C / 80 degrees F
    SFB 4 SFB2-PF-die             OK          32 degrees C / 89 degrees F
    SFB 5 Inlet2                  OK          29 degrees C / 84 degrees F
    SFB 5 Exhaust1                OK          27 degrees C / 80 degrees F
    SFB 5 Inlet1                  OK          28 degrees C / 82 degrees F
    SFB 5 Exhaust2                OK          29 degrees C / 84 degrees F
    SFB 5 SFB2-PF-local           OK          28 degrees C / 82 degrees F
    SFB 5 SFB2-PF-die             OK          34 degrees C / 93 degrees F
    SFB 6 Inlet2                  OK          33 degrees C / 91 degrees F
    SFB 6 Exhaust1                OK          32 degrees C / 89 degrees F
    SFB 6 Inlet1                  OK          32 degrees C / 89 degrees F
    SFB 6 Exhaust2                OK          34 degrees C / 93 degrees F
    SFB 6 SFB2-PF-local           OK          33 degrees C / 91 degrees F
    SFB 6 SFB2-PF-die             OK          40 degrees C / 104 degrees F
    SFB 7 Inlet2                  OK          29 degrees C / 84 degrees F
    SFB 7 Exhaust1                OK          28 degrees C / 82 degrees F
    SFB 7 Inlet1                  OK          29 degrees C / 84 degrees F
    SFB 7 Exhaust2                OK          29 degrees C / 84 degrees F
    SFB 7 SFB2-PF-local           OK          28 degrees C / 82 degrees F
    SFB 7 SFB2-PF-die             OK          33 degrees C / 91 degrees F
    FPC 0 Intake                  OK          29 degrees C / 84 degrees F
    FPC 0 Exhaust A               OK          42 degrees C / 107 degrees F
```

```
    FPC 0 Exhaust B                OK          42 degrees C / 107 degrees F
    FPC 0 XL 0 TSen                OK          38 degrees C / 100 degrees F
    FPC 0 XL 0 Chip                OK          53 degrees C / 127 degrees F
    FPC 0 XL 0 XR2 0 TSen          OK          38 degrees C / 100 degrees F
    FPC 0 XL 0 XR2 0 Chip          OK          59 degrees C / 138 degrees F
    FPC 0 XL 0 XR2 1 TSen          OK          38 degrees C / 100 degrees F
    FPC 0 XL 0 XR2 1 Chip          OK          59 degrees C / 138 degrees F
    FPC 0 XL 1 TSen                OK          30 degrees C / 86 degrees F
    FPC 0 XL 1 Chip                OK          42 degrees C / 107 degrees F
    FPC 0 XL 1 XR2 0 TSen          OK          30 degrees C / 86 degrees F
    FPC 0 XL 1 XR2 0 Chip          OK          49 degrees C / 120 degrees F
    FPC 0 XL 1 XR2 1 TSen          OK          30 degrees C / 86 degrees F
    FPC 0 XL 1 XR2 1 Chip          OK          50 degrees C / 122 degrees F
    FPC 0 XM 0 TSen                OK          42 degrees C / 107 degrees F
    FPC 0 XM 0 Chip                OK          49 degrees C / 120 degrees F
    FPC 0 XM 1 TSen                OK          42 degrees C / 107 degrees F
    FPC 0 XM 1 Chip                OK          42 degrees C / 107 degrees F
    FPC 0 XM 2 TSen                OK          42 degrees C / 107 degrees F
    FPC 0 XM 2 Chip                OK          42 degrees C / 107 degrees F
    FPC 0 XM 3 TSen                OK          42 degrees C / 107 degrees F
    FPC 0 XM 3 Chip                OK          40 degrees C / 104 degrees F
    FPC 0 PCIe Switch TSen         OK          42 degrees C / 107 degrees F
    FPC 0 PCIe Switch Chip         OK          22 degrees C / 71 degrees F
    FPC 1 Intake                   OK          29 degrees C / 84 degrees F
    FPC 1 Exhaust A                OK          52 degrees C / 125 degrees F
    FPC 1 Exhaust B                OK          44 degrees C / 111 degrees F
    FPC 1 EA0 TSen                 OK          54 degrees C / 129 degrees F
    FPC 1 EA0 Chip                 OK          47 degrees C / 116 degrees F
    FPC 1 EA0_XR0 TSen             OK          54 degrees C / 129 degrees F
    FPC 1 EA0_XR0 Chip             OK          56 degrees C / 132 degrees F
    FPC 1 EA0_XR1 TSen             OK          54 degrees C / 129 degrees F
    FPC 1 EA0_XR1 Chip             OK          53 degrees C / 127 degrees F
    FPC 1 EA1 TSen                 OK          54 degrees C / 129 degrees F
    FPC 1 EA1 Chip                 OK          49 degrees C / 120 degrees F
    FPC 1 EA1_XR0 TSen             OK          54 degrees C / 129 degrees F
    FPC 1 EA1_XR0 Chip             OK          57 degrees C / 134 degrees F
    FPC 1 EA1_XR1 TSen             OK          54 degrees C / 129 degrees F
    FPC 1 EA1_XR1 Chip             OK          58 degrees C / 136 degrees F
    FPC 1 PEX TSen                 OK          54 degrees C / 129 degrees F
    FPC 1 PEX Chip                 OK          39 degrees C / 102 degrees F
    FPC 1 EA2 TSen                 OK          43 degrees C / 109 degrees F
    FPC 1 EA2 Chip                 OK          39 degrees C / 102 degrees F
    FPC 1 EA2_XR0 TSen             OK          43 degrees C / 109 degrees F
```

```
    FPC 1 EA2_XR0 Chip            OK        45 degrees C / 113 degrees F
    FPC 1 EA2_XR1 TSen            OK        43 degrees C / 109 degrees F
    FPC 1 EA2_XR1 Chip            OK        42 degrees C / 107 degrees F
    FPC 1 EA3 TSen                OK        43 degrees C / 109 degrees F
    FPC 1 EA3 Chip                OK        40 degrees C / 104 degrees F
    FPC 1 EA3_XR0 TSen            OK        43 degrees C / 109 degrees F
    FPC 1 EA3_XR0 Chip            OK        50 degrees C / 122 degrees F
    FPC 1 EA3_XR1 TSen            OK        43 degrees C / 109 degrees F
    FPC 1 EA3_XR1 Chip            OK        46 degrees C / 114 degrees F
    FPC 1 EA0_HMC0 Logic die      OK        60 degrees C / 140 degrees F
    FPC 1 EA0_HMC0 DRAM botm      OK        57 degrees C / 134 degrees F
    FPC 1 EA0_HMC1 Logic die      OK        61 degrees C / 141 degrees F
    FPC 1 EA0_HMC1 DRAM botm      OK        58 degrees C / 136 degrees F
    FPC 1 EA0_HMC2 Logic die      OK        57 degrees C / 134 degrees F
    FPC 1 EA0_HMC2 DRAM botm      OK        54 degrees C / 129 degrees F
    FPC 1 EA1_HMC0 Logic die      OK        65 degrees C / 149 degrees F
    FPC 1 EA1_HMC0 DRAM botm      OK        62 degrees C / 143 degrees F
    FPC 1 EA1_HMC1 Logic die      OK        64 degrees C / 147 degrees F
    FPC 1 EA1_HMC1 DRAM botm      OK        61 degrees C / 141 degrees F
    FPC 1 EA1_HMC2 Logic die      OK        61 degrees C / 141 degrees F
    FPC 1 EA1_HMC2 DRAM botm      OK        58 degrees C / 136 degrees F
    FPC 1 EA2_HMC0 Logic die      OK        50 degrees C / 122 degrees F
    FPC 1 EA2_HMC0 DRAM botm      OK        47 degrees C / 116 degrees F
    FPC 1 EA2_HMC1 Logic die      OK        54 degrees C / 129 degrees F
    FPC 1 EA2_HMC1 DRAM botm      OK        51 degrees C / 123 degrees F
    FPC 1 EA2_HMC2 Logic die      OK        51 degrees C / 123 degrees F
    FPC 1 EA2_HMC2 DRAM botm      OK        48 degrees C / 118 degrees F
    FPC 1 EA3_HMC0 Logic die      OK        51 degrees C / 123 degrees F
    FPC 1 EA3_HMC0 DRAM botm      OK        48 degrees C / 118 degrees F
    FPC 1 EA3_HMC1 Logic die      OK        51 degrees C / 123 degrees F
    FPC 1 EA3_HMC1 DRAM botm      OK        48 degrees C / 118 degrees F
    FPC 1 EA3_HMC2 Logic die      OK        51 degrees C / 123 degrees F
    FPC 1 EA3_HMC2 DRAM botm      OK        48 degrees C / 118 degrees F
    FPC 7 Intake                  OK        30 degrees C / 86 degrees F
    FPC 7 Exhaust A               OK        45 degrees C / 113 degrees F
    FPC 7 Exhaust B               OK        38 degrees C / 100 degrees F
    FPC 7 QX 0 TSen               OK        48 degrees C / 118 degrees F
    FPC 7 QX 0 Chip               OK        51 degrees C / 123 degrees F
    FPC 7 LU 0 TCAM TSen          OK        48 degrees C / 118 degrees F
    FPC 7 LU 0 TCAM Chip          OK        51 degrees C / 123 degrees F
    FPC 7 LU 0 TSen               OK        48 degrees C / 118 degrees F
    FPC 7 LU 0 Chip               OK        50 degrees C / 122 degrees F
    FPC 7 MQ 0 TSen               OK        48 degrees C / 118 degrees F
```

```
      FPC 7 MQ 0 Chip                OK         54 degrees C / 129 degrees F
      FPC 7 QX 1 TSen                OK         41 degrees C / 105 degrees F
      FPC 7 QX 1 Chip                OK         42 degrees C / 107 degrees F
      FPC 7 LU 1 TCAM TSen           OK         41 degrees C / 105 degrees F
      FPC 7 LU 1 TCAM Chip           OK         43 degrees C / 109 degrees F
      FPC 7 LU 1 TSen                OK         41 degrees C / 105 degrees F
      FPC 7 LU 1 Chip                OK         46 degrees C / 114 degrees F
      FPC 7 MQ 1 TSen                OK         41 degrees C / 105 degrees F
      FPC 7 MQ 1 Chip                OK         47 degrees C / 116 degrees F
      ADC 7 Intake                   OK         32 degrees C / 89 degrees F
      ADC 7 Exhaust                  OK         39 degrees C / 102 degrees F
      ADC 7 ADC-XF1                  OK         46 degrees C / 114 degrees F
      ADC 7 ADC-XF0                  OK         54 degrees C / 129 degrees F
Fans  Fan Tray 0 Fan 1               OK         6240 RPM
      Fan Tray 0 Fan 2               OK         6120 RPM
      Fan Tray 0 Fan 3               OK         6120 RPM
      Fan Tray 0 Fan 4               OK         5760 RPM
      Fan Tray 0 Fan 5               OK         5880 RPM
      Fan Tray 0 Fan 6               OK         6000 RPM
      Fan Tray 1 Fan 1               OK         5880 RPM
      Fan Tray 1 Fan 2               OK         5880 RPM
      Fan Tray 1 Fan 3               OK         6000 RPM
      Fan Tray 1 Fan 4               OK         6000 RPM
      Fan Tray 1 Fan 5               OK         6000 RPM
      Fan Tray 1 Fan 6               OK         6000 RPM
```

show chassis environment (T320 Router)

```
user@host> show chassis environment
Class Item                   Status     Measurement
Power PEM 0                  OK
      PEM 1                  Absent
Temp  SCG 0                  OK         28 degrees C / 82 degrees F
      SCG 1                  OK         28 degrees C / 82 degrees F
      Routing Engine 0       OK         31 degrees C / 87 degrees F
      Routing Engine 1       OK         30 degrees C / 86 degrees F
      CB 0                   OK         32 degrees C / 89 degrees F
      CB 1                   OK         32 degrees C / 89 degrees F
      SIB 0                  OK         33 degrees C / 91 degrees F
      SIB 1                  OK         33 degrees C / 91 degrees F
      SIB 2                  OK         34 degrees C / 93 degrees F
```

```
      FPC 0 Top              OK         38 degrees C / 100 degrees F
      FPC 0 Bottom           OK         32 degrees C / 89 degrees F
      FPC 1 Top              OK         38 degrees C / 100 degrees F
      FPC 1 Bottom           OK         33 degrees C / 91 degrees F
      FPC 2 Top              OK         36 degrees C / 96 degrees F
      FPC 2 Bottom           OK         31 degrees C / 87 degrees F
      FPM GBUS               OK         26 degrees C / 78 degrees F
      FPM Display            OK         29 degrees C / 84 degrees F
Fans  Top Left Front fan     OK         Spinning at normal speed
      Top Left Middle fan    OK         Spinning at normal speed
      Top Left Rear fan      OK         Spinning at normal speed
      Top Right Front fan    OK         Spinning at normal speed
      Top Right Middle fan   OK         Spinning at normal speed
      Top Right Rear fan     OK         Spinning at normal speed
      Bottom Left Front fan  OK         Spinning at normal speed
      Bottom Left Middle fan OK         Spinning at normal speed
      Bottom Left Rear fan   OK         Spinning at normal speed
      Bottom Right Front fan OK         Spinning at normal speed
      Bottom Right Middle fan OK        Spinning at normal speed
      Bottom Right Rear fan  OK         Spinning at normal speed
      Rear Tray Top fan      OK         Spinning at normal speed
      Rear Tray Second fan   OK         Spinning at normal speed
      Rear Tray Middle fan   OK         Spinning at normal speed
      Rear Tray Fourth fan   OK         Spinning at normal speed
      Rear Tray Bottom fan   OK         Spinning at normal speed
Misc  CIP                    OK
      SPMB 0                 OK
      SPMB 1                 OK
```

show chassis environment (MX10003 Router)

```
user@host> show chassis environment

Class Item                           Status     Measurement
Temp  CB 0 Exhaust Temp Sensor 0x49  OK         36 degrees C / 96 degrees F
      CB 0 Inlet Temp Sensor 0x49    OK         29 degrees C / 84 degrees F
      CB 1 Exhaust Temp Sensor 0x49  OK         36 degrees C / 96 degrees F
      CB 1 Inlet Temp Sensor 0x49    OK         31 degrees C / 87 degrees F
      FPC 0 Intake Temp Sensor       OK         29 degrees C / 84 degrees F
      FPC 0 Exhaust-A Temp Sensor    OK         55 degrees C / 131 degrees F
```

```
FPC 0 Exhaust-B Temp Sensor    OK         44 degrees C / 111 degrees F
FPC 0 EA0 Chip                 OK         58 degrees C / 136 degrees F
FPC 0 EA0-XR0 Chip             OK         61 degrees C / 141 degrees F
FPC 0 EA0-XR1 Chip             OK         62 degrees C / 143 degrees F
FPC 0 EA1 Chip                 OK         67 degrees C / 152 degrees F
FPC 0 EA1-XR0 Chip             OK         71 degrees C / 159 degrees F
FPC 0 EA1-XR1 Chip             OK         72 degrees C / 161 degrees F
FPC 0 PEX Chip                 OK         75 degrees C / 167 degrees F
FPC 0 EA2 Chip                 OK         49 degrees C / 120 degrees F
FPC 0 EA2-XR0 Chip             OK         55 degrees C / 131 degrees F
FPC 0 EA2-XR1 Chip             OK         56 degrees C / 132 degrees F
FPC 0 PF Chip                  OK         68 degrees C / 154 degrees F
FPC 0 EA0_HMC0 Logic die       OK         72 degrees C / 161 degrees F
FPC 0 EA0_HMC0 DRAM botm       OK         69 degrees C / 156 degrees F
FPC 0 EA0_HMC1 Logic die       OK         72 degrees C / 161 degrees F
FPC 0 EA0_HMC1 DRAM botm       OK         69 degrees C / 156 degrees F
FPC 0 EA0_HMC2 Logic die       OK         75 degrees C / 167 degrees F
FPC 0 EA0_HMC2 DRAM botm       OK         72 degrees C / 161 degrees F
FPC 0 EA1_HMC0 Logic die       OK         81 degrees C / 177 degrees F
FPC 0 EA1_HMC0 DRAM botm       OK         78 degrees C / 172 degrees F
FPC 0 EA1_HMC1 Logic die       OK         79 degrees C / 174 degrees F
FPC 0 EA1_HMC1 DRAM botm       OK         76 degrees C / 168 degrees F
FPC 0 EA1_HMC2 Logic die       OK         82 degrees C / 179 degrees F
FPC 0 EA1_HMC2 DRAM botm       OK         79 degrees C / 174 degrees F
FPC 0 EA2_HMC0 Logic die       OK         61 degrees C / 141 degrees F
FPC 0 EA2_HMC0 DRAM botm       OK         58 degrees C / 136 degrees F
FPC 0 EA2_HMC1 Logic die       OK         62 degrees C / 143 degrees F
FPC 0 EA2_HMC1 DRAM botm       OK         59 degrees C / 138 degrees F
FPC 0 EA2_HMC2 Logic die       OK         64 degrees C / 147 degrees F
FPC 0 EA2_HMC2 DRAM botm       OK         61 degrees C / 141 degrees F
FPC 1 Intake Temp Sensor       OK         28 degrees C / 82 degrees F
FPC 1 Exhaust-A Temp Sensor    OK         58 degrees C / 136 degrees F
FPC 1 Exhaust-B Temp Sensor    OK         46 degrees C / 114 degrees F
FPC 1 EA0 Chip                 OK         64 degrees C / 147 degrees F
FPC 1 EA0-XR0 Chip             OK         67 degrees C / 152 degrees F
FPC 1 EA0-XR1 Chip             OK         68 degrees C / 154 degrees F
FPC 1 EA1 Chip                 OK         70 degrees C / 158 degrees F
FPC 1 EA1-XR0 Chip             OK         74 degrees C / 165 degrees F
FPC 1 EA1-XR1 Chip             OK         74 degrees C / 165 degrees F
FPC 1 PEX Chip                 OK         88 degrees C / 190 degrees F
FPC 1 EA2 Chip                 OK         50 degrees C / 122 degrees F
FPC 1 EA2-XR0 Chip             OK         54 degrees C / 129 degrees F
FPC 1 EA2-XR1 Chip             OK         56 degrees C / 132 degrees F
```

```
      FPC 1 PF Chip                 OK         71 degrees C / 159 degrees F
      FPC 1 EA0_HMC0 Logic die      OK         74 degrees C / 165 degrees F
      FPC 1 EA0_HMC0 DRAM botm      OK         71 degrees C / 159 degrees F
      FPC 1 EA0_HMC1 Logic die      OK         78 degrees C / 172 degrees F
      FPC 1 EA0_HMC1 DRAM botm      OK         75 degrees C / 167 degrees F
      FPC 1 EA0_HMC2 Logic die      OK         78 degrees C / 172 degrees F
      FPC 1 EA0_HMC2 DRAM botm      OK         75 degrees C / 167 degrees F
      FPC 1 EA1_HMC0 Logic die      OK         84 degrees C / 183 degrees F
      FPC 1 EA1_HMC0 DRAM botm      OK         81 degrees C / 177 degrees F
      FPC 1 EA1_HMC1 Logic die      OK         81 degrees C / 177 degrees F
      FPC 1 EA1_HMC1 DRAM botm      OK         78 degrees C / 172 degrees F
      FPC 1 EA1_HMC2 Logic die      OK         85 degrees C / 185 degrees F
      FPC 1 EA1_HMC2 DRAM botm      OK         82 degrees C / 179 degrees F
      FPC 1 EA2_HMC0 Logic die      OK         63 degrees C / 145 degrees F
      FPC 1 EA2_HMC0 DRAM botm      OK         60 degrees C / 140 degrees F
      FPC 1 EA2_HMC1 Logic die      OK         60 degrees C / 140 degrees F
      FPC 1 EA2_HMC1 DRAM botm      OK         57 degrees C / 134 degrees F
      FPC 1 EA2_HMC2 Logic die      OK         66 degrees C / 150 degrees F
      FPC 1 EA2_HMC2 DRAM botm      OK         63 degrees C / 145 degrees F
Power PEM 0                         OK
      PEM 1                         OK
      PEM 2                         OK
      PEM 3                         OK
      PEM 4                         Absent
      PEM 5                         Absent
Fans  Fan Tray 0 Fan 0              OK         Spinning at normal speed
      Fan Tray 0 Fan 1              OK         Spinning at normal speed
      Fan Tray 0 Fan 2              OK         Spinning at normal speed
      Fan Tray 0 Fan 3              OK         Spinning at normal speed
      Fan Tray 1 Fan 0              OK         Spinning at normal speed
      Fan Tray 1 Fan 1              OK         Spinning at normal speed
      Fan Tray 1 Fan 2              OK         Spinning at normal speed
      Fan Tray 1 Fan 3              OK         Spinning at normal speed
      Fan Tray 2 Fan 0              OK         Spinning at normal speed
      Fan Tray 2 Fan 1              OK         Spinning at normal speed
      Fan Tray 2 Fan 2              OK         Spinning at normal speed
      Fan Tray 2 Fan 3              OK         Spinning at normal speed
      Fan Tray 3 Fan 0              OK         Spinning at normal speed
      Fan Tray 3 Fan 1              OK         Spinning at normal speed
      Fan Tray 3 Fan 2              OK         Spinning at normal speed
      Fan Tray 3 Fan 3              OK         Spinning at normal speed
```

show chassis environment (MX10008 Router)

```
user@host> show chassis environment
Class Item                          Status     Measurement
      Routing Engine 0 CPU          OK         41 degrees C / 105 degrees F
      Routing Engine 1 CPU          OK         40 degrees C / 104 degrees F
Temp  CB 0 Intake A Temp Sensor     OK         24 degrees C / 75 degrees F
      CB 0 Intake B Temp Sensor     OK         24 degrees C / 75 degrees F
      CB 0 Exhaust A Temp Sensor    OK         28 degrees C / 82 degrees F
      CB 0 Exhaust B Temp Sensor    OK         30 degrees C / 86 degrees F
      CB 0 Middle Temp Sensor       OK         28 degrees C / 82 degrees F
      CB 1 Intake A Temp Sensor     OK         24 degrees C / 75 degrees F
      CB 1 Intake B Temp Sensor     OK         23 degrees C / 73 degrees F
      CB 1 Exhaust A Temp Sensor    OK         27 degrees C / 80 degrees F
      CB 1 Exhaust B Temp Sensor    OK         29 degrees C / 84 degrees F
      CB 1 Middle Temp Sensor       OK         28 degrees C / 82 degrees F
      FPC 0 Intake-A Temp Sensor    OK         32 degrees C / 89 degrees F
      FPC 0 Exhaust-A Temp Sensor   OK         44 degrees C / 111 degrees F
      FPC 0 Exhaust-B Temp Sensor   OK         49 degrees C / 120 degrees F
      FPC 0 EA0 Temp Sensor         OK         66 degrees C / 150 degrees F
      FPC 0 EA0_XR0 Temp Sensor     OK         69 degrees C / 156 degrees F
      FPC 0 EA0_XR1 Temp Sensor     OK         73 degrees C / 163 degrees F
      FPC 0 EA1 Temp Sensor         OK         60 degrees C / 140 degrees F
      FPC 0 EA1_XR0 Temp Sensor     OK         64 degrees C / 147 degrees F
      FPC 0 EA1_XR1 Temp Sensor     OK         63 degrees C / 145 degrees F
      FPC 0 EA2 Temp Sensor         OK         68 degrees C / 154 degrees F
      FPC 0 EA2_XR0 Temp Sensor     OK         73 degrees C / 163 degrees F
      FPC 0 EA2_XR1 Temp Sensor     OK         72 degrees C / 161 degrees F
      FPC 0 EA3 Temp Sensor         OK         63 degrees C / 145 degrees F
      FPC 0 EA3_XR0 Temp Sensor     OK         66 degrees C / 150 degrees F
      FPC 0 EA3_XR1 Temp Sensor     OK         65 degrees C / 149 degrees F
      FPC 0 EA4 Temp Sensor         OK         68 degrees C / 154 degrees F
      FPC 0 EA4_XR0 Temp Sensor     OK         71 degrees C / 159 degrees F
      FPC 0 EA4_XR1 Temp Sensor     OK         70 degrees C / 158 degrees F
      FPC 0 EA5 Temp Sensor         OK         56 degrees C / 132 degrees F
      FPC 0 EA5_XR0 Temp Sensor     OK         61 degrees C / 141 degrees F
      FPC 0 EA5_XR1 Temp Sensor     OK         63 degrees C / 145 degrees F
      FPC 0 EA0_HMC0 Logic die      OK         75 degrees C / 167 degrees F
      FPC 0 EA0_HMC0 DRAM botm      OK         72 degrees C / 161 degrees F
      FPC 0 EA0_HMC1 Logic die      OK         75 degrees C / 167 degrees F
      FPC 0 EA0_HMC1 DRAM botm      OK         72 degrees C / 161 degrees F
      FPC 0 EA0_HMC2 Logic die      OK         77 degrees C / 170 degrees F
```

```
FPC 0 EA0_HMC2 DRAM botm       OK         74 degrees C / 165 degrees F
FPC 0 EA1_HMC0 Logic die       OK         72 degrees C / 161 degrees F
FPC 0 EA1_HMC0 DRAM botm       OK         69 degrees C / 156 degrees F
FPC 0 EA1_HMC1 Logic die       OK         73 degrees C / 163 degrees F
FPC 0 EA1_HMC1 DRAM botm       OK         70 degrees C / 158 degrees F
FPC 0 EA1_HMC2 Logic die       OK         72 degrees C / 161 degrees F
FPC 0 EA1_HMC2 DRAM botm       OK         69 degrees C / 156 degrees F
FPC 0 EA2_HMC0 Logic die       OK         80 degrees C / 176 degrees F
FPC 0 EA2_HMC0 DRAM botm       OK         77 degrees C / 170 degrees F
FPC 0 EA2_HMC1 Logic die       OK         80 degrees C / 176 degrees F
FPC 0 EA2_HMC1 DRAM botm       OK         77 degrees C / 170 degrees F
FPC 0 EA2_HMC2 Logic die       OK         79 degrees C / 174 degrees F
FPC 0 EA2_HMC2 DRAM botm       OK         76 degrees C / 168 degrees F
FPC 0 EA3_HMC0 Logic die       OK         77 degrees C / 170 degrees F
FPC 0 EA3_HMC0 DRAM botm       OK         74 degrees C / 165 degrees F
FPC 0 EA3_HMC1 Logic die       OK         78 degrees C / 172 degrees F
FPC 0 EA3_HMC1 DRAM botm       OK         75 degrees C / 167 degrees F
FPC 0 EA3_HMC2 Logic die       OK         77 degrees C / 170 degrees F
FPC 0 EA3_HMC2 DRAM botm       OK         74 degrees C / 165 degrees F
FPC 0 EA4_HMC0 Logic die       OK         80 degrees C / 176 degrees F
FPC 0 EA4_HMC0 DRAM botm       OK         77 degrees C / 170 degrees F
FPC 0 EA4_HMC1 Logic die       OK         81 degrees C / 177 degrees F
FPC 0 EA4_HMC1 DRAM botm       OK         78 degrees C / 172 degrees F
FPC 0 EA4_HMC2 Logic die       OK         80 degrees C / 176 degrees F
FPC 0 EA4_HMC2 DRAM botm       OK         77 degrees C / 170 degrees F
FPC 0 EA5_HMC0 Logic die       OK         68 degrees C / 154 degrees F
FPC 0 EA5_HMC0 DRAM botm       OK         65 degrees C / 149 degrees F
FPC 0 EA5_HMC1 Logic die       OK         68 degrees C / 154 degrees F
FPC 0 EA5_HMC1 DRAM botm       OK         65 degrees C / 149 degrees F
FPC 0 EA5_HMC2 Logic die       OK         67 degrees C / 152 degrees F
FPC 0 EA5_HMC2 DRAM botm       OK         64 degrees C / 147 degrees F
FPC 2 Intake-A Temp Sensor     OK         32 degrees C / 89 degrees F
FPC 2 Exhaust-A Temp Sensor    OK         52 degrees C / 125 degrees F
FPC 2 Exhaust-B Temp Sensor    OK         50 degrees C / 122 degrees F
FPC 2 EA0 Temp Sensor          OK         71 degrees C / 159 degrees F
FPC 2 EA0_XR0 Temp Sensor      OK         75 degrees C / 167 degrees F
FPC 2 EA0_XR1 Temp Sensor      OK         78 degrees C / 172 degrees F
FPC 2 EA1 Temp Sensor          OK         64 degrees C / 147 degrees F
FPC 2 EA1_XR0 Temp Sensor      OK         67 degrees C / 152 degrees F
FPC 2 EA1_XR1 Temp Sensor      OK         65 degrees C / 149 degrees F
FPC 2 EA2 Temp Sensor          OK         75 degrees C / 167 degrees F
FPC 2 EA2_XR0 Temp Sensor      OK         80 degrees C / 176 degrees F
FPC 2 EA2_XR1 Temp Sensor      OK         80 degrees C / 176 degrees F
```

```
FPC 2 EA3 Temp Sensor          OK          66 degrees C / 150 degrees F
FPC 2 EA3_XR0 Temp Sensor      OK          69 degrees C / 156 degrees F
FPC 2 EA3_XR1 Temp Sensor      OK          69 degrees C / 156 degrees F
FPC 2 EA4 Temp Sensor          OK          75 degrees C / 167 degrees F
FPC 2 EA4_XR0 Temp Sensor      OK          76 degrees C / 168 degrees F
FPC 2 EA4_XR1 Temp Sensor      OK          75 degrees C / 167 degrees F
FPC 2 EA5 Temp Sensor          OK          60 degrees C / 140 degrees F
FPC 2 EA5_XR0 Temp Sensor      OK          64 degrees C / 147 degrees F
FPC 2 EA5_XR1 Temp Sensor      OK          64 degrees C / 147 degrees F
FPC 2 EA0_HMC0 Logic die       OK          84 degrees C / 183 degrees F
FPC 2 EA0_HMC0 DRAM botm       OK          81 degrees C / 177 degrees F
FPC 2 EA0_HMC1 Logic die       OK          85 degrees C / 185 degrees F
FPC 2 EA0_HMC1 DRAM botm       OK          82 degrees C / 179 degrees F
FPC 2 EA0_HMC2 Logic die       OK          83 degrees C / 181 degrees F
FPC 2 EA0_HMC2 DRAM botm       OK          80 degrees C / 176 degrees F
FPC 2 EA1_HMC0 Logic die       OK          76 degrees C / 168 degrees F
FPC 2 EA1_HMC0 DRAM botm       OK          73 degrees C / 163 degrees F
FPC 2 EA1_HMC1 Logic die       OK          76 degrees C / 168 degrees F
FPC 2 EA1_HMC1 DRAM botm       OK          73 degrees C / 163 degrees F
FPC 2 EA1_HMC2 Logic die       OK          76 degrees C / 168 degrees F
FPC 2 EA1_HMC2 DRAM botm       OK          73 degrees C / 163 degrees F
FPC 2 EA2_HMC0 Logic die       OK          86 degrees C / 186 degrees F
FPC 2 EA2_HMC0 DRAM botm       OK          83 degrees C / 181 degrees F
FPC 2 EA2_HMC1 Logic die       OK          87 degrees C / 188 degrees F
FPC 2 EA2_HMC1 DRAM botm       OK          84 degrees C / 183 degrees F
FPC 2 EA2_HMC2 Logic die       OK          87 degrees C / 188 degrees F
FPC 2 EA2_HMC2 DRAM botm       OK          84 degrees C / 183 degrees F
FPC 2 EA3_HMC0 Logic die       OK          80 degrees C / 176 degrees F
FPC 2 EA3_HMC0 DRAM botm       OK          77 degrees C / 170 degrees F
FPC 2 EA3_HMC1 Logic die       OK          80 degrees C / 176 degrees F
FPC 2 EA3_HMC1 DRAM botm       OK          77 degrees C / 170 degrees F
FPC 2 EA3_HMC2 Logic die       OK          80 degrees C / 176 degrees F
FPC 2 EA3_HMC2 DRAM botm       OK          77 degrees C / 170 degrees F
FPC 2 EA4_HMC0 Logic die       OK          88 degrees C / 190 degrees F
FPC 2 EA4_HMC0 DRAM botm       OK          85 degrees C / 185 degrees F
FPC 2 EA4_HMC1 Logic die       OK          89 degrees C / 192 degrees F
FPC 2 EA4_HMC1 DRAM botm       OK          86 degrees C / 186 degrees F
FPC 2 EA4_HMC2 Logic die       OK          80 degrees C / 176 degrees F
FPC 2 EA4_HMC2 DRAM botm       OK          77 degrees C / 170 degrees F
FPC 2 EA5_HMC0 Logic die       OK          72 degrees C / 161 degrees F
FPC 2 EA5_HMC0 DRAM botm       OK          69 degrees C / 156 degrees F
FPC 2 EA5_HMC1 Logic die       OK          69 degrees C / 156 degrees F
FPC 2 EA5_HMC1 DRAM botm       OK          66 degrees C / 150 degrees F
```

```
    FPC 2 EA5_HMC2 Logic die       OK         72 degrees C / 161 degrees F
    FPC 2 EA5_HMC2 DRAM botm       OK         69 degrees C / 156 degrees F
    FPC 3 Intake-A Temp Sensor     OK         30 degrees C / 86 degrees F
    FPC 3 Exhaust-A Temp Sensor    OK         48 degrees C / 118 degrees F
    FPC 3 Exhaust-B Temp Sensor    OK         44 degrees C / 111 degrees F
    FPC 3 EA0 Temp Sensor          OK         60 degrees C / 140 degrees F
    FPC 3 EA0_XR0 Temp Sensor      OK         65 degrees C / 149 degrees F
    FPC 3 EA0_XR1 Temp Sensor      OK         67 degrees C / 152 degrees F
    FPC 3 EA1 Temp Sensor          OK         54 degrees C / 129 degrees F
    FPC 3 EA1_XR0 Temp Sensor      OK         59 degrees C / 138 degrees F
    FPC 3 EA1_XR1 Temp Sensor      OK         58 degrees C / 136 degrees F
    FPC 3 EA2 Temp Sensor          OK         62 degrees C / 143 degrees F
    FPC 3 EA2_XR0 Temp Sensor      OK         66 degrees C / 150 degrees F
    FPC 3 EA2_XR1 Temp Sensor      OK         66 degrees C / 150 degrees F
    FPC 3 EA3 Temp Sensor          OK         54 degrees C / 129 degrees F
    FPC 3 EA3_XR0 Temp Sensor      OK         57 degrees C / 134 degrees F
    FPC 3 EA3_XR1 Temp Sensor      OK         56 degrees C / 132 degrees F
    FPC 3 EA4 Temp Sensor          OK         68 degrees C / 154 degrees F
    FPC 3 EA4_XR0 Temp Sensor      OK         71 degrees C / 159 degrees F
    FPC 3 EA4_XR1 Temp Sensor      OK         70 degrees C / 158 degrees F
    FPC 3 EA5 Temp Sensor          OK         55 degrees C / 131 degrees F
    FPC 3 EA5_XR0 Temp Sensor      OK         58 degrees C / 136 degrees F
    FPC 3 EA5_XR1 Temp Sensor      OK         58 degrees C / 136 degrees F
    FPC 3 EA0_HMC0 Logic die       OK         69 degrees C / 156 degrees F
    FPC 3 EA0_HMC0 DRAM botm       OK         66 degrees C / 150 degrees F
    FPC 3 EA0_HMC1 Logic die       OK         70 degrees C / 158 degrees F
    FPC 3 EA0_HMC1 DRAM botm       OK         67 degrees C / 152 degrees F
    FPC 3 EA0_HMC2 Logic die       OK         69 degrees C / 156 degrees F
    FPC 3 EA0_HMC2 DRAM botm       OK         66 degrees C / 150 degrees F
    FPC 3 EA1_HMC0 Logic die       OK         67 degrees C / 152 degrees F
    FPC 3 EA1_HMC0 DRAM botm       OK         64 degrees C / 147 degrees F
    FPC 3 EA1_HMC1 Logic die       OK         64 degrees C / 147 degrees F
    FPC 3 EA1_HMC1 DRAM botm       OK         61 degrees C / 141 degrees F
    FPC 3 EA1_HMC2 Logic die       OK         64 degrees C / 147 degrees F
    FPC 3 EA1_HMC2 DRAM botm       OK         61 degrees C / 141 degrees F
    FPC 3 EA2_HMC0 Logic die       OK         74 degrees C / 165 degrees F
    FPC 3 EA2_HMC0 DRAM botm       OK         71 degrees C / 159 degrees F
    FPC 3 EA2_HMC1 Logic die       OK         76 degrees C / 168 degrees F
    FPC 3 EA2_HMC1 DRAM botm       OK         73 degrees C / 163 degrees F
    FPC 3 EA2_HMC2 Logic die       OK         74 degrees C / 165 degrees F
    FPC 3 EA2_HMC2 DRAM botm       OK         71 degrees C / 159 degrees F
    FPC 3 EA3_HMC0 Logic die       OK         69 degrees C / 156 degrees F
    FPC 3 EA3_HMC0 DRAM botm       OK         66 degrees C / 150 degrees F
```

```
      FPC 3 EA3_HMC1 Logic die       OK         68 degrees C / 154 degrees F
      FPC 3 EA3_HMC1 DRAM botm       OK         65 degrees C / 149 degrees F
      FPC 3 EA3_HMC2 Logic die       OK         68 degrees C / 154 degrees F
      FPC 3 EA3_HMC2 DRAM botm       OK         65 degrees C / 149 degrees F
      FPC 3 EA4_HMC0 Logic die       OK         81 degrees C / 177 degrees F
      FPC 3 EA4_HMC0 DRAM botm       OK         78 degrees C / 172 degrees F
      FPC 3 EA4_HMC1 Logic die       OK         80 degrees C / 176 degrees F
      FPC 3 EA4_HMC1 DRAM botm       OK         77 degrees C / 170 degrees F
      FPC 3 EA4_HMC2 Logic die       OK         81 degrees C / 177 degrees F
      FPC 3 EA4_HMC2 DRAM botm       OK         78 degrees C / 172 degrees F
      FPC 3 EA5_HMC0 Logic die       OK         68 degrees C / 154 degrees F
      FPC 3 EA5_HMC0 DRAM botm       OK         65 degrees C / 149 degrees F
      FPC 3 EA5_HMC1 Logic die       OK         70 degrees C / 158 degrees F
      FPC 3 EA5_HMC1 DRAM botm       OK         67 degrees C / 152 degrees F
      FPC 3 EA5_HMC2 Logic die       OK         69 degrees C / 156 degrees F
      FPC 3 EA5_HMC2 DRAM botm       OK         66 degrees C / 150 degrees F
Power PEM 0                          OK         29 degrees C / 84 degrees F
      PEM 1                          OK         27 degrees C / 80 degrees F
      PEM 2                          OK         30 degrees C / 86 degrees F
      PEM 3                          Check
      PEM 4                          Check
      PEM 5                          Check
Fans  Fan Tray 0 Fan 0               OK         Spinning at normal speed
      Fan Tray 0 Fan 1               OK         Spinning at normal speed
      Fan Tray 0 Fan 2               OK         Spinning at normal speed
      Fan Tray 0 Fan 3               OK         Spinning at normal speed
      Fan Tray 0 Fan 4               Failed
      Fan Tray 0 Fan 5               Failed
      Fan Tray 0 Fan 6               OK         Spinning at normal speed
      Fan Tray 0 Fan 7               OK         Spinning at normal speed
      Fan Tray 0 Fan 8               OK         Spinning at normal speed
      Fan Tray 0 Fan 9               OK         Spinning at normal speed
      Fan Tray 0 Fan 10              OK         Spinning at normal speed
      Fan Tray 1 Fan 0               OK         Spinning at normal speed
      Fan Tray 1 Fan 1               OK         Spinning at normal speed
      Fan Tray 1 Fan 2               OK         Spinning at normal speed
      Fan Tray 1 Fan 3               OK         Spinning at normal speed
      Fan Tray 1 Fan 4               OK         Spinning at normal speed
      Fan Tray 1 Fan 5               OK         Spinning at normal speed
      Fan Tray 1 Fan 6               OK         Spinning at normal speed
      Fan Tray 1 Fan 7               OK         Spinning at normal speed
      Fan Tray 1 Fan 8               OK         Spinning at normal speed
      Fan Tray 1 Fan 9               OK         Spinning at normal speed
```

```
Fan Tray 1 Fan 10          OK        Spinning at normal speed
SFB 0 Intake-A             OK        32 degrees C / 89 degrees F
SFB 0 Intake-B             OK        21 degrees C / 69 degrees F
SFB 0 Exhaust-A            OK        27 degrees C / 80 degrees F
SFB 0 Exhaust-B            OK        32 degrees C / 89 degrees F
SFB 0 PF0                  OK        39 degrees C / 102 degrees F
SFB 0 PF1                  OK        29 degrees C / 84 degrees F
SFB 1 Intake-A             OK        43 degrees C / 109 degrees F
SFB 1 Intake-B             OK        20 degrees C / 68 degrees F
SFB 1 Exhaust-A            OK        25 degrees C / 77 degrees F
SFB 1 Exhaust-B            OK        44 degrees C / 111 degrees F
SFB 1 PF0                  OK        50 degrees C / 122 degrees F
SFB 1 PF1                  OK        29 degrees C / 84 degrees F
SFB 2 Intake-A             OK        39 degrees C / 102 degrees F
SFB 2 Intake-B             OK        20 degrees C / 68 degrees F
SFB 2 Exhaust-A            OK        25 degrees C / 77 degrees F
SFB 2 Exhaust-B            OK        38 degrees C / 100 degrees F
SFB 2 PF0                  OK        45 degrees C / 113 degrees F
SFB 2 PF1                  OK        30 degrees C / 86 degrees F
SFB 3 Intake-A             OK        36 degrees C / 96 degrees F
SFB 3 Intake-B             OK        20 degrees C / 68 degrees F
SFB 3 Exhaust-A            OK        25 degrees C / 77 degrees F
SFB 3 Exhaust-B            OK        35 degrees C / 95 degrees F
SFB 3 PF0                  OK        42 degrees C / 107 degrees F
SFB 3 PF1                  OK        29 degrees C / 84 degrees F
SFB 4 Intake-A             OK        30 degrees C / 86 degrees F
SFB 4 Intake-B             OK        20 degrees C / 68 degrees F
SFB 4 Exhaust-A            OK        25 degrees C / 77 degrees F
SFB 4 Exhaust-B            OK        31 degrees C / 87 degrees F
SFB 4 PF0                  OK        41 degrees C / 105 degrees F
SFB 4 PF1                  OK        29 degrees C / 84 degrees F
SFB 5 Intake-A             OK        30 degrees C / 86 degrees F
SFB 5 Intake-B             OK        21 degrees C / 69 degrees F
SFB 5 Exhaust-A            OK        25 degrees C / 77 degrees F
SFB 5 Exhaust-B            OK        30 degrees C / 86 degrees F
SFB 5 PF0                  OK        35 degrees C / 95 degrees F
SFB 5 PF1                  OK        34 degrees C / 93 degrees F
```

show chassis environment (MX204 Router)

```
user@host> show chassis environment

Class Item                           Status     Measurement
Temp  CB 0 Top Right Inlet Sensor    OK         35 degrees C / 95 degrees F
      CB 0 Top Left Inlet Sensor     OK         37 degrees C / 98 degrees F
      CB 0 Top Right Exhaust Sensor  OK         43 degrees C / 109 degrees F
      CB 0 Top Left Exhaust Sensor   OK         50 degrees C / 122 degrees F
      CB 0 CPU Core-0 Temp           OK         47 degrees C / 116 degrees F
      CB 0 CPU Core-1 Temp           OK         48 degrees C / 118 degrees F
      CB 0 CPU Core-2 Temp           OK         47 degrees C / 116 degrees F
      CB 0 CPU Core-3 Temp           OK         47 degrees C / 116 degrees F
      CB 0 CPU Core-4 Temp           OK         47 degrees C / 116 degrees F
      CB 0 CPU Core-5 Temp           OK         47 degrees C / 116 degrees F
      CB 0 CPU Core-6 Temp           OK         47 degrees C / 116 degrees F
      CB 0 CPU Core-7 Temp           OK         47 degrees C / 116 degrees F
      FPC 0 EA0_HMC0 Logic die       OK         77 degrees C / 170 degrees F
      FPC 0 EA0_HMC0 DRAM botm       OK         74 degrees C / 165 degrees F
      FPC 0 EA0_HMC1 Logic die       OK         81 degrees C / 177 degrees F
      FPC 0 EA0_HMC1 DRAM botm       OK         78 degrees C / 172 degrees F
      FPC 0 EA0 Chip                 OK         94 degrees C / 201 degrees F
      FPC 0 EA0-XR0 Chip             OK         64 degrees C / 147 degrees F
      FPC 0 EA0-XR1 Chip             OK         65 degrees C / 149 degrees F
Power PEM 0                          Absent
      PEM 1                          OK         48 degrees C / 118 degrees F
Fans  Fan Tray 0 Fan 0               OK         Spinning at normal speed
      Fan Tray 0 Fan 1               OK         Spinning at normal speed
      Fan Tray 1 Fan 0               OK         Spinning at normal speed
      Fan Tray 1 Fan 1               OK         Spinning at normal speed
      Fan Tray 2 Fan 0               OK         Spinning at normal speed
      Fan Tray 2 Fan 1               OK         Spinning at normal speed
```

show chassis environment (T640 Router)

```
user@host> show chassis environment
Class Item                   Status     Measurement
Temp  PEM 0                  Absent
      PEM 1                  OK         22 degrees C / 71 degrees F
      SCG 0                  OK         30 degrees C / 86 degrees F
```

```
      SCG 1                  OK         30 degrees C / 86 degrees F
      Routing Engine 0       Present
      Routing Engine 1       OK         27 degrees C / 80 degrees F
      CB 0                   Present
      CB 1                   OK         33 degrees C / 91 degrees F
      SIB 0                  Absent
      SIB 1                  Absent
      SIB 2                  Absent
      SIB 3                  Absent
      SIB 4                  Absent
      FPC 4 Top              Testing
      FPC 4 Bottom           Testing
      FPC 5 Top              Testing
      FPC 5 Bottom           Testing
      FPC 6 Top              Testing
      FPC 6 Bottom           Testing
      FPM GBUS               OK         23 degrees C / 73 degrees F
      FPM Display            Absent
Fans  Top Left Front fan     OK         Spinning at normal speed
      Top Left Middle fan    OK         Spinning at normal speed
      Top Left Rear fan      OK         Spinning at normal speed
      Top Right Front fan    OK         Spinning at normal speed
      Top Right Middle fan   OK         Spinning at normal speed
      Top Right Rear fan     OK         Spinning at normal speed
      Bottom Left Front fan  OK         Spinning at normal speed
      Bottom Left Middle fan OK         Spinning at normal speed
      Bottom Left Rear fan   OK         Spinning at normal speed
      Bottom Right Front fan OK         Spinning at normal speed
      Bottom Right Middle fan OK        Spinning at normal speed
      Bottom Right Rear fan  OK         Spinning at normal speed
      Fourth Blower from top OK         Spinning at normal speed
      Bottom Blower          OK         Spinning at normal speed
      Middle Blower          OK         Spinning at normal speed
      Top Blower             OK         Spinning at normal speed
      Second Blower from top OK         Spinning at normal speed
Misc  CIP                    OK
      SPMB 0                 OK
      SPMB 1                 OK
```

show chassis environment (T4000 Router)

```
user@host> show chassis environment
Class Item                          Status     Measurement
Temp  PEM 0                         OK         33 degrees C / 91 degrees F
      PEM 1                         Absent
      SCG 0                         OK         33 degrees C / 91 degrees F
      SCG 1                         OK         33 degrees C / 91 degrees F
      Routing Engine 0              OK         33 degrees C / 91 degrees F
      Routing Engine 0 CPU          OK         50 degrees C / 122 degrees F
      Routing Engine 1              OK         32 degrees C / 89 degrees F
      Routing Engine 1 CPU          OK         46 degrees C / 114 degrees F
      CB 0                          OK         32 degrees C / 89 degrees F
      CB 1                          OK         33 degrees C / 91 degrees F
      SIB 0                         OK         42 degrees C / 107 degrees F
      SIB 1                         OK         42 degrees C / 107 degrees F
      SIB 2                         OK         42 degrees C / 107 degrees F
      SIB 3                         OK         43 degrees C / 109 degrees F
      SIB 4                         OK         45 degrees C / 113 degrees F
      FPC 0 Fan Intake              OK         34 degrees C / 93 degrees F
      FPC 0 Fan Exhaust             OK         48 degrees C / 118 degrees F
      FPC 0 PMB                     OK         47 degrees C / 116 degrees F
      FPC 0 LMB0                    OK         50 degrees C / 122 degrees F
      FPC 0 LMB1                    OK         41 degrees C / 105 degrees F
      FPC 0 LMB2                    OK         35 degrees C / 95 degrees F
      FPC 0 PFE1 LU2                OK         46 degrees C / 114 degrees F
      FPC 0 PFE1 LU0                OK         41 degrees C / 105 degrees F
      FPC 0 PFE0 LU0                OK         57 degrees C / 134 degrees F
      FPC 0 XF1                     OK         46 degrees C / 114 degrees F
      FPC 0 XF0                     OK         52 degrees C / 125 degrees F
      FPC 0 XM1                     OK         41 degrees C / 105 degrees F
      FPC 0 XM0                     OK         50 degrees C / 122 degrees F
      FPC 0 PFE0 LU1                OK         56 degrees C / 132 degrees F
      FPC 0 PFE0 LU2                OK         45 degrees C / 113 degrees F
      FPC 0 PFE1 LU1                OK         37 degrees C / 98 degrees F
      FPC 3 Fan Intake              OK         36 degrees C / 96 degrees F
      FPC 3 Fan Exhaust             OK         51 degrees C / 123 degrees F
      FPC 3 PMB                     OK         43 degrees C / 109 degrees F
      FPC 3 LMB0                    OK         57 degrees C / 134 degrees F
      FPC 3 LMB1                    OK         54 degrees C / 129 degrees F
      FPC 3 LMB2                    OK         38 degrees C / 100 degrees F
      FPC 3 PFE1 LU2                OK         63 degrees C / 145 degrees F
```

```
      FPC 3 PFE1 LU0                OK        45 degrees C / 113 degrees F
      FPC 3 PFE0 LU0                OK        69 degrees C / 156 degrees F
      FPC 3 XF1                     OK        62 degrees C / 143 degrees F
      FPC 3 XF0                     OK        63 degrees C / 145 degrees F
      FPC 3 XM1                     OK        43 degrees C / 109 degrees F
      FPC 3 XM0                     OK        67 degrees C / 152 degrees F
      FPC 3 PFE0 LU1                OK        63 degrees C / 145 degrees F
      FPC 3 PFE0 LU2                OK        66 degrees C / 150 degrees F
      FPC 3 PFE1 LU1                OK        41 degrees C / 105 degrees F
      FPC 5 Top                     OK        39 degrees C / 102 degrees F
      FPC 5 Bottom                  OK        38 degrees C / 100 degrees F
      FPC 6 Fan Intake              OK        33 degrees C / 91 degrees F
      FPC 6 Fan Exhaust             OK        49 degrees C / 120 degrees F
      FPC 6 PMB                     OK        40 degrees C / 104 degrees F
      FPC 6 LMB0                    OK        60 degrees C / 140 degrees F
      FPC 6 LMB1                    OK        58 degrees C / 136 degrees F
      FPC 6 LMB2                    OK        40 degrees C / 104 degrees F
      FPC 6 PFE1 LU2                OK        69 degrees C / 156 degrees F
      FPC 6 PFE1 LU0                OK        45 degrees C / 113 degrees F
      FPC 6 PFE0 LU0                OK        71 degrees C / 159 degrees F
      FPC 6 XF1                     OK        58 degrees C / 136 degrees F
      FPC 6 XF0                     OK        65 degrees C / 149 degrees F
      FPC 6 XM1                     OK        39 degrees C / 102 degrees F
      FPC 6 XM0                     OK        66 degrees C / 150 degrees F
      FPC 6 PFE0 LU1                OK        69 degrees C / 156 degrees F
      FPC 6 PFE0 LU2                OK        69 degrees C / 156 degrees F
      FPC 6 PFE1 LU1                OK        42 degrees C / 107 degrees F
      FPM GBUS                      OK        24 degrees C / 75 degrees F
      FPM Display                   OK        27 degrees C / 80 degrees F
Fans  Top Left Front fan            OK        Spinning at high speed
      Top Left Middle fan           OK        Spinning at high speed
      Top Left Rear fan             OK        Spinning at high speed
      Top Right Front fan           OK        Spinning at high speed
      Top Right Middle fan          OK        Spinning at high speed
      Top Right Rear fan            OK        Spinning at high speed
      Bottom Left Front fan         OK        Spinning at high speed
      Bottom Left Middle fan        OK        Spinning at high speed
      Bottom Left Rear fan          OK        Spinning at high speed
      Bottom Right Front fan        OK        Spinning at high speed
      Bottom Right Middle fan       OK        Spinning at high speed
      Bottom Right Rear fan         OK        Spinning at high speed
      Rear Tray Top fan             OK        Spinning at high speed
      Rear Tray Second fan          OK        Spinning at high speed
```

```
      Rear Tray Third fan           OK         Spinning at high speed
      Rear Tray Fourth fan          OK         Spinning at high speed
      Rear Tray Fifth fan           OK         Spinning at high speed
      Rear Tray Sixth fan           OK         Spinning at high speed
      Rear Tray Seventh fan         OK         Spinning at high speed
      Rear Tray Bottom fan          OK         Spinning at high speed
Misc  CIP                           OK
      SPMB 0                        OK
      SPMB 1                        OK
```

show chassis environment (TX Matrix Router)

```
user@host> show chassis environment

---------------------------------------------------------------------------
Class Item                  Status     Measurement
Temp  PEM 0                 Absent
      PEM 1                 OK         29 degrees C / 84 degrees F
      Routing Engine 0      OK         34 degrees C / 93 degrees F
      Routing Engine 1      OK         34 degrees C / 93 degrees F
      CB 0                  OK         32 degrees C / 89 degrees F
      CB 1                  OK         32 degrees C / 89 degrees F
      SIB 0                 OK         44 degrees C / 111 degrees F
      SIB 0 (B)             OK         44 degrees C / 111 degrees F
      FPM GBUS              OK         27 degrees C / 80 degrees F
      FPM Display           OK         32 degrees C / 89 degrees F
Fans  Top Left Front fan    OK         Spinning at normal speed
      Top Left Middle fan   OK         Spinning at normal speed
      Top Left Rear fan     OK         Spinning at normal speed
      Top Right Front fan   OK         Spinning at normal speed
      Top Right Middle fan  OK         Spinning at normal speed
      Top Right Rear fan    OK         Spinning at normal speed
      Bottom Left Front fan OK         Spinning at normal speed
      Bottom Left Middle fan OK        Spinning at normal speed
      Bottom Left Rear fan  OK         Spinning at normal speed
      Bottom Right Front fan OK        Spinning at normal speed
      Bottom Right Middle fan OK       Spinning at normal speed
      Bottom Right Rear fan OK         Spinning at normal speed
      Rear Tray Top fan     OK         Spinning at normal speed
      Rear Tray Second fan  OK         Spinning at normal speed
      Rear Tray Third fan   OK         Spinning at normal speed
```

```
      Rear Tray Fourth fan   OK         Spinning at normal speed
      Rear Tray Fifth fan    OK         Spinning at normal speed
      Rear Tray Sixth fan    OK         Spinning at normal speed
      Rear Tray Seventh fan  OK         Spinning at normal speed
      Rear Tray Bottom fan   OK         Spinning at normal speed
Misc  CIP 0                  OK
      CIP 1                  OK
      SPMB 0                 OK
      SPMB 1                 OK

lcc0-re0:
--------------------------------------------------------------------------
Class Item                   Status     Measurement
Temp  PEM 0                  OK         29 degrees C / 84 degrees F
      PEM 1                  Absent
      SCG 0                  OK         35 degrees C / 95 degrees F
      SCG 1                  Absent
      Routing Engine 0       OK         39 degrees C / 102 degrees F
      Routing Engine 1       OK         36 degrees C / 96 degrees F
      CB 0                   OK         32 degrees C / 89 degrees F
      CB 1                   OK         32 degrees C / 89 degrees F
      SIB 0                  OK         40 degrees C / 104 degrees F
      SIB 0 (B)              OK         51 degrees C / 123 degrees F
      FPC 0 Top              OK         45 degrees C / 113 degrees F
      FPC 0 Bottom           OK         31 degrees C / 87 degrees F
      FPC 1 Top              OK         34 degrees C / 93 degrees F
      FPC 1 Bottom           OK         31 degrees C / 87 degrees F
      FPM GBUS               OK         30 degrees C / 86 degrees F
      FPM Display            OK         34 degrees C / 93 degrees F
Fans  Top Left Front fan     OK         Spinning at normal speed
      Top Left Middle fan    OK         Spinning at normal speed
      Top Left Rear fan      OK         Spinning at normal speed
      Top Right Front fan    OK         Spinning at normal speed
      Top Right Middle fan   OK         Spinning at normal speed
      Top Right Rear fan     OK         Spinning at normal speed
      Bottom Left Front fan  OK         Spinning at normal speed
      Bottom Left Middle fan OK         Spinning at normal speed
      Bottom Left Rear fan   OK         Spinning at normal speed
      Bottom Right Front fan OK         Spinning at normal speed
      Bottom Right Middle fan OK        Spinning at normal speed
      Bottom Right Rear fan  OK         Spinning at normal speed
      Rear Tray Top fan      OK         Spinning at normal speed
      Rear Tray Second fan   OK         Spinning at normal speed
```

```
      Rear Tray Third fan    OK         Spinning at normal speed
      Rear Tray Fourth fan   OK         Spinning at normal speed
      Rear Tray Fifth fan    OK         Spinning at normal speed
      Rear Tray Sixth fan    OK         Spinning at normal speed
      Rear Tray Seventh fan  OK         Spinning at normal speed
      Rear Tray Bottom fan   OK         Spinning at normal speed
Misc  CIP                    OK
      SPMB 0                 OK
      SPMB 1                 OK

lcc2-re0:
--------------------------------------------------------------------------
Class Item                   Status     Measurement
Temp  PEM 0                  OK         29 degrees C / 84 degrees F
      PEM 1                  Absent
      SCG 0                  OK         32 degrees C / 89 degrees F
      SCG 1                  Absent
      Routing Engine 0       OK         31 degrees C / 87 degrees F
      Routing Engine 1       OK         32 degrees C / 89 degrees F
      CB 0                   OK         30 degrees C / 86 degrees F
      SIB 0                  OK         38 degrees C / 100 degrees F
      SIB 0 (B)              OK         49 degrees C / 120 degrees F
      FPC 0 Top              OK         45 degrees C / 113 degrees F
      FPC 0 Bottom           OK         33 degrees C / 91 degrees F
      FPC 1 Top              OK         37 degrees C / 98 degrees F
      FPC 1 Bottom           OK         33 degrees C / 91 degrees F
      FPM GBUS               OK         30 degrees C / 86 degrees F
      FPM Display            OK         34 degrees C / 93 degrees F
Fans  Top Left Front fan     OK         Spinning at normal speed
      Top Left Middle fan    OK         Spinning at normal speed
...
```

show chassis environment (T1600 Router)

```
user@host> show chassis environment
Class Item                           Status     Measurement
Temp  PEM 0                          OK         27 degrees C / 80 degrees F
      PEM 1                          Absent
      SCG 0                          OK         31 degrees C / 87 degrees F
      SCG 1                          OK         35 degrees C / 95 degrees F
      Routing Engine 0               OK         30 degrees C / 86 degrees F
```

```
      Routing Engine 1              OK        30 degrees C / 86 degrees F
      CB 0                          OK        31 degrees C / 87 degrees F
      CB 1                          OK        31 degrees C / 87 degrees F
      SIB 0                         OK        41 degrees C / 105 degrees F
      SIB 0 (B)                     OK        34 degrees C / 93 degrees F
      SIB 1                         OK        0 degrees C / 32 degrees F
      SIB 1 (B)                     OK        0 degrees C / 32 degrees F
      SIB 2                         OK        0 degrees C / 32 degrees F
      SIB 2 (B)                     OK        0 degrees C / 32 degrees F
      SIB 3                         OK        0 degrees C / 32 degrees F
      SIB 3 (B)                     OK        0 degrees C / 32 degrees F
      SIB 4                         OK        0 degrees C / 32 degrees F
      SIB 4 (B)                     OK        0 degrees C / 32 degrees F
      FPC 0 Top                     OK        49 degrees C / 120 degrees F
      FPC 0 Bottom                  OK        50 degrees C / 122 degrees F
      FPC 1 Top                     OK        48 degrees C / 118 degrees F
      FPC 1 Bottom                  OK        49 degrees C / 120 degrees F
      FPM GBUS                      OK        27 degrees C / 80 degrees F
      FPM Display                   OK        30 degrees C / 86 degrees F
Fans  Top Left Front fan            OK        Spinning at normal speed
      Top Left Middle fan           OK        Spinning at normal speed
      Top Left Rear fan             OK        Spinning at normal speed
      Top Right Front fan           OK        Spinning at normal speed
      Top Right Middle fan          OK        Spinning at normal speed
      Top Right Rear fan            OK        Spinning at normal speed
      Bottom Left Front fan         OK        Spinning at normal speed
      Bottom Left Middle fan        OK        Spinning at normal speed
      Bottom Left Rear fan          OK        Spinning at normal speed
      Bottom Right Front fan        OK        Spinning at normal speed
      Bottom Right Middle fan       OK        Spinning at normal speed
      Bottom Right Rear fan         OK        Spinning at normal speed
      Rear Tray Top fan             OK        Spinning at normal speed
      Rear Tray Second fan          OK        Spinning at normal speed
      Rear Tray Third fan           OK        Spinning at normal speed
      Rear Tray Fourth fan          OK        Spinning at normal speed
      Rear Tray Fifth fan           OK        Spinning at normal speed
      Rear Tray Sixth fan           OK        Spinning at normal speed
      Rear Tray Seventh fan         OK        Spinning at normal speed
      Rear Tray Bottom fan          OK        Spinning at normal speed
Misc  CIP                           OK
      SPMB 0                        OK
      SPMB 1                        OK
```

show chassis environment (TX Matrix Plus Router)

```
user@host> show chassis environment

---------------------------------------------------------------------------
Class Item                          Status     Measurement
Temp  PEM 0                         OK         28 degrees C / 82 degrees F
      PEM 1                         Absent
      Routing Engine 0              OK         27 degrees C / 80 degrees F
      Routing Engine 1              OK         29 degrees C / 84 degrees F
      CB 0 Intake                   OK         26 degrees C / 78 degrees F
      CB 0 Exhaust A                OK         25 degrees C / 77 degrees F
      CB 0 Exhaust B                OK         25 degrees C / 77 degrees F
      CB 1 Intake                   OK         26 degrees C / 78 degrees F
      CB 1 Exhaust A                OK         26 degrees C / 78 degrees F
      CB 1 Exhaust B                OK         26 degrees C / 78 degrees F
      SIB F13 0                     OK         47 degrees C / 116 degrees F
      SIB F13 0 (B)                 OK         48 degrees C / 118 degrees F
      SIB F13 1                     OK         38 degrees C / 100 degrees F
      SIB F13 1 (B)                 OK         37 degrees C / 98 degrees F
      SIB F2S 0/0                   OK         27 degrees C / 80 degrees F
      SIB F2S 0/2                   OK         28 degrees C / 82 degrees F
      SIB F2S 0/4                   OK         27 degrees C / 80 degrees F
      SIB F2S 0/6                   OK         28 degrees C / 82 degrees F
      SIB F2S 1/0                   OK         26 degrees C / 78 degrees F
      SIB F2S 1/2                   OK         26 degrees C / 78 degrees F
      SIB F2S 1/4                   OK         26 degrees C / 78 degrees F
      SIB F2S 1/6                   OK         26 degrees C / 78 degrees F
      SIB F2S 2/0                   OK         25 degrees C / 77 degrees F
      SIB F2S 2/2                   OK         25 degrees C / 77 degrees F
      SIB F2S 2/4                   OK         23 degrees C / 73 degrees F
      CIP 0 Intake                  OK         23 degrees C / 73 degrees F
      CIP 0 Exhaust A               OK         24 degrees C / 75 degrees F
      CIP 0 Exhaust B               OK         24 degrees C / 75 degrees F
      CIP 1 Intake                  OK         24 degrees C / 75 degrees F
      CIP 1 Exhaust A               OK         25 degrees C / 77 degrees F
      CIP 1 Exhaust B               OK         25 degrees C / 77 degrees F
Fans  Fan Tray 0 Fan 1              OK         Spinning at normal speed
      Fan Tray 0 Fan 2              OK         Spinning at normal speed
      Fan Tray 0 Fan 3              OK         Spinning at normal speed
      Fan Tray 0 Fan 4              OK         Spinning at normal speed
      Fan Tray 0 Fan 5              OK         Spinning at normal speed
```

```
Fan Tray 0 Fan 6          OK        Spinning at normal speed
Fan Tray 1 Fan 1          OK        Spinning at normal speed
Fan Tray 1 Fan 2          OK        Spinning at normal speed
Fan Tray 1 Fan 3          OK        Spinning at normal speed
Fan Tray 1 Fan 4          OK        Spinning at normal speed
Fan Tray 1 Fan 5          OK        Spinning at normal speed
Fan Tray 1 Fan 6          OK        Spinning at normal speed
Fan Tray 2 Fan 1          OK        Spinning at normal speed
Fan Tray 2 Fan 2          OK        Spinning at normal speed
Fan Tray 2 Fan 3          OK        Spinning at normal speed
Fan Tray 2 Fan 4          OK        Spinning at normal speed
Fan Tray 2 Fan 5          OK        Spinning at normal speed
Fan Tray 2 Fan 6          OK        Spinning at normal speed
Fan Tray 2 Fan 7          OK        Spinning at normal speed
Fan Tray 2 Fan 8          OK        Spinning at normal speed
Fan Tray 2 Fan 9          OK        Spinning at normal speed
Fan Tray 3 Fan 1          OK        Spinning at normal speed
Fan Tray 3 Fan 2          OK        Spinning at normal speed
Fan Tray 3 Fan 3          OK        Spinning at normal speed
Fan Tray 3 Fan 4          OK        Spinning at normal speed
Fan Tray 3 Fan 5          OK        Spinning at normal speed
Fan Tray 3 Fan 6          OK        Spinning at normal speed
Fan Tray 3 Fan 7          OK        Spinning at normal speed
Fan Tray 3 Fan 8          OK        Spinning at normal speed
Fan Tray 3 Fan 9          OK        Spinning at normal speed
Fan Tray 4 Fan 1          OK        Spinning at normal speed
Fan Tray 4 Fan 2          OK        Spinning at normal speed
Fan Tray 4 Fan 3          OK        Spinning at normal speed
Fan Tray 4 Fan 4          OK        Spinning at normal speed
Fan Tray 4 Fan 5          OK        Spinning at normal speed
Fan Tray 4 Fan 6          OK        Spinning at normal speed
Fan Tray 4 Fan 7          OK        Spinning at normal speed
Fan Tray 4 Fan 8          OK        Spinning at normal speed
Fan Tray 4 Fan 9          OK        Spinning at normal speed
Fan Tray 5 Fan 1          OK        Spinning at normal speed
Fan Tray 5 Fan 2          OK        Spinning at normal speed
Fan Tray 5 Fan 3          OK        Spinning at normal speed
Fan Tray 5 Fan 4          OK        Spinning at normal speed
Fan Tray 5 Fan 5          OK        Spinning at normal speed
Fan Tray 5 Fan 6          OK        Spinning at normal speed
Fan Tray 5 Fan 7          OK        Spinning at normal speed
Fan Tray 5 Fan 8          OK        Spinning at normal speed
Fan Tray 5 Fan 9          OK        Spinning at normal speed
```

```
Misc  SPMB 0                        OK
      SPMB 1                        OK

lcc0-re0:
--------------------------------------------------------------------------
Class Item                          Status     Measurement
Temp  PEM 0                         OK         27 degrees C / 80 degrees F
      PEM 1                         Absent
      SCG 0                         OK         31 degrees C / 87 degrees F
      SCG 1                         OK         35 degrees C / 95 degrees F
      Routing Engine 0              OK         30 degrees C / 86 degrees F
      Routing Engine 1              OK         30 degrees C / 86 degrees F
      CB 0                          OK         31 degrees C / 87 degrees F
      CB 1                          OK         31 degrees C / 87 degrees F
      SIB 0                         OK         41 degrees C / 105 degrees F
      SIB 0 (B)                     OK         34 degrees C / 93 degrees F
      SIB 1                         OK         0 degrees C / 32 degrees F
      SIB 1 (B)                     OK         0 degrees C / 32 degrees F
      SIB 2                         OK         0 degrees C / 32 degrees F
      SIB 2 (B)                     OK         0 degrees C / 32 degrees F
      SIB 3                         OK         0 degrees C / 32 degrees F
      SIB 3 (B)                     OK         0 degrees C / 32 degrees F
      SIB 4                         OK         0 degrees C / 32 degrees F
      SIB 4 (B)                     OK         0 degrees C / 32 degrees F
      FPC 0 Top                     OK         49 degrees C / 120 degrees F
      FPC 0 Bottom                  OK         50 degrees C / 122 degrees F
      FPC 1 Top                     OK         48 degrees C / 118 degrees F
      FPC 1 Bottom                  OK         49 degrees C / 120 degrees F
      FPM GBUS                      OK         27 degrees C / 80 degrees F
      FPM Display                   OK         30 degrees C / 86 degrees F
Fans  Top Left Front fan            OK         Spinning at normal speed
      Top Left Middle fan           OK         Spinning at normal speed
      Top Left Rear fan             OK         Spinning at normal speed
      Top Right Front fan           OK         Spinning at normal speed
      Top Right Middle fan          OK         Spinning at normal speed
      Top Right Rear fan            OK         Spinning at normal speed
      Bottom Left Front fan         OK         Spinning at normal speed
      Bottom Left Middle fan        OK         Spinning at normal speed
      Bottom Left Rear fan          OK         Spinning at normal speed
      Bottom Right Front fan        OK         Spinning at normal speed
      Bottom Right Middle fan       OK         Spinning at normal speed
      Bottom Right Rear fan         OK         Spinning at normal speed
      Rear Tray Top fan             OK         Spinning at normal speed
```

```
      Rear Tray Second fan          OK         Spinning at normal speed
      Rear Tray Third fan           OK         Spinning at normal speed
      Rear Tray Fourth fan          OK         Spinning at normal speed
      Rear Tray Fifth fan           OK         Spinning at normal speed
      Rear Tray Sixth fan           OK         Spinning at normal speed
      Rear Tray Seventh fan         OK         Spinning at normal speed
      Rear Tray Bottom fan          OK         Spinning at normal speed
Misc  CIP                           OK
      SPMB 0                        OK
      SPMB 1                        OK
```

show chassis environment (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis environment

---------------------------------------------------------------------------
Class Item                          Status     Measurement
Temp  PEM 0                         Check      30 degrees C / 86 degrees F
      PEM 1                         OK         33 degrees C / 91 degrees F
      Routing Engine 0              OK         28 degrees C / 82 degrees F
      Routing Engine 0 CPU          OK         42 degrees C / 107 degrees F
      Routing Engine 1              OK         29 degrees C / 84 degrees F
      Routing Engine 1 CPU          OK         44 degrees C / 111 degrees F
      CB 0 Intake                   OK         30 degrees C / 86 degrees F
      CB 0 Exhaust A                OK         28 degrees C / 82 degrees F
      CB 0 Exhaust B                OK         30 degrees C / 86 degrees F
      CB 1 Intake                   OK         31 degrees C / 87 degrees F
      CB 1 Exhaust A                OK         27 degrees C / 80 degrees F
      CB 1 Exhaust B                OK         31 degrees C / 87 degrees F
      SIB F13 0 Board               OK         44 degrees C / 111 degrees F
      SIB F13 0 XF Junction         OK         62 degrees C / 143 degrees F
      SIB F13 3 Board               OK         45 degrees C / 113 degrees F
      SIB F13 3 XF Junction         OK         60 degrees C / 140 degrees F
      SIB F13 6 Board               OK         47 degrees C / 116 degrees F
      SIB F13 6 XF Junction         OK         62 degrees C / 143 degrees F
      SIB F2S 0/0 Board             OK         32 degrees C / 89 degrees F
      SIB F2S 0/0 XF Junction       OK         42 degrees C / 107 degrees F
      SIB F2S 0/2 Board             OK         31 degrees C / 87 degrees F
      SIB F2S 0/2 XF Junction       OK         41 degrees C / 105 degrees F
      SIB F2S 0/4 Board             OK         31 degrees C / 87 degrees F
      SIB F2S 0/4 XF Junction       OK         42 degrees C / 107 degrees F
```

```
       SIB F2S 0/6 Board             OK         31 degrees C / 87 degrees F
       SIB F2S 0/6 XF Junction       OK         41 degrees C / 105 degrees F
       SIB F2S 1/0 Board             OK         31 degrees C / 87 degrees F
       SIB F2S 1/0 XF Junction       OK         41 degrees C / 105 degrees F
       SIB F2S 1/2 Board             OK         29 degrees C / 84 degrees F
       SIB F2S 1/2 XF Junction       OK         39 degrees C / 102 degrees F
       SIB F2S 1/4 Board             OK         29 degrees C / 84 degrees F
       SIB F2S 1/4 XF Junction       OK         35 degrees C / 95 degrees F
       SIB F2S 1/6 Board             OK         30 degrees C / 86 degrees F
       SIB F2S 1/6 XF Junction       OK         41 degrees C / 105 degrees F
       SIB F2S 2/0 Board             OK         30 degrees C / 86 degrees F
       SIB F2S 2/0 XF Junction       OK         42 degrees C / 107 degrees F
       SIB F2S 2/2 Board             OK         28 degrees C / 82 degrees F
       SIB F2S 2/2 XF Junction       OK         39 degrees C / 102 degrees F
       SIB F2S 2/4 Board             OK         29 degrees C / 84 degrees F
       SIB F2S 2/4 XF Junction       OK         42 degrees C / 107 degrees F
       SIB F2S 2/6 Board             OK         29 degrees C / 84 degrees F
       SIB F2S 2/6 XF Junction       OK         41 degrees C / 105 degrees F
       CIP 0 Intake                  OK         25 degrees C / 77 degrees F
       CIP 0 Exhaust A               OK         26 degrees C / 78 degrees F
       CIP 0 Exhaust B               OK         26 degrees C / 78 degrees F
       CIP 1 Intake                  OK         26 degrees C / 78 degrees F
       CIP 1 Exhaust A               OK         27 degrees C / 80 degrees F
       CIP 1 Exhaust B               OK         27 degrees C / 80 degrees F
Fans   Fan Tray 0 Fan 1              OK         Spinning at normal speed
       Fan Tray 0 Fan 2              OK         Spinning at normal speed
       Fan Tray 0 Fan 3              OK         Spinning at normal speed
       Fan Tray 0 Fan 4              OK         Spinning at normal speed
       Fan Tray 0 Fan 5              OK         Spinning at normal speed
       Fan Tray 0 Fan 6              OK         Spinning at normal speed
       Fan Tray 1 Fan 1              OK         Spinning at normal speed
       Fan Tray 1 Fan 2              OK         Spinning at normal speed
       Fan Tray 1 Fan 3              OK         Spinning at normal speed
       Fan Tray 1 Fan 4              OK         Spinning at normal speed
       Fan Tray 1 Fan 5              OK         Spinning at normal speed
       Fan Tray 1 Fan 6              OK         Spinning at normal speed
       Fan Tray 2 Fan 1              OK         Spinning at normal speed
       Fan Tray 2 Fan 2              OK         Spinning at normal speed
       Fan Tray 2 Fan 3              OK         Spinning at normal speed
       Fan Tray 2 Fan 4              OK         Spinning at normal speed
       Fan Tray 2 Fan 5              OK         Spinning at normal speed
       Fan Tray 2 Fan 6              OK         Spinning at normal speed
       Fan Tray 2 Fan 7              OK         Spinning at normal speed
```

```
      Fan Tray 2 Fan 8              OK         Spinning at normal speed
      Fan Tray 2 Fan 9              OK         Spinning at normal speed
      Fan Tray 3 Fan 1              OK         Spinning at normal speed
      Fan Tray 3 Fan 2              OK         Spinning at normal speed
      Fan Tray 3 Fan 3              OK         Spinning at normal speed
      Fan Tray 3 Fan 4              OK         Spinning at normal speed
      Fan Tray 3 Fan 5              OK         Spinning at normal speed
      Fan Tray 3 Fan 6              OK         Spinning at normal speed
      Fan Tray 3 Fan 7              OK         Spinning at normal speed
      Fan Tray 3 Fan 8              OK         Spinning at normal speed
      Fan Tray 3 Fan 9              OK         Spinning at normal speed
      Fan Tray 4 Fan 1              OK         Spinning at normal speed
      Fan Tray 4 Fan 2              OK         Spinning at normal speed
      Fan Tray 4 Fan 3              OK         Spinning at normal speed
      Fan Tray 4 Fan 4              OK         Spinning at normal speed
      Fan Tray 4 Fan 5              OK         Spinning at normal speed
      Fan Tray 4 Fan 6              OK         Spinning at normal speed
      Fan Tray 4 Fan 7              OK         Spinning at normal speed
      Fan Tray 4 Fan 8              OK         Spinning at normal speed
      Fan Tray 4 Fan 9              OK         Spinning at normal speed
      Fan Tray 5 Fan 1              OK         Spinning at normal speed
      Fan Tray 5 Fan 2              OK         Spinning at normal speed
      Fan Tray 5 Fan 3              OK         Spinning at normal speed
      Fan Tray 5 Fan 4              OK         Spinning at normal speed
      Fan Tray 5 Fan 5              OK         Spinning at normal speed
      Fan Tray 5 Fan 6              OK         Spinning at normal speed
      Fan Tray 5 Fan 7              OK         Spinning at normal speed
      Fan Tray 5 Fan 8              OK         Spinning at normal speed
      Fan Tray 5 Fan 9              Check
Misc  SPMB 0                        OK
      SPMB 1                        OK

lcc0-re0:
--------------------------------------------------------------------------
Class Item                          Status     Measurement
Temp  PEM 0                         OK         29 degrees C / 84 degrees F
      PEM 1                         Check      29 degrees C / 84 degrees F
      SCG 0                         OK         32 degrees C / 89 degrees F
      SCG 1                         OK         33 degrees C / 91 degrees F
      Routing Engine 0              OK         32 degrees C / 89 degrees F
      Routing Engine 0 CPU          OK         51 degrees C / 123 degrees F
      Routing Engine 1              OK         32 degrees C / 89 degrees F
      Routing Engine 1 CPU          OK         49 degrees C / 120 degrees F
```

```
       CB 0                          OK         34 degrees C / 93 degrees F
       CB 1                          OK         34 degrees C / 93 degrees F
       SIB 0                         OK         39 degrees C / 102 degrees F
       SIB 0 (B)                     Absent
       SIB 1                         OK         39 degrees C / 102 degrees F
       SIB 1 (B)                     Absent
       SIB 2                         OK         39 degrees C / 102 degrees F
       SIB 2 (B)                     Absent
       FPC 4 Top                     OK         43 degrees C / 109 degrees F
       FPC 4 Bottom                  OK         43 degrees C / 109 degrees F
       FPC 7 Fan Intake              OK         35 degrees C / 95 degrees F
       FPC 7 Fan Exhaust             OK         50 degrees C / 122 degrees F
       FPC 7 PMB                     OK         50 degrees C / 122 degrees F
       FPC 7 LMB0                    OK         55 degrees C / 131 degrees F
       FPC 7 LMB1                    OK         49 degrees C / 120 degrees F
       FPC 7 LMB2                    OK         39 degrees C / 102 degrees F
       FPC 7 PFE1 LU2                OK         55 degrees C / 131 degrees F
       FPC 7 PFE1 LU0                OK         45 degrees C / 113 degrees F
       FPC 7 PFE0 LU0                OK         62 degrees C / 143 degrees F
       FPC 7 XF1                     OK         52 degrees C / 125 degrees F
       FPC 7 XF0                     OK         61 degrees C / 141 degrees F
       FPC 7 XM1                     OK         39 degrees C / 102 degrees F
       FPC 7 XM0                     OK         56 degrees C / 132 degrees F
       FPC 7 PFE0 LU1                OK         60 degrees C / 140 degrees F
       FPC 7 PFE0 LU2                OK         55 degrees C / 131 degrees F
       FPC 7 PFE1 LU1                OK         41 degrees C / 105 degrees F
       FPM GBUS                      OK         24 degrees C / 75 degrees F
       FPM Display                   OK         28 degrees C / 82 degrees F
Fans   Top Left Front fan            OK         Spinning at normal speed
       Top Left Middle fan           OK         Spinning at normal speed
       Top Left Rear fan             OK         Spinning at normal speed
       Top Right Front fan           OK         Spinning at normal speed
       Top Right Middle fan          OK         Spinning at normal speed
       Top Right Rear fan            OK         Spinning at normal speed
       Bottom Left Front fan         OK         Spinning at normal speed
       Bottom Left Middle fan        OK         Spinning at normal speed
       Bottom Left Rear fan          OK         Spinning at normal speed
       Bottom Right Front fan        OK         Spinning at normal speed
       Bottom Right Middle fan       OK         Spinning at normal speed
       Bottom Right Rear fan         OK         Spinning at normal speed
       Rear Tray fan 1 (Top)         OK         Spinning at normal speed
       Rear Tray fan 2               OK         Spinning at normal speed
       Rear Tray fan 3               OK         Spinning at normal speed
```

```
      Rear Tray fan 4                OK         Spinning at normal speed
      Rear Tray fan 5                OK         Spinning at normal speed
      Rear Tray fan 6                OK         Spinning at normal speed
      Rear Tray fan 7                OK         Spinning at normal speed
      Rear Tray fan 8                OK         Spinning at normal speed
      Rear Tray fan 9                OK         Spinning at normal speed
      Rear Tray fan 10               OK         Spinning at normal speed
      Rear Tray fan 11               OK         Spinning at normal speed
      Rear Tray fan 12               OK         Spinning at normal speed
      Rear Tray fan 13               OK         Spinning at normal speed
      Rear Tray fan 14               OK         Spinning at normal speed
      Rear Tray fan 15               OK         Spinning at normal speed
      Rear Tray fan 16 (Bottom)      OK         Spinning at normal speed
Misc  CIP                            OK
      SPMB 0                         OK
      SPMB 1                         OK
```

show chassis environment (EX4200 Standalone Switch)

```
user@switch> show chassis environment
Class Item                           Status     Measurement
Power FPC 0 Power Supply 0           OK
      FPC 0 Power Supply 1           Absent
Temp  FPC 0 CPU                      OK         41 degrees C / 105 degrees F
      FPC 0 EX-PFE1                  OK         42 degrees C / 107 degrees F
      FPC 0 EX-PFE2                  OK         46 degrees C / 114 degrees F
      FPC 0 GEPHY Front Left         OK         25 degrees C / 77 degrees F
      FPC 0 GEPHY Front Right        OK         27 degrees C / 80 degrees F
      FPC 0 Uplink Conn              OK         29 degrees C / 84 degrees F
Fans  FPC 0 Fan 1                    OK         Spinning at normal speed
      FPC 0 Fan 2                    OK         Spinning at normal speed
      FPC 0 Fan 3                    OK         Spinning at normal speed
```

show chassis environment (EX8216 Switch)

```
user@switch> show chassis environment
Class Item                           Status     Measurement
Power PSU 0                          OK
      PSU 1                          OK
      PSU 2                          OK
```

```
      PSU 3                          Check
      PSU 4                          Absent
      PSU 5                          Absent
Temp  CB 0 Intake                    OK         23 degrees C / 73 degrees F
      CB 0 Exhaust                   OK         26 degrees C / 78 degrees F
      CB 1 Intake                    OK         22 degrees C / 71 degrees F
      CB 1 Exhaust                   OK         25 degrees C / 77 degrees F
      FPC 4 Intake                   OK         49 degrees C / 120 degrees F
      FPC 4 Exhaust                  OK         59 degrees C / 138 degrees F
      SIB 5 Intake                   OK         25 degrees C / 77 degrees F
      SIB 5 Exhaust                  OK         35 degrees C / 95 degrees F
      SIB 6 Intake                   OK         25 degrees C / 77 degrees F
      SIB 6 Exhaust                  OK         38 degrees C / 100 degrees F
Fans  Top Fan 1                      OK         Spinning at normal speed
      Top Fan 2                      OK         Spinning at normal speed
      Top Fan 3                      OK         Spinning at normal speed
      Top Fan 4                      OK         Spinning at normal speed
      Top Fan 5                      OK         Spinning at normal speed
      Top Fan 6                      OK         Spinning at normal speed
      Top Fan 7                      OK         Spinning at normal speed
      Top Fan 8                      OK         Spinning at normal speed
      Top Fan 9                      OK         Spinning at normal speed
      Bottom Fan 1                   OK         Spinning at normal speed
      Bottom Fan 2                   OK         Spinning at normal speed
      Bottom Fan 3                   OK         Spinning at normal speed
      Bottom Fan 4                   OK         Spinning at normal speed
      Bottom Fan 5                   OK         Spinning at normal speed
      Bottom Fan 6                   OK         Spinning at normal speed
      Bottom Fan 7                   OK         Spinning at normal speed
      Bottom Fan 8                   OK         Spinning at normal speed
      Bottom Fan 9                   OK         Spinning at normal speed
```

show chassis environment (EX9200 Switch)

```
user@switch> show chassis environment

Class Item                           Status     Measurement
Temp  PEM 0                          Check
      PEM 1                          OK         40 degrees C / 104 degrees F
      PEM 2                          OK         40 degrees C / 104 degrees F
      PEM 3                          Absent
```

```
      Routing Engine 0              OK         35 degrees C / 95 degrees F
      Routing Engine 0 CPU          OK         33 degrees C / 91 degrees F
      Routing Engine 1              OK         38 degrees C / 100 degrees F
      Routing Engine 1 CPU          OK         33 degrees C / 91 degrees F
      CB 0 Intake                   OK         35 degrees C / 95 degrees F
      CB 0 Exhaust A                OK         33 degrees C / 91 degrees F
      CB 0 Exhaust B                OK         40 degrees C / 104 degrees F
      CB 0 ACBC                     OK         39 degrees C / 102 degrees F
      CB 0 XF A                     OK         49 degrees C / 120 degrees F
      CB 0 XF B                     OK         46 degrees C / 114 degrees F
      CB 1 Intake                   OK         37 degrees C / 98 degrees F
      CB 1 Exhaust A                OK         32 degrees C / 89 degrees F
      CB 1 Exhaust B                OK         39 degrees C / 102 degrees F
      CB 1 ACBC                     OK         41 degrees C / 105 degrees F
      CB 1 XF A                     OK         49 degrees C / 120 degrees F
      CB 1 XF B                     OK         49 degrees C / 120 degrees F
      FPC 2 Intake                  OK         37 degrees C / 98 degrees F
      FPC 2 Exhaust A               OK         40 degrees C / 104 degrees F
      FPC 2 Exhaust B               OK         34 degrees C / 93 degrees F
      FPC 2 LU 0 TCAM TSen          OK         44 degrees C / 111 degrees F
      FPC 2 LU 0 TCAM Chip          OK         48 degrees C / 118 degrees F
      FPC 2 LU 0 TSen               OK         44 degrees C / 111 degrees F
      FPC 2 LU 0 Chip               OK         60 degrees C / 140 degrees F
      FPC 2 MQ 0 TSen               OK         44 degrees C / 111 degrees F
      FPC 2 MQ 0 Chip               OK         51 degrees C / 123 degrees F
      FPC 3 Intake                  OK         39 degrees C / 102 degrees F
      FPC 3 Exhaust A               OK         51 degrees C / 123 degrees F

[...Output truncated...]

     Fans  Top Rear Fan             OK         Spinning at intermediate-speed
     Bottom Rear Fan                OK         Spinning at intermediate-speed
     Top Middle Fan                 OK         Spinning at intermediate-speed
     Bottom Middle Fan              OK         Spinning at intermediate-speed
     Top Front Fan                  OK         Spinning at intermediate-speed
     Bottom Front Fan               OK         Spinning at intermediate-speed
```

show chassis environment (EX9251 Switch)

```
user@switch> show chassis environment
```

```
Class Item                          Status     Measurement
Temp  CB 0 Top Right Inlet Sensor   OK         29 degrees C / 84 degrees F
      CB 0 Top Left Inlet Sensor    OK         29 degrees C / 84 degrees F
      CB 0 Top Right Exhaust Sensor OK         40 degrees C / 104 degrees F
      CB 0 Top Left Exhaust Sensor  OK         59 degrees C / 138 degrees F
      CB 0 CPU Core-0 Temp          OK         45 degrees C / 113 degrees F
      CB 0 CPU Core-1 Temp          OK         44 degrees C / 111 degrees F
      CB 0 CPU Core-2 Temp          OK         44 degrees C / 111 degrees F
      CB 0 CPU Core-3 Temp          OK         44 degrees C / 111 degrees F
      CB 0 CPU Core-4 Temp          OK         45 degrees C / 113 degrees F
      CB 0 CPU Core-5 Temp          OK         44 degrees C / 111 degrees F
      CB 0 CPU Core-6 Temp          OK         44 degrees C / 111 degrees F
      CB 0 CPU Core-7 Temp          OK         43 degrees C / 109 degrees F
Power PEM 0                         Check
      PEM 1                         OK         36 degrees C / 96 degrees F
Fans  Fan Tray 0 Fan 0              OK         Spinning at normal speed
      Fan Tray 0 Fan 1              OK         Spinning at normal speed
      Fan Tray 1 Fan 0              OK         Spinning at normal speed
      Fan Tray 1 Fan 1              OK         Spinning at normal speed
      Fan Tray 2 Fan 0              Absent
      Fan Tray 2 Fan 1              Absent
```

show chassis environment (EX9253 Switch)

```
user@switch> show chassis environment

Class Item                          Status     Measurement
Temp  CB 0 Exhaust Temp Sensor      OK         37 degrees C / 98 degrees F
      CB 0 Inlet Temp Sensor        OK         31 degrees C / 87 degrees F
      CB 0 CPU DIE Temp Sensor      OK         42 degrees C / 107 degrees F
      CB 1 Exhaust Temp Sensor      OK         31 degrees C / 87 degrees F
      CB 1 Inlet Temp Sensor        OK         28 degrees C / 82 degrees F
      CB 1 CPU DIE Temp Sensor      OK         42 degrees C / 107 degrees F
      FPC 0 Intake Temp Sensor      OK         31 degrees C / 87 degrees F
      FPC 0 Exhaust-A Temp Sensor   OK         58 degrees C / 136 degrees F
      FPC 0 Exhaust-B Temp Sensor   OK         47 degrees C / 116 degrees F
      FPC 1 Intake Temp Sensor      OK         29 degrees C / 84 degrees F
      FPC 1 Exhaust-A Temp Sensor   OK         59 degrees C / 138 degrees F
      FPC 1 Exhaust-B Temp Sensor   OK         48 degrees C / 118 degrees F
Power PEM 0                         OK         54 degrees C / 129 degrees F
      PEM 1                         Check
```

```
      PEM 2                          Absent
      PEM 3                          Absent
      PEM 4                          Check
      PEM 5                          OK         61 degrees C / 141 degrees F
Fans  Fan Tray 0 Fan 0               OK         Spinning at normal speed
      Fan Tray 0 Fan 1               OK         Spinning at normal speed
      Fan Tray 0 Fan 2               OK         Spinning at normal speed
      Fan Tray 0 Fan 3               OK         Spinning at normal speed
      Fan Tray 1 Fan 0               OK         Spinning at normal speed
      Fan Tray 1 Fan 1               OK         Spinning at normal speed
      Fan Tray 1 Fan 2               OK         Spinning at normal speed
      Fan Tray 1 Fan 3               OK         Spinning at normal speed
      Fan Tray 2 Fan 0               OK         Spinning at normal speed
      Fan Tray 2 Fan 1               OK         Spinning at normal speed
      Fan Tray 2 Fan 2               OK         Spinning at normal speed
      Fan Tray 2 Fan 3               OK         Spinning at normal speed
      Fan Tray 3 Fan 0               OK         Spinning at normal speed
      Fan Tray 3 Fan 1               OK         Spinning at normal speed
      Fan Tray 3 Fan 2               OK         Spinning at normal speed
      Fan Tray 3 Fan 3               OK         Spinning at normal speed
```

show chassis environment (QFX Series and OCX Series)

```
user@switch> show chassis environment
Class Item                           Status     Measurement
Temp  CB 0 Top Right Inlet Sensor    OK         29 degrees C / 84 degrees F
      CB 0 Top Left Inlet Sensor     OK         29 degrees C / 84 degrees F
      CB 0 Top Right Exhaust Sensor  OK         40 degrees C / 104 degrees F
      CB 0 Top Left Exhaust Sensor   OK         59 degrees C / 138 degrees F
      CB 0 CPU Core-0 Temp           OK         45 degrees C / 113 degrees F
      CB 0 CPU Core-1 Temp           OK         44 degrees C / 111 degrees F
      CB 0 CPU Core-2 Temp           OK         44 degrees C / 111 degrees F
      CB 0 CPU Core-3 Temp           OK         44 degrees C / 111 degrees F
      CB 0 CPU Core-4 Temp           OK         45 degrees C / 113 degrees F
      CB 0 CPU Core-5 Temp           OK         44 degrees C / 111 degrees F
      CB 0 CPU Core-6 Temp           OK         44 degrees C / 111 degrees F
      CB 0 CPU Core-7 Temp           OK         43 degrees C / 109 degrees F
Power PEM 0                          Check
      PEM 1                          OK         36 degrees C / 96 degrees F
Fans  Fan Tray 0 Fan 0               OK         Spinning at normal speed
      Fan Tray 0 Fan 1               OK         Spinning at normal speed
```

```
      Fan Tray 1 Fan 0               OK         Spinning at normal speed
      Fan Tray 1 Fan 1               OK         Spinning at normal speed
      Fan Tray 2 Fan 0               Absent
      Fan Tray 2 Fan 1               Absent
```

show chassis environment interconnect-device (QFabric System)

```
user@switch> show chassis environment interconnect-device
IC-A0004
Class Item                           Status     Measurement
      CB 0
      CB 0 L Intake                  OK         30 degrees C / 86 degrees F
      CB 0 R Intake                  OK         31 degrees C / 87 degrees F
      CB 0 L Exhaust                 OK         32 degrees C / 89 degrees F
      CB 0 R Exhaust                 OK         33 degrees C / 91 degrees F
      Routing Engine 0 CPU temp      OK         51 degrees C / 123 degrees F
      CB 1
      CB 1 L Intake                  OK         27 degrees C / 80 degrees F
      CB 1 R Intake                  OK         29 degrees C / 84 degrees F
      CB 1 L Exhaust                 OK         31 degrees C / 87 degrees F
      CB 1 R Exhaust                 OK         32 degrees C / 89 degrees F
      Routing Engine 1 CPU temp      OK         40 degrees C / 104 degrees F
      FC 0 FPC 0
      FPC 0 L Intake                 OK         25 degrees C / 77 degrees F
      FPC 0 R Intake                 OK         28 degrees C / 82 degrees F
      FPC 0 L Exhaust                OK         28 degrees C / 82 degrees F
      FPC 0 R Exhaust                OK         29 degrees C / 84 degrees F
      FC 7 FPC 7
      FPC 7 L Intake                 OK         25 degrees C / 77 degrees F
      FPC 7 R Intake                 OK         26 degrees C / 78 degrees F
      FPC 7 L Exhaust                OK         28 degrees C / 82 degrees F
      FPC 7 R Exhaust                OK         29 degrees C / 84 degrees F
      RC 0 FPC 8
      FPC 8 L Intake                 OK         25 degrees C / 77 degrees F
      FPC 8 R Intake                 OK         26 degrees C / 78 degrees F
      FPC 8 L Exhaust                OK         32 degrees C / 89 degrees F
      FPC 8 R Exhaust                OK         30 degrees C / 86 degrees F
      RC 7 FPC 15
      FPC 15 L Intake                OK         24 degrees C / 75 degrees F
      FPC 15 R Intake                OK         25 degrees C / 77 degrees F
      FPC 15 L Exhaust               OK         33 degrees C / 91 degrees F
```

```
      FPC 15 R Exhaust              OK         31 degrees C / 87 degrees F
Fans  TFT 0 Fan 0                   OK         Spinning at normal speed
Fans  TFT 0 Fan 1                   OK         Spinning at normal speed
Fans  TFT 0 Fan 2                   OK         Spinning at normal speed
Fans  TFT 0 Fan 3                   OK         Spinning at normal speed
Fans  TFT 0 Fan 4                   OK         Spinning at normal speed
Fans  TFT 0 Fan 5                   OK         Spinning at normal speed
Fans  BFT 1 Fan 0                   OK         Spinning at normal speed
Fans  BFT 1 Fan 1                   OK         Spinning at normal speed
Fans  BFT 1 Fan 2                   OK         Spinning at normal speed
Fans  BFT 1 Fan 3                   Check
Fans  BFT 1 Fan 4                   OK         Spinning at normal speed
Fans  BFT 1 Fan 5                   OK         Spinning at normal speed
Fans  SFT 0 Fan 0 Rotor 0           OK         Spinning at normal speed
Fans  SFT 0 Fan 0 Rotor 1           OK         Spinning at normal speed
Fans  SFT 0 Fan 1 Rotor 0           OK         Spinning at normal speed
Fans  SFT 0 Fan 1 Rotor 1           OK         Spinning at normal speed
Fans  SFT 0 Fan 2 Rotor 0           OK         Spinning at normal speed
Fans  SFT 0 Fan 2 Rotor 1           OK         Spinning at normal speed
Fans  SFT 0 Fan 3 Rotor 0           OK         Spinning at normal speed
Fans  SFT 0 Fan 3 Rotor 1           OK         Spinning at normal speed
Fans  SFT 1 Fan 0 Rotor 0           OK         Spinning at normal speed
Fans  SFT 1 Fan 0 Rotor 1           OK         Spinning at normal speed
Fans  SFT 1 Fan 1 Rotor 0           OK         Spinning at normal speed
Fans  SFT 1 Fan 1 Rotor 1           OK         Spinning at normal speed
Fans  SFT 1 Fan 2 Rotor 0           OK         Spinning at normal speed
Fans  SFT 1 Fan 2 Rotor 1           OK         Spinning at normal speed
Fans  SFT 1 Fan 3 Rotor 0           OK         Spinning at normal speed
Fans  SFT 1 Fan 3 Rotor 1           OK         Spinning at normal speed
Fans  SFT 2 Fan 0 Rotor 0           OK         Spinning at normal speed
Fans  SFT 2 Fan 0 Rotor 1           OK         Spinning at normal speed
Fans  SFT 2 Fan 1 Rotor 0           OK         Spinning at normal speed
Fans  SFT 2 Fan 1 Rotor 1           OK         Spinning at normal speed
Fans  SFT 2 Fan 2 Rotor 0           OK         Spinning at normal speed
Fans  SFT 2 Fan 2 Rotor 1           OK         Spinning at normal speed
Fans  SFT 2 Fan 3 Rotor 0           OK         Spinning at normal speed
Fans  SFT 2 Fan 3 Rotor 1           OK         Spinning at normal speed
Fans  SFT 3 Fan 0 Rotor 0           OK         Spinning at normal speed
Fans  SFT 3 Fan 0 Rotor 1           OK         Spinning at normal speed
Fans  SFT 3 Fan 1 Rotor 0           OK         Spinning at normal speed
Fans  SFT 3 Fan 1 Rotor 1           OK         Spinning at normal speed
Fans  SFT 3 Fan 2 Rotor 0           OK         Spinning at normal speed
Fans  SFT 3 Fan 2 Rotor 1           OK         Spinning at normal speed
```

```
Fans  SFT 3 Fan 3 Rotor 0          OK         Spinning at normal speed
Fans  SFT 3 Fan 3 Rotor 1          OK         Spinning at normal speed
Fans  SFT 4 Fan 0 Rotor 0          OK         Spinning at normal speed
Fans  SFT 4 Fan 0 Rotor 1          OK         Spinning at normal speed
Fans  SFT 4 Fan 1 Rotor 0          OK         Spinning at normal speed
Fans  SFT 4 Fan 1 Rotor 1          OK         Spinning at normal speed
Fans  SFT 4 Fan 2 Rotor 0          OK         Spinning at normal speed
Fans  SFT 4 Fan 2 Rotor 1          OK         Spinning at normal speed
Fans  SFT 4 Fan 3 Rotor 0          OK         Spinning at normal speed
Fans  SFT 4 Fan 3 Rotor 1          OK         Spinning at normal speed
Fans  SFT 5 Fan 0 Rotor 0          OK         Spinning at normal speed
Fans  SFT 5 Fan 0 Rotor 1          OK         Spinning at normal speed
Fans  SFT 5 Fan 1 Rotor 0          OK         Spinning at normal speed
Fans  SFT 5 Fan 1 Rotor 1          OK         Spinning at normal speed
Fans  SFT 5 Fan 2 Rotor 0          OK         Spinning at normal speed
Fans  SFT 5 Fan 2 Rotor 1          OK         Spinning at normal speed
Fans  SFT 5 Fan 3 Rotor 0          OK         Spinning at normal speed
Fans  SFT 5 Fan 3 Rotor 1          OK         Spinning at normal speed
Fans  SFT 6 Fan 0 Rotor 0          OK         Spinning at normal speed
Fans  SFT 6 Fan 0 Rotor 1          OK         Spinning at normal speed
Fans  SFT 6 Fan 1 Rotor 0          OK         Spinning at normal speed
Fans  SFT 6 Fan 1 Rotor 1          OK         Spinning at normal speed
Fans  SFT 6 Fan 2 Rotor 0          OK         Spinning at normal speed
Fans  SFT 6 Fan 2 Rotor 1          OK         Spinning at normal speed
Fans  SFT 6 Fan 3 Rotor 0          OK         Spinning at normal speed
Fans  SFT 6 Fan 3 Rotor 1          OK         Spinning at normal speed
Fans  SFT 7 Fan 0 Rotor 0          OK         Spinning at normal speed
Fans  SFT 7 Fan 0 Rotor 1          OK         Spinning at normal speed
Fans  SFT 7 Fan 1 Rotor 0          OK         Spinning at normal speed
Fans  SFT 7 Fan 1 Rotor 1          OK         Spinning at normal speed
Fans  SFT 7 Fan 2 Rotor 0          OK         Spinning at normal speed
Fans  SFT 7 Fan 2 Rotor 1          OK         Spinning at normal speed
Fans  SFT 7 Fan 3 Rotor 0          OK         Spinning at normal speed
Fans  SFT 7 Fan 3 Rotor 1          OK         Spinning at normal speed
Power PEM 0                        OK         30 degrees C / 86 degrees F
Power PEM 1                        OK         30 degrees C / 86 degrees F
Power PEM 2                        OK         30 degrees C / 86 degrees F
Power PEM 3                        Absent
Power PEM 4                        Absent
Power PEM 5                        Absent
```

show chassis environment node-device (QFabric System)

```
user@switch> show chassis environment node-device
node1
Class Item                          Status     Measurement
Power node1 Power Supply 0           Absent
      node1 Power Supply 1           Absent
Fans  node1 Fan Tray 0               Testing
      node1 Fan Tray 1               Testing
      node1 Fan Tray 2               Testing
```

show chassis environment pem node-device (QFabric System)

```
user@switch> show chassis environment pem node-device
node1
FPC 0 PEM 0 status:
  State                      Check
  Airflow                    Front to Back
  Temperature                OK
  AC Input:                  OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12        10            120      18
FPC 0 PEM 1 status:
  State                      Online
  Airflow                    Back to Front
  Temperature                OK
  AC Input:                  OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          11        10           110      17
```

show chassis environment (PTX5000 Packet Transport Router)

```
user@host> show chassis environment
Class Item                          Status     Measurement
Temp  PDU 0                         OK
      PDU 0 PSM 0                   OK         36 degrees C / 96 degrees F
      PDU 0 PSM 1                   OK         38 degrees C / 100 degrees F
      PDU 0 PSM 2                   OK         38 degrees C / 100 degrees F
      PDU 0 PSM 3                   OK         37 degrees C / 98 degrees F
```

```
    PDU 1                          Absent
    CCG 0                          OK         44 degrees C / 111 degrees F
    CCG 1                          OK         44 degrees C / 111 degrees F
    Routing Engine 0               OK         62 degrees C / 143 degrees F
    Routing Engine 0 CPU           OK         75 degrees C / 167 degrees F
    Routing Engine 1               OK         51 degrees C / 123 degrees F
    Routing Engine 1 CPU           OK         64 degrees C / 147 degrees F
    CB 0 Intake                    OK         38 degrees C / 100 degrees F
    CB 0 Exhaust A                 OK         46 degrees C / 114 degrees F
    CB 0 Exhaust B                 OK         42 degrees C / 107 degrees F
    CB 1 Intake                    OK         35 degrees C / 95 degrees F
    CB 1 Exhaust A                 OK         39 degrees C / 102 degrees F
    CB 1 Exhaust B                 OK         36 degrees C / 96 degrees F
    SIB 0 Exhaust                  OK         47 degrees C / 116 degrees F
    SIB 0 Junction                 OK         45 degrees C / 113 degrees F
    SIB 1 Exhaust                  OK         44 degrees C / 111 degrees F
    SIB 1 Junction                 OK         43 degrees C / 109 degrees F
    SIB 2 Exhaust                  OK         47 degrees C / 116 degrees F
    SIB 2 Junction                 OK         42 degrees C / 107 degrees F
    SIB 3 Exhaust                  OK         43 degrees C / 109 degrees F
    SIB 3 Junction                 OK         43 degrees C / 109 degrees F
    SIB 4 Exhaust                  OK         47 degrees C / 116 degrees F
    SIB 4 Junction                 OK         42 degrees C / 107 degrees F
    SIB 5 Exhaust                  OK         42 degrees C / 107 degrees F
    SIB 5 Junction                 OK         40 degrees C / 104 degrees F
    SIB 6 Exhaust                  OK         46 degrees C / 114 degrees F
    SIB 6 Junction                 OK         42 degrees C / 107 degrees F
    SIB 7 Exhaust                  OK         43 degrees C / 109 degrees F
    SIB 7 Junction                 OK         39 degrees C / 102 degrees F
    SIB 8 Exhaust                  OK         44 degrees C / 111 degrees F
    SIB 8 Junction                 OK         41 degrees C / 105 degrees F
    FPC 0 PMB                      OK         35 degrees C / 95 degrees F
    FPC 0 Intake                   OK         33 degrees C / 91 degrees F
    FPC 0 Exhaust A                OK         51 degrees C / 123 degrees F
    FPC 0 Exhaust B                OK         43 degrees C / 109 degrees F
    FPC 0 TL0                      OK         48 degrees C / 118 degrees F
    FPC 0 TQ0                      OK         53 degrees C / 127 degrees F
    FPC 0 TL1                      OK         56 degrees C / 132 degrees F
    FPC 0 TQ1                      OK         58 degrees C / 136 degrees F
    FPC 0 TL2                      OK         55 degrees C / 131 degrees F
    FPC 0 TQ2                      OK         56 degrees C / 132 degrees F
    FPC 0 TL3                      OK         59 degrees C / 138 degrees F
    FPC 0 TQ3                      OK         59 degrees C / 138 degrees F
```

```
    FPC 2 PMB                    OK         35 degrees C / 95 degrees F
    FPC 2 Intake                 OK         34 degrees C / 93 degrees F
    FPC 2 Exhaust A              OK         51 degrees C / 123 degrees F
    FPC 2 Exhaust B              OK         52 degrees C / 125 degrees F
    FPC 2 TL0                    OK         53 degrees C / 127 degrees F
    FPC 2 TQ0                    OK         53 degrees C / 127 degrees F
    FPC 2 TL1                    OK         57 degrees C / 134 degrees F
    FPC 2 TQ1                    OK         58 degrees C / 136 degrees F
    FPC 2 TL2                    OK         54 degrees C / 129 degrees F
    FPC 2 TQ2                    OK         59 degrees C / 138 degrees F
    FPC 2 TL3                    OK         60 degrees C / 140 degrees F
    FPC 2 TQ3                    OK         64 degrees C / 147 degrees F
    PIC 2/0 Ambient              OK         49 degrees C / 120 degrees F
    FPC 3 PMB                    OK         34 degrees C / 93 degrees F
    FPC 3 Intake                 OK         35 degrees C / 95 degrees F
    FPC 3 Exhaust A              OK         54 degrees C / 129 degrees F
    FPC 3 Exhaust B              OK         49 degrees C / 120 degrees F
    FPC 3 TL0                    OK         49 degrees C / 120 degrees F
    FPC 3 TQ0                    OK         55 degrees C / 131 degrees F
    FPC 3 TL1                    OK         56 degrees C / 132 degrees F
    FPC 3 TQ1                    OK         58 degrees C / 136 degrees F
    FPC 3 TL2                    OK         56 degrees C / 132 degrees F
    FPC 3 TQ2                    OK         59 degrees C / 138 degrees F
    FPC 3 TL3                    OK         62 degrees C / 143 degrees F
    FPC 3 TQ3                    OK         63 degrees C / 145 degrees F
    PIC 3/1                      Absent
    FPC 5 PMB                    OK         35 degrees C / 95 degrees F
    FPC 5 Intake                 OK         34 degrees C / 93 degrees F
    FPC 5 Exhaust A              OK         51 degrees C / 123 degrees F
    FPC 5 Exhaust B              OK         53 degrees C / 127 degrees F
    FPC 5 TL0                    OK         54 degrees C / 129 degrees F
    FPC 5 TQ0                    OK         52 degrees C / 125 degrees F
    FPC 5 TL1                    OK         61 degrees C / 141 degrees F
    FPC 5 TQ1                    OK         60 degrees C / 140 degrees F
    FPC 5 TL2                    OK         55 degrees C / 131 degrees F
    FPC 5 TQ2                    OK         55 degrees C / 131 degrees F
    FPC 5 TL3                    OK         59 degrees C / 138 degrees F
    FPC 5 TQ3                    OK         58 degrees C / 136 degrees F
    PIC 5/0 Ambient              OK         51 degrees C / 123 degrees F
    PIC 5/1 Ambient              OK         34 degrees C / 93 degrees F
    PIC 5/1 cfp-5/1/0            OK         34 degrees C / 93 degrees F
    PIC 5/1 cfp-5/1/1            OK         36 degrees C / 96 degrees F
    FPC 6 PMB                    OK         36 degrees C / 96 degrees F
```

```
      FPC 6 Intake                OK        33 degrees C / 91 degrees F
      FPC 6 Exhaust A             OK        51 degrees C / 123 degrees F
      FPC 6 Exhaust B             OK        39 degrees C / 102 degrees F
      FPC 6 TL0                   OK        44 degrees C / 111 degrees F
      FPC 6 TQ0                   OK        54 degrees C / 129 degrees F
      FPC 6 TL1                   OK        59 degrees C / 138 degrees F
      FPC 6 TQ1                   OK        58 degrees C / 136 degrees F
      FPC 6 TL2                   OK        60 degrees C / 140 degrees F
      FPC 6 TQ2                   OK        57 degrees C / 134 degrees F
      FPC 6 TL3                   OK        65 degrees C / 149 degrees F
      FPC 6 TQ3                   OK        60 degrees C / 140 degrees F
      FPC 7 PMB                   OK        35 degrees C / 95 degrees F
      FPC 7 Intake                OK        33 degrees C / 91 degrees F
      FPC 7 Exhaust A             OK        53 degrees C / 127 degrees F
      FPC 7 Exhaust B             OK        40 degrees C / 104 degrees F
      FPC 7 TL0                   OK        46 degrees C / 114 degrees F
      FPC 7 TQ0                   OK        58 degrees C / 136 degrees F
      FPC 7 TL1                   OK        53 degrees C / 127 degrees F
      FPC 7 TQ1                   OK        59 degrees C / 138 degrees F
      FPC 7 TL2                   OK        56 degrees C / 132 degrees F
      FPC 7 TQ2                   OK        61 degrees C / 141 degrees F
      FPC 7 TL3                   OK        63 degrees C / 145 degrees F
      FPC 7 TQ3                   OK        63 degrees C / 145 degrees F
      FPM I2CS                    OK        37 degrees C / 98 degrees F
Fans  Fan Tray 0 Fan 1            OK        3042 RPM
      Fan Tray 0 Fan 2            OK        3042 RPM
      Fan Tray 0 Fan 3            OK        3000 RPM
      Fan Tray 0 Fan 4            OK        3042 RPM
      Fan Tray 0 Fan 5            OK        3000 RPM
      Fan Tray 0 Fan 6            OK        3042 RPM
      Fan Tray 0 Fan 7            OK        3085 RPM
      Fan Tray 0 Fan 8            OK        3042 RPM
      Fan Tray 0 Fan 9            OK        3042 RPM
      Fan Tray 0 Fan 10           OK        3085 RPM
      Fan Tray 0 Fan 11           OK        3085 RPM
      Fan Tray 0 Fan 12           OK        3128 RPM
      Fan Tray 0 Fan 13           OK        3128 RPM
      Fan Tray 0 Fan 14           OK        3042 RPM
      Fan Tray 1 Fan 1            OK        2299 RPM
      Fan Tray 1 Fan 2            OK        2399 RPM
      Fan Tray 1 Fan 3            OK        2299 RPM
      Fan Tray 1 Fan 4            OK        2266 RPM
      Fan Tray 1 Fan 5            OK        2266 RPM
```

```
      Fan Tray 1 Fan 6              OK         2366 RPM
      Fan Tray 2 Fan 1              OK         2199 RPM
      Fan Tray 2 Fan 2              OK         2133 RPM
      Fan Tray 2 Fan 3              OK         2366 RPM
      Fan Tray 2 Fan 4              OK         2233 RPM
      Fan Tray 2 Fan 5              OK         2399 RPM
      Fan Tray 2 Fan 6              OK         2233 RPM
Misc  SPMB 0 Intake                 OK         50 degrees C / 122 degrees F
      SPMB 1 Intake                 OK         40 degrees C / 104 degrees F
```

show chassis environment (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis environment
Class Item                          Status     Measurement
Temp  PDU 0                         OK
      PDU 0 PSM 0                   OK         41 degrees C / 105 degrees F
      PDU 0 PSM 1                   Absent
      PDU 0 PSM 2                   OK         43 degrees C / 109 degrees F
      PDU 0 PSM 3                   Absent
      PDU 0 PSM 4                   OK         44 degrees C / 111 degrees F
      PDU 0 PSM 5                   Absent
      PDU 0 PSM 6                   OK         45 degrees C / 113 degrees F
      PDU 0 PSM 7                   Absent
      PDU 1                         OK
      PDU 1 PSM 0                   Absent
      PDU 1 PSM 1                   OK         45 degrees C / 113 degrees F
      PDU 1 PSM 2                   Absent
      PDU 1 PSM 3                   OK         43 degrees C / 109 degrees F
      PDU 1 PSM 4                   Absent
      PDU 1 PSM 5                   OK         46 degrees C / 114 degrees F
      PDU 1 PSM 6                   Absent
      PDU 1 PSM 7                   OK         46 degrees C / 114 degrees F
      CCG 0                         OK         27 degrees C / 80 degrees F
      CCG 1                         OK         29 degrees C / 84 degrees F
...
```

show chassis environment (PTX1000 Packet Transport Router)

```
user@host> show chassis environment
Class Item                          Status     Measurement
```

```
Power FPC 0 Power Supply 0          Absent
      FPC 0 Power Supply 1          Absent
      FPC 0 Power Supply 2          OK
      FPC 0 Power Supply 3          OK
Temp  FPC 0 Intake Temp Sensor      OK         25 degrees C / 77 degrees F
      FPC 0 Exhaust Temp Sensor     OK         35 degrees C / 95 degrees F
      FPC 0 Mezz Temp Sensor 0      OK         25 degrees C / 77 degrees F
      FPC 0 Mezz Temp Sensor 1      OK         34 degrees C / 93 degrees F
      FPC 0 PE2 Temp Sensor         OK         34 degrees C / 93 degrees F
      FPC 0 PE1 Temp Sensor         OK         32 degrees C / 89 degrees F
      FPC 0 PF0 Temp Sensor         OK         40 degrees C / 104 degrees F
      FPC 0 PE0 Temp Sensor         OK         33 degrees C / 91 degrees F
      FPC 0 PE5 Temp Sensor         OK         34 degrees C / 93 degrees F
      FPC 0 PE4 Temp Sensor         OK         34 degrees C / 93 degrees F
      FPC 0 PF1 Temp Sensor         OK         41 degrees C / 105 degrees F
      FPC 0 PE3 Temp Sensor         OK         36 degrees C / 96 degrees F
      FPC 0 CPU Die Temp Sensor     OK         40 degrees C / 104 degrees F
      FPC 0 OCXO Temp Sensor        OK         37 degrees C / 98 degrees F
Fans  FPC 0 Fan Tray 0              OK         Spinning at normal speed
      FPC 0 Fan Tray 1              OK         Spinning at normal speed
      FPC 0 Fan Tray 2              OK         Spinning at normal speed
```

show chassis environment (PTX10008 Router)

```
user@host> show chassis environment
Class Item                          Status     Measurement
      Routing Engine 0 CPU          OK         40 degrees C / 104 degrees F
      Routing Engine 1 CPU          OK         40 degrees C / 104 degrees F
Temp  CB 0 Intake Temp Sensor       OK         29 degrees C / 84 degrees F
      CB 0 Exhaust Temp Sensor      OK         33 degrees C / 91 degrees F
      CB 1 Intake Temp Sensor       OK         28 degrees C / 82 degrees F
      CB 1 Exhaust Temp Sensor      OK         32 degrees C / 89 degrees F
      FPC 0 Intake-A Temp Sensor    OK         38 degrees C / 100 degrees F
      FPC 0 Intake-B Temp Sensor    OK         34 degrees C / 93 degrees F
      FPC 0 Exhaust-A Temp Sensor   OK         36 degrees C / 96 degrees F
      FPC 0 Exhaust-B Temp Sensor   OK         37 degrees C / 98 degrees F
      FPC 0 Exhaust-C Temp Sensor   OK         40 degrees C / 104 degrees F
      FPC 0 PE0 Temp Sensor         OK         40 degrees C / 104 degrees F
      FPC 0 PE1 Temp Sensor         OK         42 degrees C / 107 degrees F
      FPC 0 PE2 Temp Sensor         OK         44 degrees C / 111 degrees F
      FPC 0 LCPU Temp Sensor        OK         41 degrees C / 105 degrees F
```

```
    FPC 1 Intake-A Temp Sensor     OK         37 degrees C / 98 degrees F
    FPC 1 Intake-B Temp Sensor     OK         33 degrees C / 91 degrees F
    FPC 1 Exhaust-A Temp Sensor    OK         37 degrees C / 98 degrees F
    FPC 1 Exhaust-B Temp Sensor    OK         38 degrees C / 100 degrees F
    FPC 1 Exhaust-C Temp Sensor    OK         40 degrees C / 104 degrees F
    FPC 1 PE0 Temp Sensor          OK         41 degrees C / 105 degrees F
    FPC 1 PE1 Temp Sensor          OK         41 degrees C / 105 degrees F
    FPC 1 PE2 Temp Sensor          OK         45 degrees C / 113 degrees F
    FPC 1 LCPU Temp Sensor         OK         40 degrees C / 104 degrees F
    FPC 2 Intake-A Temp Sensor     OK         44 degrees C / 111 degrees F
    FPC 2 Intake-B Temp Sensor     OK         30 degrees C / 86 degrees F
    FPC 2 Exhaust-A Temp Sensor    OK         52 degrees C / 125 degrees F
    FPC 2 Exhaust-B Temp Sensor    OK         54 degrees C / 129 degrees F
    FPC 2 Exhaust-C Temp Sensor    OK         52 degrees C / 125 degrees F
    FPC 2 PE0 Temp Sensor          OK         49 degrees C / 120 degrees F
    FPC 2 PE1 Temp Sensor          OK         59 degrees C / 138 degrees F
    FPC 2 PE2 Temp Sensor          OK         49 degrees C / 120 degrees F
    FPC 2 PE3 Temp Sensor          OK         60 degrees C / 140 degrees F
    FPC 2 PE4 Temp Sensor          OK         49 degrees C / 120 degrees F
    FPC 2 PE5 Temp Sensor          OK         63 degrees C / 145 degrees F
    FPC 2 LCPU Temp Sensor         OK         47 degrees C / 116 degrees F
    FPC 3 Intake-A Temp Sensor     OK         42 degrees C / 107 degrees F
    FPC 3 Intake-B Temp Sensor     OK         30 degrees C / 86 degrees F
    FPC 3 Exhaust-A Temp Sensor    OK         46 degrees C / 114 degrees F
    FPC 3 Exhaust-B Temp Sensor    OK         48 degrees C / 118 degrees F
    FPC 3 Exhaust-C Temp Sensor    OK         47 degrees C / 116 degrees F
    FPC 3 PE0 Temp Sensor          OK         47 degrees C / 116 degrees F
    FPC 3 PE1 Temp Sensor          OK         53 degrees C / 127 degrees F
    FPC 3 PE2 Temp Sensor          OK         46 degrees C / 114 degrees F
    FPC 3 PE3 Temp Sensor          OK         53 degrees C / 127 degrees F
    FPC 3 PE4 Temp Sensor          OK         48 degrees C / 118 degrees F
    FPC 3 PE5 Temp Sensor          OK         57 degrees C / 134 degrees F
    FPC 3 LCPU Temp Sensor         OK         47 degrees C / 116 degrees F
    FPC 5 Intake-A Temp Sensor     Failed
    FPC 5 Intake-B Temp Sensor     Failed
    FPC 5 Exhaust-A Temp Sensor    OK         40 degrees C / 104 degrees F
    FPC 5 Exhaust-B Temp Sensor    OK         40 degrees C / 104 degrees F
    FPC 5 Exhaust-C Temp Sensor    OK         41 degrees C / 105 degrees F
    FPC 5 PE0 Temp Sensor          OK         46 degrees C / 114 degrees F
    FPC 5 PE1 Temp Sensor          OK         48 degrees C / 118 degrees F
    FPC 5 PE2 Temp Sensor          OK         51 degrees C / 123 degrees F
    FPC 5 LCPU Temp Sensor         Failed
    FPC 6 Intake-A Temp Sensor     OK         40 degrees C / 104 degrees F
```

```
FPC 6 Intake-B Temp Sensor     OK         36 degrees C / 96 degrees F
FPC 6 Exhaust-A Temp Sensor    OK         39 degrees C / 102 degrees F
FPC 6 Exhaust-B Temp Sensor    OK         39 degrees C / 102 degrees F
FPC 6 Exhaust-C Temp Sensor    OK         39 degrees C / 102 degrees F
FPC 6 PE0 Temp Sensor          OK         44 degrees C / 111 degrees F
FPC 6 PE1 Temp Sensor          OK         45 degrees C / 113 degrees F
FPC 6 PE2 Temp Sensor          OK         50 degrees C / 122 degrees F
FPC 6 LCPU Temp Sensor         OK         40 degrees C / 104 degrees F
SIB 0 Intake-A Temp Sensor     OK         37 degrees C / 98 degrees F
SIB 0 Intake-B Temp Sensor     OK         30 degrees C / 86 degrees F
SIB 0 Exhaust-A Temp Sensor    OK         33 degrees C / 91 degrees F
SIB 0 Exhaust-B Temp Sensor    OK         38 degrees C / 100 degrees F
SIB 0 PF0 Temp Sensor          OK         46 degrees C / 114 degrees F
SIB 0 PF1 Temp Sensor          OK         39 degrees C / 102 degrees F
SIB 1 Intake-A Temp Sensor     OK         43 degrees C / 109 degrees F
SIB 1 Intake-B Temp Sensor     OK         34 degrees C / 93 degrees F
SIB 1 Exhaust-A Temp Sensor    OK         36 degrees C / 96 degrees F
SIB 1 Exhaust-B Temp Sensor    OK         44 degrees C / 111 degrees F
SIB 1 PF0 Temp Sensor          OK         54 degrees C / 129 degrees F
SIB 1 PF1 Temp Sensor          OK         41 degrees C / 105 degrees F
SIB 2 Intake-A Temp Sensor     OK         46 degrees C / 114 degrees F
SIB 2 Intake-B Temp Sensor     OK         35 degrees C / 95 degrees F
SIB 2 Exhaust-A Temp Sensor    OK         37 degrees C / 98 degrees F
SIB 2 Exhaust-B Temp Sensor    OK         47 degrees C / 116 degrees F
SIB 2 PF0 Temp Sensor          OK         55 degrees C / 131 degrees F
SIB 2 PF1 Temp Sensor          OK         42 degrees C / 107 degrees F
SIB 3 Intake-A Temp Sensor     OK         45 degrees C / 113 degrees F
SIB 3 Intake-B Temp Sensor     OK         35 degrees C / 95 degrees F
SIB 3 Exhaust-A Temp Sensor    OK         36 degrees C / 96 degrees F
SIB 3 Exhaust-B Temp Sensor    OK         45 degrees C / 113 degrees F
SIB 3 PF0 Temp Sensor          OK         54 degrees C / 129 degrees F
SIB 3 PF1 Temp Sensor          OK         42 degrees C / 107 degrees F
SIB 4 Intake-A Temp Sensor     OK         46 degrees C / 114 degrees F
SIB 4 Intake-B Temp Sensor     OK         34 degrees C / 93 degrees F
SIB 4 Exhaust-A Temp Sensor    OK         36 degrees C / 96 degrees F
SIB 4 Exhaust-B Temp Sensor    OK         46 degrees C / 114 degrees F
SIB 4 PF0 Temp Sensor          OK         54 degrees C / 129 degrees F
SIB 4 PF1 Temp Sensor          OK         41 degrees C / 105 degrees F
SIB 5 Intake-A Temp Sensor     OK         38 degrees C / 100 degrees F
SIB 5 Intake-B Temp Sensor     OK         31 degrees C / 87 degrees F
SIB 5 Exhaust-A Temp Sensor    OK         34 degrees C / 93 degrees F
SIB 5 Exhaust-B Temp Sensor    OK         39 degrees C / 102 degrees F
SIB 5 PF0 Temp Sensor          OK         44 degrees C / 111 degrees F
```

```
      SIB 5 PF1 Temp Sensor          OK         42 degrees C / 107 degrees F
Power Power Supply 0                 OK
      Power Supply 1                 OK
      Power Supply 2                 OK
      Power Supply 3                 OK
      Power Supply 4                 Check
      Power Supply 5                 OK
Fans  Fan Tray 0 Fan 0               OK         Spinning at normal speed
      Fan Tray 0 Fan 1               OK         Spinning at normal speed
      Fan Tray 0 Fan 2               OK         Spinning at normal speed
      Fan Tray 0 Fan 3               OK         Spinning at normal speed
      Fan Tray 0 Fan 4               Failed
      Fan Tray 0 Fan 5               Failed
      Fan Tray 0 Fan 6               OK         Spinning at normal speed
      Fan Tray 0 Fan 7               OK         Spinning at normal speed
      Fan Tray 0 Fan 8               OK         Spinning at normal speed
      Fan Tray 0 Fan 9               OK         Spinning at normal speed
      Fan Tray 0 Fan 10              OK         Spinning at normal speed
      Fan Tray 1 Fan 0               OK         Spinning at normal speed
      Fan Tray 1 Fan 1               OK         Spinning at normal speed
      Fan Tray 1 Fan 2               OK         Spinning at normal speed
      Fan Tray 1 Fan 3               OK         Spinning at normal speed
      Fan Tray 1 Fan 4               OK         Spinning at normal speed
      Fan Tray 1 Fan 5               OK         Spinning at normal speed
      Fan Tray 1 Fan 6               OK         Spinning at normal speed
      Fan Tray 1 Fan 7               OK         Spinning at normal speed
      Fan Tray 1 Fan 8               OK         Spinning at normal speed
      Fan Tray 1 Fan 9               OK         Spinning at normal speed
      Fan Tray 1 Fan 10              OK         Spinning at normal speed
```

show chassis environment (PTX10016 Router)

```
user@host> show chassis environment
Class Item                           Status     Measurement
      Routing Engine 0 CPU           OK         34 degrees C / 93 degrees F
      Routing Engine 1 CPU           OK         34 degrees C / 93 degrees F
Temp  CB 0 Intake Temp Sensor        OK         20 degrees C / 68 degrees F
      CB 0 Exhaust Temp Sensor       OK         24 degrees C / 75 degrees F
      CB 1 Intake Temp Sensor        OK         20 degrees C / 68 degrees F
      CB 1 Exhaust Temp Sensor       OK         23 degrees C / 73 degrees F
      FPC 1 Intake-A Temp Sensor     OK         37 degrees C / 98 degrees F
```

```
    FPC 1 Intake-B Temp Sensor     OK         32 degrees C / 89 degrees F
    FPC 1 Exhaust-A Temp Sensor    OK         37 degrees C / 98 degrees F
    FPC 1 Exhaust-B Temp Sensor    OK         36 degrees C / 96 degrees F
    FPC 1 Exhaust-C Temp Sensor    OK         36 degrees C / 96 degrees F
    FPC 1 PE0 Temp Sensor          OK         45 degrees C / 113 degrees F
    FPC 1 PE1 Temp Sensor          OK         46 degrees C / 114 degrees F
    FPC 1 PE2 Temp Sensor          OK         54 degrees C / 129 degrees F
    FPC 1 LCPU Temp Sensor         OK         35 degrees C / 95 degrees F
    FPC 3 Intake-A Temp Sensor     OK         35 degrees C / 95 degrees F
    FPC 3 Intake-B Temp Sensor     OK         31 degrees C / 87 degrees F
    FPC 3 Exhaust-A Temp Sensor    OK         36 degrees C / 96 degrees F
    FPC 3 Exhaust-B Temp Sensor    OK         35 degrees C / 95 degrees F
    FPC 3 Exhaust-C Temp Sensor    OK         33 degrees C / 91 degrees F
    FPC 3 PE0 Temp Sensor          OK         43 degrees C / 109 degrees F
    FPC 3 PE1 Temp Sensor          OK         45 degrees C / 113 degrees F
    FPC 3 PE2 Temp Sensor          OK         49 degrees C / 120 degrees F
    FPC 3 LCPU Temp Sensor         OK         35 degrees C / 95 degrees F
    FPC 6 Intake-A Temp Sensor     OK         34 degrees C / 93 degrees F
    FPC 6 Intake-B Temp Sensor     OK         31 degrees C / 87 degrees F
    FPC 6 Exhaust-A Temp Sensor    OK         35 degrees C / 95 degrees F
    FPC 6 Exhaust-B Temp Sensor    OK         35 degrees C / 95 degrees F
    FPC 6 Exhaust-C Temp Sensor    OK         35 degrees C / 95 degrees F
    FPC 6 PE0 Temp Sensor          OK         43 degrees C / 109 degrees F
    FPC 6 PE1 Temp Sensor          OK         43 degrees C / 109 degrees F
    FPC 6 PE2 Temp Sensor          OK         47 degrees C / 116 degrees F
    FPC 6 LCPU Temp Sensor         OK         35 degrees C / 95 degrees F
    FPC 8 Intake-A Temp Sensor     OK         34 degrees C / 93 degrees F
    FPC 8 Intake-B Temp Sensor     OK         31 degrees C / 87 degrees F
    FPC 8 Exhaust-A Temp Sensor    OK         37 degrees C / 98 degrees F
    FPC 8 Exhaust-B Temp Sensor    OK         37 degrees C / 98 degrees F
    FPC 8 Exhaust-C Temp Sensor    OK         38 degrees C / 100 degrees F
    FPC 8 PE0 Temp Sensor          OK         42 degrees C / 107 degrees F
    FPC 8 PE1 Temp Sensor          OK         44 degrees C / 111 degrees F
    FPC 8 PE2 Temp Sensor          OK         47 degrees C / 116 degrees F
    FPC 8 LCPU Temp Sensor         OK         33 degrees C / 91 degrees F
    FPC 9 Intake-A Temp Sensor     OK         44 degrees C / 111 degrees F
    FPC 9 Intake-B Temp Sensor     OK         28 degrees C / 82 degrees F
    FPC 9 Exhaust-A Temp Sensor    OK         51 degrees C / 123 degrees F
    FPC 9 Exhaust-B Temp Sensor    OK         52 degrees C / 125 degrees F
    FPC 9 Exhaust-C Temp Sensor    OK         48 degrees C / 118 degrees F
    FPC 9 PE0 Temp Sensor          OK         52 degrees C / 125 degrees F
    FPC 9 PE1 Temp Sensor          OK         66 degrees C / 150 degrees F
    FPC 9 PE2 Temp Sensor          OK         50 degrees C / 122 degrees F
```

```
       FPC 9 PE3 Temp Sensor          OK         65 degrees C / 149 degrees F
       FPC 9 PE4 Temp Sensor          OK         51 degrees C / 123 degrees F
       FPC 9 PE5 Temp Sensor          OK         68 degrees C / 154 degrees F
       FPC 9 LCPU Temp Sensor         OK         46 degrees C / 114 degrees F
Power  Power Supply 0                 OK         22 degrees C / 71 degrees F
       Power Supply 1                 OK         23 degrees C / 73 degrees F
       Power Supply 2                 OK         23 degrees C / 73 degrees F
       Power Supply 3                 OK         21 degrees C / 69 degrees F
       Power Supply 4                 OK         22 degrees C / 71 degrees F
       Power Supply 5                 OK         25 degrees C / 77 degrees F
       Power Supply 6                 OK         21 degrees C / 69 degrees F
       Power Supply 7                 Absent
       Power Supply 8                 Absent
       Power Supply 9                 Absent
Fans   Fan Tray 0 Fan 0               OK         Spinning at normal speed
       Fan Tray 0 Fan 1               OK         Spinning at normal speed
       Fan Tray 0 Fan 2               OK         Spinning at normal speed
       Fan Tray 0 Fan 3               OK         Spinning at normal speed
       Fan Tray 0 Fan 4               OK         Spinning at normal speed
       Fan Tray 0 Fan 5               OK         Spinning at normal speed
       Fan Tray 0 Fan 6               OK         Spinning at normal speed
       Fan Tray 0 Fan 7               OK         Spinning at normal speed
       Fan Tray 0 Fan 8               OK         Spinning at normal speed
       Fan Tray 0 Fan 9               OK         Spinning at normal speed
       Fan Tray 0 Fan 10              OK         Spinning at normal speed
       Fan Tray 0 Fan 11              OK         Spinning at normal speed
       Fan Tray 0 Fan 12              OK         Spinning at normal speed
       Fan Tray 0 Fan 13              OK         Spinning at normal speed
       Fan Tray 0 Fan 14              OK         Spinning at normal speed
       Fan Tray 0 Fan 15              OK         Spinning at normal speed
       Fan Tray 0 Fan 16              OK         Spinning at normal speed
       Fan Tray 0 Fan 17              OK         Spinning at normal speed
       Fan Tray 0 Fan 18              OK         Spinning at normal speed
       Fan Tray 0 Fan 19              OK         Spinning at normal speed
       Fan Tray 0 Fan 20              OK         Spinning at normal speed
       Fan Tray 1 Fan 0               OK         Spinning at normal speed
       Fan Tray 1 Fan 1               OK         Spinning at normal speed
       Fan Tray 1 Fan 2               OK         Spinning at normal speed
       Fan Tray 1 Fan 3               OK         Spinning at normal speed
       Fan Tray 1 Fan 4               OK         Spinning at normal speed
       Fan Tray 1 Fan 5               OK         Spinning at normal speed
       Fan Tray 1 Fan 6               OK         Spinning at normal speed
       Fan Tray 1 Fan 7               OK         Spinning at normal speed
```

```
Fan Tray 1 Fan 8              OK         Spinning at normal speed
Fan Tray 1 Fan 9              OK         Spinning at normal speed
Fan Tray 1 Fan 10             OK         Spinning at normal speed
Fan Tray 1 Fan 11             OK         Spinning at normal speed
Fan Tray 1 Fan 12             OK         Spinning at normal speed
Fan Tray 1 Fan 13             OK         Spinning at normal speed
Fan Tray 1 Fan 14             OK         Spinning at normal speed
Fan Tray 1 Fan 15             OK         Spinning at normal speed
Fan Tray 1 Fan 16             OK         Spinning at normal speed
Fan Tray 1 Fan 17             OK         Spinning at normal speed
Fan Tray 1 Fan 18             OK         Spinning at normal speed
Fan Tray 1 Fan 19             OK         Spinning at normal speed
Fan Tray 1 Fan 20             OK         Spinning at normal speed
SIB 0 Intake-A Temp Sensor    OK         20 degrees C / 68 degrees F
SIB 0 Intake-B Temp Sensor    OK         20 degrees C / 68 degrees F
SIB 0 Intake-C Temp Sensor    OK         16 degrees C / 60 degrees F
SIB 0 Exhaust-A Temp Sensor   OK         28 degrees C / 82 degrees F
SIB 0 Exhaust-B Temp Sensor   OK         28 degrees C / 82 degrees F
SIB 0 Exhaust-C Temp Sensor   OK         23 degrees C / 73 degrees F
SIB 0 PF0 Temp Sensor         OK         30 degrees C / 86 degrees F
SIB 0 PF1 Temp Sensor         OK         30 degrees C / 86 degrees F
SIB 0 PF2 Temp Sensor         OK         31 degrees C / 87 degrees F
SIB 0 PF3 Temp Sensor         OK         32 degrees C / 89 degrees F
SIB 0 PF4 Temp Sensor         OK         27 degrees C / 80 degrees F
SIB 0 PF5 Temp Sensor         OK         26 degrees C / 78 degrees F
SIB 1 Intake-A Temp Sensor    OK         22 degrees C / 71 degrees F
SIB 1 Intake-B Temp Sensor    OK         22 degrees C / 71 degrees F
SIB 1 Intake-C Temp Sensor    OK         16 degrees C / 60 degrees F
SIB 1 Exhaust-A Temp Sensor   OK         29 degrees C / 84 degrees F
SIB 1 Exhaust-B Temp Sensor   OK         31 degrees C / 87 degrees F
SIB 1 Exhaust-C Temp Sensor   OK         23 degrees C / 73 degrees F
SIB 1 PF0 Temp Sensor         OK         32 degrees C / 89 degrees F
SIB 1 PF1 Temp Sensor         OK         31 degrees C / 87 degrees F
SIB 1 PF2 Temp Sensor         OK         33 degrees C / 91 degrees F
SIB 1 PF3 Temp Sensor         OK         38 degrees C / 100 degrees F
SIB 1 PF4 Temp Sensor         OK         28 degrees C / 82 degrees F
SIB 1 PF5 Temp Sensor         OK         26 degrees C / 78 degrees F
SIB 2 Intake-A Temp Sensor    OK         24 degrees C / 75 degrees F
SIB 2 Intake-B Temp Sensor    OK         21 degrees C / 69 degrees F
SIB 2 Intake-C Temp Sensor    OK         16 degrees C / 60 degrees F
SIB 2 Exhaust-A Temp Sensor   OK         28 degrees C / 82 degrees F
SIB 2 Exhaust-B Temp Sensor   OK         32 degrees C / 89 degrees F
SIB 2 Exhaust-C Temp Sensor   OK         23 degrees C / 73 degrees F
```

```
SIB 2 PF0 Temp Sensor          OK         31 degrees C / 87 degrees F
SIB 2 PF1 Temp Sensor          OK         30 degrees C / 86 degrees F
SIB 2 PF2 Temp Sensor          OK         33 degrees C / 91 degrees F
SIB 2 PF3 Temp Sensor          OK         41 degrees C / 105 degrees F
SIB 2 PF4 Temp Sensor          OK         27 degrees C / 80 degrees F
SIB 2 PF5 Temp Sensor          OK         26 degrees C / 78 degrees F
SIB 3 Intake-A Temp Sensor     OK         22 degrees C / 71 degrees F
SIB 3 Intake-B Temp Sensor     OK         23 degrees C / 73 degrees F
SIB 3 Intake-C Temp Sensor     OK         16 degrees C / 60 degrees F
SIB 3 Exhaust-A Temp Sensor    OK         29 degrees C / 84 degrees F
SIB 3 Exhaust-B Temp Sensor    OK         31 degrees C / 87 degrees F
SIB 3 Exhaust-C Temp Sensor    OK         24 degrees C / 75 degrees F
SIB 3 PF0 Temp Sensor          OK         32 degrees C / 89 degrees F
SIB 3 PF1 Temp Sensor          OK         30 degrees C / 86 degrees F
SIB 3 PF2 Temp Sensor          OK         31 degrees C / 87 degrees F
SIB 3 PF3 Temp Sensor          OK         39 degrees C / 102 degrees F
SIB 3 PF4 Temp Sensor          OK         27 degrees C / 80 degrees F
SIB 3 PF5 Temp Sensor          OK         26 degrees C / 78 degrees F
SIB 4 Intake-A Temp Sensor     OK         22 degrees C / 71 degrees F
SIB 4 Intake-B Temp Sensor     OK         25 degrees C / 77 degrees F
SIB 4 Intake-C Temp Sensor     OK         16 degrees C / 60 degrees F
SIB 4 Exhaust-A Temp Sensor    OK         29 degrees C / 84 degrees F
SIB 4 Exhaust-B Temp Sensor    OK         32 degrees C / 89 degrees F
SIB 4 Exhaust-C Temp Sensor    OK         23 degrees C / 73 degrees F
SIB 4 PF0 Temp Sensor          OK         32 degrees C / 89 degrees F
SIB 4 PF1 Temp Sensor          OK         31 degrees C / 87 degrees F
SIB 4 PF2 Temp Sensor          OK         32 degrees C / 89 degrees F
SIB 4 PF3 Temp Sensor          OK         40 degrees C / 104 degrees F
SIB 4 PF4 Temp Sensor          OK         26 degrees C / 78 degrees F
SIB 4 PF5 Temp Sensor          OK         25 degrees C / 77 degrees F
SIB 5 Intake-A Temp Sensor     OK         21 degrees C / 69 degrees F
SIB 5 Intake-B Temp Sensor     OK         20 degrees C / 68 degrees F
SIB 5 Intake-C Temp Sensor     OK         16 degrees C / 60 degrees F
SIB 5 Exhaust-A Temp Sensor    OK         27 degrees C / 80 degrees F
SIB 5 Exhaust-B Temp Sensor    OK         27 degrees C / 80 degrees F
SIB 5 Exhaust-C Temp Sensor    OK         23 degrees C / 73 degrees F
SIB 5 PF0 Temp Sensor          OK         30 degrees C / 86 degrees F
SIB 5 PF1 Temp Sensor          OK         29 degrees C / 84 degrees F
SIB 5 PF2 Temp Sensor          OK         30 degrees C / 86 degrees F
SIB 5 PF3 Temp Sensor          OK         32 degrees C / 89 degrees F
SIB 5 PF4 Temp Sensor          OK         28 degrees C / 82 degrees F
SIB 5 PF5 Temp Sensor          OK         27 degrees C / 80 degrees F
```

show chassis environment (ACX2000 Universal Metro Router)

```
user@host> show chassis environment
Class Item                          Status     Measurement
      PCB Left                      OK         44 degrees C / 111 degrees F
      SFP+ Xcvr                     OK         50 degrees C / 122 degrees F
      FEB                           OK         70 degrees C / 158 degrees F
      PCB Up                        OK         63 degrees C / 145 degrees F
      PCB Mid                       OK         66 degrees C / 150 degrees F
      Telecom Mod                   OK         65 degrees C / 149 degrees F
      Routing Engine                OK         54 degrees C / 129 degrees F
      Heater off
```

show chassis environment (ACX4000 Universal Metro Router)

On the ACX4000 router, the MIC output of the `show chassis environment` command varies depending on the number of temperature channels present in the installed MIC.

```
user@host> show chassis environment

Class Item                          Status     Measurement
Temp  PEM 0                         OK         33 degrees C / 91 degrees F
      PEM 1                         Absent
      PCB Bottom                    OK         30 degrees C / 86 degrees F
      PCB Middle                    OK         34 degrees C / 93 degrees F
      BCM56445                      OK         33 degrees C / 91 degrees F
      SFP+ Xcvr                     OK         32 degrees C / 89 degrees F
      Fan tray inlet                OK         39 degrees C / 102 degrees F
      Exhaust                       OK         30 degrees C / 86 degrees F
      Routing Engine                OK         32 degrees C / 89 degrees F
      Heater off
Pic   PIC 0/0 Channel 0             OK         28 degrees C / 82 degrees F
      PIC 0/0 Channel 1             OK         29 degrees C / 84 degrees F
      PIC 0/0 Channel 2             OK         0 degrees C / 32 degrees F
      PIC 0/0 Channel 3             OK         0 degrees C / 32 degrees F
      PIC 0/0 Channel 4             OK         0 degrees C / 32 degrees F
      PIC 0/0 Channel 5             OK         0 degrees C / 32 degrees F
      PIC 0/0 Channel 6             OK         0 degrees C / 32 degrees F
      PIC 0/0 Channel 7             OK         0 degrees C / 32 degrees F
      PIC 0/0 Channel 8             OK         0 degrees C / 32 degrees F
```

```
      PIC 0/0 Channel 9              OK         0 degrees C / 32 degrees F
      PIC 1/0 Channel 0              OK         33 degrees C / 91 degrees F
      PIC 1/0 Channel 1              OK         31 degrees C / 87 degrees F
      PIC 1/0 Channel 2              OK         30 degrees C / 86 degrees F
      PIC 1/0 Channel 3              OK         0 degrees C / 32 degrees F
      PIC 1/0 Channel 4              OK         0 degrees C / 32 degrees F
      PIC 1/0 Channel 5              OK         0 degrees C / 32 degrees F
      PIC 1/0 Channel 6              OK         0 degrees C / 32 degrees F
      PIC 1/0 Channel 7              OK         0 degrees C / 32 degrees F
      PIC 1/0 Channel 8              OK         0 degrees C / 32 degrees F
      PIC 1/1 Channel 0              OK         31 degrees C / 87 degrees F
      PIC 1/1 Channel 1              OK         29 degrees C / 84 degrees F
      PIC 1/1 Channel 2              OK         28 degrees C / 82 degrees F
      PIC 1/1 Channel 3              OK         0 degrees C / 32 degrees F
      PIC 1/1 Channel 4              OK         0 degrees C / 32 degrees F
      PIC 1/1 Channel 5              OK         0 degrees C / 32 degrees F
      PIC 1/1 Channel 6              OK         0 degrees C / 32 degrees F
      PIC 1/1 Channel 7              OK         0 degrees C / 32 degrees F
      PIC 1/1 Channel 8              OK         0 degrees C / 32 degrees F
Fans  Fan 1                          OK         Spinning at normal speed
      Fan 2                          OK         Spinning at normal speed
```

show chassis environment (ACX5048 Router)

```
user@host> show chassis environment

Class Item                           Status     Measurement
Power FPC 0 Power Supply 0           Absent
      FPC 0 Power Supply 1           OK
Temp  FPC 0 Sensor TopMiddle E       OK         23 degrees C / 73 degrees F
      FPC 0 Sensor TopRight C        OK         18 degrees C / 64 degrees F
      FPC 0 Sensor TopLeft C         OK         21 degrees C / 69 degrees F
      FPC 0 Sensor TopRight E        OK         20 degrees C / 68 degrees F
      FPC 0 Sensor CPURight C        OK         23 degrees C / 73 degrees F
      FPC 0 Sensor CPULeft E         OK         22 degrees C / 71 degrees F
      FPC 0 Sensor CPU Die Temp      OK         39 degrees C / 102 degrees F
Fans  FPC 0 Fan Tray 0               OK         Spinning at normal speed
      FPC 0 Fan Tray 1               OK         Spinning at normal speed
      FPC 0 Fan Tray 2               OK         Spinning at normal speed
```

```
      FPC 0 Fan Tray 3               OK         Spinning at normal speed
      FPC 0 Fan Tray 4               OK         Spinning at normal speed
```

show chassis environment (ACX5096 Router)

```
user@host> show chassis environment

Class Item                           Status     Measurement
Power FPC 0 Power Supply 0           OK
      FPC 0 Power Supply 1           OK
Temp  FPC 0 Sensor TopMiddle E       OK         32 degrees C / 89 degrees F
      FPC 0 Sensor TopRight I        OK         29 degrees C / 84 degrees F
      FPC 0 Sensor TopLeft I         OK         23 degrees C / 73 degrees F
      FPC 0 Sensor TopRight E        OK         28 degrees C / 82 degrees F
      FPC 0 Sensor CPURight I        OK         30 degrees C / 86 degrees F
      FPC 0 Sensor CPULeft I         OK         29 degrees C / 84 degrees F
      FPC 0 Sensor Die Temp          OK         46 degrees C / 114 degrees F
      FPC 0 Mezz Temp                OK         23 degrees C / 73 degrees F
Fans  FPC 0 Fan Tray 0               OK         Spinning at normal speed
      FPC 0 Fan Tray 1               OK         Spinning at normal speed
      FPC 0 Fan Tray 2               OK         Spinning at normal speed
```

show chassis environment (QFX5700 Switch)

```
user@host> show chassis environment feb
FEB 0 status:
  State                      Online
  TD4 Die Temperature 0                  34 degrees C / 93 degrees F
  Intake Left Temp Sensor                24 degrees C / 75 degrees F
  Intake Right Temp Sensor               24 degrees C / 75 degrees F
  PEX Temperature Sensor                 20 degrees C / 68 degrees F
  Exhaust Left Temp Sensor               25 degrees C / 77 degrees F
  Exhaust Right Temp Sensor              26 degrees C / 78 degrees F
  Hotswap controller 0 Temp              19 degrees C / 66 degrees F
  Hotswap controller 1 Temp              22 degrees C / 71 degrees F
  DC-DC Controller 0 Internal Temp       27 degrees C / 80 degrees F
  DC-DC Controller 0 Rail 0 Temp         37 degrees C / 98 degrees F
  DC-DC Controller 0 Rail 1 Temp         34 degrees C / 93 degrees F
  DC-DC Controller 1 Temp                31 degrees C / 87 degrees F
  DC-DC Controller 2 Temp                28 degrees C / 82 degrees F
```

```
  DC-DC Controller 3 Temp                 26 degrees C / 78 degrees F
  DC-DC Controller 4 Temp                 32 degrees C / 89 degrees F
  DC-DC Controller 5 Temp                 34 degrees C / 93 degrees F
  Power 1
    BIAS3V3                     3296 mV
    FPGA_VDD1V0                 1006 mV
    FPGA_VDD1V2                 1210 mV
    FPGA_VDD3V3                 3308 mV
    PFE_VDD1V8                  1802 mV
    VDD1V2                      1202 mV
  Power 2
    Hotswap 0, 12V0            11889 mV
  Power 3
    Hotswap 1, 12V0            11942 mV
```

show chassis environment (ACX7509)

```
user@host> show chassis environment feb
FEB 0 status:
  State                      Online
  Q2C0 Die Temperature 0                  32 degrees C / 89 degrees F
  Q2C0 Die Temperature 1                  33 degrees C / 91 degrees F
  Q2C0 Die Temperature 2                  31 degrees C / 87 degrees F
  Q2C0 Die Temperature 3                  32 degrees C / 89 degrees F
  Q2C0 Die Temperature 4                  31 degrees C / 87 degrees F
  Q2C0 HBM Temperature                    32 degrees C / 89 degrees F
  Q2C1 Die Temperature 0                  35 degrees C / 95 degrees F
  Q2C1 Die Temperature 1                  33 degrees C / 91 degrees F
  Q2C1 Die Temperature 2                  34 degrees C / 93 degrees F
  Q2C1 Die Temperature 3                  33 degrees C / 91 degrees F
  Q2C1 Die Temperature 4                  34 degrees C / 93 degrees F
  Q2C1 HBM Temperature                    33 degrees C / 91 degrees F
  Intake Rear Left Temp                   28 degrees C / 82 degrees F
  PEX8725 PCIe Temp                       42 degrees C / 107 degrees F
  Intake Rear Right Temp                  24 degrees C / 75 degrees F
  Exhaust Rear Right Temp                 29 degrees C / 84 degrees F
  Q2C0 Diode Temp                         44 degrees C / 111 degrees F
  Exhaust Rear Left Temp                  30 degrees C / 86 degrees F
  Q2C1 Diode Temp                         46 degrees C / 114 degrees F
  HotSwap0 Temp                           24 degrees C / 75 degrees F
  HotSwap1 Temp                           25 degrees C / 77 degrees F
```

```
FPGA_VDD1V0 Temp                        34 degrees C / 93 degrees F
PCIE_VDD0V9 Temp                        41 degrees C / 105 degrees F
J2VDDIO_1P8 Temp                        30 degrees C / 86 degrees F
VDD3V3 Temp                             28 degrees C / 82 degrees F
VDD5V0 Temp                             29 degrees C / 84 degrees F
VDD2V5 Temp                             32 degrees C / 89 degrees F
J2TVDD_1P2 Temp                         36 degrees C / 96 degrees F
Q2C0-J2HBMVDD0C_1P2 Temp                31 degrees C / 87 degrees F
Q2C1-J2HBMVDD0C_1P2 Temp                34 degrees C / 93 degrees F
Q2C0 Core Rail 1 Temp                   34 degrees C / 93 degrees F
Q2C0 Pwr Reg Intrnl Temp                29 degrees C / 84 degrees F
Q2C1 Core Rail 0 Temp                   37 degrees C / 98 degrees F
Q2C0 Core Rail 0 Temp                   38 degrees C / 100 degrees F
Q2C1 Core Rail 1 Temp                   35 degrees C / 95 degrees F
Q2C1 Pwr Reg Intrnl Temp                29 degrees C / 84 degrees F
Power 1
  BIAS3V3                     3288 mV
  FPGA_VDD1V2                 1214 mV
  FPGA_VDD3V3                 3320 mV
  PFE0_J2HBMVDD_2P5           2528 mV
  PFE0_J2PHYVDDC_0P88          892 mV
  PFE0_J2PLLVDD_1P8           1822 mV
  PFE1_J2HBMVDD_2P5           2524 mV
  PFE1_J2PHYVDDC_0P88          894 mV
  PFE1_J2PLLVDD_1P8           1824 mV
  VDDCLK_3V3                  3304 mV
Power 2
  PFE0_J2PLLAVDD_1P8          1794 mV
  PFE0_J2PLLVDD_0P8            838 mV
  PFE1_J2PLLAVDD_1P8          1802 mV
  PFE1_J2PLLVDD_0P8            836 mV
Power 3
  Hotswap 0, 12V0            11955 mV
Power 4
  Hotswap 1, 12V0            11982 mV
Power 5
  PFE0_J2TRVDD_0P8             835 mV
  PFE0_J2VDDC_0P82             820 mV
Power 6
  PFE1_J2TRVDD_0P8             835 mV
  PFE1_J2VDDC_0P82             820 mV
Power 7
  J2VDDIO_1P8                 1795 mV
```

```
  Power 8
    VDD5V0                      4961 mV
  Power 9
    VDD2V5                      2494 mV
  Power 10
    VDD3V3                      3295 mV
  Power 11
    PCIE_VDD0V9                  906 mV
  Power 12
    J2TVDD_1P2                  1274 mV
  Power 13
    FPGA_VDD1V0                  991 mV
  Power 14
    PFE1_J2HBMVDDOC_1P2         1196 mV
  Power 15
    PFE0_J2HBMVDDOC_1P2         1196 mV
  Bus Revision               19
```

show chassis environment (ACX500 Router)

```
user@host> show chassis environment

Class Item                          Status     Measurement
      Power Mod                     OK         47 degrees C / 116 degrees F
      BCM54610                      OK         46 degrees C / 114 degrees F
      DPLL31404                     OK         45 degrees C / 113 degrees F
      CPLD                          OK         42 degrees C / 107 degrees F
      1588-FPGA                     OK         43 degrees C / 109 degrees F
      NPU                           OK         62 degrees C / 143 degrees F
      MAC sensor 1                  OK         40 degrees C / 104 degrees F
      MAC sensor 2                  OK         38 degrees C / 100 degrees F
      SFP PHY                       OK         38 degrees C / 100 degrees F
      Combo/RJ45 PHY                OK         37 degrees C / 98 degrees F
      SFP sensor 1                  OK         35 degrees C / 95 degrees F
      SFP sensor 2                  OK         33 degrees C / 91 degrees F
      SFP sensor 3                  OK         32 degrees C / 89 degrees F
      Routing Engine                OK         54 degrees C / 129 degrees F
      Heater off
```

show chassis environment (Junos OS Evolved)

```
user@host> show chassis environment
Class   Item                          Status   Measurement
Power   PSM 0                         OK       26 degrees C / 78 degrees F
        PSM 1                         OK       31 degrees C / 87 degrees F
        PSM 2                         Failed
Temp    CB 0 Intake A Temp Sensor     OK       23 degrees C / 73 degrees F
        CB 0 Intake B Temp Sensor     OK       22 degrees C / 71 degrees F
        CB 0 Exhaust A Temp Sensor    OK       26 degrees C / 78 degrees F
        CB 0 Exhaust B Temp Sensor    OK       29 degrees C / 84 degrees F
        CB 0 Middle Temp Sensor       OK       30 degrees C / 86 degrees F
        CB 1 Intake A Temp Sensor     OK       23 degrees C / 73 degrees F
        CB 1 Intake B Temp Sensor     OK       21 degrees C / 69 degrees F
        CB 1 Exhaust A Temp Sensor    OK       26 degrees C / 78 degrees F
        CB 1 Exhaust B Temp Sensor    OK       29 degrees C / 84 degrees F
        CB 1 Middle Temp Sensor       OK       29 degrees C / 84 degrees F
        Fan Tray 0 Inlet Temp Sensor  OK       24 degrees C / 75 degrees F
        Fan Tray 0 Outlet Temp Sensor OK       28 degrees C / 82 degrees F
        Fan Tray 1 Inlet Temp Sensor  OK       24 degrees C / 75 degrees F
        Fan Tray 1 Outlet Temp Sensor OK       28 degrees C / 82 degrees F
        FPC 0 BT-0 Temp Sensor 1      OK       44 degrees C / 111 degrees F
        FPC 0 BT-0 Temp Sensor 0      OK       44 degrees C / 111 degrees F
        FPC 0 BT-1 Temp Sensor 0      OK       45 degrees C / 113 degrees F
        FPC 0 BT-1 Temp Sensor 1      OK       45 degrees C / 113 degrees F
```

show chassis environment feb (ACX7509, ACX7348, ACX 7332)

```
user@host> show chassis environment feb
FEB 0 status:
  State                     Online
  PHY Die 0 Sensor 0 Temp                60 degrees C / 140 degrees F
  PHY Die 1 Sensor 0 Temp                65 degrees C / 149 degrees F
  PHY Die 2 Sensor 0 Temp                59 degrees C / 138 degrees F
  Q2C0 Die Temperature 0                 59 degrees C / 138 degrees F
  Q2C0 Die Temperature 1                 62 degrees C / 143 degrees F
  Q2C0 Die Temperature 2                 63 degrees C / 145 degrees F
  Q2C0 Die Temperature 3                 63 degrees C / 145 degrees F
  Q2C0 Die Temperature 4                 60 degrees C / 140 degrees F
  Q2C0 HBM Temperature                   60 degrees C / 140 degrees F
```

```
Intake Left Temp                        43 degrees C / 109 degrees F
Intake Right Temp                       42 degrees C / 107 degrees F
Exhaust Left Temp                       43 degrees C / 109 degrees F
Exhaust Right Temp                      44 degrees C / 111 degrees F
Q2C Diode Temp                          71 degrees C / 159 degrees F
DC-DC Controller 0 Internal Temp        54 degrees C / 129 degrees F
Q2C_TRVDD_0V8 Temp                      48 degrees C / 118 degrees F
Q2C_VDDC_0V82 Temp                      48 degrees C / 118 degrees F
Q2C_HBMVDDOC_1V2 Temp                   56 degrees C / 132 degrees F
PHY2_VDD0V8 Temp                        53 degrees C / 127 degrees F
PHY01_VDD0V8 Temp                       51 degrees C / 123 degrees F
Q2C_TVDD_1V2 Temp                       50 degrees C / 122 degrees F
PHY_AVDD1V2 Temp                        56 degrees C / 132 degrees F
VDD3V3_C Temp                           47 degrees C / 116 degrees F
VDD3V3_B Temp                           52 degrees C / 125 degrees F
CLK_VDD3V3 Temp                         51 degrees C / 123 degrees F
FPGA1V0 Temp                            51 degrees C / 123 degrees F
Hot Swap Controller 0 Temp              44 degrees C / 111 degrees F
Hot Swap Controller 1 Temp              42 degrees C / 107 degrees F
Power 1
  VDD3V3_C                     3300 mV
Power 2
  CLK_VDD3V3                   3308 mV
Power 3
  FPGA1V0                      1008 mV
Power 4
  Q2C_HBMVDDOC_1V2             1216 mV
Power 5
  Q2C_TVDD_1V2                 1213 mV
Power 6
  Q2C_TRVDD_0V8                 835 mV
  Q2C_VDDC_0V82                 860 mV
Power 7
  PHY01_VDD0V8                  800 mV
Power 8
  PHY_AVDD1V2                  1210 mV
Power 9
  VDD3V3_B                     3318 mV
Power 10
  PHY2_VDD0V8                   806 mV
Power 11
  BIAS3V3                      3316 mV
  BIAS5V0                      5010 mV
```

```
    FPGA1V2                      1206 mV
    FPGA1V8                      1804 mV
    FPGA2V5                      2504 mV
    VDD2V5                       2476 mV
    VDD3V3_A                     3300 mV
  Power 12
    Q2C_HBMVDD_2V5               2512 mV
    Q2C_PHYVDDC_0V88              888 mV
    Q2C_PLLAVDD_1V8              1802 mV
    Q2C_PLLVDD_0V8                832 mV
    Q2C_PLLVDD_1V8               1816 mV
    Q2C_VDDIO_1V8                1806 mV
  Power 13
    Hotswap 0, 12V0             12062 mV
  Power 14
    Hotswap 1, 12V0             12000 mV
  Bus Revision                 38
```

show chassis environment (MX304 with Optics)

When optics are connected to the device, the output displays transceiver name instead of the interface name.

```
user@host> show chassis environment
Class Item                           Status     Measurement
      Routing Engine 0 CPU           OK         34 degrees C / 93 degrees F
      RE 0 CPU Core-0 Temp Sensor    OK         34 degrees C / 93 degrees F
      RE 0 CPU Core-1 Temp Sensor    OK         32 degrees C / 89 degrees F
      RE 0 CPU Core-2 Temp Sensor    OK         31 degrees C / 87 degrees F
...
 FPC 0 xcvr-0/0/0               OK         26 degrees C / 78 degrees F
      FPC 0 xcvr-0/0/1               OK         28 degrees C / 82 degrees F
      FPC 0 xcvr-0/0/2               OK         27 degrees C / 80 degrees F
      FPC 0 xcvr-0/0/3               OK         31 degrees C / 87 degrees F
      FPC 0 xcvr-0/0/4               OK         29 degrees C / 84 degrees F
      FPC 0 xcvr-0/0/5               OK         35 degrees C / 95 degrees F
      FPC 0 xcvr-0/0/6               OK         33 degrees C / 91 degrees F
      FPC 0 xcvr-0/0/7               OK         36 degrees C / 96 degrees F
      FPC 0 xcvr-0/0/8               OK         28 degrees C / 82 degrees F
      FPC 0 xcvr-0/0/9               OK         36 degrees C / 96 degrees F
      FPC 0 xcvr-0/0/10              OK         28 degrees C / 82 degrees F
```

```
      FPC 0 xcvr-0/0/11              OK         29 degrees C / 84 degrees F
...
Power PEM 0                          Absent
      PEM 1                          OK         36 degrees C / 96 degrees F
Fans  Fan Tray 0 Fan 0               OK         Spinning at normal speed
      Fan Tray 0 Fan 1               OK         Spinning at normal speed
      Fan Tray 1 Fan 0               OK         Spinning at normal speed
      Fan Tray 1 Fan 1               OK         Spinning at normal speed
      Fan Tray 2 Fan 0               OK         Spinning at normal speed
      Fan Tray 2 Fan 1               OK         Spinning at normal speed
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

`monitored` option added in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

`pem` option introduced in Junos OS Release 12.3 for ACX4000 Universal Metro Routers.

`satellite` option introduced in Junos OS Release 14.2R3.

`all-members`, `local`, and `member` *member-id* options introduced in Junos OS Release 15.1 for MX2010 and MX2020 routers.

RELATED DOCUMENTATION

- show chassis environment adc
- show chassis environment cb
- show chassis environment ccg
- show chassis environment cip
- show chassis environment fpc | **17621**
- show chassis environment fpm
- show chassis environment lcc
- show chassis environment mcs
- show chassis environment monitored
- show chassis environment pcg
- show chassis environment pdu
- show chassis environment pem | **17760**

show chassis environment psm
show chassis environment psu
show chassis environment routing-engine | 17800
show chassis environment scg
show chassis environment sfb
show chassis environment sib
show chassis environment sfc

show chassis environment adc

IN THIS SECTION

- Syntax | 17553
- Description | 17553
- Options | 17554
- Required Privilege Level | 17554
- Output Fields | 17554
- Sample Output | 17555
- Release Information | 17569

Syntax

```
show chassis environment adc
<adc-slot-number>
<all-members>
<local>
<member member-id>
```

Description

Display chassis environmental information about the adapter cards.

Options

none	Display environmental information about all adapter cards.
adc-slot-number	(Optional) Display environmental information about the specified adapter card. For MX2020 routers, replace *adc-slot-number* with a value of `0` through `19`. For MX2010 and MX2008 routers, replace *adc-slot-number* with a value of `0` through `9`.
all-members	(Optional) Display chassis environmental information about the adapter cards (ADCs) in all members of the Virtual Chassis configuration.
local	(Optional) Display chassis environmental information about the ADCs in the local member of the Virtual Chassis.
member ***member-id***	(Optional) Display chassis environmental information about the ADCs in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 222 on page 17554 lists the output fields for the `show chassis environment adc` command. Output fields are listed in the approximate order in which they appear.

Table 222: show chassis environment adc Output Fields

Field Name	Field Description
`State`	Status of the adapter card. • `Online`—The adapter card is online and running. • `Offline`—Adapter card is powered down.

Table 222: show chassis environment adc Output Fields *(Continued)*

Field Name	Field Description
`Temperature`	Temperature in Celsius (C) and Fahrenheit (F) of the air flowing past the adapter card. • `Intake Temperature`—Measures the temperature of the air intake. • `Exhaust Temperature`—Measures the temperature of the hot air exhaust. • `ADC-XF1`—Measures the temperature of the ADC chipset, ADC-XF1. • `ADC-XF0`—Measures the temperature of the ADC chipset, ADC-XF0.
`Power`	Power required and measured on the adapter card. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.

Sample Output

show chassis environment adc (MX2020 Router)

```
user@host> show chassis environment adc
ADC 0 status:
  State                      Online
  Intake Temperature         39 degrees C / 102 degrees F
  Exhaust Temperature        50 degrees C / 122 degrees F
  ADC-XF1 Temperature        58 degrees C / 136 degrees F
  ADC-XF0 Temperature        64 degrees C / 147 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
```

```
ADC 1 status:
  State                      Online
  Intake Temperature         38 degrees C / 100 degrees F
  Exhaust Temperature        48 degrees C / 118 degrees F
  ADC-XF1 Temperature        59 degrees C / 138 degrees F
  ADC-XF0 Temperature        61 degrees C / 141 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 2 status:
  State                      Online
  Intake Temperature         36 degrees C / 96 degrees F
  Exhaust Temperature        50 degrees C / 122 degrees F
  ADC-XF1 Temperature        52 degrees C / 125 degrees F
  ADC-XF0 Temperature        59 degrees C / 138 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
ADC 3 status:
  State                      Online
  Intake Temperature         39 degrees C / 102 degrees F
  Exhaust Temperature        50 degrees C / 122 degrees F
  ADC-XF1 Temperature        61 degrees C / 141 degrees F
  ADC-XF0 Temperature        63 degrees C / 145 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
```

```
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 4 status:
  State                      Online
  Intake Temperature         38 degrees C / 100 degrees F
  Exhaust Temperature        49 degrees C / 120 degrees F
  ADC-XF1 Temperature        60 degrees C / 140 degrees F
  ADC-XF0 Temperature        62 degrees C / 143 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 5 status:
  State                      Online
  Intake Temperature         37 degrees C / 98 degrees F
  Exhaust Temperature        52 degrees C / 125 degrees F
  ADC-XF1 Temperature        55 degrees C / 131 degrees F
  ADC-XF0 Temperature        66 degrees C / 150 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
```

```
    LTC3880-XF1-1.5v-CH0      1500 mV
ADC 6 status:
  State                      Online
  Intake Temperature         39 degrees C / 102 degrees F
  Exhaust Temperature        51 degrees C / 123 degrees F
  ADC-XF1 Temperature        58 degrees C / 136 degrees F
  ADC-XF0 Temperature        64 degrees C / 147 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 7 status:
  State                      Online
  Intake Temperature         38 degrees C / 100 degrees F
  Exhaust Temperature        52 degrees C / 125 degrees F
  ADC-XF1 Temperature        61 degrees C / 141 degrees F
  ADC-XF0 Temperature        69 degrees C / 156 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
ADC 8 status:
  State                      Online
  Intake Temperature         38 degrees C / 100 degrees F
  Exhaust Temperature        50 degrees C / 122 degrees F
  ADC-XF1 Temperature        63 degrees C / 145 degrees F
  ADC-XF0 Temperature        64 degrees C / 147 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
```

```
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 9 status:
  State                      Online
  Intake Temperature         40 degrees C / 104 degrees F
  Exhaust Temperature        50 degrees C / 122 degrees F
  ADC-XF1 Temperature        59 degrees C / 138 degrees F
  ADC-XF0 Temperature        62 degrees C / 143 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 10 status:
  State                      Online
  Intake Temperature         46 degrees C / 114 degrees F
  Exhaust Temperature        52 degrees C / 125 degrees F
  ADC-XF1 Temperature        66 degrees C / 150 degrees F
  ADC-XF0 Temperature        65 degrees C / 149 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
```

```
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 11 status:
  State                      Online
  Intake Temperature         47 degrees C / 116 degrees F
  Exhaust Temperature        53 degrees C / 127 degrees F
  ADC-XF1 Temperature        64 degrees C / 147 degrees F
  ADC-XF0 Temperature        65 degrees C / 149 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 12 status:
  State                      Online
  Intake Temperature         48 degrees C / 118 degrees F
  Exhaust Temperature        54 degrees C / 129 degrees F
  ADC-XF1 Temperature        66 degrees C / 150 degrees F
  ADC-XF0 Temperature        65 degrees C / 149 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
```

```
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 13 status:
  State                      Online
  Intake Temperature         48 degrees C / 118 degrees F
  Exhaust Temperature        55 degrees C / 131 degrees F
  ADC-XF1 Temperature        66 degrees C / 150 degrees F
  ADC-XF0 Temperature        67 degrees C / 152 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1034 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1503 mV
ADC 14 status:
  State                      Online
  Intake Temperature         50 degrees C / 122 degrees F
  Exhaust Temperature        57 degrees C / 134 degrees F
  ADC-XF1 Temperature        68 degrees C / 154 degrees F
  ADC-XF0 Temperature        72 degrees C / 161 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1034 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1034 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 15 status:
  State                      Online
  Intake Temperature         49 degrees C / 120 degrees F
  Exhaust Temperature        56 degrees C / 132 degrees F
```

```
  ADC-XF1 Temperature        68 degrees C / 154 degrees F
  ADC-XF0 Temperature        68 degrees C / 154 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1034 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1034 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 16 status:
  State                      Online
  Intake Temperature         51 degrees C / 123 degrees F
  Exhaust Temperature        56 degrees C / 132 degrees F
  ADC-XF1 Temperature        67 degrees C / 152 degrees F
  ADC-XF0 Temperature        68 degrees C / 154 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 17 status:
  State                      Online
  Intake Temperature         51 degrees C / 123 degrees F
  Exhaust Temperature        56 degrees C / 132 degrees F
  ADC-XF1 Temperature        68 degrees C / 154 degrees F
  ADC-XF0 Temperature        69 degrees C / 156 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
```

```
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1034 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 18 status:
  State                      Online
  Intake Temperature         52 degrees C / 125 degrees F
  Exhaust Temperature        57 degrees C / 134 degrees F
  ADC-XF1 Temperature        66 degrees C / 150 degrees F
  ADC-XF0 Temperature        71 degrees C / 159 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1034 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1034 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 19 status:
  State                      Online
  Intake Temperature         49 degrees C / 120 degrees F
  Exhaust Temperature        56 degrees C / 132 degrees F
  ADC-XF1 Temperature        67 degrees C / 152 degrees F
  ADC-XF0 Temperature        70 degrees C / 158 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
```

```
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
```

show chassis environment adc (MX2010 Router)

```
user@host> show chassis environment adc
ADC 0 status:
  State                      Online
  Intake Temperature         33 degrees C / 91 degrees F
  Exhaust Temperature        42 degrees C / 107 degrees F
  ADC-XF1 Temperature        46 degrees C / 114 degrees F
  ADC-XF0 Temperature        53 degrees C / 127 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL      998 mV
    LTC3880-XF0-1.0v-CH0       998 mV
    LTC3880-XF0-1.0v-CH1      1001 mV
    LTC3880-XF0-1.5v-RAIL     1454 mV
    LTC3880-XF0-1.5v-CH0      1454 mV
    LTC3880-XF0-1.5v-CH1      1456 mV
    LTC3880-XF1-1.0v-RAIL      998 mV
    LTC3880-XF1-1.0v-CH0       998 mV
    LTC3880-XF1-1.0v-CH1      1002 mV
    LTC3880-XF1-1.5v-RAIL     1454 mV
    LTC3880-XF1-1.5v-CH0      1454 mV
    LTC3880-XF1-1.5v-CH1      1457 mV
ADC 1 status:
  State                      Online
  Intake Temperature         32 degrees C / 89 degrees F
  Exhaust Temperature        42 degrees C / 107 degrees F
  ADC-XF1 Temperature        44 degrees C / 111 degrees F
  ADC-XF0 Temperature        52 degrees C / 125 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL      998 mV
    LTC3880-XF0-1.0v-CH0       998 mV
    LTC3880-XF0-1.0v-CH1      1002 mV
    LTC3880-XF0-1.5v-RAIL     1454 mV
    LTC3880-XF0-1.5v-CH0      1454 mV
    LTC3880-XF0-1.5v-CH1      1456 mV
```

```
    LTC3880-XF1-1.0v-RAIL      999 mV
    LTC3880-XF1-1.0v-CH0       999 mV
    LTC3880-XF1-1.0v-CH1      1002 mV
    LTC3880-XF1-1.5v-RAIL     1454 mV
    LTC3880-XF1-1.5v-CH0      1454 mV
    LTC3880-XF1-1.5v-CH1      1456 mV
ADC 2 status:
  State                      Online
  Intake Temperature         35 degrees C / 95 degrees F
  Exhaust Temperature        42 degrees C / 107 degrees F
  ADC-XF1 Temperature        48 degrees C / 118 degrees F
  ADC-XF0 Temperature        54 degrees C / 129 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
LTC3880-XF1-1.5v-CH1      1502 mV
ADC 3 status:
  State                      Online
  Intake Temperature         35 degrees C / 95 degrees F
  Exhaust Temperature        40 degrees C / 104 degrees F
  ADC-XF1 Temperature        44 degrees C / 111 degrees F
  ADC-XF0 Temperature        51 degrees C / 123 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL      999 mV
    LTC3880-XF0-1.0v-CH0       999 mV
    LTC3880-XF0-1.0v-CH1      1002 mV
    LTC3880-XF0-1.5v-RAIL     1454 mV
    LTC3880-XF0-1.5v-CH0      1454 mV
    LTC3880-XF0-1.5v-CH1      1456 mV
    LTC3880-XF1-1.0v-RAIL      999 mV
    LTC3880-XF1-1.0v-CH0       999 mV
    LTC3880-XF1-1.0v-CH1      1002 mV
    LTC3880-XF1-1.5v-RAIL     1454 mV
    LTC3880-XF1-1.5v-CH0      1454 mV
```

```
    LTC3880-XF1-1.5v-CH1      1457 mV
ADC 4 status:
  State                      Online
  Intake Temperature         31 degrees C / 87 degrees F
  Exhaust Temperature        43 degrees C / 109 degrees F
  ADC-XF1 Temperature        48 degrees C / 118 degrees F
  ADC-XF0 Temperature        56 degrees C / 132 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 5 status:
  State                      Online
  Intake Temperature         33 degrees C / 91 degrees F
  Exhaust Temperature        43 degrees C / 109 degrees F
  ADC-XF1 Temperature        47 degrees C / 116 degrees F
  ADC-XF0 Temperature        54 degrees C / 129 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL      999 mV
    LTC3880-XF0-1.0v-CH0       999 mV
    LTC3880-XF0-1.0v-CH1      1002 mV
    LTC3880-XF0-1.5v-RAIL     1454 mV
    LTC3880-XF0-1.5v-CH0      1454 mV
    LTC3880-XF0-1.5v-CH1      1456 mV
    LTC3880-XF1-1.0v-RAIL      998 mV
    LTC3880-XF1-1.0v-CH0       998 mV
    LTC3880-XF1-1.0v-CH1      1002 mV
    LTC3880-XF1-1.5v-RAIL     1454 mV
 LTC3880-XF1-1.5v-CH0      1454 mV
    LTC3880-XF1-1.5v-CH1      1457 mV
ADC 6 status:
  State                      Online
  Intake Temperature         32 degrees C / 89 degrees F
  Exhaust Temperature        42 degrees C / 107 degrees F
```

```
  ADC-XF1 Temperature        47 degrees C / 116 degrees F
  ADC-XF0 Temperature        55 degrees C / 131 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 7 status:
  State                      Online
  Intake Temperature         36 degrees C / 96 degrees F
  Exhaust Temperature        43 degrees C / 109 degrees F
  ADC-XF1 Temperature        46 degrees C / 114 degrees F
  ADC-XF0 Temperature        55 degrees C / 131 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1502 mV
ADC 8 status:
  State                      Online
  Intake Temperature         35 degrees C / 95 degrees F
  Exhaust Temperature        43 degrees C / 109 degrees F
  ADC-XF1 Temperature        44 degrees C / 111 degrees F
  ADC-XF0 Temperature        51 degrees C / 123 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL      999 mV
    LTC3880-XF0-1.0v-CH0       999 mV
```

```
    LTC3880-XF0-1.0v-CH1      1002 mV
    LTC3880-XF0-1.5v-RAIL     1455 mV
    LTC3880-XF0-1.5v-CH0      1455 mV
    LTC3880-XF0-1.5v-CH1      1456 mV
    LTC3880-XF1-1.0v-RAIL      999 mV
    LTC3880-XF1-1.0v-CH0       999 mV
    LTC3880-XF1-1.0v-CH1      1002 mV
 LTC3880-XF1-1.5v-RAIL     1455 mV
    LTC3880-XF1-1.5v-CH0      1455 mV
    LTC3880-XF1-1.5v-CH1      1457 mV
ADC 9 status:
  State                      Online
  Intake Temperature         31 degrees C / 87 degrees F
  Exhaust Temperature        43 degrees C / 109 degrees F
  ADC-XF1 Temperature        48 degrees C / 118 degrees F
  ADC-XF0 Temperature        56 degrees C / 132 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1455 mV
    LTC3880-XF0-1.5v-CH0      1455 mV
    LTC3880-XF0-1.5v-CH1      1457 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1454 mV
    LTC3880-XF1-1.5v-CH0      1454 mV
    LTC3880-XF1-1.5v-CH1      1457 mV
```

show chassis environment adc (MX2008 Router)

```
user@host>show chassis environment adc
ADC 7 status:
  State                      Online
  Intake Temperature         32 degrees C / 89 degrees F
  Exhaust Temperature        39 degrees C / 102 degrees F
  ADC-XF1 Temperature        46 degrees C / 114 degrees F
  ADC-XF0 Temperature        54 degrees C / 129 degrees F
  Power
    LTC3880-XF0-1.0v-RAIL       1029 mV
```

```
    LTC3880-XF0-1.0v-CH0        1029 mV
    LTC3880-XF0-1.0v-CH1        1033 mV
    LTC3880-XF0-1.5v-RAIL       1499 mV
    LTC3880-XF0-1.5v-CH0        1499 mV
    LTC3880-XF0-1.5v-CH1        1501 mV
    LTC3880-XF1-1.0v-RAIL       1029 mV
    LTC3880-XF1-1.0v-CH0        1029 mV
    LTC3880-XF1-1.0v-CH1        1033 mV
    LTC3880-XF1-1.5v-RAIL       1499 mV
    LTC3880-XF1-1.5v-CH0        1499 mV
    LTC3880-XF1-1.5v-CH1        1502 mV
    ADC-2.5V                    2851 mV
    ADC-3.3V                    3787 mV
```

Release Information

Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 Universal Routing Platforms.

`all-members`, `local`, and `member` *`member-id`* options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.

RELATED DOCUMENTATION

show chassis adc | 17291

show chassis environment cb (Security)

IN THIS SECTION

- Syntax | 17570
- Description | 17570
- Options | 17570
- Required Privilege Level | 17570
- Output Fields | 17570
- Sample Output | 17571

Release Information | 17573

Syntax

```
show chassis environment cb
<slot>
```

Description

Display environmental information about the Control Boards (CBs) installed on SRX Series Firewalls.

Options

slot (Optional) Display environmental information about the specified CB.

Required Privilege Level

view

Output Fields

Table 223 on page 17570 lists the output fields for the **show chassis environment cb** command. Output fields are listed in the approximate order in which they appear.

Table 223: show chassis environment cb Output Fields

Field Name	Field Description
State	Status of the CB. If two CBs are installed and online, one is functioning as the primary, and the other is the standby. • **Online**—CB is online and running. • **Offline**— CB is powered down.

Table 223: show chassis environment cb Output Fields *(Continued)*

Field Name	Field Description
Temperature	Temperature in Celsius (C) and Fahrenheit (F) of the air flowing past the CB. • **Temperature Intake**—Measures the temperature of the air intake to cool the power supplies. • **Temperature Exhaust**—Measures the temperature of the hot air exhaust.
Power	Power required and measured on the CB. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
BUS Revision	Revision level of the generic bus device.
FPGA Revision	Revision level of the field-programmable gate array (FPGA).
PMBus device	Enhanced SCB on SRX Series Firewalls allows the system to save power by supplying only the amount of voltage that is required. Configurable PMBus devices are used to provide the voltage for each individual device. There is one PMBus device for each XF ASIC so that the output can be customized to each device. The following PMBus device information is displayed for devices with Enhanced MX SCB: • **Expected voltage** • **Measured voltage** • **Measured current** • **Calculated power**

Sample Output

show chassis environment cb (SRX5600 devices with SRX5K-SCB3 [SCB3] and Enhanced Midplanes)

```
user@host> show chassis environment cb node 0
node0:
--------------------------------------------------------------------------
```

```
---------------------------------------------------------------------------
CB 0 status:
  State                      Online Master
  Temperature                34 degrees C / 93 degrees F
  Power 1
    1.0 V                       1002
    1.2 V                       1198
    1.5 V                       1501
    1.8 V                       1801
    2.5 V                       2507
    3.3 V                       3300
    5.0 V                       5014
    5.0 V RE                    4982
    12.0 V                     11988
    12.0 V RE                  11930
  Power 2
    4.6 V bias MidPlane         4801
    11.3 V bias PEM            11292
    11.3 V bias FPD            11272
    11.3 V bias POE 0          11214
    11.3 V bias POE 1          11253
  Bus Revision               96
  FPGA Revision              16
  PMBus             Expected   Measured   Measured  Calculated
  device            voltage    voltage    current   power
    XF ASIC A        1033 mV    1033 mV   15500 mA   16011 mW
    XF ASIC B        1034 mV    1033 mV   15000 mA   15495 mW
```

show chassis environment cb node 1 (SRX5600 devices with SRX5K-SCB3 [SCB3] and Enhanced Midplanes)

```
user@host> show chassis environment cb node 1
node1:
---------------------------------------------------------------------------
CB 0 status:
  State                      Online Master
  Temperature                35 degrees C / 95 degrees F
  Power 1
    1.0 V                       1002
    1.2 V                       1198
    1.5 V                       1504
```

```
    1.8 V                        1801
    2.5 V                        2507
    3.3 V                        3325
    5.0 V                        5014
    5.0 V RE                     4943
    12.0 V                      12007
    12.0 V RE                   12007
  Power 2
    4.6 V bias MidPlane          4814
    11.3 V bias PEM             11272
    11.3 V bias FPD             11330
    11.3 V bias POE 0           11176
    11.3 V bias POE 1           11292
  Bus Revision                96
  FPGA Revision               16
  PMBus              Expected   Measured   Measured  Calculated
  device             voltage    voltage    current   power
    XF ASIC A          958 mV     959 mV   13500 mA   12946 mW
    XF ASIC B         1033 mV    1031 mV   16500 mA   17011 mW
```

Release Information

Command introduced in Junos OS Release 9.2.

Starting with Junos OS Release 12.1X47-D15, the SRX5K-SCBE (SCB2) is introduced and starting with Junos OS Release 15.1X49-D10, the SRX5K-SCB3 (SCB3) with enhanced midplane is introduced.

RELATED DOCUMENTATION

request chassis cb (Security) | 14440

show chassis environment cb

IN THIS SECTION

- Syntax | 17574
- Syntax (TX Matrix Routers) | 17574

- Syntax (TX Matrix Plus Routers) | **17574**
- Syntax (MX Series Routers) | **17575**
- Syntax (MX104 Universal Routing Platforms) | **17575**
- Syntax (MX2010, MX2020, MX10003, MX10004, MX204, MX2008, and MX10008 Universal Routing Platforms; EX9251 and EX9253 Switches) | **17575**
- Syntax (QFabric System) | **17575**
- Description | **17575**
- Options | **17575**
- Required Privilege Level | **17576**
- Output Fields | **17577**
- Sample Output | **17578**
- Release Information | **17613**

Syntax

```
show chassis environment cb
<slot>
```

Syntax (TX Matrix Routers)

```
show chassis environment cb
<lcc number | scc>
<slot>
```

Syntax (TX Matrix Plus Routers)

```
show chassis environment cb
<lcc number | sfc number >
<slot>
```

Syntax (MX Series Routers)

```
show chassis environment cb
<slot>
<all-members>
<local>
<member member-id>
```

Syntax (MX104 Universal Routing Platforms)

```
show chassis environment cb
```

Syntax (MX2010, MX2020, MX10003, MX10004, MX204, MX2008, and MX10008 Universal Routing Platforms; EX9251 and EX9253 Switches)

```
show chassis environment cb
<slot>
```

Syntax (QFabric System)

```
show chassis environment cb
<slot interconnect-device interconnect-device-name>
< interconnect-device interconnect-device-name slot>
```

Description

(M120, M320, MX Series, and T Series routers, EX8200 switches, and PTX Series Packet Transport Routers only) Display environmental information about the Control Boards (CBs).

Options

none Display environmental information about all CBs. For a TX Matrix router, display environmental information about all CBs on the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display environmental information about all CBs on the TX Matrix Plus router and its attached T1600 or T4000 routers.

all-members (MX Series routers only) (Optional) Display environmental information about the CBs on all the members of the Virtual Chassis configuration.

interconnect-device (QFabric systems only) Display environmental information about CBs on the Interconnect device.

lcc *number* (TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Display environmental information about the CBs on the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display environmental information about the CBs on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

scc (TX Matrix router only) (Optional) Display environmental information about the CBs in the TX Matrix router (switch-card chassis).

sfc *number* (TX Matrix Plus router only) (Optional) Display environmental information about the CBs in the TX Matrix Plus router (or switch-fabric chassis).

slot (Optional) Display environmental information about the specified CB. On routers and PTX Series Packet Transport Routers, replace ***slot*** with **0** or **1**. On EX Series switches replace ***slot*** with **0**, **1**, or **2**. On QFX Series switches, replace ***slot*** with **0** or **1**.

Required Privilege Level

view

Output Fields

Table 224 on page 17577 lists the output fields for the `show chassis environment cb` command. Output fields are listed in the approximate order in which they appear.

Table 224: show chassis environment cb Output Fields

Field Name	Field Description
`State`	Status of the CB. If two CBs are installed and online, one is functioning as the primary, and the other is the standby. • **Online**—CB is online and running. • **Offline**— CB is powered down. **NOTE**: On the EX8208 switch, the installation can include three CBs.
Temperature	Temperature in Celsius (C) and Fahrenheit (F) of the air flowing past the CB. • **Temperature Intake**—Measures the temperature of the air intake to cool the power supplies. • **Temperature Exhaust**—Measures the temperature of the hot air exhaust. **NOTE**: On the MX2010, MX2020, and MX2008 routers, the intake temperature measures the temperature of the air intake to cool the Control Board (CB). The MX2010, MX2020, and MX2008 routers include intake and exhaust temperatures for multiple zones (`Intake A`, `Intake B`, `Intake C`, `Exhaust A`, `Exhaust B`, and `TCBC`).The QFabric system and QFX5700 include intake and exhaust temperatures for `Intake 1`, `Intake 2`, `Exhaust 1`, and `Exhaust 2`.)
Power	Power required and measured on the CB. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
BUS Revision	Revision level of the generic bus device. (Not on switches.)
FPGA Revision	Revision level of the field-programmable gate array (FPGA). (Not on switches.)

Table 224: show chassis environment cb Output Fields *(Continued)*

Field Name	Field Description
PMBus device (on MX240, MX480, and MX960 routers with Enhanced MX SCB)	Enhanced SCB on MX 240, MX480, and MX960 routers allows the system to save power by supplying only the amount of voltage that is required. Configurable PMBus devices are used to provide the voltage for each individual device. There is one PMBus device for each XF ASIC so that the output can be customized to each device. The following PMBus device information is displayed for routers with Enhanced MX SCB: • **Expected voltage** • **Measured voltage** • **Measured current** • **Calculated power**

Sample Output

show chassis environment cb (M120 Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
  Temperature                33 degrees C / 91 degrees F
  Power
    1.2 V                     1214 mV
    1.5 V                     1495 mV
    2.5 V                     2494 mV
    3.3 V                     3319 mV
    5.0 V                     5085 mV
    3.3 V bias                3296 mV
  Bus Revision               12
  FPGA Revision              17
CB 1 status:
  State                      Online Standby
  Temperature                34 degrees C / 93 degrees F
  Power
    1.2 V                     1195 mV
    1.5 V                     1495 mV
```

```
    2.5 V                      2504 mV
    3.3 V                      3312 mV
    5.0 V                      5111 mV
    3.3 V bias                 3296 mV
  Bus Revision                12
  FPGA Revision               17
```

show chassis environment cb (M320 Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                           Online Master
  Temperature                  29 degrees C / 84 degrees F
  Power:
    1.8 V                    1805 mV
    2.5 V                    2501 mV
    3.3 V                    3293 mV
    4.6 V                    4725 mV
    5.0 V                    5032 mV
    12.0 V                  11975 mV
    3.3 V bias               3286 mV
    8.0 V bias               7589 mV
  BUS Revision                 40
  FPGA Revision                 7
CB 1 status:
  State                           Online Standby
  Temperature                  32 degrees C / 89 degrees F
  Power:
    1.8 V                    1802 mV
    2.5 V                    2482 mV
    3.3 V                    3289 mV
    4.6 V                    4720 mV
    5.0 V                    5001 mV
    12.0 V                  11946 mV
    3.3 V bias               3274 mV
    8.0 V bias               7562 mV
  BUS Revision                 40
  FPGA Revision                 7
```

show chassis environment cb (MX80 Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
  Temperature                36 degrees C / 96 degrees F
  Power 1
    1.0 V                     1034 mV
    1.0 V MQ                  1037 mV
    1.0 V LU                  1005 mV
    1.2 V                     1218 mV
    1.5 V                     1524 mV
    1.8 V                     1814 mV
    2.5 V                     2558 mV
    3.3 V                     3296 mV
    5.0 V                     5233 mV
    5.0 V bias                5207 mV
    12.0 V                   12162 mV
```

show chassis environment cb (MX104 Router)

```
user@host > show chassis environment cb
CB 0 status:
  State                      Online Master
  Temperature                33 degrees C / 91 degrees F
  Power 1
    0.75 V                     751 mV
    1.0 V                     1005 mV
    1.1 V                     1113 mV
    1.5 V                     1494 mV
    2.5 V                     2518 mV
    3.3 V                     3338 mV
    5.0 V                     4960 mV
    12.0 V                   12006 mV
  FPGA Revision              25
CB 1 status:
  State                      Empty
```

show chassis environment cb (MX240 Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Standby
  Temperature                37 degrees C / 98 degrees F
  Power 1
    1.2 V                     1208 mV
    1.5 V                     1521 mV
    1.8 V                     1811 mV
    2.5 V                     2513 mV
    3.3 V                     3332 mV
    5.0 V                     5059 mV
    12.0 V                   12162 mV
    1.25 V                    1260 mV
    3.3 V SM3                 3306 mV
    5.0 V RE                  5085 mV
    12.0 V RE                11872 mV
  Power 2
    11.3 V bias PEM          11272 mV
    4.6 V bias MidPlane       4827 mV
    11.3 V bias FPD          11272 mV
    11.3 V bias POE 0        11292 mV
    11.3 V bias POE 1        11253 mV
  Bus Revision               42
  FPGA Revision              1
```

show chassis environment cb (MX240 Router with Enhanced MX SCB)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Standby
  Temperature                37 degrees C / 98 degrees F
  Power 1
    1.2 V                     1208 mV
    1.5 V                     1521 mV
    1.8 V                     1811 mV
    2.5 V                     2513 mV
    3.3 V                     3332 mV
    5.0 V                     5059 mV
```

```
    12.0 V                    12162 mV
    1.25 V                     1260 mV
    3.3 V SM3                  3306 mV
    5.0 V RE                   5085 mV
    12.0 V RE                 11872 mV
  Power 2
    11.3 V bias PEM           11272 mV
    4.6 V bias MidPlane        4827 mV
    11.3 V bias FPD           11272 mV
    11.3 V bias POE 0         11292 mV
    11.3 V bias POE 1         11253 mV
  Bus Revision                42
  FPGA Revision               1
  PMBus              Expected   Measured   Measured  Calculated
  device             voltage    voltage    current   power
    XF ASIC A         1000 mV     997 mV   11031 mA   10997 mW
    XF ASIC B         1000 mV     996 mV   12125 mA   12076 mW
```

show chassis environment cb (MX480 Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                       Online Master
  Temperature                 41 degrees C / 105 degrees F
  Power 1
    1.2 V                      1202 mV
    1.5 V                      1511 mV
    1.8 V                      1798 mV
    2.5 V                      2507 mV
    3.3 V                      3312 mV
    5.0 V                      5027 mV
    12.0 V                    12200 mV
    1.25 V                     1260 mV
    3.3 V SM3                  3293 mV
    5 V RE                     5040 mV
    12 V RE                   11910 mV
  Power 2
    11.3 V bias PEM           11156 mV
    4.6 V bias MidPlane        4801 mV
    11.3 V bias FPD           11214 mV
    11.3 V bias POE 0         11098 mV
```

```
    11.3 V bias POE 1        11330 mV
  Bus Revision               42
  FPGA Revision              1
```

show chassis environment cb (MX480 Router with Enhanced MX SCB)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
  Temperature                41 degrees C / 105 degrees F
  Power 1
    1.2 V                     1202 mV
    1.5 V                     1511 mV
    1.8 V                     1798 mV
    2.5 V                     2507 mV
    3.3 V                     3312 mV
    5.0 V                     5027 mV
    12.0 V                   12200 mV
    1.25 V                    1260 mV
    3.3 V SM3                 3293 mV
    5 V RE                    5040 mV
    12 V RE                  11910 mV
  Power 2
    11.3 V bias PEM          11156 mV
    4.6 V bias MidPlane       4801 mV
    11.3 V bias FPD          11214 mV
    11.3 V bias POE 0        11098 mV
    11.3 V bias POE 1        11330 mV
  Bus Revision               42
  FPGA Revision              1
  PMBus             Expected   Measured   Measured  Calculated
  device            voltage    voltage    current   power
    XF ASIC A        1000 mV     997 mV   11031 mA   10997 mW
    XF ASIC B        1000 mV     996 mV   12125 mA   12076 mW
```

show chassis environment cb (MX960 Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
```

```
  Temperature                24 degrees C / 75 degrees F
  Power 1
    1.2 V                     1965 mV
    1.5 V                     2465 mV
    1.8 V                     2990 mV
    2.5 V                     3296 mV
    3.3 V                     3296 mV
    5.0 V                     6593 mV
    12.0 V                   13187 mV
    3.3 V bias                3296 mV
    1.25 V                    1994 mV
    3.3 V SM3                 3296 mV
    5 V RE                    6593 mV
    12 V RE                  13174 mV
  Power 2                    Sensor failure
  Bus Revision               4
  FPGA Revision              3
```

show chassis environment cb (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
  Temperature                24 degrees C / 75 degrees F
  Power 1
    1.2 V                     1965 mV
    1.5 V                     2465 mV
    1.8 V                     2990 mV
    2.5 V                     3296 mV
    3.3 V                     3296 mV
    5.0 V                     6593 mV
    12.0 V                   13187 mV
    3.3 V bias                3296 mV
    1.25 V                    1994 mV
    3.3 V SM3                 3296 mV
    5 V RE                    6593 mV
    12 V RE                  13174 mV
  Power 2                    Sensor failure
  Bus Revision               4
  FPGA Revision              3
  PMBus             Expected   Measured   Measured  Calculated
```

```
  device            voltage    voltage    current   power
    XF ASIC A        1000 mV     997 mV   11031 mA   10997 mW
    XF ASIC B        1000 mV     996 mV   12125 mA   12076 mW
```

show chassis environment cb (MX2020 Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
  IntakeA-Zone0 Temperature  44 degrees C / 111 degrees F
  IntakeB-Zone1 Temperature  34 degrees C / 93 degrees F
  IntakeC-Zone0 Temperature  45 degrees C / 113 degrees F
  ExhaustA-Zone0 Temperature 43 degrees C / 109 degrees F
  ExhaustB-Zone1 Temperature 36 degrees C / 96 degrees F
  TCBC-Zone0 Temperature     39 degrees C / 102 degrees F
  Power 1
    1.0 V                     1011 mV
    1.2 V                     1208 mV
    1.8 V                     1801 mV
    2.5 V                     2552 mV
    3.3 V                     3312 mV
    5.0 V                     5040 mV
    5.0 V RE                  4988 mV
    12.0 V                   12065 mV
    12.0 V RE                12046 mV
  Bus Revision               99
  FPGA Revision              270
CB 1 status:
  State                      Online Standby
  IntakeA-Zone0 Temperature  45 degrees C / 113 degrees F
  IntakeB-Zone1 Temperature  41 degrees C / 105 degrees F
  IntakeC-Zone0 Temperature  46 degrees C / 114 degrees F
  ExhaustA-Zone0 Temperature 44 degrees C / 111 degrees F
  ExhaustB-Zone1 Temperature 41 degrees C / 105 degrees F
  TCBC-Zone0 Temperature     45 degrees C / 113 degrees F
  Power 1
    1.0 V                     1008 mV
    1.2 V                     1208 mV
    1.8 V                     1798 mV
    2.5 V                     2539 mV
    3.3 V                     3325 mV
```

```
    5.0 V                     5033 mV
    5.0 V RE                  4950 mV
    12.0 V                   12046 mV
    12.0 V RE                11968 mV
  Bus Revision               99
  FPGA Revision              0
```

show chassis environment cb (MX2010 Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
  IntakeA-Zone0 Temperature  36 degrees C / 96 degrees F
  IntakeB-Zone1 Temperature  30 degrees C / 86 degrees F
  IntakeC-Zone0 Temperature  38 degrees C / 100 degrees F
  ExhaustA-Zone0 Temperature 36 degrees C / 96 degrees F
  ExhaustB-Zone1 Temperature 32 degrees C / 89 degrees F
  TCBC-Zone0 Temperature     34 degrees C / 93 degrees F
  Power 1
    1.0 V                     1015 mV
    1.2 V                     1205 mV
    1.8 V                     1804 mV
    2.5 V                     2552 mV
    3.3 V                     3325 mV
    5.0 V                     5020 mV
    5.0 V RE                  4988 mV
    12.0 V                   12104 mV
    12.0 V RE                12026 mV
  Bus Revision               100
  FPGA Revision              270
CB 1 status:
  State                      Online
  IntakeA-Zone0 Temperature  35 degrees C / 95 degrees F
  IntakeB-Zone1 Temperature  28 degrees C / 82 degrees F
  IntakeC-Zone0 Temperature  37 degrees C / 98 degrees F
  ExhaustA-Zone0 Temperature 34 degrees C / 93 degrees F
  ExhaustB-Zone1 Temperature 29 degrees C / 84 degrees F
  TCBC-Zone0 Temperature     32 degrees C / 89 degrees F
  Power 1
    1.0 V                     1011 mV
    1.2 V                     1208 mV
```

```
    1.8 V                      1788 mV
    2.5 V                      2526 mV
    3.3 V                      3319 mV
    5.0 V                      5046 mV
    5.0 V RE                   4975 mV
    12.0 V                    12046 mV
    12.0 V RE                 12007 mV
  Bus Revision                100
  FPGA Revision               0
```

show chassis environment cb (MX2008 Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                       Online Master
  Inlet1 Temperature          37 degrees C / 98 degrees F
  Inlet2 Temperature          45 degrees C / 113 degrees F
  Inlet3 Temperature          44 degrees C / 111 degrees F
  Inlet4 Temperature          42 degrees C / 107 degrees F
  Exhaust1 Temperature        30 degrees C / 86 degrees F
  Exhaust2 Temperature        40 degrees C / 104 degrees F
  Exhaust3 Temperature        48 degrees C / 118 degrees F
  Exhaust4 Temperature        46 degrees C / 114 degrees F
  Power 1
    1.0 V PHY                     989 mV
    1.15 V                       1150 mV
    1.2 V bias                   1189 mV
    1.5 V                        1488 mV
    1.8 V                        1772 mV
    2.5 V                        2462 mV
    3.3 V bias                   3296 mV
    VCCIO                        1028 mV
  Power 2
    1.1 V                        1099 mV
    3.3 V                        3300 mV
  Power 3
    0.95 V XL710                  949 mV
    1.05 V                       1050 mV
  Power 4
    1.2 V                        1199 mV
    5.0 V                        4999 mV
```

```
  Power 5
    1.0 V                       1000 mV
    1.2 V PHY                   1199 mV
  Bus Revision               114
  FPGA Revision              1
CB 1 status:
  State                      Online Standby
  Inlet1 Temperature         30 degrees C / 86 degrees F
  Inlet2 Temperature         31 degrees C / 87 degrees F
  Inlet3 Temperature         29 degrees C / 84 degrees F
  Inlet4 Temperature         32 degrees C / 89 degrees F
  Exhaust1 Temperature       30 degrees C / 86 degrees F
  Exhaust2 Temperature       33 degrees C / 91 degrees F
  Exhaust3 Temperature       34 degrees C / 93 degrees F
  Exhaust4 Temperature       35 degrees C / 95 degrees F
  Power 1
    1.0 V PHY                    986 mV
    1.15 V                      1153 mV
    1.2 V bias                  1195 mV
    1.5 V                       1498 mV
    1.8 V                       1798 mV
    2.5 V                       2494 mV
    3.3 V bias                  3296 mV
    VCCIO                       1034 mV
  Power 2
    1.1 V                       1100 mV
    3.3 V                       3300 mV
  Power 3
    0.95 V XL710                 949 mV
    1.05 V                      1050 mV
  Power 4
    1.2 V                       1199 mV
    5.0 V                       5000 mV
  Power 5
    1.0 V                       1000 mV
    1.2 V PHY                   1199 mV
  Bus Revision               114
  FPGA Revision              1
```

show chassis environment cb (MX10003 and MX10004 Router)

```
user@host> show chassis environment cb

CB 0 status:
  State                      Online Master
  CB 0 Exhaust Temp Sensor 0x49 35 degrees C / 95 degrees F
  CB 0 Inlet Temp Sensor 0x49 28 degrees C / 82 degrees F
  Power
    VDD1V5_PCH                  1489 mV
    VDDIO                        940 mV
    VDD3V3_PCH                  3332 mV
    VDD2V5_AB                   2508 mV
    VDD1V8_CLC                  1764 mV
    VDD3V3                      3292 mV
    VDD2V5_CD                   2508 mV
    VDD1V2_CBC_GTX                 0 mV
    VDD1V8_GLS_GTX                 0 mV
    VDD1V2_CBC                     0 mV
    VDD1V8_GLS                     0 mV
    BIAS3V3_BP                     0 mV
    VDD1V2_GH                   1200 mV
    VDD3V3_CBC                  3299 mV
    VDD1V2_CD                   1199 mV
    BIAS3V3                     3340 mV
    VDD1V2_AB                   1200 mV
    VDD5V0                      5000 mV
    VDD1V05                     1049 mV
    VDD1V05                     1050 mV
    VCORE                       1780 mV
         12V                        12272 mV    3952 mA   48984 mW
CB 1 status:
  State                      Online Standby
  CB 1 Exhaust Temp Sensor 0x49 35 degrees C / 95 degrees F
  CB 1 Inlet Temp Sensor 0x49 31 degrees C / 87 degrees F
  Power
    VDD1V5_PCH                  1489 mV
    VDDIO                        940 mV
    VDD3V3_PCH                  3351 mV
    VDD2V5_AB                   2508 mV
    VDD1V8_CLC                  1764 mV
    VDD3V3                      3312 mV
```

```
    VDD2V5_CD                   2508 mV
    VDD1V2_CBC_GTX              1195 mV
    VDD1V8_GLS_GTX              1764 mV
    VDD1V2_CBC                  1195 mV
    VDD1V8_GLS                  1783 mV
    BIAS3V3_BP                  4096 mV
    VDD1V2_GH                   1200 mV
    VDD3V3_CBC                  3300 mV
    VDD1V2_CD                   1200 mV
    BIAS3V3                     3339 mV
    VDD1V2_AB                   1200 mV
    VDD5V0                      5000 mV
    VDD1V05                     1050 mV
    VDD1V05                     1050 mV
    VCORE                       1780 mV
         12V                        12351 mV    3823 mA   45007 mW
```

show chassis environment cb (MX204 Router)

```
user@host> show chassis environment cb

CB 0 status:
  State                      Online Master
  CB 0 Top Right Inlet Sensor35 degrees C / 95 degrees F
  CB 0 Top Left Inlet Sensor 37 degrees C / 98 degrees F
  CB 0 Top Right Exhaust Sensor43 degrees C / 109 degrees F
  CB 0 Top Left Exhaust Sensor50 degrees C / 122 degrees F
  CB 0 CPU Core-0 Temp       48 degrees C / 118 degrees F
  CB 0 CPU Core-1 Temp       48 degrees C / 118 degrees F
  CB 0 CPU Core-2 Temp       48 degrees C / 118 degrees F
  CB 0 CPU Core-3 Temp       47 degrees C / 116 degrees F
  CB 0 CPU Core-4 Temp       47 degrees C / 116 degrees F
  CB 0 CPU Core-5 Temp       47 degrees C / 116 degrees F
  CB 0 CPU Core-6 Temp       47 degrees C / 116 degrees F
  CB 0 CPU Core-7 Temp       47 degrees C / 116 degrees F
  Power
    VDD1V5_PCH                  1509 mV
    VDDIO                        950 mV
    VDD3V3_PCH                  3312 mV
    VDD2V5_AB                   2508 mV
    VDD1V8_FRMR                 1813 mV
```

```
    VDD3V3                      3312 mV
    VDD2V5_CD                   2508 mV
    VDD1V8_PLL                  1813 mV
    VDD1V5                      1499 mV
    EA0_1V5                     1499 mV
    EA0_1V04                    1038 mV
    EA0_PLL_1V0                  999 mV
    EA0_2V5                     2508 mV
    BIAS3V                      3332 mV
    VDD1V2_CD                   1214 mV
    VDD1V2_AB                   1215 mV
    VDD1V05                     1050 mV
    BIAS3V3                     3309 mV
    VDD1V0                      1015 mV
    VDD1V8                      1804 mV
    VDD1V2                      1199 mV
    VDD2V5                      2504 mV
    EA0_VDD0V9                   949 mV
    EA0_HM1_VDD0V9               899 mV
    EA0_VDD0V9R2                 952 mV
    EA0_VDD1V0                  1000 mV
    VDD3V3                      3304 mV
    EA0_XR_VDD1V2               1199 mV
    EA0_XR_VDD0V9                903 mV
    EA0_HM_VDDM1V2              1199 mV
    EA0_HM_VDD1V2               1199 mV
    VDDCPU0                     1770 mV
         12V Hotswap A              11968 mV    4696 mA   55466 mW
         12V Hotswap B              12048 mV   14936 mA  180652 mW
```

show chassis environment cb (MX10008 Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
  CB 0 Intake A Temp Sensor               24 degrees C / 75 degrees F
  CB 0 Intake B Temp Sensor               24 degrees C / 75 degrees F
  CB 0 Exhaust A Temp Sensor              28 degrees C / 82 degrees F
  CB 0 Exhaust B Temp Sensor              30 degrees C / 86 degrees F
  CB 0 Middle Temp Sensor                 28 degrees C / 82 degrees F
  Power
```

```
    GESW_VDD1V0                 1000 mV
    VDD1V0                      1000 mV
    VDD1V2                      1199 mV
    VDD3V3                      3299 mV
    XL710_VCCD                   950 mV
    VDD1V05                     1050 mV
    VDD2V5                      2500 mV
    FPGA_VDD1V2                 1200 mV
    VDD1V8                      1800 mV
    VDD1V15                     1150 mV
    VDD1V1                      1099 mV
    VCCIO                        950 mV
    PHY_VDD1V0                  1000 mV
    VDD5V0                      4998 mV
    FPGA_VDD1V5                 1496 mV
    VDD1V5                      1496 mV
         12V                        12281 mV    7700 mA   92400 mW
CB 1 status:
  State                      Online Standby
  CB 1 Intake A Temp Sensor               24 degrees C / 75 degrees F
  CB 1 Intake B Temp Sensor               23 degrees C / 73 degrees F
  CB 1 Exhaust A Temp Sensor              27 degrees C / 80 degrees F
  CB 1 Exhaust B Temp Sensor              30 degrees C / 86 degrees F
  CB 1 Middle Temp Sensor                 28 degrees C / 82 degrees F
  Power
    GESW_VDD1V0                  999 mV
    VDD1V0                      1000 mV
    VDD1V2                      1199 mV
    VDD3V3                      3299 mV
    XL710_VCCD                   950 mV
    VDD1V05                     1050 mV
    VDD2V5                      2499 mV
    FPGA_VDD1V2                 1200 mV
    VDD1V8                      1799 mV
    VDD1V15                     1150 mV
    VDD1V1                      1100 mV
    VCCIO                        949 mV
    PHY_VDD1V0                   999 mV
    VDD5V0                      5000 mV
    FPGA_VDD1V5                 1502 mV
    VDD1V5                      1496 mV
         12V                        12281 mV    8002 mA   96024 mW
```

show chassis environment cb (T4000 Core Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
  Temperature                33 degrees C / 91 degrees F
  Power 1
    1.8 V                     1805 mV
    2.5 V                     2523 mV
    3.3 V                     3324 mV
    3.3 V bias                3296 mV
    4.6 V                     4680 mV
    5.0 V                     4893 mV
    8.0 V bias                7572 mV
    12.0 V                   11916 mV
  Power 2
    1.0 V                      993 mV
    1.2 V                     1210 mV
    3.3 V RE                  3330 mV
  Bus Revision               51
  FPGA Revision              5
CB 1 status:
  State                      Online Standby
  Temperature                33 degrees C / 91 degrees F
  Power 1
    1.8 V                     1810 mV
    2.5 V                     2496 mV
    3.3 V                     3308 mV
    3.3 V bias                3286 mV
    4.6 V                     4692 mV
    5.0 V                     4954 mV
    8.0 V bias                7282 mV
    12.0 V                   11926 mV
  Power 2
    1.0 V                      993 mV
    1.2 V                     1185 mV
    3.3 V RE                  3316 mV
  Bus Revision               51
  FPGA Revision              5
```

show chassis environment cb (TX Matrix Router)

```
user@host> show chassis environment cb
---------------------------------------------------------------------------
CB 0 status:
  State                            Online Master
  Temperature                   32 degrees C / 89 degrees F
  Power:
    1.8 V                     1797 mV
    2.5 V                     2477 mV
    3.3 V                     3311 mV
    4.6 V                     4727 mV
    5.0 V                     5015 mV
    12.0 V                   12185 mV
    3.3 V bias                3304 mV
    8.0 V bias                7870 mV
  BUS Revision                  40
  FPGA Revision                  1
CB 1 status:
  State                            Online Standby
...

lcc0-re0:
---------------------------------------------------------------------------
CB 0 status:
  State                            Online Master
  Temperature                   32 degrees C / 89 degrees F
  Power:
    1.8 V                     1787 mV
    2.5 V                     2473 mV
    3.3 V                     3306 mV
    4.6 V                     4793 mV
    5.0 V                     5025 mV
    12.0 V                   12156 mV
    3.3 V bias                3289 mV
    8.0 V bias                7609 mV
  BUS Revision                  40
  FPGA Revision                  5
CB 1 status:
  State                            Online Standby
 ....
  BUS Revision                  40
```

```
  FPGA Revision                 5

lcc2-re0:
--------------------------------------------------------------------------
CB 0 status:
  State                           Online Master
 ...
CB 1 status:
  State                           Online Standby
 ...
```

show chassis environment cb (TX Matrix Plus Router)

```
user@host> show chassis environment cb
sfc0-re0:
--------------------------------------------------------------------------
CB 0 status:
  State                       Online Master
  Temperature                 38 degrees C / 100 degrees F
  Power 1
    1.0 V                      1005 mV
    1.1 V                      1108 mV
    1.2 V                      1205 mV
    1.25 V                     1269 mV
    1.5 V                      1508 mV
    1.8 V                      1814 mV
    2.5 V                      2507 mV
    3.3 V                      3306 mV
    3.3 V bias                 3300 mV
    9.0 V                      9058 mV
    9.0 V RE                   9107 mV
  Power 2
    3.9 V                      3963 mV
    5.0 V                      5020 mV
    9.0 V                      9087 mV
  Bus Revision                79
  FPGA Revision               23
CB 1 status:
  State                       Online Standby
  Temperature                 39 degrees C / 102 degrees F
  Power 1
```

```
    1.0 V                     1002 mV
    1.1 V                     1105 mV
    1.2 V                     1198 mV
    1.25 V                    1276 mV
    1.5 V                     1504 mV
    1.8 V                     1804 mV
    2.5 V                     2507 mV
    3.3 V                     3300 mV
    3.3 V bias                3293 mV
    9.0 V                     9039 mV
    9.0 V RE                  9049 mV
  Power 2
    3.9 V                     3892 mV
    5.0 V                     5040 mV
    9.0 V                     9058 mV
  Bus Revision               79
  FPGA Revision              23

lcc0-re0:
--------------------------------------------------------------------------
CB 0 status:
  State                      Online Master
  Temperature                39 degrees C / 102 degrees F
  Power 1
    1.8 V                     1799 mV
    2.5 V                     2499 mV
    3.3 V                     3327 mV
    3.3 V bias                3299 mV
    4.6 V                     4673 mV
    5.0 V                     4918 mV
    8.0 V bias                7308 mV
    12.0 V                   11887 mV
  Power 2
    1.0 V                      996 mV
    1.2 V                     1199 mV
    3.3 V RE                  3319 mV
  Bus Revision               51
  FPGA Revision              3
CB 1 status:
  State                      Online Standby
  Temperature                40 degrees C / 104 degrees F
  Power 1
    1.8 V                     1800 mV
```

```
    2.5 V                     2496 mV
    3.3 V                     3322 mV
    3.3 V bias                3284 mV
    4.6 V                     4680 mV
    5.0 V                     4954 mV
    8.0 V bias                7284 mV
    12.0 V                   11902 mV
  Power 2
    1.0 V                      998 mV
    1.2 V                     1205 mV
    3.3 V RE                  3327 mV
  Bus Revision               51
  FPGA Revision              3

lcc1-re0:
--------------------------------------------------------------------------
CB 0 status:
  State                      Online Master
  Temperature                41 degrees C / 105 degrees F
  Power 1
    1.8 V                     1804 mV
    2.5 V                     2517 mV
    3.3 V                     3300 mV
    3.3 V bias                3284 mV
    4.6 V                     4681 mV
    5.0 V                     4927 mV
    8.0 V bias                7357 mV
    12.0 V                   11907 mV
  Power 2
    1.0 V                      991 mV
    1.2 V                     1202 mV
    3.3 V RE                  3301 mV
  Bus Revision               51
  FPGA Revision              3
CB 1 status:
  State                      Online Standby
  Temperature                40 degrees C / 104 degrees F
  Power 1
    1.8 V                     1805 mV
    2.5 V                     2528 mV
    3.3 V                     3324 mV
    3.3 V bias                3289 mV
    4.6 V                     4694 mV
```

```
    5.0 V                      4959 mV
    8.0 V bias                 7311 mV
    12.0 V                    11926 mV
  Power 2
    1.0 V                       998 mV
    1.2 V                      1200 mV
    3.3 V RE                   3313 mV
  Bus Revision                51
  FPGA Revision               3

lcc2-re0:
--------------------------------------------------------------------------
CB 0 status:
  State                       Online Master
  Temperature                 41 degrees C / 105 degrees F
  Power 1
    1.8 V                      1805 mV
    2.5 V                      2494 mV
    3.3 V                      3333 mV
    3.3 V bias                 3296 mV
    4.6 V                      4673 mV
    5.0 V                      4901 mV
    8.0 V bias                 7343 mV
    12.0 V                    11916 mV
  Power 2
    1.0 V                       993 mV
    1.2 V                      1213 mV
    3.3 V RE                   3328 mV
  Bus Revision                51
  FPGA Revision               3
CB 1 status:
  State                       Online Standby
  Temperature                 41 degrees C / 105 degrees F
  Power 1
    1.8 V                      1804 mV
    2.5 V                      2523 mV
    3.3 V                      3334 mV
    3.3 V bias                 3291 mV
    4.6 V                      4697 mV
    5.0 V                      4969 mV
    8.0 V bias                 7308 mV
    12.0 V                    11936 mV
  Power 2
```

```
    1.0 V                      996 mV
    1.2 V                     1200 mV
    3.3 V RE                  3328 mV
  Bus Revision               51
  FPGA Revision              3

lcc3-re0:
--------------------------------------------------------------------------
CB 0 status:
  State                      Online Master
  Temperature                37 degrees C / 98 degrees F
  Power 1
    1.8 V                     1809 mV
    2.5 V                     2510 mV
    3.3 V                     3296 mV
    3.3 V bias                3291 mV
    4.6 V                     4670 mV
    5.0 V                     4905 mV
    8.0 V bias                7211 mV
    12.0 V                   11882 mV
  Power 2
    1.0 V                      996 mV
    1.2 V                     1188 mV
    3.3 V RE                  3326 mV
  Bus Revision               51
  FPGA Revision              5
CB 1 status:
  State                      Online Standby
  Temperature                38 degrees C / 100 degrees F
  Power 1
    1.8 V                     1813 mV
    2.5 V                     2510 mV
    3.3 V                     3322 mV
    3.3 V bias                3289 mV
    4.6 V                     4692 mV
    5.0 V                     4967 mV
    8.0 V bias                7194 mV
    12.0 V                   11916 mV
  Power 2
    1.0 V                      996 mV
    1.2 V                     1205 mV
    3.3 V RE                  3273 mV
```

```
  Bus Revision                51
  FPGA Revision               5
```

show chassis environment cb (EX8200 Switch)

```
user@host> show chassis environment cb

CB 0 status:
  State                       Online Master
  Temperature Intake          20 degrees C / 68 degrees F
  Temperature Exhaust         24 degrees C / 75 degrees F
  Power 1
    1.1 V                      1086 mV
    1.2 V                      1179 mV
    1.2 V *                    1182 mV
    1.2 V *                    1182 mV
    1.25 V                     1211 mV
    1.5 V                      1472 mV
    1.8 V                      1756 mV
    2.5 V                      2449 mV
    3.3 V                      3254 mV
    3.3 V bias                 3300 mV
    5.0 V                      4911 mV
    12.0 V                    11891 mV
  Power 2
    3.3 V bias *               3615 mV
    3.3 V bias *               3615 mV
    3.3 V bias *               3567 mV
    3.3 V bias *               3664 mV
    4.3 V bias *               4224 mV
    4.3 V bias *               4215 mV
    4.3 V bias *               4224 mV
    4.3 V bias *               4205 mV
    4.3 V bias *               4195 mV
    4.3 V bias *               4215 mV
    5.0 V bias                 4920 mV
CB 1 status:
  State                       Online Standby
  Temperature Intake          19 degrees C / 66 degrees F
  Temperature Exhaust         23 degrees C / 73 degrees F
  Power 1
```

```
    1.1 V                      1082 mV
    1.2 V                      1169 mV
    1.2 V *                    1179 mV
    1.2 V *                    1179 mV
    1.25 V                     1214 mV
    1.5 V                      1482 mV
    1.8 V                      1759 mV
    2.5 V                      2481 mV
    3.3 V                      3248 mV
    3.3 V bias                 3306 mV
    5.0 V                      4911 mV
    12.0 V                    11910 mV
  Power 2
    3.3 V bias *               3644 mV
    3.3 V bias *               3664 mV
    3.3 V bias *               3586 mV
    3.3 V bias *               3654 mV
    4.3 V bias *               4224 mV
    4.3 V bias *               4215 mV
    4.3 V bias *               4224 mV
    4.3 V bias *               4205 mV
    4.3 V bias *               4244 mV
    4.3 V bias *               4215 mV
    5.0 V bias                 4930 mV
CB 2 status:
  State                       Online
  Temperature Intake          19 degrees C / 66 degrees F
  Temperature Exhaust         23 degrees C / 73 degrees F
  Power 1
    1.2 V                      1195 mV
    1.5 V                      1511 mV
    1.8 V                      1804 mV
    2.5 V                      2526 mV
    3.3 V                      3300 mV
    3.3 V bias                 3306 mV
    12.0 V                    12220 mV
```

show chassis environment cb (EX8208 Switch)

```
user@host> show chassis environment cb
CB 0 status:
```

```
  State                       Online Master
  Temperature Intake          20 degrees C / 68 degrees F
  Temperature Exhaust         24 degrees C / 75 degrees F
  Power 1
    1.1 V                      1086 mV
    1.2 V                      1179 mV
    1.2 V *                    1182 mV
    1.2 V *                    1182 mV
    1.25 V                     1211 mV
    1.5 V                      1466 mV
    1.8 V                      1759 mV
    2.5 V                      2455 mV
    3.3 V                      3261 mV
    3.3 V bias                 3300 mV
    5.0 V                      4930 mV
    12.0 V                    11891 mV
  Power 2
    3.3 V bias *               3606 mV
    3.3 V bias *               3615 mV
    3.3 V bias *               3567 mV
    3.3 V bias *               3673 mV
    4.3 V bias *               4224 mV
    4.3 V bias *               4215 mV
    4.3 V bias *               4234 mV
    4.3 V bias *               4205 mV
    4.3 V bias *               4186 mV
    4.3 V bias *               4215 mV
    5.0 V bias                 4940 mV
CB 1 status:
  State                       Online Standby
  Temperature Intake          19 degrees C / 66 degrees F
  Temperature Exhaust         23 degrees C / 73 degrees F
  Power 1
    1.1 V                      1086 mV
    1.2 V                      1169 mV
    1.2 V *                    1179 mV
    1.2 V *                    1179 mV
    1.25 V                     1211 mV
    1.5 V                      1479 mV
    1.8 V                      1759 mV
    2.5 V                      2475 mV
    3.3 V                      3235 mV
    3.3 V bias                 3306 mV
```

```
    5.0 V                      4930 mV
    12.0 V                    11891 mV
  Power 2
    3.3 V bias *               3644 mV
    3.3 V bias *               3664 mV
    3.3 V bias *               3586 mV
    3.3 V bias *               3654 mV
    4.3 V bias *               4215 mV
    4.3 V bias *               4224 mV
    4.3 V bias *               4215 mV
    4.3 V bias *               4215 mV
    4.3 V bias *               4234 mV
    4.3 V bias *               4224 mV
    5.0 V bias                 4920 mV
CB 2 status:
  State                       Online
  Temperature Intake          20 degrees C / 68 degrees F
  Temperature Exhaust         24 degrees C / 75 degrees F
  Power 1
    1.2 V                      1202 mV
    1.5 V                      1508 mV
    1.8 V                      1804 mV
    2.5 V                      2520 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    12.0 V                    12200 mV
```

show chassis environment cb (EX9251 Switch)

```
user@switch> show chassis environment cb
CB 0 status:
  State                       Online Master
  CB 0 Top Right Inlet Sensor29 degrees C / 84 degrees F
  CB 0 Top Left Inlet Sensor 28 degrees C / 82 degrees F
  CB 0 Top Right Exhaust Sensor40 degrees C / 104 degrees F
  CB 0 Top Left Exhaust Sensor59 degrees C / 138 degrees F
  CB 0 CPU Core-0 Temp        45 degrees C / 113 degrees F
  CB 0 CPU Core-1 Temp        44 degrees C / 111 degrees F
  CB 0 CPU Core-2 Temp        44 degrees C / 111 degrees F
  CB 0 CPU Core-3 Temp        44 degrees C / 111 degrees F
  CB 0 CPU Core-4 Temp        45 degrees C / 113 degrees F
```

```
  CB 0 CPU Core-5 Temp       44 degrees C / 111 degrees F
  CB 0 CPU Core-6 Temp       44 degrees C / 111 degrees F
  CB 0 CPU Core-7 Temp       43 degrees C / 109 degrees F
  Power
    VDD1V5_PCH                  1499 mV
    VDDIO                        950 mV
    VDD3V3_PCH                  3312 mV
    VDD2V5_AB                   2489 mV
    VDD1V8_FRMR                 1793 mV
    VDD3V3                      3292 mV
    VDD2V5_CD                   2508 mV
    VDD1V8_PLL                  1793 mV
    VDD1V5                      1499 mV
    EA0_1V5                     1499 mV
    EA0_1V04                     999 mV
    EA0_PLL_1V0                  999 mV
    EA0_2V5                     2508 mV
    BIAS3V                      3292 mV
    VDD1V2_CD                   1215 mV
    VDD1V2_AB                   1214 mV
    VDD1V05                     1050 mV
    BIAS3V3                     3309 mV
    VDD1V0                      1014 mV
    VDD1V8                      1805 mV
    VDD1V2                      1200 mV
    VDD2V5                      2504 mV
    EA0_VDD0V9                   949 mV
    EA0_HM1_VDD0V9               899 mV
    EA0_VDD0V9R2                 952 mV
    EA0_VDD1V0                   999 mV
    VDD3V3                      3305 mV
    EA0_XR_VDD1V2               1199 mV
    EA0_XR_VDD0V9                903 mV
    EA0_HM_VDDM1V2              1199 mV
    EA0_HM_VDD1V2               1199 mV
    VDDCPU0                     1770 mV
         12V Hotswap A              11955 mV    4861 mA   59347 mW
         12V Hotswap B              11916 mV   15046 mA  180887 mW
```

show chassis environment cb (EX9253 Switch)

```
user@switch> show chassis environment cb
CB 0 status:
  State                      Online Master
  CB 0 Exhaust Temp Sensor   38 degrees C / 100 degrees F
  CB 0 Inlet Temp Sensor     32 degrees C / 89 degrees F
  CB 0 CPU DIE Temp Sensor   43 degrees C / 109 degrees F
  Power
    VDD1V5_PCH                  1489 mV
    VDDIO                        940 mV
    VDD3V3_PCH                  3332 mV
    VDD2V5_AB                   2508 mV
    VDD1V8_CLC                  1783 mV
    VDD3V3                      3312 mV
    VDD2V5_CD                   2508 mV
    VDD1V2_CBC_GTX              1195 mV
    VDD1V8_GLS_GTX              1783 mV
    VDD1V2_CBC                  1176 mV
    VDD1V8_GLS                  1783 mV
    BIAS3V3_BP                  3978 mV
    VDD1V2_GH                   1200 mV
    VDD3V3_CBC                  3299 mV
    VDD1V2_CD                   1200 mV
    BIAS3V3                     3340 mV
    VDD1V2_AB                   1199 mV
    VDD5V0                      5000 mV
    VDD1V05                     1050 mV
    VDD1V05                     1050 mV
    VCORE                       1770 mV
         12V                        12061 mV    4806 mA   57841 mW
CB 1 status:
  State                      Offline
  CB 1 Exhaust Temp Sensor   32 degrees C / 89 degrees F
  CB 1 Inlet Temp Sensor     29 degrees C / 84 degrees F
  CB 1 CPU DIE Temp Sensor   43 degrees C / 109 degrees F
  Power                      Disabled
```

show chassis environment cb (PTX5000 Packet Transport Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
  Intake Temperature         38 degrees C / 100 degrees F
  Exhaust A Temperature      45 degrees C / 113 degrees F
  Exhaust B Temperature      42 degrees C / 107 degrees F
  Power 1
    1.2 V                     1200 mV
    1.25 V                    1250 mV
    2.5 V                     2500 mV
    3.3 V                     3300 mV
  Power 2
    1.0 V                     1000 mV
    3.3 V bias                3293 mV
    3.9 V                     3921 mV
  Bus Revision               132
  FPGA Revision              27
CB 1 status:
  State                      Online Standby
  Intake Temperature         34 degrees C / 93 degrees F
  Exhaust A Temperature      39 degrees C / 102 degrees F
  Exhaust B Temperature      36 degrees C / 96 degrees F
  Power 1
    1.2 V                     1199 mV
    1.25 V                    1250 mV
    2.5 V                     2499 mV
    3.3 V                     3299 mV
  Power 2
    1.0 V                     1000 mV
    3.3 V bias                3312 mV
    3.9 V                     3961 mV
  Bus Revision               132
  FPGA Revision              28
```

show chassis environment cb (PTX10008 Router)

```
user@host> show chassis environment cb
CB 0 status:
```

```
  State                       Online Master
  CB 0 Intake Temp Sensor    28 degrees C / 82 degrees F
  CB 0 Exhaust Temp Sensor   32 degrees C / 89 degrees F
  Power
    VDD 2.5V                    2489 mV
    Bias 3.3V                   3332 mV
    VDD 3.3V                    3292 mV
    VCC 1.8V                    1822 mV
    VDD 1.2V                    1205 mV
    VCC 1V                       999 mV
    VCC CPU 1.8V                1803 mV
    VDD 2.5V                    2489 mV
    VCC Aux 5V                  5115 mV
    VDD DDR 1.5V                1499 mV
    VTT SA CPU 0.8V              803 mV
    VTT CPU 1.05V               1048 mV
    VCC Core CPU                 901 mV
    VCC PCH 1.5V                1519 mV
    VDD 1.05V                   1048 mV
    VCC 2.5V                    2508 mV
    FORT VCCA 1V                 960 mV
    VDD .85V                     862 mV
    VTT DDRA .75V                744 mV
    VTT DDRB .75V                744 mV
         12V                        12285 mV    3779 mA   46339 mW
CB 1 status:
  State                       Online Standby
  CB 1 Intake Temp Sensor    27 degrees C / 80 degrees F
  CB 1 Exhaust Temp Sensor   32 degrees C / 89 degrees F
  Power
    VDD 2.5V                    2489 mV
    Bias 3.3V                   3332 mV
    VDD 3.3V                    3273 mV
    VCC 1.8V                    1822 mV
    VDD 1.2V                    1195 mV
    VCC 1V                       999 mV
    VCC CPU 1.8V                1783 mV
    VDD 2.5V                    2489 mV
    VCC Aux 5V                  5056 mV
    VDD DDR 1.5V                1499 mV
    VTT SA CPU 0.8V              793 mV
    VTT CPU 1.05V               1048 mV
    VCC Core CPU                 882 mV
```

```
    VCC PCH 1.5V                1509 mV
    VDD 1.05V                   1048 mV
    VCC 2.5V                    2489 mV
    FORT VCCA 1V                 960 mV
    VDD .85V                     862 mV
    VTT DDRA .75V                744 mV
    VTT DDRB .75V                744 mV
         12V                        12391 mV    3779 mA   46727 mW
```

show chassis environment cb (PTX10016 Router)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
  CB 0 Intake Temp Sensor    20 degrees C / 68 degrees F
  CB 0 Exhaust Temp Sensor   24 degrees C / 75 degrees F
  Power
    VDD 2.5V                    2508 mV
    Bias 3.3V                   3351 mV
    VDD 3.3V                    3292 mV
    VCC 1.8V                    1832 mV
    VDD 1.2V                    1205 mV
    VCC 1V                       999 mV
    VCC CPU 1.8V                1793 mV
    VDD 2.5V                    2508 mV
    VCC Aux 5V                  5056 mV
    VDD DDR 1.5V                1509 mV
    VTT SA CPU 0.8V              803 mV
    VTT CPU 1.05V               1048 mV
    VCC Core CPU                 960 mV
    VCC PCH 1.5V                1519 mV
    VDD 1.05V                   1058 mV
    VCC 2.5V                    2528 mV
    FORT VCCA 1V                 960 mV
    VDD .85V                     852 mV
    VTT DDRA .75V                744 mV
    VTT DDRB .75V                744 mV
         12V                        12259 mV    3649 mA   45173 mW
CB 1 status:
  State                      Online Standby
  CB 1 Intake Temp Sensor    20 degrees C / 68 degrees F
```

```
  CB 1 Exhaust Temp Sensor   23 degrees C / 73 degrees F
  Power
    VDD 2.5V                    2508 mV
    Bias 3.3V                   3312 mV
    VDD 3.3V                    3273 mV
    VCC 1.8V                    1822 mV
    VDD 1.2V                    1195 mV
    VCC 1V                       989 mV
    VCC CPU 1.8V                1783 mV
    VDD 2.5V                    2489 mV
    VCC Aux 5V                  5086 mV
    VDD DDR 1.5V                1499 mV
    VTT SA CPU 0.8V              803 mV
    VTT CPU 1.05V               1048 mV
    VCC Core CPU                1029 mV
    VCC PCH 1.5V                1519 mV
    VDD 1.05V                   1048 mV
    VCC 2.5V                    2528 mV
    FORT VCCA 1V                 960 mV
    VDD .85V                     862 mV
    VTT DDRA .75V                744 mV
    VTT DDRB .75V                744 mV
         12V                        12285 mV    3952 mA   48447 mW
```

show chassis environment cb (QFabric System and QFX5700)

```
user@switch> show chassis environment cb
CB 0 status:
State Online Master
Temp Sensor Inlet 1 52 degrees C / 125 degrees F
Temp Sensor Inlet 2 83 degrees C / 181 degrees F
Temp Sensor Exhaust 1 67 degrees C / 152 degrees F
Temp Sensor Exhaust 2 62 degrees C / 143 degrees F
Power 1
Volt Sensor PVTT0 1494 mV
Volt Sensor PVTT1 1504 mV
Volt Sensor VDD1V8_FPGA 1345 mV
Volt Sensor BIAS3V3 3130 mV
Volt Sensor VDD5V0 4995 mV
Volt Sensor VDD3V3_SATA 3230 mV
Volt Sensor VDD3V3_FPGA 3130 mV
```

```
Volt Sensor VDD3V3_PCH 3130 mV
Volt Sensor VDD2V5 2415 mV
Volt Sensor VDD0V9 910 mV

Volt Sensor VDD1V2_FPGA 1211 mV
Volt Sensor VDD1V8 1791 mV
Power 2
Volt Sensor VDD1V5 1494 mV
Volt Sensor VDD1V7 1704 mV
Volt Sensor VDD1V3 1345 mV
Volt Sensor VDD1V0 1002 mV
…
…
Bus Revision 99
FPGA Revision 10
CB 1 status:
State Online Master
Temp Sensor Inlet 1 52 degrees C / 125 degrees F
Temp Sensor Inlet 2 83 degrees C / 181 degrees F
Temp Sensor Exhaust 1 67 degrees C / 152 degrees F
Temp Sensor Exhaust 2 62 degrees C / 143 degrees F
Power 1
Volt Sensor PVTT0 1494 mV
Volt Sensor PVTT1 1504 mV
Volt Sensor VDD1V8_FPGA 1345 mV
Volt Sensor BIAS3V3 3130 mV
Volt Sensor VDD5V0 4995 mV
Volt Sensor VDD3V3_SATA 3230 mV
Volt Sensor VDD3V3_FPGA 3130 mV
Volt Sensor VDD3V3_PCH 3130 mV
Volt Sensor VDD2V5 2415 mV
Volt Sensor VDD0V9 910 mV
Volt Sensor VDD1V2_FPGA 1211 mV
Volt Sensor VDD1V8 1791 mV
Power 2
Volt Sensor VDD1V5 1494 mV
Volt Sensor VDD1V7 1704 mV
Volt Sensor VDD1V3 1345 mV
Volt Sensor VDD1V0 1002 mV
…
…
```

```
Bus Revision 99
FPGA Revision 10
```

```
user@switch> show chassis environment cb interconnect-device
IC-123 0
CB 0 status:
  State                      Online Master
  Left Intake Temperature    33 degrees C / 91 degrees F
  Right Intake Temperature   33 degrees C / 91 degrees F
  Left Exhaust Temperature   36 degrees C / 96 degrees F
  Right Exhaust Temperature  35 degrees C / 95 degrees F
  Power                      OK
    VDD 3V3                   3294 mV
    VDD 2V5                   2436 mV
    VDD 1V8                   1746 mV
    VDD 1V5                   1460 mV
    VDD 1V25                  1210 mV
    VDD 1V2                   1164 mV
    CPU CORE 1V2              1120 mV
    VDD 1V0                    968 mV
    VDD 5V0                   5088 mV
    CPU MP BIAS 4V3           4050 mV
    BIAS 3V3                  3180 mV
    VTT 0V9                    866 mV
```

show chassis environment cb 0 (PTX10003-80C and PTX10003-160C Routers)

```
user@router> show chassis environment cb 0

CB 0 status:
  State                      Online Master
  Power 1
    PCIe SW0 power, CORE 0.9 V     894 mV
  Power 2
    CLOCK MAIN, 2.5 V           2495 mV
    PCIe SW0 power, SERDES 0.9 V     897 mV
  Power 3
    PCIe SW1 power, CORE 0.9 V     897 mV
  Power 4
    PCIe SW1 power, SERDES 0.9 V     894 mV
```

```
  Power 5
    Ethernet SW0 power, 1.0 V     998 mV
  Power 6
    Ethernet SW1 power, CORE 1.0 V     998 mV
  Power 7
    12 V                       12034 mV
  Bus Revision               0
  FPGA Revision              0
```

Starting in Junos OS Evolved Release 19.1R1, on PTX10003-80C and PTX10003-160C devices, the `show chassis environment cb` command does not show the Bus and FPGA revision information. If you want to view the FPGA revision or version information for the CB, use the `show system firmware` command.

show chassis environment cb (ACX7100, ACX7509, ACX7348, ACX7332)

```
user@host> show chassis environment cb
CB 0 status:
  State                      Online Master
  PHY Die 0 Sensor 0 Temp                48 degrees C / 118 degrees F
  Temp Inlet Left                        34 degrees C / 93 degrees F
  Temp Inlet Right                       39 degrees C / 102 degrees F
  Temp Exhaust Left                      40 degrees C / 104 degrees F
  Temp Exhaust Right                     42 degrees C / 107 degrees F
  DC-DC Controller 0 Temp                48 degrees C / 118 degrees F
  DC-DC Controller 1 Temp                45 degrees C / 113 degrees F
  Hotswap Temp                           38 degrees C / 100 degrees F
  PEX 8714 Temp Sensor                   50 degrees C / 122 degrees F
  Power 1
    VDD3V3                      3315 mV
  Power 2
    VDD0V9                       901 mV
  Power 3
    VDD0V8                       802 mV
  Power 4
    BIAS3V3                     3296 mV
    PHY_AVDD1V2                 1200 mV
    VDD1V0                       996 mV
    VDD1V2_FPGA                 1202 mV
    VDD1V8                      1802 mV
    VDD1V8_FPGA                 1804 mV
    VDD2V5                      2504 mV
```

```
    VDD3V3_FPGA                 3336 mV
    VDD3V7                      3736 mV
    VDD5V0                      5058 mV
  Bus Revision               38
  FPGA Revision              0xa
```

Release Information

Command introduced before Junos Release 7.4.

option introduced for the TX Matrix Plus router in Junos Release 9.6.

Command introduced in Junos OS Release 18.2R1 for MX10008 Universal Routing Platforms, and EX9253 switches.

RELATED DOCUMENTATION

request chassis cb | 14441

Understanding Switching Control Board Redundancy

show chassis environment ccg

IN THIS SECTION

- Syntax | 17614
- Description | 17614
- Options | 17614
- Required Privilege Level | 17614
- Output Fields | 17614
- Sample Output | 17615
- Release Information | 17615

Syntax

```
show chassis environment ccg
<slot>
```

Description

(PTX5000 Packet Transport Routers only) Display environmental information about the Centralized Clock Generators (CCGs).

Options

none Display environmental information about all CCGs on the router.

slot (Optional) Display environmental information about the specified CCG. Replace *slot* with 0 or 1.

Required Privilege Level

view

Output Fields

Table 225 on page 17614 lists the output fields for the `show chassis environment ccg` command. Output fields are listed in the approximate order in which they appear.

Table 225: show chassis environment cb Output Fields

Field Name	Field Description
`State`	Status of the CCG: `Online - Master clock` , `Online - Standby`, or `Offline`. If two CCGs are installed and online, one is functioning as the primary clock, and the other is the standby clock.
`Temperature`	Temperature of the air flowing past the CCG.
`Power`	Power required and measured on the CCG. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.

Table 225: show chassis environment cb Output Fields *(Continued)*

Field Name	Field Description
BUS Revision	Revision level of the generic bus device.

Sample Output

show chassis environment ccg (PTX5000)

```
user@host> show chassis environment ccg
CCG 0 status:
  State                      Online -  Master clock
  Temperature                31 degrees C / 87 degrees F
  Power
    1.2 V bias                1200 mV
    1.8 V                     1799 mV
    3.3 V                     3300 mV
    3.3 V bias                3300 mV
  Bus Revision               103
CCG 1 status:
  State                      Offline
  Power                      Disabled
  Temperature                31 degrees C / 87 degrees F
  Power
    1.2 V bias                1198 mV
    1.8 V                      161 mV
    3.3 V                      451 mV
    3.3 V bias                3311 mV
  Bus Revision               103
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Clock Sources for PTX Series Packet Transport Routers

show chassis environment | 17429

show chassis environment fan

IN THIS SECTION

- Syntax | 17616
- Description | 17616
- Options | 17616
- Required Privilege Level | 17617
- Output Fields | 17617
- Sample Output | 17617
- Release Information | 17621

Syntax

```
show chassis environment fan
<all-members>
<fantray-slot-number>
<local>
<member member-id>
```

Description

(MX2020 and MX2010 routers only) Display environmental information about the fans and fan trays.

Options

none	Display environmental information about all fans and fan trays.
all-members	(Optional) Display environmental information about the fan and fan trays in all members of the Virtual Chassis configuration.

fantray-slot-number	(Optional) Display environmental information about the specified fan tray. Replace *fantray-slot-number* with a value from `0` through `3`.
local	(Optional) Display environmental information about the fans and fan trays in the local member of the Virtual Chassis.
member ***member-id***	(Optional) Display environmental information about the fans and fan trays in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 226 on page 17617 lists the output fields for the `show chassis environment fan` command. Output fields are listed in the approximate order in which they appear.

Table 226: show chassis environment fan Output Fields

Field Name	Field Description
`Fan Tray Status`	Status of the fan tray.
`Temperature`	Temperature in Celsius (C) and Fahrenheit (F) maintained by the fans.
`Power`	Power required and measured on the fan tray. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
`Max Power Consumption`	Maximum power consumed by the fan tray.

Sample Output

show chassis environment fan (MX2020 routers)

```
user@host> show chassis environment fan
Fan Tray 0 status:
```

```
  Temperature                23 degrees C / 73 degrees F
  Power
    1.2 V                     1189 mV
    3.3 V                     3293 mV
    5.0 V                     5230 mV
    5.0 V bias                5278 mV
    52.0 V A                 49944 mV
    52.0 V B                   435 mV
  Max Power Consumption       1150 Watts
Fan Tray 1 status:
  Temperature                22 degrees C / 71 degrees F
  Power
    1.2 V                     1192 mV
    3.3 V                     3300 mV
    5.0 V                     5230 mV
    5.0 V bias                5230 mV
    52.0 V A                 50205 mV
    52.0 V B                   435 mV
  Max Power Consumption       1150 Watts
Fan Tray 2 status:
  Temperature                31 degrees C / 87 degrees F
  Power
    1.2 V                     1192 mV
    3.3 V                     3287 mV
    5.0 V                     5211 mV
    5.0 V bias                5220 mV
    52.0 V A                 50031 mV
    52.0 V B                   435 mV
  Max Power Consumption       1150 Watts
Fan Tray 3 status:
  Temperature                31 degrees C / 87 degrees F
  Power
    1.2 V                     1208 mV
    3.3 V                     3306 mV
    5.0 V                     5240 mV
    5.0 V bias                5259 mV
    52.0 V A                 50553 mV
    52.0 V B                   435 mV
  Max Power Consumption       1150 Watts
```

show chassis environment fan 3 (MX2020 routers)

```
user@host> show chassis environment fan 3
Fan Tray 3 status:
  Temperature                31 degrees C / 87 degrees F
  Power
    1.2 V                     1208 mV
    3.3 V                     3306 mV
    5.0 V                     5240 mV
    5.0 V bias                5259 mV
    52.0 V A                 50553 mV
    52.0 V B                   435 mV
  Max Power Consumption       1150 Watts
```

show chassis environment fan (PTX10008 routers)

```
user@host> show chassis environment fan

Fan Tray 0  status:
         HS 0                      12338 mV    4035 mA   40623 mW
         HS 1                      12325 mV    3044 mA   36483 mW
         HS 2                      12272 mV    3374 mA   38388 mW
         HS 3                      12364 mV    2218 mA   37918 mW
  Temperature                            28 degrees C / 82 degrees
F

Fan Tray 1  status:
         HS 0                      12232 mV    3209 mA   50455 mW
         HS 1                      12311 mV    3760 mA   50196 mW
         HS 2                      12311 mV    5356 mA   47397 mW
         HS 3                      12259 mV    3264 mA   15807 mW
  Temperature                            28 degrees C / 82 degrees F
```

show chassis environment fan (PTX10016 routers)

```
user@host> show chassis environment fan

Aug 02 21:13:00
Fan Tray 0  status:
```

```
     HS 0                        12364 mV    4200 mA   48926 mW
     HS 1                        12285 mV    3264 mA   69885 mW
     HS 2                        12285 mV    4365 mA   57559 mW
     HS 3                        12259 mV    2714 mA   67768 mW
  Temperature                             25 degrees C / 77 degrees F
Fan Tray 1  status:
     HS 0                        12325 mV    1888 mA   76354 mW
     HS 1                        12325 mV    5026 mA   70779 mW
     HS 2                        12285 mV    4200 mA   50032 mW
     HS 3                        12311 mV    5246 mA   55630 mW
  Temperature                             24 degrees C / 75 degrees F
```

show chassis environment fan (MX10004 and MX10008 routers)

```
user@host> show chassis environment fan

Fan Tray 0  status:
         HS 0                        12298 mV    4200 mA   59676 mW
         HS 1                        12338 mV    4585 mA   48856 mW
         HS 2                        12325 mV    4530 mA   63863 mW
         HS 3                        12311 mV    2879 mA   37800 mW
  Temperature                             26 degrees C / 78 degrees F
Fan Tray 1  status:
         HS 0                        12298 mV    4145 mA   57559 mW
         HS 1                        12298 mV    2934 mA   60288 mW
         HS 2                        12325 mV    5026 mA   46715 mW
         HS 3                        12272 mV    1833 mA   47209 mW
  Temperature                             25 degrees C / 77 degrees F
```

show chassis environment fan (ACX7509, ACX7348, ACX7332)

```
user@host> show chassis environment fan
Fan Tray 0 status:
  Temp Hotswap Controller 0               37 degrees C / 98 degrees F
    Hotswap Controller 0       12031 mV
Fan Tray 1 status:
  Temp Hotswap Controller 1               36 degrees C / 96 degrees F
    Hotswap Controller 1       12062 mV
Fan Tray 2 status:
```

```
  Temp Hotswap Controller 2               38 degrees C / 100 degrees F
    Hotswap Controller 2       12093 mV
```

Release Information

Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 Universal Routing Platforms.

`all-members`, `local`, and `member` *`member-id`* options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.

RELATED DOCUMENTATION

show chassis environment | **17429**

show chassis fan | **18260**

show chassis environment fpc

IN THIS SECTION

- Syntax | **17622**
- Syntax (TX Matrix and TX Matrix Plus Routers) | **17622**
- Syntax (MX Series Routers) | **17622**
- Syntax (MX2010, MX10003, MX204, MX2008, and MX10008, OCX Series, PTX3000, PTX10008 devices and Junos OS Evolved platforms) | **17622**
- Syntax (MX2020 Universal Routing Platforms) | **17622**
- Syntax (QFX Series) | **17623**
- Description | **17623**
- Options | **17623**
- Required Privilege Level | **17625**
- Output Fields | **17625**
- Sample Output | **17627**
- Release Information | **17704**

Syntax

```
show chassis environment fpc
<slot>
```

Syntax (TX Matrix and TX Matrix Plus Routers)

```
show chassis environment fpc
<lcc number>
<slot>
```

Syntax (MX Series Routers)

```
show chassis environment fpc
<slot>
<all-members>
<local>
<member member-id>
```

Syntax (MX2010, MX10003, MX204, MX2008, and MX10008, OCX Series, PTX3000, PTX10008 devices and Junos OS Evolved platforms)

```
show chassis environment fpc
<slot>
```

Syntax (MX2020 Universal Routing Platforms)

```
show chassis environment fpc
<slot>
<satellite [fpc-slot slot-id |device-alias  alias-name]
```

Syntax (QFX Series)

```
show chassis environment fpc
<fpc-slot>
interconnect-device name
```

Description

(M40e, M120, M160, M320, MX Series, T Series routers, EX Series, QFX Series, and PTX Series routers only) Display environmental information about Flexible PIC Concentrators (FPCs).

Options

none

Display environmental information about all FPCs. On a TX Matrix router, display environmental information about all FPCs on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about all FPCs on the TX Matrix Plus router and its attached routers.

all-members

(MX Series routers only) (Optional) Display environmental information for the FPCs in all the members of the Virtual Chassis configuration.

interconnect-device *name*

(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.

lcc *number*

(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(MX Series routers only) (Optional) Display environmental information for the FPCs in the local Virtual Chassis member.

member *member-id*

(MX Series routers only) (Optional) Display environmental information for the FPCs in the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

satellite [fpc-slot *slot-id* | device-alias *alias-name*]

(Junos Fusion only)(Optional) Display environmental information for the FPCs in the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

slot* or *fpc-slot

(Optional) Display environmental information about an individual FPC:

- (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, if you specify the number of the T640 router by using only the `lcc` *number* option (the recommended method), replace *slot* with a value from `0` through `7`. Similarly, on a TX Matrix Plus router, if you specify the number of the router by using only the `lcc` *number* option (the recommended method), replace *slot* with a value from `0` through `7`. Otherwise, replace *slot* with a value from `0` through `31`. For example, the following commands have the same result:

```
user@host> show chassis environment fpc 1 lcc 1
user@host> show chassis environment fpc 9
```

- M120 router—Replace *slot* with a value from `0` through `5`.
- MX240 router—Replace *slot* with a value from `0` through `2`.
- MX480 router—Replace *slot* with a value from `0` through `5`.
- MX960 router—Replace *slot* with a value from `0` through `11`.
- MX2010 router—Replace *slot* with a value from `0` through `9`.
- MX2020 router—Replace *slot* with a value from `0` through `19`.
- MX2008 router—Replace *slot* with a value from `0` through `9`.
- Other routers—Replace *slot* with a value from `0` through `7`.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—Replace *slot* with 0.

    - EX4200 switches in a Virtual Chassis configuration—Replace *slot* with a value from 0 through 9 (switch's member ID).
    - EX6210 switches–Replace *slot* with a value from 0 through 3 (line card only), 4 or 5 (line card or Switch Fabric and Rotuing Engine (SRE) module), or 6 through 9 (line card only).
    - EX8208 switches—Replace *slot* with a value from 0 through 7 (line card).
    - EX8216 switches—Replace *slot* with a value from 0 through 15 (line card).
  - QFX3500 switches —Replace *fpc-slot* with 0 through 15.
  - PTX5000 Packet Transport Router—Replace *fpc-slot* with 0 through 7.
  - PTX3000 Packet Transport Router—Replace *fpc-slot* with 0 through 15.

Required Privilege Level

view

Output Fields

Table 227 on page 17626 lists the output fields for the `show chassis environment fpc` command. Output fields are listed in the approximate order in which they appear.

Table 227: show chassis environment fpc Output Fields

Field Name	Field Description
`State`	Status of the FPC: • `Unknown`—FPC is not detected by the router. • `Empty`—No FPC is present. • `Present`—FPC is detected by the chassis daemon but is either not supported by the current version of the Junos OS, or the FPC is coming up but not yet online. • `Ready`—FPC is in intermediate or transition state. • `Announce online`—Intermediate state during which the FPC is coming up but not yet online, and the chassis manager acknowledges the chassisd FPC online initiative. • `Online`—FPC is online and running. • `Offline`—FPC is powered down. • `Diagnostics`—FPC is set to operate in diagnostics mode.
`Temperature`	(M40e and M160 routers and QFX Series only) Temperature of the air flowing past the FPC.
`PMB Temperature`	(PTX Series only) Temperature of the air flowing past the PMB (bottom of the FPC). The PTX5000 Packet Transport Router with FPC2-PTX-P1A include multiple temperatures for PMB (`TEMP0` and `TEMP1`).
`PMB CPU Temperature`	(PTX5000 Packet Transport Router with FPC2-PTX-P1A only) Temperature of the air flowing past the PMB CPU.
`Temperature Intake`	(M320 routers, MX2010 routers, MX2020 routers, MX2008 routers, and PTX Series only) Temperature of the air flowing into the chassis.
`Temperature Top`	(T Series routers only) Temperature of the air flowing past the top of the FPC.

Table 227: show chassis environment fpc Output Fields *(Continued)*

Field Name	Field Description
`Temperature Exhaust`	(M120 and M320 routers, MX2010 routers, MX2020 routers, MX2008 routers, and PTX Series only) Temperature of the air flowing out of the chassis. The PTX Series Packet Transport Routers, and the MX2010, MX2020, and MX2008 routers include exhaust temperatures for multiple zones (`Exhaust A` and `Exhaust B`).
`Temperature Bottom`	(T Series routers only) Temperature of the air flowing past the bottom of the FPC.
`TL` *n* `Temperature`	(PTX Series only) Temperature of the air flowing past the specified TL area of the packet forwarding engine (PFE) on the FPC.
`TQ` *n* `Temperature`	(PTX Series only) Temperature of the air flowing past the specified TQ area of the packet forwarding engine (PFE) on the FPC.
`Temperature MMB0`	(T640 router only) Temperature of the air flowing past the type 3 FPC.
`Temperature MMB1`	(M320 and T Series routers only) Temperature of the air flowing past the type 1, type 2, and type 3 FPC.
`Power`	Information about the voltage supplied to the FPC. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
`CMB Revision` or `BUS revision`	Revision level of the chassis management bus device (M Series router) or bus (T Series routers).

Sample Output

show chassis environment fpc (M120 Router)

```
user@host> show chassis environment fpc
FPC 2 status:
  State                      Online
  Temperature Exhaust A      32 degrees C / 89 degrees F
  Temperature Exhaust B      31 degrees C / 87 degrees F
```

```
  Power A-Board
    1.2 V                     1202 mV
    1.5 V                     1508 mV
    1.8 V                     1798 mV
    2.5 V                     2507 mV
    3.3 V                     3351 mV
    5.0 V                     4995 mV
    3.3 V bias                3296 mV
    1.2 V Rocket IO           1205 mV
    1.5 V Rocket IO           1501 mV
  I2C Slave Revision         12
FPC 3 status:
  State                      Online
  Temperature Exhaust A      31 degrees C / 87 degrees F
  Temperature Exhaust B      33 degrees C / 91 degrees F
  Power A-Board
    1.2 V                     1211 mV
    1.5 V                     1501 mV
    1.8 V                     1798 mV
    2.5 V                     2471 mV
    3.3 V                     3293 mV
    5.0 V                     4930 mV
    3.3 V bias                3296 mV
    1.2 V Rocket IO           1205 mV
    1.5 V Rocket IO           1501 mV
  Power B-Board
    1.2 V                     1214 mV
    1.5 V                     1501 mV
    2.5 V                     2471 mV
    3.3 V                     3300 mV
    5.0 V                     4943 mV
    3.3 V bias                3296 mV
    1.2 V Rocket IO           1205 mV
    1.5 V Rocket IO           1501 mV
  I2C Slave Revision         12
FPC 4 status:
  State                      Online
  Temperature Exhaust A      32 degrees C / 89 degrees F
  Temperature Exhaust B      30 degrees C / 86 degrees F
  Power A-Board
    1.2 V                     1195 mV
    1.5 V                     1504 mV
    1.8 V                     1801 mV
```

```
    2.5 V                     2504 mV
    3.3 V                     3293 mV
    5.0 V                     4917 mV
    3.3 V bias                3296 mV
    1.2 V Rocket IO           1202 mV
    1.5 V Rocket IO           1492 mV
  I2C Slave Revision         12
```

show chassis environment fpc (M160 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                       Online
  Temperature              42 degrees C / 107 degrees F
  Power:
    1.5 V                1500 mV
    2.5 V                2509 mV
    3.3 V                3308 mV
    5.0 V                4991 mV
    5.0 V bias           4952 mV
    8.0 V bias           8307 mV
  CMB Revision             12
FPC 1 status:
  State                       Online
  Temperature              45 degrees C / 113 degrees F
  Power:
    1.5 V                1498 mV
    2.5 V                2501 mV
    3.3 V                3319 mV
    5.0 V                5020 mV
    5.0 V bias           5025 mV
    8.0 V bias           8307 mV
  CMB Revision             12
```

show chassis environment fpc (M320 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                           Online
  Temperature Intake           27 degrees C / 80 degrees F
```

```
  Temperature Exhaust          38 degrees C / 100 degrees F
  Temperature MMB1             31 degrees C / 87 degrees F
  Power:
    1.5 V                    1487 mV
    1.5 V *                  1494 mV
    1.8 V                    1821 mV
    2.5 V                    2533 mV
    3.3 V                    3323 mV
    5.0 V                    5028 mV
    3.3 V bias               3296 mV
    5.0 V bias               4984 mV
  CMB Revision                 16
FPC 1 status:
  State                           Online
  Temperature Intake           27 degrees C / 80 degrees F
  Temperature Exhaust          37 degrees C / 98 degrees F
  Temperature MMB1             32 degrees C / 89 degrees F
  Power:
    1.5 V                    1504 mV
    1.5 V *                  1499 mV
    1.8 V                    1820 mV
    2.5 V                    2529 mV
    3.3 V                    3328 mV
    5.0 V                    5013 mV
    3.3 V bias               3294 mV
    5.0 V bias               4984 mV
  CMB Revision                 16
FPC 2 status:
  State                           Online
  Temperature Intake           28 degrees C / 82 degrees F
  Temperature Exhaust          38 degrees C / 100 degrees F
  Temperature MMB1             32 degrees C / 89 degrees F
  Power:
    1.5 V                    1498 mV
    1.5 V *                  1487 mV
    1.8 V                    1816 mV
    2.5 V                    2531 mV
    3.3 V                    3324 mV
    5.0 V                    5025 mV
    3.3 V bias               3277 mV
    5.0 V bias               5013 mV
  CMB Revision                 17
```

```
FPC 3 status:
...
```

show chassis environment fpc (MX2020 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                      Online
  Temperature Intake         41 degrees C / 105 degrees F
  Temperature Exhaust A      48 degrees C / 118 degrees F
  Temperature Exhaust B      60 degrees C / 140 degrees F
  Temperature LU 0 TSen      56 degrees C / 132 degrees F
  Temperature LU 0 Chip      59 degrees C / 138 degrees F
  Temperature LU 1 TSen      56 degrees C / 132 degrees F
  Temperature LU 1 Chip      61 degrees C / 141 degrees F
  Temperature LU 2 TSen      56 degrees C / 132 degrees F
  Temperature LU 2 Chip      52 degrees C / 125 degrees F
  Temperature LU 3 TSen      56 degrees C / 132 degrees F
  Temperature LU 3 Chip      52 degrees C / 125 degrees F
  Temperature MQ 0 TSen      49 degrees C / 120 degrees F
  Temperature MQ 0 Chip      49 degrees C / 120 degrees F
  Temperature MQ 1 TSen      49 degrees C / 120 degrees F
  Temperature MQ 1 Chip      52 degrees C / 125 degrees F
  Temperature MQ 2 TSen      49 degrees C / 120 degrees F
  Temperature MQ 2 Chip      45 degrees C / 113 degrees F
  Temperature MQ 3 TSen      49 degrees C / 120 degrees F
  Temperature MQ 3 Chip      46 degrees C / 114 degrees F
  Power
    AS-BIAS3V3-zl2105         3299 mV
    AS-VDD1V8-zl2006          1807 mV
    AS-VDD2V5-zl2006          2512 mV
    AS-AVDD1V0-zl2004          997 mV
    AS-PCIE_1V0-zl2004         996 mV
    AS-VDD3V3-zl2004          3294 mV
    AS-VDD_1V5A-zl2004        1501 mV
    AS-VDD_1V5B-zl2004        1498 mV
    AS-LU0_1V0-zl2004          998 mV
    AS-LU1_1V0-zl2004         1002 mV
    AS-MQ0_1V0-zl2004          999 mV
    AS-MQ1_1V0-zl2004          994 mV
    AS-LU2_1V0-zl2004         1000 mV
```

```
    AS-LU3_1V0-zl2004          998 mV
    AS-MQ2_1V0-zl2004         1002 mV
    AS-MQ3_1V0-zl2004          999 mV
    AS-PMB_1V1-zl2006         1096 mV
  I2C Slave Revision         68
FPC 1 status:
  State                      Online
  Temperature Intake         39 degrees C / 102 degrees F
  Temperature Exhaust A      48 degrees C / 118 degrees F
  Temperature Exhaust B      55 degrees C / 131 degrees F
  Temperature LU 0 TSen      52 degrees C / 125 degrees F
  Temperature LU 0 Chip      54 degrees C / 129 degrees F
  Temperature LU 1 TSen      52 degrees C / 125 degrees F
  Temperature LU 1 Chip      56 degrees C / 132 degrees F
  Temperature LU 2 TSen      52 degrees C / 125 degrees F
  Temperature LU 2 Chip      49 degrees C / 120 degrees F
  Temperature LU 3 TSen      52 degrees C / 125 degrees F
  Temperature LU 3 Chip      50 degrees C / 122 degrees F
  Temperature MQ 0 TSen      48 degrees C / 118 degrees F
  Temperature MQ 0 Chip      48 degrees C / 118 degrees F
  Temperature MQ 1 TSen      48 degrees C / 118 degrees F
  Temperature MQ 1 Chip      51 degrees C / 123 degrees F
  Temperature MQ 2 TSen      48 degrees C / 118 degrees F
  Temperature MQ 2 Chip      45 degrees C / 113 degrees F
  Temperature MQ 3 TSen      48 degrees C / 118 degrees F
  Temperature MQ 3 Chip      45 degrees C / 113 degrees F
  Power
    AS-BIAS3V3-zl2105         3291 mV
    AS-VDD1V8-zl2006          1786 mV
    AS-VDD2V5-zl2006          2496 mV
    AS-AVDD1V0-zl2004         1000 mV
    AS-PCIE_1V0-zl2004        1000 mV
    AS-VDD3V3-zl2004          3294 mV
    AS-VDD_1V5A-zl2004        1500 mV
    AS-VDD_1V5B-zl2004        1498 mV
    AS-LU0_1V0-zl2004         1003 mV
    AS-LU1_1V0-zl2004         1000 mV
    AS-MQ0_1V0-zl2004         1000 mV
    AS-MQ1_1V0-zl2004          995 mV
    AS-LU2_1V0-zl2004         1002 mV
    AS-LU3_1V0-zl2004          997 mV
    AS-MQ2_1V0-zl2004         1000 mV
    AS-MQ3_1V0-zl2004          998 mV
```

```
    AS-PMB_1V1-zl2006         1096 mV
  I2C Slave Revision         68
FPC 2 status:
  State                      Online
  Temperature Intake         39 degrees C / 102 degrees F
  Temperature Exhaust A      48 degrees C / 118 degrees F
  Temperature Exhaust B      58 degrees C / 136 degrees F
  Temperature LU 0 TSen      55 degrees C / 131 degrees F
  Temperature LU 0 Chip      57 degrees C / 134 degrees F
  Temperature LU 1 TSen      55 degrees C / 131 degrees F
  Temperature LU 1 Chip      63 degrees C / 145 degrees F
  Temperature LU 2 TSen      55 degrees C / 131 degrees F
  Temperature LU 2 Chip      51 degrees C / 123 degrees F
  Temperature LU 3 TSen      55 degrees C / 131 degrees F
  Temperature LU 3 Chip      52 degrees C / 125 degrees F
  Temperature MQ 0 TSen      48 degrees C / 118 degrees F
  Temperature MQ 0 Chip      50 degrees C / 122 degrees F
  Temperature MQ 1 TSen      48 degrees C / 118 degrees F
  Temperature MQ 1 Chip      52 degrees C / 125 degrees F
  Temperature MQ 2 TSen      48 degrees C / 118 degrees F
  Temperature MQ 2 Chip      47 degrees C / 116 degrees F
  Temperature MQ 3 TSen      48 degrees C / 118 degrees F
  Temperature MQ 3 Chip      47 degrees C / 116 degrees F
  Power
    AS-BIAS3V3-zl2105         3299 mV
    AS-VDD1V8-zl2006          1805 mV
    AS-VDD2V5-zl2006          2510 mV
    AS-AVDD1V0-zl2004          999 mV
    AS-PCIE_1V0-zl2004         998 mV
    AS-VDD3V3-zl2004          3296 mV
    AS-VDD_1V5A-zl2004        1492 mV
    AS-VDD_1V5B-zl2004        1497 mV
    AS-LU0_1V0-zl2004          997 mV
    AS-LU1_1V0-zl2004         1000 mV
    AS-MQ0_1V0-zl2004          998 mV
    AS-MQ1_1V0-zl2004         1001 mV
    AS-LU2_1V0-zl2004          996 mV
    AS-LU3_1V0-zl2004          995 mV
    AS-MQ2_1V0-zl2004          998 mV
    AS-MQ3_1V0-zl2004          997 mV
    AS-PMB_1V1-zl2006         1100 mV
  I2C Slave Revision         68
FPC 3 status:
```

```
  State                      Online
  Temperature Intake         41 degrees C / 105 degrees F
  Temperature Exhaust A      48 degrees C / 118 degrees F
  Temperature Exhaust B      58 degrees C / 136 degrees F
  Temperature LU 0 TSen      56 degrees C / 132 degrees F
  Temperature LU 0 Chip      59 degrees C / 138 degrees F
  Temperature LU 1 TSen      56 degrees C / 132 degrees F
  Temperature LU 1 Chip      61 degrees C / 141 degrees F
  Temperature LU 2 TSen      56 degrees C / 132 degrees F
  Temperature LU 2 Chip      51 degrees C / 123 degrees F
  Temperature LU 3 TSen      56 degrees C / 132 degrees F
  Temperature LU 3 Chip      53 degrees C / 127 degrees F
  Temperature MQ 0 TSen      50 degrees C / 122 degrees F
  Temperature MQ 0 Chip      51 degrees C / 123 degrees F
  Temperature MQ 1 TSen      50 degrees C / 122 degrees F
  Temperature MQ 1 Chip      55 degrees C / 131 degrees F
  Temperature MQ 2 TSen      50 degrees C / 122 degrees F
  Temperature MQ 2 Chip      47 degrees C / 116 degrees F
  Temperature MQ 3 TSen      50 degrees C / 122 degrees F
  Temperature MQ 3 Chip      50 degrees C / 122 degrees F
  Power
    AS-BIAS3V3-zl2105         3305 mV
    AS-VDD1V8-zl2006          1810 mV
    AS-VDD2V5-zl2006          2508 mV
    AS-AVDD1V0-zl2004          999 mV
    AS-PCIE_1V0-zl2004        1001 mV
    AS-VDD3V3-zl2004          3294 mV
    AS-VDD_1V5A-zl2004        1500 mV
    AS-VDD_1V5B-zl2004        1498 mV
    AS-LU0_1V0-zl2004          998 mV
    AS-LU1_1V0-zl2004          998 mV
    AS-MQ0_1V0-zl2004          999 mV
    AS-MQ1_1V0-zl2004          998 mV
    AS-LU2_1V0-zl2004         1000 mV
    AS-LU3_1V0-zl2004         1001 mV
    AS-MQ2_1V0-zl2004          996 mV
    AS-MQ3_1V0-zl2004          998 mV
    AS-PMB_1V1-zl2006         1098 mV
  I2C Slave Revision         68
FPC 4 status:
...
```

show chassis environment fpc (MX2010 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                      Online
  Temperature Intake         36 degrees C / 96 degrees F
  Temperature Exhaust A      42 degrees C / 107 degrees F
  Temperature Exhaust B      51 degrees C / 123 degrees F
  Temperature LU 0 TSen      49 degrees C / 120 degrees F
  Temperature LU 0 Chip      50 degrees C / 122 degrees F
  Temperature LU 1 TSen      49 degrees C / 120 degrees F
  Temperature LU 1 Chip      54 degrees C / 129 degrees F
  Temperature LU 2 TSen      49 degrees C / 120 degrees F
  Temperature LU 2 Chip      45 degrees C / 113 degrees F
  Temperature LU 3 TSen      49 degrees C / 120 degrees F
  Temperature LU 3 Chip      46 degrees C / 114 degrees F
  Temperature MQ 0 TSen      40 degrees C / 104 degrees F
  Temperature MQ 0 Chip      41 degrees C / 105 degrees F
  Temperature MQ 1 TSen      40 degrees C / 104 degrees F
  Temperature MQ 1 Chip      44 degrees C / 111 degrees F
  Temperature MQ 2 TSen      40 degrees C / 104 degrees F
  Temperature MQ 2 Chip      38 degrees C / 100 degrees F
  Temperature MQ 3 TSen      40 degrees C / 104 degrees F
  Temperature MQ 3 Chip      41 degrees C / 105 degrees F
  Power
    AS-BIAS3V3-zl2105         3300 mV
    AS-VDD1V8-zl2006          1805 mV
    AS-VDD2V5-zl2006          2505 mV
    AS-AVDD1V0-zl2004          998 mV
    AS-PCIE_1V0-zl2004         999 mV
    AS-VDD3V3-zl2004          3303 mV
    AS-VDD_1V5A-zl2004        1497 mV
    AS-VDD_1V5B-zl2004        1497 mV
    AS-LU0_1V0-zl2004          998 mV
    AS-LU1_1V0-zl2004         1003 mV
    AS-MQ0_1V0-zl2004          998 mV
    AS-MQ1_1V0-zl2004          998 mV
    AS-LU2_1V0-zl2004          997 mV
    AS-LU3_1V0-zl2004         1001 mV
    AS-MQ2_1V0-zl2004          996 mV
    AS-MQ3_1V0-zl2004          994 mV
    AS-PMB_1V1-zl2006         1097 mV
```

```
  I2C Slave Revision         68
FPC 1 status:
  State                      Online
  Temperature Intake         34 degrees C / 93 degrees F
  Temperature Exhaust A      46 degrees C / 114 degrees F
  Temperature Exhaust B      54 degrees C / 129 degrees F
  Temperature LU 0 TSen      45 degrees C / 113 degrees F
  Temperature LU 0 Chip      55 degrees C / 131 degrees F
  Temperature LU 1 TSen      45 degrees C / 113 degrees F
  Temperature LU 1 Chip      44 degrees C / 111 degrees F
  Temperature LU 2 TSen      45 degrees C / 113 degrees F
  Temperature LU 2 Chip      50 degrees C / 122 degrees F
  Temperature LU 3 TSen      45 degrees C / 113 degrees F
  Temperature LU 3 Chip      58 degrees C / 136 degrees F
  Temperature XM 0 TSen      45 degrees C / 113 degrees F
  Temperature XM 0 Chip      51 degrees C / 123 degrees F
  Temperature XF 0 TSen      45 degrees C / 113 degrees F
  Temperature XF 0 Chip      63 degrees C / 145 degrees F
  Temperature PLX Switch TSen45 degrees C / 113 degrees F
  Temperature PLX Switch Chip47 degrees C / 116 degrees F
  Power
    MPC-BIAS3V3-zl2105        3300 mV
    MPC-VDD3V3-zl6100         3294 mV
    MPC-VDD2V5-zl6100         2505 mV
    MPC-VDD1V8-zl2004         1796 mV
    MPC-AVDD1V0-zl2004         991 mV
    MPC-VDD1V2-zl6100         1196 mV
    MPC-VDD1V5A-zl2004        1491 mV
    MPC-VDD1V5B-zl2004        1492 mV
    MPC-XF_0V9-zl2004          996 mV
    MPC-PCIE_1V0-zl6100       1003 mV
    MPC-LU0_1V0-zl2004         996 mV
    MPC-LU1_1V0-zl2004         996 mV
    MPC-LU2_1V0-zl2004         998 mV
    MPC-LU3_1V0-zl2004         994 mV
    MPC-12VA-BMR453          12031 mV
    MPC-12VB-BMR453          12003 mV
    MPC-PMB_1V1-zl2006        1104 mV
    MPC-PMB_1V2-zl2106        1194 mV
    MPC-XM_0V9-vt273m          911 mV
  I2C Slave Revision         110
FPC 8 status:
  State                      Online
```

```
  Temperature Intake         32 degrees C / 89 degrees F
  Temperature Exhaust A      44 degrees C / 111 degrees F
  Temperature Exhaust B      37 degrees C / 98 degrees F
  Temperature LU 0 TCAM TSen 41 degrees C / 105 degrees F
  Temperature LU 0 TCAM Chip 49 degrees C / 120 degrees F
  Temperature LU 0 TSen      41 degrees C / 105 degrees F
  Temperature LU 0 Chip      52 degrees C / 125 degrees F
  Temperature MQ 0 TSen      41 degrees C / 105 degrees F
  Temperature MQ 0 Chip      47 degrees C / 116 degrees F
  Temperature LU 1 TCAM TSen 39 degrees C / 102 degrees F
  Temperature LU 1 TCAM Chip 42 degrees C / 107 degrees F
  Temperature LU 1 TSen      39 degrees C / 102 degrees F
  Temperature LU 1 Chip      46 degrees C / 114 degrees F
  Temperature MQ 1 TSen      39 degrees C / 102 degrees F
  Temperature MQ 1 Chip      45 degrees C / 113 degrees F
  Power
    MPC-BIAS3V3-zl2105        3296 mV
    MPC-VDD3V3-zl2006         3298 mV
    MPC-VDD2V5-zl2006         2505 mV
    MPC-TCAM_1V0-zl2004        997 mV
    MPC-AVDD1V0-zl2006        1007 mV
    MPC-VDD1V8-zl2006         1803 mV
    MPC-PCIE_1V0-zl2006       1004 mV
    MPC-LU0_1V0-zl2004        1000 mV
    MPC-MQ0_1V0-zl2004         999 mV
    MPC-VDD_1V5-zl2004        1498 mV
    MPC-PMB_1V1-zl2006        1102 mV
    MPC-9VA-BMR453            9009 mV
    MPC-9VB-BMR453            8960 mV
    MPC-PMB_1V2-zl2105        1202 mV
    MPC-LU1_1V0-zl2004        1005 mV
    MPC-MQ1_1V0-zl2004        1000 mV
  I2C Slave Revision         70
FPC 9 status:
  State                      Online
  Temperature Intake         34 degrees C / 93 degrees F
  Temperature Exhaust A      41 degrees C / 105 degrees F
  Temperature Exhaust B      54 degrees C / 129 degrees F
  Temperature LU 0 TSen      51 degrees C / 123 degrees F
  Temperature LU 0 Chip      52 degrees C / 125 degrees F
  Temperature LU 1 TSen      51 degrees C / 123 degrees F
  Temperature LU 1 Chip      55 degrees C / 131 degrees F
  Temperature LU 2 TSen      51 degrees C / 123 degrees F
```

```
  Temperature LU 2 Chip      47 degrees C / 116 degrees F
  Temperature LU 3 TSen      51 degrees C / 123 degrees F
  Temperature LU 3 Chip      47 degrees C / 116 degrees F
  Temperature MQ 0 TSen      40 degrees C / 104 degrees F
  Temperature MQ 0 Chip      42 degrees C / 107 degrees F
  Temperature MQ 1 TSen      40 degrees C / 104 degrees F
  Temperature MQ 1 Chip      44 degrees C / 111 degrees F
  Temperature MQ 2 TSen      40 degrees C / 104 degrees F
  Temperature MQ 2 Chip      38 degrees C / 100 degrees F
  Temperature MQ 3 TSen      40 degrees C / 104 degrees F
  Temperature MQ 3 Chip      40 degrees C / 104 degrees F
  Power
    AS-BIAS3V3-zl2105         3302 mV
    AS-VDD1V8-zl2006          1808 mV
    AS-VDD2V5-zl2006          2513 mV
    AS-AVDD1V0-zl2004          997 mV
    AS-PCIE_1V0-zl2004         999 mV
    AS-VDD3V3-zl2004          3294 mV
    AS-VDD_1V5A-zl2004        1503 mV
    AS-VDD_1V5B-zl2004        1502 mV
    AS-LU0_1V0-zl2004          996 mV
    AS-LU1_1V0-zl2004          999 mV
    AS-MQ0_1V0-zl2004          997 mV
    AS-MQ1_1V0-zl2004          999 mV
    AS-LU2_1V0-zl2004          997 mV
    AS-LU3_1V0-zl2004          998 mV
    AS-MQ2_1V0-zl2004         1000 mV
    AS-MQ3_1V0-zl2004         1000 mV
    AS-PMB_1V1-zl2006         1102 mV
  I2C Slave Revision         68
```

show chassis environment fpc (MX2008 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                      Online
  Temperature Intake         29 degrees C / 84 degrees F
  Temperature Exhaust A      43 degrees C / 109 degrees F
  Temperature Exhaust B      42 degrees C / 107 degrees F
  Temperature XL 0 TSen      38 degrees C / 100 degrees F
  Temperature XL 0 Chip      53 degrees C / 127 degrees F
```

```
Temperature XL 0 XR2 0 TSen38 degrees C / 100 degrees F
Temperature XL 0 XR2 0 Chip60 degrees C / 140 degrees F
Temperature XL 0 XR2 1 TSen38 degrees C / 100 degrees F
Temperature XL 0 XR2 1 Chip60 degrees C / 140 degrees F
Temperature XL 1 TSen      30 degrees C / 86 degrees F
Temperature XL 1 Chip      43 degrees C / 109 degrees F
Temperature XL 1 XR2 0 TSen30 degrees C / 86 degrees F
Temperature XL 1 XR2 0 Chip50 degrees C / 122 degrees F
Temperature XL 1 XR2 1 TSen30 degrees C / 86 degrees F
Temperature XL 1 XR2 1 Chip50 degrees C / 122 degrees F
Temperature XM 0 TSen      42 degrees C / 107 degrees F
Temperature XM 0 Chip      49 degrees C / 120 degrees F
Temperature XM 1 TSen      42 degrees C / 107 degrees F
Temperature XM 1 Chip      42 degrees C / 107 degrees F
Temperature XM 2 TSen      42 degrees C / 107 degrees F
Temperature XM 2 Chip      42 degrees C / 107 degrees F
Temperature XM 3 TSen      42 degrees C / 107 degrees F
Temperature XM 3 Chip      40 degrees C / 104 degrees F
Temperature PCIe Switch TSen42 degrees C / 107 degrees F
Temperature PCIe Switch Chip22 degrees C / 71 degrees F
Power
  MPC-VDD_3V3-vt273m          3304 mV
  MPC-VDD_2V5-vt273m          2503 mV
  MPC-VDD_1V5-vt273m          1499 mV
  MPC-PCIE_0V9-vt273m          900 mV
  MPC-VDD_1V8-vt273m          1799 mV
  MPC-VDD_1V2-vt273m          1203 mV
  MPC-XM01_AVDD_1V0-vt273     1001 mV
  MPC-XM23_AVDD_1V0-vt273     1001 mV
  MPC-XM0_0V9-vt273m           900 mV
  MPC-XM1_0V9-vt273m           901 mV
  MPC-XM2_0V9-vt273m           903 mV
  MPC-XM3_0V9-vt273m           899 mV
  MPC-XL0_XR0_0V9-vt273m       899 mV
  MPC-XL0_XR1_0V9-vt273m       903 mV
  MPC-XL0_0V9-vt273m           899 mV
  MPC-XL0_AVDD_1V0-vt273m     1000 mV
  MPC-XL0_VDD_1V5-vt273m      1498 mV
  MPC-XL0_XR_1V2-vt273m       1200 mV
  MPC-XL1_XR0_0V9-vt273m       899 mV
  MPC-XL1_XR1_0V9-vt273m       899 mV
  MPC-XL1_0V9-vt273m           900 mV
  MPC-XL1_AVDD_1V0-vt273m     1000 mV
```

```
    MPC-XL1_VDD_1V5-vt273m      1501 mV
    MPC-XL1_XR_1V2-vt273m       1199 mV
    MPC-PMB-1V05-ltc2978        1049 mV
    MPC-PMB-1V5-ltc2978         1500 mV
    MPC-PMB-2V5-ltc2978         2500 mV
    MPC-PMB-3V3-ltc2978         3298 mV
  I2C Slave Revision         20
FPC 1 status:
  State                      Online
  Temperature Intake         29 degrees C / 84 degrees F
  Temperature Exhaust A      52 degrees C / 125 degrees F
  Temperature Exhaust B      44 degrees C / 111 degrees F
  Temperature EA0 TSen       55 degrees C / 131 degrees F
  Temperature EA0 Chip       48 degrees C / 118 degrees F
  Temperature EA0_XR0 TSen   55 degrees C / 131 degrees F
  Temperature EA0_XR0 Chip   57 degrees C / 134 degrees F
  Temperature EA0_XR1 TSen   55 degrees C / 131 degrees F
  Temperature EA0_XR1 Chip   54 degrees C / 129 degrees F
  Temperature EA1 TSen       55 degrees C / 131 degrees F
  Temperature EA1 Chip       50 degrees C / 122 degrees F
  Temperature EA1_XR0 TSen   55 degrees C / 131 degrees F
  Temperature EA1_XR0 Chip   59 degrees C / 138 degrees F
  Temperature EA1_XR1 TSen   55 degrees C / 131 degrees F
  Temperature EA1_XR1 Chip   59 degrees C / 138 degrees F
  Temperature PEX TSen       55 degrees C / 131 degrees F
  Temperature PEX Chip       39 degrees C / 102 degrees F
  Temperature EA2 TSen       43 degrees C / 109 degrees F
  Temperature EA2 Chip       39 degrees C / 102 degrees F
  Temperature EA2_XR0 TSen   43 degrees C / 109 degrees F
  Temperature EA2_XR0 Chip   45 degrees C / 113 degrees F
  Temperature EA2_XR1 TSen   43 degrees C / 109 degrees F
  Temperature EA2_XR1 Chip   43 degrees C / 109 degrees F
  Temperature EA3 TSen       43 degrees C / 109 degrees F
  Temperature EA3 Chip       41 degrees C / 105 degrees F
  Temperature EA3_XR0 TSen   43 degrees C / 109 degrees F
  Temperature EA3_XR0 Chip   50 degrees C / 122 degrees F
  Temperature EA3_XR1 TSen   43 degrees C / 109 degrees F
  Temperature EA3_XR1 Chip   46 degrees C / 114 degrees F
  Temperature EA0_HMC0 Logic die 61 degrees C / 141 degrees F
  Temperature EA0_HMC0 DRAM botm 58 degrees C / 136 degrees F
  Temperature EA0_HMC1 Logic die 62 degrees C / 143 degrees F
  Temperature EA0_HMC1 DRAM botm 59 degrees C / 138 degrees F
  Temperature EA0_HMC2 Logic die 59 degrees C / 138 degrees F
```

```
  Temperature EA0_HMC2 DRAM botm 56 degrees C / 132 degrees F
  Temperature EA1_HMC0 Logic die 67 degrees C / 152 degrees F
  Temperature EA1_HMC0 DRAM botm 64 degrees C / 147 degrees F
  Temperature EA1_HMC1 Logic die 65 degrees C / 149 degrees F
  Temperature EA1_HMC1 DRAM botm 62 degrees C / 143 degrees F
  Temperature EA1_HMC2 Logic die 63 degrees C / 145 degrees F
  Temperature EA1_HMC2 DRAM botm 60 degrees C / 140 degrees F
  Temperature EA2_HMC0 Logic die 51 degrees C / 123 degrees F
  Temperature EA2_HMC0 DRAM botm 48 degrees C / 118 degrees F
  Temperature EA2_HMC1 Logic die 55 degrees C / 131 degrees F
  Temperature EA2_HMC1 DRAM botm 52 degrees C / 125 degrees F
  Temperature EA2_HMC2 Logic die 52 degrees C / 125 degrees F
  Temperature EA2_HMC2 DRAM botm 49 degrees C / 120 degrees F
  Temperature EA3_HMC0 Logic die 51 degrees C / 123 degrees F
  Temperature EA3_HMC0 DRAM botm 48 degrees C / 118 degrees F
  Temperature EA3_HMC1 Logic die 52 degrees C / 125 degrees F
  Temperature EA3_HMC1 DRAM botm 49 degrees C / 120 degrees F
  Temperature EA3_HMC2 Logic die 52 degrees C / 125 degrees F
  Temperature EA3_HMC2 DRAM botm 49 degrees C / 120 degrees F
  Power
    MPC-EA0_0V9-vt1527mb         950 mV
    MPC-EA1_0V9-vt1527mb         950 mV
    MPC-EA2_0V9-vt1527mb         925 mV
    MPC-EA3_0V9-vt1527mb         924 mV
    MAX20751-1V0                1020 mV
    MAX20731-0V9                 891 mV
    MAX20751-EA0-AVDD1V0        1000 mV
    MAX20731-EA0-1V2            1189 mV
    MAX20731-EA0-HMC-1V2        1182 mV
    MAX20731-EA0-0V906           899 mV
    MAX20731-EA0-HMC-0V9         891 mV
    MAX20751-EA1-AVDD1V0        1000 mV
    MAX20731-EA1-1V2            1189 mV
    MAX20731-EA1-HMC-1V2        1182 mV
    MAX20731-EA1-0V906           899 mV
    MAX20731-EA1-HMC-0V9         889 mV
    MAX20751-EA2-AVDD1V0        1000 mV
    MAX20731-EA2-1V2            1186 mV
    MAX20731-EA2-HMC-1V2        1193 mV
    MAX20731-EA2-0V906           899 mV
    MAX20731-EA2-HMC-0V9         889 mV
    MAX20751-EA3-AVDD1V0        1000 mV
    MAX20731-EA3-1V2            1186 mV
```

```
    MAX20731-EA3-HMC-1V2        1193 mV
    MAX20731-EA3-0V906           897 mV
    MAX20731-EA3-HMC-0V9         894 mV
    MAX20731-3V3                3268 mV
    UCD9090_0-CH_1-EA0_PLL_     1010 mV
    UCD9090_0-CH_2-EA0_1V04     1038 mV
    UCD9090_0-CH_3-EA0_2V5      2499 mV
    UCD9090_0-CH_4-EA0_1V5      1494 mV
    UCD9090_0-CH_5-EA1_PLL_     1012 mV
    UCD9090_0-CH_6-EA1_1V04     1038 mV
    UCD9090_0-CH_7-EA1_2V5      2497 mV
    UCD9090_0-CH_8-EA1_1V5      1498 mV
    UCD9090_0-CH_9-VDD_1V8      1804 mV
    UCD9090_0-CH_10-VDD_2V5     2499 mV
    UCD9090_1-CH_1-EA2_PLL_     1017 mV
    UCD9090_1-CH_2-EA2_1V04     1041 mV
    UCD9090_1-CH_3-EA2_2V5      2499 mV
    UCD9090_1-CH_4-EA2_1V5      1503 mV
    UCD9090_1-CH_5-EA3_PLL_     1015 mV
    UCD9090_1-CH_6-EA3_1V04     1048 mV
    UCD9090_1-CH_7-EA3_2V5      2499 mV
    UCD9090_1-CH_8-EA3_1V5      1500 mV
    UCD9090_1-CH_9-VDD_1V5      1497 mV
    UCD9090_1-CH_10-VDD_1V2     1216 mV
    PMB PVCC 0.7V - 1.05V        802 mV
    PMB PVNN 0V - 1.02V          976 mV
    PMB 1.0V                    1002 mV
    PMB 1.1V                    1076 mV
    PMB 1.35V                   1347 mV
    PMB VDDQ 1.5V               1504 mV
    PMB 1.8V                    1804 mV
    PMB VDD 3.3V                3292 mV
    PMB BIAS 5.0V               5008 mV
    PMB USB 5.0V                5000 mV
    PMB 12V                    10866 mV
  I2C Slave Revision         112
FPC 7 status:
  State                      Online
  Temperature Intake         31 degrees C / 87 degrees F
  Temperature Exhaust A      46 degrees C / 114 degrees F
  Temperature Exhaust B      38 degrees C / 100 degrees F
  Temperature QX 0 TSen      49 degrees C / 120 degrees F
  Temperature QX 0 Chip      52 degrees C / 125 degrees F
```

```
  Temperature LU 0 TCAM TSen 49 degrees C / 120 degrees F
  Temperature LU 0 TCAM Chip 52 degrees C / 125 degrees F
  Temperature LU 0 TSen      49 degrees C / 120 degrees F
  Temperature LU 0 Chip      51 degrees C / 123 degrees F
  Temperature MQ 0 TSen      49 degrees C / 120 degrees F
  Temperature MQ 0 Chip      55 degrees C / 131 degrees F
  Temperature QX 1 TSen      41 degrees C / 105 degrees F
  Temperature QX 1 Chip      42 degrees C / 107 degrees F
  Temperature LU 1 TCAM TSen 41 degrees C / 105 degrees F
  Temperature LU 1 TCAM Chip 43 degrees C / 109 degrees F
  Temperature LU 1 TSen      41 degrees C / 105 degrees F
  Temperature LU 1 Chip      46 degrees C / 114 degrees F
  Temperature MQ 1 TSen      41 degrees C / 105 degrees F
  Temperature MQ 1 Chip      47 degrees C / 116 degrees F
  Power
    MPC-BIAS3V3-zl2105          3302 mV
    MPC-VDD3V3-zl2006           3307 mV
    MPC-VDD2V5-zl2006           2505 mV
    MPC-TCAM_1V0-zl2004         1000 mV
    MPC-AVDD1V0-zl2006          1006 mV
    MPC-VDD1V8-zl2006           1800 mV
    MPC-PCIE_1V0-zl2006         1000 mV
    MPC-LU0_1V0-zl2004           997 mV
    MPC-MQ0_1V0-zl2004           999 mV
    MPC-VDD_1V5-zl2004          1495 mV
    MPC-PMB_1V1-zl2006          1096 mV
    MPC-9VA-BMR453              9051 mV
    MPC-9VB-BMR453              8990 mV
    MPC-PMB_1V2-zl2106          1200 mV
    MPC-LU1_1V0-zl2004           997 mV
    MPC-MQ1_1V0-zl2004           998 mV
    MPC-QXM0_1V0-zl2006         1000 mV
    MPC-QXM1_1V0-zl2006          999 mV
  I2C Slave Revision         70
```

show chassis environment fpc (MX240 Router)

```
user@host> show chassis environment fpc
FPC 1 status:
  State                      Online
  Temperature Intake         34 degrees C / 93 degrees F
```

```
  Temperature Exhaust A      39 degrees C / 102 degrees F
  Temperature Exhaust B      53 degrees C / 127 degrees F
  Temperature I3 0 TSensor   51 degrees C / 123 degrees F
  Temperature I3 0 Chip      54 degrees C / 129 degrees F
  Temperature I3 1 TSensor   50 degrees C / 122 degrees F
  Temperature I3 1 Chip      53 degrees C / 127 degrees F
  Temperature I3 2 TSensor   48 degrees C / 118 degrees F
  Temperature I3 2 Chip      51 degrees C / 123 degrees F
  Temperature I3 3 TSensor   45 degrees C / 113 degrees F
  Temperature I3 3 Chip      48 degrees C / 118 degrees F
  Temperature IA 0 TSensor   45 degrees C / 113 degrees F
  Temperature IA 0 Chip      45 degrees C / 113 degrees F
  Temperature IA 1 TSensor   45 degrees C / 113 degrees F
  Temperature IA 1 Chip      49 degrees C / 120 degrees F
  Power
    1.5 V                     1492 mV
    2.5 V                     2507 mV
    3.3 V                     3306 mV
    1.8 V PFE 0               1801 mV
    1.8 V PFE 1               1804 mV
    1.8 V PFE 2               1798 mV
    1.8 V PFE 3               1798 mV
    1.2 V PFE 0               1169 mV
    1.2 V PFE 1               1189 mV
    1.2 V PFE 2               1182 mV
    1.2 V PFE 3               1176 mV
  I2C Slave Revision         42
FPC 2 status:
  State                      Online
  Temperature Intake         33 degrees C / 91 degrees F
  Temperature Exhaust A      41 degrees C / 105 degrees F
  Temperature Exhaust B      53 degrees C / 127 degrees F
  Temperature I3 0 TSensor   53 degrees C / 127 degrees F
  Temperature I3 0 Chip      58 degrees C / 136 degrees F
  Temperature I3 1 TSensor   52 degrees C / 125 degrees F
  Temperature I3 1 Chip      56 degrees C / 132 degrees F
  Temperature I3 2 TSensor   50 degrees C / 122 degrees F
  Temperature I3 2 Chip      52 degrees C / 125 degrees F
  Temperature I3 3 TSensor   46 degrees C / 114 degrees F
  Temperature I3 3 Chip      49 degrees C / 120 degrees F
  Temperature IA 0 TSensor   51 degrees C / 123 degrees F
  Temperature IA 0 Chip      49 degrees C / 120 degrees F
  Temperature IA 1 TSensor   48 degrees C / 118 degrees F
```

```
  Temperature IA 1 Chip      53 degrees C / 127 degrees F
  Power
    1.5 V                     1492 mV
    2.5 V                     2445 mV
    3.3 V                     3293 mV
    1.8 V PFE 0               1827 mV
    1.8 V PFE 1               1775 mV
    1.8 V PFE 2               1788 mV
    1.8 V PFE 3               1798 mV
    1.2 V PFE 0               1250 mV
    1.2 V PFE 1               1234 mV
    1.2 V PFE 2               1231 mV
    1.2 V PFE 3               1192 mV
  I2C Slave Revision         42
```

show chassis environment fpc (MX480 Router)

```
user@host> show chassis environment fpc
FPC 1 status:
  State                      Online
  Temperature Intake         36 degrees C / 96 degrees F
  Temperature Exhaust A      41 degrees C / 105 degrees F
  Temperature Exhaust B      55 degrees C / 131 degrees F
  Temperature I3 0 TSensor   55 degrees C / 131 degrees F
  Temperature I3 0 Chip      57 degrees C / 134 degrees F
  Temperature I3 1 TSensor   53 degrees C / 127 degrees F
  Temperature I3 1 Chip      53 degrees C / 127 degrees F
  Temperature I3 2 TSensor   52 degrees C / 125 degrees F
  Temperature I3 2 Chip      49 degrees C / 120 degrees F
  Temperature I3 3 TSensor   47 degrees C / 116 degrees F
  Temperature I3 3 Chip      47 degrees C / 116 degrees F
  Temperature IA 0 TSensor   54 degrees C / 129 degrees F
  Temperature IA 0 Chip      58 degrees C / 136 degrees F
  Temperature IA 1 TSensor   48 degrees C / 118 degrees F
  Temperature IA 1 Chip      53 degrees C / 127 degrees F
  Power
    1.5 V                     1479 mV
    2.5 V                     2542 mV
    3.3 V                     3319 mV
    1.8 V PFE 0               1811 mV
    1.8 V PFE 1               1804 mV
```

```
    1.8 V PFE 2               1804 mV
    1.8 V PFE 3               1814 mV
    1.2 V PFE 0               1192 mV
    1.2 V PFE 1               1202 mV
    1.2 V PFE 2               1205 mV
    1.2 V PFE 3               1189 mV
  I2C Slave Revision         40
```

show chassis environment fpc (MX960 Router MPC10E-15C-MRATE)

```
user@router> show chassis environment fpc 8

FPC 8 status:
  State                      Online
  Temperature Intake                     37 degrees C / 98 degrees F
  Temperature Exhaust A                  50 degrees C / 122 degrees F
  Temperature Exhaust B                  56 degrees C / 132 degrees F
  Temperature ZT0 Chip                   83 degrees C / 181 degrees F
  Temperature ZT1 Chip                   80 degrees C / 176 degrees F
  Temperature ZT2 Chip                   81 degrees C / 177 degrees F
  Temperature PCIE_SW Chip               64 degrees C / 147 degrees F
  Temperature ZT0 TestMacro              73 degrees C / 163 degrees F
  Temperature ZT0 hbmio_grp3             74 degrees C / 165 degrees F
  Temperature ZT0 hbmio_grp0             76 degrees C / 168 degrees F
  Temperature ZT0 gumem1                 78 degrees C / 172 degrees F
  Temperature ZT0 llm                    80 degrees C / 176 degrees F
  Temperature ZT0 wanio_sd               78 degrees C / 172 degrees F
  Temperature ZT0 fabio_sd               84 degrees C / 183 degrees F
  Temperature ZT0 flexmem                84 degrees C / 183 degrees F
  Temperature ZT1 TestMacro              70 degrees C / 158 degrees F
  Temperature ZT1 hbmio_grp3             71 degrees C / 159 degrees F
  Temperature ZT1 hbmio_grp0             74 degrees C / 165 degrees F
  Temperature ZT1 gumem1                 75 degrees C / 167 degrees F
  Temperature ZT1 llm                    78 degrees C / 172 degrees F
  Temperature ZT1 wanio_sd               76 degrees C / 168 degrees F
  Temperature ZT1 fabio_sd               78 degrees C / 172 degrees F
  Temperature ZT1 flexmem                82 degrees C / 179 degrees F
  Temperature ZT2 TestMacro              71 degrees C / 159 degrees F
  Temperature ZT2 hbmio_grp3             72 degrees C / 161 degrees F
  Temperature ZT2 hbmio_grp0             75 degrees C / 167 degrees F
  Temperature ZT2 gumem1                 76 degrees C / 168 degrees F
```

```
  Temperature ZT2 llm                    78 degrees C / 172 degrees F
  Temperature ZT2 wanio_sd               78 degrees C / 172 degrees F
  Temperature ZT2 fabio_sd               80 degrees C / 176 degrees F
  Temperature ZT2 flexmem                76 degrees C / 168 degrees F
  Temperature ZT0 HBM0                   74 degrees C / 165 degrees F
  Temperature ZT0 HBM1                   74 degrees C / 165 degrees F
  Temperature ZT1 HBM0                   74 degrees C / 165 degrees F
  Temperature ZT1 HBM1                   75 degrees C / 167 degrees F
  Temperature ZT2 HBM0                   73 degrees C / 163 degrees F
  Temperature ZT2 HBM1                   73 degrees C / 163 degrees F
  Temperature FAB RT1.0                  73 degrees C / 163 degrees F
  Temperature FAB RT2.0                  75 degrees C / 167 degrees F
  Temperature FAB RT3.0                  73 degrees C / 163 degrees F
  Temperature FAB RT4.0                  70 degrees C / 158 degrees F
  Temperature FAB RT5.0                  67 degrees C / 152 degrees F
  Temperature FAB RT6.0                  67 degrees C / 152 degrees F
  Temperature FAB RT7.0                  65 degrees C / 149 degrees F
  Temperature FAB RT8.0                  66 degrees C / 150 degrees F
  Temperature WAN RT9.0                  64 degrees C / 147 degrees F
  Temperature WAN RT9.1                  62 degrees C / 143 degrees F
  Temperature WAN RT10.0                 65 degrees C / 149 degrees F
  Temperature WAN RT10.1                 63 degrees C / 145 degrees F
  Temperature WAN RT11.0                 51 degrees C / 123 degrees F
  Temperature WAN RT11.1                 49 degrees C / 120 degrees F
  Temperature PIM4820 T1                 72 degrees C / 161 degrees F
  Temperature BMR456-12V-BRICK-A T1      83 degrees C / 181 degrees F
  Temperature BMR456-12V-BRICK-B T1      91 degrees C / 195 degrees F
  Temperature MAX20730-ZT0-AVDDH T1      72 degrees C / 161 degrees F
  Temperature MAX20730-ZT0-HBM-VDDQ T1   64 degrees C / 147 degrees F
  Temperature MAX20730-ZT0-HBM-VDDC T1   65 degrees C / 149 degrees F
  Temperature MAX20730-ZT1-AVDDH T1      65 degrees C / 149 degrees F
  Temperature MAX20730-ZT1-HBM-VDDQ T1   60 degrees C / 140 degrees F
  Temperature MAX20730-ZT1-HBM-VDDC T1   57 degrees C / 134 degrees F
  Temperature MAX20730-ZT2-AVDDH T1      65 degrees C / 149 degrees F
  Temperature MAX20730-ZT2-HBM-VDDQ T1   58 degrees C / 136 degrees F
  Temperature MAX20730-ZT2-HBM-VDDC T1   55 degrees C / 131 degrees F
  Temperature CPU0_PMB                   61 degrees C / 141 degrees F
  Temperature CPU7_PMB                   61 degrees C / 141 degrees F
  Temperature DDR4 A                     38 degrees C / 100 degrees F
  Temperature DDR4 B                     37 degrees C / 98 degrees F
  Power
    PIM4820                    56967 mV
    BMR456-12V-BRICK-A         12016 mV
```

```
    BMR456-12V-BRICK-B         12039 mV
    MAX20743-RT01-DVDD           724 mV
    MAX20743-RT234-DVDD          724 mV
    MAX20743-RT567-DVDD          724 mV
    MAX20754-ZT0-VDD             750 mV
    MAX20754-ZT0-VDDM            799 mV
    MAX20743-ZT0-AVDD            904 mV
    MAX20730-ZT0-AVDDH          1103 mV
    MAX20730-ZT0-HBM-VDDQ       1198 mV
    MAX20730-ZT0-HBM-VDDC       1202 mV
    MAX20730-VDD-1V25           1246 mV
    MAX20754-ZT1-VDD             724 mV
    MAX20754-ZT1-VDDM            800 mV
    MAX20743-ZT1-AVDD            904 mV
    MAX20730-ZT1-AVDDH          1103 mV
    MAX20730-ZT1-HBM-VDDQ       1202 mV
    MAX20730-ZT1-HBM-VDDC       1198 mV
    MAX20730-PCIE-0V9            901 mV
    MAX20754-ZT2-VDD             724 mV
    MAX20754-ZT2-VDDM            799 mV
    MAX20743-ZT2-AVDD            904 mV
    MAX20730-ZT2-AVDDH          1103 mV
    MAX20730-ZT2-HBM-VDDQ       1198 mV
    MAX20730-ZT2-HBM-VDDC       1198 mV
    MAX20730-VDD3V3             3308 mV
    MAX20754-WAN-VDD3V3         3301 mV
    MAX20754-WAN-DVDD0V8         799 mV
    MAX20743-WAN-VDD1V0A        1003 mV
    MAX20743-WAN-AVDD0V8         800 mV
    MAX20743-WAN-VDD1V0C        1003 mV
    TPS53631-1V2-VDDQ-PMB       1225 mV
    TPS53641-VCCIN-PMB          1770 mV
    TPS53641-VCCSBUS-PMB        1040 mV
    MAX20730-BIAS3P30-PMB       3308 mV
    MAX20730-BIAS5P0-PMB        5063 mV
    MAX20730-VPP-V2P5-PMB       2503 mV
    MAX20730-VDD1V2             1195 mV
    MAX20730-VDD1V5             1496 mV
    MAX20730-VDD1V8             1799 mV
    MAX20730-VDD2V5             2511 mV
    MAX20754-RT-AVDD-0V8         800 mV
    MAX20743-XGE-VDD-AVS        1012 mV
    PMB VCC1P05_PCH_SW          1048 mV
```

```
    PMB VCC1P3                  1294 mV
    PMB VCC1P5                  1485 mV
    PMB VCC1P7                  1705 mV
    PMB DDR4_VPP                2519 mV
    PMB VCC3P3                  3336 mV
    PMB VCC3P3_PCH              3332 mV
  I2C Slave Revision         124
```

show chassis environment fpc (MX960 Router)

```
user@host> show chassis environment fpc
FPC 5 status:
  State                      Online
  Temperature Intake         27 degrees C / 80 degrees F
  Temperature Exhaust A      34 degrees C / 93 degrees F
  Temperature Exhaust B      40 degrees C / 104 degrees F
  Temperature I3 0 TSensor   39 degrees C / 102 degrees F
  Temperature I3 0 Chip      41 degrees C / 105 degrees F
  Temperature I3 1 TSensor   38 degrees C / 100 degrees F
  Temperature I3 1 Chip      37 degrees C / 98 degrees F
  Temperature I3 2 TSensor   37 degrees C / 98 degrees F
  Temperature I3 2 Chip      34 degrees C / 93 degrees F
  Temperature I3 3 TSensor   32 degrees C / 89 degrees F
  Temperature I3 3 Chip      33 degrees C / 91 degrees F
  Temperature IA 0 TSensor   39 degrees C / 102 degrees F
  Temperature IA 0 Chip      44 degrees C / 111 degrees F
  Temperature IA 1 TSensor   36 degrees C / 96 degrees F
  Temperature IA 1 Chip      44 degrees C / 111 degrees F
  Power
    1.5 V                     1479 mV
    2.5 V                     2523 mV
    3.3 V                     3254 mV
    1.8 V PFE 0               1798 mV
    1.8 V PFE 1               1798 mV
    1.8 V PFE 2               1807 mV
    1.8 V PFE 3               1791 mV
    1.2 V PFE 0               1173 mV
    1.2 V PFE 1               1179 mV
    1.2 V PFE 2               1179 mV
    1.2 V PFE 3               1185 mV
  I2C Slave Revision         6
```

```
FPC 6 status:
  State                       Online
  Temperature Intake          25 degrees C / 77 degrees F
  Temperature Exhaust A       38 degrees C / 100 degrees F
  Temperature Exhaust B       38 degrees C / 100 degrees F
  Temperature I3 0 TSensor    40 degrees C / 104 degrees F
  Temperature I3 0 Chip       40 degrees C / 104 degrees F
  Temperature I3 1 TSensor    40 degrees C / 104 degrees F
  Temperature I3 1 Chip       38 degrees C / 100 degrees F
  Temperature I3 2 TSensor    37 degrees C / 98 degrees F
  Temperature I3 2 Chip       32 degrees C / 89 degrees F
  Temperature I3 3 TSensor    34 degrees C / 93 degrees F
  Temperature I3 3 Chip       33 degrees C / 91 degrees F
  Temperature IA 0 TSensor    45 degrees C / 113 degrees F
  Temperature IA 0 Chip       47 degrees C / 116 degrees F
  Temperature IA 1 TSensor    37 degrees C / 98 degrees F
  Temperature IA 1 Chip       42 degrees C / 107 degrees F
  Power
    1.5 V                      1485 mV
    2.5 V                      2510 mV
    3.3 V                      3332 mV
    1.8 V PFE 0                1801 mV
    1.8 V PFE 1                1814 mV
    1.8 V PFE 2                1804 mV
    1.8 V PFE 3                1820 mV
    1.2 V PFE 0                1192 mV
    1.2 V PFE 1                1189 mV
    1.2 V PFE 2                1202 mV
    1.2 V PFE 3                1156 mV
  I2C Slave Revision          40
```

show chassis environment fpc (MX480 Router with 100-Gigabit Ethernet CFP)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                       Online
  Temperature Intake          32 degrees C / 89 degrees F
  Temperature Exhaust A       39 degrees C / 102 degrees F
  Temperature Exhaust B       37 degrees C / 98 degrees F
  Temperature QX 0 TSen       44 degrees C / 111 degrees F
  Temperature QX 0 Chip       48 degrees C / 118 degrees F
```

```
  Temperature LU 0 TCAM TSen 44 degrees C / 111 degrees F
  Temperature LU 0 TCAM Chip 47 degrees C / 116 degrees F
  Temperature LU 0 TSen      44 degrees C / 111 degrees F
  Temperature LU 0 Chip      48 degrees C / 118 degrees F
  Temperature MQ 0 TSen      44 degrees C / 111 degrees F
  Temperature MQ 0 Chip      47 degrees C / 116 degrees F
  Power
    MPC-BIAS3V3-zl2105        3297 mV
    MPC-VDD3V3-zl2105         3306 mV
    MPC-VDD2V5-zl2105         2498 mV
    MPC-TCAM_1V0-zl2004        999 mV
    MPC-AVDD1V0-zl2006         999 mV
    MPC-VDD1V8-zl2006         1796 mV
    MPC-PCIE_1V0-zl2006       1002 mV
    MPC-LU0_1V0-zl2004         997 mV
    MPC-MQ0_1V0-zl2004         995 mV
    MPC-VDD_1V5-zl2004        1496 mV
    MPC-PMB_1V1-zl2006        1094 mV
    MPC-9VA-BMR453            9054 mV
    MPC-9VB-BMR453            9037 mV
    MPC-PMB_1V2-zl2106        1191 mV
    MPC-QXM0_1V0-zl2006       1000 mV
  I2C Slave Revision         66
FPC 1 status:
  State                      Online
  Temperature Intake         35 degrees C / 95 degrees F
  Temperature Exhaust A      50 degrees C / 122 degrees F
  Temperature Exhaust B      56 degrees C / 132 degrees F
  Temperature LU 0 TSen      46 degrees C / 114 degrees F
  Temperature LU 0 Chip      59 degrees C / 138 degrees F
  Temperature LU 1 TSen      46 degrees C / 114 degrees F
  Temperature LU 1 Chip      45 degrees C / 113 degrees F
  Temperature LU 2 TSen      46 degrees C / 114 degrees F
  Temperature LU 2 Chip      60 degrees C / 140 degrees F
  Temperature LU 3 TSen      46 degrees C / 114 degrees F
  Temperature LU 3 Chip      71 degrees C / 159 degrees F
  Temperature XM 0 TSen      46 degrees C / 114 degrees F
  Temperature XM 0 Chip      -18 degrees C / 0 degrees F
  Temperature XF 0 TSen      46 degrees C / 114 degrees F
  Temperature XF 0 Chip      76 degrees C / 168 degrees F
  Power
    MPC-BIAS3V3-zl2105        3292 mV
    MPC-VDD3V3-zl6100         3303 mV
```

```
    MPC-VDD2V5-zl6100         2501 mV
    MPC-VDD1V8-zl2004         1801 mV
    MPC-AVDD1V0-zl2006         996 mV
    MPC-VDD1V2-zl6100         1199 mV
    MPC-VDD1V5A-zl2004        1493 mV
    MPC-VDD1V5B-zl2004        1498 mV
    MPC-XF_0V9-zl2006          996 mV
    MPC-PCIE_1V0-zl6100       1000 mV
    MPC-LU0_1V0-zl2004         994 mV
    MPC-LU1_1V0-zl2004         994 mV
    MPC-LU2_1V0-zl2004         992 mV
    MPC-LU3_1V0-zl2004         993 mV
    MPC-12VA-BMR453          12003 mV
    MPC-12VB-BMR453          12043 mV
    MPC-PMB_1V1-zl2006        1091 mV
    MPC-PMB_1V2-zl2106        1196 mV
    MPC-XM_0V9-vt273m          899 mV
  I2C Slave Revision         106
```

show chassis environment fpc (MX240, MX480, MX960 with Application Services Modular Line Card

```
user@host>show chassis environment fpc 1
FPC 1 status:
  State                      Online
  Temperature Intake         36 degrees C / 96 degrees F
  Temperature Exhaust A      39 degrees C / 102 degrees F
  Temperature LU TSen        52 degrees C / 125 degrees F
  Temperature LU Chip        54 degrees C / 129 degrees F
  Temperature XM TSen        52 degrees C / 125 degrees F
  Temperature XM Chip        60 degrees C / 140 degrees F
  Temperature PCIe TSen      52 degrees C / 125 degrees F
  Temperature PCIe Chip      69 degrees C / 156 degrees F
  Power
    MPC-BIAS3V3-zl2106        3302 mV
    MPC-VDD3V3-zl6100         3325 mV
    MPC-AVDD1V0-zl6100        1007 mV
    MPC-PCIE_1V0-zl6100        904 mV
    MPC-LU0_1V0-zl2004         996 mV
    MPC-VDD_1V5-zl2004        1498 mV
    MPC-12VA-BMR453          11733 mV
```

```
    MPC-12VB-BMR453          11728 mV
    MPC-XM_0V9-vt273m          900 mV
  I2C Slave Revision         81
```

show chassis environment fpc (MX10003 Router)

```
user@host> show chassis environment fpc

FPC 0 status:
  State                      Online
  FPC 0 Intake Temp Sensor   29 degrees C / 84 degrees F
  FPC 0 Exhaust-A Temp Sensor56 degrees C / 132 degrees F
  FPC 0 Exhaust-B Temp Sensor44 degrees C / 111 degrees F
  FPC 0 EA0 Chip             58 degrees C / 136 degrees F
  FPC 0 EA0-XR0 Chip         61 degrees C / 141 degrees F
  FPC 0 EA0-XR1 Chip         62 degrees C / 143 degrees F
  FPC 0 EA1 Chip             67 degrees C / 152 degrees F
  FPC 0 EA1-XR0 Chip         72 degrees C / 161 degrees F
  FPC 0 EA1-XR1 Chip         72 degrees C / 161 degrees F
  FPC 0 PEX Chip             77 degrees C / 170 degrees F
  FPC 0 EA2 Chip             48 degrees C / 118 degrees F
  FPC 0 EA2-XR0 Chip         54 degrees C / 129 degrees F
  FPC 0 EA2-XR1 Chip         56 degrees C / 132 degrees F
  FPC 0 PF Chip              68 degrees C / 154 degrees F
  FPC 0 EA0_HMC0 Logic die   72 degrees C / 161 degrees F
  FPC 0 EA0_HMC0 DRAM botm   69 degrees C / 156 degrees F
  FPC 0 EA0_HMC1 Logic die   71 degrees C / 159 degrees F
  FPC 0 EA0_HMC1 DRAM botm   68 degrees C / 154 degrees F
  FPC 0 EA0_HMC2 Logic die   75 degrees C / 167 degrees F
  FPC 0 EA0_HMC2 DRAM botm   72 degrees C / 161 degrees F
  FPC 0 EA1_HMC0 Logic die   81 degrees C / 177 degrees F
  FPC 0 EA1_HMC0 DRAM botm   78 degrees C / 172 degrees F
  FPC 0 EA1_HMC1 Logic die   80 degrees C / 176 degrees F
  FPC 0 EA1_HMC1 DRAM botm   77 degrees C / 170 degrees F
  FPC 0 EA1_HMC2 Logic die   82 degrees C / 179 degrees F
  FPC 0 EA1_HMC2 DRAM botm   79 degrees C / 174 degrees F
  FPC 0 EA2_HMC0 Logic die   60 degrees C / 140 degrees F
  FPC 0 EA2_HMC0 DRAM botm   57 degrees C / 134 degrees F
  FPC 0 EA2_HMC1 Logic die   61 degrees C / 141 degrees F
  FPC 0 EA2_HMC1 DRAM botm   58 degrees C / 136 degrees F
  FPC 0 EA2_HMC2 Logic die   63 degrees C / 145 degrees F
```

```
FPC 0 EA2_HMC2 DRAM botm   60 degrees C / 140 degrees F
Power
  LTC3887-PF-VDD0V9-RAIL       898 mV
  LTC3887-PF-VDD0V9-DEV0-      898 mV
  LTC3887-PF-VDD0V9-DEV0-      900 mV
  LTC3887-PF-VDD0V9-DEV1-      899 mV
  LTC3887-PF-VDD0V9-DEV1-      901 mV
  LTC3887-PF-AVDD1V0-RAIL      998 mV
  LTC3887-PF-AVDD1V0-CH0       998 mV
  LTC3887-PF-AVDD1V0-CH1       999 mV
  LTC3887-ETHSW-VDD1V0        1000 mV
  LTC3887-VDD2V5              2499 mV
  LTC3887-PCIE-VDD0V9          899 mV
  LTC3887-V1P0                 999 mV
  LTC3887-PHY-VDD1V0-A         999 mV
  LTC3887-3V3                 3300 mV
  LTC3887-VDD1V8              1799 mV
  UCD9090_0-CH_1-EA0_PLL_     1005 mV
  UCD9090_0-CH_2-EA0_1V4      1049 mV
  UCD9090_0-CH_3-EA0_2V5      2499 mV
  UCD9090_0-CH_4-EA0_1V5      1499 mV
  UCD9090_0-CH_5-EA1_PLL_      999 mV
  UCD9090_0-CH_6-EA1_1V4      1037 mV
  UCD9090_0-CH_7-EA1_2V5      2499 mV
  UCD9090_0-CH_8-EA1_1V5      1510 mV
  UCD9090_0-CH_9-PVCC          797 mV
  UCD9090_0-CH_10-PVNN         991 mV
  UCD9090_1-CH_1-EA2_PLL_     1008 mV
  UCD9090_1-CH_2-EA2_1V4      1009 mV
  UCD9090_1-CH_3-EA2_2V5      2499 mV
  UCD9090_1-CH_4-EA2_1V5      1513 mV
  UCD9090_1-CH_5-1V0_PFPL     1009 mV
  UCD9090_1-CH_6-V1P1         1075 mV
  UCD9090_1-CH_7-V1P5         1531 mV
  UCD9090_1-CH_8-V1P35        1359 mV
  UCD9090_1-CH_9-VDD1V5       1511 mV
  UCD9090_1-CH_10-VDD1V2      1210 mV
  LTC3887-EA0-VDD0V9-RAIL      949 mV
  LTC3887-EA0-VDD0V9-DEV0      949 mV
  LTC3887-EA0-VDD0V9-DEV0      951 mV
  LTC3887-EA0-VDD0V9-DEV1      949 mV
  LTC3887-EA0-VDD0V9-DEV1      951 mV
  LTC3887-EA0-VDD0V9R2-RA      947 mV
```

```
    LTC3887-EA0-VDD0V9R2-CH      947 mV
    LTC3887-EA0-VDD0V9R2-CH      949 mV
    LTC3887-EA0-VDD1V0-RAIL      999 mV
    LTC3887-EA0-VDD1V0-CH0       999 mV
    LTC3887-EA0-VDD1V0-CH1      1001 mV
    LTC3887-EA0-XR-VDD0V9        900 mV
    LTC3887-EA0-XR-VDD1V2       1199 mV
    LTC3887-EA0-HM1-VDD0V9       899 mV
    LTC3887-EA0-HM-VDD1V2       1200 mV
    LTC3887-EA0-HM-VDDM1V2      1199 mV
    LTC3887-EA1-VDD0V9-RAIL      949 mV
    LTC3887-EA1-VDD0V9-DEV0      952 mV
    LTC3887-EA1-VDD0V9-DEV0      952 mV
    LTC3887-EA1-VDD0V9-DEV1      951 mV
    LTC3887-EA1-VDD0V9-DEV1      951 mV
    LTC3887-EA1-VDD0V9R2-RA      948 mV
    LTC3887-EA1-VDD0V9R2-CH      948 mV
    LTC3887-EA1-VDD0V9R2-CH      950 mV
    LTC3887-EA1-VDD1V0-RAIL     1000 mV
    LTC3887-EA1-VDD1V0-CH0      1000 mV
    LTC3887-EA1-VDD1V0-CH1      1001 mV
  I2C Slave Revision         13
FPC 1 status:
  State                      Online
  FPC 1 Intake Temp Sensor   27 degrees C / 80 degrees F
  FPC 1 Exhaust-A Temp Sensor60 degrees C / 140 degrees F
  FPC 1 Exhaust-B Temp Sensor46 degrees C / 114 degrees F
  FPC 1 EA0 Chip             63 degrees C / 145 degrees F
  FPC 1 EA0-XR0 Chip         67 degrees C / 152 degrees F
  FPC 1 EA0-XR1 Chip         68 degrees C / 154 degrees F
  FPC 1 EA1 Chip             70 degrees C / 158 degrees F
  FPC 1 EA1-XR0 Chip         75 degrees C / 167 degrees F
  FPC 1 EA1-XR1 Chip         75 degrees C / 167 degrees F
  FPC 1 PEX Chip             89 degrees C / 192 degrees F
  FPC 1 EA2 Chip             49 degrees C / 120 degrees F
  FPC 1 EA2-XR0 Chip         53 degrees C / 127 degrees F
  FPC 1 EA2-XR1 Chip         56 degrees C / 132 degrees F
  FPC 1 PF Chip              71 degrees C / 159 degrees F
  FPC 1 EA0_HMC0 Logic die   74 degrees C / 165 degrees F
  FPC 1 EA0_HMC0 DRAM botm   71 degrees C / 159 degrees F
  FPC 1 EA0_HMC1 Logic die   78 degrees C / 172 degrees F
  FPC 1 EA0_HMC1 DRAM botm   75 degrees C / 167 degrees F
  FPC 1 EA0_HMC2 Logic die   78 degrees C / 172 degrees F
```

```
FPC 1 EA0_HMC2 DRAM botm   75 degrees C / 167 degrees F
FPC 1 EA1_HMC0 Logic die   84 degrees C / 183 degrees F
FPC 1 EA1_HMC0 DRAM botm   81 degrees C / 177 degrees F
FPC 1 EA1_HMC1 Logic die   82 degrees C / 179 degrees F
FPC 1 EA1_HMC1 DRAM botm   79 degrees C / 174 degrees F
FPC 1 EA1_HMC2 Logic die   85 degrees C / 185 degrees F
FPC 1 EA1_HMC2 DRAM botm   82 degrees C / 179 degrees F
FPC 1 EA2_HMC0 Logic die   62 degrees C / 143 degrees F
FPC 1 EA2_HMC0 DRAM botm   59 degrees C / 138 degrees F
FPC 1 EA2_HMC1 Logic die   60 degrees C / 140 degrees F
FPC 1 EA2_HMC1 DRAM botm   57 degrees C / 134 degrees F
FPC 1 EA2_HMC2 Logic die   65 degrees C / 149 degrees F
FPC 1 EA2_HMC2 DRAM botm   62 degrees C / 143 degrees F
Power
  LTC3887-PF-VDD0V9-RAIL       899 mV
  LTC3887-PF-VDD0V9-DEV0-      899 mV
  LTC3887-PF-VDD0V9-DEV0-      901 mV
  LTC3887-PF-VDD0V9-DEV1-      899 mV
  LTC3887-PF-VDD0V9-DEV1-      901 mV
  LTC3887-PF-AVDD1V0-RAIL      998 mV
  LTC3887-PF-AVDD1V0-CH0       998 mV
  LTC3887-PF-AVDD1V0-CH1       999 mV
  LTC3887-ETHSW-VDD1V0         999 mV
  LTC3887-VDD2V5              2499 mV
  LTC3887-PCIE-VDD0V9          900 mV
  LTC3887-V1P0                1000 mV
  LTC3887-PHY-VDD1V0-A        1000 mV
  LTC3887-3V3                 3300 mV
  LTC3887-VDD1V8              1799 mV
  UCD9090_0-CH_1-EA0_PLL_     1004 mV
  UCD9090_0-CH_2-EA0_1V4      1004 mV
  UCD9090_0-CH_3-EA0_2V5      2499 mV
  UCD9090_0-CH_4-EA0_1V5      1511 mV
  UCD9090_0-CH_5-EA1_PLL_      999 mV
  UCD9090_0-CH_6-EA1_1V4      1008 mV
  UCD9090_0-CH_7-EA1_2V5      2499 mV
  UCD9090_0-CH_8-EA1_1V5      1510 mV
  UCD9090_0-CH_9-PVCC          839 mV
  UCD9090_0-CH_10-PVNN        1016 mV
  UCD9090_1-CH_1-EA2_PLL_     1011 mV
  UCD9090_1-CH_2-EA2_1V4      1046 mV
  UCD9090_1-CH_3-EA2_2V5      2499 mV
  UCD9090_1-CH_4-EA2_1V5      1501 mV
```

```
    UCD9090_1-CH_5-1V0_PFPL     1000 mV
    UCD9090_1-CH_6-V1P1         1037 mV
    UCD9090_1-CH_7-V1P5         1530 mV
    UCD9090_1-CH_8-V1P35        1360 mV
    UCD9090_1-CH_9-VDD1V5       1513 mV
    UCD9090_1-CH_10-VDD1V2      1217 mV
    LTC3887-EA0-VDD0V9-RAIL      949 mV
    LTC3887-EA0-VDD0V9-DEV0      949 mV
    LTC3887-EA0-VDD0V9-DEV0      951 mV
    LTC3887-EA0-VDD0V9-DEV1      949 mV
    LTC3887-EA0-VDD0V9-DEV1      952 mV
    LTC3887-EA0-VDD0V9R2-RA      947 mV
    LTC3887-EA0-VDD0V9R2-CH      947 mV
    LTC3887-EA0-VDD0V9R2-CH      949 mV
    LTC3887-EA0-VDD1V0-RAIL     1000 mV
    LTC3887-EA0-VDD1V0-CH0      1000 mV
    LTC3887-EA0-VDD1V0-CH1      1001 mV
    LTC3887-EA0-XR-VDD0V9        899 mV
    LTC3887-EA0-XR-VDD1V2       1200 mV
    LTC3887-EA0-HM1-VDD0V9       899 mV
    LTC3887-EA0-HM-VDD1V2       1199 mV
    LTC3887-EA0-HM-VDDM1V2      1199 mV
    LTC3887-EA1-VDD0V9-RAIL      948 mV
    LTC3887-EA1-VDD0V9-DEV0      950 mV
    LTC3887-EA1-VDD0V9-DEV0      950 mV
    LTC3887-EA1-VDD0V9-DEV1      951 mV
    LTC3887-EA1-VDD0V9-DEV1      951 mV
    LTC3887-EA1-VDD0V9R2-RA      947 mV
    LTC3887-EA1-VDD0V9R2-CH      947 mV
    LTC3887-EA1-VDD0V9R2-CH      949 mV
    LTC3887-EA1-VDD1V0-RAIL     1000 mV
    LTC3887-EA1-VDD1V0-CH0      1000 mV
    LTC3887-EA1-VDD1V0-CH1      1002 mV
  I2C Slave Revision         99
```

show chassis environment fpc (MX204 Router)

```
user@host> show chassis environment fpc

FPC 0 status:
  State                      Online
```

```
  FPC 0 EA0_HMC0 Logic die   77 degrees C / 170 degrees F
  FPC 0 EA0_HMC0 DRAM botm   74 degrees C / 165 degrees F
  FPC 0 EA0_HMC1 Logic die   80 degrees C / 176 degrees F
  FPC 0 EA0_HMC1 DRAM botm   77 degrees C / 170 degrees F
  FPC 0 EA0 Chip             93 degrees C / 199 degrees F
  FPC 0 EA0-XR0 Chip         63 degrees C / 145 degrees F
  FPC 0 EA0-XR1 Chip         64 degrees C / 147 degrees F
  Power
  I2C Slave Revision         0
```

show chassis environment fpc (MX10008 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                      Online
  FPC 0 Intake-A Temp Sensor              32 degrees C / 89 degrees F
  FPC 0 Exhaust-A Temp Sensor             44 degrees C / 111 degrees F
  FPC 0 Exhaust-B Temp Sensor             50 degrees C / 122 degrees F
  FPC 0 EA0 Temp Sensor                   67 degrees C / 152 degrees F
  FPC 0 EA0_XR0 Temp Sensor               69 degrees C / 156 degrees F
  FPC 0 EA0_XR1 Temp Sensor               73 degrees C / 163 degrees F
  FPC 0 EA1 Temp Sensor                   61 degrees C / 141 degrees F
  FPC 0 EA1_XR0 Temp Sensor               65 degrees C / 149 degrees F
  FPC 0 EA1_XR1 Temp Sensor               63 degrees C / 145 degrees F
  FPC 0 EA2 Temp Sensor                   69 degrees C / 156 degrees F
  FPC 0 EA2_XR0 Temp Sensor               73 degrees C / 163 degrees F
  FPC 0 EA2_XR1 Temp Sensor               72 degrees C / 161 degrees F
  FPC 0 EA3 Temp Sensor                   64 degrees C / 147 degrees F
  FPC 0 EA3_XR0 Temp Sensor               66 degrees C / 150 degrees F
  FPC 0 EA3_XR1 Temp Sensor               66 degrees C / 150 degrees F
  FPC 0 EA4 Temp Sensor                   70 degrees C / 158 degrees F
  FPC 0 EA4_XR0 Temp Sensor               72 degrees C / 161 degrees F
  FPC 0 EA4_XR1 Temp Sensor               72 degrees C / 161 degrees F
  FPC 0 EA5 Temp Sensor                   58 degrees C / 136 degrees F
  FPC 0 EA5_XR0 Temp Sensor               61 degrees C / 141 degrees F
  FPC 0 EA5_XR1 Temp Sensor               64 degrees C / 147 degrees F
  FPC 0 EA0_HMC0 Logic die                75 degrees C / 167 degrees F
  FPC 0 EA0_HMC0 DRAM botm                72 degrees C / 161 degrees F
  FPC 0 EA0_HMC1 Logic die                76 degrees C / 168 degrees F
  FPC 0 EA0_HMC1 DRAM botm                73 degrees C / 163 degrees F
```

```
FPC 0 EA0_HMC2 Logic die                77 degrees C / 170 degrees F
FPC 0 EA0_HMC2 DRAM botm                74 degrees C / 165 degrees F
FPC 0 EA1_HMC0 Logic die                72 degrees C / 161 degrees F
FPC 0 EA1_HMC0 DRAM botm                69 degrees C / 156 degrees F
FPC 0 EA1_HMC1 Logic die                73 degrees C / 163 degrees F
FPC 0 EA1_HMC1 DRAM botm                70 degrees C / 158 degrees F
FPC 0 EA1_HMC2 Logic die                72 degrees C / 161 degrees F
FPC 0 EA1_HMC2 DRAM botm                69 degrees C / 156 degrees F
FPC 0 EA2_HMC0 Logic die                80 degrees C / 176 degrees F
FPC 0 EA2_HMC0 DRAM botm                77 degrees C / 170 degrees F
FPC 0 EA2_HMC1 Logic die                81 degrees C / 177 degrees F
FPC 0 EA2_HMC1 DRAM botm                78 degrees C / 172 degrees F
FPC 0 EA2_HMC2 Logic die                80 degrees C / 176 degrees F
FPC 0 EA2_HMC2 DRAM botm                77 degrees C / 170 degrees F
FPC 0 EA3_HMC0 Logic die                77 degrees C / 170 degrees F
FPC 0 EA3_HMC0 DRAM botm                74 degrees C / 165 degrees F
FPC 0 EA3_HMC1 Logic die                78 degrees C / 172 degrees F
FPC 0 EA3_HMC1 DRAM botm                75 degrees C / 167 degrees F
FPC 0 EA3_HMC2 Logic die                77 degrees C / 170 degrees F
FPC 0 EA3_HMC2 DRAM botm                74 degrees C / 165 degrees F
FPC 0 EA4_HMC0 Logic die                80 degrees C / 176 degrees F
FPC 0 EA4_HMC0 DRAM botm                77 degrees C / 170 degrees F
FPC 0 EA4_HMC1 Logic die                81 degrees C / 177 degrees F
FPC 0 EA4_HMC1 DRAM botm                78 degrees C / 172 degrees F
FPC 0 EA4_HMC2 Logic die                80 degrees C / 176 degrees F
FPC 0 EA4_HMC2 DRAM botm                77 degrees C / 170 degrees F
FPC 0 EA5_HMC0 Logic die                69 degrees C / 156 degrees F
FPC 0 EA5_HMC0 DRAM botm                66 degrees C / 150 degrees F
FPC 0 EA5_HMC1 Logic die                68 degrees C / 154 degrees F
FPC 0 EA5_HMC1 DRAM botm                65 degrees C / 149 degrees F
FPC 0 EA5_HMC2 Logic die                68 degrees C / 154 degrees F
FPC 0 EA5_HMC2 DRAM botm                65 degrees C / 149 degrees F
Power
        12V SS 1                   12259 mV    9841 mA  120642 mW
        12V SS 2                   12259 mV   21054 mA  258104 mW
        12V SS 3                   12285 mV    9841 mA  120902 mW
        12V SS 4                   12232 mV   20968 mA  256496 mW
        12V SS 5                   12179 mV   14993 mA  182614 mW
        VDD 1.0V_A                  1000 mV   95375 mA   95375 mW
        VDD 1.0V_B                     0 mV       0 mA       0 mW
        VDD 3.3V                    3298 mV   12500 mA   41235 mW
        VDD 0.9V                     894 mV    3569 mA    3192 mW
        ETH SW 1V                    980 mV    4500 mA    4410 mW
```

```
        VDD 1.8V                    1809 mV     895 mA    1619 mW
        PVCC                         951 mV       0 mA       0 mW
        PVNN                        1009 mV       0 mA       0 mW
        V1P0                        1006 mV       0 mA       0 mW
        V1P1                        1070 mV       0 mA       0 mW
        V1P3                        1351 mV       0 mA       0 mW
        VDDQ                        1500 mV       0 mA       0 mW
        V1P8                        1816 mV       0 mA       0 mW
        VDD3V3                      3296 mV       0 mA       0 mW
        V5V0_BIAS                   5025 mV       0 mA       0 mW
        VDD12V0                    12174 mV       0 mA       0 mW
        EA0 Core 0.9V                900 mV   40625 mA   36578 mW
        EA0 AVDD 1.0V               1000 mV   32500 mA   32500 mW
        EA0 HMC Core 0.9V            894 mV   10081 mA    9017 mW
        EA0 1.2V                    1189 mV   15081 mA   17945 mW
        EA01_HMC_VDDM 1.2V          1193 mV    -151 mA    -180 mW
        EA0_XR 0.906V                905 mV   13802 mA   12496 mW
        EA1 Core 0.9V                900 mV   41000 mA   36916 mW
        EA1 AVDD 1.0V               1000 mV   28000 mA   28000 mW
        EA1 HMC Core 0.9V            897 mV    9848 mA    8835 mW
        EA1 1.2V                    1197 mV   15313 mA   18332 mW
        EA0_PLL_1V0                 1003 mV       0 mA       0 mW
        EA0_1V04                    1032 mV       0 mA       0 mW
        EA0_2V5                     2445 mV       0 mA       0 mW
        EA0_1V5                     1512 mV       0 mA       0 mW
        EA1_PLL_1V0                 1000 mV       0 mA       0 mW
        EA1_1V04                    1051 mV       0 mA       0 mW
        EA1_2V5                     2516 mV       0 mA       0 mW
        EA1_1V5                     1503 mV       0 mA       0 mW
        EA1_XR 0.906V                908 mV   14151 mA   12850 mW
        EA2 Core 0.9V                899 mV   40625 mA   36538 mW
        EA2 AVDD 1.0V               1000 mV   27250 mA   27276 mW
        EA2 HMC Core 0.9V            897 mV    9616 mA    8627 mW
        EA2 1.2V                    1193 mV   15779 mA   18832 mW
        EA23_HMC_VDDM 1.2V          1197 mV      81 mA      97 mW
        EA2_XR 0.906V                908 mV   14848 mA   13484 mW
        EA3 Core 0.9V                899 mV   40625 mA   36538 mW
        EA3 AVDD 1.0V               1000 mV   28000 mA   28000 mW
        EA3 HMC Core 0.9V            897 mV   10546 mA    9461 mW
        EA3 1.2V                    1197 mV   15895 mA   19028 mW
        EA2_PLL_1V0                 1025 mV       0 mA       0 mW
        EA2_1V04                    1048 mV       0 mA       0 mW
        EA2_2V5                     2516 mV       0 mA       0 mW
```

```
          EA2_1V5                     1500 mV       0 mA       0 mW
          EA3_PLL_1V0                 1009 mV       0 mA       0 mW
          EA3_1V04                    1032 mV       0 mA       0 mW
          EA3_2V5                     2551 mV       0 mA       0 mW
          EA3_1V5                     1496 mV       0 mA       0 mW
          EA3_XR 0.906V                908 mV   15895 mA   14434 mW
          EA4 Core 0.9V                900 mV   41000 mA   36916 mW
          EA4 AVDD 1.0V                999 mV   31250 mA   31219 mW
          EA4 HMC Core 0.9V            894 mV    9965 mA    8913 mW
          EA4 1.2V                    1197 mV   15779 mA   18889 mW
          EA45_HMC_VDDM 1.2V          1197 mV     546 mA     654 mW
          EA4_XR 0.906V                908 mV   15197 mA   13801 mW
          EA5 Core 0.9V                900 mV   39750 mA   35790 mW
          EA5 AVDD 1.0V               1000 mV   28000 mA   28000 mW
          EA5 HMC Core 0.9V            897 mV    9965 mA    8940 mW
          EA5 1.2V                    1197 mV   15546 mA   18610 mW
          EA4_PLL_1V0                 1003 mV       0 mA       0 mW
          EA4_1V04                    1041 mV       0 mA       0 mW
          EA4_2V5                     2541 mV       0 mA       0 mW
          EA4_1V5                     1506 mV       0 mA       0 mW
          EA5_PLL_1V0                 1022 mV       0 mA       0 mW
          EA5_1V04                    1048 mV       0 mA       0 mW
          EA5_2V5                     2532 mV       0 mA       0 mW
          EA5_1V5                     1509 mV       0 mA       0 mW
          VDD2V5                      2503 mV       0 mA       0 mW
          VDD1V5                      1509 mV       0 mA       0 mW
          VDD1V2                      1206 mV       0 mA       0 mW
          EA5_XR 0.906V                902 mV   14500 mA   13088 mW
FPC 2 status:
  State                       Online
  FPC 2 Intake-A Temp Sensor              33 degrees C / 91 degrees F
  FPC 2 Exhaust-A Temp Sensor             52 degrees C / 125 degrees F
  FPC 2 Exhaust-B Temp Sensor             50 degrees C / 122 degrees F
  FPC 2 EA0 Temp Sensor                   72 degrees C / 161 degrees F
  FPC 2 EA0_XR0 Temp Sensor               76 degrees C / 168 degrees F
  FPC 2 EA0_XR1 Temp Sensor               79 degrees C / 174 degrees F
  FPC 2 EA1 Temp Sensor                   64 degrees C / 147 degrees F
  FPC 2 EA1_XR0 Temp Sensor               68 degrees C / 154 degrees F
  FPC 2 EA1_XR1 Temp Sensor               66 degrees C / 150 degrees F
  FPC 2 EA2 Temp Sensor                   75 degrees C / 167 degrees F
  FPC 2 EA2_XR0 Temp Sensor               81 degrees C / 177 degrees F
  FPC 2 EA2_XR1 Temp Sensor               81 degrees C / 177 degrees F
  FPC 2 EA3 Temp Sensor                   67 degrees C / 152 degrees F
```

```
FPC 2 EA3_XR0 Temp Sensor              69 degrees C / 156 degrees F
FPC 2 EA3_XR1 Temp Sensor              69 degrees C / 156 degrees F
FPC 2 EA4 Temp Sensor                  76 degrees C / 168 degrees F
FPC 2 EA4_XR0 Temp Sensor              77 degrees C / 170 degrees F
FPC 2 EA4_XR1 Temp Sensor              76 degrees C / 168 degrees F
FPC 2 EA5 Temp Sensor                  60 degrees C / 140 degrees F
FPC 2 EA5_XR0 Temp Sensor              65 degrees C / 149 degrees F
FPC 2 EA5_XR1 Temp Sensor              65 degrees C / 149 degrees F
FPC 2 EA0_HMC0 Logic die               84 degrees C / 183 degrees F
FPC 2 EA0_HMC0 DRAM botm               81 degrees C / 177 degrees F
FPC 2 EA0_HMC1 Logic die               86 degrees C / 186 degrees F
FPC 2 EA0_HMC1 DRAM botm               83 degrees C / 181 degrees F
FPC 2 EA0_HMC2 Logic die               83 degrees C / 181 degrees F
FPC 2 EA0_HMC2 DRAM botm               80 degrees C / 176 degrees F
FPC 2 EA1_HMC0 Logic die               76 degrees C / 168 degrees F
FPC 2 EA1_HMC0 DRAM botm               73 degrees C / 163 degrees F
FPC 2 EA1_HMC1 Logic die               77 degrees C / 170 degrees F
FPC 2 EA1_HMC1 DRAM botm               74 degrees C / 165 degrees F
FPC 2 EA1_HMC2 Logic die               76 degrees C / 168 degrees F
FPC 2 EA1_HMC2 DRAM botm               73 degrees C / 163 degrees F
FPC 2 EA2_HMC0 Logic die               87 degrees C / 188 degrees F
FPC 2 EA2_HMC0 DRAM botm               84 degrees C / 183 degrees F
FPC 2 EA2_HMC1 Logic die               89 degrees C / 192 degrees F
FPC 2 EA2_HMC1 DRAM botm               86 degrees C / 186 degrees F
FPC 2 EA2_HMC2 Logic die               88 degrees C / 190 degrees F
FPC 2 EA2_HMC2 DRAM botm               85 degrees C / 185 degrees F
FPC 2 EA3_HMC0 Logic die               80 degrees C / 176 degrees F
FPC 2 EA3_HMC0 DRAM botm               77 degrees C / 170 degrees F
FPC 2 EA3_HMC1 Logic die               81 degrees C / 177 degrees F
FPC 2 EA3_HMC1 DRAM botm               78 degrees C / 172 degrees F
FPC 2 EA3_HMC2 Logic die               81 degrees C / 177 degrees F
FPC 2 EA3_HMC2 DRAM botm               78 degrees C / 172 degrees F
FPC 2 EA4_HMC0 Logic die               88 degrees C / 190 degrees F
FPC 2 EA4_HMC0 DRAM botm               85 degrees C / 185 degrees F
FPC 2 EA4_HMC1 Logic die               90 degrees C / 194 degrees F
FPC 2 EA4_HMC1 DRAM botm               87 degrees C / 188 degrees F
FPC 2 EA4_HMC2 Logic die               81 degrees C / 177 degrees F
FPC 2 EA4_HMC2 DRAM botm               78 degrees C / 172 degrees F
FPC 2 EA5_HMC0 Logic die               73 degrees C / 163 degrees F
FPC 2 EA5_HMC0 DRAM botm               70 degrees C / 158 degrees F
FPC 2 EA5_HMC1 Logic die               69 degrees C / 156 degrees F
FPC 2 EA5_HMC1 DRAM botm               66 degrees C / 150 degrees F
FPC 2 EA5_HMC2 Logic die               73 degrees C / 163 degrees F
```

```
FPC 2 EA5_HMC2 DRAM botm                70 degrees C / 158 degrees F
Power
        12V SS 1                   12285 mV    9408 mA  115582 mW
        12V SS 2                   12338 mV   20881 mA  257637 mW
        12V SS 3                   12351 mV   10317 mA  127430 mW
        12V SS 4                   12285 mV   21054 mA  258660 mW
        12V SS 5                   12153 mV   13954 mA  169591 mW
        VDD 1.0V_A                  1000 mV   91000 mA   91000 mW
        VDD 1.0V_B                     0 mV       0 mA       0 mW
        VDD 3.3V                    3298 mV    9125 mA   30101 mW
        VDD 0.9V                     897 mV    3337 mA    2993 mW
        ETH SW 1V                      0 mV       0 mA       0 mW
        VDD 1.8V                    1809 mV    1127 mA    2040 mW
        PVCC                         835 mV       0 mA       0 mW
        PVNN                        1000 mV       0 mA       0 mW
        V1P0                        1003 mV       0 mA       0 mW
        V1P1                        1070 mV       0 mA       0 mW
        V1P3                        1348 mV       0 mA       0 mW
        VDDQ                        1493 mV       0 mA       0 mW
        V1P8                        1806 mV       0 mA       0 mW
        VDD3V3                      3303 mV       0 mA       0 mW
        V5V0_BIAS                   5000 mV       0 mA       0 mW
        VDD12V0                    12116 mV       0 mA       0 mW
        EA0 Core 0.9V                900 mV   38875 mA   35002 mW
        EA0 AVDD 1.0V                999 mV   31875 mA   31843 mW
        EA0 HMC Core 0.9V            894 mV    9034 mA    8081 mW
        EA0 1.2V                    1197 mV   15430 mA   18471 mW
        EA01_HMC_VDDM 1.2V          1200 mV    -267 mA    -321 mW
        EA0_XR 0.906V                908 mV   15430 mA   14012 mW
        EA1 Core 0.9V                900 mV   38875 mA   35002 mW
        EA1 AVDD 1.0V               1000 mV   28250 mA   28250 mW
        EA1 HMC Core 0.9V            899 mV    8802 mA    7920 mW
        EA1 1.2V                    1197 mV   15081 mA   18054 mW
        EA0_PLL_1V0                 1003 mV       0 mA       0 mW
        EA0_1V04                    1048 mV       0 mA       0 mW
        EA0_2V5                     2425 mV       0 mA       0 mW
        EA0_1V5                     1483 mV       0 mA       0 mW
        EA1_PLL_1V0                 1019 mV       0 mA       0 mW
        EA1_1V04                    1019 mV       0 mA       0 mW
        EA1_2V5                     2490 mV       0 mA       0 mW
        EA1_1V5                     1480 mV       0 mA       0 mW
        EA1_XR 0.906V                908 mV   14965 mA   13590 mW
        EA2 Core 0.9V                900 mV   44000 mA   39617 mW
```

```
        EA2 AVDD 1.0V              1000 mV   28625 mA   28625 mW
        EA2 HMC Core 0.9V           891 mV   10546 mA    9404 mW
        EA2 1.2V                   1200 mV   15313 mA   18387 mW
        EA23_HMC_VDDM 1.2V         1193 mV    -267 mA    -319 mW
        EA2_XR 0.906V               908 mV   15197 mA   13801 mW
        EA3 Core 0.9V               900 mV   39750 mA   35790 mW
        EA3 AVDD 1.0V              1000 mV   27750 mA   27750 mW
        EA3 HMC Core 0.9V           897 mV    9267 mA    8314 mW
        EA3 1.2V                   1197 mV   15430 mA   18471 mW
        EA2_PLL_1V0                1009 mV       0 mA       0 mW
        EA2_1V04                   1041 mV       0 mA       0 mW
        EA2_2V5                    2496 mV       0 mA       0 mW
        EA2_1V5                    1493 mV       0 mA       0 mW
        EA3_PLL_1V0                1003 mV       0 mA       0 mW
        EA3_1V04                   1041 mV       0 mA       0 mW
        EA3_2V5                    2490 mV       0 mA       0 mW
        EA3_1V5                    1500 mV       0 mA       0 mW
        EA3_XR 0.906V               908 mV   15081 mA   13695 mW
        EA4 Core 0.9V               899 mV   45750 mA   41148 mW
        EA4 AVDD 1.0V              1000 mV   32250 mA   32250 mW
        EA4 HMC Core 0.9V           897 mV   10779 mA    9670 mW
        EA4 1.2V                   1193 mV   16011 mA   19110 mW
        EA45_HMC_VDDM 1.2V         1200 mV    -267 mA    -321 mW
        EA4_XR 0.906V               905 mV   15779 mA   14286 mW
        EA5 Core 0.9V               900 mV   38375 mA   34552 mW
        EA5 AVDD 1.0V              1000 mV   27750 mA   27777 mW
        EA5 HMC Core 0.9V           899 mV    8453 mA    7606 mW
        EA5 1.2V                   1200 mV   14732 mA   17689 mW
        EA4_PLL_1V0                1012 mV       0 mA       0 mW
        EA4_1V04                   1029 mV       0 mA       0 mW
        EA4_2V5                    2496 mV       0 mA       0 mW
        EA4_1V5                    1490 mV       0 mA       0 mW
        EA5_PLL_1V0                1003 mV       0 mA       0 mW
        EA5_1V04                   1032 mV       0 mA       0 mW
        EA5_2V5                    2503 mV       0 mA       0 mW
        EA5_1V5                    1480 mV       0 mA       0 mW
        VDD2V5                     2461 mV       0 mA       0 mW
        VDD1V5                     1490 mV       0 mA       0 mW
        VDD1V2                     1212 mV       0 mA       0 mW
        EA5_XR 0.906V               910 mV   13686 mA   12466 mW
FPC 3 status:
  State                      Online
  FPC 3 Intake-A Temp Sensor             30 degrees C / 86 degrees F
```

```
FPC 3 Exhaust-A Temp Sensor          48 degrees C / 118 degrees F
FPC 3 Exhaust-B Temp Sensor          45 degrees C / 113 degrees F
FPC 3 EA0 Temp Sensor                60 degrees C / 140 degrees F
FPC 3 EA0_XR0 Temp Sensor            65 degrees C / 149 degrees F
FPC 3 EA0_XR1 Temp Sensor            67 degrees C / 152 degrees F
FPC 3 EA1 Temp Sensor                54 degrees C / 129 degrees F
FPC 3 EA1_XR0 Temp Sensor            60 degrees C / 140 degrees F
FPC 3 EA1_XR1 Temp Sensor            58 degrees C / 136 degrees F
FPC 3 EA2 Temp Sensor                62 degrees C / 143 degrees F
FPC 3 EA2_XR0 Temp Sensor            67 degrees C / 152 degrees F
FPC 3 EA2_XR1 Temp Sensor            67 degrees C / 152 degrees F
FPC 3 EA3 Temp Sensor                55 degrees C / 131 degrees F
FPC 3 EA3_XR0 Temp Sensor            57 degrees C / 134 degrees F
FPC 3 EA3_XR1 Temp Sensor            57 degrees C / 134 degrees F
FPC 3 EA4 Temp Sensor                69 degrees C / 156 degrees F
FPC 3 EA4_XR0 Temp Sensor            71 degrees C / 159 degrees F
FPC 3 EA4_XR1 Temp Sensor            70 degrees C / 158 degrees F
FPC 3 EA5 Temp Sensor                55 degrees C / 131 degrees F
FPC 3 EA5_XR0 Temp Sensor            58 degrees C / 136 degrees F
FPC 3 EA5_XR1 Temp Sensor            59 degrees C / 138 degrees F
FPC 3 EA0_HMC0 Logic die             69 degrees C / 156 degrees F
FPC 3 EA0_HMC0 DRAM botm             66 degrees C / 150 degrees F
FPC 3 EA0_HMC1 Logic die             70 degrees C / 158 degrees F
FPC 3 EA0_HMC1 DRAM botm             67 degrees C / 152 degrees F
FPC 3 EA0_HMC2 Logic die             70 degrees C / 158 degrees F
FPC 3 EA0_HMC2 DRAM botm             67 degrees C / 152 degrees F
FPC 3 EA1_HMC0 Logic die             68 degrees C / 154 degrees F
FPC 3 EA1_HMC0 DRAM botm             65 degrees C / 149 degrees F
FPC 3 EA1_HMC1 Logic die             65 degrees C / 149 degrees F
FPC 3 EA1_HMC1 DRAM botm             62 degrees C / 143 degrees F
FPC 3 EA1_HMC2 Logic die             64 degrees C / 147 degrees F
FPC 3 EA1_HMC2 DRAM botm             61 degrees C / 141 degrees F
FPC 3 EA2_HMC0 Logic die             74 degrees C / 165 degrees F
FPC 3 EA2_HMC0 DRAM botm             71 degrees C / 159 degrees F
FPC 3 EA2_HMC1 Logic die             77 degrees C / 170 degrees F
FPC 3 EA2_HMC1 DRAM botm             74 degrees C / 165 degrees F
FPC 3 EA2_HMC2 Logic die             74 degrees C / 165 degrees F
FPC 3 EA2_HMC2 DRAM botm             71 degrees C / 159 degrees F
FPC 3 EA3_HMC0 Logic die             70 degrees C / 158 degrees F
FPC 3 EA3_HMC0 DRAM botm             67 degrees C / 152 degrees F
FPC 3 EA3_HMC1 Logic die             68 degrees C / 154 degrees F
FPC 3 EA3_HMC1 DRAM botm             65 degrees C / 149 degrees F
FPC 3 EA3_HMC2 Logic die             68 degrees C / 154 degrees F
```

```
FPC 3 EA3_HMC2 DRAM botm                65 degrees C / 149 degrees F
FPC 3 EA4_HMC0 Logic die                82 degrees C / 179 degrees F
FPC 3 EA4_HMC0 DRAM botm                79 degrees C / 174 degrees F
FPC 3 EA4_HMC1 Logic die                80 degrees C / 176 degrees F
FPC 3 EA4_HMC1 DRAM botm                77 degrees C / 170 degrees F
FPC 3 EA4_HMC2 Logic die                81 degrees C / 177 degrees F
FPC 3 EA4_HMC2 DRAM botm                78 degrees C / 172 degrees F
FPC 3 EA5_HMC0 Logic die                69 degrees C / 156 degrees F
FPC 3 EA5_HMC0 DRAM botm                66 degrees C / 150 degrees F
FPC 3 EA5_HMC1 Logic die                70 degrees C / 158 degrees F
FPC 3 EA5_HMC1 DRAM botm                67 degrees C / 152 degrees F
FPC 3 EA5_HMC2 Logic die                69 degrees C / 156 degrees F
FPC 3 EA5_HMC2 DRAM botm                66 degrees C / 150 degrees F
Power
       12V SS 1                    12259 mV    9538 mA  116927 mW
       12V SS 2                    12259 mV   20491 mA  251202 mW
       12V SS 3                    12298 mV    9711 mA  119433 mW
       12V SS 4                    12219 mV   20491 mA  250391 mW
       12V SS 5                    12206 mV   10447 mA  127520 mW
       VDD 1.0V_A                   1000 mV   42250 mA   42291 mW
       VDD 1.0V_B                    996 mV    8918 mA    8890 mW
       VDD 3.3V                     3301 mV   10375 mA   34255 mW
       VDD 0.9V                      897 mV    3569 mA    3202 mW
       ETH SW 1V                     983 mV    4267 mA    4195 mW
       VDD 1.8V                     1812 mV    1825 mA    3309 mW
       PVCC                          974 mV       0 mA       0 mW
       PVNN                         1003 mV       0 mA       0 mW
       V1P0                         1003 mV       0 mA       0 mW
       V1P1                         1070 mV       0 mA       0 mW
       V1P3                         1351 mV       0 mA       0 mW
       VDDQ                         1496 mV       0 mA       0 mW
       V1P8                         1809 mV       0 mA       0 mW
       VDD3V3                       3309 mV       0 mA       0 mW
       V5V0_BIAS                    4987 mV       0 mA       0 mW
       VDD12V0                     12212 mV       0 mA       0 mW
       EA0 Core 0.9V                 900 mV   38125 mA   34327 mW
       EA0 AVDD 1.0V                 999 mV   31125 mA   31094 mW
       EA0 HMC Core 0.9V             897 mV    9500 mA    8522 mW
       EA0 1.2V                     1193 mV   15430 mA   18416 mW
       EA01_HMC_VDDM 1.2V           1193 mV     313 mA     374 mW
       EA0_XR 0.906V                 913 mV   14965 mA   13671 mW
       EA1 Core 0.9V                 900 mV   39750 mA   35790 mW
       EA1 AVDD 1.0V                1000 mV   26000 mA   26000 mW
```

```
        EA1 HMC Core 0.9V          897 mV    8918 mA    8001 mW
        EA1 1.2V                  1200 mV   15779 mA   18946 mW
        EA0_PLL_1V0               1003 mV       0 mA       0 mW
        EA0_1V04                  1019 mV       0 mA       0 mW
        EA0_2V5                   2448 mV       0 mA       0 mW
        EA0_1V5                   1470 mV       0 mA       0 mW
        EA1_PLL_1V0               1016 mV       0 mA       0 mW
        EA1_1V04                  1035 mV       0 mA       0 mW
        EA1_2V5                   2506 mV       0 mA       0 mW
        EA1_1V5                   1483 mV       0 mA       0 mW
        EA1_XR 0.906V              908 mV   13918 mA   12639 mW
        EA2 Core 0.9V              900 mV   38625 mA   34777 mW
        EA2 AVDD 1.0V             1000 mV   26375 mA   26400 mW
        EA2 HMC Core 0.9V          897 mV    9383 mA    8418 mW
        EA2 1.2V                  1200 mV   15779 mA   18946 mW
        EA23_HMC_VDDM 1.2V        1193 mV      81 mA      97 mW
        EA2_XR 0.906V              908 mV   13918 mA   12639 mW
        EA3 Core 0.9V              899 mV   40250 mA   36201 mW
        EA3 AVDD 1.0V             1000 mV   26750 mA   26776 mW
        EA3 HMC Core 0.9V          894 mV    9267 mA    8289 mW
        EA3 1.2V                  1197 mV   16127 mA   19306 mW
        EA2_PLL_1V0                993 mV       0 mA       0 mW
        EA2_1V04                  1045 mV       0 mA       0 mW
        EA2_2V5                   2474 mV       0 mA       0 mW
        EA2_1V5                   1490 mV       0 mA       0 mW
        EA3_PLL_1V0                980 mV       0 mA       0 mW
        EA3_1V04                  1032 mV       0 mA       0 mW
        EA3_2V5                   2506 mV       0 mA       0 mW
        EA3_1V5                   1474 mV       0 mA       0 mW
        EA3_XR 0.906V              910 mV   14732 mA   13419 mW
        EA4 Core 0.9V              900 mV   42500 mA   38266 mW
        EA4 AVDD 1.0V             1000 mV   32250 mA   32281 mW
        EA4 HMC Core 0.9V          899 mV   10081 mA    9071 mW
        EA4 1.2V                  1193 mV   16360 mA   19526 mW
        EA45_HMC_VDDM 1.2V        1193 mV     662 mA     791 mW
        EA4_XR 0.906V              908 mV   15430 mA   14012 mW
        EA5 Core 0.9V              899 mV   37000 mA   33278 mW
        EA5 AVDD 1.0V             1000 mV   26125 mA   26150 mW
        EA5 HMC Core 0.9V          897 mV    9267 mA    8314 mW
        EA5 1.2V                  1197 mV   15662 mA   18750 mW
        EA4_PLL_1V0               1000 mV       0 mA       0 mW
        EA4_1V04                  1029 mV       0 mA       0 mW
        EA4_2V5                   2487 mV       0 mA       0 mW
```

```
        EA4_1V5                    1496 mV       0 mA       0 mW
        EA5_PLL_1V0                1009 mV       0 mA       0 mW
        EA5_1V04                   1032 mV       0 mA       0 mW
        EA5_2V5                    2503 mV       0 mA       0 mW
        EA5_1V5                    1496 mV       0 mA       0 mW
        VDD2V5                     2483 mV       0 mA       0 mW
        VDD1V5                     1470 mV       0 mA       0 mW
        VDD1V2                     1203 mV       0 mA       0 mW
        EA5_XR 0.906V               908 mV   14500 mA   13167 mW
```

show chassis environment fpc (T320, T640, and T1600 Routers)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                            Online
  Temperature Top              42 degrees C / 107 degrees F
  Temperature Bottom           36 degrees C / 96 degrees F
  Temperature MMB1             39 degrees C / 102 degrees F
  Power:
    1.8 V                    1959 mV
    2.5 V                    2495 mV
    3.3 V                    3344 mV
    5.0 V                    5047 mV
    1.8 V bias               1787 mV
    3.3 V bias               3291 mV
    5.0 V bias               4998 mV
    8.0 V bias               7343 mV
  BUS Revision                 40
FPC 1 status:
  State                            Online
  Temperature Top              42 degrees C / 107 degrees F
  Temperature Bottom           39 degrees C / 102 degrees F
  Temperature MMB1             40 degrees C / 104 degrees F
  Power:
    1.8 V                    1956 mV
    2.5 V                    2498 mV
    3.3 V                    3340 mV
    5.0 V                    5023 mV
    1.8 V bias               1782 mV
    3.3 V bias               3277 mV
    5.0 V bias               4989 mV
```

```
    8.0 V bias               7289 mV
  BUS Revision                 40
FPC 2 status:
  State                           Online
  Temperature Top              43 degrees C / 109 degrees F
  Temperature Bottom           39 degrees C / 102 degrees F
  Temperature MMB1             41 degrees C / 105 degrees F
  Power:
    1.8 V                    1963 mV
    2.5 V                    2503 mV
    3.3 V                    3340 mV
    5.0 V                    5042 mV
    1.8 V bias               1797 mV
    3.3 V bias               3311 mV
    5.0 V bias               5013 mV
    8.0 V bias               7221 mV
  BUS Revision                 40
```

show chassis environment fpc (T4000 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                      Online
  Fan Intake                 34 degrees C / 93 degrees F
  Fan Exhaust                48 degrees C / 118 degrees F
  PMB                        47 degrees C / 116 degrees F
  LMB0                       50 degrees C / 122 degrees F
  LMB1                       41 degrees C / 105 degrees F
  LMB2                       35 degrees C / 95 degrees F
  PFE1 LU2                   46 degrees C / 114 degrees F
  PFE1 LU0                   41 degrees C / 105 degrees F
  PFE0 LU0                   57 degrees C / 134 degrees F
  XF1                        47 degrees C / 116 degrees F
  XF0                        52 degrees C / 125 degrees F
  XM1                        41 degrees C / 105 degrees F
  XM0                        50 degrees C / 122 degrees F
  PFE0 LU1                   56 degrees C / 132 degrees F
  PFE0 LU2                   45 degrees C / 113 degrees F
  PFE1 LU1                   37 degrees C / 98 degrees F
  Power 1
    1.0 V                      991 mV
```

```
    1.2 V bias                1195 mV
    1.8 V                     1788 mV
    2.5 V                     2483 mV
    3.3 V                     3289 mV
    3.3 V bias                3299 mV
    12.0 V A                 10608 mV
    12.0 V B                 10637 mV
  Power 2
    0.9 V                      881 mV
    0.9 V PFE0                 916 mV
    0.9 V PFE1                 903 mV
    1.0 V PFE0                1012 mV
    1.0 V PFE1                1002 mV
    1.1 V                     1095 mV
    1.5 V_0                   1494 mV
    1.5 V_1                   1479 mV
  Power 3
    1.0 V PFE0                1000 mV
    1.0 V PFE1                1002 mV
    1.0 V PFE0 *               995 mV
    1.0 V PFE1 *               995 mV
    1.8 V PFE 0               1788 mV
    1.8 V PFE 1               1789 mV
    2.5 V                     2482 mV
    12.0 V                   11614 mV
  Power 4
    1.0 V PFE0 LU0            1003 mV
    1.0 V PFE1 LU0            1003 mV
    1.0 V PFE1 LU2            1004 mV
    1.0 V PFE0 LU0 *           995 mV
    1.0 V PFE1 LU0 *           998 mV
    1.0 V PFE1 LU2 *           996 mV
    12.0 V                   11643 mV
    12.0 V C                 11711 mV
  Power (Base/PMB/MMB)
    LMB0 VDD2V5               2488 mV
    LMB0 VDD1V8               1788 mV
    LMB0 VDD1V5               1496 mV
    LMB0 PFE0 LU0 AVDD1V0     1002 mV
    LMB0 PFE0 LU0 VDD1V0      1000 mV
    LMB0 VDD12V0             10752 mV
    LMB1 VDD2V5               2472 mV
    LMB1 VDD1V8               1792 mV
```

```
    LMB1 VDD1V5                1480 mV
    LMB1 PFE0 LU2 AVDD1V0       994 mV
    LMB1 PFE0 LU2 VDD1V0       1002 mV
    LMB1 VDD12V0              10800 mV
    LMB2 VDD2V5                2472 mV
    LMB2 VDD1V8                1792 mV
    LMB2 VDD1V5                1486 mV
    LMB2 PFE1 LU1 AVDD1V0       996 mV
    LMB2 PFE1 LU1 VDD1V0        998 mV
    LMB2 VDD12V0              10704 mV
    PMB 1.05v                  1049 mV
    PMB 1.5v                   1500 mV
    PMB 2.5v                   2500 mV
    PMB 3.3v                   3299 mV
  Bus Revision                113
FPC 3 status:
  State                       Online
  Fan Intake                  37 degrees C / 98 degrees F
  Fan Exhaust                 51 degrees C / 123 degrees F
  PMB                         43 degrees C / 109 degrees F
  LMB0                        57 degrees C / 134 degrees F
  LMB1                        54 degrees C / 129 degrees F
  LMB2                        38 degrees C / 100 degrees F
  PFE1 LU2                    63 degrees C / 145 degrees F
  PFE1 LU0                    45 degrees C / 113 degrees F
  PFE0 LU0                    69 degrees C / 156 degrees F
  XF1                         62 degrees C / 143 degrees F
  XF0                         63 degrees C / 145 degrees F
  XM1                         43 degrees C / 109 degrees F
  XM0                         67 degrees C / 152 degrees F
  PFE0 LU1                    63 degrees C / 145 degrees F
  PFE0 LU2                    66 degrees C / 150 degrees F
  PFE1 LU1                    41 degrees C / 105 degrees F
  Power 1
    1.0 V                      1002 mV
    1.2 V bias                 1201 mV
    1.8 V                      1785 mV
    2.5 V                      2485 mV
    3.3 V                      3288 mV
    3.3 V bias                 3285 mV
    12.0 V A                  10412 mV
    12.0 V B                  10515 mV
  Power 2
```

```
    0.9 V                      882 mV
    0.9 V PFE0                 920 mV
    0.9 V PFE1                 905 mV
    1.0 V PFE0                1015 mV
    1.0 V PFE1                1001 mV
    1.1 V                     1094 mV
    1.5 V_0                   1495 mV
    1.5 V_1                   1478 mV
  Power 3
    0.92 V PFE1                998 mV
    1.0 V PFE0                 997 mV
    1.0 V PFE0 *               992 mV
    1.0 V PFE1 *               991 mV
    1.8 V PFE 0               1780 mV
    1.8 V PFE 1               1797 mV
    2.5 V                     2492 mV
    12.0 V                   11604 mV
  Power 4
    1.0 V PFE0 LU0            1003 mV
    1.0 V PFE1 LU0            1004 mV
    1.0 V PFE1 LU2            1003 mV
    1.0 V PFE0 LU0 *          1000 mV
    1.0 V PFE1 LU0 *          1001 mV
    1.0 V PFE1 LU2 *          1003 mV
    12.0 V                   11653 mV
    12.0 V C                 11672 mV
  Power (Base/PMB/MMB)
    LMB0 VDD2V5               2512 mV
    LMB0 VDD1V8               1790 mV
    LMB0 VDD1V5               1500 mV
    LMB0 PFE0 LU0 AVDD1V0     1004 mV
    LMB0 PFE0 LU0 VDD1V0      1002 mV
    LMB0 VDD12V0             10608 mV
    LMB1 VDD2V5               2472 mV
    LMB1 VDD1V8               1788 mV
    LMB1 VDD1V5               1480 mV
    LMB1 PFE0 LU2 AVDD1V0     1000 mV
    LMB1 PFE0 LU2 VDD1V0      1004 mV
    LMB1 VDD12V0             10672 mV
    LMB2 VDD2V5               2488 mV
    LMB2 VDD1V8               1798 mV
    LMB2 VDD1V5               1494 mV
    LMB2 PFE1 LU1 AVDD1V0     1000 mV
```

```
    LMB2 PFE1 LU1 VDD1V0      1004 mV
    LMB2 VDD12V0             10528 mV
    PMB 1.05v                 1050 mV
    PMB 1.5v                  1500 mV
    PMB 2.5v                  2499 mV
    PMB 3.3v                  3299 mV
  Bus Revision               113
FPC 5 status:
  State                      Online
  Temperature Top            39 degrees C / 102 degrees F
  Temperature Bottom         38 degrees C / 100 degrees F
  Power
    1.8 V                     1804 mV
    1.8 V bias                1802 mV
    3.3 V                     3294 mV
    3.3 V bias                3277 mV
    5.0 V bias                5008 mV
    5.0 V TOP                 5067 mV
    8.0 V bias                6642 mV
  Power (Base/PMB/MMB)
    1.2 V                     1202 mV
    1.5 V                     1504 mV
    5.0 V BOT                 5079 mV
    12.0 V TOP Base          11848 mV
    12.0 V BOT Base          11780 mV
    1.1 V PMB                 1111 mV
    1.2 V PMB                 1189 mV
    1.5 V PMB                 1494 mV
    1.8 V PMB                 1819 mV
    2.5 V PMB                 2503 mV
    3.3 V PMB                 3294 mV
    5.0 V PMB                 5035 mV
    12.0  V PMB              11788 mV
    0.75 MMB TOP               766 mV
    1.5 V MMB TOP             1484 mV
    1.8 V MMB TOP             1772 mV
    2.5 V MMB TOP             2485 mV
    1.2 V MMB TOP             1137 mV
    5.0 V MMB TOP             4946 mV
    12.0 V MMB TOP           11772 mV
    3.3 V MMB TOP             3289 mV
    0.75 MMB BOT               759 mV
    1.5 V MMB BOT             1482 mV
```

```
    1.8 V MMB BOT             1792 mV
    2.5 V MMB BOT             2490 mV
    1.2 V MMB BOT             1145 mV
    5.0 V MMB BOT             4922 mV
    12.0 V MMB BOT           11625 mV
    3.3 V MMB BOT             3282 mV
    APS 00                    2495 mV
    APS 01                    3308 mV
    APS 02                    3301 mV
    5.0 V PIC 0               4967 mV
    APS 10                    2512 mV
    APS 11                    3316 mV
    APS 12                    3304 mV
    5.0 V PIC 1               5081 mV
  Bus Revision               49
FPC 6 status:
  State                      Online
  Fan Intake                 34 degrees C / 93 degrees F
  Fan Exhaust                49 degrees C / 120 degrees F
  PMB                        40 degrees C / 104 degrees F
  LMB0                       60 degrees C / 140 degrees F
  LMB1                       58 degrees C / 136 degrees F
  LMB2                       40 degrees C / 104 degrees F
  PFE1 LU2                   69 degrees C / 156 degrees F
  PFE1 LU0                   45 degrees C / 113 degrees F
  PFE0 LU0                   71 degrees C / 159 degrees F
  XF1                        58 degrees C / 136 degrees F
  XF0                        65 degrees C / 149 degrees F
  XM1                        40 degrees C / 104 degrees F
  XM0                        66 degrees C / 150 degrees F
  PFE0 LU1                   69 degrees C / 156 degrees F
  PFE0 LU2                   68 degrees C / 154 degrees F
  PFE1 LU1                   42 degrees C / 107 degrees F
  Power 1
    1.0 V                      998 mV
    1.2 V bias                1191 mV
    1.8 V                     1781 mV
    2.5 V                     2487 mV
    3.3 V                     3302 mV
    3.3 V bias                3300 mV
    12.0 V A                 10388 mV
    12.0 V B                 10388 mV
  Power 2
```

```
    0.9 V                      902 mV
    0.9 V PFE0                 921 mV
    0.9 V PFE1                 907 mV
    1.0 V PFE0                 996 mV
    1.0 V PFE1                 974 mV
    1.1 V                     1095 mV
    1.5 V_0                   1495 mV
    1.5 V_1                   1478 mV
  Power 3
    1.0 V PFE0                 997 mV
    1.0 V PFE1                 998 mV
    1.0 V PFE0 *               993 mV
    1.0 V PFE1 *               991 mV
    1.8 V PFE 0               1796 mV
    1.8 V PFE 1               1789 mV
    2.5 V                     2465 mV
    12.0 V                   11609 mV
  Power 4
    1.0 V PFE0 LU0            1003 mV
    1.0 V PFE1 LU0            1006 mV
    1.0 V PFE1 LU2            1002 mV
    1.0 V PFE0 LU0 *          1000 mV
    1.0 V PFE1 LU0 *           998 mV
    1.0 V PFE1 LU2 *           998 mV
    12.0 V                   11638 mV
    12.0 V C                 11702 mV
  Power (Base/PMB/MMB)
    LMB0 VDD2V5               2484 mV
    LMB0 VDD1V8               1780 mV
    LMB0 VDD1V5               1496 mV
    LMB0 PFE0 LU0 AVDD1V0      998 mV
    LMB0 PFE0 LU0 VDD1V0      1004 mV
    LMB0 VDD12V0             10528 mV
    LMB1 VDD2V5               2472 mV
    LMB1 VDD1V8               1776 mV
    LMB1 VDD1V5               1474 mV
    LMB1 PFE0 LU2 AVDD1V0      994 mV
    LMB1 PFE0 LU2 VDD1V0      1004 mV
    LMB1 VDD12V0             10544 mV
    LMB2 VDD2V5               2476 mV
    LMB2 VDD1V8               1790 mV
    LMB2 VDD1V5               1492 mV
    LMB2 PFE1 LU1 AVDD1V0      996 mV
```

```
    LMB2 PFE1 LU1 VDD1V0      1010 mV
    LMB2 VDD12V0             10528 mV
    PMB 1.05v                 1050 mV
    PMB 1.5v                  1499 mV
    PMB 2.5v                  2500 mV
    PMB 3.3v                  3300 mV
  Bus Revision               80
```

show chassis environment fpc lcc (TX Matrix Router)

```
user@host> show chassis environment fpc lcc
0
lcc0-re0:
--------------------------------------------------------------------------
FPC 1 status:
  State                           Online
  Temperature Top              30 degrees C / 86 degrees F
  Temperature Bottom           25 degrees C / 77 degrees F
  Temperature MMB0                Absent
  Temperature MMB1             27 degrees C / 80 degrees F
  Power:
    1.8 V                    1813 mV
    2.5 V                    2504 mV
    3.3 V                    3338 mV
    5.0 V                    5037 mV
    1.8 V bias               1797 mV
    3.3 V bias               3301 mV
    5.0 V bias               5013 mV
    8.0 V bias               7345 mV
  BUS Revision                 40
FPC 2 status:
  State                           Online
  Temperature Top              37 degrees C / 98 degrees F
  Temperature Bottom           26 degrees C / 78 degrees F
  Temperature MMB0             32 degrees C / 89 degrees F
  Temperature MMB1             27 degrees C / 80 degrees F
  Power:
    1.8 V                    1791 mV
    2.5 V                    2517 mV
    3.3 V                    3308 mV
    5.0 V                    5052 mV
```

```
    1.8 V bias                1797 mV
    3.3 V bias                3289 mV
    5.0 V bias                4991 mV
    8.0 V bias                7477 mV
  BUS Revision                  40
```

show chassis environment fpc lcc (TX Matrix Plus Router)

```
user@host> show chassis environment fpc lcc
0
lcc0-re0:
--------------------------------------------------------------------------
FPC 1 status:
  State                       Online
  Temperature Top             46 degrees C / 114 degrees F
  Temperature Bottom          47 degrees C / 116 degrees F
  Power
    1.8 V                      1788 mV
    1.8 V bias                 1787 mV
    3.3 V                      3321 mV
    3.3 V bias                 3306 mV
    5.0 V bias                 5018 mV
    5.0 V TOP                  5037 mV
    8.0 V bias                 7223 mV
  Power (Base/PMB/MMB)
    1.2 V                      1205 mV
    1.5 V                      1503 mV
    5.0 V BOT                  5084 mV
    12.0 V TOP Base           11775 mV
    12.0 V BOT Base           11794 mV
    1.1 V PMB                  1108 mV
    1.2 V PMB                  1196 mV
    1.5 V PMB                  1499 mV
    1.8 V PMB                  1811 mV
    2.5 V PMB                  2515 mV
    3.3 V PMB                  3318 mV
    5.0 V PMB                  5030 mV
    12.0  V PMB               11832 mV
    0.75 MMB TOP                752 mV
    1.5 V MMB TOP              1489 mV
    1.8 V MMB TOP              1782 mV
```

```
    2.5 V MMB TOP             2498 mV
    1.2 V MMB TOP             1155 mV
    5.0 V MMB TOP             4902 mV
    12.0 V MMB TOP           11721 mV
    3.3 V MMB TOP             3316 mV
    0.75 MMB BOT               754 mV
    1.5 V MMB BOT             1482 mV
    1.8 V MMB BOT             1758 mV
    2.5 V MMB BOT             2488 mV
    1.2 V MMB BOT             1157 mV
    5.0 V MMB BOT             4962 mV
    12.0 V MMB BOT           11691 mV
    3.3 V MMB BOT             3308 mV
    APS 00                    1484 mV
    APS 01                    2503 mV
    APS 02                    3313 mV
    5.0 V PIC 0               5025 mV
    APS 10                    1501 mV
    APS 11                    2466 mV
    APS 12                    3311 mV
    5.0 V PIC 1               5081 mV
  Bus Revision               49
```

show chassis environment fpc (QFX Series and OCX Series)

```
user@switch> show chassis environment fpc 0
FPC 0 status:
  State                      Online
  Temperature                42 degrees C / 107 degrees F
```

show chassis environment fpc interconnect-device (QFabric Systems)

```
user@switch> show chassis environment fpc interconnect-device
interconnect1 0
FC 0 FPC 0 status:
  State                      Online
  Left Intake Temperature    24 degrees C / 75 degrees F
  Right Intake Temperature   24 degrees C / 75 degrees F
  Left Exhaust Temperature   27 degrees C / 80 degrees F
  Right Exhaust Temperature  27 degrees C / 80 degrees F
```

```
  Power
    BIAS 3V3                  3330 mV
    VDD 3V3                   3300 mV
    VDD 2V5                   2502 mV
    VDD 1V5                   1496 mV
    VDD 1V2                   1194 mV
    VDD 1V0                   1000 mV
    SW0 VDD  1V0              1020 mV
    SW0 CVDD 1V025            1032 mV
    SW1 VDD  1V0              1022 mV
    SW1 CVDD 1V025            1030 mV
    VDD 12V0 DIV3_33          3414 mV
```

show chassis environment fpc 5(PTX3000 Packet Transport Router)

```
user@host> show chassis environment fpc 5
FPC 5 status:
  State                       Online
  Intake Temperature          31 degrees C / 87 degrees F
  Exhaust Temperature         41 degrees C / 105 degrees F
  Power
    FPC 12.0v                  12221 mV
    FPC VCC 0.5-1.3v            1640 mV
    FPC VNN 0.5-1.3v            1640 mV
    FPC 1.0v                    1640 mV
    FPC 1.1v                    1640 mV
    FPC 1.35v                   1640 mV
    FPC VDDQ 1.5v               1640 mV
    FPC 1.8v                    1640 mV
    FPC 3.3v                    3280 mV
    FPC 5.0v bias               5143 mV
    FPC 5.0v usb                5143 mV
    FPC VCC 12.0v              12289 mV
    FPC Vref 3.3v               3280 mV
    MAIN 12.0v-i                2265 mA
```

show chassis environment fpc 0 (PTX5000 Packet Transport Router)

```
user@host> show chassis environment fpc 0
FPC 0 status:
```

```
State                       Online
PMB Temperature             35 degrees C / 95 degrees F
Intake Temperature          33 degrees C / 91 degrees F
Exhaust A Temperature       51 degrees C / 123 degrees F
Exhaust B Temperature       43 degrees C / 109 degrees F
TL0 Temperature             48 degrees C / 118 degrees F
TQ0 Temperature             53 degrees C / 127 degrees F
TL1 Temperature             56 degrees C / 132 degrees F
TQ1 Temperature             58 degrees C / 136 degrees F
TL2 Temperature             55 degrees C / 131 degrees F
TQ2 Temperature             57 degrees C / 134 degrees F
TL3 Temperature             59 degrees C / 138 degrees F
TQ3 Temperature             59 degrees C / 138 degrees F
Power
  PMB   1.05v               1049 mV
  PMB   1.5v                1500 mV
  PMB   2.5v                2500 mV
  PMB   3.3v                3299 mV
  PFE0  1.5v                1500 mV
  PFE0  1.0v                 999 mV
  TQ0   0.9v                 900 mV
  TL0   0.9v                 900 mV
  PFE1  1.5v                1499 mV
  PFE1  1.0v                 999 mV
  TQ1   0.9v                 899 mV
  TL1   0.9v                 900 mV
  PFE2  1.5v                1500 mV
  PFE2  1.0v                1000 mV
  TQ2   0.9v                 900 mV
  TL2   0.9v                 900 mV
  PFE3  1.5v                1499 mV
  PFE3  1.0v                1000 mV
  TQ3   0.9v                 900 mV
  TL3   0.9v                 900 mV
  Bias  3.3v                3327 mV
  FPC   3.3v                3300 mV
  FPC   2.5v                2500 mV
  SAM   0.9v                 900 mV
  A    12.0v                2014 mV
  B    12.0v                2030 mV
```

show chassis environment fpc 07 (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis environment fpc 07
FPC 7 status:
  State                      Online
  PMB TEMP0 Temperature      32 degrees C / 89 degrees F
  PMB TEMP1 Temperature      28 degrees C / 82 degrees F
  PMB CPU Temperature        46 degrees C / 114 degrees F
  Intake Temperature         35 degrees C / 95 degrees F
  Exhaust A Temperature      55 degrees C / 131 degrees F
  Exhaust B Temperature      54 degrees C / 129 degrees F
  TL5 Temperature            59 degrees C / 138 degrees F
  TQ5 Temperature            57 degrees C / 134 degrees F
  TL6 Temperature            57 degrees C / 134 degrees F
  TQ6 Temperature            51 degrees C / 123 degrees F
  TL1 Temperature            76 degrees C / 168 degrees F
  TQ1 Temperature            58 degrees C / 136 degrees F
  TL2 Temperature            75 degrees C / 167 degrees F
  TQ2 Temperature            57 degrees C / 134 degrees F
  TL4 Temperature            52 degrees C / 125 degrees F
  TQ4 Temperature            66 degrees C / 150 degrees F
  TL7 Temperature            52 degrees C / 125 degrees F
  TQ7 Temperature            60 degrees C / 140 degrees F
  TL0 Temperature            72 degrees C / 161 degrees F
  TQ0 Temperature            73 degrees C / 163 degrees F
  TL3 Temperature            64 degrees C / 147 degrees F
  TQ3 Temperature            70 degrees C / 158 degrees F
  Power
    PMB   1.05v               1049 mV
    PMB   3.3v                3299 mV
    PMB   1.1v-a              1100 mV
    PMB   1.5v                1499 mV
    PMB   1.1v-b              1100 mV
    Base  3.3v                3300 mV
    FPC Base  2.5v            2499 mV
    TL1   0.9v                 897 mV
    TQ1   0.9v                 897 mV
    PFE1  1.0v                 999 mV
    PFE1  1.5v                1499 mV
    TL2   0.9v                 897 mV
    TQ2   0.9v                 897 mV
    PFE2  1.0v                 999 mV
```

```
    PFE2  1.5v                1499 mV
    FPC Base  1.0v            1000 mV
    FPC Base  1.2v            1199 mV
    TL5   0.9v                 898 mV
    TQ5   0.9v                 898 mV
    PFE5  1.0v                1000 mV
    PFE5  1.5v                1500 mV
    TL6   0.9v                 897 mV
    TQ6   0.9v                 897 mV
    PFE6  1.0v                1000 mV
    PFE6  1.5v                1499 mV
    Mezz Base  2.5v           2500 mV
    TL0   0.9v                 896 mV
    TQ0   0.9v                 896 mV
    PFE0  1.0v                 999 mV
    PFE0  1.5v                1499 mV
```

show chassis environment fpc (PTX10008 router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                      Online
  FPC 0 Intake-A Temp Sensor 37 degrees C / 98 degrees F
  FPC 0 Intake-B Temp Sensor 34 degrees C / 93 degrees F
  FPC 0 Exhaust-A Temp Sensor37 degrees C / 98 degrees F
  FPC 0 Exhaust-B Temp Sensor38 degrees C / 100 degrees F
  FPC 0 Exhaust-C Temp Sensor40 degrees C / 104 degrees F
  FPC 0 PE0 Temp Sensor      41 degrees C / 105 degrees F
  FPC 0 PE1 Temp Sensor      42 degrees C / 107 degrees F
  FPC 0 PE2 Temp Sensor      44 degrees C / 111 degrees F
  FPC 0 LCPU Temp Sensor     40 degrees C / 104 degrees F
  Power
         PE0 Core 0.9V                872 mV   28777 mA   25146 mW
         PE0 HMC0 Core 0.9V           899 mV   10359 mA    9328 mW
         PE1 Core 0.9V                896 mV   29476 mA   26414 mW
         PE1 HMC0 Core 0.9V           899 mV   10218 mA    9187 mW
         PE2 Core 0.9V                872 mV   28839 mA   25199 mW
         PE2 HMC0 Core 0.9V           900 mV   10296 mA    9265 mW
         PE0 Serdes 1.0V             1020 mV   29000 mA   29593 mW
         PE1 Serdes 1.0V             1019 mV   29109 mA   29718 mW
         PE2 Serdes 1.0V             1019 mV   28484 mA   29078 mW
```

```
        LCPU Platform 1.1V          1099 mV    3515 mA    3867 mW
        LCPU Core 1.0V              1000 mV    8750 mA    8703 mW
        PHY VDD B 1.0V              1000 mV   17062 mA   17031 mW
        PHY VDD A 1.0V               999 mV   15640 mA   15625 mW
        BCM Core 1.0V                999 mV    7054 mA    7054 mW
        BCM PEX 1.0V                 999 mV    3562 mA    3558 mW
        HMC Core 1.2V               1199 mV    1280 mA    1513 mW
        HMC Serdes 1.2V             1199 mV   32937 mA   39500 mW
        VDD 1.5V                    1500 mV    2824 mA    4234 mW
        VDD 2.5V                    2449 mV    3812 mA    9343 mW
        VDD 3.3V                    3299 mV    5085 mA   16796 mW
        12V                        12259 mV   29609 mA  368196 mW
FPC 1 status:
  State                      Online
  FPC 1 Intake-A Temp Sensor 37 degrees C / 98 degrees F
  FPC 1 Intake-B Temp Sensor 34 degrees C / 93 degrees F
  FPC 1 Exhaust-A Temp Sensor38 degrees C / 100 degrees F
  FPC 1 Exhaust-B Temp Sensor38 degrees C / 100 degrees F
  FPC 1 Exhaust-C Temp Sensor40 degrees C / 104 degrees F
  FPC 1 PE0 Temp Sensor      41 degrees C / 105 degrees F
  FPC 1 PE1 Temp Sensor      42 degrees C / 107 degrees F
  FPC 1 PE2 Temp Sensor      44 degrees C / 111 degrees F
  FPC 1 LCPU Temp Sensor     39 degrees C / 102 degrees F
  Power
        PE0 Core 0.9V                898 mV   29351 mA   26421 mW
        PE0 HMC0 Core 0.9V           899 mV    9734 mA    8750 mW
        PE1 Core 0.9V                873 mV   28539 mA   24933 mW
        PE1 HMC0 Core 0.9V           899 mV    9937 mA    8937 mW
        PE2 Core 0.9V                875 mV   28906 mA   25316 mW
        PE2 HMC0 Core 0.9V           899 mV   10140 mA    9125 mW
        PE0 Serdes 1.0V             1019 mV   28312 mA   28890 mW
        PE1 Serdes 1.0V             1020 mV   28656 mA   29234 mW
        PE2 Serdes 1.0V             1020 mV   29437 mA   30015 mW
        LCPU Platform 1.1V          1100 mV    4617 mA    5078 mW
        LCPU Core 1.0V              1000 mV    8781 mA    8781 mW
        PHY VDD B 1.0V              1000 mV   15953 mA   15984 mW
        PHY VDD A 1.0V              1000 mV   15484 mA   15484 mW
        BCM Core 1.0V                999 mV    7945 mA    7937 mW
        BCM PEX 1.0V                 999 mV    3515 mA    3515 mW
        HMC Core 1.2V               1199 mV    1269 mA    1521 mW
        HMC Serdes 1.2V             1199 mV   33000 mA   39593 mW
        VDD 1.5V                    1500 mV    2691 mA    4062 mW
        VDD 2.5V                    2449 mV    3582 mA    8781 mW
```

```
          VDD 3.3V                    3300 mV    2563 mA    8458 mW
          12V                        12311 mV   29002 mA  357577 mW
FPC 2 status:
  State                      Online
  FPC 2 Intake-A Temp Sensor 43 degrees C / 109 degrees F
  FPC 2 Intake-B Temp Sensor 30 degrees C / 86 degrees F
  FPC 2 Exhaust-A Temp Sensor50 degrees C / 122 degrees F
  FPC 2 Exhaust-B Temp Sensor52 degrees C / 125 degrees F
  FPC 2 Exhaust-C Temp Sensor51 degrees C / 123 degrees F
  FPC 2 PE0 Temp Sensor      48 degrees C / 118 degrees F
  FPC 2 PE1 Temp Sensor      56 degrees C / 132 degrees F
  FPC 2 PE2 Temp Sensor      48 degrees C / 118 degrees F
  FPC 2 PE3 Temp Sensor      57 degrees C / 134 degrees F
  FPC 2 PE4 Temp Sensor      48 degrees C / 118 degrees F
  FPC 2 PE5 Temp Sensor      60 degrees C / 140 degrees F
  FPC 2 LCPU Temp Sensor     47 degrees C / 116 degrees F
  Power
          PE0 Core 0.9V                874 mV   28117 mA   24617 mW
          PE1 Core 0.9V                899 mV   29601 mA   26632 mW
          PE0 Serdes 1.0V             1019 mV   41031 mA   41843 mW
          PE1 Serdes 1.0V             1019 mV   35656 mA   36343 mW
          PE0 HMC Core 0.9V            899 mV    8125 mA    7312 mW
          PE0,1 HMC Memory 1.2V       1199 mV     565 mA     688 mW
          PE1 HMC Core 0.9V            899 mV    7921 mA    7125 mW
          PE0,1 HMC Serdes 1.2V       1199 mV   21281 mA   25562 mW
          PE2 Core 0.9V                899 mV   29187 mA   26242 mW
          PE3 Core 0.9V                899 mV   29976 mA   27074 mW
          PE2 Serdes 1.0V             1019 mV   38562 mA   39343 mW
          PE3 Serdes 1.0V             1019 mV   34937 mA   35656 mW
          PE2 HMC Core 0.9V            899 mV    8093 mA    7281 mW
          PE2,3 HMC Memory 1.2V       1199 mV     610 mA     732 mW
          PE3 HMC Core 0.9V            899 mV    7710 mA    6937 mW
          PE2,3 HMC Serdes 1.2V       1199 mV   21500 mA   25812 mW
          VDD 3.3V                    3300 mV    7937 mA   26187 mW
          VDD 1.5V                    1499 mV    3234 mA    4851 mW
          VDD 2.5V                    2449 mV    4539 mA   11109 mW
          PE4 Core 0.9V                874 mV   29914 mA   26183 mW
          PE5 Core 0.9V                874 mV   29820 mA   26031 mW
          PE4 Serdes 1.0V             1020 mV   43968 mA   44843 mW
          PE5 Serdes 1.0V             1019 mV   27453 mA   28031 mW
          PE4 HMC Core 0.9V            900 mV    7937 mA    7140 mW
          PE4,5 HMC Memory 1.2V       1200 mV    1185 mA    1421 mW
          PE5 HMC Core 0.9V            899 mV    8718 mA    7843 mW
```

```
        PE4,5 HMC Serdes 1.2V       1199 mV   21125 mA   25343 mW
        LCPU platform 1.1V          1099 mV    3777 mA    4156 mW
        LCPU core 1.0V              1000 mV    9062 mA    9062 mW
        BCM core 1.0V               1000 mV    9328 mA    9328 mW
        BCM & PEX Serdes 1.0V        999 mV    4125 mA    4125 mW
        12V                        12311 mV   53347 mA  660345 mW
FPC 3 status:
  State                      Online
  FPC 3 Intake-A Temp Sensor 43 degrees C / 109 degrees F
  FPC 3 Intake-B Temp Sensor 30 degrees C / 86 degrees F
  FPC 3 Exhaust-A Temp Sensor48 degrees C / 118 degrees F
  FPC 3 Exhaust-B Temp Sensor49 degrees C / 120 degrees F
  FPC 3 Exhaust-C Temp Sensor47 degrees C / 116 degrees F
  FPC 3 PE0 Temp Sensor      48 degrees C / 118 degrees F
  FPC 3 PE1 Temp Sensor      55 degrees C / 131 degrees F
  FPC 3 PE2 Temp Sensor      47 degrees C / 116 degrees F
  FPC 3 PE3 Temp Sensor      54 degrees C / 129 degrees F
  FPC 3 PE4 Temp Sensor      48 degrees C / 118 degrees F
  FPC 3 PE5 Temp Sensor      58 degrees C / 136 degrees F
  FPC 3 LCPU Temp Sensor     46 degrees C / 114 degrees F
  Power
        PE0 Core 0.9V                899 mV   29695 mA   26718 mW
        PE1 Core 0.9V                899 mV   29695 mA   26710 mW
        PE0 Serdes 1.0V             1020 mV   40156 mA   40906 mW
        PE1 Serdes 1.0V             1020 mV   35281 mA   35968 mW
        PE0 HMC Core 0.9V            900 mV    7492 mA    6742 mW
        PE0,1 HMC Memory 1.2V       1199 mV     569 mA     683 mW
        PE1 HMC Core 0.9V            899 mV    7570 mA    6812 mW
        PE0,1 HMC Serdes 1.2V       1199 mV   20562 mA   24656 mW
        PE2 Core 0.9V                899 mV   29734 mA   26765 mW
        PE3 Core 0.9V                900 mV   29960 mA   26968 mW
        PE2 Serdes 1.0V             1019 mV   37718 mA   38500 mW
        PE3 Serdes 1.0V             1020 mV   35250 mA   35937 mW
        PE2 HMC Core 0.9V            899 mV    7750 mA    6976 mW
        PE2,3 HMC Memory 1.2V       1200 mV     546 mA     656 mW
        PE3 HMC Core 0.9V            899 mV    7718 mA    6945 mW
        PE2,3 HMC Serdes 1.2V       1199 mV   20625 mA   24750 mW
        VDD 3.3V                    3299 mV    5917 mA   19515 mW
        VDD 1.5V                    1499 mV    4015 mA    6015 mW
        VDD 2.5V                    2449 mV    4335 mA   10625 mW
        PE4 Core 0.9V                899 mV   29835 mA   26875 mW
        PE5 Core 0.9V                924 mV   30554 mA   28277 mW
        PE4 Serdes 1.0V             1019 mV   43281 mA   44187 mW
```

```
          PE5 Serdes 1.0V             1020 mV   27140 mA   27703 mW
          PE4 HMC Core 0.9V            899 mV    7476 mA    6726 mW
          PE4,5 HMC Memory 1.2V       1199 mV     531 mA     637 mW
          PE5 HMC Core 0.9V            899 mV    7539 mA    6781 mW
          PE4,5 HMC Serdes 1.2V       1199 mV   20375 mA   24468 mW
          LCPU platform 1.1V          1099 mV    3453 mA    3796 mW
          LCPU core 1.0V               999 mV    8984 mA    8984 mW
          BCM core 1.0V                999 mV    7929 mA    7921 mW
          BCM & PEX Serdes 1.0V       1000 mV    4046 mA    4046 mW
          12V                        12351 mV   51918 mA  644880 mW
FPC 5 status:
  State                      Online
  FPC 5 Intake-A Temp Sensor Failed
  FPC 5 Intake-B Temp Sensor Failed
  FPC 5 Exhaust-A Temp Sensor41 degrees C / 105 degrees F
  FPC 5 Exhaust-B Temp Sensor41 degrees C / 105 degrees F
  FPC 5 Exhaust-C Temp Sensor42 degrees C / 107 degrees F
  FPC 5 PE0 Temp Sensor      47 degrees C / 116 degrees F
  FPC 5 PE1 Temp Sensor      49 degrees C / 120 degrees F
  FPC 5 PE2 Temp Sensor      53 degrees C / 127 degrees F
  FPC 5 LCPU Temp Sensor     Failed
  Power
          PE0 Core 0.9V                923 mV   30976 mA   28578 mW
          PE0 HMC0 Core 0.9V           899 mV   10093 mA    9078 mW
          PE1 Core 0.9V                897 mV   29398 mA   26414 mW
          PE1 HMC0 Core 0.9V           899 mV    9734 mA    8750 mW
          PE2 Core 0.9V                922 mV   30226 mA   27886 mW
          PE2 HMC0 Core 0.9V           899 mV    9984 mA    8968 mW
          PE0 Serdes 1.0V             1019 mV   29296 mA   29890 mW
          PE1 Serdes 1.0V             1020 mV   28687 mA   29296 mW
          PE2 Serdes 1.0V             1020 mV   28187 mA   28765 mW
          LCPU Platform 1.1V          1100 mV    3664 mA    4031 mW
          LCPU Core 1.0V               999 mV    9125 mA    9125 mW
          PHY VDD B 1.0V               999 mV   15593 mA   15593 mW
          PHY VDD A 1.0V              1000 mV   15453 mA   15453 mW
          BCM Core 1.0V                999 mV    7773 mA    7765 mW
          BCM PEX 1.0V                1000 mV    3460 mA    3464 mW
          HMC Core 1.2V               1199 mV    1328 mA    1628 mW
          HMC Serdes 1.2V             1199 mV   32203 mA   38625 mW
          VDD 1.5V                    1499 mV    2675 mA    4007 mW
          VDD 2.5V                    2450 mV    3675 mA    9000 mW
          VDD 3.3V                    3300 mV    1814 mA    5980 mW
          12V                        12272 mV   29045 mA  361369 mW
```

```
FPC 6 status:
  State                      Online
  FPC 6 Intake-A Temp Sensor 41 degrees C / 105 degrees F
  FPC 6 Intake-B Temp Sensor 37 degrees C / 98 degrees F
  FPC 6 Exhaust-A Temp Sensor40 degrees C / 104 degrees F
  FPC 6 Exhaust-B Temp Sensor40 degrees C / 104 degrees F
  FPC 6 Exhaust-C Temp Sensor40 degrees C / 104 degrees F
  FPC 6 PE0 Temp Sensor      45 degrees C / 113 degrees F
  FPC 6 PE1 Temp Sensor      47 degrees C / 116 degrees F
  FPC 6 PE2 Temp Sensor      51 degrees C / 123 degrees F
  FPC 6 LCPU Temp Sensor     41 degrees C / 105 degrees F
  Power
         PE0 Core 0.9V                897 mV   30214 mA   27179 mW
         PE0 HMC0 Core 0.9V           899 mV   10000 mA    8984 mW
         PE1 Core 0.9V                873 mV   29332 mA   25601 mW
         PE1 HMC0 Core 0.9V           899 mV    9828 mA    8828 mW
         PE2 Core 0.9V                898 mV   30781 mA   27675 mW
         PE2 HMC0 Core 0.9V           899 mV   10328 mA    9296 mW
         PE0 Serdes 1.0V             1019 mV   28921 mA   29531 mW
         PE1 Serdes 1.0V             1020 mV   29437 mA   30046 mW
         PE2 Serdes 1.0V             1019 mV   29671 mA   30281 mW
         LCPU Platform 1.1V          1100 mV    3671 mA    4039 mW
         LCPU Core 1.0V              1000 mV    8218 mA    8187 mW
         PHY VDD B 1.0V              1000 mV   15984 mA   15984 mW
         PHY VDD A 1.0V               999 mV   16093 mA   16093 mW
         BCM Core 1.0V               1000 mV    8046 mA    8062 mW
         BCM PEX 1.0V                1000 mV    3500 mA    3500 mW
         HMC Core 1.2V               1199 mV    1327 mA    1579 mW
         HMC Serdes 1.2V             1199 mV   33031 mA   39593 mW
         VDD 1.5V                    1499 mV    2722 mA    4078 mW
         VDD 2.5V                    2449 mV    3539 mA    8671 mW
         VDD 3.3V                    3299 mV    8082 mA   26656 mW
         12V                        12311 mV   31124 mA  385270 mW
```

show chassis environment fpc (PTX10016 router)

```
user@host> show chassis environment fpc
FPC 1 status:
  State                      Online
  FPC 1 Intake-A Temp Sensor 36 degrees C / 96 degrees F
  FPC 1 Intake-B Temp Sensor 32 degrees C / 89 degrees F
```

```
  FPC 1 Exhaust-A Temp Sensor37 degrees C / 98 degrees F
  FPC 1 Exhaust-B Temp Sensor36 degrees C / 96 degrees F
  FPC 1 Exhaust-C Temp Sensor36 degrees C / 96 degrees F
  FPC 1 PE0 Temp Sensor      45 degrees C / 113 degrees F
  FPC 1 PE1 Temp Sensor      46 degrees C / 114 degrees F
  FPC 1 PE2 Temp Sensor      53 degrees C / 127 degrees F
  FPC 1 LCPU Temp Sensor     35 degrees C / 95 degrees F
  Power
         PE0 Core 0.9V                897 mV   28992 mA   26027 mW
         PE0 HMC0 Core 0.9V           899 mV   10156 mA    9156 mW
         PE1 Core 0.9V                871 mV   28800 mA   25164 mW
         PE1 HMC0 Core 0.9V           899 mV   10125 mA    9109 mW
         PE2 Core 0.9V                898 mV   29914 mA   26906 mW
         PE2 HMC0 Core 0.9V           899 mV   10343 mA    9296 mW
         PE0 Serdes 1.0V             1019 mV   27515 mA   28093 mW
         PE1 Serdes 1.0V             1020 mV   27968 mA   28546 mW
         PE2 Serdes 1.0V             1019 mV   27796 mA   28359 mW
         LCPU Platform 1.1V          1100 mV    3347 mA    3289 mW
         LCPU Core 1.0V              1000 mV    7960 mA    7960 mW
         PHY VDD B 1.0V              1000 mV   16437 mA   16437 mW
         PHY VDD A 1.0V               999 mV   15656 mA   15656 mW
         BCM Core 1.0V               1000 mV    7289 mA    7335 mW
         BCM PEX 1.0V                 999 mV    3453 mA    3453 mW
         HMC Core 1.2V               1199 mV    1218 mA    1453 mW
         HMC Serdes 1.2V             1199 mV   32093 mA   38562 mW
         VDD 1.5V                    1500 mV    2859 mA    4289 mW
         VDD 2.5V                    2449 mV    3875 mA    9500 mW
         VDD 3.3V                    3299 mV    2806 mA    9257 mW
         12V                        12351 mV   28569 mA  354877 mW
FPC 3 status:
  State                      Online
  FPC 3 Intake-A Temp Sensor 35 degrees C / 95 degrees F
  FPC 3 Intake-B Temp Sensor 31 degrees C / 87 degrees F
  FPC 3 Exhaust-A Temp Sensor36 degrees C / 96 degrees F
  FPC 3 Exhaust-B Temp Sensor34 degrees C / 93 degrees F
  FPC 3 Exhaust-C Temp Sensor33 degrees C / 91 degrees F
  FPC 3 PE0 Temp Sensor      43 degrees C / 109 degrees F
  FPC 3 PE1 Temp Sensor      45 degrees C / 113 degrees F
  FPC 3 PE2 Temp Sensor      49 degrees C / 120 degrees F
  FPC 3 LCPU Temp Sensor     35 degrees C / 95 degrees F
  Power
         PE0 Core 0.9V                897 mV   28832 mA   25871 mW
         PE0 HMC0 Core 0.9V           899 mV   10359 mA    9328 mW
```

```
        PE1 Core 0.9V                873 mV   28230 mA   24671 mW
        PE1 HMC0 Core 0.9V           899 mV   10468 mA    9421 mW
        PE2 Core 0.9V                898 mV   29539 mA   26539 mW
        PE2 HMC0 Core 0.9V           899 mV   10656 mA    9593 mW
        PE0 Serdes 1.0V             1020 mV   27484 mA   28031 mW
        PE1 Serdes 1.0V             1019 mV   27515 mA   28078 mW
        PE2 Serdes 1.0V             1020 mV   27625 mA   28187 mW
        LCPU Platform 1.1V          1099 mV    3050 mA    3355 mW
        LCPU Core 1.0V               999 mV    7820 mA    7804 mW
        PHY VDD B 1.0V               999 mV   15406 mA   15406 mW
        PHY VDD A 1.0V              1000 mV   14953 mA   14953 mW
        BCM Core 1.0V               1000 mV    7648 mA    7648 mW
        BCM PEX 1.0V                1000 mV    3531 mA    3531 mW
        HMC Core 1.2V               1200 mV    1234 mA    1476 mW
        HMC Serdes 1.2V             1199 mV   34671 mA   41593 mW
        VDD 1.5V                    1499 mV    3484 mA    5226 mW
        VDD 2.5V                    2449 mV    3218 mA    7890 mW
        VDD 3.3V                    3299 mV    2468 mA    8148 mW
        12V                        12311 mV   28785 mA  355950 mW
FPC 6 status:
  State                      Online
  FPC 6 Intake-A Temp Sensor 34 degrees C / 93 degrees F
  FPC 6 Intake-B Temp Sensor 31 degrees C / 87 degrees F
  FPC 6 Exhaust-A Temp Sensor34 degrees C / 93 degrees F
  FPC 6 Exhaust-B Temp Sensor35 degrees C / 95 degrees F
  FPC 6 Exhaust-C Temp Sensor35 degrees C / 95 degrees F
  FPC 6 PE0 Temp Sensor      42 degrees C / 107 degrees F
  FPC 6 PE1 Temp Sensor      43 degrees C / 109 degrees F
  FPC 6 PE2 Temp Sensor      47 degrees C / 116 degrees F
  FPC 6 LCPU Temp Sensor     34 degrees C / 93 degrees F
  Power
        PE0 Core 0.9V                922 mV   29394 mA   27160 mW
        PE0 HMC0 Core 0.9V           899 mV   10078 mA    9062 mW
        PE1 Core 0.9V                923 mV   29636 mA   27304 mW
        PE1 HMC0 Core 0.9V           899 mV    9890 mA    8890 mW
        PE2 Core 0.9V                898 mV   29734 mA   26757 mW
        PE2 HMC0 Core 0.9V           899 mV    9968 mA    8968 mW
        PE0 Serdes 1.0V             1020 mV   26968 mA   27515 mW
        PE1 Serdes 1.0V             1019 mV   27421 mA   27984 mW
        PE2 Serdes 1.0V             1019 mV   27625 mA   28171 mW
        LCPU Platform 1.1V          1099 mV    3230 mA    4742 mW
        LCPU Core 1.0V               999 mV    8171 mA    8171 mW
        PHY VDD B 1.0V              1000 mV   15671 mA   15687 mW
```

```
        PHY VDD A 1.0V              999 mV   15703 mA   15703 mW
        BCM Core 1.0V               999 mV    7500 mA    7492 mW
        BCM PEX 1.0V               1000 mV    3480 mA    3468 mW
        HMC Core 1.2V              1199 mV    1199 mA    1440 mW
        HMC Serdes 1.2V            1199 mV   31046 mA   37250 mW
        VDD 1.5V                   1499 mV    2804 mA    4203 mW
        VDD 2.5V                   2449 mV    3746 mA    9171 mW
        VDD 3.3V                   3300 mV    3173 mA   10476 mW
        12V                       12311 mV   28786 mA  355654 mW
FPC 8 status:
  State                     Online
  FPC 8 Intake-A Temp Sensor 34 degrees C / 93 degrees F
  FPC 8 Intake-B Temp Sensor 30 degrees C / 86 degrees F
  FPC 8 Exhaust-A Temp Sensor37 degrees C / 98 degrees F
  FPC 8 Exhaust-B Temp Sensor37 degrees C / 98 degrees F
  FPC 8 Exhaust-C Temp Sensor37 degrees C / 98 degrees F
  FPC 8 PE0 Temp Sensor      42 degrees C / 107 degrees F
  FPC 8 PE1 Temp Sensor      44 degrees C / 111 degrees F
  FPC 8 PE2 Temp Sensor      47 degrees C / 116 degrees F
  FPC 8 LCPU Temp Sensor     33 degrees C / 91 degrees F
  Power
        PE0 Core 0.9V               897 mV   29382 mA   26437 mW
        PE0 HMC0 Core 0.9V          899 mV   10265 mA    9250 mW
        PE1 Core 0.9V               872 mV   28867 mA   25175 mW
        PE1 HMC0 Core 0.9V          899 mV   10171 mA    9109 mW
        PE2 Core 0.9V               899 mV   30210 mA   27214 mW
        PE2 HMC0 Core 0.9V          900 mV   10187 mA    9171 mW
        PE0 Serdes 1.0V            1020 mV   27843 mA   28421 mW
        PE1 Serdes 1.0V            1020 mV   28265 mA   28828 mW
        PE2 Serdes 1.0V            1019 mV   28406 mA   29000 mW
        LCPU Platform 1.1V         1099 mV    3000 mA    3300 mW
        LCPU Core 1.0V             1000 mV    7937 mA    7937 mW
        PHY VDD B 1.0V             1000 mV   15843 mA   15843 mW
        PHY VDD A 1.0V             1000 mV   15250 mA   15250 mW
        BCM Core 1.0V               999 mV    6914 mA    6898 mW
        BCM PEX 1.0V                999 mV    3445 mA    3445 mW
        HMC Core 1.2V              1199 mV    1162 mA    1390 mW
        HMC Serdes 1.2V            1199 mV   33437 mA   40125 mW
        VDD 1.5V                   1499 mV    2851 mA    4273 mW
        VDD 2.5V                   2450 mV    3867 mA    9484 mW
        VDD 3.3V                   3300 mV    3258 mA   10753 mW
        12V                       12338 mV   28656 mA  356171 mW
FPC 9 status:
```

```
State                      Online
FPC 9 Intake-A Temp Sensor 44 degrees C / 111 degrees F
FPC 9 Intake-B Temp Sensor 28 degrees C / 82 degrees F
FPC 9 Exhaust-A Temp Sensor51 degrees C / 123 degrees F
FPC 9 Exhaust-B Temp Sensor52 degrees C / 125 degrees F
FPC 9 Exhaust-C Temp Sensor48 degrees C / 118 degrees F
FPC 9 PE0 Temp Sensor      52 degrees C / 125 degrees F
FPC 9 PE1 Temp Sensor      65 degrees C / 149 degrees F
FPC 9 PE2 Temp Sensor      50 degrees C / 122 degrees F
FPC 9 PE3 Temp Sensor      65 degrees C / 149 degrees F
FPC 9 PE4 Temp Sensor      50 degrees C / 122 degrees F
FPC 9 PE5 Temp Sensor      67 degrees C / 152 degrees F
FPC 9 LCPU Temp Sensor     45 degrees C / 113 degrees F
Power
       PE0 Core 0.9V                875 mV   28316 mA   24808 mW
       PE1 Core 0.9V                875 mV   28546 mA   24996 mW
       PE0 Serdes 1.0V             1019 mV   38906 mA   39687 mW
       PE1 Serdes 1.0V             1020 mV   33078 mA   33781 mW
       PE0 HMC Core 0.9V            899 mV    7718 mA    6945 mW
       PE0,1 HMC Memory 1.2V       1199 mV     579 mA     695 mW
       PE1 HMC Core 0.9V            899 mV    7289 mA    6570 mW
       PE0,1 HMC Serdes 1.2V       1199 mV   20187 mA   24250 mW
       PE2 Core 0.9V                924 mV   29062 mA   26894 mW
       PE3 Core 0.9V                900 mV   28914 mA   26039 mW
       PE2 Serdes 1.0V             1020 mV   36375 mA   37093 mW
       PE3 Serdes 1.0V             1019 mV   32640 mA   33296 mW
       PE2 HMC Core 0.9V            900 mV    7695 mA    6921 mW
       PE2,3 HMC Memory 1.2V       1199 mV     562 mA     674 mW
       PE3 HMC Core 0.9V            899 mV    7554 mA    6796 mW
       PE2,3 HMC Serdes 1.2V       1199 mV   20156 mA   24218 mW
       VDD 3.3V                    3300 mV    8964 mA   29609 mW
       VDD 1.5V                    1499 mV    3968 mA    5945 mW
       VDD 2.5V                    2449 mV    4414 mA   10890 mW
       PE4 Core 0.9V                900 mV   28527 mA   25679 mW
       PE5 Core 0.9V                899 mV   28902 mA   26035 mW
       PE4 Serdes 1.0V             1019 mV   41281 mA   42125 mW
       PE5 Serdes 1.0V             1019 mV   25781 mA   26328 mW
       PE4 HMC Core 0.9V            900 mV    7382 mA    6648 mW
       PE4,5 HMC Memory 1.2V       1199 mV     626 mA     750 mW
       PE5 HMC Core 0.9V            899 mV    7562 mA    6796 mW
       PE4,5 HMC Serdes 1.2V       1199 mV   20312 mA   24375 mW
       LCPU platform 1.1V          1099 mV    3687 mA    4054 mW
       LCPU core 1.0V              1000 mV    9000 mA    9000 mW
```

```
          BCM core 1.0V                 999 mV    7843 mA    7835 mW
          BCM & PEX Serdes 1.0V         999 mV    4062 mA    4062 mW
          12V                         12417 mV   51659 mA  643215 mW
```

show chassis environment FPC 1 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis environment fpc 1
FPC 1 status:
  State                       Online
  Temperature Intake          36 degrees C / 96 degrees F
  Temperature Exhaust A       39 degrees C / 102 degrees F
  Temperature LU TSen         52 degrees C / 125 degrees F
  Temperature LU Chip         54 degrees C / 129 degrees F
  Temperature XM TSen         52 degrees C / 125 degrees F
  Temperature XM Chip         60 degrees C / 140 degrees F
  Temperature PCIe TSen       52 degrees C / 125 degrees F
  Temperature PCIe Chip       69 degrees C / 156 degrees F
  Power
    MPC-BIAS3V3-zl2106         3302 mV
    MPC-VDD3V3-zl6100          3325 mV
    MPC-AVDD1V0-zl6100         1007 mV
    MPC-PCIE_1V0-zl6100         904 mV
    MPC-LU0_1V0-zl2004          996 mV
    MPC-VDD_1V5-zl2004         1498 mV
    MPC-12VA-BMR453           11733 mV
    MPC-12VB-BMR453           11728 mV
    MPC-XM_0V9-vt273m           900 mV
  I2C Slave Revision          81
```

show chassis environment FPC (Junos OS Evolved)

```
user@switch> show chassis environment fpc
FPC 0 status:
  State                       Online
  Intake Temperature          32 degrees C / 89 degrees F
  Exhaust-A Temperature       43 degrees C / 109 degrees F
  Exhaust-B Temperature       32 degrees C / 89 degrees F
  PE0 Temperature             34 degrees C / 93 degrees F
  PE1 Temperature             38 degrees C / 100 degrees F
```

```
PE2 Temperature            38 degrees C / 100 degrees F
PE3 Temperature            36 degrees C / 96 degrees F
PE4 Temperature            35 degrees C / 95 degrees F
PE5 Temperature            35 degrees C / 95 degrees F
Power 1
  RT_1  1.0v                  1018 mV
  RT_2  1.0v                  1018 mV
Power 2
  FPC 1 1.0v                   999 mV
  FPC 2 1.0v                   998 mV
Power 3
  FPC   2.5v                  2499 mV
  FPC   3.3v                  3299 mV
Power 4
  FPC   0.9v                   899 mV
  FPC   1.5v                  1499 mV
Power 5
  PE0 1 1.0v                  1039 mV
  PE0 2 1.0v                  1039 mV
Power 6
  PE0 1 0.9v                   900 mV
  PE0 2 0.9v                   900 mV
Power 7
  PE0 3 0.9v                   902 mV
  PE0 4 0.9v                   902 mV
Power 8
  PE0 H 0.9v                   899 mV
  PE0 H 1.2v                  1199 mV
Power 9
  PE1 1 1.0v                  1040 mV
  PE1 2 1.0v                  1039 mV
Power 10
  PE1 1 0.9v                   901 mV
  PE1 2 0.9v                   901 mV
Power 11
  PE1 3 0.9v                   900 mV
  PE1 4 0.9v                   900 mV
Power 12
  PE1 H 0.9v                   899 mV
  PE1 H 1.2v                  1199 mV
Power 13
  PE2 1 1.0v                  1039 mV
  PE2 2 1.0v                  1039 mV
```

```
Power 14
  PE2 1 0.9v                  900 mV
  PE2 2 0.9v                  900 mV
Power 15
  PE2 3 0.9v                  900 mV
  PE2 4 0.9v                  900 mV
Power 16
  PE2 H 0.9v                  899 mV
  PE2 H 1.2v                 1199 mV
Power 17
  PE3 1 1.0v                 1039 mV
  PE3 2 1.0v                 1039 mV
Power 18
  PE3 1 0.9v                  899 mV
  PE3 2 0.9v                  900 mV
Power 19
  PE3 3 0.9v                  899 mV
  PE3 4 0.9v                  900 mV
Power 20
  PE3 H 0.9v                  899 mV
  PE3 H 1.2v                 1199 mV
Power 21
  PE4 1 1.0v                 1039 mV
  PE4 2 1.0v                 1039 mV
Power 22
  PE4 1 0.9v                  900 mV
  PE4 2 0.9v                  900 mV
Power 23
  PE4 3 0.9v                  901 mV
  PE4 4 0.9v                  901 mV
Power 24
  PE4 H 0.9v                  899 mV
  PE4 H 1.2v                 1199 mV
Power 25
  PE5 1 1.0v                 1040 mV
  PE5 2 1.0v                 1039 mV
Power 26
  PE5 1 0.9v                  901 mV
  PE5 2 0.9v                  901 mV
Power 27
  PE5 3 0.9v                  901 mV
  PE5 4 0.9v                  901 mV
Power 28
```

```
    PE5 H 0.9v                   899 mV
    PE5 H 1.2v                  1199 mV
  Power 29
    PIC0 12.0v                 12342 mV
  Power 30
    PIC1 12.0v                 12342 mV
  Power 31
    A    12.0v                 12375 mV
    B    12.0v                  1008 mV
  Bus Revision               115
FPC 1 status:
  State                      Online
  Intake Temperature         33 degrees C / 91 degrees F
  Exhaust-A Temperature      44 degrees C / 111 degrees F
  Exhaust-B Temperature      33 degrees C / 91 degrees F
  PE0 Temperature            34 degrees C / 93 degrees F
  PE1 Temperature            38 degrees C / 100 degrees F
  PE2 Temperature            37 degrees C / 98 degrees F
  PE3 Temperature            36 degrees C / 96 degrees F
  PE4 Temperature            34 degrees C / 93 degrees F
  PE5 Temperature            36 degrees C / 96 degrees F
  Power 1
    RT_1  1.0v                   1018 mV
    RT_2  1.0v                   1018 mV
  Power 2
    FPC 1 1.0v                    999 mV
    FPC 2 1.0v                    999 mV
  Power 3
    FPC   2.5v                   2499 mV
    FPC   3.3v                   3300 mV
  Power 4
    FPC   0.9v                    899 mV
    FPC   1.5v                   1500 mV
  Power 5
    PE0 1 1.0v                   1039 mV
    PE0 2 1.0v                   1039 mV
  Power 6
    PE0 1 0.9v                    925 mV
    PE0 2 0.9v                    925 mV
  Power 7
    PE0 3 0.9v                    925 mV
    PE0 4 0.9v                    926 mV
  Power 8
```

```
    PE0 H 0.9v                    899 mV
    PE0 H 1.2v                   1199 mV
  Power 9
    PE1 1 1.0v                   1040 mV
    PE1 2 1.0v                   1039 mV
  Power 10
    PE1 1 0.9v                    900 mV
    PE1 2 0.9v                    901 mV
  Power 11
    PE1 3 0.9v                    899 mV
    PE1 4 0.9v                    900 mV
  Power 12
    PE1 H 0.9v                    899 mV
    PE1 H 1.2v                   1199 mV
  Power 13
    PE2 1 1.0v                   1040 mV
    PE2 2 1.0v                   1039 mV
  Power 14
    PE2 1 0.9v                    926 mV
    PE2 2 0.9v                    926 mV
  Power 15
    PE2 3 0.9v                    927 mV
    PE2 4 0.9v                    927 mV
  Power 16
    PE2 H 0.9v                    899 mV
    PE2 H 1.2v                   1199 mV
  Power 17
    PE3 1 1.0v                   1039 mV
    PE3 2 1.0v                   1039 mV
  Power 18
    PE3 1 0.9v                    926 mV
    PE3 2 0.9v                    927 mV
  Power 19
    PE3 3 0.9v                    925 mV
    PE3 4 0.9v                    926 mV
  Power 20
    PE3 H 0.9v                    899 mV
    PE3 H 1.2v                   1199 mV
  Power 21
    PE4 1 1.0v                   1039 mV
    PE4 2 1.0v                   1040 mV
  Power 22
    PE4 1 0.9v                    925 mV
```

```
    PE4 2 0.9v                  925 mV
  Power 23
    PE4 3 0.9v                  925 mV
    PE4 4 0.9v                  926 mV
  Power 24
    PE4 H 0.9v                  900 mV
    PE4 H 1.2v                 1199 mV
  Power 25
    PE5 1 1.0v                 1039 mV
    PE5 2 1.0v                 1039 mV
  Power 26
    PE5 1 0.9v                  898 mV
    PE5 2 0.9v                  899 mV
  Power 27
    PE5 3 0.9v                  900 mV
    PE5 4 0.9v                  900 mV
  Power 28
    PE5 H 0.9v                  899 mV
    PE5 H 1.2v                 1199 mV
  Power 29
    PIC0 12.0v                    0 mV
  Power 30
    PIC1 12.0v                12402 mV
  Power 31
    A    12.0v                12344 mV
    B    12.0v                 1008 mV
  Bus Revision               115
FPC 2 status:
  State                      Online
  Intake Temperature         31 degrees C / 87 degrees F
  Exhaust-A Temperature      38 degrees C / 100 degrees F
  Exhaust-B Temperature      28 degrees C / 82 degrees F
  PE0 Temperature            28 degrees C / 82 degrees F
  PE1 Temperature            33 degrees C / 91 degrees F
  PE2 Temperature            34 degrees C / 93 degrees F
  PE3 Temperature            31 degrees C / 87 degrees F
  Power 1
    RT_1  1.0v                 1018 mV
    RT_2  1.0v                 1018 mV
  Power 2
    FPC 1 1.0v                  999 mV
    FPC 2 1.0v                  999 mV
  Power 3
```

```
    FPC   2.5v                  2499 mV
    FPC   3.3v                  3299 mV
  Power 4
    FPC   0.9v                   899 mV
    FPC   1.5v                  1500 mV
  Power 5
    PE0 1 1.0v                  1039 mV
    PE0 2 1.0v                  1040 mV
  Power 6
    PE0 1 0.9v                   900 mV
    PE0 2 0.9v                   901 mV
  Power 7
    PE0 3 0.9v                   900 mV
    PE0 4 0.9v                   900 mV
  Power 8
    PE0 H 0.9v                   899 mV
    PE0 H 1.2v                  1199 mV
  Power 9
    PE1 1 1.0v                  1039 mV
    PE1 2 1.0v                  1039 mV
  Power 10
    PE1 1 0.9v                   875 mV
    PE1 2 0.9v                   876 mV
  Power 11
    PE1 3 0.9v                   875 mV
    PE1 4 0.9v                   875 mV
  Power 12
    PE1 H 0.9v                   899 mV
    PE1 H 1.2v                  1199 mV
  Power 13
    PE2 1 1.0v                  1039 mV
    PE2 2 1.0v                  1039 mV
  Power 14
    PE2 1 0.9v                   900 mV
    PE2 2 0.9v                   900 mV
  Power 15
    PE2 3 0.9v                   900 mV
    PE2 4 0.9v                   900 mV
  Power 16
    PE2 H 0.9v                   899 mV
    PE2 H 1.2v                  1199 mV
  Power 17
    PE3 1 1.0v                  1039 mV
```

```
    PE3 2 1.0v                  1039 mV
  Power 18
    PE3 1 0.9v                   875 mV
    PE3 2 0.9v                   875 mV
  Power 19
    PE3 3 0.9v                   875 mV
    PE3 4 0.9v                   875 mV
  Power 20
    PE3 H 0.9v                   899 mV
    PE3 H 1.2v                  1200 mV
  Power 21
    PIC0 12.0v                 12281 mV
  Power 22
    PIC1 12.0v                     0 mV
  Power 23
    A    12.0v                 12406 mV
    B    12.0v                  1006 mV
  Bus Revision               115
FPC 3 status:
  State                      Online
  Intake Temperature         33 degrees C / 91 degrees F
  Exhaust-A Temperature      44 degrees C / 111 degrees F
  Exhaust-B Temperature      30 degrees C / 86 degrees F
  PE0 Temperature            33 degrees C / 91 degrees F
  PE1 Temperature            37 degrees C / 98 degrees F
  PE2 Temperature            38 degrees C / 100 degrees F
  PE3 Temperature            34 degrees C / 93 degrees F
  PE4 Temperature            33 degrees C / 91 degrees F
  PE5 Temperature            36 degrees C / 96 degrees F
  Power 1
    RT_1  1.0v                  1018 mV
    RT_2  1.0v                  1018 mV
  Power 2
    FPC 1 1.0v                   999 mV
    FPC 2 1.0v                   999 mV
  Power 3
    FPC   2.5v                  2500 mV
    FPC   3.3v                  3299 mV
  Power 4
    FPC   0.9v                   899 mV
    FPC   1.5v                  1500 mV
  Power 5
    PE0 1 1.0v                  1039 mV
```

```
  PE0 2 1.0v                  1039 mV
Power 6
  PE0 1 0.9v                   900 mV
  PE0 2 0.9v                   900 mV
Power 7
  PE0 3 0.9v                   898 mV
  PE0 4 0.9v                   899 mV
Power 8
  PE0 H 0.9v                   899 mV
  PE0 H 1.2v                  1199 mV
Power 9
  PE1 1 1.0v                  1040 mV
  PE1 2 1.0v                  1039 mV
Power 10
  PE1 1 0.9v                   926 mV
  PE1 2 0.9v                   926 mV
Power 11
  PE1 3 0.9v                   925 mV
  PE1 4 0.9v                   925 mV
Power 12
  PE1 H 0.9v                   900 mV
  PE1 H 1.2v                  1199 mV
Power 13
  PE2 1 1.0v                  1039 mV
  PE2 2 1.0v                  1039 mV
Power 14
  PE2 1 0.9v                   873 mV
  PE2 2 0.9v                   873 mV
Power 15
  PE2 3 0.9v                   875 mV
  PE2 4 0.9v                   875 mV
Power 16
  PE2 H 0.9v                   899 mV
  PE2 H 1.2v                  1199 mV
Power 17
  PE3 1 1.0v                  1039 mV
  PE3 2 1.0v                  1039 mV
Power 18
  PE3 1 0.9v                   899 mV
  PE3 2 0.9v                   900 mV
Power 19
  PE3 3 0.9v                   899 mV
  PE3 4 0.9v                   899 mV
```

```
  Power 20
    PE3 H 0.9v                   899 mV
    PE3 H 1.2v                  1199 mV
  Power 21
    PE4 1 1.0v                  1040 mV
    PE4 2 1.0v                  1040 mV
  Power 22
    PE4 1 0.9v                   949 mV
    PE4 2 0.9v                   950 mV
  Power 23
    PE4 3 0.9v                   950 mV
    PE4 4 0.9v                   951 mV
  Power 24
    PE4 H 0.9v                   899 mV
    PE4 H 1.2v                  1199 mV
  Power 25
    PE5 1 1.0v                  1039 mV
    PE5 2 1.0v                  1039 mV
  Power 26
    PE5 1 0.9v                   900 mV
    PE5 2 0.9v                   900 mV
  Power 27
    PE5 3 0.9v                   900 mV
    PE5 4 0.9v                   900 mV
  Power 28
    PE5 H 0.9v                   899 mV
    PE5 H 1.2v                  1199 mV
  Power 29
    PIC0 12.0v                     0 mV
  Power 30
    PIC1 12.0v                     0 mV
  Power 31
    A    12.0v                 12406 mV
    B    12.0v                  1008 mV
  Bus Revision               115
FPC 6 status:
  State                      Onlining
  Bus Revision               115
```

show chassis environment FPC (ACX Series)

```
user@host> show chassis environment fpc

FPC 0 status:
  State                      Online
  et-0/0/0                               36 degrees C / 96 degrees F
  et-0/0/1                               36 degrees C / 96 degrees F
  et-0/0/2                               36 degrees C / 96 degrees F
  et-0/0/3                               39 degrees C / 102 degrees F
  et-0/0/4                               37 degrees C / 98 degrees F
  et-0/0/5                               34 degrees C / 93 degrees F
  et-0/0/6                               37 degrees C / 98 degrees F
  et-0/0/7                               40 degrees C / 104 degrees F
  et-0/0/8                               34 degrees C / 93 degrees F
  et-0/0/9                               34 degrees C / 93 degrees F
  et-0/0/10                              40 degrees C / 104 degrees F
  et-0/0/12                              37 degrees C / 98 degrees F
  et-0/0/20                              39 degrees C / 102 degrees F
  et-0/0/23                              39 degrees C / 102 degrees F
  et-0/0/24                              36 degrees C / 96 degrees F
  et-0/0/36                              36 degrees C / 96 degrees F
  et-0/0/37                              36 degrees C / 96 degrees F
  et-0/0/38                              37 degrees C / 98 degrees F
  et-0/0/39                              33 degrees C / 91 degrees F
  et-0/0/40                              38 degrees C / 100 degrees F
  et-0/0/41                              38 degrees C / 100 degrees F
  et-0/0/42                              37 degrees C / 98 degrees F
  et-0/0/43                              36 degrees C / 96 degrees F
  et-0/0/48                              36 degrees C / 96 degrees F
  et-0/0/49                              37 degrees C / 98 degrees F
  et-0/0/50                              34 degrees C / 93 degrees F
  et-0/0/51                              35 degrees C / 95 degrees F
  et-0/0/53                              30 degrees C / 86 degrees F
  Bus Revision               0
FPC 1 status:
  State                      Online
  PHY Die 0 Sensor 0 Temp                54 degrees C / 129 degrees F
  PHY Die 1 Sensor 0 Temp                48 degrees C / 118 degrees F
  et-1/0/1                               37 degrees C / 98 degrees F
  et-1/0/2                               34 degrees C / 93 degrees F
  et-1/0/3                               33 degrees C / 91 degrees F
```

```
  et-1/0/5                                28 degrees C / 82 degrees F
  Exhaust Top Temp                        38 degrees C / 100 degrees F
  Exhaust Bottom Temp                     41 degrees C / 105 degrees F
  Intake Bottom Temp                      36 degrees C / 96 degrees F
  Intake Top Temp                         41 degrees C / 105 degrees F
  Hotswap Controller 0 Temp               42 degrees C / 107 degrees F
  PHY0_1_VDD0V8 Temp                      49 degrees C / 120 degrees F
  VDD3V3 Temp                             51 degrees C / 123 degrees F
  VDD1V2 Temp                             56 degrees C / 132 degrees F
  Power 1
    BIAS3V3                     3352 mV
    VDD1V8                      1814 mV
    VDD2V5                      2538 mV
  Power 2
    Hotswap 0, 12V0            12093 mV
  Power 3
    PHY0_1_VDD0V8                798 mV
  Power 4
    VDD3V3                      3315 mV
  Power 5
    VDD1V2                      1206 mV
  Bus Revision               27
FPC 2 status:
  State                      Online
  PHY Die 0 Sensor 0 Temp                 54 degrees C / 129 degrees F
  PHY Die 1 Sensor 0 Temp                 49 degrees C / 120 degrees F
  et-2/0/0                                35 degrees C / 95 degrees F
  et-2/0/1                                32 degrees C / 89 degrees F
  et-2/0/2                                39 degrees C / 102 degrees F
  et-2/0/4                                34 degrees C / 93 degrees F
  Exhaust Top Temp                        39 degrees C / 102 degrees F
  Exhaust Bottom Temp                     42 degrees C / 107 degrees F
  Intake Bottom Temp                      35 degrees C / 95 degrees F
  Intake Top Temp                         42 degrees C / 107 degrees F
  Hotswap Controller 0 Temp               40 degrees C / 104 degrees F
  PHY0_1_VDD0V8 Temp                      51 degrees C / 123 degrees F
  VDD3V3 Temp                             53 degrees C / 127 degrees F
  VDD1V2 Temp                             59 degrees C / 138 degrees F
  Power 1
    BIAS3V3                     3332 mV
    VDD1V8                      1812 mV
    VDD2V5                      2500 mV
  Power 2
```

```
    Hotswap 0, 12V0            12031 mV
  Power 3
    PHY0_1_VDD0V8                800 mV
  Power 4
    VDD3V3                      3335 mV
  Power 5
    VDD1V2                      1214 mV
  Bus Revision               27
FPC 3 status:
  State                      Online
  PHY Die 0 Sensor 0 Temp                52 degrees C / 125 degrees F
  Exhaust Top Temp                       37 degrees C / 98 degrees F
  Exhaust Bottom Temp                    37 degrees C / 98 degrees F
  Intake Bottom Temp                     34 degrees C / 93 degrees F
  Intake Top Temp                        39 degrees C / 102 degrees F
  VDD0V8 Temp                            59 degrees C / 138 degrees F
  VDD3V3 Temp                            47 degrees C / 116 degrees F
  VDD1V2 Temp                            53 degrees C / 127 degrees F
  Hotswap 0, 12V0 Temp                   38 degrees C / 100 degrees F
  Power 1
    BIAS3V3                     3348 mV
    VDD1V8                      1830 mV
    VDD2V5                      2541 mV
  Power 2
    VDD0V8                       802 mV
  Power 3
    VDD3V3                      3315 mV
  Power 4
    VDD1V2                      1206 mV
  Bus Revision               15
```

Release Information

Command introduced before Junos OS Release 7.4.

`satellite` option introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

- request chassis fpc
- show chassis fpc

show chassis fpc-feb-connectivity
Resynchronizing FPC Sequence Numbers with Active FPCs when an FPC Comes Online
MX960 Flexible PIC Concentrator Description

show chassis environment fpm

IN THIS SECTION

- Syntax | 17705
- Syntax (TX Matrix Routers) | 17705
- Syntax (TX Matrix Plus Routers) | 17706
- Description | 17706
- Options | 17706
- Required Privilege Level | 17706
- Output Fields | 17707
- Sample Output | 17708
- Release Information | 17714

Syntax

```
show chassis environment fpm
```

Syntax (TX Matrix Routers)

```
show chassis environment fpm
<lcc number | scc>
```

Syntax (TX Matrix Plus Routers)

```
show chassis environment fpm
<lcc number | sfc number>
```

Description

(M40e, M120, M160, M320, MX Series, and T Series routers and the PTX Series Packet Transport Routers only) Display environmental information about the front panel module in the router.

Options

none (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, display environmental information about the front panel modules (craft interfaces) on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about the front panel modules (craft interfaces) on the TX Matrix Plus router and its attached routers.

lcc ***number*** —(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

scc (TX Matrix router only) (Optional) Display environmental information about the front panel module (craft interface) on the TX Matrix router (or switch-card chassis).

sfc ***number*** (TX Matrix Plus router only) (Optional) Display environmental information about the front panel module (craft interface) on the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level

view

Output Fields

Table 228 on page 17707 lists the output fields for the `show chassis environment fpm` command. Output fields are listed in the approximate order in which they appear.

Table 228: show chassis environment fpm Output Fields

Field Name	Field Description
State	FPM status: • **Online**—FPM is online and running. • **Offline**—FPM is powered down.
FPM CMB Voltage	(M40e and M160 routers only) Information about the voltage supplied to the FPM chassis management bus (CMB) device. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
FPM GBUS Voltage	(M320 and T Series routers only) Information about the voltage supplied to the FPM generic bus (GBUS) device. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
FPM I2CS Voltage	(PTX Series only) Information about the voltage supplied to the FPM generic bus (I2CS) device. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
FPM Display Voltage	Information about the voltage supplied to the FPM display. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
FPM CMB Temperature	(M40e and M160 routers only) Temperature of the air flowing past the FPM CMB device
FPM GBUS Temperature	(M320 and T Series routers only) Temperature of the air flowing past the FPM GBUS device.
FPM I2CS Temperature	(PTX Series only) Temperature of the air flowing past the FPM I2CS device.

Table 228: show chassis environment fpm Output Fields *(Continued)*

Field Name	Field Description
FPM Display Temperature	Temperature of the air flowing past the FPM display.
CMB Revision	(M40e and M160 routers only) Revision level of the CMB device.
GBUS Revision	(M320 and T Series routers only) Revision level of the GBUS device.
I2CS Revision	(MX2010 routers, MX2020 routers, MX2008 routers, and PTX Series only) Revision level of the I2CS device.

Sample Output

show chassis environment fpm (M40e and M160 Routers)

```
user@host> show chassis environment fpm
FPM status:
  State                       Online
  FPM CMB Voltage:
    5.0 V bias           5030 mV
    8.0 V bias           8083 mV
  FPM Display Voltage:
    5.0 V bias           4998 mV
  FPM CMB temperature      34 degrees C / 93 degrees F
  FPM Display temperature  35 degrees C / 95 degrees F
  CMB Revision             12
```

show chassis environment fpm (M320 Router)

```
user@host> show chassis environment fpm
FPM status:
  State                           Online
  FPM GBUS Voltage:
    5.0 V                     5006 mV
```

```
    1.8 V bias               1799 mV
    3.3 V bias               3294 mV
    5.0 V bias               4998 mV
    8.0 V bias               7682 mV
  FPM GBUS temperature         30 degrees C / 86 degrees F
  GBUS Revision                51
```

show chassis environment fpm (MX2010 Router)

```
user@host > show chassis environment fpm
FPM status:
  State                      Online
  I2CS Revision              4
```

show chassis environment fpm (MX2020 Router)

```
user@host > show chassis environment fpm
FPM status:
  State                      Online
  I2CS Revision              3
```

show chassis environment fpm (MX2008 Router)

```
user@host > show chassis environment fpm
FPM status:
  State                      Online
  I2CS Revision              5
```

show chassis environment fpm (MX240 Router)

```
user@host> show chassis environment fpm
FPM status:
  State                      Online
  I2CS Revision              41
```

show chassis environment fpm (MX480 Router)

```
user@host> show chassis environment fpm
FPM status:
  State                      Online
  I2CS Revision              41
```

show chassis environment fpm (T Series Routers)

```
user@host> show chassis environment fpm
FPM status:
  State                           Online
  FPM GBUS Voltage:
    1.8 V bias               1787 mV
    3.3 V bias               3286 mV
    5.0 V bias               4991 mV
    8.0 V bias               7162 mV
  FPM Display Voltage:
    5.0 V                    4996 mV
  FPM GBUS temperature         29 degrees C / 84 degrees F
  FPM Display temperature      26 degrees C / 78 degrees F
  GBUS Revision                37
```

show chassis environment fpm lcc (TX Matrix Router)

```
user@host> show chassis environment fpm lcc
0
lcc0-re0:
--------------------------------------------------------------------------
FPM status:
  State                           Online
  FPM GBUS Voltage:
    1.8 V bias               1797 mV
    3.3 V bias               3294 mV
    5.0 V bias               5015 mV
    8.0 V bias               7470 mV
  FPM Display Voltage:
    5.0 V                    5018 mV
  FPM GBUS temperature         25 degrees C / 77 degrees F
```

```
  FPM Display temperature      29 degrees C / 84 degrees F
  GBUS Revision                37
```

show chassis environment fpm scc (TX Matrix Router)

```
user@host> show chassis environment fpm scc
scc-re0:
--------------------------------------------------------------------------
FPM status:
  State                           Online
  FPM GBUS Voltage:
    1.8 V bias               1789 mV
    3.3 V bias               3296 mV
    5.0 V bias               5003 mV
    8.0 V bias               7592 mV
  FPM Display Voltage:
    5.0 V                    5010 mV
  FPM GBUS temperature         22 degrees C / 71 degrees F
  FPM Display temperature      27 degrees C / 80 degrees F
  GBUS Revision                37
```

show chassis environment fpm sfc (TX Matrix Plus Router)

```
user@host> show chassis environment fpm sfc

sfc0-re0:
--------------------------------------------------------------------------
FPM status:
  State                      Online
  FPM I2CS Voltage:
    3.3 V                     3300 mV
    5.0 V                     5001 mV
    9.0 V FPD                 8672 mV
  FPM I2CS temperature       33 degrees C / 91 degrees F
  I2CS Revision              69

lcc0-re0:
--------------------------------------------------------------------------
FPM status:
  State                      Online
```

```
  FPM GBUS Voltage:
    1.8 V bias                1802 mV
    3.3 V bias                3301 mV
    5.0 V bias                4984 mV
    8.0 V bias                7377 mV
  FPM Display Voltage:
    5.0 V                     5015 mV
  FPM GBUS temperature       30 degrees C / 86 degrees F
  FPM Display temperature    32 degrees C / 89 degrees F
  GBUS Revision              37

lcc1-re0:
--------------------------------------------------------------------------
FPM status:
  State                      Online
  FPM GBUS Voltage:
    1.8 V bias                1789 mV
    3.3 V bias                3311 mV
    5.0 V bias                5013 mV
    8.0 V bias                7467 mV
  FPM Display Voltage:
    5.0 V                     5015 mV
  FPM GBUS temperature       29 degrees C / 84 degrees F
  FPM Display temperature    31 degrees C / 87 degrees F
  GBUS Revision              37
```

show chassis environment fpm (T4000 Core Router)

```
user@host> show chassis environment fpm
CB 0 status:
  State                      Online Master
  Temperature                34 degrees C / 93 degrees F
  Power 1
    1.8 V                     1804 mV
    2.5 V                     2499 mV
    3.3 V                     3317 mV
    3.3 V bias                3291 mV
    4.6 V                     4663 mV
    5.0 V                     4905 mV
    8.0 V bias                7658 mV
    12.0 V                   11877 mV
```

```
  Power 2
    1.0 V                      996 mV
    1.2 V                     1207 mV
    3.3 V RE                  3354 mV
  Bus Revision               51
  FPGA Revision              5
CB 1 status:
  State                      Online Standby
  Temperature                36 degrees C / 96 degrees F
  Power 1
    1.8 V                     1791 mV
    2.5 V                     2494 mV
    3.3 V                     3321 mV
    3.3 V bias                3301 mV
    4.6 V                     4666 mV
    5.0 V                     4945 mV
    8.0 V bias                7645 mV
    12.0 V                   11897 mV
  Power 2
    1.0 V                      991 mV
    1.2 V                     1201 mV
    3.3 V RE                  3289 mV
  Bus Revision               51
  FPGA Revision              5

user@host> show chassis environment fpm
FPM status:
  State                      Online
  FPM GBUS Voltage:
    1.8 V bias                1802 mV
    3.3 V bias                3294 mV
    5.0 V bias                5003 mV
    8.0 V bias                7306 mV
  FPM Display Voltage:
    5.0 V                     5010 mV
  FPM GBUS temperature       26 degrees C / 78 degrees F
  FPM Display temperature    29 degrees C / 84 degrees F
  GBUS Revision              37
```

show chassis environment fpm (PTX5000 Packet Transport Router)

```
user@host> show chassis environment fpm

FPM status:
  State                       Online
  FPM I2CS Voltage:
    3.3 V                      3300 mV
    5.0 V                      4975 mV
  FPM I2CS temperature        37 degrees C / 98 degrees F
  I2CS Revision               109
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

request chassis fpm resync | 14489

show chassis environment mcs

IN THIS SECTION

- Syntax | 17715
- Description | 17715
- Options | 17715
- Required Privilege Level | 17715
- Output Fields | 17715
- Sample Output | 17716
- Release Information | 17717

Syntax

```
show chassis environment mcs
<slot>
```

Description

(M40e and M160 routers only) Display environmental information about the Miscellaneous Control Subsystems (MCSs).

Options

none Display environmental information about both MCSs.

slot (Optional) Display environmental information about an individual MCS. Replace ***slot*** with **0** or **1**

Required Privilege Level

view

Output Fields

Table 229 on page 17716 lists the output fields for the `show chassis environment mcs` command. Output fields are listed in the approximate order in which they appear.

Table 229: show chassis environment mcs Output Fields

Field Name	Field Description
State	Status of the MCS: • **Present**—MCS is detected by the chassis daemon but is either not supported by the current version of Junos or MCS is coming up but not yet online. • **Online**—MCS is online and running. • **Offline**—MCS is powered down. • **Empty**—No MCS is present. • **Master**—MCS is online, operating as primary. • **Standby**—MCS is online, operating as standby.
Temperature	Temperature of the air flowing past the MCS.
Power	Information about the voltage supplied to the MCS. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
BUS Revision	Revision level of the generic bus device.
FPGA Revision	Revision level of the field-programmable gate array (FPGA) revision.

Sample Output

show chassis environment mcs (M40e Router)

```
user@host> show chassis environment mcs
MCS 0 status:
  State                           Online Master
  Temperature                  45 degrees C / 113 degrees F
  Power:
    3.3 V                    3283 mV
    5.0 V                    5013 mV
    12.0 V                  11721 mV
```

```
    5.0 V bias                5025 mV
    8.0 V bias                8229 mV
  BUS Revision                  12
  FPGA Revision                 13
MCS 1 status:
  State                            Online Standby
  Temperature                   42 degrees C / 107 degrees F
  Power:
    3.3 V                    3296 mV
    5.0 V                    4971 mV
    12.0 V                  11814 mV
    5.0 V bias               4976 mV
    8.0 V bias               8241 mV
  BUS Revision                  12
  FPGA Revision                 13
```

show chassis environment mcs (M160 Router)

```
user@host> show chassis environment mcs
MCS 0 status:
  State                            Online Master
  Temperature                   50 degrees C / 122 degrees F
  Power:
    3.3 V                    3306 mV
    5.0 V                    4993 mV
    12.0 V                  11799 mV
    5.0 V bias               4993 mV
    8.0 V bias               8288 mV
  BUS Revision                  12
  FPGA Revision                 13
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

request chassis mcs | 14494

show chassis environment monitored

IN THIS SECTION

- Syntax | 17718
- Syntax (MX2020, MX2010, and MX2008 Routers) | 17718
- Description | 17718
- Options | 17719
- Required Privilege Level | 17719
- Output Fields | 17719
- Sample Output | 17720
- Release Information | 17748

Syntax

```
show chassis environment monitored
```

Syntax (MX2020, MX2010, and MX2008 Routers)

```
show chassis environment monitored
<all-members>
<local>
<member member-id>
```

Description

(PTX Series Packet Transport Routers, MX2010, MX2020, MX2008, and MX10008 routers) Display status information for monitored temperatures.

On the PTX Series Packet Transport Routers, and on MX2010, MX2020, MX2008, and MX10008 routers, you can configure which temperatures are monitored for computing temperature alarms. Use this command to display only the temperatures that are monitored. Temperatures that are not included in the temperature alarm computations are not displayed.

Options

none Display status information for monitored temperatures.

all-members (MX2010, MX2020, and MX2008 routers only) (Optional) Display chassis information for monitored temperatures in all members of the Virtual Chassis configuration.

local (MX2010, MX2020, and MX2008 routers only) (Optional) Display chassis information for monitored temperatures in the local member of the Virtual Chassis.

member *member-id* (MX2010, MX2020, and MX2008 routers only) (Optional) Display chassis information for monitored temperatures in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 230 on page 17719 lists the output fields for the `show chassis environment monitored` command. Output fields are listed in the approximate order in which they appear.

Table 230: show chassis environment monitored Output Fields

Field Name	Field Description
`Item`	Chassis component: • (PTX Series Packet Transport Routers, and MX2010, MX2020, and Mx2008 routers)—Information about the chassis, Routing Engines, Control Boards (CBs), Switch Interface Boards (SIBs), PICs, and Flexible PIC Concentrators (FPCs).
`Status`	Status of the specified item: `OK`, `Alarm`, or `Present`.
`Measurement`	Temperature of the air flowing past the specified chassis component. Temperature is displayed in degrees Celsius (C) and degrees Fahrenheit (F).

Sample Output

show chassis environment monitored (PTX3000 Packet Transport Router)

```
user@host> show chassis environment monitored
Class Item                          Status     Measurement
      Routing Engine 0 CPU          OK         54 degrees C / 129 degrees F
      Routing Engine 1 CPU          Present
      CB 0 Exhaust A                OK         25 degrees C / 77 degrees F
      CB 1 Exhaust A                OK         22 degrees C / 71 degrees F
      SIB 0 Exhaust                 OK         34 degrees C / 93 degrees F
      SIB 0 TF                      OK         42 degrees C / 107 degrees F
      SIB 1 Exhaust                 OK         31 degrees C / 87 degrees F
      SIB 1 TF                      OK         41 degrees C / 105 degrees F
      SIB 2 Exhaust                 OK         32 degrees C / 89 degrees F
      SIB 2 TF                      OK         40 degrees C / 104 degrees F
      SIB 3 Exhaust                 OK         32 degrees C / 89 degrees F
      SIB 3 TF                      OK         40 degrees C / 104 degrees F
      SIB 4 Exhaust                 OK         31 degrees C / 87 degrees F
      SIB 4 TF                      OK         40 degrees C / 104 degrees F
      SIB 5 Exhaust                 OK         31 degrees C / 87 degrees F
      SIB 5 TF                      OK         39 degrees C / 102 degrees F
      SIB 6 Exhaust                 OK         31 degrees C / 87 degrees F
      SIB 6 TF                      OK         39 degrees C / 102 degrees F
      SIB 7 Exhaust                 OK         35 degrees C / 95 degrees F
      SIB 7 TF                      OK         40 degrees C / 104 degrees F
      SIB 8 Exhaust                 OK         32 degrees C / 89 degrees F
      SIB 8 TF                      OK         40 degrees C / 104 degrees F
      FPC 2 PMB CPU                 OK         67 degrees C / 152 degrees F
      FPC 2 Exhaust                 OK         40 degrees C / 104 degrees F
      FPC 2 Intake                  OK         33 degrees C / 91 degrees F
      FPC 2 TL0                     OK         69 degrees C / 156 degrees F
      FPC 2 TQ0                     OK         60 degrees C / 140 degrees F
      FPC 2 TL1                     OK         56 degrees C / 132 degrees F
      FPC 2 TQ1                     OK         45 degrees C / 113 degrees F
      PIC Ambient                   OK         40 degrees C / 104 degrees F
      FPC 6 PMB CPU                 OK         80 degrees C / 176 degrees F
      FPC 6 Exhaust                 OK         53 degrees C / 127 degrees F
      FPC 6 Intake                  OK         36 degrees C / 96 degrees F
      FPC 6 TL0                     OK         69 degrees C / 156 degrees F
      FPC 6 TQ0                     OK         65 degrees C / 149 degrees F
      FPC 6 TL1                     OK         52 degrees C / 125 degrees F
```

```
      FPC 6 TQ1                       OK         47 degrees C / 116 degrees F
      PIC Ambient                     OK         46 degrees C / 114 degrees F
      FPC 12 PMB CPU                  OK         42 degrees C / 107 degrees F
      FPC 12 Intake                   OK         33 degrees C / 91 degrees F
      FPC 12 Exhaust                  OK         41 degrees C / 105 degrees F
      FPC 12 TL0                      OK         48 degrees C / 118 degrees F
      FPC 12 TQ0                      OK         45 degrees C / 113 degrees F
      FPC 12 TL1                      OK         58 degrees C / 136 degrees F
      FPC 12 TQ1                      OK         50 degrees C / 122 degrees F
      PIC Ambient                     OK         56 degrees C / 132 degrees F
      PIC 100G_OTN_LH-12/0/0          OK         74 degrees C / 165 degrees F
      PIC 100G_OTN_LH-12/0/1          OK         93 degrees C / 199 degrees F
```

show chassis environment monitored (PTX5000 Packet Transport Router)

```
user@host> show chassis environment monitored
Class Item                            Status     Measurement
      Routing Engine 0 CPU            OK         71 degrees C / 159 degrees F
      Routing Engine 1 CPU            OK         62 degrees C / 143 degrees F
      CB 0 Exhaust A                  OK         45 degrees C / 113 degrees F
      CB 0 Exhaust B                  OK         41 degrees C / 105 degrees F
      CB 1 Exhaust A                  OK         39 degrees C / 102 degrees F
      CB 1 Exhaust B                  OK         36 degrees C / 96 degrees F
```

show chassis environment monitored (MX2010 Router)

```
user@host > show chassis environment monitored
Class Item                            Status     Measurement
Temp  CB 0 IntakeA-Zone0              OK         37 degrees C / 98 degrees F
      CB 0 IntakeB-Zone1              OK         31 degrees C / 87 degrees F
      CB 0 IntakeC-Zone0              OK         39 degrees C / 102 degrees F
      CB 0 ExhaustA-Zone0             OK         36 degrees C / 96 degrees F
      CB 0 ExhaustB-Zone1             OK         32 degrees C / 89 degrees F
      CB 0 TCBC-Zone0                 OK         34 degrees C / 93 degrees F
      CB 1 IntakeA-Zone0              OK         36 degrees C / 96 degrees F
      CB 1 IntakeB-Zone1              OK         28 degrees C / 82 degrees F
      CB 1 IntakeC-Zone0              OK         38 degrees C / 100 degrees F
      CB 1 ExhaustA-Zone0             OK         36 degrees C / 96 degrees F
```

```
    CB 1 ExhaustB-Zone1            OK         30 degrees C / 86 degrees F
    CB 1 TCBC-Zone0                OK         33 degrees C / 91 degrees F
    SPMB 0 Intake                  OK         30 degrees C / 86 degrees F
    SPMB 1 Intake                  OK         28 degrees C / 82 degrees F
    Routing Engine 0 CPU           OK         32 degrees C / 89 degrees F
    Routing Engine 1 CPU           Present
    SFB 0 Intake-Zone0             OK         46 degrees C / 114 degrees F
    SFB 0 Exhaust-Zone1            OK         38 degrees C / 100 degrees F
    SFB 0 IntakeA-Zone0            OK         35 degrees C / 95 degrees F
    SFB 0 IntakeB-Zone1            OK         31 degrees C / 87 degrees F
    SFB 0 Exhaust-Zone0            OK         39 degrees C / 102 degrees F
    SFB 0 SFB-XF2-Zone1            OK         44 degrees C / 111 degrees F
    SFB 0 SFB-XF1-Zone0            OK         47 degrees C / 116 degrees F
    SFB 0 SFB-XF0-Zone0            OK         56 degrees C / 132 degrees F
    SFB 1 Intake-Zone0             OK         34 degrees C / 93 degrees F
    SFB 1 Exhaust-Zone1            OK         34 degrees C / 93 degrees F
    SFB 1 IntakeA-Zone0            OK         29 degrees C / 84 degrees F
    SFB 1 IntakeB-Zone1            OK         29 degrees C / 84 degrees F
    SFB 1 Exhaust-Zone0            OK         32 degrees C / 89 degrees F
    SFB 1 SFB-XF2-Zone1            OK         42 degrees C / 107 degrees F
    SFB 1 SFB-XF1-Zone0            OK         40 degrees C / 104 degrees F
    SFB 1 SFB-XF0-Zone0            OK         42 degrees C / 107 degrees F
    SFB 2 Intake-Zone0             OK         33 degrees C / 91 degrees F
    SFB 2 Exhaust-Zone1            OK         33 degrees C / 91 degrees F
    SFB 2 IntakeA-Zone0            OK         28 degrees C / 82 degrees F
    SFB 2 IntakeB-Zone1            OK         28 degrees C / 82 degrees F
    SFB 2 Exhaust-Zone0            OK         31 degrees C / 87 degrees F
    SFB 2 SFB-XF2-Zone1            OK         41 degrees C / 105 degrees F
    SFB 2 SFB-XF1-Zone0            OK         39 degrees C / 102 degrees F
    SFB 2 SFB-XF0-Zone0            OK         42 degrees C / 107 degrees F
    SFB 3 Intake-Zone0             OK         33 degrees C / 91 degrees F
    SFB 3 Exhaust-Zone1            OK         33 degrees C / 91 degrees F
    SFB 3 IntakeA-Zone0            OK         29 degrees C / 84 degrees F
    SFB 3 IntakeB-Zone1            OK         28 degrees C / 82 degrees F
    SFB 3 Exhaust-Zone0            OK         31 degrees C / 87 degrees F
    SFB 3 SFB-XF2-Zone1            OK         43 degrees C / 109 degrees F
    SFB 3 SFB-XF1-Zone0            OK         40 degrees C / 104 degrees F
    SFB 3 SFB-XF0-Zone0            OK         42 degrees C / 107 degrees F
    SFB 4 Intake-Zone0             OK         34 degrees C / 93 degrees F
    SFB 4 Exhaust-Zone1            OK         34 degrees C / 93 degrees F
    SFB 4 IntakeA-Zone0            OK         29 degrees C / 84 degrees F
    SFB 4 IntakeB-Zone1            OK         28 degrees C / 82 degrees F
    SFB 4 Exhaust-Zone0            OK         32 degrees C / 89 degrees F
```

```
    SFB 4 SFB-XF2-Zone1          OK        43 degrees C / 109 degrees F
    SFB 4 SFB-XF1-Zone0          OK        42 degrees C / 107 degrees F
    SFB 4 SFB-XF0-Zone0          OK        43 degrees C / 109 degrees F
    SFB 5 Intake-Zone0           OK        34 degrees C / 93 degrees F
    SFB 5 Exhaust-Zone1          OK        34 degrees C / 93 degrees F
    SFB 5 IntakeA-Zone0          OK        30 degrees C / 86 degrees F
    SFB 5 IntakeB-Zone1          OK        28 degrees C / 82 degrees F
    SFB 5 Exhaust-Zone0          OK        32 degrees C / 89 degrees F
    SFB 5 SFB-XF2-Zone1          OK        41 degrees C / 105 degrees F
    SFB 5 SFB-XF1-Zone0          OK        41 degrees C / 105 degrees F
    SFB 5 SFB-XF0-Zone0          OK        44 degrees C / 111 degrees F
    SFB 6 Intake-Zone0           OK        35 degrees C / 95 degrees F
    SFB 6 Exhaust-Zone1          OK        34 degrees C / 93 degrees F
    SFB 6 IntakeA-Zone0          OK        30 degrees C / 86 degrees F
    SFB 6 IntakeB-Zone1          OK        29 degrees C / 84 degrees F
    SFB 6 Exhaust-Zone0          OK        33 degrees C / 91 degrees F
    SFB 6 SFB-XF2-Zone1          OK        44 degrees C / 111 degrees F
    SFB 6 SFB-XF1-Zone0          OK        43 degrees C / 109 degrees F
    SFB 6 SFB-XF0-Zone0          OK        46 degrees C / 114 degrees F
    SFB 7 Intake-Zone0           OK        39 degrees C / 102 degrees F
    SFB 7 Exhaust-Zone1          OK        34 degrees C / 93 degrees F
    SFB 7 IntakeA-Zone0          OK        34 degrees C / 93 degrees F
    SFB 7 IntakeB-Zone1          OK        29 degrees C / 84 degrees F
    SFB 7 Exhaust-Zone0          OK        37 degrees C / 98 degrees F
    SFB 7 SFB-XF2-Zone1          OK        43 degrees C / 109 degrees F
    SFB 7 SFB-XF1-Zone0          OK        47 degrees C / 116 degrees F
    SFB 7 SFB-XF0-Zone0          OK        52 degrees C / 125 degrees F
    FPC 0 Intake                 OK        36 degrees C / 96 degrees F
    FPC 0 Exhaust A              OK        42 degrees C / 107 degrees F
    FPC 0 Exhaust B              OK        51 degrees C / 123 degrees F
    FPC 0 LU 0 TSen              OK        49 degrees C / 120 degrees F
    FPC 0 LU 0 Chip              OK        50 degrees C / 122 degrees F
    FPC 0 LU 1 TSen              OK        49 degrees C / 120 degrees F
    FPC 0 LU 1 Chip              OK        54 degrees C / 129 degrees F
    FPC 0 LU 2 TSen              OK        49 degrees C / 120 degrees F
    FPC 0 LU 2 Chip              OK        45 degrees C / 113 degrees F
    FPC 0 LU 3 TSen              OK        49 degrees C / 120 degrees F
    FPC 0 LU 3 Chip              OK        46 degrees C / 114 degrees F
    FPC 0 MQ 0 TSen              OK        40 degrees C / 104 degrees F
    FPC 0 MQ 0 Chip              OK        41 degrees C / 105 degrees F
    FPC 0 MQ 1 TSen              OK        40 degrees C / 104 degrees F
    FPC 0 MQ 1 Chip              OK        44 degrees C / 111 degrees F
    FPC 0 MQ 2 TSen              OK        40 degrees C / 104 degrees F
```

```
    FPC 0 MQ 2 Chip                OK         38 degrees C / 100 degrees F
    FPC 0 MQ 3 TSen                OK         40 degrees C / 104 degrees F
    FPC 0 MQ 3 Chip                OK         41 degrees C / 105 degrees F
    FPC 1 Intake                   OK         34 degrees C / 93 degrees F
    FPC 1 Exhaust A                OK         46 degrees C / 114 degrees F
    FPC 1 Exhaust B                OK         54 degrees C / 129 degrees F
    FPC 1 LU 0 TSen                OK         45 degrees C / 113 degrees F
    FPC 1 LU 0 Chip                OK         55 degrees C / 131 degrees F
    FPC 1 LU 1 TSen                OK         45 degrees C / 113 degrees F
    FPC 1 LU 1 Chip                OK         44 degrees C / 111 degrees F
    FPC 1 LU 2 TSen                OK         45 degrees C / 113 degrees F
    FPC 1 LU 2 Chip                OK         50 degrees C / 122 degrees F
    FPC 1 LU 3 TSen                OK         45 degrees C / 113 degrees F
    FPC 1 LU 3 Chip                OK         58 degrees C / 136 degrees F
    FPC 1 XM 0 TSen                OK         45 degrees C / 113 degrees F
    FPC 1 XM 0 Chip                OK         52 degrees C / 125 degrees F
    FPC 1 XF 0 TSen                OK         45 degrees C / 113 degrees F
    FPC 1 XF 0 Chip                OK         63 degrees C / 145 degrees F
    FPC 1 PLX Switch TSen          OK         45 degrees C / 113 degrees F
    FPC 1 PLX Switch Chip          OK         47 degrees C / 116 degrees F
    FPC 8 Intake                   OK         32 degrees C / 89 degrees F
    FPC 8 Exhaust A                OK         44 degrees C / 111 degrees F
    FPC 8 Exhaust B                OK         37 degrees C / 98 degrees F
    FPC 8 LU 0 TCAM TSen           OK         41 degrees C / 105 degrees F
    FPC 8 LU 0 TCAM Chip           OK         49 degrees C / 120 degrees F
    FPC 8 LU 0 TSen                OK         41 degrees C / 105 degrees F
    FPC 8 LU 0 Chip                OK         52 degrees C / 125 degrees F
    FPC 8 MQ 0 TSen                OK         41 degrees C / 105 degrees F
    FPC 8 MQ 0 Chip                OK         47 degrees C / 116 degrees F
    FPC 8 LU 1 TCAM TSen           OK         39 degrees C / 102 degrees F
    FPC 8 LU 1 TCAM Chip           OK         42 degrees C / 107 degrees F
    FPC 8 LU 1 TSen                OK         39 degrees C / 102 degrees F
    FPC 8 LU 1 Chip                OK         46 degrees C / 114 degrees F
    FPC 8 MQ 1 TSen                OK         39 degrees C / 102 degrees F
    FPC 8 MQ 1 Chip                OK         45 degrees C / 113 degrees F
    FPC 9 Intake                   OK         34 degrees C / 93 degrees F
    FPC 9 Exhaust A                OK         41 degrees C / 105 degrees F
    FPC 9 Exhaust B                OK         54 degrees C / 129 degrees F
    FPC 9 LU 0 TSen                OK         51 degrees C / 123 degrees F
    FPC 9 LU 0 Chip                OK         52 degrees C / 125 degrees F
    FPC 9 LU 1 TSen                OK         51 degrees C / 123 degrees F
    FPC 9 LU 1 Chip                OK         55 degrees C / 131 degrees F
    FPC 9 LU 2 TSen                OK         51 degrees C / 123 degrees F
```

```
      FPC 9 LU 2 Chip                OK         47 degrees C / 116 degrees F
      FPC 9 LU 3 TSen                OK         51 degrees C / 123 degrees F
      FPC 9 LU 3 Chip                OK         47 degrees C / 116 degrees F
      FPC 9 MQ 0 TSen                OK         40 degrees C / 104 degrees F
      FPC 9 MQ 0 Chip                OK         42 degrees C / 107 degrees F
      FPC 9 MQ 1 TSen                OK         40 degrees C / 104 degrees F
      FPC 9 MQ 1 Chip                OK         44 degrees C / 111 degrees F
      FPC 9 MQ 2 TSen                OK         40 degrees C / 104 degrees F
      FPC 9 MQ 2 Chip                OK         38 degrees C / 100 degrees F
      FPC 9 MQ 3 TSen                OK         40 degrees C / 104 degrees F
      FPC 9 MQ 3 Chip                OK         40 degrees C / 104 degrees F
      ADC 0 Intake                   OK         35 degrees C / 95 degrees F
      ADC 0 Exhaust                  OK         44 degrees C / 111 degrees F
      ADC 0 ADC-XF1                  OK         48 degrees C / 118 degrees F
      ADC 0 ADC-XF0                  OK         59 degrees C / 138 degrees F
      ADC 1 Intake                   OK         34 degrees C / 93 degrees F
      ADC 1 Exhaust                  OK         45 degrees C / 113 degrees F
      ADC 1 ADC-XF1                  OK         53 degrees C / 127 degrees F
      ADC 1 ADC-XF0                  OK         56 degrees C / 132 degrees F
      ADC 8 Intake                   OK         35 degrees C / 95 degrees F
      ADC 8 Exhaust                  OK         41 degrees C / 105 degrees F
      ADC 8 ADC-XF1                  OK         52 degrees C / 125 degrees F
      ADC 8 ADC-XF0                  OK         55 degrees C / 131 degrees F
      ADC 9 Intake                   OK         33 degrees C / 91 degrees F
      ADC 9 Exhaust                  OK         42 degrees C / 107 degrees F
      ADC 9 ADC-XF1                  OK         55 degrees C / 131 degrees F
      ADC 9 ADC-XF0                  OK         56 degrees C / 132 degrees F
```

show chassis environment monitored (MX2020 Router)

```
user@host > show chassis environment monitored
Class Item                           Status     Measurement
Temp  CB 0 IntakeA-Zone0             OK         44 degrees C / 111 degrees F
      CB 0 IntakeB-Zone1             OK         34 degrees C / 93 degrees F
      CB 0 IntakeC-Zone0             OK         46 degrees C / 114 degrees F
      CB 0 ExhaustA-Zone0            OK         44 degrees C / 111 degrees F
      CB 0 ExhaustB-Zone1            OK         36 degrees C / 96 degrees F
      CB 0 TCBC-Zone0                OK         39 degrees C / 102 degrees F
      CB 1 IntakeA-Zone0             OK         46 degrees C / 114 degrees F
      CB 1 IntakeB-Zone1             OK         43 degrees C / 109 degrees F
      CB 1 IntakeC-Zone0             OK         47 degrees C / 116 degrees F
```

```
    CB 1 ExhaustA-Zone0          OK        45 degrees C / 113 degrees F
    CB 1 ExhaustB-Zone1          OK        42 degrees C / 107 degrees F
    CB 1 TCBC-Zone0              OK        46 degrees C / 114 degrees F
    SPMB 0 Intake                OK        33 degrees C / 91 degrees F
    SPMB 1 Intake                OK        43 degrees C / 109 degrees F
    Routing Engine 0 CPU         OK        34 degrees C / 93 degrees F
    Routing Engine 1 CPU         OK        42 degrees C / 107 degrees F
    SFB 0 Intake-Zone0           OK        52 degrees C / 125 degrees F
    SFB 0 Exhaust-Zone1          OK        45 degrees C / 113 degrees F
    SFB 0 IntakeA-Zone0          OK        47 degrees C / 116 degrees F
    SFB 0 IntakeB-Zone1          OK        38 degrees C / 100 degrees F
    SFB 0 Exhaust-Zone0          OK        49 degrees C / 120 degrees F
    SFB 0 SFB-XF2-Zone1          OK        59 degrees C / 138 degrees F
    SFB 0 SFB-XF1-Zone0          OK        65 degrees C / 149 degrees F
    SFB 0 SFB-XF0-Zone0          OK        65 degrees C / 149 degrees F
    SFB 1 Intake-Zone0           OK        53 degrees C / 127 degrees F
    SFB 1 Exhaust-Zone1          OK        45 degrees C / 113 degrees F
    SFB 1 IntakeA-Zone0          OK        48 degrees C / 118 degrees F
    SFB 1 IntakeB-Zone1          OK        39 degrees C / 102 degrees F
    SFB 1 Exhaust-Zone0          OK        48 degrees C / 118 degrees F
    SFB 1 SFB-XF2-Zone1          OK        60 degrees C / 140 degrees F
    SFB 1 SFB-XF1-Zone0          OK        64 degrees C / 147 degrees F
    SFB 1 SFB-XF0-Zone0          OK        66 degrees C / 150 degrees F
    SFB 2 Intake-Zone0           OK        54 degrees C / 129 degrees F
    SFB 2 Exhaust-Zone1          OK        46 degrees C / 114 degrees F
    SFB 2 IntakeA-Zone0          OK        48 degrees C / 118 degrees F
    SFB 2 IntakeB-Zone1          OK        39 degrees C / 102 degrees F
    SFB 2 Exhaust-Zone0          OK        50 degrees C / 122 degrees F
    SFB 2 SFB-XF2-Zone1          OK        63 degrees C / 145 degrees F
    SFB 2 SFB-XF1-Zone0          OK        67 degrees C / 152 degrees F
    SFB 2 SFB-XF0-Zone0          OK        67 degrees C / 152 degrees F
    SFB 3 Intake-Zone0           OK        54 degrees C / 129 degrees F
    SFB 3 Exhaust-Zone1          OK        46 degrees C / 114 degrees F
    SFB 3 IntakeA-Zone0          OK        50 degrees C / 122 degrees F
    SFB 3 IntakeB-Zone1          OK        40 degrees C / 104 degrees F
    SFB 3 Exhaust-Zone0          OK        50 degrees C / 122 degrees F
    SFB 3 SFB-XF2-Zone1          OK        64 degrees C / 147 degrees F
    SFB 3 SFB-XF1-Zone0          OK        66 degrees C / 150 degrees F
    SFB 3 SFB-XF0-Zone0          OK        68 degrees C / 154 degrees F
    SFB 4 Intake-Zone0           OK        55 degrees C / 131 degrees F
    SFB 4 Exhaust-Zone1          OK        48 degrees C / 118 degrees F
    SFB 4 IntakeA-Zone0          OK        51 degrees C / 123 degrees F
    SFB 4 IntakeB-Zone1          OK        42 degrees C / 107 degrees F
```

```
    SFB 4 Exhaust-Zone0          OK        51 degrees C / 123 degrees F
    SFB 4 SFB-XF2-Zone1          OK        63 degrees C / 145 degrees F
    SFB 4 SFB-XF1-Zone0          OK        66 degrees C / 150 degrees F
    SFB 4 SFB-XF0-Zone0          OK        68 degrees C / 154 degrees F
    SFB 5 Intake-Zone0           OK        55 degrees C / 131 degrees F
    SFB 5 Exhaust-Zone1          OK        49 degrees C / 120 degrees F
    SFB 5 IntakeA-Zone0          OK        51 degrees C / 123 degrees F
    SFB 5 IntakeB-Zone1          OK        43 degrees C / 109 degrees F
    SFB 5 Exhaust-Zone0          OK        51 degrees C / 123 degrees F
    SFB 5 SFB-XF2-Zone1          OK        65 degrees C / 149 degrees F
    SFB 5 SFB-XF1-Zone0          OK        66 degrees C / 150 degrees F
    SFB 5 SFB-XF0-Zone0          OK        71 degrees C / 159 degrees F
    SFB 6 Intake-Zone0           OK        55 degrees C / 131 degrees F
    SFB 6 Exhaust-Zone1          OK        49 degrees C / 120 degrees F
    SFB 6 IntakeA-Zone0          OK        51 degrees C / 123 degrees F
    SFB 6 IntakeB-Zone1          OK        43 degrees C / 109 degrees F
    SFB 6 Exhaust-Zone0          OK        51 degrees C / 123 degrees F
    SFB 6 SFB-XF2-Zone1          OK        64 degrees C / 147 degrees F
    SFB 6 SFB-XF1-Zone0          OK        66 degrees C / 150 degrees F
    SFB 6 SFB-XF0-Zone0          OK        68 degrees C / 154 degrees F
    SFB 7 Intake-Zone0           OK        55 degrees C / 131 degrees F
    SFB 7 Exhaust-Zone1          OK        49 degrees C / 120 degrees F
    SFB 7 IntakeA-Zone0          OK        51 degrees C / 123 degrees F
    SFB 7 IntakeB-Zone1          OK        43 degrees C / 109 degrees F
    SFB 7 Exhaust-Zone0          OK        52 degrees C / 125 degrees F
    SFB 7 SFB-XF2-Zone1          OK        66 degrees C / 150 degrees F
    SFB 7 SFB-XF1-Zone0          OK        67 degrees C / 152 degrees F
    SFB 7 SFB-XF0-Zone0          OK        70 degrees C / 158 degrees F
    FPC 0 Intake                 OK        41 degrees C / 105 degrees F
    FPC 0 Exhaust A              OK        48 degrees C / 118 degrees F
    FPC 0 Exhaust B              OK        60 degrees C / 140 degrees F
    FPC 0 LU 0 TSen              OK        56 degrees C / 132 degrees F
    FPC 0 LU 0 Chip              OK        59 degrees C / 138 degrees F
    FPC 0 LU 1 TSen              OK        56 degrees C / 132 degrees F
    FPC 0 LU 1 Chip              OK        61 degrees C / 141 degrees F
    FPC 0 LU 2 TSen              OK        56 degrees C / 132 degrees F
    FPC 0 LU 2 Chip              OK        52 degrees C / 125 degrees F
    FPC 0 LU 3 TSen              OK        56 degrees C / 132 degrees F
    FPC 0 LU 3 Chip              OK        52 degrees C / 125 degrees F
    FPC 0 MQ 0 TSen              OK        49 degrees C / 120 degrees F
    FPC 0 MQ 0 Chip              OK        49 degrees C / 120 degrees F
    FPC 0 MQ 1 TSen              OK        49 degrees C / 120 degrees F
    FPC 0 MQ 1 Chip              OK        52 degrees C / 125 degrees F
```

```
    FPC 0 MQ 2 TSen              OK        49 degrees C / 120 degrees F
    FPC 0 MQ 2 Chip              OK        45 degrees C / 113 degrees F
    FPC 0 MQ 3 TSen              OK        49 degrees C / 120 degrees F
    FPC 0 MQ 3 Chip              OK        46 degrees C / 114 degrees F
    FPC 1 Intake                 OK        39 degrees C / 102 degrees F
    FPC 1 Exhaust A              OK        48 degrees C / 118 degrees F
    FPC 1 Exhaust B              OK        55 degrees C / 131 degrees F
    FPC 1 LU 0 TSen              OK        52 degrees C / 125 degrees F
    FPC 1 LU 0 Chip              OK        54 degrees C / 129 degrees F
    FPC 1 LU 1 TSen              OK        52 degrees C / 125 degrees F
    FPC 1 LU 1 Chip              OK        56 degrees C / 132 degrees F
    FPC 1 LU 2 TSen              OK        52 degrees C / 125 degrees F
    FPC 1 LU 2 Chip              OK        49 degrees C / 120 degrees F
    FPC 1 LU 3 TSen              OK        52 degrees C / 125 degrees F
    FPC 1 LU 3 Chip              OK        50 degrees C / 122 degrees F
    FPC 1 MQ 0 TSen              OK        48 degrees C / 118 degrees F
    FPC 1 MQ 0 Chip              OK        48 degrees C / 118 degrees F
    FPC 1 MQ 1 TSen              OK        48 degrees C / 118 degrees F
    FPC 1 MQ 1 Chip              OK        51 degrees C / 123 degrees F
    FPC 1 MQ 2 TSen              OK        48 degrees C / 118 degrees F
    FPC 1 MQ 2 Chip              OK        45 degrees C / 113 degrees F
    FPC 1 MQ 3 TSen              OK        48 degrees C / 118 degrees F
    FPC 1 MQ 3 Chip              OK        45 degrees C / 113 degrees F
    FPC 2 Intake                 OK        39 degrees C / 102 degrees F
    FPC 2 Exhaust A              OK        48 degrees C / 118 degrees F
    FPC 2 Exhaust B              OK        58 degrees C / 136 degrees F
    FPC 2 LU 0 TSen              OK        55 degrees C / 131 degrees F
    FPC 2 LU 0 Chip              OK        57 degrees C / 134 degrees F
    FPC 2 LU 1 TSen              OK        55 degrees C / 131 degrees F
    FPC 2 LU 1 Chip              OK        63 degrees C / 145 degrees F
    FPC 2 LU 2 TSen              OK        55 degrees C / 131 degrees F
    FPC 2 LU 2 Chip              OK        51 degrees C / 123 degrees F
    FPC 2 LU 3 TSen              OK        55 degrees C / 131 degrees F
    FPC 2 LU 3 Chip              OK        52 degrees C / 125 degrees F
    FPC 2 MQ 0 TSen              OK        48 degrees C / 118 degrees F
    FPC 2 MQ 0 Chip              OK        50 degrees C / 122 degrees F
    FPC 2 MQ 1 TSen              OK        48 degrees C / 118 degrees F
    FPC 2 MQ 1 Chip              OK        52 degrees C / 125 degrees F
    FPC 2 MQ 2 TSen              OK        48 degrees C / 118 degrees F
    FPC 2 MQ 2 Chip              OK        47 degrees C / 116 degrees F
    FPC 2 MQ 3 TSen              OK        48 degrees C / 118 degrees F
    FPC 2 MQ 3 Chip              OK        47 degrees C / 116 degrees F
    FPC 3 Intake                 OK        41 degrees C / 105 degrees F
```

```
    FPC 3 Exhaust A              OK        48 degrees C / 118 degrees F
    FPC 3 Exhaust B              OK        58 degrees C / 136 degrees F
    FPC 3 LU 0 TSen              OK        56 degrees C / 132 degrees F
    FPC 3 LU 0 Chip              OK        59 degrees C / 138 degrees F
    FPC 3 LU 1 TSen              OK        56 degrees C / 132 degrees F
    FPC 3 LU 1 Chip              OK        61 degrees C / 141 degrees F
    FPC 3 LU 2 TSen              OK        56 degrees C / 132 degrees F
    FPC 3 LU 2 Chip              OK        51 degrees C / 123 degrees F
    FPC 3 LU 3 TSen              OK        56 degrees C / 132 degrees F
    FPC 3 LU 3 Chip              OK        53 degrees C / 127 degrees F
    FPC 3 MQ 0 TSen              OK        50 degrees C / 122 degrees F
    FPC 3 MQ 0 Chip              OK        51 degrees C / 123 degrees F
    FPC 3 MQ 1 TSen              OK        50 degrees C / 122 degrees F
    FPC 3 MQ 1 Chip              OK        55 degrees C / 131 degrees F
    FPC 3 MQ 2 TSen              OK        50 degrees C / 122 degrees F
    FPC 3 MQ 2 Chip              OK        47 degrees C / 116 degrees F
    FPC 3 MQ 3 TSen              OK        50 degrees C / 122 degrees F
    FPC 3 MQ 3 Chip              OK        50 degrees C / 122 degrees F
    FPC 4 Intake                 OK        41 degrees C / 105 degrees F
    FPC 4 Exhaust A              OK        48 degrees C / 118 degrees F
    FPC 4 Exhaust B              OK        59 degrees C / 138 degrees F
    FPC 4 LU 0 TSen              OK        56 degrees C / 132 degrees F
    FPC 4 LU 0 Chip              OK        60 degrees C / 140 degrees F
    FPC 4 LU 1 TSen              OK        56 degrees C / 132 degrees F
    FPC 4 LU 1 Chip              OK        64 degrees C / 147 degrees F
    FPC 4 LU 2 TSen              OK        56 degrees C / 132 degrees F
    FPC 4 LU 2 Chip              OK        51 degrees C / 123 degrees F
    FPC 4 LU 3 TSen              OK        56 degrees C / 132 degrees F
    FPC 4 LU 3 Chip              OK        54 degrees C / 129 degrees F
    FPC 4 MQ 0 TSen              OK        49 degrees C / 120 degrees F
    FPC 4 MQ 0 Chip              OK        53 degrees C / 127 degrees F
    FPC 4 MQ 1 TSen              OK        49 degrees C / 120 degrees F
    FPC 4 MQ 1 Chip              OK        55 degrees C / 131 degrees F
    FPC 4 MQ 2 TSen              OK        49 degrees C / 120 degrees F
    FPC 4 MQ 2 Chip              OK        48 degrees C / 118 degrees F
    FPC 4 MQ 3 TSen              OK        49 degrees C / 120 degrees F
    FPC 4 MQ 3 Chip              OK        49 degrees C / 120 degrees F
    FPC 5 Intake                 OK        42 degrees C / 107 degrees F
    FPC 5 Exhaust A              OK        49 degrees C / 120 degrees F
    FPC 5 Exhaust B              OK        61 degrees C / 141 degrees F
    FPC 5 LU 0 TSen              OK        58 degrees C / 136 degrees F
    FPC 5 LU 0 Chip              OK        61 degrees C / 141 degrees F
    FPC 5 LU 1 TSen              OK        58 degrees C / 136 degrees F
```

```
        FPC 5 LU 1 Chip                 OK          64 degrees C / 147 degrees F
        FPC 5 LU 2 TSen                 OK          58 degrees C / 136 degrees F
        FPC 5 LU 2 Chip                 OK          56 degrees C / 132 degrees F
        FPC 5 LU 3 TSen                 OK          58 degrees C / 136 degrees F
        FPC 5 LU 3 Chip                 OK          54 degrees C / 129 degrees F
        FPC 5 MQ 0 TSen                 OK          51 degrees C / 123 degrees F
        FPC 5 MQ 0 Chip                 OK          53 degrees C / 127 degrees F
        FPC 5 MQ 1 TSen                 OK          51 degrees C / 123 degrees F
        FPC 5 MQ 1 Chip                 OK          55 degrees C / 131 degrees F
        FPC 5 MQ 2 TSen                 OK          51 degrees C / 123 degrees F
        FPC 5 MQ 2 Chip                 OK          50 degrees C / 122 degrees F
        FPC 5 MQ 3 TSen                 OK          51 degrees C / 123 degrees F
        FPC 5 MQ 3 Chip                 OK          49 degrees C / 120 degrees F
        FPC 6 Intake                    OK          42 degrees C / 107 degrees F
        FPC 6 Exhaust A                 OK          50 degrees C / 122 degrees F
        FPC 6 Exhaust B                 OK          61 degrees C / 141 degrees F
        FPC 6 LU 0 TSen                 OK          58 degrees C / 136 degrees F
        FPC 6 LU 0 Chip                 OK          62 degrees C / 143 degrees F
        FPC 6 LU 1 TSen                 OK          58 degrees C / 136 degrees F
        FPC 6 LU 1 Chip                 OK          64 degrees C / 147 degrees F
        FPC 6 LU 2 TSen                 OK          58 degrees C / 136 degrees F
        FPC 6 LU 2 Chip                 OK          56 degrees C / 132 degrees F
        FPC 6 LU 3 TSen                 OK          58 degrees C / 136 degrees F
        FPC 6 LU 3 Chip                 OK          56 degrees C / 132 degrees F
        FPC 6 MQ 0 TSen                 OK          51 degrees C / 123 degrees F
        FPC 6 MQ 0 Chip                 OK          58 degrees C / 136 degrees F
        FPC 6 MQ 1 TSen                 OK          51 degrees C / 123 degrees F
        FPC 6 MQ 1 Chip                 OK          61 degrees C / 141 degrees F
        FPC 6 MQ 2 TSen                 OK          51 degrees C / 123 degrees F
        FPC 6 MQ 2 Chip                 OK          51 degrees C / 123 degrees F
        FPC 6 MQ 3 TSen                 OK          51 degrees C / 123 degrees F
        FPC 6 MQ 3 Chip                 OK          51 degrees C / 123 degrees F
        FPC 7 Intake                    OK          42 degrees C / 107 degrees F
        FPC 7 Exhaust A                 OK          50 degrees C / 122 degrees F
        FPC 7 Exhaust B                 OK          61 degrees C / 141 degrees F
        FPC 7 LU 0 TSen                 OK          58 degrees C / 136 degrees F
        FPC 7 LU 0 Chip                 OK          59 degrees C / 138 degrees F
        FPC 7 LU 1 TSen                 OK          58 degrees C / 136 degrees F
        FPC 7 LU 1 Chip                 OK          64 degrees C / 147 degrees F
        FPC 7 LU 2 TSen                 OK          58 degrees C / 136 degrees F
        FPC 7 LU 2 Chip                 OK          54 degrees C / 129 degrees F
        FPC 7 LU 3 TSen                 OK          58 degrees C / 136 degrees F
        FPC 7 LU 3 Chip                 OK          53 degrees C / 127 degrees F
```

```
FPC 7 MQ 0 TSen                OK          51 degrees C / 123 degrees F
FPC 7 MQ 0 Chip                OK          53 degrees C / 127 degrees F
FPC 7 MQ 1 TSen                OK          51 degrees C / 123 degrees F
FPC 7 MQ 1 Chip                OK          54 degrees C / 129 degrees F
FPC 7 MQ 2 TSen                OK          51 degrees C / 123 degrees F
FPC 7 MQ 2 Chip                OK          48 degrees C / 118 degrees F
FPC 7 MQ 3 TSen                OK          51 degrees C / 123 degrees F
FPC 7 MQ 3 Chip                OK          48 degrees C / 118 degrees F
FPC 8 Intake                   OK          42 degrees C / 107 degrees F
FPC 8 Exhaust A                OK          50 degrees C / 122 degrees F
FPC 8 Exhaust B                OK          61 degrees C / 141 degrees F
FPC 8 LU 0 TSen                OK          58 degrees C / 136 degrees F
FPC 8 LU 0 Chip                OK          63 degrees C / 145 degrees F
FPC 8 LU 1 TSen                OK          58 degrees C / 136 degrees F
FPC 8 LU 1 Chip                OK          65 degrees C / 149 degrees F
FPC 8 LU 2 TSen                OK          58 degrees C / 136 degrees F
FPC 8 LU 2 Chip                OK          56 degrees C / 132 degrees F
FPC 8 LU 3 TSen                OK          58 degrees C / 136 degrees F
FPC 8 LU 3 Chip                OK          56 degrees C / 132 degrees F
FPC 8 MQ 0 TSen                OK          50 degrees C / 122 degrees F
FPC 8 MQ 0 Chip                OK          53 degrees C / 127 degrees F
FPC 8 MQ 1 TSen                OK          50 degrees C / 122 degrees F
FPC 8 MQ 1 Chip                OK          58 degrees C / 136 degrees F
FPC 8 MQ 2 TSen                OK          50 degrees C / 122 degrees F
FPC 8 MQ 2 Chip                OK          48 degrees C / 118 degrees F
FPC 8 MQ 3 TSen                OK          50 degrees C / 122 degrees F
FPC 8 MQ 3 Chip                OK          49 degrees C / 120 degrees F
FPC 9 Intake                   OK          43 degrees C / 109 degrees F
FPC 9 Exhaust A                OK          51 degrees C / 123 degrees F
FPC 9 Exhaust B                OK          61 degrees C / 141 degrees F
FPC 9 LU 0 TSen                OK          58 degrees C / 136 degrees F
FPC 9 LU 0 Chip                OK          61 degrees C / 141 degrees F
FPC 9 LU 1 TSen                OK          58 degrees C / 136 degrees F
FPC 9 LU 1 Chip                OK          63 degrees C / 145 degrees F
FPC 9 LU 2 TSen                OK          58 degrees C / 136 degrees F
FPC 9 LU 2 Chip                OK          55 degrees C / 131 degrees F
FPC 9 LU 3 TSen                OK          58 degrees C / 136 degrees F
FPC 9 LU 3 Chip                OK          54 degrees C / 129 degrees F
FPC 9 MQ 0 TSen                OK          52 degrees C / 125 degrees F
FPC 9 MQ 0 Chip                OK          53 degrees C / 127 degrees F
FPC 9 MQ 1 TSen                OK          52 degrees C / 125 degrees F
FPC 9 MQ 1 Chip                OK          54 degrees C / 129 degrees F
FPC 9 MQ 2 TSen                OK          52 degrees C / 125 degrees F
```

```
    FPC 9 MQ 2 Chip                OK         48 degrees C / 118 degrees F
    FPC 9 MQ 3 TSen                OK         52 degrees C / 125 degrees F
    FPC 9 MQ 3 Chip                OK         49 degrees C / 120 degrees F
    FPC 10 Intake                  OK         44 degrees C / 111 degrees F
    FPC 10 Exhaust A               OK         48 degrees C / 118 degrees F
    FPC 10 Exhaust B               OK         54 degrees C / 129 degrees F
    FPC 10 LU 0 TSen               OK         53 degrees C / 127 degrees F
    FPC 10 LU 0 Chip               OK         54 degrees C / 129 degrees F
    FPC 10 LU 1 TSen               OK         53 degrees C / 127 degrees F
    FPC 10 LU 1 Chip               OK         58 degrees C / 136 degrees F
    FPC 10 LU 2 TSen               OK         53 degrees C / 127 degrees F
    FPC 10 LU 2 Chip               OK         51 degrees C / 123 degrees F
    FPC 10 LU 3 TSen               OK         53 degrees C / 127 degrees F
    FPC 10 LU 3 Chip               OK         51 degrees C / 123 degrees F
    FPC 10 MQ 0 TSen               OK         49 degrees C / 120 degrees F
    FPC 10 MQ 0 Chip               OK         50 degrees C / 122 degrees F
    FPC 10 MQ 1 TSen               OK         49 degrees C / 120 degrees F
    FPC 10 MQ 1 Chip               OK         52 degrees C / 125 degrees F
    FPC 10 MQ 2 TSen               OK         49 degrees C / 120 degrees F
    FPC 10 MQ 2 Chip               OK         48 degrees C / 118 degrees F
    FPC 10 MQ 3 TSen               OK         49 degrees C / 120 degrees F
    FPC 10 MQ 3 Chip               OK         48 degrees C / 118 degrees F
    FPC 11 Intake                  OK         39 degrees C / 102 degrees F
    FPC 11 Exhaust A               OK         47 degrees C / 116 degrees F
    FPC 11 Exhaust B               OK         51 degrees C / 123 degrees F
    FPC 11 LU 0 TSen               OK         50 degrees C / 122 degrees F
    FPC 11 LU 0 Chip               OK         52 degrees C / 125 degrees F
    FPC 11 LU 1 TSen               OK         50 degrees C / 122 degrees F
    FPC 11 LU 1 Chip               OK         55 degrees C / 131 degrees F
    FPC 11 LU 2 TSen               OK         50 degrees C / 122 degrees F
    FPC 11 LU 2 Chip               OK         49 degrees C / 120 degrees F
    FPC 11 LU 3 TSen               OK         50 degrees C / 122 degrees F
    FPC 11 LU 3 Chip               OK         49 degrees C / 120 degrees F
    FPC 11 MQ 0 TSen               OK         47 degrees C / 116 degrees F
    FPC 11 MQ 0 Chip               OK         47 degrees C / 116 degrees F
    FPC 11 MQ 1 TSen               OK         47 degrees C / 116 degrees F
    FPC 11 MQ 1 Chip               OK         51 degrees C / 123 degrees F
    FPC 11 MQ 2 TSen               OK         47 degrees C / 116 degrees F
    FPC 11 MQ 2 Chip               OK         45 degrees C / 113 degrees F
    FPC 11 MQ 3 TSen               OK         47 degrees C / 116 degrees F
    FPC 11 MQ 3 Chip               OK         49 degrees C / 120 degrees F
    FPC 12 Intake                  OK         39 degrees C / 102 degrees F
    FPC 12 Exhaust A               OK         47 degrees C / 116 degrees F
```

```
        FPC 12 Exhaust B                OK          50 degrees C / 122 degrees F
        FPC 12 LU 0 TSen                OK          49 degrees C / 120 degrees F
        FPC 12 LU 0 Chip                OK          51 degrees C / 123 degrees F
        FPC 12 LU 1 TSen                OK          49 degrees C / 120 degrees F
        FPC 12 LU 1 Chip                OK          54 degrees C / 129 degrees F
        FPC 12 LU 2 TSen                OK          49 degrees C / 120 degrees F
        FPC 12 LU 2 Chip                OK          47 degrees C / 116 degrees F
        FPC 12 LU 3 TSen                OK          49 degrees C / 120 degrees F
        FPC 12 LU 3 Chip                OK          49 degrees C / 120 degrees F
        FPC 12 MQ 0 TSen                OK          47 degrees C / 116 degrees F
        FPC 12 MQ 0 Chip                OK          46 degrees C / 114 degrees F
        FPC 12 MQ 1 TSen                OK          47 degrees C / 116 degrees F
        FPC 12 MQ 1 Chip                OK          51 degrees C / 123 degrees F
        FPC 12 MQ 2 TSen                OK          47 degrees C / 116 degrees F
        FPC 12 MQ 2 Chip                OK          45 degrees C / 113 degrees F
        FPC 12 MQ 3 TSen                OK          47 degrees C / 116 degrees F
        FPC 12 MQ 3 Chip                OK          47 degrees C / 116 degrees F
        FPC 13 Intake                   OK          40 degrees C / 104 degrees F
        FPC 13 Exhaust A                OK          48 degrees C / 118 degrees F
        FPC 13 Exhaust B                OK          51 degrees C / 123 degrees F
        FPC 13 LU 0 TSen                OK          50 degrees C / 122 degrees F
        FPC 13 LU 0 Chip                OK          52 degrees C / 125 degrees F
        FPC 13 LU 1 TSen                OK          50 degrees C / 122 degrees F
        FPC 13 LU 1 Chip                OK          54 degrees C / 129 degrees F
        FPC 13 LU 2 TSen                OK          50 degrees C / 122 degrees F
        FPC 13 LU 2 Chip                OK          48 degrees C / 118 degrees F
        FPC 13 LU 3 TSen                OK          50 degrees C / 122 degrees F
        FPC 13 LU 3 Chip                OK          48 degrees C / 118 degrees F
        FPC 13 MQ 0 TSen                OK          48 degrees C / 118 degrees F
        FPC 13 MQ 0 Chip                OK          47 degrees C / 116 degrees F
        FPC 13 MQ 1 TSen                OK          48 degrees C / 118 degrees F
        FPC 13 MQ 1 Chip                OK          51 degrees C / 123 degrees F
        FPC 13 MQ 2 TSen                OK          48 degrees C / 118 degrees F
        FPC 13 MQ 2 Chip                OK          46 degrees C / 114 degrees F
        FPC 13 MQ 3 TSen                OK          48 degrees C / 118 degrees F
        FPC 13 MQ 3 Chip                OK          47 degrees C / 116 degrees F
        FPC 14 Intake                   OK          41 degrees C / 105 degrees F
        FPC 14 Exhaust A                OK          49 degrees C / 120 degrees F
        FPC 14 Exhaust B                OK          50 degrees C / 122 degrees F
        FPC 14 LU 0 TSen                OK          49 degrees C / 120 degrees F
        FPC 14 LU 0 Chip                OK          50 degrees C / 122 degrees F
        FPC 14 LU 1 TSen                OK          49 degrees C / 120 degrees F
        FPC 14 LU 1 Chip                OK          54 degrees C / 129 degrees F
```

```
    FPC 14 LU 2 TSen              OK        49 degrees C / 120 degrees F
    FPC 14 LU 2 Chip              OK        48 degrees C / 118 degrees F
    FPC 14 LU 3 TSen              OK        49 degrees C / 120 degrees F
    FPC 14 LU 3 Chip              OK        50 degrees C / 122 degrees F
    FPC 14 MQ 0 TSen              OK        48 degrees C / 118 degrees F
    FPC 14 MQ 0 Chip              OK        48 degrees C / 118 degrees F
    FPC 14 MQ 1 TSen              OK        48 degrees C / 118 degrees F
    FPC 14 MQ 1 Chip              OK        52 degrees C / 125 degrees F
    FPC 14 MQ 2 TSen              OK        48 degrees C / 118 degrees F
    FPC 14 MQ 2 Chip              OK        47 degrees C / 116 degrees F
    FPC 14 MQ 3 TSen              OK        48 degrees C / 118 degrees F
    FPC 14 MQ 3 Chip              OK        50 degrees C / 122 degrees F
    FPC 15 Intake                 OK        42 degrees C / 107 degrees F
    FPC 15 Exhaust A              OK        51 degrees C / 123 degrees F
    FPC 15 Exhaust B              OK        52 degrees C / 125 degrees F
    FPC 15 LU 0 TSen              OK        52 degrees C / 125 degrees F
    FPC 15 LU 0 Chip              OK        55 degrees C / 131 degrees F
    FPC 15 LU 1 TSen              OK        52 degrees C / 125 degrees F
    FPC 15 LU 1 Chip              OK        59 degrees C / 138 degrees F
    FPC 15 LU 2 TSen              OK        52 degrees C / 125 degrees F
    FPC 15 LU 2 Chip              OK        50 degrees C / 122 degrees F
    FPC 15 LU 3 TSen              OK        52 degrees C / 125 degrees F
    FPC 15 LU 3 Chip              OK        51 degrees C / 123 degrees F
    FPC 15 MQ 0 TSen              OK        51 degrees C / 123 degrees F
    FPC 15 MQ 0 Chip              OK        53 degrees C / 127 degrees F
    FPC 15 MQ 1 TSen              OK        51 degrees C / 123 degrees F
    FPC 15 MQ 1 Chip              OK        60 degrees C / 140 degrees F
    FPC 15 MQ 2 TSen              OK        51 degrees C / 123 degrees F
    FPC 15 MQ 2 Chip              OK        52 degrees C / 125 degrees F
    FPC 15 MQ 3 TSen              OK        51 degrees C / 123 degrees F
    FPC 15 MQ 3 Chip              OK        53 degrees C / 127 degrees F
    FPC 16 Intake                 OK        44 degrees C / 111 degrees F
    FPC 16 Exhaust A              OK        51 degrees C / 123 degrees F
    FPC 16 Exhaust B              OK        53 degrees C / 127 degrees F
    FPC 16 LU 0 TSen              OK        51 degrees C / 123 degrees F
    FPC 16 LU 0 Chip              OK        52 degrees C / 125 degrees F
    FPC 16 LU 1 TSen              OK        51 degrees C / 123 degrees F
    FPC 16 LU 1 Chip              OK        56 degrees C / 132 degrees F
    FPC 16 LU 2 TSen              OK        51 degrees C / 123 degrees F
    FPC 16 LU 2 Chip              OK        50 degrees C / 122 degrees F
    FPC 16 LU 3 TSen              OK        51 degrees C / 123 degrees F
    FPC 16 LU 3 Chip              OK        50 degrees C / 122 degrees F
    FPC 16 MQ 0 TSen              OK        51 degrees C / 123 degrees F
```

```
    FPC 16 MQ 0 Chip            OK        50 degrees C / 122 degrees F
    FPC 16 MQ 1 TSen            OK        51 degrees C / 123 degrees F
    FPC 16 MQ 1 Chip            OK        55 degrees C / 131 degrees F
    FPC 16 MQ 2 TSen            OK        51 degrees C / 123 degrees F
    FPC 16 MQ 2 Chip            OK        49 degrees C / 120 degrees F
    FPC 16 MQ 3 TSen            OK        51 degrees C / 123 degrees F
    FPC 16 MQ 3 Chip            OK        52 degrees C / 125 degrees F
    FPC 17 Intake               OK        45 degrees C / 113 degrees F
    FPC 17 Exhaust A            OK        52 degrees C / 125 degrees F
    FPC 17 Exhaust B            OK        55 degrees C / 131 degrees F
    FPC 17 LU 0 TSen            OK        54 degrees C / 129 degrees F
    FPC 17 LU 0 Chip            OK        57 degrees C / 134 degrees F
    FPC 17 LU 1 TSen            OK        54 degrees C / 129 degrees F
    FPC 17 LU 1 Chip            OK        61 degrees C / 141 degrees F
    FPC 17 LU 2 TSen            OK        54 degrees C / 129 degrees F
    FPC 17 LU 2 Chip            OK        54 degrees C / 129 degrees F
    FPC 17 LU 3 TSen            OK        54 degrees C / 129 degrees F
    FPC 17 LU 3 Chip            OK        55 degrees C / 131 degrees F
    FPC 17 MQ 0 TSen            OK        53 degrees C / 127 degrees F
    FPC 17 MQ 0 Chip            OK        53 degrees C / 127 degrees F
    FPC 17 MQ 1 TSen            OK        53 degrees C / 127 degrees F
    FPC 17 MQ 1 Chip            OK        57 degrees C / 134 degrees F
    FPC 17 MQ 2 TSen            OK        53 degrees C / 127 degrees F
    FPC 17 MQ 2 Chip            OK        51 degrees C / 123 degrees F
    FPC 17 MQ 3 TSen            OK        53 degrees C / 127 degrees F
    FPC 17 MQ 3 Chip            OK        54 degrees C / 129 degrees F
    FPC 18 Intake               OK        46 degrees C / 114 degrees F
    FPC 18 Exhaust A            OK        53 degrees C / 127 degrees F
    FPC 18 Exhaust B            OK        57 degrees C / 134 degrees F
    FPC 18 LU 0 TSen            OK        56 degrees C / 132 degrees F
    FPC 18 LU 0 Chip            OK        58 degrees C / 136 degrees F
    FPC 18 LU 1 TSen            OK        56 degrees C / 132 degrees F
    FPC 18 LU 1 Chip            OK        63 degrees C / 145 degrees F
    FPC 18 LU 2 TSen            OK        56 degrees C / 132 degrees F
    FPC 18 LU 2 Chip            OK        54 degrees C / 129 degrees F
    FPC 18 LU 3 TSen            OK        56 degrees C / 132 degrees F
    FPC 18 LU 3 Chip            OK        56 degrees C / 132 degrees F
    FPC 18 MQ 0 TSen            OK        54 degrees C / 129 degrees F
    FPC 18 MQ 0 Chip            OK        57 degrees C / 134 degrees F
    FPC 18 MQ 1 TSen            OK        54 degrees C / 129 degrees F
    FPC 18 MQ 1 Chip            OK        62 degrees C / 143 degrees F
    FPC 18 MQ 2 TSen            OK        54 degrees C / 129 degrees F
    FPC 18 MQ 2 Chip            OK        53 degrees C / 127 degrees F
```

```
    FPC 18 MQ 3 TSen            OK        54 degrees C / 129 degrees F
    FPC 18 MQ 3 Chip            OK        56 degrees C / 132 degrees F
    FPC 19 Intake               OK        49 degrees C / 120 degrees F
    FPC 19 Exhaust A            OK        56 degrees C / 132 degrees F
    FPC 19 Exhaust B            OK        62 degrees C / 143 degrees F
    FPC 19 LU 0 TSen            OK        62 degrees C / 143 degrees F
    FPC 19 LU 0 Chip            OK        63 degrees C / 145 degrees F
    FPC 19 LU 1 TSen            OK        62 degrees C / 143 degrees F
    FPC 19 LU 1 Chip            OK        69 degrees C / 156 degrees F
    FPC 19 LU 2 TSen            OK        62 degrees C / 143 degrees F
    FPC 19 LU 2 Chip            OK        61 degrees C / 141 degrees F
    FPC 19 LU 3 TSen            OK        62 degrees C / 143 degrees F
    FPC 19 LU 3 Chip            OK        62 degrees C / 143 degrees F
    FPC 19 MQ 0 TSen            OK        58 degrees C / 136 degrees F
    FPC 19 MQ 0 Chip            OK        62 degrees C / 143 degrees F
    FPC 19 MQ 1 TSen            OK        58 degrees C / 136 degrees F
    FPC 19 MQ 1 Chip            OK        64 degrees C / 147 degrees F
    FPC 19 MQ 2 TSen            OK        58 degrees C / 136 degrees F
    FPC 19 MQ 2 Chip            OK        59 degrees C / 138 degrees F
    FPC 19 MQ 3 TSen            OK        58 degrees C / 136 degrees F
    FPC 19 MQ 3 Chip            OK        60 degrees C / 140 degrees F
    ADC 0 Intake                OK        40 degrees C / 104 degrees F
    ADC 0 Exhaust               OK        50 degrees C / 122 degrees F
    ADC 0 ADC-XF1               OK        58 degrees C / 136 degrees F
    ADC 0 ADC-XF0               OK        63 degrees C / 145 degrees F
    ADC 1 Intake                OK        38 degrees C / 100 degrees F
    ADC 1 Exhaust               OK        48 degrees C / 118 degrees F
    ADC 1 ADC-XF1               OK        59 degrees C / 138 degrees F
    ADC 1 ADC-XF0               OK        61 degrees C / 141 degrees F
    ADC 2 Intake                OK        36 degrees C / 96 degrees F
    ADC 2 Exhaust               OK        50 degrees C / 122 degrees F
    ADC 2 ADC-XF1               OK        53 degrees C / 127 degrees F
    ADC 2 ADC-XF0               OK        59 degrees C / 138 degrees F
    ADC 3 Intake                OK        39 degrees C / 102 degrees F
    ADC 3 Exhaust               OK        49 degrees C / 120 degrees F
    ADC 3 ADC-XF1               OK        61 degrees C / 141 degrees F
    ADC 3 ADC-XF0               OK        62 degrees C / 143 degrees F
    ADC 4 Intake                OK        39 degrees C / 102 degrees F
    ADC 4 Exhaust               OK        49 degrees C / 120 degrees F
    ADC 4 ADC-XF1               OK        60 degrees C / 140 degrees F
    ADC 4 ADC-XF0               OK        61 degrees C / 141 degrees F
    ADC 5 Intake                OK        38 degrees C / 100 degrees F
    ADC 5 Exhaust               OK        52 degrees C / 125 degrees F
```

```
ADC 5 ADC-XF1                    OK          55 degrees C / 131 degrees F
ADC 5 ADC-XF0                    OK          65 degrees C / 149 degrees F
ADC 6 Intake                     OK          39 degrees C / 102 degrees F
ADC 6 Exhaust                    OK          51 degrees C / 123 degrees F
ADC 6 ADC-XF1                    OK          58 degrees C / 136 degrees F
ADC 6 ADC-XF0                    OK          63 degrees C / 145 degrees F
ADC 7 Intake                     OK          39 degrees C / 102 degrees F
ADC 7 Exhaust                    OK          52 degrees C / 125 degrees F
ADC 7 ADC-XF1                    OK          61 degrees C / 141 degrees F
ADC 7 ADC-XF0                    OK          68 degrees C / 154 degrees F
ADC 8 Intake                     OK          39 degrees C / 102 degrees F
ADC 8 Exhaust                    OK          50 degrees C / 122 degrees F
ADC 8 ADC-XF1                    OK          64 degrees C / 147 degrees F
ADC 8 ADC-XF0                    OK          63 degrees C / 145 degrees F
ADC 9 Intake                     OK          41 degrees C / 105 degrees F
ADC 9 Exhaust                    OK          50 degrees C / 122 degrees F
ADC 9 ADC-XF1                    OK          60 degrees C / 140 degrees F
ADC 9 ADC-XF0                    OK          62 degrees C / 143 degrees F
ADC 10 Intake                    OK          46 degrees C / 114 degrees F
ADC 10 Exhaust                   OK          53 degrees C / 127 degrees F
ADC 10 ADC-XF1                   OK          66 degrees C / 150 degrees F
ADC 10 ADC-XF0                   OK          65 degrees C / 149 degrees F
ADC 11 Intake                    OK          46 degrees C / 114 degrees F
ADC 11 Exhaust                   OK          53 degrees C / 127 degrees F
ADC 11 ADC-XF1                   OK          63 degrees C / 145 degrees F
ADC 11 ADC-XF0                   OK          64 degrees C / 147 degrees F
ADC 12 Intake                    OK          47 degrees C / 116 degrees F
ADC 12 Exhaust                   OK          53 degrees C / 127 degrees F
ADC 12 ADC-XF1                   OK          65 degrees C / 149 degrees F
ADC 12 ADC-XF0                   OK          65 degrees C / 149 degrees F
ADC 13 Intake                    OK          48 degrees C / 118 degrees F
ADC 13 Exhaust                   OK          55 degrees C / 131 degrees F
ADC 13 ADC-XF1                   OK          65 degrees C / 149 degrees F
ADC 13 ADC-XF0                   OK          67 degrees C / 152 degrees F
ADC 14 Intake                    OK          49 degrees C / 120 degrees F
ADC 14 Exhaust                   OK          57 degrees C / 134 degrees F
ADC 14 ADC-XF1                   OK          68 degrees C / 154 degrees F
ADC 14 ADC-XF0                   OK          72 degrees C / 161 degrees F
ADC 15 Intake                    OK          50 degrees C / 122 degrees F
ADC 15 Exhaust                   OK          56 degrees C / 132 degrees F
ADC 15 ADC-XF1                   OK          68 degrees C / 154 degrees F
ADC 15 ADC-XF0                   OK          68 degrees C / 154 degrees F
ADC 16 Intake                    OK          51 degrees C / 123 degrees F
```

```
      ADC 16 Exhaust                OK         57 degrees C / 134 degrees F
      ADC 16 ADC-XF1                OK         67 degrees C / 152 degrees F
      ADC 16 ADC-XF0                OK         68 degrees C / 154 degrees F
      ADC 17 Intake                 OK         51 degrees C / 123 degrees F
      ADC 17 Exhaust                OK         57 degrees C / 134 degrees F
      ADC 17 ADC-XF1                OK         69 degrees C / 156 degrees F
      ADC 17 ADC-XF0                OK         69 degrees C / 156 degrees F
      ADC 18 Intake                 OK         52 degrees C / 125 degrees F
      ADC 18 Exhaust                OK         58 degrees C / 136 degrees F
      ADC 18 ADC-XF1                OK         67 degrees C / 152 degrees F
      ADC 18 ADC-XF0                OK         72 degrees C / 161 degrees F
      ADC 19 Intake                 OK         50 degrees C / 122 degrees F
      ADC 19 Exhaust                OK         58 degrees C / 136 degrees F
      ADC 19 ADC-XF1                OK         68 degrees C / 154 degrees F
      ADC 19 ADC-XF0                OK         71 degrees C / 159 degrees F
```

show chassis environment monitored (MX2008 Router)

```
user@host> show chassis environment monitored
Class Item                          Status     Measurement
Temp  CB 0 Inlet1                   OK         37 degrees C / 98 degrees F
      CB 0 Inlet2                   OK         46 degrees C / 114 degrees F
      CB 0 Inlet3                   OK         44 degrees C / 111 degrees F
      CB 0 Inlet4                   OK         42 degrees C / 107 degrees F
      CB 0 Exhaust1                 OK         30 degrees C / 86 degrees F
      CB 0 Exhaust2                 OK         40 degrees C / 104 degrees F
      CB 0 Exhaust3                 OK         48 degrees C / 118 degrees F
      CB 0 Exhaust4                 OK         46 degrees C / 114 degrees F
      CB 1 Inlet1                   OK         30 degrees C / 86 degrees F
      CB 1 Inlet2                   OK         31 degrees C / 87 degrees F
      CB 1 Inlet3                   OK         29 degrees C / 84 degrees F
      CB 1 Inlet4                   OK         32 degrees C / 89 degrees F
      CB 1 Exhaust1                 OK         30 degrees C / 86 degrees F
      CB 1 Exhaust2                 OK         33 degrees C / 91 degrees F
      CB 1 Exhaust3                 OK         34 degrees C / 93 degrees F
      CB 1 Exhaust4                 OK         35 degrees C / 95 degrees F
      Routing Engine 0 CPU          OK         76 degrees C / 168 degrees F
      Routing Engine 1 CPU          OK         47 degrees C / 116 degrees F
      SFB 0 Inlet2                  OK         44 degrees C / 111 degrees F
      SFB 0 Exhaust1                OK         39 degrees C / 102 degrees F
      SFB 0 Inlet1                  OK         41 degrees C / 105 degrees F
```

```
    SFB 0 Exhaust2                OK          45 degrees C / 113 degrees F
    SFB 0 SFB2-PF-local           OK          45 degrees C / 113 degrees F
    SFB 0 SFB2-PF-die             OK          51 degrees C / 123 degrees F
    SFB 1 Inlet2                  OK          30 degrees C / 86 degrees F
    SFB 1 Exhaust1                OK          27 degrees C / 80 degrees F
    SFB 1 Inlet1                  OK          28 degrees C / 82 degrees F
    SFB 1 Exhaust2                OK          31 degrees C / 87 degrees F
    SFB 1 SFB2-PF-local           OK          30 degrees C / 86 degrees F
    SFB 1 SFB2-PF-die             OK          37 degrees C / 98 degrees F
    SFB 2 Inlet2                  OK          28 degrees C / 82 degrees F
    SFB 2 Exhaust1                OK          26 degrees C / 78 degrees F
    SFB 2 Inlet1                  OK          27 degrees C / 80 degrees F
    SFB 2 Exhaust2                OK          28 degrees C / 82 degrees F
    SFB 2 SFB2-PF-local           OK          27 degrees C / 80 degrees F
    SFB 2 SFB2-PF-die             OK          33 degrees C / 91 degrees F
    SFB 3 Inlet2                  OK          28 degrees C / 82 degrees F
    SFB 3 Exhaust1                OK          26 degrees C / 78 degrees F
    SFB 3 Inlet1                  OK          26 degrees C / 78 degrees F
    SFB 3 Exhaust2                OK          28 degrees C / 82 degrees F
    SFB 3 SFB2-PF-local           OK          27 degrees C / 80 degrees F
    SFB 3 SFB2-PF-die             OK          33 degrees C / 91 degrees F
    SFB 4 Inlet2                  OK          28 degrees C / 82 degrees F
    SFB 4 Exhaust1                OK          26 degrees C / 78 degrees F
    SFB 4 Inlet1                  OK          26 degrees C / 78 degrees F
    SFB 4 Exhaust2                OK          28 degrees C / 82 degrees F
    SFB 4 SFB2-PF-local           OK          27 degrees C / 80 degrees F
    SFB 4 SFB2-PF-die             OK          32 degrees C / 89 degrees F
    SFB 5 Inlet2                  OK          29 degrees C / 84 degrees F
    SFB 5 Exhaust1                OK          27 degrees C / 80 degrees F
    SFB 5 Inlet1                  OK          27 degrees C / 80 degrees F
    SFB 5 Exhaust2                OK          29 degrees C / 84 degrees F
    SFB 5 SFB2-PF-local           OK          28 degrees C / 82 degrees F
    SFB 5 SFB2-PF-die             OK          34 degrees C / 93 degrees F
    SFB 6 Inlet2                  OK          34 degrees C / 93 degrees F
    SFB 6 Exhaust1                OK          32 degrees C / 89 degrees F
    SFB 6 Inlet1                  OK          31 degrees C / 87 degrees F
    SFB 6 Exhaust2                OK          34 degrees C / 93 degrees F
    SFB 6 SFB2-PF-local           OK          34 degrees C / 93 degrees F
    SFB 6 SFB2-PF-die             OK          40 degrees C / 104 degrees F
    SFB 7 Inlet2                  OK          29 degrees C / 84 degrees F
    SFB 7 Exhaust1                OK          28 degrees C / 82 degrees F
    SFB 7 Inlet1                  OK          29 degrees C / 84 degrees F
    SFB 7 Exhaust2                OK          29 degrees C / 84 degrees F
```

```
    SFB 7 SFB2-PF-local          OK        28 degrees C / 82 degrees F
    SFB 7 SFB2-PF-die            OK        33 degrees C / 91 degrees F
    FPC 0 Intake                 OK        29 degrees C / 84 degrees F
    FPC 0 Exhaust A              OK        43 degrees C / 109 degrees F
    FPC 0 Exhaust B              OK        42 degrees C / 107 degrees F
    FPC 0 XL 0 TSen              OK        39 degrees C / 102 degrees F
    FPC 0 XL 0 Chip              OK        53 degrees C / 127 degrees F
    FPC 0 XL 0 XR2 0 TSen        OK        39 degrees C / 102 degrees F
    FPC 0 XL 0 XR2 0 Chip        OK        60 degrees C / 140 degrees F
    FPC 0 XL 0 XR2 1 TSen        OK        39 degrees C / 102 degrees F
    FPC 0 XL 0 XR2 1 Chip        OK        60 degrees C / 140 degrees F
    FPC 0 XL 1 TSen              OK        30 degrees C / 86 degrees F
    FPC 0 XL 1 Chip              OK        43 degrees C / 109 degrees F
    FPC 0 XL 1 XR2 0 TSen        OK        30 degrees C / 86 degrees F
    FPC 0 XL 1 XR2 0 Chip        OK        50 degrees C / 122 degrees F
    FPC 0 XL 1 XR2 1 TSen        OK        30 degrees C / 86 degrees F
    FPC 0 XL 1 XR2 1 Chip        OK        51 degrees C / 123 degrees F
    FPC 0 XM 0 TSen              OK        42 degrees C / 107 degrees F
    FPC 0 XM 0 Chip              OK        50 degrees C / 122 degrees F
    FPC 0 XM 1 TSen              OK        42 degrees C / 107 degrees F
    FPC 0 XM 1 Chip              OK        42 degrees C / 107 degrees F
    FPC 0 XM 2 TSen              OK        42 degrees C / 107 degrees F
    FPC 0 XM 2 Chip              OK        42 degrees C / 107 degrees F
    FPC 0 XM 3 TSen              OK        42 degrees C / 107 degrees F
    FPC 0 XM 3 Chip              OK        40 degrees C / 104 degrees F
    FPC 0 PCIe Switch TSen       OK        42 degrees C / 107 degrees F
    FPC 0 PCIe Switch Chip       OK        22 degrees C / 71 degrees F
    FPC 1 Intake                 OK        29 degrees C / 84 degrees F
    FPC 1 Exhaust A              OK        52 degrees C / 125 degrees F
    FPC 1 Exhaust B              OK        44 degrees C / 111 degrees F
    FPC 1 EA0 TSen               OK        55 degrees C / 131 degrees F
    FPC 1 EA0 Chip               OK        48 degrees C / 118 degrees F
    FPC 1 EA0_XR0 TSen           OK        55 degrees C / 131 degrees F
    FPC 1 EA0_XR0 Chip           OK        57 degrees C / 134 degrees F
    FPC 1 EA0_XR1 TSen           OK        55 degrees C / 131 degrees F
    FPC 1 EA0_XR1 Chip           OK        54 degrees C / 129 degrees F
    FPC 1 EA1 TSen               OK        55 degrees C / 131 degrees F
    FPC 1 EA1 Chip               OK        51 degrees C / 123 degrees F
    FPC 1 EA1_XR0 TSen           OK        55 degrees C / 131 degrees F
    FPC 1 EA1_XR0 Chip           OK        58 degrees C / 136 degrees F
    FPC 1 EA1_XR1 TSen           OK        55 degrees C / 131 degrees F
    FPC 1 EA1_XR1 Chip           OK        59 degrees C / 138 degrees F
    FPC 1 PEX TSen               OK        55 degrees C / 131 degrees F
```

```
    FPC 1 PEX Chip                OK         39 degrees C / 102 degrees F
    FPC 1 EA2 TSen                OK         43 degrees C / 109 degrees F
    FPC 1 EA2 Chip                OK         39 degrees C / 102 degrees F
    FPC 1 EA2_XR0 TSen            OK         43 degrees C / 109 degrees F
    FPC 1 EA2_XR0 Chip            OK         46 degrees C / 114 degrees F
    FPC 1 EA2_XR1 TSen            OK         43 degrees C / 109 degrees F
    FPC 1 EA2_XR1 Chip            OK         43 degrees C / 109 degrees F
    FPC 1 EA3 TSen                OK         43 degrees C / 109 degrees F
    FPC 1 EA3 Chip                OK         40 degrees C / 104 degrees F
    FPC 1 EA3_XR0 TSen            OK         43 degrees C / 109 degrees F
    FPC 1 EA3_XR0 Chip            OK         51 degrees C / 123 degrees F
    FPC 1 EA3_XR1 TSen            OK         43 degrees C / 109 degrees F
    FPC 1 EA3_XR1 Chip            OK         47 degrees C / 116 degrees F
    FPC 1 EA0_HMC0 Logic die      OK         61 degrees C / 141 degrees F
    FPC 1 EA0_HMC0 DRAM botm      OK         58 degrees C / 136 degrees F
    FPC 1 EA0_HMC1 Logic die      OK         62 degrees C / 143 degrees F
    FPC 1 EA0_HMC1 DRAM botm      OK         59 degrees C / 138 degrees F
    FPC 1 EA0_HMC2 Logic die      OK         59 degrees C / 138 degrees F
    FPC 1 EA0_HMC2 DRAM botm      OK         56 degrees C / 132 degrees F
    FPC 1 EA1_HMC0 Logic die      OK         66 degrees C / 150 degrees F
    FPC 1 EA1_HMC0 DRAM botm      OK         63 degrees C / 145 degrees F
    FPC 1 EA1_HMC1 Logic die      OK         65 degrees C / 149 degrees F
    FPC 1 EA1_HMC1 DRAM botm      OK         62 degrees C / 143 degrees F
    FPC 1 EA1_HMC2 Logic die      OK         63 degrees C / 145 degrees F
    FPC 1 EA1_HMC2 DRAM botm      OK         60 degrees C / 140 degrees F
    FPC 1 EA2_HMC0 Logic die      OK         51 degrees C / 123 degrees F
    FPC 1 EA2_HMC0 DRAM botm      OK         48 degrees C / 118 degrees F
    FPC 1 EA2_HMC1 Logic die      OK         55 degrees C / 131 degrees F
    FPC 1 EA2_HMC1 DRAM botm      OK         52 degrees C / 125 degrees F
    FPC 1 EA2_HMC2 Logic die      OK         52 degrees C / 125 degrees F
    FPC 1 EA2_HMC2 DRAM botm      OK         49 degrees C / 120 degrees F
    FPC 1 EA3_HMC0 Logic die      OK         51 degrees C / 123 degrees F
    FPC 1 EA3_HMC0 DRAM botm      OK         48 degrees C / 118 degrees F
    FPC 1 EA3_HMC1 Logic die      OK         51 degrees C / 123 degrees F
    FPC 1 EA3_HMC1 DRAM botm      OK         48 degrees C / 118 degrees F
    FPC 1 EA3_HMC2 Logic die      OK         52 degrees C / 125 degrees F
    FPC 1 EA3_HMC2 DRAM botm      OK         49 degrees C / 120 degrees F
    FPC 7 Intake                  OK         31 degrees C / 87 degrees F
    FPC 7 Exhaust A               OK         46 degrees C / 114 degrees F
    FPC 7 Exhaust B               OK         38 degrees C / 100 degrees F
    FPC 7 QX 0 TSen               OK         49 degrees C / 120 degrees F
    FPC 7 QX 0 Chip               OK         52 degrees C / 125 degrees F
    FPC 7 LU 0 TCAM TSen          OK         49 degrees C / 120 degrees F
```

```
      FPC 7 LU 0 TCAM Chip          OK         52 degrees C / 125 degrees F
      FPC 7 LU 0 TSen               OK         49 degrees C / 120 degrees F
      FPC 7 LU 0 Chip               OK         51 degrees C / 123 degrees F
      FPC 7 MQ 0 TSen               OK         49 degrees C / 120 degrees F
      FPC 7 MQ 0 Chip               OK         55 degrees C / 131 degrees F
      FPC 7 QX 1 TSen               OK         41 degrees C / 105 degrees F
      FPC 7 QX 1 Chip               OK         42 degrees C / 107 degrees F
      FPC 7 LU 1 TCAM TSen          OK         41 degrees C / 105 degrees F
      FPC 7 LU 1 TCAM Chip          OK         43 degrees C / 109 degrees F
      FPC 7 LU 1 TSen               OK         41 degrees C / 105 degrees F
      FPC 7 LU 1 Chip               OK         46 degrees C / 114 degrees F
      FPC 7 MQ 1 TSen               OK         41 degrees C / 105 degrees F
      FPC 7 MQ 1 Chip               OK         47 degrees C / 116 degrees F
      ADC 7 Intake                  OK         32 degrees C / 89 degrees F
      ADC 7 Exhaust                 OK         39 degrees C / 102 degrees F
      ADC 7 ADC-XF1                 OK         46 degrees C / 114 degrees F
      ADC 7 ADC-XF0                 OK         54 degrees C / 129 degrees F
```

show chassis environment monitored (MX10008 Router)

```
user@host> show chassis environment monitored
Class Item                          Status     Measurement
      Routing Engine 0 CPU
      Routing Engine 1 CPU
Temp  CB 0 Intake A Temp Sensor     OK         24 degrees C / 75 degrees F
      CB 0 Intake B Temp Sensor     OK         24 degrees C / 75 degrees F
      CB 0 Exhaust A Temp Sensor    OK         28 degrees C / 82 degrees F
      CB 0 Exhaust B Temp Sensor    OK         30 degrees C / 86 degrees F
      CB 0 Middle Temp Sensor       OK         28 degrees C / 82 degrees F
      CB 1 Intake A Temp Sensor     OK         24 degrees C / 75 degrees F
      CB 1 Intake B Temp Sensor     OK         23 degrees C / 73 degrees F
      CB 1 Exhaust A Temp Sensor    OK         27 degrees C / 80 degrees F
      CB 1 Exhaust B Temp Sensor    OK         30 degrees C / 86 degrees F
      CB 1 Middle Temp Sensor       OK         28 degrees C / 82 degrees F
      FPC 0 Intake-A Temp Sensor    OK         32 degrees C / 89 degrees F
      FPC 0 Exhaust-A Temp Sensor   OK         44 degrees C / 111 degrees F
      FPC 0 Exhaust-B Temp Sensor   OK         49 degrees C / 120 degrees F
      FPC 0 EA0 Temp Sensor         OK         67 degrees C / 152 degrees F
      FPC 0 EA0_XR0 Temp Sensor     OK         69 degrees C / 156 degrees F
      FPC 0 EA0_XR1 Temp Sensor     OK         73 degrees C / 163 degrees F
      FPC 0 EA1 Temp Sensor         OK         61 degrees C / 141 degrees F
```

```
    FPC 0 EA1_XR0 Temp Sensor      OK         65 degrees C / 149 degrees F
    FPC 0 EA1_XR1 Temp Sensor      OK         63 degrees C / 145 degrees F
    FPC 0 EA2 Temp Sensor          OK         69 degrees C / 156 degrees F
    FPC 0 EA2_XR0 Temp Sensor      OK         73 degrees C / 163 degrees F
    FPC 0 EA2_XR1 Temp Sensor      OK         72 degrees C / 161 degrees F
    FPC 0 EA3 Temp Sensor          OK         64 degrees C / 147 degrees F
    FPC 0 EA3_XR0 Temp Sensor      OK         66 degrees C / 150 degrees F
    FPC 0 EA3_XR1 Temp Sensor      OK         65 degrees C / 149 degrees F
    FPC 0 EA4 Temp Sensor          OK         71 degrees C / 159 degrees F
    FPC 0 EA4_XR0 Temp Sensor      OK         72 degrees C / 161 degrees F
    FPC 0 EA4_XR1 Temp Sensor      OK         71 degrees C / 159 degrees F
    FPC 0 EA5 Temp Sensor          OK         58 degrees C / 136 degrees F
    FPC 0 EA5_XR0 Temp Sensor      OK         62 degrees C / 143 degrees F
    FPC 0 EA5_XR1 Temp Sensor      OK         63 degrees C / 145 degrees F
    FPC 0 EA0_HMC0 Logic die       OK         75 degrees C / 167 degrees F
    FPC 0 EA0_HMC0 DRAM botm       OK         72 degrees C / 161 degrees F
    FPC 0 EA0_HMC1 Logic die       OK         76 degrees C / 168 degrees F
    FPC 0 EA0_HMC1 DRAM botm       OK         73 degrees C / 163 degrees F
    FPC 0 EA0_HMC2 Logic die       OK         77 degrees C / 170 degrees F
    FPC 0 EA0_HMC2 DRAM botm       OK         74 degrees C / 165 degrees F
    FPC 0 EA1_HMC0 Logic die       OK         73 degrees C / 163 degrees F
    FPC 0 EA1_HMC0 DRAM botm       OK         70 degrees C / 158 degrees F
    FPC 0 EA1_HMC1 Logic die       OK         73 degrees C / 163 degrees F
    FPC 0 EA1_HMC1 DRAM botm       OK         70 degrees C / 158 degrees F
    FPC 0 EA1_HMC2 Logic die       OK         72 degrees C / 161 degrees F
    FPC 0 EA1_HMC2 DRAM botm       OK         69 degrees C / 156 degrees F
    FPC 0 EA2_HMC0 Logic die       OK         80 degrees C / 176 degrees F
    FPC 0 EA2_HMC0 DRAM botm       OK         77 degrees C / 170 degrees F
    FPC 0 EA2_HMC1 Logic die       OK         81 degrees C / 177 degrees F
    FPC 0 EA2_HMC1 DRAM botm       OK         78 degrees C / 172 degrees F
    FPC 0 EA2_HMC2 Logic die       OK         80 degrees C / 176 degrees F
    FPC 0 EA2_HMC2 DRAM botm       OK         77 degrees C / 170 degrees F
    FPC 0 EA3_HMC0 Logic die       OK         77 degrees C / 170 degrees F
    FPC 0 EA3_HMC0 DRAM botm       OK         74 degrees C / 165 degrees F
    FPC 0 EA3_HMC1 Logic die       OK         78 degrees C / 172 degrees F
    FPC 0 EA3_HMC1 DRAM botm       OK         75 degrees C / 167 degrees F
    FPC 0 EA3_HMC2 Logic die       OK         77 degrees C / 170 degrees F
    FPC 0 EA3_HMC2 DRAM botm       OK         74 degrees C / 165 degrees F
    FPC 0 EA4_HMC0 Logic die       OK         80 degrees C / 176 degrees F
    FPC 0 EA4_HMC0 DRAM botm       OK         77 degrees C / 170 degrees F
    FPC 0 EA4_HMC1 Logic die       OK         81 degrees C / 177 degrees F
    FPC 0 EA4_HMC1 DRAM botm       OK         78 degrees C / 172 degrees F
    FPC 0 EA4_HMC2 Logic die       OK         80 degrees C / 176 degrees F
```

```
    FPC 0 EA4_HMC2 DRAM botm       OK         77 degrees C / 170 degrees F
    FPC 0 EA5_HMC0 Logic die       OK         69 degrees C / 156 degrees F
    FPC 0 EA5_HMC0 DRAM botm       OK         66 degrees C / 150 degrees F
    FPC 0 EA5_HMC1 Logic die       OK         68 degrees C / 154 degrees F
    FPC 0 EA5_HMC1 DRAM botm       OK         65 degrees C / 149 degrees F
    FPC 0 EA5_HMC2 Logic die       OK         68 degrees C / 154 degrees F
    FPC 0 EA5_HMC2 DRAM botm       OK         65 degrees C / 149 degrees F
    FPC 2 Intake-A Temp Sensor     OK         33 degrees C / 91 degrees F
    FPC 2 Exhaust-A Temp Sensor    OK         52 degrees C / 125 degrees F
    FPC 2 Exhaust-B Temp Sensor    OK         50 degrees C / 122 degrees F
    FPC 2 EA0 Temp Sensor          OK         71 degrees C / 159 degrees F
    FPC 2 EA0_XR0 Temp Sensor      OK         75 degrees C / 167 degrees F
    FPC 2 EA0_XR1 Temp Sensor      OK         79 degrees C / 174 degrees F
    FPC 2 EA1 Temp Sensor          OK         64 degrees C / 147 degrees F
    FPC 2 EA1_XR0 Temp Sensor      OK         68 degrees C / 154 degrees F
    FPC 2 EA1_XR1 Temp Sensor      OK         66 degrees C / 150 degrees F
    FPC 2 EA2 Temp Sensor          OK         75 degrees C / 167 degrees F
    FPC 2 EA2_XR0 Temp Sensor      OK         81 degrees C / 177 degrees F
    FPC 2 EA2_XR1 Temp Sensor      OK         81 degrees C / 177 degrees F
    FPC 2 EA3 Temp Sensor          OK         67 degrees C / 152 degrees F
    FPC 2 EA3_XR0 Temp Sensor      OK         69 degrees C / 156 degrees F
    FPC 2 EA3_XR1 Temp Sensor      OK         69 degrees C / 156 degrees F
    FPC 2 EA4 Temp Sensor          OK         75 degrees C / 167 degrees F
    FPC 2 EA4_XR0 Temp Sensor      OK         76 degrees C / 168 degrees F
    FPC 2 EA4_XR1 Temp Sensor      OK         75 degrees C / 167 degrees F
    FPC 2 EA5 Temp Sensor          OK         60 degrees C / 140 degrees F
    FPC 2 EA5_XR0 Temp Sensor      OK         64 degrees C / 147 degrees F
    FPC 2 EA5_XR1 Temp Sensor      OK         65 degrees C / 149 degrees F
    FPC 2 EA0_HMC0 Logic die       OK         84 degrees C / 183 degrees F
    FPC 2 EA0_HMC0 DRAM botm       OK         81 degrees C / 177 degrees F
    FPC 2 EA0_HMC1 Logic die       OK         85 degrees C / 185 degrees F
    FPC 2 EA0_HMC1 DRAM botm       OK         82 degrees C / 179 degrees F
    FPC 2 EA0_HMC2 Logic die       OK         83 degrees C / 181 degrees F
    FPC 2 EA0_HMC2 DRAM botm       OK         80 degrees C / 176 degrees F
    FPC 2 EA1_HMC0 Logic die       OK         76 degrees C / 168 degrees F
    FPC 2 EA1_HMC0 DRAM botm       OK         73 degrees C / 163 degrees F
    FPC 2 EA1_HMC1 Logic die       OK         77 degrees C / 170 degrees F
    FPC 2 EA1_HMC1 DRAM botm       OK         74 degrees C / 165 degrees F
    FPC 2 EA1_HMC2 Logic die       OK         76 degrees C / 168 degrees F
    FPC 2 EA1_HMC2 DRAM botm       OK         73 degrees C / 163 degrees F
    FPC 2 EA2_HMC0 Logic die       OK         87 degrees C / 188 degrees F
    FPC 2 EA2_HMC0 DRAM botm       OK         84 degrees C / 183 degrees F
    FPC 2 EA2_HMC1 Logic die       OK         88 degrees C / 190 degrees F
```

```
FPC 2 EA2_HMC1 DRAM botm       OK         85 degrees C / 185 degrees F
FPC 2 EA2_HMC2 Logic die       OK         88 degrees C / 190 degrees F
FPC 2 EA2_HMC2 DRAM botm       OK         85 degrees C / 185 degrees F
FPC 2 EA3_HMC0 Logic die       OK         81 degrees C / 177 degrees F
FPC 2 EA3_HMC0 DRAM botm       OK         78 degrees C / 172 degrees F
FPC 2 EA3_HMC1 Logic die       OK         81 degrees C / 177 degrees F
FPC 2 EA3_HMC1 DRAM botm       OK         78 degrees C / 172 degrees F
FPC 2 EA3_HMC2 Logic die       OK         81 degrees C / 177 degrees F
FPC 2 EA3_HMC2 DRAM botm       OK         78 degrees C / 172 degrees F
FPC 2 EA4_HMC0 Logic die       OK         88 degrees C / 190 degrees F
FPC 2 EA4_HMC0 DRAM botm       OK         85 degrees C / 185 degrees F
FPC 2 EA4_HMC1 Logic die       OK         90 degrees C / 194 degrees F
FPC 2 EA4_HMC1 DRAM botm       OK         87 degrees C / 188 degrees F
FPC 2 EA4_HMC2 Logic die       OK         81 degrees C / 177 degrees F
FPC 2 EA4_HMC2 DRAM botm       OK         78 degrees C / 172 degrees F
FPC 2 EA5_HMC0 Logic die       OK         72 degrees C / 161 degrees F
FPC 2 EA5_HMC0 DRAM botm       OK         69 degrees C / 156 degrees F
FPC 2 EA5_HMC1 Logic die       OK         69 degrees C / 156 degrees F
FPC 2 EA5_HMC1 DRAM botm       OK         66 degrees C / 150 degrees F
FPC 2 EA5_HMC2 Logic die       OK         73 degrees C / 163 degrees F
FPC 2 EA5_HMC2 DRAM botm       OK         70 degrees C / 158 degrees F
FPC 3 Intake-A Temp Sensor     OK         30 degrees C / 86 degrees F
FPC 3 Exhaust-A Temp Sensor    OK         48 degrees C / 118 degrees F
FPC 3 Exhaust-B Temp Sensor    OK         45 degrees C / 113 degrees F
FPC 3 EA0 Temp Sensor          OK         60 degrees C / 140 degrees F
FPC 3 EA0_XR0 Temp Sensor      OK         65 degrees C / 149 degrees F
FPC 3 EA0_XR1 Temp Sensor      OK         67 degrees C / 152 degrees F
FPC 3 EA1 Temp Sensor          OK         54 degrees C / 129 degrees F
FPC 3 EA1_XR0 Temp Sensor      OK         60 degrees C / 140 degrees F
FPC 3 EA1_XR1 Temp Sensor      OK         59 degrees C / 138 degrees F
FPC 3 EA2 Temp Sensor          OK         62 degrees C / 143 degrees F
FPC 3 EA2_XR0 Temp Sensor      OK         67 degrees C / 152 degrees F
FPC 3 EA2_XR1 Temp Sensor      OK         67 degrees C / 152 degrees F
FPC 3 EA3 Temp Sensor          OK         55 degrees C / 131 degrees F
FPC 3 EA3_XR0 Temp Sensor      OK         57 degrees C / 134 degrees F
FPC 3 EA3_XR1 Temp Sensor      OK         57 degrees C / 134 degrees F
FPC 3 EA4 Temp Sensor          OK         68 degrees C / 154 degrees F
FPC 3 EA4_XR0 Temp Sensor      OK         71 degrees C / 159 degrees F
FPC 3 EA4_XR1 Temp Sensor      OK         70 degrees C / 158 degrees F
FPC 3 EA5 Temp Sensor          OK         55 degrees C / 131 degrees F
FPC 3 EA5_XR0 Temp Sensor      OK         58 degrees C / 136 degrees F
FPC 3 EA5_XR1 Temp Sensor      OK         59 degrees C / 138 degrees F
FPC 3 EA0_HMC0 Logic die       OK         69 degrees C / 156 degrees F
```

```
      FPC 3 EA0_HMC0 DRAM botm       OK         66 degrees C / 150 degrees F
      FPC 3 EA0_HMC1 Logic die       OK         70 degrees C / 158 degrees F
      FPC 3 EA0_HMC1 DRAM botm       OK         67 degrees C / 152 degrees F
      FPC 3 EA0_HMC2 Logic die       OK         70 degrees C / 158 degrees F
      FPC 3 EA0_HMC2 DRAM botm       OK         67 degrees C / 152 degrees F
      FPC 3 EA1_HMC0 Logic die       OK         68 degrees C / 154 degrees F
      FPC 3 EA1_HMC0 DRAM botm       OK         65 degrees C / 149 degrees F
      FPC 3 EA1_HMC1 Logic die       OK         65 degrees C / 149 degrees F
      FPC 3 EA1_HMC1 DRAM botm       OK         62 degrees C / 143 degrees F
      FPC 3 EA1_HMC2 Logic die       OK         64 degrees C / 147 degrees F
      FPC 3 EA1_HMC2 DRAM botm       OK         61 degrees C / 141 degrees F
      FPC 3 EA2_HMC0 Logic die       OK         74 degrees C / 165 degrees F
      FPC 3 EA2_HMC0 DRAM botm       OK         71 degrees C / 159 degrees F
      FPC 3 EA2_HMC1 Logic die       OK         76 degrees C / 168 degrees F
      FPC 3 EA2_HMC1 DRAM botm       OK         73 degrees C / 163 degrees F
      FPC 3 EA2_HMC2 Logic die       OK         74 degrees C / 165 degrees F
      FPC 3 EA2_HMC2 DRAM botm       OK         71 degrees C / 159 degrees F
      FPC 3 EA3_HMC0 Logic die       OK         70 degrees C / 158 degrees F
      FPC 3 EA3_HMC0 DRAM botm       OK         67 degrees C / 152 degrees F
      FPC 3 EA3_HMC1 Logic die       OK         68 degrees C / 154 degrees F
      FPC 3 EA3_HMC1 DRAM botm       OK         65 degrees C / 149 degrees F
      FPC 3 EA3_HMC2 Logic die       OK         68 degrees C / 154 degrees F
      FPC 3 EA3_HMC2 DRAM botm       OK         65 degrees C / 149 degrees F
      FPC 3 EA4_HMC0 Logic die       OK         82 degrees C / 179 degrees F
      FPC 3 EA4_HMC0 DRAM botm       OK         79 degrees C / 174 degrees F
      FPC 3 EA4_HMC1 Logic die       OK         80 degrees C / 176 degrees F
      FPC 3 EA4_HMC1 DRAM botm       OK         77 degrees C / 170 degrees F
      FPC 3 EA4_HMC2 Logic die       OK         81 degrees C / 177 degrees F
      FPC 3 EA4_HMC2 DRAM botm       OK         78 degrees C / 172 degrees F
      FPC 3 EA5_HMC0 Logic die       OK         69 degrees C / 156 degrees F
      FPC 3 EA5_HMC0 DRAM botm       OK         66 degrees C / 150 degrees F
      FPC 3 EA5_HMC1 Logic die       OK         70 degrees C / 158 degrees F
      FPC 3 EA5_HMC1 DRAM botm       OK         67 degrees C / 152 degrees F
      FPC 3 EA5_HMC2 Logic die       OK         69 degrees C / 156 degrees F
      FPC 3 EA5_HMC2 DRAM botm       OK         66 degrees C / 150 degrees F
Power PEM 0                          OK         29 degrees C / 84 degrees F
      PEM 1                          OK         27 degrees C / 80 degrees F
      PEM 2                          OK         30 degrees C / 86 degrees F
      PEM 3                          Check
      PEM 4                          Check
      PEM 5                          Check
Fans  Fan Tray 0 Fan 0               OK         Spinning at normal speed
      Fan Tray 0 Fan 1               OK         Spinning at normal speed
```

```
    Fan Tray 0 Fan 2              OK         Spinning at normal speed
    Fan Tray 0 Fan 3              OK         Spinning at normal speed
    Fan Tray 0 Fan 4              Failed
    Fan Tray 0 Fan 5              Failed
    Fan Tray 0 Fan 6              OK         Spinning at normal speed
    Fan Tray 0 Fan 7              OK         Spinning at normal speed
    Fan Tray 0 Fan 8              OK         Spinning at normal speed
    Fan Tray 0 Fan 9              OK         Spinning at normal speed
    Fan Tray 0 Fan 10             OK         Spinning at normal speed
    Fan Tray 1 Fan 0              OK         Spinning at normal speed
    Fan Tray 1 Fan 1              OK         Spinning at normal speed
    Fan Tray 1 Fan 2              OK         Spinning at normal speed
    Fan Tray 1 Fan 3              OK         Spinning at normal speed
    Fan Tray 1 Fan 4              OK         Spinning at normal speed
    Fan Tray 1 Fan 5              OK         Spinning at normal speed
    Fan Tray 1 Fan 6              OK         Spinning at normal speed
    Fan Tray 1 Fan 7              OK         Spinning at normal speed
    Fan Tray 1 Fan 8              OK         Spinning at normal speed
    Fan Tray 1 Fan 9              OK         Spinning at normal speed
    Fan Tray 1 Fan 10             OK         Spinning at normal speed
    SFB 0 Intake-A                OK         33 degrees C / 91 degrees F
    SFB 0 Intake-B                OK         22 degrees C / 71 degrees F
    SFB 0 Exhaust-A               OK         27 degrees C / 80 degrees F
    SFB 0 Exhaust-B               OK         32 degrees C / 89 degrees F
    SFB 0 PF0                     OK         40 degrees C / 104 degrees F
    SFB 0 PF1                     OK         29 degrees C / 84 degrees F
    SFB 1 Intake-A                OK         43 degrees C / 109 degrees F
    SFB 1 Intake-B                OK         21 degrees C / 69 degrees F
    SFB 1 Exhaust-A               OK         25 degrees C / 77 degrees F
    SFB 1 Exhaust-B               OK         44 degrees C / 111 degrees F
    SFB 1 PF0                     OK         50 degrees C / 122 degrees F
    SFB 1 PF1                     OK         29 degrees C / 84 degrees F
    SFB 2 Intake-A                OK         39 degrees C / 102 degrees F
    SFB 2 Intake-B                OK         21 degrees C / 69 degrees F
    SFB 2 Exhaust-A               OK         25 degrees C / 77 degrees F
    SFB 2 Exhaust-B               OK         38 degrees C / 100 degrees F
    SFB 2 PF0                     OK         45 degrees C / 113 degrees F
    SFB 2 PF1                     OK         30 degrees C / 86 degrees F
    SFB 3 Intake-A                OK         37 degrees C / 98 degrees F
    SFB 3 Intake-B                OK         21 degrees C / 69 degrees F
    SFB 3 Exhaust-A               OK         26 degrees C / 78 degrees F
    SFB 3 Exhaust-B               OK         35 degrees C / 95 degrees F
    SFB 3 PF0                     OK         42 degrees C / 107 degrees F
```

```
    SFB 3 PF1                    OK         30 degrees C / 86 degrees F
    SFB 4 Intake-A               OK         31 degrees C / 87 degrees F
    SFB 4 Intake-B               OK         20 degrees C / 68 degrees F
    SFB 4 Exhaust-A              OK         25 degrees C / 77 degrees F
    SFB 4 Exhaust-B              OK         32 degrees C / 89 degrees F
    SFB 4 PF0                    OK         41 degrees C / 105 degrees F
    SFB 4 PF1                    OK         30 degrees C / 86 degrees F
    SFB 5 Intake-A               OK         30 degrees C / 86 degrees F
    SFB 5 Intake-B               OK         21 degrees C / 69 degrees F
    SFB 5 Exhaust-A              OK         26 degrees C / 78 degrees F
    SFB 5 Exhaust-B              OK         30 degrees C / 86 degrees F
    SFB 5 PF0                    OK         35 degrees C / 95 degrees F
    SFB 5 PF1                    OK         34 degrees C / 93 degrees F
```

Release Information

Command introduced in Junos OS Release 12.1.

`all-members`, `local`, and `member` *`member-id`* options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.

RELATED DOCUMENTATION

show chassis environment | 17429

Chassis-Level User Guide

show chassis environment pcg

IN THIS SECTION

- Syntax | 17749
- Description | 17749
- Options | 17749
- Required Privilege Level | 17749
- Output Fields | 17749

Sample Output | 17750
Release Information | 17752

Syntax

```
show chassis environment pcg
<slot>
```

Description

(M40e and M160 routers only) Display environmental information about the Packet Forwarding Engine clock generators (PCGs).

Options

none Display environmental information about both PCGs.

slot (Optional) Display environmental information about an individual PCG. Replace ***slot*** with 0 or 1.

Required Privilege Level

view

Output Fields

Table 231 on page 17749 lists the output fields for the `show chassis environment pcg` command. Output fields are listed in the approximate order in which they appear.

Table 231: show chassis environment pcg Output Fields

Field Name	Field Description
PCG slot status	Slot number: **0** or **1**.

Table 231: show chassis environment pcg Output Fields *(Continued)*

Field Name	Field Description
State	Status of PCG: • **Present**—PCG is detected by the chassis process but is either not supported by the current version of Junos OS or PCG is coming up but is not yet online. • **Online**—PCG is powered down. If **Online**, it can be the **Master clock** or the **Standby clock**. • **Offline**—PCG is powered down. • **Empty**—No PCG is present.
Temperature	Temperature of the air flowing past the PCG.
Frequency	Frequency setting and measurement for the PCG.
Power	Information about the voltage supplied to the PCG. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
BUS Revision	Revision level of the generic bus device.

Sample Output

show chassis environment pcg (M40e Router)

```
user@host> show chassis environment pcg
PCG 0 status:
  State                          Online -  Master clock
  Temperature                 44 degrees C / 111 degrees F
  Frequency:
   Setting                125.00 MHz
   Measurement            124.95 MHz
  Power:
    3.3 V                    3266 mV
```

```
    5.0 V bias               4964 mV
    8.0 V bias               8112 mV
  BUS Revision                 12
PCG 1 status:
  State                           Online - Standby
  Temperature                  47 degrees C / 116 degrees F
  Frequency:
   Setting                125.00 MHz
   Measurement            124.96 MHz
  Power:
    3.3 V                    3271 mV
    5.0 V bias               4979 mV
    8.0 V bias               8117 mV
  BUS Revision                 12
```

show chassis environment pcg (M160 Router)

```
user@host> show chassis environment pcg
PCG 0 status:
  State                           Online -  Master clock
  Temperature                  41 degrees C / 105 degrees F
  Frequency:
   Setting                125.00 MHz
   Measurement            125.03 MHz
  Power:
    3.3 V                    3286 mV
    5.0 V bias               5010 mV
    8.0 V bias               8183 mV
  BUS Revision                 12
PCG 1 status:
  State                           Online - Standby
  Temperature                  43 degrees C / 109 degrees F
  Frequency:
   Setting                125.00 MHz
   Measurement            125.01 MHz
  Power:
    3.3 V                    3288 mV
    5.0 V bias               4993 mV
    8.0 V bias               8197 mV
  BUS Revision                 12
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

request chassis pcg | 14501

show chassis environment pdu

IN THIS SECTION

- Syntax | 17752
- Description | 17752
- Options | 17753
- Required Privilege Level | 17753
- Output Fields | 17753
- Sample Output | 17755
- Release Information | 17760

Syntax

```
show chassis environment pdu
<none>
<slot>
```

Description

Display the environmental status information of a power distribution unit (PDU).

Starting from Junos OS Release 14.1, the `show chassis environment pdu` *`slot`* operational mode command output displays environmental status information for the new DC power supply module (PSM) and PDU that are added to provide power to the high-density FPC—FPC2-PTX-P1A.

Starting from Junos OS Release 14.2, the power management software in Junos OS tracks the PSM power capacity to identify the power available for the PTX5000 router. Each PSM is assigned a power capability value that is equal to its maximum power rating. Therefore, the total input power—power that the chassis draws from a PDU—is the sum of all the online PSMs' maximum rating. Note that to limit the PDU's output power, the power management software adjusts the maximum rating of the PSMs according to the feed—that is, 60 A, 100 A, or 150 A—selected.

Options

none Display environmental information about all PDUs.

slot (Optional) Display environmental information about an individual PDU. For the PTX5000, replace *slot* with `0` or `1`.

Required Privilege Level

view

Output Fields

Table 232 on page 17753 lists the output fields for the `show chassis environment pdu` command. Output fields are listed in the approximate order in which they appear.

Table 232: show chassis environment pdu Output Fields

Field Name	Field Description
`PDU` *slot* `status`	Number of the PDU slot.
`PDU - State`	Status of the PDU. Status can be `Online`, `Present`, or `Absent`.
`PDU - BoostConv`	Status of the booster converter.
`Feed Switch`	Status of the connected input line cord in the AC PDU. Status can be , `60A`, `100A`, or `150A`.
`PDU - Hours Used`	Number of hours the PDU has been operational.

Table 232: show chassis environment pdu Output Fields *(Continued)*

Field Name	Field Description
`PDU - Firmware Version`	Version level of the firmware running on the PDU.
`PSM` *number* `status`	PSM number. PSMs are numbered `0` through `3`.
`PSM - State`	Status of the PSM. Status can be `Online`, `Present`, or `Absent`.
`PSM - Temperature`	Temperature of the air flowing past the PSM.
`PSM - Fans`	Status of the cooling fans associated with the PSM.
`PSM - AC Input`	Status of the AC input for the specified component
`PSM - AC Output`	Status of the AC output for the specified component.
`PSM - DC input`	Status of the DC input for the specified component.
`PSM - DC output`	Status of the DC output for the specified component.
PSM-Health check status	Reason for the health check failure of the PSM.
`PSM - Hours Used`	Number of hours the PSM has been operational.
`PSM - Firmware Version`	Version level of the firmware running on the PSM.

Sample Output

show chassis environment pdu (PTX5000)

```
user@host> show chassis environment pdu 0
PDU 0 status:
  State                      Online
  Hours Used                 4281
  Firmware Version (MCU1)    00.02
  Firmware Version (MCU2)    00.02
  Firmware Version (MCU3)    00.02
  Firmware Version (MCU4)    00.02
PDU 0 PSM 0 status:
  State                      Online
  Temperature                OK   32 degrees C / 89 degrees F
  Fans                       OK
  DC Input                   OK
  DC Output                  OK
  Hours Used                 2864
  Firmware Version           00.04
PDU 0 PSM 1 status:
  State                      Online
  Temperature                OK   30 degrees C / 86 degrees F
  Fans                       OK
  DC Input                   OK
  DC Output                  OK
  Hours Used                 3540
  Firmware Version           00.04
PDU 0 PSM 2 status:
  State                      Online
  Temperature                OK   29 degrees C / 84 degrees F
  Fans                       OK
  DC Input                   OK
  DC Output                  OK
  Hours Used                 3711
  Firmware Version           00.04
PDU 0 PSM 3 status:
  State                      Online
  Temperature                OK   29 degrees C / 84 degrees F
  Fans                       OK
  DC Input                   OK
  DC Output                  OK
```

```
  Hours Used                 4243
  Firmware Version           00.04
```

show chassis environment pdu (PTX5000 Packet Transport Router with DC PSM and PDU)

```
user@host> show chassis environment pdu 1
PDU 1 status:
  State                      Online
  BoostConv                  OK
  Hours Used                 1054
  Firmware Version (MCU1)    03.05
PDU 1 PSM 0 status:
  State                      Empty
PDU 1 PSM 1 status:
  State                      Online
  Temperature                OK   45 degrees C / 113 degrees F
  Fans                       OK
  DC Input                   OK
  DC Output                  OK
  Hours Used                 1027
  Firmware Version           03.07
PDU 1 PSM 2 status:
  State                      Empty
PDU 1 PSM 3 status:
  State                      Online
  Temperature                OK   43 degrees C / 109 degrees F
  Fans                       OK
  DC Input                   OK
  DC Output                  OK
  Hours Used                 1029
  Firmware Version           03.07
PDU 1 PSM 4 status:
  State                      Empty
PDU 1 PSM 5 status:
  State                      Online
  Temperature                OK   46 degrees C / 114 degrees F
  Fans                       OK
  DC Input                   OK
  DC Output                  OK
  Hours Used                 1028
  Firmware Version           03.07
```

```
PDU 1 PSM 6 status:
  State                       Empty
PDU 1 PSM 7 status:
  State                       Online
  Temperature                 OK   46 degrees C / 114 degrees F
  Fans                        OK
  DC Input                    OK
  DC Output                   OK
  Hours Used                  1030
  Firmware Version            03.07
```

show chassis environment pdu (PTX5000 Packet Transport Router with AC PSM and PDU)

```
user@host> show chassis environment pdu 0
PDU 0 status:
  State                       Online
  BoostConv                   OK
  Feed Switch                 150 Amps
  Hours Used                  177
  Firmware Version (MCU1)     03.04
  Firmware Version (MCU2)     03.02
  Firmware Version (MCU3)     03.02
  Firmware Version (MCU4)     03.02
  Firmware Version (MCU5)     03.02
  Firmware Version (MCU6)     03.02
  Firmware Version (MCU7)     03.02
  Firmware Version (MCU8)     03.02
PDU 0 PSM 0 status:
  State                       Online
  Temperature                 OK   28 degrees C / 82 degrees F
  Fans                        OK
  AC Input                    OK
  DC Output                   OK
  Hours Used                  652
  Firmware Version            01.01
```

show chassis environment pdu (PTX5000 Packet Transport Router after a heath check failure on a PSM)

```
user@host> show chassis environment pdu 0
PDU 0 status:
  State                      Online
  BoostConv                  OK
  Feed Switch                20 Amps
  Hours Used                 16706
  Firmware Version (MCU1)    91.02
  Firmware Version (MCU2)    03.02
  Firmware Version (MCU3)    03.02
  Firmware Version (MCU4)    03.02
  Firmware Version (MCU5)    03.02
  Firmware Version (MCU6)    03.02
  Firmware Version (MCU7)    03.02
  Firmware Version (MCU8)    03.02
PDU 0 PSM 0 status:
  State                      Online
  Temperature                OK   29 degrees C / 84 degrees F
  Fans                       OK
  AC Input                   OK
  DC Output                  OK
  Hours Used                 9847
  Firmware Version           01.01
PDU 0 PSM 1 status:
  State                      Online
  Temperature                OK   29 degrees C / 84 degrees F
  Fans                       OK
  AC Input                   OK
  DC Output                  OK
  Hours Used                 5586
  Firmware Version           01.01
PDU 0 PSM 2 status:
  State                      Online
  Temperature                OK   28 degrees C / 82 degrees F
  Fans                       OK
  AC Input                   OK
  DC Output                  OK
  Hours Used                 10425
  Firmware Version           01.01
PDU 0 PSM 3 status:
```

```
  State                       Online
  Temperature                 OK   28 degrees C / 82 degrees F
  Fans                        OK
  AC Input                    OK
  DC Output                   OK
  Hours Used                  9912
  Firmware Version            01.01
PDU 0 PSM 4 status:
  State                       Empty
PDU 0 PSM 5 status:
  State                       Empty
PDU 0 PSM 6 status:
  State                       Present
  Temperature                 OK   32 degrees C / 89 degrees F
  Fans                        OK
  AC Input                    OK
  DC Output                   OK
  Health check status         Failed, PSM set to NOT OK
  Hours Used                  5770
  Firmware Version            01.01
PDU 0 PSM 7 status:
  State                       Online
  Temperature                 OK   33 degrees C / 91 degrees F
  Fans                        OK
  AC Input                    OK
  DC Output                   OK
  Hours Used                  20167
  Firmware Version            01.01
```

show chassis environment pdu (PTX10003 Router with status of AC Input and DC Output)

```
user@root> show chassis environment pdu
PDU 0 status:
  State                       Online
PDU 0 PSM 1 status:
  State                       Online
  Temperature                 25 degrees C / 77 degrees F
  Fans                        OK
  AC Input                    OK
  DC Output                   OK
```

```
Hours Used                  24378
Firmware Version            9004.9213
```

Release Information

Command introduced in Junos OS Release 12.1X48.

RELATED DOCUMENTATION

PTX5000 Packet Transport Router Hardware Guide

show chassis environment pem

IN THIS SECTION

- Syntax | 17761
- Syntax (ACX4000 Router) | 17761
- Syntax (TX Matrix Routers) | 17761
- Syntax (TX Matrix Plus Routers) | 17761
- Syntax (MX Series Router) | 17761
- Syntax (PTX Series Router) | 17761
- Syntax (MX104 Universal Routing Platforms) | 17762
- Syntax (MX10003, MX204, MX10004, MX10008, OCX Series, EX9251, and EX9253 devices) | 17762
- Syntax (QFX Series) | 17762
- Description | 17762
- Options | 17763
- Required Privilege Level | 17764
- Output Fields | 17764
- Sample Output | 17765
- Release Information | 17782

Syntax

```
show chassis environment pem
<slot>
```

Syntax (ACX4000 Router)

```
show chassis environment pem
```

Syntax (TX Matrix Routers)

```
show chassis environment pem
<lcc number | scc>
<slot>
```

Syntax (TX Matrix Plus Routers)

```
show chassis environment pem
<lcc number | sfc number>
<slot>
```

Syntax (MX Series Router)

```
show chassis environment pem
<slot>
<all-members>
<local>
<member member-id>
```

Syntax (PTX Series Router)

```
show chassis environment pem
<slot>
<all-members>
```

```
<local>
<member member-id>
```

Syntax (MX104 Universal Routing Platforms)

```
show chassis environment pem
<slot>
<satellite [fpc-slot slot-id |device-alias  alias-name]
```

Syntax (MX10003, MX204, MX10004, MX10008, OCX Series, EX9251, and EX9253 devices)

```
show chassis environment pem
<slot>
```

Syntax (QFX Series)

```
show chassis environment pem
<slot (interconnect-device name slot ) | (node-device name)>
```

Description

Display Power Entry Module (PEM) environmental status information.

> **NOTE**: The new high-capacity (4100W) enhanced DC PEM on MX960 routers includes a new design that can condition the input voltage. This results in the output voltage differing from the input voltage. The earlier generation of DC PEMs coupled the input power directly to the output, thereby making it safe to assume that the output voltage was equal to the input voltage.

Starting in Junos OS Release 23.1, on QFX10008 and QFX10016 switches, you can see the version of the primary firmware available on the power supply.

Options

none

Display environmental information about both PEMs. For the TX Matrix router, display environmental information about the PEMs, the TX Matrix router, and its attached T640 routers. For the TX Matrix Plus router, display environmental information about the PEMs, the TX Matrix Plus router, and its attached routers.

all-members

(MX Series routers only) (Optional) Display environmental information about the PEMs in all the member routers of the Virtual Chassis configuration.

interconnect-device *name*

(QFabric systems only) (Optional) Display chassis environmental information about the PEMs in the Interconnect device.

lcc *number*

(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(MX Series routers only) (Optional) Display environmental information about the PEM in the local Virtual Chassis member.

member *member-id*

(MX Series routers only) (Optional) Display environmental information about the PEM in the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device *name*

(QFabric systems only) (Optional) Display chassis environmental information about the PEMs in the Node device.

satellite [fpc-slot *slot-id* | device-alias *alias-name*]

(Junos Fusion only)(Optional) Display environmental information about the PEM in the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

scc (TX Matrix routers only) (Optional) Display environmental information about the PEM in the TX Matrix router (or switch-card chassis).

sfc (TX Matrix Plus routers only) (Optional) Display environmental information about the PEM in the TX Matrix Plus router (or switch-fabric chassis).

slot (Optional) Display environmental information about an individual PEM. Replace *slot* with 0 or 1.

Required Privilege Level

view

Output Fields

Table 233 on page 17764 lists the output fields for the `show chassis environment pem` command. Output fields are listed in the approximate order in which they appear.

Table 233: show chassis environment pem Output Fields

Field Name	Field Description
PEM*slot*status	Number of the PEM slot.
State	Status of the PEM. **NOTE**: PEM with no feeds connected at bootup on MX10004 and MX10008 devices, continues to stay in "Present" state until feeds are connected to it. PEM with feeds connected at bootup, continues to move to OFFLINE state. "Present" state is only a transitional state that can either move to OFFLINE (or) ONLINE depending on the PEM output. For example, a PEM with no feeds connected at bootup, moves through the states from "Empty" to "Present" to "Offline" state.
Temperature	Temperature of the air flowing past the PEM.
AC Input	Status of the AC input for the specified component
AC Output	Status of the AC output for the specified component.

Table 233: show chassis environment pem Output Fields *(Continued)*

Field Name	Field Description
DC input	Status of the DC input for the specified component.
DC output	Status of the DC output for the specified component.
Load	(Not available on M40e or M160 routers) Information about the load on supply, in percentage of rated current being used.
Voltage	(M120, M160, M320, T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about voltage supplied to the PEM. (MX104 routers only) Information about voltage supplied by the PEM to the system.
Current	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about the PEM current.
Power	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about the PEM power.
SCG/CB/SIB	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) SONET Clock Generator/Control Board/Switch Interface Board.
FAN	(T640, T1600, and T4000 routers with six-input DC power supply only) Information about the DC output to the fan.

Sample Output

show chassis environment pem (M40e Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                           Online
  Temperature                     OK
```

```
  AC input                        OK
  DC output                       OK
```

show chassis environment pem (M120 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                      Online
  Temperature                OK
  DC Input:                  OK
  DC Output:                 OK
  Load                       Less than 20 percent
  Voltage:
    48.0 V input             52864 mV
    48.0 V fan supply        41655 mV
    3.3 V                     3399 mV
PEM 1 status:
  State                      Online
  Temperature                OK
  DC Input:                  OK
  DC Output:                 OK
  Load                       Less than 20 percent
  Voltage:
    48.0 V input             54537 mV
    48.0 V fan supply        42910 mV
    3.3 V                     3506 mV
```

show chassis environment pem (M160 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                           Online
  Temperature                     OK
  DC input                        OK
  DC output                       OK
  Load                            Less than 20 percent
  Voltage:
    48.0 V input                54833 mV
    48.0 V fan supply           50549 mV
```

```
    8.0 V bias               8239 mV
    5.0 V bias               5006 mV
```

show chassis environment pem (M320 Router)

```
user@host> show chassis environment pem
PEM 2 status:
  State                           Online
  Temperature                     OK
  DC input                        OK
  Load                            Less than 40 percent
    48.0 V input            51853 mV
    48.0 V fan supply       48877 mV
    8.0 V bias               8449 mV
    5.0 V bias               4998 mV
PEM 3 status:
  State                           Online
  Temperature                     OK
  DC input                        OK
  Load                            Less than 40 percent
    48.0 V input            51717 mV
    48.0 V fan supply       49076 mV
    8.0 V bias               8442 mV
    5.0 V bias               4998 mV
```

show chassis environment pem (MX150)

```
user@host> show chassis environment pem
FPC 0 PEM 0 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK
```

show chassis environment pem (MX104 Router)

```
user@host> show chassis environment pem
PEM 0 status:
```

```
  State                       Online
  Temperature                 OK
  DC Output:                  OK
  Voltage:
    12.0 V output             12281 mV
    3.3 V output               3353 mV
PEM 1 status:
  State                       Empty
```

show chassis environment pem (MX240 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                       Online
  Temperature                  OK
  DC Output:                   OK
PEM 1 status:
  State                       Online
  Temperature                  OK
  DC Output:                   OK
```

show chassis environment pem (MX480 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                       Online
  Temperature                  OK
  DC Input:                    OK
  DC Output:                   OK
  Voltage:
PEM 1 status:
  State                       Online
  Temperature                  OK
  DC Input:                    OK
  DC Output:                   OK
  Voltage:
```

show chassis environment pem (MX960 Router)

```
user@host> show chassis environment pem
PEM 2 status:
  State                      Present
PEM 3 status:
  State                      Online
  Temperature                 OK
  DC Output:                  OK
```

show chassis environment pem (MX10003 Router)

```
user@host> show chassis environment pem

PEM 0 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   34 degrees C / 93 degrees F
  Temperature                OK   26 degrees C / 78 degrees F
  Temperature                OK   24 degrees C / 75 degrees F
  Firmware version           0x22
  Cooling Fan                8752 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     26            312      10
PEM 1 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   35 degrees C / 95 degrees F
  Temperature                OK   26 degrees C / 78 degrees F
  Temperature                OK   25 degrees C / 77 degrees F
  Firmware version           0x22
  Cooling Fan                8480 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     27            324      11
PEM 2 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   37 degrees C / 98 degrees F
  Temperature                OK   29 degrees C / 84 degrees F
  Temperature                OK   25 degrees C / 77 degrees F
```

```
  Firmware version           0x22
  Cooling Fan                8656 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     25            300      10
PEM 3 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   35 degrees C / 95 degrees F
  Temperature                OK   26 degrees C / 78 degrees F
  Temperature                OK   25 degrees C / 77 degrees F
  Firmware version           0x22
  Cooling Fan                8448 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     26            312      10
PEM 4 status:
  State                      Empty
PEM 5 status:
  State                      Empty
```

show chassis environment pem 0 (MX10004)

```
PEM 0 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   33 degrees C / 91 degrees F
  Firmware version           AA17AA17.AB15AB15.AC27
  Fan 0                      6848 RPM
  Fan 1                      8032 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     72            864      15
  Input                     Voltage(V) Current(A) Power(W)
  INP 1                      201.4      5.0        1015.1
  INP 2                      0.0        0.0         0.0
  Health check Information:
    Status: Scheduled
    Last Result: Not Performed Yet
    Last Execution: Empty
    Next Scheduled Run: 2022-08-31 17:29:18
```

show chassis environment pem (MX204 Router)

```
user@host> show chassis environment pem

PEM 0 status:
  State                      Empty
PEM 1 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   48 degrees C / 118 degrees F
  Temperature                OK   51 degrees C / 123 degrees F
  Fan Sensor                 5400 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          11.94     16            191      29
```

show chassis environment pem (MX10008 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                             OK   29 degrees C / 84 degrees F
  Firmware version           0x36
  Fan 0                      5880 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                        12.00       104           1248     46
PEM 1 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                             OK   27 degrees C / 80 degrees F
  Firmware version           0x36
  Fan 0                      5940 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                        12.00       104           1248     46
PEM 2 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                             OK   30 degrees C / 86 degrees F
  Firmware version           0x36
  Fan 0                      5940 RPM
```

```
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                        12.00       105           1260     46
PEM 3 status:
  State                      Present
PEM 4 status:
  State                      Present
PEM 5 status:
  State                      Present
```

show chassis environment pem (PTX10016 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   21 degrees C / 69 degrees F
  Firmware version           0x36
  Fan 0                      5760 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     51            612      22
PEM 1 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   23 degrees C / 73 degrees F
  Firmware version           0x36
  Fan 0                      5760 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     52            624      23
PEM 2 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   23 degrees C / 73 degrees F
  Firmware version           0x36
  Fan 0                      5760 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     51            612      22
PEM 3 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   21 degrees C / 69 degrees F
  Firmware version           0x36
```

```
  Fan 0                       5760 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     51            612      22
PEM 4 status:
  State                       Online
  Airflow                     Front to Back
  Temperature                 OK   22 degrees C / 71 degrees F
  Firmware version            0x36
  Fan 0                       5760 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     52            624      23
PEM 5 status:
  State                       Online
  Airflow                     Front to Back
  Temperature                 OK   24 degrees C / 75 degrees F
  Firmware version            0x36
  Fan 0                       5700 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     51            612      22
PEM 6 status:
  State                       Online
  Airflow                     Front to Back
  Temperature                 OK   21 degrees C / 69 degrees F
  Firmware version            0x36
  Fan 0                       5700 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     50            600      22
```

show chassis environment pem (T320 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                           Online
  Temperature                     OK
  DC input:                       OK
```

show chassis environment pem (T640 Router)

```
user@host> show chassis environment pem
PEM 0 status:
```

```
  State                           Online
  Temperature                 22 degrees C / 71 degrees F
  AC input: OK
  DC output:            Voltage   Current       Power    Load
       FPC 0              56875       606          34       4
       FPC 1              57016       525          29       3
       FPC 2                  0         0           0       0
       FPC 3                  0         0           0       0
       FPC 4                  0         0           0       0
       FPC 5                  0         0           0       0
       FPC 6              57158      1581          90      12
       FPC 7                  0         0           0       0
  SCG/CB/SIB              56750      1125          63       5
```

show chassis environment pem (T4000 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                      Online
  Temperature                33 degrees C / 91 degrees F
  DC Input:                  OK
                      Voltage(V)  Current(A)  Power(W)  Load(%)
      INPUT 0           54.625       9.812      535       22
      INPUT 1           54.625      10.250      559       23
      INPUT 2           55.125       0.125        6        0
      INPUT 3           54.500      10.062      548       22
      INPUT 4           54.750       9.375      513       21
      INPUT 5           54.750      10.187      557       23
  DC Output           Voltage(V)  Current(A)  Power(W)  Load(%)
      FPC 0             55.750      10.125      564       37
      FPC 1             51.625       0.000        0        0
      FPC 2             52.000       0.000        0        0
      FPC 3             55.062      10.437      574       38
      FPC 4             52.125       0.000        0        0
      FPC 5             55.000       9.375      515       34
      FPC 6             55.187       9.687      534       35
      FPC 7             51.437       0.000        0        0
      SCG/CB/SIB        55.375      15.750      872       35
      FAN               54.562      14.750      804       42
```

show chassis environment pem (T640/T1600/T4000 Routers With Six-Input DC Power Supply)

```
user@host> show chassis environment pem
PEM 1 status:
  State                      Online
  Temperature                36 degrees C / 96 degrees F
  DC Input:                  OK
                      Voltage(V)  Current(A)  Power(W)  Load(%)
      INPUT 0            0.000       0.000        0        0
      INPUT 1           54.875       3.812      209       27
      INPUT 2           55.375       3.937      218       29
      INPUT 3           54.625       3.750      204       27
      INPUT 4           55.125       3.375      186       24
      INPUT 5           55.125       3.375      186       24
  DC Output           Voltage(V)  Current(A)  Power(W)  Load(%)
      FPC 0             52.312       0.000        0        0
      FPC 1             52.687       0.000        0        0
      FPC 2             52.812       0.000        0        0
      FPC 3             55.812       7.062      394       52
      FPC 4             52.625       0.000        0        0
      FPC 5             52.625       0.000        0        0
      FPC 6             52.750       0.000        0        0
      FPC 7             52.750       0.000        0        0
      SCG/CB/SIB        55.937      11.937      667       55
      FAN               55.812       4.937      275       36
```

show chassis environment pem lcc (TX Matrix Routing Matrix)

```
user@host> show chassis environment pem 0 lcc
0
lcc0-re0:
--------------------------------------------------------------------------
PEM 0 status:
  State                           Present
  Temperature                  27 degrees C / 80 degrees F
  DC input:                       Check
  DC output:            Voltage   Current       Power    Load
        FPC 0                 0         0           0       0
        FPC 1                 0         0           0       0
```

```
        FPC 2                  0         0           0       0
        FPC 3                  0         0           0       0
        FPC 4                  0         0           0       0
        FPC 5                  0         0           0       0
        FPC 6                  0         0           0       0
        FPC 7                  0         0           0       0
  SCG/CB/SIB                   0         0           0       0
```

show chassis environment pem scc (TX Matrix Routing Matrix)

```
user@host> show chassis environment pem scc
scc-re0:
--------------------------------------------------------------------------
PEM 1 status:
  State                            Online
  Temperature                   24 degrees C / 75 degrees F
  DC input:                        OK
  DC output:             Voltage   Current       Power    Load
       SIB 0                   0         0           0       0
       SIB 1                   0         0           0       0
       SIB 2                   0         0           0       0
       SIB 3               56550         0           0       0
       SIB 4               55958      6912         386      51
```

show chassis environment pem sfc (TX Matrix Plus Routing Matrix)

```
user@host> show chassis environment pem sfc
0
sfc0-re0:
--------------------------------------------------------------------------
PEM 0 status:
  State                      Online
  Temperature                35 degrees C / 95 degrees F
  DC Input:                   OK
  DC Output              Voltage   Current       Power    Load
      Channel 0            53820     14140         761      59
      Channel 1            53550     12720         681      53
      Channel 2            53840     12930         696      54
      Channel 3            53690     14990         804      63
      Channel 4            53620     15070         808      63
```

```
      Channel 5           53900     14820         798      62
      Channel 6           54120     5020          271      21
```

show chassis environment pem lcc (TX Matrix Plus Routing Matrix)

```
user@host> show chassis environment lcc 0

lcc0-re1:
--------------------------------------------------------------------------
PEM 0 status:
  State                      Online
  Temperature                38 degrees C / 100 degrees F
  DC Input:                  OK
  DC Output             Voltage   Current       Power    Load
      FPC 0               0         0             0        0
      FPC 1               0         0             0        0
      FPC 2               0         0             0        0
      FPC 3               0         0             0        0
      FPC 4               56408     7575          427      56
      FPC 5               0         0             0        0
      FPC 6               56266     7956          447      59
      FPC 7               56283     6100          343      45
      SCG/CB/SIB          55916     8950          500      41
PEM 1 status:
  State                      Present
  Temperature                35 degrees C / 95 degrees F
  DC Input:                  Check
  DC Output             Voltage   Current       Power    Load
      FPC 0               0         0             0        0
      FPC 1               0         0             0        0
      FPC 2               0         0             0        0
      FPC 3               0         0             0        0
      FPC 4               0         0             0        0
      FPC 5               0         0             0        0
      FPC 6               0         0             0        0
      FPC 7               0         0             0        0
      SCG/CB/SIB          0         0             0        0
```

show chassis environment pem node-device (QFabric System)

```
user@switch> show chassis environment pem node-device
node1
FPC 0 PEM 0 status:
  State                      Check
  Airflow                    Front to Back
  Temperature                OK
  AC Input:                  OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12        10            120      18
FPC 0 PEM 1 status:
  State                      Online
  Airflow                    Back to Front
  Temperature                OK
  AC Input:                  OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          11        10           110      17
```

show chassis environment pem (QFX Series and OCX Series)

```
user@switch> show chassis environment pem
FPC 0 PEM 1 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK
  AC Input:                  OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12        17            204      31
```

show chassis environment pem (QFX 10016)

```
user@router> show chassis environment pem 1

PEM 1 status:
  State                      Present
  Input                      Voltage(V)  Current(A) Power(W)
  INP 1                      229.9       0.4        96.6
  INP 2                      233.7       0.4        98.2
```

```
  Health check Information:
      Status:                    Scheduled
      Last Result:               Pass
      Last Execution:            2019-04-23 15:09:54
      Next Scheduled Run:        2019-04-23 15:32:59
```

show chassis environment pem interconnect-device (QFabric System)

```
user@switch> show chassis environment pem interconnect-device
IC1 1
IC1 PEM 1 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK
  AC Input:                  OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12        18            216      33
```

show chassis environment pem (EX9251 Switches)

```
user@switch> show chassis environment pem
PEM 0 status:
  State                      Present
PEM 1 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   36 degrees C / 96 degrees F
  Temperature                OK   35 degrees C / 95 degrees F
  Fan Sensor                 5940 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                        11.85       17            201      30
```

show chassis environment pem (EX9253 Switches)

```
user@switch> show chassis environment pem
PEM 0 status:
  State                      Online
  Airflow                    Front to Back
```

```
  Temperature                OK   56 degrees C / 132 degrees F
  Temperature                OK   46 degrees C / 114 degrees F
  Temperature                OK   28 degrees C / 82 degrees F
  Firmware version           04.10
  Cooling Fan                9056 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     47            564      19
PEM 1 status:
  State                      Present
PEM 2 status:
  State                      Empty
PEM 3 status:
  State                      Empty
PEM 4 status:
  State                      Present
PEM 5 status:
  State                      Online
  Airflow                    Front to Back
  Temperature                OK   61 degrees C / 141 degrees F
  Temperature                OK   49 degrees C / 120 degrees F
  Temperature                OK   28 degrees C / 82 degrees F
  Firmware version           04.10
  Cooling Fan                8656 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          12.00     51            612      21
```

show chassis environment pem (PTX1000 Packet Transport Routers)

```
user@router> show chassis environment pem

PEM 0 status:
  State                      Online
  Airflow                    Front to Back
  Temp Sensor 0              OK   22 degrees C / 71 degrees F
  Temp Sensor 1              OK   23 degrees C / 73 degrees F
  Fan 0                      9184 RPM
  Fan 1                      7936 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                        12          24            288      18
PEM 2 status:
  State                      Online
```

```
  Airflow                    Front to Back
  Temp Sensor 0              OK   22 degrees C / 71 degrees F
  Temp Sensor 1              OK   26 degrees C / 78 degrees F
  Fan 0                      9056 RPM
  Fan 1                      7808 RPM
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                        12          24            288      18
```

On PTX1000 Packet Transport Routers, you cannot view the `show chassis environment pem` output at the PEM slot level, by using the command `show chassis environment pem` *slot*.

show chassis environment pem (ACX7100, ACX7024, ACX 7509)

```
user@host> show chassis environment pem
PSM 2 status:
  State                      Online
  Temperature                22 degrees C / 71 degrees F
  Fans                       OK
  DC Output                  OK
  Hours Used                 26801
  Firmware Version           0201.0208
  Fan 1                      8064
  Fan 2                      9024
  Health check Information:
      Status:                   Not Performed
      Last Result:              Not Performed
      Last Execution:           Empty
      Next Scheduled Run:       2024-09-26 05:03:49 PDT
PSM 3 status:
  State                      Online
  Temperature                23 degrees C / 73 degrees F
  Fans                       OK
  DC Output                  OK
  Hours Used                 26631
  Firmware Version           0201.0208
  Fan 1                      8064
  Fan 2                      9024
  Health check Information:
      Status:                   Not Performed
      Last Result:              Not Performed
      Last Execution:           Empty
```

```
        Next Scheduled Run:       2024-09-26 05:03:59 PDT
```

Release Information

Command introduced before Junos OS Release 7.4.

`satellite` option introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

show chassis hardware | 18430

show chassis environment power-supply-unit

IN THIS SECTION

- Syntax | 17782
- Description | 17782
- Required Privilege Level | 17783
- Output Fields | 17783
- Sample Output | 17783
- Release Information | 17783

Syntax

```
show chassis environment power-supply-unit
```

Description

(On EX4500 switches only) Display the state of the power supply and the direction of the airflow through the power supply.

Required Privilege Level

view

Output Fields

Table 234 on page 17783 lists the output fields for the `show chassis environment power-supply-unit` command. Output fields are listed in the approximate order in which they appear.

Table 234: show chassis environment power-supply-unit Output Fields

Field Name	Field Description
State	State of the power supply: Online or Offline.
Airflow	Direction of airflow of the power supply. One of the following: Front to back or Back to front.

Sample Output

show chassis environment power-supply-unit

```
user@switch> show chassis environment power-supply-unit

FPC 0 PSU 1 status:
  State                     Online
  Airflow                   Front to back
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

- Monitoring Chassis Information

show chassis environment psm

IN THIS SECTION

- Syntax | 17784
- Description | 17784
- Options | 17784
- Required Privilege Level | 17785
- Output Fields | 17785
- Sample Output | 17786
- Release Information | 17797

Syntax

```
show chassis environment psm
<all-members>
<local>
<member member-id>
<psm-slot-number>
```

Description

Display chassis environmental information about the power supply module (PSM).

To run system health check correctly in Junos evolved versions for PTX series, and to get correct output in the command **show environment psm**, the following prerequisites must be followed:

- System load must be enough to load a PSM to its health check threshold.
- There should be at least one redundant active PSM with no defects.

Options

none	Display environmental information about all power supply modules (PSMs).

all-members (Optional) Display chassis environmental information about the PSM in all members of the Virtual Chassis configuration.

local (Optional) Display chassis environmental information about the PSM in the local member of the Virtual Chassis.

member ***member-id*** (Optional) Display chassis environmental information about the PSM in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.

psm-slot-number (Optional) Display environmental information about the specified power supply module. For MX2020 routers, replace *psm-slot-number* with a value from 0 through 17. For MX2010 and MX2008 routers, replace *psm-slot-number* with a value from 0 through 8.

Required Privilege Level

view

Output Fields

Table 235 on page 17785 lists the output fields for the `show chassis environment psm` command. Output fields are listed in the approximate order in which they appear.

Table 235: show chassis environment psm Output Fields

Field Name	Field Description
`State`	Status of the PSM. • `Online`—The PSM is online and running. • `Offline`—PSM is powered down.
`Temperature`	The status of the temperature of the air flowing past the PSM. • `Out of range`—Displayed if the PSM detects over-temperature. • `OK`—Displayed if the temperature is within the acceptable limit.
`DC Input`	State of the DC input power feed for the specified zone at the specified amps and voltage, and load for the PSM.

Table 235: show chassis environment psm Output Fields *(Continued)*

Field Name	Field Description
DC Output	DC power output in watts (W) for the specified zone at the specified amps and voltage (A @ V), and load and percentage utilization of the maximum capacity for the PSM.
Hours Used	Number of hours the PSM has been operational.
Bootloader Version	Bootloader version for AC PSMs.

Sample Output

show chassis environment psm (MX2020 Router)

```
user@host> show chassis environment psm
PSM 2 status:
  State                      Online
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
                             INP0         50.00       18.90     945.00
                             INP1         0.00        0.00      0.00
  DC Output                  Voltage(V) Current(A)  Power(W)   Load(%)
                               51.75      16.50      853.88     40.66
  Hours Used                 6140
PSM 3 status:
  State                      Online
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
                             INP0         50.40       18.90     952.56
                             INP1         0.00        0.00      0.00
  DC Output                  Voltage(V) Current(A)  Power(W)   Load(%)
                               51.75      16.50      853.88     40.66
  Hours Used                 6140
PSM 4 status:
  State                      Online
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
                             INP0         50.40       18.90     952.56
```

```
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.00      16.75      871.00     41.48
  Hours Used                  6140
PSM 5 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         50.40       18.90     952.56
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.00      16.50      858.00     40.86
  Hours Used                  6140
PSM 6 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         50.40       18.90     952.56
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.00      16.75      871.00     41.48
  Hours Used                  6140
PSM 7 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         50.40       19.20     967.68
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.00      16.75      871.00     41.48
  Hours Used                  6140
PSM 8 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         50.00       20.40     1020.00
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.75      17.00      879.75     41.89
  Hours Used                  3380
PSM 11 status:
  State                       Online
  Temperature                 OK
```

```
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         50.40       18.30     922.32
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.00      16.25      845.00     40.24
  Hours Used                  5615
PSM 12 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         50.40       18.30     922.32
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.00      16.00      832.00     39.62
  Hours Used                  6143
PSM 13 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         50.40       18.00     907.20
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.00      16.00      832.00     39.62
  Hours Used                  6143
PSM 14 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         50.00       18.30     915.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.00      16.00      832.00     39.62
  Hours Used                  6143
PSM 15 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         48.80       18.90     922.32
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.00      16.25      845.00     40.24
  Hours Used                  6143
PSM 16 status:
```

```
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         48.80       18.90     922.32
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.00      16.25      845.00     40.24
  Hours Used                  6143
PSM 17 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         48.80       18.90     922.32
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.00      16.25      845.00     40.24
  Hours Used                  5207
```

show chassis environment psm (MX2020 Router with 240-V high-voltage DC PSMs)

```
user@host> show chassis environment psm

PSM 0 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         240.00      1.00      240.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.75      3.75       197.81     7.91
  Hours Used                  2026
PSM 1 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         240.00      0.90      216.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.75      3.25       171.44     6.86
  Hours Used                  2530
PSM 2 status:
```

```
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         240.00      0.90      216.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.75      3.50       184.62     7.38
  Hours Used                  2530
PSM 3 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         240.00      0.90      216.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.75      3.50       184.62     7.38
  Hours Used                  2026
PSM 4 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         240.00      0.90      216.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.50      3.50       183.75     7.35
  Hours Used                  2530
PSM 5 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         240.00      1.00      240.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.50      3.75       196.88     7.88
  Hours Used                  2530
PSM 6 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         240.00      1.00      240.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.50      3.75       196.88     7.88
```

```
  Hours Used                  2002
PSM 7 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         240.00      1.00      240.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.50      3.75       196.88     7.88
  Hours Used                  2146
PSM 8 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         0.00        0.00      0.00
                              INP1         240.00      1.00      240.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.50      3.75       196.88     7.88
  Hours Used                  2026
PSM 9 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         240.00      0.80      192.00
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.50      2.75       144.38     5.78
  Hours Used                  2530
PSM 10 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         240.00      0.80      192.00
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.75      2.75       145.06     5.80
  Hours Used                  682
PSM 11 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         240.00      0.70      168.00
                              INP1         0.00        0.00      0.00
```

```
  DC Output                  Voltage(V) Current(A)  Power(W)   Load(%)
                               52.75      2.75       145.06     5.80
  Hours Used                 2098
PSM 12 status:
  State                      Online
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
                             INP0         240.00      0.80      192.00
                             INP1         0.00        0.00      0.00
  DC Output                  Voltage(V) Current(A)  Power(W)   Load(%)
                               52.25      3.00       156.75     6.27
  Hours Used                 2458
PSM 13 status:
  State                      Online
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
                             INP0         240.00      0.80      192.00
                             INP1         0.00        0.00      0.00
  DC Output                  Voltage(V) Current(A)  Power(W)   Load(%)
                               52.25      2.75       143.69     5.75
  Hours Used                 2601
PSM 14 status:
  State                      Online
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
                             INP0         240.00      0.80      192.00
                             INP1         0.00        0.00      0.00
  DC Output                  Voltage(V) Current(A)  Power(W)   Load(%)
                               52.75      2.75       145.06     5.80
  Hours Used                 682
PSM 15 status:
  State                      Online
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
                             INP0         240.00      0.80      192.00
                             INP1         0.00        0.00      0.00
  DC Output                  Voltage(V) Current(A)  Power(W)   Load(%)
                               52.25      2.75       143.69     5.75
  Hours Used                 2122
PSM 16 status:
  State                      Online
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
```

```
                              INP0         240.00      0.70      168.00
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.25      2.50       130.62     5.22
  Hours Used                  2050
PSM 17 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         240.00      0.80      192.00
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                52.50      3.00       157.50     6.30
  Hours Used                  2122
```

show chassis environment psm (MX2010 Router)

```
user@host> show chassis environment psm
PSM 0 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         51.20       14.70     752.64
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.25      13.00      666.25     26.65
  Hours Used                  2056
PSM 1 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         51.20       14.35     734.72
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.25      12.75      653.44     26.14
  Hours Used                  2008
PSM 2 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
```

```
                              INP0         51.20       14.35     734.72
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.50      13.00      669.50     26.78
  Hours Used                  2032
PSM 3 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         50.40       14.35     723.24
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.00      12.75      650.25     26.01
  Hours Used                  2008
PSM 4 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         51.20       14.00     716.80
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.25      13.00      666.25     26.65
  Hours Used                  2055
PSM 5 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         51.20       14.70     752.64
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.25      13.00      666.25     26.65
  Hours Used                  2056
PSM 6 status:
  State                       Online
Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         50.80       14.70     746.76
                              INP1         0.00        0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.25      13.00      666.25     26.65
  Hours Used                  2056
PSM 7 status:
  State                       Online
```

```
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
                             INP0         50.40       14.70     740.88
                             INP1         0.00        0.00      0.00
  DC Output                  Voltage(V) Current(A)  Power(W)   Load(%)
                               51.25      13.00      666.25     26.65
  Hours Used                 2056
PSM 8 status:
  State                      Online
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
                             INP0         50.40       14.70     740.88
                             INP1         0.00        0.00      0.00
  DC Output                  Voltage(V) Current(A)  Power(W)   Load(%)
                               51.25      13.00      666.25     26.65
  Hours Used                 2056
```

show chassis environment psm (MX2008 Router)

```
user@host> show chassis environment psm
PSM 1 status:
  State                      Online
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
                             INP0         53.20       4.55      242.06
                             INP1         53.20       3.85      204.82
  DC Output                  Voltage(V) Current(A)  Power(W)   Load(%)
                               51.50      7.75       399.12     15.96
  Hours Used                 2811
PSM 2 status:
  State                      Online
  Temperature                OK
  DC Input                   Feed       Voltage(V)  Current(A) Power(W)
                             INP0         53.20       4.55      242.06
                             INP1         53.20       3.85      204.82
  DC Output                  Voltage(V) Current(A)  Power(W)   Load(%)
                               51.25      8.00       410.00     16.40
  Hours Used                 2882
PSM 3 status:
  State                      Online
  Temperature                OK
```

```
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         52.40       1.75      91.70
                              INP1         52.80       8.40      443.52
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.50      7.75       399.12     15.96
  Hours Used                  2668
PSM 4 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         52.80       0.35      18.48
                              INP1         53.20       8.40      446.88
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.50      7.75       399.12     15.96
  Hours Used                  2740
PSM 5 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         52.40       0.00      0.00
                              INP1         53.20       8.40      446.88
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.50      7.75       399.12     15.96
  Hours Used                  2932
PSM 6 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         54.00       8.40      453.60
                              INP1         53.20       0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.50      7.75       399.12     15.96
  Hours Used                  2932
PSM 7 status:
  State                       Online
  Temperature                 OK
  DC Input                    Feed       Voltage(V)  Current(A) Power(W)
                              INP0         54.00       8.40      453.60
                              INP1         53.20       0.00      0.00
  DC Output                   Voltage(V) Current(A)  Power(W)   Load(%)
                                51.25      7.75       397.19     15.89
  Hours Used                  2931
```

show chassis environment psm (ACX Series)

```
user@host> show chassis environment psm
PSM 0 status:
  State                      Online
  Temperature                35 degrees C / 95 degrees F
  Fans                       OK
  DC Output                  OK
  Hours Used                 2786
  Firmware Version           11.13
  Fan 1                      13056
  Health check Information:
      Status:                   Not Performed
      Last Result:              Not Performed
      Last Execution:           Empty
      Next Scheduled Run:       2024-09-27 09:39:12 PDT
PSM 1 status:
  State                      Online
  Temperature                34 degrees C / 93 degrees F
  Fans                       OK
  DC Output                  OK
  Hours Used                 6507
  Firmware Version           11.13
  Fan 1                      13088
  Health check Information:
      Status:                   Not Performed
      Last Result:              Not Performed
      Last Execution:           Empty
      Next Scheduled Run:       2024-09-27 09:39:23 PDT
```

Release Information

Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 Universal Routing Platforms.

`all-members`, `local`, and `member` *`member-id`* options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.

RELATED DOCUMENTATION

show chassis environment | 17429

show chassis environment psu

IN THIS SECTION

- Syntax | 17798
- Description | 17798
- Options | 17798
- Required Privilege Level | 17798
- Output Fields | 17799
- Sample Output | 17799
- Release Information | 17800

Syntax

```
show chassis environment psu
<slot-number>
```

Description

(On EX8200 switches only) Display the state of the power supply.

Options

none Display the state of the power supply for all power supplies.

slot-number (Optional) Display the state of the power supply for a specific power supply slot number (0–5).

Required Privilege Level

view

Output Fields

Table 236 on page 17799 lists the output fields for the `show chassis environment psu` command. Output fields are listed in the approximate order in which they appear.

Table 236: show chassis environment psu Output Fields

Field Name	Field Description
State	State of the power supply: Online, Offline, or Empty.
Temperature	Temperature for the online power supply: OK or Out of Range.
DC Output	DC output for the online power supply: OK or Out of Range.

Sample Output

show chassis environment psu

```
user@switch> show chassis environment psu

PSU 0 status:
  State                      Offline
PSU 1 status:
  State                      Online
  Temperature                 OK
  DC Output:                  OK
PSU 2 status:
  State                      Online
  Temperature                 OK
  DC Output:                  OK
PSU 3 status:
  State                      Offline
PSU 4 status:
  State                      Offline
PSU 5 status:
  State                      Offline
```

show chassis environment psu (for PSU 1)

```
user@switch> show chassis environment psu 1
PSU 1 status:
  State                      Online
  Temperature                 OK
  DC Output:                  OK
```

Release Information

Command introduced in Junos OS Release 10.3.

RELATED DOCUMENTATION

Verifying Power Configuration and Use

show chassis power-budget-statistics | **18681**

show chassis environment routing-engine

IN THIS SECTION

- Syntax | **17801**
- Syntax (TX Matrix Routers) | **17801**
- Syntax (TX Matrix Plus Routers) | **17801**
- Syntax (MX104, MX2010, MX2020, MX10003, MX204, and MX2008 Universal Routing Platforms) | **17801**
- Syntax (MX Series and PTX Series Devices) | **17802**
- Syntax (QFX Series and OCX Series) | **17802**
- Syntax (EX9251 and EX9253 Switches; ACX Series Routers) | **17802**
- Description | **17802**
- Options | **17802**
- Required Privilege Level | **17803**
- Output Fields | **17804**

- Sample Output | 17804
- Sample Output | 17809
- Release Information | 17810

Syntax

```
show chassis environment routing-engine
<slot>
```

Syntax (TX Matrix Routers)

```
show chassis environment routing-engine
<lcc number | scc>
<slot>
```

Syntax (TX Matrix Plus Routers)

```
show chassis environment routing-engine
<lcc number | sfc number>
<slot>
```

Syntax (MX104, MX2010, MX2020, MX10003, MX204, and MX2008 Universal Routing Platforms)

```
show chassis environment routing-engine
<slot>
<satellite [fpc-slot slot-id |device-alias  alias-name]
```

Syntax (MX Series and PTX Series Devices)

```
show chassis environment routing-engine
<slot>
<all-members>
<local>
<member member-id>
```

Syntax (QFX Series and OCX Series)

```
show chassis environment routing-engine
interconnect-device name
```

Syntax (EX9251 and EX9253 Switches; ACX Series Routers)

```
show chassis environment routing-engine
```

Description

Display Routing Engine environmental status information.

Options

none	Display environmental information about all Routing Engines. For a TX Matrix router, display environmental information about all Routing Engines on the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display environmental information about all Routing Engines on the TX Matrix Plus router and its attached routers.
all-members	(MX Series routers only) (Optional) Display environmental information about the Routing Engines in all member routers in the Virtual Chassis configuration.
interconnect-device *name*	(QFabric systems only) (Optional) Display environmental information about the Routing Engines for the Interconnect device.
lcc *number*	—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Display environmental information about the Routing Engines in the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display environmental information about the Routing Engines in the specified member in the Virtual Chassis configuration. Replace *member-id* with the value of 0 or 1.

satellite [fpc-slot *slot-id* **| device-alias** *alias-name***]** (Junos Fusion only)(Optional) Display environmental information for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

scc (TX Matrix router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix router (switch-card chassis).

sfc (TX Matrix Plus router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix Plus router (or switch-fabric chassis).

slot (Optional) Display environmental information about an individual Routing Engine. On M10i, M20, M40e, M120, M160, M320, MX Series, MX104 routers, MX2010 routers, MX2020 routers, MX2008 routers, and T Series routers, replace *slot* with 0 or 1. On M5, M7i, M10, and M40 routers, replace *slot* with 0. On EX3200 and EX4200 standalone switches, replace *slot* with 0. On EX4200 switches in a Virtual Chassis configuration and on EX8208 and EX8216 switches, replace *slot* with 0 or 1. On the QFX3500 and QFX5700, there is only one Routing Engine, so you do not need to specify the slot number. On PTX Series Packet Transport Routers, replace *slot* with 0 or 1.

Required Privilege Level

view

Output Fields

Table 237 on page 17804 lists the output fields for the `show chassis environment routing-engine` command. Output fields are listed in the approximate order in which they appear.

Table 237: show chassis environment routing-engine Output Fields

Field Name	Field Description
Routing engine *slot* status	Number of the Routing Engine slot: 0 or 1.
State	Status of the Routing Engine: • Online Primary—Routing Engine is online, operating as Primary. • Online Standby—Routing Engine is online, operating as Standby. • Offline—Routing Engine is offline.
Temperature	Temperature of the air flowing past the Routing Engines
CPU Temperature	(PTX Series and T4000 Core Routers only) Temperature of the air flowing past the Routing Engine CPU.

Sample Output

show chassis environment routing-engine (Nonredundant)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                           Online Master
  Temperature                  27 degrees C / 80 degrees
```

show chassis environment routing-engine (Redundant)

```
user@host> show chassis environment routing-engine
Route Engine 0 status:
  State:                      Online Master
```

```
  Temperature:              26 degrees C / 78 degrees F
Route Engine 1 status:
  State:                      Online Standby
  Temperature:              26 degrees C / 78 degrees F
```

show chassis enviroenment routing-engine (MX150)

```
user@ host >show chassis environment routing-engine
Routing Engine 0 status:
  State                       Online Master
  CPU Temperature             42 degrees C / 107 degrees F
```

show chassis environment routing-engine (MX104 Router)

```
user@ host >show chassis environment routing-engine
Routing Engine 0 status:
  State                       Online Master
  Temperature                 34 degrees C / 93 degrees F
  CPU Temperature             43 degrees C / 109 degrees F
Routing Engine 1 status:
  State                       Online Standby
  Temperature                 33 degrees C / 91 degrees F
  CPU Temperature             39 degrees C / 102 degrees F
```

show chassis environment routing-engine (MX2010 Router)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                       Online Master
  Temperature                 37 degrees C / 98 degrees F
  CPU Temperature             37 degrees C / 98 degrees F
Routing Engine 1 status:
  State                       Online Standby
  Temperature                 35 degrees C / 95 degrees F
  CPU Temperature             34 degrees C / 93 degrees F
```

show chassis environment routing-engine (MX2020 Router)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                      Online Master
  Temperature                35 degrees C / 95 degrees F
  CPU Temperature            34 degrees C / 93 degrees F
Routing Engine 1 status:
  State                      Online Standby
  Temperature                44 degrees C / 111 degrees F
  CPU Temperature            43 degrees C / 109 degrees F
```

show chassis environment routing-engine (MX2008 Router)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                      Online Master
  CPU Temperature            75 degrees C / 167 degrees F
Routing Engine 1 status:
  State                      Online Standby
  CPU Temperature            47 degrees C / 116 degrees F
```

show chassis environment routing-engine (TX Matrix Plus Router)

```
user@host> show chassis environment routing-engine
sfc0-re0:
--------------------------------------------------------------------------
Routing Engine 0 status:
  State                      Online Master
  Temperature                26 degrees C / 78 degrees F
Routing Engine 1 status:
  State                      Online Standby
  Temperature                28 degrees C / 82 degrees F

lcc0-re0:
--------------------------------------------------------------------------
Routing Engine 0 status:
  State                      Online Master
  Temperature                30 degrees C / 86 degrees F
```

```
Routing Engine 1 status:
  State                      Online Standby
  Temperature                29 degrees C / 84 degrees F
```

show chassis environment routing-engine (T4000 Core Router)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                      Online Master
  Temperature                33 degrees C / 91 degrees F
  CPU Temperature            50 degrees C / 122 degrees F
Routing Engine 1 status:
  State                      Online Standby
  Temperature                33 degrees C / 91 degrees F
  CPU Temperature            46 degrees C / 114 degrees F
```

show chassis environment routing-engine (QFX Series and OCX Series)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State                      Online Master
  CPU Temperature                42 degrees C / 107 degrees F
Routing Engine 1 status:
  State                      Online Standby
  CPU Temperature              34 degrees C / 93 degrees F
```

show chassis environment routing-engine interconnect-device (QFabric System)

```
user@switch> show chassis environment routing-engine
interconnect-device interconnect1
routing-engine interconnect-device interconnect1
Routing Engine 0 status:
  State                      Online Standby
  Temperature                52 degrees C / 125 degrees F
Routing Engine 1 status:
  State                      Online Master
  Temperature                57 degrees C / 134 degrees F
```

show chassis environment routing-engine (PTX5000 Packet Transport Router)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State                      Online Master
  Temperature                55 degrees C / 131 degrees F
  CPU Temperature            66 degrees C / 150 degrees F
Routing Engine 1 status:
  State                      Online Standby
  Temperature                52 degrees C / 125 degrees F
  CPU Temperature            64 degrees C / 147 degrees F
```

show chassis environment routing-engine (PTX10008 Router)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State                      Online Master
  CPU Temperature            40 degrees C / 104 degrees F
Routing Engine 1 status:
  State                      Online Standby
  CPU Temperature            40 degrees C / 104 degrees F
```

show chassis environment routing-engine (PTX10016 Router)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State                      Online Master
  CPU Temperature            33 degrees C / 91 degrees F
Routing Engine 1 status:
  State                      Online Standby
  CPU Temperature            38 degrees C / 100 degrees F
```

show chassis environment routing-engine (ACX Routers)

```
user@host> show chassis environment routing-engine

Routing Engine 0 status:
```

```
  State                       Online Master
  Temp Exhaust Left                      41 degrees C / 105 degrees F
  Temp Inlet Front Right                 35 degrees C / 95 degrees F
  Temp Inlet Front Left                  33 degrees C / 91 degrees F
  Temp Exhaust Right                     48 degrees C / 118 degrees F
  CPU Temperature                        58 degrees C / 136 degrees F
  Hotswap Temp                           42 degrees C / 107 degrees F
  Power 1
    CPU0_VCC2P5_VPP01           2502 mV
    CPU0_VCC3P3_AUX             3320 mV
    CPU0_VTT0P6_0                622 mV
    SSD_3V3                     3304 mV
    VCC5P0                      5034 mV
  Power 2
    VCC1P8                      1800 mV
    VCCIN                       1780 mV
  Power 3
    VCC1P0                      1000 mV
    VCCNAC                       835 mV
  Power 4
    VDDQ                        1230 mV
  Power 5
    VDD1P05                     1060 mV
    VNNPCH                       905 mV
```

Sample Output

show chassis environment routing-engine (PTX5000 (RE-PTX-X8-64G), MX240 (RE-S-X6-64G), MX480 (RE-S-X6-64G), MX960 (RE-S-X6-64G), MX2010 (RE-MX2K-X8-64G), MX2020 (RE-MX2K-X8-64G)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State                       Online Master
  Temperature                 37 degrees C / 98 degrees F
  CPU Temperature             52 degrees C / 125 degrees F
Routing Engine 1 status:
  State                       Online Standby
  Temperature                 37 degrees C / 98 degrees F
```

```
  CPU Temperature            51 degrees C / 123 degrees F
```

show chassis environment routing-engine (MX204 Routers)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                      Online Master
```

show chassis environment routing-engine (MX10008 Routers)

```
Routing Engine 0 status:
  State                      Online Master
  CPU Temperature                         41 degrees C / 105 degrees F
Routing Engine 1 status:
  State                      Online Standby
  CPU Temperature                         40 degrees C / 104 degrees F
```

show chassis environment routing-engine (EX9251 Switches)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State                      Online Master
```

show chassis environment routing-engine (EX9253 Switches)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State                      Online Master
Routing Engine 1 status:
  State                      Present
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced in Junos OS Release 9.6 for the TX Matrix Plus router.

RELATED DOCUMENTATION

request chassis routing-engine master

show chassis routing-engine | 18703

show chassis environment scg

IN THIS SECTION

- Syntax | 17811
- Syntax (TX Matrix and TX Matrix Plus Router) | 17811
- Description | 17812
- Options | 17812
- Required Privilege Level | 17812
- Output Fields | 17812
- Sample Output | 17813
- Release Information | 17817

Syntax

```
show chassis environment scg
<slot>
```

Syntax (TX Matrix and TX Matrix Plus Router)

```
show chassis environment scg
<lcc number>
<slot>
```

Description

Display SONET Clock Generator (SCG) environmental information.

Options

none (TX Matrix and TX Matrix Plus routers only) Display environmental information about all SCGs. On a TX Matrix router, display environmental information about all SCGs on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about all SCGs on the TX Matrix Plus router and its attached routers.

lcc *number* (TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

slot (Optional) Display environmental information about the SCG. Replace ***slot*** with **0** or **1**.

Required Privilege Level

view

Output Fields

Table 238 on page 17813 lists the output fields for the `show chassis environment scg` command. Output fields are listed in the approximate order in which they appear.

Table 238: show chassis environment scg Output Fields

Field Name	Field Description
SCG *slot* status	Number of the SCG slot: **0** or **1**.
State	Status of the SCG: • **Online**—SCG is online and running. • **Offline**—SCG is powered down. If two SCGs are installed and online, one is functioning as the primary, and the other is the standby.
Temperature	Temperature of the air flowing past the SCG.
Power	Power on the SCG. The left column displays required power, in volts. The right column displays measured power, in millivolts.
BUS Revision	Revision level of the generic bus device.

Sample Output

show chassis environment scg (T Series Routers)

```
user@host> show chassis environment scg
SCG 0 status:
  State                          Online -  Master clock
  Temperature                  29 degrees C / 84 degrees F
  Power:
    GROUND                      0 mV
    3.3 V                    3297 mV
    5.0 V                    5050 mV
    5.6 V                    5682 mV
    1.8 V bias               1787 mV
    3.3 V bias               3277 mV
    5.0 V bias               4984 mV
    8.0 V bias               8400 mV
```

```
  BUS Revision                 40
SCG 1 status:
  State                           Online - Standby
  Temperature                  28 degrees C / 82 degrees F
  Power:
    GROUND                      0 mV
    3.3 V                    3317 mV
    5.0 V                    5057 mV
    5.6 V                    5689 mV
    1.8 V bias               1794 mV
    3.3 V bias               3296 mV
    5.0 V bias               4991 mV
    8.0 V bias               8410 mV
    BUS Revision               40
```

show chassis environment scg (T4000 Core Routers)

```
user@host> show chassis environment scg
SCG 0 status:
  State                      Online -  Master clock
  Temperature                33 degrees C / 91 degrees F
  Power
    GROUND                       0 mV
    1.8 V bias                1794 mV
    3.3 V                     3310 mV
    3.3 V bias                3299 mV
    5.0 V                     5040 mV
    5.0 V bias                5003 mV
    5.6 V                     5780 mV
    8.0 V bias                7416 mV
  Bus Revision               40
SCG 1 status:
  State                      Online - Standby
  Temperature                33 degrees C / 91 degrees F
  Power
    GROUND                       0 mV
    1.8 V bias                1794 mV
    3.3 V                     3319 mV
    3.3 V bias                3286 mV
    5.0 V                     5047 mV
    5.0 V bias                5013 mV
```

```
    5.6 V                     5758 mV
    8.0 V bias                7347 mV
  Bus Revision               40
```

show chassis environment scg lcc (TX Matrix Router)

```
user@host> show chassis environment scg lcc
0 0
lcc0-re0:
--------------------------------------------------------------------------
SCG 0 status:
  State                           Online -  Master clock
  Temperature                  30 degrees C / 86 degrees F
  Power:
    GROUND                      0 mV
    3.3 V                    3321 mV
    5.0 V                    5062 mV
    5.6 V                    5682 mV
    1.8 V bias               1789 mV
    3.3 V bias               3289 mV
    5.0 V bias               4993 mV
    8.0 V bias               7807 mV
  BUS Revision                 40
```

show chassis environment scg lcc (TX Matrix Plus Router)

```
user@host> show chassis environment scg lcc
0
lcc0-re0:
--------------------------------------------------------------------------
SCG 0 status:
  State                      Online -  Master clock
  Temperature                42 degrees C / 107 degrees F
  Power
    GROUND                       0 mV
    1.8 V bias                1800 mV
    3.3 V                     3290 mV
    3.3 V bias                3304 mV
    5.0 V                     5042 mV
    5.0 V bias                4979 mV
```

```
    5.6 V                     5765 mV
    8.0 V bias                7682 mV
  Bus Revision               40
```

show chassis environment scg (TX Matrix Plus Router)

```
user@host> show chassis environment scg
lcc0-re0:
--------------------------------------------------------------------------
SCG 0 status:
  State                      Online -  Master clock
  Temperature                40 degrees C / 104 degrees F
  Power
    GROUND                       0 mV
    1.8 V bias                1800 mV
    3.3 V                     3291 mV
    3.3 V bias                3304 mV
    5.0 V                     5042 mV
    5.0 V bias                4979 mV
    5.6 V                     5765 mV
    8.0 V bias                7643 mV
  Bus Revision               40

lcc1-re0:
--------------------------------------------------------------------------
SCG 0 status:
  State                      Online -  Master clock
  Temperature                37 degrees C / 98 degrees F
  Power
    GROUND                       0 mV
    1.8 V bias                1788 mV
    3.3 V                     3305 mV
    3.3 V bias                3284 mV
    5.0 V                     5042 mV
    5.0 V bias                5010 mV
    5.6 V                     5748 mV
    8.0 V bias                7692 mV
  Bus Revision               40

lcc2-re0:
--------------------------------------------------------------------------
```

```
SCG 0 status:
  State                       Online -  Master clock
  Temperature                 39 degrees C / 102 degrees F
  Power
    GROUND                        0 mV
    1.8 V bias                 1785 mV
    3.3 V                      3306 mV
    3.3 V bias                 3301 mV
    5.0 V                      5045 mV
    5.0 V bias                 4993 mV
    5.6 V                      5765 mV
    8.0 V bias                 7838 mV
  Bus Revision                40

lcc3-re0:
--------------------------------------------------------------------------
SCG 0 status:
  State                       Online -  Master clock
  Temperature                 39 degrees C / 102 degrees F
  Power
    GROUND                        0 mV
    1.8 V bias                 1800 mV
    3.3 V                      3290 mV
    3.3 V bias                 3294 mV
    5.0 V                      5050 mV
    5.0 V bias                 4984 mV
    5.6 V                      5780 mV
    8.0 V bias                 7716 mV
  Bus Revision                40
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- request chassis scg | **14553**
- *Configuring the Clock Source*
- T320 SONET Clock Generator (SCG) Description

show chassis environment (Security)

IN THIS SECTION

- Syntax | 17818
- Description | 17818
- Options | 17818
- Required Privilege Level | 17819
- Output Fields | 17819
- Sample Output | 17819
- Release Information | 17825

Syntax

```
show chassis environment
```

Description

Display environmental information about the services gateway chassis, including the temperature and information about the fans, power supplies, and Routing Engine.

Options

none	Display environmental information about the device.
cb *slot-number*	Display chassis environmental information for the Control Board.
fpc *fpc-slot*	Display chassis environmental information for a specified Flexible PIC Concentrator.
fpm	Display chassis environmental information for the craft interface (FPM).
node	Display node specific chassis information.

pem ***slot-number***	Display chassis environmental information for the specified Power Entry Module.
routing-engine ***slot-number***	Display chassis environmental information for the specified Routing Engine.

Required Privilege Level

view

Output Fields

Table 239 on page 17819 lists the output fields for the `show chassis environment` command. Output fields are listed in the approximate order in which they appear.

Table 239: show chassis environment Output Fields

Field Name	Field Description
`Temp`	Temperature of air flowing through the chassis in degrees Celsius (C) and Fahrenheit (F).
`Fan`	Fan status: `OK`, `Testing` (during initial power-on), `Failed`, or `Absent`.

Sample Output

show chassis environment

```
user@host> show chassis environment
user@host> show chassis environment
Class Item                          Status     Measurement
Temp  PEM 0                         OK         40 degrees C / 104 degrees F
      PEM 1                         OK         40 degrees C / 104 degrees F
      PEM 2                         OK         40 degrees C / 104 degrees F
      PEM 3                         OK         45 degrees C / 113 degrees F
      Routing Engine 0              OK         31 degrees C / 87 degrees F
      Routing Engine 0 CPU          OK         27 degrees C / 80 degrees F
      Routing Engine 1              Absent
      Routing Engine 1 CPU          Absent
      CB 0 Intake                   OK         28 degrees C / 82 degrees F
```

```
    CB 0 Exhaust A                  OK          27 degrees C / 80 degrees F
    CB 0 Exhaust B                  OK          29 degrees C / 84 degrees F
    CB 0 ACBC                       OK          29 degrees C / 84 degrees F
    CB 0 SF A                       OK          36 degrees C / 96 degrees F
    CB 0 SF B                       OK          31 degrees C / 87 degrees F
    CB 1 Intake                     OK          27 degrees C / 80 degrees F
    CB 1 Exhaust A                  OK          26 degrees C / 78 degrees F
    CB 1 Exhaust B                  OK          29 degrees C / 84 degrees F
    CB 1 ACBC                       OK          27 degrees C / 80 degrees F
    CB 1 SF A                       OK          36 degrees C / 96 degrees F
    CB 1 SF B                       OK          31 degrees C / 87 degrees F
    CB 2 Intake                     Absent
    CB 2 Exhaust A                  Absent
    CB 2 Exhaust B                  Absent
    CB 2 ACBC                       Absent
    CB 2 XF A                       Absent
    CB 2 XF B                       Absent
    FPC 0 Intake                    OK          47 degrees C / 116 degrees F
    FPC 0 Exhaust A                 OK          44 degrees C / 111 degrees F
    FPC 0 Exhaust B                 OK          52 degrees C / 125 degrees F
    FPC 0 xlp0 TSen                 OK          51 degrees C / 123 degrees F
    FPC 0 xlp0 Chip                 OK          46 degrees C / 114 degrees F
    FPC 0 xlp1 TSen                 OK          51 degrees C / 123 degrees F
    FPC 0 xlp1 Chip                 OK          47 degrees C / 116 degrees F
    FPC 0 xlp2 TSen                 OK          44 degrees C / 111 degrees F
    FPC 0 xlp2 Chip                 OK          42 degrees C / 107 degrees F
    FPC 0 xlp3 TSen                 OK          48 degrees C / 118 degrees F
    FPC 0 xlp3 Chip                 OK          43 degrees C / 109 degrees F
    FPC 1 Intake                    OK          41 degrees C / 105 degrees F
    FPC 1 Exhaust A                 OK          41 degrees C / 105 degrees F
    FPC 1 Exhaust B                 OK          51 degrees C / 123 degrees F
    FPC 1 LU TSen                   OK          46 degrees C / 114 degrees F
    FPC 1 LU Chip                   OK          45 degrees C / 113 degrees F
    FPC 1 XM TSen                   OK          46 degrees C / 114 degrees F
    FPC 1 XM Chip                   OK          52 degrees C / 125 degrees F
    FPC 1 xlp0 TSen                 OK          49 degrees C / 120 degrees F
    FPC 1 xlp0 Chip                 OK          42 degrees C / 107 degrees F
    FPC 1 xlp1 TSen                 OK          49 degrees C / 120 degrees F
    FPC 1 xlp1 Chip                 OK          44 degrees C / 111 degrees F
    FPC 1 xlp2 TSen                 OK          38 degrees C / 100 degrees F
    FPC 1 xlp2 Chip                 OK          39 degrees C / 102 degrees F
    FPC 1 xlp3 TSen                 OK          44 degrees C / 111 degrees F
    FPC 1 xlp3 Chip                 OK          42 degrees C / 107 degrees F
```

```
    FPC 2 Intake                OK          29 degrees C / 84 degrees F
    FPC 2 Exhaust A             OK          34 degrees C / 93 degrees F
    FPC 2 Exhaust B             OK          40 degrees C / 104 degrees F
    FPC 2 I3 0 TSensor          OK          42 degrees C / 107 degrees F
    FPC 2 I3 0 Chip             OK          41 degrees C / 105 degrees F
    FPC 2 I3 1 TSensor          OK          40 degrees C / 104 degrees F
    FPC 2 I3 1 Chip             OK          39 degrees C / 102 degrees F
    FPC 2 I3 2 TSensor          OK          38 degrees C / 100 degrees F
    FPC 2 I3 2 Chip             OK          37 degrees C / 98 degrees F
    FPC 2 I3 3 TSensor          OK          35 degrees C / 95 degrees F
    FPC 2 I3 3 Chip             OK          35 degrees C / 95 degrees F
    FPC 2 IA 0 TSensor          OK          45 degrees C / 113 degrees F
    FPC 2 IA 0 Chip             OK          42 degrees C / 107 degrees F
    FPC 2 IA 1 TSensor          OK          41 degrees C / 105 degrees F
    FPC 2 IA 1 Chip             OK          43 degrees C / 109 degrees F
    FPC 9 Intake                OK          29 degrees C / 84 degrees F
    FPC 9 Exhaust A             OK          41 degrees C / 105 degrees F
    FPC 9 Exhaust B             OK          48 degrees C / 118 degrees F
    FPC 9 LU TSen               OK          48 degrees C / 118 degrees F
    FPC 9 LU Chip               OK          47 degrees C / 116 degrees F
    FPC 9 XM TSen               OK          48 degrees C / 118 degrees F
    FPC 9 XM Chip               OK          54 degrees C / 129 degrees F
    FPC 9 xlp0 TSen             OK          45 degrees C / 113 degrees F
    FPC 9 xlp0 Chip             OK          42 degrees C / 107 degrees F
    FPC 9 xlp1 TSen             OK          49 degrees C / 120 degrees F
    FPC 9 xlp1 Chip             OK          46 degrees C / 114 degrees F
    FPC 9 xlp2 TSen             OK          37 degrees C / 98 degrees F
    FPC 9 xlp2 Chip             OK          40 degrees C / 104 degrees F
    FPC 9 xlp3 TSen             OK          45 degrees C / 113 degrees F
    FPC 9 xlp3 Chip             OK          41 degrees C / 105 degrees F
    FPC 10 Intake               OK          32 degrees C / 89 degrees F
    FPC 10 Exhaust A            OK          44 degrees C / 111 degrees F
    FPC 10 Exhaust B            OK          53 degrees C / 127 degrees F
    FPC 10 LU 0 TSen            OK          43 degrees C / 109 degrees F
    FPC 10 LU 0 Chip            OK          52 degrees C / 125 degrees F
    FPC 10 LU 1 TSen            OK          43 degrees C / 109 degrees F
    FPC 10 LU 1 Chip            OK          44 degrees C / 111 degrees F
    FPC 10 LU 2 TSen            OK          43 degrees C / 109 degrees F
    FPC 10 LU 2 Chip            OK          50 degrees C / 122 degrees F
    FPC 10 LU 3 TSen            OK          43 degrees C / 109 degrees F
    FPC 10 LU 3 Chip            OK          58 degrees C / 136 degrees F
    FPC 10 XM 0 TSen            OK          43 degrees C / 109 degrees F
    FPC 10 XM 0 Chip            OK          53 degrees C / 127 degrees F
```

```
       FPC 10 XF 0 TSen              OK         43 degrees C / 109 degrees F
       FPC 10 XF 0 Chip              OK         64 degrees C / 147 degrees F
       FPC 10 PLX Switch TSen        OK         43 degrees C / 109 degrees F
       FPC 10 PLX Switch Chip        OK         44 degrees C / 111 degrees F
       FPC 11 Intake                 OK         32 degrees C / 89 degrees F
       FPC 11 Exhaust A              OK         41 degrees C / 105 degrees F
       FPC 11 Exhaust B              OK         56 degrees C / 132 degrees F
       FPC 11 LU 0 TSen              OK         45 degrees C / 113 degrees F
       FPC 11 LU 0 Chip              OK         50 degrees C / 122 degrees F
       FPC 11 LU 1 TSen              OK         45 degrees C / 113 degrees F
       FPC 11 LU 1 Chip              OK         47 degrees C / 116 degrees F
       FPC 11 LU 2 TSen              OK         45 degrees C / 113 degrees F
       FPC 11 LU 2 Chip              OK         52 degrees C / 125 degrees F
       FPC 11 LU 3 TSen              OK         45 degrees C / 113 degrees F
       FPC 11 LU 3 Chip              OK         60 degrees C / 140 degrees F
       FPC 11 XM 0 TSen              OK         45 degrees C / 113 degrees F
       FPC 11 XM 0 Chip              OK         56 degrees C / 132 degrees F
       FPC 11 XF 0 TSen              OK         45 degrees C / 113 degrees F
       FPC 11 XF 0 Chip              OK         65 degrees C / 149 degrees F
       FPC 11 PLX Switch TSen        OK         45 degrees C / 113 degrees F
       FPC 11 PLX Switch Chip        OK         46 degrees C / 114 degrees F
Fans   Top Fan Tray Temp             OK         34 degrees C / 93 degrees F
       Top Tray Fan 1                OK         Spinning at normal speed
       Top Tray Fan 2                OK         Spinning at normal speed
       Top Tray Fan 3                OK         Spinning at normal speed
       Top Tray Fan 4                OK         Spinning at normal speed
       Top Tray Fan 5                OK         Spinning at normal speed
       Top Tray Fan 6                OK         Spinning at normal speed
       Top Tray Fan 7                OK         Spinning at normal speed
       Top Tray Fan 8                OK         Spinning at normal speed
       Top Tray Fan 9                OK         Spinning at normal speed
       Top Tray Fan 10               OK         Spinning at normal speed
       Top Tray Fan 11               OK         Spinning at normal speed
       Top Tray Fan 12               OK         Spinning at normal speed
       Bottom Fan Tray Temp          OK         31 degrees C / 87 degrees F
       Bottom Tray Fan 1             OK         Spinning at normal speed
       Bottom Tray Fan 2             OK         Spinning at normal speed
       Bottom Tray Fan 3             OK         Spinning at normal speed
       Bottom Tray Fan 4             OK         Spinning at normal speed
       Bottom Tray Fan 5             OK         Spinning at normal speed
       Bottom Tray Fan 6             OK         Spinning at normal speed
       Bottom Tray Fan 7             OK         Spinning at normal speed
       Bottom Tray Fan 8             OK         Spinning at normal speed
```

```
      Bottom Tray Fan 9             OK         Spinning at normal speed
      Bottom Tray Fan 10            OK         Spinning at normal speed
      Bottom Tray Fan 11            OK         Spinning at normal speed
      Bottom Tray Fan 12            OK         Spinning at normal speed               OK
```

When you enter the **show chassis environment pem** command, the sample output is shown for AC PEM.

```
user@host# run show chassis environment pem node 0 | no-more
node0:
--------------------------------------------------------------------------
PEM 0 status:
  State                     Online
  Airflow                   Front to Back
  Temperature                            OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                        12          9             108      16
PEM 1 status:
  State                     Online
  Airflow                   Front to Back
  Temperature                            OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                        12          9             108      16
```

When you enter the **show chassis environment pem** command, the sample output is shown for DC PEM.

```
user@host> show chassis environment pem
node0:
--------------------------------------------------------------------------
PEM 0 status:
  State                     Online
  Temperature               OK
  DC Input:                 OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          50        12            600      24
PEM 1 status:
  State                     Online
  Temperature               OK
  DC Input:                 OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
```

```
                          50        31            1550     63


node1:
--------------------------------------------------------------------------
PEM 0 status:
  State                      Online
  Temperature                OK
  DC Input:                  OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          50        12            600      24
PEM 1 status:
  State                      Online
  Temperature                OK
  DC Input:                  OK
  DC Output           Voltage(V) Current(A)  Power(W)  Load(%)
                          49        31            1519     62
```

show chassis environment fpc (SRX5800, SRX5400, and SRX5600)

```
user@host> show chassis environment fpc
FPC 1 status:
  State                      Online
  Temperature Intake         34 degrees C / 93 degrees F
  Temperature Exhaust A      48 degrees C / 118 degrees F
  Temperature Exhaust B      48 degrees C / 118 degrees F
  Temperature CPU0 DTS       55 degrees C / 131 degrees F
  Temperature CPU1 DTS       60 degrees C / 140 degrees F
  Temperature CPU2 DTS       54 degrees C / 129 degrees F
  Temperature CPU3 DTS       70 degrees C / 158 degrees F
  Temperature Talus 0        106 degrees C / 222 degrees F
  Temperature Middle 0       40 degrees C / 104 degrees F
  Temperature Talus 1        76 degrees C / 168 degrees F
  Temperature Middle 1       67 degrees C / 152 degrees F
  Power
          TALUS0-1.20V                1199 mV   14187 mA   17010 mW
          TALUS0-0.90V                 900 mV    5000 mA    4500 mW
          BIAS0-3.30V                 3299 mV    3769 mA   12433 mW
          PIC0_CPU_memory_CD-1.20     1199 mV    3781 mA    4533 mW
          USB0-5.00V                  5000 mV     155 mA     775 mW
          PIC0_CPU_memory_AB-1.20     1200 mV    5820 mA    6984 mW
```

```
        PCH0-1.05V                  1050 mV    3582 mA    3761 mW
        TALUS1-1.20V                1199 mV   13640 mA   16354 mW
        TALUS1-0.90V                 899 mV    4679 mA    4206 mW
        BIAS1-3.30V                 3300 mV    3175 mA   10477 mW
        PIC1_CPU_memory_GH-1.20     1200 mV    4648 mA    5577 mW
        USB1-5.00V                  4999 mV     346 mA    1729 mW
        PIC1_CPU_memory_EF-1.20     1200 mV    5218 mA    6261 mW
        PCH1-1.05V                  1050 mV    3328 mA    3494 mW
        TPS53641-CPU0               1750 mV   46062 mA   80608 mW
        TPS53641-CPU1               1739 mV   47437 mA   82492 mW
        TPS53641-CPU2               1750 mV   45250 mA   79187 mW
        TPS53641-CPU3               1739 mV   46875 mA   81515 mW
        ETH-1.00V                    994 mV    2674 mA    2657 mW
        TALUS0_Core-0.85V            849 mV   37750 mA   32049 mW
        TALUS1_Core-0.85V            849 mV   26750 mA   22710 mW
        Power_Brick1-12.00V        12001 mV   21000 mA  252021 mW
        Power_Brick2-12.00V        11998 mV   23125 mA  277453 mW
        PIM4820_48V0-48.00V        58392 mV   10286 mA  600620 mW
        I2C Slave Revision         0
```

Release Information

Command introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

show chassis hardware (View) | 18457

show chassis environment sfb

IN THIS SECTION

- Syntax | 17826
- Description | 17826
- Options | 17826

- Required Privilege Level | 17827
- Output Fields | 17827
- Sample Output | 17827
- Release Information | 17843

Syntax

```
show chassis environment sfb
<all-members>
<local>
<member member-id>
<sfb-slot-number>
```

Description

Display chassis environmental information about the Switch Fabric Board (SFB).

Options

none Display environmental information about all Switch Fabric Boards.

all-members (Optional) Display chassis environmental information about the SFB in all members of the Virtual Chassis configuration.

local (Optional) Display chassis environmental information about the SFB in the local member of the Virtual Chassis.

member *member-id* (Optional) Display chassis environmental information about the SFB in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.

sfb-slot-number (Optional) Display environmental information about the specified Switch Fabric Board. For MX2020, MX2010, MX2008 and MX10008 routers, replace *sfb-slot-number* with a value from 0 through 7. MX10004 routers provide support for SF2 fabric boards with a value from 0 to 5.

Required Privilege Level

view

Output Fields

Table 240 on page 17827 lists the output fields for the `show chassis environment sfb` command. Output fields are listed in the approximate order in which they appear.

Table 240: show chassis environment sfb Output Fields

Field Name	Field Description
`State`	Status of the SFB. • `Online`—The SFB is online and running. • `Offline`— SFB is powered down.
`Temperature`	Temperature in Celsius (C) and Fahrenheit (F) of the air flowing past the SFB. • `Intake`—Measures the temperature of the air intake. • `Exhaust`—Measures the temperature of the hot air exhaust. • `SFB-XF2`—Measures the temperature of the hot air exhaust for the XF2 fabric plane. • `SFB-XF1`—Measures the temperature of the hot air exhaust for the XF1 fabric plane. • `SFB-XF0`—Measures the temperature of the hot air exhaust for the XF0 fabric plane.
`Power`	Power required and measured on the SFB. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.

Sample Output

show chassis environment sfb (MX2020 Router)

```
user@host> show chassis environment sfb
SFB 0 status:
  State                      Online
```

```
  Intake-Zone0 Temperature   51 degrees C / 123 degrees F
  Exhaust-Zone1 Temperature  44 degrees C / 111 degrees F
  IntakeA-Zone0 Temperature  46 degrees C / 114 degrees F
  IntakeB-Zone1 Temperature  37 degrees C / 98 degrees F
  Exhaust-Zone0 Temperature  48 degrees C / 118 degrees F
  SFB-XF2-Zone1 Temperature  58 degrees C / 136 degrees F
  SFB-XF1-Zone0 Temperature  65 degrees C / 149 degrees F
  SFB-XF0-Zone0 Temperature  64 degrees C / 147 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1500 mV
    LTC3880-XF2-1.5v-CH0      1500 mV
    LTC3880-XF2-1.5v-CH1      1500 mV
    LTC3880-XF2-1.0v-RAIL     1029 mV
    LTC3880-XF2-1.0v-CH0      1029 mV
    LTC3880-XF2-1.0v-CH1      1032 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-3.3v-RAIL         3300 mV
    LTC3880-3.3v-CH0          3300 mV
    LTC3880-3.3v-CH1          3299 mV
SFB 1 status:
  State                      Online
  Intake-Zone0 Temperature   52 degrees C / 125 degrees F
  Exhaust-Zone1 Temperature  44 degrees C / 111 degrees F
  IntakeA-Zone0 Temperature  47 degrees C / 116 degrees F
  IntakeB-Zone1 Temperature  37 degrees C / 98 degrees F
  Exhaust-Zone0 Temperature  47 degrees C / 116 degrees F
  SFB-XF2-Zone1 Temperature  59 degrees C / 138 degrees F
  SFB-XF1-Zone0 Temperature  63 degrees C / 145 degrees F
  SFB-XF0-Zone0 Temperature  65 degrees C / 149 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1500 mV
    LTC3880-XF2-1.5v-CH0      1500 mV
```

```
    LTC3880-XF2-1.5v-CH1      1501 mV
    LTC3880-XF2-1.0v-RAIL     1030 mV
    LTC3880-XF2-1.0v-CH0      1030 mV
    LTC3880-XF2-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1500 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-3.3v-RAIL         3300 mV
    LTC3880-3.3v-CH0          3300 mV
    LTC3880-3.3v-CH1          3299 mV
SFB 2 status:
  State                      Online
  Intake-Zone0 Temperature   52 degrees C / 125 degrees F
  Exhaust-Zone1 Temperature  44 degrees C / 111 degrees F
  IntakeA-Zone0 Temperature  47 degrees C / 116 degrees F
  IntakeB-Zone1 Temperature  37 degrees C / 98 degrees F
  Exhaust-Zone0 Temperature  49 degrees C / 120 degrees F
  SFB-XF2-Zone1 Temperature  62 degrees C / 143 degrees F
  SFB-XF1-Zone0 Temperature  66 degrees C / 150 degrees F
  SFB-XF0-Zone0 Temperature  66 degrees C / 150 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1499 mV
    LTC3880-XF2-1.5v-CH0      1499 mV
    LTC3880-XF2-1.5v-CH1      1500 mV
    LTC3880-XF2-1.0v-RAIL     1030 mV
    LTC3880-XF2-1.0v-CH0      1030 mV
    LTC3880-XF2-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
```

```
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-3.3v-RAIL         3300 mV
    LTC3880-3.3v-CH0          3300 mV
    LTC3880-3.3v-CH1          3299 mV
SFB 3 status:
  State                      Online
  Intake-Zone0 Temperature   53 degrees C / 127 degrees F
  Exhaust-Zone1 Temperature  44 degrees C / 111 degrees F
  IntakeA-Zone0 Temperature  48 degrees C / 118 degrees F
  IntakeB-Zone1 Temperature  38 degrees C / 100 degrees F
  Exhaust-Zone0 Temperature  49 degrees C / 120 degrees F
  SFB-XF2-Zone1 Temperature  62 degrees C / 143 degrees F
  SFB-XF1-Zone0 Temperature  65 degrees C / 149 degrees F
  SFB-XF0-Zone0 Temperature  68 degrees C / 154 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1500 mV
    LTC3880-XF2-1.5v-CH0      1500 mV
    LTC3880-XF2-1.5v-CH1      1500 mV
    LTC3880-XF2-1.0v-RAIL     1029 mV
    LTC3880-XF2-1.0v-CH0      1029 mV
    LTC3880-XF2-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1034 mV
    LTC3880-3.3v-RAIL         3300 mV
    LTC3880-3.3v-CH0          3300 mV
    LTC3880-3.3v-CH1          3300 mV
SFB 4 status:
  State                      Online
  Intake-Zone0 Temperature   54 degrees C / 129 degrees F
```

```
  Exhaust-Zone1 Temperature  46 degrees C / 114 degrees F
  IntakeA-Zone0 Temperature  49 degrees C / 120 degrees F
  IntakeB-Zone1 Temperature  39 degrees C / 102 degrees F
  Exhaust-Zone0 Temperature  50 degrees C / 122 degrees F
  SFB-XF2-Zone1 Temperature  61 degrees C / 141 degrees F
  SFB-XF1-Zone0 Temperature  64 degrees C / 147 degrees F
  SFB-XF0-Zone0 Temperature  67 degrees C / 152 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1500 mV
    LTC3880-XF2-1.5v-CH0      1500 mV
    LTC3880-XF2-1.5v-CH1      1500 mV
    LTC3880-XF2-1.0v-RAIL     1030 mV
    LTC3880-XF2-1.0v-CH0      1030 mV
    LTC3880-XF2-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-3.3v-RAIL         3299 mV
    LTC3880-3.3v-CH0          3299 mV
    LTC3880-3.3v-CH1          3299 mV
SFB 5 status:
  State                      Online
  Intake-Zone0 Temperature   54 degrees C / 129 degrees F
  Exhaust-Zone1 Temperature  46 degrees C / 114 degrees F
  IntakeA-Zone0 Temperature  49 degrees C / 120 degrees F
  IntakeB-Zone1 Temperature  40 degrees C / 104 degrees F
  Exhaust-Zone0 Temperature  50 degrees C / 122 degrees F
  SFB-XF2-Zone1 Temperature  63 degrees C / 145 degrees F
  SFB-XF1-Zone0 Temperature  65 degrees C / 149 degrees F
  SFB-XF0-Zone0 Temperature  70 degrees C / 158 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1500 mV
    LTC3880-XF2-1.5v-CH0      1500 mV
    LTC3880-XF2-1.5v-CH1      1500 mV
```

```
    LTC3880-XF2-1.0v-RAIL     1029 mV
    LTC3880-XF2-1.0v-CH0      1029 mV
    LTC3880-XF2-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1500 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-3.3v-RAIL         3299 mV
    LTC3880-3.3v-CH0          3299 mV
    LTC3880-3.3v-CH1          3300 mV
SFB 6 status:
  State                      Online
  Intake-Zone0 Temperature   54 degrees C / 129 degrees F
  Exhaust-Zone1 Temperature  46 degrees C / 114 degrees F
  IntakeA-Zone0 Temperature  48 degrees C / 118 degrees F
  IntakeB-Zone1 Temperature  40 degrees C / 104 degrees F
  Exhaust-Zone0 Temperature  49 degrees C / 120 degrees F
  SFB-XF2-Zone1 Temperature  62 degrees C / 143 degrees F
  SFB-XF1-Zone0 Temperature  64 degrees C / 147 degrees F
  SFB-XF0-Zone0 Temperature  68 degrees C / 154 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1499 mV
    LTC3880-XF2-1.5v-CH0      1499 mV
    LTC3880-XF2-1.5v-CH1      1501 mV
    LTC3880-XF2-1.0v-RAIL     1030 mV
    LTC3880-XF2-1.0v-CH0      1030 mV
    LTC3880-XF2-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1500 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
```

```
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-3.3v-RAIL         3300 mV
    LTC3880-3.3v-CH0          3300 mV
    LTC3880-3.3v-CH1          3300 mV
SFB 7 status:
  State                      Online
  Intake-Zone0 Temperature   53 degrees C / 127 degrees F
  Exhaust-Zone1 Temperature  46 degrees C / 114 degrees F
  IntakeA-Zone0 Temperature  49 degrees C / 120 degrees F
  IntakeB-Zone1 Temperature  40 degrees C / 104 degrees F
  Exhaust-Zone0 Temperature  50 degrees C / 122 degrees F
  SFB-XF2-Zone1 Temperature  64 degrees C / 147 degrees F
  SFB-XF1-Zone0 Temperature  66 degrees C / 150 degrees F
  SFB-XF0-Zone0 Temperature  69 degrees C / 156 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1500 mV
    LTC3880-XF2-1.5v-CH0      1500 mV
    LTC3880-XF2-1.5v-CH1      1501 mV
    LTC3880-XF2-1.0v-RAIL     1029 mV
    LTC3880-XF2-1.0v-CH0      1029 mV
    LTC3880-XF2-1.0v-CH1      1033 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1501 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1033 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1501 mV
    LTC3880-XF0-1.0v-RAIL     1030 mV
    LTC3880-XF0-1.0v-CH0      1030 mV
    LTC3880-XF0-1.0v-CH1      1033 mV
    LTC3880-3.3v-RAIL         3300 mV
    LTC3880-3.3v-CH0          3300 mV
    LTC3880-3.3v-CH1          3300 mV
```

show chassis environment sfb (MX2010 Router)

```
user@host> show chassis environment sfb
SFB 0 status:
  State                      Online
  Intake-Zone0 Temperature   31 degrees C / 87 degrees F
  Exhaust-Zone1 Temperature  22 degrees C / 71 degrees F
  IntakeA-Zone0 Temperature  21 degrees C / 69 degrees F
  IntakeB-Zone1 Temperature  16 degrees C / 60 degrees F
  Exhaust-Zone0 Temperature  23 degrees C / 73 degrees F
  SFB-XF2-Zone1 Temperature  30 degrees C / 86 degrees F
  SFB-XF1-Zone0 Temperature  28 degrees C / 82 degrees F
  SFB-XF0-Zone0 Temperature  38 degrees C / 100 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1500 mV
    LTC3880-XF2-1.5v-CH0      1500 mV
    LTC3880-XF2-1.5v-CH1      1500 mV
    LTC3880-XF2-1.0v-RAIL      949 mV
    LTC3880-XF2-1.0v-CH0       949 mV
    LTC3880-XF2-1.0v-CH1       951 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1500 mV
    LTC3880-XF1-1.0v-RAIL      949 mV
    LTC3880-XF1-1.0v-CH0       949 mV
    LTC3880-XF1-1.0v-CH1       951 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1500 mV
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1032 mV
    LTC3880-3.3v-RAIL         3300 mV
    LTC3880-3.3v-CH0          3300 mV
    LTC3880-3.3v-CH1          3299 mV
SFB 1 status:
  State                      Online
  Intake-Zone0 Temperature   32 degrees C / 89 degrees F
  Exhaust-Zone1 Temperature  20 degrees C / 68 degrees F
  IntakeA-Zone0 Temperature  25 degrees C / 77 degrees F
  IntakeB-Zone1 Temperature  15 degrees C / 59 degrees F
  Exhaust-Zone0 Temperature  24 degrees C / 75 degrees F
```

```
  SFB-XF2-Zone1 Temperature  31 degrees C / 87 degrees F
  SFB-XF1-Zone0 Temperature  31 degrees C / 87 degrees F
  SFB-XF0-Zone0 Temperature  37 degrees C / 98 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1499 mV
    LTC3880-XF2-1.5v-CH0      1499 mV
    LTC3880-XF2-1.5v-CH1      1500 mV
    LTC3880-XF2-1.0v-RAIL     1029 mV
    LTC3880-XF2-1.0v-CH0      1029 mV
    LTC3880-XF2-1.0v-CH1      1031 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1500 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1031 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1500 mV
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1032 mV
    LTC3880-3.3v-RAIL         3299 mV
    LTC3880-3.3v-CH0          3299 mV
    LTC3880-3.3v-CH1          3299 mV
SFB 2 status:
  State                      Online
  Intake-Zone0 Temperature   26 degrees C / 78 degrees F
  Exhaust-Zone1 Temperature  19 degrees C / 66 degrees F
  IntakeA-Zone0 Temperature  23 degrees C / 73 degrees F
  IntakeB-Zone1 Temperature  15 degrees C / 59 degrees F
  Exhaust-Zone0 Temperature  21 degrees C / 69 degrees F
  SFB-XF2-Zone1 Temperature  29 degrees C / 84 degrees F
  SFB-XF1-Zone0 Temperature  26 degrees C / 78 degrees F
  SFB-XF0-Zone0 Temperature  31 degrees C / 87 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1500 mV
    LTC3880-XF2-1.5v-CH0      1500 mV
    LTC3880-XF2-1.5v-CH1      1500 mV
    LTC3880-XF2-1.0v-RAIL     1029 mV
    LTC3880-XF2-1.0v-CH0      1029 mV
    LTC3880-XF2-1.0v-CH1      1031 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
```

```
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1500 mV
    LTC3880-XF1-1.0v-RAIL     1030 mV
    LTC3880-XF1-1.0v-CH0      1030 mV
    LTC3880-XF1-1.0v-CH1      1031 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1500 mV
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1032 mV
    LTC3880-3.3v-RAIL         3300 mV
    LTC3880-3.3v-CH0          3300 mV
    LTC3880-3.3v-CH1          3300 mV
SFB 3 status:
  State                      Offline
  Reason                         No power
SFB 4 status:
  State                      Online
  Intake-Zone0 Temperature   33 degrees C / 91 degrees F
  Exhaust-Zone1 Temperature  21 degrees C / 69 degrees F
  IntakeA-Zone0 Temperature  24 degrees C / 75 degrees F
  IntakeB-Zone1 Temperature  17 degrees C / 62 degrees F
  Exhaust-Zone0 Temperature  24 degrees C / 75 degrees F
  SFB-XF2-Zone1 Temperature  32 degrees C / 89 degrees F
  SFB-XF1-Zone0 Temperature  32 degrees C / 89 degrees F
  SFB-XF0-Zone0 Temperature  37 degrees C / 98 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1499 mV
    LTC3880-XF2-1.5v-CH0      1499 mV
 LTC3880-XF2-1.5v-CH1      1500 mV
    LTC3880-XF2-1.0v-RAIL      949 mV
    LTC3880-XF2-1.0v-CH0       949 mV
    LTC3880-XF2-1.0v-CH1       952 mV
    LTC3880-XF1-1.5v-RAIL     1500 mV
    LTC3880-XF1-1.5v-CH0      1500 mV
    LTC3880-XF1-1.5v-CH1      1500 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1031 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1500 mV
```

```
    LTC3880-XF0-1.0v-RAIL      949 mV
    LTC3880-XF0-1.0v-CH0       949 mV
    LTC3880-XF0-1.0v-CH1       952 mV
    LTC3880-3.3v-RAIL         3299 mV
    LTC3880-3.3v-CH0          3299 mV
    LTC3880-3.3v-CH1          3299 mV
SFB 5 status:
  State                      Online
  Intake-Zone0 Temperature   27 degrees C / 80 degrees F
  Exhaust-Zone1 Temperature  20 degrees C / 68 degrees F
  IntakeA-Zone0 Temperature  23 degrees C / 73 degrees F
  IntakeB-Zone1 Temperature  15 degrees C / 59 degrees F
  Exhaust-Zone0 Temperature  22 degrees C / 71 degrees F
  SFB-XF2-Zone1 Temperature  27 degrees C / 80 degrees F
  SFB-XF1-Zone0 Temperature  34 degrees C / 93 degrees F
  SFB-XF0-Zone0 Temperature  32 degrees C / 89 degrees F
  Power
    LTC3880-XF2-1.5v-RAIL     1500 mV
    LTC3880-XF2-1.5v-CH0      1500 mV
    LTC3880-XF2-1.5v-CH1      1500 mV
    LTC3880-XF2-1.0v-RAIL      949 mV
    LTC3880-XF2-1.0v-CH0       949 mV
    LTC3880-XF2-1.0v-CH1       951 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1500 mV
    LTC3880-XF1-1.0v-RAIL      949 mV
    LTC3880-XF1-1.0v-CH0       949 mV
    LTC3880-XF1-1.0v-CH1       951 mV
    LTC3880-XF0-1.5v-RAIL     1499 mV
    LTC3880-XF0-1.5v-CH0      1499 mV
    LTC3880-XF0-1.5v-CH1      1500 mV
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1032 mV
    LTC3880-3.3v-RAIL         3299 mV
    LTC3880-3.3v-CH0          3299 mV
    LTC3880-3.3v-CH1          3299 mV
SFB 6 status:
  State                      Online
  Intake-Zone0 Temperature   32 degrees C / 89 degrees F
  Exhaust-Zone1 Temperature  19 degrees C / 66 degrees F
  IntakeA-Zone0 Temperature  24 degrees C / 75 degrees F
```

```
  IntakeB-Zone1 Temperature  15 degrees C / 59 degrees F
   Exhaust-Zone0 Temperature  25 degrees C / 77 degrees F
   SFB-XF2-Zone1 Temperature  29 degrees C / 84 degrees F
   SFB-XF1-Zone0 Temperature  37 degrees C / 98 degrees F
   SFB-XF0-Zone0 Temperature  39 degrees C / 102 degrees F
   Power
     LTC3880-XF2-1.5v-RAIL     1500 mV
     LTC3880-XF2-1.5v-CH0      1500 mV
     LTC3880-XF2-1.5v-CH1      1500 mV
     LTC3880-XF2-1.0v-RAIL     1029 mV
     LTC3880-XF2-1.0v-CH0      1029 mV
     LTC3880-XF2-1.0v-CH1      1031 mV
     LTC3880-XF1-1.5v-RAIL     1499 mV
     LTC3880-XF1-1.5v-CH0      1499 mV
     LTC3880-XF1-1.5v-CH1      1500 mV
     LTC3880-XF1-1.0v-RAIL      949 mV
     LTC3880-XF1-1.0v-CH0       949 mV
     LTC3880-XF1-1.0v-CH1       951 mV
     LTC3880-XF0-1.5v-RAIL     1499 mV
     LTC3880-XF0-1.5v-CH0      1499 mV
     LTC3880-XF0-1.5v-CH1      1500 mV
     LTC3880-XF0-1.0v-RAIL     1029 mV
     LTC3880-XF0-1.0v-CH0      1029 mV
     LTC3880-XF0-1.0v-CH1      1032 mV
     LTC3880-3.3v-RAIL         3300 mV
     LTC3880-3.3v-CH0          3300 mV
     LTC3880-3.3v-CH1          3299 mV
SFB 7 status:
   State                      Online
   Intake-Zone0 Temperature   31 degrees C / 87 degrees F
   Exhaust-Zone1 Temperature  18 degrees C / 64 degrees F
   IntakeA-Zone0 Temperature  20 degrees C / 68 degrees F
   IntakeB-Zone1 Temperature  13 degrees C / 55 degrees F
   Exhaust-Zone0 Temperature  22 degrees C / 71 degrees F
   SFB-XF2-Zone1 Temperature  27 degrees C / 80 degrees F
   SFB-XF1-Zone0 Temperature  26 degrees C / 78 degrees F
   SFB-XF0-Zone0 Temperature  39 degrees C / 102 degrees F
   Power
     LTC3880-XF2-1.5v-RAIL     1499 mV
     LTC3880-XF2-1.5v-CH0      1499 mV
     LTC3880-XF2-1.5v-CH1      1500 mV
     LTC3880-XF2-1.0v-RAIL     1029 mV
     LTC3880-XF2-1.0v-CH0      1029 mV
```

```
    LTC3880-XF2-1.0v-CH1      1031 mV
    LTC3880-XF1-1.5v-RAIL     1499 mV
    LTC3880-XF1-1.5v-CH0      1499 mV
    LTC3880-XF1-1.5v-CH1      1500 mV
    LTC3880-XF1-1.0v-RAIL     1029 mV
    LTC3880-XF1-1.0v-CH0      1029 mV
    LTC3880-XF1-1.0v-CH1      1031 mV
    LTC3880-XF0-1.5v-RAIL     1500 mV
    LTC3880-XF0-1.5v-CH0      1500 mV
    LTC3880-XF0-1.5v-CH1      1500 mV
    LTC3880-XF0-1.0v-RAIL     1029 mV
    LTC3880-XF0-1.0v-CH0      1029 mV
    LTC3880-XF0-1.0v-CH1      1031 mV
    LTC3880-3.3v-RAIL         3300 mV
    LTC3880-3.3v-CH0          3300 mV
    LTC3880-3.3v-CH1          3299 mV
```

show chassis environment sfb (MX2008 Router)

```
user@host> show chassis environment sfb

SFB 0 status:
  State                       Online
  Inlet1 Temperature          26 degrees C / 78 degrees F
  Inlet2 Temperature          27 degrees C / 80 degrees F
  Exhaust1 Temperature        26 degrees C / 78 degrees F
  Exhaust2 Temperature        28 degrees C / 82 degrees F
  SFB2-PF-local Temperature   27 degrees C / 80 degrees F
  SFB2-PF-die Temperature     33 degrees C / 91 degrees F
  Power
    SFB2-PF0-1.0V-PLL           1008 mV
    SFB2-1.0V                    988 mV
    SFB2-10.8V                  10727 mV
    SFB2-1.5V                    1491 mV
    SFB2-PF0-1.8V                1782 mV
    SFB2-2.5V                    2475 mV
    SFB2-Bias-5V                 5535 mV
    SFB2-3.3V-Main               3294 mV
    SFB2-1.8V                    1798 mV
    MAX20751-PF0-0.9v             974 mV
```

```
    MAX20751-PF0-1.0v           1020 mV
SFB 1 status:
  State                       Online
  Inlet1 Temperature          26 degrees C / 78 degrees F
  Inlet2 Temperature          29 degrees C / 84 degrees F
  Exhaust1 Temperature        26 degrees C / 78 degrees F
  Exhaust2 Temperature        29 degrees C / 84 degrees F
  SFB2-PF-local Temperature   29 degrees C / 84 degrees F
  SFB2-PF-die Temperature     41 degrees C / 105 degrees F
  Power
    SFB2-PF0-1.0V-PLL           1014 mV
    SFB2-1.0V                    994 mV
    SFB2-10.8V                  10746 mV
    SFB2-1.5V                   1495 mV
    SFB2-PF0-1.8V               1814 mV
    SFB2-2.5V                   2478 mV
    SFB2-Bias-5V                5594 mV
    SFB2-3.3V-Main              3306 mV
    SFB2-1.8V                   1774 mV
    MAX20751-PF0-0.9v            974 mV
    MAX20751-PF0-1.0v           1020 mV
SFB 2 status:
  State                       Online
  Inlet1 Temperature          29 degrees C / 84 degrees F
  Inlet2 Temperature          36 degrees C / 96 degrees F
  Exhaust1 Temperature        27 degrees C / 80 degrees F
  Exhaust2 Temperature        34 degrees C / 93 degrees F
  SFB2-PF-local Temperature   38 degrees C / 100 degrees F
  SFB2-PF-die Temperature     40 degrees C / 104 degrees F
  Power
    SFB2-PF0-1.0V-PLL           1010 mV
    SFB2-1.0V                    994 mV
    SFB2-10.8V                  10805 mV
    SFB2-1.5V                   1495 mV
    SFB2-PF0-1.8V               1804 mV
    SFB2-2.5V                   2475 mV
    SFB2-Bias-5V                5562 mV

    SFB2-3.3V-Main              3302 mV
    SFB2-1.8V                   1788 mV
    MAX20751-PF0-0.9v            974 mV
    MAX20751-PF0-1.0v           1020
```

show chassis environment sfb 0 (MX10004)

```
user@host> show chassis environment sfb 0
SFB 0 status:
  State                      Online
  Intake-A Temp Sensor Temperature        44 degrees C / 111 degrees F
  Intake-B Temp Sensor Temperature        31 degrees C / 87 degrees F
  Exhaust-A Temp Sensor Temperature       38 degrees C / 100 degrees F
  Exhaust-B Temp Sensor Temperature       41 degrees C / 105 degrees F
  ZF0 Temp Sensor Temperature             47 degrees C / 116 degrees F
  ZF1 Temp Sensor Temperature             41 degrees C / 105 degrees F
  Power
    ZF0 Core 0.9V                  899 mV
    ZF0 AVDD 1V                    999 mV
    ZF1 Core 0.9V                  975 mV
    ZF1 AVDD 1V                    999 mV
    12V                          12311 mV
...
```

show chassis environment sfb (MX10008 Router)

```
user@host> show chassis environment sfb
SFB 0 status:
  State                      Online
  Intake-A Temperature                    33 degrees C / 91 degrees F
  Intake-B Temperature                    22 degrees C / 71 degrees F
  Exhaust-A Temperature                   27 degrees C / 80 degrees F
  Exhaust-B Temperature                   32 degrees C / 89 degrees F
  PF0 Temperature                         38 degrees C / 100 degrees F
  PF1 Temperature                         29 degrees C / 84 degrees F
  Power
    PF0 Core 0.9V                  925 mV
    PF0 AVDD 1V                    999 mV
    PF1 Core 0.9V                  922 mV
    PF1 AVDD 1V                   1000 mV
    12V                          12311 mV
SFB 1 status:
  State                      Online
  Intake-A Temperature                    43 degrees C / 109 degrees F
  Intake-B Temperature                    21 degrees C / 69 degrees F
```

```
  Exhaust-A Temperature                  25 degrees C / 77 degrees F
  Exhaust-B Temperature                  43 degrees C / 109 degrees F
  PF0 Temperature                        49 degrees C / 120 degrees F
  PF1 Temperature                        29 degrees C / 84 degrees F
  Power
    PF0 Core 0.9V                 951 mV
    PF0 AVDD 1V                  1000 mV
    PF1 Core 0.9V                 924 mV
    PF1 AVDD 1V                  1000 mV
    12V                         12311 mV
SFB 2 status:
  State                      Online
  Intake-A Temperature                   39 degrees C / 102 degrees F
  Intake-B Temperature                   21 degrees C / 69 degrees F
  Exhaust-A Temperature                  25 degrees C / 77 degrees F
  Exhaust-B Temperature                  38 degrees C / 100 degrees F
  PF0 Temperature                        44 degrees C / 111 degrees F
  PF1 Temperature                        30 degrees C / 86 degrees F
  Power
    PF0 Core 0.9V                 923 mV
    PF0 AVDD 1V                  1000 mV
    PF1 Core 0.9V                 925 mV
    PF1 AVDD 1V                   999 mV
    12V                         12364 mV
SFB 3 status:
  State                      Online
  Intake-A Temperature                   37 degrees C / 98 degrees F
  Intake-B Temperature                   21 degrees C / 69 degrees F
  Exhaust-A Temperature                  26 degrees C / 78 degrees F
  Exhaust-B Temperature                  34 degrees C / 93 degrees F
  PF0 Temperature                        41 degrees C / 105 degrees F
  PF1 Temperature                        29 degrees C / 84 degrees F
  Power
    PF0 Core 0.9V                 925 mV
    PF0 AVDD 1V                  1000 mV
    PF1 Core 0.9V                 923 mV
    PF1 AVDD 1V                  1000 mV
    12V                         12285 mV
SFB 4 status:
  State                      Online
  Intake-A Temperature                   30 degrees C / 86 degrees F
  Intake-B Temperature                   20 degrees C / 68 degrees F
  Exhaust-A Temperature                  25 degrees C / 77 degrees F
```

```
  Exhaust-B Temperature                   31 degrees C / 87 degrees F
  PF0 Temperature                         40 degrees C / 104 degrees F
  PF1 Temperature                         29 degrees C / 84 degrees F
  Power
    PF0 Core 0.9V                   951 mV
    PF0 AVDD 1V                    1000 mV
    PF1 Core 0.9V                   948 mV
    PF1 AVDD 1V                     999 mV
    12V                           12193 mV
SFB 5 status:
  State                       Online
  Intake-A Temperature                    30 degrees C / 86 degrees F
  Intake-B Temperature                    21 degrees C / 69 degrees F
  Exhaust-A Temperature                   26 degrees C / 78 degrees F
  Exhaust-B Temperature                   30 degrees C / 86 degrees F
  PF0 Temperature                         34 degrees C / 93 degrees F
  PF1 Temperature                         33 degrees C / 91 degrees F
  Power
    PF0 Core 0.9V                   951 mV
    PF0 AVDD 1V                    1000 mV
    PF1 Core 0.9V                   948 mV
    PF1 AVDD 1V                     999 mV
    12V                           12166 mV
```

Release Information

Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 Universal Routing Platforms.

`all-members`, `local`, and `member` *`member-id`* options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.

Command introduced in Junos OS Release 18.2 for MX10008 Universal Routing Platforms.

RELATED DOCUMENTATION

request chassis sfb | **14555**

show chassis sfb | **18823**

show chassis environment sfm

IN THIS SECTION

- Syntax | 17844
- Description | 17844
- Options | 17844
- Required Privilege Level | 17844
- Output Fields | 17845
- Sample Output | 17845
- Release Information | 17848

Syntax

```
show chassis environment sfm
<slot>
```

Description

(M40e and M160 routers only) Display Switching and Forwarding Module (SFM) environmental information.

Options

none Display environmental information about all SFMs.

slot (Optional) Display environmental information about an individual SFM. Replace ***slot*** with a value from **0** through **3**.

Required Privilege Level

view

Output Fields

Table 241 on page 17845 lists the output fields for the `show chassis environment sfm` command. Output fields are listed in the approximate order in which they appear.

Table 241: show chassis environment sfm Output Fields

Field Name	Field Description
SFM *slot* status	SFM slot number: **0** or **1** on an M40e router, or **0**, **1**, **2**, or **3** on an M160 router.
State	Status of the SFM: • **Online**—SFM is online and running. • **Offline**—SFM is powered down. If two SFMs are installed and online, one is functioning as the primary, and the other is marked as the **Standby**.
SPP Temperature	Temperature of the air flowing past the Switch Plane Processor card.
SPR Temperature	Temperature of the air flowing past the Switch Plane Router card.
SPP Power	Information about the voltage supplied to the Switch Plane Processor card. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
SPR Power	Information about the voltage supplied to the Switch Plane Router. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
CMB Revision	Revision level of the Chassis Management Bus (CMB) device.

Sample Output

show chassis environment sfm (M40e Router)

```
user@host> show chassis environment sfm
SFM 0 status:
```

```
  State                           Online
  SPP temperature              40 degrees C / 104 degrees F
  SPR temperature              44 degrees C / 111 degrees F
  SPP Power:
    1.5 V                    1501 mV
    2.5 V                    2472 mV
    3.3 V                    3293 mV
    5.0 V                    5028 mV
    5.0 V bias               4964 mV
  SPR Power:
    1.5 V                    1501 mV
    2.5 V                    2483 mV
    3.3 V                    3308 mV
    5.0 V                    5035 mV
    5.0 V bias               4981 mV
    8.0 V bias               8239 mV
  CMB Revision                 12
SFM 1 status:
  State                           Online - Standby
  SPP temperature              43 degrees C / 109 degrees F
  SPR temperature              45 degrees C / 113 degrees F
  SPP Power:
    1.5 V                    1503 mV
    2.5 V                    2483 mV
    3.3 V                    3284 mV
    5.0 V                    5045 mV
    5.0 V bias               4993 mV
  SPR Power:
    1.5 V                    1498 mV
    2.5 V                    2472 mV
    3.3 V                    3284 mV
    5.0 V                    5035 mV
    5.0 V bias               4991 mV
    8.0 V bias               8231 mV
  CMB Revision                 12
```

show chassis environment sfm (M160 Router)

```
user@host> show chassis environment sfm
SFM 0 status:
  State                           Online
```

```
  SPP temperature              43 degrees C / 109 degrees F
  SPR temperature              44 degrees C / 111 degrees F
  SPP Power:
    1.5 V                    1504 mV
    2.5 V                    2474 mV
    3.3 V                    3290 mV
    5.0 V                    5015 mV
    5.0 V bias               4962 mV
  SPR Power:
    1.5 V                    1498 mV
    2.5 V                    2482 mV
    3.3 V                    3299 mV
    5.0 V                    5020 mV
    5.0 V bias               4971 mV
    8.0 V bias               8229 mV
  CMB Revision                 12
SFM 1 status:
  State                           Online
  SPP temperature              47 degrees C / 116 degrees F
  SPR temperature              50 degrees C / 122 degrees F
  SPP Power:
    1.5 V                    1499 mV
    2.5 V                    2466 mV
    3.3 V                    3274 mV
    5.0 V                    5025 mV
    5.0 V bias               4984 mV
  SPR Power:
    1.5 V                    1496 mV
    2.5 V                    2470 mV
    3.3 V                    3279 mV
    5.0 V                    5020 mV
    5.0 V bias               4993 mV
    8.0 V bias               8222 mV
  CMB Revision                 12
SFM 2 status:
  State                           Online
  SPP temperature              50 degrees C / 122 degrees F
  SPR temperature              52 degrees C / 125 degrees F
  SPP Power:
    1.5 V                    1504 mV
    2.5 V                    2471 mV
    3.3 V                    3294 mV
    5.0 V                    5045 mV
```

```
    5.0 V bias              4981 mV
  SPR Power:
    1.5 V                   1496 mV
    2.5 V                   2470 mV
    3.3 V                   3293 mV
    5.0 V                   5028 mV
    5.0 V bias              4971 mV
    8.0 V bias              8214 mV
  CMB Revision                12
SFM 3 status:
  State                          Online
  SPP temperature             49 degrees C / 120 degrees F
  SPR temperature             48 degrees C / 118 degrees F
  SPP Power:
    1.5 V                   1505 mV
    2.5 V                   2484 mV
    3.3 V                   3296 mV
    5.0 V                   5040 mV
    5.0 V bias              4984 mV
  SPR Power:
    1.5 V                   1503 mV
    2.5 V                   2488 mV
    3.3 V                   3302 mV
    5.0 V                   5037 mV
    5.0 V bias              4993 mV
    8.0 V bias              8249 mV
  CMB Revision                12
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

request chassis sfm | **14558**

request chassis sfm master switch | **14559**

Configuring SFM Redundancy on M40e and M160 Routers

Switching the Global Primary and Backup Roles in a Virtual Chassis Configuration

show chassis environment sib

IN THIS SECTION

- Syntax | 17849
- Syntax (TX Matrix router) | 17849
- Syntax (TX Matrix Plus Router) | 17849
- Description | 17850
- Options | 17850
- Required Privilege Level | 17851
- Output Fields | 17851
- Sample Output | 17852
- Release Information | 17895

Syntax

```
show chassis environment sib
<slot>
```

Syntax (TX Matrix router)

```
show chassis environment sib
<lcc number / scc | slot>
```

Syntax (TX Matrix Plus Router)

```
show chassis environment sib
<sib-slot | lcc number | sfc number | f13 sib-slot | f2s sib-slot/sib-f2s-slot-number>
```

Description

Display Switch Interface Board (SIB) environmental information.

Options

none Display environmental information about all SIBs. On a TX Matrix router, display environmental information about all SIBs on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about all SIBs on the TX Matrix Plus router and its attached routers .

f13 ***sib-slot*** (TX Matrix Plus routers only) (Optional) Display SIB F13 environmental information only. Replace ***sib-slot*** with one of the following values: **0**, **1**, **3**, **4**, **6**, **7**, **8**, **9**, **11**, or **12**. (Slots 2, 5, 10, 13, 14, and 15 are unused).

f2s ***sib-slot/ sib-f2s-slot-number*** (TX Matrix Plus routers only) (Optional) Display SIB F2s environmental information only. Replace *sib-slot* with a value from **0** through **4**, followed by a *sib-f2s-slot-number* value of **0**, **2**, **4** or **6**.

lcc ***number*** (TX Matrix router, and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

scc (TX Matrix routers only) (Optional) Display environmental information about the SIB in the TX Matrix router (switch-card chassis).

sfc (TX Matrix Plus routers only) (Optional) On a TX Matrix Plus router, display environmental information about the SIB in the TX Matrix Plus router (switch-fabric chassis).

sib-slot (Optional) Display environmental information about the specified SIB. For the M320 router, replace ***sib-slot*** with a value from **0** through **3**. For the T640, T1600, T4000, and TX Matrix routers, replace ***sib-slot*** with a value from **0** through **4**. For the TX Matrix Plus router, see

f13 *sib-slot* and f2s *sib-slot/ sib-f2s-slot-number*. For the T320 router, replace ***sib-slot*** with a value from **0** through **2**. For the PTX5000 Packet Transport Router, replace ***sib-slot*** with a value from **0** through **8**.

Required Privilege Level

view

Output Fields

Table 242 on page 17851 lists the output fields for the `show chassis environment sib` command. Output fields are listed in the approximate order in which they appear.

Table 242: show chassis environment sib Output Fields

Field Name	Field Description
SIB slot status	SIB slot number: • **0** through **3** on an M320 router. • **0** or **2** on a T320 router. • **0** through **4** on a T640, T1600, T4000, or TX Matrix router. • **0**, **1**, **3**, **4**, **6**, **7**, **8**, **9**, **11**, or **12** for F13 SIBs on a TX Matrix Plus router. (Slots 2, 5, 10, 13, 14, and 15 are unused). • **0** through **4**, followed by **0**, **2**, **4**, or **6** for F2S SIBs on a TX Matrix Plus router. For example, **SIB F2S 0/4**. • **0** through **8** on a PTX5000 Packet Transport Router.

Table 242: show chassis environment sib Output Fields *(Continued)*

Field Name	Field Description
State	Status of the SIB: • **Online**—SIB is online and running. • **Offline**—SIB is powered down. • **Spare** (T640, T1600, T4000, and TX Matrix routers only)—SIB is redundant and will move to active state if one of the working SIBs fails. Only four of the SIBs are active at any time. The fifth one is marked **Spare**. It is activated if there is a fault on one of the active SIBs. Online standby (TX Matrix Plus router only).
Temperature	Temperature of the air flowing past the SIB. On PTX Series Packet Transport Routers, separate temperatures are displayed for **Intake**, **Exhaust**, and **Junction**.
Power	Information about the voltage supplied to the SIB. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.

Sample Output

show chassis environment sib (M320 Router)

```
user@host> show chassis environment sib
SIB 0 status:
  State                           Online
  Temperature                 34 degrees C / 93 degrees F
  Power:
    GROUND                         0 mV
    1.8 V                       1805 mV
    2.5 V                       2498 mV
    3.3 V                       3306 mV
    1.8 V bias                  1789 mV
    3.3 V bias                  3299 mV
    5.0 V bias                  5003 mV
```

```
    8.0 V bias                7374 mV
SIB 1 status:
  State                           Online
  Temperature                   35 degrees C / 95 degrees F
  Power:
    GROUND                       0 mV
    1.8 V                     1814 mV
    2.5 V                     2477 mV
    3.3 V                     3319 mV
    1.8 V bias                1792 mV
    3.3 V bias                3291 mV
    5.0 V bias                4981 mV
    8.0 V bias                7335 mV
SIB 2 status:
  State                           Online
  Temperature                   33 degrees C / 91 degrees F
  Power:
    GROUND                       0 mV
    1.8 V                     1811 mV
    2.5 V                     2489 mV
    3.3 V                     3330 mV
    1.8 V bias                1797 mV
    3.3 V bias                3304 mV
    5.0 V bias                5025 mV
    8.0 V bias                7330 mV
SIB 3 status:
  State                           Online
  Temperature                   37 degrees C / 98 degrees F
  Power:
    GROUND                       0 mV
    1.8 V                     1798 mV
    2.5 V                     2481 mV
    3.3 V                     3328 mV
    1.8 V bias                1792 mV
    3.3 V bias                3313 mV
    5.0 V bias                5013 mV
    8.0 V bias                7467 mV
```

show chassis environment sib 1 (T640 Router)

```
user@host> show chassis environment sib 1
SIB 1 status:
  State                            Online
  Temperature                   39 degrees C / 102 degrees F
  Power:
    GROUND                        0 mV
    1.8 V                      1809 mV
    2.5 V                      2478 mV
    3.3 V                      3308 mV
    1.8 V bias                 1794 mV
    3.3 V bias                 3274 mV
    5.0 V bias                 4996 mV
    8.0 V bias                 7247 mV
```

show chassis environment sib 1 (T4000 Router)

```
user@host> show chassis environment sib 1
SIB 1 status:
  State                        Online
  Temperature                  42 degrees C / 107 degrees F
  Power
    8.0 V bias                  8100 mV
    3.3 V bias                  3284 mV
    0.9 V bias                   904 mV
    1.1 V bias                  1090 mV
    1.5 V bias                  1488 mV
    2.5 V bias                  2504 mV
    9.0 V                       8940 mV
    3.3 V                       3288 mV
    XF0 1.0 V                    998 mV
    XF0 1.0 V LDO                994 mV
    PCIe SW 1.0 V                990 mV
    XF0 1.8 V                   1788 mV
    XF1 1.0 V                   1002 mV
    XF2 1.0 V                   1002 mV
    XF3 1.0 V                    998 mV
    1.2 V                       1194 mV
    XF1 1.0 V LDO               1000 mV
```

```
    XF2 1.0 V LDO              998 mV
    XF3 1.0 V LDO              998 mV
    XF1 1.8 V                 1798 mV
    XF2 1.8 V                 1800 mV
    XF3 1.8 V                 1794 mV
    1.5 V                     1488 mV
    SW 3.3 V                  3320 mV
```

show chassis environment sib scc (TX Matrix Router)

```
user@host> show chassis environment sib scc
scc-re0:
--------------------------------------------------------------------------
SIB 3 status:
  State                           Offline
  Reason                          Offlined by button press
  Temperature                   0 degrees C / 32 degrees F
  Power:
    GROUND                      0 mV
    1.8 V                       0 mV
    2.5 V                       0 mV
    3.3 V                       0 mV
    1.8 V bias                  0 mV
    3.3 V bias                  0 mV
    5.0 V bias                  0 mV
    8.0 V bias                  0 mV
SIB 4 status:
  State                           Online
  Temperature                  42 degrees C / 107 degrees F
  Temperature (B)              41 degrees C / 105 degrees F
  Power:
    GROUND                      0 mV
    1.8 V                    1787 mV
    2.5 V                    2488 mV
    3.3 V                    3294 mV
    1.8 V bias               1787 mV
    3.3 V bias               3306 mV
    5.0 V bias               5010 mV
    8.0 V bias               7418 mV
  Power (B):
    GROUND                      0 mV
```

```
    1.8 V                    1785 mV
    2.5 V                    2485 mV
    3.3 V                    3289 mV
    1.8 V bias               1799 mV
    3.3 V bias               3284 mV
    5.0 V bias               4979 mV
    8.0 V bias               7882 mV
```

show chassis environment sib (TX Matrix Plus Router)

```
user@host> show chassis environment sib
sfc0-re0:
--------------------------------------------------------------------------
SIB F13 0 status:
  State                      Online - Standby
  Temperature                54 degrees C / 129 degrees F
  Temperature (B)            50 degrees C / 122 degrees F
  Power
    1.2 V_0                   1205 mV
    1.2 V_1                   1202 mV
    1.2 V_2                   1205 mV
    1.2 V_3                   1208 mV
    1.5 V_0                   1501 mV
    1.5 V_1                   1508 mV
    1.8 V                     1798 mV
    2.5 V                     2510 mV
    3.3 V                     3312 mV
    9.0 V                     8991 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2510 mV
    3.3 V                     3318 mV
    9.0 V                     9024 mV
SIB F13 1 status:
  State                      Online - Standby
  Temperature                45 degrees C / 113 degrees F
  Temperature (B)            42 degrees C / 107 degrees F
  Power
    1.2 V_0                   1202 mV
    1.2 V_1                   1198 mV
    1.2 V_2                   1202 mV
```

```
    1.2 V_3                   1202 mV
    1.5 V_0                   1498 mV
    1.5 V_1                   1501 mV
    1.8 V                     1811 mV
    2.5 V                     2504 mV
    3.3 V                     3292 mV
    9.0 V                     8991 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2507 mV
    3.3 V                     3306 mV
    9.0 V                     8970 mV
SIB F13 3 status:
  State                      Online
  Temperature                48 degrees C / 118 degrees F
  Temperature (B)            44 degrees C / 111 degrees F
  Power
    1.2 V_0                   1205 mV
    1.2 V_1                   1202 mV
    1.2 V_2                   1202 mV
    1.2 V_3                   1202 mV
    1.5 V_0                   1508 mV
    1.5 V_1                   1504 mV
    1.8 V                     1798 mV
    2.5 V                     2520 mV
    3.3 V                     3300 mV
    9.0 V                     9009 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2504 mV
    3.3 V                     3312 mV
    9.0 V                     9006 mV
SIB F13 4 status:
  State                      Online
  Temperature                44 degrees C / 111 degrees F
  Temperature (B)            40 degrees C / 104 degrees F
Power
    1.2 V_0                   1205 mV
    1.2 V_1                   1205 mV
    1.2 V_2                   1202 mV
    1.2 V_3                   1205 mV
    1.5 V_0                   1508 mV
    1.5 V_1                   1508 mV
```

```
    1.8 V                     1811 mV
    2.5 V                     2510 mV
    3.3 V                     3312 mV
    9.0 V                     8970 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2513 mV
    3.3 V                     3318 mV
    9.0 V                     9048 mV
SIB F13 6 status:
  State                      Online
  Temperature                50 degrees C / 122 degrees F
  Temperature (B)            46 degrees C / 114 degrees F
  Power
    1.2 V_0                   1195 mV
    1.2 V_1                   1205 mV
    1.2 V_2                   1202 mV
    1.2 V_3                   1202 mV
    1.5 V_0                   1495 mV
    1.5 V_1                   1495 mV
    1.8 V                     1801 mV
    2.5 V                     2494 mV
    3.3 V                     3300 mV
    9.0 V                     8991 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2500 mV
    3.3 V                     3300 mV
    9.0 V                     9006 mV
SIB F13 7 status:
  State                      Online
  Temperature                52 degrees C / 125 degrees F
  Temperature (B)            49 degrees C / 120 degrees F
  Power
    1.2 V_0                   1202 mV
    1.2 V_1                   1202 mV
    1.2 V_2                   1198 mV
    1.2 V_3                   1185 mV
    1.5 V_0                   1501 mV
    1.5 V_1                   1492 mV
    1.8 V                     1795 mV
    2.5 V                     2491 mV
    3.3 V                     3286 mV
```

```
    9.0 V                     8892 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2507 mV
    3.3 V                     3306 mV
    9.0 V                     8952 mV
SIB F13 8 status:
  State                      Online
  Temperature                55 degrees C / 131 degrees F
  Temperature (B)            50 degrees C / 122 degrees F
  Power
    1.2 V_0                   1208 mV
    1.2 V_1                   1205 mV
    1.2 V_2                   1205 mV
    1.2 V_3                   1211 mV
    1.5 V_0                   1514 mV
    1.5 V_1                   1508 mV
    1.8 V                     1807 mV
    2.5 V                     2516 mV
    3.3 V                     3324 mV
    9.0 V                     9027 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2520 mV
    3.3 V                     3318 mV
    9.0 V                     9066 mV
SIB F13 9 status:
  State                      Online
  Temperature                46 degrees C / 114 degrees F
  Temperature (B)            41 degrees C / 105 degrees F
  Power
    1.2 V_0                   1208 mV
    1.2 V_1                   1202 mV
    1.2 V_2                   1208 mV
    1.2 V_3                   1202 mV
    1.5 V_0                   1504 mV
    1.5 V_1                   1504 mV
    1.8 V                     1817 mV
    2.5 V                     2516 mV
    3.3 V                     3312 mV
    9.0 V                     9009 mV
    9.0 V bias                   0 mV
  Power (B)
```

```
    2.5 V                     2510 mV
    3.3 V                     3312 mV
    9.0 V                     9024 mV
SIB F13 11 status:
  State                      Online
  Temperature                47 degrees C / 116 degrees F
  Temperature (B)            42 degrees C / 107 degrees F
  Power
    1.2 V_0                   1202 mV
    1.2 V_1                   1205 mV
    1.2 V_2                   1202 mV
    1.2 V_3                   1202 mV
    1.5 V_0                   1501 mV
    1.5 V_1                   1501 mV
    1.8 V                     1801 mV
    2.5 V                     2510 mV
    3.3 V                     3312 mV
    9.0 V                     8979 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2252 mV
    3.3 V                     5014 mV
    9.0 V                     9954 mV
SIB F13 12 status:
  State                      Online
  Temperature                45 degrees C / 113 degrees F
  Temperature (B)            40 degrees C / 104 degrees F
  Power
    1.2 V_0                   1211 mV
    1.2 V_1                   1208 mV
    1.2 V_2                   1205 mV
    1.2 V_3                   1205 mV
    1.5 V_0                   1511 mV
    1.5 V_1                   1501 mV
    1.8 V                     1817 mV
    2.5 V                     2504 mV
    3.3 V                     3318 mV
    9.0 V                     9027 mV
    9.0 V bias                   0 mV
 Power (B)
    2.5 V                     2520 mV
    3.3 V                     3338 mV
    9.0 V                     9006 mV
```

```
SIB F2S 0/0 status:
  State                       Online - Standby
  Temperature                 40 degrees C / 104 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1198 mV
    1.2 V_ASF_B                1198 mV
    1.2 V_ASF_D                1202 mV
    1.5 V                      1498 mV
    1.8 V                      1814 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3286 mV
    9.0 V                      8250 mV
SIB F2S 0/2 status:
  State                       Online - Standby
  Temperature                 40 degrees C / 104 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1198 mV
    1.2 V_ASF_B                1195 mV
    1.2 V_ASF_D                1202 mV
    1.5 V                      1498 mV
    1.8 V                      1807 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3286 mV
    9.0 V                      8250 mV
SIB F2S 0/4 status:
  State                       Online - Standby
  Temperature                 40 degrees C / 104 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1202 mV
    1.2 V_ASF_B                1198 mV
    1.2 V_ASF_D                1202 mV
    1.5 V                      1504 mV
    1.8 V                      1817 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3306 mV
    9.0 V                      8250 mV
SIB F2S 0/6 status:
```

```
  State                       Online - Standby
  Temperature                 39 degrees C / 102 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1202 mV
    1.2 V_ASF_B                1198 mV
    1.2 V_ASF_D                1202 mV
    1.5 V                      1495 mV
    1.8 V                      1814 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3280 mV
    9.0 V                      8250 mV
SIB F2S 1/0 status:
  State                       Online
  Temperature                 39 degrees C / 102 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1195 mV
    1.2 V_ASF_B                1192 mV
    1.2 V_ASF_D                1195 mV
    1.5 V                      1488 mV
    1.8 V                      1798 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3280 mV
    9.0 V                      8250 mV
SIB F2S 1/2 status:
  State                       Online
  Temperature                 39 degrees C / 102 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1205 mV
    1.2 V_ASF_B                1202 mV
    1.2 V_ASF_D                1205 mV
    1.5 V                      1501 mV
    1.8 V                      1820 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3306 mV
    9.0 V                      8250 mV
SIB F2S 1/4 status:
  State                       Online
```

```
  Temperature                 39 degrees C / 102 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1198 mV
    1.2 V_ASF_B                1195 mV
    1.2 V_ASF_D                1195 mV
    1.5 V                      1498 mV
    1.8 V                      1811 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3300 mV
    9.0 V                      8250 mV
SIB F2S 1/6 status:
  State                       Online
  Temperature                 39 degrees C / 102 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1195 mV
    1.2 V_ASF_B                1195 mV
    1.2 V_ASF_D                1198 mV
    1.5 V                      1498 mV
    1.8 V                      1807 mV
    3.3 V                      3306 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3292 mV
    9.0 V                      8250 mV
SIB F2S 2/0 status:
  State                       Online
  Temperature                 39 degrees C / 102 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1195 mV
    1.2 V_ASF_B                1195 mV
    1.2 V_ASF_D                1198 mV
    1.5 V                      1498 mV
    1.8 V                      1804 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3286 mV
    9.0 V                      8250 mV
SIB F2S 2/2 status:
  State                       Online
Temperature                 38 degrees C / 100 degrees F
```

```
  Power
    1.2 V_1                      0 mV
    1.2 V_ASF                 1195 mV
    1.2 V_ASF_B               1195 mV
    1.2 V_ASF_D               1198 mV
    1.5 V                     1495 mV
    1.8 V                     1807 mV
    3.3 V                     3300 mV
    3.3 V bias                3300 mV
    3.3 V ASF                 3300 mV
    9.0 V                     8250 mV
SIB F2S 2/4 status:
  State                      Online
  Temperature                38 degrees C / 100 degrees F
  Power
    1.2 V_1                      0 mV
    1.2 V_ASF                 1198 mV
    1.2 V_ASF_B               1195 mV
    1.2 V_ASF_D               1198 mV
    1.5 V                     1501 mV
    1.8 V                     1804 mV
    3.3 V                     3286 mV
    3.3 V bias                3292 mV
    3.3 V ASF                 3300 mV
    9.0 V                     8230 mV
SIB F2S 2/6 status:
  State                      Online
  Temperature                38 degrees C / 100 degrees F
  Power
    1.2 V_1                      0 mV
    1.2 V_ASF                 1202 mV
    1.2 V_ASF_B               1198 mV
    1.2 V_ASF_D               1202 mV
    1.5 V                     1501 mV
    1.8 V                     1817 mV
    3.3 V                     3300 mV
    3.3 V bias                3300 mV
    3.3 V ASF                 3318 mV
    9.0 V                     8250 mV
SIB F2S 3/0 status:
  State                      Online
  Temperature                38 degrees C / 100 degrees F
  Power
```

```
    1.2 V_1                      0 mV
    1.2 V_ASF                 1195 mV
    1.2 V_ASF_B               1195 mV
    1.2 V_ASF_D               1198 mV
    1.5 V                     1501 mV
    1.8 V                     1814 mV
    3.3 V                     3300 mV
    3.3 V bias                3300 mV
    3.3 V ASF                 3274 mV
    9.0 V                     8250 mV
SIB F2S 3/2 status:
  State                      Online
  Temperature                37 degrees C / 98 degrees F
  Power
    1.2 V_1                      0 mV
    1.2 V_ASF                 1202 mV
    1.2 V_ASF_B               1195 mV
    1.2 V_ASF_D               1195 mV
    1.5 V                     1495 mV
    1.8 V                     1804 mV
    3.3 V                     3300 mV
    3.3 V bias                3300 mV
    3.3 V ASF                 3286 mV
    9.0 V                     8250 mV
SIB F2S 3/4 status:
  State                      Online
  Temperature                37 degrees C / 98 degrees F
  Power
    1.2 V_1                      0 mV
    1.2 V_ASF                 1205 mV
    1.2 V_ASF_B               1198 mV
    1.2 V_ASF_D               1202 mV
    1.5 V                     1501 mV
    1.8 V                     1811 mV
    3.3 V                     3300 mV
    3.3 V bias                3300 mV
    3.3 V ASF                 3318 mV
    9.0 V                     8250 mV
SIB F2S 3/6 status:
  State                      Online
  Temperature                37 degrees C / 98 degrees F
  Power
    1.2 V_1                      0 mV
```

```
    1.2 V_ASF                 1205 mV
    1.2 V_ASF_B               1202 mV
    1.2 V_ASF_D               1202 mV
    1.5 V                     1511 mV
    1.8 V                     1820 mV
    3.3 V                     3306 mV
    3.3 V bias                3306 mV
    3.3 V ASF                 3318 mV
    9.0 V                     8265 mV
SIB F2S 4/0 status:
  State                      Online
  Temperature                36 degrees C / 96 degrees F
  Power
    1.2 V_1                      0 mV
    1.2 V_ASF                 1198 mV
    1.2 V_ASF_B               1198 mV
    1.2 V_ASF_D               1198 mV
    1.5 V                     1501 mV
    1.8 V                     1814 mV
    3.3 V                     3292 mV
    3.3 V bias                3292 mV
    3.3 V ASF                 3312 mV
    9.0 V                     8230 mV
SIB F2S 4/2 status:
  State                      Online
  Temperature                37 degrees C / 98 degrees F
  Power
    1.2 V_1                      0 mV
    1.2 V_ASF                 1198 mV
    1.2 V_ASF_B               1192 mV
    1.2 V_ASF_D               1195 mV
    1.5 V                     1495 mV
    1.8 V                     1807 mV
    3.3 V                     3300 mV
    3.3 V bias                3300 mV
    3.3 V ASF                 3300 mV
    9.0 V                     8250 mV
SIB F2S 4/4 status:
  State                      Online
  Temperature                36 degrees C / 96 degrees F
  Power
    1.2 V_1                      0 mV
    1.2 V_ASF                 1202 mV
```

```
    1.2 V_ASF_B               1195 mV
    1.2 V_ASF_D               1202 mV
   1.5 V                     1501 mV
    1.8 V                     1814 mV
    3.3 V                     3300 mV
    3.3 V bias                3300 mV
    3.3 V ASF                 3312 mV
    9.0 V                     8250 mV
SIB F2S 4/6 status:
  State                      Online
  Temperature                36 degrees C / 96 degrees F
  Power
    1.2 V_1                      0 mV
    1.2 V_ASF                 1198 mV
    1.2 V_ASF_B               1195 mV
    1.2 V_ASF_D               1198 mV
    1.5 V                     1498 mV
    1.8 V                     1820 mV
    3.3 V                     3292 mV
    3.3 V bias                3292 mV
    3.3 V ASF                 3286 mV
    9.0 V                     8230 mV

lcc0-re0:
--------------------------------------------------------------------------
SIB 0 status:
  State                      Online - Standby
  Temperature                49 degrees C / 120 degrees F
  Temperature (B)            42 degrees C / 107 degrees F
  Power
    1.2 V                     1204 mV
    1.5 V                     1484 mV
    2.5 V                     2500 mV
    3.3 V                     3312 mV
    3.3 V bias                3312 mV
    5.0 V bias                4956 mV
    8.0 V bias                7740 mV
    9.0 V                     8880 mV
  Power (B)
    1.2 V                     1206 mV
    2.5 V                     2500 mV
    3.3 V                     3316 mV
    9.0 V                     8988 mV
```

```
SIB 1 status:
  State                      Online
  Temperature                49 degrees C / 120 degrees F
  Temperature (B)            42 degrees C / 107 degrees F
  Power
    1.2 V                     1202 mV
    1.5 V                     1482 mV
    2.5 V                     2500 mV
    3.3 V                     3296 mV
    3.3 V bias                3288 mV
    5.0 V bias                4986 mV
    8.0 V bias                7800 mV
    9.0 V                     8868 mV
  Power (B)
    1.2 V                     1206 mV
    2.5 V                     2512 mV
    3.3 V                     3312 mV
    9.0 V                     8952 mV
SIB 2 status:
  State                      Online
  Temperature                49 degrees C / 120 degrees F
  Temperature (B)            42 degrees C / 107 degrees F
  Power
    1.2 V                     1202 mV
  1.5 V                      1480 mV
    2.5 V                     2476 mV
    3.3 V                     3292 mV
    3.3 V bias                3308 mV
    5.0 V bias                5010 mV
    8.0 V bias                7800 mV
    9.0 V                     8880 mV
  Power (B)
    1.2 V                     1204 mV
    2.5 V                     2516 mV
    3.3 V                     3308 mV
    9.0 V                     8988 mV
SIB 3 status:
  State                      Online
  Temperature                48 degrees C / 118 degrees F
  Temperature (B)            42 degrees C / 107 degrees F
  Power
    1.2 V                     1204 mV
    1.5 V                     1480 mV
```

```
    2.5 V                     2500 mV
    3.3 V                     3292 mV
    3.3 V bias                3292 mV
    5.0 V bias                4986 mV
    8.0 V bias                7812 mV
    9.0 V                     8892 mV
  Power (B)
    1.2 V                     1198 mV
    2.5 V                     2512 mV
    3.3 V                     3308 mV
    9.0 V                     8892 mV
SIB 4 status:
  State                      Online
  Temperature                48 degrees C / 118 degrees F
  Temperature (B)            42 degrees C / 107 degrees F
  Power
    1.2 V                     1206 mV
    1.5 V                     1482 mV
    2.5 V                     2484 mV
    3.3 V                     3324 mV
    3.3 V bias                3340 mV
    5.0 V bias                4980 mV
    8.0 V bias                7764 mV
    9.0 V                     8784 mV
  Power (B)
    1.2 V                     1202 mV
    2.5 V                     2504 mV
    3.3 V                     3308 mV
    9.0 V                     8820 mV
lcc1-re0:
--------------------------------------------------------------------------
SIB 0 status:
  State                      Online - Standby
  Temperature                49 degrees C / 120 degrees F
  Temperature (B)            43 degrees C / 109 degrees F
  Power
    1.2 V                     1206 mV
    1.5 V                     1506 mV
    2.5 V                     2496 mV
    3.3 V                     3308 mV
    3.3 V bias                3296 mV
    5.0 V bias                4974 mV
    8.0 V bias                7884 mV
```

```
    9.0 V                     8820 mV
  Power (B)
    1.2 V                     1200 mV
    2.5 V                     2508 mV
    3.3 V                     3292 mV
    9.0 V                     8892 mV
...
```

show chassis environment sib sfc (TX Matrix Plus Router)

```
user@host> show chassis environment sib sfc
sfc0-re0:
--------------------------------------------------------------------------
SIB F13 0 status:
  State                      Online - Standby
  Temperature                54 degrees C / 129 degrees F
  Temperature (B)            50 degrees C / 122 degrees F
  Power
    1.2 V_0                   1205 mV
    1.2 V_1                   1205 mV
    1.2 V_2                   1208 mV
    1.2 V_3                   1208 mV
    1.5 V_0                   1501 mV
    1.5 V_1                   1508 mV
    1.8 V                     1804 mV
    2.5 V                     2504 mV
    3.3 V                     3312 mV
    9.0 V                     8991 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2516 mV
    3.3 V                     3318 mV
    9.0 V                     9048 mV
SIB F13 1 status:
  State                      Online - Standby
  Temperature                45 degrees C / 113 degrees F
  Temperature (B)            42 degrees C / 107 degrees F
  Power
    1.2 V_0                   1202 mV
    1.2 V_1                   1205 mV
    1.2 V_2                   1198 mV
```

```
    1.2 V_3                   1205 mV
    1.5 V_0                   1498 mV
    1.5 V_1                   1495 mV
    1.8 V                     1801 mV
    2.5 V                     2507 mV
    3.3 V                     3306 mV
    9.0 V                     8970 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2507 mV
    3.3 V                     3306 mV
    9.0 V                     8970 mV
SIB F13 3 status:
  State                      Online
  Temperature                48 degrees C / 118 degrees F
  Temperature (B)            43 degrees C / 109 degrees F
  Power
    1.2 V_0                   1208 mV
    1.2 V_1                   1195 mV
    1.2 V_2                   1202 mV
    1.2 V_3                   1198 mV
    1.5 V_0                   1504 mV
    1.5 V_1                   1504 mV
    1.8 V                     1801 mV
    2.5 V                     2510 mV
    3.3 V                     3312 mV
    9.0 V                     8970 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2500 mV
    3.3 V                     3332 mV
    9.0 V                     8970 mV
SIB F13 4 status:
  State                      Online
  Temperature                44 degrees C / 111 degrees F
  Temperature (B)            40 degrees C / 104 degrees F
  Power
    1.2 V_0                   1205 mV
    1.2 V_1                   1202 mV
    1.2 V_2                   1205 mV
    1.2 V_3                   1202 mV
    1.5 V_0                   1508 mV
    1.5 V_1                   1511 mV
```

```
    1.8 V                     1811 mV
    2.5 V                     2510 mV
    3.3 V                     3312 mV
    9.0 V                     8952 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2510 mV
    3.3 V                     3306 mV
    9.0 V                     9024 mV
SIB F13 6 status:
  State                      Online
  Temperature                49 degrees C / 120 degrees F
  Temperature (B)            46 degrees C / 114 degrees F
  Power
    1.2 V_0                   1195 mV
    1.2 V_1                   1198 mV
    1.2 V_2                   1202 mV
    1.2 V_3                   1202 mV
    1.5 V_0                   1501 mV
    1.5 V_1                   1495 mV
    1.8 V                     1801 mV
    2.5 V                     2507 mV
    3.3 V                     3306 mV
    9.0 V                     8979 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2497 mV
    3.3 V                     3318 mV
    9.0 V                     9006 mV
SIB F13 7 status:
  State                      Online
  Temperature                52 degrees C / 125 degrees F
  Temperature (B)            48 degrees C / 118 degrees F
  Power
    1.2 V_0                   1198 mV
    1.2 V_1                   1198 mV
    1.2 V_2                   1202 mV
    1.2 V_3                   1189 mV
    1.5 V_0                   1498 mV
    1.5 V_1                   1498 mV
    1.8 V                     1804 mV
    2.5 V                     2491 mV
    3.3 V                     3292 mV
```

```
    9.0 V                     8904 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2500 mV
    3.3 V                     3306 mV
    9.0 V                     8952 mV
SIB F13 8 status:
  State                      Online
  Temperature                54 degrees C / 129 degrees F
  Temperature (B)            49 degrees C / 120 degrees F
  Power
    1.2 V_0                   1211 mV
    1.2 V_1                   1208 mV
    1.2 V_2                   1208 mV
    1.2 V_3                   1211 mV
    1.5 V_0                   1508 mV
    1.5 V_1                   1511 mV
    1.8 V                     1801 mV
    2.5 V                     2513 mV
    3.3 V                     3324 mV
    9.0 V                     9048 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2516 mV
    3.3 V                     3318 mV
    9.0 V                     9102 mV
SIB F13 9 status:
  State                      Online
  Temperature                46 degrees C / 114 degrees F
  Temperature (B)            41 degrees C / 105 degrees F
  Power
    1.2 V_0                   1205 mV
    1.2 V_1                   1202 mV
    1.2 V_2                   1205 mV
    1.2 V_3                   1198 mV
    1.5 V_0                   1504 mV
    1.5 V_1                   1504 mV
    1.8 V                     1817 mV
    2.5 V                     2507 mV
    3.3 V                     3306 mV
    9.0 V                     8991 mV
    9.0 V bias                   0 mV
  Power (B)
```

```
    2.5 V                     2510 mV
    3.3 V                     3332 mV
    9.0 V                     9006 mV
SIB F13 11 status:
  State                      Online
  Temperature                47 degrees C / 116 degrees F
  Temperature (B)            42 degrees C / 107 degrees F
  Power
    1.2 V_0                   1202 mV
    1.2 V_1                   1205 mV
    1.2 V_2                   1202 mV
    1.2 V_3                   1198 mV
    1.5 V_0                   1501 mV
    1.5 V_1                   1504 mV
    1.8 V                     1807 mV
    2.5 V                     2510 mV
    3.3 V                     3306 mV
    9.0 V                     8991 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2249 mV
    3.3 V                     4994 mV
    9.0 V                     9936 mV
SIB F13 12 status:
  State                      Online
  Temperature                44 degrees C / 111 degrees F
  Temperature (B)            40 degrees C / 104 degrees F
  Power
    1.2 V_0                   1208 mV
    1.2 V_1                   1202 mV
    1.2 V_2                   1208 mV
    1.2 V_3                   1205 mV
    1.5 V_0                   1511 mV
    1.5 V_1                   1508 mV
    1.8 V                     1814 mV
    2.5 V                     2507 mV
    3.3 V                     3318 mV
    9.0 V                     9039 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2516 mV
    3.3 V                     3344 mV
    9.0 V                     9006 mV
```

```
SIB F2S 0/0 status:
  State                       Online - Standby
  Temperature                 40 degrees C / 104 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1198 mV
    1.2 V_ASF_B                1198 mV
    1.2 V_ASF_D                1202 mV
    1.5 V                      1498 mV
    1.8 V                      1814 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3286 mV
    9.0 V                      8250 mV
SIB F2S 0/2 status:
  State                       Online - Standby
  Temperature                 40 degrees C / 104 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1198 mV
    1.2 V_ASF_B                1195 mV
    1.2 V_ASF_D                1202 mV
    1.5 V                      1498 mV
    1.8 V                      1807 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3292 mV
    9.0 V                      8250 mV
SIB F2S 0/4 status:
  State                       Online - Standby
  Temperature                 40 degrees C / 104 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1198 mV
    1.2 V_ASF_B                1195 mV
    1.2 V_ASF_D                1202 mV
    1.5 V                      1501 mV
    1.8 V                      1817 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3306 mV
    9.0 V                      8250 mV
SIB F2S 0/6 status:
```

```
  State                       Online - Standby
  Temperature                 39 degrees C / 102 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1202 mV
    1.2 V_ASF_B                1198 mV
    1.2 V_ASF_D                1198 mV
    1.5 V                      1495 mV
    1.8 V                      1814 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3280 mV
    9.0 V                      8250 mV
SIB F2S 1/0 status:
  State                       Online
  Temperature                 39 degrees C / 102 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1195 mV
    1.2 V_ASF_B                1192 mV
    1.2 V_ASF_D                1195 mV
    1.5 V                      1492 mV
    1.8 V                      1798 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3280 mV
    9.0 V                      8250 mV
SIB F2S 1/2 status:
  State                       Online
  Temperature                 39 degrees C / 102 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1205 mV
    1.2 V_ASF_B                1202 mV
    1.2 V_ASF_D                1205 mV
    1.5 V                      1504 mV
    1.8 V                      1820 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3306 mV
    9.0 V                      8250 mV
SIB F2S 1/4 status:
  State                       Online
```

```
  Temperature                 39 degrees C / 102 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1202 mV
    1.2 V_ASF_B                1195 mV
    1.2 V_ASF_D                1198 mV
    1.5 V                      1498 mV
    1.8 V                      1811 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3300 mV
    9.0 V                      8250 mV
SIB F2S 1/6 status:
  State                       Online
  Temperature                 39 degrees C / 102 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1195 mV
    1.2 V_ASF_B                1192 mV
    1.2 V_ASF_D                1198 mV
    1.5 V                      1498 mV
    1.8 V                      1807 mV
    3.3 V                      3306 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3292 mV
    9.0 V                      8250 mV
SIB F2S 2/0 status:
  State                       Online
  Temperature                 38 degrees C / 100 degrees F
  Power
    1.2 V_1                       0 mV
    1.2 V_ASF                  1195 mV
    1.2 V_ASF_B                1195 mV
    1.2 V_ASF_D                1198 mV
    1.5 V                      1498 mV
    1.8 V                      1804 mV
    3.3 V                      3300 mV
    3.3 V bias                 3300 mV
    3.3 V ASF                  3292 mV
    9.0 V                      8250 mV
...
```

show chassis environment sib f13 (TX Matrix Plus Router)

```
user@host> show chassis environment sib f13 0
SIB F13 0 status:
  State                      Online - Standby
  Temperature                54 degrees C / 129 degrees F
  Temperature (B)            50 degrees C / 122 degrees F
  Power
    1.2 V_0                   1202 mV
    1.2 V_1                   1202 mV
    1.2 V_2                   1208 mV
    1.2 V_3                   1208 mV
    1.5 V_0                   1501 mV
    1.5 V_1                   1504 mV
    1.8 V                     1801 mV
    2.5 V                     2504 mV
    3.3 V                     3318 mV
    9.0 V                     8991 mV
    9.0 V bias                   0 mV
  Power (B)
    2.5 V                     2510 mV
    3.3 V                     3318 mV
    9.0 V                     9024 mV
```

show chassis environment sib f2s (TX Matrix Plus Router)

```
user@host> show chassis environment sib f2s 0/2
SIB F2S 0/2 status:
  State                      Online - Standby
  Temperature                40 degrees C / 104 degrees F
  Power
    1.2 V_1                      0 mV
    1.2 V_ASF                 1198 mV
    1.2 V_ASF_B               1195 mV
    1.2 V_ASF_D               1202 mV
    1.5 V                     1501 mV
    1.8 V                     1807 mV
    3.3 V                     3300 mV
    3.3 V bias                3300 mV
```

```
    3.3 V ASF                 3286 mV
    9.0 V                     8250 mV
```

show chassis environment sib (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis environment sib
sfc0-re0:
--------------------------------------------------------------------------
SIB F13 0 status:
  State                      Online
  Board Temperature          44 degrees C / 111 degrees F
  XF Junction Temperature    62 degrees C / 143 degrees F
  Power
    XF F1 LCC0 1.0 V           999 mV
    PCIe Switch 1.0 V         1000 mV
    XF F3 LCC0 1.0 V          1000 mV
    XF F1/F3 LCC0 1.2 V       1199 mV
    XF F1 LCC1 1.0 V          1000 mV
    XF F1/F3 LCC1 1.2 V       1199 mV
    XF F3 LCC1 1.0 V          1000 mV
    XF F1/F3 1.5 V            1499 mV
    XF RC LCC0 Base 1.0       1000 mV
    XF RC Base 1.2 V          1200 mV
    XF RC LCC1 Base 1.0       1000 mV
    XF RC Base 1.5 V          1499 mV
    3.3 V Base                3300 mV
    VSC8248 Base 1.8V         1796 mV
    FPGA Core 0.9 V            899 mV
    2.5 V Base                2500 mV
    ltc3880-3.3v-bias         3343 mV
    CXP Base 4.0 V            3999 mV
    XF RC LCC0 Mezz 1.0       1000 mV
    XF RC Mezz 1.2 V          1199 mV
    XF RC LCC1 Mezz 1.0        999 mV
    XF RC Mezz 1.5 V          1499 mV
    3.3 V Mezz                3299 mV
    VSC8248 Mezz 1.8V         1800 mV
    CXP Mezz 4.0 V            3999 mV
 [...Output Truncated...]
SIB F2S 0/0 status:
  State                      Online
```

```
  Board Temperature          32 degrees C / 89 degrees F
  XF Junction Temperature    41 degrees C / 105 degrees F
  Power
    XF F2S 1.8 V LD0          1775 mV
    XF F2S 1.0 V AN            992 mV
    XF F2S 1.0 V              1002 mV
    XF F2S 1.5 V              1488 mV
    1.5 V Base                2500 mV
    3.3 V bias                3306 mV
    3.3 V Base                3280 mV
    12.0 V Base              11928 mV
 [...Output Truncated...]
SIB F2S 2/6 status:
  State                      Online
  Board Temperature          28 degrees C / 82 degrees F
  XF Junction Temperature    40 degrees C / 104 degrees F
  Power
    XF F2S 1.8 V LD0          1782 mV
    XF F2S 1.0 V AN            999 mV
    XF F2S 1.0 V              1005 mV
    XF F2S 1.5 V              1498 mV
    1.5 V Base                2510 mV
    3.3 V bias                3292 mV
    3.3 V Base                3292 mV
    12.0 V Base              12024 mV

lcc0-re0:
--------------------------------------------------------------------------
SIB 0 status:
  State                      Online
  Temperature                41 degrees C / 105 degrees F
  Temperature (B)            Absent
  Max Jn Temperature         48 degrees C / 118 degrees F
  Power
    8.0 V bias                8156 mV
    3.3 V bias                3284 mV
    FPGA 0.9 V bias            908 mV
    FPGA 1.1 V bias           1086 mV
    FPGA 1.5 V bias           1487 mV
    FPGA 2.5 V bias           2525 mV
    3.3 V                     3282 mV
    1.5 V                     1487 mV
    XF HSS 1.5 V              1501 mV
```

```
    XF1 1.0 V                 1001 mV
    XF2 1.0 V                 1003 mV
    XF3 1.0 V                  998 mV
    XF1 1.8 V LD0             1782 mV
    XF2 1.8 V LDO             1792 mV
    XF3 1.8 V LDO             1782 mV
    CLK  BUF 2.5 V LDO        2493 mV
    XF1 1.0 V LDO              991 mV
    XF2 1.0 V LDO              991 mV
    XF3 1.0 V LDO              991 mV
    PCIe SW 3.3 V             3274 mV
    PCIe 1.0 V                 996 mV
    RETIMER 1.2 V             1174 mV
    RETIMER IO 1.8 V          1770 mV
                                 0 mV
  Power (B)
    1.2 V                        0 mV
    2.5 V                        0 mV
    3.3 V                        0 mV
    9.0 V                        0 mV
 [...Output Truncated...]
lcc2-re0:
--------------------------------------------------------------------------
SIB 0 status:
  State                      Online
  Temperature                42 degrees C / 107 degrees F
  Temperature (B)            Absent
  Max Jn Temperature         51 degrees C / 123 degrees F
  Power
    8.0 V bias                8146 mV
    3.3 V bias                3277 mV
    FPGA 0.9 V bias            903 mV
    FPGA 1.1 V bias           1089 mV
    FPGA 1.5 V bias           1479 mV
    FPGA 2.5 V bias           2515 mV
    3.3 V                     3277 mV
    1.5 V                     1482 mV
    XF HSS 1.5 V              1501 mV
    XF1 1.0 V                 1001 mV
    XF2 1.0 V                 1003 mV
    XF3 1.0 V                  998 mV
    XF1 1.8 V LD0             1787 mV
    XF2 1.8 V LDO             1792 mV
```

```
    XF3 1.8 V LDO             1792 mV
    CLK  BUF 2.5 V LDO        2481 mV
    XF1 1.0 V LDO              986 mV
    XF2 1.0 V LDO              993 mV
    XF3 1.0 V LDO              991 mV
    PCIe SW 3.3 V             3279 mV
    PCIe 1.0 V                 991 mV
    RETIMER 1.2 V             1179 mV
    RETIMER IO 1.8 V          1772 mV
                                 0 mV
  Power (B)
    1.2 V                        0 mV
    2.5 V                        0 mV
    3.3 V                        0 mV
    9.0 V                        0 mV
 [...Output Truncated...]
```

show chassis environment sib (PTX5000 Packet Transport Router with SIB-I-8S)

```
user@host> show chassis environment sib
SIB 0 status:
  State                      Online
  Exhaust Temperature        31 degrees C / 87 degrees F
  Junction Temperature       40 degrees C / 104 degrees F
  Power
    1.0 V                       1000 mV
    1.5 V                       1499 mV
    1.2 V                       1200 mV
    3.3 V                       3300 mV
    0.9 V                        900 mV
    2.5 V                       2499 mV
    3.3 V bias                  3313 mV
    12.0 V                     12296 mV
    12.0 V i                    2908 mA
SIB 1 status:
  State                      Online
  Exhaust Temperature        31 degrees C / 87 degrees F
  Junction Temperature       40 degrees C / 104 degrees F
  Power
    1.0 V                       1000 mV
    1.5 V                       1499 mV
```

```
    1.2 V                       1200 mV
    3.3 V                       3299 mV
    0.9 V                        900 mV
    2.5 V                       2500 mV
    3.3 V bias                  3313 mV
    12.0 V                     12312 mV
    12.0 V i                    2979 mA
SIB 2 status:
  State                      Online
  Exhaust Temperature        31 degrees C / 87 degrees F
  Junction Temperature       38 degrees C / 100 degrees F
  Power
    1.0 V                        999 mV
    1.5 V                       1499 mV
    1.2 V                       1200 mV
    3.3 V                       3300 mV
    0.9 V                        900 mV
    2.5 V                       2500 mV
    3.3 V bias                  3324 mV
    12.0 V                     12312 mV
    12.0 V i                    2990 mA
SIB 3 status:
  State                      Online
  Exhaust Temperature        31 degrees C / 87 degrees F
  Junction Temperature       39 degrees C / 102 degrees F
  Power
    1.0 V                        999 mV
    1.5 V                       1500 mV
    1.2 V                       1200 mV
    3.3 V                       3300 mV
    0.9 V                        900 mV
    2.5 V                       2499 mV
    3.3 V bias                  3307 mV
    12.0 V                     12296 mV
    12.0 V i                    3144 mA
SIB 4 status:
  State                      Online
  Exhaust Temperature        31 degrees C / 87 degrees F
  Junction Temperature       38 degrees C / 100 degrees F
  Power
    1.0 V                       1000 mV
    1.5 V                       1500 mV
    1.2 V                       1200 mV
```

```
    3.3 V                       3300 mV
    0.9 V                        899 mV
    2.5 V                       2499 mV
    3.3 V bias                  3341 mV
    12.0 V                     12328 mV
    12.0 V i                    2836 mA
SIB 5 status:
  State                       Online
  Exhaust Temperature         31 degrees C / 87 degrees F
  Junction Temperature        38 degrees C / 100 degrees F
  Power
    1.0 V                       1000 mV
    1.5 V                       1499 mV
    1.2 V                       1199 mV
    3.3 V                       3299 mV
    0.9 V                        900 mV
    2.5 V                       2499 mV
    3.3 V bias                  3327 mV
    12.0 V                     12296 mV
    12.0 V i                    2919 mA
SIB 6 status:
  State                       Online
  Exhaust Temperature         31 degrees C / 87 degrees F
  Junction Temperature        38 degrees C / 100 degrees F
  Power
    1.0 V                       1000 mV
    1.5 V                       1500 mV
    1.2 V                       1200 mV
    3.3 V                       3299 mV
    0.9 V                        899 mV
    2.5 V                       2500 mV
    3.3 V bias                  3294 mV
    12.0 V                     12296 mV
    12.0 V i                    2825 mA
SIB 7 status:
  State                       Online
  Exhaust Temperature         30 degrees C / 86 degrees F
  Junction Temperature        36 degrees C / 96 degrees F
  Power
    1.0 V                       1000 mV
    1.5 V                       1500 mV
    1.2 V                       1200 mV
    3.3 V                       3300 mV
```

```
    0.9 V                        899 mV
    2.5 V                       2499 mV
    3.3 V bias                  3316 mV
    12.0 V                     12312 mV
    12.0 V i                    2844 mA
SIB 8 status:
  State                      Online
  Exhaust Temperature        31 degrees C / 87 degrees F
  Junction Temperature       38 degrees C / 100 degrees F
  Power
    1.0 V                       1000 mV
    1.5 V                       1499 mV
    1.2 V                       1199 mV
    3.3 V                       3299 mV
    0.9 V                        900 mV
    2.5 V                       2500 mV
    3.3 V bias                  3333 mV
    12.0 V                     12328 mV
    12.0 V i                    2900 mA
```

show chassis environment sib (PTX5000 Packet Transport Router with SIB-I-8SE)

```
user@host> show chassis environment sib
SIB 0 status:
  State                      Online
  Exhaust Temperature        29 degrees C / 84 degrees F
  Junction Temperature       42 degrees C / 107 degrees F
  Power
    1.0 V                        999 mV
    1.5 V                       1500 mV
    1.2 V                       1200 mV
    3.3 V                       3300 mV
    0.9 V                        900 mV
    2.5 V                       2500 mV
    3.3 V bias                  3289 mV
    12.0 V                     12484 mV
    2.5 V LDO                   2502 mV
    12.0 V i                    3494 mA
SIB 1 status:
  State                      Online
  Exhaust Temperature        28 degrees C / 82 degrees F
```

```
  Junction Temperature       33 degrees C / 91 degrees F
  Power
    1.0 V                       1000 mV
    1.5 V                       1500 mV
    1.2 V                       1199 mV
    3.3 V                       3300 mV
    0.9 V                        899 mV
    2.5 V                       2500 mV
    3.3 V bias                  3313 mV
    12.0 V                     12484 mV
    2.5 V LDO                   2513 mV
    12.0 V i                    3099 mA
SIB 2 status:
  State                        Online
  Exhaust Temperature        29 degrees C / 84 degrees F
  Junction Temperature       39 degrees C / 102 degrees F
  Power
    1.0 V                       1000 mV
    1.5 V                       1500 mV
    1.2 V                       1199 mV
    3.3 V                       3300 mV
    0.9 V                        900 mV
    2.5 V                       2500 mV
    3.3 V bias                  3307 mV
    12.0 V                     12484 mV
    2.5 V LDO                   2512 mV
    12.0 V i                    3336 mA
SIB 3 status:
  State                        Online
  Exhaust Temperature        28 degrees C / 82 degrees F
  Junction Temperature       34 degrees C / 93 degrees F
  Power
    1.0 V                        999 mV
    1.5 V                       1500 mV
    1.2 V                       1200 mV
    3.3 V                       3300 mV
    0.9 V                        900 mV
    2.5 V                       2500 mV
    3.3 V bias                  3299 mV
    12.0 V                     12500 mV
    2.5 V LDO                   2494 mV
    12.0 V i                    3479 mA
SIB 4 status:
```

```
  State                       Online
  Exhaust Temperature         28 degrees C / 82 degrees F
  Junction Temperature        38 degrees C / 100 degrees F
  Power
    1.0 V                        1000 mV
    1.5 V                        1500 mV
    1.2 V                        1199 mV
    3.3 V                        3299 mV
    0.9 V                         900 mV
    2.5 V                        2499 mV
    3.3 V bias                   3309 mV
    12.0 V                      12484 mV
    2.5 V LDO                    2499 mV
    12.0 V i                     3159 mA
SIB 5 status:
  State                       Online
  Exhaust Temperature         27 degrees C / 80 degrees F
  Junction Temperature        33 degrees C / 91 degrees F
  Power
    1.0 V                        1000 mV
    1.5 V                        1500 mV
    1.2 V                        1200 mV
    3.3 V                        3299 mV
    0.9 V                         900 mV
    2.5 V                        2499 mV
    3.3 V bias                   3308 mV
    12.0 V                      12468 mV
    2.5 V LDO                    2495 mV
    12.0 V i                     3054 mA
SIB 6 status:
  State                       Online
  Exhaust Temperature         28 degrees C / 82 degrees F
  Junction Temperature        38 degrees C / 100 degrees F
  Power
    1.0 V                         999 mV
    1.5 V                        1500 mV
    1.2 V                        1199 mV
    3.3 V                        3299 mV
    0.9 V                         899 mV
    2.5 V                        2499 mV
    3.3 V bias                   3313 mV
    12.0 V                      12468 mV
    2.5 V LDO                    2510 mV
```

```
    12.0 V i                    3122 mA
SIB 7 status:
  State                       Online
  Exhaust Temperature         26 degrees C / 78 degrees F
  Junction Temperature        34 degrees C / 93 degrees F
  Power
    1.0 V                        999 mV
    1.5 V                       1499 mV
    1.2 V                       1199 mV
    3.3 V                       3300 mV
    0.9 V                        900 mV
    2.5 V                       2499 mV
    3.3 V bias                  3278 mV
    12.0 V                     12468 mV
    2.5 V LDO                   2504 mV
    12.0 V i                    3234 mA
SIB 8 status:
  State                       Online
  Exhaust Temperature         26 degrees C / 78 degrees F
  Junction Temperature        34 degrees C / 93 degrees F
  Power
    1.0 V                        999 mV
    1.5 V                       1499 mV
    1.2 V                       1199 mV
    3.3 V                       3300 mV
    0.9 V                        900 mV
    2.5 V                       2499 mV
    3.3 V bias                  3278 mV
    12.0 V                     12468 mV
    2.5 V LDO                   2504 mV
    12.0 V i                    3234 mA
```

show chassis environment sib (PTX10008 Router)

```
user@host> show chassis environment sib
SIB 0 status:
  State                       Online
  SIB 0 Intake-A Temp Sensor 37 degrees C / 98 degrees F
  SIB 0 Intake-B Temp Sensor 31 degrees C / 87 degrees F
  SIB 0 Exhaust-A Temp Sensor34 degrees C / 93 degrees F
  SIB 0 Exhaust-B Temp Sensor38 degrees C / 100 degrees F
```

```
  SIB 0 PF0 Temp Sensor      47 degrees C / 116 degrees F
  SIB 0 PF1 Temp Sensor      40 degrees C / 104 degrees F
  Power
         PF0 Core 0.9V                925 mV   28929 mA   26835 mW
         PF0 AVDD 1V                 1000 mV   32656 mA   32718 mW
         PF1 Core 0.9V                922 mV   28351 mA   26210 mW
         PF1 AVDD 1V                 1000 mV   30890 mA   30937 mW
         12V                        12311 mV   10880 mA  134744 mW
SIB 1 status:
  State                      Online
  SIB 1 Intake-A Temp Sensor 43 degrees C / 109 degrees F
  SIB 1 Intake-B Temp Sensor 35 degrees C / 95 degrees F
  SIB 1 Exhaust-A Temp Sensor37 degrees C / 98 degrees F
  SIB 1 Exhaust-B Temp Sensor45 degrees C / 113 degrees F
  SIB 1 PF0 Temp Sensor      54 degrees C / 129 degrees F
  SIB 1 PF1 Temp Sensor      42 degrees C / 107 degrees F
  Power
         PF0 Core 0.9V                951 mV   30468 mA   29007 mW
         PF0 AVDD 1V                  999 mV   32921 mA   32921 mW
         PF1 Core 0.9V                924 mV   29046 mA   26828 mW
         PF1 AVDD 1V                 1000 mV   30796 mA   30796 mW
         12V                        12298 mV   11356 mA  140682 mW
SIB 2 status:
  State                      Online
  SIB 2 Intake-A Temp Sensor 45 degrees C / 113 degrees F
  SIB 2 Intake-B Temp Sensor 36 degrees C / 96 degrees F
  SIB 2 Exhaust-A Temp Sensor38 degrees C / 100 degrees F
  SIB 2 Exhaust-B Temp Sensor46 degrees C / 114 degrees F
  SIB 2 PF0 Temp Sensor      55 degrees C / 131 degrees F
  SIB 2 PF1 Temp Sensor      43 degrees C / 109 degrees F
  Power
         PF0 Core 0.9V                924 mV   28273 mA   26140 mW
         PF0 AVDD 1V                 1000 mV   32562 mA   32625 mW
         PF1 Core 0.9V                925 mV   29351 mA   27265 mW
         PF1 AVDD 1V                  999 mV   31812 mA   31859 mW
         12V                        12351 mV   11010 mA  137481 mW
SIB 3 status:
  State                      Online
  SIB 3 Intake-A Temp Sensor 44 degrees C / 111 degrees F
  SIB 3 Intake-B Temp Sensor 36 degrees C / 96 degrees F
  SIB 3 Exhaust-A Temp Sensor37 degrees C / 98 degrees F
  SIB 3 Exhaust-B Temp Sensor45 degrees C / 113 degrees F
  SIB 3 PF0 Temp Sensor      54 degrees C / 129 degrees F
```

```
  SIB 3 PF1 Temp Sensor      43 degrees C / 109 degrees F
  Power
         PF0 Core 0.9V                924 mV   28820 mA   26695 mW
         PF0 AVDD 1V                  999 mV   32265 mA   32281 mW
         PF1 Core 0.9V                923 mV   28933 mA   26757 mW
         PF1 AVDD 1V                  999 mV   31281 mA   31265 mW
         12V                        12285 mV   11096 mA  137296 mW
SIB 4 status:
  State                      Online
  SIB 4 Intake-A Temp Sensor 45 degrees C / 113 degrees F
  SIB 4 Intake-B Temp Sensor 36 degrees C / 96 degrees F
  SIB 4 Exhaust-A Temp Sensor37 degrees C / 98 degrees F
  SIB 4 Exhaust-B Temp Sensor46 degrees C / 114 degrees F
  SIB 4 PF0 Temp Sensor      54 degrees C / 129 degrees F
  SIB 4 PF1 Temp Sensor      43 degrees C / 109 degrees F
  Power
         PF0 Core 0.9V                950 mV   29757 mA   28281 mW
         PF0 AVDD 1V                  999 mV   32109 mA   32093 mW
         PF1 Core 0.9V                948 mV   29460 mA   27921 mW
         PF1 AVDD 1V                  999 mV   30687 mA   30671 mW
         12V                        12193 mV   11183 mA  137352 mW
SIB 5 status:
  State                      Online
  SIB 5 Intake-A Temp Sensor 38 degrees C / 100 degrees F
  SIB 5 Intake-B Temp Sensor 32 degrees C / 89 degrees F
  SIB 5 Exhaust-A Temp Sensor35 degrees C / 95 degrees F
  SIB 5 Exhaust-B Temp Sensor39 degrees C / 102 degrees F
  SIB 5 PF0 Temp Sensor      45 degrees C / 113 degrees F
  SIB 5 PF1 Temp Sensor      43 degrees C / 109 degrees F
  Power
         PF0 Core 0.9V                951 mV   29546 mA   28156 mW
         PF0 AVDD 1V                  999 mV   31468 mA   31468 mW
         PF1 Core 0.9V                948 mV   30007 mA   28515 mW
         PF1 AVDD 1V                  999 mV   31812 mA   31859 mW
         12V                        12179 mV   11399 mA  139720 mW
```

show chassis environment sib (PTX10016 Router)

```
user@host> show chassis environment sib
SIB 0 status:
  State                      Online
```

```
  SIB 0 Intake-A Temp Sensor 22 degrees C / 71 degrees F
  SIB 0 Intake-B Temp Sensor 21 degrees C / 69 degrees F
  SIB 0 Intake-C Temp Sensor 17 degrees C / 62 degrees F
  SIB 0 Exhaust-A Temp Sensor29 degrees C / 84 degrees F
  SIB 0 Exhaust-B Temp Sensor30 degrees C / 86 degrees F
  SIB 0 Exhaust-C Temp Sensor25 degrees C / 77 degrees F
  SIB 0 PF0 Temp Sensor      31 degrees C / 87 degrees F
  SIB 0 PF1 Temp Sensor      31 degrees C / 87 degrees F
  SIB 0 PF2 Temp Sensor      32 degrees C / 89 degrees F
  SIB 0 PF3 Temp Sensor      33 degrees C / 91 degrees F
  SIB 0 PF4 Temp Sensor      29 degrees C / 84 degrees F
  SIB 0 PF5 Temp Sensor      27 degrees C / 80 degrees F
  Power
         PF0 Core 0.9V                924 mV   24429 mA   22664 mW
         PF0 AVDD 1V                  999 mV   19515 mA   19531 mW
         PF1 Core 0.9V                924 mV   24531 mA   22679 mW
         PF1 AVDD 1V                  999 mV   17796 mA   17812 mW
         PF2 Core 0.9V                924 mV   24308 mA   22503 mW
         PF2 AVDD 1V                  999 mV   16250 mA   16265 mW
         PF3 Core 0.9V                925 mV   24414 mA   22601 mW
         PF3 AVDD 1V                  999 mV   15023 mA   15023 mW
         PF4 Core 0.9V                924 mV   24089 mA   22285 mW
         PF4 AVDD 1V                 1000 mV   15148 mA   15156 mW
         PF5 Core 0.9V                898 mV   23601 mA   21273 mW
         PF5 AVDD 1V                  999 mV   15453 mA   15453 mW
         12V                        12311 mV   11399 mA  141366 mW
         12V_1                      12325 mV   11053 mA  137204 mW
SIB 1 status:
  State                      Online
  SIB 1 Intake-A Temp Sensor 23 degrees C / 73 degrees F
  SIB 1 Intake-B Temp Sensor 23 degrees C / 73 degrees F
  SIB 1 Intake-C Temp Sensor 18 degrees C / 64 degrees F
  SIB 1 Exhaust-A Temp Sensor30 degrees C / 86 degrees F
  SIB 1 Exhaust-B Temp Sensor32 degrees C / 89 degrees F
  SIB 1 Exhaust-C Temp Sensor25 degrees C / 77 degrees F
  SIB 1 PF0 Temp Sensor      33 degrees C / 91 degrees F
  SIB 1 PF1 Temp Sensor      32 degrees C / 89 degrees F
  SIB 1 PF2 Temp Sensor      34 degrees C / 93 degrees F
  SIB 1 PF3 Temp Sensor      39 degrees C / 102 degrees F
  SIB 1 PF4 Temp Sensor      29 degrees C / 84 degrees F
  SIB 1 PF5 Temp Sensor      28 degrees C / 82 degrees F
  Power
         PF0 Core 0.9V                949 mV   24101 mA   22910 mW
```

```
        PF0 AVDD 1V                 1000 mV   19046 mA   19062 mW
        PF1 Core 0.9V                924 mV   24375 mA   22546 mW
        PF1 AVDD 1V                  999 mV   17843 mA   17859 mW
        PF2 Core 0.9V                923 mV   24183 mA   22355 mW
        PF2 AVDD 1V                  999 mV   16109 mA   16117 mW
        PF3 Core 0.9V                949 mV   24246 mA   23097 mW
        PF3 AVDD 1V                 1000 mV   14632 mA   14632 mW
        PF4 Core 0.9V                925 mV   23613 mA   21882 mW
        PF4 AVDD 1V                 1000 mV   14742 mA   14710 mW
        PF5 Core 0.9V                924 mV   24328 mA   22515 mW
        PF5 AVDD 1V                 1000 mV   14640 mA   14664 mW
        12V                        12351 mV   11096 mA  138036 mW
        12V_1                      12219 mV   11226 mA  137500 mW
SIB 2 status:
  State                      Online
  SIB 2 Intake-A Temp Sensor 25 degrees C / 77 degrees F
  SIB 2 Intake-B Temp Sensor 22 degrees C / 71 degrees F
  SIB 2 Intake-C Temp Sensor 18 degrees C / 64 degrees F
  SIB 2 Exhaust-A Temp Sensor29 degrees C / 84 degrees F
  SIB 2 Exhaust-B Temp Sensor33 degrees C / 91 degrees F
  SIB 2 Exhaust-C Temp Sensor24 degrees C / 75 degrees F
  SIB 2 PF0 Temp Sensor      33 degrees C / 91 degrees F
  SIB 2 PF1 Temp Sensor      31 degrees C / 87 degrees F
  SIB 2 PF2 Temp Sensor      34 degrees C / 93 degrees F
  SIB 2 PF3 Temp Sensor      42 degrees C / 107 degrees F
  SIB 2 PF4 Temp Sensor      28 degrees C / 82 degrees F
  SIB 2 PF5 Temp Sensor      27 degrees C / 80 degrees F
  Power
        PF0 Core 0.9V                899 mV   24046 mA   21656 mW
        PF0 AVDD 1V                  999 mV   19265 mA   19250 mW
        PF1 Core 0.9V                900 mV   24234 mA   21867 mW
        PF1 AVDD 1V                  999 mV   18000 mA   18015 mW
        PF2 Core 0.9V                900 mV   23250 mA   20953 mW
        PF2 AVDD 1V                  999 mV   16328 mA   16343 mW
        PF3 Core 0.9V                899 mV   23976 mA   21570 mW
        PF3 AVDD 1V                  999 mV   14976 mA   15007 mW
        PF4 Core 0.9V                924 mV   23718 mA   21976 mW
        PF4 AVDD 1V                 1000 mV   14781 mA   14765 mW
        PF5 Core 0.9V                899 mV   23265 mA   20937 mW
        PF5 AVDD 1V                  999 mV   15125 mA   15132 mW
        12V                        12298 mV   10750 mA  133523 mW
        12V_1                      12245 mV   10880 mA  134041 mW
SIB 3 status:
```

```
  State                      Online
  SIB 3 Intake-A Temp Sensor 23 degrees C / 73 degrees F
  SIB 3 Intake-B Temp Sensor 25 degrees C / 77 degrees F
  SIB 3 Intake-C Temp Sensor 17 degrees C / 62 degrees F
  SIB 3 Exhaust-A Temp Sensor30 degrees C / 86 degrees F
  SIB 3 Exhaust-B Temp Sensor32 degrees C / 89 degrees F
  SIB 3 Exhaust-C Temp Sensor25 degrees C / 77 degrees F
  SIB 3 PF0 Temp Sensor      33 degrees C / 91 degrees F
  SIB 3 PF1 Temp Sensor      31 degrees C / 87 degrees F
  SIB 3 PF2 Temp Sensor      32 degrees C / 89 degrees F
  SIB 3 PF3 Temp Sensor      40 degrees C / 104 degrees F
  SIB 3 PF4 Temp Sensor      28 degrees C / 82 degrees F
  SIB 3 PF5 Temp Sensor      27 degrees C / 80 degrees F
  Power
         PF0 Core 0.9V                924 mV   24558 mA   22734 mW
         PF0 AVDD 1V                  999 mV   19500 mA   19515 mW
         PF1 Core 0.9V                925 mV   24570 mA   22750 mW
         PF1 AVDD 1V                  999 mV   17609 mA   17625 mW
         PF2 Core 0.9V                899 mV   23144 mA   20832 mW
         PF2 AVDD 1V                 1000 mV   16375 mA   16390 mW
         PF3 Core 0.9V                925 mV   24203 mA   22414 mW
         PF3 AVDD 1V                 1000 mV   15039 mA   15023 mW
         PF4 Core 0.9V                899 mV   23523 mA   21183 mW
         PF4 AVDD 1V                  999 mV   15273 mA   15296 mW
         PF5 Core 0.9V                924 mV   24125 mA   22367 mW
         PF5 AVDD 1V                 1000 mV   14953 mA   14968 mW
         12V                        12245 mV   10880 mA  133652 mW
         12V_1                      12259 mV   11053 mA  136464 mW
SIB 4 status:
  State                      Online
  SIB 4 Intake-A Temp Sensor 23 degrees C / 73 degrees F
  SIB 4 Intake-B Temp Sensor 26 degrees C / 78 degrees F
  SIB 4 Intake-C Temp Sensor 18 degrees C / 64 degrees F
  SIB 4 Exhaust-A Temp Sensor30 degrees C / 86 degrees F
  SIB 4 Exhaust-B Temp Sensor33 degrees C / 91 degrees F
  SIB 4 Exhaust-C Temp Sensor24 degrees C / 75 degrees F
  SIB 4 PF0 Temp Sensor      34 degrees C / 93 degrees F
  SIB 4 PF1 Temp Sensor      32 degrees C / 89 degrees F
  SIB 4 PF2 Temp Sensor      33 degrees C / 91 degrees F
  SIB 4 PF3 Temp Sensor      41 degrees C / 105 degrees F
  SIB 4 PF4 Temp Sensor      28 degrees C / 82 degrees F
  SIB 4 PF5 Temp Sensor      27 degrees C / 80 degrees F
  Power
```

```
          PF0 Core 0.9V                925 mV   24644 mA   22824 mW
          PF0 AVDD 1V                  999 mV   19375 mA   19390 mW
          PF1 Core 0.9V                900 mV   24109 mA   21703 mW
          PF1 AVDD 1V                  999 mV   17687 mA   17695 mW
          PF2 Core 0.9V                899 mV   24085 mA   21710 mW
          PF2 AVDD 1V                  999 mV   16578 mA   16570 mW
          PF3 Core 0.9V                949 mV   24652 mA   23410 mW
          PF3 AVDD 1V                 1000 mV   14445 mA   14453 mW
          PF4 Core 0.9V                924 mV   23902 mA   22097 mW
          PF4 AVDD 1V                  999 mV   14750 mA   14742 mW
          PF5 Core 0.9V                925 mV   24082 mA   22308 mW
          PF5 AVDD 1V                  999 mV   14671 mA   14671 mW
          12V                        12338 mV   11139 mA  138277 mW
          12V_1                      12325 mV   11356 mA  140978 mW
SIB 5 status:
  State                      Online
  SIB 5 Intake-A Temp Sensor 22 degrees C / 71 degrees F
  SIB 5 Intake-B Temp Sensor 22 degrees C / 71 degrees F
  SIB 5 Intake-C Temp Sensor 18 degrees C / 64 degrees F
  SIB 5 Exhaust-A Temp Sensor28 degrees C / 82 degrees F
  SIB 5 Exhaust-B Temp Sensor28 degrees C / 82 degrees F
  SIB 5 Exhaust-C Temp Sensor25 degrees C / 77 degrees F
  SIB 5 PF0 Temp Sensor      32 degrees C / 89 degrees F
  SIB 5 PF1 Temp Sensor      31 degrees C / 87 degrees F
  SIB 5 PF2 Temp Sensor      32 degrees C / 89 degrees F
  SIB 5 PF3 Temp Sensor      33 degrees C / 91 degrees F
  SIB 5 PF4 Temp Sensor      29 degrees C / 84 degrees F
  SIB 5 PF5 Temp Sensor      28 degrees C / 82 degrees F
  Power
          PF0 Core 0.9V                924 mV   25093 mA   23210 mW
          PF0 AVDD 1V                  999 mV   19781 mA   19796 mW
          PF1 Core 0.9V                899 mV   24113 mA   21753 mW
          PF1 AVDD 1V                 1000 mV   17968 mA   17984 mW
          PF2 Core 0.9V                925 mV   24218 mA   22437 mW
          PF2 AVDD 1V                 1000 mV   16539 mA   16531 mW
          PF3 Core 0.9V                898 mV   23511 mA   21164 mW
          PF3 AVDD 1V                  999 mV   15015 mA   15023 mW
          PF4 Core 0.9V                975 mV   25328 mA   24718 mW
          PF4 AVDD 1V                 1000 mV   14578 mA   14601 mW
          PF5 Core 0.9V                923 mV   24175 mA   22367 mW
          PF5 AVDD 1V                 1000 mV   14789 mA   14765 mW
```

```
        12V                         12259 mV   11053 mA  136464 mW
        12V_1                       12272 mV   11226 mA  138221 mW
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced in Junos OS Release 9.6. for the TX Matrix Plus router.

RELATED DOCUMENTATION

- request chassis sib | **14562**
- show chassis sibs | **18835**
- Configuring Junos OS to Upgrade and Downgrade Switch Interface Boards on a TX Matrix Router
- M320 SIB Description

show chassis errors active

IN THIS SECTION

- Syntax | **17896**
- Description | **17896**
- Options | **17896**
- Required Privilege Level | **17896**
- Output Fields | **17896**
- Sample Output | **17898**
- show chassis errors active detail (MX-Series Devices with major error) | **17899**
- Release Information | **17900**

Syntax

```
show chassis errors active
<detail>
<fpc-slot slot>
<detail fpc-slot slot>
```

Description

Display the details of chassis errors. In a high availability application, use these commands on primary RE. If you use these commands on secondary RE, the output shows a warning message to run the commands on primary RE.

Options

detail	(Optional) Displays detailed system error information.
fpc-slot *slot*	(Optional) Displays the active errors on the selected FPC slot.
detail fpc-slot *slot*	(Optional) Displays the active errors on the selected FPC slot.

Required Privilege Level

view

Output Fields

Table 243 on page 17896 lists the output fields for the `show chassis errors active detail` command. Output fields are listed in the approximate order in which they appear.

Table 243: Output Fields of `show chassis errors active detail` Command

Field Name	Field Description
`Location`	The FPC number.
`Identifier`	The unique identifier of the error.

Table 243: Output Fields of `show chassis errors active detail` Command *(Continued)*

Field Name	Field Description
`Error`	The error type.
`Scope`	Scope of the error. The following scopes are available: • `board` • `pfe`
`Category`	Error categories include the following: • functional • `io` • `storage` • `memory` • `processing` • `switch`
`Severity`	The severity of the error. It can be configured as follows: • `fatal`—Fatal error on FPC • `major`—Major error on FPC • `minor`—Minor error on FPC
`Details`	Details of the error
`Count`	Number of error instances occurred.
`support`	Possible fix for the error

Sample Output

show chassis errors active (Major error)

```
user@host> show chassis errors active

Slot 0

--------------------------------------------------
Active Minor Errors: 0
Active Major Errors: 1
Active Fatal Errors: 0
```

NOTE:

clear chassis fpc errors error-id <*error_id*>error-id

show chassis errors active detail (MX Series Devices)

```
user@host> show chassis errors active detail
Slot 0

--------------------------------------------------
Location      : FPC 0
Identifier    : /fpc/0/pfe/0/cm/0/BTCHIP/0/BTCHIP_CMERROR_TEMP_MONITOR
Error         : BTCHIP_CMERROR_TEMP_MONITOR
Scope         : pfe
Category      : functional
Severity      : Minor
Details       : TEMP_MONITOR: BT Temperature monitor throttled bandwidth
Count         : 1
Support       : No help info provided
--------------------------------------------------
Location      : FPC 0
Identifier    : /fpc/0/pfe/0/cm/0/BTCHIP/1/BTCHIP_CMERROR_TEMP_MONITOR
Error         : BTCHIP_CMERROR_TEMP_MONITOR
Scope         : pfe
Category      : functional
Severity      : Minor
Details       : TEMP_MONITOR: BT Temperature monitor throttled bandwidth
```

```
Count         : 1
Support       : No help info provided

Slot 1

--------------------------------------------------
Location      : FPC 1
Identifier    : /fpc/1/pfe/0/cm/0/PIC0/0/PIC_CMERROR_STORAGE_SMARTD_ERROR
Error         : PIC_CMERROR_STORAGE_SMARTD_ERROR
Scope         : board
Category      : functional
Severity      : Major
Details       : PIC STORAGE SMARTD ERROR
Count         : 1
Support       : No help info provided
--------------------------------------------------
Location      : FPC 1
Identifier    : /fpc/1/pfe/0/cm/0/PIC1/0/PIC_CMERROR_DDR_CORRECTABLE_MINOR
Error         : PIC_CMERROR_DDR_CORRECTABLE_MINOR
Scope         : board
Category      : functional
Severity      : Minor
Details       : PIC DDR Correctable Minor Error
Count         : 1
Support       : No help info provided
```

show chassis errors active detail (MX-Series Devices with major error)

```
user@host> show chassis errors active detail

Slot 0

--------------------------------------------------
Location      : FPC 0
/fpc/0/pfe/0/cm/0/Eth_Port_Error/0/ETH_CMERROR_MAJOR_0
Error         : ETH_CMERROR_MAJOR_0
Scope         : board
Category      : switch
Oc Category   : default
Severity      : Major
```

```
Details       : eth0 major error
Count         : 1
Support       : No help info provided
--------------------------------------------------
```

show chassis errors active detail fpc-slot

```
user@router> show chassis errors active detail fpc-slot 3

Slot 3

--------------------------------------------------
Location      : FPC 3
Identifier    : /fpc/3/pfe/0/cm/0/XQSS(0)/0/XQSS_CMERROR_CPQW_ERR_INT_FSET_FAST_DEQ_DRY_ERR
Error         : XQSS_CMERROR_CPQW_ERR_INT_FSET_FAST_DEQ_DRY_ERR
Scope         : pfe
Category      : functional
Severity      : Major
Details       : XQSS CPQW FAST DEQ DRY
Count         : 1
Support       : Help for XQSS_CMERROR_CPQW_ERR_INT_FSET_FAST_DEQ_DRY_ERR
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

error | 6441

clear chassis fpc errors | 3468

show chassis ethernet-switch (Security)

IN THIS SECTION

- Syntax | 17901
- Description | 17901
- Required Privilege Level | 17901
- Output Fields | 17901
- Sample Output | 17902
- Release Information | 17907

Syntax

```
show chassis ethernet-switch
<errors>
<statistics>
```

Description

Display information about the ports on the Control Board (CB) Ethernet switch and accumulated Ethernet switch errors on an SRX Series Firewall.

Required Privilege Level

view

Output Fields

Table 244 on page 17902 lists the output fields for the **show chassis ethernet-switch** command. Output fields are listed in the approximate order in which they appear.

Table 244: show chassis ethernet-switch Output Fields

Field Name	Field Description
Link is good on port n connected to device or **Link is good on Fast Ethernet port n connected to device**	Information about the link between each port on the CB's Ethernet switch and one of the following devices: • FPC0 (Flexible PIC Concentrator 0) through FPC7 • Local controller • Routing Engine • Other Routing Engine (on a system with two Routing Engines) • SPMB (Switch Processor Mezzanine Board)
Speed is	Speed at which the Ethernet link is running.
Duplex is	Duplex type of the Ethernet link: **full** or **half**.
Autonegotiate is Enabled (or Disabled)	By default, built-in Fast Ethernet ports on a PIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode.

Sample Output

show chassis ethernet-switch

```
user@host> show chassis ethernet-switch
node0:
--------------------------------------------------------------------------
Displaying summary for switch 0
Link is good on GE port 0 connected to device: FPC0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
```

```
Link is good on GE port 1 connected to device: FPC1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 2 connected to device: FPC2
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 3 connected to device: FPC3
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 4 connected to device: FPC4
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 5 connected to device: FPC5

Link is down on GE port 6 connected to device: FPC6

Link is good on GE port 7 connected to device: FPC7
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 8 connected to device: FPC8
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
```

```
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 9 connected to device: FPC9
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 10 connected to device: FPC10

Link is down on GE port 11 connected to device: FPC11

Link is good on GE port 12 connected to device: Other RE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 13 connected to device: RE-GigE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 14 connected to device: Debug-GigE

node1:
--------------------------------------------------------------------------
Displaying summary for switch 0
Link is good on GE port 0 connected to device: FPC0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 1 connected to device: FPC1
  Speed is 1000Mb
  Duplex is full
```

```
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 2 connected to device: FPC2
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 3 connected to device: FPC3
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 4 connected to device: FPC4
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 5 connected to device: FPC5

Link is down on GE port 6 connected to device: FPC6

Link is good on GE port 7 connected to device: FPC7
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 8 connected to device: FPC8
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
```

```
Link is good on GE port 9 connected to device: FPC9
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 10 connected to device: FPC10

Link is down on GE port 11 connected to device: FPC11

Link is good on GE port 12 connected to device: Other RE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 13 connected to device: RE-GigE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 14 connected to device: Debug-GigE
```

show chassis ethernet-switch errors (SRX5600)

```
user@host> show chassis ethernet-switch errors
Displaying error for switch 0
Error count not available for port 0 connected to device FPC0
Error count not available for port 1 connected to device FPC1

Error count not available for port 2 connected to device FPC2
Error count not available for port 3 connected to device FPC3
Error count not available for port 4 connected to device FPC4
Error count not available for port 5 connected to device FPC5
Error count not available for port 6 connected to device FPC6
Error count not available for port 7 connected to device FPC7
Error count not available for port 8 connected to device FPC8
```

```
Error count not available for port 9 connected to device FPC9
Error count not available for port 10 connected to device FPC10
Error count not available for port 11 connected to device FPC11

Error count not available for port 12 connected to device Other RE
Accumulated error counts for port 13 connected to device RE-GigE:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  0
  Disconnects    0
  FX mode        0
Accumulated error counts for port 14 connected to device Debug-GigE:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  0
  Disconnects    0
  FX mode        0
Error count not available for port 15 connected to device Nothing
Error count not available for port 16 connected to device Nothing
Error count not available for port 17 connected to device Debug-GigE
Error count not available for port 18 connected to device Nothing
Error count not available for port 19 connected to device Nothing
Error count not available for port 20 connected to device Nothing
Error count not available for port 21 connected to device Nothing
Error count not available for port 22 connected to device Nothing
Error count not available for port 23 connected to device Nothing
Error count not available for port 24 connected to device External-Ethernet
Error count not available for port 25 connected to device External-Ethernet
Error count not available for port 26 connected to device External-Ethernet
Error count not available for port 27 connected to device External-Ethernet
```

Release Information

Command introduced in Junos OS Release 9.2.

show chassis ethernet-switch

IN THIS SECTION

- Syntax | 17908
- Syntax (EX8200 Switch) | 17908
- Syntax (SRX Series Firewalls) | 17909
- Syntax (T4000 Router) | 17909
- Syntax (TX Matrix Router) | 17909
- Syntax (TX Matrix Plus Router) | 17909
- Syntax (MX Series Router) | 17909
- Syntax (MX2010, MX2020, and MX2008 Universal Routing Platforms) | 17910
- Syntax (MX10008 Universal Routing Platforms) | 17910
- Syntax (PTX Series Packet Transport Routers) | 17910
- Description | 17910
- Options | 17910
- Required Privilege Level | 17912
- Output Fields | 17912
- Sample Output | 17919
- Release Information | 17960

Syntax

```
show chassis ethernet-switch
<errors <port>>
```

Syntax (EX8200 Switch)

```
show chassis ethernet-switch
<statistics <port> | switch <number>
```

Syntax (SRX Series Firewalls)

```
show chassis ethernet-switch
<errors <port> | statistics <port>>
```

Syntax (T4000 Router)

```
show chassis ethernet-switch
<errors <port> | statistics <port>>
```

Syntax (TX Matrix Router)

```
show chassis ethernet-switch
<errors <port> | statistics <port>>
<lcc <number> | scc>
```

Syntax (TX Matrix Plus Router)

```
show chassis ethernet-switch
<errors <port> | switch <number>
<lcc number | sfc number>
<statistics <port> | switch <number>
```

Syntax (MX Series Router)

```
show chassis ethernet-switch
<all-members>
<errors <port>>
<local>
<member member-id>
```

Syntax (MX2010, MX2020, and MX2008 Universal Routing Platforms)

```
show chassis ethernet-switch
<errors <port> | statistics <port>>
<old-rom-packet-count>
```

Syntax (MX10008 Universal Routing Platforms)

```
statistics <port>>
```

Syntax (PTX Series Packet Transport Routers)

```
show chassis ethernet-switch
<errors <port>>
<statistics <port>>
<port-state  <port>>
```

Description

(M10i, M40e, M120, M160, M320, MX Series, and T Series routers and EX8200 and PTX Series routers only) Display information about the ports on the Control Board (CB) Ethernet switch. Display information about the status of Ethernet switch ports when FPC failures occur and the management Ethernet interface fails on the master RE (MX10008, MX10016, PTX10008, PTX10016, QFX10008, and QFX10016 Switches).

Options

none Display information about each connected port on the Ethernet switch. On a TX Matrix router, display information about each connected port on the Ethernet switch on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display information about each connected port on the Ethernet switch on the TX Matrix Plus router and its attached routers.

all-members (MX Series routers only) (Optional) Display information about the ports on the CB Ethernet switch on all the members of the Virtual Chassis configuration.

errors

(Optional) Display the numbers and types of errors accumulated on all ports of the Ethernet switch.

errors *port*

(Optional) Display the numbers and types of errors accumulated on the specified port (0 through 15) of the Ethernet switch. On the TX Matrix router, replace *port* with a value from 0 through 15. On the TX Matrix Plus router and EX8200 switch, replace *port* with a value from 0 through 27. On the PTX Series Packet Transport Routers, replace *port* with a value from 0 through 25. On the T4000 routers, MX2020 routers, MX2010 routers, and MX2008 routers, replace *port* with a value from 0 through 27.

errors switch *number*

(TX Matrix Plus router only) (Optional) Display the numbers and types of errors accumulated on the specified switch. Replace *number* with a value from 0 through 2.

lcc *number*

(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(MX Series routers only) (Optional) Display information about the ports on the CB Ethernet switch on the local Virtual Chassis member.

member *member-id*

(MX Series routers only) (Optional) Display information about the ports on the CB Ethernet switch on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

old-rom-packet-count

(MX 2020 Routers only) (Optional) Display information about installed linecards. A non-zero number indicates that the bootrom on that linecard needs to be updated.

port-state

(PTX Series only) (Optional) Display information about current port operation (`Blocking`, `Listening`, or `Disabled`).

scc

(TX Matrix router only) (Optional) Display information about the ports on the CB's Ethernet switch on the TX Matrix router (switch-card chassis).

sfc *number* (TX Matrix Plus router only) (Optional) Display information about the ports on the CB's Ethernet switch on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

statistics (Optional) Display traffic statistics for each connected port on the Ethernet switch.

statistics *port* (Optional) Display traffic statistics for the specified port on the Ethernet switch. On the TX Matrix router, replace *port* with a value from 0 through 25. On the TX Matrix Plus router or EX8200 switch, replace *port* with a value from 0 through 27. On the PTX Series Packet Transport Routers, replace *port* with a value from 0 through 25. On the T4000 routers, MX2020 routers, MX2010 routers, and MX2008 routers, replace *port* with a value from 0 through 27.

statistics switch *number* (TX Matrix Plus routers and EX8200 switch only) (Optional) Display traffic statistics for the specified Ethernet switch number. On the TX Matrix Plus router and EX8216 switch, replace *number* with a value from 0 through 2. On the EX8208 switch, replace *number* with a value from 0 through 1.

Required Privilege Level

view

Output Fields

Table 245 on page 17913 lists the output fields for the `show chassis ethernet-switch` command. Output fields are listed in the approximate order in which they appear.

Table 245: show chassis ethernet-switch Output Fields

Field Name	Field Description

`Link is good on port n connected to device`
or
`Link is good on Fast Ethernet port n connected to device`
or
`Link is good on Gigabit Ethernet port n connected to device`
or
`Link is down on Gigabit Ethernet port connected to device`
	Information about the link between each port on the CB's Ethernet switch and one of the following devices:
	• FPC0 (Flexible PIC Concentrator 0) through FPC7
	• Local controller
	• Routing Engine
	• Other Routing Engine (on a system with two Routing Engines)
	• SPMB (Switch Processor Mezzanine Board)
	• (TX Matrix router only) LCC0 (line-card chassis 0) through LCC3

`Speed is`	Speed at which the Ethernet link is running: `10 Mb` or `100 Mb`. When the device is `RE` or `Other RE` on the TX Matrix router, the speed is `1000 Mb`. **NOTE**: Irrespective of the device, the speed is `1000` Mb on the MX2010, MX2020, and MX2008 routers.
`Duplex is`	Duplex type of the Ethernet link: `full` or `half`.
`Autonegotiate is Enabled (or Disabled)`	By default, built-in Fast Ethernet ports on a PIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the `no-concatenate` statement at the `[edit chassis]` hierarchy level, as described in the *Junos OS System Basics Configuration Guide*).

Table 245: show chassis ethernet-switch Output Fields *(Continued)*

Field Name	Field Description
`Flow Control TX is Enabled (or Disabled)`	(MX2010 routers, MX2020 routers, MX2008 routers, and PTX Series) Flow control in the transmit direction is enabled (or disabled). Flow control regulates the flow of packets from the switch to the remote side of the connection.
`Flow Control RX is Enabled (or Disabled)`	(MX2010 routers, MX2020 routers, MX2008 routers, and PTX Series) Flow control in the receive direction is enabled (or disabled). Flow control regulates the flow of packets from the remote side of the connection to the switch.
`MLT3`	Number of multilevel threshold-3 (MLT-3) Fast Ethernet errors detected.
Accumulated error counts for port *n* connected to device FPC*n*: (error output only)	
`Lock`	Number of lock errors detected.
`Xmit`	Number of transmission errors detected.
`ESD`	Number of electrostatic discharge (ESD) errors detected.
`False Carrier`	Number of false carrier errors detected. This number is increased by one if a FRU is removed.
`Disconnects`	Number of disconnect errors detected.
`FX mode`	Number of errors detected on an Ethernet link over optical fiber.
Statistics for port *n* connected to device FPC*n* (statistics output only)	
`TX Packets 64 Octets`	(MX2010, MX2020, and MX2008 routers) Number of packets of size 64 octets transmitted.
`TX Packets 65 - 127 Octets`	(MX2010, MX2020, and MX2008 routers) Number of packets of size 65 through 127 octets transmitted.

Table 245: show chassis ethernet-switch Output Fields *(Continued)*

Field Name	Field Description
TX Packets 128 - 255 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 128 through 255 octets transmitted.
TX Packets 256 - 511 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 256 through 511 octets transmitted.
TX Packets 512 - 1023 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 512 through 1023 octets transmitted.
TX Packets 1024 - 1518 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 1024 through 1518 octets transmitted.
TX Packets 1519 - 2047 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 1519 through 2047 octets transmitted.
TX Packets 2048 - 4095 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 2048 through 4095 octets transmitted.
TX Packets 4096 - 9216 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 4096 through 9216 octets transmitted.
TX 1519 - 1522 Good Vlan frms	(MX2010, MX2020, and MX2008 routers) Number of transmitted frames of size 1519 through 1522 octets that are good VLAN frames.
TX Octets	Number of octets sent.
TX Unicast packets	Number of unicast packets sent.
TX Multicast packets	Number of multicast packets sent.
TX Broadcast packets	Number of broadcast packets sent.

Table 245: show chassis ethernet-switch Output Fields *(Continued)*

Field Name	Field Description
`TX Single Collision frames`	(MX2010, MX2020, and MX2008 routers) Number of packets sent after one collision.
`TX Mult. Collision frames`	(MX2010, MX2020, and MX2008 routers) Number of packets sent after multiple collisions.
`TX Late collisions`	Number of packets aborted during sending because of collisions after 64 bytes.
`TX Excessive collisions`	Number of packets not sent because of too many collisions.
`TX Dropped packets`	Number of transmitted packets that were dropped.
`TX PAUSEMAC Ctrl Frames`	Number of Media Access Control (MAC) frames containing PAUSE commands that were sent.
`TX Oversize Packets`	Number of oversize packets that were sent.
`TX FCS Error Counter`	Number of packets discarded because of frame check sequence errors.
`TX Fragment Counter`	Number of fragmented packets sent.
`TX Byte Counter`	Number of bytes sent.
`TX Packet OK Counter`	Number of viable packets sent.
`TX Pause Packet Counter`	Number of PAUSE packets sent.
`RX Packets 64 Octets`	(MX2010, MX2020, and MX2008 routers) Number of packets of size 64 octets received.

Table 245: show chassis ethernet-switch Output Fields *(Continued)*

Field Name	Field Description
RX Packets 65 - 127 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 65 through 127 octets received.
RX Packets 128 - 255 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 128 through 255 octets received.
RX Packets 256 - 511 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 256 through 511 octets received.
RX Packets 512 - 1023 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 512 through 1023 octets received.
RX Packets 1024 - 1518 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 1024 through 1518 octets received.
RX Packets 1519 - 2047 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 1519 through 2047 octets received.
RX Packets 2048 - 4095 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 2048 through 4095 octets received.
RX Packets 4096 - 9216 Octets	(MX2010, MX2020, and MX2008 routers) Number of packets of size 4096 through 9216 octets received.
RX Octets	Number of octets received.
RX Unicast packets	Number of unicast packets received.
RX Multicast packets	Number of multicast packets received.
RX Broadcast packets	Number of broadcast packets received.

Table 245: show chassis ethernet-switch Output Fields *(Continued)*

Field Name	Field Description
RX FCS Errors	Number of packets discarded because of frame check sequence errors.
RX Alignment Errors	Number of incomplete octets received.
RX Dropped Packets	Number of incoming packets that were dropped.
RX Fragments	Number of fragmented packets received.
RX Symbol Errors	Number of symbols received that the router did not correctly decode.
RX MAC Control	Number of Media Access Control (MAC) packets received.
RX Oversize Packets	Number of oversize packets received.
RX Undersize Packets	Number of undersize packets received.
RX Jabbers	Total number of frames received that exceed the maximum byte count and contain CRC errors .
RX Control Frame Counter	Number of control frames received.
RX Pause Frame Counter	Number of pause frames received.
RX Byte Counter	Number of bytes received.
RX Packet OK Counter	Number of viable packets received.

Sample Output

show chassis ethernet-switch

```
user@host> show chassis ethernet-switch
Link is good on port 0 connected to device: FPC0
  Speed is 100 MB
  Duplex is full

Link is good on port 1 connected to device: FPC1
  Speed is 100 MB
  Duplex is full

Link is good on port 2 connected to device: FPC2
  Speed is 100 MB
  Duplex is full

Link is good on port 3 connected to device: FPC3
  Speed is 100 MBb
  Duplex is full

Link is good on port 7 connected to device: Local controller
  Speed is 100 MB
  Duplex is full

Link is good on port 9 connected to device: SPMB
  Speed is 100 MB
  Duplex is full

Link is good on port 13 connected to device: FPC5
  Speed is 100 MB
  Duplex is full
```

show chassis ethernet-switch (MX2020 Router)

```
user@host > show chassis ethernet-switch
Displaying summary for switch 0
Link is good on GE port 0 connected to device: FPC0
  Speed is 1000Mb
```

```
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 1 connected to device: FPC1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 2 connected to device: FPC3
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 3 connected to device: FPC2
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 4 connected to device: FPC5
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 5 connected to device: FPC4
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 6 connected to device: FPC6
  Speed is 1000Mb
  Duplex is full
```

```
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 7 connected to device: FPC7
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 8 connected to device: FPC8
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 9 connected to device: FPC9
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 10 connected to device: FPC10
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 11 connected to device: FPC11
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 12 connected to device: FPC13
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
```

```
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 13 connected to device: FPC12
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 14 connected to device: FPC14
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 15 connected to device: FPC15
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 16 connected to device: FPC17
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 17 connected to device: FPC16
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 18 connected to device: FPC18
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
```

```
  Flow Control RX is Disabled

Link is good on GE port 19 connected to device: FPC19
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 20 connected to device: Other RE-GigE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 21 connected to device: RE-GigE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 22 connected to device: Debug-GigE

Link is good on GE port 23 connected to device: SPMB
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 24 connected to device: SFP+ 0

Link is down on XE port 25 connected to device: SFP+ 1

Link is down on XE port 26 connected to device: RE-10GigE

Link is down on XE port 27 connected to device: Other RE-10GigE
```

show chassis ethernet-switch statistics (MX2020 Router)

```
user@host > show chassis ethernet-switch statistics
Displaying port statistics for switch 0
Statistics for port 0 connected to device FPC0:
  TX Packets 64 Octets        1468564
  TX Packets 65-127 Octets    153896
  TX Packets 128-255 Octets   237
  TX Packets 256-511 Octets   286
  TX Packets 512-1023 Octets  599
  TX Packets 1024-1518 Octets  22803
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   1646385
  TX Multicast Packets        6
  TX Broadcast Packets        970939
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             130470290
  RX Packets 64 Octets        180266
  RX Packets 65-127 Octets    519030
  RX Packets 128-255 Octets   1390
  RX Packets 256-511 Octets   42857
  RX Packets 512-1023 Octets  3482
  RX Packets 1024-1518 Octets  8147
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   755172
```

```
  RX Multicast Packets        0
  RX Broadcast Packets        42822
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             75374021
Statistics for port 1 connected to device FPC1:
  TX Packets 64 Octets        1493739
  TX Packets 65-127 Octets    126996
  TX Packets 128-255 Octets   241
  TX Packets 256-511 Octets   283
  TX Packets 512-1023 Octets  604
  TX Packets 1024-1518 Octets  33687
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   1655550
  TX Multicast Packets        6
  TX Broadcast Packets        969032
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
```

```
  TX Fragment Counter         0
  TX Byte Counter             141832690
  RX Packets 64 Octets        155655
  RX Packets 65-127 Octets    545561
  RX Packets 128-255 Octets   1394
  RX Packets 256-511 Octets   42811
  RX Packets 512-1023 Octets  3514
  RX Packets 1024-1518 Octets  8171
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   757106
  RX Multicast Packets        0
  RX Broadcast Packets        44509
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             75691392
Statistics for port 2 connected to device FPC3:
  TX Packets 64 Octets        1465749
  TX Packets 65-127 Octets    152849
  TX Packets 128-255 Octets   238
  TX Packets 256-511 Octets   289
  TX Packets 512-1023 Octets  602
  TX Packets 1024-1518 Octets  38903
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   1658630
  TX Multicast Packets        6
  TX Broadcast Packets        968873
```

```
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             147427010
  RX Packets 64 Octets        181636
  RX Packets 65-127 Octets    517526
  RX Packets 128-255 Octets   1405
  RX Packets 256-511 Octets   42806
  RX Packets 512-1023 Octets  3515
  RX Packets 1024-1518 Octets  8168
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   755056
  RX Multicast Packets        0
  RX Broadcast Packets        44490
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             75381869
Statistics for port 3 connected to device FPC2:
  TX Packets 64 Octets        1473828
```

```
TX Packets 65-127 Octets    145643
TX Packets 128-255 Octets   253
TX Packets 256-511 Octets   285
TX Packets 512-1023 Octets  612
TX Packets 1024-1518 Octets  26603
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan frms  0
TX Octets                   1647224
TX Multicast Packets        6
TX Broadcast Packets        968925
TX Single Collision frames  0
TX Mult. Collision frames   0
TX Late Collisions          0
TX Excessive Collisions     0
TX Collision frames         0
TX PAUSEMAC Ctrl Frames     0
TX MAC ctrl frames          0
TX Frame deferred Xmns      0
TX Frame excessive deferl   0
TX Oversize Packets         0
TX Jabbers                  0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             134293832
RX Packets 64 Octets        174230
RX Packets 65-127 Octets    525756
RX Packets 128-255 Octets   1404
RX Packets 256-511 Octets   42815
RX Packets 512-1023 Octets  3530
RX Packets 1024-1518 Octets  8176
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                   755911
RX Multicast Packets        0
RX Broadcast Packets        44499
RX FCS Errors               0
RX Align Errors             0
RX Fragments                0
RX Symbol errors            0
RX Unsupported opcodes      0
```

```
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             75517355
Statistics for port 4 connected to device FPC5:
  TX Packets 64 Octets        1466664
  TX Packets 65-127 Octets    151155
  TX Packets 128-255 Octets   238
  TX Packets 256-511 Octets   277
  TX Packets 512-1023 Octets  615
  TX Packets 1024-1518 Octets  54674
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   1673623
  TX Multicast Packets        6
  TX Broadcast Packets        968610
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             164247790
  RX Packets 64 Octets        180006
  RX Packets 65-127 Octets    518217
  RX Packets 128-255 Octets   1406
  RX Packets 256-511 Octets   42787
  RX Packets 512-1023 Octets  3515
```

```
  RX Packets 1024-1518 Octets  8164
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                    754095
  RX Multicast Packets         0
  RX Broadcast Packets         44457
  RX FCS Errors                0
  RX Align Errors              0
  RX Fragments                 0
  RX Symbol errors             0
  RX Unsupported opcodes       0
  RX Out of Range Length       0
  RX False Carrier Errors      0
  RX Undersize Packets         0
  RX Oversize Packets          0
  RX Jabbers                   0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter        0
  RX Control Frame Counter     0
  RX Pause Frame Counter       0
  RX Byte Counter              75311970
Statistics for port 5 connected to device FPC4:
  TX Packets 64 Octets         1464770
  TX Packets 65-127 Octets     154498
  TX Packets 128-255 Octets    225
  TX Packets 256-511 Octets    280
  TX Packets 512-1023 Octets   637
  TX Packets 1024-1518 Octets  26355
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                    1646765
  TX Multicast Packets         6
  TX Broadcast Packets         968730
  TX Single Collision frames   0
  TX Mult. Collision frames    0
  TX Late Collisions           0
  TX Excessive Collisions      0
  TX Collision frames          0
  TX PAUSEMAC Ctrl Frames      0
  TX MAC ctrl frames           0
```

```
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             134058606
  RX Packets 64 Octets        169269
  RX Packets 65-127 Octets    515285
  RX Packets 128-255 Octets   1527
  RX Packets 256-511 Octets   42804
  RX Packets 512-1023 Octets  3521
  RX Packets 1024-1518 Octets  9142
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   741548
  RX Multicast Packets        0
  RX Broadcast Packets        44470
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             75498393
Statistics for port 6 connected to device FPC6:
  TX Packets 64 Octets        1475260
  TX Packets 65-127 Octets    143324
  TX Packets 128-255 Octets   260
  TX Packets 256-511 Octets   274
  TX Packets 512-1023 Octets  603
  TX Packets 1024-1518 Octets  40631
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
```

```
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan frms  0
TX Octets                    1660352
TX Multicast Packets         6
TX Broadcast Packets         968466
TX Single Collision frames   0
TX Mult. Collision frames    0
TX Late Collisions           0
TX Excessive Collisions      0
TX Collision frames          0
TX PAUSEMAC Ctrl Frames      0
TX MAC ctrl frames           0
TX Frame deferred Xmns       0
TX Frame excessive deferl    0
TX Oversize Packets          0
TX Jabbers                   0
TX FCS Error Counter         0
TX Fragment Counter          0
TX Byte Counter              149212764
RX Packets 64 Octets         172275
RX Packets 65-127 Octets     526519
RX Packets 128-255 Octets    1394
RX Packets 256-511 Octets    42777
RX Packets 512-1023 Octets   3514
RX Packets 1024-1518 Octets  8161
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                    754640
RX Multicast Packets         0
RX Broadcast Packets         44443
RX FCS Errors                0
RX Align Errors              0
RX Fragments                 0
RX Symbol errors             0
RX Unsupported opcodes       0
RX Out of Range Length       0
RX False Carrier Errors      0
RX Undersize Packets         0
RX Oversize Packets          0
RX Jabbers                   0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter        0
```

```
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             75386517
Statistics for port 7 connected to device FPC7:
  TX Packets 64 Octets        1472361
  TX Packets 65-127 Octets    145646
  TX Packets 128-255 Octets   251
  TX Packets 256-511 Octets   250
  TX Packets 512-1023 Octets  580
  TX Packets 1024-1518 Octets  49530
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   1668618
  TX Multicast Packets        6
  TX Broadcast Packets        968317
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             158689814
  RX Packets 64 Octets        174618
  RX Packets 65-127 Octets    523421
  RX Packets 128-255 Octets   1393
  RX Packets 256-511 Octets   42764
  RX Packets 512-1023 Octets  3514
  RX Packets 1024-1518 Octets  8158
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   753868
  RX Multicast Packets        0
  RX Broadcast Packets        44429
```

```
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             75309863
Statistics for port 8 connected to device FPC8:
...
```

show chassis ethernet-switch (MX2020 Router with MPC4E)

```
user@ host > show chassis ethernet-switch
Displaying summary for switch 0
Link is good on GE port 0 connected to device: FPC0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 1 connected to device: FPC1

Link is down on GE port 2 connected to device: FPC3

Link is down on GE port 3 connected to device: FPC2

Link is down on GE port 4 connected to device: FPC5

Link is down on GE port 5 connected to device: FPC4

Link is down on GE port 6 connected to device: FPC6
```

```
Link is down on GE port 7 connected to device: FPC7

Link is down on GE port 8 connected to device: FPC8

Link is good on GE port 9 connected to device: FPC9
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 10 connected to device: FPC10
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 11 connected to device: FPC11

Link is down on GE port 12 connected to device: FPC13

Link is down on GE port 13 connected to device: FPC12

Link is good on GE port 14 connected to device: FPC14
  Speed is 1000Mb
  Duplex is full
Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 15 connected to device: FPC15

Link is down on GE port 16 connected to device: FPC17

Link is down on GE port 17 connected to device: FPC16

Link is down on GE port 18 connected to device: FPC18

Link is good on GE port 19 connected to device: FPC19
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
```

```
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 20 connected to device: Other RE-GigE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 21 connected to device: RE-GigE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 22 connected to device: Debug-GigE

Link is good on GE port 23 connected to device: SPMB
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 24 connected to device: SFP+ 0

Link is down on XE port 25 connected to device: SFP+ 1

Link is down on XE port 26 connected to device: RE-10GigE

Link is down on XE port 27 connected to device: Other RE-10GigE
```

show chassis ethernet-switch (MX2008 Router)

```
user@host> show chassis ethernet-switch
Displaying summary for switch 0
Link is good on GE port 0 connected to device: FPC0
  Speed is 1000Mb
  Duplex is full
```

```
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 1 connected to device: FPC1

Link is good on GE port 2 connected to device: FPC3
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 3 connected to device: FPC2

Link is good on GE port 4 connected to device: FPC5
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 5 connected to device: FPC4

Link is down on GE port 6 connected to device: FPC6

Link is good on GE port 7 connected to device: FPC7
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 8 connected to device: FPC8

Link is good on GE port 9 connected to device: FPC9
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 20 connected to device: CB-to-CB-GigE 1
```

```
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 21 connected to device: CB-to-CB-GigE 2
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 22 connected to device: (null)

Link is down on GE port 23 connected to device: (null)

Link is good on XE port 24 connected to device: Other RE-10GigE
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 25 connected to device: RE-10GigE
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 26 connected to device: SFP+ 0

Link is down on XE port 27 connected to device: SFP+ 1
```

show chassis ethernet-switch statistics (Mx10008 Router)

```
user@host> show chassis ethernet-switch statistics
Switch Status: Online
Link is Disabled on port connected to QPHY_0
Link is Disabled on port connected to QPHY_1
```

```
Link is Down on port connected to PTP_FPGA
Link is Disabled on port connected to Unused
Link is Up on port connected to LC0
        Speed      : 10G
        Duplexity  : FD
        Autoneg    : No
        tx_packets : 2835539
        rx_packets : 2624197
        tx_errors  : 0
        rx_errors  : 0
Link is Down on port connected to LC1
Link is Up on port connected to LC2
        Speed      : 10G
        Duplexity  : FD
        Autoneg    : No
        tx_packets : 2889426
        rx_packets : 2441270
        tx_errors  : 0
        rx_errors  : 0
Link is Up on port connected to LC3
        Speed      : 10G
        Duplexity  : FD
        Autoneg    : No
        tx_packets : 2776323
        rx_packets : 2322320
        tx_errors  : 0
        rx_errors  : 0
Link is Disabled on port connected to LC8
Link is Down on port connected to LC4
Link is Disabled on port connected to LC12
Link is Disabled on port connected to LC9
Link is Down on port connected to LC5
Link is Disabled on port connected to LC13
Link is Disabled on port connected to LC10
Link is Down on port connected to LC6
Link is Disabled on port connected to LC14
Link is Disabled on port connected to LC11
Link is Down on port connected to LC7
Link is Disabled on port connected to LC15
Link is Disabled on port connected to OCB_SW
Link is Disabled on port connected to Unused
Link is Disabled on port connected to Fortville_1
Link is Up on port connected to Fortville_0
```

```
        Speed      : 10G
        Duplexity  : FD
        Autoneg    : Yes
        tx_packets : 7387765
        rx_packets : 8348292
        tx_errors  : 0
        rx_errors  : 0
```

show chassis ethernet-switch (TX Matrix Router)

```
user@host> show chassis ethernet-switch
scc-re0:
--------------------------------------------------------------------------
Link is good on FE port 4 connected to device: LCC0
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled

Link is good on FE port 6 connected to device: LCC2
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled

Link is good on FE port 8 connected to device: SPMB
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled

lcc0-re0:
--------------------------------------------------------------------------
Link is good on FE port 1 connected to device: FPC1
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled

Link is good on FE port 2 connected to device: FPC2
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled

Link is good on FE port 8 connected to device: SPMB
```

```
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled

Link is good on FE port 10 connected to device: SCC
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled

lcc2-re0:
--------------------------------------------------------------------------
Link is good on FE port 0 connected to device: FPC0
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled

Link is good on FE port 1 connected to device: FPC1
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled

Link is good on FE port 2 connected to device: FPC2
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled

Link is good on FE port 8 connected to device: SPMB
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled

Link is good on FE port 10 connected to device: SCC
  Speed is 100 MB
  Duplex is full
  Autonegotiate is Enabled
```

show chassis ethernet-switch errors

```
user@host> show chassis ethernet-switch errors
Accumulated error counts for port 0 connected to device FPC0:
  MLT3           2
```

```
  Lock           0
  Xmit           0
  ESD            0
  False carrier  2
  Disconnects    0
  FX mode        0
Accumulated error counts for port 1 connected to device FPC1:
  MLT3           2
  Lock           0
  Xmit           0
  ESD            0
  False carrier  2
  Disconnects    0
  FX mode        0
Accumulated error counts for port 2 connected to device FPC2:
  MLT3           2
  Lock           0
  Xmit           0
  ESD            0
  False carrier  3
  Disconnects    0
  FX mode        0
Accumulated error counts for port 3 connected to device FPC3:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  0
  Disconnects    0
Accumulated error counts for port 4 connected to device Nothing:
   MLT3           0
   Lock           0
   Xmit           0
   ESD            0
   False carrier  0
   Disconnects    0
   FX mode        0
...
```

show chassis ethernet-switch errors (TX Matrix Plus Router)

```
user@host> show chassis ethernet-switch errors
sfc0-re0:
--------------------------------------------------------------------------
Displaying error for switch 0

Displaying error for switch 1
Accumulated error counts for port 0 connected to device LCC0:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  0
  Disconnects    0
  FX mode        0

lcc0-re0:
--------------------------------------------------------------------------
Displaying error for switch 0
Accumulated error counts for port 6 connected to device FPC0:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  5
  Disconnects    0
  FX mode        0
Accumulated error counts for port 7 connected to device FPC1:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  7
  Disconnects    0
  FX mode        0
Accumulated error counts for port 19 connected to device Other RE:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  0
```

```
  Disconnects    0
  FX mode        0
Accumulated error counts for port 20 connected to device SFC0:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  0
  Disconnects    0
  FX mode        0
```

show chassis ethernet-switch sfc errors (TX Matrix Plus Router)

```
user@host> show chassis ethernet-switch errors
switch sfc
sfc0-re0:
--------------------------------------------------------------------------
Displaying error for switch 1
Accumulated error counts for port 0 connected to device LCC0:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  0
  Disconnects    0
  FX mode        0
Accumulated error counts for port 2 connected to device LCC1:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  0
  Disconnects    0
  FX mode        0
Accumulated error counts for port 4 connected to device LCC2:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  0
  Disconnects    0
```

```
  FX mode        0
Accumulated error counts for port 6 connected to device LCC3:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  0
  Disconnects    0
  FX mode        0

lcc0-re0:
--------------------------------------------------------------------------
error: command is not valid on the t1600

lcc1-re0:
--------------------------------------------------------------------------
error: command is not valid on the t1600

lcc2-re0:
--------------------------------------------------------------------------
error: command is not valid on the t1600

lcc3-re0:
--------------------------------------------------------------------------
error: command is not valid on the t1600
```

show chassis ethernet-switch statistics (TX Matrix Plus Router)

```
user@host> show chassis ethernet-switch statistics
sfc0-re0:
--------------------------------------------------------------------------
Displaying port statistics for switch 0
Statistics for port 1 connected to device 1GSW:
  TX Packets 64 Octets        5183577
  TX Packets 65-127 Octets    67820
  TX Packets 128-255 Octets   772
  TX Packets 256-511 Octets   136
  TX Packets 512-1023 Octets  68
  TX Packets 1024-1518 Octets  10881
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
```

```
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   5263254
  TX Multicast Packets        16
  TX Broadcast Packets        723403
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             349922253
  TX Packet OK Counter        5263254
  TX Pause Packet Counter     0
  TX Unicast Counter          4539835
  RX Packets 64 Octets        6513629
  RX Packets 65-127 Octets    88761
  RX Packets 128-255 Octets   6382
  RX Packets 256-511 Octets   22027
  RX Packets 512-1023 Octets  4319
  RX Packets 1024-1518 Octets  49922
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   6685040
  RX Multicast Packets        4
  RX Broadcast Packets        2137376
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             509224602
  RX Unicast Frame Count      4547660
  RX Packet OK Count          6685040
Statistics for port 9 connected to device RE1:
  TX Packets 64 Octets        2500318
  TX Packets 65-127 Octets    443
  TX Packets 128-255 Octets   0
  TX Packets 256-511 Octets   0
```

```
  TX Packets 512-1023 Octets  0
  TX Packets 1024-1518 Octets  0
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   2500761
  TX Multicast Packets        4
  TX Broadcast Packets        2500757
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             160049670
  TX Packet OK Counter        0
  TX Pause Packet Counter     0
  TX Unicast Counter          0
  RX Packets 64 Octets        701191
  RX Packets 65-127 Octets    5882
  RX Packets 128-255 Octets   2
  RX Packets 256-511 Octets   0
  RX Packets 512-1023 Octets  17965
  RX Packets 1024-1518 Octets  7
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   725047
  RX Multicast Packets        8
  RX Broadcast Packets        2500757
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             62402656
  RX Unicast Frame Count      0
  RX Packet OK Count          0
Statistics for port 17 connected to device RE0:
```

```
TX Packets 64 Octets        7214818
TX Packets 65-127 Octets    94640
TX Packets 128-255 Octets   6384
TX Packets 256-511 Octets   22027
TX Packets 512-1023 Octets  22284
TX Packets 1024-1518 Octets  49929
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX Packets 9217-16383 Octets  0
TX Octets                   7410082
TX Multicast Packets        12
TX Broadcast Packets        2497247
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets         0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             571626932
TX Packet OK Counter        0
TX Pause Packet Counter     0
TX Unicast Counter          0
RX Packets 64 Octets        4823701
RX Packets 65-127 Octets    67812
RX Packets 128-255 Octets   772
RX Packets 256-511 Octets   136
RX Packets 512-1023 Octets  68
RX Packets 1024-1518 Octets  10881
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Packets 9217-16383 Octets  0
RX Octets                   4903370
RX Multicast Packets        8
RX Broadcast Packets        2497247
RX FCS Errors               0
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0
```

```
  RX Byte Counter             326889517
  RX Unicast Frame Count      0
  RX Packet OK Count          0

Displaying port statistics for switch 1
Statistics for port 0 connected to device LCC0:
  TX Packets 64 Octets        5053443
  TX Packets 65-127 Octets    59737
  TX Packets 128-255 Octets   768
  TX Packets 256-511 Octets   87
  TX Packets 512-1023 Octets  68
  TX Packets 1024-1518 Octets  85
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   5114188
  TX Multicast Packets        16
  TX Broadcast Packets        1125742
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             329291449
  RX Packets 64 Octets        5640175
  RX Packets 65-127 Octets    79875
  RX Packets 128-255 Octets   6338
  RX Packets 256-511 Octets   165
  RX Packets 512-1023 Octets  4317
  RX Packets 1024-1518 Octets  10
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   5730880
```

```
  RX Multicast Packets        4
  RX Broadcast Packets        1735007
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
RX Pause Frame Counter      0
  RX Byte Counter             371282850
Statistics for port 18 connected to device SPMB:
  TX Packets 64 Octets        2990326
  TX Packets 65-127 Octets    8572
  TX Packets 128-255 Octets   4
  TX Packets 256-511 Octets   49
  TX Packets 512-1023 Octets  0
  TX Packets 1024-1518 Octets  10793
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   3009744
  TX Multicast Packets        20
  TX Broadcast Packets        2458322
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
```

```
  TX Fragment Counter         0
  TX Byte Counter             203712524
  RX Packets 64 Octets        873454
  RX Packets 65-127 Octets    8886
  RX Packets 128-255 Octets   44
  RX Packets 256-511 Octets   21862
  RX Packets 512-1023 Octets  2
  RX Packets 1024-1518 Octets  49912
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   954160
  RX Multicast Packets        0
  RX Broadcast Packets        402369
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             137941752
...
```

show chassis ethernet-switch (T4000 Router)

```
user@host> show chassis ethernet-switch
Displaying summary for switch 0
Link is good on GE port 6 connected to device: FPC0
  Speed is 100Mb
  Duplex is full
  Autonegotiate is Enabled
  False carrier sense count = 04
```

```
Link is good on GE port 9 connected to device: FPC3
  Speed is 100Mb
  Duplex is full
  Autonegotiate is Enabled
  False carrier sense count = 03

Link is good on GE port 11 connected to device: FPC5
  Speed is 100Mb
  Duplex is full
  Autonegotiate is Enabled
  False carrier sense count = 03

Link is good on GE port 12 connected to device: FPC6
  Speed is 100Mb
  Duplex is full
  Autonegotiate is Enabled
  False carrier sense count = 03

Link is good on GE port 14 connected to device: SPMB
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled

Link is good on GE port 18 connected to device: RE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled

Link is good on GE port 19 connected to device: Other RE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
```

show chassis ethernet-switch errors (T4000 Router)

```
user@host> show chassis ethernet-switch errors

Displaying error for switch 0
Accumulated error counts for port 6 connected to device FPC0:
  MLT3           0
  Lock           0
```

```
  Xmit           0
  ESD            0
  False carrier  4
  Disconnects    0
  FX mode        0
Accumulated error counts for port 9 connected to device FPC3:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  3
  Disconnects    0
  FX mode        0
Accumulated error counts for port 11 connected to device FPC5:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  3
  Disconnects    0
  FX mode        0
Accumulated error counts for port 12 connected to device FPC6:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  3
  Disconnects    0
  FX mode        0
Accumulated error counts for port 19 connected to device Other RE:
  MLT3           0
  Lock           0
  Xmit           0
  ESD            0
  False carrier  0
  Disconnects    0
  FX mode        0
```

show chassis ethernet-switch (PTX5000 Packet Transport Router)

```
user@host> show chassis ethernet-switch
Displaying summary for switch 0
Link is good on XE port 2 connected to device: SPMB
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 11 connected to device: FPC7
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 12 connected to device: FPC6
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 13 connected to device: FPC5
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 15 connected to device: FPC3
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 16 connected to device: FPC2
  Speed is 1000Mb
  Duplex is full
```

```
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 18 connected to device: FPC0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 19 connected to device: OTHER RE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 20 connected to device: RE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
```

show chassis ethernet-switch statistics (PTX5000 Packet Transport Router)

```
user@host> show chassis ethernet-switch statistics
Displaying port statistics for switch 0
Statistics for port 2 connected to device SPMB:
  TX Packets 64 Octets        10942
  TX Packets 65-127 Octets    843
  TX Packets 128-255 Octets   2
  TX Packets 256-511 Octets   2
  TX Packets 512-1023 Octets  0
  TX Packets 1024-1518 Octets  6862
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   18651
```

```
  TX Multicast Packets        6
  TX Broadcast Packets        10331
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             8105166
  TX Packet OK Counter        0
  TX Pause Packet Counter     0
  TX Unicast Counter          0
  RX Packets 64 Octets        8679
  RX Packets 65-127 Octets    2364
  RX Packets 128-255 Octets   531
  RX Packets 256-511 Octets   112
  RX Packets 512-1023 Octets  26
  RX Packets 1024-1518 Octets  8
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   11720
  RX Multicast Packets        0
  RX Broadcast Packets        10331
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             938105
  RX Unicast Frame Count      0
  RX Packet OK Count          0
Statistics for port 11 connected to device FPC7:
  TX Packets 64 Octets        14492
  TX Packets 65-127 Octets    3542
  TX Packets 128-255 Octets   6
  TX Packets 256-511 Octets   45
  TX Packets 512-1023 Octets  60
```

Continued...

```
Statistics for port 18 connected to device FPC0:
  TX Packets 64 Octets        15212
  TX Packets 65-127 Octets    3810
  TX Packets 128-255 Octets   6
  TX Packets 256-511 Octets   43
  TX Packets 512-1023 Octets  66
  TX Packets 1024-1518 Octets  169
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   19306
  TX Multicast Packets        0
  TX Broadcast Packets        10886
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             1569412
  TX Packet OK Counter        0
  TX Pause Packet Counter     0
  TX Unicast Counter          0
  RX Packets 64 Octets        17994
  RX Packets 65-127 Octets    8006
  RX Packets 128-255 Octets   230
  RX Packets 256-511 Octets   19
  RX Packets 512-1023 Octets  53
  RX Packets 1024-1518 Octets  11
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   26313
  RX Multicast Packets        0
  RX Broadcast Packets        10886
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
```

```
  RX Control Frame Counter    2
  RX Pause Frame Counter      2
  RX Byte Counter             1836287
  RX Unicast Frame Count      0
  RX Packet OK Count          0
Statistics for port 19 connected to device OTHER RE:
  TX Packets 64 Octets        10234
  TX Packets 65-127 Octets    162
  TX Packets 128-255 Octets   0
  TX Packets 256-511 Octets   0
  TX Packets 512-1023 Octets  0
  TX Packets 1024-1518 Octets  0
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   10396
  TX Multicast Packets        8
  TX Broadcast Packets        10317
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             666260
  TX Packet OK Counter        0
  TX Pause Packet Counter     0
  TX Unicast Counter          0
  RX Packets 64 Octets        4073
  RX Packets 65-127 Octets    325
  RX Packets 128-255 Octets   1
  RX Packets 256-511 Octets   0
  RX Packets 512-1023 Octets  0
  RX Packets 1024-1518 Octets  72
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   4471
  RX Multicast Packets        0
  RX Broadcast Packets        10317
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
```

```
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             387333
  RX Unicast Frame Count      0
  RX Packet OK Count          0
Statistics for port 20 connected to device RE:
  TX Packets 64 Octets        658856
  TX Packets 65-127 Octets    45535
  TX Packets 128-255 Octets   1900
  TX Packets 256-511 Octets   532
  TX Packets 512-1023 Octets  372
  TX Packets 1024-1518 Octets  191
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   707386
  TX Multicast Packets        0
  TX Broadcast Packets        10421
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             46608676
  TX Packet OK Counter        0
  TX Pause Packet Counter     0
  TX Unicast Counter          0
  RX Packets 64 Octets        27394
  RX Packets 65-127 Octets    20271
  RX Packets 128-255 Octets   78
  RX Packets 256-511 Octets   215
  RX Packets 512-1023 Octets  269
  RX Packets 1024-1518 Octets  253370
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   301597
  RX Multicast Packets        8
```

```
  RX Broadcast Packets        10421
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             275043436
  RX Unicast Frame Count      0
  RX Packet OK Count          0

Continued ...
```

show chassis ethernet-switch port-state (PTX5000 Packet Transport Router)

```
user@host> show chassis ethernet-switch port-state
Displaying port state for switch 0
Port       : 02
Target     : SPMB

Error reading port 2 connected to device: SPMB
```

Release Information

Command introduced before Junos OS Release 7.4.

The `sfc` option introduced in Junos OS Release 9.6 for the TX Matrix Plus router.

show chassis ethernet-switch interconnect-device cb

IN THIS SECTION

Syntax | 17961

- Description | 17961
- Options | 17961
- Required Privilege Level | 17961
- Output Fields | 17962
- Sample Output | 17967
- Release Information | 17985

Syntax

```
show chassis ethernet-switch interconnect-device name cb
<detail>
<port number>
<slotnumber>
```

Description

(QFX3000-G QFabric systems only) Display Ethernet switch information for the Control Board (CB) ports in an Interconnect device.

Options

none	Display Ethernet switch information about each connected port on each online CB in the Interconnect device.
detail	(Optional) Display detailed status information for all CBs or for the CB in the specified slot in the Interconnect device.
port *number*	(Optional) Display Ethernet switch information about a specific port on a CB in the Interconnect device.
slot *number*	(Optional) Display Ethernet switch information about a CB in a specific slot in the Interconnect device.

Required Privilege Level

view

Output Fields

Table 246 on page 17962 lists the output fields for the `show chassis ethernet-switch interconnect-device cb` command. Output fields are listed in the approximate order in which they appear.

Table 246: show chassis ethernet-switch interconnect-device fpc Output Fields

Field Name	Field Description
Link is good on port n connected to device	Information about the link between each port on the FPC's Ethernet switch and one of the following devices: • FWD-SWITCH-0 • FWD-SWITCH-1 • CB0 • CB1
Speed is	Speed at which the Ethernet link is running: **10 Mb**When the device is **RE** or **Other RE** on the TX Matrix router, the speed is **1000 Mb**.
Duplex is	Duplex type of the Ethernet link: **full** or **half**.
Autonegotiate is Enabled (or Disabled)	By default, built-in Fast Ethernet ports on a PIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the `no-concatenate` statement at the `[edit chassis]` hierarchy level, as described in the *Junos OS System Basics Configuration Guide*).
Flow Control TX is Enabled (or Disabled)	Flow control in the transmit direction is enabled (or disabled). Flow control regulates the flow of packets from the switch to the remote side of the connection.
Flow Control RX is Enabled (or Disabled)	Flow control in the receive direction is enabled (or disabled). Flow control regulates the flow of packets from the remote side of the connection to the switch.
TX Octets	Number of octets sent.

Table 246: show chassis ethernet-switch interconnect-device fpc Output Fields *(Continued)*

Field Name	Field Description
TX Packets 64 Octets	Number of transmitted packets of size 64 octets.
TX Packets 65-127 Octets	Number of transmitted frames of size 65 through 127 octets.
TX Packets 128-255 Octets	Number of transmitted frames of size 128 through 255 octets.
TX Packets 256-511 Octets	Number of transmitted frames of size 256 through 511 octets.
TX Packets 512-1023 Octets	Number of transmitted frames of size 512 through 1023 octets.
TX Packets 1024-1518 Octets	Number of transmitted frames of size 1024 through 1518 octets.
TX Packets 1519-2047 Octets	Number of transmitted frames of size 1519 through 2047 octets.
TX Packets 2048-4095 Octets	Number of transmitted frames of size 2048 through 4095 octets.
TX Packets 4096-9216 Octets	Number of transmitted frames of size 4096 through 9216 octets.
TX Packets 9217-16383 Octets	Number of transmitted frames of size 9217 through 16383 octets.
TX Multicast packets	Number of multicast packets sent.

Table 246: show chassis ethernet-switch interconnect-device fpc Output Fields *(Continued)*

Field Name	Field Description
TX Broadcast packets	Number of broadcast packets sent.
TX Single Collision frames	Number of packets sent after one collision.
TX Mult. Collision frames	Number of packets sent after multiple collisions.
TX Late Collision Frames	Number of packets terminated during sending because of collisions after 64 bytes.
TX Excessive collisions	Number of packets not sent because of too many collisions.
TX Collision frames	Number of collision packets sent.
TX PAUSEMAC Ctrl Frames	Number of Media Access Control (MAC) frames containing PAUSE commands sent.
TX MAC ctrl frames	Number of MAC control packets sent.
TX Frame deferred Xmns	Number of frames deferred in *x* milliseconds.
TX Oversize Packets	Number of oversized packets sent.
TX Jabbers	Total number of frames sent that exceed the maximum byte count and contain CRC errors .
TX FCS Error Counter	Number of packets discarded because of frame check sequence errors.

Table 246: show chassis ethernet-switch interconnect-device fpc Output Fields *(Continued)*

Field Name	Field Description
TX Fragment Counter	Number of fragmented packets sent.
TX Byte Counter	Number of bytes sent.
RX Octets	Number of octets received.
RX Packets 64 Octets	Number of received packets of size 64 octets.
RX Packets 65-127 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 128-255 Octets	Number of received packets of size 128 through 255 octets.
RX Packets 256-511 Octets	Number of received packets of size 256 through 511 octets.
RX Packets 512-1023 Octets	Number of received packets of size 512 through 1023 octets.
RX Packets 1024-1518 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 1519-2047 Octets	Number of received packets of size 1519 through 2047 octets.
RX Packets 2048-4095 Octets	Number of received packets of size 2048 through 4095 octets.
RX Packets 4096-9216 Octets	Number of received packets of size 4096 through 9216 octets.

Table 246: show chassis ethernet-switch interconnect-device fpc Output Fields *(Continued)*

Field Name	Field Description
RX Multicast Packets	Number of multicast packets received.
RX Broadcast Packets	Number of broadcast packets received.
RX FCS Errors	Number of packets discarded because of frame check sequence errors.
RX Align Errors	Number of incomplete octets received.
RX Fragments	Number of fragmented packets received.
RX Symbol errors	Number of symbols received that the router did not correctly decode.
RX Unsupported opcodes	Number of packets received with unsupported op codes.
RX Out of Range Length	Number of packets received with an out of range length.
RX False Carrier Errors	Number of packets received with false carrier errors.
RX Undersize Packets	Number of undersized packets received.
RX Oversize Packets	Number of oversized packets received.
RX Jabbers	Total number of frames received that exceed the maximum byte count and contain CRC errors .
RX 1519-1522 Good Vlan frms	

Table 246: show chassis ethernet-switch interconnect-device fpc Output Fields *(Continued)*

Field Name	Field Description
RX MTU Exceed Counter	Number of packets received that exceed the MTU.
RX Control Frame Counter	Number of control frames received.
RX Pause Frame Counter	Number of pause frames received.
RX Byte Counter	Number of bytes received.

Sample Output

show chassis ethernet-switch interconnect-device cb

```
user@switch> show chassis ethernet-switch interconnect-device IC-WS001 cb
Displaying summary for switch 0
Link is down on XE port 1 connected to device: FPC7
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 2 connected to device: FPC6
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 3 connected to device: FPC5
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 5 connected to device: FPC4
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 7 connected to device: FPC3
  Flow Control TX is Disabled
```

```
  Flow Control RX is Disabled

Link is down on XE port 9 connected to device: FPC2
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 10 connected to device: FPC1
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  TX Octets                   326358
  RX Octets                   237947

Link is good on XE port 11 connected to device: FPC0
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  TX Octets                   548249
  RX Octets                   386013

Link is down on XE port 20 connected to device: SFP3
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 21 connected to device: SFP2
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on XE port 22 connected to device: SFP1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  TX Octets                   1
  RX Octets                   11704758

Link is good on XE port 23 connected to device: SFP0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  TX Octets                   1500022
  RX Octets                   11629453

Link is good on XE port 24 connected to device: VCCPD
```

```
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  TX Octets                   23332467
  RX Octets                   1500023

Link is good on GE port 25 connected to device: SFI
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  TX Octets                   643918
  RX Octets                   894548
```

show chassis ethernet-switch interconnect-device cb detail

```
user@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 cb detail
Port statistics for CB switch

Link is down on XE port 1 connected to device: FPC7

Link is down on XE port 2 connected to device: FPC6

Link is down on XE port 3 connected to device: FPC5

Link is down on XE port 5 connected to device: FPC4

Link is down on XE port 7 connected to device: FPC3

Link is down on XE port 9 connected to device: FPC2
Statistics for port 10 connected to device FPC1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   319293
  TX Packets 256-511 Octets   5043
  TX Packets 512-1023 Octets  2072
  TX Packets 1024-1518 Octets  6
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   326415
```

```
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             71659246
  TX Packet OK Counter        326415
  TX Pause Packet Counter     0
  TX Unicast Counter          326414
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   235428
  RX Packets 256-511 Octets   2134
  RX Packets 512-1023 Octets  420
  RX Packets 1024-1518 Octets  6
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   237988
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             55821504
  RX Unicast Frame Count      237988
  RX Packet OK Count          237988
Statistics for port 11 connected to device FPC0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   535483
  TX Packets 256-511 Octets   9289
  TX Packets 512-1023 Octets  3564
  TX Packets 1024-1518 Octets  5
  TX Packets 1519-2047 Octets  0
```

```
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   548342
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             120498414
  TX Packet OK Counter        548342
  TX Pause Packet Counter     0
  TX Unicast Counter          548341
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   382931
  RX Packets 256-511 Octets   2762
  RX Packets 512-1023 Octets  386
  RX Packets 1024-1518 Octets  3
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   386082
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             90717369
  RX Unicast Frame Count      386082
  RX Packet OK Count          386082

Link is down on XE port 20 connected to device: SFP3

Link is down on XE port 21 connected to device: SFP2
```

```
Statistics for port 22 connected to device SFP1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    0
  TX Packets 128-255 Octets   0
  TX Packets 256-511 Octets   0
  TX Packets 512-1023 Octets  0
  TX Packets 1024-1518 Octets  1
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan Frms  0
  TX Octets                   1
  TX Multicast Packets        1
  TX Broadcast Packets        0
  TX Single Collision Frames  0
  TX Mult. Collision Frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision Frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC Ctrl Frames          0
  TX Frame Deferred Xmns      0
  TX Frame Excessive Deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             1422
  RX Packet OK Count          1
  RX Packets 64 Octets        230013
  RX Packets 65-127 Octets    174529
  RX Packets 128-255 Octets   286735
  RX Packets 256-511 Octets   343412
  RX Packets 512-1023 Octets  172152
  RX Packets 1024-1518 Octets  10500065
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   11706906
  RX Multicast Packets        11672320
  RX Broadcast Packets        34460
  RX FCS Errors               0
  RX Align Errors             0
```

```
  RX Fragments                0
  RX Symbol Errors            0
  RX Unsupported Opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan Frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             2379464164
  RX Packet OK Count          11706906
Statistics for port 23 connected to device SFP0:
  TX Packets 64 Octets        3
  TX Packets 65-127 Octets    484733
  TX Packets 128-255 Octets   219112
  TX Packets 256-511 Octets   129014
  TX Packets 512-1023 Octets  503
  TX Packets 1024-1518 Octets  666958
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan Frms  0
  TX Octets                   1500323
  TX Multicast Packets        794098
  TX Broadcast Packets        1040
  TX Single Collision Frames  0
  TX Mult. Collision Frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision Frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC Ctrl Frames          0
  TX Frame Deferred Xmns      0
  TX Frame Excessive Deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             1065466891
  RX Packet OK Count          1500323
```

```
  RX Packets 64 Octets        341563
  RX Packets 65-127 Octets    430810
  RX Packets 128-255 Octets   318279
  RX Packets 256-511 Octets   347147
  RX Packets 512-1023 Octets  184798
  RX Packets 1024-1518 Octets  10008993
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   11631590
  RX Multicast Packets        10878484
  RX Broadcast Packets        33420
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol Errors            0
  RX Unsupported Opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan Frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             1720484325
  RX Packet OK Count          11631591
Statistics for port 24 connected to device VCCPD:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1176546
  TX Packets 128-255 Octets   604988
  TX Packets 256-511 Octets   690561
  TX Packets 512-1023 Octets  356942
  TX Packets 1024-1518 Octets  20507438
  TX Packets 1519-2047 Octets  278
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan Frms  278
  TX Octets                   23336753
  TX Multicast Packets        22549383
  TX Broadcast Packets        67862
  TX Single Collision Frames  0
```

```
  TX Mult. Collision Frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision Frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC Ctrl Frames          0
  TX Frame Deferred Xmns      0
  TX Frame Excessive Deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             4191296788
  RX Packet OK Count          23336753
  RX Packets 64 Octets        3
  RX Packets 65-127 Octets    484673
  RX Packets 128-255 Octets   219074
  RX Packets 256-511 Octets   129100
  RX Packets 512-1023 Octets  516
  RX Packets 1024-1518 Octets  666959
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   1500325
  RX Multicast Packets        794099
  RX Broadcast Packets        1040
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol Errors            0
  RX Unsupported Opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan Frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             1071469739
  RX Packet OK Count          1500325
Statistics for port 25 connected to device SFI:
```

```
TX Packets 64 Octets        12
TX Packets 65-127 Octets    1
TX Packets 128-255 Octets   618363
TX Packets 256-511 Octets   4896
TX Packets 512-1023 Octets  806
TX Packets 1024-1518 Octets  19950
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan Frms  0
TX Octets                   644028
TX Multicast Packets        4
TX Broadcast Packets        19954
TX Single Collision Frames  0
TX Mult. Collision Frames   0
TX Late Collisions          0
TX Excessive Collisions     0
TX Collision Frames         0
TX PAUSEMAC Ctrl Frames     0
TX MAC Ctrl Frames          0
TX Frame Deferred Xmns      0
TX Frame Excessive Deferl   0
TX Oversize Packets         0
TX Jabbers                  0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             167039705
RX Packet OK Count          644028
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   854776
RX Packets 256-511 Octets   14332
RX Packets 512-1023 Octets  5636
RX Packets 1024-1518 Octets  19954
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                   894698
RX Multicast Packets        0
RX Broadcast Packets        19943
RX FCS Errors               0
RX Align Errors             0
RX Fragments                0
```

```
  RX Symbol Errors            0
  RX Unsupported Opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan Frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             212658920
  RX Packet OK Count          894698
```

show chassis ethernet-switch interconnect-device cb detail slot port

```
user@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 cb slot 1 port
1
re0:
--------------------------------------------------------------------------
Port statistics for CB switch

Link is down on XE port 1 connected to device: FPC7

Link is down on XE port 2 connected to device: FPC6

Link is down on XE port 3 connected to device: FPC5

Link is down on XE port 5 connected to device: FPC4

Link is down on XE port 7 connected to device: FPC3

Link is down on XE port 9 connected to device: FPC2
Statistics for port 10 connected to device FPC1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   319366
  TX Packets 256-511 Octets   5043
  TX Packets 512-1023 Octets  2072
  TX Packets 1024-1518 Octets  6
  TX Packets 1519-2047 Octets  0
```

```
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   326488
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             71675330
  TX Packet OK Counter        326488
  TX Pause Packet Counter     0
  TX Unicast Counter          326487
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   235481
  RX Packets 256-511 Octets   2134
  RX Packets 512-1023 Octets  420
  RX Packets 1024-1518 Octets  6
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   238041
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             55834224
  RX Unicast Frame Count      238041
  RX Packet OK Count          238041
Statistics for port 11 connected to device FPC0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   535606
```

```
TX Packets 256-511 Octets   9289
TX Packets 512-1023 Octets  3564
TX Packets 1024-1518 Octets  5
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX Packets 9217-16383 Octets  0
TX Octets                   548465
TX Multicast Packets        0
TX Broadcast Packets        1
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets         0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             120525524
TX Packet OK Counter        548465
TX Pause Packet Counter     0
TX Unicast Counter          548464
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   383018
RX Packets 256-511 Octets   2762
RX Packets 512-1023 Octets  386
RX Packets 1024-1518 Octets  3
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Packets 9217-16383 Octets  0
RX Octets                   386169
RX Multicast Packets        0
RX Broadcast Packets        0
RX FCS Errors               0
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             90738249
RX Unicast Frame Count      386169
RX Packet OK Count          386169
```

```
Link is down on XE port 20 connected to device: SFP3

Link is down on XE port 21 connected to device: SFP2
Statistics for port 22 connected to device SFP1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    0
  TX Packets 128-255 Octets   0
  TX Packets 256-511 Octets   0
  TX Packets 512-1023 Octets  0
  TX Packets 1024-1518 Octets  1
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan Frms  0
  TX Octets                   1
  TX Multicast Packets        1
  TX Broadcast Packets        0
  TX Single Collision Frames  0
  TX Mult. Collision Frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision Frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC Ctrl Frames          0
  TX Frame Deferred Xmns      0
  TX Frame Excessive Deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             1422
  RX Packet OK Count          1
  RX Packets 64 Octets        230071
  RX Packets 65-127 Octets    174571
  RX Packets 128-255 Octets   286812
  RX Packets 256-511 Octets   343500
  RX Packets 512-1023 Octets  172203
  RX Packets 1024-1518 Octets  10502544
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   11709701
```

```
  RX Multicast Packets        11675110
  RX Broadcast Packets        34465
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol Errors            0
  RX Unsupported Opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan Frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             2383079858
  RX Packet OK Count          11709701
Statistics for port 23 connected to device SFP0:
  TX Packets 64 Octets        3
  TX Packets 65-127 Octets    485048
  TX Packets 128-255 Octets   219200
  TX Packets 256-511 Octets   129053
  TX Packets 512-1023 Octets  503
  TX Packets 1024-1518 Octets  667127
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan Frms  0
  TX Octets                   1500934
  TX Multicast Packets        794300
  TX Broadcast Packets        1040
  TX Single Collision Frames  0
  TX Mult. Collision Frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision Frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC Ctrl Frames          0
  TX Frame Deferred Xmns      0
  TX Frame Excessive Deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
```

```
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             1065764997
  RX Packet OK Count          1500934
  RX Packets 64 Octets        341648
  RX Packets 65-127 Octets    431183
  RX Packets 128-255 Octets   318367
  RX Packets 256-511 Octets   347225
  RX Packets 512-1023 Octets  184849
  RX Packets 1024-1518 Octets  10011311
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   11634583
  RX Multicast Packets        10881071
  RX Broadcast Packets        33425
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol Errors            0
  RX Unsupported Opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan Frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             1723893006
  RX Packet OK Count          11634583
Statistics for port 24 connected to device VCCPD:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1177102
  TX Packets 128-255 Octets   605153
  TX Packets 256-511 Octets   690727
  TX Packets 512-1023 Octets  357044
  TX Packets 1024-1518 Octets  20512235
  TX Packets 1519-2047 Octets  278
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan Frms  278
```

```
TX Octets                   23342539
TX Multicast Packets        22554760
TX Broadcast Packets        67872
TX Single Collision Frames  0
TX Mult. Collision Frames   0
TX Late Collisions          0
TX Excessive Collisions     0
TX Collision Frames         0
TX PAUSEMAC Ctrl Frames     0
TX MAC Ctrl Frames          0
TX Frame Deferred Xmns      0
TX Frame Excessive Deferl   0
TX Oversize Packets         0
TX Jabbers                  0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             4198344167
RX Packet OK Count          23342539
RX Packets 64 Octets        3
RX Packets 65-127 Octets    484985
RX Packets 128-255 Octets   219164
RX Packets 256-511 Octets   129139
RX Packets 512-1023 Octets  516
RX Packets 1024-1518 Octets  667128
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                   1500935
RX Multicast Packets        794301
RX Broadcast Packets        1040
RX FCS Errors               0
RX Align Errors             0
RX Fragments                0
RX Symbol Errors            0
RX Unsupported Opcodes      0
RX Out of Range Length      0
RX False Carrier Errors     0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX 1519-1522 Good Vlan Frms 0
RX MTU Exceed Counter       0
RX Control Frame Counter    0
```

```
  RX Pause Frame Counter      0
  RX Byte Counter             1071770147
  RX Packet OK Count          1500935
Statistics for port 25 connected to device SFI:
  TX Packets 64 Octets        12
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   618503
  TX Packets 256-511 Octets   4896
  TX Packets 512-1023 Octets  806
  TX Packets 1024-1518 Octets  19950
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan Frms  0
  TX Octets                   644168
  TX Multicast Packets        4
  TX Broadcast Packets        19954
  TX Single Collision Frames  0
  TX Mult. Collision Frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision Frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC Ctrl Frames          0
  TX Frame Deferred Xmns      0
  TX Frame Excessive Deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             167073305
  RX Packet OK Count          644168
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   854972
  RX Packets 256-511 Octets   14332
  RX Packets 512-1023 Octets  5636
  RX Packets 1024-1518 Octets  19954
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   894894
  RX Multicast Packets        0
```

```
RX Broadcast Packets        19943
RX FCS Errors               0
RX Align Errors             0
RX Fragments                0
RX Symbol Errors            0
RX Unsupported Opcodes      0
RX Out of Range Length      0
RX False Carrier Errors     0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX 1519-1522 Good Vlan Frms 0
RX MTU Exceed Counter       0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             212702114
RX Packet OK Count          894894
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

chassis (QFabric System) | 3303

show chassis environment cb | 17573

show chassis ethernet-switch interconnect-device fpc | 17985

show chassis ethernet-switch interconnect-device fpc

IN THIS SECTION

- Syntax | 17986
- Description | 17986
- Options | 17986

- Required Privilege Level | 17986
- Output Fields | 17987
- Sample Output | 17992
- Release Information | 18021

Syntax

```
show chassis ethernet-switch interconnect-device name fpc
<detail>
<port number>
<slot number>
```

Description

(QFX3000-G QFabric systems only) Display Ethernet switch information for the front card Flexible Port Concentrators (FPCs) in an Interconnect device.

Options

none Display Ethernet switch information about each connected port on each online FPC in the Interconnect device.

detail (Optional) Display detailed status information for all FPCs or for the FPC in the specified slot in the Interconnect device.

port *number* (Optional) Display Ethernet switch information about a specific port on an FPC in the Interconnect device.

slot *number* (Optional) Display Ethernet switch information about an FPC in a specific slot in the Interconnect device.

Required Privilege Level

view

Output Fields

Table 247 on page 17987 lists the output fields for the `show chassis ethernet-switch interconnect-device fpc` command. Output fields are listed in the approximate order in which they appear.

Table 247: show chassis ethernet-switch interconnect-device fpc Output Fields

Field Name	Field Description
Link is good on port n connected to device	Information about the link between each port on the FPC's Ethernet switch and one of the following devices: • FWD-SWITCH-0 • FWD-SWITCH-1 • CB0 • CB1
Speed is	Speed at which the Ethernet link is running: **10 Mb**When the device is **RE** or **Other RE** on the TX Matrix router, the speed is **1000 Mb**.
Duplex is	Duplex type of the Ethernet link: **full** or **half**.
Autonegotiate is Enabled (or Disabled)	By default, built-in Fast Ethernet ports on a PIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the `no-concatenate` statement at the `[edit chassis]` hierarchy level, as described in the *Junos OS System Basics Configuration Guide*).
TX Octets	Number of octets sent.
TX Packets 64 Octets	Number of transmitted frames of size 64 octets.
TX Packets 65-127 Octets	Number of transmitted frames of size 65 through 127 octets.

Table 247: show chassis ethernet-switch interconnect-device fpc Output Fields *(Continued)*

Field Name	Field Description
TX Packets 128-255 Octets	Number of transmitted frames of size 128 through 255 octets.
TX Packets 256-511 Octets	Number of transmitted frames of size 256 through 511 octets.
TX Packets 512-1023 Octets	Number of transmitted frames of size 512 through 1023 octets.
TX Packets 1024-1518 Octets	Number of transmitted frames of size 1024 through 1518 octets.
TX Packets 1519-2047 Octets	Number of transmitted frames of size 1519 through 2047 octets.
TX Packets 2048-4095 Octets	Number of transmitted frames of size 2048 through 4095 octets.
TX Packets 4096-9216 Octets	Number of transmitted frames of size 4096 through 9216 octets.
TX Packets 9217-16383 Octets	Number of transmitted frames of size 9217 through 16383 octets.
TX Multicast packets	Number of multicast packets sent.
TX Broadcast packets	Number of broadcast packets sent.
TX Single Collision frames	Number of packets sent after one collision.

Table 247: show chassis ethernet-switch interconnect-device fpc Output Fields *(Continued)*

Field Name	Field Description
TX Mult. Collision frames	Number of packets sent after multiple collisions.
TX Late Collision Frames	Number of packets terminated during sending because of collisions after 64 bytes.
TX Excessive collisions	Number of packets not sent because of too many collisions.
TX Collision frames	Number of collision packets sent.
TX PAUSEMAC Ctrl Frames	Number of Media Access Control (MAC) frames containing PAUSE commands sent.
TX MAC ctrl frames	Number of MAC control packets sent.
TX Frame deferred Xmns	Number of frames deferred in *x* milliseconds.
TX Oversize Packets	Number of oversized packets sent.
TX Jabbers	Total number of frames sent that exceed the maximum byte count and contain CRC errors .
TX FCS Error Counter	Number of packets discarded because of frame check sequence errors.
TX Fragment Counter	Number of fragmented packets sent.
TX Byte Counter	Number of bytes sent.
RX Octets	Number of octets received.

Table 247: show chassis ethernet-switch interconnect-device fpc Output Fields *(Continued)*

Field Name	Field Description
RX Packets 64 Octets	Number of received packets of size 64 octets.
RX Packets 65-127 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 128-255 Octets	Number of received packets of size 128 through 255 octets.
RX Packets 256-511 Octets	Number of received packets of size 256 through 511 octets.
RX Packets 512-1023 Octets	Number of received packets of size 512 through 1023 octets.
RX Packets 1024-1518 Octets	Number of received packets of size 65 through 127 octets.
RX Packets 1519-2047 Octets	Number of received packets of size 1519 through 2047 octets.
RX Packets 2048-4095 Octets	Number of received packets of size 2048 through 4095 octets.
RX Packets 4096-9216 Octets	Number of received packets of size 4096 through 9216 octets.
RX Multicast Packets	Number of multicast packets received.
RX Broadcast Packets	Number of broadcast packets received.

Table 247: show chassis ethernet-switch interconnect-device fpc Output Fields *(Continued)*

Field Name	Field Description
RX FCS Errors	Number of packets discarded because of frame check sequence errors.
RX Align Errors	Number of incomplete octets received.
RX Fragments	Number of fragmented packets received.
RX Symbol errors	Number of symbols received that the router did not correctly decode.
RX Unsupported opcodes	Number of packets received with unsupported op codes.
RX Out of Range Length	Number of packets received with an out of range length.
RX False Carrier Errors	Number of packets received with false carrier errors.
RX Undersize Packets	Number of undersized packets received.
RX Oversize Packets	Number of oversized packets received.
RX Jabbers	Total number of frames received that exceed the maximum byte count and contain CRC errors .
RX 1519-1522 Good Vlan frms	Number of transmitted frames of size 1519 through 1522 octets that are good VLAN frames.
RX MTU Exceed Counter	Number of packets received that exceed the MTU.
RX Control Frame Counter	Number of control frames received.

Table 247: show chassis ethernet-switch interconnect-device fpc Output Fields *(Continued)*

Field Name	Field Description
RX Pause Frame Counter	Number of pause frames received.
RX Byte Counter	Number of bytes received.

Sample Output

show chassis ethernet-switch interconnect-device fpc

```
user@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 fpc
Summary for switch on FC0
Link is good on GE port 2 connected to device: FWD-SWITCH-0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   124638
  RX Octets                   86496

Link is good on GE port 4 connected to device: FWD-SWITCH-1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   82191
  RX Octets                   58979

Link is good on XE port 28 connected to device: CB0
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
```

```
  TX Octets                   145475
  RX Octets                   206828

Link is good on XE port 29 connected to device: CB1
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   1
  RX Octets                   0

Summary for switch on FC1
Link is good on GE port 2 connected to device: FWD-SWITCH-0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   82290
  RX Octets                   59443

Link is good on GE port 4 connected to device: FWD-SWITCH-1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   40900
  RX Octets                   30013

Link is good on XE port 28 connected to device: CB0
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   89456
  RX Octets                   123189

Link is good on XE port 29 connected to device: CB1
  Speed is 10000Mb
  Duplex is full
```

```
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   1
  RX Octets                   0

root@qfabric> show chassis ethernet-switch interconnect-device IC-WS001 fpc
Summary for switch on FC0
Link is good on GE port 2 connected to device: FWD-SWITCH-0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   124697
  RX Octets                   86535

Link is good on GE port 4 connected to device: FWD-SWITCH-1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   82229
  RX Octets                   59009

Link is good on XE port 28 connected to device: CB0
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   145544
  RX Octets                   206925

Link is good on XE port 29 connected to device: CB1
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   1
  RX Octets                   0
```

```
Summary for switch on FC1
Link is good on GE port 2 connected to device: FWD-SWITCH-0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   82327
  RX Octets                   59472

Link is good on GE port 4 connected to device: FWD-SWITCH-1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   40918
  RX Octets                   30028

Link is good on XE port 28 connected to device: CB0
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   89500
  RX Octets                   123244

Link is good on XE port 29 connected to device: CB1
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   1
  RX Octets                   0
```

show chassis ethernet-switch interconnect-device fpc detail

```
user@host> show chassis ethernet-switch interconnect-device IC-WS001 fpc detail
Port statistics for FC0 switch
Statistics for port 2 connected to device FWD-SWITCH-0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   121716
  TX Packets 256-511 Octets   2200
  TX Packets 512-1023 Octets  823
  TX Packets 1024-1518 Octets  2
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   124742
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             27391588
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   85924
  RX Packets 256-511 Octets   555
  RX Packets 512-1023 Octets  86
  RX Packets 1024-1518 Octets  1
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   86566
```

```
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             20380581
Statistics for port 4 connected to device FWD-SWITCH-1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   80374
  TX Packets 256-511 Octets   1347
  TX Packets 512-1023 Octets  532
  TX Packets 1024-1518 Octets  3
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   82257
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
```

```
  TX Fragment Counter         0
  TX Byte Counter             18146746
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   58410
  RX Packets 256-511 Octets   522
  RX Packets 512-1023 Octets  96
  RX Packets 1024-1518 Octets  2
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   59030
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             13882179
Statistics for port 28 connected to device CB0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    0
  TX Packets 128-255 Octets   144334
  TX Packets 256-511 Octets   1077
  TX Packets 512-1023 Octets  182
  TX Packets 1024-1518 Octets  3
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   145596
  TX Multicast Packets        0
  TX Broadcast Packets        0
```

```
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             34262760
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    1
  RX Packets 128-255 Octets   202090
  RX Packets 256-511 Octets   3547
  RX Packets 512-1023 Octets  1355
  RX Packets 1024-1518 Octets  5
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   206998
  RX Multicast Packets        0
  RX Broadcast Packets        1
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             45538262
Statistics for port 29 connected to device CB1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   0
  TX Packets 256-511 Octets   0
  TX Packets 512-1023 Octets  0
  TX Packets 1024-1518 Octets  0
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   1
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX PAUSEMAC Ctrl Frames     0
```

```
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             72
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   0
  RX Packets 256-511 Octets   0
  RX Packets 512-1023 Octets  0
  RX Packets 1024-1518 Octets  0
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   0
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             0

Port statistics for FC1 switch
Statistics for port 2 connected to device FWD-SWITCH-0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   80560
  TX Packets 256-511 Octets   1279
  TX Packets 512-1023 Octets  514
  TX Packets 1024-1518 Octets  3
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   82357
  TX Multicast Packets        0
  TX Broadcast Packets        1
```

```
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             18059906
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   58733
  RX Packets 256-511 Octets   639
  RX Packets 512-1023 Octets  119
  RX Packets 1024-1518 Octets  3
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   59494
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             13994432
Statistics for port 4 connected to device FWD-SWITCH-1:
  TX Packets 64 Octets        0
```

```
TX Packets 65-127 Octets    1
TX Packets 128-255 Octets   39971
TX Packets 256-511 Octets   668
TX Packets 512-1023 Octets  290
TX Packets 1024-1518 Octets  3
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan frms  0
TX Octets                   40933
TX Multicast Packets        0
TX Broadcast Packets        1
TX Single Collision frames  0
TX Mult. Collision frames   0
TX Late Collisions          0
TX Excessive Collisions     0
TX Collision frames         0
TX PAUSEMAC Ctrl Frames     0
TX MAC ctrl frames          0
TX Frame deferred Xmns      0
TX Frame excessive deferl   0
TX Oversize Packets         0
TX Jabbers                  0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             9050841
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   29767
RX Packets 256-511 Octets   225
RX Packets 512-1023 Octets  44
RX Packets 1024-1518 Octets  3
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                   30039
RX Multicast Packets        0
RX Broadcast Packets        0
RX FCS Errors               0
RX Align Errors             0
RX Fragments                0
RX Symbol errors            0
RX Unsupported opcodes      0
```

```
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             7043738
Statistics for port 28 connected to device CB0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    0
  TX Packets 128-255 Octets   88500
  TX Packets 256-511 Octets   864
  TX Packets 512-1023 Octets  163
  TX Packets 1024-1518 Octets  6
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   89533
  TX Multicast Packets        0
  TX Broadcast Packets        0
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             21038170
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    1
  RX Packets 128-255 Octets   120531
  RX Packets 256-511 Octets   1947
  RX Packets 512-1023 Octets  804
  RX Packets 1024-1518 Octets  6
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   123289
  RX Multicast Packets        0
  RX Broadcast Packets        1
  RX FCS Errors               0
```

```
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             27110675
Statistics for port 29 connected to device CB1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   0
  TX Packets 256-511 Octets   0
  TX Packets 512-1023 Octets  0
  TX Packets 1024-1518 Octets  0
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   1
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             72
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   0
  RX Packets 256-511 Octets   0
  RX Packets 512-1023 Octets  0
  RX Packets 1024-1518 Octets  0
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   0
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Fragments                0
```

```
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             0

Port statistics for FC2 switch
Empty fpc slot number 2

Port statistics for FC3 switch
Empty fpc slot number 3

Port statistics for FC4 switch
Empty fpc slot number 4

Port statistics for FC5 switch
Empty fpc slot number 5

Port statistics for FC6 switch
Empty fpc slot number 6

Port statistics for FC7 switch
Empty fpc slot number 7
```

show chassis ethernet-switch fpc detail slot

```
user@qfabric> show chassis ethernet-switch fpc detail 0
re0:
--------------------------------------------------------------------------
Port statistics for FC0 switch
Statistics for port 2 connected to device FWD-SWITCH-0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   121823
  TX Packets 256-511 Octets   2200
  TX Packets 512-1023 Octets  823
  TX Packets 1024-1518 Octets  2
  TX Packets 1519-2047 Octets  0
```

```
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan frms  0
TX Octets                   124849
TX Multicast Packets        0
TX Broadcast Packets        1
TX Single Collision frames  0
TX Mult. Collision frames   0
TX Late Collisions          0
TX Excessive Collisions     0
TX Collision frames         0
TX PAUSEMAC Ctrl Frames     0
TX MAC ctrl frames          0
TX Frame deferred Xmns      0
TX Frame excessive deferl   0
TX Oversize Packets         0
TX Jabbers                  0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             27414524
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   85998
RX Packets 256-511 Octets   557
RX Packets 512-1023 Octets  86
RX Packets 1024-1518 Octets  1
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                   86642
RX Multicast Packets        0
RX Broadcast Packets        0
RX FCS Errors               0
RX Align Errors             0
RX Fragments                0
RX Symbol errors            0
RX Unsupported opcodes      0
RX Out of Range Length      0
RX False Carrier Errors     0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX 1519-1522 Good Vlan frms 0
```

```
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             20398564
Statistics for port 4 connected to device FWD-SWITCH-1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   80443
  TX Packets 256-511 Octets   1347
  TX Packets 512-1023 Octets  532
  TX Packets 1024-1518 Octets  3
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   82326
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             18161734
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   58460
  RX Packets 256-511 Octets   523
  RX Packets 512-1023 Octets  96
  RX Packets 1024-1518 Octets  2
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   59081
  RX Multicast Packets        0
```

```
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             13894171
Statistics for port 28 connected to device CB0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    0
  TX Packets 128-255 Octets   144458
  TX Packets 256-511 Octets   1080
  TX Packets 512-1023 Octets  182
  TX Packets 1024-1518 Octets  3
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   145723
  TX Multicast Packets        0
  TX Broadcast Packets        0
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             34292735
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    1
  RX Packets 128-255 Octets   202266
  RX Packets 256-511 Octets   3547
  RX Packets 512-1023 Octets  1355
  RX Packets 1024-1518 Octets  5
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
```

```
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   207174
  RX Multicast Packets        0
  RX Broadcast Packets        1
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             45576186
Statistics for port 29 connected to device CB1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   0
  TX Packets 256-511 Octets   0
  TX Packets 512-1023 Octets  0
  TX Packets 1024-1518 Octets  0
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   1
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             72
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   0
  RX Packets 256-511 Octets   0
  RX Packets 512-1023 Octets  0
  RX Packets 1024-1518 Octets  0
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
```

```
  RX Packets 9217-16383 Octets  0
  RX Octets                   0
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             0

Port statistics for FC1 switch
Statistics for port 2 connected to device FWD-SWITCH-0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   80629
  TX Packets 256-511 Octets   1279
  TX Packets 512-1023 Octets  514
  TX Packets 1024-1518 Octets  3
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   82426
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
```

```
  TX Byte Counter             18074790
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   58785
  RX Packets 256-511 Octets   640
  RX Packets 512-1023 Octets  119
  RX Packets 1024-1518 Octets  3
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   59547
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             14006842
Statistics for port 4 connected to device FWD-SWITCH-1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   40004
  TX Packets 256-511 Octets   668
  TX Packets 512-1023 Octets  290
  TX Packets 1024-1518 Octets  3
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   40966
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX Single Collision frames  0
```

```
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             9058102
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   29794
  RX Packets 256-511 Octets   225
  RX Packets 512-1023 Octets  44
  RX Packets 1024-1518 Octets  3
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   30066
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             7050000
Statistics for port 28 connected to device CB0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    0
```

```
  TX Packets 128-255 Octets   88579
  TX Packets 256-511 Octets   865
  TX Packets 512-1023 Octets  163
  TX Packets 1024-1518 Octets  6
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   89613
  TX Multicast Packets        0
  TX Broadcast Packets        0
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             21056842
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    1
  RX Packets 128-255 Octets   120633
  RX Packets 256-511 Octets   1947
  RX Packets 512-1023 Octets  804
  RX Packets 1024-1518 Octets  6
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   123391
  RX Multicast Packets        0
  RX Broadcast Packets        1
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             27132820
Statistics for port 29 connected to device CB1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   0
```

```
  TX Packets 256-511 Octets   0
  TX Packets 512-1023 Octets  0
  TX Packets 1024-1518 Octets  0
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   1
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             72
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   0
  RX Packets 256-511 Octets   0
  RX Packets 512-1023 Octets  0
  RX Packets 1024-1518 Octets  0
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   0
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             0

Port statistics for FC2 switch
Empty fpc slot number 2

Port statistics for FC3 switch
```

```
Empty fpc slot number 3

Port statistics for FC4 switch
Empty fpc slot number 4

Port statistics for FC5 switch
Empty fpc slot number 5

Port statistics for FC6 switch
Empty fpc slot number 6

Port statistics for FC7 switch
Empty fpc slot number 7
```

show chassis ethernet-switch fpc interconnect-device port

```
user@qfabric> show chassis ethernet-switch fpc interconnect-device IC-WS001 port 2
Summary for switch on FC0
Link is good on GE port 2 connected to device: FWD-SWITCH-0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   319466
  RX Octets                   221869

Link is good on GE port 4 connected to device: FWD-SWITCH-1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   210295
  RX Octets                   151164

Link is good on XE port 28 connected to device: CB0
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
```

```
  Flow Control RX is Disabled
  TX Octets                   373033
  RX Octets                   529760

Link is good on XE port 29 connected to device: CB1
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   1
  RX Octets                   0

Summary for switch on FC1
Link is good on GE port 2 connected to device: FWD-SWITCH-0
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   210760
  RX Octets                   152617

Link is good on GE port 4 connected to device: FWD-SWITCH-1
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Disabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   104587
  RX Octets                   77315

Link is good on XE port 28 connected to device: CB0
  Speed is 10000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   229932
  RX Octets                   315346

Link is good on XE port 29 connected to device: CB1
  Speed is 10000Mb
```

```
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled
  TX Octets                   1
  RX Octets                   0
```

show chassis ethernet-switch fpc interconnect-device detail port

```
user@qfabric> show chassis ethernet-switch fpc interconnect-device IC-WS001 detail port
2
Port statistics for FC0 switch
Statistics for port 2 connected to device FWD-SWITCH-0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   311974
  TX Packets 256-511 Octets   5552
  TX Packets 512-1023 Octets  2084
  TX Packets 1024-1518 Octets  2
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   319613
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             70091196
  RX Packets 64 Octets        0
```

```
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   220284
  RX Packets 256-511 Octets   1486
  RX Packets 512-1023 Octets  198
  RX Packets 1024-1518 Octets  1
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   221969
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             52192002
Statistics for port 4 connected to device FWD-SWITCH-1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   205595
  TX Packets 256-511 Octets   3426
  TX Packets 512-1023 Octets  1366
  TX Packets 1024-1518 Octets  3
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX 1519-1522 Good Vlan frms  0
  TX Octets                   210391
  TX Multicast Packets        0
  TX Broadcast Packets        1
  TX Single Collision frames  0
  TX Mult. Collision frames   0
  TX Late Collisions          0
```

```
  TX Excessive Collisions     0
  TX Collision frames         0
  TX PAUSEMAC Ctrl Frames     0
  TX MAC ctrl frames          0
  TX Frame deferred Xmns      0
  TX Frame excessive deferl   0
  TX Oversize Packets         0
  TX Jabbers                  0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             46380018
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    0
  RX Packets 128-255 Octets   149866
  RX Packets 256-511 Octets   1194
  RX Packets 512-1023 Octets  173
  RX Packets 1024-1518 Octets  2
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Octets                   151235
  RX Multicast Packets        0
  RX Broadcast Packets        0
  RX FCS Errors               0
  RX Align Errors             0
  RX Fragments                0
  RX Symbol errors            0
  RX Unsupported opcodes      0
  RX Out of Range Length      0
  RX False Carrier Errors     0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX 1519-1522 Good Vlan frms 0
  RX MTU Exceed Counter       0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             35496911
Statistics for port 28 connected to device CB0:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    0
  TX Packets 128-255 Octets   370150
  TX Packets 256-511 Octets   2680
```

```
  TX Packets 512-1023 Octets  371
  TX Packets 1024-1518 Octets  3
  TX Packets 1519-2047 Octets  0
  TX Packets 2048-4095 Octets  0
  TX Packets 4096-9216 Octets  0
  TX Packets 9217-16383 Octets  0
  TX Octets                   373204
  TX Multicast Packets        0
  TX Broadcast Packets        0
  TX PAUSEMAC Ctrl Frames     0
  TX Oversize Packets         0
  TX FCS Error Counter        0
  TX Fragment Counter         0
  TX Byte Counter             87688913
  RX Packets 64 Octets        0
  RX Packets 65-127 Octets    1
  RX Packets 128-255 Octets   517569
  RX Packets 256-511 Octets   8978
  RX Packets 512-1023 Octets  3450
  RX Packets 1024-1518 Octets  5
  RX Packets 1519-2047 Octets  0
  RX Packets 2048-4095 Octets  0
  RX Packets 4096-9216 Octets  0
  RX Packets 9217-16383 Octets  0
  RX Octets                   530003
  RX Multicast Packets        0
  RX Broadcast Packets        1
  RX FCS Errors               0
  RX Fragments                0
  RX MAC Control Packets      0
  RX Out of Range Length      0
  RX Undersize Packets        0
  RX Oversize Packets         0
  RX Jabbers                  0
  RX Control Frame Counter    0
  RX Pause Frame Counter      0
  RX Byte Counter             116471142
Statistics for port 29 connected to device CB1:
  TX Packets 64 Octets        0
  TX Packets 65-127 Octets    1
  TX Packets 128-255 Octets   0
  TX Packets 256-511 Octets   0
  TX Packets 512-1023 Octets  0
```

```
TX Packets 1024-1518 Octets  0
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX Packets 9217-16383 Octets  0
TX Octets                   1
TX Multicast Packets        0
TX Broadcast Packets        1
TX PAUSEMAC Ctrl Frames     0
TX Oversize Packets         0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             72
RX Packets 64 Octets        0
RX Packets 65-127 Octets    0
RX Packets 128-255 Octets   0
RX Packets 256-511 Octets   0
RX Packets 512-1023 Octets  0
RX Packets 1024-1518 Octets  0
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Packets 9217-16383 Octets  0
RX Octets                   0
RX Multicast Packets        0
RX Broadcast Packets        0
RX FCS Errors               0
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             0
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

chassis (QFabric System) | 3303

show chassis environment fpc | 17621

show chassis ethernet-switch interconnect-device cb | 17960

show chassis fabric connectivity

IN THIS SECTION

- Syntax | 18022
- Description | 18022
- Options | 18022
- Required Privilege Level | 18023
- Output Fields | 18023
- Sample Output | 18023
- Release Information | 18031

Syntax

```
show chassis fabric connectivity (device | slot)
```

Description

(QFabric systems only) Display the status of the data plane connections in your QFabric system.

Options

none Display the status of all data plane connections in your QFabric system.

device Display the status of the data plane connections for a specific device.

slot Display the status of the data plane connections for a specific Flexible PIC Concentrator (FPC) slot on a specific device.

Required Privilege Level

admin

Output Fields

Table 248 on page 18023 lists the output fields for the `show chassis fabric connectivity` command. Output fields are listed in the approximate order in which they appear.

Table 248: show chassis fabric connectivity Output Fields

Field Name	Field Description
Device ID	Hardware serial identifier of the QFabric system component.
Type	Model number of the QFabric system component. Values include `qfxc08-3008` (QFX3008-I Interconnect device), `qfx3600-I` (QFX3600-I Interconnect device), and supported Node device names such as `qfx3500` (QFX3500 Node device) and `qfx3600-16q` (QFX3600 Node device).
Fabric: Incoming links	Displays inbound data plane (fte-) connections between Node devices and Interconnect devices, and their status (such as `Ok`).
Fabric: Outgoing links	Displays outbound data plane (fte-) connections between Node devices and Interconnect devices, and their status (such as `Ok`).

Sample Output

show chassis fabric connectivity

```
user@qfabric> show chassis fabric connectivity
Device ID: ED1487, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-0/0/1                   -> ED1487:fte-0/1/2                   Ok
```

```
  A0010:fte-3/0/1                   -> ED1487:fte-0/1/3                   Ok
  Fabric: Outgoing links:
  ED1487:fte-0/1/2                  -> A0010:fte-0/0/1                    Ok
  ED1487:fte-0/1/3                  -> A0010:fte-3/0/1                    Ok
Device ID: ED3683, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-0/0/2                   -> ED3683:fte-0/1/2                   Ok
  A0010:fte-3/0/2                   -> ED3683:fte-0/1/3                   Ok
  Fabric: Outgoing links:
  ED3683:fte-0/1/2                  -> A0010:fte-0/0/2                    Ok
  ED3683:fte-0/1/3                  -> A0010:fte-3/0/2                    Ok
Device ID: ED3705, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-0/0/0                   -> ED3705:fte-0/1/2                   Ok
  A0010:fte-3/0/0                   -> ED3705:fte-0/1/3                   Ok
  Fabric: Outgoing links:
  ED3705:fte-0/1/2                  -> A0010:fte-0/0/0                    Ok
  ED3705:fte-0/1/3                  -> A0010:fte-3/0/0                    Ok
Device ID: ED3707, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-0/0/8                   -> ED3707:fte-0/1/2                   Ok
  A0010:fte-3/0/8                   -> ED3707:fte-0/1/3                   Ok
  Fabric: Outgoing links:
  ED3707:fte-0/1/2                  -> A0010:fte-0/0/8                    Ok
  ED3707:fte-0/1/3                  -> A0010:fte-3/0/8                    Ok
Device ID: ED3711, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-0/0/9                   -> ED3711:fte-0/1/2                   Ok
  A0010:fte-3/0/9                   -> ED3711:fte-0/1/3                   Ok
  Fabric: Outgoing links:
  ED3711:fte-0/1/2                  -> A0010:fte-0/0/9                    Ok
  ED3711:fte-0/1/3                  -> A0010:fte-3/0/9                    Ok
Device ID: ED3702, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-0/0/10                  -> ED3702:fte-0/1/2                   Ok
  A0010:fte-3/0/10                  -> ED3702:fte-0/1/3                   Ok
  Fabric: Outgoing links:
  ED3702:fte-0/1/2                  -> A0010:fte-0/0/10                   Ok
  ED3702:fte-0/1/3                  -> A0010:fte-3/0/10                   Ok
Device ID: BBAK8737, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-0/0/11                  -> BBAK8737:fte-0/1/1                 Ok
  Fabric: Outgoing links:
```

```
  BBAK8737:fte-0/1/1                -> A0010:fte-0/0/11                   Ok
Device ID: BBAK8777, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-0/0/5                   -> BBAK8777:fte-0/1/0                 Ok
  A0010:fte-0/0/6                   -> BBAK8777:fte-0/1/1                 Ok
  Fabric: Outgoing links:
  BBAK8777:fte-0/1/0                -> A0010:fte-0/0/5                    Ok
  BBAK8777:fte-0/1/1                -> A0010:fte-0/0/6                    Ok
Device ID: BBAK8866, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-0/0/3                   -> BBAK8866:fte-0/1/0                 Ok
  A0010:fte-0/0/4                   -> BBAK8866:fte-0/1/1                 Ok
  Fabric: Outgoing links:
  BBAK8866:fte-0/1/0                -> A0010:fte-0/0/3                    Ok
  BBAK8866:fte-0/1/1                -> A0010:fte-0/0/4                    Ok
Device ID: BBAK8810, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-3/0/5                   -> BBAK8810:fte-0/1/0                 Ok
  A0010:fte-3/0/6                   -> BBAK8810:fte-0/1/1                 Ok
  Fabric: Outgoing links:
  BBAK8810:fte-0/1/0                -> A0010:fte-3/0/5                    Ok
  BBAK8810:fte-0/1/1                -> A0010:fte-3/0/6                    Ok
Device ID: BBAK8854, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-3/0/3                   -> BBAK8854:fte-0/1/0                 Ok
  A0010:fte-3/0/4                   -> BBAK8854:fte-0/1/1                 Ok
  Fabric: Outgoing links:
  BBAK8854:fte-0/1/0                -> A0010:fte-3/0/3                    Ok
  BBAK8854:fte-0/1/1                -> A0010:fte-3/0/4                    Ok
Device ID: BBAK8885, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-0/0/14                  -> BBAK8885:fte-0/1/0                 Ok
  A0010:fte-0/0/15                  -> BBAK8885:fte-0/1/1                 Ok
  Fabric: Outgoing links:
  BBAK8885:fte-0/1/0                -> A0010:fte-0/0/14                   Ok
  BBAK8885:fte-0/1/1                -> A0010:fte-0/0/15                   Ok
Device ID: BBAK8864, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-3/0/7                   -> BBAK8864:fte-0/1/0                 Ok
  A0010:fte-3/0/11                  -> BBAK8864:fte-0/1/1                 Ok
  Fabric: Outgoing links:
  BBAK8864:fte-0/1/0                -> A0010:fte-3/0/7                    Ok
  BBAK8864:fte-0/1/1                -> A0010:fte-3/0/11                   Ok
```

```
Device ID: BBAK8759, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-3/0/12                  -> BBAK8759:fte-0/1/0                 Ok
  A0010:fte-3/0/13                  -> BBAK8759:fte-0/1/1                 Ok
  Fabric: Outgoing links:
  BBAK8759:fte-0/1/0                -> A0010:fte-3/0/12                   Ok
  BBAK8759:fte-0/1/1                -> A0010:fte-3/0/13                   Ok
Device ID: BBAK8704, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-0/0/12                  -> BBAK8704:fte-0/1/0                 Ok
  A0010:fte-0/0/13                  -> BBAK8704:fte-0/1/1                 Ok
  Fabric: Outgoing links:
  BBAK8704:fte-0/1/0                -> A0010:fte-0/0/12                   Ok
  BBAK8704:fte-0/1/1                -> A0010:fte-0/0/13                   Ok
Device ID: BBAK8714, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-3/0/14                  -> BBAK8714:fte-0/1/0                 Ok
  A0010:fte-3/0/15                  -> BBAK8714:fte-0/1/1                 Ok
  Fabric: Outgoing links:
  BBAK8714:fte-0/1/0                -> A0010:fte-3/0/14                   Ok
  BBAK8714:fte-0/1/1                -> A0010:fte-3/0/15                   Ok
Device ID: A0010, Type: qfxc08-3008
  Front  Card 0   : Incoming links:
  ED3705:fte-0/1/2                  -> A0010:fte-0/0/0                    Ok
  ED1487:fte-0/1/2                  -> A0010:fte-0/0/1                    Ok
  ED3683:fte-0/1/2                  -> A0010:fte-0/0/2                    Ok
  BBAK8866:fte-0/1/0                -> A0010:fte-0/0/3                    Ok
  BBAK8866:fte-0/1/1                -> A0010:fte-0/0/4                    Ok
  BBAK8777:fte-0/1/0                -> A0010:fte-0/0/5                    Ok
  BBAK8777:fte-0/1/1                -> A0010:fte-0/0/6                    Ok
  ED3707:fte-0/1/2                  -> A0010:fte-0/0/8                    Ok
  ED3711:fte-0/1/2                  -> A0010:fte-0/0/9                    Ok
  ED3702:fte-0/1/2                  -> A0010:fte-0/0/10                   Ok
  BBAK8737:fte-0/1/1                -> A0010:fte-0/0/11                   Ok
  BBAK8704:fte-0/1/0                -> A0010:fte-0/0/12                   Ok
  BBAK8704:fte-0/1/1                -> A0010:fte-0/0/13                   Ok
  BBAK8885:fte-0/1/0                -> A0010:fte-0/0/14                   Ok
  BBAK8885:fte-0/1/1                -> A0010:fte-0/0/15                   Ok
  Front  Card 0   : Outgoing links:
  A0010:fte-0/0/8                   -> ED3707:fte-0/1/2                   Ok
  A0010:fte-0/0/9                   -> ED3711:fte-0/1/2                   Ok
  A0010:fte-0/0/10                  -> ED3702:fte-0/1/2                   Ok
  A0010:fte-0/0/11                  -> BBAK8737:fte-0/1/1                 Ok
```

```
A0010:fte-0/0/12                -> BBAK8704:fte-0/1/0                Ok
A0010:fte-0/0/13                -> BBAK8704:fte-0/1/1                Ok
A0010:fte-0/0/14                -> BBAK8885:fte-0/1/0                Ok
A0010:fte-0/0/15                -> BBAK8885:fte-0/1/1                Ok
A0010:fte-0/0/0                 -> ED3705:fte-0/1/2                  Ok
A0010:fte-0/0/1                 -> ED1487:fte-0/1/2                  Ok
A0010:fte-0/0/2                 -> ED3683:fte-0/1/2                  Ok
A0010:fte-0/0/3                 -> BBAK8866:fte-0/1/0                Ok
A0010:fte-0/0/4                 -> BBAK8866:fte-0/1/1                Ok
A0010:fte-0/0/5                 -> BBAK8777:fte-0/1/0                Ok
A0010:fte-0/0/6                 -> BBAK8777:fte-0/1/1                Ok
Front  Card 3   : Incoming links:
ED3705:fte-0/1/3                -> A0010:fte-3/0/0                   Ok
ED1487:fte-0/1/3                -> A0010:fte-3/0/1                   Ok
ED3683:fte-0/1/3                -> A0010:fte-3/0/2                   Ok
BBAK8854:fte-0/1/0              -> A0010:fte-3/0/3                   Ok
BBAK8854:fte-0/1/1              -> A0010:fte-3/0/4                   Ok
BBAK8810:fte-0/1/0              -> A0010:fte-3/0/5                   Ok
BBAK8810:fte-0/1/1              -> A0010:fte-3/0/6                   Ok
BBAK8864:fte-0/1/0              -> A0010:fte-3/0/7                   Ok
ED3707:fte-0/1/3                -> A0010:fte-3/0/8                   Ok
ED3711:fte-0/1/3                -> A0010:fte-3/0/9                   Ok
ED3702:fte-0/1/3                -> A0010:fte-3/0/10                  Ok
BBAK8864:fte-0/1/1              -> A0010:fte-3/0/11                  Ok
BBAK8759:fte-0/1/0              -> A0010:fte-3/0/12                  Ok
BBAK8759:fte-0/1/1              -> A0010:fte-3/0/13                  Ok
BBAK8714:fte-0/1/0              -> A0010:fte-3/0/14                  Ok
BBAK8714:fte-0/1/1              -> A0010:fte-3/0/15                  Ok
Front  Card 3   : Outgoing links:
A0010:fte-3/0/8                 -> ED3707:fte-0/1/3                  Ok
A0010:fte-3/0/9                 -> ED3711:fte-0/1/3                  Ok
A0010:fte-3/0/10                -> ED3702:fte-0/1/3                  Ok
A0010:fte-3/0/11                -> BBAK8864:fte-0/1/1                Ok
A0010:fte-3/0/12                -> BBAK8759:fte-0/1/0                Ok
A0010:fte-3/0/13                -> BBAK8759:fte-0/1/1                Ok
A0010:fte-3/0/14                -> BBAK8714:fte-0/1/0                Ok
A0010:fte-3/0/15                -> BBAK8714:fte-0/1/1                Ok
A0010:fte-3/0/0                 -> ED3705:fte-0/1/3                  Ok
A0010:fte-3/0/1                 -> ED1487:fte-0/1/3                  Ok
A0010:fte-3/0/2                 -> ED3683:fte-0/1/3                  Ok
A0010:fte-3/0/3                 -> BBAK8854:fte-0/1/0                Ok
A0010:fte-3/0/4                 -> BBAK8854:fte-0/1/1                Ok
A0010:fte-3/0/5                 -> BBAK8810:fte-0/1/0                Ok
```

```
  A0010:fte-3/0/6                   -> BBAK8810:fte-0/1/1                 Ok
  A0010:fte-3/0/7                   -> BBAK8864:fte-0/1/0                 Ok
```

show chassis fabric connectivity device

```
user@qfabric> show chassis fabric connectivity device BBAK8714
Device ID: BBAK8714, Type: qfx3500
  Fabric: Incoming links:
  A0010:fte-3/0/14                  -> BBAK8714:fte-0/1/0                 Ok
  A0010:fte-3/0/15                  -> BBAK8714:fte-0/1/1                 Ok
  Fabric: Outgoing links:
  BBAK8714:fte-0/1/0                -> A0010:fte-3/0/14                   Ok
  BBAK8714:fte-0/1/1                -> A0010:fte-3/0/15                   Ok
```

command-name

```
user@qfabric> show chassis fabric connectivity device A0010
Device ID: A0010, Type: qfxc08-3008
  Front  Card 0   : Incoming links:
  ED3705:fte-0/1/2                  -> A0010:fte-0/0/0                    Ok
  ED1487:fte-0/1/2                  -> A0010:fte-0/0/1                    Ok
  ED3683:fte-0/1/2                  -> A0010:fte-0/0/2                    Ok
  BBAK8866:fte-0/1/0                -> A0010:fte-0/0/3                    Ok
  BBAK8866:fte-0/1/1                -> A0010:fte-0/0/4                    Ok
  BBAK8777:fte-0/1/0                -> A0010:fte-0/0/5                    Ok
  BBAK8777:fte-0/1/1                -> A0010:fte-0/0/6                    Ok
  ED3707:fte-0/1/2                  -> A0010:fte-0/0/8                    Ok
  ED3711:fte-0/1/2                  -> A0010:fte-0/0/9                    Ok
  ED3702:fte-0/1/2                  -> A0010:fte-0/0/10                   Ok
  BBAK8737:fte-0/1/1                -> A0010:fte-0/0/11                   Ok
  BBAK8704:fte-0/1/0                -> A0010:fte-0/0/12                   Ok
  BBAK8704:fte-0/1/1                -> A0010:fte-0/0/13                   Ok
  BBAK8885:fte-0/1/0                -> A0010:fte-0/0/14                   Ok
  BBAK8885:fte-0/1/1                -> A0010:fte-0/0/15                   Ok
  Front  Card 0   : Outgoing links:
  A0010:fte-0/0/8                   -> ED3707:fte-0/1/2                   Ok
  A0010:fte-0/0/9                   -> ED3711:fte-0/1/2                   Ok
  A0010:fte-0/0/10                  -> ED3702:fte-0/1/2                   Ok
  A0010:fte-0/0/11                  -> BBAK8737:fte-0/1/1                 Ok
  A0010:fte-0/0/12                  -> BBAK8704:fte-0/1/0                 Ok
```

```
A0010:fte-0/0/13                -> BBAK8704:fte-0/1/1                Ok
A0010:fte-0/0/14                -> BBAK8885:fte-0/1/0                Ok
A0010:fte-0/0/15                -> BBAK8885:fte-0/1/1                Ok
A0010:fte-0/0/0                 -> ED3705:fte-0/1/2                  Ok
A0010:fte-0/0/1                 -> ED1487:fte-0/1/2                  Ok
A0010:fte-0/0/2                 -> ED3683:fte-0/1/2                  Ok
A0010:fte-0/0/3                 -> BBAK8866:fte-0/1/0                Ok
A0010:fte-0/0/4                 -> BBAK8866:fte-0/1/1                Ok
A0010:fte-0/0/5                 -> BBAK8777:fte-0/1/0                Ok
A0010:fte-0/0/6                 -> BBAK8777:fte-0/1/1                Ok
Front  Card 3   : Incoming links:
ED3705:fte-0/1/3                -> A0010:fte-3/0/0                   Ok
ED1487:fte-0/1/3                -> A0010:fte-3/0/1                   Ok
ED3683:fte-0/1/3                -> A0010:fte-3/0/2                   Ok
BBAK8854:fte-0/1/0              -> A0010:fte-3/0/3                   Ok
BBAK8854:fte-0/1/1              -> A0010:fte-3/0/4                   Ok
BBAK8810:fte-0/1/0              -> A0010:fte-3/0/5                   Ok
BBAK8810:fte-0/1/1              -> A0010:fte-3/0/6                   Ok
BBAK8864:fte-0/1/0              -> A0010:fte-3/0/7                   Ok
ED3707:fte-0/1/3                -> A0010:fte-3/0/8                   Ok
ED3711:fte-0/1/3                -> A0010:fte-3/0/9                   Ok
ED3702:fte-0/1/3                -> A0010:fte-3/0/10                  Ok
BBAK8864:fte-0/1/1              -> A0010:fte-3/0/11                  Ok
BBAK8759:fte-0/1/0              -> A0010:fte-3/0/12                  Ok
BBAK8759:fte-0/1/1              -> A0010:fte-3/0/13                  Ok
BBAK8714:fte-0/1/0              -> A0010:fte-3/0/14                  Ok
BBAK8714:fte-0/1/1              -> A0010:fte-3/0/15                  Ok
Front  Card 3   : Outgoing links:
A0010:fte-3/0/8                 -> ED3707:fte-0/1/3                  Ok
A0010:fte-3/0/9                 -> ED3711:fte-0/1/3                  Ok
A0010:fte-3/0/10                -> ED3702:fte-0/1/3                  Ok
A0010:fte-3/0/11                -> BBAK8864:fte-0/1/1                Ok
A0010:fte-3/0/12                -> BBAK8759:fte-0/1/0                Ok
A0010:fte-3/0/13                -> BBAK8759:fte-0/1/1                Ok
A0010:fte-3/0/14                -> BBAK8714:fte-0/1/0                Ok
A0010:fte-3/0/15                -> BBAK8714:fte-0/1/1                Ok
A0010:fte-3/0/0                 -> ED3705:fte-0/1/3                  Ok
A0010:fte-3/0/1                 -> ED1487:fte-0/1/3                  Ok
A0010:fte-3/0/2                 -> ED3683:fte-0/1/3                  Ok
A0010:fte-3/0/3                 -> BBAK8854:fte-0/1/0                Ok
A0010:fte-3/0/4                 -> BBAK8854:fte-0/1/1                Ok
A0010:fte-3/0/5                 -> BBAK8810:fte-0/1/0                Ok
A0010:fte-3/0/6                 -> BBAK8810:fte-0/1/1                Ok
```

```
  A0010:fte-3/0/7                   -> BBAK8864:fte-0/1/0                 Ok
  Ok
  A0010:fte-3/0/3                   -> BBAK8854:fte-0/1/0                 Ok
  A0010:fte-3/0/4                   -> BBAK8854:fte-0/1/1                 Ok
  A0010:fte-3/0/5                   -> BBAK8810:fte-0/1/0                 Ok
  A0010:fte-3/0/6                   -> BBAK8810:fte-0/1/1                 Ok
  A0010:fte-3/0/7                   -> BBAK8864:fte-0/1/0                 Ok
```

show chassis fabric connectivity device device-name slot

```
user@qfabric> show chassis fabric connectivity device A0010 slot 3
Device ID: A0010, Type: qfxc08-3008
  Front  Card 3   : Incoming links:
  ED3705:fte-0/1/3                  -> A0010:fte-3/0/0                    Ok
  ED1487:fte-0/1/3                  -> A0010:fte-3/0/1                    Ok
  ED3683:fte-0/1/3                  -> A0010:fte-3/0/2                    Ok
  BBAK8854:fte-0/1/0                -> A0010:fte-3/0/3                    Ok
  BBAK8854:fte-0/1/1                -> A0010:fte-3/0/4                    Ok
  BBAK8810:fte-0/1/0                -> A0010:fte-3/0/5                    Ok
  BBAK8810:fte-0/1/1                -> A0010:fte-3/0/6                    Ok
  BBAK8864:fte-0/1/0                -> A0010:fte-3/0/7                    Ok
  ED3707:fte-0/1/3                  -> A0010:fte-3/0/8                    Ok
  ED3711:fte-0/1/3                  -> A0010:fte-3/0/9                    Ok
  ED3702:fte-0/1/3                  -> A0010:fte-3/0/10                   Ok
  BBAK8864:fte-0/1/1                -> A0010:fte-3/0/11                   Ok
  BBAK8759:fte-0/1/0                -> A0010:fte-3/0/12                   Ok
  BBAK8759:fte-0/1/1                -> A0010:fte-3/0/13                   Ok
  BBAK8714:fte-0/1/0                -> A0010:fte-3/0/14                   Ok
  BBAK8714:fte-0/1/1                -> A0010:fte-3/0/15                   Ok
  Front  Card 3   : Outgoing links:
  A0010:fte-3/0/8                   -> ED3707:fte-0/1/3                   Ok
  A0010:fte-3/0/9                   -> ED3711:fte-0/1/3                   Ok
  A0010:fte-3/0/10                  -> ED3702:fte-0/1/3                   Ok
  A0010:fte-3/0/11                  -> BBAK8864:fte-0/1/1                 Ok
  A0010:fte-3/0/12                  -> BBAK8759:fte-0/1/0                 Ok
  A0010:fte-3/0/13                  -> BBAK8759:fte-0/1/1                 Ok
  A0010:fte-3/0/14                  -> BBAK8714:fte-0/1/0                 Ok
  A0010:fte-3/0/15                  -> BBAK8714:fte-0/1/1                 Ok
  A0010:fte-3/0/0                   -> ED3705:fte-0/1/3                   Ok
  A0010:fte-3/0/1                   -> ED1487:fte-0/1/3                   Ok
  A0010:fte-3/0/2                   -> ED3683:fte-0/1/3                   Ok
```

```
A0010:fte-3/0/3                    -> BBAK8854:fte-0/1/0                  Ok
A0010:fte-3/0/4                    -> BBAK8854:fte-0/1/1                  Ok
A0010:fte-3/0/5                    -> BBAK8810:fte-0/1/0                  Ok
A0010:fte-3/0/6                    -> BBAK8810:fte-0/1/1                  Ok
A0010:fte-3/0/7                    -> BBAK8864:fte-0/1/0                  Ok
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

request chassis fabric fpc | **14465**

show chassis fabric device | **18066**

Understanding Interconnect Devices

show chassis fabric degradation

IN THIS SECTION

- Syntax | **18031**
- Description | **18032**
- Required Privilege Level | **18032**
- Sample Output | **18032**

Syntax

```
show chassis fabric degradation
```

Description

(PTX10004, PTX10008, PTX10016, and PTX10002-36QDD) Display the state of fabric degradation for all FPCs.

Required Privilege Level

view

The following table lists the output fields for the `show chassis fabric degradation` command. Output fields are listed in the approximate order in which they appear.

Table 249: show chassis fabric degradation Output Fields

Field Name	Field Description
FPC	FPC number.
State	FPC status: Online, Offline, or Empty.
Reqd/Curr Planes	Required number of functioning planes / Current number of functioning planes.
Configured Degrad(%)	Configured degradation percentage.
Time Last Action Initiated	Time stamp when last FPC action was initiated.

Sample Output

show chassis fabric degradation

```
 user@root> show chassis fabric degradation
                 Reqd/Curr  Configured          Current     Time Last
FPC    State     Planes     Degrad(%),action    Degrad(%)   Action Initiated
0      Online    21/15      n/a,none            28
1      Online    21/15      n/a,none            28
2      Online    21/15      n/a,none            28
```

```
3       Online     21/15       n/a,none            28
4       Empty
```

RELATED DOCUMENTATION

show chassis fabric degradation actions | 18033

show chassis fabric degradation actions

IN THIS SECTION

- Syntax (PTX10004, PTX10008, PTX10016) | 18033
- Description | 18033
- Options | 18033
- Required Privilege Level | 18034
- Sample Output | 18034

Syntax (PTX10004, PTX10008, PTX10016)

```
show chassis fabric degradation actions
```

Description

Display the state of fabric degradation actions for degraded FPCs.

Options

actions Displays information about the fabric degradation actions for all FPCs.

Required Privilege Level

view

The following table lists the output fields for the `show chassis fabric degradation actions` command. Output fields are listed in the approximate order in which they appear.

Table 250: show chassis fabric degradation action Output Fields

Field Name	Field Description
FPC	Disabled or degraded FPC number.
Time Action Started	Time stamp when last FPC action was initiated.
Last Action	Last FPC action.
Reason	Reason for FPC degradation or FPC offline.
Total Actions	Total number of actions initiated on the FPC
Time Action Completed	Time stamp when last FPC action was completed.

Sample Output

show chassis fabric degradation actions

```
 user@root> show chassis fabric degradation actions

FPC Time Action              Last      Reason       Total    Time Action
    Started                  Action                 Actions  Completed

 0  2020-10-07 10:32:21 PDT  offline   unreachable     2     2020-10-07 10:33:53 PDT
```

RELATED DOCUMENTATION

show chassis fabric degradation | 18031

show chassis fabric degraded-fabric-reachability

IN THIS SECTION

- Syntax | 18035
- Description | 18035
- Additional Information | 18035
- Required Privilege Level | 18035
- Output Fields | 18036
- Sample Output | 18036
- Release Information | 18037

Syntax

```
show chassis fabric degraded-fabric-reachability
```

Description

Display the current state of reachability between the Packet Forwarding Engines in the system.

Additional Information

Required Privilege Level

view

Output Fields

Table 251 on page 18036 lists the output fields for the `show chassis fabric degraded-fabric-reachability` command. Output fields are listed in the approximate order in which they appear.

Table 251: show chassis fabric degraded-fabric-reachability Output Fields

Field Name	Field Description
FPC	Display fabric reachability for the displayed FPC slot. PTX10016 supports 16 FPCs and output detail for each FPC slot.
PFE	Display fabric reachability for the displayed PFE slot on a per SIB and plane basis.
SIBx_Plane y	Display the SIB (x) and plane (y) where link errors occurred.
Link errors FPC/PFEs	Display the list of FPC and PFE slots that are unreachable for the displayed SIB and plane due to link errors.

Sample Output

show chassis fabric degraded-fabric-reachability

```
user@host> show chassis fabric degraded-fabric-reachability
Degraded Fabric reachability Information:
FPC #0
    PFE #0
      SIB0_Plane 0
           Link errors  FPC/PFEs    2/0 5/0 5/1 5/2 5/3
      SIB0_Plane 1
           Link errors  FPC/PFEs    2/0 5/0
    PFE #1
      SIB0_Plane 0
           Link errors  FPC/PFEs    2/0 5/0 5/1 5/2 5/3
      SIB0_Plane 1
           Link errors  FPC/PFEs    2/0 5/0
```

Release Information

Command introduced in Junos OS Release 12.1X48R4.

RELATED DOCUMENTATION

show chassis fabric errors | 18068

show chassis fabric reachability | 18168

degraded | 5236

show chassis fabric destinations

IN THIS SECTION

- Syntax | 18037
- Syntax (MX240, MX480, MX960 , MX2010, MX2020, MX10003, MX10008, and MX2008 Universal Routing Platforms) | 18038
- Description | 18038
- Options | 18038
- Required Privilege Level | 18038
- Output Fields | 18038
- Sample Output | 18039
- Release Information | 18066

Syntax

```
show chassis fabric destinations
```

Syntax (MX240, MX480, MX960 , MX2010, MX2020, MX10003, MX10008, and MX2008 Universal Routing Platforms)

```
show chassis fabric destinations fpc <fpc-slot-number>
<extended>
```

Description

Display the state of fabric destinations for all FPCs.

Options

none	Display information about the fabric destinations of all FPCs.
fpc-slot-number	(Optional) Display information about the specified FPC. For MX2020 routers, replace *fpc-slot-number* with a value from 0 through 19. For MX2010 and MX2008 routers, replace *fpc-slot-number* with a value from 0 through 9. For MX10003, replace *fpc-slot-number* with a value from 0 through 1.
extended	(Optional) (MX2020, MX2010, and MX2008 routers with SFB2 only) Display information about the fabric destination of all 24 planes for each FPC.

Required Privilege Level

view

Output Fields

Table 252 on page 18039 lists the output fields for the `show chassis fabric destinations` command. Output fields are listed in the approximate order in which they appear.

Table 252: show chassis fabric destinations Output Fields

Field Name	Field Description
Fabric destinations state	Indicates the state of the fabric destinations: • 0—Destination is non-existent. • 2—Destination is enabled. • 3—Destination is disabled. • 6—Destination is in erroneous state and is disabled.
Flexible PIC Concentrator (FPC) number	Source FPC number.
Packet Forwarding Engine number	Source Packet Forwarding Engine number.
Plane number	Source plane number.

Sample Output

show chassis fabric destinations fpc 1 (MX240 Router)

In the output, the values followed by the plane number denote multiple quadruples. The first quadruple specifies FPC1, the second quadrapule specifies FPC2 and so on. Each quadruple specifies the states of the fabric plane to the Packet Forwarding Engines.

```
user@host> show chassis fabric destinations
fpc 1

    Fabric destinations state:
       0: non-existent
       2: enabled
       3: disabled
       6: dest-err and disabled

    FPC 1
```

```
     PFE 0
      Plane 0   0000 3300 3333
      Plane 1   0000 2200 2222
      Plane 2   0000 2200 2222
      Plane 3   0000 2200 2222
      Plane 4   0000 2200 2222
      Plane 5   0000 3300 3333
      Plane 6   0000 3300 3333
      Plane 7   0000 3300 3333
     PFE 1
      Plane 0   0000 3300 3333
      Plane 1   0000 2200 2222
      Plane 2   0000 2200 2222
      Plane 3   0000 2200 2222
      Plane 4   0000 2200 2222
      Plane 5   0000 3300 3333
      Plane 6   0000 3300 3333
      Plane 7   0000 3300 3333
```

show chassis fabric destinations (MX10008 Router)

```
user@host> show chassis fabric destinations

Fabric destinations state:
   0: non-existent
   2: enabled
   3: disabled
   6: dest-err and disabled

FPC 1
 PFE 0
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
 PFE 1
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
```

```
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
 PFE 2
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
 PFE 3
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
 PFE 4
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
 PFE 5
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
FPC 5
 PFE 0
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
 PFE 1
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
```

```
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
 PFE 2
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
 PFE 3
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
 PFE 4
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
 PFE 5
  Plane 0   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 1   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 2   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 3   000000 222222 000000   000000 000000 222222   000000 000000
  Plane 4   000000 333333 000000   000000 000000 333333   000000 000000
  Plane 5   000000 333333 000000   000000 000000 333333   000000 000000
```

show chassis fabric destinations fpc 4 (MX480 Router with MPC4E)

```
user@host > show chassis fabric destinations fpc 4
Fabric destinations state:
   0: non-existent
   2: enabled
   3: disabled
   6: dest-err and disabled
```

```
FPC 4
 PFE 0
  Plane 0   2200 2222 0000   2000 2200 0000
  Plane 1   2200 2222 0000   2000 2200 0000
  Plane 2   2200 2222 0000   2000 2200 0000
  Plane 3   2200 2222 0000   2000 2200 0000
  Plane 4   3300 3333 0000   3000 3300 0000
  Plane 5   3300 3333 0000   3000 3300 0000
  Plane 6   3300 3333 0000   3000 3300 0000
  Plane 7   3300 3333 0000   3000 3300 0000
 PFE 1
  Plane 0   2200 2222 0000   2000 2200 0000
  Plane 1   2200 2222 0000   2000 2200 0000
  Plane 2   2200 2222 0000   2000 2200 0000
  Plane 3   2200 2222 0000   2000 2200 0000
  Plane 4   3300 3333 0000   3000 3300 0000
  Plane 5   3300 3333 0000   3000 3300 0000
  Plane 6   3300 3333 0000   3000 3300 0000
  Plane 7   3300 3333 0000   3000 3300 0000
```

show chassis fabric destinations (MX960 Router)

```
user@host> show chassis fabric destinations

    Fabric destinations state:
       0: non-existent
       2: enabled
       3: disabled
       6: dest-err and disabled

    FPC 1
     PFE 0
      Plane 0   0000 3300 3333
      Plane 1   0000 2200 2222
      Plane 2   0000 2200 2222
      Plane 3   0000 2200 2222
      Plane 4   0000 2200 2222
      Plane 5   0000 3300 3333
      Plane 6   0000 3300 3333
      Plane 7   0000 3300 3333
     PFE 1
```

```
    Plane 0   0000 3300 3333
    Plane 1   0000 2200 2222
    Plane 2   0000 2200 2222
    Plane 3   0000 2200 2222
    Plane 4   0000 2200 2222
    Plane 5   0000 3300 3333
    Plane 6   0000 3300 3333
    Plane 7   0000 3300 3333
  FPC 2
   PFE 0
    Plane 0   0000 3300 3333
    Plane 1   0000 2200 2222
    Plane 2   0000 2200 2222
    Plane 3   0000 2200 2222
    Plane 4   0000 2200 2222
    Plane 5   0000 3300 3333
    Plane 6   0000 3300 3333
    Plane 7   0000 3300 3333
   PFE 1
    Plane 0   0000 3300 3333
    Plane 1   0000 2200 2222
    Plane 2   0000 2200 2222
    Plane 3   0000 2200 2222
    Plane 4   0000 2200 2222
    Plane 5   0000 3300 3333
    Plane 6   0000 3300 3333
    Plane 7   0000 3300 3333
   PFE 2
    Plane 0   0000 3300 3333
    Plane 1   0000 2200 2222
    Plane 2   0000 2200 2222
    Plane 3   0000 2200 2222
    Plane 4   0000 2200 2222
    Plane 5   0000 3300 3333
    Plane 6   0000 3300 3333
    Plane 7   0000 3300 3333
   PFE 3
    Plane 0   0000 3300 3333
    Plane 1   0000 2200 2222
    Plane 2   0000 2200 2222
    Plane 3   0000 2200 2222
    Plane 4   0000 2200 2222
    Plane 5   0000 3300 3333
```

```
      Plane 6   0000 3300 3333
      Plane 7   0000 3300 3333
```

show chassis fabric destinations fpc 1 (MX2020 Router)

```
user@host> show chassis fabric destinations
fpc 1
Fabric destinations state:
   0: non-existent
   2: enabled
   3: disabled
   6: dest-err and disabled

FPC 1
 PFE 0
  Plane 0   3333 3333 3333   3333 3333 3333   3333 3333 3333   3333 3333 3333   3333 3333 3333
3333 3333 3333   3333 3333
  Plane 1   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 2   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 3   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 4   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 5   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 6   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 7   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
 PFE 1
  Plane 0   3333 3333 3333   3333 3333 3333   3333 3333 3333   3333 3333 3333   3333 3333 3333
3333 3333 3333   3333 3333
  Plane 1   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 2   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 3   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
```

```
  Plane 4   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 5   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 6   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 7   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
 PFE 2
  Plane 0   3333 3333 3333   3333 3333 3333   3333 3333 3333   3333 3333 3333   3333 3333 3333
3333 3333 3333   3333 3333
  Plane 1   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 2   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 3   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 4   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 5   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 6   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 7   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
 PFE 3
  Plane 0   3333 3333 3333   3333 3333 3333   3333 3333 3333   3333 3333 3333   3333 3333 3333
3333 3333 3333   3333 3333
  Plane 1   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 2   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 3   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 4   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 5   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 6   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
  Plane 7   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222   2222 2222 2222
2222 2222 2222   2222 2222
```

show chassis fabric destinations extended (MX2020 Router with SFB2)

```
user@host> show chassis fabric destinations extended
Fabric destinations state:
   0: non-existent
   2: enabled
   3: disabled
   6: dest-err and disabled

FPC 0
 PFE 0
  Plane 0   2000 2222 2200 2200 0000 0000 2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 1   2000 2222 2200 2200 0000 0000 2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 2   2000 2222 2200 2200 0000 0000 2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 3   6000 3333 3300 3300 0000 0000 3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 4   2000 2222 2200 2200 0000 0000 2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 5   2000 2222 2200 2200 0000 0000 2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 6   2000 2222 2200 2200 0000 0000 2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 7   2000 2222 2200 2200 0000 0000 2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 8   6000 3333 3300 3300 0000 0000 3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 9   6000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 10  6000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 11  6000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 12  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 13  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 14  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 15  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
```

```
0000 2200 2200   0000 2222
  Plane 16  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 17  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 18  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 19  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 20  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 21  6000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 22  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 23  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
FPC 1
 PFE 0
  Plane 0   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 1   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 2   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 3   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 2200 2200   0000 2222
  Plane 4   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 2200 2200   0000 2222
  Plane 5   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 2200 2200   0000 2222
  Plane 6   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 7   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 8   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333 Plane 9   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000
0000 0000   3300 3300 0000   0000 3300 3300   0000 3333
  Plane 10  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 11  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 12  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
```

```
0000 2200 2200   0000 2222
  Plane 13  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 14  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 15  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 16  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 17  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 18  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 19  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 20  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 21  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 22  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 23  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
 PFE 1
  Plane 0   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 1   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 2   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 3   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 2200 2200   0000 2222
  Plane 4   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 2200 2200   0000 2222
  Plane 5   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 2200 2200   0000 2222
  Plane 6   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 7   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 8   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 9   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
```

```
0000 3300 3300   0000 3333
  Plane 10  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 11  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 12  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 13  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 14  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 15  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 16  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 17  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 18  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 19  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 20  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 21  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 22  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 23  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
 PFE 2
  Plane 0   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 1   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 2   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 3   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 2200 2200   0000 2222 Plane 4   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000
0000 0000   3300 3300 0000   0000 2200 2200   0000 2222
  Plane 5   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 2200 2200   0000 2222
  Plane 6   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
```

```
  Plane 7   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 8   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 9   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 10  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 11  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 12  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 13  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 14  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 15  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 16  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 17  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 18  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 19  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 20  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 21  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 22  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 23  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
 PFE 3
  Plane 0   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 1   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 2   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 3   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 2200 2200   0000 2222
```

```
  Plane 4   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 2200 2200   0000 2222
  Plane 5   3000 3333 3300   2200 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 2200 2200   0000 2222
  Plane 6   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 7   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 8   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 9   3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 10  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 11  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 12  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 13  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000
2222
  Plane 14  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 15  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 16  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 17  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 18  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 19  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 20  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 21  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 22  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 23  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
....
FPC 19
```

```
 PFE 0
  Plane 0   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 1   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 2   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 3   3000 2222 2200   2200 0000 0000   3000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 4   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 5   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 6   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 7   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 8   3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 9   3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 10  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 11  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 12  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 13  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 14  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 15  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 16  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 17  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 18  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 19  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 20  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
```

```
  Plane 21  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 22  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 23  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
 PFE 1
  Plane 0   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 1   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 2   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 3   3000 2222 2200   2200 0000 0000   3000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 4   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 5   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 6   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 7   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 8   3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 9   3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 10  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 11  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 12  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 13  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 14  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 15  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 16  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 17  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
```

```
  Plane 18  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 19  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 20  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 21  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 22  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 23  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
PFE 2
  Plane 0   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 1   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 2   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 3   3000 2222 2200   2200 0000 0000   3000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 4   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 5   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 6   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 7   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 8   3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 9   3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 10  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 11  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 12  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 13  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 14  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
```

```
  Plane 15  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 16  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 17  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 18  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 19  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 20  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 21  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 22  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 23  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
 PFE 3
  Plane 0   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 1   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 2   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 3   3000 2222 2200   2200 0000 0000   3000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 4   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 5   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 6   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 7   2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 8   3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 9   3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 10  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
  Plane 11  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3200 0000
0000 2200 2200   0000 2222
```

```
  Plane 12  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 13  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 14  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 15  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 16  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 17  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 18  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
 Plane 19  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 20  3000 3333 3300   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 3300 0000
0000 3300 3300   0000 3333
  Plane 21  3000 2222 2200   3300 0000 0000   3000 3000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 22  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
  Plane 23  2000 2222 2200   2200 0000 0000   2000 2000 0000   0000 0000 0000   3300 2200 0000
0000 2200 2200   0000 2222
```

show chassis fabric destinations (MX2010 Router)

```
user@host> show chassis fabric destinations
Fabric destinations state:
   0: non-existent
   2: enabled
   3: disabled
   6: dest-err and disabled

FPC 0
 PFE 0
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
```

```
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
 PFE 1
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
FPC 1
 PFE 0
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
FPC 2
 PFE 0
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
 PFE 1
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200

FPC 3
```

```
 PFE 0
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
 PFE 1
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
 PFE 2
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
 PFE 3
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
FPC 4
 PFE 0
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
```

```
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
FPC 5
 PFE 0
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
 PFE 1
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
FPC 6
 PFE 0
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
 PFE 1
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
 PFE 2
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
```

```
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
 PFE 3
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
FPC 7
 PFE 0
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
 PFE 1
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
FPC 8
 PFE 0
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
```

```
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
FPC 9
 PFE 0
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
 PFE 1
  Plane 0   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 1   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 2   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 3   3300 3000 3300   3333 3000 3300   3333 3300 3000   3300
  Plane 4   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 5   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 6   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
  Plane 7   2200 2000 2200   2222 2000 2200   2222 2200 2000   2200
```

show chassis fabric destinations (MX10003 Router)

```
user@host> show chassis fabric destinations

Fabric destinations state:
   0: non-existent
   2: enabled
   3: disabled
   6: dest-err and disabled

FPC 0
 PFE 0
  Plane 0   2220 2220
  Plane 1   2220 2220
  Plane 2   2220 2220
  Plane 3   2220 2220
  Plane 4   2220 2220
  Plane 5   2220 2220
  Plane 6   2220 2220
```

```
  Plane 7   2220 2220
  Plane 8   2220 2220
  Plane 9   2220 2220
  Plane 10  2220 2220
  Plane 11  2220 2220
  Plane 12  2220 2220
  Plane 13  2220 2220
  Plane 14  2220 2220
  Plane 15  2220 2220
  Plane 16  2220 2220
  Plane 17  2220 2220
  Plane 18  2220 2220
  Plane 19  2220 2220
  Plane 20  2220 2220
  Plane 21  2220 2220
 PFE 1
  Plane 0   2220 2220
  Plane 1   2220 2220
  Plane 2   2220 2220
  Plane 3   2220 2220
  Plane 4   2220 2220
  Plane 5   2220 2220
  Plane 6   2220 2220
  Plane 7   2220 2220
  Plane 8   2220 2220
  Plane 9   2220 2220
  Plane 10  2220 2220
  Plane 11  2220 2220
  Plane 12  2220 2220
  Plane 13  2220 2220
  Plane 14  2220 2220
  Plane 15  2220 2220
  Plane 16  2220 2220
  Plane 17  2220 2220
  Plane 18  2220 2220
  Plane 19  2220 2220
  Plane 20  2220 2220
  Plane 21  2220 2220
 PFE 2
  Plane 0   2220 2220
  Plane 1   2220 2220
  Plane 2   2220 2220
  Plane 3   2220 2220
```

```
  Plane 4   2220 2220
  Plane 5   2220 2220
  Plane 6   2220 2220
  Plane 7   2220 2220
  Plane 8   2220 2220
  Plane 9   2220 2220
  Plane 10  2220 2220
  Plane 11  2220 2220
  Plane 12  2220 2220
  Plane 13  2220 2220
  Plane 14  2220 2220
  Plane 15  2220 2220
  Plane 16  2220 2220
  Plane 17  2220 2220
  Plane 18  2220 2220
  Plane 19  2220 2220
  Plane 20  2220 2220
  Plane 21  2220 2220
FPC 1
 PFE 0
  Plane 0   2220 2220
  Plane 1   2220 2220
  Plane 2   2220 2220
  Plane 3   2220 2220
  Plane 4   2220 2220
  Plane 5   2220 2220
  Plane 6   2220 2220
  Plane 7   2220 2220
  Plane 8   2220 2220
  Plane 9   2220 2220
  Plane 10  2220 2220
  Plane 11  2220 2220
  Plane 12  2220 2220
  Plane 13  2220 2220
  Plane 14  2220 2220
  Plane 15  2220 2220
  Plane 16  2220 2220
  Plane 17  2220 2220
  Plane 18  2220 2220
  Plane 19  2220 2220
  Plane 20  2220 2220
  Plane 21  2220 2220
 PFE 1
```

```
  Plane 0   2220 2220
  Plane 1   2220 2220
  Plane 2   2220 2220
  Plane 3   2220 2220
  Plane 4   2220 2220
  Plane 5   2220 2220
  Plane 6   2220 2220
  Plane 7   2220 2220
  Plane 8   2220 2220
  Plane 9   2220 2220
  Plane 10  2220 2220
  Plane 11  2220 2220
  Plane 12  2220 2220
  Plane 13  2220 2220
  Plane 14  2220 2220
  Plane 15  2220 2220
  Plane 16  2220 2220
  Plane 17  2220 2220
  Plane 18  2220 2220
  Plane 19  2220 2220
  Plane 20  2220 2220
  Plane 21  2220 2220
 PFE 2
  Plane 0   2220 2220
  Plane 1   2220 2220
  Plane 2   2220 2220
  Plane 3   2220 2220
  Plane 4   2220 2220
  Plane 5   2220 2220
  Plane 6   2220 2220
  Plane 7   2220 2220
  Plane 8   2220 2220
  Plane 9   2220 2220
  Plane 10  2220 2220
  Plane 11  2220 2220
  Plane 12  2220 2220
  Plane 13  2220 2220
  Plane 14  2220 2220
  Plane 15  2220 2220
  Plane 16  2220 2220
  Plane 17  2220 2220
  Plane 18  2220 2220
  Plane 19  2220 2220
```

```
Plane 20  2220 2220
Plane 21  2220 2220
```

Release Information

Command introduced in Junos OS Release 12.1 for MX240, MX480, and MX960 routers.

Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 Universal Routing Platforms.

`extended` option introduced in Junos OS Release 16.1 for MX2010 and MX2020 Routers.

show chassis fabric device

IN THIS SECTION

- Syntax | 18066
- Description | 18066
- Options | 18067
- Required Privilege Level | 18067
- Output Fields | 18067
- Sample Output | 18068
- Release Information | 18068

Syntax

```
show chassis fabric device device-name
```

Description

(QFabric systems only) Display the fabric management status of devices in your QFabric system.

Options

none Display the fabric management status for all devices in your QFabric system.

device-name Display the fabric management status for a specific device in your QFabric system. You can enter either the alias name or the serial number of the device.

Required Privilege Level

admin

Output Fields

Table 253 on page 18067 lists the output fields for the `show chassis fabric device` command. Output fields are listed in the approximate order in which they appear.

Table 253: show chassis fabric device Output Fields

Field Name	Field Description
Device ID	Hardware serial identifier of the QFabric system component.
Type	Model number of the QFabric system component. Values include `qfxc08-3008` (QFX3008-I Interconnect device), `qfx3600-I` (QFX3600-I Interconnect device), and supported Node device names such as `qfx3500` (QFX3500 Node device) and `qfx3600-16q` (QFX3600 Node device).
Management status	Displays keepalive status for fabric management processes on a specific device. Values include `On` and `Off`.
Hardware status	Displays operational status of the device participating in fabric management (such as `Ok`).

Sample Output

show chassis fabric device

```
user@qfabric> show chassis fabric device

Device ID: node2, Type: qfx3500
Management status: On, Hardware status: Ok

Device ID: node3, Type: qfx3500
Management status: On, Hardware status: Ok

Device ID: node0, Type: qfx3500
Management status: On, Hardware status: Ok

Device ID: node1, Type: qfx3500
Management status: On, Hardware status: Ok
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

- request chassis fabric fpc | 14465
- show chassis fabric connectivity | 18022
- show fabric administration inventory | 19878
- *Understanding Interconnect Devices*

show chassis fabric errors

IN THIS SECTION

- Syntax | 18069

- Syntax (PTX Series Packet Transport Routers) | 18069
- Description | 18069
- Options | 18069
- Required Privilege Level | 18071
- Output Fields | 18071
- Sample Output | 18073
- Release Information | 18075

Syntax

```
show chassis fabric errors
<autoheal>
<fpc slot-number lcc number>
<sib (slot | f13 sib-slot | f2s sib-slot/sib-f2s-slot-number | lcc number)>
```

Syntax (PTX Series Packet Transport Routers)

```
show chassis fabric errors
(autoheal | fpc slot-number | sib sib-slot)
```

Description

Display the first ten and last ten fabric errors for the FPC or Switch Interface Boards (SIBs).

NOTE: This command can only be issued on a primary Routing Engine.

Options

autoheal (TX Matrix Plus routers and PTX Series Packet Transport Routers only) Show an error log of the first 100 autoheal actions taken on the system.

fpc *slot-number* Show error log of the first ten and last ten errors for the specified FPC.

(PTX3000 routers only)—Replace *slot-number* with an FPC slot number: **0**, **2**, **4**, **6**, **8**, **10**, **12**, or **14**.

(PTX5000 routers only)—Replace *slot-number* with a value from **0** through **7**.

(TX Matrix Plus routers only)—Replace `fpc slot-number` with the following values depending on the LCC configuration:

- On a TX Matrix Plus router with the TXP-T1600 configuration, if you specify the number of a T1600 LCC by using the `lcc number` option (the recommended method), replace `fpc slot-number` with a value from `0` through `7`. Otherwise, use a value from `0` through `31`.
- On a TX Matrix Plus router with the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 LCC by using the `lcc number` option (the recommended method), replace `fpc slot-number` with a value from `0` through `7`. Otherwise, use a value from `0` through `63`.
- `lcc number`—Show error log of the first ten and last ten errors for the specified FPC on a specific network device (line-card chassis) that is part of the routing matrix.

  Replace `lcc number` with the following values depending on the LCC configuration:

  - 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
  - 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
  - 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
  - 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

  If you specify the number of the network device by using only the `lcc number` option (the recommended method), replace *slot-number* with a value from **0** through **7**. Otherwise, replace *slot-number* with a value from **0** through **31**. For example, the following commands have the same result:

```
user@host> show chassis fabric errors fpc 1 lcc 1
user@host> show chassis fabric errors fpc 9
```

sib

Show error log of the first ten and last ten errors for the specified SIB. This option has the following suboptions:

- (TX Matrix Plus routers only)*sib-slot*—Specify a value ranging from **0** through **4**.
- (PTX3000 and PTX5000 routers) *sib-slot*—Specify a value ranging from **0** through **8**.
- (TX Matrix Plus routers only) `f13` *`sib-slot`*—(Optional) Show SIB F13 errors. Specify a valid SIB value number: `0`, `1`, `3`, `4`, `6`, `7`, `8`, `9`, `11`, or `12`.
- (TX Matrix Plus routers only) `f2s` *`sib-slot`*/*`sib-f2s-slot-number`*—(Optional) Show SIB F2S errors. Replace *sib-slot* with a value from `0` through `4`, followed by a *sib-f2s-slot-number* value `0`, `2`, `4` or `6`.
- (TX Matrix Plus routers only) `lcc` *`number`*—(Optional) Show error log of the first ten and last ten SIB errors for the specified network device (line-card chassis).

  Replace *number* with the following values depending on the LCC configuration:

  - 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
  - 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
  - 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
  - 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

  NOTE: The `lcc` *`number`* suboption is mandatory when using the following format for the command: `show chassis fabric errors sib lcc` *`number`* `sib` *`slot-number`*. For instance, issuing `show chassis fabric errors sib lcc` *`2 3`* displays errors detected on LCC 2, SIB 3.

  This suboption is not required when the `f13` or `f2s` suboptions are used with the `sib` *`slot-number`* option.

Required Privilege Level

view

Output Fields

Table 254 on page 18072 lists the output fields for the `show chassis fabric errors` command. Output fields are listed in the approximate order in which they appear.

Table 254: show chassis fabric errors Output Fields

Field Name	Field Description
`Time`	Time the error was logged. (TX Matrix Plus routers and PTX Series Packet Transport Routers only) For the `autoheal` option, shows the timestamp when autoheal was attempted on a SIB that was in fault state.
`Error log of first 10 errors`	List of the first ten errors.
`Error log of last 10 errors`	List of the last ten errors.
`Error log of first 100 errors`	Indicates the autoheal action taken on the SIB. The following actions can occur: • Req—A SIB autoheal request was made on a faulty SIB. • Action—Autohealing (taking the SIB offline and then online) is initiated. • Denied—Autohealing (taking the SIB offline and then online) is denied because the SIB went to a fault state before the autoheal configuration period completed. • Set info—Setting information to force skipping autoheal on the SIB so that no further attempts to autoheal the faulty SIB are made. • Clear info— if a user takes a SIB offline and then online, then the autoheal information of the SIB is cleared. If the SIB goes to a fault state, autoheal is attempted on the SIB. • (PTX3000 and PTX5000 routers only) Completed—Autohealing is completed. This message is displayed whether autohealing is succesful or not.
`fpc slot number`	(PTX5000 Packet Transport Router only)—Range is 0 through 7.
`sib slot number`	(PTX Series Packet Transport Routers only)—Range is 0 through 8.
`lcc number`	Not supported on PTX Series Packet Transport Routers.

Sample Output

show chassis fabric errors (F13 SIB Errors on a TX Matrix Plus Router)

```
user@host> show chassis fabric errors sib f13
11Can you provide CLI output for QFX,
depending on which options are available?

Time                          Error log of first 10 errors
2009-10-06 02:21:17 PDT       LOS on Cable-D(1,0)
```

show chassis fabric errors (F2S SIB Errors on a TX Matrix Plus Router)

```
user@host> show chassis fabric errors sib f2s
0/0

Time                          Error log of first 10 errors
2009-10-06 13:51:42 PDT       Cell drop errors on CLOS F2 SF 0 Port 0 link
```

show chassis fabric errors (SIB Errors Specific to an LCC Connected to a TX Matrix Plus Router)

```
user@host> show chassis fabric errors sib 1
lcc 0
lcc0-re0:

--------------------------------------------------------------------------

Time                          Error log of first 10 errors

2009-10-06 02:23:16 PDT       Cell drop errors on FPC7_T link

2009-10-06 02:23:16 PDT       Cell drop errors on FPC7_B link
```

show chassis fabric errors (FPC Errors Specific to an LCC Connected to a TX Matrix Plus Router)

```
user@host> show chassis fabric errors fpc 5
lcc 0
lcc0-re0:

---------------------------------------------------------------------------

Time                          Error log of first 10 errors

2009-10-06 13:56:59 PDT      PFE_T has link error on plane 1
```

show chassis fabric errors (SIB Errors Specific to an LCC Connected to a TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric errors sib 1
lcc 0
lcc0-re0:
---------------------------------------------------------------------------
Time                          Error log of first 10 errors
2013-02-11 04:46:42 PST      CRC errors on XC link SIB01_XF3#11,0
```

show chassis fabric errors fpc or sib (PTX Series Packet Transport Routers)

```
user@host> show chassis fabric errors fpc 1
Time                          Error log of first 10 errors
2012-01-06 16:27:03 PST      Link errs on PFE 2, SIB 0, Plane 0

user@host> show chassis fabric errors sib 1
Time                          Error log of first 10 errors
2015-011-16 15:34:33 PST      Link errs on PFE 0, FPC 0, Plane 2
2015-011-16 15:44:33 PST      CM set ASIC 1 to FAULT (Fault due to PIO errors)
```

show chassis fabric errors autoheal (PTX Series Packet Transport Routers)

```
user@host> show chassis fabric errors autoheal
Mar 30 01:43:00
Time                      Error log of first 100 errors
2016-03-29 23:46:23 PDT   Req: sib 0
2016-03-29 23:46:23 PDT   Action: SIB 0 (autohealing)
2016-03-29 23:54:52 PDT   Completed: SIB 0  (autoheal)
```

show chassis fabric errors autoheal (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric errors autoheal
Time                      Error log of first 100 errors
2013-03-25 00:16:10 PDT   Req: Plane 3 F13 8 Cbl 4 (tx) LCC0-SIB3 Cbl 4 (rx)
2013-03-25 00:16:12 PDT   Action: Plane 3 F13 8 Cbl 4 (autohealing)
2013-03-25 00:17:24 PDT   Req: Plane 3 F13 8 Cbl 4 (tx) LCC0-SIB3 Cbl 4 (rx)
2013-03-25 00:17:24 PDT   Denied: Plane 3 F13 8 Cbl 4 (time < configured)
2013-03-25 00:17:24 PDT   Set info: Plane 3 F13 8 Cbl 4 (skip autoheal)
2013-03-25 01:20:17 PDT   Clear info: Plane 3
```

Release Information

Command introduced in Junos OS Release 10.0.

show chassis fabric faults recovery-actions

IN THIS SECTION

- Syntax | 18076
- Description | 18076
- Required Privilege Level | 18076
- Output Fields | 18076
- Sample Output | 18077
- Release Information | 18077

Syntax

```
show chassis fabric faults recovery-actions
```

Description

Display the last 64 recovery actions related to a fabric null route condition.

Required Privilege Level

view

Output Fields

Table 255 on page 18076 lists the output fields for the `show chassis fabric faults recovery-actions` command. Output fields are listed in the approximate order in which they appear.

Table 255: show chassis fabric faults recovery-actions Output Fields

Field Name	Field Description
`Fault Name`	The name of fault detected. The fault name appears in a message such as: `CLOS LINK ERROR` detected on `F2S or F13 slot` *`slot number`* `xf chip` *`chip number`* and `Link Number` *`link number`*
`Recovery Start Time`	The time when recovery actions were initiated.

Table 255: show chassis fabric faults recovery-actions Output Fields *(Continued)*

Field Name	Field Description
`Recovery Action`	The recovery action that was used to recover from the mentioned fault. Recovery options depend on the type of faults and can include: • `SFC SIB Reboot`: The SFC SIB was rebooted. • `LCC SIB Reboot`: The LCC SIB was rebooted. • `FPC Reboot`: The FPC was rebooted. • `Destination Reprogramming`: Reenabling the data flow between Packet Forwarding Engines. • `Interchassis Link Retraining`: Retraining of optical links between an LCC SIB and an SFC SIB.

Sample Output

show chassis fabric faults recovery-actions

```
user@host> show chassis fabric faults recovery-actions
Fault Name          :CLOS LINK ERROR on F2SSlot3Chip0LinkNum2
Recovery Start Time :2014-03-25 19:52:50 PDT
Recovery Action     :F2S slot 3 Reboot
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

auto-recovery-disable

show chassis fabric feb

IN THIS SECTION

- Syntax | 18078
- Description | 18078
- Options | 18078
- Required Privilege Level | 18078
- Output Fields | 18078
- Sample Output | 18079
- Release Information | 18079

Syntax

```
show chassis fabric feb
```

Description

(M120 router only) Display the state of the electrical and optical switching fabric links between the Forwarding Engine Boards (FEBs) and the fabric planes, as interpreted by the FEB.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 256 on page 18079 lists the output fields for the `show chassis fabric feb` command.

Table 256: show chassis fabric feb Output Fields

Field Name	Field Description
Fabric management FEB state	State of the switching fabric link between each FEB and fabric plane: **desalination error**, **disabled**, **enabled**, **link error**, **link ok**, or **unused**.

Sample Output

show chassis fabric feb

```
user@host> show chassis fabric feb
Fabric management       FEB state
FEB 0                   Plane 0: Plane enabled
                        Plane 1: Plane enabled
                        Plane 2: Plane enabled
                        Plane 3: Plane enabled
FEB 4
                        Plane 0: Plane enabled
                        Plane 1: Plane enabled
                        Plane 2: Plane enabled
                        Plane 3: Plane enabled
```

Release Information

Command introduced in Junos OS Release 8.0.

show chassis fabric fpcs

IN THIS SECTION

- Syntax | 18080
- Syntax (MX Series Routers) | 18080
- Syntax (PTX Series Routers) | 18080

- Syntax (QFX Series Switches) | 18081
- Description | 18081
- Options | 18081
- Required Privilege Level | 18081
- Output Fields | 18082
- Sample Output | 18084
- Release Information | 18091

Syntax

```
show chassis fabric fpcs
<extended>
```

Syntax (MX Series Routers)

```
show chassis fabric fpcs
<all-members>
<extended>
<local>
<member member-id>
```

Syntax (PTX Series Routers)

```
show chassis fabric fpcs
<extended>
<slot fpc-slot>
```

Syntax (QFX Series Switches)

```
show chassis fabric fpcs
<extended>
<slot fpc-slot>
```

Description

Display the state of the electrical switch fabric links between the Flexible PIC Concentrators (FPCs) and the Switch Interface Boards (SIBs). (PTX10002-36QDD devices only) Display peer FPC / Packet Forwarding Engine (PFE) details because the PFEs are directly connected to each other.

Options

none	Display the switch fabric link state.
extended	(Optional) Display the fabric link state for all fabric planes.
all-members	(MX Series routers only) (Optional) Display the switching fabric link states for the FPCs in all members of the Virtual Chassis configuration.
local	(MX Series routers only) (Optional) Display the switching fabric link states for the FPCs in the local Virtual Chassis member.
member ***member-id***	(MX Series routers only) (Optional) Display the switching fabric link states for the FPCs in the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.
slot *fpc-slot* **(Optional)**	(PTX Series Packet Transport Routers and QFX Series switches only) (Optional) Display the fabric state of the specified FPC slot. If no value is provided, display the status of all FPCs. The FPC slots range from 0-7 in PTX10008 and 0-15 in PTX10016.

Required Privilege Level

view

Output Fields

Table 257 on page 18083 lists the output fields for the `show chassis fabric fpcs` command. Output fields are listed in the approximate order in which they appear.

Table 257: show chassis fabric fpcs Output Fields

Field Name	Field Description

`Fabric management FPC state`	Switching fabric link (link from SIB to FPC) state for each FPC:
	• `Unused`—FPC is not present.
	(On MX240 and MX480 routers with AS- MLC modular carrier card or MPC4E only) the fabric plane from the pair that share physical links (1 and 5, and 3 and 7) is inactive.
	• `Destination error on PFEs` *`list of PFE numbers`*—Destination errors to the listed Packet Forwarding Engines. Indicates that the link is not carrying traffic to the listed Packet Forwarding Engines.
	NOTE: In Junos OS Release 9.6 and later, the list of Packet Forwarding Engines with destination errors is displayed in the output.
	In Junos OS Releases before 9.6, the output only indicates that there are destination errors. However, the list of Packet Forwarding Engines with destination errors is not displayed.
	• `Links ok`—Link between the spare SIB and FPC is eligible to carry traffic.
	• `Link error`—Link between the SIB and FPC has CRC errors. However, the link is still eligible to carry traffic.
	• `Plane disabled`—Fabric plane has been disabled for the following reasons:
	  • Destination errors have exceeded the thresholds.
	  • Run-time link errors have exceeded the thresholds.
	  • Initialization time link errors detected, and link training was unsuccessful.
	  • `Plane Disabled, Links Error` (PTX Series Packet Transport Routers and QFX Series switches only)—The plane is disabled because of link errors detected at the FPC RX.
	• `Plane Disabled, Links Down` (PTX Series Packet Transport Routers and QFX Series switches only)—The plane is disabled because of link errors detected at the SIB RX.
	• `Plane enabled`—Link between the active SIB and FPC is eligible to carry traffic.

Table 257: show chassis fabric fpcs Output Fields *(Continued)*

Field Name	Field Description
	NOTE: On the Enhanced MX SCB with MPC, a maximum of 4 planes are operational and running. On all the other SCBs with MPC, all the planes are operational and running. • `Plane Enabled, Links OK` (PTX Series Packet Transport Routers and QFX Series switches only)—The FPC CCL RX link is eligible to carry traffic.

Sample Output

show chassis fabric fpcs (MX Series Routers)

```
user@host> show chassis fabric fpcs
FPC 0
  PFE #0
       Plane 0: Plane enabled
       Plane 1: Plane enabled
       Plane 2: Plane enabled
       Plane 3: Plane enabled
       Plane 4: Links ok
       Plane 5: Links ok
  PFE #1
       Plane 0: Plane enabled
       Plane 1: Plane enabled
       Plane 2: Plane enabled
       Plane 3: Plane enabled
       Plane 4: Links ok
       Plane 5: Links ok
  PFE #2
       Plane 0: Plane enabled
       Plane 1: Plane enabled
       Plane 2: Plane enabled
       Plane 3: Plane enabled
       Plane 4: Links ok
       Plane 5: Links ok
  PFE #3
       Plane 0: Plane enabled
```

```
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Links ok
        Plane 5: Links ok
FPC 1
  PFE #0
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
  PFE #1
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
FPC 2
  PFE #0
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Links ok
        Plane 5: Links ok
  PFE #1
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Links ok
        Plane 5: Links ok
  PFE #2
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Links ok
...
```

show chassis fabric fpcs (PTX Series Routers)

```
user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC #0
    PFE #0
    SIB0_Asic0_Fcore0 (plane 0)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB0_Asic0_Fcore0 (plane 1)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB0_Asic1_Fcore0 (plane 2)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB0_Asic1_Fcore0 (plane 3)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB0_Asic2_Fcore0 (plane 4)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB0_Asic2_Fcore0 (plane 5)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic0_Fcore0 (plane 6)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic0_Fcore0 (plane 7)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic1_Fcore0 (plane 8)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic1_Fcore0 (plane 9)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic2_Fcore0 (plane 10)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic2_Fcore0 (plane 11)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB2_Asic0_Fcore0 (plane 12)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB2_Asic0_Fcore0 (plane 13)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB2_Asic1_Fcore0 (plane 14)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB2_Asic1_Fcore0 (plane 15)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB2_Asic2_Fcore0 (plane 16)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB2_Asic2_Fcore0 (plane 17)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB3_Asic0_Fcore0 (plane 18)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB3_Asic0_Fcore0 (plane 19)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB3_Asic1_Fcore0 (plane 20)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB3_Asic1_Fcore0 (plane 21)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB3_Asic2_Fcore0 (plane 22)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB3_Asic2_Fcore0 (plane 23)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB4_Asic0_Fcore0 (plane 24)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB4_Asic0_Fcore0 (plane 25)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB4_Asic1_Fcore0 (plane 26)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB4_Asic1_Fcore0 (plane 27)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB4_Asic2_Fcore0 (plane 28)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB4_Asic2_Fcore0 (plane 29)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB5_Asic0_Fcore0 (plane 30)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB5_Asic0_Fcore0 (plane 31)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB5_Asic1_Fcore0 (plane 32)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB5_Asic1_Fcore0 (plane 33)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB5_Asic2_Fcore0 (plane 34)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB5_Asic2_Fcore0 (plane 35)  Plane Enabled, Links ok Fabric OAM Runtime success
```

```
    PFE #1
    SIB0_Asic0_Fcore0 (plane 0)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB0_Asic0_Fcore0 (plane 1)  Plane Enabled, Links ok Fabric OAM Runtime success
...
```

show chassis fabric fpcs (PTX10002-36QDD Routers)

```
user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC #0
    PFE #0
        FPC0_PFE1 (plane 0)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 1)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 2)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 3)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 4)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 5)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 6)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 7)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 8)  Plane Enabled, Links ok
…
    PFE #1
        FPC0_PFE0 (plane 0)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 1)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 2)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 3)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 4)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 5)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 6)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 7)  Plane Enabled, Links ok
…
    PFE #2
        FPC0_PFE0 (plane 80)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 0)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 1)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 2)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 3)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 4)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 5)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 6)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 7)  Plane Enabled, Links ok
```

```
        FPC0_PFE1 (plane 8)  Plane Enabled, Links ok
        FPC0_PFE1 (plane 9)  Plane Enabled, Links ok
…
    PFE #3
        FPC0_PFE0 (plane 0)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 1)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 2)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 3)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 4)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 5)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 6)  Plane Enabled, Links ok
        FPC0_PFE0 (plane 7)  Plane Enabled, Links ok
…
```

show chassis fabric fpcs (QFX10008 Switch)

```
user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC #1
    PFE #0
    SIB0_FASIC0 (plane  0)   Plane Enabled, Links OK
    SIB0_FASIC1 (plane  1)   Plane Enabled, Links OK
    SIB1_FASIC0 (plane  2)   Plane Enabled, Links OK
    SIB1_FASIC1 (plane  3)   Plane Enabled, Links OK
    SIB2_FASIC0 (plane  4)   Plane Enabled, Links OK
    SIB2_FASIC1 (plane  5)   Plane Enabled, Links OK
    SIB3_FASIC0 (plane  6)   Plane Enabled, Links OK
    SIB3_FASIC1 (plane  7)   Plane Enabled, Links OK
    SIB4_FASIC0 (plane  8)   Plane Enabled, Links OK
    SIB4_FASIC1 (plane  9)   Plane Enabled, Links OK
    SIB5_FASIC0 (plane 10)   Plane Enabled, Links OK
    SIB5_FASIC1 (plane 11)   Plane Enabled, Links OK
    PFE #1
    SIB0_FASIC0 (plane  0)   Plane Enabled, Links OK
    SIB0_FASIC1 (plane  1)   Plane Enabled, Links OK
    SIB1_FASIC0 (plane  2)   Plane Enabled, Links OK
    SIB1_FASIC1 (plane  3)   Plane Enabled, Links OK
    SIB2_FASIC0 (plane  4)   Plane Enabled, Links OK
    SIB2_FASIC1 (plane  5)   Plane Enabled, Links OK
    SIB3_FASIC0 (plane  6)   Plane Enabled, Links OK
    SIB3_FASIC1 (plane  7)   Plane Enabled, Links OK
```

```
    SIB4_FASIC0 (plane  8)   Plane Enabled, Links OK
    SIB4_FASIC1 (plane  9)   Plane Enabled, Links OK
    SIB5_FASIC0 (plane 10)   Plane Enabled, Links OK
    SIB5_FASIC1 (plane 11)   Plane Enabled, Links OK
    PFE #2
    SIB0_FASIC0 (plane  0)   Plane Enabled, Links OK
    SIB0_FASIC1 (plane  1)   Plane Enabled, Links OK
    SIB1_FASIC0 (plane  2)   Plane Enabled, Links OK
    SIB1_FASIC1 (plane  3)   Plane Enabled, Links OK
    SIB2_FASIC0 (plane  4)   Plane Enabled, Links OK
    SIB2_FASIC1 (plane  5)   Plane Enabled, Links OK
    SIB3_FASIC0 (plane  6)   Plane Enabled, Links OK
    SIB3_FASIC1 (plane  7)   Plane Enabled, Links OK
    SIB4_FASIC0 (plane  8)   Plane Enabled, Links OK
    SIB4_FASIC1 (plane  9)   Plane Enabled, Links OK
    SIB5_FASIC0 (plane 10)   Plane Enabled, Links OK
    SIB5_FASIC1 (plane 11)   Plane Enabled, Links OK
FPC #2
    PFE #0
    SIB0_FASIC0 (plane  0)   Plane Enabled, Links OK
    SIB0_FASIC1 (plane  1)   Plane Enabled, Links OK
    SIB1_FASIC0 (plane  2)   Plane Enabled, Links OK
    SIB1_FASIC1 (plane  3)   Plane Enabled, Links OK
    SIB2_FASIC0 (plane  4)   Plane Enabled, Links OK
    SIB2_FASIC1 (plane  5)   Plane Enabled, Links OK
    SIB3_FASIC0 (plane  6)   Plane Enabled, Links OK
    SIB3_FASIC1 (plane  7)   Plane Enabled, Links OK
    SIB4_FASIC0 (plane  8)   Plane Enabled, Links OK
    SIB4_FASIC1 (plane  9)   Plane Enabled, Links OK
    SIB5_FASIC0 (plane 10)   Plane Enabled, Links OK
    SIB5_FASIC1 (plane 11)   Plane Enabled, Links OK
    PFE #1
    SIB0_FASIC0 (plane  0)   Plane Enabled, Links OK
    SIB0_FASIC1 (plane  1)   Plane Enabled, Links OK
    SIB1_FASIC0 (plane  2)   Plane Enabled, Links OK
    SIB1_FASIC1 (plane  3)   Plane Enabled, Links OK
    SIB2_FASIC0 (plane  4)   Plane Enabled, Links OK
    SIB2_FASIC1 (plane  5)   Plane Enabled, Links OK
    SIB3_FASIC0 (plane  6)   Plane Enabled, Links OK
    SIB3_FASIC1 (plane  7)   Plane Enabled, Links OK
    SIB4_FASIC0 (plane  8)   Plane Enabled, Links OK
    SIB4_FASIC1 (plane  9)   Plane Enabled, Links OK
    SIB5_FASIC0 (plane 10)   Plane Enabled, Links OK
```

```
    SIB5_FASIC1 (plane 11)   Plane Enabled, Links OK
    PFE #2
    SIB0_FASIC0 (plane  0)   Plane Enabled, Links OK
    SIB0_FASIC1 (plane  1)   Plane Enabled, Links OK
    SIB1_FASIC0 (plane  2)   Plane Enabled, Links OK
    SIB1_FASIC1 (plane  3)   Plane Enabled, Links OK
    SIB2_FASIC0 (plane  4)   Plane Enabled, Links OK
    SIB2_FASIC1 (plane  5)   Plane Enabled, Links OK
    SIB3_FASIC0 (plane  6)   Plane Enabled, Links OK
    SIB3_FASIC1 (plane  7)   Plane Enabled, Links OK
    SIB4_FASIC0 (plane  8)   Plane Enabled, Links OK
    SIB4_FASIC1 (plane  9)   Plane Enabled, Links OK
    SIB5_FASIC0 (plane 10)   Plane Enabled, Links OK
    SIB5_FASIC1 (plane 11)   Plane Enabled, Links OK
FPC #4
    PFE #0
    SIB0_FASIC0 (plane  0)   Plane Enabled, Links OK
    SIB0_FASIC1 (plane  1)   Plane Enabled, Links OK
    SIB1_FASIC0 (plane  2)   Plane Enabled, Links OK
    SIB1_FASIC1 (plane  3)   Plane Enabled, Links OK
    SIB2_FASIC0 (plane  4)   Plane Enabled, Links OK
    SIB2_FASIC1 (plane  5)   Plane Enabled, Links OK
    SIB3_FASIC0 (plane  6)   Plane Enabled, Links OK
    SIB3_FASIC1 (plane  7)   Plane Enabled, Links OK
    SIB4_FASIC0 (plane  8)   Plane Enabled, Links OK
    SIB4_FASIC1 (plane  9)   Plane Enabled, Links OK
    SIB5_FASIC0 (plane 10)   Plane Enabled, Links OK
    SIB5_FASIC1 (plane 11)   Plane Enabled, Links OK
    PFE #1
    SIB0_FASIC0 (plane  0)   Plane Enabled, Links OK
    SIB0_FASIC1 (plane  1)   Plane Enabled, Links OK
    SIB1_FASIC0 (plane  2)   Plane Enabled, Links OK
    SIB1_FASIC1 (plane  3)   Plane Enabled, Links OK
    SIB2_FASIC0 (plane  4)   Plane Enabled, Links OK
    SIB2_FASIC1 (plane  5)   Plane Enabled, Links OK
    SIB3_FASIC0 (plane  6)   Plane Enabled, Links OK
    SIB3_FASIC1 (plane  7)   Plane Enabled, Links OK
    SIB4_FASIC0 (plane  8)   Plane Enabled, Links OK
    SIB4_FASIC1 (plane  9)   Plane Enabled, Links OK
    SIB5_FASIC0 (plane 10)   Plane Enabled, Links OK
    SIB5_FASIC1 (plane 11)   Plane Enabled, Links OK
```

Release Information

Command introduced before Junos OS Release 7.4.

`extended` option introduced in JunosOS Release 16.1.

RELATED DOCUMENTATION

request chassis fpc | 14475

Displaying Information About DPCs or FPCs in an MX Series Router

show chassis fabric map

IN THIS SECTION

- Syntax | 18091
- Syntax (MX Series Router) | 18092
- Description | 18092
- Options | 18092
- Required Privilege Level | 18093
- Output Fields | 18093
- Sample Output | 18094
- Release Information | 18101

Syntax

```
show chassis fabric map
plane <plane-number>
```

Syntax (MX Series Router)

```
show chassis fabric map
<all-members>
<local>
<member member-id>
<plane plane-number>
```

Description

(M120 and MX Series routers and EX8200 switches only) On the M120 router, display the state of the switching fabric map for connections from the Forwarding Engine Boards (FEBs) to the ports on the fabric planes, as interpreted by the fabric plane. On the MX Series router and the EX8200 switch, display the state of the switching fabric map for connections from each Packet Forwarding Engine on the Dense Port Concentrators (DPCs) to the ports on the fabric planes, as interpreted by the fabric plane. For information about the meaning of “fabric plane”, “DPCs”, and “SIBs” on the switches, see the hardware documentation for your switch.

Options

none Display the switching fabric map state for the M120 or MX Series router or EX8200 switch.

all-members (MX Series routers only) (Optional) Display the switching fabric map state for all the members of the Virtual Chassis configuration.

local (MX Series routers only) (Optional) Display the switching fabric map state for the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display the switching fabric map state for the specified member of the Virtual Chassis configuration. Replace the ***member-id*** with a value of 0 or 1.

plane *plane-number* (Optional) Display the state of the fabric link for the specified plane number.

- For the M120 router, replace ***plane-number*** with a value from 0 through 3.
- For the MX480 and MX240 routers, replace ***plane-number*** with a value from 0 through 7.
- For the MX960 router, replace ***plane-number*** with a value from 0 through 5.

- For the EX8208 switch, replace *plane-number* with a value from 0 through 11.
- For the EX8216 switch, replace *plane-number* with a value from 0 through 7.

Required Privilege Level

view

Output Fields

Table 258 on page 18093 lists the output fields for the `show chassis fabric map` command. Output fields are listed in the approximate order in which they appear.

Table 258: show chassis fabric map Output Fields

Field Name	Field Description
in-links	Fabric map for receive side links.
out-links	Fabric map for transmit side links.

state	State of the fabric link:
	• **RESET**—Link between SIB and FPC/DPC is powered down on purpose. This is done in all non-dual PFE based boards.
	• **UP**—Link between SIB and FPC/DPC is up and running.
	• **DOWN**—Link between SIB and FPC/DPC is powered down.
	• **FAULT**—SIB is in alarmed state where the SIB's plane is not operational for the following reasons:
	  • On-board F-chip is not operational.
	  • Fiber optic connector faults.
	  • FPC connector faults.
	  • SIB midplane connector faults.

Sample Output

show chassis fabric map (MX480 Router)

This sample applies to Junos OS Releases 18.2R2, 18.3R2, and 18.4R1 and later, in which the output of the CLI command `show chassis fabric map`, on MX series routers carrying MPCs with SF and XF ASICs, uses the label FPC instead of DPC (for example, 'FPC5PFE0').

```
user@host> show chassis fabric map plane 0

FPC5PFE0->CB0F0_02_0    Up         CB0F0_02_0->FPC5PFE0    Up
FPC5PFE1->CB0F0_02_1    Up         CB0F0_02_1->FPC5PFE1    Up
FPC5PFE2->CB0F0_02_2    Down       CB0F0_02_2->FPC5PFE2    Down
FPC5PFE3->CB0F0_02_3    Down       CB0F0_02_3->FPC5PFE3    Down
FPC5PFE0->CB0F0_03_0    Up         CB0F0_03_0->FPC5PFE0    Up
FPC5PFE1->CB0F0_03_1    Up         CB0F0_03_1->FPC5PFE1    Up
FPC5PFE2->CB0F0_03_2    Down       CB0F0_03_2->FPC5PFE2    Down
FPC5PFE3->CB0F0_03_3    Down       CB0F0_03_3->FPC5PFE3    Down
FPC3PFE0->CB0F0_04_0    Up         CB0F0_04_0->FPC3PFE0    Up
FPC3PFE1->CB0F0_04_1    Up         CB0F0_04_1->FPC3PFE1    Up
FPC3PFE2->CB0F0_04_2    Down       CB0F0_04_2->FPC3PFE2    Down
FPC3PFE3->CB0F0_04_3    Down       CB0F0_04_3->FPC3PFE3    Down
FPC3PFE0->CB0F0_05_0    Up         CB0F0_05_0->FPC3PFE0    Up
FPC3PFE1->CB0F0_05_1    Up         CB0F0_05_1->FPC3PFE1    Up
FPC3PFE2->CB0F0_05_2    Down       CB0F0_05_2->FPC3PFE2    Down
FPC3PFE3->CB0F0_05_3    Down       CB0F0_05_3->FPC3PFE3    Down
FPC1PFE0->CB0F0_06_0    Down       CB0F0_06_0->FPC1PFE0    Down
FPC1PFE1->CB0F0_06_1    Down       CB0F0_06_1->FPC1PFE1    Down
FPC1PFE2->CB0F0_06_2    Down       CB0F0_06_2->FPC1PFE2    Down
FPC1PFE3->CB0F0_06_3    Down       CB0F0_06_3->FPC1PFE3    Down
[...Output truncated...]
```

show chassis fabric map (M120 Router)

```
user@host> show chassis fabric map
FEB0->CB0F0_00 up CB0F0_08->FEB7 Down

FEB1->CB0F0_01 Down CB0F0_09->FEB6 Down

FEB6->CB0F0_02 Down CB0F0_10->FEB1 Down
```

```
FEB2->CB0F0_03 Down CB0F0_11->FEB0 up

FEB3->CB0F0_04 Down CB0F0_12->FEB3 Down

FEB4->CB0F0_05 up CB0F0_13->FEB2 Down

FEB7->CB0F0_06 Down CB0F0_14->FEB5 Down

FEB5->CB0F0_07 Down CB0F0_15->FEB4 up:
```

show chassis fabric map (MX Series Routers)

```
user@host> show chassis fabric map
DPC4PFE0->CB0F0_00_0    up         CB0F0_00_0->DPC4PFE0    up
DPC4PFE1->CB0F0_00_1    up         CB0F0_00_1->DPC4PFE1    up
DPC4PFE2->CB0F0_00_2    up         CB0F0_00_2->DPC4PFE2    up
DPC4PFE3->CB0F0_00_3    up         CB0F0_00_3->DPC4PFE3    up
DPC7PFE0->CB0F0_01_0    Down       CB0F0_01_0->DPC7PFE0    Down
DPC7PFE1->CB0F0_01_1    Down       CB0F0_01_1->DPC7PFE1    Down
DPC7PFE2->CB0F0_01_2    Down       CB0F0_01_2->DPC7PFE2    Down
DPC7PFE3->CB0F0_01_3    Down       CB0F0_01_3->DPC7PFE3    Down
DPC3PFE0->CB0F0_03_0    Down       CB0F0_03_0->DPC3PFE0    Down
DPC3PFE1->CB0F0_03_1    Down       CB0F0_03_1->DPC3PFE1    Down
DPC3PFE2->CB0F0_03_2    Down       CB0F0_03_2->DPC3PFE2    Down
DPC3PFE3->CB0F0_03_3    Down       CB0F0_03_3->DPC3PFE3    Down
DPC8PFE0->CB0F0_05_0    Down       CB0F0_05_0->DPC8PFE0    Down
DPC8PFE1->CB0F0_05_1    Down       CB0F0_05_1->DPC8PFE1    Down
DPC8PFE2->CB0F0_05_2    Down       CB0F0_05_2->DPC8PFE2    Down
DPC8PFE3->CB0F0_05_3    Down       CB0F0_05_3->DPC8PFE3    Down
DPC1PFE0->CB0F0_06_0    Down       CB0F0_06_0->DPC1PFE0    Down
DPC1PFE1->CB0F0_06_1    Down       CB0F0_06_1->DPC1PFE1    Down
DPC1PFE2->CB0F0_06_2    Down       CB0F0_06_2->DPC1PFE2    Down
DPC1PFE3->CB0F0_06_3    Down       CB0F0_06_3->DPC1PFE3    Down
DPC10PFE0->CB0F0_07_0   Down       CB0F0_07_0->DPC10PFE0   Down
DPC10PFE1->CB0F0_07_1   Down       CB0F0_07_1->DPC10PFE1   Down
DPC10PFE2->CB0F0_07_2   Down       CB0F0_07_2->DPC10PFE2   Down
DPC10PFE3->CB0F0_07_3   Down       CB0F0_07_3->DPC10PFE3   Down
DPC11PFE0->CB0F0_08_0   Down       CB0F0_08_0->DPC11PFE0   Down
DPC11PFE1->CB0F0_08_1   Down       CB0F0_08_1->DPC11PFE1   Down
DPC11PFE2->CB0F0_08_2   Down       CB0F0_08_2->DPC11PFE2   Down
```

```
DPC11PFE3->CB0F0_08_3   Down      CB0F0_08_3->DPC11PFE3   Down
DPC0PFE0->CB0F0_09_0    Down      CB0F0_09_0->DPC0PFE0    Down
DPC0PFE1->CB0F0_09_1    Down      CB0F0_09_1->DPC0PFE1    Down
DPC0PFE2->CB0F0_09_2    Down      CB0F0_09_2->DPC0PFE2    Down
DPC0PFE3->CB0F0_09_3    Down      CB0F0_09_3->DPC0PFE3    Down
DPC9PFE0->CB0F0_11_0    Down      CB0F0_11_0->DPC9PFE0    Down
DPC9PFE1->CB0F0_11_1    Down      CB0F0_11_1->DPC9PFE1    Down
DPC9PFE2->CB0F0_11_2    Down      CB0F0_11_2->DPC9PFE2    Down
DPC9PFE3->CB0F0_11_3    Down      CB0F0_11_3->DPC9PFE3    Down
DPC2PFE0->CB0F0_13_0    up        CB0F0_13_0->DPC2PFE0    up
DPC2PFE1->CB0F0_13_1    up        CB0F0_13_1->DPC2PFE1    up
DPC2PFE2->CB0F0_13_2    up        CB0F0_13_2->DPC2PFE2    up
DPC2PFE3->CB0F0_13_3    up        CB0F0_13_3->DPC2PFE3    up
DPC6PFE0->CB0F0_14_0    Down      CB0F0_14_0->DPC6PFE0    Down
DPC6PFE1->CB0F0_14_1    Down      CB0F0_14_1->DPC6PFE1    Down
DPC6PFE2->CB0F0_14_2    Down      CB0F0_14_2->DPC6PFE2    Down
DPC6PFE3->CB0F0_14_3    Down      CB0F0_14_3->DPC6PFE3    Down
DPC5PFE0->CB0F0_15_0    Down      CB0F0_15_0->DPC5PFE0    Down
DPC5PFE1->CB0F0_15_1    Down      CB0F0_15_1->DPC5PFE1    Down
DPC5PFE2->CB0F0_15_2    Down      CB0F0_15_2->DPC5PFE2    Down
DPC5PFE3->CB0F0_15_3    Down      CB0F0_15_3->DPC5PFE3    Down
DPC4PFE0->CB0F1_00_0    up        CB0F1_00_0->DPC4PFE0    up
DPC4PFE1->CB0F1_00_1    up        CB0F1_00_1->DPC4PFE1    up
DPC4PFE2->CB0F1_00_2    up        CB0F1_00_2->DPC4PFE2    up
DPC4PFE3->CB0F1_00_3    up        CB0F1_00_3->DPC4PFE3    up
DPC7PFE0->CB0F1_01_0    Down      CB0F1_01_0->DPC7PFE0    Down
DPC7PFE1->CB0F1_01_1    Down      CB0F1_01_1->DPC7PFE1    Down
DPC7PFE2->CB0F1_01_2    Down      CB0F1_01_2->DPC7PFE2    Down
DPC7PFE3->CB0F1_01_3    Down      CB0F1_01_3->DPC7PFE3    Down
DPC3PFE0->CB0F1_03_0    Down      CB0F1_03_0->DPC3PFE0    Down
DPC3PFE1->CB0F1_03_1    Down      CB0F1_03_1->DPC3PFE1    Down
DPC3PFE2->CB0F1_03_2    Down      CB0F1_03_2->DPC3PFE2    Down
DPC3PFE3->CB0F1_03_3    Down      CB0F1_03_3->DPC3PFE3    Down
DPC8PFE0->CB0F1_05_0    Down      CB0F1_05_0->DPC8PFE0    Down
DPC8PFE1->CB0F1_05_1    Down      CB0F1_05_1->DPC8PFE1    Down
DPC8PFE2->CB0F1_05_2    Down      CB0F1_05_2->DPC8PFE2    Down
DPC8PFE3->CB0F1_05_3    Down      CB0F1_05_3->DPC8PFE3    Down
DPC1PFE0->CB0F1_06_0    Down      CB0F1_06_0->DPC1PFE0    Down
DPC1PFE1->CB0F1_06_1    Down      CB0F1_06_1->DPC1PFE1    Down
DPC1PFE2->CB0F1_06_2    Down      CB0F1_06_2->DPC1PFE2    Down
DPC1PFE3->CB0F1_06_3    Down      CB0F1_06_3->DPC1PFE3    Down
DPC10PFE0->CB0F1_07_0   Down      CB0F1_07_0->DPC10PFE0   Down
DPC10PFE1->CB0F1_07_1   Down      CB0F1_07_1->DPC10PFE1   Down
```

```
DPC10PFE2->CB0F1_07_2   Down      CB0F1_07_2->DPC10PFE2   Down
DPC10PFE3->CB0F1_07_3   Down      CB0F1_07_3->DPC10PFE3   Down
DPC11PFE0->CB0F1_08_0   Down      CB0F1_08_0->DPC11PFE0   Down
DPC11PFE1->CB0F1_08_1   Down      CB0F1_08_1->DPC11PFE1   Down
DPC11PFE2->CB0F1_08_2   Down      CB0F1_08_2->DPC11PFE2   Down
DPC11PFE3->CB0F1_08_3   Down      CB0F1_08_3->DPC11PFE3   Down
DPC0PFE0->CB0F1_09_0    Down      CB0F1_09_0->DPC0PFE0    Down
DPC0PFE1->CB0F1_09_1    Down      CB0F1_09_1->DPC0PFE1    Down
DPC0PFE2->CB0F1_09_2    Down      CB0F1_09_2->DPC0PFE2    Down
DPC0PFE3->CB0F1_09_3    Down      CB0F1_09_3->DPC0PFE3    Down
DPC9PFE0->CB0F1_11_0    Down      CB0F1_11_0->DPC9PFE0    Down
DPC9PFE1->CB0F1_11_1    Down      CB0F1_11_1->DPC9PFE1    Down
DPC9PFE2->CB0F1_11_2    Down      CB0F1_11_2->DPC9PFE2    Down
DPC9PFE3->CB0F1_11_3    Down      CB0F1_11_3->DPC9PFE3    Down
DPC2PFE0->CB0F1_13_0    up        CB0F1_13_0->DPC2PFE0    up
DPC2PFE1->CB0F1_13_1    up        CB0F1_13_1->DPC2PFE1    up
DPC2PFE2->CB0F1_13_2    up        CB0F1_13_2->DPC2PFE2    up
DPC2PFE3->CB0F1_13_3    up        CB0F1_13_3->DPC2PFE3    up
DPC6PFE0->CB0F1_14_0    Down      CB0F1_14_0->DPC6PFE0    Down
DPC6PFE1->CB0F1_14_1    Down      CB0F1_14_1->DPC6PFE1    Down
DPC6PFE2->CB0F1_14_2    Down      CB0F1_14_2->DPC6PFE2    Down
DPC6PFE3->CB0F1_14_3    Down      CB0F1_14_3->DPC6PFE3    Down
DPC5PFE0->CB0F1_15_0    Down      CB0F1_15_0->DPC5PFE0    Down
DPC5PFE1->CB0F1_15_1    Down      CB0F1_15_1->DPC5PFE1    Down
DPC5PFE2->CB0F1_15_2    Down      CB0F1_15_2->DPC5PFE2    Down
DPC5PFE3->CB0F1_15_3    Down      CB0F1_15_3->DPC5PFE3    Down
DPC4PFE0->CB1F0_00_0    up        CB1F0_00_0->DPC4PFE0    up
DPC4PFE1->CB1F0_00_1    up        CB1F0_00_1->DPC4PFE1    up
DPC4PFE2->CB1F0_00_2    up        CB1F0_00_2->DPC4PFE2    up
DPC4PFE3->CB1F0_00_3    up        CB1F0_00_3->DPC4PFE3    up
DPC7PFE0->CB1F0_01_0    Down      CB1F0_01_0->DPC7PFE0    Down
DPC7PFE1->CB1F0_01_1    Down      CB1F0_01_1->DPC7PFE1    Down
DPC7PFE2->CB1F0_01_2    Down      CB1F0_01_2->DPC7PFE2    Down
DPC7PFE3->CB1F0_01_3    Down      CB1F0_01_3->DPC7PFE3    Down
DPC3PFE0->CB1F0_03_0    Down      CB1F0_03_0->DPC3PFE0    Down
DPC3PFE1->CB1F0_03_1    Down      CB1F0_03_1->DPC3PFE1    Down
DPC3PFE2->CB1F0_03_2    Down      CB1F0_03_2->DPC3PFE2    Down
DPC3PFE3->CB1F0_03_3    Down      CB1F0_03_3->DPC3PFE3    Down
DPC8PFE0->CB1F0_05_0    Down      CB1F0_05_0->DPC8PFE0    Down
DPC8PFE1->CB1F0_05_1    Down      CB1F0_05_1->DPC8PFE1    Down
DPC8PFE2->CB1F0_05_2    Down      CB1F0_05_2->DPC8PFE2    Down
DPC8PFE3->CB1F0_05_3    Down      CB1F0_05_3->DPC8PFE3    Down
DPC1PFE0->CB1F0_06_0    Down      CB1F0_06_0->DPC1PFE0    Down
```

```
DPC1PFE1->CB1F0_06_1     Down        CB1F0_06_1->DPC1PFE1     Down
DPC1PFE2->CB1F0_06_2     Down        CB1F0_06_2->DPC1PFE2     Down
DPC1PFE3->CB1F0_06_3     Down        CB1F0_06_3->DPC1PFE3     Down
DPC10PFE0->CB1F0_07_0    Down        CB1F0_07_0->DPC10PFE0    Down
DPC10PFE1->CB1F0_07_1    Down        CB1F0_07_1->DPC10PFE1    Down
DPC10PFE2->CB1F0_07_2    Down        CB1F0_07_2->DPC10PFE2    Down
DPC10PFE3->CB1F0_07_3    Down        CB1F0_07_3->DPC10PFE3    Down
DPC11PFE0->CB1F0_08_0    Down        CB1F0_08_0->DPC11PFE0    Down
DPC11PFE1->CB1F0_08_1    Down        CB1F0_08_1->DPC11PFE1    Down
DPC11PFE2->CB1F0_08_2    Down        CB1F0_08_2->DPC11PFE2    Down
DPC11PFE3->CB1F0_08_3    Down        CB1F0_08_3->DPC11PFE3    Down
DPC0PFE0->CB1F0_09_0     Down        CB1F0_09_0->DPC0PFE0     Down
DPC0PFE1->CB1F0_09_1     Down        CB1F0_09_1->DPC0PFE1     Down
DPC0PFE2->CB1F0_09_2     Down        CB1F0_09_2->DPC0PFE2     Down
DPC0PFE3->CB1F0_09_3     Down        CB1F0_09_3->DPC0PFE3     Down
DPC9PFE0->CB1F0_11_0     Down        CB1F0_11_0->DPC9PFE0     Down
DPC9PFE1->CB1F0_11_1     Down        CB1F0_11_1->DPC9PFE1     Down
DPC9PFE2->CB1F0_11_2     Down        CB1F0_11_2->DPC9PFE2     Down
DPC9PFE3->CB1F0_11_3     Down        CB1F0_11_3->DPC9PFE3     Down
DPC2PFE0->CB1F0_13_0     up          CB1F0_13_0->DPC2PFE0     up
DPC2PFE1->CB1F0_13_1     up          CB1F0_13_1->DPC2PFE1     up
DPC2PFE2->CB1F0_13_2     up          CB1F0_13_2->DPC2PFE2     up
DPC2PFE3->CB1F0_13_3     up          CB1F0_13_3->DPC2PFE3     up
DPC6PFE0->CB1F0_14_0     Down        CB1F0_14_0->DPC6PFE0     Down
DPC6PFE1->CB1F0_14_1     Down        CB1F0_14_1->DPC6PFE1     Down
DPC6PFE2->CB1F0_14_2     Down        CB1F0_14_2->DPC6PFE2     Down
DPC6PFE3->CB1F0_14_3     Down        CB1F0_14_3->DPC6PFE3     Down
DPC5PFE0->CB1F0_15_0     Down        CB1F0_15_0->DPC5PFE0     Down
DPC5PFE1->CB1F0_15_1     Down        CB1F0_15_1->DPC5PFE1     Down
DPC5PFE2->CB1F0_15_2     Down        CB1F0_15_2->DPC5PFE2     Down
DPC5PFE3->CB1F0_15_3     Down        CB1F0_15_3->DPC5PFE3     Down
DPC4PFE0->CB1F1_00_0     up          CB1F1_00_0->DPC4PFE0     up
DPC4PFE1->CB1F1_00_1     up          CB1F1_00_1->DPC4PFE1     up
DPC4PFE2->CB1F1_00_2     up          CB1F1_00_2->DPC4PFE2     up
DPC4PFE3->CB1F1_00_3     up          CB1F1_00_3->DPC4PFE3     up
DPC7PFE0->CB1F1_01_0     Down        CB1F1_01_0->DPC7PFE0     Down
DPC7PFE1->CB1F1_01_1     Down        CB1F1_01_1->DPC7PFE1     Down
DPC7PFE2->CB1F1_01_2     Down        CB1F1_01_2->DPC7PFE2     Down
DPC7PFE3->CB1F1_01_3     Down        CB1F1_01_3->DPC7PFE3     Down
DPC3PFE0->CB1F1_03_0     Down        CB1F1_03_0->DPC3PFE0     Down
DPC3PFE1->CB1F1_03_1     Down        CB1F1_03_1->DPC3PFE1     Down
DPC3PFE2->CB1F1_03_2     Down        CB1F1_03_2->DPC3PFE2     Down
DPC3PFE3->CB1F1_03_3     Down        CB1F1_03_3->DPC3PFE3     Down
```

```
DPC8PFE0->CB1F1_05_0    Down        CB1F1_05_0->DPC8PFE0    Down
DPC8PFE1->CB1F1_05_1    Down        CB1F1_05_1->DPC8PFE1    Down
DPC8PFE2->CB1F1_05_2    Down        CB1F1_05_2->DPC8PFE2    Down
DPC8PFE3->CB1F1_05_3    Down        CB1F1_05_3->DPC8PFE3    Down
DPC1PFE0->CB1F1_06_0    Down        CB1F1_06_0->DPC1PFE0    Down
DPC1PFE1->CB1F1_06_1    Down        CB1F1_06_1->DPC1PFE1    Down
DPC1PFE2->CB1F1_06_2    Down        CB1F1_06_2->DPC1PFE2    Down
DPC1PFE3->CB1F1_06_3    Down        CB1F1_06_3->DPC1PFE3    Down
DPC10PFE0->CB1F1_07_0   Down        CB1F1_07_0->DPC10PFE0   Down
DPC10PFE1->CB1F1_07_1   Down        CB1F1_07_1->DPC10PFE1   Down
DPC10PFE2->CB1F1_07_2   Down        CB1F1_07_2->DPC10PFE2   Down
DPC10PFE3->CB1F1_07_3   Down        CB1F1_07_3->DPC10PFE3   Down
DPC11PFE0->CB1F1_08_0   Down        CB1F1_08_0->DPC11PFE0   Down
DPC11PFE1->CB1F1_08_1   Down        CB1F1_08_1->DPC11PFE1   Down
DPC11PFE2->CB1F1_08_2   Down        CB1F1_08_2->DPC11PFE2   Down
DPC11PFE3->CB1F1_08_3   Down        CB1F1_08_3->DPC11PFE3   Down
DPC0PFE0->CB1F1_09_0    Down        CB1F1_09_0->DPC0PFE0    Down
DPC0PFE1->CB1F1_09_1    Down        CB1F1_09_1->DPC0PFE1    Down
DPC0PFE2->CB1F1_09_2    Down        CB1F1_09_2->DPC0PFE2    Down
DPC0PFE3->CB1F1_09_3    Down        CB1F1_09_3->DPC0PFE3    Down
DPC9PFE0->CB1F1_11_0    Down        CB1F1_11_0->DPC9PFE0    Down
DPC9PFE1->CB1F1_11_1    Down        CB1F1_11_1->DPC9PFE1    Down
DPC9PFE2->CB1F1_11_2    Down        CB1F1_11_2->DPC9PFE2    Down
DPC9PFE3->CB1F1_11_3    Down        CB1F1_11_3->DPC9PFE3    Down
DPC2PFE0->CB1F1_13_0    up          CB1F1_13_0->DPC2PFE0    up
DPC2PFE1->CB1F1_13_1    up          CB1F1_13_1->DPC2PFE1    up
DPC2PFE2->CB1F1_13_2    up          CB1F1_13_2->DPC2PFE2    up
DPC2PFE3->CB1F1_13_3    up          CB1F1_13_3->DPC2PFE3    up
DPC6PFE0->CB1F1_14_0    Down        CB1F1_14_0->DPC6PFE0    Down
DPC6PFE1->CB1F1_14_1    Down        CB1F1_14_1->DPC6PFE1    Down
DPC6PFE2->CB1F1_14_2    Down        CB1F1_14_2->DPC6PFE2    Down
DPC6PFE3->CB1F1_14_3    Down        CB1F1_14_3->DPC6PFE3    Down
DPC5PFE0->CB1F1_15_0    Down        CB1F1_15_0->DPC5PFE0    Down
DPC5PFE1->CB1F1_15_1    Down        CB1F1_15_1->DPC5PFE1    Down
DPC5PFE2->CB1F1_15_2    Down        CB1F1_15_2->DPC5PFE2    Down
DPC5PFE3->CB1F1_15_3    Down        CB1F1_15_3->DPC5PFE3    Down
plane 4 is not up
plane 5 is not up
```

show chassis fabric map plane 1 (EX8200 Switch)

```
user@host> show chassis fabric map plane 1
user@host> show chassis fabric map plane 1
DPC6PFE0->CB0F0_00_0    Down       CB0F0_00_0->DPC6PFE0    Down
DPC6PFE1->CB0F0_00_1    Down       CB0F0_00_1->DPC6PFE1    Down
DPC6PFE2->CB0F0_00_2    Down       CB0F0_00_2->DPC6PFE2    Down
DPC6PFE3->CB0F0_00_3    Down       CB0F0_00_3->DPC6PFE3    Down
DPC0PFE0->CB0F0_01_0    Down       CB0F0_01_0->DPC0PFE0    Down
DPC0PFE1->CB0F0_01_1    Down       CB0F0_01_1->DPC0PFE1    Down
DPC0PFE2->CB0F0_01_2    Down       CB0F0_01_2->DPC0PFE2    Down
DPC0PFE3->CB0F0_01_3    Down       CB0F0_01_3->DPC0PFE3    Down
DPC5PFE0->CB0F0_02_0    Down       CB0F0_02_0->DPC5PFE0    Down
DPC5PFE1->CB0F0_02_1    Down       CB0F0_02_1->DPC5PFE1    Down
DPC5PFE2->CB0F0_02_2    Down       CB0F0_02_2->DPC5PFE2    Down
DPC5PFE3->CB0F0_02_3    Down       CB0F0_02_3->DPC5PFE3    Down
DPC3PFE0->CB0F0_03_0    Down       CB0F0_03_0->DPC3PFE0    Down
DPC3PFE1->CB0F0_03_1    Down       CB0F0_03_1->DPC3PFE1    Down
DPC3PFE2->CB0F0_03_2    Down       CB0F0_03_2->DPC3PFE2    Down
DPC3PFE3->CB0F0_03_3    Down       CB0F0_03_3->DPC3PFE3    Down
DPC4PFE0->CB0F0_04_0    Down       CB0F0_04_0->DPC4PFE0    Down
DPC4PFE1->CB0F0_04_1    Down       CB0F0_04_1->DPC4PFE1    Down
DPC4PFE2->CB0F0_04_2    Down       CB0F0_04_2->DPC4PFE2    Down
DPC4PFE3->CB0F0_04_3    Down       CB0F0_04_3->DPC4PFE3    Down
DPC2PFE0->CB0F0_05_0    Down       CB0F0_05_0->DPC2PFE0    Down
DPC2PFE1->CB0F0_05_1    Down       CB0F0_05_1->DPC2PFE1    Down
DPC2PFE2->CB0F0_05_2    Down       CB0F0_05_2->DPC2PFE2    Down
DPC2PFE3->CB0F0_05_3    Down       CB0F0_05_3->DPC2PFE3    Down
DPC7PFE0->CB0F0_06_0    Down       CB0F0_06_0->DPC7PFE0    Down
DPC7PFE1->CB0F0_06_1    Down       CB0F0_06_1->DPC7PFE1    Down
DPC7PFE2->CB0F0_06_2    Down       CB0F0_06_2->DPC7PFE2    Down
DPC7PFE3->CB0F0_06_3    Down       CB0F0_06_3->DPC7PFE3    Down
DPC1PFE0->CB0F0_07_0    Down       CB0F0_07_0->DPC1PFE0    Down
DPC1PFE1->CB0F0_07_1    Down       CB0F0_07_1->DPC1PFE1    Down
DPC1PFE2->CB0F0_07_2    Down       CB0F0_07_2->DPC1PFE2    Down
DPC1PFE3->CB0F0_07_3    Down       CB0F0_07_3->DPC1PFE3    Down
DPC0PFE0->CB0F0_08_0    Down       CB0F0_08_0->DPC0PFE0    Down
DPC0PFE1->CB0F0_08_1    Down       CB0F0_08_1->DPC0PFE1    Down
DPC0PFE2->CB0F0_08_2    Down       CB0F0_08_2->DPC0PFE2    Down
DPC0PFE3->CB0F0_08_3    Down       CB0F0_08_3->DPC0PFE3    Down
DPC7PFE0->CB0F0_09_0    Down       CB0F0_09_0->DPC7PFE0    Down
DPC7PFE1->CB0F0_09_1    Down       CB0F0_09_1->DPC7PFE1    Down
```

```
DPC7PFE2->CB0F0_09_2    Down        CB0F0_09_2->DPC7PFE2    Down
DPC7PFE3->CB0F0_09_3    Down        CB0F0_09_3->DPC7PFE3    Down
DPC1PFE0->CB0F0_10_0    Down        CB0F0_10_0->DPC1PFE0    Down
DPC1PFE1->CB0F0_10_1    Down        CB0F0_10_1->DPC1PFE1    Down
DPC1PFE2->CB0F0_10_2    Down        CB0F0_10_2->DPC1PFE2    Down
DPC1PFE3->CB0F0_10_3    Down        CB0F0_10_3->DPC1PFE3    Down
DPC4PFE0->CB0F0_11_0    Down        CB0F0_11_0->DPC4PFE0    Down
DPC4PFE1->CB0F0_11_1    Down        CB0F0_11_1->DPC4PFE1    Down
DPC4PFE2->CB0F0_11_2    Down        CB0F0_11_2->DPC4PFE2    Down
DPC4PFE3->CB0F0_11_3    Down        CB0F0_11_3->DPC4PFE3    Down
DPC2PFE0->CB0F0_12_0    Down        CB0F0_12_0->DPC2PFE0    Down
DPC2PFE1->CB0F0_12_1    Down        CB0F0_12_1->DPC2PFE1    Down
DPC2PFE2->CB0F0_12_2    Down        CB0F0_12_2->DPC2PFE2    Down
DPC2PFE3->CB0F0_12_3    Down        CB0F0_12_3->DPC2PFE3    Down
DPC5PFE0->CB0F0_13_0    Down        CB0F0_13_0->DPC5PFE0    Down
DPC5PFE1->CB0F0_13_1    Down        CB0F0_13_1->DPC5PFE1    Down
DPC5PFE2->CB0F0_13_2    Down        CB0F0_13_2->DPC5PFE2    Down
DPC5PFE3->CB0F0_13_3    Down        CB0F0_13_3->DPC5PFE3    Down
DPC3PFE0->CB0F0_14_0    Down        CB0F0_14_0->DPC3PFE0    Down
DPC3PFE1->CB0F0_14_1    Down        CB0F0_14_1->DPC3PFE1    Down
DPC3PFE2->CB0F0_14_2    Down        CB0F0_14_2->DPC3PFE2    Down
DPC3PFE3->CB0F0_14_3    Down        CB0F0_14_3->DPC3PFE3    Down
DPC6PFE0->CB0F0_15_0    Down        CB0F0_15_0->DPC6PFE0    Down
DPC6PFE1->CB0F0_15_1    Down        CB0F0_15_1->DPC6PFE1    Down
DPC6PFE2->CB0F0_15_2    Down        CB0F0_15_2->DPC6PFE2    Down
DPC6PFE3->CB0F0_15_3    Down        CB0F0_15_3->DPC6PFE3    Down
```

Release Information

Command introduced in Junos OS Release 8.0.

show chassis fabric mesh

IN THIS SECTION

- Syntax | **18102**
- Description | **18102**

- Options | 18102
- Required Privilege Level | 18102
- Output Fields | 18102
- Sample Output | 18103
- Release Information | 18104

Syntax

```
show chassis fabric mesh < feb | topology >
```

Description

(M120 router only) Display the state of the electrical and optical switching fabric links between the Forwarding Engine Boards (FEBs) and the fabric planes, as interpreted by the FEB.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 259: show chassis fabric mesh feb Output Fields

Field Name	Field Description
Fabric management FEB state	State of the switching fabric mesh between each FEB and fabric **plane**: **desalination error**, **disabled**, **enabled**, **link error**, **link ok**, or **unused**.

Table 260: show chassis fabric mesh topology Output Fields

Field Name	Field Description
Output Fields	See "show chassis fabric topology" on page 18220 output fields.

Sample Output

show chassis fabric mesh feb

```
user@host> show chassis fabric mesh feb

        Fabric mesh management Feb state:
FEB #0
  PFE #0
      Asic0_Fcore0 (plane 0)  plane enabled, links ok
      Asic0_Fcore0 (plane 1)  plane enabled, links ok
      Asic0_Fcore0 (plane 2)  plane enabled, links ok
      Asic0_Fcore0 (plane 3)  plane enabled, links ok
      Asic0_Fcore0 (plane 4)  plane enabled, links ok
      Asic0_Fcore0 (plane 5)  plane enabled, links ok
```

show chassis fabric mesh topology

```
user@host> show chassis fabric mesh topology
Mesh Topology: asci connectivity

 In-link  : FEB# FE# (TX PM_ID#, TX Serdes_tx  #) ->
            FEB# FE# (RX PM_ID#, RX Serdes_rx  #)
 Out-link : FEB# FE# (TX PM_ID#, TX Serdes_tx  #) ->
            FEB# FE# (RX PM_ID#, RX Serdes_rx  #)

FEB 0  ASIC 0 FCORE 0 :

           In-links               State            Out-links              State
--------------------------------------------------------------------------------
 FEB0FE1(33,44)-->FEB0FE0(28,0)   up      FEB0FE0(28,0)-->FEB0FE1(33,44)   up
 FEB0FE1(33,45)-->FEB0FE0(28,1)   up      FEB0FE0(28,1)-->FEB0FE1(33,45)   up
```

```
FEB0FE1(33,46)-->FEB0FE0(28,2)   up      FEB0FE0(28,2)-->FEB0FE1(33,46)   up
FEB0FE1(33,47)-->FEB0FE0(28,3)   up      FEB0FE0(28,3)-->FEB0FE1(33,47)   up
FEB0FE1(33,40)-->FEB0FE0(28,4)   up      FEB0FE0(28,4)-->FEB0FE1(33,40)   up
FEB0FE1(33,41)-->FEB0FE0(28,5)   up      FEB0FE0(28,5)-->FEB0FE1(33,41)   up
```

Release Information

Command introduced in Junos OS Release 21.3

RELATED DOCUMENTATION

- request chassis feb | 14472
- show chassis feb | 18278
- show chassis fabric feb | 18078
- show chassis fpc-feb-connectivity | 18397

show chassis fabric plane (Security)

IN THIS SECTION

- Syntax | 18104
- Description | 18105
- Required Privilege Level | 18105
- Output Fields | 18105
- Sample Output | 18108
- Release Information | 18113

Syntax

```
show chassis fabric plane
```

Description

Show state of fabric management plane.

Required Privilege Level

view

Output Fields

Table 261 on page 18105 lists the output fields for the **show chassis fabric plane** command. Output fields are listed in the approximate order in which they appear.

Table 261: show chassis fabric plane Output Fields

Field Name	Field Description	Level of output
Plane	Number of the plane.	none

Plane state	State of each plane:
	• **ACTIVE**—SIB is operational and running.
	• **FAULTY**— SIB is in alarmed state where the SIB's plane is not operational for the following reasons:
	  • On-board fabric ASIC is not operational.
	  • Fiber-optic connector faults.
	  • FPC connector faults.
	  • SIB midplane connector faults.
		none

FPC	Slot number of each Flexible PIC Concentrator (FPC).	none

Table 261: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output
PFE	Slot number of each Packet Forwarding Engine and the state of the links to the FPC: • **Links ok**: Link between SIB and FPC is active. • **Link error**: Link between SIB and FPC is not operational. • **Unused**: No FPC is present.	none

Table 261: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output

State	State of the fabric plane:
	• **Online**: Fabric plane is operational and running and links on the SIB are operational.
	• **Offline**: Fabric plane state is **Offline** because the plane does not have four or more F2S and one F13 online.
	• **Empty**: Fabric plane state is **Empty** if all SIBs in the plane are absent.
	• **Spare**: Fabric plane is redundant and can be operational if the operational fabric plane encounters an error.
	• **Check**: Fabric plane is in alarmed state due to the following reason and the cause of the error must be resolved:
	  • One or more SIBs (belonging to the fabric plane) in the **Online** or **Spare** states has transitioned to the **Check** state.
	  Check state of the SIB can be caused by link errors or destination errors.
	• **Fault**: Fabric plane is in alarmed state if one or more SIBs belonging to the plane are in the **Fault** state. A SIB can be in the **Fault** state because of the following reasons:
	  • On-board fabric ASIC is not operational.
	  • Fiber-optic connector faults.
	  • FPC connector faults.
	  • SIB midplane connector faults.
	  • Link errors have exceeded the threshold.
		none

Sample Output

show chassis fabric plane (SRX5600 and SRX5800 Devices with SRX5000 Line SCB II [SRX5K-SCBE] and SRX5K-RE-1800X4)

```
user@host> show chassis fabric plane
node0:
--------------------------------------------------------------------------
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
      FPC 3
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
      FPC 7
          PFE 0 :Links ok
      FPC 8
          PFE 0 :Links ok
      FPC 9
          PFE 0 :Links ok
      FPC 10
          PFE 0 :Links ok
Plane 1
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
      FPC 3
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
      FPC 7
          PFE 0 :Links ok
      FPC 8
          PFE 0 :Links ok
      FPC 9
```

```
            PFE 0 :Links ok
        FPC 10
            PFE 0 :Links ok
Plane 2
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
      FPC 3
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
      FPC 7
          PFE 0 :Links ok
      FPC 8
          PFE 0 :Links ok
      FPC 9
          PFE 0 :Links ok
      FPC 10
          PFE 0 :Links ok
Plane 3
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
      FPC 3
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
      FPC 7
          PFE 0 :Links ok
      FPC 8
          PFE 0 :Links ok
      FPC 9
          PFE 0 :Links ok
      FPC 10
          PFE 0 :Links ok
Plane 4
  Plane state: SPARE
      FPC 0
          PFE 0 :Links ok
```

```
      FPC 2
          PFE 0 :Links ok
      FPC 3
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
      FPC 7
          PFE 0 :Links ok
      FPC 8
          PFE 0 :Links ok
      FPC 9
          PFE 0 :Links ok
      FPC 10
          PFE 0 :Links ok
Plane 5
  Plane state: SPARE
      FPC 0
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
      FPC 3
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
      FPC 7
          PFE 0 :Links ok
      FPC 8
          PFE 0 :Links ok
      FPC 9
          PFE 0 :Links ok
      FPC 10
          PFE 0 :Links ok

node1:
--------------------------------------------------------------------------
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
      FPC 1
          PFE 0 :Links ok
      FPC 2
```

```
            PFE 0 :Links ok
        FPC 3
            PFE 0 :Links ok
        FPC 4
            PFE 0 :Links ok
        FPC 7
            PFE 0 :Links ok
        FPC 8
            PFE 0 :Links ok
        FPC 10
            PFE 0 :Links ok
Plane 1
  Plane state: ACTIVE
        FPC 0
            PFE 0 :Links ok
        FPC 1
            PFE 0 :Links ok
        FPC 2
            PFE 0 :Links ok
        FPC 3
            PFE 0 :Links ok
        FPC 4
            PFE 0 :Links ok
        FPC 7
            PFE 0 :Links ok
        FPC 8
            PFE 0 :Links ok
        FPC 10
            PFE 0 :Links ok
Plane 2
  Plane state: ACTIVE
        FPC 0
            PFE 0 :Links ok
        FPC 1
            PFE 0 :Links ok
        FPC 2
            PFE 0 :Links ok
        FPC 3
            PFE 0 :Links ok
        FPC 4
            PFE 0 :Links ok
        FPC 7
            PFE 0 :Links ok
```

```
      FPC 8
          PFE 0 :Links ok
      FPC 10
          PFE 0 :Links ok
Plane 3
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
      FPC 1
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
      FPC 3
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
      FPC 7
          PFE 0 :Links ok
      FPC 8
          PFE 0 :Links ok
      FPC 10
          PFE 0 :Links ok
Plane 4
  Plane state: SPARE
      FPC 0
          PFE 0 :Links ok
      FPC 1
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
      FPC 3
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
      FPC 7
          PFE 0 :Links ok
      FPC 8
          PFE 0 :Links ok
      FPC 10
          PFE 0 :Links ok
Plane 5
  Plane state: SPARE
      FPC 0
```

```
            PFE 0 :Links ok
        FPC 1
            PFE 0 :Links ok
        FPC 2
            PFE 0 :Links ok
        FPC 3
            PFE 0 :Links ok
        FPC 4
            PFE 0 :Links ok
        FPC 7
            PFE 0 :Links ok
        FPC 8
            PFE 0 :Links ok
        FPC 10
            PFE 0 :Links ok
```

Release Information

Command introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

show chassis fabric plane-location (Security) | 18154

show chassis fabric plane

IN THIS SECTION

- Syntax | 18114
- Syntax (TX Matrix Plus Router) | 18114
- Syntax (MX Series Routers) | 18114
- Description | 18114
- Options | 18115
- Required Privilege Level | 18116
- Output Fields | 18116

Sample Output | 18127
Release Information | 18154

Syntax

```
show chassis fabric plane
```

Syntax (TX Matrix Plus Router)

```
show chassis fabric plane
<detail | extensive | terse>
<lcc number | sfc number>
```

Syntax (MX Series Routers)

```
show chassis fabric plane
<extended>
<detail | extensive | terse>
<all-members>
<local>
<member member-id>
```

Description

(TX Matrix Plus router, T4000, T1600, M120, and MX Series routers and EX8200 switches only) On the M120 router, display the state of all fabric plane connections to the Forwarding Engine Boards (FEBs). On MX Series routers, display the state of all fabric plane connections to the Dense Port Concentrators (DPCs) and Packet Forwarding Engines (PFEs) on the Flexible PIC Concentrators (FPCs). On the TX Matrix Plus router, and on T1600 or T4000 routers in a routing matrix, display the state of the fabric management plane and the logical planes on the switch-fabric chassis (SFC) and line-card chassis (LCC). On EX8200 switches, display the state of all fabric planes. This command can be used on the primary Routing Engine only.

Options

none (MX2010, MX2020, and MX2008 Routers only) (Optional) Display the state of the fabric management plane.

extended (MX2020, MX2010, and MX2008 Routers only) (Optional) Display the state of the fabric management planes (all 24 fabric planes).

detail (TX Matrix Plus routers, T1600 or T4000 routers in a routing matrix, and MX Series routers only) (Optional) Display detailed output for the fabric management plane. Show Switch Interface Board (SIB) states for the TXP-F13 SIB and the TXP-F2S SIB.

extensive (TX Matrix Plus routers, T1600 or T4000 routers in a routing matrix, and MX Series routers only) (Optional) Display extensive output for the fabric management plane.

terse (TX Matrix Plus routers and MX Series routers only) (Optional) Display terse output for the fabric management plane.

all-members (MX Series routers only) (Optional) Display the state of all fabric plane connections on all members of the Virtual Chassis configuration.

lcc *number* (TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Display the state of all fabric plane connections on the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display the state of all fabric plane connections on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

sfc *number* (TX Matrix Plus router only) (Optional) Show information about the TX Matrix Plus router (SFC). Replace *number* with 0.

Required Privilege Level

view

Output Fields

Table 262 on page 18116 lists the output fields for the `show chassis fabric plane` command. Output fields are listed in the approximate order in which they appear.

Table 262: show chassis fabric plane Output Fields

Field Name	Field Description	Level of output
Plane	(TX Matrix Plus, MX Series routers, M120 routers, and EX8200 switches only) Number of the plane.	none

Table 262: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output
Plane state	(MX Series and M120 routers and EX8200 switches only) State of each plane: • ACTIVE—SIB is operational and running. **NOTE**: On the Enhanced MX SCB with MPCs, a maximum of 4 planes are operational and running. On all the other SCBs with MPCs, all the planes are operational and running. • FAULTY— SIB is in alarmed state where the SIB's plane is not operational for the following reasons: • On-board fabric ASIC is not operational. • Fiber optic connector faults. • FPC connector faults. • SIB midplane connector faults. (MX2010, MX2020, MX10003, and MX2008 Routers only) State of each plane: • ACTIVE—SFB is operational and running. • OFFLINE— SFB is in offline.	none
FEB	(M120 routers only) FEB number and state of links to each FEB: • Link error—Link between SIB and FPC is not operational. • Links ok—Link between SIB and FPC is active. • Unused—No FPC is present.	none
FPC	(MX Series routers only) Slot number of each Dense Port Concentrator (DPC) or Flexible PIC Concentrator (FPC). An FPC occupies two DPC slots on an MX Series router. The interface corresponds to the lowest numbered DPC slot for which the FPC is installed.	none

Table 262: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output
PFE	(MX Series and M120 routers only) Slot number of each Packet Forwarding Engine and the state of the links to the DCP: **Links ok**, **Link error**, or **Unused**. Each DPC includes four Packet Forwarding Engines. • **Links ok**: Link between SIB and FPC is active. • **Link error**: Link between SIB and FPC is not operational. • **Unused**: No FPC is present. (On MX240 and MX480 routers with AS MLC modular carrier card and MPC4E only) Indicates that the link between the fabric plane and the hardware link on the modular carrier card or MPC4E is not operational. (MX2010, MX2020, and MX2008 routers only) Slot number of each Packet Forwarding Engine and the state of the links to the DPC: **Links ok**, **Link error**, or **Unused**. Each DPC includes four Packet Forwarding Engines. • **Links ok**: Link between SFB and FPC is active. • **Link error**: Link between SFB and FPC is not operational. • **Unused**: No FPC is present.	none

Table 262: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output

State	(TX Matrix Plus, and T1600 or T4000 routers in a routing matrix only)—State of the fabric plane:
	• **Online**: Fabric plane is operational and running and links on the SIB are operational.
	• **Offline**: Fabric plane state is **Offline** because the plane does not have four or more F2S and one F13 online.
	• Empty: Fabric plane state is **Empty** if all SIBs in the plane are absent.
	• **Spare**: Fabric plane is redundant and can be operational if the operational fabric plane encounters an error.
	• **Check**: Fabric plane is in alarmed state due to the following reason and the cause of the error must be resolved:
	  • One or more SIBs (belonging to the fabric plane) in the **Online** or **Spare** states has transitioned to the **Check** state.
	  Check state of the SIB can be caused by link errors or destination errors.
	• **Fault**: Fabric plane is in alarmed state if one or more SIBs belonging to the plane are in the **Fault** state. A SIB can be in the **Fault** state because of the following reasons:
	  • On-board fabric ASIC is not operational.
	  • Fiber optic connector faults.
	  • FPC connector faults.
	  • SIB midplane connector faults.
	  • Link errors have exceeded the threshold.
		none

Link Errors	(TX Matrix Plus routers with 3D SIBs only) indicate the number of links which are marked faulty because the errors on them have crossed threshold.	none

Table 262: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output
Cable Errors	(TX Matrix Plus routers with 3D SIBs only) Indicate the number of mandatory cables that are not connected, or in up state for that plane	none
Destination Errors	(TX Matrix Plus routers with 3D SIBs only) Indicates the number of destinations that are not reachable on this plane.	none
Uptime	(TX Matrix Plus, and T1600 or T4000 routers in a routing matrix only)—Time the fabric plane has been up and running.	none
Fabric Management Plane State Output Fields for the `show chassis fabric plane extensive` Command on a TX Matrix Plus Router		

Table 262: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output

PLANE *number*	State of the fabric plane:
	• **Online**: Fabric plane is operational and running and links on the SIB are operational.
	• **Offline**: Fabric plane state is **Offline** because the plane does not have 4 or more F2S and 1 F13 online.
	• Empty: Fabric plane state is **Empty** if all SIBs in the plane are absent.
	• **Spare**: Fabric plane is redundant and can be operational if the operational fabric plane encounters an error.
	• **Check**: Fabric plane is in alarmed state due to the following reasons and the cause of the error must be resolved:
	  • One or more SIBs (belonging to the fabric plane) in the **Online** or **Spare** states has transitioned to the **Check** state.
	  Check state of the SIB can be caused because of link errors or destination errors.
	• **Fault**: Fabric plane is in alarmed state if one or more SIBs belonging to the plane are in the **Fault** state. A SIB can be in the **Fault** state because of the following reasons:
	  • On-board fabric ASIC is not operational.
	  • Fiber optic connector faults.
	  • FPC connector faults.
	  • SIB midplane connector faults.
	  • Link errors have exceeded the threshold.
		extensive

Table 262: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output

SIB F13/F2S *slot-number*	State of the TXP-F13 SIB or TXP-F2S SIB:
	• **Activating**—Transitional state when the SIB is transitioning to the **Online** or **Spare** state.
	• **Deactivating**—Transitional state when the SIB is going offline.
	• **Online**—SIB is operational and running.
	• **Offline**—SIB is powered down.
	• **Spare**—SIB is redundant and will move to active state if one of the working SIBs fails to pass traffic.
	• **Empty**—No SIB is present.
	• **Fault**— SIB is in alarmed state because of the following reasons and the cause of the error must be resolved:
	  • On-board fabric ASIC is not operational.
	  • Fiber optic connector faults.
	  • FPC connector faults.
	  • SIB midplane connector faults.
	  • Link errors have exceeded the threshold
	• **Check**—SIB is in alarmed state where the SIB is partially operational because of link or destination errors. Only a SIB that is **Online** or **Spare** can transition to the **Check** state.
	  NOTE: If a SIB is not inserted properly, the SIB cannot transition to the **Online** or **Spare** state, and therefore cannot transition to the **Check** state.
		extensive

Table 262: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output
SIB F13 ***slot-number Odd/Even***	State of the TXP-F13 SIB even and odd port connection optical links from the TX Matrix Plus router (SFC) to the router (LCC) in the routing matrix . The left four ports on the SFC are labeled **Even** and provide connections to one even-numbered LCC—LCC0 or LCC2. The right four ports on the SFC are labeled **Odd** and provide connections to one odd-numbered LCC—LCC1 or LCC3.	**extensive**
LCC ***number*****, SIB** ***slot-number***	State of the SIB on the LCC that is connected to the **Even** or **Odd** port on the TXP-F13 SIB faceplate: • **Links ok**—Links between the TXP-F13 SIB on the SFC and the LCC are active. • **Links error**—One or more links between the TXP-F13 SIB on the SFC and the LCC, have experienced an error, but the affected links remain operational. • **Unused**—No SIB is present.	**extensive**
SG ***number*** **Port** ***number***	State of the SG chip ports on the LCC: • **Links ok**—Link is active. • **Link error**—Link is operational with errors. • `Link error crc saturated`—CRC has exceeded the rate threshold and reached saturation without optical issues—that is, a cable has not been cut, removed, or otherwise experienced an error. • `Link error crc saturated with optical errors`—CRC has exceeded the rate threshold and reached saturation with optical issues—that is, a cable has been cut, removed, or otherwise experienced an error. • **Unused**—Port is not in use.	**extensive**
SIB F2S ***slot-number***	State of the intra-chassis links between the TXP-F2S and TXP-F13 SIBs.	**extensive**

Table 262: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output
Fabric Management SIB State Output Fields for the `show chassis fabric plane extensive` Command on a TX Matrix Plus Router		

Table 262: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output
SIB *slot-number*	State of the SIBs on the T1600/T4000 router (LCC) in the routing matrix: • **Activating**—Transitional state when the SIB is coming online. • **Deactivating**—Transitional state when the SIB is going offline. • **Connected**—SIBS on an LCC are connected and trained, but are either not online or are spare, because the plane on the the TX Matrix Plus router (SFC) is still offline. The LCC SIB transitions to the **Connected** state when the F13 SIB to which it connects is online but the SFC plane (to which the LCC SIB connects) is offline for some reason; for instance, when there are insufficient number of F2 SIBs in the plane. • **Disconnected**—If an F13 SIB on the TX Matrix Plus router (SFC) goes offline, then the SIBs on the LCCs connected to the F13 SIB get disconnected. On the TX Matrix Plus router with 3D SIBs, the LCC SIB is also disconnected if the F13 SIB is online, but none of the cables are connected or trained. The **Disconnected** state is valid only for SIBs on an LCC. An LCC SIB transitions to the **Disconnected** state when the F13 SIB to which it connects goes **Offline**, irrespective of the state of the SFC plane. **SFC Error**—If an F13 SIB on the TX Matrix Plus router (SFC) transitions to the **Fault** state (because of link errors, for instance), and if an LCC SIB connected to the F13 SIB comes online, the LCC SIB transitions to the **SFC Error** state. This state indicates that the F13 SIB to which the LCC SIB is connected has errors. **NOTE**: The **Connected**, **Disconnected**, and **SFC Error** states are applicable only to the SIBs on an LCC. • **Online**—SIB is operational and running.	**extensive**

Table 262: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output

	• **Offline**—SIB is powered down.
	• **Spare**—SIB is redundant and will move to active state if one of the working SIBs fails to pass traffic.
	• **Empty**—No SIB is present.
	• **Fault**— SIB is in alarmed state where the SIB's plane is not operational for the following reasons:
	  • On-board fabric ASIC is not operational.
	  • Fiber optic connector faults.
	  • FPC connector faults.
	  • SIB midplane connector faults.
	  • Link errors have exceeded the threshold
	• **Check**—SIB is in alarmed state where the SIB is partially operational because of link or destination errors. Only a SIB that is **Online** or **Spare** can transition to the **Check** state.
	  NOTE: If a SIB is not inserted properly, the SIB cannot transition to the **Online** or **Spare** state, and therefore cannot transition to the **Check** state.	

LCC SIB Link State	State of the LCC SIB link: • **Links ok**—Link is active. • **Links error**—A link error has occurred, but the link remains operational. • **Unused**—SIB is not in use.	**extensive**

Table 262: show chassis fabric plane Output Fields *(Continued)*

Field Name	Field Description	Level of output
SG *number* Port *number*	State of the SG chip ports on the LCC: • **Links ok**—Link is active. • **Link error**—Link is operational with errors. • `Link error crc saturated`—CRC has exceeded the rate threshold and reached saturation without optical issues—that is, a cable has not been cut, removed, or otherwise experienced an error. • `Link error crc saturated with optical errors`—CRC has exceeded the rate threshold and reached saturation with optical issues—that is, a cable has been cut, removed, or otherwise experienced an error. • **Unused**—Port is not in use.	**extensive**

Sample Output

show chassis fabric plane (M120 Router)

```
user@host> show chassis fabric plane
Fabric management PLANE state
Plane 0
Plane state: ACTIVE
FEB 0: Links ok
FEB 1: Links ok
FEB 2: Links ok
FEB 3: Links ok
FEB 4: Links ok
FEB 5: Links ok
Plane 1
Plane state: ACTIVE
FEB 0: Links ok
FEB 1: Links ok
FEB 2: Links ok
FEB 3: Links ok
```

```
FEB 4: Links ok
FEB 5: Links ok
Plane 2
Plane state: ACTIVE
FEB 0: Links ok
FEB 1: Links ok
FEB 2: Links ok
FEB 3: Links ok
FEB 4: Links ok
FEB 5: Links ok
Plane 3
Plane state: ACTIVE
FEB 0: Links ok
FEB 1: Links ok
FEB 2: Links ok
FEB 3: Links ok
FEB 4: Links ok
FEB 5: Links ok
```

show chassis fabric plane (MX10004 and MX10008 Router)

```
user@host> show chassis fabric plane

Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
          PFE 3 :Links ok
          PFE 4 :Links ok
          PFE 5 :Links ok
      FPC 5
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
          PFE 3 :Links ok
          PFE 4 :Links ok
          PFE 5 :Links ok
Plane 1
```

```
  Plane state: ACTIVE
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
          PFE 3 :Links ok
          PFE 4 :Links ok
          PFE 5 :Links ok
      FPC 5
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
          PFE 3 :Links ok
          PFE 4 :Links ok
          PFE 5 :Links ok
Plane 2
  Plane state: ACTIVE
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
          PFE 3 :Links ok
          PFE 4 :Links ok
          PFE 5 :Links ok
      FPC 5
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
          PFE 3 :Links ok
          PFE 4 :Links ok
          PFE 5 :Links ok
Plane 3
  Plane state: ACTIVE
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
          PFE 3 :Links ok
          PFE 4 :Links ok
          PFE 5 :Links ok
      FPC 5
          PFE 0 :Links ok
          PFE 1 :Links ok
```

```
            PFE 2 :Links ok
            PFE 3 :Links ok
            PFE 4 :Links ok
            PFE 5 :Links ok
```

show chassis fabric plane (MX240 with AS MLC Modular Carrier Card)

In the following output, FPC 1 is the AS MLC modular carrier card (AS MCC).

```
user@host>show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
      FPC 1
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
Plane 1
  Plane state: ACTIVE
     Plane state: ACTIVE
      FPC 1
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
Plane 2
  Plane state: ACTIVE
      FPC 2
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
          PFE 2 :Links ok
      FPC 5
          PFE 0 :Links ok
Plane 3
  Plane state: ACTIVE
     Plane state: ACTIVE
      FPC 1
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
Plane 4
```

```
  Plane state: ACTIVE
      Plane state: ACTIVE
      FPC 1
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
Plane 5
  Plane state: ACTIVE
      FPC 1
          PFE 0 :Unused
      FPC 2
          PFE 0 :Links ok
Plane 6
  Plane state: ACTIVE
      FPC 1
          PFE 0 :Links ok
      FPC 2
          PFE 0 :Links ok
Plane 7
  Plane state: ACTIVE
      FPC 1
          PFE 0 :Unused
      FPC 2
          PFE 0 :Links ok
```

show chassis fabric plane (TX Matrix Plus Router)

```
user@host> show chassis fabric plane
sfc0-re0:
---------------------------------------------------------------------------
Plane  State              Link errors  Destination errors  Uptime
 0     Spare                 NONE             NONE
 1     Online                NONE             NONE         10 hours, 16 seconds
 2     Online                NONE             NONE         10 hours, 13 seconds
 3     Online                NONE             NONE         10 hours, 9 seconds
 4     Online                NONE             NONE         10 hours, 7 seconds

lcc0-re0:
---------------------------------------------------------------------------
SIB    State              Link errors  Destination errors  Uptime
```

```
 0     Spare                NONE            NONE
 1     Online               NONE            NONE         10 hours, 16 seconds
 2     Online               NONE            NONE         10 hours, 13 seconds
 3     Online               NONE            NONE         10 hours, 9 seconds
 4     Online               NONE            NONE         10 hours, 7 seconds
lcc2-re0:
--------------------------------------------------------------------------
SIB    State             Link errors  Destination errors  Uptime
 0     Spare                NONE            NONE
 1     Online               NONE            NONE         10 hours, 16 seconds
 2     Online               NONE            NONE         10 hours, 12 seconds
 3     Online               NONE            NONE         10 hours, 9 seconds
 4     Online               NONE            NONE         10 hours, 7 seconds
```

show chassis fabric plane (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric plane
sfc0-re0:
--------------------------------------------------------------------------
Plane  State        Cable errors   Link errors   Destination errors  Uptime
 0     Spare           NONE           NONE              NONE
 1     Online          NONE           NONE              NONE         5 hours, 11 minutes, 3
seconds
 2     Online          NONE           NONE              NONE         8 hours, 4 minutes, 24
seconds
 3     Online          NONE           NONE              NONE         8 hours, 3 minutes, 16
seconds
 4     Online          NONE           NONE              NONE         8 hours, 2 minutes, 12
seconds

lcc2-re0:
--------------------------------------------------------------------------
SIB    State        Cable errors   Link errors   Destination errors  Uptime
 0     Spare           NONE           NONE              NONE
 1     Online          NONE           NONE              NONE         5 hours, 11 minutes, 3
seconds
 2     Online          NONE           NONE              NONE         8 hours, 4 minutes, 57
seconds
 3     Online          NONE           NONE              NONE         8 hours, 3 minutes, 53
seconds
 4     Online          NONE           NONE              NONE         8 hours, 2 minutes, 45
```

```
seconds

lcc4-re0:
--------------------------------------------------------------------------
SIB    State        Cable errors   Link errors   Destination errors  Uptime
 0     Spare            NONE          NONE              NONE
 1     Online           NONE          NONE              NONE         5 hours, 11 minutes, 12
seconds
 2     Online           NONE          NONE              NONE         8 hours, 4 minutes, 24
seconds
 3     Online           NONE          NONE              NONE         8 hours, 3 minutes, 16
seconds
 4     Online           NONE          NONE              NONE         8 hours, 2 minutes, 12
seconds

lcc5-re0:
--------------------------------------------------------------------------
SIB    State        Cable errors   Link errors   Destination errors  Uptime
 0     Spare            NONE          NONE              NONE
 1     Online           NONE          NONE              NONE         5 hours, 11 minutes, 12
seconds
 2     Online           NONE          NONE              NONE         8 hours, 4 minutes, 24
seconds
 3     Online           NONE          NONE              NONE         8 hours, 3 minutes, 15
seconds
 4     Online           NONE          NONE              NONE         8 hours, 2 minutes, 11
seconds
```

show chassis fabric plane detail (TX Matrix Plus Router)

```
user@host> show chassis fabric plane detail
sfc0-re0:
--------------------------------------------------------------------------
Fabric Management PLANE State:
PLANE 0:    Spare
   SIB F13 0   :    Spare
   SIB F13 1   :    Empty
   SIB F2S 0/0 :    Spare
   SIB F2S 0/2 :    Spare
   SIB F2S 0/4 :    Spare
   SIB F2S 0/6 :    Spare
```

```
PLANE 1:    Online
   SIB F13 3   :    Online
   SIB F13 4   :    Empty
   SIB F2S 1/0 :    Online
   SIB F2S 1/2 :    Online
   SIB F2S 1/4 :    Online
   SIB F2S 1/6 :    Online
PLANE 2:    Online
   SIB F13 6   :    Online
   SIB F13 7   :    Empty
   SIB F2S 2/0 :    Online
   SIB F2S 2/2 :    Online
   SIB F2S 2/4 :    Online
   SIB F2S 2/6 :    Online
PLANE 3:    Online
   SIB F13 8   :    Online
   SIB F13 9   :    Online
   SIB F2S 3/0 :    Online
   SIB F2S 3/2 :    Online
   SIB F2S 3/4 :    Online
   SIB F2S 3/6 :    Online
PLANE 4:    Online
   SIB F13 11  :    Online
   SIB F13 12  :    Online
   SIB F2S 4/0 :    Online
   SIB F2S 4/2 :    Online
   SIB F2S 4/4 :    Online
   SIB F2S 4/6 :    Online

lcc0-re0:
--------------------------------------------------------------------------
Fabric Management SIB State:
   SIB     0   :    Spare
   SIB     1   :    Online
   SIB     2   :    Online
   SIB     3   :    Online
   SIB     4   :    Online

lcc1-re0:
--------------------------------------------------------------------------
Fabric Management SIB State:
   SIB     0   :    Spare
   SIB     1   :    Online
```

```
   SIB     2   :    Online
   SIB     3   :    Online
   SIB     4   :    Online
...
```

show chassis fabric plane extensive (TX Matrix Plus Router)

```
user@host> show chassis fabric plane extensive
sfc0-re0:
--------------------------------------------------------------------------
Fabric Management PLANE State:
PLANE 0:    Spare
   SIB F13 0   :    Spare
   SIB F13 1   :    Empty
   SIB F2S 0/0 :    Spare
   SIB F2S 0/2 :    Spare
   SIB F2S 0/4 :    Spare
   SIB F2S 0/6 :    Spare
   SIB F13 0 Even:
        LCC 0, SIB 0 : Links ok
            SG 0
                Port 0    : Links ok
                Port 1    : Links ok
                Port 2    : Links ok
                Port 3    : Links ok
            SG 1
                Port 0    : Links ok
                Port 1    : Links ok
                Port 2    : Links ok
                Port 3    : Links ok
            SG 2
                Port 0    : Links ok
                Port 1    : Links ok
                Port 2    : Links ok
                Port 3    : Links ok
            SG 3
                Port 0    : Links ok
                Port 1    : Links ok
                Port 2    : Links ok
                Port 3    : Links ok
   SIB F13 0 Odd:
```

```
        LCC 1, SIB 0 : Links ok
            SG 0
                Port 0    : Links ok
                Port 1    : Links ok
                Port 2    : Links ok
                Port 3    : Links ok
            SG 1
                Port 0    : Links ok
                Port 1    : Links ok
                Port 2    : Links ok
                Port 3    : Links ok
            SG 2
                Port 0    : Links ok
                Port 1    : Links ok
                Port 2    : Links ok
                Port 3    : Links ok
            SG 3
                Port 0    : Links ok
                Port 1    : Links ok
                Port 2    : Links ok
                Port 3    : Links ok
        SIB F2S 0/0: Links ok
        SIB F2S 0/2: Links ok
        SIB F2S 0/4: Links ok
        SIB F2S 0/6: Links ok
    SIB F13 1 Even:
        LCC 2, SIB 0 : Unused
            SG 0
                Port 0    : Unused
                Port 1    : Unused
                Port 2    : Unused
                Port 3    : Unused
            SG 1
                Port 0    : Unused
                Port 1    : Unused
                Port 2    : Unused
                Port 3    : Unused
            SG 2
                Port 0    : Unused
                Port 1    : Unused
                Port 2    : Unused
                Port 3    : Unused
            SG 3
```

```
                Port 0    : Unused
                Port 1    : Unused
                Port 2    : Unused
                Port 3    : Unused
   SIB F13 1 Odd:
        LCC 3, SIB 0 : Unused
            SG 0
                Port 0    : Unused
                Port 1    : Unused
                Port 2    : Unused
                Port 3    : Unused
            SG 1
                Port 0    : Unused
                Port 1    : Unused
                Port 2    : Unused
                Port 3    : Unused
            SG 2
                Port 0    : Unused
                Port 1    : Unused
                Port 2    : Unused
                Port 3    : Unused
            SG 3
                Port 0    : Unused
                Port 1    : Unused
                Port 2    : Unused
                Port 3    : Unused
        SIB F2S 0/0: Unused
        SIB F2S 0/2: Unused
        SIB F2S 0/4: Unused
        SIB F2S 0/6: Unused
PLANE 1:    Online
   SIB F13 3   :    Online
   SIB F13 4   :    Empty
   SIB F2S 1/0 :    Online
   SIB F2S 1/2 :    Online
   SIB F2S 1/4 :    Online
   SIB F2S 1/6 :    Online
   SIB F13 3 Even:
...
```

show chassis fabric plane extensive (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric plane extensive
sfc0-re0:
---------------------------------------------------------------------------
Fabric Management PLANE State:
PLANE 0:    Online
   SIB F13 0   :    Empty
   SIB F13 1   :    Online
   SIB F2S 0/0 :    Online
   SIB F2S 0/2 :    Online
   SIB F2S 0/4 :    Online
   SIB F2S 0/6 :    Online
   SIB F13 0
        LCC 0, SIB 0 : Unused
            PFE 0   : Unused
            PFE 1   : Unused
            PFE 2   : Unused
            PFE 3   : Unused
            PFE 4   : Unused
            PFE 5   : Unused
            PFE 6   : Unused
            PFE 7   : Unused
            PFE 8   : Unused
            PFE 9   : Unused
            PFE 10  : Unused
            PFE 11  : Unused
            PFE 12  : Unused
            PFE 13  : Unused
            PFE 14  : Unused
            PFE 15  : Unused
        LCC 1, SIB 0 : Unused
            PFE 0   : Unused
            PFE 1   : Unused
            PFE 2   : Unused
            PFE 3   : Unused
            PFE 4   : Unused
            PFE 5   : Unused
            PFE 6   : Unused
            PFE 7   : Unused
            PFE 8   : Unused
            PFE 9   : Unused
```

```
            PFE 10  : Unused
            PFE 11  : Unused
            PFE 12  : Unused
            PFE 13  : Unused
            PFE 14  : Unused
            PFE 15  : Unused
        LCC 2, SIB 0 : Unused
            PFE 0   : Unused
            PFE 1   : Unused
            PFE 2   : Unused
            PFE 3   : Unused
            PFE 4   : Unused
            PFE 5   : Unused
            PFE 6   : Unused
            PFE 7   : Unused
            PFE 8   : Unused
            PFE 9   : Unused
            PFE 10  : Unused
...
lcc5-re0:
--------------------------------------------------------------------------
Fabric Management SIB State:
   SIB     0   :    Online
        LCC SIB Link State : Links ok
           PFE 0   : Links ok
           PFE 1   : Links ok
           PFE 2   : Links ok
           PFE 3   : Links ok
           PFE 4   : Links ok
           PFE 5   : Links ok
           PFE 6   : Links ok
           PFE 7   : Links ok
           PFE 8   : Links ok
           PFE 9   : Links ok
           PFE 10  : Links ok
           PFE 11  : Links ok
           PFE 12  : Links ok
           PFE 13  : Links ok
           PFE 14  : Links ok
           PFE 15  : Links ok
           FPC 1
               PFE 0     : Links ok
           FPC 2
```

```
                PFE 0     : Links ok
            FPC 3
                PFE 0     : Links ok
                PFE 1     : Links ok
            FPC 4
                PFE 0     : Links ok
    SIB     1   :    Online
         LCC SIB Link State : Links ok
            PFE 0   : Links ok
            PFE 1   : Links ok
            PFE 2   : Links ok
            PFE 3   : Links ok
            PFE 4   : Links ok
            PFE 5   : Links ok
            PFE 6   : Links ok
            PFE 7   : Links ok
            PFE 8   : Links ok
            PFE 9   : Links ok
            PFE 10  : Links ok
            PFE 11  : Links ok
            PFE 12  : Links ok
            PFE 13  : Links ok
            PFE 14  : Links ok
            PFE 15  : Links ok
            FPC 1
                PFE 0     : Links ok
            FPC 2
                PFE 0     : Links ok
            FPC 3
                PFE 0     : Links ok
                PFE 1     : Links ok
            FPC 4
                PFE 0     : Links ok
```

show chassis fabric plane terse (TX Matrix Plus Router)

```
user@host> show chassis fabric plane terse
sfc0-re0:
--------------------------------------------------------------------------
Plane  State              Link errors  Destination errors  Uptime
 0     Spare                 NONE             NONE
```

```
 1     Online                NONE             NONE         18 minutes, 37 seconds
 2     Online                NONE             NONE         18 minutes, 36 seconds
 3     Online                NONE             NONE         18 minutes, 33 seconds
 4     Online                NONE             NONE         18 minutes, 31 seconds

lcc1-re0:
--------------------------------------------------------------------------
SIB    State              Link errors  Destination errors  Uptime
 0     Spare                 NONE             NONE
 1     Online                NONE             NONE         18 minutes, 37 seconds
 2     Online                NONE             NONE
 3     Online                NONE             NONE
 4     Empty                 NONE             NONE

lcc2-re0:
--------------------------------------------------------------------------
SIB    State              Link errors  Destination errors  Uptime
 0     Spare                 NONE             NONE
 1     Online                NONE             NONE         18 minutes, 37 seconds
 2     Online                NONE             NONE         18 minutes, 36 seconds
 3     Online                NONE             NONE         18 minutes, 32 seconds
 4     Online                NONE             NONE         18 minutes, 31 seconds
```

show chassis fabric plane terse (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric plane terse
sfc0-re0:
--------------------------------------------------------------------------
Plane  State        Cable errors   Link errors   Destination errors  Uptime
 0     Offline          NONE          NONE              NONE
 1     Online           NONE          NONE              NONE         1 day, 18 hours, 14
minutes, 26 seconds
 2     Offline          NONE          NONE              NONE
 3     Offline          NONE          NONE              NONE
 4     Offline          NONE          NONE              NONE

lcc2-re0:
--------------------------------------------------------------------------
SIB    State        Cable errors   Link errors   Destination errors  Uptime
 0     Offline          NONE          NONE              NONE
 1     Online           NONE          NONE              NONE         1 day, 18 hours, 17 minutes
```

```
 2     Offline          NONE          NONE              NONE
 3     Offline          NONE          NONE              NONE
 4     Offline          NONE          NONE              NONE

lcc4-re0:
--------------------------------------------------------------------------
SIB    State        Cable errors   Link errors   Destination errors  Uptime
 0     Offline          NONE          NONE              NONE
 1     Online           NONE          NONE              NONE         1 day, 18 hours, 14
minutes, 38 seconds
 2     Offline          NONE          NONE              NONE
 3     Offline          NONE          NONE              NONE
 4     Offline          NONE          NONE              NONE

lcc5-re0:
--------------------------------------------------------------------------
SIB    State        Cable errors   Link errors   Destination errors  Uptime
 0     Offline          NONE          NONE              NONE
 1     Online           NONE          NONE              NONE         1 day, 18 hours, 14
minutes, 34 seconds
 2     Offline          NONE          NONE              NONE
 3     Offline          NONE          NONE              NONE
 4     Offline          NONE          NONE              NONE
```

show chassis fabric plane lcc (TX Matrix Plus Router)

```
user@host> show chassis fabric plane lcc 1
lcc1-re0:
--------------------------------------------------------------------------
SIB    State              Link errors  Destination errors  Uptime
 0     Spare                 NONE             NONE
 1     Online                NONE             NONE         25 minutes, 17 seconds
 2     Disconnected          NONE             NONE
 3     Disconnected          NONE             NONE
 4     Empty                 NONE             NONE
```

show chassis fabric plane lcc (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric plane lcc 2
lcc2-re0:
```

```
---------------------------------------------------------------------------
SIB    State        Cable errors   Link errors   Destination errors  Uptime
 0     Offline          NONE          NONE              NONE
 1     Online           NONE          NONE              NONE         1 day, 18 hours, 16
minutes, 44 seconds
 2     Offline          NONE          NONE              NONE
 3     Offline          NONE          NONE              NONE
 4     Offline          NONE          NONE              NONE
```

show chassis fabric plane sfc (TX Matrix Plus Router)

```
user@host> show chassis fabric plane sfc 0
sfc0-re0:
---------------------------------------------------------------------------
Plane  State              Link errors  Destination errors  Uptime
 0     Spare                 NONE             NONE
 1     Online                NONE             NONE         27 minutes, 7 seconds
 2     Online                NONE             NONE         27 minutes, 6 seconds
 3     Online                NONE             NONE         27 minutes, 3 seconds
 4     Online                NONE             NONE         27 minutes, 1 second
```

show chassis fabric plane sfc (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric plane sfc 0
sfc0-re0:
---------------------------------------------------------------------------
Plane  State        Cable errors   Link errors   Destination errors  Uptime
 0     Offline          NONE          NONE              NONE
 1     Online           NONE          NONE              NONE         1 day, 18 hours, 14
minutes, 20 seconds
 2     Offline          NONE          NONE              NONE
 3     Offline          NONE          NONE              NONE
 4     Offline          NONE          NONE              NONE
```

show chassis fabric plane (T1600 Router)

```
user@host> show chassis fabric plane
Plane  State               Uptime
```

```
 0    Online               15 hours, 42 minutes, 9 seconds
 1    Online               15 hours, 42 minutes, 9 seconds
 2    Fault
 3    Online               15 hours, 42 minutes, 9 seconds
 4    Online               15 hours, 42 minutes, 9 seconds
```

show chassis fabric plane extensive (T1600 Router)

```
user@host> show chassis fabric plane extensive
Fabric Management PLANE State:
PLANE 0:    Online
   ST-SIB-L  0: Links ok
     SG 0
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
     SG 1
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
     SG 2
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
     SG 3
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
   ST-SIB-L 0
       FPC 4
           PFE 0: Links ok
           PFE 1: Links ok
       FPC 6
           PFE 0: Links ok
           PFE 1: Links ok
       FPC 7
           PFE 0: Links ok
```

```
PLANE 1:    Online
   ST-SIB-L  1: Links ok
     SG 0
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
     SG 1
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
     SG 2
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
     SG 3
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
   ST-SIB-L 1
       FPC 4
           PFE 0: Links ok
           PFE 1: Links ok
       FPC 6
           PFE 0: Links ok
           PFE 1: Links ok
       FPC 7
           PFE 0: Links ok
PLANE 2:    Online
   ST-SIB-L  2: Links ok
     SG 0
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
     SG 1
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
```

```
      SG 2
          Port 0    : Links ok
          Port 1    : Links ok
          Port 2    : Links ok
          Port 3    : Links ok
      SG 3
          Port 0    : Links ok
          Port 1    : Links ok
          Port 2    : Links ok
          Port 3    : Links ok
   ST-SIB-L 2
       FPC 4
           PFE 0: Links ok
           PFE 1: Links ok
       FPC 6
           PFE 0: Links ok
           PFE 1: Links ok
       FPC 7
           PFE 0: Links ok
PLANE 3:    Spare
   ST-SIB-L  3: Links ok
      SG 0
          Port 0    : Links ok
          Port 1    : Links ok
          Port 2    : Links ok
          Port 3    : Links ok
      SG 1
          Port 0    : Links ok
          Port 1    : Links ok
          Port 2    : Links ok
          Port 3    : Links ok
      SG 2
          Port 0    : Links ok
          Port 1    : Links ok
          Port 2    : Links ok
          Port 3    : Links ok
      SG 3
          Port 0    : Links ok
          Port 1    : Links ok
          Port 2    : Links ok
          Port 3    : Links ok
   ST-SIB-L 3
       FPC 4
```

```
            PFE 0: Links ok
            PFE 1: Links ok
        FPC 6
            PFE 0: Links ok
            PFE 1: Links ok
        FPC 7
            PFE 0: Links ok
PLANE 4:    Online
   ST-SIB-L  4: Links ok
     SG 0
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
     SG 1
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
     SG 2
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
     SG 3
         Port 0    : Links ok
         Port 1    : Links ok
         Port 2    : Links ok
         Port 3    : Links ok
   ST-SIB-L 4
        FPC 4
            PFE 0: Links ok
            PFE 1: Links ok
        FPC 6
            PFE 0: Links ok
            PFE 1: Links ok
        FPC 7
            PFE 0: Links ok
```

show chassis fabric plane detail (T1600 Router)

```
user@host> show chassis fabric plane detail
Fabric Management PLANE State:
PLANE 0:    Online
PLANE 1:    Online
PLANE 2:    Online
PLANE 3:    Spare
PLANE 4:    Online
```

show chassis fabric plane (EX8200 Switch)

```
user@host> show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
Plane 1
  Plane state: ACTIVE
Plane 2
  Plane state: ACTIVE
Plane 3
  Plane state: ACTIVE
Plane 4
  Plane state: SPARE
Plane 5
  Plane state: SPARE
Plane 6
  Plane state: SPARE
Plane 7
  Plane state: SPARE
Plane 8
  Plane state: ACTIVE
Plane 9
  Plane state: ACTIVE
Plane 10
  Plane state: ACTIVE
Plane 11
  Plane state: ACTIVE
```

show chassis fabric plane (EX9253 Switch)

```
user@switch> show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 1
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 2
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 3
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
```

```
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 4
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 5
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 6
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 7
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 8
```

```
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 9
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 10
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 11
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
Plane 12
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
```

```
            PFE 1 :Links ok
            PFE 2 :Links ok
        FPC 1
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
Plane 13
  Plane state: ACTIVE
        FPC 0
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
        FPC 1
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
Plane 14
  Plane state: ACTIVE
        FPC 0
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
        FPC 1
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
Plane 15
  Plane state: ACTIVE
        FPC 0
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
        FPC 1
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
Plane 16
  Plane state: ACTIVE
        FPC 0
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
        FPC 1
```

```
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
Plane 17
  Plane state: ACTIVE
      FPC 0
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
      FPC 1
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
Plane 18
  Plane state: ACTIVE
      FPC 0
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
      FPC 1
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
Plane 19
  Plane state: ACTIVE
      FPC 0
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
      FPC 1
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
Plane 20
  Plane state: ACTIVE
      FPC 0
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
      FPC 1
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
```

```
Plane 21
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
```

Release Information

Command introduced in Junos OS Release 8.0.

detail, extensive, lcc, sfc, and terse options introduced for the TX Matrix Plus router in Junos OS Release 9.6.

`extended` option introduced in Junos OS Release 16.1 for MX2020 and MX2010 Routers.

RELATED DOCUMENTATION

- request chassis fabric plane | **14468**
- show chassis fabric plane-location | **18156**
- Routing Matrix with a TX Matrix Plus Router Solutions Page

show chassis fabric plane-location (Security)

IN THIS SECTION

- Syntax | **18155**
- Description | **18155**
- Required Privilege Level | **18155**
- Output Fields | **18155**
- Sample Output | **18155**

Release Information | 18156

Syntax

```
show chassis fabric plane-location
```

Description

Show fabric plane location.

Required Privilege Level

view

Output Fields

Table 263 on page 18155 lists the output fields for the **show chassis fabric plane-location** command. Output fields are listed in the approximate order in which they appear.

Table 263: show chassis fabric plane-location Output Fields

Field Name	Field Description
Plane *n*	Plane number.
Control Board *n*	Control Board number.

Sample Output

show chassis fabric plane-location (SRX5600 and SRX5800 Devices with SRX5000 Line SCB II [SRX5K-SCBE] and SRX5K-RE-1800X4)

```
user@host> show chassis fabric plane-location
node0:
```

```
---------------------------------------------------------------------------
------------Fabric Plane Locations-------------
Plane 0                         Control Board 0
Plane 1                         Control Board 0
Plane 2                         Control Board 1
Plane 3                         Control Board 1
Plane 4                         Control Board 2
Plane 5                         Control Board 2

node1:
---------------------------------------------------------------------------
------------Fabric Plane Locations-------------
Plane 0                         Control Board 0
Plane 1                         Control Board 0
Plane 2                         Control Board 1
Plane 3                         Control Board 1
Plane 4                         Control Board 2
Plane 5                         Control Board 2
```

Release Information

Command introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

show chassis fabric plane (Security) | 18104

show chassis fabric plane-location

IN THIS SECTION

- Syntax | 18157
- Syntax (MX Series Routers) | 18157
- Syntax (MX2010, MX2020, MX10003, and MX2008 Universal Routing Platforms) | 18157
- Description | 18157

- Options | 18158
- Required Privilege Level | 18158
- Output Fields | 18158
- Sample Output | 18159
- Release Information | 18168

Syntax

```
show chassis fabric plane-location
```

Syntax (MX Series Routers)

```
show chassis fabric plane-location
<all-members>
<local>
<member member-id>
```

Syntax (MX2010, MX2020, MX10003, and MX2008 Universal Routing Platforms)

```
show chassis fabric plane-location
<extended>
```

Description

(M120, MX Series routers, and EX8200 switches only) Display the Control Board (CB) location of each plane. This command can be used on the primary Routing Engine or the backup Routing Engine. For information about the meaning of “CBs” and “fabric plane” on the switches, see the hardware documentation for your switch.

(TX Matrix Plus routers only) Display the SIB location of each fabric plane.

(PTX Series Packet Transport Routers and QFX Series switches only) Display the fabric plane location of each SIB.

(MX2010, MX2020, and MX2008 Routers only) Display the fabric plane location of each Switch Fabric Board (SFB).

Options

all-members (MX Series routers only) (Optional) Display the CB location of each fabric plane on the Routing Engines in all member routers in the Virtual Chassis configuration.

local (MX Series routers only) (Optional) Display the CB location of each fabric plane on the Routing Engines in the local Virtual Chassis member.

member ***member-id*** (MX Series routers only) (Optional) Display the CB location of each fabric plane on the Routing Engines in the specified member in the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

extended (MX2020, MX2010, and MX2008 routers only) (Optional) Display the fabric plane location of all 3 planes of each Switch Fabric Board (SFB) or enhanced Switch Fabric Board (SFB2).

Required Privilege Level

view

Output Fields

Table 264 on page 18158 lists the output fields for the `show chassis fabric plane-location` command. Output fields are listed in the approximate order in which they appear.

Table 264: show chassis fabric plane-location Output Fields

Field Name	Field Description
Plane *n*	Plane number. (PTX Series Packet Transport Routers and QFX Series switches) Plane numbers associated with the SIB. (MX2010, MX2020, and MX2008 Routers only) Plane numbers associated with the SFB.
Control Board *n*	Control board number.

Table 264: show chassis fabric plane-location Output Fields *(Continued)*

Field Name	Field Description
SFC ABS-SIB-F13	(TX Matrix Plus routers only) Switch Interface Board (SIB) slot number on the F13 SIB.
SFC ABS-SIB-F2S	(TX Matrix Plus routers only) SIB slot number on the F2S SIB.
LCC ST-SIB-L	(TX Matrix Plus routers only) Line-card chassis (LCC) SIB slot number.
SFC SIB F13	(TX Matrix Plus routers with 3D SIBs only) Switch Interface Board (SIB) slot number on the F13 SIB.
SFC SIB F2S	(TX Matrix Plus routers with 3D SIBs only) SIB slot number on the F2S SIB.
LCC SIB	(TX Matrix Plus routers with 3D SIBs only) Line-card chassis (LCC) SIB slot number.
SIB	(PTX Series Packet Transport Routers and QFX Series switches) SIB number.
Switch Fabric Board*n*	(MX2010, MX2020, and MX2008 Routers only) SFB number.

Sample Output

show chassis fabric plane-location (M120 Router)

```
user@host> show chassis fabric plane-location
------------Fabric Plane Locations-------------
Plane 0                         Control Board 0
Plane 1                         Control Board 0
Plane 2                         Control Board 1
Plane 3                         Control Board 1
```

show chassis fabric plane-location (MX240 and MX480 Routers)

```
user@host> show chassis fabric plane-location
------------Fabric Plane Locations-------------
Plane 0                         Control Board 0
Plane 1                         Control Board 0
Plane 2                         Control Board 0
Plane 3                         Control Board 0
Plane 4                         Control Board 1
Plane 5                         Control Board 1
Plane 6                         Control Board 1
Plane 7                         Control Board 1
```

show chassis fabric plane-location (MX960 Router)

```
user@host> show chassis fabric plane-location
------------Fabric Plane Locations-------------
Plane 0                         Control Board 0
Plane 1                         Control Board 0
Plane 2                         Control Board 1
Plane 3                         Control Board 1
Plane 4                         Control Board 2
Plane 5                         Control Board 2
```

show chassis fabric plane-location (MX10008 Router)

```
user@host> show chassis fabric plane-location

------------Fabric Plane Locations-------------
Plane 0              Switch Fabric Board 0
Plane 1              Switch Fabric Board 0
Plane 2              Switch Fabric Board 0
Plane 3              Switch Fabric Board 0
Plane 4              Switch Fabric Board 1
Plane 5              Switch Fabric Board 1
Plane 6              Switch Fabric Board 1
Plane 7              Switch Fabric Board 1
Plane 8              Switch Fabric Board 2
Plane 9              Switch Fabric Board 2
```

```
Plane 10             Switch Fabric Board 2
Plane 11             Switch Fabric Board 2
Plane 12             Switch Fabric Board 3
Plane 13             Switch Fabric Board 3
Plane 14             Switch Fabric Board 3
Plane 15             Switch Fabric Board 3
Plane 16             Switch Fabric Board 4
Plane 17             Switch Fabric Board 4
Plane 18             Switch Fabric Board 4
Plane 19             Switch Fabric Board 4
Plane 20             Switch Fabric Board 5
Plane 21             Switch Fabric Board 5
Plane 22             Switch Fabric Board 5
Plane 23             Switch Fabric Board 5
```

show chassis fabric plane-location (MX2020 Router)

```
user@host> show chassis fabric plane-location
------------Fabric Plane Locations-------------
Plane 0              Switch Fabric Board 0
Plane 1              Switch Fabric Board 1
Plane 2              Switch Fabric Board 2
Plane 3              Switch Fabric Board 3
Plane 4              Switch Fabric Board 4
Plane 5              Switch Fabric Board 5
Plane 6              Switch Fabric Board 6
Plane 7              Switch Fabric Board 7
```

show chassis fabric plane-location (MX2020 Router with SFB2)

```
user@host> show chassis fabric plane-location extended
------------Fabric Plane Locations-------------
Plane 0              Switch Fabric Board 0
Plane 1              Switch Fabric Board 0
Plane 2              Switch Fabric Board 0
Plane 3              Switch Fabric Board 1
Plane 4              Switch Fabric Board 1
Plane 5              Switch Fabric Board 1
Plane 6              Switch Fabric Board 2
Plane 7              Switch Fabric Board 2
```

```
Plane 8              Switch Fabric Board 2
Plane 9              Switch Fabric Board 3
Plane 10             Switch Fabric Board 3
Plane 11             Switch Fabric Board 3
Plane 12             Switch Fabric Board 4
Plane 13             Switch Fabric Board 4
Plane 14             Switch Fabric Board 4
Plane 15             Switch Fabric Board 5
Plane 16             Switch Fabric Board 5
Plane 17             Switch Fabric Board 5
Plane 18             Switch Fabric Board 6
Plane 19             Switch Fabric Board 6
Plane 20             Switch Fabric Board 6
Plane 21             Switch Fabric Board 7
Plane 22             Switch Fabric Board 7
Plane 23             Switch Fabric Board 7
```

show chassis fabric plane-location (MX10003 Router)

```
user@host> show chassis fabric plane-location

------------Fabric Plane Locations-------------
Plane 0
         FPC 0
         FPC 1
Plane 1
         FPC 0
         FPC 1
Plane 2
         FPC 0
         FPC 1
Plane 3
         FPC 0
         FPC 1
Plane 4
         FPC 0
         FPC 1
Plane 5
         FPC 0
         FPC 1
Plane 6
```

```
          FPC 0
          FPC 1
Plane 7
          FPC 0
          FPC 1
Plane 8
          FPC 0
          FPC 1
Plane 9
          FPC 0
          FPC 1
Plane 10
          FPC 0
          FPC 1
Plane 11
          FPC 0
          FPC 1
Plane 12
          FPC 0
          FPC 1
Plane 13
          FPC 0
          FPC 1
Plane 14
          FPC 0
          FPC 1
Plane 15
          FPC 0
          FPC 1
Plane 16
          FPC 0
          FPC 1
Plane 17
          FPC 0
          FPC 1
Plane 18
          FPC 0
          FPC 1
Plane 19
          FPC 0
          FPC 1
Plane 20
          FPC 0
```

```
         FPC 1
Plane 21
         FPC 0
         FPC 1
```

show chassis fabric plane-location (TX Matrix Plus Router)

```
user@host> show chassis fabric plane-location
Fabric Plane Locations :
Plane       SFC ABS-SIB-F13        SFC ABS-SIB-F2        LCC ST-SIB-L
 0             0, 1              0/0, 0/2, 0/4, 0/6           0
 1             3, 4              1/0, 1/2, 1/4, 1/6           1
 2             6, 7              2/0, 2/2, 2/4, 2/6           2
 3             8, 9              3/0, 3/2, 3/4, 3/6           3
 4             11, 12            4/0, 4/2, 4/4, 4/6           4
```

show chassis fabric plane-location (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric plane-location
sfc0-re0
-------------------------------------------------------------------------------
------------Fabric Plane Locations------------
Plane       SFC SIB F13            SFC SIB F2               LCC SIB
 0             0, 1              0/0, 0/2, 0/4, 0/6           0
 1             3, 4              1/0, 1/2, 1/4, 1/6           1
 2             6, 7              2/0, 2/2, 2/4, 2/6           2
 3             8, 9              3/0, 3/2, 3/4, 3/6           3
 4             11, 12            4/0, 4/2, 4/4, 4/6           4
```

show chassis fabric plane-location (EX8200 Switch)

```
user@host> show chassis fabric plane-location

------------Fabric Plane Locations-------------
Plane  0                        Control Board 0
Plane  1                        Control Board 0
Plane  2                        Control Board 0
Plane  3                        Control Board 0
```

```
Plane  4                        Control Board 1
Plane  5                        Control Board 1
Plane  6                        Control Board 1
Plane  7                        Control Board 1
Plane  8                        Control Board 2
Plane  9                        Control Board 2
Plane 10                        Control Board 2
Plane 11                        Control Board 2
```

show chassis fabric plane-location (EX9253 Switch)

```
user@switch> show chassis fabric plane-location

------------Fabric Plane Locations-------------
Plane 0
         FPC 0
         FPC 1
Plane 1
         FPC 0
         FPC 1
Plane 2
         FPC 0
         FPC 1
Plane 3
         FPC 0
         FPC 1
Plane 4
         FPC 0
         FPC 1
Plane 5
         FPC 0
         FPC 1
Plane 6
         FPC 0
         FPC 1
Plane 7
         FPC 0
         FPC 1
Plane 8
         FPC 0
         FPC 1
```

```
Plane 9
         FPC 0
         FPC 1
Plane 10
         FPC 0
         FPC 1
Plane 11
         FPC 0
         FPC 1
Plane 12
         FPC 0
         FPC 1
Plane 13
         FPC 0
         FPC 1
Plane 14
         FPC 0
         FPC 1
Plane 15
         FPC 0
         FPC 1
Plane 16
         FPC 0
         FPC 1
Plane 17
         FPC 0
         FPC 1
Plane 18
         FPC 0
         FPC 1
Plane 19
         FPC 0
         FPC 1
Plane 20
         FPC 0
         FPC 1
Plane 21
         FPC 0
         FPC 1
```

show chassis fabric plane-location (PTX Series Packet Transport Routers)

```
user@host> show chassis fabric plane-location
------------Fabric Plane Locations-------------
SIB            Planes
 0             0    1
 1             2    3
 2             4    5
 3             6    7
 4             8    9
 5            10   11
 6            12   13
 7            14   15
 8            16   17
```

show chassis fabric plane-location (PTX10008 Routers)

```
user@host> show chassis fabric plane-location
------------Fabric Plane Locations-------------
SIB            Planes
 0             0    1
 1             2    3
 2             4    5
 3             6    7
 4             8    9
 5            10   11
```

show chassis fabric plane-location (QFX 10008 Switch)

```
user@host> show chassis fabric plane-location
------------Fabric Plane Locations-------------
SIB            Planes
 0             0    1
 1             2    3
 2             4    5
 3             6    7
 4             8    9
 5            10   11
```

Release Information

Command introduced in Junos OS Release 8.0.

`extended` option introduced in Junos OS Release 16.1R1 for MX2020 and MX2010 Universal Routing Platforms.

show chassis fabric reachability

IN THIS SECTION

- Syntax | 18168
- Description | 18168
- Options | 18169
- Required Privilege Level | 18169
- Output Fields | 18169
- Sample Output | 18175
- Release Information | 18181

Syntax

```
show chassis fabric reachability
<detail>
<extended>
```

Description

(M320, MX240, MX480, MX960, PTX Series, and T Series routers only) Display the current state of fabric destination reachability. Additionally, display the details of the automated actions taken by the system to stop null-route filtering and attempt healing, and the final resolution of the actions.

Options

none Display the state of fabric destination reachability for M320, MX240, MX480, MX960, PTX Series, T640, T1600, and TX Matrix routers, based on periodic reachability checks. Display the system's action phase sequences to stop the null route and attempt healing, and the final resolution.

detail (Optional) Display the details of the actions carried out by the system in the different action phases and the final resolution.

extended (MX2020 and MX2010 Routers only) (Optional) Display the state of fabric destination reachability for MX2010 and MX2020 routers.

> **NOTE**: Fabric mode cannot be changed dynamically for MX-Series devices. Please reboot the MX device after switching the fabric mode.

Required Privilege Level

view

Output Fields

The table lists the output fields for the `show chassis fabric reachability` command. Output fields are listed in the approximate order in which they appear.

Table 265: show chassis fabric reachability Output Fields

Field Name	Field Description	Level of Output
Fabric reachability status	Display the reachability status of the fabric. • `Enabled destinations transitioned to unreachable, Fabric down action in progress`—Some enabled destinations that were originally reachable have become unreachable. The system is trying to stop the fabric down condition and attempt healing. • `Enabled destinations reachable`—The enabled destinations are reachable. • `Unreachable destinations healed`—The unreachable destinations are healed and are reachable. • `Unreachable destinations removed`—The unreachable destinations are removed. • `Unreachable destinations present`—Unreachable destinations are present in the system. • `Unreachable destinations present due to FPC restart disable configuration`—Unreachable destinations are present as a result of user configuration set to disable FPC restart.	All levels
Unreachable destinations	Number of FPCs that have unreachable destinations.	All levels
Detected on	Date and time when unreachable destinations are detected.	All levels
Reason	Reason for the destination turning unreachable. • `Single FPC error`—A single bad FPC is not reachable over the fabric. • `Fabric plane error`—Multiple FPCs are not able to forward traffic over the fabric planes.	All levels

Table 265: show chassis fabric reachability Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Fabric reachability action	Action taken to handle the unreachable destination. • `Plane Action`—The healing is attempted only for the fabric planes. • `SIB Action`—(PTX Series system only) The healing is attempted only for the SIBs. • `Plane and FPC Action`—The healing is attempted both for the fabric planes and the FPCs. • `SIB and FPC Action`—(PTX Series system only) The healing is attempted both for the SIBs and the FPCs. • `FPC Action`—The healing is attempted only for the bad FPCs.	All levels
Acting on	Current action is being performed on: • `Single FPC error`—The current operation is for healing the single bad FPC. • `Fabric Plane error`—The current operation is for healing the fabric planes.	All levels
Initial phase	Starting phase for the healing action. • `Plane restart`—The fabric planes are restarted. • `SIB restart`—(PTX Series system only) The SIBs are restarted. • `Plane and FPC restart`—Both the fabric planes and affected FPCs are restarted. • `SIB and FPC restart`—(PTX Series system only) SIBs and affected FPCs are restarted.	All levels

Table 265: show chassis fabric reachability Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Current phase	Current phase for the healing action. • `Plane restart`—The fabric planes are restarted. • `SIB restart`—(PTX Series system only) The SIBs are restarted. • `Plane and FPC restart`—Both the fabric planes and affected FPCs are restarted. • `SIB and FPC restart`—(PTX Series system only) Both the SIBs and affected FPCs are restarted. • `FPC offline`—The FPCs are turned offline because the previously mentioned healing processes have failed.	All levels
Action started	Date and time when the system fabric down healing attempt is started.	All levels
Plane restart phase	The status of the plane restart phase. • `Completed`—The plane restart phase is completed. • `In progress`—The plane restart phase is in progress.	`detail`
Phase started	Date and time when the plane restart phase is started.	`detail`
Planes restarted	List of plane numbers restarted by the system.	`detail`
Planes timed out	List of plane numbers that have timed out waiting to be restarted by the system.	`detail`
Planes being offlined / onlined	Planes that are turned offline or turned online by the system, with date and time.	`detail`
Phase completed	Date and time when the plane restart phase is completed.	`detail`

Table 265: show chassis fabric reachability Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Plane and FPC Restart Phase	Status of the plane and FPC restart phase. • `Completed`—The plane and FPC restart phase is completed. • `In progress`—The plane and FPC restart phase is in progress.	`detail`
Phase started	Date and time when the plane and FPC restart phase is started.	`detail`
FPC Offline Started	Date and time when the FPC offline action is started.	`detail`
Offlined FPCs	List of FPCs that are turned offline by the system.	`detail`
FPCs timed out	List of FPCs that have timed out waiting to be turned offline by the system.	`detail`
FPC being offlined	FPC that is being turned offline by the system, with date and time.	`detail`
FPC Offline completed	Date and time when the FPC offline action is completed.	`detail`
Plane restarting started	Date and time when the plane restart action is started.	`detail`
Planes restarted	List of planes restarted by the system.	`detail`
Planes being offlined / onlined	Planes that are currently being turned offline or turned online by the system, with date and time.	`detail`
Plane restarting completed	Date and time when the plane restarting action is completed.	`detail`
FPC online started	Date and time when FPC online action is started.	`detail`

Table 265: show chassis fabric reachability Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Onlined FPCs	List of FPCs that are turned online by the system.	`detail`
FPCs timed out	FPCs that have timed out waiting to be turned online by the system.	`detail`
FPC being onlined	FPC that is being turned online by the system, with date and time.	`detail`
FPC Online completed	Date and time when the action of turning the FPCs online is completed.	`detail`
Phase Completed	Date and time when the plane and FPC restart phase is completed.	`detail`
Phase started	Date and time when the plane and FPC restart phase is started.	`detail`
FPC restart time	Date and time when the FPC restart action is started.	`detail`
FPC restarted	FPC that is restarted by the system, with date and time.	`detail`
Phase Completed	Date and time when the plane and FPC restart phase is completed.	`detail`
FPC Offline Phase	Status of the FPC offline phase. • `Completed`— The FPC offline phase is completed. • `In progress`—The FPC offline phase is currently in progress.	`detail`
Phase started	Date and time when the FPC offline phase is started.	`detail`
FPC Offline started	Date and time when the FPC offline action is started.	`detail`
Offlined FPCs	List of FPCs turned offline by the system.	`detail`

Table 265: show chassis fabric reachability Output Fields *(Continued)*

Field Name	Field Description	Level of Output
FPCs timed out	List of FPCs that have timed out waiting to be turned offline by the system.	`detail`
FPC being offlined	FPC that is being turned offline by the system, with date and time.	`detail`
FPC Offline completed	Date and time when the FPC offline action is completed.	`detail`
Phase Completed	Date and time when the FPC offline phase is completed.	`detail`
Action Completed	Date and time when the system fabric down healing attempt is completed.	All levels
Fabric reachability resolution	Status after the healing actions are performed. • `Unreachable destinations healed after` *`phase name`*—The unreachable destinations are healed after the healing actions are performed. The phase name indicates the last healing phase. • `Unreachable destinations removed by FPCs` *`FPC number`* `offline`—The unreachable destinations are removed by turning the FPCs offline. • `Unreachable destinations present on FPC/PFE` *`FPC/PFE number`*—The unreachable destinations are present on the FPCs or Packet Forwarding Engines and need to be acted upon.	All levels

Sample Output

show chassis fabric reachability (T640 and T1600 routers)

```
user@host> show chassis fabric reachability
Fabric reachability status: Unreachable destinations removed

Fabric reachability detection:
```

```
    Unreachable destinations                : Present on 3 FPCs
        Detected on                         : 2010-11-22 15:19:45 PST
        Reason                              : Fabric plane error

Fabric reachability action:
    Fabric reachability action              : FPC action
    Acting on                               : Fabric plane error
    Initial phase                           : Plane restart
    Current phase                           : FPC offline is completed
    Action started                          : 2010-11-22 15:08:05 PST
    Action completed                        : 2010-11-22 15:19:45 PST

Fabric reachability resolution: Unreachable destinations removed by FPCs 2, 3, 5
 offline
```

show chassis fabric reachability detail (T640 and T1600 routers)

```
user@host> show chassis fabric reachability
detail
Fabric reachability status: Unreachable destinations removed
Fabric reachability detection:
    Unreachable destinations                : Present on 3 FPCs
        Detected on                         : 2010-11-15 15:50:32 PST
        Reason                              : Fabric plane error

Fabric reachability action:
    Fabric reachability action              : FPC action
    Acting on                               : Fabric plane error
    Initial phase                           : Plane restart
    Current phase                           : FPC offline is completed
    Action started                          : 2010-11-15 15:41:47 PST
        Plane restart phase                 : Completed
            Phase started                   : 2010-11-15 15:41:47 PST
                Planes restarted            : 0, 1, 2, 3, 4, 0
            Phase completed                 : 2010-11-15 15:42:14 PST
        Plane and FPC Restart Phase         : Completed
            Phase started                   : 2010-11-15 15:45:52 PST
                FPC Offline Started         : 2010-11-15 15:45:52 PST
                    Offlined FPCs           : 2, 3, 5, 7
                FPC Offline completed       : 2010-11-15 15:45:52 PST
                Plane restarting started    : 2010-11-15 15:45:52 PST
```

```
                    Planes restarted        : 0, 1, 2, 3, 4, 0
                Plane restarting completed  : 2010-11-15 15:46:11 PST
                FPC online started          : 2010-11-15 15:46:11 PST
                    Onlined FPCs            : 2, 3, 5, 7
                FPC online completed        : 2010-11-15 15:46:50 PST
            Phase completed                 : 2010-11-15 15:46:50 PST
        FPC offline phase                   : Completed
            Phase started                   : 2010-11-15 15:50:32 PST
                FPC offline started         : 2010-11-15 15:50:32 PST
                    Offlined FPCs           : 2, 3, 5
                FPC offline completed       : 2010-11-15 15:50:32 PST
            Phase completed                 : 2010-11-15 15:50:32 PST
    Action completed                        : 2010-11-15 15:50:32 PST

Fabric reachability resolution: Unreachable destinations removed by FPCs 2, 3, 5
 offline
```

show chassis fabric reachability (PTX5000 system)

```
user@host> show chassis fabric reachability
Fabric reachability status: Enabled destinations transitioned to unreachable, Fabric down action
in progress

Fabric reachability detection:
    Unreachable destinations                : Present on 5 FPCs
        Detected on                         : 2012-11-14 15:53:00 PST
        Reason                              : Fabric plane error

Fabric reachability action:
    Fabric reachability action              : SIB action
    Acting on                               : Fabric plane error
    Initial phase                           : SIB restart
    Current phase                           : SIB restart is in progress
    Action started                          : 2012-11-14 15:53:00 PST
```

show chassis fabric reachability (MX2020 Router with SFB2)

```
user@host > show chassis fabric reachability
Fabric reachability status: No Fabric degradation detected now
```

show chassis fabric reachability detail (MX10003 Router)

```
user@host > show chassis fabric reachability detail

May 23 23:52:27

Fabric reachability status: Fabric degradation condition healed
        Detected on                         : 2017-05-23 23:49:54 PDT
        Reason                              : Fabric Degradation due to Plane faults (fabric
error)

Fabric reachability action:
    Fabric reachability action              : Plane action
    Current phase                           : Plane Restart Phase is completed
    Action started                          : 2017-05-23 23:50:04 PDT
    Action completed                        : 2017-05-23 23:52:22 PDT
        Plane restart phase                 : Completed
            Phase started                   : 2017-05-23 23:50:04 PDT
                Planes restarted            : 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14,
15, 16, 17, 18, 19, 20, 21
            Phase completed                 : 2017-05-23 23:52:22 PDT
```

show chassis fabric reachability (TX Matrix router)

```
user@host> show chassis fabric reachability
Fabric reachability status: Enabled destinations transitioned to unreachable, Fabric down action
in progress

Fabric reachability detection:
    Unreachable destinations                : Present on 14 FPCs
        Detected on                         : 2013-08-29 02:09:16 PDT
        Reason                              : Fabric plane error

Fabric reachability action:
    Fabric reachability action              : Plane action
    Acting on                               : Fabric plane error
    Initial phase                           : Plane restart
    Current phase                           : Plane restart is in progress
    Action started                          : 2013-08-29 02:09:16 PDT
```

show chassis fabric reachability detail (TX Matrix router)

```
user@host> show chassis fabric reachability
detail
Fabric reachability status: Enabled destinations transitioned to unreachable, Fabric down action
in progress

Fabric reachability detection:
    Unreachable destinations                : Present on 14 FPCs
        Detected on                         : 2013-08-29 02:09:16 PDT
        Reason                              : Fabric plane error

Fabric reachability action:
    Fabric reachability action              : Plane action
    Acting on                               : Fabric plane error
    Initial phase                           : Plane restart
    Current phase                           : Plane restart is in progress
    Action started                          : 2013-08-29 02:09:16 PDT
        Plane restart phase                 : In progress
            Phase started                   : 2013-08-29 02:09:16 PDT
                Planes restarted            : 0, 2, 3
                Planes being offlined       : 4 : 2013-08-29 02:10:11 PDT
```

show chassis fabric reachability detail (MX960 router)

After triggering a phase action from an MPC, the `show chassis fabric reachability` command includes output from DPCs, MPCs, or FPCs.

```
user@host> show chassis fabric reachability
detail
Fabric reachability status: Fabric degradation condition healed
        Detected on                         : 2018-03-14 22:28:03 PDT
        Reason                              : Fabric Degradation due to grant timeouts seen by
DPCs, MPCs, or FPCs

Fabric reachability action:
    Fabric reachability action              : Plane action
    Current phase                           : Plane Restart Phase is completed
    Action started                          : 2018-03-14 22:28:17 PDT
    Action completed                        : 2018-03-14 22:29:28 PDT
```

```
        Plane restart phase                 : Completed
            Phase started                   : 2018-03-14 22:28:17 PDT
                Planes restarted            : 0, 1, 2, 3
            Phase completed                 : 2018-03-14 22:29:28 PDT

Fabric reachability resolution: Fabric degradation healed after phase Plane restart
```

show chassis fabric reachability (PTX10004, PTX10008, and PTX100016 with PTX10K-LC1202-36MR linecard)

```
user@host> show chassis fabric reachability
Fabric reachability status: No Fabric degradation detected now

Fabric reachability status: Unreachable destinations removed

Fabric reachability detection:
    Unreachable destinations                : Present on 3 FPCs
        Detected on                         : 2010-11-22 15:19:45 PST
        Reason                              : Fabric plane error

Fabric reachability action:
    Fabric reachability action              : FPC action
    Acting on                               : Fabric plane error
    Initial phase                           : Plane restart
    Current phase                           : FPC offline is completed
    Action started                          : 2010-11-22 15:08:05 PST
    Action completed                        : 2010-11-22 15:19:45 PST

Fabric reachability resolution: Unreachable destinations removed by FPCs 2, 3, 5 offline
```

show chassis fabric reachability detail (PTX10004, PTX10008, and PTX100016 with PTX10K-LC1202-36MR line card)

```
user@host> show chassis fabric reachability detail
Fabric reachability status: Fabric degradation condition healed
        Detected on                         : 2023-09-20 14:28:45 PDT
        Reason                              : Fabric Degradation due to FPC/PFE faults
Fabric reachability action:
    Fabric reachability action              : PFE Restart action
```

```
    Current phase                           : Completed
    Action started                          : 2023-09-20 14:28:45 PDT
    Action completed                        : 2023-09-20 14:35:42 PDT
        PFE restart phase                   : Completed
            Phase started                   : 2023-09-20 14:28:48 PDT
                PFEs restarted              : 0/0, 0/4
            Phase completed                 : 2023-09-20 14:30:42 PDT
Fabric reachability resolution: Fabric degradation condition healed  after PFE restart phase
```

Release Information

Command introduced before Junos OS Release 11.4.

`extended` option introduced in Junos OS Release 16.1R1 for MX2010 and MX2020 Universal Routing Platforms.

RELATED DOCUMENTATION

show chassis fabric unreachable-destinations | 18257

show chassis fabric redundancy-mode

IN THIS SECTION

- Syntax | 18182
- Description | 18182
- Options | 18182
- Required Privilege Level | 18182
- Output Fields | 18182
- Sample Output | 18182
- Release Information | 18183

Syntax

```
show chassis fabric redundancy-mode
```

Description

(MX240, MX480, and MX960 routers only) Display whether redundancy mode is configured for active control boards to enable increased fabric bandwidth usage.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 266 on page 18182 lists the output fields for the `show chassis fabric redundancy-mode` command. Output fields are listed in the approximate order in which they appear.

Table 266: show chassis fabric redundancy mode Output Fields

Field name	Field Description
Fabric redundancy mode	Currently configured mode of the fabric

Sample Output

show chassis fabric redundancy-mode

```
user@host> show chassis fabric redundancy-mode
Fabric redundancy mode: Redundant Fabric
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Detection and Corrective Actions of Line Cards with Degraded Fabric on MX Series Routers

Detection and Recovery of Fabric-Related Failures Caused by Loss of Connectivity on MX Series Routers

MX Series Routers Fabric Resiliency

redundancy-mode | **14191**

Configuring Fabric Redundancy Mode for Active Control Boards on MX Series Routers

show chassis fabric sibs

IN THIS SECTION

- Syntax | **18183**
- Description | **18184**
- Options | **18184**
- Required Privilege Level | **18184**
- Output Fields | **18184**
- Sample Output | **18187**
- Release Information | **18202**

Syntax

```
show chassis fabric sibs
<lcc number | scc>
<slot slot-number>
```

Description

(TX Matrix routers only) Display the state of the electrical and optical switch fabric link between the SIBs in the TX Matrix router (TX-SIBs) and the SIBs in the T640 routers (T640 LCC SIBs).

(M320, T640, T1600, T4000 and PTX Series routers and QFX Series switches) Display the state of the electrical switch fabric link between the SIBs and the FPCs.

Options

none (TX Matrix routers only) Display the state of the electrical and optical switch fabric link between the SIBs in the TX Matrix router (TX-SIBs) and the SIBs in the T640 routers (T640 LCC SIBs).

(M320, T640, T1600, T4000 and PTX Series routers and QFX Series switches) Display the state of the electrical switch fabric link between the SIBs and the FPCs.

lcc *number* (Optional) Display the switching fabric link state for the T640 SIBs on a specified T640 router (line-card chassis) connected to a TX Matrix router.

scc (Optional) Display the switching fabric link state for the TX-SIBs on the TX Matrix router (switch-card chassis).

slot *slot-number* (Optional) Display the state of the electrical switch fabric link between the specified SIB slot and the FPCs.

Required Privilege Level

view

Output Fields

Table 267 on page 18185 lists the output fields for the `show chassis fabric sibs` command. Output fields are listed in the approximate order in which they appear.

Table 267: show chassis fabric sibs Output Fields

Field Name	Field Description
`Fabric management SIB state`	Switching fabric link (link from FPC to SIB) state for each SIB: • `Unused`—SIB is not present. • `Links ok`—Link between the SIB and the FPC is active. • `Link error`—Link between the SIB and the FPC is not operational.

Table 267: show chassis fabric sibs Output Fields *(Continued)*

Field Name	Field Description

`Plane state`	Possible plane state of the M320 SIB, TX-SIB or T640 SIB:
	• `S_ACTIVE`—Links on the SIB are operational, and the fabric plane (SIB) is operational and running.
	• `S_SPARE`—Links on the SIB are operational and the fabric plane (SIB) is redundant and can be operational if any of the fabric planes in the `S_ACTIVE` state encounters an error.
	NOTE: If the plane is unusable by any of the Packet Forwarding Engines, the command output displays an additional string, `plane has link errors on # pfes`, where, #indicates the total number of links (both from SIB to FPC, and from FPC to SIB) having link errors (detected either during initialization time or runtime) in this particular plane. This does not count links having destination errors.
	• `S_EMPTY`—No links are present on the SIB, and the fabric plane (SIB) is powered down.
	• `S_ACTIVATING`—Links on the SIB are coming online; this is a transitional state.
	• `S_DEACTIVATING`—Links on the SIB are going offline; this is a transitional state.
	• `S_FAULTING`—Links on the SIB are being marked faulty, and the fabric plane (SIB) is not operational.
	• `S_FAULT`—Links on the SIB are in an alarmed state, and the fabric plane (SIB) is not operational for the following reasons:
	  • On-board F-chip is not operational.
	  • Fiber optic connector faults.
	  • FPC connector faults.
	Possible plane state of the QFX Series SIB:
	  • `Active`—Links on the SIB are operational, and the fabric plane (SIB) is operational and running.

Table 267: show chassis fabric sibs Output Fields *(Continued)*

Field Name	Field Description

	• `Spare`—Links on the SIB are operational and the fabric plane (SIB) is redundant and can be operational if any of the fabric planes in the `S_ACTIVE` state encounters an error.
	• `Empty`—No links are present on the SIB, and the fabric plane (SIB) is powered down.
	• `Activating`—Links on the SIB are coming online; this is a transitional state.
	• `Deactivating`—Links on the SIB are going offline; this is a transitional state.
	• `Faulting`—Links on the SIB are being marked faulty, and the fabric plane (SIB) is not operational.
	• `Fault`—Links on the SIB are in an alarmed state, and the fabric plane (SIB) is not operational for the following reasons:
	  • On-board F-chip is not operational.
	  • Fiber optic connector faults.
	  • FPC connector faults.

Sample Output

show chassis fabric sibs (M320 Router)

```
user@host> show chassis fabric sibs
Fabric management SIB state:
SIB #0
    plane state: S_ACTIVE
    FPC #0
        PFE #1  : Links ok
    FPC #1
        PFE #1  : Links ok
    FPC #2
        PFE #1  : Links ok
```

```
    FPC #3
        PFE #1  : Links ok
SIB #1
    plane state: S_ACTIVE
    FPC #0
        PFE #1  : Links ok
    FPC #1
        PFE #1  : Links ok
    FPC #2
        PFE #1  : Links ok
    FPC #3
        PFE #1  : Links ok
SIB #2
    plane state: S_ACTIVE
    FPC #0
        PFE #1  : Links ok
    FPC #1
        PFE #1  : Links ok
    FPC #2
        PFE #1  : Links ok
    FPC #3
        PFE #1  : Links ok
SIB #3
    plane state: S_ACTIVE
    FPC #0
        PFE #1  : Links ok
    FPC #1
        PFE #1  : Links ok
    FPC #2
        PFE #1  : Links ok
    FPC #3
        PFE #1  : Links ok
```

show chassis fabric sibs (T640 Router)

```
user@host> show chassis fabric sibs
Fabric management SIB state:
SIB #0
    plane state: S_SPARE
    FPC #0
        PFE #1  : Links ok
```

```
    FPC #2
        PFE #1  : Links ok
    FPC #3
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #1
    plane state: S_ACTIVE
    FPC #0
        PFE #1  : Links ok
    FPC #2
        PFE #1  : Links ok
    FPC #3
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #2
    plane state: S_ACTIVE
    FPC #0
        PFE #1  : Links ok
    FPC #2
        PFE #1  : Links ok
    FPC #3
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #3
    plane state: S_ACTIVE
    FPC #0
        PFE #1  : Links ok
    FPC #2
        PFE #1  : Links ok
    FPC #3
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #4
    plane state: S_ACTIVE
    FPC #0
        PFE #1  : Links ok
    FPC #2
        PFE #1  : Links ok
    FPC #3
        PFE #0  : Links ok
        PFE #1  : Links ok
```

show chassis fabric sibs (T1600 Router)

```
user@host> show chassis fabric sibs
SIB #0
    plane state: S_SPARE
    FPC #0
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #1
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #2
        PFE #0  : Links ok
    FPC #4
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #5
        PFE #0  : Links ok
    FPC #6
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #7
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #1
    plane state: S_ACTIVE , plane has link errors on 2 pfes
    FPC #0
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #1
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #3
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #4
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #5
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #7
```

```
        PFE #0  : Links ok
        PFE #1  : Links okSIB #2
    plane state: S_ACTIVE
SIB #2
        plane state: S_ACTIVE
FPC #0
        PFE #0  : Links ok
        PFE #1  : Links ok
FPC #1
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #2
        PFE #0  : Links ok
    FPC #4
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #5
        PFE #0  : Links ok
    FPC #6
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #7
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #3
    plane state: S_ACTIVE
    FPC #0
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #1
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #2
        PFE #0  : Links ok
    FPC #4
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #5
        PFE #0  : Links ok
    FPC #6
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #7
```

```
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #4
    plane state: S_ACTIVE
    FPC #0
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #1
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #2
        PFE #0  : Links ok
    FPC #4
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #5
        PFE #0  : Links ok
    FPC #6
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #7
        PFE #0  : Links ok
        PFE #1  : Links ok
```

show chassis fabric sibs (T4000 Core Router)

```
user@host> show chassis fabric sibs
Fabric management SIB state:
SIB #0
    plane state: S_SPARE
    FPC #2
        PFE #0  : Links ok
    FPC #3
        PFE #0  : Links ok
    FPC #5
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #6
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #1
```

```
    plane state: S_ACTIVE
    FPC #2
        PFE #0  : Links ok
    FPC #3
        PFE #0  : Links ok
    FPC #5
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #6
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #2
    plane state: S_ACTIVE
    FPC #2
        PFE #0  : Links ok
    FPC #3
        PFE #0  : Links ok
    FPC #5
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #6
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #3
    plane state: S_ACTIVE
    FPC #2
        PFE #0  : Links ok
    FPC #3
        PFE #0  : Links ok
    FPC #5
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #6
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #4
    plane state: S_ACTIVE
    FPC #2
        PFE #0  : Links ok
    FPC #3
        PFE #0  : Links ok
    FPC #5
        PFE #0  : Links ok
```

```
        PFE #1  : Links ok
    FPC #6
        PFE #0  : Links ok
        PFE #1  : Links ok
```

show chassis fabric sibs (TX Matrix Router)

```
user@host> show chassis fabric sibs
scc-re0:
---------------------------------------------------------------------------
Fabric management SIB state:
SIB #1
    plane state: S_ACTIVE , plane has link errors on 2 pfes
    FPC #0
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #1
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #3
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #4
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #5
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #7
        PFE #0  : Links ok
        PFE #1  : Links ok
SIB #2
    plane state: S_ACTIVE
    LCC #0      : Links ok
    LCC #1      : Links ok
SIB #3
    plane state: S_ACTIVE
    LCC #0      : Links ok
    LCC #1      : Links ok
SIB #4
    plane state: S_ACTIVE
```

```
    LCC #0      : Links ok
    LCC #1      : Links ok

lcc0-re0:
--------------------------------------------------------------------------
Fabric management SIB state:
SIB #1
    plane state: S_ACTIVE
    FPC #0
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #1
        PFE #1  : Links ok
    FPC #2
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #3
        PFE #1  : Links ok
    FPC #4
        PFE #1  : Links ok
    FPC #5
        PFE #0  : Links ok
    FPC #6
        PFE #1  : Links ok
    FPC #7
        PFE #1  : Links ok
    SCC         : Links ok
SIB #2
    plane state: S_ACTIVE
    FPC #0
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #1
        PFE #1  : Links ok
    FPC #2
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #3
        PFE #1  : Links ok
    FPC #4
        PFE #1  : Links ok
    FPC #5
        PFE #0  : Links ok
```

```
    FPC #6
        PFE #1  : Links ok
    FPC #7
        PFE #1  : Links ok
    SCC         : Links ok
SIB #3
    plane state: S_ACTIVE
    FPC #0
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #1
        PFE #1  : Links ok
    FPC #2
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #3
        PFE #1  : Links ok
    FPC #4
        PFE #1  : Links ok
    FPC #5
        PFE #0  : Links ok
    FPC #6
        PFE #1  : Links ok
    FPC #7
        PFE #1  : Links ok
    SCC         : Links ok
SIB #4
    plane state: S_ACTIVE
    FPC #0
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #1
        PFE #1  : Links ok
    FPC #2
        PFE #0  : Links ok
        PFE #1  : Links ok
    FPC #3
        PFE #1  : Links ok
    FPC #4
        PFE #1  : Links ok
    FPC #5
        PFE #0  : Links ok
    FPC #6
```

```
        PFE #1  : Links ok
    FPC #7
        PFE #1  : Links ok
    SCC         : Links o
```

show chassis fabric sibs lcc (TX Matrix Router)

```
user@host> show chassis fabric sibs lcc 0
lcc1-re0:
--------------------------------------------------------------------------
Fabric management SIB state:
SIB #1
    plane state: S_ACTIVE
    FPC #0
        PFE #0  : Links ok
    FPC #2
        PFE #1  : Links ok
    FPC #4
        PFE #0  : Links ok
    FPC #5
        PFE #1  : Links ok
    FPC #7
        PFE #0  : Links ok
    SCC         : Links ok
SIB #2
    plane state: S_ACTIVE
    FPC #0
        PFE #0  : Links ok
    FPC #2
        PFE #1  : Links ok
    FPC #4
        PFE #0  : Links ok
    FPC #5
        PFE #1  : Links ok
    FPC #7
        PFE #0  : Links ok
    SCC         : Links ok
SIB #3
    plane state: S_ACTIVE
    FPC #0
        PFE #0  : Links ok
```

```
    FPC #2
        PFE #1  : Links ok
    FPC #4
        PFE #0  : Links ok
    FPC #5
        PFE #1  : Links ok
    FPC #7
        PFE #0  : Links ok
    SCC         : Links ok
SIB #4
    plane state: S_ACTIVE
    FPC #0
        PFE #0  : Links ok
    FPC #2
        PFE #1  : Links ok
    FPC #4
        PFE #0  : Links ok
    FPC #5
        PFE #1  : Links ok
    FPC #7
        PFE #0  : Links ok
    SCC         : Links ok
```

show chassis fabric sibs scc (TX Matrix Router)

```
user@host> show chassis fabric sibs scc
scc-re0:
--------------------------------------------------------------------------
Fabric management SIB state:
SIB #1
    plane state: S_ACTIVE
    LCC #0      : Links ok
    LCC #1      : Links ok
SIB #2
    plane state: S_ACTIVE
    LCC #0      : Links ok
    LCC #1      : Links ok
SIB #3
    plane state: S_ACTIVE
    LCC #0      : Links ok
    LCC #1      : Links ok
```

```
SIB #4
    plane state: S_ACTIVE
    LCC #0      : Links ok
    LCC #1      : Links ok
```

show chassis fabric sibs slot (PTX3000 Router)

```
user@host> show chassis fabric sibs slot 0
Fabric management SIB state:
SIB #0  Online
    Fcore #0 (plane 0) Active
        FPC #8
            PFE #0  : OK
            PFE #1  : OK
        FPC #12
            PFE #0  : OK
            PFE #1  : OK
    Fcore #1 (plane 1) Active
        FPC #8
            PFE #0  : OK
            PFE #1  : OK
        FPC #12
            PFE #0  : OK
            PFE #1  : OK
```

show chassis fabric sibs slot (Junos OS Evolved PTX10008 and PTX10016)

The output of the `show chassis fabric sibs` command is modified to include ASIC information. In the earlier releases, this command displayed ASICs as FCOREs. With this modification, for every FPC and SIB, information about ASIC and FCORE is presented separately in the output. Each fabric chip hosts 2 planes (plane 0, 1) in PTX10008. Each fabric chip hosts a single plane (plane 0) in PTX10016.

PTX10008 supports 0-7 FPCs. PTX10016 supports 0-15 FPCs.

```
user@host> show chassis fabric sibs slot 0
Fabric management SIB state:
SIB #0  Online
    Asic #0 Fcore #0 (plane 0) Empty
            FPC #0
                PFE #0  : Links ok
```

```
                PFE #1  : Links ok
                PFE #2  : Links ok
                PFE #3  : Links ok
```

show chassis fabric sibs (PTX10008 Router)

```
user@host> show chassis fabric sibs
Fabric management SIB state:
SIB #0  Online
    FASIC #0 (plane 0) Active
        FPC #0
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
        FPC #5
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
            PFE #3  : OK
            PFE #4  : OK
            PFE #5  : OK
        FPC #6
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
            PFE #3  : OK
            PFE #4  : OK
            PFE #5  : OK
    FASIC #1 (plane 1) Active
        FPC #0
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
        FPC #5
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
            PFE #3  : OK
            PFE #4  : OK
            PFE #5  : OK
        FPC #6
```

```
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
            PFE #3  : OK
            PFE #4  : OK
            PFE #5  : OK
SIB #1  Online
    FASIC #0 (plane 2) Active
        FPC #0
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
        FPC #5
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
            PFE #3  : OK
            PFE #4  : OK
            PFE #5  : OK
        FPC #6
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
            PFE #3  : OK
            PFE #4  : OK
            PFE #5  : OK
    FASIC #1 (plane 3) Active
        FPC #0
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
        FPC #5
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
            PFE #3  : OK
            PFE #4  : OK
            PFE #5  : OK
        FPC #6
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
            PFE #3  : OK
```

```
            PFE #4  : OK
            PFE #5  : OK
SIB #2  Empty
SIB #3  Empty
SIB #4  Empty
SIB #5  Empty
```

show chassis fabric sibs (QFX10008 Switch)

```
user@host> show chassis fabric sibs
Fabric management SIB state:
SIB #0  Online
    FASIC #0 (plane 0) Active
        FPC #0
            PFE #0  : OK
            PFE #1  : OK
            PFE #2  : OK
            PFE #3  : OK
        FPC #1
            PFE #0  : OK
            PFE #1  : OK
    FASIC #1 (plane 1) Active
        FPC #0
            PFE #0  : OK
            PFE #1  : OK
        FPC #12
            PFE #0  : OK
            PFE #1  : OK
SIB #1  Empty
SIB #2  Empty
SIB #3  Empty
SIB #4  Empty
SIB #5  Empty
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced on QFX Series switches in Junos OS Release 15.1X53-D30

RELATED DOCUMENTATION

- Monitoring the SIBs
- Redundant SIBs Overview

show chassis fabric stream-info

IN THIS SECTION

- Syntax | 18203
- Description | 18203
- Additional Information | 18203
- Required Privilege Level | 18203
- Output Fields | 18204
- Sample Output | 18204
- Release Information | 18208

Syntax

```
show chassis fabric stream-info
```

Description

Display the fabric stream information.

Additional Information

This command provides the fabric stream numbering for all PFEs across all FPC slots in the chassis. The number of PFEs depends on the line card in the chassis.

Required Privilege Level

view

Output Fields

Table 268 on page 18204 lists the output fields for the `show chassis fabric stream-info` command.

The output is displayed independent of the line card and status of the FPC or PFE. The output is based on the chassis and the fabric card. Output displays the maximum PFEs defined for that chassis (for example 8 PFEs for MX2xxx series and 16 PFEs for MX1xxxx series).

Table 268: show chassis fabric stream-info Output Fields

Field Name	Field Description
FPC	FPC slot number of the Flexible PIC Concentrator (FPC).
Low/High Priority Fabric Streams	Display fabric stream ID for the PFE ID. Stream is a fabric connection from a source PFE to a destination PFE. There are two fabric streams from every source PFE to the destination PFE, low priority stream and high priority stream. The numbers (values) for low priority streams and high priority streams depend on chassis and the SFB type. The first row displays low priority fabric stream numbers. • For MX10004, MX10008 and MX10016 devices with SFB2, the numbers are from 0 to 127. • In MX2010 and MX2020 chassis with SFB2, the numbers are from 1 to 127 and 256 to 383. The second row displays high priority fabric stream numbers. • For MX10004, MX10008 and MX10016 devices with SFB2, the numbers are from 128 to 255. • In MX2010 and MX2020 chassis with SFB2, the numbers are from 128 to 255 and 384 to 511. Local PFE ID = Unique identifier for a PFE in an FPC.

Sample Output

show chassis fabric stream-info (MX10008)

```
user@host> show chassis fabric stream-info
Fabric stream info:
    First Row:  Low  Stream Numbers
```

```
    Second Row: High Stream Numbers
FPC                                     Low/High Priority Fabric Streams
        PFE0  PFE1  PFE2  PFE3  PFE4  PFE5  PFE6  PFE7  PFE8  PFE9  PFE10  PFE11  PFE12  PFE13
PFE14  PFE15
-------------------------------------------------------------------------------------------------
---------
0         0     1     2     3     4     5     6     7     8     9    10     11     12     13
14     15
        128   129   130   131   132   133   134   135   136   137   138    139    140    141
142    143
1        16    17    18    19    20    21    22    23    24    25    26     27     28     29
30     31
        144   145   146   147   148   149   150   151   152   153   154    155    156    157
158    159
2        32    33    34    35    36    37    38    39    40    41    42     43     44     45
46     47
        160   161   162   163   164   165   166   167   168   169   170    171    172    173
174    175
3        48    49    50    51    52    53    54    55    56    57    58     59     60     61
62     63
        176   177   178   179   180   181   182   183   184   185   186    187    188    189
190    191
4        64    65    66    67    68    69    70    71    72    73    74     75     76     77
78     79
        192   193   194   195   196   197   198   199   200   201   202    203    204    205
206    207
5        80    81    82    83    84    85    86    87    88    89    90     91     92     93
94     95
        208   209   210   211   212   213   214   215   216   217   218    219    220    221
222    223
6        96    97    98    99   100   101   102   103   104   105   106    107    108    109
110    111
        224   225   226   227   228   229   230   231   232   233   234    235    236    237
238    239
7       112   113   114   115   116   117   118   119   120   121   122    123    124    125
126    127
        240   241   242   243   244   245   246   247   248   249   250    251    252    253
254    255
```

show chassis fabric stream-info (MX10004)

```
user@host> show chassis fabric stream-info

Fabric stream info:
    First Row:  Low  Stream Numbers
    Second Row: High Stream Numbers
-----------------------------------------------------------------------------------------------
---------
FPC                                     Low/High Priority Fabric Streams
        PFE0  PFE1  PFE2  PFE3  PFE4  PFE5  PFE6  PFE7  PFE8  PFE9  PFE10  PFE11  PFE12  PFE13
PFE14  PFE15
-----------------------------------------------------------------------------------------------
---------
0         0     1     2     3     4     5     6     7     8     9    10     11     12     13
14     15
        128   129   130   131   132   133   134   135   136   137   138    139    140    141
142    143
1        16    17    18    19    20    21    22    23    24    25    26     27     28     29
30     31
        144   145   146   147   148   149   150   151   152   153   154    155    156    157
158    159
2        32    33    34    35    36    37    38    39    40    41    42     43     44     45
46     47
        160   161   162   163   164   165   166   167   168   169   170    171    172    173
174    175
3        48    49    50    51    52    53    54    55    56    57    58     59     60     61
62     63
        176   177   178   179   180   181   182   183   184   185   186    187    188    189
190    191
-----------------------------------------------------------------------------------------------
---------
```

show chassis fabric stream-info (MX2020 Routers)

```
user@host> show chassis fabric stream-info
Fabric stream info:
    First Row:  Low  Stream Numbers
    Second Row: High Stream Numbers
-------------------------------------------------------------
```

```
FPC             Low/High Priority Fabric Streams
        PFE0  PFE1  PFE2  PFE3  PFE4  PFE5  PFE6  PFE7
-------------------------------------------------------
0         0     1     2     3   256   257   258   259
        128   129   130   131   384   385   386   387
1         4     5     6     7   260   261   262   263
        132   133   134   135   388   389   390   391
2         8     9    10    11   264   265   266   267
        136   137   138   139   392   393   394   395
3        12    13    14    15   268   269   270   271
        140   141   142   143   396   397   398   399
4        16    17    18    19   272   273   274   275
        144   145   146   147   400   401   402   403
5        20    21    22    23   276   277   278   279
        148   149   150   151   404   405   406   407
6        24    25    26    27   280   281   282   283
        152   153   154   155   408   409   410   411
7        28    29    30    31   284   285   286   287
        156   157   158   159   412   413   414   415
8        32    33    34    35   288   289   290   291
        160   161   162   163   416   417   418   419
9        36    37    38    39   292   293   294   295
        164   165   166   167   420   421   422   423
10       40    41    42    43   296   297   298   299
        168   169   170   171   424   425   426   427
11       44    45    46    47   300   301   302   303
        172   173   174   175   428   429   430   431
12       48    49    50    51   304   305   306   307
        176   177   178   179   432   433   434   435
13       52    53    54    55   308   309   310   311
        180   181   182   183   436   437   438   439
14       56    57    58    59   312   313   314   315
        184   185   186   187   440   441   442   443
15       60    61    62    63   316   317   318   319
        188   189   190   191   444   445   446   447
16       64    65    66    67   320   321   322   323
        192   193   194   195   448   449   450   451
17       68    69    70    71   324   325   326   327
        196   197   198   199   452   453   454   455
18       72    73    74    75   328   329   330   331
        200   201   202   203   456   457   458   459
19       76    77    78    79   332   333   334   335
```

```
        204   205   206   207   460   461   462   463
-------------------------------------------------------
```

show chassis fabric stream-info (MX2010 Routers)

```
user@host> show chassis fabric stream-info
Fabric stream info:
    First Row:  Low  Stream Numbers
    Second Row: High Stream Numbers
-------------------------------------------------------
FPC             Low/High Priority Fabric Streams
        PFE0  PFE1  PFE2  PFE3  PFE4  PFE5  PFE6  PFE7
-------------------------------------------------------
0         0     1     2     3   256   257   258   259
        128   129   130   131   384   385   386   387
1         4     5     6     7   260   261   262   263
        132   133   134   135   388   389   390   391
2         8     9    10    11   264   265   266   267
        136   137   138   139   392   393   394   395
3        12    13    14    15   268   269   270   271
        140   141   142   143   396   397   398   399
4        16    17    18    19   272   273   274   275
        144   145   146   147   400   401   402   403
5        20    21    22    23   276   277   278   279
        148   149   150   151   404   405   406   407
6        24    25    26    27   280   281   282   283
        152   153   154   155   408   409   410   411
7        28    29    30    31   284   285   286   287
        156   157   158   159   412   413   414   415
8        32    33    34    35   288   289   290   291
        160   161   162   163   416   417   418   419
9        36    37    38    39   292   293   294   295
        164   165   166   167   420   421   422   423
-------------------------------------------------------
```

Release Information

Command introduced in Junos OS Release 21.4R1.

show chassis fabric summary (Security)

IN THIS SECTION

- Syntax | 18209
- Description | 18209
- Options | 18209
- Required Privilege Level | 18209
- Output Fields | 18209
- Sample Output | 18211
- Release Information | 18212

Syntax

```
show chassis fabric summary
```

Description

Show summary fabric management state.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 269 on page 18210 lists the output fields for the **show chassis fabric summary** command. Output fields are listed in the approximate order in which they appear.

Table 269: show chassis fabric summary Output Fields

Field Name	Field Description
Plane	Plane number.

State	State of the SIB or FPC:
	• **Online**—Switch Interface Board (SIB) is operational and running.
	• **Empty**—SIB is powered down.
	• **Check**—SIB is in the **Check** state because of the following reasons:
	  • SIB is not inserted properly.
	  • Some destination errors are detected on the SIB. In this case, the Packet Forwarding Engine stops using the SIB to send traffic to the affected destination Packet Forwarding Engine.
	  • Some link errors are detected on the channel between the SIB and a Packet Forwarding Engine. Link errors can be detected at initialization time or runtime:
	  • Link errors caused by a link training failure at initialization time—The Packet Forwarding Engine does not use the SIB to send traffic. The **show chassis fabric fpcs** command shows **Plane disabled** as status for this link.
	  • Link errors caused by CRC errors detected at runtime—The Packet Forwarding Engine continues to use the SIB to send traffic. The **show chassis fabric fpcs** command shows **Link error** as the status for this link.
	For information about link and destination errors, issue the **show chassis fabric fpc** commands.
	• **Spare**—SIB is redundant and will move to active state if one of the working SIBs fails.

Table 269: show chassis fabric summary Output Fields *(Continued)*

Field Name	Field Description
Errors	Indicates whether there is any error on the SIB. • **None**—No errors • **Link Errors**—Fabric link errors were found on the SIB RX link. • **Cell drops**—Fabric cell drops were found on the SIB ASIC. • **Link, Cell drops**—Both link errors and cell drops were detected on at least one of the FPC's fabric links. **NOTE**: The **Errors** column is empty only when the FPC or SIB is offline.
Uptime	Elapsed time the plane has been online.

Sample Output

show chassis fabric summary (SRX5600 and SRX5800 devices with SRX5000 line SCB II (SRX5K-SCBE) and SRX5K-RE-1800X4)

```
user@host> show chassis fabric summary
node0:
--------------------------------------------------------------------------
Plane   State    Uptime
 0      Online   14 minutes, 10 seconds
 1      Online   14 minutes, 5 seconds
 2      Online   14 minutes
 3      Online   13 minutes, 55 seconds
 4      Spare    13 minutes, 50 seconds
 5      Spare    13 minutes, 44 seconds

node1:
--------------------------------------------------------------------------
Plane   State    Uptime
 0      Online   14 minutes, 7 seconds
 1      Online   14 minutes, 2 seconds
 2      Online   13 minutes, 57 seconds
```

```
3       Online    13 minutes, 51 seconds
4       Spare     13 minutes, 46 seconds
5       Spare     13 minutes, 41 seconds
```

Release Information

Command introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

show chassis fabric plane (Security) | 18104

show chassis fabric plane-location (Security) | 18154

show chassis fabric summary

IN THIS SECTION

- Syntax | 18212
- Description | 18213
- Options | 18213
- Required Privilege Level | 18213
- Output Fields | 18213
- Sample Output - show chassis fabric summary | 18216
- Sample Output - show chassis fabric summary extended | 18218
- Release Information | 18219

Syntax

```
show chassis fabric summary <extended>
```

Description

(MX Series routers and EX8200 switches only) Display the state of all fabric planes and the elapsed uptime.

(QFX Series devices) Display the state of all fabric planes.

(PTX Series devices) Display the state of all fabric planes. Each fabric chip hosts a single plane in PTX10016. PTX10008 hosts 2 planes per fabric chip.

Options

`extended` —(Optional) Display the extended summary of fabric planes. PTX Series devices do not support the "extended" option.

Required Privilege Level

view

Output Fields

Table 270 on page 18213 lists the output fields for the `show chassis fabric summary` command. Output fields are listed in the approximate order in which they appear.

Table 270: show chassis fabric summary Output Fields

Field Name	Field Description
`Plane`	(MX Series, MX2020, MX2010, and MX2008 Routers only) Plane number.

Table 270: show chassis fabric summary Output Fields *(Continued)*

Field Name	Field Description

`State`	(MX Series and QFX Series) State of the SIB or FPC:
	• `Online`—Switch Interface Board (SIB) is operational and running.
	NOTE: On the Enhanced MX SCB with Trio MPC, a maximum of 4 planes are operational and running. On all the other SCBs with Trio MPC, all the planes are operational and running.
	• `Fault`—Switch Interface Board (SIB) has a fault.
	• `Empty`—SIB is powered down.
	• `Check`—SIB is in the `Check` state because of the following reasons:
	  • SIB is not inserted properly.
	  • Some destination errors are detected on the SIB. In this case, the Packet Forwarding Engine stops using the SIB to send traffic to the affected destination Packet Forwarding Engine.
	  • Some link errors are detected on the channel between the SIB and a Packet Forwarding Engine. Link errors can be detected at initialization time or runtime:
	  • Link errors caused by a link training failure at initialization time—The Packet Forwarding Engine does not use the SIB to send traffic. The `show chassis fabric fpcs` command shows `Plane disabled` as status for this link.
	  • Link errors caused by CRC errors detected at runtime—The Packet Forwarding Engine continues to use the SIB to send traffic. The `show chassis fabric fpcs` command shows `Link error` as the status for this link.
	NOTE: The `Check` state does not apply to PTX Series Packet Transport Routers because there are no SIBs in the Check state.
	For information about link and destination errors, issue the `show chassis fabric fpcs` commands.
	• `Spare`—SIB is redundant and will move to active state if one of the working SIBs fails.

Table 270: show chassis fabric summary Output Fields *(Continued)*

Field Name	Field Description
	NOTE: `Spare` does not apply to PTX Series Packet Transport Routers because there are no spare SIBs in the device. (MX2010, MX2020, and MX2008 Routers) State of the SFB. • `Online`—Switch Fabric Board (SFB) is operational and running. • `Offline`—Switch Fabric Board (SFB) is powered down. • `Check`—Switch Fabric Board (SFB) is in the check state.
`Errors`	(PTX Series and QFX Series) Indicates whether there is any error on the SIB. • `None`—No errors • `Link Errors`—Fabric link errors were found on the SIB RX link. • `Cell drops`—Fabric cell drops were found on the SIB ASIC. • `Link, Cell drops`—Both Link errors and cell drops were detected on at least one of the FPC's fabric links. • `Asic Errors`—A fault affecting one of the ASICs on the SIB is detected. It can be an IO error or an internal error signaled by the ASIC. **NOTE**: The `Errors` column is empty only when the FPC or SIB is offline.
`Uptime`	(MX Series, MX2010, MX2020, and MX2008 Routers) Elapsed time the plane has been online.
`Link Error`	Fabric link errors were found on the SIB RX link.
`Link TF`	Fabric link training failure has occurred.
`Destination errors`	• `Local`—Destination error detected on the FPC or PFE's own self-stream. • `Remote`—Destination error detected on the FPC or PFE's non-self-streams.

Sample Output - show chassis fabric summary

show chassis fabric summary (EX and MX Series)

```
user@host> show chassis fabric summary
Plane   State    Uptime
 0      Online   8 hours, 45 minutes, 29 seconds
 1      Online   8 hours, 45 minutes, 28 seconds
 2      Online   8 hours, 45 minutes, 28 seconds
 3      Online   8 hours, 45 minutes, 28 seconds
 4      Spare    8 hours, 45 minutes, 28 seconds
 5      Spare    8 hours, 45 minutes, 28 seconds
 6      Spare    8 hours, 45 minutes, 28 seconds
 7      Check    6 hours, 10 minutes, 12 seconds
```

show chassis fabric summary (MX Series with MPC4E)

```
user@host > show chassis fabric summary
Plane   State    Uptime
 0      Online   6 hours, 57 minutes, 44 seconds
 1      Online   6 hours, 57 minutes, 40 seconds
 2      Online   6 hours, 57 minutes, 39 seconds
 3      Online   6 hours, 57 minutes, 34 seconds
 4      Spare    6 hours, 57 minutes, 34 seconds
 5      Spare    6 hours, 57 minutes, 29 seconds
 6      Spare    6 hours, 57 minutes, 29 seconds
 7      Spare    6 hours, 57 minutes, 24 seconds
 Note:
 For FPC slots with MPC Type 4 or MCC:
   Fabric planes 1 and 5, 3 and 7 use shared physical links.
   Those slots may run in a reduced bandwidth in case both
   plane 1 and 5, or both 3 and 7 are active.
```

show chassis fabric summary (Junos OS Evolved)

```
user@host> show chassis fabric summary
Plane   State      Link   Link  Reachability errors  Uptime
```

```
                    Error  TF    Local / Remote
0      Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 24 seconds
1      Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 24 seconds
2      Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 24 seconds
3      Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 24 seconds
4      Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 24 seconds
5      Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 24 seconds
6      Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 14 seconds
7      Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 14 seconds
8      Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 14 seconds
9      Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 14 seconds
10     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 14 seconds
11     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 14 seconds
12     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 16 seconds
13     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 16 seconds
14     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 16 seconds
15     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 16 seconds
16     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 16 seconds
17     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 16 seconds
18     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 23 seconds
19     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 23 seconds
20     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 23 seconds
21     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 23 seconds
22     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 23 seconds
23     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 23 seconds
24     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 11 seconds
25     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 11 seconds
26     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 11 seconds
27     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 11 seconds
28     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 11 seconds
29     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 11 seconds
30     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 27 seconds
31     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 27 seconds
32     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 27 seconds
33     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 27 seconds
34     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 27 seconds
35     Online     NO     NO     YES  / YES         2 days, 17 hours, 2 minutes, 27 seconds
```

show chassis fabric summary (PTX and QFX Series)

```
user@host> show chassis fabric summary
FRU         State          Errors

SIB0        Online         None
SIB1        Online         Link Errors
SIB2        Online         None
SIB3        Online         Cell drops
SIB4        Offline
SIB5        Online         None
SIB6        Online         Link, Cell drops
SIB7        Online         None
SIB8        Online         Link, Cell drops

FPC0        Online         None
FPC1        Online         Link Errors
FPC2        Online         None
FPC3        Offline
FPC4        Online         None
FPC5        Online         None
FPC6        Empty
FPC7        Empty
```

Sample Output - show chassis fabric summary extended

show chassis fabric summary extended (MX Series)

```
user@host> show chassis fabric summary extended

Plane   State      Link   Link  Destination errors  Uptime
                   Error  TF    Local / Remote
 0      Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 1      Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 2      Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 3      Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
```

```
 4      Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 5      Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 6      Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 7      Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 8      Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 9      Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 10     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 11     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 12     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 13     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 14     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 15     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 16     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 17     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 18     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 19     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 20     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
 21     Online     NO     NO        NO/  NO         1 day, 10 hours, 14 minutes,        26
seconds
```

Release Information

Command introduced in Junos OS Release 8.4.

`extended` option added in Junos OS Release 14.1R2.

show chassis fabric topology

IN THIS SECTION

- Syntax (TX Matrix Router) | 18220
- Syntax (TX Matrix Plus Router) | 18220
- Syntax (T4000 Core Router) | 18220
- Syntax (PTX Series Packet Transport Routers and QFX Series Switches) | 18221
- Description | 18221
- Options | 18221
- Required Privilege Level | 18222
- Output Fields | 18222
- Sample Output | 18229
- Release Information | 18256

Syntax (TX Matrix Router)

```
show chassis fabrictopology
<lcc number | scc>
<sib-slot>
```

Syntax (TX Matrix Plus Router)

```
show chassis fabric topology
<lcc number | sfc number>
<sib-slot>
```

Syntax (T4000 Core Router)

```
show chassis fabric topology
<sib-slot>
```

Syntax (PTX Series Packet Transport Routers and QFX Series Switches)

```
show chassis fabric topology
```

Description

(TX Matrix routers only) Display the state of the switching fabric topology for the Switch Interface Board (SIB) connection between the TX Matrix router and the T640 routers.

(TX Matrix Plus routers only) Display the state of the switching fabric topology for the SIB connection between the TX Matrix Plus router and the connected routers.

(T320, T640, T1600, and T4000 routers only) Display the state of the switching fabric topology for the connection between the Switch Interface Board (SIB) and the FPCs.

(PTX Series Packet Transport Routers and QFX Series switches) Display the input-output link topology.

(PTX10002-36QDD routers only) Display physical link connectivity.

Options

none

(TX Matrix routers only) Display the state of the switching fabric topology for the Switch Interface Board (SIB) connection between the TX Matrix router and the T640 routers.

(TX Matrix Plus routers only) Display the state of the switching fabric topology for the SIB connection between the TX Matrix Plus router and the connected routers.

(T320, T640, T1600, and T4000 routers only) Display the state of the switching fabric topology for the connection between the Switch Interface Board (SIB) and the FPCs.

For PTX10008 devices, fabric ASIC value ranges from 0-2 and FPC numbering ranges from 0-7.

For PTX10016 devices, fabric ASIC ranges from 0-5 and FPC numbering ranges from 0-15.

(QFX Series switches) Display the input-output link topology.

lcc *number*

(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the fabric topology state for a specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display the fabric topology state for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

scc (TX Matrix routers only) (Optional) Display the fabric topology state for the TX Matrix router (or switch-card chassis).

sfc ***number*** (TX Matrix Plus routers only) (Optional) Display the fabric topology for the switch-fabric chassis. Replace *number* with `0`.

sib-slot (Optional) Display the fabric topology state for a specified SIB slot. Replace *sib-slot* with a value from `0` through `4`. On a TX Matrix Plus router, replace *sib-slot* with a value from `0` through `15`.

Required Privilege Level

view

Output Fields

Table 271 on page 18222 lists the output fields for the `show chassis fabric topology` command. Output fields are listed in the approximate order in which they appear.

Table 271: show chassis fabric topology Output Fields

Field Name	Field Description
`in-links`	Fabric topology for receive side links.
`out-links`	Fabric topology for transmit side links.

Table 271: show chassis fabric topology Output Fields *(Continued)*

Field Name	Field Description

`state`	State of the fabric link:
	• `RESET`—Link between the SIB and the FPC/DPC is powered down on purpose. This is done in all non-dual Packet Forwarding Engine–based boards.
	• `UP`—Link between the SIB and the FPC/DCP is up and running.
	• `DOWN`—Link between the SIB and the FPC/DCP is powered down.
	• `FAULT`—The SIB is in the alarmed state, in which the SIB's plane is not operational for one or more of the following reasons:
	  • On-board F-chip is not operational.
	  • Fiber-optic connector faults.
	  • FPC connector faults.
	  • SIB midplane connector faults.
	  NOTE: The following state descriptions are applicable only to PTX Series Packet Transport Routers.
	• `OK`—The link between the SIB and the FPC is operational.
	• `Down`—The link between the SIB and the FPC is powered down.
	• `Error`—The CCL link between the SIB and FPC is not operational for one or more of the following reasons:
	  • FPC midplane connector failure.
	  • SIB midplane connector failure.
	  • CCL link CRC error.

Table 271: show chassis fabric topology Output Fields *(Continued)*

`Out-Links:` and `In-Links` (TX Matrix	State of the links from the F13 SIB to the LCC or vice-versa. Out-Links indicate Tx links. In-Links indicate an Rx link. The following additional fields are displayed for each SIB:

Plus router only)	• `VCSEL Status`—Optical (VCSEL channel) link status for the corresponding electrical (HSL2) link. The states include:
	  • `OK`—Optical signal power is good.
	  • `Error`—Internal error.
	  • `LOS`—Loss of Signal detected.
	  • `High Cur`—The Tx Bias-current is higher than threshold on this channel. This is applicable only to Tx Channels.
	  • `Low Cur` —The Tx Bias-current is lower than threshold on this channel. This is applicable only to Tx Channels.
	• `HSL2 Channel`—HSL2 is the electrical link used to connect ASICs to the in-link and out-link. The channel number corresponds to the link and varies based on the ASIC or configuration.

- `HSL2 Status` —The status of the HSL2 Channel. Includes the following states:
- `Up`—Channel is up.
- `Down`—Channel is down.
- `Reset`—Channel has been reset.
- `Fault`—Channel has faults.

The following is a representation of display output for links originating from the SIBs (LCC or SFC)

```
SF_[1|3]_port#_FB_[A-D](VCSEL#, fiber)
```

- `SF_[1|3]`—Name of the ASIC, with Fabric F1 or F3 mode.
- `port#`—HSL2 port number on the SF ASIC in the LCC.
- `FB_[A-D]`—via fiber bundle A, B, C or D.
- `VCSEL#`—VCSEL module number on SIB.
- `fiber`—Fiber channel number.

The following is a representation of display output for links originating from the SIBs (LCC or SFC)

```
SF_[1|3]_port#_FB_[A-D](VCSEL#, fiber)
```

- `SF_[1|3]`—Name of the ASIC, with Fabric F1 or F3 mode.
- `port#`—HSL2 port number on the SF ASIC in the LCC.
- `FB_[A-D]`—via fiber bundle A, B, C or D.
- `VCSEL#`—VCSEL module number on SIB.
- `fiber`—Fiber channel number.

The following is a sample output with description of the fields displayed in the output for Out-Links:

```
Out-Links:
=========
SF_30_13_FB_A(21,09) -> FPC7_B_SG(3,3,6)_FB_A(18,09)      OK       203     Up
```

- `SF_30_13`—Name of the ASIC, with Fabric F1 or F3 mode. In this case, 3 is the F3 direction and is used in the Tx path and 0 identifies the serial link on the SF chip (in this case, link goes to sf-3 chip number 0). You can also have F1 mode and Rx path instead.
- `FB_A (21, 09)`—Fiber bundle A, with VCSEL unit number 21 within the SIB, and channel number 9 within the unit number.
- `FPC7_B_SG(3,3,6`—FPC 7.with bottom Packet Forwarding Engine (T for top PFE and B for bottom PFE), SG ASIC, with number 3 and port number 3, with HSL2 link number with the SIB as 6.
- `FB_A(18, 09)`—Fiber Bundle, with VCSEL unit number 18 within the SIB, and VCSEL channel number 9 within the unit number.

The following is a representation of display output for links originating from the FPCs (In-Links)

```
FPC#_[T|B]_SG(ASIC#, port#, HSL2_bit)_FB_[A-D](VCSEL#, fiber)
```

- `FPC#`—FPC number with PFE (0 or 1).
- `T`—Top Packet Forwarding Engine.
- `B`—Bottom Packet Forwarding Engine.
- `SG(ASIC#, port#, HSL2_bit)`—SG ASIC information (ASIC 0-3, ASIC 0-5 (PTX10016),port 0-3, HSL2_bit 0-7).
- `FB_[A-D]`—via fiber bundle A, B, C or D.
- `VCSEL#`—VCSEL module number on SIB.
- `fiber`—Fiber channel number.

The following is a representation of display output for links originating from the FPCs (In-Links)

```
FPC#_[T|B]_SG(ASIC#, port#, HSL2_bit)_FB_[A-D](VCSEL#, fiber)
```

- `FPC#`—FPC number with PFE (0 or 1).
- `T`—Top Packet Forwarding Engine.
- `B`—Bottom Packet Forwarding Engine.
- `SG(ASIC#, port#, HSL2_bit)`—SG ASIC information (ASIC 0-3, ASIC 0-5 (PTX10016) port 0-3, HSL2_bit 0-7).
- `FB_[A-D]`—via fiber bundle A, B, C or D.

- VCSEL#—VCSEL module number on SIB.
- fiber—Fiber channel number.

The following is a sample output with description of the fields displayed in the output for In-Links:

```
In-Links :
========
FPC0_T_SG(0,0,0)_FB_D(04,11)  -> SF_10_00_FB_D(01,11)     OK       0       Up
```

- FPC0—FPC 0.
- T—Top Packet Forwarding Engine.
- SG (0, 0, 0)—SG ASIC with port number 0 and link 0.
- FB_D (04,11)—Fiber Bundle D with VCSEL 4, channel 11.
- SF_10—Indicates F1 mode chip number 0 and Rx path.
- SF_10_00_FB_D(01,11) —Indicates F1 mode chip number 0 and Rx path with port 0, fiber bundle D, with VCSEL 1, channel 11.

`Out-links` and `In-links` (TX Matrix Plus router with 3D SIBs only)	State of the links from the F13 SIB to the SFC/LCC or vice-versa. Out-Links indicate Tx links. In-Links indicate an Rx link. The following additional fields are displayed for each SIB:

- Description of the fields displayed in the output for `In-links` and `Out-links` for SFC:

```
-----------------------------------------------------------------------
           In-links               State             Out-links              State
-------------------------------------------------------------------------------
CXP0_Evn->F13_SIB0_XF2,04_0      Up      F13_SIB0_XF2,04_0->CXP0_Evn      Up
```

  - `CXP0_Evn`—CXP optics with type of port bits such as even or odd. In this case, it indicates CXP optics with even port bit number 0.
  - `F13_SIB0`—Name of the SFC data plane SIB with the SIB number. In this case, it indicates F13 SIB with number 0.
  - `XF2,04_0`—Name of the ASIC with port and subchannel number. In this case, it Indicates XF2 chip with port number 4 and subchannel number 0.

- Description of the fields displayed in the output for `In-links` and `Out-links` for LCC:

```
-----------------------------------------------------------------------
            In-links               State             Out-links              State
-------------------------------------------------------------------------------
CXP0_Evn->LCC_SIB0_XF3,10_0      Up      LCC_SIB0_XF3,10_0->CXP0_Evn      Up
```

  - `CXP0_Evn`—CXP optics with the type of port bits such as even or odd. In this case, it indicates CXP optics with even port bit number 0.
  - `LCC_SIB0`—LCC SIB number. In this case, it indicates LCC SIB with number 0.
  - `XF3,10_0`—Name of the ASIC with port and subchannel number. In this case, it Indicates XF3 with port number 10 and subchannel number 0.

Sample Output

show chassis fabric topology scc (TX Matrix Router)

```
user@host> show chassis fabric topology scc
scc-re1:
--------------------------------------------------------------------------------
fchip (mode)
in-links                                state         out-links
state
--------------------------------------------------------------------------------
Sib #0 :
---------
SIB0_F0 (F2 ):
LCC0_SIB-L0_F0,03->SIB-S0_F0,00     UP                                SIB-S0_F0,00->LCC0_SIB-
L0_F1,00  UP
LCC1_SIB-L0_F0,03->SIB-S0_F0,01     UP                                SIB-S0_F0,01->LCC1_SIB-
L0_F1,08  UP
LCC2_SIB-L0_F0,03->SIB-S0_F0,02     RESET                            SIB-S0_F0,02->LCC2_SIB-
L0_F1,08  UP
LCC3_SIB-L0_F0,03->SIB-S0_F0,03     RESET                            SIB-S0_F0,03->LCC3_SIB-
L0_F1,00  UP
LCC0_SIB-L0_F0,02->SIB-S0_F0,04  UP                                 SIB-S0_F0,04->LCC0_SIB-
L0_F1,01  UP
LCC1_SIB-L0_F0,02->SIB-S0_F0,05  UP                                 SIB-S0_F0,05->LCC1_SIB-
L0_F1,09  UP
LCC2_SIB-L0_F0,02->SIB-S0_F0,06  RESET                             SIB-S0_F0,06->LCC2_SIB-
L0_F1,09  UP
LCC3_SIB-L0_F0,02->SIB-S0_F0,07  RESET                             SIB-S0_F0,07->LCC3_SIB-
L0_F1,01  UP
LCC0_SIB-L0_F0,07->SIB-S0_F0,08  UP                                 SIB-S0_F0,08->LCC0_SIB-
L0_F1,04  UP
LCC1_SIB-L0_F0,07->SIB-S0_F0,09  UP                                 SIB-S0_F0,09->LCC1_SIB-
L0_F1,12  UP
LCC2_SIB-L0_F0,07->SIB-S0_F0,10  RESET                             SIB-S0_F0,10->LCC2_SIB-
L0_F1,12  UP
LCC3_SIB-L0_F0,07->SIB-S0_F0,11  RESET                             SIB-S0_F0,11->LCC3_SIB-
L0_F1,04  UP
LCC0_SIB-L0_F0,06->SIB-S0_F0,12  UP                                 SIB-S0_F0,12->LCC0_SIB-
L0_F1,05  UP
LCC1_SIB-L0_F0,06->SIB-S0_F0,13  UP                                 SIB-S0_F0,13->LCC1_SIB-
L0_F1,13  UP
```

```
LCC2_SIB-L0_F0,06->SIB-S0_F0,14  RESET                          SIB-S0_F0,14->LCC2_SIB-
L0_F1,13  UP
LCC3_SIB-L0_F0,06->SIB-S0_F0,15  RESET                          SIB-S0_F0,15->LCC3_SIB-
L0_F1,05  UP
SIB0_F1 (F2 ):
LCC0_SIB-L0_F0,11->SIB-S0_F1,00  UP                               SIB-S0_F1,00->LCC0_SIB-
L0_F1,08  UP
LCC1_SIB-L0_F0,11->SIB-S0_F1,01  UP                               SIB-S0_F1,01->LCC1_SIB-
L0_F1,00  UP
LCC2_SIB-L0_F0,11->SIB-S0_F1,02  RESET                          SIB-S0_F1,02->LCC2_SIB-
L0_F1,00  UP
LCC3_SIB-L0_F0,11->SIB-S0_F1,03  RESET                          SIB-S0_F1,03->LCC3_SIB-
L0_F1,08  UP
LCC0_SIB-L0_F0,10->SIB-S0_F1,04  UP                               SIB-S0_F1,04->LCC0_SIB-
L0_F1,09  UP
LCC1_SIB-L0_F0,10->SIB-S0_F1,05  UP                               SIB-S0_F1,05->LCC1_SIB-
L0_F1,01  UP
LCC2_SIB-L0_F0,10->SIB-S0_F1,06  RESET                          SIB-S0_F1,06->LCC2_SIB-
L0_F1,01  UP
LCC3_SIB-L0_F0,10->SIB-S0_F1,07  RESET                          SIB-S0_F1,07->LCC3_SIB-
L0_F1,09  UP
LCC0_SIB-L0_F0,15->SIB-S0_F1,08  UP                               SIB-S0_F1,08->LCC0_SIB-
L0_F1,12  UP
LCC1_SIB-L0_F0,15->SIB-S0_F1,09  UP                               SIB-S0_F1,09->LCC1_SIB-
L0_F1,04  UP
LCC2_SIB-L0_F0,15->SIB-S0_F1,10  RESET                          SIB-S0_F1,10->LCC2_SIB-
L0_F1,04  UP
LCC3_SIB-L0_F0,15->SIB-S0_F1,11  RESET                          SIB-S0_F1,11->LCC3_SIB-
L0_F1,12  UP
LCC0_SIB-L0_F0,14->SIB-S0_F1,12  UP                               SIB-S0_F1,12->LCC0_SIB-
L0_F1,13  UP
LCC1_SIB-L0_F0,14->SIB-S0_F1,13  UP                               SIB-S0_F1,13->LCC1_SIB-
L0_F1,05  UP
LCC2_SIB-L0_F0,14->SIB-S0_F1,14  RESET                             SIB-S0_F1,14->LCC2_SIB-
L0_F1,05  UP
LCC3_SIB-L0_F0,14->SIB-S0_F1,15  RESET                             SIB-S0_F1,15->LCC3_SIB-
L0_F1,13  UP
SIB0_F2 (F2 ):
LCC3_SIB-L0_F0,13->SIB-S0_F2,00         RESET                   SIB-S0_F2,00->LCC3_SIB-
L0_F1,14  UP
LCC2_SIB-L0_F0,13->SIB-S0_F2,01  RESET                             SIB-S0_F2,01->LCC2_SIB-
L0_F1,06  UP
LCC1_SIB-L0_F0,13->SIB-S0_F2,02  UP                               SIB-S0_F2,02->LCC1_SIB-
```

```
L0_F1,06  UP
LCC0_SIB-L0_F0,13->SIB-S0_F2,03  UP                                SIB-S0_F2,03->LCC0_SIB-
L0_F1,14  UP
LCC3_SIB-L0_F0,12->SIB-S0_F2,04  RESET                            SIB-S0_F2,04->LCC3_SIB-
L0_F1,15  UP
LCC2_SIB-L0_F0,12->SIB-S0_F2,05  RESET                           SIB-S0_F2,05->LCC2_SIB-
L0_F1,07  UP
LCC1_SIB-L0_F0,12->SIB-S0_F2,06  UP                                SIB-S0_F2,06->LCC1_SIB-
L0_F1,07  UP
LCC0_SIB-L0_F0,12->SIB-S0_F2,07  UP                                SIB-S0_F2,07->LCC0_SIB-
L0_F1,15  UP
LCC3_SIB-L0_F0,09->SIB-S0_F2,08  RESET                            SIB-S0_F2,08->LCC3_SIB-
L0_F1,10  UP
LCC2_SIB-L0_F0,09->SIB-S0_F2,09  RESET                            SIB-S0_F2,09->LCC2_SIB-
L0_F1,02  UP
LCC1_SIB-L0_F0,09->SIB-S0_F2,10  UP                                SIB-S0_F2,10->LCC1_SIB-
L0_F1,02  UP
LCC0_SIB-L0_F0,09->SIB-S0_F2,11  UP                                SIB-S0_F2,11->LCC0_SIB-
L0_F1,10  UP
LCC3_SIB-L0_F0,08->SIB-S0_F2,12  RESET                            SIB-S0_F2,12->LCC3_SIB-
L0_F1,11  UP
LCC2_SIB-L0_F0,08->SIB-S0_F2,13  RESET                            SIB-S0_F2,13->LCC2_SIB-
L0_F1,03  UP
LCC1_SIB-L0_F0,08->SIB-S0_F2,14     UP                                SIB-S0_F2,14->LCC1_SIB-
L0_F1,03  UP
LCC0_SIB-L0_F0,08->SIB-S0_F2,15  UP                                SIB-S0_F2,15->LCC0_SIB-
L0_F1,11  UP
SIB0_F3 (F2 ):
LCC3_SIB-L0_F0,05->SIB-S0_F3,00  RESET                            SIB-S0_F3,00->LCC3_SIB-
L0_F1,06  UP
LCC2_SIB-L0_F0,05->SIB-S0_F3,01  RESET                            SIB-S0_F3,01->LCC2_SIB-
L0_F1,14  UP
LCC1_SIB-L0_F0,05->SIB-S0_F3,02  UP                                SIB-S0_F3,02->LCC1_SIB-
L0_F1,14  UP
LCC0_SIB-L0_F0,05->SIB-S0_F3,03  UP                                SIB-S0_F3,03->LCC0_SIB-
L0_F1,06  UP
LCC3_SIB-L0_F0,04->SIB-S0_F3,04  RESET                            SIB-S0_F3,04->LCC3_SIB-
L0_F1,07  UP
LCC2_SIB-L0_F0,04->SIB-S0_F3,05  RESET                            SIB-S0_F3,05->LCC2_SIB-
L0_F1,15  UP
LCC1_SIB-L0_F0,04->SIB-S0_F3,06  UP                                SIB-S0_F3,06->LCC1_SIB-
L0_F1,15  UP
LCC0_SIB-L0_F0,04->SIB-S0_F3,07  UP                                SIB-S0_F3,07->LCC0_SIB-
```

```
L0_F1,07  UP
LCC3_SIB-L0_F0,01->SIB-S0_F3,08  RESET                              SIB-S0_F3,08->LCC3_SIB-
L0_F1,02  UP
LCC2_SIB-L0_F0,01->SIB-S0_F3,09  RESET                              SIB-S0_F3,09->LCC2_SIB-
L0_F1,10  UP
LCC1_SIB-L0_F0,01->SIB-S0_F3,10  UP                                  SIB-S0_F3,10->LCC1_SIB-
L0_F1,10  UP
LCC0_SIB-L0_F0,01->SIB-S0_F3,11  UP                                  SIB-S0_F3,11->LCC0_SIB-
L0_F1,02  UP
LCC3_SIB-L0_F0,00->SIB-S0_F3,12  RESET                              SIB-S0_F3,12->LCC3_SIB-
L0_F1,03  UP
LCC2_SIB-L0_F0,00->SIB-S0_F3,13  RESET                              SIB-S0_F3,13->LCC2_SIB-
L0_F1,11  UP
LCC1_SIB-L0_F0,00->SIB-S0_F3,14  UP                                  SIB-S0_F3,14->LCC1_SIB-
L0_F1,11  UP
LCC0_SIB-L0_F0,00->SIB-S0_F3,15  UP                                  SIB-S0_F3,15->LCC0_SIB-
L0_F1,03  UP
Sib #1 :
---------
SIB1_F0 (F2 ):
LCC0_SIB-L1_F0,03->SIB-S1_F0,00 RESET                             SIB-S1_F0,00->LCC0_SIB-
L1_F1,00  UP
LCC1_SIB-L1_F0,03->SIB-S1_F0,01 RESET                             SIB-S1_F0,01->LCC1_SIB-
L1_F1,08  UP
LCC2_SIB-L1_F0,03->SIB-S1_F0,02 RESET                             SIB-S1_F0,02->LCC2_SIB-
L1_F1,08  UP
LCC3_SIB-L1_F0,03->SIB-S1_F0,03 RESET                             SIB-S1_F0,03->LCC3_SIB-
L1_F1,00  UP
LCC0_SIB-L1_F0,02->SIB-S1_F0,04 RESET                             SIB-S1_F0,04->LCC0_SIB-
L1_F1,01  UP
LCC1_SIB-L1_F0,02->SIB-S1_F0,05 RESET                             SIB-S1_F0,05->LCC1_SIB-
L1_F1,09  UP
LCC2_SIB-L1_F0,02->SIB-S1_F0,06 RESET                             SIB-S1_F0,06->LCC2_SIB-
L1_F1,09  UP
LCC3_SIB-L1_F0,02->SIB-S1_F0,07 RESET                             SIB-S1_F0,07->LCC3_SIB-
L1_F1,01  UP
LCC0_SIB-L1_F0,07->SIB-S1_F0,08 RESET                             SIB-S1_F0,08->LCC0_SIB-
L1_F1,04  UP
LCC1_SIB-L1_F0,07->SIB-S1_F0,09 RESET                             SIB-S1_F0,09->LCC1_SIB-
L1_F1,12  UP
LCC2_SIB-L1_F0,07->SIB-S1_F0,10 RESET                             SIB-S1_F0,10->LCC2_SIB-
L1_F1,12     UP
LCC3_SIB-L1_F0,07->SIB-S1_F0,11 RESET                             SIB-S1_F0,11->LCC3_SIB-
```

```
L1_F1,04     UP
LCC0_SIB-L1_F0,06->SIB-S1_F0,12 RESET                              SIB-S1_F0,12->LCC0_SIB-
L1_F1,05  UP
LCC1_SIB-L1_F0,06->SIB-S1_F0,13 RESET                              SIB-S1_F0,13->LCC1_SIB-
L1_F1,13  UP
LCC2_SIB-L1_F0,06->SIB-S1_F0,14 RESET                              SIB-S1_F0,14->LCC2_SIB-
L1_F1,13  UP
LCC3_SIB-L1_F0,06->SIB-S1_F0,15 RESET                              SIB-S1_F0,15->LCC3_SIB-
L1_F1,05  UP
SIB1_F1 (F2 ):
LCC0_SIB-L1_F0,11->SIB-S1_F1,00 RESET                              SIB-S1_F1,00->LCC0_SIB-
L1_F1,08  UP
LCC1_SIB-L1_F0,11->SIB-S1_F1,01 RESET                              SIB-S1_F1,01->LCC1_SIB-
L1_F1,00  UP
LCC2_SIB-L1_F0,11->SIB-S1_F1,02 RESET                              SIB-S1_F1,02->LCC2_SIB-
L1_F1,00  UP
LCC3_SIB-L1_F0,11->SIB-S1_F1,03 RESET                              SIB-S1_F1,03->LCC3_SIB-
L1_F1,08  UP
LCC0_SIB-L1_F0,10->SIB-S1_F1,04 RESET                              SIB-S1_F1,04->LCC0_SIB-
L1_F1,09     UP
LCC1_SIB-L1_F0,10->SIB-S1_F1,05 RESET                                 SIB-S1_F1,05->LCC1_SIB-
L1_F1,01  UP
LCC2_SIB-L1_F0,10->SIB-S1_F1,06 RESET                              SIB-S1_F1,06->LCC2_SIB-
L1_F1,01  UP
LCC3_SIB-L1_F0,10->SIB-S1_F1,07 RESET                              SIB-S1_F1,07->LCC3_SIB-
L1_F1,09  UP
LCC0_SIB-L1_F0,15->SIB-S1_F1,08 RESET                              SIB-S1_F1,08->LCC0_SIB-
L1_F1,12  UP
LCC1_SIB-L1_F0,15->SIB-S1_F1,09 RESET                              SIB-S1_F1,09->LCC1_SIB-
L1_F1,04  UP
LCC2_SIB-L1_F0,15->SIB-S1_F1,10 RESET                              SIB-S1_F1,10->LCC2_SIB-
L1_F1,04  UP
LCC3_SIB-L1_F0,15->SIB-S1_F1,11 RESET                              -S1_F1,11->LCC3_SIB-
L1_F1,12,05  UP
LCC0_SIB-L1_F0,14->SIB-S1_F1,12 RESET                              SIB-S1_F1,12->LCC0_SIB-
L1_F1,13  UP
LCC1_SIB-L1_F0,14->SIB-S1_F1,13 RESET                              SIB-S1_F1,13->LCC1_SIB-
L1_F1,05  UP
LCC2_SIB-L1_F0,14->SIB-S1_F1,14 RESET                              SIB-S1_F1,14->LCC2_SIB-
L1_F1,05  UP
```

show chassis fabric topology lcc

```
user@host> show chassis fabric topology lcc
0
lcc0-re0:
--------------------------------------------------------------------------
   fchip (mode)
in-links                state         out-links               state
-------------------------------------------------------------------
Sib #2 :
---------
SIB2_F0  (F1 ):
FPC0_T->SIB-L2_F0,00    DOWN       SIB-L2_F0,00->SIB-S2_F3,15 DOWN
FPC0_B->SIB-L2_F0,01    UP         SIB-L2_F0,01->SIB-S2_F3,11 DOWN
FPC1_T->SIB-L2_F0,02    DOWN       SIB-L2_F0,02->SIB-S2_F0,04 DOWN
FPC1_B->SIB-L2_F0,03    DOWN       SIB-L2_F0,03->SIB-S2_F0,00 DOWN
FPC2_T->SIB-L2_F0,04    DOWN       SIB-L2_F0,04->SIB-S2_F3,07 DOWN
FPC2_B->SIB-L2_F0,05    DOWN       SIB-L2_F0,05->SIB-S2_F3,03 DOWN
FPC3_T->SIB-L2_F0,06    DOWN       SIB-L2_F0,06->SIB-S2_F0,12 DOWN
FPC3_B->SIB-L2_F0,07    DOWN       SIB-L2_F0,07->SIB-S2_F0,08 DOWN
FPC4_T->SIB-L2_F0,08    DOWN       SIB-L2_F0,08->SIB-S2_F2,15 DOWN
FPC4_B->SIB-L2_F0,09    DOWN       SIB-L2_F0,09->SIB-S2_F2,11 DOWN
FPC5_T->SIB-L2_F0,10    DOWN       SIB-L2_F0,10->SIB-S2_F1,04 DOWN
FPC5_B->SIB-L2_F0,11    DOWN       SIB-L2_F0,11->SIB-S2_F1,00 DOWN
FPC6_T->SIB-L2_F0,12    DOWN       SIB-L2_F0,12->SIB-S2_F2,07 DOWN
FPC6_B->SIB-L2_F0,13    UP         SIB-L2_F0,13->SIB-S2_F2,03 DOWN
FPC7_T->SIB-L2_F0,14    DOWN       SIB-L2_F0,14->SIB-S2_F1,12 DOWN
FPC7_B->SIB-L2_F0,15    DOWN       SIB-L2_F0,15->SIB-S2_F1,08 DOWN
SIB2_F1  (F3 ):
SIB-S2_F0,00->SIB-L2_F1,00 UP      SIB-L2_F1,00->FPC7_B    DOWN
SIB-S2_F0,04->SIB-L2_F1,01 UP      SIB-L2_F1,01->FPC7_T    DOWN
SIB-S2_F3,11->SIB-L2_F1,02 UP      SIB-L2_F1,02->FPC6_B    DOWN
SIB-S2_F3,15->SIB-L2_F1,03 UP      SIB-L2_F1,03->FPC6_T    DOWN
SIB-S2_F0,08->SIB-L2_F1,04 UP      SIB-L2_F1,04->FPC5_B    DOWN
SIB-S2_F0,12->SIB-L2_F1,05 UP      SIB-L2_F1,05->FPC5_T    DOWN
SIB-S2_F3,03->SIB-L2_F1,06 UP      SIB-L2_F1,06->FPC4_B    DOWN
SIB-S2_F3,07->SIB-L2_F1,07 UP      SIB-L2_F1,07->FPC4_T    DOWN
SIB-S2_F1,00->SIB-L2_F1,08 UP      SIB-L2_F1,08->FPC3_B    DOWN
SIB-S2_F1,04->SIB-L2_F1,09 UP      SIB-L2_F1,09->FPC3_T    DOWN
SIB-S2_F2,11->SIB-L2_F1,10 UP      SIB-L2_F1,10->FPC2_B    DOWN
SIB-S2_F2,15->SIB-L2_F1,11 UP      SIB-L2_F1,11->FPC2_T    DOWN
SIB-S2_F1,08->SIB-L2_F1,12 UP      SIB-L2_F1,12->FPC1_B    DOWN
```

```
SIB-S2_F1,12->SIB-L2_F1,13 UP      SIB-L2_F1,13->FPC1_T    DOWN
SIB-S2_F2,03->SIB-L2_F1,14 UP      SIB-L2_F1,14->FPC0_B    DOWN
SIB-S2_F2,07->SIB-L2_F1,15 UP      SIB-L2_F1,15->FPC0_T    DOWN
Sib #4 :
---------
SIB4_F0  (F1 ):
FPC0_T->SIB-L4_F0,00    RESET      SIB-L4_F0,00->SIB-S4_F3,15 UP
FPC0_B->SIB-L4_F0,01    UP         SIB-L4_F0,01->SIB-S4_F3,11 UP
FPC1_T->SIB-L4_F0,02    RESET      SIB-L4_F0,02->SIB-S4_F0,04 UP
FPC1_B->SIB-L4_F0,03    RESET      SIB-L4_F0,03->SIB-S4_F0,00 UP
FPC2_T->SIB-L4_F0,04    RESET      SIB-L4_F0,04->SIB-S4_F3,07 UP
FPC2_B->SIB-L4_F0,05    RESET      SIB-L4_F0,05->SIB-S4_F3,03 UP
FPC3_T->SIB-L4_F0,06    RESET      SIB-L4_F0,06->SIB-S4_F0,12 UP
FPC3_B->SIB-L4_F0,07    RESET      SIB-L4_F0,07->SIB-S4_F0,08 UP
FPC4_T->SIB-L4_F0,08    RESET      SIB-L4_F0,08->SIB-S4_F2,15 UP
FPC4_B->SIB-L4_F0,09    RESET      SIB-L4_F0,09->SIB-S4_F2,11 UP
FPC5_T->SIB-L4_F0,10    RESET      SIB-L4_F0,10->SIB-S4_F1,04 UP
FPC5_B->SIB-L4_F0,11    RESET      SIB-L4_F0,11->SIB-S4_F1,00 UP
FPC6_T->SIB-L4_F0,12    RESET      SIB-L4_F0,12->SIB-S4_F2,07 UP
FPC6_B->SIB-L4_F0,13    UP         SIB-L4_F0,13->SIB-S4_F2,03 UP
FPC7_T->SIB-L4_F0,14    RESET      SIB-L4_F0,14->SIB-S4_F1,12 UP
FPC7_B->SIB-L4_F0,15    RESET      SIB-L4_F0,15->SIB-S4_F1,08 UP
SIB4_F1  (F3 ):
SIB-S4_F0,00->SIB-L4_F1,00 UP      SIB-L4_F1,00->FPC7_B    UP
SIB-S4_F0,04->SIB-L4_F1,01 UP      SIB-L4_F1,01->FPC7_T    UP
SIB-S4_F3,11->SIB-L4_F1,02 UP      SIB-L4_F1,02->FPC6_B    UP
SIB-S4_F3,15->SIB-L4_F1,03 UP      SIB-L4_F1,03->FPC6_T    UP
SIB-S4_F0,08->SIB-L4_F1,04 UP      SIB-L4_F1,04->FPC5_B    UP
SIB-S4_F0,12->SIB-L4_F1,05 UP      SIB-L4_F1,05->FPC5_T    UP
SIB-S4_F3,03->SIB-L4_F1,06 UP      SIB-L4_F1,06->FPC4_B    UP
SIB-S4_F3,07->SIB-L4_F1,07 UP      SIB-L4_F1,07->FPC4_T    UP
SIB-S4_F1,00->SIB-L4_F1,08 UP      SIB-L4_F1,08->FPC3_B    UP
SIB-S4_F1,04->SIB-L4_F1,09 UP      SIB-L4_F1,09->FPC3_T    UP
SIB-S4_F2,11->SIB-L4_F1,10 UP      SIB-L4_F1,10->FPC2_B    UP
SIB-S4_F2,15->SIB-L4_F1,11 UP      SIB-L4_F1,11->FPC2_T    UP
SIB-S4_F1,08->SIB-L4_F1,12 UP      SIB-L4_F1,12->FPC1_B    UP
SIB-S4_F1,12->SIB-L4_F1,13 UP      SIB-L4_F1,13->FPC1_T    UP
SIB-S4_F2,03->SIB-L4_F1,14 UP      SIB-L4_F1,14->FPC0_B    UP
SIB-S4_F2,07->SIB-L4_F1,15 UP      SIB-L4_F1,15->FPC0_T    UP
```

show chassis fabric topology (TX Matrix Plus Router)

```
user@host> show chassis fabric topology
sfc0-re0:
--------------------------------------------------------------------------

F13_SIB0
=========

Out-Links:
=========
SFC0_F13_SIB_00      -> LCC00_ST_SIB_L00                VCSEL    HSL2    HSL2
                                                        Status   Channel Status
==================================================================================
SF_30_00_FB_D(04,11) -> FPC0_T_SG(0,0,0)_FB_D(01,11)    OK       112     Up
SF_30_00_FB_D(04,10) -> FPC0_T_SG(0,0,1)_FB_D(01,10)    OK       112     Up
SF_30_00_FB_D(04,09) -> FPC0_T_SG(0,0,2)_FB_D(01,09)    OK       112     Up
SF_30_00_FB_D(04,08) -> FPC0_T_SG(0,0,3)_FB_D(01,08)    OK       112     Up
SF_30_00_FB_D(04,07) -> FPC0_T_SG(0,0,4)_FB_D(01,07)    OK       112     Up
SF_30_00_FB_D(04,06) -> FPC0_T_SG(0,0,5)_FB_D(01,06)    OK       112     Up
SF_30_00_FB_D(04,05) -> FPC0_T_SG(0,0,6)_FB_D(01,05)    OK       112     Up
SF_30_00_FB_D(04,04) -> FPC0_T_SG(0,0,7)_FB_D(01,04)    OK       112     Up
SF_30_01_FB_B(16,11) -> FPC4_T_SG(2,0,0)_FB_B(13,11)    OK       119     Up
SF_30_01_FB_B(16,10) -> FPC4_T_SG(2,0,1)_FB_B(13,10)    OK       119     Up
SF_30_01_FB_B(16,09) -> FPC4_T_SG(2,0,2)_FB_B(13,09)    OK       119     Up
SF_30_01_FB_B(16,08) -> FPC4_T_SG(2,0,3)_FB_B(13,08)    OK       119     Up
SF_30_01_FB_B(16,07) -> FPC4_T_SG(2,0,4)_FB_B(13,07)    OK       119     Up
SF_30_01_FB_B(16,06) -> FPC4_T_SG(2,0,5)_FB_B(13,06)    OK       119     Up
SF_30_01_FB_B(16,05) -> FPC4_T_SG(2,0,6)_FB_B(13,05)    OK       119     Up
SF_30_01_FB_B(16,04) -> FPC4_T_SG(2,0,7)_FB_B(13,04)    OK       119     Up
SF_30_02_FB_D(05,08) -> FPC1_T_SG(0,2,0)_FB_D(02,08)    OK       126     Up
SF_30_02_FB_D(05,07) -> FPC1_T_SG(0,2,1)_FB_D(02,07)    OK       126     Up
SF_30_02_FB_D(05,06) -> FPC1_T_SG(0,2,2)_FB_D(02,06)    OK       126     Up
SF_30_02_FB_D(05,05) -> FPC1_T_SG(0,2,3)_FB_D(02,05)    OK       126     Up
SF_30_02_FB_D(05,03) -> FPC1_T_SG(0,2,4)_FB_D(02,03)    OK       126     Up
SF_30_02_FB_D(05,02) -> FPC1_T_SG(0,2,5)_FB_D(02,02)    OK       126     Up
SF_30_02_FB_D(05,01) -> FPC1_T_SG(0,2,6)_FB_D(02,01)    OK       126     Up
SF_30_02_FB_D(05,00) -> FPC1_T_SG(0,2,7)_FB_D(02,00)    OK       126     Up
SF_30_03_FB_B(17,08) -> FPC5_T_SG(2,2,0)_FB_B(14,08)    OK       133     Up
SF_30_03_FB_B(17,07) -> FPC5_T_SG(2,2,1)_FB_B(14,07)    OK       133     Up
SF_30_03_FB_B(17,06) -> FPC5_T_SG(2,2,2)_FB_B(14,06)    OK       133     Up
SF_30_03_FB_B(17,05) -> FPC5_T_SG(2,2,3)_FB_B(14,05)    OK       133     Up
```

```
SF_30_03_FB_B(17,03) -> FPC5_T_SG(2,2,4)_FB_B(14,03)     OK        133      Up
SF_30_03_FB_B(17,02) -> FPC5_T_SG(2,2,5)_FB_B(14,02)     OK        133      Up
SF_30_03_FB_B(17,01) -> FPC5_T_SG(2,2,6)_FB_B(14,01)     OK        133      Up
SF_30_03_FB_B(17,00) -> FPC5_T_SG(2,2,7)_FB_B(14,00)     OK        133      Up
SF_30_04_FB_C(10,11) -> FPC2_T_SG(1,0,0)_FB_C(07,11)     OK        140      Up
SF_30_04_FB_C(10,10) -> FPC2_T_SG(1,0,1)_FB_C(07,10)     OK        140      Up
SF_30_04_FB_C(10,09) -> FPC2_T_SG(1,0,2)_FB_C(07,09)     OK        140      Up
SF_30_04_FB_C(10,08) -> FPC2_T_SG(1,0,3)_FB_C(07,08)     OK        140      Up
SF_30_04_FB_C(10,07) -> FPC2_T_SG(1,0,4)_FB_C(07,07)     OK        140      Up
SF_30_04_FB_C(10,06) -> FPC2_T_SG(1,0,5)_FB_C(07,06)     OK        140      Up
SF_30_04_FB_C(10,05) -> FPC2_T_SG(1,0,6)_FB_C(07,05)     OK        140      Up
SF_30_04_FB_C(10,04) -> FPC2_T_SG(1,0,7)_FB_C(07,04)     OK        140      Up
SF_30_05_FB_A(22,11) -> FPC6_T_SG(3,0,0)_FB_A(19,11)     OK        147      Up
SF_30_05_FB_A(22,10) -> FPC6_T_SG(3,0,1)_FB_A(19,10)     OK        147      Up
SF_30_05_FB_A(22,09) -> FPC6_T_SG(3,0,2)_FB_A(19,09)     OK        147      Up
SF_30_05_FB_A(22,08) -> FPC6_T_SG(3,0,3)_FB_A(19,08)     OK        147      Up
SF_30_05_FB_A(22,07) -> FPC6_T_SG(3,0,4)_FB_A(19,07)     OK        147      Up
SF_30_05_FB_A(22,06) -> FPC6_T_SG(3,0,5)_FB_A(19,06)     OK        147      Up
SF_30_05_FB_A(22,05) -> FPC6_T_SG(3,0,6)_FB_A(19,05)     OK        147      Up
SF_30_05_FB_A(22,04) -> FPC6_T_SG(3,0,7)_FB_A(19,04)     OK        147      Up
SF_30_06_FB_C(11,08) -> FPC3_T_SG(1,2,0)_FB_C(08,08)     OK        154      Up
SF_30_06_FB_C(11,07) -> FPC3_T_SG(1,2,1)_FB_C(08,07)     OK        154      Up
SF_30_06_FB_C(11,06) -> FPC3_T_SG(1,2,2)_FB_C(08,06)     OK        154      Up
SF_30_06_FB_C(11,05) -> FPC3_T_SG(1,2,3)_FB_C(08,05)     OK        154      Up
SF_30_06_FB_C(11,03) -> FPC3_T_SG(1,2,4)_FB_C(08,03)     OK        154      Up
SF_30_06_FB_C(11,02) -> FPC3_T_SG(1,2,5)_FB_C(08,02)     OK        154      Up
SF_30_06_FB_C(11,01) -> FPC3_T_SG(1,2,6)_FB_C(08,01)     OK        154      Up
SF_30_06_FB_C(11,00) -> FPC3_T_SG(1,2,7)_FB_C(08,00)     OK        154      Up
…
```

show chassis fabric topology sfc (TX Matrix Plus Router)

```
user@host> show chassis fabric topology sfc 0
sfc0-re0:
---------------------------------------------------------------------------

F13_SIB0
=========

Out-Links:
=========
```

```
SFC0_F13_SIB_00      -> LCC00_ST_SIB_L00                 VCSEL    HSL2    HSL2
                                                         Status   Channel Status
================================================================================
SF_30_00_FB_D(04,11) -> FPC0_T_SG(0,0,0)_FB_D(01,11)     OK       112     Up
SF_30_00_FB_D(04,10) -> FPC0_T_SG(0,0,1)_FB_D(01,10)     OK       112     Up
SF_30_00_FB_D(04,09) -> FPC0_T_SG(0,0,2)_FB_D(01,09)     OK       112     Up
SF_30_00_FB_D(04,08) -> FPC0_T_SG(0,0,3)_FB_D(01,08)     OK       112     Up
SF_30_00_FB_D(04,07) -> FPC0_T_SG(0,0,4)_FB_D(01,07)     OK       112     Up
SF_30_00_FB_D(04,06) -> FPC0_T_SG(0,0,5)_FB_D(01,06)     OK       112     Up
SF_30_00_FB_D(04,05) -> FPC0_T_SG(0,0,6)_FB_D(01,05)     OK       112     Up
SF_30_00_FB_D(04,04) -> FPC0_T_SG(0,0,7)_FB_D(01,04)     OK       112     Up
SF_30_01_FB_B(16,11) -> FPC4_T_SG(2,0,0)_FB_B(13,11)     OK       119     Up
SF_30_01_FB_B(16,10) -> FPC4_T_SG(2,0,1)_FB_B(13,10)     OK       119     Up
SF_30_01_FB_B(16,09) -> FPC4_T_SG(2,0,2)_FB_B(13,09)     OK       119     Up
SF_30_01_FB_B(16,08) -> FPC4_T_SG(2,0,3)_FB_B(13,08)     OK       119     Up
SF_30_01_FB_B(16,07) -> FPC4_T_SG(2,0,4)_FB_B(13,07)     OK       119     Up
SF_30_01_FB_B(16,06) -> FPC4_T_SG(2,0,5)_FB_B(13,06)     OK       119     Up
SF_30_01_FB_B(16,05) -> FPC4_T_SG(2,0,6)_FB_B(13,05)     OK       119     Up
SF_30_01_FB_B(16,04) -> FPC4_T_SG(2,0,7)_FB_B(13,04)     OK       119     Up
SF_30_02_FB_D(05,08) -> FPC1_T_SG(0,2,0)_FB_D(02,08)     OK       126     Up
SF_30_02_FB_D(05,07) -> FPC1_T_SG(0,2,1)_FB_D(02,07)     OK       126     Up
SF_30_02_FB_D(05,06) -> FPC1_T_SG(0,2,2)_FB_D(02,06)     OK       126     Up
SF_30_02_FB_D(05,05) -> FPC1_T_SG(0,2,3)_FB_D(02,05)     OK       126     Up
SF_30_02_FB_D(05,03) -> FPC1_T_SG(0,2,4)_FB_D(02,03)     OK       126     Up
SF_30_02_FB_D(05,02) -> FPC1_T_SG(0,2,5)_FB_D(02,02)     OK       126     Up
SF_30_02_FB_D(05,01) -> FPC1_T_SG(0,2,6)_FB_D(02,01)     OK       126     Up
SF_30_02_FB_D(05,00) -> FPC1_T_SG(0,2,7)_FB_D(02,00)     OK       126     Up
SF_30_03_FB_B(17,08) -> FPC5_T_SG(2,2,0)_FB_B(14,08)     OK       133     Up
SF_30_03_FB_B(17,07) -> FPC5_T_SG(2,2,1)_FB_B(14,07)     OK       133     Up
SF_30_03_FB_B(17,06) -> FPC5_T_SG(2,2,2)_FB_B(14,06)     OK       133     Up
SF_30_03_FB_B(17,05) -> FPC5_T_SG(2,2,3)_FB_B(14,05)     OK       133     Up
SF_30_03_FB_B(17,03) -> FPC5_T_SG(2,2,4)_FB_B(14,03)     OK       133     Up
SF_30_03_FB_B(17,02) -> FPC5_T_SG(2,2,5)_FB_B(14,02)     OK       133     Up
SF_30_03_FB_B(17,01) -> FPC5_T_SG(2,2,6)_FB_B(14,01)     OK       133     Up
SF_30_03_FB_B(17,00) -> FPC5_T_SG(2,2,7)_FB_B(14,00)     OK       133     Up
SF_30_04_FB_C(10,11) -> FPC2_T_SG(1,0,0)_FB_C(07,11)     OK       140     Up
SF_30_04_FB_C(10,10) -> FPC2_T_SG(1,0,1)_FB_C(07,10)     OK       140     Up
SF_30_04_FB_C(10,09) -> FPC2_T_SG(1,0,2)_FB_C(07,09)     OK       140     Up
SF_30_04_FB_C(10,08) -> FPC2_T_SG(1,0,3)_FB_C(07,08)     OK       140     Up
SF_30_04_FB_C(10,07) -> FPC2_T_SG(1,0,4)_FB_C(07,07)     OK       140     Up
SF_30_04_FB_C(10,06) -> FPC2_T_SG(1,0,5)_FB_C(07,06)     OK       140     Up
SF_30_04_FB_C(10,05) -> FPC2_T_SG(1,0,6)_FB_C(07,05)     OK       140     Up
SF_30_04_FB_C(10,04) -> FPC2_T_SG(1,0,7)_FB_C(07,04)     OK       140     Up
```

```
SF_30_05_FB_A(22,11) -> FPC6_T_SG(3,0,0)_FB_A(19,11)     OK        147     Up
SF_30_05_FB_A(22,10) -> FPC6_T_SG(3,0,1)_FB_A(19,10)     OK        147     Up
SF_30_05_FB_A(22,09) -> FPC6_T_SG(3,0,2)_FB_A(19,09)     OK        147     Up
SF_30_05_FB_A(22,08) -> FPC6_T_SG(3,0,3)_FB_A(19,08)     OK        147     Up
SF_30_05_FB_A(22,07) -> FPC6_T_SG(3,0,4)_FB_A(19,07)     OK        147     Up
SF_30_05_FB_A(22,06) -> FPC6_T_SG(3,0,5)_FB_A(19,06)     OK        147     Up
SF_30_05_FB_A(22,05) -> FPC6_T_SG(3,0,6)_FB_A(19,05)     OK        147     Up
SF_30_05_FB_A(22,04) -> FPC6_T_SG(3,0,7)_FB_A(19,04)     OK        147     Up
SF_30_06_FB_C(11,08) -> FPC3_T_SG(1,2,0)_FB_C(08,08)     OK        154     Up
SF_30_06_FB_C(11,07) -> FPC3_T_SG(1,2,1)_FB_C(08,07)     OK        154     Up
SF_30_06_FB_C(11,06) -> FPC3_T_SG(1,2,2)_FB_C(08,06)     OK        154     Up
SF_30_06_FB_C(11,05) -> FPC3_T_SG(1,2,3)_FB_C(08,05)     OK        154     Up
SF_30_06_FB_C(11,03) -> FPC3_T_SG(1,2,4)_FB_C(08,03)     OK        154     Up
SF_30_06_FB_C(11,02) -> FPC3_T_SG(1,2,5)_FB_C(08,02)     OK        154     Up
SF_30_06_FB_C(11,01) -> FPC3_T_SG(1,2,6)_FB_C(08,01)     OK        154     Up
SF_30_06_FB_C(11,00) -> FPC3_T_SG(1,2,7)_FB_C(08,00)     OK        154     Up
…
```

show chassis fabric topology lcc (TX Matrix Plus Router)

```
user@host> show chassis fabric topology lcc 0
lcc0-re0:
--------------------------------------------------------------------------

SIB0
=========

Out-Links:
=========
LCC00_ST_SIB_L00             -> SFC0_F13_SIB_00          VCSEL    HSL2    HSL2
                                                         Status   Channel Status
================================================================================
FPC0_T_SG(0,0,0)_FB_D(04,11) -> SF_10_00_FB_D(01,11)     OK       12      Up
FPC0_T_SG(0,0,1)_FB_D(04,10) -> SF_10_00_FB_D(01,10)     OK       12      Up
FPC0_T_SG(0,0,2)_FB_D(04,09) -> SF_10_00_FB_D(01,09)     OK       12      Up
FPC0_T_SG(0,0,3)_FB_D(04,08) -> SF_10_00_FB_D(01,08)     OK       12      Up
FPC0_T_SG(0,0,4)_FB_D(04,07) -> SF_10_00_FB_D(01,07)     OK       12      Up
FPC0_T_SG(0,0,5)_FB_D(04,06) -> SF_10_00_FB_D(01,06)     OK       12      Up
FPC0_T_SG(0,0,6)_FB_D(04,05) -> SF_10_00_FB_D(01,05)     OK       12      Up
FPC0_T_SG(0,0,7)_FB_D(04,04) -> SF_10_00_FB_D(01,04)     OK       12      Up
FPC0_B_SG(0,1,0)_FB_D(03,07) -> SF_10_10_FB_D(00,07)     OK       15      Up
```

```
FPC0_B_SG(0,1,1)_FB_D(03,06) -> SF_10_10_FB_D(00,06)     OK        15       Up
FPC0_B_SG(0,1,2)_FB_D(03,05) -> SF_10_10_FB_D(00,05)     OK        15       Up
FPC0_B_SG(0,1,3)_FB_D(03,04) -> SF_10_10_FB_D(00,04)     OK        15       Up
FPC0_B_SG(0,1,4)_FB_D(03,03) -> SF_10_10_FB_D(00,03)     OK        15       Up
FPC0_B_SG(0,1,5)_FB_D(03,02) -> SF_10_10_FB_D(00,02)     OK        15       Up
FPC0_B_SG(0,1,6)_FB_D(03,01) -> SF_10_10_FB_D(00,01)     OK        15       Up
FPC0_B_SG(0,1,7)_FB_D(03,00) -> SF_10_10_FB_D(00,00)     OK        15       Up
FPC1_T_SG(0,2,0)_FB_D(05,08) -> SF_10_02_FB_D(02,08)     OK        18       Up
FPC1_T_SG(0,2,1)_FB_D(05,07) -> SF_10_02_FB_D(02,07)     OK        18       Up
FPC1_T_SG(0,2,2)_FB_D(05,06) -> SF_10_02_FB_D(02,06)     OK        18       Up
FPC1_T_SG(0,2,3)_FB_D(05,05) -> SF_10_02_FB_D(02,05)     OK        18       Up
FPC1_T_SG(0,2,4)_FB_D(05,03) -> SF_10_02_FB_D(02,03)     OK        18       Up
FPC1_T_SG(0,2,5)_FB_D(05,02) -> SF_10_02_FB_D(02,02)     OK        18       Up
FPC1_T_SG(0,2,6)_FB_D(05,01) -> SF_10_02_FB_D(02,01)     OK        18       Up
FPC1_T_SG(0,2,7)_FB_D(05,00) -> SF_10_02_FB_D(02,00)     OK        18       Up
FPC1_B_SG(0,3,0)_FB_D(04,03) -> SF_10_11_FB_D(01,03)     OK        21       Up
FPC1_B_SG(0,3,1)_FB_D(04,02) -> SF_10_11_FB_D(01,02)     OK        21       Up
FPC1_B_SG(0,3,2)_FB_D(04,01) -> SF_10_11_FB_D(01,01)     OK        21       Up
FPC1_B_SG(0,3,3)_FB_D(04,00) -> SF_10_11_FB_D(01,00)     OK        21       Up
FPC1_B_SG(0,3,4)_FB_D(03,11) -> SF_10_11_FB_D(00,11)     OK        21       Up
FPC1_B_SG(0,3,5)_FB_D(03,10) -> SF_10_11_FB_D(00,10)     OK        21       Up
FPC1_B_SG(0,3,6)_FB_D(03,09) -> SF_10_11_FB_D(00,09)     OK        21       Up
FPC1_B_SG(0,3,7)_FB_D(03,08) -> SF_10_11_FB_D(00,08)     OK        21       Up
FPC2_T_SG(1,0,0)_FB_C(10,11) -> SF_10_04_FB_C(07,11)     OK        12       Up
FPC2_T_SG(1,0,1)_FB_C(10,10) -> SF_10_04_FB_C(07,10)     OK        12       Up
FPC2_T_SG(1,0,2)_FB_C(10,09) -> SF_10_04_FB_C(07,09)     OK        12       Up
FPC2_T_SG(1,0,3)_FB_C(10,08) -> SF_10_04_FB_C(07,08)     OK        12       Up
FPC2_T_SG(1,0,4)_FB_C(10,07) -> SF_10_04_FB_C(07,07)     OK        12       Up
FPC2_T_SG(1,0,5)_FB_C(10,06) -> SF_10_04_FB_C(07,06)     OK        12       Up
FPC2_T_SG(1,0,6)_FB_C(10,05) -> SF_10_04_FB_C(07,05)     OK        12       Up
FPC2_T_SG(1,0,7)_FB_C(10,04) -> SF_10_04_FB_C(07,04)     OK        12       Up
FPC2_B_SG(1,1,0)_FB_C(09,07) -> SF_10_14_FB_C(06,07)     OK        15       Up
FPC2_B_SG(1,1,1)_FB_C(09,06) -> SF_10_14_FB_C(06,06)     OK        15       Up
FPC2_B_SG(1,1,2)_FB_C(09,05) -> SF_10_14_FB_C(06,05)     OK        15       Up
FPC2_B_SG(1,1,3)_FB_C(09,04) -> SF_10_14_FB_C(06,04)     OK        15       Up
FPC2_B_SG(1,1,4)_FB_C(09,03) -> SF_10_14_FB_C(06,03)     OK        15       Up
FPC2_B_SG(1,1,5)_FB_C(09,02) -> SF_10_14_FB_C(06,02)     OK        15       Up
FPC2_B_SG(1,1,6)_FB_C(09,01) -> SF_10_14_FB_C(06,01)     OK        15       Up
FPC2_B_SG(1,1,7)_FB_C(09,00) -> SF_10_14_FB_C(06,00)     OK        15       Up
FPC3_T_SG(1,2,0)_FB_C(11,08) -> SF_10_06_FB_C(08,08)     OK        18       Up
FPC3_T_SG(1,2,1)_FB_C(11,07) -> SF_10_06_FB_C(08,07)     OK        18       Up
FPC3_T_SG(1,2,2)_FB_C(11,06) -> SF_10_06_FB_C(08,06)     OK        18       Up
FPC3_T_SG(1,2,3)_FB_C(11,05) -> SF_10_06_FB_C(08,05)     OK        18       Up
```

```
FPC3_T_SG(1,2,4)_FB_C(11,03) -> SF_10_06_FB_C(08,03)     OK       18      Up
FPC3_T_SG(1,2,5)_FB_C(11,02) -> SF_10_06_FB_C(08,02)     OK       18      Up
FPC3_T_SG(1,2,6)_FB_C(11,01) -> SF_10_06_FB_C(08,01)     OK       18      Up
...
```

show chassis fabric topology (T4000 Core Router)

```
user@host> show chassis fabric topology 0
 fchip (mode)
          In-links               State             Out-links              State
--------------------------------------------------------------------------------

SIB0 :
---------

Onboard Links
-------------
SIB0_XF1,14_0->SIB0_XF,00_0      Up      SIB0_XF,00_0->SIB0_XF1,14_0      Up
SIB0_XF,00_0->SIB0_XF1,14_0      Up      SIB0_XF1,14_0->SIB0_XF,00_0      Up
SIB0_XF1,13_0->SIB0_XF,01_0      Up      SIB0_XF,01_0->SIB0_XF1,13_0      Up
SIB0_XF,01_0->SIB0_XF1,13_0      Up      SIB0_XF1,13_0->SIB0_XF,01_0      Up
SIB0_XF1,12_0->SIB0_XF,02_0      Up      SIB0_XF,02_0->SIB0_XF1,12_0      Up
SIB0_XF,02_0->SIB0_XF1,12_0      Up      SIB0_XF1,12_0->SIB0_XF,02_0      Up
SIB0_XF1,11_0->SIB0_XF,03_0      Up      SIB0_XF,03_0->SIB0_XF1,11_0      Up
SIB0_XF,03_0->SIB0_XF1,11_0      Up      SIB0_XF1,11_0->SIB0_XF,03_0      Up
SIB0_XF1,10_0->SIB0_XF,04_0      Up      SIB0_XF,04_0->SIB0_XF1,10_0      Up
SIB0_XF,04_0->SIB0_XF1,10_0      Up      SIB0_XF1,10_0->SIB0_XF,04_0      Up
SIB0_XF1,09_0->SIB0_XF,05_0      Up      SIB0_XF,05_0->SIB0_XF1,09_0      Up
SIB0_XF,05_0->SIB0_XF1,09_0      Up      SIB0_XF1,09_0->SIB0_XF,05_0      Up
SIB0_XF2,14_0->SIB0_XF,06_0      Up      SIB0_XF,06_0->SIB0_XF2,14_0      Up
SIB0_XF,06_0->SIB0_XF2,14_0      Up      SIB0_XF2,14_0->SIB0_XF,06_0      Up
SIB0_XF2,13_0->SIB0_XF,07_0      Up      SIB0_XF,07_0->SIB0_XF2,13_0      Up
SIB0_XF,07_0->SIB0_XF2,13_0      Up      SIB0_XF2,13_0->SIB0_XF,07_0      Up
SIB0_XF2,12_0->SIB0_XF,08_0      Up      SIB0_XF,08_0->SIB0_XF2,12_0      Up
SIB0_XF,08_0->SIB0_XF2,12_0      Up      SIB0_XF2,12_0->SIB0_XF,08_0      Up
SIB0_XF2,11_0->SIB0_XF,09_0      Up      SIB0_XF,09_0->SIB0_XF2,11_0      Up
SIB0_XF,09_0->SIB0_XF2,11_0      Up      SIB0_XF2,11_0->SIB0_XF,09_0      Up
SIB0_XF2,10_0->SIB0_XF,10_0      Up      SIB0_XF,10_0->SIB0_XF2,10_0      Up
SIB0_XF,10_0->SIB0_XF2,10_0      Up      SIB0_XF2,10_0->SIB0_XF,10_0      Up
SIB0_XF2,09_0->SIB0_XF,11_0      Up      SIB0_XF,11_0->SIB0_XF2,09_0      Up
SIB0_XF,11_0->SIB0_XF2,09_0      Up      SIB0_XF2,09_0->SIB0_XF,11_0      Up
```

```
SIB0_XF3,13_0->SIB0_XF,12_0      Up      SIB0_XF,12_0->SIB0_XF3,13_0      Up
SIB0_XF,12_0->SIB0_XF3,13_0      Up      SIB0_XF3,13_0->SIB0_XF,12_0      Up
SIB0_XF3,12_0->SIB0_XF,13_0      Up      SIB0_XF,13_0->SIB0_XF3,12_0      Up
SIB0_XF,13_0->SIB0_XF3,12_0      Up      SIB0_XF3,12_0->SIB0_XF,13_0      Up
SIB0_XF3,11_0->SIB0_XF,14_0      Up      SIB0_XF,14_0->SIB0_XF3,11_0      Up
SIB0_XF,14_0->SIB0_XF3,11_0      Up      SIB0_XF3,11_0->SIB0_XF,14_0      Up
SIB0_XF3,10_0->SIB0_XF,15_0      Up      SIB0_XF,15_0->SIB0_XF3,10_0      Up
SIB0_XF,15_0->SIB0_XF3,10_0      Up      SIB0_XF3,10_0->SIB0_XF,15_0      Up

PFE Links
---------------
FPC2PFE0->SIB0_XF1,05_0          Up      SIB0_XF1,05_0->FPC2PFE0          Up
FPC3PFE0->SIB0_XF2,15_0          Up      SIB0_XF2,15_0->FPC3PFE0          Up
FPC5PFE0->SIB0_XF2,05_0          Up      SIB0_XF2,05_0->FPC5PFE0          Up
FPC5PFE1->SIB0_XF2,07_0          Up      SIB0_XF2,07_0->FPC5PFE1          Up
FPC6PFE0->SIB0_XF3,01_0          Up      SIB0_XF3,01_0->FPC6PFE0          Up
FPC6PFE0->SIB0_XF3,01_1          Up      SIB0_XF3,01_1->FPC6PFE0          Up
FPC6PFE0->SIB0_XF3,02_0          Up      SIB0_XF3,02_0->FPC6PFE0          Up
FPC6PFE1->SIB0_XF3,03_0          Up      SIB0_XF3,03_0->FPC6PFE1          Up
FPC6PFE1->SIB0_XF3,03_1          Up      SIB0_XF3,03_1->FPC6PFE1          Up
FPC6PFE1->SIB0_XF3,02_1          Up      SIB0_XF3,02_1->FPC6PFE1          Up
```

show chassis fabric topology lcc (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric topology lcc
6
lcc6-re0:
--------------------------------------------------------------------------
  fchip (mode)
          In-links               State             Out-links              State
-------------------------------------------------------------------------------
SIB0 :
---------
CXP0_Evn->LCC_SIB0_XF3,10_0      Up      LCC_SIB0_XF3,10_0->CXP0_Evn      Up
CXP0_Odd->LCC_SIB0_XF3,11_0      Up      LCC_SIB0_XF3,11_0->CXP0_Odd      Up
CXP1_Evn->LCC_SIB0_XF3,12_0      Up      LCC_SIB0_XF3,12_0->CXP1_Evn      Up
CXP1_Odd->LCC_SIB0_XF3,13_0      Up      LCC_SIB0_XF3,13_0->CXP1_Odd      Up
CXP2_Evn->LCC_SIB0_XF2,09_0      Up      LCC_SIB0_XF2,09_0->CXP2_Evn      Up
CXP2_Odd->LCC_SIB0_XF2,10_0      Up      LCC_SIB0_XF2,10_0->CXP2_Odd      Up
CXP3_Evn->LCC_SIB0_XF2,11_0      Up      LCC_SIB0_XF2,11_0->CXP3_Evn      Up
CXP3_Odd->LCC_SIB0_XF2,12_0      Up      LCC_SIB0_XF2,12_0->CXP3_Odd      Up
```

```
CXP4_Evn->LCC_SIB0_XF2,13_0      Up      LCC_SIB0_XF2,13_0->CXP4_Evn      Up
CXP4_Odd->LCC_SIB0_XF1,09_0      Up      LCC_SIB0_XF1,09_0->CXP4_Odd      Up
CXP5_Evn->LCC_SIB0_XF2,14_0      Up      LCC_SIB0_XF2,14_0->CXP5_Evn      Up
CXP5_Odd->LCC_SIB0_XF1,10_0      Up      LCC_SIB0_XF1,10_0->CXP5_Odd      Up
CXP6_Evn->LCC_SIB0_XF1,11_0      Up      LCC_SIB0_XF1,11_0->CXP6_Evn      Up
CXP6_Odd->LCC_SIB0_XF1,12_0      Up      LCC_SIB0_XF1,12_0->CXP6_Odd      Up
CXP7_Evn->LCC_SIB0_XF1,13_0      Up      LCC_SIB0_XF1,13_0->CXP7_Evn      Up
CXP7_Odd->LCC_SIB0_XF1,14_0      Up      LCC_SIB0_XF1,14_0->CXP7_Odd      Up
SIB1 :
---------
SIB2 :
---------
CXP0_Evn->LCC_SIB2_XF3,10_0      Up      LCC_SIB2_XF3,10_0->CXP0_Evn      Up
CXP0_Odd->LCC_SIB2_XF3,11_0      Up      LCC_SIB2_XF3,11_0->CXP0_Odd      Up
CXP1_Evn->LCC_SIB2_XF3,12_0      Up      LCC_SIB2_XF3,12_0->CXP1_Evn      Up
CXP1_Odd->LCC_SIB2_XF3,13_0      Up      LCC_SIB2_XF3,13_0->CXP1_Odd      Up
CXP2_Evn->LCC_SIB2_XF2,09_0      Up      LCC_SIB2_XF2,09_0->CXP2_Evn      Up
CXP2_Odd->LCC_SIB2_XF2,10_0      Up      LCC_SIB2_XF2,10_0->CXP2_Odd      Up
CXP3_Evn->LCC_SIB2_XF2,11_0      Up      LCC_SIB2_XF2,11_0->CXP3_Evn      Up
CXP3_Odd->LCC_SIB2_XF2,12_0      Up      LCC_SIB2_XF2,12_0->CXP3_Odd      Up
CXP4_Evn->LCC_SIB2_XF2,13_0      Up      LCC_SIB2_XF2,13_0->CXP4_Evn      Up
CXP4_Odd->LCC_SIB2_XF1,09_0      Up      LCC_SIB2_XF1,09_0->CXP4_Odd      Up
CXP5_Evn->LCC_SIB2_XF2,14_0      Up      LCC_SIB2_XF2,14_0->CXP5_Evn      Up
CXP5_Odd->LCC_SIB2_XF1,10_0      Up      LCC_SIB2_XF1,10_0->CXP5_Odd      Up
CXP6_Evn->LCC_SIB2_XF1,11_0      Up      LCC_SIB2_XF1,11_0->CXP6_Evn      Up
CXP6_Odd->LCC_SIB2_XF1,12_0      Up      LCC_SIB2_XF1,12_0->CXP6_Odd      Up
CXP7_Evn->LCC_SIB2_XF1,13_0      Up      LCC_SIB2_XF1,13_0->CXP7_Evn      Up
CXP7_Odd->LCC_SIB2_XF1,14_0      Up      LCC_SIB2_XF1,14_0->CXP7_Odd      Up
SIB3 :
---------
CXP0_Evn->LCC_SIB3_XF3,10_0      Up      LCC_SIB3_XF3,10_0->CXP0_Evn      Up
CXP0_Odd->LCC_SIB3_XF3,11_0      Up      LCC_SIB3_XF3,11_0->CXP0_Odd      Up
CXP1_Evn->LCC_SIB3_XF3,12_0      Up      LCC_SIB3_XF3,12_0->CXP1_Evn      Up
CXP1_Odd->LCC_SIB3_XF3,13_0      Up      LCC_SIB3_XF3,13_0->CXP1_Odd      Up
CXP2_Evn->LCC_SIB3_XF2,09_0      Up      LCC_SIB3_XF2,09_0->CXP2_Evn      Up
CXP2_Odd->LCC_SIB3_XF2,10_0      Up      LCC_SIB3_XF2,10_0->CXP2_Odd      Up
CXP3_Evn->LCC_SIB3_XF2,11_0      Up      LCC_SIB3_XF2,11_0->CXP3_Evn      Up
CXP3_Odd->LCC_SIB3_XF2,12_0      Up      LCC_SIB3_XF2,12_0->CXP3_Odd      Up
CXP4_Evn->LCC_SIB3_XF2,13_0      Up      LCC_SIB3_XF2,13_0->CXP4_Evn      Up
CXP4_Odd->LCC_SIB3_XF1,09_0      Up      LCC_SIB3_XF1,09_0->CXP4_Odd      Up
CXP5_Evn->LCC_SIB3_XF2,14_0      Up      LCC_SIB3_XF2,14_0->CXP5_Evn      Up
CXP5_Odd->LCC_SIB3_XF1,10_0      Up      LCC_SIB3_XF1,10_0->CXP5_Odd      Up
CXP6_Evn->LCC_SIB3_XF1,11_0      Up      LCC_SIB3_XF1,11_0->CXP6_Evn      Up
```

```
CXP6_Odd->LCC_SIB3_XF1,12_0      Up      LCC_SIB3_XF1,12_0->CXP6_Odd      Up
CXP7_Evn->LCC_SIB3_XF1,13_0      Up      LCC_SIB3_XF1,13_0->CXP7_Evn      Up
CXP7_Odd->LCC_SIB3_XF1,14_0      Up      LCC_SIB3_XF1,14_0->CXP7_Odd      Up
SIB4 :
---------
CXP0_Evn->LCC_SIB4_XF3,10_0      Up      LCC_SIB4_XF3,10_0->CXP0_Evn      Up
CXP0_Odd->LCC_SIB4_XF3,11_0      Up      LCC_SIB4_XF3,11_0->CXP0_Odd      Up
CXP1_Evn->LCC_SIB4_XF3,12_0      Up      LCC_SIB4_XF3,12_0->CXP1_Evn      Up
CXP1_Odd->LCC_SIB4_XF3,13_0      Up      LCC_SIB4_XF3,13_0->CXP1_Odd      Up
CXP2_Evn->LCC_SIB4_XF2,09_0      Up      LCC_SIB4_XF2,09_0->CXP2_Evn      Up
CXP2_Odd->LCC_SIB4_XF2,10_0      Up      LCC_SIB4_XF2,10_0->CXP2_Odd      Up
CXP3_Evn->LCC_SIB4_XF2,11_0      Up      LCC_SIB4_XF2,11_0->CXP3_Evn      Up
CXP3_Odd->LCC_SIB4_XF2,12_0      Up      LCC_SIB4_XF2,12_0->CXP3_Odd      Up
CXP4_Evn->LCC_SIB4_XF2,13_0      Up      LCC_SIB4_XF2,13_0->CXP4_Evn      Up
CXP4_Odd->LCC_SIB4_XF1,09_0      Up      LCC_SIB4_XF1,09_0->CXP4_Odd      Up
CXP5_Evn->LCC_SIB4_XF2,14_0      Up      LCC_SIB4_XF2,14_0->CXP5_Evn      Up
CXP5_Odd->LCC_SIB4_XF1,10_0      Up      LCC_SIB4_XF1,10_0->CXP5_Odd      Up
CXP6_Evn->LCC_SIB4_XF1,11_0      Up      LCC_SIB4_XF1,11_0->CXP6_Evn      Up
CXP6_Odd->LCC_SIB4_XF1,12_0      Up      LCC_SIB4_XF1,12_0->CXP6_Odd      Up
CXP7_Evn->LCC_SIB4_XF1,13_0      Up      LCC_SIB4_XF1,13_0->CXP7_Evn      Up
CXP7_Odd->LCC_SIB4_XF1,14_0      Up      LCC_SIB4_XF1,14_0->CXP7_Odd      Up
```

show chassis fabric topology (PTX5000 Router)

```
user@host> show chassis fabric topology
In-link  : FPC# FE# TQ# (TQ-TX sub-chnl #) ->
            SIB# TF#_FCORE# (TF-RX port#, TF-RX sub-chn#, TF-RX inst#)

 Out-link : SIB# TF#_FCORE# (TF-TX port#, TF-TX sub-chn#, TF-TX inst#) ->
            FPC# FE# TQ# (TQ-RX sub-chnl #)
(6, 4, 06) in FPC02FE0TQ0(02)->S01F0_0(6,4,06) will be TF Rx Port 6, TF CCL Rx Sub-Channel 4, TF
CCL Rx Instance 6.
(2, 7, 10) in S01F0_0(2,7,10)->FPC02FE0TQ0(02) will be TF-Tx Port 2, TF CCL Tx Sub-channel 7, TF
CCL Tx Instance 10.
SIB 0 FCHIP 0 FCORE 0 :
-----------------------
          In-links               State             Out-links              State
--------------------------------------------------------------------------------
FPC00FE0TQ0(00)->S00F0_0(7,4,07) OK      S00F0_0(3,7,11)->FPC00FE0TQ0(00) OK
FPC00FE1TQ1(00)->S00F0_0(7,6,07) OK      S00F0_0(3,5,11)->FPC00FE1TQ1(00) OK
FPC00FE2TQ2(00)->S00F0_0(7,5,07) OK      S00F0_0(3,6,11)->FPC00FE2TQ2(00) OK
```

```
FPC00FE3TQ3(00)->S00F0_0(7,7,07) OK      S00F0_0(3,4,11)->FPC00FE3TQ3(00) OK
FPC01FE0TQ0(00)->S00F0_0(7,0,07) OK      S00F0_0(3,3,11)->FPC01FE0TQ0(00) OK
FPC01FE1TQ1(00)->S00F0_0(7,1,07) OK      S00F0_0(3,1,11)->FPC01FE1TQ1(00) OK
FPC01FE2TQ2(00)->S00F0_0(7,2,07) OK      S00F0_0(3,2,11)->FPC01FE2TQ2(00) Error
FPC01FE3TQ3(00)->S00F0_0(7,3,07) OK      S00F0_0(3,0,11)->FPC01FE3TQ3(00) OK
FPC02FE0TQ0(00)->S00F0_0(6,4,06) OK      S00F0_0(2,7,10)->FPC02FE0TQ0(00) OK
FPC02FE1TQ1(00)->S00F0_0(6,5,06) OK      S00F0_0(2,5,10)->FPC02FE1TQ1(00) OK
FPC02FE2TQ2(00)->S00F0_0(6,6,06) OK      S00F0_0(2,6,10)->FPC02FE2TQ2(00) OK
FPC02FE3TQ3(00)->S00F0_0(6,7,06) OK      S00F0_0(2,4,10)->FPC02FE3TQ3(00) OK
FPC03FE0TQ0(00)->S00F0_0(6,0,06) Down    S00F0_0(2,3,10)->FPC03FE0TQ0(00) Down
FPC03FE1TQ1(00)->S00F0_0(6,1,06) Down    S00F0_0(2,0,10)->FPC03FE1TQ1(00) Down
FPC03FE2TQ2(00)->S00F0_0(6,2,06) Down    S00F0_0(2,2,10)->FPC03FE2TQ2(00) Down
FPC03FE3TQ3(00)->S00F0_0(6,3,06) Down    S00F0_0(2,1,10)->FPC03FE3TQ3(00) Down
FPC04FE0TQ0(00)->S00F0_0(5,4,05) OK      S00F0_0(1,7,09)->FPC04FE0TQ0(00) OK
FPC04FE1TQ1(00)->S00F0_0(5,5,05) OK      S00F0_0(1,6,09)->FPC04FE1TQ1(00) OK
FPC04FE2TQ2(00)->S00F0_0(5,6,05) OK      S00F0_0(1,4,09)->FPC04FE2TQ2(00) OK
FPC04FE3TQ3(00)->S00F0_0(5,7,05) OK      S00F0_0(1,5,09)->FPC04FE3TQ3(00) OK
FPC05FE0TQ0(00)->S00F0_0(5,0,05) OK      S00F0_0(1,3,09)->FPC05FE0TQ0(00) OK
FPC05FE1TQ1(00)->S00F0_0(5,1,05) OK      S00F0_0(1,0,09)->FPC05FE1TQ1(00) OK
FPC05FE2TQ2(00)->S00F0_0(5,2,05) OK      S00F0_0(1,2,09)->FPC05FE2TQ2(00) OK
FPC05FE3TQ3(00)->S00F0_0(5,3,05) OK      S00F0_0(1,1,09)->FPC05FE3TQ3(00) OK
FPC06FE0TQ0(00)->S00F0_0(4,4,04) Down    S00F0_0(0,7,08)->FPC06FE0TQ0(00) Down
FPC06FE1TQ1(00)->S00F0_0(4,5,04) Down    S00F0_0(0,5,08)->FPC06FE1TQ1(00) Down
FPC06FE2TQ2(00)->S00F0_0(4,6,04) Down    S00F0_0(0,6,08)->FPC06FE2TQ2(00) Down
FPC06FE3TQ3(00)->S00F0_0(4,7,04) Down    S00F0_0(0,4,08)->FPC06FE3TQ3(00) Down
FPC07FE0TQ0(00)->S00F0_0(4,2,04) Down    S00F0_0(0,3,08)->FPC07FE0TQ0(00) Down
FPC07FE1TQ1(00)->S00F0_0(4,0,04) Down    S00F0_0(0,0,08)->FPC07FE1TQ1(00) Down
FPC07FE2TQ2(00)->S00F0_0(4,1,04) Down    S00F0_0(0,1,08)->FPC07FE2TQ2(00) Down
FPC07FE3TQ3(00)->S00F0_0(4,3,04) Down    S00F0_0(0,2,08)->FPC07FE3TQ3(00) Down

SIB 0 FCHIP 0 FCORE 1 :
-----------------------
          In-links               State            Out-links              State
--------------------------------------------------------------------------------
FPC00FE0TQ0(01)->S00F0_1(3,4,11) OK      S00F0_1(7,6,07)->FPC00FE0TQ0(01) OK
FPC00FE1TQ1(01)->S00F0_1(3,5,11) OK      S00F0_1(7,4,07)->FPC00FE1TQ1(01) OK
FPC00FE2TQ2(01)->S00F0_1(3,6,11) OK      S00F0_1(7,7,07)->FPC00FE2TQ2(01) OK
FPC00FE3TQ3(01)->S00F0_1(3,7,11) OK      S00F0_1(7,5,07)->FPC00FE3TQ3(01) OK
FPC01FE0TQ0(01)->S00F0_1(3,0,11) OK      S00F0_1(7,2,07)->FPC01FE0TQ0(01) OK
FPC01FE1TQ1(01)->S00F0_1(3,1,11) OK      S00F0_1(7,0,07)->FPC01FE1TQ1(01) OK
FPC01FE2TQ2(01)->S00F0_1(3,2,11) OK      S00F0_1(7,3,07)->FPC01FE2TQ2(01) OK
FPC01FE3TQ3(01)->S00F0_1(3,3,11) OK      S00F0_1(7,1,07)->FPC01FE3TQ3(01) OK
FPC02FE0TQ0(01)->S00F0_1(2,4,10) OK      S00F0_1(6,5,06)->FPC02FE0TQ0(01) OK
```

```
FPC02FE1TQ1(01)->S00F0_1(2,5,10) OK      S00F0_1(6,4,06)->FPC02FE1TQ1(01) OK
FPC02FE2TQ2(01)->S00F0_1(2,6,10) OK      S00F0_1(6,7,06)->FPC02FE2TQ2(01) OK
FPC02FE3TQ3(01)->S00F0_1(2,7,10) OK      S00F0_1(6,6,06)->FPC02FE3TQ3(01) OK
FPC03FE0TQ0(01)->S00F0_1(2,0,10) Down    S00F0_1(6,1,06)->FPC03FE0TQ0(01) Down
FPC03FE1TQ1(01)->S00F0_1(2,1,10) Down    S00F0_1(6,0,06)->FPC03FE1TQ1(01) Down
FPC03FE2TQ2(01)->S00F0_1(2,2,10) Down    S00F0_1(6,3,06)->FPC03FE2TQ2(01) Down
FPC03FE3TQ3(01)->S00F0_1(2,3,10) Down    S00F0_1(6,2,06)->FPC03FE3TQ3(01) Down
FPC04FE0TQ0(01)->S00F0_1(1,4,09) OK      S00F0_1(5,5,05)->FPC04FE0TQ0(01) OK
FPC04FE1TQ1(01)->S00F0_1(1,5,09) OK      S00F0_1(5,4,05)->FPC04FE1TQ1(01) OK
FPC04FE2TQ2(01)->S00F0_1(1,6,09) OK      S00F0_1(5,7,05)->FPC04FE2TQ2(01) OK
FPC04FE3TQ3(01)->S00F0_1(1,7,09) OK      S00F0_1(5,6,05)->FPC04FE3TQ3(01) OK
FPC05FE0TQ0(01)->S00F0_1(1,0,09) OK      S00F0_1(5,1,05)->FPC05FE0TQ0(01) OK
FPC05FE1TQ1(01)->S00F0_1(1,1,09) OK      S00F0_1(5,0,05)->FPC05FE1TQ1(01) OK
FPC05FE2TQ2(01)->S00F0_1(1,2,09) OK      S00F0_1(5,3,05)->FPC05FE2TQ2(01) OK
FPC05FE3TQ3(01)->S00F0_1(1,3,09) OK      S00F0_1(5,2,05)->FPC05FE3TQ3(01) OK
FPC06FE0TQ0(01)->S00F0_1(0,4,08) Down    S00F0_1(4,7,04)->FPC06FE0TQ0(01) Down
FPC06FE1TQ1(01)->S00F0_1(0,5,08) Down    S00F0_1(4,0,04)->FPC06FE1TQ1(01) Down
FPC06FE2TQ2(01)->S00F0_1(0,6,08) Down    S00F0_1(4,6,04)->FPC06FE2TQ2(01) Down
FPC06FE3TQ3(01)->S00F0_1(0,7,08) Down    S00F0_1(4,1,04)->FPC06FE3TQ3(01) Down
FPC07FE0TQ0(01)->S00F0_1(0,0,08) Down    S00F0_1(4,3,04)->FPC07FE0TQ0(01) Down
FPC07FE1TQ1(01)->S00F0_1(0,1,08) Down    S00F0_1(4,4,04)->FPC07FE1TQ1(01) Down
FPC07FE2TQ2(01)->S00F0_1(0,2,08) Down    S00F0_1(4,2,04)->FPC07FE2TQ2(01) Down
FPC07FE3TQ3(01)->S00F0_1(0,3,08) Down    S00F0_1(4,5,04)->FPC07FE3TQ3(01) Down

SIB 1 FCHIP 0 FCORE 0 :
-----------------------
          In-links               State             Out-links              State
--------------------------------------------------------------------------------
FPC00FE0TQ0(02)->S01F0_0(7,4,07) Error   S01F0_0(3,7,11)->FPC00FE0TQ0(02) Down
FPC00FE1TQ1(02)->S01F0_0(7,6,07) OK      S01F0_0(3,5,11)->FPC00FE1TQ1(02) OK
FPC00FE2TQ2(02)->S01F0_0(7,5,07) OK      S01F0_0(3,6,11)->FPC00FE2TQ2(02) OK
FPC00FE3TQ3(02)->S01F0_0(7,7,07) OK      S01F0_0(3,4,11)->FPC00FE3TQ3(02) OK
FPC01FE0TQ0(02)->S01F0_0(7,0,07) OK      S01F0_0(3,3,11)->FPC01FE0TQ0(02) OK
FPC01FE1TQ1(02)->S01F0_0(7,1,07) OK      S01F0_0(3,1,11)->FPC01FE1TQ1(02) OK
FPC01FE2TQ2(02)->S01F0_0(7,2,07) OK      S01F0_0(3,2,11)->FPC01FE2TQ2(02) OK
FPC01FE3TQ3(02)->S01F0_0(7,3,07) OK      S01F0_0(3,0,11)->FPC01FE3TQ3(02) OK
FPC02FE0TQ0(02)->S01F0_0(6,4,06) OK      S01F0_0(2,7,10)->FPC02FE0TQ0(02) OK
FPC02FE1TQ1(02)->S01F0_0(6,5,06) OK      S01F0_0(2,5,10)->FPC02FE1TQ1(02) OK
FPC02FE2TQ2(02)->S01F0_0(6,6,06) OK      S01F0_0(2,6,10)->FPC02FE2TQ2(02) OK
FPC02FE3TQ3(02)->S01F0_0(6,7,06) OK      S01F0_0(2,4,10)->FPC02FE3TQ3(02) OK
FPC03FE0TQ0(02)->S01F0_0(6,0,06) Down    S01F0_0(2,3,10)->FPC03FE0TQ0(02) Down
FPC03FE1TQ1(02)->S01F0_0(6,1,06) Down    S01F0_0(2,0,10)->FPC03FE1TQ1(02) Down
FPC03FE2TQ2(02)->S01F0_0(6,2,06) Down    S01F0_0(2,2,10)->FPC03FE2TQ2(02) Down
```

```
FPC03FE3TQ3(02)->S01F0_0(6,3,06) Down    S01F0_0(2,1,10)->FPC03FE3TQ3(02) Down
FPC04FE0TQ0(02)->S01F0_0(5,4,05) OK      S01F0_0(1,7,09)->FPC04FE0TQ0(02) OK
FPC04FE1TQ1(02)->S01F0_0(5,5,05) OK      S01F0_0(1,6,09)->FPC04FE1TQ1(02) OK
FPC04FE2TQ2(02)->S01F0_0(5,6,05) OK      S01F0_0(1,4,09)->FPC04FE2TQ2(02) OK
FPC04FE3TQ3(02)->S01F0_0(5,7,05) OK      S01F0_0(1,5,09)->FPC04FE3TQ3(02) OK
FPC05FE0TQ0(02)->S01F0_0(5,0,05) OK      S01F0_0(1,3,09)->FPC05FE0TQ0(02) OK
FPC05FE1TQ1(02)->S01F0_0(5,1,05) OK      S01F0_0(1,0,09)->FPC05FE1TQ1(02) OK
FPC05FE2TQ2(02)->S01F0_0(5,2,05) OK      S01F0_0(1,2,09)->FPC05FE2TQ2(02) OK
FPC05FE3TQ3(02)->S01F0_0(5,3,05) OK      S01F0_0(1,1,09)->FPC05FE3TQ3(02) OK
FPC06FE0TQ0(02)->S01F0_0(4,4,04) Down    S01F0_0(0,7,08)->FPC06FE0TQ0(02) Down
FPC06FE1TQ1(02)->S01F0_0(4,5,04) Down    S01F0_0(0,5,08)->FPC06FE1TQ1(02) Down
FPC06FE2TQ2(02)->S01F0_0(4,6,04) Down    S01F0_0(0,6,08)->FPC06FE2TQ2(02) Down
FPC06FE3TQ3(02)->S01F0_0(4,7,04) Down    S01F0_0(0,4,08)->FPC06FE3TQ3(02) Down
FPC07FE0TQ0(02)->S01F0_0(4,2,04) Down    S01F0_0(0,3,08)->FPC07FE0TQ0(02) Down
FPC07FE1TQ1(02)->S01F0_0(4,0,04) Down    S01F0_0(0,0,08)->FPC07FE1TQ1(02) Down
FPC07FE2TQ2(02)->S01F0_0(4,1,04) Down    S01F0_0(0,1,08)->FPC07FE2TQ2(02) Down
FPC07FE3TQ3(02)->S01F0_0(4,3,04) Down    S01F0_0(0,2,08)->FPC07FE3TQ3(02) Down

SIB 1 FCHIP 0 FCORE 1 :
-----------------------
          In-links               State             Out-links              State
--------------------------------------------------------------------------------
FPC00FE0TQ0(03)->S01F0_1(3,4,11) OK      S01F0_1(7,6,07)->FPC00FE0TQ0(03) OK
FPC00FE1TQ1(03)->S01F0_1(3,5,11) OK      S01F0_1(7,4,07)->FPC00FE1TQ1(03) OK
FPC00FE2TQ2(03)->S01F0_1(3,6,11) OK      S01F0_1(7,7,07)->FPC00FE2TQ2(03) OK
FPC00FE3TQ3(03)->S01F0_1(3,7,11) OK      S01F0_1(7,5,07)->FPC00FE3TQ3(03) OK
FPC01FE0TQ0(03)->S01F0_1(3,0,11) OK      S01F0_1(7,2,07)->FPC01FE0TQ0(03) OK
FPC01FE1TQ1(03)->S01F0_1(3,1,11) OK      S01F0_1(7,0,07)->FPC01FE1TQ1(03) OK
FPC01FE2TQ2(03)->S01F0_1(3,2,11) OK      S01F0_1(7,3,07)->FPC01FE2TQ2(03) OK
FPC01FE3TQ3(03)->S01F0_1(3,3,11) OK      S01F0_1(7,1,07)->FPC01FE3TQ3(03) OK
FPC02FE0TQ0(03)->S01F0_1(2,4,10) OK      S01F0_1(6,5,06)->FPC02FE0TQ0(03) OK
FPC02FE1TQ1(03)->S01F0_1(2,5,10) OK      S01F0_1(6,4,06)->FPC02FE1TQ1(03) OK
FPC02FE2TQ2(03)->S01F0_1(2,6,10) OK      S01F0_1(6,7,06)->FPC02FE2TQ2(03) OK
FPC02FE3TQ3(03)->S01F0_1(2,7,10) OK      S01F0_1(6,6,06)->FPC02FE3TQ3(03) OK
FPC03FE0TQ0(03)->S01F0_1(2,0,10) Down    S01F0_1(6,1,06)->FPC03FE0TQ0(03) Down
FPC03FE1TQ1(03)->S01F0_1(2,1,10) Down    S01F0_1(6,0,06)->FPC03FE1TQ1(03) Down
FPC03FE2TQ2(03)->S01F0_1(2,2,10) Down    S01F0_1(6,3,06)->FPC03FE2TQ2(03) Down
FPC03FE3TQ3(03)->S01F0_1(2,3,10) Down    S01F0_1(6,2,06)->FPC03FE3TQ3(03) Down
FPC04FE0TQ0(03)->S01F0_1(1,4,09) OK      S01F0_1(5,5,05)->FPC04FE0TQ0(03) OK
FPC04FE1TQ1(03)->S01F0_1(1,5,09) OK      S01F0_1(5,4,05)->FPC04FE1TQ1(03) OK
FPC04FE2TQ2(03)->S01F0_1(1,6,09) OK      S01F0_1(5,7,05)->FPC04FE2TQ2(03) OK
FPC04FE3TQ3(03)->S01F0_1(1,7,09) OK      S01F0_1(5,6,05)->FPC04FE3TQ3(03) OK
FPC05FE0TQ0(03)->S01F0_1(1,0,09) OK      S01F0_1(5,1,05)->FPC05FE0TQ0(03) OK
```

```
FPC05FE1TQ1(03)->S01F0_1(1,1,09) OK      S01F0_1(5,0,05)->FPC05FE1TQ1(03) OK
FPC05FE2TQ2(03)->S01F0_1(1,2,09) OK      S01F0_1(5,3,05)->FPC05FE2TQ2(03) OK
FPC05FE3TQ3(03)->S01F0_1(1,3,09) OK      S01F0_1(5,2,05)->FPC05FE3TQ3(03) OK
FPC06FE0TQ0(03)->S01F0_1(0,4,08) Down    S01F0_1(4,7,04)->FPC06FE0TQ0(03) Down
FPC06FE1TQ1(03)->S01F0_1(0,5,08) Down    S01F0_1(4,0,04)->FPC06FE1TQ1(03) Down
FPC06FE2TQ2(03)->S01F0_1(0,6,08) Down    S01F0_1(4,6,04)->FPC06FE2TQ2(03) Down
FPC06FE3TQ3(03)->S01F0_1(0,7,08) Down    S01F0_1(4,1,04)->FPC06FE3TQ3(03) Down
FPC07FE0TQ0(03)->S01F0_1(0,0,08) Down    S01F0_1(4,3,04)->FPC07FE0TQ0(03) Down
FPC07FE1TQ1(03)->S01F0_1(0,1,08) Down    S01F0_1(4,4,04)->FPC07FE1TQ1(03) Down
FPC07FE2TQ2(03)->S01F0_1(0,2,08) Down    S01F0_1(4,2,04)->FPC07FE2TQ2(03) Down
FPC07FE3TQ3(03)->S01F0_1(0,3,08) Down    S01F0_1(4,5,04)->FPC07FE3TQ3(03) Down
```

show chassis fabric topology (PTX10002-36QDD Routers)

```
user@host> show chassis fabric topology
In-link  : FPC# FE# (TX inst#, TX sub-chnl #) ->
           FPC# FE# (RX port#, RX sub-chn#, RX inst#)

Out-link : FPC# FE# (TX port#, TX sub-chn#, TX inst#) ->
           FPC# FE# (RX inst#, RX sub-chnl #)
FPC 0 FPFE 0 :
-----------------------
          In-links               State            Out-links              State
--------------------------------------------------------------------------------
FPC00FE2(8,08)->FPC00FE0(8,08)   UP      FPC00FE0(8,08)->FPC00FE2(8,08)   UP
FPC00FE3(8,00)->FPC00FE0(8,00)   UP      FPC00FE0(8,00)->FPC00FE3(8,00)   UP
FPC00FE3(8,01)->FPC00FE0(8,01)   UP      FPC00FE0(8,01)->FPC00FE3(8,01)   UP
FPC00FE3(8,02)->FPC00FE0(8,02)   UP      FPC00FE0(8,02)->FPC00FE3(8,02)   UP
FPC00FE3(8,03)->FPC00FE0(8,03)   UP      FPC00FE0(8,03)->FPC00FE3(8,03)   UP
FPC00FE3(8,04)->FPC00FE0(8,04)   UP      FPC00FE0(8,04)->FPC00FE3(8,04)   UP
FPC00FE3(8,05)->FPC00FE0(8,05)   UP      FPC00FE0(8,05)->FPC00FE3(8,05)   UP
FPC00FE3(8,06)->FPC00FE0(8,06)   UP      FPC00FE0(8,06)->FPC00FE3(8,06)   UP
FPC00FE3(8,07)->FPC00FE0(8,07)   UP      FPC00FE0(8,07)->FPC00FE3(8,07)   UP
…
FPC 0 FPFE 1 :
-----------------------
          In-links               State            Out-links              State
--------------------------------------------------------------------------------
FPC00FE2(0,00)->FPC00FE1(0,00)   UP      FPC00FE1(0,00)->FPC00FE2(0,00)   UP
FPC00FE2(0,01)->FPC00FE1(0,01)   UP      FPC00FE1(0,01)->FPC00FE2(0,01)   UP
FPC00FE2(0,02)->FPC00FE1(0,02)   UP      FPC00FE1(0,02)->FPC00FE2(0,02)   UP
```

```
FPC00FE2(0,03)->FPC00FE1(0,03)   UP      FPC00FE1(0,03)->FPC00FE2(0,03)   UP
FPC00FE2(0,04)->FPC00FE1(0,04)   UP      FPC00FE1(0,04)->FPC00FE2(0,04)   UP
FPC00FE2(0,05)->FPC00FE1(0,05)   UP      FPC00FE1(0,05)->FPC00FE2(0,05)   UP
FPC00FE2(0,06)->FPC00FE1(0,06)   UP      FPC00FE1(0,06)->FPC00FE2(0,06)   UP
…
FPC 0 FPFE 2 :
-----------------------
          In-links               State             Out-links              State
--------------------------------------------------------------------------------
FPC00FE0(8,08)->FPC00FE2(8,08)   UP      FPC00FE2(8,08)->FPC00FE0(8,08)   UP
FPC00FE1(8,00)->FPC00FE2(8,00)   UP      FPC00FE2(8,00)->FPC00FE1(8,00)   UP
FPC00FE1(8,01)->FPC00FE2(8,01)   UP      FPC00FE2(8,01)->FPC00FE1(8,01)   UP
FPC00FE1(8,02)->FPC00FE2(8,02)   UP      FPC00FE2(8,02)->FPC00FE1(8,02)   UP
FPC00FE1(8,03)->FPC00FE2(8,03)   UP      FPC00FE2(8,03)->FPC00FE1(8,03)   UP
FPC00FE1(8,04)->FPC00FE2(8,04)   UP      FPC00FE2(8,04)->FPC00FE1(8,04)   UP
FPC00FE1(8,05)->FPC00FE2(8,05)   UP      FPC00FE2(8,05)->FPC00FE1(8,05)   UP
FPC00FE1(8,06)->FPC00FE2(8,06)   UP      FPC00FE2(8,06)->FPC00FE1(8,06)   UP
FPC00FE1(8,07)->FPC00FE2(8,07)   UP      FPC00FE2(8,07)->FPC00FE1(8,07)   UP
FPC00FE1(7,01)->FPC00FE2(7,01)   UP      FPC00FE2(7,01)->FPC00FE1(7,01)   UP
FPC00FE1(7,02)->FPC00FE2(7,02)   UP      FPC00FE2(7,02)->FPC00FE1(7,02)   UP
…
FPC 0 FPFE 3 :
-----------------------
          In-links               State             Out-links              State
--------------------------------------------------------------------------------
FPC00FE0(0,00)->FPC00FE3(0,00)   UP      FPC00FE3(0,00)->FPC00FE0(0,00)   UP
FPC00FE0(0,01)->FPC00FE3(0,01)   UP      FPC00FE3(0,01)->FPC00FE0(0,01)   UP
FPC00FE0(0,02)->FPC00FE3(0,02)   UP      FPC00FE3(0,02)->FPC00FE0(0,02)   UP
FPC00FE0(0,03)->FPC00FE3(0,03)   UP      FPC00FE3(0,03)->FPC00FE0(0,03)   UP
FPC00FE0(0,04)->FPC00FE3(0,04)   UP      FPC00FE3(0,04)->FPC00FE0(0,04)   UP
FPC00FE0(0,05)->FPC00FE3(0,05)   UP      FPC00FE3(0,05)->FPC00FE0(0,05)   UP
FPC00FE0(0,06)->FPC00FE3(0,06)   UP      FPC00FE3(0,06)->FPC00FE0(0,06)   UP
FPC00FE0(0,07)->FPC00FE3(0,07)   UP      FPC00FE3(0,07)->FPC00FE0(0,07)   UP
…
```

show chassis fabric topology (PTX10016 Router)

```
user@host> show chassis fabric topology

             In-link  : FPC# FE# (TX inst#, TX sub-chnl #) ->
                        SIB# ASIC#_FCORE# (RX port#,RX sub-chn#, RX inst#)
```

```
             Out-link : SIB# ASIC#_FCORE#(TX port#, TX sub-chn#, TX inst#) ->
                        FPC# FE# (RX inst#, RX sub-chnl #)
            SIB 0 FCHIP 0 FCORE 0 :
            -----------------------
                      In-links               State             Out-links              State
            -------------------------------------------------------------------------------
            FPC00FE1(0,05)->S00F0_0(46,0,46) DOWN      S00F0_0(46,0,46)->FPC00FE1(0,05) FAULT
            FPC00FE1(0,07)->S00F0_0(46,1,46) UP      S00F0_0(46,1,46)->FPC00FE1(0,07) UP

user@host> show chassis fabric topology
 In-link  : FPC# FE# ASIC# (TX inst#, TX sub-chnl #) ->
            SIB# ASIC#_FCORE# (RX port#, RX sub-chn#, RX inst#)

 Out-link : SIB# ASIC#_FCORE# (TX port#, TX sub-chn#, TX inst#) ->
            FPC# FE# ASIC# (RX inst#, RX sub-chnl #)
SIB 0 FCHIP 0 FCORE 0 :
-----------------------
          In-links               State             Out-links              State
-------------------------------------------------------------------------------
FPC00FE0(1,17)->S00F0_0(01,0,01) OK      S00F0_0(00,0,00)->FPC00FE0(1,09) OK
FPC00FE0(1,09)->S00F0_0(02,0,02) OK      S00F0_0(00,1,00)->FPC00FE0(1,17) OK
FPC00FE0(1,07)->S00F0_0(02,2,02) OK      S00F0_0(00,2,00)->FPC00FE0(1,07) OK
FPC00FE1(1,12)->S00F0_0(01,1,01) OK      S00F0_0(00,3,00)->FPC00FE1(1,06) OK
FPC00FE1(1,06)->S00F0_0(01,2,01) OK      S00F0_0(01,1,01)->FPC00FE1(1,12) OK
FPC00FE1(1,10)->S00F0_0(01,3,01) OK      S00F0_0(01,3,01)->FPC00FE1(1,10) OK
FPC00FE2(1,16)->S00F0_0(00,4,00) OK      S00F0_0(00,4,00)->FPC00FE2(1,08) OK
FPC00FE2(1,08)->S00F0_0(01,6,01) OK      S00F0_0(00,5,00)->FPC00FE2(1,16) OK
FPC00FE2(1,06)->S00F0_0(01,7,01) OK      S00F0_0(00,6,00)->FPC00FE2(1,06) OK
FPC05FE0(1,07)->S00F0_0(05,5,05) OK      S00F0_0(05,2,05)->FPC05FE0(1,17) OK
FPC05FE0(1,09)->S00F0_0(05,7,05) OK      S00F0_0(06,4,06)->FPC05FE0(1,07) OK
FPC05FE0(1,17)->S00F0_0(09,3,09) OK      S00F0_0(06,7,06)->FPC05FE0(1,09) OK
FPC05FE1(1,06)->S00F0_0(06,1,06) OK      S00F0_0(06,0,06)->FPC05FE1(1,06) OK
FPC05FE1(1,08)->S00F0_0(06,3,06) OK      S00F0_0(06,2,06)->FPC05FE1(1,08) OK
FPC05FE1(1,16)->S00F0_0(09,7,09) OK      S00F0_0(09,6,09)->FPC05FE1(1,16) OK
FPC05FE2(1,10)->S00F0_0(09,0,09) OK      S00F0_0(05,0,05)->FPC05FE2(1,06) OK
FPC05FE2(1,06)->S00F0_0(09,1,09) OK      S00F0_0(05,1,05)->FPC05FE2(1,10) OK
FPC05FE2(1,12)->S00F0_0(09,2,09) OK      S00F0_0(05,3,05)->FPC05FE2(1,12) OK
FPC05FE3(1,11)->S00F0_0(09,4,09) OK      S00F0_0(09,4,09)->FPC05FE3(1,07) OK
FPC05FE3(1,07)->S00F0_0(09,5,09) OK      S00F0_0(09,5,09)->FPC05FE3(1,11) OK
FPC05FE3(1,13)->S00F0_0(09,6,09) OK      S00F0_0(09,7,09)->FPC05FE3(1,13) OK
FPC05FE4(1,16)->S00F0_0(05,3,05) OK      S00F0_0(05,4,05)->FPC05FE4(1,06) OK
FPC05FE4(1,06)->S00F0_0(06,5,06) OK      S00F0_0(05,6,05)->FPC05FE4(1,08) OK
```

```
FPC05FE4(1,08)->S00F0_0(06,7,06) OK      S00F0_0(09,2,09)->FPC05FE4(1,16) OK
FPC05FE5(1,10)->S00F0_0(05,0,05) OK      S00F0_0(09,0,09)->FPC05FE5(1,06) OK
FPC05FE5(1,06)->S00F0_0(05,1,05) OK      S00F0_0(09,1,09)->FPC05FE5(1,10) OK
FPC05FE5(1,12)->S00F0_0(05,2,05) OK      S00F0_0(09,3,09)->FPC05FE5(1,12) OK
FPC06FE0(1,17)->S00F0_0(05,6,05) OK      S00F0_0(06,6,06)->FPC06FE0(1,17) OK
FPC06FE0(1,07)->S00F0_0(07,0,07) OK      S00F0_0(08,0,08)->FPC06FE0(1,07) OK
FPC06FE0(1,09)->S00F0_0(07,2,07) OK      S00F0_0(08,2,08)->FPC06FE0(1,09) OK
FPC06FE1(1,16)->S00F0_0(06,2,06) OK      S00F0_0(06,3,06)->FPC06FE1(1,16) OK
FPC06FE1(1,06)->S00F0_0(07,4,07) OK      S00F0_0(07,4,07)->FPC06FE1(1,06) OK
FPC06FE1(1,08)->S00F0_0(07,6,07) OK      S00F0_0(07,6,07)->FPC06FE1(1,08) OK
FPC06FE2(1,06)->S00F0_0(05,4,05) OK      S00F0_0(06,5,06)->FPC06FE2(1,06) OK
FPC06FE2(1,10)->S00F0_0(07,1,07) OK      S00F0_0(08,1,08)->FPC06FE2(1,10) OK
FPC06FE2(1,12)->S00F0_0(07,3,07) OK      S00F0_0(08,3,08)->FPC06FE2(1,12) OK
FPC06FE3(1,07)->S00F0_0(06,0,06) OK      S00F0_0(06,1,06)->FPC06FE3(1,07) OK
FPC06FE3(1,11)->S00F0_0(07,5,07) OK      S00F0_0(07,5,07)->FPC06FE3(1,11) OK
FPC06FE3(1,13)->S00F0_0(07,7,07) OK      S00F0_0(07,7,07)->FPC06FE3(1,13) OK
FPC06FE4(1,16)->S00F0_0(06,6,06) OK      S00F0_0(05,7,05)->FPC06FE4(1,16) OK
FPC06FE4(1,06)->S00F0_0(08,0,08) OK      S00F0_0(07,0,07)->FPC06FE4(1,06) OK
FPC06FE4(1,08)->S00F0_0(08,2,08) OK      S00F0_0(07,2,07)->FPC06FE4(1,08) OK
FPC06FE5(1,06)->S00F0_0(06,4,06) OK      S00F0_0(05,5,05)->FPC06FE5(1,06) OK
FPC06FE5(1,10)->S00F0_0(08,1,08) OK      S00F0_0(07,1,07)->FPC06FE5(1,10) OK
FPC06FE5(1,12)->S00F0_0(08,3,08) OK      S00F0_0(07,3,07)->FPC06FE5(1,12) OK
...
...
FPC15FE5(1,01)->S01F0_0(08,3,08) OK      S01F0_0(07,3,07)->FPC06FE5(1,01) OK

SIB 0 FCHIP 1 FCORE 0 :
-----------------------
          In-links               State             Out-links              State
--------------------------------------------------------------------------------
FPC00FE0(1,15)->S00F1_0(15,4,15) OK      S00F1_0(16,4,16)->FPC00FE0(1,15) OK
FPC00FE0(1,11)->S00F1_0(17,4,17) OK      S00F1_0(18,4,18)->FPC00FE0(1,11) OK
FPC00FE0(1,13)->S00F1_0(17,6,17) OK      S00F1_0(18,6,18)->FPC00FE0(1,13) OK
FPC00FE1(1,08)->S00F1_0(15,6,15) OK      S00F1_0(16,6,16)->FPC00FE1(1,08) OK
FPC00FE1(1,14)->S00F1_0(17,5,17) OK      S00F1_0(18,5,18)->FPC00FE1(1,14) OK
FPC00FE1(1,16)->S00F1_0(17,7,17) OK      S00F1_0(18,7,18)->FPC00FE1(1,16) OK
FPC00FE2(1,14)->S00F1_0(16,0,16) OK      S00F1_0(16,0,16)->FPC00FE2(1,14) OK
FPC00FE2(1,10)->S00F1_0(18,0,18) OK      S00F1_0(18,0,18)->FPC00FE2(1,10) OK
FPC00FE2(1,12)->S00F1_0(18,2,18) OK      S00F1_0(18,2,18)->FPC00FE2(1,12) OK
FPC05FE0(1,11)->S00F1_0(02,0,02) OK      S00F1_0(02,1,02)->FPC05FE0(1,13) OK
FPC05FE0(1,13)->S00F1_0(02,2,02) OK      S00F1_0(02,3,02)->FPC05FE0(1,11) OK
FPC05FE0(1,15)->S00F1_0(04,7,04) OK      S00F1_0(03,6,03)->FPC05FE0(1,15) OK
FPC05FE1(1,10)->S00F1_0(02,4,02) OK      S00F1_0(02,5,02)->FPC05FE1(1,12) OK
```

```
FPC05FE1(1,12)->S00F1_0(02,6,02) OK      S00F1_0(02,7,02)->FPC05FE1(1,10) OK
FPC05FE1(1,14)->S00F1_0(04,3,04) OK      S00F1_0(04,2,04)->FPC05FE1(1,14) OK
FPC05FE2(1,16)->S00F1_0(04,4,04) OK      S00F1_0(03,4,03)->FPC05FE2(1,08) OK
FPC05FE2(1,08)->S00F1_0(04,5,04) OK      S00F1_0(03,5,03)->FPC05FE2(1,16) OK
FPC05FE2(1,14)->S00F1_0(04,6,04) OK      S00F1_0(03,7,03)->FPC05FE2(1,14) OK
FPC05FE3(1,17)->S00F1_0(04,0,04) OK      S00F1_0(04,0,04)->FPC05FE3(1,09) OK
FPC05FE3(1,09)->S00F1_0(04,1,04) OK      S00F1_0(04,1,04)->FPC05FE3(1,17) OK
FPC05FE3(1,15)->S00F1_0(04,2,04) OK      S00F1_0(04,3,04)->FPC05FE3(1,15) OK
FPC05FE4(1,10)->S00F1_0(03,0,03) OK      S00F1_0(03,1,03)->FPC05FE4(1,12) OK
FPC05FE4(1,12)->S00F1_0(03,2,03) OK      S00F1_0(03,3,03)->FPC05FE4(1,10) OK
FPC05FE4(1,14)->S00F1_0(03,7,03) OK      S00F1_0(04,6,04)->FPC05FE4(1,14) OK
FPC05FE5(1,16)->S00F1_0(03,4,03) OK      S00F1_0(04,4,04)->FPC05FE5(1,08) OK
FPC05FE5(1,08)->S00F1_0(03,5,03) OK      S00F1_0(04,5,04)->FPC05FE5(1,16) OK
FPC05FE5(1,14)->S00F1_0(03,6,03) OK      S00F1_0(04,7,04)->FPC05FE5(1,14) OK
FPC06FE0(1,15)->S00F1_0(01,0,01) OK      S00F1_0(00,3,00)->FPC06FE0(1,15) OK
FPC06FE0(1,11)->S00F1_0(02,1,02) OK      S00F1_0(01,0,01)->FPC06FE0(1,13) OK
FPC06FE0(1,13)->S00F1_0(02,3,02) OK      S00F1_0(01,2,01)->FPC06FE0(1,11) OK
FPC06FE1(1,14)->S00F1_0(01,4,01) OK      S00F1_0(00,7,00)->FPC06FE1(1,14) OK
FPC06FE1(1,10)->S00F1_0(02,5,02) OK      S00F1_0(01,4,01)->FPC06FE1(1,12) OK
FPC06FE1(1,12)->S00F1_0(02,7,02) OK      S00F1_0(01,6,01)->FPC06FE1(1,10) OK
FPC06FE2(1,08)->S00F1_0(01,2,01) OK      S00F1_0(00,1,00)->FPC06FE2(1,08) OK
FPC06FE2(1,16)->S00F1_0(15,0,15) OK      S00F1_0(01,5,01)->FPC06FE2(1,16) OK
FPC06FE2(1,14)->S00F1_0(15,2,15) OK      S00F1_0(01,7,01)->FPC06FE2(1,14) OK
FPC06FE3(1,09)->S00F1_0(01,6,01) OK      S00F1_0(00,5,00)->FPC06FE3(1,09) OK
FPC06FE3(1,17)->S00F1_0(19,4,19) OK      S00F1_0(02,4,02)->FPC06FE3(1,17) OK
FPC06FE3(1,15)->S00F1_0(19,6,19) OK      S00F1_0(02,6,02)->FPC06FE3(1,15) OK
FPC06FE4(1,14)->S00F1_0(01,7,01) OK      S00F1_0(01,3,01)->FPC06FE4(1,14) OK
FPC06FE4(1,10)->S00F1_0(03,1,03) OK      S00F1_0(02,0,02)->FPC06FE4(1,12) OK
FPC06FE4(1,12)->S00F1_0(03,3,03) OK      S00F1_0(02,2,02)->FPC06FE4(1,10) OK
FPC06FE5(1,08)->S00F1_0(01,5,01) OK      S00F1_0(01,1,01)->FPC06FE5(1,08) OK
FPC06FE5(1,16)->S00F1_0(19,0,19) OK      S00F1_0(03,0,03)->FPC06FE5(1,16) OK
FPC06FE5(1,14)->S00F1_0(19,2,19) OK      S00F1_0(03,2,03)->FPC06FE5(1,14) OK
...
...
FPC15FE5(1,01)->S01F0_0(08,3,08) OK      S01F0_0(07,3,07)->FPC06FE5(1,01) OK

SIB 1 FCHIP 0 FCORE 0 :
-----------------------
          In-links               State             Out-links              State
--------------------------------------------------------------------------------
FPC00FE0(1,05)->S01F0_0(01,0,01) OK      S01F0_0(00,0,00)->FPC00FE0(0,01) OK
FPC00FE0(0,01)->S01F0_0(02,0,02) OK      S01F0_0(00,1,00)->FPC00FE0(1,05) OK
FPC00FE0(0,07)->S01F0_0(02,2,02) OK      S01F0_0(00,2,00)->FPC00FE0(0,07) OK
```

```
FPC00FE1(1,04)->S01F0_0(00,4,00) OK      S01F0_0(00,4,00)->FPC00FE1(0,00) OK
FPC00FE1(0,00)->S01F0_0(01,6,01) OK      S01F0_0(00,5,00)->FPC00FE1(1,04) OK
FPC00FE1(0,06)->S01F0_0(01,7,01) OK      S01F0_0(00,6,00)->FPC00FE1(0,06) OK
FPC00FE2(1,01)->S01F0_0(01,1,01) OK      S01F0_0(00,3,00)->FPC00FE2(0,06) OK
FPC00FE2(0,06)->S01F0_0(01,2,01) OK      S01F0_0(01,1,01)->FPC00FE2(1,01) OK
FPC00FE2(1,00)->S01F0_0(01,3,01) OK      S01F0_0(01,3,01)->FPC00FE2(1,00) OK
FPC05FE0(0,07)->S01F0_0(05,5,05) OK      S01F0_0(05,2,05)->FPC05FE0(1,05) OK
FPC05FE0(0,01)->S01F0_0(05,7,05) OK      S01F0_0(06,4,06)->FPC05FE0(0,07) OK
FPC05FE0(1,05)->S01F0_0(09,3,09) OK      S01F0_0(06,7,06)->FPC05FE0(0,01) OK
FPC05FE1(1,00)->S01F0_0(09,0,09) OK      S01F0_0(05,0,05)->FPC05FE1(0,06) OK
FPC05FE1(0,06)->S01F0_0(09,1,09) OK      S01F0_0(05,1,05)->FPC05FE1(1,00) OK
FPC05FE1(1,01)->S01F0_0(09,2,09) OK      S01F0_0(05,3,05)->FPC05FE1(1,01) OK
FPC05FE2(0,06)->S01F0_0(06,1,06) OK      S01F0_0(06,0,06)->FPC05FE2(0,06) OK
FPC05FE2(0,00)->S01F0_0(06,3,06) OK      S01F0_0(06,2,06)->FPC05FE2(0,00) OK
FPC05FE2(1,04)->S01F0_0(09,7,09) OK      S01F0_0(09,6,09)->FPC05FE2(1,04) OK
FPC05FE3(1,00)->S01F0_0(09,4,09) OK      S01F0_0(09,4,09)->FPC05FE3(0,06) OK
FPC05FE3(0,06)->S01F0_0(09,5,09) OK      S01F0_0(09,5,09)->FPC05FE3(1,00) OK
FPC05FE3(1,01)->S01F0_0(09,6,09) OK      S01F0_0(09,7,09)->FPC05FE3(1,01) OK
FPC05FE4(0,04)->S01F0_0(05,3,05) OK      S01F0_0(05,4,05)->FPC05FE4(0,14) OK
FPC05FE4(0,14)->S01F0_0(06,5,06) OK      S01F0_0(05,6,05)->FPC05FE4(0,16) OK
FPC05FE4(0,16)->S01F0_0(06,7,06) OK      S01F0_0(09,2,09)->FPC05FE4(0,04) OK
FPC05FE5(1,00)->S01F0_0(05,0,05) OK      S01F0_0(09,0,09)->FPC05FE5(0,06) OK
FPC05FE5(0,06)->S01F0_0(05,1,05) OK      S01F0_0(09,1,09)->FPC05FE5(1,00) OK
FPC05FE5(1,01)->S01F0_0(05,2,05) OK      S01F0_0(09,3,09)->FPC05FE5(1,01) OK
FPC06FE0(1,05)->S01F0_0(05,6,05) OK      S01F0_0(06,6,06)->FPC06FE0(1,05) OK
FPC06FE0(0,07)->S01F0_0(07,0,07) OK      S01F0_0(08,0,08)->FPC06FE0(0,07) OK
FPC06FE0(0,01)->S01F0_0(07,2,07) OK      S01F0_0(08,2,08)->FPC06FE0(0,01) OK
FPC06FE1(0,06)->S01F0_0(05,4,05) OK      S01F0_0(06,5,06)->FPC06FE1(0,06) OK
FPC06FE1(1,00)->S01F0_0(07,1,07) OK      S01F0_0(08,1,08)->FPC06FE1(1,00) OK
FPC06FE1(1,01)->S01F0_0(07,3,07) OK      S01F0_0(08,3,08)->FPC06FE1(1,01) OK
FPC06FE2(1,04)->S01F0_0(06,2,06) OK      S01F0_0(06,3,06)->FPC06FE2(1,04) OK
FPC06FE2(0,06)->S01F0_0(07,4,07) OK      S01F0_0(07,4,07)->FPC06FE2(0,06) OK
FPC06FE2(0,00)->S01F0_0(07,6,07) OK      S01F0_0(07,6,07)->FPC06FE2(0,00) OK
FPC06FE3(0,06)->S01F0_0(06,0,06) OK      S01F0_0(06,1,06)->FPC06FE3(0,06) OK
FPC06FE3(1,00)->S01F0_0(07,5,07) OK      S01F0_0(07,5,07)->FPC06FE3(1,00) OK
FPC06FE3(1,01)->S01F0_0(07,7,07) OK      S01F0_0(07,7,07)->FPC06FE3(1,01) OK
FPC06FE4(0,04)->S01F0_0(06,6,06) OK      S01F0_0(05,7,05)->FPC06FE4(0,04) OK
FPC06FE4(0,14)->S01F0_0(08,0,08) OK      S01F0_0(07,0,07)->FPC06FE4(0,14) OK
FPC06FE4(0,16)->S01F0_0(08,2,08) OK      S01F0_0(07,2,07)->FPC06FE4(0,16) OK
FPC06FE5(0,06)->S01F0_0(06,4,06) OK      S01F0_0(05,5,05)->FPC06FE5(0,06) OK
FPC06FE5(1,00)->S01F0_0(08,1,08) OK      S01F0_0(07,1,07)->FPC06FE5(1,00) OK
FPC06FE5(1,01)->S01F0_0(08,3,08) OK      S01F0_0(07,3,07)->FPC06FE5(1,01) OK
...
```

```
...
FPC15FE5(1,01)->S01F0_0(08,3,08) OK      S01F0_0(07,3,07)->FPC06FE5(1,01) OK

SIB 1 FCHIP 1 FCORE 0 :
-----------------------
          In-links               State             Out-links              State
--------------------------------------------------------------------------------
FPC00FE0(1,03)->S01F1_0(15,4,15) OK      S01F1_0(16,4,16)->FPC00FE0(1,03) OK
FPC00FE0(0,02)->S01F1_0(17,4,17) OK      S01F1_0(18,4,18)->FPC00FE0(0,02) OK
FPC00FE0(0,03)->S01F1_0(17,6,17) OK      S01F1_0(18,6,18)->FPC00FE0(0,03) OK
FPC00FE1(1,02)->S01F1_0(16,0,16) OK      S01F1_0(16,0,16)->FPC00FE1(1,02) OK
FPC00FE1(1,00)->S01F1_0(18,0,18) OK      S01F1_0(18,0,18)->FPC00FE1(1,00) OK
FPC00FE1(1,01)->S01F1_0(18,2,18) OK      S01F1_0(18,2,18)->FPC00FE1(1,01) OK
FPC00FE2(0,00)->S01F1_0(15,6,15) OK      S01F1_0(16,6,16)->FPC00FE2(0,00) OK
FPC00FE2(1,02)->S01F1_0(17,5,17) OK      S01F1_0(18,5,18)->FPC00FE2(1,02) OK
FPC00FE2(1,04)->S01F1_0(17,7,17) OK      S01F1_0(18,7,18)->FPC00FE2(1,04) OK
FPC05FE0(0,02)->S01F1_0(02,0,02) OK      S01F1_0(02,1,02)->FPC05FE0(0,03) OK
FPC05FE0(0,03)->S01F1_0(02,2,02) OK      S01F1_0(02,3,02)->FPC05FE0(0,02) OK
FPC05FE0(1,03)->S01F1_0(04,7,04) OK      S01F1_0(03,6,03)->FPC05FE0(1,03) OK
FPC05FE1(1,04)->S01F1_0(04,4,04) OK      S01F1_0(03,4,03)->FPC05FE1(0,00) OK
FPC05FE1(0,00)->S01F1_0(04,5,04) OK      S01F1_0(03,5,03)->FPC05FE1(1,04) OK
FPC05FE1(1,02)->S01F1_0(04,6,04) OK      S01F1_0(03,7,03)->FPC05FE1(1,02) OK
FPC05FE2(1,00)->S01F1_0(02,4,02) OK      S01F1_0(02,5,02)->FPC05FE2(1,01) OK
FPC05FE2(1,01)->S01F1_0(02,6,02) OK      S01F1_0(02,7,02)->FPC05FE2(1,00) OK
FPC05FE2(1,02)->S01F1_0(04,3,04) OK      S01F1_0(04,2,04)->FPC05FE2(1,02) OK
FPC05FE3(1,04)->S01F1_0(04,0,04) OK      S01F1_0(04,0,04)->FPC05FE3(0,00) OK
FPC05FE3(0,00)->S01F1_0(04,1,04) OK      S01F1_0(04,1,04)->FPC05FE3(1,04) OK
FPC05FE3(1,02)->S01F1_0(04,2,04) OK      S01F1_0(04,3,04)->FPC05FE3(1,02) OK
FPC05FE4(0,10)->S01F1_0(03,0,03) OK      S01F1_0(03,1,03)->FPC05FE4(0,12) OK
FPC05FE4(0,12)->S01F1_0(03,2,03) OK      S01F1_0(03,3,03)->FPC05FE4(0,10) OK
FPC05FE4(0,08)->S01F1_0(03,7,03) OK      S01F1_0(04,6,04)->FPC05FE4(0,08) OK
FPC05FE5(1,04)->S01F1_0(03,4,03) OK      S01F1_0(04,4,04)->FPC05FE5(0,00) OK
FPC05FE5(0,00)->S01F1_0(03,5,03) OK      S01F1_0(04,5,04)->FPC05FE5(1,04) OK
FPC05FE5(1,02)->S01F1_0(03,6,03) OK      S01F1_0(04,7,04)->FPC05FE5(1,02) OK
FPC06FE0(1,03)->S01F1_0(01,0,01) OK      S01F1_0(00,3,00)->FPC06FE0(1,03) OK
FPC06FE0(0,02)->S01F1_0(02,1,02) OK      S01F1_0(01,0,01)->FPC06FE0(0,03) OK
FPC06FE0(0,03)->S01F1_0(02,3,02) OK      S01F1_0(01,2,01)->FPC06FE0(0,02) OK
FPC06FE1(0,00)->S01F1_0(01,2,01) OK      S01F1_0(00,1,00)->FPC06FE1(0,00) OK
FPC06FE1(1,04)->S01F1_0(15,0,15) OK      S01F1_0(01,5,01)->FPC06FE1(1,04) OK
FPC06FE1(1,02)->S01F1_0(15,2,15) OK      S01F1_0(01,7,01)->FPC06FE1(1,02) OK
FPC06FE2(1,02)->S01F1_0(01,4,01) OK      S01F1_0(00,7,00)->FPC06FE2(1,02) OK
FPC06FE2(1,00)->S01F1_0(02,5,02) OK      S01F1_0(01,4,01)->FPC06FE2(1,01) OK
FPC06FE2(1,01)->S01F1_0(02,7,02) OK      S01F1_0(01,6,01)->FPC06FE2(1,00) OK
```

```
FPC06FE3(0,00)->S01F1_0(01,6,01) OK      S01F1_0(00,5,00)->FPC06FE3(0,00) OK
FPC06FE3(1,04)->S01F1_0(19,4,19) OK      S01F1_0(02,4,02)->FPC06FE3(1,04) OK
FPC06FE3(1,02)->S01F1_0(19,6,19) OK      S01F1_0(02,6,02)->FPC06FE3(1,02) OK
FPC06FE4(0,08)->S01F1_0(01,7,01) OK      S01F1_0(01,3,01)->FPC06FE4(0,08) OK
FPC06FE4(0,10)->S01F1_0(03,1,03) OK      S01F1_0(02,0,02)->FPC06FE4(0,12) OK
FPC06FE4(0,12)->S01F1_0(03,3,03) OK      S01F1_0(02,2,02)->FPC06FE4(0,10) OK
FPC06FE5(0,00)->S01F1_0(01,5,01) OK      S01F1_0(01,1,01)->FPC06FE5(0,00) OK
FPC06FE5(1,04)->S01F1_0(19,0,19) OK      S01F1_0(03,0,03)->FPC06FE5(1,04) OK
FPC06FE5(1,02)->S01F1_0(19,2,19) OK      S01F1_0(03,2,03)->FPC06FE5(1,02) OK
...
...
FPC15FE5(1,01)->S01F0_0(08,3,08) OK      S01F0_0(07,3,07)->FPC06FE5(1,01) OK

SIB 2
Not Online

SIB 3
Not Online

SIB 4
Not Online

SIB 5
Not Online
```

show chassis fabric topology (QFX10008 Switch)

```
user@host> show chassis fabric topology
 In-link  : FPC# FE# ASIC# (TX inst#, TX sub-chnl #) ->
            SIB# ASIC#_FCORE# (RX port#, RX sub-chn#, RX inst#)

 Out-link : SIB# ASIC#_FCORE# (TX port#, TX sub-chn#, TX inst#) ->
            FPC# FE# ASIC# (RX inst#, RX sub-chnl #)
SIB 0 FCHIP 0 FCORE 0 :
-----------------------
          In-links               State             Out-links              State
--------------------------------------------------------------------------------
FPC00FE0(1,17)->S00F0_0(01,0,01) OK      S00F0_0(00,0,00)->FPC00FE0(1,09) OK
FPC00FE0(1,09)->S00F0_0(02,0,02) OK      S00F0_0(00,1,00)->FPC00FE0(1,17) OK
FPC00FE0(1,07)->S00F0_0(02,2,02) OK      S00F0_0(00,2,00)->FPC00FE0(1,07) OK
FPC00FE1(1,12)->S00F0_0(01,1,01) OK      S00F0_0(00,3,00)->FPC00FE1(1,06) OK
```

```
FPC00FE1(1,06)->S00F0_0(01,2,01) OK      S00F0_0(01,1,01)->FPC00FE1(1,12) OK
FPC00FE1(1,10)->S00F0_0(01,3,01) OK      S00F0_0(01,3,01)->FPC00FE1(1,10) OK
FPC00FE2(1,16)->S00F0_0(00,4,00) OK      S00F0_0(00,4,00)->FPC00FE2(1,08) OK
FPC00FE2(1,08)->S00F0_0(01,6,01) OK      S00F0_0(00,5,00)->FPC00FE2(1,16) OK
FPC00FE2(1,06)->S00F0_0(01,7,01) OK      S00F0_0(00,6,00)->FPC00FE2(1,06) OK

SIB 0 FCHIP 1 FCORE 0 :
-----------------------
          In-links               State             Out-links              State
--------------------------------------------------------------------------------
FPC00FE0(1,15)->S00F1_0(15,4,15) OK      S00F1_0(16,4,16)->FPC00FE0(1,15) OK
FPC00FE0(1,11)->S00F1_0(17,4,17) OK      S00F1_0(18,4,18)->FPC00FE0(1,11) OK
FPC00FE0(1,13)->S00F1_0(17,6,17) OK      S00F1_0(18,6,18)->FPC00FE0(1,13) OK
FPC00FE1(1,08)->S00F1_0(15,6,15) OK      S00F1_0(16,6,16)->FPC00FE1(1,08) OK
FPC00FE1(1,14)->S00F1_0(17,5,17) OK      S00F1_0(18,5,18)->FPC00FE1(1,14) OK
FPC00FE1(1,16)->S00F1_0(17,7,17) OK      S00F1_0(18,7,18)->FPC00FE1(1,16) OK
FPC00FE2(1,14)->S00F1_0(16,0,16) OK      S00F1_0(16,0,16)->FPC00FE2(1,14) OK
FPC00FE2(1,10)->S00F1_0(18,0,18) OK      S00F1_0(18,0,18)->FPC00FE2(1,10) OK
FPC00FE2(1,12)->S00F1_0(18,2,18) OK      S00F1_0(18,2,18)->FPC00FE2(1,12) OK

SIB 1
Not Online

SIB 2
Not Online

SIB 3
Not Online

SIB 4
Not Online

SIB 5
Not Online
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

Layer 2 Wholesale Network Topology Overview

show chassis fabric unreachable-destinations

IN THIS SECTION

- Syntax | 18257
- Description | 18257
- Required Privilege Level | 18257
- Output Fields | 18257
- Sample Output | 18258
- Release Information | 18260

Syntax

```
show chassis fabric unreachable-destinations
```

Description

(M320 and T Series routers only) Display the list of destinations that have transitioned from a reachable state to an unreachable state.

Required Privilege Level

view

Output Fields

The table lists the output fields for the `show chassis fabric unreachable-destinations` command. Output fields are listed in the approximate order in which they appear.

Table 272: show chassis fabric unreachable-destinations Output Fields

Field Name	Field Description
Flexible PIC Concentrator (FPC) number	Source FPC number where unreachable destinations are present.
Packet Forwarding Engine number	Source Packet Forwarding Engine number where unreachable destinations are present.
Destination error on Packet Forwarding Engine	List of destination FPCs *FPC number*/Packet Forwarding Engines *Packet Forwarding Engine number* that are not reachable from the source FPCs *FPC number*/Packet Forwarding Engines *Packet Forwarding Engine number* over the fabric.

Sample Output

show chassis fabric unreachable-destinations(T640 and T1600 routers)

```
user@host> show chassis fabric unreachable-destinations
Fabric management unreachable destinations:
FPC 2
  PFE 0
      Destination error on PFEs    2/0 3/0 3/1 7/0
FPC 3
  PFE 0
      Destination error on PFEs    2/0 3/0 3/1 7/0
FPC 3
  PFE 1
      Destination error on PFEs    2/0 3/0 3/1 7/0
FPC 7
  PFE 0
      Destination error on PFEs    2/0 3/0 3/1 7/0
```

show chassis fabric unreachable-destinations(TX Matrix routers)

```
user@host> show chassis fabric unreachable-destinations
Fabric management unreachable destinations:
FPC 10
  PFE 0
      Destination error on PFEs    10/0 16/0 16/1 17/0 17/1 19/0 20/1 21/1 22/1 24/0
       26/0 27/0 27/1 28/1 29/1 31/1
FPC 12
  PFE 0
      Destination error on PFEs    12/0 16/0 16/1 17/0 17/1 19/0 20/1 21/1 22/1 24/0
       26/0 27/0 27/1 28/1 29/1 31/1
FPC 16
  PFE 0
      Destination error on PFEs    10/0 12/0
FPC 16
  PFE 1
      Destination error on PFEs    10/0 12/0
FPC 17
  PFE 0
      Destination error on PFEs    10/0 12/0
FPC 17
  PFE 1
      Destination error on PFEs    10/0 12/0
FPC 19
  PFE 0
      Destination error on PFEs    10/0 12/0
FPC 20
  PFE 1
      Destination error on PFEs    10/0 12/0
FPC 21
  PFE 1
      Destination error on PFEs    10/0 12/0
FPC 22
  PFE 1
      Destination error on PFEs    10/0 12/0
FPC 24
  PFE 0
      Destination error on PFEs    10/0 12/0
FPC 26
  PFE 0
      Destination error on PFEs    10/0 12/0
```

```
FPC 27
  PFE 0
      Destination error on PFEs    10/0 12/0
FPC 27
  PFE 1
      Destination error on PFEs    10/0
FPC 28
  PFE 1
      Destination error on PFEs    10/0 12/0
FPC 29
  PFE 1
      Destination error on PFEs    10/0 12/0
FPC 31
  PFE 1
      Destination error on PFEs    10/0 12/0
```

Release Information

Command introduced before Junos OS Release 11.4.

RELATED DOCUMENTATION

show chassis fabric reachability | 18168

show chassis fan

IN THIS SECTION

- Syntax | 18261
- Syntax (MX Series Routers) | 18261
- Syntax (MX104, MX204, MX2010, MX2020, MX2008, and MX10003 Universal Routing Platform) | 18261
- Syntax (QFX Series) | 18261
- Syntax (TX Matrix Router) | 18262
- Syntax (TX Matrix Plus Router) | 18262

- Description | 18262
- Options | 18262
- Required Privilege Level | 18263
- Output Fields | 18263
- Sample Output | 18264
- Release Information | 18278

Syntax

```
show chassis fan
```

Syntax (MX Series Routers)

```
show chassis fan
<all-members>
<local>
<member member-id>
```

Syntax (MX104, MX204, MX2010, MX2020, MX2008, and MX10003 Universal Routing Platform)

```
show chassis fan
<satellite [slot-id slot-id |device-alias alias-name]>
```

Syntax (QFX Series)

```
show chassis fan
<interconnect-device name>
```

Syntax (TX Matrix Router)

```
show chassis fan
<lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
show chassis fan
<lcc number | sfc number>
```

Description

(T Series routers, TX Matrix routers, TX Matrix Plus routers, M120 routers, M320 routers, MX104 routers, MX2010 routers, MX2020 routers, MX2008 routers, MX Series 5G Universal Routing Platforms, QFX3008-I Interconnect devices, QFX Series, OCX Series, EX Series switches, and PTX Series Packet Transport Routers only) Show information about the fan tray and fans.

Options

all-members (MX Series routers only) (Optional) Display information about the fan tray and fans for all members of the Virtual Chassis configuration.

local (MX Series routers only) (Optional) Display information about the fan tray and fans for the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display information about the fan tray and fans for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* variable with a value 0 or 1.

interconnect-device *name* (QFX3000-G QFabric systems only) (Optional) Display information about the fan tray and fans for the specified QFX3008-I Interconnect device.

lcc *number* (TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display information about the fan tray and fans for the specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display information about the fan tray and fans for the specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

satellite [slot-id *slot-id* | device-alias *alias-name*] (Junos Fusion only) (Optional) Display information about the fan tray and fans for the specified satellite device or devices in a Junos Fusion, or for all satellite devices if no satellite devices are specified.

scc (TX Matrix routers only) (Optional) Display information about the fan tray and fans for the TX Matrix router (switch-card chassis).

sfc *number* (TX Matrix Plus routers only) (Optional) Display information about the fan tray and fans for the TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with `0`.

Required Privilege Level

view

Output Fields

Table 273 on page 18263 lists the output fields for the `show chassis fan` command. Output fields are listed in the approximate order in which they appear.

Table 273: show chassis fan Output Fields

Field Name	Field Description
`Item`	Fan item identifier.

Table 273: show chassis fan Output Fields *(Continued)*

Field Name	Field Description
`Status`	Status of the fan: • `OK`—Fan is running properly and within the normal range. • `Check`—Fan is in `Check` state because of some fault or alarm condition.
`RPM`	(T Series routers, TX Matrix routers, TX Matrix Plus routers, MX Series 5G Universal Routing Platforms, QFX3108 Interconnect devices, and EX Series switches only) Fan speed in revolutions per minute (RPM).
% RPM	(PTX10003, MX2010 routers, MX2020 routers, MX2008 routers, and PTX Series Packet Transport Routers only) Percentage of the fan speed being used.
`Measurement`	(T Series routers, TX Matrix routers, TX Matrix Plus routers, MX Series 5G Universal Routing Platforms, QFX3108 Interconnect devices, and EX Series switches only) Fan speed status based on different chassis cooling requirements: • Spinning at high speed • Spinning at intermediate speed • Spinning at normal speed • Spinning at low speed (except EX Series switches) (PTX10003, MX2010 routers, MX2020 routers, MX2008 routers, and PTX Series Packet Transport Routers only) Fan speed in revolutions per minute (RPM) for each fan in the fan tray.

Sample Output

show chassis fan

```
user@host> show chassis fan

      Item                      Status   RPM     Measurement
      Top Tray Fan 1            OK       3790    Spinning at normal speed
```

```
    Top Tray Fan 2            OK       3769    Spinning at normal speed
    Top Tray Fan 3            OK       3769    Spinning at normal speed
    Top Tray Fan 4            OK       3790    Spinning at normal speed
    Top Tray Fan 5            OK       3790    Spinning at normal speed
    Top Tray Fan 6            OK       3769    Spinning at normal speed
    Top Tray Fan 7            OK       3790    Spinning at normal speed
    Top Tray Fan 8            OK       3769    Spinning at normal speed
    Top Tray Fan 9            OK       3769    Spinning at normal speed
    Top Tray Fan 10           OK       3790    Spinning at normal speed
    Top Tray Fan 11           OK       3790    Spinning at normal speed
    Top Tray Fan 12           OK       3769    Spinning at normal speed
    Bottom Tray Fan 1         OK       2880    Spinning at normal speed
    Bottom Tray Fan 2         OK       2912    Spinning at normal speed
    Bottom Tray Fan 3         OK       2928    Spinning at normal speed
    Bottom Tray Fan 4         OK       2896    Spinning at normal speed
    Bottom Tray Fan 5         OK       2896    Spinning at normal speed
    Bottom Tray Fan 6         OK       2928    Spinning at normal speed
```

show chassis fan (QFabric Systems)

```
user@host> show chassis fan interconnect-device interconnect1

    Item                      Status   RPM     Measurement
    TFT 0 Fan 0               OK       2849    Spinning at normal speed
    TFT 0 Fan 1               OK       2821    Spinning at normal speed
    TFT 0 Fan 2               OK       2735    Spinning at normal speed
    TFT 0 Fan 3               OK       2815    Spinning at normal speed
    TFT 0 Fan 4               OK       2828    Spinning at normal speed
    TFT 0 Fan 5               OK       2863    Spinning at normal speed
    BFT 1 Fan 0               OK       2941    Spinning at normal speed
    BFT 1 Fan 1               OK       3008    Spinning at normal speed
    BFT 1 Fan 2               OK       3073    Spinning at normal speed
    BFT 1 Fan 3               OK       2925    Spinning at normal speed
    BFT 1 Fan 4               OK       2863    Spinning at normal speed
    BFT 1 Fan 5               OK       2933    Spinning at normal speed
    SFT 0 Fan 0 Rotor 0       OK       15472   Spinning at normal speed
    SFT 0 Fan 0 Rotor 1       OK       14477   Spinning at normal speed
    SFT 0 Fan 1 Rotor 0       OK       15561   Spinning at normal speed
    SFT 0 Fan 1 Rotor 1       OK       14210   Spinning at normal speed
    SFT 0 Fan 2 Rotor 0       OK       16167   Spinning at normal speed
    SFT 0 Fan 2 Rotor 1       OK       14248   Spinning at normal speed
```

```
        SFT 0 Fan 3 Rotor 0       OK       16463   Spinning at normal speed
        SFT 0 Fan 3 Rotor 1       OK       14099   Spinning at normal speed
        SFT 1 Fan 0 Rotor 0       OK       15083   Spinning at normal speed
        SFT 1 Fan 0 Rotor 1       OK       13533   Spinning at normal speed
        SFT 1 Fan 1 Rotor 0       OK       16071   Spinning at normal speed
        SFT 1 Fan 1 Rotor 1       OK       14400   Spinning at normal speed
        SFT 1 Fan 2 Rotor 0       OK       15517   Spinning at normal speed
        SFT 1 Fan 2 Rotor 1       OK       14210   Spinning at normal speed
        SFT 1 Fan 3 Rotor 0       OK       16413   Spinning at normal speed
        SFT 1 Fan 3 Rotor 1       OK       14400   Spinning at normal speed
        SFT 2 Fan 0 Rotor 0       OK       15297   Spinning at normal speed
        SFT 2 Fan 0 Rotor 1       OK       14634   Spinning at normal speed
        SFT 2 Fan 1 Rotor 0       OK       15561   Spinning at normal speed
        SFT 2 Fan 1 Rotor 1       OK       14285   Spinning at normal speed
        SFT 2 Fan 2 Rotor 0       OK       15835   Spinning at normal speed
        SFT 2 Fan 2 Rotor 1       OK       14400   Spinning at normal speed
        SFT 2 Fan 3 Rotor 0       OK       15789   Spinning at normal speed
        SFT 2 Fan 3 Rotor 1       OK       14323   Spinning at normal speed
        SFT 3 Fan 0 Rotor 0       OK       16314   Spinning at normal speed
        SFT 3 Fan 0 Rotor 1       OK       14876   Spinning at normal speed
        SFT 3 Fan 1 Rotor 0       OK       15835   Spinning at normal speed
        SFT 3 Fan 1 Rotor 1       OK       14323   Spinning at normal speed
        SFT 3 Fan 2 Rotor 0       OK       16265   Spinning at normal speed
        SFT 3 Fan 2 Rotor 1       OK       14594   Spinning at normal speed
        SFT 3 Fan 3 Rotor 0       OK       16071   Spinning at normal speed
        SFT 3 Fan 3 Rotor 1       OK       14323   Spinning at normal speed
        SFT 4 Fan 0 Rotor 0       OK       15652   Spinning at normal speed
        SFT 4 Fan 0 Rotor 1       OK       14438   Spinning at normal speed
        SFT 4 Fan 1 Rotor 0       OK       16167   Spinning at normal speed
        SFT 4 Fan 1 Rotor 1       OK       14555   Spinning at normal speed
        SFT 4 Fan 2 Rotor 0       OK       16023   Spinning at normal speed
        SFT 4 Fan 2 Rotor 1       OK       14361   Spinning at normal speed
        SFT 4 Fan 3 Rotor 0       OK       16216   Spinning at normal speed
        SFT 4 Fan 3 Rotor 1       OK       14438   Spinning at normal speed
        SFT 5 Fan 0 Rotor 0       OK       15297   Spinning at normal speed
        SFT 5 Fan 0 Rotor 1       OK       14173   Spinning at normal speed
        SFT 5 Fan 1 Rotor 0       OK       15472   Spinning at normal speed
        SFT 5 Fan 1 Rotor 1       OK       13846   Spinning at normal speed
        SFT 5 Fan 2 Rotor 0       OK       15340   Spinning at normal speed
        SFT 5 Fan 2 Rotor 1       OK       13917   Spinning at normal speed
        SFT 5 Fan 3 Rotor 0       OK       15835   Spinning at normal speed
        SFT 5 Fan 3 Rotor 1       OK       13917   Spinning at normal speed
        SFT 6 Fan 0 Rotor 0       OK       15743   Spinning at normal speed
```

```
      SFT 6 Fan 0 Rotor 1       OK       14594   Spinning at normal speed
      SFT 6 Fan 1 Rotor 0       OK       16167   Spinning at normal speed
      SFT 6 Fan 1 Rotor 1       OK       14634   Spinning at normal speed
      SFT 6 Fan 2 Rotor 0       OK       16167   Spinning at normal speed
      SFT 6 Fan 2 Rotor 1       OK       14516   Spinning at normal speed
      SFT 6 Fan 3 Rotor 0       OK       16666   Spinning at normal speed
      SFT 6 Fan 3 Rotor 1       OK       14438   Spinning at normal speed
      SFT 7 Fan 0 Rotor 0       OK       15517   Spinning at normal speed
      SFT 7 Fan 0 Rotor 1       OK       14438   Spinning at normal speed
      SFT 7 Fan 1 Rotor 0       OK       15517   Spinning at normal speed
      SFT 7 Fan 1 Rotor 1       OK       14361   Spinning at normal speed
      SFT 7 Fan 2 Rotor 0       OK       16167   Spinning at normal speed
      SFT 7 Fan 2 Rotor 1       OK       14555   Spinning at normal speed
      SFT 7 Fan 3 Rotor 0       OK       15697   Spinning at normal speed
      SFT 7 Fan 3 Rotor 1       OK       14361   Spinning at normal speed
```

show chassis fan (EX Series Switches)

```
user@host> show chassis fan

      Item                      Status   RPM     Measurement
      Fan 1                     OK       3477    Spinning at normal speed
      Fan 2                     OK       3477    Spinning at normal speed
      Fan 3                     OK       3479    Spinning at normal speed
      Fan 4                     OK       3508    Spinning at normal speed
      Fan 5                     OK       3517    Spinning at normal speed
      Fan 6                     OK       3531    Spinning at normal speed
      Fan 7                     OK       3439    Spinning at normal speed
      Fan 8                     OK       3424    Spinning at normal speed
      Fan 9                     OK       3413    Spinning at normal speed
      Fan 10                    OK       3439    Spinning at normal speed
      Fan 11                    OK       3446    Spinning at normal speed
      Fan 12                    OK       3432    Spinning at normal speed
```

show chassis fan (T4000 Core Router)

```
user@host> show chassis fan

      Item                      Status   RPM     Measurement
      Top Left Front fan        OK       5190    Spinning at high speed
```

```
      Top Left Middle fan       OK       5220    Spinning at high speed
      Top Left Rear fan         OK       5190    Spinning at high speed
      Top Right Front fan       OK       5160    Spinning at high speed
      Top Right Middle fan      OK       5190    Spinning at high speed
      Top Right Rear fan        OK       5160    Spinning at high speed
      Bottom Left Front fan     OK       6030    Spinning at high speed
      Bottom Left Middle fan    OK       6090    Spinning at high speed
      Bottom Left Rear fan      OK       6090    Spinning at high speed
      Bottom Right Front fan    OK       6030    Spinning at high speed
      Bottom Right Middle fan   OK       6060    Spinning at high speed
      Bottom Right Rear fan     OK       6060    Spinning at high speed
      Rear Tray Top fan         OK       10000   Spinning at high speed
      Rear Tray Second fan      OK       10000   Spinning at high speed
      Rear Tray Third fan       OK       10000   Spinning at high speed
      Rear Tray Fourth fan      OK       10000   Spinning at high speed
      Rear Tray Fifth fan       OK       10000   Spinning at high speed
      Rear Tray Sixth fan       OK       10000   Spinning at high speed
      Rear Tray Seventh fan     OK       10000   Spinning at high speed
      Rear Tray Bottom fan      OK       10000   Spinning at high speed
```

show chassis fan (TX Matrix Router)

```
user@host> show chassis fan
scc-re0:
--------------------------------------------------------------------------
      Item                      Status   RPM     Measurement
      Top Left Front fan        OK       3420    Spinning at normal speed
      Top Left Middle fan       OK       3390    Spinning at normal speed
      Top Left Rear fan         OK       3420    Spinning at normal speed
      Top Right Front fan       OK       3390    Spinning at normal speed
      Top Right Middle fan      OK       3420    Spinning at normal speed
      Top Right Rear fan        OK       3390    Spinning at normal speed
      Bottom Left Front fan     OK       3420    Spinning at normal speed
      Bottom Left Middle fan    OK       3450    Spinning at normal speed
      Bottom Left Rear fan      OK       3420    Spinning at normal speed
      Bottom Right Front fan    OK       3420    Spinning at normal speed
      Bottom Right Middle fan   OK       3420    Spinning at normal speed
      Bottom Right Rear fan     OK       3420    Spinning at normal speed
      Rear Tray Top fan         OK       3420    Spinning at normal speed
      Rear Tray Second fan      OK       5190    Spinning at normal speed
      Rear Tray Third fan       OK       5190    Spinning at normal speed
```

```
      Rear Tray Fourth fan      OK       5190    Spinning at normal speed
      Rear Tray Fifth fan       OK       3420    Spinning at normal speed
      Rear Tray Sixth fan       OK       3420    Spinning at normal speed
      Rear Tray Seventh fan     OK       3420    Spinning at normal speed
      Rear Tray Bottom fan      OK       3420    Spinning at normal speed

lcc2-re0:
--------------------------------------------------------------------------
      Item                      Status   RPM     Measurement
      Top Left Front fan        OK       3420    Spinning at normal speed
      Top Left Middle fan       OK       3420    Spinning at normal speed
      Top Left Rear fan         OK       3450    Spinning at normal speed
      Top Right Front fan       OK       3420    Spinning at normal speed
      Top Right Middle fan      OK       3450    Spinning at normal speed
      Top Right Rear fan        OK       3360    Spinning at normal speed
      Bottom Left Front fan     OK       3420    Spinning at normal speed
      Bottom Left Middle fan    OK       3480    Spinning at normal speed
      Bottom Left Rear fan      OK       3420    Spinning at normal speed
      Bottom Right Front fan    OK       3420    Spinning at normal speed
      Bottom Right Middle fan   OK       3390    Spinning at normal speed
      Bottom Right Rear fan     OK       3420    Spinning at normal speed
      Rear Tray Top fan         OK       3420    Spinning at normal speed
      Rear Tray Second fan      OK       3420    Spinning at normal speed
      Rear Tray Third fan       OK       3420    Spinning at normal speed
      Rear Tray Fourth fan      OK       3420    Spinning at normal speed
      Rear Tray Fifth fan       OK       3420    Spinning at normal speed
      Rear Tray Sixth fan       OK       3420    Spinning at normal speed
      Rear Tray Seventh fan     OK       3420    Spinning at normal speed
      Rear Tray Bottom fan      OK       3420    Spinning at normal speed
```

show chassis fan (TX Matrix Plus Router)

```
user@host> show chassis fan
  sfc0-re0:
--------------------------------------------------------------------------
      Item                      Status   RPM     Measurement
      Fan Tray 0 Fan 1          OK       4350    Spinning at normal speed
      Fan Tray 0 Fan 2          OK       4380    Spinning at normal speed
      Fan Tray 0 Fan 3          OK       4410    Spinning at normal speed
      Fan Tray 0 Fan 4          OK       4380    Spinning at normal speed
      Fan Tray 0 Fan 5          OK       4350    Spinning at normal speed
```

```
Fan Tray 0 Fan 6          OK       4380    Spinning at normal speed
Fan Tray 1 Fan 1          OK       4410    Spinning at normal speed
Fan Tray 1 Fan 2          OK       4380    Spinning at normal speed
Fan Tray 1 Fan 3          OK       4410    Spinning at normal speed
Fan Tray 1 Fan 4          OK       4380    Spinning at normal speed
Fan Tray 1 Fan 5          OK       4410    Spinning at normal speed
Fan Tray 1 Fan 6          OK       4410    Spinning at normal speed
Fan Tray 2 Fan 1          OK       4380    Spinning at normal speed
Fan Tray 2 Fan 2          OK       4380    Spinning at normal speed
Fan Tray 2 Fan 3          OK       4380    Spinning at normal speed
Fan Tray 2 Fan 4          OK       4410    Spinning at normal speed
Fan Tray 2 Fan 5          OK       4380    Spinning at normal speed
Fan Tray 2 Fan 6          OK       4410    Spinning at normal speed
Fan Tray 2 Fan 7          OK       4410    Spinning at normal speed
Fan Tray 2 Fan 8          OK       4380    Spinning at normal speed
Fan Tray 2 Fan 9          OK       4380    Spinning at normal speed
Fan Tray 3 Fan 1          OK       4350    Spinning at normal speed
Fan Tray 3 Fan 2          OK       4380    Spinning at normal speed
Fan Tray 3 Fan 3          OK       4410    Spinning at normal speed
Fan Tray 3 Fan 4          OK       4440    Spinning at normal speed
Fan Tray 3 Fan 5          OK       4380    Spinning at normal speed
Fan Tray 3 Fan 6          OK       4410    Spinning at normal speed
Fan Tray 3 Fan 7          OK       4410    Spinning at normal speed
Fan Tray 3 Fan 8          OK       4380    Spinning at normal speed
Fan Tray 3 Fan 9          OK       4410    Spinning at normal speed
Fan Tray 4 Fan 1          OK       4410    Spinning at normal speed
Fan Tray 4 Fan 2          OK       4410    Spinning at normal speed
Fan Tray 4 Fan 3          OK       4380    Spinning at normal speed
Fan Tray 4 Fan 4          OK       4380    Spinning at normal speed
Fan Tray 4 Fan 5          OK       4410    Spinning at normal speed
Fan Tray 4 Fan 6          OK       4410    Spinning at normal speed
Fan Tray 4 Fan 7          OK       4410    Spinning at normal speed
Fan Tray 4 Fan 8          OK       4410    Spinning at normal speed
Fan Tray 4 Fan 9          OK       4410    Spinning at normal speed
Fan Tray 5 Fan 1          OK       4350    Spinning at normal speed
Fan Tray 5 Fan 2          OK       4380    Spinning at normal speed
Fan Tray 5 Fan 3          OK       4380    Spinning at normal speed
Fan Tray 5 Fan 4          OK       4350    Spinning at normal speed
Fan Tray 5 Fan 5          OK       4380    Spinning at normal speed
Fan Tray 5 Fan 6          OK       4410    Spinning at normal speed
Fan Tray 5 Fan 7          OK       4410    Spinning at normal speed
Fan Tray 5 Fan 8          OK       4380    Spinning at normal speed
Fan Tray 5 Fan 9          OK       4410    Spinning at normal speed
```

```
lcc0-re0:
--------------------------------------------------------------------------
      Item                      Status   RPM     Measurement
      Top Left Front fan        OK       3420    Spinning at normal speed
      Top Left Middle fan       OK       3420    Spinning at normal speed
      Top Left Rear fan         OK       3420    Spinning at normal speed
      Top Right Front fan       OK       3450    Spinning at normal speed
      Top Right Middle fan      OK       3420    Spinning at normal speed
      Top Right Rear fan        OK       3420    Spinning at normal speed
      Bottom Left Front fan     OK       3420    Spinning at normal speed
      Bottom Left Middle fan    OK       3420    Spinning at normal speed
      Bottom Left Rear fan      OK       3390    Spinning at normal speed
      Bottom Right Front fan    OK       3420    Spinning at normal speed
      Bottom Right Middle fan   OK       3390    Spinning at normal speed
      Bottom Right Rear fan     OK       3390    Spinning at normal speed
      Rear Tray Top fan         OK       7050    Spinning at normal speed
      Rear Tray Second fan      OK       7050    Spinning at normal speed
      Rear Tray Third fan       OK       7050    Spinning at normal speed
      Rear Tray Fourth fan      OK       7050    Spinning at normal speed
      Rear Tray Fifth fan       OK       7050    Spinning at normal speed
      Rear Tray Sixth fan       OK       7050    Spinning at normal speed
      Rear Tray Seventh fan     OK       7050    Spinning at normal speed
      Rear Tray Bottom fan      OK       7050    Spinning at normal speed
```

show chassis fan (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fan
sfc0-re0:
--------------------------------------------------------------------------
      Item                      Status   RPM     Measurement
      Fan Tray 0 Fan 1          OK       4830    Spinning at normal speed
      Fan Tray 0 Fan 2          OK       4860    Spinning at normal speed
      Fan Tray 0 Fan 3          OK       4830    Spinning at normal speed
      Fan Tray 0 Fan 4          OK       4800    Spinning at normal speed
      Fan Tray 0 Fan 5          OK       4830    Spinning at normal speed
      Fan Tray 0 Fan 6          OK       4770    Spinning at normal speed
      Fan Tray 1 Fan 1          OK       4800    Spinning at normal speed
      Fan Tray 1 Fan 2          OK       4770    Spinning at normal speed
      Fan Tray 1 Fan 3          OK       4800    Spinning at normal speed
      Fan Tray 1 Fan 4          OK       4770    Spinning at normal speed
```

```
      Fan Tray 1 Fan 5          OK       4770    Spinning at normal speed
      Fan Tray 1 Fan 6          OK       4800    Spinning at normal speed
      Fan Tray 2 Fan 1          OK       4800    Spinning at normal speed
      Fan Tray 2 Fan 2          OK       4800    Spinning at normal speed
      Fan Tray 2 Fan 3          OK       4830    Spinning at normal speed
      Fan Tray 2 Fan 4          OK       4830    Spinning at normal speed
      Fan Tray 2 Fan 5          OK       4830    Spinning at normal speed
      Fan Tray 2 Fan 6          OK       4830    Spinning at normal speed
      Fan Tray 2 Fan 7          OK       4800    Spinning at normal speed
      Fan Tray 2 Fan 8          OK       4830    Spinning at normal speed
      Fan Tray 2 Fan 9          OK       4800    Spinning at normal speed
      Fan Tray 3 Fan 1          OK       4860    Spinning at normal speed
      Fan Tray 3 Fan 2          OK       4860    Spinning at normal speed
      Fan Tray 3 Fan 3          OK       4800    Spinning at normal speed
      Fan Tray 3 Fan 4          OK       4830    Spinning at normal speed
      Fan Tray 3 Fan 5          OK       4830    Spinning at normal speed
      Fan Tray 3 Fan 6          OK       4830    Spinning at normal speed
      Fan Tray 3 Fan 7          OK       4830    Spinning at normal speed
      Fan Tray 3 Fan 8          OK       4800    Spinning at normal speed
      Fan Tray 3 Fan 9          OK       4800    Spinning at normal speed
      Fan Tray 4 Fan 1          OK       4830    Spinning at normal speed
      Fan Tray 4 Fan 2          OK       4830    Spinning at normal speed
      Fan Tray 4 Fan 3          OK       4830    Spinning at normal speed
      Fan Tray 4 Fan 4          OK       4830    Spinning at normal speed
      Fan Tray 4 Fan 5          OK       4830    Spinning at normal speed
      Fan Tray 4 Fan 6          OK       4860    Spinning at normal speed
      Fan Tray 4 Fan 7          OK       4800    Spinning at normal speed
      Fan Tray 4 Fan 8          OK       4860    Spinning at normal speed
      Fan Tray 4 Fan 9          OK       4770    Spinning at normal speed
      Fan Tray 5 Fan 1          OK       4830    Spinning at normal speed
      Fan Tray 5 Fan 2          OK       4830    Spinning at normal speed
      Fan Tray 5 Fan 3          OK       4830    Spinning at normal speed
      Fan Tray 5 Fan 4          OK       4800    Spinning at normal speed
      Fan Tray 5 Fan 5          OK       4800    Spinning at normal speed
      Fan Tray 5 Fan 6          OK       4800    Spinning at normal speed
      Fan Tray 5 Fan 7          OK       4830    Spinning at normal speed
      Fan Tray 5 Fan 8          OK       4830    Spinning at normal speed
      Fan Tray 5 Fan 9          Check    2010

lcc0-re0:
--------------------------------------------------------------------------
      Item                      Status   RPM     Measurement
      Top Left Front fan        OK       3420    Spinning at normal speed
```

```
      Top Left Middle fan       OK       3390    Spinning at normal speed
      Top Left Rear fan         OK       3390    Spinning at normal speed
      Top Right Front fan       OK       3420    Spinning at normal speed
      Top Right Middle fan      OK       3420    Spinning at normal speed
      Top Right Rear fan        OK       3450    Spinning at normal speed
      Bottom Left Front fan     OK       3420    Spinning at normal speed
      Bottom Left Middle fan    OK       3390    Spinning at normal speed
      Bottom Left Rear fan      OK       3420    Spinning at normal speed
      Bottom Right Front fan    OK       3420    Spinning at normal speed
      Bottom Right Middle fan   OK       3390    Spinning at normal speed
      Bottom Right Rear fan     OK       3420    Spinning at normal speed
      Rear Tray fan 1 (Top)     OK       7740    Spinning at normal speed
      Rear Tray fan 2           OK       7740    Spinning at normal speed
      Rear Tray fan 3           OK       7740    Spinning at normal speed
      Rear Tray fan 4           OK       7740    Spinning at normal speed
      Rear Tray fan 5           OK       7740    Spinning at normal speed
      Rear Tray fan 6           OK       7740    Spinning at normal speed
      Rear Tray fan 7           OK       7740    Spinning at normal speed
      Rear Tray fan 8           OK       7740    Spinning at normal speed
      Rear Tray fan 9           OK       7740    Spinning at normal speed
      Rear Tray fan 10          OK       7740    Spinning at normal speed
      Rear Tray fan 11          OK       7740    Spinning at normal speed
      Rear Tray fan 12          OK       7740    Spinning at normal speed
      Rear Tray fan 13          OK       7740    Spinning at normal speed
      Rear Tray fan 14          OK       7740    Spinning at normal speed
      Rear Tray fan 15          OK       7740    Spinning at normal speed
      Rear Tray fan 16 (Bottom) OK       7740    Spinning at normal speed

lcc2-re0:
--------------------------------------------------------------------------
      Item                      Status   RPM     Measurement
      Top Left Front fan        OK       3420    Spinning at normal speed
      Top Left Middle fan       OK       3390    Spinning at normal speed
      Top Left Rear fan         OK       3420    Spinning at normal speed
      Top Right Front fan       OK       3420    Spinning at normal speed
      Top Right Middle fan      OK       3420    Spinning at normal speed
      Top Right Rear fan        OK       3450    Spinning at normal speed
      Bottom Left Front fan     OK       3420    Spinning at normal speed
      Bottom Left Middle fan    OK       3390    Spinning at normal speed
      Bottom Left Rear fan      OK       3420    Spinning at normal speed
      Bottom Right Front fan    OK       3420    Spinning at normal speed
      Bottom Right Middle fan   OK       3390    Spinning at normal speed
      Bottom Right Rear fan     OK       3420    Spinning at normal speed
```

```
      Rear Tray fan 1 (Top)     OK       7740    Spinning at normal speed
      Rear Tray fan 2           OK       7740    Spinning at normal speed
      Rear Tray fan 3           OK       7740    Spinning at normal speed
      Rear Tray fan 4           OK       7740    Spinning at normal speed
      Rear Tray fan 5           OK       7740    Spinning at normal speed
      Rear Tray fan 6           OK       7740    Spinning at normal speed
      Rear Tray fan 7           OK       7740    Spinning at normal speed
      Rear Tray fan 8           OK       7740    Spinning at normal speed
      Rear Tray fan 9           OK       7740    Spinning at normal speed
      Rear Tray fan 10          OK       7740    Spinning at normal speed
      Rear Tray fan 11          OK       7740    Spinning at normal speed
      Rear Tray fan 12          OK       7740    Spinning at normal speed
      Rear Tray fan 13          OK       7740    Spinning at normal speed
      Rear Tray fan 14          OK       7740    Spinning at normal speed
      Rear Tray fan 15          OK       7740    Spinning at normal speed
      Rear Tray fan 16 (Bottom) OK       7740    Spinning at normal speed
```

show chassis fan (PTX5000 Packet Transport Router)

```
user@host> show chassis fan
user@host> show chassis fan
      Item                      Status   % RPM     Measurement
      Fan Tray 0 Fan 1          OK       29%       2700 RPM
      Fan Tray 0 Fan 2          OK       29%       2700 RPM
      Fan Tray 0 Fan 3          OK       29%       2742 RPM
      Fan Tray 0 Fan 4          OK       29%       2700 RPM
      Fan Tray 0 Fan 5          OK       30%       2828 RPM
      Fan Tray 0 Fan 6          OK       30%       2828 RPM
      Fan Tray 0 Fan 7          OK       29%       2700 RPM
      Fan Tray 0 Fan 8          OK       30%       2785 RPM
      Fan Tray 0 Fan 9          OK       30%       2828 RPM
      Fan Tray 0 Fan 10         OK       30%       2828 RPM
      Fan Tray 0 Fan 11         OK       30%       2785 RPM
      Fan Tray 0 Fan 12         OK       30%       2828 RPM
      Fan Tray 0 Fan 13         OK       31%       2871 RPM
      Fan Tray 0 Fan 14         OK       30%       2828 RPM
      Fan Tray 1 Fan 1          OK       42%       3033 RPM
      Fan Tray 1 Fan 2          OK       42%       3066 RPM
      Fan Tray 1 Fan 3          OK       43%       3099 RPM
      Fan Tray 1 Fan 4          OK       43%       3166 RPM
      Fan Tray 1 Fan 5          OK       45%       3266 RPM
```

```
      Fan Tray 1 Fan 6          OK       43%       3133 RPM
      Fan Tray 2 Fan 1          OK       29%       2099 RPM
      Fan Tray 2 Fan 2          OK       30%       2199 RPM
      Fan Tray 2 Fan 3          OK       30%       2166 RPM
      Fan Tray 2 Fan 4          OK       33%       2399 RPM
      Fan Tray 2 Fan 5          OK       29%       2133 RPM
      Fan Tray 2 Fan 6          OK       32%       2366 RPM
```

show chassis fan (MX150)

```
user@host > show chassis fan
Item                      Status   RPM     Measurement
      FPC 0 Tray 0 Fan 0        OK       7419    Spinning at normal speed
      FPC 0 Tray 1 Fan 0        OK       7419    Spinning at normal speed
```

show chassis fan (MX104 Router)

```
user@host > show chassis fan
Item                      Status   RPM     Measurement
      Fan 1                     OK       5640    Spinning at normal speed
      Fan 2                     OK       5640    Spinning at normal speed
      Fan 3                     OK       5760    Spinning at normal speed
      Fan 4                     OK       5640    Spinning at normal speed
      Fan 5                     OK       5640    Spinning at normal speed
```

show chassis fan (MX2010 Router)

```
user@host > show chassis fan
Item                      Status   % RPM     Measurement
      Fan Tray 0 Fan 1          OK       37%       3360 RPM
      Fan Tray 0 Fan 2          OK       38%       3480 RPM
      Fan Tray 0 Fan 3          OK       37%       3360 RPM
      Fan Tray 0 Fan 4          OK       37%       3360 RPM
      Fan Tray 0 Fan 5          OK       38%       3480 RPM
      Fan Tray 0 Fan 6          OK       37%       3360 RPM
      Fan Tray 1 Fan 1          OK       38%       3480 RPM
      Fan Tray 1 Fan 2          OK       40%       3600 RPM
      Fan Tray 1 Fan 3          OK       38%       3480 RPM
```

```
      Fan Tray 1 Fan 4          OK       38%       3480 RPM
      Fan Tray 1 Fan 5          OK       38%       3480 RPM
      Fan Tray 1 Fan 6          OK       38%       3480 RPM
      Fan Tray 2 Fan 1          OK       38%       3480 RPM
      Fan Tray 2 Fan 2          OK       41%       3720 RPM
      Fan Tray 2 Fan 3          OK       38%       3480 RPM
      Fan Tray 2 Fan 4          OK       38%       3480 RPM
      Fan Tray 2 Fan 5          OK       38%       3480 RPM
      Fan Tray 2 Fan 6          OK       38%       3480 RPM
      Fan Tray 3 Fan 1          OK       38%       3480 RPM
      Fan Tray 3 Fan 2          OK       40%       3600 RPM
      Fan Tray 3 Fan 3          OK       40%       3600 RPM
      Fan Tray 3 Fan 4          OK       40%       3600 RPM
      Fan Tray 3 Fan 5          OK       40%       3600 RPM
      Fan Tray 3 Fan 6          OK       38%       3480 RPM
```

show chassis fan (ACX4000 Router)

```
user@host > show chassis fan
      Item                      Status   RPM     Measurement
      Fan 1                     OK       4140    Spinning at normal speed
      Fan 2                     OK       4200    Spinning at normal speed
```

show chassis fan (ACX5048 Router)

```
user@host > show chassis fan
      Item                      Status   RPM     Measurement
      FPC 0 Tray 0 Fan 0        OK       18305   Spinning at normal speed
      FPC 0 Tray 0 Fan 1        OK       15743   Spinning at normal speed
      FPC 0 Tray 1 Fan 0        OK       18305   Spinning at normal speed
      FPC 0 Tray 1 Fan 1        OK       15606   Spinning at normal speed
      FPC 0 Tray 2 Fan 0        OK       19014   Spinning at normal speed
      FPC 0 Tray 2 Fan 1        OK       16167   Spinning at normal speed
      FPC 0 Tray 3 Fan 0        OK       18947   Spinning at normal speed
      FPC 0 Tray 3 Fan 1        OK       16265   Spinning at normal speed
      FPC 0 Tray 4 Fan 0        OK       18120   Spinning at normal speed
      FPC 0 Tray 4 Fan 1        OK       15743   Spinning at normal speed
```

show chassis fan (QFX5100 Switch and OCX Series)

```
user@switch > show chassis fan
Item                      Status   RPM     Measurement
FPC 0 Tray 0 Fan 0        OK       6428    Spinning at normal speed
FPC 0 Tray 0 Fan 1        OK       5515    Spinning at normal speed
FPC 0 Tray 1 Fan 0        OK       6360    Spinning at normal speed
FPC 0 Tray 1 Fan 1        OK       5532    Spinning at normal speed
```

show chassis fan (EX9251 switches)

```
user@switch > show chassis fan

       Item                      Status   % RPM     Measurement
       Fan Tray 0 Fan 0          OK       40%       9600 RPM
       Fan Tray 0 Fan 1          OK       40%       8832 RPM
       Fan Tray 1 Fan 0          OK       40%       9728 RPM
       Fan Tray 1 Fan 1          OK       40%       9088 RPM
       Fan Tray 2                Absent
```

show chassis fan (EX9253 switches)

```
user@switch > show chassis fan

       Item                      Status   % RPM    Measurement
       Fan Tray 0 Fan 0          OK       40%       7552 RPM
       Fan Tray 0 Fan 1          OK       40%       6272 RPM
       Fan Tray 0 Fan 2          OK       40%       7552 RPM
       Fan Tray 0 Fan 3          OK       40%       6272 RPM
       Fan Tray 1 Fan 0          OK       40%       7552 RPM
       Fan Tray 1 Fan 1          OK       40%       6272 RPM
       Fan Tray 1 Fan 2          OK       40%       7552 RPM
       Fan Tray 1 Fan 3          OK       40%       6272 RPM
       Fan Tray 2 Fan 0          OK       40%       7552 RPM
       Fan Tray 2 Fan 1          OK       40%       6400 RPM
       Fan Tray 2 Fan 2          OK       40%       7552 RPM
       Fan Tray 2 Fan 3          OK       40%       6272 RPM
       Fan Tray 3 Fan 0          OK       40%       7552 RPM
       Fan Tray 3 Fan 1          OK       40%       6400 RPM
```

```
        Fan Tray 3 Fan 2          OK        40%        7552 RPM
        Fan Tray 3 Fan 3          OK        40%        6272 RPM
```

show chassis fan (Junos OS Evolved)

```
user@device> show chassis fan
      Item                      Status   % RPM     Measurement
      Fan Tray 1 Fan 1          Ok       48%       6597 RPM
      Fan Tray 1 Fan 2          Ok       49%       5649 RPM
      Fan Tray 2 Fan 1          Ok       49%       6687 RPM
      Fan Tray 2 Fan 2          Ok       49%       5649 RPM
      Fan Tray 3 Fan 1          Ok       49%       6642 RPM
      Fan Tray 3 Fan 2          Ok       49%       5649 RPM
```

Release Information

Command introduced in Junos OS Release 10.0 on MX Series 5G Universal Routing Platforms, M120 routers, and M320 routers, T320 routers, T640 routers, T1600 routers, TX Matrix Routers, and TX Matrix Plus routers.

`satellite` option introduced in Junos OS Release 14.2R3.

Command output introduced for Junos OS Evolved Release 19.1R1.

show chassis feb

IN THIS SECTION

- Syntax | 18279
- Description | 18279
- Options | 18279
- Required Privilege Level | 18279
- Output Fields | 18279
- Sample Output | 18281

Release Information | 18282

Syntax

```
show chassis feb
```

Description

(ACX1000, ACX7509 and QFX5700 devices only) Display Forwarding Engine Board (FEB) status information. Use this command to identify the memory allocated for each PFE (Packet Forwarding Engine).

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 274 on page 18280 lists the output fields for the `show chassis feb` command. Output fields are listed in the approximate order in which they appear.

Table 274: show chassis feb

Field Name	Field Description

State	State of the FEB:
	• **Offline**—FEB is powered down.
	• **Online**—FEB is operational and running.
	• **Check**—FEB is in alarmed state where the Switch Interface Board (SIB) plane is partially operational for the following reasons:
	  • FEB is not inserted properly.
	  • Two or more links between the FEB and Packet Forwarding Engine fail.

Temp (C) or Intake temperature	Temperature of the air passing by the FEB, in degrees Celsius or in both degrees Celsius and degrees Fahrenheit.
CPU Utilization (%)	Percentage of CPU being used: • **Total**—Total percentage of CPU being used by the FEB processor. • **Interrupt**—Of the total CPU being used by the FEB processor, the percentage being used for interrupts.
Memory DRAM (MB)	Total DRAM, in megabytes, available to the FEB processor.
Utilization (%)	Percentage of memory utilization: • **Heap**—Percentage of heap space (dynamic memory) being used by the FEB processor. If this number exceeds 80 percent, you might experience a software problem (memory leak). • **Buffer**—Percentage of buffer space being used by the FPC processor for buffering internal messages.
Exhaust A temperature	Temperature of the air flowing past Exhaust A.

Table 274: show chassis feb *(Continued)*

Field Name	Field Description
Exhaust B temperature	Temperature of the air flowing past Exhaust B.
Total DDR DRAM	Amount of double data rate dynamic random access memory (DDR DRAM) available to the FEB CPU.
Total RLDRAM	Amount of reduced latency dynamic random access memory (RLDRAM) available to the FEB CPU.
Start time (Detail output only)	Time when the Routing Engine detected that the FEB was running.
Uptime (Detail output only)	How long the Routing Engine has been connected to the FEB, and therefore, how long the Flexible PIC Concentrator (PIC) has been up and running.

Sample Output

show chassis feb detail (ACX1000 Universal Metro Router)

```
user@host> show chassis feb detail
FEB status:
Slot 0 information:
  State                                 Online
  Temperature                        46 degrees C / 114 degrees F
  CPU utilization                    15 percent
  Interrupt utilization               5 percent
  Heap utilization                   45 percent
  Buffer utilization                 37 percent
  Total CPU DRAM                    256 MB
  Start time:                           2012-06-05 19:51:53 PDT
  Uptime:                               19 minutes, 6 seconds
```

show chassis feb (ACX7509 Universal Metro Router)

```
user@host> show chassis feb
Slot 0 information:
  State                                   Online
  Start time:                             2022-06-13 21:30:00 PDT
  Uptime:                                 2 hours, 58 minutes, 40 seconds
Slot 1 information:
 State                                   Online Standby
 Start time:                             2022-06-13 21:30:00 PDT
 Uptime:                                 2 hours, 58 minutes, 40 seconds
```

show chassis feb (QFX5700 Switch)

```
user@host> show chassis feb
Slot 0 information:
  State                                   Online
  Start time:                             2024-06-04 05:41:18 UTC
  Uptime:                                 1 day, 3 hours, 42 minutes, 3 seconds
Slot 1 information:
  State                                   Offline Standby
  Start time:                             2024-06-04 05:41:01 UTC
  Uptime:                                 0 second
  Reason:                                 None

{master}
```

Release Information

Command introduced before Junos OS Release 7.4.

Command support added for ACX7509 platform in Junos OS Evolved Release 22.2. Detail option is not supported on ACX7509 platform.

RELATED DOCUMENTATION

Understanding Switching Control Board Redundancy

feb (Assigning a FEB to a Redundancy Group) | **6981**

request chassis feb | 14472
show chassis fabric feb | 18078
show chassis fpc-feb-connectivity | 18397
show chassis routing-engine | 18703

show chassis firmware

IN THIS SECTION

- Syntax | 18283
- Syntax (TX Matrix Routers) | 18284
- Syntax (TX Matrix Plus Routers) | 18284
- Syntax (MX Series Routers) | 18284
- Syntax (MX104, MX204, MX2010, MX2020, MX10003, and MX2008 Universal Routing Platforms) | 18284
- Syntax (QFX Series) | 18284
- Syntax (ACX5048 and ACX5096 Routers) | 18285
- Syntax (EX Series Switches) | 18285
- Description | 18285
- Options | 18285
- Required Privilege Level | 18287
- Output Fields | 18287
- Sample Output | 18288
- Release Information | 18304

Syntax

```
show chassis firmware
```

Syntax (TX Matrix Routers)

```
show chassis firmware
<lcc number | scc>
```

Syntax (TX Matrix Plus Routers)

```
show chassis firmware
<lcc number | sfc number>
```

Syntax (MX Series Routers)

```
show chassis firmware
<all-members>
<local>
<member member-id>
```

Syntax (MX104, MX204, MX2010, MX2020, MX10003, and MX2008 Universal Routing Platforms)

```
show chassis firmware
<satellite [slot-id slot-id |device-alias  alias-name]>
```

Syntax (QFX Series)

```
show chassis firmware
interconnect-device name
node-device name
```

Syntax (ACX5048 and ACX5096 Routers)

```
show chassis firmware
interconnect-device name
node-device name
```

Syntax (EX Series Switches)

```
show chassis firmware
<detail>
<satellite [slot-id slot-id |device-alias  alias-name]>
```

Description

On routers and switches, display the version levels of the firmware running on the System Control Board (SCB), Switching and Forwarding Module (SFM), System and Switch Board (SSB), Forwarding Engine Board (FEB), Flexible PIC Concentrators (FPCs), and Routing Engines. On a TX Matrix Plus router, display the version levels of the firmware running on the FPCs and the Switch Processor Mezzanine Board (SPMBs).

On EX2200, EX3200, EX4200, QFX Series, and OCX Series switches, display the version levels of the firmware running on the switch. On an EX8208 switch, display the version levels of the firmware running on the Switch Fabric and Routing Engine (SRE) modules and on the line cards (shown as FPCs). On an EX8216 switch, display the version levels of the firmware running on the Routing Engine (RE) modules and on the line cards (shown as FPCs).

Starting in Junos OS Release 23.1, on QFX10008 and QFX10016 switches, you can see the version of the primary firmware available on the power supply.

Options

none Display the version levels of the firmware running. For an EX4200 switch that is a member of a Virtual Chassis, display version levels for all members. For a TX Matrix router, display version levels for the firmware on the TX Matrix router and on all the T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, display version levels for the firmware on the TX Matrix Plus router and on all the routers connected to the TX Matrix Plus router.

all-members (MX Series routers only) (Optional) Display the version levels of the firmware running for all members of the Virtual Chassis configuration.

interconnect-device *name* (QFabric systems) (Optional) Display the version levels of the firmware running on the Interconnect device.

lcc *number* (TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display version levels for the firmware on a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the version levels for the firmware on a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Display the version levels of the firmware running for the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display the version levels of the firmware running for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device (QFabric systems only) (Optional) Display the version levels of the firmware running on the Node device.

satellite [slot-id *slot-id* | device-alias *alias-name*] (Junos Fusion only) (Optional) Display version levels of the firmware running for the specified satellite device or devices in a Junos Fusion, or for all satellite devices if no satellite devices are specified.

scc (TX Matrix router only) (Optional) Display version levels for the firmware on the TX Matrix router (switch-card chassis).

sfc *number* (TX Matrix Plus router only) (Optional) Display version levels for the firmware on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

detail (EX3200, EX3300, EX4200, and EX4500 standalone and Virtual Chassis member switches only) (Optional) Display version levels of the firmware running on the switch for its programmable hardware components.

Required Privilege Level

view

Output Fields

Table 275 on page 18287 lists the output fields for the show chassis firmware command. Output fields are listed in the approximate order in which they appear.

Table 275: show chassis firmware Output Fields

Field Name	Field Description
`Part`	(MX Series, MX2010, MX2020, and MX2008 routers) Chassis part name.
`Type`	(MX Series, MX2010, MX2020, and MX2008 routers) Type of firmware: On routers: ROM or O/S. On switches: uboot or loader.
`Version`	(MX Series, MX2010, MX2020, and MX2008 routers) Version of firmware running on the chassis part.
`FPC`	(*detail* option only) Number of FPC. For a standalone switch, the value is 0. For a Virtual Chassis configuration, value in the range of 0-9; refers to the member ID assigned to the switch.
`AFEB`	(MX104 routers) Version of the compact Forwarding Engine Board.
`Boot`	(*detail* option only) Version of the SYSPLD.
`PoE`	(*detail* option only) Version of the PoE firmware.

Table 275: show chassis firmware Output Fields *(Continued)*

Field Name	Field Description
PFE-*<number>*	(*detail* option only) Version of the Packet Forwarding Engine used in the switch.
PHY-	(*detail* option only) Version of the physical layer device (PHY) used in the switch.
microcode	(*detail* option only) Microcode of the physical layer devices (PHY) used in the switch.
uboot	(*detail* option only) Version of the u-boot used in the switch.
loader	(*detail* option only) Version of the loader used in the switch.

Sample Output

show chassis firmware (M10 Router)

```
user@host> show chassis firmware
Part                     Type       Version
Forwarding engine board  ROM        Juniper ROM Monitor Version 4.1b2
                         O/S        Version 4.1I1 by usera on 2000-04-24 11:27
```

show chassis firmware (M20 Router)

```
user@host> show chassis firmware
Part                     Type       Version
System switch board      ROM        Juniper ROM Monitor Version 3.4b26
                         O/S        Version 3.4I16 by userc on 2000-02-29 2
FPC 1                    ROM        Juniper ROM Monitor Version 3.0b1
                         O/S        Version 3.4I4 by userc on 2000-02-25 21
FPC 2                    ROM        Juniper ROM Monitor Version 3.0b1
                         O/S        Version 3.4I4 by userc on 2000-02-25 21
```

show chassis firmware (M40 Router)

```
user@host> show chassis firmware
Part                     Type       Version
System control board     ROM        Juniper ROM Monitor Version 2.0i126Copyri
                         O/S        Version 2.0i1 by root on Thu Jul 23 00:51
FPC 5                    ROM        Juniper ROM Monitor Version 2.0i49Copyrig
                         O/S        Version 2.0i1 by root on Thu Jul 23 00:59
```

show chassis firmware (M160 Router)

```
user@host> show chassis firmware
Part                     Type       Version
SFM 0                    ROM        Juniper ROM Monitor Version 4.0b2
                         O/S        Version 4.0I1 by usera on 2000-02-29 11:50
SFM 1                    ROM        Juniper ROM Monitor Version 4.0b2
                         O/S        Version 4.0I1 by usera on 2000-02-29 11:50
FPC 0                    ROM        Juniper ROM Monitor Version 4.0b2
                         O/S        Version 4.0I1 by usera on 2000-02-29 11:56
FPC 1                    ROM        Juniper ROM Monitor Version 4.0b2
                         O/S        Version 4.0I1 by usera on 2000-02-29 11:56
FPC 2                    ROM        Juniper ROM Monitor Version 4.0b3
                         O/S        Version 4.0I1 by usera on 2000-02-29 11:56
```

show chassis firmware (MX150)

```
user@host > show chassis firmware
Part                     Type       Version
FPC                      ROM        PC Bios
                         O/S        Version 17.2I20170220_0929_rohitn by rohitn on 2017-02-20
09:38:59 UTC
```

show chassis firmware (MX104 Router)

```
user@host > show chassis firmware
Part                     Type       Version
FPC 0                    ROM        Juniper ROM Monitor Version 13.1b24
```

```
                          O/S        Version 13.2-20130514.1 by userb on 2013-
FPC 1                     ROM        Juniper ROM Monitor Version 13.1b24
                          O/S        Version 13.2-20130514.1 by userb on 2013-
FPC 2                     ROM        Juniper ROM Monitor Version 13.1b24
                          O/S        Version 13.2-20130514.1 by userb on 2013-
AFEB                      ROM        Juniper ROM Monitor Version 13.1b24
                          O/S        Version 13.2-20130514.1 by userb on 2013-
```

show chassis firmware (MX240 Router)

```
user@host> show chassis firmware
Part                      Type       Version
FPC 1                     ROM        Juniper ROM Monitor Version 8.3b1
                          O/S        Version 9.0-20080103.0 by userb on 2008-0
FPC 2                     ROM        Juniper ROM Monitor Version 8.3b1
                          O/S        Version 9.0-20080103.0 by userb on 2008-0
```

show chassis firmware (MX480 Router)

```
user@host> show chassis firmware
Part                      Type       Version
FPC 1                     ROM        Juniper ROM Monitor Version 8.3b1
                          O/S        Version 9.0-20070916.3 by userb on 2007-0
```

show chassis firmware (MX960 Router)

```
user@host> show chassis firmware
Part                      Type       Version
FPC 4                     ROM        Juniper ROM Monitor Version 8.0b8
                          O/S        Version 8.2I59 by user3 on 2006-10-31 19:22
FPC 7                     ROM        Juniper ROM Monitor Version 8.2b1
                          O/S        Version 8.2-20061026.1 by userb on 2006-1
```

show chassis firmware (MX2020 Router)

```
user@host> show chassis firmware
Part                    Type       Version
FPC 0                   ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 1                   ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 2                   ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 3                   ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 4                   ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 5                   ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 6                   ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 7                   ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 8                   ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 9                   ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 10                  ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 11                  ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 12                  ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 13                  ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 14                  ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 15                  ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 16                  ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 17                  ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
FPC 18                  ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
```

```
FPC 19                  ROM        Juniper ROM Monitor Version 10.0b39
                        O/S        Version 12.3-20130415.0 by userb on 2013-
SPMB 0                  ROM        Juniper ROM Monitor Version 12.1b1
                        O/S        Version 12.3-20130415.0 by userb on 2013-
SPMB 1                  ROM        Juniper ROM Monitor Version 12.1b1
                        O/S        Version 12.3-20130415.0 by userb on 2013-
```

show chassis firmware (MX204 Router)

```
user@host> show chassis firmware

Part                    Type       Version
RE 0                    PRI BIOS   CBEP_P_SUM1_00.11.01
                        RE-FPGA    300
FPC                     ROM        PC Bios
                        O/S        Version 17.4I20171105_0609_aahluwalia by aahluwalia on
2017-11-05 06:09:28 UTC
```

show chassis firmware (MX10008 Router)

```
user@host> show chassis firmware
Part                    Type       Version
RE 0                    PRI BIOS   CBEP_P_VAL0_00.14.1
                        FPGA       264.0
                        RE-FPGA    41.0
                        RE-SSD1    SF-SBR12050
                        RE-SSD2    SF-SBR12050
                        i40e-NVM   6.01
RE 1                    PRI BIOS   CBEP_P_VAL0_00.13.01
                        FPGA       261.0
                        RE-FPGA    41.0
                        RE-SSD1    SF-SBR12034
                        RE-SSD2    SF-SBR12034
                        i40e-NVM   5.02
FPC 0                   ROM        PC Bios
                        O/S        Version 18.4-20180716_dev_common.0 by builder on 2018-07-16
00:43:35 UTC
                        ROM Monitor 0 9.14.0
                        PCIE Sw(0) 1.0.0
                        MPCS(0)    0.2.0
```

```
                        I2CS CPLD  0.4.0
                        BOOT CPLD  0.4.0
FPC 2                   ROM        PC Bios
                        O/S        Version 18.4-20180716_dev_common.0 by builder on 2018-07-16
00:43:35 UTC
                        ROM Monitor 0 9.14.0
                        PCIE Sw(0) 1.0.0
                        MPCS(0)    0.2.0
                        I2CS CPLD  0.4.0
                        BOOT CPLD  0.4.0
FPC 3                   ROM        PC Bios
                        O/S        Version 18.4-20180716_dev_common.0 by builder on 2018-07-16
00:43:35 UTC
                        ROM Monitor 0 9.14.0
                        PCIE Sw(0) 1.0.0
                        MPCS(0)    0.4.0
                        I2CS CPLD  0.8.0
                        BOOT CPLD  0.8.0
FPM                     FPGA       1.9
FTC 0                   FPGA       2.0
FTC 1                   FPGA       2.0
SFB 0                   FPGA       3.0
SFB 1                   FPGA       3.0
SFB 2                   FPGA       3.0
SFB 3                   FPGA       3.0
SFB 4                   FPGA       3.0
SFB 5                   FPGA       3.0
```

show chassis firmware (MX240, MX480, MX960 Router with Application Services Modular Line Card)

```
user@host> show chassis firmware

Part                    Type       Version
FPC 1                   ROM        Juniper ROM Monitor Version 12.1b1
                        O/S        Version 12.2I21 by user1 on 2012-06-19 17:
```

show chassis firmware (EX4200 Switch)

```
user@switch> show chassis firmware

Part                    Type       Version
FPC 0                   uboot      U-Boot 1.1.6 (Feb  6 2008 - 11:27:42)
                        loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.1
FPC 1                   uboot      U-Boot 1.1.6 (Feb  6 2008 - 11:27:42)
                        loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.1
FPC 2                   uboot      U-Boot 1.1.6 (Feb  6 2008 - 11:27:42)
                        loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.1
```

show chassis firmware (EX8200 Switch)

```
user@switch> show chassis firmware

Part                      Type       Version
FPC 0                     U-Boot     U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0
                          loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 3                     U-Boot     U-Boot 1.1.6 (Dec  4 2009 - 13:17:34) 3.1.0
                          loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 5                     U-Boot     U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0
                          loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 7                     U-Boot     U-Boot 1.1.6 (Feb  6 2009 - 05:31:46) 2.4.0
                          loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.2
Routing Engine 0          U-Boot     U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0
                                     FreeBSD/PowerPC U-Boot bootstrap loader 2.2
                          loader
Routing Engine 1          U-Boot     U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0
                          loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.2
```

show chassis firmware (EX9200 Switch)

```
user@switch> show chassis firmware

 Part                      Type       Version
 FPC 2                     ROM        Juniper ROM Monitor Version 11.4b2
                           O/S        Version 14.1I20140312_0741 by userd o
```

```
FPC 3                     ROM        Juniper ROM Monitor Version 10.4b1
                          O/S        Version 14.1I20140312_0741 by userd o
```

show chassis firmware (EX9251 Switch)

```
user@switch> show chassis firmware

Part                      Type       Version
RE 0                      PRI BIOS   CBEP_P_SUM1_00.11.01
                          RE-FPGA    301
FPC                       ROM        PC Bios
                          O/S        Version 18.1R1.4 by builder on 2018-03-06 00:31:54 UTC
```

show chassis firmware (EX9253 Switch)

```
user@switch> show chassis firmware

Part                      Type       Version
RE 0                      PRI BIOS   CBEP_P_SUM1_00.11.01
                          RE-FPGA    402
RE 1                      PRI BIOS   CBEP_P_SUM1_00.11.01
                          RE-FPGA    402
FPC 0                     ROM        PC Bios
                          O/S        Version 18.2-20180129_dev_common.1 by builder on 2018-01-29
13:35:11 UTC
FPC 1                     ROM        PC Bios
                          O/S        Version 18.2-20180129_dev_common.1 by builder on 2018-01-29
13:35:11 UTC
```

show chassis firmware lcc (TX Matrix Router)

```
user@host> show chassis firmware lcc 0
lcc0-re0:
--------------------------------------------------------------------------
Part                      Type       Version
FPC 1                     ROM        Juniper ROM Monitor Version 6.4b18
                          O/S        Version 7.0-20040804.0 by userb on 2004-0
FPC 2                     ROM        Juniper ROM Monitor Version 6.4b20
```

```
                          O/S        Version 7.0-20040804.0 by userb on 2004-0
SPMB 0                    ROM        Juniper ROM Monitor Version 6.4b18
                          O/S        Version 7.0-20040804.0 by userb on 2004-0
```

show chassis firmware scc (TX Matrix Router)

```
user@host> show chassis firmware scc

scc-re0:
--------------------------------------------------------------------------
Part                      Type       Version
SPMB 0                    ROM        Juniper ROM Monitor Version 6.4b18
                          O/S        Version 7.0-20040804.0 by userb on 2004-0
```

show chassis firmware (TX Matrix Plus Router)

```
user@host> show chassis firmware
sfc0-re0:
--------------------------------------------------------------------------
Part                      Type       Version
Global FPC 4
Global FPC 6
Global FPC 7
Global FPC 12
Global FPC 14
Global FPC 15
Global FPC 20
Global FPC 21
Global FPC 22
Global FPC 23
Global FPC 24
Global FPC 25
Global FPC 26
Global FPC 28
Global FPC 29
Global FPC 31
SPMB 0                    ROM        Juniper ROM Monitor Version 9.5b1
                          O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 1                    ROM        Juniper ROM Monitor Version 9.5b1
                          O/S        Version 9.6-20090507.0 by userb on 2009-0
```

```
lcc0-re1:
--------------------------------------------------------------------------
Part                     Type       Version
FPC 4                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 6                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 7                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 0                   ROM        Juniper ROM Monitor Version 9.5b1
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 1                   ROM        Juniper ROM Monitor Version 9.5b1
                         O/S        Version 9.6-20090507.0 by userb on 2009-0

lcc1-re1:
--------------------------------------------------------------------------
Part                     Type       Version
FPC 4                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 6                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 7                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 0                   ROM        Juniper ROM Monitor Version 9.5b1
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 1                   ROM        Juniper ROM Monitor Version 9.5b1
                         O/S        Version 9.6-20090507.0 by userb on 2009-0

lcc2-re1:
--------------------------------------------------------------------------
Part                     Type       Version
FPC 4                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 5                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 6                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 7                    ROM        Juniper ROM Monitor Version 7.5b4
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 0                   ROM        Juniper ROM Monitor Version 9.5b1
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 1                   ROM        Juniper ROM Monitor Version 9.5b1
```

```
                         O/S        Version 9.6-20090507.0 by userb on 2009-0

lcc3-re1:
--------------------------------------------------------------------------
Part                     Type       Version
FPC 0                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 1                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 2                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 4                    ROM        Juniper ROM Monitor Version 7.5b4
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 5                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 7                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 0                   ROM        Juniper ROM Monitor Version 9.5b1
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 1                   ROM        Juniper ROM Monitor Version 9.5b1
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
```

show chassis firmware lcc (TX Matrix Plus Router)

```
user@host> show chassis firmware lcc 0
lcc0-re1:
--------------------------------------------------------------------------
Part                     Type       Version
FPC 4                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 6                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
FPC 7                    ROM        Juniper ROM Monitor Version 9.0b2
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 0                   ROM        Juniper ROM Monitor Version 9.5b1
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 1                   ROM        Juniper ROM Monitor Version 9.5b1
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
```

show chassis firmware sfc (TX Matrix Plus Router)

```
user@host> show chassis firmware sfc 0
sfc0-re0:
--------------------------------------------------------------------------
Part                     Type       Version
Global FPC 4
Global FPC 6
Global FPC 7
Global FPC 12
Global FPC 14
Global FPC 15
Global FPC 20
Global FPC 21
Global FPC 22
Global FPC 23
Global FPC 24
Global FPC 25
Global FPC 26
Global FPC 28
Global FPC 29
Global FPC 31
SPMB 0                   ROM        Juniper ROM Monitor Version 9.5b1
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
SPMB 1                   ROM        Juniper ROM Monitor Version 9.5b1
                         O/S        Version 9.6-20090507.0 by userb on 2009-0
```

show chassis firmware (QFX Series and OCX Series)

```
user@switch> show chassis firmware
Part                     Type       Version
FPC 0
Routing Engine 0         U-Boot     U-Boot 1.1.6 (Sep 15 2010 - 02:11:11) 1.0.5
                         loader     FreeBSD/MIPS U-Boot bootstrap loader 0.1
```

show chassis firmware (PTX1000 Packet Transport Routers)

```
user@host> show chassis firmware
Part                     Type       Version
```

```
FPC 0                    U-Boot      ***
                         loader           FreeBSD/i386 bootstrap loader 1.2
                         BIOS       V0018.2U
                         EC FPGA    2.0
                         MAIN_CPLD  1.f
                         MEZZ_CPLD  1.f
                         RE FPGA    2.3
```

show chassis firmware (PTX10008 Routers)

```
user@host> show chassis firmware
Part                     Type       Version
RE 0                     PRI BIOS   QFXS_SFP_00.31_01.01
                         GDN BIOS   QFXS_SFP_00.31_01.01
                         FPGA       2.4
                         RE-FPGA    3.2
RE 1                     PRI BIOS   QFXS_SFP_00.31_01.01
                         GDN BIOS   QFXS_SFP_00.31_01.01
                         FPGA       2.3
                         RE-FPGA    3.2
FPC 0                    U-Boot     Bank A: U-Boot 2011.12-gfbea47a (Feb 26 2016 - 22:56:52)
                         CTRL FPGA  4.1
                         PORT FPGA  2.0
FPC 5                    U-Boot     Bank A: U-Boot 2011.12-gfbea47a (Feb 26 2016 - 22:56:52)
                         CTRL FPGA  3.1
                         PORT FPGA  2.0
FPC 6                    U-Boot     Bank B: U-Boot 2011.12-gfbea47a (Feb 26 2016 - 22:56:52)
                         CTRL FPGA  3.1
                         PORT FPGA  2.0
FPM                      FPGA       1.9
FTC 0                    FPGA       2.0
FTC 1                    FPGA       2.0
SIB 0                    FPGA       3.0
SIB 1                    FPGA       3.0
```

show chassis firmware interconnect-device (QFabric System)

```
user@switch> show chassis firmware interconnect-device
interconnect1
Part                     Type       Version
```

```
Routing Engine 0         U-Boot     U-Boot 1.1.6 (May 10 2011 - 04:52:59) 1.1.1
                         loader     FreeBSD/MIPS U-Boot bootstrap loader 0.1
Routing Engine 1         U-Boot     U-Boot 1.1.6 (May 10 2011 - 04:52:59) 1.1.1
                         loader     FreeBSD/MIPS U-Boot bootstrap loader 0.1
```

show chassis firmware (ACX2000 Universal Metro Router)

```
user@switch> show chassis firmware
Part                     Type       Version
FPC                      O/S        Version 12.2I13 by user2 on 2012-05-29 06:
FEB                      O/S        Version 12.2I13 by user2 on 2012-05-29 06:
```

show chassis firmware detail (EX3300 Switch)

```
user@switch> show chassis firmware detail
FPC 0
   Boot SYSPLD             3
   PoE firmware            4.1.6
   PFE-0                   3
   PFE-1                   3
   PHY
      microcode            0x514
   Boot Firmware
      uboot                U-Boot 1.1.6 (Aug 21 2011 - 01:45:26)   1.0.0
      loader               FreeBSD/arm U-Boot loader 1.0
```

show chassis firmware (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis firmware
Part                     Type       Version
FPC 1                    ROM        Juniper ROM Monitor Version 12.1b1
                         O/S        Version 12.2I21 by user1 on 2012-06-19 17:
```

show chassis firmware (ACX5048 Router)

```
user@host> show chassis firmware
Part                      Type       Version
FPC                       loader           FreeBSD/i386 bootstrap loader 1.2
                          BIOS       V0018.7
                          TMC FPGA   6.d8
                          PIC0 CPLD0 7.b
                          PIC0 CPLD1 7.b
                          PIC0 CPLD2 7.b
                          PIC0 CPLD3 7.b
                          PIC0 CPLD4 7.b
                          PIC0 CPLD5 7.b
                          PIC0 CPLD6 6.a
                          MRE        17.9
                          Power CPLD 3.a
```

show chassis firmware (ACX5096 Router)

```
user@host> show chassis firmware
Part                      Type       Version
FPC                       loader           FreeBSD/i386 bootstrap loader 1.2
                          BIOS       V0018.7
                          TMC FPGA   3000001.5
                          PIC0 CPLD0 7.b
                          PIC0 CPLD1 7.b
                          PIC0 CPLD2 7.b
                          PIC0 CPLD3 7.b
                          PIC0 CPLD4 7.b
                          PIC0 CPLD5 7.b
                          PIC0 CPLD6 c6.a
                          PIC0 CPLD7 -NA-
                          PIC0 CPLD8 7.b
                          PIC0 CPLD9 7.b
                          PIC0 CPLD10 7.b
                          PIC0 CPLD11 7.b
                          PIC0 CPLD12 7.b
                          PIC0 CPLD13 7.b
                          PIC0 CPLD14 c6.a
```

```
                        MRE        7.5
                        Power CPLD 4.1
```

show chassis firmware (ACX500 Router)

```
user@host> show chassis firmware

Part                    Type       Version
FPC                     O/S        Version 15.2-20150815_dev_rbu_1_16q1.0 by userb on
2015-08-15 04:18:02 UTC
FEB                     O/S        Version 15.2-20150815_dev_rbu_1_16q1.0 by userb on
2015-08-15 04:18:02 UTC
```

show chassis firmware (ACX7100 Router)

```
user@host> show chassis firmware

Part             Type              Tag Current               Available             Status
                                       version               version
PSM 1            Psm MCU AC AFO    0   21.41                 21.43
OK
RE 0             BIOS ROM          0   13.01                 13.01
OK
RE 0             XMC FPGA          8   0x002A                0x002A
OK
RE 0             RE FPGA           2   0x0008                0x0008
OK
RE 0             SSD Primary       15  A20722J               A20722J
OK
RE 0             SSD Secondary     16  A20722J               A20722J
OK
RE 0             I210 NIC          12  3.25                  3.25
OK
RE 0             IDT PLL           17  4.8.15                4.8.15
OK
CB 0             Optics CPLD 0     1   0x010C                0x010C
OK
CB 0             Optics CPLD 1     2   0x010C                0x010C
OK
CB 0             Fan CPLD          3   0xB3                  0xB3
```

```
OK
FPC 0            BCM Jericho 2     1   2.50020              2.50020
OK
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced in Junos OS Release 9.6 for the TX Matrix Plus router.

Command introduced for EX8200 switches in Junos OS Release 10.2 for EX Series switches.

`satellite` option introduced in Junos OS Release 14.2R3.

show chassis forwarding

IN THIS SECTION

- Syntax | 18304
- Description | 18304
- Options | 18305
- Required Privilege Level | 18305
- Output Fields | 18305
- Sample Output | 18306
- Release Information | 18306

Syntax

```
show chassis forwarding
```

Description

Display status of the forwarding process (fwdd). This command is supported on Branch SRX Series Firewalls.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 276 on page 18305 lists the output fields for the `show chassis forwarding` command. Output fields are listed in the approximate order in which they appear.

Table 276: show chassis forwarding Output Fields

Field Name	Field Description

FWWD status	Forwarding status:
	• **State**:
	  • **Online**—FWDD is operational and running.
	  • **Offline**—FWDD is not running.
	• **Microkernel CPU utilization**—Percentage of microkernel CPU being used by the forwarding process.
	• **Real-time threads CPU utilization**—Percentage of CPU being used by the forwarding process.
	• **Heap utilization**—Percentage of heap space (dynamic memory) being used by the forwarding process. If this number exceeds 80 percent, there may be a software problem (memory leak).
	• **Buffer utilization**—Percentage of buffer space being used by the forwarding process for buffering internal messages.
	• **Uptime**—How long the forwarding process has been up and running.

Sample Output

show chassis forwarding

```
user@host> show chassis forwarding
FWDD status:
  State                                 Online
  Microkernel CPU utilization        10 percent
  Real-time threads CPU utilization   4 percent
  Heap utilization                   26 percent
  Buffer utilization                  0 percent
  Uptime:                              1 day, 1 hour, 30 minutes, 11 seconds
```

Release Information

Current—Command introduced before Junos OS Release 7.4.

Now—Command introduced in Junos OS Release 7.4. Support for Branch SRX Series added in Junos OS Release 10.1

show chassis forwarding-options

IN THIS SECTION

- Syntax | 18307
- Description | 18307
- Options | 18307
- Required Privilege Level | 18307
- Output Fields | 18307
- Sample Output | 18309
- Release Information | 18311

Syntax

```
show chassis forwarding-options
```

Description

Display the configuration for the Unified Forwarding Table.

Options

There are no options for this command.

> **NOTE**: Starting in Junos OS Releases 17.3R2, for QFX5200 Virtual Chassis, information about memory banks are displayed only for the Primary, not for the other members. Values remain the same across all members. All configuration changes for the Unified Forwarding Table are made through the Primary.

Required Privilege Level

view

Output Fields

Table 277 on page 18308 lists the output fields for the `show chassis forwarding-options` command. Output fields are listed in the approximate order in which they appear.

Table 277: show chassis forwarding-options Output Fields

Field Name	Field Description
profile name	Name of profile configured: • `custom-profile` (QFX5200 only) • `l2-profile-one` • `l2-profile-three` (default) • `l2-profile-two` • `l3-profile` • `lpm-profile` • `vxlan-gbp-profile` (EX4100, EX4400, EX4650, and QFX5120-32C and QFX5120-48Y switches) • `vxlan-gbp-l2-profile` (EX4400, EX4650, and QFX5120-32C and QFX5120-48Y switches) • `vxlan-gbp-l3-profile` (EX4400, EX4650, and QFX5120-32C and QFX5120-48Y switches)
MAC	Maximum amount of memory allocated for Layer 2 entries.
L3-host	Maximum amount of memory allocated for Layer 3 host entries.
LPM	Maximum amount of memory allocated for longest match prefix (LPM) entries.
num-65-127-prefix	Maximum amount of memory allocated in LPM table for IP prefixes with lengths in the range /65 through /127.
Total scale(K)	(QFX5200 only) Maximum amount of memory allocated for each address type. This amount includes the amount configured plus the amount allocated through the dedicated hash table.
Bank details for various types of entries	(QFX5200 only) Maximum amount of memory configured by address type for each of the four shared memory banks and the dedicated hash table.

Table 277: show chassis forwarding-options Output Fields *(Continued)*

Field Name	Field Description
Entry type	(QFX5200 only) Type of forwarding-table entry: `L2(mac)`; `L3 (unicast and multicast)`; `Exact Match`; and `Longest Prefix Match (lpm)`
Dedicated bank size(K)	(QFX5200 only) Maximum amount of memory allocated for each address type in the dedicated hash table.
Shared bank size(K)	(QFX5200 only) Default Maximum amount of memory allocated for each address type in the shared memory banks.

Sample Output

show chassis forwarding-options (vxlan-gbp-l2-profile on EX4400 Series switch)

```
user@host> show chasiss forwarding-options
UFT Configuration:
vxlan-gbp-l2-profile
num-65-127-prefix = 1K
--------------Bank details for various types of entries------------------
Entry type                        Dedicated Bank Size(K)     Shared Bank Size(K)
L2(mac)                           16                         16 * num shared banks
L3(unicast & multicast)           16                         16 * num shared banks
Exact Match                       0                          8 * num shared banks
```

show chassis forwarding-options (l2-profile-three)

```
user@host> show chassis forwarding-options
UFT Configuration:
l2-profile-three. (MAC: 160K L3-host: 144K LPM: 16K) (default)
num-65-127-prefix = none

{master:0}
```

show chassis forwarding-options (custom-profile on QFX5200 Series switch)

```
user@host> show chasiss forwarding-options
UFT Configuration:
custom-profile
Configured custom scale:
Entry type                 Total scale(K)
L2(mac)                         8
L3 (unicast & multicast)        72
Exact Match                     0
Longest Prefix Match (lpm)      80
num-65-127-prefix = 1K
-------------Bank details for various types of entries------------
Entry type                   Dedicated Bank Size(K)      Shared Bank Size(K)
L2 (mac)                     8                          32 * num shared banks
L3 (unicast & multicast      8                          32 * num shared banks
Exact match                  0                          16 * num shared banks
Longest Prefix match(lpm)    16                         32 * num shared banks
```

show chassis forwarding-options (QFX5200 Virtual Chassis)

```
user@host> show chassis forwarding-options
localre:
-
UFT Configuration:
l2-profile-three.(default)
num-65-127-prefix = 1K
-Bank details for various types of entries-
Entry type                        Dedicated Bank Size(K)     Shared Bank Size(K)
L2(mac)                           8                          32 * num shared banks
L3(unicast & multicast)           8                          32 * num shared banks
Exact Match                       0                          16 * num shared banks
Longest Prefix Match(lpm)         16                         32 * num shared banks

fpc1:
-
UFT Configurtion:
l2-profile-three.(default)
num-65-127-prefix = 1K
```

Release Information

Command introduced in Junos OS Release 13.2

Support added to QFX5200 switches in Junos OS Release 15.1X53-D30

RELATED DOCUMENTATION

Configuring the Unified Forwarding Table on Switches

Example: Configuring a Unified Forwarding Table Custom Profile

show chassis fpc

IN THIS SECTION

- Syntax | **18312**
- Syntax (EX Series Switches, PTX Series Routers, ACX Universal Series Routers, ACX500, ACX7509, TX Matrix and TX Matrix Plus Routers) | **18312**
- Syntax (T4000 Routers) | **18312**
- Syntax (MX Series Routers, EX Series, and PTX Series switches) | **18313**
- Syntax (PTX10008 and PTX10001-36MR) | **18313**
- Syntax (MX104, MX204, MX2010, MX2020, MX10003, and MX2008 Universal Routing Platforms) | **18313**
- Syntax (MX304) | **18313**
- Syntax (MX10008 Universal Routing Platforms) | **18313**
- Syntax (MX10004) | **18314**
- Syntax (QFX Series) | **18314**
- Syntax (OCX Series) | **18314**
- Syntax (ACX500 Routers) | **18314**
- Syntax (SRX Series) | **18315**
- Syntax (Junos OS Evolved) | **18315**
- Description | **18315**
- Options | **18315**
- Required Privilege Level | **18323**

- Output Fields | 18323
- Sample Output (Junos OS) | 18330
- Sample Output (Junos OS Evolved) | 18374
- Sample Output (Node Slicing) | 18377
- show chassis fpc (MX10004 and MX10008 Devices) | 18379
- Sample Output | 18381
- Sample Output | 18381
- Sample Output | 18385
- Usage | 18386
- Release Information | 18386

Syntax

```
show chassis fpc
<detail <slot>> | <pic-status <slot>> | <pfe-instance <pfe-id>>
all
```

Syntax (EX Series Switches, PTX Series Routers, ACX Universal Series Routers, ACX500, ACX7509, TX Matrix and TX Matrix Plus Routers)

```
show chassis fpc
<detail <fpc-slot>> | <pic-status <fpc-slot>>
<fpc-slot> | <pfe <pfe-id>>
```

Syntax (T4000 Routers)

```
show chassis fpc
<detail <fpc-slot>>
<pic-status <fpc-slot>>
```

Syntax (MX Series Routers, EX Series, and PTX Series switches)

```
show chassis fpc
< <slot>> | <pic-status <slot>>
<all-members>
<local>
<member member-id>
```

Syntax (PTX10008 and PTX10001-36MR)

```
show chassis fpc
<slot> | <pic-status <slot>> | <pfe-instance <pfe-id>>
all
```

Syntax (MX104, MX204, MX2010, MX2020, MX10003, and MX2008 Universal Routing Platforms)

```
show chassis fpc
<slot> detail | <detail <slot>> | <pic-status <slot>>
<fpc-slot>
```

Syntax (MX304)

```
show chassis fpc
<fpc-slot>
<detail>
<errors>
<global-pfe-id>
mic-status <fpc-slot>
pic-status <fpc-slot>
```

Syntax (MX10008 Universal Routing Platforms)

```
show chassis fpc
<detail>
```

```
<errors>
<fpc-slot>
pic-status <fpc-slot>
```

Syntax (MX10004)

```
show chassis fpc
<fpc-slot>
<detail>
<errors>
<global-pfe-id>
pic-status <fpc-slot>
```

Syntax (QFX Series)

```
show chassis fpc
<detail>
<interconnect-device name <fpc-slot fpc-slot>>
<node-device name>
```

Syntax (OCX Series)

```
show chassis fpc
<detail>
```

Syntax (ACX500 Routers)

```
show chassis fpc
<fpc-slot>
detail <fpc-slot>
pic-status <fpc-slot>
```

Syntax (SRX Series)

```
show chassis fpc
<fpc-slot>
detail <fpc-slot> | <node (node-id | local |primary)>>
pic-status <fpc-slot | <node (node-id | local |primary)>>
```

Syntax (Junos OS Evolved)

```
show chassis fpc <fpc-slot>
<detail | pic-status | <pfe-instance <pfe-id>> | errors> all
```

Description

Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.

Options

none

Display status information for all FPCs. On a TX Matrix router, display status information for all FPCs on the attached T640 routers in the routing matrix. On a TX Matrix Plus router, display status information for all FPCs on the attached routers in the routing matrix.

> **NOTE**: In EX8200 switches, line cards initialize Packet Forwarding Engine during startup. If an error occurs during hardware initialization, the FPCs with bad hardware parts power down after transferring the debug information to the Routing Engine. The Routing Engine marks the FPC offline, logs the error `communication failure with /fpc8/hwdfpc/...` in system log messages (/var/log/messages), and generates an alarm to inform the user.
>
> See the following sample output:

```
user@host> show chassis fpc
                    Temp  CPU Utilization (%)   Memory
Utilization (%)
Slot State           (C)  Total  Interrupt      DRAM (MB) Heap
```

```
Buffer
  0  Empty
  1  Empty
  2  Empty
  3  Empty
  4  Offline         ---Configured power off---
  5  Offline         ---Hard FPC error---
  6  Empty
  7  Online            26      4          0       1024
0         32
     error: communication failure with /fpc8/hwdfpc/
```

The following sample output shows the alarm raised for the failed FPCs:

```
user@host> show chassis alarms
4 alarms currently active
Alarm time               Class  Description
2011-03-24 00:52:51 UTC  Major  FPC 5 Hard errors
2011-03-24 00:52:31 UTC  Major  Fan Tray Failure
2011-03-24 00:52:31 UTC  Major  Fan Tray Failure
2011-03-24 00:51:26 UTC  Minor  Loss of communication with Backup RE
```

NOTE: On T4000 routers, when you include the `enhanced-mode` statement at the `[edit chassis network-services]` hierarchy level and reboot the system, only the T4000 Type 5 FPCs present on the router become online while the remaining FPCs are offline, and FPC misconfiguration alarms are generated. The `show chassis alarm` command output displays FPC misconfiguration (`FPC` *`fpc-slot`* `misconfig`) as the reason for the generation the alarms.

The following sample output shows the FPC status after the `enhanced-mode` statement is configured on the T4000 router. The T4000 Type 5 FPC present in slot 5 becomes online while the remaining FPCs are offline.

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   Memory
Utilization (%)
```

```
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap
Buffer
  0  offline         ---FPC misconfiguration---
  1  offline         ---FPC misconfiguration---
  2  offline         ---FPC misconfiguration---
  3  Empty
  4  Empty
  5  Online            66      50          0       2816
29      27
```

The following sample output shows FPC misconfiguration alarms:

```
user@host> show chassis alarms
3 alarms currently active
Alarm time               Class  Description
2011-03-24 00:52:51 PST  Major  FPC 1 misconfig
2011-03-24 00:52:31 PST  Major  FPC 2 misconfig
2011-03-24 00:52:31 PST  Major  FPC 3 misconfig
```

The following sample output shows FPC communication errors and error messages:

```
user@host> run show chassis fpc
                              Temp  CPU Utilization (%)   CPU
Utilization (%)  Memory    Utilization (%)
         Slot State            (C)  Total  Interrupt      1min
5min   15min  DRAM (MB) Heap     Buffer
           3  Online            57      3          0        3
3      3    31561      16          0
           8  Offlining       ---Configured power off---
          12  Online            49      2          0        2
2      2    31561      15          0
          13  Online            49      2          0        2
2      2    31561      14          0
          14  Offline         ---Configured power off---
          15  Offline         ---Configured power off---
 error: communication failure with /fpc8/hwdfpc/
 1. The producer app is either down or unresponsive
 Run "show system processes node <node> | grep <app>" to check if
```

```
app is running
 Run "show system application app <app>" to check the state of the
app
 2. If there is a change in mastership recently, please wait for
the switchover to complete for the app to be online
 Run "request chassis routing-engine master switch check" to check
the status of switchover
```

detail (Optional) Display detailed status information for all FPCs or for the FPC in the specified slot (see *fpc-slot* or *slot*).

all-members (MX Series routers, EX Series switches, PTX10016 only) (Optional) Display status information for all FPCs on all members of the Virtual Chassis configuration.

interconnect-device *name* (QFabric systems only) (Optional) Display status information for all FPCs on the Interconnect device.

fpc-slot (Optional) FPC slot number:

- (TX Matrix and TX Matrix Plus routers only)—On a TX Matrix router, if you specify the number of the T640 router (line-card chassis) by using the `lcc` *number* option (the recommended method), replace *fpc-slot* with a value from `0` through `7`. Otherwise, replace *fpc-slot* with a value from `0` through `31`. Likewise, on a TX Matrix Plus router, if you specify the number of the specified router (line-card chassis) by using the `lcc` *number* option (the recommended method), replace *fpc-slot* with a value from `0` through `7`. Otherwise, replace *fpc-slot* with a value from `0` through `31`. For example, the following commands have the same result:

  ```
  user@host> show chassis fpc detail 1 lcc 1
  user@host> show chassis fpc detail 9
  ```

- M120 router—Replace *fpc-slot* with a value from `0` through `5`.
- MX80 router—Replace *fpc-slot* with a value from `0` through `1`.
- MX104 and MX104-40G routers—Replace *fpc-slot* with a value from `0` through `2`.
- MX240 router—Replace *fpc-slot* with a value from `0` through `2`.
- MX480 router—Replace *fpc-slot* with a value from `0` through `5`.

- MX-960 router—Replace *fpc-slot* with a value from 0 through 11.
- MX2010 router—Replace *fpc-slot-number* with a value from 0 through 9.
- MX2008 router—Replace *fpc-slot-number* with a value from 0 through 9.
- MX2020 router—Replace *fpc-slot-number* with a value from 0 through 19.
- Other routers—Replace *fpc-slot* with a value from 0 through 7.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—Replace *fpc-slot* with 0.
  - EX4200 switches in a Virtual Chassis configuration—Replace *fpc-slot* with a value from 0 through 9.
  - EX6210 switches—Replace *fpc-slot* with a value from 0 through 9.
  - EX8208 switches—Replace *fpc-slot* with a value from 0 through 7.
  - EX8216 switches—Replace *fpc-slot* with a value from 0 through 15.
  - EX9204 switches—Replace *fpc-slot* with a value from 0 through 2.
  - EX9208 switches—Replace *fpc-slot* with a value from 0 through 5.
  - EX9214 switches—Replace *fpc-slot* with a value from 0 through 11.
- QFX Series:
  - QFXSeries and OCX Series switches—Replace *fpc-slot* with 0.
  - QFabric systems—Replace *fpc-slot* with 0 through 31 on the Interconnect device.
- PTX Series Packet Transport Routers:
  - PTX5000 Packet Transport Router—Replace *fpc-slot* with a value from 0 through 7.
  - PTX10016 Packet Transport Router—Replace *fpc-slot* with a value from 0 through 15.
- ACX Universal Metro Routers, ACX7509:

- ACX1000 and ACX2000 Universal Metro Routers—Replace *fpc-slot* with 0.

- SRX Series:

  - none—Display status information for all FPCs.

  - detail—(Optional) Display detailed FPC status information.

  - *fpc-slot*—(Optional) Display information about the FPC in this slot.

  - node—(Optional) For chassis cluster configurations, display status information for all FPCs or for the specified FPC on a specific node (device) in the cluster.

    - *node-id*—Identification number of the node. It can be 0 or 1.

    - local—Display information about the local node.

    - primary—Display information about the primary node.

  - pic-status—(Optional) Display status information for all FPCs or for the FPC in the specified slot (see *fpc-slot*).

interconnect-device *name*

(QFabric systems only) (Optional) Display status information for all FPCs on the Interconnect device.

lcc *number*

(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.

- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on the local Virtual Chassis member.

member *member-id*

(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device *name* (QFabric systems only) (Optional) Display status information for each Node device. Each Node device is equivalent to an FPC.

mic-status (MX304 devices only) (Optional) Show Modular Interface Cards installed in FPC.

pic-status (Optional) Display status information for all PICs or for the PIC in the specified slot (see *fpc-slot*).

NOTE: On T1600 routers, Type 4 FPCs with ASICs based on the SL2.0 chipset do not support the 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (10x10GE [LAN/WAN] SFPP). If you issue the `show chassis fpc` command with the `pic-status` option, the CLI displays the string "Not Supported" for 10x10GE(LAN/WAN) SFPP PICs installed on such FPCs. The following is a sample output:

```
user@host> show chassis fpc pic-status
Slot 0   Online       E2-FPC Type 1
  PIC 0  Online       1x G/E SFP, 1000 BASE
  PIC 1  Online       Adaptive Services-II
  PIC 2  Online       1x G/E IQ, 1000 BASE
  PIC 3  Online       1x G/E IQ, 1000 BASE
Slot 1   Online       FPC Type 3-ES
  PIC 0  Present      UNUSED- Not Supported
Slot 2   Online       FPC Type 4-ES
  PIC 0  Offline      4x OC-192 SONET XFP
  PIC 1  Present      10x10GE(LAN/WAN) SFPP- Not Supported
<<<<<<
Slot 4   Offline      FPC Type 1-ES
Slot 5   Offline      FPC Type 2-ES
Slot 6   Online       E2-FPC Type 3
  PIC 0  Online       1x OC-192 SONET XFP
  PIC 1  Online       4x OC-48 SONET
  PIC 2  Online       4x OC-48 SONET
  PIC 3  Online       MultiServices 500
Slot 7   Online       FPC Type 4-ES
  PIC 0  Online       4x 10GE (LAN/WAN) XFP
  PIC 1  Online       4x 10GE (LAN/WAN) XFP
```

In addition, an entry is logged in the system log messages (/var/log/messages) that the PIC is not supported. The following is a sample message logged in the system log:

```
Apr  5 08:47:36  router1 chassisd[2770]: CHASSISD_UNSUPPORTED_PIC:
PIC 1 in FPC 2 (type 763, version 257) is not supported
```

If you see this issue, contact Juniper Networks Technical Assistance Center (JTAC) for a possible fix.

NOTE: When there is a double-bit ECC error in a network processor's memory, the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP or Channelized E1/T1 Circuit Emulation MIC is switched to the offline state.

```
user@host> show chassis fpc pic-status
Slot 1   Online       MPC Type 2 3D Q
PIC 0  Offline      1xCOC12/4xCOC3 CH-CE- ECC error detected
```

pfe-instance

- *pfe-id*— (PTX Series Devices) (Optional) Display status information for the specific PFE instance on the FPCs.
- `all`—(PTX10008 and PTX10001-36MR) (Optional) Display status information for all PFE instances on the FPCs.

lcc *number*

(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

lcc *number*

(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.

Required Privilege Level

view

Output Fields

Table 278 on page 18324 lists the output fields for the `show chassis fpc` command. Output fields are listed in the approximate order in which they appear.

Table 278: show chassis fpc Output Fields

Field Name	Field Description	Level of Output

`Slot` or `Slot State`	Slot number and state. The state can be one of the following conditions:
	• `Dead`—Held in reset because of errors.
	• `Diag`—Slot is being ignored while the FPC is running diagnostics.
	• `Dormant`—Held in reset.
	• `Empty`—No FPC is present.
	• `Offline`—(PTX Series Packet Transport Routers only) One of the following two states is displayed:
	  • `FPC offlined due to unreachable destinations`
	  • `FPC Offlined due to degraded FPC action`
	• `Online`—FPC is online and running.
	• `Present`—FPC is detected by the chassis daemon but either is not supported by the current version of Junos OS or is inserted in the wrong slot. The output also states either `Hardware Not Supported` or `Hardware Not In Right Slot`. The FPC is coming up but not yet online.
	• `Probed`—Probe is complete; awaiting restart of the Packet Forwarding Engine.
	• `Probe-wait`—Waiting to be probed.
	• `Unknown`—FPC is present but the state is unknown.
	• `Present`—FPC is plugged in. The FPC is not powered on or operational.
	• `Onlining`—FPC is in the process of going online. ASIC and rest of the hardware is initializing.
	• `Offlining`—FPC is in the process of going offline. ASIC and rest of the hardware is being shutdown down to take the offline gracefully.
	• `Fault`—FPC is in an alarmed state in which none of the PICs are operational.
		all levels

Table 278: show chassis fpc Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Fault-off`—FPC is powered off due to a fault. • `Spare`—FPC is redundant and will move to active state if one of the working FPCs fails to pass traffic.	
`Logical slot`	Slot number.	all levels
`Temp (C) or Temperature`	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.	all levels all levels
`Temperature (PTX Series)`	On PTX Series Packet Transport Routers, temperature details are provided in degrees Celsius and Fahrenheit. Output includes: • Temperature (PMB)—Temperature of the air passing by the Processor Mezzanine Board (PMB) at the bottom of the FPC. • Temperature (Intake)—Temperature of the air flowing into the chassis. • Temperature (Exhaust)—Exhaust temperatures for multiple zones (Exhaust A and Exhaust B). • Temperature (TL*n*)-–Temperature of the specified Lookup ASIC (TL) of the packet forwarding engine on the FPC. • Temperature (TQ*n*)-–Temperature of the specified Queuing and Memory Interface ASIC (TQ) of the packet forwarding engine on the FPC.	`detail`
`Total CPU Utilization (%)`	Total percentage of CPU being used by the FPC's processor.	all levels
`Interrupt CPU Utilization (%)`	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.	none specified

Table 278: show chassis fpc Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`1 min CPU utilization (%)` **NOTE**: Supported only on MX240, MX480, MX960, MX2010, MX2020, and MX2008.	Information about the Routing Engine's CPU utilization in the past 1 minute.	none specified
`5 min CPU utilization (%)` **NOTE**: Supported only on MX240, MX480, MX960, MX2010, MX2020, and MX2008.	Information about the Routing Engine's CPU utilization in the past 5 minutes.	none specified
`15 min CPU utilization (%)` **NOTE**: Supported only on MX240, MX480, MX960, MX2010, MX2020, and MX2008.	Information about the Routing Engine's CPU utilization in the past 15 minutes.	none specified

Table 278: show chassis fpc Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Memory DRAM (MB)`	Total DRAM, in megabytes, available to the FPC's processor.	none specified
`Heap Utilization (%)`	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak). **NOTE**: On MX Series routers and EX Series switches in a broadband edge environment, heap utilization levels higher than 70 percent can affect unified ISSU, router stability, or scaling capability.	none specified
`Buffer Utilization (%)`	Percentage of buffer space being used by the FPC's processor for buffering internal messages.	none specified
`Total CPU DRAM`	Amount of DRAM available to the FPC's CPU.	`detail`
`Total RLDRAM`	Amount of reduced latency dynamic random access memory (RLDRAM) available to the FPC CPU.	`detail`
`Total DDR DRAM`	Amount of double data rate dynamic random access memory (DDR DRAM) available to the FPC CPU.	`detail`
`Total SRAM`	Amount of static RAM (SRAM) used by the FPC's CPU.	`detail`
`Total SDRAM`	Total amount of memory used for storing packets and notifications.	`detail`
`I/O Manager ASICs information`	I/O Manager version number, manufacturer, and part number.	`detail`
`Start time`	Time when the Routing Engine detected that the FPC was running.	`detail`
`Uptime`	How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running.	`detail`

Table 278: show chassis fpc Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`PIC type`	(`pic-status` output only) Type of PIC.	none specified
`PFE-Instance` (PTX10008 and PTX10001-36MR)	The PFE instance ID for the specific FPC.	
`PFE-State` (PTX10008 and PTX10001-36MR)	The state of PFE, The possible PFE-States are: **ONLINE:** PFE is online. **Configured power off:** PFE Offlined through configuration. **Offlined by CLI:** PFE is offlined through the request command in CLI. **Offlined by CmError:**PFE offlined through CM-Error action. **Transition Offline:** Status During PFE Transition. **Fault:** If disable-pfe cmerror action is taken by system on Major/Fatal errors.	
GNF (Node slicing)	GNF identifier associated with each line card. (`pic-status` output only) GNF identifier associated with each PIC.	all levels

Table 278: show chassis fpc Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Temperature Performance Throttle` (only for PTX10004, PTX10008, PTX10001-36MR, and PTX100016 with PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards, and PTX10001-36MR FPC Junos Evolved)	Temperature performance throttle status associated with an FPC and PFE ID : "Enabled" or "Disabled". This feature is enabled by default on all PFE instances.	`pfe` *`pfe_id`*
`Temperature voltage reduction` (only for PTX10004, PTX10008, PTX10001-36MR and PTX100016 with PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards, and PTX10001-36MR FPC Junos Evolved)	Temperature voltage reduction status associated with an FPC and PFE ID: "Enabled" or "Disabled". This feature is enabled by default on all PFE instances.	`pfe` *`pfe_id`*
PFE	All PFE instances or single PFE number (0 to 3)	`pfe` *`pfe_id`*
`PFE-Instance`	(`PFE-instance` output only) PFE instance ID.	none specified

Table 278: show chassis fpc Output Fields *(Continued)*

Field Name	Field Description	Level of Output
PFE-State	State of PFE. (Enabled or Disabled)	all levels
PPFE-ID	Per PFE ID.	pfe *pfe_id*

Sample Output (Junos OS)

show chassis fpc (MX10004 Universal Routing Platform)

```
user@host> show chassis fpc
                    Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State           (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  0  Empty
  1  Online           25      2          0        2      2      2    32768      33          2
  2  Online           30     20          0       21     21     21    3168        9         22
  3  Empty
```

show chassis fpc (EX6210 Switch)

```
user@switch> show chassis fpc
                    Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State           (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Empty
  1  Online            7      5          0       1024        0         32
  2  Empty
  3  Empty
  4  Online           25     17          2       2048        0         30
  5  Online           25      3          0       2048        0         24
  6  Online            6      5          0       1024        0         32
  7  Empty
  8  Empty
  9  Online            8      7          0       1024        0         32
```

show chassis fpc (M20 Router)

```
user@host> show chassis fpc
FPC status:
               Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State      (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Empty        0      0          0         0         0          0
  1  Online      38      0          0         8         0          4
  2  Online      35      0          0         8         0          3
  3  Empty        0      0          0         0         0          0
```

show chassis fpc detail (M Series Routers)

```
user@host> show chassis fpc detail 1
Slot 1 information:
  State                                 Online
  Temperature                        48 degrees C
  Total CPU DRAM                     32 MB
  Total SRAM                          4 MB
  Total SDRAM                       256 MB
  I/O Manager ASICs information         Version 2.0, Foundry IBM, Part number 0
  I/O Manager ASICs information         Version 2.0, Foundry IBM, Part number 0
  Start time                            2000-02-08 02:18:49 UTC
  Uptime                                14 hours, 41 minutes, 41 seconds
```

show chassis fpc detail (MX150)

```
user@host> show chassis fpc detail
Slot 0 information:
  State                                 Online
  Temperature                        42 degrees C / 107 degrees F
  Total CPU DRAM                   2048 MB
  Total RLDRAM                       10 MB
  Total DDR DRAM                      0 MB
  Start time                            2017-04-04 04:44:04 PDT
  Uptime                                7 days, 19 hours, 45 minutes, 50 seconds
```

show chassis fpc (MX104 Router)

```
user@host> show chassis fpc
Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online            32     15          5       2048       22         13
  1  Online            32     15          5       2048       22         13
  2  Online            32     15          5       2048       22         13
```

show chassis fpc detail (MX104 Router)

```
user@host> show chassis fpc detail
Slot 0 information:
  State                                 Online
  Temperature                         32 (C)
  Total CPU DRAM                    2048 MB
  Total SRAM                         403 MB
  Total SDRAM                       1316 MB
  Start time                            2013-05-23 14:39:18 IST
  Uptime                                1 hour, 20 minutes, 22 seconds
Slot 1 information:
  State                                 Online
  Temperature                         32 (C)
  Total CPU DRAM                    2048 MB
  Total SRAM                         403 MB
  Total SDRAM                       1316 MB
  Start time                            2013-05-23 14:39:18 IST
  Uptime                                1 hour, 20 minutes, 22 seconds
Slot 2 information:
  State                                 Online
  Temperature                         32 (C)
  Total CPU DRAM                    2048 MB
  Total SRAM                         403 MB
  Total SDRAM                       1316 MB
  Start time                            2013-05-23 14:39:18 IST
  Uptime                                1 hour, 20 minutes, 22 seconds
```

show chassis fpc pic-status (MX104 Router)

```
user@host> show chassis fpc pic-status
Slot 0   Online
Slot 1   Online
  PIC 0  Online       10x 1GE(LAN) -E  SFP
  PIC 1  Online       10x 1GE(LAN) -E  SFP
Slot 2   Online
  PIC 0  Online       4x 10GE(LAN) SFP+
```

show chassis fpc (MX240 Router)

```
user@host> show chassis fpc
               Temp    CPU Utilization (%)     Memory       Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Empty
  1  Online            34      6          0       1024       18         30
  2  Online            33      9          0       1024       24         30
```

If you have installed a Switch Control Board (SCB1) instead of a line card (DPC0) in the multifunction slot on the MX240, the `show chassis fpc` output shows the slot `0` as empty.

show chassis fpc detail (EX9200 Switch)

```
user@switch> show chassis fpc detail

Slot 2 information:
  State                                 Online
  Temperature                        37
  Total CPU DRAM                   2048 MB
  Total RLDRAM                      331 MB
  Total DDR DRAM                   1536 MB
  Start time:                           2014-03-12 15:35:28 UTC
  Uptime:                               1 hour, 4 minutes, 29 seconds
  Max Power Consumption             239 Watts
Slot 3 information:
  State                                 Online
  Temperature                        39
  Total CPU DRAM                   2048 MB
```

```
  Total RLDRAM                    1036 MB
  Total DDR DRAM                  6656 MB
  Start time:                          2014-03-12 15:00:18 UTC
  Uptime:                              1 hour, 39 minutes, 39 seconds
  Max Power Consumption            520 Watts
```

show chassis fpc (MX480 Router)

```
user@host> show chassis fpc
               Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  0  Online                    1          0        1      2      3    1024        4         56
  1  Online                    1          0        2      2      3    1024        4         56
```

show chassis fpc (MX480 Router with 100-Gigabit Ethernet CFP)

```
user@host> show chassis fpc
        Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State     (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
 0  Online      33   4      0              2048      10       13
 1  Online      36   7      0              2048      16       13
 2  Online      29   6      0              1024      27       29
 3  Online      33   0      0                 0       0        0
 4  Online      36   7      0              2048      19       13
 5  Online      34  31      11             2048      14       13
```

show chassis fpc pic-status (MX480 Router with 100-Gigabit Ethernet CFP)

```
user@host> show chassis fpc pic-status
Slot 1   Online       MPC Type 3
  PIC 2  Online       1X100GE CFP
 Slot 2  Online       DPCE 40x 1GE R EQ
  PIC 0  Online       10x 1GE(LAN) EQ
  PIC 1  Online       10x 1GE(LAN) EQ
  PIC 2  Online       10x 1GE(LAN) EQ
  PIC 3  Online       10x 1GE(LAN) EQ
 Slot 3  Online       MPC Type 3
  PIC 0  Online       1X100GE CFP
```

```
  PIC 2  Online       1X100GE CFP
 Slot 4  Online       MPC Type 3
  PIC 0  Online       1X100GE CFP
  PIC 2  Online       1X100GE CFP
 Slot 5  Online       MPC Type 2 3D EQ
  PIC 0  Online       2x 10GE  XFP
  PIC 1  Online       2x 10GE  XFP
  PIC 2  Online       10x 1GE(LAN) SFP
  PIC 3  Online       10x 1GE(LAN) SFP
```

show chassis fpc pic-status (EX Series Switch)

```
user@host> show chassis fpc pic-status
Slot 1   Online       EX9200 32x10G SFP
  PIC 0  Online       8X10GE SFPP
  PIC 1  Online       8X10GE SFPP
  PIC 2  Online       8X10GE SFPP
  PIC 3  Online       8X10GE SFPP
Slot 2   Online       EX9200 32x10G SFP
  PIC 0  Online       8X10GE SFPP
  PIC 1  Online       8X10GE SFPP
  PIC 2  Online       8X10GE SFPP
  PIC 3  Online       8X10GE SFPP
```

show chassis fpc (MX960 Router with MPC5EQ)

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online            38     16          0       3584        7         13
  1  Online            31     15          0       2048       17         13
  2  Empty
  3  Online            31     14          0       2048       20         13
  4  Online            34     16          0       3584        7         13
  5  Online            34     16          0       3584        7         13
  6  Empty
  7  Online            32      9          0       2048       18         14
  8  Online            36     19          0       3584        7         13
  9  Online            31      9          0       2048       13         13
```

```
 10  Online            35     14          0       3584        7          13
 11  Online            33     11          0       2048       18          14
```

show chassis fpc (MX240, MX480, MX960 with Application Services Modular Line Card)

```
user@host> show chassis fpc 1 detail
Slot 1 information:
  State                                 Online
  Temperature                        34
  Total CPU DRAM                   3072 MB
  Total RLDRAM                      259 MB
  Total DDR DRAM                   4864 MB
  Start time:                           2012-06-19 10:51:43 PDT
  Uptime:                               16 minutes, 48 seconds
  Max Power Consumption             550 Watts
```

show chassis fpc (MX240, MX480, MX960, MX2010, MX2020, and MX2008 Universal Routing Platforms with Dynamic Power Management)

```
user@host> show chassis fpc 2 detail

Slot 2 information:
  State                                 Online
  Temperature                        37
  Total CPU DRAM                   3584 MB
  Total XR2                         275 MB
  Total DDR DRAM                   20352 MB
  Start time:                           2014-07-18 02:51:23 PDT
  Uptime:                               5 minutes, 19 seconds
  Max MPC Base Power Consumption   485 Watts
  Max MIC0 Power Consumption        50 Watts
  Max MIC1 Power Consumption        50 Watts
  Max MPC Total Power Consumption  585 Watts
```

show chassis fpc (MX2010 Router with Fabric Grant Bypass Enabled)

Following is the output of the `show chassis fpc` command on an MX2010 router with Switch Fabric Board (SFB), where fabric grant bypass is enabled by default. All MPCs power on.

```
user@host> show chassis fpc
Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online            34     20          0       2048        9         14
  1  Offline           33     22                0       2048        9         14
  2  Online            33     17          0       2048        9         14
  3  Offline           34     25                0       2048        9         14
  4  Online            32     27          0       2048        9         14
  5  Offline           32     26                0            2048        9         14
  6  Empty
  7  Empty
  8  Empty
  9  Empty
```

show chassis fpc (MX2010 Router with Fabric Grant Bypass Disabled)

Following is the output of the `show chassis fpc` command on an MX2010 router with Switch Fabric Board (SFB), where fabric grant bypass has been disabled. MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFP) do not power on after you disable fabric grant bypass and reboot the router. Also, FPC misconfiguration alarms are generated.

```
user@host> show chassis fpc
 Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online            34     20          0       2048        9         14
  1  Offline         ---FPC misconfiguration---
  2  Online            33     17          0       2048        9         14
  3  Offline         ---FPC misconfiguration---
  4  Online            32     27          0       2048        9         14
  5  Offline         ---FPC misconfiguration---
  6  Empty
  7  Empty
  8  Empty
  9  Empty
```

show chassis fpc pic-status (MX2010 Router with Fabric Grant Bypass Enabled)

Following is the output of the `show chassis fpc pic-status` command on an MX2010 router with Switch Fabric Board (SFB), where fabric grant bypass has been enabled by default. All MPCs power on.

```
user@host> show chassis fpc pic-status
Slot 0   Present      MPCE Type 3 3D
Slot 1   Present      MPC Type 2 3D EQ
Slot 2   Present      MPCE Type 3 3D
Slot 3   Present      MPC 3D 16x 10GE
Slot 4   Present      MPCE Type 3 3D
Slot 5   Present      MPCE Type 1 3D Q
```

show chassis fpc pic-status (MX2010 Router with Fabric Grant Bypass Disabled)

Following is the output of the `show chassis fpc pic-status` command on an MX2010 router with Switch Fabric Board (SFB), where fabric grant bypass has been disabled. MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFP) do not power on after you disable fabric grant bypass mode and reboot the router.

```
user@host> show chassis fpc pic-status
Slot 0   Present      MPCE Type 3 3D
Slot 1   Offline      MPC Type 2 3D EQ
Slot 2   Present      MPCE Type 3 3D
Slot 3   Offline      MPC 3D 16x 10GE
Slot 4   Present      MPCE Type 3 3D
Slot 5   Offline      MPCE Type 1 3D Q
```

show chassis fpc (MX2020 Routers)

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online            10     12          0       2048       18         13
  1  Online             8      9          0       2048       18         13
  2  Online             7      9          0       2048       18         13
  3  Online             8     10          0       2048       18         13
  4  Online             9     10          0       2048       18         13
  5  Online             8      9          0       2048       18         13
```

```
  6  Online            8     10          0       2048       18         13
  7  Online            9      9          0       2048       18         13
  8  Online            9     10          0       2048       18         13
  9  Online           10      9          0       2048       18         13
 10  Online           16      8          0       2048       18         13
 11  Online           11     10          0       2048       18         13
 12  Online           10     10          0       2048       18         13
 13  Online           11      9          0       2048       18         13
 14  Online           12     10          0       2048       18         13
 15  Online           13      9          0       2048       18         13
 16  Online           13      9          0       2048       18         13
 17  Online           12      9          0       2048       18         13
 18  Online           12      8          0       2048       18         13
 19  Online           14     10          0       2048       18         13
```

show chassis fpc (MX10003 Router)

```
user@host> show chassis fpc

                     Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  0  Online            59     25          0       25     24     23    3136       12         11
  1  Online            62     29          0       26     24     23    3136       12         11
```

show chassis fpc detail (MX10003 Router)

```
user@host> show chassis fpc detail

Slot 0 information:
  State                                 Online
  Total CPU DRAM                  3136 MB
  Total RLDRAM                     771 MB
  Total DDR DRAM                  18432 MB
  Temperature                        60 degrees C / 140 degrees F
  Start time                            2017-07-19 20:49:58 PDT
  Uptime                                2 hours, 29 minutes, 22 seconds
  Max MPC base power consumption   910 Watts
  Max MIC1 power consumption        95 Watts
  Max MPC total power consumption 1005 Watts
Slot 1 information:
```

```
  State                                Online
  Total CPU DRAM                 3136 MB
  Total RLDRAM                    771 MB
  Total DDR DRAM                 18432 MB
  Temperature                       63 degrees C / 145 degrees F
  Start time                           2017-07-19 20:48:01 PDT
  Uptime                               2 hours, 31 minutes, 19 seconds
  Max MPC base power consumption   910 Watts
  Max MIC1 power consumption       155 Watts
  Max MPC total power consumption 1065 Watts
```

show chassis fpc <fpc-slot> (MX10003 Router)

```
user@host> show chassis fpc 0

                     Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  0  Online            49     26          0       22     22     23    3136       12         11
```

show chassis fpc (MX204 Router)

```
user@host> show chassis fpc

                    Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  0  Online           Absent   8          0        8      8      8    3136        8          8
```

show chassis fpc detail (MX204 Router)

```
user@host> show chassis fpc detail
Slot 0 information:
  State                                Online
  Total CPU DRAM                 3136 MB
  Total RLDRAM                    257 MB
  Total DDR DRAM                 4096 MB
  Temperature                       Absent
```

```
  Start time                          2017-11-05 22:14:01 PST
  Uptime                              2 days, 8 hours, 5 minutes, 55 seconds
```

show chassis fpc <fpc-slot> (MX204 Router)

```
user@host> show chassis fpc 0

                     Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  0  Online           Absent   8          0        8      8      8    3136        8          8
```

show chassis fpc (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online            31     20          0       3584        7         13
  1  Online            28     19          0       2048       17         13
  2  Online            27     10          0       2048       18         14
  3  Online            26     10          0       2048       13         13
  4  Online            29     19          0       3584        7         13
  5  Online            28     68          0       2048       20         13
  6  Empty
  7  Empty
  8  Empty
  9  Online            36     19          0       3584       10         13
 10  Online            37     26          0       3584       10         13
 11  Empty
 12  Empty
 13  Empty
 14  Empty
 15  Empty
 16  Empty
 17  Online            28     43          0       3584       10         13
 18  Online            29     19          0       3584        7         13
 19  Online            31     19          0       3584        7         13
```

show chassis fpc detail (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis fpc detail
Slot 0 information:
  State                                 Online
  Temperature                        31
  Total CPU DRAM                   3584 MB
  Total XR2                         291 MB
  Total DDR DRAM                  24960 MB
  Start time:                           2014-04-22 23:33:19 PDT
  Uptime:                               6 minutes, 24 seconds
  Max Power Consumption             607 Watts
Slot 1 information:
  State                                 Online
  Temperature                        28
  Total CPU DRAM                   2048 MB
  Total RLDRAM                     1036 MB
  Total DDR DRAM                   6656 MB
  Start time:                           2014-04-22 23:33:24 PDT
  Uptime:                               6 minutes, 19 seconds
  Max Power Consumption             520 Watts
Slot 2 information:
  State                                 Online
  Temperature                        27
  Total CPU DRAM                   2048 MB
  Total RLDRAM                     1036 MB
  Total DDR DRAM                  11264 MB
  Start time:                           2014-04-22 23:33:34 PDT
  Uptime:                               6 minutes, 9 seconds
  Max Power Consumption             608 Watts
Slot 3 information:
  State                                 Online
  Temperature                        26
  Total CPU DRAM                   2048 MB
  Total RLDRAM                      734 MB
  Total DDR DRAM                   3108 MB
  Start time:                           2014-04-22 23:33:39 PDT
  Uptime:                               6 minutes, 4 seconds
  Max Power Consumption             368 Watts
Slot 4 information:
  State                                 Online
  Temperature                        29
```

```
  Total CPU DRAM                    3584 MB
  Total XR2                          291 MB
  Total DDR DRAM                    24960 MB
  Start time:                            2014-04-22 23:33:51 PDT
  Uptime:                                5 minutes, 52 seconds
  Max Power Consumption              607 Watts
Slot 5 information:
  State                                  Online
  Temperature                         28
  Total CPU DRAM                    2048 MB
  Total RLDRAM                      1324 MB
  Total DDR DRAM                    5120 MB
  Start time:                            2014-04-22 23:33:57 PDT
  Uptime:                                5 minutes, 46 seconds
  Max Power Consumption              440 Watts
Slot 9 information:
  State                                  Online
 Temperature                          25
  Total CPU DRAM                    3584 MB
  Total XR2                          518 MB
  Total DDR DRAM                    49920 MB
  Start time:                            2014-04-22 23:31:20 PDT
  Uptime:                                8 minutes, 23 seconds
  Max Power Consumption             1130 Watts
Slot 10 information:
  State                                  Online
  Temperature                         32
  Total CPU DRAM                    3584 MB
  Total XR2                          518 MB
  Total DDR DRAM                    49920 MB
  Start time:                            2014-04-22 23:31:25 PDT
  Uptime:                                8 minutes, 18 seconds
  Max Power Consumption             1130 Watts
Slot 17 information:
  State                                  Online
  Temperature                         25
  Total CPU DRAM                    3584 MB
  Total XR2                          518 MB
  Total DDR DRAM                    49920 MB
  Start time:                            2014-04-22 23:31:29 PDT
  Uptime:                                8 minutes, 14 seconds
  Max Power Consumption             1130 Watts
Slot 18 information:
```

```
  State                                 Online
  Temperature                      29
  Total CPU DRAM                 3584 MB
  Total XR2                       291 MB
  Total DDR DRAM                 24960 MB
  Start time:                           2014-04-22 23:34:11 PDT
  Uptime:                               5 minutes, 32 seconds
  Max Power Consumption           607 Watts
Slot 19 information:
  State                                 Online
  Temperature                      32
  Total CPU DRAM                 3584 MB
  Total XR2                       291 MB
  Total DDR DRAM                 24960 MB
  Start time:                           2014-04-22 23:34:20 PDT
  Uptime:                               5 minutes, 23 seconds
  Max Power Consumption           607 Watts
```

show chassis fpc pic-status (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis fpc pic-status
Slot 0   Online       MPC5E 3D Q 24XGE+6XLGE
  PIC 0  Online       12X10GE SFPP OTN
  PIC 1  Online       12X10GE SFPP OTN
  PIC 2  Offline      3X40GE QSFPP
  PIC 3  Offline      3X40GE QSFPP
Slot 1   Online       MPCE Type 3 3D
  PIC 0  Online       10X10GE SFPP
  PIC 2  Online       1X100GE CXP
Slot 2   Online       MPC4E 3D 2CGE+8XGE
  PIC 0  Online       4x10GE SFPP
  PIC 1  Online       1X100GE CFP
  PIC 2  Online       4x10GE SFPP
  PIC 3  Online       1X100GE CFP
Slot 3   Online       MPCE Type 2 3D P
  PIC 0  Online       2x 10GE  XFP
  PIC 1  Online       2x 10GE  XFP
Slot 4   Online       MPC5E 3D Q 2CGE+4XGE
  PIC 0  Online       2X10GE SFPP OTN
  PIC 1  Online       1X100GE CFP2 OTN
  PIC 2  Online       2X10GE SFPP OTN
```

```
  PIC 3  Online       1X100GE CFP2 OTN
Slot 5   Online       MPC 3D 16x 10GE
  PIC 0  Online       4x 10GE(LAN) SFP+
  PIC 1  Online       4x 10GE(LAN) SFP+
  PIC 2  Online       4x 10GE(LAN) SFP+
  PIC 3  Online       4x 10GE(LAN) SFP+
Slot 9   Online       MPC6E 3D
  PIC 0  Online       2X100GE CFP2 OTN
  PIC 1  Online       2X100GE CFP2 OTN
Slot 10  Online       MPC6E 3D
  PIC 0  Online       24X10GE SFPP OTN
  PIC 1  Online       4X100GE CXP
Slot 17  Online       MPC6E 3D
  PIC 0  Online       24X10GE SFPP
  PIC 1  Online       4X100GE CXP
Slot 18  Online       MPC5E 3D Q 24XGE+6XLGE
  PIC 0  Offline      12X10GE SFPP OTN
  PIC 1  Offline      12X10GE SFPP OTN
  PIC 2  Online       3X40GE QSFPP
  PIC 3  Online       3X40GE QSFPP
Slot 19  Online       MPC5E 3D Q 24XGE+6XLGE
  PIC 0  Online       12X10GE SFPP OTN
  PIC 1  Offline      12X10GE SFPP OTN
  PIC 2  Offline      3X40GE QSFPP
  PIC 3  Online       3X40GE QSFPP
```

show chassis fpc detail (MX Series Routers)

```
user@host> show chassis fpc detail 2
Slot 0 information:
  State                                 Online
  Temperature                         36 degrees C / 96 degrees F
  Total CPU DRAM                    1024 MB
  Total RLDRAM                       256 MB
  Total DDR DRAM                    4096 MB
  Start time:                           2009-08-11 21:20:30 PDT
  Uptime:                               2 hours, 8 minutes, 50 seconds
  Max Power Consumption              335 Watts
```

show chassis fpc detail (EX Series Switches)

```
user@host> show chassis fpc detail 2
Slot 1 information:
  State                                 Online
  Temperature                        41
  Total CPU DRAM                   2048 MB
  Total RLDRAM                     1036 MB
  Total DDR DRAM                   11264 MB
  Start time:                           2013-04-02 00:04:52 PDT
  Uptime:                               7 days, 9 hours, 47 minutes, 46 seconds
  Max Power Consumption             610 Watts
Slot 2 information:
  State                                 Online
  Temperature                        41
  Total CPU DRAM                   2048 MB
  Total RLDRAM                     1036 MB
  Total DDR DRAM                   11264 MB
  Start time:                           2013-04-02 00:04:56 PDT
  Uptime:                               7 days, 9 hours, 47 minutes, 42 seconds
  Max Power Consumption             610 Watts
```

show chassis fpc detail (EX9251 Switches)

```
user@switch> show chassis fpc detail 2
Slot 0 information:
  State                                 Online
  Total CPU DRAM                 3136 MB
  Total RLDRAM                    257 MB
  Total DDR DRAM                 4096 MB
  Temperature                      Absent
  Start time                            2018-03-12 14:59:49 PDT
  Uptime                                1 day, 1 hour, 10 minutes, 48 seconds
```

show chassis fpc detail (EX9253 Switches)

```
user@switch> show chassis fpc detail 1
Slot 1 information:
  State                                 Online
```

```
  Total CPU DRAM                 3136 MB
  Total RLDRAM                    771 MB
  Total DDR DRAM                 18432 MB
  Temperature                      59 degrees C / 138 degrees F
  Start time                          2018-03-04 14:20:42 PST
  Uptime                              3 days, 10 hours, 40 minutes, 57 seconds
  Max MPC base power consumption   910 Watts
  Max MIC1 power consumption        95 Watts
  Max MPC total power consumption 1005 Watts
```

show chassis fpc (Hardware Not Supported)

```
user@host> show chassis fpc
show chassis fpc
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
 Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online          -------------------- CPU less FPC --------------------
  1  Present         ------ Hardware Not In Right Slot ------
  2  Online                    0          0         0         0          0
  3  Present         ------ Hardware Not Supported ------
  4  Empty
  5  Empty
  6  Online                    0          0         0         0          0
```

show chassis fpc detail (Hardware Not Supported)

```
user@host> show chassis fpc detail
Slot 0 information:
  State                                 Online
  Total CPU DRAM                        ---- CPU less FPC ----
  Start time                            2006-07-07 03:21:00 UTC
  Uptime                                27 minutes, 51 seconds
 Slot 1 information:
  State                                 Present
  Reason                                --- Hardware Not In Right Slot ---
 Slot 2 information:
  State                                 Online
  Total CPU DRAM                    32 MB
  Start time                            2006-07-07 03:20:59 UTC
  Uptime                                27 minutes, 52 seconds
```

```
 Slot 3 information:
  State                                 Present
  Reason                                --- Hardware Not Supported ---
  Total CPU DRAM                    0 MB
 Slot 6 information:
  State                                 Online
  Total CPU DRAM                   32 MB
  Start time                            2006-07-07 03:21:01 UTC
  Uptime                                27 minutes, 50 seconds
```

show chassis fpc pic-status

```
user@host> show chassis fpc pic-status
Slot 0 Online
  PIC 1    1x OC-12 ATM, MM
  PIC 2    1x OC-12 ATM, MM
  PIC 3    1x OC-12 ATM, MM
Slot 1 Online
  PIC 0    1x OC-48 SONET, SMIR
Slot 2 Online
  PIC 0    1x OC-192 SONET, SMSR
```

show chassis fpc pic-status (M Series Routers)

```
user@host> show chassis fpc pic-status
Slot 1   Online       FPC Type 1
  PIC 0  Present      2x OC-3 ATM, MM- Hardware Error
  PIC 1  Online       4x OC-3 SONET, SMIR
Slot 2   Online       E-FPC Type 2
  PIC 0  Online       4x G/E, 1000 BASE-SX
  PIC 1  Online       2x G/E SFP, 1000 BASE
  PIC 3  Online       1x Tunnel
Slot 3   Online       E-FPC Type 1
  PIC 0  Online       1x G/E IQ, 1000 BASE
  PIC 2  Online       1x G/E SFP, 1000 BASE
Slot 4   Online       E-FPC Type 2
  PIC 0  Online       4x G/E SFP, 1000 BASE
  PIC 1  Online       4x G/E SFP, 1000 BASE
  PIC 2  Online       4x G/E SFP, 1000 BASE
```

```
  PIC 3  Online       4x G/E SFP, 1000 BASE
Slot 5   Online       FPC Type 2
...
```

show chassis fpc pic-status (M120 Router)

```
user@host> show chassis fpc pic-status
Slot 1   Online       M120 CFPC 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) XFP
Slot 3   Online       M120 FPC Type 2 (proto)
  PIC 0  Online       2x G/E IQ, 1000 BASE
  PIC 1  Online       4x OC-3 SONET, SMIR
  PIC 2  Online       2x G/E IQ, 1000 BASE
  PIC 3  Online       8x 1GE(LAN), IQ2
Slot 4   Online       M120 FPC Type 3 (proto)
  PIC 0  Online       10x 1GE(LAN), 1000 BASE
Slot 5   Online       M120 FPC Type 1 (proto)
  PIC 0  Present      1x G/E, 1000 BASE-LX- Not Supported
  PIC 1  Online       1x CHOC3 IQ SONET, SMLR
  PIC 2  Online       4x CHDS3 IQ
  PIC 3  Online       1x G/E SFP, 1000 BASE
```

show chassis fpc pic-status (MX240, MX480, and MX960 Routers with Application Services Modular Line Card)

In the following output **Slot 1 and Slot 5** are the Application Services Modular Carrier Cards (AS MCC), **PIC 0** is the Application Services Modular Storage Card (AS MSC), and **PIC 2** is the Application Services Modular Processing Card (AS MXC).

```
user@host>show chassis fpc pic-status
Slot 2   Online       MPC Type 1 3D Q
  Slot 1   Online       AS-MCC
  PIC 0  Online       AS-MSC
  PIC 2  Online       AS-MXC
Slot 4   Offline      MPC 3D 16x 10GE
Slot 5   Offline      AS-MCC
```

show chassis fpc lcc (TX Matrix Router)

```
user@host> show chassis fpc lcc 0
lcc0-re0:
--------------------------------------------------------------------------
                     Temp  CPU                        Utilization (%)     Memory    Utilization
(%)
Slot State                     (C)  Total  Interrupt        DRAM (MB)
Heap      Buffer
  0  Empty
  1  Online                     27      2          0             256         8         44
  2  Online                     27      3          0           256        15         44
  3  Empty
  4  Empty
  5  Empty
  6  Empty
  7  Empty
```

show chassis fpc pic-status (TX Matrix Router)

```
user@host> show chassis fpc pic-status
lcc0-re0:
--------------------------------------------------------------------------
Slot 0   Online       FPC Type 3
  PIC 0  Online       1x OC-192 SM SR1
  PIC 1  Online       1x OC-192 SM SR2
  PIC 2  Online       1x OC-192 SM SR1
  PIC 3  Online       1x Tunnel
Slot 1   Online       FPC Type 2
  PIC 0  Online       1x OC-48 SONET, SMSR
  PIC 1  Online       1x OC-48 SONET, SMSR

lcc1-re0:
--------------------------------------------------------------------------

lcc2-re0:
--------------------------------------------------------------------------
Slot 1   Online       FPC Type 3
  PIC 0  Online       1x OC-192 SM SR1
Slot 5   Online       FPC Type 2
```

```
  PIC 0  Online       1x OC-48 SONET, SMSR
  PIC 1  Online       2x G/E, 1000 BASE-LX
  PIC 2  Online       2x G/E, 1000 BASE-LX
  PIC 3  Online       1x OC-48 SONET, SMSR

lcc3-re0:
--------------------------------------------------------------------------
```

show chassis fpc pic-status lcc (TX Matrix Router)

```
user@host> show chassis fpc pic-status lcc 0
lcc0-re0:
--------------------------------------------------------------------------
Slot 0   Online       FPC Type 3
  PIC 0  Online       1x OC-192 SM SR2
Slot 1   Online       FPC Type 2
  PIC 0  Online       2x OC-12 ATM2 IQ, MM
  PIC 1  Online       1x OC-48 SONET, SMSR
  PIC 2  Online       1x OC-48 SONET, SMSR
  PIC 3  Online       4x G/E, 1000 BASE-SX
```

show chassis fpc (TX Matrix Plus Router)

```
user@host> show chassis fpc
lcc0-re0:
--------------------------------------------------------------------------
                      Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State             (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Empty
  1  Online            38      4          0       2048        3         24
  2  Online            43      8          0       2048        6         24
  3  Empty
  4  Online            43      6          0       2048        6         24
  5  Empty
  6  Online            42     13          0       2048        6         24
  7  Online            45      7          0       2048        3         24

lcc2-re0:
--------------------------------------------------------------------------
                      Temp  CPU Utilization (%)   Memory    Utilization (%)
```

```
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online            42     10          0       2048        6         24
  1  Empty
  2  Online            42     11          0       2048        6         24
  3  Online            40      5          0       2048        3         24
  4  Online            33     26          0       1024        8         49
  5  Empty
  6  Online            43      8          0       2048        6         24
  7  Online            46      6          0       2048        3         24

lcc3-re0:
--------------------------------------------------------------------------
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Empty
  1  Empty
  2  Online            39     30          0       2048        7         24
  3  Empty
  4  Online            41      8          0       2048        6         24
  5  Online            41     12          0       2048        6         24
  6  Online            40      8          0       2048        6         24
  7  Online            42      4          0       2048        3         24
```

show chassis fpc lcc (TX Matrix Plus Router)

```
user@host> show chassis fpc lcc 0
lcc0-re0:
--------------------------------------------------------------------------
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Empty
  1  Online            38      4          0       2048        3         24
  2  Online            43      8          0       2048        6         24
  3  Empty
  4  Online            43      6          0       2048        6         24
  5  Empty
  6  Online            42     14          0       2048        6         24
  7  Online            45      6          0       2048        3         24
```

show chassis fpc detail (TX Matrix Plus Router)

```
user@host> show chassis fpc details

lcc0-re0:
--------------------------------------------------------------------------
Slot 1 information:
  State                                 Online
  Temperature                        38 degrees C / 100 degrees F
  Total CPU DRAM                   2048 MB
  Total SRAM                         64 MB
  Total SDRAM                      1280 MB
  Start time                            2010-10-04 20:06:22 PDT
  Uptime                                1 hour, 32 minutes, 51 seconds
Slot 2 information:
  State                                 Online
  Temperature                        43 degrees C / 109 degrees F
  Total CPU DRAM                   2048 MB
  Total SRAM                        128 MB
  Total SDRAM                      2560 MB
  Start time                            2010-10-04 20:06:37 PDT
  Uptime                                1 hour, 32 minutes, 36 seconds
Slot 4 information:
  State                                 Online
  Temperature                        43 degrees C / 109 degrees F
  Total CPU DRAM                   2048 MB
  Total SRAM                        128 MB
  Total SDRAM                      2560 MB
  Start time                            2010-10-04 20:06:40 PDT
  Uptime                                1 hour, 32 minutes, 33 seconds
Slot 6 information:
  State                                 Online
  Temperature                        42 degrees C / 107 degrees F
  Total CPU DRAM                   2048 MB
  Total SRAM                        128 MB
  Total SDRAM                      2560 MB
  Start time                            2010-10-04 20:06:42 PDT
  Uptime                                1 hour, 32 minutes, 31 seconds
Slot 7 information:
  State                                 Online
  Temperature                        45 degrees C / 113 degrees F
  Total CPU DRAM                   2048 MB
```

```
  Total SRAM                       64 MB
  Total SDRAM                    1280 MB
  Start time                          2010-10-04 20:06:43 PDT
  Uptime                              1 hour, 32 minutes, 30 seconds

lcc2-re0:
--------------------------------------------------------------------------
Slot 0 information:
  State                               Online
  Temperature                      42 degrees C / 107 degrees F
  Total CPU DRAM                 2048 MB
  Total SRAM                      128 MB
  Total SDRAM                    2560 MB
  Start time                          2010-10-04 20:06:35 PDT
  Uptime                              1 hour, 32 minutes, 38 seconds
Slot 2 information:
  State                               Online
  Temperature                      42 degrees C / 107 degrees F
  Total CPU DRAM                 2048 MB
  Total SRAM                      128 MB
  Total SDRAM                    2560 MB
  Start time                          2010-10-04 20:06:37 PDT
  Uptime                              1 hour, 32 minutes, 36 seconds
Slot 3 information:
  State                               Online
  Temperature                      40 degrees C / 104 degrees F
  Total CPU DRAM                 2048 MB
  Total SRAM                       64 MB
  Total SDRAM                    1280 MB
  Start time                          2010-10-04 20:06:28 PDT
  Uptime                              1 hour, 32 minutes, 45 seconds
Slot 4 information:
  State                               Online
  Temperature                      33 degrees C / 91 degrees F
  Total CPU DRAM                 1024 MB
  Total SRAM                       64 MB
  Total SDRAM                    1280 MB
  Start time                          2010-10-04 20:08:03 PDT
  Uptime                              1 hour, 31 minutes, 10 seconds
Slot 6 information:
  State                               Online
  Temperature                      43 degrees C / 109 degrees F
  Total CPU DRAM                 2048 MB
```

```
  Total SRAM                      128 MB
  Total SDRAM                    2560 MB
  Start time                           2010-10-04 20:06:44 PDT
  Uptime                               1 hour, 32 minutes, 29 seconds
Slot 7 information:
  State                                Online
  Temperature                     46 degrees C / 114 degrees F
  Total CPU DRAM                 2048 MB
  Total SRAM                       64 MB
  Total SDRAM                    1280 MB
  Start time                           2010-10-04 20:06:46 PDT
  Uptime                               1 hour, 32 minutes, 27 seconds

lcc3-re0:
--------------------------------------------------------------------------
Slot 2 information:
  State                                Online
  Temperature                     38 degrees C / 100 degrees F
  Total CPU DRAM                 2048 MB
  Total SRAM                      128 MB
  Total SDRAM                    2560 MB
  Start time                           2010-10-04 20:17:31 PDT
  Uptime                               1 hour, 21 minutes, 42 seconds
Slot 4 information:
  State                                Online
  Temperature                     41 degrees C / 105 degrees F
  Total CPU DRAM                 2048 MB
  Total SRAM                      128 MB
  Total SDRAM                    2560 MB
  Start time                           2010-10-04 20:17:34 PDT
  Uptime                               1 hour, 21 minutes, 39 seconds
Slot 5 information:
  State                                Online
  Temperature                     41 degrees C / 105 degrees F
  Total CPU DRAM                 2048 MB
  Total SRAM                      128 MB
  Total SDRAM                    2560 MB
  Start time                           2010-10-04 20:17:36 PDT
  Uptime                               1 hour, 21 minutes, 37 seconds
Slot 6 information:
  State                                Online
  Temperature                     40 degrees C / 104 degrees F
  Total CPU DRAM                 2048 MB
```

```
  Total SRAM                      128 MB
  Total SDRAM                    2560 MB
  Start time                           2010-10-04 20:17:39 PDT
  Uptime                               1 hour, 21 minutes, 34 seconds
Slot 7 information:
  State                                Online
  Temperature                       42 degrees C / 107 degrees F
  Total CPU DRAM                 2048 MB
  Total SRAM                       64 MB
  Total SDRAM                    1280 MB
  Start time                           2010-10-04 20:17:41 PDT
  Uptime                               1 hour, 21 minutes, 32 seconds
```

show chassis fpc pic-status (TX Matrix Plus Router)

```
user@host> show chassis fpc pic-status

lcc0-re0:
--------------------------------------------------------------------------
Slot 1   Online       FPC Type 2-ES
  PIC 0  Online       8x 1GE(LAN), IQ2
Slot 2   Online       FPC Type 4-ES
  PIC 0  Online       4x 10GE (LAN/WAN) XFP
Slot 4   Online       FPC Type 4-ES
  PIC 0  Online       4x 10GE (LAN/WAN) XFP
Slot 6   Online       FPC Type 4-ES
  PIC 0  Online       4x 10GE (LAN/WAN) XFP
  PIC 1  Online       4x 10GE (LAN/WAN) XFP
Slot 7   Online       FPC Type 3-ES
  PIC 0  Online       10x 1GE(LAN), 1000 BASE
  PIC 2  Online       1x OC-192 SM SR2
  PIC 3  Online       10x 1GE(LAN), 1000 BASE

lcc2-re0:
--------------------------------------------------------------------------
Slot 0   Online       FPC Type 4-ES
  PIC 0  Online       4x 10GE (LAN/WAN) XFP
Slot 2   Online       FPC Type 4-ES
  PIC 0  Online       4x 10GE (LAN/WAN) XFP
  PIC 1  Online       4x 10GE (LAN/WAN) XFP
Slot 3   Online       FPC Type 2-ES
```

```
  PIC 0  Online       8x 1GE(LAN), IQ2
Slot 4   Online       FPC Type 4
  PIC 0  Online       10x10GE(LAN/WAN) SFPP
Slot 6   Online       FPC Type 4-ES
  PIC 0  Online       4x OC-192 SONET XFP
Slot 7   Online       FPC Type 3-ES
  PIC 0  Online       10x 1GE(LAN), 1000 BASE
  PIC 1  Offline      1x 10GE(LAN/WAN) IQ2E
  PIC 2  Online       1x OC-192 SM SR2
  PIC 3  Online       1x Tunnel

lcc3-re0:
--------------------------------------------------------------------------
Slot 2   Online       FPC Type 4-ES
  PIC 0  Online       10x10GE(LAN/WAN) SFPP
Slot 4   Online       FPC Type 4-ES
  PIC 0  Online       4x OC-192 SONET XFP
Slot 5   Online       FPC Type 4-ES
  PIC 0  Online       4x OC-192 SONET XFP
  PIC 1  Online       4x 10GE (LAN/WAN) XFP
Slot 6   Online       FPC Type 4-ES
  PIC 1  Online       4x 10GE (LAN/WAN) XFP
Slot 7   Online       FPC Type 3-ES
  PIC 0  Online       10x 1GE(LAN), 1000 BASE
  PIC 1  Online       8x 1GE(TYPE3), IQ2E
  PIC 2  Online       4x OC-48 SONET
```

show chassis fpc (T1600 Router)

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Empty
  1  Empty
  2  Online            49      3          0       2048        3         24
  3  Online            46      6          0       2048        6         24
  4  Empty
  5  Online            46      5          0       2048        3         24
  6  Empty
  7  Online            44      8          0       1024        7         49
```

show chassis fpc detail (T1600 Router)

```
user@host> show chassis fpc detail

show chassis fpc detail
Slot 2 information:
  State                                 Online
  Temperature                        49 degrees C / 120 degrees F
  Total CPU DRAM                   2048 MB
  Total SRAM                         64 MB
  Total SDRAM                      1280 MB
  Start time                            2010-10-04 21:12:52 PDT
  Uptime                                32 minutes, 9 seconds
Slot 3 information:
  State                                 Online
  Temperature                        47 degrees C / 116 degrees F
  Total CPU DRAM                   2048 MB
  Total SRAM                        128 MB
  Total SDRAM                      2560 MB
  Start time                            2010-10-04 21:13:06 PDT
  Uptime                                31 minutes, 55 seconds
Slot 5 information:
  State                                 Online
  Temperature                        46 degrees C / 114 degrees F
  Total CPU DRAM                   2048 MB
  Total SRAM                         64 MB
  Total SDRAM                      1280 MB
  Start time                            2010-10-04 21:12:56 PDT
  Uptime                                32 minutes, 5 seconds
Slot 7 information:
  State                                 Online
  Temperature                        44 degrees C / 111 degrees F
  Total CPU DRAM                   1024 MB
  Total SRAM                         64 MB
  Total SDRAM                      1280 MB
  Start time                            2010-10-04 21:14:34 PDT
  Uptime                                30 minutes, 27 seconds
```

show chassis fpc <fpc-slot> (EX Series Switch)

```
user@host> show chassis fpc 2

                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  2  Online            40     12          0       2048       19         14
```

show chassis fpc slot (T1600 Router)

```
user@host> show chassis fpc slot 2

                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  2  Online            49      3          0       2048        3         24
```

show chassis fpc pic-status (T1600 Router)

```
user@host> show chassis fpc pic-status

Slot 2   Online       FPC Type 1-ES
  PIC 0  Online       Load Type 1
  PIC 1  Online       4x 1GE(LAN), IQ2E
  PIC 3  Online       1x OC-12-3 SFP
Slot 3   Online       FPC Type 4-ES
  PIC 0  Online       4x 10GE (LAN/WAN) XFP
  PIC 1  Online       4x OC-192 SONET XFP
Slot 5   Online       FPC Type 2-ES
  PIC 0  Online       Load Type 2
  PIC 1  Online       8x 1GE(LAN), IQ2E
  PIC 2  Online       8x 1GE(LAN), IQ2E
  PIC 3  Online       1x OC-48-12-3 SFP
Slot 7   Online       FPC Type 4
  PIC 0  Online       4x 10GE (LAN/WAN) XFP
```

show chassis fpc (T4000 Router)

```
user@host> show chassis fpc

                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online            48     15          0       2816       21         27
  1  Empty
  2  Empty
  3  Online            51     15          0       2816       21         27
  4  Empty
  5  Online            39      8          0       2048        6         23
  6  Online            49     15          0       2816       21         27
  7  Empty
```

show chassis fpc detail (T4000 Router)

```
user@host> show chassis fpc detail
Slot 0 information:
  State                                 Online
  Temperature                         48 degrees C / 118 degrees F
  Total CPU DRAM                    2816 MB
  Total SRAM                        1554 MB
  Total SDRAM                       10752 MB
  Start time                            2012-02-09 22:56:25 PST
  Uptime                                2 hours, 40 minutes, 52 seconds
Slot 3 information:
  State                                 Online
  Temperature                         51 degrees C / 123 degrees F
  Total CPU DRAM                    2816 MB
  Total SRAM                        1554 MB
  Total SDRAM                       10752 MB
  Start time                            2012-02-09 22:56:22 PST
  Uptime                                2 hours, 40 minutes, 55 seconds
Slot 5 information:
  State                                 Online
  Temperature                         39 degrees C / 102 degrees F
  Total CPU DRAM                    2048 MB
  Total SRAM                         128 MB
  Total SDRAM                       2560 MB
```

```
  Start time                          2012-02-09 22:51:27 PST
  Uptime                              2 hours, 45 minutes, 50 seconds
Slot 6 information:
  State                               Online
  Temperature                      49 degrees C / 120 degrees F
  Total CPU DRAM                 2816 MB
  Total SRAM                     1554 MB
  Total SDRAM                    10752 MB
  Start time                          2012-02-09 22:56:29 PST
  Uptime                              2 hours, 40 minutes, 48 seconds
```

show chassis fpc pic-status (T4000 Router)

```
user@host> show chassis fpc pic-status
Slot 0   Online       FPC Type 5-3D
  PIC 0  Online       12x10GE (LAN/WAN) SFPP
  PIC 1  Online       12x10GE (LAN/WAN) SFPP
Slot 3   Online       FPC Type 5-3D
  PIC 0  Online       1x100GE
  PIC 1  Online       12x10GE (LAN/WAN) SFPP
Slot 5   Online       FPC Type 4-ES
  PIC 0  Online       100GE
  PIC 1  Online       100GE CFP
Slot 6   Online       FPC Type 5-3D
  PIC 0  Online       12x10GE (LAN/WAN) SFPP
  PIC 1  Online       12x10GE (LAN/WAN) SFPP
```

show chassis fpc (MX10008 Router with MX10K-LC9600 and SFB)

MX10K-LC9600 supports SFB2 only. The following output indicates that MX10K-LC9600 doest support SFB.

```
user@host> show chassis fpc
Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
0  Offline             Offlined due to Unsupported fabric
```

show chassis fpc (QFX Series and OCX Series)

```
user@switch> show chassis fpc
Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
0  Online            26      2          0       2820        0         49
```

show chassis fpc detail (QFX3500 Switches)

```
user@switch> show chassis fpc detail
Slot 0 information:
  State                                 Online
  Temperature                       28 degrees C / 82 degrees F
  Total CPU DRAM                  2820 MB
  Total SRAM                         0 MB
  Total SDRAM                        0 MB
  Start time                            2010-09-20 01:34:13 PDT
  Uptime                                3 days, 3 hours, 31 minutes, 48 seconds
```

show chassis fpc pic-status (QFX3500 Switches)

```
user@switch> show chassis fpc pic-status
Slot 0   Online       QFX 48x10G 4x40G Switch
  PIC 0  Online       48x 10G-SFP+
  PIC 1  Online       15x 10G-SFP+
```

show chassis fpc interconnect-device (QFabric System)

```
user@switch> show chassis fpc interconnect-device
interconnect1
FPC status:
                     Temp
Slot State            (C)
  0  Online             0
  1  Online             0
  2  Online             0
  3  Online             0
```

```
  4  Online             0
  5  Online             0
  6  Online             0
  7  Online             0
  8  Online             0
  9  Online             0
 10  Online             0
 11  Online             0
 12  Online             0
 13  Online             0
 14  Online             0
 15  Online             0
```

show chassis fpc interconnect-device (QFabric System)

```
user@switch> show chassis fpc interconnect-device
interconnect1 3
FPC status:
                     Temp
Slot State            (C)
  3  Online             0
```

show chassis fpc interconnect-device detail (QFabric System)

```
user@switch> show chassis fpc interconnect-device
interconnect1 3 detail
Slot 3 information:
  State                                Online
  Temperature                      0 degrees C / 32 degrees F
  Start time                           2011-08-18 10:45:04 PDT
  Uptime                               1 minute, 49 seconds
```

show chassis fpc pic-status interconnect-device (QFabric System)

```
user@switch> show chassis fpc pic-status interconnect-device
interconnect1
Slot 0   Online       QFX 16-port QSFP+ Front Card
  PIC 0  Online       16x 40G-QSFP+
```

```
  PIC 1  Online       16x 40G-GE
Slot 1   Online       QFX 16-port QSFP+ Front Card
  PIC 0  Online       16x 40G-QSFP+
 PIC 1  Online       16x 40G-GE
Slot 2   Online       QFX 16-port QSFP+ Front Card
  PIC 0  Online       16x 40G-QSFP+
  PIC 1  Online       16x 40G-GE
Slot 3   Online       QFX 16-port QSFP+ Front Card
  PIC 0  Online       16x 40G-QSFP+
  PIC 1  Online       16x 40G-GE
Slot 4   Online       QFX 16-port QSFP+ Front Card
  PIC 0  Online       16x 40G-QSFP+
  PIC 1  Online       16x 40G-GE
Slot 5   Online       QFX 16-port QSFP+ Front Card
  PIC 0  Online       16x 40G-QSFP+
  PIC 1  Online       16x 40G-GE
Slot 6   Online       QFX 16-port QSFP+ Front Card
  PIC 0  Online       16x 40G-QSFP+
  PIC 1  Online       16x 40G-GE
Slot 7   Online       QFX 16-port QSFP+ Front Card
  PIC 0  Online       16x 40G-QSFP+
  PIC 1  Online       16x 40G-GE
Slot 8   Online       QFX Fabric Rear Card
  PIC 0  Online       16x 40G-GE
Slot 9   Online       QFX Fabric Rear Card
  PIC 0  Online       16x 40G-GE
Slot 10  Online       QFX Fabric Rear Card
  PIC 0  Online       16x 40G-GE
Slot 11  Online       QFX Fabric Rear Card
  PIC 0  Online       16x 40G-GE
Slot 12  Online       QFX Fabric Rear Card
  PIC 0  Online       16x 40G-GE
Slot 13  Online       QFX Fabric Rear Card
  PIC 0  Online       16x 40G-GE
Slot 14  Online       QFX Fabric Rear Card
  PIC 0  Online       16x 40G-GE
Slot 15  Online       QFX Fabric Rear Card
  PIC 0  Online       16x 40G-GE
```

show chassis fpc pic-status node-device (QFabric System)

```
user@switch> show chassis fpc pic-status node-device
node1
Slot node1 Online      QFX 48x10G 4x40G Switch
  PIC 0  Online       48x 10G-SFP+
  PIC 1  Online       4x 40G-QSFP+
```

show chassis fpc (PTX5000 Packet Transport Router)

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Empty
  1  Empty
  2  Online            50      6          0       2816        5         27
  3  Empty
  4  Empty
  5  Online            48      9          0       2816        5         27
  6  Empty
  7  Online            49      8          0       2816        5         27
```

show chassis fpc detail (PTX5000 Packet Transport Router)

```
user@host> show chassis fpc detail
Slot 2 information:
  State                                 Online
  Temperature                       35 degrees C / 95 degrees F (PMB)
  Temperature                       35 degrees C / 95 degrees F (Intake)
  Temperature                       50 degrees C / 122 degrees F (Exhaust A)
  Temperature                       54 degrees C / 129 degrees F (Exhaust B)
  Temperature                       54 degrees C / 129 degrees F (TL0)
  Temperature                       52 degrees C / 125 degrees F (TQ0)
  Temperature                       61 degrees C / 141 degrees F (TL1)
  Temperature                       58 degrees C / 136 degrees F (TQ1)
  Temperature                       57 degrees C / 134 degrees F (TL2)
  Temperature                       58 degrees C / 136 degrees F (TQ2)
  Temperature                       62 degrees C / 143 degrees F (TL3)
```

```
  Temperature                      61 degrees C / 141 degrees F (TQ3)
  Total CPU DRAM                2816 MB
  Total SRAM                       0 MB
  Total SDRAM                      0 MB
  Start time                         2012-01-12 12:05:42 PST
  Uptime                             3 hours, 14 minutes, 7 seconds
Slot 5 information:
  State                              Online
  Temperature                      35 degrees C / 95 degrees F (PMB)
  Temperature                      34 degrees C / 93 degrees F (Intake)
  Temperature                      48 degrees C / 118 degrees F (Exhaust A)
  Temperature                      53 degrees C / 127 degrees F (Exhaust B)
  Temperature                      54 degrees C / 129 degrees F (TL0)
  Temperature                      52 degrees C / 125 degrees F (TQ0)
  Temperature                      69 degrees C / 156 degrees F (TL1)
  Temperature                      56 degrees C / 132 degrees F (TQ1)
  Temperature                      54 degrees C / 129 degrees F (TL2)
  Temperature                      56 degrees C / 132 degrees F (TQ2)
  Temperature                      59 degrees C / 138 degrees F (TL3)
  Temperature                      60 degrees C / 140 degrees F (TQ3)
  Total CPU DRAM                2816 MB
  Total SRAM                       0 MB
  Total SDRAM                      0 MB
  Start time                         2012-01-12 12:05:43 PST
  Uptime                             3 hours, 14 minutes, 6 seconds
Slot 7 information:
  State                              Online
  Temperature                      35 degrees C / 95 degrees F (PMB)
  Temperature                      33 degrees C / 91 degrees F (Intake)
  Temperature                      50 degrees C / 122 degrees F (Exhaust A)
  Temperature                      55 degrees C / 131 degrees F (Exhaust B)
  Temperature                      56 degrees C / 132 degrees F (TL0)
  Temperature                      56 degrees C / 132 degrees F (TQ0)
  Temperature                      61 degrees C / 141 degrees F (TL1)
  Temperature                      57 degrees C / 134 degrees F (TQ1)
  Temperature                      55 degrees C / 131 degrees F (TL2)
  Temperature                      59 degrees C / 138 degrees F (TQ2)
  Temperature                      62 degrees C / 143 degrees F (TL3)
  Temperature                      62 degrees C / 143 degrees F (TQ3)
  Total CPU DRAM                2816 MB
  Total SRAM                       0 MB
  Total SDRAM                      0 MB
```

```
  Start time                          2012-01-12 12:05:44 PST
  Uptime                              3 hours, 14 minutes, 5 seconds
```

show chassis fpc pic-status (PTX5000 Packet Transport Router)

```
user@host> show chassis fpc pic-status
Slot 2   Online       FPC
  PIC 0  Online       24x 10GE(LAN) SFP+
  PIC 1  Online       24x 10GE(LAN) SFP+
Slot 5   Online       FPC
  PIC 0  Online       24x 10GE(LAN) SFP+
  PIC 1  Online       2x 40GE CFP
Slot 7   Online       FPC
  PIC 0  Online       24x 10GE(LAN) SFP+
  PIC 1  Online       2x 40GE CFP
```

show chassis fpc (PTX10008)

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  0  Online            38     26          2       26     26     26    1953       20         32
  1  Empty
  2  Empty
  3  Empty
  4  Empty
  5  Online            67     26          2       26     26     26    1953       25         32
  6  Online            52     26          2       26     26     26    1953       25         32
  7  Empty
```

show chassis fpc pfe-instance all (PTX10008 and PTX10001-36MR Router)

To view the PFE state information for all FPCs and PFEs on the device:

```
user@host> show chassis fpc pfe-instance all

FPC 0
PFE-Instance   PFE-State
```

```
0               Configured power off
1               Configured power off
2               Configured power off
3               Configured power off
4               Configured power off

FPC 2
PFE-Instance   PFE-State
0             Configured power off
1             Configured power off
2             Configured power off
3             Configured power off
4             TRANSITION_OFFLINE

FPC 3
PFE-Instance   PFE-State
0             ONLINE
1             ONLINE
2             Offlined by CmError
3             ONLINE
4             Offlined by CLI
```

show chassis fpc detail (PTX10008)

```
user@host> show chassis fpc detail

Slot 6 information:
  State                                 Online
  Total CPU DRAM                  8192 MB
  Temperature                       42 degrees C / 107 degrees F
  Start time                           2018-09-17 02:42:16 PDT
  Uptime                               53 minutes, 40 seconds
  Max power consumption             675 Watts
Slot 7 information:
  State                                 Online
  Total CPU DRAM                  8192 MB
  Temperature                       51 degrees C / 123 degrees F
  Start time                           2018-09-17 02:42:26 PDT
  Uptime                               53 minutes, 30 seconds
  Max power consumption            1150 Watts
```

show chassis fpc (PTX10016 Router)

```
user@host> show chassis fpc

                     Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  0  Empty
  1  Online            36     27          2       27     27     27    1953       22         32
  2  Empty
  3  Online            36     27          2       27     27     27    1953       22         32
  4  Empty
  5  Empty
  6  Online            35     27          2       27     27     27    1953       22         32
  7  Empty
  8  Online            34     27          2       27     27     27    1953       22         32
  9  Online            46     24          2       24     24     24    1953       26         32
 10  Empty
 11  Empty
 12  Empty
 13  Empty
 14  Empty
 15  Empty
```

show chassis fpc detail (PTX10016 Router)

```
user@host> show chassis fpc detail

Slot 0 information:
  State                                 Online
  Total CPU DRAM                  8192 MB
  Temperature                       44 degrees C / 111 degrees F
  Start time                            2018-09-10 07:01:09 PDT
  Uptime                                6 days, 23 hours, 17 minutes, 9 seconds
  Max power consumption            1150 Watts
Slot 4 information:
  State                                 Online
  Total CPU DRAM                  8192 MB
  Temperature                       40 degrees C / 104 degrees F
  Start time                            2018-09-10 07:01:17 PDT
  Uptime                                6 days, 23 hours, 17 minutes, 1 second
```

```
  Max power consumption           1150 Watts
Slot 6 information:
  State                               Online
  Total CPU DRAM                 8192 MB
  Temperature                      42 degrees C / 107 degrees F
  Start time                          2018-09-10 07:01:27 PDT
  Uptime                              6 days, 23 hours, 16 minutes, 51 seconds
  Max power consumption           1150 Watts
Slot 7 information:
  State                               Online
  Total CPU DRAM                 8192 MB
  Temperature                      41 degrees C / 105 degrees F
  Start time                          2018-09-10 07:01:32 PDT
  Uptime                              6 days, 23 hours, 16 minutes, 46 seconds
  Max power consumption           1150 Watts
Slot 9 information:
  State                               Online
  Total CPU DRAM                 16384 MB
  Temperature                      42 degrees C / 107 degrees F
  Start time                          2018-09-10 07:01:45 PDT
  Uptime                              6 days, 23 hours, 16 minutes, 33 seconds
  Max power consumption           1150 Watts
Slot 10 information:
  State                               Online
  Total CPU DRAM                 8192 MB
  Temperature                      41 degrees C / 105 degrees F
  Start time                          2018-09-10 07:01:46 PDT
  Uptime                              6 days, 23 hours, 16 minutes, 32 seconds
  Max power consumption           1150 Watts
Slot 11 information:
  State                               Online
  Total CPU DRAM                 16384 MB
  Temperature                      40 degrees C / 104 degrees F
  Start time                          2018-09-10 07:01:55 PDT
  Uptime                              6 days, 23 hours, 16 minutes, 23 seconds
  Max power consumption           1150 Watts
Slot 14 information:
  State                               Online
  Total CPU DRAM                 8192 MB
  Temperature                      42 degrees C / 107 degrees F
  Start time                          2018-09-10 07:01:54 PDT
  Uptime                              6 days, 23 hours, 16 minutes, 24 seconds
  Max power consumption           1150 Watts
```

```
Slot 15 information:
  State                                 Online
  Total CPU DRAM                  16384 MB
  Temperature                        41 degrees C / 105 degrees F
  Start time                            2018-09-10 07:02:03 PDT
  Uptime                                6 days, 23 hours, 16 minutes, 15 seconds
  Max power consumption            1150 Watts
```

show chassis fpc (ACX2000 Universal Metro Router, ACX7900 Router)

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online            61     17          6       512        21         37
```

show chassis fpc 0 (ACX2000 Universal Metro Router, ACX7900 Router)

```
user@host> show chassis fpc 0
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online            61     17          6       512        21         37
```

show chassis fpc detail (ACX2000 Universal Metro Router, ACX7900 Router)

```
user@host> show chassis fpc detail
Slot 0 information:
  State                                 Online
  Temperature                        61 degrees C / 141 degrees F
  Total CPU DRAM                    512 MB
  Start time                            2012-05-29 02:52:06 PDT
  Uptime                                27 minutes, 17 seconds
```

show chassis fpc pic-status (ACX2000 Universal Metro Router , ACX7900 Router)

```
user@host> show chassis fpc pic-status
Slot 0   Online
```

```
  PIC 0  Online       16x CHE1T1, RJ48
  PIC 1  Online       8x 1GE(LAN) RJ45
  PIC 2  Online       2x 1GE(LAN) SFP
  PIC 3  Online       2x 10GE(LAN) SFP+
```

show chassis FPC 1 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis fpc 1
                  Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  1  Online            34      5          0       3072        5         13
```

show chassis FPC 1 detail (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis fpc 1 detail
Slot 1 information:
  State                                 Online
  Temperature                        34
  Total CPU DRAM                   3072 MB
  Total RLDRAM                      259 MB
  Total DDR DRAM                   4864 MB
  Start time:                           2012-06-19 10:51:43 PDT
  Uptime:                               16 minutes, 48 seconds
  Max Power Consumption             550 Watts
```

show chassis fpc pic-status (PTX5000 Router)

```
user@router>show chassis fpc pic-status

Slot 0   Online       FPC-P2
  PIC 0  Online       Universal pic 96x10_24x40
  PIC 1  Online       Universal pic 96x10_24x40
Slot 1   Online       FPC-P2
  PIC 1  Online       Universal pic 96x10_24x40
```

show chassis fpc detail (MX304)

```
user@router> show chassis fpc detail
Slot 0 information:
  State                                 Online
  Total CPU DRAM                  32768 MB
  Total HBM                       65536 MB
  FIPS Capable                          True
  Temperature                        52 degrees C / 125 degrees F
  Start time                            2022-06-06 16:47:31 IST
  Uptime                                19 hours, 2 minutes, 33 seconds
  Max MPC base power consumption  1105 Watts
  Max MIC0 power consumption        95 Watts
  Max MIC1 power consumption        95 Watts
  Max MPC total power consumption 1295 Watts

PFE Information:

  PFE  Power ON/OFF  Bandwidth         SLC
  0    ON            800G
  1    ON            800G
  2    ON            800G
  3    ON            800G
  4    OFF           _
  5    OFF           _
```

show chassis fpc mic-status (MX304 with LMIC)

```
user@host> show chassis fpc mic-status 0
Slot 0   Online       FPC-BUILTIN
 MIC 0   Online       MRATE LMIC 16x100G/4x400G
  PIC 0  Online       MRATE LMIC 16x100G/4x400G
 MIC 1   Online       MRATE LMIC 16x100G/4x400G
  PIC 1  Online       MRATE LMIC 16x100G/4x400G
 MIC 2   Online       MRATE LMIC 16x100G/4x400G
  PIC 2  Online       MRATE LMIC 16x100G/4x400G
```

Sample Output (Junos OS Evolved)

show chassis fpc (Junos OS Evolved)

```
user@router> show chassis fpc
                    Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State           (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  1  Offline        ---Configured power off---
  4  Online           54      2          0        2      2      2    16384      37          1
  5  Offline        ---Configured power off---
```

> **NOTE**: The `show chassis fpc` command does not display the CPU and memory utilization details on fixed configuration routers, such as PTX10001-36MR, PTX10003-80C, PTX10003-160C, and so on, as the FPCs in these routers do not have dedicated CPUs and memory.

On PTX10003-80C and PTX10003-160C devices, if you have configured a field-replaceable unit (FRU) to be powered off and rebooted the device, the `show chassis fpc` output will display the state of the FRU as Present and reason as `Configured power off`. See the following output:

```
user@router> show chassis fpc

                    Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State           (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  0  Present        ---Configured power off---
  1  Online
```

The `show chassis fpc` output displays the FPC offline reason as `Offlined by cli command` when the FPC becomes offline because of an fpc major or fatal error. This behavior is applicable only if you have configured the offline action for major or fatal error severity, by using the command `set chassis error major action offline`. See the following output:

```
user@router> show chassis fpc
                    Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State           (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  0  Offline        ---Offlined by cli command---
  1  Online
```

show chassis fpc 0 pfe all (PTX10002-36QDD)

```
user@host> show chassis fpc 0 pfe all
FPC 0
PFE         PFE-Instance        PFE-State
0             0              Configured power off
1             0              Configured power off
2             1              ONLINE
3             1              ONLINE
```

show chassis fpc pfe-instance (PTX10008, Junos OS Evolved)

```
user@host> run show chassis fpc pfe-instance all
FPC 0
PFE-Instance                PFE-State
    0                       Configured power off     <<< PFE Offline via
Config
    1                       ONLINE
    2                       ONLINE
    3                       ONLINE
    4                       ONLINE
```

show chassis fpc pic-status (PTX10016, Junos OS Evolved)

```
user@host> show chassis fpc pic-status
Slot 0   Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 1  Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 2   Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 3  Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 4   Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 5  Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 6   Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
```

```
Slot 7  Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 8   Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 9  Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 10   Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 11  Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 12   Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 13  Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 14   Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
Slot 15  Online       JNP10K-LC1201
  PIC 0  Online       JNP10K-36QDD-LC-PIC
```

show chassis fpc pfe (PTX10004, PTX10008, PTX100016 with PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards, and PTX10001-36MR)

```
user@router> show chassis fpc 0 pfe-instance all
    FPC 0
PFE-Instance      PFE-State
   0              ONLINE
   1              ONLINE
   2              ONLINE
```

```
user@router> show chassis fpc 0 pfe-instance all detail
    FPC 0
    PFE-Instance    PFE-State     PPFE-ID      Temp-Perf-Throttle       Temp-Volt-Reduction
       0            ONLINE        0, 1         Enabled                  Enabled|Nominal
       1            ONLINE        2, 3         Enabled                  Enabled|Nominal
       2            ONLINE        4, 5         Enabled                  Enabled|Nominal
```

```
        3            ONLINE        6, 7         Enabled                  Enabled|Optimized
        4            ONLINE        8            Enabled                  Enabled|Nominal
```

```
user@router> show chassis fpc 0 pfe-instance 0 detail
    FPC 0
    PFE-Instance    PFE-State     PPFE-ID      Temp-Perf-Throttle       Temp-Volt-Reduction
       0            ONLINE        0, 1         Enabled                  Enabled
```

show chassis fpc (ACX7000 Series)

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap     Buffer
  0  Online
```

```
user@host> show chassis fpc detail
Slot 0 information:
  State                                 Online
  Start time                            2024-08-06 23:58:02 PDT
  Uptime                                7 hours, 38 seconds
  Max power consumption            825 Watts
```

```
user@host> show chassis fpc pic-status
Slot 0   Online       ACX7100-48L
  PIC 0  Online       MRATE- 48xSFP56 + 6xQSFP56-DD
```

Sample Output (Node Slicing)

show chassis fpc (Node Slicing)

```
user@router>show chassis fpc

                     Temp  CPU Utilization (%)   CPU Utilization (%)  Memory    Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM (MB) Heap
Buffer    GNF
```

```
  0  Online            45     12          0       12     12     12    3584        6
25       3
  1  Online            57     22          0       20     20     20    3136       16
22       2
  2  Online            50     19          0       17     17     16    3584        6
25       3
  3  Online            28     10          0       11     11     11    2048       10
20       6
  4  Online            42     20          0       20     19     19    3584        8
25       6
  5  Online            58     22          0       21     20     20    3136       16
22       4
  6  Online            49     17          0       15     16     16    3136       13
20       1
  7  Online            44     11          0       10     10     10    3584        6
25       5
  8  Online            40     19          0       18     18     18    3584        8
25       5
  9  Online            44     19          0       20     20     20    3584        8
25       5
```

show chassis fpc pic-status (Node Slicing)

```
user@router> show chassis fpc pic-status

Slot 0   Online       MPC5E 3D 24XGE+6XLGE                          GNF 3
  PIC 0  Online       12X10GE SFPP OTN
  PIC 1  Offline      12X10GE SFPP OTN
  PIC 2  Offline      3X40GE QSFPP
  PIC 3  Online       3X40GE QSFPP
Slot 1   Online       MPC9E 3D                                      GNF 2
  PIC 1  Online       MRATE-12xQSFPP-XGE-XLGE-CGE
Slot 2   Online       MPC5E 3D Q 2CGE+4XGE                          GNF 3
  PIC 0  Online       2X10GE SFPP OTN
  PIC 1  Online       1X100GE CFP2 OTN
  PIC 2  Online       2X10GE SFPP OTN
  PIC 3  Online       1X100GE CFP2 OTN
Slot 3   Online       MPCE Type 2 3D EQ                             GNF 6
Slot 4   Online       MPC6E 3D                                      GNF 6
  PIC 0  Online       24X10GE SFPP
  PIC 1  Online       2X100GE CFP2 OTN
```

```
Slot 5   Online       MPC9E 3D                                      GNF 4
  PIC 0  Online       MRATE-12xQSFPP-XGE-XLGE-CGE
Slot 6   Online       MPC7E 3D MRATE-12xQSFPP-XGE-XLGE-CGE          GNF 1
  PIC 0  Online       MRATE-6xQSFPP-XGE-XLGE-CGE
  PIC 1  Online       MRATE-6xQSFPP-XGE-XLGE-CGE
Slot 7   Online       MPC5E 3D 2CGE+4XGE                            GNF 5
  PIC 0  Online       2X10GE SFPP OTN
  PIC 1  Online       1X100GE CFP2 OTN
  PIC 2  Online       2X10GE SFPP OTN
  PIC 3  Online       1X100GE CFP2 OTN
Slot 8   Online       MPC6E 3D                                      GNF 5
  PIC 0  Online       24X10GE SFPP OTN
Slot 9   Online       MPC6E 3D                                      GNF 5
  PIC 0  Online       24X10GE SFPP
  PIC 1  Online       4X100GE CXP
```

show chassis fpc (MX10004 and MX10008 Devices)

On MX10004 and MX10008 devices, the JNP10K-LC4800 line cards are powered ON only when the following compatible components are present:

- Fan Trays : JNP10004-FAN2, JNP10008-FAN2
- PSM or PEM : JNP10K-PWR-AC2, JNP10K-PWR-AC3, JNP10K-PWR-DC2
- Enhanced Switch Fabric Board (SFB2)

The `show chassis fpc` command output displays the following error message when compatible fan trays are not present:

```
user@host> show chassis fpc 3
                    Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            I  Total  Interrupt      DRAM (MB) Heap     Buffer
3  Offline         ---FPC incompatible with FAN-TRAYs ---
```

The `show chassis fpc` command output displays the following error message when compatible PSMs or PEMs are not present:

```
user@router> show chassis sfb slot 0
Slot  State              Uptime
0    Offline              ---SFB incompatible with PEMs ---
```

show chassis fpc (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc
                      Temp      CPU Utilization (%)       Memory    Utilization (%)
Slot State            (C)      Total  Interrupt            DRAM (MB) Heap     Buffer
  0  Empty
  1  Empty
  2  Empty
  3  Online            37          3          0             1024        7     42
  4  Empty
  5  Empty
  6  Online            30           8         0               1024        23    30
  7  Empty
  8  Empty
  9  Empty
 10  Empty
 11  Empty
```

show chassis fpc (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3)

```
user@host> show chassis fpc
                       Temp  CPU Utilization (%)   CPU Utilization (%)
Memory
                              Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM
(MB)
                         Heap     Buffer
  0  Online            36     20          0       20     19     19
1024
                                4         26
  1  Online            35      8          0        8      8      8
```

```
2048
                                12          14
  2  Online            40     21           0       20     20     20
3584
                                 5          13
```

Sample Output

show chassis fpc detail 2

```
user@host> show chassis fpc detail 2
Slot 2 information:
  State                                 Online
  Temperature                            37
  Total CPU DRAM                      1024 MB
  Total RLDRAM                         0 MB
  Total DDR DRAM                       0 MB
  Start time:                           2012-07-18 07:18:50 PDT
  Uptime:                               4 days, 21 hours, 51 minutes, 59 seconds
  Max Power Consumption                0 Watts
```

Sample Output

show chassis fpc pic-status (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc pic-status
Slot 3   Online       SRX5k SPC
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
Slot 6   Online       SRX5k DPC 4x 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SPC2)

```
user@host> show chassis fpc pic-status

Slot 0   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 2   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 3   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SRX5K-MPC)

```
user@host> show chassis fpc pic-status

Slot 0   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Online       SRX5k DPC 4X 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
```

```
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
Slot 6   Offline      SRX5k SPC II
Slot 9   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 10  Online       SRX5k IOC II
  PIC 0  Online       10x 10GE SFP+
  PIC 2  Online       1x 100GE CFP
Slot 11  Online       SRX5k IOC II
  PIC 0  Online       1x 100GE CFP
  PIC 2  Online       2x 40GE QSFP+
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices when Express Path [formerly known as services offloading] is configured)

```
user@host> show chassis fpc pic-status

Slot 0   Offline      SRX5k DPC 40x 1GE
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Offline      SRX5k SPC
Slot 4   Online       SRX5k IOC3 24XGE+6XLG
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 5   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 1  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
```

show chassis fpc pic-status (with 20-Gigabit Ethernet MIC with SFP)

```
user@host> show chassis fpc pic-status
```

```
node0:
--------------------------------------------------------------------------
Slot 0   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 1   Offline      SRX5k SPC II
Slot 2   Online       SRX5k DPC 4X 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
Slot 9   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP
  PIC 1  Online       10x 1GE(LAN) SFP
  PIC 2  Online       10x 1GE(LAN) SFP
  PIC 3  Online       10x 1GE(LAN) SFP
Slot 10  Online       SRX5k IOC II
  PIC 0  Online       10x 10GE SFP+
  PIC 2  Online       1x 100GE CFP
Slot 11  Offline      SRX5k IOC II
```

show chassis fpc pic-status (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3 and when Express Path [formerly known as services offloading] is configured)

```
user@host> show chassis fpc pic-status
Slot 0   Offline      SRX5k DPC 40x 1GE
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Offline      SRX5k SPC
Slot 4   Online       SRX5k IOC3 24XGE+6XLG
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 5   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 1  Online       10x 1GE(LAN) SFP- np-cache/services-offload
```

```
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
```

Sample Output

show chassis fpc pic-status for HA (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc pic-status
node0:
--------------------------------------------------------------
Slot 4   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow

node1:
---------------------------------------------------------------
Slot 4   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow
```

show chassis fpc pic-status for HA (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3)

```
user@host> show chassis fpc pic-status
user@host> show chassis fpc pic-status
node0:
--------------------------------------------------------------------------
Slot 2   Online       SRX5k IOC3 24XGE+6XLG
  PIC 0  Online       12x 10GE SFP+
```

```
  PIC 1  Online       12x 10GE SFP+
  PIC 2  Offline      3x 40GE QSFP+
  PIC 3  Offline      3x 40GE QSFP+
Slot 4   Online       SRX5k IOC II
  PIC 2  Online       10x 10GE SFP+
Slot 5   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Offline
  PIC 3  Offline

node1:
--------------------------------------------------------------------------
Slot 2   Online       SRX5k IOC3 24XGE+6XLG
  PIC 0  Online       12x 10GE SFP+
  PIC 1  Online       12x 10GE SFP+
  PIC 2  Offline      3x 40GE QSFP+
  PIC 3  Offline      3x 40GE QSFP+
Slot 4   Online       SRX5k IOC II
  PIC 2  Online       10x 10GE SFP+
Slot 5   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Offline
  PIC 3  Offline
```

Usage

- To display all PFE information along with temperature performance throttle status for FPC slot 1 (FPC 0), use the command `show chassis fpc 0 pfe-instance all detail`.
- To display only PFE 0 information along with temperature performance throttle status for FPC slot 1 (FPC 0), use the command `show chassis fpc 0 pfe-instance 0 detail`.

Release Information

Command introduced before Junos OS Release 7.4.

pfe-instance option added in Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

request chassis fpc | 14475

show chassis fabric fpcs | 18079

show chassis fpc errors

IN THIS SECTION

- Syntax | 18387
- Description | 18387
- Options | 18388
- Required Privilege Level | 18388
- Output Fields | 18388
- Sample Output | 18389
- Release Information | 18392

Syntax

```
show chassis fpc errors;
```

Description

Display chassis error information including FPC number, severity of error, number of error occurred, cleared, threshold, and corresponding action.

Error Severity Level	Default Threshold	Default Action
Fatal	1	Restart the FPC
Major	1	Get the current state of the FPC and raise an alarm.

(Continued)

Error Severity Level	Default Threshold	Default Action
Minor	10	Write a log for the event.

NOTE: The MX104 router does not support the `show chassis fpc error` CLI command. This command is not available in MX104 in the Junos OS releases 13.3R7, 15.1R2, 14.1R5, 14.2R4, 13.3R8, and later.

Options

detail Display the entire list of errors that contribute to a failure for all the FPCs.

fpc-slot Display the list of errors pertaining to the FPC installed in a particular slot

Required Privilege Level

view

Output Fields

Table 279 on page 18388 lists the output fields for the `show chassis fpc errors` command. Output fields are listed in the approximate order in which they appear.

Table 279: show chassis fpc errors Output Fields

Field Name	Field Description
FPC	The FPC number.

Table 279: show chassis fpc errors Output Fields *(Continued)*

Field Name	Field Description
`Level`	The severity of the error. It can be configured as follows: • `fatal`—Fatal error on FPC • `major`—Major error on FPC • `minor`—Minor error on FPC
`Occurred`	Number of error instances occurred.
`Cleared`	Number of error instances cleared.
`Threshold`	Configured threshold value. The associated detection and recovery actions are triggered when this threshold value is crossed.
Action-Taken	The number of instances of action taken.
`Action`	The detection and recovery actions that are triggered when the threshold value is crossed. • Restart the FPC. • Get the current state of the FPC and raise an alarm. • Write a log for the event. • Disable PFE. • Reset PFE.

Sample Output

show chassis fpc errors

```
user@host> show chassis fpc errors

FPC  Level Occurred Cleared Threshold Action-Taken Action
```

```
0   Minor      0      0     10          LOG|
    Major      0      0      1        GET STATE|ALARM|
    Fatal      0      0      1         TRAP
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED | pfe-4 -
ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
1   Minor      0      0     10          LOG|
    Major      0      0      1        GET STATE|ALARM|
    Fatal      0      0      1         TRAP
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED | pfe-4 -
ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
2   Minor      0      0     10          LOG|
    Major      0      0      1        GET STATE|ALARM|
    Fatal      0      0      1         TRAP
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED | pfe-4 -
ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
4   Minor      0      0     10          LOG|
    Major      0      0      1        GET STATE|ALARM|
    Fatal      0      0      1         TRAP
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED | pfe-4 -
ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
5   Minor      0      0     10          LOG|
    Major      0      0      1        GET STATE|ALARM|
    Fatal      0      0      1         TRAP
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED | pfe-4 -
ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
6   Minor      0      0     10          LOG|
    Major      0      0      1        GET STATE|ALARM|
    Fatal      0      0      1         TRAP
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED | pfe-4 -
ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
7   Minor      0      0     10          LOG|
    Major      0      0      1        GET STATE|ALARM|
    Fatal      0      0      1         TRAP
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED | pfe-4 -
ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
```

show chassis fpc errors (MX Series Routers)

```
user@host> show chassis fpc errors 0
FPC Scope    Category   Level Occurred Cleared Threshold Action-Taken Action
  0 board    functional Minor      1        0       1         2       LOG|CM ALARM|
```

```
                        Major      0        0       1         0       GET STATE|CM ALARM|
                        Fatal      0        0       1         0       GET STATE|RESET
             memory     Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|
                        Fatal      0        0       1         0       GET STATE|RESET
             io         Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|
                        Fatal      0        0       1         0       GET STATE|RESET
             storage    Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|
                        Fatal      0        0       1         0       GET STATE|RESET
             switch     Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|
                        Fatal      0        0       1         0       GET STATE|RESET
             processing Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|
                        Fatal      0        0       1         0       GET STATE|RESET
             internal   Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|
                        Fatal      0        0       1         0       GET STATE|RESET
    pfe      functional Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|DISABLE
PFE|RESET PFE
                        Fatal      0        0       1         0       GET STATE|RESET
             memory     Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|DISABLE
PFE|RESET PFE
                        Fatal      0        0       1         0       GET STATE|RESET
             io         Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|DISABLE
PFE|RESET PFE
                        Fatal      0        0       1         0       GET STATE|RESET
             storage    Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|DISABLE
PFE|RESET PFE
                        Fatal      0        0       1         0       GET STATE|RESET
             switch     Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|DISABLE
PFE|RESET PFE
                        Fatal      0        0       1         0       GET STATE|RESET
             processing Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|DISABLE
PFE|RESET PFE
```

```
                        Fatal      0        0       1         0       GET STATE|RESET
             internal   Minor      0        0       1         0       LOG|CM ALARM|
                        Major      0        0       1         0       GET STATE|CM ALARM|DISABLE
PFE|RESET PFE
                        Fatal      0        0       1         0       GET STATE|RESET
                        Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -
ENABLED | pfe-4 -ENABLED
```

Release Information

Command introduced in Junos OS Release 12.3.

Reset PFE output option added in Junos OS Release 21.4R3.

RELATED DOCUMENTATION

error | 6441

show chassis fpc ether-types

IN THIS SECTION

- Syntax | 18393
- Description | 18393
- Options | 18393
- Required Privilege Level | 18393
- Output Fields | 18393
- Sample Output | 18394
- Release Information | 18397

Syntax

```
show chassis fpc ether-types
<slot-number>
```

Description

Display the current EtherType values in ternary content addressable memory (TCAM) on a device for a specified FPC slot or for all currently active (online) FPCs on the device.

Ethernet frame headers contain an EtherType field to identify the protocol in the frame's payload so that a receiving device knows how to process incoming Layer 2 traffic. Devices store the EtherType values they can process in TCAM for fast access. You can configure custom EtherType values in place of the default values in the TCAM EtherTypes table, although some values are reserved and can't be changed or configured. For details on how to configure custom EtherType values, see the **ether-type** configuration statement at the **[edit chassis fpc *slot*]** hierarchy.

Options

none	Display the EtherType values in TCAM for all active (online) FPC slots.
slot-number	(Optional) Display the EtherType values in TCAM for the specified FPC slot number.

Required Privilege Level

view

Output Fields

Table 280 on page 18393 lists the output fields for the `show chassis fpc ether-types` command.

Table 280: show chassis fpc ether-types Output Fields

Field Name	Field Description
FPC	FPC slot number.
INDEX	Index in the EtherTypes TCAM table. **NOTE**: Indexes 0-5 and 28-31 are reserved. You can't set custom EtherType values at those locations.

Table 280: show chassis fpc ether-types Output Fields *(Continued)*

Field Name	Field Description
DEFAULT	Default EtherType stored in TCAM at this index.
CONFIGURABLE	EtherType stored at this TCAM index is configurable (YES), or is reserved and can't be configured or changed (NO).
ACTIVE	Currently active EtherType value for this index. If you configure a custom value at a particular index, this column shows that value.

Sample Output

show chassis fpc ether-types

```
user@host> show chassis fpc ether-types
```

```
FPC      INDEX         DEFAULT        CONFIGURABLE       ACTIVE
  2        0             0x800              NO             0x800
           1             0x86dd             NO             0x86dd
           2             0x806              NO             0x806
           3             0x88cc             NO             0x88cc
           4             0x8809             NO             0x8809
           5             0x8035             NO             0x8035
           6             0x8847            YES             0x8847
           7             0x8848            YES             0x8848
           8             0x8906            YES             0x8906
           9             0x8914            YES             0x8914
          10             0x8902            YES             0x8902
          11             0x88f7            YES             0x88f7
          12             0x88b7            YES             0x88b7
          13             0x88f5            YES             0x88f5
          14             0x88f6            YES             0x88f6
          15             0x888e            YES             0x888e
          16             0x893f            YES             0x893f
          17             0x8940            YES             0x8940
          18             0x894b            YES             0x894b
          19             0x886f            YES             0x886f
```

	20	0x8922	YES	0x8922
	21	0x8850	YES	0x8850
	22	0x8808	YES	**0x88B5**
	23	0x8915	YES	0x8915
	24	0x889a	YES	0x889a
	25	0x88e5	YES	0x88e5
	26	0x88de	YES	0x88de
	27	0x88f8	YES	0x88f8
	28	0x8100	NO	0x8100
	29	0x88a8	NO	0x88a8
	30	0x9100	NO	0x9100
	31	0x9200	NO	0x9200
3	0	0x800	NO	0x800
	1	0x86dd	NO	0x86dd
	2	0x806	NO	0x806
	3	0x88cc	NO	0x88cc
	4	0x8809	NO	0x8809
	5	0x8035	NO	0x8035
	6	0x8847	YES	0x8847
	7	0x8848	YES	0x8848
	8	0x8906	YES	0x8906
	9	0x8914	YES	0x8914
	10	0x8902	YES	0x8902
	11	0x88f7	YES	0x88f7
	12	0x88b7	YES	0x88b7
	13	0x88f5	YES	0x88f5
	14	0x88f6	YES	0x88f6
	15	0x888e	YES	0x888e
	16	0x893f	YES	0x893f
	17	0x8940	YES	0x8940
	18	0x894b	YES	0x894b
	19	0x886f	YES	0x886f
	20	0x8922	YES	0x8922
	21	0x8850	YES	0x8850
	22	0x8808	YES	0x8808
	23	0x8915	YES	0x8915
	24	0x889a	YES	0x889a
	25	0x88e5	YES	0x88e5
	26	0x88de	YES	0x88de
	27	0x88f8	YES	0x88f8
	28	0x8100	NO	0x8100
	29	0x88a8	NO	0x88a8

```
            30            0x9100             NO             0x9100
            31            0x9200             NO             0x9200
```

show chassis fpc ether-types *slot-number*

```
user@host> show chassis fpc ether-types 0
```

```
FPC       INDEX          DEFAULT        CONFIGURABLE        ACTIVE
  0          0             0x800              NO              0x800
             1             0x86dd             NO              0x86dd
             2             0x806              NO              0x806
             3             0x88cc             NO              0x88cc
             4             0x8809             NO              0x8809
             5             0x8035             NO              0x8035
             6             0x8847            YES              0x8847
             7             0x8848            YES              0x8848
             8             0x8906            YES              0x8906
             9             0x8914            YES              0x8914
            10             0x8902            YES              0x8902
            11             0x88f7            YES              0x88f7
            12             0x88b7            YES              0x88b7
            13             0x88f5            YES              0x88f5
            14             0x88f6            YES              0x88f6
            15             0x888e            YES              0x888e
            16             0x893f            YES              0x893f
            17             0x8940            YES              0x8940
            18             0x894b            YES              0x894b
            19             0x886f            YES              0x886f
            20             0x8922            YES              0x8922
            21             0x8850            YES              0x8850
            22             0x8808            YES              0x8808
            23             0x8915            YES              0x8915
            24             0x889a            YES              0x8890
            25             0x88e5            YES              0x88e5
            26             0x88de            YES              0x88de
            27             0x88f8            YES              0x88f8
            28             0x8100             NO              0x8100
            29             0x88a8             NO              0x88a8
```

```
            30            0x9100            NO            0x9100
            31            0x9200            NO            0x9200
```

Release Information

Introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

ether-type (Chassis FPC) | 6477

show chassis fpc-feb-connectivity

IN THIS SECTION

- Syntax | 18397
- Description | 18397
- Options | 18398
- Required Privilege Level | 18398
- Output Fields | 18398
- Sample Output | 18399
- Release Information | 18400

Syntax

```
show chassis fpc-feb-connectivity
```

Description

(M120 router only) Display the Flexible PIC Concentrator (FPC) and Forwarding Engine Board (FEB) mapping and their respective states.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 281 on page 18398 lists the output fields for the `show chassis fpc-feb-connectivity` command. Output fields are listed in the approximate order in which they appear.

Table 281: show chassis fpc-feb-connectivity Output Fields

Field Name	Field Description
FPC	Slot number of the Flexible PIC Concentrator (FPC).
FPC type	Type of FPC: **Type 1**, **Type 2**, **Type 3**, or **cFPC**.
FPC state	State of the FPC. State can be any of the following: • **Announce offline**—Intermediate state where FPC is going down but is not offline and the Chassis manager acknowledges that the FPC is in the process of going offline. • **Announce online**—Intermediate state where FPC is coming up but is not online and the Chassis manager acknowledges that the FPC is in the process of coming online. • **Empty**—No FPC is present. • **Offline**—FPC is powered down. • **Online**—FPC is online and running. • **Present**—The chassis process has detected the FPC, but the FPC is either not supported by the current version of the Junos OS or FPC is coming up but is not online. • **Ready**—FPC is in transition state.
Connected FEB	Slot number of the Forwarding Engine Board (FEB) connected to the FPC or **None** if the FPC is not connected to a FEB.

Table 281: show chassis fpc-feb-connectivity Output Fields *(Continued)*

Field Name	Field Description
FEB state	State of the FEB. State can be any of the following: • **Announce offline**—Intermediate state where FEB is going down but is not offline and the Chassis manager acknowledges that the FEB is in the process of going offline. • **Announce online**—Intermediate state where FEB is coming up but is not online and the Chassis manager acknowledges that the FEB is in the process of coming online. • **Empty**—No FEB is present. • **Offline**—FEB is powered down. • **Online**—FEB is online and running. • **Present**—The chassis process has detected the FEB, but the FEB is either not supported by the current version of the Junos OS or FEB is coming up but is not online. • **Ready**—FEB is in transition state.
Link status	Status of the link connecting the R-FEB and R-FPC: • **Error** • **Misconfiguration**—Configuration between the R-FEB and the F-FPC is incorrect. • **OK**

Sample Output

show chassis fpc-feb-connectivity

```
user@host> show chassis fpc-feb-connectivity

FPC  FPC type  FPC state       Connected FEB  FEB state       Link status
0    cFPC      Online          0              Empty
1    cFPC      Online          1              Online          OK
2    Type 3    Online          3              Online          OK
3    Type 2    Online          None
4    Type 1    Online          4              Online          OK
```

```
5    Type 3    Online          None
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 8 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Active alarms  : PLL, LOS, LINK
  Active defects : PLL, LOF, LOS, SEF, LOP, BERR-SF, PLM-P, LINK
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         3
  MAC statistics:                      Receive         Transmit
    Total octets                           0                0
    Total packets                          0                0
```

Release Information

Command introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

- request chassis fpc | **14475**
- show chassis fpc | **18311**
- show chassis fabric fpcs | **18079**
- *Resynchronizing FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
- MX960 Flexible PIC Concentrator Description

show chassis fpc global pfe-id

IN THIS SECTION

- Syntax | **18401**
- Description | **18401**

- Additional Information | 18401
- Required Privilege Level | 18401
- Output Fields | 18401
- Sample Output | 18402
- Release Information | 18404

Syntax

```
show chassis fpc global pfe-id
```

Description

Display the global pfe-id. The devices supported are MX10004, MX10008, MX10016, MX2010, and MX2020. The outputs are independent of the line card and the status of the FPC or PFE. This output is based on the chassis and the fabric card.

Additional Information

The output is displayed independent of the line card and the status of the FPC or PFE. This output is based on the chassis and the fabric card.

Required Privilege Level

view

Output Fields

Table 282 on page 18401 lists the output fields for the `show chassis fpc global pfe-id` command.

Table 282: show chassis fpc global pfe-id Output Fields

Field Name	Field Description
FPC	FPC slot number of the Flexible PIC Concentrator (FPC).

Table 282: show chassis fpc global pfe-id Output Fields *(Continued)*

Field Name	Field Description
Global PFE ID	Unique Identifier for a PFE in a chassis. Local PFE ID = Unique identifier for a PFE in an FPC.

Sample Output

show chassis fpc global pfe-id (MX10008)

```
user@host> show chassis fpc global pfe-id
FPC                                                Global PFE ID
     PFE0  PFE1  PFE2  PFE3  PFE4  PFE5  PFE6  PFE7  PFE8  PFE9  PFE10  PFE11  PFE12  PFE13
PFE14  PFE15
-----------------------------------------------------------------------------------------------
---------
0      0     1     2     3     4     5     6     7     8     9    10     11     12     13
14     15
1     16    17    18    19    20    21    22    23    24    25    26     27     28     29
30     31
2     32    33    34    35    36    37    38    39    40    41    42     43     44     45
46     47
3     48    49    50    51    52    53    54    55    56    57    58     59     60     61
62     63
4     64    65    66    67    68    69    70    71    72    73    74     75     76     77
78     79
5     80    81    82    83    84    85    86    87    88    89    90     91     92     93
94     95
6     96    97    98    99   100   101   102   103   104   105   106    107    108    109
110    111
7    112   113   114   115   116   117   118   119   120   121   122    123    124    125
126    127
-----------------------------------------------------------------------------------------------
---------
```

show chassis fpc global pfe-id (MX10004)

```
user@host> show chassis fpc global-pfe-id
-----------------------------------------------------------------------------------------------
---------
FPC                                                Global PFE ID
   PFE0  PFE1  PFE2  PFE3  PFE4  PFE5  PFE6  PFE7  PFE8  PFE9  PFE10  PFE11  PFE12  PFE13
PFE14  PFE15
-----------------------------------------------------------------------------------------------
---------
0   0     1     2     3     4     5     6     7     8     9    10     11     12     13
14     15
1  16    17    18    19    20    21    22    23    24    25    26     27     28     29
30     31
2  32    33    34    35    36    37    38    39    40    41    42     43     44     45
46     47
3  48    49    50    51    52    53    54    55    56    57    58     59     60     61
62     63
-----------------------------------------------------------------------------------------------
---------
```

show chassis fpc global pfe-id (MX2020 Routers)

```
user@host> show chassis fpc global pfe-id
FPC                       Global PFE ID
        PFE0  PFE1  PFE2  PFE3  PFE4  PFE5  PFE6  PFE7
-------------------------------------------------------
0          0     1     2     3    80    81    82    83
1          4     5     6     7    84    85    86    87
2          8     9    10    11    88    89    90    91
3         12    13    14    15    92    93    94    95
4         16    17    18    19    96    97    98    99
5         20    21    22    23   100   101   102   103
6         24    25    26    27   104   105   106   107
7         28    29    30    31   108   109   110   111
8         32    33    34    35   112   113   114   115
9         36    37    38    39   116   117   118   119
10        40    41    42    43   120   121   122   123
11        44    45    46    47   124   125   126   127
12        48    49    50    51   128   129   130   131
```

```
13        52    53    54    55   132   133   134   135
14        56    57    58    59   136   137   138   139
15        60    61    62    63   140   141   142   143
16        64    65    66    67   144   145   146   147
17        68    69    70    71   148   149   150   151
18        72    73    74    75   152   153   154   155
19        76    77    78    79   156   157   158   159
-------------------------------------------------------
```

show chassis fpc global pfe-id (MX2010 Routers)

```
user@host> show chassis fpc global pfe-id
FPC                       Global PFE ID
        PFE0  PFE1  PFE2  PFE3  PFE4  PFE5  PFE6  PFE7
-------------------------------------------------------
0          0     1     2     3    80    81    82    83
1          4     5     6     7    84    85    86    87
2          8     9    10    11    88    89    90    91
3         12    13    14    15    92    93    94    95
4         16    17    18    19    96    97    98    99
5         20    21    22    23   100   101   102   103
6         24    25    26    27   104   105   106   107
7         28    29    30    31   108   109   110   111
8         32    33    34    35   112   113   114   115
9         36    37    38    39   116   117   118   119
-------------------------------------------------------
```

Release Information

Command introduced in Junos OS Release 21.4R1.

show chassis fpc (GNF)

IN THIS SECTION

Syntax | 18405

- Description | 18405
- Required Privilege Level | 18405
- Output Fields | 18405
- Sample Output | 18407
- Release Information | 18407

Syntax

```
show chassis fpc
```

Description

Display information about the Flexible PIC Concentrators (fpcs) assigned to the guest network function (GNF).

Required Privilege Level

view

Output Fields

Table 283 on page 18406 lists the output fields for the `show chassis fpc` command. Output fields are listed in the approximate order in which they appear.

Table 283: show chassis fpc Output Fields

Field Name	Field Description
`Slot` or `Slot State`	Slot number and state. The state can be one of the following conditions: • `Dead`—Held in reset because of errors. • `Diag`—Slot is being ignored while the FPC is running diagnostics. • `Dormant`—Held in reset. • `Empty`—No FPC is present. • `Online`—FPC is online and running. • `Present`—FPC is detected by the chassis daemon but either is not supported by the current version of Junos OS or is inserted in the wrong slot. The output also states either `Hardware Not Supported` or `Hardware Not In Right Slot`. The FPC is coming up but not yet online. • `Probed`—Probe is complete; awaiting restart of the Packet Forwarding Engine. • `Probe-wait`—Waiting to be probed.
`Temp (C)` or `Temperature`	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.
`Total CPU Utilization (%)`	Total percentage of CPU being used by the FPC's processor.
`Interrupt CPU Utilization (%)`	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.
`1 min CPU utilization (%)`	Information about the Routing Engine's CPU utilization in the past 1 minute.
`5 min CPU utilization (%)`	Information about the Routing Engine's CPU utilization in the past 5 minutes.

Table 283: show chassis fpc Output Fields *(Continued)*

Field Name	Field Description
`15 min CPU utilization (%)`	Information about the Routing Engine's CPU utilization in the past 15 minutes.
`Memory DRAM (MB)`	Total DRAM, in megabytes, available to the FPC's processor.
`Heap Utilization (%)`	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there might be a software problem (memory leak). **NOTE**: On MX Series routers and EX Series switches in a broadband edge environment, heap utilization levels higher than 70 percent can affect unified ISSU, router stability, or scaling capability.
`Buffer Utilization (%)`	Percentage of buffer space being used by the FPC's processor for buffering internal messages.
GNF	GNF identifier associated with each FPC. (`pic-status` output only) GNF identifier associated with each PIC.

Sample Output

command-name

```
user@router>show chassis fpc
           Temp CPU Utilization(%) CPU Utilization(%) Memory Utilization(%)
Slot State  (C)  Total Interrupt    1min 5min  15min  DRAM(MB)  Heap  Buffer  GNF
  4  Online  42     20         0     19   19    19    3584        8     25     2
  6  Online  46     12         0     11   11    11    3136        8     19     2
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show chassis hardware (GNF) | 18453
show chassis fpc pic-status (GNF) | 18411
show chassis adc (GNF) | 17295

show chassis fpc (Node Slicing)

IN THIS SECTION

- Syntax | 18408
- Description | 18408
- Required Privilege Level | 18408
- Output Fields | 18408
- Sample Output | 18410
- Release Information | 18411

Syntax

```
show chassis fpc
```

Description

Display information about Flexible PIC Concentrators (fpcs) assigned to different guest network functions (GNFs).

Required Privilege Level

view

Output Fields

Table 284 on page 18409 lists the output fields for the `show chassis fpc` command. Output fields are listed in the approximate order in which they appear.

Table 284: show chassis fpc Output Fields

Field Name	Field Description
`Slot` or `Slot State`	Slot number and state. The state can be one of the following conditions: • `Dead`—Held in reset because of errors. • `Diag`—Slot is being ignored while the FPC is running diagnostics. • `Dormant`—Held in reset. • `Empty`—No FPC is present. • `Online`—FPC is online and running. • `Present`—FPC is detected by the chassis daemon but either is not supported by the current version of Junos OS or is inserted in the wrong slot. The output also states either `Hardware Not Supported` or `Hardware Not In Right Slot`. The FPC is coming up but not yet online. • `Probed`—Probe is complete; awaiting restart of the Packet Forwarding Engine. • `Probe-wait`—Waiting to be probed.
`Temp (C)` or `Temperature`	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.
`Total CPU Utilization (%)`	Total percentage of CPU being used by the FPC's processor.
`Interrupt CPU Utilization (%)`	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.
`1 min CPU Utilization (%)`	Information about the Routing Engine's CPU utilization in the past 1 minute.
`5 min CPU Utilization (%)`	Information about the Routing Engine's CPU utilization in the past 5 minutes.

Table 284: show chassis fpc Output Fields *(Continued)*

Field Name	Field Description
`15 min CPU Utilization (%)`	Information about the Routing Engine's CPU utilization in the past 15 minutes.
`Memory DRAM (MB)`	Total DRAM, in megabytes, available to the FPC's processor.
`Heap Utilization (%)`	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak). **NOTE**: On MX Series routers and EX Series switches in a broadband edge environment, heap utilization levels higher than 70 percent can affect unified ISSU, router stability, or scaling capability.
`Buffer Utilization (%)`	Percentage of buffer space being used by the FPC's processor for buffering internal messages.
`GNF`	GNF identifier associated with each FPC. (`pic-status` output only) GNF identifier associated with each PIC.

Sample Output

command-name

user@router> show chassis fpc

```
                Temp  CPU               CPU                  Memory
                      Utilization (%)   Utilization(%)       Utilization (%)
 Slot State     (C)   Total  Interrupt  1min   5min   15min DRAM(MB) Heap  Buffer GNF
 0  Online      45     12          0    12     12     12    3584      6      25   3
 1  Online      57     22          0    20     20     20    3136     16      22   2
 2  Online      50     19          0    17     17     16    3584      6      25   3
 3  Online      28     10          0    11     11     11    2048     10      20   6
 4  Online      42     20          0    20     19     19    3584      8      25   6
 5  Online      58     22          0    21     20     20    3136     16      22   4
 6  Online      49     17          0    15     16     16    3136     13      20   1
```

```
7  Online   44     11          0    10     10     10    3584     6      25   5
8  Online   40     19          0    18     18     18    3584     8      25   5
9  Online   44     19          0    20     20     20    3584     8      25   5
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show chassis network-slices | 18575
- show chassis fpc pic-status (Node Slicing) | 18413
- show chassis adc (Node Slicing) | 17297
- show chassis network-slices fpcs | 18579

show chassis fpc pic-status (GNF)

IN THIS SECTION

- Syntax | 18411
- Description | 18412
- Required Privilege Level | 18412
- Sample Output | 18412
- Release Information | 18412

Syntax

```
show chassis fpc pic-status
```

Description

Display the status of the physical interface cards (PICs) of each Flexible PIC Concentrator (FPC) assigned to the guest network function (GNF).

Required Privilege Level

view

Sample Output

command-name

```
user@router> show chassis fpc pic-status


Slot 4   Online       MPC6E 3D                                      GNF 2
  PIC 0  Online       24X10GE SFPP
  PIC 1  Online       4X100GE CXP
Slot 6   Online       MPC7E 3D MRATE-12xQSFPP-XGE-XLGE-CGE          GNF 2
  PIC 0  Online       MRATE-6xQSFPP-XGE-XLGE-CGE
  PIC 1  Online       MRATE-6xQSFPP-XGE-XLGE-CGE
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show chassis hardware (GNF) | **18453**

show chassis fpc (GNF) | **18404**

show chassis adc (GNF) | **17295**

show chassis fpc pic-status (Node Slicing)

IN THIS SECTION

- Syntax | 18413
- Description | 18413
- Required Privilege Level | 18413
- Sample Output | 18413
- Release Information | 18414

Syntax

```
show chassis fpc pic-status
```

Description

Display the status of the physical interface cards (PICs) of each Flexible PIC Concentrator (FPC) assigned to different guest network functions (GNFs).

Required Privilege Level

view

Sample Output

user@router> show chassis fpc pic-status

command-name

```
Slot 0   Online       MPC5E 3D 24XGE+6XLGE                          GNF 3
  PIC 0  Online       12X10GE SFPP OTN
  PIC 1  Offline      12X10GE SFPP OTN
  PIC 2  Offline      3X40GE QSFPP
```

```
  PIC 3  Online       3X40GE QSFPP
Slot 1   Online       MPC9E 3D                                    GNF 2
  PIC 1  Online       MRATE-12xQSFPP-XGE-XLGE-CGE
Slot 2   Online       MPC5E 3D Q 2CGE+4XGE                        GNF 3
  PIC 0  Online       2X10GE SFPP OTN
  PIC 1  Online       1X100GE CFP2 OTN
  PIC 2  Online       2X10GE SFPP OTN
  PIC 3  Online       1X100GE CFP2 OTN
Slot 3   Online       MPCE Type 2 3D EQ                           GNF 6
Slot 4   Online       MPC6E 3D                                    GNF 6
  PIC 0  Online       24X10GE SFPP
  PIC 1  Online       2X100GE CFP2 OTN
Slot 5   Online       MPC9E 3D                                    GNF 4
  PIC 0  Online       MRATE-12xQSFPP-XGE-XLGE-CGE
Slot 6   Online       MPC7E 3D MRATE-12xQSFPP-XGE-XLGE-CGE        GNF 1
  PIC 0  Online       MRATE-6xQSFPP-XGE-XLGE-CGE
  PIC 1  Online       MRATE-6xQSFPP-XGE-XLGE-CGE
Slot 7   Online       MPC5E 3D 2CGE+4XGE                          GNF 5
  PIC 0  Online       2X10GE SFPP OTN
  PIC 1  Online       1X100GE CFP2 OTN
  PIC 2  Online       2X10GE SFPP OTN
  PIC 3  Online       1X100GE CFP2 OTN
Slot 8   Online       MPC6E 3D                                    GNF 5
  PIC 0  Online       24X10GE SFPP OTN
Slot 9   Online       MPC6E 3D                                    GNF 5
  PIC 0  Online       24X10GE SFPP
  PIC 1  Online       4X100GE CXP
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show chassis network-slices | 18575
- show chassis fpc (Node Slicing) | 18408
- show chassis adc (Node Slicing) | 17297
- show chassis network-slices fpcs | 18579

show chassis fpc (View)

IN THIS SECTION

- Syntax | 18415
- Description | 18415
- Options | 18415
- Required Privilege Level | 18416
- Output Fields | 18416
- Sample Output | 18418
- Sample Output | 18419
- Sample Output | 18420
- Sample Output | 18423
- Release Information | 18425

Syntax

```
show chassis fpc
<detail <fpc-slot >| <node (node-id | local | primary)>> |
<node (node-id | local | primary)> |
<pic-status <fpc-slot >| <node (node-id | local | primary)>>
```

Description

Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.

Options

- none—Display status information for all FPCs.
- `detail`—(Optional) Display detailed FPC status information.
- *`fpc-slot`* —(Optional) Display information about the FPC in this slot.

- `node`—(Optional) For chassis cluster configurations, display status information for all FPCs or for the specified FPC on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.
- `pic-status`—(Optional) Display status information for all FPCs or for the FPC in the specified slot (see *`fpc-slot`*).

Required Privilege Level

view

Output Fields

Table 285 on page 18417 lists the output fields for the `show chassis fpc` command. Output fields are listed in the approximate order in which they appear.

Table 285: show chassis fpc Output Fields

Field Name	Field Description
`Slot or Slot State`	Slot number and state. The state can be one of the following conditions: • `Dead`—Held in reset because of errors. • `Diag`—Slot is being ignored while the device is running diagnostics. • `Dormant`—Held in reset. • `Empty`—No FPC is present. • `Online`—FPC is online and running. • `Offline`—FPC configured offline. • `Present`—FPC is detected by the device, but is either not supported by the current version of Junos OS or inserted in the wrong slot. The output also states either `Hardware Not Supported` or `Hardware Not In Right Slot`. FPC is coming up but not yet online. • `Probed`—Probe is complete; awaiting restart of the Packet Forwarding Engine (PFE). • `Probe-wait`—Waiting to be probed.
`Temp (C) or Temperature`	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.
`Total CPU Utilization (%)`	Total percentage of CPU being used by the FPC's processor.
`Interrupt CPU Utilization (%)`	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.
`Memory DRAM (MB)`	Total DRAM, in megabytes, available to the FPC's processor.
`Heap Utilization (%)`	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak).

Table 285: show chassis fpc Output Fields *(Continued)*

Field Name	Field Description
`Buffer Utilization (%)`	Percentage of buffer space being used by the FPC's processor for buffering internal messages.
`Start Time`	Time when the Routing Engine detected that the FPC was running.
`Uptime`	How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running.
`PIC type`	(`pic-status` output only) Type of FPC.

Sample Output

show chassis fpc

```
user@host> show chassis fpc
                              Temp      CPU Utilization (%)     Memory    Utilization (%)
     Slot State         (C)       Total  Interrupt        DRAM (MB) Heap  Buffer
      0  Online          -------------------- CPU less FPC --------------------
      1  Online          ---------------------  Not Usable --------------------
      2  Online          -------------------- CPU less FPC --------------------
```

show chassis fpc (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc
                      Temp      CPU Utilization (%)       Memory    Utilization (%)
Slot State            (C)      Total  Interrupt            DRAM (MB) Heap     Buffer
  0  Empty
  1  Empty
  2  Empty
  3  Online            37          3          0             1024        7     42
  4  Empty
  5  Empty
  6  Online            30           8         0               1024        23    30
```

```
  7  Empty
  8  Empty
  9  Empty
 10  Empty
 11  Empty
```

show chassis fpc (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3)

```
user@host> show chassis fpc
                      Temp  CPU Utilization (%)   CPU Utilization (%)
Memory
                             Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM
(MB)
                        Heap     Buffer
  0  Online            36     20          0       20     19     19
1024
                              4         26
  1  Online            35      8          0        8      8      8
2048
                             12         14
  2  Online            40     21          0       20     20     20
3584
                              5         13
```

Sample Output

show chassis fpc detail 2

```
user@host> show chassis fpc detail 2
Slot 2 information:
  State                                 Online
  Temperature                            37
  Total CPU DRAM                      1024 MB
  Total RLDRAM                         0 MB
  Total DDR DRAM                       0 MB
  Start time:                           2012-07-18 07:18:50 PDT
```

```
  Uptime:                                4 days, 21 hours, 51 minutes, 59 seconds
  Max Power Consumption               0 Watts
```

Sample Output

show chassis fpc pic-status (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc pic-status
Slot 3   Online       SRX5k SPC
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
Slot 6   Online       SRX5k DPC 4x 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SPC2)

```
user@host> show chassis fpc pic-status

Slot 0   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 2   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 3   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SRX5K-MPC)

```
user@host> show chassis fpc pic-status

Slot 0   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Online       SRX5k DPC 4X 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
Slot 6   Offline      SRX5k SPC II
Slot 9   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 10  Online       SRX5k IOC II
  PIC 0  Online       10x 10GE SFP+
  PIC 2  Online       1x 100GE CFP
Slot 11  Online       SRX5k IOC II
  PIC 0  Online       1x 100GE CFP
  PIC 2  Online       2x 40GE QSFP+
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices when Express Path [formerly known as services offloading] is configured)

```
user@host> show chassis fpc pic-status

Slot 0   Offline      SRX5k DPC 40x 1GE
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
```

```
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Offline      SRX5k SPC
Slot 4   Online       SRX5k IOC3 24XGE+6XLG
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 5   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 1  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
```

show chassis fpc pic-status (with 20-Gigabit Ethernet MIC with SFP)

```
user@host> show chassis fpc pic-status

node0:
--------------------------------------------------------------------------
Slot 0   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 1   Offline      SRX5k SPC II
Slot 2   Online       SRX5k DPC 4X 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
Slot 9   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP
  PIC 1  Online       10x 1GE(LAN) SFP
  PIC 2  Online       10x 1GE(LAN) SFP
  PIC 3  Online       10x 1GE(LAN) SFP
Slot 10  Online       SRX5k IOC II
  PIC 0  Online       10x 10GE SFP+
  PIC 2  Online       1x 100GE CFP
Slot 11  Offline      SRX5k IOC II
```

show chassis fpc pic-status (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3 and when Express Path [formerly known as services offloading] is configured)

```
user@host> show chassis fpc pic-status
Slot 0   Offline      SRX5k DPC 40x 1GE
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Offline      SRX5k SPC
Slot 4   Online       SRX5k IOC3 24XGE+6XLG
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 5   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 1  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
```

Sample Output

show chassis fpc pic-status for HA (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc pic-status
node0:
--------------------------------------------------------------
Slot 4   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow

node1:
---------------------------------------------------------------
Slot 4   Online       SRX5k DPC 40x 1GE
```

```
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow
```

show chassis fpc pic-status for HA (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3)

```
user@host> show chassis fpc pic-status
user@host> show chassis fpc pic-status
node0:
--------------------------------------------------------------------------
Slot 2   Online       SRX5k IOC3 24XGE+6XLG
  PIC 0  Online       12x 10GE SFP+
  PIC 1  Online       12x 10GE SFP+
  PIC 2  Offline      3x 40GE QSFP+
  PIC 3  Offline      3x 40GE QSFP+
Slot 4   Online       SRX5k IOC II
  PIC 2  Online       10x 10GE SFP+
Slot 5   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Offline
  PIC 3  Offline

node1:
--------------------------------------------------------------------------
Slot 2   Online       SRX5k IOC3 24XGE+6XLG
  PIC 0  Online       12x 10GE SFP+
  PIC 1  Online       12x 10GE SFP+
  PIC 2  Offline      3x 40GE QSFP+
  PIC 3  Offline      3x 40GE QSFP+
Slot 4   Online       SRX5k IOC II
  PIC 2  Online       10x 10GE SFP+
Slot 5   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
```

```
PIC 2  Offline
PIC 3  Offline
```

Release Information

Command modified in Junos OS Release 9.2.

Starting with Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) are introduced.

> **NOTE**: On SRX5K-MPC3-40G10G (IOC3), all four PICs cannot be powered on. A maximum of two PICs can be powered on at the same time. By default, PIC0 and PIC1 are online.

Use the `set chassis fpc <slot> pic <pic> power off` command to choose the PICs you want to power on.

When you use the `set chassis fpc <slot> pic <pic> power off` command to power off PIC0 and PIC1, PIC2 and PIC3 are automatically turned on.

When you switch from one set of PICs to another set of PICs using the `set chassis fpc <slot> pic <pic> power off` command again, ensure that there is 60 seconds duration between the two actions, otherwise core files are seen during the configuration.

The Table 286 on page 18425 summarizes the SRX5K-MPC3-40G10G (IOC3) PICs selected for various configuration scenarios.

Table 286: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary

CLI Configuration	PIC Selection
Default (i.e. no CLI configuration)	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-1, PIC-2 and PIC-3 powered OFF	Online: PIC-0 Offline: PIC-1, PIC-2, PIC-3
PIC-0, PIC-2 and PIC-3 powered OFF	Online: PIC-1 Offline: PIC-0, PIC-2, PIC-3

Table 286: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary *(Continued)*

CLI Configuration	PIC Selection
PIC-0, PIC-1 and PIC-3 powered OFF	Online: PIC-2 Offline: PIC-0, PIC-1, PIC-3
PIC-0, PIC-1 and PIC-2 powered OFF	Online: PIC-3 Offline: PIC-0, PIC-1, PIC-2
PIC-2 and PIC-3 powered OFF	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-2 and PIC-3 powered OFF	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-1 and PIC-2 powered OFF	Online: PIC-0, PIC-3 Offline: PIC-1, PIC-2
PIC-0 and PIC-3 powered OFF	Online: PIC-2, PIC-1 Offline: PIC-0, PIC-3
PIC-0 and PIC-1 powered OFF	Online: PIC-2, PIC-3 Offline: PIC-0, PIC-1
All other combinations of PICs being powered OFF (Invalid)	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3 Default PICs will be selected for the invalid combinations. Also, a system log message will be displayed to indicate the invalid combination PIC selection.

RELATED DOCUMENTATION

Understanding Interfaces

show chassis gnf

IN THIS SECTION

- Syntax | 18427
- Description | 18427
- Required Privilege Level | 18427
- Output Fields | 18427
- Sample Output | 18429
- Release Information | 18429

Syntax

```
show chassis gnf
```

Description

Display information about the guest network function (GNF) you logged in.

Required Privilege Level

view

Output Fields

Table 287 on page 18428 lists the output fields for the `show chassis gnf` command. Output fields are listed in the approximate order in which they appear.

Table 287: show chassis gnf Output Fields

Field Name	Field Description
GNF ID	Shows the GNF ID.
`GNF State`	Status of the GNF. • `Online`—GNF online and running. • `Offline`—GNF is offline.
GNF description	Description of the guest network function.
FPCs assigned	The FPC slot numbers assigned to the GNF.
FPCs online	The FPCs that are online.
BSYS	Name of the BSYS.
BSYS sw version	The Junos software version used in the BSYS.
GNF sw version	The Junos software version used in the GNF.
BSYS primary RE	Primary Routing Engine slot.
`GNF uptime`	Duration for which the GNF has been up and running.
GNF Routing Engine Status:	The details of the Routing Engines in the slot 0 and 1. The details include the Current State, (primary or backup), Routing Engine model, and GNF host name.

Sample Output

show chassis gnf

```
user@router> show chassis gnf

GNF ID           1
GNF description  NA
GNF state        Online
FPCs assigned    8 9
FPCs online      8 9
BSYS             router(mx960)
BSYS sw version  18.2-20180321_0948_bsys
GNF  sw version  18.2-20180314_1035_gnf
Chassis          mx960
BSYS master RE   0
GNF uptime       54 minutes, 37 seconds
GNF Routing Engine Status:
Slot 0:
    Current state   Master
    Model           RE-GNF-2100x4
    GNF host name   gnf-host0
Slot 1:
    Current state   Backup
    Model           RE-GNF-2100x4
    GNF host name   gnf-host1
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

- show chassis hardware (GNF) | 18453
- show chassis fpc pic-status (GNF) | 18411
- show chassis adc (GNF) | 17295
- show chassis network-slices | 18575

show chassis hardware

IN THIS SECTION

- Syntax | 18430
- Syntax (EX Series Devices, MX104, MX204, MX2010, MX2020, MX10003, MX10008, and MX2008 Universal Routing Platforms) | 18430
- Syntax (TX Matrix Router) | 18431
- Syntax (TX Matrix Plus Router) | 18431
- Syntax (MX Series Devices) | 18431
- Syntax (QFX Series) | 18431
- Description | 18432
- Options | 18432
- Additional Information | 18434
- Required Privilege Level | 18434
- Output Fields | 18435
- Sample Output | 18438
- Release Information | 18453

Syntax

```
show chassis hardware
<detail | extensive>
<clei-models>
<models>
```

Syntax (EX Series Devices, MX104, MX204, MX2010, MX2020, MX10003, MX10008, and MX2008 Universal Routing Platforms)

```
show chassis hardware
<clei-models>
<detail | extensive>
```

```
<models>
<satellite [slot-id slot-id |device-alias  alias-name]>
```

Syntax (TX Matrix Router)

```
show chassis hardware
<clei-models>
<detail | extensive>
<models>
<lcc number |  scc>
```

Syntax (TX Matrix Plus Router)

```
show chassis hardware
<clei-models>
<detail | extensive>
<models>
<lcc number |  sfc number>
```

Syntax (MX Series Devices)

```
show chassis hardware
<detail | extensive>
<clei-models>
<models>
<all-members>
<local>
<member member-id>
```

Syntax (QFX Series)

```
show chassis hardware
<detail | extensive>
<clei-models>
<interconnect-device name>
```

```
<node-device name>
<models>
```

Description

Display a list of all Flexible PIC Concentrators (FPCs) and PICs installed in the router or switch chassis, including the hardware version level and serial number.

In the EX Series switch command output, FPC refers to the following:

- On EX2200 switches, EX3200 switches, EX4200 standalone switches, and EX4500 switches—Refers to the switch; FPC *number* is always 0.
- On EX4200 switches in a Virtual Chassis configuration—Refers to the member of a Virtual Chassis; FPC *number* equals the member ID, from 0 through 9.
- On EX8208 and EX8216 switches—Refers to a line card; FPC *number* equals the slot number for the line card.

On QFX3500, QFX5100, and OCX Series standalone switches, and PTX1000 devices both the FPC and FPC *number* are always `0`.

On T4000 Type 5 FPCs, there are no `top temperature sensor` or `bottom temperature sensor` parameters. Instead, `fan intake temperature sensor` and `fan exhaust temperature sensors` parameters are displayed.

Starting from Junos OS Release 11.4, the output of the `show chassis hardware models` operational mode command displays the enhanced midplanes FRU model numbers (CHAS-BP3-MX240-S, CHAS-BP3-MX480-S or CHAS-BP3-MX960-S) based on the router. Prior to release 11.4, the FRU model numbers are left blank when the router has enhanced midplanes. Note that the enhanced midplanes are introduced through the Junos OS Release 13.3, but can be supported on all Junos OS releases.

Starting with Junos OS Release 14.1, the output of the `show chassis hardware detail | extensive | clei-models | models` operational mode command displays the new DC power supply module (PSM) and power distribution unit (PDU) that are added to provide power to the high-density FPC (FPC2-PTX-P1A) and other components in a PTX5000 Packet Transport Router.

Options

none Display information about hardware. For a TX Matrix router, display information about the TX Matrix router and its attached T640 devices. For a TX Matrix Plus router, display information about the TX Matrix Plus router and its attached devices.

clei-models (Optional) Display Common Language Equipment Identifier (CLEI) barcode and model number for orderable field-replaceable units (FRUs).

detail (Optional) Include RAM and disk information in output.

extensive (Optional) Display ID EEPROM information.

all-members (MX Series devices only) (Optional) Display hardware-specific information for all the members of the Virtual Chassis configuration.

interconnect-device *name* (QFabric systems only) (Optional) Display hardware-specific information for the Interconnect device.

lcc *number* (TX Matrix devices and TX Matrix Plus router only) (Optional) On a TX Matrix router, display hardware information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display hardware information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 devices are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 devices are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 devices are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 devices are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series devices only) (Optional) Display hardware-specific information for the local Virtual Chassis members.

member *member-id* (MX Series devices and EX Series switches) (Optional) Display hardware-specific information for the specified member of the Virtual Chassis configuration. Replace *member-id* variable with a value 0 or 1.

models (Optional) Display model numbers and part numbers for orderable FRUs and, for components that use ID EEPROM format v2, the CLEI code.

node-device *name* (QFabric systems only) (Optional) Display hardware-specific information for the Node device.

satellite [slot-id *slot-id* | device-alias *alias-name*]	(Junos Fusion only) (Optional) Display hardware information for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.
scc	(TX Matrix router only) (Optional) Display hardware information for the TX Matrix router (switch-card chassis).
sfc *number*	(TX Matrix Plus router only) (Optional) Display hardware information for the TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with `0`.

Additional Information

The `show chassis hardware detail` command now displays DIMM information for the following Routing Engines, as shown in Table 288 on page 18434.

Table 288: Routing Engines Displaying DIMM Information

Routing Engines	Devices
RE-S-1800x2 and RE-S-1800x4	MX240, MX480, and MX960 devices
RE-A-1800x2	M120 and M320 devices

In Junos OS Release 11.4 and later, the output for the `show chassis hardware models` operational mode command for MX Series devices display the enhanced midplanes FRU model numbers—CHAS-BP3-MX240-S, CHAS-BP3-MX480-S, or CHAS-BP3-MX960-S—based on the router. In releases before Junos OS Release 11.4, the FRU model numbers are left blank when the router has enhanced midplanes. Note that the enhanced midplanes are introduced through Junos OS Release 13.3, but can be supported on all Junos OS releases.

Starting with Junos OS Release 17.3R1, the output of the `show chassis hardware` command displays the mode in which vMX is running (performance mode or lite mode) in the part number field for the FPC. `RIOT-PERF` indicates performance mode and `RIOT-LITE` indicates lite mode.

Starting with Junos OS Release 22.2R1, the RE-S-X6-128G-K Routing Engine (RE) is supported for MX240, MX480, and MX960 devices. View the details of the RE in the command output.

Required Privilege Level

view

Output Fields

Table 289 on page 18435 lists the output fields for the `show chassis hardware` command. Output fields are listed in the approximate order in which they appear.

Table 289: show chassis hardware Output Fields

Field Name	Field Description	Level of Output
`Item`	Show information about the device hardware.	All levels
`Version`	Revision level of the chassis component.	All levels
`Part number`	Part number of the chassis component.	All levels
`Serial number`	Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis.	All levels
`Assb ID` or `Assembly ID`	(`extensive` keyword only) Identification number that describes the FRU hardware.	`extensive`
`Assembly Version`	(`extensive` keyword only) Version number of the FRU hardware.	`extensive`
`Assembly Flags`	(`extensive` keyword only) Flags.	`extensive`
`FRU model number`	(`clei-models`, `extensive`, and `models` keyword only) Model number of the FRU hardware component.	none specified
`CLEI code`	(`clei-models` and `extensive` keyword only) Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.	none specified
`EEPROM Version`	ID EEPROM version used by the hardware component: `0x00` (version 0), `0x01` (version 1), or `0x02` (version 2).	`extensive`

Table 289: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Description`	Brief description of the hardware item: • Type of power supply. • Type of PIC. If the PIC type is not supported on the current software release, the output states `Hardware Not Supported`. • Type of FPC: `FPC Type 1`, `FPC Type 2`, `FPC Type 3`, `FPC Type 4` , or `FPC TypeOC192`. On EX Series switches, a brief description of the FPC. The following list shows the PIM abbreviation in the output and the corresponding PIM name. • `2x FE`—Either two built-in Fast Ethernet interfaces (fixed PIM) or dual-port Fast Ethernet PIM • `4x FE`—4-port Fast Ethernet ePIM • `1x GE Copper`—Copper Gigabit Ethernet ePIM (one 10-Mbps, 100-Mbps, or 1000-Mbps port) • `1x GE SFP`—SFP Gigabit Ethernet ePIM (one fiber port) • `2x Serial`—Dual-port serial PIM • `2x T1`—Dual-port T1 PIM • `2x E1`—Dual-port E1 PIM • `2x CT1E1`—Dual-port channelized T1/E1 PIM • `1x T3`—T3 PIM (one port) • `1x E3`—E3 PIM (one port) • `4x BRI S/T`—4-port ISDN BRI S/T PIM • `4x BRI U`—4-port ISDN BRI U PIM • `1x ADSL Annex A`—ADSL 2/2+ Annex A PIM (one port, for POTS)	All levels

Table 289: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `1x ADSL Annex B`—ADSL 2/2+ Annex B PIM (one port, for ISDN) • `2x SHDSL (ATM)`—G SHDSL PIM (2-port two-wire module or 1-port four-wire module) • `1x TGM550`—TGM550 Telephony Gateway Module (Avaya VoIP gateway module with one console port, two analog `LINE` ports, and two analog `TRUNK` ports) • `1x DS1 TIM510`—TIM510 E1/T1 Telephony Interface Module (Avaya VoIP media module with one E1 or T1 trunk termination port and ISDN PRI backup) • `4x FXS, 4xFX0, TIM514`—TIM514 Analog Telephony Interface Module (Avaya VoIP media module with four analog `LINE` ports and four analog `TRUNK` ports) • `4x BRI TIM521`—TIM521 BRI Telephony Interface Module (Avaya VoIP media module with four ISDN BRI ports) • `Crypto Accelerator Module`—For enhanced performance of cryptographic algorithms used in IP Security (IPsec) services • `MPC M 16x 10GE`—16-port 10-Gigabit Module Port Concentrator that supports SFP+ optical transceivers. (Not on EX Series switches.) • For hosts, the Routing Engine type. • For small form-factor pluggable transceiver (SFP) modules, the type of fiber: `LX`, `SX`, `LH`, or `T`. • LCD description for EX Series switches (except EX2200 switches). • `MPC2`—1-port MPC2 that supports two separate slots for MICs. • `MPC3E`—1-port MPC3E that supports two separate slots for MICs (MIC-3D-1X100GE-CFP and MIC-3D-20GE-SFP) on MX960, MX480, and MX240 devices. The MPC3E maps one MIC to one PIC (1 MIC, 1 PIC), which differs from the mapping of legacy MPCs. • 100GBASE-LR4, pluggable CFP optics	

Table 289: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• Supports the Enhanced MX Switch Control Board with fabric redundancy and existing SCBs without fabric redundancy. • Interoperates with existing MX Series line cards, including Flexible Port Concentrators (FPC), Dense Port Concentrators (DPCs), and Modular Port Concentrators (MPCs). • `MPC4E`—Fixed configuration MPC4E that is available in two flavors: MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE on MX2020, MX960, MX480, and MX240 devices. • LCD description for MX Series devices	

Sample Output

show chassis hardware (MX240, MX480, MX960 Devices)

```
user@host> show chassis hardware

Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1230686AFB      MX480
Midplane         REV 05   710-017414   ACRB7717          MX480 Midplane
FPM Board        REV 02   710-017254   CADF2017          Front Panel Display
PEM 0            Rev 03   740-022697   QCS1142C0HJ       PS 1.2-1.7kW; 100-240V AC in
PEM 1            Rev 01   740-022697   QCS1035C0CM       PS 1.2-1.7kW; 100-240V AC in
Routing Engine 0 REV 06   711-145342   CASE0160          RE-S-X6-128G-K
CB 0             REV 05   750-055976   CAES7891          Enhanced MX SCB 2
CB 1             REV 12   750-062572   CALM3310          Enhanced MX SCB 2
Fan Tray                                                  Enhanced Left Fan Tray
```

show chassis hardware (MX10008 Router)

```
user@host> show chassis hardware

Hardware inventory:
```

```
Item             Version  Part number  Serial number     Description
Chassis                               DE538             JNP10008 [MX10008]
Midplane         REV 27   750-054097   ACPD6954          Midplane 8
Routing Engine 0          BUILTIN      BUILTIN           RE X10
Routing Engine 1          BUILTIN      BUILTIN           RE X10 128
CB 0             REV 10   750-079562   CAKF2158          Control Board
CB 1             REV 05   711-065897   CAJG2680          Control Board
FPC 1            REV 04   750-084779   CAKN5706          JNP10K-LC2101
  CPU            REV 05   750-073391   CAKJ2864          LC 2101 PMB
  PIC 0                   BUILTIN      BUILTIN           4xQSFP28 SYNCE
    Xcvr 0       REV 01   740-046565   XXL0BQM           QSFP+-40G-SR4
    Xcvr 1       REV 01   740-032986   QB350242          QSFP+-40G-SR4
    Xcvr 2       REV 01   740-054053   QE408285          QSFP+-4X10G-SR
    Xcvr 3       REV 01   740-046565   QF3300Z9          QSFP+-40G-SR4
  PIC 1                   BUILTIN      BUILTIN           4xQSFP28 SYNCE
    Xcvr 0       REV 01   740-067442   QJ2200LD          QSFP+-40G-SR4
    Xcvr 1       REV 01   740-038153   APF170500382DP    QSFP+-40G-CU3M
    Xcvr 2       REV 01   740-067442   QI4302LC          QSFP+-40G-SR4
  PIC 2                   BUILTIN      BUILTIN           4xQSFP28 SYNCE
    Xcvr 0       REV 01   740-067442   1ACP1335119       QSFP+-40G-SR4
    Xcvr 1       REV 01   740-067442   1ACP1313156       QSFP+-40G-SR4
    Xcvr 2       REV 01   740-067442   QK050040          QSFP+-40G-SR4
    Xcvr 3       REV 01   740-067442   QJ2201BG          QSFP+-40G-SR4
  PIC 3                   BUILTIN      BUILTIN           4xQSFP28 SYNCE
    Xcvr 0       REV 01   740-058734   1ECQ12400CS       QSFP-100GBASE-SR4
    Xcvr 1       REV 01   740-046565   QF3300ZX          QSFP+-40G-SR4
    Xcvr 2       REV 01   740-061405   1ECQ12510FH       QSFP-100G-SR4-T2
    Xcvr 3       REV 01   740-032986   QB491182          QSFP+-40G-SR4
  PIC 4                   BUILTIN      BUILTIN           4xQSFP28 SYNCE
    Xcvr 0       REV 01   740-067442   QJ2200D5          QSFP+-40G-SR4
    Xcvr 1       REV 01   740-054053   XXS0L95           QSFP+-4X10G-SR
  PIC 5                   BUILTIN      BUILTIN           4xQSFP28 SYNCE
    Xcvr 0       REV 01   740-054053   QE251550          QSFP+-4X10G-SR
    Xcvr 1       REV 01   740-054053   XZB01D5           QSFP+-4X10G-SR
    Xcvr 3       REV 01   740-046565   QI1402F9          QSFP+-40G-SR4
FPD Board        REV 07   711-054687   ACPF2896          Front Panel Display
PEM 1            REV 02   740-049388   1EDL62102PR       Power Supply AC
PEM 2            REV 02   740-049388   1EDL60300H2       Power Supply AC
PEM 4            REV 02   740-049388   1EDL603003Z       Power Supply AC
PEM 5            REV 01   740-049388   1EDL339001B       Power Supply AC
FTC 0            REV 14   750-050108   ACNW3344          Fan Controller 8
FTC 1            REV 14   750-050108   ACPE3978          Fan Controller 8
Fan Tray 0       REV 09   760-054372   ACNV5507          Fan Tray 8
```

```
Fan Tray 1       REV 09   760-054372   ACNV5371          Fan Tray 8
SFB 0            REV 25   750-050058   ACPH6821          Switch Fabric (SIB) 8
SFB 1
SFB 2            REV 24   750-050058   ACNZ0641          Switch Fabric (SIB) 8
SFB 3            REV 27   750-050058   ACPH9127          Switch Fabric (SIB) 8
SFB 5            REV 24   750-050058   ACNX7396          Switch Fabric (SIB) 8
```

show chassis hardware (MX960 Router with MPC10E-10C-MRATE, MPC10E-15C-MRATE line cards)

```
user@host> show chassis hardware

Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1203BEAAFB      MX480
Midplane         REV 05   710-017414   ACAB6444          MX480 Midplane
FPM Board        REV 02   710-017254   CACG2646          Front Panel Display
PEM 0            Rev 02   740-063046   QCS2207U0AJ       PS 1.4-2.52kW; 90-264V AC in
PEM 1            Rev 02   740-063046   QCS2207U075       PS 1.4-2.52kW; 90-264V AC in
Routing Engine 0 REV 06   750-072925   EAAG9532          RE-S-2X00x6
Routing Engine 1 REV 06   750-072925   EAAD4894          RE-S-2X00x6
CB 0             REV 40   750-070866   EAAV1128          Enhanced MX SCB 3
CB 1             REV 40   750-070866   EAAT2273          Enhanced MX SCB 3
FPC 0            REV 60   750-070395   EAAT9462          MPC10E 3D MRATE-15xQSFPP
  CPU            REV 21   750-072571   EAAL4646          FMPC PMB
  PIC 0                   BUILTIN      BUILTIN           MRATE-5xQSFPP
  PIC 1                   BUILTIN      BUILTIN           MRATE-5xQSFPP
    Xcvr 4       REV 01   740-157132   2E1CZFA728016     QSFP56-DD-400G-ZR
  PIC 2                   BUILTIN      BUILTIN           MRATE-5xQSFPP
FPC 1            REV 51   750-078633   EAAY6176          MPC10E 3D MRATE-10xQSFPP
  CPU            REV 21   750-072571   EAAY5228          FMPC PMB
  PIC 0                   BUILTIN      BUILTIN           MRATE-5xQSFPP
Xcvr 4       REV 01   740-114884   1T1TZFA740012     QSFP56-DD-400G-ZR
  PIC 1                   BUILTIN      BUILTIN           MRATE-5xQSFPP
    Xcvr 4       REV 01   740-157132   2E1CZFA728022     QSFP56-DD-400G-ZR
FPC 2            REV 60   750-070395   EAAY6079          MPC10E 3D MRATE-15xQSFPP
  CPU            REV 21   750-072571   EAAB6925          FMPC PMB
  PIC 0                   BUILTIN      BUILTIN           MRATE-5xQSFPP
    Xcvr 4       REV 01   740-114884   1T1TZFA740010     QSFP56-DD-400G-ZR
  PIC 1                   BUILTIN      BUILTIN           MRATE-5xQSFPP
    Xcvr 4       REV 01   740-114884   1T1TZFA740005     QSFP56-DD-400G-ZR
```

```
  PIC 2                   BUILTIN      BUILTIN           MRATE-5xQSFPP
    Xcvr 4       REV 01   740-114884   1T1TZFA74000Z     QSFP56-DD-400G-ZR
```

show chassis hardware (PTX10008 Router)

```
root@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                DK404             JNP10008 [PTX10008]
Midplane         REV 28   750-054097   ACPP2394          Midplane 8
Routing Engine 0          BUILTIN      BUILTIN           RE-PTX-2X00x4
Routing Engine 1          BUILTIN      BUILTIN           RE-PTX-2X00x4
CB 0             REV 04   750-068820   ACPT5303          Control Board
CB 1             REV 04   750-068820   ACPR1627          Control Board
FPC 6            REV 13   750-068822   ACPB2753          LC1102 - 12C / 36Q /
14                                                                                                  4X
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           12x100GE/36x40GE/
144x10                                                                                                 GE
    Xcvr 0       REV 01   740-067442   XV304N6           QSFP+-40G-SR4
    Xcvr 1       REV 01   740-067442   XV30A5M           QSFP+-40G-SR4
    Xcvr 2       REV 01   740-067442   XV300HC           QSFP+-40G-SR4
    Xcvr 3       REV 01   740-067443   XU20L17           QSFP+-40G-SR4
    Xcvr 4       REV 01   740-067442   XV303XG           QSFP+-40G-SR4
    Xcvr 5       REV 01   740-067443   XV306QC           QSFP+-40G-SR4
    Xcvr 6       REV 01   740-067442   XV303Y7           QSFP+-40G-SR4
    Xcvr 7       REV 01   740-067443   XX60DMR           QSFP+-40G-SR4
    Xcvr 9       REV 01   740-067443   XX60DNY           QSFP+-40G-SR4
    Xcvr 10      REV 01   740-054053   QF4605WF          QSFP+-4X10G-SR
    Xcvr 13      REV 01   740-058734   1ECQ115007D       QSFP-100GBASE-SR4
    Xcvr 15      REV 01   740-046565   QH06035R          QSFP+-40G-SR4
    Xcvr 16      REV 01   740-046565   QH0602KC          QSFP+-40G-SR4
    Xcvr 17      REV 01   740-046565   QH0507PA          QSFP+-40G-SR4
    Xcvr 18      REV 01   740-046565   QH06035M          QSFP+-40G-SR4
    Xcvr 24      REV 01   740-046565   QH0507QL          QSFP+-40G-SR4
    Xcvr 25      REV 01   740-067443   XV20CWP           QSFP+-40G-SR4
    Xcvr 34      REV 01   740-046565   QH06035U          QSFP+-40G-SR4
    Xcvr 35      REV 01   740-067443   XX60DN9           QSFP+-40G-SR4
FPC 7            REV 41   750-051357   ACPL3446          LC1101 - 30C / 30Q / 96X
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           30x100GE/30x40GE/96x10GE
```

```
    Xcvr 0       REV 01   740-067443   XX60DPC           QSFP+-40G-SR4
    Xcvr 1       REV 01   740-054053   QF4605W7          QSFP+-4X10G-SR
    Xcvr 2       REV 01   740-067443   XX60DP8           QSFP+-40G-SR4
    Xcvr 3       REV 01   740-067442   XV30FYM           QSFP+-40G-SR4
    Xcvr 4       REV 01   740-067442   1ACP133406Z       QSFP+-40G-SR4
    Xcvr 5       REV 01   740-067443   XX60DP5           QSFP+-40G-SR4
    Xcvr 8       REV 01   740-046565   QH060355          QSFP+-40G-SR4
    Xcvr 12      REV 01   740-058734   1ECQ115008C       QSFP-100GBASE-SR4
    Xcvr 15      REV 01   740-046565   QH0602KG          QSFP+-40G-SR4
    Xcvr 16      REV 01   740-046565   QH0602LG          QSFP+-40G-SR4
    Xcvr 17      REV 01   740-046565   QH06035S          QSFP+-40G-SR4
    Xcvr 18      REV 01   740-046565   QH0602KS          QSFP+-40G-SR4
    Xcvr 24      REV 01   740-067443   QI2902DP          QSFP+-40G-SR4
Power Supply 2   REV 02   740-049388   1EDL70200NP       Power Supply AC
Power Supply 3   REV 02   740-049388   1EDL603005C       Power Supply AC
Power Supply 4   REV 02   740-049388   1EDL70200P1       Power Supply AC
Power Supply 5   REV 02   740-049388   1EDL70200B7       Power Supply AC
FTC 0            REV 16   750-050108   ACPK8682          Fan Controller 8
FTC 1            REV 16   750-050108   ACPR9530          Fan Controller 8
Fan Tray 0       REV 10   760-054372   ACPR9509          Fan Tray 8
Fan Tray 1       REV 10   760-054372   ACPV7260          Fan Tray 8
SIB 1            REV 28   750-050058   ACPV6306          Switch Fabric 8
SIB 2            REV 28                ACPR2569          Switch Fabric 8
SIB 3            REV 28   750-05       ACPW7402          Switch Fabric 8
SIB 4            REV 28   750-050058   ACPR2577          Switch Fabric 8
FPD Board        REV 07   711-054687   ACPM4965          Front Panel Display
```

show chassis hardware (PTX10016 Router Junos OS Evolved Release)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                DY814             JNP10016 [PTX10016]
Midplane 0       REV 02   750-085608   BCAW3258          Midplane 16
FPM 0            REV 01   711-086964   BCAR0775          Front Panel Display
PSM 0            Rev 03   740-069994   1F269200046       JNP10K 5500W AC/HVDC Power Supply Unit
PSM 1            Rev 03   740-069994   1F269170144       JNP10K 5500W AC/HVDC Power Supply Unit
PSM 2            REV 02   740-049388   1EDL60300C6      JNP10K 5500W AC/HVDC Power Supply Unit
PSM 3            REV 02   740-049388   1EDL603005X      JNP10K 5500W AC/HVDC Power Supply Unit
PSM 4            REV 02   740-049388   1EDL6170275      JNP10K 5500W AC/HVDC Power Supply Unit
PSM 5            REV 02   740-049388   1EDL61701WD      JNP10K 5500W AC/HVDC Power Supply Unit
```

```
Routing Engine 0          BUILTIN      BUILTIN           JNP10K-RE1-E
CB 0             REV 15   750-079562   BCAW3941          Control Board
Routing Engine 1          BUILTIN      BUILTIN           JNP10K-RE1-E
CB 1             REV 15   750-079562   BCAW3942          Control Board
FPC 0            REV 07   750-093524   BCAY8271          JNP10K-LC1201
CPU              REV 14   710-001726   HM1084            FPC CPU
FPC 10           REV 07   750-093524   BCAY8277          JNP10K-LC1201CPU              REV 05
710-010169   HZ3219            FPC CPU-Enhanced
SIB 0            REV 02   750-083426   BCAV7680          SIB-JNP10016
SIB 1            REV 02   750-083426   BCAV7682          SIB-JNP10016
SIB 2            REV 02   750-083426   BCAV7681          SIB-JNP10016
SIB 3            REV 02   750-083426   BCAV7684          SIB-JNP10016
SIB 4            REV 02   750-083426   BCAV7683          SIB-JNP10016
SIB 5            REV 02   750-083426   BCAV7685          SIB-JNP10016
FTC 0            REV 10   750-086270   BCAV0609          Fan Controller 16
Fan Tray 0       REV 02   750-103311   BCAY1793          Fan Tray 16
Fan Tray 1       REV 02   750-103311   BCAY1797          Fan Tray 16
```

show chassis hardware clei-models (PTX10016 Router)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 24   750-077138   CMMUN00ARA        JNP10016
CB 0             REV 04   711-065897   PROTOXCLEI        PROTO-ASSEMBLY
CB 1             REV 05   711-065897   PROTOXCLEI        PROTO-ASSEMBLY
FPC 2
  PIC 0                   BUILTIN
FPC 4            REV 35   750-071976   CMUIANABAA        JNP10K-LC1101
  PIC 0                   BUILTIN
FPC 5            REV 13   750-068822   CMUIAM9BAC        QFX10000-36Q
  PIC 0                   BUILTIN
FPC 6            REV 41   750-071976   CMUIANABAB        JNP10K-LC1101
  PIC 0                   BUILTIN
FPC 7            REV 35   750-071976   CMUIANABAA        JNP10K-LC1101
  PIC 0                   BUILTIN
FPC 8            REV 35   750-071976   CMUIANABAA        JNP10K-LC1101
  PIC 0                   BUILTIN
FPC 9            REV 41   750-071976   CMUIANABAB        JNP10K-LC1101
  PIC 0                   BUILTIN
FPC 10           REV 35   750-071976   CMUIANABAA        JNP10K-LC1101
```

```
  PIC 0                   BUILTIN
FPC 11           REV 35   750-071976   CMUIANABAA        JNP10K-LC1101
  PIC 0                   BUILTIN
FPC 13           REV 41   750-071976   CMUIANABAB        JNP10K-LC1101
  PIC 0                   BUILTIN
FPC 15           REV 37   750-071976   CMUIANABAA        JNP10K-LC1101
  PIC 0                   BUILTIN
Power Supply 0   REV 01   740-073147   CMUPADPBAA        JNP10K-PWR-DC
Power Supply 1   REV 01   740-073147   CMUPADPBAA        JNP10K-PWR-DC
Power Supply 2   REV 01   740-073147   CMUPADPBAA        JNP10K-PWR-DC
Power Supply 3   REV 01   740-073147   CMUPADPBAA        JNP10K-PWR-DC
Power Supply 4   REV 01   740-073147   CMUPADPBAA        JNP10K-PWR-DC
Power Supply 5   REV 01   740-073147   CMUPADPBAA        JNP10K-PWR-DC
Power Supply 6   REV 01   740-073147   CMUPADPBAA        JNP10K-PWR-DC
Power Supply 7   REV 01   740-073147   CMUPADPBAA        JNP10K-PWR-DC
Power Supply 8   REV 01   740-073147   CMUPADPBAA        JNP10K-PWR-DC
Power Supply 9   REV 01   740-073147   CMUPADPBAA        JNP10K-PWR-DC
Fan Tray 0                                              QFX5100-FAN-AFO
Fan Tray 1                                              QFX5100-FAN-AFO
SIB 0            REV 15   750-077140   CMUCAH6CAA        JNP10016-SF
SIB 1            REV 15   750-077140   CMUCAH6CAA        JNP10016-SF
SIB 2            REV 15   750-077140   CMUCAH6CAA        JNP10016-SF
SIB 3            REV 15   750-077140   CMUCAH6CAA        JNP10016-SF
SIB 4            REV 15   750-077140   CMUCAH6CAA        JNP10016-SF
SIB 5            REV 15   750-077140   CMUCAH6CAA        JNP10016-SF
FPD Board        REV 07   711-054687
```

show chassis hardware models (PTX10003 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
FPM 0            REV 01   711-078358   CAJK7987
PDU 0            REV 02   711-078357   BCAB0555          PROTO-ASSEMBLY
PSM 1            REV 01   740-073765   1GE27460076       JPSU-3000W-AC-AFO
PSM 2            REV 01   740-073765   1GE27500091
Routing Engine 0 REV 02   750-080695   BCAC2431          PROTO-ASSEMBLY
CB 0             REV 02   750-077001   BCAB9511
FPC 0            REV 11   750-077005   BCAJ2707          JNP10003-FPC
  PIC 0                   BUILTIN      BUILTIN           JNP10003-PIC
```

```
  PIC 1                   BUILTIN      BUILTIN           JNP10003-PIC
FPC 1                     BUILTIN      BUILTIN           JNP10003-FPC
  PIC 0                   BUILTIN      BUILTIN           JNP10003-PIC
  PIC 1                   BUILTIN      BUILTIN           JNP10003-PIC
SIB 0                     BUILTIN      BUILTIN           JNP10003-SIB
SIB 1                     BUILTIN      BUILTIN           JNP10003-SIB
Fan Tray 1       REV 01   711-078356   BCAC5263
Fan Tray 2       REV 01   711-078356   BCAC5277
Fan Tray 3       REV 01   711-078356   BCAC5224
```

show chassis hardware (QFX5700 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item               Version        Part number   Serial number   Description
Chassis                           EC819         JNP5700         [QFX5700]
…..
…..
…..
Routing Engine 0                  BUILTIN       BUILTIN         JNP5700-RCB
Routing Engine 1                  BUILTIN       BUILTIN         JNP5700-RCB
CB 0               REV 15         750-079562    BCAV5526        Control Board
CB 1               REV 15         750-079562    BCAW0403        Control Board
```

show chassis hardware (SRX5800 Router)

```
user@host> show chassis hardware
show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1080B50AFA      MX960
Midplane         REV 02   710-013698   CC6226            MX960 Backplane
Fan Extender     REV 02   710-018051   JY5235            Extended Cable Manager
FPM Board        REV 01   710-014974   JS4207            Front Panel Display
PEM 0            Rev 01   740-080280   1F238510008       MX960-UNIVERSAL-HV-PSM
PEM 1            Rev 01   740-080280   1F238510030       MX960-UNIVERSAL-HV-PSM
PEM 2            Rev 01   740-080280   1F238510023       MX960-UNIVERSAL-HV-PSM
PEM 3            Rev 01   740-080280   1F238510047       MX960-UNIVERSAL-HV-PSM
```

show chassis hardware interconnect-device (QFabric Systems)

```
user@switch> show chassis hardware interconnect-device interconnect1
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis          REV 07                                  QFX_olive
Midplane         REV 07   750-021261   BH0208188289      QFX Midplane
CB 0             REV 07   750-021261   BH0208188289      QFXIC08-CB4S
```

show chassis hardware models (MX2010 Router)

```
user@host > show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
FPM Board        REV 06   711-032349   ZX8744            711-032349
PSM 4            REV 0C   740-033727   VK00254           0000000000000000000000
PSM 5            REV 0B   740-033727   VG00015           0000000000000000000000
PSM 6            REV 0B   740-033727   VH00097           0000000000000000000000
PSM 7            REV 0C   740-033727   VJ00151           0000000000000000000000
PSM 8            REV 0C   740-033727   VJ00149           0000000000000000000000
PDM 0            REV 0B   740-038109   WA00008
PDM 1            REV 0B   740-038109   WA00014
Routing Engine 0 REV 02   740-041821   9009094134        RE-S-1800X4-16G-S
Routing Engine 1 REV 02   740-041821   9009094141        RE-S-1800X4-16G-S
CB 0             REV 08   750-040257   CAAB3491          750-040257
CB 1             REV 08   750-040257   CAAB3489          750-040257
SFB 0            REV 06   711-032385   ZV1828            711-032385
SFB 1            REV 07   711-032385   ZZ2568            711-032385
SFB 2            REV 07   711-032385   ZZ2563            711-032385
SFB 3            REV 07   711-032385   ZZ2564            711-032385
SFB 4            REV 07   711-032385   ZZ2580            711-032385
SFB 5            REV 07   711-032385   ZZ2579            711-0323856
SFB 6            REV 07   711-032385   CAAB4882          711-044170
SFB 7            REV 07   711-032385   CAAB4898          711-044170
FPC 0            REV 33   750-028467   CAAB1919          MPC-3D-16XGE-SFPP
FPC 1            REV 21   750-033205   ZG5027            MX-MPC3-3D
  MIC 0          REV 03   750-033307   ZV6299            MIC3-3D-10XGE-SFPP
  MIC 1          REV 03   750-033307   ZV6268            MIC3-3D-10XGE-SFPP
FPC 8            REV 22   750-031089   ZT9746            MX-MPC2-3D
  MIC 0          REV 26   750-028392   ABBS1150          MIC-3D-20GE-SFP
```

```
  MIC 1          REV 26   750-028387   ABBR9582          MIC-3D-4XGE-XFP
FPC 9            REV 11   750-036284   ZL3591            MPCE-3D-16XGE-SFPP
ADC 0            REV 05   750-043596   CAAC2073          750-043596
ADC 1            REV 01   750-043596   ZV4117            750-043596
ADC 8            REV 01   750-043596   ZV4107            750-043596
ADC 9            REV 02   750-043596   ZW1555            750-043596
Fan Tray 0       REV 2A   760-046960   ACAY0015
Fan Tray 1       REV 2A   760-046960   ACAY0019
Fan Tray 2       REV 2A   760-046960   ACAY0020
Fan Tray 3       REV 2A   760-046960   ACAY0021
```

show chassis hardware node-device (QFabric Systems)

```
user@switch> show chassis hardware node-device node1
Routing Engine 0   BUILTIN      BUILTIN           QFX Routing Engine
node1            REV 05   711-032234   ED3694            QFX3500-48S4Q-AFI

  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           48x 10G-SFP+
    Xcvr 8       REV 01   740-030658   AD0946A028B       SFP+-10G-USR
```

show chassis hardware (ACX Series)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                YW0220320039      JNP7100-48L
[ACX7100-48
  L]
PSM 1            REV 04   740-085431   1ED79520221       PSU 1600W AC, Front
to                                                                                    Back
Airflow
Routing Engine 0 REV 04   611-112446   YY0220320013      RE-JNP-7100
CB 0             REV 05   650-113149   YW0220320039      Control Board
FPC 0                     BUILTIN      BUILTIN           ACX7100-48L
  PIC 0                   BUILTIN      BUILTIN           MRATE- 48xSFP56 +
6xQSF                                                                                   P56-
DD
    Xcvr 0       REV 01   740-068639   1A1M31A311008     SFP28-25G-BASE-SR
    Xcvr 1       REV 01   740-030658   AA1230AZYWW       SFP+-10G-USR
```

```
    Xcvr 2       REV 01   740-031980   AP42G0C           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   ARN2FS9           SFP+-10G-SR
    Xcvr 4       REV 01   740-031980   AP4150P           SFP+-10G-SR
    Xcvr 5       REV 01   740-031980   123363A01134      SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   B11E02539         SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   ARQ0WRX           SFP+-10G-SR
    Xcvr 12      REV 01   740-031980   193363A00707      SFP+-10G-SR
    Xcvr 15      REV 01   740-031980   AMS15RT           SFP+-10G-SR
    Xcvr 18      REV 01   740-068639   1A1M31A5370MX     SFP28-25G-BASE-SR
    Xcvr 19      REV 01   740-068639   1A1M31A5370MS     SFP28-25G-BASE-SR
    Xcvr 21      REV 01   740-068639   1A1M31A5370MW     SFP28-25G-BASE-SR
    Xcvr 22      REV 01   740-068639   1A1M31A5370MT     SFP28-25G-BASE-SR
    Xcvr 43      REV 01   740-031980   CH09KN1H4         SFP+-10G-SR
    Xcvr 44      REV 01   740-031980   AMB0TC1           SFP+-10G-SR
    Xcvr 46      REV 01   740-031980   AP40WQN           SFP+-10G-SR
    Xcvr 47      REV 01   740-031980   APR1BCS           SFP+-10G-SR
    Xcvr 48      REV 01   740-054053   XXH0KH1           QSFP+-4X10G-SR
    Xcvr 49      REV 01   740-058734   1ECQ144605L       QSFP-100GBASE-SR4
    Xcvr 50      REV 01   740-054053   XWP0Q7D           QSFP+-4X10G-SR
    Xcvr 51      REV 01   740-058734   1ECQ144604H       QSFP-100GBASE-SR4
    Xcvr 52      REV 01   740-054053   XXH0KGC           QSFP+-4X10G-SR
    Xcvr 53      REV 01   740-058734   1ECQ144605K       QSFP-100GBASE-SR4
Fan Tray 0                                               JNP7100 Fan Tray,
Front                                                                                                    to
Back Airflow - AFO
Fan Tray 1                                               JNP7100 Fan Tray,
Front                                                                                                    to
Back Airflow - AFO
Fan Tray 2                                               JNP7100 Fan Tray,
Front                                                                                                    to
Back Airflow - AFO
Fan Tray 3                                               JNP7100 Fan Tray,
Front                                                                                                    to
Back Airflow - AFO
Fan Tray 4                                               JNP7100 Fan Tray,
Front                                                                                                    to
Back Airflow - AFO
Fan Tray 5                                               JNP7100 Fan Tray,
Front                                                                                                    to
Back Airflow - AFO
```

show chassis hardware (MX204)

The `show chassis hardware` command output does not include information about the fan tray serial number and part number of MX204 routers. This does not affect the fan's functionality.

```
user@host> show chassis hardware
Hardware inventory:Item             Version  Part number  Serial number     Description
Chassis                                         CV190             JNP204 [MX204]
Routing Engine 0          BUILTIN      BUILTIN           RE-S-1600x8
CB 0             REV 40   750-069579   CARN0542          JNP204 [MX204]
FPC 0                     BUILTIN      BUILTIN           MPC
  PIC 0                   BUILTIN      BUILTIN           4XQSFP28 PIC
    Xcvr 0       REV 01   740-067443   1ACP16161C7       QSFP+-40G-SR4
    Xcvr 1       REV 01   740-067443   1ACP16161E7       QSFP+-40G-SR4
    Xcvr 2       REV 01   740-067443   1ACP16161DV       QSFP+-40G-SR4
    Xcvr 3       REV 01   740-099582   1A1CQ1B6350DC     QSFP-100GBASE-SR4
  PIC 1                   BUILTIN      BUILTIN           8XSFPP PIC
    Xcvr 0       REV 01   740-031980   AJJ036G           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AN20CP8           SFP+-10G-SR
PEM 1            REV 04   740-043886   1GA4A380672       JPSU-650W-DC-AFO
Fan Tray 0                                               Fan Tray, Front to Back Airflow - AFO
Fan Tray 1                                               Fan Tray, Front to Back Airflow - AFO
Fan Tray 2                                               Fan Tray, Front to Back Airflow -
AFO                                                                                          to
Back Airflow - AFO
```

show chassis hardware (SRX5800)

```
user@host> show chassis hardware

Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1080B50AFA      MX960
Midplane         REV 02   710-013698   CC6226            MX960 Backplane
Fan Extender     REV 02   710-018051   JY5235            Extended Cable Manager
FPM Board        REV 01   710-014974   JS4207            Front Panel Display
PEM 0            Rev 01   740-080280   1F238510008       SRX5800-UNIVERSAL-HV-PSM
PEM 1            Rev 01   740-080280   1F238510030       SRX5800-UNIVERSAL-HV-PSM
PEM 2            Rev 01   740-080280   1F238510023       SRX5800-UNIVERSAL-HV-PSM
PEM 3            Rev 01   740-080280   1F238510047       SRX5800-UNIVERSAL-HV-PSM
```

show chassis hardware (PTX10001-36MR) (400G ZR and 400G ZR-M optics)

```
user@host show chassis hardware
Hardware inventory:
    Item             Version  Part number  Serial number     Description
    Chassis                                EM436             JNP10001-36MR [PTX10001-36MR]
    PSM 0            REV 07   740-073765   1GE29391099       AC AFO 3000W PSU
    Routing Engine 0 REV 08   750-100243   BCBH4788          RE-JNP10001-36MR
    CB 0             REV 16   750-099260   BCBJ1424          Control Board
    FPC 0                     BUILTIN      BUILTIN           FPC-JNP10001-36MR
      PIC 0                   BUILTIN      BUILTIN           8X400GE-MR + 4X100GE-MR
        Xcvr 0       REV 01   740-032986   ECX40K80042       QSFP+-40G-SR4
        Xcvr 2       REV 01   740-032986   QF4804ZM          QSFP+-40G-SR4
        Xcvr 3       REV 01   740-046565   QD190447          QSFP+-40G-SR4
        Xcvr 4       REV 01   740-038624   S1810376855-1     QSFP+-40G-CU3M
        Xcvr 5       REV 01   740-032986   ECX40K80028       QSFP+-40G-SR4
        Xcvr 10      REV 01   740-131169   1T1TZFA70500D     QSFP56-DD-400G-ZR-M
      PIC 1                   BUILTIN      BUILTIN           8X400GE-MR + 4X100GE-MR
        Xcvr 0       REV 01   740-032986   XV303SF           QSFP+-40G-SR4
        Xcvr 1       REV 01   740-067442   QI1202B5          QSFP+-40G-SR4
        Xcvr 2       XXXX     NON-JNPR     204053396         QSFP56-DD-400G-ZR
        Xcvr 4       REV 01   740-061405   1ECQ13150KE       QSFP-100GBASE-SR4-T2
        Xcvr 10      REV 01   740-131169   1T1TZFA70500C     QSFP56-DD-400G-ZR-M
      PIC 2                   BUILTIN      BUILTIN           8X400GE-MR + 4X100GE-MR
        Xcvr 0       REV 01   740-032986   QD481959          QSFP+-40G-SR4
        Xcvr 1       REV 01   740-038624   C1909242760-1     QSFP+-40G-CU3M
        Xcvr 3       REV 01   740-032986   ECX02IC0146       QSFP+-40G-SR4
        Xcvr 4       REV 01   740-065631   1ACS12510R3       QSFP28-100G-AOC-3M
        Xcvr 5       REV 01   740-065631   1ACS1251023       QSFP28-100G-AOC-3M
        Xcvr 9       REV 01   740-061405   1ECQ13150A5       QSFP-100GBASE-SR4-T2
        Xcvr 10      XXXX     NON-JNPR     204053387         QSFP56-DD-400G-ZR
        Xcvr 11      REV 01   740-038623   MOC14256230865    QSFP+-40G-CU1M
    SIB 0                     BUILTIN      BUILTIN           SIB-JNP10001-36MR
    Fan Tray 0                                               JNP10001 Fan Tray, Front to Back
Airflow - AFO
    Fan Tray 1                                               JNP10001 Fan Tray, Front to Back
Airflow - AFO
    Fan Tray 2                                               JNP10001 Fan Tray, Front to Back
Airflow - AFO
    Fan Tray 3                                               JNP10001 Fan Tray, Front to Back
Airflow - AFO
    Fan Tray 4                                               JNP10001 Fan Tray, Front to Back
```

```
Airflow - AFO
    Fan Tray 5                                          JNP10001 Fan Tray, Front to Back
Airflow - AFO
```

show chassis hardware (MX304)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                FG907             JNP304 [MX304]
Routing Engine 0 REV 11   750-123749   BCDA4102          RE 2700 8C 128G
Routing Engine 1 REV 11   750-123749   BCDA4049          RE 2700 8C 128G
CB 0             REV 30   750-123404   BCDA2318          Control Board
FPC 0                     BUILTIN      BUILTIN           FPC-BUILTIN
  CPU            REV 10   750-122877   BCCX0173          JNP304 PMB
  MIC 0          REV 21   750-122718   BCDB5606          MRATE LMIC 16x100G/4x400G
    PIC 0                 BUILTIN      BUILTIN           MRATE LMIC 16x100G/4x400G
      Xcvr 1     REV 01   740-067442   1ACP15500HW       QSFP+-40G-SR4
      Xcvr 3     REV 01   740-054053   QL35006T          QSFP+-4X10G-SR
      Xcvr 4     REV 01   740-061405   1F1CQ1A52008G     QSFP-100GBASE-SR4
      Xcvr 6     REV 01   740-061405   1ACQ151800V       QSFP-100GBASE-SR4
  MIC 1          REV 28   750-122718   BCER0799          MRATE LMIC 16x100G/4x400G
    PIC 1                 BUILTIN      BUILTIN           MRATE LMIC 16x100G/4x400G
      Xcvr 1     REV 01   740-061405   1ECQ144505J       QSFP-100GBASE-SR4
      Xcvr 15    REV 01   740-061405   1ACQ15200LD       QSFP-100GBASE-SR4
PEM 0            Rev 02   740-110419   1F27C030036       AC AFO 2200W Power Supply
PEM 1            Rev 02   740-110419   1F27C030151       AC AFO 2200W Power Supply
Fan Tray 0       REV 04   760-126744   BCCW1082          JNP304 Fan Tray, Front to Back Airflow
Fan Tray 1       REV 04   760-126744   BCCW1091          JNP304 Fan Tray, Front to Back Airflow
Fan Tray 2       REV 04   760-126744   BCCW1081          JNP304 Fan Tray, Front to Back Airflow
SFB 0            REV 07   750-122847   BCCW0194          Switch Fabric Board
TIB              REV 06   750-126514   BCCW7956          Timing Interface Board
```

show chassis hardware (ACX7509-FPC-20Y)

```
user@host show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
FPC 5            REV 04   750-120787   CARH7889          JNP-FPC-20Y
  PIC 0                   BUILTIN      BUILTIN           20x1/10/25/50-SFP56
```

```
    Xcvr 0       REV 01   740-071562   1A1M3AA5480F8     SFP28-25G-BASE-LR
    Xcvr 1       REV 01   720-120775   1P1C54A628HVN     SFP56-50G-DAC-3M
    Xcvr 2       REV 01   720-120775   1P1C54A629JCQ     SFP56-50G-DAC-3M
    Xcvr 5       REV 01   740-031980   A55B20R           SFP+-10G-SR
    Xcvr 6       REV 01   720-120775   1P1C54A629JCS     SFP56-50G-DAC-3M
    Xcvr 7       REV 01   720-120775   1P1C54A628HWR     SFP56-50G-DAC-3M
    Xcvr 8       REV 01   740-068639   1A1M31A31105D     SFP28-25G-BASE-SR
    Xcvr 9       REV 01   740-068639   1A1M31A31105A     SFP28-25G-BASE-SR
    Xcvr 10      REV 01   720-120775   1P1C54A628HZR     SFP56-50G-DAC-3M
    Xcvr 11      REV 01   720-120775   1P1C54A628J00     SFP56-50G-DAC-3M
    Xcvr 12      REV 01   720-120775   1P1C54A628HW4     SFP56-50G-DAC-3M
    Xcvr 16      REV 01   740-031980   MXA0NHJ           SFP+-10G-SR
    Xcvr 17      REV 01   740-031980   AP41HCV           SFP+-10G-SR
    Xcvr 18      REV 01   740-031980   A54BA6S           SFP+-10G-SR
    Xcvr 19      REV 01   740-031980   AP40WRE           SFP+-10G-SR
```

show chassis hardware (QFX5230-64CD devices with logical FPC and PIC) (Junos OS Evolved Release)

```
user@root> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                AH36022884        QFX5230-64CD
PSM 0            REV 08   740-073765   1GE2A302501       AC AFO 3000W PSU
PSM 1            REV 08   740-073765   1GE2A392826       AC AFO 3000W PSU
Routing Engine 0          BUILTIN      BUILTIN           QFX5230-64CD
CB 0             REV 01   650-086244   AH36022884        QFX5230-64CD
FPC 0                     BUILTIN      BUILTIN           QFX5230-64CD
  PIC 0                   BUILTIN      BUILTIN           64X400G-QSFP-DD
    Xcvr 0       REV 01   740-102183   49876             QSFP56-DD LPBK
    Xcvr 1       REV 01   740-102183   49942             QSFP56-DD LPBK
    Xcvr 2       REV 01   740-102183   49895             QSFP56-DD LPBK
    Xcvr 3       REV 01   740-102183   49915             QSFP56-DD LPBK
    Xcvr 9       REV 01   740-090166   2H2TSBA535003     QSFP56-DD-400G-AOC-5M
    Xcvr 15      REV 01   740-102183   49864            QSFP56-DD LPBK
    Xcvr 18      REV 01   740-102183   49869             QSFP56-DD LPBK
...
...
...
    Xcvr 62               740-102183   49860             QSFP56-DD LPBK
```

```
    Xcvr 64      REV 01   740-021308   MSP0BWW           SFP+-10G-SR
    Xcvr 65      0        NON-JNPR                       UNKNOWN
...
...
...
Fan Tray 5                                               Unsupported
```

Release Information

Command introduced before Junos OS Release 7.4.

`models` option introduced in Junos OS Release 8.2.

`sfc` option introduced in Junos OS Release 9.6 for the TX Matrix Plus router.

Information for `disk` and `usb` introduced in Junos OS Release 15.1X53-D60 for QFX10002, QFX10008, and QFX10016 switches.

`pem` detail output added in Junos OS Release 21.4R1 for SRX5800.

NOTE: Devices and routing platforms use the basic syntax, unless otherwise listed. For example, the EX Series has an additional satellite parameter available.

RELATED DOCUMENTATION

show chassis power

show chassis hardware (GNF)

IN THIS SECTION

- Syntax | 18454
- Description | 18454
- Required Privilege Level | 18454
- Output Fields | 18454
- Sample Output | 18455

Release Information | 18456

Syntax

```
show chassis hardware
```

Description

Display a list of all hardware components of the chassis, including the hardware version level and serial number, the GNF Routing Engine details, and the FPCs assigned to the GNF.

Required Privilege Level

view

Output Fields

Table 290 on page 18454 lists the output fields for the `show chassis hardware` command. Output fields are listed in the approximate order in which they appear.

Table 290: show chassis hardware Output Fields

Field Name	Field Description
`Item`	Chassis component: • Information about the backplane, Routing Engine, Power Entry Modules (PEMs), and fan trays. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs), Modular Port Concentrators (MPCs) and associated Modular Interface Cards (MICs), or Dense Port Concentrators (DPCs).
`Version`	Revision level of the chassis component.
`Part number`	Part number of the chassis component.

Table 290: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
`Serial number`	Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis.
`Description`	Brief description of the hardware item.

Sample Output

command-name

```
user@router> show chassis hardware

bsys-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11C9CDDAFK      MX2010
Midplane         REV 35   750-044636   ABAB9184          Lower Backplane
Midplane 1       REV 02   711-044557   ABAB9048          Upper Backplane
PMP              REV 04   711-032426   ACAJ2622          Power Midplane
FPM Board        REV 09   760-044634   ABCF2618          Front Panel Display
PSM 0            REV 01   740-050037   1EDB3130084       DC 52V Power Supply Module
PSM 1            REV 01   740-050037   1EDB313001Z       DC 52V Power Supply Module
PSM 2            REV 01   740-050037   1EDB321018D       DC 52V Power Supply Module
PSM 3            REV 01   740-050037   1EDB32101AZ       DC 52V Power Supply Module
PSM 4            REV 01   740-050037   1EDB32202C2       DC 52V Power Supply Module
PSM 5            REV 01   740-050037   1EDB32100TC       DC 52V Power Supply Module
PSM 6            REV 01   740-050037   1EDB3210166       DC 52V Power Supply Module
PSM 7            REV 01   740-050037   1EDB3210165       DC 52V Power Supply Module
PSM 8            REV 01   740-050037   1EDB3210163       DC 52V Power Supply Module
PDM 0            REV 03   740-045234   1EGA3170177       DC Power Dist Module
Routing Engine 0 REV 08   750-055814   CAFV5537          RE-S-2X00x8
CB 0             REV 08   750-055087   CAFN3426          MX2K Enhanced SCB
  Xcvr 0         REV 01   740-031980   ALM0HC7           SFP+-10G-SR
  Xcvr 1         REV 01   740-031980   123363A00418      SFP+-10G-SR
```

```
CB 1            REV 08   750-055087   CAFN3423          MX2K Enhanced SCB
SPMB 0          REV 05   711-041855   CAEZ5998          PMB Board
SPMB 1          REV 05   711-041855   CAEZ5993          PMB Board
SFB 0           REV 06   711-044466   ABCD6742          Switch Fabric Board
SFB 1           REV 06   711-044466   ABCG5627          Switch Fabric Board
SFB 2           REV 06   711-044466   ABCG5659          Switch Fabric Board
SFB 3           REV 06   711-044466   ABCG5653          Switch Fabric Board
SFB 4           REV 06   711-044466   ABCG5611          Switch Fabric Board
SFB 5           REV 06   711-044466   ABCG5635          Switch Fabric Board
SFB 6           REV 06   711-044466   ABCG5638          Switch Fabric Board
SFB 7           REV 06   711-044466   ABCG3650          Switch Fabric Board
FPC 8           REV 68   750-044130   ABCY5967          MPC6E 3D
  CPU           REV 12   711-045719   ABCY9696          RMPC PMB
Fan Tray 0      REV 06   760-046960   ACAY0428          172mm FanTray - 6 Fans
Fan Tray 1      REV 06   760-046960   ACAY0800          172mm FanTray - 6 Fans
Fan Tray 2      REV 06   760-046960   ACAY0797          172mm FanTray - 6 Fans
Fan Tray 3      REV 06   760-046960   ACAY1047          172mm FanTray - 6 Fans

gnf2-re0:
--------------------------------------------------------------------------
Chassis                                GN59081553B0      MX2010-GNF     <<<
Routing Engine 0                                         RE-GNF-1700x4
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show chassis fpc (GNF) | **18404**

show chassis fpc pic-status (GNF) | **18411**

show chassis adc (GNF) | **17295**

show chassis hardware (View)

IN THIS SECTION

- Syntax | 18457
- Description | 18457
- Options | 18458
- Required Privilege Level | 18458
- Output Fields | 18459
- show chassis hardware | 18464
- show chassis hardware (SRX4200) | 18474
- show chassis hardware (vSRX Virtual Firewall 3.0) | 18475
- show chassis hardware clei-models | 18475
- Release Information | 18476

Syntax

```
show chassis hardware
<clei-models>
<detail | extensive>
<models>
<node(node-id | all | local | primary)>
```

Description

Display chassis hardware information.

Starting in Junos OS Release 20.1R1, when vSRX Virtual Firewall 3.0 performs resource management, the vCPUs and RAM available to the instance are assigned based on what has been allocated prior to launching the instance. A maximum of 32 cores will be assigned to SRXPFE, for flow processing. Any allocation of cores in excess of 32 will automatically be assigned to the Routing Engine. For example, if 36 cores are allocated to the VM during the creation process, 32 cores are assigned for flow processing and 4 cores will be assigned to the RE. For memory allocations, up to 64G of vRAM would be used by the SRXPFE. Any allocated memory in excess of 64G would be assigned to system memory and would not be used for maintaining flow sessions information.

Table 291: Recommended vCPU and vRAM Combinations

vCPU Number	vRAM Size (G)
2	4
5	8
9	16
17	32

On a deployed vSRX Virtual Firewall, only memory scale up is supported. Scaling down memory on a deployed vSRX Virtual Firewall, is not supported. If you need to scale down memory, then a fresh install is required.

Options

- `clei-models`—(Optional) Display Common Language Equipment Identifier Code (CLEI) barcode and model number for orderable field-replaceable units (FRUs).
- `detail` | `extensive`—(Optional) Display the specified level of output.
- `models`—(Optional) Display model numbers and part numbers for orderable FRUs.
- `node`—(Optional) For chassis cluster configurations, display chassis hardware information on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 292 on page 18459 lists the output fields for the `show chassis hardware` command. Output fields are listed in the approximate order in which they appear.

Table 292: show chassis hardware Output Fields

Field Name	Field Description
`Item`	Chassis component—Information about the backplane; power supplies; fan trays; Routing Engine; each Physical Interface Module (PIM)—reported as FPC and PIC—and each fan, blower, and impeller.
`Version`	Revision level of the chassis component.
`Part Number`	Part number for the chassis component.
`Serial Number`	Serial number of the chassis component. The serial number of the backplane is also the serial number of the device chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the device chassis.
`Assb ID or Assembly ID`	Identification number that describes the FRU hardware.
`FRU model number`	Model number of FRU hardware component.
`CLEI code`	Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.
`EEPROM Version`	ID EEPROM version used by hardware component: 0x01 (version 1) or 0x02 (version 2).

Table 292: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
`Description`	Brief description of the hardware item: • Type of power supply. • Switch Control Board (SCB) Starting with Junos OS Release 12.1X47-D15 and Junos OS Release 17.3R1, the SRX5K-SCBE (SCB2) is introduced. • There are three SCB slots in SRX5800 devices. The third slot can be used for an SCB or an FPC. When an SRX5K-SCB was used , the third SCB slot was used as an FPC. SCB redundancy is provided in chassis cluster mode. • With an SCB2, a third SCB is supported. If a third SCB is plugged in, it provides intra-chassis fabric redundancy. • The Ethernet switch in the SCB2 provides the Ethernet connectivity among all the FPCs and the Routing Engine. The Routing Engine uses this connectivity to distribute forwarding and routing tables to the FPCs. The FPCs use this connectivity to send exception packets to the Routing Engine. • Fabric connects all FPCs in the data plane. The Fabric Manager executes on the Routing Engine and controls the fabric system in the chassis. Packet Forwarding Engines on the FPC and fabric planes on the SCB are connected through HSL2 channels. • SCB2 supports HSL2 with both 3.11 Gbps and 6.22 Gbps (SerDes) link speed and various HSL2 modes. When an FPC is brought online, the link speed and HSL2 mode are determined by the type of FPC. Starting with Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the SRX5K-SCB3 (SCB3) with enhanced midplane is introduced. • All existing SCB software that is supported by SCB2 is supported on SCB3. • SRX5K-RE-1800X4 mixed Routing Engine use is not supported.

Table 292: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
	• SCB3 works with the SRX5K-MPC (IOC2), SRX5K-MPC3-100G10G (IOC3), SRX5K-MPC3-40G10G (IOC3), and SRX5K-SPC-4-15-320 (SPC2) with current midplanes and the new enhanced midplanes. • Mixed SCB use is not supported. If an SCB2 and an SCB3 are used, the system will only power on the primary Routing Engine's SCB and will power off the other SCBs. Only the SCB in slot 0 is powered on and a system log is generated. • SCB3 supports up to 400 Gbps per slot with old midplanes and up to 500 Gbps per slot with new midplanes. • SCB3 supports fabric intra-chassis redundancy. • SCB3 supports the same chassis cluster function as the SRX5K-SCB (SCB1) and the SRX5K-SCBE (SCB2), except for in-service software upgrade (ISSU) and in-service hardware upgrade (ISHU). • SCB3 has a second external Ethernet port. • Fabric bandwidth increasing mode is not supported. Starting in Junos OS 19.3R1, SRX5K-SCB4 is supported on SRX5600 and SRX5800 devices along with SRX5K-SPC3. SRX5K-SCB4: • Interoperate with SRX5K-RE3-128G, SRX5K-RE-1800X4, IOC2, IOC3, IOC4, SPC2, and SPC3. SCB4 is compatible with all midplanes and interoperate with existing PEMs, fan trays, and front panel displays. • Does not interoperate with SCB, SCB2, and SCB3. • Supports 480-Gbps link speed per slot. • Supports 1-Gigabit Ethernet interfaces speed with SRX5K-RE-1800X4 and 1-Gigabit, 2.5-Gigabit, and 10-Gigabit Ethernet speeds with SRX5K-RE3-128G. • Support ISHU and ISSU in chassis cluster.

Table 292: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
	• Supports fabric bandwidth mode and redundant fabric mode on SRX5600 and SRX5800 devices. The bandwidth mode is the new default mode which is necessary to configure redundant mode in setting up the chassis cluster successfully. • Type of Flexible PIC Concentrator (FPC), Physical Interface Card (PIC), Modular Interface Cards (MICs), and PIMs. • IOCs Starting with Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) are introduced. • IOC3 has two types of IOC3 MPCs, which have different built-in MICs: the 24x10GE + 6x40GE MPC and the 2x100GE + 4x10GE MPC. • IOC3 supports SCB3 and SRX5000 line backplane and enhanced backplane. • IOC3 can only work with SRX5000 line SCB2 and SCB3. If an SRX5000 line SCB is detected, IOC3 is offline, an FPC misconfiguration alarm is raised, and a system log message is generated. • IOC3 interoperates with SCB2 and SCB3. • IOC3 interoperates with the SRX5K-SPC-4-15-320 (SPC2) and the SRX5K-MPC (IOC2). • The maximum power consumption for one IOC3 is 645W. An enhanced power module must be used. • The IOC3 does not support the following command to set a PIC to go offline or online: `request chassis pic fpc-slot <fpc-slot> pic-slot <pic-slot> <offline | online>` . • IOC3 supports 240 Gbps of throughput with the enhanced SRX5000 line backplane. • Chassis cluster functions the same as for the SRX5000 line IOC2.

Table 292: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
	• IOC3 supports intra-chassis and inter-chassis fabric redundancy mode. • IOC3 supports ISSU and ISHU in chassis cluster mode. • IOC3 supports intra-FPC and and Inter-FPC Express Path (previously known as *services offloading*) with IPv4. • NAT of IPv4 and IPv6 in normal mode and IPv4 for Express Path mode. • All four PICs on the 24x10GE + 6x40GE cannot be powered on. A maximum of two PICs can be powered on at the same time. Use the `set chassis fpc <slot> pic <pic> power off` command to choose the PICs you want to power on. Fabric bandwidth increasing mode is not supported on IOC3. • SRX Clustering Module (SCM) • Fan tray Starting in Junos OS Release 19.3R1, the SRX5K-IOC4-10G and SRX5K-IOC4-MRAT line cards are supported along with SRX5K-SPC3 on the SRX5000 line devices. SRX5K-IOC4-10G: • Interoperates with SCB3, SCB4, SRX5K-RE-1800X4, SRX5K-RE3-128G, SPC2, SPC3, IOC2,IOC3, and IOC4. • Supports 480-Gbps speed. • Supports 40X10GE Interfaces with SCB3. • 40 10-Gigabit Ethernet port provides 10-Gigabit Ethernet MACsec support. • Supports reth and aggregated interfaces on the chassis cluster. • Supports ISSU and logical system on the chassis cluster. • Does not support SCB2.

Table 292: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
	• SRX5K-IOC4-MRAT with SCB3 supports 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet Interfaces. • For hosts, the Routing Engine type. Starting with Junos OS Release 12.1X47-D15 and Junos OS Release 17.3R1, the SRX5K-RE-1800X4 Routing Engine is introduced. • The SRX5K-RE-1800X4 has an Intel Quad core Xeon processor, 16 GB of DRAM, and a 128-GB solid-state drive (SSD). The number 1800 refers to the speed of the processor (1.8 GHz). The maximum required power for this Routing Engine is 90W. **NOTE**: The SRX5K-RE-1800X4 provides significantly better performance than the previously used Routing Engine, even with a single core. Starting in Junos OS Release 19.3R1, SRX5K-RE3-128G Routing Engine is supported along with SRX5K-SPC3 on the SRX5000 line devices. SRX5K-RE3-128G: • Provides improved control plane performance and scalability. SRX5K-RE3-128G has Intel's Haswell-EP based processor with six cores. • Supports two 200G SSDs to store log files and 128-GB of memory for storing routing and forwarding tables and for other routing engines. • Interoperates with SCB3, SCB4, SRX5K-RE3-128G, SPC2, SPC3, IOC2, IOC3, and IOC4. • Does not support SCB2 and SRX5K-RE-1800X4.

show chassis hardware

show chassis hardware (SRX5800)

```
user@host> show chassis hardware
node0:
--------------------------------------------------------------------------
```

```
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1267B0FAGA      SRX5800
Midplane         REV 42   760-063937   ACRL3065          Enhanced SRX5800
Backpl
                   ane
FPM Board        REV 05   760-061272   CAHE4860          Front Panel Display
PDM              Rev 01   740-063049   QCS2209509D       Power Distribution
Modu
                 le
PEM 0            Rev 04   740-034724   QCS171002016      PS 4.1kW; 200-240V AC
i
              n
PEM 1            Rev 11   740-027760   QCS1825N07S       PS 4.1kW; 200-240V AC
i
              n
Routing Engine 0 REV 01   750-095568   CALK8884          SRX5k RE-2000x6
Routing Engine 1 REV 01   750-095568   CADZ9076          SRX5k RE-2000x6
CB 0             REV 26   750-031391   CALV3002          SRX5k SCB4
CB 1             REV 26   750-031391   CALV3009          SRX5k SCB4
FPC 2            REV 28   750-073435   CALS4630          SPC3
  CPU                     BUILTIN      BUILTIN           SRX5k vCPP Broadwell
  PIC 0                   BUILTIN      BUILTIN           SPU Cp-Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
FPC 3            REV 17   750-044175   CABE7777          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 4            REV 08   750-061262   CAFD8147          SRX5k IOC II
  CPU            REV 02   711-061263   CAFV7488          SRX5k MPC PMB
  MIC 0          REV 03   750-055732   CAFV9369          20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 02   740-011613   PNB1GJR           SFP-SX
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
FPC 5            REV 10   750-062242   CAKX2328          SRX5k IOC3 2CGE+4XGE
  PIC 0                   BUILTIN      BUILTIN           2x 10GE SFP+
    Xcvr 0       REV 01   740-021308   ANA07RE           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AQF0RBJ           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1x 100GE CFP2
  PIC 2                   BUILTIN      BUILTIN           2x 10GE SFP+
    Xcvr 0       REV 01   740-031980   AA1650304RF       SFP+-10G-SR
```

```
    Xcvr 1       REV 01   740-021308   AQ93BDK           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           1x 100GE CFP2
FPC 8            REV 46   750-056519   CALC4514          SRX5k IOC4 MRAT
  CPU            REV 21   750-057177   CALC3494          SMPC PMB
  PIC 0                   BUILTIN      BUILTIN           MRATE-6xQSFPP-XGE-XLGE-CGE
    Xcvr 0       REV 01   740-059437   000T20128         QSFP28-LPBK
    Xcvr 1       REV 01   740-067443   1ACP13450KH       QSFP+-40G-SR4
  PIC 1                   BUILTIN      BUILTIN           MRATE-6xQSFPP-XGE-XLGE-CGE
    Xcvr 0       REV 01   740-059437   0000T3443         QSFP28-LPBK
Fan Tray 0       REV 06   740-035409   ACAE9390          Enhanced Fan Tray
Fan Tray 1       REV 06   740-035409   ACAE9386          Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1267B01AGA      SRX5800
Midplane         REV 42   760-063937   ACRL3068          Enhanced SRX5800 Backplane
FPM Board        REV 05   760-061272   CAJX9988          Front Panel Display
PDM              Rev 01   740-063049   QCS2209507A       Power Distribution Module
PEM 0            Rev 11   740-027760   QCS1822N0EY       PS 4.1kW; 200-240V AC in
PEM 1            Rev 03   740-034724   QCS17020203F      PS 4.1kW; 200-240V AC in
Routing Engine 0 REV 01   750-095568   CALK8904          SRX5k RE-2000x6
Routing Engine 1 REV 01   750-095568   CADZ9076          SRX5k RE-2000x6
CB 0             REV 26   750-031391   CALV3010          SRX5k SCB4
CB 1             REV 26   750-031391   CALV3000          SRX5k SCB4
FPC 2            REV 28   750-073435   CAKZ9620          SPC3
  CPU                     BUILTIN      BUILTIN           SRX5k vCPP Broadwell
  PIC 0                   BUILTIN      BUILTIN           SPU Cp-Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
FPC 3            REV 18   750-054877   CACH4082          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 4            REV 08   750-061262   CAFD8165          SRX5k IOC II
  CPU            REV 02   711-061263   CAFV7507          SRX5k MPC PMB
  MIC 0          REV 03   750-055732   CAFV6603          20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-011613   AM0805S8M4N       SFP-SX
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
FPC 5            REV 03   750-062242   CAFZ2748          SRX5k IOC3 2CGE+4XGE
```

```
  PIC 0                   BUILTIN      BUILTIN           2x 10GE SFP+
    Xcvr 0       REV 01   740-021308   11T511100788      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AS92WJ0           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1x 100GE CFP2
  PIC 2                   BUILTIN      BUILTIN           2x 10GE SFP+
    Xcvr 0       REV 01   740-031980   AA1650304EZ       SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   ANS0EAR           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           1x 100GE CFP2
FPC 8            REV 46   750-056519   CALC4526          SRX5k IOC4 MRAT
  CPU            REV 21   750-057177   CALF5727          SMPC PMB
  PIC 0                   BUILTIN      BUILTIN           MRATE-6xQSFPP-XGE-XLGE-CGE
    Xcvr 1       REV 01   740-067443   1ACP13450L9       QSFP+-40G-SR4
  PIC 1                   BUILTIN      BUILTIN           MRATE-6xQSFPP-XGE-XLGE-CGE
Fan Tray 0       REV 06   740-035409   ACAE9298          Enhanced Fan Tray
Fan Tray 1       REV 06   740-035409   ACAE9314          Enhanced Fan Tray

{primary:node0}
```

show chassis hardware (SRX5600 and SRX5800 devices for SRX5K-MPC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN12170EAAGA      SRX 5800
Midplane         REV 01   710-041799   ACAX3849          SRX 5800 Backplane
FPM Board        REV 01   710-024632   CAAX7297          Front Panel Display
PDM              Rev 03   740-013110   QCS170250DU       Power Distribution Modu
le
PEM 0            Rev 03   740-034724   QCS17020203F      PS 4.1kW; 200-240V AC i
n
PEM 1            Rev 03   740-034724   QCS17020203C      PS 4.1kW; 200-240V AC i
n
PEM 2            Rev 04   740-034724   QCS17100200A      PS 4.1kW; 200-240V AC i
n
PEM 3            Rev 03   740-034724   QCS17080200M      PS 4.1kW; 200-240V AC i
n
Routing Engine 0 REV 11   740-023530   9012047437        SRX5k RE-13-20
CB 0             REV 09   710-024802   CAAX7202          SRX5k SCB
CB 1             REV 09   710-024802   CAAX7157          SRX5k SCB
FPC 0            REV 07   750-044175   CAAD0791          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
```

```
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 07   750-044175   CAAD0751          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 28   750-020751   CAAW1817          SRX5k DPC 4X 10GE
  CPU            REV 04   710-024633   CAAZ5269          SRX5k DPC PMB
  PIC 0                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
    Xcvr 0       REV 02   740-014289   T10A00404         XFP-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
  PIC 2                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
  PIC 3                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
FPC 6            REV 02   750-044175   ZY2552            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
FPC 9            REV 10   750-044175   CAAP5932          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 10           REV 22   750-043157   ZH8192            SRX5k IOC II   CPU            REV 08
711-043360   YX3879            SRX5k MPC PMB
  MIC 0          REV 01   750-049488   YZ2084            10x 10GE SFP+
    PIC 0                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 0     REV 01   740-031980   AMB0HG3           SFP+-10G-SR
      Xcvr 1     REV 01   740-031980   AM20B6F           SFP+-10G-SR
  MIC 1          REV 19   750-049486   CAAH3504          1x 100GE CFP
    PIC 2                 BUILTIN      BUILTIN           1x 100GE CFP
      Xcvr 0     REV 01   740-035329   X000D375          CFP-100G-SR10
FPC 11           REV 07.04.07 750-043157 CAAJ8771        SRX5k IOC II   CPU            REV 08
711-043360   CAAJ3881          SRX5k MPC PMB
  MIC 0          REV 19   750-049486   CAAH0979          1x 100GE CFP
    PIC 0                 BUILTIN      BUILTIN           1x 100GE CFP
      Xcvr 0     REV 01   740-035329   UP1020Z           CFP-100G-SR10
  MIC 1          REV 08   750-049487   CAAM1160          2x 40GE QSFP+
    PIC 2                 BUILTIN      BUILTIN           2x 40GE QSFP+
      Xcvr 0     REV 01   740-032986   QB151094          QSFP+-40G-SR4
      Xcvr 1     REV 01   740-032986   QB160509          QSFP+-40G-SR4
```

```
Fan Tray 0       REV 04   740-035409   ACAE0875          Enhanced Fan Tray
Fan Tray 1       REV 04   740-035409   ACAE0876          Enhanced Fan Tray
```

show chassis hardware (with 20-Gigabit Ethernet MIC with SFP)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN108DA5AAGA      SRX 5800
Midplane         REV 02   710-013698   TR0037            SRX 5600 Midplane
FPM Board        REV 02   710-014974   JY4635            Front Panel Display
PDM              Rev 02   740-013110   QCS10465005       Power Distribution Module
PEM 0            Rev 03   740-023514   QCS11154040       PS 1.7kW; 200-240VAC in
PEM 2            Rev 02   740-023514   QCS10504014       PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 05   740-015113   1000681023        RE-S-1300
CB 0             REV 05   710-013385   JY4775            SRX5k SCB
FPC 1            REV 17   750-020751   WZ6349            SRX5k DPC 4X 10GE
  CPU            REV 02   710-024633   WZ0718            SRX5k DPC PMB
  PIC 0                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
    Xcvr 0                NON-JNPR     C724XM088         XFP-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
    Xcvr 0       REV 02   740-011571   C831XJ08S         XFP-10G-SR
  PIC 2                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
  PIC 3                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
FPC 3            REV 22   750-043157   ZH8189            SRX5k IOC II
  CPU            REV 06   711-043360   YX3912            SRX5k MPC PMB
  MIC 0          REV 01   750-055732   CACF9115          20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 2     REV 02   740-013111   B358549           SFP-T
      Xcvr 9     REV 02   740-011613   PNB1FQS           SFP-SX
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 9     REV 02   740-011613   PNB1FFF           SFP-SX
FPC 5            REV 01   750-027945   JW9665            SRX5k FIOC
  CPU
FPC 8            REV 08   750-023996   XA7234            SRX5k SPC
  CPU            REV 02   710-024633   XA1599            SRX5k DPC PMB
  PIC 0                   BUILTIN      BUILTIN           SPU Cp-Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
Fan Tray 0       REV 03   740-014971   TP0902            Fan Tray
Fan Tray 1       REV 01   740-014971   TP0121            Fan Tray
```

show chassis hardware (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE [SCB2] and SRX5K-RE-1800X4 [RE2])

```
user@host> show chassis hardware
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1251EA1AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2657          Enhanced SRX5600 Midplane
FPM Board        REV 01   710-024631   CABY3551          Front Panel Display
PEM 0            Rev 03   740-034701   QCS13380901P      PS 1.4-2.6kW; 90-264V AC in
PEM 1            Rev 03   740-034701   QCS133809019      PS 1.4-2.6kW; 90-264V AC in
Routing Engine 0 REV 02   740-056658   9009210105        SRX5k RE-1800X4
Routing Engine 1 REV 02   740-056658   9013115551        SRX5k RE-1800X4
CB 0             REV 01   750-062257   CADW3663          SRX5k SCB3
CB 1             REV 01   750-062257   CADZ3263          SRX5k SCB3
FPC 0            REV 18   750-054877   CABG6043          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 01   750-062243   CAEE5918          SRX5k IOC3 24XGE+6XLG
  CPU            REV 02   711-062244   CADX8509          RMPC PMB
  PIC 0                   BUILTIN      BUILTIN           12x 10GE SFP+
    Xcvr 0       REV 01   740-031980   273363A01891      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   273363A01915      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   ANA0BK6           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AP407GA           SFP+-10G-SR
    Xcvr 9       REV 01   740-021308   MUC20G1           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 2                   BUILTIN      BUILTIN           3x 40GE QSFP+
  PIC 3                   BUILTIN      BUILTIN           3x 40GE QSFP+
  WAN MEZZ       REV 15   750-049136   CAEE5845          MPC5E 24XGE OTN Mezz
FPC 3            REV 11   750-043157   CACL7452          SRX5k IOC II
  CPU            REV 04   711-043360   CACP1977          SRX5k MPC PMB
  MIC 0          REV 04   750-049488   CABL4759          10x 10GE SFP+
    PIC 0                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 0     REV 01   740-021308   CF36KM0SY         SFP+-10G-SR
      Xcvr 1     REV 01   740-021308   MUC0MF2           SFP+-10G-SR
      Xcvr 2     REV 01   740-021308   CF36KM01S         SFP+-10G-SR
```

```
      Xcvr 3     REV 01   740-021308   MUC229N           SFP+-10G-SR
FPC 5            REV 07   750-044175   CAAD0764          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                  Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN124FE77AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2970          Enhanced SRX5600 Midplane
FPM Board        REV 01   710-024631   CABY3552          Front Panel Display
PEM 0            Rev 03   740-034701   QCS133809028      PS 1.4-2.6kW; 90-264V AC in
PEM 1            Rev 03   740-034701   QCS133809027      PS 1.4-2.6kW; 90-264V AC in
Routing Engine 0 REV 02   740-056658   9009218294        SRX5k RE-1800X4
Routing Engine 1 REV 02   740-056658   9013104758        SRX5k RE-1800X4
CB 0             REV 01   750-062257   CAEB8180          SRX5k SCB3
CB 1             REV 01   750-062257   CADZ3334          SRX5k SCB3
FPC 0            REV 18   750-054877   CACJ9834          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 01   750-062243   CAEB0981          SRX5k IOC3 24XGE+6XLG
  CPU            REV 02   711-062244   CAEA4644          RMPC PMB
  PIC 0                   BUILTIN      BUILTIN           12x 10GE SFP+
    Xcvr 0       REV 01   740-031980   AP41BLH           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AQ400SL           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AP422LJ           SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AMG0RBT           SFP+-10G-SR
    Xcvr 9       REV 01   740-021308   MUC2FRG           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 2                   BUILTIN      BUILTIN           3x 40GE QSFP+
  PIC 3                   BUILTIN      BUILTIN           3x 40GE QSFP+
  WAN MEZZ       REV 15   750-049136   CAEA4837          MPC5E 24XGE OTN Mezz
FPC 3            REV 11   750-043157   CACA8784          SRX5k IOC II
  CPU            REV 04   711-043360   CACA8820          SRX5k MPC PMB
  MIC 0          REV 05   750-049488   CADF0521          10x 10GE SFP+
```

```
    PIC 0                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 0     REV 01   740-030658   AD1130A00PV       SFP+-10G-USR
      Xcvr 1     REV 01   740-031980   AN40MVV           SFP+-10G-SR
      Xcvr 2     REV 01   740-021308   CF36KM37B         SFP+-10G-SR
      Xcvr 3     REV 01   740-021308   AD153830DSZ       SFP+-10G-SR
  MIC 1          REV 01   750-049487   CABB5961          2x 40GE QSFP+
    PIC 2                 BUILTIN      BUILTIN           2x 40GE QSFP+
      Xcvr 1     REV 01   740-032986   QB160513          QSFP+-40G-SR4
FPC 5            REV 02   750-044175   ZY2569            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                  Enhanced Fan Tray
```

show chassis hardware (SRX5400, SRX5600, and SRX5800 devices with SRX5000 line SRX5K-SCB3 [SCB3] with enhanced midplanes and SRX5K-MPC3-100G10G [IOC3] or SRX5K-MPC3-40G10G [IOC3])

```
user@host> show chassis hardware
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1250870AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2578          Enhanced SRX5600
Midplane

FPM Board        REV 02   710-017254   KD9027            Front Panel Display
PEM 0            Rev 03   740-034701   QCS13090900T      PS 1.4-2.6kW; 90-264V
A
                     C in
PEM 1            Rev 03   740-034701   QCS13090904T      PS 1.4-2.6kW; 90-264V
A
                     C in
Routing Engine 0 REV 01   740-056658   9009196496        SRX5k RE-1800X4
CB 0             REV 01   750-062257   CAEC2501          SRX5k SCB3
FPC 0            REV 10   750-056758   CADC8067          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
```

```
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 01   750-062243   CAEE5924          SRX5k IOC3 24XGE+6XLG
  CPU            REV 01   711-062244   CAEB4890          SRX5k IOC3 PMB
  PIC 0                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 1                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 2                   BUILTIN      BUILTIN           3x 40GE QSFP+
    Xcvr 0       REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
    Xcvr 2       REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
  PIC 3                   BUILTIN      BUILTIN           3x 40GE QSFP+
  WAN MEZZ       REV 01   750-062682   CAEE5817          24x 10GE SFP+ Mezz
FPC 4            REV 11   750-043157   CACY1595          SRX5k IOC II
  CPU            REV 04   711-043360   CACZ8879          SRX5k MPC PMB
  MIC 1          REV 04   750-049488   CACM6062          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 7     REV 01   740-021308   AD1439301TU       SFP+-10G-SR
      Xcvr 8     REV 01   740-021308   AD1439301SD       SFP+-10G-SR
      Xcvr 9     REV 01   740-021308   AD1439301TS       SFP+-10G-SR
FPC 5            REV 05   750-044175   ZZ1371            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                  Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN124FEC0AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2946          Enhanced SRX5600 Midplane
FPM Board        test     710-017254   test              Front Panel Display
PEM 0            Rev 01   740-038514   QCS114111003      DC 2.6kW Power Entry Module
PEM 1            Rev 01   740-038514   QCS12031100J      DC 2.6kW Power Entry Module
Routing Engine 0 REV 01   740-056658   9009186342        SRX5k RE-1800X4
CB 0             REV 01   750-062257   CAEB8178          SRX5k SCB3
FPC 0            REV 07   750-044175   CAAD0769          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
```

```
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 4            REV 11   750-043157   CACY1592          SRX5k IOC II
  CPU            REV 04   711-043360   CACZ8831          SRX5k MPC PMB
  MIC 1          REV 04   750-049488   CACN0239          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 7     REV 01   740-031980   ARN23HW           SFP+-10G-SR
      Xcvr 8     REV 01   740-031980   ARN2FVW           SFP+-10G-SR
      Xcvr 9     REV 01   740-031980   ARN2YVM           SFP+-10G-SR
FPC 5            REV 10   750-056758   CADA8736          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                  Enhanced Fan Tray
```

show chassis hardware (SRX4200)

command-name

```
user@host> show chassis hardware

Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                DK2816AR0020      SRX4200
Mainboard        REV 01   650-071675   16061032317       SRX4200
Routing Engine 0          BUILTIN      BUILTIN           SRX Routing Engine
FPC 0                     BUILTIN      BUILTIN           FEB
  PIC 0                   BUILTIN      BUILTIN           8x10G-SFP
    Xcvr 0       REV 01   740-038153   MOC11511530020    SFP+-10G-CU3M
    Xcvr 1       REV 01   740-038153   MOC11511530020    SFP+-10G-CU3M
    Xcvr 2       REV 01   740-038153   MOC11511530020    SFP+-10G-CU3M
    Xcvr 3       REV 01   740-038153   MOC11511530020    SFP+-10G-CU3M
    Xcvr 4       REV 01   740-021308   04DZ06A00364      SFP+-10G-SR
    Xcvr 5       REV 01   740-031980   233363A03066      SFP+-10G-SR
    Xcvr 6       REV 01   740-021308   AL70SWE           SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   ALN0N6C           SFP+-10G-SR
    Xcvr 8       REV 01   740-030076   APF16220018NK1    SFP+-10G-CU1M
Power Supply 0   REV 04   740-041741   1GA26241849       JPSU-650W-AC-AFO
Power Supply 1   REV 04   740-041741   1GA26241846       JPSU-650W-AC-AFO
Fan Tray 0                                              SRX4200 0, Front to Back Airflow - AFO
```

```
Fan Tray 1                                                SRX4200 1, Front to Back Airflow - AFO
Fan Tray 2                                                SRX4200 2, Front to Back Airflow - AFO
Fan Tray 3                                                SRX4200 3, Front to Back Airflow - AFO
```

show chassis hardware (vSRX Virtual Firewall 3.0)

command-name

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                806dddb1a141      VSRX
Midplane
System IO
Routing Engine                                            VSRX-2CPU-8G memory
FPC 0                                                     FPC
  PIC 0                                                   VSRX DPDK GE
Power Supply 0
```

show chassis hardware clei-models

show chassis hardware clei-models (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE [SCB2] and SRX5K-RE-1800X4 [RE2])

```
user@host> show chassis hardware clei-models node 1
node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 01   710-024803                     SRX5800-BP-A
FPM Board        REV 01   710-024632                     SRX5800-CRAFT-A
PEM 0            Rev 04   740-034724                     SRX5800-PWR-4100-AC
PEM 1            Rev 05   740-034724                     SRX5800-PWR-4100-AC
Routing Engine 0 REV 01   740-056658   COUCATTBAA        SRX5K-RE-1800X4
CB 0             REV 01   750-056587   COUCATSBAA        SRX5K-SCBE
CB 1             REV 01   750-056587   COUCATSBAA        SRX5K-SCBE
CB 2             REV 01   750-056587   COUCATSBAA        SRX5K-SCBE
FPC 0            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
```

```
FPC 1            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 2            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 3            REV 11   750-043157   COUIBCWBAA        SRX5K-MPC
  MIC 0          REV 05   750-049486   COUIBCYBAA        SRX-MIC-1X100G-CFP
  MIC 1          REV 04   750-049488   COUIBCXBAA        SRX-MIC-10XG-SFPP
FPC 4            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 7            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 8            REV 11   750-043157   COUIBCWBAA        SRX5K-MPC
  MIC 0          REV 05   750-049486   COUIBCYBAA        SRX-MIC-1X100G-CFP
FPC 9            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 10           REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
Fan Tray 0       REV 04   740-035409                     SRX5800-HC-FAN
Fan Tray 1       REV 04   740-035409                     SRX5800-HC-FAN
```

Release Information

Command introduced in Junos OS Release 9.2. Command modified in Junos OS Release 9.2 to include `node` option.

RELATED DOCUMENTATION

- Understanding Traffic Processing on Security Devices
- Interface Naming Conventions

show chassis high-availability data-plane statistics

IN THIS SECTION

- Syntax | 18477
- Description | 18477

- Required Privilege Level | 18477
- Output Fields | 18477
- Sample Output | 18480
- Release Information | 18482

Syntax

```
show chassis high-availability data-plane statistics
```

Description

Display Multinode High Availability data plane statistics.

Required Privilege Level

view

Output Fields

Table 293: show chassis cluster data-plane statistics Output Fields

Field Name	Field Description
`Services Synchronized`	• `Service name`—Name of the service. • `Translation context`—Messages synchronizing Network Address Translation (NAT) translation context. • `Incoming NAT`—Messages synchronizing incoming Network Address Translation (NAT) service. • `Resource manager`—Messages synchronizing resource manager groups and resources. • `DS-LITE create`—Messages synchronizing DS-LITE create sessions. • `Session create`—Messages synchronizing session creation. • `IPv6 session create`—Messages synchronizing IPv6 session create sessions. • `IPv4/6 session RTO ACK`—Messages synchronizing IPv4/6 session RTO ACK sessions. • `Session close`—Messages synchronizing session close. • `IPv6 session close`—Messages synchronizing IPv6 session close sessions. • `Session change`—Messages synchronizing session change. • `IPv6 session change`—Messages synchronizing IPv6 session change sessions. • `ALG Support Library`—Messages synchronizing ALG Support Library sessions. • `Gate create`—Messages synchronizing creation of pinholes (temporary openings in the firewall). • `Session ageout refresh request`—Messages synchronizing request session after age-out. • `IPv6 session ageout refresh requests`—Messages synchronizing IPv6 session ageout refresh requests. • `Session ageout refresh replies`—Messages synchronizing reply session after age-out.

Table 293: show chassis cluster data-plane statistics Output Fields *(Continued)*

Field Name	Field Description
	• `IPv6 session ageout refresh replies`—Messages synchronizing IPv6 session ageout refresh replies sessions. • `IPsec VPN`—Messages synchronizing VPN session. • `Firewall user authentication`—Messages synchronizing firewall user authentication session. • `MGCP ALG`—Messages synchronizing MGCP ALG sessions. • `H323 ALG`—Messages synchronizing H.323 ALG sessions. • `SIP ALG`—Messages synchronizing SIP ALG sessions. • `SCCP ALG`—Messages synchronizing SCCP ALG sessions. • `PPTP ALG`—Messages synchronizing PPTP ALG sessions. • `JSF PPTP ALG`—Messages synchronizing JSF PPTP ALG sessions. • `RPC ALG`—Messages synchronizing RPC ALG sessions. • `RTSP ALG`—Messages synchronizing RTSP ALG sessions. • `RAS ALG`—Messages synchronizing RAS ALG sessions. • `MAC address learning`—Messages synchronizing MAC address learning sessions. • `GPRS GTP`—Messages synchronizing GPRS GTP sessions. • `GPRS SCTP`—Messages synchronizing GPRS SCTP sessions. • `GPRS FRAMEWORK`—Messages synchronizing GPRS FRAMEWORK sessions. • `JSF RTSP ALG`—Messages synchronizing JSF RTSP ALG sessions. • `JSF SUNRPC MAP`—Messages synchronizing JSF SUNRPC MAP sessions. • `JSF MSRPC MAP`—Messages synchronizing JSF MSRPC MAP sessions. • `DS-LITE delete`—Messages synchronizing DS-LITE delete sessions. • `JSF SLB`—Messages synchronizing JSF SLB sessions.

Table 293: show chassis cluster data-plane statistics Output Fields *(Continued)*

Field Name	Field Description
	• APPID—Messages synchronizing APPID sessions. • JSF MGCP MAP—Messages synchronizing JSF MGCP MAP sessions. • JSF H323 ALG—Messages synchronizing JSF H323 ALG sessions. • JSF RAS ALG—Messages synchronizing JSF RAS ALG sessions. • JSF SCCP MAP—Messages synchronizing JSF SCCP MAP sessions. • JSF SIP MAP—Messages synchronizing JSF SIP MAP sessions. • PST_NAT_CREATE—Messages synchronizing PST NAT CREATE sessions. • PST_NAT_CLOSE—Messages synchronizing PST NAT CLOSE sessions. • PST_NAT_UPDATE—Messages synchronizing PST NAT UPDATE sessions. • JSF TCP STACK—Messages synchronizing JSF TCP STACK sessions. • JSF IKE ALG—Messages synchronizing JSF IKE ALG sessions.

Sample Output

show chassis high-availability data-plane statistics

```
user@host> show chassis high-availability data-plane
statistics
Services Synchronized:
    Service name                              RTOs sent    RTOs received
    Translation context                       0            0
    Incoming NAT                              0            0
    Resource manager                          0            0
    DS-LITE create                            0            0
    Session create                            0            0
    IPv6 session create                       0            0
    IPv4/6 session RTO ACK                    0            0
    Session close                             0            0
    IPv6 session close                        0            0
```

```
    Session change                          0           0
    IPv6 session change                     0           0
    ALG Support Library                     0           0
    Gate create                             0           0
    Session ageout refresh requests         0           0
    IPv6 session ageout refresh requests    0           0
    Session ageout refresh replies          0           0
    IPv6 session ageout refresh replies     0           0
    IPSec VPN                               0           0
    Firewall user authentication            0           0
    MGCP ALG                                0           0
    H323 ALG                                0           0
    SIP ALG                                 0           0
    SCCP ALG                                0           0
    PPTP ALG                                0           0
    JSF PPTP ALG                            0           0
    RPC ALG                                 0           0
    RTSP ALG                                0           0
    RAS ALG                                 0           0
    MAC address learning                    0           0
    GPRS GTP                                0           0
    GPRS SCTP                               0           0
    GPRS FRAMEWORK                          0           0
    JSF RTSP ALG                            0           0
    JSF SUNRPC MAP                          0           0
    JSF MSRPC MAP                           0           0
    DS-LITE delete                          0           0
    JSF SLB                                 0           0
    APPID                                   0           0
    JSF MGCP MAP                            0           0
    JSF H323 ALG                            0           0
    JSF RAS ALG                             0           0
    JSF SCCP MAP                            0           0
    JSF SIP MAP                             0           0
    PST_NAT_CREATE                          0           0
    PST_NAT_CLOSE                           0           0
    PST_NAT_UPDATE                          0           0
    JSF TCP STACK                           0           0
    JSF IKE ALG                             0           0
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

clear chassis high-availability data-plane statistics | 3472

show chassis high-availability information

IN THIS SECTION

- Syntax | 18482
- Description | 18482
- Required Privilege Level | 18482
- Output Fields | 18483
- Sample Output | 18486
- Release Information | 18492

Syntax

```
show chassis high-availability information
```

Description

Display details of the Multinode High Availability status on your security device including health status of the peer node.

Required Privilege Level

view

Output Fields

Table 294 on page 18483 lists the output fields for the `show chassis high-availability information` command. Output fields are listed in the approximate order in which they appear.

Table 294: show chassis high-availability information

Field Name	Field Description
`Node failure codes`	Node failure codes are: • HW Hardware monitoring • MB Mbuf monitoring • CS Cold Sync monitoring • LB Loopback monitoring • SP SPU monitoring
`Node Status`	• status of the node • Local-id—local identifier • Local-IP—local IP address of the node • Local Forwarding IP—Source IP address used to route the traffic (local interchassis datalink[ICD])

Table 294: show chassis high-availability information *(Continued)*

Field Name	Field Description
`HA Peer Information`	• Description—peer information • Peer Id—peer identifier • IP address—peer IP address • Interface—interface name • Peer Forwarding IP—Destination IP address used to route the traffic (peer node's ICD) • Interface—Interface used for ICD. • Peer ICD Conn State—Status of ICD. • Routing Instance—routing instance name • Encrypted—data encrypted status • Cold Sync Status—cold sync status of the node.
`HA Hardware Upgrade Events`	Message on additional SPC3 installation status.
`Services Redundancy Group`	• Current State—current state of the node • Peer Information—peer information • Peer Id—peer identifier
`SRG failure event codes`	• BF BFD monitoring—monitor Bidirectional Forwarding Detection. • IP IP monitoring—monitor IP address • CP Control Plane monitoring—monitor control plane state

Table 294: show chassis high-availability information *(Continued)*

Field Name	Field Description

`Services Redundancy Group`	• Status—node status
	• Activeness Priority—the node with the higher activeness priority become active for the service redundancy group.
	• Services—Associated services for the SRG.
	  • IPsec.
	• Process Packet In Backup State—packet processing in Backup state.
	• Control Plane State—Displays the Control plane cold sync readiness after taking over the backup role. Means the control plane finished syncing all the control plane data from the new active node.
	  The options are:
	  • READY—Active mode
	  • READY/NOT READY—Backup node based on the actual state
	  • N/A—in other states
	• System Integrity Check—Displays if the hold timer is running and system integrity check is under progress. Applicable only in BACKUP and INELLIGIBLE state.
	• Failure Events—Displays local service redundancy group related attribute failure events.
	• Peer Information—Displays peer node information.
	  • Peer Id— Peer node identification number.
	  • Status—Displays status of the peer node. The options are Active/Backup/Ineligible/Unknown.
	  • Health Status—Displays health status of the peer node as advertised last
	  • HEALTHY—Peer node SRG state is healthy
	  • UNHEALTHY—Peer node SRG state is unhealthy

Table 294: show chassis high-availability information *(Continued)*

Field Name	Field Description
	• UNKNOWN—Peer link is down • SRG NOT CONFIGURED—SRG not configured at peer node. • Failover Readiness—Displays Failover readiness of the peer node and applicable only in Active state. The options are: • READY—Peer's healthy, control and data plane cold sync is complete and hence peer is ready for a failover. A manual failover can be done • NOT READY—Peer is either unhealthy, or data/control plane cold sync is pending • UNKNOWN—Peer's failover readiness is awaited or if peer link is down • N/A—In all states other than Active.

Sample Output

show chassis high-availability information

```
user@host> show chassis high-availability information
Node failure codes:
    HW  Hardware monitoring    LB  Loopback monitoring
    MB  Mbuf monitoring        SP  SPU monitoring
    CS  Cold Sync monitoring   SU  Software Upgrade

Node Status: ONLINE
Local-id: 2
Local-IP: 10.22.0.2
HA Peer Information:

    Peer Id: 1        IP address: 10.22.0.1     Interface: ge-0/0/2.0
    Routing Instance: default
    Encrypted: YES    Conn State: DOWN
    Cold Sync Status: IN PROGRESS
```

```
Services Redundancy Group: 0
        Current State: ONLINE
        Peer Information:
          Peer Id: 1

SRG failure event codes:
    BF  BFD monitoring
    IP  IP monitoring
    IF  Interface monitoring
    CP  Control Plane monitoring

Services Redundancy Group: 1
        Deployment Type: ROUTING
        Status: ACTIVE
        Activeness Priority: 1
        Preemption: DISABLED
        Process Packet In Backup State: NO
        Control Plane State: READY
        System Integrity Check: N/A
        Failure Events: NONE
        Peer Information:
          Peer Id: 1
          Status : BACKUP
          Health Status: HEALTHY
          Failover Readiness: READY
```

show chassis high-availability information detail

```
user@host> show chassis high-availability information detail
Node level Information:
        Node Status: ONLINE
        Local-id: 2
        Local-IP: 10.22.0.2
HA Peer Information:

   Peer-ID: 1        IP address: 10.22.0.1     Interface: ge-0/0/2.0
   Routing Instance: default
   Encrypted: YES    Conn State: UP
   Cold Sync Status: COMPLETE
   Internal Interface: st0.16000
   Internal Local-IP: 180.100.1.2
```

```
   Internal Peer-IP: 180.100.1.1
   Internal Routing-instance: __juniper_private1__
Packet Statistics:
        Receive Error : 0        Send Error : 0

        Packet-type            Sent       Received
        SRG Status Msg            9              8
        SRG Status Ack            8              7
        Attribute Msg             3              2
        Attribute Ack             2              2

HA Peer Conn events:
        Dec 11 22:22:18.696 : HA Peer 180.100.1.1 BFD conn came up
        Dec 11 22:22:25.269 : HA Peer 180.100.1.1 BFD conn went down
        Dec 11 22:22:31.619 : HA Peer 180.100.1.1 BFD conn came up

Cold Synchronization:
    Status:
        Cold synchronization completed for: N/A
        Cold synchronization failed for: N/A
        Cold synchronization not known for: N/A
        Current Monitoring Weight: 0

    Progress:
        CS Prereq               1 of 1 SPUs completed
           1. if_state sync          1 SPUs completed
           2. ha peer conn           1 SPUs completed
           3. policy data sync       1 SPUs completed
           4. cp ready               1 SPUs completed
           5. VPN data sync          1 SPUs completed
           6. IPID data sync         1 SPUs completed
           7. All SPU ready          1 SPUs completed
           8. AppID ready            1 SPUs completed
           9. Tunnel Sess ready      1 SPUs completed
        CS RTO sync             1 of 1 SPUs completed
        CS Postreq              1 of 1 SPUs completed

    Statistics:
        Number of cold synchronization completed: 0
        Number of cold synchronization failed: 0

    Events:
        Dec 11 22:22:23.870 : Cold sync for PFE  is Post-req check in process
```

```
        Dec 11 22:22:24.872 : Cold sync for PFE  is Completed
        Dec 11 22:22:37.293 : Cold sync for PFE  is Post-req check in process
        Dec 11 22:22:38.291 : Cold sync for PFE  is Completed

SPU monitoring:
    Status: Enabled
        Current monitoring weight: 0



    Statistics:
        SPU up count: 1
        NPC up count: 0
        SPU down count: 0
        NPC down count: 0
        Chassis info processing error count: 0

Loopback Information:

    PIC Name        Loopback        Nexthop     Mbuf
    ------------------------------------------------
                    Success         Success     Success

Hardware monitoring:
    Status:
        Activation status: Enabled
        Ctrl Plane Hardware errors: 0
        Data Plane Hardware errors: 0

SRGS Information:

Services Redundancy Group: 1
        Deployment Type: ROUTING
        Status: BACKUP
        Activeness Priority: 1
        Hold Timer: 1
        Services: [ IPSEC ]
        Process Packet In Backup State: NO
        Control Plane State: READY
        System Integrity Check: COMPLETE
        Peer Information:
        Failure Events: NONE
          Peer Id: 1
          Last Advertised HA Status: ACTIVE
```

```
          Last Advertised Health Status: HEALTHY
          Failover Readiness: N/A

        Signal Route Info:
          Active Signal Route:
          IP: 10.39.1.1
          Routing Instance: default
          Status: NOT INSTALLED

          Backup Signal Route:
          IP: 10.39.1.2
          Routing Instance: default
          Status: INSTALLED

        Split-brain Prevention Probe Info:
          DST-IP: 10.111.0.1
          SRC-IP: 10.11.0.1
          Routing Instance: default
          Status: NOT RUNNING
          Result: N/A           Reason: N/A


SRG State Change Events:
        Dec 11 22:21:36.379 : SRG[1] state UNKNOWN -> HOLD, Reason: State machine start
        Dec 11 22:23:28.889 : SRG[1] state HOLD -> ACTIVE, Reason: Peer state Backup adv received
        Dec 11 22:23:31.152 : SRG[1] state ACTIVE -> BACKUP, Reason: HA state preemption

        BFD Monitoring:
          Status: UNKNOWN

          SRC-IP: 10.5.0.1     DST-IP: 10.5.0.2
          Routing Instance: default
          Type: SINGLE-HOP
              IFL Name: ge-0/0/3.0
          State: INSTALLED


        Interface Monitoring:
          Status: UP

          IF Name: ge-0/0/4     State: Up
```

```
          IF Name: ge-0/0/3     State: Up
```

show chassis high-availability information (Junos OS 23.4R1)

```
user@host> show chassis high-availability information
Node failure codes:
    HW  Hardware monitoring    LB  Loopback monitoring
    MB  Mbuf monitoring        SP  SPU monitoring
    CS  Cold Sync monitoring   SU  Software Upgrade

Node Status: ONLINE
Local-id: 1
Local-IP: 10.22.0.1
Local Forwarding IP: 100.1.1.1
HA Peer Information:

    Peer Id: 2        IP address: 10.22.0.2     Interface: ge-0/0/2.0
    Routing Instance: default
    Encrypted: YES    Conn State: UP
    Configured BFD Detection Time: 5 * 400ms
    Cold Sync Status: COMPLETE
    Peer Forwarding IP: 200.1.1.1               Interface: lo0.0
    Peer ICD Conn State: UP

Services Redundancy Group: 0
        Current State: ONLINE
        Peer Information:
          Peer Id: 2

SRG failure event codes:
    BF  BFD monitoring
    IP  IP monitoring
    IF  Interface monitoring
    CP  Control Plane monitoring

Services Redundancy Group: 1
        Deployment Type: ROUTING
        Status: ACTIVE
        Activeness Priority: 200
        Preemption: ENABLED
```

```
        Process Packet In Backup State: NO
        Control Plane State: READY
        System Integrity Check: N/A
        Failure Events: NONE
        Peer Information:
          Peer Id: 2
          Status : BACKUP
          Health Status: HEALTHY
          Failover Readiness: READY
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Example: Configure Multinode High Availability in a Layer 3 Network

show chassis high-availability peer-info

IN THIS SECTION

- Syntax | 18493
- Description | 18493
- Required Privilege Level | 18493
- Output Fields | 18493
- Sample Output | 18494
- Release Information | 18494

Syntax

```
show chassis high-availability peer-info
```

Description

Display details of the peer node in a Multinode High Availability setup. Use this command to gather details of peer node, connection details, and packet statistics.

Required Privilege Level

view

Output Fields

Table 295 on page 18493 lists the output fields for the `show chassis high-availability information` command. Output fields are listed in the approximate order in which they appear.

Table 295: show chassis high-availability peer-info

Field Name	Field Description
`HA Peer Information`	• Description—peer information • Peer Id—peer identifier • IP address—peer IP address • Interface—interface name • Routing Instance—routing instance name
Internal connection details	Details related to internal traffic including internal interface, internal local IP address, internal peer IP address, and internal routing instance)
Packet Statistics	Details of packets sent and received. The detail includes: number of inbound/outbound errors and number of packets sent and received (SRG messages and attribute messages).

Sample Output

show chassis high-availability peer-info

```
user@host> show chassis high-availability peer-info
HA Peer Information:

   Peer-ID: 2        IP address: 10.22.0.2     Interface: ge-0/0/2.0
   Routing Instance: default
   Encrypted: YES    Conn State: UP
   Cold Sync Status: COMPLETE
   Internal Interface: st0.16000
   Internal Local-IP: 180.100.1.1
   Internal Peer-IP: 180.100.1.2
   Internal Routing-instance: __juniper_private1__
Packet Statistics:
        Receive Error : 0        Send Error : 0

        Packet-type            Sent       Received

        SRG Status Msg            4           3

        SRG Status Ack            3           4

        Attribute Msg             1           1

        Attribute Ack             1           1
```

Release Information

Command introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Example: Configure Multinode High Availability in a Layer 3 Network

show chassis high-availability prefix-srgid-table

IN THIS SECTION

- Syntax | 18495
- Description | 18495
- Required Privilege Level | 18495
- Output Fields | 18495
- Sample Output | 18496
- Release Information | 18496

Syntax

```
show chassis high-availability prefix-srgid-table
```

Description

Display mapping of prefix list with SRGs.

Required Privilege Level

view

Output Fields

Table 296 on page 18495 lists the output fields for the `show chassis high-availability information` command. Output fields are listed in the approximate order in which they appear.

Table 296: show chassis high-availability prefix-srgid-table

Field Name	Field Description
SRGID	SRG identification number.

Table 296: show chassis high-availability prefix-srgid-table *(Continued)*

Field Name	Field Description
IP Prefix	Listing of IP prefixes that represent a set of routes used as local address of IKE gateway.
Routing Table	Name of the routing table.

Sample Output

show chassis high-availability peer-info

```
user@host> show chassis high-availability prefix-srgid-table
IP SRGID Table:
        SRGID    IP Prefix                                  Routing Table
        1        10.11.0.1/32                                default

        1        10.19.0.1/32                                default

        1        10.20.0.1/32                                default

        2        10.11.1.1/32                                default

        2        10.19.1.1/32                                default

        2        10.20.1.1/32                                default
```

Release Information

Command introduced in Junos OS Release 22.4R1.

show chassis high-availability services-redundancy-group

IN THIS SECTION

- Syntax | 18497
- Description | 18497
- Required Privilege Level | 18497
- Output Fields | 18497
- Sample Output | 18503
- Release Information | 18507

Syntax

```
show chassis high-availability services-redundancy-group services-redundancy-group-id | monitor-object object-name
```

Description

Display the service redundancy group information in a Multinode High Availability setup.

Required Privilege Level

view

Output Fields

Table 297 on page 18498 lists the output fields for the `show chassis high-availability services-redundancy-group` command. Output fields are listed in the approximate order in which they appear.

Table 297: show chassis high-availability services-redundancy-group

Field Name	Field Description
`SRG failure event codes`	SRG failure event codes are: • BF BFD monitoring • IP IP monitoring • IF Interface Monitoring • CP Control Plane monitoring

Table 297: show chassis high-availability services-redundancy-group *(Continued)*

Field Name	Field Description

`Services Redundancy Group`	• Deployment Type—Multinode High Availability deployent type - Routing, switching, hybrid, or cloud
	• Status—node status
	• Activeness Priority—the node with the higher activeness priority become active for the service redundancy group.
	• Process Packet In Backup State—packet processing in Backup state.
	• Control Plane State—Displays the Control plane cold sync readiness after taking over the backup role. Means the control plane finished syncing all the control plane data from the new active node.
	The options are:
	  • READY—Active mode
	  • READY/NOT READY—Backup node based on the actual state
	  • N/A—in other states
	• System Integrity Check—Displays if the hold timer is running and system integrity check is under progress. Applicable only in BACKUP and INELLIGIBLE state.
	• Failure Events—Displays local service redundancy group related attribute failure events.
	• Monitoring Object Failures: —Displays failed monitor-object information.
	• Peer Information—Displays peer node information.
	  • Peer Id— Peer node identification number.

Table 297: show chassis high-availability services-redundancy-group *(Continued)*

Field Name	Field Description

	• Status—Displays status of the peer node. The options are Active/Backup/Ineligible/Unknown.
	• Health Status—Displays health status of the peer node as advertised last
	  • HEALTHY—Peer node SRG state is healthy
	  • UNHEALTHY—Peer node SRG state is unhealthy
	  • UNKNOWN—Peer link is down
	  • SRG NOT CONFIGURED—SRG not configured at peer node.
	• Failover Readiness—Displays Failover readiness of the peer node and applicable only in Active state.
	  The options are:
	  • READY—Peer's healthy, control and data plane cold sync is complete and hence peer is ready for a failover. A manual failover can be done
	  • NOT READY—Peer is either unhealthy, or data/control plane cold sync is pending
	  • UNKNOWN—Peer's failover readiness is awaited or if peer link is down
	  • N/A—In all states other than Active.
	• Signal Route Info
	  • Active Signal Route—Active signal route required for activeness enforcement

Table 297: show chassis high-availability services-redundancy-group *(Continued)*

Field Name	Field Description

	• Backup Signal Route—Backup signal route required for activeness enforcement
	• Routing Instannce—Routing instance of routes
	• Status—Status of the routes - **INSTALLED** or **NOT INSTALLED**
	• Split-brain Prevention Probe Info
	  • DST-IP—IP addresses of the upstream routers as the destination IP address for the activeness determination probe.
	  • SRC-IP—floating IP address as source IP address for activeness determination probe.
	  • Routing Instance—Routing instance of routes
	  • Type - Type of probing. Available options are BFD-based probing or ICMP-based probing.
	  For BFD-based probing:
	  • Interval: Configured intervals for probes.
	  • Multiplier: Multipliers for probes.
	  • Status—Status of the probes - **RUNNING** or **NOT RUNNUNG**
	• BFD Monitoring
	  • Status—State of the BFD session: Up, Down, Init or failing
	  • SRC-IP—Source IP address for outgoing BFD packets

Table 297: show chassis high-availability services-redundancy-group *(Continued)*

Field Name	Field Description

	• DST-IP—Destination IP address for outgoing BFD packets
	• Routing Instance—Routing instance of routes on which BFD monitoring is configured.
	• Type—Session type (single hop or multihop)
	• IFL Name—
	Interface on which the BFD session is active
	• Status—Status of the probes - **RUNNING** or **NOT RUNNUNG**
	• Object Monitoring
	  • Object Name—Name of the object-monitor.
	  • Object Status—Status of the object-monitor. Indicates if the object-monitoring is running or not (Up or Down).
	  • Object Monitored Entries: Monitoring options created under object-monitor entity. Available options are IP monitoring, BFD monitoring, or interface monitoring.
	  • Object Failures: Failed monitoring entity (IP monitoring, BFD monitoring, or interface monitoring)
	  • Object Threshold: Threshold value set for object-monitor.
	  • Object Current Weight: Weight of object-monitor at present.

Sample Output

show chassis high-availability services-redundancy-group

For SRG0:

```
user@host> show chassis high-availability services-redundancy-group 0
Services Redundancy Group: 0
        Current State: ONLINE
        Peer Information:
          Peer Id: 2
```

For SRG1:

```
user@host> show chassis high-availability services-redundancy-group 1
SRG failure event codes:
    BF  BFD monitoring
    IP  IP monitoring
    IF  Interface monitoring
    CP  Control Plane monitoring

Services Redundancy Group: 1
        Deployment Type: ROUTING
        Status: ACTIVE
        Activeness Priority: 200
        Preemption: ENABLED
        Process Packet In Backup State: NO
        Control Plane State: READY
        System Integrity Check: N/A
        Failure Events: NONE
        Peer Information:
          Peer Id: 2
          Status : BACKUP
          Health Status: HEALTHY
          Failover Readiness: READY

        Signal Route Info:
          Active Signal Route:
          IP: 10.39.1.1
          Routing Instance: default
          Status: INSTALLED
```

```
          Backup Signal Route:
          IP: 10.39.1.2
          Routing Instance: default
          Status: NOT INSTALLED

        Split-brain Prevention Probe Info:
          DST-IP: 10.111.0.1
          SRC-IP: 10.11.0.1
          Routing Instance: default
          Status: NOT RUNNING
          Result: N/A           Reason: N/A


        BFD Monitoring:
          Status: UP

          SRC-IP: 10.4.0.1     DST-IP: 10.4.0.2
          Routing Instance: default
          Type: SINGLE-HOP
              IFL Name: ge-0/0/4.0
          State: UP
```

show chassis high-availability services-redundancy-group (Junos OS Release 23.4R1)

For SRG0:

```
user@host> show chassis high-availability services-redundancy-group 0
SRG failure event codes:
    BF  BFD monitoring
    IP  IP monitoring
IF  Interface monitoring
PM  Path Monitoring

Services Redundancy Group: 0
        Current State: ISOLATED [ Node Failure ]
        Failure Events: [ PM ]
        Monitoring Object Failures: [ upstream ]
        Peer Information:
          Peer Id: 2
        Shut-on-failures interfaces:
```

```
        xe-1/0/4
        xe-1/0/5
```

For SRG1:

```
user@host> show chassis high-availability services-redundancy-group 1
SRG failure event codes:
    BF  BFD monitoring
    IP  IP monitoring
    IF  Interface monitoring
    CP  Control Plane monitoring

Services Redundancy Group: 1
        Deployment Type: ROUTING
        Status: BACKUP
        Activeness Priority: 1
        Preemption: ENABLED
        Process Packet In Backup State: NO
        Control Plane State: READY
        System Integrity Check: COMPLETE
        Failure Events: NONE
        Peer Information:
          Peer Id: 2
          Status : N/A
          Health Status: SRG NOT CONFIGURED
          Failover Readiness: N/A

        Signal Route Info:
          Active Signal Route:
          IP: 39.1.1.3
          Routing Instance: default
          Status: NOT INSTALLED

          Backup Signal Route:
          IP: 39.1.1.4
          Routing Instance: default
          Status: INSTALLED

        Split-brain Prevention Probe Info:
          DST-IP: 111.0.0.1
          SRC-IP: 11.0.1.1
          Routing Instance: default
```

```
          Type: ICMP Probe
          Status: NOT RUNNING
          Result: N/A           Reason: N/A


        SRG Path Monitor Info:
          SRG Monitor Status: UP
          SRG Monitor Threshold: 300
          SRG Monitor Weight: 0
          SRG Monitor Failed Objects: NONE

          Object Name: m1
          Object Status: UP
          Object Monitored Entries: [ BFD ]
          Object Failures: NONE
          Object Threshold: 10
          Object Current Weight: 0

          Object Name: lan
          Object Status: UP
          Object Monitored Entries: [ IP IF ]
          Object Failures: NONE
          Object Threshold: 300
          Object Current Weight: 0

IP SRGID Table:
        SRGID    IP Prefix                                    Routing Table
        2        NULL                                         NULL
```

show chassis high-availability services-redundancy-group *group-id* monitor-object (Junos OS Release 23.4R1)

For SRG1:

```
user@host> show chassis high-availability services-redundancy-group 1 monitor-object m1

        Object Name: m1
```

```
            Object Status: UP
            Object Monitored Entries: [ BFD ]
            Object Failures: NONE
            Object Threshold: 10
            Object Current Weight: 0

            BFD Monitoring:
                Threshold: 10
                Current Weight: 0
                Status:  UP
                State               Weight    Object
                INSTALLED           10        111.0.0.2

                SRC-IP: 4.0.0.2      DST-IP: 111.0.0.2
                Routing Instance: default
                Type: MULTI-HOP
                IFL Name: N/A
                State: INSTALLED
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Example: Configure Multinode High Availability in a Layer 3 Network

show chassis in-service-upgrade

IN THIS SECTION

- Syntax | 18508
- Description | 18508
- Options | 18508
- Required Privilege Level | 18508

- Output Fields | 18508
- Sample Output | 18509
- Release Information | 18512

Syntax

```
show chassis in-service-upgrade
```

Description

Display the status of Flexible PIC Concentrators (FPCs) and their corresponding PICs after the most recent unified in-service software upgrade (ISSU). This command must be issued on the primary Routing Engine.

> **NOTE**: Only Intelligent Queuing (IQ) PICs are displayed by this command output. Unified ISSU status for other PIC types is controlled internally by the FPC.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 298 on page 18509 lists the output fields for the `show chassis in-service-upgrade` command. Output fields are listed in the approximate order in which they appear.

Table 298: show chassis in-service-upgrade Output Fields

Field Name	Field Description
`Item`	Flexible PIC Concentrator (FPC) slot number.
`Status`	FPC and corresponding PIC state. State can be either of the follwing: • `Online`—FPC is online and running. • `Offline`—FPC is powered down.
`Reason`	Reason for the state (if offline).

Sample Output

show chassis in-service-upgrade

```
user@host> show chassis in-service-upgrade
  Item           Status                  Reason
  FPC 0          Online
  FPC 1          Online
  FPC 2          Online
    PIC 0        Online
    PIC 1        Online
  FPC 3          Offline                 Offlined by CLI command
  FPC 4          Online
    PIC 1        Online
  FPC 5          Online
    PIC 0        Online
  FPC 6          Online
    PIC 3        Online
  FPC 7          Online
```

show chassis in-service-upgrade (MX2010 Router)

```
user@host> show chassis in-service-upgrade
Item           Status                  Reason
```

```
  FPC 0          Online
  FPC 1          Online
  FPC 8          Online
  FPC 9          Online
```

show chassis in-service-upgrade (MX2020 Router)

```
user@host> show chassis in-service-upgrade
  Item           Status                   Reason
  FPC 0          Online
  FPC 1          Online
  FPC 2          Online
  FPC 3          Online
  FPC 4          Online
  FPC 5          Online
  FPC 6          Online
  FPC 7          Online
  FPC 8          Online
  FPC 9          Online
  FPC 10         Online
  FPC 11         Online
  FPC 12         Online
  FPC 13         Online
  FPC 14         Online
  FPC 15         Online
  FPC 16         Online
  FPC 17         Online
  FPC 18         Online
  FPC 19         Online
```

show chassis in-service-upgrade (MX2008 Router)

```
user@host> show chassis in-service-upgrade
  Item           Status                   Reason
  FPC 0          Online
  FPC 3          Online
  FPC 5          Online
  FPC 7          Online
  FPC 9          Online
```

show chassis in-service-upgrade (TX Matrix Plus Router)

```
user@host> show chassis in-service-upgrade
lcc0-re0:
--------------------------------------------------------------------------
  Item           Status                   Reason
  FPC 1          Online
    PIC 0        Online
  FPC 2          Online
  FPC 3          Online
    PIC 1        Online
  FPC 4          Online
  FPC 6          Online
  FPC 7          Online

lcc1-re0:
--------------------------------------------------------------------------
  Item           Status                   Reason
  FPC 0          Online
    PIC 3        Online
  FPC 1          Online
  FPC 2          Online
  FPC 4          Online
  FPC 6          Online
  FPC 7          Online

lcc2-re0:
--------------------------------------------------------------------------
  Item           Status                   Reason
  FPC 0          Online
  FPC 2          Online
  FPC 3          Online
    PIC 0        Online
  FPC 4          Online
  FPC 6          Online
  FPC 7          Online
    PIC 1        Online

lcc3-re0:
--------------------------------------------------------------------------
  Item           Status                   Reason
  FPC 0          Online
```

```
    PIC 0        Online
  FPC 1          Online
  FPC 2          Online
  FPC 3          Online
    PIC 2        Online
  FPC 4          Online
  FPC 5          Online
  FPC 6          Online
  FPC 7          Online
    PIC 1        Online
```

show chassis in-service-upgrade (QFX5100 Switch)

```
user@switch> show chassis in-service-upgrade
  Item           Status                   Reason
  FPC 0          Online (ISSU)
```

show chassis in-service-upgrade (EX9253 Switch)

```
user@switch> show chassis in-service-upgrade
 Item           Status                   Reason
  FPC 0          Online
  FPC 1          Online
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

request system software abort | 15212

request system software in-service-upgrade | 15262

Getting Started with Unified In-Service Software Upgrade

Example: Performing a Unified ISSU

show chassis jnu satellite

IN THIS SECTION

- Syntax | 18513
- Description | 18513
- Options | 18513
- Required Privilege Level | 18513
- Output Fields | 18513
- Sample Output | 18514
- Release Information | 18514

Syntax

```
show chassis jnu satellite [device-list value]
```

Description

Display the status, model, and Junos OS version of all the JNU Satellites that joined the JNU Controller in Junos Node Unifier (JNU) topology. You can also designate a specific device name to obtain its details.

Options

device-list *value* (optional) Specify a device value from the available list

Required Privilege Level

view

Output Fields

This section lists the output fields for the `show chassis jnu satellites` command.

Table 299: show chassis jnu satellites Output Fields

Field Name	Field Description
Satellite	Shows the name(s) of the satellite nodes
Alive	Shows the status of the satellite nodes
Model	Displays the satellite node model
Version	Display the Junos OS version

Sample Output

request jnu satellite sync

```
user@host> show chassis jnu satellites
```

```
Satellite            Alive      Model      Version
----------------------------------------------------
satellite1           up         vSRX       24.4R1
```

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

show chassis lcd

IN THIS SECTION

- show chassis lcd (EX Series) | 18515
- show chassis lcd (QFX Series) | 18515
- show chassis lcd (OCX Series) | 18515
- Description | 18516
- Options | 18516
- Required Privilege Level | 18517
- Output Fields | 18517
- Sample Output | 18519
- Release Information | 18534

show chassis lcd (EX Series)

```
show chassis lcd
<fpc-slot fpc-slot-number>
<menu <(all-members | local | member member-id)>>
```

show chassis lcd (QFX Series)

```
show chassis lcd
<fpc-slot fpc-slot-number>
<interconnect-device device-id>
<node-device device-id>
```

show chassis lcd (OCX Series)

```
show chassis lcd
<fpc-slot fpc-slot-number>
```

Description

Display the information that appears on the LCD panel of EX3200, EX3300, EX4200, EX4500, EX6200, and EX8200 switches, XRE200 External Routing Engines, QFX Series standalone switches, OCX Series switches, and Interconnect devices and Node devices within a QFabric system. Display the status of the currently selected port parameter of the Status LED for each network port on the device.

Options

none

Display the information that appears on the LCD panel (for any EX Series member switch in a Virtual Chassis or for XRE200 External Routing Engines, display the information for all Virtual Chassis members). Display the status of the currently selected port parameter of the Status LED for each network port.

fpc-slot <*fpc-slot-number*>

(Optional) Display the information as follows:

- (EX3200, EX3300, EX4200, and EX4500 switches, QFX Series, or OCX Series) Display the information that appears on the LCD panel for either an FPC slot with no *fpc-slot-number* value specified or for the FPC slot specified by `fpc-slot 0`. `fpc-slot` refers to the switch itself and `0` is the only valid value for *fpc-slot-number*. Output for these options is the same as for the `none` option.

  Also display the status of the currently selected port parameter of the Status LED for each network port.

- (EX Series Virtual Chassis member switches or XRE200 External Routing Engines) If no *fpc-slot-number* value is specified, display the information that appears on the LCD panel for all members of the Virtual Chassis. Output for this option is the same as for the `none` option. If the *fpc-slot-number* value is specified (it equals the *member-id* value), display the information for the specified member.

  Also display the status of the currently selected port parameter of the Status LED for each network port.

- (EX6200 or EX8200 switches)—Display the information that appears on the LCD panel for the line card in the line-card slot specified by the *fpc-slot-number* value.

  Also display the status of the currently selected port parameter of the Status LED for each network port.

interconnect-device *device-id*

(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Interconnect device.

menu (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel.

menu all-members (EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for all Virtual Chassis members.

menu local (EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the Virtual Chassis member from which you issued the command.

menu member ***member-id*** (EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the specified Virtual Chassis member.

node-device ***device-id*** (QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Node device.

Required Privilege Level

view

Output Fields

Table 300 on page 18517 lists the output fields for the `show chassis lcd` command. Output fields are listed in the approximate order in which they appear.

Table 300: show chassis lcd Output Fields

Field Name	Field Description
`member` *number* (XRE200 External Routing Engine)	Member ID of the device whose content is being displayed.

Table 300: show chassis lcd Output Fields *(Continued)*

Field Name	Field Description
`Front panel contents for slot` `Front panel contents` (EX6200, EX8200 switch, XRE200 External Routing Engine, and QFX Series)	FPC slot number of the switch whose content is being displayed. The number is always `0`, except for EX4200 switches in a Virtual Chassis, where it is the member ID value. On EX6200 switches, EX8200 switches, and XRE200 External Routing Engines, no slot number is displayed. On XRE200 External Routing Engines, this field appears under the `member` *`number`* field for each member device in the EX8200 Virtual Chassis.
`LCD screen`	The first line displays the hostname (for Virtual Chassis members, displays the member ID, the current role, and hostname; for EX8200 switches, displays `RE` and the hostname). The second line displays the currently selected port parameter of the Status LED and the alarms counter. The Status LED port parameters are: • `ADM`—Administrative • `SPD`—Speed • `DPX`—Duplex • `POE`—Power over Ethernet (EX3200 and EX4200 switches only)
`LEDs status`	Current state of the Alarms, System, and Master LEDs (chassis status LEDs).
`Interface`	Names of the interfaces on the switch.

Table 300: show chassis lcd Output Fields *(Continued)*

Field Name	Field Description
LED (ADM/SPD/DPX/POE)	State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are: **NOTE**: The XRE200 External Routing Engine always displays the NA parameter. The QFX Series products do not have any of the port parameters listed below. • ADM—Administrative • SPD—Speed • DPX—Duplex • NA—Not applicable. • POE—Power over Ethernet
fpc*x*	On standalone EX Series and QFX Series switches, always 0. On EX Series Virtual Chassis member switches, member ID of the Virtual Chassis member whose LCD menu is displayed.

Sample Output

show chassis lcd (Two-Member EX4200 Virtual Chassis)

```
user@switch> show chassis lcd
Front panel contents for slot: 0
--------------------------------
LCD screen:
    00:BK switch1
    LED:SPD ALARM 00
LEDs status:
    Alarms LED: Off
    System LED: Green
    Master LED: Off
Interface       LED(ADM/SPD/DPX/POE)
-------------------------------------
ge-0/0/0        Off
ge-0/0/1        Off
```

```
ge-0/0/2        Off
ge-0/0/3        Off
ge-0/0/4        Off
ge-0/0/5        Off
ge-0/0/6        Off
ge-0/0/7        Off
ge-0/0/8        Off
ge-0/0/9        Off
ge-0/0/10       Off
ge-0/0/11       Off
ge-0/0/12       Off
ge-0/0/13       Off
ge-0/0/14       Off
ge-0/0/15       Off
ge-0/0/16       Off
ge-0/0/17       Off
ge-0/0/18       Off
ge-0/0/19       Off
ge-0/0/20       Off
ge-0/0/21       Off
ge-0/0/22       Off
ge-0/0/23       Off
Front panel contents for slot: 1
---------------------------------
LCD screen:
    01:RE switch2
    LED:SPD ALARM 01
LEDs status:
    Alarms LED: Yellow
    System LED: Green
    Master LED: Green
Interface       LED(ADM/SPD/DPX/POE)
--------------------------------------
ge-1/0/0        Off
ge-1/0/1        Off
ge-1/0/2        Off
ge-1/0/3        Off
ge-1/0/4        Off
ge-1/0/5        Off
ge-1/0/6        Off
ge-1/0/7        Off
ge-1/0/8        Off
ge-1/0/9        Off
```

```
ge-1/0/10       Off
ge-1/0/11       Off
ge-1/0/12       Off
ge-1/0/13       Off
ge-1/0/14       Off
ge-1/0/15       Off
ge-1/0/16       Off
ge-1/0/17       Off
ge-1/0/18       Off
ge-1/0/19       Off
ge-1/0/20       Off
ge-1/0/21       Off
ge-1/0/22       Off
ge-1/0/23       Off
```

The output for the `show chassis lcd fpc-slot` command is the same as the output for the `show chassis lcd` command.

show chassis lcd fpc-slot 1 (EX4200 Virtual Chassis)

```
user@switch> show chassis lcd fpc-slot 1
Front panel contents for slot: 1
--------------------------------
LCD screen:
    01:RE switch2
    LED:SPD ALARM 01
LEDs status:
    Alarms LED: Yellow
    System LED: Green
    Master LED: Green
Interface       LED(ADM/SPD/DPX/POE)
-------------------------------------
ge-1/0/0        Off
ge-1/0/1        Off
ge-1/0/2        Off
ge-1/0/3        Off
ge-1/0/4        Off
ge-1/0/5        Off
ge-1/0/6        Off
ge-1/0/7        Off
ge-1/0/8        Off
```

```
ge-1/0/9        Off
ge-1/0/10       Off
ge-1/0/11       Off
ge-1/0/12       Off
ge-1/0/13       Off
ge-1/0/14       Off
ge-1/0/15       Off
ge-1/0/16       Off
ge-1/0/17       Off
ge-1/0/18       Off
ge-1/0/19       Off
ge-1/0/20       Off
ge-1/0/21       Off
ge-1/0/22       Off
ge-1/0/23       Off
```

show chassis lcd (EX8200 Switch)

```
user@switch> show chassis lcd
Front panel contents:
---------------------
LCD screen:
    RE st-8200-r
    LED:ADM ALARM 01
LEDs status:
    Alarms LED: Yellow
    System LED: Yellow
    Master LED: Green
Interface       LED(ADM/SPD/DPX)
-----------------------------------------
ge-0/0/0        Off
ge-0/0/1        Off
ge-0/0/2        Off
ge-0/0/3        Off
ge-0/0/4        Off
ge-0/0/5        Off
ge-0/0/6        Off
ge-0/0/7        Off
ge-0/0/8        Off
ge-0/0/9        Off
ge-0/0/10       Off
```

```
ge-0/0/11       Off
ge-0/0/12       Off
ge-0/0/13       Off
ge-0/0/14       Off
ge-0/0/15       Off
ge-0/0/16       Off
ge-0/0/17       Off
ge-0/0/18       Off
ge-0/0/19       Off
ge-0/0/20       Off
ge-0/0/21       Off
ge-0/0/22       Off
ge-0/0/23       Off
ge-0/0/24       Off
ge-0/0/25       Off
ge-0/0/26       Off
ge-0/0/27       Off
ge-0/0/28       Off
ge-0/0/29       Off
ge-0/0/30       Off
ge-0/0/31       Off
ge-0/0/32       Off
ge-0/0/33       Off
ge-0/0/34       Off
ge-0/0/35       Off
ge-0/0/36       Off
ge-0/0/37       Off
ge-0/0/38       Off
ge-0/0/39       Off
ge-0/0/40       Off
ge-0/0/41       Off
ge-0/0/42       Off
ge-0/0/43       Off
ge-0/0/44       Off
ge-0/0/45       Off
ge-0/0/46       Off
ge-0/0/47       Off
xe-2/0/0        Off
xe-2/0/1        Off
xe-2/0/2        Off
xe-2/0/3        Off
xe-2/0/4        Off
xe-2/0/5        Off
```

```
xe-2/0/6        Off
xe-2/0/7        Off
xe-3/0/0        Off
xe-3/0/1        Off
xe-3/0/2        Off
xe-3/0/3        Off
xe-3/0/4        Off
xe-3/0/5        Off
xe-3/0/6        Off
xe-3/0/7        Off
xe-5/0/0        Off
xe-5/0/1        Off
xe-5/0/2        Off
xe-5/0/3        Off
xe-5/0/4        Off
xe-5/0/5        Off
xe-5/0/6        On
xe-5/0/7        On
xe-7/0/5        Off
```

show chassis lcd fpc-slot 2 (EX8200 Switch)

```
show chassis lcd fpc-slot 2

Interface       LED(ADM/SPD/DPX)
-----------------------------------------
xe-2/0/0        Off
xe-2/0/1        Off
xe-2/0/2        Off
xe-2/0/3        Off
xe-2/0/4        Off
xe-2/0/5        Off
xe-2/0/6        Off
xe-2/0/7        Off
```

show chassis lcd menu (EX4200 Switch)

```
user@switch> show chassis lcd menu
fpc0:
--------------------------------------------------------------------------
```

```
status-menu
status-menu vcp-status
status-menu power-status
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu vc-uplink-config
maintenance-menu factory-default
```

On an EX4200 switch in a Virtual Chassis, the output for the `show chassis lcd menu all-members` command is the same as the output for the `show chassis lcd menu` command.

show chassis lcd menu (EX8200 Switch)

```
user@switch> show chassis lcd menu
status-menu
status-menu sf-status1-menu
status-menu sf-status2-menu
status-menu psu-status1-menu
status-menu psu-status2-menu
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu factory-default
```

show chassis lcd (QFX3500 Switches)

```
user@switch> show chassis lcd
Front panel contents for slot: 0
--------------------------------
LCD screen:
00:RE switch
ALARM 01
LEDs status:
```

```
Status/Beacon LED: Yellow Blinking
Interface STATUS LED ACTIVITY LED
-------------------------------------------
fte-0/1/0 Off Off
```

show chassis lcd (XRE200 External Routing Engine in EX8200 Virtual Chassis)

```
user@external-routing-engine> show chassis lcd
member0:
--------------------------------------------------------------------------
Front panel contents:
---------------------
LCD screen:
    RE ex8200-member0
    LED:ADM ALARM 04
LEDs status:
    Alarms LED: Red
    System LED: Yellow
    Master LED: Green

member1:
--------------------------------------------------------------------------

member8:
--------------------------------------------------------------------------
Front panel contents:
---------------------
LCD screen:
    BACKUP

member9:
--------------------------------------------------------------------------
Front panel contents:
---------------------
LCD screen:
    09:RE xre200-member9
    LED: NA ALARM 01
Interface       LED(ADM/SPD/DPX/POE)
-----------------------------------------
ge-0/0/0        On
ge-0/0/1        On
```

```
ge-0/0/2        On
ge-0/0/3        On
ge-0/0/4        Off
ge-0/0/5        Off
ge-0/0/6        Off
ge-0/0/7        Off
ge-0/0/8        Off
ge-0/0/9        Off
ge-0/0/10       On
ge-0/0/11       Off
ge-0/0/12       Off
ge-0/0/13       Off
ge-0/0/14       Off
ge-0/0/15       Off
ge-0/0/16       Off
ge-0/0/17       Off
ge-0/0/18       Off
ge-0/0/19       Off
ge-0/0/20       Off
ge-0/0/21       Off
ge-0/0/22       Off
ge-0/0/23       Off
ge-0/0/24       Off
ge-0/0/25       Off
ge-0/0/26       Off
ge-0/0/27       Off
ge-0/0/28       Off
ge-0/0/29       Off
ge-0/0/30       Off
ge-0/0/31       Off
ge-0/0/32       Off
ge-0/0/33       Off
ge-0/0/34       Off
ge-0/0/35       Off
ge-0/0/36       Off
ge-0/0/37       Off
ge-0/0/38       Off
ge-0/0/39       Off
ge-0/0/40       On
ge-0/0/41       On
ge-0/0/42       On
ge-0/0/43       On
ge-0/0/44       On
```

```
ge-0/0/45       On
ge-0/0/46       On
ge-0/0/47       On
ge-16/0/0       On
ge-16/0/1       Off
ge-16/0/2       On
ge-16/0/3       Off
ge-16/0/4       On
ge-16/0/5       Off
ge-16/0/6       On
ge-16/0/7       Off
ge-16/0/8       Off
ge-16/0/9       Off
ge-16/0/10      Off
ge-16/0/11      Off
ge-16/0/12      Off
ge-16/0/13      On
ge-16/0/14      Off
ge-16/0/15      On
ge-16/0/16      Off
ge-16/0/17      On
ge-16/0/18      On
ge-16/0/19      On
ge-16/0/20      On
ge-16/0/21      Off
ge-16/0/22      On
ge-16/0/23      Off
ge-16/0/24      Off
ge-16/0/25      Off
ge-16/0/26      On
ge-16/0/27      Off
ge-16/0/28      Off
ge-16/0/29      Off
ge-16/0/30      On
ge-16/0/31      Off
ge-16/0/32      On
ge-16/0/33      On
ge-16/0/34      On
ge-16/0/35      Off
ge-16/0/36      On
ge-16/0/37      Off
ge-16/0/38      Off
ge-16/0/39      Off
```

```
ge-16/0/40      Off
ge-16/0/41      Off
ge-16/0/42      On
ge-16/0/43      Off
ge-16/0/44      Off
ge-16/0/45      Off
ge-16/0/46      Off
ge-16/0/47      Off
xe-19/0/0       Off
xe-19/0/1       On
xe-19/0/2       On
xe-19/0/3       On
xe-19/0/4       On
xe-19/0/5       On
ge-22/0/0       Off
ge-22/0/1       Off
ge-22/0/2       On
ge-22/0/3       Off
ge-22/0/4       On
ge-22/0/5       On
ge-22/0/6       On
ge-22/0/7       On
ge-22/0/8       Off
ge-22/0/9       Off
ge-22/0/10      Off
ge-22/0/11      Off
ge-22/0/12      Off
ge-22/0/13      Off
ge-22/0/14      Off
ge-22/0/15      Off
ge-22/0/16      On
ge-22/0/17      Off
ge-22/0/18      On
ge-22/0/19      Off
ge-22/0/20      On
ge-22/0/21      Off
ge-22/0/22      On
ge-22/0/23      Off
ge-22/0/24      On
ge-22/0/25      Off
ge-22/0/26      Off
ge-22/0/27      Off
ge-22/0/28      Off
```

```
ge-22/0/29      Off
ge-22/0/30      Off
ge-22/0/31      Off
ge-22/0/32      On
ge-22/0/33      Off
ge-22/0/34      On
ge-22/0/35      Off
ge-22/0/36      Off
ge-22/0/37      Off
ge-22/0/38      Off
ge-22/0/39      Off
ge-22/0/40      Off
ge-22/0/41      Off
ge-22/0/42      Off
ge-22/0/43      Off
ge-22/0/44      Off
ge-22/0/45      Off
ge-22/0/46      Off
ge-22/0/47      Off
```

show chassis lcd interconnect-device (QFabric Systems)

```
show chassis lcd interconnect-device IC-F1012
                Front Panel Module Information
                ---------------------------------
                LCD screen:
                IC-F1012        3 Alarms active
LEDs status:
    Status LED: Green
    Power LED : Green
    Major Alarm LED: off
    Minor Alarm LED: Yellow
    Fan 0 LED : Green
    Fan 1 LED : Green
    Fan 2 LED : Green
    Fan 3 LED : Green
    Fan 4 LED : Green
    Fan 5 LED : Green
    Fan 6 LED : Green
    Fan 7 LED : Green
    Fan 8 LED : Green
```

```
    Fan 9 LED : Green
    PEM 0 LED : Green
    PEM 1 LED : Green
    PEM 2 LED : Green
    PEM 3 LED : off
    PEM 4 LED : off
    PEM 5 LED : off

                LED info for: CB - 0
                ---------------------------------
LEDs status:
    Status LED: Green
    Mastership LED: Green

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
IC-F1012:pme0 :         Green           N/A
IC-F1012:pme1 :         Green           N/A
IC-F1012:pme2 :         off             N/A
IC-F1012:pme3 :         off             N/A

                LED info for: CB - 1
                ---------------------------------
LEDs status:
    Status LED: Green
    Mastership LED: Amber

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
IC-F1012:pme0 :         Green           N/A
IC-F1012:pme1 :         Green           N/A
IC-F1012:pme2 :         off             N/A
IC-F1012:pme3 :         off             N/A

                LED info for: FC 0 FPC - 0
                ---------------------------------
LEDs status:
    Status LED: Green

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
IC-F1012:fte-0/0/0      Green           N/A
IC-F1012:fte-0/0/1      Green           N/A
```

```
IC-F1012:fte-0/0/2      Green           N/A
IC-F1012:fte-0/0/3      Green           N/A
IC-F1012:fte-0/0/4      Green           N/A

                LED info for: FC 1 FPC - 1
                ---------------------------------
LEDs status:
    Status LED: Green

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
IC-F1012:fte-1/0/0      Green           N/A
IC-F1012:fte-1/0/1      Green           N/A
IC-F1012:fte-1/0/2      Green           N/A
IC-F1012:fte-1/0/3      Green           N/A
IC-F1012:fte-1/0/4      Green           N/A

                LED info for: RC 0 FPC - 8
                ---------------------------------
LEDs status:
    Status LED: Green

                LED info for: RC 1 FPC - 9
                ---------------------------------
LEDs status:
    Status LED: Green

                LED info for: RC 2 FPC - 10
                ---------------------------------
LEDs status:
    Status LED: Green

                LED info for: RC 3 FPC - 11
                ---------------------------------
LEDs status:
    Status LED: Green

                LED info for: RC 4 FPC - 12
                ---------------------------------
LEDs status:
    Status LED: Green

                LED info for: RC 5 FPC - 13
```

```
                --------------------------------
LEDs status:
    Status LED: Green

                LED info for: RC 6 FPC - 14
                --------------------------------
LEDs status:
    Status LED: Green

                LED info for: RC 7 FPC - 15
                --------------------------------
LEDs status:
    Status LED: Green
```

show chassis lcd node-device (QFabric Systems)

```
show chassis lcd node-device P3774-C
                Front panel contents for: P3774-C
                ----------------------------------
                LCD screen:
                P3774-C
LEDs status:
    Status/Beacon LED: Yellow Blinking

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
P3774-C:xe-0/0/6        Green           Green
P3774-C:xe-0/0/7        Green           Green
P3774-C:ge-0/0/10       Green           Green
P3774-C:ge-0/0/11       Green           Green Blinking
P3774-C:ge-0/0/12       Green           Off
P3774-C:ge-0/0/13       Green           Green Blinking
P3774-C:ge-0/0/20       Green           Green
P3774-C:ge-0/0/21       Green           Green
P3774-C:ge-0/0/22       Green           Green Blinking
P3774-C:ge-0/0/23       Green           Off
P3774-C:ge-0/0/30       Green           Green
P3774-C:ge-0/0/31       Green           Green
P3774-C:ge-0/0/32       Green           Green Blinking
P3774-C:ge-0/0/33       Green           Green Blinking
P3774-C:fte-0/1/0       Green           Green
```

```
P3774-C:fte-0/1/1       Green           Green Blinking
P3774-C:fte-0/1/2       Green           Green Blinking
P3774-C:fte-0/1/3       Green           Green
```

Release Information

Command introduced in Junos OS Release 9.0.

`menu` option introduced in Junos OS Release 10.2 for EX Series switches.

RELATED DOCUMENTATION

- LCD Panel in EX3200 Switches
- *LCD Panel in EX4200 Switches*
- LCD Panel in EX4500 Switches
- LCD Panel in an EX8200 Switch
- LCD Panel in an XRE200 External Routing Engine
- *Configuring the LCD Panel on EX Series Switches (CLI Procedure)*
- set chassis display message | **16865**

show chassis led

IN THIS SECTION

- show chassis led (EX Series) | **18535**
- show chassis led (QFX Series) | **18535**
- Syntax (OCX Series) | **18535**
- Description | **18535**
- Options | **18535**
- Required Privilege Level | **18536**
- Output Fields | **18536**
- Sample Output | **18541**
- Release Information | **18550**

show chassis led (EX Series)

```
show chassis led
<fpc-slot <fpc-slot-number>>
```

show chassis led (QFX Series)

```
show chassis led
<fpc-slot <fpc-slot-number>>
interconnect-device name
node-device name
```

Syntax (OCX Series)

```
show chassis led
<fpc-slot <fpc-slot-number>>
```

Description

Display the status and colors of the chassis LEDs on the front panel of the switch. A major alarm (red) indicates a critical error condition that requires immediate action. A minor alarm (yellow) indicates a noncritical condition that requires monitoring or maintenance. A minor alarm that is left unchecked might cause interruption in service or performance degradation.

Options

none

Display the status of the chassis status LEDs (for EX4200 switches configured as a Virtual Chassis, display the information for all Virtual Chassis members; for EX4300, display the LED status only for ports in PIC 0 and PIC 1).

fpc-slot <*fpc-slot-number*>

(Optional) (Not on EX2200 switches) Display the information as follows:

- (EX3200, standalone EX4200, standalone QFX3500,EX4500, and OCX Series switches) Display the status of the chassis status LEDs for either an FPC slot with no ***fpc-slot-number*** value specified or for the FPC slot specified by **fpc-slot**

0. **fpc-slot** refers to the switch itself and **0** is the only valid value for ***fpc-slot-number***. Output for these options is the same as for the **none** option.

- (EX4200 switches in a Virtual Chassis with two or more members) If no ***fpc-slot-number*** value is specified, display the status of the chassis status LEDs for all members of the Virtual Chassis. Output for this option is the same as for the **none** option. If the ***fpc-slot-number*** value is specified (it equals the ***member-id*** value), display the status of the chassis status LEDs for the specified member.
- (EX8200 switches)—Display the status of the chassis status LEDs for the line card in the line-card slot specified by the ***fpc-slot-number*** value.

interconnect-device *name*

.

(QFabric systems only) (Optional) Display the status of the chassis and interface status LEDs for the Interconnect device.

node-device *name*

(QFabric systems only) (Optional) Display the status of the chassis and interface status LEDs for the Node device.

Required Privilege Level

view

Output Fields

Table 301 on page 18536 lists the output fields for the `show chassis led` command. Output fields are listed in the approximate order in which they appear.

Table 301: show chassis led Output Fields

Field Name	Field Description
Front panel contents for slot **Front panel contents** (EX8200 Switches) **Front Panel Module Information** (QFabric system QFX3008-I Interconnect device) **Front panel contents for** (QFabric system Node devices and QFX3600-I Interconnect devices)	FPC slot number of the device whose content is being displayed. The number is always **0**, except for EX4200 switches in a Virtual Chassis, where it is the member ID value. On EX8200 switches, no slot number is displayed. On QFabric system Node devices, the name of the Node device whose content is being displayed.

Table 301: show chassis led Output Fields *(Continued)*

Field Name	Field Description
Alarms LED	(EX Series switches only) Displays status of the ALM LED: • Off—No alarm has been configured. • Green—No alarm has been triggered. • Red—Major alarm. • Yellow—Minor alarm
System LED	(EX Series switches only) Displays status of the SYS LED: • Off—Switch is powered off. • Green—Switch is operating normally. • Yellow—Switch is booting.

Master LED:	Displays status of the MST LED (on EX3200, EX4200, and EX8200 switches):
	• Green—On an EX4200 Virtual Chassis switch, indicates the switch is the primary in the Virtual Chassis configuration. On other switches, indicates that the Routing Engine is operational.
	• Off
	  • On an EX4200 Virtual Chassis switch, indicates that this switch is not the primary in the Virtual Chassis configuration.
	  • On EX3200, standalone EX4200, and EX8200 switches, indicates that the Routing Engine is not operational.

Table 301: show chassis led Output Fields *(Continued)*

Field Name	Field Description
Mode LED:	(EX Series switches only) On an EX2200 switch only, displays the currently selected port parameter of the Status LED: • **ADM**—Administrative • **SPD**—Speed • **DPX**—Duplex • **POE**—Power over Ethernet
Status/Beacon LED	(QFX Series and OCX Series) Displays the system status as indicated by the Status LED on the chassis. For more information, see: • Chassis Status LEDs on a QFX3500 Device • Chassis Status LEDs in the QFX3600 and QFX3600-I Device
LINK/SPEED LED	(QFX Series and OCX Series) Displays the link status and speed of a management port. For more information, see: • Management Port LEDs on a QFX3500 Device • Management Port LEDs in the QFX3600 and QFX3600-I Device
ACTIVITY LED	(QFX Series and OCX Series) Displays the activity status of a management port. For more information, see: • Management Port LEDs on a QFX3500 Device • Management Port LEDs in the QFX3600 and QFX3600-I Device

Table 301: show chassis led Output Fields *(Continued)*

Field Name	Field Description
STATUS LED	(QFX Series and OCX Series) Displays the link status of an interface as indicated by the ST LED. For more information, see: • Control Board LEDs on a QFX3008-I Interconnect Device • Access Port and Uplink Port LEDs on a QFX3500 Device • Access Port and Uplink Port LEDs on a QFX3600 or QFX3600-I Device
LINK/ACTIVITY LED	(QFX Series and OCX Series) Displays link activity or faults on an interface as indicated by the LA LED. For more information, see: • Access Port and Uplink Port LEDs on a QFX3500 Device • Access Port and Uplink Port LEDs on a QFX3600 or QFX3600-I Device
Status LED	(QFX3008-I Interconnect device only) • Displays the system status as indicated by the STATUS LED on the front panel of the chassis. For more information, see Chassis Status LEDs on a QFX3008-I Interconnect Device. • Displays the status of a Control Board as indicated by the STATUS LED on the Control Board. For more information, see Control Board LEDs on a QFX3008-I Interconnect Device.
Power LED	(QFX3008-I Interconnect device only) Displays the status of system power on the device. For more information, see Chassis Status LEDs on a QFX3008-I Interconnect Device.
Major Alarm LED	(QFX3008-I Interconnect device only) Displays whether a critical error condition that requires immediate action exists on the device. For more information, see Chassis Status LEDs on a QFX3008-I Interconnect Device.

Table 301: show chassis led Output Fields *(Continued)*

Field Name	Field Description
Minor Alarm LED	(QFX3008-I Interconnect device only) Displays whether a noncritical condition that requires monitoring or maintenance exists on the device. For more information, see Chassis Status LEDs on a QFX3008-I Interconnect Device.
Fan 0 LED **Fan 1 LED** **Fan 2 LED** **Fan 3 LED** **Fan 4 LED** **Fan 5 LED** **Fan 6 LED** **Fan 7 LED** **Fan 8 LED**	(QFX3008-I Interconnect device only) Displays the status of fan trays on the device. For more information, see Chassis Status LEDs on a QFX3008-I Interconnect Device.
PEM 0 LED **PEM 1 LED** **PEM 2 LED** **PEM 3 LED** **PEM 4 LED**	(QFX3008-I Interconnect device only) Displays the status of power supplies on the device. For more information, see Chassis Status LEDs on a QFX3008-I Interconnect Device.
LED info for	(QFX3008-I Interconnect device only) Displays the LED information for a Control Board.
Mastership LED	(QFX3008-I Interconnect device only) Displays status of the MASTER LED on a Control Board. For more information, see Control Board LEDs on a QFX3008-I Interconnect Device.
Interface	Names of the interfaces on the device.

Table 301: show chassis led Output Fields *(Continued)*

Field Name	Field Description
LED (ADM/SPD/DPX/POE)	(EX Series switches only) State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are: **NOTE**: EX4500 and EX8200 switches do not have the POE port parameter. • **ADM**—Administrative • **SPD**—Speed • **DPX**—Duplex • **POE**—Power over Ethernet

Sample Output

show chassis led (EX2200 Switch)

```
user@switch> show chassis led

Front panel contents for slot: 0
--------------------------------
LEDs status:
    Alarms LED: Amber
    System LED: Green
    Mode LED  : Duplex
Interface    LED(ADM/SPD/DPX/POE)
-------------------------------------
ge-0/0/0        Off
ge-0/0/1        Full Duplex
ge-0/0/2        Full Duplex
ge-0/0/3        Off
ge-0/0/4        Off
ge-0/0/5        Full Duplex
ge-0/0/6        Full Duplex
ge-0/0/7        Full Duplex
ge-0/0/8        Full Duplex
ge-0/0/9        Full Duplex
```

```
ge-0/0/10       Full Duplex
ge-0/0/11       Full Duplex
ge-0/0/12       Full Duplex
ge-0/0/13       Full Duplex
ge-0/0/14       Full Duplex
ge-0/0/15       Full Duplex
ge-0/0/16       Full Duplex
ge-0/0/17       Full Duplex
ge-0/0/18       Full Duplex
ge-0/0/19       Full Duplex
ge-0/0/20       Full Duplex
ge-0/0/21       Full Duplex
ge-0/0/22       Off
ge-0/0/23       Off
ge-0/0/24       Full Duplex
ge-0/0/25       Full Duplex
ge-0/0/26       Off
ge-0/0/27       Off
ge-0/0/28       Full Duplex
ge-0/0/29       Full Duplex
```

show chassis led

```
user@switch> show chassis led

Front panel contents for slot: 0
--------------------------------
LEDs status:
    Alarms LED: Off
    System LED: Green
    Master LED: Green
Interface       LED(ADM/SPD/DPX/POE)
-------------------------------------
ge-0/0/0        Off
ge-0/0/1        Full Duplex
ge-0/0/2        Full Duplex
ge-0/0/3        Off
ge-0/0/4        Off
ge-0/0/5        Full Duplex
ge-0/0/6        Full Duplex
ge-0/0/7        Full Duplex
```

```
ge-0/0/8        Full Duplex
ge-0/0/9        Full Duplex
ge-0/0/10       Full Duplex
ge-0/0/11       Full Duplex
ge-0/0/12       Full Duplex
ge-0/0/13       Full Duplex
ge-0/0/14       Full Duplex
ge-0/0/15       Full Duplex
ge-0/0/16       Full Duplex
ge-0/0/17       Full Duplex
ge-0/0/18       Full Duplex
ge-0/0/19       Full Duplex
ge-0/0/20       Full Duplex
ge-0/0/21       Full Duplex
ge-0/0/22       Off
ge-0/0/23       Off
ge-0/0/24       Full Duplex
ge-0/0/25       Full Duplex
ge-0/0/26       Off
ge-0/0/27       Off
ge-0/0/28       Full Duplex
ge-0/0/29       Full Duplex
```

show chassis led fpc-slot 0

```
user@switch> show chassis led fpc-slot 0
Front panel contents for slot: 0
---------------------------------
LEDs status:
    Alarms LED: Red
    System LED: Green
    Master LED: Green
Interface       LED(ADM/SPD/DPX/POE)
-------------------------------------
ge-0/0/0        Off
ge-0/0/1        Off
ge-0/0/2        Off
ge-0/0/3        Off
ge-0/0/4        Off
ge-0/0/5        Off
ge-0/0/6        Off
```

```
ge-0/0/7        Off
ge-0/0/8        Off
ge-0/0/9        Off
ge-0/0/10       Off
ge-0/0/11       Off
ge-0/0/12       Off
ge-0/0/13       Off
ge-0/0/14       Off
ge-0/0/15       Off
ge-0/0/16       Off
ge-0/0/17       Off
ge-0/0/18       Off
ge-0/0/19       Off
ge-0/0/20       Off
ge-0/0/21       Off
ge-0/0/22       Off
ge-0/0/23       Off
```

show chassis led (EX Series)

```
user@switch> show chassis led
Front panel contents for slot: 0
--------------------------------
LEDs status:
Alarms LED: Amber
Status LED: Green
Mode LED : Duplex
Interface LED(ADM/SPD/DPX/POE)
-------------------------------------
ge-0/0/0 Off
ge-0/0/1 Full Duplex
ge-0/0/2 Full Duplex
ge-0/0/3 Off
ge-0/0/4 Off
ge-0/0/5 Full Duplex
ge-0/0/6 Full Duplex
ge-0/0/7 Full Duplex
ge-0/0/8 Full Duplex
ge-0/0/9 Full Duplex
ge-0/0/10 Full Duplex
ge-0/0/11 Full Duplex
```

```
ge-0/0/12 Full Duplex
ge-0/0/13 Full Duplex
ge-0/0/14 Full Duplex
ge-0/0/15 Full Duplex
ge-0/0/16 Full Duplex
ge-0/0/17 Full Duplex
ge-0/0/18 Full Duplex
ge-0/0/19 Full Duplex
ge-0/0/20 Full Duplex
ge-0/0/21 Full Duplex
ge-0/0/22 Off
ge-0/0/23 Off
ge-0/0/24 Full Duplex
ge-0/0/25 Full Duplex
ge-0/0/26 Off
ge-0/0/27 Off
ge-0/0/28 Full Duplex
ge-0/0/29 Full Duplex
```

show chassis led node-device (QFabric System Node Device)

```
user@switch> show chassis led node-device node1
 Front panel contents for: node1
LEDs status:
    Status/Beacon LED: Yellow Blinking

Interface              LINK/SPEED LED      ACTIVITY LED
-------------------------------------------------------
node1:me5              Green               N/A
node1:me6              Green               N/A

Interface              STATUS LED      LINK/ACTIVITY LED
-----------------------------------------------------
node1:xe-0/0/8         Green           Green
node1:ge-0/0/10        Green           Green
node1:ge-0/0/12        Green           Green
node1:ge-0/0/24        Green           Green
node1:ge-0/0/25        Green           Green
node1:ge-0/0/26        Green           Green
node1:ge-0/0/27        Green           Green
node1:ge-0/0/28        Green           Green
```

```
node1:ge-0/0/29        Green           Green
node1:ge-0/0/30        Green           Green
node1:ge-0/0/31        Green           Green
node1:ge-0/0/32        Green           Green
node1:ge-0/0/33        Green           Green
node1:ge-0/0/34        Green           Green
node1:ge-0/0/35        Green           Green
node1:ge-0/0/36        Green           Green
node1:ge-0/0/37        Green           Green
node1:ge-0/0/38        Green           Green
node1:ge-0/0/39        Green           Green
node1:fte-0/1/0        Green           Green Blinking
node1:fte-0/1/2        Green           Green Blinking
```

show chassis led interconnect-device (QFabric System - QFX3600-I Interconnect Device)

```
user@switch> show chassis led interconnect-device
IC2
                Front panel contents for: FPC 0
                ----------------------------------
LEDs status:
    Status/Beacon LED: Yellow Blinking

Interface               LINK/SPEED LED      ACTIVITY LED
-------------------------------------------------------
IC-EG0712:me5           Green               N/A
IC-EG0712:me6           Green               N/A

Interface               STATUS LED      LINK/ACTIVITY LED
---------------------------------------------------------
IC2:fte-0/1/0     Green           Green
IC2:fte-0/1/1     Green           Green Blinking
IC2:fte-0/1/2     Green           Green
IC2:fte-0/1/3     Green           Green Blinking
IC2:fte-0/1/4     Green           Green
IC2:fte-0/1/5     Green           Green Blinking
IC2:fte-0/1/6     Green           Green
IC2:fte-0/1/7     Green           Green
IC2:fte-0/1/8     Green           Green Blinking
```

```
IC2:fte-0/1/9     Green           Green Blinking
IC2:fte-0/1/10    Green           Green Blinking
```

show chassis led interconnect-device (QFabric System - QFX3008-I Interconnect Device)

```
user@switch> show chassis led interconnect-device
IC2
                Front Panel Module Information
                --------------------------------
LEDs status:
    Status LED: Green
    Power LED : Yellow Blinking
    Major Alarm LED: Red
    Minor Alarm LED: Yellow
    Fan 0 LED : Green
    Fan 1 LED : Green
    Fan 2 LED : Green
    Fan 3 LED : Green
    Fan 4 LED : Green
    Fan 5 LED : Green
    Fan 6 LED : Green
    Fan 7 LED : Green
    Fan 8 LED : Green
    Fan 9 LED : Green
    PEM 0 LED : Green
    PEM 1 LED : Green
    PEM 2 LED : Green
    PEM 3 LED : off
    PEM 4 LED : Yellow Blinking
    PEM 5 LED : off

                LED info for: CB - 0
                --------------------------------
LEDs status:
    Status LED: Green
    Mastership LED: Green

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
IC2:pme0 :         Green           N/A
IC2:pme1 :         off             N/A
```

```
IC2:pme2 :         off             N/A
IC2:pme3 :         off             N/A

                LED info for: CB - 1
                --------------------------------
LEDs status:
    Status LED: Green
    Mastership LED: Amber

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
IC2:pme0 :         Green           N/A
IC2:pme1 :         off             N/A
IC2:pme2 :         off             N/A
IC2:pme3 :         off             N/A

                LED info for: FC 0 FPC - 0
                --------------------------------
LEDs status:
    Status LED: Green

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
IC2:fte-0/0/0      Green           N/A
IC2:fte-0/0/1      Green           N/A
IC2:fte-0/0/2      Green           N/A
IC2:fte-0/0/3      Green           N/A
IC2:fte-0/0/4      Green           N/A
IC2:fte-0/0/5      Green           N/A
IC2:fte-0/0/6      Green           N/A
IC2:fte-0/0/7      Green           N/A
IC2:fte-0/0/8      Green           N/A
IC2:fte-0/0/9      Green           N/A
IC2:fte-0/0/10     Green           N/A
IC2:fte-0/0/11     Green           N/A
IC2:fte-0/0/12     Green           N/A
IC2:fte-0/0/13     Green           N/A
IC2:fte-0/0/14     Green           N/A
IC2:fte-0/0/15     Green           N/A

                LED info for: FC 1 FPC - 1
                --------------------------------
LEDs status:
```

```
    Status LED: Green

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
IC2:fte-1/0/0      Green           N/A
IC2:fte-1/0/1      Green           N/A

                LED info for: RC 2 FPC - 10
                --------------------------------
LEDs status:
    Status LED: Green

                LED info for: RC 3 FPC - 11
                --------------------------------
LEDs status:
    Status LED: Green
```

show chassis led (ACX Series)

```
user@host> show chassis led
LEDs status:
    Alarm LED : Red/Yellow Blink
    Beacon LED: On
    System LED: Green

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
et-0/0/0                N/A             Off
et-0/0/1                N/A             Off
et-0/0/2                N/A             Off
et-0/0/3                N/A             Off
et-0/0/4                N/A             Green
et-0/0/5                N/A             Green
et-0/0/6                N/A             Green
et-0/0/7                N/A             Green
et-0/0/8                N/A             Off
et-0/0/9                N/A             Off
et-0/0/10               N/A             Off
et-0/0/11               N/A             Off
et-0/0/12               N/A             Off
et-0/0/13               N/A             Off
```

```
et-0/0/14              N/A            Off
et-0/0/15              N/A            Off
et-0/0/16              N/A            Off
et-0/0/17              N/A            Off
et-0/0/18              N/A            Off
et-0/0/19              N/A            Off
et-0/0/20              N/A            Off
et-0/0/21              N/A            Off
et-0/0/22              N/A            Off
et-0/0/23              N/A            Off
et-0/0/24              N/A            Off
et-0/0/25              N/A            Green
et-0/0/26              N/A            Off
et-0/0/27              N/A            Off
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

- Chassis Status LEDs in EX2200 Switches
- Chassis Status LEDs in EX3200 Switches
- *Chassis Status LEDs in EX4200 Switches*
- Chassis Status LEDs in EX4500 Switches
- Chassis Status LEDs in an EX8200 Switch
- Chassis Status LEDs on a QFX3500 Device
- Chassis Status LEDs in the QFX3600 and QFX3600-I Device
- Management Port LEDs on a QFX3500 Device
- Management Port LEDs in the QFX3600 and QFX3600-I Device
- Chassis Status LEDs on a QFX3008-I Interconnect Device
- Control Board LEDs on a QFX3008-I Interconnect Device

show chassis led satellite

IN THIS SECTION

- Syntax | 18551
- Description | 18551
- Options | 18551
- Required Privilege Level | 18552
- Output Fields | 18552
- Sample Output | 18555
- Release Information | 18556

Syntax

```
show chassis led satellite [slot-id slot-id |device-alias  alias-name]
```

Description

Display the status and colors of the chassis LEDs of the satellite devices in a Junos Fusion. A major alarm (red) indicates a critical error condition that requires immediate action. A minor alarm (yellow) indicates a noncritical condition that requires monitoring or maintenance. A minor alarm that is left unchecked might cause interruption in service or performance degradation.

Options

none	Display the status of the chassis status LEDs of every satellite device in the Junos Fusion.
slot-id *slot-id*	(Optional) Display the status of the chassis status LEDs of the satellite device using the specified FPC slot identifier in the Junos Fusion. The *slot-id* is the FPC slot ID number.
device-alias *alias-name*	(Optional) Display the status of the chassis status LEDs of the satellite device using the specified alias in the Junos Fusion.

Required Privilege Level

view

Output Fields

Table 302 on page 18552 lists the output fields for the `show chassis led satellite` command. Output fields are listed in the approximate order in which they appear.

Table 302: show chassis led Output Fields

Field Name	Field Description
`Beacon LED`	(Applies when QFX5100, QFX5110, and QFX5200 switches are in an satellite device role only) Indicates if the beacon feature is on or off. The beacon feature is always off in a Junos Fusion. The `Beacon LED` output maps to the `ID-Identification` LED state. For more information, see: • Chassis Status LEDs on a QFX5100 Device • QFX5110 Chassis Status LEDs • QFX5200 Chassis Status LEDs
`System LED`	Indicates the state of the System (`SYS`) LED on the satellite device. For more information, see: • *Chassis Status LEDs on EX4300 Switches* • Chassis Status LEDs on a QFX5100 Device • QFX5110 Chassis Status LEDs • QFX5200 Chassis Status LEDs

Table 302: show chassis led Output Fields *(Continued)*

Field Name	Field Description
`Master LED`	Indicates the state of the Primary (`MST`) LED on the satellite device. For more information, see: • *Chassis Status LEDs on EX4300 Switches* • Chassis Status LEDs on a QFX5100 Device • QFX5110 Chassis Status LEDs • QFX5200 Chassis Status LEDs
`Alarm LED`	Indicates the state of the Alarm (`ALM`) LED on the satellite device. For more information, see: • *Chassis Status LEDs on EX4300 Switches* • Chassis Status LEDs on a QFX5100 Device • QFX5110 Chassis Status LEDs • QFX5200 Chassis Status LEDs
`Mgmt Port0 LED`	(Applies when QFX5100, QFX5110, and QFX5200 switches are in an satellite device role only) Indicates the state of the management port 0 (`em0`) LED status on the satellite device. This port is always off in a Junos Fusion. For more information, see: • Management Port LEDs on a QFX5100 Device • QFX5110 Management Port LEDs • QFX5200 Management Port LEDs

Table 302: show chassis led Output Fields *(Continued)*

Field Name	Field Description
`Mgmt Port1 LED`	(Applies when QFX5100, QFX5110, and QFX5200 switches are in an satellite device role only) Indicates the state of the management port 1(`em0`) LED status on the satellite device. This port is always off in a Junos Fusion. For more information, see: • Management Port LEDs on a QFX5100 Device • QFX5110 Management Port LEDs • QFX5200 Management Port LEDs
`Interface`	The interface name on the satellite device. For more information, see: • Management Port LEDs on a QFX5100 Device • QFX5110 Management Port LEDs • QFX5200 Management Port LEDs
`Status LED`	The state of the Status LED for the particular interface on the satellite device. For more information, see: • *Network Port, Built-In QSFP+ Port, Uplink Port, and Uplink Module Port LEDs on EX4300 Switches* • Access Port and Uplink Port LEDs on a QFX5100 Device • QFX5110 Network Port LEDs • QFX5200 Access Port and Uplink Port LEDs

Table 302: show chassis led Output Fields *(Continued)*

Field Name	Field Description
`Link/Activity LED`	The state of the Link/Activity LED for the particular interface on the satellite device. For more information, see: • *Network Port, Built-In QSFP+ Port, Uplink Port, and Uplink Module Port LEDs on EX4300 Switches* • Access Port and Uplink Port LEDs on a QFX5100 Device • QFX5110 Network Port LEDs • QFX5200 Access Port and Uplink Port LEDs

Sample Output

show chassis led satellite

```
user@aggregation-device> show chassis led satellite

                LED status for: FPC 101
                -----------------------------------
LEDs status:
    Beacon LED: OFF
    System LED: GREEN
    Master LED: OFF
    Alarm LED : YELLOW
    Mgmt Port0 LED: OFF
    Mgmt Port1 LED: OFF

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
xe-101/0/0              green
xe-101/0/1              green
xe-101/0/10             off
xe-101/0/48:0           green
xe-101/0/48:1           green
xe-101/0/48:2           green
```

```
xe-101/0/48:3           green

                LED status for: FPC 102
                -----------------------------------
LEDs status:
    Beacon LED: OFF
    System LED: GREEN
    Master LED: OFF
    Alarm LED : YELLOW
    Mgmt Port0 LED: OFF
    Mgmt Port1 LED: OFF

Interface               STATUS LED      LINK/ACTIVITY LED
-------------------------------------------------------
xe-102/0/0              green
xe-102/0/1              green
xe-102/0/10             off
xe-102/0/48:0           green
xe-102/0/48:1           green
xe-102/0/48:2           green
xe-102/0/48:3           green
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Satellite Device Alarm Handling Using an Environment Monitoring Satellite Policy in a Junos Fusion

Understanding Junos Fusion Provider Edge Components

Understanding Junos Fusion Enterprise Components

show chassis location

IN THIS SECTION

- Syntax | 18557
- Syntax (TX Matrix Router) | 18557
- Syntax (TX Matrix Plus Router) | 18557
- Syntax (MX Series Router) | 18558
- Syntax (QFX Series) | 18558
- Description | 18558
- Options | 18558
- Required Privilege Level | 18560
- Output Fields | 18560
- Sample Output | 18561
- Release Information | 18562

Syntax

```
show chassis location
```

Syntax (TX Matrix Router)

```
show chassis location
<fpc | interface (by-name name | by-slot fpc number  lcc number) | lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
show chassis location
<fpc | interface (by-name name | by-slot fpc number  lcc number) | lcc number | sfc number>
```

Syntax (MX Series Router)

```
show chassis location
<all-members>
<local>
<member member-id>
```

Syntax (QFX Series)

```
show chassis location
<interconnect-device name>
<node-device name>
```

Description

Display the physical location of the chassis. This command can only be used on the primary Routing Engine.

Options

none	Display all information about the physical location of the chassis. On a TX Matrix router, display all information about the physical location of the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display all information about the physical location of the TX Matrix Plus router and its attached routers.
all-members	(MX Series routers only) (Optional) Display the physical location of the chassis for all the member routers in the Virtual Chassis configuration.
fpc	(TX Matrix router and TX Matrix Plus router only) (Optional) Display the physical location of all Flexible PIC Concentrators (FPCs).
interconnect-device *name*	(QFabric systems only) (Optional) Display the physical location of the Interconnect device.
interface by-name *name*	(TX Matrix and TX Matrix Plus routers only) (Optional) Display the physical location of a specified interface name. On a TX Matrix router, this option displays the FPC number and T640 router (line-card chassis) number associated with the specified interface. On a TX Matrix Plus router, this option displays the FPC number and router (line-card chassis) number associated with the specified interface.

interface by-slot fpc *number* lcc *number*

(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the global FPC number of an interface by specifying its local FPC number and T640 router (line-card chassis) number. On a TX Matrix Plus router, display the global FPC number of an interface by specifying its local FPC number and router (line-card chassis) number.

- The global FPC number is the FPC slot number when all the FPC slots in the routing matrix are considered: **0** through **31**. On TX Matrix Plus router with 3D SIBs, the value is `0` through `63`. The local FPC number is the FPC slot number on a particular T640 router.
- For **fpc**, replace ***number*** with a value from **0** through **7**.
- For **lcc**, replace ***number*** with a value from **0** through **7**.

lcc *number*

(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the physical location of a specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router , display the physical location of a specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(MX Series routers only) (Optional) Display the physical location of the chassis for the local Virtual Chassis member.

member *member-id*

(MX Series routers only) (Optional) Display the physical location of the chassis for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

node-device *name*

(QFabric systems only) (Optional) Display the physical location of the Node device.

scc (TX Matrix routers only) (Optional) Display the physical location of the TX Matrix router (switch-card chassis).

sfc (TX Matrix Plus routers only) (Optional) Display the physical location of the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level

view

Output Fields

Table 303 on page 18560 lists the output fields for the `show chassis location` command. Output fields are listed in the approximate order in which they appear.

Table 303: show chassis location Output Fields

Field Name	Field Description
country-code	Country code information.
postal-code	Postal code information.
Building	Building information.
Floor	Floor information.
Global FPC	Global FPC number. The FPC slot number, when all FPC slots in the routing matrix are considered. The range of values is **0** through **31**. On TX Matrix Plus router with 3D SIBs the value is `0` through `63`.
`LATA`	Local access transport area information.
LCC	Line-card chassis number. On a TX Matrix router, the number of a particular T640 router connected to the TX Matrix router. On a TX Matrix Plus router, the number of a particular router connected to the TX Matrix Plus router.

Table 303: show chassis location Output Fields *(Continued)*

Field Name	Field Description
Local FPC	Local FPC number. On a TX Matrix router, the FPC slot number on a particular T640 router. On a TX Matrix Plus router, the FPC slot number on a particular router.

Sample Output

show chassis location

```
user@host> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

show chassis location fpc (TX Matrix Router)

```
user@host> show chassis location fpc
Global FPC    LCC     Local FPC
    17          2         1
    21          2         5
```

show chassis location interface by-slot (TX Matrix Router)

```
user@host> show chassis location interface by-slot
fpc 1 lcc 1
Global FPC: 9
```

show chassis location fpc (TX Matrix Plus Router)

```
user@host> show chassis location fpc
Global FPC    LCC     Local FPC
```

```
    0          0         0
    1          0         1
```

show chassis location interface by-slot (TX Matrix Plus Router)

```
user@host> show chassis location interface by-slot
fpc 2 lcc 1
Global FPC: 10
```

show chassis location (QFX Series and OCX Series)

```
user@switch> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

show chassis location (QFabric Systems)

```
user@switch> show chassis location interconnect-device
interconnect1
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Displaying Chassis Physical Locations for a Routing Matrix with a TX Matrix Plus Router

show chassis mac-addresses

IN THIS SECTION

- Syntax | 18563
- Syntax (TX Matrix Router) | 18563
- Syntax (TX Matrix Plus Router) | 18563
- Syntax (MX Series Router) | 18564
- Syntax (QFX Series) | 18564
- Description | 18564
- Options | 18564
- Required Privilege Level | 18565
- Output Fields | 18565
- Sample Output | 18567
- Release Information | 18571

Syntax

```
show chassis mac-addresses
```

Syntax (TX Matrix Router)

```
show chassis mac-addresses
<lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
show chassis mac-addresses
<lcc number | sfc number>
```

Syntax (MX Series Router)

```
show chassis mac-addresses
<all-members>
<local>
<member member-id>
```

Syntax (QFX Series)

```
show chassis mac-addresses
<interconnect-device name>
<node-group name>
```

Description

Display the media access control (MAC) addresses for the router, switch chassis, or switch.

Options

none

(TX Matrix, TX Matrix Plus routers, QFX Series, and OCX Series Switches) Display the MAC addresses for the router chassis or switch. On a TX Matrix router, display MAC addresses on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display MAC addresses on the TX Matrix Plus router and its attached routers.

all-members

(MX Series routers only) (Optional) Display the MAC addresses for all the member routers of the Virtual Chassis configuration.

interconnect-device *name*

(QFabric switches only) (Optional) Display the MAC addresses for the Interconnect device.

lcc *number*

(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display MAC addresses for a specified T640 router (or line-card chassis) that is connected to the TX Matrix Plus router. On a TX Matrix Plus router, display MAC addresses for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Display the MAC addresses for the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display the MAC addresses for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-group *name* (QFabric switches only) (Optional) Display the MAC addresses for the specified Node group.

scc (TX Matrix routers only) (Optional) Display MAC addresses for the TX Matrix router (or switch-card chassis).

sfc *number* (TX Matrix Plus routers only) (Optional) Display MAC addresses for the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level

view

Output Fields

Table 304 on page 18565 lists the output fields for the `show chassis mac-addresses` command. Output fields are listed in the approximate order in which they appear.

Table 304: show chassis mac-addresses Output Fields

Field Name	Field Description
MAC address information	

Table 304: show chassis mac-addresses Output Fields *(Continued)*

Field Name	Field Description
Public base address	Base address of the MAC addresses allocated to this router or switch, for example 00:24:dc:18:09:40. This address is also printed on the box of the device. Public MAC addresses are those addresses your switch/router assigns to Ethernet interfaces. The public base address is the first MAC address your device assigns to an interface. Consecutive public count MAC addresses are reserved for additional interfaces.
Public count	Number of allocated public addresses, for example 64. Public addresses are calculated starting with the public base address. If the public base address is 00:24:dc:18:09:40, then The MAC address of ge-0/0/0 on this device is 00:24:dc:18:09:40, and ge-0/0/1 is 00:24:dc:18:09:41, and so on, up to 64 available addresses.
Private base address	Base address of the private MAC addresses allocated to this router or switch. The private base address is the first MAC address after the allocated consecutive public count addresses. For example, if 00:24:dc:18:09:40 is the public base and 0x40 is the public count in hex, then 00:24:dc:18:09:80 would be the private base.
Private count	Number of allocated private addresses.

NOTE: Starting in Junos OS Release 21.4R3, the SRX380 device increments the MAC address of all the revenue ports (for example, ge-0/0/0) by one upon upgrade. This resolves the MAC address duplication issue between fxp0 and revenue ports, aligning with the SRX340/345 Series MAC address pool design. If external applications or software rely on the old public base address or revenue ports MACs, they need to extract and refresh the MAC address from the ports after the upgrade. If the public base address is used for device identification, subtract one from the previously recorded public base address to access the same device.

Sample Output

show chassis mac-addresses

```
user@host> show chassis mac-addresses
MAC address information
  Public base address   0:90:69:0:4:0
  Public count       1008
  Private base address   0:90:69:0:7:f0
  Private count       16
```

show chassis mac-addresses (MX2010 Router)

```
user@host> show chassis mac-addresses
MAC address information:
  Public base address     64:87:88:04:50:00
  Public count            1984
  Private base address    64:87:88:04:57:c0
  Private count           64
```

show chassis mac-addresses (PTX10008 Router)

```
user@host> show chassis mac-addresses
MAC address information:
  Public base address      30:b6:4f:0a:7a:bb
  Public count             1856
  Private base address     30:b6:4f:0a:81:fb
  Private count            192
```

show chassis mac-addresses (TX Matrix Router)

```
user@host> show chassis mac-addresses
scc-re0:
--------------------------------------------------------------------------
MAC address information:
  Public base address     00:05:85:9e:cc:00
  Public count            8064
```

```
  Private base address    00:05:85:9e:eb:80
  Private count           128
lcc0-re0:
--------------------------------------------------------------------------
MAC address information:
  Public base address     00:05:85:68:98:00
  Public count            2032
  Private base address    00:05:85:68:9f:f0
  Private count           16
lcc2-re0:
--------------------------------------------------------------------------
MAC address information:
  Public base address     00:05:85:68:78:00
  Public count            2032
  Private base address    00:05:85:68:7f:f0
  Private count           16
```

show chassis mac-addresses (TX Matrix Plus Router)

```
user@host> show chassis mac-addresses
sfc0-re0:
--------------------------------------------------------------------------
MAC address information:
  Public base address     00:1d:b5:14:00:00
  Public count            65023
  Private base address    00:1d:b5:14:fd:ff
  Private count           512

lcc0-re0:
--------------------------------------------------------------------------
MAC address information:
  Public base address     00:1f:12:7a:84:00
  Public count            2032
  Private base address    00:1f:12:7a:8b:f0
  Private count           16

lcc1-re0:
--------------------------------------------------------------------------
MAC address information:
  Public base address     00:22:83:42:48:00
  Public count            2032
```

```
  Private base address    00:22:83:42:4f:f0
  Private count           16

lcc2-re0:
--------------------------------------------------------------------------
MAC address information:
  Public base address     00:1f:12:c3:58:00
  Public count            2032
  Private base address    00:1f:12:c3:5f:f0
  Private count           16

lcc3-re0:
--------------------------------------------------------------------------
MAC address information:
  Public base address     00:21:59:ef:b8:00
  Public count            2032
  Private base address    00:21:59:ef:bf:f0
  Private count           16
```

show chassis mac-addresses (QFX Series and OCX Series)

```
user@switch> show chassis mac-addresses
MAC address information:
Public base address 02:00:08:00:00:00
Public count 512
Private base address 02:00:00:00:00:00
Private count 64
```

show chassis mac-addresses interconnect-device (QFabric Switches)

```
user@switch> show chassis mac-addresses interconnect-device
interconnect1
MAC address information:
  Public base address     00:1f:12:30:9c:c0
  Public count            58
  Private base address    00:1f:12:30:9c:fa
  Private count           6
```

show chassis mac-addresses node-group (QFabric Switches)

```
user@switch> show chassis mac-addresses node-group
NW-NG-0
MAC address information:
------------------------
RE:
  FC MAC base    00:11:00:00:00:00
  FC MAC count   2
  VLAN MAC       00:11:00:00:00:09
EC6007
  Base address   00:00:01:76:00:00
  Count          64
EC6008
  Base address   00:22:83:22:52:ae
  Count          260
```

show chassis mac-addresses (ACX5048 and ACX5096 Routers)

```
user@host> show chassis mac-addresses
FPC 0
    Base address   64:64:9b:5e:0a:00
    Count          1280
```

show chassis mac-addresses (ACX500 and ACX Series)

```
user@host> show chassis mac-addresses
MAC address information:
  Public base address     f0:1c:2d:1b:60:80
  Public count            112
  Private base address    f0:1c:2d:1b:60:f0
  Private count           16
```

show chassis mac-addresses (EX9251 Switche)

```
user@switch> show chassis mac-addresses
 MAC address information:
```

```
  Public base address      4c:16:fc:90:68:00
  Public count             2032
  Private base address     4c:16:fc:90:6f:f0
  Private count            16
```

show chassis mac-addresses (EX9253 Switche)

```
user@switch> show chassis mac-addresses
 MAC address information:
  Public base address      38:4f:49:8f:00:b8
  Public count             2330
  Private base address     38:4f:49:8f:09:d2
  Private count            1766
```

Release Information

Command introduced before JUNOS Release 7.4.

sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6.

RELATED DOCUMENTATION

ACX2000 and ACX2100 Routers Hardware and CLI Terminology Mapping

show chassis network-services

IN THIS SECTION

- Syntax | 18572
- Description | 18572
- Options | 18572
- Required Privilege Level | 18572
- Output Fields | 18573
- Sample Output | 18573

Release Information | 18575

Syntax

```
show chassis network-services
```

Description

Display the network services mode that the router is configured to run in—IP Network Services mode, Ethernet Network Services mode, Enhanced IP Network Services mode, Enhanced Ethernet Network Services mode, or Enhanced mode.

Options

all-ethernet	All Ethernet network services
all-ip	All IP network services
enhanced-ethernet	Enhanced ethernet network services
enhanced-ip	Enhanced IP network services
enhanced-mode	Enhanced network services
ethernet	Ethernet network services
ip	IP network services
lan	Ethernet LAN services

Required Privilege Level

view

Output Fields

Table 305 on page 18573 lists the output fields for the `show chassis network services` command. Output fields are listed in the approximate order in which they appear.

Table 305: show chassis network services Output Fields

Field Name	Field Description
Network Services Mode	Network services mode configured for the router: • **IP**—IP Network Services mode. • **Ethernet**—Ethernet Network Services mode. • **enhanced-ip**—Enhanced IP Network Services mode • **enhanced-ethernet**—Enhanced Ethernet Network Services mode • `Enhanced-Mode`—Enhanced mode for PTX Series routers that have third-generation FPCs installed. See "enhanced-mode" on page 6398.

Sample Output

show chassis network services

```
user@host> show chassis network services
Network Services Mode: IP
```

show chassis network services (MX104 Router)

```
user@host> show chassis network services
Network Services Mode: IP
```

show chassis network services (MX2010 Router)

```
user@host> show chassis network services
Network Services Mode: Enhanced-IP
```

show chassis network services (MX2020 Router)

```
user@host> show chassis network services
Network Services Mode: Enhanced-IP
```

show chassis network services (MX2008 Router)

```
user@host> show chassis network-services
Network Services Mode: Enhanced-IP
```

show chassis network-services (MX10003 Router)

```
user@host> show chassis network-services

Network Services Mode: Enhanced-IP
```

show chassis network-services (MX204 Router)

```
user@host> show chassis network-services

Network Services Mode: Enhanced-IP
```

show chassis network-services (MX10008 Router)

```
user@host> show chassis network-services
Network Services Mode: Enhanced-IP
```

show chassis network services (PTX Router with third-generation FPCs)

```
user@host> show chassis network services
Network Services Mode: Enhanced-Mode
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

enhanced-mode (Network Services) | 6398

show chassis network-slices

IN THIS SECTION

- Syntax | 18575
- Description | 18575
- Required Privilege Level | 18575
- Output Fields | 18576
- Sample Output | 18577
- Release Information | 18579

Syntax

```
show chassis network-slices
<gnf gnf-id>
```

Description

Display Junos node slicing information for the guest network functions (GNFs) configured on the base system (BSYS). The `gnf` *`gnf-id`* option displays the information about a particular GNF.

Required Privilege Level

view

Output Fields

Table 306 on page 18576 lists the output fields for the `show chassis network-slices` command. Output fields are listed in the approximate order in which they appear.

Table 306: show chassis network-slices Output Fields

Field Name	Field Description
`GNF`	GNF identifier for each partition.
`Description`	Description of the guest network function.
State	Status of the GNF. • `Online`—GNF online and running. • `Offline`—GNF is offline.
Uptime	Duration for which the GNFs have been up and running. In case of a GNF Routing Engine switchover, the uptime value is reset as the new primary GNF Routing Engine establishes a fresh connection with the BSYS. This behavior is specific to GNFs and does not have any functional impact.
GNF ID	Shows the GNF ID.
GNF description	Shows the description of GNF.
FPCs assigned	Shows the FPC slot numbers assigned to the GNF.
FPCs online	Shows the FPCs that are online.
BSYS	Shows the name of the BSYS.
BSYS sw version	Shows the Junos software version used in the BSYS.
GNF sw version	Shows the Junos software version used in the GNF.

Table 306: show chassis network-slices Output Fields *(Continued)*

Field Name	Field Description
BSYS primary RE	Primary Routing Engine slot.
GNF uptime	Duration for which the GNF has been up and running. In case of a GNF Routing Engine switchover, the GNF uptime value is reset as the new primary GNF Routing Engine establishes a fresh connection with the BSYS. This behavior is specific to GNFs and does not have any functional impact.
GNF Routing Engine Status:	Show the details of the Routing Engines in the slot 0 and 1. It includes the Current State, (primary or backup), Routing Engine model, and GNF host name.

Sample Output

show chassis network-slices

```
user@router> show chassis network-slices
guest-network-functions:
GNF    Description    State          Uptime
 1     gnf-a          Online         12 hours, 46 minutes, 11 seconds
 2     gnf-b          Online         12 hours, 13 minutes, 57 seconds
 3     gnf-c          Online         12 hours, 3 minutes, 55 seconds
 4     gnf-d          Online         12 hours, 8 minutes, 20 seconds
 5     gnf-e          Online         12 hours, 2 minutes, 46 seconds
 6     gnf-f          Online         11 hours, 56 minutes, 29 seconds
```

show chassis network-slices gnf

```
user@router> show chassis network-slices gnf

GNF ID           1
GNF description  NA
GNF state        Online
FPCs assigned    7 8
FPCs online      7 8
```

```
BSYS             router(mx960)
BSYS sw version  18.2-20180321_0948_bsys
GNF  sw version  18.2-20180314_gnf
Chassis          mx960
BSYS master RE   0
GNF uptime       4 days, 23 hours, 55 minutes, 1 second
GNF Routing Engine Status:
Slot 0:
    Current state   Master
    Model           RE-GNF-2400x4
    GNF host name   gnf-host0
Slot 1:
    Current state   Backup
    Model           RE-GNF-2400x4
    GNF host name   gnf-host1
GNF ID           2
GNF description  NA
GNF state        Online
FPCs assigned    NA
FPCs online      NA
BSYS             router(mx960)
BSYS sw version  18.2-20180321_0948_bsys
GNF  sw version  18.2-20180216_gnf
Chassis          mx960
BSYS master RE   0
GNF uptime       4 days, 23 hours, 53 minutes, 54 seconds
GNF Routing Engine Status:
Slot 0:
    Current state   Master
    Model           RE-GNF-2400x4
    GNF host name   gnf-host2
Slot 1:
    Current state   Backup
    Model           RE-GNF-2400x4
    GNF host name   gnf-host3
```

show chassis network-slices gnf <gnf-id>

```
user@router> show chassis network-slices gnf 2


GNF ID           2
```

```
GNF description  NA
GNF state        Online
FPCs assigned    NA
FPCs online      NA
BSYS             router(mx960)
BSYS sw version  18.2-20180321_0948_bsys
GNF  sw version  18.2-20180216_gnf
Chassis          mx960
BSYS master RE   0
GNF uptime       4 days, 23 hours, 53 minutes, 54 seconds
GNF Routing Engine Status:
Slot 0:
    Current state   Master
    Model           RE-GNF-2400x4
    GNF host name   gnf-host2
Slot 1:
    Current state   Backup
    Model           RE-GNF-2400x4
    GNF host name   gnf-host3
```

Release Information

Command introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

- show chassis fpc pic-status (Node Slicing) | 18413
- show chassis fpc (Node Slicing) | 18408
- show chassis adc (Node Slicing) | 17297
- show chassis network-slices fpcs | 18579

show chassis network-slices fpcs

IN THIS SECTION

- Syntax | 18580

- Description | 18580
- Required Privilege Level | 18580
- Output Fields | 18580
- Sample Output | 18581
- Release Information | 18581

Syntax

```
show chassis network-slices fpcs
```

Description

Display information about the FPCs associated with different guest network functions (GNFs).

Required Privilege Level

view

Output Fields

Table 307 on page 18580 lists the output fields for the `show chassis network-slices fpcs` command. Output fields are listed in the approximate order in which they appear.

Table 307: show chassis network-slices fpcs Output Fields

Field Name	Field Description
GNF	GNF ID.
FPCs	Slot numbers of FPCs associated with the GNF ID.

Sample Output

command-name

```
user@router> show chassis network-slices fpcs

guest-network-functions:
GNF     FPCs
 1      6
 2      1
 3      0 2
 4      5
 5      7 8 9
 6      3 4
```

Release Information

Command introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

- show chassis network-slices | 18575
- show chassis fpc pic-status (Node Slicing) | 18413
- show chassis fpc (Node Slicing) | 18408
- show chassis adc (Node Slicing) | 17297

show chassis nonstop-upgrade

IN THIS SECTION

- Syntax | 18582
- Description | 18582
- Required Privilege Level | 18582
- Output Fields | 18582

Sample Output | 18583
Release Information | 18584

Syntax

```
show chassis nonstop-upgrade
```

Description

(EX6200 switches, EX8200 switches, EX8200 Virtual Chassis, QFX3500 and QFX3600 Virtual Chassis, and Virtual Chassis Fabric only) Display the status of the line cards or Virtual Chassis members in the linecard role after the most recent nonstop software upgrade (NSSU). You must issue this command on the primary Routing Engine.

Required Privilege Level

view

Output Fields

Table 308 on page 18582 lists the output fields for the `show chassis nonstop-upgrade` command. Output fields are listed in the approximate order in which they appear.

Table 308: show chassis nonstop-upgrade Output Fields

Field Name	Field Description
Item	Line card slot number.

Table 308: show chassis nonstop-upgrade Output Fields *(Continued)*

Field Name	Field Description
Status	State of line card: • **Error**—Line card is in an error state. • **Offline**—Line card is powered down. • **Online**—Line card is online and running.
Reason	Reason for the state (if the line card is offline).

Sample Output

show chassis nonstop-upgrade (EX8200 Switch)

```
user@switch> show chassis nonstop-upgrade
  Item           Status                  Reason
  FPC 0          Online
  FPC 1          Online
  FPC 2          Online
  FPC 3          Offline                 Offlined by CLI command
  FPC 4          Online
  FPC 5          Online
  FPC 6          Online
  FPC 7          Online
```

show chassis nonstop-upgrade (EX8200 Virtual Chassis)

```
user@external-routing-engine> show chassis nonstop-upgrade
member0:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online
  FPC 1          Online
  FPC 2          Online
  FPC 5          Online
```

```
member1:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online
  FPC 1          Offline                 Offlined due to config
  FPC 2          Online
  FPC 3          Online
  FPC 4          Online
  FPC 5          Online
  FPC 7          Online
```

show chassis nonstop-upgrade (Virtual Chassis Fabric)

```
 Item           Status                  Reason
  FPC 0          Online
  FPC 1          Online
  FPC 2          Online
  FPC 3          Online
  FPC 4          Online
  FPC 5          Online
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- request system software nonstop-upgrade | **15319**
- Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)
- Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade
- *Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade*
- Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)

show chassis nonstop-upgrade node-group

IN THIS SECTION

- Syntax | 18585
- Description | 18585
- Required Privilege Level | 18585
- Output Fields | 18585
- Sample Output | 18586
- Release Information | 18586

Syntax

```
show chassis nonstop-upgrade node-group node-group-name
```

Description

Display the status of the Node group after the most recent nonstop software upgrade (NSSU).

Required Privilege Level

view

Output Fields

Table 309 on page 18585 lists the output fields for the `show chassis nonstop-upgrade node-group` command. Output fields are listed in the approximate order in which they appear.

Table 309: show chassis nonstop-upgrade node-group Output Fields

Field Name	Field Description
`Item`	Node device slot number.

Table 309: show chassis nonstop-upgrade node-group Output Fields *(Continued)*

Field Name	Field Description
`Status`	State of Node device: • `Error`—Node device is in an error state. • `Offline`—Node device is powered down. • `Online`—Node device is online and running.
`Reason`	Reason for the state (if the line card is offline).

Sample Output

show chassis nonstop-upgrade node-group

```
user@qfabric> show chassis nonstop-upgrade node-group NW-NG-0
 Item           Status                  Reason
  P1550-C        Online
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Performing a Nonstop Software Upgrade on the QFabric System

request system software nonstop-upgrade (QFabric Systems) | **15311**

show chassis oss-map

IN THIS SECTION

- Syntax | 18587
- Description | 18587
- Required Privilege Level | 18587
- Output Fields | 18587
- Sample Output | 18588
- Release Information | 18588

Syntax

```
show chassis oss-map
```

Description

(T4000 routers only) Display the operations support systems (OSS) mapping details.

Required Privilege Level

view

Output Fields

Table 310 on page 18587 lists the output fields for the `show chassis oss-map` command. Output fields are listed in the approximate order in which they appear.

Table 310: show chassis oss-map Output Fields

Field Name	Field Description
`Chassis type`	Displays the original chassis type.

Table 310: show chassis oss-map Output Fields *(Continued)*

Field Name	Field Description
Oss-map	Displays the mapped chassis type.

Sample Output

command-name

```
user@T4000# show chassis oss-map
Chassis type             Oss-map
T4000                    T640
```

Release Information

Command introduced in Junos OS Release 12.3R3, 13.1R2, and 13.2R1.

RELATED DOCUMENTATION

Configuring OSS Mapping to Represent a T4000 Chassis as a T1600 or a T640 Chassis

Example: Configuring a T4000 Chassis to Represent a T640 Chassis

oss-map | **11944**

show chassis pic

IN THIS SECTION

- Syntax | **18589**
- Syntax (TX Matrix and TX Matrix Plus Routers) | **18589**
- Syntax (MX Series Routers and EX Series Switches) | **18589**

- Syntax (PTX Series Packet Transport Router and MX240, MX480, MX960, MX2010, and MX2020 Routers) | **18589**
- Syntax (QFX Series) | **18590**
- Syntax (ACX500, ACX5048, ACX5096, and ACX7509 Routers) | **18590**
- Description | **18590**
- Options | **18590**
- Required Privilege Level | **18592**
- Output Fields | **18593**
- Sample Output | **18597**
- Release Information | **18628**

Syntax

```
show chassis pic fpc-slot slot-number  pic-slot slot-number
```

Syntax (TX Matrix and TX Matrix Plus Routers)

```
show chassis pic fpc-slot slot-number  pic-slot slot-number
<lcc number>
```

Syntax (MX Series Routers and EX Series Switches)

```
show chassis pic fpc-slot slot-number   pic-slot slot-number
<all-members>
<local>
<member member-id>
```

Syntax (PTX Series Packet Transport Router and MX240, MX480, MX960, MX2010, and MX2020 Routers)

```
show chassis pic transport fpc-slot  slot-number pic-slot slot-number
```

Syntax (QFX Series)

```
show chassis pic fpc-slot slot-number  pic-slot slot-number
<interconnect-device name (fpc-slot slot-number | pic-slot slot-number)>
<node-device name pic-slot slot-number>
```

Syntax (ACX500, ACX5048, ACX5096, and ACX7509 Routers)

```
show chassis pic
(fpc-slot slot-number | pic-slot slot-number)
```

Description

Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.

Options

fpc-slot *slot-number*

Display information about the PIC in this particular FPC slot:

- On a TX Matrix router, if you specify the number of the T640 router by using the `lcc number` option (the recommended method), replace *slot-number* with a value from `0` through `7`. Otherwise, replace *slot-number* with a value from `0` through `31`.

  Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 router by using the `lcc number` option (the recommended method), replace *slot-number* with a value from `0` through `7`. Otherwise, replace *slot-number* with a value from `0` through `31`. For example, the following commands have the same result:

  ```
  user@host> show chassis pic fpc-slot 1 lcc 1 pic-slot 1
  user@host> show chassis pic fpc-slot 9 pic-slot 1
  ```

- M120 routers only—Replace *slot-number* with a value from `0` through `5`.
- MX80 routers only—Replace *slot-number* with a value from `0` through `1`.
- MX104 routers only—Replace *slot-number* with a value from `0` through `2`.

- MX240 routers only—Replace *slot-number* with a value from 0 through 2.
- MX480 routers only—Replace *slot-number* with a value from 0 through 5.
- MX960 routers only—Replace *slot-number* with a value from 0 through 11.
- MX2010 routers only—Replace *slot-number* with a value from 0 through 9.
- MX2020 routers only—Replace *slot-number* with a value from 0 through 19.
- MX2008 routers only—Replace *slot-number* with a value from 0 through 9.
- MX10003 routers only—Replace *slot-number* with a value from 0 through 1.
- Other routers—Replace *slot-number* with a value from 0 through 7.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—Replace *slot-number* with 0.
  - EX4200 switches in a Virtual Chassis configuration—Replace *slot-number* with a value from 0 through 9 (switch's member ID).
  - EX8208 switches—Replace *slot-number* with a value from 0 through 7 (line card).
  - EX8216 switches—Replace *slot-number* with a value from 0 through 15 (line card).
- QFX Series:
  - QFX3500, QFX3600, QFX5100, and OCX Series standalone switches—Replace *slot-number* with 0. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.
  - QFabric systems—Replace *slot-number* with any number between 0 and 15. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.

all-members (MX Series routers and EX Series switches only) (Optional) Display PIC information for all member routers in the Virtual Chassis configuration.

interconnect-device *name* (QFabric systems only) (Optional) Display PIC information for a specified Interconnect device.

lcc *number*

(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display PIC information for a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display PIC information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(MX Series routers and EX Series switches only) (Optional) Display PIC information for the local Virtual Chassis member.

member *member-id*

(MX Series routers and EX Series switches only) (Optional) Display PIC information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device *name*

(QFabric systems only) (Optional) Display PIC information for a specified Node device.

pic-slot *slot-number*

Display information about the PIC in this particular PIC slot. For routers, replace *slot-number* with a value from 0 through 3. For EX3200 and EX4200 switches, replace *slot-number* with 0 for built-in network interfaces and 1 for interfaces on uplink modules. For EX8208 and EX8216 switches, replace *slot-number* with 0. For the QFX3500 standalone switch and the QFabric system, replace *slot-number* with 0 or 1.

transport

Display PIC information for optical transport network.

Required Privilege Level

view

Output Fields

Table 311 on page 18593 lists the output fields for the `show chassis pic` command. Output fields are listed in the approximate order in which they appear.

Table 311: show chassis pic Output Fields

Field Name	Field Description
`Type`	PIC type. **NOTE**: On the 1-port OC192/STM64 MICs with the SDH framing mode, the type is displayed as `MIC-3D-1STM64-XFP` and with the SONET framing mode, the type is displayed as `MIC-3D-1OC192-XFP`. By default, the 1-port OC192/STM64 MICs displays the type as `MIC-3D-1OC192-XFP`.
`Account Layer2 Overhead`	(MX Series routers) Indicates whether functionality to count the Layer 2 overhead bytes in the interface statistics at the PIC level is enabled or disabled.
`ASIC type`	Type of ASIC on the PIC.
`State`	Status of the PIC. State is displayed only when a PIC is in the slot. • `Online`— PIC is online and running. • `Offline`—PIC is powered down. • `Empty`—No PIC is present. • `Present`—PIC is plugged in. The PIC is not powered on or operational. • `Onlining`—PIC is in the process of going online. PICs and rest of the hardware is initializing. • `Offlining`—PIC is in the process of going offline. PIC and rest of the hardware is being shutdown down to take the offline gracefully. • `Fault`—PIC is in an alarmed state and the PIC is not operational.

Table 311: show chassis pic Output Fields *(Continued)*

Field Name	Field Description
Substate	Display the internal state for the following transient states to help understand the FRU progress: • PowerOnWait • PowerOnPermission • PowerOnStart • TRANSITION_ONLINE
`PIC version`	PIC hardware version.
`Uptime`	How long the PIC has been online.
`Package`	(Multiservices interfaces cards only) The following services package are supported: • `Layer-2` (not applicable to the MS-MPC and MS-MIC.) • `Layer-3` (not applicable to the MS-MPC and MS-MIC.) • `extension-provider` (PICs of MS-MPC and MS-MIC support only this package.)
`Port Number`	Port number for the PIC.
`Cable Type`	Type of cable connected to the port: `LH`, `LX`, or `SX`.

Table 311: show chassis pic Output Fields *(Continued)*

Field Name	Field Description
`PIC Port Information (MX480 Router 100-Gigabit Ethernet CFP)`	Port-level information for the PIC. • Port—Port number • Cable type—Type of optical transceiver installed. • Fiber type—Type of fiber. SM is single-mode. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. Uplinks and downlinks are always 1550 nm. There is a separate fiber for each direction • Xcvr Firmware—Transceiver firmware version.

`PIC Port Information (MX960 Router Bidirectional Optics )`	Port-level information for the PIC.
	• Port—Port number
	• Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. Uplink interfaces display -U. Down link interfaces display -D.
	• Fiber type—Type of fiber. SM is single-mode.
	• Xcvr vendor—Transceiver vendor name.
	• Xcvr vendor part number—Transceiver vendor part number.
	  • BX10-10-km bidirectional optics.
	  • BX40-40-km bidirectional optics.
	  • SFP-LX-40-km SFP optics.
	• Wavelength—Wavelength of the transmitted signal. Uplinks are always 1310 nm. Downlinks are either 1490 nm or 1550 nm.

Table 311: show chassis pic Output Fields *(Continued)*

Field Name	Field Description
`PIC Port Information (Next-Generation SONET/SDH SFP)`	Port-level information for the next-generation SONET/SDH SFP PIC. • Port—Port number. • Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. • Fiber type—Type of fiber: `SM` (single-mode) or `MM` (multimode). • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. Next-generation SONET/SDH SFPs use 1310 nm.
`PIC port information (MX104 router)`	Port-level information for the PIC. • Port—Port number • Cable type—Type of optical transceiver installed. • Fiber type—Type of fiber. SM is single-mode. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. • Xcvr Firmware—Firmware version of the transceiver.
MSA Version	Multi-source Agreements (MSA) specification version that the specified optics is compliant to. Values supported are: • SFP+/SFP28—SFF 8472 (Versions 9.3 through 12,3) • QSFP+/QSFP28—SFF 8363 (Versions 1.3 through 2.10) • QSFP-DD—CMIS (Versions 3.0, 4.0, and 5.0)

Table 311: show chassis pic Output Fields *(Continued)*

Field Name	Field Description
Port speed information	Information pertaining to port speed: • Port—Port number. • PFE—Packet Forwarding Engine slot number. • Capable Port Speed—Speed supported by each port.
`Multirate Mode`	Rate-selectabilty status for the MIC: `Enabled` or `Disabled`.
`Channelization`	Indicates whether channelization is enabled or disabled on the DS3/E3 MIC.
`Administrative State`	Indicates the administrative state of the PIC. Possible values are: In Service (Default) and Out of Service.
`Operational State`	Indicates the operational state of the PIC. Possible values are: Normal and Fault.
Fault Type	Category of fault. Possible values are defined in **enum FaultCategory** in **Fru.evl**.
Fault Time	Millisecond resolution time instant when the FRU was marked as faulty. This helps correlate with the logs.

Sample Output

show chassis pic fpc-slot pic-slot

```
user@host> show chassis pic fpc-slot 2 pic-slot 0
PIC fpc slot 2 pic slot 0 information:
  Type                            10x 1GE(LAN), 1000 BASE
  ASIC type                       H chip
  State                           Online
  PIC version                     1.1
  Uptime                          1 day, 50 minutes, 58 seconds
```

```
PIC Port Information:
  Port         Cable           Xcvr               Xcvr Vendor
  Number       Type            Vendor Name        Part Number
  0            GIGE 1000EX     FINISAR CORP.      FTRJ8519P1BNL-J3
  1            GIGE 1000EX     FINISAR CORP.      FTRJ-8519-7D-JUN
```

show chassis pic fpc-slot pic-slot (PIC Offline)

```
user@host> show chassis pic fpc-slot 1 pic-slot 0
PIC fpc slot 1 pic slot 0 information:
  State                           Offline
```

show chassis pic fpc-slot pic-slot (FPC Offline)

```
user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC 1 is not online
```

show chassis pic fpc-slot pic-slot (FPC Not Present)

```
user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4 is empty
```

show chassis pic fpc-slot pic-slot (PIC Not Present)

```
user@host> show chassis pic fpc-slot 5 pic-slot 2
FPC 5, PIC 2 is empty
```

show chassis pic fpc-slot 3 pic-slot 0 (M120 Router)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
PC slot 3, PIC slot 0 information:
  Type                            2x G/E IQ, 1000 BASE
  ASIC type                       IQ GE 2 VLAN-TAG FPGA
  State                           Online
  PIC version                 1.16
```

```
  Uptime                        3 hours, 3 minutes

PIC Port Information:
  Port        Cable             Xcvr               Xcvr Vendor
  Number      Type              Vendor Name        Part Number
  0           GIGE 1000SX       FINISAR CORP.      FTRJ8519P1BNL-J3
  1           GIGE 1000SX       FINISAR CORP.      FTRJ-8519-7D-JUN
```

show chassis pic fpc-slot pic-slot (MX150)

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                              Virtual
  State                             Online
  PIC version                  0.0
  Uptime                          7 days, 19 hours, 44 minutes, 40 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  10   GIGE 1000T        n/a   Methode Elec.      SP7041-M1-JN      n/a      0.0
  11   GIGE 1000T        n/a   Methode Elec.      SP7041-M1-JN      n/a      0.0
```

show chassis pic pic-slot 0 fpc-slot 1 (MX960 Router with MPC10E-10C-MRATE, MPC10E-15C-MRATE line cards)

```
user@host> show chassis pic pic-slot 0 fpc-slot 1

FPC slot 0, PIC slot 2 information:
  Type                              MRATE-5xQSFPP
  State                             Online
  PIC version                       0.0
  Uptime               18 hours, 41 minutes, 27 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-
Xcvr          JNPR     MSA
  Port Cable type        type  Xcvr vendor        part number       length
Firmware      Rev      Version
```

```
  4    400G-ZR           SM    JUNIPER-2E1        740-157132        1528.77 nm - 1568.77 nm
2.3           REV 01   CMIS 5.0

Port speed information:

  Port  PFE      Capable Port Speeds
  0      2       4x10GE, 4x25GE, 40GE, 100GE
  1      2       4x10GE, 4x25GE, 40GE, 100GE
  2      2       4x10GE, 4x25GE, 40GE, 100GE
  3      2       4x10GE, 4x25GE, 40GE, 100GE
  4      2       4x10GE, 4x25GE, 40GE, 100GE, 4x100GE, 400GE
```

show chassis pic fpc-slot pic-slot (MX960 Router with Bidirectional Optics)

```
user@host> show chassis pic fpc-slot 4 pic-slot 1
FPC slot 4, PIC slot 1 information:
  Type                           10x 1GE(LAN)
  Account Layer2 Overhead        Enabled
  State                          Online
  PIC version                    0.0
  Uptime                         18 days, 5 hours, 41 minutes, 54 seconds

PIC port information:
                           Fiber                    Xcvr vendor
  Port  Cable type         type  Xcvr vendor        part number       Wavelength
  0     SFP-1000BASE-BX10-D SM  SumitomoElectric   SBP6H44-J3-BW-49  1490 nm
  1     SFP-1000BASE-BX10-D SM  SumitomoElectric   SBP6H44-J3-BW-49  1490 nm
  2     SFP-1000BASE-BX10-D SM  SumitomoElectric   SBP6H44-J3-BW-49  1490 nm
  3     SFP-1000BASE-BX10-D SM  OCP                TRXBG1LXDBVM2-JW  1490 nm
  4     SFP-1000BASE-BX10-D SM  OCP                TRXBG1LXDBVM2-JW  1490 nm
  5     SFP-1000BASE-BX10-U SM  SumitomoElectric   SBP6H44-J3-BW-31  1310 nm
  6     SFP-1000BASE-BX10-U SM  SumitomoElectric   SBP6H44-J3-BW-31  1310 nm
  7     SFP-1000BASE-BX10-U SM  OCP                TRXBG1LXDBBMH-J1  1310 nm
  8     SFP-1000BASE-BX10-U SM  OCP                TRXBG1LXDBBMH-J1  1310 nm
  9     SFP-1000BASE-BX10-U SM  SumitomoElectric   SBP6H44-J3-BW-31  1310 nm
```

show chassis pic fpc-slot pic-slot (MX480 Router with 100-Gigabit Ethernet MIC)

```
user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
```

```
  Type                        1X100GE CFP
  State                       Online
  PIC version                 2.10
  Uptime                      4 minutes, 48 seconds

PIC port information:
        Fiber                    Xcvr vendor
  Port  Cable type        type  Xcvr vendor        part number       Wavelength
  0     100GBASE LR4      SM    FINISAR CORP.      FTLC1181RDNS-J3   1310 nm

        Xcvr vendor
        firmware version
        1.8
```

show chassis pic fpc-slot pic-slot (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```
user@host>show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                              AS-MXC
  State                             Online
  PIC version                  1.0
  Uptime                         11 hours, 18 minutes, 3 seconds
```

show chassis pic fpc-slot pic-slot (MX960 Router with MPC5EQ)

```
user@host> show chassis pic fpc-slot 0 pic-slot 3
FPC slot 0, PIC slot 3 information:
  Type                              1X100GE CFP2 OTN
  State                             Online
  PIC version                  0.0
  Uptime                         1 hour, 22 minutes, 42 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    100GBASE LR4      n/a   Oclaro Inc.        TRB5E20FNF-LF150  1309 nm  1.0
```

show chassis pic fpc-slot pic-slot (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type                              1X100GE DWDM CFP2-ACO
  State                             Online
  PIC version                  1.3
  Uptime                         9 hours, 4 minutes, 43 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    100G LH           SM    OCLARO             TRB100AJ-01       1528.77 nm - 1568.36 nm 20.10
```

show chassis pic fpc-slot pic-slot (MX10003 Routers)

```
user@host > show chassis pic fpc-slot 0 pic-slot 0

FPC slot 0, PIC slot 1 information:
  Type                              MIC1
  State                             Online
  PIC version                  1.5
  Uptime                         13 hours, 54 minutes, 33 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU2   850 nm   0.0
  11   40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU2   850 nm   0.0

Port speed information:

  Port  PFE      Capable Port Speeds
  0      0       4x10GE, 40GE, 100GE
  1      0       4x10GE, 40GE, 100GE
  2      0       4x10GE, 40GE, 100GE
  3      0       4x10GE, 40GE, 100GE
  4      1       4x10GE, 40GE, 100GE
  5      1       4x10GE, 40GE, 100GE
```

```
  6      1       4x10GE, 40GE, 100GE
  7      1       4x10GE, 40GE, 100GE
  8      2       4x10GE, 40GE, 100GE
  9      2       4x10GE, 40GE, 100GE
  10     2       4x10GE, 40GE, 100GE
  11     2       4x10GE, 40GE, 100GE
```

show chassis pic fpc-slot pic-slot (PTX1000 and PTX10000)

```
user@host > show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                              288X10GE/72X40GE/24X100GE
  State                             Online
  PIC version                 1.18
  Uptime                         9 day, 5 hours, 10 minutes, 56 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  13   100GBASE LR4      SM    JUNIPER-SOURCE     SPQCELRCDFBJ2     1302 nm  0.0
  25   100GBASE LR4      SM    JUNIPER-SOURCE     SPQCELRCDFAJ2     1302 nm  0.0
  36   40GBASE LR4       SM    FINISAR CORP.      FTL4C1QE1C-J1     1301 nm  0.0
  37   40GBASE LR4       SM    FINISAR CORP.      FTL4C1QE1C-J1     1301 nm  0.0
  54   40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU1   850 nm   0.0

Port speed information:
 Port  PFE      Capable Port Speeds
  0              4x10GE, 40GE
  1              4x10GE, 40GE, 100GE
  2              4x10GE, 40GE
  3              4x10GE, 40GE
  4              4x10GE, 40GE
  5              4x10GE, 40GE, 100GE
  6              4x10GE, 40GE
  7              4x10GE, 40GE, 100GE
  8              4x10GE, 40GE
  9              4x10GE, 40GE
  10             4x10GE, 40GE
  11             4x10GE, 40GE, 100GE
  12             4x10GE, 40GE
  13             4x10GE, 40GE, 100GE
```

```
14              4x10GE, 40GE
15              4x10GE, 40GE
16              4x10GE, 40GE
17              4x10GE, 40GE, 100GE
18              4x10GE, 40GE
19              4x10GE, 40GE, 100GE
20              4x10GE, 40GE
21              4x10GE, 40GE
22              4x10GE, 40GE
23              4x10GE, 40GE, 100GE
24              4x10GE, 40GE
25              4x10GE, 40GE, 100GE
26              4x10GE, 40GE
27              4x10GE, 40GE
28              4x10GE, 40GE
29              4x10GE, 40GE, 100GE
30              4x10GE, 40GE
31              4x10GE, 40GE, 100GE
32              4x10GE, 40GE
33              4x10GE, 40GE
34              4x10GE, 40GE
35              4x10GE, 40GE, 100GE
36              4x10GE, 40GE
37              4x10GE, 40GE, 100GE
38              4x10GE, 40GE
39              4x10GE, 40GE
40              4x10GE, 40GE
41              4x10GE, 40GE, 100GE
42              4x10GE, 40GE
43              4x10GE, 40GE, 100GE
44              4x10GE, 40GE
45              4x10GE, 40GE
46              4x10GE, 40GE
47              4x10GE, 40GE, 100GE
48              4x10GE, 40GE
49              4x10GE, 40GE, 100GE
50              4x10GE, 40GE
51              4x10GE, 40GE
52              4x10GE, 40GE
53              4x10GE, 40GE, 100GE
54              4x10GE, 40GE
55              4x10GE, 40GE, 100GE
56              4x10GE, 40GE
```

```
  57             4x10GE, 40GE
  58             4x10GE, 40GE
  59             4x10GE, 40GE, 100GE
  60             4x10GE, 40GE
  61             4x10GE, 40GE, 100GE
  62             4x10GE, 40GE
  63             4x10GE, 40GE
  64             4x10GE, 40GE
  65             4x10GE, 40GE, 100GE
  66             4x10GE, 40GE
  67             4x10GE, 40GE, 100GE
  68             4x10GE, 40GE
  69             4x10GE, 40GE
  70             4x10GE, 40GE
```

show chassis pic fpc-slot pic-slot (MSA Version)

```
user@host > show chassis pic fpc-slot pic-slot
..
PIC port information:
                     Fiber                Xcvr vendor     Wave-   Xcvr   JNPR   MSA
Port Cable type      type  Xcvr vendor    part number    length Firmware Rev   Version
0   2x100GBASE-SR4   MM   JUNIPER-1X1     740-084673      850 nm  6.0   REV 01 CMIS 3.0
1   2x100GBASE-SR4   MM   JUNIPER-1X1     740-084673      850 nm  6.0   REV 01 CMIS 3.0
3   10GBASE SR       MM   FINISAR CORP.   FTLX8571D3BNL-J1850 nm  0.0   REV 01 SFF-8472 ver 10.2
4   100GBASE SR4 T2  MM   JUNIPER-FINISAR FTLC9551REPM-J1 850 nm  0.0   REV 01 SFF-8636 ver 2.7
```

show chassis pic fpc-slot pic-slot (PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host > show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
  Type                              5X100GE DWDM CFP2-ACO
  State                             Online
  PIC version                 1.17
  Uptime                         1 day, 5 hours, 15 minutes, 17 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
```

```
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    100G LH           SM    MULTILANE SAL      ML4030-ACO-2      1528.77 nm - 1568.36 nm
1.0
  1    100G LH           SM    MULTILANE SAL      ML4030-ACO-2      1528.77 nm - 1568.36 nm
1.0
  2    100G LH           SM    JUNIPER-FUJITSU    FIM38500/222      1528.77 nm - 1568.36 nm
1.16
  3    100G LH           SM    FUJITSU            FIM38500/222      1528.77 nm - 1568.36 nm
1.16
  4    100G LH           SM    FUJITSU            FIM38500/222      1528.77 nm - 1568.36 nm 1.16
```

show chassis pic fpc-slot pic-slot (MX480 Router with MPC4E)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type                              4x10GE SFPP
  State                             Online
  PIC version                  0.0
  Uptime                         41 seconds

PIC port information:
                         Fiber                     Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    10GBASE SR        MM    OPNEXT, INC.       TRS2001EM-0014    850 nm   0.0
  1    10GBASE SR        MM    OPNEXT, INC.       TRS2001EM-0014    850 nm   0.0
```

show chassis pic fpc-slot pic-slot (MX480 router with OTN Interface)

```
user@host> show chassis pci fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
  Type                              12X10GE SFPP OTN
  State                             Online
  PIC version                  0.0
  Uptime                         5 hours, 28 minutes, 23 seconds

PIC port information:
                         Fiber                     Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
```

```
  1    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
  2    10GBASE SR        MM    OPNEXT, INC.       TRS2001EM-0014    850 nm   0.0
```

show chassis pic fpc-slot pic-slot (MX2010 Router with OTN Interfaces)

```
user@host> show chassis pic fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
  Type                              2X100GE CFP2 OTN
  State                             Online
  PIC version                  1.9
  Uptime                         3 hours, 56 minutes, 16 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    100GBASE LR4-D    SM    FUJITSU            FIM37300/222      1310 nm  1.3
  1    100GBASE SR10     MM    AVAGO              AFBR-8420Z        n/a      1.0
```

show chassis pic fpc-slot pic-slot (MX2010 Router)

```
user@host> show chassis pic fpc-slot 9 pic-slot 3
FPC slot 9, PIC slot 3 information:
  Type                              1X100GE CFP
  Account Layer2 Overhead           Enabled
  State                             Online
  PIC version                       0.0
  Uptime                            14 hours, 51 seconds
```

show chassis pic fpc-slot pic-slot (MX2020 Router)

```
user@host> show chassis pic fpc-slot 19 pic-slot 3
FPC slot 19, PIC slot 3 information:
  Type                              4x 10GE(LAN) SFP+
  Account Layer2 Overhead           Enabled
  State                             Online
  PIC version                  0.0
  Uptime                         1 day, 11 hours, 26 minutes, 36 seconds
```

```
PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    10GBASE SR        MM    SumitomoElectric   SPP5200SR-J6-M    850 nm   0.0
  1    10GBASE SR        MM    SumitomoElectric   SPP5200SR-J6-M    850 nm   0.0
  2    10GBASE SR        MM    SumitomoElectric   SPP5200SR-J6-M    850 nm   0.0
  3    10GBASE SR        MM    SumitomoElectric   SPP5200SR-J6-M    850 nm   0.0
```

show chassis pic fpc-slot pic-slot (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis pic fpc-slot 18 pic-slot 2
FPC slot 18, PIC slot 2 information:
  Type                              3X40GE QSFPP
  State                             Online
  PIC version                  0.0
  Uptime                         6 minutes, 31 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    40GBASE SR4       MM    AVAGO              AFBR-79E4Z-D-JU2  850 nm   0.0
  1    40GBASE SR4       MM    AVAGO              AFBR-79E4Z-D-JU2  850 nm   0.0
  2    40GBASE SR4       MM    AVAGO              AFBR-79E4Z-D-JU2  850 nm   0.0
```

show chassis pic fpc-slot pic-slot (MX2020 Router with MPC6E and OTN MIC)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 0, PIC slot 1 information:
  Type                              24X10GE SFPP OTN
  State                             Online
  PIC version                  1.1
  Uptime                         1 hour, 33 minutes, 59 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  7    10GBASE SR        MM    SumitomoElectric   SPP5200SR-J6-M    850 nm   0.0
  9    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  12   10GBASE LR        SM    FINISAR CORP.      FTLX1472M3BNL-J3  1310 nm  0.0
  20   10GBASE ZR        SM    FINISAR CORP.      FTLX1871M3BNL-J3  1550 nm  0.0
```

```
  21   10GBASE ER        SM    FINISAR CORP.      FTLX1671D3BTL-J4  1550 nm  0.0
  22   10GBASE LR        SM    SOURCEPHOTONICS    SPP10SLREDFCJNP   1310 nm  0.0
  23   10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BNL-J1  1310 nm  0.0
```

show chassis pic fpc-slot pic-slot (MX2020 Router with MPC4E)

```
user@host> show chassis pic fpc-slot 14 pic-slot 0
FPC slot 14, PIC slot 2 information:
  Type                              4x10GE SFPP
  State                             Online
  PIC version                  0.0
  Uptime                         1 day, 14 hours, 49 minutes, 9 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
  1    10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
  3    10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
```

show chassis pic fpc-slot pic-slot (MX2010 Router)

```
user@host> show chassis pic fpc-slot 9 pic-slot 3
FPC slot 9, PIC slot 3 information:
  Type                              1X100GE CFP
  Account Layer2 Overhead           Enabled
  State                             Online
  PIC version                       0.0
  Uptime                            14 hours, 51 seconds
```

show chassis pic fpc-slot pic-slot (T1600 Router with 100-Gigabit Ethernet PIC)

```
user@host> run show chassis pic fpc-slot 3 pic-slot 1
FPC slot 3, PIC slot 1 information:
  Type                              100GE SLOT1
  ASIC type                         Brooklyn 100GE FPGA
  State                             Online
  PIC version                  1.3
```

```
  Uptime                         10 minutes, 44 seconds

PIC port information:
                           Fiber                    Xcvr vendor
  Port  Cable type         type  Xcvr vendor        part number       Wavelength
  0     100GBASE LR4       SM    Opnext Inc.        TRC5E20ENFSF000F  1310 nm
```

show chassis pic fpc-slot pic-slot lcc (TX Matrix Router)

```
user@host> show chassis pic fpc-slot 1 pic-slot 1 lcc 0
lcc0-re0:
--------------------------------------------------------------------------
PIC fpc slot 1 pic slot 1 information:
  Type                             4x OC-3 SONET, SMIR
  ASIC type                        D chip
  State                            Online
  PIC version                  1.2
  Uptime                        5 days, 2 hours, 12 minutes, 8 seconds
```

show chassis pic fpc-slot pic-slot lcc (TX Matrix Plus Router)

```
user@host> show chassis pic pic-slot 0 fpc-slot 8
lcc0-re0:
--------------------------------------------------------------------------
FPC slot 8, PIC slot 0 information:
  Type                             1x 10GE(LAN/WAN)
  State                            Online
  Uptime                         2 hours, 46 minutes, 23 seconds

PIC port information:
                           Fiber
  Port  Cable type         type  Xcvr vendor         part number       Wavelength
  0     10GBASE ZR         SM    Opnext Inc.         TRF7061BN-LF150   1550 nm
  0     10GBASE ZR         SM    FINISAR CORP.       FTRX-1811-3-J2    1550 nm
```

show chassis pic fpc-slot pic-slot (Next-Generation SONET/SDH SFP)

```
user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
   Type                              4x OC-3 1x OC-12 SFP
   ASIC type                         D FPGA
   State                             Online
   PIC version                  1.3
   Uptime                        1 day, 50 minutes, 4 seconds

 PIC port information:
                           Fiber                    Xcvr vendor
   Port  Cable type        type  Xcvr vendor        part number       Wavelength
   0     OC48 short reach  SM    FINISAR CORP.      FTRJ1321P1BTL-J2  1310 nm
   1     OC3 short reach   MM    OCP                TRPA03MM3BAS-JE   1310 nm
   2     OC3 short reach   MM    OCP                TRXA03MM3BAS-JW   1310 nm
   3     OC12 inter reach  SM    FINISAR CORP.      FTLF1322P1BTR     1310 nm
```

show chassis pic fpc-slot pic-slot (12-Port T1/E1)

```
user@host> show chassis pic fpc-slot 0 pic-slot 3
FPC slot 0, PIC slot 3 information:
  Type                              12x T1/E1 CE
  State                             Online
  PIC version                  1.1
  CPU load average                1 percent
  Interrupt load average          0 percent
  Total DRAM size               128 MB
  Memory buffer utilization     100 percent
  Memory heap utilization         4 percent
  Uptime                          1 day, 22 hours, 28 minutes, 12 seconds
  Internal Clock Synchronization   Normal
```

show chassis pic fpc-slot 0 pic-slot 1 (4x CHOC3 SONET CE SFP)

```
user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
  Type                              4x CHOC3 SONET CE SFP
  State                             Online
```

```
  PIC version                  1.3
  CPU load average               1 percent
  Interrupt load average         0 percent
  Total DRAM size              128 MB
  Memory buffer utilization     99 percent
  Memory heap utilization        4 percent
  Uptime                         1 day, 22 hours, 55 minutes, 37 seconds
  Internal Clock Synchronization   Normal

PIC port information:
                           Fiber                    Xcvr vendor
  Port  Cable type         type  Xcvr vendor        part number       Wavelength
  0     OC3 short reach    MM    AVAGO              HFBR-57E0P-JU2    n/a
  1     OC3 short reach    MM    AVAGO              HFBR-57E0P-JU2    n/a
  3     OC3 long reach     SM    OPNEXT INC         TRF5456AVLB314    1310 nm
```

show chassis pic fpc-slot 0 pic-slot 0 (SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                             MIC-3D-8OC3OC12-4OC48
  State                            Online
  PIC version                  1.8
  Uptime                         3 days, 22 hours, 3 minutes, 50 seconds

PIC port information:
                           Fiber                    Xcvr vendor
  Port  Cable type         type  Xcvr vendor        part number       Wavelength
  1     OC12 inter reach    SM   FINISAR CORP       FTRJ1322P1BTR-J3  1310 nm
  7     OC12 inter reach    SM   FINISAR CORP       FTRJ1322P1BTR-J3  1310 nm
Multirate Mode                     Enabled
```

show chassis pic fpc-slot 3 pic-slot 0 (8-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type                             MIC-3D-8CHOC3-4CHOC12
  State                            Online
```

```
  PIC version                        1.9
  Uptime                               1 hour, 21 minutes, 24 seconds

PIC port information:
                           Fiber                    Xcvr vendor
  Port  Cable type         type  Xcvr vendor        part number       Wavelength
  0     OC12 short reach   SM    FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  1     OC12 short reach   SM    FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  2     OC12 inter reach   SM    FINISAR CORP.      FTRJ1322P1BTR-J2  1310 nm
  4     OC12 short reach   SM    FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  5     OC12 short reach   SM    FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  6     OC12 short reach   SM    FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  7     OC12 short reach   SM    FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
```

show chassis pic fpc-slot 5 pic-slot 0 (4-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
FPC slot 5, PIC slot 0 information:
  Type                         MIC-3D-4CHOC3-2CHOC12
  State                        Online
  PIC version                  1.9
  Uptime                       1 hour, 21 minutes

PIC port information:
                           Fiber                    Xcvr vendor
  Port  Cable type         type  Xcvr vendor        part number       Wavelength
  1     OC12 inter reach   SM    FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  2     OC12 inter reach   SM    FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  3     OC12 short reach   SM    FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
```

show chassis pic fpc-slot 1 pic-slot 0 (1-port OC192/STM64 MIC with XFP)

```
user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
  Type                             MIC-3D-1OC192-XFP
  State                            Online
  PIC version                      1.2
  Uptime                           1 day, 11 hours, 4 minutes, 6 seconds
```

```
PIC port information:
                           Fiber                    Xcvr vendor
  Port  Cable type         type  Xcvr vendor        part number       Wavelength
  0     OC192 short reach n/a   FINISAR CORP.      FTLX1412M3BCL-J3  1310 nm
```

show chassis pic fpc-slot 1 pic-slot 2 (8-port DS3/E3 MIC)

```
user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                              MIC-3D-8DS3-E3
  State                             Online
  PIC version                          1.10
  Uptime                            4 days, 1 hour, 29 minutes, 19 seconds
  Channelization Mode               Disabled
```

show chassis pic fpc-slot pic-slot (OTN)

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
PIC fpc slot 5 pic slot 0 information:
     Type                              1x10GE(LAN),OTN
     ASIC type                         H chip
     State                             Online
     PIC version                    1.0
     Uptime                         5 minutes, 50 seconds
```

show chassis pic fpc-slot pic-slot (QFX3500 Switch)

```
user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
Type 48x 10G-SFP+ Builtin
State Online
Uptime 3 days, 3 hours, 5 minutes, 20 seconds
```

show chassis pic fpc-slot pic-slot (QFX5100 Switches and OCX Series)

```
user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                             Unknown Builtin
  State                            Online
  Uptime                         1 day, 17 hours, 5 minutes, 9 seconds
```

show chassis pic interconnect-device fpc-slot pic-slot (QFabric Systems)

```
user@switch> show chassis pic interconnect-device interconnect1 fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
  Type                             16x 40G-GEBuiltin
  State                            Online
  Uptime                         2 hours, 47 minutes, 40 seconds
```

show chassis pic node-device fpc-slot pic-slot (QFabric System)

```
user@switch> show chassis pic node-device node1 pic-slot 0
FPC slot node1, PIC slot 0 information:
  Type                             48x 10G-SFP+Builtin
  State                            Online
  Uptime                         2 hours, 52 minutes, 37 seconds

PIC port information:
                          Fiber                    Xcvr vendor
  Port  Cable type        type  Xcvr vendor        part number       Wavelength
  0     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  1     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  2     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  3     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  4     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  5     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  6     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  7     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  8     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  9     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  10    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  11    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
```

```
  12    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  13    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  14    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  15    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  16    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  17    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  18    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  19    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  20    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  21    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  22    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  23    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  24    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  25    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  26    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  27    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  28    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  29    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  30    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  31    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  32    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  33    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  34    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  35    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  36    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  37    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  38    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  39    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  40    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  41    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  42    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  43    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  44    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  45    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  46    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
  47    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3      850 nm
```

show chassis pic fpc-slot 0 pic-slot 1 (ACX2000 Universal Access Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
```

```
  Type                              8x 1GE(LAN) RJ45 Builtin
  State                             Online
  Uptime                          6 days, 2 hours, 51 minutes, 11 seconds
```

show chassis pic FPC-slot 1 PIC-slot 0 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
  Type                              AS-MSC
  State                             Online
  PIC version                   1.6
  Uptime                          11 hours, 17 minutes, 56 seconds
```

show chassis pic FPC slot 1, PIC slot 2 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis pic fpc-slot 1 pic-slot 2
Type                              AS-MXC
  State                             Online
  PIC version                   1.0
  Uptime                          11 hours, 18 minutes, 3 seconds
```

show chassis pic transport fpc-slot pic-slot (PTX Series Packet Transport Routers)

```
user@host> show chassis pic transport fpc-slot 2 pic-slot 0
Administrative State:      In Service
Operational    State:      Normal
```

show chassis pic transport fpc-slot pic-slot (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show chassis pic transport fpc-slot 3 pic-slot 0
Administrative State:      In Service
Operational    State:      Normal
```

show chassis pic fpc-slot 0 pic-slot 0 (ACX5096 Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                              96x10G-8x40G
  State                             Online
  PIC version                  2.9
  Uptime                         21 hours, 28 minutes, 13 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
  1    10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BCL-J1  1310 nm  0.0
  3    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  4    10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
  5    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  6    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  7    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  8    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  9    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  10   10GBASE SR        MM    OPNEXT, INC.       TRS2001EN-0014    850 nm   0.0
  11   10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
  12   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  13   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
  14   10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
  15   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  16   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  17   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  18   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
  19   10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BCL-J1  1310 nm  0.0
  20   10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BNL-J1  1310 nm  0.0
  21   10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
  22   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  23   10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
  24   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
  25   10GBASE USR       MM    FINISAR CORP.      FTLX8570D3BCL-J1  850 nm   0.0
  26   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  27   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
  28   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  29   GIGE 1000SX       MM    FINISAR CORP.      FTLF8519P3BNL-J1  850 nm   0.0
  31   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
```

```
32   10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
33   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
34   10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
35   10GBASE USR       MM    FINISAR CORP.      FTLX8570D3BCL-J1  850 nm   0.0
36   10GBASE USR       MM    FINISAR CORP.      FTLX8570D3BCL-J1  850 nm   0.0
37   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
38   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
40   GIGE 1000LX10     SM    FINISAR CORP.      FTLF1318P2BTL-J1  1310 nm  0.0
41   10GBASE LR        SM    OPNEXT,INC         TRS5021EN-S201    1310 nm  0.0
42   10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BCL-J1  1310 nm  0.0
43   10GBASE LR        SM    SumitomoElectric   SPP5100LR-J3      1310 nm  0.0
44   10GBASE LR        SM    SumitomoElectric   SPP5100LR-J3      1310 nm  0.0
45   10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BCL-J1  1310 nm  0.0
46   10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BCL-J1  1310 nm  0.0
47   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
48   10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
49   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
50   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
51   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
52   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
53   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
54   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
55   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
56   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
57   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
58   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
59   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
60   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
61   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
62   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
63   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
64   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
65   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
66   10GBASE SR        MM    SumitomoElectric   SPP5200SR-J6-M    850 nm   0.0
67   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
68   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
69   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
70   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
71   10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BNL-J1  1310 nm  0.0
72   10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BCL-J1  1310 nm  0.0
73   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
74   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
75   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
```

```
  76   10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
  77   10GBASE USR       MM    OPNEXT, INC.       TRS20A0EN-0014    850 nm   0.0
  78   10GBASE USR       MM    OPNEXT, INC.       TRS20A0EN-0014    850 nm   0.0
  79   10GBASE LRM       MM    OPNEXT INC         TRS5001EN-0014    1310 nm  0.0
  80   10GBASE LRM       MM    OPNEXT INC         TRS5001EN-0014    1310 nm  0.0
  81   10GBASE USR       MM    OPNEXT, INC.       TRS20A0EN-0014    850 nm   0.0
  82   10GBASE USR       MM    OPNEXT, INC.       TRS20A0EN-0014    850 nm   0.0
  83   10GBASE USR       MM    OPNEXT, INC.       TRS20A0EN-0014    850 nm   0.0
  84   10GBASE USR       MM    OPNEXT, INC.       TRS20A0EN-0014    850 nm   0.0
  85   10GBASE LR        SM    OPNEXT,INC         TRS5021EN-S201    1310 nm  0.0
  86   10GBASE ER        SM    OPNEXT,INC         TRS7050EN-S201    1550 nm  0.0
  87   10GBASE LRM       MM    OPNEXT INC         TRS5001EN-0014    1310 nm  0.0
  88   10GBASE LRM       MM    OPNEXT INC         TRS5001EN-0014    1310 nm  0.0
  89   10GBASE LRM       MM    OPNEXT INC         TRS5001EN-0014    1310 nm  0.0
  90   10GBASE LRM       MM    OPNEXT INC         TRS5001EN-0014    1310 nm  0.0
  91   10GBASE USR       MM    FINISAR CORP.      FTLX8570D3BCL-J1  850 nm   0.0
  92   10GBASE USR       MM    FINISAR CORP.      FTLX8570D3BCL-J1  850 nm   0.0
  93   10GBASE LR        SM    SumitomoElectric   SPP5100LR-J3      1310 nm  0.0
  94   10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BNL-J1  1310 nm  0.0
  95   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  96   40GBASE SR4       MM    AVAGO              AFBR-79E4Z-D-JU1  850 nm   0.0
  97   40GBASE SR4       MM    AVAGO              AFBR-79E4Z-D-JU1  850 nm   0.0
  98   40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU1   850 nm   0.0
  99   40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU1   850 nm   0.0
  100  40GBASE CU 1M     n/a   Molex Inc.         1110409055        n/a      0.0
  101  40GBASE CU 1M     n/a   Molex Inc.         1110409055        n/a      0.0
  102  40GBASE CU 1M     n/a   Molex Inc.         1110409055        n/a      0.0
  103  40GBASE CU 1M     n/a   Molex Inc.         1110409055        n/a      0.0
```

show chassis pic fpc-slot 0 pic-slot 0 (ACX5048 Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                              96x10G-8x40G
  State                             Online
  PIC version                  2.9
  Uptime                         1 day, 5 hours, 27 minutes, 25 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr
```

```
  Port Cable type        type  Xcvr vendor        part number       length   Firmware
  0    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
  10   GIGE 1000SX       MM    FINISAR CORP.      FTLF8519P3BNL-J1  850 nm   0.0
  14   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  20   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0
  30   GIGE 1000SX       MM    FINISAR CORP.      FTLF8519P2BNL-J1  850 nm   0.0
  41   10GBASE SR        MM    OPNEXT, INC.       TRS2001EN-0014    850 nm   0.0
  46   GIGE 1000SX       MM    FINISAR CORP.      FTLF8519P2BNL-J1  850 nm   0.0
  64   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
  78   GIGE 1000SX       MM    AVAGO              AFBR-5715PZ-JU2   850 nm   0.0
  96   40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU1   850 nm   0.0
  99   40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU1   850 nm   0.0
  100  40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU1   850 nm   0.0
```

show chassis pic fpc-slot 0 pic-slot 0 (ACX500 Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                              2x 1GE(LAN) SFP Builtin
  State                             Online
  Uptime                           17 hours, 54 minutes, 45 seconds
```

show chassis pic fpc-slot 0 pic-slot 1 (ACX500 Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
  Type                              4x 1GE(LAN) RJ45, SFP Builtin
  State                             Online
  Uptime                           17 hours, 54 minutes, 45 seconds
```

show chassis pic transport fpc-slot pic-slot (PTX Series Packet Transport Routers)

```
user@host> show chassis pic transport fpc-slot 2 pic-slot 0
Administrative State:      In Service
Operational    State:      Normal
```

show chassis pic transport fpc-slot pic-slot (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show chassis pic transport fpc-slot 3 pic-slot 0
Administrative State:      In Service
Operational    State:      Normal
```

show chassis pic fpc-slot 7 pic-slot 1 (MX960 Router MPC10E-15C-MRATE Line Card)

```
user@router> show chassis pic fpc-slot 7 pic-slot 1

FPC slot 7, PIC slot 1 information:
  Type                              MRATE-5xQSFPP
  State                             Online
  PIC version                  0.0
  Uptime                         3 hours, 33 minutes, 21 seconds

PIC port information:
                         Fiber                    Xcvr vendor       Wave-    Xcvr         JNPR
  Port Cable type        type  Xcvr vendor        part number       length   Firmware     Rev
  0    100GBASE LR4      SM    JUNIPER-FINISAR    FTLC1151RDPL-J3   1302 nm  0.0          REV 01

Port speed information:

  Port  PFE      Capable Port Speeds
  0      1       4x10GE, 40GE, 100GE
  1      1       4x10GE, 40GE, 100GE
  2      1       4x10GE, 40GE, 100GE
  3      1       4x10GE, 40GE, 100GE
  4      1       4x10GE, 40GE, 100GE
```

show chassis pic fpc-slot 0 pic-slot 1 (PTX10001-36MR) (400G ZR and 400G ZR-M optics)

```
user@router> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
  Type                              8X400GE-MR + 4X100GE-MR
  State                             Online
  PIC version               255.09
  Uptime                    2 hours, 37 minutes, 52 seconds
```

```
PIC port information:
                        Fiber                    Xcvr vendor       Wave-
Xcvr          JNPR     MSA
  Port Cable type        type  Xcvr vendor        part number       length
Firmware      Rev      Version
  0    40GBASE SR4       MM    FINISAR CORP       FTL410QE3C-J1     850 nm
0.0           REV 01
  1    40GBASE SR4       MM    AVAGO              AFBR-79EQPZ-JU1   850 nm
0.0           REV 01
  2    400G-ZR           SM    Acacia Comm Inc.   DP04QSDD-E20-00B  1528.77 nm - 1568.77 nm
161.20        XXXX
  4    100GBASE SR4 T2   MM    JUNIPER-AVAGO      AFBR-89CDDZ-JU1   850 nm
0.0           REV 01
  10   400G-ZR-M         SM    JUNIPER-1T1        740-131169        1528.77 nm - 1568.77 nm
61.20         REV 01

Port speed information:

  Port  PFE      Capable Port Speeds
  0      1       1x10G 4x10G 1x40G 4x25G 1x100G 2x50G 8x25G 8x50G 2x100G 1x200G 4x100G 2x200G
1x400G
  1      1       1x10G 4x10G 1x40G 4x25G 1x100G 2x50G 8x25G 8x50G 2x100G 1x200G 4x100G 2x200G
1x400G
  2      1       1x10G 4x10G 1x40G 4x25G 1x100G 2x50G 8x25G 8x50G 2x100G 1x200G 4x100G 2x200G
1x400G
  3      1       1x10G 4x10G 1x40G 4x25G 1x100G 2x50G 8x25G 8x50G 2x100G 1x200G 4x100G 2x200G
1x400G
  4      1       1x10G 4x10G 1x40G 4x25G 1x100G
  5      1       1x10G 1x100G
  6      1       1x10G 4x10G 1x40G 4x25G 1x100G
  7      1       1x10G 1x100G
  8      1       1x10G 4x10G 1x40G 4x25G 1x100G 2x50G 8x25G 8x50G 2x100G 1x200G 4x100G 2x200G
1x400G
  9      1       1x10G 4x10G 1x40G 4x25G 1x100G 2x50G 8x25G 8x50G 2x100G 1x200G 4x100G 2x200G
1x400G
  10     1       1x10G 4x10G 1x40G 4x25G 1x100G 2x50G 8x25G 8x50G 2x100G 1x200G 4x100G 2x200G
1x400G
  11     1       1x10G 4x10G 1x40G 4x25G 1x100G 2x50G 8x25G 8x50G 2x100G 1x200G 4x100G 2x200G
1x400G
```

show chassis pic fpc-slot 7 pic-slot 0 (ACX7509-FPC-20Y, ACX Series)

```
user@router> show chassis pic fpc-slot 7 pic-slot 0
FPC slot 7, PIC slot 0 information:
  Type                              20x1/10/25/50-SFP56
  State                             Online
  PIC version               255.09
  Uptime                         2 minutes, 16 seconds
PIC port information:
                         Fiber                    Xcvr vendor       Wave-
Xcvr          JNPR     MSA
  Port Cable type        type  Xcvr vendor        part number       length
Firmware      Rev      Version
  2    10GBASE SR        MM    AVAGO              AFBR-709ASMZ-JU3  850 nm
0.0           REV 01   SFF-8472 ver 10.2
  3    10GBASE SR        MM    AVAGO              AFBR-709ASMZ-JU3  850 nm
0.0           REV 01   SFF-8472 ver 10.2
  4    10GBASE SR        MM    OPNEXT, INC.       TRS2001EM-0014    850 nm
0.0           REV 01   SFF-8472 ver 10.2
  5    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm
0.0           REV 01   SFF-8472 ver 10.2
  6    SFP28-END 100G-4x25G DAC BO 1M n/a JUNIPER-AMPHENOL NDAQGF-J101 n/a
0.0           REV 01   SFF-8472 ver n/a
  7    SFP28-END 100G-4x25G DAC BO 1M n/a JUNIPER-AMPHENOL NDAQGF-J101 n/a
0.0           REV 01   SFF-8472 ver n/a
  8    SFP28-END 100G-4x25G DAC BO 1M n/a JUNIPER-AMPHENOL NDAQGF-J101 n/a
0.0           REV 01   SFF-8472 ver n/a
  9    SFP28-END 100G-4x25G DAC BO 1M n/a JUNIPER-AMPHENOL NDAQGF-J101 n/a
0.0           REV 01   SFF-8472 ver n/a
  10   GIGE 1000EX       n/a   SOURCEPHOTONICS    SPGBEXIDFCJNP     1310 nm
0.0           REV 01   SFF-8472 ver 9.5
  11   SFP28-25G-BASE-SR SM    JUNIPER-1A1        1A131A            850 nm
0.0           REV 01   SFF-8472 ver 12.3
  12   SFP28-25G-BASE-SR SM    JUNIPER-1A1        1A131A            850 nm
0.0           REV 01   SFF-8472 ver 12.3
  13   SFP28-25G-BASE-SR SM    JUNIPER-1A1        1A131A            850 nm
0.0           REV 01   SFF-8472 ver 12.3
  14   SFP28-25G-BASE-SR SM    JUNIPER-1A1        1A131A            850 nm
0.0           REV 01   SFF-8472 ver 12.3
  15   SFP28-25G-BASE-SR SM    JUNIPER-1A1        1A131A            850 nm
0.0           REV 01   SFF-8472 ver 12.3
  16   GIGE 1000SX       n/a   AVAGO              AFBR-5715PZ-JU3   850 nm
```

```
0.0           REV 01   SFF-8472 ver 9.3
  18   10GBASE CU 1M     n/a   Amphenol           584990001         n/a
0.0           REV 01   SFF-8472 ver n/a
  19   10GBASE CU 1M     n/a   Amphenol           584990001         n/a
0.0           REV 01   SFF-8472 ver n/a
Port speed information:
  Port  PFE      Capable Port Speeds
  0      1       1x1G 1x10G 1x25G
  1      1       1x1G 1x10G 1x25G
  2      1       1x1G 1x10G 1x25G
  3      1       1x1G 1x10G 1x25G
  4      1       1x1G 1x10G 1x25G
  5      1       1x1G 1x10G 1x25G
  6      1       1x1G 1x10G 1x25G
  7      1       1x1G 1x10G 1x25G
  8      1       1x1G 1x10G 1x25G
  9      1       1x1G 1x10G 1x25G
  10     1       1x1G 1x10G 1x25G
  11     1       1x1G 1x10G 1x25G
  12     1       1x1G 1x10G 1x25G
  13     1       1x1G 1x10G 1x25G
  14     1       1x1G 1x10G 1x25G
  15     1       1x1G 1x10G 1x25G
  16     1       1x1G 1x10G 1x25G
  17     1       1x1G 1x10G 1x25G
  18     1       1x1G 1x10G 1x25G
  19     1       1x1G 1x10G 1x25G
{master}
```

show chassis pic fpc-slot 0 pic-slot 0 (ACX7509-FPC-16C)

```
user@router> show chassis pic fpc-slot 0 pic-slot 0
Port speed information:
  Port  PFE      Capable Port Speeds
  0      0      1x40G 1x100G 4x25G 4x10G
  1      0      1x40G 1x100G 4x25G 4x10G
  4      0      1x40G 1x100G 4x25G 4x10G
  5      0      1x40G 1x100G 4x25G 4X10G
  8      0      1x40G 1x100G 4x25G 4x10G
  9      0      1x40G 1x100G 4x25G 4x10G
```

```
 12      0      1x40G 1x100G 4x25G 4x10G
 13      0      1x40G 1x100G 4x25G 4x10G
```

show chassis pic fpc-slot 0 pic-slot 0 (ACX7509-FPC-4CD)

```
user@router> show chassis pic fpc-slot 0 pic-slot 0
Port speed information:
  Port  PFE      Capable Port Speeds
  0      1       1x400G 4x100G 8x50G
  1      1       1x400G 4x100G 8x50G
  2      1       1x400G 4x100G 8x50G
  3      1       1x400G 4x100G 8x50G
```

show chassis pic fpc-slot pic-slot (MX304)

```
user@router>show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                            MRATE LMIC 16x100G/4x400G
  State                           Online
  PIC version                     0.0
  Uptime                          2 hours, 39 minutes, 16 seconds

PIC port information:
                          Fiber                    Xcvr vendor       Wave-
Xcvr          JNPR     MSA
  Port Cable type         type  Xcvr vendor        part number       length
Firmware      Rev      Version
  0    100GBASE LR4       SM    JUNIPER-1F3        1F3QAA            1302 nm
0.0           REV 01   SFF-8636 ver 2.7
  1    100GBASE SR4       MM    JUNIPER-1F1        1F1Q1A            850 nm
0.0           REV 01   SFF-8636 ver 2.7
  2    100G AOC 3M        MM    JUNIPER-INNO       TF-FC003-NJC      850 nm
0.0           REV 01   SFF-8636 ver 2.7
  3    100G CWDM4         SM    JUNIPER-1F2        1F2Q5A            1301 nm
0.0           REV 01   SFF-8636 ver 2.7
  4    100GBASE LR4       SM    JUNIPER-1F3        1F3QAA            1302 nm
0.0           REV 01   SFF-8636 ver 2.7
  5    100GBASE DR        SM    JUNIPER-2J1        2J1Q8A            1311 nm
0.0           REV 01   SFF-8636 ver 2.10
  6    100GBASE DR        SM    JUNIPER-2J1        2J1Q8A            1311 nm
```

```
0.0          REV 01   SFF-8636 ver 2.10
  7    100G FR          SM    JUNIPER-1W2        1W2Q9A            1311 nm
0.0          REV 01   SFF-8636 ver 2.10
  8    100G PSM4        SM    JUNIPER-1F1        1F1Q3A            1310 nm
0.0          REV 02   SFF-8636 ver 2.7
  9    100G FR          SM    JUNIPER-1W2        1W2Q9A            1311 nm
0.0          REV 01   SFF-8636 ver 2.10
  10   100G AOC 3M      MM    JUNIPER-INNO       TF-FC003-NJC      850 nm
0.0          REV 01   SFF-8636 ver 2.7
  11   100GBASE ER4 Lite SM    JUNIPER-FUJITSU    FIM37801/223      1302 nm
0.0          REV 02   SFF-8636 ver 2.7
  12   100G AOC 3M      MM    JUNIPER-INNO       TF-FC003-NJC      850 nm
0.0          REV 01   SFF-8636 ver 2.7
  13   100GBASE CU 3M   n/a   JUNIPER-AMPHENOL   NDAAFJ-J102       n/a
0.0          REV 01   SFF-8636 ver 2.7
  14   100GBASE LR4     SM    JUNIPER-1G3        1G3QAA            1302 nm
0.0          REV 01   SFF-8636 ver 2.10
  15   100GBASE LR4     SM    JUNIPER-1G3        1G3QAA            1302 nm
0.0          REV 01   SFF-8636 ver 2.10

Port speed information:

  Port  PFE      Capable Port Speeds
  0      1       1x1GE, 1x10GE, 1x25GE, 8x50GE, 100GE, 2x100GE, 4x100GE, 400GE
  1      1       1x1GE, 1x10GE, 1x25GE, 4x1GE, 4x10GE, 4x25GE, 40GE, 100GE
  2      1       1x1GE, 1x10GE, 1x25GE, 100GE
  3      1       1x1GE, 1x10GE, 1x25GE, 4x1GE, 4x10GE, 4x25GE, 40GE, 100GE
  4      1       1x1GE, 1x10GE, 1x25GE, 100GE
  5      1       1x1GE, 1x10GE, 1x25GE, 4x1GE, 4x10GE, 4x25GE, 40GE, 100GE
  6      1       1x1GE, 1x10GE, 1x25GE, 8x50GE, 100GE, 2x100GE, 4x100GE, 400GE
  7      1       1x1GE, 1x10GE, 1x25GE, 4x1GE, 4x10GE, 4x25GE, 40GE, 100GE
  8      0       1x1GE, 1x10GE, 1x25GE, 8x50GE, 100GE, 2x100GE, 4x100GE, 400GE
  9      0       1x1GE, 1x10GE, 1x25GE, 4x1GE, 4x10GE, 4x25GE, 40GE, 100GE
  10     0       1x1GE, 1x10GE, 1x25GE, 100GE
  11     0       1x1GE, 1x10GE, 1x25GE, 4x1GE, 4x10GE, 4x25GE, 40GE, 100GE
  12     0       1x1GE, 1x10GE, 1x25GE, 100GE
  13     0       1x1GE, 1x10GE, 1x25GE, 4x1GE, 4x10GE, 4x25GE, 40GE, 100GE
  14     0       1x1GE, 1x10GE, 1x25GE, 8x50GE, 100GE, 2x100GE, 4x100GE, 400GE
  15     0       1x1GE, 1x10GE, 1x25GE, 4x1GE, 4x10GE, 4x25GE, 40GE, 100GE
```

Release Information

Command introduced before Junos OS Release 7.4.

`transport` option introduced in Junos OS Release 16.1R1 for MX Series Routers.

RELATED DOCUMENTATION

- request chassis pic
- show chassis hardware | 18430
- *100-Gigabit Ethernet Type 4 PIC with CFP Overview*

show chassis pic-mode

IN THIS SECTION

- Syntax | 18628
- Description | 18628
- Options | 18629
- Required Privilege Level | 18629
- Sample Output | 18629
- Release Information | 18629

Syntax

```
show chassis pic-mode
<all-members>
<local>
<member member-id>
```

Description

Display the PIC mode on an EX4500 switch.

The PIC mode must be set to **virtual-chassis** if the EX4500 switch is cabled into a Virtual Chassis using the dedicated Virtual Chassis ports (VCPs) on the Virtual Chassis module.

> **NOTE**: This PIC mode setting is not used on an EX4550 switch.

The default PIC mode setting is **virtual-chassis**.

Use the `request chassis pic-mode` command to configure the PIC mode setting.

Options

none	Display the PIC mode on the member switch where the command is issued.
all-members	(Optional) Display the PIC mode for all members of the Virtual Chassis configuration.
local	(Optional) Display the PIC mode on the member switch where the command is issued.
member ***member-id***	(Optional) Display the PIC mode of the specified member of the Virtual Chassis configuration.

Required Privilege Level

view

Sample Output

show chassis pic-mode

```
user@switch> show chassis pic-mode
fpc0:
--------------------------------------------------------------------------
    Pic Mode: Intraconnect
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

request chassis pic-mode | 14509

Resetting the EX4500 Switch to Run as a Standalone Switch (CLI Procedure)

Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)

show chassis pic (Security)

IN THIS SECTION

- Syntax | 18630
- Description | 18630
- Options | 18630
- Required Privilege Level | 18631
- Output Fields | 18631
- Sample Output | 18632
- Release Information | 18633

Syntax

```
show chassis pic fpc-slot slot-number pic-slot slot-number
```

Description

Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.

Options

fpc-slot *slot-number*	Display information about the FPC in the slot.
pic-slot *slot-number*	Display information about the PIC in this particular FPC slot.

Required Privilege Level

view

Output Fields

Table 312 on page 18631 lists the output fields for the `show chassis pic` command. Output fields are listed in the approximate order in which they appear.

Table 312: show chassis pic Output Fields

Field Name	Field Description
`Type`	PIC type.
`State`	Status of the PIC. State is displayed only when a PIC is in the slot. • `Online`— PIC is online and running. • `Offline`—PIC is powered down.
`PIC version`	PIC hardware version.
`Uptime`	How long the PIC has been online.
`Port Number`	Port number for the PIC.
`Cable Type`	Type of cable connected to the port: `LH`, `LX`, or `SX`.

Table 312: show chassis pic Output Fields *(Continued)*

Field Name	Field Description
PIC Port Information	Port-level information for the PIC. • Port—Port number • Cable type—Type of transceiver installed. • Fiber type—Type of fiber. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal.

Sample Output

show chassis pic fpc-slot pic-slot

```
user@host> show chassis pic fpc-slot 10 pic-slot 0
FPC slot 10, PIC slot 0 information:
  Type                            10x 10GE SFP+
  State                           Online
  PIC version                  1.1
  Uptime                          6 days, 7 hours, 29 minutes, 28 seconds

PIC port information:
                           Fiber                    Xcvr vendor
  Port  Cable type         type  Xcvr vendor        part number       Wavelength
  0     10GBASE SR         MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm

        Xcvr vendor
        firmware version
        0.0

PIC port information:
                           Fiber                    Xcvr vendor
  Port  Cable type         type  Xcvr vendor        part number       Wavelength
  1     10GBASE SR         MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm
```

```
Xcvr vendor
firmware version
0.0
```

Release Information

Command introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Interfaces User Guide for Security Devices

Understanding Traffic Processing on Security Devices

show chassis power

IN THIS SECTION

- Syntax | **18634**
- Syntax (MX Series Router) | **18634**
- Syntax (MX320, MX2010, MX2020, MX10003, MX10004, MX2008, PTX Series, ACX7024, and SRX5400 routers; EX9200 and EX9253 Switches) | **18634**
- Syntax (MX10008 Universal Routing Platforms) | **18634**
- Description | **18634**
- Options | **18636**
- Required Privilege Level | **18636**
- Output Fields | **18636**
- Sample Output | **18641**
- show chassis power detail (MX10004) | **18655**
- show chassis power (SRX5800) | **18673**
- show chassis power (ACX Series) | **18674**
- show chassis power detail (ACX Series) | **18675**
- Release Information | **18676**

Syntax

```
show chassis power
```

Syntax (MX Series Router)

```
show chassis power
<all-members>
<local>
<member member-id>
<detail>
```

Syntax (MX320, MX2010, MX2020, MX10003, MX10004, MX2008, PTX Series, ACX7024, and SRX5400 routers; EX9200 and EX9253 Switches)

```
show chassis power
<detail>
```

Syntax (MX10008 Universal Routing Platforms)

```
show chassis power
<detail>
<sequence>
```

Description

(EX9200 switches, MX Series 5G Universal Routing Platforms and PTX Series Packet Transport Routers only) Display power limits and usage information for the AC or DC power sources.

- On EX9200 switches, MX Series, ACX7024, and SRX5800 5G Universal Routing Platforms, power is supplied by Power Entry Modules (PEMs).

> **NOTE**: The new high-capacity (4100 W) enhanced DC PEM on MX960 routers includes a new design that can condition the input voltage. This results in the output voltage

differing from the input voltage. The earlier generation of DC PEMs coupled the input power directly to the output, thereby making it safe to assume that the output voltage was equal to the input voltage.

- On the MX2020 Universal Routing Platforms, the power system consists of three components: the power supply modules (PSMs), the power distribution module (PDM), and the power midplane. The power feed is connected to the PDM. The PDM delivers power to the power midplane. The power midplane supplies power to the PSMs. The MX2020 router chassis provides 3+3 (2500W/80A) or 4+4 (2100W/60A) PSM redundancy for the critical FRUs with two power zones.

- On the MX2010 and MX2008 Universal Routing Platforms, the power system consists of three components: the power supply modules (PSMs), the power distribution module (PDM), and the power midplane. The power feed is connected to the PDM. The PDM delivers power to the power midplane. The power midplane supplies power to the PSMs. Unlike the MX2020 router chassis, the MX2010 and MX2008 router chassis does not provide redundancy for the critical FRUs because there is only one power zone.

  Starting from Junos OS Release 17.3, MX10003 routers support DC (1100W), AC High (1600W), and AC Low (800W) power supply units.

- Starting from Junos OS Release 21.4R1, SRX5800 power system consists of high-capacity second-generation AC power supply module (PSM). This single or dual PSM provides a maximum output power of 5100 W. In the single-feed mode, the PSM provides power at a reduced capacity (2550 W). In dualfeed mode, the PSM provides power at full capacity (5100W). The PSM supports 1+1 redundancy.

- Starting from Junos OS Evolved Release 22.3R1, ACX7024 supports single PSM mode and PSM 1+1 redundancy.

- On the PTX Series Packet Transport Routers, power is supplied by power supply modules (PSMs). On PTX5000 routers, the power feeds connect to the power distribution units (PDUs).

- Starting with Junos OS Release 14.1, the `show chassis power <detail>` operational mode command output displays power usage information for the new DC power supply module (PSM) and power distribution unit (PDU) that are added to provide power to the high-density FPC (FPC2-PTX-P1A) and other components in a PTX5000 Packet Transport Router. The output also displays power usage information for each PIC that is connected to the router. This command also displays power usage information for MX Series routers that have a MPC5EQ MPC installed.

- Starting in Junos OS Release 18.4R1, on MX240 routers and EX9204 switches, the enhanced AC PEM in high-line power configuration supplies a power output of 2400 W. On Junos OS versions prior to 18.4R1, this PEM supplied only 2050 W of power.

- Starting in Junos OS release 21.4R1, on SRX5400 routers the AC redundancy mode for high capacity high line PEMs is increased from 1+1 to 2+2 with a capacity of 4100W from earlier 2050W.

NOTE: The MX204 routers do not support the `show chassis power` command.

Options

none Display basic power usage information for the AC and DC power sources.

all-members (MX Series routers only) (Optional) Display power usage information for all members of the Virtual Chassis configuration.

detail (Optional) Include power usage for specific FRUs.

local (MX Series routers only) (Optional) Display power usage information for the local Virtual Chassis member.

member ***member-id*** (MX Series routers only) (Optional) Display power usage information for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

Required Privilege Level

view

Output Fields

Table 313 on page 18637 lists the output fields for the `show chassis power` command. Output fields are listed in the approximate order in which they appear.

Table 313: show chassis power Output Fields:

Field Name	Field Description	Level of Output

`PEM` *`number`*	(EX9200 switches, SRX5800 and MX Series routers only) AC or DC PEM number on the chassis. The following output fields are displayed for the PEM:
	• `State`—State of the PEM:
	  • `Online`—PEM is present in the slot and online.
	  • `Empty`—PEM is not present in the slot.
	  • `Present`—PEM is present in the slot, but not online.
	• *`AC/DC`* `Input`—`OK` or `Check`—State of the AC or DC input power feed with the number of active and expected feeds (one or two).
	  For a DC input power feed, this output field also displays the reference voltage input with maximum input voltage displayed in mV (in parentheses) for the AC or DC PEM.
	• `Capacity`—Actual power input capacity with maximum capacity displayed (in parentheses) in watts.
	  NOTE: The maximum actual power capacity for AC and DC PEMs is:
	  • MX960 AC PEM—4100 W if two feeds are connected. 1700 W if one feed is connected.
	  • MX960 DC PEM—4100 W if two feeds are connected. 1700 W if one feed is connected.
	  • MX480 AC PEM—2520 W if it is high-line. 1450 W if it is low-line.
	  • MX480 DC PEM—2240 W if the DIP switch is off. 2440 W if the DIP switch is on.
	  • MX240 AC PEM—2240 W if it is high-line. 2440 W if it is low-line.
	  • MX240 DC PEM—2240 W if the DIP switch is off. 2440 W if the DIP switch is on.
	  • EX9204 AC PEM—2050 W if it is high-line. 1167 W if it is low-line.
		All levels

Table 313: show chassis power Output Fields: ***(Continued)***

Field Name	Field Description	Level of Output
	• EX9204 DC PEM—2400 W if the DIP switch is off. 2600 W if the DIP switch is on. • EX9208 AC PEM—2050 W if it is high-line. 1167 W if it is low-line. • EX9208 DC PEM—2400 W if the DIP switch is off. 2600 W if the DIP switch is on. • EX9214 AC PEM—4100 W if two feeds are connected. 1700 W if one feed is connected. • EX9214 DC PEM—4100 W if two feeds are connected. 1700 W if one feed is connected. • `DC Output`—DC power output in Watts for the specified zone, at the specified amps and voltage (*A* @ *V*), and load and percentage utilization of the maximum capacity) for the zone.	

Table 313: show chassis power Output Fields: *(Continued)*

Field Name	Field Description	Level of Output

`System`	(EX9200 switches, MX Series, MX2020 routers, MX2010 routers, and MX2008 routers only) Overall power statistics for the system zone.
	
	The following output fields are displayed for MX Series routers:
	
	• Zone *number*:
	  • `Capacity`—Maximum power capacity applicable for the zone, in watts.
	  • `Allocated power`—Actual capacity allocated for the zone, in watts, with remaining power displayed in parentheses.
	  • `Actual usage`—Actual power usage for the zone, in watts.
	• `Total system capacity`—Cumulative power capacity of all the zones, in watts.
	• `Total remaining capacity`—Difference between the `Total system capacity` and cumulative `Allocated power` of all the zones, in watts.
	
	The following output fields are displayed for MX2010, MX2020, and MX2008 routers:
	
	• `Capacity`—Maximum power capacity applicable for the zone, in watts.
	• `Allocated power`—Actual capacity allocated for the zone, in watts, with remaining power displayed in parentheses.
	• `Actual usage`—Actual power usage for the zone, in watts.
	
	NOTE: For MX2020 routers, there are two power subsystems (Lower Zone and Upper Zone) and the listed output fields are displayed for each zone.
		All levels

`Total Power`	(PTX Series only) Total power used by the switch (displayed in watts).	All levels
`PDU` *`number`*	(PTX5000 only) ID number of the power distribution unit (PDU) on the chassis.	All levels

Table 313: show chassis power Output Fields: *(Continued)*

Field Name	Field Description	Level of Output

`PSM` *`number`*	(PTX Series, MX2020 routers, MX2010 routers, and MX2008 routers only) ID number of the power supply module..
	
	(PTX Series) The following output fields are displayed for each PSM:
	
	• `Input (V)`—Voltage supplied to the PSM.
	
	• `Used (W)`—Actual power usage for the PSM (measured in watts).
	
	NOTE: Starting with Junos OS Release 14.1, no output is displayed for `Input (v)` and `Used (W)` fields for missing PSMs; unlike in earlier releases where 0 was displayed for missing PSMs.
	
	(MX2010, MX2020, and MX2008 routers) The following output fields are displayed for each PSM:
	
	• `State`—State of the PSM:
	  • `Online`—PSM is present in the slot and online.
	  • `Empty`—PSM is not present in the slot.
	  • `Present`—PSM is present in the slot but not online.
	
	• `DC Input`—State of the DC input power feed with the number of active or expected feeds (in parantheses).
	
	• `Capacity`—Actual power input capacity and maximum capacity (in parantheses) displayed in watts.
	  NOTE: The maximum capacity for AC and DC PSMs is:
	  • MX2010/MX2020/MX2008 AC PSM—2500 W.
	  • MX2010/MX2020/MX2008 DC PSM—2100 W if the DIP switch is at 60A settings. 2500 W if the DIP switch is at 80A settings.
	
	• `DC Output`—DC power output in watts for the specified zone at the specified amperes and voltage (A at V), and load and percentage utilization of the maximum capacity for the zone.
		All levels

Table 313: show chassis power Output Fields: *(Continued)*

Field Name	Field Description	Level of Output
`Item`	Actual power usage (measured in watts) for the following FRUs: • `Fan Tray` *n*—Power usage for the specified fan tray. • `RE`*n*`/CB`*n*—Power usage for the specified Routing Engines and Control Boards • `SIB/CCG/FPD`—Power usage for the Switch Interface Board, Centralized Clock Generator (PTX5000 only), and Front Panel Display (craft interface). • `FPC` *n*—Power usage for the FPC in the slot specified. **NOTE**: MX Series routers must have a MPC5EQ MPC installed to view FRU power usage with the `detail` command.	`detail`

Sample Output

show chassis power (MX960 Router with DC PEM)

```
user@host> show chassis power
PEM 0:
  State:     Online
  DC input:  OK (2 feed expected, 2 feed connected)
  DC input:  48.0 V input (57000 mV)
  Capacity:  4100 W (maximum 4100 W)
  DC output: 513 W (zone 0, 9 A at 57 V, 12% of capacity)

PEM 1:
  State:     Online
  DC input:  OK (2 feed expected, 2 feed connected)
  DC input:  48.0 V input (57000 mV)
  Capacity:  4100 W (maximum 4100 W)
  DC output: 228 W (zone 1, 4 A at 57 V, 5% of capacity)

PEM 2:
  State:     Online
  DC input:  OK (2 feed expected, 2 feed connected)
  DC input:  48.0 V input (57000 mV)
```

```
  Capacity:  4100 W (maximum 4100 W)
  DC output: 513 W (zone 0, 9 A at 57 V, 12% of capacity)

PEM 3:
  State:     Online
  DC input:  OK (2 feed expected, 2 feed connected)
  DC input:  48.0 V input (57000 mV)
  Capacity:  4100 W (maximum 4100 W)
  DC output: 342 W (zone 1, 6 A at 57 V, 8% of capacity)

System:
  Zone 0:
      Capacity:          4100 W (maximum 4100 W)
      Allocated power:   1680 W (2420 W remaining)
      Actual usage:      1026 W
  Zone 1:
      Capacity:          4100 W (maximum 4100 W)
      Allocated power:   1263 W (2837 W remaining)
      Actual usage:      570 W
  Total system capacity: 8200 W (maximum 8200 W)
  Total remaining power: 5257 W
```

show chassis power (MX960 Router with AC PEM)

```
user@host> show chassis power
PEM 0:
  State:     Online
  AC input:  OK (2 feed expected, 2 feed connected)
  Capacity:  4100 W (maximum 4100 W)
  DC output: 0 W (zone 0, 0 A at 56 V, 0% of capacity)

PEM 1:
  State:     Present
  AC input:  Check (2 feed expected, 1 feed connected)
  Capacity:  1700 W (maximum 4100 W)

PEM 2:
  State:     Empty
  Input:     Absent

PEM 3:
```

```
  State:     Online
  AC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  1700 W (maximum 1700 W)

System:
  Zone 0:
      Capacity:          4100 W (maximum 4100 W)
      Allocated power:   540 W (3560 W remaining)
      Actual usage:      0 W
  Zone 1:
      Capacity:          0 W (maximum 0 W)
      Allocated power:   0 W (0 W remaining)
      Actual usage:      0 W
  Total system capacity: 4100 W (maximum 4100 W)
  Total remaining power: 3560 W
```

show chassis power detail (MX960 Router with MPC5EQ)

```
user@host> show chassis power detail
PEM 0:
  State:     Online
  DC input:  OK (2 feed expected, 2 feed connected)
  DC input:  48.0 V input (57500 mV)
  Capacity:  4100 W (maximum 4100 W)
  DC output: 798 W (zone 0, 14 A at 57 V, 19% of capacity)

PEM 1:
  State:     Online
  DC input:  OK (2 feed expected, 2 feed connected)
  DC input:  48.0 V input (57500 mV)
  Capacity:  4100 W (maximum 4100 W)
  DC output: 1311 W (zone 1, 23 A at 57 V, 31% of capacity)

PEM 2:
  State:     Online
  DC input:  OK (2 feed expected, 2 feed connected)
  DC input:  48.0 V input (57500 mV)
  Capacity:  4100 W (maximum 4100 W)
  DC output: 855 W (zone 0, 15 A at 57 V, 20% of capacity)

PEM 3:
```

```
  State:     Online
  DC input:  OK (2 feed expected, 2 feed connected)
  DC input:  48.0 V input (57500 mV)
  Capacity:  4100 W (maximum 4100 W)
  DC output: 912 W (zone 1, 16 A at 57 V, 22% of capacity)

System:
  Zone 0:
      Capacity:          4100 W (maximum 4100 W)
      Allocated power:   2497 W (1603 W remaining)
      Actual usage:      1653 W
  Zone 1:
      Capacity:          4100 W (maximum 4100 W)
      Allocated power:   3336 W (764 W remaining)
      Actual usage:      2223 W
  Total system capacity: 8200 W (maximum 8200 W)
  Total remaining power: 2367 W

Item                 Used(W)
  FPC 0                 255
  FPC 10                341
```

show chassis power (MX480 Router with AC PEM)

```
user@host> show chassis power
PEM 0:
  State:     Online
  AC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  2520 W (maximum 2520 W)
  DC output: 472 W (zone 0, 8 A at 59 V, 18% of capacity)

PEM 1:
  State:     Online
  AC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  2520 W (maximum 2520 W)
  DC output: 472 W (zone 0, 8 A at 59 V, 18% of capacity)

PEM 2:
  State:     Online
  AC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  2520 W (maximum 2520 W)
```

```
  DC output: 118 W (zone 0, 2 A at 59 V, 4% of capacity)

PEM 3:
  State:     Empty
  Input:     Absent

System:
  Maximum capacity:   5040 W
  Allocated capacity: 1675 W (33% of maximum)
  Remaining capacity: 3365 W
  Actual usage:       1062 W
```

show chassis power (MX2020 Router)

```
user@host > show chassis power
PSM 0:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 858.44 W (Lower Zone, 16.75 A at 51.25 V, 34.34% of capacity)

PSM 1:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 854.25 W (Lower Zone, 16.75 A at 51.00 V, 34.17% of capacity)

PSM 2:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 858.44 W (Lower Zone, 16.75 A at 51.25 V, 34.34% of capacity)

PSM 3:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 867.00 W (Lower Zone, 17.00 A at 51.00 V, 34.68% of capacity)

PSM 4:
  State:     Online
```

```
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 871.25 W (Lower Zone, 17.00 A at 51.25 V, 34.85% of capacity)

PSM 5:
  State:     Empty
  Input:     Absent

PSM 6:
  State:     Empty
  Input:     Absent

PSM 7:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 867.00 W (Lower Zone, 17.00 A at 51.00 V, 34.68% of capacity)

PSM 8:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 879.75 W (Lower Zone, 17.25 A at 51.00 V, 35.19% of capacity)

PSM 9:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 624.75 W (Upper Zone, 12.25 A at 51.00 V, 29.75% of capacity)

PSM 10:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 615.00 W (Upper Zone, 12.00 A at 51.25 V, 29.29% of capacity)

PSM 11:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 624.75 W (Upper Zone, 12.25 A at 51.00 V, 29.75% of capacity)

PSM 12:
```

```
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 624.75 W (Upper Zone, 12.25 A at 51.00 V, 29.75% of capacity)

PSM 13:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 612.00 W (Upper Zone, 12.00 A at 51.00 V, 29.14% of capacity)

PSM 14:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 627.81 W (Upper Zone, 12.25 A at 51.25 V, 29.90% of capacity)

PSM 15:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 627.81 W (Upper Zone, 12.25 A at 51.25 V, 29.90% of capacity)

PSM 16:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 615.00 W (Upper Zone, 12.00 A at 51.25 V, 29.29% of capacity)

PSM 17:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 624.75 W (Upper Zone, 12.25 A at 51.00 V, 29.75% of capacity)

System:
  Upper Zone:
      Capacity:          18900 W (maximum 22500 W)
      Allocated power:   12900 W (6000 W remaining)
      Actual usage:      5596.62 W
  Lower Zone:
      Capacity:          17500 W (maximum 17500 W)
      Allocated power:   12900 W (4600 W remaining)
```

```
      Actual usage:      6056.12 W
  Total system capacity: 36400 W (maximum 40000 W)
  Total remaining power: 10600 W
```

show chassis power detail (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis power detail
PSM 0:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 540.75 W (Lower Zone, 10.50 A at 51.50 V, 25.75% of capacity)

PSM 1:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 538.12 W (Lower Zone, 10.50 A at 51.25 V, 25.62% of capacity)

PSM 2:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 538.12 W (Lower Zone, 10.50 A at 51.25 V, 25.62% of capacity)

PSM 3:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 540.75 W (Lower Zone, 10.50 A at 51.50 V, 25.75% of capacity)

PSM 4:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 515.00 W (Lower Zone, 10.00 A at 51.50 V, 24.52% of capacity)

PSM 5:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
```

```
  DC output: 527.88 W (Lower Zone, 10.25 A at 51.50 V, 25.14% of capacity)

PSM 6:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 540.75 W (Lower Zone, 10.50 A at 51.50 V, 25.75% of capacity)

PSM 7:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 527.88 W (Lower Zone, 10.25 A at 51.50 V, 25.14% of capacity)

PSM 8:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 530.44 W (Lower Zone, 10.25 A at 51.75 V, 25.26% of capacity)

PSM 9:
  State:     Empty
  Input:     Absent

PSM 10:
  State:     Empty
  Input:     Absent

PSM 11:
  State:     Empty
  Input:     Absent

PSM 12:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 517.50 W (Upper Zone, 10.00 A at 51.75 V, 24.64% of capacity)

PSM 13:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 517.50 W (Upper Zone, 10.00 A at 51.75 V, 24.64% of capacity)
```

```
PSM 14:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 515.00 W (Upper Zone, 10.00 A at 51.50 V, 24.52% of capacity)

PSM 15:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 527.88 W (Upper Zone, 10.25 A at 51.50 V, 25.14% of capacity)

PSM 16:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 517.50 W (Upper Zone, 10.00 A at 51.75 V, 24.64% of capacity)

PSM 17:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2100 W (maximum 2500 W)
  DC output: 515.00 W (Upper Zone, 10.00 A at 51.50 V, 24.52% of capacity)

System:
  Upper Zone:
      Capacity:          12600 W (maximum 15000 W)
      Allocated power:   9436 W (3164 W remaining)
      Actual usage:      3110.38 W
  Lower Zone:
      Capacity:          18900 W (maximum 22500 W)
      Allocated power:   10842 W (8058 W remaining)
Actual usage:      4799.69 W
  Total system capacity: 31500 W (maximum 37500 W)
  Total remaining power: 11222 W

Item                 Used(W)
  FPC 0                    0
  FPC 4                    0
  FPC 9                  719
  FPC 10                 681
  FPC 17                 656
```

```
  FPC 18                  0
  FPC 19                  0
```

show chassis power (MX2020 Router with 240-V high-voltage DC PSMs and PDMs)

```
user@host> show chassis power

PSM 0:
  State:     Online
  DC input:  OK (INP1 feed expected, INP1 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 197.81 W (Lower Zone, 3.75 A at 52.75 V, 7.91% of capacity)

PSM 1:
  State:     Online
  DC input:  OK (INP1 feed expected, INP1 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 171.44 W (Lower Zone, 3.25 A at 52.75 V, 6.86% of capacity)

PSM 2:
  State:     Online
  DC input:  OK (INP1 feed expected, INP1 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 184.62 W (Lower Zone, 3.50 A at 52.75 V, 7.38% of capacity)

PSM 3:
  State:     Online
  DC input:  OK (INP1 feed expected, INP1 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 184.62 W (Lower Zone, 3.50 A at 52.75 V, 7.38% of capacity)

PSM 4:
  State:     Online
  DC input:  OK (INP1 feed expected, INP1 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 183.75 W (Lower Zone, 3.50 A at 52.50 V, 7.35% of capacity)

PSM 5:
  State:     Online
  DC input:  OK (INP1 feed expected, INP1 feed connected)
```

```
  Capacity:  2500 W (maximum 2500 W)
  DC output: 196.88 W (Lower Zone, 3.75 A at 52.50 V, 7.88% of capacity)

PSM 6:
  State:     Online
  DC input:  OK (INP1 feed expected, INP1 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 183.75 W (Lower Zone, 3.50 A at 52.50 V, 7.35% of capacity)

PSM 7:
  State:     Online
  DC input:  OK (INP1 feed expected, INP1 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 196.88 W (Lower Zone, 3.75 A at 52.50 V, 7.88% of capacity)

PSM 8:
  State:     Online
  DC input:  OK (INP1 feed expected, INP1 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 196.88 W (Lower Zone, 3.75 A at 52.50 V, 7.88% of capacity)

PSM 9:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 144.38 W (Upper Zone, 2.75 A at 52.50 V, 5.78% of capacity)

PSM 10:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 145.06 W (Upper Zone, 2.75 A at 52.75 V, 5.80% of capacity)

PSM 11:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 131.88 W (Upper Zone, 2.50 A at 52.75 V, 5.28% of capacity)

PSM 12:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
```

```
  DC output: 143.69 W (Upper Zone, 2.75 A at 52.25 V, 5.75% of capacity)

PSM 13:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 143.69 W (Upper Zone, 2.75 A at 52.25 V, 5.75% of capacity)

PSM 14:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 145.06 W (Upper Zone, 2.75 A at 52.75 V, 5.80% of capacity)

PSM 15:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 144.38 W (Upper Zone, 2.75 A at 52.50 V, 5.78% of capacity)

PSM 16:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 130.62 W (Upper Zone, 2.50 A at 52.25 V, 5.22% of capacity)

PSM 17:
  State:     Online
  DC input:  OK (INP0 feed expected, INP0 feed connected)
  Capacity:  2500 W (maximum 2500 W)
  DC output: 157.50 W (Upper Zone, 3.00 A at 52.50 V, 6.30% of capacity)

System:
  Upper Zone:
      Capacity:          22500 W (maximum 22500 W)
      Allocated power:   6757 W (15743 W remaining)
      Actual usage:      1286.25 W
  Lower Zone:
      Capacity:          22500 W (maximum 22500 W)
      Allocated power:   7240 W (15260 W remaining)
      Actual usage:      1696.62 W
```

```
  Total system capacity: 45000 W (maximum 45000 W)
  Total remaining power: 31003 W
```

show chassis power (MX10004 and MX10008)

```
user@host> show chassis power
PEM 0:
  State:     Online
  Capacity:  2700 W (maximum 2700 W)
  AC input:  OK (Both feed expected, Both feed connected)
  DC output: 1248 W (zone 0, 104 A at 12 V, 46% of capacity)

PEM 1:
  State:     Online
  Capacity:  2700 W (maximum 2700 W)
  AC input:  OK (Both feed expected, Both feed connected)
  DC output: 1248 W (zone 0, 104 A at 12 V, 46% of capacity)

PEM 2:
  State:     Online
  Capacity:  2700 W (maximum 2700 W)
  AC input:  OK (Both feed expected, Both feed connected)
  DC output: 1260 W (zone 0, 105 A at 12 V, 46% of capacity)

PEM 3:
  State:     Present
  Capacity:  2700 W (maximum 2700 W)
  AC input:  Not ready

PEM 4:
  State:     Present
  Capacity:  2700 W (maximum 2700 W)
  AC input:  Not ready

PEM 5:
  State:     Present
  Capacity:  2700 W (maximum 2700 W)
  AC input:  Not ready

System:
  Zone 0:
```

```
      Capacity:          8100 W (maximum 8100 W)
      Allocated power:   7105 W (995 W remaining)
      Actual usage:      3756 W
  Total system capacity: 8100 W (maximum 8100 W)
  Total remaining power: 995 W
```

show chassis power detail (MX10004)

```
PEM 0:
  State:     Online
  Capacity:  2800 W (maximum 5500 W)
  AC input:  OK (INP2 feed expected, INP2 feed connected)
  Feed2:     AC input
  DC output: 744 W (zone 0, 62 A at 12 V, 13% of capacity)

PEM 1:
  State:     Online
  Capacity:  2800 W (maximum 5500 W)
  AC input:  OK (INP2 feed expected, INP2 feed connected)
  Feed2:     AC input
  DC output: 756 W (zone 0, 63 A at 12 V, 13% of capacity)

PEM 2:
  State:     Online
  Capacity:  2800 W (maximum 5500 W)
  AC input:  OK (INP2 feed expected, INP2 feed connected)
  Feed2:     AC input
  DC output: 744 W (zone 0, 62 A at 12 V, 13% of capacity)

System:
  Zone 0:
      Capacity:          8400 W (maximum 8400 W)
      Allocated power:   6197 W (2203 W remaining)
      Actual usage:      2244 W
  Total system capacity: 8400 W (maximum 8400 W)
  Total remaining power: 2203 W

Item                 Used(W)
  Fan Tray 0            138
  Fan Tray 1            136
  RE0/CB0                86
```

```
  RE1/CB1                 96

Item                  Used(W)          Max(W)
  SFB 0                   94             225
  SFB 1                   91             225
  SFB 2                   91             225
  SFB 3                   93             225
  SFB 4                   90             225
  SFB 5                   94             225

Item                  Used(W)
  FPC 0                  775

Item                  Used(W)          Max(W)
  FPC 3                  522            1770
```

show chassis power (PTX5000 Packet Transport Router)

```
user@host> show chassis power
Chassis Power         Input(V)      Used(W)

Total Power                          4006

  PDU 0                              1986
    PSM 0
      Input 1             54          149
    PSM 1
      Input 1             54          377
    PSM 2
      Input 1             54          745
    PSM 3
      Input 1             54          715

  PDU 1                              2020
    PSM 0
      Input 1             54          246
    PSM 1
      Input 1             54          332
    PSM 2
      Input 1             54          721
```

```
    PSM 3
      Input 1            54          721
```

show chassis power (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis power

Chassis Power        Input(V)      Used(W)

Total Power                         4402

  PDU 0                             2104
    PSM 0
      Input 1            53          229
      Input 2            53          375
    PSM 1
    PSM 2
      Input 1            53          248
      Input 2            53          323
    PSM 3
    PSM 4
      Input 1            53          206
      Input 2            53          255
    PSM 5
    PSM 6
      Input 1            53          206
      Input 2            53          262
    PSM 7

  PDU 1                             2298
    PSM 0
    PSM 1
      Input 1            53          289
      Input 2            53          267
    PSM 2
    PSM 3
      Input 1            53          309
      Input 2            53          315
    PSM 4
    PSM 5
      Input 1            53          335
```

```
      Input 2            53          220
    PSM 6
    PSM 7
      Input 1            53          308
      Input 2            53          255
```

show chassis power detail (PTX5000 Packet Transport Router)

```
user@host> show chassis power detail
Chassis Power        Input(V)      Used(W)

Total Power                         3997

  PDU 0                             1975
    PSM 0
      Input 1            54          136
    PSM 1
      Input 1            54          377
    PSM 2
      Input 1            54          741
    PSM 3
      Input 1            54          721

  PDU 1                             2022
    PSM 0
      Input 1            54          235
    PSM 1
      Input 1            54          332
    PSM 2
      Input 1            54          726
    PSM 3
      Input 1            54          729

Item                 Used(W)
  Fan Tray 0             49
  Fan Tray 1            127
  Fan Tray 2            117
  RE0/CB0               109
  RE1/CB1               100
  SIB/CCG/FPD           375
  FPC 0                 381
```

```
  FPC 1                    0
  FPC 2                  447
  FPC 3                  560
  FPC 4                    0
  FPC 5                  448
  FPC 6                  379
  FPC 7                  388
```

show chassis power detail (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis power detail
Chassis Power        Input(V)      Used(W)

Total Power                         4394

  PDU 0                             2095
    PSM 0
      Input 1            53          222
      Input 2            53          368
    PSM 1
    PSM 2
      Input 1            53          248
      Input 2            53          329
    PSM 3
    PSM 4
      Input 1            53          212
      Input 2            53          248
    PSM 5
    PSM 6
      Input 1            53          206
      Input 2            53          262
    PSM 7

  PDU 1                             2299
    PSM 0
    PSM 1
      Input 1            53          296
      Input 2            53          260
    PSM 2
    PSM 3
      Input 1            53          309
```

```
      Input 2            53          315
    PSM 4
    PSM 5
      Input 1            53          342
      Input 2            53          214
    PSM 6
    PSM 7
      Input 1            53          308
      Input 2            53          255

Item                 Used(W)
  Fan Trays             252
  RE0/CB0                93
  RE1/CB1                92
  SIB                   360
  FPC 0                 369
      PIC 0              16
      PIC 1               0
  FPC 1                   0
  FPC 2                 437
      PIC 0              44
      PIC 1              38
  FPC 3                 740
      PIC 0              41
      PIC 1              46
  FPC 4                 732
      PIC 0              74
      PIC 1              37
  FPC 5                   0
  FPC 6                   0
  FPC 7                   0
```

show chassis power detail (PTX10008 Router)

```
user@host> show chassis power detail

PEM 0:
  State:     Online
  Capacity:  2700 W (maximum 2700 W)
  AC input:  OK (No feed expected, Both feed connected)
  DC output: 1164 W (zone 0, 97 A at 12 V, 43% of capacity)
```

```
PEM 1:
  State:     Online
  Capacity:  2700 W (maximum 2700 W)
  AC input:  OK (Both feed expected, Both feed connected)
  DC output: 1188 W (zone 0, 99 A at 12 V, 44% of capacity)

PEM 2:
  State:     Online
  Capacity:  2700 W (maximum 2700 W)
  AC input:  OK (No feed expected, Both feed connected)
  DC output: 1188 W (zone 0, 99 A at 12 V, 44% of capacity)

PEM 3:
  State:     Empty
  Input:     Absent

PEM 4:
  State:     Empty
  Input:     Absent

PEM 5:
  State:     Empty
  Input:     Absent

System:
  Zone 0:
      Capacity:          8100 W (maximum 8100 W)
      Allocated power:   7160 W (940 W remaining)
      Actual usage:      3540 W
  Total system capacity: 8100 W (maximum 8100 W)
  Total remaining power: 940 W

Item                 Used(W)
  Fan Tray 0            475
  Fan Tray 1            475
  RE0/CB0                42
  RE1/CB1                46
```

show chassis power detail (PTX10016 Router)

```
user@host> show chassis power detail

PEM 0:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (Both feed expected, Both feed connected)
  DC output: 432 W (zone 0, 36 A at 12 V, 17% of capacity)

PEM 1:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (Both feed expected, Both feed connected)
  DC output: 456 W (zone 0, 38 A at 12 V, 18% of capacity)

PEM 2:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (Both feed expected, Both feed connected)
  DC output: 432 W (zone 0, 36 A at 12 V, 17% of capacity)

PEM 3:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (Both feed expected, Both feed connected)
  DC output: 432 W (zone 0, 36 A at 12 V, 17% of capacity)

PEM 4:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (INP2 feed expected, INP2 feed connected)
  DC output: 432 W (zone 0, 36 A at 12 V, 17% of capacity)

PEM 5:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (INP2 feed expected, INP2 feed connected)
  DC output: 432 W (zone 0, 36 A at 12 V, 17% of capacity)

PEM 6:
  State:     Online
```

```
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (INP2 feed expected, INP2 feed connected)
  DC output: 420 W (zone 0, 35 A at 12 V, 16% of capacity)

PEM 7:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (INP2 feed expected, INP2 feed connected)
  DC output: 432 W (zone 0, 36 A at 12 V, 17% of capacity)

PEM 8:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (INP2 feed expected, INP2 feed connected)
  DC output: 432 W (zone 0, 36 A at 12 V, 17% of capacity)

PEM 9:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (INP2 feed expected, INP2 feed connected)
  DC output: 420 W (zone 0, 35 A at 12 V, 16% of capacity)

System:
  Zone 0:
      Capacity:          25000 W (maximum 25000 W)
      Allocated power:   9550 W (15450 W remaining)
      Actual usage:      4320 W
  Total system capacity: 25000 W (maximum 25000 W)
  Total remaining power: 15450 W

Item                 Used(W)
  Fan Tray 0            975
  Fan Tray 1            975
  RE0/CB0                42
  RE1/CB1                46
```

show chassis power detail (EX9208 Switch)

```
user@host> show chassis power detail
PEM 0:
  State:     Present
```

```
  AC input:  Out of range (1 feed expected, 1 feed connected)
  Capacity:  2050 W (maximum 2050 W)
  DC output: 0 W (zone 0, 0 A at 0 V, 0% of capacity)

PEM 1:
  State:     Present
  AC input:  Out of range (1 feed expected, 1 feed connected)
  Capacity:  2050 W (maximum 2050 W)
  DC output: 0 W (zone 0, 0 A at 0 V, 0% of capacity)

PEM 2:
  State:     Present
  AC input:  Out of range (1 feed expected, 1 feed connected)
  Capacity:  2050 W (maximum 2050 W)
  DC output: 0 W (zone 0, 0 A at 0 V, 0% of capacity)

PEM 3:
  State:     Online
  AC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  2050 W (maximum 2050 W)
  DC output: 754 W (zone 0, 13 A at 58 V, 36% of capacity)

System:
  Zone 0:
      Capacity:          2050 W (maximum 2050 W)
      Allocated power:   1478 W (572 W remaining)
      Actual usage:      754 W
  Total system capacity: 2050 W (maximum 2050 W)
  Total remaining power: 572 W
```

show chassis power detail (EX9253 Switch)

```
user@switch> show chassis power detail
PEM 0:
  State:     Online
  Capacity:  1600 W (maximum 1600 W)
  AC input:  OK (1 feed expected, 1 feed connected)
  DC output: 564 W (zone 0, 47 A at 12 V, 35% of capacity)

PEM 1:
  State:     Present
```

```
  Input:     Absent

PEM 2:
  State:     Empty
  Input:     Absent

PEM 3:
  State:     Empty
  Input:     Absent

PEM 4:
  State:     Present
  Input:     Absent

PEM 5:
  State:     Online
  Capacity:  1600 W (maximum 1600 W)
  AC input:  OK (1 feed expected, 1 feed connected)
  DC output: 612 W (zone 0, 51 A at 12 V, 38% of capacity)

System:
  Zone 0:
      Capacity:          3200 W (maximum 3200 W)
      Allocated power:   2157 W (1043 W remaining)
      Actual usage:      1176 W
  Total system capacity: 3200 W (maximum 3200 W)
  Total remaining power: 1043 W

Item                 Used(W)
  FPC 0                 555
  FPC 1                 543
  Fan Tray 0             12
  Fan Tray 1             13
  Fan Tray 2             12
  Fan Tray 3             12
  RE0/CB0                55
```

show chassis power (PTX10008 with JNP10K-PWR-DC2 Power Supply Unit - Junos OS Evolved)

```
user@router> show chassis power
Chassis Power        Voltage(V)    Power(W)

Total Input Power                  10176
  PSM 0
    State: Online
    PS0                  51          912
         A0              52
         B0              52
    PS1                  51          936
         A1              52
         B1              52
    Output            12.29        1752.57   (80A input select)
  PSM 1
    State: Online
    PS0                  51          912
         A0              52
         B0              52
    PS1                  51          936
         A1              52
         B1              52
    Output            12.25         1747     (80A input select)
  PSM 2
    State: Online
    PS0                  51          888
         A0              52
         B0              52
    PS1                  51          912
         A1              52
         B1              52
    Output            12.51        1766.71   (80A input select)
  PSM 3
    State: Online
    PS0                  51          936
         B0              50
    PS1                  51          984
         A1              51
         B1              51
    Output            12.49        1807.91   (80A input select)
```

```
  PSM 4
    State: Online
    PS0                  51          864
         A0              51
         B0              52
    PS1                  51          936
         A1              51
         B1              52
    Output            12.45        1763.9    (80A input select)
  PSM 5
    State: Fault
    PS0                  51          960
         A0              52
         B0              52
    PS1                  52            0
         A1              52
         B1              52
    Output            12.43        913.38    (80A input select)

System:
  Zone 0:
      Capacity:          27500 W (maximum 27500 W)
      Allocated power:   19840 W (7660 W remaining)
      Actual usage:      9216 W
  Total system capacity: 27500 W (maximum 27500 W)
  Total remaining power: 7660 W
```

show chassis power detail (PTX10008 with JNP10K-PWR-DC2 Power Supply Unit - Junos OS Evolved)

```
user@router> show chassis power detail
Chassis Power        Voltage(V)    Power(W)

Total Input Power                  10176
  PSM 0
    State: Online
    PS0                  52          912
         A0              52
         B0              52
    PS1                  52          936
         A1              52
```

```
        B1              52
  Output            12.29        1741.75   (80A input select)
  Capacity           5500 W (maximum 5500 W)
PSM 1
  State: Online
  PS0                 51          912
        A0              52
        B0              52
  PS1                 51          936
        A1              52
        B1              52
  Output            12.25         1747     (80A input select)
  Capacity           5500 W (maximum 5500 W)
PSM 2
  State: Online
  PS0                 52          888
        A0              52
        B0              52
  PS1                 51          912
        A1              52
        B1              52
  Output            12.51        1766.71   (80A input select)
  Capacity           5500 W (maximum 5500 W)
PSM 3
  State: Online
  PS0                 51          936
        B0              50
  PS1                 51          984
        A1              52
        B1              52
  Output            12.49        1785.93   (80A input select)
  Capacity           5500 W (maximum 5500 W)
PSM 4
  State: Online
  PS0                 51          864
        A0              52
        B0              52
  PS1                 51          936
        A1              52
        B1              52
  Output            12.47        1766.67   (80A input select)
  Capacity           5500 W (maximum 5500 W)
PSM 5
```

```
    State: Fault
    PS0                  52          960
         A0              52
         B0              52
    PS1                  52            0
         A1              52
         B1              52
    Output            12.43        896.97    (80A input select)
    Capacity              0 W (maximum 0 W)

Item                 Used(W)
  CB 0                   80
  CB 1                   73
  FPC 0                1116
  FPC 1                1153
  FPC 2                 457
  FPC 3                 476
  FPC 4                 468
  FPC 5                 446
  FPC 6                1100
  FPC 7                 477
  SIB 0                 392
  SIB 1                 374
  SIB 2                 395
  SIB 3                 385
  SIB 4                 380
  SIB 5                 389
  Fan Tray 0            629
  Fan Tray 1            609

System:
  Zone 0:
      Capacity:          27500 W (maximum 27500 W)
      Allocated power:   19840 W (7660 W remaining)
      Actual usage:      9264 W
  Total system capacity: 27500 W (maximum 27500 W)
  Total remaining power: 7660 W
```

show chassis power (PTX10008 with JNP10K-PWR-AC2 Power Supply - Junos OS Evolved)

```
user@router> show chassis power
Chassis Power        Voltage(V)    Power(W)

Total Input Power                   7248
  PSM 0
    State: Online
    INP1                203          600
    INP2                203          600
    Output            12.25        1105.35   (30A input select)
  PSM 1
    State: Online
    INP1                203          600
    INP2                201          600
    Output            12.29        1092.65   (30A input select)
  PSM 2
    State: Online
    INP1                205          600
    INP2                201          600
    Output            12.27        1117.92   (30A input select)
  PSM 3
    State: Online
    INP1                205          624
    INP2                201          624
    Output            12.27        1155.72   (30A input select)
  PSM 4
    State: Online
    INP1                203          600
    INP2                203          600
    Output            12.25        1094.57   (30A input select)
  PSM 5
    State: Online
    INP1                201          600
    INP2                201          600
    Output            12.25        1094.57   (30A input select)

System:
  Zone 0:
      Capacity:          33000 W (maximum 33000 W)
      Allocated power:   18810 W (14190 W remaining)
      Actual usage:      7248 W
```

```
  Total system capacity: 33000 W (maximum 33000 W)
  Total remaining power: 14190 W
```

show chassis power detail (PTX10008 with JNP10K-PWR-AC2 Power Supply - Junos OS Evolved)

```
user@router> show chassis power detail
Chassis Power        Voltage(V)    Power(W)

Total Input Power                   7248
  PSM 0
    State: Online
    INP1                201          600
    INP2                205          600
    Output            12.25        1094.57   (30A input select)
    Capacity           5500 W (maximum 5500 W)
  PSM 1
    State: Online
    INP1                203          600
    INP2                201          600
    Output            12.29        1087.24   (30A input select)
    Capacity           5500 W (maximum 5500 W)
  PSM 2
    State: Online
    INP1                201          600
    INP2                201          600
    Output            12.27        1112.52   (30A input select)
    Capacity           5500 W (maximum 5500 W)
  PSM 3
    State: Online
    INP1                203          624
    INP2                203          624
    Output            12.27        1155.72   (30A input select)
    Capacity           5500 W (maximum 5500 W)
  PSM 4
    State: Online
    INP1                203          600
    INP2                201          600
    Output            12.25        1094.57   (30A input select)
    Capacity           5500 W (maximum 5500 W)
  PSM 5
```

```
    State: Online
    INP1                201          600
    INP2                203          600
    Output            12.25        1099.96   (30A input select)
    Capacity           5500 W (maximum 5500 W)

Item                 Used(W)
  CB 0                    86
  CB 1                    75
  FPC 0                  966
  FPC 2                  995
  FPC 3                  793
  FPC 5                  809
  FPC 7                  856
  SIB 0                  353
  SIB 1                  330
  SIB 2                  190
  SIB 3                  333
  SIB 4                    0
  SIB 5                  343
  Fan Tray 0             304
  Fan Tray 1             305

System:
  Zone 0:
      Capacity:          33000 W (maximum 33000 W)
      Allocated power:   18810 W (14190 W remaining)
      Actual usage:      7248 W
  Total system capacity: 33000 W (maximum 33000 W)
  Total remaining power: 14190 W
```

show chassis power (SRX5400)

```
user@router> show chassis power
PEM 0:
  State:     Online
  AC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  2050 W (maximum 2050 W)
  DC output: 342 W (zone 0, 6 A at 57 V, 16% of capacity)

PEM 1:
```

```
  State:     Online
  AC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  2050 W (maximum 2050 W)
  DC output: 232 W (zone 0, 4 A at 58 V, 11% of capacity)

PEM 2:
  State:     Online
  AC input:  OK (1 feed expected, 1 feed connected)
Capacity:  2050 W (maximum 2050 W)
  DC output: 228 W (zone 0, 4 A at 57 V, 11% of capacity)

PEM 3:
  State:     Online
  AC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  2050 W (maximum 2050 W)
  DC output: 171 W (zone 0, 3 A at 57 V, 8% of capacity)

System:
  Zone 0:
      Capacity:          4100 W (maximum 4100 W)
      Allocated power:   1442 W (2658 W remaining)
      Actual usage:      973 W
  Total system capacity: 4100 W (maximum 4100 W)
  Total remaining power: 2658 W
```

show chassis power (SRX5800)

```
user@root> show chassis power

PEM 0:
  State:     Online
  AC input:  Check (2 feed expected, 2 feed connected)
  Capacity:  5100 W (maximum 5100 W)
  DC output: 114 W (zone 0, 2 A at 57 V, 2% of capacity)

PEM 1:
  State:     Online
  AC input:  Check (2 feed expected, 2 feed connected)
  Capacity:  5100 W (maximum 5100 W)
  DC output: 1197 W (zone 1, 21 A at 57 V, 29% of capacity)
```

```
PEM 2:
  State:     Online
  AC input:  OK (2 feed expected, 2 feed connected)
  Capacity:  5100 W (maximum 5100 W)
  DC output: 171 W (zone 0, 3 A at 57 V, 4% of capacity)

PEM 3:
  State:     Empty
  Input:     Absent

System:
  Zone 0:
      Capacity:          5100 W (maximum 5100 W)
      Allocated power:   595 W (4505 W remaining)
      Actual usage:      285 W
  Zone 1:
      Capacity:          5100 W (maximum 5100 W)
      Allocated power:   1815 W (3285 W remaining)
      Actual usage:      1197 W
  Total system capacity: 10200 W (maximum 10200 W)
  Total remaining power: 7790 W
```

show chassis power (ACX Series)

The command output is as follows when the default PSU 1+1 mode is active.

```
user@root> show chassis power
Chassis Power        Voltage(V)    Power(W)
Total Input Power                     109
  PSM 0
    State: Online
    Input               229           61
    Output            12.15         51.75
    Capacity           400 W (maximum 400 W)
  PSM 1
    State: Online
    Input               230           48
    Output            12.17        38.79
    Capacity            400 W (maximum 400 W)

System:
```

```
  Zone 0:
      Capacity:          800 W (maximum 800 W)
      Actual usage:      109 W
  Total system capacity: 800 W (maximum 800 W)
```

show chassis power detail (ACX Series)

```
user@host> show chassis power detail

Chassis Power        Voltage(V)    Power(W)

Total Input Power                    454
  PSM 0
    State: Online
    Input                53          454
    Output            12.07        421.97
    Capacity           2200 W (maximum 2200 W)
  PSM 1
    State: Offline
    Input                 0            0
    Output                0            0
    Capacity              0 W (maximum 0 W)

Item                 Used(W)
  Routing Engine 0       35
  Routing Engine 1       34
  CB 0                   18
  FPC 1                  48
  FPC 2                  39
  FPC 3                  24
  Fan Tray 0              5
  Fan Tray 1              6
  Fan Tray 2              4
  Fan Tray 3              6
  FEB 0                 200

System:
  Zone 0:
      Capacity:          2200 W (maximum 2200 W)
      Allocated power:   1815 W (385 W remaining)
      Actual usage:      454 W
```

```
Total system capacity: 2200 W (maximum 2200 W)
Total remaining power: 385 W
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- show chassis power sequence | 18688
- Checklist for Monitoring Power Supplies

show chassis power (Security)

IN THIS SECTION

- Syntax | 18676
- Description | 18676
- Options | 18677
- Required Privilege Level | 18677
- Output Fields | 18677
- Sample Output | 18678
- Release Information | 18681

Syntax

```
show chassis power
```

Description

Display power limits and usage information for the Power Entry Modules (PEMs).

Options

node Displays information specific to the chassis.

sequence Shows the chassis fru power on sequence.

Required Privilege Level

view

Output Fields

Table 314 on page 18677 lists the output fields for the `show chassis power` command. Output fields are listed in the approximate order in which they appear.

Table 314: show chassis power Output Fields

Field Name	Field Description

`PEM number`	AC or DC PEM number on the chassis. The following output fields are displayed for the PEM:
	• State—State of the PEM:
	  • Online—PEM is present in the slot and online.
	  • Empty—PEM is not present in the slot.
	  • Present—PEM is present in the slot, but not online.
	• AC Input —State of the AC input power feed with the number of active and expected feeds (1 or 2).
	• Capacity—Actual power input capacity with maximum capacity displayed (in parentheses) in watts.
	• DC Output—DC power output, in watts, for the specified zone, at the specified amps and voltage (A @ V), and load and percentage utilization of the maximum capacity for the zone.

Table 314: show chassis power Output Fields *(Continued)*

Field Name	Field Description

`System`	Overall power statistics for the system zone:
	• Zone number:
	  • Capacity—Maximum power capacity available for the zone, in watts.
	  • Allocated power—Actual capacity allocated for the zone, in watts, with remaining power displayed in parentheses.
	  • Actual usage—Actual power usage for the zone, in watts.
	• Total system capacity—Cumulative power capacity of all the zones, in watts.
	• Total remaining capacity—Difference between the total system capacity and cumulative allocated power of all zones, in watts.

Sample Output

show chassis power

When you enter the **show chassis power** command, the sample output is shown for DC PEM.

```
user@host> show chassis power
node0:
--------------------------------------------------------------------------

PEM 0:
  State:     Online
  DC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  2440 W (maximum 2440 W)
  DC output: 600 W (zone 0, 12 A at 50 V, 24% of capacity)

PEM 1:
  State:     Online
  DC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  2440 W (maximum 2440 W)
  DC output: 1550 W (zone 1, 31 A at 50 V, 63% of capacity)
```

```
PEM 2:
  State:     Empty
  Input:     Absent

PEM 3:
  State:     Empty
  Input:     Absent

System:
  Zone 0:
      Capacity:          2440 W (maximum 2440 W)
      Allocated power:   1050 W (1390 W remaining)
      Actual usage:      600 W
  Zone 1:
      Capacity:          2440 W (maximum 2440 W)
      Allocated power:   2310 W (130 W remaining)
      Actual usage:      1550 W
  Total system capacity: 4880 W (maximum 4880 W)
  Total remaining power: 1520 W

node1:
--------------------------------------------------------------------------

PEM 0:
  State:     Online
  DC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  2440 W (maximum 2440 W)
  DC output: 600 W (zone 0, 12 A at 50 V, 24% of capacity)

PEM 1:
  State:     Online
  DC input:  OK (1 feed expected, 1 feed connected)
  Capacity:  2440 W (maximum 2440 W)
  DC output: 1519 W (zone 1, 31 A at 49 V, 62% of capacity)

PEM 2:
  State:     Empty
  Input:     Absent

PEM 3:
  State:     Empty
  Input:     Absent
```

```
System:
  Zone 0:
      Capacity:          2440 W (maximum 2440 W)
      Allocated power:   1050 W (1390 W remaining)
      Actual usage:      600 W
  Zone 1:
      Capacity:          2440 W (maximum 2440 W)
      Allocated power:   2310 W (130 W remaining)
      Actual usage:      1519 W
  Total system capacity: 4880 W (maximum 4880 W)
  Total remaining power: 1520 W
```

show chassis power detail

When you enter the **show chassis power detail** command, the sample output is shown for DC PEM.

```
user@host> show chassis power detail
PEM 0:
  State:     Online
  Capacity:  1600 W (maximum 1600 W)
  DC output: 240 W (zone 0, 20 A at 12 V, 8% of capacity)

PEM 1:
  State:     Online
  Capacity:  1600 W (maximum 1600 W)
  DC output: 204 W (zone 0, 17 A at 12 V, 7% of capacity)

System:
  Zone 0:
      Capacity:          3200 W (maximum 3200 W)
      Allocated power:   1000 W (2200 W remaining)
      Actual usage:      444 W
  Total system capacity: 3200 W (maximum 3200 W)
  Total remaining power: 2200 W

Item                 Used(W)
  Fan Tray 0              0
  Fan Tray 1              0
  Fan Tray 2              0
  Fan Tray 3              0
```

```
Fan Tray 4              0
RE0/CB0                 0
FPC 0                   0
FPC 1                   0
```

Release Information

Command modified in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show chassis power sequence

show chassis power-budget-statistics

IN THIS SECTION

- Syntax | 18681
- Description | 18681
- Required Privilege Level | 18682
- Output Fields | 18682
- Sample Output | 18685
- Release Information | 18687

Syntax

```
show chassis power-budget-statistics
```

Description

Display the power budget of the device.

Required Privilege Level

view

Output Fields

Table 315 on page 18682 lists the output fields for the `show chassis power-budget-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 315: show chassis power-budget-statistics Output Fields

Field Name	Field Description
`PSU` *`n`* `(`*`supply type`*`)`	Capacity rating of the power supply and whether the power supply is currently operating (`Online`) or not (`Offline`). If a power supply is offline, the capacity is shown as 0 W.
`Total Power supplied by all Online PSUs`	Total number of watts supplied by all currently operating power supplies.
`Power Redundancy Configuration`	Configured power redundancy setting, either *N*+1 or *N* +*N*.
`Base power reserved`	Total number of watts reserved for the device.
`Non-PoE power being consumed`	The amount of power, in W, currently being consumed for PoE.

Table 315: show chassis power-budget-statistics Output Fields *(Continued)*

Field Name	Field Description
`Power Reserved for the Chassis`	Power reserved for the chassis: • For an EX6200 switch, 500 W. • For an EX8208 switch: 1600 W in an *N*+1 configuration; 1200 W in an *N*+*N* configuration • For an EX8216 switch: 2400 W in an *N*+1 configuration; 1800 W in an *N*+*N* configuration The power reserved for the chassis includes the maximum power requirements for the fan tray and Switch Fabric and Routing Engine (SRE), Routing Engine (RE), and Switch Fabric (SF) modules in both base and redundant configurations.
`Fan Tray Statistics`	(EX6200 switch only) Information about the fan tray: • `Base power`—Power allocated to the fan tray in the power budget. This allocation is included in `Power Reserved for the Chassis`. • `Power Used`—Actual power being used by the fan tray. This value is for informational purposes only: the power budget for the switch is based on allocated power (the theoretical maximum the fan tray might use) rather than used power.

Table 315: show chassis power-budget-statistics Output Fields *(Continued)*

Field Name	Field Description
`FPC` *`n`* `(`*`card type`*	Information about the line card installed in slot *n*. For EX6200 switches, information about the SRE modules in slot 4 and slot 5 is also shown. • `Base power`—For line cards without PoE ports, the total power allocated to the line card. For line cards with PoE ports, the power allocated to the line card before the PoE power budget is allocated. The base power includes 37 W of PoE power that is always allocated to line cards that support PoE. • `Power Used`—(EX6200 switch only) The actual power being consumed by the line card or SRE module, including PoE power. This value is for informational purposes only: the power budget for the switch is based on allocated power (the theoretical maximum the line card might use) rather than used power. • `PoE power`—For line cards with PoE ports, the PoE power budget allocated to the line card. This value includes the 37 W of PoE power that is always part of the base power allocation for line cards that support PoE. For line cards without PoE ports, the value is always 0 W. • The power priority assigned to the line card slot.
`Total (non-PoE) Power allocated`	Power budgeted for all the components in the switch, excluding the PoE power budget allocated to line cards. This value is equal to the power reserved for the chassis plus the base power allocations of all online line cards.
`Total Power allocated for PoE`	The total of the PoE power budgets allocated to the line cards in the switch. This figure includes the 37 W of PoE power always included in the base allocation for each line card that supports PoE.

Table 315: show chassis power-budget-statistics Output Fields *(Continued)*

Field Name	Field Description
`Total PoE power reserved`	The amount of power that has been reserved by PoE.
`Total PoE power remaining`	The amount of available power remaining that can be used for PoE.
`Power Available (Redundant case)`	Unused power available to the switch in the power budget, not including the power reserved for redundancy. If power is insufficient to meet the *N*+1 or *N*+*N* redundancy requirements, this value is 0. PoE power allocations are not included in the calculation of this value.
`Total Power Available`	Unused power available to the switch in the power budget. This value is derived by subtracting all power allocations, including PoE power allocations, from the total power available on the switch (the `Total Power supplied by all Online PSUs` value).

Sample Output

show chassis power-budget-statistics (EX6200 Switch)

```
user@switch> show chassis power-budget-statistics
        PSU  0     (EX6200-PWR-AC2500)                :    2500 W   Online
        PSU  1     (EX6200-PWR-AC2500)                :    2500 W   Online
        PSU  2     (EX6200-PWR-AC2500)                :    2500 W   Online
        PSU  3     (EX6200-PWR-AC2500)                :    2500 W   Online
        Total Power supplied by all Online PSUs  :   10000 W
        Power Redundancy Configuration           :     N+1
        Power Reserved for the Chassis           :     500 W
  Fan Tray Statistics            Base power    Power Used
  FTC  0                         :     300 W      43.04 W
  FPC Statistics                  Base power   Power Used   PoE power   Priority
  FPC  1   (EX6200-48P)          :     220 W      49.47 W      1440 W       1
  FPC  2   (EX6200-48P)          :     220 W      47.20 W       800 W       2
```

```
  FPC  3   (EX6200-48P)          :     220 W    1493.57 W      1440 W       0
  FPC  4   (EX6200-SRE64-4XS)    :     100 W      51.38 W         0 W       0
  FPC  5   (EX6200-SRE64-4XS)    :     100 W      50.28 W         0 W       0
  FPC  6   (EX6200-48P)          :     220 W      49.38 W       800 W       6
  FPC  8   (EX6200-48P)          :     220 W      61.41 W      1440 W       9
  FPC  9   (EX6200-48T)          :     150 W      12.49 W         0 W       9

        Total (non-PoE) Power allocated          :    1750 W
        Total Power allocated for PoE            :    5920 W
        Power Available (Redundant case)         :    5750 W
        Total Power Available                    :    2515 W
```

show chassis power-budget-statistics (EX8200 Switch)

```
user@switch> show chassis power-budget-statistics
    PSU  0     (EX8200-AC2K)                 :    2000 W  Online
    PSU  1     (EX8200-AC2K)                 :    2000 W  Online
    PSU  2     (EX8200-AC2K)                 :    2000 W  Online
    PSU  3     (EX8200-AC2K)                 :    2000 W  online
    PSU  4     (EX8200-AC2K)                 :    2000 W  Online
    Total Power supplied by all Online PSUs  :   10000 W
    Power Redundancy Configuration           :     N+1
    Power Reserved for the Chassis           :    2400 W
    FPC Statistics                            Base power    PoE power   Priority
    FPC  1   (EX8200-48T)                    :     350 W          0 W     15
    FPC  5   (EX8200-2XS-40P)                :     387 W        792 W      0
    FPC  9   (EX8200-48PL)                   :     267 W        915 W     15
    FPC  10  (EX8200-2XS-40T)                :     350 W          0 W      1
    FPC  12  (EX8200-48T)                    :     350 W          0 W     15

    Total (non-PoE) Power allocated          :    4104 W
    Total Power allocated for PoE            :    1707 W
    Power Available (Redundant case)         :    3896 W
    Total Power Available                    :    4263 W
```

show chassis power-budget-statistics (SRX380 device)

```
user@host> show chassis power-budget-statistics
            PSU 0    (JPSU-600W-AC-AFO   )           :     600 W   Online
            PSU 1    (JPSU-600W-AC-AFO   )           :     600 W   Online
```

```
            Power redundancy configuration           :     N+N
            Total power supplied by all online PSUs  :     600 W
            Base power reserved                      :     300 W
            Non-PoE power being consumed             :     300 W
            Total power allocated for PoE            :     300 W
            Total PoE power consumed                 :     289 W
            Total PoE power remaining                :      11 W
```

show chassis power-budget-statistics (EX4400-48XP device)

```
user@host> show chassis power-budget-statistics
            PSU 0    (JPSU-2000-C-AC-AFO )           :    1950 W   Online
            PSU 1                        )           :       0 W   Online
            Power redundancy configuration           :     N+0
            Total power supplied by all online PSUs  :    1950 W
            Base power reserved                      :     300 W
            Non-PoE power being consumed             :      85 W
            Total power allocated for PoE            :    1650 W
            Total PoE power reserved                 :      43 W
            Total PoE power remaining                :    1607 W
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Verifying Power Configuration and Use

Configuring the Power Priority of Line Cards (CLI Procedure)

Configuring Power Supply Redundancy (CLI Procedure)

show chassis power sequence

IN THIS SECTION

- Syntax | 18688
- Description | 18688
- Options | 18688
- Required Privilege Level | 18688
- Output Fields | 18689
- Sample Output | 18689
- Release Information | 18690

Syntax

```
show chassis power sequence
```

Description

(MX Series 5G Universal Routing Platforms only) Show power-on sequence for the chassis Dense Port Concentrators (DPCs).

(PTX Series Packet Transport Routers, MX2010 and MX2020 routers only) Show power-on sequence for FPCs installed in the chassis.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 316 on page 18689 lists the output fields for the `show chassis power sequence` command. Output fields are listed in the approximate order in which they appear.

Table 316: show chassis power sequence Output Fields

Field Name	Field Description
Chassis FRU Power Sequence	(MX Series) Power-on sequence for the DPCs in the chassis. The numbers indicate the slot number of the DPCs. (PTX Series, MX2010 and MX2020 routers only) Power-on sequence for the FPCs in the chassis. The numbers indicate the slot number of the FPC.

Sample Output

show chassis power sequence (MX Series and ACX Series)

```
user@host> show chassis power sequence
Chassis FRU Power Sequence: 3 4 5 6 7 8 9 10 11 0 1 2
```

show chassis power sequence (MX2010 Routers)

```
user@host > show chassis power sequence
Chassis FRU Power On Sequence: 0 1 2 3 4 5 6 7 8 9
```

show chassis power sequence (MX2020 Routers)

```
user@host > show chassis power sequence
Chassis FRU Power On Sequence: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19
```

show chassis power sequence (PTX5000 Packet Transport Router)

```
user@host> show chassis power sequence
Chassis FRU Power On Sequence: 0 1 2 3 4 5 6 7
```

Release Information

Command introduced in Junos OS Release 10.0.

Command introduces in Junos OS Release 12.1 for PTX Series Packet Transport Routers.

RELATED DOCUMENTATION

show chassis power | 18633

show chassis psd

IN THIS SECTION

- Syntax | 18690
- Description | 18691
- Options | 18691
- Additional Information | 18691
- Required Privilege Level | 18691
- Output Fields | 18691
- Sample Output | 18692
- Release Information | 18692

Syntax

```
show chassis psd
```

Description

(Root System Domain [RSD] only) Display information about Protected System Domains (PSDs). A PSD is initially created by the RSD configuration. An RSD and PSDs are supported on a T320 or T640 router, or a T1600 routing node, or a TX Matrix Plus Platform that is interconnected with the JCS1200 platform.

> **NOTE**: RSD configuration is not supported on a routing matrix based on TX Matrix Plus router with 3D SIBs.

Options

This command has no options.

Additional Information

For more information about PSDs, RSDs, and the JCS1200 platform, see the *Junos OS Protected System Domain User Guide for Routing Devices*.

Required Privilege Level

view

Output Fields

Table 317 on page 18691 lists the output fields for the `show chassis psd` command. Output fields are listed in the approximate order in which they appear.

Table 317: show chassis psd Output Fields

Field Name	Field Description
Slot Description	PSD identification.
State	PSD status: • **Online**—PSD is online and running. • **Offline**—PSD is powered down.

Table 317: show chassis psd Output Fields *(Continued)*

Field Name	Field Description
Uptime	Length of time that the PSD has been up and running.

Sample Output

show chassis psd

```
{master}

user@host> show chassis psd
Slot Description         State               Uptime
 1                       Online              12 hours, 19 minutes, 51 seconds
 2                       Online              2 hours, 18 minutes, 17 seconds
 3                       Online              12 hours, 19 minutes, 51 seconds
```

Release Information

Command introduced in Junos OS Release 9.1.

show chassis realtime-ukern-thread

IN THIS SECTION

- Syntax | 18693
- Description | 18693
- Options | 18693
- Required Privilege Level | 18693
- Output Fields | 18693
- Sample Output | 18693
- Release Information | 18693

Syntax

```
show chassis realtime-ukern-thread
```

Description

Display whether real-time Bidirectional Forwarding Detection (BFD) is enabled or disabled. If real-time BFD is enabled, the output of the show command displays the value of the `realtime Ukern thread Status` field as `Enabled`.

Options

This command has no options.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

show chassis realtime-ukern-thread

```
user@host> show chassis realtime-ukern-thread
realtime Ukern thread Status: Enabled
```

Release Information

Command introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Enabling Dedicated and Real-Time BFD

realtime-ukern-thread (BFD) | **14099**

show chassis redundancy feb

IN THIS SECTION

- Syntax | 18694
- Description | 18694
- Options | 18694
- Required Privilege Level | 18695
- Output Fields | 18695
- Sample Output | 18696
- Release Information | 18698

Syntax

```
show chassis redundancy feb
<errors>
<redundancy-group group-name>
```

Description

(M120 routers only) Display information about the status of configured Forwarding Engine Board (FEB) redundancy groups.

Options

none	Display information about the status of all configured FEB redundancy groups.
redundancy-group *group-name*	(Optional) Display information about the specified configured redundancy group.

errors (Optional) Display information about any errors encountered on the components in configured redundancy groups or on links between a FEB and a Flexible PIC Concentrator (FPC).

Required Privilege Level

view

Output Fields

Table 318 on page 18695 lists the output fields for the `show chassis redundancy feb` command. Output fields are listed in the approximate order in which they appear.

Table 318: show chassis redundancy feb Output Fields

Field name	Field Description
Group	Name of configured redundancy group.
FEB	Slot number of each FEB included in redundancy groups.
State	State of each FEB: • **Online**—FEB is online and running. • **Offline**—FEB is powered down.
Priority	(Standard and **redundancy-group** option) Status of FEB in the redundancy group: **Backup**, **Primary**, **Other**, or null.
Connected FPCs	(Standard and **redundancy-group** option) Slot number of each FPC connected to the FEB. The status **Check** is displayed when an error might have occurred.

Table 318: show chassis redundancy feb Output Fields *(Continued)*

Field name	Field Description
Redundancy State	(Standard and **redundancy-group** option) Status of the FEB: • **Active**—FEB is currently active. • **Ready**—Backup FEB is ready for a switchover • **Not Ready**—Backup FEB is not ready for a switchover.
Auto-failover	(Standard and **redundancy-group** option) Automatic failover status of redundancy group: **Enabled** or **Disabled**.
Switch-reason	(Standard and **redundancy-group** option) Reason a switchover occurred to the backup FEB in the redundancy group.
Hard error: Yes	(**errors** option only) Displayed when a hard error occurs on a FEB.
FPC	(**errors** option only) Slot number and status of FPC: **link ok** or **link error**.
Fabric plane	(**errors** option only) Slot number and status of fabric plane.

Sample Output

show chassis redundancy feb

```
user@host> show chassis redundancy feb
Group:        cfpc
  FEB  State            Priority  Connected FPCs  Redundancy state
  0    Offline          Backup                    Not ready
  1    Online                     5               Active
  Auto-failover:  Enabled
 Group:        grp0
  FEB  State            Priority  Connected FPCs  Redundancy state
  3    Offline          Backup                    Not ready
```

```
  5    Online            Primary   0               Active
  Auto-failover:  Enabled
```

show chassis redundancy feb redundancy-group grp1

```
user@host> show chassis redundancy feb redundancy-group grp1
Group:        grp1
  FEB  State             Priority  Connected FPCs  Redundancy state
  0    Online            Other     0               Active
  1    Online            Other     1               Active
  4    Online            Primary   4               Active
  5    Online            Backup                    Ready
  Autofailover:  Enabled
  Switch-reason: Switchover from CLI
```

show chassis redundancy feb redundancy-group grp0 errors

```
user@host> show chassis redundancy feb redundancy-group grp0 errors
Group: grp0
    FEB: 0    State: Online
        FPC 0 link OK
        Fabric plane 0 OK
        Fabric plane 1 OK
        Fabric plane 2 OK
        Fabric plane 3 OK
    FEB: 1    State: Online
        FPC 0 link OK
        Fabric plane 0 OK
        Fabric plane 1 OK
        Fabric plane 2 OK
        Fabric plane 3 OK
    FEB: 2    State: Online
        FPC 2 link OK
        Fabric plane 0 OK
        Fabric plane 1 OK
        Fabric plane 2 OK
        Fabric plane 3 OK
    FEB: 3    State: Online
        FPC 3 link OK
        Fabric plane 0 OK
```

```
        Fabric plane 1 OK
        Fabric plane 2 OK
        Fabric plane 3 OK
    FEB: 4    State: Online
        FPC 4 link OK
        Fabric plane 0 OK
        Fabric plane 1 OK
        Fabric plane 2 OK
        Fabric plane 3 OK
    FEB: 5    State: Online
        FPC 5 link OK
        Fabric plane 0 OK
        Fabric plane 1 OK
        Fabric plane 2 OK
        Fabric plane 3 OK
```

Release Information

Command introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

request chassis redundancy feb slot | **14518**

Configuring FEB Redundancy on the M120 Router

Understanding Switching Control Board Redundancy

show chassis redundant-power-system

IN THIS SECTION

- Syntax | **18699**
- Description | **18699**
- Required Privilege Level | **18699**
- Output Fields | **18699**
- Sample Output | **18700**

Release Information | 18701

Syntax

```
show chassis redundant-power-system
```

Description

Display information about the Redundant Power Systems (RPS) connected to the switch.

Required Privilege Level

view

Output Fields

Table 319 on page 18699 lists the output fields for the `show chassis redundant-power-system` command. Output fields are listed in the approximate order in which they appear.

Table 319: show chassis redundant-power-system Output Fields

Field Name	Field Description	Level of Output
Member	Member number of the switch connected to the RPS—For a switch that has never been configured in a Virtual Chassis, the value is always zero. For a Virtual Chassis member, the range is zero through the maximum number of members in the Virtual Chassis.	All levels

Table 319: show chassis redundant-power-system Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Status	Status of the RPS: • ARMED—The switch is ready to get backup power from the RPS if power supply fails on the switch. • OFF—The switch has zero and is not configured to receive backup power from the RPS. • BACKED-UP—The switch is receiving power backup from the RPS. • OVER-SUBSCRIBED—The switch cannot receive backup power from the RPS even if you set the .	All levels
RPS	Serial number of the RPS.	
Port	Number of the switch connector on the RPS that is connected to a switch.	All levels

Sample Output

show chassis redundant-power-system (Standalone Switch)

```
user@switch> show chassis redundant-power-system

Member Status      RPS            Port
  0    Armed       CG0209121807   0
```

show chassis redundant-power-system (Virtual Chassis member)

```
user@switch> show chassis redundant-power-system

Member Status      RPS            Port
```

```
0     Armed       CG0209121814   5
2     Armed       CG0209121815   4
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Determining and Setting Priority for Switches Connected to an EX Series RPS

show chassis reswatch

IN THIS SECTION

- Syntax | 18701
- Description | 18701
- Required Privilege Level | 18702
- Output Fields | 18702
- Sample Output | 18703
- Release Information | 18703

Syntax

```
show chassis reswatch-config
```

Description

Display resource usage statistics report. By default, reswatch only logs resource usages when the monitored resource entered its orange zone.

Required Privilege Level

Output Fields

Table 320 on page 18702 lists the output fields for the `show chassis reswatch-config` command. Output fields are listed in the approximate order in which they appear.

Table 320: Show chassis reswatch-config

Field Name	Field Description
`Reswatch`	`Enable` or `Disable`.
`CPU statistics count`	CPU statistics count value.
`Junos statistics count`	Junos kernel statistics count value.
`Red Zone`	`Enable` or `Disable`.
`Red Zone CPU threshold`	Red zone CPU threshold value
`Red Zone Junos threshold`	Red zone Junos threshold vlaue
`Red Zone SNMP`	`Enable` or `Disable`.
`Orange Zone`	`Enable` or `Disable`.
`Orange Zone threshold`	Orange zone threshold value.

Sample Output

show chassis reswatch-config

```
user@host> show chassis reswatch-config

  Reswatch:                  Enable
  CPU statistics count:      0
  Junos statistics count:    0
  Red Zone:                  Enable
  Red Zone CPU threshold:    15
  Red Zone Junos threshold: 60
  Red Zone SNMP:             Enable
  Orange Zone:               Enable
  Orange Zone threshold:     30
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

resource-watch | 15622

show chassis routing-engine

IN THIS SECTION

- Syntax | 18704
- Syntax (ACX Series, PTX Series, and MX104 Universal Routing Platforms.) | 18704
- Syntax (EX Series Switches) | 18704
- Syntax (QFX Series) | 18705
- Syntax (MX Series Routers) | 18705
- Syntax (MX204 and MX10003 Universal Routing Platforms) | 18705

- Syntax (TX Matrix Routers) | **18705**
- Syntax (TX Matrix Plus Routers) | **18706**
- Syntax (Junos OS Evolved) | **18706**
- Description | **18706**
- Options | **18706**
- Required Privilege Level | **18707**
- Output Fields | **18708**
- Sample Output | **18712**
- show chassis routing-engine (MX Series Routers) | **18713**
- show chassis routing-engine (QFX5700 Switches) (Junos OS Evolved Release) | **18715**
- show chassis routing-engine (QFX5230-64CD devices with logical FPC and PIC, ACX Series) (Junos OS Evolved Release) | **18716**
- show chassis routing-engine bios (ACX Series) | **18717**
- show chassis routing-engine hard-disk-test disk /dev/sdb status (ACX Series) | **18717**
- Usage | **18720**
- Release Information | **18720**

Syntax

```
show chassis routing-engine
<bios | slot>
```

Syntax (ACX Series, PTX Series, and MX104 Universal Routing Platforms.)

```
show chassis routing-engine
```

Syntax (EX Series Switches)

```
show chassis routing-engine
<slot>
<satellite [slot-id slot-id |device-alias alias-name>
```

Syntax (QFX Series)

```
show chassis routing-engine
<interconnect-device name>
<node-device name>
<slot>
<bios>
<errors>
```

Syntax (MX Series Routers)

```
show chassis routing-engine
<all-members>
<bios | slot>
<local>
<member member-id>
<satellite [slot-id slot-id |device-alias alias-name>
```

Syntax (MX204 and MX10003 Universal Routing Platforms)

```
show chassis routing-engine
<slot>
<bios>
<errors>
```

Syntax (TX Matrix Routers)

```
show chassis routing-engine
<bios | slot>
<lcc number  | scc>
```

Syntax (TX Matrix Plus Routers)

```
show chassis routing-engine
<bios | slot>
<lcc number  | sfc number>
```

Syntax (Junos OS Evolved)

```
show chassis routing-engine
<slot>
<bios>
<hard-disk-test>
```

Description

Display the status of the Routing Engine.

Options

none	Display information about one or more Routing Engines. On a TX Matrix router, display information about all Routing Engines on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display information about all Routing Engines on the TX Matrix Plus router and its attached routers.
all-members	(MX Series routers only) (Optional) Display Routing Engine information for all members of the Virtual Chassis configuration.
bios	(Optional) Display the (BIOS) firmware version.
errors	(Optional) Display routing engine errors.
hard-disk-test	(Junos OS Evolved only) (Optional) Display the health of the hard disk. Use `disk /dev/disk-name status` to display the status of a particular disk.
interconnect-device *number*	(QFabric systems only) (Optional) Display Routing Engine information for a specified Interconnect device.
lcc *number*	(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Routing Engine information for a specified T640 router (line-card chassis)

that is connected to the TX Matrix router. On a TX Matrix Plus router, display Routing Engine information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Display Routing Engine information for the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display Routing Engine information for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-device *number* (QFabric systems only) (Optional) Display Routing Engine information for a specified Node device.

satellite [slot-id *slot-id* **|device-alias** *alias-name***]** (Junos Fusion only) (Optional) Display Routing Engine information for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

scc (TX Matrix routers only) (Optional) Display Routing Engine information for the TX Matrix router (switch-card chassis).

sfc *number* (TX Matrix Plus routers only) (Optional) Display Routing Engine information for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

slot (Systems with multiple Routing Engines) (Optional) Display information for an individual Routing Engine. Replace *slot* with 0 or 1. For QFX3500 switches, there is only one Routing Engine, so you do not need to specify the slot number.

Required Privilege Level

view

Output Fields

Table 321 on page 18708 lists the output fields for the `show chassis routing-engine` command. Output fields are listed in the approximate order in which they appear.

Table 321: show chassis routing-engine Output Fields

Field Name	Field Description
`Slot`	(Systems with single and multiple Routing Engines) Slot number.
`Current state`	(Systems with multiple Routing Engines) Current state of the Routing Engine: `Master`, `Backup`, or `Disabled`.
`Election priority`	(Systems with multiple Routing Engines) Election priority for the Routing Engine: `Master` or `Backup`.
`Temperature`	Temperature of the air flowing past the Routing Engine.
`CPU Temperature`	Temperature of the CPU.
`DRAM`	Total DRAM available to the Routing Engine's processor. **NOTE**: When the chassis has two Routing Engines, the amount of DRAM should be the same on both. A DRAM size mismatch error can result when the Routing Engines have different amounts of DRAM. Starting with Junos OS Release 12.3R1, the DRAM field displays both available memory and installed memory.
`Memory utilization`	Percentage of Routing Engine memory being used. **NOTE**: For platforms running Junos OS with upgraded FreeBSD, the way memory utilization is calculated has changed. Starting in Junos OS Release 15.1R1, inactive memory is no longer included in calculating memory utilization. Inactive memory is considered as free. The value for used memory decreases and results in more memory to be available for other processes. For platforms that run Junos OS with upgraded FreeBSD, see Release Information for Junos OS with Upgraded FreeBSD.

Table 321: show chassis routing-engine Output Fields *(Continued)*

Field Name	Field Description
`CPU utilization`	Information about the Routing Engine's CPU utilization: • `User`—Percentage of CPU time being used by user processes. • `Background`—Percentage of CPU time being used by background processes. • `Kernel`—Percentage of CPU time being used by kernel processes. • `Interrupt`—Percentage of CPU time being used by interrupts. • `Idle`—Percentage of CPU time that is idle.
`5 sec CPU Utilization`	Information about the Routing Engine's CPU utilization in the past 5 seconds: • `User`—Percentage of CPU time being used by user processes. • `Background`—Percentage of CPU time being used by background processes. • `Kernel`—Percentage of CPU time being used by kernel processes. • `Interrupt`—Percentage of CPU time being used by interrupts. • `Idle`—Percentage of CPU time that is idle.
`1 min CPU Utilization`	Information about the Routing Engine's CPU utilization in the past 1 minute: • `User`—Percentage of CPU time being used by user processes. • `Background`—Percentage of CPU time being used by background processes. • `Kernel`—Percentage of CPU time being used by kernel processes. • `Interrupt`—Percentage of CPU time being used by interrupts. • `Idle`—Percentage of CPU time that is idle.

Table 321: show chassis routing-engine Output Fields *(Continued)*

Field Name	Field Description
`5 min CPU Utilization`	Information about the Routing Engine's CPU utilization in the past 5 minutes: • `User`—Percentage of CPU time being used by user processes. • `Background`—Percentage of CPU time being used by background processes. • `Kernel`—Percentage of CPU time being used by kernel processes. • `Interrupt`—Percentage of CPU time being used by interrupts. • `Idle`—Percentage of CPU time that is idle.
`15 min CPU Utilization`	Information about the Routing Engine's CPU utilization in the past 15 minutes: • `User`—Percentage of CPU time being used by user processes. • `Background`—Percentage of CPU time being used by background processes. • `Kernel`—Percentage of CPU time being used by kernel processes. • `Interrupt`—Percentage of CPU time being used by interrupts. • `Idle`—Percentage of CPU time that is idle.
`Model`	Routing Engine model number.
`Serial ID`	(Systems with multiple Routing Engines) Identification number of the Routing Engine in this slot.
`Start time`	Time at which the Routing Engine started running.
`Uptime`	How long the Routing Engine has been running.
Routing Engine BIOS Version	BIOS version being run by the Routing Engine.

Table 321: show chassis routing-engine Output Fields *(Continued)*

Field Name	Field Description
`Last reboot reason`	Reason for last reboot, including: • `power cycle/failure`—Halt of the Routing Engine using the `halt` command, powering down using the power button on the chassis or any other method (such as removal of the control board or Routing Engine), and then powering back the Routing Engine. A halt of the operating system also occurs if you enter the `request system halt` command. You can enter this command to halt the system operations on the chassis or specific Routing Engines. To restart the software, press any key on the keyboard. • `watchdog`—Reboot due to a hardware watchdog. A watchdog is a hardware monitoring process that examines the health and performance of the router to enable the device to recover from failures. A watchdog checks for problems at certain intervals, and reboots the routing engine if a problem is encountered. • `reset-button reset`—(Not available on the EX Series switch) Reboot due to pressing of the reset button on the Routing Engine. • `power-button hard power off`—Reboot due to pressing of the power button on the chassis. A powering down of the software also occurs if you enter the `request system power-off` command. You can enter this command to power down the chassis or specific Routing Engines; you can then restart the software. • `misc hardware reason`—Reboot due to miscellaneous hardware reasons. • `thermal shutdown`—Reboot due to the router or switch reaching a critical temperature at which point it is unsafe to continue operations. • `hard disk failure`—Reboot due to a hard disk or solid-state drive (SSD) failure. • `reset from debugger`—Reboot due to reset from the debugger. • `chassis control reset`—Restart the chassis process that manages PICs, FPCs, and other hardware components. The chassis control module that runs the Routing Engine performs management and monitoring functions, and it provides a single access point for operational and maintenance functions. A reset of the chassis management process occurs when you enter the `restart chassis-control` command. • `bios auto recovery reset`—Reboot due to a BIOS auto-recovery reset. • `could not be determined`—Reboot due to an undetermined reason.

Table 321: show chassis routing-engine Output Fields *(Continued)*

Field Name	Field Description
	• `Router rebooted after a normal shutdown`—Reboot due to a normal shutdown. This reason is displayed if the Routing Engine is powered down by pushing and holding the online/offline button on the Routing Engine faceplate for 30 seconds, and then powered back. A reboot of the software also occurs if you enter the `request system reboot` command. You can enter this command to reboot the chassis or specific Routing Engines. • `Low battery detected` This reason is displayed if the Routing Engine is powered down due to detection of low battery and powered back when battery is charged. • `Catastrophic error` This reason is displayed if the Routing Engine is powered down due to catastrophic error (for example, link failure between FPC line card and PFE(packet forwarding engine)) and powered back. QFX5700 displays the following additional reboot reasons: • `FPGA reset` Reboot due to FPGA reset. • `MSMI error/hw misc` Reboot due to hardware MSMI (Microsoft Mail Connector interchange) Service Returned Service-Specific error. • `swizzle reset` Reboot due to swizzle reset, power off and turning back on. • `PCH cold reset` Reboot due to complete removal of power and restart of Intel PCH (Platform Controller Hub). • `software reboot` Reboot of system due to software reboot process. • `Straight to s5 reset` Reboot due to system in S5 state and does not retain memory state. • `iTCO watchdog` Linux TCO hardware watchdog reboots the system automatically when there is a systems pause.
`Load averages`	Routing Engine load averages for the last 1, 5, and 15 minutes.

Sample Output

The sample outputs for show chassis routing engine are as follows:

show chassis routing-engine (MX Series Routers)

```
user@host> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state                  Master
    Election priority              Master (default)
    Temperature                 46 degrees C / 114 degrees F
    CPU temperature             45 degrees C / 113 degrees F
    DRAM                       49089 MB (49152 MB installed)
    Memory utilization           4 percent
    5 sec CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
      Idle                      99 percent
    1 min CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
      Idle                      99 percent
    5 min CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
      Idle                      99 percent
    15 min CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
      Idle                      99 percent
    Model                          RE X10
    Start time                     2024-03-28 09:19:21 PDT
    Uptime                         5 days, 18 hours, 43 seconds
    Last reboot reason             0x2000:hypervisor reboot
    Load averages:                 1 minute   5 minute  15 minute
                                       0.19       0.14       0.14
Routing Engine status:
```

```
Slot 1:
  Current state                  Backup
  Election priority              Backup (default)
  Temperature                 44 degrees C / 111 degrees F
  CPU temperature             44 degrees C / 111 degrees F
  DRAM                       49089 MB (49152 MB installed)
  Memory utilization          5 percent
  5 sec CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                 0 percent
    Idle                    100 percent
  1 min CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                 0 percent
    Idle                     99 percent
  5 min CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                 0 percent
    Idle                     99 percent
  15 min CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                 0 percent
    Idle                     99 percent
  Model                          RE X10
  Serial ID                      BCBS9737
  Start time                     2024-03-28 09:20:49 PDT
  Uptime                         5 days, 17 hours, 59 minutes, 5 seconds
  Last reboot reason             0x10:misc hardware reason
  Load averages:                 1 minute   5 minute  15 minute
                                     0.04       0.14       0.10
```

show chassis routing-engine (QFX5700 Switches) (Junos OS Evolved Release)

```
user@switch> show chassis routing-engine
Routing Engine status:
  Slot 0:
  Current state         Master
  Election priority     Master
  Temperature           35 degrees C / 95 degrees F
  CPU temperature       39 degrees C / 102 degrees F
  DRAM                  25825 MB (31566 MB installed)
  Memory utilization    18 percent
  5 sec CPU utilization:
  User                   2 percent
  Background             0 percent
  Kernel                 2 percent
  Interrupt              0 percent
  Idle                   94 percent
  1 min CPU utilization:
  User                   2 percent
  Background             0 percent
  Kernel                 2 percent
  Interrupt              0 percent
  Idle                   94 percent
  5 min CPU utilization:
  User                   2 percent
  Background             0 percent
  Kernel                 2 percent
  Interrupt              0 percent
  Idle                   94 percent
  15 min CPU utilization:
   User                   2 percent
   Background             0 percent
   Kernel                 2 percent
   Interrupt              0 percent
   Idle                   94 percent
   Model                  QFX5700-RE
  Serial ID               BUILTIN
  Start time 2021-06-09 02:50:08 PDT
  Uptime 42 minutes, 48 seconds
  Load averages: 1 minute 5 minute 15 minute
  1.48 1.71 1.83
```

```
  Last reboot reason      power cycle
```

show chassis routing-engine (QFX5230-64CD devices with logical FPC and PIC, ACX Series) (Junos OS Evolved Release)

```
user@root> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state                  Master
    Election priority              Master (default)
    Temperature                 0 degrees C / 32 degrees F
    DRAM                       59026 MB (65536 MB installed)
    Memory utilization           9 percent
    5 sec CPU utilization:
      User                       1 percent
      Background                 0 percent
      Kernel                     1 percent
      Interrupt                  1 percent
      Idle                      98 percent
    1 min CPU utilization:
      User                       1 percent
      Background                 0 percent
      Kernel                     2 percent
      Interrupt                  1 percent
      Idle                      95 percent
    5 min CPU utilization:
      User                       1 percent
      Background                 0 percent
      Kernel                     2 percent
      Interrupt                  1 percent
      Idle                      95 percent
    15 min CPU utilization:
      User                       1 percent
      Background                 0 percent
      Kernel                     2 percent
      Interrupt                  0 percent
      Idle                      78 percent
    Model                          QFX5230-64CD
    Serial ID                      BUILTIN
    Start time                     2022-04-08 18:07:00 UTC
```

```
    Uptime                          25 minutes, 5 seconds
    Load averages:                  1 minute   5 minute  15 minute
                                        0.44       0.92       1.28
    Last reboot reason              Unknown
```

show chassis routing-engine bios (ACX Series)

```
user@host> show chassis routing-engine bios
Routing Engine BIOS Version: 13.01
```

show chassis routing-engine hard-disk-test disk /dev/sdb status (ACX Series)

```
user@host> show chassis routing-engine hard-disk-test disk /dev/sda status

smartctl 7.3 2022-02-28 r5338 [x86_64-linux-5.15.76-yocto-standard-juniper-02183-g46c537862907]
(local build)
Copyright (C) 2002-22, Bruce Allen, Christian Franke, www.smartmontools.org

=== START OF INFORMATION SECTION ===
Model Family:     SwissBit Slim Sata SSD
Device Model:     SFSA100GM2AK4TO-I-6B-66A-JUN
Serial Number:    0000602219249A000046
Firmware Version: SBR10021
User Capacity:    100,030,242,816 bytes [100 GB]
Sector Size:      512 bytes logical/physical
Rotation Rate:    Solid State Device
Form Factor:      2.5 inches
TRIM Command:     Available, deterministic, zeroed
Device is:        In smartctl database 7.3/5319
ATA Version is:   ACS-3 (minor revision not indicated)
SATA Version is:  SATA 3.2, 6.0 Gb/s (current: 6.0 Gb/s)
Local Time is:    Sat Sep 21 03:35:14 2024 PDT
SMART support is: Available - device has SMART capability.
SMART support is: Enabled

=== START OF READ SMART DATA SECTION ===
SMART overall-health self-assessment test result: PASSED

General SMART Values:
```

```
Offline data collection status:  (0x02)    Offline data collection activity
                    was completed without error.
                    Auto Offline Data Collection: Disabled.
Self-test execution status:      (   0)    The previous self-test routine completed
                    without error or no self-test has ever
                    been run.
Total time to complete Offline
data collection:         (    0) seconds.
Offline data collection
capabilities:              (0x53) SMART execute Offline immediate.
                    Auto Offline data collection on/off support.
                    Suspend Offline collection upon new
                    command.
                    No Offline surface scan supported.
                    Self-test supported.
                    No Conveyance Self-test supported.
                    Selective Self-test supported.
SMART capabilities:            (0x0003)    Saves SMART data before entering
                    power-saving mode.
                    Supports SMART auto save timer.
Error logging capability:        (0x01)    Error logging supported.
                    General Purpose Logging supported.
Short self-test routine
recommended polling time:      (   2) minutes.
Extended self-test routine
recommended polling time:      (  15) minutes.
SCT capabilities:            (0x003d)    SCT Status supported.
                    SCT Error Recovery Control supported.
                    SCT Feature Control supported.
                    SCT Data Table supported.

SMART Attributes Data Structure revision number: 1
Vendor Specific SMART Attributes with Thresholds:
ID# ATTRIBUTE_NAME          FLAG     VALUE WORST THRESH TYPE      UPDATED  WHEN_FAILED RAW_VALUE
  1 Raw_Read_Error_Rate     0x000b   100   100   000    Pre-fail  Always       -       0
  5 Reallocated_Sector_Ct   0x0013   100   100   000    Pre-fail  Always       -       0
  9 Power_On_Hours          0x0012   100   100   000    Old_age   Always       -       7739
 12 Power_Cycle_Count       0x0012   100   100   000    Old_age   Always       -       136
 16 Unknown_Attribute       0x0112   094   094   001    Old_age   Always       -       1958
 17 Unknown_Attribute       0x0012   100   100   000    Old_age   Always       -       30000
160 Uncorrect_Sect_Cnt_RW   0x0002   100   100   000    Old_age   Always       -       0
161 Spare_Block_Count       0x0012   100   100   000    Old_age   Always       -       61539243
163 Initial_Invalid_Blocks  0x0003   100   100   000    Pre-fail  Always       -       10
```

```
164 Total_Erase_Count       0x0012   100   100   000    Old_age   Always       -       593606
165 Maximum_Erase_Count     0x0002   100   100   000    Old_age   Always       -       449
166 Minimum_Erase_Count     0x0012   100   100   000    Old_age   Always       -       348
167 Average_Erase_Count     0x0012   087   087   000    Old_age   Always       -       408
168 MaxSpecified_Erase_Count0x0012   100   100   000    Old_age   Always       -       3000
169 PowerOn_UECC_Count      0x0003   100   100   000    Pre-fail  Always       -       102
193 Dynamic_Remaps          0x0012   100   100   000    Old_age   Always       -       0
194 Temperature_Celsius     0x0023   100   100   000    Pre-fail  Always       -       3871780
195 Hardware_ECC_Recovered  0x0012   100   100   000    Old_age   Always       -       12849
196 Reallocated_Event_Count 0x0012   000   000   000    Old_age   Always       -       0
198 Uncorrectable_Sect_Count0x0012   100   100   000    Old_age   Always       -       0
199 SATA_Phy_CRC_Error_Count0x000b   100   100   000    Pre-fail  Always       -       0
215 Trim_Count              0x0012   100   100   000    Old_age   Always       -       1816
235 TotalFlashLBA_Written   0x0012   100   100   000    Old_age   Always       -
233412999168
237 TotalFlashLBA_Written_E 0x0012   100   100   000    Old_age   Always       -       0
241 TotalHostLBA_Write      0x0012   100   100   000    Old_age   Always       -
167245610753
242 TotalHostLBA_Read       0x0012   100   100   000    Old_age   Always       -
19887124523
243 TotalHostLBA_Written_E  0x0012   100   100   000    Old_age   Always       -       0
244 TotalHostLBA_Read_E     0x0012   100   100   000    Old_age   Always       -       0
248 SSD_Life_Left           0x0112   087   087   001    Old_age   Always       -       87
249 Spare_Block_Life_Left   0x1913   100   100   025    Pre-fail  Always       -       100

SMART Error Log Version: 1
No Errors Logged

SMART Self-test log structure revision number 1
Num  Test_Description    Status                  Remaining  LifeTime(hours)  LBA_of_first_error
# 1  Extended offline    Completed without error       00%      7720         -

SMART Selective self-test log data structure revision number 1
 SPAN  MIN_LBA  MAX_LBA  CURRENT_TEST_STATUS
    1        0        0  Not_testing
    2        0        0  Not_testing
    3        0        0  Not_testing
    4        0        0  Not_testing
    5        0        0  Not_testing
Selective self-test flags (0x0):
  After scanning selected spans, do NOT read-scan remainder of disk.
If Selective self-test is pending on power-up, resume after 0 minute delay.
```

Usage

The `show chassis routing-engine` and `show chassis feb` commands can be used to find the memory allocated for each RE (Routing Engine) and Packet Forwarding Engine components.

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced in Junos OS Release in 9.6 for the TX Matrix Plus router.

`5 sec CPU Utilization`, `1 min CPU Utilization`, `5 min CPU Utilization`, and `15 min CPU Utilization` output fields introduced in Junos OS Release 11.3R1.

`satellite` option introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

- request chassis routing-engine master
- Configuring Routing Engine Redundancy
- Switching the Global Primary and Backup Roles in a Virtual Chassis Configuration
- show chassis feb | **18278**

show chassis routing-engine (View)

IN THIS SECTION

- Syntax | **18721**
- Description | **18721**
- Required Privilege Level | **18721**
- Output Fields | **18721**
- Sample Output | **18723**
- Sample Output | **18723**
- Sample Output | **18723**
- Sample Output | **18724**

Sample Output | 18725
Release Information | 18726

Syntax

```
show chassis routing-engine
```

Description

Display the Routing Engine status of the chassis cluster.

Required Privilege Level

view

Output Fields

Table 322 on page 18721 lists the output fields for the `show chassis routing-engine` command. Output fields are listed in the approximate order in which they appear.

Table 322: show chassis routing-engine Output Fields

Field Name	Field Description
`Temperature`	Routing Engine temperature. (Not available for vSRX Virtual Firewall deployments.)
`CPU temperature`	CPU temperature. (Not available for vSRX Virtual Firewall deployments.)

Table 322: show chassis routing-engine Output Fields *(Continued)*

Field Name	Field Description
`Total memory`	Total memory available on the system. **NOTE**: Starting with Junos OS Release 15.1x49-D70 and Junos OS Release 17.3R1, there is a change in the method for calculating the memory utilization by a Routing Engine. The inactive memory is now subtracted from the total available memory. There is thus, a decrease in the reported value for used memory; as the inactive memory is now considered as free.
`Control plane memory`	Memory available for the control plane.
`Data plane memory`	Memory reserved for data plane processing.
`CPU utilization`	Current CPU utilization statistics on the control plane core.
`User`	Current CPU utilization in user mode on the control plane core.
`Background`	Current CPU utilization in nice mode on the control plane core.
`Kernel`	Current CPU utilization in kernel mode on the control plane core.
`Interrupt`	Current CPU utilization in interrupt mode on the control plane core.
`Idle`	Current CPU utilization in idle mode on the control plane core.
`Model`	Routing Engine model.
`Start time`	Routing Engine start time.
`Uptime`	Length of time the Routing Engine has been up (running) since the last start.
`Last reboot reason`	Reason for the last reboot of the Routing Engine.

Table 322: show chassis routing-engine Output Fields *(Continued)*

Field Name	Field Description
Load averages	The average number of threads waiting in the run queue or currently executing over 1-, 5-, and 15-minute periods.

Sample Output

Sample Output

show chassis routing-engine (vSRX Virtual Firewall)

```
user@host> show chassis routing-engine
Routing Engine status:
    Total memory              1024 MB Max   358 MB used ( 35 percent)
      Control plane memory    1024 MB Max   358 MB used ( 35 percent)
    5 sec CPU utilization:
      User                       2 percent
      Background                 0 percent
      Kernel                     4 percent
      Interrupt                  6 percent
      Idle                      88 percent
    Model                          VSRX RE
    Start time                     2015-03-03 07:04:18 UTC
    Uptime                         2 days, 11 hours, 51 minutes, 11 seconds
    Last reboot reason             Router rebooted after a normal shutdown.
    Load averages:                 1 minute   5 minute  15 minute
                                       0.07       0.04       0.06
```

Sample Output

show chassis routing-engine (SRX5400)

```
user@host> show chassis routing-engine
Routing Engine status:
  Slot 0:
```

```
    Current state                   Master
    Election priority               Master (default)
    Temperature                 31 degrees C / 87 degrees F
    CPU temperature             31 degrees C / 87 degrees F
    DRAM                      16323 MB (16384 MB installed)
    Memory utilization          10 percent
    5 sec CPU utilization:
      User                       2 percent
      Background                 0 percent
      Kernel                    12 percent
      Interrupt                  1 percent
      Idle                      85 percent
    Model                           RE-S-1800x4
    Serial ID                       9016272401
    Start time                      2019-07-08 01:17:10 PDT
    Uptime                          2 minutes, 43 seconds
    Last reboot reason              Router rebooted after a normal shutdown.
    Load averages:                  1 minute   5 minute  15 minute
                                        0.63       0.52       0.24
```

Sample Output

show chassis routing-engine (SRX4100)

```
user@host> show chassis routing-engine
Routing Engine status:
    Temperature                 29 degrees C / 84 degrees F
    CPU temperature             29 degrees C / 84 degrees F
    Total memory              1954 MB Max   567 MB used ( 29 percent)
    Memory utilization          24 percent
    5 sec CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
      Idle                     100 percent
    1 min CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
```

```
      Idle                     100 percent
    5 min CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
      Idle                     100 percent
    15 min CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
      Idle                     100 percent
    Model                          SRX Routing Engine
    Serial ID                      BUILTIN
    Uptime                         17 days, 5 hours, 1 minute, 52 seconds
    Last reboot reason             0x4000:VJUNOS reboot
    Load averages:                 1 minute   5 minute  15 minute
                                       0.00       0.00       0.00
```

The Total memory 64 GB is distributed between the routing engine in the form of virtual machine for the TVP platforms (SRX1500, SRX4100, SRX4200) and the rest for the packet forwarding engine (PFE). TVP has a different architecture differentiating PFE from Junos and additional API compatibility. The above mentioned devices are the only ones with this TVP architecture in SRX. The `show chassis routing-engine` command displays only the Routing Engine memory.

Sample Output

show chassis routing-engine (SRX1500)

```
user@host> show chassis routing-engine
Routing Engine status:
    Temperature                 42 degrees C / 107 degrees F
    CPU temperature             42 degrees C / 107 degrees F
    Total memory              1954 MB Max   528 MB used ( 27 percent)
    Memory utilization          23 percent
    5 sec CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
```

```
    Idle                      100 percent
  1 min CPU utilization:
    User                        0 percent
    Background                  0 percent
    Kernel                      0 percent
    Interrupt                   0 percent
    Idle                       99 percent
  5 min CPU utilization:
    User                        0 percent
    Background                  0 percent
    Kernel                      0 percent
    Interrupt                   0 percent
    Idle                       99 percent
  15 min CPU utilization:
    User                        0 percent
    Background                  0 percent
    Kernel                      0 percent
    Interrupt                   0 percent
    Idle                       96 percent
  Model                           SRX Routing Engine
  Serial ID                       BUILTIN
  Uptime                          52 minutes, 27 seconds
  Last reboot reason              0x4000:VJUNOS reboot
  Load averages:                  1 minute   5 minute  15 minute
                                      0.00       0.00       0.00
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

request system snapshot (Maintenance) | 15208

show chassis satellite

IN THIS SECTION

- Syntax | 18727
- Description | 18727
- Options | 18727
- Required Privilege Level | 18728
- Output Fields | 18728
- Sample Output | 18737
- Sample Output | 18738
- Sample Output | 18739
- Sample Output | 18739
- Release Information | 18743

Syntax

```
show chassis satellite
[device-alias device-alias | fpc-slot fpc-slot | cluster cluster-name]
[brief | detail | extensive | terse]
<since time>
```

Description

Display the status of the satellite device connections in a Junos Fusion.

Options

none	(Same as **brief**) Display satellite device connection information
device-alias *device-alias*	(Optional) Display satellite device connection information for the satellite device using the specified device alias only.

fpc-slot *fpc-slot* (Optional) Display satellite device connection information for the satellite device using the specified FPC slot number only.

cluster *cluster-name* (Optional) Display satellite device connection information for the satellite devices in the specified satellite device cluster only.

brief | detail | extensive | terse (Optional) Display the specified level of output.

since *time* (Optional) Display the satellite devices that have been added to the Junos Fusion on or after a certain date or time, in *YYYY-MM-DD.HH:MM:SS* format.

To display all satellite devices added since a specified date, enter the specific date. For instance, to display all satellite devices added on or after December 22nd, 2015, enter **2015-12-22** as the *time*.

To display all satellite devices added since a specified time, enter the specific date and time. For instance, to display all satellite devices added on or after 11:01AM on December 22nd, 2015, enter **2015-12-22.11:01:00** as the *time*.

Required Privilege Level

view

Output Fields

Table 323 on page 18728 lists the output fields for the `show chassis satellite` command. Output fields are listed in the approximate order in which they appear.

Table 323: show chassis satellite Output Fields

Field Name	Field Description	Level of Output
Fields for Interface		
`Alias`	The satellite device's alias. The satellite device's alias is configured using the **set chassis satellite-management fpc** *slot-id* **alias** *alias* statement.	brief extensive none

Table 323: show chassis satellite Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Slot`	The slot number of the satellite device. The slot number can be configured using the **set chassis satellite-management fpc *slot-id*** statement..	brief terse extensive none

Table 323: show chassis satellite Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Device State`	The state of the satellite device within the Junos Fusion. The most common device states: • **Online**—the satellite device is online and active. This is the satellite device state during normal operating procedure. • **Offline**—the satellite device is offline and not detected. This state is typically seen when the satellite device has been disconnected from the aggregation device, or when all cascade or uplink ports connecting the satellite device to the aggregation device are down. • **Present**—the satellite device is recognized by the aggregation device, but is not online. This state is typically seen before a satellite device goes online, or while satellite device configuration is in progress or finalizing. • **Rebooting**—the satellite device is rebooting. • **Disable**—the satellite device has been disabled. • **Misconfig**—the satellite device is not properly configured. This state is typically seen when the system ID, cascade port, or FPC slot ID defined for the satellite device has a misconfiguration. • **Miswire**—the satellite device is miswired. This state is typically seen when a satellite device is wired to two aggregation devices but is not configured for multihoming. Use **show chassis satellite detail** to gather more information on the issue when the device state is **Miswire**. Other less common device states include: • **ModeChanging**—the device is converting from a standalone device to a satellite device, or from a satellite device to a standalone device.	brief terse extensive none

Table 323: show chassis satellite Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **ModeChangeFail**—the mode change operation failed. • **MinorUpgradeOn**—A minor satellite software upgrade is in progress. • **MajorUpgradeOn**—A major satellite software upgrade is in progress. • **Upgrade-pending**—the satellite device is waiting for a satellite software upgrade. • **ProvSessionDn**—the provisioning session is down. • **ReconcileState**—the satellite provisioning daemon has restarted and is reconciling the satellite device state.	
`Cascade Ports`	The cascade port or ports. A cascade port is a port on the aggregation device that connects to a satellite device in a Junos Fusion.	brief extensive none

Table 323: show chassis satellite Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Port State`	The state of the cascade port on the aggregation device. Port states include: • **online**—the cascade port is online and active. This is the port state during normal operating procedure. • **txUpRxDn**—Tx or Rx forwarding is disabled on the cascade port. This state is often seen when a second aggregation device is added to a Junos Fusion topology, and the devices in the Junos Fusion are synchronizing to the new topology. • **miswire**—the cascade port is miswired. This state is typically seen when a satellite device is interconnected to two aggregation devices but multihoming is not configured. Use **show chassis satellite detail** to gather more information on the issue when the device state is **Miswire**. • **present**—The cascade port recognized the satellite device and is up. • **misconfig**—the cascade port is assigned, but this interface is not working correctly due to a misconfiguration. • **down**—the cascade port is down. • **offline**—the satellite device was previously recognized from this interface, but is no longer present. • **absent**—the cascade port is configured but no satellite device is detected on the interface.	brief extensive none

Table 323: show chassis satellite Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Extended Ports Total`	The total number of extended ports on the satellite device. An extended port is a network-facing port on the satellite device that sends and receives network traffic for the Junos Fusion.	brief none terse
`Extended Ports Up`	The number of active extended ports.	brief none terse
`Model`	The hardware model of the satellite device.	terse
`Version`	The version of satellite device software running on the satellite device.	terse
`Satellite Alias`	The satellite device's alias. The satellite device's alias is configured using the **set chassis satellite-management fpc *slot-id* alias *alias*** statement.	detail
`FPC slot`	The FPC slot number of the satellite device. The slot number can be configured using the **set chassis satellite-management fpc *slot-id*** statement.	detail
`Operational State`	The operational state of the satellite device. The state UFDDown indicates that uplink failure detection disabled the satellite device's extended ports due to an uplink port failure.	detail
`Product Model`	The hardware model of the satellite device.	detail

Table 323: show chassis satellite Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Product Family`	The product family of the satellite device.	detail
`Serial number`	The serial number of the satellite device.	detail
`System ID`	The system ID of the satellite device. The system ID is also the satellite device's MAC address.	detail
`Software package version`	The satellite software version running on the satellite device.	detail
`Host software version`	The host operating system software version running on the satellite device.	detail
`Management Address`	The management IP address of the satellite device. This management IP address belongs to an internal routing instance. This management address is assigned by the control plane internally based on FPC slot ID and is used for the control plane traffic between the aggregation device and satellite device. All management in a Junos Fusion should be done through the aggregation device. The management IP address of the satellite device is useful for debugging purposes by expert users only.	detail
`UFD config state`	Uplink failure detection configuration state.	detail
`Minimum link`	Uplink failure detection minimum active uplink port setting.	detail

Table 323: show chassis satellite Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Holddown timer (seconds)`	Uplink failure detection holddown timer setting, in seconds.	detail
`UFD operational state`	Uplink failure detection operational state.	detail
`Candidate uplink interfaces (pic/port)`	Uplink failure detection candidate uplink interfaces.	detail
`Extended Ports`	The number of extended ports for the satellite device. The number on the left is the total number of extended ports, and the number on the right is the total number of extended ports currently in the up state.	extensive
`When`	The date and time of the event.	extensive
`Event`	The event.	extensive
`Action`	The actions that resulted from the event.	extensive
Fields for Cascade interfaces		
`Interface Name`	The name of the cascade interface on the aggregation device.	detail
`State`	The state of the cascade interface.	detail

Table 323: show chassis satellite Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Uplink Interface`	The name of the uplink interface on the satellite device.	detail
`Adjacency state`	The adjacency state of the cascade to uplink interface link.	detail
`Last transition`	The amount of time that has passed since the last transition of the cascade to uplink interface link.	detail
`Adjacency down count (Interface Name)`	The number of times the cascade to uplink interface link has gone into the down state.	detail
`RX Packet`	The number of packets received on the cascade interface.	detail
`Last received packet`	The amount of time that has passed since the last packet was received on the cascade interface.	detail
`Peer adjacency information`	The amount of time that has passed since the last peer adjacency transition.	detail
`Adjacency down count (Peer adjacency information)`	The number of times the cascade to uplink interface link has gone into the down state.	detail
`Last down cause`	The cause of the last adjacency failure.	detail

Table 323: show chassis satellite Output Fields *(Continued)*

Field Name	Field Description	Level of Output
SDPD restart detected	The number of times that the satellite device protocol process has restarted.	detail
Fields for Process information		
Process Name	The name of the process.	detail
PID	The process identification number of the process.	detail
State	The current state of the process.	detail
Number of restart detected	The number of times the process has restarted.	detail
Uptime	The amount of time that the process has been running.	detail

Sample Output

show chassis satellite

```
user@aggregation-device> show chassis satellite
                        Device          Cascade      Port       Extended Ports
Alias            Slot   State           Ports        State      Total/Up
qfx5100-24q-01   100    Online          xe-0/0/1     online     9/2
                                        xe-0/3/0     online
qfx5100-24q-02   101    Online          xe-0/0/2     online     20/12
                                        xe-0/3/1     online
qfx5100-24q-03   102    Online          xe-0/0/3     online     16/6
                                        xe-0/3/2     online
qfx5100-24q-04   103    Online          xe-0/0/4     online     16/4
                                        xe-0/3/3     online
qfx5100-24q-05   104    Online          xe-0/0/5     online     13/3
```

```
                                        xe-0/3/4     online
qfx5100-24q-06   105    Online          xe-0/0/6     online     24/15
                                        xe-0/3/5     online
qfx5100-24q-07   106    Online          xe-0/0/7     online     24/15
                                        xe-0/3/6     online
qfx5100-24q-08   107    Online          xe-0/0/8     online     21/12
                                        xe-0/3/7     online
ex4300-01        109    Online          xe-1/0/1     online     49/2
ex4300-02        110    Online          xe-1/0/2     online     49/2
ex4300-03        111    Online          xe-1/0/3     online     49/2
ex4300-04        112    Online          xe-1/0/4     online     49/11
ex4300-05        113    Online          xe-1/0/5     online     49/11
ex4300-06        114    Online          xe-1/0/6     online     49/11
ex4300-07        115    Online          xe-1/0/7     online     49/11
ex4300-08        116    Online          xe-1/1/0     online     49/11
ex4300-09        117    Online          xe-1/1/1     online     49/11
ex4300-10        118    Online          xe-1/1/2     online     49/11
ex4300-11        119    Online          xe-1/1/3     online     49/11
ex4300-12        120    Online          xe-1/1/4     online     49/11
ex4300-13        121    Online          xe-1/1/5     online     49/11
ex4300-14        122    Online          xe-1/1/6     online     49/11
ex4300-15        123    Online          xe-1/1/7     online     49/11
ex4300-16        124    Online          xe-1/2/1     online     49/11
ex4300-17        125    Online          xe-1/2/2     online     49/11
ex4300-18        126    Online          xe-1/2/3     online     49/2
ex4300-19        127    Online          xe-1/2/4     online     49/1
ex4300-20        128    Online          xe-1/2/5     online     49/1
ex4300-21        129    Online          xe-1/2/6     online     49/1
ex4300-22        130    Online          xe-1/2/7     online     49/1
```

Sample Output

show chassis satellite device-alias

```
user@aggregation-device> show chassis satellite device-alias ex4300-22
                         Device          Cascade      Port       Extended Ports
Alias            Slot    State           Ports        State      Total/Up
    ex4300-22    130     Online          xe-0/2/7     online     49/1
```

Sample Output

show chassis satellite fpc-slot 130

```
user@aggregation-device> show chassis satellite fpc-slot 130
                        Device          Cascade      Port       Extended Ports
Alias            Slot   State           Ports        State      Total/Up
ex4300-22        101    Online          xe-0/0/2     online     20/12
                                        xe-0/3/1     online
```

Sample Output

show chassis satellite terse

```
user@aggregation-device> show chassis satellite terse
       Device                            Extended Ports
Slot   State            Model             Total/Up    Version
101    Online           QFX5100-48S-6Q    7/7         3.0R1.1
102    Online           QFX5100-48S-6Q    7/7         3.0R1.1
103    Online           QFX5100-48S-6Q    6/5         3.0R1.1
104    Online           QFX5100-48S-6Q    14/14       3.0R1.1
105    Online           QFX5100-48S-6Q    18/18       3.0R1.1
106    Online           QFX5100-48S-6Q    17/16       3.0R1.1
107    Online           EX4300-48T        52/6        3.0R1.1
108    Online           EX4300-48T        52/15       3.0R1.1
109    Online           EX4300-48T        51/14       3.0R1.1
110    Online           EX4300-48T        51/14       3.0R1.1
111    Online           EX4300-48T        51/13       3.0R1.1
112    Online           EX4300-48T        51/12       3.0R1.1
113    Online           EX4300-48T        51/13       3.0R1.1
114    Online           QFX5100-24Q-2P    17/13       3.0R1.1
```

show chassis satellite detail

```
user@aggregation-device> show chassis satellite detail
Satellite Alias: qfx5100-48s-02
FPC Slot: 101
Operational State: Online
```

```
Product Model: QFX5100-48S-6Q
Product Family: i386
Serial number: ABC123DEF456
System id: 00:11:22:aa:bb:cc
Software package version: 3.0R1.1
Host software version: 1.0.0
Management Address: 172.16.0.101/32
Cascade interfaces:
    Interface Name: xe-0/0/2 State: online
        Uplink Interface: xe-001/0/48:0
        Adjacency state: Two-Way
        Last transition: 00:10:22
        Adjacency down count: 0
        Rx Packet: 65 Last received packet: 00:00:02
        Peer adjacency information: 00:10:22
            Adjacency down count: 3
            Last down cause: Interface Down
            SDPD restart detected: 3
    Interface Name: xe-0/2/1 State: online
        Uplink Interface: xe-001/0/48:1
        Adjacency state: Two-Way
        Last transition: 00:10:22
        Adjacency down count: 0
        Rx Packet: 64 Last received packet: 00:00:02
        Peer adjacency information: 00:10:22
            Adjacency down count: 3
            Last down cause: Interface Down
            SDPD restart detected: 3
    Interface Name: xe-2/0/0 State: online
        Uplink Interface: xe-001/0/48:2
        Adjacency state: Two-Way
        Last transition: 00:10:22
        Adjacency down count: 0
        Rx Packet: 65 Last received packet: 00:00:02
        Peer adjacency information: 00:10:22
            Adjacency down count: 3
            Last down cause: Interface Down
            SDPD restart detected: 3
    Interface Name: xe-2/1/6 State: online
        Uplink Interface: xe-001/0/48:3
        Adjacency state: Two-Way
        Last transition: 00:10:22
        Adjacency down count: 0
```

```
        Rx Packet: 65 Last received packet: 00:00:02
        Peer adjacency information: 00:10:22
            Adjacency down count: 3
            Last down cause: Hold timer expire
            SDPD restart detected: 3
Process information:
    Process Name: Provisioning PID: 6716 State: Running
        Number of restart detected: 0
        Uptime: 00:10:22
    Process Name: PFE PID: 3194 State: Running
        Number of restart detected: 0
        Uptime: 00:10:22
UFD config state: Enable (persist), Minimum link: 1,
Holdddown timer (seconds): 6
UFD operational state: Enable
Candidate uplink interfaces (pic/port):
        1/0
        1/1
        1/2
        1/3
        2/0
        2/1
        2/2
        2/3

Satellite Alias: qfx5100-48s-03
FPC Slot: 102
Operational State: Online
Product Model: QFX5100-48S-6Q
Product Family: i386
Serial number: ABCDEFG12345
System id: 00:11:22:aa:ba:cc
Software package version: 3.0R1.1
Host software version: 1.0.0
Management Address: 172.16.0.102/32
Cascade interfaces:
    Interface Name: xe-0/0/3 State: online
        Uplink Interface: xe-002/0/48:0
        Adjacency state: Two-Way
        Last transition: 00:10:22
        Adjacency down count: 0
        Rx Packet: 65 Last received packet: 00:00:02
        Peer adjacency information: 00:10:22
```

```
            Adjacency down count: 3
            Last down cause: Interface Down
            SDPD restart detected: 3
    Interface Name: xe-0/2/2 State: online
        Uplink Interface: xe-002/0/48:1
        Adjacency state: Two-Way
        Last transition: 00:10:22
        Adjacency down count: 0
        Rx Packet: 65 Last received packet: 00:00:02
        Peer adjacency information: 00:10:22
            Adjacency down count: 3
            Last down cause: Interface Down
            SDPD restart detected: 3
    Interface Name: xe-2/0/1 State: online
        Uplink Interface: xe-002/0/48:2
        Adjacency state: Two-Way
        Last transition: 00:10:22
        Adjacency down count: 0
        Rx Packet: 65 Last received packet: 00:00:02
        Peer adjacency information: 00:10:22
            Adjacency down count: 3
            Last down cause: Interface Down
            SDPD restart detected: 3
    Interface Name: xe-2/1/7 State: online
        Uplink Interface: xe-002/0/48:3
        Adjacency state: Two-Way
        Last transition: 00:10:22
        Adjacency down count: 0
        Rx Packet: 65 Last received packet: 00:00:02
        Peer adjacency information: 00:10:22
            Adjacency down count: 3
            Last down cause: Interface Down
            SDPD restart detected: 3
Process information:
    Process Name: Provisioning PID: 6667 State: Running
        Number of restart detected: 0
        Uptime: 00:10:22
    Process Name: PFE PID: 3155 State: Running
        Number of restart detected: 0
        Uptime: 00:10:22
<additional output removed for brevity>
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

show chassis satellite-cluster

IN THIS SECTION

- Syntax | 18743
- Description | 18743
- Options | 18744
- Required Privilege Level | 18744
- Output Fields | 18744
- Sample Output | 18750
- Sample Output | 18750
- Release Information | 18751

Syntax

```
show chassis satellite-cluster
[cluster cluster-name]
[brief | detail | extensive | terse]
```

Description

Display the status of the satellite device clusters in a Junos Fusion.

Options

none	(Same as **brief**) Display satellite device cluster information for satellite device clusters in the Junos Fusion.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
cluster *cluster-name*	Display satellite device cluster information for the specified satellite device cluster only.

Required Privilege Level

view

Output Fields

Table 324 on page 18744 lists the output fields for the `show chassis satellite-cluster` command. Output fields are listed in the approximate order in which they appear.

Table 324: show chassis satellite-cluster Output Fields

Field Name	Field Description	Level of Output
`cluster`	The name of the satellite device cluster. The name of the satellite device cluster is assigned using the **set chassis satellite-management cluster *cluster-name*** statement.	terse
`Provision`	The number of provisioned satellite devices in the specified satellite device cluster.	terse
`Present`	The number of present satellite devices in the specified satellite device cluster.	terse
`Online`	The number of online satellite devices in the specified satellite device cluster.	terse

Table 324: show chassis satellite-cluster Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Unprovision`	The number of unprovisioned satellite devices in the specified satellite device cluster.	terse
`Misconfig`	The number of misconfigured satellite devices in the specified satellite device cluster.	terse
`Cluster Name`	The name of the satellite device cluster. The name of the satellite device cluster is assigned using the **set chassis satellite-management cluster *cluster-name*** statement.	none detail extensive
`Number of devices provisioned`	The number of provisioned satellite devices in the specified satellite device cluster.	none detail extensive
`Number of devices present`	The number of present satellite devices in the specified satellite device cluster.	none detail extensive
`Number of devices unprovisioned`	The number of unprovisioned satellite devices in the specified satellite device cluster.	none detail extensive
`Number of devices misconfig/ miswired`	The number of misconfigured or miswired satellite devices in the specified satellite device cluster.	none detail extensive

Table 324: show chassis satellite-cluster Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Number of devices online`	The number of online satellite devices in the specified satellite device cluster.	none detail extensive
`Number of devices offline`	The number of offline satellite devices in the specified satellite device cluster.	none detail extensive
`Slot`	The slot number of the satellite device. The slot number can be configured using the **set chassis satellite-management fpc *slot-id*** statement.	none detail extensive

Table 324: show chassis satellite-cluster Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Device State`	The device state of the satellite device in the satellite device cluster. The most common device states: • **Online**—the satellite device is online and active. This is the satellite device state during normal operating procedure. • **Offline**—the satellite device is offline and not detected. This state is typically seen when the satellite device has been disconnected from the aggregation device, or when all cascade or uplink ports connecting the satellite device to the aggregation device are down. • **Present**—the satellite device is recognized by the aggregation device. In a satellite device cluster, this state is seen during normal operation for all satellite devices that are not directly cabled to the aggregation device. • **Rebooting**—the satellite device is rebooting. • **Disable**—the satellite device has been disabled. • **Misconfig**—the satellite device is not properly configured. This state is typically seen when the system ID, cascade port, or FPC slot ID defined for the satellite device has a misconfiguration. • **Miswire**—the satellite device is miswired. This state is typically seen when a satellite device is wired to two aggregation devices but is not configured for multihoming. Use **show chassis satellite detail** to gather more information on the issue when the device state is **Miswire**. Other less common device states include: • **ModeChanging**—the device is converting from a standalone device to a satellite device, or from a satellite device to a standalone device.	none detail extensive

Table 324: show chassis satellite-cluster Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **ModeChangeFail**—the mode change operation failed. • **MinorUpgradeOn**—A minor satellite software upgrade is in progress. • **MajorUpgradeOn**—A major satellite software upgrade is in progress. • **Upgrade-pending**—the satellite device is waiting for a satellite software upgrade. • **ProvSessionDn**—the provisioning session is down. • **ReconcileState**—the satellite provisioning daemon has restarted and is reconciling the satellite device state.	
`Distance`	The distance the satellite device is away from the aggregation device. In this output, each hop to get from the satellite device to the aggregation device is counted as 4. The `via` output provides the slot number of the satellite device in the satellite device cluster that is directly connected to the aggregation device and passing traffic for the specified satellite device.	none detail extensive
`Local Interface`	A local interface on the specified satellite device.	none detail extensive
`Remote Interface`	A remote interface on the specified satellite device.	none detail extensive

Table 324: show chassis satellite-cluster Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface State`	The state of the local and remote interface connection. Port states include: • **online**—the cascade port is online and active. This is the port state during normal operating procedure. • **txUpRxDn**—Tx or Rx forwarding is disabled on the cascade port. This state is often seen when a second aggregation device is added to a Junos Fusion topology, and the devices in the Junos Fusion are synchronizing to the new topology. • **miswire**—the cascade port is miswired. This state is typically seen when a satellite device is interconnected to two aggregation devices but multihoming is not configured. Use **show chassis satellite detail** to gather more information on the issue when the device state is **Miswire**. • **present**—The cascade port recognized the satellite device and is up. • **misconfig**—the cascade port is assigned, but this interface is not working correctly due to a misconfiguration. • **down**—the cascade port is down. • **offline**—the satellite device was previously recognized from this interface, but is no longer present. • **absent**—the cascade port is configured but no satellite device is detected on the interface.	none detail extensive
`Adj Up/Dn Count`	The number of times the satellite device in the satellite device cluster has transitioned to the up or down state.	detail extensive

Table 324: show chassis satellite-cluster Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Last Transition`	The time of the last transition to the up or down state.	detail extensive
`When`	The date and time of the event.	extensive
`Event`	The event.	extensive
`Action`	The actions that resulted from the event.	extensive

Sample Output

show chassis satellite-cluster terse

```
user@aggregation-device> show chassis satellite-cluster terse

Cluster     Provision   Present   Online  Unprovision  Misconfig   Offline
cl1         4           4         4       0            0           0
```

Sample Output

show chassis satellite-cluster

```
user@aggregation-device> show chassis satellite-cluster

Cluster Name: cluster1
    Number of devices provisioned: 4
    Number of devices present: 4
    Number of devices unprovisioned: 0
    Number of devices misconfig/miwired: 0
    Number of devices online: 4
    Number of devices offline: 0
```

```
        Device                      Local        Remote       Interface
Slot    State         Distance      Interface    Interface    State
101     Online        0             xe-101/0/0   ge-0/0/0     online
                                    xe-101/0/1   ge-0/0/1     online
                                    xe-101/0/4   xe-102/0/0   present
                                    xe-101/0/5   xe-102/0/1   present
102     Online        4   [via 101] xe-102/0/0   xe-101/0/4   present
                                    xe-102/0/1   xe-101/0/5   present
                                    xe-102/0/2   xe-103/0/0   present
                                    xe-102/0/3   xe-103/0/1   present
103     Online        8   [via 101] xe-103/0/0   xe-102/0/2   present
                                    xe-103/0/1   xe-102/0/3   present
                                    xe-103/0/2   xe-104/0/4   present
                                    xe-103/0/3   xe-104/0/5   present
104     Online        12  [via 101] xe-104/0/4   xe-103/0/2   present
                                    xe-104/0/5   xe-103/0/3   present
```

Release Information

Command introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

show chassis satellite-cluster route

IN THIS SECTION

- Syntax | 18752
- Description | 18752
- Options | 18752
- Required Privilege Level | 18752
- Output Fields | 18752
- Sample Output | 18753

- Release Information | 18754

Syntax

```
show chassis satellite-cluster route
[cluster cluster-name]
[fpc-slot slot-id]
```

Description

Display information about the route to the aggregation device for a satellite device in a satellite device cluster in a Junos Fusion.

Options

cluster *cluster-name*	Display route information for all satellite devices in the specified cluster.
fpc-slot *slot-id*	Display route information for the satellite device using the specified FPC slot ID.

Required Privilege Level

view

Output Fields

Table 325 on page 18753 lists the output fields for the `show chassis satellite-cluster route` command. Output fields are listed in the approximate order in which they appear.

Table 325: show chassis satellite-cluster route Output Fields

Field Name	Field Description
`Cluster Name`	The name of the satellite device cluster. The name of the satellite device cluster is assigned using the **set chassis satellite-management cluster *cluster-name*** statement.
`Slot`	The slot number of the satellite device. The slot number and the FPC ID are the same number in this context.
`Interface`	Interface on the satellite device that is either interconnected to an aggregation device or another satellite device in the satellite device cluster.
`Transit slot`	The slot number of the satellite device in the satellite device cluster that is passing traffic to the aggregation device for the specified satellite device.
`Distance`	The distance the satellite device is away from the aggregation device. A satellite device in a satellite device cluster that has an uplink port connection to an aggregation device has a distance of 0. The number given in this output counts each hop to a satellite device in the satellite device cluster as 4. For instance, a satellite device that is one aggregation device away from the satellite device with the uplink port has a distance of 4. The `via` output provides the slot number of the satellite device in the satellite device cluster that is directly connected to the aggregation device and passing traffic for the specified satellite device.

Sample Output

show chassis satellite-cluster route

```
user@aggregation-device> show chassis satellite-cluster route

Cluster Name: cluster1

                           Transit
```

```
Slot     Interface        Slot      Distance
101      ge-0/0/1         direct    0
         ge-0/0/0         direct    0
102      ge-0/0/1         101       4
         ge-0/0/0         101       4
103      ge-0/0/1         101       8
         ge-0/0/0         101       8
104      ge-0/0/1         101       12
         ge-0/0/0         101       12
```

Release Information

Command introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

show chassis satellite-cluster statistics

IN THIS SECTION

- Syntax | 18755
- Description | 18755
- Options | 18755
- Required Privilege Level | 18755
- Output Fields | 18755
- Sample Output | 18757
- Release Information | 18759

Syntax

```
show chassis satellite-cluster statistics
[cluster cluster-name]
[fpc-slot slot-id]
```

Description

Display satellite device cluster statistics for satellite devices in satellite device clusters in a Junos Fusion.

Options

cluster *cluster-name* Display satellite device cluster statistics for all satellite devices in the specified cluster.

fpc-slot *slot-id* Display satellite device cluster statistics for the satellite device using the specified FPC slot ID.

Required Privilege Level

view

Output Fields

Table 326 on page 18755 lists the output fields for the `show chassis satellite-cluster statistics` command. Output fields are listed in the approximate order in which they appear.

Table 326: show chassis satellite-cluster statistics Output Fields

Field Name	Field Description
`Cluster Name`	The name of the satellite device cluster. The name of the satellite device cluster is assigned using the **set chassis satellite-management cluster *cluster-name*** statement.
`FPC Slot`	The FPC slot ID number of the satellite device.

Table 326: show chassis satellite-cluster statistics Output Fields *(Continued)*

Field Name	Field Description
`Message Type`	The message type.
`Received`	The number of times the specified message type has been received on the satellite device in the satellite device cluster.
`Sent`	The number of times the specified message type has been sent on the satellite device in the satellite device cluster.
`Rx errors`	The number of times the specified message type has experienced a receive error on the satellite device in the satellite device cluster.
`Adjacency Up`	The total number of Adjacency Up events on the satellite device in the satellite device cluster.
`Adjacency Down`	The total number of Adjacency Down events on the satellite device in the satellite device cluster.
`LLDP Adjacency Up`	The total number of LLDP Adjacency Up events on the satellite device in the satellite device cluster.
`LLDP Adjacency Down`	The total number of LLDP Adjacency Down events on the satellite device in the satellite device cluster.
`SD Adjacency Up`	The total number of satellite device Adjacency Up events on the satellite device in the satellite device cluster.
`SD Adjacency Down`	The total number of satellite device Adjacency Down events on the satellite device in the satellite device cluster.
`Route Add`	The total number of Route Add events on the satellite device in the satellite device cluster.

Table 326: show chassis satellite-cluster statistics Output Fields *(Continued)*

Field Name	Field Description
Route Change	The total number of Route change events on the satellite device in the satellite device cluster.
Route Delete	The total number of Route Delete events on the satellite device in the satellite device cluster.
Provisioned event	The total number of Provisioned events on the satellite device in the satellite device cluster.
Unprovisioned event	The total number of Unprovisioned events on the satellite device in the satellite device cluster.
Delete event	The total number of delete events on the satellite device in the satellite device cluster.
Protocol Session Up	The total number of Protocol Session Up events on the satellite device in the satellite device cluster.
Protocol Session Down	The total number of Protocol Session Down events on the satellite device in the satellite device cluster.

Sample Output

show chassis satellite-cluster statistics cluster cluster1

```
user@aggregation-device> show chassis satellite-cluster statistics cluster cluster1

Cluster Name: cluster1

FPC Slot: 101

Message type                        Received       Sent   Rx errors
Open                                0              1      0
Adjacency UP                        4              0      0
Adjacency Down                      2              0      0
```

```
Sync Complete                       0              0      0
Provision                           0              1      0
Unprovision                         0              0      0
Remote SD Route                     7              0      0
Mode change Request                 0              0      0
Mode change Cancel                  0              0      0
Set Policy                          0              0      0
Reset Policy                        0              0      0
Msg Ack                             0              7      0
Keepalive                           2568           2567   0

General Statistics:

Adjacency Up               1
Adjacency Down             0
LLDP Adjacency Up          2
LLDP Adjacency Down        0
SD Adjacency Up            2
SD Adjacency Down          0
Route Add                  1
Route Change               2
Route Delete               0
Provisioned event          1
Unprovisioned event        0
Delete event               0
Protocol Session Up        1
Protocol Session Down      0

FPC Slot: 102

Message type                        Received       Sent   Rx errors
Open                                0              2      0
Adjacency UP                        8              0      0
Adjacency Down                      0              0      0
Sync Complete                       0              0      0
Provision                           0              2      0
Unprovision                         0              0      0
Remote SD Route                     0              0      0
Mode change Request                 0              0      0
Mode change Cancel                  0              0      0
Set Policy                          0              0      0
Reset Policy                        0              0      0
Msg Ack                             0              0      0
```

```
Keepalive                          2566           2568   0

General Statistics:

Adjacency Up              1
Adjacency Down            0
LLDP Adjacency Up         0
LLDP Adjacency Down       0
SD Adjacency Up           6
SD Adjacency Down         2
Route Add                 2
Route Change              6
Route Delete              1
Provisioned event         1
Unprovisioned event       0
Delete event              0
Protocol Session Up       2
Protocol Session Down     1
```

Release Information

Command introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

show chassis satellite extended-port

IN THIS SECTION

- Syntax | **18760**
- Description | **18760**
- Options | **18760**
- Required Privilege Level | **18761**

- Output Fields | 18761
- Sample Output | 18763
- Release Information | 18764

Syntax

```
show chassis satellite extended-port interface-name
<fpc-slot fpc-slot>
<interface-name interface-name>
[brief | detail | extensive | terse]
<since time>
```

Description

Display the status of the extended ports on the satellite devices in a Junos Fusion.

The extended ports are the network-facing ports on satellite devices that send and receive network traffic for a Junos Fusion.

Options

none	(Same as **brief** and **terse**) Display extended port status information.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
fpc *fpc-slot*	Display extended port status information for the specified FPC slot only. In a Junos Fusion, one FPC slot ID is assigned to each satellite device, so you can use this option to display extended port status information for all extended ports on one satellite device.
interface-name *interface-name*	Display extended port status information for the extended port interface only.
history	Display extended port history.
statistics	Display extended port statistics.

since *time*

(Optional) Display extended port status information for the satellite devices that have been added to the Junos Fusion on or after a certain date or time, which is entered in the *YYYY-MM-DD.HH:MM:SS* format.

To display extended port status information for all satellite devices added since a specified date, enter the specific date as the time as the *time*. For instance, **2015-12-22**.

To display extended port status information for all satellite devices added since a specified time, enter the specific date and time as the *time*. For instance, **2015-12-22.11:01:00**.

Required Privilege Level

view

Output Fields

Table 327 on page 18761 lists the output fields for the `show chassis satellite extended-port` command. Output fields are listed in the approximate order in which they appear.

Table 327: show chassis satellite extended-port Output Fields

Field Name	Field Description	Level of Output
Fields for Interface		
`Name`	The interface name of the extended port.	brief terse detail extensive none
`State`	The state of the extended port.	brief terse detail extensive none

Table 327: show chassis satellite extended-port Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Rx Request State`	The receive request state of the extended port.	brief terse detail extensive none
`Tx Request State`	The transmit request state of the extended port.	brief terse detail extensive none
`Admin State`	The administrative state of the extended port.	brief terse detail extensive none
`Op State`	The operational state of the extended port.	brief terse detail extensive none
`IFD Idx`	The internal interface index.	brief terse detail extensive none
`PCID`	The port's E-channel identifier (ECID), abbreviated as PCID.	brief terse detail extensive none

Table 327: show chassis satellite extended-port Output Fields *(Continued)*

Field Name	Field Description	Level of Output
When	The date and time of the event.	detail extensive
Event	The event.	detail extensive
Action	The actions that resulted from the event.	detail extensive

Sample Output

show chassis satellite extended-port

```
user@aggregation-device> show chassis satellite extended-port
Legend for interface types:
   * -- Uplink interface
                              Rx            Tx            Admin/Op IFD
Name           State          Request State Request State State    Idx   PCID
et-100/0/2     AddComplete    None          Ready         Up/Dn    838   110
et-104/0/2     AddComplete    None          Ready         Up/Dn    813   110
et-107/0/23    AddComplete    None          Ready         Up/Up    544   194
ge-109/0/0     AddComplete    None          Ready         Up/Up    402   115
ge-109/0/1     AddComplete    None          Ready         Up/Dn    403   114
ge-109/0/10    AddComplete    None          Ready         Up/Dn    412   113
ge-109/0/11    AddComplete    None          Ready         Up/Dn    413   112
ge-109/0/12    AddComplete    None          Ready         Up/Dn    414   123
ge-109/0/13    AddComplete    None          Ready         Up/Dn    415   122
ge-109/0/14    AddComplete    None          Ready         Up/Dn    416   125
ge-109/0/15    AddComplete    None          Ready         Up/Dn    417   124
ge-109/0/16    AddComplete    None          Ready         Up/Dn    418   131
ge-109/0/17    AddComplete    None          Ready         Up/Dn    419   130
ge-109/0/18    AddComplete    None          Ready         Up/Dn    420   133
ge-109/0/19    AddComplete    None          Ready         Up/Dn    421   132
ge-109/0/2     AddComplete    None          Ready         Up/Dn    404   117
ge-109/0/20    AddComplete    None          Ready         Up/Dn    422   127
```

```
ge-109/0/21     AddComplete     None            Ready           Up/Dn     423    126
ge-109/0/22     AddComplete     None            Ready           Up/Dn     424    129
ge-109/0/23     AddComplete     None            Ready           Up/Dn     425    128
ge-109/0/24     AddComplete     None            Ready           Up/Dn     426    103
ge-109/0/25     AddComplete     None            Ready           Up/Dn     427    102
ge-109/0/26     AddComplete     None            Ready           Up/Dn     428    105
ge-109/0/27     AddComplete     None            Ready           Up/Dn     429    104
ge-109/0/28     AddComplete     None            Ready           Up/Dn     430    107
ge-109/0/29     AddComplete     None            Ready           Up/Dn     431    106
ge-109/0/3      AddComplete     None            Ready           Up/Dn     405    116
ge-109/0/30     AddComplete     None            Ready           Up/Dn     432    109
ge-109/0/31     AddComplete     None            Ready           Up/Dn     433    108
ge-109/0/32     AddComplete     None            Ready           Up/Dn     434    135
ge-109/0/33     AddComplete     None            Ready           Up/Dn     435    134
ge-109/0/34     AddComplete     None            Ready           Up/Dn     436    137
ge-109/0/35     AddComplete     None            Ready           Up/Dn     437    136
ge-109/0/36     AddComplete     None            Ready           Up/Dn     438    144
ge-109/0/37     AddComplete     None            Ready           Up/Dn     439    143
ge-109/0/38     AddComplete     None            Ready           Up/Dn     440    146
ge-109/0/39     AddComplete     None            Ready           Up/Dn     441    145
ge-109/0/4      AddComplete     None            Ready           Up/Dn     406    119
ge-109/0/40     AddComplete     None            Ready           Up/Dn     442    140
ge-109/0/41     AddComplete     None            Ready           Up/Dn     443    139
ge-109/0/42     AddComplete     None            Ready           Up/Dn     444    142
ge-109/0/43     AddComplete     None            Ready           Up/Dn     445    141
ge-109/0/44     AddComplete     None            Ready           Up/Dn     446    148
ge-109/0/45     AddComplete     None            Ready           Up/Dn     447    147
ge-109/0/46     AddComplete     None            Ready           Up/Dn     448    150
ge-109/0/47     AddComplete     None            Ready           Up/Dn     449    149
ge-109/0/5      AddComplete     None            Ready           Up/Dn     407    118
ge-109/0/6      AddComplete     None            Ready           Up/Dn     408    121
ge-109/0/7      AddComplete     None            Ready           Up/Dn     409    120
ge-109/0/8      AddComplete     None            Ready           Up/Dn     410    111
ge-109/0/9      AddComplete     None            Ready           Up/Dn     411    110
ge-110/0/0      AddComplete     None            Ready           Up/Up     728    115
ge-110/0/1      AddComplete     None            Ready           Up/Dn     729    114
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

show chassis satellite interface

IN THIS SECTION

- Syntax | 18765
- Description | 18765
- Options | 18766
- Required Privilege Level | 18766
- Output Fields | 18766
- Sample Output | 18768
- Sample Output | 18771
- Release Information | 18772

Syntax

```
show chassis satellite interface
<interface-name>
[ brief | detail | extensive ]
<since time>
```

Description

Display the status of the cascade ports as well as the internal satellite interfaces in a Junos Fusion.

You might see **sd** interfaces in the output of this command. These are internal interfaces for the Junos Fusion.

Options

interface-name Specify the name of the interface.

none (Same as **brief**) Display aggregation device interface information.

brief | detail | extensive (Optional) Display the specified level of output.

since *time* (Optional) Display interface status information for the satellite devices that have been added to the Junos Fusion on or after a certain date or time, which is entered in the *YYYY-MM-DD.HH:MM:SS* format.

To display extended port status information for all satellite devices added since a specified date, enter the specific date as the time as the *time*. For instance, **2015-12-22**.

To display extended port status information for all satellite devices added since a specified time, enter the specific date and time as the *time*. For instance, **2015-12-22.11:01:00**.

Required Privilege Level

view

Output Fields

Table 328 on page 18766 lists the output fields for the `show chassis satellite interface` command. Output fields are listed in the approximate order in which they appear.

Table 328: show chassis satellite interface Output Fields

Field Name	Field Description	Level of Output
Fields for Interface		
`Interface`	The interface name.	brief detail extensive none

Table 328: show chassis satellite interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	The state of the interface.	brief detail extensive none
Type	The type of interface.	brief detail extensive none
DF-Role	(Junos Fusion Data Center with EVPN-VXLAN architecture) The designated forwarder (DF) role: • NA—Not applicable. • NON-DF—This aggregation device is not the designated forwarder for the satellite device • DF—This aggregation device is the designated forwarder for the satellite device.	brief detail extensive none
Provisioned Addresses	The provisioned IP addresses for the Junos Fusion. This information is primarily useful for debugging purposes by expert users.	detail extensive
Operational Addresses	The operational IP addresses for the Junos Fusion. This information is primarily useful for debugging purposes by expert users.	detail extensive
When	The date and time of the event.	detail extensive

Table 328: show chassis satellite interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Event	The event.	detail extensive
Action	The actions that resulted from the event.	detail extensive

Sample Output

show chassis satellite interface

```
user@aggregation-device> show chassis satellite interface


Interface         State        Type
lo0               Up           Loopback

sd-101/0/0        Up           Satellite

sd-102/0/0        Up           Satellite

sd-103/0/0        Up           Satellite

sd-104/0/0        Up           Satellite

sd-105/0/0        Up           Satellite

sd-106/0/0        Up           Satellite

sd-107/0/0        Up           Satellite

sd-108/0/0        Up           Satellite

sd-109/0/0        Up           Satellite

sd-110/0/0        Up           Satellite
```

```
sd-111/0/0        Up          Satellite

sd-112/0/0        Up          Satellite

sd-113/0/0        Up          Satellite

sd-114/0/0        Up          Satellite

xe-0/0/1          Up          Cascade

xe-0/0/2          Up          Cascade

xe-0/0/3          Up          Cascade

xe-0/0/4          Up          Cascade

xe-0/0/5          Up          Cascade

xe-0/0/6          Up          Cascade

xe-0/0/7          Up          Cascade

xe-0/0/8          Up          Cascade

xe-0/0/9          Up          Cascade

xe-0/2/0          Up          Cascade

xe-0/2/1          Up          Cascade

xe-0/2/2          Up          Cascade

xe-0/2/3          Up          Cascade

xe-0/2/4          Up          Cascade

xe-0/2/5          Up          Cascade

xe-0/2/6          Up          Cascade

xe-0/2/7          Up          Cascade
```

```
xe-1/0/1          Up          Cascade

xe-1/0/2          Up          Cascade

xe-1/0/3          Up          Cascade

xe-1/2/1          Up          Cascade

xe-1/2/2          Up          Cascade

xe-1/2/3          Up          Cascade

xe-2/0/0          Up          Cascade

xe-2/0/1          Up          Cascade

xe-2/0/2          Up          Cascade

xe-2/0/3          Up          Cascade

xe-2/0/4          Up          Cascade

xe-2/0/5          Up          Cascade

xe-2/0/6          Up          Cascade

xe-2/0/7          Up          Cascade

xe-2/1/0          Up          Cascade

xe-2/1/1          Up          Cascade

xe-2/1/2          Up          Cascade

xe-2/1/3          Up          Cascade

xe-2/1/4          Up          Cascade

xe-2/1/5          Up          Cascade

xe-2/1/6          Up          Cascade

xe-2/1/7          Up          Cascade
```

```
xe-2/2/0          Up           Cascade

xe-2/2/1          Up           Cascade

xe-2/2/2          Up           Cascade

xe-2/2/3          Up           Cascade

xe-2/2/4          Up           Cascade

xe-2/2/5          Up           Cascade

xe-2/2/6          Up           Cascade

xe-2/2/7          Up           Cascade

xe-2/3/0          Up           Cascade

xe-2/3/3          Dn           Cascade

xe-2/3/4          Up           Cascade

xe-2/3/5          Up           Cascade

xe-2/3/6          Up           Cascade

xe-2/3/7          Up           Cascade
```

Sample Output

show chassis satellite interface (Junos Fusion Data Center with EVPN-VXLAN)

```
user@aggregation-device> show chassis satellite interface


Interface         State        Type               DF-Role
lo0               Up           Loopback           NA

sd-101/0/0        Up           Satellite          Non-DF
```

```
sd-102/0/0        Up          Satellite         Non-DF

sd-103/0/0        Up          Satellite         DF

xe-0/0/1          Up          Cascade           NA

xe-0/0/2          Up          Cascade           NA

xe-0/0/3          Up          Cascade           NA

xe-0/0/4          Up          Cascade           NA

xe-0/0/5          Up          Cascade           NA
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

show chassis satellite neighbor

IN THIS SECTION

- Syntax | 18773
- Description | 18773
- Options | 18773
- Required Privilege Level | 18773
- Output Fields | 18773
- Sample Output | 18777
- Release Information | 18779

Syntax

```
show chassis satellite neighbor
[interface-name]
[ brief | detail | extensive | terse]
<since time>
```

Description

Display the status of the satellite device to aggregation device links in a Junos Fusion.

Options

interface-name	Specify the name of the cascade port on the aggregation device.
none	(Same as **terse**) Display satellite device connection information.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
since *time*	(Optional) Display satellite device connection information for the satellite devices that have been added to the Junos Fusion on or after a certain date or time, which is entered in the *YYYY-MM-DD.HH:MM:SS* format. To display satellite device connection information for all satellite devices added since a specified date, enter the specific date as the time as the *time*. For instance, **2015-12-22**. To display satellite device connection information for all satellite devices added since a specified time, enter the specific date and time as the *time*. For instance, **2015-12-22.11:01:00**.

Required Privilege Level

view

Output Fields

Table 329 on page 18774 lists the output fields for the `show chassis satellite neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 329: show chassis satellite neighbor Output Fields

Field Name	Field Description	Level of Output
Fields for Interface		
`Interface`	A cascade port interface on the aggregation device in the Junos Fusion. A cascade port interface on an aggregation device connects to a satellite device in a Junos Fusion.	brief terse detail extensive none
`State`	The state of the interface.	brief terse detail extensive none
`Port Info`	The uplink port interface on the satellite device. An uplink port interface on a satellite device connects the satellite device to an aggregation device in a Junos Fusion.	brief terse detail extensive none
`System Name`	The system name, or alias, of the satellite device. The satellite device's alias is configured using the **set chassis satellite-management fpc *slot-id* alias *alias*** statement.	brief terse detail extensive none
`Model`	The hardware model of the satellite device.	brief terse detail extensive none

Table 329: show chassis satellite neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`SW Version`	The version of satellite software running on the satellite device.	brief terse detail extensive none
`Adjacency up-down transition count`	The number of times that the adjacency has transitioned between up and down.	brief detail extensive
`Last transition`	The last transition of the adjacency state.	brief detail extensive
`Device Serial Number`	The serial number of the satellite device.	brief detail extensive
`Chassis ID`	The chassis ID of the satellite device. The chassis ID of the satellite device is the satellite's device's MAC address. The chassis ID is also specified as the system ID in some Junos Fusion configuration tasks.	brief detail extensive
`Device Family Name`	The device family name.	brief detail extensive
`Version Sequence Number`	The version sequence number.	brief detail extensive
`System Description`	A plain-text description of the hardware and software currently running on the satellite device.	brief detail extensive

Table 329: show chassis satellite neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Build date`	The date and time that the satellite software was built.	brief detail extensive
`Hello interval`	The current hello interval configuration.	brief detail extensive
`Satellite hello interval`	The current satellite device hello interval configuration.	brief detail extensive
`Local-end (Local assigned primary address)`	The local-end cascade port IP address.	brief detail extensive
`Remote-end (Local assigned primary address)`	The remote-end uplink port IP address.	brief detail extensive
`Cause (Adjacency Down History)`	The cause of the last adjacency down event.	brief detail extensive
`Timestamp (Adjacency Down History)`	The date and time of the last adjacency down event.	brief detail extensive
`Information (Adjacency Down History)`	Information related to the last adjacency down event.	brief detail extensive
`When`	The date and time of the event.	detail extensive

Table 329: show chassis satellite neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Event	The event.	detail extensive
Action	The actions that resulted from the event.	detail extensive

Sample Output

show chassis satellite neighbor

```
user@aggregation-device> show chassis satellite neighbor
Interface   State      Port Info   System Name  Model           SW Version
xe-2/3/7    Init
xe-2/3/6    Init
xe-2/3/5    Init
xe-2/3/4    Init
xe-2/3/3    Dn
xe-2/3/0    Two-Way    xe-0/2/2        ex4300-29 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-2/2/7    Two-Way    xe-0/2/2        ex4300-28 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-2/2/6    Two-Way    xe-0/2/2        ex4300-27 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-2/2/5    Two-Way    xe-0/2/2        ex4300-26 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-2/2/4    Init
xe-2/2/3    Init
xe-2/2/2    Two-Way    xe-0/0/48:3 qfx5100-48s-06 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-2/2/1    Two-Way    xe-0/0/48:3 qfx5100-48s-05 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-2/2/0    Init
xe-2/1/7    Init
xe-2/1/6    Init
xe-2/1/5    Two-Way    xe-0/0/4:2  qfx5100-24q-09 QFX5100-24Q-2P 0.1I20150224_18
```

```
27_dc-builder
xe-2/1/4    Two-Way    xe-0/2/1        ex4300-31 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-2/1/3    Two-Way    xe-0/2/1        ex4300-30 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-2/1/2    Two-Way    xe-0/2/1        ex4300-29 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-2/1/1    Two-Way    xe-0/2/1        ex4300-28 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-2/1/0    Init
xe-2/0/7    Two-Way    xe-0/2/1        ex4300-26 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-2/0/6    Init
xe-2/0/5    Init
xe-2/0/4    Init
xe-2/0/3    Init
xe-2/0/2    Two-Way    xe-0/0/48:2 qfx5100-48s-04 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-2/0/1    Two-Way    xe-0/0/48:2 qfx5100-48s-03 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-2/0/0    Init
xe-1/2/3    Two-Way    xe-0/0/0:0  qfx5100-24q-09 QFX5100-24Q-2P 0.1I20150224_18
27_dc-builder
xe-1/2/2    Two-Way    xe-0/2/0        ex4300-31 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-1/2/1    Two-Way    xe-0/2/0        ex4300-30 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-1/0/3    Two-Way    xe-0/2/0        ex4300-29 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-1/0/2    Two-Way    xe-0/2/0        ex4300-28 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-1/0/1    Two-Way    xe-0/2/0        ex4300-27 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-0/2/7    Two-Way    xe-0/0/0:1  qfx5100-24q-09 QFX5100-24Q-2P 0.1I20150224_18
27_dc-builder
xe-0/2/6    Init
xe-0/2/5    Init
xe-0/2/4    Two-Way    xe-0/0/48:1 qfx5100-48s-05 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-0/2/3    Two-Way    xe-0/0/48:1 qfx5100-48s-04 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-0/2/2    Two-Way    xe-0/0/48:1 qfx5100-48s-03 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
```

```
xe-0/2/1    Init
xe-0/2/0    Init
xe-0/0/9    Two-Way    xe-0/2/0        ex4300-26 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-0/0/8    Two-Way    xe-0/2/0        ex4300-25 EX4300-48T     0.1I20150224_182
7_dc-builder
xe-0/0/7    Two-Way    xe-0/0/48:0 qfx5100-48s-07 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-0/0/6    Two-Way    xe-0/0/48:0 qfx5100-48s-06 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-0/0/5    Two-Way    xe-0/0/48:0 qfx5100-48s-05 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-0/0/4    Two-Way    xe-0/0/48:0 qfx5100-48s-04 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-0/0/3    Two-Way    xe-0/0/48:0 qfx5100-48s-03 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-0/0/2    Two-Way    xe-0/0/48:0 qfx5100-48s-02 QFX5100-48S-6Q 0.1I20150224_18
27_dc-builder
xe-0/0/1    Init
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

show chassis satellite power-budget-statistics

IN THIS SECTION

- Syntax | 18780
- Description | 18780
- Options | 18780

- Required Privilege Level | 18780
- Output Fields | 18780
- Sample Output | 18782
- Release Information | 18783

Syntax

```
show chassis satellite power-budget-statistics
<slot-id slot-id-number>
```

Description

Display the power budget statistics of a satellite device or devices in a Junos Fusion.

Options

none	Display power budget statistics for all satellite devices in the Junos Fusion.
slot-id *slot-id-number*	Display power budget statistics for the specified satellite device only. The *slot-id-number* and the FPC ID are the same number in a Junos Fusion.

Required Privilege Level

view

Output Fields

Table 330 on page 18781 lists the output fields for the `show chassis satellite-management power-budget-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 330: show chassis satellite-management power-budget-statistics Output Fields

Field Name	Field Description
`FPC` *n*	The FPC slot ID number in the Junos Fusion, where *n* Is the FPC slot ID. The FPC slot ID and the satellite device number are the same thing in a Junos Fusion.
`PSU` *n* (*supply type*)	Capacity rating of the power supply and whether the power supply is currently operating (`Online`) or not (`Offline`). If a power supply is offline, the capacity is shown as 0 W.
`Total Power supplied by all Online PSUs`	Total number of watts supplied by all currently operating power supplies for the satellite device.
`Power Redundancy Configuration`	Configured power redundancy setting, either *N*+1 or *N* +*N*.
`Base power reserved`	Total number of watts reserved for the satellite device.
`Non-PoE power being consumed`	The amount of power, in W, currently being consumed for functions other than PoE by the satellite device.
`Total Power allocated for PoE`	The total of the PoE power budgets allocated to the satellite device.
`Total PoE power consumed`	The amount of power that has been consumed by PoE by the satellite device.
`Total PoE power remaining`	The amount of available power remaining that can be used for PoE on the satellite device.

Sample Output

show chassis satellite power-budget-statistics

```
user@aggregation-device> show chassis satellite power-budget-statistics
fpc 100:
--------------------------------------------------------------------------
            PSU 0    (JPSU-550-DC-AFI    )           :     550 W   Online
            PSU 1    (JPSU-550-DC-AFO    )           :     550 W   Online
            Power redundancy configuration           :     N+N
            Total power supplied by all online PSUs  :     522 W
            Base power reserved                      :     175 W
            Non-PoE power being consumed             :      82 W
            Total power allocated for PoE            :     347 W
            Total PoE power consumed                 :       0 W
            Total PoE power remaining                :     347 W
fpc 120:
--------------------------------------------------------------------------
            Power redundancy configuration           :     N+N
            Total power supplied by all online PSUs  :     170 W
            Base power reserved                      :       0 W
            Non-PoE power being consumed             :       0 W
fpc 128:
--------------------------------------------------------------------------
            Power redundancy configuration           :     N+N
            Total power supplied by all online PSUs  :       0 W
            Base power reserved                      :       0 W
            Non-PoE power being consumed             :       0 W
fpc 133:
--------------------------------------------------------------------------
            PSU 0                        )           :       0 W  Offline
            PSU 1    (JPSU-1100-AC-AFO   )           :    1100 W   Online
            Power redundancy configuration           :     N+N
            Total power supplied by all online PSUs  :    1100 W
            Base power reserved                      :     175 W
            Non-PoE power being consumed             :      74 W
            Total power allocated for PoE            :     925 W
            Total PoE power consumed                 :       0 W
            Total PoE power remaining                :     925 W
fpc 240:
--------------------------------------------------------------------------
            Power redundancy configuration           :     N+N
```

```
            Total power supplied by all online PSUs  :       0 W
            Base power reserved                      :       0 W
            Non-PoE power being consumed             :       0 W
```

Release Information

Command introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

> *Understanding Power over Ethernet in a Junos Fusion*

show chassis satellite redundancy-group

IN THIS SECTION

- Syntax | **18783**
- Description | **18784**
- Options | **18784**
- Required Privilege Level | **18784**
- Output Fields | **18784**
- Sample Output | **18786**
- Release Information | **18786**

Syntax

```
show chassis satellite redundancy-group
[brief | detail | extensive | terse]
<since time>
```

Description

Display the status of the redundancy group in a Junos Fusion.

Redundancy groups are used in a Junos Fusion to configure, monitor, and maintain a topology using two aggregation devices.

Options

none (Same as **terse** and **brief**) Display redundancy group information for the Junos Fusion.

brief | detail | extensive | terse (Optional) Display the specified level of output.

since *time* (Optional) Display redundancy group information on or after a certain date or time, in *YYYY-MM-DD.HH:MM:SS* format.

To display all output since a specified date, enter the specific date. For instance, enter **2015-12-22** as the *time*.

To display all output since a specified time, enter the specific date and time. For instance, to display all information on or after 11:01AM on December 22nd, 2015, enter **2015-12-22.11:01:00** as the *time*.

Required Privilege Level

view

Output Fields

Table 331 on page 18785 lists the output fields for the `show chassis satellite redundancy-groups` command. Output fields are listed in the approximate order in which they appear.

Table 331: show chassis satellite redundancy-groups Output Fields

Field Name	Field Description	Level of Output
`Name`	The name of the redundancy group. The redundancy group name is set using the **set chassis satellite-management redundancy-groups** ***redundancy-group-name*** statement.	none terse brief detail extensive
`Cluster State`	The cluster state.	none terse brief detail extensive
`Peer Chassis ID`	The chassis ID of the peer chassis. In a Junos Fusion using redundancy groups, the peer chassis ID is the chassis ID of the other aggregation device and is required to create an interchassis link (ICL). The chassis ID is set using the **set** ***redundancy-group-name*** **chassis-id** ***chassis-id-number*** statement.	none terse brief detail extensive
`Peer Chassis SN`	The serial number of the peer chassis. In a Junos Fusion using redundancy groups, the peer chassis serial number is the serial number of the other aggregation device.	none terse brief detail extensive

Table 331: show chassis satellite redundancy-groups Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Device Count`	The device count.	none terse brief detail extensive
`When`	The date and time of the event.	detail extensive
`Event`	The event.	detail extensive
`Action`	The actions that resulted from the event.	detail extensive

Sample Output

show chassis satellite redundancy-group

```
user@aggregation-device> show chassis satellite redundancy-group

            Cluster              Peer          Peer               Device
Name        State                Chassis ID    Chassis SN         Count
gr1         Online               2             DC334              143/143/150
```

Release Information

Command introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

show chassis satellite redundancy-group devices

IN THIS SECTION

- Syntax | 18787
- Description | 18787
- Options | 18787
- Required Privilege Level | 18788
- Output Fields | 18788
- Sample Output | 18789
- Release Information | 18790

Syntax

```
show chassis satellite redundancy-group devices
[brief | detail | extensive | terse]
<history>
```

Description

Display the status of the devices in a redundancy group in a Junos Fusion.

Redundancy groups are used in a Junos Fusion to configure, monitor, and maintain a topology using two aggregation devices.

Options

none	(Same as **terse** and **brief**) Display redundancy group device information for the Junos Fusion.

brief | detail | extensive | terse	(Optional) Display the specified level of output.
history	(Optional) Display historical output.

Required Privilege Level

view

Output Fields

Table 332 on page 18788 lists the output fields for the `show chassis satellite redundancy-groups` command. Output fields are listed in the approximate order in which they appear.

Table 332: show chassis satellite redundancy-groups Output Fields

Field Name	Field Description	Level of Output
`Cluster Name`	The name of the redundancy group. The redundancy group name is set using the **set chassis satellite-management redundancy-groups** ***redundancy-group-name*** statement.	none terse brief detail extensive
`Slot ID`	The FPC slot ID of the satellite device.	none terse brief detail extensive

Table 332: show chassis satellite redundancy-groups Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Local State`	The local state of the satellite device.	none terse brief detail extensive
`Peer State`	The peer state.	none terse brief detail extensive
`When`	The date and time of the event.	detail extensive
`Event`	The event.	detail extensive
`Action`	The actions that resulted from the event.	detail extensive

Sample Output

show chassis satellite redundancy-group devices

```
user@aggregation-device> show chassis satellite redundancy-group devices

Cluster                         Local           Peer
```

```
name            Slot-ID          State           State
gr1             100              online          online
gr1             101              online          online
gr1             102              online          online
gr1             103              not-provisioned online
gr1             104              not-provisioned online
gr1             105              not-provisioned online
gr1             106              not-provisioned online
gr1             107              not-provisioned online
gr1             108              not-provisioned online
gr1             109              not-provisioned online
```

Release Information

Command introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

show chassis satellite software

IN THIS SECTION

- Syntax | 18791
- Description | 18791
- Options | 18791
- Required Privilege Level | 18791
- Output Fields | 18791
- Sample Output | 18792
- Sample Output | 18793
- Release Information | 18793

Syntax

```
show chassis satellite software
[ brief | detail ]
```

Description

Display information related to the satellite software in the Junos Fusion.

Options

none (Same as **brief**) Display satellite device software information.

brief | detail (Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 333 on page 18791 lists the output fields for the `show chassis satellite neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 333: show chassis satellite software Output Fields

Field Name	Field Description	Level of Output
Fields for Interface		
`Version`	The versions of satellite software that are installed and associated with a software upgrade group.	brief none
`Platforms`	The hardware platform information.	brief none

Table 333: show chassis satellite software Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Group`	The name of the assigned satellite software group or groups, if assigned.	brief none
`Software Package Version`	The satellite software package version.	detail
`Platform`	The platform type.	detail
`Host Version`	The host version of software for the platform.	detail
`Current Groups`	The name or names of the satellite software upgrade groups that are using the software package. This output only appears if the software package is associated with a satellite software upgrade group.	detail
`Former Groups`	The name or names of satellite software upgrade groups that were previously using the software package. This output only appears if the software package was previously associated with a satellite software upgrade group.	detail

Sample Output

show chassis satellite software

```
user@aggregation-device> show chassis satellite software
Version                         Platforms           Group
3.0R1.1                         i386 ppc            group0
```

Sample Output

show chassis satellite software detail

```
user@aggregation-device> show chassis satellite software detail
Software package version: 3.0R1.6
Platforms supported by package: i386 ppc arm arm563xx
  Platform      Host Version  Models Supported
  i386          3.0.3         QFX5100-24Q-2P
                              QFX5100-48C-6Q
                              QFX5100-48S-6Q
                              QFX5100-48T-6Q
                              QFX5100-96S-8Q
                              QFX5100-48SH-6Q
                              QFX5100-48TH-6Q
  ppc           1.1.2         EX4300-24P
                              EX4300-24T
                              EX4300-48P
                              EX4300-48T
                              EX4300-48T-BF
                              EX4300-48T-DC
                              EX4300-48T-DC-BF
  arm           1.0.0         EX2300-24P
                              EX2300-24T-DC
                              EX2300-C-12T
                              EX4300-C-12P
  arm563xx      1.0.0         EX3400-24P
                              EX3400-24T
                              EX3400-48T
                              EX3400-48P
Current Groups: group1
                group2
                group3
                group4
                group5
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

show chassis satellite statistics

IN THIS SECTION

- Syntax | 18794
- Description | 18794
- Options | 18795
- Required Privilege Level | 18795
- Output Fields | 18795
- Sample Output | 18796
- Sample Output | 18800
- Sample Output | 18801
- Release Information | 18801

Syntax

```
show chassis satellite statistics
<device-alias device-alias>
<fpc-slot fpc-slot>
<cluster cluster-name>
```

Description

Display statistics for satellite devices in a Junos Fusion.

Options

device-alias Display output for the specified satellite device, which is identified by the device alias, only.

fpc-slot Display output for the specified satellite device, which is identified by the FPC slot ID, only.

cluster-name Display output for the satellite devices in the specified satellite device cluster only.

Required Privilege Level

view

Output Fields

Table 334 on page 18795 lists the output fields for the `show chassis satellite statistics` command. Output fields are listed in the approximate order in which they appear.

Table 334: show chassis satellite statistics Output Fields

Field Name	Field Description
Fields for Interface	
`Serial Number`	The serial number of the satellite device.
`Slot-ID`	The FPC slot ID of the satellite device.
`CSP down transition count`	The number of times that the Control and Status Protocol (CSP) session has gone down.
`Last transition (CSP down transition count)`	The last time that the Control and Status Protocol (CSP) session transitioned.
`Reachability down transition count`	The number of times the satellite device has been in the reachability down state.

Table 334: show chassis satellite statistics Output Fields *(Continued)*

Field Name	Field Description
`Reachability change transition count (Reachability down transition count)`	The number of times that the satellite device's reachability state has transitioned.
`S/W image update count`	The number of times that the satellite software has been updated on the satellite device.
`Extended Port add/delete/up/down request/response`	The number of times an extended port—a network-facing port on the satellite device—has been added, deleted, placed in the up position, received a down request, or received a response.
`Extended Port Params change request`	The number of times that an extended port—a network-facing port on the satellite device—has had a change request.
`Extended Port up/down operational state transition`	The number of times that an extended port—a network-facing port on the satellite device—has had an operational state transition to up or down.
`Rx sync complete`	The number of times the receive synchronization state has been completed.
`Uplink ready rx count`	The number of times the uplink port—the port on the satellite device that connects to the aggregation device—has been placed in the ready-to-receive state.
`Uplink ready tx count`	The number of times the uplink port—the port on the satellite device that connects to the aggregation device—has been placed in the ready-to-transmit state.

Sample Output

show chassis satellite statistics

```
user@aggregation-device> show chassis satellite statistics
Serial Number: TA3714160468 Slot-ID: 101
```

```
   CSP down transition count: 0 Last transition: 05:23:56
   Reachability down transition count: 0
   Reachability change transition count: 4 Last transition: 05:23:16
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 7/0/7/2 7/0/7/2
   Extended Port Params change request: 0
   Extended Port up/down operational state transition: 7/0
   Rx sync complete: 1
   Uplink ready rx count: 4
   Uplink ready tx count: 4
Serial Number: TA3714160046 Slot-ID: 102
   CSP down transition count: 0 Last transition: 05:23:55
   Reachability down transition count: 0
   Reachability change transition count: 4 Last transition: 05:23:19
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 7/0/7/2 7/0/7/2
   Extended Port Params change request: 0
   Extended Port up/down operational state transition: 7/0
   Rx sync complete: 1
   Uplink ready rx count: 4
   Uplink ready tx count: 4
Serial Number: TA3714140404 Slot-ID: 103
   CSP down transition count: 0 Last transition: 05:23:57
   Reachability down transition count: 0
   Reachability change transition count: 4 Last transition: 05:23:14
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 6/0/5/3 6/0/5/3
   Extended Port Params change request: 0
   Extended Port up/down operational state transition: 5/0
   Rx sync complete: 1
   Uplink ready rx count: 4
   Uplink ready tx count: 4
Serial Number: TA3714141327 Slot-ID: 104
   CSP down transition count: 0 Last transition: 05:23:57
   Reachability down transition count: 0
   Reachability change transition count: 4 Last transition: 05:23:15
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 14/0/14/2 14/0/14/2
   Extended Port Params change request: 0
   Extended Port up/down operational state transition: 14/0
   Rx sync complete: 1
   Uplink ready rx count: 4
   Uplink ready tx count: 4
```

```
Serial Number: TA3714140200 Slot-ID: 105
   CSP down transition count: 0 Last transition: 05:23:59
   Reachability down transition count: 0
   Reachability change transition count: 4 Last transition: 05:23:15
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 18/0/18/2 18/0/18/2
   Extended Port Params change request: 6
   Extended Port up/down operational state transition: 18/0
   Rx sync complete: 1
   Uplink ready rx count: 4
   Uplink ready tx count: 4
Serial Number: TA3714140904 Slot-ID: 106
   CSP down transition count: 0 Last transition: 05:23:57
   Reachability down transition count: 0
   Reachability change transition count: 4 Last transition: 05:23:16
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 17/0/16/3 17/0/16/3
   Extended Port Params change request: 2
   Extended Port up/down operational state transition: 16/0
   Rx sync complete: 1
   Uplink ready rx count: 4
   Uplink ready tx count: 4
Serial Number: PE3714040197 Slot-ID: 107
   CSP down transition count: 0 Last transition: 05:24:32
   Reachability down transition count: 0
   Reachability change transition count: 4 Last transition: 05:23:18
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 52/0/7/50 52/0/7/50
   Extended Port Params change request: 0
   Extended Port up/down operational state transition: 7/1
   Rx sync complete: 1
   Uplink ready rx count: 4
   Uplink ready tx count: 4
Serial Number: PE3714080398 Slot-ID: 108
   CSP down transition count: 0 Last transition: 05:24:32
   Reachability down transition count: 0
   Reachability change transition count: 4 Last transition: 05:23:18
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 52/0/15/40 52/0/15/40
   Extended Port Params change request: 0
   Extended Port up/down operational state transition: 15/0
   Rx sync complete: 1
   Uplink ready rx count: 4
```

```
   Uplink ready tx count: 4
Serial Number: PE3714080103 Slot-ID: 109
   CSP down transition count: 0 Last transition: 05:23:22
   Reachability down transition count: 0
   Reachability change transition count: 3 Last transition: 05:23:19
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 51/0/14/37 51/0/14/37
   Extended Port Params change request: 51
   Extended Port up/down operational state transition: 14/0
   Rx sync complete: 1
   Uplink ready rx count: 3
   Uplink ready tx count: 3
Serial Number: PE3714090246 Slot-ID: 110
   CSP down transition count: 0 Last transition: 05:23:22
   Reachability down transition count: 0
   Reachability change transition count: 3 Last transition: 05:23:19
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 51/0/14/37 51/0/14/37
   Extended Port Params change request: 42
   Extended Port up/down operational state transition: 14/0
   Rx sync complete: 1
   Uplink ready rx count: 3
   Uplink ready tx count: 3
Serial Number: PE3714080417 Slot-ID: 111
   CSP down transition count: 0 Last transition: 05:23:22
   Reachability down transition count: 0
   Reachability change transition count: 3 Last transition: 05:23:19
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 51/0/13/38 51/0/13/38
   Extended Port Params change request: 51
   Extended Port up/down operational state transition: 13/0
   Rx sync complete: 1
   Uplink ready rx count: 3
   Uplink ready tx count: 3
Serial Number: PE3714080018 Slot-ID: 112
   CSP down transition count: 0 Last transition: 05:23:22
   Reachability down transition count: 0
   Reachability change transition count: 2 Last transition: 05:23:18
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 51/0/12/39 51/0/12/39
   Extended Port Params change request: 51
   Extended Port up/down operational state transition: 12/0
   Rx sync complete: 1
```

```
   Uplink ready rx count: 2
   Uplink ready tx count: 2
Serial Number: PE3714080030 Slot-ID: 113
   CSP down transition count: 0 Last transition: 05:23:22
   Reachability down transition count: 0
   Reachability change transition count: 3 Last transition: 05:23:18
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 51/0/13/38 51/0/13/38
   Extended Port Params change request: 51
   Extended Port up/down operational state transition: 13/0
   Rx sync complete: 1
   Uplink ready rx count: 3
   Uplink ready tx count: 3
Serial Number: TB3714070145 Slot-ID: 114
   CSP down transition count: 0 Last transition: 05:23:58
   Reachability down transition count: 0
   Reachability change transition count: 4 Last transition: 05:23:15
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 17/0/13/7 17/0/13/7
   Extended Port Params change request: 0
   Extended Port up/down operational state transition: 13/0
   Rx sync complete: 1
   Uplink ready rx count: 4
   Uplink ready tx count: 4
```

Sample Output

show chassis satellite statistics device-alias qfx5100-48s-02

```
user@aggregation-device> show chassis satellite statistics device-alias qfx5100-48s-02
Serial Number: TA3714160468 Slot-ID: 101
   CSP down transition count: 0 Last transition: 05:52:44
   Reachability down transition count: 0
   Reachability change transition count: 4 Last transition: 05:52:04
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 7/0/7/2 7/0/7/2
   Extended Port Params change request: 0
   Extended Port up/down operational state transition: 7/0
   Rx sync complete: 1
```

```
   Uplink ready rx count: 4
   Uplink ready tx count: 4
```

Sample Output

show chassis satellite statistics fpc-slot 101

```
user@aggregation-device> show chassis satellite statistics fpc-slot 101
Serial Number: TA3714160468 Slot-ID: 101
   CSP down transition count: 0 Last transition: 05:52:44
   Reachability down transition count: 0
   Reachability change transition count: 4 Last transition: 05:52:04
   S/W image update count: 0
   Extended Port add/delete/up/down request/response: 7/0/7/2 7/0/7/2
   Extended Port Params change request: 0
   Extended Port up/down operational state transition: 7/0
   Rx sync complete: 1
   Uplink ready rx count: 4
   Uplink ready tx count: 4
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

show chassis satellite unprovision

IN THIS SECTION

Syntax | 18802

- Description | 18802
- Options | 18802
- Required Privilege Level | 18803
- Output Fields | 18803
- Sample Output | 18806
- Sample Output | 18806
- Release Information | 18808

Syntax

```
show chassis satellite unprovision
[brief | detail | extensive]
[cluster cluster-name]
<since time>
```

Description

Display information about unprovisioned satellite devices in a Junos Fusion.

An unprovisioned satellite device in a Junos Fusion is a satellite device that is recognized by the aggregation device, but is not participating in a Junos Fusion.

No output appears when this command is entered when a Junos Fusion contains no unprovisioned satellite devices.

This command is helpful in identifying satellite devices that are not participating in a Junos Fusion due to configuration issues. Notably, a satellite device that has not been associated with an FPC ID in a Junos Fusion becomes an unprovisioned satellite device. See *Configuring Junos Fusion Provider Edge* or Configuring or Expanding a Junos Fusion Enterprise for information on associating an FPC ID with a Junos Fusion.

Options

none	(Same as **brief**) Display unprovisioned satellite device information.
brief | detail | extensive	(Optional) Display the specified level of output.

cluster *cluster-name*
(Optional) Display unprovisioned satellite device information for the specified satellite device cluster only.

since *time*
(Optional) Display unprovisioned satellite device information for the satellite devices that have been unprovisioned from a Junos Fusion on or after a certain date or time, which is entered in the *YYYY-MM-DD.HH:MM:SS* format.

To display unprovisioned satellite device information for all satellite devices unprovisioned since a specified date, enter the specific date as the time as the *time*. For instance, **2015-12-22**.

To display unprovisioned satellite device information for all satellite devices added since a specified time, enter the specific date and time as the *time*. For instance, **2015-12-22.11:01:00**.

Required Privilege Level

view

Output Fields

Table 335 on page 18803 lists the output fields for the `show chassis satellite unprovision` command. Output fields are listed in the approximate order in which they appear.

Table 335: show chassis satellite unprovision Output Fields

Field Name	Field Description	Level of Output
`System-Id`	The MAC address of the satellite device.	brief extensive none
`Serial-Number`	The serial number of the satellite device.	brief extensive none
`Device State`	The device state of the unprovisioned satellite device.	brief extensive none

Table 335: show chassis satellite unprovision Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Cascade Ports`	The cascade ports on the aggregation device that are connected to the satellite device.	brief extensive none
`Port State`	The port state of the cascade port.	brief extensive none
`Operational State`	The operational state of the satellite device.	detail
`Product Model`	The product model of the satellite device.	detail
`Product Family`	The product family of the satellite device.	detail
`Serial number`	The serial number of the satellite device.	detail
`System id`	The MAC address of the satellite device.	detail
`Software package version`	The satellite software package version running on the satellite device.	detail
`Host software version`	The host software version.	detail
Fields for Cascade interfaces		
`Interface Name`	The interface name of the cascade port on the aggregation device.	detail
`State`	The state of the cascade port.	detail

Table 335: show chassis satellite unprovision Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Uplink Interface`	The uplink interface name. The uplink interface is the interface on the satellite device that connects to the aggregation device.	detail
`Adjacency State`	The adjacency state of the uplink interface to cascade port link.	detail
`Last transition`	The amount of time that has passed since the last link transition.	detail
`Adjacency down count`	The number of times that the uplink interface to cascade port link has gone into the adjacency down count.	detail
`Rx Packet`	The number of received packets.	detail
`Last received packet`	The amount of time that has passed since the last received packet.	detail
`Peer adjacency information`	The amount of time that the adjacency has been active.	detail
`Last down cause`	The cause of the last time the adjacency went down.	detail
`SDPD restart detected`	The number of times that the SDPD has restarted.	detail
Fields for process information		
`Process Name`	The name of the process.	detail
`PID`	The PID of the process.	detail

Table 335: show chassis satellite unprovision Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	The current state of the process.	detail
Number of restart detected	The number of times that the process has restarted.	detail
Uptime	The amount of time that the process has been active.	detail
When	The date and time of the event.	extensive
Event	The event.	extensive
Action	The actions that resulted from the event.	extensive

Sample Output

show chassis satellite unprovision

```
user@aggregation-device> show chassis satellite unprovision
                                    Device          Cascade    Port
System-Id           Serial-Number   State           Ports      State
AA:BB:CC:aa:bb:cc   TABCDE111111    Present         xe-0/0/1   present
                                                    xe-0/1/2   present
AA:BB:CC:aa:bb:zz   PABCDE111111    Present         xe-0/0/2   present
                                                    xe-0/3/2   present
```

Sample Output

show chassis satellite unprovision detail

```
user@aggregation-device> show chassis satellite unprovision detail
Operational State: Present
```

```
Product Model: QFX5100-48S-6Q
Product Family: i386
Serial number: TABCDE111111
System id: AA:BB:CC:aa:bb:cc
Software package version: 3.0R1
Host software version: 0.2.3
Cascade interfaces:
    Interface Name: xe-0/0/1 State: present
        Uplink Interface: xe-0/0/25
        Adjacency state: Two-Way
        Last transition: 3d 22:06:55
        Adjacency down count: 0
        Rx Packet: 33875 Last received packet: 00:00:09
        Peer adjacency information: 3d 22:06:55
            Adjacency down count: 3
            Last down cause: TTL is 0
            SDPD restart detected: 3
    Interface Name: xe-0/1/2 State: present
        Uplink Interface: xe-0/0/24
        Adjacency state: Two-Way
        Last transition: 3d 22:06:58
        Adjacency down count: 0
        Rx Packet: 33875 Last received packet: 00:00:09
        Peer adjacency information: 3d 22:06:58
            Adjacency down count: 5
            Last down cause: TTL is 0
            SDPD restart detected: 3
Process information:
    Process Name: Provisioning PID: 2488 State: Running
        Number of restart detected: 0
        Uptime: 3d 22:06:58
    Process Name: PFE PID: 2631 State: Running
        Number of restart detected: 0
        Uptime: 3d 22:06:58
Operational State: Present
Product Model: EX4300-48T
Product Family: ppc
Serial number: PABCDE111111
System id: AA:BB:CC:aa:bb:zz
Software package version: 3.0R1
Host software version: 0.2.4
Cascade interfaces:
    Interface Name: xe-0/0/2 State: present
```

```
        Uplink Interface: xe-0/2/1
        Adjacency state: Two-Way
        Last transition: 3d 22:06:56
        Adjacency down count: 0
        Rx Packet: 33876 Last received packet: 00:00:05
        Peer adjacency information: 3d 22:06:56
            Adjacency down count: 1
            Last down cause: TTL is 0
            SDPD restart detected: 2
    Interface Name: xe-0/3/2 State: present
        Uplink Interface: xe-0/2/0
        Adjacency state: Two-Way
        Last transition: 3d 22:06:57
        Adjacency down count: 0
        Rx Packet: 33876 Last received packet: 00:00:05
        Peer adjacency information: 3d 22:06:57
            Adjacency down count: 3
            Last down cause: TTL is 0
            SDPD restart detected: 2
Process information:
    Process Name: Provisioning PID: 1603 State: Running
        Number of restart detected: 0
        Uptime: 3d 22:06:57
    Process Name: PFE PID: 1615 State: Running
        Number of restart detected: 0
        Uptime: 3d 22:06:57
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

show chassis satellite upgrade-group

IN THIS SECTION

- Syntax | 18809
- Description | 18809
- Options | 18810
- Required Privilege Level | 18810
- Output Fields | 18810
- Sample Output | 18812
- Sample Output | 18812
- Release Information | 18813

Syntax

```
show chassis satellite upgrade-group
<upgrade-group-name>
[ brief | detail | extensive | terse]
```

Description

Display information about the satellite software upgrade groups for the Junos Fusion.

A satellite software upgrade group is a group of satellite devices that are updated at the same time to the same version of the satellite software. One Junos Fusion can contain multiple software upgrade groups, and multiple software upgrade groups should be configured in most Junos Fusions to avoid network downtimes during satellite software installations.

A satellite software upgrade group that contains all satellite devices in a satellite device cluster is automatically created when a satellite device cluster is configured. The software upgrade group name for these automatically created software upgrade groups is the cluster name.

Options

none	(Same as **brief** and **terse**) Display satellite software upgrade group information for all satellite software upgrade groups.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
upgrade-group-name	Display satellite software upgrade group information for the specified satellite software upgrade group only. The satellite software upgrade group name is set using the **set chassis satellite-management upgrade-groups *upgrade-group-name*** statement for standalone satellite devices and is the cluster name for satellite device clusters.

Required Privilege Level

view

Output Fields

Table 336 on page 18810 lists the output fields for the `show chassis satellite upgrade-group` command. Output fields are listed in the approximate order in which they appear.

Table 336: show chassis satellite upgrade-group Output Fields

Field Name	Field Description	Level of Output
Fields for Interface		
`Group`	The satellite software upgrade group name. The satellite software upgrade group name is the name of the satellite device cluster when used with a satellite device cluster. A satellite software upgrade group with the name of the satellite device cluster is created automatically when a satellite device cluster is configured. The satellite software upgrade group name is set using the **set chassis satellite-management upgrade-groups *upgrade-group-name*** statement for standalone satellite devices.	brief terse extensive none

Table 336: show chassis satellite upgrade-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Sw-Version`	The version of satellite software associated with the satellite software upgrade group.	brief terse extensive none
`Group State`	The state of the satellite software upgrade group.	brief terse extensive none
`Slot`	The FPC slot identification number of the satellite device that is a member of the satellite software upgrade group.	brief terse detail extensive none
`Device State`	The state of the satellite software for the specified member of the satellite software upgrade group. The `version-in-sync` output appears when the satellite device is running the satellite software version that is associated with the satellite software upgrade group.	brief terse detail extensive none
`Software upgrade group`	The name of the satellite software upgrade group.	detail
`Software package version`	The satellite software package associated with the satellite software upgrade group.	detail
`Previous software package version`	The satellite software package that was previously associated with the satellite software upgrade group. This output only appears if the satellite software upgrade group was previously associated with another version of satellite software.	detail

Sample Output

show chassis satellite upgrade-group

```
user@aggregation-device> show chassis satellite upgrade-group
                                                Group              Device
Group         Sw-Version                        State      Slot    State
__ungrouped__
ex4300        3.0R1.0                           in-sync    107     version-in-sync
                                                           108     version-in-sync
                                                           109     version-in-sync
                                                           110     version-in-sync
                                                           111     version-in-sync
                                                           112     version-in-sync
                                                           113     version-in-sync
qfx           3.0R1.0                           in-sync    102     version-in-sync
                                                           103     version-in-sync
                                                           104     version-in-sync
                                                           105     version-in-sync
                                                           106     version-in-sync
                                                           114     version-in-sync
```

Sample Output

show chassis satellite upgrade-group detail

```
user@aggregation-device> show chassis satellite upgrade-group detail
Software upgrade group: ex4300
Software package version: 3.0R1.0
Previous software package version: 3.0R1.1
  Slot    Device State
  107     version-in-sync
  108     version-in-sync
  109     version-in-sync
  110     version-in-sync
  111     version-in-sync
  112     version-in-sync
  113     version-in-sync

Software upgrade group: qfx
```

```
Software package version: 3.0R1.0
  Slot    Device State
  102     version-in-sync
  103     version-in-sync
  104     version-in-sync
  105     version-in-sync
  106     version-in-sync
  114     version-in-sync
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

show chassis scb

IN THIS SECTION

- Syntax | 18814
- Description | 18814
- Options | 18814
- Required Privilege Level | 18814
- Output Fields | 18814
- Sample Output | 18815
- Release Information | 18816

Syntax

```
show chassis scb
```

Description

(M40 router only) Display System Control Board (SCB) status information.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 337 on page 18814 lists the output fields for the `show chassis scb` command. Output fields are listed in the approximate order in which they appear.

Table 337: show chassis scb Output Fields

Field Name	Field Description
Temperature	Temperature of the air passing by the SCB, in degrees Celsius.
CPU utilization	Total percentage of CPU being used by the SCB's processor.
Interrupt utilization	Of the total CPU being used by the SCB's processor, the percentage being used for interrupts.
Heap utilization	Percentage of heap space being used by the SCB's processor.
Buffer utilization	Percentage of buffer space being used by the SCB's processor.
DRAM	Total DRAM available to the SCB's processor.

Table 337: show chassis scb Output Fields *(Continued)*

Field Name	Field Description
Start time	Time when the SCB started running.
Uptime	How long the SCB has been running.
Internet Processor memory	Information about the memory of the Internet Processor ASIC on the SCB: • **IP routes**—Number of IP routes known to the Internet Processor. • **MPLS routes**—Number of MPLS routes known to the Internet Processor. • **SRAM banks enabled**—Which SRAM banks are enabled. • **SRAM size**—Size of SCB SRAM, in bytes. • **SRAM used**—Amount of SRAM used, in bytes. • **SRAM utilization**—Percentage of SRAM used.

Sample Output

show chassis scb

```
user@host> show chassis scb
SCB status:
  Temperature:                30 Centigrade
  CPU utilization:             5 percent
  Interrupt utilization:       0 percent
  Heap utilization:            0 percent
  Buffer utilization:          2 percent
  DRAM:                       64 Mbytes
  Start time:                    1998-10-28 18:35:46 UTC
  Uptime:                        6 minutes, 16 seconds
  Internet Processor memory:
    IP routes:                16
    MPLS routes:               1
    SRAM banks enabled:      [ 1 1 1 1 ]
    SRAM size:                 4 Mbytes
```

```
    SRAM used:                256 bytes
    SRAM utilization:           0 percent
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Checklist for Monitoring the SCB

show chassis scb-ha-port

IN THIS SECTION

- Syntax | 18816
- Description | 18816
- Required Privilege Level | 18817
- Output Fields | 18817
- Sample Output | 18819
- Release Information | 18821

Syntax

```
show chassis scb-ha-port statistics
```

Description

On SRX5400, SRX5600, and SRX5800 devices, this command displays information about switch control board (SCB) HA control port 0 small form-factor pluggable (SFP) of the local node in chassis cluster. The supported SCB versions are SRX5K SCB3 and SCB4 and SFP types are SFPP-10GE-LR and SFPP-10GE-SR.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show chassis scb-ha-port statistics` command.

Table 338: SCB HA Port Statistics

Field Name	Field Description
`Control ports optics basic information`	• `Port`—Specifies control port number. • `Cable type`—Specifies the type of cable. • `Fiber type`—Specifies the fiber type. • `Xcvr vendor`—Specifies the vendor details. • `Xcvr vendor part number`—Specifies the part number. • `Wave-length`—Specifies the wave length. • `Xcvr Firmware`—Specifies the firmware version. • `JNPR Rev`—Specifies the revision number. • `MSA Version`—Specifies the MSA version.

Table 338: SCB HA Port Statistics *(Continued)*

Field Name	Field Description
`Control ports optics diagnostics information: Control port 0`	• `Laser bias current`—Specifies the power details. • `Laser output power`—Specifies the power details. • `Module temperature`—Specifies the power details. • `Module voltage`—Specifies the power details. • `Laser receiver power`—Specifies the power details. • `Alarm/warning flags implemented`—Displays alarm On or Off status. • `Laser bias current high alarm / low alarm / high warning / low warning`—Displays alarm On or Off status. • `Laser output power high alarm / low alarm / high warning / low warning`—Displays alarm On or Off status • `Module temperature high alarm / low alarm / high warning / low warning`—Displays alarm On or Off status. • `Module voltage high alarm / low alarm / high warning / low warning`—Displays alarm On or Off status. • `Laser rx power high alarm / low alarm / high warning / low warning`—Displays alarm On or Off status. • `Laser bias current high alarm threshold`—Specifies the power details. • `Laser bias current low alarm threshold`—Specifies the power details. • `Laser bias current high warning threshold`—Specifies the power details. • `Laser bias current low warning threshold`—Specifies the power details. • `Laser output power high alarm threshold`—Specifies the power details. • `Laser output power low alarm threshold`—Specifies the power details. • `Laser output power high warning threshold`—Specifies the power details. • `Laser output power low warning threshold`—Specifies the power details.

Table 338: SCB HA Port Statistics *(Continued)*

Field Name	Field Description
	• `Module temperature high alarm threshold`—Specifies the power details. • `Module temperature low alarm threshold`—Specifies the power details. • `Module temperature high warning threshold`—Specifies the power details. • `Module temperature low warning threshold`—Specifies the power details. • `Module voltage high alarm threshold`—Specifies the power details. • `Module voltage low alarm threshold`—Specifies the power details. • `Module voltage high warning threshold`—Specifies the power details. • `Module voltage low warning threshold`—Specifies the power details. • `Laser rx power high alarm threshold`—Specifies the power details. • `Laser rx power low alarm threshold`—Specifies the power details. • `Laser rx power high warning threshold`—Specifies the power details. • `Laser rx power low warning threshold`—Specifies the power details.

Sample Output

show chassis scb-ha-port statistics

```
user@host> show chassis scb-ha-port statistics
```

```
Control ports optics basic information:
                        Fiber                    Xcvr vendor       Wave-
Xcvr          JNPR     MSA
  Port Cable type        type  Xcvr vendor        part number       length
Firmware      Rev      Version
  0    10G Base-SR       MM    FINISAR CORP.      FTLX8571D3BCL-J1  850nm
0.0           REV 01   SFF-8472 ver 10.2
```

```
Control ports optics diagnostics information:
Control port 0:
Laser bias current                                                  :  7.778 mA
Laser output power                                                  :  0.5914 mW / -2.28 dBm
Module temperature                                                  :  32.57 degrees C /
90.63 degrees F
Module voltage                                                      :  3.2880 V
Laser receiver power                                                :  0.4990 mW / -3.02 dBm
Alarm/warning flags implemented                                     :  On
Laser bias current high alarm / low alarm / high warning / low warning  :  Off / Off / Off / Off
Laser output power high alarm / low alarm / high warning / low warning  :  Off / Off / Off / Off
Module temperature high alarm / low alarm / high warning / low warning  :  Off / Off / Off / Off
Module voltage high alarm / low alarm / high warning / low warning      :  Off / Off / Off / Off
Laser rx power high alarm / low alarm / high warning / low warning      :  Off / Off / Off / Off
Laser bias current high alarm threshold                             :  13.200 mA
Laser bias current low alarm threshold                              :  4.000 mA
Laser bias current high warning threshold                           :  12.600 mA
Laser bias current low warning threshold                            :  5.000 mA
Laser output power high alarm threshold                             :  1.0000 mW / 0.00 dBm
Laser output power low alarm threshold                              :  0.2512 mW / -6.00 dBm
Laser output power high warning threshold                           :  0.7943 mW / -1.00 dBm
Laser output power low warning threshold                            :  0.3162 mW / -5.00 dBm
Module temperature high alarm threshold                             :  78.00 degrees C /
172.40 degrees F
Module temperature low alarm threshold                              :  -13.00 degrees C /
8.60 degrees F
Module temperature high warning threshold                           :  73.00 degrees C /
163.40 degrees F
Module temperature low warning threshold                            :  -8.00 degrees C /
17.60 degrees F
Module voltage high alarm threshold                                 :  3.7000 V
Module voltage low alarm threshold                                  :  2.9000 V
Module voltage high warning threshold                               :  3.6000 V
Module voltage low warning threshold                                :  3.0000 V
Laser rx power high alarm threshold                                 :  1.0000 mW / 0.00 dBm
Laser rx power low alarm threshold                                  :  0.0100 mW / -20.00 dBm
Laser rx power high warning threshold                               :  0.7943 mW / -1.00 dBm
Laser rx power low warning threshold                                :  0.0158 mW / -18.01 dBm
```

Release Information

Command introduced in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

Chassis Cluster Dual Control Links

show chassis security uefi variables

IN THIS SECTION

- Syntax | 18821
- Description | 18821
- Required Privilege Level | 18821
- Sample Output | 18822
- Release Information | 18822

Syntax

```
show chassis security uefi variables
```

Description

Show information about platform Unified Extensible Firmware Interface (UEFI) variables.

Required Privilege Level

view

Sample Output

show chassis security uefi variables

```
user@host> show chassis security uefi variables
Secure Bios UEFI Variables:
---------------------------
Variable PK has no entries
Variable KEK has no entries
Variable db has no entries
Variable dbx, length 2929
dbx: List 0, type X509
    Signature 0, size 1078, owner 00000000-0000-0000-0000-000000000000
        Subject:
            CN=SecureBoot, DC=db1
        Issuer:
            CN=SecureBoot, DC=db1
dbx: List 1, type X509
    Signature 0, size 1078, owner 00000000-0000-0000-0000-000000000000
        Subject:
            CN=SecureBoot, DC=db2
        Issuer:
            CN=SecureBoot, DC=db2
dbx: List 2, type SHA256
    Signature 0, size 32, owner 605dab50-e046-4300-abb6-3dd810dd8b23
        Hash:0000000000000000000000000bd9f36d1d60e653bb469dd2d24f34406ad87f6eb
dbx: List 3, type X509
    Signature 0, size 489, owner 00000000-0000-0000-0000-000000000000
        Subject:
            CN=sfTest
        Issuer:
            CN=sfTest
dbx: List 4, type SHA256
    Signature 0, size 32, owner 605dab50-e046-4300-abb6-3dd810dd8b23
        Hash:0000000000000000000000000495db49f2553505cb3bc338dc8d816f730b2dc67
Variable MokList has no entries
```

Release Information

Statement introduced in Junos OS Evolved Release 18.4R1.

show chassis sfb

IN THIS SECTION

- Syntax | 18823
- Syntax (MX10004, MX10008 Universal Edge Routers) | 18823
- Description | 18823
- Options | 18824
- Required Privilege Level | 18824
- Output Fields | 18824
- Sample Output | 18825
- show chassis sfb (MX10004 and MX10008 Devices) | 18827
- Release Information | 18828

Syntax

```
show chassis sfb
<all-members>
<local>
<member member-id>
<slot sfb-slot-number>
```

Syntax (MX10004, MX10008 Universal Edge Routers)

```
show chassis sfb
<errors>
<slot sfb-slot-number>
```

Description

Display chassis information about the Switch Fabric Board (SFB, SFB2).

Options

none	Display chassis information about all Switch Fabric Boards.
all-members	(Optional) Display chassis information about the SFB in all members of the Virtual Chassis configuration.
local	(Optional) Display chassis information about the SFB in the local member of the Virtual Chassis.
member *member-id*	(Optional) Display chassis information about the SFB in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.
errors	Show ASIC errors on sfb.
sfb-slot-number	(Optional) Display chassis information about the specified Switch Fabric Board, and specified SFB2 for MX10004. For MX2020, MX2010, and MX2008 routers, replace *sfb-slot-number* with a value from 0 through 7. For MX10004, value from 0 through 5.

Required Privilege Level

view

Output Fields

Table 339 on page 18824 lists the output fields for the `show chassis sfb` command. Output fields are listed in the approximate order in which they appear.

Table 339: show chassis sfb Output Fields

Field Name	Field Description
`Slot`	Slot number.
`State`	Status of the SFB. • `Online`—The SFB is online and running. • `Offline`— SFB is powered down.

Table 339: show chassis sfb Output Fields *(Continued)*

Field Name	Field Description
`Uptime`	How long the Routing Engine has been connected to the SFB and, therefore, how long the SFB has been up and running.

Sample Output

show chassis sfb (MX2020 Router)

```
user@host> show chassis sfb
Slot  State              Uptime
 0    Online              6 hours, 11 minutes, 33 seconds
 1    Online              6 hours, 11 minutes, 27 seconds
 2    Online              6 hours, 11 minutes, 21 seconds
 3    Online              6 hours, 11 minutes, 15 seconds
 4    Online              6 hours, 11 minutes, 8 seconds
 5    Online              6 hours, 11 minutes, 2 seconds
 6    Online              6 hours, 10 minutes, 57 seconds
 7    Online              6 hours, 10 minutes, 51 seconds
```

show chassis sfb (MX2010 Router)

```
user@host> show chassis sfb
Slot  State              Uptime
 0    Online              6 hours, 48 minutes, 28 seconds
 1    Online              6 hours, 48 minutes, 23 seconds
 2    Online              6 hours, 48 minutes, 17 seconds
 3    Offline              --- Restarting unresponsive board ---
 4    Online              6 hours, 48 minutes, 12 seconds
 5    Online              6 hours, 48 minutes, 6 seconds
 6    Online              6 hours, 48 minutes
 7    Online              6 hours, 47 minutes, 55 seconds
```

show chassis sfb (MX2008 Router)

```
user@host> show chassis sfb
Slot  State              Uptime
 0    Online              10 days, 6 hours, 52 minutes, 38 seconds
 1    Online              10 days, 6 hours, 52 minutes, 32 seconds
 2    Online              10 days, 6 hours, 52 minutes, 26 seconds
 3    Online              10 days, 6 hours, 52 minutes, 20 seconds
 4    Online              10 days, 6 hours, 52 minutes, 15 seconds
 5    Online              10 days, 6 hours, 52 minutes, 9 seconds
 6    Online              10 days, 6 hours, 52 minutes, 3 seconds
 7    Online              10 days, 6 hours, 51 minutes, 58 seconds
```

show chassis sfb (MX10004 Router)

```
user@router> show chassis sfb
Slot  State              Uptime
 0    Online              3 days, 5 hours, 19 minutes, 1 second
 1    Online              3 days, 5 hours, 18 minutes, 31 seconds
 2    Online              3 days, 5 hours, 18 minutes, 2 seconds
 3    Online              3 days, 5 hours, 17 minutes, 31 seconds
 4    Online              3 days, 5 hours, 17 minutes, 1 second
 5    Online              3 days, 5 hours, 16 minutes, 31 seconds
```

show chassis sfb slot (MX10004 Router)

MX10004 router supports chassis slots from 0 to 5.

```
user@router> show chassis sfb slot 0
slot  State              Uptime
 0    Online              3 days, 5 hours, 27 minutes, 57 seconds
```

```
user@router> show chassis sfb slot 2
slot  State              Uptime
 2    Offline             --- Offlined due to unsupported fabric ---
```

show chassis sfb (MX10008 Router)

```
user@host> show chassis sfb
Slot  State              Uptime
 0    Online              21 hours, 25 minutes, 50 seconds
 1    Online              21 hours, 25 minutes, 22 seconds
 2    Online              21 hours, 24 minutes, 55 seconds
 3    Online              21 hours, 24 minutes, 27 seconds
 4    Online              21 hours, 23 minutes, 56 seconds
 5    Online              21 hours, 23 minutes, 23 seconds
```

show chassis sfb (MX10004 and MX10008 Devices)

On MX10004 and MX10008 devices, the JNP10K-LC4800 line cards are powered ON only when the following compatible components are present:

- Fan Trays : JNP10004-FAN2, JNP10008-FAN2
- PSM or PEM : JNP10K-PWR-AC2, JNP10K-PWR-AC3, JNP10K-PWR-DC2
- Enhanced Switch Fabric Board (SFB2)

The `show chassis sfb slot` command output displays the following error message when compatible SFB2 is not present:

```
user@host> show chassis sfb slot 0
Slot  State              Uptime
0    Offline              ---SFB incompatible with FAN-TRAYs ---
```

The `show chassis sfb slot` command output displays the following error message when compatible PSMs or PEMs are not present:

```
user@router> show chassis fpc 3
                    Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State            I  Total  Interrupt      DRAM (MB) Heap     Buffer
3  Offline         ---FPC incompatible with PEMs ---
```

Release Information

Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 Universal Routing Platforms.

`all-members`, `local`, and `member` *`member-id`* options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.

RELATED DOCUMENTATION

request chassis sfb | 14555

show chassis sfb errors

IN THIS SECTION

- Syntax | 18828
- Description | 18828
- Options | 18829
- Required Privilege Level | 18829
- Output Fields | 18829
- Sample Output | 18830
- Release Information | 18830

Syntax

```
show chassis sfb errors
< slot sfb-slot-number>
```

Description

Display information about errors pertaining to Switch Fabric Boards (SFBs).

Options

none Display information about the errors pertaining to all the SFBs in the router.

sfb-slot-number (Optional) Display information about error for the specified Switch Fabric Board.

- **Range:** 0–5

Required Privilege Level

view

Output Fields

Table 340 on page 18829 lists the output fields for the `show chassis sfb errors` command. Output fields are listed in the approximate order in which they appear.

Table 340: show chassis sfb errors Output Fields

Field Name	Field Description
`Slot`	Displays the SFB slot number.
`Name`	Displays the name of the SFB. In MX10008 routers, SFB is identified as SIB.
`DeviceID`	Displays device ID.
Threshold	Displays the error threshold.
Active	Displays the number of active error instances.
Occurred	Displays the number of error instances occurred.
Cleared	Displays the number of error instances cleared.
Description	Displays error description.

Sample Output

show chassis sfb errors slot 1(MX10008 and MX10004 Router)

```
user@host> show chassis sfb errors slot 1

Slot Name              DeviceId    Threshold  Active  Occured  Cleared  Description
1    SIB 1_LTC3880_0   09204F8B00  1          0       0        0        PF0 Core 0.9V, voltage
                       09204FFFFF  1          0       0        0        PF0 Core 0.9V, I2C Access
     SIB 1_LTC3880_1   09204D8B00  1          0       0        0        PF0 AVDD 1V, voltage
                       09204DFFFF  1          0       0        0        PF0 AVDD 1V, I2C Access
     SIB 1_LTC3880_2   09204B8B00  1          0       0        0        PF1 Core 0.9V, voltage
                       09204BFFFF  1          0       0        0        PF1 Core 0.9V, I2C Access
     SIB 1_LTC3880_3   0920498B00  1          0       0        0        PF1 AVDD 1V, voltage
                       092049FFFF  1          0       0        0        PF1 AVDD 1V, I2C Access
     SIB 1_VT7505_0    0910408800  1          0       0        0        HOTSWAP 12V, voltage
                       091040FFFF  1          0       0        0        HOTSWAP 12V, I2C Access
     SIB 1_TEMP_0      0901000000  1          0       0        0        Intake-A Temp Sensor
                       0901000100  1          0       0        0        Intake-B Temp Sensor
                       0901002300  1          0       0        0        PF0 Temp Sensor
     SIB 1_TEMP_1      0902000000  1          0       0        0        Exhaust-A Temp Sensor
                       0902000100  1          0       0        0        Exhaust-B Temp Sensor
                       0902002300  1          0       0        0        PF1 Temp Sensor
```

Release Information

Command introduced in Junos OS Release 18.2.

RELATED DOCUMENTATION

request chassis sfb | 14555

show chassis sfm

IN THIS SECTION

- Syntax | 18831
- Description | 18831
- Options | 18831
- Required Privilege Level | 18831
- Output Fields | 18832
- Sample Output | 18833
- Release Information | 18835

Syntax

```
show chassis sfm
<detail <sfm-slot>>
```

Description

(M40e and M160 routers only) Display Switching and Forwarding Module (SFM) status information.

Options

none Display standard status information about all SFMs.

detail (Optional) Display detailed SFM status information.

sfm-slot (Optional) Display status information about the SFM in the specified slot only. For the M40e router, replace ***sfm-slot*** with **0** or **1**. For the M160 router, replace ***sfm-slot*** with a value from **0** through **3**.

Required Privilege Level

view

Output Fields

Table 341 on page 18832 lists the output fields for the `show chassis sfm` command. Output fields are listed in the approximate order in which they appear.

Table 341: show chassis sfm Output Fields

Field Name	Field Description	Level of Output
Slot	Slot number.	All levels
State	Status of the SFM. State can be any of the following: • **Online**—SFM is online and running. • **Online-Standby** (M40e router only)—SFM is online, operating as Standby. • **Offline**—SFM is powered down. • **Empty**—No SFM is present.	All levels
Reason	If the status is **Offline**, reason for this state.	All levels
Temp	Temperature of air passing by the SFM, in degrees Celsius.	none specified
CPU Utilization (%)	Information about CPU usage.	none specified
Total	Total percentage of the CPU being used by the SFM's processor.	All levels
Interrupt	Of the total CPU being used by the SFM's processor, the percentage being used for interrupts.	All levels
Memory Utilization	Information about memory usage.	none specified
DRAM	Total DRAM available to the SFM's processor, in megabytes (MB).	All levels

Table 341: show chassis sfm Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Heap	Percentage of heap space (dynamic memory) being used by the SFM's processor. If this number exceeds 80 percent, it might indicate a software problem (memory leak).	All levels
Buffer	Percentage of buffer space being used by the SFM's processor for buffering internal messages.	All levels
SPP Temperature	Temperature of air passing by the Switch Plane Processor card, in degrees Celsius and Fahrenheit	**detail**
SPR Temperature	Temperature of air passing by the Switch Plane Router card, in degrees Celsius and Fahrenheit.	**detail**
Total CPU DRAM	Total amount of CPU DRAM being used by the SFM's processor.	**detail**
Total SSRAM	Total amount of SSRAM being used by the SFM's processor.	**detail**
Internet processor II	(M160 router only) Processor type.	**detail**
Start time	Time this SFM became active.	**detail**
Uptime	How long the SFM has been up and running.	**detail**
Packet scheduling mode	(M160 router only) Enabled or disabled.	**detail**

Sample Output

show chassis sfm (M160 Router)

```
user@host> show chassis sfm
```

```
SFM status:
               Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State      (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online      39      0          0        64         0          6
  1  Online      43      0          0        64         0          6
  2  Empty        0      0          0         0         0          0
  3  Empty        0      0          0         0         0          0
```

show chassis sfm detail (M40e Router)

```
user@host> show chassis sfm detail
Slot 0 information:
  State                                 Offline
  Reason:                               - power configured off
Slot 1 information:
  State                                 Present
  SPP temperature                    0 degrees C / 32 degrees F
  SPR temperature                    0 degrees C / 32 degrees F
  Total CPU DRAM                      0 MB
  Total SSRAM                         0 MB
```

show chassis sfm detail (M160 Router)

```
user@host> show chassis sfm detail
Slot 0 information:
  State                                 Online
  SPP temperature                   37 degrees C / 98 degrees F
  SPR temperature                   39 degrees C / 102 degrees F
  Total CPU DRAM                    64 MB
  Total SSRAM                        8 MB
  Internet Processor II                 Version 1, Foundry IBM, Part number 9
  Start time:                           2004-08-17 09:23:08 PDT
  Uptime:                              72 days, 1 hour, 15 minutes, 57 seconds
Slot 1 information:
  State                                 Online
  SPP temperature                   36 degrees C / 96 degrees F
  SPR temperature                   37 degrees C / 98 degrees F
  Total CPU DRAM                    64 MB
  Total SSRAM                        8 MB
  Internet Processor II                 Version 1, Foundry IBM, Part number 9
```

```
  Start time:                          2004-08-17 09:23:08 PDT
  Uptime:                             72 days, 1 hour, 15 minutes, 57 seconds
Slot 2 information:
....
Packet scheduling mode : Disabled
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- request chassis sfm | 14558
- request chassis sfm master switch | 14559
- *Switching the Global Primary and Backup Roles in a Virtual Chassis Configuration*

show chassis sibs

IN THIS SECTION

- Syntax | 18836
- Syntax (TX Matrix Router) | 18836
- Syntax (TX Matrix Plus Router) | 18836
- Syntax (PTX Series Packet Transport Routers) | 18836
- Description | 18836
- Options | 18836
- Required Privilege Level | 18837
- Output Fields | 18837
- Sample Output | 18843
- Release Information | 18852

Syntax

```
show chassis sibs
```

Syntax (TX Matrix Router)

```
show chassis sibs
<lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
show chassis sibs
<lcc number | sfc number>
```

Syntax (PTX Series Packet Transport Routers)

```
show chassis sibs
<detail>
<slot>
```

Description

(M320,T Series routers, TX Matrix routers, TX Matrix Plus routers, and PTX Series routers only) Display Switch Interface Boards (SIBs) status information.

Options

none (TX Matrix routers and TX Matrix Plus routers only) On a TX Matrix router, display the SIB status for the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display the SIB status for the TX Matrix Plus router and its attached routers.

detail (PTX Series) (Optional) Display detailed SIB status information.

lcc *number* (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display SIB status information for a specified T640 router (line-card chassis or LCC) that is

connected to the TX Matrix router. On a TX Matrix Plus router, display SIB status information for a specified T1600 or T4000 router (LCC) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

scc (TX Matrix routers only) (Optional) Display SIB status information for the TX Matrix router (switch-card chassis).

sfc ***number*** (TX Matrix Plus routers only) (Optional) Display SIB status information for the TX Matrix Plus router (switch-fabric chassis or SFC). Replace ***number*** with **0**.

slot (PTX Series) (Optional) Display status information about the SIB in the specified slot only. The range of values is 0 through 8.

Required Privilege Level

view

Output Fields

Table 342 on page 18837 lists the output fields for the `show chassis sibs` command. Output fields are listed in the approximate order in which they appear.

Table 342: show chassis sibs Output Fields

Field Name	Field Description
Slot	SIB slot number.

Table 342: show chassis sibs Output Fields *(Continued)*

Field Name	Field Description
Type	(TX Matrix Plus router only) SIB type.
Uptime	How long the SIB has been up and running.

Table 342: show chassis sibs Output Fields *(Continued)*

Field Name	Field Description
State	SIB status: • **Activating**—SIB is coming online; this is a transitional state. • **Deactivating**—SIB is going offline; this is a transitional state. • **Connected**—SIBs on a T1600 router are connected and trained but are either not online or are spare, because the plane on the TX Matrix Plus router (or switch-fabric chassis) is still offline. • **Disconnected**—SIBs on all T640 routers on the TX Matrix router (switch-card chassis) are in the **Disconnected** state, because a SIB on the SCC has gone offline. Likewise, SIBs on all T1600 routers on the TX Matrix Plus router (or switch-fabric chassis) are in the **Disconnected** state, because a SIB on the SFC has gone offline. On the TX Matrix Plus router with 3D SIBs, the LCC SIB is also disconnected if the F13 SIB is online, but none of the cables are connected or trained. • `Onlining`—SIB is in the process of going online. ASICs , retimers and the rest of the hardware is initializing. • **Online**—SIB is operational and running. • `Offlining`—SIB is in the process of going offline. ASICs , retimers and rest of the hardware is being shutdown down to take the offline gracefully. • **Offline**—SIB is powered down. **NOTE**: If a SIB transitions to the **Offline** state, the command displays an appropriate reason in the output. For instance, if the SIB is taken offline using the `request chassis sib` command, the `show chassis sibs` command displays **--- Offlined by cli command ---** in the output. • **Spare**—SIB is redundant and will move to active state if one of the working SIBs fails to pass traffic. **NOTE**: **Spare** does not apply to PTX Series Packet Transport Routers, as there are no spare SIBs. • **Empty**—No SIB is present.

Table 342: show chassis sibs Output Fields *(Continued)*

Field Name	Field Description

	• **Fault**—SIB is in an alarmed state in which none of the SIB's planes are operational for one of the following reasons:
	  • All onboard fabric ASICs are not operational.
	  • Fiber-optic connector faults.
	  • FPC connector faults.
	  • SIB midplane connector faults.
	• `Fault-off`—SIB is powered off due to a fault in SIB.
	• **Check**—SIB is in an alarmed state due to link errors or destination errors. A SIB can transition to the **Check** state from the online or spare state.
	  The **Check** state can be caused by the following reasons:
	  • Unsupported FPC installed on a router.
	  • SIB not inserted properly (such as bent pins).
	  • Destination errors are detected on the SIB. In this case, the Packet Forwarding Engine stops using the SIB to send traffic to the affected destination Packet Forwarding Engine. When a Packet Forwarding Engine cannot be reached on that plane or SIB, a destination error is reported against that SIB.
	  NOTE: For SIBs in the **Check** state, the output displays some additional information:
	  • In Junos OS Release 9.6 and later, the **Check** state message shows the number of Packet Forwarding Engines in the plane having destination errors. For example, **Check (10 destination errors)** indicates 10 Packet Fowarding Engines cannot be reached on that particular SIB. If there are no destination errors, and if the SIB transitions to the **Check** state because of link errors only, the **Check** state message shows **Check (0 destination errors)**.

Table 342: show chassis sibs Output Fields *(Continued)*

Field Name	Field Description
	• In Junos OS Release 9.5 and earlier, the **Check** state message shows **Check (destination errors)** if there are Packet Forwarding Engines with destination errors in this plane. However, it does not show the number of Packet Forwarding Engines having destination errors. If there are no destination errors and if the SIB transitions to the **Check** state because of link errors only, the **Check** state message shows **Check (no destination errors)**. If the SIB is in a **Check** state, because of destination errors, the CLI displays an additional line in the output, **use "show chassis fabric fpcs" and "show chassis fabric sibs" for more details**. • Link errors are detected on the channel between the SIB and a Packet Forwarding Engine. Link errors can be detected at initialization time or runtime: • Link errors caused by a link training failure at initialization time—The Packet Forwarding Engine does not use the SIB to send traffic. The `show chassis fabric fpcs` command shows **Plane disabled** as status for this link. • Link errors caused by CRC errors detected at runtime—The Packet Forwarding Engine continues to use the SIB to send traffic. The `show chassis fabric fpcs` command shows **Link error** as the status for this link. **NOTE**: The **Check** state does not apply to PTX Series Packet Transport Routers. • **SFC Error**—If an F13 SIB on the TX Matrix Plus router (SFC) transitions to the **Fault** state (for instance, because of link errors), and then if an LCC SIB (connected to the F13 SIB) comes online, the LCC SIB transitions to the **SFC Error** state. This state indicates that the F13 SIB to which the LCC SIB is connected has errors. **NOTE**: The **Connected**, **Disconnected**, and **SFC Error** states are only applicable to the SIBs on an LCC. • **Invalid**—The specific SIB slot is not valid for 4-LCC chassis configuration. See the *TX Matrix Plus Hardware Guide* for more information about the supported SIB slots.

Table 342: show chassis sibs Output Fields *(Continued)*

Field Name	Field Description
	NOTE: The **Invalid** state is applicable to TX Matrix Plus routers only. • `Unknown`—SIB is present but the state is unknown. • `Present`—SIB is plugged in. The SIB is not powered on or operational.
Fabric links	Indicates status of fabric links on the SIB. • **Active**—All fabric links on SIB are active. Errors detected on the SIB's fabric links, if any, are reported in the Errors column. • **Unused**—All fabric links on the SIB are not used for fabric traffic.
Errors	Indicates if there is any error on the SIB. • **None**—No errors • **Link Errors**—Fabric link errors were found on SIB RX link. • **Cell drops**—Fabric cell drops were found on the SIB ASIC. • **Link Errors, Cell drops**—Both link errors and cell drops were detected on at least one of the SIB's fabric links. • `Asic Errors`—A fault affecting one of the ASICs on the SIB is detected. It can be an IO error or an internal error signaled by the ASIC.
Link Errors	indicate the number of links which are marked faulty because the errors on them have crossed threshold.
Cable Errors	Indicate the number of mandatory cables that are not connected, or in up state for that plane
Destination Errors	Indicate the number of destinations that are not reachable on this plane.

Sample Output

show chassis sibs (T640 Router)

```
user@host> show chassis sibs

Slot  State                       Uptime
 0    Empty
 1    Offline                          --- Offlined by cli command ---
 2    Check (21 destination errors)    1 day, 1 hour, 32 minutes, 55 seconds
 3    Check (0 destination errors)     1 day, 1 hour, 32 minutes, 45 seconds
 4    Empty

use "show chassis fabric fpcs" and "show chassis fabric sibs" for more details
```

show chassis sibs (T4000 Router)

```
user@host> show chassis sibs
Slot  State                           Uptime
 0    Spare
 1    Online                           3 hours, 48 minutes, 38 seconds
 2    Online                           3 hours, 48 minutes, 22 seconds
 3    Online                           3 hours, 48 minutes, 5 seconds
 4    Online                           3 hours, 47 minutes, 49 seconds
```

show chassis sibs (TX Matrix Router)

```
user@host> show chassis sibs
scc-re0:
--------------------------------------------------------------------------
Slot  State               Uptime
 0    Empty
 1    Empty
 2    Offline             --- Offlined by cli command ---
 3    Offline
 4    Online              7 days, 21 hours, 50 minutes, 4 seconds
lcc0-re0:
--------------------------------------------------------------------------
```

```
Slot  State                           Uptime
 0    Offline                         --- Offlined by cli command ---
 1    Empty
 2    Check (21 destination errors)   1 day, 1 hour, 32 minutes, 55 seconds
 3    Check (0 destination errors)    1 day, 1 hour, 32 minutes, 45 seconds
 4    Empty

use "show chassis fabric fpcs" and "show chassis fabric sibs" for more details
```

show chassis sibs (T1600 Router)

```
user@host> show chassis sibs

Slot
 Slot  State                           Uptime
  0    Check (destination errors)      2 hours, 23 minutes, 2 seconds
  1    Offline                         --- Offlined by cli command ---
  2    Check (destination errors)      2 hours, 23 minutes, 3 seconds
  3    Check (destination errors)      2 hours, 23 minutes, 3 seconds
  4    Check (destination errors)      2 hours, 23 minutes, 3 seconds

 use "show chassis fabric fpcs" and "show chassis fabric sibs" for more details
```

show chassis sibs (TX Matrix Plus Router)

```
user@host> show chassis sibs
sfc0-re0:
--------------------------------------------------------------------------
Slot  State              Type          Link errors  Destination errors  Uptime
 0    Spare              SIB F13          NONE             NONE
 1    Empty                               NONE             NONE
 2    Invalid                             NONE             NONE
 3    Online             SIB F13          NONE             NONE         1 hour, 53 minutes, 19
seconds
 4    Empty                               NONE             NONE
 5    Invalid                             NONE             NONE
 6    Online             SIB F13          NONE             NONE         1 hour, 53 minutes, 8
seconds
 7    Empty                               NONE             NONE
```

```
 8    Online            SIB F13          NONE             NONE         1 hour, 52 minutes, 57
seconds
 9    Empty                              NONE             NONE
10    Invalid                            NONE             NONE
11    Online            SIB F13          NONE             NONE         1 hour, 52 minutes, 46
seconds
12    Empty                              NONE             NONE
13    Invalid                            NONE             NONE
14    Invalid                            NONE             NONE
15    Invalid                            NONE             NONE
0/0   Spare             SIB F2S          NONE             NONE
0/2   Spare             SIB F2S          NONE             NONE
0/4   Spare             SIB F2S          NONE             NONE
0/6   Spare             SIB F2S          NONE             NONE
1/0   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 29
seconds
1/2   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 28
seconds
1/4   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 27
seconds
1/6   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 26
seconds
2/0   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 18
seconds
2/2   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 17
seconds
2/4   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 16
seconds
2/6   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 14
seconds
3/0   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 7
seconds
3/2   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 5
seconds
3/4   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 4
seconds
3/6   Online            SIB F2S          NONE             NONE         1 hour, 53 minutes, 3
seconds
4/0   Online            SIB F2S          NONE             NONE         1 hour, 52 minutes, 56
seconds
4/2   Online            SIB F2S          NONE             NONE         1 hour, 52 minutes, 54
seconds
4/4   Online            SIB F2S          NONE             NONE         1 hour, 52 minutes, 53
```

```
seconds
4/6   Online             SIB F2S          NONE             NONE         1 hour, 52 minutes, 52
seconds

lcc0-re0:
--------------------------------------------------------------------------
Slot  State              Link errors  Destination errors  Uptime
 0    Spare                 NONE             NONE
 1    Online                NONE             NONE         1 hour, 53 minutes, 31 seconds
 2    Online                NONE             NONE         1 hour, 53 minutes, 27 seconds
 3    Online                NONE             NONE         1 hour, 53 minutes, 23 seconds
 4    Online                NONE             NONE         1 hour, 53 minutes, 19 seconds
```

show chassis sibs (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis sibs
sfc0-re0:
--------------------------------------------------------------------------
Slot  State              Type          Cable errors  Link errors  Destination errors  Uptime
 0    Online             SIB F13            6           NONE             NONE         21 hours,
54 minutes, 28 seconds
 1    Online             SIB F13            8           NONE             NONE         21 hours,
54 minutes, 12 seconds
 2    Invalid                              NONE         NONE             NONE
 3    Online             SIB F13            6           NONE             NONE         21 hours,
57 minutes, 6 seconds
 4    Online             SIB F13            8            1               NONE         21 hours,
56 minutes, 49 seconds
 5    Invalid                              NONE         NONE             NONE
 6    Online             SIB F13            6           NONE             NONE         21 hours,
56 minutes, 25 seconds
 7    Online             SIB F13            8           NONE             NONE         21 hours,
56 minutes, 8 seconds
 8    Online             SIB F13            6           NONE             NONE         21 hours,
55 minutes, 43 seconds
 9    Online             SIB F13            8           NONE             NONE         21 hours,
55 minutes, 26 seconds
10    Invalid                              NONE         NONE             NONE
11    Empty                                NONE         NONE             NONE
12    Empty                                NONE         NONE             NONE
13    Invalid                              NONE         NONE             NONE
```

```
14    Invalid                              NONE         NONE             NONE
15    Invalid                              NONE         NONE             NONE
0/0   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
55 minutes, 16 seconds
0/2   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
54 minutes, 49 seconds
0/4   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
54 minutes, 47 seconds
0/6   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
54 minutes, 45 seconds
1/0   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
57 minutes, 29 seconds
1/2   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
57 minutes, 27 seconds
1/4   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
57 minutes, 25 seconds
1/6   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
57 minutes, 23 seconds
2/0   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
56 minutes, 48 seconds
2/2   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
56 minutes, 46 seconds
2/4   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
56 minutes, 43 seconds
2/6   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
56 minutes, 41 seconds
3/0   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
56 minutes, 6 seconds
3/2   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
56 minutes, 4 seconds
3/4   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
56 minutes, 2 seconds
3/6   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
56 minutes
4/0   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
55 minutes, 24 seconds
4/2   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
55 minutes, 22 seconds
4/4   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
55 minutes, 20 seconds
4/6   Online            SIB F2S            -n/a-        NONE             NONE         21 hours,
55 minutes, 18 seconds
```

```
lcc0-re0:
--------------------------------------------------------------------------
Slot  State            Cable errors  Link errors  Destination errors  Uptime
 0    Online                 6           NONE             NONE         21 hours, 47 minutes, 29
seconds
 1    Online                 6           NONE             NONE         21 hours, 47 minutes, 50
seconds
 2    Online                 6           NONE             NONE         21 hours, 47 minutes, 43
seconds
 3    Online                 6           NONE             NONE         21 hours, 47 minutes, 36
seconds
 4    Empty                 NONE         NONE             NONE
use "show chassis fabric fpcs" and "show chassis fabric sibs" for more details

lcc4-re0:
--------------------------------------------------------------------------
Slot  State            Cable errors  Link errors  Destination errors  Uptime
 0    Online                 6           NONE             NONE         21 hours, 57 minutes, 1
second
 1    Online                 6           NONE             NONE         21 hours, 57 minutes, 21
seconds
 2    Online                 6           NONE             NONE         21 hours, 57 minutes, 14
seconds
 3    Online                 6           NONE             NONE         21 hours, 57 minutes, 7
seconds
 4    Empty                 NONE         NONE             NONE
use "show chassis fabric fpcs" and "show chassis fabric sibs" for more details

lcc7-re0:
--------------------------------------------------------------------------
Slot  State            Cable errors  Link errors  Destination errors  Uptime
 0    Online                 2           NONE             NONE         21 hours, 56 minutes, 54
seconds
 1    Online                 2           NONE             NONE         21 hours, 57 minutes, 21
seconds
 2    Online                 2           NONE             NONE         21 hours, 57 minutes, 12
seconds
 3    Online                 2           NONE             NONE         21 hours, 57 minutes, 3
seconds
 4    Empty                 NONE         NONE             NONE
use "show chassis fabric fpcs" and "show chassis fabric sibs" for more details
```

show chassis sibs sfc (TX Matrix Plus Router))

```
user@host> show chassis sibs sfc 0
sfc0-re0:
--------------------------------------------------------------------------
Slot  State              Type          Link errors  Destination errors  Uptime
 0    Spare              SIB F13          NONE             NONE
 1    Empty                               NONE             NONE
 2    Invalid                             NONE             NONE
 3    Online             SIB F13          NONE             NONE         12 hours, 6 minutes, 22
seconds
 4    Empty                               NONE             NONE
 5    Invalid                             NONE             NONE
 6    Online             SIB F13          NONE             NONE         12 hours, 6 minutes, 11
seconds
 7    Empty                               NONE             NONE
 8    Online             SIB F13          NONE             NONE         12 hours, 6 minutes
 9    Empty                               NONE             NONE
10    Invalid                             NONE             NONE
11    Online             SIB F13          NONE             NONE         12 hours, 5 minutes, 49
seconds
12    Empty                               NONE             NONE
13    Invalid                             NONE             NONE
14    Invalid                             NONE             NONE
15    Invalid                             NONE             NONE
0/0   Spare              SIB F2S          NONE             NONE
0/2   Spare              SIB F2S          NONE             NONE
0/4   Spare              SIB F2S          NONE             NONE
0/6   Spare              SIB F2S          NONE             NONE
1/0   Online             SIB F2S          NONE             NONE         12 hours, 6 minutes, 32
seconds
1/2   Online             SIB F2S          NONE             NONE         12 hours, 6 minutes, 31
seconds
1/4   Online             SIB F2S          NONE             NONE         12 hours, 6 minutes, 30
seconds
1/6   Online             SIB F2S          NONE             NONE         12 hours, 6 minutes, 29
seconds
2/0   Online             SIB F2S          NONE             NONE         12 hours, 6 minutes, 21
seconds
2/2   Online             SIB F2S          NONE             NONE         12 hours, 6 minutes, 20
seconds
2/4   Online             SIB F2S          NONE             NONE         12 hours, 6 minutes, 19
```

```
seconds
2/6   Online            SIB F2S          NONE             NONE         12 hours, 6 minutes, 17
seconds
3/0   Online            SIB F2S          NONE             NONE         12 hours, 6 minutes, 10
seconds
3/2   Online            SIB F2S          NONE             NONE         12 hours, 6 minutes, 9
seconds
3/4   Online            SIB F2S          NONE             NONE         12 hours, 6 minutes, 7
seconds
3/6   Online            SIB F2S          NONE             NONE         12 hours, 6 minutes, 6
seconds
4/0   Online            SIB F2S          NONE             NONE         12 hours, 5 minutes, 59
seconds
4/2   Online            SIB F2S          NONE             NONE         12 hours, 5 minutes, 57
seconds
4/4   Online            SIB F2S          NONE             NONE         12 hours, 5 minutes, 56
seconds
4/6   Online            SIB F2S          NONE             NONE         12 hours, 5 minutes, 55
seconds
```

show chassis sibs lcc (TX Matrix Plus Router)

```
user@host> show chassis sibs lcc 0
lcc0-re0:
--------------------------------------------------------------------------
Slot  State              Link errors  Destination errors  Uptime
 0    Online                NONE             NONE         20 hours, 14 minutes, 50 seconds
 1    Fault                 NONE             NONE
 2    Online                NONE             NONE         20 hours, 15 minutes, 2 seconds
 3    Online                NONE             NONE         20 hours, 14 minutes, 58 seconds
 4    Online                NONE             NONE         20 hours, 14 minutes, 54 seconds
```

show chassis sibs lcc (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis sibs lcc 0
lcc0-re0:
--------------------------------------------------------------------------
Slot  State            Cable errors  Link errors  Destination errors  Uptime
 0    Disconnected         NONE         NONE             NONE         17 hours, 2 minutes, 37
seconds
```

```
 1    Online               NONE         NONE             NONE         17 hours, 3 minutes, 6
seconds
 2    Online               NONE         NONE             NONE         17 hours, 2 minutes, 59
seconds
 3    Online               NONE         NONE             NONE         17 hours, 2 minutes, 52
seconds
 4    Online               NONE         NONE             NONE         17 hours, 2 minutes, 44
seconds
```

show chassis sibs (M320 Router)

```
user@host> show chassis sibs

 0    Online             1 hour, 18 minutes, 3 seconds
 1    Offline            --- Offlined by cli command ---
 2    Online             1 hour, 18 minutes, 18 seconds
 3    Online             1 hour, 18 minutes, 3 seconds
```

show chassis sibs (PTX Series)

```
user@host> show chassis sibs
Slot  State         Fabric links        Errors
0    Online             Active         Asic Errors
1    Online             Active         Link Errors
2    Online             Active         None
3    Online             Active         Cell drops
4    Offline            Unused         None
5    Online             Active         None
6    Online             Active         None
7    Online             Active         None
8    Online             Active         None
```

show chassis sibs (PTX Series)

```
user@host> show chassis sibs detail
Slot 4 information
  State                                 Offline
  Reason                                Offlined by cli command
```

```
  Fabric links                        Unused
  Errors                              None
```

show chassis sibs (PTX5000)

The `show chassis sibs` command displays only the non-empty slots.

```
user@host> show chassis sibs
Slot  State         Fabric links    Errors
 1    Online             Unused     None
 2    Online             Unused     None
 3    Online             Unused     None
 4    Online             Unused     None
 6    Fault              Unused     None
 7    Online             Unused     None
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

detail and ***sib-slot*** options introduced for the PTX Packet Transport Router in Junos OS Release 12.1

Command introduced in Junos OS Release 17.2 for PTX10008 Routers.

RELATED DOCUMENTATION

- Monitoring the SIBs
- M320 SIB Description
- Routing Matrix with a TX Matrix Plus Router Solutions Page

show chassis spmb

IN THIS SECTION

- Syntax | 18853
- Syntax (MX Series Routers) | 18853
- Syntax (T4000 Routers) | 18853
- Syntax (TX Matrix Routers) | 18854
- Syntax (TX Matrix Plus Routers) | 18854
- Description | 18854
- Options | 18854
- Required Privilege Level | 18856
- Output Fields | 18856
- Sample Output | 18857
- Release Information | 18867

Syntax

```
show chassis spmb
```

Syntax (MX Series Routers)

```
show chassis spmb
<all-members>
<local>
<member member-id>
```

Syntax (T4000 Routers)

```
show chassis spmb
<sibs>
```

Syntax (TX Matrix Routers)

```
show chassis spmb
<sibs>
<lcc number | scc>
```

Syntax (TX Matrix Plus Routers)

```
show chassis spmb
<sibs>
<lcc number | sfc number>
```

Description

(T Series routers, MX2010, MX2020, and MX2008 routers only) Display Switch Processor Mezzanine Board (SPMB) status information.

Options

none (TX Matrix, TX Matrix Plus router, MX2010, MX2020, and MX2008 routers only) On a TX Matrix router, display SPMB status for the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display SPMB status for the TX Matrix Plus router and its attached routers. On MX2010, MX2020, and MX2008 routers, display the SPMB status for the routers.

all-members (MX2010, MX2020, and MX2008 routers only) (Optional) Display status information for the SPMB in all members of the Virtual Chassis configuration.

lcc *number* (TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display information about the SPMB on a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display information about the SPMB on a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.

- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX2010, MX2020, and MX2008 routers only) (Optional) Display status information for the SPMB in the local member of the Virtual Chassis.

member *member-id* (MX2010, MX2020 , and MX2008 routers only) (Optional) Display status information for the SPMB in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.

scc (TX Matrix routers only) (Optional) Display information about the SPMB on the TX Matrix router (switch-card chassis).

sfc *number* (TX Matrix Plus router only) (Optional) Display information about the SPMB on the TX Matrix Plus router (switch-fabric chassis). Replace *number* with `0`.

sibs (TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix Plus router, display information about the SIBs on the TX Matrix router (switch-card chassis). On a TX Matrix Plus router, display information about the SIBs on The TX Matrix Plus router (switch-fabric chassis). The `sibs` option has the following sub-options:

`lcc` *number* (TX Matrix, TX Matrix Plus router only) (Optional) On a TX Matrix router, display information about the SIBs on a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display information about the SIBs on a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

`scc` *number*—(TX Matrix routers only) (Optional) Display information about the SIBs on the TX Matrix router (switch-card chassis). Replace *number* variable with `0`.

`sfc` *number*—(TX Matrix Plus router only) (Optional) Display information about the SIBs on the TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with `0`.

Required Privilege Level

view

Output Fields

Table 343 on page 18856 lists the output fields for the `show chassis spmb` command. Output fields are listed in the approximate order in which they appear.

Table 343: show chassis spmb Output Fields

Field Name	Field Description
`Slot`	SPMB slot number: `0` or `1`.
`State`	SPMB status: • `Online`—SPMB is operational and running. • `Offline`—SPMB is powered down.
`Total CPU Utilization (%)`	Total percentage of CPU being used by the SPMB processor.
`Interrupt CPU Utilization (%)`	Of the total CPU being used by the SPMB processor, the percentage being used for interrupts.
`Memory Heap Utilization (%)`	Percentage of heap space (dynamic memory) being used by the FPC processor. If this number exceeds 80 percent, there may be a software problem (memory leak).
`Buffer Utilization (%)`	Percentage of buffer space being used by the SPMB processor for buffering internal messages.

Table 343: show chassis spmb Output Fields *(Continued)*

Field Name	Field Description
`Start time`	Time at which the SPMB last came online.
`Uptime`	How long the SPMB has been up and running.

Sample Output

show chassis spmb

```
user@host> show chassis spmb
Slot 0 information:
  State                          Online
  Total CPU Utilization             1%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               40%
  Start time:                    2001-08-27 14:05:04 PDT
  Uptime:                         46 minutes, 36 seconds
```

show chassis spmb (MX2010 Router)

```
user@host> show chassis spmb
Slot 0 information:
  State                          Online
  Total CPU Utilization            12%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           1%
  Buffer Utilization               22%
  Start time:                    2012-10-04 15:34:29 PDT
  Uptime:                        7 hours, 10 minutes, 15 seconds
Slot 1 information:
  State                          Online - Standby
  Total CPU Utilization             1%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
```

```
  Buffer Utilization               22%
  Start time:                    2012-10-02 14:34:54 PDT
  Uptime:                        2 days, 8 hours, 9 minutes, 50 seconds
```

show chassis spmb (MX2020 Router)

```
user@host> show chassis spmb
Slot 0 information:
  State                          Online
  Total CPU Utilization           100%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           3%
  Buffer Utilization               22%
  Start time:                    2012-10-03 14:58:26 PDT
  Uptime:                        1 day, 12 hours, 16 minutes, 14 seconds
Slot 1 information:
  State                          Online - Standby
  Total CPU Utilization             0%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               22%
  Start time:                    2012-10-03 14:58:27 PDT
  Uptime:                        1 day, 12 hours, 16 minutes, 13 seconds
```

show chassis spmb (MX2008 Router)

```
user@host> show chassis spmb
Slot 0 information:
  State                          Online
  Start time:                    2017-05-04 02:53:36 PDT
  Uptime:                        10 days, 7 hours, 1 minute, 14 seconds
Slot 1 information:
  State                          Online - Standby
  Start time:                    2017-05-04 02:53:36 PDT
  Uptime:                        10 days, 7 hours, 1 minute, 14 seconds
```

show chassis spmb (T4000 Router)

```
user@host> show chassis spmb

Slot 0 information:
  State                          Online
  Total CPU Utilization            18%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               22%
  Start time:                    2012-02-09 22:51:09 PST
  Uptime:                        2 hours, 25 minutes, 45 seconds
Slot 1 information:
  State                          Online - Standby
  Total CPU Utilization             0%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               22%
  Start time:                    2012-02-09 22:51:10 PST
  Uptime:                        2 hours, 25 minutes, 44 seconds
```

show chassis spmb lcc (TX Matrix Router)

```
user@host> show chassis spmb lcc 0
lcc0-re0:
--------------------------------------------------------------------------
Slot 0 information:
  State                          Online
  Total CPU Utilization             0%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               42%
  Start time:                    2004-08-05 18:43:38 PDT
  Uptime:                        8 days, 55 minutes, 52 seconds
```

show chassis spmb scc (TX Matrix Router)

```
user@host> show chassis spmb scc
scc-re0:
```

```
---------------------------------------------------------------------------
Slot 0 information:
  State                           Online
  Total CPU Utilization              1%
  Interrupt CPU Utilization          0%
  Memory Heap Utilization            0%
  Buffer Utilization                42%
  Start time:                     2004-08-05 18:43:37 PDT
  Uptime:                         8 days, 1 hour, 6 minutes, 51 seconds
```

show chassis spmb (T1600 Router)

```
user@host> show chassis spmb
Slot 0 information:
  State                           Online
  Total CPU Utilization              2%
  Interrupt CPU Utilization          0%
  Memory Heap Utilization            0%
  Buffer Utilization                24%
  Start time:                     2009-05-07 22:34:03 PDT
  Uptime:                         3 days, 4 hours, 14 minutes, 33 seconds
Slot 1 information:
  State                           Online - Standby
  Total CPU Utilization              0%
  Interrupt CPU Utilization          0%
  Memory Heap Utilization            0%
  Buffer Utilization                24%
  Start time:                     2009-05-07 22:34:02 PDT
  Uptime:                         3 days, 4 hours, 14 minutes, 34 seconds
```

show chassis spmb sibs (T1600 Router)

```
user@host> show chassis spmb sibs
Slot  State               Uptime
 0    Check              3 days, 4 hours, 11 minutes, 59 seconds
 1    Disconnected       3 days, 4 hours, 12 minutes, 36 seconds
 2    Disconnected       3 days, 4 hours, 12 minutes, 26 seconds
 3    Disconnected       3 days, 4 hours, 12 minutes, 17 seconds
 4    Disconnected       3 days, 4 hours, 12 minutes, 8 seconds
```

show chassis spmb (TX Matrix Plus Router)

```
user@host> show chassis spmb
sfc0-re0:
--------------------------------------------------------------------------
Slot 0 information:
  State                          Online
  Total CPU Utilization            84%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:20 PDT
  Uptime:                        46 minutes, 6 seconds
Slot 1 information:
  State                          Online - Standby
  Total CPU Utilization             0%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:20 PDT
  Uptime:                        46 minutes, 6 seconds

lcc0-re1:
--------------------------------------------------------------------------
Slot 0 information:
  State                          Online - Standby
  Total CPU Utilization             0%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:09 PDT
  Uptime:                        46 minutes, 24 seconds
Slot 1 information:
  State                          Online
  Total CPU Utilization             5%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:08 PDT
  Uptime:                        46 minutes, 25 seconds

lcc1-re1:
```

```
-------------------------------------------------------------------------
Slot 0 information:
  State                          Online - Standby
  Total CPU Utilization             1%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:09 PDT
  Uptime:                        46 minutes, 24 seconds
Slot 1 information:
  State                          Online
  Total CPU Utilization             5%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:10 PDT
  Uptime:                        46 minutes, 23 seconds

lcc2-re1:
-------------------------------------------------------------------------
Slot 0 information:
  State                          Online - Standby
  Total CPU Utilization             0%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:08 PDT
  Uptime:                        46 minutes, 25 seconds
Slot 1 information:
  State                          Online
  Total CPU Utilization             5%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:10 PDT
  Uptime:                        46 minutes, 23 seconds

lcc3-re1:
-------------------------------------------------------------------------
Slot 0 information:
  State                          Online - Standby
  Total CPU Utilization             1%
  Interrupt CPU Utilization         0%
```

```
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:10 PDT
  Uptime:                        46 minutes, 23 seconds
Slot 1 information:
  State                          Online
  Total CPU Utilization             5%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:09 PDT
  Uptime:                        46 minutes, 24 seconds
```

show chassis spmb lcc (TX Matrix Plus Router)

```
user@host> show chassis spmb lcc 2
lcc2-re1:
--------------------------------------------------------------------------
Slot 0 information:
  State                          Online - Standby
  Total CPU Utilization             0%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:08 PDT
  Uptime:                        45 minutes, 18 seconds
Slot 1 information:
  State                          Online
  Total CPU Utilization             6%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:10 PDT
  Uptime:                        45 minutes, 16 seconds
```

show chassis spmb scc (TX Matrix Plus Router)

```
user@host> show chassis spmb sfc 0
sfc0-re0:
--------------------------------------------------------------------------
```

```
Slot 0 information:
  State                          Online
  Total CPU Utilization            87%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:20 PDT
  Uptime:                        43 minutes, 32 seconds
Slot 1 information:
  State                          Online - Standby
  Total CPU Utilization             0%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-05-11 01:25:20 PDT
  Uptime:                        43 minutes, 32 seconds
```

show chassis spmb sibs (TX Matrix Plus Router)

```
user@host> show chassis spmb sibs
sfc0-re0:
--------------------------------------------------------------------------
Slot  State              Type          Uptime
 0    Online             SIB F13        1 hour, 18 minutes, 54 seconds
 1    Online             SIB F13        1 hour, 18 minutes, 45 seconds
 2    Invalid
 3    Online             SIB F13        1 hour, 20 minutes, 21 seconds
 4    Online             SIB F13        1 hour, 20 minutes, 18 seconds
 5    Invalid
 6    Online             SIB F13        1 hour, 19 minutes, 51 seconds
 7    Fault              SIB F13
 8    Online             SIB F13        1 hour, 19 minutes, 17 seconds
 9    Online             SIB F13        1 hour, 19 minutes, 13 seconds
10    Invalid
11    Online             SIB F13        1 hour, 17 minutes, 54 seconds
12    Online             SIB F13        1 hour, 17 minutes, 51 seconds
13    Invalid
14    Invalid
15    Invalid
0/0   Online             SIB F2S        1 hour, 18 minutes, 52 seconds
0/2   Online             SIB F2S        1 hour, 18 minutes, 51 seconds
```

```
0/4   Online             SIB F2S        1 hour, 18 minutes, 49 seconds
0/6   Online             SIB F2S        1 hour, 18 minutes, 48 seconds
1/0   Online             SIB F2S        1 hour, 20 minutes, 16 seconds
1/2   Online             SIB F2S        1 hour, 20 minutes, 15 seconds
1/4   Online             SIB F2S        1 hour, 20 minutes, 14 seconds
1/6   Online             SIB F2S        1 hour, 20 minutes, 13 seconds
2/0   Online             SIB F2S        1 hour, 19 minutes, 48 seconds
2/2   Online             SIB F2S        1 hour, 19 minutes, 47 seconds
2/4   Online             SIB F2S        1 hour, 19 minutes, 46 seconds
2/6   Online             SIB F2S        1 hour, 19 minutes, 44 seconds
3/0   Online             SIB F2S        1 hour, 19 minutes, 24 seconds
3/2   Online             SIB F2S        1 hour, 19 minutes, 22 seconds
3/4   Online             SIB F2S        1 hour, 19 minutes, 21 seconds
3/6   Online             SIB F2S        1 hour, 19 minutes, 20 seconds
4/0   Online             SIB F2S        1 hour, 18 minutes, 2 seconds
4/2   Online             SIB F2S        1 hour, 18 minutes
4/4   Online             SIB F2S        1 hour, 17 minutes, 58 seconds
4/6   Online             SIB F2S        1 hour, 17 minutes, 58 seconds

lcc0-re1:
--------------------------------------------------------------------------
Slot  State               Uptime
 0    Online             1 hour, 18 minutes, 58 seconds
 1    Online             1 hour, 20 minutes, 25 seconds
 2    Fault
 3    Online             1 hour, 18 minutes, 30 seconds
 4    Online             1 hour, 18 minutes, 28 seconds

lcc1-re1:
--------------------------------------------------------------------------
Slot  State               Uptime
 0    Online             1 hour, 18 minutes, 58 seconds
 1    Online             1 hour, 20 minutes, 26 seconds
 2    Fault
 3    Online             1 hour, 18 minutes, 22 seconds
 4    Online             1 hour, 18 minutes, 20 seconds

lcc2-re1:
--------------------------------------------------------------------------
Slot  State               Uptime
 0    Online             1 hour, 18 minutes, 19 seconds
 1    Online             1 hour, 20 minutes, 25 seconds
 2    Fault
```

```
 3    Online             1 hour, 18 minutes, 17 seconds
 4    Online             1 hour, 18 minutes, 15 seconds

lcc3-re1:
--------------------------------------------------------------------------
Slot  State               Uptime
 0    Online             1 hour, 18 minutes, 27 seconds
 1    Online             1 hour, 20 minutes, 24 seconds
 2    Fault
 3    Online             1 hour, 18 minutes, 25 seconds
 4    Online             1 hour, 18 minutes, 23 seconds
```

show chassis spmb lcc (TX Matrix Plus router with 3D SIBs)

```
user@host > show chassis spmb lcc 0
lcc0-re1:
--------------------------------------------------------------------------
Slot 0 information:
  State                          Online - Standby
  Total CPU Utilization             0%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization                0%
  Start time:                    2013-02-08 00:57:20 PST
  Uptime:                        19 minutes, 43 seconds
Slot 1 information:
  State                          Online
  Total CPU Utilization             0%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               22%
  Start time:                    2013-02-08 00:56:59 PST
  Uptime:                        20 minutes, 4 seconds
```

show chassis spmb sfc (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis spmb sfc o
sfc0-re0:
--------------------------------------------------------------------------
Slot 0 information:
```

```
  State                           Online
  Total CPU Utilization              0%
  Interrupt CPU Utilization          0%
  Memory Heap Utilization            0%
  Buffer Utilization                 0%
  Start time:                     2013-02-06 19:16:55 PST
  Uptime:                         1 day, 6 hours, 2 minutes, 59 seconds
Slot 1 information:
  State                           Online - Standby
  Total CPU Utilization              0%
  Interrupt CPU Utilization          0%
  Memory Heap Utilization            0%
  Buffer Utilization                 0%
  Start time:                     2013-02-06 19:16:53 PST
  Uptime:                         1 day, 6 hours, 3 minutes, 1 second
```

Release Information

Command introduced before Junos OS Release 7.4.

`sibs` option introduced for the T1600 and TX Matrix Plus routers in Junos OS Release 9.6.

`all-members`, `local`, and `member` *`member-id`* options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.

RELATED DOCUMENTATION

- request chassis sib | 14562
- request chassis spmb restart | 14565
- show chassis spmb sibs | 18867

show chassis spmb sibs

IN THIS SECTION

- Syntax | 18868

- Syntax (TX Matrix Router) | 18868
- Syntax (TX Matrix Plus Router) | 18868
- Description | 18868
- Options | 18869
- Additional Information | 18869
- Required Privilege Level | 18869
- Output Fields | 18870
- Sample Output | 18872
- Release Information | 18875

Syntax

```
show chassis spmb sibs
```

Syntax (TX Matrix Router)

```
show chassis spmb sibs
<lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
show chassis spmb sibs
<lcc number | sfc number>
```

Description

(T Series routers only) Display Switch Processor Mezzanine Board (SPMB) Switch Interface Board (SIB) status information.

Options

none (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, display the SIB status for the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display the SIB status for the TX Matrix Plus router and its attached routers.

lcc ***number*** (TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display SIB status information for a specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display SIB status information for a specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

scc (TX Matrix router only) (Optional) Display SIB status information for the TX Matrix router (switch-card chassis).

sfc (TX Matrix Plus router only) (Optional) Display SIB status information for the TX Matrix Plus router (or switch-fabric chassis).

Additional Information

On a T Series router, you can use either this command or the `show chassis sibs` command to produce the same output. The `show chassis sibs` command is supported on the M320 router and on the T Series routers.

Required Privilege Level

view

Output Fields

Table 344 on page 18870 lists the output fields for the `show chassis spmb sibs` command. Output fields are listed in the approximate order in which they appear.

Table 344: show chassis spmb sibs Output Fields

Field Name	Field Description

Slot	SIB slot number:
	• T640 router, T1600 router or TX Matrix router—**0** through **4**
	• TX Matrix Plus router:
	  • TXP-F13 SIB Slots—**0** through **16**
	  • TXP-F2S SIB Slots —**0** – **4**/[**0**|**2**|**4**|**6**]
	• T320 router—**0** through **2**

Table 344: show chassis spmb sibs Output Fields *(Continued)*

Field Name	Field Description

State	SIB status:
	• **Disconnected**—SIBs on all T640 routers on the TX Matrix router (switch-card chassis) are in the **Disconnected** state, because a SIB on the SCC has gone offline. Likewise, SIBs on all T1600 or T4000 routers on the TX Matrix Plus router (or switch-fabric chassis) are in the **Disconnected** state, because a SIB on the SFC has gone offline.
	• **Online**—SPMB is operational and running.
	• **Offline**—SPMB is powered down.
	• **Spare**—SIB is redundant and will move to active state if one of the working SIBs fail to pass traffic.
	• **Empty**—No SPMB is present.
	• **Fault**— SIB is in alarmed state where the SIB's plane is not operational for the following reasons:
	  • On-board F-chip is not operational.
	  • Fiber optic connector faults.
	  • FPC connector faults.
	  • SIB midplane connector faults.
	• **Check**—SIB is in alarmed state where the SIB's plane is partially operational for the following reasons:
	  • SIB is not inserted properly.
	  • Two or more links between the SIB and PFE fails.

Uptime	How long the SIB has been up and running.

Sample Output

show chassis spmb sibs (T320 Router)

```
user@host> show chassis spmb sibs
Slot  State
 0    Spare
 1    Online
 2    Online
```

show chassis-spmb-sibs (T1600 Router)

```
user@host> show chassis spmb sibs
Slot  State
 0    Spare
 1    Online
 2    Empty
 3    Online
 4    Offline
```

show chassis spmb sibs (T4000 Router)

```
user@host> show chassis spmb sibs

Slot  State                           Uptime
 0    Spare
 1    Online                          2 hours, 28 minutes, 13 seconds
 2    Online                          2 hours, 27 minutes, 57 seconds
 3    Online                          2 hours, 27 minutes, 40 seconds
 4    Online                          2 hours, 27 minutes, 24 seconds
```

show chassis spmb sibs (TX Matrix Router)

```
user@host> show chassis spmb sibs
Slot  State
 0    Online
 1    Online
```

```
 2    Empty
 3    Online
 4    Offline
```

show chassis spmb sibs lcc (TX Matrix Router)

```
user@host> show chassis spmb sibs lcc 0
lcc0-re0:
--------------------------------------------------------------------------
Slot  State               Uptime
 0    Empty
 1    Empty
 2    Empty
 3    Disconnected           8 days, 48 minutes, 58 seconds
 4    Online                 8 days, 48 minutes, 57 seconds
```

show chassis spmb sibs scc (TX Matrix Router)

```
user@host> show chassis spmb sibs scc
scc-re0:
--------------------------------------------------------------------------
Slot  State                         Uptime
 0    Empty
 1    Empty
 2    Empty
 3    Offline
 4    Online             8 days, 54 minutes, 1 second
```

show chassis spmb sibs (TX Matrix Plus Router)

```
user@host> show chassis spmb sibs
sfc0-re0:
--------------------------------------------------------------------------
Slot  State              Type          Uptime
 0    Online             SIB F13        1 hour, 52 minutes, 55 seconds
 1    Empty
 2    Invalid
 3    Online             SIB F13        1 hour, 53 minutes, 3 seconds
```

```
 4    Empty
 5    Invalid
 6    Empty
 7    Empty
 8    Empty
 9    Empty
10    Invalid
11    Empty
12    Empty
13    Invalid
14    Invalid
15    Invalid
0/0   Online             SIB F2S        1 hour, 53 minutes, 2 seconds
0/2   Online             SIB F2S        1 hour, 53 minutes, 1 second
0/4   Online             SIB F2S        1 hour, 52 minutes, 59 seconds
0/6   Online             SIB F2S        1 hour, 52 minutes, 58 seconds
1/0   Online             SIB F2S        1 hour, 53 minutes, 10 seconds
1/2   Online             SIB F2S        1 hour, 53 minutes, 8 seconds
1/4   Online             SIB F2S        1 hour, 53 minutes, 7 seconds
1/6   Online             SIB F2S        1 hour, 53 minutes, 6 seconds
2/0   Empty
2/2   Empty
2/4   Empty
2/6   Empty
3/0   Empty
3/2   Empty
3/4   Empty
3/6   Empty
4/0   Empty
4/2   Empty
4/4   Empty
4/6   Empty

lcc0-re0:
--------------------------------------------------------------------------
Slot  State               Uptime
 0    Online             1 hour, 53 minutes, 1 second
 1    Online             1 hour, 53 minutes, 3 seconds
 2    Empty
 3    Empty
 4    Empty

lcc1-re1:
```

```
---------------------------------------------------------------------------
Slot  State                Uptime
 0    Online              1 hour, 47 minutes, 13 seconds
 1    Online              1 hour, 47 minutes, 15 seconds
 2    Empty
 3    Empty
 4    Empty
```

show chassis spmb sibs sfc (TX Matrix Plus Router)

```
user@host> show chassis spmb sibs sfc 0
sfc0-re0:
---------------------------------------------------------------------------
Slot 0 information:
  State                          Online
  Total CPU Utilization            16%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-06-17 20:59:47 PDT
  Uptime:                        1 hour, 56 minutes, 30 seconds
Slot 1 information:
  State                          Online - Standby
  Total CPU Utilization             0%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               24%
  Start time:                    2009-06-17 20:59:48 PDT
  Uptime:                        1 hour, 56 minutes, 29 seconds
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

show chassis sibs | 18835

request chassis sib | 14562
request chassis spmb restart | 14565

show chassis ssb

IN THIS SECTION

- Syntax | 18876
- Description | 18876
- Options | 18876
- Required Privilege Level | 18877
- Output Fields | 18877
- Sample Output | 18878
- Release Information | 18878

Syntax

```
show chassis ssb
<slot>
```

Description

(M20 routers only) Display status information about the System and Switch Board (SSB).

Options

none Display information about all SSBs.

slot (Optional) Display information about the SSB in the specified slot. Replace ***slot*** with **0** or **1**.

Required Privilege Level

view

Output Fields

Table 345 on page 18877 lists the output fields for the`show chassis ssb` command. Output fields are listed in the approximate order in which they appear.

Table 345: show chassis ssb Output Fields

Field Name	Field Description
Failover	Number of times primary role has changed.
Slot	SSB slot number.
State	Current state of the SSB in this slot. State can be any one of the following: • **Master**—SSB is online, operating as primary. • **Backup**—SSB running as backup. • **Empty**—No SSB is present.
Temperature	Temperature of the air passing by the SSB, in degrees Celsius.
CPU utilization	Total percentage of the CPU being used by the SSB's processor.
Interrupt utilization	Of the total CPU being used by the SSB's processor, the percentage being used for interrupts.
Heap utilization	Percentage of heap space being used by the SSB's processor.
Buffer utilization	Percentage of buffer space being used by the SSB's processor.
DRAM	Total DRAM available to the SSB's processor.

Table 345: show chassis ssb Output Fields *(Continued)*

Field Name	Field Description
Start time	Time when the SSB started running.
Uptime	How long the SSB has been up and running.

Sample Output

show chassis ssb

```
user@host> show chassis ssb
SSB status:
  Failover:                      0 time
  Slot 0:
    State:                         Master
    Temperature:                  33 Centigrade
    CPU utilization:               0 percent
    Interrupt utilization:         0 percent
    Heap utilization:              0 percent
    Buffer utilization:            6 percent
    DRAM:                         64 Mbytes
    Start time:                    1999-01-15 22:05:36 UTC
    Uptime:                        21 hours, 21 minutes, 22 seconds
...
```

Release Information

Command introduced before Junos OS Release 7.4.

show chassis synchronization

IN THIS SECTION

- Syntax | 18879
- Syntax (on ACX500 Series) | 18879
- Syntax (on EX9251 Switches) | 18880
- Description | 18880
- Options | 18880
- Required Privilege Level | 18880
- Output Fields | 18881
- Sample Output | 18883
- Release Information | 18893

Syntax

```
show chassis synchronization
<clock-module> <(re0 | re1 | routing-engine (backup | both | local | master | other))>>
<extensive>
<backup | master>
<interface interface-name>
```

Syntax (on ACX500 Series)

```
show chassis synchronization
<gnss>
<extensive>
<backup | master>
```

Syntax (on EX9251 Switches)

```
show chassis synchronization
<clock-module
<extensive>
<interface interface-name>
```

Description

(ACX Series, M320, M40e,M120, T320, T640, and T1600 routers and PTX Series Packet Transport Routers only) Display information about the external clock source currently used for chassis synchronization.

Options

clock-module Display clock module information. You can optionally specify one of these Routing Engine qualifiers:

re0	Routing Engine 0
re1	Routing Engine 1
routing-engine (backup | both | local | master | other)	Routing Engine type

extensive (Optional) Display clock synchronization information in detail.

interface *interface-name* (Optional) Display clock synchronization information for the specified interface.

backup (Optional) Display clock synchronization information about the backup clock.

master (Optional) Display clock synchronization information about the primary clock.

Required Privilege Level

maintenance

Output Fields

Table 346 on page 18881 lists the output fields for the `show chassis synchronization` command. Output fields are listed in the approximate order in which they appear."show chassis synchronization" on page 18879"show chassis synchronization" on page 18879"show chassis synchronization" on page 18879

Table 346: show chassis synchronization Output Fields

Field Name	Field Description
Current state	Indicates current status of external clock sources: • backup—Source is currently the backup clock source. • master—Source is currently the primary clock source. • Online-Master—(PTX Series Packet Transport Routers) Source is the primary clock. Source is online. • Online-Standby—(PTX Series Packet Transport Routers) Source is the standby (backup) clock. Source is online.
Current clock state	Indicates current source of external synchronization: • internal—Source is providing its own clocking. • locked to master CB—(M320, M40e, and M120 routers) Source is locked to primary clock source. • locked to master SCG—(T320, T640, and T1600 routers) Source is locked to primary clock source. • locked to master CCG—(PTX Series Packet Transport Routers) Source is locked to primary clock source.
SNMP trap status	Denotes the SNMP trap generation status (Enabled or Disabled) on ACX Series routers.
Selected for	Number of seconds this clock has been the primary or backup clock source.
Selected since	Timestamp for establishment as primary or backup clock source.

Table 346: show chassis synchronization Output Fields *(Continued)*

Field Name	Field Description
Deviation (in ppm)	Difference in clock timing, in parts per million (ppm). **NOTE**: Starting in Junos OS Release 16.1R2, if the clock source is locked and the deviation exceeds 10 ppm, the status of the configured source is displayed in the command output as `in-use`. This status indicates that the clock source is locked even though the clock has significant deviation.
Last deviation (in ppm)	Previous difference in clock timing, if any, in ppm.
Configured sources	Information about clock sources eligible for selection as primary clock.
Source	Information about external clock sources.
Priority	Indicates priority of external clock sources: • primary—Source is a primary reference. • secondary—Source is a secondary reference.
Deviation (in ppm)	Current difference in clock timing, in ppm: • measuring—Establishing source deviation. • *number*—Deviation in ppm.
Last deviation (in ppm)	Previous difference in clock timing, in ppm: • *number*—Deviation in ppm.
Status	Indicates status of external sources: • present—Source is configured and present. • qualified—Source is eligible for synchronization source.

Sample Output

show chassis synchronization

```
user@host> show chassis synchronization
Clock Synchronization Status :
  Clock module on CB 0
    Current state            : master
    Current clock state      : internal
      Selected for           : 18 hours, 12 minutes, 43 seconds
      Selected since         : 2008-09-10 03:27:47 PDT
      Deviation (in ppm)     : +0.00
      Last deviation (in ppm): +0.00
Clock Synchronization Status :
  Clock module on CB 1
    Current state            : backup
    Current clock state      : locked to master CB
      Selected for           : 1 day, 12 hours, 49 minutes, 20 seconds
      Selected since         : 2008-09-09 08:51:10 PDT
```

show chassis synchronization primary

```
user@host> show chassis synchronization master
Clock Synchronization Status :
  Clock module on CB 0
    Current state            : master
    Current clock state      : internal
      Selected for           : 8 days, 21 minutes, 12 seconds
      Selected since         : 2008-08-27 21:05:40 PDT
      Deviation (in ppm)     : +0.00
      Last deviation (in ppm): +0.00
```

show chassis synchronization backup

```
user@host> show chassis synchronization backup
Clock Synchronization Status :
  Clock module on CB 1
    Current state            : backup
```

```
    Current clock state      : locked to master CB
      Selected for           : 34 days, 20 hours, 17 minutes, 8 seconds
      Selected since         : 2008-08-01 01:22:16 PDT
```

show chassis synchronization extensive

```
user@host> show chassis synchronization extensive
Clock Synchronization Status :
  Clock module on CB 0
    Current state            : master
    Current clock state      : internal
      Selected for           : 8 days, 36 minutes, 29 seconds
      Selected since         : 2008-08-27 21:05:40 PDT
      Deviation (in ppm)     : +0.00
      Last deviation (in ppm): +0.00
Clock Synchronization Status :
  Clock module on CB 1
    Current state            : backup
    Current clock state      : locked to master CB
      Selected for           : 34 days, 20 hours, 19 minutes, 53 seconds
      Selected since         : 2008-08-01 01:22:16 PDT
```

show chassis synchronization (T320, T640, and T1600 Routers)

```
user@host> show chassis synchronization
    Clock Synchronization Status :
      Clock module on SCG 0
        Current state            : master
        Current clock state      : locked to external-a
          Selected for           : 2 hours, 28 minutes, 4 seconds
          Selected since         : 2006-02-17 01:12:58 PST
        Configured sources
          Source      Priority   Deviation    Last deviation   Status
                                 (in ppm)     (in ppm)
          external-a  primary    measuring    -0.10            in-use
          external-b  secondary  -0.10        -0.10            qualified
    Clock Synchronization Status :
      Clock module on SCG 1
        Current state            : backup
        Current clock state      : locked to master SCG
```

```
          Selected for           : 19 hours, 49 minutes, 14 seconds
          Selected since         : 2006-02-16 07:51:48 PST
        Configured sources
          Source      Priority   Deviation    Last deviation   Status
                                 (in ppm)     (in ppm)
          external-a  primary    -0.25        -0.25            qualified
          external-b  secondary  -0.25        -0.25            qualified
```

show chassis synchronization (PTX Series Packet Transport Routers)

```
user@host> show chassis synchronization
Clock Synchronization Status :
  Clock module on CCG 0
    Current state            : Online - Master
    Current clock state      : internal
      Selected for           : 1 hour, 24 minutes, 21 seconds
      Selected since         : 2011-03-21 15:59:37 PDT
      Deviation (in ppm)     : +0.51
      Last deviation (in ppm): +0.51
Clock Synchronization Status :
  Clock module on CCG 1
    Current state            : Online - Standby
    Current clock state      : locked to master CCG
      Selected for           : 1 hour, 39 minutes, 12 seconds
      Selected since         : 2011-03-21 15:44:46 PDT
```

show chassis synchronization clock-module (PTX Series Packet Transport Routers)

```
user@host> show chassis synchronization clock-module
re0:
--------------------------------------------------------------------------
Clock Synchronization Status :
  Clock module on CCG 0
    Current state            : Online - Master
    Current clock state      : locked to bits-a
      Selected for           : 1 minute, 24 seconds
      Selected since         : 2015-06-22 15:01:33 PDT
      Deviation (in ppm)     : unknown
      Last deviation (in ppm): unknown
    Configured sources
```

```
      Source      Priority   Deviation    Last deviation   Status
ro                              (in ppm)     (in ppm)
      bits-a      primary    unknown      unknown          unknown
      fpc-2       secondary  unknown      unknown          unknown
Clock Synchronization Status :
  Clock module on CCG 1
    Current state            : Online - Standby
    Current clock state      : locked to master CCG
      Selected for           : 1 hour, 39 minutes, 52 seconds
      Selected since         : 2015-06-22 13:23:05 PDT
      Deviation (in ppm)     : unknown
      Last deviation (in ppm): unknown
    Configured sources
      Source      Priority   Deviation    Last deviation   Status
                             (in ppm)     (in ppm)
      bits-a      primary    unknown      unknown          unknown
      fpc-2       secondary  unknown      unknown          unknown
```

This following sample output displays the status of configured sources as `in-use#`. This status indicates that the clock source is locked even though the clock has significant deviation in timing.

```
user@host> show chassis synchronization clock-module
re0:
--------------------------------------------------------------------------
Clock Synchronization Status :
  Clock module on CB 0
    Current state            : Online - Master
    Current clock state      : locked to fpc-12
      Selected for           : 7 seconds
      Selected since         : 2016-08-16 09:39:07 PDT
      Deviation (in ppm)     : 10.00
      Last deviation (in ppm): 8.75
    Configured sources
      Source      Priority   Deviation    Last deviation   Status
                             (in ppm)     (in ppm)
      fpc-12      primary    10.00        8.75             in-use#
```

show chassis synchronization clock-module extensive (PTX10008)

```
user@host> show chassis synchronization clock-module extensive

Clock module on SCB-1
  Current role     : master
  Current state    : locked to et-6/0/16:2
    State for      : 0 days, 00 hrs, 07 mins, 36 secs
    State since    : Thu May  6 18:50:50 2021
  Monitored clock sources
    Interface      Type         Status
    et-6/0/16:2    syncE-hybrid qualified-selected
```

show chassis synchronization extensive (ACX Series Routers)

```
user@host> show chassis synchronization extensive

Current clock status  : Locked
Clock locked to       : Primary
SNMP trap status : Enabled

Configured sources:

Interface           : ge-0/0/7
Status              : Secondary     Index       : 136
Clock source state  : Clk qualified Priority    : 3
Configured QL       : SEC           ESMC QL     : PRC
Clock source type   : ifd           Clock Event : Clock qualified
Interface State     : Up,sec,ESMC Rx(SSM 0x2),ESMC TX(QL PRC/SSM 0x2),

Interface           : ge-0/1/1
Status              : Primary       Index       : 138
Clock source state  : Clk qualified Priority    : 2
Configured QL       : SEC           ESMC QL     : PRC
Clock source type   : ifd           Clock Event : Clock locked
Interface State     : Up,pri,ESMC Rx(SSM 0x2),ESMC TX(QL DNU/SSM 0xf)
```

show chassis synchronization extensive with extended ESMC QL

```
user@host>show chassis synchronization extensive
Current clock status  : Locked
Clock locked to       : Primary
SNMP trap status  : Disabled

Configured sources:

Interface              : et-0/0/20
Status                 : Secondary     Index       : 1023
Clock source state     : Clk qualified Priority    : 2
Configured QL          : PRTC          ESMC QL     : PRTC
Clock source type      : ifd           Clock Event : Clock qualified
Wait-to-restore        : 0  min        Hold-off    : 1000 ms
Interface State        : Up,sec,ESMC Rx(QL PRTC/SSM 0x20),ESMC TX(QL ePRTC/SSM 0x21),

Interface              : et-0/0/21
Status                 : Primary       Index        : 1024
Clock source state     : Clk qualified Priority     : 1
Configured QL          : ePRTC         ESMC QL      : ePRTC
Clock source type      : ifd           Clock Event  : Clock locked
Wait-to-restore        : 0  min        Hold-off     : 1000 ms
Interface State        : Up,pri,ESMC Rx(QL ePRTC/SSM 0x21),ESMC TX(QL DNU/SSM 0xf)
```

show chassis synchronization extensive (ACX500 Series Routers)

```
user@host> show chassis synchronization extensive
Current clock status  : Locked
Clock locked to           : Primary

Configured ports:

Name                : gnss
Current ToD         : Mon Aug 10 08:50:52 2015
Last ToD update     : Mon Aug 10 08:50:51 2015
GPS receiver status : Synchronized
UTC Pending         : FALSE
UTC Offset          : 36
```

```
One PPS status : Active
```

show chassis synchronization extensive (PTX10008)

```
user@host> show chassis synchronization extensive

Current clock status  : Locked
Clock locked to       : Primary
SNMP trap status      : Disabled

Configured sources:

Interface           : et-6/0/16:2
Status              : Primary       Index       : 1035
Clock source state  : Clk qualified Priority    : Default(8)
Configured QL       : PRC           ESMC QL     : PRC
Clock source type   : ifd           Clock Event : Clock locked
Wait-to-restore     : 0  min        Hold-off    : 1000 ms
Interface State     : Up,pri,ESMC Rx(QL PRC/SSM 0x2),ESMC TX(QL DNU/SSM 0xf)
```

show chassis synchronization gnss (ACX500 Series Routers)

```
user@host> show chassis synchronization gnss

Receiver Status          : Good
Constellation            : GPS & GLONASS
Cable delay compensation : 0
Antenna Status           : GPS Antenna Good
```

show chassis synchronization gnss extensive (ACX500, ACX7024, ACX7332, ACX7348)

```
user@host> show chassis synchronization gnss extensive

Receiver Status          : Good
Constellation            : GPS & GLONASS
Cable delay compensation : 0
```

```
Antenna Status           : GPS Antenna Good
Position                 : 12 56' 35.740410'' N   : 77 41' 32.228191'' E
Altitude                 : 841 meters
Number of satellites     : 17
Satellite List:

Sat No    Signal Level     Status        Type          Mode(T=Timing/P=Position)
17        34    dBHZ       Acquired      GPS           T,P
24        38    dBHZ       Acquired      GPS           T,P
12        31    dBHZ       Acquired      GPS           T,P
2         37    dBHZ       Acquired      GPS           T,P
5         48    dBHZ       Acquired      GPS           T,P
20        35    dBHZ       Acquired      GPS           T,P
28        42    dBHZ       Acquired      GPS           T,P
6         40    dBHZ       Acquired      GPS           T,P
71        41    dBHZ       Acquired      GLONASS       T,P
70        46    dBHZ       Acquired      GLONASS       T,P
76        36    dBHZ       Acquired      GLONASS       T,P
75        45    dBHZ       Acquired      GLONASS       T,P
72        40    dBHZ       Acquired      GLONASS       T,P
15        18    dBHZ       Acquired      GPS           NONE
13        36    dBHZ       Acquired      GPS           NONE
74        41    dBHZ       Acquired      GLONASS       NONE
77        19    dBHZ       Acquired      GLONASS       NONE
```

show chassis synchronization clock-module (EX9251 Switches)

```
user@switch> show chassis synchronization clock-module
re0:
--------------------------------------------------------------------------

Clock module on SCB0
  Current role     : master
```

Show chassis synchronization (external 1pps) (ACX710 Router)

```
user@host> show chassis synchronization extensive
Current clock status  : Locked
Clock locked to       : Primary
SNMP trap status : Enabled
```

```
Configured sources:

Interface           : ge-0/1/1
Status              : Primary       Index       : 138
Clock source state  : Clk qualified Priority    : 2
Configured QL       : SEC           ESMC QL     : PRC
Clock source type   : ifd           Clock Event : Clock locked
Interface State     : Up,pri,ESMC Rx(SSM 0x2),ESMC TX(QL DNU/SSM 0xf)
```

Show chassis synchronization gru extensive (ACX710 Router)

```
user@host> Show chassis synchronization gru extensive
Receiver Status          : Good
Constellation            : GPS & GLONASS
Cable delay compensation : 0
+ 1PPS Loopback          : {Enabled|Not enabled}
+ 1PPS Counts            : 42
+ One PPS status         : Active
+ ToD Protocol           : {No-protocol|ericsson-combined}
+ Reset counts           : 1
+ Status                 : {Free running|1PPS Not synchronized|GPS locked}
+ Status last received   : 21:05:02 21 May 2019
+ Antenna DC Status      : 2 (Nominal current detected)
+ GPS Quality            : {Positioning not started|Standalone|Differential}
+ Self Survey Length     : 2000 sec
+ Antenna Cable Delay    : 0 ns
+ Current ToD            : Mon Aug 10 08:50:52 2015
+ Current TAI            : Mon Aug 10 08:55:52 2015
+ Current UTC            : Mon Aug 10 08:50:52 2015
+ Client                 : D821713120/2 R6K (In Control)
+ Client in control      : D821713120 R6K
+ Pps SM mode            : 2 (Position Hold)
+ Product designation    : GRU 04 01
+ Product No.            : NCD 901 65/1
+ Product Rev.           : R1C
+ Additional Information : E640000065.1851
+ GRM hardware ID        : 3052
+ GRM serial number      : 0
+ GRM manf date          : 12.19.2018
```

```
+ GRM firmware version   : 0.08.00
+  Received sentences:
+    PERC, GPppr : 2811
+    PERC, GPsts : 2811
+    PERC, GPver : 47
+    PERC, GPrst : 0
+    PERC, GPssl : 14
+    PERC, GPdly : 14
+    PERC, GPavp : 47
+    PERC, GPtsp : 42
+    GPGSV       : 141
+    GPGSA       : 47
+    BDGSV       : 0
+    BDGSA       : 0
+    GLGSV       : 130
+    GLGSA       : 47
+    GAGSV       : 0
+    GAGSA       : 0
+    QZGSV       : 0
+    QZGSA       : 0
+    GPGGA       : 47
+    GPCTR       : 57
+    Other       : 2886
+Protocol Errors:
+  Unexpected character   : 0
+  Checksum error         : 0
+  Format error           : 0
+  Other error            : 0
Number of satellites     : 17
Satellite List:
Sat No    Signal Level     Status        Type          Mode(T=Timing/P=Position)
17        34    dBHZ       Acquired      GPS           T,P
24        38    dBHZ       Acquired      GPS           T,P
12        31    dBHZ       Acquired      GPS           T,P
2         37    dBHZ       Acquired      GPS           T,P
5         48    dBHZ       Acquired      GPS           T,P
20        35    dBHZ       Acquired      GPS           T,P
28        42    dBHZ       Acquired      GPS           T,P
6         40    dBHZ       Acquired      GPS           T,P
71        41    dBHZ       Acquired      GLONASS       T,P
70        46    dBHZ       Acquired      GLONASS       T,P
76        36    dBHZ       Acquired      GLONASS       T,P
75        45    dBHZ       Acquired      GLONASS       T,P
```

```
72        40    dBHZ       Acquired      GLONASS       T,P
15        18    dBHZ       Acquired      GPS           NONE
13        36    dBHZ       Acquired      GPS           NONE
74        41    dBHZ       Acquired      GLONASS       NONE
77        19    dBHZ       Acquired      GLONASS       NONE
```

Show chassis synchronization gru (ACX710 Router)

```
user@host> Show chassis synchronization gru
Current ToD         : Mon Aug 10 08:50:52 2015
Current TAI         : Mon Aug 10 08:55:52 2015
Current UTC         : Mon Aug 10 08:50:52 2015

One PPS status : Active
```

show chassis synchronization clock-module (ACX7100 Router)

```
user@host> Show chassis synchronization clock-module
Clock module on SCB0
  Current role     : master
  Current state    : freerun
    State for      : 0 days, 09 hrs, 40 mins, 54 secs
    State since    : Tue Aug  6 23:57:51 2024
```

show chassis synchronization clock-module (ACX Series)

```
user@host> show chassis synchronization clock-module
Clock module on SCB0
  Current role     : master
  Current state    : freerun
    State for      : 0 days, 09 hrs, 40 mins, 54 secs
    State since    : Tue Aug  6 23:57:51 2024
```

Release Information

Command introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

request chassis synchronization switch | 14574

Configuring Clock Synchronization Interface on MX Series Routers

Supported Time Synchronization Standard

Configuring External Clock Synchronization for ACX Series Routers

show chassis synchronization (MX Series Routers)

IN THIS SECTION

- Syntax | 18894
- Description | 18895
- Options | 18896
- Required Privilege Level | 18896
- Output Fields | 18896
- Sample Output | 18902
- Sample Output | 18903
- Sample Output | 18903
- Sample Output | 18904
- Sample Output | 18904
- Sample Output | 18905
- Sample Output | 18905
- Sample Output | 18907
- Sample Output | 18907
- Sample Output | 18907
- Release Information | 18910

Syntax

```
show chassis synchronization
<clock-module <(re0 | re1 | routing-engine (backup | both | local | master | other))>>
```

```
<extensive>
<interface interface-name>
```

Description

Display information about clocks used for chassis synchronization.

NOTE: In hybrid mode, the EEC in the MPC derives frequency synchronization from Synchronous Ethernet and the phase and time of day from PTP; however, the `show chassis synchronization extensive` operational mode command output displays the lock status that is derived from the EEC located on the SCB.

NOTE: The Switch Control Board (SCB) framer in MX Series routers supports only the first-generation Synchronization Status Message (SSM) format. Therefore, whenever the router needs to transmit an SSM value of st3e or tnc via an external interface, an SSM value of st3 is transmitted. However, on a Synchronous Ethernet interface, an ESMC packet with the unadjusted SSM is transmitted. The term *unadjusted* here means:

- If the `receive-quality` statement at the `[edit chassis synchronization selection-mode]` hierarchy level is configured, the originally received SSM value st3e or tnc (corresponding to the currently active Synchronous Ethernet clock interface) is transmitted.
- If the `configured quality` statement at the `[edit chassis synchronization selection-mode]` hierarchy level is configured, the originally configured SSM value of st3e or tnc (corresponding to the currently active Synchronous Ethernet clock interface) is transmitted.

Note that when the external interface receives an SSM value of either st3e or tnc, the SCB framer does not recognize either of these SSM codes, and therefore, it reports that the Do Not Use (DNU) quality value has been received.

SCBE2 does not support SSM when you configure the following framing format values on the external interface at the `[edit chassis synchronization interfaces external]` hierarchy level:

- sf for the `t1-options` statement
- g704-no-crc for the `e1-options` statement

Options

clock-module (MX240, MX480, MX960, MX10003, MX204, MX2010, MX2020, and MX2008 routers with Enhanced MX Switch Control Board;) (Optional) Display clock module information. You can optionally specify one of these Routing Engine qualifiers:

re0	Routing Engine 0
re1	Routing Engine 1
routing-engine (backup | both | local | master | other)	Routing Engine type

extensive (Optional) Display clock synchronization information in detail.

interface ***interface-name*** (Optional) Display clock synchronization information for the specified interface.

Required Privilege Level

maintenance

Output Fields

Table 347 on page 18896 lists the output fields for the `show chassis synchronization` command. Output fields are listed in the approximate order in which they appear.

Table 347: show chassis synchronization Output Fields

Field Name	Field Description	Level of Output
`Current clock status`	Indicates the current status of chassis synchronization: • `Locked`—Clock is operational. • `Holdover`—Clock is not operational. • `Freerun`—Clock is locked to the free-run local oscillator.	none
`Clock locked to`	Indicates whether the clock is locked to either the `primary` source or the `secondary` source.	none

Table 347: show chassis synchronization Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Configured sources`	Heading for the list of interfaces configured for chassis synchronization and their subsequent status indicators.	none
`Source name`	Indicates the configured interface that is the source. The `external` source name indicates the external clock interface.	none
`Configured Priority`	Indicates the priority configured for the interface.	none
`Interface Status`	Indicates the status of the interface as `primary`, `secondary`, or `n/a` (external).	none
`Configured quality`	Indicates the configured quality of the interface. • `prs` —Primary reference source—Stratum 1 • `st2`—Stratum 2 • `tnc`—Transit node clock • `st3e`—Stratum 3E • `st3`—Stratum 3 • `smc`—SONET minimum clock • `st4`—Stratum 4 • `prc`—Primary reference clock • `ssu-a`—Synchronization supply unit A • `ssu-b`—Synchronization supply unit B • `sec`—SDH equipment clock	none
`Interface`	Indicates the configured interface: • `ge-`*fpc*/*pic*/*port*—Indicates the interface type and which FPC, PIC, and port are configured.	`extensive`

Table 347: show chassis synchronization Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Status`	Indicates the synchronization status of the indicated interface, as follows: • `Primary`—This interface is the selected primary chassis clock source. • `Secondary`—This interface is the selected secondary chassis clock source. • `n/a`—This interface is not a selected clock source.	`extensive`
`Index`	Unique numeric identifier for the established Synchronous Ethernet configuration.	`extensive`
`Clock source state`	Indicates the status of the Synchronous Ethernet clock source: • `Clk qualified`—The Synchronous Ethernet clock source is qualified. • `n/a`—The Synchronous Ethernet clock source is not qualified.	`extensive`
`Priority`	Indicates the configured priority. The range is from 1 through 5. The following values indicate whether the parameter is not specified or undefined: • `Default(8)`—The parameter is not specified. • `-` —The parameter is undefined or out of range.	`extensive`
`Configured QL`	Indicates the configured source interface quality level (QL), which is dependent on the source interface and option. The following quality levels are supported and the configured QL is indicated: • `prc | ssu-a | ssu-b | sec`—Network option-1 QLs • `prs | st2 | tnc | st3e | st3 | smc | st4`—Network option-2 QLs	`extensive`
`ESMC QL`	Indicates the configured Ethernet Synchronization Message Channel (ESMC) quality level: • `DNU`—Network option-1 source • `DSU`—Network option-2 source	`extensive`

Table 347: show chassis synchronization Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Clock source type`	Indicates that the configured chassis synchronization clock source is one of the following types: • `ifd`—Uses the free-run local oscillator. • `extern`—Uses a configured qualified clock source.	`extensive`
`Clock Event`	Indicates the event clock status: • `Clock locked`—Clock is established. • `n/a`—Clock is not established.	`extensive`
Configuration flags	Indicates `Ext` for external interface configuration	
`Wait-to-restore time`	Indicates the amount of time in minutes for the port signal to be up before the port is opened to receive and transmit Ethernet Synchronization messages.	`extensive`
`Hold-off time`	Indicates the amount of time in milliseconds for hold off time for Synchronous Ethernet interfaces and external clock source interfaces to prevent rapid successive switching. If an interface goes down, hold-off time delays short signal failures from being sent to the clock selection process.	`extensive`
`Kernel flags`	Indicates the Synchronous Ethernet software operational status: • `Up`—The Synchronous Ethernet software is operational for the configured interface. • `pri`—The source is the selected primary clock source. • `Dn`—The Synchronous Ethernet software is not operational for the configured interface.	`extensive`

Table 347: show chassis synchronization Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Ineligibility reason`	Indicates the reason the interface is ineligible for the Synchronous Ethernet operation, including the following: • `Link Down`—The link between the Synchronous Ethernet interfaces is not operational. • `Not restored`—The Synchronous Ethernet link has not yet been restored because it is waiting for the specified `wait-to-restore` time to elapse. • `Forbidden slot`—Slot 10 is not supported. • `Interface unit missing`—The unit parameter is not set or is invalid. • `Locked`—The paired interface is not available. • `No cfg`—Synchronous Ethernet is not configured. • `RX Disabled`—The receiving interface is disabled. • `Undefined/invalid QL`—The QL mode is not specified in the configuration or, if specified, is not supported. **NOTE**: When GPS is configured as the clock source and the `quality-mode-enable` statement is configured, the `Current Lock Status` is displayed as `Locked` but `Undefined/invalid QL` is displayed as the Ineligibility reason. As a workaround, configure the `no-ssm-support` statement at the `[edit chassis synchronization source interfaces interface-name]` hierarchy level. • `System initialization in progress`—The remote system is performing initialization and not currently available for synchronization. • `Unsupported interface`—The configured interface does not support Synchronous Ethernet.	`extensive`
`Clock module on`	Indicates whether the clock module is on the Switch Control Board `SCB0` or `SCB1`.	`clock-module`

Table 347: show chassis synchronization Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Current role`	Indicates the role of the clock module: • `master`—The clock module is on the primary SCB, which is the active chassis clock source. • `backup`—The clock module is on the backup SCB, which mirrors the state of the active clock.	`clock-module`
`Current state`	Indicates the state of the clock module: • `freerun`—The clock module is in free-run mode. When the system starts up, the default clock module state is free-run. • `acquiring-lock on`—The clock module is attempting to acquire a lock on the specified clock source. • `locked to`—The clock module is locked to the specified clock source. • `holdover on`—The clock module is in holdover mode on the specified clock source. Prior to the specified clock source becoming invalid, the clock module was locked on the source and holdover data was collected. • `holdover`—The clock module has transitioned into holdover prior to locking on a valid clock source and collecting holdover data.	`clock-module`
`Monitored clock sources`	Displays information about monitored clock sources.	`clock-module`
`Interface`	Indicates the interface type and which FPC, PIC, and port are configured: • `external`—External clock source • `ge-`*`fpc`*`/`*`pic`*`/`*`port`*—Line Synchronous Ethernet or PTP client • `xe-`*`fpc`*`/`*`pic`*`/`*`port`*—Line Synchronous Ethernet or PTP client	`clock-module`

Table 347: show chassis synchronization Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Type`	Indicates the type of clock source: • `t1`—BITS T1 framed • `e1`—BITS E1 framed • `2048khz`—BITS unframed 2048 KHz frequency source • `syncE`—Synchronous Ethernet frequency source • `ptp`—PTP client source • `ptp-hybrid`—PTP client source using Synchronous Ethernet for frequency	`clock-module`
`Status`	Indicates the status of the clock source: • `failed`—The clock source is in the failed state. • `qualifying`—The clock source is being qualified. • `qualified`—The clock source is qualified and can be selected as the chassis clock source. • `qualified-selected`—The clock source is qualified and selected as the chassis clock source.	`clock-module`

Sample Output

show chassis synchronization

```
user@host> show chassis synchronization
Current clock status: Locked
Clock locked to : Primary

Configured sources
Source      Configured       Interface       Configured
Name        Priority         Status          Quality
```

```
ge-1/0/0       -             Primary         PRC
```

Sample Output

show chassis synchronization extensive

```
user@host> show chassis synchronization extensive
Current clock status: Locked
Clock locked to            : Primary

Configured sources:
Interface           : ge-1/0/0
Status              : Primary         Index : 143
Clock source state  : Clk qualified   Priority : -
Configured QL       : PRC             ESMC QL : DNU
Clock source type   : ifd             Clock Event : Clock locked
Kernel flags        : Up,pri,
Wait-to-restore time            : 1
Hold-off time             :  1200
Interface State           : Up,ESMC TX(QL SEC/SSM 0xb),
```

Sample Output

show chassis synchronization extensive (Synchronous Ethernet with link down)

```
user@host> show chassis synchronization extensive
Current clock status  : Holdover
Configured sources:

Interface           : ge-1/0/2
Status              : n/a           Index       : 142
Clock source state  : n/a           Priority    : Default(8)
Configured QL       : SSU-B         ESMC QL     : DNU
Clock source type   : ifd           Clock Event : n/a
Kernel flags        : Dn,
Ineligibility reason: Link Down,
Wait-to-restore time            : 1
```

```
Hold-off time             :  1200
Interface State           : Up,ESMC TX(QL SEC/SSM 0xb),
```

Sample Output

show chassis synchronization extensive (Synchronous Ethernet with physical interface not restored)

```
user@host> show chassis synchronization extensive
Current clock status  : Holdover
Configured sources:

Interface           : ge-1/0/2
Status              : n/a           Index       : 142
Clock source state  : n/a           Priority    : Default(8)
Configured QL       : SSU-B         ESMC QL     : DNU
Clock source type   : ifd           Clock Event : n/a
Kernel flags        : Restoring in 13s,ESMC TX(QL DNU/SSM 0xf),
Ineligibility reason: Not restored,
Wait-to-restore time            : 1
Hold-off time             :  1200
Interface State           : Up,ESMC TX(QL SEC/SSM 0xb),
```

Sample Output

show chassis synchronization extensive (Synchronous Ethernet configured on ineligible slot 10)

```
user@host> show chassis synchronization extensive
Current clock status  : Holdover
Configured sources:

Interface           : ge-10/0/2  # Note: configuration 10/x/y (slot 10), which does not support
Synchronous Ethernet
Status              : n/a           Index       : 142
Clock source state  : n/a           Priority    : Default(8)
Configured QL       : SSU-B         ESMC QL     : DNU
Clock source type   : ifd           Clock Event : n/a
Kernel flags        : Up,
```

```
Ineligibility reason: Forbidden slot,
Wait-to-restore time            : 1
Hold-off time             :  1200
Interface State           : Up,ESMC TX(QL SEC/SSM 0xb),
```

Sample Output

show chassis synchronization extensive (on MX104 Router)

```
user@host> show chassis synchronization extensive
Current clock status  : Locked
Clock locked to       : Primary

Configured interfaces:

Name         : bits
Signal type  : e1 (g704 hdb3 sa4)
Rx status    : active
Tx status    : active

Configured sources:

Interface           : bits
Status              : Primary       Index       : 2
Clock source state  : Clk qualified Priority    : Default(7)
Configured QL       : SEC           ESMC QL     : SEC
Clock source type   : extern        Clock Event : Clock locked
Wait-to-restore     : 5  min        Hold-off    : 1000 ms
Interface State     : Up,pri,ESMC Rx(SSM 0xb)
```

Sample Output

show chassis synchronization clock-module (MX10003 Router)

```
user@host> show chassis synchronization clock-module
re0:
--------------------------------------------------------------------------

Clock module on SCB0
```

```
  Current role     : master
  Current state    : holdover
    State for      : 11 days, 01 hrs, 20 mins, 33 secs
    State since    : Sun Jan 29 20:11:48 2023

user@host> show chassis synchronization extensive

Current clock status : Locked
Clock locked to : Primary
SNMP trap status : Disabled

Configured sources:

Interface : et-0/1/3
Status : Secondary Index : 563
Clock source state : Clk qualified Priority : 1
Configured QL : PRC ESMC QL : DNU
Clock source type : ifd Clock Event : Clock qualified
Wait-to-restore : 0 min Hold-off : 1000 ms
Interface State : Up,sec,ESMC TX(QL PRC/SSM 0x2),

Interface : et-0/1/7
Status : n/a Index : 567
Clock source state : n/a Priority : 3
Configured QL : PRC ESMC QL : DNU
Clock source type : ifd Clock Event : n/a
Wait-to-restore : 0 min Hold-off : 1000 ms
Interface State : Up,ESMC TX(QL PRC/SSM 0x2),

Interface : et-0/1/8
Status : Primary Index : 568
Clock source state : Clk qualified Priority : 1
Configured QL : PRC ESMC QL : DNU
Clock source type : ifd Clock Event : Clock locked
Wait-to-restore : 0 min Hold-off : 1000 ms
Interface State : Up,pri,ESMC TX(QL DNU/SSM 0xf),

Interface : et-1/1/1
Status : n/a Index : 555
Clock source state : n/a Priority : 5
Configured QL : PRC ESMC QL : DNU
Clock source type : ifd Clock Event : n/a
Wait-to-restore : 0 min Hold-off : 1000 ms
```

```
Interface State : Dn,
Ineligibility reason: Link Down
```

Sample Output

show chassis synchronization clock-module (MX204 Router)

```
user@host> show chassis synchronization clock-module

re0:
--------------------------------------------------------------------------

Clock module on SCB0
  Current role     : master
```

Sample Output

show chassis synchronization interface

```
user@host> show chassis synchronization interface ge-1/0/2
Current clock status  : Locked
Clock locked to       : Primary
```

Sample Output

show chassis synchronization clock-module

```
user@host> show chassis synchronization clock-module
Clock module on SCB0
  Current role     : master
  Current state    : locked to ge-4/1/0
    State for      : 0 days, 00 hrs, 00 mins, 15 secs
    State since    : Mon Jun  6 07:29:40 2011
  Monitored clock sources
    Interface      Type         Status
    ge-4/1/0       syncE        qualified-selected
```

```
    ge-4/3/0       syncE        qualified
```

show chassis synchronization clock-module (MX10008 with MX10k-LC2101 line cards)

```
user@host> show chassis synchronization clock-module
Clock module on SCB0
Current role : master
Current state : locked to external-0
State for : 0 days, 00 hrs, 01 mins, 16 secs
State since : Tue May 17 02:44:36 2022
Monitored clock sources
Interface Type Status
et-0/3/1 syncE-hybrid qualified
external-0 bits-hybrid qualified-selected
```

show chassis synchronization (configured external clock interface)

```
user@host> show chassis synchronization
Current clock status  : Free-run

Configured interfaces:
Name            Signal type                Rx status      Tx status
external        e1 (g704 ami sa4)          loss of signal squelched

Configured outputs:
Interface       Tx status   Minimum QL   Tx QL
external        squelched   SEC          DNU

Configured sources:
Source         Configured    Interface   Configured
 Name           Priority     Status      Quality
external       Default(6)    n/a         SSU-A
```

show chassis synchronization clock-module (configured external clock interface)

```
user@host> show chassis synchronization clock-module
re0:
--------------------------------------------------------------------------

Clock module on SCB0
  Current role     : master
  Current state    : freerun
    State for      : 2 days, 06 hrs, 16 mins, 57 secs
    State since    : Wed Nov 14 08:02:07 2012
  Monitored clock sources
    Interface      Type         Status
    external       e1           failed
```

show chassis synchronization clock-module(configured external clock interfaces)

```
user@host> show chassis synchronization clock-module
re0:
--------------------------------------------------------------------------

Clock module on SCB0
Current role : master
Current state : locked to external-0
State for : 0 days, 00 hrs, 08 mins, 27 secs
State since : Tue May 17 02:44:36 2022
Monitored clock sources
Interface Type Status
external-0 bits-hybrid qualified-selected
```

show chassis synchronization extensive (configured external clock interface)

```
user@host> show chassis synchronization extensive
Current clock status  : Free-run

Configured interfaces:

Name         : external
```

```
Signal type  : e1 (g704 ami sa4)
Rx status    : loss of signal
Tx status    : squelched
LED color    : red

Configured outputs:

Interface      : external
Tx status      : squelched (holdover data invalid)
Minimum QL     : SEC             Tx QL         : DNU
Holdover mode  : enabled         Wander filter : disabled
Source mode    : chassis         Source Tx DNU : disabled
Holdover data  : invalid
Current state  : holdover
  State for    : 2 days, 06 hrs, 03 mins, 46 secs
  State since  : Wed Nov 14 08:02:09 2012

Configured sources:

Interface           : external
Status              : n/a           Index       : 0
Clock source state  : n/a           Priority    : Default(6)
Configured QL       : SSU-A         ESMC QL     : DNU
Clock source type   : extern        Clock Event : n/a
Interface State     : Dn,ESMC Rx(SSM 0xf),
Ineligibility reason: Link Down,
```

Release Information

Command introduced in Junos OS Release 10.4.

`clock-module` option introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring an External Clock Synchronization Interface for MX Series Routers

Configuring External Clock Synchronization for ACX Series Routers

Example: Configuring Framing Mode for Synchronous Ethernet on MX Series Routers with 10-Gigabit Ethernet MIC

request chassis synchronization mode | **14570**

show chassis synchronization | 18879
synchronization (MX Series) | 28795
Synchronous Ethernet Overview

show chassis temperature-thresholds (Security)

IN THIS SECTION

- Syntax | 18911
- Description | 18911
- Required Privilege Level | 18911
- Output Fields | 18911
- Sample Output | 18913
- Release Information | 18913

Syntax

```
show chassis temperature-thresholds
```

Description

Display chassis temperature threshold settings, in degrees Celsius.

Required Privilege Level

View

Output Fields

Table 348 on page 18912 lists the output fields for the `show chassis temperature-thresholds` command. Output fields are listed in the approximate order in which they appear.

Table 348: show chassis temperature-thresholds Output Fields

Field name	Field Description
Item	Chassis component. If per FRU per slot thresholds are configured, the components about which information is displayed include the chassis, the Routing Engines, FPCs, and FEBs. If per FRU per slot thresholds are not configured, the components about which information is displayed include the chassis and the Routing Engines.
Chassis Default	Default values of the chassis in degrees Celsius.
Routing Engine	Default values of the routing engine in degrees Celsius.
Fan speed	The fan speed changes at the threshold when going from a low speed to a higher speed. When the fan speed changes from a higher speed to a lower speed, the temperature changes two degrees below the threshold. Temperature threshold settings, in degrees Celsius, for the fans to operate at normal and high speeds. • **Normal**—The fans operate at normal speed if the component is at or below this temperature and all the fans are present and functioning normally. • **High**—The fans operate at high speed if the component has exceeded this temperature or a fan has failed or is missing. An alarm is not triggered until the temperature exceeds the threshold settings for a yellow alarm or a red alarm.
Yellow alarm	Temperature threshold settings, in degrees Celsius, that trigger a yellow alarm. • **Normal**—The temperature that must be exceeded on the component to trigger a yellow alarm when the fans are running at full speed. • **Bad fan**—The temperature that must be exceeded on the component to trigger a yellow alarm when one or more fans have failed or are missing.

Table 348: show chassis temperature-thresholds Output Fields *(Continued)*

Field name	Field Description
Red alarm	Temperature threshold settings, in degrees Celsius, that trigger a red alarm. • **Normal**—The temperature that must be exceeded on the component to trigger a red alarm when the fans are running at full speed. • **Bad fan**—The temperature that must be exceeded on the component to trigger a red alarm when one or more fans have failed or are missing.
Fire shutdown	Temperature threshold settings, in degrees Celsius, for the device to shut down.

Sample Output

show chassis temperature-thresholds

```
user@host> show chassis temperature-thresholds
                                        Fan speed      Yellow alarm      Red alarm      Fire
Shutdown
                                       (degrees C)      (degrees C)     (degrees C)
(degrees C)
Item                                  Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Chassis default                           48    54       65       55       75       65        100
Routing Engine 0                          70    80       95       95      110      110        112
FPC 0                                     55    60       75       65       90       80         95
FPC 2                                     55    60       75       65       90       80         95
```

Release Information

Command introduced in Junos OS Release 9.0

show chassis temperature-thresholds

IN THIS SECTION

- Syntax | 18914
- Syntax (TX Matrix Routers) | 18914
- Syntax (TX Matrix Plus Routers) | 18914
- Syntax (MX Series Routers) | 18915
- Syntax (QFX Series) | 18915
- Description | 18915
- Options | 18915
- Required Privilege Level | 18916
- Output Fields | 18916
- Sample Output | 18918
- Release Information | 18945

Syntax

```
show chassis temperature-thresholds
```

Syntax (TX Matrix Routers)

```
show chassis temperature-thresholds
<lcc number | scc>
```

Syntax (TX Matrix Plus Routers)

```
show chassis temperature-thresholds
<lcc number | sfc number>
```

Syntax (MX Series Routers)

```
show chassis temperature-thresholds
<all-members>
<local>
<member member-id>
<satellite [slot-id slot-ID |device-alias alias-name]>
```

Syntax (QFX Series)

```
show chassis temperature-thresholds
<interconnect-device name>
<node-device name>
```

Description

Display chassis temperature threshold settings, in degrees Celsius. You can run the `show chassis enhanced-temperature-thresholds` command on a SRX4700 device to view the bad fan temperatures along with other temperature details.

Options

none Display the temperature threshold details.

all-members (MX Series routers only) (Optional) Display the chassis temperature threshold settings of all member routers in the Virtual Chassis configuration.

interconnect-device *name* (QFabric systems only) (Optional) Display the chassis temperature threshold settings of the Interconnect device.

lcc *number* (TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the temperature threshold details of a specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display the temperature threshold details of a specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Display the chassis temperature threshold settings of the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display the chassis temperature threshold settings of the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device *name* (QFabric systems only) (Optional) Display the chassis temperature threshold settings of the Node device.

satellite [slot-id *slot-ID* | device-alias *alias-name*] (Junos Fusion only) (Optional) Display the chassis temperature threshold settings for the specified satellite device or devices in a Junos Fusion, or for all satellite devices if no satellite devices are specified.

scc (TX Matrix routers only) (Optional) Display the temperature threshold details of the TX Matrix router (switch-card chassis).

sfc *number* (TX Matrix Plus routers only) (Optional) On TX Matrix Plus routers, display the temperature threshold details of the TX Matrix Plus router, which is the switch-fabric chassis. Replace *number* with 0.

Required Privilege Level

view

Output Fields

Table 349 on page 18917 lists the output fields for the `show chassis temperature-thresholds` command. Output fields are listed in the approximate order in which they appear.

Table 349: show chassis temperature-thresholds Output Fields

Field name	Field Description
Item	Chassis component. If per FRU per slot thresholds are configured, the components about which information is displayed include the chassis, the Routing Engines, CPUs, FPCs, and FEBs. If per FRU per slot thresholds are not configured, the components about which information is displayed include the chassis and the Routing Engines.
Fan speed	**NOTE**: On the QFX3500 switch and QFX3600 switch, there are four fan speeds: `low`, `medium-low`, `medium-high`, and `high`. The fan speed changes at the threshold when going from a low speed to a higher speed. When the fan speed changes from a higher speed to a lower speed, the temperature changes two degrees below the threshold. Temperature threshold settings, in degrees Celsius, for the fans to operate at normal and high speeds. • Normal—The fans operate at normal speed if the component is at or below this temperature and all the fans are present and functioning normally. **NOTE**: On a TX Matrix Plus router with 3D SIBs, the threshold temperature at the XF junction is set to 70°C for `Normal` fan speed, which is less than or equal to 4800 RPM. • High—The fans operate at high speed if the component has exceeded this temperature or a fan has failed or is missing. **NOTE**: On a TX Matrix Plus router with 3D SIBs, the threshold temperature at the XF junction is set to 75°C for `High` fan speed, which is greater than or equal to 5000 RPM. **NOTE**: For MX480 Routers, there are three fan speeds: Low, Medium, and High. An alarm is not triggered until the temperature exceeds the threshold settings for a yellow alarm or a red alarm.
Yellow alarm	Temperature threshold settings, in degrees Celsius, that trigger a yellow alarm. • Normal—The temperature that must be exceeded on the component to trigger a yellow alarm when the fans are running at full speed. • Bad fan—The temperature that must be exceeded on the component to trigger a yellow alarm when one or more fans have failed or are missing.

Table 349: show chassis temperature-thresholds Output Fields *(Continued)*

Field name	Field Description
Red alarm	Temperature threshold settings, in degrees Celsius, that trigger a red alarm. • Normal—The temperature that must be exceeded on the component to trigger a red alarm when the fans are running at full speed. • Bad fan—The temperature that must be exceeded on the component to trigger a red alarm when one or more fans have failed or are missing.
Fire Shutdown	(T4000 routers, TX Matrix Plus router with 3D SIBs, and PTX Series Packet Transport Routers only)—Temperature threshold settings, in degrees Celsius, for the network device to shut down.

Sample Output

show chassis temperature-thresholds

```
user@host> show chassis temperature-thresholds

                          Fan speed        Yellow alarm         Red alarm
                         (degrees C)        (degrees C)        (degrees C)
Item                    Normal   High    Normal   Bad fan    Normal   Bad fan
Chassis default             48     54        65        55        75        65
Routing Engine 0            70     80        95        95       110       110
Routing Engine 1            70     80        95        95       110       110
FPC 0                       55     60        75        65        90        80
FPC 1                       55     60        75        65        90        80
FPC 2                       55     60        75        65        90        80
FPC 3                       55     60        75        65        90        80
FPC 4                       55     60        75        65        90        80
FPC 5                       55     60        75        65        90        80
FPC 6                       55     60        75        65        90        80
FPC 7                       55     60        75        65        90        80
FPC 8                       55     60        75        65        90        80
FPC 9                       55     60        75        65        90        80
FPC 10                      55     60        75        65        90        80
FPC 11                      55     60        75        65        90        80
```

show chassis temperature-thresholds (MX150)

```
user@host> show chassis temperature-thresholds
Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                          (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                     Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
FPC 0 Sensor 1               43    65       68       68       70       70       72
FPC 0 Sensor 2               43    65       68       68       70       70       72
FPC 0 Coretemp               78    94      100      100      105      105      110
```

show chassis temperature-thresholds (MX104 Router)

```
user@host> show chassis temperature-thresholds
                            Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                          (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                     Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Chassis default              48    54       65       55       75       65      100
Routing Engine 0             55    80       95       95      105      100      108
```

show chassis temperature-thresholds (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```
user@host> show chassis temperature-thresholds
Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                          (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                     Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Chassis default              48    54       65       55       75       65      100
Routing Engine 0             70    80       95       95      110      110      112
Routing Engine 1             70    80       95       95      110      110      112
FPC 0                        55    60       75       65       90       80       95
FPC 1                        55    60       75       65       90       80       95
FPC 2                        55    60       75       65       90       80       95
FPC 4                        55    60       75       65       90       80       95
FPC 5                        55    60       75       65       90       80       95
```

show chassis temperature-thresholds (MX480 Router with MPC4E)

```
user@ host> show chassis temperature-thresholds
                            Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                          (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                     Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Chassis default              48    54       65       55       75       65      100
Routing Engine 0             70    80       95       95      110      110      112
Routing Engine 1             70    80       95       95      110      110      112
FPC 2                        55    60       75       65       95       80      100
FPC 3                        55    60       75       65       95       80      100
FPC 4                        55    60       75       65       90       80       95
```

show chassis temperature-thresholds (MX2010 Router with MPC7E, MPC8E, and MPC9E)

```
user@ host> show chassis temperature-thresholds
                                        Fan speed      Yellow alarm      Red alarm      Fire
Shutdown
                                       (degrees C)      (degrees C)     (degrees C)
(degrees C)
Item                                  Normal  High   Normal  Bad fan   Normal  Bad fan     Normal

FPC 3 Intake                              53    59       72       67       80       75         85
FPC 3 Exhaust A                           77    85       98       93      103       98        108
FPC 3 Exhaust B                           54    62       80       75      103       98        108
FPC 3 EA0 Chip                            64    72       90       90      100      100        105
FPC 3 EA0_XR0 Chip                        79    87      102      102      106      106        108
FPC 3 EA0_XR1 Chip                        79    87      102      102      106      106        108
FPC 3 EA1 Chip                            64    72       90       90      100      100        105
FPC 3 EA1_XR0 Chip                        79    87      102      102      106      106        108
FPC 3 EA1_XR1 Chip                        79    87      102      102      106      106        108
FPC 3 PEX Chip                            74    82      100      100      105      105        110
FPC 3 EA2 Chip                            64    72       90       90      100      100        105
FPC 3 EA2_XR0 Chip                        79    87      102      102      106      106        108
FPC 3 EA2_XR1 Chip                        79    87      102      102      106      106        108
FPC 3 EA3 Chip                            64    72       90       90      100      100        105
FPC 3 EA3_XR0 Chip                        79    87      102      102      106      106        108
FPC 3 EA3_XR1 Chip                        79    87      102      102      106      106        108
FPC 3 EA0_HMC0 Logic die                  81    89      103      103      107      107        111
FPC 3 EA0_HMC0 DRAM botm                  76    84       98       98      102      102        106
```

FPC 3 EA0_HMC1 Logic die	81	89	103	103	107	107	111
FPC 3 EA0_HMC1 DRAM botm	76	84	98	98	102	102	106
FPC 3 EA0_HMC2 Logic die	81	89	103	103	107	107	111
FPC 3 EA0_HMC2 DRAM botm	76	84	98	98	102	102	106
FPC 3 EA1_HMC0 Logic die	81	89	103	103	107	107	111
FPC 3 EA1_HMC0 DRAM botm	76	84	98	98	102	102	106
FPC 3 EA1_HMC1 Logic die	81	89	103	103	107	107	111
FPC 3 EA1_HMC1 DRAM botm	76	84	98	98	102	102	106
FPC 3 EA1_HMC2 Logic die	81	89	103	103	107	107	111
FPC 3 EA1_HMC2 DRAM botm	76	84	98	98	102	102	106
FPC 3 EA2_HMC0 Logic die	81	89	103	103	107	107	111
FPC 3 EA2_HMC0 DRAM botm	76	84	98	98	102	102	106
FPC 3 EA2_HMC1 Logic die	81	89	103	103	107	107	111
FPC 3 EA2_HMC1 DRAM botm	76	84	98	98	102	102	106
FPC 3 EA2_HMC2 Logic die	81	89	103	103	107	107	111
FPC 3 EA2_HMC2 DRAM botm	76	84	98	98	102	102	106
FPC 3 EA3_HMC0 Logic die	81	89	103	103	107	107	111
FPC 3 EA3_HMC0 DRAM botm	76	84	98	98	102	102	106
FPC 3 EA3_HMC1 Logic die	81	89	103	103	107	107	111
FPC 3 EA3_HMC1 DRAM botm	76	84	98	98	102	102	106
FPC 3 EA3_HMC2 Logic die	81	89	103	103	107	107	111
FPC 3 EA3_HMC2 DRAM botm	76	84	98	98	102	102	106
FPC 4 Intake	46	55	65	60	81	76	90
FPC 4 Exhaust A	61	70	80	75	100	95	110
FPC 4 Exhaust B	61	70	80	75	95	90	105
FPC 4 EA0 Chip	86	95	105	100	117	112	123
FPC 4 EA0_XR0 Chip	86	95	105	100	110	105	116
FPC 4 EA0_XR1 Chip	86	95	105	100	115	110	121
FPC 4 EA1 Chip	86	95	105	100	117	112	123
FPC 4 EA1_XR0 Chip	86	95	105	100	110	105	116
FPC 4 EA1_XR1 Chip	86	95	105	100	115	110	121
FPC 4 PCIE_SW Chip	81	90	105	100	115	110	121
FPC 4 EA0_HMC0 DRAM botm	86	95	105	100	115	110	121
FPC 4 EA0_HMC1 DRAM botm	86	95	105	100	115	110	121
FPC 4 EA1_HMC0 DRAM botm	86	95	105	100	115	110	121
FPC 4 EA1_HMC1 DRAM botm	86	95	105	100	115	110	121
FPC 7 Intake	53	59	72	67	80	75	85
FPC 7 Exhaust A	77	85	98	93	103	98	108
FPC 7 Exhaust B	54	62	80	75	103	98	108
FPC 7 EA0 Chip	64	72	90	90	100	100	105
FPC 7 EA0_XR0 Chip	79	87	102	102	106	106	108
FPC 7 EA0_XR1 Chip	79	87	102	102	106	106	108
FPC 7 EA1 Chip	64	72	90	90	100	100	105

```
FPC 7 EA1_XR0 Chip                    79    87     102     102     106     106      108
FPC 7 EA1_XR1 Chip                    79    87     102     102     106     106      108
FPC 7 PEX Chip                        74    82     100     100     105     105      110
FPC 7 EA2 Chip                        64    72      90      90     100     100      105
FPC 7 EA2_XR0 Chip                    79    87     102     102     106     106      108
FPC 7 EA2_XR1 Chip                    79    87     102     102     106     106      108
FPC 7 EA3 Chip                        64    72      90      90     100     100      105
FPC 7 EA3_XR0 Chip                    79    87     102     102     106     106      108
FPC 7 EA3_XR1 Chip                    79    87     102     102     106     106      108
FPC 7 EA0_HMC0 Logic die              81    89     103     103     107     107      111
FPC 7 EA0_HMC0 DRAM botm              76    84      98      98     102     102      106
FPC 7 EA0_HMC1 Logic die              81    89     103     103     107     107      111
FPC 7 EA0_HMC1 DRAM botm              76    84      98      98     102     102      106
FPC 7 EA0_HMC2 Logic die              81    89     103     103     107     107      111
FPC 7 EA0_HMC2 DRAM botm              76    84      98      98     102     102      106
FPC 7 EA1_HMC0 Logic die              81    89     103     103     107     107      111
FPC 7 EA1_HMC0 DRAM botm              76    84      98      98     102     102      106
FPC 7 EA1_HMC1 Logic die              81    89     103     103     107     107      111
FPC 7 EA1_HMC1 DRAM botm              76    84      98      98     102     102      106
FPC 7 EA1_HMC2 Logic die              81    89     103     103     107     107      111
FPC 7 EA1_HMC2 DRAM botm              76    84      98      98     102     102      106
FPC 7 EA2_HMC0 Logic die              81    89     103     103     107     107      111
FPC 7 EA2_HMC0 DRAM botm              76    84      98      98     102     102      106
FPC 7 EA2_HMC1 Logic die              81    89     103     103     107     107      111
FPC 7 EA2_HMC1 DRAM botm              76    84      98      98     102     102      106
FPC 7 EA2_HMC2 Logic die              81    89     103     103     107     107      111
FPC 7 EA2_HMC2 DRAM botm              76    84      98      98     102     102      106
FPC 7 EA3_HMC0 Logic die              81    89     103     103     107     107      111
FPC 7 EA3_HMC0 DRAM botm              76    84      98      98     102     102      106
FPC 7 EA3_HMC1 Logic die              81    89     103     103     107     107      111
FPC 7 EA3_HMC1 DRAM botm              76    84      98      98     102     102      106
FPC 7 EA3_HMC2 Logic die              81    89     103     103     107     107      111
FPC 7 EA3_HMC2 DRAM botm              76    84      98      98     102     102      106
```

As per the above output, the MPC7E, MPC8E, and MPC9E are installed in the FPC slots 4, 7, and 3, respectively.

show chassis temperature-thresholds (MX2020 Router with MPC4E)

```
user@host> show chassis temperature-thresholds
                     Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                      (degrees C)      (degrees C)     (degrees C)      (degrees C)
```

Item	Normal	High	Normal	Bad fan	Normal	Bad fan	Normal
Routing Engine 0	70	80	95	95	110	110	112
Routing Engine 1	70	80	95	95	110	110	112
CB 0 IntakeA-Zone0	60	65	78	75	85	80	95
CB 0 IntakeB-Zone1	60	65	78	75	85	80	95
CB 0 IntakeC-Zone0	60	65	78	75	85	80	95
CB 0 ExhaustA-Zone0	60	65	78	75	85	80	95
CB 0 ExhaustB-Zone1	60	65	78	75	85	80	95
CB 0 TCBC-Zone0	60	65	78	75	85	80	95
CB 1 IntakeA-Zone0	60	65	78	75	85	80	95
CB 1 IntakeB-Zone1	60	65	78	75	85	80	95
CB 1 IntakeC-Zone0	60	65	78	75	85	80	95
CB 1 ExhaustA-Zone0	60	65	78	75	85	80	95
CB 1 ExhaustB-Zone1	60	65	78	75	85	80	95
CB 1 TCBC-Zone0	60	65	78	75	85	80	95
SPMB 0 Intake	56	62	75	63	83	76	95
SPMB 1 Intake	56	62	75	63	83	76	95
SFB 0 Intake-Zone0	56	62	70	70	85	85	89
SFB 0 Exhaust-Zone1	56	62	70	70	85	85	89
SFB 0 IntakeA-Zone0	56	62	70	70	85	85	89
SFB 0 IntakeB-Zone1	56	62	70	70	85	85	89
SFB 0 Exhaust-Zone0	56	62	70	70	85	85	89
SFB 0 SFB-XF2-Zone1	70	75	90	85	95	90	100
SFB 0 SFB-XF1-Zone0	70	75	90	85	95	90	100
SFB 0 SFB-XF0-Zone0	70	75	90	85	95	90	100
SFB 1 Intake-Zone0	56	62	70	70	85	85	89
SFB 1 Exhaust-Zone1	56	62	70	70	85	85	89
SFB 1 IntakeA-Zone0	56	62	70	70	85	85	89
SFB 1 IntakeB-Zone1	56	62	70	70	85	85	89
SFB 1 Exhaust-Zone0	56	62	70	70	85	85	89
SFB 1 SFB-XF2-Zone1	70	75	90	85	95	90	100
SFB 1 SFB-XF1-Zone0	70	75	90	85	95	90	100
SFB 1 SFB-XF0-Zone0	70	75	90	85	95	90	100
SFB 2 Intake-Zone0	56	62	70	70	85	85	89
SFB 2 Exhaust-Zone1	56	62	70	70	85	85	89
SFB 2 IntakeA-Zone0	56	62	70	70	85	85	89
SFB 2 IntakeB-Zone1	56	62	70	70	85	85	89
SFB 2 Exhaust-Zone0	56	62	70	70	85	85	89
SFB 2 SFB-XF2-Zone1	70	75	90	85	95	90	100
SFB 2 SFB-XF1-Zone0	70	75	90	85	95	90	100
SFB 2 SFB-XF0-Zone0	70	75	90	85	95	90	100
SFB 3 Intake-Zone0	56	62	70	70	85	85	89
SFB 3 Exhaust-Zone1	56	62	70	70	85	85	89

SFB 3 IntakeA-Zone0	56	62	70	70	85	85	89
SFB 3 IntakeB-Zone1	56	62	70	70	85	85	89
SFB 3 Exhaust-Zone0	56	62	70	70	85	85	89
SFB 3 SFB-XF2-Zone1	70	75	90	85	95	90	100
SFB 3 SFB-XF1-Zone0	70	75	90	85	95	90	100
SFB 3 SFB-XF0-Zone0	70	75	90	85	95	90	100
SFB 4 Intake-Zone0	56	62	70	70	85	85	89
SFB 4 Exhaust-Zone1	56	62	70	70	85	85	89
SFB 4 IntakeA-Zone0	56	62	70	70	85	85	89
SFB 4 IntakeB-Zone1	56	62	70	70	85	85	89
SFB 4 Exhaust-Zone0	56	62	70	70	85	85	89
SFB 4 SFB-XF2-Zone1	70	75	90	85	95	90	100
SFB 4 SFB-XF1-Zone0	70	75	90	85	95	90	100
SFB 4 SFB-XF0-Zone0	70	75	90	85	95	90	100
SFB 5 Intake-Zone0	56	62	70	70	85	85	89
SFB 5 Exhaust-Zone1	56	62	70	70	85	85	89
SFB 5 IntakeA-Zone0	56	62	70	70	85	85	89
SFB 5 IntakeB-Zone1	56	62	70	70	85	85	89
SFB 5 Exhaust-Zone0	56	62	70	70	85	85	89
SFB 5 SFB-XF2-Zone1	70	75	90	85	95	90	100
SFB 5 SFB-XF1-Zone0	70	75	90	85	95	90	100
SFB 5 SFB-XF0-Zone0	70	75	90	85	95	90	100
SFB 6 Intake-Zone0	56	62	70	70	85	85	89
SFB 6 Exhaust-Zone1	56	62	70	70	85	85	89
SFB 6 IntakeA-Zone0	56	62	70	70	85	85	89
SFB 6 IntakeB-Zone1	56	62	70	70	85	85	89
SFB 6 Exhaust-Zone0	56	62	70	70	85	85	89
SFB 6 SFB-XF2-Zone1	70	75	90	85	95	90	100
SFB 6 SFB-XF1-Zone0	70	75	90	85	95	90	100
SFB 6 SFB-XF0-Zone0	70	75	90	85	95	90	100
SFB 7 Intake-Zone0	56	62	70	70	85	85	89
SFB 7 Exhaust-Zone1	56	62	70	70	85	85	89
SFB 7 IntakeA-Zone0	56	62	70	70	85	85	89
SFB 7 IntakeB-Zone1	56	62	70	70	85	85	89
SFB 7 Exhaust-Zone0	56	62	70	70	85	85	89
SFB 7 SFB-XF2-Zone1	70	75	90	85	95	90	100
SFB 7 SFB-XF1-Zone0	70	75	90	85	95	90	100
SFB 7 SFB-XF0-Zone0	70	75	90	85	95	90	100
FPC 0	55	60	75	65	90	80	95
FPC 9	55	60	75	65	90	80	95
FPC 10	55	60	75	65	90	80	95
FPC 14	55	60	75	65	95	80	100
FPC 19	55	60	75	65	90	80	95

```
ADC 0 Intake            50    55       60       60       65       65       80
ADC 0 Exhaust           50    55       60       60       65       65       80
ADC 0 ADC-XF1           70    75       90       85       95       90      100
ADC 0 ADC-XF0           70    75       90       85       95       90      100
ADC 9 Intake            50    55       60       60       65       65       80
ADC 9 Exhaust           50    55       60       60       65       65       80
ADC 9 ADC-XF1           70    75       90       85       95       90      100
ADC 9 ADC-XF0           70    75       90       85       95       90      100
ADC 10 Intake           50    55       60       60       65       65       80
ADC 10 Exhaust          50    55       60       60       65       65       80
ADC 10 ADC-XF1          70    75       90       85       95       90      100
ADC 10 ADC-XF0          70    75       90       85       95       90      100
ADC 14 Intake           50    55       60       60       65       65       80
ADC 14 Exhaust          50    55       60       60       65       65       80
ADC 14 ADC-XF1          70    75       90       85       95       90      100
ADC 14 ADC-XF0          70    75       90       85       95       90      100
ADC 19 Intake           50    55       60       60       65       65       80
ADC 19 Exhaust          50    55       60       60       65       65       80
ADC 19 ADC-XF1          70    75       90       85       95       90      100
ADC 19 ADC-XF0          70    75       90       85       95       90      100
```

show chassis temperature-thresholds (MX2008 Routers)

```
user@host> show chassis temperature-thresholds
                           Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                          (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                     Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Routing Engine 0 CPU         58    63       78       75       93       90       98
Routing Engine 1 CPU         58    63       78       75       93       90       98
CB 0 Inlet1                  55    60       65       62       75       72       85
CB 0 Inlet2                  45    50       61       58       80       77       90
CB 0 Inlet3                  57    62       68       65       80       77       90
CB 0 Inlet4                  55    60       80       77       90       87       95
CB 0 Exhaust1                55    60       65       62       75       72       85
CB 0 Exhaust2                50    55       60       57       80       77       90
CB 0 Exhaust3                70    75       81       78       91       88       96
CB 0 Exhaust4                75    80       90       87      100       97      105
CB 1 Inlet1                  55    60       65       62       75       72       85
CB 1 Inlet2                  45    50       61       58       80       77       90
CB 1 Inlet3                  57    62       68       65       80       77       90
CB 1 Inlet4                  55    60       80       77       90       87       95
```

CB 1 Exhaust1	55	60	65	62	75	72	85
CB 1 Exhaust2	50	55	60	57	80	77	90
CB 1 Exhaust3	70	75	81	78	91	88	96
CB 1 Exhaust4	75	80	90	87	100	97	105
SFB 0 Inlet1	49	54	62	59	76	73	81
SFB 0 Inlet2	65	70	71	68	83	80	88
SFB 0 Exhaust1	45	50	61	58	75	72	80
SFB 0 Exhaust2	60	65	69	66	80	77	85
SFB 0 SFB2-PF-local	65	70	75	72	95	92	100
SFB 0 SFB2-PF-die	88	93	98	95	118	115	120
SFB 1 Inlet1	49	54	62	59	76	73	81
SFB 1 Inlet2	65	70	71	68	83	80	88
SFB 1 Exhaust1	45	50	61	58	75	72	80
SFB 1 Exhaust2	60	65	69	66	80	77	85
SFB 1 SFB2-PF-local	65	70	75	72	95	92	100
SFB 1 SFB2-PF-die	88	93	98	95	118	115	120
SFB 2 Inlet1	49	54	62	59	76	73	81
SFB 2 Inlet2	65	70	71	68	83	80	88
SFB 2 Exhaust1	45	50	61	58	75	72	80
SFB 2 Exhaust2	60	65	69	66	80	77	85
SFB 2 SFB2-PF-local	65	70	75	72	95	92	100
SFB 2 SFB2-PF-die	88	93	98	95	118	115	120
SFB 3 Inlet1	49	54	62	59	76	73	81
SFB 3 Inlet2	65	70	71	68	83	80	88
SFB 3 Exhaust1	45	50	61	58	75	72	80
SFB 3 Exhaust2	60	65	69	66	80	77	85
SFB 3 SFB2-PF-local	65	70	75	72	95	92	100
SFB 3 SFB2-PF-die	88	93	98	95	118	115	120
SFB 4 Inlet1	49	54	62	59	76	73	81
SFB 4 Inlet2	65	70	71	68	83	80	88
SFB 4 Exhaust1	45	50	61	58	75	72	80
SFB 4 Exhaust2	60	65	69	66	80	77	85
SFB 4 SFB2-PF-local	65	70	75	72	95	92	100
SFB 4 SFB2-PF-die	88	93	98	95	118	115	120
SFB 5 Inlet1	49	54	62	59	76	73	81
SFB 5 Inlet2	65	70	71	68	83	80	88
SFB 5 Exhaust1	45	50	61	58	75	72	80
SFB 5 Exhaust2	60	65	69	66	80	77	85
SFB 5 SFB2-PF-local	65	70	75	72	95	92	100
SFB 5 SFB2-PF-die	88	93	98	95	118	115	120
SFB 6 Inlet1	49	54	62	59	76	73	81
SFB 6 Inlet2	65	70	71	68	83	80	88
SFB 6 Exhaust1	45	50	61	58	75	72	80

SFB 6 Exhaust2	60	65	69	66	80	77	85
SFB 6 SFB2-PF-local	65	70	75	72	95	92	100
SFB 6 SFB2-PF-die	88	93	98	95	118	115	120
SFB 7 Inlet1	49	54	62	59	76	73	81
SFB 7 Inlet2	65	70	71	68	83	80	88
SFB 7 Exhaust1	45	50	61	58	75	72	80
SFB 7 Exhaust2	60	65	69	66	80	77	85
SFB 7 SFB2-PF-local	65	70	75	72	95	92	100
SFB 7 SFB2-PF-die	88	93	98	95	118	115	120
FPC 0	55	60	75	65	90	80	95
FPC 3	55	60	75	65	105	80	110
FPC 5	55	60	75	65	105	80	110
FPC 7	55	60	75	65	90	80	95
FPC 9 Intake	60	65	75	75	85	85	95
FPC 9 Exhaust A	60	65	75	75	85	85	95
FPC 9 Exhaust B	60	65	75	75	85	85	95
FPC 9 XL 0 Chip	70	75	85	85	102	102	110
FPC 9 XL 0 XR2 0 Chip	75	80	90	90	105	105	115
FPC 9 XL 0 XR2 1 Chip	75	80	90	90	105	105	115
FPC 9 XL 1 Chip	70	75	85	85	102	102	110
FPC 9 XL 1 XR2 0 Chip	75	80	90	90	105	105	115
FPC 9 XL 1 XR2 1 Chip	75	80	90	90	105	105	115
FPC 9 XM 0 Chip	70	75	85	85	100	100	110
FPC 9 XM 1 Chip	70	75	85	85	100	100	110
FPC 9 XM 2 Chip	70	75	85	85	100	100	110
FPC 9 XM 3 Chip	70	75	85	85	100	100	110
FPC 9 PCIe Switch Chip	80	85	95	95	105	105	120
ADC 0 Intake	50	55	65	65	75	75	80
ADC 0 Exhaust	50	55	65	65	75	75	80
ADC 0 ADC-XF1	70	75	90	85	95	90	100
ADC 0 ADC-XF0	70	75	90	85	95	90	100
ADC 3 Intake	50	55	65	65	75	75	80
ADC 3 Exhaust	50	55	65	65	75	75	80
ADC 3 ADC-XF1	70	75	90	85	95	90	100
ADC 3 ADC-XF0	70	75	90	85	95	90	100
ADC 5 Intake	50	55	65	65	75	75	80
ADC 5 Exhaust	50	55	65	65	75	75	80
ADC 5 ADC-XF1	70	75	90	85	95	90	100
ADC 5 ADC-XF0	70	75	90	85	95	90	100
ADC 7 Intake	50	55	65	65	75	75	80
ADC 7 Exhaust	50	55	65	65	75	75	80

```
ADC 7 ADC-XF1                70     75       90       85       95       90      100
ADC 7 ADC-XF0                70     75       90       85       95       90      100
```

show chassis temperature-thresholds (MX204 Router)

```
user@host> show chassis temperature-thresholds

                                        Fan speed      Yellow alarm      Red alarm      Fire
Shutdown
                                       (degrees C)      (degrees C)     (degrees C)
(degrees C)
Item                                  Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Routing Engine                            48    54       85       85      100      100        102
CB Top Right Inlet Sensor                 35    40       63       63       85       85         95
CB Top Left Inlet Sensor                  40    45       65       65       85       85         95
CB Top Right Exhaust Sensor               45    50       68       68       85       85         95
CB Top Left Exhaust Sensor                65    70       78       78       85       85         95
CB CPU Core-0 Temp                        65    70       80       80       90       90        100
CB CPU Core-1 Temp                        65    70       80       80       90       90        100
CB CPU Core-2 Temp                        65    70       80       80       90       90        100
CB CPU Core-3 Temp                        65    70       80       80       90       90        100
CB CPU Core-4 Temp                        65    70       80       80       90       90        100
CB CPU Core-5 Temp                        65    70       80       80       90       90        100
CB CPU Core-6 Temp                        65    70       80       80       90       90        100
CB CPU Core-7 Temp                        65    70       80       80       90       90        100
FPC EA0_HMC0 Logic die                    85    90       95       95      105      105        110
FPC EA0_HMC0 DRAM botm                    80    85       90       90      105      105        110
FPC EA0_HMC1 Logic die                    85    90       95       95      105      105        110
FPC EA0_HMC1 DRAM botm                    80    85       90       90      105      105        110
FPC EA0 Chip                              92    97      103      103      109      109        115
FPC EA0-XR0 Chip                          85    90       98       98      103      103        110
FPC EA0-XR1 Chip                          85    90       98       98      103      103        110
```

show chassis temperature-thresholds (MX10004 and MX10008 Router)

```
 user@host> show chassis temperature-thresholds
                                        Fan speed      Yellow alarm      Red alarm      Fire
Shutdown
                                       (degrees C)      (degrees C)     (degrees C)
(degrees C)
```

```
Item                                Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
SFB 0 Intake-Top Temp                   58    63      100      100      105      105        115
SFB 0 Intake-Middle-High Temp           63    68      100      100      105      105        115
SFB 0 Intake-Middle-Low Temp            66    71      100      100      105      105        115
SFB 0 Intake-Bottom Temp                65    70      100      100      105      105        115
SFB 0 Exhaust-Top Temp                  82    87      100      100      105      105        115
SFB 0 Exhaust-Middle Temp               76    81      100      100      105      105        115
SFB 0 Exhaust-Bottom Temp               77    82      100      100      105      105        115
SFB 0 PCIe Switch Temp                  70    75      110      110      115      115        125
SFB 0 ZF0 External Temp                 85    90       95       95      116      116        121
SFB 0 ZF1 External Temp                 84    89       95       95      116      116        121
SFB 0 ZF0 Internal Main                 85    90       95       95      116      116        121
SFB 0 ZF0 Internal Remote 0             87    92       95       95      116      116        121
SFB 0 ZF0 Internal Remote 1             89    94       95       95      116      116        121
SFB 0 ZF0 Internal Remote 2             86    91       95       95      116      116        121
SFB 0 ZF0 Internal Remote 3             88    93       95       95      116      116        121
SFB 0 ZF0 Internal Remote 4             88    93       95       95      116      116        121
SFB 0 ZF0 Internal Remote 5             87    92       95       95      116      116        121
SFB 0 ZF0 Internal Remote 6             86    91       95       95      116      116        121
SFB 0 ZF1 Internal Main                 84    89       95       95      116      116        121
SFB 0 ZF1 Internal Remote 0             85    90       95       95      116      116        121
SFB 0 ZF1 Internal Remote 1             88    93       95       95      116      116        121
SFB 0 ZF1 Internal Remote 2             86    91       95       95      116      116        121
SFB 0 ZF1 Internal Remote 3             88    93       95       95      116      116        121
SFB 0 ZF1 Internal Remote 4             87    92       95       95      116      116        121
SFB 0 ZF1 Internal Remote 5             85    90       95       95      116      116        121
SFB 0 ZF1 Internal Remote 6             87    92       95       95      116      116        121
```

show chassis temperature-thresholds (PTX10008 Routers)

```
user@host> show chassis temperature-thresholds
                           Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                          (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                     Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Routing Engine 0             48    54       85       85      100      100      102
Routing Engine 1             48    54       85       85      100      100      102
CB 0 Intake Temp Sensor      30    35       80       80       85       85       95
CB 0 Exhaust Temp Sensor     30    35       80       80       85       85       95
CB 0 CPU Die Temp Sensor     40    45       95       95      100      100      110
CB 1 Intake Temp Sensor      30    35       80       80       85       85       95
CB 1 Exhaust Temp Sensor     30    35       80       80       85       85       95
```

CB 1 CPU Die Temp Sensor	40	45	95	95	100	100	110
FPC 0 Intake-A Temp Sensor	30	35	80	80	85	85	95
FPC 0 Intake-B Temp Sensor	30	35	80	80	85	85	95
FPC 0 Exhaust-A Temp Sensor	30	35	80	80	85	85	95
FPC 0 Exhaust-B Temp Sensor	30	35	80	80	85	85	95
FPC 0 Exhaust-C Temp Sensor	30	35	80	80	85	85	95
FPC 0 PE0 Temp Sensor	40	45	100	100	105	105	115
FPC 0 PE1 Temp Sensor	40	45	100	100	105	105	115
FPC 0 PE2 Temp Sensor	40	45	100	100	105	105	115
FPC 0 LCPU Temp Sensor	40	45	95	95	100	100	110
FPC 5 Intake-A Temp Sensor	30	35	80	80	85	85	95
FPC 5 Intake-B Temp Sensor	30	35	80	80	85	85	95
FPC 5 Exhaust-A Temp Sensor	30	35	80	80	85	85	95
FPC 5 Exhaust-B Temp Sensor	30	35	80	80	85	85	95
FPC 5 Exhaust-C Temp Sensor	30	35	80	80	85	85	95
FPC 5 PE0 Temp Sensor	40	45	100	100	105	105	115
FPC 5 PE1 Temp Sensor	40	45	100	100	105	105	115
FPC 5 PE2 Temp Sensor	40	45	100	100	105	105	115
FPC 5 PE3 Temp Sensor	40	45	100	100	105	105	115
FPC 5 PE4 Temp Sensor	40	45	100	100	105	105	115
FPC 5 PE5 Temp Sensor	40	45	100	100	105	105	115
FPC 5 LCPU Temp Sensor	40	45	95	95	100	100	110
FPC 6 Intake-A Temp Sensor	30	35	80	80	85	85	95
FPC 6 Intake-B Temp Sensor	30	35	80	80	85	85	95
FPC 6 Exhaust-A Temp Sensor	30	35	80	80	85	85	95
FPC 6 Exhaust-B Temp Sensor	30	35	80	80	85	85	95
FPC 6 Exhaust-C Temp Sensor	30	35	80	80	85	85	95
FPC 6 PE0 Temp Sensor	40	45	100	100	105	105	115
FPC 6 PE1 Temp Sensor	40	45	100	100	105	105	115
FPC 6 PE2 Temp Sensor	40	45	100	100	105	105	115
FPC 6 PE3 Temp Sensor	40	45	100	100	105	105	115
FPC 6 PE4 Temp Sensor	40	45	100	100	105	105	115
FPC 6 PE5 Temp Sensor	40	45	100	100	105	105	115
FPC 6 LCPU Temp Sensor	40	45	95	95	100	100	110
SIB 0 Intake-A Temp Sensor	40	45	90	90	95	95	105
SIB 0 Intake-B Temp Sensor	40	45	90	90	95	95	105
SIB 0 Exhaust-A Temp Sensor	40	45	90	90	95	95	105
SIB 0 Exhaust-B Temp Sensor	40	45	90	90	95	95	105
SIB 0 PF0 Temp Sensor	50	55	100	100	105	105	115
SIB 0 PF1 Temp Sensor	50	55	100	100	105	105	115
SIB 1 Intake-A Temp Sensor	40	45	90	90	95	95	105
SIB 1 Intake-B Temp Sensor	40	45	90	90	95	95	105
SIB 1 Exhaust-A Temp Sensor	40	45	90	90	95	95	105

```
SIB 1 Exhaust-B Temp Sensor   40   45       90       90       95       95      105
SIB 1 PF0 Temp Sensor        50    55      100      100      105      105      115
SIB 1 PF1 Temp Sensor        50    55      100      100      105      105      115
```

show chassis temperature-thresholds (T4000 Core Routers)

```
user@host> show chassis temperature-thresholds

                     Fan speed      Yellow alarm       Red alarm   Fire Shutdown
                    (degrees C)      (degrees C)      (degrees C)    (degrees C)
Item                Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Chassis default        48    54       65       55       75       65      100
Routing Engine 0       55    65       85       85      100      100      102
Routing Engine 1       55    65       85       85      100      100      102
FPC 0                  63    68       75       70       90       83       95
FPC 3                  63    68       75       70       90       83       95
FPC 5                  56    62       75       63       83       76       95
FPC 6                  63    68       75       70       90       83       95
SIB 0                  64    70       76       72       87       84       95
SIB 1                  64    70       76       72       87       84       95
SIB 2                  64    70       76       72       87       84       95
SIB 3                  64    70       76       72       87       84       95
SIB 4                  64    70       76       72       87       84       95
```

show chassis temperature-thresholds (TX Matrix Plus Router)

```
user@host> show chassis temperature-thresholds
sfc0-re0:
--------------------------------------------------------------------------
                           Fan speed        Yellow alarm         Red alarm
                          (degrees C)        (degrees C)        (degrees C)
Item                     Normal   High    Normal   Bad fan    Normal   Bad fan
Chassis default              48     54        65        55        75        65
Routing Engine 0             55     65        85        85       100       100
Routing Engine 1             55     65        85        85       100       100
SIB F13 0                    64     70        76        72        90        84
SIB F13 3                    64     70        76        72        90        84
SIB F13 6                    64     70        76        72        90        84
SIB F13 8                    64     70        76        72        90        84
SIB F13 11                   64     70        76        72        90        84
```

```
SIB F13 12                    64     70        76        72        90        84
SIB F2S 16                    64     70        76        72        90        84
SIB F2S 17                    64     70        76        72        90        84
SIB F2S 18                    64     70        76        72        90        84
SIB F2S 19                    64     70        76        72        90        84
SIB F2S 20                    64     70        76        72        90        84
SIB F2S 21                    64     70        76        72        90        84
SIB F2S 22                    64     70        76        72        90        84
SIB F2S 23                    64     70        76        72        90        84
SIB F2S 24                    64     70        76        72        90        84
SIB F2S 25                    64     70        76        72        90        84
SIB F2S 26                    64     70        76        72        90        84
SIB F2S 27                    64     70        76        72        90        84
SIB F2S 28                    64     70        76        72        90        84
SIB F2S 29                    64     70        76        72        90        84
SIB F2S 30                    64     70        76        72        90        84
SIB F2S 31                    64     70        76        72        90        84
SIB F2S 32                    64     70        76        72        90        84
SIB F2S 33                    64     70        76        72        90        84
SIB F2S 34                    64     70        76        72        90        84
SIB F2S 35                    64     70        76        72        90        84

lcc0-re0:
--------------------------------------------------------------------------
                           Fan speed        Yellow alarm         Red alarm
                          (degrees C)        (degrees C)        (degrees C)
Item                     Normal   High    Normal   Bad fan    Normal   Bad fan
Chassis default              48     54        65        55        75        65
Routing Engine 0             55     65        85        85       100       100
Routing Engine 1             55     65        85        85       100       100
FPC 1                        56     62        75        63        83        76
FPC 3                        56     62        75        63        83        76
FPC 4                        56     62        75        63        83        76
FPC 6                        56     62        75        63        83        76
FPC 7                        56     62        75        63        83        76
SIB 0                        48     54        65        60        80        75
SIB 1                        48     54        65        60        80        75
SIB 2                        48     54        65        60        80        75
SIB 3                        48     54        65        60        80        75
SIB 4                        48     54        65        60        80        75

lcc1-re0:
--------------------------------------------------------------------------
```

```
                           Fan speed        Yellow alarm         Red alarm
                          (degrees C)        (degrees C)        (degrees C)
Item                     Normal   High    Normal   Bad fan    Normal   Bad fan
Chassis default              48     54        65        55        75        65
Routing Engine 0             55     65        85        85       100       100
Routing Engine 1             55     65        85        85       100       100
FPC 1                        56     62        75        63        83        76
FPC 3                        56     62        75        63        83        76
FPC 4                        56     62        75        63        83        76
FPC 6                        56     62        75        63        83        76
...
```

show chassis temperature-thresholds lcc (TX Matrix Plus Router)

```
user@host> show chassis temperature-thresholds lcc 1
lcc1-re0:
--------------------------------------------------------------------------
                           Fan speed        Yellow alarm         Red alarm
                          (degrees C)        (degrees C)        (degrees C)
Item                     Normal   High    Normal   Bad fan    Normal   Bad fan
Chassis default              48     54        65        55        75        65
Routing Engine 0             55     65        85        85       100       100
Routing Engine 1             55     65        85        85       100       100
FPC 1                        56     62        75        63        83        76
FPC 3                        56     62        75        63        83        76
FPC 4                        56     62        75        63        83        76
FPC 6                        56     62        75        63        83        76
SIB 0                        48     54        65        60        80        75
SIB 1                        48     54        65        60        80        75
SIB 2                        48     54        65        60        80        75
SIB 3                        48     54        65        60        80        75
SIB 4                        48     54        65        60        80        75
```

show chassis temperature-thresholds sfc (TX Matrix Plus Router)

```
user@host> show chassis temperature-thresholds sfc 0
sfc0-re0:
--------------------------------------------------------------------------
                           Fan speed        Yellow alarm         Red alarm
                          (degrees C)        (degrees C)        (degrees C)
```

```
Item                      Normal   High    Normal   Bad fan    Normal   Bad fan
Chassis default               48     54        65        55        75        65
Routing Engine 0              55     65        85        85       100       100
Routing Engine 1              55     65        85        85       100       100
SIB F13 0                     64     70        76        72        90        84
SIB F13 3                     64     70        76        72        90        84
SIB F13 6                     64     70        76        72        90        84
SIB F13 8                     64     70        76        72        90        84
SIB F13 11                    64     70        76        72        90        84
SIB F13 12                    64     70        76        72        90        84
SIB F2S 16                    64     70        76        72        90        84
SIB F2S 17                    64     70        76        72        90        84
SIB F2S 18                    64     70        76        72        90        84
SIB F2S 19                    64     70        76        72        90        84
SIB F2S 20                    64     70        76        72        90        84
SIB F2S 21                    64     70        76        72        90        84
SIB F2S 22                    64     70        76        72        90        84
SIB F2S 23                    64     70        76        72        90        84
SIB F2S 24                    64     70        76        72        90        84
SIB F2S 25                    64     70        76        72        90        84
SIB F2S 26                    64     70        76        72        90        84
SIB F2S 27                    64     70        76        72        90        84
SIB F2S 28                    64     70        76        72        90        84
SIB F2S 29                    64     70        76        72        90        84
SIB F2S 30                    64     70        76        72        90        84
SIB F2S 31                    64     70        76        72        90        84
SIB F2S 32                    64     70        76        72        90        84
SIB F2S 33                    64     70        76        72        90        84
SIB F2S 34                    64     70        76        72        90        84
SIB F2S 35                    64     70        76        72        90        84
```

show chassis temperature-thresholds (TX Matrix Plus routers with 3D SIBs)

```
user@host> show chassis temperature-thresholds
sfc0-re0:
--------------------------------------------------------------------------
                           Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                          (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                      Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Chassis default               48    54       65       55       75       65       100
Routing Engine 0              70    75       90       87      102       97       115
```

```
Routing Engine 1            70    75       90       87      102       97      115
SIB F13 0 Board             60    65       78       75       85       80       95
SIB F13 0 XF Junction       70    75       82       74      105      100      107
SIB F13 4 Board             60    65       78       75       85       80       95
SIB F13 4 XF Junction       70    75       82       74      105      100      107
SIB F13 6 Board             60    65       78       75       85       80       95
SIB F13 6 XF Junction       70    75       82       74      105      100      107
SIB F2S 16 Board            60    65       78       75       85       80       95
SIB F2S 16 XF Junction      70    75       82       74      105      100      107
SIB F2S 17 Board            60    65       78       75       85       80       95
SIB F2S 17 XF Junction      70    75       82       74      105      100      107
SIB F2S 18 Board            60    65       78       75       85       80       95
SIB F2S 18 XF Junction      70    75       82       74      105      100      107
SIB F2S 19 Board            60    65       78       75       85       80       95
SIB F2S 19 XF Junction      70    75       82       74      105      100      107
SIB F2S 24 Board            60    65       78       75       85       80       95
SIB F2S 24 XF Junction      70    75       82       74      105      100      107
SIB F2S 25 Board            60    65       78       75       85       80       95
SIB F2S 25 XF Junction      70    75       82       74      105      100      107
SIB F2S 26 Board            60    65       78       75       85       80       95
SIB F2S 26 XF Junction      70    75       82       74      105      100      107
SIB F2S 27 Board            60    65       78       75       85       80       95
SIB F2S 27 XF Junction      70    75       82       74      105      100      107

lcc0-re0:
--------------------------------------------------------------------------
                          Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                         (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                    Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Chassis default             48    54       65       55       75       65      100
Routing Engine 0            55    65       85       85      100      100      102
FPC 0                       63    68       75       70       90       83       95
FPC 1                       56    62       75       63       83       76       95
FPC 7                       56    62       75       63       83       76       95
SIB 0                       64    70       76       72       87       84       95
SIB 0 ASIC Junction         63    68       75       70      105      100      107
SIB 2                       64    70       76       72       87       84       95
SIB 2 ASIC Junction         63    68       75       70      105      100      107
SIB 3                       64    70       76       72       87       84       95
SIB 3 ASIC Junction         63    68       75       70      105      100      107
```

show chassis temperature-thresholds (QFX3500 Switch and QFX3600)

```
user@switch> show chassis temperature-thresholds
                           Fan speed      Yellow alarm      Red alarm
                          (degrees C)      (degrees C)     (degrees C)
Item                     Normal  High   Normal  Bad fan   Normal  Bad fan
rmal
FPC Sensor TopLeft I         48    56       53       43       56       46
FPC Sensor TopRight I        46    54       51       41       54       44
FPC Sensor TopLeft E         58    65       62       52       65       55
FPC Sensor TopRight E        56    64       61       51       64       54
FPC Sensor TopMiddle I       58    64       61       51       64       54
FPC Sensor TopMiddle E       67    74       71       61       74       64
FPC Sensor Bottom I          59    67       64       54       67       57
FPC Sensor Bottom E          66    73       70       60       73       63
FPC Sensor Die Temp          69    75       72       62       75       65
FPC Sensor Mgmnt Brd I       46    54       51       41       54       44
FPC Sensor Switch I          56    63       60       50       63       53
```

show chassis temperature-thresholds interconnect-device (QFabric System)

```
user@switch> show chassis temperature-thresholds interconnect-device
interconnect1
temperature-thresholds interconnect-device interconnect1
                           Fan speed        Yellow alarm         Red alarm
Item                     Normal   High    Normal   Bad fan    Normal   Bad fan
Chassis default              48     54        65        55        75        65
```

show chassis temperature-thresholds (QFX5700)

```
user@switch> show chassis temperature-thresholds
                           Fan speed        Yellow alarm         Red alarm
Item                     Normal   High    Normal   Bad fan    Normal   Bad fan

Routing Engine 0 CPU Temperature          48    54       85       85      100      100        102
Routing Engine 1 CPU Temperature          48    54       85       85      100      100        102
CB 0 Intake 1 Temp Sensor                 60    65       75       75       85       85         95
CB 0 Intake 2 Temp Sensor                 60    65       75       75       85       85         95
CB 0 Middle Temp Sensor                   83    90       98       98      105      105        110
CB 1 Intake 1 Temp Sensor                 60    65       75       75       85       85         95
```

```
CB 1 Intake 2 Temp Sensor                60    65       75       75       85       85         95
CB 1 Exhaust 1 Temp Sensor               83    90       98       98      105      105        110
CB 1 Exhaust 2 Temp Sensor               83    90       98       98      105      105        110
CB 1 Middle Temp Sensor                  83    90       98       98      105      105        110
```

show chassis temperature-thresholds (PTX5000 Packet Transport Router)

```
user@switch> show chassis temperature-thresholds
user@switch> show chassis temperature-thresholds

                          Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                         (degrees C)     (degrees C)     (degrees C)      (degrees C)
Item                    Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Routing Engine 0            80    90       95       85      105       95      115
CB 0 Exhaust A              60    65       78       75       85       80       95
CB 0 Exhaust B              60    65       78       75       85       80       95
CB 1 Exhaust A              60    65       78       75       85       80       95
CB 1 Exhaust B              60    65       78       75       85       80       95
FPC 3 Exhaust A             80    90       95       85      105       95      115
FPC 3 Exhaust B             80    90       95       85      105       95      115
FPC 3 TL5                   80    90       95       85      105       95      115
FPC 3 TQ5                   80    90       95       85      105       95      115
FPC 3 TL6                   80    90       95       85      105       95      115
FPC 3 TQ6                   80    90       95       85      105       95      115
FPC 3 TL1                   80    90       95       85      105       95      115
FPC 3 TQ1                   80    90       95       85      105       95      115
FPC 3 TL2                   80    90       95       85      105       95      115
FPC 3 TQ2                   80    90       95       85      105       95      115
FPC 3 TL4                   80    90       95       85      105       95      115
FPC 3 TQ4                   80    90       95       85      105       95      115
FPC 3 TL7                   80    90       95       85      105       95      115
FPC 3 TQ7                   80    90       95       85      105       95      115
FPC 3 TL0                   80    90       95       85      105       95      115
FPC 3 TQ0                   80    90       95       85      105       95      115
FPC 3 TL3                   80    90       95       85      105       95      115
FPC 3 TQ3                   80    90       95       85      105       95      115
SIB 0 Exhaust               60    65       78       75       85       80       95
SIB 0 Junction              75    80       90       85      105       95      115
SIB 1 Exhaust               60    65       78       75       85       80       95
SIB 1 Junction              75    80       90       85      105       95      115
SIB 2 Exhaust               60    65       78       75       85       80       95
```

```
SIB 2 Junction              75    80       90       85      105       95      115
SIB 3 Exhaust               60    65       78       75       85       80       95
SIB 3 Junction              75    80       90       85      105       95      115
SIB 4 Exhaust               60    65       78       75       85       80       95
SIB 4 Junction              75    80       90       85      105       95      115
SIB 5 Exhaust               60    65       78       75       85       80       95
SIB 5 Junction              75    80       90       85      105       95      115
SIB 6 Exhaust               60    65       78       75       85       80       95
SIB 6 Junction              75    80       90       85      105       95      115
SIB 7 Exhaust               60    65       78       75       85       80       95
SIB 7 Junction              75    80       90       85      105       95      115
SIB 8 Exhaust               60    65       78       75       85       80       95
SIB 8 Junction              75    80       90       85      105       95      115
```

show chassis temperature-thresholds (PTX1000 Packet Transport Router)

```
user@host> show chassis temperature-thresholds

                           Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                          (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                     Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
FPC 0 Intake Temp Sensor     30    65       65       65       70       70       75
FPC 0 Exhaust Temp Sensor    30    65       65       65       70       70       75
FPC 0 Mezz Temp Sensor 0     30    65       65       65       70       70       75
FPC 0 Mezz Temp Sensor 1     30    65       65       65       70       70       75
FPC 0 PE2 Temp Sensor        50    90       90       90      100      100      103
FPC 0 PE1 Temp Sensor        50    90       90       90      100      100      103
FPC 0 PF0 Temp Sensor        50    90       90       90      100      100      103
FPC 0 PE0 Temp Sensor        50    90       90       90      100      100      103
FPC 0 PE5 Temp Sensor        50    90       90       90      100      100      103
FPC 0 PE4 Temp Sensor        50    90       90       90      100      100      103
FPC 0 PF1 Temp Sensor        50    90       90       90      100      100      103
FPC 0 PE3 Temp Sensor        50    90       90       90      100      100      103
FPC 0 CPU Die Temp Sensor    50    90       90       90      100      100      103
FPC 0 OCXO Temp Sensor       50    90       90       90      100      100      103
```

show chassis temperature-thresholds (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis temperature-thresholds
  Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                         (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                    Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Chassis default             48    54       65       55       75       65      100
Routing Engine 0            70    80       95       95      110      110      112
Routing Engine 1            70    80       95       95      110      110      112
FPC 0                       55    60       75       65       90       80       95
FPC 1                       55    60       75       65       90       80       95
FPC 2                       55    60       75       65       90       80       95
FPC 4                       55    60       75       65       90       80       95
FPC 5                       55    60       75       65       90       80       95
```

show chassis temperature-thresholds (EX4300)

```
user@switch> show chassis temperature-thresholds

                                      Fan speed      Yellow alarm      Red alarm      Fire
Shutdown
                                      (degrees C)      (degrees C)     (degrees C)
(degrees C)
Item                                 Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
FPC 0 CPU                                35    85       77       77       85       85         95
FPC 0 NW-PFE                             35    60       59       59       67       67         70
FPC 0 SE-PFE                             35    71       70       70       78       78         81
FPC 0 PHY-2/3                            35    65       64       64       72       72         75
FPC 0 MGMT PHY                           35    73       72       72       80       80         84
FPC 0 PHY-4/5                            35    70       69       69       77       77         80
FPC 1 CPU                                35    85       77       77       85       85         95
FPC 1 NW-PFE                             35    60       59       59       67       67         70
FPC 1 SE-PFE                             35    71       70       70       78       78         81
FPC 1 PHY-2/3                            35    65       64       64       72       72         75
FPC 1 MGMT PHY                           35    73       72       72       80       80         84
FPC 1 PHY-4/5                            35    70       69       69       77       77         80
FPC 2 CPU                                35    85       77       77       85       85         95
FPC 2 NW-PFE                             35    60       59       59       67       67         70
FPC 2 SE-PFE                             35    71       70       70       78       78         81
```

```
FPC 2 PHY-2/3                          35    65       64       64       72       72         75
FPC 2 MGMT PHY                         35    73       72       72       80       80         84
FPC 2 PHY-4/5                          35    70       69       69       77       77         80
```

show chassis temperature-thresholds (EX9251 Switches)

```
user@switch> show chassis temperature-thresholds

                          Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                          (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                     Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Routing Engine                         48    54       85       85      100      100        102
CB Top Right Inlet Sensor              35    40       63       63       85       85         95
CB Top Left Inlet Sensor               40    45       65       65       85       85         95
CB Top Right Exhaust Sensor            45    50       68       68       85       85         95
CB Top Left Exhaust Sensor             65    70       78       78       85       85         95
CB CPU Core-0 Temp                     65    70       80       80       90       90        100
CB CPU Core-1 Temp                     65    70       80       80       90       90        100
CB CPU Core-2 Temp                     65    70       80       80       90       90        100
CB CPU Core-3 Temp                     65    70       80       80       90       90        100
CB CPU Core-4 Temp                     65    70       80       80       90       90        100
CB CPU Core-5 Temp                     65    70       80       80       90       90        100
CB CPU Core-6 Temp                     65    70       80       80       90       90        100
CB CPU Core-7 Temp                     65    70       80       80       90       90        100
FPC EA0_HMC0 Logic die                 85    90       95       95      105      105        110
FPC EA0_HMC0 DRAM botm                 80    85       90       90      105      105        110
FPC EA0_HMC1 Logic die                 85    90       95       95      105      105        110
FPC EA0_HMC1 DRAM botm                 80    85       90       90      105      105        110
FPC EA0 Chip                           92    97      103      103      109      109        115
FPC EA0-XR0 Chip                       85    90       98       98      103      103        110
FPC EA0-XR1 Chip                       85    90       98       98      103      103        110
```

show chassis temperature-thresholds (EX9253 Switches)

```
user@switch> show chassis temperature-thresholds

                          Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                          (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                     Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Routing Engine 0                       48    54       85       85      100      100        102
```

CB 0 Exhaust Temp Sensor	60	65	75	75	85	85	95
CB 0 Inlet Temp Sensor	60	65	75	75	85	85	95
CB 0 CPU DIE Temp Sensor	83	90	98	98	105	105	110
CB 1 Exhaust Temp Sensor	60	65	75	75	85	85	95
CB 1 Inlet Temp Sensor	60	65	75	75	85	85	95
CB 1 CPU DIE Temp Sensor	83	90	98	98	105	105	110
FPC 0 Intake Temp Sensor	40	45	75	70	85	80	95
FPC 0 Exhaust-A Temp Sensor	55	60	85	80	90	90	100
FPC 0 Exhaust-B Temp Sensor	55	60	85	80	90	90	100
FPC 0 EA0 Chip	87	92	97	97	105	105	110
FPC 0 EA0-XR0 Chip	88	93	98	98	120	120	125
FPC 0 EA0-XR1 Chip	88	93	98	98	120	120	125
FPC 0 EA1 Chip	87	92	97	97	105	105	110
FPC 0 EA1-XR0 Chip	88	93	98	98	120	120	125
FPC 0 EA1-XR1 Chip	88	93	98	98	120	120	125
FPC 0 EA2 Chip	87	92	97	97	105	105	110
FPC 0 EA2-XR0 Chip	88	93	98	98	120	120	125
FPC 0 EA2-XR1 Chip	88	93	98	98	120	120	125
FPC 0 PF Chip	89	94	104	104	120	120	120
FPC 0 EA0_HMC0 Logic die	88	93	103	103	120	120	125
FPC 0 EA0_HMC0 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA0_HMC1 Logic die	88	93	103	103	120	120	125
FPC 0 EA0_HMC1 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA0_HMC2 Logic die	88	93	103	103	120	120	125
FPC 0 EA0_HMC2 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA1_HMC0 Logic die	88	93	103	103	120	120	125
FPC 0 EA1_HMC0 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA1_HMC1 Logic die	88	93	103	103	120	120	125
FPC 0 EA1_HMC1 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA1_HMC2 Logic die	88	93	103	103	120	120	125
FPC 0 EA1_HMC2 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA2_HMC0 Logic die	88	93	103	103	120	120	125
FPC 0 EA2_HMC0 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA2_HMC1 Logic die	88	93	103	103	120	120	125
FPC 0 EA2_HMC1 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA2_HMC2 Logic die	88	93	103	103	120	120	125
FPC 0 EA2_HMC2 DRAM botm	83	88	98	98	120	120	125
FPC 1 Intake Temp Sensor	40	45	75	70	85	80	95
FPC 1 Exhaust-A Temp Sensor	55	60	85	80	90	90	100
FPC 1 Exhaust-B Temp Sensor	55	60	85	80	90	90	100
FPC 1 EA0 Chip	87	92	97	97	105	105	110
FPC 1 EA0-XR0 Chip	88	93	98	98	120	120	125
FPC 1 EA0-XR1 Chip	88	93	98	98	120	120	125

```
FPC 1 EA1 Chip                          87    92       97       97      105      105        110
FPC 1 EA1-XR0 Chip                      88    93       98       98      120      120        125
FPC 1 EA1-XR1 Chip                      88    93       98       98      120      120        125
FPC 1 EA2 Chip                          87    92       97       97      105      105        110
FPC 1 EA2-XR0 Chip                      88    93       98       98      120      120        125
FPC 1 EA2-XR1 Chip                      88    93       98       98      120      120        125
FPC 1 PF Chip                           89    94      104      104      120      120        120
FPC 1 EA0_HMC0 Logic die                88    93      103      103      120      120        125
FPC 1 EA0_HMC0 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA0_HMC1 Logic die                88    93      103      103      120      120        125
FPC 1 EA0_HMC1 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA0_HMC2 Logic die                88    93      103      103      120      120        125
FPC 1 EA0_HMC2 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA1_HMC0 Logic die                88    93      103      103      120      120        125
FPC 1 EA1_HMC0 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA1_HMC1 Logic die                88    93      103      103      120      120        125
FPC 1 EA1_HMC1 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA1_HMC2 Logic die                88    93      103      103      120      120        125
FPC 1 EA1_HMC2 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA2_HMC0 Logic die                88    93      103      103      120      120        125
FPC 1 EA2_HMC0 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA2_HMC1 Logic die                88    93      103      103      120      120        125
FPC 1 EA2_HMC1 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA2_HMC2 Logic die                88    93      103      103      120      120        125
FPC 1 EA2_HMC2 DRAM botm                83    88       98       98      120      120        125
```

show chassis temperature-thresholds (ACX Series)

```
user@host> show chassis temperature-thresholds

                          Fan speed      Yellow alarm      Red alarm      Fire Shutdown
                          (degrees C)      (degrees C)     (degrees C)      (degrees C)
Item                     Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Routing Engine 0                         48    54       85       85      100      100        102
CB 0 Exhaust Temp Sensor                 60    65       75       75       85       85         95
CB 0 Inlet Temp Sensor                   60    65       75       75       85       85         95
CB 0 CPU DIE Temp Sensor                 83    90       98       98      105      105        110
CB 1 Exhaust Temp Sensor                 60    65       75       75       85       85         95
CB 1 Inlet Temp Sensor                   60    65       75       75       85       85         95
CB 1 CPU DIE Temp Sensor                 83    90       98       98      105      105        110
FPC 0 Intake Temp Sensor                 40    45       75       70       85       80         95
```

FPC 0 Exhaust-A Temp Sensor	55	60	85	80	90	90	100
FPC 0 Exhaust-B Temp Sensor	55	60	85	80	90	90	100
FPC 0 EA0 Chip	87	92	97	97	105	105	110
FPC 0 EA0-XR0 Chip	88	93	98	98	120	120	125
FPC 0 EA0-XR1 Chip	88	93	98	98	120	120	125
FPC 0 EA1 Chip	87	92	97	97	105	105	110
FPC 0 EA1-XR0 Chip	88	93	98	98	120	120	125
FPC 0 EA1-XR1 Chip	88	93	98	98	120	120	125
FPC 0 EA2 Chip	87	92	97	97	105	105	110
FPC 0 EA2-XR0 Chip	88	93	98	98	120	120	125
FPC 0 EA2-XR1 Chip	88	93	98	98	120	120	125
FPC 0 PF Chip	89	94	104	104	120	120	120
FPC 0 EA0_HMC0 Logic die	88	93	103	103	120	120	125
FPC 0 EA0_HMC0 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA0_HMC1 Logic die	88	93	103	103	120	120	125
FPC 0 EA0_HMC1 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA0_HMC2 Logic die	88	93	103	103	120	120	125
FPC 0 EA0_HMC2 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA1_HMC0 Logic die	88	93	103	103	120	120	125
FPC 0 EA1_HMC0 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA1_HMC1 Logic die	88	93	103	103	120	120	125
FPC 0 EA1_HMC1 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA1_HMC2 Logic die	88	93	103	103	120	120	125
FPC 0 EA1_HMC2 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA2_HMC0 Logic die	88	93	103	103	120	120	125
FPC 0 EA2_HMC0 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA2_HMC1 Logic die	88	93	103	103	120	120	125
FPC 0 EA2_HMC1 DRAM botm	83	88	98	98	120	120	125
FPC 0 EA2_HMC2 Logic die	88	93	103	103	120	120	125
FPC 0 EA2_HMC2 DRAM botm	83	88	98	98	120	120	125
FPC 1 Intake Temp Sensor	40	45	75	70	85	80	95
FPC 1 Exhaust-A Temp Sensor	55	60	85	80	90	90	100
FPC 1 Exhaust-B Temp Sensor	55	60	85	80	90	90	100
FPC 1 EA0 Chip	87	92	97	97	105	105	110
FPC 1 EA0-XR0 Chip	88	93	98	98	120	120	125
FPC 1 EA0-XR1 Chip	88	93	98	98	120	120	125
FPC 1 EA1 Chip	87	92	97	97	105	105	110
FPC 1 EA1-XR0 Chip	88	93	98	98	120	120	125
FPC 1 EA1-XR1 Chip	88	93	98	98	120	120	125
FPC 1 EA2 Chip	87	92	97	97	105	105	110
FPC 1 EA2-XR0 Chip	88	93	98	98	120	120	125
FPC 1 EA2-XR1 Chip	88	93	98	98	120	120	125
FPC 1 PF Chip	89	94	104	104	120	120	120

```
FPC 1 EA0_HMC0 Logic die                88    93      103      103      120      120        125
FPC 1 EA0_HMC0 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA0_HMC1 Logic die                88    93      103      103      120      120        125
FPC 1 EA0_HMC1 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA0_HMC2 Logic die                88    93      103      103      120      120        125
FPC 1 EA0_HMC2 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA1_HMC0 Logic die                88    93      103      103      120      120        125
FPC 1 EA1_HMC0 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA1_HMC1 Logic die                88    93      103      103      120      120        125
FPC 1 EA1_HMC1 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA1_HMC2 Logic die                88    93      103      103      120      120        125
FPC 1 EA1_HMC2 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA2_HMC0 Logic die                88    93      103      103      120      120        125
FPC 1 EA2_HMC0 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA2_HMC1 Logic die                88    93      103      103      120      120        125
FPC 1 EA2_HMC1 DRAM botm                83    88       98       98      120      120        125
FPC 1 EA2_HMC2 Logic die                88    93      103      103      120      120        125
FPC 1 EA2_HMC2 DRAM botm                83    88       98       98      120      120        125
```

show chassis temperature-thresholds (MX304 with Optics)

When optics are connected to the device, the transceiver name displays instead of the interface name.

```
user@host> show chassis temperature-thresholds | match xcvr
FPC 0 xcvr-0/0/0                        57    62       72       72       75       75         79
FPC 0 xcvr-0/0/1                        52    57       67       67       70       70         74
FPC 0 xcvr-0/0/2                        52    57       67       67       70       70         74
FPC 0 xcvr-0/0/3                        55    60       70       70       73       73         76
FPC 0 xcvr-0/0/4                        57    62       72       72       75       75         79
FPC 0 xcvr-0/0/5                        55    60       70       70       73       73         76
FPC 0 xcvr-0/0/6                        55    60       70       70       73       73         76
FPC 0 xcvr-0/0/7                        55    60       70       70       73       73         76
FPC 0 xcvr-0/0/8                        55    60       70       70       73       73         76
FPC 0 xcvr-0/0/9                        55    60       70       70       73       73         76
FPC 0 xcvr-0/0/10                       52    57       67       67       70       70         74
FPC 0 xcvr-0/0/11                       57    62       72       72       75       75         79
FPC 0 xcvr-0/0/12                       52    57       67       67       70       70         74
FPC 0 xcvr-0/0/14                       57    62       72       72       75       75         79
FPC 0 xcvr-0/0/15                       57    62       72       72       75       75         79
FPC 0 xcvr-0/1/0                        55    60       70       70       73       73         76
FPC 0 xcvr-0/1/1                        52    57       67       67       70       70         73
```

```
FPC 0 xcvr-0/1/2                    52    57    67    67    70    70    75
FPC 0 xcvr-0/1/3                    52    57    67    67    70    70    73
FPC 0 xcvr-0/1/4                    52    57    67    67    70    70    73
FPC 0 xcvr-0/1/5                    52    57    67    67    70    70    75
FPC 0 xcvr-0/1/6                    55    60    70    70    73    73    76
FPC 0 xcvr-0/1/7                    52    57    67    67    70    70    75
FPC 0 xcvr-0/1/8                    55    60    70    70    73    73    76
FPC 0 xcvr-0/1/9                    52    57    67    67    70    70    75
FPC 0 xcvr-0/1/10                   67    72    82    82    85    85    88
FPC 0 xcvr-0/1/11                   52    57    67    67    70    70    75
FPC 0 xcvr-0/1/14                   55    60    70    70    73    73    76
FPC 0 xcvr-0/1/15                   52    57    67    67    70    70    73
FPC 0 xcvr-0/2/2                    55    60    70    70    73    73    76
FPC 0 xcvr-0/2/3                    52    57    67    67    70    70    75
FPC 0 xcvr-0/2/5                    52    57    67    67    70    70    75
FPC 0 xcvr-0/2/7                    52    57    67    67    70    70    75
FPC 0 xcvr-0/2/9                    52    57    67    67    70    70    75
FPC 0 xcvr-0/2/10                   52    57    67    67    70    70    75
FPC 0 xcvr-0/2/11                   52    57    67    67    70    70    75
FPC 0 xcvr-0/2/12                   67    72    82    82    85    85   100
FPC 0 xcvr-0/2/13                   67    72    82    82    85    85   100
FPC 0 xcvr-0/2/14                   75    80    90    90    93    93   110
FPC 0 xcvr-0/2/15                   55    60    70    70    73    73    76
```

Release Information

Command introduced in Junos OS Release 8.0.

`sfc` command introduced in Junos OS Release 9.6 for the TX Matrix Plus router.

`satellite` option introduced in Junos OS Release 14.2R3.

Command introduced in Junos OS Release 18.2R1 for MX10008 Routers and EX9253 Switches.

show chassis usb storage

IN THIS SECTION

Syntax | **18946**

- Description | 18946
- Required Privilege Level | 18946
- Sample Output | 18946
- Release Information | 18947

Syntax

```
show chassis usb storage
```

Description

This command displays the current status of any USB storage device and whether the USB ports are enabled or disabled.

Required Privilege Level

view

Sample Output

show chassis hardware detail

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                BV4911AA0005      SRX240H2-POE
Routing Engine   REV 01   750-043613   AAEC1923          RE-SRX240H2-POE
  usb0 (addr 1)  DWC OTG root hub 0    vendor 0x0000     uhub0
  usb0 (addr 2)  product 0x005a 90     vendor 0x0409     uhub1
  usb0 (addr 3)  ST72682  High Speed Mode 64218 STMicroelectronics umass0
  usb0 (addr 4)  Mass Storage Device 4096 JetFlash       umass1
FPC 0                                                    FPC
  PIC 0                                                  16x GE Base PIC
Power Supply 0
```

show chassis usb storage

```
user@host> show chassis usb storage
USB Disabled
```

Release Information

Command introduced in Junos OS Release 11.4 R2.

RELATED DOCUMENTATION

Installing Software on SRX Series Devices

show chassis zones

IN THIS SECTION

- Syntax | 18947
- Syntax (MX Series Routers) | 18948
- Syntax (QFX Series) | 18948
- Description | 18948
- Options | 18948
- Required Privilege Level | 18949
- Output Fields | 18949
- Sample Output | 18950
- Release Information | 18958

Syntax

```
show chassis zones
<detail>
```

Syntax (MX Series Routers)

```
show chassis zones
<detail>
<all-members>
<local>
<member member-id>
```

Syntax (QFX Series)

```
show chassis zones
<detail>
<interconnect-device name>
```

Description

(QFabric systems only) Display the status of the two cooling system zones on the Interconnect device. Zone 1 consists of eight (0 – 7) front cards, which are cooled by two fan trays. Zone 2 consists of two control boards and eight rear cards, which are cooled by eight (0 – 7) fan trays.

(MX2010, MX2020, and MX2008 routers only) Display the status of the cooling system zones of the chassis. Zone 0 consists of the Control Board, ten (0–9) FPCs, and their respective PICs, Switch Fabric Boards, and Adapter Cards. Zone 1 consists of the Routing Engine, Control Board, and Switch Processor Mezzanine Boards.

(PTX5000 Packet Transport Router only) Display the status of the two cooling system zones of the chassis. Zone 0 consists of the Routing Engine, Control Board, SIB, PMB, and the CCG, and is cooled by the vertical fan tray. Zone 1 consists of the eight (0–7) FPCs, and their respective PICs, and is cooled by the horizontal fan trays. The vertical fan tray is located at the front of the chassis. One horizontal fan tray is located at the front top of the chassis, and another is located at the front bottom of the chassis.

Options

all-members (MX2010, MX2020, and MX2008 routers only) (Optional) Display the status of the cooling system zones in all members of the Virtual Chassis configuration.

detail (MX2010, MX2020, and MX2008 routers only) (Optional) Display detailed status of the cooling system zones.

detail ***device-name*** (QFabric systems only) (Optional) Display detailed status of the two cooling systems on the Interconnect device.

interconnect-device ***name*** (QFabric systems only) (Optional) Display the status of the cooling zones on the Interconnect device.

local (MX2010, MX2020, and MX2008 routers only) (Optional) Display the status of the cooling system zones in the local member of the Virtual Chassis.

member ***member-id*** (MX2010, MX2020, and MX2008 routers only) (Optional) Display the status of the cooling system zones in the specified member of the Virtual Chassis. Replace *member-id* with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 350 on page 18949 lists the output fields for the `show chassis zones` command. Output fields are listed in the approximate order in which they appear.

Table 350: show chassis zones Output Fields

Field Name	Field Description
`Slot`	FPC slot number of the device whose content is being displayed. On QFX3500 standalone switches, the number is always `0`.
`Beacon State`	Status of the beacon state: • Off—The beacon is `OFF`. • On—The beacon is `ON`.
`show chassis zones` command output fields for MX2020, MX2010, and MX2008 routers:	
`Driving FRU`	Field replacable unit (FRU).

Table 350: show chassis zones Output Fields *(Continued)*

Field Name	Field Description
`Temperature`	Temperature of the specified FRU in degrees Celsius and degrees Fahrenheit.
`Condition`	Condition of the specified FRU. Condition can be `HIGH TEMP`, `WARM TEMP`, `OK`, and `Offline`.
`Num Fans Missing`	Number of fans or fan trays missing.
`Num Fans Failed`	Number of fans or fan trays that have failed.
`Fan Duty Cycle`	Fan duty cycle value.
`show chassis zones detail` command output fields for MX2020, MX2010, and MX2008 routers:	
`Item`	Chassis component: • Information about the chassis, Routing Engines, Control Boards (CBs), Switch Fabric Boards (SFBs), PICs, Flexible PIC Concentrators (FPCs), and Adapter Cards (ADCs).
`Measurement`	Fan tray speed utilization in percentage.
`Status`	Status of the specified item. Status can be `OK`, `Absent`, or `Offline`.

Sample Output

show chassis zones interconnect-device (QFabric System)

```
user@switch> show chassis zones interconnect-device
interconnect1
Slot            Beacon State
            FPC     0           OFF
```

show chassis zones (MX2010 Router)

```
user@host> show chassis zones
ZONE 0 Status
    Driving FRU                 FPC 6
    Temperature                 81 degrees C / 177 degrees F
    Condition                   HIGH TEMP
    Num Fans Missing            0
    Num Fans Failed             0
    Fan Duty Cycle              30

ZONE 1 Status
    Driving FRU                 SFB 0 Exhaust-Zone1
    Temperature                 71 degrees C / 159 degrees F
    Condition                   WARM TEMP
    Num Fans Missing            0
    Num Fans Failed             0
    Fan Duty Cycle              30
```

show chassis zones detail (MX2010 Router)

```
user@host > show chassis zones
ZONE 0 Status
Item                    Status                Measurement
CB 0                    WARM TEMP
CB 1                    WARM TEMP
FPC 0                   HIGH TEMP
FPC 1                   HIGH TEMP
FPC 2                   WARM TEMP
FPC 3                   HIGH TEMP
FPC 4                   HIGH TEMP
FPC 5                   HIGH TEMP
FPC 6                   HIGH TEMP
FPC 7                   HIGH TEMP
FPC 8                   HIGH TEMP
FPC 9                   HIGH TEMP
ADC 0                   WARM TEMP
ADC 1                   WARM TEMP
ADC 2                   WARM TEMP
ADC 3                   WARM TEMP
```

```
ADC 4                   WARM TEMP
ADC 5                   WARM TEMP
ADC 6                   WARM TEMP
ADC 7                   WARM TEMP
ADC 8                   WARM TEMP
ADC 9                   WARM TEMP
SFB 0                   WARM TEMP
SFB 1                   WARM TEMP
SFB 2                   WARM TEMP
SFB 3                   Offline
SFB 4                   HIGH TEMP
SFB 5                   WARM TEMP
SFB 6                   HIGH TEMP
SFB 7                   WARM TEMP
Fan Tray 0              OK                      Spinning at 98% fan tray speed
Fan Tray 1              OK                      Spinning at 98% fan tray speed

ZONE 1 Status
Item                    Status                  Measurement
CB 0                    WARM TEMP
CB 1                    WARM TEMP
Routing Engine 0        OK
Routing Engine 1        OK
SFB 0                   WARM TEMP
SFB 1                   WARM TEMP
SFB 2                   WARM TEMP
SFB 3                   Offline
SFB 4                   HIGH TEMP
SFB 5                   WARM TEMP
SFB 6                   HIGH TEMP
SFB 7                   WARM TEMP
SPMB 0                  OK
SPMB 1                  OK
Fan Tray 2              OK                      Spinning at 64% fan tray speed
Fan Tray 3              OK                      Spinning at 64% fan tray speed
```

show chassis zones (MX2020 Router)

```
user@host> show chassis zones
ZONE 0 Status
    Driving FRU                 FPC 0
```

```
    Temperature                 31 degrees C / 87 degrees F
    Condition                   OK
    Num Fans Missing            0
    Num Fans Failed             0
    Fan Duty Cycle              30

ZONE 1 Status
    Driving FRU                 FPC 19
    Temperature                 32 degrees C / 89 degrees F
    Condition                   OK
    Num Fans Missing            0
    Num Fans Failed             0
    Fan Duty Cycle              30
```

show chassis zones detail (MX2020 Router)

```
user@host> show chassis zones detail
ZONE 0 Status
Item                    Status                Measurement
CB 0                    OK
CB 1                    OK
FPC 0                   OK
FPC 1                   OK
FPC 2                   OK
FPC 3                   OK
FPC 4                   OK
FPC 5                   OK
FPC 6                   OK
FPC 7                   OK
FPC 8                   OK
FPC 9                   OK
ADC 0                   OK
ADC 1                   OK
ADC 2                   OK
ADC 3                   OK
ADC 4                   OK
ADC 5                   OK
ADC 6                   OK
ADC 7                   OK
ADC 8                   OK
ADC 9                   OK
```

```
SFB 0                   OK
SFB 1                   OK
SFB 2                   OK
SFB 3                   OK
SFB 4                   OK
SFB 5                   OK
SFB 6                   OK
SFB 7                   OK
Fan Tray 0              OK                      Spinning at 38% fan tray speed
Fan Tray 1              OK                      Spinning at 37% fan tray speed

ZONE 1 Status
Item                    Status                  Measurement
CB 0                    OK
CB 1                    OK
Routing Engine 0        OK
Routing Engine 1        OK
FPC 10                  OK
FPC 11                  OK
FPC 12                  OK
FPC 13                  OK
FPC 14                  OK
FPC 15                  OK
FPC 16                  OK
FPC 17                  OK
FPC 18                  OK
FPC 19                  OK
ADC 10                  OK
ADC 11                  OK
ADC 12                  OK
ADC 13                  OK
ADC 14                  OK
ADC 15                  OK
ADC 16                  OK
ADC 17                  OK
ADC 18                  OK
ADC 19                  OK
SFB 0                   OK
SFB 1                   OK
SFB 2                   OK
SFB 3                   OK
SFB 4                   OK
SFB 5                   OK
```

```
SFB 6                   OK
SFB 7                   OK
SPMB 0                  OK
SPMB 1                  OK
Fan Tray 2              OK                      Spinning at 38% fan tray speed
Fan Tray 3              OK                      Spinning at 38% fan tray speed
```

show chassis zones (MX2008 Router)

```
user@host> show chassis zones

ZONE 0 Status
    Driving FRU                 Routing Engine 0
    Temperature                 67 degrees C / 152 degrees F
    Condition                   WARM TEMP
    Num Fans Missing            0
    Num Fans Failed             0
    Fan Duty Cycle              27
```

show chassis zones detail (MX2008 Router)

```
user@host> show chassis zones detail
ZONE 0 Status
Item                    Status                Measurement
CB 0                    OK
CB 1                    OK
Routing Engine 0        OK
Routing Engine 1        OK
FPC 0                   OK
FPC 1                   Absent
FPC 2                   Absent
FPC 3                   OK
FPC 4                   Absent
FPC 5                   OK
FPC 6                   Absent
FPC 7                   OK
FPC 8                   Absent
FPC 9                   OK
ADC 0                   OK
ADC 1                   Absent
```

```
ADC 2                   Absent
ADC 3                   OK
ADC 4                   Absent
ADC 5                   OK
ADC 6                   Absent
ADC 7                   OK
ADC 8                   Absent
ADC 9                   Absent
SFB 0                   OK
SFB 1                   OK
SFB 2                   OK
SFB 3                   OK
SFB 4                   OK
SFB 5                   OK
SFB 6                   OK
SFB 7                   OK
SPMB 0                  OK
SPMB 1                  OK
Fan Tray 0              OK                      Spinning at 60% fan tray speed
Fan Tray 1              OK                      Spinning at 58% fan tray speed
```

show chassis beacon interconnect-device (QFabric System)

```
user@switch> show chassis beacon interconnect-device
interconnect1
Chassis              OFF
CB 0                 OFF
CB 1                 OFF
FC 0 FPC 0           OFF
FC 1 FPC 1           OFF
RC 0 FPC 8           OFF
RC 1 FPC 9           OFF
```

show chassis beacon interconnect-device fpc (QFabric System)

```
user@switch> show chassis beacon interconnect-device
interconnect1 fpc 0
FPC 0                ON
```

show chassis beacon node-device (QFabric System)

```
user@switch> show chassis beacon node-device
node1
node1                ON
```

show chassis beacon node-device fpc (QFabric System)

```
user@switch> show chassis beacon node-device
node1 fpc 0
FPC 0               ON
```

show chassis zones (PTX5000 Packet Transport Router)

```
user@host> show chassis zones
ZONE 0 Status
    Driving FRU                 Routing Engine 1
    Temperature                 62 degrees C / 143 degrees F
    Condition                   OK
    Num Fans Missing            0
    Num Fans Failed             0
    Fan Duty Cycle              0

ZONE 1 Status
    Driving FRU                 FPC 0 TL0
    Temperature                 71 degrees C / 159 degrees F
    Condition                   OK
    Num Fans Missing            0
    Num Fans Failed             0
    Fan Duty Cycle              0
```

show chassis zones detail (PTX5000 Packet Transport Router)

```
user@host> show chassis zones detail
ZONE 0 Status
Item                    Status                Measurement
CB 0                    OK
```

```
CB 1                    OK
Routing Engine 0        OK
Routing Engine 1        OK
SIB 0                   OK
SIB 1                   OK
SIB 2                   OK
SIB 3                   OK
SIB 4                   OK
SIB 5                   Absent
SIB 6                   Absent
SIB 7                   Absent
SIB 8                   Absent
Fan Tray 0              OK                      Spinning at 30% fan tray speed

ZONE 1 Status
Item                    Status                  Measurement
FPC 0                   OK
FPC 1                   OK
FPC 2                   OK
FPC 3                   OK
FPC 4                   OK
FPC 5                   Absent
FPC 6                   Offline
FPC 7                   OK
Fan Tray 1              OK                      Spinning at 33% fan tray speed
Fan Tray 2              OK                      Spinning at 36% fan tray speed
```

Release Information

Command introduced in Junos OS Release 11.3.

`all-members`, `local`, and `member` *`member-id`* options introduced in Junos OS Release 15.1 for MX2020 and MX2010 routers.

RELATED DOCUMENTATION

show chassis fan | **18260**

show chassis temperature-thresholds | **18914**

show class-of-service

IN THIS SECTION

- Syntax | 18959
- Description | 18959
- Options | 18959
- Required Privilege Level | 18959
- Output Fields | 18959
- Sample Output | 18962
- Release Information | 18965

Syntax

```
show class-of-service
```

Description

Display the entire class-of-service (CoS) configuration, including system-chosen defaults. Executing this command is equivalent to executing all `show class-of-service` commands in succession.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 351 on page 18960 lists all possible output fields for the `show class-of-service` command. The output fields that appear can vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

Table 351: show class-of-service Output Fields

Field Name	Field Description	Level of Output
`Forwarding class`	The forwarding class configuration: • `Forwarding class`—Name of the forwarding class. • `ID`—Forwarding class ID. • `Queue`—Queue number.	All levels
`Code point type`	The type of code-point alias: • `dscp`—Aliases for DiffServ code point (DSCP) values. • `ieee-802.1`—Aliases for IEEE 802.1p values. • `exp`—Aliases for MPLS EXP values.	All levels
`Alias`	Names given to CoS values.	All levels
`Bit pattern`	Set of bits associated with an alias.	All levels
`Classifier`	Name of the classifier.	All levels
`Code point`	Code-point values.	All levels
`Loss priority`	Loss priority assigned to specific CoS values and aliases of the classifier.	All levels
`Rewrite rule`	Name of the rewrite rule if one has been configured.	All levels
`Drop profile`	Name of the drop profile.	All levels
`Type`	Type of drop profile.	All levels
`Fill level`	Percentage of queue buffer fullness in a drop profile at which packets begin to drop during periods of congestion.	All levels

Table 351: show class-of-service Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Scheduler map`	Name of the scheduler map.	All levels
`Scheduler`	Name of the scheduler.	All levels
`Transmit rate`	Transmission rate of the scheduler.	All levels
`Excess rate`	Percentage of excess bandwidth traffic to share.	All levels
`Buffer size`	Delay buffer size in the queue.	All levels
`Drop profiles`	Drop profiles configured for the specified scheduler.	All levels
`Protocol`	Transport protocol corresponding to the drop profile.	All levels
`Name`	Name of the drop profile.	All levels
`Queues supported`	Number of queues that can be configured on the interface.	All levels
`Queues in use`	Number of queues currently configured.	All levels
`Physical interface`	Name of the physical interface.	All levels
`Scheduler map`	Name of the scheduler map.	All levels
`Congestion-notification`	Enabled if a congestion notification profile is applied to the interface; disabled if no congestion notification profile is applied to the interface.	All levels
`Forwarding class set`	Name of the forwarding class set (priority group). **NOTE**: Only on systems that support enhanced transmission selection (ETS) hierarchical port scheduling.	

Table 351: show class-of-service Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Index	Internal index of an object.	All levels

Sample Output

show class-of-service

```
user@host> show class-of-service
Forwarding class                      ID      Queue  Restricted queue  Fabric priority
Policing priority   SPU priority
  best-effort                          0       0          0             low
normal            low
  expedited-forwarding                 1       1          1             low
normal            low
  assured-forwarding                   2       2          2             low
normal            low
  network-control                      3       3          3             low
normal            low

Code point type: dscp
  Alias              Bit pattern
  af11               001010
  af12               001100
  af13               001110
...
Code point type: dscp-ipv6
  Alias              Bit pattern
  af11               001010
  af12               001100
  af13               001110
...
Code point type: exp
  Alias              Bit pattern
  af11               100
  af12               101
  be                 000
```

```
...
Code point type: ieee-802.1
  Alias              Bit pattern
  af11               100
  af12               101
  be                 000
...
Code point type: inet-precedence
  Alias              Bit pattern
  af11               001
  af21               010
  af31               011
...
Code point type: ieee-802.1ad
  Alias              Bit pattern
  af11               1000
  af11-dei           1001
  af12               1010
...
Code point type: inet6-precedence
  Alias              Bit pattern
  af11               001
  af21               010
  af31               011
...
Classifier: dscp-default, Code point type: dscp, Index: 6
  Code point         Forwarding class                    Loss priority
  000000             best-effort                         low
  000001             best-effort                         low
  000010             best-effort                         low
....
Classifier: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 7
  Code point         Forwarding class                    Loss priority
  000000             best-effort                         low
  000001             best-effort                         low
  000010             best-effort                         low
...
Loss-priority-map: frame-relay-de-default, Code point type: frame-relay-de, Index: 12
  Code point         Loss priority
  0                  low
  1                  high

Rewrite rule: dscp-default, Code point type: dscp, Index: 23
```

```
  Forwarding class                    Loss priority       Code point
  best-effort                         low                 000000
  best-effort                         high                000000
  expedited-forwarding                low                 101110
...
Rewrite rule: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 24
  Forwarding class                    Loss priority       Code point
  best-effort                         low                 000000
  best-effort                         high                000000
...
....
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100

Scheduler map: <default>, Index: 2

  Scheduler: <default-be>, Forwarding class: best-effort, Index: 16
    Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent, Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium low      any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>
...
Fabric priority: low
  Scheduler: <default-fabric>, Index: 24
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP         1    <default-drop-profile>
      Low             TCP             1    <default-drop-profile>
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

Fabric priority: high
  Scheduler: <default-fabric>, Index: 24
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP         1    <default-drop-profile>
      Low             TCP             1    <default-drop-profile>
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>
```

```
Physical interface: fe-0/0/0, Index: 137
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2

  Logical interface: fe-0/0/0.0, Index: 69
    Object            Name                   Type                    Index
    Adaptive-shaper   fr-shaper                                      35320
    Classifier        ipprec-compatibility   ip                         11

Physical interface: fe-0/0/1, Index: 138
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2
...
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service adjustment-control-profile

IN THIS SECTION

- Syntax | 18966
- Description | 18966
- Options | 18966
- Required Privilege Level | 18966
- Output Fields | 18966
- Sample Output | 18967
- Release Information | 18967

Syntax

```
show class-of-service adjustment-control-profile
<profile-name>
```

Description

For MPC/MIC interfaces only, display the adjustment control profiles.

Options

none	Display all profiles.
profile-name	(Optional) Display information about a single profile.

Required Privilege Level

view

Output Fields

Table 352 on page 18966 describes the output fields for the `show class-of-service adjustment-control-profile` command. Output fields are listed in the approximate order in which they appear.

Table 352: show class-of-service adjustment-control-profile Output Fields

Field Name	Field Description
`Name`	Name of the adjusting application. Possible values: • RADIUS-CoA—RADIUS CoA application. • ANCP—ANCP application. • PPPoE IA tags—PPPoE IA tag application. • DHCP tags—DHCP application.

Table 352: show class-of-service adjustment-control-profile Output Fields *(Continued)*

Field Name	Field Description
`Priority`	Priority of the adjusting application. Possible values are 1 through 10; 1 is the highest priority. The lower the priority value, the higher the priority
`Algorithm`	Algorithm the adjusting application uses to make adjustments. • adjust-never—Never perform rate adjustments. • adjust-always—Adjust the shaping rate unconditionally. • adjust-less—Adjust the shaping rate if it is less than the configured value. • adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value. • adjust-greater—Adjust the shaping rate if it is greater than the configured value. • adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.

Sample Output

show class-of-service adjustment-control-profile

```
user@host> show class-of-service adjustment-control-profile
user@host> show class-of-service adjustment-control-profile acp1
  name: ANCP, priority: 1, algorithm: less
  name: RADIUS CoA, priority: 1, algorithm: always
  name: PPPoE IA tags, priority: 2, algorithm: less
  name: DHCP tags, priority: 2, algorithm: less
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Verifying the CoS Adjustment Control Profile Configuration

show class-of-service application-traffic-control counter

IN THIS SECTION

- Syntax | 18968
- Description | 18968
- Required Privilege Level | 18968
- Output Fields | 18968
- Sample Output | 18969
- Release Information | 18974

Syntax

```
show class-of-service application-traffic-control counter
```

Description

Display the details of the sessions after applying application quality-of-service (AppQoS). The output includes AppQoS DSCP marking and honoring statistics with the number of sessions processed, marked, and honored.

Required Privilege Level

view

Output Fields

Table 353 on page 18969 lists the output fields for the `show class-of-service application-traffic-control counter` command. Output fields are listed in the approximate order in which they appear.

Table 353: show class-of-service application-traffic-control counter Output Fields

Field Name	Field Description
`pic`	PIC number of the accumulated statistics. **NOTE**: The PIC number is always displayed as `0` for SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, and SRX1600 firewalls.
`Sessions processed`	The number of sessions where the class of service was checked.
`Sessions marked`	The number of sessions marked based on application-aware DSCP marking.
`Sessions honored`	The number of sessions honored based on application-aware traffic honoring.
`Sessions rate limited`	The number of sessions that have been rate limited.
`Client-to-server flows rate limited`	The number of client-to-server flows that have been rate limited.
`Server-to-client flows rate limited`	The number of server-to-client flows that have been rate limited.

Sample Output

show class-of-service application-traffic-control counter

```
user@host> show class-of-service application-traffic-control counter
pic: 2/1
  Counter type                                          Value
  Sessions processed                                    300
  Sessions marked                                       200
  Sessions honored                                      0
  Sessions rate limited                                 100
  Client-to-server flows rate limited                   100
  Server-to-client flows rate limited                   70
```

```
pic: 2/0
  Counter type                                          Value
  Sessions processed                                    400
  Sessions marked                                       300
  Sessions honored                                      0
  Sessions rate limited                                 200
  Client-to-server flows rate limited                   200
  Server-to-client flows rate limited                   100
```

show class-of-service application-traffic-control counter (Unified Policies)

```
user@host> show class-of-service application-traffic-control counter
pic: 0/0
  Counter type                                                  Value
 Sessions processed                                             2
 Sessions marked                                                1
 Sessions honored                                               1
 Sessions rate limited                                          1
 Client-to-server flows rate limited                            0
 Server-to-client flows rate limited                            1
 Session default ruleset hit                                                    1
 Session ignored no default ruleset                             1
```

show class-of-service application-traffic-control counter logical-system LSYS1

```
user@host>show class-of-service application-traffic-control counter logical-system LSYS1
Logical System: LSYS1

pic: 0/0
  Counter type                                                  Value
 Sessions processed                                             1
 Sessions marked                                                0
 Sessions honored                                               0
 Sessions rate limited                                          0
 Client-to-server flows rate limited                            0
 Server-to-client flows rate limited                            0
 Session default ruleset hit                                    0
 Session ignored no default ruleset                             0
```

show class-of-service application-traffic-control counter logical-system all

```
user@host>show class-of-service application-traffic-control counter logical-system all
Logical System: root-logical-system

pic: 0/0
  Counter type                                                    Value
 Sessions processed                                               0
 Sessions marked                                                  0
 Sessions honored                                                 0
 Sessions rate limited                                            0
 Client-to-server flows rate limited                              0
 Server-to-client flows rate limited                              0
 Session default ruleset hit                                      0
 Session ignored no default ruleset                               0

Logical System: LSYS0

pic: 0/0
  Counter type                                                    Value
 Sessions processed                                               0
 Sessions marked                                                  0
 Sessions honored                                                 0
 Sessions rate limited                                            0
 Client-to-server flows rate limited                              0
 Server-to-client flows rate limited                              0
 Session default ruleset hit                                      0
 Session ignored no default ruleset                               0

Logical System: LSYS1

pic: 0/0
  Counter type                                                    Value
 Sessions processed                                               1
 Sessions marked                                                  0
 Sessions honored                                                 0
 Sessions rate limited                                            0
 Client-to-server flows rate limited                              0
 Server-to-client flows rate limited                              0
 Session default ruleset hit                                      0
 Session ignored no default ruleset                               0
```

```
Logical System: LSYS2

pic: 0/0
  Counter type                                                    Value
 Sessions processed                                               0
 Sessions marked                                                  0
 Sessions honored                                                 0
 Sessions rate limited                                            0
 Client-to-server flows rate limited                              0
 Server-to-client flows rate limited                              0
 Session default ruleset hit                                      0
 Session ignored no default ruleset                               0
```

show class-of-service application-traffic-control counter tenant TSYS1

```
user@host>show class-of-service application-traffic-control counter tenant TSYS1
Tenant System: TSYS1

pic: 0/0
  Counter type                                                    Value
 Sessions processed                                               1
 Sessions marked                                                  0
 Sessions honored                                                 0
 Sessions rate limited                                            0
 Client-to-server flows rate limited                              0
 Server-to-client flows rate limited                              0
 Session default ruleset hit                                      0
 Session ignored no default ruleset                               0
```

show class-of-service application-traffic-control counter tenant all

```
user@host>show class-of-service application-traffic-control counter tenant all
Tenant System: root-logical-system

pic: 0/0
  Counter type                                                    Value
 Sessions processed                                               0
 Sessions marked                                                  0
 Sessions honored                                                 0
 Sessions rate limited                                            0
```

```
 Client-to-server flows rate limited                                0
 Server-to-client flows rate limited                                0
 Session default ruleset hit                                        0
 Session ignored no default ruleset                                 0

Tenant System: TSYS0

pic: 0/0
  Counter type                                                      Value
 Sessions processed                                                 0
 Sessions marked                                                    0
 Sessions honored                                                   0
 Sessions rate limited                                              0
 Client-to-server flows rate limited                                0
 Server-to-client flows rate limited                                0
 Session default ruleset hit                                        0
 Session ignored no default ruleset                                 0

Tenant System: TSYS1

pic: 0/0
  Counter type                                                      Value
 Sessions processed                                                 1
 Sessions marked                                                    0
 Sessions honored                                                   0
 Sessions rate limited                                              0
 Client-to-server flows rate limited                                0
 Server-to-client flows rate limited                                0
 Session default ruleset hit                                        0
 Session ignored no default ruleset                                 0

Tenant System: TSYS2

pic: 0/0
  Counter type                                                      Value
 Sessions processed                                                 0
 Sessions marked                                                    0
 Sessions honored                                                   0
 Sessions rate limited                                              0
 Client-to-server flows rate limited                                0
 Server-to-client flows rate limited                                0
```

```
Session default ruleset hit                                    0
Session ignored no default ruleset                             0
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring AppQoS

show class-of-service application-traffic-control rate-limiters

IN THIS SECTION

- Syntax | 18974
- Description | 18975
- Options | 18975
- Additional Information | 18975
- Required Privilege Level | 18975
- Output Fields | 18976
- Sample Output | 18978
- Release Information | 18986

Syntax

```
show class-of-service application-traffic-control rate-limiters (rl-all | rl-name | summary)
<all-logical-systems-tenants>
<logical-system (logical-system-name | all)>
<root-logical-system>
<tenant (tenant-name | all)>
```

Description

Displays the application traffic control rate limiters information.

Options

rl-all	(Optional) Displays all application traffic control rate limiters information.
rl-name	(Optional) Displays the application traffic control rate limiters information by name.
summary	(Optional) Displays the application traffic control rate limiters information summary.
all-logical-systems-tenants	(Optional) Displays application traffic control rate limiters information of all logical systems and tenant systems.
logical-system *logical-system-name*	(Optional) Displays application traffic control rate limiters information of a specified logical system.
logical-system all	(Optional) Displays application traffic control rate limiters information of all logical systems.
root-logical-system	(Optional) Displays application traffic control rate limiters information of the root logical system as default.
tenant *tenant-name*	(Optional) Displays application traffic control rate limiters information of a specified tenant system.
logical-system all	(Optional) Displays application traffic control rate limiters information of all tenant systems.

Additional Information

Required Privilege Level

view

Output Fields

Table 354 on page 18976 lists the output fields for the `show class-of-service application-traffic-control rate-limiters` command. Output fields are listed in the approximate order in which they appear.

Table 354: show class-of-service application-traffic-control rate-limiters Output Fields

Field Name	Field Description
`LSYS/Tenant`	Name of the logical system or tenant system.
`ID`	ID of the rate limiter.
`BW-Limit`	The bandwidth of the rate limiter.
`Burst-Limit`	The burst limit of the rate limiter.
`PerIP-Burst-Limit`	Burst limit per IP.
`Strict-BW-hit`	Number of strict bandwidth hit.
`Total Session processed`	Total number of sessions processed.
`Total PerIP_RLs used`	Total number of rate limiters used per IP.
`PerIP-RLs`	Active rate limiters per IP.
`Sessions`	Number of active sessions.
`Pkts-Trans`	Number of packets transferred.
`Pkts-Drop`	Number of packets dropped.
`Rate-Trans(bps)`	Rate at which packets are transferred.

Table 354: show class-of-service application-traffic-control rate-limiters Output Fields *(Continued)*

Field Name	Field Description
`Rate-Drop(bps)`	Rate at which packets are dropped.
`Rate-LP(bps)`	Rate at which packets with LP are dropped.
`RL-Name`	Name of the rate limiter
`Counter type`	Type of the counter.
`Rate Limiters`	Number of rate limiters.
`Sessions Processed`	Number of sessions processed.
`PerIP RLimiters used`	Number of rate limiters used per IP.
`Active Sessions`	Number of active sessions.
`Active perIP RLimiters`	Number of active rate limiters per IP.
`Sessions discarded strict-bw`	Number of sessions discarded based on strict bandwidth.
`Sessions discarded no perIP`	Number of sessions discarded per IP.
`Sessions discarded no memory`	Number of sessions discarded due to no memory.
`Maximum Rate-Limiters`	Number of maximum rate limiters.
`Maximum perIP-RLimiters`	Number of maximum rate limiters per IP.
`Total Pkts Tx`	Number of total packets transferred.

Table 354: show class-of-service application-traffic-control rate-limiters Output Fields *(Continued)*

Field Name	Field Description
Total Pkts discarded	Number of total packets discarded.
Total Pkts high LP	Number of total high LP packets.
Rate Tx (bps)	Rate at which packets are transferred.
Rate Discard (bps)	Rate at which packets are dropped.
Rate LP (bps)	Rate at which packets with LP are dropped.

Sample Output

show class-of-service application-traffic-control rate-limiters rl-all

```
root@host> show class-of-service application-traffic-control rate-limiters rl-all
LSYS/Tenant ID   BW-Limit   PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP Rate-Trans(bps) Rate-
Drop(bps) Rate-LP(bps) RL-Name
root        3    1000       0         0        0          0         0       0
0              0            R1
```

show class-of-service application-traffic-control rate-limiters rl-name R1

```
root@host> show class-of-service application-traffic-control rate-limiters rl-name R1
Logical System: root-logical-system
Counter type                              Value
Name                                      R1
BW-Mode                                   -
ID                                        3
BW-Limit                                  1000
Burst-Limit                               128000
PerIP-BW-Limit                            0
PerIP-Burst-Limit                         0
```

```
Strict-BW-hit                           0
Total Session processed                 0
Total PerIP_RLs used                    0
Active-Sessions                         0
Active-PerIP-RL                         0
Total-Pkt-transmit                      0
Total-Pkt-Drop                          0
Total-Pkt-LP                            0
Total-Rate-Transmit(bps)                0
Total-Rate-Drop(bps)                    0
Total-Rate-LP(bps)                      0
RL-Pkt-transmit                         0
RL-Pkt-Drop                             0
RL-Pkt-LP                               0
RL-Rate-Transmit(bps)                   0
RL-Rate-Drop(bps)                       0
RL-Rate-LP(bps)                         0
```

show class-of-service application-traffic-control rate-limiters rl-all all-logical-systems-tenants

```
root@host> show class-of-service application-traffic-control rate-limiters rl-all all-logical-
systems-tenants
LSYS/Tenant ID BW-Limit PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP  Rate-Trans(bps) Rate-
Drop(bps)  Rate-LP(bps) RL-Name
root        3  1000     0         0        0          0         0        0
0               0            R1
TSYS1       1  1000     0         0        0          0         0        0
0               0            rt11
TSYS2       2  1500     0         0        0          0         0        0
0               0            rt12
```

show class-of-service application-traffic-control rate-limiters rl-all logical-system all

```
root@host> show class-of-service application-traffic-control rate-limiters rl-all logical-system
all
LSYS/Tenant ID  BW-Limit PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP Rate-Trans(bps) Rate-
Drop(bps) Rate-LP(bps) RL-Name
root        3   1000     0         0        0          0         0       0
0              0            R1
```

show class-of-service application-traffic-control rate-limiters rl-all tenant all

```
root@host> show class-of-service application-traffic-control rate-limiters rl-all tenant all
LSYS/Tenant ID BW-Limit  PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP Rate-Trans(bps) Rate-
Drop(bps)  Rate-LP(bps) RL-Name
TSYS1         1  1000      0         0        0          0         0       0
0               0            rt11
TSYS2         2  1500      0         0        0          0         0       0
0               0            rt12
```

show class-of-service application-traffic-control rate-limiters rl-all tenant TSYS1

```
root@host> show class-of-service application-traffic-control rate-limiters rl-all tenant TSYS1
LSYS/Tenant ID BW-Limit PerIP-RLs Sessions Pkts-Trans   Pkts-Drop  Pkts-LP  Rate-Trans(bps) Rate-
Drop(bps) Rate-LP(bps) RL-Name
TSYS1       1  1000     0         0        0            0          0        0
0              0            rt11
```

show class-of-service application-traffic-control rate-limiters summary all-logical-systems-tenants

```
root@host> show class-of-service application-traffic-control rate-limiters summary all-logical-
systems-tenants
Logical System: root-logical-system
Counter type                              Value
Rate Limiters                             1
Sessions Processed                        0
PerIP RLimiters used                      0
Active Sessions                           0
Active perIP RLimiters                    0
Sessions discarded strict-bw              0
Sessions discarded no perIP               0
Sessions discarded no memory              0
Maximum Rate-Limiters                     1000
Maximum perIP-RLimiters                   24000
Total Pkts Tx                             0
Total Pkts discarded                      0
Total Pkts high LP                        0
Rate Tx (bps)                             0
```

```
Rate Discard (bps)                        0
Rate LP (bps)                             0

Logical System: TSYS1
Counter type                              Value
Rate Limiters                             1
Sessions Processed                        0
PerIP RLimiters used                      0
Active Sessions                           0
Active perIP RLimiters                    0
Sessions discarded strict-bw              0
Sessions discarded no perIP               0
Sessions discarded no memory              0
Maximum Rate-Limiters                     1000
Maximum perIP-RLimiters                   24000
Total Pkts Tx                             0
Total Pkts discarded                      0
Total Pkts high LP                        0
Rate Tx (bps)                             0
Rate Discard (bps)                        0
Rate LP (bps)                             0

Logical System: TSYS2
Counter type                              Value
Rate Limiters                             1
Sessions Processed                        0
PerIP RLimiters used                      0
Active Sessions                           0
Active perIP RLimiters                    0
Sessions discarded strict-bw              0
Sessions discarded no perIP               0
Sessions discarded no memory              0
Maximum Rate-Limiters                     1000
Maximum perIP-RLimiters                   24000
Total Pkts Tx                             0
Total Pkts discarded                      0
Total Pkts high LP                        0
Rate Tx (bps)                             0
Rate Discard (bps)                        0
Rate LP (bps)                             0
```

show class-of-service application-traffic-control rate-limiters summary logical-system all

```
root@host> show class-of-service application-traffic-control rate-limiters summary logical-
system all
Logical System: root-logical-system
Counter type                              Value
Rate Limiters                             1
Sessions Processed                        0
PerIP RLimiters used                      0
Active Sessions                           0
Active perIP RLimiters                    0
Sessions discarded strict-bw              0
Sessions discarded no perIP               0
Sessions discarded no memory              0
Maximum Rate-Limiters                     1000
Maximum perIP-RLimiters                   24000
Total Pkts Tx                             0
Total Pkts discarded                      0
Total Pkts high LP                        0
Rate Tx (bps)                             0
Rate Discard (bps)                        0
Rate LP (bps)                             0
```

show class-of-service application-traffic-control rate-limiters summary tenant TSYS1

```
root@host> show class-of-service application-traffic-control rate-limiters summary tenant TSYS1
Logical System: TSYS1
Counter type                              Value
Rate Limiters                             1
Sessions Processed                        0
PerIP RLimiters used                      0
Active Sessions                           0
Active perIP RLimiters                    0
Sessions discarded strict-bw              0
Sessions discarded no perIP               0
Sessions discarded no memory              0
Maximum Rate-Limiters                     1000
Maximum perIP-RLimiters                   24000
Total Pkts Tx                             0
Total Pkts discarded                      0
```

```
Total Pkts high LP                        0
Rate Tx (bps)                             0
Rate Discard (bps)                        0
Rate LP (bps)                             0
```

show class-of-service application-traffic-control rate-limiters rl-all

```
root@host:TSYS1> show class-of-service application-traffic-control rate-limiters rl-all
LSYS/Tenant ID BW-Limit PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP Rate-Trans(bps) Rate-
Drop(bps) Rate-LP(bps) RL-Name
TSYS1         1  1000     0         0        0          0         0       0
0               0           R1
```

show class-of-service application-traffic-control rate-limiters rl-name rt11

```
root@host:TSYS1> show class-of-service application-traffic-control rate-limiters rl-name rt11
Logical System: TSYS1
Counter type                              Value
Name                                      rt11
BW-Mode                                   -
ID                                        1
BW-Limit                                  1000
Burst-Limit                               20000
PerIP-BW-Limit                            0
PerIP-Burst-Limit                         0
Strict-BW-hit                             0
Total Session processed                   0
Total PerIP_RLs used                      0
Active-Sessions                           0
Active-PerIP-RL                           0
Total-Pkt-transmit                        0
Total-Pkt-Drop                            0
Total-Pkt-LP                              0
Total-Rate-Transmit(bps)                  0
Total-Rate-Drop(bps)                      0
Total-Rate-LP(bps)                        0
RL-Pkt-transmit                           0
RL-Pkt-Drop                               0
RL-Pkt-LP                                 0
RL-Rate-Transmit(bps)                     0
```

```
RL-Rate-Drop(bps)                        0
RL-Rate-LP(bps)                          0
```

show class-of-service application-traffic-control rate-limiters summary

```
root@host:TSYS1> show class-of-service application-traffic-control rate-limiters summary
Logical System: TSYS1
Counter type                             Value
Rate Limiters                            1
Sessions Processed                       0
PerIP RLimiters used                     0
Active Sessions                          0
Active perIP RLimiters                   0
Sessions discarded strict-bw             0
Sessions discarded no perIP              0
Sessions discarded no memory             0
Maximum Rate-Limiters                    1000
Maximum perIP-RLimiters                  24000
Total Pkts Tx                            0
Total Pkts discarded                     0
Total Pkts high LP                       0
Rate Tx (bps)                            0
Rate Discard (bps)                       0
Rate LP (bps)                            0
```

show class-of-service application-traffic-control rate-limiters rl-all

```
root@host:LSYS1> show class-of-service application-traffic-control rate-limiters rl-all
LSYS/Tenant ID BW-Limit PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP Rate-Trans(bps) Rate-
Drop(bps) Rate-LP(bps) RL-Name
LSYS1       1  1000     0         0        0          0         0       0
0              0            rt1
```

show class-of-service application-traffic-control rate-limiters rl-name rt1

```
root@host:LSYS1> show class-of-service application-traffic-control rate-limiters rl-name rt1
Logical System: LSYS1
Counter type                             Value
```

```
Name                                    rt1
BW-Mode                                 -
ID                                      1
BW-Limit                                1000
Burst-Limit                             128000
PerIP-BW-Limit                          0
PerIP-Burst-Limit                       0
Strict-BW-hit                           0
Total Session processed                 0
Total PerIP_RLs used                    0
Active-Sessions                         0
Active-PerIP-RL                         0
Total-Pkt-transmit                      0
Total-Pkt-Drop                          0
Total-Pkt-LP                            0
Total-Rate-Transmit(bps)                0
Total-Rate-Drop(bps)                    0
Total-Rate-LP(bps)                      0
RL-Pkt-transmit                         0
RL-Pkt-Drop                             0
RL-Pkt-LP                               0
RL-Rate-Transmit(bps)                   0
RL-Rate-Drop(bps)                       0
RL-Rate-LP(bps)                         0
```

show class-of-service application-traffic-control rate-limiters summary

```
root@host:LSYS1> show class-of-service application-traffic-control rate-limiters summary
Logical System: LSYS1
Counter type                            Value
Rate Limiters                           1
Sessions Processed                      0
PerIP RLimiters used                    0
Active Sessions                         0
Active perIP RLimiters                  0
Sessions discarded strict-bw            0
Sessions discarded no perIP             0
Sessions discarded no memory            0
Maximum Rate-Limiters                   1000
Maximum perIP-RLimiters                 24000
Total Pkts Tx                           0
```

```
Total Pkts discarded                    0
Total Pkts high LP                      0
Rate Tx (bps)                           0
Rate Discard (bps)                      0
Rate LP (bps)                           0
```

Release Information

Command introduced in Junos OS Release 11.4.

`all-logical-systems-tenants`, `logical-system`, `root-logical-system`, and `tenant` options introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

clear class-of-service application-traffic-control rate-limiters | 3478

show class-of-service application-traffic-control statistics rate-limiter

IN THIS SECTION

- Syntax | 18986
- Description | 18987
- Required Privilege Level | 18987
- Output Fields | 18987
- Sample Output | 18988
- Release Information | 18989

Syntax

```
show class-of-service application-traffic-control statistics rate-limiter
```

Description

Display AppQoS real-time run information about application rate limiting of current or recent sessions.

Required Privilege Level

view

Output Fields

Table 355 on page 18987 lists the output fields for the `show class-of-service application-traffic-control statistics rate-limiter` command. Output fields are listed in the approximate order in which they appear.

Table 355: show class-of-service application-traffic-control statistics rate-limiter Output Fields

Field Name	Field Description
`pic`	PIC number. **NOTE**: The PIC number is always displayed as `0` for SRX300, SRX320, SRX340, SRX345, and SRX550M firewalls.
`Ruleset`	The rule set applied on the session.
`Application`	The application match for applying the rule set.
`Client-to-server`	The rate limiter applied from client to server.
`Rate(kbps)`	The rate in the client-to-server direction
`Server-to-client`	The rate limiter applied from server to client.
`Rate(kbps)`	The rate in the server-to-client direction.

Sample Output

show class-of-service application-traffic-control statistics rate-limiter

```
user@host> show class-of-service application-traffic-control statistics rate-limiter
pic: 2/1
  Ruleset       Application  Client-to-server   Rate(kbps)   Server-to-client   Rate(kbps)
  my-ruleset-1  HTTP         my-http-c2s-rl     10000000    my-http-s2c-rl     20000000
  my-ruleset-2  HTTP         my-http-c2s-rl-2   20000000    my-http-s2c-rl-2   30000000
  my-ruleset-2  FTP          my-ftp-c2s-rl      50000       my-ftp-s2c-rl      50000
  ...

pic: 2/0
  Ruleset       Application  Client-to-server   Rate(kbps)   Server-to-client   Rate(kbps)
  my-ruleset-1  HTTP         my-http-c2s-rl     10000000    my-http-s2c-rl     20000000
  my-ruleset-2  HTTP         my-http-c2s-rl-2   20000000    my-http-s2c-rl-2   30000000
  my-ruleset-2  FTP          my-ftp-c2s-rl      50000       my-ftp-s2c-rl      50000
```

show class-of-service application-traffic-control statistics rate-limiter logical-system LSYS1

```
user@host>show class-of-service application-traffic-control statistics rate-limiter logical-
system LSYS1
Logical System: LSYS1

pic: 0/0
```

show class-of-service application-traffic-control statistics rate-limiter logical-system all

```
user@host>show class-of-service application-traffic-control statistics rate-limiter logical-
system all
 class-of-service application-traffic-control statistics rate-limiter logical-system all
Logical System: root-logical-system

pic: 0/0

Logical System: LSYS0

pic: 0/0
```

```
Logical System: LSYS1

pic: 0/0

Logical System: LSYS2

pic: 0/0
```

show class-of-service application-traffic-control statistics rate-limiter tenant TSYS1

```
user@host>show class-of-service application-traffic-control statistics rate-limiter tenant TSYS1
Tenant System: LSYS1

pic: 0/0
```

show class-of-service application-traffic-control statistics rate-limiter tenant all

```
user@host>show class-of-service application-traffic-control statistics rate-limiter tenant all
Tenant System: root-logical-system

pic: 0/0
Tenant System: TSYS0

pic: 0/0
Tenant System: TSYS1

pic: 0/0
Tenant System: TSYS2

pic: 0/0
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring AppQoS

show class-of-service application-traffic-control statistics rule

IN THIS SECTION

- Syntax | 18990
- Description | 18990
- Required Privilege Level | 18990
- Output Fields | 18990
- Sample Output | 18991
- Release Information | 18993

Syntax

```
show class-of-service application-traffic-control statistics rule
```

Description

Display the number of time an AppQoS rule was applied.

Required Privilege Level

view

Output Fields

Table 356 on page 18991 lists the output fields for the `show class-of-service application-traffic-control statistics rule` command. Output fields are listed in the approximate order in which they appear.

Table 356: show class-of-service application-traffic-control statistics rule Output Fields

Field Name	Field Description
`pic`	PIC number where the rule is applied. **NOTE**: The PIC number is always displayed as `0` for for SRX300, SRX320, SRX340, SRX345, and SRX550M firewalls.
`Ruleset`	The rule set containing the rule.
`Rule`	The rule to which the statistic applies.
`Hits`	The number of times a match for the rule was encountered.

Sample Output

show class-of-service application-traffic-control statistics rule

```
user@host> show class-of-service application-traffic-control statistics rule
 class-of-service application-traffic-control statistics rule
pic: 2/0
  Ruleset          Rule           Hits
  my-ruleset-1     ftp-rule       100
  my-ruleset-1     httpp-rule     100
  my-ruleset-2     telnet-rule    300
  my-ruleset-2     smtp-rule      300
  ...

pic: 2/1
  Ruleset          Rule           Hits
  my-ruleset-1     ftp-rule       200
  my-ruleset-1     httpp-rule     300
  my-ruleset-2     telnet-rule    400
  my-ruleset-2     smtp-rule      500
```

show class-of-service application-traffic-control statistics rule logical-system LSYS1

```
user@host>show class-of-service application-traffic-control statistics rule logical-system LSYS1
Logical System: LSYS1

pic: 0/0
```

show class-of-service application-traffic-control statistics rule logical-system all

```
user@host>show class-of-service application-traffic-control statistics rule logical-system all
Logical System: root-logical-system

pic: 0/0

Logical System: LSYS0

pic: 0/0

Logical System: LSYS1

pic: 0/0

Logical System: LSYS2

pic: 0/0
```

show class-of-service application-traffic-control statistics rule tenant TSYS1

```
user@host>show class-of-service application-traffic-control statistics rule tenant TSYS1
Tenant System: TSYS1

pic: 0/0
```

show class-of-service application-traffic-control statistics rule tenant all

```
user@host>show class-of-service application-traffic-control statistics rule tenant all
Tenant System: root-logical-system
```

```
pic: 0/0
Tenant System: TSYS0

pic: 0/0
Tenant System: TSYS1

pic: 0/0
Tenant System: TSYS2

pic: 0/0
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring AppQoS

show class-of-service classifier

IN THIS SECTION

- Syntax | 18994
- Description | 18994
- Options | 18994
- Required Privilege Level | 18994
- Output Fields | 18994
- Sample Output | 18995
- Release Information | 18998

Syntax

```
show class-of-service classifier
<name name>
<type dscp | type dscp-ipv6 | type exp | type ieee-802.1 | type inet-precedence>
```

Description

For each class-of-service (CoS) classifier, display the mapping of code point value to forwarding class and loss priority.

Options

none	Display all classifiers.
name *name*	(Optional) Display named classifier.
type dscp	(Optional) Display all classifiers of the Differentiated Services code point (DSCP) type.
type dscp-ipv6	(Optional) Display all classifiers of the DSCP for IPv6 type.
type exp	(Optional) Display all classifiers of the MPLS experimental (EXP) type.
type ieee-802.1	(Optional) Display all classifiers of the ieee-802.1 type.
type inet-precedence	(Optional) Display all classifiers of the inet-precedence type.

Required Privilege Level

view

Output Fields

Table 357 on page 18995 describes the output fields for the `show class-of-service classifier` command. Output fields are listed in the approximate order in which they appear.

Table 357: show class-of-service classifier Output Fields

Field Name	Field Description
`Classifier`	Name of the classifier.
`Code point type`	Type of the classifier: `exp` (not on EX Series switch), `dscp`, `dscp-ipv6` (not on EX Series switch), `ieee-802.1`, or `inet-precedence`.
`Index`	Internal index of the classifier.
`Code point`	Code point value used for classification
`Forwarding class`	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
`Loss priority`	Loss priority value used for classification. For most platforms, the value is `high` or `low`. For some platforms, the value is `high`, `medium-high`, `medium-low`, or `low`.

Sample Output

show class-of-service classifier

```
user@host>show class-of-service classifier
Classifier: dscp-default, Code point type: dscp, Index: 7
  Code point       Forwarding class       Loss priority
  000010           best-effort            low
  000011           assured-forwarding     high
  110111           network-control        low
  111111           expedited-forwarding   high

Classifier: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 8
  Code point       Forwarding class       Loss priority
  000010           best-effort            low
  000011           assured-forwarding     high
  110111           network-control        low
  111111           expedited-forwarding   high
```

```
Classifer: exp-default, Code point type: exp, Index: 10
  Code point       Forwarding class       Loss priority
  010              best-effort            low
  011              assured-forwarding     high
  110              network-control        low
  111              expedited-forwarding   high

Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
  Code point       Forwarding class       Loss priority
  010              best-effort            low
  011              assured-forwarding     high
  110              network-control        low
  111              expedited-forwarding   high

Classifier: ipprec-default, Code point type: inet-precedence, Index: 12
  Code point       Forwarding class       Loss priority
  010              best-effort            low
  011              assured-forwarding     low
  110              network-control        low
  111              expedited-forwarding   high

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
  Code point       Forwarding class       Loss priority
  010              best-effort            low
  110              network-control        high

Classifier: ieee8021p-untrust, Code point type: ieee-802.1, Index: 16
  Code point       Forwarding class       Loss priority
  000              best-effort            low
  001              best-effort            low

Classifier: dscp-mcast, Code point type: dscp, Index: 17
  Code point       Forwarding class       Loss priority
  000010           mcast-be               low
  000011           mcast-be               low
  110011           mcast-nc               low
  111000           mcast-nc               low

Classifier: ieee8021ad-default, Code point type: ieee-802.1ad, Index: 42
  Code point       Forwarding class       Loss priority
  0000             best-effort            low
  0110             best-effort            high
```

```
  0111             assured-forwarding     low
  1100             network-control        low
  1110             expedited-forwarding   high

Classifier: ieee-mcast, Code point type: ieee-802.1, Index: 47
  Code point       Forwarding class       Loss priority
  000              mcast-be               low
  010              mcast-be               low
  110              mcast-nc               low
  111              mcast-nc               low

Classifier: dscp-v6-mcast, Code point type: dscp-ipv6, Index: 46
  Code point       Forwarding class       Loss priority
  000010           mcast-be               low
  000011           mcast-be               low
  110011           mcast-nc               low
  111000           mcast-nc               low
```

show class-of-service classifier type ieee-802.1

```
user@host> show class-of-service classifier type ieee-802.1
Classifier: ieee802.1-default, Code point type: ieee-802.1, Index: 3
Code Point         Forwarding Class                  Loss priority
  000              best-effort                         low
  001              best-effort                         high
  010              expedited-forwarding                low
  011              expedited-forwarding                high
  100              assured-forwarding                  low
  101              assured-forwarding                  medium-high
  110              network-control                     low
  111              network-control                     high

Classifier: users-ieee802.1, Code point type: ieee-802.1
  Code point         Forwarding class                    Loss priority
  100                expedited-forwarding                low
```

show class-of-service classifier type ieee-802.1 (QFX Series)

```
                  user@switch>                  show class-of-service classifier type
ieee-802.1
```

```
Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
  Code point         Forwarding class                    Loss priority
  000                best-effort                         low
  001                best-effort                         low
  010                best-effort                         low
  011                fcoe                                low
  100                no-loss                             low
  101                best-effort                         low
  110                network-control                     low
  111                network-control                     low

Classifier: ieee8021p-untrust, Code point type: ieee-802.1, Index: 16
  Code point         Forwarding class                    Loss priority
  000                best-effort                         low
  001                best-effort                         low
  010                best-effort                         low
  011                best-effort                         low
  100                best-effort                         low
  101                best-effort                         low
  110                best-effort                         low
  111                best-effort                         low

Classifier: ieee-mcast, Code point type: ieee-802.1, Index: 46
  Code point         Forwarding class                    Loss priority
  000                mcast                               low
  001                mcast                               low
  010                mcast                               low
  011                mcast                               low
  100                mcast                               low
  101                mcast                               low
  110                mcast                               low
  111                mcast                               low
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service code-point-aliases

IN THIS SECTION

- Syntax | 18999
- Description | 18999
- Options | 18999
- Required Privilege Level | 19000
- Output Fields | 19000
- Sample Output | 19000
- Release Information | 19001

Syntax

```
show class-of-service code-point-aliases
<dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence>
```

Description

Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns.

Options

none	Display code point aliases of all code point types.
dscp	(Optional) Display Differentiated Services code point (DSCP) aliases.
dscp-ipv6	(Optional) Display IPv6 DSCP aliases.
exp	(Optional) Display MPLS EXP code point aliases.
ieee-802.1	(Optional) Display IEEE-802.1 code point aliases.
inet-precedence	(Optional) Display IPv4 precedence code point aliases.

Required Privilege Level

view

Output Fields

Table 358 on page 19000 describes the output fields for the `show class-of-service code-point-aliases` command. Output fields are listed in the approximate order in which they appear.

Table 358: show class-of-service code-point-aliases Output Fields

Field Name	Field Description
`Code point type`	Type of the code points displayed: `dscp`, `dscp-ipv6` (not on EX Series switch), `exp` (not on EX Series switch or the QFX Series), `ieee-802.1`, or `inet-precedence` (not on the QFX Series).
`Alias`	Alias for a bit pattern.
`Bit pattern`	Bit pattern for which the alias is displayed.

Sample Output

show class-of-service code-point-aliases exp

```
user@host> show class-of-service code-point-aliases exp
Code point type: exp
  Alias              Bit pattern
  af11               100
  af12               101
  be                 000
  be1                001
  cs6                110
  cs7                111
  ef                 010
  ef1                011
  nc1                110
  nc2                111
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service congestion-notification

IN THIS SECTION

- Syntax | 19001
- Description | 19001
- Options | 19001
- Required Privilege Level | 19002
- Output Fields | 19002
- Sample Output | 19003
- Sample Output | 19004
- Release Information | 19006

Syntax

```
show class-of-service congestion-notification <profile-name>
```

Description

Display whether priority-based flow control (PFC) is enabled for each IEEE 802.1p code point or for each DSCP value for DSCP-based PFC.

Options

none Display the PFC state for all IEEE 802.1p code points or DSCP values.

profile-name Display the PFC state for all IEEE 802.1p code points or DSCP values for the specified congestion notification profile.

Required Privilege Level

view

Output Fields

Table 359 on page 19002 describes the output fields for the `show class-of-service congestion-notification` command. Output fields are listed in the approximate order in which they appear.

Table 359: show class-of-service congestion-notification Output Fields

Field Name	Field Description
`Type`	Type of interfaces on which congestion notification is applied. Congestion notification is applied on input interfaces.
`Index`	Index of this congestion notification profile.
`Name`	Name of the congestion notification profile.
`Cable Length`	Length of the attached physical cable in meters. The default value is 100 meters.
`Priority`	IEEE 802.1p code points, or for DSCP-based PFC, DSCP values.
`PFC`	State of PFC for the corresponding code point, either `enabled` or `disabled`.
`MRU`	Maximum receive unit of the interface in bytes. (Incoming traffic that exceeds the MRU size of an interface is dropped.) The default values are: • 2500 bytes for priority 3 traffic • 9216 bytes for priority 4 traffic **NOTE**: If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not configure an MRU value, the default MRU value is 2500 bytes.

Table 359: show class-of-service congestion-notification Output Fields *(Continued)*

Field Name	Field Description
`Flow-Control-Queues`	Output queue mapping to IEEE 802.1p code points (priorities). Explicit output queue to priority mapping overwrites the default configuration, and only explicitly mapped queues are displayed in the output. Flow control is only enabled on a queue when you enable PFC on the corresponding priority in the input stanza of the congestion notification profile.

Sample Output

show class-of-service congestion-notification

```
user@switch> show class-of-service congestion-notification
Name: fcoe_p3_cnp, Index: 12037
Type: Input
Cable Length: 100 m
  Priority     PFC          MRU
  000          Disabled
  001          Disabled
  010          Disabled
  011          Enabled      2500
  100          Enabled      9216
  101          Disabled
  110          Disabled
  111          Disabled
Type: Output
  Priority     Flow-Control-Queues
  000
               0
  001
               1
  010
               2
  011
               3
  100
               4
  101
               5
```

```
  110
                6
  111
                7

Name: fcoe_p3_p5_cnp, Index: 46484
Type: Input
Cable Length: 100 m
  Priority     PFC          MRU
  000          Disabled
  001          Disabled
  010          Disabled
  011          Enabled      2240
  100          Disabled
  101          Enabled      2240
  110          Disabled
  111          Disabled
Type: Output
  Priority     Flow-Control-Queues
  011
                3
  101
                5
```

Sample Output

show class-of-service congestion-notification (QFX Series with DSCP-based PFC and a specified congestion notification profile)

```
user@switch> show class-of-service congestion-notification dscp_cnp
Name: dscp_cnp, Index: 13825
Type: Input
Cable Length: 100 m
  Priority     PFC          MRU
  000000       Disabled
  000001       Disabled
  000010       Disabled
  000011       Disabled
  000100       Disabled
  000101       Disabled
  000110       Disabled
```

```
    000111       Disabled
    001000       Disabled
    001001       Disabled
    001010       Disabled
    001011       Disabled
    001100       Disabled
    001101       Disabled
    001110       Disabled
    001111       Disabled
    010000       Disabled
    010001       Disabled
    010010       Disabled
    010011       Disabled
    010100       Disabled
    010101       Disabled
    010110       Disabled
    010111       Disabled
    011000       Disabled
    011001       Disabled
    011010       Disabled
    011011       Disabled
    011100       Disabled
    011101       Disabled
    011110       Disabled
    011111       Disabled
    100000       Disabled
    100001       Disabled
    100010       Disabled
    100011       Disabled
    100100       Disabled
    100101       Disabled
    100110       Disabled
    100111       Disabled
    101000       Disabled
    101001       Disabled
    101010       Disabled
    101011       Disabled
    101100       Disabled
    101101       Disabled
    101110       Disabled
    101111       Disabled
    110000       Enabled      3000
    110001       Disabled
```

```
  110010       Disabled
  110011       Disabled
  110100       Disabled
  110101       Disabled
  110110       Disabled
  110111       Disabled
  111000       Enabled      2000
  111001       Disabled
  111010       Disabled
  111011       Disabled
  111100       Enabled      4000
  111101       Disabled
  111110       Disabled
  111111       Disabled
Type: Output
  Priority     Flow-Control-Queues
  111          1
```

Release Information

Command introduced in Junos OS Release 11.1.

Output for DSCP values introduced for DSCP-based PFC in Junos OS Release 17.4R1 for the QFX Series.

RELATED DOCUMENTATION

Configuring CoS PFC (Congestion Notification Profiles)

Example: Configuring CoS PFC for FCoE Traffic

Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)

Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface

Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces

Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)

Example: Configuring PFC Across Layer 3 Interfaces

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Understanding PFC Using DSCP at Layer 3 for Untagged Traffic

show class-of-service dedicated-buffer-profile

IN THIS SECTION

- Syntax | 19007
- Description | 19007
- Options | 19007
- Required Privilege Level | 19007
- Output Fields | 19008
- Sample Output | 19008
- Release Information | 19008

Syntax

```
show class-of-service dedicated-buffer-profile
<profile name>
```

Description

Display information about all configured dedicated buffer profiles.

Options

none	Display all dedicated buffer profiles.
profile name	Display information for a specific dedicated buffer profile.

Required Privilege Level

view

Output Fields

Table 360 on page 19008 lists the output fields for the `show class-of-service dedicated-buffer-profile` command.

Table 360: show class-of-service dedicated-buffer-profile Output Fields

Field Name	Field Description
Dedicated Buffer Profile	Configured name of the dedicated buffer profile.
Index	Internal index of the dedicated buffer profile.
Ingress Buffer Size	Size of the ingress buffer, in cells.
Egress Buffer Size	Size of the egress buffer, in cells.

Sample Output

show class-of-service dedicated-buffer-profile

```
user@host> show class-of-service dedicated-buffer-profile
```

```
Dedicated Buffer Profile: dbp1, Index: 1
Ingress Buffer Size: None
Egress Buffer Size: 1000
```

Release Information

Command introduced in Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Configuring Ingress and Egress Dedicated Buffers

dedicated-buffer-profile | 5172

show class-of-service drop-profile

IN THIS SECTION

- Syntax | 19009
- Description | 19009
- Options | 19009
- Required Privilege Level | 19010
- Output Fields | 19010
- Sample Output | 19010
- Release Information | 19011

Syntax

```
show class-of-service drop-profile
<profile-name profile-name>
```

Description

Display data points for each class-of-service (CoS) random early detection (RED) drop profile.

Options

none	Display all drop profiles.
profile-name *profile-name*	(Optional) Display the specified profile only.

Required Privilege Level

view

Output Fields

Table 361 on page 19010 describes the output fields for the `show class-of-service drop-profile` command. Output fields are listed in the approximate order in which they appear.

Table 361: show class-of-service drop-profile Output Fields

Field Name	Field Description
Drop profile	Name of a drop profile.
Type	Type of drop profile: • **discrete** (default) • **interpolated**
Index	Internal index of this drop profile.
Fill Level	Percentage fullness of a queue.
Drop probability	Drop probability at this fill level.

Sample Output

show class-of-service drop-profile

```
user@host> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: user-drop-profile, Type: interpolated, Index: 2989
  Fill level    Drop probability
           0                   0
```

```
            1                    1
            2                    2
            4                    4
            5                    5
            6                    6
            8                    8
           10                   10
           12                   15
           14                   20
           15                   23
        ... 64 entries total
           90                   96
           92                   96
           94                   97
           95                   98
           96                   98
           98                   99
           99                   99
          100                  100
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service fabric scheduler-map

IN THIS SECTION

- Syntax | 19012
- Description | 19012
- Options | 19012
- Required Privilege Level | 19012
- Output Fields | 19012
- Sample Output | 19013
- Release Information | 19013

Syntax

```
show class-of-service fabric scheduler-map
```

Description

(M320 routers, MX Series routers, T Series routers and EX Series switches only) Display the mapping of class-of-service (CoS) schedulers to switch fabric traffic priorities and a summary of scheduler parameters for each priority.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 362 on page 19012 describes the output fields for the `show class-of-service fabric scheduler-map` command. Output fields are listed in the approximate order in which they appear.

Table 362: show class-of-service fabric scheduler-map Output Fields

Field Name	Field Description
`Fabric priority`	Indicates the fabric traffic priority. Currently, two priorities are supported: `low` and `high`.
`Scheduler`	Name of the scheduler.
`Index`	Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.

Table 362: show class-of-service fabric scheduler-map Output Fields *(Continued)*

Field Name	Field Description
`Drop profiles`	Display the assignment of drop profile by name and index to a given loss priority and protocol pair: • `Loss priority`—Packet loss priority for drop profile assignment. • `Protocol`—Transport protocol for drop profile assignment. • `Name`—Name of the drop profile.

Sample Output

show class-of-service fabric scheduler-map

```
 user@host> show class-of-service fabric scheduler-map
Fabric priority: low
  Scheduler: fab-ef-scheduler, Index: 60211
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP     44321    fab-ef-profile
      Low             TCP         44321    fab-ef-profile
      High            non-TCP     44321    fab-ef-profile
      High            TCP         44321    fab-ef-profile

Fabric priority: high
  Scheduler: fab-ef-scheduler, Index: 60211
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP     44321    fab-ef-profile
      Low             TCP         44321    fab-ef-profile
      High            non-TCP     44321    fab-ef-profile
      High            TCP         44321    fab-ef-profile
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service fabric statistics

IN THIS SECTION

- Syntax | 19014
- Description | 19014
- Options | 19015
- Required Privilege Level | 19015
- Output Fields | 19015
- Sample Output | 19017
- Release Information | 19019

Syntax

```
show class-of-service fabric statistics
<destination fpc-number>
<detail>
<source fpc-number>
<summary>
```

Description

Display class-of-service (CoS) switch fabric queue statistics.

NOTE: On the Switch Control Board (SCB) and the SCBE on MX Series routers, the ratio between high-priority queue and low-priority queue for traffic scheduled to enter the fabric is 85:15. However, on the SCBE2, this ratio is 97:3.

NOTE: After an FPC restart, executing this command can return an `Error = Operation timed out` message for up to a minute even though the FPC is back online. No statistics are lost during this time, however.

Options

none	Same as summary.
destination *fpc-number*	(Optional) Display details for the specified destination Flexible PIC Concentrator (FPC). The FPC number is a value from 0 through 7 (depending on platform).
detail	(Optional) Display detailed statistics at the PFE level.
source *fpc-number*	(Optional) Display details for the specified source FPC. The FPC number is a value from 0 through 7 (depending on platform).
summary	(Optional) Display all switch fabric statistics.

detail Options

none	Display details for all FPCs and PFEs.
destination-fpc *fpc-number*	(Optional) Display details for the specified destination FPC. The FPC number is a value from 0 through 7 (depending on platform).
destination-pfe *pfe-number*	(Optional) Display details for the specified destination packet forwarding engine (PFE). The PFE number is a value from 0 through 15 (depending on platform).
source-fpc *fpc-number*	(Optional) Display details for the specified source FPC. The FPC number is a value from 0 through 7 (depending on platform).
source-pfe *pfe-number*	(Optional) Display details for the specified source PFE. The PFE number is a value from 0 through 15 (depending on platform).

Required Privilege Level

view

Output Fields

Table 363 on page 19016 describes the output fields for the `show class-of-service fabric statistics` command. Output fields are listed in the approximate order in which they appear.

Table 363: show class-of-service fabric statistics Output Fields

Field Name	Field Description
`Destination FPC Index`	Index number associated with the destination FPC.
Destination Pfe Index	Index number associated with the destination PFE.
`Source PFC Index`	Index number associated with the source FPC.
Source Pfe Index	Index number associated with the source PFE.
`Total statistics`	Fabric queue statistic totals: • `Packets`—Total packet count for high-priority and low-priority queues. • `Bytes`—Total byte count for high-priority and low-priority queues. • `Pps`—Total packets-per-second count for high-priority and low-priority queues. • `bps`—Total bits-per-second count for high-priority and low-priority queues.
`Tx statistics`	Fabric queue statistics for transmitted traffic: • `Packets`—Transmitted packet count for high-priority and low-priority queues. • `Bytes`—Transmitted byte count for high-priority and low-priority queues. • `Pps`—Transmitted packets-per-second count for high-priority and low-priority queues. • `bps`—Transmitted bits-per-second count for high-priority and low-priority queues.

Table 363: show class-of-service fabric statistics Output Fields *(Continued)*

Field Name	Field Description
`Drop statistics`	Fabric queue statistics for dropped traffic, including packets dropped because of internal error: • `Packets`—Dropped packet count for high-priority and low-priority queues. • `Bytes`—Dropped byte count for high-priority and low-priority queues. • `Pps`—Dropped packets-per-second count for high-priority and low-priority queues. • `bps`—Dropped bits-per-second count for high-priority and low-priority queues.
`Qdepth statistics`	Fabric queue depth statistics • `Average`—Average queue depth in bytes. • `Current`—Current queue depth in bytes. • `Peak`—Peak queue depth in bytes. • `Max`—Maximum queue depth in bytes.

Sample Output

show class-of-service fabric statistics

```
 user@host> show class-of-service fabric statistics
Destination FPC Index: 0, Source FPC Index: 0
 Total statistics:   High priority          Low priority
    Packets:                    0                     0
    Bytes  :                    0                     0
    Pps    :                    0                     0
    bps    :                    0                     0
 Tx statistics:      High priority          Low priority
    Packets:                    0                     0
    Bytes  :                    0                     0
    Pps    :                    0                     0
    bps    :                    0                     0
 Drop statistics:    High priority          Low priority
```

```
    Packets:                    0                     0
    Bytes  :                    0                     0
    Pps    :                    0                     0
    bps    :                    0                     0

Destination FPC Index: 0, Source FPC Index: 1
 Total statistics:   High priority          Low priority
    Packets:                    0                     0
    Bytes  :                    0                     0
    Pps    :                    0                     0
    bps    :                    0                     0
 Tx statistics:      High priority          Low priority
    Packets:                    0                     0
    Bytes  :                    0                     0
    Pps    :                    0                     0
    bps    :                    0                     0
 Drop statistics:    High priority          Low priority
    Packets:                    0                     0
    Bytes  :                    0                     0
...
```

show class-of-service fabric statistics detail

```
 user@host> show class-of-service fabric statistics detail
Destination FPC Index: 4, Destination Pfe Index: 0, Source FPC Index: 4, Source Pfe Index: 0
 Total statistics:   High priority          Low priority
    Packets:                28953                     0
    Bytes  :             14823936                     0
    Pps    :                   19                     0
    bps    :                81024                     0
 Tx statistics:      High priority          Low priority
    Packets:                28953                     0
    Bytes  :             14823936                     0
    Pps    :                   19                     0
    bps    :                81024                     0
 Drop statistics:    High priority          Low priority
    Packets:                    0                     0
    Bytes  :                    0                     0
    Pps    :                    0                     0
    bps    :                    0                     0
 Qdepth statistics:    High priority           Low priority
```

```
    Average:                    0                    0 b
    Current:                    0                    0 b
    Peak   :                    0                    0 b
    Max    :            1367343104            1367343104 b
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-class

IN THIS SECTION

- Syntax | 19019
- Description | 19019
- Required Privilege Level | 19019
- Output Fields | 19020
- Sample Output (Junos OS) | 19021
- Sample Output (Junos OS Evolved) | 19023
- Release Information | 19024

Syntax

```
show class-of-service forwarding-class
```

Description

Display information about forwarding classes, including the mapping of forwarding classes to queue numbers.

Required Privilege Level

view

Output Fields

The following table lists all possible output fields for the `show class-of-service forwarding-class` command. The fields that appear may vary depending on the platform and software release.

Table 364: show class-of-service forwarding-class Output Fields

Field Name	Field Description
`Forwarding class`	Name of the forwarding class.
`ID`	Forwarding class identifier. For DSCP-based PFC, the forwarding class ID is assigned from (and should be the same as) the configured PFC priority for the forwarding class. See *Configuring DSCP-based PFC for Layer 3 Untagged Traffic* for details.
`Queue`	CoS output queue mapped to the forwarding class.
`Restricted queue`	Restricted queue number.
`Fabric priority`	Fabric priority for the forwarding class, either `high` or `low`. Determines the priority of packets entering the switch fabric.
`Policing priority`	Layer 2 policing priority, either `premium` or `normal`.
`SPU priority`	Services Processing Unit (SPU) priority, either `high` or `low`.
`No-Loss`	Packet loss attribute to differentiate lossless forwarding classes from lossy forwarding classes: • Disabled—Lossless transport is not configured on the forwarding class (packet drop attribute is `drop`). • Enabled—Lossless transport is configured on the forwarding class (packet drop attribute is `no-loss`).

Table 364: show class-of-service forwarding-class Output Fields *(Continued)*

Field Name	Field Description
PFC Priority	For DSCP-based PFC, the explicitly configured PFC priority configured for the forwarding class. The DSCP value on which PFC is enabled maps to this priority, and this priority is used in PFC pause frames sent to the peer to request to pause traffic on the mapped DSCP value when the link becomes congested. The forwarding class ID is assigned from and should match this value in the output of this command. See *Configuring DSCP-based PFC for Layer 3 Untagged Traffic* for details.

Sample Output (Junos OS)

show class-of-service forwarding-class (ACX Series, EX Series, MX Series, PTX Series)

```
user@host> show class-of-service forwarding-class
Forwarding class        ID  Queue  Restricted queue  Fabric priority  Policing priority  SPU
priority
  best-effort            0   0          0             low                normal           low
  expedited-forwarding   1   1          1             low                normal           high
  assured-forwarding     2   2          2             low                normal           low
  network-control        3   3          3             low                normal           low
```

show class-of-service forwarding-class (SRX Series)

```
user@host> show class-of-service forwarding-class
Forwarding class        ID  Queue  Restricted queue  Fabric priority  Policing priority  SPU
priority
  best-effort            0   0          0             low                normal           low
  expedited-forwarding   1   1          1             low                normal           high
  assured-forwarding     2   2          2             low                normal           low
  network-control        3   3          3             low                normal           low
```

show class-of-service forwarding-class (EX4100, EX4300, and EX4400 Switch)

```
user@switch> show class-of-service forwarding-class
Forwarding class                      ID      Queue  Restricted queue  Fabric priority
Policing priority   SPU priority
  best-effort                          0       0          0             low
normal            low
  expedited-forwarding                 1       1          1             low
normal            low
  assured-forwarding                   2       2          2             low
normal            low
  network-control                      3       3          3             low
normal            low
  mcast-be                             8       8          8             low
normal            low
  mcast-ef                             9       9          9             low
normal            low
  mcast-af                            10      10         10             low
normal            low
  mcast-nc                            11      11         11             low
normal            low
```

show class-of-service forwarding-class (QFX Series)

```
user@switch> show class-of-service forwarding-class
Forwarding class                      ID      Queue  Policing priority  No-Loss   PFC priority
  best-effort                         0         0         normal        Disabled
  fcoe                                1         3         normal        Enabled
  no-loss                             2         4         normal        Enabled
  network-control                     3         7         normal        Disabled
  VIP-FC                              4         5         normal        Disabled
  mcast                               8         8         normal        Disabled
```

show class-of-service forwarding-class (QFX Series with DSCP-based PFC)

```
user@switch> show class-of-service forwarding-class
Forwarding class        ID      Queue  Policing priority  No-Loss     PFC priority
  best-effort           0         0        normal         Disabled
  fcoe                  1         3        normal         Enabled
```

```
  no-loss                2          4          normal          Enabled
  fc2                    3          2          normal          Enabled          3
  network-control        5          7          normal          Disabled
  fc1                    7          1          normal          Enabled          7
  mcast                  8          8          normal          Disabled
```

On switches that do not use different forwarding classes and output queues for unicast and multidestination (multicast, broadcast, destination lookup fail) traffic, there is no `mcast` forwarding class and there is no queue 8. (Switches that use different forwarding classes and output queues for unicast and multidestination traffic support 12 forwarding classes and output queues, of which four of each are dedicated to multidestination traffic. Switches that use the same forwarding classes and output queues for unicast and multidestination traffic support eight forwarding classes and eight output queues.)

Sample Output (Junos OS Evolved)

show class-of-service forwarding-class (ACX Series, QFX Series)

```
user@host> show class-of-service forwarding-class
Forwarding class        ID  Queue  Restricted queue  Fabric priority  Policing priority  SPU
priority
  best-effort            0   0          0             low                normal           low
  expedited-forwarding   1   1          1             low                normal           high
  assured-forwarding     2   2          2             low                normal           low
  network-control        3   3          3             low                normal           low
```

show class-of-service forwarding-class (PTX Series)

NOTE: You can see more advanced details about forwarding classes by logging into the Packet Forwarding Engine and entering the command `show class-of-service forwarding-classes`.

```
user@host> show class-of-service forwarding-class
Forwarding class                      ID       Queue      No-Loss
  assured-forwarding                  2          2        disabled
  best-effort                         0          0        disabled
  expedited-forwarding                1          1        disabled
  network-control                     3          3        disabled
```

Release Information

Command introduced in Junos OS Release 9.0.

show class-of-service forwarding-class-set

IN THIS SECTION

- Syntax | 19024
- Description | 19024
- Options | 19024
- Required Privilege Level | 19025
- Output Fields | 19025
- Sample Output | 19025
- Release Information | 19026

Syntax

```
show class-of-service forwarding-class-set
<forwarding-class-set-name>
```

Description

Display the forwarding classes associated with each forwarding class set.

Options

none	Display all forwarding class sets.
forwarding-class-set-name	(Optional) Display the forwarding classes associated with the specified forwarding class set.

Required Privilege Level

view

Output Fields

Table 365 on page 19025 describes the output fields for the `show class-of-service forwarding-class-set` command. Output fields are listed in the approximate order in which they appear.

Table 365: show class-of-service forwarding-class-set Output Fields

Field Name	Field Description
`Forwarding class set`	Name of the forwarding class set.
`Type`	Internal Junos OS type.
`Forwarding class set index`	Index of this forwarding class set.
`Forwarding class`	Name of a forwarding class.
`Index`	Index of this forwarding class.

Sample Output

show class-of-service forwarding-class-set

```
user@switch> show class-of-service forwarding-class-set
Forwarding class set: san_fcset, Type: normal-type, Forwarding class set index: 37839
  Forwarding class                      Index
  fcoe                                    1

Forwarding class set: lan_fcset, Type: normal-type, Forwarding class set index: 37840
  Forwarding class                      Index
  best-effort                             0

Forwarding class set: multicast_fcset, Type: normal-type, Forwarding class set index: 37841
```

```
Forwarding class                        Index
mcast                                     8
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Understanding CoS Forwarding Class Sets (Priority Groups)

Defining CoS Forwarding Class Sets

Example: Configuring Forwarding Class Sets

show class-of-service forwarding-table

IN THIS SECTION

- Syntax | **19026**
- Syntax (TX Matrix and TX Matrix Plus Router) | **19027**
- Description | **19027**
- Options | **19027**
- Required Privilege Level | **19027**
- Output Fields | **19027**
- Sample Output | **19028**
- Release Information | **19031**

Syntax

```
show class-of-service forwarding-table
```

Syntax (TX Matrix and TX Matrix Plus Router)

```
show class-of-service forwarding-table
<lcc number> | <sfc number>
```

Description

Display the entire class-of-service (CoS) configuration as it exists in the forwarding table. Executing this command is equivalent to executing all `show class-of-service forwarding-table` commands in succession.

Options

lcc ***number***

(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the forwarding table configuration for a specific T640 router (or line-card chassis) configured in a routing matrix. On a TX Matrix Plus router, display the forwarding table configuration for a specific router (or line-card chassis) configured in the routing matrix.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

sfc ***number***

(TX Matrix Plus routers only) (Optional) Display the forwarding table configuration for the TX Matrix Plus router. Replace *number* with `0`.

Required Privilege Level

view

Output Fields

See the output field descriptions for `show class-of-service forwarding-table` commands:

- "show class-of-service forwarding-table classifier" on page 19031

- "show class-of-service forwarding-table classifier mapping" on page 19033
- "show class-of-service forwarding-table drop-profile" on page 19035
- "show class-of-service forwarding-table fabric scheduler-map" on page 19038
- "show class-of-service forwarding-table rewrite-rule" on page 19040
- "show class-of-service forwarding-table rewrite-rule mapping" on page 19042
- "show class-of-service forwarding-table scheduler-map" on page 19044

Sample Output

show class-of-service forwarding-table

```
user@host> show class-of-service forwarding-table
Classifier table index: 9, # entries: 8, Table type: EXP
Entry #   Code point   Forwarding-class #   PLP
   0          000        0                 0
   1          001        0                 1
   2          010        1                 0
   3          011        1                 1
   4          100        2                 0
   5          101        2                 1
   6          110        3                 0
   7          111        3                 1

                       Table Index/
Interface      Index      Q num      Table type
sp-0/0/0.1001     66         11      IPv4 precedence
sp-0/0/0.2001     67         11      IPv4 precedence
sp-0/0/0.16383    68         11      IPv4 precedence
fe-0/0/0.0        69         11      IPv4 precedence

Interface: sp-0/0/0 (Index: 129, Map index: 2, Map type: FINAL,
 Num of queues: 2):
  Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
    Tx rate: 0 Kb (95%), Buffer size: 95 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
  Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
    Tx rate: 0 Kb (5%), Buffer size: 5 percent
```

```
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/0 (Index: 137, Map index: 2, Map type: FINAL,
 Num of queues: 2):
  Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
    Tx rate: 0 Kb (95%), Buffer size: 95 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
  Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
    Tx rate: 0 Kb (5%), Buffer size: 5 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
Interface: fe-0/0/1 (Index: 138, Map index: 2, Map type: FINAL,
 Num of queues: 2):
  Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
    Tx rate: 0 Kb (95%), Buffer size: 95 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
  Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
    Tx rate: 0 Kb (5%), Buffer size: 5 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

...

RED drop profile index: 1, # entries: 1
                             Drop
Entry      Fullness(%)   Probability(%)
    0              100              100
```

show class-of-service forwarding-table lcc (TX Matrix Plus Router)

```
user@host> show class-of-service forwarding-table lcc 0
lcc0-re0:
--------------------------------------------------------------------------

Classifier table index: 9, # entries: 64, Table type: IPv6 DSCP
Entry #   Code point   Forwarding-class #   PLP
   0        000000        0       0
   1        000001        0       0
```

```
  2        000010        0       0
  3        000011        0       0
  4        000100        0       0
  5        000101        0       0
  6        000110        0       0
  7        000111        0       0
  8        001000        0       0
  9        001001        0       0
 10        001010        0       0
 11        001011        0       0
 12        001100        0       0
 13        001101        0       0
 14        001110        0       0
 15        001111        0       0
 16        010000        0       0
 17        010001        0       0
 18        010010        0       0
 19        010011        0       0
 20        010100        0       0
 21        010101        0       0
 22        010110        0       0
 23        010111        0       0
 24        011000        0       0
 25        011001        0       0
 26        011010        0       0
 27        011011        0       0
 28        011100        0       0
 29        011101        0       0
 30        011110        0       0
 31        011111        0       0
 32        100000        0       0
 33        100001        0       0
 34        100010        0       0
 35        100011        0       0
 36        100100        0       0
 37        100101        0       0
 38        100110        0       0
 39        100111        0       0
 40        101000        0       0
 41        101001        0       0
 42        101010        0       0
 43        101011        0       0
 44        101100        0       0
```

```
  45        101101        0       0
  46        101110        0       0
...
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table classifier

IN THIS SECTION

- Syntax | 19031
- Description | 19031
- Options | 19031
- Required Privilege Level | 19032
- Output Fields | 19032
- Sample Output | 19032
- Release Information | 19033

Syntax

```
show class-of-service forwarding-table classifier
```

Description

Display the mapping of code point value to queue number and loss priority for each classifier as it exists in the forwarding table.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 366 on page 19032 describes the output fields for the `show class-of-service forwarding-table classifier` command. Output fields are listed in the approximate order in which they appear.

Table 366: show class-of-service forwarding-table classifier Output Fields

Field Name	Field Description
`Classifier table index`	Index of the classifier table.
`entries`	Total number of entries.
`Table type`	Type of code points in the table: `DSCP`, `EXP` (not on the QFX Series), `IEEE 802.1`, `IPv4 precedence` (not on the QFX Series), or `IPv6 DSCP`.
`Entry #`	Entry number.
`Code point`	Code point value used for classification.
`Forwarding-class #`	Forwarding class to which the code point is assigned.
`PLP`	Packet loss priority value set by classification. For most platforms, the value can be `0` or `1`. For some platforms, the value is `0`, `1`, `2`, or `3`. The value `0` represents low PLP. The value `1` represents `high` PLP. The value `2` represents medium-low PLP. The value `3` represents medium-high PLP.

Sample Output

show class-of-service forwarding-table classifier

```
user@host> show class-of-service forwarding-table classifier
Classifier table index: 62436, # entries: 64, Table type: DSCP
```

```
Entry #   Code point   Forwarding-class #   PLP
   0        000000       0                  0
   1        000001       0                  0
   2        000010       0                  0
   3        000011       0                  0
   4        000100       0                  0
   5        000101       0                  0
   6        000110       0                  0
   7        000111       0                  0
   8        001000       0                  0
   9        001001       0                  0
  10        001010       1                  1
  11        001011       0                  0
...
  60        111100       0                  0
  61        111101       0                  0
  62        111110       0                  0
  63        111111       0                  0
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table classifier mapping

IN THIS SECTION

- Syntax | 19034
- Description | 19034
- Options | 19034
- Required Privilege Level | 19034
- Output Fields | 19034
- Sample Output | 19035
- Release Information | 19035

Syntax

```
show class-of-service forwarding-table classifier mapping
```

Description

For each logical interface, display either the table index of the classifier for a given code point type or the queue number (if it is a fixed classification) in the forwarding table.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 367 on page 19034 describes the output fields for the `show class-of-service forwarding-table classifier mapping` command. Output fields are listed in the approximate order in which they appear.

Table 367: show class-of-service forwarding-table classifier mapping Output Fields

Field Name	Field Description
`Table index/ Q num`	If the table type is `Fixed`, the number of the queue to which the interface is mapped. For all other types, this value is the classifier index number.
`Interface`	Name of the logical interface. This field can also show the physical interface (QFX Series).
`Index`	Logical interface index.
`Table type`	Type of code points in the table: `DSCP`, `EXP` (not on the QFX Series), `Fixed`, `IEEE 802.1`, `IPv4 precedence` (not on the QFX Series),or `IPv6 DSCP`. `none` if `no-default` option set.

Sample Output

show class-of-service forwarding-table classifier mapping

```
user@host> show class-of-service forwarding-table classifier mapping
                        Table index/
Interface       Index     Q num      Table type
so-5/0/0.0         10     62436      DSCP
so-0/1/0.0         11     62436      DSCP
so-0/2/0.0         12         1      Fixed
so-0/2/1.0         13     62436      DSCP
so-0/2/1.0         13     62437      IEEE 802.1
so-0/2/2.0         14     62436      DSCP
so-0/2/2.0         14     62438      IPv4 precedence
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table drop-profile

IN THIS SECTION

- Syntax | 19036
- Description | 19036
- Options | 19036
- Required Privilege Level | 19036
- Output Fields | 19036
- Sample Output | 19037
- Release Information | 19038

Syntax

```
show class-of-service forwarding-table drop-profile
```

Description

Display the data points of all random early detection (RED) drop profiles as they exist in the forwarding table.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 368 on page 19036 describes the output fields for the `show class-of-service forwarding-table drop-profile` command. Output fields are listed in the approximate order in which they appear.

Table 368: show class-of-service forwarding-table drop-profile Output Fields

Field Name	Field Description
RED drop profile index	Index of this drop profile.
# entries	Number of entries in a particular RED drop profile index.
Entry	Drop profile entry number.
Fullness(%)	Percentage fullness of a queue.
Drop probability(%)	Drop probability at this fill level.

Sample Output

show class-of-service forwarding-table drop-profile

```
user@host> show class-of-service forwarding-table drop-profile
RED drop profile index: 4, # entries: 1
                             Drop
Entry      Fullness(%)   Probability(%)
    0              100              100

RED drop profile index: 8742, # entries: 3
                             Drop
Entry      Fullness(%)   Probability(%)
    0               10               10
    1               20               20
    2               30               30

RED drop profile index: 24627, # entries: 64
                             Drop
Entry      Fullness(%)   Probability(%)
    0                0                0
    1                1                1
    2                2                2
    3                4                4
 ...
   61               98               99
   62               99               99
   63              100              100

RED drop profile index: 25393, # entries: 64
                             Drop
Entry      Fullness(%)   Probability(%)
    0                0                0
    1                1                1
    2                2                2
    3                4                4
 ...
   61               98               98
   62               99               99
   63              100              100
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table fabric scheduler-map

IN THIS SECTION

- Syntax | 19038
- Description | 19038
- Options | 19038
- Additional Information | 19039
- Required Privilege Level | 19039
- Output Fields | 19039
- Sample Output | 19039
- Release Information | 19040

Syntax

```
show class-of-service forwarding-table fabric scheduler-map
```

Description

(M320 routers, MX Series routers, T Series routers and EX Series switches only) Display the scheduler map information as it exists in the forwarding table for switch fabric.

Options

This command has no options.

Additional Information

For information about how packet loss priority is assigned to packets, see *Managing Congestion by Setting Packet Loss Priority for Different Traffic Flows*.

Required Privilege Level

view

Output Fields

Table 369 on page 19039 describes the output fields for the`show class-of-service forwarding-table fabric scheduler-map` command. Output fields are listed in the approximate order in which they appear.

Table 369: show class-of-service forwarding-table fabric scheduler-map Output Fields

Field Name	Field Description
`Fabric priority`	Fabric traffic priority: `low` and `high`.
`Scheduler index`	Index of the scheduler applied to a fabric traffic priority.
`PLP high`	Drop profile index for high-packet-loss-priority (PLP) packets.
`PLP low`	Drop profile index for low-PLP packets.
`TCP PLP high`	Drop profile index for low-PLP and Transmission Control Protocol (TCP) packets.
`TCP PLP low`	Drop profile index for high-PLP and TCP packets.

Sample Output

show class-of-service forwarding-table fabric scheduler-map

```
user@host> show class-of-service forwarding-table fabric scheduler-map
Fabric priority: low
```

```
  Scheduler index: 60211
    PLP high: 44321, PLP low: 44321, TCP PLP high: 44321, TCP PLP low: 44321

Fabric priority: high
  Scheduler index: 60211
    PLP high: 44321, PLP low: 44321, TCP PLP high: 44321, TCP PLP low: 44321
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table rewrite-rule

IN THIS SECTION

- Syntax | 19040
- Description | 19040
- Options | 19041
- Required Privilege Level | 19041
- Output Fields | 19041
- Sample Output | 19042
- Release Information | 19042

Syntax

```
show class-of-service forwarding-table rewrite-rule
```

Description

Display mapping of queue number and loss priority to code point value for each rewrite rule as it exists in the forwarding table.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 370 on page 19041 describes the output fields for the `show class-of-service forwarding-table rewrite-rule` command. Output fields are listed in the approximate order in which they appear.

Table 370: show class-of-service forwarding-table rewrite-rule Output Fields

Field Name	Field Description
Rewrite table index	Index for this rewrite rule.
# entries	Number of entries in this rewrite rule.
Table type	Type of table: **DSCP**, **EXP** (not on the QFX Series), **EXP-PUSH-3** (not on the QFX Series), **IEEE 802.1**,**IPv4 precedence** (not on the QFX Series), **IPv6 DSCP**, or **Fixed**.
Q#	Queue number to which this entry is assigned.
Low bits	Code point value for low-priority loss profile.
State	State of this code point: **enabled**, **rewritten**, or **disabled**.
High bits	Code point value for high-priority loss profile.

Sample Output

show class-of-service forwarding-table rewrite-rule

```
user@host> show class-of-service forwarding-table rewrite-rule
Rewrite table index: 3753, # entries: 4, Table type: DSCP
Q#     Low bits  State      High bits  State
 0      000111  Enabled      001010   Enabled
 2      000000  Disabled     001100   Enabled
 1      101110  Enabled      110111   Enabled
 3      110000  Enabled      111000   Enabled
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table rewrite-rule mapping

IN THIS SECTION

- Syntax | 19042
- Description | 19043
- Options | 19043
- Required Privilege Level | 19043
- Output Fields | 19043
- Sample Output | 19043
- Release Information | 19044

Syntax

```
show class-of-service forwarding-table rewrite-rule mapping
```

Description

For each logical interface, display the table identifier of the rewrite rule map for each code point type.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 371 on page 19043 describes the output fields for the `show class-of-service forwarding-table rewrite-rule mapping` command. Output fields are listed in the approximate order in which they appear.

Table 371: show class-of-service forwarding-table rewrite-rule mapping Output Fields

Field Name	Field Description
Interface	Name of the logical interface. This field can also show the physical interface (QFX Series).
Index	Logical interface index.
Table index	Rewrite table index.
Type	Type of classifier: **DSCP**, **EXP** (not on the QFX Series), **EXP-PUSH-3** (not on the QFX Series), **EXP-SWAP-PUSH-2** (not on the QFX Series), **IEEE 802.1**, **IPv4 precedence** (not on the QFX Series), **IPv6 DSCP**, or **Fixed**.

Sample Output

show class-of-service forwarding-table rewrite-rule mapping

```
user@host> show class-of-service forwarding-table rewrite-rule mapping
Interface      Index   Table index  Type
```

```
so-5/0/0.0        10       3753  DSCP
so-0/1/0.0        11       3753  DSCP
so-0/2/0.0        12       3753  DSCP
so-0/2/1.0        13       3753  DSCP
so-0/2/2.0        14       3753  DSCP
so-0/2/3.0        15       3753  DSCP
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table scheduler-map

IN THIS SECTION

- Syntax | 19044
- Description | 19044
- Options | 19045
- Required Privilege Level | 19045
- Output Fields | 19045
- Sample Output | 19046
- Release Information | 19047

Syntax

```
show class-of-service forwarding-table scheduler-map
```

Description

For each physical interface, display the scheduler map information as it exists in the forwarding table.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 372 on page 19045 describes the output fields for the `show class-of-service forwarding-table scheduler-map` command. Output fields are listed in the approximate order in which they appear.

Table 372: show class-of-service forwarding-table scheduler-map Output Fields

Field Name	Field Description
Interface	Name of the physical interface.
Index	Physical interface index.
Map index	Scheduler map index.
Num of queues	Number of queues defined in this scheduler map.
Entry	Number of this entry in the scheduler map.
Scheduler index	Scheduler policy index.
Forwarding-class #	Forwarding class number to which this entry is applied.
Tx rate	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword **remainder**, which indicates that the scheduler receives the remaining bandwidth of the interface.

Table 372: show class-of-service forwarding-table scheduler-map Output Fields *(Continued)*

Field Name	Field Description
Max buffer delay	Amount of transmit delay (in milliseconds) or buffer size of the queue. This amount is a percentage of the total interface buffer allocation or the keyword **remainder**, which indicates that the buffer is sized according to what remains after other scheduler buffer allocations.
Priority	• **high**—Queue priority is high. • **low**—Queue priority is low.
PLP high	Drop profile index for a high packet loss priority profile.
PLP low	Drop profile index for a low packet loss priority profile.
PLP medium-high	Drop profile index for a medium-high packet loss priority profile.
PLP medium-low	Drop profile index for a medium-low packet loss priority profile.
TCP PLP high	Drop profile index for a high TCP packet loss priority profile.
TCP PLP low	Drop profile index for a low TCP packet loss priority profile.
Policy is exact	If this line appears in the output, exact rate limiting is enabled. Otherwise, no rate limiting is enabled.

Sample Output

show class-of-service forwarding-table scheduler-map

```
user@host> show class-of-service forwarding-table scheduler-map
Interface: so-5/0/0 (Index: 9, Map index: 17638, Num of queues: 2):
  Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
    Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
    Priority low
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low:8742
```

```
    Policy is exact
  Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
    Priority high
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

Interface: at-6/1/0 (Index: 10, Map index: 17638, Num of queues: 2):
  Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
    Priority high
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
  Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
    Priority low
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table slice

IN THIS SECTION

- Syntax | 19048
- Description | 19048
- Options | 19048
- Required Privilege Level | 19048
- Output Fields | 19048
- Sample Output | 19049
- Release Information | 19049

Syntax

```
show class-of-service forwarding-table slice
```

Description

Show slices associated with the forwarding table.

Options

- mapping - show mapping of traffic control profiles to slices

Required Privilege Level

view

Output Fields

Table 373 on page 19048 lists the output fields for the `show class-of-service forwarding-table slice` command.

Table 373: show class-of-service forwarding-table slice Output Fields

Field Name	Field Description
Slice	The name of the slice
Index	The index of the slice

Sample Output

show class-of-service forwarding-table slice

```
user@host> show class-of-service forwarding-table slice
```

```
user@host# run show class-of-service forwarding-table slice
Apr 06 01:09:35
  Slice: s1, Index: 1

  Slice: s2, Index: 2
```

Release Information

Junos 22.2R1

show class-of-service forwarding-table slice mapping

IN THIS SECTION

- Syntax | 19049
- Description | 19050
- Required Privilege Level | 19050
- Output Fields | 19050
- Sample Output | 19050
- Release Information | 19051

Syntax

```
show class-of-service forwarding-table slice mapping
```

Description

Show mapping of traffic control profiles to slices.

Required Privilege Level

view

Output Fields

Table 374 on page 19050 lists the output fields for the `show class-of-service forwarding-table slice mapping` command.

Table 374: show class-of-service forwarding-table slice mapping Output Fields

Field Name	Field Description
Index	Slice index
Traffic control profile index	The index of the traffic control profile attached to the slice.

Sample Output

show class-of-service forwarding-table slice mapping

```
user@host> show class-of-service forwarding-table slice mapping
```

```
user@host# show class-of-service forwarding-table slice mapping
Interface: xe-4/0/5, (Index: 1, Map index: 2, Map type: FINAL, Num of queues: 2):
Traffic control profile index: 33290
Traffic control profile instance index: 0
Shaping rate: 1000000000, Shaping rate burst: 4294967295
Guaranteed rate: 0, Guaranteed rate burst: 0
Overhead bytes: 0, Adjust minimum shaping rate: 0, Traffic control profile flags: 0x1
Traffic control profile level: 1
  Entry 0 (Scheduler index: 0, Forwarding class #: 0, Queue #: 0):
    Tx rate: 0 Kb (95%), Buffer size: 95 percent
```

```
    Priority low
    Excess rate: proportion 0 Excess priority: low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
  Entry 1 (Scheduler index: 0, Forwarding class #: 3, Queue #: 3):
    Tx rate: 0 Kb (5%), Buffer size: 5 percent
    Priority low
    Excess rate: proportion 0 Excess priority: low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
```

Release Information

Junos 22.2R1

show class-of-service forwarding-table traffic-class-map

IN THIS SECTION

- Syntax | 19051
- Description | 19051
- Options | 19052
- Required Privilege Level | 19052
- Output Fields | 19052
- Sample Output | 19053
- Sample Output | 19053
- Release Information | 19054

Syntax

```
show class-of-service forwarding-table traffic-class-map <mapping>
```

Description

Display the mapping of code point value to the traffic class map as it exists in the forwarding table.

Options

`mapping`—Display the mapping of interfaces to traffic class maps.

Required Privilege Level

view

Output Fields

Table 375 on page 19052 describes the output fields for the `show class-of-service forwarding-table traffic-class-map` command. Output fields are listed in the approximate order in which they appear.

Table 375: show class-of-service forwarding-table traffic-class-map Output Fields

Field Name	Field Description
`Traffic-class-map table index`	Index of the traffic class map table.
`entries`	Total number of entries.
`Table type`	Type of code points in the table: `DSCP`, `EXP` , `IEEE 802.1`, `IEEE 802.1ad`, or `INET-precedence`
`Entry #`	Entry number.
`Code point`	Code point value used for classification.
`Traffic-class`	Traffic class to which the code point is assigned.

Table 376 on page 19053 describes the output fields for the `show class-of-service forwarding-table traffic-class-map mapping` command. Output fields are listed in the approximate order in which they appear.

Table 376: show class-of-service forwarding-table traffic-class-map mapping Output Fields

Field Name	Field Description
Interface	Interface to which the traffic class map is assigned.
Index	Internal index of the traffic class map.
Table Index	Index of the traffic class map table.
Table type	Type of code points in the table: DSCP, EXP , IEEE 802.1, IEEE 802.1ad, or INET-precedence

Sample Output

show class-of-service forwarding-table traffic-class-map

```
user@host> show class-of-service forwarding-table traffic-class-map
Traffic-class-map table index: 44231, # entries: 6, Table type: INET-Precedence
        Entry #   Code point   Traffic-class
           0           000      real-time
           1           001      real-time
           2           010      network-control
           3           011      network-control
           4           100      best-effort
           5           101      best-effort
```

Sample Output

show class-of-service forwarding-table traffic-class-map mapping

```
user@host> show class-of-service forwarding-table traffic-class-map mapping
Interface      Index    Table Index   Table type
        xe-4/0/0         210        44231     INET-Precedence
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

- traffic-class-map | 30003
- *Managing Ingress Oversubscription at the PFE*
- *Configuring Traffic Class Maps to Manage Ingress Oversubscription*
- *Example: Configuring Traffic Class Maps*
- show class-of-service traffic-class-map | 19142

show class-of-service fragmentation-map

IN THIS SECTION

- Syntax | 19054
- Description | 19054
- Options | 19055
- Required Privilege Level | 19055
- Output Fields | 19055
- Sample Output | 19055
- Release Information | 19056

Syntax

```
show class-of-service fragmentation-map
```

Description

For Multiservices and Services PIC link services IQ interfaces (`lsq`) only, display fragmentation properties for specific forwarding classes.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 377 on page 19055 describes the output fields for the `show class-of-service fragmentation-map` command. Output fields are listed in the approximate order in which they appear.

Table 377: show class-of-service fragmentation-map Output Fields

Field Name	Field Description
`Fragmentation map`	Name of the class of service (CoS) fragmentation map.
`Index`	Index number of the CoS fragmentation map.
`Forwarding class`	Name of the associated forwarding class.
`Fragmentation threshold`	Maximum size of each multilink fragment.
`No Fragmentation`	Packets of this class are not fragmented.
`Multilink Class`	For multilink multiclass PPP only, the multilink class number corresponding to the forwarding class.

Sample Output

show class-of-service fragmentation-map

```
user@host> show class-of-service fragmentation-map
  Fragmentation map: fragmap2, Index: 19801
    Forwarding class: fcDefault
```

```
    No Fragmentation

     Forwarding class: fcCopper
    Fragmentation threshold: 64, Multilink Class: 1

    Forwarding class: fcSilver
    Fragmentation threshold: 100, Multilink Class: 0

    Forwarding class: fcCritical
    Fragmentation threshold: 64, Multilink Class: 0

  Fragmentation map: fragmap, Index: 23147
    Forwarding class: fcDefault
    No Fragmentation

    Forwarding class: fcSilver
    Fragmentation threshold: 100

    Forwarding class: fcCritical
    Fragmentation threshold: 100
```

Release Information

Command introduced in Junos OS Release 7.5.

show class-of-service interface

IN THIS SECTION

- Syntax | 19057
- Description | 19057
- Options | 19057
- Required Privilege Level | 19058
- Output Fields | 19058
- Sample Output | 19078

Syntax

```
show class-of-service interface <comprehensive | detail> <interface-name>
```

Description

Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.

> **NOTE**: This topic lists all possible options and output fields for the `show class-of-service interface` command. Options and output fields can vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

> **NOTE**: On routing platforms with dual Routing Engines, running this command on the backup Routing Engine, with or without any of the available options, is not supported and produces the following error message:
>
> `error: the class-of-service subsystem is not running`

Options

none Display CoS associations for all physical and logical interfaces.

comprehensive (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.

detail (Optional) Display QoS and CoS information based on the interface.

If the `interface` *interface-name* is a physical interface, the output includes:

- Brief QoS information about the physical interface
- Brief QoS information about the logical interface
- CoS information about the physical interface
- Brief information about filters or policers of the logical interface
- Brief CoS information about the logical interface

If the `interface` *interface-name* is a logical interface, the output includes:

- Brief QoS information about the logical interface
- Information about filters or policers for the logical interface
- CoS information about the logical interface

interface-name (Optional) Display class-of-service (CoS) associations for the specified interface.

none Display CoS associations for all physical and logical interfaces.

Required Privilege Level

view

Output Fields

Table 378 on page 19058 describes the output fields for the `show class-of-service interface` command. Output fields are listed in the approximate order in which they appear.

Table 378: show class-of-service interface Output Fields

Field Name	Field Description
`Physical interface`	Name of a physical interface.
`Index`	Index of this interface or the internal index of this object. (Enhanced subscriber management) Index values for dynamic CoS traffic control profiles and dynamic scheduler maps are larger for enhanced subscriber management than they are for legacy subscriber management.
`Dedicated Queues`	Status of dedicated queues configured on an interface. (Enhanced subscriber management) This field is not displayed for enhanced subscriber management.
`Maximum usable queues`	Number of queues you can configure on the interface.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Maximum usable queues`	Maximum number of queues you can use.
`Total non-default queues created`	Number of queues created in addition to the default queues. (Enhanced subscriber management) This field is not displayed for enhanced subscriber management.
`Rewrite Input IEEE Code-point`	IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value.
`Shaping rate`	Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the `Shaping rate` field is displayed for either the physical interface or the logical interface.
`Scheduler map`	Name of the output scheduler map associated with this interface. (Enhanced subscriber management) The name of the dynamic scheduler map object is associated with a generated UID (for example, `SMAP-1_UID1002`) instead of with a subscriber interface.
`Scheduler map forwarding class sets`	Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.
`Input shaping rate`	Maximum transmission rate on the input interface.
`Input scheduler map`	Name of the input scheduler map associated with this interface.
`Chassis scheduler map`	Name of the scheduler map associated with the packet forwarding component queues.
`Rewrite`	Name and type of the rewrite rules associated with this interface.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Traffic-control-profile`	Name of the associated traffic control profile. (Enhanced subscriber management) The name of the dynamic traffic control profile object is associated with a generated UID (for example, `TC_PROF_100_199_SERIES_UID1006`) instead of with a subscriber interface.
`Classifier`	Name and type of classifiers associated with this interface.
`Forwarding-class-map`	Name of the forwarding map associated with this interface.
`Congestion-notification`	Congestion notification state, `enabled` or `disabled`.
Dedicated Buffer Profile	Name of the dedicated buffer profile associated with the interface.
`Monitoring Profile Name`	Name of the monitoring profile defined to monitor the peak queue length for virtual output queues (VOQs) for the interface.
`Logical interface`	Name of a logical interface.
`Object`	Category of an object: `Classifier`, `Fragmentation-map` (for LSQ interfaces only), `Scheduler-map`, `Rewrite`, `Translation Table`, or `traffic-class-map`.
`Name`	Name of an object.
`Type`	Type of an object: `dscp`, `dscp-ipv6`, `exp`, `ieee-802.1`, `ip`, `inet-precedence`, or `ieee-802.1ad`.
`Link-level type`	Encapsulation on the physical interface.
`MTU`	MTU size on the physical interface.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Speed`	Speed at which the interface is running.
`Loopback`	Whether loopback is enabled and the type of loopback.
`Source filtering`	Whether source filtering is enabled or disabled.
`Flow control`	Whether flow control is enabled or disabled.
`Auto-negotiation`	Whether autonegotiation is enabled or disabled.
`Remote-fault`	Remote fault status. • `Online`—Autonegotiation is manually configured as online. • `Offline`—Autonegotiation is manually configured as offline.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Device flags`	The `Device flags` field provides information about the physical device and displays one or more of the following values: • `Down`—Device has been administratively disabled. • `Hear-Own-Xmit`—Device receives its own transmissions. • `Link-Layer-Down`—The link-layer protocol has failed to connect with the remote endpoint. • `Loopback`—Device is in physical loopback. • `Loop-Detected`—The link layer has received frames that it sent, thereby detecting a physical loopback. • `No-Carrier`—On media that support carrier recognition, no carrier is currently detected. • `No-Multicast`—Device does not support multicast traffic. • `Present`—Device is physically present and recognized. • `Promiscuous`—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media. • `Quench`—Transmission on the device is quenched because the output buffer is overflowing. • `Recv-All-Multicasts`—Device is in multicast promiscuous mode and therefore provides no multicast filtering. • `Running`—Device is active and enabled.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

`Interface flags`	The `Interface flags` field provides information about the physical interface and displays one or more of the following values:
	• `Admin-Test`—Interface is in test mode and some sanity checking, such as loop detection, is disabled.
	• `Disabled`—Interface is administratively disabled.
	• `Down`—A hardware failure has occurred.
	• `Hardware-Down`—Interface is nonfunctional or incorrectly connected.
	• `Link-Layer-Down`—Interface keepalives have indicated that the link is incomplete.
	• `No-Multicast`—Interface does not support multicast traffic.
	• `No-receive No-transmit`—Passive monitor mode is configured on the interface.
	• `Point-To-Point`—Interface is point-to-point.
	• `Pop all MPLS labels from packets of depth`—MPLS labels are removed as packets arrive on an interface that has the `pop-all-labels` statement configured. The depth value can be one of the following:
	  • `1`—Takes effect for incoming packets with one label only.
	  • `2`—Takes effect for incoming packets with two labels only.
	  • `[ 1 2 ]`—Takes effect for incoming packets with either one or two labels.
	• `Promiscuous`—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.
	• `Recv-All-Multicasts`—Interface is in multicast promiscuous mode and provides no multicast filtering.
	• `SNMP-Traps`—SNMP trap notifications are enabled.
	• `Up`—Interface is enabled and operational.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Flags`	The `Logical interface flags` field provides information about the logical interface and displays one or more of the following values: • `ACFC Encapsulation`—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer). • `Device-down`—Device has been administratively disabled. • `Disabled`—Interface is administratively disabled. • `Down`—A hardware failure has occurred. • `Clear-DF-Bit`—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit. • `Hardware-Down`—Interface protocol initialization failed to complete successfully. • `PFC`—Protocol field compression is enabled for the PPP session. • `Point-To-Point`—Interface is point-to-point. • `SNMP-Traps`—SNMP trap notifications are enabled. • `Up`—Interface is enabled and operational.
`Encapsulation`	Encapsulation on the logical interface.
`Admin`	Administrative state of the interface (`Up` or `Down`)
`Link`	Status of physical link (`Up` or `Down`).
`Proto`	Protocol configured on the interface.
`Input Filter`	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
`Output Filter`	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Link flags`	Provides information about the physical link and displays one or more of the following values: • `ACFC`—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option. • `Give-Up`—Link protocol does not continue connection attempts after repeated failures. • `Loose-LCP`—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational. • `Loose-LMI`—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational. • `Loose-NCP`—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational. • `Keepalives`—Link protocol keepalives are enabled. • `No-Keepalives`—Link protocol keepalives are disabled. • `PFC`—Protocol field compression is configured. The PPP session negotiates the PFC option.
`Hold-times`	Current interface hold-time up and hold-time down, in milliseconds.
`CoS queues`	Number of CoS queues configured.
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second`*`:`*`timezone`* `(`*`hour`*`:`*`minute`*`:`*`second`* `ago)`. For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Statistics last cleared`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.
Exclude Overhead Bytes	Exclude the counting of overhead bytes from aggregate queue statistics. • `Disabled`—Default configuration. Includes the counting of overhead bytes in aggregate queue statistics. • `Enabled`—Excludes the counting of overhead bytes from aggregate queue statistics for just the physical interface. • `Enabled for hierarchy`—Excludes the counting of overhead bytes from aggregate queue statistics for the physical interface as well as all child interfaces, including logical interfaces and interface sets.
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Input errors`	Input errors on the interface. The labels are explained in the following list: • `Errors`—Sum of the incoming frame terminations and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Giants`—Number of frames received that are larger than the giant threshold. • `Bucket Drops`—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • `L3 incompletes`—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` statement. • `L2 channel errors`—Number of times the software did not find a valid logical interface for an incoming frame. • `L2 mismatch timeouts`—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • `HS link CRC errors`—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • `HS link FIFO overflows`—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Output errors`	Output errors on the interface. The labels are explained in the following list: • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • `Errors`—Sum of the outgoing frame terminations and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Aged packets`—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • `HS link FIFO underflows`—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces. • `MTU errors`—Number of packets whose size exceeds the MTU of the interface.
`Egress queues`	Total number of egress Maximum usable queues on the specified interface.
`Queue counters`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.
`SONET alarms` `SONET defects`	SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: `SONET PHY`, `SONET section`, `SONET line`, and `SONET path`.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
SONET PHY	Counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. The SONET PHY field has the following subfields: • PLL Lock—Phase-locked loop • PHY Light—Loss of optical signal
SONET section	Counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. The SONET section field has the following subfields: • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOS—Loss of signal • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section)

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`SONET line`	Active alarms and defects, plus counts of specific SONET errors with detailed information. • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. A state other than `OK` indicates a problem. The `SONET line` field has the following subfields: • `BIP-B2`—Bit interleaved parity for SONET line overhead • `REI-L`—Remote error indication (near-end line) • `RDI-L`—Remote defect indication (near-end line) • `AIS-L`—Alarm indication signal (near-end line) • `BERR-SF`—Bit error rate fault (signal failure) • `BERR-SD`—Bit error rate defect (signal degradation) • `ES-L`—Errored seconds (near-end line) • `SES-L`—Severely errored seconds (near-end line) • `UAS-L`—Unavailable seconds (near-end line) • `ES-LFE`—Errored seconds (far-end line) • `SES-LFE`—Severely errored seconds (far-end line) • `UAS-LFE`—Unavailable seconds (far-end line)

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`SONET path`	Active alarms and defects, plus counts of specific SONET errors with detailed information. • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. A state other than `OK` indicates a problem. The `SONET path` field has the following subfields: • `BIP-B3`—Bit interleaved parity for SONET section overhead • `REI-P`—Remote error indication • `LOP-P`—Loss of pointer (path) • `AIS-P`—Path alarm indication signal • `RDI-P`—Path remote defect indication • `UNEQ-P`—Path unequipped • `PLM-P`—Path payload (signal) label mismatch • `ES-P`—Errored seconds (near-end STS path) • `SES-P`—Severely errored seconds (near-end STS path) • `UAS-P`—Unavailable seconds (near-end STS path) • `ES-PFE`—Errored seconds (far-end STS path) • `SES-PFE`—Severely errored seconds (far-end STS path) • `UAS-PFE`—Unavailable seconds (far-end STS path)

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

`Received SONET overhead`
 	
`Transmitted SONET overhead`
	Values of the received and transmitted SONET overhead:
	
	• `C2`—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P.
	
	• `F1`—Section user channel byte. This byte is set aside for the purposes of users.
	
	• `K1` and `K2`—These bytes are allocated for APS signaling for the protection of the multiplex section.
	
	• `J0`—Section trace. This byte is defined for STS-1 number 1 of an STS-*N* signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter.
	
	• `S1`—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-*N* signal.
	
	• `Z3` and `Z4`—Allocated for future use.

`Received path trace`
 	
`Transmitted path trace`
	SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.

`HDLC configuration`	Information about the HDLC configuration. • `Policing bucket`—Configured state of the receiving policer. • `Shaping bucket`—Configured state of the transmitting shaper. • `Giant threshold`—Giant threshold programmed into the hardware. • `Runt threshold`—Runt threshold programmed into the hardware.
`Packet Forwarding Engine configuration`	Information about the configuration of the Packet Forwarding Engine: • `Destination slot`—FPC slot number. • `PLP byte`—Packet Level Protocol byte.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`CoS information`	Information about the CoS queue for the physical interface. • `CoS transmit queue`—Queue number and its associated user-configured forwarding class name. • `Bandwidth %`—Percentage of bandwidth allocated to the queue. • `Bandwidth bps`—Bandwidth allocated to the queue (in bps). • `Buffer %`—Percentage of buffer space allocated to the queue. • `Buffer usec`—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • `Priority`—Queue priority: `low` or `high`. • `Limit`—Displayed if rate limiting is configured for the queue. Possible values are `none` and `exact`. If `exact` is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If `none` is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.
`Forwarding classes`	Total number of forwarding classes supported on the specified interface.
`Egress queues`	Total number of egress Maximum usable queues on the specified interface.
`Queue`	Queue number.
`Forwarding classes`	Forwarding class name.
`Queued Packets`	Number of packets queued to this queue.
`Queued Bytes`	Number of bytes queued to this queue. The byte counts vary by PIC type.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Transmitted Packets`	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the `Packet Forwarding Engine Chassis Queues` field) shows the prefragmentation values.
`Transmitted Bytes`	Number of bytes transmitted by this queue. The byte counts vary by PIC type.
`Tail-dropped packets`	Number of packets dropped because of tail drop.

`RED-dropped packets`	Number of packets dropped because of random early detection (RED).
	• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories:
	  • `Low, non-TCP`—Number of low-loss priority non-TCP packets dropped because of RED.
	  • `Low, TCP`—Number of low-loss priority TCP packets dropped because of RED.
	  • `High, non-TCP`—Number of high-loss priority non-TCP packets dropped because of RED.
	  • `High, TCP`—Number of high-loss priority TCP packets dropped because of RED.
	• (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories:
	  • `Low`—Number of low-loss priority packets dropped because of RED.
	  • `Medium-low`—Number of medium-low loss priority packets dropped because of RED.
	  • `Medium-high`—Number of medium-high loss priority packets dropped because of RED.
	  • `High`—Number of high-loss priority packets dropped because of RED.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

`RED-dropped bytes`	Number of bytes dropped because of RED. The byte counts vary by PIC type.
	• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories:
	  • `Low, non-TCP`—Number of low-loss priority non-TCP bytes dropped because of RED.
	  • `Low, TCP`—Number of low-loss priority TCP bytes dropped because of RED.
	  • `High, non-TCP`—Number of high-loss priority non-TCP bytes dropped because of RED.
	  • `High, TCP`—Number of high-loss priority TCP bytes dropped because of RED.

`Transmit rate`	Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.
`Rate Limit`	Rate limiting configuration of the queue. Possible values are : • `None`—No rate limit. • `exact`—Queue transmits at the configured rate.
`Buffer size`	Delay buffer size in the queue.
`Priority`	Scheduling priority configured as `low` or `high`.
`Excess Priority`	Priority of the excess bandwidth traffic on a scheduler: `low`, `medium-low`, `medium-high`, `high`, or `none`.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Drop profiles`	Display the assignment of drop profiles. • `Loss priority`—Packet loss priority for drop profile assignment. • `Protocol`—Transport protocol for drop profile assignment. • `Index`—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • `Name`—Name of the drop profile. • `Type`—Type of the drop profile: `discrete` or `interpolated`. • `Fill Level`—Percentage fullness of a queue. • `Drop probability`—Drop probability at this fill level.
`Excess Priority`	Priority of the excess bandwidth traffic on a scheduler.
`Drop profiles`	Display the assignment of drop profiles. • `Loss priority`—Packet loss priority for drop profile assignment. • `Protocol`—Transport protocol for drop profile assignment. • `Index`—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • `Name`—Name of the drop profile. • `Type`—Type of the drop profile: `discrete` or `interpolated`. • `Fill Level`—Percentage fullness of a queue. • `Drop probability`—Drop probability at this fill level.

Table 378: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

`Adjustment information`	Display the assignment of shaping-rate adjustments on a scheduler node or queue.
	• `Adjusting application`—Application that is performing the shaping-rate adjustment.
	  • The adjusting application can appear as `ancp LS-0`, which is the Junos OS Access Node Control Profile process (`ancpd`) that performs shaping-rate adjustments on schedule nodes.
	  • The adjusting application can appear as `DHCP`, which adjusts the shaping-rate and overhead-accounting class-of-service attributes based on DSL Forum VSA conveyed in DHCP option 82, suboption 9 (Vendor Specific Information). The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).
	  • The adjusting application can also appear as `pppoe`, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).
	• `Adjustment type`—Type of adjustment: `absolute` or `delta`.
	• `Configured shaping rate`—Shaping rate configured for the scheduler node or queue.
	• `Adjustment value`—Value of adjusted shaping rate.
	• `Adjustment target`—Level of shaping-rate adjustment performed: `node` or `queue`.
	• `Adjustment overhead-accounting mode`—Configured shaping mode: `frame` or `cell`.
	• `Adjustment overhead bytes`—Number of bytes that the ANCP agent adds to or subtracts from the actual downstream frame overhead before reporting the adjusted values to CoS.
	• `Adjustment target`—Level of shaping-rate adjustment performed: `node` or `queue`.
	• `Adjustment multicast index`—

Sample Output

show class-of-service interface (Physical)

```
user@host> show class-of-service interface et-1/0/4
Physical interface: et-1/0/4, Index: 1098
Maximum usable queues: 8, Queues in use: 4
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: default, Index: 0
  Congestion-notification: Disabled
  Dedicated Buffer Profile: dbp1
  Monitoring Profile Name: XYZ

  Logical interface: et-1/0/4.16386, Index: 1057
```

show class-of-service interface (aggregated Ethernet)

```
user@host> show class-of-service ae0
Physical interface: ae0, Index: 7860
Maximum usable queues: 8, Queues in use: 4
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: default, Index: 0
  Congestion-notification: Disabled
  Monitoring Profile Name: mp1 [shared]
```

show class-of-service interface (Logical)

```
user@host> show class-of-service interface so-0/2/3.0
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
    Shaping rate: 32000
    Object                   Name                   Type             Index
    Scheduler-map            <default>                                27
    Rewrite                  exp-default            exp               21
    Classifier               exp-default            exp               5
```

```
    Classifier               ipprec-compatibility   ip                8
    Forwarding-class-map     exp-default            exp               5
```

show class-of-service interface (ANCP)

```
user@host> show class-of-service interface pp0.1073741842
  Logical interface: pp0.1073741842, Index: 341
Object                  Name                   Type                    Index
Traffic-control-profile TCP-CVLAN              Output                  12408
Classifier              dscp-ipv6-compatibility dscp-ipv6                  9
Classifier              ipprec-compatibility   ip                         13

    Adjusting application: ancp LS-0
      Adjustment type: absolute
      Configured shaping rate: 4000000
      Adjustment value: 11228000
      Adjustment overhead-accounting mode: Frame Mode
      Adjustment overhead bytes: 50
      Adjustment target: node
```

show class-of-service interface (PPPoE Interface)

```
user@host> show class-of-service interface pp0.1
Logical interface: pp0.1, Index: 85
    Object                  Name                   Type           Index
    Traffic-control-profile tcp-pppoe.o.pp0.1      Output         2726446535
    Classifier              ipprec-compatibility   ip             13

    Adjusting application: PPPoE
      Adjustment type: absolute
      Adjustment value: 5000000
      Adjustment overhead-accounting mode: cell
      Adjustment target: node
```

show class-of-service interface (DHCP Interface)

```
user@host> show class-of-service interface demux0.1
Logical interface: pp0.1, Index: 85
```

```
    Object                  Name                    Type           Index
    Traffic-control-profile tcp-dhcp.o.demux0.1     Output         2726446535
    Classifier              ipprec-compatibility    ip             13

    Adjusting application: DHCP
      Adjustment type: absolute
      Adjustment value: 5000000
      Adjustment overhead-accounting mode: cell
      Adjustment target: node
```

show class-of-service interface detail

```
user@host> show class-of-service interface ge-0/3/0 detail

Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, Loopback: Disabled, Source filtering:
Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000

  Physical interface: ge-0/3/0, Index: 138
  Maximum usable queues: 4, Queues in use: 5
  Shaping rate: 50000 bps
  Scheduler map: interface-schedular-map, Index: 58414
  Input shaping rate: 10000 bps
  Input scheduler map: schedular-map, Index: 15103
  Chassis scheduler map: <default-chassis>, Index: 4
  Congestion-notification: Disabled

  Logical interface ge-0/3/0.0
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    inet
    mpls
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.0      up    up   inet
                           mpls
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.0      up    up   inet
                           mpls

  Logical interface: ge-0/3/0.0, Index: 68
```

```
    Object                  Name                    Type                Index
    Rewrite                 exp-default             exp (mpls-any)        33
    Classifier              exp-default             exp                   10
    Classifier              ipprec-compatibility    ip                    13


  Logical interface ge-0/3/0.1
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    inet
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.1      up    up   inet
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.1      up    up   inet

  Logical interface: ge-0/3/0.1, Index: 69
    Object                  Name                    Type             Index
    Classifier              ipprec-compatibility    ip                13
```

show class-of-service interface comprehensive

```
user@host> show class-of-service interface ge-0/3/0 comprehensive
Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 601, Generation: 141
  Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Schedulers     : 256
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:14:f6:f4:b4:5d, Hardware address: 00:14:f6:f4:b4:5d
  Last flapped   : 2010-09-07 06:35:22 PDT (15:14:42 ago)
  Statistics last cleared: Never  Exclude Overhead Bytes: Disabled
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 total statistics:
    Input  bytes  :                   0
```

```
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Drop   bytes  :                    0                    0 bps
   Drop   packets:                    0                    0 pps
  Label-switched interface (LSI) traffic statistics:
   Input  bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 4 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 af3                           0                    0                    0
    1 af2                           0                    0                    0
    2 ef2                           0                    0                    0
    3 ef1                           0                    0                    0
  Egress queues: 4 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 af3                           0                    0                    0
    1 af2                           0                    0                    0
    2 ef2                           0                    0                    0
    3 ef1                           0                    0                    0
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
```

```
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                      0
    Output packet pad count                                  0
    Output packet error count                                0
    CAM destination filters: 0, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault: OK
    Local resolution:
      Flow control: Symmetric, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 0
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    2 ef2                    39          19500     0            120     high    none
    Direction : Input
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 af3                    30           3000    45              0      low    none

Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 601
Forwarding classes: 16 supported, 5 in use
Ingress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                      0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                     0 pps
```

```
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: af2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Forwarding classes: 16 supported, 5 in use
Egress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
```

```
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: af2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
```

```
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available
Queue: 1, Forwarding classes: af2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets              :                108546                    0 pps
    Bytes                :              12754752                  376 bps
  Transmitted:
    Packets              :                108546                    0 pps
    Bytes                :              12754752                  376 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available

Physical interface: ge-0/3/0, Index: 138
Maximum usable queues: 4, Queues in use: 5
  Shaping rate: 50000 bps
```

```
Scheduler map: interface-schedular-map, Index: 58414

  Scheduler: ef2, Forwarding class: ef2, Index: 39155
    Transmit rate: 39 percent, Rate Limit: none, Buffer size: 120 us, Buffer Limit: none,
Priority: high
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
  Input shaping rate: 10000 bps
  Input scheduler map: schedular-map

Scheduler map: schedular-map, Index: 15103

  Scheduler: af3, Forwarding class: af3, Index: 35058
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 45 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         40582    green
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any         18928    yellow
Drop profile: green, Type: discrete, Index: 40582
  Fill level    Drop probability
          50                   0
         100                 100
```

```
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: yellow, Type: discrete, Index: 18928
  Fill level    Drop probability
          50                   0
         100                 100
  Chassis scheduler map: < default-drop-profile>
Scheduler map: < default-drop-profile>, Index: 4

  Scheduler: < default-drop-profile>, Forwarding class: af3, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none,
Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100

  Scheduler: < default-drop-profile>, Forwarding class: af2, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none,
Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
```

```
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100

  Scheduler: < default-drop-profile>, Forwarding class: ef2, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none,
Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100

  Scheduler: < default-drop-profile>, Forwarding class: ef1, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none,
Priority: low
    Excess Priority: low
```

```
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: , Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
  Congestion-notification: Disabled
Forwarding class                       ID      Queue  Restricted queue  Fabric priority
Policing priority
  af3                                   0       0          0             low
normal
  af2                                   1       1          1             low
normal
  ef2                                   2       2          2             high
normal
  ef1                                   3       3          3             high
normal
  af1                                   4       4          0             low
normal

  Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152) (Generation 159)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
```

```
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
    Protocol inet, MTU: 1500, Generation: 172, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Input Filters: filter-in-ge-0/3/0.0-i,
      Policer: Input: p1-ge-0/3/0.0-inet-i
    Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 173, Route table: 0
      Flags: Is-Primary
      Output Filters: exp-filter,,,,,

Logical interface ge-1/2/0.0 (Index 347) (SNMP ifIndex 638) (Generation 156)
Forwarding class ID  Queue  Restricted queue  Fabric priority  Policing priority   SPU priority
best-effort      0   0      0                 low              normal               low

Aggregate Forwarding-class statistics per forwarding-class
  Aggregate Forwarding-class statistics:
   Forwarding-class statistics:

Forwarding-class best-effort statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0
```

```
IPv4 protocol forwarding-class statistics:
  Forwarding-class statistics:
  Forwarding-class best-effort statistics:

    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

    IPv6  protocol forwarding-class statistics:
    Forwarding-class statistics:
     Forwarding-class best-effort statistics:

    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
```

```
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0

Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.0      up    up   inet  filter-in-ge-0/3/0.0-i
                           mpls                       exp-filter
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.0      up    up
                           inet  p1-ge-0/3/0.0-inet-i
                           mpls

Filter: filter-in-ge-0/3/0.0-i
Counters:
Name                                                Bytes              Packets
count-filter-in-ge-0/3/0.0-i                            0                    0

Filter: exp-filter
Counters:
Name                                                Bytes              Packets
count-exp-seven-match                                   0                    0
count-exp-zero-match                                    0                    0
Policers:
Name                                               Packets
p1-ge-0/3/0.0-inet-i                                    0


  Logical interface: ge-0/3/0.0, Index: 68
    Object                  Name                   Type                    Index
    Rewrite                 exp-default            exp (mpls-any)             33

Rewrite rule: exp-default, Code point type: exp, Index: 33
  Forwarding class                     Loss priority       Code point
  af3                                  low                 000
```

```
  af3                                high                 001
  af2                                low                  010
  af2                                high                 011
  ef2                                low                  100
  ef2                                high                 101
  ef1                                low                  110
  ef1                                high                 111
    Object                  Name                    Type                      Index
    Classifier              exp-default             exp                          10

Classifier: exp-default, Code point type: exp, Index: 10
  Code point         Forwarding class                    Loss priority
  000                af3                                 low
  001                af3                                 high
  010                af2                                 low
  011                af2                                 high
  100                ef2                                 low
  101                ef2                                 high
  110                ef1                                 low
  111                ef1                                 high
    Object                  Name                    Type                      Index
    Classifier              ipprec-compatibility    ip                           13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
  Code point         Forwarding class                    Loss priority
  000                af3                                 low
  001                af3                                 high
  010                af3                                 low
  011                af3                                 high
  100                af3                                 low
  101                af3                                 high
  110                ef1                                 low
  111                ef1                                 high
Forwarding class                       ID      Queue  Restricted queue  Fabric priority
Policing priority
  af3                                   0       0          0             low
normal
  af2                                   1       1          1             low
normal
  ef2                                   2       2          2             high
normal
  ef1                                   3       3          3             high
normal
```

```
  af1                                   4       4          0             low
normal

  Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154) (Generation 160)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 1500, Generation: 174, Route table: 0
      Flags: Sendbcast-pkt-to-re

  Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0


Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.1      up    up   mpls
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.1      up    up
                           mpls

  Logical interface: ge-0/3/0.1, Index: 69
    Object                  Name                   Type                     Index
    Classifier              ipprec-compatibility   ip                          13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
  Code point         Forwarding class                    Loss priority
  000                af3                                 low
  001                af3                                 high
```

```
  010                af3                                low
  011                af3                                high
  100                af3                                low
  101                af3                                high
  110                ef1                                low
  111                ef1                                high
Forwarding class                       ID      Queue  Restricted queue  Fabric priority
Policing priority
  af3                                   0       0          0             low
normal
  af2                                   1       1          1             low
normal
  ef2                                   2       2          2             high
normal
  ef1                                   3       3          3             high
normal
  af1                                   4       4          0             low
normal
```

show class-of-service interface (ACX Series Routers)

```
user@host-g11# show class-of-service interface
Physical interface: at-0/0/0, Index: 130
Maximum usable queues: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

  Logical interface: at-0/0/0.0, Index: 69

  Logical interface: at-0/0/0.32767, Index: 70

Physical interface: at-0/0/1, Index: 133
Maximum usable queues: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

  Logical interface: at-0/0/1.0, Index: 71

  Logical interface: at-0/0/1.32767, Index: 72
```

```
Physical interface: ge-0/1/0, Index: 146
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Rewrite                 dscp-default            dscp                       31
Classifier              d1                      dscp                    11331
Classifier              ci                      ieee8021p                 583

  Logical interface: ge-0/1/0.0, Index: 73
Object                  Name                    Type                    Index
Rewrite                 custom-exp              exp (mpls-any)          46413

  Logical interface: ge-0/1/0.1, Index: 74

  Logical interface: ge-0/1/0.32767, Index: 75

Physical interface: ge-0/1/1, Index: 147
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

  Logical interface: ge-0/1/1.0, Index: 76

Physical interface: ge-0/1/2, Index: 148
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Rewrite                 ri                      ieee8021p (outer)       35392
Classifier              ci                      ieee8021p                 583

Physical interface: ge-0/1/3, Index: 149
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

  Logical interface: ge-0/1/3.0, Index: 77
```

```
Object                  Name                    Type                    Index
Rewrite                 custom-exp2             exp (mpls-any)          53581

Physical interface: ge-0/1/4, Index: 150
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

Physical interface: ge-0/1/5, Index: 151
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

Physical interface: ge-0/1/6, Index: 152
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

Physical interface: ge-0/1/7, Index: 153
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              d1                      dscp                    11331

Physical interface: ge-0/2/0, Index: 154
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

Physical interface: ge-0/2/1, Index: 155
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
```

```
Classifier              ipprec-compatibility   ip                        13

  Logical interface: ge-0/2/1.0, Index: 78

  Logical interface: ge-0/2/1.32767, Index: 79

Physical interface: xe-0/3/0, Index: 156
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                   Type                    Index
Classifier              ipprec-compatibility   ip                        13

  Logical interface: xe-0/3/0.0, Index: 80

Physical interface: xe-0/3/1, Index: 157
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                   Type                    Index
Classifier              ipprec-compatibility   ip                        13

  Logical interface: xe-0/3/1.0, Index: 81

[edit]
user@host-g11#
```

show class-of-service interface (PPPoE Subscriber Interface for Enhanced Subscriber Management)

```
user@host> show class-of-service interface pp0.3221225474
  Logical interface: pp0.3221225475, Index: 3221225475
Object                  Name                   Type                    Index
Traffic-control-profile TC_PROF_100_199_SERIES_UID1006 Output      4294967312
Scheduler-map           SMAP-1_UID1002         Output              4294967327
Rewrite-Output          ieee-rewrite           ieee8021p               60432
Rewrite-Output          rule1                  ip                      50463

    Adjusting application: PPPoE IA tags
      Adjustment type: absolute
      Configured shaping rate: 11000000
```

```
      Adjustment value: 5000000
      Adjustment target: node

    Adjusting application: ucac
      Adjustment type: delta
      Configured shaping rate: 5000000
      Adjustment value: 100000
      Adjustment target: node
```

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

show class-of-service interface-set

IN THIS SECTION

- Syntax | **19100**
- Description | **19101**
- Options | **19101**
- Required Privilege Level | **19101**
- Output Fields | **19101**
- Sample Output | **19103**
- Release Information | **19103**

Syntax

```
show class-of-service interface-set
<interface-set-name>
```

Description

Display the configured shaping rate and the adjusted shaping rate for each logical interface set configured for hierarchical class of service (CoS).

Options

none	Display CoS associations for all logical interface sets.
interface-set *interface-set-name*	(Optional) Display CoS associations for the specified interface set.

Required Privilege Level

view

Output Fields

Table 379 on page 19101 describes the output fields for the `show class-of-service interface-set` command. Output fields are listed in the approximate order in which they appear.

Table 379: show class-of-service interface-set Output Fields

Field Name	Field Description
`Interface-set`	Name of a logical interface set composed of one or more logical interfaces for which hierarchical scheduling is enabled.
`Index`	Index number of this interface set or the internal index number of this object.
`Physical interface`	Name of a physical interface.
`Queues supported`	Number of queues you can configure on the interface.
`Queues in use`	Number of queues currently configured.
`Output traffic control profile`	Name of the output traffic control profile attached to the logical interface set.

Table 379: show class-of-service interface-set Output Fields *(Continued)*

Field Name	Field Description
`Output traffic control profile remaining`	(Enhanced subscriber management for MX Series routers) For dynamic subscriber management, name of the output traffic control profile for remaining traffic attached to the logical interface set.
`Adjusting application`	Name of the application that communicates shaping-rate adjustment information to the Junos OS class-of-service process (**cosd**) on the broadband services router (BSR). The BSR uses the information from this application to perform shaping-rate adjustments on the scheduler node that manages the interface set. The adjusting application appears as **ancp LS-0** which is the Junos OS Access Node Control Profile process (**ancpd**) that performs shaping-rate adjustments on schedule nodes. The nodes are logical interface sets configured to represent subscriber local loops. When the synchronization speed of the DSL line changes, **ancpd** communicates the local loop speed to **cosd** over the default logical system, **LS-0**, and then the BSR throttles the shaping rate on the scheduler node to the loop speed. The adjusting application can also appear as **PPPoE**, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual data rate downstream attribute. The overhead accounting value is based on the access loop encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).
`Adjustment type`	Type of shaping-rate adjustment performed by the BSR on the scheduler node. The type of adjustment appears as **Adjustment type**, meaning that the configured shaping rate is adjusted by an absolute value as opposed to by a percentage of the configured rate.
`Configured shaping rate`	The maximum transmission rate on the physical interface as configured by the output traffic-control profile attached to the scheduler node.
`Adjustment value`	Value of the shaping-rate adjustment information sent by the adjusting application to **cosd**.
`Adjustment overhead-accounting mode`	Configured shaping mode: `frame` or `cell`.

Sample Output

show class-of-service interface-set

```
user@host> show class-of-service interface-set example-ifset-ge-4/0/0-7
Interface-set: example-ifset-ge-4/0/0-7, Index: 8
Physical interface: ge-4/0/0, Index: 270
Queues supported: 8, Queues in use: 8
  Output traffic control profile: example-tcp-basic-rate, Index: 11395
Adjusting application: ancp LS-0
  Adjustment type: absolute
  Configured shaping rate: 50000000
  Adjustment value: 888000
  Adjustment overhead-accounting mode: cell
```

show class-of-service interface-set (Enhanced Subscriber Management)

```
user@host> show class of service interface-set
Interface-set: ge-1/0/0-201-201, Index: 1
Physical interface: ge-1/0/0, Index: 142
Queues supported: 8, Queues in use: 4
  Output traffic control profile: LEVEL_2_UID1001, Index: 4294967307
  Output traffic control profile remaining: TCP_REMAIN_UID1003, Index: 4294967308
```

Release Information

Command introduced in Junos OS Release 9.4.

show class-of-service l2tp-session

IN THIS SECTION

- Syntax | 19104
- Description | 19104

- Options | 19104
- Required Privilege Level | 19104
- Output Fields | 19104
- Sample Output | 19105
- Release Information | 19106

Syntax

```
show class-of-service l2tp-session session-id
```

Description

Display CoS objects associated with an L2TP session on M7i, M10i, and M120 routers.

Options

session-id L2TP session number for which you want to display a summary of CoS attributes.

Required Privilege Level

view

Output Fields

Table 380 on page 19104 lists the output fields for the `show class-of-service l2tp-session` command. Output fields are listed in the approximate order in which they appear.

Table 380: show class-of-service l2tp-session Output Fields

Field Name	Field Description
`L2TP Session Username`	Username associated with the L2TP session.

Table 380: show class-of-service l2tp-session Output Fields *(Continued)*

Field Name	Field Description
Index	Session index identification number.
Physical interface	Physical interface on which the tunnel session is established.
Index	Index ID associated with the physical interface on which the tunnel session is established.
Queues supported	Number of scheduler queues supported for the L2TP session.
Queues in use	Number of scheduler queues active on the L2TP session.
Scheduler map	Scheduler map name associated with the session.
Index	Scheduler map index number associated with the session.
Shaping rate	Maximum bandwidth configured for the session. Each active queue on the session receives a maximum of the configured amount of absolute bandwidth or the configured percentage of bandwidth, even if more bandwidth is available.

Sample Output

show class-of-service l2tp-session

```
user@host> show class-of-service l2tp-session 123
L2TP Session Username: user1@example.com, Index: 12553
Physical interface: ge-2/0/0, Index: 130
Queues supported: 8, Queues in use: 4
  Scheduler map: GEN-SCHED-MAP-EF-65%, Index: 5212
  Shaping rate: 200000 bps
```

Release Information

Command introduced in Junos OS Release 8.2.

show class-of-service loss-priority-rewrite

IN THIS SECTION

- Syntax | 19106
- Description | 19106
- Options | 19106
- Required Privilege Level | 19107
- Output Fields | 19107
- Sample Output | 19107
- Release Information | 19108

Syntax

```
show class-of-service loss-priority-rewrite
<name name>
<type frame-relay-de>
```

Description

Display the mapping of the code-point value to the loss priority rewrite rule.

Options

none	Display all loss priority rewrite maps.
name	(Optional) Display the specified loss priority rewrite.
frame-relay-de	(Optional) Display the Frame Relay discard eligibility code-point information.

Required Privilege Level

view

Output Fields

This table describes the output fields for the `show class-of-service loss-priority-rewrite` command. Output fields are listed in the approximate order in which they appear.

Table 381: show class-of-service loss-priority-rewrite Output Fields

Field Name	Field Description
Loss-priority-rewrite	Name of the loss priority rewrite.
Code point type	Type: **frame-relay-de**.
Index	Internal index.
Loss priority	Loss priority of **low**, **medium-low**, **medium-high**, or **high**.
Code point	Code-point value.

Sample Output

show class-of-service loss-priority-rewrite

```
user@host> show class-of-service loss-priority-rewrite
Loss-priority-rewrite: frame-relay-de-default, Code point type: frame-relay-de, Index: 38
  Loss priority      Code point
  low                 0
  high                1
  medium-low          0
  medium-high         1
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

frame-relay-de (Defining Loss Priority Rewrites) | 7563

show class-of-service monitoring-profile

IN THIS SECTION

- Syntax | 19108
- Description | 19108
- Options | 19108
- Required Privilege Level | 19109
- Output Fields | 19109
- Sample Output | 19109
- Release Information | 19110

Syntax

```
show class-of-service monitoring-profile <name>
```

Description

Display configuration information for any monitoring profile you define to monitor peak queue length for virtual output queues (VOQs).

Options

none Display all monitoring profiles.

name (Optional) Display named monitoring profile.

Required Privilege Level

view

Output Fields

Table 382: show class-of-service monitoring-profile Output Fields

Field Name	Field Description
`Monitoring profile`	Name of the monitoring profile.
`Export filter`	Name of an export filter defined for the monitoring profile. A monitoring profile can have one or more export filters.
`Queue Number`	The number of the queue in the export filter. Each export filter can have one or more queues.
`Peak Queue Length`	The peak queue length percentage defined for the export filter. If the peak queue length exceeds this percentage, Junos exports VOQ telemetry data for all member queues of the export filter.

Sample Output

show class-of-service monitoring-profile

```
user@host> show class-of-service monitoring-profile
```

```
Monitoring profile: mp1
  Export filter         Queue Number                    Peak Queue Length
  ef1                    0                                 50%
  ef1                    1                                 50%
```

```
ef2                    2                                0%
ef2                    3                                0%
```

Release Information

Command introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

VOQ Queue-depth Monitoring

monitoring-profile | **10959**

show class-of-service multi-destination

IN THIS SECTION

- Syntax | **19110**
- Description | **19110**
- Options | **19111**
- Required Privilege Level | **19111**
- Output Fields | **19111**
- Sample Output | **19111**
- Release Information | **19112**

Syntax

```
show class-of-service multi-destination
```

Description

For each class-of-service (CoS) multidestination classifier, display the classifier type.

Options

none Display all multidestination classifiers.

Required Privilege Level

view

Output Fields

Table 383 on page 19111 describes the output fields for the `show class-of-service multi-destination` command. Output fields are listed in the approximate order in which they appear.

Table 383: show class-of-service multi-destination Output Fields

Field Name	Field Description
`Family ethernet`	Family to which the classifier belongs.
`Classifier Name`	Name of the classifier.
`Classifier Type`	Type of the classifier: `dscp` or `ieee-802.1`.
`Classifier Index`	Internal index of the classifier.

Sample Output

show class-of-service multi-destination

```
user@switch>                    show class-of-service multi-destination

Family ethernet:
Classifier Name                 Classifier Type                 Classifier Index
  ba-mcast-classifier             ieee-802.1                      62376
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers

Understanding CoS Classifiers

Understanding CoS Classifiers

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

show class-of-service policy-map

IN THIS SECTION

- Syntax | 19112
- Description | 19113
- Options | 19113
- Required Privilege Level | 19113
- Output Fields | 19113
- Sample Output | 19114
- Release Information | 19114

Syntax

```
show class-of-service policy-map
<policy-map-name>
<type (configured | reserved)>
```

Description

Display class of service (CoS) policy map information.

Options

policy-map-name (Optional) Enter the name of a policy map to show only the information for that policy map. Otherwise, information for all policy maps is displayed.

type (Optional) Display information about different types of policy maps.

- **configured** Display user configured policy maps.
- **reserved** Display reserved policy map name-to-ID mapping.

Required Privilege Level

view

Output Fields

Table 384 on page 19113 describes the output fields for the `show class-of-service policy-map` command. Output fields are listed in the approximate order in which they appear.

Table 384: show class-of-service policy-map Output Fields

Field Name	Field Description
`Type`	The type of packet marking to rewrite.
`Code Point`	The code point the packet marking should be rewritten to.
`Option`	The type of the traffic the packet marking should be rewritten for.

Sample Output

show class-of-service policy-map

```
user@host> show class-of-service policy-map
Policy-map: P-1, Index: 1
  Type              Code Point  Option
  inet-precedence        110    (proto-ip)
  inet-precedence        110    (proto-mpls)
  dscp-ipv6           101010    (proto-ip)
  dscp-ipv6           101010    (proto-mpls)
  exp                    110    (all-label)
  exp                    111    (outer-label)
  ieee-802.1ad          0110    (outer-and-inner)
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

policy-map (Class of Service) | 12686

Assigning Rewrite Rules on a Per-Customer Basis Using Policy Maps Overview

Configuring Policy Maps to Assign Rewrite Rules on a Per-Customer Basis

show class-of-service rewrite-rule

IN THIS SECTION

- Syntax | 19115
- Description | 19115
- Options | 19115
- Required Privilege Level | 19115
- Output Fields | 19116

Sample Output | 19116
Release Information | 19117

Syntax

```
show class-of-service rewrite-rule
<name name>
<type type>
```

Description

Display the mapping of forwarding classes and loss priority to code point values.

Options

none Display all rewrite rules.

name *name* (Optional) Display the specified rewrite rule.

type *type* (Optional) Display the rewrite rule of the specified type. The rewrite rule type can be one of the following:

- **dscp**—For IPv4 traffic.
- **dscp-ipv6**—For IPv6 traffic.
- **exp**—For MPLS traffic.
- **frame-relay-de**—(SRX Series only) For Frame Relay traffic.
- **ieee-802.1**—For Layer 2 traffic.
- **inet-precedence**—For IPv4 traffic.

Required Privilege Level

view

Output Fields

Table 385 on page 19116 describes the output fields for the `show class-of-service rewrite-rule` command. Output fields are listed in the approximate order in which they appear.

Table 385: show class-of-service rewrite-rule Output Fields

Field Name	Field Description
Rewrite rule	Name of the rewrite rule.
Code point type	Type of rewrite rule: **dscp**, **dscp-ipv6**, **exp**, **frame-relay-de**, or **inet-precedence**.
Forwarding class	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router or switch.
Index	Internal index for this particular rewrite rule.
Loss priority	Loss priority for rewriting.
Code point	Code point value to rewrite.

Sample Output

show class-of-service rewrite-rule type dscp

```
user@host> show class-of-service rewrite-rule type dscp
Rewrite rule: dscp-default, Code point type: dscp
  Forwarding class                    Loss priority       Code point
  gold                                high                000000
  silver                              low                 110000
  silver                              high                111000
  bronze                              low                 001010
  bronze                              high                001100
  lead                                high                101110

Rewrite rule: abc-dscp-rewrite, Code point type: dscp, Index: 3245
Forwarding class                    Loss priority       Code point
```

```
gold                    low         000111
gold                    high        001010
silver                  low         110000
silver                  high        111000
bronze                  high        001100
lead                    low         101110
lead                    high        110111
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Rewrite Rules Overview

show class-of-service routing-instance

IN THIS SECTION

- Syntax | 19117
- Description | 19118
- Options | 19118
- Required Privilege Level | 19118
- Output Fields | 19118
- Sample Output | 19119
- Release Information | 19119

Syntax

```
show class-of-service routing-instance
<routing-instance-name>
```

Description

Display mapping of class of service (CoS) objects to routing instances.

Options

routing-instance-name (Optional) Name of a routing instance.

Required Privilege Level

view

Output Fields

Table 386 on page 19118 describes the output fields for the `show class-of-service routing-instance` command. Output fields are listed in the approximate order in which they appear.

Table 386: show class-of-service routing-instance Output Fields

Field Name	Field Description
`Index`	Internal index.
`Name`	Name of an object.
`Object`	Category of an object: `Classifier`.
`Routing instance`	Name of a routing instance.
`Type`	Type: `exp`.

Sample Output

show class-of-service routing-instance

```
user@host> show class-of-service routing-instance
  Routing Instance : vpn1
    Object            Name                  Type                    Index
    Classifier        exp-default           exp                         8

  Routing Instance : vpn2
    Object            Name                  Type                    Index
    Classifier        test2                 exp                     57507
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service scheduler-hierarchy interface

IN THIS SECTION

- Syntax | 19119
- Description | 19120
- Options | 19120
- Required Privilege Level | 19120
- Output Fields | 19120
- Sample Output | 19121
- Release Information | 19121

Syntax

```
show class-of-service scheduler-hierarchy interface interface-name <detail>
```

Description

For MPC/MIC interfaces only, display the scheduler hierarchy as well as the shaping rate, guaranteed rate, priorities, and queue weight information for each forwarding class at each hierarchy level.

Options

detail (Optional) Display scheduler hierarchies based on the interface set.

interface-name Display information about a specific interface.

Required Privilege Level

view

Output Fields

Table 387 on page 19120 describes the output fields for the `show class-of-service scheduler-hierarchy interface` command. Output fields are listed in the approximate order in which they appear.

Table 387: show class-of-service scheduler-hierarchy interface Output Fields

Field Name	Field Description
`interface`	Interface name
`resource`	Traffic resource associated with the logical interface
`shaping-rate`	Shaping rate in bits per second
`guaranteed rate`	Guaranteed rate in bits per second
`guaranteed priority`	Queue priority in the guaranteed region (high, low, or none)
`excess priority`	Queue priority in the excess region (high, low, or none)
`queue weight`	Queue weight for excess CoS weighted round-robin

Table 387: show class-of-service scheduler-hierarchy interface Output Fields *(Continued)*

Field Name	Field Description
excess weight	Interface unit per priority weights for excess weighted round-robin

Sample Output

show class-of-service scheduler-hierarchy interface

```
user@host> show class-of-service scheduler-hierarchy interface xe-1/0/0
--------------------------------------------------------------------------------
Interface/                    shaping guaranteed  guaranteed/  queue   excess
resource name                    rate      rate       excess  weight   weight
                                kbits     kbits     priority          high/low
--------------------------------------------------------------------------------
xe-1/0/0                        12000                                            <<< L1
  xe-1/0/0 RTP                  12000         0                         1    1
    best-effort                 12000         0     Low  Low     950
    network-control             12000         0     Low  Low      50
  ifset1                        12000         0                       500  500   <<< L2
    ifset1 RTP                  12000         0                         1    1
      be1                         720         0     Low  Low     250
      nc1                       12000         0     Low  Low     250
    demux0.96                    3000         0                         1    1   <<< L3
      demux0.96 RTP              3000         0                       500  500
        be1                      1000         0     Low  Low     250
        nc1                      3000         0     Low  Low     250
      pp0.81                     2000         0                         1    1   <<< L4
         be1                     1000         0     Low  Low     250
         nc1                     2000         0     Low  Low     250
```

Release Information

Command introduced in Junos OS Release 13.3.

Support for up to four hierarchy levels added in Junos OS Release 16.1.

NOTE: Before Junos OS R19.2, the shaping rate would incorrectly display as 90% of the guaranteed rate.

RELATED DOCUMENTATION

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) | 8062

show class-of-service scheduler-hierarchy interface-set

IN THIS SECTION

- Syntax | 19122
- Description | 19122
- Options | 19122
- Required Privilege Level | 19123
- Output Fields | 19123
- Sample Output | 19124
- Release Information | 19124

Syntax

```
show class-of-service scheduler-hierarchy interface-set interface-set-name <detail>
```

Description

For MPC/MIC interface sets only, display the scheduler hierarchy.

Options

detail (Optional) Display scheduler hierarchies based on the interface-set.

interface-set-name Display information about a specific interface-set.

Required Privilege Level

view

Output Fields

Table 388 on page 19123 describes the output fields for the `show class-of-service scheduler-hierarchy interface-set` command. Output fields are listed in the approximate order in which they appear.

Table 388: show class-of-service scheduler-hierarchy interface-set Output Fields

Field Name	Field Description
`interface`	Type of interface
`resource`	Traffic resource associated with the logical interface
`shaping-rate`	Actual shaping rate in bits per second
`guaranteed rate`	Actual guaranteed rate in bits per second
`guaranteed priority`	Actual queue priority in the guaranteed region (high, low, or none)
`excess priority`	Actual queue priority in the excess region (high, low, or none)
`queue weight`	Actual queue weight for excess CoS weighted round-robin
`excess weight`	Actual interface-set per priority weights for excess weighted round-robin

Sample Output

show class-of-service scheduler-hierarchy interface-set

```
user@host> show class-of-service scheduler-hierarchy interface-set ifset
--------------------------------------------------------------------------------
Interface/                    shaping guarnteed  guaranteed/   queue   excess
resource name                    rate      rate       excess  weight   weight
                                kbits     kbits     priority          high/low
--------------------------------------------------------------------------------
ge-1/0/0                       100000
  ge-1/0/0 RTP                 100000         0                           1   1
    be                         100000      1000     Low  Low       1
    da                           9000      2000  Medium High       1
    vi                         100000      3000  Medium None     626
    vo                         100000      4000    High High     373
    gt                         100000         0    High High       1
  ge-1/0/0.20                   50000     40000                       750 750
    be                          50000      1000     Low  Low       1
    da                           9000      2000  Medium High       1
    vi                          50000      3000  Medium None     626
    vo                          50000      4000    High High     373
    gt                          50000  Disabled    High High       1
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

- show interfaces queue

show class-of-service scheduler-map

IN THIS SECTION

- Syntax | 19125
- Description | 19125
- Options | 19125
- Required Privilege Level | 19125
- Output Fields | 19126
- Sample Output | 19128
- Release Information | 19130

Syntax

```
show class-of-service scheduler-map
<name>
```

Description

Display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.

Options

none Display all scheduler maps.

name (Optional) Display a summary of scheduler parameters for each forwarding class to which the named scheduler is assigned.

Required Privilege Level

view

Output Fields

The following table lists all possible output fields for the `show class-of-service scheduler-map` command. The fields that appear may vary depending on the platform and software release.

Table 389: show class-of-service scheduler-map Output Fields

Field Name	Field Description
`Scheduler map`	Name of the scheduler map. The name of the dynamic scheduler map object is associated with a generated UID (for example, `SMAP-1_UID1002`) instead of with a subscriber interface.
`Index`	Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles. Index values for dynamic CoS traffic control profiles are larger for enhanced subscriber management than they are for legacy subscriber management.
`Scheduler`	Name of the scheduler.
`Forwarding class`	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
`Transmit rate`	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword `remainder`, which indicates that the scheduler receives the remaining bandwidth of the interface.
`Rate Limit`	Rate limiting configuration of the queue. Possible values are `none`, meaning no rate limiting, and `exact`, meaning the queue only transmits at the configured rate.
`Buffer Rate`	The buffer rate is the target rate of a VOQ, which is the intended egress queue rate during typical congestion. The target rate is an import parameter for VOQs, as the ratio of the drain rate to the target rate determines how much buffer adjustment happens at ingress.
`Maximum buffer delay`	Amount of transmit delay (in milliseconds) or the buffer size of the queue. The buffer size is shown as a percentage of the total interface buffer allocation, or by the keyword `remainder` to indicate that the buffer is sized according to what remains after other scheduler buffer allocations.

Table 389: show class-of-service scheduler-map Output Fields *(Continued)*

Field Name	Field Description
`Buffer size`	The size of the buffer as a percent of the total buffer size for the port. The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent. On QFX10000 switches, the buffer size is the amount of time in milliseconds of port bandwidth that a queue can use to continue to transmit packets during periods of congestion, before the buffer runs out and packets begin to drop.
`Buffer dynamic threshold`	The alpha value of the shared-buffer pool at the egress buffer partition for each queue.
`Priority`	Scheduling priority: `low` or `high`.
`Excess priority`	Priority of excess bandwidth: `low`, `medium-low`, `medium-high`, `high`, or `none`.
`Explicit Congestion Notification`	Explicit congestion notification (ECN) state: • Disable—ECN is disabled on the specified scheduler • Enable—ECN is enabled on the specified scheduler ECN is disabled by default.
`Adjust minimum`	Minimum shaping rate for an adjusted queue, in bps.
`Adjust percent`	Bandwidth adjustment applied to a queue, in percent.
`Drop profiles`	Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.
`Loss priority`	Packet loss priority for drop profile assignment.
`Protocol`	Transport protocol for drop profile assignment.

Table 389: show class-of-service scheduler-map Output Fields *(Continued)*

Field Name	Field Description
Name	Name of the drop profile.

Sample Output

show class-of-service scheduler-map (MX Series)

```
user@host> show class-of-service scheduler-map
Scheduler map: dd-scheduler-map, Index: 84

  Scheduler: aa-scheduler, Forwarding class: aa-forwarding-class, Index: 8721
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 95 percent, Buffer Limit: none,
Priority: high
    Excess Priority: low, Queue Depth Monitoring: disabled
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         8724     aa-drop-profile
      Medium low      any         9874     bb-drop-profile
      Medium high     any         8833     cc-drop-profile
      High            any         8484     dd-drop-profile

  Scheduler: bb-scheduler, Forwarding class: aa-forwarding-class, Index: 223
    Transmit rate: 40 percent, Rate limit: none, Buffer size: 5 percent, Buffer Limit: none,
Priority: low
    Excess Priority: low, Queue Depth Monitoring: disabled
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         8724     aa-drop-profile
      Medium low      any         9874     bb-drop-profile
      Medium high     any         8833     cc-drop-profile
      High            any         8484     dd-drop-profile
```

show class-of-service scheduler-map (EX4100 and EX4400)

```
user@host> show class-of-service scheduler-map
Scheduler map: <default>, Index: 2

  Scheduler: <default-be>, Forwarding-class: best-effort, Index: 21
    Transmit rate: 75 percent, Rate Limit: none, Buffer size: 75 percent, Buffer Limit: none,
Priority: low
    Excess Priority: low, Explicit Congestion Notification: disable, ECN Security Log Rate
Limit: 65535, Queue Depth Monitoring: disabled
    Drop profiles:
      Loss priority    Protocol    Index    Name
      Low              non-TCP         1    <default-drop-profile>
      Low              TCP             1    <default-drop-profile>
      Medium low       any             1    <default-drop-profile>
      Medium high      any             1    <default-drop-profile>
      High             non-TCP         1    <default-drop-profile>
      High             TCP             1    <default-drop-profile>

  Scheduler: mcast-nc-scheduler, Forwarding-class: mcast-nc, Index: 48
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer Limit: none,
Priority: low
    Excess Priority: low, Explicit Congestion Notification: disable, ECN Security Log Rate
Limit: 65535, Queue Depth Monitoring: disabled
    Drop profiles:
      Loss priority    Protocol    Index    Name
      Low              non-TCP         1    <default-drop-profile>
      Low              TCP             1    <default-drop-profile>
      Medium low       any             1    <default-drop-profile>
      Medium high      any             1    <default-drop-profile>
      High             non-TCP         1    <default-drop-profile>
      High             TCP             1    <default-drop-profile>
```

show class-of-service scheduler-map (PTX10001-36MR, PTX10004, PTX10008, PTX10016)

```
user@host# show class-of-service scheduler-map smap
Scheduler map: smap, Index: 4

  Scheduler: s2, Forwarding class: best-effort, Index: 5
    Transmit rate: unspecified, Rate Limit: none, Buffer Rate: 40000000000 bps, Buffer Limit:
```

```
none, Buffer dynamic threshold: unspecified, Priority: low
    Excess Priority: unspecified, Excess rate: unspecified, Explicit Congestion Notification:
disable, ECN pfc no assist: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             0    default-drop-profile
      Medium low      any             0    default-drop-profile
      Medium high     any             0    default-drop-profile
      High            any             0    default-drop-profile
```

show class-of-service scheduler-map (QFX Series)

```
user@switch# show class-of-service scheduler-map
Scheduler map: be-map, Index: 12240

  Scheduler:be-sched, Forwarding class: best-effort, Index: 115
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable

    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any          3312    lan-dp
      Medium-high     any          2714    be-dp1
      High            any          3178    be-dp2
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

show class-of-service scheduler-map-forwarding-class-sets

IN THIS SECTION

- Syntax | 19131
- Description | 19131
- Options | 19131
- Required Privilege Level | 19131
- Output Fields | 19132
- Sample Output | 19133
- Release Information | 19134

Syntax

```
show class-of-service scheduler-map-forwarding-class-sets
<scheduler-map-fcset-name>
```

Description

Display the mapping of fabric schedulers to fabric fc-sets and a summary of scheduler parameters for each entry on Interconnect devices.

Options

none	Display all fabric scheduler maps.
scheduler-map-fcset-name	(Optional) Display a summary of fabric scheduler parameters for each fabric fc-set to which the specified scheduler is assigned.

Required Privilege Level

view

Output Fields

Table 390 on page 19132 describes the output fields for the `show class-of-service scheduler-map fc-set` command. Output fields are listed in the approximate order in which they appear.

Table 390: show class-of-service scheduler-map-fcset Output Fields

Field Name	Field Description
`Scheduler map fcset`	Name of the fabric scheduler map.
`Index`	Internal index of the indicated object. Objects having indexes in this output include fabric scheduler maps, schedulers, and drop profiles.
`Scheduler`	Name of the scheduler.
`Forwarding class set`	Name of the fabric forwarding class set (fabric fc-set) mapped to the scheduler.
`Transmit rate`	Configured transmit rate of the scheduler (percentage). The transmit rate of a fabric scheduler determines the proportion of class group (priority group) bandwidth that the associated fabric fc-set receives. The sum of the fabric fc-set transmit rates in a class group determines the percentage of available port bandwidth allocated to the class group.
`Rate Limit`	Rate limiting configuration of the queue. Possible values are `none`, meaning no rate limiting, and `exact`, meaning the queue only transmits at the configured rate.
`Buffer size`	Egress queue buffer size (percentage).
`Buffer limit`	Buffer limiting configuration of the queue. Possible values are `none`, meaning no buffer limiting, and `exact`, meaning the buffer is limited to the configured buffer size.
`Priority`	Scheduling priority: `low` or `high` (default value only; value is not user-configurable).
`Excess priority`	Priority of excess bandwidth: `low`, `medium-low`, `medium-high`, `high`, or `none` (default value only; value is not user-configurable).
`Shaping rate`	Maximum bandwidth allowed to traffic mapped to the scheduler.

Table 390: show class-of-service scheduler-map-fcset Output Fields *(Continued)*

Field Name	Field Description
Drop profiles	Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.
Loss priority	Packet loss priority for drop profile assignment.
Protocol	Transport protocol for drop profile assignment.
Name	Name of the drop profile.

Sample Output

show class-of-service scheduler-map-forwarding-class-sets

```
user@host> show class-of-service scheduler-map-forwarding-class-sets
Scheduler map fcset: fab-be-map, Index: 2

  Scheduler: fab-be-sched, Forwarding class set: fabric_fcset_be, Index: 21
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Shaping rate: 100 percent
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         55387    fab-dp-be-low
      Medium high     any             1    <default-drop-profile>
      High            any          4369    fab-dp-be-high

Scheduler map fcset: fab-fcoe-map, Index: 3

  Scheduler: fab-fcoe-sched, Forwarding class set: fabric_fcset_noloss1, Index: 23
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 30 percent,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Shaping rate: 100 percent
```

```
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>
```

Release Information

Command introduced in Junos OS Release 13.1R2.

RELATED DOCUMENTATION

Understanding CoS Scheduling Across the QFabric System

Understanding CoS Fabric Forwarding Class Sets

Example: Configuring CoS Scheduling Across the QFabric System

show class-of-service scheduler-resources fpc

IN THIS SECTION

- Syntax | 19134
- Description | 19135
- Options | 19135
- Required Privilege Level | 19135
- Output Fields | 19135
- Sample Output | 19136
- Release Information | 19137

Syntax

```
show class-of-service scheduler-resources fpc slot slot-number
```

Description

Display scheduler resources for a line card.

Options

slot *slot-number* Display scheduler resources for a specific line card. `show class-of-service scheduler-resources fpc` displays scheduler resources for all line cards.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show class-of-service scheduler-resources fpc` command.

Table 391: show class-of-service scheduler-resources Output Fields

Field Name	Field Description
`Slot`	The slot number for the line card.
`PFE`	The packet forwarding engine (PFE) on the line card.
`Scheduler`	The PFE scheduler.
`Used`	The number of schedulers in use per scheduling hierarchy level.
`Total`	The total schedulers available per scheduling hierarchy level.
`Resources Unavailable for (seconds)`	The number of seconds a scheduling resource has been unavailable.

Sample Output

show class-of-service scheduler-resources fpc slot *slot-number*

```
user@host> show class-of-service scheduler-resources fpc slot 0
```

```
  Slot: 0
    PFE: 0
      Scheduler: 0
                                                  Resource Unavailable
                        Used        Total         for (seconds)
                L1        11           63
                L2        11         4095
                L3        11        16383
                Q         88       131064
      Scheduler: 1
                                                  Resource Unavailable
                        Used        Total         for (seconds)
                L1        11           63
                L2        11         4095
                L3        11        16383
                Q         88       131064
    PFE: 1
      Scheduler: 0
                                                  Resource Unavailable
                        Used        Total         for (seconds)
                L1         2           63
                L2         2         4095
                L3         2        16383
                Q         16       131064
      Scheduler: 1
                                                  Resource Unavailable
                        Used        Total         for (seconds)
                L1         2           63
                L2         2         4095
                L3         2        16383
                Q         16       131064
```

Release Information

Command introduced in Junos OS Release 19.3R1.

show class-of-service shared-buffer

IN THIS SECTION

- Syntax | 19137
- Description | 19137
- Options | 19138
- Required Privilege Level | 19138
- Output Fields | 19138
- Sample Output | 19139
- Release Information | 19140

Syntax

```
show class-of-service shared-buffer
<egress | ingress>
```

Description

Display the shared buffer allocation and partitioning configuration.

> **NOTE**: Due to QFX5200 cross-point architecture, all buffer usage counters are maintained separately. When usage counters are displayed with the command show class-of-service shared-buffer on QFX5200, various pipe counters are displayed separately.

Options

none Display ingress and egress shared buffer settings.

egress (Optional) Display the egress shared buffer settings.

ingress (Optional) Display the ingress shared buffer settings.

Required Privilege Level

view

Output Fields

Table 392 on page 19138 describes the output fields for the `show class-of-service shared-buffer` command. Output fields are listed in the approximate order in which they appear.

Table 392: show class-of-service shared-buffer Output Fields

Field Name	Field Description
`Ingress`	Ingress shared buffer configuration.
`Total Buffer`	Total buffer space available to the ports in KB. This is the combined dedicated buffer pool and shared buffer pool.
`Dedicated Buffer`	Buffer space allocated to the dedicated buffer pool in KB.
`Shared Buffer`	Buffer space allocated to the shared buffer pool in KB.
`Lossless`	Buffer space allocated to the lossless traffic buffer pool in KB.
`Lossless Headroom`	Buffer space allocated to the lossless headroom traffic buffer pool to support priority-based flow control (PFC) and Ethernet PAUSE in KB. (Ingress ports only.) **NOTE**: OCX Series switches do not support PFC.
`Lossy`	Buffer space allocated to the lossy (best-effort) traffic buffer pool in KB.

Table 392: show class-of-service shared-buffer Output Fields *(Continued)*

Field Name	Field Description
Lossless Headroom Utilization	Utilization of the ingress lossless headroom buffer pool. (These fields can help you to determine how much headroom buffer space you need to reserve to support PFC and Ethernet PAUSE for lossless flows.) **NOTE**: OCX Series switches do not support PFC.
Node Device	Index number that identifies the switch. On a QFX3500 switch, this field always has a value of zero (0).
Total	Size of the lossless headroom ingress buffer pool in KB.
Used	Amount in KB of lossless headroom ingress buffer used.
Free	Amount in KB of lossless headroom ingress buffer free (unused).
Egress	Egress shared buffer configuration.
Multicast	Buffer space allocated to the multicast traffic buffer pool in KB. (Egress ports only.)

Sample Output

show class-of-service shared-buffer

```
user@switch> show class-of-service shared-buffer
Ingress:
  Total Buffer     :  9360.00 KB
  Dedicated Buffer :  2158.00 KB
  Shared Buffer    :  7202.00 KB
    Lossless          :  648.18 KB
    Lossless Headroom :  3240.90 KB
    Lossy             :  3312.92 KB

  Lossless Headroom Utilization:
  Node Device         Total          Used                  Free
```

```
  0                    3240.90 KB     0.00 KB               3240.90 KB

Egress:
  Total Buffer     :  9360.00 KB
  Dedicated Buffer :  2704.00 KB
  Shared Buffer    :  6656.00 KB
    Lossless          :  3328.00 KB
    Multicast         :  1264.64 KB
    Lossy             :  2063.36 KB
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic

Configuring Global Ingress and Egress Shared Buffers

Understanding CoS Buffer Configuration

show class-of-service slice <slice_name> interface <interface_name>

IN THIS SECTION

- Syntax | 19141
- Description | 19141
- Required Privilege Level | 19141
- Output Fields | 19141
- Sample Output | 19142

Release Information | 19142

Syntax

```
show class-of-service slice <slice_name> interface <interface_name>
```

Description

Show traffic control profile(s) attached to a slice under an interface.

Required Privilege Level

view

Output Fields

Table 393 on page 19141 lists the output fields for the `show class-of-service slice <slice_name> interface <interface_name>` command.

Table 393: show class-of-service slice <slice_name> interface <interface_name> Output Fields

Field Name	Field Description
Slice	Slice name
Index	Slice index
Output traffic control profile	The name of the output traffic control profile
Index	The index of the traffic control profile attached to the slice under this interface.

Sample Output

show class-of-service slice <slice_name> interface <interface_name>

```
user@host> show class-of-service slice <slice_name> interface <interface_name>
```

```
user@host# show class-of-service slice s1 interface xe-4/0/5
  Slice: s1, Index: 1
  Output traffic control profile: tcp1, Index: 33290
  Slice scale-factor: 1
    Scheduler map: <default>, Index: 2
    Shaping rate: 1000000000 bps
    Guaranteed rate: 0 bps
```

Release Information

Junos 22.2R1

show class-of-service traffic-class-map

IN THIS SECTION

- Syntax | 19143
- Description | 19143
- Options | 19143
- Required Privilege Level | 19143
- Output Fields | 19143
- Sample Output | 19144
- Release Information | 19144

Syntax

```
show class-of-service traffic-class-map
<name traffic-class-map-name>
<type (dscp | exp | ieee-802.1 | ieee-802.1ad |inet-precedence)>
```

Description

For each traffic class map, display the mapping of the code point value to the input traffic class.

Options

none	Display all the mappings.
name *name*	(Optional) Display the named traffic class map.
type dscp	(Optional) Display all traffic class maps of the Differentiated Services code point (DSCP) type.
type exp	(Optional) Display all traffic class maps of the MPLS EXP type.
type ieee-802.1	(Optional) Display all traffic class maps of the IEEE 802.1 type.
type ieee-802.1ad	(Optional) Display all traffic class maps of the IEEE 802.1ad type.
type inet-precedence	(Optional) Display all traffic class maps of the IPv4 precedence type.

Required Privilege Level

view

Output Fields

Table 394 on page 19144 describes the output fields for the `show class-of-service traffic-class-map` command. Output fields are listed in the approximate order in which they appear.

Table 394: show class-of-service traffic-class-map Output Fields

Field Name	Field Description
`Traffic-class-map`	Name of the traffic class map.
`Code point type`	Type of the traffic class map: `exp`, `dscp`, `ieee-802.1`, `ieee-802.1ad`, or `inet-precedence`.
`Index`	Internal index of the traffic class map.
`Code point`	Code point value used for classification.
`Traffic class`	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.

Sample Output

show class-of-service traffic-class-map

```
user@host> show class-of-service traffic-class-map
Traffic-class-map: inet-precedence, Code-point type: inet-precedence, Index: 44231
          Code point         Traffic class
          000                real-time
          001                real-time
          010                network-control
          011                network-control
          100                best-effort
          101                best-effort
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

traffic-class-map | 30003

Managing Ingress Oversubscription at the PFE

Configuring Traffic Class Maps to Manage Ingress Oversubscription

Example: Configuring Traffic Class Maps

show class-of-service forwarding-table traffic-class-map | 19051

show class-of-service traffic-control-profile

IN THIS SECTION

- Syntax | 19145
- Description | 19145
- Options | 19145
- Required Privilege Level | 19146
- Output Fields | 19146
- Sample Output | 19148
- Release Information | 19150

Syntax

```
show class-of-service traffic-control-profile
<profile-name>
```

Description

Display information on traffic shaping and scheduling profiles.

Options

none	Display all profiles.

profile-name (Optional) Display information about a single profile.

Required Privilege Level

view

Output Fields

Table 395 on page 19146 describes the possible output fields for the `show class-of-service traffic-control-profile` command. The output fields that display depend on the platform and software device of your device. Output fields are listed in the approximate order in which they appear.

Table 395: show class-of-service traffic-control-profile Output Fields

Field Name	Field Description
`Traffic control profile`	Name of the traffic control profile. (Enhanced subscriber management for MX Series routers) The name of the dynamic traffic control profile object is associated with a generated UID (for example, `TC_PROF_100_199_SERIES_UID1000`) instead of with a subscriber interface.
`Index`	Index number of the traffic control profile. (Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles are larger for enhanced subscriber management than they are for legacy subscriber management.
`ATM Service`	(MX Series routers with ATM Multi-Rate CE MIC) Configured category of ATM service. Possible values: • cbr—Constant bit rate. • rtvbr—Real time variable bit rate. • nrtvbr—Non real time variable bit rate. • ubr—Unspecified bit rate.
`Maximum Burst Size`	Configured maximum burst size, in cells.

Table 395: show class-of-service traffic-control-profile Output Fields *(Continued)*

Field Name	Field Description
`Peak rate`	Configured peak rate, in cps.
`Sustained rate`	Configured sustained rate, in cps.
`Shaping rate`	Configured shaping rate, in bps. **NOTE**: (MX Series routers with ATM Multi-Rate CE MIC) Configured peak rate, in cps.
`Shaping rate burst`	Configured burst size for the shaping rate, in bytes. **NOTE**: (MX Series routers with ATM Multi-Rate CE MIC) Configured maximum burst rate, in cells.
`Shaping rate priority high`	Configured shaping rate for high-priority traffic, in bps.
`Shaping rate priority medium`	Configured shaping rate for medium-priority traffic, in bps.
`Shaping rate priority low`	Configured shaping rate for low-priority traffic, in bps.
`Shaping rate excess high`	Configured shaping rate for high-priority excess traffic, in bps.
`Shaping rate excess low`	Configured shaping rate for low-priority excess traffic, in bps.
`Scheduler map`	Name of the associated scheduler map. (Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, `SMAP-1_UID1002`) instead of with a subscriber interface.
`Delay Buffer rate`	Configured delay buffer rate, in bps.
`Excess rate`	Configured excess rate, in percent or proportion.

Table 395: show class-of-service traffic-control-profile Output Fields *(Continued)*

Field Name	Field Description
`Excess rate high`	Configured excess rate for high priority traffic, in percent or proportion.
`Excess rate low`	Configured excess rate for low priority traffic, in percent or proportion.
`Guaranteed rate`	Configured guaranteed rate, in bps or cps. **NOTE**: (MX Series routers with ATM Multi-Rate CE MIC) This value depends on the ATM service category chosen. Possible values: • `cbr`—Guaranteed rate is equal to the configured peak rate in cps. • `rtvbr`—Guaranteed rate is equal to the configured sustained rate in cps. • `nrtvbr`—Guaranteed rate is equal to the configured sustained rate in cps.
`Guaranteed rate burst`	Configured burst size for the guaranteed rate, in bytes.
`adjust-minimum`	Configured minimum shaping rate for an adjusted queue, in bps.
`overhead accounting mode`	Configured shaping mode: `Frame Mode` or `Cell Mode`.
`Overhead bytes`	Configured byte adjustment value.
`Adjust parent`	Configured shaping-rate adjustment for parent scheduler nodes. If enabled, this field appears. `flow-aware` indicates that the parent scheduler node is adjusted only once per multicast channel.

Sample Output

show class-of-service traffic-control-profile

```
user@host> show class-of-service traffic-control-profile
Traffic control profile: Profile1, Index: 57625
```

```
  Scheduler map: m1
  Delay Buffer rate: 500000
  Guaranteed rate: 1000000
  Overhead bytes: 20

Traffic control profile: Profile2, Index: 57624
  Scheduler map: m2
  Delay Buffer rate: 600000
  Guaranteed rate: 2000000

Traffic control profile: Profile3, Index: 57627
  Scheduler map: m3
  Delay Buffer rate: 800000
  Guaranteed rate: 3000000
 .Excess rate high: proportion 4

Traffic control profile: Profile4, Index: 57626
  Scheduler map: m4
  Delay Buffer rate: 750000
  Guaranteed rate: 4000000
..adjust-minimum 20000000

 Traffic control profile: foo, Index: 57626
   Shaping rate: 100000000
   Scheduler map: <default>
   Overhead accounting mode: Frame Mode
   Frame mode overhead accounting bytes: -12
   Adjust parent: flow-aware
```

show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC)

```
user@host> show class-of-service traffic-control-profile
Traffic control profile: at-vbr1, Index: 11395
  ATM Service: RTVBR
  Scheduler map: m3
  overhead accounting mode: Frame Mode
  Shaping rate: 1000 cps
  Shaping rate burst: 500 cells
  Delay Buffer rate: 2000 cps
  Guaranteed rate: 1000 cps
```

```
Traffic control profile: foo, Index: 38286
  ATM Service: UBR
  Scheduler map: m3
  overhead accounting mode: Frame Mode
```

show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces)

```
user@host> show class-of-service traffic-control-profile
Traffic control profile: foo, Index: 38286
  ATM Service: RTVBR
  Shaping rate: 2000 cps
  Shaping rate burst: 200 cells
  Scheduler map: <default>
  Delay Buffer rate: 1000 cps
  Guaranteed rate: 1700 cps
```

show class-of-service traffic-control-profile (Enhanced Subscriber Management)

```
user@host> show class-of-service traffic-control-profile
Traffic control profile: TC_PROF_100_199_SERIES_UID1000, Index: 4294967313
  Shaping rate: 11000000
  Shaping rate burst: 1 bytes
  Scheduler map: SMAP-1_UID1002
  Delay Buffer rate: 5000000
  Overhead accounting mode: Cell Mode
  Frame mode overhead accounting bytes: -4
  Cell mode overhead accounting bytes: 20
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

traffic-control-profiles | 30009

show class-of-service translation-table

IN THIS SECTION

- Syntax | 19151
- Description | 19151
- Options | 19151
- Required Privilege Level | 19152
- Output Fields | 19152
- Sample Output | 19153
- Release Information | 19157

Syntax

```
show class-of-service translation-table
<name translation-table-name> |
<type (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp | to-inet-precedence-
from-inet-precedence)>
```

Description

Display the mapping of class-of-service (CoS) translation table code points to corresponding bit patterns.

Options

none Display translation table code points for all translation tables.

name (Optional) Display information for the named translation table.

type (Optional) Display information for a certain translation table type:

to-dscp-from-dscp	Display DSCP translation table information.
to-dscp-ipv6-from-dscp-ipv6	Display DSCP IPv6 translation table information.
to-exp-from-exp	Display MPLS EXP translation table information.
to-inet-precedence-from-intet-precedence	Display Internet precedence translation table information.

Required Privilege Level

view

Output Fields

Table 396 on page 19152 describes the output fields for the `show class-of-service translation-table` command. Output fields are listed in the approximate order in which they appear.

Table 396: show class-of-service translation-table Output Fields

Field Name	Field Description
`Translation Table`	Name of the translation table.
`Translation table type`	Type of the translation table.
`Index`	Internal index number of the translation table.
`From Code Point`	Value of code point received.
`To Code Point`	Value of translated code point.

Sample Output

show class-of-service translation-table

```
user@host> show class-of-service translation-table
Translation Table: inet-trans-table, Translation table type: inet-to-inet, Index: 61075
  From Code point    To Code Point
  000                            101
  001                            111
  010                            101
  011                            111
  100                            101
  101                            101
  110                            001
  111                            000

Translation Table: dscp-trans-table, Translation table type: dscp-to-dscp, Index: 6761
  From Code point    To Code Point
  000000                      000111
  000001                      000111
  000010                      000111
  000011                      000111
  000100                      000111
  000101                      000111
  000110                      000111
  000111                      111000
  001000                      000111
  001001                      000111
  001010                      000111
  001011                      000111
  001100                      000111
  001101                      000111
  001110                      000111
  001111                      000111
  010000                      000111
  010001                      000111
  010010                      000111
  010011                      000111
  010100                      000111
  010101                      000111
  010110                      000111
  010111                      000111
```

```
011000                    000111
011001                    000111
011010                    000111
011011                    000111
011100                    000111
011101                    000111
011110                    000111
011111                    000111
100000                    000111
100001                    000111
100010                    000111
100011                    000111
100100                    000111
100101                    000111
100110                    000111
100111                    111000
101000                    000111
101001                    000111
101010                    000111
101011                    000111
101100                    000111
101101                    000111
101110                    000111
101111                    000111
110000                    000111
110001                    000111
110010                    000111
110011                    000111
110100                    000111
110101                    000111
110110                    000111
110111                    000111
111000                    000111
111001                    000111
111010                    000111
111011                    000111
111100                    000111
111101                    000111
111110                    000001
111111                    000000
```

show class-of-service translation-table name exp-trans-table

```
user@host> show class-of-service translation-table name exp-trans-table
Translation Table: exp-trans-table, Translation table type: exp-to-exp, Index: 9048
  From Code point    To Code Point
  000                            101
  001                            111
  010                            101
  011                            111
  100                            101
  101                            101
  110                            001
  111                            000
```

show class-of-service translation-table type to-dscp-ipv6-from-dscp-ipv6

```
user@host> show class-of-service translation-table type to-dscp-ipv6-from-dscp-ipv6
Translation Table: dscp-ipv6-trans-table, Translation table type: dscp-ipv6-to-dscp-ipv6, Index:
64704
  From Code point    To Code Point
  000000                      000111
  000001                      000111
  000010                      000111
  000011                      000111
  000100                      000111
  000101                      000111
  000110                      000111
  000111                      111000
  001000                      000111
  001001                      000111
  001010                      000111
  001011                      000111
  001100                      000111
  001101                      000111
  001110                      000111
  001111                      000111
  010000                      000111
  010001                      000111
  010010                      000111
```

```
010011                    000111
010100                    000111
010101                    000111
010110                    000111
010111                    000111
011000                    000111
011001                    000111
011010                    000111
011011                    000111
011100                    000111
011101                    000111
011110                    000111
011111                    000111
100000                    000111
100001                    000111
100010                    000111
100011                    000111
100100                    000111
100101                    000111
100110                    000111
100111                    111000
101000                    000111
101001                    000111
101010                    000111
101011                    000111
101100                    000111
101101                    000111
101110                    000111
101111                    000111
110000                    000111
110001                    000111
110010                    000111
110011                    000111
110100                    000111
110101                    000111
110110                    000111
110111                    000111
111000                    000111
111001                    000111
111010                    000111
111011                    000111
111100                    000111
111101                    000111
```

```
111110                    000001
111111                    000000
```

Release Information

Command introduced in Junos OS Release 9.3.

show cli

IN THIS SECTION

- Syntax | 19157
- Description | 19157
- Options | 19157
- Required Privilege Level | 19158
- Output Fields | 19158
- Sample Output | 19159
- Release Information | 19159

Syntax

```
show cli
```

Description

Display configured CLI settings.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 397 on page 19158 lists the output fields for the `show cli` command. Output fields are listed in the approximate order in which they appear.

Table 397: show cli Output Fields

Field Name	Field Description
`CLI complete-on-space`	Capability to complete a partial command entry when you type a space or a tab: `on` or `off`.
`CLI idle-timeout`	Maximum time that an individual session can be idle before the user is logged out from the router or switch. When this feature is enabled, the number of minutes is displayed. Otherwise, the state is `disabled`.
`CLI restart-on-upgrade`	CLI is set to prompt you to restart the router or switch after upgrading the software: `on` or `off`.
`CLI screen-length`	Number of lines of text that the terminal screen displays.
`CLI screen-width`	Number of characters in a line on the terminal screen.
`CLI terminal`	Terminal type.
`CLI is operating in`	Mode: `enhanced`.
`CLI timestamp`	Date and time format for the timestamp. If the timestamp is not set, the state is `disabled`.
`CLI working directory`	Pathname of the working directory.

Sample Output

show cli

```
user@host> show cli
CLI complete-on-space set to on
CLI idle-timeout disabled
CLI restart-on-upgrade set to on
CLI screen-length set to 47
CLI screen-width set to 132
CLI terminal is 'vt100'
CLI is operating in enhanced mode
CLI timestamp disabled
CLI working directory is '/var/tmp'
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show cli authorization | 19159
- show cli directory | 19164
- show cli history | 19166

show cli authorization

IN THIS SECTION

- Syntax | 19160
- Description | 19160
- Options | 19160
- Required Privilege Level | 19160
- Output Fields | 19160

Sample Output | 19163
Release Information | 19164

Syntax

```
show cli authorization
```

Description

Display the permissions for the current user.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 398 on page 19160 lists the output fields for the `show cli authorization` command. In the table, all possible permissions are displayed and output fields are listed in alphabetical order.

Table 398: show cli authorization Output Fields

Field Name	Field Description
`access`	Can view access configuration information.
`access-control`	Can modify access configuration.
`admin`	Can view user account information.

Table 398: show cli authorization Output Fields *(Continued)*

Field Name	Field Description
admin-control	Can modify user account information.
clear	Can clear learned network information.
configure	Can enter configuration mode.
control	Can modify any configuration.
edit	Can edit configuration files.
field	Reserved for field (debugging) support.
firewall	Can view firewall configuration information.
firewall-control	Can modify firewall configuration information.
floppy	Can read from and write to removable media.
flow-tap	Can view flow-tap configuration information.
flow-tap-control	Can configure flow-tap configuration information.
idp-profiler-operation	Can configure Profiler data.
interface	Can view interface configuration information.
interface-control	Can modify interface configuration information.
maintenance	Can perform system maintenance.

Table 398: show cli authorization Output Fields *(Continued)*

Field Name	Field Description
network	Can access the network by entering the ping, ssh, telnet, and traceroute commands.
pgcp-session-mirroring	Can view Packet Gateway Control Protocol session mirroring configuration.
pgcp-session-mirroring-control	Can modify Packet Gateway Control Protocol session mirroring configuration all-control.
reset	Can reset or restart interfaces and system processes.
rollback	Can roll back to previous configurations.
routing	Can view routing configuration information.
routing-control	Can modify routing configuration information.
secret	Can view passwords and authentication keys in the configuration.
secret-control	Can modify passwords and authentication keys in the configuration.
security	Can view security configuration information.
security-control	Can modify security configuration information.
shell	Can start a local shell.
snmp	Can view SNMP configuration information.
snmp-control	Can modify SNMP configuration information.

Table 398: show cli authorization Output Fields *(Continued)*

Field Name	Field Description
system	Can view system configuration information.
system-control	Can modify system configuration information.
trace	Can view trace file settings information.
trace-control	Can modify trace file settings information.
view	Can view current values and statistics.
view-configuration	Can view all configuration information (not including secrets).

Sample Output

show cli authorization

```
user@host> show cli authorization
Current user: 'remote' login: 'user' class ''
Permissions:
    admin       -- Can view user accounts
    admin-control-- Can modify user accounts
    clear       -- Can clear learned network information
    configure   -- Can enter configuration mode
    control     -- Can modify any configuration
    edit        -- Can edit full files
    field       -- Special for field (debug) support
    floppy      -- Can read and write from the floppy
    interface   -- Can view interface configuration
    interface-control-- Can modify interface configuration
    network     -- Can access the network
    reset       -- Can reset/restart interfaces and daemons
    routing     -- Can view routing configuration
    routing-control-- Can modify routing configuration
```

```
    shell       -- Can start a local shell
    snmp        -- Can view SNMP configuration
    snmp-control-- Can modify SNMP configuration
    system      -- Can view system configuration
    system-control-- Can modify system configuration
    trace       -- Can view trace file settings
    trace-control-- Can modify trace file settings
    view        -- Can view current values and statistics
    maintenance -- Can become the super-user
    firewall    -- Can view firewall configuration
    firewall-control-- Can modify firewall configuration
    secret      -- Can view secret configuration
    secret-control-- Can modify secret configuration
    rollback    -- Can rollback to previous configurations
    security    -- Can view security configuration
    security-control-- Can modify security configuration
    access      -- Can view access configuration
    access-control-- Can modify access configuration
    view-configuration-- Can view all configuration (not including secrets)
    flow-tap    -- Can view flow-tap configuration
    flow-tap-control-- Can configure flow-tap service
Individual command authorization:
    Allow regular expression: none
    Deny regular expression: none
    Allow configuration regular expression: none
    Deny configuration regular expression: none
```

Release Information

Command introduced before Junos OS Release 7.4.

show cli directory

IN THIS SECTION

- Syntax | 19165
- Description | 19165

- Options | 19165
- Required Privilege Level | 19165
- Release Information | 19165

Syntax

```
show cli directory
```

Description

Display the current working directory.

```
user@host> show cli directory
Current directory: /var/home/user
```

Options

This command has no options.

Required Privilege Level

view

Release Information

Command introduced before Junos OS Release 7.4.

show cli history

IN THIS SECTION

- Syntax | 19166
- Description | 19166
- Options | 19166
- Required Privilege Level | 19166
- Release Information | 19167

Syntax

```
show cli history
<count>
```

Description

Display a list of previous CLI commands.

```
user@host> show cli history
  11:14:14 -- show arp
  11:22:10 -- show cli authorization
  11:27:12 -- show cli history
```

Options

none Display all previous CLI commands.

count (Optional) Maximum number of commands to display.

Required Privilege Level

view

Release Information

Command introduced before Junos OS Release 7.4.

show configuration

IN THIS SECTION

- Syntax | 19167
- Description | 19167
- Options | 19167
- Additional Information | 19169
- Required Privilege Level | 19169
- Output Fields | 19169
- Sample Output | 19170
- Release Information | 19171

Syntax

```
show configuration
<statement-path>
```

Description

Display the configuration that currently is running on the router or switch, which is the last committed configuration.

Options

none Display the entire configuration.

statement-path

(Optional) Display one of the following hierarchies in a configuration. (Each *statement-path* option has additional suboptions not described here. See the appropriate user guide or EX Series switch documentation for more information.)

- `access`—Network access configuration.
- `access-profile`—Access profile configuration.
- `accounting-options`—Accounting data configuration.
- `applications`—Applications defined by protocol characteristics.
- `apply-groups`—Groups from which configuration data is inherited.
- `chassis`—Chassis configuration.
- `chassis network-services`—Current running mode.
- `class-of-service`—Class-of-service configuration.
- `diameter`—Diameter base protocol layer configuration.
- `ethernet-switching-options`—(EX Series switch only) Ethernet switching configuration.
- `event-options`—Event processing configuration.
- `firewall`—Firewall configuration.
- `forwarding-options`—Options that control packet sampling.
- `groups`—Configuration groups.
- `interfaces`—Interface configuration.
- `jsrc`—JSRC partition configuration.
- `jsrc-partition`—JSRC partition configuration.
- `logical-systems`—Logical system configuration.
- `poe`—(EX Series switch only) Power over Ethernet configuration.
- `policy-options`—Routing policy option configuration.
- `protocols`—Routing protocol configuration.
- `routing-instances`—Routing instance configuration.

- routing-options—Protocol-independent routing option configuration.
- security—Security configuration.
- services—Service PIC applications configuration.
- snmp—Simple Network Management Protocol configuration.
- system—System parameters configuration.
- virtual-chassis—(EX Series switch only) Virtual Chassis configuration.
- vlans—(EX Series switch only) VLAN configuration.

Additional Information

The portions of the configuration that you can view depend on the user class that you belong to and the corresponding permissions. If you do not have permission to view a portion of the configuration, the text ACCESS-DENIED is substituted for that portion of the configuration. If you do not have permission to view authentication keys and passwords in the configuration, because the secret permission bit is not set for your user account, the text SECRET-DATA is substituted for that portion of the configuration. If an identifier in the configuration contains a space, the identifier is displayed in quotation marks.

Likewise, when you issue the show configuration command with the | display set pipe option to view the configuration as set commands, those portions of the configuration that you do not have permissions to view are substituted with the text ACCESS-DENIED.

Required Privilege Level

view—To execute the command in the CLI mode. In addition to view, to display the contents under a configuration hierarchy, user should have the required privilege level to view the configuration under that configuration hierarchy.

Output Fields

This command displays information about the current running configuration.

Sample Output

show configuration

```
user@host> show configuration
## Last commit: 2006-10-31 14:13:00 PST by user1 version "8.2I0 [userb]"; ## last changed:
2006-10-31 14:05:53 PST
system {
    host-name exhost;
    domain-name ex1.net;
    backup-router 198.51.100.254;
   time-zone America/Los_Angeles;
    default-address-selection;
    name-server {
        192.0.2.254;
        192.0.2.249;
        192.0.2.176;
    }
    services {
        telnet;
    }
    tacplus-server {
        10.2.3.4 {
            secret /* SECRET-DATA */;
            ...
        }
    }
}
interfaces {
    ...
}
protocols {
        isis {
            export "direct routes";
        }
}
policy-options {
    policy-statement "direct routes" {
        from protocol direct;
        then accept;
```

```
    }
}
```

show configuration policy-options

```
user@host> show configuration policy-options
policy-options {
    policy-statement "direct routes" {
        from protocol direct;
        then accept;
    }
}
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Displaying the Current Configuration

Overview of Operational Mode Commands

show configuration chassis cluster traceoptions

IN THIS SECTION

- Syntax | 19172
- Description | 19172
- Required Privilege Level | 19172
- Output Fields | 19172
- Sample Output | 19172
- Release Information | 19173

Syntax

```
show configuration chassis cluster traceoptions
```

Description

Display tracing options for the chassis cluster redundancy process.

Required Privilege Level

view

Output Fields

Table 399 on page 19172 lists the output fields for the `show configuration chassis cluster traceoptions` command. Output fields are listed in the approximate order in which they appear.

Table 399: show configuration chassis cluster traceoptions Output Fields

Field Name	Field Description
`file`	Name of the file that receives the output of the tracing operation.
`size`	Size of each trace file.
`files`	Maximum number of trace files.

Sample Output

show configuration chassis cluster traceoptions

```
user@host> show configuration chassis cluster traceoptions
file chassis size 10k files 300;
level all;
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

traceoptions (Chassis Cluster) | 29506

show configuration services analytics sensor ddos

IN THIS SECTION

- Syntax | 19173
- Description | 19173
- Sample Output | 19173
- Release Information | 19174

Syntax

```
show configuration services analytics sensor ddos
```

Description

This command shows the DDOS configuration.

Sample Output

show configuration services analytics sensor ddos

```
user@host> show configuration services analytics sensor ddos
server-name sserver;
```

```
        export-name eprofile;
        resource /junos/system/linecard/ddos/;
    streaming-server sserver {
        remote-address 10.0.0.1;
        remote-port 2000;
    }
    export-profile eprofile {
        local-address 10.0.0.10;
        local-port 21111;
        reporting-rate 3000;
        format gpb;
    }
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

export-profile (Junos Telemetry Interface) | 6669

sensor (Junos Telemetry Interface) | 16564

streaming-server (Junos Telemetry Interface) | 28640

show configuration smtp

IN THIS SECTION

- Syntax | 19175
- Description | 19175
- Options | 19175
- Required Privilege Level | 19175
- Output Fields | 19175
- Sample Output | 19176
- Release Information | 19176

Syntax

```
show configuration smtp
```

Description

Display complete SMTP information.

Options

- apply-groups—Groups from which SMTP inherits configuration data.
- apply-groups-except—Groups from which SMTP restricts inheriting configuration data.

Required Privilege Level

view

Output Fields

Table 400 on page 19175 describes the output fields for the `show configuration smtp` command.

Table 400: show configuration smtp

Field Name	Field Description	Level of Output
address	SMTP server's IPv4 address	All levels
login	Configure a mail sender account to the server	All levels
password	Default sender password for user authentication	All levels

Sample Output

show configuration smtp

```
user@host> show configuration smtp
primary-server {
    address 218.102.48.213;
    login "dayone@example.com" {
        password "$ABC123"; ## SECRET-DATA
    }
}
```

Release Information

Command introduced in Junos OS Release 10.0 .

RELATED DOCUMENTATION

Content Security | 4589

show connections

IN THIS SECTION

- Syntax | 19177
- Syntax (EX Series Switches) | 19177
- Description | 19177
- Options | 19177
- Required Privilege Level | 19178
- Output Fields | 19178
- Sample Output | 19180
- Release Information | 19180

Syntax

```
show connections
<brief | extensive>
<all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch | remote-
interface-switch>
<down | up | up-down>
<history>
<labels>
<logical-system (all | logical-system-name)>
<name>
<status>
```

Syntax (EX Series Switches)

```
show connections
<brief | extensive>
<all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch | remote-
interface-switch>
<down | up | up-down>
<history>
<labels>
<name>
<status>
```

Description

Display information about the configured circuit cross-connect (CCC) connections.

Options

none	Display the standard level of output for all configured CCC connections.
all	(Optional) Display all connections.
brief | extensive	(Optional) Display the specified level of output. Use history to display information about connection history. Use labels to display labels used for transmit and receive LSPs. Use status to display information about the connection and interface status.

interface-switch	(Optional) Display interface switch connections only.
lsp-switch	(Optional) Display LSP switch connections only.
p2mp-receive-switch	(Optional) Display point-to-multipoint LSP to local interfaces switch connections only.
p2mp-transmit-switch	(Optional) Display local interface to point-to-multipoint LSP switch connections only.
remote-interface-switch	(Optional) Display remote interface switch connections only.
down | up | up-down	(Optional) Display nonoperational, operational, or both kinds of connections.
history	(Optional) Display information about connection history.
labels	(Optional) Display labels used for transmit and receive.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
name	(Optional) Display information about the specified connection only.
status	(Optional) Display information about the connection and interface status.

Required Privilege Level

view

Output Fields

Table 401 on page 19178 describes the output fields for the `show connections` command. Output fields are listed in the approximate order in which they appear.

Table 401: show connections Output Fields

Field Name	Field Description
`CCC and TCC connections [Link Monitoring On I Off]`	Whether link monitoring is enabled: **On** or **Off**.

Table 401: show connections Output Fields *(Continued)*

Field Name	Field Description
Legend for Status (St)	Connection or circuit status. See the output's legend for an explanation of the status field values.
Legend for connection types	Type of connection: • **if-sw**—Layer 2 switching cross-connect. • **rmt-if**—Remote interface switch. While graceful restart is in progress, **rmt-if** will display a state **(St)** of **Restart**. • **lsp-sw**—LSP stitching cross-connect. While graceful restart is in progress, **lsp-sw** will display a state **(St)** of **Restart**.
Legend for circuit types	Type of circuits: • **intf**—Interface circuit. • **tlsp**—Transmit LSP circuit. • **rlsp**—Receive LSP circuit.
Connection/Circuit	Name of the configured CCC connection.
Type	Type of connection.
St	State of the connection.
Time last up	Time that the connection or circuit last transitioned to the **Up** (operational) state.
# Up trans	Number of times that the connection or circuit has transitioned to the **Up** (operational) state.

Sample Output

show connections

```
user@switch> show connections
CCC and TCC connections [Link Monitoring On]
  Legend for status (St)              Legend for connection types
  UN -- uninitialized                 if-sw:  interface switching
  NP -- not present                   rmt-if: remote interface switching
  WE -- wrong encapsulation           lsp-sw: LSP switching
  DS -- disabled
  Dn -- down                          Legend for circuit types
  -> -- only outbound conn is up      intf -- interface
  <- -- only inbound  conn is up      tlsp -- transmit LSP
  Up -- operational                   rlsp -- receive LSP
  RmtDn -- remote CCC down
  Restart -- restarting

CCC Graceful restart : Restarting

Connection/Circuit                Type    St      Time last up     # Up trans
IFSW-ed                           if-sw   Up      Aug  5 15:39:15           1
  so-1/0/2.0                        intf  Up
  t1-0/1/2.0                        intf  Up
SW-db                             rmt-if  Restart                           0
  so-1/0/3.0                        intf  Up
  pro4-ca                           tlsp  Dn
  pro4-ac                           rlsp  NP
```

Release Information

Command introduced before Junos OS Release 7.4.

show cos halp flow queue-stats

IN THIS SECTION

- Syntax | 19181
- Description | 19181
- Options | 19181
- Required Privilege Level | 19181
- Output Fields | 19182
- Sample Output | 19182
- Release Information | 19183

Syntax

```
show cos halp flow queue-stats <slice-index>
```

Description

Show statistics for a slice. Note that statistics for a slice are shown at the PFE level. The <slice-index> is assigned to the slice. It is like the IFD-index which is dynamically assigned by the system. You can use “show class-of-service forwarding-table slice” CLI to get the slice-index value for a given slice name.

Options

- Number - display flow queue statistics for a slice index
- All - display all flow queue statistics

Required Privilege Level

view

Output Fields

Table 402 on page 19182 lists the output fields for the `show cos halp flow queue-stats <slice-index>` command.

Table 402: show cos halp flow queue-stats <slice-index> Output Fields

Field Name	Field Description
flow-index	The slice index
ifd-index	Index assigned to the physical interface

Sample Output

show cos halp flow queue-stats <slice-index>

```
show cos halp flow queue-stats 1
```

```
xyz4_vty# show cos halp flow queue-stats 1
HALP FLOW STATS:

flow-index: 1   ifd-index: 195

----------------------------------------
  FLOW: 1
  PFE: 0
 Queue: 0
  Last Enqueue time     : No packets
Queued              :
  Packets           :      455271897         491606 pps
  Bytes             :   191214196740      206474752 Bps
Transmitted         :
  Packets           :      130206172         148568 pps
  Bytes             :    54686592240       62398752 Bps
  Tail-dropped pkts :      325065725         343038 pps
  Tail-dropped bytes:   136527604500      144076000 Bps
  RED-dropped pkts  :
```

```
    Low             :              0              0 pps
    Medium-low      :              0              0 pps
    Medium-high     :              0              0 pps
    High            :              0              0 pps
  RED-dropped bytes :
    Low             :              0              0 Bps
    Medium-low      :              0              0 Bps
    Medium-high     :              0              0 Bps
    High            :              0              0 Bps
  RL-dropped pkts   :              0              0 pps
  RL-dropped bytes  :              0              0 Bps
  Queue depths      :
    Average         :       29884140
    Current         :       29884140
    Peak till now   :       29884140
    Maximum         :       30015488
```

Release Information

Junos 22.2R1

show database-replication statistics

IN THIS SECTION

- Syntax | 19184
- Description | 19184
- Options | 19184
- Required Privilege Level | 19184
- Output Fields | 19184
- Sample Output | 19185
- Release Information | 19185

Syntax

```
show database-replication statistics
```

Description

Display statistics regarding the replication of the subscriber management session database.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 403 on page 19184 lists the output fields for the `show database-replication statistics` command. Output fields are listed in the approximate order in which they appear.

Table 403: show database-replication statistics Output Fields

Field Name	Field Description
General	Number of dropped connections and the maximum buffer count.
Message Received	Total size of messages received and the number of received messages that have been processed.
Message Sent	Total size of messages sent and the number of sent messages that have been processed.
Message Queue	Number of messages in the queue and the maximum size of the queue.

Sample Output

show database-replication statistics

```
user@host> show database-replication statistics

General:
    Dropped connections        0
    Max buffer count           0
Message received:
    Size (bytes)               0
    Processed                  0
Message sent:
    Size (bytes)               0
    Processed                  0
Message queue:
    Queue full                 0
    Max queue size             0
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show database-replication summary | 19185

show database-replication summary

IN THIS SECTION

- Syntax | 19186
- Description | 19186
- Options | 19186

- Required Privilege Level | 19186
- Output Fields | 19186
- Sample Output | 19188
- Release Information | 19189

Syntax

```
show database-replication summary
```

Description

Display summary information regarding database replication for the subscriber management session database.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 404 on page 19187 lists the output fields for the `show database-replication summary` command. Output fields are listed in the approximate order in which they appear.

Table 404: show database-replication summary Output Fields

Field Name	Field Description
`Graceful Restart`	State of graceful Routing Engine switchover (GRES): • `Enabled` • `Disabled`
`Mastership`	State of the Routing Engine primary role: • `Master` • `Standby`
`Connection`	State of the connection: • `Up` • `Down`
`Disconnection Reason`	Reason the Routing Engines are disconnected. • RE DRAM Size Mismatch—Displayed when the amount of DRAM is different on the primary and standby Routing Engines.
`Database`	State of the subscriber management database: • `Available` • `Unavailable` • `Synchronized`

Table 404: show database-replication summary Output Fields *(Continued)*

Field Name	Field Description
`Message Queue`	State of the message queue: • `Full` • `Init` • `Not Ready` • `Ready`

Sample Output

show database-replication summary

```
user@host> show database-replication summary
General:
    Graceful Restart     Enabled
    Mastership           Standby
    Connection           Up
    Database             Available
    Message Queue        Ready
```

show database-replication summary (DRAM Size Mismatch Error)

```
user@host> show database-replication summary
General:
    Graceful Restart     Enabled
    Mastership           Standby
    Connection                 Down
    Disconnection Reason       RE DRAM Size Mismatch
    Database                   Available
    Message Queue              Not Ready
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show database-replication statistics | 19183

show dcbx

IN THIS SECTION

- Syntax | 19189
- Description | 19189
- Required Privilege Level | 19189
- Output Fields | 19190
- Sample Output | 19190
- Release Information | 19190

Syntax

```
show dcbx
```

Description

List DCBX status (enabled or disabled) and the interfaces on which DCBX is enabled.

Required Privilege Level

view

Output Fields

Table 405 on page 19190 lists the output fields for the `show dcbx` command. Output fields are listed in the approximate order in which they appear.

Table 405: show dcbx output fields

Field Name	Field Description
DCBX	Status of DCBX on the switch or for the specified interface: • Enabled—DCBX is enabled on the switch or on the specified interface • Disabled—DCBX is disabled on the switch or on the specified interface
Interface	Name of the interface

Sample Output

show dcbx

```
user@switch> show dcbx
DCBX                    : Enabled
Interface       DCBX
xe-0/0/9.0      enabled
xe-0/0/32.0     enabled
xe-0/0/36.0     enabled
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

show dcbx neighbors | **19191**

Configuring DCBX Autonegotiation

show dcbx neighbors

IN THIS SECTION

- Syntax | 19191
- Description | 19191
- Options | 19191
- Required Privilege Level | 19191
- Output Fields | 19192
- Sample Output | 19209
- Release Information | 19221

Syntax

```
show dcbx neighbors
<interface interface-name>
<terse>
```

Description

Display information about Data Center Bridging Capability Exchange protocol (DCBX) neighbor interfaces.

Options

none	Display information about all DCBX neighbor interfaces.
interface-name	(Optional) Display information for the specified interface.
terse	Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 406 on page 19192 lists the output fields for the `show dcbx neighbors` command. Output fields are listed in the approximate order in which they appear.

Table 406: show dcbx neighbors Output Fields

Field Name	Field Description
`Interface`	Name of the interface.
`Parent Interface`	Name of the link aggregation group (LAG) interface to which the DCBX interface belongs.
`Active-application-map`	Name of the application map applied to the interface.
`Protocol-Mode`	(QFX Series) DCBX protocol mode the interface uses: • IEEE DCBX Version—The interface uses IEEE DCBX mode. • DCBX Version 1.01—The interface uses DCBX version 1.01. **NOTE**: On interfaces that use the IEEE DCBX mode, the `show dcbx neighbors interface` *`interface-name`* operational command does not include application, PFC, or ETS operational state in the output.
`Protocol-State`	(DCBX Version 1.01 only) DCBX protocol state synchronization status: • `in-sync`—The local interface received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface. • `ack-pending`—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name		Field Description
Local-Advertisement		(DCBX Version 1.01 only) Status of advertisements that the local interface sends to the peer.
	Operational version	Version of the DCBX standard used.
	sequence-number	Number of state change messages sent to the peer. If the interface Protocol-State value is in-sync, this number should match the acknowledge-id number in the Peer-Advertisement section. If the interface Protocol-State value is ack-pending, this number does not match the acknowledge-id number in the Peer-Advertisement section.
	acknowledge-id	Number of acknowledge messages received from the peer. If the Protocol-State value is in-sync, this number should match the sequence-number value in the Peer-Advertisement section. If the Protocol-State value is ack-pending, this number does not match the sequence-number value in the Peer-Advertisement section.
Peer-Advertisement		(DCBX Version 1.01 only) Status of advertisements that the peer sends to the local interface.
	Operational version	Version of the DCBX standard used.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name		Field Description
	`sequence-number`	Number of state change messages the peer sent to the local interface. If this number matches the `acknowledge-id` number in the `Local-Advertisement` field, this indicates that the local interface has acknowledged all of the peer's state change messages and is synchronized. If this number does not match the `acknowledge-id` number in the `Local-Advertisement` field, this indicates that the peer has not yet received an acknowledgment for a state change message from the local interface.
	`acknowledge-id`	Number of acknowledge messages the peer has received from the local interface. If this number matches the `sequence-number` value in the `Local-Advertisement` field, this indicates that the peer has acknowledged all of the local interface's state change messages and is in synchronization. If this number does not match the `sequence-number` value in the `Local-Advertisement` field, this indicates that the peer has not yet sent an acknowledgment for a state change message from the local interface.
`Feature: PFC`		Priority-based flow control (PFC) feature DCBX state information.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
	`Protocol-State`		(DCBX Version 1.01 only) DCBX protocol state synchronization status: • `ack-pending`—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface. • `in-sync`—The local interface received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface. • `not-applicable`—PFC autonegotiation is disabled.
	`Operational State`		(DCBX Version 1.01 only) Operational state of the feature: `enabled` or `disabled`.
	`Local-Advertisement`		Status of advertisements that the local interface sends to the peer.
		`Enable`	(DCBX Version 1.01 only) State that the local interface advertises to the peer: • `Yes`—The feature is enabled. • `No`—The feature is disabled.
		`Willing`	Willingness of the local interface to learn the PFC configuration from the peer using DCBX: • `Yes`—The local interface is willing to learn the PFC configuration from the peer. • `No`—The local interface is not willing to learn the PFC configuration from the peer.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		`Mac auth Bypass Capability`	(IEEE DCBX only) (QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. This is not supported, so the only value seen in the local advertisement field is `no`.
		`Error`	(DCBX Version 1.01 only) Configuration compatibility error status: • `No`—No error detected. Local and peer configuration are compatible. • `Yes`—Error detected. Local and peer configuration are not compatible.
		`Operational State`	PFC operational state on the interface: • Enabled—PFC is enabled on the interface • Disabled—PFC is disabled on the interface
		`Maximum Traffic Classes capable to support PFC`	Largest number of traffic classes the local interface supports for PFC: • `6` (EX Series switches) • `6` (QFX Series)
		`Code Point`	PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.
		`Admin Mode`	PFC administrative state for each code point on the local interface: • `Enabled`—PFC is enabled for the code point. • `Disabled`—PFC is disabled for the code point.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		`Operational Mode`	(QFX Series) PFC operational mode for each code point: • Enable—PFC is enabled on the code point. • Disable—PFC is disabled on the code point.
	`Peer-Advertisement`		Status of advertisements that the peer sends to the local interface.
		`Enable`	(DCBX Version 1.01 only) State that the peer advertises to the local interface: • `Yes`—The feature is enabled. • `No`—The feature is disabled.
		`Willing`	Willingness of the peer to learn the PFC configuration from the local interface using DCBX: • `Yes`—The peer is willing to learn the PFC configuration from the local interface. • `No`—The peer is not willing to learn the PFC configuration from the local interface.
		`Error`	(DCBX Version 1.01 only) Configuration compatibility error status: • `No`—No error detected. Local and peer configuration are compatible. • `Yes`—Error detected. Local and peer configuration are not compatible.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		`Operational State`	PFC operational state on the interface: • Enabled—PFC is enabled on the interface • Disabled—PFC is disabled on the interface
		`Mac auth Bypass Capability`	(IEEE DCBX only) (QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. Although the QFX Series does not support this feature, the connected peer might support it. This field reports the peer state: • `Yes`—The connected peer supports MAC authentication bypass. • `No`—The connected peer does not support MAC authentication bypass.
		`Maximum Traffic Classes capable to support PFC`	Largest number of traffic classes the peer supports for PFC: • `6` (EX Series switches) • `8` (QFX Series)
		`Code Point`	PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.
		`Admin Mode`	PFC administrative state for each code point on the peer: • `Enabled`—PFC is enabled for the code point. • `Disabled`—PFC is disabled for the code point.
`Feature: Application`			State information for the DCBX application.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
	`Protocol-State`		(DCBX Version 1.01 only) DCBX protocol state synchronization status: • `in-sync`—The local interface received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface. • `ack-pending`—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface. • `not-applicable`—The local interface is set to `no-auto-negotiation` (autonegotiation is disabled). If the interface is associated with an FCoE forwarding class, the interface advertises FCoE capability even if the connected peer does not advertise FCoE capability.
	`Local-Advertisement`		Status of advertisements that the local interface sends to the peer. If the local interface is set to `no-auto-negotiation` (autonegotiation is disabled), the local advertisement portion of the output is not shown.
		`Enable`	(DCBX Version 1.01 only) State that the local interface advertises to the peer: • `Yes`—The feature is enabled. • `No`—The feature is disabled.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		`Willing`	(DCBX Version 1.01 only) Willingness of the local interface to learn the FCoE interface state from the peer using DCBX: • `Yes`—The local interface is willing to learn the FCoE interface state from the peer. • `No`—The local interface is not willing to learn the FCoE interface state from the peer.
		`Error`	(DCBX Version 1.01 only) Configuration compatibility error status: • `No`—No error detected. The local and peer configuration are compatible. • `Yes`—Error detected. The local and peer configuration are not compatible.
		`Appl-Name`	Name of the application:
		`Ethernet-Type`	(DCBX Version 1.01 only) Ethernet type (EtherType) of the application. For example, `0x8906` indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.
		`Socket-Number`	Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		`Priority-Field` or `Priority-Map`	Priority assigned to the application. For EX Series switches, the priority of the FCoE application is determined by the PFC congestion notification profile that has been configured and associated with the FCoE interface. For other applications, the priority is based on the application map.
		`Status`	(DCBX Version 1.01 only) Local status when autonegotiation is enabled: • `Enabled`—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.) • `Disabled`—The local configuration and the peer configuration do not match. **NOTE**: If there is a configuration mismatch in one application between the switch and the peer, all the other applications including FCoE are disabled.
	`Peer-Advertisement`		Status of advertisements that the peer sends to the local interface.
		`Enable`	(DCBX Version 1.01 only) State that the peer advertises to the local interface: • `Yes`—The feature is enabled. • `No`—The feature is disabled.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		`Willing`	(DCBX Version 1.01 only) Willingness of the peer to learn the FCoE interface state from the local interface using DCBX: • `Yes`—The peer is willing to learn the FCoE interface state from the local interface. • `No`—The peer is not willing to learn the FCoE interface state from the local interface.
		`Error`	(DCBX Version 1.01 only) Configuration compatibility error status: • `No`—No error detected. Local and peer configuration are compatible. • `Yes`—Error detected. Local and peer configuration are not compatible.
		`Appl-Name`	Name of the application: • `FCoE`—Fibre Channel over Ethernet
		`Ethernet-Type`	Ethernet type (EtherType) of the application. For example, `0x8906` indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket-Number (for Layer 4 applications) of the application is displayed in the output.
		`Socket-Number`	Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.
		`Priority-Field` or `Priority-Map`	Priority assigned to the application.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		`Status`	(DCBX Version 1.01 only) Peer interface status: • `Enabled`—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.) • `Disabled`—The local configuration and the peer configuration do not match.
`Feature: ETS`			Enhanced Transmission Selection (ETS) DCBX state information.
	`Protocol-State`		(DCBX Version 1.01 only) ETS protocol state synchronization status: • `in-sync`—The local interface received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface. • `ack-pending`—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.
	`Operational State`		(DCBX Version 1.01 only) Operational state of the feature, `enabled` or `disabled`.
	`Local-Advertisement`		Status of advertisements that the local interface sends to the peer.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		`Enable`	(DCBX Version 1.01 only) State that the local interface advertises to the peer: • `Yes`—The feature is enabled. • `No`—The feature is disabled.
		`TLV Type`	(IEEE DCBX only) Type of ETS TLV: • Configuration—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration. • Recommendation—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is "willing," configures the peer interface to match the local ETS configuration. • Recommendation-or-Configuration—Advertises both TLVs.
		`Willing`	Willingness of the local interface to learn the ETS state from the peer using DCBX (EX Series switches always advertise `No` for this field): • `Yes`—Local interface is willing to learn the ETS state from the peer. • `No`—Local interface is not willing to learn the ETS state from the peer.
		`Credit Based Shaper`	(IEEE DCBX only) Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always `No`.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		`Error`	(DCBX Version 1.01 only) Configuration error status: • `No`—No error. This should always be the switch ETS error state. • `Yes`—Error detected.
		`Maximum Traffic Classes capable to support PFC`	(DCBX Version 1.01 only) Largest number of traffic classes the local interface supports for PFC.
		`Maximum Traffic Classes supported`	(IEEE DCBX only) Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)
		`Code Point`	PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.
		`Priority-Group`	Class-of-service (CoS) priority group (forwarding class set) identification number.
		`Percentage B/W`	Configured minimum percentage of link bandwidth allocated to the priority group. Only explicitly configured values appear in this output column. If the link bandwidth is the default percentage, it is not shown. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		Transmission Selection Algorithm	(IEEE DCBX only) The transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is ETS.
	Peer-Advertisement		Status of advertisements that the peer sends to the local interface.
		Enable	(DCBX Version 1.01 only) State that the peer advertises to the local interface: • Yes—The feature is enabled. • No—The feature is disabled.
		TLV Type	(IEEE DCBX only) Type of ETS TLV: • Configuration—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration. • Recommendation—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration. • Configuration/Recommendation—Advertises both TLVs.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		`Willing`	Willingness of the peer to learn the ETS state from the local interface using DCBX: • `Yes`—Peer is willing to learn the ETS state from the local interface. • `No`—Peer is not willing to learn the ETS state from the local interface.
		`Credit Based Shaper`	(IEEE DCBX only) Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always `No`.
		`Error`	(DCBX Version 1.01 only) Configuration error status of the peer: • `No`—No error in peer ETS TLV. • `Yes`—Error in peer ETS TLV.
		`Maximum Traffic Classes capable to support PFC`	(DCBX Version 1.01 only) Largest number of traffic classes the local interface supports for PFC.
		`Maximum Traffic Classes supported`	(IEEE DCBX only) Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)
		`Code Point`	PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name			Field Description
		`Priority-Group`	CoS priority group (forwarding class set) identification number.
		`Percentage B/W`	Configured minimum percentage of link bandwidth allocated to the priority group. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)
		`Transmission Selection Algorithm`	(IEEE DCBX only) Transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is `ETS`.
`PFC`			(QFX Series, `terse` option only) DCBX TLV advertisement state for PFC: • Disabled—PFC configuration matches the configuration on the connected peer and PFC is disabled • Enabled—PFC configuration matches the configuration on the connected peer and PFC is enabled • Not Advt—Interface does not advertise PFC to the connected peer
`ETS`			(`terse` option only) Local DCBX TLV advertisement state for ETS: • Advt—Interface advertises ETS TLVs • Disabled—ETS is disabled on the interface (interface does not advertise ETS)

Table 406: show dcbx neighbors Output Fields *(Continued)*

Field Name	Field Description
ETS Rec	(terse option only) DCBX TLV peer advertisement state for ETS (state received from the connected DCBX peer): • Advt—Peer interface advertises ETS TLVs • Not Advt—Peer interface does not advertise ETS **NOTE**: When the DCBX mode is DCBX version 1.01, no peer information is displayed.
Version	(terse option only) The DCBX version used on the interface and whether the DCBX version was autonegotiated or explicitly configured: • IEEE—The interface uses IEEE DCBX. • 1.01—The interface uses DCBX version 1.01. When the DCBX version used is the result of autonegotiation, the term (Auto) appears next to the version. For example, IEEE (Auto) indicates that the interface autonegotiated with the connected peer to use IEEE DCBX. Autonegotiation is enabled by default.

Sample Output

show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode)

```
user@switch> show dcbx neighbors interface xe-0/0/0
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
    Active-application-map: app-map-1
    Protocol-State: in-sync
    Protocol-Mode: DCBX Version 1.01

    Local-Advertisement:
        Operational version: 1
        sequence-number: 130, acknowledge-id: 102
```

```
Peer-Advertisement:
    Operational version: 1
    sequence-number: 102, acknowledge-id: 130

Feature: PFC, Protocol-State: in-sync

  Operational State: Enabled

  Local-Advertisement:
    Enable: Yes, Willing: No, Error: No
    Maximum Traffic Classes capable to support PFC: 8

    Code Point           Admin Mode   Operational Mode
      000                 Disabled    Disable
      001                 Disabled    Disable
      010                 Disabled    Disable
      011                 Enabled     Enable
      100                 Enabled     Enable
      101                 Disabled    Disable
      110                 Disabled    Disable
      111                 Disabled    Disable

  Peer-Advertisement:
    Enable: Yes, Willing: No, Error: No
    Maximum Traffic Classes capable to support PFC: 8

    Code Point           Admin Mode
      000                 Disabled
      001                 Disabled
      010                 Disabled
      011                 Enabled
      100                 Enabled
      101                 Disabled
      110                 Disabled
      111                 Disabled

Feature: Application, Protocol-State: in-sync

  Local-Advertisement:
    Enable: Yes, Willing: No, Error: No

    Appl-Name    Ethernet-Type    Socket-Number    Priority-Map    Status
      FCoE         0x8906                            00001110        Enabled
```

```
        iSCSI                           3260           10000000        Enabled

    Peer-Advertisement:
      Enable: Yes, Willing: Yes, Error: No

      Appl-Name    Ethernet-Type    Socket-Number    Priority-Map    Status
        FCoE         0x8906             N/A            00001110        Enabled

  Feature: ETS, Protocol-State: in-sync

    Operational State: Enabled

    Local-Advertisement:
      Enable: Yes, Willing: No, Error: No
      Maximum Traffic Classes capable to support PFC: 8

      Code Point           Priority-Group
        000                  0
        001                  7
        010                  7
        011                  7
        100                  0
        101                  1
        110                  1
        111                  7

      Priority-Group       Percentage B/W
        0                    40%
        1                    5%

    Peer-Advertisement:
      Enable: Yes, Willing: No, Error: No
      Maximum Traffic Classes capable to support PFC: 8

      Code Point           Priority-Group
        000                  0
        001                  7
        010                  7
        011                  7
        100                  0
        101                  1
        110                  1
        111                  7
```

```
        Priority-Group       Percentage B/W
          0                    40%
          1                    5%
```

show dcbx neighbors interface (QFX Series, IEEE DCBX Mode)

```
user@switch> show dcbx neighbors interface xe-0/0/0
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
    Active-application-map: app-map-1
    Protocol-Mode: IEEE-DCBX Version

    Feature: PFC

      Local-Advertisement:
        Willing: No
        Mac auth Bypass Capability: No
        Operational State: Enabled

        Maximum Traffic Classes capable to support PFC: 8

        Code Point           Admin Mode
          000                 Disabled
          001                 Disabled
          010                 Disabled
          011                 Enabled
          100                 Enabled
          101                 Disabled
          110                 Disabled
          111                 Disabled

      Peer-Advertisement:
        Willing: No
        Mac auth Bypass Capability: No
        Operational State: Enabled

        Maximum Traffic Classes capable to support PFC: 8

        Code Point          Admin Mode
```

```
        000                 Disabled
        001                 Disabled
        010                 Disabled
        011                 Enabled
        100                 Enabled
        101                 Disabled
        110                 Disabled
        111                 Disabled

    Feature: Application

      Local-Advertisement:

        Appl-Name    Ethernet-Type    Socket-Number    Priority-field
          FCoE         0x8906                            00001110
          iSCSI                            3260          10000000

      Peer-Advertisement:
        Appl-Name    Ethernet-Type    Socket-Number    Priority-field
          FCoE         0x8906             N/A            00001110

    Feature: ETS

      Local-Advertisement:
        TLV Type: Configuration/Recommendation
        Willing: No
        Credit Based Shaper: No
        Maximum Traffic Classes supported: 3

        Code Point           Priority-Group
          000                  0
          001                  7
          010                  7
          011                  7
          100                  0
          101                  1
          110                  1
          111                  7

        Priority-Group       Percentage B/W
          0                    40%
          1                    5%
```

```
    Priority-Group   Transmission Selection Algorithm
      0              Enhanced Transmission Selection
      1              Enhanced Transmission Selection

  Peer-Advertisement:
    TLV Type: Configuration
    Willing: No
    Credit Based Shaper: No

    Code Point           Priority-Group
      000                  0
      001                  7
      010                  7
      011                  7
      100                  0
      101                  1
      110                  1
      111                  7

    Priority-Group       Percentage B/W
      0                    40%
      1                    5%

    Priority-Group   Transmission Selection Algorithm
      0              Enhanced Transmission Selection
      1              Enhanced Transmission Selection

  Peer-Advertisement:

    TLV Type: Recommendation

    Code Point           Priority-Group
      000                  0
      001                  7
      010                  7
      011                  7
      100                  0
      101                  1
      110                  1
      111                  7

    Priority-Group       Percentage B/W
      0                    40%
```

```
          1                    5%

        Priority-Group   Transmission Selection Algorithm
          0              Enhanced Transmission Selection
          1              Enhanced Transmission Selection
```

show dcbx neighbors terse (QFX Series)

```
user@switch> show dcbx neighbors terse
Interface   Parent      PFC        ETS        ETS        Version
            Interface                                    Rec
xe-0/0/8.0  -           Enabled    Advt       Advt       IEEE (Auto)
xe-0/0/9.0  -           Disabled   Disabled              1.01
xe-0/0/11.0 ae0.0       Enabled    Advt       Advt       IEEE (Auto)
xe-0/0/12.0 ae0.0       Enabled    Advt       Advt       IEEE (Auto)
xe-0/0/32.0 -           Enabled    Advt       Not Advt   IEEE
xe-0/0/36.0 -           Not Advt   Advt       Advt       IEEE
```

show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly)

```
user@switch> show dcbx neighbors interface xe-0/0/14

   Interface : xe-0/0/14.0 - Parent Interface: ae0.0
    Protocol-State: in-sync

    Local-Advertisement:
        Operational version: 0
        sequence-number: 6, acknowledge-id: 6

    Peer-Advertisement:
        Operational version: 0
        sequence-number: 6, acknowledge-id: 6

    Feature: PFC, Protocol-State: in-sync
```

```
      Operational State: Enabled

      Local-Advertisement:
        Enable: Yes, Willing: No, Error: No
        Maximum Traffic Classes capable to support PFC: 6

         Code Point            Admin Mode
          000                 Disabled
          001                 Disabled
          010                 Disabled
          011                 Enabled
          100                 Disabled
          101                 Disabled
          110                 Disabled
          111                 Disabled



      Peer-Advertisement:
        Enable: Yes, Willing: No, Error: No
        Maximum Traffic Classes capable to support PFC: 6

        Code Point             Admin Mode
          000                 Disabled
          001                 Disabled
          010                 Disabled
          011                 Enabled
          100                 Disabled
          101                 Disabled
          110                 Disabled
          111                 Disabled



    Feature: Application, Protocol-State: in-sync

      Local-Advertisement:
        Enable: Yes, Willing: No, Error: No   <<< Error bit will not be set as there is no miss
configuration between local and peer.


        Appl-Name    Ethernet-Type    Socket-Number    Priority-Map    Status
        FCoE          0x8906                            00001000        Enabled
```

```
      Peer-Advertisement:
        Enable: Yes, Willing: No, Error: No

                 Appl-Name    Ethernet-Type    Socket-Number    Priority-Map    Status
                 FCoE          0x8906                           00001000        Enabled
```

show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application)

```
user@switch> show dcbx neighbors interface xe-0/0/14

  Interface : xe-0/0/14.0 - Parent Interface: ae0.0

    Protocol-State: in-sync

    Active-application-map: iscsi-map

     Local-Advertisement:
        Operational version: 0
        sequence-number: 9, acknowledge-id: 12


    Peer-Advertisement:
        Operational version: 0
        sequence-number: 12, acknowledge-id: 9


    Feature: PFC, Protocol-State: in-sync
      Operational State: Enabled
       Local-Advertisement:
        Enable: Yes, Willing: No, Error: No
        Maximum Traffic Classes capable to support PFC: 6


     Code Point              Admin Mode
          000                 Disabled
          001                 Disabled
```

```
     010                 Disabled
     011                 Enabled
     100                 Disabled
     101                 Disabled
     110                 Disabled
     111                 Disabled

 Peer-Advertisement:
   Enable: Yes, Willing: No, Error: No
   Maximum Traffic Classes capable to support PFC: 6

Code Point              Admin Mode
     000                 Disabled
     001                 Disabled
     010                 Disabled
     011                 Enabled
     100                 Disabled
     101                 Disabled
     110                 Disabled
     111                 Disabled

Feature: Application, Protocol-State: in-sync
 Local-Advertisement:
   Enable: Yes, Willing: No, Error: No

        Appl-Name    Ethernet-Type    Socket-Number   Priority-Map    Status
         FCoE         0x8906                           00001000        Enabled
         iscsi                          3260           00100000        Enabled


 Peer-Advertisement:
   Enable: Yes, Willing: No, Error: No

        Appl-Name    Ethernet-Type    Socket-Number   Priority-Map    Status
         FCoE         0x8906                           00001000        Enabled
         iscsi                          3260           00100000        Enabled
```

show dcbx neighbors (EX4500 Switch: Includes ETS)

```
user@switch> show dcbx neighbors interface xe-0/0/3

   Interface : xe-0/0/3.0
    Protocol-State: in-sync
    Active-application-map: map_iscsi

    Local-Advertisement:
        Operational version: 0
        sequence-number: 1, acknowledge-id: 5

    Peer-Advertisement:
        Operational version: 0
        sequence-number: 5, acknowledge-id: 1

    Feature: PFC, Protocol-State: in-sync

      Operational State: Enabled

      Local-Advertisement:
        Enable: Yes, Willing: No, Error: No
        Maximum Traffic Classes capable to support PFC: 6

        Code Point          Admin Mode
          000                 Enabled
          001                 Enabled
          010                 Disabled
          011                 Disabled
          100                 Disabled
          101                 Disabled
          110                 Disabled
          111                 Disabled

      Peer-Advertisement:
        Enable: Yes, Willing: Yes, Error: No
        Maximum Traffic Classes capable to support PFC: 8

        Code Point          Admin Mode
          000                 Enabled
          001                 Disabled
          010                 Disabled
```

```
        011                 Disabled
        100                 Enabled
        101                 Disabled
        110                 Disabled
        111                 Disabled

    Feature: Application, Protocol-State: in-sync

      Local-Advertisement:
        Enable: Yes, Willing: No, Error: No

        Appl-Name    Ethernet-Type    Socket-Number    Priority-Map    Status
          FCoE         0x8906                            00000001        Enabled
          iscsi                          3260            00000010        Enabled

      Peer-Advertisement:
        Enable: Yes, Willing: Yes, Error: No

        Appl-Name    Ethernet-Type    Socket-Number    Priority-Map    Status
          FCoE         0x8906                            00001000        Enabled
          iscsi                          3260            00010000        Enabled

    Feature: ETS, Protocol-State: in-sync

      Operational State: Enabled

      Local-Advertisement:
        Enable: Yes, Willing: No, Error: No
        Maximum Traffic Classes supported : 3

        Code Point           Priority-Group
          000                  7
          001                  7
          010                  7
          011                  7
          100                  7
          101                  7
          110                  7
          111                  7

        Priority-Group       Percentage B/W
          7                    100%
```

```
      Peer-Advertisement:
        Enable: Yes, Willing: Yes, Error: No
        Maximum Traffic Classes supported : 8

        Code Point           Priority-Group
          000                  0
          001                  1
          010                  0
          011                  0
          100                  2
          101                  0
          110                  0
          111                  0

        Priority-Group       Percentage B/W
          0                    30%
          1                    40%
          2                    30%
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- *Configuring DCBX Autonegotiation*
- *Example: Configuring DCBX Application Protocol TLV Exchange*
- Example: Configuring an FCoE Transit Switch
- Example: Configuring DCBX to Support an iSCSI Application
- *Understanding DCB Features and Requirements*
- Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches
- dcbx | **4858**

show ddos-protection protocols

IN THIS SECTION

- Syntax | 19222
- Description | 19222
- Options | 19222
- Required Privilege Level | 19223
- Output Fields | 19223
- Sample Output | 19230
- Release Information | 19239

Syntax

```
show ddos-protection protocols <protocol-group (aggregate | packet-type)>
```

Description

Display control plane DDoS protection configuration and statistics for supported protocol groups or individual packet types.

Options

none	Display information for all packet types in all protocol groups.
aggregate	(Optional) Display control plane DDoS protection information for the aggregate policer. The `aggregate` option is available for all supported protocol groups.
packet-type	(Optional) Display control plane DDoS protection information for the specified packet type in the specified protocol group. The available packet types vary by protocol group, and only some protocol groups can have policers for individual packet types.
protocol-group	(Optional) Display control plane DDoS protection information for a protocol group.

See the following configuration statements for the list of available *protocol-group* and *packet-type* options on different devices that you can use with this command, which are the same as the supported options you use to change default policer configurations:

- For routing devices except PTX Series routers, see "protocols (DDoS)" on page 13751.
- For PTX Series routers and QFX Series switches, see "protocols (DDoS) (ACX Series, PTX Series, and QFX Series)" on page 13765.

Required Privilege Level

view

Output Fields

Table 407 on page 19223 lists the output fields for the `show ddos-protection protocols` command. Output fields are listed in the approximate order in which they appear.

Table 407: show ddos-protection protocols Output Fields

Field Name	Field Description
`Packet types`	Number of packet types
`Modified`	Number of packets for which policer values have been modified from the default.
`Received traffic`	Number of traffic flows received.
`Currently violated`	Number of flows that are currently violating the flow bandwidth limit.
`Currently tracked flows`	Number of active flows that are being tracked as culprit flows by flow detection.
`Total detected flows`	Total number of culprit flows that have been detected, including those that have recovered or timed out.
`Protocol Group`	Name of protocol group.

Table 407: show ddos-protection protocols Output Fields *(Continued)*

Field Name	Field Description
`Packet type`	Name of packet type in protocol group.
`Bandwidth`	Bandwidth policer value; number of packets per second that is allowed before a violation is declared.
`Burst`	Burst policer value; the maximum number of packets that is allowed in a burst before a violation is declared.
`Priority`	Priority of the packet type for individual packet policers that enables more important traffic to pass through in the event of traffic congestion: `low`, `medium`, or `high`. Lower priority packets can be dropped when insufficient bandwidth is available.
`Recover time`	Time that must pass since the last violation before the traffic flow is considered to have recovered from the attack. A notification is generated when the timer expires.
`Enabled`	State of the policer: • `Yes`—The policer is enabled on both the Routing Engine and the FPC (line card). This is the default state. • `No`—The policer is disabled on both the Routing Engine and the FPC by global configuration. It is not disabled by the packet type level configuration. • `No*`—The policer is disabled on both the Routing Engine and the FPC. The asterisk (*) indicates that one or both of these instances is disabled at the packet type level; it may also be disabled globally. • `Partial`—The policer is disabled on either the Routing Engine or the FPC, but not both. It is disabled by global configuration. It is not disabled by the packet type level configuration. • `Partial*`—The policer is disabled on either the Routing Engine or the FPC, but not both. The asterisk (*) indicates that the instance is disabled by the packet type level configuration; it may also be disabled globally. Disabling can occur globally for all packet types at the `[edit system ddos-protection global]` hierarchy level, for a specific packet type at the `[edit system ddos-protection protocols` *`protocol-group`* `(aggregate | `*`packet-type`*`]` hierarchy level, or at both levels.

Table 407: show ddos-protection protocols Output Fields *(Continued)*

Field Name	Field Description
`Bypass aggregate`	State of the bypass aggregate configuration: • Yes—The aggregate policer is bypassed. • No—The aggregate policer is enforced. This field appears only for individual policers.

`Flow detection configuration`	State of flow detection configured on the router:
	• Detection mode—Mode of operation for suspicious flow detection: automatic, off, or on.
	• Log flows—State of automatic logging of suspicious traffic flows: on (`Yes`) or off (`No`).
	• Timeout flows—State of culprit flow timeout behavior: flow is suppressed for a configured timeout period (`Yes`) or flow is suppressed until it is no longer in violation (`No`).
	• Detect time—Time in seconds that must pass before a suspicious flow that has exceeded the bandwidth allowed for the packet type is considered to be a culprit flow.
	• Recover time—Time in seconds that must pass before a culprit flow is considered to have returned to normal. The period starts when the flow drops below the threshold that triggered the last violation.
	• Timeout time—Time in seconds that a culprit flow is suppressed, if timeouts have been enabled.
	• Flow aggregation level configuration—Flow detection mode, flow control mode, and flow bandwidth for traffic at each of the traffic flow aggregation levels: subscriber, logical interface, and physical interface.
	  • Detection mode—State of flow detection: automatic, off, or on.
	  Control mode—Mode of controlling culprit traffic: dropped, kept, or policed back to within the allowed bandwidth.
	  Flow rate—Bandwidth allowed for the control traffic in packets per second.

Table 407: show ddos-protection protocols Output Fields *(Continued)*

Field Name	Field Description
`System-wide information`	The following information collected for the router: • A message indicates whether the policer has been violated. • No. of FPCs currently receiving excess traffic—Number of cards that are currently in violation of a policer. • No. of FPCs that have received excess traffic—Number of cards that have at some point been in violation of a policer. • Violation first detected at—Timestamp of the first violation. • Violation last seen at—Timestamp of the last observed violation. • Duration of violation—Length of the violation. • Number of violations—Number of times the violation has occurred. • Received—Number of packets received at all card slots and the Routing Engine. • Dropped—Number of packets dropped regardless of where they were dropped. • Arrival rate—Current traffic rate for packets arriving from all cards and at the Routing Engine. • Max arrival rate—Highest traffic rate for packets arriving from all cards and at the Routing Engine.

Table 407: show ddos-protection protocols Output Fields *(Continued)*

Field Name	Field Description

`Routing Engine information`	The following information collected for the Routing Engine:
	• Bandwidth—Maximum number of packets per second that is allowed.
	• Burst—Maximum number of packets that is allowed in a burst.
	• State of the policer:
	  • `enabled`—The Routing Engine policer is enabled. This is the default state.
	  • `disabled`—The Routing Engine policer is disabled globally. It is not disabled by the packet type level configuration.
	  • `disabled*`—The Routing Engine policer is disabled by the packet type level configuration; it may also be disabled globally.
	• A message indicates whether the policer has been violated; the policer might be passed at the individual cards, but the combined rate of packets arriving at the Routing Engine can exceed the configured policer value.
	• Violation first detected at—Timestamp of the first violation.
	• Violation last seen at—Timestamp of the last observed violation.
	• Duration of violation—Length of the violation.
	• Number of violations—Number of times the violation has occurred.
	• Received—Number of packets received at the Routing Engine from all cards.
	• Dropped—Number of packets dropped at the Routing Engine; includes packets dropped by the aggregate policer and by individual protocol policers.
	• Arrival rate—Current traffic rate for packets arriving at the Routing Engine from all cards.
	• Max arrival rate—Highest traffic rate for packets arriving at the Routing Engine from all cards.
	• Dropped by aggregate policer—Number of packets dropped by the aggregate policer.
	• Dropped by individual policers—Number of packets dropped by individual policer.

Table 407: show ddos-protection protocols Output Fields *(Continued)*

Field Name	Field Description

`FPC slot information`	The following information collected for the card in the indicated slot:
	• Bandwidth—Bandwidth scaling percentage and the number of packets per second that is allowed before a violation is declared.
	• Burst—Burst scaling percentage and the maximum number of packets that is allowed in a burst before a violation is declared.
	• State of the policer:
	  • `enabled`—The FPC policer is enabled. This is the default state.
	  • `disabled`—The FPC policer is disabled globally. It is not disabled by the packet type level configuration.
	  • `disabled*`—The FPC policer is disabled by the packet type level configuration; it may also be disabled globally.
	• A message indicates whether the policer has been violated.
	• Violation first detected at—Timestamp of the first violation.
	• Violation last seen at—Timestamp of the last observed violation.
	• Duration of violation—Length of the violation.
	• Number of violations—Number of times the violation has occurred.
	• Received—Number of packets received on the line card.
	• Dropped—Number of packets dropped at the line card; includes packets dropped by the aggregate policer and by individual protocol policers.
	• Arrival rate—Current traffic rate for packets arriving at the line card.
	• Max arrival rate—Highest traffic rate for packets arriving at the line card.
	• Dropped by this policer—Number of packets dropped by the individual policer.
	• Dropped by aggregate policer—Number of packets dropped by the aggregate policer.
	NOTE:

Table 407: show ddos-protection protocols Output Fields *(Continued)*

Field Name	Field Description
	• On MX Series routers with built-in MPCs—the MX5, MX10, MX40, MX80, and MX104 routers—this field actually displays information for tfeb0 because these routers have no Flexible PIC Concentrator (FPC) slots. Instead, the Packet Forwarding Engine has two "pseudo" FPCs (FPC 0 and FPC1).
`Bypass aggr.`	State of the bypass aggregate configuration: • Yes—The aggregate policer configuration is bypassed. • No—The aggregate policer configuration is enforced. Dashes indicate that the bypass aggregate configuration is not available; this is possible only for aggregate policers.
`FPC Mod`	Indicates whether configuration has changed from the default for any line cards. • No—The default configuration has not changed from the default for the packet type. • Yes—The default configuration has changed from the default for the packet type
`Op mode`	Mode of operation for suspicious flow detection for the packet type: always-on (`on`), (`auto`), or disabled (`off`).
`Policer BW (pps)`	Bandwidth policer value; number of packets per second that is allowed before a violation is declared.
`Aggr level Op:Fc:Bwidth (pps)`	Flow operation mode, flow control mode, and flow bandwidth for traffic of the packet type at each traffic flow aggregation level: subscriber (`sub`), logical interface (`ifl`), and physical interface (`ifd`).
`Log flow`	State of automatic logging of suspicious traffic flows for the packet type: on (`Yes`) or off (`No`).
`Time out`	State of culprit flow timeout behavior for the packet type: flow is suppressed or monitored for a configured timeout period (`Yes`) or flow is suppressed or monitored until it is no longer in violation (`No`).

Sample Output

show ddos-protection protocols

```
user@host> show ddos-protection protocols

Packet types: 190, Modified: 0, Received traffic: 12, Currently violated: 3
Currently tracked flows: 0, Total detected flows: 0
* = User configured value

Protocol Group: IPv4-Unclassified

  Packet type: aggregate (Aggregate for unclassified host-bound IPv4 traffic)
    Aggregate policer configuration:
      Bandwidth:        2000 pps
      Burst:            10000 packets
      Recover time:     300 seconds
      Enabled:          Yes
    Flow detection configuration:
      Detection mode: Automatic  Detect time:  3 seconds
      Log flows:      No         Recover time: 60 seconds
      Timeout flows:  No         Timeout time: 300 seconds
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          10 pps
        Logical interface   Automatic       Drop          10 pps
        Physical interface  Automatic       Drop          2000 pps
    System-wide information:
      Aggregate bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Bandwidth: 2000 pps, Burst: 10000 packets, enabled
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by individual policers: 0
    FPC slot 1 information:
      Bandwidth: 100% (2000 pps), Burst: 100% (10000 packets), enabled
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
```

```
        Dropped by individual policers: 0
        Dropped by flow suppression:    0

…

Protocol Group: PPPoE

  Packet type: aggregate (Aggregate for all PPPoE control traffic)
    Aggregate policer configuration:
      Bandwidth:        2000 pps
      Burst:            2000 packets
      Recover time:     300 seconds
      Enabled:          Yes
    Flow detection configuration:
      Detection mode: Automatic  Detect time:  3 seconds
      Log flows:      No         Recover time: 60 seconds
      Timeout flows:  No         Timeout time: 300 seconds
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          10 pps
        Logical interface   Automatic       Drop          10 pps
        Physical interface  Automatic       Drop          2000 pps
    System-wide information:
      Aggregate bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Bandwidth: 2000 pps, Burst: 2000 packets, enabled
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by individual policers: 0
    FPC slot 1 information:
      Bandwidth: 100% (2000 pps), Burst: 100% (2000 packets), enabled
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by individual policers: 0
        Dropped by flow suppression:    0

  Packet type: padi (PPPoE PADI)
    Individual policer configuration:
      Bandwidth:        500 pps
```

```
      Burst:            500 packets
      Priority:         Low
      Recover time:     300 seconds
      Enabled:          Yes
      Bypass aggregate: No
    Flow detection configuration:
      Detection mode: Automatic  Detect time:  3 seconds
      Log flows:      No         Recover time: 60 seconds
      Timeout flows:  No         Timeout time: 300 seconds
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          10 pps
        Logical interface   Automatic       Drop          10 pps
        Physical interface  Automatic       Drop          500 pps
    System-wide information:
      Bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Bandwidth: 500 pps, Burst: 500 packets, enabled
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Bandwidth: 100% (500 pps), Burst: 100% (500 packets), enabled
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
        Dropped by flow suppression:  0
...
```

show ddos-protection protocols (Specific Packet Type with Flow Detection Disabled)

```
user@host> show ddos-protection protocols pppoe padi
Currently tracked flows: 0, Total detected flows: 0
* = User configured value

Protocol Group: PPPoE
```

```
  Packet type: padi (PPPoE PADI)
    Individual policer configuration:
      Bandwidth:        500 pps
      Burst:            500 packets
      Priority:         Low
      Recover time:     300 seconds
      Enabled:          Yes
      Bypass aggregate: No
    Flow detection configuration:
      Detection mode: Off*       Detect time:  3 seconds
      Log flows:      No         Recover time: 60 seconds
      Timeout flows:  No         Timeout time: 300 seconds
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          10 pps
        Logical interface   Automatic       Drop          10 pps
        Physical interface  Automatic       Drop          500 pps
    System-wide information:
      Bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Bandwidth: 500 pps, Burst: 500 packets, enabled
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Bandwidth: 100% (500 pps), Burst: 100% (500 packets), enabled
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
        Dropped by flow suppression:  0
```

show ddos-protection protocols (Specific Packet Type with Flow Detection Enabled and Automatic)

```
user@host> show ddos-protection protocols pppoe padi
Currently tracked flows: 0, Total detected flows: 0
* = User configured value
```

```
Protocol Group: PPPoE

  Packet type: padi (PPPoE PADI)
    Individual policer configuration:
      Bandwidth:        500 pps
      Burst:            500 packets
      Priority:         Low
      Recover time:     300 seconds
      Enabled:          Yes
      Bypass aggregate: No
    Flow detection configuration:
      Detection mode: Automatic  Detect time:  3 seconds
      Log flows:      No         Recover time: 60 seconds
      Timeout flows:  No         Timeout time: 300 seconds
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          10 pps
        Logical interface   Automatic       Drop          10 pps
        Physical interface  Automatic       Drop          500 pps
    System-wide information:
      Bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Bandwidth: 500 pps, Burst: 500 packets, enabled
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Bandwidth: 100% (500 pps), Burst: 100% (500 packets), enabled
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
        Dropped by flow suppression:  0
```

show ddos-protection protocols (Specific Packet Type with Bandwidth Violation)

```
user@host> show ddos-protection protocols bfd
Packet types: 1, Modified: 0, Received traffic: 1, Currently violated: 1
Currently tracked flows: 1, Total detected flows: 1
* = User configured value

Protocol Group: BFD

  Packet type: aggregate (Aggregate for all bfd traffic)
    Aggregate policer configuration:
      Bandwidth:        20000 pps
      Burst:            20000 packets
      Recover time:     300 seconds
      Enabled:          Yes
    Flow detection configuration:
      Detection mode: Automatic  Detect time:  3 seconds
      Log flows:      No         Recover time: 60 seconds
      Timeout flows:  No         Timeout time: 300 seconds
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          10 pps
        Logical interface   Automatic       Drop          10 pps
        Physical interface  Automatic       Drop          20000 pps
    System-wide information:
      Aggregate bandwidth is being violated!
        No. of FPCs currently receiving excess traffic: 1
        No. of FPCs that have received excess traffic:  1
        Violation first detected at: 2012-10-24 23:40:20 EDT
        Violation last seen at:      2012-10-25 10:25:48 EDT
        Duration of violation: 10:45:28 Number of violations: 1
      Received:  1173471731          Arrival rate:     30304 pps
      Dropped:   399135607           Max arrival rate: 30331 pps
      Flow counts:
        Aggregation level     Current       Total detected
        Subscriber            1             1
        Total                 1             1
    Routing Engine information:
      Bandwidth: 20000 pps, Burst: 20000 packets, enabled
      Aggregate policer is never violated
      Received:  366831604           Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 9522 pps
```

```
        Dropped by individual policers: 0
    FPC slot 1 information:
      Bandwidth: 100% (20000 pps), Burst: 100% (20000 packets), enabled
      Aggregate policer is currently being violated!
        Violation first detected at: 2012-10-24 23:40:21 EDT
        Violation last seen at:      2012-10-25 10:25:48 EDT
        Duration of violation: 10:45:27 Number of violations: 1
      Received:  1173471731          Arrival rate:     30304 pps
      Dropped:   399135607           Max arrival rate: 30331 pps
        Dropped by individual policers: 0
        Dropped by aggregate policer:   398854530
        Dropped by flow suppression:    281077
      Flow counts:
        Aggregation level     Current       Total detected   State
        Subscriber            1             1                Active
        Logical-interface     0             0                Active
        Physical-interface    0             0                Active
        Total                 1             1
```

show ddos-protection protocols (ARP broadcast)

```
user@host> show ddos-protection protocols arp bcast
Currently tracked flows: 0, Total detected flows: 0
* = User configured value
Protocol Group: ARP

  Packet type: bcast (Arp broadcast)
    Aggregate policer configuration:
      Bandwidth:        10000 pps
      Burst:            10000 packets
      Priority:         Low
      Recover time:     300 seconds
      Enabled:          Yes
    Flow detection configuration:
      Flow detection system is off
      Detection mode: Automatic  Detect time:  3 seconds
      Log flows:      Yes        Recover time: 60 seconds
      Timeout flows:  No         Timeout time: 300 seconds
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          10 pps
```

```
        Logical interface   Automatic       Drop          10 pps
        Physical interface  Automatic       Drop          10000 pps
    System-wide information:
      Aggregate bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Bandwidth: 10000 pps, Burst: 10000 packets, enabled
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
Dropped by individual policers: 0
    FPC slot 2 information:
      Bandwidth: 100% (10000 pps), Burst: 100% (10000 packets), enabled
      Hostbound queue 2
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by individual policers: 0
        Dropped by flow suppression:    0
```

show ddos-protection protocols (ARP unicast)

```
user@host> show ddos-protection protocols arp ucast
Currently tracked flows: 0, Total detected flows: 0
* = User configured value
Protocol Group: ARP

  Packet type: ucast (Arp unicast)
    Aggregate policer configuration:
      Bandwidth:        10000 pps
      Burst:            10000 packets
      Priority:         High
      Recover time:     300 seconds
      Enabled:          Yes
      Bypass aggregate: No
    Flow detection configuration:
      Flow detection system is off
      Detection mode: Automatic  Detect time:  3 seconds
      Log flows:      Yes        Recover time: 60 seconds
      Timeout flows:  No         Timeout time: 300 seconds
```

```
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          10 pps
        Logical interface   Automatic       Drop          10 pps
        Physical interface  Automatic       Drop          10000 pps
    System-wide information:
      Aggregate bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Bandwidth: 10000 pps, Burst: 10000 packets, enabled
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
Dropped by individual policers: 0
    FPC slot 2 information:
      Bandwidth: 100% (10000 pps), Burst: 100% (10000 packets), enabled
      Hostbound queue 3
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
        Dropped by flow suppression: 0
```

show ddos-protection protocols ip-options parameters

```
user@host> show ddos-protection protocols ip-options parameters
Packet types: 1, Modified: 0
* = User configured value
Protocol Group: IP-Options

  Packet type: aggregate (Aggregate for all options traffic)
    Aggregate policer configuration:
      Bandwidth:        100 pps
      Burst:            100 packets
      Priority:         Medium
      Recover time:     300 seconds
      Enabled:          Yes
    Routing Engine information:
      Bandwidth: 100 pps, Burst: 100 packets, enabled
    FPC slot 0 information:
```

```
      Bandwidth: 100% (100 pps), Burst: 100% (100 packets), enabled
      Hostbound queue 255
    FPC slot 1 information:
      Bandwidth: 100% (100 pps), Burst: 100% (100 packets), enabled
      Hostbound queue 255
    FPC slot 7 information:
      Bandwidth: 100% (100 pps), Burst: 100% (100 packets), enabled
      Hostbound queue 255
```

Release Information

Command introduced in Junos OS Release 11.2.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

Support for ARP broadcast and unicast protocol added in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

Control Plane Distributed Denial-of-Service (DDoS) Protection Overview

clear ddos-protection protocols | **3482**

show ddos-protection protocols culprit-flows | **19239**

show ddos-protection protocols flow-detection | **19249**

show ddos-protection protocols parameters | **19265**

show ddos-protection protocols statistics | **19276**

show ddos-protection protocols violations | **19292**

show ddos-protection protocols culprit-flows

IN THIS SECTION

- Syntax | **19240**
- Description | **19240**
- Options | **19240**
- Required Privilege Level | **19241**

- Output Fields | 19241
- Sample Output | 19242
- Release Information | 19248

Syntax

```
show ddos-protection protocols <protocol-group (aggregate | packet-type)> culprit-flows
```

Description

Display culprit flow information for protocol groups or individual packet types.

Options

none Display information for all protocol groups and packet types.

brief | detail (Optional) Display the specified level of output.

fpc-slot (Optional) Display information for the specified Flexible PIC Concentrator (FPC) slot.

- **Default:** system-wide, that is; include all the FPC slots.
- **Range:** 0 through 2

summary (Optional) Display flow information summary.

aggregate (Optional) Display DDoS protection information for the aggregate policer. The `aggregate` option is available for all protocol groups.

packet-type (Optional) Display information for the specified packet type in the protocol group. The available packet types vary by protocol group.

See "show ddos-protection protocols" on page 19222 for a list of available packet types.

protocol-group (Optional) Display information for a particular protocol group.

See "show ddos-protection protocols" on page 19222 for a list of available groups.

Required Privilege Level

view

Output Fields

Table 408 on page 19241 lists the output fields for the `show ddos-protection protocols culprit-flows` command. Output fields are listed in the approximate order in which they appear.

Table 408: show ddos-protection protocols culprit-flows Output Fields

Field Name	Field Description	Level of Output
`Currently tracked flows`	Number of active flows that are being tracked as culprit flows by flow detection.	All levels
`Total detected flows`	Total number of culprit flows that have been detected, including those that have recovered or timed out.	All levels
`Protocol Group`	Name of protocol group.	`detail`
`Packet type`	Name of packet type in protocol group.	`detail`
`Arriving Interface`	Logical interface on which the traffic flow arrived.	`detail`
`Aggr Flow Id level`	Shows the flow_id, such as flow_id 0001000000000022	`detail`
`Source Address MAC or IP`	Source address of the traffic flow, either a MAC address or an IP address.	`detail`
`Destination Address MAC or IP`	Destination address of the traffic flow, either a MAC address or an IP address.	`detail`
`Source Port`	Source port number.	`detail`

Table 408: show ddos-protection protocols culprit-flows Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Destination Port`	Destination port number.	`detail`
`pps`	Rate of the traffic flow in packets per second.	`brief`
`Rate`	Rate of the traffic flow in packets per second.	`detail`
`pkts`	Number of packets received in the traffic flow.	`brief`
`received packets`	Number of packets received in the traffic flow.	`detail`
Additional information	Flow ID numbers automatically assigned to flow, with embedded slot ID. The flow ID is prefixed by `sub`, `ifl`, or `ifd`, which indicate the subscriber, logical interface, and physical interface flow aggregation levels. Timestamp that identifies when the flow arrived on the interface.	`detail`

Sample Output

show ddos-protection protocols culprit-flows brief

```
user@host> show ddos-protection protocols culprit-flows brief
Currently tracked flows: 1000, Total detected flows: 1000
Protocol Packet Arriving Source Address
group type Interface MAC or IP
ndpv6 router-adv ge-1/1/0.0

2001:db8::03d4 sub:0001000000000384 2015-03-13 00:21:07 PDT pps:72 pkts:547072
ndpv6 router-adv ge-1/1/0.0
2001:db8::013f
sub:0001000000000385 2015-03-13 00:21:07 PDT pps:72 pkts:552704
ndpv6 router-adv ge-1/1/0.0
2001:db8::02e4
sub:0001000000000386 2015-03-13 00:21:07 PDT pps:72 pkts:726784
```

```
ndpv6 router-adv ge-1/1/0.0
2001:0db8::0102
 sub:0001000000000387 2015-03-13 00:21:07 PDT pps:72 pkts:762880
```

show ddos-protection protocols culprit-flows for all protocols

```
user@host> show ddos-protection protocols culprit-flows
Currently tracked flows: 1003, Total detected flows: 1003
Protocol group    Packet type    Arriving Interface    Source Address MAC or IP
pppoe            padi                ge-1/3/0                00:10:94:00:00:02
    flow_id:0001000000000003 2017-09-12 16:48:58 PDT pps:2000    pkts:153606295
dhcpv4            discover            ge-1/2/0.100        -- -- --
    flow_id:0001000000000000 2017-09-12 16:48:56 PDT pps:1000    pkts:76805613
dhcpv4            discover            ge-1/2/0.100            192.85.1.2
    flow_id:0001000000000001 2017-09-12 16:48:56 PDT pps:1000    pkts:76805603
bfd                aggregate        ge-1/2/0.100            192.85.1.2
    flow_id:0001000000000002 2017-09-12 16:48:57 PDT pps:30    pkts:2303747286
bfd                aggregate        ge-1/2/0.100            192.85.2.249
    flow_id:0001000000000004 2017-09-13 14:08:53 PDT pps:30    pkts:203
bfd    a            ggregate            ge-1/2/0.100            192.85.1.36
    flow_id:0001000000000005 2017-09-13 14:08:53 PDT pps:30    pkts:204
bfd                aggregate        ge-1/2/0.100            192.85.1.211
    flow_id:0001000000000006 2017-09-13 14:08:53 PDT pps:30    pkts:204
bfd                aggregate        ge-1/2/0.100            192.85.4.79
    flow_id:0001000000000007 2017-09-13 14:08:53 PDT pps:30    pkts:205
bfd                aggregate        ge-1/2/0.100            192.85.4.219
    flow_id:0001000000000008 2017-09-13 14:08:53 PDT pps:30    pkts:204
bfd                aggregate        ge-1/2/0.100            192.85.2.134
    flow_id:0001000000000009 2017-09-13 14:08:53 PDT pps:30    pkts:204
```

show ddos-protection protocols culprit-flows detail (Specific Protocol Group)

```
user@host> show ddos-protection protocols pppoe culprit-flows detail
Currently tracked flows: 2, Total detected flows: 1000
Protocol group        Packet type        Arriving    Interface    Aggr    Flow Id level
pppoe                    padi                ge-1/1/0.1                flow_id 0001000000000022
Ethertype:    0x0    outer-vlan:    100    inner-vlan: ---
Source Address:        00:10:94:00:00:02
Destination Address: FF:FF:FF:FF:FF:FF
Found at: 2017-10-07 07:11:27 PDT
```

```
Last Violation: 2017-10-07 07:43:24 PDT
Rate:    9995 pps    received packets: 18546724

ppoe                    padi                ge-1/1/0.1                flow_id 000100000000031c
Ethertype:    0x0    outer-vlan: 100 inner-vlan: ---
Source Address:        00:10:94:00:00:03
Destination Address: FF:FF:FF:FF:FF:FF
Found at: 2017-10-07 07:11:27 PDT
Last Violation: 2017-10-07 07:43:24 PDT
Rate:    9995 pps    received packets: 18546715

user@host> show ddos-protection protocols pppoe culprit-flows detail
Currently tracked flows: 1, Total detected flows: 1000
Protocol Packet Arriving Aggr Flow Id
group type Interface level
pppoe padi ge-1/1/0.1 sub 0001000000000022
Ethertype: 0x0 outer-vlan: 100 inner-vlan: ---
Source Address: 2001:db8::02
Destination Address: 2001:db8::FF
Found at: 2014-10-07 07:11:27 PDT
Last Violation: 2014-10-07 07:43:24 PDT
Rate: 9995 pps received packets: 18546724
user@host> show ddos-protection protocols ndpv6 culprit-flows detail
Currently tracked flows: 1, Total detected flows: 1
Protocol Packet Arriving Aggr Flow Id
group type Interface level
ndpv6 router-sol ge-1/1/0.2 sub 0001000000000001
Source Address: 2001:db8::03
Destination Address: 2001:0db8::0111
Type: 133 Code: 0
Found at: 2014-10-23 11:55:20 PDT
Last Violation: 2014-10-23 11:55:21 PDT
Rate: 30000 pps received packets: 43469
```

show expanded format for dhcpv4 discover packet type

```
user@host> show ddos-protection protocols dhcpv4 discover
Currently tracked flows: 0, Total detected flows: 0
* = User configured value Protocol Group: DHCPv4

Packet type: discover (DHCPv4 DHCPDISCOVER) Individual policer configuration:
```

```
    Bandwidth:    500 pps
    Burst:    500 packets
    Priority:    Low
    Recover time:    300 seconds Enabled:    Yes
    Bypass aggregate: No
Flow detection configuration:
    Detection mode: Automatic    Detect time:    3 seconds
    Log flows:    Yes
    Recover time: 60 seconds
    Timeout flows:    No
    Timeout time: 300 seconds
    Flow aggregation level configuration:
Aggregation level        Detection mode        Control mode    Flow rate
    Subscriber                Automatic            Drop                10 pps
    Logical interface        Automatic            Drop                10 pps
    Physical interface    Automatic            Drop            500 pps
System-wide information: Bandwidth is never violated
    Received:    0
    Arrival rate:    0 pps
    Dropped:    0
    Max arrival rate: 0 pps
Routing Engine information:
    Bandwidth: 500 pps, Burst: 500 packets, enabled Policer is never violated
    Received:    0    Arrival rate:    0 pps
    Dropped:    0    Max arrival rate: 0 pps Dropped by aggregate policer: 0
FPC slot 1 information:
    Bandwidth: 100% (500 pps), Burst: 100% (500 packets), enabled Policer is never violated
    Received:    0    Arrival rate:    0 pps
    Dropped:    0    Max arrival rate: 0 pps Dropped by aggregate policer: 0
    Dropped by flow suppression: 0
```

show dhcpv4 flow detection information

```
user@host> show ddos-protection protocols dhcpv4 flow-detection
Packet types: 19, Modified: 0
* = User configured value Protocol Group: DHCPv4
Packet type: aggregate
Flow detection configuration:
Detection mode: Automatic
Detect time:    3 seconds
Log flows:    Yes
```

```
Recover time: 60 seconds
Timeout flows:    No
Timeout time: 300 seconds
Flow aggregation level configuration:
Aggregation level        Detection mode        Control mode    Flow rate
Subscriber                Automatic            Drop                10 pps
Logical interface        Automatic            Drop                10 pps
Physical interface        Automatic            Drop                5000 pps

Packet type: unclassified
Flow detection configuration:
Detection mode: Automatic
Detect time:    3 seconds
Log flows:    Yes
Recover time: 60 seconds
Timeout flows:    No
Timeout time: 300 seconds
Flow aggregation level configuration:
Aggregation level        Detection mode        Control mode    Flow rate
Subscriber                Automatic            Drop                10 pps
Logical interface        Automatic            Drop                10 pps
Physical interface        Automatic            Drop                300 pps

Packet type: discover
Flow detection configuration:
Detection mode: Automatic
Detect time:    3 seconds
Log flows:    Yes
Recover time: 60 seconds
Timeout flows:    No
Timeout time: 300 seconds
Flow aggregation level configuration:
Aggregation level        Detection mode        Control mode    Flow rate
Subscriber                Automatic            Drop                10 pps
Logical interface        Automatic            Drop                10 pps
Physical interface        Automatic            Drop                500 pps

Packet type: offer
Flow detection configuration:
Detection mode: Automatic
Detect time:    3 seconds
Log flows:    Yes
Recover time: 60 seconds
```

```
Timeout flows:    No
Timeout time: 300 seconds
Flow aggregation level configuration:
Aggregation level        Detection mode        Control mode    Flow rate
Subscriber                Automatic            Drop                10 pps
Logical interface        Automatic            Drop                10 pps
```

show dhcpv4 flow detection information in brief format

```
user@host> show ddos-protection protocols dhcpv4 flow-detection brief
Packet types: 19, Modified: 0
* = User configured value

Detection mode(Op): a = automatic Flow control mode(Fc): d = drop o = on    k = keep x = off
p = police

Protocol    Packet        Op        Policer    Aggr lvl Op:Fc:BWidth(pps)Log     Time
group        type            mode    BW(pps)    sub        ifl        ifd        flow    out
-------------------------------------------------------------------
dhcpv4        aggregate        auto    5000    a:d:10    a:d:10    a:d:5000    Yes    No
dhcpv4        unclass..        auto    300    a:d:10    a:d:10    a:d:300    Yes    No
dhcpv4        discover            auto    500    a:d:10    a:d:10    a:d:500    Yes    No
dhcpv4        offer                auto    1000    a:d:10    a:d:10    a:d:1000    Yes    No
dhcpv4        request            auto    1000    a:d:10    a:d:10    a:d:1000    Yes    No
dhcpv4        decline            auto    500    a:d:10    a:d:10    a:d:500    Yes    No
dhcpv4        ack                auto    500    a:d:10    a:d:10    a:d:500    Yes    No
dhcpv4        nak                auto    500    a:d:10    a:d:10    a:d:500    Yes    No
dhcpv4        release            auto    2000    a:d:10    a:d:10    a:d:2000    Yes    No
dhcpv4        inform            auto    500    a:d:10    a:d:10    a:d:500    Yes    No
dhcpv4        renew                auto    2000    a:d:10    a:d:10    a:d:2000    Yes    No
dhcpv4        forcerenew        auto    2000    a:d:10    a:d:10    a:d:2000    Yes    No
dhcpv4        leasequery        auto    2000    a:d:10    a:d:10    a:d:2000    Yes    No
dhcpv4        leaseuna..        auto    2000    a:d:10    a:d:10    a:d:2000    Yes    No
dhcpv4        leaseunk..        auto    2000    a:d:10    a:d:10    a:d:2000    Yes    No
dhcpv4        leaseact..        auto    2000    a:d:10    a:d:10    a:d:2000    Yes    No
dhcpv4        bootp                auto    300    a:d:10    a:d:10    a:d:300    Yes    No
dhcpv4        no-msgtype        auto    1000    a:d:10    a:d:10    a:d:1000    Yes    No
dhcpv4        bad-pack..        auto    0        a:d:10    a:d:10    a:d:0        Yes    No
```

show global statistics

```
user@host> show ddos-protection statistics
DDOS protection global statistics:
    Policing on routing engine:    Yes
    Policing on FPC:    Yes
    Flow detection:    No
    Logging:    Yes
    Policer violation report rate:    100
    Flow report rate:    100
    Currently violated packet types:    0
    Packet types have seen violations: 0
    Total violation counts:    0
    Currently tracked flows:    0
    Total detected flows:    0
```

show ddos-protection protocols culprit-flows fpc-slot

```
user@host> show ddos-protection protocols ndpv6 culprit-flows fpc-slot 1
Currently tracked flows: 2, Total detected flows: 2
```

Release Information

Command introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

clear ddos-protection protocols | 3482

show ddos-protection protocols | 19222

show ddos-protection protocols flow-detection | 19249

show ddos-protection protocols parameters | 19265

show ddos-protection protocols statistics | 19276

show ddos-protection protocols violations | 19292

show ddos-protection protocols flow-detection

IN THIS SECTION

- Syntax | 19249
- Description | 19249
- Options | 19249
- Required Privilege Level | 19250
- Output Fields | 19250
- Sample Output | 19252
- Release Information | 19254

Syntax

```
show ddos-protection protocols <protocol-group> flow-detection
<brief | detail | terse>
```

Description

Display flow detection information for all protocol groups or for a particular protocol group.

Options

none Display information for all protocol groups.

brief | detail | terse (Optional) Display the specified level of output.

- `brief`—Display basic function information.
- `detail`—Add information to the `brief` output; it is identical to the output displayed when you choose no option. The `brief` and `detail` options display information for all protocol groups, which can be a long list.
- `terse`—Display the same level of information as the `brief` option but only for active protocol groups.

protocol-group (Optional) Display information for a particular protocol group. See "show ddos-protection protocols" on page 19222 for a list of available groups.

Required Privilege Level

view

Output Fields

Table 409 on page 19250 lists the output fields for the `show ddos-protection protocols flow-detection` command. Output fields are listed in the approximate order in which they appear.

Table 409: show ddos-protection protocols flow-detection Output Fields

Field Name	Field Description	Level of Output
`Packet types`	Number of packet types.	All levels
`Modified`	Number of packets for which policer values have been modified from the default.	All levels
`Protocol Group`	Name of protocol group.	All levels
`Packet type`	Name of packet type in protocol group.	All levels
`Flow detection configuration`	Configuration of flow detection at the packet level.	`detail` none
`Detection mode` or `Op mode`	Mode of operation for flow detection at the packet level: • `Automatic` or `a`—Search flows only when a policer is being violated. • `Off` or `x`—Never search flows even when a policer is being violated. • `On` or `o`—Search flows even when no policer is being violated.	All levels
Policer BW (pps)	Bandwidth allowed at the packet level.	`briefterse`

Table 409: show ddos-protection protocols flow-detection Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Detect time`	Time in seconds that a suspicious flow that has exceeded the bandwidth allowed for the packet type must remain in violation to be confirmed as a culprit flow.	`detail` none
`Log flows` or `Log flow`	State of automatic logging of suspicious traffic flows for the packet type: on (`Yes`) or off (`No`).	All levels
`Recover time`	Time in seconds that must pass before a culprit flow for the packet type is considered to have returned to normal. The period starts when the flow drops below the threshold that triggered the last violation.	`detail` none
`Timeout flows` or `Time out`	State of timeout enabling for culprit flows: • `Yes`—Enabled; flows can time out (released from suppression) when a timeout period expires, regardless of whether flow is still in violation. • `No`—Disabled; flows are not allowed to time out.	All levels
`Timeout time`	Time in seconds that a culprit flow is suppressed. On expiration, the flow times out even if it is still violating the bandwidth limit.	`detail` none
`Flow aggregation level configuration`	Configuration of flow detection for each flow aggregation level.	`detail` none
`Aggregation level` or `Agg level`	One of three levels of flow aggregation • `Subscriber` or `sub` • `Logical interface` or `ifl` • `Physical interface` or `ifd`	All levels

Table 409: show ddos-protection protocols flow-detection Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Detection mode` or `Op`	Mode of operation for flow detection at the flow aggregation level: • `Automatic`—Search flows only when a policer is being violated. • `Off`—Never search flows even when a policer is being violated. • `On`—Search flows even when no policer is being violated.	All levels
`Control mode` or `Fc`	Mode by which traffic in a culprit flow is handled. • `drop`—Drop all traffic in flow. • `keep`—Keep all traffic in flow. • `police`—Police the traffic to within its allowed bandwidth.	All levels
`Flow rate` or BWidth (pps)	Bandwidth allowed at the flow aggregation level.	`briefterse`

Sample Output

show ddos-protection protocols flow-detection

```
user@host> show ddos-protection protocols flow-detection
Packet types: 190, Modified: 2
* = User configured value

Protocol Group: IPv4-Unclassified

  Packet type: aggregate
    Flow detection configuration:
      Detection mode: Automatic  Detect time:  3 seconds
      Log flows:      No         Recover time: 60 seconds
      Timeout flows:  No         Timeout time: 300 seconds
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          10 pps
```

```
        Logical interface   Automatic       Drop          10 pps
        Physical interface  Automatic       Drop          2000 pps


Protocol Group: IPv6-Unclassified

  Packet type: aggregate
    Flow detection configuration:
      Detection mode: Automatic  Detect time:  3 seconds
      Log flows:      No         Recover time: 60 seconds
      Timeout flows:  No         Timeout time: 300 seconds
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          10 pps
        Logical interface   Automatic       Drop          10 pps
        Physical interface  Automatic       Drop          2000 pps

...
```

show ddos-protection protocols flow-detection brief (Parameters for a Specific Protocol)

```
user@host> show ddos-protection protocols dhcpv4 flow-detection brief
Packet types: 19, Modified: 1
* = User configured value

Detection mode(Op): a = automatic    Flow control mode(Fc): d = drop
                    o = on                                  k = keep
                    x = off                                 p = police

Protocol    Packet      Op   Policer  Aggr level Op:Fc:BWidth(pps)   Log  Time
group       type        mode BW(pps)  sub       ifl       ifd        flow out
dhcpv4      aggregate   auto 5000     a:d:10    a:d:10    a:d:5000   No   No
dhcpv4      unclass..   auto 300      a:d:10    a:d:10    a:d:300    No   No
dhcpv4      discover    auto 777*     a:d:10    a:d:10    a:d:500    No   No
dhcpv4      offer       auto 1000     a:d:10    a:d:10    a:d:1000   No   No
dhcpv4      request     auto 1000     a:d:10    a:d:10    a:d:1000   No   No
dhcpv4      decline     auto 500      a:d:10    a:d:10    a:d:500    No   No
dhcpv4      ack         auto 500      a:d:10    a:d:10    a:d:500    No   No
dhcpv4      nak         auto 500      a:d:10    a:d:10    a:d:500    No   No
dhcpv4      release     auto 2000     a:d:10    a:d:10    a:d:2000   No   No
dhcpv4      inform      auto 500      a:d:10    a:d:10    a:d:500    No   No
```

```
dhcpv4      renew       auto 2000     a:d:10    a:d:10    a:d:2000   No   No
dhcpv4      forcerenew  auto 2000     a:d:10    a:d:10    a:d:2000   No   No
dhcpv4      leasequery  auto 2000     a:d:10    a:d:10    a:d:2000   No   No
dhcpv4      leaseuna..  auto 2000     a:d:10    a:d:10    a:d:2000   No   No
dhcpv4      leaseunk..  auto 2000     a:d:10    a:d:10    a:d:2000   No   No
dhcpv4      leaseact..  auto 2000     a:d:10    a:d:10    a:d:2000   No   No
dhcpv4      bootp       auto 300      a:d:10    a:d:10    a:d:300    No   No
dhcpv4      no-msgtype  auto 0        a:d:10    a:d:10    a:d:0      No   No
dhcpv4      bad-pack..  auto 0        a:d:10    a:d:10    a:d:0      No   No
```

Release Information

Command introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

- clear ddos-protection protocols | 3482
- show ddos-protection protocols | 19222
- show ddos-protection protocols culprit-flows | 19239
- show ddos-protection protocols parameters | 19265
- show ddos-protection protocols statistics | 19276
- show ddos-protection protocols violations | 19292

show ddos-protection protocols isis

IN THIS SECTION

- Syntax | 19255
- Description | 19255
- Options | 19255
- Required Privilege Level | 19255
- Output Fields | 19255

Sample Output | 19262
Release Information | 19265

Syntax

```
show ddos-protection protocols <protocol-group> isis
 <aggregate | isis-data | isis-hello>
```

Description

Display the ISIS data traffic information for all protocol groups or individual packet types.

Options

none	Display information for all protocol groups and packet types.
protocol-group	(Optional) Display control plane DDoS protection information for a protocol group.
packet-type	(Optional) Display control plane DDoS protection information for the specified packet type in the specified protocol group. The available packet types vary by protocol group, and only some protocol groups can have policers for individual packet types.
aggregate	(Optional) Display control plane DDoS protection information for the aggregate policer. The `aggregate` option is available for all ISIS data traffic information.
isis-data	(Optional) Display ISIS Data traffic information.
isis-hello	(Optional) Display ISIS Hello traffic information.

Required Privilege Level

view

Output Fields

Table 410 on page 19256 lists the output fields for the `show ddos-protection protocols isis` command. Output fields are listed in the approximate order in which they appear.

Table 410: show ddos-protection protocols isis Output Fields

Field Name	Field Description
`Packet types`	Number of packet types
`Modified`	Number of packets for which policer values have been modified from the default.
`Received traffic`	Number of traffic flows received.
`Currently violated`	Number of flows that are currently violating the flow bandwidth limit.
`Currently tracked flows`	Number of active flows that are being tracked as culprit flows by flow detection.
`Total detected flows`	Total number of culprit flows that have been detected, including those that have recovered or timed out.
`Protocol Group`	Name of protocol group.
`Packet type`	Name of packet type in protocol group.
`Bandwidth`	Bandwidth policer value; number of packets per second that is allowed before a violation is declared.
`Burst`	Burst policer value; the maximum number of packets that is allowed in a burst before a violation is declared.
`Priority`	Priority of the packet type for individual packet policers that enables more important traffic to pass through in the event of traffic congestion: `low`, `medium`, or `high`. Lower priority packets are dropped when insufficient bandwidth is available.
`Recover time`	Time in seconds that must pass before the traffic flow is considered to have recovered from the attack. A notification is generated when the timer expires.

Table 410: show ddos-protection protocols isis Output Fields *(Continued)*

Field Name	Field Description
`Enabled`	State of the policer: • `Yes`—The policer is enabled on both the Routing Engine and the FPC (line card). This is the default state. • `No`—The policer is disabled on both the Routing Engine and the FPC by global configuration. The policer is not disabled by the packet type level configuration. • `No*`—The policer is disabled on both the Routing Engine and the FPC. The asterisk (*) indicates that one or both of these instances are disabled at the packet type level; the policer can also be disabled globally. • `Partial`—The policer is disabled on either the Routing Engine or the FPC, but not both. It is disabled by global configuration. The policer is not disabled by the packet type level configuration. • `Partial*`—The policer is disabled on either the Routing Engine or the FPC, but not both. The asterisk (*) indicates that the instance is disabled by the packet type level configuration; the policer can also be disabled globally. Disabling to occurs globally for all packet types at the `[edit system ddos-protection global]` hierarchy level, for a specific packet type at the `[edit system ddos-protection protocols` *`protocol-group`* `(aggregate | `*`packet-type`*`]` hierarchy level, or at both levels.
`Bypass aggregate`	State of the bypass aggregate configuration: • Yes—The aggregate policer is bypassed. • No—The aggregate policer is enforced. This field appears only for individual policers.

Table 410: show ddos-protection protocols isis Output Fields *(Continued)*

Field Name	Field Description

`Flow detection configuration`	State of flow detection configured on the router:
	• Detection mode—Mode of operation for suspicious flow detection: automatic, off, or on.
	• Log flows—State of automatic logging of suspicious traffic flows: on (`Yes`) or off (`No`).
	• Timeout flows—State of culprit flow timeout behavior: flow is suppressed for a configured timeout period (`Yes`) or flow is suppressed until it is no longer in violation (`No`).
	• Detect time—Time in seconds that must pass before a suspicious flow that has exceeded the bandwidth allowed for the packet type is considered to be a culprit flow.
	• Recover time—Time in seconds that must pass before a culprit flow is considered to have returned to normal. The period starts when the flow drops below the threshold that triggered the last violation.
	• Timeout time—Time in seconds that a culprit flow is suppressed, if timeouts have been enabled.
	• Flow aggregation level configuration—Flow detection mode, flow control mode, and flow bandwidth for traffic at each of the traffic flow aggregation levels: subscriber, logical interface, and physical interface.
	  • Aggregation level— Flow detection mode, flow control mode, and flow bandwidth for traffic at each of the traffic flow aggregation levels: subscriber, logical interface, and physical interface.
	  • Detection mode—State of flow detection: automatic, off, or on.
	  • Control mode—Mode of controlling culprit traffic: dropped, kept, or policed back to within the allowed bandwidth.
	  • Flow rate—Bandwidth allowed for the control traffic in packets per second.

Table 410: show ddos-protection protocols isis Output Fields *(Continued)*

Field Name	Field Description
`System-wide information`	The following information is collected for the router: • A message indicates whether the policer has been violated. • No. of FPCs currently receiving excess traffic—Number of cards that are currently in violation of a policer. • No. of FPCs that have received excess traffic—Number of cards that have at some point been in violation of a policer. • Violation first detected at—Timestamp of the first violation. • Violation last seen at—Timestamp of the last observed violation. • Duration of violation—Length of the violation. • Number of violations—Number of times the violation has occurred. • Received—Number of packets received at all card slots and the Routing Engine. • Dropped—Number of packets dropped regardless of where they were dropped. • Arrival rate—Current traffic rate for packets arriving from all cards and at the Routing Engine. • Max arrival rate—Highest traffic rate for packets arriving from all cards and at the Routing Engine.

Table 410: show ddos-protection protocols isis Output Fields *(Continued)*

Field Name	Field Description

`Routing Engine information`	The following information is collected for the Routing Engine:
	• Bandwidth—Maximum number of packets per second that is allowed.
	• Burst—Maximum number of packets that is allowed in a burst.
	• State of the policer:
	  • `enabled`—The Routing Engine policer is enabled. This is the default state.
	  • `disabled`—The Routing Engine policer is disabled globally. It is not disabled by the packet type level configuration.
	  • `disabled*`—The Routing Engine policer is disabled by the packet type level configuration; it can also be disabled globally.
	• A message indicates whether the policer has been violated; the policer might be passed at the individual cards, but the combined rate of packets arriving at the Routing Engine can exceed the configured policer value.
	• Violation first detected at—Timestamp of the first violation.
	• Violation last seen at—Timestamp of the last observed violation.
	• Duration of violation—Length of the violation.
	• Number of violations—Number of times the violation has occurred.
	• Received—Number of packets received at the Routing Engine from all cards.
	• Dropped—Number of packets dropped at the Routing Engine; includes packets dropped by the aggregate policer and by individual protocol policers.
	• Arrival rate—Current traffic rate for packets arriving at the Routing Engine from all cards.
	• Max arrival rate—Highest traffic rate for packets arriving at the Routing Engine from all cards.
	• Dropped by aggregate policer—Number of packets dropped by the aggregate policer.
	• Dropped by individual policers—Number of packets dropped by individual policer.

Table 410: show ddos-protection protocols isis Output Fields *(Continued)*

Field Name	Field Description

`FPC slot information`	The following information is collected for the line card in the indicated slot:
	• Bandwidth—Bandwidth scaling percentage and the number of packets per second that is allowed before a violation is declared.
	• Burst—Burst scaling percentage and the maximum number of packets that is allowed in a burst before a violation is declared.
	• State of the policer:
	  • `enabled`—The FPC policer is enabled. This is the default state.
	  • `disabled`—The FPC policer is disabled globally. It is not disabled by the packet type level configuration.
	  • `disabled*`—The FPC policer is disabled by the packet type level configuration; it may also be disabled globally.
	• A message indicates whether the policer has been violated.
	• Violation first detected at—Timestamp of the first violation.
	• Violation last seen at—Timestamp of the last observed violation.
	• Duration of violation—Length of the violation.
	• Number of violations—Number of times the violation has occurred.
	• Received—Number of packets received on the line card.
	• Dropped—Number of packets dropped at the line card; includes packets dropped by the aggregate policer and by individual protocol policers.
	• Arrival rate—Current traffic rate for packets arriving at the line card.
	• Max arrival rate—Highest traffic rate for packets arriving at the line card.
	• Dropped by this policer—Number of packets dropped by the individual policer.
	• Dropped by aggregate policer—Number of packets dropped by the aggregate policer.

Sample Output

show ddos-protection protocols isis

```
user@host> show ddos-protection protocols isis

Packet types: 3, Modified: 0, Received traffic: 0, Currently violated: 0
Currently tracked flows: 0, Total detected flows: 0
* = User configured value

Protocol Group: ISIS

  Packet type: aggregate (Aggregate for isis traffic)
  Packet type: isis-data (All isis-data traffic)
  Packet type: isis-hello (All isis hello traffic)
...
```

show ddos-protection protocols isis aggregate

```
user@host> show ddos-protection protocols isis aggregate

Currently tracked flows: 0, Total detected flows: 0
* = User configured value
Protocol Group: ISIS

  Packet type: aggregate (Aggregate for isis traffic)
    Aggregate policer configuration:
      Bandwidth:        6000 pps
      Burst:            6000 packets
      Recover time:     300 seconds
      Enabled:          Yes
    System-wide information:
      Aggregate bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Bandwidth: 6000 pps, Burst: 6000 packets, enabled
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
```

```
        Dropped by individual policers: 0
    FPC slot 0 information:
      Bandwidth: 100% (6000 pps), Burst: 100% (6000 packets), enabled
      Hostbound queue 255
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by individual policers: 0
        Dropped by flow suppression:    0
```

show ddos-protection protocols isis isis-data

```
user@host> show ddos-protection protocols isis isis-data

Currently tracked flows: 0, Total detected flows: 0
* = User configured value
Protocol Group: ISIS

  Packet type: isis-data (All isis-data traffic)
    Individual policer configuration:
      Bandwidth:        5000 pps
      Burst:            4096 packets
      Priority:         High
      Recover time:     300 seconds
      Enabled:          Yes
      Bypass aggregate: No
    Flow detection configuration:
      Flow detection system is off
      Detection mode: Automatic  Detect time:  0 seconds
      Log flows:      Yes        Recover time: 0 seconds
      Timeout flows:  No         Timeout time: 0 seconds
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          0  pps
        Logical interface   Automatic       Drop          0  pps
        Physical interface  Automatic       Drop          5000 pps
    System-wide information:
      Bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
```

```
      Bandwidth: 5000 pps, Burst: 4096 packets, enabled
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
    FPC slot 0 information:
      Bandwidth: 100% (5000 pps), Burst: 100% (4096 packets), enabled
      Hostbound queue 255
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
        Dropped by flow suppression:  0
```

show ddos-protection protocols isis isis-hello

```
user@host> show ddos-protection protocols isis isis-hello

Currently tracked flows: 0, Total detected flows: 0
* = User configured value
Protocol Group: ISIS

  Packet type: isis-hello (All isis hello traffic)
    Individual policer configuration:
      Bandwidth:        1000 pps
      Burst:            2048 packets
      Priority:         High
      Recover time:     300 seconds
      Enabled:          Yes
      Bypass aggregate: No
    Flow detection configuration:
      Flow detection system is off
      Detection mode: Automatic  Detect time:  0 seconds
      Log flows:      Yes        Recover time: 0 seconds
      Timeout flows:  No         Timeout time: 0 seconds
      Flow aggregation level configuration:
        Aggregation level   Detection mode  Control mode  Flow rate
        Subscriber          Automatic       Drop          0  pps
        Logical interface   Automatic       Drop          0  pps
        Physical interface  Automatic       Drop          1000 pps
    System-wide information:
```

```
    Bandwidth is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
  Routing Engine information:
    Bandwidth: 1000 pps, Burst: 2048 packets, enabled
    Policer is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
      Dropped by aggregate policer: 0
  FPC slot 0 information:
    Bandwidth: 100% (1000 pps), Burst: 100% (2048 packets), enabled
    Hostbound queue 255
    Policer is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
      Dropped by aggregate policer: 0
      Dropped by flow suppression:  0
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Control Plane DDoS Protection Flow Detection Overview

protocols (DDoS) (ACX Series, PTX Series, and QFX Series) | **13765**

show ddos-protection protocols parameters

IN THIS SECTION

- Syntax | **19266**
- Description | **19266**
- Options | **19266**
- Required Privilege Level | **19267**

- Output Fields | 19267
- Sample Output | 19269
- Release Information | 19275

Syntax

```
show ddos-protection protocols <protocol-group> parameters
 <brief | detail | terse>
```

Description

Display DDoS protection configuration information for all protocol groups or for a particular protocol group.

Starting in Junos OS Release 22.3R1, on MX Series and EX9200 Series devices, we've updated the default bandwidth value from 20000 to 100 pps and burst policer value from 20000 to 100 packets. This enhancement avoids the CPU usage of `eventd` and `snmpd` reaching more than 100%. Earlier to this release, when the system receives a violated traffic for SNMP along with other protocols traffic, the CPU usage of `eventd` and `snmpd` was reaching more than 100% with an error.

Starting in Junos OS Evolved Release 23.2R2, on PTX Series devices, the `show ddos-protection protocols statistics` displays the `Max arrival rate` and `Arrival rate` output values as expected. Earlier to this release, the `Max arrival rate` and `Arrival rate` output values were displayed larger than expected.

Options

none Display information for all protocol groups.

brief | detail | terse (Optional) Display the specified level of output.

- `brief`—Display basic function information.
- `detail`—Add information to the `brief` output; it is identical to the output displayed when you choose no option. The `brief` and `detail` options display information for all protocol groups, which can be a long list.
- `terse`—Display the same level of information as the `brief` option but only for active protocol groups—groups that show traffic in the `Received (packets)` column.

protocol-group (Optional) Display information for a particular protocol group. See "show ddos-protection protocols" on page 19222 for a list of available groups.

Required Privilege Level

view

Output Fields

Table 411 on page 19267 lists the output fields for the `show ddos-protection protocols parameters` command. Output fields are listed in the approximate order in which they appear.

Table 411: show ddos-protection protocols parameters Output Fields

Field Name	Field Description	Level of Output
`Protocol Group`	Name of protocol group.	All levels
`Packet type`	Name of packet type in protocol group.	All levels
`Bandwidth`	Bandwidth policer value; number of packets per second that is allowed before a violation is declared. In the `brief` output, an asterisk indicates the value has been modified from the default.	All levels
`Burst`	Burst policer value; the maximum number of packets that is allowed in a burst before a violation is declared. In the `brief` output, an asterisk indicates the value has been modified from the default.	All levels
`Priority`	Priority of the packet type in the event of traffic congestion: `low`, `medium`, or `high`. Lower priority packets can be dropped when insufficient bandwidth is available. In the `brief` output, an asterisk indicates the value has been modified from the default.	All levels

Table 411: show ddos-protection protocols parameters Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Recover time`	Time that must pass since the last violation before the traffic flow is considered to have recovered from the attack. A notification is generated when the timer expires. In the `brief` output, an asterisk indicates the value has been modified from the default.	All levels
`Enabled`	State of the policer, enabled (`Yes`) or disabled (`No`).	`detail` none
`Bypass aggregate`	State of the bypass aggregate configuration: • Yes—The aggregate policer is bypassed. • No—The aggregate policer is enforced. This field appears only for individual policers.	`detail` none
`FPC slot information`	The following configuration information for the card in the indicated slot: • Bandwidth—Bandwidth scale and the number of packets per second that is allowed before a violation is declared • Burst—Burst scale and the maximum number of packets that is allowed in a burst before a violation is declared • `enabled` or `disabled`—State of the line card policer	`detail` none
`Number of policers modified`	Number of policers that have been changed from the default configuration. An asterisk by a particular value indicates that value has been modified.	`briefterse`
`Policer Enabled`	State of the policer, enabled (`Yes`), disabled (`No`), or partially disabled (`part.`); `part.` indicates that only some of the policer instances are disabled for the policer.	`briefterse`

Table 411: show ddos-protection protocols parameters Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Bypass aggr.`	State of the bypass aggregate configuration: • Yes—The aggregate policer is bypassed. • No—The aggregate policer is enforced. Dashes indicate that the bypass aggregate configuration is not available; this is possible only for aggregate policers.	`briefterse`
`FPC Mod`	Indicates whether configuration has changed from the default for any line cards. • No—The default configuration has not changed from the default for the packet type. • Yes—The default configuration has changed from the default for the packet type	`briefterse`

Sample Output

show ddos-protection protocols parameters

```
user@host> show ddos-protection protocols parameters
Protocol Group: IPv4-Unclassified

  Packet type: aggregate (Aggregate for unclassified host-bound IPv4 traffic)
    Aggregate policer configuration:
      Bandwidth:        20000 pps
      Burst:            20000 packets
      Priority:         medium
      Recover time:     300 seconds
      Enabled:          Yes
    FPC slot 1 information:
      Bandwidth: 100% (20000 pps), Burst: 100% (20000 packets), enabled

Protocol Group: IPv6-Unclassified

  Packet type: aggregate (Aggregate for unclassified host-bound IPv6 traffic)
```

```
    Aggregate policer configuration:
      Bandwidth:        20000 pps
      Burst:            20000 packets
      Priority:         medium
      Recover time:     300 seconds
      Enabled:          Yes
    FPC slot 1 information:
      Bandwidth: 100% (20000 pps), Burst: 100% (20000 packets), enabled

...

Protocol Group: PPPoE

  Packet type: aggregate (Aggregate for all PPPoE control traffic)
    Aggregate policer configuration:
      Bandwidth:        800 pps
      Burst:            2000 packets
      Priority:         medium
      Recover time:     300 seconds
      Enabled:          Yes
    FPC slot 1 information:
      Bandwidth: 100% (800 pps), Burst: 100% (2000 packets), enabled

  Packet type: padi (PPPoE PADI)
    Individual policer configuration:
      Bandwidth:        500 pps
      Burst:            500 packets
      Priority:         low
      Recover time:     300 seconds
      Enabled:          Yes
      Bypass aggregate: No
    FPC slot 1 information:
      Bandwidth: 100% (500 pps), Burst: 100% (500 packets), enabled

  Packet type: pado (PPPoE PADO)
    Individual policer configuration:
      Bandwidth:        0 pps
      Burst:            0 packets
      Priority:         low
      Recover time:     300 seconds
      Enabled:          Yes
      Bypass aggregate: No
    FPC slot 1 information:
```

```
      Bandwidth: 100% (0 pps), Burst: 100% (0 packets), enabled

  Packet type: padr (PPPoE PADR)
    Individual policer configuration:
      Bandwidth:        500 pps
      Burst:            500 packets
      Priority:         medium
      Recover time:     300 seconds
      Enabled:          Yes
      Bypass aggregate: No
    FPC slot 1 information:
      Bandwidth: 100% (500 pps), Burst: 100% (500 packets), enabled
```

show ddos-protection protocols parameters brief

```
user@host> show ddos-protection protocols parameters brief
Number of policers modified: 3
Protocol    Packet      Bandwidth Burst  Priority Recover   Policer Bypass FPC
group       type        (pps)     (pkts)          time(sec) enabled aggr.  mod
ipv4-uncls  aggregate   20000     20000  medium   300       yes     --     no
ipv6-uncls  aggregate   20000     20000  medium   300       yes     --     no
dynvlan     aggregate   1000      500    low      300       yes     --     no
ppp         aggregate   16000     16000  medium   300       yes     --     no
ppp         unclass     1000      500    low      300       yes     no     no
ppp         lcp         12000     12000  low      300       yes     no     no
ppp         auth        2000      2000   medium   300       yes     no     no
ppp         ipcp        2000      2000   high     300       yes     no     no
ppp         ipv6cp      2000      2000   high     300       yes     no     no
ppp         mplscp      2000      2000   high     300       yes     no     no
ppp         isis        2000      2000   high     300       yes     no     no
pppoe       aggregate   800*      2000   medium   300       part.*  --     no
pppoe       padi        500       500    low      300       part.   no     no
pppoe       pado        0         0      low      300       part.   no     no
pppoe       padr        500       500    medium   300       part.   no     no
pppoe       pads        0         0      low      300       part.   no     no
pppoe       padt        1000      1000   high     300       part.   no     no
pppoe       padm        0         0      low      300       part.   no     no
pppoe       padn        0         0      low      300       part.   no     no
dhcpv4      aggregate   669*      5000   medium   300       yes     --     no
dhcpv4      unclass..   300       150    low      300       yes     no     no
dhcpv4      discover    100*      500    low      300       yes     no     no
```

```
dhcpv4      offer       1000      1000   low      300       yes     no     no
dhcpv4      request     1000      1000   medium   300       yes     no     no
dhcpv4      decline     500       500    low      300       yes     no     no
dhcpv4      ack         500       500    medium   300       yes     no     no
dhcpv4      nak         500       500    low      300       yes     no     no
dhcpv4      release     2000      2000   high     300       yes     no     no
dhcpv4      inform      500       500    low      300       yes     no     no
dhcpv4      renew       2000      2000   high     300       yes     no     no
dhcpv4      forcerenew  2000      2000   high     300       yes     no     no
dhcpv4      leasequery  2000      2000   high     300       yes     no     no
dhcpv4      leaseuna..  2000      2000   high     300       yes     no     no
dhcpv4      leaseunk..  2000      2000   high     300       yes     no     no
dhcpv4      leaseact..  2000      2000   high     300       yes     no     no
dhcpv4      bootp       300       300    low      300       yes     no     no
dhcpv4      no-msgtype  0         0      low      300       yes     no     no
dhcpv4      bad-pack..  0         0      low      300       yes     no     no

...

icmp        aggregate   20000     20000  high     300       yes     --     no
igmp        aggregate   20000     20000  high     300       yes     --     no
ospf        aggregate   20000     20000  high     300       yes     --     no
rsvp        aggregate   20000     20000  high     300       yes     --     no
pim         aggregate   20000     20000  high     300       yes     --     no
rip         aggregate   20000     20000  high     300       yes     --     no
ptp         aggregate   20000     20000  high     300       yes     --     no
bfd         aggregate   20000     20000  high     300       yes     --     no
lmp         aggregate   20000     20000  high     300       yes     --     no
ldp         aggregate   20000     20000  high     300       yes     --     no
msdp        aggregate   20000     20000  high     300       yes     --     no
bgp         aggregate   20000     20000  low      300       yes     --     no
vrrp        aggregate   20000     20000  high     300       yes     --     no
telnet      aggregate   20000     20000  low      300       yes     --     no
ftp         aggregate   20000     20000  low      300       yes     --     no
ssh         aggregate   20000     20000  low      300       yes     --     no
snmp        aggregate   20000     20000  low      300       yes     --     no
ancp        aggregate   20000     20000  low      300       yes     --     no

...
```

show ddos-protection protocols dhcpv4 parameters brief

```
user@host> show ddos-protection protocols dhcpv4 parameters brief
Number of policers modified: 2
Protocol    Packet      Bandwidth Burst  Priority Recover   Policer Bypass FPC
group       type        (pps)     (pkts)          time(sec) enabled aggr.  mod
dhcpv4      aggregate   669*      5000   medium   300       yes     --     no
dhcpv4      unclass..   300       150    low      300       yes     no     no
dhcpv4      discover    100*      500    low      300       yes     no     no
dhcpv4      offer       1000      1000   low      300       yes     no     no
dhcpv4      request     1000      1000   medium   300       yes     no     no
dhcpv4      decline     500       500    low      300       yes     no     no
dhcpv4      ack         500       500    medium   300       yes     no     no
dhcpv4      nak         500       500    low      300       yes     no     no
dhcpv4      release     2000      2000   high     300       yes     no     no
dhcpv4      inform      500       500    low      300       yes     no     no
dhcpv4      renew       2000      2000   high     300       yes     no     no
dhcpv4      forcerenew  2000      2000   high     300       yes     no     no
dhcpv4      leasequery  2000      2000   high     300       yes     no     no
dhcpv4      leaseuna..  2000      2000   high     300       yes     no     no
dhcpv4      leaseunk..  2000      2000   high     300       yes     no     no
dhcpv4      leaseact..  2000      2000   high     300       yes     no     no
dhcpv4      bootp       300       300    low      300       yes     no     no
dhcpv4      no-msgtype  0         0      low      300       yes     no     no
dhcpv4      bad-pack..  0         0      low      300       yes     no     no
```

show ddos-protection protocols dhcpv4 parameters terse

```
user@host> show ddos-protection protocols dhcpv4 parameters terse
Number of policers modified: 2
Protocol    Packet      Bandwidth Burst  Priority Recover   Policer Bypass FPC
group       type        (pps)     (pkts)          time(sec) enabled aggr.  mod
dhcpv4      aggregate   669*      5000   medium   300       yes     --     no
dhcpv4      discover    100*      500    low      300       yes     no     no
```

show ddos-protection protocols dhcpv4 parameters

```
user@host> show ddos-protection protocols dhcpv4 parameters
Protocol Group: DHCPv4
```

```
Packet type: aggregate (aggregate for all DHCPv4 traffic)
  Aggregate policer configuration:
    Bandwidth:        669 pps
    Burst:            5000 packets
    Priority:         medium
    Recover time:     300 seconds
    Enabled:          Yes
  FPC slot 1 information:
    Bandwidth: 100% (669 pps), Burst: 100% (5000 packets), enabled

Packet type: unclassified (Unclassified DHCPv4 traffic)
  Individual policer configuration:
    Bandwidth:        300 pps
    Burst:            150 packets
    Priority:         low
    Recover time:     300 seconds
    Enabled:          Yes
    Bypass aggregate: No
  FPC slot 1 information:
    Bandwidth: 100% (300 pps), Burst: 100% (150 packets), enabled

Packet type: discover (DHCPv4 DHCPDISCOVER)
  Individual policer configuration:
    Bandwidth:        100 pps
    Burst:            500 packets
    Priority:         low
    Recover time:     300 seconds
    Enabled:          Yes
    Bypass aggregate: No
  FPC slot 1 information:
    Bandwidth: 100% (100 pps), Burst: 100% (500 packets), enabled

Packet type: offer (DHCPv4 DHCPOFFER)
  Individual policer configuration:
    Bandwidth:        1000 pps
    Burst:            1000 packets
    Priority:         low
    Recover time:     300 seconds
    Enabled:          Yes
    Bypass aggregate: No
  FPC slot 1 information:
    Bandwidth: 100% (1000 pps), Burst: 100% (1000 packets), enabled
```

```
  Packet type: request (DHCPv4 DHCPREQUEST)
    Individual policer configuration:
      Bandwidth:        1000 pps
      Burst:            1000 packets
      Priority:         medium
      Recover time:     300 seconds
      Enabled:          Yes
      Bypass aggregate: No
    FPC slot 1 information:
      Bandwidth: 100% (1000 pps), Burst: 100% (1000 packets), enabled

...
```

show ddos-protection protocols snmp parameters (Starting in Junos OS Release 22.3R1)

```
Packet types: 1, Modified: 0
* = User configured value
Protocol Group: SNMP

  Packet type: aggregate (Aggregate for all snmp traffic)
    Aggregate policer configuration:
      Bandwidth:        100 pps
      Burst:            100 packets
      Priority:         Low
      Recover time:     300 seconds
      Enabled:          Yes
    Routing Engine information:
      Bandwidth: 100 pps, Burst: 100 packets, enabled
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

clear ddos-protection protocols | 3482

show ddos-protection protocols | 19222

show ddos-protection protocols culprit-flows | 19239

show ddos-protection protocols flow-detection | 19249
show ddos-protection protocols statistics | 19276
show ddos-protection protocols violations | 19292

show ddos-protection protocols statistics

IN THIS SECTION

- Syntax | 19276
- Description | 19276
- Options | 19277
- Required Privilege Level | 19277
- Output Fields | 19277
- Sample Output | 19281
- Release Information | 19291

Syntax

```
show ddos-protection protocols <protocol-group> statistics
 <brief | detail | terse>
```

Description

Display traffic statistics and DDoS policer violation statistics for all protocol groups or for a particular protocol group.

NOTE: DDoS protection policers act on the system's traffic queues. The QFX5100 and QFX5200 lines of switches manage traffic for more protocols than the number of queues, so the system often must map more than one protocol to the same queue. When traffic for one protocol shares a queue with other protocols and violates DDoS protection policer limits, this command reports a violation on that queue for all mapped

protocols because the system doesn't distinguish which protocol's traffic specifically caused the violation. You can use what you know about the types of traffic flowing through your network to identify which of the reported protocols actually triggered the violation.

Options

none Display information for all protocol groups.

brief | detail | terse (Optional) Display the specified level of output.

- `brief`—Display basic function information.
- `detail`—Add information to the `brief` output; it is identical to the output displayed when you choose no option. The `brief` and `detail` options display information for all protocol groups, which can be a long list.
- `terse`—Display the same level of information as the `brief` option but only for active protocol groups—groups that show traffic in the `Received (packets)` column.

protocol-group (Optional) Display information for a particular protocol group. See "show ddos-protection protocols" on page 19222 for a list of available groups.

Required Privilege Level

view

Output Fields

Table 412 on page 19277 lists the output fields for the `show ddos-protection protocols statistics` command. Output fields are listed in the approximate order in which they appear.

Table 412: show ddos-protection protocols statistics Output Fields

Field Name	Field Description	Level of Output
`Protocol Group`	Name of protocol group.	All levels
`Packet type`	Name of packet type in protocol group.	All levels

Table 412: show ddos-protection protocols statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`System-wide information`	The following information collected for the router: • A message indicates whether the policer has been violated. • No. of FPCs currently receiving excess traffic—Number of cards that are currently in violation of a policer. • No. of FPCs that have received excess traffic—Number of cards that have at some point been in violation of a policer. • Violation first detected at—Timestamp of the first violation. • Violation last seen at—Timestamp of the last observed violation. • Duration of violation—Length of the violation. • Number of violations—Number of times the violation has occurred. • Received—Number of packets received at all card slots and the Routing Engine. • Dropped—Number of packets dropped regardless of where they were dropped. • Arrival rate—Current traffic rate for packets arriving from all cards and at the Routing Engine. • Max arrival rate—Highest traffic rate for packets arriving from all cards and at the Routing Engine.	`detail` none

Table 412: show ddos-protection protocols statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Routing Engine information`	The following information collected for the Routing Engine: • A message indicates whether the policer has been violated; the policer might be passed at the individual cards, but the combined rate of packets arriving at the Routing Engine can exceed the configured policer value. • Violation first detected at—Timestamp of the first violation. • Violation last seen at—Timestamp of the last observed violation. • Duration of violation—Length of the violation. • Number of violations—Number of times the violation has occurred. • Received—Number of packets received at the Routing Engine from all cards. • Dropped—Number of packets dropped at the Routing Engine; includes packets dropped by the aggregate policer and by individual protocol policers. • Arrival rate—Current traffic rate for packets arriving at the Routing Engine from all cards. • Max arrival rate—Highest traffic rate for packets arriving at the Routing Engine from all cards. • Dropped by aggregate policer—Number of packets dropped by the aggregate policer. • Dropped by individual policers—Number of packets dropped by individual policer.	`detail` none

Table 412: show ddos-protection protocols statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`FPC slot information`	The following information collected for the card in the indicated slot: • A message indicates whether the policer has been violated • Violation first detected at—Timestamp of the first violation • Violation last seen at—Timestamp of the last observed violation • Duration of violation—Length of the violation • Number of violations—Number of times the violation has occurred • Received—Number of packets received on the line card • Dropped—Number of packets dropped at the line card; includes packets dropped by the aggregate policer and by individual protocol policers • Arrival rate—Current traffic rate for packets arriving at the line card • Max arrival rate—Highest traffic rate for packets arriving at the line card • Dropped by this policer—Number of packets dropped by the individual policer • Dropped by aggregate policer—Number of packets dropped by the aggregate policer	`detail` none
`Received (packets)`	Number of packets of this packet type or protocol group received at all cards and the Routing Engine.	`briefterse`
`Dropped (packets)`	Number of packets dropped for this packet type or protocol group, regardless of where the packets were dropped.	`briefterse`
`Rate (pps)`	Highest observed traffic rate for this packet type or protocol group.	`briefterse`
`Violation counts`	Number of violations of the policer bandwidth.	`briefterse`

Table 412: show ddos-protection protocols statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Violation state of the packet type: • ok—Policer has not been violated for this packet type • viol—Policer has been violated for this packet type	briefterse

Sample Output

show ddos-protection protocols statistics

```
user@host> show ddos-protection protocols statistics
Protocol Group: IPv4-Unclassified

  Packet type: aggregate
    System-wide information:
      Aggregate bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by individual policers: 0
    FPC slot 1 information:
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by individual policers: 0

Protocol Group: IPv6-Unclassified

  Packet type: aggregate
    System-wide information:
      Aggregate bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
```

```
    Routing Engine information:
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by individual policers: 0
    FPC slot 1 information:
      Aggregate policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by individual policers: 0


Protocol Group: PPPoE

  Packet type: aggregate
    System-wide information:
      Aggregate bandwidth is never violated
      Received:  61961244            Arrival rate:     4000 pps
      Dropped:   0                   Max arrival rate: 4002 pps
    Routing Engine information:
      Aggregate policer is never violated
      Received:  15488871            Arrival rate:     1001 pps
      Dropped:   0                   Max arrival rate: 1011 pps
        Dropped by individual policers: 0
    FPC slot 1 information:
      Aggregate policer is never violated
      Received:  61961244            Arrival rate:     4000 pps
      Dropped:   46473017            Max arrival rate: 4002 pps
        Dropped by individual policers: 46473017

  Packet type: padi
    System-wide information:
      Bandwidth is being violated!
        No. of FPCs currently receiving excess traffic: 1
        No. of FPCs that have received excess traffic:  1
        Violation first detected at: 2011-04-19 08:23:17 PDT
        Violation last seen at:      2011-04-19 12:41:23 PDT
        Duration of violation: 04:18:06 Number of violations: 1
      Received:  30980622            Arrival rate:     2000 pps
      Dropped:   23236505            Max arrival rate: 2001 pps
    Routing Engine information:
      Policer is never violated
      Received:  7744433             Arrival rate:     500 pps
```

```
      Dropped:   0                   Max arrival rate: 505 pps
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Policer is currently being violated!
        Violation first detected at: 2011-04-19 08:23:17 PDT
        Violation last seen at:      2011-04-19 12:41:23 PDT
        Duration of violation: 04:18:06 Number of violations: 1
      Received:  30980622            Arrival rate:     2000 pps
      Dropped:   23236505            Max arrival rate: 2001 pps
        Dropped by this policer: 23236505
        Dropped by aggregate policer: 0

  Packet type: pado
    System-wide information:
      Bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0

  Packet type: padr
    System-wide information:
      Bandwidth is being violated!
        No. of FPCs currently receiving excess traffic: 1
        No. of FPCs that have received excess traffic:  1
        Violation first detected at: 2011-04-19 08:23:17 PDT
        Violation last seen at:      2011-04-19 12:43:23 PDT
        Duration of violation: 04:20:06 Number of violations: 1
      Received:  31220846            Arrival rate:     2000 pps
      Dropped:   23416690            Max arrival rate: 2001 pps
    Routing Engine information:
      Policer is never violated
      Received:  7806417             Arrival rate:     499 pps
      Dropped:   0                   Max arrival rate: 506 pps
        Dropped by aggregate policer: 0
```

```
    FPC slot 1 information:
      Policer is currently being violated!
        Violation first detected at: 2011-04-19 08:23:17 PDT
        Violation last seen at:      2011-04-19 12:43:23 PDT
        Duration of violation: 04:20:06 Number of violations: 1
      Received:  31220846            Arrival rate:     2000 pps
      Dropped:   23416690            Max arrival rate: 2001 pps
        Dropped by this policer: 23416690
        Dropped by aggregate policer: 0

  Packet type: pads
    System-wide information:
      Bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0

  Packet type: padt
    System-wide information:
      Bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0

  Packet type: padm
```

```
    System-wide information:
      Bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0

  Packet type: padn
    System-wide information:
      Bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    Routing Engine information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0

...
```

show ddos-protection protocols statistics brief

```
user@host> show ddos-protection protocols statistics brief

Protocol    Packet      Received        Dropped        Rate     Violation State
group       type        (packets)       (packets)      (pps)    counts
ipv4-uncls  aggregate   0               0              0        0         ok
ipv6-uncls  aggregate   0               0              0        0         ok
```

```
dynvlan     aggregate   0           0           0       0       ok
ppp         aggregate   0           0           0       0       ok
ppp         unclass     0           0           0       0       ok
ppp         lcp         0           0           0       0       ok
ppp         auth        0           0           0       0       ok
ppp         ipcp        0           0           0       0       ok
ppp         ipv6cp      0           0           0       0       ok
ppp         mplscp      0           0           0       0       ok
ppp         isis        0           0           0       0       ok
pppoe       aggregate   61561238    0           4000    0       ok
pppoe       padi        30780619    23086506    2000    1       viol
pppoe       pado        0           0           0       0       ok
pppoe       padr        30780619    23086499    2000    1       viol
pppoe       pads        0           0           0       0       ok
pppoe       padt        0           0           0       0       ok
pppoe       padm        0           0           0       0       ok
pppoe       padn        0           0           0       0       ok
dhcpv4      aggregate   0           0           0       0       ok
dhcpv4      unclass..   0           0           0       0       ok
dhcpv4      discover    0           0           0       0       ok
dhcpv4      offer       0           0           0       0       ok
dhcpv4      request     0           0           0       0       ok
dhcpv4      decline     0           0           0       0       ok
dhcpv4      ack         0           0           0       0       ok
dhcpv4      nak         0           0           0       0       ok
dhcpv4      release     0           0           0       0       ok
dhcpv4      inform      0           0           0       0       ok
dhcpv4      renew       0           0           0       0       ok
dhcpv4      forcerenew  0           0           0       0       ok
dhcpv4      leasequery  0           0           0       0       ok
dhcpv4      leaseuna..  0           0           0       0       ok
dhcpv4      leaseunk..  0           0           0       0       ok
dhcpv4      leaseact..  0           0           0       0       ok
dhcpv4      bootp       0           0           0       0       ok
dhcpv4      no-msgtype  0           0           0       0       ok
dhcpv4      bad-pack..  0           0           0       0       ok



...

icmp        aggregate   0           0           0       0       ok
igmp        aggregate   0           0           0       0       ok
ospf        aggregate   0           0           0       0       ok
```

```
rsvp        aggregate   0               0               0        0         ok
pim         aggregate   0               0               0        0         ok
rip         aggregate   0               0               0        0         ok
ptp         aggregate   0               0               0        0         ok
bfd         aggregate   0               0               0        0         ok
lmp         aggregate   0               0               0        0         ok
ldp         aggregate   0               0               0        0         ok
msdp        aggregate   0               0               0        0         ok
bgp         aggregate   0               0               0        0         ok
vrrp        aggregate   0               0               0        0         ok
telnet      aggregate   0               0               0        0         ok

...
```

show ddos-protection protocols statistics terse

```
user@host> show ddos-protection protocols statistics terse
Protocol    Packet      Received        Dropped         Rate     Violation State
group       type        (packets)       (packets)       (pps)    counts
ipv4-uncls  aggregate   241             0               0        0         ok
icmp        aggregate   20              0               0        0         ok
igmp        aggregate   55              0               0        0         ok
ospf        aggregate   956             0               0        0         ok
rsvp        aggregate   784             0               0        0         ok
ldp         aggregate   2984            0               0        0         ok
bgp         aggregate   312             0               0        0         ok
lacp        aggregate   1744            0               0        0         ok
stp         aggregate   9791            0               0        0         ok
arp         aggregate   19              0               0        0         ok
pvstp       aggregate   393             0               0        0         ok
mlp         aggregate   624774          0               0        0         ok
mlp         packets     1714371         223937          0        3         ok
mcast-copy  aggregate   3018038         0               0        0         ok
igmp-snoop  aggregate   43              0               0        0         ok
fw-host     aggregate   95547           0               0        0         ok
uncls       aggregate   10000           0               0        0         ok
```

show ddos-protection protocols pppoe statistics

```
user@host> show ddos-protection protocols pppoe statistics
Protocol Group: PPPoE

  Packet type: aggregate
    System-wide information:
      Aggregate bandwidth is never violated
      Received:  60381200            Arrival rate:     4000 pps
      Dropped:   0                   Max arrival rate: 4002 pps
    Routing Engine information:
      Aggregate policer is never violated
      Received:  15095242            Arrival rate:     1001 pps
      Dropped:   0                   Max arrival rate: 1011 pps
        Dropped by individual policers: 0
    FPC slot 1 information:
      Aggregate policer is never violated
      Received:  60381200            Arrival rate:     4000 pps
      Dropped:   45287921            Max arrival rate: 4002 pps
        Dropped by individual policers: 45287921

  Packet type: padi
    System-wide information:
      Bandwidth is being violated!
        No. of FPCs currently receiving excess traffic: 1
        No. of FPCs that have received excess traffic:  1
        Violation first detected at: 2011-04-19 08:23:17 PDT
        Violation last seen at:      2011-04-19 12:34:48 PDT
        Duration of violation: 04:11:31 Number of violations: 1
      Received:  30190600            Arrival rate:     2000 pps
      Dropped:   22643960            Max arrival rate: 2001 pps
    Routing Engine information:
      Policer is never violated
      Received:  7547621             Arrival rate:     499 pps
      Dropped:   0                   Max arrival rate: 505 pps
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Policer is currently being violated!
        Violation first detected at: 2011-04-19 08:23:17 PDT
        Violation last seen at:      2011-04-19 12:34:48 PDT
        Duration of violation: 04:11:31 Number of violations: 1
      Received:  30190600            Arrival rate:     2000 pps
```

```
    Dropped:   22643960            Max arrival rate: 2001 pps
      Dropped by this policer: 22643960
      Dropped by aggregate policer: 0

Packet type: pado
  System-wide information:
    Bandwidth is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
  Routing Engine information:
    Policer is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
      Dropped by aggregate policer: 0
  FPC slot 1 information:
    Policer is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
      Dropped by aggregate policer: 0

Packet type: padr
  System-wide information:
    Bandwidth is being violated!
      No. of FPCs currently receiving excess traffic: 1
      No. of FPCs that have received excess traffic:  1
      Violation first detected at: 2011-04-19 08:23:17 PDT
      Violation last seen at:      2011-04-19 12:34:48 PDT
      Duration of violation: 04:11:31 Number of violations: 1
    Received:  30190600            Arrival rate:     2000 pps
    Dropped:   22643961            Max arrival rate: 2001 pps
  Routing Engine information:
    Policer is never violated
    Received:  7547621             Arrival rate:     501 pps
    Dropped:   0                   Max arrival rate: 506 pps
      Dropped by aggregate policer: 0
  FPC slot 1 information:
    Policer is currently being violated!
      Violation first detected at: 2011-04-19 08:23:17 PDT
      Violation last seen at:      2011-04-19 12:34:48 PDT
      Duration of violation: 04:11:31 Number of violations: 1
    Received:  30190600            Arrival rate:     2000 pps
    Dropped:   22643961            Max arrival rate: 2001 pps
      Dropped by this policer: 22643961
```

```
      Dropped by aggregate policer: 0

Packet type: pads
  System-wide information:
    Bandwidth is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
  Routing Engine information:
    Policer is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
      Dropped by aggregate policer: 0
  FPC slot 1 information:
    Policer is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
      Dropped by aggregate policer: 0

Packet type: padt
  System-wide information:
    Bandwidth is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
  Routing Engine information:
    Policer is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
      Dropped by aggregate policer: 0
  FPC slot 1 information:
    Policer is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
      Dropped by aggregate policer: 0

Packet type: padm
  System-wide information:
    Bandwidth is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
  :
    Policer is never violated
    Received:  0                   Arrival rate:     0 pps
    Dropped:   0                   Max arrival rate: 0 pps
```

```
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0

  Packet type: padn
    System-wide information:
      Bandwidth is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
    :
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
    FPC slot 1 information:
      Policer is never violated
      Received:  0                   Arrival rate:     0 pps
      Dropped:   0                   Max arrival rate: 0 pps
        Dropped by aggregate policer: 0
```

show ddos-protection protocols pppoe statistics brief

```
user@host> show ddos-protection protocols pppoe statistics brief
Protocol    Packet      Received        Dropped        Rate     Violation State
group       type        (packets)       (packets)      (pps)    counts
pppoe       aggregate   60901227        0              4000     0         ok
pppoe       padi        30450613        22838981       2000     1         viol
pppoe       pado        0               0              0        0         ok
pppoe       padr        30450614        22838977       2000     1         viol
pppoe       pads        0               0              0        0         ok
pppoe       padt        0               0              0        0         ok
pppoe       padm        0               0              0        0         ok
pppoe       padn        0               0              0        0         ok
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

clear ddos-protection protocols | 3482
show ddos-protection protocols | 19222
show ddos-protection protocols culprit-flows | 19239
show ddos-protection protocols flow-detection | 19249
show ddos-protection protocols parameters | 19265
show ddos-protection protocols violations | 19292

show ddos-protection protocols violations

IN THIS SECTION

- Syntax | 19292
- Description | 19292
- Options | 19293
- Required Privilege Level | 19293
- Output Fields | 19293
- Sample Output | 19294
- Release Information | 19295

Syntax

```
show ddos-protection protocols <protocol-group> violations
```

Description

Display information about control plane DDoS protection policer violations for all protocol groups or for a particular protocol group.

NOTE: Control plane DDoS protection policers act on the system's traffic queues. The QFX5100 and QFX5200 lines of switches manage traffic for more protocols than the number of queues, so the system often must map more than one protocol to the same queue. When traffic for one protocol shares a queue with other protocols and violates DDoS protection policer limits, this command reports a violation on that queue for all mapped protocols because the system doesn't distinguish which protocol's traffic specifically caused the violation. You can use what you know about the types of traffic flowing through your network to identify which of the reported protocols actually triggered the violation.

Options

none Display information for all protocol groups.

protocol-group (Optional) Name of a particular protocol group. See "show ddos-protection protocols" on page 19222 for a list of available groups.

Required Privilege Level

view

Output Fields

Table 413 on page 19293 lists the output fields for the `show ddos-protection protocols violations` command. Output fields are listed in the approximate order in which they appear.

Table 413: show ddos-protection protocols violations Output Fields

Field Name	Field Description
`Number of packet types that are being violated`	Number of individual policers and aggregate policers that are currently being violated
`Protocol Group`	Name of protocol group
`Packet type`	Name of packet type in protocol group

Table 413: show ddos-protection protocols violations Output Fields *(Continued)*

Field Name	Field Description
Bandwidth (pps)	Policer bandwidth
Arrival rate (pps)	Current traffic rate for packets arriving from all cards and at the Routing Engine
Peak rate (pps)	Highest traffic rate for packets arriving from all cards and at the Routing Engine
Policer bandwidth violation detected at	Timestamp of the policer violation
Detected on	Slot number of the card on which the violation was detected

Sample Output

show ddos-protection protocols violations

```
user@host> show ddos-protection protocols violations
Number of packet types that are being violated: 2
Protocol    Packet      Bandwidth  Arrival   Peak      Policer bandwidth
group       type        (pps)      rate(pps) rate(pps) violation detected at
pppoe       padi        500        2000      2001      2011-04-19 08:23:17 PDT
          Detected on: FPC-1
pppoe       padr        500        1999      2001      2011-04-19 08:23:17 PDT
          Detected on: FPC-1
```

show ddos-protection protocols lldp violations

```
user@host> show ddos-protection protocols lldp violations
Number of packet types that are being violated: 0
```

show ddos-protection protocols pppoe violations

```
user@host> show ddos-protection protocols pppoe violations
Number of packet types that are being violated: 2
Protocol    Packet      Bandwidth  Arrival   Peak      Policer bandwidth
group       type        (pps)      rate(pps) rate(pps) violation detected at
pppoe       padi        500        2000      2001      2011-04-19 08:23:17 PDT
          Detected on: FPC-1
pppoe       padr        500        1999      2001      2011-04-19 08:23:17 PDT
          Detected on: FPC-1
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

clear ddos-protection protocols | 3482

show ddos-protection protocols | 19222

show ddos-protection protocols culprit-flows | 19239

show ddos-protection protocols flow-detection | 19249

show ddos-protection protocols parameters | 19265

show ddos-protection protocols statistics | 19276

show ddos-protection statistics

IN THIS SECTION

- Syntax | 19296
- Description | 19296
- Options | 19296
- Required Privilege Level | 19296
- Output Fields | 19296

Sample Output | 19297
Release Information | 19298

Syntax

```
show ddos-protection statistics
```

Description

Display DDoS protection global statistics for bandwidth violations.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 414 on page 19296 lists the output fields for the `show ddos-protection statistics` command. Output fields are listed in the approximate order in which they appear.

Table 414: show ddos-protection statistics Output Fields

Field Name	Field Description
Policing on routing engine	Shows whether or not policing is enabled on the Routing Engine.
Policing on FPC	Shows whether or not policing is enabled on the line card.
Flow detection	Shows whether or not flow detection in enabled.

Table 414: show ddos-protection statistics Output Fields *(Continued)*

Field Name	Field Description
Logging	Shows whether or not DDoS event logging is enabled.
Policer violation report rate	Shows the violation report rate as a percentage.
Flow report rate	Shows the flow report rate as a percentage.
Default flow detection mode	Flow detection and tracking mode configured at the global level for all protocol groups and packet types.
Default flow level detection mode	Flow detection and tracking mode configured at the flow aggregation level for all protocol groups and packet types.
Default flow level control mode	Default behavior configured for how traffic in detected flows is controlled for all protocol groups and packet types.
`Currently violated packet types`	Number of packet types currently experiencing a bandwidth violation.
`Packet types have seen violations`	Number of packet types that have experienced a bandwidth violation since statistics were cleared.
`Total violation counts`	Total number of bandwidth violations.

Sample Output

show ddos-protection statistics

```
user@host> show ddos-protection statistics
DDOS protection global statistics:
  Policing on routing engine:         Yes
```

```
Policing on FPC:                    Yes
Flow detection:                     No
Logging:                            Yes
Policer violation report rate:      100
Flow report rate:                   100
Default flow detection mode         Automatic
Default flow level detection mode   Automatic
Default flow level control mode     Drop
Currently violated packet types:    2
Packet types have seen violations:  4
Total violation counts:             4
Currently tracked flows:            0
Total detected flows:               0
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

clear ddos-protection protocols | 3482
show ddos-protection protocols | 19222
show ddos-protection version | 19298

show ddos-protection version

IN THIS SECTION

- Syntax | 19299
- Description | 19299
- Options | 19299
- Required Privilege Level | 19299
- Output Fields | 19299
- Sample Output | 19300

Release Information | 19300

Syntax

```
show ddos-protection version
```

Description

Display the DDoS protection version and the total numbers of protocol groups and packet types that this version can be configured in this version.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 415 on page 19299 lists the output fields for the `show ddos-protection version` command. Output fields are listed in the approximate order in which they appear.

Table 415: show ddos-protection version Output Fields

Field Name	Field Description
`Version`	Version number of the DDoS protection code.
`Total protocol groups`	Number of protocol groups configured with DDoS protection.
`Total tracked packet types`	Number of protocol packet types configured with DDoS protection.

Sample Output

show ddos-protection version

```
user@host> show ddos-protection version
DDOS protection, Version 1.0
  Total protocol groups       = 83
  Total tracked packet types  = 154
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

- clear ddos-protection protocols | 3482
- show ddos-protection protocols | 19222
- show ddos-protection statistics | 19295

show dhcp client binding

IN THIS SECTION

- Syntax | 19301
- Description | 19301
- Options | 19301
- Required Privilege Level | 19302
- Output Fields | 19302
- Sample Output | 19303
- Release Information | 19304

Syntax

```
show dhcp client binding
[<address> |interface <interface-name>]
 routing-instance <routing-instance name>
[brief | detail | summary ]
logical-system
tenant
```

Description

Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

Options

address	(Optional) Display DHCP binding information for a specific client identified by one of the following entries: • ip-address—The specified IP address. • mac-address—The specified MAC address.
routing-instance <routing-instance name>	(Optional) Display DHCP binding information for DHCP clients on the specified routing instance.
interface <interface-name>	(Optional) Perform this operation on the specified interface.
brief	(Optional) Display brief information about the active client bindings.
detail	(Optional) Display detailed client binding information.
summary	(Optional) Display a summary of DHCP client information.
logical-system	(Optional) Displays the DHCP binding information for DHCP clients on the specified logical system.
tenant	(Optional) Displays the DHCP binding information for DHCP clients on the specified tenant system.

Required Privilege Level

view

Output Fields

Table 416 on page 19302 lists the output fields for the `show dhcp client binding` command. Output fields are listed in the approximate order in which they appear.

Table 416: show dhcp client binding Output Fields

Field Name	Field Description
IP address	IP address of the DHCP client.
Hardware address	Hardware address of the DHCP client.
Server	IP address of the DHCP server.
Expires	Number of seconds in which the lease expires.
State	State of the address binding table on the DHCP local server.
Interface	Interface on which the request was received.
Lease Expires	Date and time at which the client's IP address lease expires.
Lease Expires in	Number of seconds in which the lease expires.
Lease Start	Date and time at which the client's IP address lease started.
Vendor Identifier	Vendor identifier.
Server Identifier	IP address of the DHCP server.
Client IP Address	IP address of the DHCP client.

Sample Output

show dhcp client binding

```
user@host> show dhcp client binding
2 clients, (2 bound,  0 init, 0 discover, 0 renew, 0 rebind)

             IP address      Hardware address     Server      Expires     State      Interface
             10.1.1.89       00:0a:12:00:12:12    10.1.1.1    348         BOUND      fe-0/0/1.0
             20.1.1.90       00:0a:12:00:12:34    20.1.1.1    568         BOUND      fe-0/0/2.0
```

command-name

```
user@host> show dhcp client binding interface fe-0/0/1.0 detail
Client Interface: fe-0/0/1.0
               Hardware address:       00:0a:12:00:12:12
               State:                  BOUND
               Lease Expires:          2010-09-16 14:45:41 UTC
               Lease Expires in:       528 seconds
               Lease Start:            2010-09-16 14:35:41 UTC
               Vendor Identifier:      ether
               Server Identifier:      10.1.1.1
               Client IP Address:      10.1.1.89
               update server           enabled

           DHCP Options :
               Name: name-server, Value: [ 10.209.194.131, 198.51.110.2, 192.0.2.3 ]
               Name: server-identifier, Value: 10.1.1.1
               Name: router, Value: [ 10.1.1.80 ]
               Name: domain-name, Value: example-50
```

command-name

```
user@host> show dhcp client binding 10.1.1.89
IP address      Hardware address     Server      Expires     State      Interface
         10.1.1.89       00:0a:12:00:12:12    10.1.1.1    348         BOUND      fe-0/0/1.0
```

command-name

```
user@host> show dhcp client binding tenant TSYS1 routing-instance R1
IP address        Hardware address   Expires     State      Interface
       33.33.33.3        00:50:56:b0:b8:21  1628        BOUND      ge-0/0/3.0
```

command-name

```
user@host> show dhcp client binding detail tenant TSYS1 routing-instance R1
Client Interface/Id: ge-0/0/3.0
     Hardware Address:             00:50:56:b0:b8:21
     State:                        BOUND(LOCAL_CLIENT_STATE_BOUND)
     Lease Expires:                2018-04-26 16:24:34 UTC
     Lease Expires in:             1626 seconds
     Lease Start:                  2018-04-26 15:51:14 UTC
     Server Identifier:            11.11.11.1
     Client IP Address:            33.33.33.3
     Update Server                 No

    DHCP options:
    Name: dhcp-lease-time, Value: 33 minutes, 20 seconds
    Name: server-identifier, Value: 11.11.11.1
    Name: router, Value: [ 33.33.33.1 ]
    Name: subnet-mask, Value: 255.255.255.0
```

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

The `logical-system` and `tenant` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

clear dhcp client binding | 3484

show dhcp client statistics

IN THIS SECTION

- Syntax | 19305
- Description | 19305
- Options | 19305
- Required Privilege Level | 19306
- Output Fields | 19306
- Sample Output | 19307
- Release Information | 19308

Syntax

```
show dhcp client statistics
<routing-instance routing-instance-name >
logical-system
tenant
```

Description

Displays the DHCP client statistics.

Options

routing-instance routing-instance-name	(Optional) Display the statistics for DHCP clients on the specified routing instance.
logical-system	(Optional) Displays the DHCP statistics information for DHCP clients on the specified logical system.
tenant	(Optional) Displays the DHCP statistics information for DHCP clients on the specified tenant system.

Required Privilege Level

view

Output Fields

Table 417 on page 19306 lists the output fields for the `show dhcp client statistics` command. Output fields are listed in the approximate order in which they appear.

Table 417: show dhcp client statistics

Field Name	Field Description
Packets dropped	Number of packets discarded by the DHCP local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.
Messages received	Number of DHCP messages received. • BOOTREPLY—Number of BOOTP protocol data units (PDUs) received • DHCPOFFER—Number of DHCP PDUs of type OFFER received • DHCPACK—Number of DHCP PDUs of type ACK received • DHCPNACK—Number of DHCP PDUs of type NACK received • DHCPFORCERENEW—Number of DHCP PDUs of type FORCERENEW received

Table 417: show dhcp client statistics *(Continued)*

Field Name	Field Description
Messages sent	Number of DHCP messages sent. • BOOTREQUEST—Number of BOOTP protocol data units (PDUs) transmitted • DHCPDECLINE—Number of DHCP PDUs of type DECLINE transmitted • DHCPDISCOVER—Number of DHCP PDUs of type DISCOVER transmitted • DHCPREQUEST—Number of DHCP PDUs of type REQUEST transmitted • DHCPINFORM—Number of DHCP PDUs of type INFORM transmitted • DHCPRELEASE—Number of DHCP PDUs of type RELEASE transmitted • DHCPRENEW—Number of DHCP PDUs of type RENEW transmitted • DHCPREBIND—Number of DHCP PDUs of type REBIND transmitted

Sample Output

show dhcp client statistics

```
user@host> show dhcp client statistics
Packets dropped:
    Total                     0
Messages received:
    BOOTREPLY                 0
    DHCPOFFER                 0
    DHCPACK                   0
    DHCPNAK                   0
    DHCPFORCERENEW            0
```

```
Messages sent:
    BOOTREQUEST                0
    DHCPDECLINE                0
    DHCPDISCOVER               0
    DHCPREQUEST                0
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPRENEW                  0
    DHCPREBIND                 0
```

show dhcp client statistics tenant TSYS1 routing-instance R1

```
user@host> show dhcp client statistics tenant TSYS1 routing-instance R1
Packets dropped:
    Total                      0

   Messages received:
    BOOTREPLY                  14
    DHCPOFFER                  4
    DHCPACK                    10
    DHCPNAK                    0
    DHCPFORCERENEW             0

  Messages sent:
    BOOTREQUEST                17
    DHCPDECLINE                0
    DHCPDISCOVER               4
    DHCPREQUEST                10
    DHCPINFORM                 0
    DHCPRELEASE                3
    DHCPRENEW                  6
    DHCPREBIND                 0
```

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

The `logical-system` and `tenant` option is introduced in Junos OS Release 18.4R1.

NOTE: This command is not supported in Junos OS Evolved because it is not needed. Stateless DHCPv4 and DHCPv6 relay are supported and enabled by default, so Junos OS Evolved does not maintain states or records of statistics or bindings for DHCP clients.

RELATED DOCUMENTATION

clear dhcp client statistics | 3486

show dhcp relay active-leasequery

IN THIS SECTION

- Syntax | 19309
- Description | 19310
- Options | 19310
- Required Privilege Level | 19310
- Output Fields | 19310
- Sample Output | 19313
- Release Information | 19315

Syntax

```
show dhcp relay active-leasequery
<details | summary>
<interface interface-name>
<logical-system logical-system-name>
<peer ip-address>
<routing-instance routing-instance-name>
<statistics>
```

Description

Display information about DHCPv4 active leasequery peer relay agents.

Options

details (Optional) Display the topology discovery translation table for the specified peer. You must also specify the `peer` *`ip-address`* option.

summary (Optional) Display summary information for all active leasequery peers. The `summary` option produces the same output as not specifying any option. You can also specify the `logical-system` *`logical-system-name`* option or the `routing-instance` *`routing-instance-name`* option with the `summary` option. You cannot specify any other option with the `summary` option.

interface ***interface-name*** (Optional) Display active leasequery statistics for a specific access interface. You must also specify the `statistics` option.

logical-system ***logical-system-name*** (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.

peer ***ip-address*** (Optional) Display information about active leasequery peer relay agents. You must also specify either the `details` option or the `statistics` option.

routing-instance ***routing-instance-name*** (Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.

statistics (Optional) Display active leasequery statistics for a specific active leasequery peer relay agent or a specific access interface. You must also specify either the `peer` *`ip-address`* option or the `interface` *`interface-name`* option.

Required Privilege Level

view

Output Fields

Table 418 on page 19311 lists the output fields for the `show dhcp relay active-leasequery` command. Output fields are listed in the approximate order in which they appear.

Table 418: show dhcp relay active-leasequery Output Fields

Field Name	Field Description
`peer`	IP address of topology discovery peer relay agent.
`Connected Peers`	Number of access interfaces that have completed the topology discovery process and are associated with the peer. Each interface is in the Connected state.
`Connecting Peers`	Number of access interfaces that are configured with topology discovery but have not completed the topology discovery process. Consequently these interfaces are not yet associated with any peer. Each interface is in the Connecting state. When topology discovery completes and the interface is associated with a peer, the state moves to Connected.
`Local Circuit ID`	The local access interface for the specified peer.
`Remote Circuit ID`	The remote access interface on another peer relay agent that corresponds to the specified peer's local access interface.
`Local Interface Address`	IP address of the local access interface.
`State`	State of the topology discovery process. • `Done`—Topology discovery has completed for the specified peer.
`Redundancy State`	M:N redundancy state associated with the peer or interface. • `Backup`—The specified peer or interface is currently in backup mode. This means that the BNG hosting the relay agent is the current backup BNG for the group. • `Master`—The specified peer or interface is currently in primary mode. This means that the BNG hosting the relay agent is the current primary BNG for the group. • `Unknown`—The redundancy state of the peer or interface is unknown.
`xid`	Randomly generated, temporary transaction ID for the topology discovery query sent for the local access interface.

Table 418: show dhcp relay active-leasequery Output Fields *(Continued)*

Field Name	Field Description
`Remote ALQ Status`	State of the active leasequery process. • `Done`—The active leasequery process has completed. Subscriber state and binding information has been synchronized for subscriber groups that use the local access interface. • `Queued`—The active leasequery is queued. Subscriber state and binding information are not yet synchronized/ • `Unknown`—Active leasequery process state is unknown.
`Interface`	Name of the specified address for which statistics are displayed.
`Topology-Discover Configured`	Indicates whether topology discovery has been configured on the specified peer or interface, `Yes` or `No`.
`Bindings Sent`	Number of DHCP bindings sent based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings sent over the interface. For the peer, it's the count of all bindings sent over all interfaces that belong to the peer.
`Bindings Received`	Number of DHCP bindings received based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings received over the interface. For the peer, it's the count of all bindings received over all interfaces that belong to the peer.
`Bindings Installed Successfully`	Number of DHCP bindings successfully installed based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings successfully installed over the interface. For the peer, it's the count of all bindings successfully installed over all interfaces that belong to the peer.

Table 418: show dhcp relay active-leasequery Output Fields *(Continued)*

Field Name	Field Description
`Bindings Failed to Install`	Number of DHCP bindings that failed to install based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings that failed over the interface. For the peer, it's the count of all bindings that failed over all interfaces that belong to the peer.
`Last Synchronization Time`	Time stamp for the last time synchronization for subscribers occurred for the specified interface or peer.
`ALQ Transmit Buffer count`	The maximum TCP received and transmitted bytes buffer for the jdhcpd process. The value is nonconfigurable and is always 65,535 (ffff).
`Max Leasequery Transmit Rate`	The transmit rate of bulk leasequery replies per second. It is the maximum number of bulk leasequery replies per one-second interval. The value is nonconfigurable and is always 60. For example, if 600 leases are synchronized to the peer, it takes 10 seconds to complete.
`Local Interface Count`	Number of local access interfaces for the specified peer or interface. For the interface, the value is always 1.
`Remote Interface Count`	Number of remote access interfaces for the specified peer or interface. For the interface, the value is always 1.

Sample Output

show dhcp relay active-leasequery (Summary)

```
user@host> show dhcp relay active-leasequery summary

            Active Lease-query peer List:

            peer                  Connected Peers          Connecting Peers
```

```
            192.0.2.2             1000                    0
            198.51.100.2          1000                    0
            10:80:3::2            2                       0
```

show dhcp relay active-leasequery (IPv4 Peer Details)

```
user@host> show dhcp relay active-leasequery peer 192.0.2.2 details

 Local Circuit-ID   Remote Circuit-ID   Local Interface Address   State   Redundancy State
xid        Remote ALQ Status
 ge-2/1/1.1020      ge-0/0/6.1020       192.0.2.235               Done    Backup
0xb9dc35   Done
```

show dhcp relay active-leasequery (IPv6 Peer Details)

```
user@host> show dhcp relay active-leasequery peer 2001:db8:1:3::20 details

Local Circuit-ID    Remote Circuit-ID   Local Interface Address   State   Redundancy State
xid         Remote ALQ Status
 ps3.0               ps3.0              2001:db8:80:1::1          Done    Master
0x663345    Done
 ps4.0               ps4.0              2001:db8:90:1::2          Done    Master
0xb68623    Done
```

show dhcp relay active-leasequery (Peer Details Pseudowire Interfaces)

```
user@host> show dhcp relay active-leasequery statistics interface peer 198.51.100.1 details

 Local Circuit-ID    Remote Circuit-ID   Local Interface Address   State   Redundancy State
xid         Remote ALQ Status
 ps1.0               ps1.0               198.51.100.1              Done    Backup
0x5da771    Done
 ps2.0               ps2.0               198.51.100.1              Done    Backup
0x7c7859    Unknown
```

show dhcp relay active-leasequery (Peer Statistics)

```
user@host> show dhcp relay active-leasequery statistics peer 192.0.2.2 routing-instance RI-test-vrf

     peer : 192.0.2.2
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 111
     Bindings Received                      : 0
     Bindings Installed Successfully        : 0
     Bindings Failed to install             : 0
     Last Synchronization Time              : 2019-02-15 16:28:36 IST
     ALQ Transmit Buffer count              : ffff
     Max Leasequery Transmit Rate           : 60
     Local Interface count                  : 4
     Remote Interface count                 : 4
```

show dhcp relay active-leasequery (Interface Statistics)

```
user@host> show dhcp relay active-leasequery statistics interface ge-2/1/1.1020 routing-instance R1-test-vrf

     Interface : ge-2/1/1.1020
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 400
     Bindings Received                      : 0
     Bindings Installed Successfully        : 0
     Bindings Failed to install             : 0
     Last Synchronization Time              : 2019-02-15 16:20:05 IST
     ALQ Transmit Buffer count              : ffff
     Max Leasequery Transmit Rate           : 60
     Local Interface count                  : 1
     Remote Interface count                 : 1
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear dhcp relay active-leasequery statistics | 3488

show dhcpv6 relay active-leasequery | 19398

show dhcp relay binding (Security)

IN THIS SECTION

- Syntax | 19316
- Description | 19316
- Options | 19316
- Required Privilege Level | 19317
- Output Fields | 19317
- Sample Output | 19318
- Release Information | 19319

Syntax

```
Show dhcp relay binding
[<address> |interface <interface-name>]
routing-instance <routing-instance name>
[brief | detail | summary]
```

Description

Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) relay client table.

Options

address (Optional) Display DHCP binding information for a specific client identified by one of the following entries:

- ip-address—The specified IP address.
- mac-address—The specified MAC address.

routing-instance <routing-instance name>	(Optional) Display DHCP binding information on the specified routing instance.
interface <interface-name>	(Optional) Perform this operation on the specified interface.
brief	(Optional) Display brief information about the active client bindings.
detail	(Optional) Display detailed client binding information.
summary	(Optional) Display a summary of DHCP client information.

Required Privilege Level

view

Output Fields

Table 419 on page 19317 lists the output fields for the `show dhcp relay binding` command. Output fields are listed in the approximate order in which they appear.

Table 419: show dhcp relay binding Output Fields

Field Name	Field Description
IP address	IP address of the DHCP client.
Hardware address	Hardware address of the DHCP client.
Request received on	Interface on which the request was received.
Type	Type of DHCP packet processing performed on the device.
Obtained at	Date and time at which the client's IP address lease started.
Expires at	Date and time at which the client's IP address lease expires.

Table 419: show dhcp relay binding Output Fields *(Continued)*

Field Name	Field Description
State	State of the address binding table on the DHCP local server.

Sample Output

show dhcp relay binding

```
user@host> show dhcp relay binding detail
IP address       Hardware address   Type     Lease expires            State
       100.20.32.1      90:00:00:01:00:01  active   2007-01-17 11:38:47 PST  rebind
       100.20.32.3      90:00:00:02:00:01  active   2007-01-17 11:38:41 PST  rebind
       100.20.32.4      90:00:00:03:00:01  active   2007-01-17 11:38:01 PST  rebind
       100.20.32.5      90:00:00:04:00:01  active   2007-01-17 11:38:07 PST  rebind
       100.20.32.6      90:00:00:05:00:01  active   2007-01-17 11:38:47 PST  rebind
```

command-name

```
user@host> show dhcp relay binding 100.20.32.1
Active binding information:
           IP address               100.20.32.1
           Hardware address         90:00:00:01:00:01

       Lease information:
          Type                DHCP
          Obtained at         2007-01-17 11:28:47 PST
          Expires at          2007-01-17 11:38:47 PST

       > show dhcp relay binding 100.20.32.1 detail
       Active binding information:
           IP address              100.20.32.1
           Hardware address        90:00:00:01:00:01
           Request received on     fe-0/0/2.0, relayed by 100.20.32.2

       Lease information:
```

```
            Type                DHCP
            Obtained at         2007-01-17 11:28:47 PST
            Expires at          2007-01-17 11:38:47 PST
            State               rebind
```

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

clear dhcp relay binding (Security) | 3491

show dhcp relay binding

IN THIS SECTION

- Syntax | 19319
- Description | 19320
- Options | 19320
- Required Privilege Level | 19321
- Output Fields | 19321
- Sample Output | 19324
- Release Information | 19327

Syntax

```
show dhcp relay binding
<address>
<brief>
<detail>
<interface interface-name>
```

```
<interface-tag interface-tag-name>
<interfaces-vlan>
<interfaces-wildcard>
<ip-address | mac-address>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<summary>
```

Description

Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

Options

address	(Optional) Display DHCP binding information for a specific client identified by one of the following entries: • *ip-address*—The specified IP address. • *mac-address*—The specified MAC address. • *session-id*—The specified session ID.
brief	(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as `show dhcp relay binding`.
detail	(Optional) Display detailed client binding information.
interface *interface-name*	(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.
interface-tag *interface-tag-name*	(Optional) Displays the binding information for the specified interface-tag.
interfaces-vlan	(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.
interfaces-wildcard	(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).
logical-system *logical-system-name*	(Optional) Perform this operation on the specified logical system.
routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance.

summary (Optional) Display a summary of DHCP client information.

Required Privilege Level

view

Output Fields

Table 420 on page 19321 lists the output fields for the `show dhcp relay binding` command. Output fields are listed in the approximate order in which they appear.

Table 420: show dhcp relay binding Output Fields

Field Name	Field Description	Level of Output
number clients,(*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* rebinding, *number* releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary
IP address	IP address of the DHCP client.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Generated Remote ID	Remote ID generated by the Option 82 Agent Remote ID (suboption 1)	detail
Hardware address	Hardware address of the DHCP client.	brief detail
Expires	Number of seconds in which the lease expires.	brief detail

Table 420: show dhcp relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the DHCP relay address binding table on the DHCP client: • BOUND—Client has an active IP address lease. • INIT—Initial state. • REBINDING—Client is broadcasting a request to renew the IP address lease. • RELEASE—Client is releasing the IP address lease. • RENEWING—Client is sending a request to renew the IP address lease. • REQUESTING—Client is requesting a DHCP server. • SELECTING—Client is receiving offers from DHCP servers.	brief detail
Interface	Incoming client interface.	brief
Lease Expires	Date and time at which the client's IP address lease expires.	detail
Lease Expires in	Number of seconds in which the lease expires.	detail
Lease Start	Date and time at which the client's IP address lease started.	detail
Lease time violated	Lease time violation has occurred.	detail
Incoming Client Interface	Client's incoming interface.	detail
Interface tag	Name of the Interface tag.	detail

Table 420: show dhcp relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Server IP Address`	IP address of the DHCP server.	detail
`Server Interface`	Interface of the DHCP server.	detail
`Bootp Relay Address`	IP address of BOOTP relay.	detail
`Type`	Type of DHCP packet processing performed on the router: • `active`—Router actively processes and relays DHCP packets. • `passive`—Router passively snoops DHCP packets passing through the router.	All levels
`Lease expires at`	Date and time at which the client's IP address lease expires.	All levels
`Dual Stack Group`	Name of dual stack that is configured with the DHCP binding.	detail
`Dual Stack Peer Prefix`	Prefix of dual stack DHCPv6 peer.	detail
`Dual Stack Peer Address`	Address of the dual stack DHCPv6 peer.	detail

NOTE: When you configure dhcp relay and EVPN multihoming, you may notice that the CLI output for the `show dhcp relay binding` command may not display the physical interface and vlan id information.

Sample Output

show dhcp relay binding

```
user@host> show dhcp relay binding
IP address        Session Id  Hardware address   Expires     State      Interface
192.51.100.11     41          00:00:5e:00:53:01  86371       BOUND      ge-1/0/0.0
192.51.100.12     42          00:00:5e:00:53:02  86371       BOUND      ge-1/0/0.0
192.51.100.13     43          00:00:5e:00:53:03  86371       BOUND      ge-1/0/0.0
192.51.100.14     44          00:00:5e:00:53:04  86371       BOUND      ge-1/0/0.0
192.51.100.15     45          00:00:5e:00:53:05  86371       BOUND      ge-1/0/0.0
```

show dhcp relay binding detail

```
user@host> show dhcp relay binding detail



Client IP Address:  192.51.100.11
     Hardware Address:             00:00:5e:00:53:01
     State:                        BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 11:00:06 PDT
     Lease Expires in:             86361 seconds
     Lease Start:                  2009-07-20 11:00:06 PDT
     Lease time violated:          yes
     Last Packet Received:         2009-07-20 11:00:06 PDT
     Incoming Client Interface:    ge-1/0/0.0
     Interface tag:                None
     Server Ip Address:            192.51.100.22
     Server Interface:             none
     Bootp Relay Address:          192.51.100.32
     Session Id:                   41
     Dual Stack Group:             dual-stack-retail6
     Dual Stack Peer Prefix:       2001:db8:0:4::/64
     Dual Stack Peer Address:      2001:db8:1:0:8003::1/128

Client IP Address:  192.51.100.12
     Hardware Address:             00:00:5e:00:53:02
     State:                        BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 11:00:06 PDT
     Lease Expires in:             86361 seconds
```

```
    Lease Start:                  2009-07-20 11:00:06 PDT
    Last Packet Received:         2009-07-20 11:00:06 PDT
    Incoming Client Interface:    ge-1/0/0.0
    Interface tag:                None
    Server Ip Address:            192.51.100.22
    Server Interface:             none
    Bootp Relay Address:          192.51.100.32
    Session Id:                   42
    Generated Remote ID           host:ge-1/0/0:100
```

show dhcp relay binding interface

```
user@host> show dhcp relay binding interface fe-0/0/2

IP address       Hardware address   Type     Lease expires at
192.51.100.1      00:00:5e:00:53:01  active   2007-03-27 15:06:20 EDT
```

show dhcp relay binding interface vlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:100

IP address        Session Id  Hardware address   Expires     State      Interface
192.51.100.15      6           00:00:5e:00:53:94  86124       BOUND      ge-1/1/0:100
```

show dhcp relay binding interface svlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:10-100

IP address        Session Id  Hardware address   Expires     State      Interface
192.51.100.16      7           00:00:5e:00:53:92  86124       BOUND      ge-1/1/0:10-100
```

show dhcp relay binding interface-tag <*interface-tag-name*>

```
user@host> show dhcp relay binding interface-tag sample_tag
```

```
IP address        Session Id  Hardware address   Expires     State      Interface
192.2.1.1          2           00:10:94:00:00:01     740       BOUND     ge-0/0/1.3221225472
```

show dhcp relay binding interface-tag *interface-tag-name* detail

```
user@host> show dhcp relay binding interface-tag AGF_IFD detail
     Client IP Address:  192.168.0.100
     Hardware Address:             00:34:44:44:44:44
     State:                        BOUND(RELAY_STATE_BOUND)
     Lease Expires:                2023-05-02 11:22:44 PDT
     Lease Expires in:             2575 seconds
     Lease Start:                  2023-05-02 09:52:44 PDT
     Last Packet Received:         2023-05-02 10:22:44 PDT
     Incoming Client Interface:    demux0.3221225472
     Interface tag name:           AGF_IFD
     Client Interface Svlan Id:    1
     Client Interface Vlan Id:     1
     Demux Interface:              demux0.3221225491
     Server Ip Address:            192.168.0.50
     Server Interface:             none
     Bootp Relay Address:          192.168..1
     Session Id:                   8
     Client Profile Name:          dhcp-profile
     Generated Circuit ID:         aci1
     Generated Remote ID:          ari1
     Relay Id Length:              31
     Relay Id:                     /0x00020000/0x00000583/0x01000000/0x00000000
     Relay Id:                     /0x30303a32/0x363a3838/0x3a64373a/0x66373a
```

show dhcp relay binding ip-address

```
user@host> show dhcp relay binding 192.51.100.13
IP address        Session Id  Hardware address   Expires     State      Interface
192.51.100.13      43          00:00:5e:00:53:03  86293       BOUND      ge-1/0/0.0
```

show dhcp relay binding mac-address

```
user@host> show dhcp relay binding 00:00:5e:00:53:05
IP address        Session Id  Hardware address   Expires     State      Interface
192.51.100.15      45          00:00:5e:00:53:05  86279       BOUND      ge-1/0/0.0
```

show dhcp relay binding session-id

```
user@host> show dhcp relay binding 41
 IP address        Session Id  Hardware address   Expires     State      Interface
192.51.100.11      41          00:00:5e:00:53:53  86305       BOUND      ge-1/0/0.0
```

show dhcp relay binding <interfaces-vlan>

```
user@host> show dhcp relay binding ge-1/0/0:100-200
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.17       42          00:00:5e:00:53:02  86346       BOUND      ge-1/0/0.1073741827
192.168.0.16       41          00:00:5e:00:53:01  86346       BOUND      ge-1/0/0.1073741827
```

show dhcp relay binding <interfaces-wildcard>

```
user@host> show dhcp relay binding ge-1/3/*
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.9        24          00:00:5e:00:53:04  86361       BOUND      ge-1/3/0.110
192.168.0.8        23          00:00:5e:00:53:03  86361       BOUND      ge-1/3/0.110
192.168.0.7        22          00:00:5e:00:53:02  86361       BOUND      ge-1/3/0.110
```

show dhcp relay binding summary

```
user@host> show dhcp relay binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 rebinding, 0 releasing)
```

Release Information

Command introduced in Junos OS Release 8.3.

Options `interface` and *`mac-address`* added in Junos OS Release 8.4.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings

clear dhcp relay binding | 3493

show dhcp relay lockout-entries

IN THIS SECTION

- Syntax | 19328
- Description | 19328
- Options | 19329
- Required Privilege Level | 19329
- Output Fields | 19329
- Sample Output | 19330
- Release Information | 19331

Syntax

```
show dhcp relay lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv4 relay agent lockout database.

Options

all Display all client entries in the lockout database.

index Number identifying a client entry to be displayed.

Required Privilege Level

view

Output Fields

Table 421 on page 19329 lists the output fields for the `show dhcp relay lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 421: show dhcp relay lockout-entries Output Fields

Field Name	Field Description	Level of Output
`Index`	Number identifying a specific entry in the lockout database.	`all` and `index`
`Key`	Client identifier for the client in the lockout database.	`all` and `index`
`State`	Type of lockout period for the entry: • `Grace`—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. • `Lockout`—Client is currently locked out; attempts to establish a session are rejected.	`all` and `index`
`Expires (s)`	Number of seconds until the current lockout period expires.	`all` only

Table 421: show dhcp relay lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Elapsed (s)`	Number of seconds since the current lockout or grace timer started.	`all` only
`Count`	Number of consecutive times the client has been locked out.	`all` only
`Expires`	Date and time when the current lockout period ends.	`index` only
`Expires in`	Number of seconds until the current period expires.	`index` only
`Lockout count`	Number of consecutive times client has been locked out.	`index` only
`Next lockout time`	Duration of the next lockout period for this client.	`index` only
`Min lockout time`	Minimum duration for a lockout period; the initial lockout time.	`index` only
`Lockout reason`	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	`index` only

Sample Output

show dhcp relay lockout-entries (All Entries)

```
user@host> show dhcp relay lockout-entries all
Index     Key               State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

show dhcp relay lockout-entries (Specific Entry)

```
user@host> show dhcp relay lockout-entries index 2
     Index:                                  2
     Key:                                    default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                                  Lockout
     Expires:                                2018-03-29 19:06:17 IST
     Expires in:                             87
     Lockout count:                          1
     Next lockout time:                      200
     Min lockout time:                       100
     Lockout reason:                         181
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

- clear dhcp relay lockout-entries | 3496
- show dhcp server lockout-entries | 19377
- show dhcpv6 relay lockout-entries | 19417
- show dhcpv6 server lockout-entries | 19446

show dhcp relay statistics (Security)

IN THIS SECTION

- Syntax | 19332
- Description | 19332
- Options | 19332
- Required Privilege Level | 19332
- Output Fields | 19332

Sample Output | 19333
Release Information | 19334

Syntax

```
show dhcp relay statistics
[<routing-instance>]
```

Description

Display Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

routing-instance (Optional) Display the DHCP relay statistics on the specified routing instance.

Required Privilege Level

view

Output Fields

Table 422 on page 19333 lists the output fields for the `show dhcp relay statistics` command. Output fields are listed in the approximate order in which they appear.

Table 422: show dhcp relay statistics

Field Name	Field Description
Messages received	Number of DHCP messages sent. • BOOTREQUEST—Number of BOOTP protocol data units (PDUs) received • DHCPDECLINE—Number of DHCP PDUs of type DECLINE received • DHCPDISCOVER—Number of DHCP PDUs of type DISCOVER received • DHCPREQUEST—Number of DHCP PDUs of type REQUEST received • DHCPINFORM—Number of DHCP PDUs of type INFORM received • DHCPRELEASE—Number of DHCP PDUs of type RELEASE received
Messages sent	Number of DHCP messages received. • BOOTREPLY—Number of BOOTP PDUs transmitted • DHCPOFFER—Number of DHCP PDUs of type OFFER transmitted • DHCPACK—Number of DHCP PDUs of type ACK transmitted • DHCPNACK—Number of DHCP PDUs of type NACK transmitted • DHCPFORCERENEW—Number of DHCP PDUs of type FORCERENEW transmitted

Sample Output

show dhcp relay statistics

```
user@host> show dhcp relay statistics
Messages received:
    BOOTREQUEST                0
    DHCPDECLINE                0
    DHCPDISCOVER               0
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                0
```

```
Messages sent:
    BOOTREPLY                  0
    DHCPOFFER                  0
    DHCPACK                    0
    DHCPNAK                    0
    DHCPFORCERENEW             0
```

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

clear dhcp relay statistics (Security) | 3499

show dhcp relay statistics

IN THIS SECTION

- Syntax | 19334
- Syntax | 19335
- Description | 19335
- Options | 19335
- Required Privilege Level | 19335
- Output Fields | 19335
- Sample Output | 19339
- Release Information | 19342

Syntax

```
show dhcp relay statistics
<bulk-leasequery-connections>
<leasequery>
```

```
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Syntax

Syntax for EX Series switches:

```
show dhcp relay statistics
<routing-instance routing-instance-name>
```

Description

Display Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

bulk-leasequery-connections	(Optional) Display information about DHCP relay bulk leasequery statistics.
leasequery	(Optional) Display information about DHCP relay individual leasequery statistics.
logical-system *logical-system-name*	(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.
routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 423 on page 19336 lists the output fields for the `show dhcp relay statistics` command. Output fields are listed in the approximate order in which they appear.

Table 423: show dhcp relay statistics Output Fields

Field Name	Field Description
`Packets dropped`	Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the `Packets dropped` output. When all of the Packets dropped statistics are 0 (zero), only the `Total` field appears. • `Total`—Total number of packets discarded by the extended DHCP relay agent application. • `Leasequery reply with older timestamp`— Number of packets dropped when the binding information on the device has latest information compared to that of received in ALQ. • `Leasequery reply aborted`— Number of packets dropped when events like release or authorize disconnect are received during lease query processing. • `Bad hardware address`—Number of packets discarded because an invalid hardware address was specified. • `Bad opcode`—Number of packets discarded because an invalid operation code was specified. • `Bad options`—Number of packets discarded because invalid options were specified. • `Invalid server address`—Number of packets discarded because an invalid server address was specified. • `Lease Time Violation`—Number of packets discarded because of a lease time violation • `No available addresses`—Number of packets discarded because there were no addresses available for assignment. • `No interface match`—Number of packets discarded because they did not belong to a configured interface. • `No routing instance match`—Number of packets discarded because they did not belong to a configured routing instance. • `No valid local address`—Number of packets discarded because there was no valid local address. • `Packet too short`—Number of packets discarded because they were too short. • `Read error`—Number of packets discarded because of a system read error. • `Send error`— Number of packets that the extended DHCP relay application could not send to the receivers through the FreeBSD sendmsg ( ) function.

Table 423: show dhcp relay statistics Output Fields *(Continued)*

Field Name	Field Description
	• `Option 60`—Number of packets discarded containing DHCP option 60 vendor-specific information. • `Option 82`—Number of packets discarded because DHCP option 82 information could not be added.
`Messages received`	Number of DHCP messages received. • `BOOTREQUEST`—Number of BOOTP protocol data units (PDUs) received • `DHCPDECLINE`—Number of DHCP PDUs of type DECLINE received • `DHCPDISCOVER`—Number of DHCP PDUs of type DISCOVER received • `DHCPINFORM`—Number of DHCP PDUs of type INFORM received • `DHCPRELEASE`—Number of DHCP PDUs of type RELEASE received • `DHCPREQUEST`—Number of DHCP PDUs of type REQUEST received • `DHCPLEASEACTIVE`—Number of active DHCP leases • `DHCPLEASEUNASSIGNED`—Number of DHCP leases that are managed by the server but have not yet been assigned • `DHCPLEASEUNKNOWN`—Number of unknown DHCP leases • `DHCPLEASEQUERYDONE`—The leasequery is complete

Table 423: show dhcp relay statistics Output Fields *(Continued)*

Field Name	Field Description
`Messages sent`	Number of DHCP messages sent. • `BOOTREPLY`—Number of BOOTP PDUs transmitted • `DHCPOFFER`—Number of DHCP OFFER PDUs transmitted • `DHCPACK`—Number of DHCP ACK PDUs transmitted • `DHCPNACK`—Number of DHCP NACK PDUs transmitted • `DHCPFORCERENEW`—Number of DHCP FORCERENEW PDUs transmitted • `DHCPLEASEQUERY`—Number of DHCP leasequery messages transmitted • `DHCPLEASEBULKLEASEQUERY`—Number of DHCP bulk leasequery messages transmitted
`External Server Response`	State of the external DHCP server responsiveness.
`Packets forwarded`	Number of packets forwarded. • `BOOTREQUEST`—Number of BOOTREQUEST protocol data units (PDUs) forwarded • `BOOTREPLY`—Number of BOOTREPLY protocol data units (PDUs) forwarded
`External Server Response`	State of the external DHCP server responsiveness.
`Total Requested Servers`	Total number of servers with which the DHCP relay agent has requested a bulk leasequery connection.
`Total Attempted Servers`	Total number of servers with which the DHCP relay agent has attempted to create a bulk leasequery connection.
`Total Connected`	Total number of servers that have formed a bulk leasequery connection with the DHCP relay agent.

Table 423: show dhcp relay statistics Output Fields *(Continued)*

Field Name	Field Description
Total Terminated by Server	Total number of servers that have terminated a bulk leasequery connection with the DHCP relay agent.
Total Max Attempted	Total number of servers where the DHCP relay agent reached the maximum retry limit when it attempted to create a bulk leasequery connection.
Total Closed due to Errors	Total number of bulk leasequery connections that closed due to an internal error on the DHCP relay agent.
In-Flight Connected	Number of current bulk leasequery connections on the DHCP relay agent.
Bulk Leaasequery Reply Packet Retries	Number of bulk leasequery reply packets that the DHCP relay agent has retried.

Sample Output

show dhcp relay statistics

```
user@host> show dhcp relay statistics
Packets dropped:
    Total                       34
    Leasequery reply with older timestamp      7
    Leasequery reply aborted                   12
    Bad hardware address       1
    Bad opcode                 1
    Bad options                3
    Invalid server address     5
    Lease Time Violation       1
    No available addresses     1
    No interface match         2
    No routing instance match  9
    No valid local address     4
    Packet too short           2
    Read error                 1
```

```
    Send error                  1
    Option 60                   1
    Option 82                   2

Messages received:
    BOOTREQUEST                 116
    DHCPDECLINE                 0
    DHCPDISCOVER                11
    DHCPINFORM                  0
    DHCPRELEASE                 0
    DHCPREQUEST                 105
    DHCPLEASEACTIVE             0
    DHCPLEASEUNASSIGNED         0
    DHCPLEASEUNKNOWN            0
    DHCPLEASEQUERYDONE          0

Messages sent:
    BOOTREPLY                   0
    DHCPOFFER                   2
    DHCPACK                     1
    DHCPNAK                     0
    DHCPFORCERENEW              0
    DHCPLEASEQUERY              0
    DHCPBULKLEASEQUERY          0

Packets forwarded:
    Total                       4
    BOOTREQUEST                 2
    BOOTREPLY                   2

External Server Response:
    State                       Responding
```

show dhcp relay statistics routing-instance (ACX Series)

```
user@root> show dhcp relay statistics routing-instance tenant-a
Packets dropped:
    Total                       19
    Leasequery reply with older timestamp      7
    Leasequery reply aborted                   12
```

```
Messages received:
    BOOTREQUEST                80
    DHCPDECLINE                0
    DHCPDISCOVER               40
    DHCPINFORM                 0
    DHCPRELEASE                3
    DHCPREQUEST                37
    DHCPLEASEACTIVE            17
    DHCPLEASEUNASSIGNED        0
    DHCPLEASEUNKNOWN           101
    DHCPLEASEQUERYDONE         0
    DHCPACTIVELEASEQUERY       0

Messages sent:
    BOOTREPLY                  179
    DHCPOFFER                  17
    DHCPACK                    22
    DHCPNAK                    7
    DHCPFORCERENEW             0
    DHCPLEASEQUERY             0
    DHCPBULKLEASEQUERY         0
    DHCPLEASEACTIVE            32
    DHCPLEASEUNASSIGNED        0
    DHCPLEASEUNKNOWN           101
    DHCPLEASEQUERYDONE         0
    DHCPACTIVELEASEQUERY       0
```

show dhcp relay statistics bulk-leasequery-connections

```
user@host> show dhcp relay statistics bulk-leasequery-connections

Total Requested Servers:    0
Total Attempted Servers:    0
Total Connected:            0
Total Terminated by Server: 0
Total Max Attempted:        0
Total Closed due to Errors: 0
In-Flight Connected:        0
Bulk Leasequery Reply Packet Retries:   0
```

Release Information

Command introduced in Junos OS Release 8.3.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

clear dhcp relay statistics | 3500

show dhcp-security arp inspection statistics

IN THIS SECTION

- Syntax | 19342
- Description | 19342
- Required Privilege Level | 19342
- Output Fields | 19343
- Sample Output | 19343
- Release Information | 19343

Syntax

```
show dhcp-security arp inspection statistics
```

Description

Display Address Resolution Protocol (ARP) inspection statistics.

Required Privilege Level

view

Output Fields

Table 424 on page 19343 lists the output fields for the `show dhcp-security arp inspection statistics` command. Output fields are listed in the approximate order in which they appear.

The IP source guard database table shows the untrusted access interfaces in VLANs that have been enabled for IP source guard. The entries include the VLAN 802.1Q tag IDs if there are any, and the IP addresses and MAC addresses that are bound to one another.

Table 424: show dhcp-security arp inspection statistics Output Fields

Field Name	Field Description	Level of Output
`Interface`	Interface on which ARP inspection has been applied.	All levels
`Packets received`	Total number of packets that underwent ARP inspection.	All levels
`ARP inspection pass`	Total number of packets that passed ARP inspection.	All levels
`ARP inspection fail`	Total number of packets that failed ARP inspection.	All levels

Sample Output

show dhcp-security arp inspection statistics

```
user@device> show dhcp-security arp inspection statistics

Interface      Packets received   ARP inspection pass   ARP inspection fail
ge-0/0/30.0    7                  7                     0
ge-0/0/4.0     3                  3                     0
ge-0/0/6.0     72                 4                     68
```

Release Information

Command introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

show dhcp-security binding | 19344

clear dhcp-security binding | 3503

clear interfaces statistics | 3651

Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing

Example: Configuring IP Source Guard and Dynamic ARP Inspection on a Specified Bridge Domain to Protect the Devices Against Attacks

Port Security Features

show dhcp-security binding

IN THIS SECTION

- Syntax | 19344
- Description | 19345
- Options | 19345
- Required Privilege Level | 19345
- Output Fields | 19345
- Sample Output | 19346
- Release Information | 19348

Syntax

```
show dhcp-security binding
<detail>
<interface interface-name>
<ip-address ip-address>
<ip-source-guard ip-sg-name>
<statistics>
<vlan vlan-name>
```

Description

Display the DHCP snooping database information.

Options

detail	(Optional) Display the DHCP snooping database information along with various DHCP client options. The option information associated with a binding is displayed in Code-Length-Value triplet. The option values are displayed as a hex dump.
interface *interface-name*	(Optional) Display the DHCP snooping database information for an interface.
ip-address *ip-address*	(Optional) Display the DHCP snooping database information for an IP address.
vlan *vlan-name*	(Optional) Display the DHCP snooping database information for a VLAN.

Required Privilege Level

view

Output Fields

Table 425 on page 19345 lists the output fields for the `show dhcp-security binding` command. Output fields are listed in the approximate order in which they appear.

Table 425: show dhcp-security binding Output Fields

Field Name	Field Description	Level of Output
`IP Address`	IP address of the network device; bound to the MAC address.	All levels
`MAC address`	MAC address of the network device; bound to the IP address.	All levels
`VLAN`	VLAN name of the network device whose MAC address is shown.	All levels
`Expires`	The time, in seconds, remaining before the lease of the IP address to the MAC address expires. This field is `0` for static entries.	All levels

Table 425: show dhcp-security binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the address binding: • BOUND: IP address is leased to the MAC address for a limited period of time. • INIT: Initial state. • RELEASE: Client is releasing IP address lease. • REQUESTING: Client is requesting a server. • SELECTING: Client is receiving offers from servers. • RENEWING: Client is sending a request to renew the IP address lease. • RECONFIGURE: Client has received a reconfigure message from the server. • STATIC: IP address is attached to a fixed MAC address.	All levels
Interface	Interface address (port).	All levels

Sample Output

show dhcp-security binding

```
user@device> show dhcp-security binding

IP address        MAC address         Vlan     Expires State   Interface
10.1.1.10         00:10:00:20:00:01   vlan20   0       STATIC  ge-0/0/4.0
10.1.1.18         00:10:94:00:00:34   vlan20   86287   BOUND   ge-0/0/6.0
10.1.1.15         00:10:94:00:00:55   vlan20   86265   BOUND   ge-0/0/4.0
10.1.1.16         00:10:94:00:00:56   vlan20   86265   BOUND   ge-0/0/4.0
10.1.1.19         00:10:94:00:00:5b   vlan20   86287   BOUND   ge-0/0/6.0
10.1.1.20         00:10:94:00:00:5c   vlan20   86287   BOUND   ge-0/0/6.0
10.1.1.21         00:10:94:00:00:5d   vlan20   86287   BOUND   ge-0/0/6.0
10.1.1.17         00:10:94:00:00:68   vlan20   86265   BOUND   ge-0/0/4.0
```

show dhcp-security binding detail

```
user@device> show dhcp-security binding detail

IP address: 20.1.1.16
MAC address: 00:10:94:00:00:03
Vlan: snoop-vlan
Lease Expires in: 589
State: BOUND
Interface: ge-0/0/13.0
DHCP Client Options:
  Option: 53, Length: 1, Value: 03
  Option: 57, Length: 2, Value: 02 40
  Option: 61, Length: 7, Value: 01 00 10 94 00 00 03
```

show dhcp-security binding interface

```
user@device> show dhcp-security binding interface ge-0/0/6

IP address        MAC address         Vlan     Expires State   Interface
10.1.1.18         00:10:94:00:00:34   vlan20   86282   BOUND   ge-0/0/6.0
10.1.1.19         00:10:94:00:00:5b   vlan20   86282   BOUND   ge-0/0/6.0
10.1.1.20         00:10:94:00:00:5c   vlan20   86282   BOUND   ge-0/0/6.0
10.1.1.21         00:10:94:00:00:5d   vlan20   86282   BOUND   ge-0/0/6.0
```

show dhcp-security binding ip-address

```
user@device> show dhcp-security binding ip-address

IP address     MAC address              Vlan         Expires     State      Interface
10.1.1.18      00:10:94:00:00:34     vlan20       86282       BOUND      ge-0/0/6.0
```

show dhcp-security binding vlan

```
user@device> show dhcp-security binding vlan vlan20
```

```
IIP address     MAC address              Vlan         Expires     State      Interface
10.1.1.18       00:10:94:00:00:34    vlan20       86282       BOUND      ge-0/0/6.0
```

Release Information

Command introduced in Junos OS Release 12.1.

`detail` option introduced in Junos OS Release 23.1.

RELATED DOCUMENTATION

- show dhcp-security binding ip-source-guard | **19348**
- clear dhcp-security binding | **3503**
- *Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing*
- *Example: Configuring IP Source Guard and Dynamic ARP Inspection on a Specified Bridge Domain to Protect the Devices Against Attacks*
- *Port Security Features*

show dhcp-security binding ip-source-guard

IN THIS SECTION

- Syntax | **19349**
- Description | **19349**
- Required Privilege Level | **19349**
- Output Fields | **19349**
- Sample Output | **19350**
- Release Information | **19350**

Syntax

```
show dhcp-security binding ip-source-guard
```

Description

Display IP source guard database table.

Required Privilege Level

view

Output Fields

Table 426 on page 19349 lists the output fields for the `show dhcp-security binding ip-source-guard` command. Output fields are listed in the approximate order in which they appear.

The IP source guard database table shows the untrusted access interfaces in VLANs that have been enabled for IP source guard. The entries include the IP addresses and MAC addresses that are bound to one another.

Table 426: show dhcp-security binding ip-source-guard Output Fields

Field Name	Field Description	Level of Output
`IP Address`	IP address of the network device; bound to the MAC address.	All levels
`MAC address`	MAC address of the network device; bound to the IP address.	All levels
`VLAN`	VLAN name of the network device whose MAC address is shown.	All levels
`Expires`	The time, in seconds, remaining before the lease of the IP address to the MAC address expires.	All levels

Table 426: show dhcp-security binding ip-source-guard Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Specifies whether the IP address is: • BOUND: Temporarily leased to the MAC address for a limited period of time. • STATIC: Attached to a fixed MAC address.	All levels
Interface	Interface address (port).	All levels

Sample Output

show dhcp-security binding ip-source-guard

```
user@device> show dhcp-security binding ip-source-guard

IP address        MAC address         Vlan     Expires State   Interface
10.1.1.10         00:10:00:20:00:01   vlan20   0       STATIC  ge-0/0/4.0
10.1.1.18         00:10:94:00:00:34   vlan20   86276   BOUND   ge-0/0/6.0
10.1.1.15         00:10:94:00:00:55   vlan20   86254   BOUND   ge-0/0/4.0
10.1.1.16         00:10:94:00:00:56   vlan20   86254   BOUND   ge-0/0/4.0
10.1.1.19         00:10:94:00:00:5b   vlan20   86276   BOUND   ge-0/0/6.0
10.1.1.20         00:10:94:00:00:5c   vlan20   86276   BOUND   ge-0/0/6.0
10.1.1.21         00:10:94:00:00:5d   vlan20   86276   BOUND   ge-0/0/6.0
10.1.1.17         00:10:94:00:00:68   vlan20   86254   BOUND   ge-0/0/4.0
```

Release Information

Command introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

show dhcp-security binding | 19344

clear dhcp-security binding | 3503

Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing

Example: Configuring IP Source Guard and Dynamic ARP Inspection on a Specified Bridge Domain to Protect the Devices Against Attacks

Port Security Features

show dhcp-security ipv6 binding

IN THIS SECTION

- Syntax | 19351
- Description | 19351
- Options | 19352
- Required Privilege Level | 19352
- Output Fields | 19352
- Sample Output | 19353
- Release Information | 19355

Syntax

```
show dhcp-security ipv6 binding
<detail>
<interface interface-name>
<ipv6-address ipv6-address>
<vlan vlan-name>
```

Description

Display bindings between IPv6 addresses and MAC addresses (IP-MAC bindings) along with other DHCP lease information, also known as the DHCPv6 binding table or DHCPv6 snooping database.

Options

detail (Optional) Display the DHCPv6 snooping table along with various DHCP client options. The option information associated with a binding is displayed in Code-Length-Value triplet. The option values are displayed as a hex dump.

> **NOTE**: DHCP client options are displayed for DHCPv6 only when the snoop entry type is IA_NA.

interface *interface-name* (Optional) Display the DHCPv6 snooping table for the specified interface.

ipv6-address *ipv6-address* (Optional) Display the DHCPv6 snooping table for the specified IPv6 address.

vlan *vlan-name* (Optional) Display the DHCPv6 snooping table for a VLAN.

Required Privilege Level

view

Output Fields

Table 427 on page 19352 lists the output fields for the `show dhcp-security ipv6 binding` command. Output fields are listed in the approximate order in which they appear.

The DHCPv6 binding table shows the untrusted access interfaces in VLANs that have been enabled for DHCPv6 snooping. The entries include the IPv6 addresses and MAC addresses that are bound to one another.

Table 427: show dhcp-security ipv6 binding Output Fields

Field Name	Field Description	Level of Output
`IPv6 address`	IPv6 addresses of the network device; bound to the MAC address. There are two entries added for each client: one with the link-local IPv6 address, which is used by the client for DHCP transactions, and another with the IPv6 address assigned by the server. The link-local address always has the prefix fe80::/10.	All levels
`MAC address`	MAC address of the network device; bound to the IPv6 address.	All levels

Table 427: show dhcp-security ipv6 binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
VLAN	VLAN name of the network device whose MAC address is shown.	All levels
Expires	The time, in seconds, remaining before the lease of the IPv6 address to the MAC address expires. This field is 0 for static entries.	All levels
State	Specifies whether the IPv6 address is: • BOUND: Temporarily leased to the MAC address for a limited period of time. • STATIC: Attached to a fixed MAC address.	All levels
Interface	Interface address (port).	All levels
Type	Shows the snoop entry type, whenever one is present. Snoop entry types can be one of the following: • IA_NA: The snoop entry related to the IA Address • IA_PD: The snoop entry related to the IA Prefix • Link-local: The snoop entry corresponding to the source address of the requesting client. • Static: The snoop entry created through configuration by a user or admin.	All levels

Sample Output

show dhcp-security ipv6 binding

```
user@switch> show dhcp-security ipv6 binding
IPv6 address               MAC address         Vlan     Expires State   Interface
Type
fe80::210:94ff:fe00:7      00:10:94:00:00:07   vlan-iapd 86394  BOUND   ge-0/0/13.0
Link local
```

```
2002::e                    00:10:94:00:00:07   vlan-iapd 86394  BOUND   ge-0/0/13.0           IA-
NA
2002::c                    00:10:94:00:00:02   vlan-iapd 0      STATIC  ge-0/0/13.0
Static
2000:1::5:0/112            00:10:94:00:00:07   vlan-iapd 86394  BOUND   ge-0/0/13.0           IA-
PD
```

show dhcp-security ipv6 binding detail

```
user@device> show dhcp-security ipv6 binding detail

IPv6 Address: fe80::210:94ff:fe00:7
MAC address: 00:10:94:00:00:02
Vlan: snoop-vlan
Lease Expires in: 91
State: BOUND
Interface:  ge-0/0/13.0
Type: Link local

IPv6 Address: 2001:db8:fe16::/112
MAC address: 00:10:94:00:00:02
Vlan: snoop-vlan
Lease Expires in: 91
State: BOUND
Interface:  ge-0/0/13.0
Type: IA-PD

IPv6 Address: 2001:db8::10:10:2
MAC address: 00:10:94:00:00:02
Vlan: snoop-vlan
Lease Expires in: 91
State: BOUND
Interface:  ge-0/0/13.0
Type: IA-NA
DHCP Client Options:
  Option: 1, Length: 14, Value: 00 01 00 01 64 35 42 89 00 10 94 00 00 02
  Option: 2, Length: 26, Value: 00 02 00 00 05 83 34 34 3a 61 61 3a 35 30
  3a 63 32 3a 63 33 3a 32 30 00 00 00
  Option: 3, Length: 40, Value: 00 00 00 00 00 00 00 3c 00 00 00 60 00 05
  00 18 20 00 00 00 00 00 00 00 00 10 00 10 00 00 00 02 00 00 00 78 00 00 00 78
```

show dhcp-security ipv6 binding interface

```
user@switch> show dhcp-security ipv6 binding interface ge-0/0/4.0
IPv6 address             MAC address         Vlan     Expires State   Interface
2001:db8:fe16::          00:10:00:20:00:01   vlan20   0       STATIC  ge-0/0/4.0
fe80::210:94ff:fe00:4    00:10:00:20:00:01   vlan20   0       STATIC  ge-0/0/4.0
```

Release Information

Command introduced in Junos OS Release 13.2X51-D20.

`detail` option introduced in Junos OS Release 23.1.

RELATED DOCUMENTATION

show dhcp-security ipv6 statistics | 19355

clear dhcp-security ipv6 binding | 3505

Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing

show dhcp-security ipv6 statistics

IN THIS SECTION

- Syntax | 19356
- Description | 19356
- Required Privilege Level | 19356
- Output Fields | 19356
- Sample Output | 19358
- Release Information | 19359

Syntax

```
show dhcp-security ipv6 statistics
```

Description

Display DHCPv6 statistics.

Required Privilege Level

view

Output Fields

Table 428 on page 19357 lists the output fields for the `show dhcp-security ipv6 statistics` command. Output fields are listed in the approximate order in which they appear.

Table 428: show dhcp-security ipv6 statistics Output Fields

Field Name	Field Description
`DHCPv6 messages`	Number of DHCPv6 messages exchanged. • `Total`—Total number of DHCPv6 messages exchanged. • `Solicit`—Number of DHCPv6 messages of type Solicit. A client sends a Solicit message to locate servers. • `Advertise`—Number of DHCPv6 messages of type Advertise. A server sends an Advertise message, in reponse to a Solicit message, to indicate that it is available for DHCPv6 service. • `Request`—Number of DHCPv6 messages of type Request. A client sends a Request message to request configuration parameters from a server. • `Reply`—Number of DHCPv6 messages of type Reply. A server sends a Reply message in response to a Solicit, Request, Renew, Rebind, Confirm, Information Request, Release, or Decline message. • `Confirm`—Number of DHCPv6 messages of type Confirm. A client sends a Confirm message to any available server to determine whether the addresses it was assigned are still appropriate for the link to which the client is connected. • `Decline`—Number of DHCPv6 messages of type Decline. A client sends a Decline message to a server to indicate that one or more of the addresses assigned by the server are already in use on the link to which the client is connected. • `Release`—Number of DHCPv6 messages of type Release. A client sends a Release message to the server to indicate that the client will no longer use one or more of the assigned addresses. • `Renew`—Number of DHCPv6 messages of type Renew. A client sends a Renew message to the server to extend the lifetimes on the addresses assigned to the client by that server and to update other configuration parameters received by that server. • `Rebind`—Number of DHCPv6 messages of type Rebind. A client sends a Rebind message to any available server after receiving no reply to a Renew message. • `Relay-forward`—Number of DHCPv6 messages of type Relay-forward. A relay agent sends a Relay-forward message to relay messages to servers, either directly or through another relay agent. The received message is encapsulated in an option in the Relay-forward message.

Table 428: show dhcp-security ipv6 statistics Output Fields *(Continued)*

Field Name	Field Description
	• `Relay-reply`—Number of DHCPv6 messages of type Relay-reply. A server sends a Replay-reply message to a relay agent containing a message that the relay agent delivers to a client. • `Information-request`—Number of DHCPv6 messages of type Information-request. A client sends an Information-request message to a server to request configuration parameters without the assignment of any IP addresses to the client. • `Reconfigure`—Number of DHCPv6 messages of type Reconfigure. A server sends a Reconfigure message to a client to inform the client that the server has new or updated configuration parameters, and that the client needs to initiate a Renew/Reply or Information-request/Reply transaction with the server in order to receive the updated information.
`Packets dropped`	Number of packets not considered for DHCPv6 snooping because of errors. • `Total`—Total number of packets discarded by DHCPv6 snooping. • `No configuration`—Number of packets discarded because they did not have a valid configuration. • `No VLAN`—Number of packets discarded because they did not belong to a valid VLAN. • `No interface`—Number of packets discarded because they did not belong to a valid interface. • `Request on trusted port`—Number of packets discarded because a Request message was received on a trusted port.

Sample Output

show dhcp-security ipv6 statistics

```
user@host> show dhcp-security ipv6 statistics
DHCPv6 messages:
    Total                      32
    Solicit                     1
    Advertise                   1
    Request                     3
```

```
    Reply                       5
    Confirm                     1
    Decline                     2
    Release                     9
    Renew                       4
    Rebind                      2
    Relay forward               1
    Relay reply                 1
    Information request         1
    Reconfigure                 2

Packets dropped:
    Total                       0
    No configuration            0
    No VLAN                     0
    No interface                0
    Request on trusted port     0
```

Release Information

Command introduced in Junos OS Release 13.2X51-D20.

RELATED DOCUMENTATION

show dhcp-security ipv6 binding | 19351

show dhcp-security neighbor-discovery-inspection statistics | 19359

show dhcp-security neighbor-discovery-inspection statistics

IN THIS SECTION

- Syntax | 19360
- Description | 19360
- Options | 19360

- Required Privilege Level | 19360
- Output Fields | 19360
- Sample Output | 19361
- Sample Output | 19361
- Release Information | 19362

Syntax

```
show dhcp-security neighbor-discovery-inspection statistics
<interface interface-name>
```

Description

Display IPv6 neighbor discovery inspection statistics to determine whether there is IPv6 address spoofing on the network.

Options

interface *interface-name* (Optional) Display neighbor discovery inspection statistics for the specified interface.

Required Privilege Level

view

Output Fields

Table 429 on page 19361 lists the output fields for the `show dhcp-security neighbor-discovery-inspection statistics` command. Output fields are listed in the approximate order in which they appear.

Table 429: show dhcp-security neighbor-discovery-inspection statistics Output Fields

Field Name	Field Description	Level of Output
`Interface`	Interface on which neighbor discovery inspection has been applied.	All levels
`Packets received`	Total number of packets that underwent neighbor discovery inspection.	All levels
`ND inspection pass`	Total number of packets that passed neighbor discovery inspection.	All levels
`ND inspection fail`	Total number of packets that failed neighbor discovery inspection.	All levels

Sample Output

show dhcp-security neighbor-discovery-inspection statistics

```
user@switch> show dhcp-security neighbor-discovery-inspection statistics

Interface     ND Packets received   ND inspection pass    ND inspection failed
ge-0/0/1.0                      7                    5                       2
ge-0/0/2.0                     10                   10                       0
ge-0/0/3.0                     12                   12                       0
```

Sample Output

show dhcp-security neighbor-discovery-inspection statistics interface

```
user@switch> show dhcp-security neighbor-discovery-inspection statistics interface
ge-0/0/1.0

Interface     ND Packets received   ND inspection pass    ND inspection failed
ge-0/0/1.0                      7                    5                       2
```

Release Information

Command introduced in Junos OS Release 13.2X51-D20.

RELATED DOCUMENTATION

show dhcp-security ipv6 binding | 19351

IPv6 Neighbor Discovery Inspection

Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing

show dhcp server active-leasequery

IN THIS SECTION

- Syntax | 19362
- Description | 19362
- Required Privilege Level | 19363
- Output Fields | 19363
- Sample Output | 19363
- Release Information | 19363

Syntax

```
show dhcp server active-leasequery summary
```

Description

Display the active leasequery status summary of the DHCP local server.

Required Privilege Level

view

Output Fields

Table 430 on page 19363 lists the output fields for the `show dhcp server active-leasequery summary` command.

Table 430: show dhcp server active-leasequery summary Output Fields

Field Name	Field Description
`Connected Peers`	Number of peer DHCP servers connected.
`Connecting Peers`	Number of DHCP server peers trying to connect or reconnect after failure.
`peer`	IP address of the peer DHCP server.

Sample Output

show dhcp server active-leasequery summary

```
user@host> show dhcp server active-leasequery summary
```

```
Active Lease-query peer List:

peer         Connected Peers   Connecting Peers
192.0.2.0    2                 0
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

M:N Subscriber Service Redundancy on DHCP Server

show dhcp server active-leasequery statistics

IN THIS SECTION

- Syntax | 19364
- Description | 19364
- Options | 19364
- Required Privilege Level | 19365
- Output Fields | 19365
- Sample Output | 19366
- Release Information | 19366

Syntax

```
show dhcp server active-leasequery statistics
<peer peer-address>
<interface interface-name>
```

Description

Display the active leasequery statistics of the DHCP local server.

Options

peer *peer-address*	IP address of the DHCP server on which you want to view the active leasequery statistics.
interface *interface-name*	Interface name of the DHCP server on which you want to view the active leasequery statistics.

Required Privilege Level

view

Output Fields

Table 431 on page 19365 lists the output fields for the `show dhcp server active-leasequery statistics` command.

Table 431: show dhcp server active-leasequery statistics Output Fields

Field Name	Field Description
`ALQ Transmit Buffer count`	Numbers of packets the active leasequery transmitted.
`Bindings Failed to install`	Number of binding installation attempt failed.
`Bindings Installed Successfully`	Number of successful binding installation.
`Bindings Received`	Number of binding request received.
`Bindings Sent`	Number of binding request sent.
`Last Synchronization Time`	Last synchronization time.
`Local Interface count`	Number of local interfaces.
`Max Leasequery Transmit Rate`	Maximum leasequery transmit rate.
`peer`	IP address of the peer DHCP server.
`Remote Interface count`	Number of remote interfaces.
`State`	Status of the topology discovery configuration.

Table 431: show dhcp server active-leasequery statistics Output Fields *(Continued)*

Field Name	Field Description
Topology-Discover Configured	Status of the topology discovery.

Sample Output

show dhcp server active-leasequery peer 192.0.2.0

```
user@host> show dhcp server active-leasequery peer 192.0.2.0
```

```
    peer : 192.0.2.0
    Topology-Discover Configured          : Yes
    State                                 : Done
    Bindings Sent                         : 0
    Bindings Received                     : 6
    Bindings Installed Successfully       : 0
    Bindings Failed to install            : 0
    Last Synchronization Time             : 2021-04-18 19:02:16 IST
    ALQ Transmit Buffer count             : 0x ffff
    Max Leasequery Transmit Rate          : 60
    Local Interface count                 : 1
    Remote Interface count                : 1
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

clear dhcp server active-leasequery statistics | 3506

show dhcp server binding

IN THIS SECTION

- Syntax | 19367
- Description | 19367
- Options | 19368
- Required Privilege Level | 19368
- Output Fields | 19368
- Sample Output | 19373
- Release Information | 19377

Syntax

```
show dhcp server binding
<address>
<interfaces-vlan><brief | detail | summary>
<interface interface-name>
<interface-tag interface-tag-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol (DHCP) local server.

NOTE: If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

Options

address	(Optional) Display DHCP binding information for a specific client identified by one of the following entries: • *ip-address*—The specified IP address. • *mac-address*—The specified MAC address. • *session-id*—The specified session ID.
brief | detail | summary	(Optional) Display the specified level of output about active client bindings. The default is `brief`, which produces the same output as `show dhcp server binding`.
interface *interface-name*	(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.
interface-tag *interface-tag-name*	(Optional) Displays the binding information for the specified interface-tag.
interfaces-vlan	(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.
interfaces-wildcard	(Optional) The set of interfaces on which to show the binding state information. This option supports the use of the wildcard character (*).
logical-system *logical-system-name*	(Optional) Display information about active client bindings for DHCP clients on the specified logical system.
routing-instance *routing-instance-name*	(Optional) Display information about active client bindings for DHCP clients on the specified routing instance.

Required Privilege Level

view

Output Fields

Table 432 on page 19369 lists the output fields for the `show dhcp server binding` command. Output fields are listed in the approximate order in which they appear.

Table 432: show dhcp server binding Output Fields

Field Name	Field Description	Level of Output
number clients, (*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary
IP address	IP address of the DHCP client.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Hardware address	Hardware address of the DHCP client.	brief detail
Expires	Number of seconds in which lease expires.	brief detail
State	State of the address binding table on the extended DHCP local server: • BOUND—Client has active IP address lease. • FORCERENEW—Client has received forcerenew message from server. • INIT—Initial state. • RELEASE—Client is releasing IP address lease. • RENEWING—Client sending request to renew IP address lease. • REQUESTING—Client requesting a DHCP server. • SELECTING—Client receiving offers from DHCP servers.	brief detail

Table 432: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface`	Interface on which the request was received.	`brief`
`Lease Expires`	Date and time at which the client's IP address lease expires.	`detail`
`Lease Expires in`	Number of seconds in which lease expires.	`detail`
`Lease Start`	Date and time at which the client's IP address lease started.	`detail`
`Lease time violated`	Lease time violation has occurred.	`detail`
`Last Packet Received`	Date and time at which the router received the last packet.	`detail`
`Incoming Client Interface`	Client's incoming interface.	`detail`
`Interface tag`	Name of the Interface tag.	`detail`
`Client Interface Svlan Id`	S-VLAN ID of the client's incoming interface.	`detail`
`Client Interface Vlan Id`	VLAN ID of the client's incoming interface.	`detail`
`Demux Interface`	Name of the IP demultiplexing (demux) interface.	`detail`
`Server IP Address` or `Server Identifier`	IP address of DHCP server.	`detail`
`Server Interface`	Interface of DHCP server.	`detail`

Table 432: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Client Pool Name`	Name of address pool used to assign client IP address lease.	`detail`

Table 432: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Liveness Detection State`	State of the liveness detection status for a subscriber's Bidirectional Forwarding Detection (BFD) protocol session: **NOTE**: This output field displays status only when liveness detection has been explicitly configured for a subscriber and the liveness detection protocol is actively functioning for that subscriber. • `DOWN`—Liveness detection has been enabled for a subscriber but the broadband network gateway (BNG) detects that the liveness detection session for the BFD protocol is in the `DOWN` state. A liveness detection session that was previously in an `UP` state has transitioned to a `DOWN` state, beginning with a liveness detection failure, and ending with the deletion of the client binding. The `DOWN` state is reported only during this transition period of time. • `UNKNOWN`—Liveness detection has been enabled for a subscriber but the actual liveness detection state has not yet been determined. The `UNKNOWN` state is reported after a DHCP subscriber initially logs in while the underlying liveness detection protocol handshake, such as BFD, is still processing and the BFD session has not yet reached the `UP` state. • `UP`—Liveness detection has been enabled for a subscriber, and the BNG and the subscriber or client have *both* determined that the liveness detection session for the BFD protocol is in the `UP` state. • `WENT_DOWN`—State is functionally equivalent to the `DOWN` state. A liveness detection session that was previously in an `UP` state has transitioned to a `DOWN` state implying a liveness detection failure.	`detail`

Table 432: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	The WENT_DOWN state applies to the internal distribution of the liveness detection mechanism between the Junos DHCP Daemon for Subscriber Services (JDHCPd), the BFD plug-in within the Broadband Edge Subscriber Management Daemon (BBE-SMGD), and the Packet Forwarding Engine.	
ACI Interface Set Name	Internally generated name of the dynamic agent circuit identifier (ACI) interface set.	detail
ACI Interface Set Index	Index number of the dynamic ACI interface set.	detail
ACI Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.	detail
Client Profile Name	DHCP client profile name.	detail
Dual Stack Group	DHCP server profile name.	detail
Dual Stack Peer Prefix	IPv6 prefix of peer.	detail
Dual Stack Peer Address	IPv6 address of peer.	detail

Sample Output

show dhcp server binding

```
user@host> show dhcp server binding
IP address        Session Id  Hardware address   Expires     State      Interface
192.51.100.15     6           00:00:5e:00:53:01  86180       BOUND      ge-1/0/0.0
192.51.100.16     7           00:00:5e:00:53:02  86180       BOUND      ge-1/0/0.0
192.51.100.17     8           00:00:5e:00:53:03  86180       BOUND      ge-1/0/0.0
```

```
192.51.100.18     9           00:00:5e:00:53:04  86180       BOUND      ge-1/0/0.0
192.51.100.19     10          00:00:5e:00:53:05  86180       BOUND      ge-1/0/0.0
```

show dhcp server binding detail

```
user@host> show dhcp server binding detail
Client IP Address:  192.51.100.15
     Hardware Address:             00:00:5e:00:53:01
     State:                        BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 10:10:25 PDT
     Lease Expires in:             86151 seconds
     Lease Start:                  2009-07-20 10:10:25 PDT
     Incoming Client Interface:    ge-1/0/0.0
     Interface tag:                None
     Server Ip Address:            192.51.100.9
     Server Interface:             none
     Session Id:                   6
     Client Pool Name:             6
     Liveness Detection State:     UP
 Client IP Address:            192.51.100.16
     Hardware Address:             00:00:5e:00:53:02
     State:                        BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 10:10:25 PDT
     Lease Expires in:             86151 seconds
     Lease Start:                  2009-07-20 10:10:25 PDT
     Lease time violated:          yes
     Incoming Client Interface:    ge-1/0/0.0
     Server Ip Address:            192.51.100.9
     Server Interface:             none
     Session Id:                   7
     Client Pool Name:             7
     Liveness Detection State:     UP
```

When DHCP binding is configured with dual-stack, we get the following output:

```
user@host> show dhcp server binding detail
Client IP Address:  192.51.100.10
     Hardware Address:             00:00:64:03:01:02
     State:                        BOUND(LOCAL_SERVER_STATE_BOUND)
     Protocol-Used:                DHCP
```

```
     Lease Expires:                2016-11-07 08:30:39 PST
     Lease Expires in:             43706 seconds
     Lease Start:                  2016-11-04 11:00:37 PDT
     Last Packet Received:         2016-11-06 09:00:39 PST
     Incoming Client Interface:    ae0.3221225472
     Interface tag:                None
     Client Interface Svlan Id:    2000
     Client Interface Vlan Id:     1
     Server Ip Address:            192.51.100.2
     Session Id:                   2
     Client Pool Name:             my-v4-pool
     Client Profile Name:          dhcp-retail
     Dual Stack Group:             my-dual-stack
     Dual Stack Peer Prefix:       2001:db8:ffff:0:4::/64
     Dual Stack Peer Address:      2001:db8:0:8003::1/128
```

show dhcp server binding detail (ACI Interface Set Configured)

```
user@host> show dhcp server binding detail
Client IP Address:  192.51.100.14
     Hardware Address:             00:00:5e:00:53:02
     State:                        BOUND(LOCAL_SERVER_STATE_BOUND)
     Lease Expires:                2012-03-13 09:53:32 PDT
     Lease Expires in:             82660 seconds
     Lease Start:                  2012-03-12 10:23:32 PDT
     Last Packet Received:         2012-03-12 10:23:32 PDT
     Incoming Client Interface:    demux0.1073741827
     Interface tag:                None
     Client Interface Svlan Id:    1802
     Client Interface Vlan Id:     302
     Demux Interface:              demux0.1073741832
     Server Identifier:            192.51.100.202
     Session Id:                   11
     Client Pool Name:             poolA
     Client Profile Name:          DEMUXprofile
     Liveness Detection State:     UP
     ACI Interface Set Name:       aci-1002-demux0.1073741827
     ACI Interface Set Index:      2
     ACI Interface Set Session ID: 6
```

show dhcp server binding interface <vlan-id>

```
user@host> show dhcp server binding interface ge-1/1/0:100
IP address        Session Id  Hardware address   Expires     State      Interface
192.51.100.15      6           00:00:5e:00:53:01  86124       BOUND      ge-1/1/0:100
```

show dhcp server binding interface <svlan-id>

```
user@host> show dhcp server binding interface ge-1/1/0:10-100
IP address        Session Id  Hardware address   Expires     State      Interface
192.51.100.16      7          00:00:5e:00:53:02  86124       BOUND      ge-1/1/0:10-100
```

show dhcp server binding <ip-address>

```
user@host> show dhcp server binding 192100.19
IP address        Session Id  Hardware address   Expires     State      Interface
192.51.100.19      10          00:00:5e:00:53:05  86081       BOUND      ge-1/0/0.0
```

show dhcp server binding <session-id>

```
user@host> show dhcp server binding 6
IP address        Session Id  Hardware address   Expires     State      Interface
192.51.100.15      6           00:00:5e:00:53:01  86124       BOUND      ge-1/0/0.0
```

show dhcp server binding summary

```
user@host> show dhcp server binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

show dhcp server binding <interfaces-vlan>

```
user@host> show dhcp server binding ge-1/0/0:100-200
IP address        Session Id  Hardware address   Expires     State      Interface
```

```
192.168.0.17      42          00:00:5e:00:53:02  86346       BOUND      ge-1/0/0.1073741827
192.168.0.16      41          00:00:5e:00:53:01  86346       BOUND      ge-1/0/0.1073741827
```

show dhcp server binding <interfaces-wildcard>

```
user@host> show dhcp server binding ge-1/3/*
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.9       24          00:00:5e:00:53:04  86361       BOUND      ge-1/3/0.110
192.168.0.8       23          00:00:5e:00:53:03  86361       BOUND      ge-1/3/0.110
192.168.0.7       22          00:00:5e:00:53:02  86361       BOUND      ge-1/3/0.110
```

Release Information

Command introduced in Junos OS Release 9.0.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

clear dhcp server binding | 3508

show dhcp server lockout-entries

IN THIS SECTION

- Syntax | 19378
- Description | 19378
- Options | 19378
- Required Privilege Level | 19378
- Output Fields | 19378
- Sample Output | 19380

Release Information | 19380

Syntax

```
show dhcp server lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv4 local lockout database.

Options

all Display all client entries in the lockout database.

index *index* Display detailed information for the specified client.

Required Privilege Level

view

Output Fields

Table 433 on page 19378 lists the output fields for the `show dhcp server lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 433: show dhcp server lockout-entries Output Fields

Field Name	Field Description	Level of Output
`Index`	Number identifying a specific entry in the lockout database.	`all` and `index`

Table 433: show dhcp server lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Key`	Client identifier for the client in the lockout database.	`all` and `index`
`State`	Type of lockout period for the entry: • `Grace`—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. • `Lockout`—Client is currently locked out; attempts to establish a session are rejected.	`all` and `index`
`Expires (s)`	Number of seconds until the current lockout period expires.	`all` only
`Elapsed (s)`	Number of seconds since the current lockout or grace timer started.	`all` only
`Count`	Number of consecutive times the client has been locked out.	`all` only
`Expires`	Date and time when the current lockout period ends.	`index` only
`Expires in`	Number of seconds until the current period expires.	`index` only
`Lockout count`	Number of consecutive times client has been locked out.	`index` only
`Next lockout time`	Duration of the next lockout period for this client.	`index` only
`Min lockout time`	Minimum duration for a lockout period; the initial lockout time.	`index` only

Table 433: show dhcp server lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Lockout reason`	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	`index` only

Sample Output

show dhcp server lockout-entries (All Entries)

```
user@host> show dhcp server lockout-entries all
Index     Key               State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

show dhcp server lockout-entries (Specific Entry)

```
user@host> show dhcp server lockout-entries index 2
     Index:                                  2
     Key:                                    default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                                  Lockout
     Expires:                                2018-03-29 19:06:17 IST
     Expires in:                             87
     Lockout count:                          1
     Next lockout time:                      200
     Min lockout time:                       100
     Lockout reason:                         181
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

clear dhcp server lockout-entries | 3512
show dhcp relay lockout-entries | 19328
show dhcpv6 relay lockout-entries | 19417
show dhcpv6 server lockout-entries | 19446

show dhcp server statistics

IN THIS SECTION

- Syntax | 19381
- Description | 19381
- Options | 19382
- Required Privilege Level | 19382
- Output Fields | 19382
- Sample Output | 19386
- Release Information | 19387

Syntax

```
show dhcp server statistics
<bulk-leasequery-connections>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options

bulk-leasequery-connections	(Optional) Display information about DHCP local server bulk leasequery statistics.
logical-system ***logical-system-name***	(Optional) Display information about extended DHCP local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.
routing-instance ***routing-instance-name***	(Optional) Display information about extended DHCP local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 434 on page 19383 lists the output fields for the `show dhcp server statistics` command. Output fields are listed in the approximate order in which they appear.

Table 434: show dhcp server statistics Output Fields

Field Name	Field Description
`Packets dropped`	Number of packets discarded by the extended DHCP local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears. • `Total`—Total number of packets discarded by the extended DHCP local server • `Authentication`—Number of packets discarded because they could not be authenticated • `Bad hardware address`—Number of packets discarded because an invalid hardware address was specified • `Bad opcode`—Number of packets discarded because an invalid operation code was specified • `Bad options`—Number of packets discarded because invalid options were specified • `Dynamic profile`—Number of packets discarded due to dynamic profile information • `Invalid server address`—Number of packets discarded because an invalid server address was specified • `Lease Time Violation`—Number of packets discarded because of a lease time violation • `No available addresses`—Number of packets discarded because there were no addresses available for assignment • `No interface match`—Number of packets discarded because they did not belong to a configured interface • `No routing instance match`—Number of packets discarded because they did not belong to a configured routing instance • `No valid local address`—Number of packets discarded because there was no valid local address • `Packet too short`—Number of packets discarded because they were too short • `Read error`—Number of packets discarded because of a system read error • `Send error`—Number of packets that the extended DHCP local server could not send

Table 434: show dhcp server statistics Output Fields *(Continued)*

Field Name	Field Description
`Offer Delay`	Number of DHCPv4 offer messages delayed. • `DELAYED`—Number of DHCPv4 offer packets that have been sent after being delayed. • `INPROGRESS`—Number of DHCPv4 offer packets that are in the delay queue. • `TOTAL`—Total number of delayed DHCPv4 offer messages; sum of `DELAYED` and `INPROGRESS`.
`Messages received`	Number of DHCP messages received. • `BOOTREQUEST`—Number of BOOTP protocol data units (PDUs) received • `DHCPDECLINE`—Number of DHCP PDUs of type DECLINE received • `DHCPDISCOVER`—Number of DHCP PDUs of type DISCOVER received • `DHCPINFORM`—Number of DHCP PDUs of type INFORM received • `DHCPRELEASE`—Number of DHCP PDUs of type RELEASE received • `DHCPREQUEST`—Number of DHCP PDUs of type REQUEST received • `DHCPLEASEQUERY`—Number of DHCP leasequery messages received. • `DHCPBULKLEASEQUERY`—Number of DHCP bulk leasequery messages received. • `DHCPRENEW`—Number of DHCP renew messages received; subset of `DHCPREQUEST` counter. • `DHCPREBIND`—Number of DHCP rebind messages received; subset of `DHCPREQUEST` counter.

Table 434: show dhcp server statistics Output Fields *(Continued)*

Field Name	Field Description
`Messages sent`	Number of DHCP messages sent. • `BOOTREPLY`—Number of BOOTP PDUs transmitted • `DHCPOFFER`—Number of DHCP OFFER PDUs transmitted • `DHCPACK`—Number of DHCP ACK PDUs transmitted • `DHCPNACK`—Number of DHCP NACK PDUs transmitted • `DHCPFORCERENEW`—Number of DHCP FORCERENEW PDUs transmitted • `DHCPLEASEUNASSIGNED`—Number of DHCP leases that are managed by the server but have not yet been assigned • `DHCPLEASEUNKNOWN`—Number of unknown DHCP leases • `DHCPLEASEACTIVE`—Number of active DHCP leases • `DHCPLEASEQUERYDONE`—The leasequery is complete
`Total Accepted Connections`	Total number of bulk leasequery connections accepted by the server.
`Total Not-Accepted Connections`	Total number of bulk leasequery connections not accepted by the server.
`Connections Closed due to Errors`	Number of bulk leasequery connections that the server closed due to an internal error.
`Connections Closed due to max-empty-replies`	Number of bulk leasequery connections that the server closed because the maximum number of empty replies was reached.
`In-flight Connections`	Number of bulk leasequery connections on the server.

Sample Output

show dhcp server statistics

```
user@host> show dhcp server statistics
Packets dropped:
    Total                      1
    Lease Time Violation       1

Offer Delay:
    DELAYED                    3
    INPROGRESS                 9
    TOTAL                     12

Messages received:
    BOOTREQUEST                25
    DHCPDECLINE                0
    DHCPDISCOVER               10
    DHCPINFORM                 0
    DHCPRELEASE                4
    DHCPREQUEST                10
    DHCPRENEW                  4
    DHCPREBIND                 2

Messages sent:
    BOOTREPLY                  20
    DHCPOFFER                  10
    DHCPACK                    10
    DHCPNAK                    0
    DHCPFORCERENEW             0
```

show dhcp server statistics

```
user@host> show dhcp server statistics verbose
Packets dropped:
    Total                      0

Messages received:
    BOOTREQUEST                238
    DHCPDECLINE                0
```

```
    DHCPDISCOVER                1
    DHCPINFORM                  0
    DHCPRELEASE                 0
    DHCPREQUEST                 237
    DHCPRENEW                   236
    DHCPREBIND                  0

Messages sent:
    BOOTREPLY                   20
    DHCPOFFER                   10
    DHCPACK                     10
    DHCPNAK                     0
    DHCPFORCERENEW              0
```

Release Information

Command introduced in Junos OS Release 9.0.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

clear dhcp server statistics | 3515

Table 417 | 19306

show dhcp snooping binding

IN THIS SECTION

- Syntax | 19388
- Description | 19388
- Options | 19388
- Required Privilege Level | 19388
- Output Fields | 19388
- Sample Output | 19389

Release Information | 19389

Syntax

```
show dhcp snooping binding
<interface interface-name>
<vlan vlan-name>
```

Description

Display the DHCP snooping database information.

Options

interface *interface-name* (Optional) Display the DHCP snooping database information for an interface.

vlan *vlan-name* (Optional) Display the DHCP snooping database information for a VLAN.

Required Privilege Level

view

Output Fields

Table 435 on page 19388 lists the output fields for the `show dhcp snooping binding` command. Output fields are listed in the approximate order in which they appear.

Table 435: show dhcp snooping binding Output Fields

Field Name	Field Description	Level of Output
MAC Address	MAC address of the network device; bound to the IP address.	All levels
IP Address	IP address of the network device; bound to the MAC address.	All levels

Table 435: show dhcp snooping binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Lease	Lease granted to the IP address.	All levels
Type	How the MAC address was acquired.	All levels
VLAN	VLAN name of the network device whose MAC address is shown.	All levels
Interface	Interface address (port).	All levels

Sample Output

show dhcp snooping binding

```
user@switch> show dhcp snooping binding

DHCP Snooping Information:
MAC Address             IP Address Lease   Type     VLAN    Interface
-----------------       ---------- -----   -------  ----    ---------
00:00:01:00:00:03       192.0.2.0  640     dynamic  guest    ge-0/0/12.0
00:00:01:00:00:04       192.0.2.1  720     dynamic  guest    ge-0/0/12.0
00:00:01:00:00:05       192.0.2.5  800     dynamic  guest    ge-0/0/13.0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

clear dhcp snooping binding | 3517

Example: Configuring Port Security (non-ELS)

Enabling DHCP Snooping (non-ELS)

show dhcp snooping statistics

IN THIS SECTION

- Syntax | 19390
- Description | 19390
- Required Privilege Level | 19390
- Output Fields | 19390
- Sample Output | 19391
- Release Information | 19391

Syntax

```
show dhcp snooping statistics
```

Description

Display statistics for read and write operations to the DHCP snooping database.

Required Privilege Level

view

Output Fields

Table 436 on page 19390 lists the output fields for the `show dhcp snooping statistics` command. Output fields are listed in the approximate order in which they appear.

Table 436: show dhcp snooping statistics Output Fields

Field Name	Field Description
Successful Transfers	Number of entries successfully transferred from memory to the DHCP snooping database.

Table 436: show dhcp snooping statistics Output Fields *(Continued)*

Field Name	Field Description
Successful Reads	Number of entries successfully read from memory to the DHCP snooping database.
Successful Writes	Number of entries successfully written from memory to the DHCP snooping database.
Failed Transfers	Number of entries that failed being transferred from memory to the DHCP snooping database.
Failed Reads	Number of entries that failed being read from memory to the DHCP snooping database.
Failed Writes	Number of entries that failed being written from memory to the DHCP snooping database.

Sample Output

show dhcp snooping statistics

```
user@switch>                  show dhcp snooping statistics

Successful Transfers :        0   Failed Transfers :       21
Successful Reads     :        0   Failed Reads     :        0
Successful Writes    :        0   Failed Writes    :       21
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

clear dhcp snooping statistics | 3519

Understanding DHCP Snooping (non-ELS)

show dhcpv6 client binding

IN THIS SECTION

- Syntax | 19392
- Description | 19392
- Options | 19392
- Required Privilege Level | 19393
- Output Fields | 19393
- Sample Output | 19394
- Release Information | 19394

Syntax

```
show dhcpv6 client binding
interface interface-name
routing-instance <routing-instance-name>
[brief | detail | summary]
```

Description

Display the address bindings in the Dynamic Host Configuration Protocol version 6 (DHCPv6) client table.

Options

interface *interface-name*	(Optional) Perform this operation on the specified interface.
routing-instance *routing-instance-name*	(Optional) Display DHCPv6 binding information for DHCPv6 clients on the specified routing instance.
brief	(Optional) Display brief information about the active client bindings.
detail	(Optional) Display detailed client binding information.

summary (Optional) Display a summary of DHCPv6 client information.

Required Privilege Level

view

Output Fields

Table 437 on page 19393 lists the output fields for the `show dhcpv6 client binding` command. Output fields are listed in the approximate order in which they appear.

Table 437: show dhcpv6 client binding Output Fields

Field Name	Field Description
Hardware Address	Hardware address of the DHCPv6 client.
State	State of the address-binding table on the DHCPv6 local server.
Lease Expires	Date and time at which the client's IP address lease expires.
Lease Expires in	Number of seconds until the lease expires.
Lease Start	Date and time at which the client's IP address lease started.
Client DUID	The DHCPv6 client's unique identifier.
Bind type	The bind type.
Client Type	The type of DHCPv6 client. The client type can be autoconfig or stateful.
Rapid Commit	Two-message exchange option for address assignment.
Server IP Address	IP address of the DHCPv6 server.
Client IP Address	IP address of the DHCPv6 client.

Sample Output

show dhcpv6 client binding

```
user@host> show dhcpv6 client binding
IP prefix                 Expires         ClientType     State  Interface       Client DUID
    2001:db8::b2b7:8631:d968:8d5e/128     96              STATEFUL      BOUND ge-0/0/1.0
LL_TIME0x3-0x0-2c:6b:f5:62:39:c1
```

command-name

```
user@host> show dhcpv6 client binding detail
Client Interface: ge-0/0/1.0
             Hardware Address:             2c:6b:f5:62:39:c1
             State:                        BOUND(DHCPV6_CLIENT_STATE_BOUND)
             Lease Expires:                2012-08-07 15:52:19 UTC
             Lease Expires in:             116 seconds
                Lease Start:                  2012-08-07 15:50:19 UTC
             Client DUID                   VENDOR0x00000583-0x3000103f
                Bind Type:                        IA_NA
                ClientType :                     STATEFUL
             Rapid Commit                  Off
             Server Ip Address:            fe80::230:48ff:fe5d:5bf7
             Client IP Address:            2001:db8::655b:3c80:2deb:1a3/128

        DHCP options:
            Name: server-identifier, Value: LL_TIME0x1-0x17acddab-00:30:48:5d:5b:f7
            Name: vendor-opts, Value: 000005830002aaaa
            Name: sip-server-list, Value: 2000::300 2000::302 2000::303 2000::304
            Name: dns-recursive-server, Value: 2000::ff2000::fe
            Name: domain-search-list, Value: 076578616d706c6503636f6d00
```

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

clear dhcpv6 client binding | 3521

show dhcpv6 client statistics

IN THIS SECTION

- Syntax | 19395
- Description | 19395
- Options | 19395
- Required Privilege Level | 19395
- Output Fields | 19396
- Sample Output | 19397
- Release Information | 19398

Syntax

```
show dhcpv6 client statistics
routing-instance<routing-instance-name>
```

Description

Display Dynamic Host Configuration Protocol (DHCPv6) client statistics.

Options

routing-instance <routing-instance-name>	(Optional) Display the statistics for DHCPv6 clients on the specified routing instance.

Required Privilege Level

view

Output Fields

Table 438 on page 19396 lists the output fields for the `show dhcpv6 client statistics` command. Output fields are listed in the approximate order in which they appear.

Table 438: show dhcpv6 client statistics Output Fields

Field Name	Field Description
Dhcpv6 Packets dropped	Number of packets discarded by the DHCPv6 local server because of errors. Only nonzero statistics appear in the DHCPv6 Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.
Messages sent	Number of DHCPv6 messages sent. • DHCPV6_DECLINE—Number of DHCPv6 PDUs of type DECLINE transmitted • DHCPV6_SOLICIT—Number of DHCPv6 PDUs of type SOLICIT transmitted • DHCPV6_INFORMATION_REQUEST—Number of DHCPv6 PDUs of type INFORMATION REQUEST transmitted • DHCPV6_RELEASE—Number of DHCPv6 PDUs of type RELEASE transmitted • DHCPV6_REQUEST—Number of DHCPv6 PDUs of type REQUEST transmitted • DHCPV6_CONFIRM—Number of DHCPv6 PDUs of type CONFIRM transmitted • DHCPV6_RENEW—Number of DHCPv6 PDUs of type RENEW transmitted • DHCPV6_REBIND—Number of DHCPv6 PDUs of type REBIND transmitted

Table 438: show dhcpv6 client statistics Output Fields *(Continued)*

Field Name	Field Description
Messages received	Number of DHCPv6 messages received. • DHCPV6_ADVERTISE—Number of DHCPv6 PDUs of type ADVERTISE received • DHCPV6_REPLY—Number of DHCPv6 PDUs of type REPLY received • DHCPV6_RECONFIGURE—Number of DHCPv6 PDUs of type RECONFIGURE received

Sample Output

show dhcpv6 client statistics

```
user@host> show dhcpv6 client statistics
Dhcpv6 Packets dropped:
            Total               0

            Messages sent:
            DHCPV6_DECLINE             0
            DHCPV6_SOLICIT             3
            DHCPV6_INFORMATION_REQUEST 6
            DHCPV6_RELEASE             1
            DHCPV6_REQUEST             2
            DHCPV6_CONFIRM             0
            DHCPV6_RENEW               0
            DHCPV6_REBIND              0

            Messages received:
            DHCPV6_ADVERTISE           3
            DHCPV6_REPLY               3
            DHCPV6_RECONFIGURE         0
```

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

> **NOTE**: This command is not supported in Junos OS Evolved because it is not needed. Stateless DHCPv4 and DHCPv6 relay are supported and enabled by default, so Junos OS Evolved does not maintain states or records of statistics or bindings for DHCP clients.

RELATED DOCUMENTATION

clear dhcpv6 client statistics | 3522

show dhcpv6 relay active-leasequery

IN THIS SECTION

- Syntax | 19398
- Description | 19399
- Options | 19399
- Required Privilege Level | 19399
- Output Fields | 19400
- Sample Output | 19403
- Release Information | 19404

Syntax

```
show dhcpv6 relay active-leasequery
<details | summary>
<interface interface-name>
<logical-system logical-system-name>
<peer ipv6-address>
```

```
<routing-instance routing-instance-name>
<statistics>
```

Description

Display information about DHCPv6 active leasequery peer relay agents.

Options

details (Optional) Display detailed information about active leasequery peers or interfaces. You must also specify either the `peer` *`ip-address`* option or the `interface` *`interface-name`* option.

summary (Optional) Display summary information about active leasequery peers. The `summary` option produces the same output as not specifying any option. You can also specify the `logical-system` *`logical-system-name`* option or the `routing-instance` *`routing-instance-name`* option with the `summary` option. You cannot specify any other option with the `summary` option.

interface ***interface-name*** (Optional) Display active leasequery statistics for a specific access interface. You must also specify the `statistics` option.

logical-system ***logical-system-name*** (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.

peer ***ipv6-address*** (Optional) Display information about active leasequery peer relay agents. You must also specify either the `details` option or the `statistics` option.

routing-instance ***routing-instance-name*** (Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.

statistics (Optional) Display active leasequery statistics for a specific active leasequery peer relay agent or a specific access interface. You must also specify either the `peer` *`ip-address`* option or the `interface` *`interface-name`* option.

Required Privilege Level

view

Output Fields

Table 439 on page 19400 lists the output fields for the `show dhcpv6 relay active-leasequery` command. Output fields are listed in the approximate order in which they appear.

Table 439: show dhcpv6 relay active-leasequery Output Fields

Field Name	Field Description
`peer`	IP address of topology discovery peer relay agent.
`Connected Peers`	Number of access interfaces that have completed the topology discovery process and are associated with the peer. Each interface is in the Connected state.
`Connecting Peers`	Number of access interfaces that are configured with topology discovery but have not completed the topology discovery process. Consequently these interfaces are not yet associated with any peer. Each interface is in the Connecting state. When topology discovery completes and the interface is associated with a peer, the state moves to Connected.
`Local Circuit ID`	The local access interface for the specified peer.
`Remote Circuit ID`	The remote access interface on another peer relay agent that corresponds to the specified peer's local access interface.
`Local Interface Address`	IP address of the local access interface.
`State`	State of the topology discovery process. • `Done`—Topology discovery has completed for the specified peer.
`Redundancy State`	VRRP redundancy state associated with the peer or interface. • `Backup`—The specified peer or interface is currently in backup mode. This means that the BNG hosting the relay agent is the current backup BNG for the group. • `Master`—The specified peer or interface is currently in primary mode. This means that the BNG hosting the relay agent is the current primary BNG for the group. • `Unknown`—The VRRP redundancy state of the peer or interface is unknown.

Table 439: show dhcpv6 relay active-leasequery Output Fields *(Continued)*

Field Name	Field Description
`xid`	Randomly generated, temporary transaction ID for the topology discovery query sent for the local access interface.
`Remote ALQ Status`	State of the active leasequery process. • `Done`—The active leasequery process has completed. Subscriber state and binding information has been synchronized for subscriber groups that use the local access interface. • `Queued`—The active leasequery is queued. Subscriber state and binding information are not yet synchronized/ • `Unknown`—Active leasequery process state is unknown.
`Interface`	Name of the specified address for which statistics are displayed.
`Topology-Discover Configured`	Indicates whether topology discovery has been configured on the specified peer or interface, `Yes` or `No`.
`Bindings Sent`	Number of DHCP bindings sent based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings sent over the interface. For the peer, it's the count of all bindings sent over all interfaces that belong to the peer.
`Bindings Received`	Number of DHCP bindings received based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings received over the interface. For the peer, it's the count of all bindings received over all interfaces that belong to the peer.

Table 439: show dhcpv6 relay active-leasequery Output Fields *(Continued)*

Field Name	Field Description
`Bindings Installed Successfully`	Number of DHCP bindings successfully installed based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings successfully installed over the interface. For the peer, it's the count of all bindings successfully installed over all interfaces that belong to the peer.
`Bindings Failed to Install`	Number of DHCP bindings that failed to install based on active leasequery for the specified interface or peer. For the interface, it's the count of all bindings that failed over the interface. For the peer, it's the count of all bindings that failed over all interfaces that belong to the peer.
`Last Synchronization Time`	Time stamp for the last time synchronization for subscribers occurred for the specified interface or peer.
`ALQ Transmit Buffer count`	The maximum TCP received and transmitted bytes buffer for the jdhcpd process. The value is nonconfigurable and is always 65,535 (ffff).
`Max Leasequery Transmit Rate`	The transmit rate of bulk leasequery replies per second. It is the maximum number of bulk leasequery replies per one-second interval. The value is nonconfigurable and is always 60. For example, if 600 leases are synchronized to the peer, it takes 10 seconds to complete.
`Local Interface Count`	Number of local access interfaces for the specified peer or interface. For the interface, the value is always 1.
`Remote Interface Count`	Number of remote access interfaces for the specified peer or interface. For the interface, the value is always 1.

Sample Output

show dhcpv6 relay active-leasequery (Summary)

```
user@host> show dhcpv6 relay active-leasequery summary

            Active Lease-query peer List:

            peer                  Connected Peers          Connecting Peers
            2001:db8::fa:2        1                        0
            2001:db8::fe:3        1                        0
```

show dhcpv6 relay active-leasequery (Peer Details)

```
user@host> show dhcpv6 relay active-leasequery peer 2001:db8::fa:2 details

 Local Circuit-ID     Remote Circuit-ID    Local Interface Address   State    Redundancy State
xid          Remote ALQ Status
 ge-2/0/3.50          ge-1/3/1.35          2001:db8::8:10            Done     Backup
0x521bc8    Done
```

show dhcpv6 relay active-leasequery (Peer Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics peer 2001:db8::fa:2 routing-instance
RI-test-vrf

     peer : 2001:db8::fa:2
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 420
     Bindings Received                      : 0
     Bindings Installed Successfully        : 0
     Bindings Failed to install             : 0
     Last Synchronization Time              : 2019-03-22 09:15:15 IST
     ALQ Transmit Buffer count              : ffff
     Max Leasequery Transmit Rate           : 60
     Local Interface count                  : 4
     Remote Interface count                 : 4
```

show dhcpv6 relay active-leasequery (Interface Statistics)

```
user@host> show dhcpv6 relay active-leasequery statistics interface ge-0/0/0.1 routing-instance
R1-test-vrf

     Interface : ge-0/0/0.1
     Topology-Discover Configured           : Yes
     State                                  : Done
     Bindings Sent                          : 400
     Bindings Received                      : 0
     Bindings Installed Successfully        : 0
     Bindings Failed to install             : 0
     Last Synchronization Time              : 2019-02-28 14:11:32 IST
     ALQ Transmit Buffer count              : ffff
     Max Leasequery Transmit Rate           : 60
     Local Interface count                  : 1
     Remote Interface count                 : 1
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear dhcpv6 relay active-leasequery statistics | 3524

show dhcp relay active-leasequery | 19309

show dhcpv6 relay binding

IN THIS SECTION

- Syntax | 19405
- Description | 19405
- Options | 19405
- Required Privilege Level | 19406

- Output Fields | 19406
- Sample Output | 19409
- Release Information | 19416

Syntax

```
show dhcpv6 relay binding
<address>
<brief>
<detail>
<interface interface-name>
<interface-tag interface-tag-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<summary>
```

Description

Display the DHCPv6 address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

Options

address (Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:

- *CID*—The specified Client ID (CID).
- *ipv6-prefix*—The specified IPv6 prefix.
- *session-id*—The specified session ID.

brief (Optional) Display brief information about the active client bindings. This is the default, and produces the same output as `show dhcpv6 relay binding`.

detail (Optional) Display detailed client binding information.

interface *interface-name*	(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and S-VLAN ID.
interface-tag *interface-tag-name*	(Optional) Displays the binding information for the specified interface-tag.
interfaces-vlan	(Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.
interfaces-wildcard	(Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).
logical-system *logical-system-name*	(Optional) Perform this operation on the specified logical system.
routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance.
summary	(Optional) Display a summary of DHCPv6 client information.

Required Privilege Level

view

Output Fields

Table 440 on page 19406 lists the output fields for the `show dhcpv6 relay binding` command. Output fields are listed in the approximate order in which they appear.

Table 440: show dhcpv6 relay binding Output Fields

Field Name	Field Description	Level of Output
number clients,(*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* rebinding, *number* releasing)	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	summary
Client IPv6 Prefix	Prefix of the DHCPv6 client.	brief detail
Client IPv6 Excluded Prefix	IPv6 Prefix of the DHCP client excluded.	detail

Table 440: show dhcpv6 relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Client DUID`	DHCP for IPv6 Unique Identifier (DUID) of the client.	`brief detail`
`Client IPv6 Address`	IPv6 address assigned to the subscriber.	`detail`
`Session Id`	Session ID of the subscriber session.	`brief detail`
`Expires`	Number of seconds in which the lease expires.	`brief detail`
`State`	State of the DHCPv6 relay address binding table on the DHCPv6 client: • `BOUND`—Client has an active IP address lease. • `INIT`—Initial state. • `REBINDING`—Client is broadcasting a request to renew the IP address lease. • `RECONFIGURE`—Client is broadcasting a request to reconfigure the IP address lease. • `RELEASE`—Client is releasing the IP address lease. • `RENEWING`—Client is sending a request to renew the IP address lease. • `REQUESTING`—Client is requesting a DHCPv6 server. • `SELECTING`—Client is receiving offers from DHCPv6 servers.	`brief detail`
`Interface`	Incoming client interface.	`brief`
`Interface tag`	Name of the Interface tag.	detail

Table 440: show dhcpv6 relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Lease Expires`	Date and time at which the client's IP address lease expires.	`detail`
`Lease Expires in`	Number of seconds in which the lease expires.	`detail`
`Preferred Lease Expires`	Date and UTC time at which the client's IPv6 prefix expires.	`detail`
`Preferred Lease Expires in`	Number of seconds at which the client's IPv6 prefix expires.	`detail`
`Lease Start`	Date and time at which the client's IP address lease started.	`detail`
`Lease time violated`	Lease time violation has occurred.	`detail`
`Incoming Client Interface`	Client's incoming interface.	`detail`
`Server Address`	IP address of the DHCPv6 server. Displays `unknown` for a DHCPv6 relay agent in a multi-relay topology that is not directly adjacent to the DHCPv6 server and does not detect the IP address of the server. In that case, the output instead displays the `Next Hop Server Facing Relay` field.	`detail`
`Next Hop Server Facing Relay`	Next-hop address in the direction of the DHCPv6 server.	`detail`
`Server Interface`	Interface of the DHCPv6 server.	`detail`
`Relay Address`	IP address of the relay.	`detail`

Table 440: show dhcpv6 relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Client Pool Name	Address pool that granted the client lease.	detail
Client ID Length	Length of client ID.	All levels
Client Id	Client ID.	All levels
Generated Circuit ID	Circuit ID generated by the DHCPv6 Interface-ID option (option 18)	detail
Generated Remote ID Enterprise Number	The Juniper Networks IANA private enterprise number	detail
Generated Remote ID	Remote ID generated by the DHCPv6 Remote-ID option (option 37)	detail
Dual Stack Group	Name of the dual-stack group for the DHCPv6 binding.	detail
Dual Stack Peer Address	Address of the dual-stack DHCPv4 peer.	detail

Sample Output

show dhcpv6 relay binding

```
user@host> show dhcpv6 relay binding
Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:db8:3c4d:15::/64   1           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
2001:db8:3c4d:16::/64   2           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:db8:3c4d:17::/64   3           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:db8:3c4d:18::/64   4           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
```

```
2001:db8:3c4d:19::/64    5           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:db8:3c4d:20::/64    6           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06
```

show dhcpv6 relay binding (Address)

```
user@host> show dhcp6 relay binding 2001:db8:1111:2222::/64 detail
Session Id:  1
     Client IPv6 Prefix:                    2001:db8:3c4d:15::/64
     Client DUID:                           LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
     State:                                 BOUND(RELAY_STATE_BOUND)
     Lease Expires:                         2011-05-25 07:12:09 PDT
     Lease Expires in:                      77115 seconds
     Preferred Lease Expires:               2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2011-05-24 07:12:09 PDT
     Incoming Client Interface:             ge-1/0/0.0
     Interface tag:                         None
     Server Address:                        2001:db8:aaaa:bbbb::1
     Server Interface:                      none
     Relay Address:                         2001:db8:1111:2222::
     Client Pool Name:                      pool-25
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x4bfa26af/0x00109400/0x0001
```

show dhcpv6 relay binding detail (Client ID)

```
user@host> show dhcpv6 relay binding 14/0x00010001/0x4bfa26af/0x00109400/0x0001 detail
Session Id:  1
     Client IPv6 Prefix:                    2001:db8:3c4d:15::/64
     Client DUID:                           LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
     State:                                 BOUND(RELAY_STATE_BOUND)
     Lease Expires:                         2011-05-25 07:12:09 PDT
     Lease Expires in:                      77115 seconds
     Preferred Lease Expires:               2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2011-05-24 07:12:09 PDT
     Lease time violated:                   yes
```

```
     Incoming Client Interface:             ge-1/0/0.0
     Interface tag:                         None
     Server Address:                        2001:db8:aaaa:bbbb::1
     Server Interface:                      none
     Relay Address:                         2001:db8:1111:2222::
     Client Pool Name:                      pool-25
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x4bfa26af/0x00109400/0x0001
```

show dhcpv6 relay binding detail

```
user@host> show dhcpv6 relay binding detail
Session Id:  1
     Client IPv6 Prefix:                    2001:db8:3c4d:15::/64
     Client DUID:                           LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
     State:                                 BOUND(RELAY_STATE_BOUND)
     Lease Expires:                         2011-05-25 07:12:09 PDT
     Lease Expires in:                      77115 seconds
     Preferred Lease Expires:               2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2011-05-24 07:12:09 PDT
     Lease time violated:                   yes
     Incoming Client Interface:             ge-1/0/0.0
     Interface tag:                         None
     Server Address:                        2001:db8:aaaa:bbbb::1
     Server Interface:                      none
     Relay Address:                         2001:db8:1111:2222::
     Client Pool Name:                      pool-25
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x4bfa26af/0x00109400/0x0001
     Generated Remote ID Enterprise Number: 1411
     Generated Remote ID:                   host:ge-1/0/0:100
```

show dhcpv6 relay binding detail (Dual-Stack)

```
user@host> show dhcpv6 relay binding detail
Session Id:  2
     Client IPv6 Prefix:                    2001:db8:ffff:0:4::/64
     Client IPv6 Address:                   2001:db8:3000:8003::1/128
     Client DUID:                           LL0x1-00:00:64:01:01:02
```

```
    State:                                  BOUND(DHCPV6_RELAY_STATE_BOUND)
    Lease Expires:                          2016-10-17 07:39:25 PDT
    Lease Expires in:                       3450 seconds
    Lease Start:                            2016-10-17 06:39:25 PDT
    Last Packet Received:                   2016-10-17 06:39:25 PDT
    Incoming Client Interface:              ae0.3221225472
    Interface tag:                          None
    Client Interface Svlan Id:              2000
    Client Interface Vlan Id:               1
    Server Ip Address:                      2001:db8:3000::2
    Server Interface:                       none
    Client Profile Name:                    my-dual-stack
    Client Id Length:                       10
    Client Id:                              /0x00030001/0x00006401/0x0102
    Dual Stack Group:                       group1
    Dual Stack Peer Address:                192.0.2.4
```

show dhcpv6 relay binding detail (Multi-Relay Topology)

```
user@host > show dhcpv6 relay binding detail
Session Id:  13
    Client IPv6 Prefix:                     2001:db8:3000:0:8001::5/128
    Client DUID:                            LL0x1-00:00:65:03:01:02
    State:                                  BOUND(DHCPV6_RELAY_STATE_BOUND)
    Lease Expires:                          2011-11-21 06:14:50 PST
    Lease Expires in:                       293 seconds
    Preferred Lease Expires:                2012-07-24 00:18:14 UTC
    Preferred Lease Expires in:             600 seconds
    Lease Start:                            2011-11-21 06:09:50 PST
    Incoming Client Interface:              ge-1/0/0.0
    Interface tag:                          None
    Server Address:                         unknown
    Next Hop Server Facing Relay:           2001:db8:4000::2
    Server Interface:                       none
    Client Id Length:                       10
    Client Id:                              /0x00030001/0x00006503/0x0102
```

show dhcpv6 relay binding (Session ID)

```
user@host> show dhcpv6 relay binding 41
Prefix                  Session Id   Expires   State   Interface    Client DUID
2001:db8:3c4d:15::/64   41           78837     BOUND   ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
```

show dhcpv6 relay binding (Subscriber with Multiple Addresses)

```
user@host> show dhcpv6 relay binding
Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:db8:1001::1:24/128          23          593      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:1001::1:1c/128          23          393      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:1001::1:14/128          23          193      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:3001::300/120           23          293      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:3001::200/120           23          193      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:3001::100/120           23          93       BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
```

When DHCPv6 relay binding is configured with prefix exclude option, we get the following output:

```
user@host> show dhcpv6 relay binding detail
Session Id:  6
     Hardware Address:                      00:10:94:00:00:01
     Client IPv6 Address:                   7001:2:3::d/128
     Lease Expires:                         2017-12-11 07:45:27 IST
     Lease Expires in:                      9999952 seconds
     Preferred Lease Expires:               2017-12-11 07:45:27 IST
     Preferred Lease Expires in:            9999952 seconds
     Client IPv6 Prefix:                    7001::1000:0:0:0/68
     Client IPv6 Excluded Prefix:           7001::1fff:ffff:ffff:ff00/120
     Lease Expires:                         2017-12-11 07:45:27 IST
     Lease Expires in:                      9999952 seconds
     Preferred Lease Expires:               2017-12-11 07:45:27 IST
     Preferred Lease Expires in:            9999952 seconds
```

```
    Client DUID:                          LL_TIME0x1-0x599553b0-00:10:94:00:00:01
    State:                                BOUND(DHCPV6_RELAY_STATE_BOUND)
    Lease Start:                          2017-08-17 13:58:33 IST
    Last Packet Received:                 2017-08-17 13:58:48 IST
    Incoming Client Interface:            ge-0/0/0.100
    Interface tag:                        None
    Client Interface Vlan Id:             100
    Server Ip Address:                    7002::1
    Server Interface:                     none
    Client Id Length:                     14
    Client Id:                            /0x00010001/0x599553b0/0x00109400/0x0001
    Generated Circuit ID:                 ge-0/0/0:100
```

show dhcpv6 relay binding detail (Subscriber with Multiple Addresses)

```
user@host> show dhcpv6 relay binding detail
Session Id:  3
    Client IPv6 Address:                  2001:db8:1001::1:2/128
    Lease Expires:                        2015-05-15 02:34:51 PDT
    Lease Expires in:                     24 seconds
    Preferred Lease Expires:              2015-05-15 02:34:51 PDT
    Preferred Lease Expires in:           24 seconds
    Client IPv6 Address:                  2001:db8:1001::1:12/128
    Lease Expires:                        2015-05-15 02:41:31 PDT
    Lease Expires in:                     424 seconds
    Preferred Lease Expires:              2015-05-15 02:41:31 PDT
    Preferred Lease Expires in:           424 seconds
    Client IPv6 Address:                  2001:db8:1001::1:a/128
    Lease Expires:                        2015-05-15 02:38:11 PDT
    Lease Expires in:                     224 seconds
    Preferred Lease Expires:              2015-05-15 02:38:11 PDT
    Preferred Lease Expires in:           224 seconds
    Client IPv6 Prefix:                   2001:db8:3001::/120
    Lease Expires:                        2015-05-15 02:34:51 PDT
    Lease Expires in:                     24 seconds
    Preferred Lease Expires:              2015-05-15 02:34:51 PDT
    Preferred Lease Expires in:           24 seconds
    Client IPv6 Prefix:                   2001:db8:3001::200/120
    Lease Expires:                        2015-05-15 02:38:11 PDT
    Lease Expires in:                     224 seconds
    Preferred Lease Expires:              2015-05-15 02:38:11 PDT
```

```
    Preferred Lease Expires in:            224 seconds
    Client IPv6 Prefix:                    2001:db8:3001::100/120
    Lease Expires:                         2015-05-15 02:36:31 PDT
    Lease Expires in:                      124 seconds
    Preferred Lease Expires:               2015-05-15 02:36:31 PDT
    Preferred Lease Expires in:            124 seconds
    Client DUID:                           LL_TIME0x1-0x55554c6e-00:10:94:00:00:02
    State:                                 BOUND(DHCPV6_RELAY_STATE_BOUND)
    Lease Start:                           2015-05-15 02:34:21 PDT
    Last Packet Received:                  2015-05-15 02:34:22 PDT
    Incoming Client Interface:             ge-9/0/9.0
    Interface tag:                         None
    Client Interface Vlan Id:              111
    Demux Interface:                       demux0.3221225475
    Server Ip Address:                     2001:db8:5001::1
    Server Interface:                      none
    Client Profile Name:                   DHCP-IPDEMUX-PROF
    Client Id Length:                      14
    Client Id:                             /0x00010001/0x55554c6e/0x00109400/0x0002
    Generated Circuit ID:                  ge-9/0/9:111
    Generated Remote ID Enterprise Number: 1411
    Generated Remote ID:                   ge-9/0/9:111
```

show dhcpv6 relay binding (Interfaces VLAN)

```
user@host> show dhcpv6 relay binding ge-1/0/0:100-200
Prefix            Session Id  Expires  State    Interface             Client DUID
2001:DB8::/32        11        87583    BOUND    ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:19::/32     12       87583    BOUND    ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 relay binding (Interfaces Wildcard)

```
user@host> show dhcpv6 relay binding demux0
Prefix            Session Id  Expires  State    Interface           Client DUID
2001:DB8::/32        30       79681    BOUND    demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:19::/32     31       79681    BOUND    demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

```
2001:DB8:C9::/32     32       79681    BOUND    demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 relay binding (Interfaces Wildcard)

```
user@host> show dhcpv6 relay binding ge-1/3/*
Prefix            Session Id  Expires  State    Interface      Client DUID
2001:DB8::/32        22       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:19::/32     33       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:C9::/32     24       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 relay binding summary

```
user@host> show dhcpv6 relay binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

show dhcpv6 relay binding interface-tag <*interface-tag-name*>

```
user@host> show dhcpv6 relay binding interface-tag sample_tag
IP address               Session Id     Hardware address     Expires     State     Interface
2001:db8:1001::1:a/128      77          00:10:94:00:00:01      740       BOUND
ge-0/0/1.3221225472
```

Release Information

Command introduced in Junos OS Release 11.4.

interfaces-vlan and *interfaces-wildcard* options introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings

clear dhcpv6 relay binding | 3527

show dhcpv6 relay lockout-entries

IN THIS SECTION

- Syntax | 19417
- Description | 19417
- Options | 19417
- Required Privilege Level | 19417
- Output Fields | 19418
- Sample Output | 19419
- Release Information | 19420

Syntax

```
show dhcpv6 relay lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv6 relay agent lockout database.

Options

all	Display all client entries in the lockout database.
index *index*	Display detailed information for the specified client.

Required Privilege Level

view

Output Fields

Table 441 on page 19418 lists the output fields for the `show dhcpv6 relay lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 441: show dhcpv6 relay lockout-entries Output Fields

Field Name	Field Description	Level of Output
`Index`	Number identifying a specific entry in the lockout database.	`all` and `index`
`Key`	DUID identifying the client in the lockout database.	`all` and `index`
`State`	Type of lockout period for the entry: • `Grace`—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. • `Lockout`—Client is currently locked out; attempts to establish a session are rejected.	`all` and `index`
`Expires (s)`	Number of seconds until the current lockout period expires.	`all` only
`Elapsed (s)`	Number of seconds since the current lockout or grace timer started.	`all` only
`Count`	Number of consecutive times the client has been locked out.	`all` only
`Expires`	Date and time when the current lockout period ends.	`index` only
`Expires in`	Number of seconds until the current period expires.	`index` only

Table 441: show dhcpv6 relay lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Lockout count`	Number of consecutive times client has been locked out.	`index` only
`Next lockout time`	Duration of the next lockout period for this client.	`index` only
`Min lockout time`	Minimum duration for a lockout period; the initial lockout time.	`index` only
`Lockout reason`	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	`index` only

Sample Output

show dhcpv6 relay lockout-entries (All Entries)

```
user@host> show dhcpv6 relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

show dhcpv6 relay lockout-entries (Specific Entry)

```
user@host> show dhcpv6 relay lockout-entries index 2
     Index:                                 2
     Key:                                   default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                                 Lockout
     Expires:                               2018-03-29 19:06:17 IST
     Expires in:                            87
     Lockout count:                         1
     Next lockout time:                     200
```

```
    Min lockout time:                    100
    Lockout reason:                      181
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

- clear dhcpv6 relay lockout-entries | 3531
- show dhcp relay lockout-entries | 19328
- show dhcp server lockout-entries | 19377
- show dhcpv6 server lockout-entries | 19446

show dhcpv6 relay statistics

IN THIS SECTION

- Syntax | 19420
- Description | 19421
- Options | 19421
- Required Privilege Level | 19421
- Output Fields | 19421
- Sample Output | 19425
- Release Information | 19426

Syntax

```
show dhcpv6 relay statistics
<bulk-leasequery-connections>
<leasequery>
```

```
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.

Options

bulk-leasequery-connections	(Optional) Display DHCPv6 relay bulk leasequery statistics.
leasequery	(Optional) Display information about DHCPv6 relay individual leasequery statistics.
logical-system ***logical-system-name***	(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.
routing-instance ***routing-instance-name***	(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 442 on page 19422 lists the output fields for the `show dhcpv6 relay statistics` command. Output fields are listed in the approximate order in which they appear.

Table 442: show dhcpv6 relay statistics Output Fields

Field Name	Field Description
`DHCPv6 Packets dropped`	Number of packets discarded by the extended DHCPv6 relay agent application due to errors. Only nonzero statistics appear in the `Packets dropped` output. When all of the Packets dropped statistics are 0 (zero), only the `Total` field appears. • `Total`—Total number of packets discarded by the DHCPV6 relay agent application. • `Bad options`—Number of packets discarded because invalid options were specified. • `Bad send`—Number of packets that the extended DHCP relay application could not send. • `Bad src address`—Number of packets discarded because the family type was not AF_INET6. • `Client MAC validation`—Number of packets discarded because validation of the client MAC address failed. • `No client id`—Number of packets discarded because they could not be matched to a client. • `Lease Time Violation`—Number of packets discarded because of a lease time violation • `No safd`—Number of packets discarded because they arrived on an unconfigured interface. • `Short packet`—Number of packets discarded because they were too short. • `Relay hop count`—Number of packets discarded because the hop count in the packet exceeded 32.

Table 442: show dhcpv6 relay statistics Output Fields *(Continued)*

Field Name	Field Description
`Messages received`	Number of DHCPv6 messages received. • `DHCPv6_DECLINE`—Number of DHCPv6 PDUs of type DECLINE received • `DHCPV6_SOLICIT`—Number of DHCPv6 PDUs of type SOLICIT received • `DHCPV6_INFORMATION_REQUEST`—Number of DHCPv6 PDUs of type INFORMATION-REQUEST received • `DHCPV6_RELEASE`—Number of DHCPv6 PDUs of type RELEASE received • `DHCPV6_REQUEST`—Number of DHCPv6 PDUs of type REQUEST received • `DHCPV6_CONFIRM`—Number of DHCPv6 PDUs of type CONFIRM received `DHCPV6_RENEW`—Number of DHCPv6 PDUs of type RENEW received `DHCPV6_REBIND`—Number of DHCPv6 PDUs of type REBIND received `DHCPV6_RELAY_REPL`—Number of DHCPv6 PDUs of type RELAY-REPL received • `DHCPV6_RELAY_FORW`—Number of DHCPv6 RELAY-FORW PDUs received. • `DHCPV6_LEASEQUERY_REPLY`—Number of DHCPv6 replies received from the DHCPv6 sever • `DHCPV6_LEASEQUERY_DATA`—xxxx • `DHCPV6_LEASEQUERY_DONE`—The leasequery is complete
`Messages sent`	Number of DHCPv6 messages sent. • `DHCPV6_ADVERTISE`—Number of DHCPv6 ADVERTISE PDUs transmitted • `DHCP_REPLY`—Number of DHCPv6 REPLY PDUs transmitted • `DHCP_RECONFIGURE`—Number of DHCPv6 RECONFIGURE PDUs transmitted • `DHCP_RELAY_FORW`—Number of DHCPv6 RELAY-FORW PDUs transmitted • `DHCPV6_RELAY_REPL`—Number of DHCPv6 RELAY-REPL PDUs transmitted. • `DHCP6_LEASEQUERY`—Number of DHCP leasequery messages transmitted

Table 442: show dhcpv6 relay statistics Output Fields *(Continued)*

Field Name	Field Description
`Packets forwarded`	Number of packets forwarded by the extended DHCPv6 relay agent application. • `FWD REQUEST`—Number of DHCPv6 REQUEST packets forwarded • `FWD REPLY`—Number of DHCPv6 REPLY packets forwarded
`External Server Response`	State of the external DHCP server responsiveness.
`Total Requested Servers`	Total number of servers with which the DHCP relay agent has requested a bulk leasequery connection.
`Total Attempted Servers`	Total number of servers with which the DHCP relay agent has attempted to create a bulk leasequery connection.
`Total Connected`	Total number of servers that have formed a bulk leasequery connection with the DHCP relay agent.
`Total Terminated by Server`	Total number of servers that have terminated a bulk leasequery connection with the DHCP relay agent.
`Total Max Attempted`	Total number of servers where the DHCP relay agent reached the maximum retry limit when it attempted to create a bulk leasequery connection.
`Total Closed due to Errors`	Total number of bulk leasequery connections that closed due to an internal error on the DHCP relay agent.
`In-Flight Connected`	Number of current bulk leasequery connections on the DHCP relay agent.
`Bulk Leaasequery Reply Packet Retries`	Number of bulk leasequery reply packets that the DHCP relay agent has retried.

Sample Output

show dhcpv6 relay statistics

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                       2
    Lease Time Violation        1
    Client MAC validation       1

Messages received:
    DHCPV6_DECLINE              0
    DHCPV6_SOLICIT              10
    DHCPV6_INFORMATION_REQUEST  0
    DHCPV6_RELEASE              0
    DHCPV6_REQUEST              10
    DHCPV6_CONFIRM              0
    DHCPV6_RENEW                0
    DHCPV6_REBIND               0
    DHCPV6_RELAY_FORW           0
    DHCPV6_LEASEQUERY_REPLY     0
    DHCPV6_LEASEQUERY_DATA      0
    DHCPV6_LEASEQUERY_DONE      0

Messages sent:
    DHCPV6_ADVERTISE            0
    DHCPV6_REPLY                0
    DHCPV6_RECONFIGURE          0
    DHCPV6_RELAY_REPL           0
    DHCPV6_LEASEQUERY           0

Packets forwarded:
    Total                       4
    FWD REQUEST                 2
    FWD REPLY                   2

External Server Response:
    State                      Responding
```

show dhcpv6 relay statistics bulk-leasequery-connections

```
user@host> show dhcpv6 relay statistics bulk-leasequery-connections

Total Requested Servers:    0
Total Attempted Servers:    0
Total Connected:            0
Total Terminated by Server: 0
Total Max Attempted:        0
Total Closed due to Errors: 0
In-Flight Connected:        0
Bulk Leasequery Reply Packet Retries:   0
```

Release Information

Command introduced in Junos OS Release 11.4.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

clear dhcpv6 relay statistics | 3534

DHCPv6 Client MAC Address Validation to Prevent Session Hijacking

show dhcpv6 server active-leasequery

IN THIS SECTION

- Syntax | 19427
- Description | 19427
- Required Privilege Level | 19427
- Output Fields | 19427
- Sample Output | 19427
- Release Information | 19428

Syntax

```
show dhcpv6 server active-leasequery summary
```

Description

Display the active leasequery status summary of the DHCPv6 local server.

Required Privilege Level

view

Output Fields

Table 443 on page 19427 lists the output fields for the `show dhcpv6 server active-leasequery summary` command.

Table 443: show dhcpv6 server active-leasequery summary Output Fields

Field Name	Field Description
`Connected Peers`	Number of peer DHCP servers connected.
`Connecting Peers`	Number of DHCP server peers trying to connect or reconnect after failure.
`peer`	IP address of the peer DHCPserver.

Sample Output

show dhcp server active-leasequery summary

```
user@host> show dhcp server active-leasequery summary
```

```
            Active Lease-query peer List:
```

```
            peer                    Connected Peers          Connecting Peers
            192.0.2.0               1                        0
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

M:N Subscriber Service Redundancy on DHCP Server

show dhcpv6 server active-leasequery statistics

IN THIS SECTION

- Syntax | 19428
- Description | 19428
- Options | 19429
- Required Privilege Level | 19429
- Output Fields | 19429
- Sample Output | 19430
- Release Information | 19430

Syntax

```
show dhcpv6 server active-leasequery statistics
<peer peer-address>
<interface interface-name>
```

Description

Display the active leasequery statistics of the DHCPv6 local server.

Options

peer *peer-address*	IP address of the peer DHCP server on which you want to view the active leasequery statistics.
interface *interface-name*	Interface name of the DHCP server on which you want to view the active leasequery statistics.

Required Privilege Level

view

Output Fields

Table 444 on page 19429 lists the output fields for the `show dhcpv6 server active-leasequery statistics` command.

Table 444: show dhcpv6 server active-leasequery statistics Output Fields

Field Name	Field Description
`ALQ Transmit Buffer count`	Numbers of packets the active leasequery transmitted.
`Bindings Failed to install`	Number of binding installation attempt failed.
`Bindings Installed Successfully`	Number of successful binding installation.
`Bindings Received`	Number of binding request received.
`Bindings Sent`	Number of binding request sent.
`Last Synchronization Time`	Last synchronization time.
`Local Interface count`	Number of local interfaces.
`Max Leasequery Transmit Rate`	Maximum leasequery transmit rate.

Table 444: show dhcpv6 server active-leasequery statistics Output Fields *(Continued)*

Field Name	Field Description
peer	IP address of the peer DHCP server.
Remote Interface count	Number of remote interfaces.
State	Status of the topology discovery configuration.
Topology-Discover Configured	Status of the topology discovery.

Sample Output

show dhcpv6 server active-leasequery peer ip-address

```
user@host> show dhcpv6 server active-leasequery peer 192.0.2.0
```

```
peer : 192.0.2.0
Topology-Discover Configured           : Yes
State                                  : Done
Bindings Sent                          : 0
Bindings Received                      : 4
Bindings Installed Successfully        : 0
Bindings Failed to install             : 0
Last Synchronization Time              : 2021-11-22 17:53:52 IST
ALQ Transmit Buffer count              : 65535
Max Leasequery Transmit Rate           : 60
Local Interface count                  : 1
Remote Interface count                 : 1
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

clear dhcpv6 server active-leasequery statistics | 3536

show dhcpv6 server binding

IN THIS SECTION

- Syntax | 19431
- Description | 19431
- Options | 19432
- Required Privilege Level | 19432
- Output Fields | 19432
- Sample Output | 19435
- Release Information | 19440

Syntax

```
show dhcpv6 server binding
<address>
<brief | detail | summary>
<interface interface-name>
<interface-tag interface-tag-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server.

Options

address	(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show: • *CID*—The specified Client ID (CID). • *ipv6-prefix*—The specified IPv6 prefix. • *session-id*—The specified session ID.
brief | detail | summary	(Optional) Display the specified level of output about active client bindings. The default is `brief`, which produces the same output as `show dhcpv6 server binding`.
interface *interface-name*	(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.
interface-tag *interface-tag-name*	(Optional) Displays the binding information for the specified interface-tag.
interfaces-vlan	(Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.
interfaces-wildcard	(Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).
logical-system *logical-system-name*	(Optional) Display information about active client bindings for DHCPv6 clients on the specified logical system.
routing-instance *routing-instance-name*	(Optional) Display information about active client bindings for DHCPv6 clients on the specified routing instance.

Required Privilege Level

view

Output Fields

Table 445 on page 19433 lists the output fields for the `show dhcpv6 server binding` command. Output fields are listed in the approximate order in which they appear.

Table 445: show dhcpv6 server binding Output Fields

Field Name	Field Description	Level of Output
number clients, (*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* releasing)	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	summary
Prefix	Client's DHCPv6 prefix, or prefix used to support multiple address assignment.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Expires	Number of seconds in which lease expires.	brief detail
State	State of the address binding table on the extended DHCPv6 local server: • BOUND—Client has active IP address lease. • INIT—Initial state. • RECONFIGURE—Server has sent reconfigure message to client. • RELEASE—Client is releasing IP address lease. • RENEWING—Client sending request to renew IP address lease. • REQUESTING—Client requesting a DHCPv6 server. • SELECTING—Client receiving offers from DHCPv6 servers.	brief detail
Interface	Interface on which the DHCPv6 request was received.	brief
Client IPv6 Address	Client's IPv6 address.	detail

Table 445: show dhcpv6 server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Client IPv6 Prefix`	Client's IPv6 prefix.	`detail`
`Client IPv6 Excluded Prefix`	IPv6 Prefix of the DHCP client excluded.	`detail`
`Client DUID`	Client's DHCP Unique Identifier (DUID).	`brief detail`
`Lease expires`	Date and time at which the client's IP address lease expires.	`detail`
`Lease expires in`	Number of seconds in which lease expires.	`detail`
`Preferred Lease Expires`	Date and UTC time at which the client's IPv6 prefix expires.	`detail`
`Preferred Lease Expires in`	Number of seconds at which client's IPv6 prefix expires.	`detail`
`Lease Start`	Date and time at which the client's address lease was obtained.	`detail`
`Lease time violated`	Lease time violation has occurred.	`detail`
`Incoming Client Interface`	Client's incoming interface.	`detail`
`Server IP Address`	IP address of DHCPv6 server.	`detail`
`Interface tag`	Name of the Interface tag.	`detail`
`Server Interface`	Interface of DHCPv6 server.	`detail`

Table 445: show dhcpv6 server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Client Pool Name	Address pool used to assign IPv6 address.	detail
Client Prefix Pool Name	Address pool used to assign IPv6 prefix.	detail
Client Id length	Length of the DHCPv6 client ID, in bytes.	detail
Client Id	ID of the DHCPv6 client.	detail
Server Id	DHCP unique identifier (DUID) for the DHCPv6 server.	detail
Client Interface Svlan Id	S-VLAN ID of the client's incoming interface.	detail
Client Interface Vlan Id	VLAN ID of the client's incoming interface.	detail
Dual Stack Group	DHCPv6 server profile name.	detail
Dual Stack Peer Address	DHCPv6 Peer IP address.	detail

Sample Output

show dhcpv6 server binding

```
user@host> show dhcpv6 server binding
Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:db8:1111:2222::/64 6           86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:01
2001:db8:1111:2222::/64 7           86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
```

```
2001:db8:1111:2222::/64 8           86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
2001:db8:1111:2222::/64 9           86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:04
2001:db8:1111:2222::/64 10          86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:05
2001:db8:2002::1/74 11              86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:06
```

show dhcpv6 server binding detail

```
user@host> show dhcpv6 server binding detail
Session Id:  2
     Client IPv6 Prefix:                    2001:db8:ffff:0:4::/64
     Client IPv6 Address:                   2001:db8:0:8003::1/128
     Client DUID:                           LL0x1-00:00:64:01:01:02
     State:                                 BOUND(DHCPV6_LOCAL_SERVER_STATE_BOUND)
     Lease Expires:                         2016-11-07 08:30:39 PST
     Lease Expires in:                      43706 seconds
     Preferred Lease Expires:               2016-11-07 08:30:39 PST
     Preferred Lease Expires in:            43706 seconds
     Lease Start:                           2016-11-04 11:00:37 PDT
     Last Packet Received:                  2016-11-06 09:00:39 PST
     Incoming Client Interface:             ae0.3221225472
     Interface tag:                         None
     Client Interface Svlan Id:             2000
     Client Interface Vlan Id:              1
     Server Ip Address:                     2001:db8::2
     Server Interface:                      none
     Client Profile Name:                   my-dual-stack
     Client Id Length:                      10
     Client Id:                             /0x00030001/0x00006401/0x0102
     Dual Stack Group:                      my-dual-stack
     Dual Stack Peer Address:               192.0.2.10
```

command-name

When DHCPv6 binding is configured with prefix exclude option, we get the following output:

```
user@host> show dhcpv6 server binding detail
Session Id:  5
     Client IPv6 Address:                  2001:db8:2:3::d/128
     Lease Expires:                        2017-12-11 07:45:15 IST
     Lease Expires in:                     9999995 seconds
     Preferred Lease Expires:              2017-12-11 07:45:15 IST
     Preferred Lease Expires in:           9999995 seconds
     Client IPv6 Prefix:                   2001:db8::1000:0:0/68
       Client IPv6 Excluded Prefix:          2001:db8::1fff:ffff:ff00/120
     Lease Expires:                        2017-12-11 07:45:15 IST
     Lease Expires in:                     9999995 seconds
     Preferred Lease Expires:              2017-12-11 07:45:15 IST
     Preferred Lease Expires in:           9999995 seconds
     Client DUID:                          LL_TIME0x1-0x599553b0-00:10:94:00:00:01
     State:                                BOUND(DHCPV6_LOCAL_SERVER_STATE_BOUND)
     Lease Start:                          2017-08-17 13:58:32 IST
     Last Packet Received:                 2017-08-17 13:58:36 IST
     Incoming Client Interface:            ge-0/0/0.0
     Interface tag:                        None
     Client Interface Vlan Id:             100
     Client Pool Name:                     ia_na_pool
     Client Prefix Pool Name:              prefix_delegate_pool
     Client Id Length:                     14
     Client Id:                            /0x00010001/0x599553b0/0x00109400/0x0001
     Relay Id Length:                      31
     Relay Id:                             /0x00020000/0x05830130/0x303a3035/0x3a38363a
     Relay Id:                             /0x34343a65/0x323a6330/0x00000000/0x000000
```

show dhcpv6 server binding interface

```
user@host> show dhcpv6 server binding interface ge-1/0/0:10-101
Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:db8:1111:2222::/64 1           86055    BOUND    ge-1/0/0.100
LL_TIME0x1-0x4b0a53b9-00:10:94:00:00:01
```

show dhcpv6 server binding interface detail

```
user@host> show dhcpv6 server binding interface ge-1/0/0:10-101 detail
Session Id:  7
     Client IPv6 Prefix:                    2001:db8:1111:2222::/64
     Client DUID:                           LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
     State:                                 BOUND(bound)
     Lease Expires:                         2009-07-21 10:41:15 PDT
     Lease Expires in:                      86136 seconds
     Preferred Lease Expires:               2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2009-07-20 10:41:15 PDT
     Incoming Client Interface:             ge-1/0/0.0
     Interface tag:                         None
     Server Ip Address:                     0.0.0.0
     Server Interface:                      none
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x02e159c0/0x00109400/0x0002
```

show dhcpv6 server binding (IPv6 Prefix)

```
user@host> show dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005 detail
Session Id:  7
     Client IPv6 Prefix:                    2001:db8:1111:2222::/64
     Client DUID:                           LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
     State:                                 BOUND(bound)
     Lease Expires:                         2009-07-21 10:41:15 PDT
     Lease Expires in:                      86136 seconds
     Preferred Lease Expires:               2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2009-07-20 10:41:15 PDT
     Incoming Client Interface:             ge-1/0/0.0
     Interface tag:                         None
     Server Ip Address:                     0.0.0.0
     Server Interface:                      none
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x02e159c0/0x00109400/0x0002
```

show dhcpv6 server binding (Session ID)

```
user@host> show dhcpv6 server binding 8
Prefix           Session Id  Expires  State    Interface    Client DUID
2001:db8::/32    8           86235    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
```

show dhcpv6 server binding (Interfaces VLAN)

```
user@host> show dhcpv6 server binding ge-1/0/0:100-200
Prefix            Session Id  Expires  State    Interface             Client DUID
2001:db8::/32     11          87583    BOUND    ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:19::/32     12       87583    BOUND    ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 server binding (Interfaces Wildcard)

```
user@host> show dhcpv6 server binding demux0
Prefix            Session Id  Expires  State    Interface           Client DUID
2001:db8::/32     30          79681    BOUND    demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:19::/32     31       79681    BOUND    demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:C9::/32     32       79681    BOUND    demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 server binding (Interfaces Wildcard)

```
user@host> show dhcpv6 server binding ge-1/3/*
Prefix            Session Id  Expires  State    Interface      Client DUID
2001:db8::/32     22          79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:19::/32     33       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:C9::/32     24       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 server binding summary

```
user@host> show dhcpv6 server binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

show dhcpv6 server binding interface-tag <*interface-tag-name*>

```
user@host> show dhcpv6 server binding interface-tag sample_tag
IP address Session Id Hardware address Expires State Interface
2001:db8:1001::1:a/128 00:10:94:00:00:01 740 BOUND ge-0/0/1.3221225472
```

Release Information

Command introduced in Junos OS Release 9.6.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings

clear dhcpv6 server binding | 3538

show dhcpv6 server binding (View)

IN THIS SECTION

- Syntax | 19441
- Description | 19441
- Options | 19441
- Required Privilege Level | 19441
- Output Fields | 19441
- Sample Output | 19443

Release Information | **19446**

Syntax

```
show dhcpv6 server binding
<brief | detail | summary>
<interface interface-name>
<routing-instance routing-instance-name>
```

Description

Display the address bindings in the client table for DCHPv6 local server.

Options

- brief | detail | summary—(Optional) Display the specified level of output about active client bindings. The default is `brief`, which produces the same output as `show dhcpv6 server binding`.
- interface *interface-name*—(Optional) Display information about active client bindings on the specified interface.
- routing-instance *routing-instance-name*—(Optional) Display information about active client bindings for DHCPv6 clients on the specified routing instance.
-

Required Privilege Level

view

Output Fields

Table 446 on page 19442 lists the output fields for the `show dhcpv6 server binding` command. Output fields are listed in the approximate order in which they appear.

Table 446: show dhcv6p server binding Output Fields

Field Name	Field Description	Level of Output
number clients, (*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* releasing)	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	summary
Prefix	Client's DHCPv6 prefix.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Expires	Number of seconds in which lease expires.	brief detail
State	State of the address binding table on the DHCPv6 local server: • BOUND—Client has active IP address lease. • INIT—Initial state. • RELEASE—Client is releasing IP address lease. • RECONFIGURE—Client has received reconfigure message from server. • RENEWING—Client sending request to renew IP address lease. • REQUESTING—Client requesting a DHCPv6 server. • SELECTING—Client receiving offers from DHCPv6 servers.	brief detail
Interface	Interface on which the DHCPv6 request was received.	brief

Table 446: show dhcv6p server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Client DUID	Client's DHCP Unique Identifier (DUID).	brief
Lease expires	Date and time at which the client's IP address lease expires.	detail
Lease expires in	Number of seconds in which lease expires.	detail
Lease Start	Date and time at which the client's address lease was obtained.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server IP Address	IP address of DHCPv6 server.	detail
Server Interface	Interface of DHCPv6 server.	detail
Client Id length	Length of the DHCPv6 client ID, in bytes.	detail
Client Id	ID of the DHCPv6 client.	detail
Server Id	ID type and ID of the DHCPv6 server.	detail

Sample Output

show dhcpv6 server binding

```
user@host> show dhcpv6 server binding



Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:bd8:1111:2222::/64 6           86321    BOUND    ge-1/0/0.0
```

```
LL_TIME0x1-0x2e159c0-00:10:94:00:00:01
2001:bd8:1111:2222::/64 7           86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
2001:bd8:1111:2222::/64 8           86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
2001:bd8:1111:2222::/64 9           86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:04
2001:bd8:1111:2222::/64 10          86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:05
```

show dhcpv6 server binding detail

```
user@host> show dhcpv6 server binding detail
Session Id:  6
     Client IPv6 Prefix:                    2001:bd8:1111:2222::/64
     Client DUID:                           LL_TIME0x1-0x2e159c0-00:10:94:00:00:01
     State:                                 BOUND(bound)
     Lease Expires:                         2009-07-21 10:41:15 PDT
     Lease Expires in:                      86308 seconds
     Lease Start:                           2009-07-20 10:41:15 PDT
     Incoming Client Interface:             ge-1/0/0.0
     Server Ip Address:                     0.0.0.0
     Server Interface:                      none
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x02e159c0/0x00109400/0x0001      Server
Id:                             <VENDOR 2198142976/4a4e313132414343374146430000000000000000>

Session Id:  7
     Client IPv6 Prefix:                    2001:bd8:1111:2222::/64
     Client DUID:                           LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
     State:                                 BOUND(bound)
     Lease Expires:                         2009-07-21 10:41:15 PDT
     Lease Expires in:                      86308 seconds
     Lease Start:                           2009-07-20 10:41:15 PDT
     Incoming Client Interface:             ge-1/0/0.0
     Server Ip Address:                     0.0.0.0
     Server Interface:                      none
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x02e159c0/0x00109400/0x0002      Server
Id:                             <VENDOR 2198142976/4a4e313132414343374146430000000000000000>
```

show dhcpv6 server binding interface

```
user@host> show dhcp6 server binding interface ge-1/0/0:10-101
Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:bd8:1111:2222::/64 1           86055    BOUND    ge-1/0/0.100
LL_TIME0x1-0x4b0a53b9-00:10:94:00:00:01
```

show dhcpv6 server binding interface detail

```
user@host> show dhcp6 server binding interface ge-1/0/0:10-101 detail
Session Id:  7
     Client IPv6 Prefix:                    2001:bd8:1111:2222::/64
     Client DUID:                           LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
     State:                                 BOUND(bound)
     Lease Expires:                         2009-07-21 10:41:15 PDT
     Lease Expires in:                      86136 seconds
     Lease Start:                           2009-07-20 10:41:15 PDT
     Incoming Client Interface:             ge-1/0/0.0
     Server Ip Address:                     0.0.0.0
     Server Interface:                      none
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x02e159c0/0x00109400/0x0002      Server
Id:                             <VENDOR 2198142976/4a4e313132414343374146430000000000000000>
```

show dhcpv6 server binding prefix

```
user@host> show dhcp6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005 detail
Session Id:  7
     Client IPv6 Prefix:                    2001:bd8:1111:2222::/64
     Client DUID:                           LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
     State:                                 BOUND(bound)
     Lease Expires:                         2009-07-21 10:41:15 PDT
     Lease Expires in:                      86136 seconds
     Lease Start:                           2009-07-20 10:41:15 PDT
     Incoming Client Interface:             ge-1/0/0.0
     Server Ip Address:                     0.0.0.0
     Server Interface:                      none
```

```
    Client Id Length:                    14
    Client Id:                           /0x00010001/0x02e159c0/0x00109400/0x0002
```

show dhcpv6 server binding session-id

```
user@host> show dhcpv6 server binding 8
Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:bd8:1111:2222::/64 8           86235    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
```

show dhcpv6 server binding summary

```
user@host> show dhcpv6 server binding summary

5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

clear dhcpv6 server binding (Local Server) | 3542

show dhcpv6 server lockout-entries

IN THIS SECTION

- Syntax | 19447
- Description | 19447
- Options | 19447

- Required Privilege Level | 19447
- Output Fields | 19447
- Sample Output | 19449
- Release Information | 19450

Syntax

```
show dhcpv6 server lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv6 local server lockout database.

Options

all	Display all client entries in the lockout database.
index *index*	Display detailed information for the specified client.

Required Privilege Level

view

Output Fields

Table 447 on page 19448 lists the output fields for the `show dhcpv6 server lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 447: show dhcpv6 server lockout-entries Output Fields

Field Name	Field Description	Level of Output
`Index`	Number identifying a specific entry in the lockout database.	`all` and `index`
`Key`	DUID identifying the client in the lockout database.	`all` and `index`
`State`	Type of lockout period for the entry: • `Grace`—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. • `Lockout`—Client is currently locked out; attempts to establish a session are rejected.	`all` and `index`
`Expires (s)`	Number of seconds until the current lockout period expires.	`all` only
`Elapsed (s)`	Number of seconds since the current lockout or grace timer started.	`all` only
`Count`	Number of consecutive times the client has been locked out.	`all` only
`Expires`	Date and time when the current lockout period ends.	`index` only
`Expires in`	Number of seconds until the current period expires.	`index` only
`Lockout count`	Number of consecutive times client has been locked out.	`index` only
`Next lockout time`	Duration of the next lockout period for this client.	`index` only

Table 447: show dhcpv6 server lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Min lockout time`	Minimum duration for a lockout period; the initial lockout time.	`index` only
`Lockout reason`	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	`index` only

Sample Output

show dhcpv6 server lockout-entries (All Entries)

```
user@host> show dhcpv6 server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

show dhcpv6 server lockout-entries (Specific Entry)

```
user@host> show dhcpv6 server lockout-entries index 2
     Index:                                 2
     Key:                                   default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                                 Lockout
     Expires:                               2018-03-29 19:06:17 IST
     Expires in:                            87
     Lockout count:                         1
     Next lockout time:                     200
     Min lockout time:                      100
     Lockout reason:                        181
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

- clear dhcpv6 server lockout-entries | 3543
- show dhcp relay lockout-entries | 19328
- show dhcp server lockout-entries | 19377
- show dhcpv6 relay lockout-entries | 19417

show dhcpv6 server statistics

IN THIS SECTION

- Syntax | 19450
- Description | 19451
- Options | 19451
- Required Privilege Level | 19451
- Output Fields | 19451
- Sample Output | 19454
- Release Information | 19455

Syntax

```
show dhcpv6 server statistics
<bulk-leasequery-connections>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.

Options

bulk-leasequery-connections	(Optional) Display information about DHCPv6 local server bulk leasequery statistics.
logical-system ***logical-system-name***	(Optional) Display information about extended DHCPv6 local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.
routing-instance ***routing-instance-name***	(Optional) Display information about extended DHCPv6 local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 448 on page 19452 lists the output fields for the `show dhcpv6 server statistics` command. Output fields are listed in the approximate order in which they appear.

Table 448: show dhcpv6 server statistics Output Fields

Field Name	Field Description
`Packets dropped`	Number of packets discarded by the extended DHCPv6 local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears. • `Total`—Total number of packets discarded by the extended DHCPv6 local server • `Strict Reconfigure`—Number of solicit messages discarded because the client does not support reconfiguration • `Bad hardware address`—Number of packets discarded because an invalid hardware address was specified • `Bad opcode`—Number of packets discarded because an invalid operation code was specified • `Bad options`—Number of packets discarded because invalid options were specified • `Client MAC validation`—Number of packets discarded because validation of the client MAC address failed. • `Invalid server address`—Number of packets discarded because an invalid server address was specified • `Lease Time Violation`—Number of packets discarded because of a lease time violation • `No available addresses`—Number of packets discarded because there were no addresses available for assignment • `No interface match`—Number of packets discarded because they did not belong to a configured interface • `No routing instance match`—Number of packets discarded because they did not belong to a configured routing instance • `No valid local address`—Number of packets discarded because there was no valid local address • `Packet too short`—Number of packets discarded because they were too short • `Read error`—Number of packets discarded because of a system read error • `Send error`—Number of packets that the extended DHCPv6 local server could not send

Table 448: show dhcpv6 server statistics Output Fields *(Continued)*

Field Name	Field Description
`Advertise Delay`	Number of DHCP advertise messages delayed. • `DELAYED`—Number of DHCPv6 advertise packets that have been sent after being delayed. • `INPROGRESS`—Number of DHCPv6 advertise packets that are in the delay queue. • `TOTAL`—Total number of delayed DHCPv6 advertise messages; sum of `DELAYED` and `INPROGRESS`.
`Messages received`	Number of DHCPv6 messages received. • `DHCPV6_CONFIRM`—Number of DHCPv6 CONFIRM PDUs received. • `DHCPV6_DECLINE`—Number of DHCPv6 DECLINE PDUs received. • `DHCPV6_INFORMATION_REQUEST`—Number of DHCPv6 INFORMATION-REQUEST PDUs received. • `DHCPV6_REBIND`—Number of DHCPv6 REBIND PDUs received. • `DHCPV6_RELAY_FORW`—Number of DHCPv6 RELAY-FORW PDUs received. • `DHCPV6_RELAY_REPL`—Number of DHCPv6 RELAY-REPL PDUs received. • `DHCPV6_RELEASE`—Number of DHCPv6 RELEASE PDUs received. • `DHCPV6_RENEW`—Number of DHCPv6 RENEW PDUs received. • `DHCPV6_REQUEST`—Number of DHCPv6 REQUEST PDUs received. • `DHCPV6_SOLICIT`—Number of DHCPv6 SOLICIT PDUs received. • `DHCPV6_LEASEQUERY`—Number of DHCPv6 leasequery messages received.

Table 448: show dhcpv6 server statistics Output Fields *(Continued)*

Field Name	Field Description
`Messages sent`	Number of DHCPv6 messages sent. • `DHCPV6_ADVERTISE`—Number of DHCPv6 ADVERTISE PDUs transmitted. • `DHCPV6_REPLY`—Number of DHCPv6 ADVERTISE PDUs transmitted. • `DHCPV6_LOGICAL_NAK`—Number of logical NAK messages sent, signifying T1 and T2 timers with values of zero; subset of `DHCPV6_REPLY` counter. (Displays only at `verbose` level. • `DHC6_RECONFIGURE`—Number of DHCPv6 RECONFIGURE PDUs transmitted. • `DHCPV6_RELAY_REPL`—Number of DHCPv6 RELAY-REPL PDUs transmitted. • `DHCPV6_RELAY_FORW`—Number of DHCPv6 RELAY-FORW PDUs transmitted. • `DHCPV6_LEASEQUERY_REPLY`—Number of DHCPv6 leasequery replies transmitted to the DHCPv6 relay agent. • `DHCPV6_LEASEQUERY_DATA`—Number of DHCPv6 LEASEQUERY-DATA packets transmitted. • `DHCPV6_LEASEQUERY_DONE`—Number of DHCPv6 LEASEQUERY-DONE packets sent.

Sample Output

show dhcpv6 server statistics

```
user@host> show dhcpv6 server statistics
Dhcpv6 Packets dropped:
    Total                      2
    Lease Time Violation       1
    Client MAC validation      1

Advertise Delay:
    DELAYED                    3
    INPROGRESS                 9
    TOTAL                     12

Messages received:
    DHCPV6_DECLINE             0
```

```
    DHCPV6_SOLICIT             9
    DHCPV6_INFORMATION_REQUEST 0
    DHCPV6_RELEASE             0
    DHCPV6_REQUEST             5
    DHCPV6_CONFIRM             0
    DHCPV6_RENEW               0
    DHCPV6_REBIND              0
    DHCPV6_RELAY_FORW          0

    DHCPV^_LEASEQUERY          0

Messages sent:
    DHCPV6_ADVERTISE           9
    DHCPV6_REPLY               5
    DHCPV6_RECONFIGURE         0
    DHCPV6_RELAY_REPL          0
    DHCPV6_LEASEQUERY_REPLY    0
    DHCPV6_LEASEQUERY_DATA     0
    DHCPV6_LEASEQUERY_DONE     0
```

show dhcpv6 server statistics bulk-leasequery-connections

```
user@host> show dhcpv6 server statistics bulk-leasequery-connections

Total Accepted Connections:                     0
Total Not-Accepted Connections:                 0
Connections Closed due to Errors:               0
Connections Closed due to max-empty-replies:    0
In-flight Connections:                          0
```

Release Information

Command introduced in Junos OS Release 9.6.

`bulk-leasequery-connections` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

clear dhcpv6 server statistics | 3546

show dhcp client statistics | 19305

show dhcpv6 server statistics (View)

IN THIS SECTION

- Syntax | 19456
- Description | 19456
- Options | 19456
- Required Privilege Level | 19457
- Output Fields | 19457
- Sample Output | 19459
- Release Information | 19460

Syntax

```
show dhcpv6 server statistics
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display DHCPv6 local server statistics.

Options

logical-system *logical-system-name*	(Optional) Display information about extended DHCPv6 local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.
routing-instance *routing-instance-name*	(Optional) Display information about DHCPv6 local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 449 on page 19458 lists the output fields for the `show dhcpv6 server statistics` command. Output fields are listed in the approximate order in which they appear.

Table 449: show dhcpv6 server statistics Output Fields

Field Name	Field Description
`Dhcpv6 Packets dropped`	Number of packets discarded by the DHCPv6 local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears. • `Total`—Total number of packets discarded by the DHCPv6 local server • `Authentication`—Number of packets discarded because they could not be authenticated • `Strict Reconfigure`—Number of solicit messages discarded because the client does not support reconfiguration • `Bad hardware address`—Number of packets discarded because an invalid hardware address was specified • `Bad opcode`—Number of packets discarded because an invalid operation code was specified • `Bad options`—Number of packets discarded because invalid options were specified • `Invalid server address`—Number of packets discarded because an invalid server address was specified • `No available addresses`—Number of packets discarded because there were no addresses available for assignment • `No interface match`—Number of packets discarded because they did not belong to a configured interface • `No routing instance match`—Number of packets discarded because they did not belong to a configured routing instance • `No valid local address`—Number of packets discarded because there was no valid local address • `Packet too short`—Number of packets discarded because they were too short • `Read error`—Number of packets discarded because of a system read error • `Send error`—Number of packets that the DHCPv6 local server could not send

Table 449: show dhcpv6 server statistics Output Fields *(Continued)*

Field Name	Field Description
`Messages received`	Number of DHCPv6 messages received. • `DHCPV6_CONFIRM`—Number of DHCPv6 CONFIRM PDUs received. • `DHCPV6_DECLINE`—Number of DHCPv6 DECLINE PDUs received. • `DHCPV6_INFORMATION_REQUEST`—Number of DHCPv6 INFORMATION-REQUEST PDUs received. • `DHCPV6_REBIND`—Number of DHCPv6 REBIND PDUs received. • `DHCPV6_RELAY_FORW`—Number of DHCPv6 RELAY-FORW PDUs received from a relay by the DHCPv6 server. • `DHCPV6_RELEASE`—Number of DHCPv6 RELEASE PDUs received. • `DHCPV6_RENEW`—Number of DHCPv6 RENEW PDUs received. • `DHCPV6_REQUEST`—Number of DHCPv6 REQUEST PDUs received. • `DHCPV6_SOLICIT`—Number of DHCPv6 SOLICIT PDUs received.
`Messages sent`	Number of DHCPv6 messages sent. • `DHCPV6_ADVERTISE`—Number of DHCPv6 ADVERTISE PDUs transmitted. • `DHCPV6_REPLY`—Number of DHCPv6 ADVERTISE PDUs transmitted. • `DHC6_RECONFIGURE`—Number of DHCPv6 RECONFIGURE PDUs transmitted. • `DHCPV6_RELAY_REPL`—Number of DHCPv6 RELAY-REPL PDUs sent from DHCPv6 server to DHCPv6 relay.

Sample Output

show dhcpv6 server statistics

```
user@host> show dhcpv6 server statistics
Dhcpv6 Packets dropped:
    Total               0
```

```
Messages received:
    DHCPV6_DECLINE             0
    DHCPV6_SOLICIT             9
    DHCPV6_INFORMATION_REQUEST 0
    DHCPV6_RELEASE             0
    DHCPV6_REQUEST             5
    DHCPV6_CONFIRM             0
    DHCPV6_RENEW               0
    DHCPV6_REBIND              0
    DHCPV6_RELAY_FORW          0
   Messages sent:
    DHCPV6_ADVERTISE           9
    DHCPV6_REPLY               5
    DHCPV6_RECONFIGURE         0
    DHCPV6_RELAY_REPL          0
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

clear dhcpv6 server statistics (Local Server) | 3548

show dhcpv6 snooping binding

IN THIS SECTION

- Syntax | 19461
- Description | 19461
- Options | 19461
- Required Privilege Level | 19461
- Output Fields | 19461
- Sample Output | 19462

Release Information | 19462

Syntax

```
show dhcpv6 snooping binding
<interface interface-name>
<vlan vlan-name>
```

Description

Display the DHCPv6 snooping database information.

Options

interface *interface-name*	(Optional) Display the DHCPv6 snooping database information for an interface.
vlan *vlan-name*	(Optional) Display the DHCPv6 snooping database information for a VLAN.

Required Privilege Level

view

Output Fields

Table 450 on page 19461 lists the output fields for the `show dhcpv6 snooping binding` command. Output fields are listed in the approximate order in which they appear.

Table 450: show dhcp snooping binding Output Fields

Field Name	Field Description	Level of Output
MAC Address	MAC address of the network device; bound to the IP address.	All levels

Table 450: show dhcp snooping binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IP Address	IP address of the network device; bound to the MAC address.	All levels
Lease	Lease granted to the IP address.	All levels
Type	How the MAC address was acquired.	All levels
VLAN	VLAN name of the network device whose MAC address is shown.	All levels
Interface	Interface address (port).	All levels

Sample Output

show dhcpv6 snooping binding

```
user@switch> show dhcpv6 snooping binding
DHCP Snooping Information:
MAC address        IP address             Lease (seconds) Type     VLAN  Interface
00:10:94:00:00:01  2001:db8::10:10        3599992         dynamic  v1       ge-0/0/0.0
00:10:94:00:00:01  fe80::210:94ff:fe00:1  3599992         dynamic  v1    ge-0/0/0.0
00:10:94:00:00:02  2001:db8::10:11        3599992         dynamic  v1    ge-0/0/0.0
00:10:94:00:00:02  fe80::210:94ff:fe00:2  3599992         dynamic  v1    ge-0/0/0.0
00:10:94:00:00:03  2001:db8::10:12        3599992         dynamic  v1    ge-0/0/0.0
00:10:94:00:00:03  fe80::210:94ff:fe00:3  3599992         dynamic  v1    ge-0/0/0.0
00:10:94:00:00:04  2001:db8::10:13        3599992         dynamic  v1    ge-0/0/0.0
00:10:94:00:00:04  fe80::210:94ff:fe00:4  3599992         dynamic  v1    ge-0/0/0.0
00:10:94:00:00:05  2001:db8::10:14        3599992         dynamic  v1    ge-0/0/0.0
00:10:94:00:00:05  fe80::210:94ff:fe00:5  3599992         dynamic  v1    ge-0/0/0.0
```

Release Information

Command introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

clear dhcp snooping binding | 3517

Example: Configuring Port Security (non-ELS)

Enabling DHCP Snooping (non-ELS)

show dhcpv6 snooping statistics

IN THIS SECTION

- Syntax | 19463
- Description | 19463
- Required Privilege Level | 19463
- Output Fields | 19463
- Sample Output | 19464
- Release Information | 19464

Syntax

```
show dhcpv6 snooping statistics
```

Description

Display statistics for read and write operations performed on the DHCPv6 snooping database.

Required Privilege Level

view

Output Fields

Table 451 on page 19464 lists the output fields for the `show dhcpv6 snooping statistics` command. Output fields are listed in the approximate order in which they appear.

Table 451: show dhcpv6 snooping statistics Output Fields

Field Name	Field Description
Successful Transfers	Number of entries successfully transferred from memory to the DHCPv6 snooping database.
Successful Reads	Number of entries successfully read from memory to the DHCPv6 snooping database.
Successful Writes	Number of entries successfully written from memory to the DHCPv6 snooping database.
Failed Transfers	Number of entries that failed being transferred from memory to the DHCPv6 snooping database.
Failed Reads	Number of entries that failed being read from memory to the DHCPv6 snooping database.
Failed Writes	Number of entries that failed being written from memory to the DHCPv6 snooping database.

Sample Output

show dhcpv6 snooping statistics

```
user@switch>                  show dhcpv6 snooping statistics
DHCP Snoop Persistence statistics
Successful Remote Transfers: 0          Failed Remote Transfers: 0
Successful Record Reads    : 0          Failed Record Reads    : 0
Successful Record Writes   : 0          Failed Record Writes   : 0
```

Release Information

Command introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

clear dhcp snooping statistics | 3519

Understanding DHCP Snooping (non-ELS)

show diagnostics tdr

IN THIS SECTION

- Syntax | 19465
- Description | 19465
- Options | 19466
- Required Privilege Level | 19466
- Output Fields | 19466
- Sample Output | 19470
- Release Information | 19472

Syntax

```
show diagnostics tdr
<interface interface-name>
```

Description

Display the results of a time domain reflectometry (TDR) diagnostic test run on an interface. A TDR test characterizes and locates faults on twisted-pair Ethernet cables. For example, it can detect a broken twisted pair and provide the approximate distance to the break. It can also detect polarity swaps, pair swaps, and excessive skew.

Use Time domain reflectometry (TDR) to confirm platform and release support for TDR.

Use the `request diagnostics tdr` command to request a TDR test on a specified interface. Use the `show diagnostic tdr` command to display:

- The last TDR test results for a specified interface.
- The last TDR test results for all network interfaces on the switch that support the TDR test.

Options

none	Show summarized last results for all interfaces on the switch that support the TDR test.
interface ***interface-name***	(Optional) Show detailed last results for the specified interface or a range of interfaces. Specify a range of interfaces by entering the beginning and ending interface in the range, separated by a dash—for example, **ge-0/0/15-ge-0/0/20**.

Required Privilege Level

view

Output Fields

The following table lists all possible options for the `show diagnostics tdr` command. The options that appear may vary depending on the platform and software release. Output fields are listed in the approximate order in which they appear.

Table 452: show diagnostics tdr Output Fields

Field Name	Field Description
Interface name or **Interface**	Name of interface for which you report TDR test results.

Table 452: show diagnostics tdr Output Fields *(Continued)*

Field Name	Field Description
Test status	Status of TDR test: • **Aborted**—Test was terminated by operator before it was complete. • **Failed**—Test was not completed successfully. • **Interface *interface-name* not found**—Specified interface does not exist. • **Not Started**—No TDR test results are available for the interface. • **Passed**—Test completed successfully. The cable, however, might still have a fault—see the **Cable status** field for information about the cable. • **Started**—Test is currently running and not yet complete. • **VCT not supported on *interface-name***—TDR test is not supported on the interface.
Link status	Operating status of link: **UP** or **Down**.
MDI pair	Twisted pair for which test results are being reported, identified by pin numbers. (Displayed only when the **interface** option is used.)

Table 452: show diagnostics tdr Output Fields *(Continued)*

Field Name	Field Description

Cable status	When detailed information is displayed, status for a twisted pair:
	• **Failed**—TDR test failed on the cable pair.
	• **Impedance Mismatch**—Impedance on the twisted pair is not correct. Possible reasons for an impedance mismatch include:
	  • The twisted pair is not connected properly.
	  • The twisted pair is damaged.
	  • The connector is faulty.
	• **Normal**—No cable fault detected for the twisted pair.
	• **Open**—Lack of continuity between the pins at each end of the twisted-pair.
	• **Short on Pair-*n***—A short-circuit was detected on the twisted pair.
	When summary information for all interfaces is displayed, status for the cable as a whole:
	• **Fault**—A fault was detected on one or more of the twisted-pairs.
	• **OK**—No fault was detected on any of the twisted pairs.

Distance fault or **Max distance fault**	Distance to the fault in whole meters. If there is no fault, this value is 0. When summary information for all interfaces is displayed, this value is the distance to the most distant fault if there is more than one twisted pair with a fault.

Table 452: show diagnostics tdr Output Fields *(Continued)*

Field Name	Field Description
Polarity swap	Indicates the polarity status of the twisted pair: • **Normal**—Polarity is normal. Each conductor in the twisted pair has been connected the same pins at the both ends of the connection. For example, a conductor connected to pin 1 at the near end of the connection is connected to pin 1 at the far end. • **Reversed**—Polarity has been reversed. For the twisted pair, the conductors have switched which pins they are connected to at the near and far ends of the connection. For example, the conductor connected to pin 1 at the near end is connected to pin 2 at the far end. (Not available on EX8200 switches.) (Displayed only when the **interface** option is used)
Skew time	Difference in nanoseconds between the propagation delay on this twisted pair and the twisted pair with the shortest propagation delay. (Not available on EX8200 switches.) (Displayed only when the **interface** option is used.)
Channel Pair	Number of the 10/100BASE-T transmit/receive pair being reported on.
Pair Swap	Indicates whether or not the twisted pairs are swapped: • **MDI**—The pairs are not swapped (straight-through cable). • **MDIX**—The pairs are swapped (cross-over cable). (Displayed only when the **interface** option is used.)
Downshift	Indicates whether the connection speed is being downshifted: • **No Downshift**—No downshifting of connection speed. • **Downshift occurs**—Connection speed is downshifted to 10 or 100 Mbps. This occurs if the cable is a two-pair cable rather than the four-pair cable that is required by GbE. (Displayed only when the **interface** option is used.)

Sample Output

show diagnostics tdr interface ge-0/0/19 (Normal Cable)

```
user@switch> show diagnostics tdr interface ge-0/0/19
Interface TDR detail:
Interface name                  : ge-0/0/19
Test status                     : Passed
Link status                     :  UP
MDI pair                        : 1-2
  Cable status                  : Normal
  Distance fault                : 0 Meters
  Polartiy swap                 : Normal
  Skew time                     : 0 ns
MDI pair                        : 3-6
  Cable status                  : Normal
  Distance fault                : 0 Meters
  Polartiy swap                 : Normal
  Skew time                     : 8 ns
MDI pair                        : 4-5
  Cable status                  : Normal
  Distance fault                : 0 Meters
  Polartiy swap                 : Normal
  Skew time                     : 8 ns
MDI pair                        : 7-8
  Cable status                  : Normal
  Distance fault                : 0 Meters
  Polartiy swap                 : Normal
  Skew time                     : 8 ns
Channel pair                    : 1
  Pair swap                     : MDI
Channel pair                    : 2
  Pair swap                     : MDI
Downshift                       : No Downshift
```

show diagnostics tdr interface ge-2/0/2 (Faulty Cable)

```
user@switch> show diagnostics tdr interface ge-2/0/2
Interface TDR detail:
Interface name                  : ge-2/0/2
```

```
Test status                     : Passed
Link status                     : Down
MDI Pair                        : 1-2
   Cable status                 : 1-2
   Distance fault               : 2 Meters
   Polartiy swap                : N/A
   Skew time                    : N/A
MDI Pair                        : 3-6
   Cable status                 : Impedance Mismatch
   Distance fault               : 3 Meters
   Polartiy swap                : N/A
   Skew time                    : N/A
MDI Pair                        : 4-5
   Cable status                 : Impedance Mismatch
   Distance fault               : 3 Meters
   Polartiy swap                : N/A
   Skew time                    : N/A
MDI Pair                        : 7-8
   Cable status                 : Short on Pair-2
   Distance fault               : 3 Meters
   Polartiy swap                : N/A
   Skew time                    : N/A
Channel pair                    : 1
   Pair swap                    : N/A
Channel pair                    : 2
   Pair swap                    : N/A
Downshift                       : N/A
```

show diagnostics tdr (All Supported Interfaces)

```
user@switch> show diagnostics tdr
Interface   Test status    Link status  Cable status  Max distance fault
 ge-0/0/0   Not Started    N/A          N/A           N/A
 ge-0/0/1   Not Started    N/A          N/A           N/A
 ge-0/0/2       Started    N/A          N/A           N/A
 ge-0/0/3       Started    N/A          N/A           N/A
 ge-0/0/4        Passed     UP           OK             0
 ge-0/0/5        Passed     UP          Fault         173
 ge-0/0/6        Passed     UP           OK             0
 ge-0/0/7        Passed     UP           OK             0
 ge-0/0/8        Passed     UP           OK             0
```

```
 ge-0/0/9        Passed     UP          OK            0
ge-0/0/10        Passed     UP          OK            0
ge-0/0/11        Passed     UP          OK            0
ge-0/0/12        Passed     UP          OK            0
ge-0/0/13        Passed     UP          OK            0
ge-0/0/14        Passed     UP          OK            0
ge-0/0/15        Passed     UP          OK            0
ge-0/0/16        Passed     UP          OK            0
ge-0/0/17        Passed     UP          OK            0
ge-0/0/18        Passed     UP          OK            0
ge-0/0/19        Passed     UP          OK            0
ge-0/0/20        Passed    Down       Fault          0
ge-0/0/21        Passed    Down       Fault          5
ge-0/0/22        Passed     UP          OK            0
ge-0/0/23        Passed     UP          OK            0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

request diagnostics tdr | **14611**

Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure)

show diameter

IN THIS SECTION

- Syntax | **19473**
- Description | **19473**
- Options | **19473**
- Required Privilege Level | **19473**
- Output Fields | **19473**
- Sample Output | **19479**

Release Information | 19481

Syntax

```
show diameter
<brief | detail | summary>
```

Description

Display information about the Diameter node.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes Diameter node status. The `brief` output adds summary information about functions, instances, network elements, and peers. The `detail` output adds summary information about routes.

Required Privilege Level

view

Output Fields

Table 453 on page 19473 lists the output fields for the `show diameter` command. Output fields are listed in the approximate order in which they appear.

Table 453: show diameter Output Fields

Field Name	Field Description	Level of Output
`Diameter process id`	ID number of the Diameter process.	All levels

Table 453: show diameter Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Functions`	Number of functions associated with Diameter.	All levels
`Connected functions`	Number of functions with active Diameter connections.	All levels
`Instances`	Number of configured Diameter instances.	All levels
`Network elements (NEs)`	Number of configured Diameter network elements.	All levels
`Connected NEs`	Number of Diameter network elements with active connections.	All levels
`Peers`	Number of Diameter peer nodes.	All levels
`Activated peers`	Number of Diameter peers with active connections.	All levels
`Open peers`	Number of peers in the open state, without active network element connections but available for a connection.	All levels
`Transports`	Number of transports configured.	All levels
`Requests queued for network transmit`	Number of requests waiting to be sent to the Diameter peers.	All levels
`Answers queued for network transmit`	Number of replies waiting to be sent to the Diameter peers.	All levels
`Expected answers from network`	Number of replies expected to be received from the Diameter peers.	All levels

Table 453: show diameter Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Requests queued for function transmit`	Number of requests waiting to be sent to the functions associated with Diameter.	All levels
`Answers queued for function transmit`	Number of replies waiting to be sent to the functions associated with Diameter.	All levels
`Expected answers from functions`	Number of replies expected to be received from the functions associated with Diameter.	All levels
`Memory used by network transmit queues`	Amount of memory consumed by network transmit queues.	All levels
`Memory used by function transmit queues`	Amount of memory consumed by function transmit queues.	All levels
`Origin-state-id`	Value of the Origin-State-ID AVP.	All levels
`Function`	Name of the function for which information is displayed.	`briefdetail`
`State`	State of the Diameter connection with the function: `Connected` or `Disconnec` (disconnected).	`briefdetail`
`Upstream Transaction Utilization`	Percent of upstream traffic used for this function.	`briefdetail`
`Downstream Transaction Utilization`	Percent of downstream traffic used for this function.	`briefdetail`

Table 453: show diameter Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Net Queue Buffer Utilization	Percent of network transmission buffer used for this function.	briefdetail
Func Queue Buffer Utilization	Percent of function transmission buffer used for this function.	briefdetail
Routed Dests	Number of destinations that have this function associated with their routes.	briefdetail
Name	Name of the Diameter instance.	briefdetail
Origin-Realm	Value of Origin-Realm attribute-value pair (AVP).	briefdetail
Origin Host	Value of Origin-Host AVP.	briefdetail
NE-Total	Number of configured network elements.	briefdetail
NE-Connected	Number of network elements with active Diameter connections.	briefdetail
Name	Name of the Diameter network element.	briefdetail
Instance	Name of the Diameter instance in which the network element is configured.	briefdetail

Table 453: show diameter Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the network element: • Connecting—None of the network element peers are in the open state and available for connection. • Selecting—One network element peer is connected and the network element is waiting for another peer to reach the open state so that it can be connected. • Partially-Connected—One network element peer is in the open state and connected. • Post-selection-delay—Three or more peers are in the open state and the network element is waiting to deactivate the peers in excess of two. • Fully-connected—Two network element peers are in the open state and connected.	briefdetail
Primary Peer	Primary peer for the network element, based on the configured peer priority.	briefdetail
Secondary Peer	Secondary peer for the network element, based on the configured peer priority.	briefdetail
Peer	Name of the peer.	briefdetail
Instance	Name of the Diameter instance in which the peer is configured.	briefdetail

Table 453: show diameter Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the peer: • Bad-Config—Misconfiguration. • Bad-Remote—Remote side does not conform to one of the decisions or is sending malformed messages. • Closed—Normal disconnect due to a request from the remote site or due to excessive watchdog timeouts. • Destructing—Peer to be deleted on the next timer tick. Until then, it performs no actions. • Disabled—Peer is administratively disabled. • Internal-error—Internal error has been detected and the peer is in the process of restarting. • No-Activation—Peer is not used by any Diameter network element. • Rejected—Connection was rejected by remote side of the connection. • Suspended—All other reasons to be suspended.	briefdetail
NE-Count	Number of network elements associated with the peer.	briefdetail
Activated Count	Activation status of the peer: • 1—Peer is activated. • 0—Peer is not activated.	briefdetail
Primary Count	Status of the peer: primary (1) or secondary (0).	briefdetail
Secondary Count	Status of the peer: secondary (0) or primary (1).	briefdetail
Route	Name of the Diameter route.	detail

Table 453: show diameter Output Fields *(Continued)*

Field Name	Field Description	Level of Output
NE	Name of the Diameter network element in which the route is configured	detail
Instance	Name of the Diameter instance in which the route is configured.	detail
Valid	Determination of whether the route is valid: yes or no.	detail
Up	State of the route: yes for an active route, no for an inactive route.	detail

Sample Output

show diameter brief

```
user@host> show diameter brief

Diameter node:
  Diameter process id                     :      1446
  Functions                               :   4
  Connected functions                     :   2
  Instances                               :   1
  Network elements(NEs)                   :   1
  Connected NEs                           :   0
  Peers                                   :   2
  Activated peers                         :   1
  Open peers                              :   0
  Transports                              :   1
  Requests queued for network transmit    :         0
  Answers queued for network transmit     :         0
  Expected answers from network           :         0
  Requests queued for function transmit   :         0
  Answers queued for function transmit    :         0
  Expected answers from functions         :         0
  Memory used by network transmit queues  :         0
  Memory used by function transmit queues :         0
```

```
  Origin-state-id                        :         0

Diameter function list:
                      Upstream    Downstream  Net Queue   Func Queue
                      Transaction Transaction Buffer      Buffer      Routed
                      Utilization Utilization Utilization Utilization Dests
  Function  State     %           %           %           %
  charging- Disconnec   0           0           0           0             0
  gx-plus   Connected   0           0           0           0             1
  jsrc      Connected   0           0           0           0             0
  packet-tr Disconnec   0           0           0           0             0

Diameter instances:
  Name         Origin-Realm         Origin-Host          NE-Total  NE-Connected
  master       orrr                 ohhh                         1         0

Diameter network-elements:
                                            Primary      Secondary
  Name         Instance     State           Peer         Peer
  n0           master       Connecting      <NONE>       <NONE>

Diameter peer list:
                                                      Activated Primary Secondary
  Peer         Instance     State          NE-Count Count     Count   Count
  p0           master       Suspended        1        1         0       0
  p100         master       No-Activation    0        0         0       0
```

show diameter detail

```
user@host> show diameter detail
...
Diameter routes:
  Route        NE           Instance     Valid Up
  dne-route1   dne1         master       yes   no
```

show diameter summary

```
user@host> show diameter summary
Diameter node:
  Diameter process id                    :      1446
```

```
    Functions                                :   4
    Connected functions                      :   2
    Instances                                :   1
    Network elements(NEs)                    :   1
    Connected NEs                            :   0
    Peers                                    :   2
    Activated peers                          :   1
    Open peers                               :   0
    Transports                               :   1
    Requests queued for network transmit     :         0
    Answers queued for network transmit      :         0
    Expected answers from network            :         0
    Requests queued for function transmit    :         0
    Answers queued for function transmit     :         0
    Expected answers from functions          :         0
    Memory used by network transmit queues   :         0
    Memory used by function transmit queues  :         0
    Origin-state-id                          :         0
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- clear diameter function statistics | 3553
- clear diameter peer | 3554
- show diameter function | 19482
- show diameter instance | 19493
- show diameter network-element | 19495
- show diameter peer | 19504
- show diameter route | 19519

show diameter function

IN THIS SECTION

- Syntax | 19482
- Description | 19482
- Options | 19482
- Required Privilege Level | 19483
- Output Fields | 19483
- Sample Output | 19485
- Release Information | 19487

Syntax

```
show diameter function
<brief | detail | summary>
<function-name>
```

Description

Display information about all functions associated with Diameter instances or only the specified function.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic function information. The `brief` output displays the summary information in a different format. The `detail` output adds information to the `brief` output.

function-name (Optional) Display information for only the specified function. Gx-Plus, JSRC, and packet-triggered-subscribers are supported functions.

Required Privilege Level

view

Output Fields

Table 454 on page 19483 lists the output fields for the `show diameter function` command. Output fields are listed in the approximate order in which they appear.

Table 454: show diameter function Output Fields

Field Name	Field Description	Level of Output
`Function name`	Name of the function for which information is displayed.	All levels
`State`	State of the Diameter connection with the function.	All levels
`Upstream transaction utilization`	Percent of upstream traffic used for this function.	All levels
`Downstream transaction utilization`	Percent of downstream traffic used for this function.	All levels
`Network transmit buffer utilization`	Percent of network transmission buffer used for this function.	All levels
`Function transmit buffer utilization`	Percent of function transmission buffer used for this function.	All levels
`Routed destinations`	Number of destinations that have this function associated with their routes.	All levels

Table 454: show diameter function Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Requests queued for network tx`	Number of requests waiting to be sent to the Diameter peers for this function.	`detail`
`Pending answers from network`	Number of replies expected from the Diameter peers for this function.	`detail`
`Answers queued for function tx`	Number of replies waiting to be sent to this function.	`detail`
`Total upstream transactions pending`	Total number of messages queued for this function.	`detail`
`Upstream transactions limit`	Total number of messages queued for this function.	`detail`
`Requests queued for function tx`	Number of requests waiting to be sent to this function.	`detail`
`Pending answers from function`	Number of replies expected to be received from this function.	`detail`
`Answers queued for network tx`	Number of replies waiting to be sent to this function.	`detail`
`Total downstream transactions pending`	Total number of messages queued for the Diameter peers.	`detail`
`Downstream transactions limit`	Maximum number of messages that can be queued for the Diameter peers.	`detail`

Table 454: show diameter function Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Buffers used by network tx queue`	Number of buffers used by messages queued for the Diameter peers.	`detail`
`Limit on network tx queue buffers`	Maximum buffer capacity available for messages queued for the Diameter peers.	`detail`
`Buffers used by function tx queue`	Number of buffers used by messages queued for this function.	`detail`
`Limit on function tx queue buffers`	Maximum buffer capacity available for messages queued for this function.	`detail`
`Origin-state-id`	Value of the Origin-State-ID AVP.	`detail`

Sample Output

show diameter function

```
user@host> show diameter function

Diameter function list:
                      Upstream    Downstream  Net Queue   Func Queue
                      Transaction Transaction Buffer      Buffer      Routed
                      Utilization Utilization Utilization Utilization Dests
  Function  State     %           %           %           %
  jsrc      Disconnec   0           0           0           0             0
```

show diameter function brief

```
user@host> show diameter function brief
```

```
Diameter function:
  Function name                          : gx-plus
  State                                  : Connected
  Upstream transaction utilization       :   0 %
  Downstream transaction utilization     :   0 %
  Network transmit buffer utilization    :   0 %
  Function transmit buffer utilization   :   0 %
  Routed destinations                    :         1

  Function name                          : jsrc
  State                                  : Disconnected
  Upstream transaction utilization       :   0 %
  Downstream transaction utilization     :   0 %
  Network transmit buffer utilization    :   0 %
  Function transmit buffer utilization   :   0 %
  Routed destinations                    :         0
```

show diameter function detail (JSRC)

```
user@host> show diameter function detail

Diameter function:
  Function name                          : jsrc
  State                                  : Disconnected
  Upstream transaction utilization       :   0 %
  Downstream transaction utilization     :   0 %
  Network transmit buffer utilization    :   0 %
  Function transmit buffer utilization   :   0 %
  Routed destinations                    :         0
  Requests queued for network tx         :         0
  Pending answers from network           :         0
  Answers queued for function tx         :         0
  Total upstream transactions pending    :         0
  Upstream transactions limit            :      1024
  Requests queued for function tx        :         0
  Pending answers from function          :         0
  Answers queued for network tx          :         0
  Total downstream transactions pending  :         0
  Downstream transactions limit          :      1024
  Buffers used by network tx queue       :         0
```

```
  Limit on network tx queue buffers      :  10485760
  Buffers used by function tx queue      :         0
  Limit on function tx queue buffers     :  10485760
```

show diameter function detail (Gx-Plus)

```
user@host> show diameter function gx-plus detail

Diameter function:
  Function name                          : gx-plus
  State                                  : Connected
  Upstream transaction utilization       :   0 %
  Downstream transaction utilization     :   0 %
  Network transmit buffer utilization    :   0 %
  Function transmit buffer utilization   :   0 %
  Routed destinations                    :         1
  Requests queued for network tx         :         0
  Pending answers from network           :         0
  Answers queued for function tx         :         0
  Total upstream transactions pending    :         0
  Upstream transactions limit            :      1024
  Requests queued for function tx        :         0
  Pending answers from function          :         0
  Answers queued for network tx          :         0
  Total downstream transactions pending  :         0
  Downstream transactions limit          :      1024
  Buffers used by network tx queue       :         0
  Limit on network tx queue buffers      :  10485760
  Buffers used by function tx queue      :         0
  Limit on function tx queue buffers     :  10485760
  Origin-state-id                        :         0
```

Release Information

Command introduced in Junos OS Release 9.6.

Support for PTSP introduced in Junos OS Release 10.2.

Support for Gx-Plus introduced in Junos OS Release 11.2.

Support for PTSP discontinued in Junos OS Release 13.1.

RELATED DOCUMENTATION

clear diameter function statistics | 3553

show diameter | 19472

show diameter function statistics | 19488

show diameter function statistics

IN THIS SECTION

- Syntax | 19488
- Description | 19488
- Options | 19488
- Required Privilege Level | 19489
- Output Fields | 19489
- Sample Output | 19491
- Release Information | 19492

Syntax

```
show diameter function statistics
<brief | detail | summary>
<function-name>
```

Description

Display statistics about all functions associated with Diameter instances or only the specified function.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic function statistics. The `brief` output displays the summary

information in a different format and adds numbers accumulated since the Diameter node was started. The `detail` output adds information to the `brief` output.

function-name (Optional) Display information for only the specified function. Gx-Plus, JSRC, and packet-triggered-subscribers are supported functions. When you specify a function, the `brief` output is displayed by default, even when you explicitly specify `summary`.

Required Privilege Level

view

Output Fields

Table 455 on page 19489 lists the output fields for the `show diameter function statistics` command. Output fields are listed in the approximate order in which they appear.

Table 455: show diameter function statistics Output Fields

Field Name	Field Description	Level of Output
`Function`	Name of the function for which information is displayed.	All levels
`Delivered Requests`	Number of requests delivered by Diameter to the application.	All levels
`Delivered Answers`	Number of answers delivered by Diameter to the application.	All levels
`Delivered Messages`	Total number of messages delivered by Diameter to the application.	All levels
`Forwarded Requests`	Number of requests sent by Diameter to the network.	All levels
`Forwarded Answers`	Number of answers sent by Diameter to the network.	All levels
`Forwarded Messages`	Number of messages sent by Diameter to the network.	All levels

Table 455: show diameter function statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Function name`	Name of the function for which information is displayed.	All levels
`Over-limit network requests`	Number of requests sent to Diameter peers that exceeded the limit on the network transmit queue.	`detail`
`Over-limit network answers`	Number of answers sent to Diameter peers that exceeded the limit on the network transmit queue.	`detail`
`Over-limit network messages`	Total number of messages sent to Diameter peers that exceeded the limit on the network transmit queue.	`detail`
`Failed to deliver requests`	Number of requests sent by Diameter to its application that were not successfully delivered.	`detail`
`Failed to deliver answers`	Number of answers sent by Diameter to its application that were not successfully delivered.	`detail`
`Failed to deliver messages`	Total number of messages sent by Diameter to its application that were not successfully delivered.	`detail`
`Over-limit function requests`	Number of requests sent to Diameter peers that exceeded the limit on the function transmit queue.	`detail`
`Over-limit function answers`	Number of answers sent to Diameter peers that exceeded the limit on the function transmit queue.	`detail`
`Over-limit function messages`	Total number of messages sent to Diameter peers that exceeded the limit on the function transmit queue.	`detail`
`Failed to forward requests`	Number of requests that were not successfully sent by Diameter to the network.	`detail`

Table 455: show diameter function statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Failed to forward answers`	Number of answers that were not successfully sent by Diameter to the network.	`detail`
`Failed to forward messages`	Total number of messages that were not successfully sent by Diameter to the network.	`detail`

Sample Output

show diameter function statistics

```
user@host> show diameter function statistics
Diameter function statistics:
           Delivered Delivered Delivered Forwarded Forwarded Forwarded
  Function Requests  Answers   Messages  Requests  Answers   Messages
  jsrc             0         0         0         0         0         0
```

show diameter function statistics brief

```
user@host> show diameter function statistics brief

Diameter function statistics:
  Function name                : jsrc
  Delivered requests           :         0         0
  Delivered answers            :         0         0
  Delivered messages           :         0         0
  Forwarded requests           :         0         0
  Forwarded answers            :         0         0
  Forwarded messages           :         0         0
```

show diameter function statistics detail

```
user@host> show diameter function statistics detail

Diameter function statistics:
  Function name                : jsrc
  Delivered requests           :         0         0
  Delivered answers            :         0         0
  Delivered messages           :         0         0
  Forwarded requests           :         0         0
  Forwarded answers            :         0         0
  Forwarded messages           :         0         0
  Over-limit network requests  :         0         0
  Over-limit network answers   :         0         0
  Over-limit network messages  :         0         0
  Failed to deliver requests   :         0         0
  Failed to deliver answers    :         0         0
  Failed to deliver messages   :         0         0
  Over-limit function requests :         0         0
  Over-limit function answers  :         0         0
  Over-limit function messages :         0         0
  Failed to forward requests   :         0         0
  Failed to forward answers    :         0         0
  Failed to forward messages   :         0         0
```

Release Information

Command introduced in Junos OS Release 9.6.

Support for PTSP introduced in Junos OS Release 10.2.

Support for Gx-Plus introduced in Junos OS Release 11.2.

Support for PTSP discontinued in Junos OS Release 13.1.

RELATED DOCUMENTATION

clear diameter function statistics | 3553

show diameter | 19472

show diameter function | 19482

show diameter instance

IN THIS SECTION

- Syntax | 19493
- Description | 19493
- Options | 19493
- Required Privilege Level | 19493
- Output Fields | 19494
- Sample Output | 19494
- Release Information | 19495

Syntax

```
show diameter instance
<brief | detail | summary>
<instance-name>
```

Description

Display information about all Diameter instances or only the specified instance.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic instance information. The `brief` output displays the summary information in a different format. The `detail` output is the same as the `brief` output.

instance-name (Optional) Display information for only the specified Diameter instance.

Required Privilege Level

view

Output Fields

Table 456 on page 19494 lists the output fields for the `show diameter instance` command. Output fields are listed in the approximate order in which they appear.

Table 456: show diameter instance Output Fields

Field Name	Field Description	Level of Output
`name`	Name of the Diameter instance.	`summary`
`Origin-realm`	Value of Origin-Realm AVP.	`summary`
`Origin-host`	Value of Origin-Host AVP.	`summary`
`NE-total`	Total number of network elements configured for this instance.	`summary`
`NE-connected`	Number of network elements with active Diameter connections.	`summary`
`Instance name`	Name of the Diameter instance.	`brief detail`
`Origin realm`	Value of Origin-Realm AVP.	`brief detail`
`Origin host`	Value of Origin-Host AVP.	`brief detail`
`NEs`	Total number of network elements configured for this instance.	`brief detail`
`Connected NEs`	Number of network elements with active Diameter connections.	`brief detail`

Sample Output

show diameter instance

```
user@host> show diameter instance
```

```
Diameter instances:
  Name         Origin-Realm         Origin-Host          NE-Total  NE-Connected
  master       rrrr                 hhhh                        1            1
```

show diameter instance detail

```
user@host> show diameter instance detail

Diameter instance:
  Instance name : master
  Origin realm  : rrrr
  Origin host   : hhhh
  NEs           :   1
  Connected NEs :   1
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

show diameter | 19472

show diameter network-element

IN THIS SECTION

- Syntax | 19496
- Description | 19496
- Options | 19496
- Required Privilege Level | 19496
- Output Fields | 19496

Sample Output | 19498
Release Information | 19499

Syntax

```
show diameter network-element
<brief | detail | summary>
<element-name>
```

Description

Display information about all Diameter network elements or only the specified network element.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic network element information. The `brief` output displays the summary information in a different format. The `detail` output adds information to the `brief` output.

element-name (Optional) Display information for only the specified network element.

Required Privilege Level

view

Output Fields

Table 457 on page 19497 lists the output fields for the `show diameter network-element` command. Output fields are listed in the approximate order in which they appear.

Table 457: show diameter network-element Output Fields

Field Name	Field Description	Level of Output
`Name`	Name of the Diameter network element.	`summary`
`Instance`	Name of the Diameter instance in which the network element is configured.	`summary`
`State`	State of the network element: • Connecting—None of the network element peers are in the open state and available for connection. • Selecting—One network element peer is connected and the network element is waiting for another peer to reach the open state so that it can be connected. • Partially-Connected—One network element peer is in the open state and connected. • Post-selection-delay—Three or more peers are in the open state and the network element is waiting to deactivate the peers in excess of two. • Fully-connected—Two network element peers are in the open state and connected.	All levels
`Primary peer`	Primary peer for the network element, based on the configured peer priority.	All levels
`Secondary peer`	Secondary peer for the network element, based on the configured peer priority.	All levels
`NE name`	Name of the Diameter network element.	`brief detail`
`Instance name`	Name of the Diameter instance in which the network element is configured.	`brief detail`
`Peers`	Number of configured peers.	`brief detail`

Table 457: show diameter network-element Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Activated peers	Number of peers that have been activated.	brief detail
Open peers	Number of peers in the open state, without active network element connections but available for a connection.	brief detail
Routes	Number of routes configured for the network element.	brief detail
Invalid routes	Number of routes that are invalid because they lack one or more of the following: application and partition, Diameter instance, or destination realm.	brief detail
Activation delay	Period in milliseconds between peer activations by the network element.	brief detail
First selection delay	Period in milliseconds that the network element waited after connecting to the first peer to allow other peers to reach the open state.	brief detail
Post selection delay	Period in milliseconds that the network element waited after having two peers in the open state before deactivating all lower-priority peers.	brief detail

Sample Output

show diameter network-element

```
user@host> show diameter network-element

Diameter network-elements:
                                          Primary      Secondary
  Name          Instance     State           Peer          Peer
  ne0           master       Fully-connected p0            p1
```

show diameter network-element detail

```
user@host> show diameter network-element detail

Diameter network-element:
  NE name                : ne0
  Instance name          : master
  State                  : Fully-connected
  Primary peer           : p0
  Secondary peer         : p1
  Peers                  :   5
  Activated peers        :   4
  Open peers             :   2
  Routes                 :   1
  Invalid routes         :   0
  Activation delay       :     10000 ms
  First selection delay  :         0 ms
  Post selection delay   :     30000 ms
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

show diameter | 19472

show diameter function | 19482

show diameter network-element map | 19500

show diameter peer | 19504

show diameter route | 19519

show diameter network-element map

IN THIS SECTION

- Syntax | 19500
- Description | 19500
- Options | 19500
- Required Privilege Level | 19500
- Output Fields | 19501
- Sample Output | 19502
- Release Information | 19503

Syntax

```
show diameter network-element map
<brief | detail | summary>
<element-name>
```

Description

Display network-element-to-peer mapping information for all Diameter network elements or only the specified network element.

Options

brief | detail | summary	(Optional) Display the specified level of output. The `summary` output is displayed by default. The `brief` output and `detail` output display the summary information in a different format.
element-name	(Optional) Display information for only the specified network element.

Required Privilege Level

view

Output Fields

Table 458 on page 19501 lists the output fields for the `show diameter network-element map` command. Output fields are listed in the approximate order in which they appear.

Table 458: show diameter network-element map Output Fields

Field Name	Field Description	Level of Output
`Name`	Name of the Diameter network element.	`summary`
`Instance`	Name of the Diameter instance in which the network element is configured.	`summary`
`Peer`	Name of the peer.	All levels
`Priority`	Priority configured for the peer. A lower number indicates a higher priority.	All levels
`State`	State of the peer: • Activated—Peer has been activated (selected) by the network element. • Not-Activated—Peer has not been selected by the network element. • Primary—Peer that is connected to the network element and has the higher priority of the two connected peers. • Secondary—Peer that is connected to the network element and has the lower priority of the two connected peers.	`summary`
`NE name`	Name of the Diameter network element.	`brief detail`
`Instance name`	Name of the Diameter instance in which the network element is configured.	`brief detail`

Table 458: show diameter network-element map Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Usage	State of the peer: • Activated—Peer has been activated (selected) by the network element. • Not-Activated—Peer has not been selected by the network element. • Primary—Peer that is connected to the network element and has the higher priority of the two connected peers. • Secondary—Peer that is connected to the network element and has the lower priority of the two connected peers.	brief detail

Sample Output

show diameter network-element map

```
user@host> show diameter network-element map

Diameter network-element peers:
  Name         Instance     Peer         Priority  State
  ne0          master       p288               30 Activated
  ne0          master       p0                 20 Primary
  ne0          master       pA                 15 Activated
  ne0          master       p1                 10 Secondary
  ne0          master       pB                  5 Not-Activated
```

show diameter network-element map detail

```
user@host> show diameter network-element map detail

Diameter network-element peers:
  NE name               : ne0
  Instance name         : master
  Peer                  : p288
```

```
    Priority              :        30
    Usage                 : Activated

    NE name               : ne0
    Instance name         : master
    Peer                  : p0
    Priority              :        20
    Usage                 : Primary

    NE name               : ne0
    Instance name         : master
    Peer                  : pA
    Priority              :        15
    Usage                 : Activated

    NE name               : ne0
    Instance name         : master
    Peer                  : p1
    Priority              :        10
    Usage                 : Secondary

    NE name               : ne0
    Instance name         : master
    Peer                  : pB
    Priority              :         5
    Usage                 : Not-Activated
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

show diameter | 19472

show diameter network-element | 19495

show diameter peer

IN THIS SECTION

- Syntax | 19504
- Description | 19504
- Options | 19504
- Required Privilege Level | 19504
- Output Fields | 19505
- Sample Output | 19508
- Release Information | 19510

Syntax

```
show diameter peer
<brief | detail | summary>
<peer-name>
```

Description

Display information about all peers associated with Diameter instances or only the specified peer.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic peer information. The `brief` output displays the summary information in a different format. The `detail` output adds information to the `brief` output.

peer-name (Optional) Display information for only the specified peer.

Required Privilege Level

view

Output Fields

Table 459 on page 19505 lists the output fields for the `show diameter peer` command. Output fields are listed in the approximate order in which they appear.

Table 459: show diameter peer Output Fields

Field Name	Field Description	Level of Output
`Peer`	Name of the peer.	`brief summary`
`Instance`	Name of the Diameter instance in which the peer is configured.	`brief summary`
`State`	State of the peer: • Bad-Config—Misconfiguration. • Bad-Remote—Remote side does not conform to one of the decisions or is sending malformed messages. • Closed—Normal disconnect due to a request from the remote site or due to excessive watchdog timeouts. • Destructing—Peer to be deleted on the next timer tick; until then, it performs no actions. • Disabled—Peer is administratively disabled. • Internal-error—Internal error has been detected and the peer is in the process of restarting. • No-Activation—Peer is not used by any Diameter network element. • Rejected—Connection was rejected by remote side of the connection. • Reopen–Connection has been unexpectedly closed and Diameter is attempting to reopen the connection. • Suspended—All other reasons to be suspended.	All levels
`NE-Count`	Number of network elements associated with the peer.	`brief summary`

Table 459: show diameter peer Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Activated Count`	Activation status of the peer: • `1`—Peer is activated. • `0`—Peer is not activated.	All levels
`Primary Count`	Status of the peer, primary (`1`) or secondary (`0`).	All levels
`Secondary Count`	Secondary (`0`) versus Primary (`1`) status of the peer.	All levels
`Peer name`	Name of the peer.	`detail`
`NEs`	Number of network elements associated with the peer.	`detail`
`Vrf`	Logical system:routing instance of the configuration.	`detail`
`Remote address`	Remote IP address of the peer.	`detail`
`Remote port`	Remote port on the peer on which the connection is made.	`detail`
`Remote end origin realm`	Name of the realm of the Diameter node that originates messages to the peer.	`detail`
`Remote end origin host`	Name of the host of the Diameter node that originates messages to the peer.	`detail`
`Local address`	Local IP address on the Diameter origin node.	`detail`
`Local port`	Local port on the Diameter origin node.	`detail`
Local transport	Number of transports configured.	`detail`

Table 459: show diameter peer Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Time since last enable`	Period since peer was enabled in *hh:mm.ss* format.	`detail`
`In state time`	Period that peer has been in present state in *hh:mm.ss* format.	`detail`
`Remaining in state time`	Period that peer will remain in present state in *hh:mm.ss* format.	`detail`
`Missing wd events`	Number of missed watchdog events.	`detail`
`Tx queue length`	Number of messages in the transmit queue.	`detail`
`Answer waiting count`	Number of answers on which the peer is waiting.	`detail`
`Time since last rx`	Number of milliseconds since the last message was received by the peer.	`detail`
`Time until wd timeout`	Time remaining until next watchdog event.	`detail`
`Operation timeout`	Watchdog timeout period.	`detail`
`Suspended timeout base`	Base timeout period in suspended states (suspended, rejected, bad-remonte, bad-config). This timeout doubles after each consecutive suspension, until the maximum value of 600 seconds is reached.	`detail`
`Closed timeout`	Timeout period in normal closed state, such as when an external peer requested a disconnect.	`detail`
`Connection timeout`	Timeout period for establishing a connection.	`detail`

Table 459: show diameter peer Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Waiting origin state id`	Whether the peer is waiting for the Origin-State-Id AVP, `yes` or `no`.	`detail`

Sample Output

show diameter peer

```
user@host> show diameter peer

Diameter peer list:
                                                        Activated Primary Secondary
  Peer         Instance     State          NE-Count Count     Count   Count
  p0           master       I-Open           1        1         1       0
  p1           master       I-Open           1        1         0       1
  p288         master       Suspended        1        1         0       0
  pA           master       Suspended        1        1         0       0
  pB           master       No-Activation    1        0         0       0
  pc           master       No-Activation    0        0         0       0
  pd           master       No-Activation    0        0         0       0
```

show diameter peer detail

```
user@host> show diameter peer detail

Diameter peer:
  Peer name               : p0
  State                   : I-Open
  NEs                     :   1
  Activated count         :   1
  Primary count           :   1
  Secondary count         :   0
  Vrf                     : default:master
  Remote address          : 203.0.113.158
```

```
Remote port             : 62917
Remote end origin realm : rrrrA
Remote end origin host  : hhhhA
Local address           : 203.0.113.155
Local port              : 57095
Local transport         : <NO-TRANSPORT>
Time since last enable  : 08:56.200
In state time           : 08:56.200
Remaining in state time : no limit
Missed wd events        :         0
Tx queue length         :         0
Answer waiting count    :         0
Time since last rx      :      2200 ms
Time until wd timeout   :      3800 ms
Operation timeout       :      6000 ms
Suspended timeout base  :     30000 ms
Closed timeout          :     30000 ms
Connection timeout      :      6000 ms
Waiting origin state id : no

Peer name               : p1
State                   : I-Open
NEs                     :   1
Activated count         :   1
Primary count           :   0
Secondary count         :   1
Vrf                     : default:master
Remote address          : 203.0.113.158
Remote port             : 58490
Remote end origin realm : rrrrA
Remote end origin host  : hhhhB
Local address           : 203.0.113.155
Local port              : 49293
Local transport         : <NO-TRANSPORT>
Time since last enable  : 08:56.200
In state time           : 08:36.000
Remaining in state time : no limit
Missed wd events        :         0
Tx queue length         :         0
Answer waiting count    :         0
Time since last rx      :         0 ms
Time until wd timeout   :      6000 ms
Operation timeout       :      6000 ms
```

```
Suspended timeout base  :     30000 ms
Closed timeout          :     30000 ms
Connection timeout      :      6000 ms
Waiting origin state id : no
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- clear diameter peer | 3554
- show diameter | 19472
- show diameter peer map | 19510
- show diameter peer statistics | 19514

show diameter peer map

IN THIS SECTION

- Syntax | 19511
- Description | 19511
- Options | 19511
- Required Privilege Level | 19511
- Output Fields | 19511
- Sample Output | 19512
- Release Information | 19514

Syntax

```
show diameter peer map
<brief | detail | summary>
<peer-name>
```

Description

Display peer-to-network-element mapping information for all peers associated with Diameter instances or with the specified peer.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic peer information. The `brief` output displays the summary information in a different format. The `detail` output adds information to the `brief` output.

peer-name (Optional) Display mapping information for only the specified peer.

Required Privilege Level

view

Output Fields

Table 460 on page 19511 lists the output fields for the `show diameter peer map` command. Output fields are listed in the approximate order in which they appear.

Table 460: show diameter peer map Output Fields

Field Name	Field Description	Level of Output
`Peer`	Name of the peer.	All levels
`Instance`	Name of the Diameter instance in which the network element is configured.	All levels

Table 460: show diameter peer map Output Fields *(Continued)*

Field Name	Field Description	Level of Output
NE	Name of the Diameter network element.	All levels
Priority	Priority configured for the peer. A lower number indicates a higher priority.	All levels
State	State of the peer: • Activated—Peer has been activated (selected) by the network element. • Not-Activated—Peer has not been selected by the network element. • Primary—Peer that is connected to the network element and has the higher priority of the two connected peers. • Secondary—Peer that is connected to the network element and has the lower priority of the two connected peers.	All levels
Instance name	Name of the Diameter instance in which the network element is configured.	brief detail
NE name	Name of the Diameter network element.	brief detail
Usage	Role of the peer for the network element, Primary or Secondary.	brief detail

Sample Output

show diameter peer map

```
user@host> show diameter peer map

Diameter peer usage by network elements:
  Peer         Instance     NE              Priority State
  p0           master       ne0                   20 Primary
  p1           master       ne0                   10 Secondary
```

```
  p288         master       ne0                    30 Activated
  pA           master       ne0                    15 Activated
  pB           master       ne0                     5 Not-Activated
```

show diameter peer map detail

```
user@host> show diameter peer map detail

Diameter network-element peers:
  Peer                  : p0
  Instance name         : master
  NE name               : ne0
  Priority              :        20
  Usage                 : Primary

  Peer                  : p1
  Instance name         : master
  NE name               : ne0
  Priority              :        10
  Usage                 : Secondary

  Peer                  : p288
  Instance name         : master
  NE name               : ne0
  Priority              :        30
  Usage                 : Activated

  Peer                  : pA
  Instance name         : master
  NE name               : ne0
  Priority              :        15
  Usage                 : Activated

  Peer                  : pB
  Instance name         : master
  NE name               : ne0
  Priority              :         5
  Usage                 : Not-Activated
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

clear diameter peer | 3554
show diameter | 19472
show diameter peer | 19504
show diameter peer statistics | 19514

show diameter peer statistics

IN THIS SECTION

- Syntax | 19514
- Description | 19514
- Options | 19515
- Required Privilege Level | 19515
- Output Fields | 19515
- Sample Output | 19516
- Release Information | 19518

Syntax

```
show diameter peer statistics
<brief | detail | summary>
<peer-name>
```

Description

Display statistics about all peers associated with Diameter instances or only the specified peer.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic function statistics. The `brief` output displays the summary information in a different format and adds numbers accumulated since the peer was connected. The `detail` output adds information to the `brief` output.

peer-name (Optional) Display information for only the specified peer. When you specify a peer, the `brief` output is displayed by default, even when you explicitly specify `summary`.

Required Privilege Level

view

Output Fields

Table 461 on page 19515 lists the output fields for the `show diameter peer statistics` command. Output fields are listed in the approximate order in which they appear.

Table 461: show diameter peer statistics Output Fields

Field Name	Field Description	Level of Output
`Peer`	Name of the peer.	`summary brief`
`Instance`	Name of the Diameter instance in which the network element is configured.	`summary brief`
`Rx`	Total number of messages received.	`summary brief`
`Rx-Peer`	Number of messages received by the peer.	`summary brief`
`Rx-node`	Number of messages received by the Diameter node.	`summary brief`
`Forw`	Total number of forwarded messages.	`summary brief`
`Tx-Peer`	Number of messages transmitted by the peer.	`summary brief`

Table 461: show diameter peer statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Tx	Total number of transmitted messages.	summary brief
Peer name	Name of the peer.	detail
Instance name	Name of the Diameter instance in which the network element is configured.	detail

Sample Output

show diameter peer statistics

```
user@host> show diameter peer statistics

Diameter peer statistics:
  Peer         Instance     Rx        Rx-Peer   Rx-Node   Forw      Tx-Peer   Tx
  p0           master            113       113         0         0       113       113
  p1           master            110       110         0         0       110       110
  p288         master              0         0         0         0         0         0
  pA           master              0         0         0         0         0         0
  pB           master              0         0         0         0         0         0
  pc           master              0         0         0         0         0         0
  pd           master              0         0         0         0         0         0
```

show diameter peer statistics detail

```
user@host> show diameter peer statistics detail

Diameter peer statistics:
  Peer name               : p0
  Instance name           : master
                            Current       Since last enable
  Rx errors               :         0             0
```

```
Rx messages             :       114           114
Rx handled by peer      :       114           114
Rx dropped msgs         :         0             0
Rx unmatched answers    :         0             0
Rx answers              :         0             0
Rx requests             :         0             0
Rx total                :         0             0
Forw to connection      :         0             0
Forw to peer            :         0             0
Forw to routed dest     :         0             0
Total forwarding        :         0             0
Forwarding failures     :         0             0
Forwarding success      :         0             0
Moved-in messages       :         0             0
Moved-out messages      :         0             0
Rerouted messages       :         0             0
Dropped tx messages     :         0             0
Tx by peer              :       114           114
Tx errors               :         0             0
Tx total                :       114           114
Connection attempts     :         0             1
Connection fails        :         0             0
Connections             :         0             1
Passive teminations     :         0             0
Active terminations     :         0             0
Passive disconnects     :         0             0
Active disconnects      :         0             0
Rx block requests       :         0             0
Rx block timeoutss      :         0             0
Connection management messages
                          Rx current    Rx since      Tx current    Tx since
                                        last enable                 last enable
CER                     :         0             0             1             1
CEA                     :         1             1             0             0
DWR                     :         0             0           113           113
DWA                     :       113           113             0             0
DPR                     :         0             0             0             0
DPA                     :         0             0             0             0

Peer name               : p1
Instance name           : master
                          Current       Since last enable
Rx errors               :         0             0
```

```
    Rx messages             :       110           110
    Rx handled by peer      :       110           110
    Rx dropped msgs         :         0             0
    Rx unmatched answers    :         0             0
    Rx answers              :         0             0
    Rx requests             :         0             0
    Rx total                :         0             0
    Forw to connection      :         0             0
    Forw to peer            :         0             0
    Forw to routed dest     :         0             0
    Total forwarding        :         0             0
    Forwarding failures     :         0             0
    Forwarding success      :         0             0
    Moved-in messages       :         0             0
    Moved-out messages      :         0             0
    Rerouted messages       :         0             0
    Dropped tx messages     :         0             0
    Tx by peer              :       110           110
    Tx errors               :         0             0
    Tx total                :       110           110
    Connection attempts     :         0             1
    Connection fails        :         0             0
    Connections             :         0             1
    Passive teminations     :         0             0
    Active terminations     :         0             0
    Passive disconnects     :         0             0
    Active disconnects      :         0             0
    Rx block requests       :         0             0
    Rx block timeoutss      :         0             0
    Connection management messages
                              Rx current    Rx since      Tx current    Tx since
                                            last enable                 last enable
    CER                     :         0             0             1             1
    CEA                     :         1             1             0             0
    DWR                     :         0             0           109           109
    DWA                     :       109           109             0             0
    DPR                     :         0             0             0             0
    DPA                     :         0             0             0             0
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

clear diameter peer | 3554
show diameter | 19472
show diameter peer | 19504
show diameter peer map | 19510

show diameter route

IN THIS SECTION

- Syntax | 19519
- Description | 19519
- Options | 19520
- Required Privilege Level | 19520
- Output Fields | 19520
- Sample Output | 19521
- Release Information | 19522

Syntax

```
show diameter route
<brief | detail | summary>
<route-name>
```

Description

Display information about all routes associated with Diameter instances or only the specified route.

Options

brief | detail | summary (Optional) Display the specified level of output. The `summary` output is displayed by default and includes basic function information. The `brief` output displays the summary information in a different format. The `detail` output adds information to the `brief` output.

route-name (Optional) Display information for only the specified route.

Required Privilege Level

view

Output Fields

Table 462 on page 19520 lists the output fields for the `show diameter route` command. Output fields are listed in the approximate order in which they appear.

Table 462: show diameter route Output Fields

Field Name	Field Description	Level of Output
`Route`	Name of the route.	`summary brief`
`NE`	Name of the network element associated with the route.	`summary brief`
`Instance`	Name of the Diameter instance in which the route is configured.	`summary brief`
`NE name`	Name of the network element associated with the route.	`brief detail`
`Instance name`	Name of the Diameter instance in which the route is configured.	`brief detail`
`Valid`	Determination whether the route is valid, `yes` or `no`.	All levels
`Up`	State of the route, `yes` (up) or `no` (down).	All levels

Table 462: show diameter route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Function	Name of the function associated with the route.	brief detail
Partition	Partition associated with the function.	brief detail
Dest-realm	Destination realm configured for the route.	brief detail
Dest-host	Destination hostname configured for the route.	brief detail
Metric	Metric associated with the destination and function to create the route.	brief detail
Score	Value that represents how a route is configured. The basic score is 0. Points are added according to the following scheme: • Function is specified—Add 3. • Function partition is specified—Add 1. • Destination realm is specified—Add 1. • Destination host is specified—Add 1.	brief detail

Sample Output

show diameter route

```
user@host> show diameter route

Diameter routes:
  Route        NE           Instance     Valid Up
  rA           ne0          master       yes   yes
```

show diameter route detail

```
user@host> show diameter route detail

Diameter route:
  Route name        : rA
  NE name           : ne0
  Instance mame     : master
  Valid             : yes
  Up                : yes
  Function          : jsrc
  Partition         : jsrc-a
  Dest-realm        : outer-realm
  Dest-host         : outer-host
  Metric            :        50
  Score             :         6
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

show diameter | 19472

show diameter network-element | 19495

show | display

IN THIS SECTION

- Syntax | 19523
- Description | 19523
- Options | 19523
- Required Privilege Level | 19523

Sample Output | 19524
Release Information | 19526

Syntax

```
show | display <changed | commit-scripts |detail | inheritance | json | mark-changed | max-depth
depth | max-version version | omit | set | translation-scripts | xml>
```

Description

Display the configuration as a series of configuration mode commands required to re-create the configuration from the top level of the hierarchy.

Options

changed	Tag changes with junos:changed attribute (XML only).
commit-scripts	Show data after commit scripts have been applied.
detail	Show configuration data detail.
inheritance	Show inherited configuration data and source group.
json	Show output in JSON format.
mark-changed	Tag changes with junos:mark-changed attribute (XML only).
max-depth *depth*	Specify the maximum value for the depth of configuration data to be displayed.
max-version *version*	Specify the maximum value for the version of the configuration data to be displayed.
omit	Emit configuration statements with the `omit` option.
set	Show `set` commands that create configuration.
translation-scripts	Show data after translation scripts have been applied.
xml	Show output as XML tags.

Required Privilege Level

view

Sample Output

show | display max-depth

```
user@host# show interfaces | display max-depth 3
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.1.12.1/24;
        }
    }
}
```

```
user@host# show interfaces | display max-depth 2
ge-0/0/0 {
    unit 0 {
        family inet { ... }
    }
}
```

```
user@host# show interfaces | display max-depth 1
ge-0/0/0 {
    unit 0 { ... }
}
```

show configuration | display detail

```
user@host> show configuration | display detail
## Last commit: 2023-04-13 02:04:36 PDT by regress
##
## version: Software version information
## YANG module: junos-conf-root@2023-01-01.yang
## require: system
##
version 23.1R1.8;
##
## groups: Configuration groups
```

```
## YANG module: junos-conf-root@2023-01-01.yang
## require: configuration
##
groups {
    ##
    ## Group name
    ## range: 1 .. 254
    ##
    global {
        ##
        ## system: System parameters
        ## YANG module: junos-conf-system@2023-01-01.yang
        ## require: admin system
        ## Daemon notify list: : sdk-vmmd
        ##
        system {
            ##
            ## root-authentication: Authentication information for the root login
            ## YANG module: junos-conf-system@2023-01-01.yang
            ## require: admin
            ## Daemon notify list: : sdk-vmmd spmd
            ## package: junos-base
            ##
            root-authentication {
                ##
                ## encrypted-password: Encrypted password string
                ## YANG module: junos-conf-system@2023-01-01.yang
                ## range: 1 .. 128
                ##
                encrypted-password "$1$ZUlES4dp$OUwWo1g7cLoV/aMWpHUnC/"; ## SECRET-DATA
            }
            ##
            ## login: Names, login classes, and passwords for users
            ## YANG module: junos-conf-system@2023-01-01.yang
            ## require: admin
            ## Daemon notify list: : fsad spmd
            ## package: junos-base
            ##
            login {
                ##
                ## Login class name
                ## match: Must be a string of alphanumericals, dashes or underscores
                ## match (regex): ^[[:alnum:]_-]+$
```

```
                ## package: junos-base
                ##
                class wheel {
                    ##
                    ## permissions: Set of permitted operation categories
                    ## YANG module: junos-conf-system@2023-01-01.yang
                    ##
                    permissions [ admin clear field floppy interface maintenance network reset
routing shell snmp system trace view ];
                }
                ##
                ## Login class name
                ## match: Must be a string of alphanumericals, dashes or underscores
                ## match (regex): ^[[:alnum:]_-]+$
                ## package: junos-base
                ##
                class readonly {
                    ##
                    ## permissions: Set of permitted operation categories
                    ## YANG module: junos-conf-system@2023-01-01.yang
                    ##
                    permissions [ interface network routing system trace view ];
                }
                ##
                ## User name (login)
                ## match: Must contain characters (alphanumerics, underscores or hyphens)
beginning with an alphanumeric or an underscore character.
                ## match (regex): ^[[:alnum:]_]{1,}[.]{0,1}[[:alnum:]_-]{0,}[.]{0,1}[[:alnum:]_-]
{0,}$
                ## package: junos-base
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show | 16994

show | display commit-scripts

IN THIS SECTION

- Syntax | 19527
- Description | 19527
- Options | 19527
- Required Privilege Level | 19527
- Sample Output | 19528
- Release Information | 19529

Syntax

```
show | display commit-scripts
<no-transients | view>
```

Description

Show data after commit scripts have been applied. Commit scripts contain instructions that enforce custom configuration rules and are invoked during the commit process before the standard Junos OS validity checks are performed.

Options

no-transients Do not display transient changes in the configuration.

view Display commit-script-style xml data. Display XML format of the post-inheritance candidate configuration in the CLI.

Required Privilege Level

view

Sample Output

show system login | display commit-scripts view

```
user@host# show system login | display commit-scripts view<rpc-reply xmlns:junos="http://
xml.juniper.net/junos/23.1R0/junos">
    <configuration junos:changed-seconds="1682095681" junos:changed-localtime="2023-04-21
09:48:01 PDT" junos:changed="changed">
        <system junos:group="global">
            <login junos:group="global">
                <class junos:group="global">
                    <name junos:group="global">wheel</name>
                    <permissions junos:group="global">admin</permissions>
                    <permissions junos:group="global">clear</permissions>
                    <permissions junos:group="global">field</permissions>
                    <permissions junos:group="global">floppy</permissions>
                    <permissions junos:group="global">interface</permissions>
                    <permissions junos:group="global">maintenance</permissions>
                    <permissions junos:group="global">network</permissions>
                    <permissions junos:group="global">reset</permissions>
                    <permissions junos:group="global">routing</permissions>
                    <permissions junos:group="global">shell</permissions>
                    <permissions junos:group="global">snmp</permissions>
                    <permissions junos:group="global">system</permissions>
                    <permissions junos:group="global">trace</permissions>
                    <permissions junos:group="global">view</permissions>
                </class>
                <class junos:group="global">
                    <name junos:group="global">readonly</name>
                    <permissions junos:group="global">interface</permissions>
                    <permissions junos:group="global">network</permissions>
                    <permissions junos:group="global">routing</permissions>
                    <permissions junos:group="global">system</permissions>
                    <permissions junos:group="global">trace</permissions>
                    <permissions junos:group="global">view</permissions>
                </class>
                <user junos:group="global">
                    <name junos:group="global">regress</name>
                    <uid junos:group="global">928</uid>
                    <class junos:group="global">superuser</class>
                    <undocumented><shell junos:group="global">csh</shell></undocumented>
                    <authentication junos:group="global">
```

```
                        <encrypted-password junos:group="global">
$1$kPU..$w.4FGRAGanJ8U4Yq6sbj7.</encrypted-password>
                    </authentication>
                </user>
                <user junos:group="global">
                    <name junos:group="global">remote</name>
                    <uid junos:group="global">9999</uid>
                    <class junos:group="global">superuser</class>
                    <undocumented><shell junos:group="global">csh</shell></undocumented>
                </user>
            </login>
        </system>
    </configuration>
    <cli>
        <banner>[edit]</banner>
    </cli>
</rpc-reply>
```

Release Information

Command introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

How Commit Scripts Work

show | display inheritance

IN THIS SECTION

- Syntax | 19530
- Description | 19530
- Options | 19530
- Required Privilege Level | 19530
- Output Fields | 19531

Sample Output | 19531
Release Information | 19532

Syntax

```
show | display inheritance
<brief | no-comments | terse>
<defaults>
<when>
```

Description

Show the inherited configuration data and information about the source group from which the configuration has been inherited. Show interface ranges configuration data in expanded format and information about the source interface-range from which the configuration has been expanded.

Options

brief Display brief output for the command.

defaults Display the defaults that have been applied to the configuration.

no-comments Display configuration information without in-line comments marked with ##.

terse Display terse output with inheritance details as an in-line comment.

when Display inherited data for a specified parameter that is defined as conditions under the configuration "groups" on page 7916

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

show system login class readonly | display inheritance

```
user@host#show system login class readonly | display inheritance

##
## 'interface' was inherited from group 'global'
## 'network' was inherited from group 'global'
## 'routing' was inherited from group 'global'
## 'system' was inherited from group 'global'
## 'trace' was inherited from group 'global'
## 'view' was inherited from group 'global'
##
permissions [ interface network routing system trace view ];
```

show system login class readonly | display inheritance brief

```
user@host#show system login class readonly | display inheritance brief

## 'interface' was inherited from group 'global'
## 'network' was inherited from group 'global'
## 'routing' was inherited from group 'global'
## 'system' was inherited from group 'global'
## 'trace' was inherited from group 'global'
## 'view' was inherited from group 'global'
permissions [ interface network routing system trace view ];
```

show system ports | display inheritance defaults

```
user@host# show system ports | display inheritance defaults

## 'console' was inherited from group 'junos-defaults'
```

```
## 'vt100' was inherited from group 'junos-defaults'
## console type vt100;
```

show system login class readonly | display inheritance no-comments

```
user@host# show system login class readonly | display inheritance no-comments
permissions [ interface network routing system trace view ];
```

show system login class readonly | display inheritance terse

```
user@host# show system login class readonly | display inheritance terse
permissions [ interface network routing system trace view ]; ## inherited from group 'global';
inherited from group 'global'; inherited from group 'global'; inherited from group 'global';
inherited from group 'global'; inherited from group 'global'
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Using the junos-defaults Configuration Group

show | display json

IN THIS SECTION

- Syntax | 19533
- Description | 19533
- Options | 19533
- Required Privilege Level | 19533
- Sample Output | 19533

Release Information | 19535

Syntax

```
show | display json
<groups | interface-ranges | no-export-path>
```

Description

Show output in JavaScript Object Notation (JSON) format. Devices running Junos OS support JSON representation of the operational command output and the Junos OS configuration hierarchy. On the Junos OS CLI, to display the command output or configuration in JSON instead of in the default formatted ASCII text, append the | display json option to the command.

Options

groups	Tag inherited data with the source group name.
interface-ranges	Tag inherited data with the source interface-range name
no-export-path	Do not export parent path in the data.

Required Privilege Level

view

Sample Output

show chassis hardware | display json

```
user@host# show chassis hardware | display json
{
    "chassis-inventory" : [
    {
        "attributes" : {
            "xmlns" : "http://xml.juniper.net/junos/16.1R1/junos-chassis"
        },
```

```
        "chassis" : [
        {
            "attributes" : {"junos:style" : "inventory"},
            "name" : [
            {
                "data" : "Chassis"
            }
            ],
            "serial-number" : [
            {
                "data" : "serial-number"
            }
            ],
            "description" : [
            {
                "data" : "MX80-48T"
            }
            ],
            "chassis-module" : [
            {
                "name" : [
                {
                    "data" : "Midplane"
                }
                ],
                "version" : [
                {
                    "data" : "REV 11"
                }
                ],
                "part-number" : [
                {
                    "data" : "711-031603"
                }
                ],
                "serial-number" : [
                {
                    "data" : "serial-number"
                }
                ],
                "description" : [
                {
                    "data" : "MX80-48T"
```

```
                }
                ],
                "clei-code" : [
                {
                    "data" : "CMMAX10BRD"
                }
                ],
                "model-number" : [
                {
                    "data" : "CHAS-MX80-48T-S"
                }
                ]
            },
            //* additional JSON objects *//
            ]
        }
        ]
    }
    ]
}
```

Release Information

Command introduced in Junos OS Release 14.2R1

RELATED DOCUMENTATION

Map Junos OS Command Output to JSON in the CLI

show | display omit

IN THIS SECTION

- Syntax | **19536**
- Description | **19536**
- Required Privilege Level | **19536**

Release Information | 19536

Syntax

```
show | display omit
```

Description

Display configuration statements (including those marked as hidden by the `apply-flags omit` configuration statement).

```
user@host# show | display omit
    system {
        apply-flags omit;
        login {
            message lengthy-login-message;
        }
    }
```

Following is an example that shows how to set omit:

```
user@host#set system apply-flags omit
    [edit]
user@host# commit
     commit complete
```

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

show | 16994

show | display rfc5952

IN THIS SECTION

- Syntax | 19537
- Description | 19537
- Required Privilege Level | 19537
- Sample Output | 19537
- Release Information | 19538

Syntax

```
show | display rfc5952
```

Description

Display IPv6 addresses as per RFC 5952 specifications. For RFC information, go to: http://tools.ietf.org/html/rfc5952.

Required Privilege Level

View

Sample Output

command-name

```
user@host> show configuration interfaces ge-0/0/0 |display rfc5952
unit 0 {
```

```
    family inet6 {
        address 2012::2:1/64;
        address 2001::1111:2222:0:0:1/96;
    }
}
```

Release Information

Command introduced before Junos OS Release 11.4.

RELATED DOCUMENTATION

show | 16994

show | display set

IN THIS SECTION

- Syntax | 19538
- Description | 19539
- Options | 19539
- Required Privilege Level | 19539
- Sample Output | 19539
- Release Information | 19540

Syntax

```
show | display set <explicit> <relative>
```

Description

Display the configuration as a series of configuration mode commands required to re-create the configuration from the top level of the hierarchy as `set` commands.

Options

explicit Display explicitly, as a series of set commands, all the configuration statements that the system internally creates when you configure certain statements from the top level of the hierarchy.

For example, assume you issue the `set interfaces ge-0/0/0.0 family inet` configuration mode command. You then show the resulting configuration with the `show interfaces ge-0/0/0 | display set` command. The output displays the same set command you entered. If you include the `explicit` argument, the output also shows the configuration statements needed to create the hierarchy where the `family inet` statement is specified. Specifically for this example, the output therefore includes the `set interfaces ge-0/0/0 unit 0` statement in addition to the `set interfaces ge-0/0/0.0 family inet` statement.

relative Display a series of set commands relative to the current edit path.

For example, assume you issue the `set interfaces ge-0/0/0.0 family inet` configuration mode command. If you include the `relative` argument with the `show interfaces ge-0/0/0 | display set` command, the output shows the configuration statements relative to the `interfaces` path. Specifically for this example, the output displays the `set unit 0 family inet` statement and `set unit 0 family inet address 10.0.1.254/24` statement in the output.

Required Privilege Level

view

Sample Output

Sample output for the `show | display set` and `show | display set <explicit>` commands:

```
[edit interfaces ge-0/0/0]
user@host# show | display set
set interfaces ge-0/0/0 unit 0 family inet address 10.0.1.254/24

[edit interfaces ge-0/0/0]
user@host# show | display set explicit
set interfaces ge-0/0/0 unit 0 family inet address 10.0.1.254/24
```

```
set interfaces ge-0/0/0 unit 0 family inet
set interfaces ge-0/0/0 unit 0

[edit interfaces ge-0/0/0]
user@host# show | display set relative
set unit 0 family inet address 10.0.1.254/24
set unit 0 family inet
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `explicit` command introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

show | 16994

Displaying set Commands from the Device Configuration

show | display set relative

IN THIS SECTION

- Syntax | 19541
- Description | 19541
- Options | 19541
- Required Privilege Level | 19541
- Sample Output | 19541
- Release Information | 19542

Syntax

```
show | display set relative <explicit>
```

Description

Display the configuration as a series of configuration mode commands required to re-create the configuration from the current hierarchy level.

Options

explicit Display explicitly, as a series of commands, all the configurations that the system internally creates when you configure certain statements from the current hierarchy level.

Required Privilege Level

view

Sample Output

Sample output for the `show | display set relative <explicit>` command:

```
[edit interfaces fe-0/0/0]
user@host# show
unit 0 {
family inet {
address 192.107.1.230/24;
}
family iso;
family mpls;
}
inactive: unit 1 {
family inet {
address 10.0.0.1/8;
}
}
user@host# show | display set relative
set unit 0 family inet address 192.107.1.230/24
set unit 0 family iso
```

```
set unit 0 family mpls
set unit 1 family inet address 10.0.0.1/8
deactivate unit 1

user@host# show | display set relative | explicit
set unit 0 family inet address 192.168.1.230/24
set unit 0 family inet
set unit 0 family iso
set unit 0 family mpls
set unit 0
set unit 1 family inet address 10.0.0.1/8
set unit 1 family inet
deactivate unit 1
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `explicit` command introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

Displaying set Commands from the Device Configuration

show | display translation-scripts

IN THIS SECTION

- Syntax | 19543
- Description | 19543
- Options | 19543
- Required Privilege Level | 19543
- Sample Output | 19543
- Release Information | 19545

Syntax

```
show | display tanslation-scripts
<configured-delta | translated-config | translated-delta>
```

Description

Show data after translation scripts have been applied. Translation scripts convert the configuration data corresponding to the nonnative data models into Junos OS syntax and add the translated configuration data as a transient change in the checkout configuration during the commit operation.

Options

configured-delta	Display the input configuration to the translation script.
translated-config	Display the translated configuration.
translated-delta	Display the output configuration from the translation script.

Required Privilege Level

view

Sample Output

show display | translation-scripts

The following truncated output displays the complete post-inheritance configuration and includes the nonnative configuration data as well as the translation of that data.

```
user@host>show display | translation-scripts
## Last changed: 2016-05-13 16:37:42 PDT
  version "16.1R1;
  system {
      host-name host;
      domain-name example.com;
      ...
      /* Translated data */
      scripts {
          op {
```

```
            file test.slax;
        }
    }
    ...
}
...
/* Nonnative configuration data */
myconfig:myscript {
    op {
        filename test.slax;
    }
}
```

show display | translation-scripts translated-config

```
user@host>show | display translation-scripts translated-config
 system {
     scripts {
         op {
             file test.slax;
         }
     }
 }
```

show | display translation-scripts configured-delta

```
user@host>show | display translation-scripts configured-delta
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/16.1R1/junos">
      <configuration operation="delete">
      </configuration>
      <configuration operation="create">
          <myscript xmlns="http://jnpr.net/yang/myscript" operation="create">
              <op>
                  <filename>test2.slax</filename>
              </op>
          </myscript>
      </configuration>
      <cli>
          <banner>[edit]</banner>
```

```
      </cli>
  </rpc-reply>
```

show | display translation-scripts translated-delta

```
user@host>show | display translation-scripts translated-delta
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/R1/junos">
      <configuration xmlns:junos="http://xml.juniper.net/junos/*/junos">
          <system>
                <scripts>
                  <op>
                    <file>
                      <name>test2.slax</name>
                    </file>
                  </op>
                </scripts>
          </system>
      </configuration>
      <!-- EOF -->
      <cli>
          <banner>[edit]</banner>
      </cli>
  </rpc-reply>
```

Release Information

Command introduced in Junos OS Release 16.1R1

RELATED DOCUMENTATION

> *Committing and Displaying Configuration Data for Nonnative YANG Modules*

show | display xml

IN THIS SECTION

- Syntax | 19546
- Description | 19546
- Options | 19546
- Required Privilege Level | 19546
- Sample Output | 19547
- Release Information | 19548

Syntax

```
show | display xml <groups | interface-ranges | no-export-path>
```

Description

Show output as XML tags. On the Junos OS CLI, to display command output as Junos XML tag elements instead of as the default formatted ASCII text, include the `| display xml` option after the command.

Options

groups	Tag inherited data with the source group name.
interface-ranges	Tag inherited data with the source interface-range name
no-export-path	Do not export parent path in the data.

Required Privilege Level

view

Sample Output

show interfaces | display xml

```
user@host> show interfaces | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/23.1R0/junos">
<rpc-reply xmlns:junos=http://xml.juniper.net/junos/21.1R0/junos>
    <configuration junos:changed-seconds="1682580061" junos:changed-localtime="2023-04-27
07:21:01 UTC">
            <interfaces>
                <interface>
                    <name>ge-0/0/0</name>
                    <unit>
                        <name>0</name>
                        <family>
                            <inet>
                                <address>
                                    <name>10.1.12.1/24</name>
                                </address>
                            </inet>
                        </family>
                    </unit>
                </interface>

            </interfaces>
    </configuration>
    <cli>
        <banner>[edit]</banner>
    </cli>
</rpc-reply>
```

show interfaces | display xml no-export-path

```
user@host>show interfaces | display xml no-export-path
<rpc-reply xmlns:junos=http://xml.juniper.net/junos/21.1R0/junos>
    <interfaces>
            <interface>
                <name>ge-0/0/0</name>
                <unit>
```

```
                    <name>0</name>
                    <family>
                        <inet>
                            <address>
                                <name>10.1.12.1/24</name>
                            </address>
                        </inet>
                    </family>
                </unit>
            </interface>
    </interfaces>
    <cli>
        <banner>[edit]</banner>
    </cli>
</rpc-reply>
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Map Junos OS Commands and Command Output to Junos XML Tag Elements

show dot1x

IN THIS SECTION

- Syntax | 19549
- Description | 19549
- Options | 19549
- Required Privilege Level | 19549
- Output Fields | 19549
- Sample Output | 19556
- Release Information | 19557

Syntax

```
show dot1x
<brief | detail>
<interface interface-name>
```

Description

Display the current operational state of all ports with the list of connected users.

This command displays the list of connected supplicants received from the RADIUS authentication server regardless of the session state—that is, for both authenticated supplicants and for supplicants that attempted authentication.

Options

none	Display information for all authenticator ports.
brief | detail	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display information for the specified port with a list of connected supplicants.

Required Privilege Level

view

Output Fields

Table 463 on page 19549 lists the output fields for the `show dot1x` command. Output fields are listed in the approximate order in which they appear.

Table 463: show dot1x Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of a port.	All levels

Table 463: show dot1x Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`MAC address`	The MAC address of the connected supplicant on the port.	All levels
`Role`	The 802.1X authentication role of the interface. When 802.1X is enabled on an interface, the role is **Authenticator**. As **Authenticator**, the interface blocks LAN access until a supplicant is authenticated through 802.1X or MAC RADIUS authentication.	`brief`, `detail`
`State`	The state of the port: • **Authenticated**—The supplicant has been authenticated through the RADIUS server or has been permitted access through server fail fallback. • **Authenticating**—The supplicant is authenticating through the RADIUS server. • **Held**—An action has been triggered through server fail fallback during a RADIUS server timeout. A supplicant is denied access, permitted access through a specified VLAN, or maintains the authenticated state granted to it before the RADIUS server timeout occurred.	`brief`
`User`	The username of the connected supplicant.	`brief`
`Administrative state`	The administrative state of the port: • **auto**—Traffic is allowed through the port based on the authentication result (by default). • **force-authorize**—All traffic flows through the port irrespective of the authentication result. This state is not allowed on an interface whose VLAN membership has been set to **dynamic**. • **force-unauthorize**—All traffic drops on the port irrespective of the authentication result. This state is not allowed on an interface whose VLAN membership has been set to **dynamic**.	`detail`

Table 463: show dot1x Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Supplicant`	The mode for the supplicant: • **single**—Only the first supplicant is authenticated. All other supplicants who connect later to the port are allowed full access without any further authentication. They effectively "piggyback" on the first supplicant's authentication. • **single-secure**—Only one supplicant is allowed to connect to the port. No other supplicant is allowed to connect until the first supplicant logs out. • **multiple**—Multiple supplicants are allowed to connect to the port. Each supplicant is authenticated individually.	`detail`
`Quiet period`	The number of seconds the port waits following a failed authentication exchange with the supplicant before reattempting the authentication. The default value is 60 seconds. The range is 0 through 65,535 seconds.	`detail`
`Transmit period`	The number of seconds the port waits before retransmitting the initial EAPOL PDUs to the supplicant. The default value is 30 seconds. The range is 1 through 65,535 seconds.	`detail`
`MAC radius`	MAC RADIUS authentication: • **enabled**—The switch sends an EAPOL request to the connecting host to attempt 802.1X authentication and if the connecting host is unresponsive, the switch tries to authenticate the host by using the MAC address. • **disabled**—The default. The switch does not attempt to authenticate the MAC address of the connecting host.	`detail`

Table 463: show dot1x Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`MAC radius authentication protocol`	MAC RADIUS authentication protocol: • **EAP-MD5**—The EAP-MD5 protocol is used for MAC RADIUS authentication. EAP-MD5 is an authentication method belonging to the Extensible Authentication Protocol (EAP) authentication framework. EAP-MD5 is the default authentication protocol. • **PAP**—The Password Authentication Protocol (PAP) authentication protocol is used for MAC RADIUS authentication.	`detail`
`MAC radius restrict`	The authentication method is restricted to MAC RADIUS only. 802.1X authentication is not enabled.	`detail`
`Reauthentication`	The reauthentication state: • **disable**—Periodic reauthentication of the client is disabled. • **interval**—Sets the periodic reauthentication time interval. The default value is 3600 seconds. The range is 1 through 65,535 seconds.	`detail`
`Supplicant timeout`	The number of seconds the port waits for a response when relaying a request from the authentication server to the supplicant before resending the request. The default value is 30 seconds. The range is 1 through 60 seconds.	`detail`
`Server timeout`	The number of seconds the port waits for a reply when relaying a response from the supplicant to the authentication server before timing out. The default value is 30 seconds. The range is 1 through 60 seconds.	`detail`
`Maximum EAPOL requests`	The maximum number of times an EAPOL request packet is retransmitted to the supplicant before the authentication session times out. The default value is 2. The range is 1 through 10.	`detail`

Table 463: show dot1x Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Number of clients bypassed because of authentication`	The number of non-802.1X clients granted access to the LAN by means of static MAC bypass. The following fields are displayed: • **Client**—MAC address of the client. • **vlan** —The name of the VLAN to which the client is connected.	`detail`
`Guest VLAN member`	The VLAN to which a supplicant is connected when the supplicant is authenticated using a guest VLAN. If a guest VLAN is not configured on the interface, this field displays <**not configured**>.	`detail`
`Multi domain data session count`	The number of data sessions that have been authenticated on a multi-domain authentication interface.	`detail`
`Number of connected supplicants`	The number of supplicants connected to a port.	`detail`
`Supplicant`	The username and MAC address of the connected supplicant.	`detail`

Table 463: show dot1x Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Authentication method`	The authentication method used for a supplicant:
	• `CWA Authentication`—A supplicant is authenticated by the central Web authentication (CWA) server.
	• `Fail`—Authentication failed and supplicant is in `Held` state.
	• `Guest VLAN`—A supplicant is connected to the LAN through the guest VLAN.
	• `MAC RADIUS`—A nonresponsive host is authenticated based on its MAC address. The MAC address is configured as permitted on the RADIUS server, the RADIUS server lets the switch know that the MAC address is a permitted address, and the switch opens LAN access to the nonresponsive host on the interface to which it is connected.
	• `RADIUS`—A supplicant is configured on the RADIUS server, the RADIUS server communicates this to the switch, and the switch opens LAN access on the interface to which the supplicant is connected.
	• `Server-fail`—One of the following fallback actions is in effect because the RADIUS server is unreachable. Indicates whether EAPOL block is in effect, and the amount of time remaining for EAPOL block (in seconds).
	  • `deny`—The supplicant is denied access to the LAN, preventing traffic from flowing from the supplicant through the interface. This is the default server fail fallback action.
	  • `permit`—The supplicant is permitted access to the LAN as if the supplicant had been successfully authenticated by the RADIUS server.
	  • `use-cache`—In the event that the RADIUS server times out when the supplicant is attempting reauthentication, the supplicant is reauthenticated only if it was previously authenticated; otherwise, the supplicant is denied LAN access.
		`detail`

Table 463: show dot1x Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `VLAN`—The supplicant is configured to be moved to a specified VLAN if the RADIUS server is unavailable to reauthenticate the supplicant. (The VLAN must already exist on the switch.) • `Server-reject VLAN`—The supplicant received a RADIUS access-reject message from the authentication server and was moved to a server-reject VLAN, a specified VLAN already configured on the switch.	
`Authenticated VLAN`	The VLAN to which the supplicant is connected.	`detail`
`Dynamic filter`	User policy filter sent by the RADIUS server.	`detail`
`Session Reauth interval`	The configured reauthentication interval.	`detail`
`Reauthentication due in`	The number of seconds in which reauthentication will occur again for the connected supplicant.	`detail`
`Session Accounting Interim Interval`	The number of seconds between interim RADIUS accounting messages.	`detail`
`Accounting Update due in`	The number of seconds until the next interim RADIUS accounting update is due.	`detail`
`CWA Redirect URL`	The URL used to redirect the supplicant to a central Web server for authentication.	`detail`
`Eapol Block`	Shows whether EAPOL block is in effect or not in effect.	`detail`

Sample Output

show dot1x interface brief

```
user@switch> show dot1x interface brief
802.1X Information:
Interface     Role           State           MAC address          User
ge-0/0/1      Authenticator  Authenticated   00:a0:d2:18:1a:c8    user1
ge-0/0/2      Authenticator  Connecting
ge-0/0/3      Authenticator  Held            00:a6:55:f2:94:ae    user3
```

show dot1x interface detail

```
user@switch> show dot1x interface ge-0/0/16.0 detail

 ge-0/0/16.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Disabled
  Mac Radius Authentication Protocol: PAP
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 2
    Supplicant: abc, 00:30:48:8C:66:BD
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: v200
      Session Reauth interval: 3600 seconds
      Reauthentication due in 3587 seconds
      Eapol-Block: Not In Effect
   Supplicant: 000303030303, 00:03:03:03:03:03
```

```
        Operational state: Authenticated
        Backend Authentication state: Idle
        Authentication method: Mac Radius
        Authenticated VLAN: dyn_vlan2
        Session Reauth interval: 3600 seconds
        Reauthentication due in 3587 seconds
        Eapol-Block: In Effect
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

clear dot1x | 3564

Example: Setting Up 802.1X for Single-Supplicant or Multiple-Supplicant Configurations on an EX Series Switch

Example: Configuring 802.1X Authentication Options When the RADIUS Server Is Unavailable to an EX Series Switch

Example: Configuring Fallback Options on EX Series Switches for EAP-TTLS Authentication and Odyssey Access Clients

Filtering 802.1X Supplicants by Using RADIUS Server Attributes

Verifying 802.1X Authentication

show dot1x accounting attribute

IN THIS SECTION

- Syntax | 19558
- Description | 19558
- Required Privilege Level | 19558
- Output Fields | 19558
- Sample Output | 19560
- Release Information | 19560

Syntax

```
show dot1x accounting attribute
```

Description

Display the RADIUS accounting attributes sent by the switch, operating as the network access server (NAS), to the RADIUS accounting server. RADIUS accounting attributes convey information that is used to account for a service provided to an authenticated user. The user session statistics are recorded by the accounting server in an accounting log file.

RADIUS accounting attributes are included in Accounting-Request messages sent from the switch to the accounting server. Attribute information is created only if the data for the attribute is available.

Required Privilege Level

view

Output Fields

Table 464 on page 19559 lists the output fields for the `show dot1x accounting-attributes` command. Output fields are listed in the approximate order in which they appear.

Table 464: show dot1x accounting attribute Output Fields

Field Name	Field Description
`Accounting attributes`	Shows the value for the RADIUS accounting attributes sent from the NAS to the server. An attribute is displayed only if data is available for that attribute value. The following RADIUS accounting attributes are supported: • User-Name—The name of the authenticated user. • NAS-Port—The physical port number of the NAS which is authenticating the user. • Framed-IP-Address—The IP address of the authenticated user. • Filter-ID—The name of the filter list for the authenticated user. • Framed-MTU—The maximum transmission unit that can be configured for user. • Client-System-Name—This is a vendor-specific attribute (VSA) used to indicate the client host name. Supported for LLDP-capable devices only. • Session-Timeout—The maximum number of seconds that a session will stay active before termination of the session or prompt. • Called-Station-ID—Allows the NAS to send the phone number that the user called, using Dialed Number Identification (DNIS) or similar technology. • Calling-Station-ID—Allows the NAS to send the phone number that the call came from, using Dialed Number Identification (DNIS) or similar technology. • NAS-Identifier—Contains a string identifying the NAS originating the Accounting-Request. • Acct-Status-Type—Indicates whether this Accounting-Request marks the beginning of the user session (Start) or the end (Stop). Can also be used for an interim update (Interim-Update). • Acct-Authentic—Indicates whether the user was authenticated locally (Local), by the RADIUS server (RADIUS), or by another remote authentication protocol (Remote). • Acct-Session-ID—A unique ID for a specific accounting session that can be used to match start and stop records for a session in the log file. • Event-Timestamp—Records the time an event occurred. • NAS-Port-ID—The port of the NAS that is authenticating the user. • Framed-IPv6-Address—The IPv6 address of the authenticated user.

Sample Output

show dot1x accounting-attributes

```
user@switch> show dot1x accounting-attributes

Accounting Attribute:
      Calling Station Id:              88-e0-f3-1f-c5-e0
      Called station Id:               00-10-94-00-00-02
      Framed Ipv6 Address              :2001:db8:0:1:2a0:a514:0:24d
      Accounting Session ID:           802.1x812f00250002dcc6
      Client System Name:              AVX149485
      Session-Timeout:                 120s
      Framed-MTU:                         492
      Acct-Authentic:                  RADIUS
      Nas-Port-ID                      ge-0/0/5.0
```

Release Information

Command introduced in JUNOS Release 16.1 .

RELATED DOCUMENTATION

show dot1x | **19548**

clear dot1x | **3564**

Example: Configuring Static MAC Bypass of 802.1X and MAC RADIUS Authentication on an EX Series Switch

Example: Setting Up 802.1X for Single-Supplicant or Multiple-Supplicant Configurations on an EX Series Switch

Filtering 802.1X Supplicants by Using RADIUS Server Attributes

show dot1x authentication-bypassed-users

IN THIS SECTION

- Syntax | 19561
- Description | 19561
- Required Privilege Level | 19561
- Output Fields | 19561
- Sample Output | 19562
- Release Information | 19562

Syntax

```
show dot1x authentication-bypassed-users
```

Description

Display the supplicants (users) that have bypassed 802.1X authentication.

Required Privilege Level

view

Output Fields

Table 465 on page 19561 lists the output fields for the show dot1x authentication-bypassed-users command. Output fields are listed in the approximate order in which they appear.

Table 465: show dot1x authentication-bypassed-users Output Fields

Field Name	Field Description	Level of Output
Interface	The MAC address configured to bypass 802.1X authentication.	**all**

Table 465: show dot1x authentication-bypassed-users Output Fields *(Continued)*

Field Name	Field Description	Level of Output
MAC address	The MAC address configured statically on the interface.	**all**
VLAN	The VLAN that is configured to bypass 802.1X authentication.	**all**

Sample Output

show dot1x authentication-bypassed-users

```
user@host> show dot1x authentication-bypassed-users

MAC address           Interface      VLAN
00:50:56:85:66:0f     ge-0/0/0.0     vlan6
00:50:56:9e:56:42     ge-0/0/1.0     vlan6
```

Release Information

Command introduced in Junos OS Release 15.1X49-D80.

RELATED DOCUMENTATION

show dot1x authentication-failed-users | **19562**

dot1x | **5917**

show dot1x authentication-failed-users

IN THIS SECTION

Syntax | **19563**

- Description | 19563
- Required Privilege Level | 19563
- Output Fields | 19563
- Sample Output | 19564
- Release Information | 19564

Syntax

```
show dot1x authentication-failed-users
```

Description

Display the supplicants (users) that have failed 802.1X authentication.

Required Privilege Level

view

Output Fields

Table 466 on page 19563 lists the output fields for the `show dot1x authentication-failed-users` command. Output fields are listed in the approximate order in which they appear.

Table 466: show dot1x authentication-failed-users Output Fields

Field Name	Field Description	Level of Output
Interface	The MAC address configured to bypass 802.1X authentication.	**all**
MAC address	The MAC address configured statically on the interface.	**all**
User	The user that is configured on the RADIUS server and that has failed 802.1X authentication.	**all**

Table 466: show dot1x authentication-failed-users Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Failure Count	The number of times that 802.1X authentication has failed on the interface.	**all**

Sample Output

show dot1x authentication-failed-users

```
user@host> show dot1x authentication-failed-users

Interface      MAC address          User            Failure Count
ge-0/0/0.0     00:50:56:85:66:0f    00505685660f    1
ge-0/0/1.0     00:50:56:9e:56:42    0050569e5642    1
```

Release Information

Command introduced in Junos OS Release 15.1X49-D80.

RELATED DOCUMENTATION

show dot1x authentication-bypassed-users | **19561**

dot1x | **5917**

show dot1x firewall

IN THIS SECTION

- Syntax | **19565**
- Description | **19565**
- Options | **19565**

- Required Privilege Level | 19565
- Output Fields | 19565
- Sample Output | 19566
- Release Information | 19566

Syntax

```
show dot1x firewall <interface interface-name>
```

Description

Display information about the firewall filters for each user or nonresponsive host that is authenticated on each 802.1X-enabled interface that is configured for multiple supplicants. For example, if the firewall filter is configured with a term for counters, the command shows the count for each user.

Options

none	Display information for all interfaces.
interface *interface-names*	(Optional) Display information for the specified interface.

Required Privilege Level

view

Output Fields

Output fields include any action modifier that is specified in firewall filters.

Sample Output

show dot1x firewall

(Showing counter action)

```
user@switch> show dot1x firewall
Filter: dot1x-filter-ge-0/0/3
Counters
 counter1_dot1x_ge-0/0/3_user1   342
 counter1_dot1x_ge-0/0/3_user2   857
```

show dot1x firewall

(Showing policer action)

```
user@switch> show dot1x firewall
Filter: dot1x_ge-0/0/0
Counters
 p1-t1   494946
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

clear dot1x | 3564

show dot1x interface

IN THIS SECTION

- Syntax | 19567
- Description | 19567
- Options | 19567
- Required Privilege Level | 19568
- Output Fields | 19568
- Sample Output | 19573
- Release Information | 19576

Syntax

```
show dot1x interface interface-name
<brief | detail | extensive>
```

Description

Display the current operational state of all ports with the list of connected users.

This command displays the list of connected supplicants received from the RADIUS authentication server regardless of the session state—that is, for both authenticated supplicants and for supplicants that attempted authentication.

Options

none	Display information for all authenticator ports.
brief | detail | extensive	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display information for the specified interface with a list of connected supplicants.

Required Privilege Level

view

Output Fields

Table 467 on page 19568 lists the output fields for the `show dot1x interface` command. Output fields are listed in the approximate order in which they appear.

Table 467: show dot1x interface Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of a port.	**all**
`MAC address`	The MAC address of the connected supplicant on the port.	**all**
`Role`	The 802.1X authentication role of the interface. When 802.1X is enabled on an interface, the role is **Authenticator**. As **Authenticator**, the interface blocks LAN access until a supplicant is authenticated through 802.1X or MAC RADIUS authentication.	`brief`, `detail`
`State`	The state of the port: • **Authenticated**—The supplicant has been authenticated through the RADIUS server or has been permitted access through server fail fallback. • **Authenticating**—The supplicant is authenticating through the RADIUS server. • **Held**—An action has been triggered through server fail fallback during a RADIUS server timeout. A supplicant is denied access, permitted access through a specified VLAN, or maintains the authenticated state granted to it before the RADIUS server timeout occurred.	`brief`, `extensive`
`User`	The username of the connected supplicant.	`brief`, `extensive`

Table 467: show dot1x interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Administrative state`	The administrative state of the port: • **auto**—Traffic is allowed through the port based on the authentication result (by default). • **force-authorize**—All traffic flows through the port irrespective of the authentication result. This state is not allowed on an interface whose VLAN membership has been set to **dynamic**. • **force-unauthorize**—All traffic drops on the port irrespective of the authentication result. This state is not allowed on an interface whose VLAN membership has been set to **dynamic**.	`detail`
`Supplicant`	The mode for the supplicant: • **single**—Only the first supplicant is authenticated. All other supplicants that connect later to the port are allowed full access without any further authentication. They effectively *piggyback* on the first supplicant's authentication. • **single-secure**—Only one supplicant is allowed to connect to the port. No other supplicant is allowed to connect until the first supplicant logs out. • **multiple**—Multiple supplicants are allowed to connect to the port. Each supplicant is authenticated individually.	`detail`
`Quiet period`	The number of seconds the port waits before reattempting authentication after a failed authentication exchange with the supplicant.	`detail`
`Transmit period`	The number of seconds the port waits before retransmitting the initial EAPOL PDUs to the supplicant.	`detail`

Table 467: show dot1x interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`MAC Radius`	MAC RADIUS authentication: • **enabled**—The device sends an EAPOL request to the connecting host to attempt 802.1X authentication and if the connecting host is unresponsive, the device tries to authenticate the host by using the MAC address. • **disabled**—The default. The device does not attempt to authenticate the MAC address of the connecting host.	`detail`
`MAC Radius authentication protocol`	MAC RADIUS authentication protocol: • **EAP-MD5**—The EAP-MD5 protocol is used for MAC RADIUS authentication. EAP-MD5 is an authentication method belonging to the Extensible Authentication Protocol (EAP) authentication framework. EAP-MD5 is the default authentication protocol. • **PAP**—The Password Authentication Protocol (PAP) authentication protocol is used for MAC RADIUS authentication.	`detail`
`MAC Radius restrict`	The authentication method is restricted to MAC RADIUS. 802.1X authentication is not enabled.	`detail`
`Reauthentication`	The reauthentication state: • **disable**—Periodic reauthentication of the client is disabled. • **interval**—Sets the periodic reauthentication time interval.	`detail`
`Supplicant timeout`	The number of seconds the port waits for a response when relaying a request from the authentication server to the supplicant before resending the request.	`detail`
`Server timeout`	The number of seconds the port waits for a reply when relaying a response from the supplicant to the authentication server before timing out.	`detail`

Table 467: show dot1x interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Maximum EAPOL requests`	The maximum number of times an EAPOL request packet is retransmitted to the supplicant before the authentication session times out.	`detail`
`Number of clients bypassed because of authentication`	The number of non-802.1X clients granted access to the LAN by means of static MAC bypass. The following fields are displayed: • **Client**—MAC address of the client. • **vlan**—The name of the VLAN to which the client is connected.	`detail`
`Guest VLAN member`	The VLAN to which a supplicant is connected when the supplicant is authenticated using a guest VLAN. If a guest VLAN is not configured on the interface, this field displays <**not configured**>.	`detail`
`Retain mac aged session`	The behavior of the IEEE 802.1X clients when their MAC addresses age out: • **Enabled**—Keeps the IEEE 802.1X clients active even after their MAC addresses age out, and relearns the MAC addresses again. • **Disabled**—The default. MAC addresses are not relearnt when the MAC addresses age out.	`detail`
`Number of connected supplicants`	The number of supplicants connected to a port.	`detail`
`Supplicant`	The username and MAC address of the connected supplicant.	`detail`

Table 467: show dot1x interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Authentication method`
`
	The authentication method used for a supplicant:
	• `CWA Authentication`—A supplicant is authenticated by the central Web authentication (CWA) server.
	• `Guest VLAN`—A supplicant is connected to the LAN through the guest VLAN.
	• `MAC RADIUS`—A nonresponsive host is authenticated based on its MAC address. The MAC address is configured as permitted on the RADIUS server. The RADIUS server lets the device know that the MAC address is a permitted address, and the device opens LAN access to the nonresponsive host on the interface to which it is connected.
	• `RADIUS`—A supplicant is configured on the RADIUS server, the RADIUS server communicates this to the device, and the device opens LAN access on the interface to which the supplicant is connected.
	• `Server-fail`—One of the following fallback actions is in effect because the RADIUS server is unreachable. Indicates whether EAPOL block is in effect, and the amount of time remaining for EAPOL block (in seconds).
	  • `deny`—The supplicant is denied access to the LAN, preventing traffic from flowing from the supplicant through the interface. This is the default server fail fallback action.
	  • `permit`—The supplicant is permitted access to the LAN as if the supplicant had been successfully authenticated by the RADIUS server.
	  • `use-cache`—In the event that the RADIUS server times out when the supplicant is attempting reauthentication, the supplicant is reauthenticated only if it was previously authenticated; otherwise, the supplicant is denied LAN access.
	  • `VLAN`—The supplicant is configured to be moved to a specified VLAN if the RADIUS server is unavailable to reauthenticate the supplicant. (The VLAN must already exist on the device.)
		`detail`, `extensive`

Table 467: show dot1x interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Authenticated VLAN`	The VLAN to which the supplicant is connected.	`detail`, `extensive`
`Dynamic filter`	User policy filter sent by the RADIUS server.	`detail`
`Group Based Policy (GBP) ID`	The configured GBP tag received by the Juniper-Switching-Filter VSA or the Juniper-Group-Based-Policy-Id VSA.	`detail`
`Session Reauth interval`	The configured reauthentication interval.	`detail`
`Reauthentication due in`	The number of seconds in which reauthentication occurs again for the connected supplicant.	`detail`
`Session Accounting Interim Interval`	The number of seconds between interim RADIUS accounting messages.	`detail`
`Accounting Update due in`	The number of seconds until the next interim RADIUS accounting update is due.	`detail`
`CWA Redirect URL`	The URL used to redirect the supplicant to a central Web server for authentication.	`detail`
Authenticated VoIP VLAN	The VoIP VLAN to which the supplicant is connected.	`detail`, `extensive`

Sample Output

show dot1x interface brief

```
user@host> show dot1x interface brief
802.1X Information:
Interface     Role           State           MAC address             User
```

```
ge-0/0/1      Authenticator  Connecting      2001:db8:56:85:66:0F    00505685660f
ge-0/0/2      Authenticator  Authenticated   2001:db8:56:9E:56:42    0050569e5642
```

show dot1x interface detail (with GBP configured for QFX and EX switches)

```
user@host> show dot1x interface detail
ge-0/0/0.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 2 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Disabled
  Mac Radius Authentication Protocol: EAP-MD5
  Reauthentication: Enabled
  Reauthentication interval: 120 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: not configured
  Retain mac aged session: Enabled
  Number of connected supplicants: 1
    Supplicant: 001094001122, 00:10:94:00:11:22
      Operational state: Authenticated
      Backend Authentication state: Idle
      Authentication method: Mac Radius
      Authenticated VLAN: VLAN_1
      Group Based Policy Id: 100
      Session Reauth interval: 120 seconds
      Reauthentication due in 51 seconds
      Eapol-Block: Not In Effect
      Domain: Data
```

```
user@host> show dot1x interface detail

 ge-0/0/0.0
  Role: Authenticator
  Administrative state: Auto
```

```
  Supplicant mode: Multiple
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 30 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: not configured
  Retain mac aged session: Enabled
  Number of connected supplicants: 1
    Supplicant: 00505685660f, 00:50:56:85:66:0F
      Operational state: Authenticated
      Backend Authentication state: Idle
      Authentication method: Server-Reject Vlan
      Authenticated VLAN: visitor-vlan
      Session Reauth interval: 30 seconds
      Reauthentication due in 20 seconds
ge-0/0/1.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Multiple
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 30 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: not configured
  Number of connected supplicants: 1
    Supplicant: 0050569e5642, 00:50:56:9E:56:42
      Operational state: Authenticated
      Backend Authentication state: Idle
      Authentication method: Server-Reject Vlan
      Authenticated VLAN: visitor-vlan
```

```
      Session Reauth interval: 30 seconds
      Reauthentication due in 24 seconds
```

show dot1x interface extensive

```
user@host> show dot1x interface extensive

802.1X Information:
Interface  State         MAC address          Method             Vlan User
ge-0/0/6.0 Authenticated 2001:db8:94:00:00:01 Server-Reject Vlan 1400 Test12345
```

Release Information

Command introduced in Junos OS Release 15.1X49-D80.

extensive option introduced in Junos OS Release 19.4R1 to display the additional fields when compared to **brief** option. The additional fields are `authentication method` and `vlan-id`.

show dot1x interface detail (to check authenticated VoIP VLAN)

```
user@host> show dot1x interface detail
ge-0/0/0.0
 xe-0/0/37.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Multiple
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Enabled
  Mac Radius Authentication Protocol: EAP-MD5
  Reauthentication: Enabled
  Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: not configured
  Last Mac-Learn Request: 00:10:94:00:00:34
  Last Mac-Learn Time: 2024-09-12 10:19:53.254451
```

```
  Number of connected supplicants: 1
    Supplicant: 001094000034, 00:10:94:00:00:34
      Operational state: Authenticated
      Backend Authentication state: Idle
      Authentication method: Mac Radius
      Authenticated VLAN: vlan-10
      Authenticated Voip VLAN: vlan-20
      Session Reauth interval: 3600 seconds
      Reauthentication due in 3594 seconds
      Eapol-Block: Not In Effect
      Domain: Voip
```

```
user@host> show dot1x interface extensive

 802.1X Information:
Interface     State          MAC address         Method             Data  Voip  User
                                                                    vlan  vlan
xe-0/0/37.0   Authenticated  00:10:94:00:00:34   Mac Radius         10    20    001094000034
```

RELATED DOCUMENTATION

show dot1x authentication-bypassed-users | 19561

dot1x | 5917

show dot1x static-mac-address

IN THIS SECTION

- Syntax | 19578
- Description | 19578
- Options | 19578
- Required Privilege Level | 19578
- Output Fields | 19578
- Sample Output | 19579

Release Information | 19579

Syntax

```
show dot1x static-mac-address <interface interface-name>
```

Description

Display all the static MAC addresses of interfaces that are configured to bypass 802.1X authentication.

Options

none	Display static MAC addresses for all interfaces.
`interface` *`interface-name`*	(Optional) Display static MAC addresses for a specific interface.

Required Privilege Level

view

Output Fields

Table 468 on page 19578 lists the output fields for the `show dot1x static-mac-address` command. Output fields are listed in the approximate order in which they appear.

Table 468: show dot1x static-mac-address Output Fields

Field Name	Field Description	Level of Output
MAC address/ prefix	The MAC address of the device that is configured to bypass 802.1X authentication.	**all**
VLAN-Assignment	The name of the VLAN to which the device is assigned.	**all**

Table 468: show dot1x static-mac-address Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Interface	The name of the interface on which authentication is bypassed for a given MAC address.	**all**

Sample Output

show dot1x static-mac-address

```
user@host> show dot1x static-mac-address

MAC address/prefix            VLAN-Assignment           Interface
00:50:56:85:66:0f/48          vlan6                     ge-0/0/0.0
00:50:56:9e:56:42/48          vlan6                     ge-0/0/1.0
```

show dot1x static-mac-address interface (Specific Interface)

```
user@host> show dot1x static-mac-address interface ge-0/0/0

MAC address/prefix            VLAN-Assignment           Interface
00:50:56:85:66:0f/48          vlan6                     ge-0/0/0.0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D80.

RELATED DOCUMENTATION

show dot1x authentication-bypassed-users | **19561**

dot1x | **5917**

show dot1x statistics

IN THIS SECTION

- Syntax | 19580
- Description | 19580
- Options | 19580
- Required Privilege Level | 19580
- Output Fields | 19581
- Sample Output | 19582
- Release Information | 19582

Syntax

```
show dot1x statistics
<interface interface>
```

Description

Display the number of EAPOL messages transmitted or received on all interfaces or specific interfaces.

Options

none	Displays statistical information for all interfaces.
interface *interface-name*	(Optional) Displays statistical information for the specified interface.

Required Privilege Level

view

Output Fields

Table 469 on page 19581 lists the output fields for the `show dot1x statistics` command. Output fields are listed in the approximate order in which they appear.

Table 469: show dot1x statistics Output Fields

Field Name	Field Description
TxReqId	The number of EAP-Request/Identity messages transmitted on the interface.
TxReq	The number of transmitted EAP-Request frames that were not EAP-Request/Identity.
TxTotal	The total number of EAPOL messages transmitted on the interface.
RxStart	The number of EAPOL-Start messages received on the interface.
RxLogoff	The number of EAP-Logoff messages received on the interface.
RxRespId	The number of EAP-Response/Identity frames received on the interface.
RxResp	The number of EAP-Response messages received that were not EAP-Response/Identity.
CoA-Request	The number of Change of Authorization (CoA) Request messages received on the interface.
CoA-Ack	The number of CoA-Ack messages transmitted on the interface.
CoA-Nak	The number of CoA-Nak messages transmitted on the interface.
RxInvalid	The number of invalid EAPOL messages received on the interface.
RxLenErr	The number of EAPOL messages with incorrect length received on the interface.
RxTotal	The total number of EAPOL messages creceived on the interface.

Table 469: show dot1x statistics Output Fields *(Continued)*

Field Name	Field Description
LastRxVersion	The version number of the last EAPOL message received on the interface.
LastRxSrcMac	The source MAC address in the last EAPOL message received on the interface.
PortBounceReqRx	The number of port bounce requests received on the port.

Sample Output

show dot1x statistics interface

```
user@host> show dot1x statistics interface ge-0/0/0

Interface: ge-0/0/0.0
  TxReqId = 4 TxReq = 0 TxTotal = 4
  RxStart = 0 RxLogoff = 0 RxRespId = 0 RxResp = 0
  CoA-Request = 0 CoA-Ack = 0 CoA-Nak = 0
  RxInvalid = 0 RxLenErr = 0 RxTotal = 0
  LastRxVersion = 0 LastRxSrcMac = 00:50:56:85:66:0f
  PortBounceReqRx = 0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show dot1x | 19548

clear dot1x | 3564

Example: Setting Up 802.1X for Single-Supplicant or Multiple-Supplicant Configurations on an EX Series Switch

Filtering 802.1X Supplicants by Using RADIUS Server Attributes

show dual-phase-bootup-status

IN THIS SECTION

- Syntax | 19583
- Description | 19583
- Required Privilege Level | 19583
- Sample Output | 19584
- Release Information | 19584

Syntax

```
show dual-phase-bootup-status
```

Description

Check the current status of a dual-phase bootup or see if dual-phase bootup is enabled. There are four possible statuses:

- Enabled : This state means the dual-phase-bootup feature is enabled and it is waiting to load the device with the rescue configuration file (**rescue.conf)** on the next reboot.
- Reboot activated: During a reboot the device was able to find the rescue configuration. Post reboot, if the device is still in this state, then there was a problem loading the rescue configuration on the device.
- rescue.conf operational: The device arrives at this state when the reboot is successful, and the device comes up with the rescue configuration **rescue.conf**.
- Enabled: This state is achieved again when the device is finished and successful with the internal commit of the last-run user configuration. Check the log messages to see if everything ran as expected during this commit of the last-run user configuration.

Required Privilege Level

system

Sample Output

show dual-phase-bootup-status

```
user@host> show dual-phase-bootup-status
```

```
Dual phase bootup status : Enabled
```

Release Information

Command introduced in Junos OS Release 23.4R1.

show dvmrp interfaces

IN THIS SECTION

- Syntax | 19584
- Description | 19585
- Options | 19585
- Required Privilege Level | 19585
- Output Fields | 19585
- Sample Output | 19586
- Release Information | 19586

Syntax

```
show dvmrp interfaces
<logical-system (all | logical-system-name)>
```

Description

Display information about Distance Vector Multicast Routing Protocol (DVMRP)–enabled interfaces.

Options

none	(Same as **logical-system all**) Display information about DVMRP-enabled interfaces.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 470 on page 19585 describes the output fields for the `show dvmrp interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 470: show dvmrp interfaces Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: **up** or **down**.
Leaf	Whether the interface is a leaf (that is, whether it has no neighbors) or whether it has neighbors.
Metric	Interface metric: a value from **1** through **31**.
Announce	Number of routes the interface is announcing.

Table 470: show dvmrp interfaces Output Fields *(Continued)*

Field Name	Field Description
Mode	DVMRP mode: • **Forwarding**—DVMRP does both the routing and the multicast data forwarding. • **Unicast-routing**—DVMRP does only the routing. Forwarding of the multicast data packets can be done by enabling PIM on the interface.

Sample Output

show dvmrp interfaces

```
user@host> show dvmrp interfaces
Interface State Leaf Metric Announce Mode
fxp0.0   Up   N   1   4 Forwarding
fxp1.0   Up   N   1   4 Forwarding
fxp2.0   Up   N   1   3 Forwarding
lo0.0    Up   Y   1   0 Unicast-routing
```

Release Information

NOTE: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Command introduced before Junos OS Release 7.4.

show dvmrp neighbors

IN THIS SECTION

- Syntax | 19587
- Description | 19587
- Options | 19587
- Required Privilege Level | 19587
- Output Fields | 19588
- Sample Output | 19589
- Release Information | 19589

Syntax

```
show dvmrp neighbors
<logical-system (all | logical-system-name)>
```

Description

Display information about Distance Vector Multicast Routing Protocol (DVMRP) neighbors.

Options

none	(Same as **logical-system all**) Display information about DVMRP neighbors.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 471 on page 19588 describes the output fields for the `show dvmrp neighbors` command. Output fields are listed in the approximate order in which they appear.

Table 471: show dvmrp neighbors Output Fields

Field Name	Field Description
Neighbor	Address of the neighboring DVMRP router.
Interface	Interface through which the neighbor is reachable.
Version	Version of DVMRP that the neighbor is running, in the format ***majorminor***.
Flags	Information about the neighbor: • **1**—One way. The local router has seen the neighbor, but the neighbor has not seen the local router. • **G**—Neighbor supports generation ID. • **L**—Neighbor is a leaf router. • **M**—Neighbor supports mtrace. • **N**—Neighbor supports netmask in prune messages and graft messages. • **P**—Neighbor supports pruning. • **S**—Neighbor supports SNMP.
Routes	Number of routes learned from the neighbor.
Timeout	How long until the DVMRP neighbor information times out, in seconds.
Transitions	Number of generation ID changes that have occurred since the local router learned about the neighbor.

Sample Output

show dvmrp neighbors

```
user@host> show dvmrp neighbors
Neighbor        Interface        Version Flags   Routes Timeout Transitions
192.168.1.1    ipip.0             3.255  PGM         3      28           1
```

Release Information

> **NOTE**: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Command introduced before Junos OS Release 7.4.

show dvmrp prefix

IN THIS SECTION

- Syntax | 19590
- Description | 19590
- Options | 19590
- Required Privilege Level | 19590
- Output Fields | 19590
- Sample Output | 19591
- Release Information | 19592

Syntax

```
show dvmrp prefix
<brief | detail>
<logical-system (all | logical-system-name)>
<prefix>
```

Description

Display information about Distance Vector Multicast Routing Protocol (DVMRP) prefixes.

Options

none	Display standard information about all DVMRP prefixes.
brief | detail	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
prefix	(Optional) Display information about specific prefixes.

Required Privilege Level

view

Output Fields

Table 472 on page 19590 describes the output fields for the `show dvmrp prefix` command. Output fields are listed in the approximate order in which they appear.

Table 472: show dvmrp prefix Output Fields

Field Name	Field Description	Level of Output
Prefix	DVMRP route.	All levels
Next hop	Next hop from which the route was learned.	All levels

Table 472: show dvmrp prefix Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Age	Last time that the route was refreshed.	All levels
multicast-group	Multicast group address.	**detail**
Prunes sent	Number of prune messages sent to the multicast group.	**detail**
Grafts sent	Number of grafts sent to the multicast group.	**detail**
Cache lifetime	Lifetime of the group in the multicast cache, in seconds.	**detail**
Prune lifetime	Lifetime remaining and total lifetime of prune messages, in seconds.	**detail**

Sample Output

show dvmrp prefix

```
user@host> show dvmrp prefix
Prefix              Next hop        Age
10.38.0.0      /30 10.38.0.1       00:06:17
10.38.0.4      /30 10.38.0.5       00:06:13
10.38.0.8      /30 10.38.0.2       00:00:04
10.38.0.12     /30 10.38.0.6       00:00:04
10.255.14.114  /32 10.255.14.114   00:06:17
10.255.14.142  /32 10.38.0.2       00:00:04
10.255.14.144  /32 10.38.0.2       00:00:04
10.255.70.15   /32 10.38.0.6       00:00:04
192.168.14.0   /24 192.168.14.114  00:06:17
192.168.195.40 /30 192.168.195.41  00:06:17
192.168.195.92 /30 10.38.0.2       00:00:04
```

show dvmrp prefix brief

The output for the `show dvmrp prefix brief` command is identical to that for the `show dvmrp prefix` command.

show dvmrp prefix detail

```
user@host> show dvmrp prefix detail
Prefix              Next hop        Age
10.38.0.0      /30 10.38.0.1       00:06:28
10.38.0.4      /30 10.38.0.5       00:06:24
10.38.0.8      /30 10.38.0.2       00:00:15
10.38.0.12     /30 10.38.0.6       00:00:15
10.255.14.114  /32 10.255.14.114   00:06:28
10.255.14.142  /32 10.38.0.2       00:00:15
10.255.14.144  /32 10.38.0.2       00:00:15
10.255.70.15   /32 10.38.0.6       00:00:15
192.168.14.0   /24 192.168.14.114  00:06:28
192.168.195.40 /30 192.168.195.41  00:06:28
192.168.195.92 /30 10.38.0.2       00:00:15
```

Release Information

NOTE: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Command introduced before Junos OS Release 7.4.

show dvmrp prunes

IN THIS SECTION

- Syntax | 19593
- Description | 19593
- Options | 19593
- Required Privilege Level | 19593
- Output Fields | 19593

Sample Output | 19594
Release Information | 19594

Syntax

```
show dvmrp prunes
<all | rx | tx>
<logical-system (all | logical-system-name)>
```

Description

Display information about active Distance Vector Multicast Routing Protocol (DVMRP) prune messages.

Options

none	Display received and transmitted DVMRP prune information.
all	(Optional) Display information about all received and transmitted prune messages.
rx	(Optional) Display information about received prune messages.
tx	(Optional) Display information about transmitted prune messages.
logical-system (all | *logical-system-name)*	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 473 on page 19594 describes the output fields for the `show dvmrp prunes` command. Output fields are listed in the approximate order in which they appear.

Table 473: show dvmrp prunes Output Fields

Field Name	Field Description
Group	Group address.
Source prefix	Prefix for the prune.
Timeout	How long until the prune message expires, in seconds.
Neighbor	Neighbor to which the prune was sent or from which the prune was received.

Sample Output

show dvmrp prunes

```
user@host> show dvmrp prunes
Group           Source prefix      Timeout Neighbor
224.0.1.1       128.112.0.0    /12    7077 192.168.1.1
224.0.1.32      160.0.0.0      /3     7087 192.168.1.1
224.2.123.4     136.0.0.0      /5     6955 192.168.1.1
224.2.127.1     129.0.0.0      /8     7046 192.168.1.1
224.2.135.86    128.102.128.0  /17    7071 192.168.1.1
224.2.135.86    129.0.0.0      /8     7074 192.168.1.1
224.2.135.86    130.0.0.0      /7     7071 192.168.1.1
...
```

Release Information

NOTE: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Command introduced before Junos OS Release 7.4.

show dynamic-profile session

IN THIS SECTION

- Syntax | 19595
- Description | 19595
- Options | 19596
- Required Privilege Level | 19596
- Output Fields | 19596
- Sample Output | 19596
- Release Information | 19601

Syntax

```
show dynamic-profile session
<client-id client-id>
<profile-name profile-name>
<service-id service-id>
```

Description

Display dynamic profile (client or service) information for all subscribers or for subscribers specified by client ID or service session ID. You can filter the output by also specifying a dynamic profile.

NOTE:

- The output does not display the variable stanzas defined in the dynamic profile configuration.
- The variables in the profile configuration are replaced with subscriber specific values.
- If the conditional variable in the dynamic profile is evaluated as NULL, the subscriber value for the variable is displayed as `NONE` in the command output.

- The variable is also displayed as `NONE` when the variable (any variable and not necessarily conditional) in the dynamic profile has no value associated with it.
- The format in which the configuration is displayed looks similar, but not exactly the same as the format of the `show configuration dynamic-profiles` command.

Options

client-id *client-id*	Display dynamic profile information for subscribers associated with the specified client.
profile-name *profile-name*	(Optional) Display dynamic profile information for the specified subscriber or service profile.
service-id *service-id*	Display dynamic profile information for subscribers associated with the specified service session.

Required Privilege Level

view

Output Fields

This command displays the dynamic client or service profile configuration for each subscriber.

Sample Output

show dynamic-profile session client-id (Client ID)

```
user@host>show dynamic-profile session client-id 20
pppoe {
    interfaces {
        pp0 {
            unit 1073741831 {
                ppp-options {
                    chap;
                    pap;
                }
                pppoe-options {
```

```
                    underlying-interface ge-2/0/0.0;
                    server;
                }
                family {
                    inet {
                        unnumbered-address lo0.0;
                    }
                }
            }
        }
    }
    class-of-service {
        traffic-control-profiles {
            tcp1 {
                scheduler-map smap1_UID1024;
                shaping-rate 100m;
            }
        }
        interfaces {
            pp0 {
                unit 1073741831 {
                    output-traffic-control-profile tcp1;
                }
            }
        }
        scheduler-maps {
            smap1_UID1024 {
                forwarding-class best-effort scheduler sch1_UID1023;
            }
        }
        schedulers {
            sch1_UID1023 {
                transmit-rate percent 40;
                buffer-size percent 40;
                priority low;
            }
        }
    }
}
filter-service {
    interfaces {
        pp0 {
            unit 1073741831 {
```

```
                family {
                    inet {
                        filter {
                            input input-filter_UID1026 precedence 50;
                            output output-filter_UID1027 precedence 50;
                        }
                    }
                }
            }
        }
    }
    firewall {
        family {
            inet {
                filter input-filter_UID1026 {
                    interface-specific;
                    term t1 {
                        then {
                            policer policer1_UID1025;
                            service-accounting;
                        }
                    }
                    term rest {
                        then accept;
                    }
                }
                filter output-filter_UID1027 {
                    interface-specific;
                    term rest {
                        then accept;
                    }
                }
            }
        }
        policer policer1_UID1025 {
            if-exceeding {
                bandwidth-limit 1m;
                burst-size-limit 15k;
            }
            then discard;
        }
    }
}
```

```
cos-service {
    class-of-service {
        scheduler-maps {
            smap2_UID1029 {
                forwarding-class assured-forwarding scheduler sch2_UID1028;
            }
        }
        schedulers {
            sch2_UID1028 {
                transmit-rate percent 60;
                buffer-size percent 60;
                priority high;
            }
        }
    }
}
```

show dynamic-profile session client-id profile-name (Client ID and Dynamic Profile)

```
user@host>show dynamic-profile session client-id 20 profile-name cos-service
cos-service {
    class-of-service {
        scheduler-maps {
            smap2_UID1029 {
                forwarding-class assured-forwarding scheduler sch2_UID1028;
            }
        }
        schedulers {
            sch2_UID1028 {
                transmit-rate percent 60;
                buffer-size percent 60;
                priority high;
            }
        }
    }
}
```

show dynamic-profile session service-id (Service Session)

```
user@host>show dynamic-profile session service-id 21
filter-service {
    interfaces {
        pp0 {
            unit 1073741831 {
                family {
                    inet {
                        filter {
                            input input-filter_UID1026 precedence 50;
                            output output-filter_UID1027 precedence 50;
                        }
                    }
                }
            }
        }
    }
    firewall {
        family {
            inet {
                filter input-filter_UID1026 {
                    interface-specific;
                    term t1 {
                        then {
                            policer policer1_UID1025;
                            service-accounting;
                        }
                    }
                    term rest {
                        then accept;
                    }
                }
                filter output-filter_UID1027 {
                    interface-specific;
                    term rest {
                        then accept;
                    }
                }
            }
        }
        policer policer1_UID1025 {
```

```
            if-exceeding {
                bandwidth-limit 1m;
                burst-size-limit 15k;
            }
            then discard;
        }
    }
}
```

Release Information

Command introduced in Junos OS Release 13.3.

show dynamic-tunnels database

IN THIS SECTION

- Syntax | 19601
- Description | 19602
- Options | 19602
- Required Privilege Level | 19602
- Output Fields | 19602
- Sample Output | 19604
- Release Information | 19609

Syntax

```
show dynamic-tunnels database
<destination>
<logical-system (all | logical-system-name) >
<table routing-table-name>
```

Description

Display dynamic tunnel database information.

Options

none	Display dynamic tunnel database information for all destinations and routing tables.
destination	(Optional) Display database entries for the specified IP address (with optional destination prefix length) only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
table *routing-table-name*	(Optional) Display database entries for the specified table only.

Required Privilege Level

view

Output Fields

Table 474 on page 19602 lists the output fields for the `show dynamic-tunnels database` command. Output fields are listed in the approximate order in which they appear.

Table 474: show dynamic-tunnels database Output Fields

Field Name	Field Description
`Table`	Name of the routing table (for example, `inet.0`).
`Destination-network`	Destination IP address and subnet.
`Tunnel to`	Destination IP address and prefix of the tunnel.
`State`	State of the tunnel: `Up`, `Up (expires in` *nn*:*nn*:*nn*`seconds)`, or `Dn` (down).

Table 474: show dynamic-tunnels database Output Fields *(Continued)*

Field Name	Field Description
`Reference count`	Number of routes across the dynamic tunnel that are currently being resolved.
`Next-hop type`	Type of tunnel: GRE or UDP (BGP-Signal). • GRE, IPoIP, or UDP (BGP signal) • SR-TE—Segment routing traffic-engineered LSP. • GREv6, • UDPv6 • IPIPv6.
`Source address`	Source IP address of the tunnel.
`Next-hop`	IP address of the destination interface.
VPN Label	The label provided by the peer device to identify the VPN through which the packet needs to go. This label is used to identify the VRF for route lookup.
Ingress Route	The IGP route along with the corresponding metric that has been selected for forwarding the tunnel-encapsulated packet.
`Localized PFE`	Packet Forwarding Engine interface which is the anchor Packet Forwarding Engine for the localized next-hop-based dynamic tunnels. When the anchor Packet Forwarding Engine of the tunnel goes down, it is represented by a # near the Packet Forwarding Engine name.
LSP template name	Name of the segment routing traffic-engineered template configured for dynamic creation of segment routing LSPs.
`State`	State of the destination interface: `Up`, `Dn`, or `Dn (no tunnel pic)`.

Table 474: show dynamic-tunnels database Output Fields *(Continued)*

Field Name	Field Description
Status	Status of the dynamic segment routing LSP.
Tunnel Encapsulation	State of the encapulated packet.
Traffic Statistics	Number of packets transmitted.

Sample Output

show dynamic-tunnels database (Tunnel Is Up)

```
user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32
  Reference count: 4
  Next-hop type: UDP
    Source address: 10.255.120.92
    Next hop: tunnel-composite, 0x31132f64, nhid 3406
      VPN Label: Push 120 Reference count: 3
      Ingress Route: 10.255.120.94/32, via metric 2
      Traffic Statistics: Packets 241367951, Bytes 356741831578
      State: Up
```

show dynamic-tunnels database (No Tunnel PIC)

```
user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Dn
  Reference count: 2
  Next-hop type: gre
```

```
    Source address: 10.255.120.92
    State: Dn (no tunnel pic)
```

show dynamic-tunnels database (Tunnel Is Expiring)

```
user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up (expires in 00:14:56 seconds)
  Reference count: 0
  Next-hop type: gre
    Source address: 10.255.120.92
    Next hop: gr-4/3/0.32769
    State: Up
```

show dynamic-tunnels database (Destination Specified)

```
user@host> show dynamic-tunnels database 10.255.120.94
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up
  Reference count: 2
  Next-hop type: gre
    Source address: 10.255.120.92
    Next hop: gr-4/3/0.32769
    State: Up
```

show dynamic-tunnels database (Localization)

```
user@host> show dynamic-tunnels database
Destination-network: 10.0.0.0/8
Tunnel to: 10.1.1.6/32
  Reference count: 5
  Next-hop type: UDP
    Source address: 10.1.1.2
    Next hop: tunnel-composite, 0xc807930, nhid 1016
```

```
    Localized PFE: pfe-1/0/0
     VPN Label: Push 299808 Reference count: 4
 Ingress Route: 10.1.1.6/32, via metric 2
 Traffic Statistics: Packets 0, Bytes 0
 State: Up
```

show dynamic-tunnels database (MPLS-over-UDP Dynamic Tunnels on PTX Series Routers and QFX Series Switches)

```
user@host> show dynamic-tunnels database
*- Signal Tunnels #- PFE-down
Table: inet.3

Destination-network: 10.33.0.0/16

Destination-network: 10.33.44.0/24
```

show dynamic-tunnels database (Segment Routing LSPs)

```
user@host> show dynamic-tunnels database
Table: inetcolor.0

Destination-network: 10.33.44.0:0/24
Tunnel to: 10.33.44.55:124/64
  Reference count: 2
  Next-hop type: SRTE
  LSP template name: 10.33.44.55:7c:dt-srte-tunnel1
  Status: Initiated/Established
```

show dynamic-tunnels database (Dynamic Tunnels for IPv6 underlay on PTX Series Routers)

```
user@host> show dynamic-tunnels database
*- Signal Tunnels #- PFE-down
Table: inet6.3

Destination-network: 100::6464:6402/128

Tunnel to: 100::6464:6402/128
```

```
  Reference count: 3
  Next-hop type: UDPv6 (forwarding-nexthop)
    Source address: 1010::1
    Next hop: tunnel-composite, 0x7b42614, nhid 559
      Reference count: 2
      Ingress Route: [OSPF3] 100::6464:6402/128, via metric 1
      Traffic Statistics: Packets 0, Bytes 0
      State: Up
  Aggregate Traffic Statistics:
    Tunnel Encapsulation: Dest 100::6464:6402, Src 1010::1, UDPv6, Tunnel-Id 2
    Traffic Statistics: Packets 0, Bytes 0
```

show dynamic-tunnels database summary

```
user@host> show dynamic-tunnels database summary
Dynamic Tunnels,  Total 1 displayed
GRE Tunnel:
Active Tunnel Mode, Next Hop Base
  IFL Based,  Total 0 displayed, Up 0, Down 0
  Nexthop Based,  Total 0 displayed, Up 0, Down 0
RSVP Tunnel:
  Total 0 displayed
SPRING-TE tunnel:
  Total 0 displayed, Up 0, Down 0
UDP Tunnel:
  Total 0 displayed, Up 0, Down 0
IPIP Tunnel:
  Total 0 displayed, Up 0, Down 0
GREv6 Tunnel:
Active Tunnel Mode, Next Hop Base
  IFL Based,  Total 0 displayed, Up 0, Down 0
  Nexthop Based,  Total 0 displayed, Up 0, Down 0
UDPv6 Tunnel:
  Total 1 displayed, Up 1, Down 0
IPIPv6 Tunnel:
  Total 0 displayed, Up 0, Down 0
```

show route table inet6.3 extensive

```
user@host> show route table inet6.3 extensive
inet6.3: 1 destinations, 2 routes (1 active, 0 holddown, 0 hidden)
100::6464:6402/128 (2 entries, 1 announced)
        *Tunnel Preference: 5
                Next hop type: Tunnel Composite, Next hop index: 0
                Address: 0x7b3eb04
                Next-hop reference count: 2, key opaque handle: 0x0, non-key opaque handle: 0x0
                Tunnel type: UDPv6, (extended-attr), Reference count: 1, nhid: 0
                Destination address: 100::6464:6402, Source address: 1010::1
                Tunnel encap routing table id : 0, Tunnel decap routing table id : 0
                State: <Active>
                Local AS:  1000
                Age: 3:18:04    Metric: 1
                Validation State: unverified
                Area: 0.0.0.0
                Task: DYN_TUNNEL
                Announcement bits (3): 0-Resolve tree 2 1-Resolve_IGP_FRR task 2-Resolve tree 4
                AS path: I
                Thread: junos-main
         Tunnel Preference: 5
                Next hop type: Tunnel, Next hop index: 0
                Address: 0x7b3e7f4
                Next-hop reference count: 1, key opaque handle: 0x0, non-key opaque handle: 0x0
                Inactive reason: Route Preference
                Local AS:  1000
                Age: 3:18:04
                Validation State: unverified
                Task: DYN_TUNNEL
                AS path: I
                Thread: junos-main
```

show dynamic-tunnels database <destination> terse tunnel-type ?

```
user@host> show dynamic-tunnels database 100::6464:6402 terse tunnel-type udpv6 tunnel-id 2
Warning: License key missing; requires 'mplsoudp' license

*- Signal Tunnels #- PFE-down
```

```
Table: inet6.3

Destination-network: 100::6464:6402/128
Destination                     Source          Next-hop                Type       Status
100::6464:6402/128              1010::1         0x7b42614 nhid 559      UDPv6      Up (via
metric 1)
```

show dynamic-tunnels database statistics tunnel-type ?

```
user@host>show dynamic-tunnels database statistics tunnel-type ?
Possible completions:
  gre                  GRE tunnel type
  grev6                GRE v6 tunnel type
  ipoip                IPIP tunnel type
  ipv6                 IPv6 tunnel type
  udp                  UDP tunnel type
  udpv6                UDP v6 tunnel type
```

Release Information

Command introduced before Junos OS Release 7.4.

show dynamic-tunnels pfe-tunnel-localization

IN THIS SECTION

- Syntax | 19610
- Description | 19610
- Options | 19610
- Required Privilege Level | 19610
- Output Fields | 19610
- Sample Output | 19612
- Release Information | 19616

Syntax

```
show dynamic-tunnels pfe-tunnel-localization
```

Description

Display dynamic tunnel Packet Forwarding Engine (PFE) localization information such as tunnel count of an anchor PFE, the GENCFG blobs that are sent to the kernel for each IPv6 source, the aggregate routes connected to the anchor PFE, the IPv6 routes that are mapped to this anchor PFE. When the PFE goes down BGP withdraws all the aggregate routes and the IPv6 sources.

NOTE: PFE localization needs to be configured to display output for this command. This command throws back an error if there are no PFE localization entries to display.

Options

none Display dynamic tunnel localization information for the Packet Forwarding Engine tunnel.

Required Privilege Level

view

Output Fields

Table 475 on page 19610 lists the output fields for the `show dynamic-tunnels pfe-tunnel-localization` command. Output fields are listed in the approximate order in which they appear.

Table 475: show dynamic-tunnels pfe-tunnel-localization Output Fields

Field Name	Field Description
`Anchor PFE Name`	Name of the PFE in pfe-x/y/z format.
`Reference count`	Number of the dynamic tunnels that are currently anchored to a particular PFE. Each tunnel prefix and each Gencfg takes one reference count.

Table 475: show dynamic-tunnels pfe-tunnel-localization Output Fields *(Continued)*

Field Name	Field Description
`Gencfg keyid`	The key ID of a Gencfg blob that is returned by the kernel. A Gencfg blob stores the configuration values.
`Gencfg tunnel type: UDP`	Tunnel count anchored to a particular anchor Packet Forwarding Engine (reference count).
`Gencfg last index`	The last know index value of the Gencfg blob.
`Gencfg current index`	The current index value of the Gencfg blob.
`Gencfg last state`	Last state of the tunnel: `Up`, or `Dn` (down).
`Gencfg current state`	Current state of the tunnel: `Up`, or `Dn` (down).
`Gencfg V6 source address`	Source IP address of the Gencfg blob IPv6 tunnel.
`Gencfg backup lock`	Backup count of the Gencfg blob. This value is displayed in the backup routing protocol process. Each gencfg takes a backup reference count, which is decremented when tunnel is deleted or after NSR switchover and backup becomes new primary.
`Gencfg kqp`	Reference count of Gencfg entries in the Kernel Routing table (KRT) queue. KRT is the module within the Routing Process Daemon (RPD) that synchronizes the routing tables with the forwarding tables in the kernel. This table is then synchronized to the PFE through the KRT queue.
`Gencfg Reference count`	Each prefix using a particular IPv6 source takes reference count on the Gencfg corresponding to that V6 source. If the Gencfg is added to the kernel and it has key ID, then it has additional reference count of 1. When Gencfg operation is queued in KRTQ then krt q entry takes one ref count on Gencfg.

Table 475: show dynamic-tunnels pfe-tunnel-localization Output Fields *(Continued)*

Field Name	Field Description
Aggregate	All aggregate routes that are connected with this anchor PFE. When the PFE goes down, BGP withdraws all these aggregates.
Aggregate refcnt	Reference count of aggregate routes.
V6 source address	All IPv6 source routes that are connected to this anchor PFE. When the PFE goes down, BGP withdraws all these IPv6 sources.
V6 source refcnt	Reference count of IPv6 source routes.

Sample Output

show dynamic-tunnels pfe-tunnel-localization (Primary)

```
user@host> show dynamic-tunnels pfe-tunnel-localization
Anchor PFE Name: pfe-0/0/0
    Reference count: 8

    Gencfg keyid: 1
      Gencfg last index: 145
      Gencfg current index: 145
      Gencfg last state: UP
      Gencfg current state: UP
      Gencfg V6 source address: 66:66:00:00:00:00:00:00:00:00:00:00:00:00:66:66
      Gencfg backup lock: 0
      Gencfg kqp: 0x0
      Gencfg Reference count: 0x2

    Gencfg keyid: 2
      Gencfg last index: 145
      Gencfg current index: 145
      Gencfg last state: UP
      Gencfg current state: UP
      Gencfg V6 source address: 77:77:00:00:00:00:00:00:00:00:00:00:00:00:77:77
```

```
    Gencfg backup lock: 0
    Gencfg kqp: 0x0
    Gencfg Reference count: 0x2

  Gencfg keyid: 3
    Gencfg last index: 145
    Gencfg current index: 145
    Gencfg last state: UP
    Gencfg current state: UP
    Gencfg V6 source address: 88:88:00:00:00:00:00:00:00:00:00:00:00:00:88:88
    Gencfg backup lock: 0
    Gencfg kqp: 0x0
    Gencfg Reference count: 0x2

  Gencfg keyid: 4
    Gencfg last index: 145
    Gencfg current index: 145
    Gencfg last state: UP
    Gencfg current state: UP
    Gencfg V6 source address: 99:99:00:00:00:00:00:00:00:00:00:00:00:00:99:99
    Gencfg backup lock: 0
    Gencfg kqp: 0x0
    Gencfg Reference count: 0x2

  Aggregate: 4.4.4.0/24
  Aggregate refcnt: 1

  Aggregate: 1.1.1.0/24
  Aggregate refcnt: 1

  Aggregate: 2.2.2.0/24
  Aggregate refcnt: 1

  Aggregate: 3.3.3.0/24
  Aggregate refcnt: 1

  V6 source address: 9999::9999/128
  V6 source refcnt: 1

  V6 source address: 8888::8888/128
  V6 source refcnt: 1

  V6 source address: 7777::7777/128
```

```
    V6 source refcnt: 1

    V6 source address: 6666::6666/128
    V6 source refcnt: 1

{master}
regress@10.102.171.225>
```

show dynamic-tunnels pfe-tunnel-localization (Backup)

```
user@host> show dynamic-tunnels pfe-tunnel-localization
Anchor PFE Name: pfe-0/0/0
    Reference count: 16

    Gencfg keyid: 1
      Gencfg last index: 145
      Gencfg current index: 145
      Gencfg last state: UP
      Gencfg current state: UP
      Gencfg V6 source address: 66:66:00:00:00:00:00:00:00:00:00:00:00:00:66:66
      Gencfg backup lock: 1
      Gencfg kqp: 0x0
      Gencfg Reference count: 0x3

    Gencfg keyid: 2
      Gencfg last index: 145
      Gencfg current index: 145
      Gencfg last state: UP
      Gencfg current state: UP
      Gencfg V6 source address: 77:77:00:00:00:00:00:00:00:00:00:00:00:00:77:77
      Gencfg backup lock: 1
      Gencfg kqp: 0x0
      Gencfg Reference count: 0x3

    Gencfg keyid: 3
      Gencfg last index: 145
      Gencfg current index: 145
      Gencfg last state: UP
      Gencfg current state: UP
      Gencfg V6 source address: 88:88:00:00:00:00:00:00:00:00:00:00:00:00:88:88
      Gencfg backup lock: 1
```

```
      Gencfg kqp: 0x0
      Gencfg Reference count: 0x3

    Gencfg keyid: 4
      Gencfg last index: 145
      Gencfg current index: 145
      Gencfg last state: UP
      Gencfg current state: UP
      Gencfg V6 source address: 99:99:00:00:00:00:00:00:00:00:00:00:00:00:99:99
      Gencfg backup lock: 1
      Gencfg kqp: 0x0
      Gencfg Reference count: 0x3

    Aggregate: 1.1.1.0/24
    Aggregate refcnt: 1

    Aggregate: 2.2.2.0/24
    Aggregate refcnt: 1

    Aggregate: 3.3.3.0/24
    Aggregate refcnt: 1

    Aggregate: 4.4.4.0/24
    Aggregate refcnt: 1

    V6 source address: 8888::8888/128
    V6 source refcnt: 1

    V6 source address: 7777::7777/128
    V6 source refcnt: 1

    V6 source address: 6666::6666/128
    V6 source refcnt: 1

    V6 source address: 9999::9999/128
    V6 source refcnt: 1
```

show dynamic-tunnels pfe-tunnel-localization (Localization)

```
user@host> show dynamic-tunnels pfe-tunnel-localization
Anchor PFE Name: pfe-0/0/0
```

```
    Reference count: 12
    Gencfg keyid: 1
      Gencfg tunnel type: UDP
      Gencfg last index: 142
      Gencfg current index: 142
      Gencfg last state: UP
      Gencfg current state: UP
      Gencfg backup lock: 0
      Gencfg kqp: 0x0
      Gencfg Reference count: 14
```

Release Information

Command introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

- dynamic-tunnels | **6161**
- extended-nexthop | **6700**
- tunnel-attributes | **30194**
- show v4ov6-tunnels | **27544**
- *Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP*

show edge-virtual-bridging

IN THIS SECTION

- Syntax | **19617**
- Description | **19617**
- Options | **19617**
- Required Privilege Level | **19617**
- Output Fields | **19617**
- Sample Output | **19619**
- Release Information | **19620**

Syntax

```
show edge-virtual-bridging
<detail>
<edge-control-protocol statistics <interface interface-name>>
<firewall>
<interface interface-name>
vsi-profiles <interface interface-name>
```

Description

Display information about edge virtual bridging (EVB).

Options

none	Display EVB parameters for all interfaces configured with EVB.
detail	(Optional) Display EVB parameters and virtual station interface (VSI) profiles associated with each interface.
edge-control-protocol statistics <interface *<nterface-name>*	(Optional) Display Edge Control Protocol (ECP)l statistics for all configured EVB interfaces or for the specified interface.
firewall	Display the firewall filters created by EVB.
interface *interface-name*	(Optional) Display EVB parameters for the specified interface.
vsi-profiles <interface *interface-name*>	(Optional) Display VSI profiles associated on each interface or for the specified interface.

Required Privilege Level

view

Output Fields

Table 476 on page 19618 lists the output fields for the `show edge-virtual-bridging` command. Output fields are listed in the approximate order in which they appear.

Table 476: show edge-virtual-bridging Output Field Descriptions

Field Name	Field Description
Interface	Switch interface configured for EVB.
Interface input ECP Packets	Number of ECP packets received by the switch. ECP is a Layer 2 protocol that is used to carry VSI Discovery and Configuration Protocol (VDP) messages.
Interface output ECP Packets	Number of ECP packets sent by the switch. ECP is a Layer 2 protocol that is used to carry VDP messages.
Forwarding Mode	Mode by which packets are forwarded to their destination. The value for forwarding mode is either **Standard** (meaning the forwarding is done through 802.1Q) or **Reflective-relay**, meaning that both the source and destination addresses are located on the same VM server.
RTE	Retransmission timer exponent (RTE) is an EVB interface attribute used to calculate the minimum VDP protocol data unit (PDU) retransmission time.
Number of VSIs	Number of virtual station interfaces on the switch connected to the VEPA.
Protocols	EVB protocols currently enabled. The values can be **VDP**, **ECP** or **RTE**. Protocols are configured during the capabilities exchange via an EVB type, length, and value (TLV) carried by the Link Layer Discovery Protocol (LLDP) between the switch and the server.
VSI profile	EVB profile including parameters that uniquely identify each VSI entry (VSI manager, VSI type, VSI version, VSI instance, VSI state).
Filter Name	Name of the filter defined in the firewall stanza.
Counters	Number of packets and bytes that have satisfied the match conditions defined by the filter.

Sample Output

show edge-virtual-bridging

```
user@switch#show edge-virtual-bridging
Interface     Forwarding Mode      RTE   Number of VSIs  Protocols
ge-0/0/20.0   Reflective-relay     25    400             ECP, VDP, RTE
```

show edge-virtual-bridging interface

```
user@switch#show edge-virtual-bridging interface ge-0/0/20.0
Interface: ge-0/0/20.0, Forwarding mode: Reflective-relay RTE: 25, Number of VSIs: 400,
Protocols: ECP, VDP, RTE

VSI profiles:
 Manager: 97, Type: 997, Version: 3, VSI State: Associate
 Instance: 09b11c53-8b5c-4eeb-8f00-c84ebb0bb997
      MAC                       VLAN
      00:10:94:00:00:04
```

show edge-virtual-bridging edge-control-protocol statistics

```
user@switch#show edge-virtual-bridging edge-control-protocol-statistics
Interface: ge-0/0/20.0
      Input ECP packets: 302
      Output ECP packets: 303
```

show edge-virtual-bridging vsi-profiles

```
user@switch#show edge-virtual-bridging vsi-profiles
Interface: ge-0/0/20.0
Manager: 97, Type: 997, Version: 3, VSI State: Associate
 Instance: 09b11c53-8b5c-4eeb-8f00-c84ebb0bb997
      MAC                       VLAN
      00:10:94:00:00:04         3
```

show edge-virtual-bridging vsi-profiles interface

```
user@switch#show edge-virtual-bridging vsi-profiles interface ge-0/0/20.0
Interface: ge-0/0/20.0
Manager: 97, Type: 997, Version: 3, VSI State: Associate
 Instance: 09b11c53-8b5c-4eeb-8f00-c84ebb0bb997
          MAC
          VLAN   00:10:94:00:00:04         3
```

show edge-virtual-bridging firewall

```
user@switch#show edge-virtual-bridging firewall
Filter name: evb_filter_ge-0/0/20
Counters:
  Name: evb_filter_term_3_00:10:94:00:00:04_default
          Bytes: 0, Packets: 0
  Name: f3_accept__evb_filter_term_3_00:10:94:00:00:04-f3-t1
          Bytes: 1028, Packets: 14
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on an EX Series Switch

show ephemeral-configuration

IN THIS SECTION

- Syntax | **19621**
- Syntax (Junos OS Release 18.1 and Earlier) | **19621**
- Description | **19621**

- Options | 19621
- Required Privilege Level | 19622
- Sample Output | 19622
- Release Information | 19623

Syntax

```
show ephemeral-configuration (instance instance-name | merge)
```

Syntax (Junos OS Release 18.1 and Earlier)

```
show ephemeral-configuration
<instance-name>
```

Description

Display configuration data committed to the ephemeral configuration database.

Options

none In Junos OS Release 18.1 and earlier, display the configuration committed to the default instance of the ephemeral configuration database.

instance-name (Optional) Name of a user-defined ephemeral instance for which to display the committed ephemeral configuration data.

instance *instance-name* Display the configuration committed to an instance of the ephemeral configuration database.

- To display the configuration data in the default ephemeral instance, set the instance name to `default`.
- To display the configuration data for sensors that have been provisioned by an external collector to export data through gRPC, set the instance name to `junos-analytics`.

- To display the configuration data in a user-defined instance, specify the name of an instance configured at the `[edit system configuration-database ephemeral instance]` hierarchy level.

merge Display the configuration data in all instances of the ephemeral configuration database merged with the complete post-inheritance view of the static configuration database.

> **NOTE**: In Junos OS Release 18.1 and earlier, to display the configuration data in all instances of the ephemeral configuration database merged with the complete post-inheritance view of the static configuration database, use the `show ephemeral-configuration | display merge` command.

Required Privilege Level

view

Sample Output

show ephemeral-configuration (Junos OS Release 18.1 or earlier)

```
user@host> show ephemeral-configuration
## Last changed: 2017-02-12 17:15:48 PDT
protocols {
    mpls {
        label-switched-path to-cust1 {
            to 198.51.100.1;
        }
    }
}
```

show ephemeral-configuration eph1 (Junos OS Release 18.1 or earlier)

```
user@host> show ephemeral-configuration eph1
## Last changed: 2017-02-10 13:20:32 PDT
protocols {
    mpls {
        label-switched-path to-hastings {
            to 192.0.2.1;
```

```
        }
    }
}
```

show ephemeral-configuration instance eph1 (Junos OS Release 18.2R1 or later)

```
user@host> show ephemeral-configuration instance eph1
## Last changed: 2017-02-10 13:20:32 PDT
protocols {
    mpls {
        label-switched-path to-hastings {
            to 192.0.2.1;
        }
    }
}
```

Release Information

Command introduced in Junos OS Release 16.2R2 and Junos OS Evolved Release 22.1R1.

`instance` and `merge` options added in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Enabling and Configuring Instances of the Ephemeral Configuration Database

show ethernet-switching evpn nd-statistics

IN THIS SECTION

- Syntax | 19624
- Description | 19624
- Options | 19624
- Required Privilege Level | 19624

- Output Fields | 19624
- Sample Output | 19625
- Release Information | 19626

Syntax

```
show ethernet-switching evpn nd-statistics
<brief |count | detail | extensive | summary | display xml>
```

Description

Show Neighbor Discovery (ND) proxy table statistics for integrated routing and bridging (IRB) interfaces participating in the Ethernet VPN (EVPN). ND proxy is a kernel module that implements IPv6 Neighbor Discovery proxying over Ethernet-like access networks.

Options

none	Display information about ND proxy table statistics for all IRB interfaces participating in the EVPN.
brief | count | detail | extensive | summary | display xml	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 477 on page 19625 lists the output fields for the `show ethernet-switching evpn nd-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 477: show ethernet-switching evpn nd-statistics Output Fields

Field Name	Field Description	Level of Output
Interface	IRB interface on which the statistics has been applied	All levels
EVPN	Name of the EVPN	All levels
Bridge domain	Name of the bridge domain for the EVPN	All levels
ND routes add received	Total number of additional ND routes received	All levels
ND routes del received	Total number of deleted ND routes received	All levels
ND routes dropped	Total number of ND routes dropped	All levels
MAC+IPv6s sent to peer	Total number of MAC and IPv6 addresses sent to peer device	All levels

Sample Output

show ethernet-switching evpn nd-statistics

```
user@host> show ethernet-switching evpn nd-statistics
interface irb.0

Interface : irb.0       EVPN : evpn1    Bridge domain: vlan10
ND routes add received          : 2
ND routes del received          : 0
ND routes dropped               : 0
MAC+IPv6s sent to peer          : 3
```

Release Information

Command introduced in Junos OS Release 17.3R1.

show ethernet-switching evpn nd-table

IN THIS SECTION

- Syntax | 19626
- Description | 19626
- Options | 19626
- Required Privilege Level | 19627
- Output Fields | 19627
- Sample Output | 19628
- Release Information | 19630

Syntax

```
show ethernet-switching evpn nd-table
<brief | detail | extensive>
<count>
<instance-name>
<mac-address>
<vlan-name>
```

Description

Display information about INET entries associated with MAC addresses learned through Network Discovery Protocol (NDP).

Options

none	Display information for all INET entries.

brief | detail | extensive	(Optional) Display the specified level of output.
count	(Optional) Display the number of INET addresses learned in a routing instance.
instance	(Optional) Display information for a specified instance.
mac-adress	(Optional) Display information for a specified MAC address.
vlan-name (all)	(Optional) Display information for a specified VLAN or for all VLANs.

Required Privilege Level

view

Output Fields

Table 478 on page 19627 lists the output fields for the `show ethernet-switching evpn nd-table` command. Output fields are listed in the approximate order in which they appear.

Table 478: show ethernet-switching evpn nd-table Output Fields

Field Name	Field Description	Level of Output
INET address	The INET address related to the INET entries that are added to the NDP table.	All levels
MAC address	MAC addresses learned through NDP.	**brief**, **detail**, **extensive**, **instance**, **mac-addressvlan-name**,,
Logical Interface	Logical interface associated with the routing instance in which the NDP INET address is learned.	**brief**, **instance**, **mac-addressvlan-name**,,
Routing instance	Routing instance in which the NDP INET address is learned.	all levels
Bridging domain	Bridging domain in which the NDP INET address is learned.	all levels
Learning interface	Interface on which the NDP INET address is learned.	**detail**, **extensive**

Table 478: show ethernet-switching evpn nd-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Count	Indicates the number of NDP INET addresses learned in a routing instance in a bridge domain.	**count**

Sample Output

show ethernet-switching evpn nd-table brief

```
user@switch> show ethernet-switching evpn nd-table
brief
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
8002::2             00:05:86:a0:d5:00  irb.0         evpn1      vlan10
```

show ethernet-switching evpn nd-table detail

```
user@switch> show ethernet-switching evpn nd-table
detail
INET address: 8002::2
MAC address: 00:05:86:a0:d5:00
  Routing instance: evpn1
   Bridging domain: vlan10
   Learning interface: irb.0
```

show ethernet-switching evpn nd-table extensive

```
user@switch> show ethernet-switching evpn nd-table
extensive
INET address: 8002::2
MAC address: 00:05:86:a0:d5:00
  Routing instance: evpn1
   Bridging domain: vlan10
   Learning interface: irb.0
```

show ethernet-switching evpn nd-table count

```
user@switch> show ethernet-switching evpn nd-table
count
1 ND INET addresses learned in routing instance evpn1 bridge domain vlan10
```

show ethernet-switching evpn nd-table instance evpn1

```
user@switch> show ethernet-switching evpn nd-table
instance evpn1
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
8002::2
            00:05:86:a0:d5:00
  irb.0         evpn1        vlan10
```

show ethernet-switching evpn nd-table instance 00:05:86:90:bd:f0 (MAC address)

```
user@switch> show ethernet-switching evpn nd-table
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
8002::2
            00:05:86:a0:d5:00
  irb.0         evpn1        __evpn1__
```

show ethernet-switching evpn nd-table vlan-name vlan10

```
user@switch>show ethernet-switching evpn nd-table
vlan-name vlan10

INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
8002::2            00:05:86:a0:d5:00
  irb.0         evpn1        vlan10
```

Release Information

Command introduced in Junos OS Release 17.4R2.

> **NOTE**: Starting with Junos OS Release 17.4R2, the `show ethernet-switching evpn nd-table` command replaces the `show vlans evpn nd-table` command.

show ethernet-switching flood

IN THIS SECTION

- Syntax | 19630
- Description | 19630
- Options | 19631
- Required Privilege Level | 19631
- Sample Output | 19632
- Release Information | 19636

Syntax

```
show ethernet-switching flood
<brief | detail | extensive>
<event-queue>
<instance instance-name>
<logical-system logical-system-name>
<route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-flood)>
<vlan-name vlan-name>
```

Description

(EX Series switches and QFX Series switches only) Display Ethernet-switching flooding information.

Options

none	Display all Ethernet-switching flooding information for all VLANs.
brief | detail | extensive	(Optional) Display the specified level of output.
event-queue	(Optional) Display the queue of pending Ethernet-switching flood events.
instance *instance-name*	(Optional) Display Ethernet-switching flooding information for the specified routing instance.
logical-system *logical-system-name*	(Optional) Display Ethernet-switching flooding information for the specified logical system.
route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-flood)	(Optional) Display the following: • **all-ce-flood**—Display the route for flooding traffic to all customer edge routers or switches if **no-local-switching** is enabled. • **all-ve-flood**—Display the route for flooding traffic to all VPLS edge routers or switches if **no-local-switching** is enabled. • **alt-root-flood**—Display the Spanning Tree Protocol (STP) alt-root flooding route used for the interface. • **bd-flood**—Display the route for flooding traffic of a VLAN if **no-local-switching** is not enabled. • **mlp-flood**—Display the route for flooding traffic to MAC learning chips. • **re-flood**—Display the route for Routing Engine flooding to all interfaces.
vlan-name *vlan-name*	(Optional) Display Ethernet-switching flooding information for the specified VLAN.

Required Privilege Level

view

Sample Output

show ethernet-switching flood

```
user@host> show ethernet-switching flood
Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
Flood Routes:
  Prefix    Type          Owner                 NhType          NhIndex
  0x3057b/51 FLOOD_GRP_COMP_NH __all_ces__      comp            12866
  0x30004/51 FLOOD_GRP_COMP_NH __re_flood__     comp            12863
VLAN Name: VLAN102
Flood Routes:
  Prefix    Type          Owner                 NhType          NhIndex
  0x3057c/51 FLOOD_GRP_COMP_NH __all_ces__      comp            12875
  0x30005/51 FLOOD_GRP_COMP_NH __re_flood__     comp            12872
VLAN Name: VLAN103
Flood Routes:
  Prefix    Type          Owner                 NhType          NhIndex
  0x3057d/51 FLOOD_GRP_COMP_NH __all_ces__      comp            12884
  0x30006/51 FLOOD_GRP_COMP_NH __re_flood__     comp            12881
```

show ethernet-switching flood brief

```
user@host> show ethernet-switching flood brief
Name                Active CEs         Active VEs
__juniper_private1__ 0                 0
default-switch      9                  0
```

show ethernet-switching flood detail

```
user@host> show ethernet-switching flood detail
Name: __juniper_private1__
```

```
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
Flood Routes:
  Prefix    Type          Owner                 NhType          NhIndex
  0x3057b/51 FLOOD_GRP_COMP_NH __all_ces__      comp            12866
  0x30004/51 FLOOD_GRP_COMP_NH __re_flood__     comp            12863
VLAN Name: VLAN102
Flood Routes:
  Prefix    Type          Owner                 NhType          NhIndex
  0x3057c/51 FLOOD_GRP_COMP_NH __all_ces__      comp            12875
  0x30005/51 FLOOD_GRP_COMP_NH __re_flood__     comp            12872
VLAN Name: VLAN103
Flood Routes:
  Prefix    Type          Owner                 NhType          NhIndex
  0x3057d/51 FLOOD_GRP_COMP_NH __all_ces__      comp            12884
  0x30006/51 FLOOD_GRP_COMP_NH __re_flood__     comp            12881
```

show ethernet-switching flood extensive

```
user@host> show ethernet-switching flood extensive
Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
  Flood route prefix: 0x3057b/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __all_ces__
  Flood group name: __all_ces__
  Flood group index: 1
  Nexthop type: comp
  Nexthop index: 12866
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12860
```

```
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605

  Flood route prefix: 0x30004/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __re_flood__
  Flood group name: __re_flood__
  Flood group index: 65534
  Nexthop type: comp
  Nexthop index: 12863
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12860
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605
VLAN Name: VLAN102

  Flood route prefix: 0x3057c/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __all_ces__
  Flood group name: __all_ces__
  Flood group index: 1
  Nexthop type: comp
  Nexthop index: 12875
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12869
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605

  Flood route prefix: 0x30005/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __re_flood__
  Flood group name: __re_flood__
  Flood group index: 65534
  Nexthop type: comp
  Nexthop index: 12872
```

```
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12869
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605
VLAN Name: VLAN103

  Flood route prefix: 0x3057d/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __all_ces__
  Flood group name: __all_ces__
  Flood group index: 1
  Nexthop type: comp
  Nexthop index: 12884
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12878
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605

  Flood route prefix: 0x30006/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __re_flood__
  Flood group name: __re_flood__
  Flood group index: 65534
  Nexthop type: comp
  Nexthop index: 12881
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12878
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605
VLAN Name: VLAN104
```

Release Information

Command introduced in Junos OS Release 12.3R2.

show ethernet-switching interface

IN THIS SECTION

- Syntax | 19636
- Description | 19636
- Options | 19636
- Required Privilege Level | 19637
- Output Fields | 19637
- Sample Output | 19638
- Release Information | 19640

Syntax

```
show ethernet-switching interface
<brief | detail | extensive>
<interface-name>
```

Description

Display Layer 2 learning information for all the interfaces.

Options

none	Display Ethernet-switching information for all interfaces.
brief | detail | extensive	(Optional) Display the specified level of output.
interface-name	(Optional) Display Ethernet-switching information for the specified interface.

Required Privilege Level

view

Output Fields

Table 479 on page 19637 describes the output fields for the `show ethernet-switching interface` command. Output fields are listed in the approximate order in which they appear.

Table 479: show ethernet-switching interface Output Fields

Field Name	Field Description
Logical interface	Name of the logical interface.
VLAN members	VLANs associated with this interface.
Tag	VLAN ID.
MAC limit	Number of MAC addresses that can be associated with the interface.
STP state	Spanning Tree protocol (STP) state.
Logical interface flags	Status of Layer 2 learning properties for each interface: • `DL`—MAC learning is disabled. • `LH`—MAC interface limit has been reached. • `AD`—Packets are dropped after the MAC interface limit is reached. • `DN`—The MAC interface is down. • `MMAS`—The MAC interface is disabled after a MAC address move. • `SCTL`—The MAC interface is disabled after a configured storm-control level is exceeded. **NOTE**: If the physical interface is shutdown due to storm control, all logical interfaces on the shutdown interface display the SCTL logical interface flag.

Table 479: show ethernet-switching interface Output Fields *(Continued)*

Field Name	Field Description
Tagging	Tagging state of the VLAN.

Sample Output

show ethernet switching interface (Specific Interface)

```
user@host> show ethernet-switching inerface ae10.0
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down)

Logical      Vlan       TAG   MAC      STP          Logical         Tagging
interface    members          limit    state        interface flags
ae10.0                        8192                                   tagged
             VLAN70..   701   1024     Forwarding

             VLAN70..   702   1024     Forwarding

             VLAN70..   703   1024     Forwarding

             VLAN70..   704   1024     Forwarding

             VLAN70..   705   1024     Forwarding

             VLAN70..   706   1024     Forwarding

             VLAN70..   707   1024     Forwarding

             VLAN70..   708   1024     Forwarding

             VLAN70..   709   1024     Forwarding

             VLAN71..   710   1024     Forwarding

             VLAN71..   711   1024     Forwarding

             VLAN71..   712   1024     Forwarding
```

```
             VLAN71..   713   1024     Forwarding

             VLAN71..   714   1024     Forwarding

             VLAN71..   715
[...output truncated...]
```

show ethernet switching interface (Storm Control in Effect)

```
user@host> show ethernet-switching interface ge-0/0/2.0
Logical Interface flags (DL - disable learning, AD - packet action drop, LH -
                         MAC limit hit, DN - interface down, MMAS - Mac-move
                         action shutdown, AS - Autostate-exclude enabled, SCTL
                         - shutdown by Storm-control, MI - MAC+IP limit hit)

Logical      Vlan       TAG   MAC    MAC+IP STP         Logical          Tagging
interface    members          limit  limit  state       interface flags
ge-0/0/2.0                    524287 8192               DN,SCTL          untagged
             VLAN1      100   65535  1024   Forwarding                   untagged
```

show ethernet-switching interface detail

```
user@host> show ethernet-switching interface detail

Information for interface family:
Name: ge-1/0/3.0
  Type: IFF                             Handle: 0x8bba280
  Index: 331                            Generation: 159
                                        Flags: UP,
  IFD index: 141                        Routing/Vlan index: 4
  IFL index: 331                        Address family: 50
  Sequence number: 0                    MAC sequence number: 0
  MAC limit: 65535                      MACs learned: 0
  Static MACs learned: 0                Non configured static MACs learned: 0
Name: ge-1/0/3.0
  Type: IFBD (static)                   Handle: 0x8bb6e00
  Index:                                Generation: 129
  Trunk id: 0                           Flags: UP,
  IFD index:                            Routing/Vlan index: 2
```

```
  IFL index:                            Address family:
  Sequence number: 1                    MAC sequence number: 1
  MAC limit: 65535                      MACs learned: 0
  Static MACs learned: 0                Non configured static MACs learned: 0
  VSTP index: 11                        Rewrite op:
Name: ge-1/0/3.0
  Type: IFBD (static)                   Handle: 0x8bb6f00
  Index:                                Generation: 130
  Trunk id: 0                           Flags: UP,
  IFD index:                            Routing/Vlan index: 3
  IFL index:                            Address family:
  Sequence number: 1                    MAC sequence number: 1
  MAC limit: 65535                      MACs learned: 0
  Static MACs learned: 0                Non configured static MACs learned: 0
  VSTP index: 11                        Rewrite op:
```

show ethernet-switching interface xe-0/0/2.0 (autostate-exclude enabled on QFX5100 switch)

```
user@switch> show ethernet-switching interface xe-0/0/2.0
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown, AS - Autostate-exclude enabled
                         SCTL - shutdown by Storm-control)

Logical          Vlan          TAG     MAC         STP         Logical           Tagging
interface        members               limit       state       interface flags
xe-0/0/2.0                             294912                   AS                 tagged
                 v100          100     294912      Forwarding                      tagged
```

Release Information

Command introduced in Junos OS Release 12.3R2.

show ethernet-switching interfaces

IN THIS SECTION

- Syntax | 19641
- Description | 19641
- Options | 19641
- Required Privilege Level | 19641
- Output Fields | 19642
- Sample Output for QFX Series Switches, QFabric, NFX Series, EX4600 and OCX1100 | 19646
- Sample Output for EX Series Switches | 19649
- Release Information | 19651

Syntax

```
show ethernet-switching interfaces
<brief | detail | summary>
<interface interface-name>
```

Description

Display information about switched Ethernet interfaces.

Options

none	(Optional) Display brief information for Ethernet-switching interfaces.
brief | detail | summary	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display Ethernet-switching information for a specific interface.

Required Privilege Level

view

Output Fields

For QFX Series, QFabric, NFX Series, EX4600 and OCX1100:

Table 480 on page 19642 lists the output fields for the `show ethernet-switching interfaces` command on QFX Series, QFabric, NFX Series, EX4600 and OCX1100. Output fields are listed in the approximate order in which they appear.

Table 480: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	All levels
State	Interface state. Values are **up** or **down**.	none, **brief**, **detail**, **summary**
VLAN members	Name of a VLAN.	none, **brief**, **detail**, **summary**
Blocking	Forwarding state of the interface: • **blocked**—Traffic is not being forwarded on the interface. • **unblocked**—Traffic is forwarded on the interface. • **MAC limit exceeded**—The interface is temporarily disabled because of a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **MAC move limit exceeded**—The interface is temporarily disabled because of a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **Storm control in effect** —The interface is temporarily disabled because of a storm control error. The disabled interface is automatically restored to service when the disable timeout expires. • **Storm control shutdown in effect** —The interface is temporarily disabled because of a storm control shutdown error. The disabled interface is automatically restored to service when the disable timeout expires.	none, **brief**, **detail**, **summary**

Table 480: show ethernet-switching interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Index	VLAN index internal to Junos OS software.	**detail**
untagged | tagged	Specifies whether the interface forwards IEEE802.1Q-tagged or untagged traffic.	**detail**

Output fields for EX Series:

Table 481 on page 19643 lists the output fields for the `show ethernet-switching interfaces` command on EX Series switches. Output fields are listed in the approximate order in which they appear.

Table 481: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	none, **brief**, **detail**, **summary**
Index	VLAN index internal to Junos OS.	**detail**
State	Interface state. Values are **up** and **down**.	none, **brief**, **detail**
Port mode	The **access** mode is the port mode default and works with a single VLAN. Port mode can also be **trunk**, which accepts tagged packets from multiple VLANs on other switches. The third port mode value is **tagged-access**, which accepts tagged packets from access devices.	**detail**
Reflective Relay Status	Reflective relay allows packets to use the same interface for both upstream and downstream traffic. When reflective relay has been configured, the status displayed is always **enabled** . When reflective relay is not configured, this entry does not appear in the command output.	**detail**

Table 481: show ethernet-switching interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Ether type for the interface	Ether type is a two-octet field in an Ethernet frame used to indicate which protocol is encapsulated in the payload of an incoming Ethernet packet. Both 802.1Q packets and Q-in-Q packets use this field. The output displayed for this particular field indicates the interface's Ether type, which is used to match the Ether type of incoming 802.1Q packets and Q-in-Q packets. The indicated Ether type field is also added to the interface's outgoing 802.1Q and Q-in-Q packets.	**detail**
VLAN membership	Names of VLANs that belong to this interface.	none, **brief**, **detail**,
Tag	Number of the 802.1Q tag.	none, **brief**, **detail**,
Tagging	Specifies whether the interface forwards 802.1Q **tagged** or **untagged** traffic.	none, **brief**, **detail**,

Table 481: show ethernet-switching interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Blocking	The forwarding state of the interface: • **unblocked**—Traffic is forwarded on the interface. • **blocked**—Traffic is not being forwarded on the interface. • **Disabled by bpdu control**—The interface is disabled due to receiving BPDUs on a protected interface. If the `disable-timeout` statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires. • **blocked by RTG**—The specified redundant trunk group is disabled. • **blocked by STP**—The interface is disabled due to a spanning-tree protocol error. • **MAC limit exceeded**—The interface is temporarily disabled due to a MAC limit error. The disabled interface is automatically restored to service when the disable timeout expires. • **MAC move limit exceeded**—The interface is temporarily disabled due to a MAC move limit error. The disabled interface is automatically restored to service when the disable timeout expires. • **Storm control in effect**—The interface is temporarily disabled due to a storm control error. The disabled interface is automatically restored to service when the disable timeout expires.	none, **brief**, **detail**,
Number of MACs learned on IFL	Number of MAC addresses learned by this interface.	**detail**

Table 481: show ethernet-switching interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
mapping	When mapping is configured, the status is one of the following C-VLAN to S-VLAN mapping types: • **dot1q-tunneled**—The interface maps all traffic to the S-VLAN (all-in-one bundling). • **native**—The interface maps untagged and priority tagged packets to the S-VLAN. • **push**—The interface maps packets to a firewall filter to an S-VLAN. • **policy-mapped**—The interface maps packets to a specifically defined S-VLAN. • *integer*—The interface maps packets to the specified S-VLAN. When mapping is not configured, this entry does not appear in the command output.	**detail**

Sample Output for QFX Series Switches, QFabric, NFX Series, EX4600 and OCX1100

show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces

Interface   State    VLAN members           Blocking
xe-0/0/0.0  up       T1122                  unblocked
xe-0/0/1.0  down     default                — MAC limit exceeded
xe-0/0/2.0  down     default                — MAC move limit exceeded
xe-0/0/3.0  down     default                — Storm control in effect
xe-0/0/4.0  down     default                unblocked
xe-0/0/5.0  down     default                unblocked
xe-0/0/6.0  down     default                unblocked
xe-0/0/7.0  down     default                unblocked
xe-0/0/8.0  down     default                unblocked
xe-0/0/9.0  up       T111                   unblocked
xe-0/0/10.0 down     default                unblocked
xe-0/0/11.0 down     default                unblocked
```

```
xe-0/0/12.0 down     default                unblocked
xe-0/0/13.0 down     default                unblocked
xe-0/0/14.0 down     default                unblocked
xe-0/0/15.0 down     default                unblocked
xe-0/0/16.0 down     default                unblocked
xe-0/0/17.0 down     default                unblocked
xe-0/0/18.0 down     default                unblocked
xe-0/0/19.0 up       T111                   unblocked
xe-0/1/0.0  down     default                unblocked
xe-0/1/1.0  down     default                unblocked
xe-0/1/2.0  down     default                unblocked
xe-0/1/3.0  down     default                unblocked
```

show ethernet-switching interfaces summary

```
user@switch> show ethernet-switching interfaces summary
xe-0/0/0.0
xe-0/0/1.0
xe-0/0/2.0
xe-0/0/3.0
xe-0/0/8.0
xe-0/0/10.0
xe-0/0/11.0
```

show ethernet-switching interfaces brief

```
user@switch> show ethernet-switching interfaces brief
Interface   State    VLAN members           Blocking
xe-0/0/0.0  down     default                unblocked
xe-0/0/1.0  down     employee-vlan          unblocked
xe-0/0/2.0  down     employee-vlan          unblocked
xe-0/0/3.0  down     employee-vlan          unblocked
xe-0/0/8.0  down     employee-vlan          unblocked
xe-0/0/10.0 down     default                unblocked
xe-0/0/11.0 down     employee-vlan          unblocked
```

show ethernet-switching interfaces detail

```
user@switch> show ethernet-switching interfaces detail
Interface: xe-0/0/0.0 Index: 65
  State: down
  VLANs:
    default                untagged     unblocked

Interface: xe-0/0/1.0 Index: 66
  State: down
  VLANs:
    employee-vlan          untagged     unblocked

Interface: xe-0/0/2.0 Index: 67
  State: down
  VLANs:
    employee-vlan          untagged     unblocked

Interface: xe-0/0/3.0 Index: 68
  State: down
  VLANs:
    employee-vlan          untagged     unblocked

Interface: xe-0/0/8.0 Index: 69
  State: down
  VLANs:
    employee-vlan          untagged     unblocked

Interface: xe-0/0/10.0 Index: 70
  State: down
  VLANs:
    default                untagged     unblocked

Interface: xe-0/0/11.0 Index: 71
  State: down
  VLANs:
    employee-vlan          tagged       unblocked
```

show ethernet-switching interfaces interface-name

```
user@switch> show ethernet-switching interfaces xe-0/0/0.0
 Interface   State    VLAN members           Blocking
xe-0/0/0.0  down     default                unblocked
```

Sample Output for EX Series Switches

show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces

Interface    State  VLAN members        Tag   Tagging  Blocking

ae0.0        up      default                   untagged unblocked
ge-0/0/2.0   up      vlan300            300    untagged blocked by RTG (rtggroup)
ge-0/0/3.0   up      default                            blocked by STP
ge-0/0/4.0   down    default                            MAC limit exceeded
ge-0/0/5.0   down    default                            MAC move limit exceeded
ge-0/0/6.0   down    default                            Storm control in effect
ge-0/0/7.0   down    default                            unblocked
ge-0/0/13.0  up      default                   untagged unblocked
ge-0/0/14.0  up      vlan100             100   tagged   unblocked
                     vlan200             200   tagged   unblocked
ge-0/0/15.0  up      vlan100             100   tagged   blocked by STP
                     vlan200             200   tagged   blocked by STP
ge-0/0/16.0  down    default                   untagged unblocked
ge-0/0/17.0  down    vlan100             100   tagged   Disabled by bpdu-control
                     vlan200             200   tagged   Disabled by bpdu-control
```

show ethernet-switching interfaces ge-0/0/15 brief

```
user@switch> show ethernet-switching interfaces ge-0/0/15 brief

Interface   State  VLAN members   Tag   Tagging        Blocking

ge-0/0/15.0  up     vlan100        100   tagged         blocked by STP
                    vlan200        200   tagged         blocked by STP
```

show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup)

```
user@switch> show ethernet-switching interfaces ge-0/0/2 detail

Interface: ge-0/0/2.0, Index: 65, State: up, Port mode: Access
Ether type for the interface: 0X8100
VLAN membership:
    vlan300, 802.1Q Tag: 300, untagged, msti-id: 0, blocked by RTG(rtggroup)
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP)

```
user@switch> show ethernet-switching interfaces ge-0/0/15 detail

Interface: ge-0/0/15.0, Index: 70, State: up, Port mode: Trunk
Ether type for the interface: 0X8100
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 0, blocked by STP
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 0, blocked by STP

Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control)

```
user@switch> show ethernet-switching interfaces ge-0/0/17 detail

Interface: ge-0/0/17.0, Index: 71, State: down, Port mode: Trunk
Ether type for the interface: 0X8100
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 1, Disabled by bpdu-control
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 2, Disabled by bpdu-control
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping)

```
user@switch>show ethernet-switching interfaces ge-0/0/6.0 detail
Interface: ge-0/0/6.0, Index: 73, State: up, Port mode: Access
Ether type for the interface: 0X8100
VLAN membership:
    map, 802.1Q Tag: 134, Mapped Tag: native, push, dot1q-tunneled, unblocked
    map, 802.1Q Tag: 134, Mapped Tag: 20, push, dot1q-tunneled, unblocked
```

show ethernet-switching interfaces detail (Reflective Relay Is Configured)

```
user@switch1> show ethernet-switching interfaces ge-7/0/2 detail
Interface: ge-7/0/2, Index: 66, State: down, Port mode: Tagged-access
Ether type for the interface: 0X8100
Reflective Relay Status: Enabled
Ether type for the interface: 0x8100
VLAN membership:
    VLAN_Purple VLAN_Orange VLAN_Blue, 802.1Q Tag: 450, tagged, unblocked
Number of MACs learned on IFL: 0
```

Release Information

Command introduced in Junos OS Release 9.0.

In Junos OS Release 9.6 for EX Series switches, the following updates were made:

- **Blocking** field output was updated.
- The default view was updated to include information about 802.1Q tags.
- The **detail** view was updated to include information on VLAN mapping.

In Junos OS Release 11.1 for EX Series switches, the **detail** view was updated to include reflective relay information.

RELATED DOCUMENTATION

Troubleshooting Ethernet Switching

Understanding Bridging and VLANs on Switches

Example: Setting Up Basic Bridging and a VLAN on Switches

Example: Setting Up Bridging with Multiple VLANs

Understanding FCoE

Interfaces Overview for Switches

show ethernet-switching mac-learning-log | **19670**

show ethernet-switching table | **19700**

Configuring Autorecovery for Port Security Events

show ethernet-switching layer2-protocol-tunneling interface

IN THIS SECTION

- Syntax | **19652**
- Description | **19652**
- Options | **19653**
- Required Privilege Level | **19653**
- Output Fields | **19653**
- Sample Output | **19653**
- Release Information | **19654**

Syntax

```
show ethernet-switching-layer2-protocol-tunneling interface
<interface-name>
```

Description

Display information about Layer 2 protocol tunneling (L2PT) on interfaces that have been configured for L2PT.

Options

none	Display L2PT information about all interfaces on which L2PT is enabled.
interface-name	(Optional) Display L2PT information for the specified interface.

Required Privilege Level

view

Output Fields

Table 482 on page 19653 lists the output fields for the `show ethernet-switching layer2-protocol-tunneling interface` command. Output fields are listed in the approximate order in which they appear.

Table 482: show ethernet-switching layer2-protocol-tunneling interface Output Fields

Field Name	Field Description
Interface	Name of an interface on the switch.
Operation	Type of operation being performed on the interface. Values are **Encapsulation** and **Decapsulation**.
State	State of the interface. Values are **active** and **shutdown**.
Description	If the interface state is **shutdown**, displays why the interface is shut down. If the description says **Loop detected**, it means that the interface is an access interface that has received L2PT-enabled PDUs. Access interfaces should not receive L2PT-enabled PDUs. This scenario might mean that there is a loop in the network.

Sample Output

show ethernet-switching layer2-protocol-tunneling interface

```
user@switch> show ethernet-switching layer2-protocol-tunneling interface

Layer2 Protocol Tunneling information:
Interface     Operation       State      Description
```

```
xe-0/0/0.0     Encapsulation   Shutdown   Shutdown threshold exceeded
xe-0/0/1.0     Decapsulation   Shutdown   Loop detected
xe-0/0/2.0     Decapsulation   Active
```

show ethernet-switching layer2-protocol-tunneling interface xe-0/0/0.0

```
user@switch> show ethernet-switching layer2-protocol-tunneling interface xe-0/0/0.0

Layer2 Protocol Tunneling information:
Interface      Operation       State      Description
xe-0/0/0.0     Encapsulation   Shutdown   Shutdown threshold exceeded
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show ethernet-switching layer2-protocol-tunneling statistics | **19654**

show ethernet-switching layer2-protocol-tunneling vlan | **19658**

Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

show ethernet-switching layer2-protocol-tunneling statistics | **19654**

show ethernet-switching layer2-protocol-tunneling vlan | **19658**

show ethernet-switching layer2-protocol-tunneling statistics

IN THIS SECTION

- Syntax | **19655**
- Description | **19655**
- Options | **19655**
- Required Privilege Level | **19655**
- Output Fields | **19655**

Sample Output | 19656
Release Information | 19657

Syntax

```
show ethernet-switching-layer2-protocol-tunneling statistics
<interface interface-name>
<vlan vlan-name>
```

Description

Display Layer 2 protocol tunneling (L2PT) statistics for Layer 2 PDU packets received by the switch.

NOTE: The `show ethernet-switching-layer2-protocol-tunneling statistics` command does not display L2PT statistics for Layer 2 PDU packets transmitted from the switch.

Options

none	Display L2PT statistics for all interfaces on which you enabled L2PT.
interface *interface-name*	(Optional) Display L2PT statistics for the specified interface.
vlan *vlan-name*	(Optional) Display L2PT statistics for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 483 on page 19656 lists the output fields for the `show ethernet-switching layer2-protocol-tunneling statistics` command. Output fields are listed in the approximate order in which they appear.

Table 483: show ethernet-switching layer2-protocol-tunneling statistics Output Fields

VLAN	Field Description
VLAN	Name of a VLAN on which L2PT has been configured.
Interface	Name of an interface on which L2PT has been configured.
Protocol	Name of a protocol for which L2PT has been enabled. Values are **all**, **802.1x**, **802.3ah**, **cdp**, **e-lmi**, **gvrp**, **lacp**, **lldp**, **mmrp**, **mvrp**, **stp**, **udld**, **vstp**, and **vtp**.
Operation	Type of operation being performed on the interface. Values are **Encapsulation** and **Decapsulation**.
Packets	Number of packets that have been encapsulated or de-encapsulated.
Drops	Number of packets that have exceeded the drop threshold and have been dropped.
Shutdowns	Number of times that packets have exceeded the shutdown threshold and the interface has been shut down.

Sample Output

show ethernet-switching layer2-protocol-tunneling statistics

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics

Layer2 Protocol Tunneling Statistics:
VLAN   Interface  Protocol  Operation     Packets Drops      Shutdowns
v1     xe-0/0/0.0 mvrp      Encapsulation 0       0          0
v1     xe-0/0/1.0 mvrp      Decapsulation 0       0          0
v1     xe-0/0/2.0 mvrp      Decapsulation 60634   0          0
v2     xe-0/0/0.0 cdp       Encapsulation 0       0          0
v2     xe-0/0/0.0 gvrp      Encapsulation 0       0          0
v2     xe-0/0/0.0 lldp      Encapsulation 0       0          0
```

show ethernet-switching layer2-protocol-tunneling statistics interface xe-0/0/0.0

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics interface xe-0/0/0.0

Layer2 Protocol Tunneling Statistics:
VLAN    Interface  Protocol Operation      Packets  Drops     Shutdowns
v1      xe-0/0/0.0 mvrp     Encapsulation  0        0          0
v2      xe-0/0/0.0 cdp      Encapsulation  0        0          0
v2      xe-0/0/0.0 gvrp     Encapsulation  0        0          0
v2      xe-0/0/0.0 lldp     Encapsulation  0        0          0
v2      xe-0/0/0.0 mvrp     Encapsulation  0        0          0
v2      xe-0/0/0.0 stp      Encapsulation  0        0          0
v2      xe-0/0/0.0 vtp      Encapsulation  0        0          0
v2      xe-0/0/0.0 vstp     Encapsulation  0        0          0
```

show ethernet-switching layer2-protocol-tunneling statistics vlan v2

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics vlan v2

Layer2 Protocol Tunneling Statistics:
VLAN    Interface   Protocol Operation       Packets  Drops   Shutdowns
v2      xe-0/0/0.0  cdp      Encapsulation   0        0       0
v2      xe-0/0/0.0  gvrp     Encapsulation   0        0       0
v2      xe-0/0/0.0  lldp     Encapsulation   0        0       0
v2      xe-0/0/0.0  mvrp     Encapsulation   0        0       0
v2      xe-0/0/0.0  stp      Encapsulation   0        0       0
v2      xe-0/0/0.0  vtp      Encapsulation   0        0       0
v2      xe-0/0/0.0  vstp     Encapsulation   0        0       0
v2      xe-0/0/1.0  cdp      Decapsulation   0        0       0
v2      xe-0/0/1.0  gvrp     Decapsulation   0        0       0
v2      xe-0/0/1.0  lldp     Decapsulation   0        0       0
v2      xe-0/0/1.0  mvrp     Decapsulation   0        0       0
v2      xe-0/0/1.0  stp      Decapsulation   0        0       0
v2      xe-0/0/1.0  vtp      Decapsulation   0        0       0
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

clear ethernet-switching layer2-protocol-tunneling statistics | 3577

show ethernet-switching layer2-protocol-tunneling interface | 19652

show ethernet-switching layer2-protocol-tunneling vlan | 19658

Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

show vlans | 27738

show ethernet-switching layer2-protocol-tunneling vlan

IN THIS SECTION

- Syntax | 19658
- Description | 19658
- Options | 19659
- Required Privilege Level | 19659
- Output Fields | 19659
- Sample Output | 19659
- Release Information | 19660

Syntax

```
show ethernet-switching-layer2-protocol-tunneling vlan <vlan-name>
```

Description

Display information about Layer 2 protocol tunneling (L2PT) on VLANs that have been configured for L2PT.

Options

none Display information about L2PT for the VLANs on which you have configured L2PT.

vlan-name (Optional) Display information about L2PT for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 484 on page 19659 lists the output fields for the `show ethernet-switching layer2-protocol-tunneling vlan` command. Output fields are listed in the approximate order in which they appear.

Table 484: show ethernet-switching layer2-protocol-tunneling vlan Output Fields

Field Name	Field Description
VLAN	Name of the VLAN on which L2PT has been configured.
Protocol	Name of a protocol for which L2PT has been enabled. Values are **all**, **802.1x**, **802.3ah**, **cdp**, **e-lmi**, **gvrp**, **lacp**, **lldp**, **mmrp**, **mvrp**, **stp**, **vstp**, and **vtp**.
Drop Threshold	Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the VLAN before the switch begins dropping the Layer 2 PDUs.
Shutdown Threshold	Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the VLAN before the interface is disabled.

Sample Output

show ethernet-switching layer2-protocol-tunneling vlan

```
user@switch> show ethernet-switching layer2-protocol-tunneling vlan

Layer2 Protocol Tunneling VLAN information:
VLAN          Protocol      Drop       Shutdown
```

```
                              Threshold  Threshold
v1            mvrp            100        200
v2            cdp             0          0
v2            cdp             0          0
v2            gvrp            0          0
```

show ethernet-switching layer2-protocol-tunneling vlan v2

```
user@switch> show ethernet-switching layer2-protocol-tunneling vlan v2

Layer2 Protocol Tunneling VLAN information:
VLAN          Protocol      Drop       Shutdown
                            Threshold  Threshold
v2            cdp           0          0
v2            cdp           0          0
v2            gvrp          0          0
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- show ethernet-switching layer2-protocol-tunneling interface | **19652**
- show ethernet-switching layer2-protocol-tunneling statistics | **19654**
- *Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support*
- *Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support*
- show vlans | **27738**

show ethernet-switching mac-ip-table

IN THIS SECTION

- Syntax | **19661**

- Description | 19661
- Options | 19661
- Required Privilege Level | 19661
- Output Fields | 19662
- Sample Output | 19665
- Release Information | 19669

Syntax

```
show ethernet-switching mac-ip-table
<brief | count | detail | detail | extensive>
<instance>
<kernel >
<vlan-name (all | vlan-name)>
```

Description

Displays the MAC-IP address for all IPv4 (ARP) and IPv6 (ND) bindings for VLANs in routing instances where the instance-type is ethernet-switching.

Options

brief | detail |extensive | summary	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display the MAC-IP address information for the specified routing instance.
kernel	(Optional) Display the MAC-IP address entries in the kernel
vlan-name (all | *vlan-name*)	(Optional) Display the MAC-IP address entries for all VLANs or for a specified VLAN.

Required Privilege Level

maintenance

Output Fields

Table 485 on page 19662 lists the output fields for the `show ethernet-switching mac-ip-table` command. Output fields are listed in the approximate order in which they appear.

Table 485: show ethernet-switching mac-ip-table Output Fields

Field Name	Field Description	Level of Output
IP address	IP address associated with the ethernet switching instance type.	All levels
MAC address	MAC address	All levels

Table 485: show ethernet-switching mac-ip-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	MAC IP flags: Identifies : statically installed MAC-IP entries. • S—Statically installed MAC-IP entries. • D—Dynamic installed MAC-IP address entries. • L—Locally learned MAC-IP address entries • R— MAC-IP address entries that are learnt from a remote device via the control plane. • Lp—Local proxy MAC IP address entries. In multihoming networks, a device that learns of a MAC address as a DLR (local and remote) and also learns the MAC-IP address binding remotely will send a MAC-IP address notification message (as if it was generated by traffic). This entry is identified as a local proxy address. The remote multihomed peer identifies the corresponding MAC-IP entry as remote proxy address. • Rp—Remote proxy MAC IP address entries. See Lp flag for more information • K—MAC-IP address entries that are installed in the kernel.	Brief, detail

Table 485: show ethernet-switching mac-ip-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• RT—Destination Route that corresponds to an installed entry. • AD—Advertised to remote device. • NAD—MAC-IP address entries on disabled IRB interfaces. These entries are not advertised. • RE—Specifies temporary ARP/ND entries while the device is waiting for a response from an arp or network solicitation request. The entries are removed after three unsuccessful request attempts. • RO—Router. This applies only to IPv6 MAC-IP addresses. Refer to RFC 4861 for more information on the Router flag. • OV—Override.This applies only to IPv6 MAC-IP addresses. Refer to RFC 4861 for more information on the Override flag. • UR—Unresolved MAC-IP address entries. This flag is specific to the Junos EVO environment • RTS—Device skipped adding an EVPN Type 2 route in favor of an EVPN Type 5 route with a matching prefix (extensive output displays dest_route_skipped).	

Table 485: show ethernet-switching mac-ip-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
GBP Tag	Assigned Group Based Policy (GBP) from 1 through 65535.	All levels
Logical Interface	The logical interface associated with the routing instance.	Brief, detail
Active Source	The source of the learned MAC-IP address entry. Displays the ESI or router ID.	All levels
Routing Instance	Information on the Routing instance where the MAC-IP address entry was learned. • Bridging domain • MAC-IP local mask • MAC-IP remote mask • MAC-IP RPD flags • MAC-IP flags	Extensive
BD ID	Bridge Domain ID	Brief

Sample Output

show ethernet-switching mac-ip-table

```
user@host> show ethernet-switching mac-ip-table
MAC IP flags  (S - Static, D - Dynamic, L - Local , R - Remote, Lp - Local Proxy,
               Rp - Remote Proxy, K - Kernel, RT - Dest Route, (N)AD - (Not) Advt to remote,
               RE - Re-ARP/ND, RO - Router, OV - Override, Ur - Unresolved,
               RTS - Dest Route Skipped, RGw - Remote Gateway, GBP - Group Based Policy,
               RTF - Dest Route Forced)
```

```
 Routing instance : default-switch
 Bridging domain : vlan1000
   IP                          MAC                  Flags              GBP
Logical            Active
   address                     address                                 Tag
Interface          source
   10.1.1.2                   c8:e7:f0:4b:d1:00    S,K                       irb.1000
   10.1.1.1                   dc:38:e1:e0:30:c0    DL,K,RT,AD          300   ae0.0
   192.168.1.1                00:21:59:aa:77:f0    SR,K,RT,RGw               vtep.32769
18.18.18.18
   192.168.1.100              00:00:5e:00:01:01    S,K
```

show ethernet-switching mac-ip-table

```
user@host> show ethernet-switching mac-ip-table
vlan-name v100
MAC IP flags  (S - Static, D - Dynamic, L - Local , R - Remote, K - Kernel, RT - Dest Route,
               AD - Advt to remote, RE - Re-ARP/ND, RO - Router, OV - Override)
 Routing instance : default-
 Bridging domain : v100
   IP                          MAC                  Flags             Logical            Active
   address                     address                                Interface          source
   10.1.1.1                    00:11:00:00:00:01    DR,K,RT           vtep.32769
101.1.1.1
   2001::192:100:1:1           00:00:10:00:11:00    DR,K,RT           vtep.32769
101.1.1.1
   fe80::205:8600:6471:e600    00:00:10:00:11:00    DR,K,RT           vtep.32769
101.1.1.1
   10.0.1.2                    00:00:10:00:22:00    S,K
   2001::192:100:1:2           00:00:10:00:22:00    S,K
   fe80::205:8600:6471:d900    00:00:10:00:22:00    S,K
   10.0.1.3                    00:00:10:00:33:00    DR,K,RT           vtep.32771
103.1.1.1
   2001::192:100:1:3           00:00:10:00:33:00    DR,K,RT           vtep.32771
103.1.1.1
   fe80::205:8600:6471:7400    00:00:10:00:33:00    DR,K,RT           vtep.32771
103.1.1.1
   10.0.1.4                    00:00:10:00:44:00    DR,K,RT           vtep.32770
104.1.1.1
```

```
   2001::192:100:1:4            00:00:10:00:44:00    DR,K,RT           vtep.32770
104.1.1.1
   fe80::205:8600:6471:c000     00:00:10:00:44:00    DR,K,RT           vtep.32770
104.1.1.1
   10.0.1.254                   40:00:10:11:11:00    S,K               esi.1874
05:00:00:00:64:00:00:00:64:00
   2001::192:100:1:254          60:00:10:11:11:00    S,K               esi.1874
05:00:00:00:64:00:00:00:64:00
```

show ethernet-switching mac-ip-table extensive

```
user@host> show ethernet-switching mac-ip-table
extensive
IP address: 10.1.1.1
  MAC address: 00:11:00:00:00:01
  Routing instance: default
   Bridging domain: v100
   MAC-IP local mask: 0x0000000000000000
   MAC-IP remote mask: 0x8000000000000000
   MAC-IP RPD flags: 0x00000081
   MAC-IP flags: remote,kernel
```

show ethernet-switching mac-ip-table kernel

```
user@host> show ethernet-switching mac-ip-table
kernel
    BD       IP                          MAC                 Logical           Flags
    id       address                     address             Interface
    2        192.168.1.1                 00:21:59:aa:77:f0   irb.1             0x00000209
    2        192.168.1.100               00:00:5e:00:01:01   irb.1             0x00000609
    3        192.168.2.1                 00:21:59:aa:77:f0   irb.2             0x00000209
    3        192.168.2.100               00:00:5e:00:01:01   irb.2             0x00000609
    4        192.168.3.1                 00:21:59:aa:77:f0   irb.3             0x00000209
    4        192.168.3.100               00:00:5e:00:01:01   irb.3             0x00000609
    5        192.168.4.1                 00:21:59:aa:77:f0   irb.4             0x00000209
    5        192.168.4.100               00:00:5e:00:01:01   irb.4             0x00000609
```

show ethernet-switching mac-ip-table vlan-name v1 (with EVPN Type 2 and EVPN Type 5 route coexistence)

```
user@device> show ethernet-switching mac-ip-table vlan-name v1

MAC IP flags  (S - Static, D - Dynamic, L - Local , R - Remote, Lp - Local Proxy,
               Rp - Remote Proxy, K - Kernel, RT - Dest Route, (N)AD - (Not) Advt to remote,
               RE - Re-ARP/ND, RO - Router, OV - Override, Ur - Unresolved,
               RTS - Dest Route Skipped, RGw - Remote Gateway)
 Routing instance : default-switch
 Bridging domain : v1
   IP                          MAC                 Flags            Logical            Active
   address                     address                              Interface          source
   10.1.1.254                  10:10:10:00:00:01   S,K              irb.1
   2001:db8::a:1:1:fffe        20:20:20:00:00:01   S,K,RTS          irb.1
   10.1.1.2                    c8:e7:f0:4b:d1:00   DR,K,RTS         vtep.32769
192.168.22.22
   2001:db8::a:1:1:2           c8:e7:f0:4b:d1:00   DR,K,RTS         vtep.32769
192.168.22.22
   fe80::cae7:f000:14b:d100    c8:e7:f0:4b:d1:00   DR,K,RT          vtep.32769
192.168.22.22
   10.1.1.1                    dc:38:e1:e0:30:c0   S,K              irb.1
   2001:db8::a:1:1:1           dc:38:e1:e0:30:c0   S,K              irb.1
   fe80::de38:e100:1e0:30c0    dc:38:e1:e0:30:c0   S,K              irb.1
```

show ethernet-switching mac-ip-table extensive (with EVPN Type 2 and EVPN Type 5 route coexistence)

```
user@device> show ethernet-switching mac-ip-table extensive c8:e7:f0:4b:d1:00

IP address: 10.1.1.2
  MAC address: c8:e7:f0:4b:d1:00
  Routing instance: default-switch
   Bridging domain: v1
   Logical interface: vtep.32769
   IP                          MAC                 Flags            Logical            Active
   address                     address                              Interface          source

192.168.22.22
   MAC-IP local mask: 0x0000000000000000
```

```
   MAC-IP remote mask: 0x0000000000000000
   MAC-IP remote-proxy mask: 0x0000000000000000
   MAC-IP RPD flags: 0x08000081
   MAC-IP flags: remote,kernel,dest_route_skipped

IP address: 2001:db8::a:1:1:2
  MAC address: c8:e7:f0:4b:d1:00
  Routing instance: default-switch
   Bridging domain: v1
   Logical interface: vtep.32769
   IP                          MAC                 Flags             Logical           Active
   address                     address                               Interface         source

192.168.22.22
   MAC-IP local mask: 0x0000000000000000
   MAC-IP remote mask: 0x0000000000000000
   MAC-IP remote-proxy mask: 0x0000000000000000
   MAC-IP RPD flags: 0x08000101
   MAC-IP flags: remote,kernel,dest_route_skipped

IP address: fe80::cae7:f000:14b:d100
  MAC address: c8:e7:f0:4b:d1:00
  Routing instance: default-switch
   Bridging domain: v1
   Logical interface: vtep.32769
   IP                          MAC                 Flags             Logical           Active
   address                     address                               Interface         source

192.168.22.22
   MAC-IP local mask: 0x0000000000000000
   MAC-IP remote mask: 0x0000000000000000
   MAC-IP remote-proxy mask: 0x0000000000000000
   MAC-IP RPD flags: 0x08000101
   MAC-IP flags: remote,kernel,dest_route
. . .
```

Release Information

Command introduced in Junos OS Release 17.4R2.

`GBP Tag` option introduced in Junos OS Release 22.4R1.

show ethernet-switching mac-learning-log

IN THIS SECTION

- Syntax | 19670
- Description | 19670
- Required Privilege Level | 19670
- Output Fields | 19670
- Sample Output | 19673
- Release Information | 19676

Syntax

```
show ethernet-switching mac-learning-log
```

Description

Displays the event log of learned MAC addresses.

Required Privilege Level

view

Output Fields

Output fields for EX Series switches:

The following table lists the output fields for the `show ethernet-switching mac-learning-log` command. Output fields are listed in the approximate order in which they appear.

Table 486: show ethernet-switching mac-learning-log Output Fields

Field Name	Field Description
Date and Time	Timestamp when the MAC address was added or deleted from the log.
vlan_name	VLAN name. A value defined by the user for all user-configured VLANs.
MAC	Learned MAC address.
Deleted | Added	MAC address deleted or added to the MAC learning log.
Blocking	The forwarding state of the interface: • **blocked**—Traffic is not being forwarded on the interface. • **unblocked**—Traffic is forwarded on the interface.
`Flags`	Displays the MAC address flags in which the MAC event occurred. This option is for debugging purposes.

Output fields for QFX Series switches, QFabric, NFX Series devices and EX4600:

Table 487 on page 19671 lists the output fields for the `show ethernet-switching mac-learning-log` command. Output fields are listed in the approximate order in which they appear.

Table 487: show ethernet-switching mac-learning-log Output Fields

Field Name	Field Description
Date and Time	Timestamp in UTC when the MAC operation occurred.
vlan_name	VLAN name. A value defined by the user for all user-configured VLANs. The name of the VLAN on which the MAC is learned.
MAC	Learned MAC address.

Table 487: show ethernet-switching mac-learning-log Output Fields *(Continued)*

Field Name	Field Description
`Event op`	MAC address that are added, learned, deleted, changed or moved from one interface to another interface.
`Interface Name`	The name of the interface on which the MAC address is learned. When a MAC address is moved, there is another field with the name of the interface. The log displays the name of the interface from where the MAC address moved, and the name of the interface to where the MAC address moved.
`Flags`	Displays the MAC address flags in which the MAC event occurred. This option is for debugging purposes.

Output fields for SRX Series Firewalls:

Table 488 on page 19672 lists the output fields for the show ethernet-switching mac-learning-log command on SRX Series Firewalls. Output fields are listed in the approximate order in which they appear.

Table 488: show ethernet-switching-mac-learning-log Output Fields

Field Name	Field Description
`Date and Time`	Timestamp when the MAC address was added or deleted from the log.
`VLAN-IDX`	VLAN index. An internal value assigned by Junos OS for each VLAN.
`MAC`	Learned MAC address.
`Deleted | Added`	MAC address deleted or added to the MAC learning log.
`Blocking`	The forwarding state of the interface: • blocked—Traffic is not being forwarded on the interface. • unblocked—Traffic is forwarded on the interface.

Sample Output

show ethernet-switching mac-learning-log (EX Series switch)

```
user@switch> show ethernet-switching mac-learning-log
Mon Feb 25 08:07:05 2008
 vlan_name v1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v9 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v3 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v12 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v13 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name sales_vlan mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v3 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:05:00:00:05 was learned
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:30:48:90:54:89 was learned
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:5e:00:01:00 was learned
Mon Feb 25 08:07:05 2008
 vlan_name sales_vlan mac 00:00:5e:00:01:08 was learned
[output truncated]
```

show ethernet-switching mac-learning-log (QFX Series Switches, QFabric, NFX Series Devices and EX4600)

```
user@switch> show ethernet-switching mac-learning-log
Mon Jun 30 13:49:49 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was learned on ge-1/0/22.0 with
flags: 0x2001f     << MAC address that as dynamically learned
Mon Jun 30 13:50:29 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was deleted from ge-1/0/22.0
with flags: 0x1080  << MAC address that was deleted
Mon Jun 30 13:51:28 2014 vlan_name v11+11 mac 00:00:00:01:01:01 was added to ge-1/0/22.0 with
flags: 0x2013f        << Static MAC address that was added
Mon Jun 30 13:51:46 2014 vlan_name v11+11 mac 00:00:00:01:01:01 was deleted from ge-1/0/22.0
with flags: 0x1120   << delete of Static MAC address that was deleted
Mon Jun 30 13:52:03 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was learned on ge-1/0/22.0 with
flags: 0x2001f    << MAC address that was dynamically learned
Mon Jun 30 13:52:11 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was moved from ge-1/0/22.0 to
ge-1/0/21.0 with flags: 0x2101f    << MAC address that was moved
Mon Jun 30 13:54:24 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was changed on ge-1/0/21.0 with
flags: 0x2113f   << MAC address that changed from a dynamic address to a static address
```

show ethernet-switching mac-learning-log (SRX Series Firewalls)

```
user@host> show ethernet-switching mac-learning-log
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:00 was deleted
```

```
Wed Mar 18 08:07:05 2009
vlan_idx 4 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 18 mac 00:00:5E:00:53:AA was learned
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:00:5E:00:53:AB was learned
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:5E:00:53:AC was learned
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:AD was learned
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:AE was learned
Wed Mar 18 08:07:05 2009
vlan_idx 8 mac 00:00:5E:00:53:AF was learned
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:AG was learned
[output truncated]
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- show ethernet-switching table | 19700
- show ethernet-switching interfaces | 19641
- show ethernet-switching table | 19700
- show ethernet-switching interfaces | 19641
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- Example: Connecting an EX Series Access Switch to a Distribution Switch

show ethernet-switching mgrp-policy

IN THIS SECTION

- Syntax | 19676
- Description | 19677
- Options | 19677
- Required Privilege Level | 19678
- Output Fields | 19678
- Sample Output | 19679
- Release Information | 19682

Syntax

```
show ethernet-switching mgrp-policy
<instance instance-name>
    <mgrp-id mgrp-id>
```

Description

Display information about logical mesh groups, which are part of a mechanism that Juniper Networks uses to control the flooding of broadcast, unknown unicast, and multicast (BUM) traffic.

Checking mesh groups helps you to monitor the state of the seamless EVPN-VXLAN stitching feature that interconnects EVPN-VXLAN points of delivery (PODs) in a data center, or EVPN-VXLAN data centers. In each of these use cases, Juniper Networks devices that support seamless EVPN-VXLAN stitching act as a spine or super spine device that interconnects the PODs or data centers.

Juniper Networks supports the following default mesh groups for a VLAN. All interfaces that you configure on a spine or super spine device can be assigned to a particular mesh group.

- CE mesh group, to which all directly connected interfaces are assigned.
- VE mesh group, to which all virtual tunnel endpoint (VTEP) interfaces for the local POD or data center are assigned.
- WAN mesh group, to which all remote VTEP interfaces within a POD or data center, or interconnected PODs or data centers are assigned.
- RE mesh group, to which interfaces are typically not assigned.

When a BUM packet enters an interface on the spine or super spine device and that packet needs to be flooded, the mesh group to which the interface belongs is identified. The mesh group associated with a particular VLAN determines the outgoing groups to which the packet should be flooded. For example, if a BUM packet enters an interface assigned to the CE mesh group, the outgoing flood groups might the following:

- All interfaces in the CE mesh group except the interface through which the packet entered.
- The VE and WAN mesh groups.

Additional factors that determine the outgoing flood groups include the following:

- Whether the spine or super spine device that receives a packet acts as a designated forwarder (DF) or non-DF.
- The local bias feature, also known as split horizon, which eliminates duplication of traffic to its destination in another POD or data center.

Options

none	Display the flood groups associated with all mesh groups.

instance ***instance-name***	(Optional) Display the flood groups associated with the mesh groups for a specified routing or switching instance.
mgrp-id ***mgrp-id***	(Optional) Display interfaces associated with a specified mesh group for a particular routing or switching instance.

Required Privilege Level

admin

Output Fields

Table 489 on page 19678 lists the output fields for the `show ethernet-switching mgrp-policy` command. Table 490 on page 19678 lists the output fields for the `show ethernet-switching mgrp-policy instance` command. Table 491 on page 19679 lists the output fields for the `show ethernet-switching mgrp-policy instance mgrp-id` command. Output fields are listed in the approximate order in which they appear.

Table 489: show ethernet-switching mgrp-policy Output Fields

Field Name	Field Description
Role	Possible roles include either DF (designated forwarder) or NDF (non-designated forwarder). The output is organized by role. Mesh groups listed under each role provide a summary of the groups to which a BUM packet received by a spine or super spine device that acts as a DF or non-DF will be flooded.
Mesh Group	Name of a mesh group. For each mesh group, the output displays the outgoing flood groups to which a BUM packet will be flooded.

Table 490: show ethernet-switching mgrp-policy instance Output Fields

Field Name	Field Description
Role	Possible roles include either DF (designated forwarder) or NDF (non-designated forwarder). The output displays the role for each mesh group.

Table 490: show ethernet-switching mgrp-policy instance Output Fields *(Continued)*

Field Name	Field Description
Mesh Group	Name of a mesh group. For each mesh group, the output displays: • Flags associated with the mesh group. • The outgoing flood groups to which a BUM packet will be flooded.

Table 491: show ethernet-switching mgrp-policy instance mgrp-id Output Fields

Field Name	Field Description
Mesh Group	Name of a mesh group. For the mesh group, the output displays associated flags.
Interfaces	Interfaces associated with the mesh group. When a spine or super spine device floods a BUM packet to the next hop, they will do so through these interfaces.

Sample Output

show ethernet-switching mgrp-policy

```
user@switch> show ethernet-switching mgrp-policy
 Role: DF

Mesh Group:    __all_ces__
    __all_ces__
    __ves__
    __wan_flood__
    __icl_flood__
    __etree_leaf__

Mesh Group:    __esi_flood__

Mesh Group:    __etree_leaf__
    __all_ces__
    __ves__
```

```
Mesh Group:    __icl_flood__
    __all_ces__

Mesh Group:    __mlp_flood__

Mesh Group:    __re_flood__
    __all_ces__
    __ves__
    __wan_flood__
    __icl_flood__
    __etree_leaf__

Mesh Group:    __ves__
    __all_ces__
    __wan_flood__
    __etree_leaf__

Mesh Group:    __wan_flood__
    __all_ces__
    __ves__

 Role: NDF

Mesh Group:    __all_ces__
    __all_ces__
    __ves__
    __wan_flood__
    __icl_flood__

Mesh Group:    __ar_flood__
    __all_ces__
    __ar_flood__
    __wan_flood__

Mesh Group:    __esi_flood__

Mesh Group:    __etree_leaf__

Mesh Group:    __icl_flood__

Mesh Group:    __mlp_flood__

Mesh Group:    __re_flood__
```

```
    __all_ces__
    __ves__
    __wan_flood__

Mesh Group:    __ves__

Mesh Group:    __wan_flood__
    __all_ces__
```

show ethernet-switching mgrp-policy instance

```
user@switch> show ethernet-switching mgrp-policy
instance evpn-vxlan
Role: DF
Mesh Group:    __all_ces__Flags: 0xa ( split-horizon, shared-RT)
    __all_ces__
    __ves__
    __wan_flood__

Role: DF
Mesh Group:    __mlp_flood__Flags: 0xc ( shared-RT, no-local-flooding)

Role: DF
Mesh Group:    __re_flood__Flags: 0x4c ( shared-RT, no-local-flooding, permanent-comp-nh)
    __all_ces__
    __ves__
    __wan_flood__

Role: DF
Mesh Group:    __ves__Flags: 0xc ( shared-RT, no-local-flooding)
    __all_ces__
    __wan_flood__

Role: DF
Mesh Group:    __wan_flood__Flags: 0xc ( shared-RT, no-local-flooding)
    __all_ces__
    __ves__

Role: Non-DF
Mesh Group:    __ves__Flags: 0xc ( shared-RT, no-local-flooding)
```

```
Role: Non-DF
Mesh Group:    __wan_flood__Flags: 0xc ( shared-RT, no-local-flooding)
    __all_ces__
```

show ethernet-switching mgrp-policy instance mgrp-id

```
user@switch> show ethernet-switching mgrp-policy
instance evpn-vxlan mgrp-id 0
Mesh Group:    __ves__Flags: 0xc ( shared-RT, no-local-flooding)
Interfaces
     vtep.32769
     vtep.32770
     vtep.32771
     vtep.32772
```

Release Information

Command introduced in Junos OS Release 20.3R1.

show ethernet-switching redundancy-groups

IN THIS SECTION

- Syntax | 19683
- Description | 19683
- Options | 19683
- Required Privilege Level | 19683
- Output Fields | 19683
- Sample Output | 19686
- Release Information | 19687

Syntax

```
show ethernet-switching redundancy-groups
<redundancy-group-id [0 to 4294967294]>
arp-statistics
nd-statistics
remote-macs
```

Description

Display ARP statistics, Neighbor Discovery statistics, or remote MAC addresses for the Multi-Chassis Aggregated Ethernet (MC-AE) nodes for all or specified redundancy groups on a router or switch. Note that the Redundancy Group ID is inherited by the bridging domain or VLAN from member AE interfaces.

Options

redundancy-group-id	(Optional) The redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate the routing or switching devices contained in a redundancy group.
arp-statistics	(Optional) Count of ARP packets sent and received by the two MC-AE nodes.
nd-statistics	(Optional) Count of Neighbor Discovery packets sent and received by the two MC-AE nodes.
remote-macs	(Optional) List of remote MAC addresses in the “Installed” state, as learned from the remote MC-AE node.

Required Privilege Level

view

Output Fields

Output fields are listed in the approximate order in which they appear.

Table 492: show ethernet-switching redundancy-groups arp-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
MCLAG ARP Statistics Group ID	ARP statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.
ARP Rx Count From Line	Total number of ARPs received from the Line.
ARP Tx Count To Peer	Total number of ARPs sent to the peer.
ARP Rx Count From Peer	Total number of ARPs received from the peer.
ARP Drop Count received from line	Total number of ARPs sent by the peer that were received.
ARP Drop Count received from peer	Total number of ARPs sent by the peer that were dropped
ARP Install Count	ARP Install Count

Table 493: show ethernet-switching redundancy-groups nd-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
MCLAG ND Statistics Group ID	Neighbor Discovery statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.

Table 493: show ethernet-switching redundancy-groups nd-statistics Output Fields *(Continued)*

Field Name	Field Description
ND Rx Count From Line	Total number of Neighbor Discovery packets received from the Line.
ND Tx Count To Peer	Total number of Neighbor Discovery packets sent to the peer.
ND Rx Count From Peer	Total number of Neighbor Discovery packets received from the peer.
ND Drop Count received from line	Total number of Neighbor Discovery packets sent by the peer that were received.
ND Drop Count received from peer	Total number of Neighbor Discovery packets sent by the peer that were dropped
ND Install Count	ND Install Count

Table 494: show ethernet-switching redundancy-groups remote-macs Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
Service ID	Service ID (configured at the routing instance level).
Peer-Addr	IP address of the remote peer.
VLAN	Virtual LAN identifier associated with the redundancy group.
MAC	Hardware media access control address associated with the redundancy group.
MCAE-ID	ID number of the MC-AE used by the redundancy group.

Table 494: show ethernet-switching redundancy-groups remote-macs Output Fields *(Continued)*

Field Name	Field Description
Flags	Connection state: local connect or Remote connect. If no flag is shown, the redundancy group may not be connected.
Status	Installation state: Installed or Not Installed.

Sample Output

show ethernet-switching redundancy-groups arp-statistics

```
user@host> show ethernet-switching redundancy-groups arp-statistics

Redundancy Group ID : 1       Flags : Local Connect,Remote Connect


    MCLAG ARP Statistics
    Group ID                                  : 1
    ARP Rx Count From Line                    : 3493
    ARP Tx Count To Peer                      : 647
    ARP Rx Count From Peer                    : 0
    ARP Install Count                         : 0
    ARP Drop Count received from line         : 2846
    ARP Drop Count received from peer         : 0
```

show ethernet-switching redundancy-groups nd-statistics

```
user@host> show ethernet-switching redundancy-groups nd-statistics

 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID                                   : 1
ND Rx Count From Line                     : 52
ND Tx Count To Peer                       : 15
```

```
ND Rx Count From Peer                    : 39
ND Install Count                         : 34
ND Drop Count received from line         : 37
ND Drop Count received from peer         : 5
```

show ethernet-switching redundancy-groups remote-macs

```
user@host> show ethernet-switching redundancy-groups <redundancy-group-id> remote-macs

Redundancy Group ID : 1       Flags : Local Connect,Remote Connect


    Service-id  Peer-Addr    VLAN        MAC            MCAE-ID  Subunit  Opcode   Flags
Status
    10          10.3.3.2     100    80:ac:ac:1f:10:a1   1        0        1        0
Installed
```

show ethernet-switching redundancy-groups group-id

```
user@host> show ethernet-switching redundancy-groups 1

Redundancy Group ID : 1       Flags : Local Connect,Remote Connect
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

- Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up

show ethernet-switching statistics

IN THIS SECTION

- Syntax | 19688
- Description | 19688
- Options | 19688
- Required Privilege Level | 19689
- Sample Output | 19689
- Release Information | 19691

Syntax

```
show ethernet-switching statistics
<instance instance-name>
<logical-system logical-system-name>
<vlan-name vlan-name>
```

Description

Display Ethernet-switching statistics.

> **NOTE**: On some devices running Junos OS Evolved, results from the show ethernet-switching statistics command are limited. Only the Current MAC count statistic is displayed.

Options

none	Display Ethernet-switching statistics for all VLANs in all routing instances.
instance *instance-name*	(Optional) Display statistics for the specified routing instance.
logical-system *logical-system-name*	(Optional) Display Ethernet-switching statistics information for the specified logical system.

vlan-name *vlan-name* (Optional) Display statistics for the specified VLAN.

Required Privilege Level

view

Sample Output

show ethernet-switching statistics

```
user@host> show ethernet-switching statistics
 Local interface: ae1.0, Index: 1035
   Broadcast packets:                  220
   Broadcast bytes  :                13720
   Multicast packets:                  130
   Multicast bytes  :                11700
   Flooded packets  :                    0
   Flooded bytes    :                    0
   Unicast packets  :                    0
   Unicast bytes    :                    0
   Current MAC count:                    0 (Limit 1024)
 Local interface: vt-3/3/10.1048576, Index: 1280
   Broadcast packets:                    0
   Broadcast bytes  :                    0
   Multicast packets:                    0
   Multicast bytes  :                    0
   Flooded packets  :                    2
   Flooded bytes    :                  128
   Unicast packets  :                  632
   Unicast bytes    :                39184
   Current MAC count:                    2
 Local interface: ge-3/1/2.0, Index: 1258
   Broadcast packets:                  100
   Broadcast bytes  :                 6800
   Multicast packets:                  200
   Multicast bytes  :                18000
   Flooded packets  :                    0
   Flooded bytes    :                    0
   Unicast packets  :                  632
   Unicast bytes    :                39184
```

```
  Current MAC count:                     2 (Limit 1024)
Local interface: ae3.0, Index: 1043
  Broadcast packets:                     0
  Broadcast bytes  :                     0
  Multicast packets:                     0
  Multicast bytes  :                     0
  Flooded packets  :                     0
  Flooded bytes    :                     0
  Unicast packets  :                     0
  Unicast bytes    :                     0
  Current MAC count:                     0 (Limit 1024)
Local interface: ge-3/3/8.0, Index: 1276
  Broadcast packets:                     0
  Broadcast bytes  :                     0
  Multicast packets:                     0
  Multicast bytes  :                     0
  Flooded packets  :                     0
  Flooded bytes    :                     0
  Unicast packets  :                     0
  Unicast bytes    :                     0
  Current MAC count:                     0 (Limit 8192)
Local interface: ae5.0, Index: 1045
  Broadcast packets:                     0
  Broadcast bytes  :                     0
  Multicast packets:                     0
  Multicast bytes  :                     0
  Flooded packets  :                     0
  Flooded bytes    :                     0
  Unicast packets  :                     0
  Unicast bytes    :                     0
  Current MAC count:                     0 (Limit 8192)
Local interface: ae4.0, Index: 1044
  Broadcast packets:                   200
  Broadcast bytes  :                 13600
  Multicast packets:                     0
  Multicast bytes  :                     0
  Flooded packets  :                     0
  Flooded bytes    :                     0
  Unicast packets  :                     0
  Unicast bytes    :                     0
  Current MAC count:                     0 (Limit 8192)
Local interface: ae26.0, Index: 1042
  Broadcast packets:                     0
```

```
    Broadcast bytes  :                    0
    Multicast packets:                    0
    Multicast bytes  :                    0
    Flooded packets  :                    0
    Flooded bytes    :                    0
    Unicast packets  :                    0
    Unicast bytes    :                    0
    Current MAC count:                    0 (Limit 8192)
  Local interface: ae25.0, Index: 1041
    Broadcast packets:                  133
    Broadcast bytes  :                 7980
    Multicast packets:               369934
    Multicast bytes  :             59207572
    Flooded packets  :                    0
    Flooded bytes    :                    0
    Unicast packets  :                 1433
    Unicast bytes    :               119930
    Current MAC count:                    3 (Limit 8192)
  Local interface: ae23.0, Index: 1040
    Broadcast packets:                  226
    Broadcast bytes  :                14464
    Multicast packets:               585668
    Multicast bytes  :            153464476
    Flooded packets  :                    0
    Flooded bytes    :                    0
    Unicast packets  :                26552
    Unicast bytes    :              1947627
    Current MAC count:                    7 (Limit 8192)
  Local interface: ae20.0, Index: 1039
    Broadcast packets:                  115
    Broadcast bytes  :                 6900
    Multicast packets:               395113
    Multicast bytes  :             61622869
    Flooded packets  :                    0
    Flooded bytes    :                    0
    Unicast packets  :                 1419
    Unicast bytes    :               117924
    Current MAC count:                    4 (Limit 8192)
```

Release Information

Command introduced in Junos OS Release 12.3R2.

show ethernet-switching statistics aging

IN THIS SECTION

- Syntax | 19692
- Description | 19692
- Options | 19692
- Required Privilege Level | 19692
- Output Fields | 19693
- Sample Output | 19693
- Release Information | 19694

Syntax

```
show ethernet-switching statistics aging
```

Description

Display media access control (MAC) aging statistics.

The command is not supported on MX Series routers and on products that support the Enhanced Layer 2 Software (ELS).

Options

none	(Optional) Display MAC aging statistics.
brief | detail	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 495 on page 19693 lists the output fields for the `show ethernet-switching statistics aging` command. Output fields are listed in the approximate order in which they appear.

Table 495: show ethernet-switching statistics aging Output Fields

Field Name	Field Description	Level of Output
Total age messages received	Total number of aging messages received from the hardware.	All levels
Immediate aging	Aging message indicating that the entry should be removed immediately.	All levels
MAC address seen	Aging message indicating that the MAC address has been detected by hardware and that the aging timer should be stopped.	All levels
MAC address not seen	Aging message indicating that the MAC address has not been detected by the hardware and that the aging timer should be started.	All levels
Error age messages	The received aging message contains the following errors: • **Invalid VLAN**—The VLAN of the packet does not exist. • **No such entry**—The MAC address and VLAN pair provided by the aging message does not exist. • **Static entry**—An unsuccessful attempt was made to age out a static MAC entry.	All levels

Sample Output

show ethernet-switching statistics aging

```
user@switch> show ethernet-switching statistics aging

Total age messages received: 0
  Immediate aging: 0, MAC address seen: 0, MAC address not seen: 0
```

```
Error age messages: 0
  Invalid VLAN: 0, No such entry: 0, Static entry: 0
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

show ethernet-switching statistics mac-learning | 19694

Configuring MAC Table Aging on Switches

show ethernet-switching statistics mac-learning

IN THIS SECTION

- Syntax | 19694
- Description | 19695
- Options | 19695
- Required Privilege Level | 19695
- Output Fields | 19695
- Sample Output | 19697
- Release Information | 19699

Syntax

```
show ethernet-switching statistics mac-learning
<brief | detail>
<interface interface-name>
```

Description

Display media access control (MAC) learning statistics.

> **NOTE**: For the QFX Series, this command is not supported in Enhanced Layer 2 Software (ELS).

Options

none	(Optional) Display MAC learning statistics for all interfaces.
brief | detail	(Optional) Display the specified level of output. The default is `brief`.
interface *interface-name*	(Optional) Display MAC learning statistics for the specified interface.

Required Privilege Level

view

Output Fields

Table 496 on page 19695 lists the output fields for the `show ethernet-switching statistics mac-learning` command. Output fields are listed in the approximate order in which they appear.

Table 496: show ethernet-switching statistics mac-learning Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface for which statistics are being reported. (Displayed in the output under the heading **Interface**.)	All levels
Learning message from local packets	MAC learning message generated due to packets coming in on the management interface. (Displayed in the output under the heading **Local pkts**.)	All levels
Learning message from transit packets	MAC learning message generated due to packets coming in on network interfaces. (Displayed in the output under the heading **Transit pkts**.)	All levels

Table 496: show ethernet-switching statistics mac-learning Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Learning message with error	MAC learning messages received with errors (Displayed under the heading **Error**): • **Invalid VLAN**—The VLAN of the packet does not exist. • **Invalid MAC**—The MAC address is either NULL or a multicast MAC address. • **Security violation**—The MAC address is not an allowed MAC address. • **Interface down**—The MAC address is learned on an interface that is down. • **Incorrect membership**—The MAC address is learned on an interface that is not a member of the VLAN. • **Interface limit**—The number of MAC addresses learned on the interface has exceeded the limit. • **MAC move limit**—This MAC address has moved among multiple interfaces too many times in a given interval. • **VLAN limit**—The number of MAC addresses learned on the VLAN has exceeded the limit. • **VLAN membership limit**—The number of MAC addresses learned on the interface as a member of the specified VLAN (VLAN membership MAC limit) has exceeded the limit. • **Invalid VLAN index**—The VLAN of the packet, although configured, does not yet exist in the kernel. • **Interface not learning**—The MAC address is learned on an interface that does not yet allow learning—for example, the interface is blocked. • **No nexthop**—The MAC address is learned on an interface that does not have a unicast next hop. • **MAC learning disabled**—The MAC address is learned on an interface on which MAC learning has been disabled.	All levels

Table 496: show ethernet-switching statistics mac-learning Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **Others**—The message contains some other error.	

Sample Output

show ethernet-switching statistics mac-learning

```
user@switch> show ethernet-switching statistics mac-learning

Learning stats: 0 learn msg rcvd, 0 error
  Interface         Local pkts        Transit pkts      Error
  ge-0/0/0.0        0                 0                 0
  ge-0/0/1.0        0                 0                 0
  ge-0/0/2.0        0                 0                 0
  ge-0/0/3.0        0                 0                 0
```

show ethernet-switching statistics mac-learning detail

```
user@switch> show ethernet-switching statistics mac-learning detail
Learning stats: 0 learn msg rcvd, 0 error

    Interface: ge-0/0/0.0
    Learning message from local packets:   0
    Learning message from transit packets: 1
    Learning message with error:           0
      Invalid VLAN:               0       Invalid MAC:                  0
      Security violation:         0       Interface down:               0
      Incorrect membership:       0       Interface limit:              0
      MAC move limit:             0       VLAN limit:                   0
      Invalid VLAN index:         0       Interface not learning:       0
      No nexthop:                 0       MAC learning disabled:        0
      Others:                     0

    Interface: ge-0/0/1.0
    Learning message from local packets:   0
```

```
    Learning message from transit packets: 2
    Learning message with error:           0
      Invalid VLAN:               0       Invalid MAC:                  0
      Security violation:         0       Interface down:               0
      Incorrect membership:       0       Interface limit:              0
      MAC move limit:             0       VLAN limit:                   0
      Invalid VLAN index:         0       Interface not learning:       0
      No nexthop:                 0       MAC learning disabled:        0
      Others:                     0
```

show ethernet-switching statistics mac-learning interface ge-0/0/28 detail

```
user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/28
detail

    Interface: ge-0/0/28.0
    Learning message from local packets:   0
    Learning message from transit packets: 5
    Learning message with error:           0
      Invalid VLAN:               0       Invalid MAC:                  0
      Security violation:         0       Interface down:               0
      Incorrect membership:       0       Interface limit:              0
      MAC move limit:             0       VLAN limit:                   0
                                          VLAN membership limit:        20
      Invalid VLAN index:         0       Interface not learning:       0
      No nexthop:                 0       MAC learning disabled:        0
      Others:                     0
```

show ethernet-switching statistics mac-learning interface

```
user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/1

Interface         Local pkts        Transit pkts      Error
ge-0/0/1.0        0                 1                 1
```

show ethernet-switching statistics mac-learning detail (QFX Series)

```
user@switch> show ethernet-switching statistics mac-learning detail
Learning stats: 0 learn msg rcvd, 0 error

    Interface: xe-0/0/0.0
    Learning message from local packets:   0
    Learning message from transit packets: 1
    Learning message with error:           0
      Invalid VLAN:               0       Invalid MAC:                  0
      Security violation:         0       Interface down:               0
      Incorrect membership:       0       Interface limit:              0
      MAC move limit:             0       VLAN limit:                   0
      Invalid VLAN index:         0       Interface not learning:       0
      No nexthop:                 0       MAC learning disabled:        0
      Others:                     0

    Interface: xe-0/0/1.0
    Learning message from local packets:   0
    Learning message from transit packets: 2
    Learning message with error:           0
      Invalid VLAN:               0       Invalid MAC:                  0
      Security violation:         0       Interface down:               0
      Incorrect membership:       0       Interface limit:              0
      MAC move limit:             0       VLAN limit:                   0
      Invalid VLAN index:         0       Interface not learning:       0
      No nexthop:                 0       MAC learning disabled:        0
      Others:                     0
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

show ethernet-switching table | 19700

show ethernet-switching interfaces | 19641

show ethernet-switching mac-learning-log | 19670

show ethernet-switching table | 19700

show ethernet-switching interfaces | 19641

Example: Setting Up Basic Bridging and a VLAN on Switches

show ethernet-switching table

IN THIS SECTION

- Syntax (Products that support the Enhanced Layer 2 Software (ELS)) | 19700
- Syntax (Products that do not support ELS, Except MX Series Routers) | 19701
- Description | 19701
- Options | 19701
- Additional Information | 19702
- Required Privilege Level | 19702
- Output Fields | 19703
- Sample Output | 19706
- Release Information | 19723

Syntax (Products that support the Enhanced Layer 2 Software (ELS))

```
show ethernet-switching table
<brief | count | detail | extensive | summary>
<address>
<instance instance-name>
<interface interface-name>
isid isid
<logical-system logical-system-name>
<persistent-learning (interface interface-name | mac mac-address)>
<vlan-id (all-vlan | vlan-id)>
<vlan-name (all | vlan-name)>
```

Syntax (Products that do not support ELS, Except MX Series Routers)

```
show ethernet-switching table
<brief | detail | extensive | summary>
<interface interface-name>
<management-vlan>
<persistent-mac <interface interface-name>>
<sort-by (name | tag)>
<vlan vlan-name>
```

> **NOTE**: MX Series routers support the `show bridge mac-table` command in place of this command. For the syntax on MX Series routers, see "show bridge mac-table" on page 17264.

Description

Displays the Ethernet switching table.
Displays Layer 2 MAC address information.

Options

For products that support ELS:

none	Display all learned Layer 2 MAC address information.
brief | count | detail | extensive | summary	(Optional) Display the specified level of output.
address	(Optional) Display the specified learned Layer 2 MAC address information.
instance *instance-name*	(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.
interface *interface-name*	(Optional) Display learned Layer 2 MAC addresses for the specified interface.
isid *isid*	(Optional) Display learned Layer 2 MAC addresses for the specified ISID.
logical-system *logical-system-name*	(Optional) Display Ethernet-switching statistics information for the specified logical system.

persistent-learning (interface ***interface-name*** **| mac** ***mac-address*****)**	(Optional) Display dynamically learned MAC addresses that are retained despite device restarts and interface failures for a specified interface, or information about a specified MAC address.
vlan-id (all-vlan | ***vlan-id*****)**	(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.
vlan-name (all | ***vlan-name*****)**	(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

For products that do not support ELS:

none	(Optional) Display brief information about the Ethernet switching table.
brief | detail | extensive | summary	(Optional) Display the specified level of output.
interface ***interface-name***	(Optional) Display the Ethernet switching table for a specific interface.
management-vlan	(Optional) Display the Ethernet switching table for a management VLAN.
persistent-mac <interface ***interface-name*****>**	(Optional) Display the persistent MAC addresses learned for all interfaces or a specified interface. You can use this command to view entries that you want to clear for an interface that you intentionally disabled.
sort-by (*name* **|** ***tag*****)**	(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.
vlan ***vlan-name***	(Optional) Display the Ethernet switching table for a specific VLAN.

Additional Information

When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.

Required Privilege Level

view

Output Fields

For products that support ELS: The table describes the output fields for the `show ethernet-switching table` command on products that support ELS. Output fields are listed in the approximate order in which they appear.

Table 497: show ethernet-switching table Output fields on Products That Support ELS

Field Name	Field Description
Routing instance	Name of the routing instance.
VLAN name	Name of the VLAN.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: • **S**—Static MAC address is configured. • **D**—Dynamic MAC address is configured. • **L**—Locally learned MAC address is configured. • **SE**—MAC accounting is enabled. • **NM**—Non-configured MAC. • **R**—Remote PE MAC address is configured. • **O**—OVSDB learned MAC address is configured. • **B**—Blocked duplicate MAC address.
Age	This field is not supported.
Logical interface	Name of the logical interface. Name of the VTEP logical interface learned over remote VTEP.
GBP Tag	Assigned Group Based Policy (GBP) from 1 through 65535.

Table 497: show ethernet-switching table Output fields on Products That Support ELS *(Continued)*

Field Name	Field Description
SVLBNH/VENH Index	**NOTE**: This field appears on QFX5*XXX* switches that support dynamic load balancing in an EVPN-VXLAN network. Next-hop index number associated with the MAC address of a multihomed remote device in an EVPN-VXLAN network. This index number appears when the Logical Interface column displays esi.*nnnn*. The index number can be an SVLBNH, a VENH, or a remote virtual tunnel endpoint (VTEP). To get more information about SVLBNHs, VENHs, and remote VTEPs, see "show ethernet-switching vxlan-tunnel-end-point svlbnh" on page 19733.
Active source	IP address or Ethernet segment identifier (ESI) of remote entity on which MAC address is learned.
MAC count	Number of MAC addresses learned on the specific routing instance or interface.
Learning interface	Name of the logical interface on which the MAC address was learned.
Learning VLAN	VLAN ID of the routing instance or VLAN in which the MAC address was learned.
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning-tree-protocolepoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

For products that do not support ELS: The following table lists the output fields for the `show ethernet-switching table` command on products that do not support ELS. Output fields are listed in the approximate order in which they appear.

Table 498: show ethernet-switching table Output Fields on Products That Do Not Support ELS

Field Name	Field Description	Level of Output
VLAN	The name of a VLAN.	All levels
Tag	The VLAN ID tag name or number.	**extensive**
MAC or **MAC address**	The MAC address associated with the VLAN.	All levels
Type	The type of MAC address. Values are: • **static**—The MAC address is manually created. • **learn**—The MAC address is learned dynamically from a packet's source MAC address. • **flood**—The MAC address is unknown and flooded to all members. • **persistent**—The learned MAC addresses that will persist across restarts of the switch or interface-down events.	All levels except **persistent-mac**
Type	The type of MAC address. Values are: • **installed**—addresses that are in the Ethernet switching table. • **uninstalled**—addresses that could not be installed in the table or were uninstalled in an interface-down event and will be reinstalled in the table when the interface comes back up.	**persistent-mac**
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels
Interfaces	Interface associated with learned MAC addresses or **All-members** (flood entry).	All levels
Learned	For learned entries, the time which the entry was added to the Ethernet switching table.	**detail**, **extensive**

Table 498: show ethernet-switching table Output Fields on Products That Do Not Support ELS *(Continued)*

Field Name	Field Description	Level of Output
Nexthop index	The next-hop index number.	**detail**, **extensive**
persistent-mac	**installed** indicates MAC addresses that are in the Ethernet switching table and **uninstalled** indicates MAC addresses that could not be installed in the table or were uninstalled in an interface-down event (and will be reinstalled in the table when the interface comes back up).	

Sample Output

show ethernet-switching table (Products that support ELS)

```
user@switch> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan1               b0:c6:9a:ca:3c:01   D             -   ae1.0
    vlan1               b0:c6:9a:ca:3c:03   D             -   ae1.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan10              b0:c6:9a:ca:3c:01   D             -   ae1.0
    vlan10              b0:c6:9a:ca:3c:03   D             -   ae1.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
```

```
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan2               b0:c6:9a:ca:3c:01   D             -   ae1.0
    vlan2               b0:c6:9a:ca:3c:03   D             -   ae1.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan3               b0:c6:9a:ca:3c:01   D             -   ae1.0
    vlan3               b0:c6:9a:ca:3c:03   D             -   ae1.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan4               b0:c6:9a:ca:3c:01   D             -   ae1.0
    vlan4               b0:c6:9a:ca:3c:03   D             -   ae1.0
```

show ethernet-switching table (Products that support ELS)

```
user@host> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
```

```
    VLAN101             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN102             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN103             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN104             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1101            00:1f:12:32:f5:c1   D             -   ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
```

```
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1102            00:1f:12:32:f5:c1   D             -   ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1103            00:1f:12:32:f5:c1   D             -   ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1104            00:1f:12:32:f5:c1   D             -   ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1105            00:1f:12:32:f5:c1   D             -   ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1106            00:1f:12:32:f5:c1   D             -   ae0.0
[...output truncated...]
```

show ethernet-switching table brief (Products that support ELS)

```
user@host> show ethernet-switching table brief
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN101             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN102             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN103             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN104             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)
```

```
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1101            00:1f:12:32:f5:c1   D             -   ae0.0
[...output truncated...]
```

show ethernet-switching table (QFX5120-32CD, QFX5120-48Y, QFX5130-32CD, QFX5130-48C, QFXC5130-48CM, QFX5700, EX4100, and EX4400-48P devices, Blocked (B) MAC Address Field)

```
user@host> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC
           B - Blocked)


Ethernet switching table : 3 entries, 3 learned
Routing instance : MACVRF_Blue
   Vlan                            MAC                 MAC      Logical
SVLBNH/      Active
   name                          address                flags      interface              VENH
Index   source
   VLAN-10             aa:bb:cc:dd:ee:ff  DLR      ae6.1010
   VLAN-10             gg:hh:ii:jj:kk:ll   DLB         ae6.1010
   VLAN-10             mm:nn:oo:pp:qq:rr    DRB
vtep-58.32778                            192.168.1.0
```

show ethernet-switching table count (Products that support ELS)

```
user@host> show ethernet-switching table count
0 MAC address learned in routing instance default-switch VLAN VLAN1000
ae26.0:1000

1 MAC address learned in routing instance default-switch VLAN VLAN101
ae20.0:101

  MAC address count per learn VLAN within routing instance:
```

```
    Learn VLAN ID      MAC count      Static MAC count
            101              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN102
ae20.0:102

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
            102              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN103
ae20.0:103

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
            103              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN104
ae20.0:104

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
            104              1                 0

0 MAC address learned in routing instance default-switch VLAN VLAN105
ae20.0:105

0 MAC address learned in routing instance default-switch VLAN VLAN106
ae20.0:106

0 MAC address learned in routing instance default-switch VLAN VLAN107
ae20.0:107

0 MAC address learned in routing instance default-switch VLAN VLAN108
ae20.0:108

0 MAC address learned in routing instance default-switch VLAN VLAN109
ae20.0:109

0 MAC address learned in routing instance default-switch VLAN VLAN110
ae20.0:110

1 MAC address learned in routing instance default-switch VLAN VLAN1101
```

```
ae0.0:1101

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
          1101              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN1102
ae0.0:1102

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
          1102              1                 0
[...output truncated...]
```

show ethernet-switching table extensive (Products that support ELS)

```
user@host> show ethernet-switching table extensive

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 101
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 2
   Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 102
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 2
   Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 103
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 2
   Learning mask: 0x00000008
```

```
MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 104
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                           Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1101
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                           Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1102
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                           Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1103
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                           Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1104
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                           Sequence number: 2
   Learning mask: 0x00000008
```

show ethernet-switching table (EX Series switches, GBP Tag field)

```
user@switch>  show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)
Ethernet switching table : 4 entries, 4 learned
Routing instance : default-switch
   Vlan                MAC                 MAC       GBP    Logical                SVLBNH/
Active
   name                address             flags     tag    interface              VENH Index
source
   vlan1000            c8:e7:f0:4b:d1:00   DRP              esi.1699               1698
05:00:00:fd:e9:00:00:03:e8:00
   vlan1000            dc:38:e1:e0:30:c0   D          300   ae0.0
   vlan1000            00:21:59:aa:77:f0   DR               vtep.32769
18.18.18.18
   vlan2000            00:00:5e:00:01:01   DR               vtep.32769
18.18.18.18
```

show ethernet-switching table (QFX Series switches, SVLBNH/VENH field)

```
user@switch> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 16 entries, 16 learned Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                SVLBNH/      Active
   name                address             flags    interface              VENH Index   source
   vlanBLACK           00:00:5e:00:53:01   DR       esi.1773               1782
05:00:00:02:9a:00:00:00:68:00
   vlanBLACK           00:00:5e:00:53:02   DR       esi.1773               1782
05:00:00:02:9a:00:00:00:68:00
   vlanBLACK           00:00:5e:00:53:80   D        vtep.32772
10.0.0.1
   vlanBLACK           00:00:5e:00:53:00   D        vtep.32769
10.0.1.1
   vlanBLUE            00:00:5e:00:53:01   DR       esi.1772               1782
05:00:00:02:9a:00:00:00:66:00
   vlanBLUE            00:00:5e:00:53:02   DR       esi.1772               1782
```

```
05:00:00:02:9a:00:00:00:66:00
   vlanBLUE            00:00:5e:00:53:80   D        vtep.32772
10.0.0.1
   vlanBLUE            00:00:5e:00:53:00   D        vtep.32769
10.0.1.1
   vlanGREEN           00:00:5e:00:53:01   DR       esi.1774               1782
05:00:00:02:9a:00:00:00:67:00
   vlanGREEN           00:00:5e:00:53:02   DR       esi.1774               1782
05:00:00:02:9a:00:00:00:67:00
   vlanGREEN           00:00:5e:00:53:80   D        vtep.32772
10.0.0.1
   vlanGREEN           00:00:5e:00:53:00   D        vtep.32769
10.0.1.1
   vlanRED             00:00:5e:00:53:01   DR       esi.1771               1782
05:00:00:02:9a:00:00:00:65:00
   vlanRED             00:00:5e:00:53:02   DR       esi.1771               1782
05:00:00:02:9a:00:00:00:65:00
   vlanRED             00:00:5e:00:53:80   D        vtep.32772
10.0.0.1
```

show ethernet-switching table persistent-mac interface ge-0/0/16.0 (EX Series switches)

```
VLAN              MAC address       Type       Interface
default           00:10:94:00:05:02 uninstalled ge-0/0/16.0
default           00:10:94:00:06:03 uninstalled ge-0/0/16.0
default           00:10:94:00:07:04 uninstalled ge-0/0/16.0
```

show ethernet-switching table (Products that do not support ELS)

```
user@switch> show ethernet-switching table
Ethernet-switching table: 57 entries, 17 learned
  VLAN              MAC address       Type         Age Interfaces
  F2                *                 Flood          - All-members
  F2                00:00:05:00:00:03 Learn          0 xe-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static         - Router
  Linux             *                 Flood          - All-members
  Linux             00:19:e2:50:7d:e0 Static         - Router
  Linux             00:30:48:90:54:89 Learn          0 xe-0/0/47.0
  T1                *                 Flood          - All-members
  T1                00:00:05:00:00:01 Learn          0 xe-0/0/46.0
```

```
  T1                00:00:5e:00:01:00 Static         - Router
  T1                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static         - Router
  T10               *                 Flood          - All-members
  T10               00:00:5e:00:01:09 Static         - Router
  T10               00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static         - Router
  T111              *                 Flood          - All-members
  T111              00:19:e2:50:63:e0 Learn          0 xe-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static         - Router
  T111              00:19:e2:50:ac:00 Learn          0 xe-0/0/15.0
  T2                *                 Flood          - All-members
  T2                00:00:5e:00:01:01 Static         - Router
  T2                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static         - Router
  T3                *                 Flood          - All-members
  T3                00:00:5e:00:01:02 Static         - Router
  T3                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static         - Router
  T4                *                 Flood          - All-members
  T4                00:00:5e:00:01:03 Static         - Router
  T4                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
[output truncated]
```

show ethernet-switching table (Private VLANs on products that do not support ELS)

```
user@switch> show ethernet-switching table
Ethernet-switching table: 10 entries, 3 learned
  VLAN              MAC address       Type         Age Interfaces
  pvlan             *                 Flood          - All-members
  pvlan             00:10:94:00:00:02 Replicated        - xe-0/0/28.0
  pvlan             00:10:94:00:00:35 Replicated        - xe-0/0/46.0
  pvlan             00:10:94:00:00:46 Replicated        - xe-0/0/4.0
  c2                *                 Flood          - All-members
  c2                00:10:94:00:00:02 Learn          0 xe-0/0/28.0
  c1                *                 Flood          - All-members
  c1                00:10:94:00:00:46 Learn          0 xe-0/0/4.0
  __pvlan_pvlan_xe-0/0/46.0__ *       Flood          - All-members
  __pvlan_pvlan_xe-0/0/46.0__ 00:10:94:00:00:35 Learn        0 xe-0/0/46.0
```

show ethernet-switching table detail (Products that do not support ELS)

```
user@switch> show ethernet-switching table detail
Ethernet-switching table: 57 entries, 17 learned
  F2, *
    Interface(s): xe-0/0/44.0
    Type: Flood
    Nexthop index: 0

  F2, 00:00:05:00:00:03
    Interface(s): xe-0/0/44.0
    Type: Learn, Age: 0, Learned: 2:03:09
    Nexthop index: 0

  F2, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  Linux, *
    Interface(s): xe-0/0/47.0
    Type: Flood
    Nexthop index: 0

  Linux, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  Linux, 00:30:48:90:54:89
    Interface(s): xe-0/0/47.0
    Type: Learn, Age: 0, Learned: 2:03:08
    Nexthop index: 0

  T1, *
    Interface(s): xe-0/0/46.0
    Type: Flood
    Nexthop index: 0

  T1, 00:00:05:00:00:01
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:07
```

```
    Nexthop index: 0

  T1, 00:00:5e:00:01:00
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T1, 00:19:e2:50:63:e0
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:07
    Nexthop index: 0

  T1, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T10, *
    Interface(s): xe-0/0/46.0
    Type: Flood
    Nexthop index: 0

  T10, 00:00:5e:00:01:09
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T10, 00:19:e2:50:63:e0
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:08
    Nexthop index: 0

  T10, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T111, *
    Interface(s): xe-0/0/15.0
    Type: Flood
    Nexthop index: 0
[output truncated]
```

show ethernet-switching table extensive (Products that do not support ELS)

```
user@switch> show ethernet-switching table extensive
Ethernet-switching table: 57 entries, 17 learned
  F2, *
    Interface(s): xe-0/0/44.0
    Type: Flood
    Nexthop index: 0

  F2, 00:00:05:00:00:03
    Interface(s): xe-0/0/44.0
    Type: Learn, Age: 0, Learned: 2:03:09
    Nexthop index: 0

  F2, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  Linux, *
    Interface(s): xe-0/0/47.0
    Type: Flood
    Nexthop index: 0

  Linux, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  Linux, 00:30:48:90:54:89
    Interface(s): xe-0/0/47.0
    Type: Learn, Age: 0, Learned: 2:03:08
    Nexthop index: 0

  T1, *
    Interface(s): xe-0/0/46.0
    Type: Flood
    Nexthop index: 0

  T1, 00:00:05:00:00:01
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:07
```

```
    Nexthop index: 0

  T1, 00:00:5e:00:01:00
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T1, 00:19:e2:50:63:e0
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:07
    Nexthop index: 0

  T1, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T10, *
    Interface(s): xe-0/0/46.0
    Type: Flood
    Nexthop index: 0

  T10, 00:00:5e:00:01:09
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T10, 00:19:e2:50:63:e0
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:08
    Nexthop index: 0

  T10, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T111, *
    Interface(s): xe-0/0/15.0
    Type: Flood
    Nexthop index: 0
[output truncated]
```

show ethernet-switching table interface (Products that do not support ELS)

```
user@switch> show ethernet-switching table interface xe-0/0/1
Ethernet-switching table: 1 unicast entries
  VLAN              MAC address       Type         Age Interfaces
  V1                *                 Flood          - All-members
  V1                00:00:05:00:00:05 Learn          0 xe-0/0/1.0
```

show ethernet-switching table persistent-mac (Products that do not support ELS)

```
user@switch> show ethernet-switching table persistent-mac
VLAN              MAC address       Type       Interface
default           00:10:94:00:00:02 installed    ge-0/0/42.0
default           00:10:94:00:00:03 installed    ge-0/0/42.0
default           00:10:94:00:00:04 installed    ge-0/0/42.0
default           00:10:94:00:00:05 installed    ge-0/0/42.0
default           00:10:94:00:00:06 installed    ge-0/0/42.0
default           00:10:94:00:05:02 uninstalled ge-0/0/16.0
default           00:10:94:00:06:03 uninstalled ge-0/0/16.0
default           00:10:94:00:07:04 uninstalled ge-0/0/16.0
```

show ethernet-switching table vlan-name v100 (Products that support ELS)

```
user@host> show ethernet-switching table vlan-name v100
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Pseudo mac
Ethernet switching table : 11 entries, 11 learned
Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                Active
   name                address             flags    interface              source
   v100                00:00:00:00:02:80   S,NM     vtep.1074102274
112.1.1.1
   v100                00:00:00:00:03:80   S,NM     vtep.1074102275
115.1.1.1
   v100                00:00:00:00:04:80   S,NM     vtep.1074102276
116.1.1.1
   v100                00:00:00:00:05:80   S,NM     vtep.1074102277
```

```
114.1.1.1
   v100                00:00:00:00:06:80   S,NM     vtep.1074102278
101.1.1.1
   v100                00:00:00:00:07:80   S,NM     vtep.1074102279
102.1.1.1
   v100                00:00:00:00:08:80   S,NM     vtep.1074102280
103.1.1.1
   v100                00:00:00:00:09:80   S,NM     vtep.1074102281
104.1.1.1
   v100                00:00:00:00:0a:80   S,NM     vtep.1074102282
113.1.1.1
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 9.5 for SRX Series.

Options **summary**, **management-vlan**, and **vlan** ***vlan-name*** introduced in Junos OS Release 9.6 for EX Series switches.

Option **sort-by** and field name **tag** introduced in Junos OS Release 10.1 for EX Series switches.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Output for private VLANs introduced in Junos OS Release 12.1 for the QFX Series.

Option **persistent-mac** introduced in Junos OS Release 11.4 for EX Series switches.

Command introduced in Junos OS Release 12.3R2.

ELS commands introduced in Junos OS Release 12.3R2 for EX Series switches.

Options `logical-system`, `persistent-learning`, and `summary` introduced in Junos OS Release 13.2X50-D10 (ELS).

Output for shared VXLAN load balancing next hop (SVLBNH) and VXLAN encapsulated next hop (VENH) introduced in Junos OS Release 20.3R1 for QFX Series switches.

Output for pseudo VTEP logical interfaces introduced in Junos OS Release 20.3R1 for QFX Series switches.

`GBP Tag` option introduced in Junos OS Release 22.4R1 for supported QFX5120, EX4100, EX4400, and EX4650 Series switches.

RELATED DOCUMENTATION

- Example: Setting Up Basic Bridging and a VLAN on Switches
- Example: Setting Up Bridging with Multiple VLANs
- Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch
- Example: Setting Up Bridging with Multiple VLANs for EX Series Switches
- Example: Setting Up Q-in-Q Tunneling on EX Series Switches
- *Dynamic Load Balancing in an EVPN-VXLAN Network*
- clear ethernet-switching table
- show ethernet-switching mac-learning-log

show ethernet-switching vxlan-tunnel-end-point esi

IN THIS SECTION

- Syntax | **19724**
- Description | **19725**
- Options | **19725**
- Required Privilege Level | **19725**
- Output Fields | **19725**
- Sample Output | **19726**
- Release Information | **19727**

Syntax

```
show ethernet-switching vxlan-tunnel-end-point esi
<esi-identifier esi-identifier>
<instance instance>
<svlbnh>
```

Description

Displays the ESI (Ethernet Segment Identifier) information on VXLAN tunnel endpoints.

Options

none	Display brief information about VXLAN tunnel endpoint.
esi-identifier *esi-identifier*	(Optional) Display information for a specified ethernet segment identified by the Ethernet segment identifier (ESI) value.
instance *instance*	(Optional) Display information for a specified instance .
svlbnh	(Optional) Display the SVLBNH index for a given ESI-identifier and instance. You must use this option in conjunction with both the `esi-identifier` and `instance` options. For detailed information on a particular SVLBNH index, see "." on page 19733.

Required Privilege Level

admin

Output Fields

Table 499: show ethernet-switching vxlan-tunnel-end-point esi Output Fields

Field Name	Field Description
ESI	Ethernet Segment Identifier.
RTT	Routing instance in which you have configured the ESI.
VLBNH	VXLAN load balance next hop.
INH	Indirect next hop.
ESI-IFL	ESI logical interface.
LOC-IFL and LOCAL-IFL	Local logical interface name for a given Ethernet Segment

Table 499: show ethernet-switching vxlan-tunnel-end-point esi Output Fields *(Continued)*

Field Name	Field Description
ELP-NH	Next hop index for an associated egress link protection (ELP) reroute tunnel. **NOTE**: If you enable ELP tunnels for an Ethernet segment, this command lists information about the ELP tunnels as remote VTEPs.
RVTEP-NH-LIST	List of remote VTEP next hops.
#RVTEPS	Number of remote VTEPs.
RVTEP-IP	IP address of the remote VTEP. **NOTE**: If you enable ELP tunnels for an Ethernet segment, you'll also see any ELP tunnels listed as remote VTEPs under this field heading.
RVTEP-IFL	Index number of the remote VTEP.
VENH	VXLAN encapsulated next hop.
MASK-ID	Position of the RVTEP in the ESI mask. Each RVTEP is assigned a unique position.
FLAGS	Internal flags used for debugging. They help to identify the state of the RVTEP.
MAC-COUNT	Count of MACs learned on RVTEP for a given ESI.
SVLBNH	Index number of the SVLBNH.

Sample Output

show ethernet-switching vxlan-tunnel-end-point esi

```
user@host> show ethernet-switching vxlan-tunnel-end-point esi

PESI                           RTT                      VLNBH INH     ESI-IFL   LOC-IFL   #RVTEPs
00:10:11:12:13:14:15:16:17:11 default-switch           1758  131077  esi.1758            2
Aliasing
```

```
    RVTEP-IP             RVTEP-IFL      VENH     MASK-ID   FLAGS        MAC-COUNT
    111.1.1.1            vtep.32769     1753     0         2            0
    112.1.1.1            vtep.32773     1757     1         2            0
ESI                           RTT                      VLNBH INH     ESI-IFL   LOC-IFL   #RVTEPs
00:10:11:12:13:14:15:16:17:22 default-switch           1768  131080  esi.1768  ae0.0,    1
Aliasing
    RVTEP-IP             RVTEP-IFL      VENH     MASK-ID   FLAGS        MAC-COUNT
    113.1.1.1            vtep.32774     1767     0         2            0
ESI                           RTT                      VLNBH INH     ESI-IFL   LOC-IFL   #RVTEPs
05:00:00:00:64:00:00:00:64:00 default-switch           1762  131079  esi.1762            3
    RVTEP-IP             RVTEP-IFL      VENH     MASK-ID   FLAGS        MAC-COUNT
    101.1.1.1            vtep.32770     1754     0         2            2
    103.1.1.1            vtep.32771     1755     1         2            2
    102.1.1.1            vtep.32772     1756     2         2            2
ESI                           RTT                      VLNBH INH     ESI-IFL   LOC-IFL   #RVTEPs
05:00:00:00:64:00:00:00:c8:00 default-switch           1760  131078  esi.1760            3
    RVTEP-IP             RVTEP-IFL      VENH     MASK-ID   FLAGS        MAC-COUNT
    101.1.1.1            vtep.32770     1754     0         2            2
    103.1.1.1            vtep.32771     1755     1         2            2
    102.1.1.1            vtep.32772     1756     2         2            2
```

show ethernet-switching vxlan-tunnel-end-point esi esi-identifier instance svlbnh

This command is used to retrieve SVLBNH index for a given esi-identifier and instance.

```
user@host> show ethernet-switching vxlan-tunnel-end-point esi esi-identifier
00:20:21:22:23:24:25:26:27:01 instance default-switch svlbnh
RTT                     ESI                            ESI-IFL   VLNBH INH     SVLBNH
default-switch          00:20:21:22:23:24:25:26:27:01 esi.1815  1815  131094  1816
```

Release Information

Command introduced in Junos OS Release 16.1R1.

`svlbnh` option introduced in Junos OS Release 21.2R1.

show ethernet-switching vxlan-tunnel-end-point remote

IN THIS SECTION

- Syntax | 19728
- Description | 19728
- Required Privilege Level | 19728
- Output Fields | 19728
- Sample Output | 19730
- Release Information | 19730

Syntax

```
show ethernet-switching vxlan-tunnel-end-point remote
```

Description

Display information about remote VXLAN tunnel endpoints (VTEPs) for a source VTEP.

Required Privilege Level

admin

Output Fields

Table 500: show ethernet-switching vxlan-tunnel-end-point remote Output Fields

Field Name	Field Description
Logical System Name	Name of the logical system on which you have configured the related MAC-VRF routing instance or other EVPN instance type.
Id	Logical system ID

Table 500: show ethernet-switching vxlan-tunnel-end-point remote Output Fields *(Continued)*

Field Name	Field Description
L2-RTT	Layer 2 routing instance name
SVTEP-IP	IP address of the source VTEP
IFL	Logical interface name
L3-Idx	System index number of the associated Layer 3 interface
SVTEP-Mode	Source VTEP mode
ELP-SVTEP-IP	IP address of an associated egress link protection (ELP) tunnel for the source VTEP (if applicable)
RVTEP-IP	Remote VTEP IP address
L2-RTT	Layer 2 routing instance name
IFL-Idx	Index number of the remote VTEP interface
Interface	VTEP interface name
NH-Id	Remote VTEP next hop ID
RVTEP Mode	Remote VTEP mode
ELP-IP	IP address of an associated ELP tunnel for the remote VTEP (if applicable). This field contains "ELP" instead of an IP address in RVTEP-IP rows that correspond to the ELP tunnel VTEP for an RVTEP.
Flags	Internal flags used for debugging. They help to identify the state of the RVTEP.
VNID	Virtual network ID mapped to the VLANs associated with the remote VTEP.

Table 500: show ethernet-switching vxlan-tunnel-end-point remote Output Fields *(Continued)*

Field Name	Field Description
MC-Group-IP	Multicast group address (if applicable)

Sample Output

show-ethernet-switching-vxlan-tunnel-end-point-remote

```
user@device> show ethernet-switching vxlan-tunnel-end-point remote
Logical System Name       Id  SVTEP-IP         IFL   L3-Idx    SVTEP-Mode    ELP-SVTEP-IP
<default>                 0   192.168.0.1      lo0.0    0                    192.168.102.1
 RVTEP-IP      IFL-Idx     Interface    NH-Id   RVTEP-Mode  ELP-IP        Flags
 192.168.0.2     2756      vtep.32830   11821   RNVE        192.168.201.1
 192.168.201.1   2757      vtep.32831   11822   RNVE        ELP
    VNID          MC-Group-IP
    530           0.0.0.0
    531           0.0.0.0
```

Release Information

Command introduced in Junos OS Release 21.3R1.

show ethernet-switching vxlan-tunnel-end-point source

IN THIS SECTION

- Syntax | 19731
- Description | 19731
- Required Privilege Level | 19731
- Output Fields | 19731
- Sample Output | 19732

Release Information | 19733

Syntax

```
show ethernet-switching vxlan-tunnel-end-point source
```

Description

Display information about source VXLAN tunnel endpoints (VTEPs), including source VTEP IP address, interface, mode, and any associated egress link protection (ELP) reroute tunnel, bridge domains (VLANs), and translated VNI information for data center interconnect (DCI) (if applicable).

Required Privilege Level

admin

Output Fields

Table 501: show ethernet-switching vxlan-tunnel-end-point source Output Fields

Field Name	Field Description
Logical System Name	Name of the logical system on which you have configured the related MAC-VRF routing instance or other EVPN instance type.
Id	Logical system ID
L2-RTT	Layer 2 routing instance name
SVTEP-IP	IP address of the source VTEP
IFL	Logical interface name
L3-Idx	System index number of the associated Layer 3 interface

Table 501: show ethernet-switching vxlan-tunnel-end-point source Output Fields *(Continued)*

Field Name	Field Description
SVTEP-Mode	Source VTEP mode
ELP-SVTEP-IP	IP address of an associated ELP tunnel for the source VTEP
L2-RTT	Layer 2 routing instance name
Bridge Domain	Associated bridge domains (VLANs)
VNID	VLAN network ID. This is the original VNID in the data center
Translation-VNID	Translation VLAN network ID. This is the translated VNID that is being used in the packet as it exits the data center.
MC-Group-IP	Multicast group IP address

Sample Output

show ethernet-switching vxlan-tunnel-end-point source

```
user@host> show ethernet-switching vxlan-tunnel-end-point source
Logical System Name       Id  SVTEP-IP         IFL   L3-Idx    SVTEP-Mode    ELP-SVTEP-IP
<default>                 0   193.168.4.1      lo0.0    0
    L2-RTT                              Bridge Domain            VNID   Translation-VNID    MC-
Group-IP
    MACVRF-mac-vrf-ep-t2-stchd-transl-1 EP-TYPE-2-VLAN-1000+1000 301000 920200       0.0.0.0
    MACVRF-mac-vrf-ep-t2-stchd-transl-1 EP-TYPE-2-VLAN-1001+1001 301001 920201       0.0.0.0
    MACVRF-mac-vrf-ep-t2-stchd-transl-1 EP-TYPE-2-VLAN-1002+1002 301002 920202       0.0.0.0
    MACVRF-mac-vrf-ep-t2-stchd-transl-1 EP-TYPE-2-VLAN-1003+1003 301003 920203       0.0.0.0
    MACVRF-mac-vrf-ep-t2-stchd-transl-1 EP-TYPE-2-VLAN-1004+1004 301004 920204       0.0.0.0
    MACVRF-mac-vrf-ep-t2-stchd-transl-1 EP-TYPE-2-VLAN-1005+1005 301005 920205       0.0.0.0
```

Release Information

Command introduced in Junos OS Release 21.3R1.

show ethernet-switching vxlan-tunnel-end-point svlbnh

IN THIS SECTION

- Syntax | 19733
- Description | 19733
- Options | 19733
- Required Privilege Level | 19734
- Output Fields | 19734
- Sample Output | 19735
- Sample Output | 19736
- Release Information | 19736

Syntax

```
show ethernet-switching vxlan-tunnel-end-point svlbnh
<brief | detail | extensive>
<overflow>
<svlbnh-index index-number>
```

Description

Display information about a shared VXLAN load balancing next hop (SVLBNH) and the virtual tunnel endpoints (VTEPs) associated with the SVLBNH.

Options

none (Same as `brief`) Display information for all SVLBNHs.

brief | detail | extensive	(Optional) Display the specified level of output.
overflow	(Optional) Display the pending requests for SVLBNH allocation.
svlbnh-index *index-number*	(Optional) Display information about the specified SVLBNH index.

Required Privilege Level

admin

Output Fields

Table 502 on page 19734 lists the output fields for the `show ethernet-switching vxlan-tunnel-end-point svlbnh` command. Output fields are listed in the approximate order in which they appear.

Table 502: show ethernet-switching vxlan-tunnel-end-point svlbnh Output

Field Name	Field Description
SVLBNH *index-number*	Index number of the SVLBNH.
RVTEP count	Number of remote VTEPs associated with the SVLBNH.
ESI count	Number of Ethernet segment identifiers (ESIs) associated with the SVLBNH.
RVTEP-IP	IP address of the remote VTEP.
RVTEP-IFL	Index number of remote VTEP.
VENH	VXLAN encapsulated next hop (VENH) associated with the SVLBNH.
ESI	ESIs associated with the SVLBNH.
RTT	Routing instance in which you have configured the ESI.

Sample Output

show ethernet-switching vxlan-tunnel-end-point svlbnh detail

```
user@switch> show ethernet-switching vxlan-tunnel-end-point
svlbnh detail
SVLBNH: 1882        RVTEP count: 2           ESI count: 1
    RVTEP-IP             RVTEP-IFL            VENH
    10.0.5.1             vtep.32770           1859
    10.0.3.1             vtep.32769           1835
    ESI                           RTT
    00:00:00:00:00:00:00:00:01:03 default-switch

SVLBNH: 1879        RVTEP count: 2           ESI count: 1
    RVTEP-IP             RVTEP-IFL            VENH
    10.0.2.1             vtep.32771           1869
    10.0.3.1             vtep.32769           1835
    ESI                           RTT
    00:00:00:00:00:00:00:00:01:01 default-switch

SVLBNH: 1881        RVTEP count: 2           ESI count: 1
    RVTEP-IP             RVTEP-IFL            VENH
    10.0.2.1             vtep.32771           1869
    10.0.5.1             vtep.32770           1859
    ESI                           RTT
    00:00:00:00:00:00:00:00:01:02 default-switch

SVLBNH: 1870        RVTEP count: 3           ESI count: 5
    RVTEP-IP             RVTEP-IFL            VENH
    10.0.2.1             vtep.32771           1869
    10.0.5.1             vtep.32770           1859
    10.0.3.1             vtep.32769           1835
    ESI                           RTT
    00:00:00:00:00:00:00:00:01:04 default-switch
    05:00:00:02:9a:00:00:00:65:00 default-switch
    05:00:00:02:9a:00:00:00:66:00 default-switch
    05:00:00:02:9a:00:00:00:67:00 default-switch
    05:00:00:02:9a:00:00:00:68:00 default-switch
```

Sample Output

show ethernet-switching vxlan-tunnel-end-point svlbnh (Specific SVLBNH)

```
user@switch> show ethernet-switching vxlan-tunnel-end-point
svlbnh svlbnh-index 1882
SVLBNH: 1882        RVTEP count: 2           ESI count: 1
    RVTEP-IP             RVTEP-IFL            VENH
    10.0.5.1             vtep.32770           1859
    10.0.3.1             vtep.32769           1835
    ESI                           RTT
    00:00:00:00:00:00:00:00:01:03 default-switch
```

Release Information

Command introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Dynamic Load Balancing in an EVPN-VXLAN Network

show ethernet-switching table | 19700

show evpn arp-table

IN THIS SECTION

- Syntax | 19737
- Description | 19737
- Options | 19737
- Required Privilege Level | 19737
- Output Fields | 19738
- Sample Output | 19738
- Release Information | 19740

Syntax

```
show evpn arp-table
<address>
<brief |count | detail | extensive>
<instance instance-name>
<logical-system logical-system-name>
```

Description

Show Ethernet VPN (EVPN) Address Resolution Protocol (ARP) entries associated with learned MAC addresses.

Routing platforms support the following command that displays the same information and includes the same options: `show bridge evpn arp-table`.

> **NOTE**: Devices learn the MAC addresses associated with IPv6 addresses using the Network Discovery Protocol (NDP) instead of with ARP. As a result, for IPv6 traffic, use "show evpn nd-table" on page 19818 to see MAC addresses the device learns from NDP.
>
> Routing platforms also support the `show bridge evpn nd-table` command that displays similar information and includes the same options as the `show evpn nd-table` command.

Options

none	Display brief information about the EVPN ARP table.
address	(Optional) Display ARP information for the specified MAC address.
brief | count | detail |extensive	(Optional) Display the specified level of output.
instance <*instance-name*>	oot(Optional) Display ARP information for the specified routing instance .
logical-system <*logical-system-name*>	(Optional) Display ARP information for the specified logical system or all logical systems.

Required Privilege Level

view

Output Fields

Table 503 on page 19738 lists the output fields for the `show evpn arp-table` command. Output fields are listed in the approximate order in which they appear.

Table 503: show evpn arp-table Output Fields

Field Name	Field Description	Level of Output
INET address	The INET address related to the INET entries that are added to the ARP table.	All levels
MAC address	MAC addresses learned through ARP.	**brief**, **detail**, **extensive**, **instance**, **mac-address**,,
Logical Interface	Logical interface associated with the routing instance in which the ARP INET address is learned.	**brief**, **instance**, **mac-address**,,
Routing instance	Routing instance in which the ARP INET address is learned.	all levels
Bridging domain	Bridging domain in which the ARP INET address is learned.	all levels
Learning interface	Interface on which the ARP INET address is learned.	**detail**, **extensive**
Count	Indicates the number of ARP INET addresses learned in a routing instance in a bridge domain.	**count**

Sample Output

show evpn arp-table

```
user@host> show evpn arp-table
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
172.16.113.2        00:05:86:a0:dc:f0  irb.0         evpn1        __evpn1__
```

show evpn arp-table 00:05:86:a0:dc:f0 (MAC address)

```
user@host> show evpn arp-table 00:05:86:a0:dc:f0
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
172.16.113.2        00:05:86:a0:dc:f0  irb.0         evpn1        __evpn1__
```

show evpn arp-table brief

```
user@host> show evpn arp-table brief
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
172.16.113.2        00:05:86:a0:dc:f0  irb.0         evpn1        __evpn1__
```

show evpn arp-table detail

```
user@host> show evpn arp-table detail

INET address: 172.16.113.2
MAC address: 00:05:86:a0:dc:f0
  Routing instance: evpn1
   Bridging domain: __evpn1__
   Learning interface: irb.0
```

show evpn arp-table count

```
user@switch> show evpn arp-table count
1 ARP INET addresses learned in routing instance evpn1 bridge domain __evpn1__
```

show evpn arp-table extensive

```
user@host> show evpn arp-table extensive

INET address: 172.16.113.2
MAC address: 00:05:86:a0:dc:f0
  Routing instance: evpn1
```

```
    Bridging domain: __evpn1__
    Learning interface: irb.0
```

show evpn arp-table instance evpn1

```
user@host> show evpn arp-table instance evpn1
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
172.16.113.2        00:05:86:a0:dc:f0  irb.0         evpn1        __evpn1__
```

Release Information

Command introduced in Junos OS Release 15.1.

logical system option introduced in Junos OS Release 18.1R1.

show evpn database

IN THIS SECTION

- Syntax | 19741
- Syntax for mac-vrf routing | 19741
- Description | 19741
- Options | 19742
- Required Privilege Level | 19742
- Output Fields | 19742
- Sample Output | 19744
- Release Information | 19747

Syntax

```
show evpn database <extensive>
<esi esi>
<instance instance-name>
<interface interface-name>
<logical-system logical-system-name>
<mac-address address>
<neighbor neighbor-name>
<origin origin-name>
<state state-name>
<vlan-id vlan-id>
```

Syntax for mac-vrf routing

```
show mac-vrf routing database <extensive>
<esi esi>
<instance instance-name>
<interface interface-name>
<logical-system logical-system-name>
<mac-address address>
<neighbor neighbor-name>
<origin origin-name>
<state state-name>
<vlan-id vlan-id>
```

Description

Show Ethernet VPN (EVPN) database information.

NOTE: When you configure OAM Connectivity Fault Management (CFM) with an up maintenance endpoint (MEP) interface, the **show evpn database instance** command displays the CFM MAC local interface with the naming convention ".local..*number*" (for example, .local..8).

Options

`none`	Display brief information about the EVPN database.
`extensive`	(Optional) Display detailed information about the EVPN database.
`esi` *esi*	(Optional) Display EVPN database information associated with the specified Ethernet segment ID, including local and remotely learned database entries.
`instance` *instance-name*	(Optional) Display MAC addresses from the specified routing instance.
`interface` *interface-name*	(Optional) Display the MAC address learned locally from the specified interface.
`logical-system` <*logical-system-name*>	(Optional) Display database information for the specified logical system or all logical systems.
`mac-address` *address*	(Optional) Display the specified MAC address.
`neighbor` *neighbor-name*	(Optional) Display the MAC address learned from the specified neighbor.
`origin` *origin-name*	(Optional) Display the MAC address with the specified origin.
`state` *state-name*	(Optional) Display the MAC address with the specified state.
`vlan-id` *vlan-id*	(Optional) Display the MAC address with the specified VLAN.

Required Privilege Level

view

Output Fields

Table 504 on page 19742 lists the output fields for the `show evpn database` command. Output fields are listed in the approximate order in which they appear.

Table 504: show evpn database Output Fields

Field Name	Field Description	Level of Output
`Instance`	Name of the routing instances.	All Levels

Table 504: show evpn database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`VNI`	VXLAN network identifier.	All Levels
`MAC address`	MAC address of the routing instance.	All Levels
`Origin`	Specific origin of the MAC address.	All Levels
`Timestamp`	Month, day and time of generated command output.	All Levels
`IP address`	IP address.	All Levels
`Prefix`	Prefix address of the VXLAN network identifier.	`extensive`
`Source`	Source address of the VXLAN network identifier.	`extensive`
`Rank`	Level number of the VXLAN network identifier.	`extensive`
`Status`	Status of the of the VXLAN network identifier.	`extensive`
`Remote origin`	Specific remote origin of the VXLAN network identifier.	`extensive`
Mobility History	Information on a MAC address mobility history: • Mobility event time—Date and time of the MAC mobility event • Type—Origin of the learned MAC address (Local or Remote). • Source—Source of the learned MAC address (Local Interface, Remote PE IP address, or ESI) • Seq num—Sequence number associated with the mobility event	`extensive`

Table 504: show evpn database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Mac label	Label advertised by the remote PE device when forwarding unicast traffic. In EVPN-MPLS, it is the MPLS label and in EVPN-VXLAN, it is the VNI .	extensive
Mobility Sequence number	Current sequence number associated with the MAC address. The minimum origin address is the IP address of the PE device with the lowest IP address that is sending a MAC advertisement.	extensive
State	Information collected in the EVPN database from different flag states including Duplicate-Detected, Local-Pinned, and Remote-Pinned.	extensive
IP address	IPv4 or IPv6 address for the MAC address.	extensive
L3 route	Route installed on the EVPN IRB interface.	extensive
L3 context	Name of the routing instance that has the Layer 3 routes installed for an EVPN IRB interface, typically a virtual routing and forward (VRF) routing instance.	extensive

Sample Output

show evpn database

```
user@host> show evpn database
Instance: bd_red
VNI  MAC address        Origin             Timestamp        IP address
50    1a:1b:1b:1b:1b:1b  192.0.2.3          Jun 12 21:57:17
50    1a:1c:1c:1c:1c:1c  ge-2/3/0.0         Jun 12 22:05:37
50    b0:c6:9a:e9:cd:d8  192.0.2.3          Jun 12 21:57:17
50    b0:c6:9a:ea:87:42  ge-2/3/0.0         Jun 12 22:05:38

Instance: evpn-0
VNI   MAC address        Origin           Timestamp        IP address
```

```
4000  1a:1a:1a:1a:01:01  ge-2/3/0.10      Jun 12 21:53:01
4000  1a:1a:1a:1a:01:02  ge-2/3/0.10      Jun 12 21:53:01
```

show evpn database extensive

```
user@host> show evpn database extensive
Instance: ALPHA

VN Identifier: 9100, Prefix: 10.0.0.0/24
  Source: 10.255.0.2, Rank: 1, Status: Active
    Timestamp: May 22 17:16:12 (0x555fc6cc)

VN Identifier: 9100, Prefix: 198.51.100.0/24
  Source: 05:00:00:00:64:00:00:23:8c:00, Rank: 1, Status: Active
    Remote origin: 10.255.0.2
    Remote origin: 10.255.0.3
    Timestamp: May 22 17:14:20 (0x555fc65c)

VN Identifier: 9100, Prefix: 192.0.2.0/24
  Source: irb.0, Rank: 1, Status: Active
    Timestamp: May 22 17:16:12 (0x555fc6cc)

VN Identifier: 9100, Prefix: 203.0.113.0/24
  Source: 10.255.0.3, Rank: 1, Status: Active
    Timestamp: May 22 17:16:12 (0x555fc6cc)
```

show evpn database extensive instance

```
user@PE1> show evpn database instance ALPHA mac-address 00:00:00:00:00:01
extensive
Instance: ALPHA

VLAN ID: 100, MAC address: 00:00:00:00:00:01
  Nexthop ID: 1048575
  Mobility history
    Mobility event time     Type    Source                                  Seq num
    Jul 10 16:26:14.920136  Remote  10.255.0.3                              13
    Jul 10 16:26:16.174769  Local   ge-0/0/2.0                              14
    Jul 10 16:26:17.868187  Remote  10.255.0.3                              15
    Jul 10 16:26:19.129879  Local   ge-0/0/2.0                              16
```

```
    Jul 10 16:26:25.972747  Remote  10.255.0.3                              17
  Source: 10.255.0.3, Rank: 1, Status: Active
    MAC label: 299776
    Mobility sequence number: 17 (minimum origin address 10.255.0.3)
    Timestamp: Jul 10 16:26:25 (0x59640d21)
    State: <Remote-To-Local-Adv-Done>
    IP address: 10.0.0.3
      L3 route: 10.0.0.3/32, L3 context: DELTA (irb.0)
```

show evpn database extensive (Duplicate MAC Address)

```
user@PE1> show evpn database instance ALPHA mac-address 00:00:00:00:00:012extensive
Instance: ALPHA

VLAN ID: 100, MAC address: 00:00:00:00:00:02
  State: 0x1 <Duplicate-Detected>
  Mobility history
    Mobility event time     Type    Source                                  Seq num
    Aug 03 17:22:28.585619  Local   ge-0/0/2.0                              31
    Aug 03 17:22:30.307198  Remote  10.255.0.3                              32
    Aug 03 17:22:37.611786  Local   ge-0/0/2.0                              33
    Aug 03 17:22:39.289357  Remote  10.255.0.3                              34
    Aug 03 17:22:45.609449  Local   ge-0/0/2.0                              35
  Source: ge-0/0/2.0, Rank: 1, Status: Active
    Mobility sequence number: 35 (minimum origin address 10.255.0.2)
    Timestamp: Aug 03 17:22:44 (0x5983be54)
    State: <Local-MAC-Only Local-To-Remote-Adv-Allowed>
    MAC advertisement route status: Not created (duplicate MAC suppression)
    IP address: 10.0.0.2
  Source: 10.255.0.3, Rank: 2, Status: Inactive
    MAC label: 300176
    Mobility sequence number: 34 (minimum origin address 10.255.0.3)
    Timestamp: Aug 03 17:22:39 (0x5983be4f)
    State: <>
    MAC advertisement route status: Not created (inactive source)
    IP address: 10.0.0.3
```

show evpn database instance extensive mac-address extensive (Duplicate MAC address)

```
user@PE1> show evpn database instance ALPHA mac-address aa:bb:cc:dd:ee:ff extensive
Instance: ALPHA

VLAN ID: 100, MAC address: aa:bb:cc:dd:ee:ff
  State: 0x1 <Duplicate-Detected>
  Mobility history
    Mobility event time     Type    Source                                  Seq num
    Aug 03 17:22:28.585619  Local   ge-0/0/2.0                              31
    Aug 03 17:22:30.307198  Remote  10.0.255.3                              32
    Aug 03 17:22:37.611786  Local   ge-0/0/2.0                              33
    Aug 03 17:22:39.289357  Remote  10.0.255.3                              34
    Aug 03 17:22:45.609449  Local   ge-0/0/2.0                              35
  Source: ge-0/0/2.0, Rank: 1, Status: Active
    Mobility sequence number: 35 (minimum origin address 10.255.0.2)
    Timestamp: Aug 03 17:22:44 (0x5983be54)
    State: <Local-MAC-Only Local-To-Remote-Adv-Allowed>
    MAC advertisement route status: Not created (duplicate MAC suppression)
    IP address: 10.1.1.2
  Source: 10.0.255.3, Rank: 2, Status: Inactive
    MAC label: 300176
    Mobility sequence number: 34 (minimum origin address 10.0.255.3)
    Timestamp: Aug 03 17:22:39 (0x5983be4f)
    State: <>
    MAC advertisement route status: Not created (inactive source)
    IP address: 10.1.1.3
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Example: Configuring an EVPN with IRB Solution on EX9200 Switches

show evpn flood

IN THIS SECTION

- Syntax | 19748
- Description | 19748
- Options | 19748
- Required Privilege Level | 19749
- Release Information | 19749

Syntax

```
show evpn flood
event-queue
<instance instance-name>
<logical-system logical-system-name>
<route route-name>
```

Description

Show Ethernet VPN (EVPN) flooding information.

Options

none	Display brief information about EVPN flooding.
brief | detail |extensive | summary	(Optional) Display the specified level of output.
event-queue	(Optional) Display the queue of pending EVPN flood events.
instance *instance-name*	(Optional) Display flooding information for the specified routing instance.
logical-system <*logical-system-name*>	(Optional) Display flooding information for the specified logical system or all logical systems.

`route` *`route-name`*	(Optional) Display flooding information for the specified route.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Example: Configuring an EVPN with IRB Solution on EX9200 Switches

show evpn gbp-src-tag

IN THIS SECTION

- Syntax | 19749
- Description | 19750
- Options | 19750
- Required Privilege Level | 19750
- Output Fields | 19750
- Sample Output | 19751
- Release Information | 19752

Syntax

```
show evpn gbp-src-tag filter-bind routing-instance [instance]
show evpn gbp-src-tag ip-prefix inet | inet6
```

Description

Show the associations (bindings) between GBP filters and routing instances as well as the GBP tags being advertised in EVPN Type 5 IP prefix routes.

Options

filter-bind routing-instance [instance] Display the associations (bindings) between GBP filters and routing instances. If a routing-instance is specified, then only the associations for that routing-instance are shown. If no routing-instance is specified, then the associations for all routing-instances are shown.

ip-prefix inet | inet6 Display the GBP tags propagated in EVPN Type 5 advertisements for the associated IPv4 or IPv6 prefixes. Output is limited to those routes with a GBP tag associated. If a route does not have an associated GBP tag, then the route is not displayed.

Required Privilege Level

view

Output Fields

Table 505 on page 19750 lists the output fields for the `show evpn gbp-src-tag filter-bind routing-instance` command.

Table 505: show evpn gbp-src-tag filter-bind routing-instance Output Fields

Field Name	Field Description
Filter-name	The name of the GBP firewall filter being bound.
Filter-bind-type	The entity type that the GBP filter is bound to. The only valid type is routing-instance.
Filter-bind-name	The name of the routing-instance.

Table 506 on page 19751 lists the output fields for the `show evpn gbp-src-tag ip-prefix inet | inet6` command.

Table 506: show evpn gbp-src-tag ip-prefix inet | inet6 Output Fields

Field Name	Field Description
Prefix-type	Either IPv4 or IPv6.
Prefix-value	The IP prefix in the EVPN Type 5 advertisement.
Src-GBP-Tag	The GBP tag being propagated in the EVPN Type 5 advertisement for the displayed IP prefix.
Prefix-Vrf-Name	The name of the routing-instance.
Filter-Name	The name of the GBP firewall filter that contains the match condition assigning the displayed GBP tag to the displayed IP prefix.

Sample Output

show evpn gbp-src-tag filter-bind routing-instance

```
user@host> show evpn gbp-src-tag filter-bind routing-instance VRF-100
```

```
     Firewall filter association information

    Filter-name      :  TESTD
    Filter-bind-type :  routing-instance
    Filter-bind-name :  VRF-100
```

show evpn gbp-src-tag ip-prefix inet

```
user@host> show evpn gbp-src-tag ip-prefix inet
```

```
     IPV4-Prefix lookup db while advertise prefix in type5

    Prefix-type            :  ipv4
    Prefix-value           :  100.1.0.0/22
```

```
    Src-GBP-Tag            :  100
    Prefix-Vrf-Name        :  VRF-100
    Filter-Name            :  TESTB
```

Release Information

Command introduced in Junos OS Release 24.2R1.

show evpn igmp-snooping database

IN THIS SECTION

- Syntax | 19752
- Syntax for mac-vrf routing | 19753
- Description | 19753
- Options | 19753
- Required Privilege Level | 19753
- Output Fields | 19753
- Sample Output | 19755
- Release Information | 19756

Syntax

```
show evpn igmp-snooping database
<extensive>
<esi id>
<group group-address>
<instance id>
<interface name>
<l2-domain-id id>
<logical-system name | all>
```

Syntax for mac-vrf routing

```
show mac-vrf routing igmp-snooping database
<extensive>
<esi id>
<group group-address>
<instance id>
<interface name>
<l2-domain-id id>
```

Description

Display IGMP snooping information in an EVPN network.

Options

none	Display detailed information.
extensive	(Optional) Display more extensive output.
esi ***esi-id***	(Optional) Display output only for a specified Ethernet segment ID (ESI).
group ***group-address***	(Optional) Display output only for the specified multicast group address range.
instance ***instance-name***	(Optional) Display information for the specified EVPN instance (EVI).
interface ***name***	(Optional) Display information about the specified interface.
l2domain ***l2-domain-id***	(Optional) Display output for the specified bridging domain, VLAN, or virtual network identifier (VNI).

Required Privilege Level

view

Output Fields

Table 507 on page 19754 lists the output fields for the `show evpn igmp-snooping database` command. Output fields are listed in the approximate order in which they appear.

Table 507: show evpn igmp-snooping database Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the EVPN routing instance for which information is displayed.	All levels
VN Identifier	Tunnel network identifier.	All levels
Group IP	Multicast IP address of the multicast group.	All levels
Source IP	IP address of the multicast data source for the group.	All levels
Access OIF Count	Number of outgoing interface (OIF) list entries on the access side.	All levels
Access OIF Count (more information)	More OIF Entry Information: • `Interface`—Interface name • `ESI`—Associated EVPN Ethernet Segment ID for the interface • `Local`—Local receiver. • `Remote`—Remote receiver.	`extensive`
Remote OIF	Information about outgoing remote interfaces: • `Count`—Number of OIF list entries • `Core NH`—Next hop ID corresponding to the next hop interface. OIF Entry Information: • `Interface`—Interface name • `Nbr`—Neighbor ID.	`extensive`

Sample Output

show evpn igmp-snooping database extensive (default switch instance)

```
user@device> show evpn igmp-snooping database extensive
Instance: default-switch
  VLAN ID: 7000
        Group IP: 233.252.0.2
            Access OIF Count: 1
                  Interface                 ESI             Local    Remote
                xe-0/0/1.0    00:01:02:03:04:05:06:07:08:09 0           1
            Remote OIF Count: 4, Core NH: 1048676
                 Interface           Nbr
                vtep.32768    9.255.255.2
                vtep.32768    9.255.255.3
                vtep.32768    9.255.255.4
                vtep.32768    9.255.255.5
```

show evpn igmp-snooping database instance instance-name extensive

```
user@device> show evpn igmp-snooping database instance evpn-B extensive
Instance: evpn-B
  VN Identifier: 477
        Group IP: 233.252.0.1, Source IP: 0.0.0.0
            Access OIF Count: 1
                  Interface                 ESI             Local    Remote
                ae477.477     00:22:44:66:88:00:22:44:66:99 1           0
        Group IP: 233.252.0.2, Source IP: 0.0.0.0
            Access OIF Count: 1
                  Interface                 ESI             Local    Remote
                ae477.477     00:22:44:66:88:00:22:44:66:99 1           0
        Group IP: 233.252.0.3, Source IP: 0.0.0.0
            Access OIF Count: 1
                  Interface                 ESI             Local    Remote
                ae477.477     00:22:44:66:88:00:22:44:66:99 1           0
        Group IP: 233.252.0.4, Source IP: 0.0.0.0
            Access OIF Count: 1
                  Interface                 ESI             Local    Remote
                ae477.477     00:22:44:66:88:00:22:44:66:99 1           0
```

show evpn igmp-snooping database instance l2-domain-id domain-ID

```
user@device> show evpn igmp-snooping database l2-domain-id 100201
Instance: default-switch
    VN Identifier: 100201
        Group IP: 233.252.0.1, Source IP: 0.0.0.0, Access OIF Count: 2
        Group IP: 233.252.0.2, Source IP: 0.0.0.0, Access OIF Count: 2
```

Release Information

Introduced in Junos OS Release 17.2R1 on QFX10000 switches.

Introduced in Junos OS Release 18.2R1 on MX Series routers and EX9200 switches.

Introduced in Junos OS Release 19.4R2 on ACX Series routers.

RELATED DOCUMENTATION

Configure Multicast Forwarding with IGMP Snooping in an EVPN-MPLS Environment

Example: Preserving Bandwidth with IGMP Snooping in an EVPN-VXLAN Environment

Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment

Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment

show evpn igmp-snooping proxy

IN THIS SECTION

- Syntax | 19757
- Syntax for mac-vrf routing | 19757
- Description | 19757
- Options | 19757
- Required Privilege Level | 19758
- Output Fields | 19758

Sample Output | 19759
Release Information | 19760

Syntax

```
show evpn igmp-snooping proxy
<deterministic-ar>
<extensive>
<group group-address>
<instance instance-name>
<l2-domain-id id>
<logical-system name | all>
```

Syntax for mac-vrf routing

```
show mac-vrf routing igmp-snooping proxy
<deterministic-ar>
<extensive>
<group group-address>
<instance instance-name>
<l2-domain-id id>
```

Description

Display EVPN IGMP snooping proxy information.

Options

deterministic-ar (Optional) Display extensive output with additional details on an assisted replication (AR) leaf device when you have configured the leaf device to deterministically steer multicast flows to specified AR replicator devices. See *Deterministic Load Balancing and Traffic Steering to AR Replicators* for more on how this feature works.

extensive	(Optional) Display more extensive output. You can use this option with other options such as `group`, `instance`, and `l2-domain-id`.
group ***group-address***	(Optional) Display output only for the specified multicast group address range.
instance ***instance-name***	(Optional) Display information for the specified EVPN instance.
l2-domain-id ***id***	(Optional) Display output for the specified bridging domain, VLAN, or virtual network identifier (VNI).

Required Privilege Level

view

Output Fields

Table 508 on page 19758 lists the output fields for the show evpn instance command. Output fields are listed in the approximate order in which they appear.

Table 508: show evpn igmp snooping output fields

Field Name	Field Description
Instance	EVPN instance name.
VN Identifier	VXLAN network identifier.
Group	Multicast group IP address.
Source	Multicast source IP address.
Local	Indicates whether the access interface has a listener. A value of 1 indicates there is an interested listener.
Remote	Number of remote interested listeners.
Corenh	Core next hop identifier.

Table 508: show evpn igmp snooping output fields *(Continued)*

Field Name	Field Description
Flood	Indicates whether flooding is used for this multicast group. A value of 0 indicates that selective multicast forwarding is used in the traffic flow. A value of 1 indicates that inclusive multicast forwarding (flooding) is being used in the traffic flow.

Deterministic-AR (Mode/Preferred replicator/Fallback)	Additional information when you enable deterministic load-balancing of multicast traffic with AR. See *Deterministic Load Balancing and Traffic Steering to AR Replicators* for details on that feature. Output fields are similar to those displayed with the extensive option with the following additional output fields as values separated by / characters):
	• Mode—Deterministic AR mode
	  • Strict: Device sends multicast traffic for matching flows only to a specified preferred AR replicator. If that AR replicator is down, the AR leaf device drops the traffic.
	  • Loose: Device can send multicast traffic for matching flows to a fallback AR replicator if the preferred AR replicator is down.
	• Preferred replicator—Preferred AR replicator address
	• Fallback—Address of a fallback AR replicator the device can use in loose mode. Always displays 0.0.0.0 in this position if the mode is "Strict".

Sample Output

show evpn igmp-snooping proxy extensive

```
user@pe1> show evpn igmp-snooping proxy extensive
Instance: default-switch
  VN Identifier: 100
    Group       Source   Local   Remote    Corenh     Flood
    235.1.1.2   0.0.0.0    0       1         131083     0
    235.1.1.4   0.0.0.0    0       1         131080     0
```

show evpn igmp-snooping proxy deterministic-ar (Strict Mode)

```
user@ar-leaf1> show evpn igmp-snooping proxy deterministic-ar
Instance: evpn-inst1
  VN Identifier: 1110
    Group         Source   Local   Remote    Corenh     Flood    Deterministic-AR(Mode/Preferred
replicator/Fallback)
    233.252.0.2   0.0.0.0    0       1        4385        0      N/A
    233.252.0.4   0.0.0.0    0       1        4381        0      Strict/192.168.102.1/0.0.0.0
```

show evpn igmp-snooping proxy deterministic-ar (Loose Mode)

```
user@ar-leaf1> show evpn igmp-snooping proxy deterministic-ar
Instance: evpn-inst1
  VN Identifier: 1110
    Group         Source   Local   Remote    Corenh     Flood    Deterministic-AR(Mode/Preferred
replicator/Fallback)
    233.252.0.3   0.0.0.0    0       1        4447        0      N/A
    233.252.0.5   0.0.0.0    0       1        4453        0      Loose/
192.168.102.1/192.168.102.2
```

Release Information

Introduced in Junos OS Release 18.3R1.

Deterministic-AR output fields added in Junos OS Evolved Release 23.4R2.

RELATED DOCUMENTATION

- *Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment*
- *Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment*
- *Overview of Selective Multicast Forwarding*

show evpn instance

IN THIS SECTION

- Syntax | 19761
- Syntax for mac-vrf routing | 19761
- Description | 19762
- Options | 19762
- Required Privilege Level | 19763
- Output Fields | 19763
- Sample Output | 19768
- Release Information | 19779

Syntax

```
show evpn instance
<brief | extensive>
<backup-forwarder>
<dci>
<designated-forwarder>
<esi esi>
<esi-info esi>
<instance-name>
<neighbor neighbor-address>
<neighbor-info neighbor-address>
```

Syntax for mac-vrf routing

```
show mac-vrf routing instance
<brief | extensive>
<backup-forwarder>
<dci>
<designated-forwarder>
<esi esi>
```

```
<esi-info esi>
<instance-name>
<neighbor neighbor-address>
<neighbor-info neighbor-address>
```

Description

Show Ethernet VPN (EVPN) routing instance information.

When up MEP is configured for an EVPN instance (EVI), the **show evpn instance** displays a default interface without any configuration with a naming convention as .local..*number* (for example, .local..8).

Options

none	Display brief information about the EVPN routing instance.
brief | extensive	(Optional) Display the specified level of output.
backup-forwarder	(Optional) Display IP addresses of all the backup designated forwarders for the Ethernet segment.
bridge-domains	(Optional) Display information on the bridge domain.
dci	(Optional) Display data center interconnect information.
designated-forwarder	(Optional) Display the IP address of the designated forwarder for the Ethernet segment.
esi ***esi***	(Optional) Display brief information about the routing instance associated with the specified Ethernet segment identifier (ESI) value, including information on the route distinguisher, bridge domain, IRB, and other information associated with the ESI.
esi-info ***esi***	(Optional) Display only the ESI information about the routing instance associated with the specified ESI value.
instance-name	(Optional) Display information about the specified routing instance.
neighbor ***neighbor-address***	(Optional) Display the IP address of the EVPN neighbor along with information on the connected bridge domain, route distinguisher, and IRB.
neighbor-info ***neighbor-addres***	(Optional) Display only the IP address of the EVPN neighbor for a routing instance.

Required Privilege Level

view

Output Fields

Table 509 on page 19763 lists the output fields for the `show evpn instance` command. Output fields are listed in the approximate order in which they appear.

Table 509: show evpn instance Output Fields

Field Name	Field Description	Level of Output
Instance	Names of the routing instances.	All levels
Intfs	Total number of interfaces participating in each routing instance, and number of interfaces that are up.	`brief`
IRB intfs	Statistics on the number of integrated routing and bridging (IRB) interfaces for each routing instance: • `Total`—Total number of IRB interfaces. • `Up`—Number of active IRB interfaces. • `Nbrs`—Number of neighbor IRB interfaces.	`brief`
MH ESIs	Number of Ethernet segments per routing instance that connect to a multihomed customer site.	`brief`
MAC addresses	Number of local and remote MAC addresses for each routing instance.	`brief`
Route Distinguisher	Unique route distinguisher associated with this routing instance.	none
VLAN ID	VLAN identifier.	none
Label allocation mode	Label allocation policy for the routing instance.	none

Table 509: show evpn instance Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Encapsulation type	Encapsulation type for VXLAN EVPN instances (EVIs).	`extensive`
Per-instance MAC route label	Label of MAC route for each routing instance.	none
Flow-label	Displays whether a device is configured to support flow-aware transport (FAT) flow labels. A device with `flow-label` enabled advertises that it is capable of support FAT flow labels.	`extensive`
Flow-label-static	Displays whether a device is configured to support FAT flow labels. A device with `flow-label-static` enabled supports FAT Flow labels but does not advertise its support.	`extensive`
Duplicate MAC detection threshold	Number of MAC mobility events detected for a given MAC address before it is identified as a duplicate MAC address.	`extensive`
Duplicate MAC detection window	The time interval used in detecting a duplicate MAC address.	`extensive`
Duplicate MAC auto-recovery time	Length of time a device suppresses a duplicate MAC address. At the end of this duration, MAC address updates will resume.	`extensive`
DF Election preference	Preference value used for the designated forwarder (DF) election: • Highest preference—Default DF election preference. • Lowest preference—Based on the configuration of the `designated-forwarder-preference-least` statement at the `[edit routing-instance` *`routing-instance-name`* `protocols evpn]` hierarchy level.	`extensive`
Per-instance multicast route label	Label of the multicast route for each routing instance.	none

Table 509: show evpn instance Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Total MAC addresses	Total number of local and remote MAC addresses received for each routing instance.	`extensive`
Default gateway MAC addresses	Number of local and remote MAC addresses serving as a default gateway in this routing instance.	`extensive`
Number of local interfaces	Number of local interfaces belonging to this routing instance.	`extensive`
Number of local interfaces up	Number of active local interfaces belonging to this routing instance.	`extensive`
Interface name	Name of interfaces that belong to this routing instance.	`extensive`
ESI	Ethernet segment identifier (ESI) value of the interfaces belonging to this routing instance.	`extensive`
DF Election Algorithm	DF election type: • MOD based—Default DF election algorithm based on the modulo operation. • Preference based—DF election based on the configured preference values for an ESI.	`extensive`
ESI granularity	Type: Per ESI.	extensive
LACP OOS on NDF	Enabled or Disabled	extensive
`Preference`	Preference value for EVPN Multihoming DF election.	`extensive`

Table 509: show evpn instance Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Mode	Mode of operation for each routing instance: • `single-homed`—Default mode and does not require Ethernet segment values to be configured. • `single-active`—EVPN active-standby multihoming mode of operation.	`extensive`
SH label	Split horizon label used for the active-standby multihoming mode of operation.	`extensive`
Number of IRB interfaces	Number of IRB interfaces that belong to this routing instance.	`extensive`
Number of protect interfaces	Number of protect interfaces that belong to this routing instance.	`extensive`
Inteface name	Name of the primary interface.	`extensive`
Protect Interface name	Name of the protect interface.	`extensive`
Status	Protected status of the primary interface. The status is either Protect-inactive or Protect-active.	`extensive`
L3 context	Names of routing instances that have the Layer 3 routes installed for an EVPN IRB interface, typically a virtual routing and forwarding (VRF) routing instance.	`extensive`
Number of bridge domains	Number of bridge domain for each EVPN instance.	`extensive`
MAC label	Locally allocated label advertised per bridge domain for forwarding unicast traffic.	extensive

Table 509: show evpn instance Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Aliasing label	Aliasing is the ability of a remote PE device to load balance Layer 2 unicast traffic on all the other PE devices that have the same Ethernet segment toward a CE device.	extensive
Number of neighbors	Number of neighbors connected to this routing instance and their IP addresses.	`extensive`
MAC address advertisement	Number of MAC address advertisements received from the neighbor.	`extensive`
MAC+IP address advertisement	Number of MAC and IP address advertisements received from the neighbor.	`extensive`
Inclusive multicast	Number of inclusive multicast routes received from the neighbor.	`extensive`
Ethernet auto-discovery	Number of autodiscovery routes per Ethernet segment received from the neighbor.	`extensive`
Number of ethernet segments	Total number of Ethernet segments for the routing instance.	`extensive`
Designated forwarder	IP address of the designated forwarder (DF) for the Ethernet segment.	`extensive`
Backup forwarder	IP address of all the backup designated forwarders or routers that are not designated forwarders for the Ethernet segment. **NOTE**: Immediately after an EVPN interface Ethernet segment identifier value is changed and the new configuration is committed, the `Designated forwarder` information changes to `DF not elected yet` and the backup forwarder information is not displayed until after the election is complete.	`extensive`
SMET Forwarding	Identifes whether a device is using selective mulicast forwarding.	`extensive`

Table 509: show evpn instance Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Nexthop Limit	The number of SMET next hop that is supported.	extensive
Nexthop Usage	The number of SMET next hop currently in use.	extensive
Domain ID	Virtual Network ID (VNI) in the VXLAN network.	none
Trans Domain ID	Translation Virtual Network ID (VNI) in the data center interconnect.	none
Inclusive Multicast Label	Label for the inclusive multicast route.	none
IPv4 Inclusive Multicast Core Nexthop	The identifier for the next hop on the .inclusive multicast route	none

Sample Output

show evpn instance brief

```
user@host> show evpn instance brief
                         Intfs       IRB intfs         MH      MAC addresses
Instance                 Total   Up  Total   Up  Nbrs  ESIs    Local  Remote
ALPHA                        2    2      1    1     2     1        3       4
BETA                         2    2      1    1     2     1        2       4
__default_evpn__             0    0      0    0     1     0        0       0
```

show evpn instance

```
user@host> show evpn instance
Instance: black
  Route Distinguisher: 101:101
  VLAN ID: 100
```

```
  Label allocation mode: Per-instance
  Per-instance MAC route label: 299776
  Per-instance multicast route label: 299792
  Number of local interfaces: 1
  Number of local interfaces up: 1
    Interface name    Static MACs    ESI
    cbp-0.0                  0    0
  Number of neighbors: 1
    192.0.2.1
      Received routes
        MAC address advertisement:          1
        Ethernet auto-discovery:            0
        Inclusive multicast:                1
```

show evpn instance extensive (In Junos OS Release 16.1 and earlier)

```
user@host> show evpn instance extensive
Instance: ALPHA
  Route Distinguisher: 10.255.0.1:100
  Encapsulation type: VXLAN
  Per-instance MAC route label: 300144
  Per-instance multicast route label: 300160
  MAC database status                Local  Remote
    Total MAC addresses:                 3       4
    Default gateway MAC addresses:       1       2
  Number of local interfaces: 2 (2 up)
    Interface name  ESI                             Mode             SH label
    ae0.0           00:11:22:33:44:55:66:77:88:99  single-active
    ge-0/0/2.0      00:00:00:00:00:00:00:00:00:00  single-homed
  Number of IRB interfaces: 1 (1 up)
    Interface name  L3 context
    irb.0           DELTA
  Number of neighbors: 2
    10.255.0.2
      Received routes
        MAC address advertisement:              2
        MAC+IP address advertisement:           3
        Inclusive multicast:                    1
        Ethernet auto-discovery:                1
    10.255.0.3
      Received routes
```

```
        MAC address advertisement:              2
        MAC+IP address advertisement:           2
        Inclusive multicast:                    1
        Ethernet auto-discovery:                0
  Number of ethernet segments: 1
    ESI: 00:11:22:33:44:55:66:77:88:99
      Designated forwarder: 10.255.0.1
      Backup forwarder: 10.255.0.2

Instance: BETA
  Route Distinguisher: 10.255.0.1:300
  Encapsulation type: VXLAN
  VLAN ID: 300
  Per-instance MAC route label: 300176
  Per-instance multicast route label: 300192
  MAC database status                Local  Remote
    Total MAC addresses:                 3       4
    Default gateway MAC addresses:       1       2
  Number of local interfaces: 2 (2 up)
    Interface name  ESI                            Mode             SH label
    ae1.0           00:00:00:00:00:00:00:00:00:00  single-homed
    ge-0/0/4.0      00:22:44:66:88:00:22:44:66:88  single-active
  Number of IRB interfaces: 1 (1 up)
    Interface name  L3 context
    irb.1           DELTA
  Number of neighbors: 2
    10.255.0.2
      Received routes
        MAC address advertisement:              2
        MAC+IP address advertisement:           3
        Inclusive multicast:                    1
        Ethernet auto-discovery:                1
    10.255.0.3
      Received routes
        MAC address advertisement:              2
        MAC+IP address advertisement:           2
        Inclusive multicast:                    1
        Ethernet auto-discovery:                0
  Number of ethernet segments: 1
    ESI: 00:22:44:66:88:00:22:44:66:88
      Designated forwarder: 10.255.0.1
      Backup forwarder: 10.255.0.2
```

```
Instance: __default_evpn__
  Route Distinguisher: 10.255.0.1:0
  Encapsulation type: VXLAN
  VLAN ID: 0
  Per-instance MAC route label: 300208
  Per-instance multicast route label: 300224
  MAC database status                Local  Remote
    Total MAC addresses:                 0       0
    Default gateway MAC addresses:       0       0
  Number of local interfaces: 0 (0 up)
  Number of IRB interfaces: 0 (0 up)
  Number of neighbors: 1
    10.255.0.2
      Received routes
        Ethernet auto-discovery:                0
        Ethernet Segment:                       2
  Number of ethernet segments: 0
```

show evpn instance extensive (In Junos OS Release 16.2 and later)

```
user@host> show evpn instance extensive

user@host> show evpn instance extensive

Instance: ALPHA
  Route Distinguisher: 10.255.0.1:100
  Encapsulation type: VXLAN
  Per-instance MAC route label: 300144
  Per-instance multicast route label: 300160
  MAC database status                Local  Remote
    Total MAC addresses:                 3       4
    Default gateway MAC addresses:       1       2
  Number of local interfaces: 2 (2 up)
    Interface name  ESI                            Mode             SH label
    ae0.0           00:11:22:33:44:55:66:77:88:99  single-active
    ge-0/0/2.0      00:00:00:00:00:00:00:00:00:00  single-homed
  Number of IRB interfaces: 1 (1 up)
    Interface name  L3 context
    irb.0           DELTA
  Number of neighbors: 2
    Address               MAC    MAC+IP        AD        IM        ES
```

```
    10.255.0.2              2         3         1         1         0
    10.255.0.3              2         2         0         1         0
  Number of ethernet segments: 1
    ESI: 00:11:22:33:44:55:66:77:88:99
      Designated forwarder: 10.255.0.1
      Backup forwarder: 10.255.0.2

Instance: BETA
  Route Distinguisher: 10.255.0.1:300
  Encapsulation type: VXLAN
  VLAN ID: 300
  Per-instance MAC route label: 300176
  Per-instance multicast route label: 300192
  MAC database status                Local  Remote
    Total MAC addresses:                 3       4
    Default gateway MAC addresses:       1       2
  Number of local interfaces: 2 (2 up)
    Interface name  ESI                            Mode             SH label
    ae1.0           00:00:00:00:00:00:00:00:00:00  single-homed
    ge-0/0/4.0      00:22:44:66:88:00:22:44:66:88  single-active
  Number of IRB interfaces: 1 (1 up)
    Interface name  L3 context
    irb.1           DELTA
  Number of neighbors: 2
    Address               MAC    MAC+IP        AD        IM        ES
    10.255.0.2              2         3         1         1         0
    10.255.0.3              2         2         0         1         0
  Number of ethernet segments: 1
    ESI: 00:22:44:66:88:00:22:44:66:88
      Designated forwarder: 10.255.0.1
      Backup forwarder: 10.255.0.2

Instance: __default_evpn__
  Route Distinguisher: 10.255.0.1:0
  Encapsulation type: VXLAN
  VLAN ID: 0
  Per-instance MAC route label: 300208
  Per-instance multicast route label: 300224
  MAC database status                Local  Remote
    Total MAC addresses:                 0       0
    Default gateway MAC addresses:       0       0
  Number of local interfaces: 0 (0 up)
  Number of IRB interfaces: 0 (0 up)
```

```
  Number of neighbors: 1
    Address                MAC    MAC+IP        AD        IM        ES
    10.255.0.2               0         0         0         0         2
  Number of ethernet segments: 0
```

show evpn instance extensive (Preference-based DF Election)

```
user@host> show evpn instance EVPN_1 extensive
Instance: EVPN_1
  Route Distinguisher: 1:101
  Per-instance MAC route label: 299792
+ DF Election preference: Lowest preference
  MAC database status                     Local  Remote
    MAC advertisements:                       7      15
    MAC+IP advertisements:                    7      12
    Default gateway MAC advertisements:       4       8
  Number of local interfaces: 2 (2 up)
    Interface name  ESI                            Mode             Status
    ae101.0         00:11:11:11:11:11:11:11:11:01  all-active       Up
    ae102.0         00:11:11:11:11:11:11:11:11:02  all-active       Up
  Number of IRB interfaces: 4 (4 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.101         101            Up     master
    irb.102         102            Up     master
    irb.103         103            Up     master
    irb.104         104            Up     master
  Number of bridge domains: 4
    VLAN  Domain ID   Intfs / up    IRB intf   Mode             MAC sync  IM route label
    101                  1    1     irb.101    Extended         Enabled   299776
    102                  1    1     irb.102    Extended         Enabled   299776
    103                  1    1     irb.103    Extended         Enabled   299776
    104                  1    1     irb.104    Extended         Enabled   299776
  Number of neighbors: 2
    100.0.0.2
      Received routes
        MAC address advertisement:              7
        MAC+IP address advertisement:           4
        Inclusive multicast:                    4
        Ethernet auto-discovery:                4
    100.0.0.4
      Received routes
```

```
        MAC address advertisement:              8
        MAC+IP address advertisement:           8
        Inclusive multicast:                    4
        Ethernet auto-discovery:                0
  Number of ethernet segments: 2
    ESI: 00:11:11:11:11:11:11:11:11:11
      Status: Resolved by IFL ae0.110
      Local interface: ae0.110, Status: Up/Forwarding
      Number of remote PEs connected: 2
        Remote PE        MAC label  Aliasing label  Mode
        100.100.100.3    0          0               single-active
        100.100.100.2    0          0               single-active
      DF Election Algorithm: Preference based
      Designated forwarder: 100.100.100.3, Preference: 200
      Backup forwarder: 100.100.100.1, Preference: 800
      Backup forwarder: 100.100.100.2, Preference: 400
    ESI: 00:11:11:11:11:11:11:11:11:02
      Status: Resolved by IFL ae102.0
      Local interface: ae102.0, Status: Up/Forwarding
      Number of remote PEs connected: 1
        Remote PE        MAC label  Aliasing label  Mode
        100.0.0.2        299792     299792          all-active
      DF Election Algorithm: MOD based
      Designated forwarder: 100.0.0.2
      Backup forwarder: 100.0.0.1
      Last designated forwarder update: Feb 21 22:23:32
      Advertised split horizon label: 300800
```

show evpn instance extensive (Duplicate MAC Address)

```
user@host> show evpn instance ALPHA extensive
Instance: ALPHA
  Route Distinguisher: 10.255.0.2:100
  Per-instance MAC route label: 304192
  Duplicate MAC detection threshold: 5
  Duplicate MAC detection window: 180
  Duplicate MAC auto-recovery time: 10

...
```

show evpn instance extensive (Protected Interface)

```
user@host> show evpn instance blue extensive
Instance: blue
  Route Distinguisher: 10.255.255.1:100
  Per-instance MAC route label: 299776
  MAC database status                     Local  Remote
     MAC advertisements:                      0       0
     MAC+IP advertisements:                   0       0
     Default gateway MAC advertisements:      0       0
  Number of local interfaces: 5 (5 up)
     Interface name  ESI                            Mode             Status     AC-Role
     ae0.0           00:11:22:33:44:55:66:77:88:99  all-active       Up         Root
     ge-0/0/3.0      00:00:00:00:00:00:00:00:00:00  single-homed     Up         Root
     ge-0/0/4.0      00:11:11:11:44:55:66:77:88:99  all-active       Up         Root
     ge-0/0/4.1      00:22:22:22:44:55:66:77:88:99  all-active       Up         Root
     ge-0/0/4.50     00:00:00:00:00:00:00:00:00:00  single-homed     Up         Root
  Number of IRB interfaces: 1 (0 up)
     Interface name  VLAN   VNI    Status  L3 context
     irb.1           25             Down   vrf
  Number of protect interfaces: 1
     Interface name  protect-interface   active-interface
     ge-0/0/3.1       ge-0/0/4.50            ge-0/0/3.1
```

show evpn instance extensive

The following output shows a partial output snippet with information on a bridge domain. The output shows the MAC labels that have been allocated for each bridge domain.

```
user@host> show evpn instance extensive
Instance: evpn-A
.
.
.
  Number of bridge domains: 1
    VLAN  Domain-ID Intfs/up   IRB-intf  Mode            MAC-sync IM-label  MAC-Label v4-SG-sync
IM-core-NH v6-SG-sync IM-core-NH Trans-ID
    601                1  1    irb.0                   Extended    Enabled      300880
300656      Disabled              Disabled
    602                1  1                            Extended      Enabled      300896
```

```
300672      Disabled              Disabled
  Number of neighbors: 3
    Address                MAC    MAC+IP        AD        IM        ES Leaf-label Remote-DCI-Peer
    81.3.3.3                 0         0         2         1         0
    81.4.4.4                 0         0         2         1         0
    81.5.5.5                 1         1         2         1         0
  Number of ethernet segments: 1
    ESI: 00:01:02:03:04:05:06:07:08:09
      Status: Resolved by NH 47084
      Number of remote PEs connected: 3
        Remote-PE        MAC-label  Aliasing-label  Mode
        81.4.4.4         0          18              all-active
        81.3.3.3         0          18              all-active
        81.5.5.5         299920     299920          all-active
  SMET Forwarding: Disabled
```

show evpn instance esi-info

```
user@host> show evpn instance esi-info esi 00:11:22:33:44:55:66:77:88:99
Instance: ALPHA
  Number of ethernet segments: 1
    ESI: 00:11:22:33:44:55:66:77:88:99
      Status: Resolved by IFL ae0.0
      Local interface: ae0.0, Status: Up/Forwarding
      Number of remote PEs connected: 1
        Remote PE        MAC label  Aliasing label  Mode
        10.255.0.2       0          299856          all-active
      DF Election Algorithm: MOD based
      Designated forwarder: 10.255.0.1
      Backup forwarder: 10.255.0.2
      Last designated forwarder update: Feb 04 11:52:47
      Advertised MAC label: 299856
      Advertised aliasing label: 299856
      Advertised split horizon label: 299872
```

show evpn instance esi backup-forwarder (Instance Name with Ethernet Segment Identifier)

```
user@host> show evpn instance ALPHA esi 00:11:22:33:44:55:66:77:88:99
backup-forwarder
Instance: ALPHA
```

```
  Number of ethernet segments: 1
    ESI: 00:11:22:33:44:55:66:77:88:99
      Backup forwarder: 10.255.0.2
```

show evpn instance esi designated-forwarder (Instance Name with Ethernet Segment Identifier)

```
user@host> show evpn instance ALPHA esi 00:11:22:33:44:55:66:77:88:99
designated-forwarder
Instance: ALPHA
  Number of ethernet segments: 1
    ESI: 00:11:22:33:44:55:66:77:88:99
      Designated forwarder: 10.255.0.1
```

show evpn instance (CFM up)

The following shows the partial output listing the interface for show evpn instance command.

```
user@host > show evpn instance evpn_1
Number of local interfaces: 3 (3 up)
    Interface name  ESI                             Mode             Status     AC-Role
    .local..8       00:00:00:00:00:00:00:00:00:00  single-homed     Up         Root
    xe-0/0/11:0.0   00:00:00:00:00:00:00:00:00:00  single-homed     Up         Root
```

show evpn instance (SMET)

```
user@host > show evpn instance evpn_1
Instance: default-switch
  Route Distinguisher: 10.255.0.3:100
  Encapsulation type: VXLAN
  Duplicate MAC detection threshold: 5
  Duplicate MAC detection window: 180
  MAC database status                     Local  Remote
    MAC advertisements:                       4       1
    MAC+IP advertisements:                    4       1
    Default gateway MAC advertisements:       2       0
  Number of local interfaces: 3 (3 up)
    Interface name  ESI                             Mode             Status     AC-Role
```

```
    .local..5       00:00:00:00:00:00:00:00:00:00  single-homed     Up         Root
    xe-0/0/1.0      00:00:00:00:00:00:00:00:00:00  single-homed     Up         Root
    xe-0/0/2.0      00:00:00:00:00:00:00:00:00:00  single-homed     Up         Root
  Number of IRB interfaces: 2 (2 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.0                  100     Up     DELTA
    irb.1                  200     Up     DELTA
  Number of protect interfaces: 0
  Number of bridge domains: 2
    VLAN  Domain ID   Intfs / up    IRB intf   Mode      MAC sync  IM route label  IPv4 SG sync
IPv4 IM core nexthop  IPv6 SG sync  IPv6 IM core nexthop
    100   100            1    1     irb.0      Extended         Enabled   100
Enabled   131074            Disabled
    200   200            1    1     irb.1      Extended         Enabled   200
Enabled   131077            Disabled
  Number of neighbors: 3
    Address               MAC    MAC+IP        AD        IM        ES Leaf-label
    10.255.0.2              0         0         0         2         0
    10.255.0.4              1         1         0         2         0
    10.255.0.5              0         0         0         2         0
  Number of ethernet segments: 2
    ESI: 05:00:00:00:64:00:00:00:64:00
      Local interface: irb.0, Status: Up/Forwarding
    ESI: 05:00:00:00:64:00:00:00:c8:00
      Local interface: irb.1, Status: Up/Forwarding
  Router-ID: 10.255.0.3
  SMET Forwarding: Enabled: Nexthop Limit: 10000 Nexthop Usage: 7
```

show evpn instance (Per ESI)

```
user@host > show evpn instance ALPHA esi 00:11:11:11:11:11:11:11:11:11
Instance: ALPHA
  Number of ethernet segments: 1
    ESI: 00:11:11:11:11:11:11:11:11:11
      Status: Unresolved
      Local interface: ae0.210, Status: Down
      Number of remote PEs connected: 1
        Remote PE        MAC label  Aliasing label  Mode
        10.10.10.2    0          300256          single-active
      DF Election Algorithm: MOD based
      ESI granularity: Per ESI
```

```
        LACP OOS on NDF: Enabled
        Designated forwarder: 10.255.0.1

        Last designated forwarder update: Apr 16 17:08:33
        Advertised MAC label: 301248
        Advertised aliasing label: 301248
        Advertised split horizon label: 300032
```

show evpn instance bridge-domains

```
user@PE3> show evpn instance bridge-domains
VLAN ID: 701
Domain ID: 599
Trans Domain ID: 599
Inclusive Multicast Label: 599
IPv4 Inclusive Multicast Core Nexthop: 1048624
```

Release Information

Command introduced in Junos OS Release 14.1.

show evpn ip-prefix-database

IN THIS SECTION

- Syntax | 19780
- Description | 19780
- Options | 19780
- Required Privilege Level | 19781
- Output Fields | 19781
- Sample Output | 19782
- Release Information | 19789

Syntax

```
show evpn ip-prefix-database
<extensive>
<direction (imported | exported)>
<esi number>
<ethernet-tag number>
<family (inet | inet6)>
<gateway ip-address>
<l3-context routing-instance-name>
<logical-system logical-system-name>
<nexthop ip-address>
<prefix ip-prefix>
```

Description

Display Ethernet VPN (EVPN) database information for imported and exported IP prefixes.

Options

none	Display standard information for all EVPN imported and exported IP prefixes.
extensive	Display the specified level of output.
direction (imported | exported)	(Optional) Display the specified subset of IP prefixes.
esi *number*	(Optional) Display the IP prefix associated with the specified Ethernet segment identifier.
ethernet-tag *number*	(Optional) Display the IP prefix associated with the specified Ethernet tag.
family (inet | inet6)	(Optional) Display IP prefixes for the specified protocol family.
gateway *ip-address*	(Optional) Display the IP Prefix associated with the overlay gateway IP address.
l3-context *routing-instance-name*	(Optional) Display EVPN IP prefix information for the specified Layer 3 virtual routing and forwarding (VRF) instance.

`logical-system <logical-system-name>`	(Optional) Display EVPN IP prefix information for the specified logical system or all logical systems.
nexthop *ip-address*	(Optional) Display EVPN IP prefix information for the specified underlay next-hop IP address.
prefix *ip-prefix*	(Optional) Display EVPN database information for the specified IP prefix.

Required Privilege Level

view

Output Fields

Table 510 on page 19781 lists the output fields for the `show evpn ip-prefix-database` command. Output fields are listed in the approximate order in which they appear.

Table 510: show evpn ip-prefix-database Output Fields

Field Name	Field Description
L3 context	Name of virtual routing and forwarding (VRF) instance.
EVPN Exported Prefixes Prefix	List of exported IP prefixes.
EVPN route status	For exported prefixes only, status of route: `Created`.
EVPN imported Prefixes	List of imported EVPN IP prefixes.
Etag	Ethernet tag
Route distinguisher	IP address identifier for the IP prefix route.
VNI	Virtual network identifier for the Layer 3 virtual and routing forwarding (VRF) for the customer or tenant domain.

Table 510: show evpn ip-prefix-database Output Fields *(Continued)*

Field Name	Field Description
Router MAC	MAC address associated with the IP prefix.
Nexthop/Overlay GW/ESI	For imported IP prefixes, next-hop IP address.
Change flags	Trace flags.
Advertisement mode	For exported IP prefixes, type of next-hop address.
Encapsulation	For exported IP prefixes, type of encapsulation
Remote Advertisements	For imported IP prefixes, route distinguisher identifier, MAC address, virtual network identifier, and BGP next-hop address to remote destination.
Route-status	Status of EVPN route in the IP prefix database: • Accepted • Rejected
Reject-Reason	The reason why a route was rejected. If VNI symmetry is enforced, then the output would display the reason as `Asymmetric-vni rejected`.

Sample Output

show evpn ip-prefix-database

```
user@host > show evpn ip-prefix-database
L3 context: VRF-100

IPv4->EVPN Exported Prefixes
Prefix                                        EVPN route status
100.1.0.0/22                                  Created
100.1.4.0/22                                  Created
```

```
100.1.8.0/22                                Created
100.1.12.0/22                               Created
100.1.16.0/22                               Created

IPv6->EVPN Exported Prefixes
Prefix                                      EVPN route status
1234:100:1::/64                             Created
1234:100:1:4::/64                           Created
1234:100:1:8::/64                           Created
1234:100:1:12::/64                          Created
1234:100:1:16::/64                          Created

EVPN->IPv4 Imported Prefixes
Prefix                                      Etag      IP route status
100.2.0.0/22                                0         Created
  Route distinguisher    VNI/Label  Router MAC         Nexthop/Overlay GW/ESI
  10.255.2.1:100         9101       00:05:86:e1:03:f0  10.255.2.1
  10.255.2.2:100         9101       00:05:86:d0:a6:f0  10.255.2.2
100.2.4.0/22                                0         Created
  Route distinguisher    VNI/Label  Router MAC         Nexthop/Overlay GW/ESI
  10.255.2.1:100         9101       00:05:86:e1:03:f0  10.255.2.1
  10.255.2.2:100         9101       00:05:86:d0:a6:f0  10.255.2.2
100.2.8.0/22                                0         Created
  Route distinguisher    VNI/Label  Router MAC         Nexthop/Overlay GW/ESI
  10.255.2.1:100         9101       00:05:86:e1:03:f0  10.255.2.1
  10.255.2.2:100         9101       00:05:86:d0:a6:f0  10.255.2.2
100.2.12.0/22                               0         Created
  Route distinguisher    VNI/Label  Router MAC         Nexthop/Overlay GW/ESI
  10.255.2.1:100         9101       00:05:86:e1:03:f0  10.255.2.1
  10.255.2.2:100         9101       00:05:86:d0:a6:f0  10.255.2.2
100.2.16.0/22                               0         Created
  Route distinguisher    VNI/Label  Router MAC         Nexthop/Overlay GW/ESI
  10.255.2.1:100         9101       00:05:86:e1:03:f0  10.255.2.1
  10.255.2.2:100         9101       00:05:86:d0:a6:f0  10.255.2.2

EVPN->IPv6 Imported Prefixes
Prefix                                      Etag      IP route status
1234:100:2::/64                             0         Created
  Route distinguisher    VNI/Label  Router MAC         Nexthop/Overlay GW/ESI
  10.255.2.1:100         9101       00:05:86:e1:03:f0  10.255.2.1
  10.255.2.2:100         9101       00:05:86:d0:a6:f0  10.255.2.2
1234:100:2:4::/64                           0         Created
  Route distinguisher    VNI/Label  Router MAC         Nexthop/Overlay GW/ESI
```

```
  10.255.2.1:100         9101       00:05:86:e1:03:f0  10.255.2.1
  10.255.2.2:100         9101       00:05:86:d0:a6:f0  10.255.2.2
1234:100:2:8::/64                           0         Created
  Route distinguisher    VNI/Label  Router MAC         Nexthop/Overlay GW/ESI
  10.255.2.1:100         9101       00:05:86:e1:03:f0  10.255.2.1
  10.255.2.2:100         9101       00:05:86:d0:a6:f0  10.255.2.2
1234:100:2:12::/64                          0         Created
  Route distinguisher    VNI/Label  Router MAC         Nexthop/Overlay GW/ESI
  10.255.2.1:100         9101       00:05:86:e1:03:f0  10.255.2.1
  10.255.2.2:100         9101       00:05:86:d0:a6:f0  10.255.2.2
1234:100:2:16::/64                          0         Created
  Route distinguisher    VNI/Label  Router MAC         Nexthop/Overlay GW/ESI
  10.255.2.1:100         9101       00:05:86:e1:03:f0  10.255.2.1
  10.255.2.2:100         9101       00:05:86:d0:a6:f0  10.255.2.2
```

show evpn ip-prefix-database extensive

```
user@host> show evpn ip-prefix-database extensive
L3 context: VRF-100

IPv4->EVPN Exported Prefixes

Prefix: 100.1.0.0/22
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0

Prefix: 100.1.4.0/22
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0

Prefix: 100.1.8.0/22
  EVPN route status: Created
  Change flags: 0x0
```

```
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0

Prefix: 100.1.12.0/22
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0

Prefix: 100.1.16.0/22
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0

IPv6->EVPN Exported Prefixes

Prefix: 1234:100:1::/64
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0

Prefix: 1234:100:1:4::/64
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0

Prefix: 1234:100:1:8::/64
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
```

```
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0

Prefix: 1234:100:1:12::/64
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0

Prefix: 1234:100:1:16::/64
  EVPN route status: Created
  Change flags: 0x0
  Advertisement mode: Direct nexthop
  Encapsulation: VXLAN
  VNI: 9100
  Router MAC: 00:05:86:28:90:f0

EVPN->IPv4 Imported Prefixes

Prefix: 100.2.0.0/22, Ethernet tag: 0
  IP route status: Created
  Change flags: 0x0
  Remote advertisements:
    Route Distinguisher: 10.255.2.1:100
      VNI: 9101
      Router MAC: 00:05:86:e1:03:f0
      BGP nexthop address: 10.255.2.1
    Route Distinguisher: 10.255.2.2:100
      VNI: 9101
      Router MAC: 00:05:86:d0:a6:f0
      BGP nexthop address: 10.255.2.2

Prefix: 100.2.4.0/22, Ethernet tag: 0
  IP route status: Created
  Change flags: 0x0
  Remote advertisements:
    Route Distinguisher: 10.255.2.1:100
      VNI: 9101
      Router MAC: 00:05:86:e1:03:f0
      BGP nexthop address: 10.255.2.1
```

```
    Route Distinguisher: 10.255.2.2:100
      VNI: 9101
      Router MAC: 00:05:86:d0:a6:f0
      BGP nexthop address: 10.255.2.2

Prefix: 100.2.8.0/22, Ethernet tag: 0
  IP route status: Created
  Change flags: 0x0
  Remote advertisements:
    Route Distinguisher: 10.255.2.1:100
      VNI: 9101
      Router MAC: 00:05:86:e1:03:f0
      BGP nexthop address: 10.255.2.1
    Route Distinguisher: 10.255.2.2:100
      VNI: 9101
      Router MAC: 00:05:86:d0:a6:f0
      BGP nexthop address: 10.255.2.2

Prefix: 100.2.12.0/22, Ethernet tag: 0
  IP route status: Created
  Change flags: 0x0
  Remote advertisements:
    Route Distinguisher: 10.255.2.1:100
      VNI: 9101
      Router MAC: 00:05:86:e1:03:f0
      BGP nexthop address: 10.255.2.1
    Route Distinguisher: 10.255.2.2:100
      VNI: 9101
      Router MAC: 00:05:86:d0:a6:f0
      BGP nexthop address: 10.255.2.2

Prefix: 100.2.16.0/22, Ethernet tag: 0
  IP route status: Created
  Change flags: 0x0
  Remote advertisements:
    Route Distinguisher: 10.255.2.1:100
      VNI: 9101
      Router MAC: 00:05:86:e1:03:f0
      BGP nexthop address: 10.255.2.1
    Route Distinguisher: 10.255.2.2:100
      VNI: 9101
      Router MAC: 00:05:86:d0:a6:f0
      BGP nexthop address: 10.255.2.2
```

```
EVPN->IPv6 Imported Prefixes

Prefix: 1234:100:2::/64, Ethernet tag: 0
  IP route status: Created
  Change flags: 0x0
  Remote advertisements:
    Route Distinguisher: 10.255.2.1:100
      VNI: 9101
      Router MAC: 00:05:86:e1:03:f0
      BGP nexthop address: 10.255.2.1
    Route Distinguisher: 10.255.2.2:100
      VNI: 9101
      Router MAC: 00:05:86:d0:a6:f0
      BGP nexthop address: 10.255.2.2

Prefix: 1234:100:2:4::/64, Ethernet tag: 0
  IP route status: Created
  Change flags: 0x0
  Remote advertisements:
    Route Distinguisher: 10.255.2.1:100
      VNI: 9101
      Router MAC: 00:05:86:e1:03:f0
      BGP nexthop address: 10.255.2.1
    Route Distinguisher: 10.255.2.2:100
      VNI: 9101
      Router MAC: 00:05:86:d0:a6:f0
      BGP nexthop address: 10.255.2.2

Prefix: 1234:100:2:8::/64, Ethernet tag: 0
  IP route status: Created
  Change flags: 0x0
  Remote advertisements:
    Route Distinguisher: 10.255.2.1:100
      VNI: 9101
      Router MAC: 00:05:86:e1:03:f0
      BGP nexthop address: 10.255.2.1
    Route Distinguisher: 10.255.2.2:100
      VNI: 9101
      Router MAC: 00:05:86:d0:a6:f0
      BGP nexthop address: 10.255.2.2

Prefix: 1234:100:2:12::/64, Ethernet tag: 0
```

```
  IP route status: Created
  Change flags: 0x0
  Remote advertisements:
    Route Distinguisher: 10.255.2.1:100
      VNI: 9101
      Router MAC: 00:05:86:e1:03:f0
      BGP nexthop address: 10.255.2.1
    Route Distinguisher: 10.255.2.2:100
      VNI: 9101
      Router MAC: 00:05:86:d0:a6:f0
      BGP nexthop address: 10.255.2.2

Prefix: 1234:100:2:16::/64, Ethernet tag: 0
  IP route status: Created
  Change flags: 0x0
  Remote advertisements:
    Route Distinguisher: 10.255.2.1:100
      VNI: 9101
      Router MAC: 00:05:86:e1:03:f0
      BGP nexthop address: 10.255.2.1
    Route Distinguisher: 10.255.2.2:100
      VNI: 9101
      Router MAC: 00:05:86:d0:a6:f0
      BGP nexthop address: 10.255.2.2
```

Release Information

Command introduced in Junos OS Release 15.1X53-D30.

RELATED DOCUMENTATION

show evpn l3-context | 19790

show route table | 23789

show evpn l3-context

IN THIS SECTION

- Syntax | 19790
- Description | 19790
- Options | 19790
- Additional Information | 19791
- Required Privilege Level | 19791
- Output Fields | 19791
- Sample Output | 19792
- Release Information | 19793

Syntax

```
show evpn l3-context
<brief | extensive>
<logical-system logical-system-name>
<routing-instance-name>
```

Description

Display EVPN database information for Layer 3 virtual routing and forwarding (VRF) instances.

Options

none	Display standard EVPN database information for all Layer 3 VRF instances.
brief | extensive	(Optional) Display the specified level of output.
`logical-system <logical-system-name>`	(Optional) Display Layer 3 information for the specified logical system or all logical systems.
routing-instance-name	(Optional) Display EVPN database information for the specified Layer 3 VRF instance.

Additional Information

Required Privilege Level

view

Output Fields

Table 511 on page 19791 lists the output fields for the `show evpn l3-context` command. Output fields are listed in the approximate order in which they appear.

Table 511: show evpn l3-context Output Fields

Field Name	Field Description	Level of Output
L3 context	Name of VRF instance.	All levels
Type	Status of VRF instance	All levels
Fwd	Type of forwarding route	level-of-output none
Encap	Type of encapsulation.	level-of-output none
VNI	Virtual network identifier for VRF instance.	All levels
Router MAC/GW intf	For a pure type-5 route, the router MAC provides next-hop reachability information and the MAC address of the sending device For a standard type-5 route, the GW interface provides the next-hop route.	All levels
Advertisement mode	Type of forwarding route.	level-of-output Extensive
IPv4 source VTEP address	IPv4 source address for the Virtual Extensible LAN (VXLAN) tunnel.	level-of-output Extensive

Table 511: show evpn l3-context Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IPv6 source VTEP adddress	IPv6 source address for the Virtual Extensible LAN (VXLAN) tunnel.	level-of-output Extensive
IP->EVPN export policy	Name of export routing policy applied to forwarded IP routes.	level-of-output extensive
Flags	Enabled trace flags.	level-of-output extensive
Change flags	Changed trace flags.	level-of-output extensive
Composite nexthop	Status of next-hop route: `Enabled` or `Disabled`.	level-of-output extensive
Route distinguisher	Route distinguisher identifier of the VRF instance	level-of-output extensive
Reference count		level-of-output extensive

Sample Output

show evpn l3-context

```
user@DC1_SPINE1_RE> show evpn l3-context
L3 context                      Type  Fwd   Encap  VNI/Label  Router MAC/GW intf
VRF-100                         Cfg   Sym   VXLAN  9100       00:05:86:28:90:f0
VRF-200                         Cfg   Sym   VXLAN  9200       00:05:86:28:90:f0
```

show evpn l3-context extensive

```
user@DC1_SPINE1_RE> show evpn l3-context extensive
L3 context: VRF-100
Type: Configured
  Advertisement mode: Direct nexthop, Router MAC: 00:05:86:28:90:f0
  Encapsulation: VXLAN, VNI: 9100
  IPv4 source VTEP address: 10.255.1.1
  IPv6 source VTEP address: abcd::128:102:242:145
  IP->EVPN export policy: EVPN-TYPE5-EXPORT-VRF-100
  Flags: 0x19 <Configured IRB-MAC New>
  Change flags: 0xe <Export-Policy Fwd-Mode Encap>
  Composite nexthop support: Enabled
  Route Distinguisher: 10.255.1.1:100
  Reference count: 21

L3 context: VRF-200
  Type: Configured
  Advertisement mode: Direct nexthop, Router MAC: 00:05:86:28:90:f0
  Encapsulation: VXLAN, VNI: 9200
  IPv4 source VTEP address: 10.255.1.1
  IPv6 source VTEP address: abcd::128:102:242:145
  IP->EVPN export policy: EVPN-TYPE5-EXPORT-VRF-200
  Flags: 0x19 <Configured IRB-MAC New>
  Change flags: 0xe <Export-Policy Fwd-Mode Encap>
  Composite nexthop support: Enabled
  Route Distinguisher: 10.255.1.1:200
  Reference count: 21
```

Release Information

Statement introduced in Junos OS Release 15.1X53-D30.

RELATED DOCUMENTATION

show evpn ip-prefix-database | 19779

show evpn mac-ip-table

IN THIS SECTION

- Syntax | 19794
- Description | 19794
- Options | 19794
- Required Privilege Level | 19795
- Output Fields | 19795
- Sample Output | 19797
- Release Information | 19798

Syntax

```
show evpn mac-ip-table
<brief | count | detail | detail | extensive>
<instance>
<kernel >
<logical-system (all | logical-system-name)>
```

Description

Displays the MAC-IP address for all IPv4 (ARP) and IPv6 (ND) bindings for routing instances where the instance-type is evpn.

Options

brief | detail |extensive | summary	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display the MAC-IP address information for the specified routing instance.
kernel	(Optional) Display the MAC-IP address entries in the kernel

`logical-system <logical-system-name>`	(Optional) Display the MAC-IP address information for the specified logical system or all logical systems.

Required Privilege Level

maintenance

Output Fields

Table 512 on page 19795 lists the output fields for the `show evpn mac-ip-table` command. Output fields are listed in the approximate order in which they appear.

Table 512: show evpn mac-ip-table Output Fields

Field Name	Field Description	Level of Output
IP address	IP address associated with the EVPN routing instance.	All levels
MAC address	MAC address	All levels

Table 512: show evpn mac-ip-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	MAC IP flags: Identifies : statically installed MAC-IP entries. • S—Statically installed MAC-IP entries. • D—Dynamic installed MAC-IP address entries. • L—Locally learned MAC-IP address entries • R— MAC-IP address entries that are learnt from a remote device via the control plane. • K—MAC-IP address entries that are installed in the kernel. • RT—Destination Route that corresponds to an installed entry. • AD—Advertised to remote device. • RE—ARP/ND entry has expired and another arp or network solicitation request has been sent. • R0—Router • OV—Override	Brief, detail
Logical Interface	The logical interface associated with the routing instance.	Brief, detail
Active Source	The source of the learned MAC-IP address entry. Displays the ESI or router ID.	All levels

Table 512: show evpn mac-ip-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Routing Instance	Information on the Routing instance where the MAC-IP address entry was learned. • Bridging domain • MAC-IP local mask • MAC-IP remote mask • MAC-IP RPD flags • MAC-IP flags	Extensive
BD ID	Bridge Domain ID	Brief

Sample Output

show evpn mac-ip-table

```
user@host> show evpn mac-ip-table
MAC IP flags  (S - Static, D - Dynamic, L - Local , R - Remote, K - Kernel, RT - Dest Route,
               AD - Advt to remote, RE - Re-ARP/ND, RO - Router, OV - Override)
 Routing instance : evpn
 Bridging domain : bd1
   IP                          MAC                 Flags            Logical            Active
   address                     address                              Interface          source
   10.1.1.1                    00:11:00:00:00:01   DR,K
81.1.1.1
   10.1.1.2                    00:22:00:00:00:01   DR,K
81.2.2.2
   10.1.1.3                    00:33:00:00:00:01   DL,K,AD          ge-0/0/2.1
   10.1.1.4                    00:44:00:00:00:01   DR,K
81.4.4.4
```

show evpn mac-ip-table extensive

```
user@host> show evpn mac-ip-table extensive
IP address: 10.1.1.1
  MAC address: 00:11:00:00:00:01
  Routing instance: evpn
   Bridging domain: bd1
   MAC-IP local mask: 0x0000000000000000
   MAC-IP remote mask: 0x8000000000000000
   MAC-IP RPD flags: 0x00000081
   MAC-IP flags: remote,kernel
```

show evpn mac-ip-table kernel

```
user@host> show evpn mac-ip-table kernel
    BD       IP                          MAC                 Logical          Flags
    id       address                     address             Interface
    2        192.168.1.1                 00:21:59:aa:77:f0   irb.1            0x00000209
    2        192.168.1.100               00:00:5e:00:01:01   irb.1            0x00000609
    3        192.168.2.1                 00:21:59:aa:77:f0   irb.2            0x00000209
    3        192.168.2.100               00:00:5e:00:01:01   irb.2            0x00000609
    4        192.168.3.1                 00:21:59:aa:77:f0   irb.3            0x00000209
    4        192.168.3.100               00:00:5e:00:01:01   irb.3            0x00000609
    5        192.168.4.1                 00:21:59:aa:77:f0   irb.4            0x00000209
    5        192.168.4.100               00:00:5e:00:01:01   irb.4            0x00000609
```

Release Information

Command introduced in Junos OS Release 17.4R2.

show evpn mac-table

IN THIS SECTION

Syntax | 19799

- Description | 19799
- Options | 19799
- Required Privilege Level | 19800
- Release Information | 19800

Syntax

```
show evpn mac-table
<age>
<address>
<brief |count | detail | extensive | summary>
<instance instance-name>
<interface interface-name>
<isid <isid>>
<logical-system <logical-system-name>>
<vlan-id vlan-id>
```

Description

Show Ethernet VPN (EVPN) MAC table information.

Options

none	Display brief information about the EVPN MAC table.
age	(Optional) Display age of a single mac-address.
address	(Optional) Display MAC table information for the specified MAC address.
brief | count | detail |extensive | summary	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display MAC table information for a specific routing instance.
isid *<isid>*	(Optional) Display MAC table information for the specified ISID or all ISIDs.

`logical-system <logical-system-name>`	(Optional) Display MAC table information for the specified logical system or all logical systems.
`vlan-id vlan-id`	(Optional) Display MAC table information for the specified VLAN.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Example: Configuring an EVPN with IRB Solution on EX9200 Switches

show evpn mld-snooping database

IN THIS SECTION

- Syntax | 19801
- Syntax for mac-vrf routing | 19801
- Description | 19801
- Options | 19801
- Required Privilege Level | 19802
- Output Fields | 19802
- Sample Output | 19803
- Release Information | 19804

Syntax

```
show evpn mld-snooping database
<extensive>
<esi id>
<group group-address>
<instance id>
<interface name>
<l2-domain-id id>
```

Syntax for mac-vrf routing

```
show mac-vrf routing mld-snooping database
<extensive>
<esi id>
<group group-address>
<instance id>
<interface name>
<l2-domain-id id>
```

Description

Display MLD snooping information for IPv6 multicast forwarding in an EVPN network.

Options

none	Display detailed information.
extensive	(Optional) Display more extensive output.
esi	(Optional) Display output only for a specified EVPN Ethernet segment ID (ESI).
group *group-address*	(Optional) Display output only for the specified IPv6 multicast group address range.
instance *instance-name*	(Optional) Display information for the specified EVPN instance (EVI).
interface *name*	(Optional) Display information about the specified interface.

l2-domain-id *id* (Optional) Display output for the specified bridging domain, VLAN, or virtual network identifier (VNI).

Required Privilege Level

view

Output Fields

Table 513 on page 19802 lists the output fields for the `show evpn mld-snooping database` command. Output fields are listed in the approximate order in which they appear.

Table 513: show evpn mld-snooping database Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the EVPN routing instance for which information is displayed.	All levels
VN Identifier	Tunnel network identifier.	All levels
Group IP	Multicast IP address of the multicast group.	All levels
Source IP	IP address of the multicast data source for the group.	All levels
Access OIF Count	Number of outgoing interface (OIF) list entries on the access side.	All levels
Access OIF Count (more information)	More OIF Entry Information: • `Interface`—Interface name • `ESI`—Associated EVPN Ethernet Segment ID for the interface • `Local`—Local receiver. • `Remote`—Remote receiver.	extensive

Sample Output

show evpn mld-snooping database

```
user@device> show evpn mld-snooping database
Instance: EVPN-2
    VN Identifier: 100
        Group IP: ff1e::25:25:1, Source IP: ::, Access OIF Count: 1
        Group IP: ff1e::33:33:1, Source IP: ::, Access OIF Count: 1

Instance: EVPN-4
    VN Identifier: 200
        Group IP: ff1e::25:25:1, Source IP: ::, Access OIF Count: 1
        Group IP: ff1e::33:33:1, Source IP: ::, Access OIF Count: 1
```

show evpn mld-snooping database instance instance-name

```
user@device> show evpn mld-snooping database
EVPN-4
Instance: EVPN-4
    VN Identifier: 200
        Group IP: ff1e::25:25:1, Source IP: ::, Access OIF Count: 1
        Group IP: ff1e::33:33:1, Source IP: ::, Access OIF Count: 1
```

show evpn mld-snooping database extensive

```
user@device> show evpn mld-snooping database
extensive
Instance: EVPN-2
  VN Identifier: 100
        Group IP: ff1e::25:25:1, Source IP: ::
            Access OIF Count: 1
                 Interface                  ESI             Local    Remote
                ae3.0         00:22:22:22:22:22:22:22:22:22 0           1
        Group IP: ff1e::33:33:1, Source IP: ::
            Access OIF Count: 1
```

```
                 Interface                  ESI             Local    Remote
                ae3.0         00:22:22:22:22:22:22:22:22:22 1           0

Instance: EVPN-4
  VN Identifier: 200
        Group IP: ff1e::25:25:1, Source IP: ::
            Access OIF Count: 1
                 Interface                  ESI             Local    Remote
                ae4.0         00:44:44:44:44:44:44:44:44:44 0           1
        Group IP: ff1e::33:33:1, Source IP: ::
            Access OIF Count: 1
                 Interface                  ESI             Local    Remote
                ae4.0         00:44:44:44:44:44:44:44:44:44 1           0
```

Release Information

Command introduced in Junos OS Release 18.4R1.

Introduced in Junos OS Release 19.4R2 on ACX Series routers.

RELATED DOCUMENTATION

Configure Multicast Forwarding with MLD Snooping in an EVPN-MPLS Environment

Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment

Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment

show evpn multicast-snooping assisted-replication multihomed-peers

IN THIS SECTION

- Syntax | **19805**
- Description | **19805**
- Options | **19805**
- Required Privilege Level | **19806**

- Output Fields | 19806
- Sample Output | 19806
- Release Information | 19807

Syntax

```
show evpn multicast-snooping assisted-replication multihomed-peers
<extensive>
<instance evpn-instance-name>
```

Description

Display multihoming peer neighbor addresses and corresponding multihomed Ethernet segment information for an assisted replication (AR) leaf device.

When devices with AR enabled operate in extended AR mode (see *Extended AR Mode for Multihomed Ethernet Segments*), an ingress AR leaf device uses the information this command shows to replicate and forward multicast traffic to all other AR leaf devices that are its multihoming peers (have attached access devices with the same multihomed Ethernet segment identifier [ESI]). The ingress AR leaf device also sends the traffic to an AR replicator device that replicates it to other devices in the EVPN network that are not multihoming peers with the AR leaf device.

> **NOTE**: QFX Series devices must use extended AR mode to support AR in an EVPN network with a Virtual Extensible LAN (VXLAN) overlay and multihomed Ethernet segments.

Output for this command is only available when IGMP snooping is enabled.

Options

extensive	Display more details about multihoming peers.
instance *evpn-instance-name*	(Optional) Display information only for the specified EVPN instance.

Required Privilege Level

view

Output Fields

Table 514 on page 19806 describes the output fields for the `show evpn multicast-snooping assisted-replication multihomed-peers` command. Output fields are listed in the approximate order in which they appear.

Table 514: show evpn multicast-snooping assisted-replication multihomed-peers Output Fields

Field Name	Field Description
Instance	EVPN instance name.
Neighbor Address	Multihoming peer IP address.
Nexthop Index	Ingress replication next hop for the multihoming peer AR leaf device in the `Neighbor Address` output field.
Interface	Interface name to reach the multihoming peer device in the `Neighbor Address` output field.
Local Multi-homed ESIs	(Displayed with `extensive` option only) Multihomed Ethernet segments shared with the multihoming peer in the `Neighbor Address` output field.

Sample Output

show evpn multicast-snooping assisted-replication multihomed-peers

```
user@device> show evpn multicast-snooping assisted-replication
multihomed-peers
Instance: default-switch
  Neighbor Address   Nexthop Index    Interface
  10.0.1.2        1841             vtep.32782
```

show evpn multicast-snooping assisted-replication multihomed-peers (extensive)

```
user@device> show evpn multicast-snooping assisted-replication
multihomed-peers extensive
Instance: default-switch

  Neighbor address: 10.0.1.2
    Nexthop Index: 1841
    Interface: vtep.32782
      Local Multi-homed ESIs
      00:11:22:33:44:55:66:77:88:99
      00:11:22:33:44:55:66:77:88:aa
```

Release Information

Command introduced in Junos OS Release 18.4R2 and 19.4R1.

RELATED DOCUMENTATION

Assisted Replication Multicast Optimization in EVPN Networks

Extended AR Mode for Multihomed Ethernet Segments

show evpn multicast-snooping assisted-replication next-hops

IN THIS SECTION

- Syntax | 19808
- Description | 19808
- Options | 19808
- Required Privilege Level | 19809
- Output Fields | 19809
- Sample Output | 19810
- Release Information | 19810

Syntax

```
show evpn multicast-snooping assisted-replication next-hops
<index next-hop-ID>
<instance evpn-instance-name>
<l2-domain ID>
```

Description

Display next-hop information on an assisted replication (AR) leaf device that the AR leaf uses for load-balancing among available AR replicators.

- AR leaf devices in the QFX5000 line of switches perform load-balancing among the available AR replicators by designating a particular AR replicator device for a VLAN or VXLAN network identifier (VNI).
- AR leaf devices in the QFX10000 line of switches load-balance among the available AR replicators actively based on traffic flow levels within a VNI.

An AR device will have an AR VXLAN encapsulation next hop and an ingress replication VXLAN encapsulation next hop towards a particular AR replicator. Load-balancing next hops are AR VXLAN encapsulation next hops towards the AR replicator.

Output for this command is only available when IGMP snooping is enabled.

Options

index *nh-index* (Optional) Display information only for the specified next-hop.

instance *evpn-instance-name* (Optional) Display information only for the specified EVPN instance.

l2-domain-id *ID* (Optional) Display information only for the specified Layer 2 domain ID, which can be a VLAN ID, a VXLAN network identifier (VNI), or a service identifier (I-SID) within the range from 1 through 16777214.

> **NOTE**: For QFX Series switches in an EVPN-VXLAN environment, you must specify the Layer 2 domain ID as a VNI.

Required Privilege Level

view

Output Fields

Table 515 on page 19809 describes the output fields for the `show evpn multicast-snooping assisted-replication next-hops` command. Output fields are listed in the approximate order in which they appear.

Table 515: show evpn multicast-snooping assisted-replication next-hops Output Fields

Field Name	Field Description
Instance	EVPN instance name.
VN Identifier	Lists AR load-balancing next-hop information about this VNI.
Load Balance Nexthop Index	Index of the load-balancing next hop (top level).
Load balance to	Lists next-hop information for available AR replicators.
Nexthop Index	AR next hop for the AR replicator in the `AR IP` output field.
Interface	AR tunnel interface name to this AR replicator.
AR IP	IP address of this AR replicator (the configured secondary loopback interface address for AR).
(Designated Node)	(AR leaf devices that are EX4650 switches or switches in the QFX5000 line only) Designated AR replicator for a VLAN or VNI for load-balancing AR operation among the available AR replicators. **NOTE**: AR leaf devices in the QFX10000 line of switches load-balance among the available AR replicators based on traffic flow levels within a VNI, and don't show a designated replicator in the output of this command.

Sample Output

show evpn multicast-snooping assisted-replication next-hops

```
user@device> show evpn multicast-snooping assisted-replication
next-hops
Instance: default-switch

  VN Identifier: 100
    Load Balance Nexthop Index: 131091
      Load balance to:
      Nexthop Index    Interface        AR IP
      1797             vtep.32772       192.168.6.6
      1772             vtep.32770       192.168.5.5 (Designated Node)

  VN Identifier: 101
    Load Balance Nexthop Index: 131092
      Load balance to:
      Nexthop Index    Interface        AR IP
      1797             vtep.32772       192.168.6.6 (Designated Node)
      1772             vtep.32770       192.168.5.5
```

Release Information

Command introduced in Junos OS Release 18.4R2 and 19.4R1.

RELATED DOCUMENTATION

Assisted Replication Multicast Optimization in EVPN Networks

show evpn multicast-snooping assisted-replication replicators

IN THIS SECTION

- Syntax | 19811
- Description | 19811
- Options | 19811
- Required Privilege Level | 19812
- Output Fields | 19812
- Sample Output | 19814
- Release Information | 19815

Syntax

```
show evpn multicast-snooping assisted-replication replicators
<instance evpn-instance-name>
<l2-domain-id ID>
```

Description

Displays information about the assisted replication (AR) replicator devices advertised in an EVPN network and the mode in which they operate.

Extended AR mode enables AR to work with AR replicator devices that can't retain the ingress AR leaf device IP address as the multicast source when replicating and forwarding the traffic into the EVPN core, usually due to restrictions with the overlay tunnel type. QFX series AR devices require this mode to support AR in an EVPN network with a Virtual Extensible LAN (VXLAN) overlay that might have multihomed Ethernet segments.

Output for this command is only available when IGMP snooping is enabled.

Options

instance *evpn-instance-name*	(Optional) Display information only for the specified EVPN instance.

l2-domain-id *ID*

(Optional) Display information only for the specified Layer 2 domain ID, which can be a VLAN ID, a VXLAN network identifier (VNI), or a service identifier (I-SID) within the range from 1 through 16777214.

> **NOTE**: For QFX Series switches in an EVPN-VXLAN environment, you must specify the Layer 2 domain ID as a VNI.

Required Privilege Level

view

Output Fields

Table 516 on page 19812 describes the output fields for the `show evpn multicast-snooping assisted-replication replicators` command. Output fields are listed in the approximate order in which they appear.

Table 516: show evpn multicast-snooping assisted-replication replicators Output Fields

Field Name	Field Description
Instance	EVPN instance name.
AR Role	AR role of the device where you run this command: • AR Leaf • AR Replicator
VN Identifier	Lists information about AR replicators available to this VNI.

Table 516: show evpn multicast-snooping assisted-replication replicators Output Fields *(Continued)*

Field Name	Field Description
Operational Mode	AR mode in which this AR device is operating based on conditions in the network for AR that the device detects: • `Misconfiguration/Ingress Replication`—The AR device detected a misconfiguration in which AR won't operate correctly, such as when an AR replicator and AR leaf share a multihomed Ethernet segment. The AR leaf device will not use AR, and defaults to using regular ingress replication. • `No replicators/Ingress Replication`—The AR leaf device has not learned about any AR replicator devices, and defaults to using regular ingress replication instead of AR. • `Extended AR`—Ingress AR leaf devices should replicate multicast traffic to all other AR leaf devices that are its multihoming peers (have attached access devices in the same multihomed Ethernet segment). This AR replicator device will not duplicate traffic to those multihoming peers. • `Regular AR`—When ingress AR leaf devices send multicast traffic to this AR replicator device, they will not send the traffic to their multihoming peers. This AR replicator device will do the replication and forwarding to the other devices in the EVPN network, including the ingress AR leaf device's multihoming peers.
Replicator IP	IP address of an available AR replicator (the configured loopback address for AR).
Nexthop Index	Next hops index for the AR replicator in the `Replicator IP` output field.
Interface	Tunnel interface name for the AR replicator in the `Replicator IP` output field. **NOTE**: An AR device will have an AR VXLAN encapsulation next hop and an ingress replication VXLAN encapsulation next hop towards a particular AR replicator. This field displays the AR VXLAN encapsulation next hop towards the AR replicator.

Table 516: show evpn multicast-snooping assisted-replication replicators Output Fields *(Continued)*

Field Name	Field Description
Mode	AR mode advertised by the AR replicator in the `Replicator IP` output field: • `Extended AR`—Ingress AR leaf devices should replicate multicast traffic to all other AR leaf devices that are its multihoming peers (have attached access devices in the same multihomed Ethernet segment). This AR replicator device will not duplicate traffic to those multihoming peers. • `Regular AR`—When ingress AR leaf devices send multicast traffic to this AR replicator device, they will not send the traffic to their multihoming peers. This AR replicator device will do the replication and forwarding to the other devices in the EVPN network, including the ingress AR leaf device's multihoming peers.

Sample Output

show evpn multicast-snooping assisted-replication replicators (on AR replicator device)

```
user@device> show evpn multicast-snooping assisted-replication
replicators
Instance: default-switch
AR Role: AR Replicator

  VN Identifier: 100
  Operational Mode: Extended AR
    Replicator IP  Nexthop Index    Interface        Mode
    192.168.6.6    0                Self             Extended AR
    192.168.5.5    1875             vtep.32780       Extended AR

  VN Identifier: 101
  Operational Mode: Extended AR
    Replicator IP  Nexthop Index    Interface        Mode
    192.168.6.6    0                Self             Extended AR
    192.168.5.5    1875             vtep.32780       Extended AR
```

show evpn multicast-snooping assisted-replication replicators (on AR leaf device)

```
user@device> show evpn multicast-snooping assisted-replication
replicators
Instance: default-switch
AR Role: AR Leaf

  VN Identifier: 100
  Operational Mode: Extended AR
    Replicator IP  Nexthop Index    Interface        Mode
    192.168.5.5    1772             vtep.32770       Extended AR
    192.168.6.6    1797             vtep.32772       Extended AR

  VN Identifier: 101
  Operational Mode: Extended AR
    Replicator IP  Nexthop Index    Interface        Mode
    192.168.5.5    1772             vtep.32770       Extended AR
    192.168.6.6    1797             vtep.32772       Extended AR
```

Release Information

Command introduced in Junos OS Release 18.4R2 and 19.4R1.

RELATED DOCUMENTATION

Assisted Replication Multicast Optimization in EVPN Networks

Extended AR Mode for Multihomed Ethernet Segments

show evpn multicast-snooping assisted-replication next-hops | **19807**

show evpn multicast-snooping status

IN THIS SECTION

- Syntax | **19816**
- Description | **19816**

- Options | 19816
- Required Privilege Level | 19816
- Output Fields | 19817
- Sample Output | 19817
- Release Information | 19818

Syntax

```
show evpn multicast-snooping status
<extensive>
```

Description

Display IGMP snooping (for IPv4 multicast traffic) status or MLD snooping (for IPv6 multicast traffic) status in the `default-switch` instance on a device in an EVPN fabric.

You can enable IGMP snooping (or MLD snooping) in an EVPN network to help optimize multicast traffic flow in the fabric. The devices in the fabric use the following EVPN route types to synchronize the multicast state between multihoming peer provider edge (PE) devices:

- EVPN Type 7 Join Sync routes
- EVPN Type 8 Leave Sync routes

Options

extensive	Display more detailed output.

Required Privilege Level

view

Output Fields

Table 517: show evpn multicast-snooping status Output Fields

Field Name	Field Description
Instance	Layer 2 routing instance. The instance corresponds to the `default-switch` instance on QFX Series devices in EVPN-VXLAN fabrics.
VLAN ID	VLAN ID.
Multicast Address Family	Show enabled or disabled multicast snooping status for: • INET—IGMP snooping (IPv4 multicast traffic). • INET6—MLD snooping (IPv6 multicast traffic).
SG Sync	EVPN multicast snooping status for the indicated multicast traffic family—`Enabled` or `Disabled`. Extensive output shows the following additional statistics with multicast snooping status `Enabled`: • Number of local join-sync SGs • Number of remote join-sync SGs • Number of local proxy SGs • Number of remote proxy SGs

Sample Output

show evpn multicast-snooping status

```
user@host> show evpn multicast-snooping status
Instance: __default_evpn__

Instance: default-switch
  VLAN ID: 1
    Multicast Address Family: INET
```

```
      SG Sync: Enabled
    Multicast Address Family: INET6
      SG Sync: Disabled
```

show evpn multicast-snooping status—With IGMP Snooping and MLD Snooping Enabled)

```
user@host> show evpn multicast-snooping status
Instance: __default_evpn__

Instance: default-switch
  VLAN ID: 200
    Multicast Address Family: INET
      SG Sync: Enabled
    Multicast Address Family: INET6
      SG Sync: Enabled
  VLAN ID: 201
    Multicast Address Family: INET
      SG Sync: Enabled
    Multicast Address Family: INET6
      SG Sync: Enabled
```

Release Information

RELATED DOCUMENTATION

Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment

Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment

show evpn nd-table

IN THIS SECTION

Syntax | 19819

- Description | 19819
- Options | 19819
- Required Privilege Level | 19819
- Output Fields | 19820
- Sample Output | 19820
- Release Information | 19822

Syntax

```
show evpn nd-table
<address>
<brief |count | detail | extensive>
<instance instance-name>
```

Description

Show Ethernet VPN (EVPN) Network Discovery Protocol (NDP) entries associated with learned MAC addresses.

Routing platforms support the `show bridge evpn nd-table` command, which displays the same information and includes the same options.

Options

none	Display brief information about the EVPN NDP table.
address	(Optional) Display NDP information for the specified MAC address.
brief | count | detail |extensive	(Optional) Display the specified level of output.
instance <*instance-name*>	(Optional) Display NDP information for the specified routing instance .

Required Privilege Level

view

Output Fields

Table 518 on page 19820 lists the output fields for the `show evpn nd-table` command. Output fields are listed in the approximate order in which they appear.

Table 518: show evpn nd-table Output Fields

Field Name	Field Description	Level of Output
INET address	The INET address related to the INET entries that are added to the NDP table.	All levels
MAC address	MAC addresses learned through NDP.	**brief**, **detail**, **extensive**, **instance**, **mac-address**,,
Logical Interface	Logical interface associated with the routing instance in which the NDP INET address is learned.	**brief**, **instance**, **mac-address**,,
Routing instance	Routing instance in which the NDP INET address is learned.	all levels
Bridging domain	Bridging domain in which the NDP INET address is learned.	all levels
Learning interface	Interface on which the NDP INET address is learned.	**detail**, **extensive**
Count	Indicates the number of NDP INET addresses learned in a routing instance in a bridge domain.	**count**

Sample Output

show evpn nd-table

```
user@host> show evpn nd-table
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
2001:db8::2         00:05:86:a0:dc:f0  irb.0         evpn_1       __evpn_1__
```

show evpn nd-table 00:05:86:a0:dc:f0 (MAC address)

```
user@host> show evpn nd-table 00:05:86:a0:dc:f0
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
2001:db8::2         00:05:86:a0:dc:f0  irb.0         evpn1        __evpn1__
```

show evpn nd-table brief

```
user@host> show evpn nd-table brief
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
2001:db8::2         00:05:86:a0:dc:f0  irb.0         evpn1        __evpn1__
```

show evpn nd-table detail

```
user@host> show evpn nd-table detail

INET address: 2001:db8::2
MAC address: 00:05:86:a0:dc:f0
  Routing instance: evpn1
   Bridging domain: __evpn1__
   Learning interface: irb.0
```

show evpn nd-table count

```
user@switch> show evpn nd-table count
1 NDP INET addresses learned in routing instance evpn1 bridge domain __evpn1__
```

show evpn nd-table extensive

```
user@host> show evpn nd-table extensive

INET address: 2001:db8::2
MAC address: 00:05:86:a0:dc:f0
  Routing instance: evpn1
```

```
    Bridging domain: __evpn1__
    Learning interface: irb.0
```

show evpn ndp-table instance evpn1

```
user@host> show evpn arp-table instance evpn1
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
2001:db8::2         00:05:86:a0:dc:f0  irb.0         evpn1        __evpn1__
```

Release Information

Command introduced in Junos OS Release 16.2.

show evpn oism

IN THIS SECTION

- Syntax | 19822
- Description | 19823
- Options | 19823
- Required Privilege Level | 19823
- Output Fields | 19823
- Sample Output | 19824
- Release Information | 19825

Syntax

```
show evpn oism
<extensive>
<l3-context l3-vrf-name>
```

Description

Display information about optimized intersubnet multicast (OISM) elements configured on the devices in an EVPN-VXLAN fabric.

Options

`extensive`	(Optional) Display more detailed output. Without this option, the command displays brief output by default.
`l3-context` *`l3-vrf-name`*	(Optional) Display information only for the specified Layer 3 virtual routing and forwarding (VRF) routing instance. Without this option, this command displays information for all Layer 3 VRF routing instances where you enabled OISM.

Required Privilege Level

view

Output Fields

Table 519: show evpn oism Output Fields

Field Name and Output Level	Field Description
L3 context (default) or EVPN L3 context (extensive)	VRF routing instance name.
SBD (default) or OISM SBD interface (extensive)	IRB interface of the supplemental bridge domain (SBD) that you configure in the corresponding VRF routing instance when you enable OISM. If you don't enable OISM in the VRF routing instance, this command displays `(null)` in this field.

Table 519: show evpn oism Output Fields *(Continued)*

Field Name and Output Level	Field Description
PEG-DF-ELECTION (default) or PEG DF Election (extensive)	PIM EVPN gateway (PEG) device designated forwarder (DF) election method enabled for the L3 context. By default, peer OISM border leaf devices configured as PEG devices use native PIM DF election. You can enable OISM PEG DF election instead to customize the DF election method the PEG devices use. To do this, configure the `peg-df-election` statement at the `[edit <routing-instances name> protocols evpn oism pim-evpn-gateway]` hierarchy level.
PEG DF Algorithm (extensive)	PEG DF election method configured for the L3 context: • MOD—Mod-based DF election • PREFERENCE—Preference-based DF election
OISM Mode (extensive)	OISM operating mode: • Regular OISM—Original symmetric bridge domains (VLANs) version of OISM that requires you to configure all revenue VLANS symmetrically on all OISM leaf devices in the fabric. • Enhanced OISM—Asymmetric bridge domains (VLANs) enhanced version of OISM, where on each leaf device you can save resources by configuring only the revenue VLANs hosted by that device.

Sample Output

show evpn oism

```
user@host> show evpn oism
L3 context                      SBD        PEG-DF-ELECTION
DELTA-VRF                       irb.901    Enabled
GAMMA-VRF                       irb.902    Enabled
```

show evpn oism l3-context extensive

```
user@host> show evpn oism l3-context DELTA extensive
EVPN L3 context: DELTA
  OISM SBD interface: irb.901
  PEG DF Election: Enabled
  PEG DF Algorithm: MOD
  OISM Mode: Regular OISM
```

show evpn oism extensive (Enhanced OISM)

```
regress@SL1> show evpn oism extensive
EVPN L3 context: DELTA
  OISM SBD interface: irb.901
  OISM Mode: Enhanced OISM

EVPN L3 context: GAMMA
  OISM SBD interface: irb.902
  OISM Mode: Enhanced OISM
```

Release Information

Command introduced in Junos OS Release 21.2R1.

PEG DF election output information added in Junos OS and Junos OS Evolved Release 23.4R1.

OISM Mode output field added in Junos OS and Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Optimized Inter-Subnet Multicast in EVPN Networks

show evpn oism peg-df-status

IN THIS SECTION

- Syntax | 19826
- Description | 19826
- Options | 19826
- Required Privilege Level | 19827
- Output Fields | 19827
- Sample Output | 19828
- Release Information | 19829

Syntax

```
show evpn oism peg-df-status
<extensive>
<l2-domain-id domain-id>
<l3-context l3-vrf-name>
<logical-system (name | all)>
```

Description

Display information about the designated forwarder (DF) status of PIM Ethernet VPN (EVPN) gateway (PEG) devices in an EVPN-VXLAN network running optimized intersubnet multicast (OISM).

Options

none (default output)	Display summary output for all Layer 3 (L3) virtual routing and forwarding (VRF) instances and Layer 2 (L2) domains.
extensive	Display more extensive output.
l2-domain-id *domain-id*	Display output for the specified L2 domain identifier, which in this case is a VXLAN virtual network identifier (VNI).

l3-context *l3-vrf-name* Display output for the specified L3 VRF instance.

logical-system (*name* **| all)** Display output for the specified logical system or all logical systems.

Required Privilege Level

admin

Output Fields

Table 520: show evpn oism peg-df-status Output Fields

Field Name	Field Description	Level of Output
EVPN L3 context	L3 VRF instance name.	default
VN Identifier	VNI (L2 domain identifier).	default
IPv4 Multicast DF Status	DF status for the VNI: • DF: Elected designated forwarder. • nDF: Not elected the designated forwarder.	default
DF Address	IP address of the elected DF for the VNI.	default
DF Candidates List	DF candidates for DF election for the VNI. The list includes the ordinal number in the candidates list and the IP address of each DF candidate.	extensive
PIM Update TS	Timestamp of most recent PIM DF update for the VNI . **NOTE**: The device communicates elected PEG DF information to the PIM protocol, which builds the multicast routes, even when the device is using PEG DF election rather than native PIM DF election.	extensive
DF Compute TS	Timestamp of most recent PEG DF election computation for the VNI.	extensive

Table 520: show evpn oism peg-df-status Output Fields *(Continued)*

Field Name	Field Description	Level of Output
History db	History listing of event information related to the DF candidates and the DF election for the VNI. You can set the number of DF election history database entries you want the device to maintain per VNI using the `peg-df-election-history` statement at the `[edit protocols evpnevpn]` hierarchy level.	extensive

Sample Output

show evpn oism peg-df-status

```
user@host> show evpn oism peg-df-status
EVPN L3 context: VRF-1

    VN Identifier: 1001

    VN Identifier: 2
    IPv4 Multicast DF Status: DF, DF Address: 10.1.1.1

    VN Identifier: 3
    IPv4 Multicast DF Status: NDF, DF Address: 10.2.2.1

    VN Identifier: 4
    IPv4 Multicast DF Status: DF, DF Address: 10.1.1.1
```

show evpn oism peg-df-status extensive

```
user@host> show evpn oism peg-df-status extensive
EVPN L3 context: VRF-1

    VN Identifier: 1001
    History db: <No entries>

    VN Identifier: 2
        IPv4 Multicast DF Status: DF, DF Address: 10.1.1.1
```

```
            DF Candidates List[Total: 2]
             0: 1.1.1.1
             1: 10.2.2.1
            PIM Update TS: Aug 14 11:25:17, DF Compute TS: Aug 14 11:25:17
    History db:
      Time                        Event
      Oct 14 11:25:17.358 2023   Added DF candidate 10.1.1.1 df type: 0, pref 32767 dp: 0 gran:
0 ord_num 0 total 1
      Oct 14 11:25:17.358 2023   Added DF candidate 10.2.2.1 df type: 0, pref 0 dp: 0 gran: 0
ord_num 1 total 2
      Oct 14 11:25:17.359 2023   Completed DF election, 2 total candidate(s),
max_candidate_count 2, mod 0, l2_domain_ID 2, new DF PE address 1.1.1.1
```

Release Information

Command introduced in Junos OS 23.4R1 and Junos OS Evolved 23.4R1.

RELATED DOCUMENTATION

Optimized Inter-Subnet Multicast in EVPN Networks

PEG DF Election

peg-df-election | **12395**

show evpn oism spmsi-ad

IN THIS SECTION

- Syntax | **19830**
- Description | **19830**
- Options | **19831**
- Required Privilege Level | **19831**
- Output Fields | **19831**
- Sample Output | **19832**
- Release Information | **19832**

Syntax

```
show evpn oism spmsi-ad
<extensive>
<group multicast-group-address>
<l3-context l3-vrf-name>
<source multicast-source-address>
```

Description

Display multicast source and group (S,G) information corresponding to Selective P-router Multicast Service Interface (S-PMSI) Auto-Discovery (A-D) routes in an EVPN-VXLAN network running enhanced optimized intersubnet multicast (OISM). Enhanced OISM devices use S-PMSI A-D routes to perform PIM source registration only for multicast sources inside the EVPN-VXLAN network.

See *PIM Registration with Enhanced OISM for Internal Sources Based on EVPN Type 10 S-PMSI A-D Routes* for details on how OISM leaf devices use these routes.

With enhanced OISM, also called the asymmetric bridge domains design, the PEG devices receive traffic on the OISM supplemental bridge domain (SBD) from both external and internal multicast sources. Because OISM PEG devices should only perform PIM registration to the PIM rendezvous point (RP) for internal sources, the enhanced OISM design must distinguish between internal and external multicast sources. With enhanced OISM, the ingress leaf devices advertise S-PMSI A-D routes for internal multicast sources. The PEG devices then send a PIM register to the PIM RP only for the sources that correspond to received S-PMSI A-D routes.

NOTE: You can see more details about received EVPN Type 10 S-PMSI A-D route advertisements in the EVPN instance multicast routing table using `show route table` commands such as the following:

```
user@host> show route table default-switch.evpn.0 match-prefix 10*
```

Or:

```
user@host> show route table evpn-instance-name.evpn-mcsn.1 match 10* extensive
```

Options

Basic brief output (no `extensive` option)	Display summary output for all Layer 3 (L3) virtual routing and forwarding (VRF) instances and all EVPN instances.
`extensive`	(Optional) Display more detailed output.
`group` *`multicast-group-address`*	(Optional) Display information for the specified multicast group.
`l3-context` *`l3-vrf-name`*	(Optional) Display information only for the specified Layer 3 virtual routing and forwarding (VRF) routing instance. Without this option, this command displays information for all Layer 3 VRF routing instances where you enabled OISM.
`source` *`multicast-source-address`*	(Optional) Display information for the specified multicast source address.

Required Privilege Level

view

Output Fields

Table 521: show evpn oism spmsi-ad Output Fields

Field Name	Field Description
Instance	EVPN Instance.
VN Identifier	VXLAN virtual network identifier (VNI), which is the L2 domain identifier.
Group IP (basic) or Group (extensive)	Multicast group address.
Source IP (basic) or Source (extensive)	Multicast source address.
Local (extensive only)	Set to 1 if the device *originated* an EVPN Type 10 S-PMSI A-D route.

Table 521: show evpn oism spmsi-ad Output Fields *(Continued)*

Field Name	Field Description
Remote (extensive only)	Set to 1 if the device *received* an EVPN Type 10 S-PMSI A-D route.

Sample Output

show evpn oism spmsi-ad (Brief output with other options)

```
user@host> show evpn oism spmsi-ad group 233.252.0.4 source 10.0.1.3 l3-context DELTA
Instance: DELTA
  VN Identifier: 901
   Group IP: 233.252.0.4, Source IP: 10.0.1.1
```

show evpn oism spmsi-ad extensive

```
user@host> show evpn oism spmsi-ad extensive
Instance evpn-vxlan-A
  VN Identifier: 902
   Group        Source    Local   Remote
233.252.0.4     10.0.1.2   1       0

Instance evpn-vxlan-B
  VN Identifier: 903
   Group        Source    Local   Remote
233.252.0.4     10.0.1.2   1       0
```

Release Information

Command introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

show evpn oism | 19822

show evpn p2mp

IN THIS SECTION

- Syntax | 19833
- Description | 19833
- Options | 19833
- Required Privilege Level | 19834
- Output Fields | 19834
- Sample Output | 19835
- Release Information | 19836

Syntax

```
show evpn p2mp
<brief | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Displays EVPN P2MP information.

Options

brief | detail | extensive	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display routing instance information for the specified instance only.

logical-system (all | *logical-system-name*)	(Optional) Display the EVPN P2MP information of all logical systems or a particular logical system.

Required Privilege Level

View

Output Fields

Table 522 on page 19834 lists the output fields for the show evpn instance command. Output fields are listed in the approximate order in which they appear.

Table 522: show evpn p2mp Output Fields

Field Name	Field Description
Instance	Names of the routing instances.
VLAN	VLAN identifier. The following information is displayed for each VLAN when the `extensive` option is used. • Multicast Nexthop—The Inclusive Multicast (IM) nexthop that is used to send BUM traffic for a VLAN. • P2MP Tunnel—The name of the P2MP tunnel. • PMSI—The provider multicast service interface attributes as advertised in EVPN IM route. This is described in RFC 7432. • ESI—Ethernet segment identifier (ESI) value of the interfaces. For each ESI, the SH nexhop identifier that is used.

Table 522: show evpn p2mp Output Fields *(Continued)*

Field Name	Field Description
Neighbor Address	IP address of the EVPN neighbor. The following information is displayed for each neighbor when the `extensive` option is used. • Transport Labels—The LDP/RSVP transport label for the P2MP LSP rooted at this neighbor • PMSI—The provider multicast service interface attributes as advertised in EVPN IM route. This is described in RFC 7432.
Multicast Nexthop	Multicast nexthop identifier. This is the Inclusive Multicast (IM) nexthop that is used to send BUM traffic for a VLAN
E-tree Leaf Nexthop	E-tree leaf nexthop identifier.
Remote PE	IP address of the remote provider edge router.
Transport Labels	The LDP/RSVP transport label for the P2MP LSP for the remote PE.

Sample Output

show evpn p2mp

```
user@host> show evpn p2mp
regress@PE1_re> show evpn p2mp
Instance: ALPHA
  VLAN       Multicast Nexthop   E-tree Leaf Nexthop
  100        1048576
    Remote PE         VLAN     Transport Labels
    10.255.0.2        100      300448
    10.255.0.3        100      300560
    10.255.0.4        100      300608
```

show evpn p2mp extensive

```
user@host> show evpn p2mp extensive
Instance: ALPHA
  VLAN ID: 100
    Multicast Nexthop: 1048576
      P2MP Tunnel: 10.255.0.1:100:ldp-p2mp:evpn:100:ALPHA
      PMSI: Flags 0x0: Label 0: LDP-P2MP: Root 10.255.0.1 , lsp-id 16777217
        ESI: 00:11:22:33:44:55:66:77:88:99 SH Nexthop: 1048577
  Neighbor Address: 10.255.0.2, VLAN ID: 100
    Transport Labels: 300448
      PMSI: Flags 0x0: Label 0: LDP-P2MP: Root 10.255.0.2 , lsp-id 16777217
  Neighbor Address: 10.255.0.3, VLAN ID: 100
    Transport Labels: 300560
      PMSI: Flags 0x0: Label 0: LDP-P2MP: Root 10.255.0.3 , lsp-id 16777217
  Neighbor Address: 10.255.0.4, VLAN ID: 100
    Transport Labels: 300608
      PMSI: Flags 0x0: Label 0: LDP-P2MP: Root 10.255.0.4 , lsp-id 16777217
```

Release Information

Command introduced in Junos OS Release 18.2.

show evpn peer-gateway-macs

IN THIS SECTION

- Syntax | 19837
- Description | 19837
- Options | 19837
- Required Privilege Level | 19837
- Release Information | 19837

Syntax

```
show evpn peer-gateway-macs
<address>
<instance instance-name>
```

Description

Show Ethernet VPN (EVPN) peer gateway MAC information.

Options

none	Display brief information about the EVPN peer gateway MAC.
address	(Optional) Display peer gateway information for the specified MAC address.
instance *instance-name*	(Optional) Display peer gateway MAC information for the specified routing instance.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

> *Example: Configuring an EVPN with IRB Solution on EX9200 Switches*

show evpn prefix

IN THIS SECTION

- Syntax | 19838
- Description | 19838
- Required Privilege Level | 19838
- Output Fields | 19838
- Sample Output | 19839
- Release Information | 19839

Syntax

```
show evpn prefix
```

Description

Display Ethernet VPN (EVPN) information for imported and exported prefixes.

Required Privilege Level

view

Output Fields

The table 2 lists the output fields for the `show evpn ip-prefix-database` command. Output fields are listed in the approximate order in which they appear.

Table 523: show evpn prefix Output Fields

Field Name	Field Description
VLAN	VLAN identifier of the Layer 2 circuit.

Table 523: show evpn prefix Output Fields *(Continued)*

Field Name	Field Description
VNI	VXLAN network identifier for the Layer 3 virtual and routing forwarding (VRF) for the customer or tenant domain.
Prefix	MAC address associated with the IP prefix.
Active Source	The IP address or the ESI value or the IRB interface name of the source.
Timestamp	Month, day and time of generated command output.

Sample Output

show evpn prefixes

```
user@host > show evpn prefixes
Instance: ALPHA
VLAN  VNI    Prefix                    Active source                  Timestamp
      9100   10.0.0.0/24               10.255.0.2                     May 22 17:16:12
      9100   20.0.0.0/24               00:11:22:33:44:55:66:77:88:99  May 22 22:33:35
      9100   30.0.0.0/24               irb.0                          May 22 17:13:31
      9100   40.0.0.0/24               10.255.0.3                     May 22 17:14:20
```

Release Information

Command introduced in Junos OS Release 17.1.

show evpn vpws-instance

IN THIS SECTION

- Syntax | 19840
- Description | 19840
- Options | 19840
- Required Privilege Level | 19840
- Output Fields | 19840
- Sample Output | 19845
- Release Information | 19848

Syntax

```
show evpn vpws-instance
evpn-vpws-instance-name
```

Description

Display the details of the VPWS instance of the EVPN.

Options

evpn-vpws-instance-name	(Optional) Name of the VPWS instance of the EVPN.

Required Privilege Level

view

Output Fields

Table 524 on page 19841 describes the output fields for the `show evpn vpws-instance` command. Output fields are listed in the approximate order in which they appear.

Table 524: show evpn vpws-instance Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Route Distinguisher	Route distinguisher as configured for the routing instance on the node.
Number of local interfaces	Number of interfaces configured to be part of the routing instance.
Interface name	Name of the interface.
ESI	Ethernet segment identifier (ESI) configured for the interface.
Mode	Single-homed, single-active, or all-active form of multihoming.
Role	Primary or backup, depending on the role. The role, depending on the mode, is as follows: • Single-active multihoming–Role is primary if the interface is the designated forwarder(DF) and backup if the interface is the non-DF. • All-active multihoming–Role is always primary. • Single-homed–Role is always primary.
Status	Status of the interface is either Up or Down. Interface is available if the status is Up and interface is not available if the interface is Down.
Control-Word	Whether the device has negotiated to use the control word for MPLS pseudowire traffic (Yes) or not (No, the default).
Flow-Label-Tx	Flow-aware transport of pseudowires (FAT) flow labels is enabled (Yes) or is not enabled (No, the default) in the routing instance to push flow labels in pseudowire packets sent to a remote provider edge (PE) router.
Flow-Label-Rx	Flow-aware transport of pseudowires (FAT) flow labels is enabled (Yes) or is not enabled (No, the default) in the routing instance to pop flow labels from pseudowire packets from a remote PE router.

Table 524: show evpn vpws-instance Output Fields *(Continued)*

Field Name	Field Description
Local SID	Local service identifier configured on the interface.
Advertised Label	Label advertised for the service identifier as part of the autodiscovery route per the EVPN instance (EVI).
Srv6 SID mode	• Classic SID—Indicates classic SID SRv6 routing instance. • Micro SID—Indicates micro-SID SRv6 routing instance.
PE routers hosting the local service identifier as configured on the interface	
PE addr	IP address of the PE router advertising the autodiscovered route per the EVI.
ESI	ESI encoded in the autodiscovered route per the EVI.
Label	Label as encoded in the autodiscovered route per the EVI.
Mode	Single-homed, single-active, or all-active form of multihoming.
Role	Primary or backup depending on the role. The role depending on the mode is as follows: • Single-active multihoming–Role is primary if the interface is designated forwarder(DF) and backup if the interface is non-DF. • All-active multihoming–Role is always primary. • Single-homed–Role is always primary.
TS	Timestamp of the receipt of the autodiscovered route per the EVI. This is used as tie breaker in case there are two PE routers advertising the autodiscovered route per the EVI with the P bit set.

Table 524: show evpn vpws-instance Output Fields *(Continued)*

Field Name	Field Description
Status	The status can be as follows: • Resolved—The autodiscovered route per the ESI was received for the ESI from the PE router. • Unresolved—The autodiscovered route per the ESI was not received for the ESI from the PE router.
PE routers hosting the remote service identifier as configured on the interface	
Remote SID	Remote service identifier configured on the interface.
PE addr	IP address of the PE router advertising the autodiscovered route per the EVI.
ESI	ESI encoded in the autodiscovered route per the EVI.
Label	Label as encoded in the autodiscovery route per the EVI.
Mode	Single-homed, single-active, or all-active form of multihoming.
Role	Primary or backup depending on the role. The role depending on the mode is as follows: • Single-active multihoming–Role is primary if the interface is designated forwarder(DF) and backup if the interface is non-DF. • All-active multihoming–Role is always primary. • Single-homed–Role is always primary.
TS	Timestamp of the receipt of the autodiscovered route per the EVI. This is used as tie breaker in case there are two PE routers advertising the autodiscovered route per the EVI with the P bit set.

Table 524: show evpn vpws-instance Output Fields *(Continued)*

Field Name	Field Description
Status	The status can be as follows: • Resolved—The autodiscovered route per the ESI was received for the ESI from the PE router. • Unresolved—The autodiscovered route per the ESI was not received for the ESI from the PE router.
Fast Convergence Information	Details of the ESI, the PE router, and the advertised service identifier used during fast convergence.
ESI	ESI as advertised in the autodiscovered route per the ESI.
Number of PE nodes	Number of PE nodes available.
PE	IP address of the PE router advertising the autodiscovered route per the ESI.
Advertised SID	List of service identifiers that get impacted if the PE router withdraws the autodiscovered route per the ESI.
DF Election Information for Single-Active ESI	Details of the DF election information for the single-active ESI.
ESI	ESI as advertised in the Ethernet segment route.
SID used for DF Election	Minimum service identifier configured for the ESI in the EVI.
ESI granularity	Type: Per ESI.
LACP OOS on NDF	Enabled or Disabled
Primary PE	PE router that is the DF for the ESI.

Table 524: show evpn vpws-instance Output Fields *(Continued)*

Field Name	Field Description
Backup PE	PE router that is the backup DF for the ESI.
Last DF Election	Last time when the DF election algorithm was executed on the PE router.

Sample Output

show evpn vpws-instance *instance-name*

```
user@host> show evpn vpws-instance vpws1004
Instance: vpws1004
  Route Distinguisher: 100.100.100.4:1004
  Number of local interfaces: 1 (1 up)

    Interface name  ESI                             Mode            Role       Status
    ge-0/0/1.1004   00:00:00:00:00:00:00:00:00:00  single-homed    Primary    Up
        Local SID: 1004 Advertised Label: 301360
        Remote SID: 2004
            PE addr         ESI                           Label  Mode           Role
TS                      Status
            100.100.100.2   00:11:11:11:11:11:11:11:11:11 301888 all-active     Primary
2017-01-11 21:57:31.916 Resolved
            100.100.100.1   00:11:11:11:11:11:11:11:11:11 301952 all-active     Primary
2017-01-11 21:59:28.491 Resolved
            100.100.100.3   00:11:11:11:11:11:11:11:11:11 301984 all-active     Primary
2017-01-11 21:59:28.522 Resolved

    Fast Convergence Information
    ESI: 00:11:11:11:11:11:11:11:11:11 Number of PE nodes: 3
        PE: 100.100.100.2
            Advertised SID: 2004
        PE: 100.100.100.1
            Advertised SID: 2004
        PE: 100.100.100.3
            Advertised SID: 2004
```

show evpn vpws-instance *instance-name*

```
user@host>show evpn vpws-instance vpws1003
Instance: vpws1003
  Route Distinguisher: 100.100.100.4:1003
  Number of local interfaces: 2 (2 up)

    Interface name  ESI                             Mode            Role       Status
    ae10.2003       00:44:44:44:44:44:44:44:44:44  all-active      Primary    Up
        Local SID: 2003 Advertised Label: 301312
            PE addr         ESI                           Label  Mode           Role
TS                      Status
            100.100.100.3   00:44:44:44:44:44:44:44:44:44 301792 all-active     Primary
2017-01-11 21:59:28.498 Resolved
        Remote SID: 1003
            Local Interface: ge-0/0/1.1003 (Up)

    Interface name  ESI                             Mode            Role       Status
    ge-0/0/1.1003   00:00:00:00:00:00:00:00:00:00  single-homed    Primary    Up
        Local SID: 1003 Advertised Label: 301296
        Remote SID: 2003
            Local Interface: ae10.2003 (Up)
            PE addr         ESI                           Label  Mode           Role
TS                      Status
            100.100.100.3   00:44:44:44:44:44:44:44:44:44 301792 all-active     Primary
2017-01-11 21:59:28.498 Resolved

    Fast Convergence Information
    ESI: 00:44:44:44:44:44:44:44:44:44 Number of PE nodes: 1
        PE: 100.100.100.3
            Advertised SID: 2003
```

show evpn vpws-instance *instance-name*(Per ESI)

```
user@host>show evpn vpws-instance vpws1001
    DF Election Information for Single-Active ESI
    ESI: 00:11:11:11:11:11:11:11:11:11
    DF Election Algorithm: MOD based
    SID used for DF Election: 1001
    ESI granularity: Per ESI
```

```
    LACP OOS on NDF: Enabled
    Primary PE: 10.10.10.2
    Backup PE: 10.10.10.1
    Last DF Election: 2020-04-16 16:04:32
```

show evpn vpws-instance *instance-name* (with FAT flow label load balancing enabled)

```
user@host> show evpn vpws-instance
Instance: VPWS-MH-SA, Instance type: EVPN VPWS
  Route Distinguisher: 10.255.0.1:200
  Number of local interfaces: 1 (1 up)

    Interface name  ESI                           Mode          Role     Status   Control-Word
Flow-Label-Tx  Flow-Label-Rx
    ge-0/0/2.200    00:00:00:00:00:00:00:22:22:22 single-active Primary  Up       No
Yes            Yes
        Local SID: 102 Advertised Label: 299808
            PE addr      ESI                           Label  Mode           Role
TS                      Status
            10.255.0.2   00:00:00:00:00:00:00:22:22:22 299808 single-active  Backup   2021-03-08
17:19:02.107 Resolved
        Remote SID: 202
            PE addr      ESI                           Label  Mode           Role
TS                      Status
            10.255.0.3   00:00:00:00:00:00:00:00:00:00 299808 single-homed   Primary  2021-03-08
17:17:53.121 Resolved
  Number of protect interfaces: 0

    Fast Convergence Information
    ESI: 00:00:00:00:00:00:00:22:22:22 Number of PE nodes: 1
        PE: 10.255.0.2
            Advertised SID: 102

    DF Election Information for Single-Active ESI
    ESI: 00:00:00:00:00:00:00:22:22:22
    DF Election Algorithm: MOD based
    SID used for DF Election: 102
    Primary PE: 10.255.0.1
    Backup PE: 10.255.0.2
    Last DF Election: 2021-03-08 17:19:02
```

show evpn vpws-instance (Micro-SID)

```
user@host> show evpn vpws-instance evpn-vpws-mh
Instance: evpn-vpws-mh, Instance type: EVPN VPWS, Encapsulation type: SRv6
  Route Distinguisher: 10.1.1.1:12
  Number of local interfaces: 1 (1 up)

    Interface name  ESI                             Mode          Role       Status     Control-
Word    Flow-Label-Tx    Flow-Label-Rx
    ae0.0           00:12:12:12:12:12:12:12:12:12 single-active   Primary    Up
No              No               No
        Local SID: 103 Advertised Label: 3 Advertised End.Dx2 SID: fcbb:bb01:100:e000:: Srv6 SID
mode: Micro-SID
        Remote SID: 301
            PE addr         ESI                             Label  End.Dx2 SID     Mode
Role     TS                      Status
            2001:db8:a3:1:: 00:00:00:00:00:00:00:00:00:00 0      fcbb:cc01:100:e000:: single-
homed Primary 2023-12-05 12:28:05.760 Resolved
  Number of protect interfaces: 0

    DF Election Information for Single-Active ESI
    ESI: 00:12:12:12:12:12:12:12:12:12
    DF Election Algorithm: MOD based
    SID used for DF Election: 103
    LACP OOS on NDF: Enabled
    Primary PE: abcd::128:205:181:57
    ESI granularity: Per ESI
    Last DF Election: 2023-12-05 19:28:40
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Overview of VPWS with EVPN Signaling Mechanisms

show extensible-subscriber-services counters

IN THIS SECTION

- Syntax | 19849
- Description | 19849
- Required Privilege Level | 19849
- Output Fields | 19849
- Sample Output | 19850
- Release Information | 19851

Syntax

```
show extensible-subscriber-services counters
```

Description

Display number of times various events were processed by Extensible Subscriber Services Manager.

Required Privilege Level

view

Output Fields

Table 525 on page 19849 describes the output fields for the `show extensible-subscriber-services counters` command. Output fields are listed in the approximate order in which they appear.

Table 525: show extensible-subscriber-services counters Output Fields

Field Name	Field Description
Total Sessions	Number of extensible-subscriber-service control sessions that are currently in the system

Table 525: show extensible-subscriber-services counters Output Fields *(Continued)*

Field Name	Field Description
Total Services	Number of extensible-subscriber-services that are currently in the system
Active Services	Number of extensible-subscriber-services that are currently in active state
Op-script execution	Total number of operational scripts executed: • Successful—Operational scripts that are executed without error • Unsuccessful—Operational scripts that encountered errors
Commit execution	Total number of commit scripts executed: • Successful—Commit scripts that are executed without error • Unsuccessful—Commit scripts that encountered errors
Application execution	Number of applications executed: • Successful—Applications that are executed without error • Unsuccessful—Applications that encountered errors
Service requests	Number of service requests received by the daemon. It has counters for total number of requests received, acknowledged (Acked), and negatively acknowledged (Nacked).

Sample Output

command-name

```
# show extensible-subscriber-services counters

Total Sessions:         7
Total Services:         14
Active Services:        14

Op-script execution:
```

```
    Successful          30
    Unsuccessful        0
Commit execution:
    Successful          4
    Unsuccessful        0
Application execution:
    Successful          0
    Unsuccessful        0
Service requests:
    Received            30
    Acked               30
    Nacked              0
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

clear extensible-subscriber-services counters | 3593

show extensible-subscriber-services debug-information

IN THIS SECTION

- Syntax | 19852
- Description | 19852
- Required Privilege Level | 19852
- Output Fields | 19852
- Sample Output | 19852
- Release Information | 19852

Syntax

```
show extensible-subscriber-services debug-information
```

Description

Write service session information to a file.

Required Privilege Level

view

Output Fields

Table 526 on page 19852 lists the output fields for the `show extensible-subscriber-services debug-information` command.

Table 526: show extensible-subscriber-services debug-information Output Fields

Field Name	Field Description
File Name	Name of the file that contains the service session information

Sample Output

command-name

```
#show extensible-subscriber-services debug-information
File Name: /var/tmp/dump_session_record
```

Release Information

Command introduced in Junos OS Release 15.1.

show extensible-subscriber-services dictionary

IN THIS SECTION

- Syntax | 19853
- Description | 19853
- Required Privilege Level | 19853
- Output Fields | 19853
- Sample Output | 19855
- Release Information | 19858

Syntax

```
show extensible-subscriber-services dictionary
```

Description

Display contents of the dictionary including attributes and services.

Required Privilege Level

view

Output Fields

Table 527 on page 19854 lists the output fields for the `show extensible-subscriber-services dictionary` command. Output fields are listed in the approximate order in which they appear.

Table 527: show extensible-subscriber-services dictionary Output Fields

Field Name	Field Description
Acct-Session-Id	Unique accounting identifier that makes it easy to match start and stop records in a log file. The identifier can be in one of the following formats: • `decimal`—For example, `435264` • `description`—In the generic format, `jnpr` *`interface-specifier`*:*`subscriber-session-id`*; For example, `jnpr fastEthernet 3/2.6:1010101010101`
ERX-Service-Activate	Service to be activated for the subscriber.
Service-Type	Type of service the user has requested or the type of service to be provided.
ADSL-Agent-Remote-Id	Identifier for the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request.
ERX-Med-Port-Number	UDP port in the mediation device to which mirrored traffic is forwarded
NAS-Port-Id	Text string that identifies the physical interface of the NAS that is authenticating the user.
ERX-Med-Dev-Handle	Identifier that associates mirrored traffic with a specific subscriber
ERX-Service-Deactivate	Service to be deactivated for the subscriber.
NAS-IP-Address	IP address of the network access server (NAS) that is requesting authentication of the user.
NAS-Identifier	Name of the NAS that originated the authentication or accounting request.
ERX-Service-Acct-Interval	Amount of time between accounting updates for the service type

Table 527: show extensible-subscriber-services dictionary Output Fields *(Continued)*

Field Name	Field Description
ADSL-Agent-Circuit-Id	Identifier for subscriber's access node and the digital subscriber line (DSL) on the access node.
ERX-LI-Action	Traffic mirroring action. For dynamic CoA, VSA 26-58 changes the action on the mirrored traffic identified by VSA 26–59. CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26–58, 26–59, 26–60, or 26–61) must always include all four VSAs. If the CoA action is to stop mirroring (VSA 26–58 value is 0), then the values of the other three attributes in the CoA message must match the existing attribute values, or the action fails.
ERX-Med-Ip-Address	IP address of content destination device to which mirrored traffic is forwarded.

Sample Output

command-name

```
root@ce-bras-mx240-g> show extensible-subscriber-services dictionary

RADIUS DICTIONARY

  Attribute Name              Vendor ID   Attribute Code    Attribute Type    Has Tag   Config
Type
  Class                       0           25                string            FALSE
INVALID
  Acct-Session-Id             0           44                string            FALSE
INVALID
  ERX-Service-Activate        4874        65                string            TRUE
SERVICE_ACTIVATE
  Service-Type                4874        173               integer           TRUE
INVALID
  ADSL-Agent-Remote-Id        3561        2                 string            FALSE
INVALID
```

```
  ERX-Med-Port-Number         4874        61                OctetString       FALSE
INVALID
  NAS-Port-Id                 0           87                text              FALSE
INVALID
  ERX-Med-Dev-Handle          4874        59                OctetString       FALSE
INVALID
  ERX-Service-Deactivate      4874        66                string            TRUE
SERVICE_DEACTIVATE
  NAS-IP-Address              0           4                 address           FALSE
INVALID
  NAS-Identifier              0           32                string            FALSE
INVALID
  ERX-Service-Acct-Interval   4874        140               integer           TRUE
INVALID
  ADSL-Agent-Circuit-Id       3561        1                 string            FALSE
INVALID
  ERX-LI-Action               4874        58                OctetString       FALSE
SERVICE_ACTIVATE
  ERX-Med-Ip-Address          4874        60                OctetString       FALSE
INVALID

  Services List

  Service Name: ngcoco
    Service Attribute Name      : ERX-Service-Activate

    PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_ngcoco_add

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Activate

    DE-PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_del

      PARAMETERS
        Acct-Session-Id
```

```
        NAS-Port-Id
        ERX-Service-Activate

Service Name: dhcprelay
  Service Attribute Name      : ERX-Service-Activate

  PROVISION ACTION
    Action Type               : OP-SCRIPT
    Action Version            : 1
    Action Name               : iceaaa_dhcprelay_add

    PARAMETERS
      Acct-Session-Id
      NAS-Port-Id
      ERX-Service-Activate

  DE-PROVISION ACTION
    Action Type               : OP-SCRIPT
    Action Version            : 1
    Action Name               : iceaaa_del

    PARAMETERS
      Acct-Session-Id
      NAS-Port-Id
      ERX-Service-Activate

Service Name: default
  Service Attribute Name      : ERX-LI-Action

  PROVISION ACTION
    Action Type               : APPLICATION
    Action Version            : 1
    Action Name               : LI

    PARAMETERS
      ERX-Med-Dev-Handle
      ERX-Med-Ip-Address
      ERX-Med-Port-Number

Service Name: dhcprelay
  Service Attribute Name      : ERX-Service-Deactivate

  DE-PROVISION ACTION
```

```
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_del

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Deactivate

  Service Name: ngcoco
    Service Attribute Name      : ERX-Service-Deactivate

    DE-PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_del

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Deactivate
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

- dictionary | 5703
- request services extensible-subscriber-services reload-dictionary | 14942
- show extensible-subscriber-services dictionary attributes | 19859
- show extensible-subscriber-services dictionary services | 19863
- *Understanding the Dictionary File*

show extensible-subscriber-services dictionary attributes

IN THIS SECTION

- Syntax | 19859
- Description | 19859
- Required Privilege Level | 19859
- Output Fields | 19859
- Sample Output | 19861
- Release Information | 19862

Syntax

```
show extensible-subscriber-services dictionary attributes
```

Description

Display contents of the dictionary file. It shows only the parameter attributes in the dictionary.

Required Privilege Level

view

Output Fields

Table 528 on page 19860lists the output fields for the `show extensible-subscriber-services dictionary attributes` command. Output fields are listed in the approximate order in which they appear.

Table 528: show extensible-subscriber-services dictionary attributes Output Fields

Field Name	Field Description
Acct-Session-Id	Unique accounting identifier that makes it easy to match start and stop records in a log file. The identifier can be in one of the following formats: • `decimal`—For example, `435264` • `description`—In the generic format, `jnpr` *`interface-specifier`*:*`subscriber-session-id`*; For example, `jnpr fastEthernet 3/2.6:1010101010101`
ERX-Service-Activate	Service to be activated for the subscriber.
Service-Type	Type of service the user has requested or the type of service to be provided.
ADSL-Agent-Remote-Id	Identifier for the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request.
ERX-Med-Port-Number	UDP port in the mediation device to which mirrored traffic is forwarded
NAS-Port-Id	Text string that identifies the physical interface of the NAS that is authenticating the user.
ERX-Med-Dev-Handle	Identifier that associates mirrored traffic with a specific subscriber
ERX-Service-Deactivate	Service to be deactivated for the subscriber.
NAS-IP-Address	IP address of the network access server (NAS) that is requesting authentication of the user.
NAS-Identifier	Name of the NAS that originated the authentication or accounting request.
ERX-Service-Acct-Interval	Amount of time between accounting updates for the service type

Table 528: show extensible-subscriber-services dictionary attributes Output Fields *(Continued)*

Field Name	Field Description
ADSL-Agent-Circuit-Id	Identifier for subscriber's access node and the digital subscriber line (DSL) on the access node.
ERX-LI-Action	Traffic mirroring action. For dynamic CoA, VSA 26-58 changes the action on the mirrored traffic identified by VSA 26–59. CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26–58, 26–59, 26–60, or 26–61) must always include all four VSAs. If the CoA action is to stop mirroring (VSA 26–58 value is 0), then the values of the other three attributes in the CoA message must match the existing attribute values, or the action fails.
ERX-Med-Ip-Address	IP address of content destination device to which mirrored traffic is forwarded.

Sample Output

command-name

```
root@ce-bras-mx240-g> show extensible-subscriber-services dictionary attributes

RADIUS DICTIONARY

  Attribute Name              Vendor ID   Attribute Code   Attribute Type  Has Tag  Config Type
  Class                       0           25               string          FALSE
INVALID
  Acct-Session-Id             0           44               string          FALSE
INVALID
  ERX-Service-Activate        4874        65               string          TRUE
SERVICE_ACTIVATE
  Service-Type                4874        173              integer         TRUE
INVALID
  ADSL-Agent-Remote-Id        3561        2                string          FALSE
INVALID
  ERX-Med-Port-Number         4874        61               OctetString     FALSE
```

```
INVALID
  NAS-Port-Id                 0           87                text            FALSE
INVALID
  ERX-Med-Dev-Handle          4874        59                OctetString     FALSE
INVALID
  ERX-Service-Deactivate      4874        66                string          TRUE
SERVICE_DEACTIVATE
  NAS-IP-Address              0           4                 address         FALSE
INVALID
  NAS-Identifier              0           32                string          FALSE
INVALID
  ERX-Service-Acct-Interval   4874        140               integer         TRUE
INVALID
  ADSL-Agent-Circuit-Id       3561        1                 string          FALSE
INVALID
  ERX-LI-Action               4874        58                OctetString     FALSE
SERVICE_ACTIVATE
  ERX-Med-Ip-Address          4874        60                OctetString     FALSE
INVALID
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

- dictionary | **5703**
- request services extensible-subscriber-services reload-dictionary | **14942**
- show extensible-subscriber-services dictionary | **19853**
- show extensible-subscriber-services dictionary services | **19863**
- *Understanding the Dictionary File*

show extensible-subscriber-services dictionary services

IN THIS SECTION

- Syntax | 19863
- Description | 19863
- Required Privilege Level | 19863
- Output Fields | 19863
- Sample Output | 19864
- Release Information | 19866

Syntax

```
show extensible-subscriber-services dictionary services
```

Description

Display services configured in the dictionary file and their attributes.

Required Privilege Level

view

Output Fields

Table 529 on page 19863lists the output fields for the `show extensible-subscriber-services dictionary services` command. Output fields are listed in the approximate order in which they appear.

Table 529: show extensible-subscriber-services dictionary services Output Fields

Field Name	Field Description
Service Name	Name of the service specified in the dictionary.

Table 529: show extensible-subscriber-services dictionary services Output Fields *(Continued)*

Field Name	Field Description
Service Attribute Name	Name of the service attribute.
PROVISION ACTION	Name, type, and version of the provisioning action
DE-PROVISION ACTION	Name, type, and version of the deprovisioning action
PARAMETERS	Parameters to be processed for the provisioning or deprovisioning action.

Sample Output

command-name

```
root@ce-bras-mx240-g> show extensible-subscriber-services dictionary services

RADIUS DICTIONARY

  Services List

  Service Name: ngcoco
    Service Attribute Name      : ERX-Service-Activate

    PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_ngcoco_add

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Activate

    DE-PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
```

```
      Action Name               : iceaaa_del

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Activate

  Service Name: dhcprelay
    Service Attribute Name      : ERX-Service-Activate

    PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_dhcprelay_add

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Activate

    DE-PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_del

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Activate
  Service Name: default
    Service Attribute Name      : ERX-LI-Action

    PROVISION ACTION
      Action Type               : APPLICATION
      Action Version            : 1
      Action Name               : LI

      PARAMETERS
        ERX-Med-Dev-Handle
        ERX-Med-Ip-Address
        ERX-Med-Port-Number

  Service Name: dhcprelay
```

```
    Service Attribute Name      : ERX-Service-Deactivate

    DE-PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_del

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Deactivate

  Service Name: ngcoco
    Service Attribute Name      : ERX-Service-Deactivate

    DE-PROVISION ACTION
      Action Type               : OP-SCRIPT
      Action Version            : 1
      Action Name               : iceaaa_del

      PARAMETERS
        Acct-Session-Id
        NAS-Port-Id
        ERX-Service-Deactivate
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

- dictionary | 5703
- request services extensible-subscriber-services reload-dictionary | 14942
- show extensible-subscriber-services dictionary | 19853
- show extensible-subscriber-services dictionary attributes | 19859
- *Understanding the Dictionary File*

show extensible-subscriber-services service

IN THIS SECTION

- Syntax | 19867
- Description | 19867
- Required Privilege Level | 19867
- Output Fields | 19867
- Sample Output | 19868
- Release Information | 19868

Syntax

```
show extensible-subscriber-services service
```

Description

Display service information. It displays the service name specified in the dictionary, scripts used for provisioning and deprovisioning, and the number of services. A service can have more than one version. Each version is grouped using an action set. The number of action sets is equal to the number of versions for the service. Number of services in all the versions is displayed in **Total Services**.

Required Privilege Level

view

Output Fields

Table 530 on page 19868 lists the output fields for the `show extensible-subscriber-services service` command. Output fields are listed in the approximate order in which they appear.

Table 530: show extensible-subscriber-services service Output Fields

Field Name	Field Description
Service Name	Name of the service
Action	Name and number of services provisioned or deprovisioned.
Total Services	Total number of services provisioned or deprovisioned.

Sample Output

command-name

```
#show extensible-subscriber-services services
Service Name: ngcoco
    Action:
        Provision:            iceaaa_ngcoco_add_1
        Deprovision:          iceaaa_del_1
        Services:                 1301

    Total Services:              1301

Service Name: dhcprelay
    Action:
        Provision:            iceaaa_dhcprelay_add_1
        Deprovision:          iceaaa_del_1
        Services:                 1301

    Total Services:             1301
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding the Dictionary File

show extensible-subscriber-services sessions

IN THIS SECTION

- Syntax | 19869
- Description | 19869
- Options | 19869
- Required Privilege Level | 19869
- Output Fields | 19870
- Sample Output | 19870
- Release Information | 19871

Syntax

```
show extensible-subscriber-services sessions accounting-session-id
```

Description

Display session information. It displays a list of the services applied and their current state. Specify the accounting session ID to view information about a specific session. If an accounting session ID is not specified in the command, the command displays details of all the sessions.

Options

accounting-session-id	Identifier of the session you want information about.

Required Privilege Level

view

Output Fields

Table 531 on page 19870 describes the output fields for the `show extensible-subscriber-services sessions` command. Output fields are listed in the approximate order in which they appear.

Table 531: show extensible-subscriber-services sessions Output Fields

Field Name	Field Description
Session ID	Session ID. Accounting session ID must be enclosed in quotation marks.
Timestamp	Time when the session was started
Service Name	Name of the service
Commit State	Commit state: `Init`, `Queued`, or `Success`.
Service ID	Service ID

Sample Output

command-name

```
#show extensible-subscriber-services sessions
Session ID: jnpr demux0.1073762028:46422
    Timestamp: Fri Mar  8 04:58:27 2013
    Service Name: ngcoco
    Service Name: dhcprelay

Total Sessions: 1

# show extensible-subscriber-services sessions "jnpr demux0.1073762028:46422"

Service ID: jnpr demux0.1073762028:46422:46425
    Timestamp: Fri Mar  8 04:58:27 2013
    Service Name: ngcoco
    Commit State: Success
```

```
Service ID: jnpr demux0.1073762028:46422:46426
    Timestamp: Fri Mar  8 04:58:27 2013
    Service Name: dhcprelay
    Commit State: Success

Total Services: 2
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

clear extensible-subscriber-services sessions | 3594

show extension-service request-response clients

IN THIS SECTION

- Syntax | 19871
- Description | 19872
- Options | 19872
- Required Privilege Level | 19872
- Output Fields | 19872
- Sample Output | 19873
- Release Information | 19874

Syntax

```
show extension-service request-response clients (detail | brief)
```

Description

Display the status of the request-response clients connected to the device.

Options

client-id The client identifier.

brief (Default) Display a summary of the information.

detail Display detailed information.

Required Privilege Level

view

Output Fields

Table 532 on page 19872 lists the output fields for the `show extension-service request-response clients` command.

Table 532: show extension-service request-response clients Output Fields

Field Name	Field Description
`Client ID`	Client identifier.
`Socket Address`	Address of the socket.
`Client Type`	Type of the client.
`Client Login Time`	The most recent login time of the remote procedure call (gRPC) client. This is when the authentication request was received for the channel.
`Channel Count`	The number of channels.

Table 532: show extension-service request-response clients Output Fields *(Continued)*

Field Name	Field Description
User Name	The user name for which the session was authenticated in a gRPC session. If authentication is not required, this field displays as “No User.” This helps you identify which users have requested programmable operations.

Sample Output

show extension-service request-response clients

```
user@device> show extension-service request-response clients

Client ID                      Socket Address                Client Type Client Login Time
Channel Count
ipv6:::ffff:10.209.0.224:45888 ipv6:::ffff:10.209.0.224:45888 gRPC       2020-02-24 04:08:11
UTC    1

unix::20               unix::20                           gRPC             2020-02-23 15:23:47
UTC      1
```

show extension-service request-response clients detail

```
user@device> show extension-service request-response clients detail
Channel information:
  Client ID: ipv6:::ffff:10.209.0.224:45888
  Socket Address: ipv6:::ffff:10.209.0.224:45888
  Client Type: gRPC
  Client Login Time : 2020-02-24 04:08:11 UTC
  Channel Count: 1

    Channel target: unix:/var/run/japi_mgd-api
    Channel status: GRPC_CHANNEL_READY
    User Name: root

Channel information:
  Client ID: unix::20
```

```
  Socket Address: unix::20
  Client Type: gRPC
  Client Login Time : 2020-02-23 15:23:47 UTC
  Channel Count: 1

    Channel target: unix:/var/run/japi_na-grpcd
    Channel status: GRPC_CHANNEL_READY
    User Name: No User
```

Release Information

Command introduced before Junos OS Release 18.2R1.

`Client Login Time` and `User Name` output fields introduced in Junos OS Release 20.4R1.

show extension-service request-response servers

IN THIS SECTION

- Syntax | 19874
- Description | 19874
- Required Privilege Level | 19875
- Output Fields | 19875
- Sample Output | 19875
- Release Information | 19875

Syntax

```
show extension-service request-response servers
```

Description

Display the status of the request-response servers connected to the device.

Required Privilege Level

view

Output Fields

Table 533 on page 19875 lists the output fields for the `show extension-service request-response servers` command.

Table 533: show extension-service request-response servers Output Fields

Field Name	Field Description
`Max connections`	The maximum number of simultaneous connections for request-response that can be attached to jsd.
`Address`	The address of the server.
`Status`	The status of the server.

Sample Output

show extension-service request-response servers

```
user@device> show extension-service request-response servers
gRPC server information:
  Max connections: 5, Skip-authentication: Disabled

  Address: unix:/var/run/japi_jsd
  Status: Up, Type: Clear-text
```

Release Information

Command introduced before Junos OS Release 18.2R1.

show extension-service status

IN THIS SECTION

- Syntax | 19876
- Description | 19876
- Options | 19876
- Required Privilege Level | 19876
- Output Fields | 19877
- Sample Output | 19877
- Release Information | 19878

Syntax

```
show extension-service status (application-name | all)
```

Description

Display the status of all JET applications.

> **NOTE**: The `show extension-service status` operational command is limited to use with Python applications only.

Options

application-name	Display information for a single application.
all	Display information for all JET applications running on the system.

Required Privilege Level

view

Output Fields

Table 534 on page 19877 lists the output fields for the `show extension-service status` command.

Table 534: show extension-service status Output Fields

Field Name	Field Description
`Name`	Name of the application.
`Arguments`	Arguments passed to the application.
`Process-id`	Process ID.
`Stack-Segment-Size`	Size of the stack segment memory.
`Data-Segment-Size`	Size of the data segment memory.

Sample Output

show extension-service status

```
user@host> show extension-service status application-one
Extension service application details:
Name : application-one
Arguments: -arg1 foo -arg2 goo
Process-id: 52592
Stack-Segment-Size: 16777216B
Data-Segment-Size: 134217728B
```

show extension-service status all

```
user@host> show extension-service status all
Extension service application details:
Name : application-name1
Arguments: -arg1 foo -arg2 goo
```

```
Process-id: 54834
Stack-Segment-Size: 16777216B
Data-Segment-Size: 134217728B
Name : application-name2
Arguments: -arg1 foo -arg2 goo
Process-id: 55011
Stack-Segment-Size: 16777216B
Data-Segment-Size: 134217728B
```

show extension-service status all (when no applications are active)

```
user@host> show extension-service status all
warning: No active extension-services
```

Release Information

Command introduced in Junos OS Release 16.1.

show fabric administration inventory

IN THIS SECTION

- Syntax | 19879
- Description | 19879
- Options | 19879
- Required Privilege Level | 19880
- Output Fields | 19880
- Sample Output | 19882
- Release Information | 19885

Syntax

```
show fabric administration inventory
<brief | detail | summary | terse>
<director-group (status)>
<infrastructure  (fabric-controls | fabric-managers | diagnostic-routing-engines)>
<interconnect-devices interconnect-device-name>
<node-devices node-device-name>
<node-groups node-group-name>
<summary>
```

Description

(QFabric systems only) Display all devices that belong to the QFabric system. You can narrow the level of output by specifying a device type.

> **NOTE**: If your Node devices do not appear in the output of the `show fabric administration inventory` command, check the cabling of your system.

Options

none	Display all devices within a QFabric system.
brief | detail | summary | terse	(Optional) Display the specified level of output.
director-group (status)	(Optional) Display the status for the Director group within a QFabric system.
infrastructure (fabric-controls | fabric-managers | diagnostic-routing-engines)	(Optional) Display information for the fabric control Routing Engine, fabric manager Routing Engine, and diagnostic Routing Engine running on the Director group for the QFabric system.
interconnect-devices *interconnect-device-name*	(Optional) Display a specific Interconnect device within a QFabric system.
node-devices *node-device-name*	(Optional) Display a specific Node device within a QFabric system.
node-groups *node-group-name*	(Optional) Display a specific Node group within a QFabric system.

Required Privilege Level

admin

Output Fields

Table 535 on page 19880 lists the output fields for the `show fabric administration inventory` command. Output fields are listed in the approximate order in which they appear.

Table 535: show fabric administration inventory Output Fields

Field Name	Field Description	Level of Outpu
Item	Type of QFabric system component being viewed. Possible values include `Node group`, `Interconnect device`, `Fabric control`, `Fabric manager`, `Diagnostic routing engine`, `Director group`, and `Ungrouped Node device`.	detail none
Identifier	Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.	detail none
Connection	Status of a QFabric system component: either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for the listed component.	detail none
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be `Configured`, `Failed` (with details about the failure), `Pending` (in the process of being written or retried), or `Unknown`.	detail none
Node group	Name of the Node groups associated with the QFabric system, and the Node devices assigned to each Node group. The group can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for the devices in the group. This field also displays the serial ID for the Node group and the status for the Node group.	detail none
Interconnect device	Name of the Interconnect devices associated with the Node group. The device can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for the device. This field also displays the serial ID and configuration status for the Interconnect device.	detail none

Table 535: show fabric administration inventory Output Fields *(Continued)*

Field Name	Field Description	Level of Outpu
Fabric manager	Name of the primary virtual Junos Routing Engine associated with the QFabric system. The fabric manager Routing Engine can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric manager Routing Engine.	detail none
Fabric control	Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition. The fabric control Routing Engine can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric control Routing Engine.	detail none
Diagnostic routing engine	Name of the virtual Junos Routing Engine responsible for troubleshooting and diagnostic utilities within a QFabric system partition. The diagnostic Routing Engine can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the diagnostic Routing Engine.	detail none
Director group	Identifier for the Director devices that are part of the Director group in a QFabric system. Each Director device can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for the device.	detail none
Connection uptime	Length of time the component has been operational. The time is listed in days, hours, minutes, and seconds (`D+HH:MM:SS`).	detail
Network domain	Indicates whether a Node group is a network Node group (`Yes`) or a server Node group (`No`).	detail
Member id	Member identification number for a Node device within a Node group.	detail
Node group master	Indicates whether or not a Node device acts as a primary device within a Node group. Values for this field are `Yes` or `No`.	detail

Table 535: show fabric administration inventory Output Fields *(Continued)*

Field Name	Field Description	Level of Outpu
Configuration checkout failed	Provides troubleshooting details about a failed configuration on a component. This field includes the following output: • `Edit path`—Displays the configuration hierarchy level where the problem occurs. • `Statement`—Displays the configuration statement that causes the problem. • `Message`—Provides a suggested workaround to resolve the problem.	detail
Member	Name of a Routing Engine allocated to an Interconnect device.	detail
Master	Indicates whether or not a Routing Engine acts as a primary device within an Interconnect device. Values are `Yes` or `No`.	detail
Total Node devices	Number of connected and disconnected Node devices in the QFabric system.	summary
Total connected Node devices	Number of available Node devices in the QFabric system.	summary
Total disconnected Node devices	Number of unavailable Node devices in the QFabric system.	summary
Total Interconnect devices	Number of Interconnect devices in the QFabric system.	summary

Sample Output

show fabric administration inventory

```
user@qfabric> show fabric administration inventory

Item                    Identifier              Connection      Configuration
Ungrouped Node device
  Node6                 BBAK8979                Disconnected
Node group
```

```
  P3359-C                                       Connected       Configured
    P3359-C                                     Connected
  P3865-C                                       Connected       Configured
    P3865-C                                     Connected
  RSNG-1                                        Connected       Configured
    Node-3              BBAK8276                Connected
    Node-4              BBAK8273                Connected
  NW-NG-0                                       Connected       Configured
    Node-0              BBAK8309                Connected
    Node-1              BBAK8283                Connected
Interconnect device
  IC-F1032                                      Connected       Configured
    F1032/RE0                                   Connected
    F1032/RE1                                   Connected
  IC-F1092                                      Connected       Configured
    F1092/RE0                                   Connected
    F1092/RE1                                   Connected
Fabric manager
  FM-0                                          Connected       Configured
Fabric control
  FC-0                                          Connected       Configured
  FC-1                                          Connected       Configured
Diagnostic routing engine
  DRE-0                                         Connected       Configured
Director group
  0281112011000023                              Connected
  0281112011000082                              Connected
```

show fabric administration inventory detail

```
user@qfabric> show fabric administration inventory detail
Node group: NW-NG-0, Connected, Configured
  Connection uptime: 2+22:40:46
  Network domain: Yes

  Node device: node01, Connected,
    Member id: 0

Node group: RSNG, Connected, Configured
  Connection uptime: 1:20:22
```

```
  Node device: node02, Connected,
    Member id: 0

  Node device: node03, Connected,
    Member id: 1
    Node group master: Yes

Node group: BBAK0423, Connected, Failed (invalid configuration)
  Connection uptime: 0:01:06

  Configuration checkout failed:
    Edit path: edit ethernet-switching-options
    Statement: analyzer
      Message: Vlan vlan1 which is configured as output vlan for analyzer session an1 contains
untagged interface.
               The analyzer output vlan should only contain tagged interface.

  Node device: node0, Connected,
    Member id: 0
    Node group master: Yes

Interconnect device: IC-F4912, Connected, Configured
  Connection uptime: 0:08:05

  Member: F4912/RE0, Connected,
    Connection uptime: 0:08:05
    Master: Yes

Fabric manager: FM-0, Connected, Configured
  Connection uptime: 2+22:40:47

Fabric control: FC-0, Connected, Configured
  Connection uptime: 2+22:17:38

Fabric control: FC-1, Connected, Configured
  Connection uptime: 2+22:17:57

Diagnostic routing engine: DRE-0, Connected, Configured
  Connection uptime: 2+22:39:56
```

show fabric administration inventory summary

```
user@qfabric> show fabric administration inventory summary
Total Node devices: 3
Total connected Node devices: 3
Total disconnected Node devices: 0
Total Interconnect devices: 1
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Performing the QFabric System Initial Setup on a QFX3100 Director Group

Configuring Aliases for the QFabric System

Configuring Node Groups for the QFabric System

show fabric administration inventory infrastructure | **19893**

show fabric administration inventory director-group status | **19885**

show fabric administration inventory node-devices | **19900**

show fabric administration inventory node-groups | **19903**

show fabric administration inventory interconnect-devices | **19898**

show fabric inventory | **19908**

show fabric administration inventory director-group status

IN THIS SECTION

- Syntax | **19886**
- Description | **19886**
- Options | **19886**
- Required Privilege Level | **19886**

- Output Fields | 19886
- Sample Output | 19888
- Release Information | 19893

Syntax

```
show fabric administration inventory director-group status (target director-device-name | all)
```

Description

(QFabric systems only) Display the status of Director devices that belong to a QFabric system Director group.

Options

none	Display the status of all Director devices within a QFabric system.
all	Display the status of all Director devices within a QFabric system.
target *director-device-name*	Display the status of a specific Director device within a QFabric system.

Required Privilege Level

admin

Output Fields

Table 536 on page 19887 lists the output fields for the `show fabric administration inventory director-group status` command. Output fields are listed in the approximate order in which they appear.

Table 536: show fabric administration inventory director-group status Output Fields

Field Name	Field Description
Director Group Status	Timestamp for the Director group status report.
Member	Name of the Director device.
Status	Current operational mode of the Director device: `online` or `offline`.
Role	High availability operational role of the Director device: `master` or `standby`.
Mgmt Address	Management IP address of the Director device.
CPU	Percentage of CPU processing memory being used by the Director device.
Free Memory	Available storage memory on the Director device.
VMs	Number of virtual machines operating on the Director device. A Routing Engine issue exists on a particular Director device when the system displays a star (`*`) in the `VMs` column along with the following message: `Error in retrieving VM Count in dg`*X*.
Up Time	Length of time the Director device has been operating.
Device Id/Alias	Name or identifier of the Director device.
Master Services	Operational status of the database server, load balancer, and QFabric partition address.
Director Group Managed Services	Operational status of the shared file system, network file system, virtual machine server, and DHCP load balancer.

Table 536: show fabric administration inventory director-group status Output Fields *(Continued)*

Field Name	Field Description
Hard Drive Status	Operational status of the Director device hard drive, including information about the volume identifier, physical identifiers, and SCSI identifiers. There is also status information for the drive partitions, including directory size, available and used drive space, utilization, directory locations, and drive resynchronization progress. A hard drive issue exists when the system displays one of the following messages: `Error in retrieving Hard disk status` or `Error in retrieving Hard disk storage status`.
Director Group Processes	Operational status of the Director group processes, such as device managers, SSH, NFS, FTP, system log messages, and SNMP.
Interface Link Status	Operational status of the Director device interfaces, such as the management interface, the control plane bridge interface, the control plane LAG, the control plane links (where the first number represents the Ethernet module [0 or 1] and the second number represents the port [0 - 3]), and the inter-Director crossover LAG. The state of the interfaces can be `up` or `down`. The `Active` designation indicates that a control plane link is active.

Sample Output

show fabric administration inventory director-group status

```
user@qfabric> show fabric administration inventory director-group status
Director Group Status Thu Aug  2 17:36:34 UTC 2012

 Member Status Role     Mgmt Address    CPU Free Memory VMs Up Time
 ------ ------ -------- --------------- --- ----------- --- -------------
 dg0    online master   10.94.215.38    8%  15191684k   4   6 days, 06:24 hrs
 dg1    online backup   10.94.215.39    7%  17733160k   3   6 days, 06:24 hrs

 Member Device Id/Alias  Status  Role
 ------ ---------------- ------- ---------
 dg0    0281052011000001 online  master

  Master Services
  ---------------
  Database Server                online
  Load Balancer Director         online
```

```
QFabric Partition Address      online

Director Group Managed Services
-------------------------------
Shared File System             online
Network File System            online
Virtual Machine Server         online
Load Balancer/DHCP             online

Hard Drive Status
----------------
Volume ID:0C72369323C3837B     optimal
Physical ID:0                  online
Physical ID:1                  online
Resync Progress Remaining:0    0%
Resync Progress Remaining:1    0%


Size  Used Avail Used% Mounted on
----  ---- ----- ----- ----------
423G 9.4G 391G  3%   /
99M  16M  79M   17%  /boot
93G  11G  83G   12%  /pbdata

Director Group Processes
------------------------
Director Group Manager         online
Partition Manager              online
Software Mirroring             online
Shared File System master      online
Secure Shell Process           online
Network File System            online
DHCP Server master             online     master
FTP Server                     online
Syslog                         online
Distributed Management         online
SNMP Trap Forwarder            online
SNMP Process                   online
Platform Management            online

Interface Link Status
---------------------
Management Interface           up
```

```
 Control Plane Bridge           up
 Control Plane LAG              up
 CP Link [0/2]                  up         Active
 CP Link [0/1]                  up         Active
 CP Link [0/0]                  up         Active
 CP Link [1/2]                  down
 CP Link [1/1]                  down
 CP Link [1/0]                  down
 Crossover LAG                  up
 CP Link [0/3]                  up
 CP Link [1/3]                  up

Member Device Id/Alias  Status  Role
------ ---------------- ------- ---------
dg1    0281052011000032 online  backup

 Director Group Managed Services
 -------------------------------
 Shared File System             online
 Network File System            online
 Virtual Machine Server         online
 Load Balancer/DHCP             online

 Hard Drive Status
 ----------------
 Volume ID:0C9BA4D5F21511B9     optimal
 Physical ID:0                  online
 Physical ID:1                  online
 Resync Progress Remaining:0    0%
 Resync Progress Remaining:1    0%


 Size  Used Avail Used% Mounted on
 ----  ---- ----- ----- ----------
 423G 9.8G 391G  3%   /
 99M  16M  79M   17%  /boot
 93G  11G  83G   12%  /pbdata

 Director Group Processes
 ------------------------
 Director Group Manager         online
 Partition Manager              online
 Software Mirroring             online
```

```
  Shared File System master      online
  Secure Shell Process           online
  Network File System            online
  DHCP Server master             online     backup
  FTP Server                     online
  Syslog                         online
  Distributed Management         online
  SNMP Trap Forwarder            online
  SNMP Process                   offline
  Platform Management            online

  Interface Link Status
  ---------------------
  Management Interface           up
  Control Plane Bridge           up
  Control Plane LAG              up
  CP Link [0/2]                  up         Active
  CP Link [0/1]                  up         Active
  CP Link [0/0]                  up         Active
  CP Link [1/2]                  down
  CP Link [1/1]                  down
  CP Link [1/0]                  down
  Crossover LAG                  up
  CP Link [0/3]                  up
  CP Link [1/3]                  up
```

show fabric administration inventory director-group status target

```
user@qfabric> show fabric administration inventory director-group status target 0281052011000004
Director Group Status Thu Aug 16 02:25:37 UTC 2012

 Member Status Role     Mgmt Address    CPU Free Memory VMs Up Time
 ------ ------ -------- --------------- --- ----------- --- -------------
 dg0    online backup   10.94.195.109   7%  10009364k   3   23:44 hrs
 dg1    online master   10.94.195.110   9%  6120712k    4   1 day, 33:17 mins

 Member Device Id/Alias  Status  Role
 ------ ---------------- ------- ---------
 dg1    0281052011000004 online  master

  Master Services
```

```
---------------
Database Server                online
Load Balancer Director         online
QFabric Partition Address      online

Director Group Managed Services
------------------------------
Shared File System             online
Network File System            online
Virtual Machine Server         online
Load Balancer/DHCP             online

Hard Drive Status
----------------
Physical ID:0                  online
Physical ID:1                  online


Size  Used Avail Used% Mounted on
----  ---- ----- ----- ----------
423G 7.2G 394G  2%   /
99M  20M  75M   21%  /boot
93G  8.8G 85G   10%  /pbdata

Director Group Processes
------------------------
Director Group Manager         online
Partition Manager              online
Software Mirroring             online
Shared File System master      online
Secure Shell Process           online
Network File System            online
DHCP Server master             online     master
FTP Server                     online
Syslog                         online
Distributed Management         online
SNMP Trap Forwarder            online
SNMP Process                   online
Platform Management            online

Interface Link Status
---------------------
Management Interface           up
```

```
Control Plane Bridge          up
Control Plane LAG             up
CP Link [0/2]                 up        Active
CP Link [0/1]                 up        Active
CP Link [0/0]                 up        Active
CP Link [1/2]                 up
CP Link [1/1]                 up
CP Link [1/0]                 up
Crossover LAG                 up
CP Link [0/3]                 up
CP Link [1/3]                 up
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

- request fabric administration director-group change-master | 14616
- *Performing the QFabric System Initial Setup on a QFX3100 Director Group*
- *Understanding the Director Group*

show fabric administration inventory infrastructure

IN THIS SECTION

- Syntax | 19894
- Description | 19894
- Options | 19894
- Required Privilege Level | 19894
- Output Fields | 19894
- Sample Output | 19896
- Release Information | 19897

Syntax

```
show administrator inventory infrastructure
<brief | detail>
(fabric-controls | fabric-managers | diagnostic-routing-engines)
```

Description

(QFabric systems only) Display the services running on the Director group for the QFabric system. These services can include external Routing Engines that are used to support QFabric system operations, such as partitioning and fabric control.

Options

none	Display all services running on the Director group, which are used to support the QFabric system.
brief | detail	(Optional) Display the specified level of output.
fabric-managers	Display information for the fabric manager Routing Engine running on the Director group, which is used to support all partitions in the QFabric system.
fabric-controls	Display information for the fabric control Routing Engine running on the Director group, which is used to support route information in the QFabric system.
diagnostic-routing-engines	Display information for the diagnostic Routing Engine running on the Director group, which is responsible for troubleshooting and diagnostic utilities within a QFabric system partition.

Required Privilege Level

admin

Output Fields

Table 537 on page 19895 lists the output fields for the `show fabric administration inventory infrastructure` command. Output fields are listed in the approximate order in which they appear.

Table 537: show fabric administration inventory infrastructure Output Fields

Field Name	Field Description
Routing Engine Type	Type of virtual Junos Routing Engine being viewed. Examples include the network Node group, fabric control, fabric manager, and diagnostic Routing Engines.
Hostname	Name of the QFabric system component.
PID	Process identifier for the component.
CPU-Use (%)	Percentage of CPU processing memory being used by the component.
Fabric control	Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition. With the `fabric-controls` option, the fabric control Routing Engine can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for this virtual device. This field also displays the identifier and configuration status for the fabric control Routing Engine.
Fabric manager	Name of the virtual Junos Routing Engine that manages the QFabric system. With the `fabric-managers` option, the fabric manager Routing Engine can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for this virtual device. This field also displays the identifier and configuration status for the fabric manager Routing Engine.
Network Node group	Name of the virtual Junos Routing Engine instance that handles routing processes for a network Node group.
Diagnostic	Name of the virtual Junos Routing Engine responsible for troubleshooting and diagnostic utilities within a QFabric system partition. With the `diagnostic-routing-engines` option, the diagnostic Routing Engine can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the diagnostic Routing Engine.
Item	Type of QFabric system component being viewed.

Table 537: show fabric administration inventory infrastructure Output Fields *(Continued)*

Field Name	Field Description
Identifier	Hardware serial identifier of a QFabric system component.
Connection	Status of a QFabric system component: either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for the listed component.
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be `Configured`, `Failed` (unsuccessful), `Pending` (in the process of being written or retried), or `Unknown`.

Sample Output

show fabric administration inventory infrastructure

```
user@qfabric> show fabric administration inventory infrastructure
dg0:
Routing Engine Type     Hostname                      PID      CPU-Use(%)
---------------------------------------------------------------------------
Fabric manager          FM-0                          9832      1.0
Network Node group      QFabric_default_NW-NG-1_RE1   24633     4.2
Fabric control          QFabric_default_FC-1_RE0      25374     1.8
Diagnostic              QFabric_DRE                   6789      1.3

dg1:
Routing Engine Type     Hostname                      PID      CPU-Use(%)
---------------------------------------------------------------------------
Fabric manager          FM-1                          572       1.6
Network Node group      QFabric_default_NW-NG-0_RE0   19217     7.8
Fabric control          QFabric_default_FC-0_RE0      20071     1.9
```

show fabric administration inventory infrastructure fabric-controls

```
user@qfabric> show fabric administration inventory infrastructure fabric-controls fabric-controls
Item                    Identifier              Connection      Configuration
Fabric control
  FC-0                                          Connected       Configured
  FC-1                                          Connected       Configured
```

show fabric administration inventory infrastructure fabric-managers

```
user@qfabric> show fabric administration inventory infrastructure fabric-managers
Item                    Identifier              Connection      Configuration
Fabric manager
  FM-0                                          Connected       Configured
```

show fabric administration inventory infrastructure diagnostic-routing-engines

```
user@qfabric> show fabric administration inventory infrastructure diagnostic-routing-engines
Item                    Identifier              Connection      Configuration
Diagnostic routing engine
  DRE-0                                         Connected       Configured
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

show fabric administration inventory | 19878

show fabric administration inventory interconnect-devices

IN THIS SECTION

- Syntax | 19898
- Description | 19898
- Options | 19898
- Required Privilege Level | 19898
- Output Fields | 19898
- Sample Output | 19899
- Release Information | 19900

Syntax

```
show fabric administration inventory interconnect-devices interconnect-device-name
```

Description

(QFabric systems only) Display the Interconnect devices that belong to a QFabric system.

Options

none	Display all Interconnect devices within a QFabric system.
interconnect-device-name	Display a specific Interconnect device within a QFabric system.

Required Privilege Level

admin

Output Fields

Table 538 on page 19899 lists the output fields for the `show fabric administration inventory interconnect-devices` command. Output fields are listed in the approximate order in which they appear.

Table 538: show fabric administration inventory interconnect-devices Output Fields

Field Name	Field Description
Interconnect device	Name of the Interconnect devices associated with the partition.
Item	Type of QFabric system component being viewed. Interconnect devices either display the alias name (if configured) or the hardware serial identifier and Control Board Routing Engine numbers.
Identifier	Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.
Connection	Status of a QFabric system component: either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for the listed component.
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be `Configured`, `Failed` (unsuccessful), `Pending` (in the process of being written or retried), or `Unknown`.

Sample Output

show fabric administration inventory interconnect-devices

```
user@qfabric> show fabric administration inventory interconnect-devices
Item                    Identifier              Connection      Configuration
Interconnect device
  IC-YW3781                                     Connected       Configured
    YW3781/RE0                                  Connected
    YW3781/RE1                                  Connected
  IC-YW3798                                     Connected       Configured
    YW3798/RE0                                  Connected
    YW3798/RE1                                  Connected
```

show fabric administration inventory interconnect-devices device-name

```
user@qfabric> show fabric administration inventory interconnect-devices IC-YW3781
Item                    Identifier              Connection      Configuration
Interconnect device
  IC-YW3781                                     Connected       Configured
    YW3781/RE0                                  Connected
    YW3781/RE1                                  Connected
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

- request chassis fabric fpc | 14465
- show chassis fabric connectivity | 18022
- show chassis fabric device | 18066
- show fabric administration inventory | 19878
- *Understanding Interconnect Devices*

show fabric administration inventory node-devices

IN THIS SECTION

- Syntax | 19901
- Description | 19901
- Options | 19901
- Required Privilege Level | 19901
- Output Fields | 19901
- Sample Output | 19902
- Release Information | 19902

Syntax

```
show fabric administration inventory node-devices node-device-name
```

Description

(QFabric systems only) Display the Node devices that belong to the QFabric system.

> **NOTE**: If your Node devices do not appear in the output of the `show fabric administration inventory node-devices` command, check the cabling of your system.

Options

none	Display all Node devices within the QFabric system.
node-device-name	Display a specific Node device within the QFabric system.

Required Privilege Level

admin

Output Fields

Table 539 on page 19901 lists the output fields for the `show fabric administration inventory node-devices` command. Output fields are listed in the approximate order in which they appear.

Table 539: show fabric administration inventory node-devices Output Fields

Field Name	Field Description
Item	Type of QFabric system component being viewed.
Identifier	Hardware serial identifier of a QFabric system component.

Table 539: show fabric administration inventory node-devices Output Fields *(Continued)*

Field Name	Field Description
Connection	Status of a QFabric system component: either `Connected` or `Not Connected`, depending on whether or not the Director software has detected keepalive messages for the listed component.
Node device	Name of the Node devices associated with the Node group. The device can be either `Connected` or `Not Connected`, depending on whether or not the Director software has detected keepalive messages for the device. This field also displays the serial ID and configuration status for the Node device.

Sample Output

show fabric administration inventory node-devices

```
user@qfabric> show fabric administration inventory node-devices
Item                    Identifier              Connection
Node device
  node1                 P3749-C                 Connected
  node2                 P3767-C                 Connected
  node3                 P3850-C                 Connected
  node4                 P3947-C                 Connected
```

show fabric administration inventory node-devices device-name

```
user@qfabric> show fabric administration inventory node-devices node0
Item                    Identifier              Connection
Node device
  node1                 P3749-C                 Connected
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Aliases for the QFabric System

Configuring Node Groups for the QFabric System

show fabric administration inventory node-groups | **19903**

show fabric administration inventory | **19878**

Understanding Node Devices

show fabric administration inventory node-groups

IN THIS SECTION

- Syntax | **19903**
- Description | **19903**
- Options | **19903**
- Required Privilege Level | **19904**
- Output Fields | **19904**
- Sample Output | **19905**
- Release Information | **19905**

Syntax

```
show fabric administration inventory node-groups node-group-name
```

Description

(QFabric systems only) Display the Node groups and the corresponding Node devices that belong to the QFabric system.

Options

none Display all Node groups within the QFabric system.

node-group-name Display information for a specific Node group within the QFabric system.

Required Privilege Level

admin

Output Fields

Table 540 on page 19904 lists the output fields for the `show fabric administration inventory node-groups` command. Output fields are listed in the approximate order in which they appear.

Table 540: show fabric administration inventory node-groups Output Fields

Field Name	Field Description
Node group	Name of the Node groups associated with the partition.
Item	Type of QFabric system component being viewed. • Autogenerated Node groups display the hardware serial identifier for both the name of the Node group and the name of the included Node device. • User-configured Node groups either display the alias name (if configured) or the hardware serial identifier for each Node device contained in the Node group.
Identifier	Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.
Connection	Status of a QFabric system component: either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for the listed component.
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be `Configured`, `Failed` (unsuccessful), `Pending` (in the process of being written or retried), or `Unknown`.

Sample Output

show fabric administration inventory node-groups

```
user@qfabric> show fabric administration inventory node-groups
Item                    Identifier              Connection      Configuration
Node group
  BBAK8891                                      Connected       Configured
    BBAK8891                                    Connected
  BBAK8868                                      Connected       Configured
    BBAK8868                                    Connected
  RSNG-1                                        Connected       Configured
    Node-3              BBAK8276                Connected
    Node-4              BBAK8273                Connected
  NW-NG-0                                       Connected       Configured
    Node-0              BBAK8309                Connected
    Node-1              BBAK8283                Connected
```

show fabric administration inventory node-groups node-group-name

```
user@qfabric> show fabric administration inventory node-groups RSNG-1
Item                    Identifier              Connection      Configuration
Node group
  RSNG-1                                        Connected       Configured
    Node-3              BBAK8276                Connected
    Node-4              BBAK8273                Connected
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Configuring Node Groups for the QFabric System

show fabric administration inventory node-devices | **19900**

show fabric administration inventory | **19878**

show fabric administration system mac-pool

IN THIS SECTION

- Syntax | 19906
- Description | 19906
- Options | 19906
- Required Privilege Level | 19906
- Output Fields | 19906
- Sample Output | 19907
- Release Information | 19907

Syntax

```
show fabric administration system mac-pool
```

Description

(QFabric systems only) Display the MAC addresses that belong to a QFabric Director group.

Options

There are no options for this command.

Required Privilege Level

admin

Output Fields

Table 541 on page 19907 lists the output fields for the `show fabric administration system mac-pool` command. Output fields are listed in the approximate order in which they appear.

Table 541: show fabric administration system mac-pool Output Fields

Field Name	Field Description
MAC Block Base	Starting MAC address for the pool assigned to the QFabric system.
Total MACs	Total number of MAC addresses assigned to the QFabric system.
Available MACs	Number of available MAC addresses from the total.

Sample Output

show fabric administration system mac-pool

```
user@qfabric> show fabric administration system mac-pool
 Mac Block Base         Total MACs      Available MACs
         00:11:00:00:00:00       4096            4084
         02:00:00:11:22:00       10              10
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

request fabric administration system mac-pool add | 14626

request fabric administration system mac-pool delete | 14628

show fabric inventory

IN THIS SECTION

- Syntax | 19908
- Description | 19908
- Options | 19909
- Required Privilege Level | 19909
- Output Fields | 19909
- Sample Output | 19910
- Release Information | 19912

Syntax

```
show fabric inventory
<brief | detail | terse>
<infrastructure fabric-controls <FC-0 | FC-1>>
<node-devices node-device-name>
<node-groups node-group-name>
```

Description

(QFabric systems only) Display Node devices, Node groups, and fabric control Routing Engines that belong to the QFabric system. You can narrow the level of output by specifying a device type.

NOTE:

- If you have administrator privileges, issue the `show fabric administration inventory` command to view all devices in your QFabric system (including Interconnect devices and Director devices).
- If your Node devices do not appear in the output of the `show fabric inventory` command, check the cabling of your system.

Options

none	Display all devices within a QFabric system.
brief | detail | terse	(Optional) Display the specified level of output.
infrastructure fabric-controls <FC-0 | FC-1>	(Optional) Display information for all fabric control Routing Engines running on the Director group within the QFabric system, or the individual fabric control Routing Engine you specify (either FC-0 or FC-1).
node-devices ***node-device-name***	(Optional) Display a specific Node device within a QFabric system.
node-groups ***node-group-name***	(Optional) Display a specific Node group within a QFabric system.

Required Privilege Level

admin

Output Fields

Table 542 on page 19909 lists the output fields for the `show fabric inventory` command. Output fields are listed in the approximate order in which they appear.

Table 542: show fabric inventory Output Fields

Field Name	Field Description
Item	Type of QFabric system component being viewed. Possible values include `Node device`, `Node group`, `Fabric control`, and `Ungrouped Node device`.
Identifier	Hardware serial identifier of a QFabric system component.
Connection	Status of a QFabric system component: either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for the listed component.

Table 542: show fabric inventory Output Fields *(Continued)*

Field Name	Field Description
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be `Configured`, `Failed` (unsuccessful), `Pending` (in the process of being written or retried), or `Unknown`.
Node group	Name of the Node groups associated with the QFabric system, and the Node devices assigned to each Node group. The group can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for the devices in the group. This field also displays the serial ID for the Node group and the status for the Node group.
Node device	Name of the Node devices associated with the Node group. The device can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for the device. This field also displays the serial ID and configuration status for the Node device.
Fabric control	Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition. The fabric control Routing Engine can be either `Connected` or `Disconnected`, depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric control Routing Engine.

Sample Output

show fabric inventory

```
user@qfabric> show fabric inventory

Item                    Identifier              Connection      Configuration
Ungrouped Node device
  P3747-C                                       Disconnected
Node group
  NW-NG-0                                       Connected       Configured
    Node-1              P4093-C                 Connected
  RSNG-1                                        Connected       Configured
    Node-2              P4514-C                 Connected
```

```
    Node-3              P3917-C                 Connected
Fabric control
  FC-0                                          Connected       Configured
  FC-1                                          Connected       Configured
```

show fabric inventory infrastructure fabric-controls

```
user@qfabric> show fabric inventory infrastructure fabric-controls

Item                    Identifier              Connection      Configuration
Fabric control
  FC-0                                          Connected       Configured
  FC-1                                          Connected       Configured
```

show fabric inventory node-devices

```
user@qfabric> show fabric inventory node-devices

Item                    Identifier              Connection      Configuration
Node device
  Node-1                P4093-C                 Connected
  Node-2                P4514-C                 Connected
  Node-3                P3917-C                 Connected
```

show fabric inventory node-groups

```
user@qfabric> show fabric inventory node-groups

Item                    Identifier              Connection      Configuration
Node group
  NW-NG-0                                       Connected       Configured
    Node-1              P4093-C                 Connected
  RSNG-1                                        Connected       Configured
    Node-2              P4514-C                 Connected
    Node-3              P3917-C                 Connected
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Performing the QFabric System Initial Setup on a QFX3100 Director Group

Configuring Aliases for the QFabric System

Configuring Node Groups for the QFabric System

show fabric administration inventory | **19878**

show fabric session-host

IN THIS SECTION

- Syntax | **19912**
- Description | **19912**
- Options | **19913**
- Required Privilege Level | **19913**
- Output Fields | **19913**
- Sample Output | **19913**
- Release Information | **19913**

Syntax

```
show fabric session-host
```

Description

(QFabric systems only) Display the Director device within the Director group that hosts the QFabric CLI session.

Options

none Display the Director device hosting the QFabric CLI session.

Required Privilege Level

admin

Output Fields

Table 543 on page 19913 lists the output fields for the `show fabric session-host` command. Output fields are listed in the approximate order in which they appear.

Table 543: show fabric session-host Output Fields

Field Name	Field Description
Identifier	Hardware serial identifier of the Director device that hosts the SSH QFabric CLI session.

Sample Output

show fabric session-host

```
user@qfabric> show fabric session-host
Identifier: 0281052011000032
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Understanding the Director Group

show fabric administration inventory director-group status | 19885

show fib-local-accounting ip

IN THIS SECTION

- Syntax | 19914
- Description | 19914
- Required Privilege Level | 19914
- Sample Output | 19914
- Release Information | 19915

Syntax

```
show fib-local-accounting ip
```

Description

Display the number of packets that were sent to an anchor MPC due to FIB localization.

Required Privilege Level

view

Sample Output

show fib-local-accounting ip

```
user@host> show fib-local-accounting ip
PFE 0
         fe_addr         packets         bytes
         28                      0                     0
         29                      0                     0
         30                      0                     0
         31                      0                     0
PFE 1
```

```
        fe_addr         packets         bytes
        28                      0                       0
        29                      0                       0
        30                      0                       0
        31                      0                       0
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

fib-remote | 7009

fib-local | 7005

Example: Configuring Packet Forwarding Engine FIB Localization

show fib-streaming

IN THIS SECTION

- Syntax | 19916
- Description | 19916
- Options | 19916
- Required Privilege Level | 19916
- Output Fields | 19916
- Sample Output | 19918
- Release Information | 19921

Syntax

```
show fib-streaming (next-hop-groups | next-hops)
<all>
<index>
```

Description

Show information about Forwarding Information Base (FIB) streaming. You can use Junos telemetry interface (JTI) and remote procedure calls (gRPC) services to stream FIB operational state statistics to outside collectors.

Options

next-hops	Show information about next hop entries present in the routing table.
next-hop-groups	Show information about next hop group entries present in the routing table.
all	Show all next hops or next hop groups.
index	Show only the specified next hop or next hop group.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show fib-streaming` command.

Table 544: show-fib-streaming Output Fields

Field Name	Field Description
`next-hop`	For each next hop, shows the following information: • `index`—Interface index. • `lspid`—ID for the label switched path used to reach the next hop. • `lspname`—Name of the label switched path used to reach the next hop. • `encap-header`— The type of encapsulation performed when forwarding a packet to the specified next hop. • `programmed-index`—The index set by configuration for the next hop. • `ip-address`—IP address. • `mac-address`—MAC address. • `pushed-mpls-label-stack`—The MPLS label stack imposed when forwarding packets to the ne The order of entries is such that the first entry in the list is the label at the bottom of the sta be pushed. • `interface`—Interface name.

Table 544: show-fib-streaming Output Fields *(Continued)*

Field Name	Field Description
`next-hop-group`	For each next hop group, shows the following information: • `next-hop-group - id`—ID for the next hop group. • `programmed-id`—The ID set by configuration for each next hop. • `next-hop index`—ID for the next hop. • `weight`—Weight applied to the next hop within the group. If there are no indexed next hops, such as with CoS-based forwarding, the output will display th following fields: • `next-hop-group - id`—ID for the next hop group. • `condition-id`—Unique ID under the next-hop-group for the condition. • `dscp`—DSCP value. • `id`—IFL index of the interface. • `interface`—Name of the interface. • `subinterface`—Sub-unit of the interface. If a valid back up path is configured, the `backup-next-hop-group - id` is displayed.

Sample Output

show fib-streaming next-hops all

```
user@host> show fib-streaming next-hops all
next-hop index : 1
                lspid            : 0
                encap-header     : 5
                programmed-index : 0
                interface        : .local..1
next-hop index : 21
                lspid            : 0
                encap-header     : 5
```

```
                programmed-index : 0
                interface        : em1.0
next-hop index : 114
                lspid            : 0
                encap-header     : 1
                programmed-index : 0
                interface        : ip-0/0/10.0
```

show fib-streaming next-hops index 79

```
user@host> show fib-streaming next-hops index 79
next-hop index : 79
        lspid            : 5
        lspname          : HL4-01_TO_HL3-01
        programmed-index : 0
        ip-address       : 1.0.0.2
        mac-address      : 2c:6b:f5:78:ef:00
        interface        : ge-0/0/1.0
```

show fib-streaming next-hop-groups all

```
user@host> show fib-streaming next-hop-groups all
next-hop-group - id: 1
                        programmed-id: 0
   next-hop   index: 1          weight: 10000
next-hop-group - id: 2
                        programmed-id: 0
   next-hop   index: 2          weight: 10000
next-hop-group - id: 3
                        programmed-id: 0
   next-hop   index: 3          weight: 10000
next-hop-group - id: 4
                        programmed-id: 0
   next-hop   index: 4          weight: 10000
next-hop-group - id: 5
                        programmed-id: 0
   next-hop   index: 5          weight: 10000
```

show fib-streaming next-hop-groups index 105

```
user@host> show fib-streaming next-hop-groups index 105
next-hop-group - id: 105
   next-hop   index: 105        weight: 2500
   next-hop   index: 106        weight: 2500
   next-hop   index: 107        weight: 2500
   next-hop   index: 108        weight: 2500
```

```
user@host> show fib-streaming next-hop-groups index 110
next-hop-group - id: 110
     id: 358         interface:  ge-0/0/0           subinterface: 0
     id: 359         interface:  ge-0/0/0           subinterface: 1
     id: 360         interface:  ge-0/0/0           subinterface: 2
     id: 361         interface:  ge-0/0/0           subinterface: 3
     id: 362         interface:  ge-0/0/0           subinterface: 4
     id: 363         interface:  ge-0/0/0           subinterface: 5
     id: 364         interface:  ge-0/0/0           subinterface: 6
     id: 365         interface:  ge-0/0/0           subinterface: 7
     id: 366         interface:  ge-0/0/0           subinterface: 8
     id: 367         interface:  ge-0/0/0           subinterface: 9
     id: 368         interface:  ge-0/0/0           subinterface: 10
     id: 369         interface:  ge-0/0/0           subinterface: 11
     id: 370         interface:  ge-0/0/0           subinterface: 12

   condition id : 67305485
                     dscp :  48
                     dscp :  49
                     dscp :  50
                     dscp :  51
                     dscp :  52
                     dscp :  53
                     dscp :  54
                     dscp :  55
                     dscp :  56
                     dscp :  57
                     dscp :  58
                     dscp :  59
                     dscp :  60
                     dscp :  61
```

```
                    dscp :  62
                    dscp :  63
```

show fib-streaming next-hop-groups index 756 debug

```
next-hop-group - id: 756
                        programmed-id: 0
                        flags: 0x10         fc_bitmask: 0x0
   next-hop   index: 834        weight: 50.00, Label-Count: 0       , Interface: ge-1/0/0.0
                                         IP-addr: 11.1.1.8
                                         IP-addr: 11.1.1.8
                fc_bitmask: 0x0                     LSP-ID:  0
   next-hop   index: 835        weight: 50.00, Label-Count: 0       , Interface: ge-1/0/0.0
                                         IP-addr: 11.1.1.9
                                         IP-addr: 11.1.1.9
                fc_bitmask: 0x0                     LSP-ID:  0
backup-next-hop-group - id: 757
```

Release Information

Command introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

show fib-streaming routes | 19924

show fib-streaming route-tables

IN THIS SECTION

- Syntax | 19922
- Description | 19922
- Options | 19922
- Required Privilege Level | 19922

- Output Fields | 19922
- Sample Output | 19923
- Release Information | 19924

Syntax

```
show fib-streaming route-tables
<all>
<summary>
```

Description

Show information about Forwarding Information Base (FIB) streaming routing tables. You can use Junos telemetry interface (JTI) and remote procedure calls (gRPC) services to stream FIB operational state statistics to outside collectors.

Options

all	Show all routing tables.
summary	Show routing tables and a summary of the route entries in each table.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show fib-streaming route-tables` command.

Table 545: show fib-streaming route-tables Output Fields

Field Name	Field Description
`Table-ID`	Routing table ID.
`Table name`	Name of the routing table.
`count`	Number of route entries in the table by type: • `v4`—IPv4 unicast • `v6`—IPv6 unicast • `mpls`—MPLS

Sample Output

show fib-streaming route-tables summary

```
user@host> show fib-streaming route-tables summary
Table-ID   Table name
0          DEFAULT       v4 count:52  v6 count:9   mpls count:38
1          __juniper_1__ v4 count:21  v6 count:19  mpls count:0
2          __juniper_2__ v4 count:6   v6 count:0   mpls count:0
3          __juniper_3__ v4 count:10  v6 count:0   mpls count:0
36736      __juniper_4__ v4 count:10  v6 count:0   mpls count:0
```

show fib-streaming route-tables all

```
user@host> show fib-streaming route-tables all
Table-ID   Table name
0          DEFAULT
1          __juniper_1__
2          __juniper_2__
3          __juniper_3__
36736      __juniper_4__
```

Release Information

Command introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

show fib-streaming | 19915

show fib-streaming routes

IN THIS SECTION

- Syntax | 19924
- Description | 19924
- Options | 19925
- Required Privilege Level | 19925
- Output Fields | 19925
- Sample Output | 19926
- Release Information | 19927

Syntax

```
show fib-streaming routes (ipv4-unicast | ipv6-unicast | mpls)
<all>
<prefix>
<label>
<table name>
```

Description

Show information about Forwarding Information Base (FIB) streaming routes. You can use Junos telemetry interface (JTI) and remote procedure calls (gRPC) services to stream FIB operational state statistics to outside collectors.

Options

ipv4-unicast	Show IPv4 unicast route entries present in the routing table.
ipv6-unicast	Show IPv6 unicast route entries present in the routing table.
mpls	Show MPLS route entries present in the routing table.
all	Show all routes.
prefix	Prefix of the route.
label	For MPLS routes only, the label used to represent the route.
table *name*	Name of the routing table.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show fib-streaming routes` command.

Table 546: show fib-streaming routes Output Fields

Field Name	Field Description
`prefix`	Prefix of the route.
`next-hop-group`	ID of the next hop group.
`table`	Name of the routing table.
`label`	MPLS label.

Sample Output

show fib-streaming routes ipv4-unicast table DEFAULT

```
user@host> show fib-streaming routes ipv4-unicast table DEFAULT
prefix                next-hop-group     table
0.0.0.0/0             48                 DEFAULT
224.0.0.0/4           49                 DEFAULT
10.0.0.0/8            50                 DEFAULT
172.16.0.0/12         50                 DEFAULT
66.129.0.0/16         50                 DEFAULT
192.168.0.0/16        50                 DEFAULT
```

show fib-streaming routes ipv4-unicast prefix 100.0.0.1

```
user@host> show fib-streaming routes ipv4-unicast prefix 100.0.0.1
prefix                next-hop-group     table
100.0.0.1/32           75                 DEFAULT
```

show fib-streaming routes mpls table DEFAULT

```
user@host> show fib-streaming routes mpls table DEFAULT
Label          next-hop-group      table
0              11                  DEFAULT
1              12                  DEFAULT
2              13                  DEFAULT
13             12                  DEFAULT
16             14                  DEFAULT
17             15                  DEFAULT
18             16                  DEFAULT
19             17                  DEFAULT
20             18                  DEFAULT
21             19                  DEFAULT
22             20                  DEFAULT
```

show fib-streaming routes mpls label 50

```
user@host> show fib-streaming routes mpls label 50
Label        next-hop-group    table
50           47                DEFAULT
```

Release Information

Command introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

show fib-streaming | 19915

show fib-streaming state

IN THIS SECTION

- Syntax | 19927
- Description | 19928
- Required Privilege Level | 19928
- Output Fields | 19928
- Sample Output | 19929
- Release Information | 19929

Syntax

```
show fib-streaming state
```

Description

Displays the forwarding information base (FIB) telemetry streaming state.

Required Privilege Level

view

Output Fields

Table 547 on page 19928 lists the output fields for the `show fib-streaming state` command.

Table 547: show fib-streaming state Output Fields

Field Name	Field Description
FIB states resync	Number of times the forwarding information base (FIB) has synchronized.
RPD client session state	Session state for the routing protocol process (daemon) (RPD) client.
LSP name init sync done	The INIT_SYNC status.
Resync done	Indicates if a resynchronization is complete.
Operation mode	Operation mode. If prefix-filtering is enabled, the mode is `prefix-filter`. If prefix-filtering is not configured, mode is `non-prefix-filter`.

Sample Output

show fib-streaming state

```
user@host> show fib-streaming state
```

```
Jun 03 22:45:43
            FIB states resync: 1
            RPD client session state: Connected
            LSP name init sync done: 1
            Resync done: 1
            Operation mode: non-prefix-filter
```

Release Information

Command introduced in Junos OS Release 24.2R1 for MX2020 routers.

RELATED DOCUMENTATION

prefix-list (Junos Telemetry Interface) | 13093

Configuring IP-AFT Prefix Filtering

show fibre-channel fabric

IN THIS SECTION

- Syntax | 19930
- Description | 19930
- Options | 19930
- Required Privilege Level | 19930
- Output Fields | 19930
- Sample Output | 19931

Release Information | 19932

Syntax

```
show fibre-channel fabric
<extensive | summary>
<fabric-name>
<sort-by (name | fabric-id)>
```

Description

Display Fibre Channel fabric information.

Options

fabric-name	(Optional) Display output only for the specified fabric.
extensive | summary	(Optional) Display the specified level of output.
sort-by (name | fabric-id)	(Optional) Sort output by fabric name or fabric ID.

Required Privilege Level

view

Output Fields

Table 548 on page 19930 lists the output fields for the `show fibre-channel fabric` command. Output fields are listed in the approximate order in which they appear.

Table 548: show fibre-channel fabric Output Fields

Field Name	Field Description	Level of Output
Fabric	Name of the fabric.	All

Table 548: show fibre-channel fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Fabric-ID	Identification number of the fabric.	All
Type	Type of fabric. All fabrics are PROXY fabrics.	All
Interfaces	Native Fibre Channel interfaces and FCoE interfaces assigned to the fabric.	All
Created at	Date and time the fabric was created.	extensive
Internal Index	Fabric index internal to Junos OS.	extensive
Origin	Origin information internal to Junos OS.	extensive
Description	Text description of the fabric.	extensive
Fabric WWN	Unique WWN of the fabric generated by the FCF.	extensive
Login sessions	Number of FIP login sessions currently running on the fabric.	extensive
Configured max login sessions	Configured maximum number of FIP login sessions permitted on the fabric.	extensive

Sample Output

show fibre-channel fabric

```
user@switch> show fibre-channel fabric
Fabric         Fabric-ID     Type     Interfaces

proxy2         200           PROXY
```

```
                                        fc-0/0/0.0
                                        fc-0/0/1.0
```

show fibre-channel fabric extensive

```
user@switch> show fibre-channel fabric extensive
Fabric: proxy2, Created at: Mon Apr 19 14:02:58 2010
Fabric-ID: 200, Internal index: 2, Origin: Static
Description: srv-fabric,  Type: PROXY,  Fabric WWN: 10:00:00:05:33:51:d7:cd
Login sessions: 200,  Configured max login sessions: 500
      fc-0/0/0.0, (untagged)
      fc-0/0/1.0, (untagged)
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

fc-fabrics | 6954

Configuring an FCoE-FC Gateway Fibre Channel Fabric

show fibre-channel fc2 sessions

IN THIS SECTION

- Syntax | 19933
- Description | 19933
- Options | 19933
- Required Privilege Level | 19933
- Output Fields | 19933
- Sample Output | 19934
- Release Information | 19935

Syntax

```
show fibre-channel fc2 sessions
<fabric fabric-name>
<brief | detail>
```

Description

Display Fibre Channel FC-2 information.

> **NOTE**: A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions.

Options

fabric *fabric-name*	(Optional) Display output only for the specified fabric.
brief | detail	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 549 on page 19933 lists the output fields for the `show fibre-channel fc2 sessions` command. Output fields are listed in the approximate order in which they appear.

Table 549: show fibre-channel fc2 sessions Output Fields

Field Name	Field Description	Level of Output
`Fabric`	Name of the fabric.	All
`Fabric-id`	Identification number of the fabric.	All

Table 549: show fibre-channel fc2 sessions Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface Name`	Name of the interface.	All
`Local FCID`	Address of the local end of the connection.	All
`Far FCID`	Address of the far (remote) end of the connection.	All
`# Pending Exchanges`	Number of pending exchanges for the session.	All
`Flags`	Flags internal to Junos OS.	`detail`
`Refcount`	Reference count internal to Junos OS.	`detail`
`Users`	Information internal to Junos OS.	`detail`

Sample Output

show fibre-channel fc2 sessions

```
user@switch> show fibre-channel fc2 sessions
Fabric: fip-proxy, Fabric-id: 1
Interface      Local     Far         # Pending
Name           FCID      FCID        Exchanges
fc-0/0/0.0     *         0xfffffe    0
fc-0/0/1.0     *         0xfffffe    0
fc-0/0/2.0     *         0xfffffe    0
```

show fibre-channel fc2 sessions detail

```
user@switch> show fibre-channel fc2 sessions detail
Fabric: fip-proxy, Fabric-id: 1
Interface Name      fc-0/0/0.0
Local FCID:         *
Far FCID:           0xfffffe
Exchanges:          0
Flags:              SELF_LOCK USER_SYNCHED
Refcount:           2
Users:              1

Interface Name      fc-0/0/1.0
Local FCID:         *
Far FCID:           0xfffffe
Exchanges:          0
Flags:              SELF_LOCK USER_SYNCHED
Refcount:           2

Interface Name      fc-0/0/2.0
Local FCID:         *
Far FCID:           0xfffffe
Exchanges:          0
Flags:              SELF_LOCK USER_SYNCHED
Refcount:           2
Users:              1
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fc2 statistics | 19936

clear fibre-channel fc2 statistics | 3596

show fibre-channel fc2 statistics

IN THIS SECTION

- Syntax | 19936
- Description | 19936
- Options | 19936
- Required Privilege Level | 19936
- Output Fields | 19936
- Sample Output | 19938
- Release Information | 19938

Syntax

```
show fibre-channel fc2 statistics
<fabric fabric-name>
```

Description

Display Fibre Channel FC-2 statistics.

Options

fabric ***fabric-name***	(Optional) Display output only for the specified fabric.

Required Privilege Level

view

Output Fields

Table 550 on page 19937 lists the output fields for the `show fibre-channel fc2 statistics` command. Output fields are listed in the approximate order in which they appear.

Table 550: show fibre-channel fc2 statistics Output Fields

Field Name	Field Description
`Global statistics`	Statistics for all fabrics.
`Frame buffers allocated`	Number of frame buffers currently allocated to all fabrics.
`Frame buffers freed`	Number of frame buffers freed.
`Frames dropped`	Number of dropped frames.
`Fabric statistics`	Fabric-specific statistics.
`Fabric`	Name of the fabric.
`Fabric-id`	Identification number of the fabric.
`Tx-FRJTs`	Number of fabric frame rejects (F_RJTs).
`Tx-PRJTs`	Number of port frame rejects (P_RJTs).
`Tx-LSRJTs`	Number of link service rejections.
`Tx-ABTS`	Number of terminate sequence frames sent.
`Rx-Drops`	Number of received frames dropped.
`Rx-ABTS`	Number of terminate sequence frames received.

Sample Output

show fibre-channel fc2 statistics

```
user@switch> show fibre-channel fc2 statistics
Global statistics:

Frame buffers allocated: 60
Frame buffers freed:     60
Frames dropped:          0

Fabric statistics:

Fabric : fip-proxy, Fabric-id: 1
Tx-FRJTs:    0
Tx-PRJTs:    0
Tx-LSRJTs:   0
Tx-ABTS:     0
Rx-Drops:    0
Rx-ABTS:     0
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fc2 sessions | 19932

clear fibre-channel fc2 statistics | 3596

show fibre-channel fip

IN THIS SECTION

Syntax | 19939

- Description | 19939
- Options | 19939
- Required Privilege Level | 19939
- Output Fields | 19939
- Sample Output | 19943
- Release Information | 19945

Syntax

```
show fibre-channel fip
<brief | detail>
```

Description

Display Fibre Channel over Ethernet Initialization Protocol (FIP) information.

Options

brief | detail	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 551 on page 19940 lists the output fields for the `show fibre-channel fip` command. Output fields are listed in the approximate order in which they appear. A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions.

Table 551: show fibre-channel fip Output Fields

Field Name	Field Description	Level of Output
Configured max FIP sessions per Node Device	Configured maximum number of FIP sessions permitted on the Node device. For QFabric systems, this is the maximum number of FIP sessions permitted on each Node device in the fabric. For QFX3500 devices, this is the maximum number of FIP sessions permitted on the device.	`detail`
`Node Device`	Node device identifier.	`detail`
`Total FIP sessions`	Total number of FIP sessions on the FCoE-FC gateway switch.	`detail`
`Total FCoE filters`	Total number of FIP filters on the FCoE-FC gateway switch.	`detail`
`Fabric Name`	Name of the fabric and in parentheses the fabric ID.	All
`FC-MAP`	FCoE mapped address prefix of the FCoE forwarder for the fabric.	`detail`
`FKA-ADV-PERIOD`	Period of time in milliseconds between FIP keepalive advertisements configured for the FC fabric.	`detail`
`MAX-SESSIONS-PER-ENODE`	Maximum number of concurrent sessions (FLOGI and FDISC combined) that each ENode can instantiate.	`detail`

Table 551: show fibre-channel fip Output Fields *(Continued)*

Field Name		Field Description	Level of Output
FCoE trusted		Whether ports on the FC fabric are trusted or untrusted: • Yes—Ports on the FC fabric are trusted; FIP snooping is turned off. • No—Ports on the FC fabric are not trusted; FIP snooping is turned on.	detail
Member		Information about an FCF that is a member of the fabric.	All
	• FCF-MAC	MAC address used in discovery advertisements.	All
	• FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC interface.	detail
	• FKA-ADV-D-BIT	Disable FIP keepalive advertisement monitoring bit. The state is always off.	detail
	• Type	Type of interface: • VF_Port Capable—Interface can act as a VF_Port interface.	detail
	• Priority	Priority value associated with the switch FCF-MAC. Converged network adapters (CNAs) use the priority value to determine the switch with which they will perform FIP FLOGI. The lower the value, the higher the priority. Value range: 0 through 255.	detail

Table 551: show fibre-channel fip Output Fields *(Continued)*

Field Name		Field Description	Level of Output
	• State	FIP state on the fabric: • Enable—FIP is enabled on the fabric. • Disable—FIP is disabled on the fabric.	detail
ENode		Information about a connected FCoE node (ENode).	All
	• ENode-MAC	MAC address of the connected ENode.	All
	• Enode State	Login state internal to Junos OS.	All
	• Configured ENode timer	User-configured FIP keepalive advertisement interval in milliseconds.	detail
	• Running ENode timer	Runtime interval in milliseconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	detail
	• Active FIP Sessions	Number of active FIP sessions on the ENode.	detail
	• VN-Port-MAC	MAC address of a VN_Port on the ENode.	All
	• Session State	Session state internal to Junos OS.	detail
	• Configured FKA-ADV	User-configured FIP keepalive advertisement interval in milliseconds.	detail
	• Running FKA-ADV	Runtime interval in milliseconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	detail

Table 551: show fibre-channel fip Output Fields *(Continued)*

Field Name		Field Description	Level of Output
	• Configured VN-Port Timer	Configured state of the VN_Port keepalive timer in milliseconds. This value is always 90 and is not user-configurable.	detail
	• Running VN-Port Timer	Running state of the VN_Port keepalive timer in milliseconds.	detail
	• FCID	Fibre Channel ID of the VN_Port.	detail
	• WWN	Unique worldwide name of the VN_Port.	detail

Sample Output

show fibre-channel fip

```
user@switch> show fibre-channel fip
Fabric Name : proxy2 (200)
  Member
  FCF-MAC : 00:30:48:b0:ee:d2 (Interface vlan.100)
    Enode
    Enode-MAC : 00:10:94:00:00:02       State : Logged-in
        Session
        VN-Port-MAC              : 0e:fc:00:03:00:02
        VN-Port-MAC              : 0e:fc:00:03:00:01
    Enode-MAC : 00:10:94:00:00:03       State : Logged-in
        Session
        VN-Port-MAC              : 0e:fc:00:03:00:04
        VN-Port-MAC              : 0e:fc:00:03:00:03
```

show fibre-channel fip detail

```
user@switch> show fibre-channel fip detail
Configured max FIP sessions per Node Device: 2500
```

```
Node Device: 0  Total FIP sessions: 4  Total FCoE filters: 4

Fabric Name : proxy2 (200)
  FC-MAP         : 0e:fc:00
  FKA-ADV-PERIOD : 90000        MAX-SESSIONS-PER-ENODE : 32
  FCoE trusted   : No

    Member
    FCF-MAC: 00:30:48:b0:ee:d2  (Interface vlan.100)
      FKA-ADV-PERIOD  : 90000    FKA-ADV-D-BIT-bit : Off
      Type : VF_Port Capable
      Priority : 86              State : Enable

        ENode
        Enode-MAC : 00:10:94:00:00:02   ENode State : Logged-in
         Configured ENode timer: 8000    Running ENode timer: 12226
         Active FIP Sessions : 2

          Session details
          VN-Port-MAC              : 0e:fc:00:03:00:02
          Session state            : Up
          Configured FKA-ADV       : 90000
          Running FKA-ADV          : 0
          Configured VN-Port Timer : 90000
          Running VN-Port Timer    : 213193
          FCID                     : 0x2c1a01
          WWN                      : 10:00:00:00:c9:a4:a3:cf

          VN-Port-MAC              : 0e:fc:00:03:00:01
          Session state            : Up
          Configured FKA-ADV       : 90000
          Running FKA-ADV          : 0
          Configured VN-Port Timer : 90000
          Running VN-Port Timer    : 213632
          FCID                     : 0x2c1a02
          WWN                      : 10:00:00:00:d9:b4:e3:df

        ENode
        Enode-MAC : 00:10:94:00:00:03   ENode State : Logged-in
         Configured ENode timer: 8000    Running ENode timer: 12254
         Active FIP Sessions : 2

          Session details
```

```
        VN-Port-MAC              : 0e:fc:00:03:00:04
        Session state            : Up
        Configured FKA-ADV       : 90000
        Running FKA-ADV          : 0
        Configured VN-Port Timer : 90000
        Running VN-Port Timer    : 213480
        FCID                     : 0x2c1a03
        WWN                      : 21:00:00:c0:dd:11:09:13

        VN-Port-MAC              : 0e:fc:00:03:00:03
        Session state            : Up
        Configured FKA-ADV       : 90000
        Running FKA-ADV          : 0
        Configured VN-Port Timer : 90000
        Running VN-Port Timer    : 214004
        FCID                     : 0x2c1a04
        WWN                      : 21:00:00:c0:df:12:08:14
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring FIP on an FCoE-FC Gateway

show fibre-channel fip enode | **19946**

show fibre-channel fip fabric | **19951**

show fibre-channel fip fcf | **19956**

show fibre-channel fip interface | **19961**

show fibre-channel fip statistics | **19966**

clear fibre-channel fip statistics | **3599**

show fibre-channel fip enode

IN THIS SECTION

- Syntax | 19946
- Description | 19946
- Options | 19946
- Required Privilege Level | 19946
- Output Fields | 19947
- Sample Output | 19950
- Release Information | 19951

Syntax

```
show fibre-channel fip enode enode-mac
<brief | detail>
<vn-port-mac vn-port-mac>
```

Description

Display Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) information for a specified ENode or a specified VN_Port on an ENode.

Options

brief | detail	(Optional) Display the specified level of output.
enode-mac	Display information for the ENode specified by the MAC address.
vn-port-mac *vn-port-mac*	(Optional) Display information only for the specified VN_Port.

Required Privilege Level

view

Output Fields

Table 552 on page 19947 lists the output fields for the `show fibre-channel fip enode` command. Output fields are listed in the approximate order in which they appear. A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions.

Table 552: show fibre-channel fip enode Output Fields

Field Name		Field Description	Level of Output
`Fabric Name`		Name of the fabric and in parentheses the fabric ID.	All
`FC-MAP`		FCoE mapped address prefix of the FCoE forwarder for the fabric.	`detail`
`FKA-ADV-PERIOD`		Period of time in milliseconds between FIP keepalive advertisements configured for the FC fabric.	`detail`
`MAX-SESSIONS-PER-ENODE`		Maximum number of concurrent sessions (FLOGI and FDISC combined) that each ENode can instantiate.	`detail`
`FCoE trusted`		Whether ports on the FC fabric are trusted or untrusted: • `Yes`—Ports on the FC fabric are trusted; FIP snooping is turned off. • `No`—Ports on the FC fabric are not trusted; FIP snooping is turned on.	`detail`
`Member`		Information about an FCF that is a member of the fabric.	All
	• `FCF-MAC`	MAC address used in discovery advertisements.	All

Table 552: show fibre-channel fip enode Output Fields *(Continued)*

Field Name		Field Description	Level of Output
	• FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC interface.	detail
	• FKA-ADV-D-BIT	Disable FIP keepalive advertisement monitoring bit. The state is always off.	detail
	• Type	Type of interface: • VF_Port Capable—Interface can act as a VF_Port interface.	detail
	• Priority	Priority value associated with the switch FCF-MAC. Converged network adapters (CNAs) use the priority value to determine the switch with which they will perform FIP FLOGI. The lower the value, the higher the priority. Value range: 0 through 255.	detail
	• State	FIP state on the fabric: • Enable—FIP is enabled on the fabric. • Disable—FIP is disabled on the fabric.	detail
ENode		Information about a connected FCoE node (ENode).	All
	• ENode-MAC	MAC address of the connected ENode.	All
	• ENode State	Login state internal to Junos OS.	All
	• Configured ENode timer	User-configured FIP keepalive advertisement interval in milliseconds.	detail

Table 552: show fibre-channel fip enode Output Fields *(Continued)*

Field Name	Field Description	Level of Output
• `Running ENode timer`	Runtime interval in milliseconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	`detail`
• `Active FIP Sessions`	Number of active FIP sessions on the ENode.	`detail`
• `VN-Port-MAC`	MAC address of a VN_Port on the ENode.	`detail`
• `Session State`	Session state internal to Junos OS.	`detail`
• `Configured FKA-ADV`	User-configured FIP keepalive advertisement interval in milliseconds.	`detail`
• `Running FKA-ADV`	Runtime interval in seconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	`detail`
• `Configured VN-Port Timer`	Configured state of the VN_Port keepalive timer in seconds. This value is always 90 and is not user-configurable.	`detail`
• `Running VN-Port Timer`	Running state of the VN_Port keepalive timer in seconds.	`detail`
• `FCID`	Fibre Channel ID of the VN_Port.	`detail`
• `WWN`	Unique worldwide name of the VN_Port.	`detail`

Sample Output

show fibre-channel fip enode

```
user@switch> show fibre-channel fip enode 00:10:94:00:00:02
Fabric Name : proxy2 (200)
  Member
  FCF-MAC : 00:30:48:b0:ee:d2 (Interface vlan.100)
    Enode
    Enode-MAC : 00:10:94:00:00:02       State : Logged-in
        Session
        VN-Port-MAC              : 0e:fc:00:03:00:02
        VN-Port-MAC              : 0e:fc:00:03:00:01
```

show fibre-channel fip enode detail

```
user@switch> show fibre-channel fip enode 00:10:94:00:00:02 detail
Fabric Name : proxy2 (200)
  FC-MAP         : 0e:fc:00
  FKA-ADV-PERIOD : 90000        MAX-SESSIONS-PER-ENODE : 32
  FCoE trusted   : No

    Member
    FCF-MAC: 00:30:48:b0:ee:d2  (Interface vlan.100)
      FKA-ADV-PERIOD  : 90000    FKA-ADV-D-BIT-bit : Off
      Type : VF_Port Capable
      Priority : 86             State : Enable

        ENode
        Enode-MAC : 00:10:94:00:00:02   ENode State : Logged-in
         Configured ENode timer: 8000    Running ENode timer: 12226
         Active FIP Sessions : 2

          Session details
          VN-Port-MAC              : 0e:fc:00:03:00:02
          Session state            : Up
          Configured FKA-ADV       : 90000
          Running FKA-ADV          : 0
          Configured VN-Port Timer : 90000
          Running VN-Port Timer    : 213193
```

```
        FCID                      : 0x2c1a01
        WWN                       : 10:00:00:00:c9:a4:a3:cf

        VN-Port-MAC               : 0e:fc:00:03:00:01
        Session state             : Up
        Configured FKA-ADV        : 90000
        Running FKA-ADV           : 0
        Configured VN-Port Timer  : 90000
        Running VN-Port Timer     : 213632
        FCID                      : 0x2c1a02
        WWN                       : 10:00:00:00:d9:b4:e3:df
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring FIP on an FCoE-FC Gateway

show fibre-channel fip | **19938**

show fibre-channel fip fabric | **19951**

show fibre-channel fip fcf | **19956**

show fibre-channel fip interface | **19961**

show fibre-channel fip statistics | **19966**

clear fibre-channel fip enode | **3597**

show fibre-channel fip fabric

IN THIS SECTION

- Syntax | **19952**
- Description | **19952**
- Options | **19952**
- Required Privilege Level | **19952**

- Output Fields | 19952
- Sample Output | 19954
- Release Information | 19956

Syntax

```
show fibre-channel fip fabric fabric-name
<brief | detail>
```

Description

Display Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) information for a specified Fibre Channel fabric.

Options

brief | detail	(Optional) Display the specified level of output.
fabric-name	Display information for the specified fabric.

Required Privilege Level

view

Output Fields

Table 553 on page 19953 lists the output fields for the `show fibre-channel fip fabric` command. Output fields are listed in the approximate order in which they appear.

Table 553: show fibre-channel fip fabric Output Fields

Field Name		Field Description	Level of Output
`Fabric Name`		Name of the fabric and in parentheses the fabric ID.	All
`FC-MAP`		FCoE mapped address prefix of the FCoE forwarder for the fabric.	`detail`
`FKA-ADV-PERIOD`		Period of time in milliseconds between FIP keepalive advertisements configured for the FC fabric.	`detail`
`Member`		Information about an FCF that is a member of the fabric.	All
	• `FCF-MAC`	MAC address used in discovery advertisements.	All
	• `FKA-ADV-PERIOD`	Period of time in milliseconds between FIP keepalive advertisements configured for the FC interface.	`detail`
	• `FKA-ADV-D-BIT`	Disable FIP keepalive advertisement monitoring bit. The state is always `off`.	`detail`
	• `Type`	Type of interface: • `VF_Port Capable`—Interface can act as a VF_Port interface.	`detail`
`ENode`		Information about a connected FCoE node (ENode).	All
	• `ENode-MAC`	MAC address of the connected ENode.	All
	• `State`	Login state internal to Junos OS.	All

Table 553: show fibre-channel fip fabric Output Fields *(Continued)*

Field Name		Field Description	Level of Output
	• VN-Port-MAC	MAC address of a VN_Port on the ENode.	detail
	• Session State	Session state internal to Junos OS.	detail
	• Configured FKA-ADV	User-configured FIP keepalive advertisement interval in milliseconds.	detail
	• Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	detail
	• Configured VN-Port Timer	Configured state of the VN_Port keepalive timer in seconds. This value is always 90 and is not user-configurable.	detail
	• Running VN-Port Timer	Running state of the VN_Port keepalive timer in seconds.	detail

Sample Output

show fibre-channel fip fabric proxy2

```
user@switch> show fibre-channel fip fabric proxy2
Fabric Name : proxy2 (200)
  Member
  FCF-MAC : 00:30:48:b0:ee:d2 (Interface vlan.100)
    Enode
    Enode-MAC : 00:10:94:00:00:02       State : Logged-in
    Enode-MAC : 00:10:94:00:00:03       State : Logged-in
```

show fibre-channel fip fabric detail

```
user@switch> show fibre-channel fip fabric proxy2 detail
Fabric Name : proxy2 (200)
  FC-MAP         : 0e:fc:00
  FKA-ADV-PERIOD : 90000

    Member
    FCF-MAC: 00:30:48:b0:ee:d2  (Interface vlan.100)
      FKA-ADV-PERIOD  : 90000     FKA-ADV-D-bit : Off
      Type : VF_Port Capable

        ENode
        Enode-MAC : 00:10:94:00:00:02   State : Logged-in

          Session details
          VN-Port-MAC              : 0e:fc:00:03:00:02
          Session state            : Up
          Configured FKA-ADV       : 90000
          Running FKA-ADV          : 0
          Configured VN-Port Timer : 90
          Running VN-Port Timer    : 0

          VN-Port-MAC              : 0e:fc:00:03:00:01
          Session state            : Up
          Configured FKA-ADV       : 90000
          Running FKA-ADV          : 0
          Configured VN-Port Timer : 90
          Running VN-Port Timer    : 0

        ENode
        Enode-MAC : 00:10:94:00:00:03   State : Logged-in

          Session details
          VN-Port-MAC              : 0e:fc:00:03:00:04
          Session state            : Up
          Configured FKA-ADV       : 90000
          Running FKA-ADV          : 0
          Configured VN-Port Timer : 90
          Running VN-Port Timer    : 0

          VN-Port-MAC              : 0e:fc:00:03:00:03
```

```
            Session state            : Up
            Configured FKA-ADV       : 90000
            Running FKA-ADV          : 0
            Configured VN-Port Timer : 90
            Running VN-Port Timer    : 0
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring FIP on an FCoE-FC Gateway

show fibre-channel fip | **19938**

show fibre-channel fip enode | **19946**

show fibre-channel fip fcf | **19956**

show fibre-channel fip interface | **19961**

show fibre-channel fip statistics | **19966**

show fibre-channel fip fcf

IN THIS SECTION

- Syntax | **19957**
- Description | **19957**
- Options | **19957**
- Required Privilege Level | **19957**
- Output Fields | **19957**
- Sample Output | **19959**
- Release Information | **19961**

Syntax

```
show fibre-channel fip fcf fcf-mac
<brief | detail>
<fabric fabric-name>
```

Description

Display Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) information for a specified FCoE forwarder (FCF).

Options

brief | detail	(Optional) Display the specified level of output.
fabric *fabric-name*	(Optional) Display FCF information only for the specified fabric.
fcf-mac	Display information for the FCF specified by the MAC address.

Required Privilege Level

view

Output Fields

Table 554 on page 19957 lists the output fields for the `show fibre-channel fip fcf` command. Output fields are listed in the approximate order in which they appear.

Table 554: show fibre-channel fip fcf Output Fields

Field Name	Field Description	Level of Output
`Fabric Name`	Name of the fabric and in parentheses the fabric ID.	All
`FC-MAP`	FCoE mapped address prefix of the FCoE forwarder for the fabric.	`detail`

Table 554: show fibre-channel fip fcf Output Fields *(Continued)*

Field Name		Field Description	Level of Output
FKA-ADV-PERIOD		Period of time in milliseconds between FIP keepalive advertisements configured for the FC fabric.	detail
Member		Information about an FCF that is a member of the fabric.	All
	• FCF-MAC	MAC address used in discovery advertisements.	All
	• FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC interface.	detail
	• FKA-ADV-D-BIT	Disable FIP keepalive advertisement monitoring bit. The state is always off.	detail
	• Type	Type of interface: • VF_Port Capable—Interface can act as a VF_Port interface.	detail
ENode		Information about a connected FCoE node (ENode).	All
	• ENode-MAC	MAC address of the connected ENode.	All
	• State	Login state internal to Junos OS.	All
	• VN-Port-MAC	MAC address of a VN_Port on the ENode.	detail
	• Session State	Session state internal to Junos OS.	detail

Table 554: show fibre-channel fip fcf Output Fields *(Continued)*

Field Name		Field Description	Level of Output
	• `Configured FKA-ADV`	User-configured FIP keepalive advertisement interval in milliseconds.	`detail`
	• `Running FKA-ADV`	Runtime interval in seconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	`detail`
	• `Configured VN-Port Timer`	Configured state of the VN_Port keepalive timer in seconds. This value is always 90 and is not user-configurable.	`detail`
	• `Running VN-Port Timer`	Running state of the VN_Port keepalive timer in seconds.	`detail`

Sample Output

show fibre-channel fip fcf

```
user@switch> show fibre-channel fip fcf 00:30:48:b0:ee:d2
Fabric Name : proxy2 (200)
  Member
  FCF-MAC : 00:30:48:b0:ee:d2 (Interface vlan.100)
    Enode
    Enode-MAC : 00:10:94:00:00:02       State : Logged-in
    Enode-MAC : 00:10:94:00:00:03       State : Logged-in
```

show fibre-channel fip fcf detail

```
user@switch> show fibre-channel fip fcf 00:30:48:b0:ee:d2 detail
Fabric Name : proxy2 (200)
  FC-MAP         : 0e:fc:00
  FKA-ADV-PERIOD : 90000
```

```
Member
FCF-MAC: 00:30:48:b0:ee:d2  (Interface vlan.100)
  FKA-ADV-PERIOD  : 90000     FKA-ADV-D-bit : Off
  Type : VF_Port Capable

    ENode
    Enode-MAC : 00:10:94:00:00:02   State : Logged-in

      Session details
      VN-Port-MAC              : 0e:fc:00:03:00:02
      Session state            : Up
      Configured FKA-ADV       : 90000
      Running FKA-ADV          : 0
      Configured VN-Port Timer : 90
      Running VN-Port Timer    : 0

      VN-Port-MAC              : 0e:fc:00:03:00:01
      Session state            : Up
      Configured FKA-ADV       : 90000
      Running FKA-ADV          : 0
      Configured VN-Port Timer : 90
      Running VN-Port Timer    : 0

    ENode
    Enode-MAC : 00:10:94:00:00:03   State : Logged-in

      Session details
      VN-Port-MAC              : 0e:fc:00:03:00:04
      Session state            : Up
      Configured FKA-ADV       : 90000
      Running FKA-ADV          : 0
      Configured VN-Port Timer : 90
      Running VN-Port Timer    : 0

      VN-Port-MAC              : 0e:fc:00:03:00:03
      Session state            : Up
      Configured FKA-ADV       : 90000
      Running FKA-ADV          : 0
      Configured VN-Port Timer : 90
      Running VN-Port Timer    : 0
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring FIP on an FCoE-FC Gateway

show fibre-channel fip | 19938

show fibre-channel fip enode | 19946

show fibre-channel fip fabric | 19951

show fibre-channel fip interface | 19961

show fibre-channel fip statistics | 19966

show fibre-channel fip interface

IN THIS SECTION

- Syntax | 19961
- Description | 19962
- Options | 19962
- Required Privilege Level | 19962
- Output Fields | 19962
- Sample Output | 19964
- Release Information | 19965

Syntax

```
show fibre-channel fip interface interface-name
<brief | detail>
<enode enode-mac>
<fabric fabric-name>
<vn-port vn-port-mac>
```

Description

Display Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) information for a specified interface.

Options

brief | detail	(Optional) Display the specified level of output.
enode-mac	MAC address of the ENode.
fabric *fabric-name*	(Optional) Display interface information only for the specified fabric.
interface-name	Display information for the specified interface.
vn-port-mac	MAC address of the VN_Port.

Required Privilege Level

view

Output Fields

Table 555 on page 19962 lists the output fields for the `show fibre-channel fip interface` command. Output fields are listed in the approximate order in which they appear.

Table 555: show fibre-channel fip interface Output Fields

Field Name	Field Description	Level of Output
`Fabric Name`	Name of the fabric and in parentheses the fabric ID.	All
`FC-MAP`	FCoE mapped address prefix of the FCoE forwarder for the fabric.	`detail`
`FKA-ADV-PERIOD`	Period of time in milliseconds between FIP keepalive advertisements configured for the FC fabric.	`detail`

Table 555: show fibre-channel fip interface Output Fields *(Continued)*

Field Name		Field Description	Level of Output
Member		Information about an FCF that is a member of the fabric.	All
	• FCF-MAC	MAC address used in discovery advertisements.	All
	• FKA-ADV-PERIOD	Period of time in milliseconds between FIP keepalive advertisements configured for the FC interface.	detail
	• FKA-ADV-D-BIT	Disable FIP keepalive advertisement monitoring bit. The state is always off.	detail
	• Type	Type of interface: • VF_Port Capable—Interface can act as a VF_Port interface.	detail
ENode		Information about a connected FCoE node (ENode).	All
	• ENode-MAC	MAC address of the connected ENode.	All
	• State	Login state internal to Junos OS.	All
	• VN-Port-MAC	MAC address of a VN_Port on the ENode.	detail
	• Session State	Session state internal to Junos OS.	detail
	• Configured FKA-ADV	User-configured FIP keepalive advertisement interval in milliseconds.	detail

Table 555: show fibre-channel fip interface Output Fields *(Continued)*

Field Name		Field Description	Level of Output
	• Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement received. This value changes every time an FKA_ADV is received.	detail
	• Configured VN-Port Timer	Configured state of the VN_Port keepalive timer in seconds. This value is always 90 and is not user-configurable.	detail
	• Running VN-Port Timer	Running state of the VN_Port keepalive timer in seconds.	detail

Sample Output

show fibre-channel fip interface

```
user@switch> show fibre-channel fip interface vlan.100
Fabric Name : proxy2 (200)
  Member
  FCF-MAC : 00:30:48:b0:ee:d2 (Interface vlan.100)
    Enode
    Enode-MAC : 00:10:94:00:00:02       State : Logged-in
    Enode-MAC : 00:10:94:00:00:03       State : Logged-in
```

show fibre-channel fip interface detail

```
user@switch> show fibre-channel fip interface vlan.100 detail
Fabric Name : proxy2 (200)
  FC-MAP         : 0e:fc:00
  FKA-ADV-PERIOD : 90000

    Member
    FCF-MAC: 00:30:48:b0:ee:d2  (Interface vlan.100)
      FKA-ADV-PERIOD  : 90000     FKA-ADV-D-bit : Off
      Type : VF_Port Capable
```

```
ENode
Enode-MAC : 00:10:94:00:00:02   State : Logged-in

  Session details
  VN-Port-MAC              : 0e:fc:00:03:00:02
  Session state            : Up
  Configured FKA-ADV       : 90000
  Running FKA-ADV          : 0
  Configured VN-Port Timer : 90
  Running VN-Port Timer    : 0

  VN-Port-MAC              : 0e:fc:00:03:00:01
  Session state            : Up
  Configured FKA-ADV       : 90000
  Running FKA-ADV          : 0
  Configured VN-Port Timer : 90
  Running VN-Port Timer    : 0

ENode
Enode-MAC : 00:10:94:00:00:03   State : Logged-in

  Session details
  VN-Port-MAC              : 0e:fc:00:03:00:04
  Session state            : Up
  Configured FKA-ADV       : 90000
  Running FKA-ADV          : 0
  Configured VN-Port Timer : 90
  Running VN-Port Timer    : 0

  VN-Port-MAC              : 0e:fc:00:03:00:03
  Session state            : Up
  Configured FKA-ADV       : 90000
  Running FKA-ADV          : 0
  Configured VN-Port Timer : 90
  Running VN-Port Timer    : 0
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring FIP on an FCoE-FC Gateway

show fibre-channel fip | 19938

show fibre-channel fip enode | 19946

show fibre-channel fip fabric | 19951

show fibre-channel fip fcf | 19956

show fibre-channel fip statistics | 19966

clear fibre-channel fip vn-port | 3600

show fibre-channel fip statistics

IN THIS SECTION

- Syntax | 19966
- Description | 19966
- Options | 19967
- Required Privilege Level | 19967
- Output Fields | 19967
- Sample Output | 19970
- Release Information | 19970

Syntax

```
show fibre-channel fip statistics
<fabric fabric-name>
```

Description

Display Fibre Channel over Ethernet Initialization Protocol (FIP) statistics.

Options

`fabric` *`fabric-name`*	(Optional) Display output only for the specified fabric.

Required Privilege Level

view

Output Fields

Table 556 on page 19967 lists the output fields for the `show fibre-channel fip statistics` command. Output fields are listed in the approximate order in which they appear.

Table 556: show fibre-channel fip statistics Output Fields

Field Name		Field Description
`Fabric name`		Name of the fabric.
`Interface name`		Name of the FCoE VLAN interface.
`FIP Message Type`		Type of FIP message for the displayed row of statistics..
	• `MDS`	Number of multicast discovery solicitations.
	• `UDS`	Number of unicast discovery solicitations.
	• `FLOGI`	Number of fabric login (FLOGI) messages.
	• `FDISC`	Number of fabric discovery (FDISC) messages.
	• `LOGO`	Number of fabric logout (LOGO) messages.

Table 556: show fibre-channel fip statistics Output Fields *(Continued)*

Field Name		Field Description
	• ENODE KA	Number of ENode keepalive messages.
	• VN_Port KA	Number of VN_Port keepalive messages.
	• MDA	Number of multicast discovery advertisements.
	• UDA	Number of unicast discovery advertisements.
	• FLOGI ACC	Number of fabric login requests accepted.
	• FLOGI RJT	Number of fabric login requests rejected.
	• FDISC ACC	Number of fabric discovery requests accepted.
	• FDISC RJT	Number of fabric discovery requests rejected.
	• LOGO ACC	Number of logout requests accepted.
	• LOGO RJT	Number of logout requests rejected.
	• CVL	Number of clear virtual links (CVL) messages.
	• CVL ALL	Number of CVL all messages.
Received		Number of messages received.
Sent		Number of messages sent.

Table 556: show fibre-channel fip statistics Output Fields *(Continued)*

Field Name		Field Description
`Rx errors`		Number of receive errors.
`Dropped`		Number of dropped messages. **NOTE**: One cause of dropped messages is that the system limits the number of discovery solicitations (MDS and UDS) it accepts to a maximum of 100 outstanding requests at any given time. If the system has 100 discovery solicitations outstanding, the system does not respond to new discovery solicitations. Instead, the system drops new discovery solicitations and reports the number of dropped discovery solicitations in this field. When there are fewer than 100 outstanding discovery solicitations, the system responds to new requests as usual with a discovery advertisement.
`General Statistics`	`Number of frames recvd with invalid src-mac`	Number of frames received that have an invalid source media access control (MAC) address.
	`Number of frames recvd with invalid version`	Number of FIP frames received with an Invalid FIP version.
	`Number of frames recvd with invalid opcode`	Number of FIP validation descriptors with an invalid opcode received.
	`Number of frames recvd with invalid subcode`	Number of FIP validation descriptors with an invalid subcode received.
	`Number of frames recvd on inactive FCF`	Number of frames received on a logical interface if FIP is not active on that logical interface (for example, if a WWN is not allocated to that logical interface).

Sample Output

show fibre-channel fip statistics

```
user@switch> show fibre-channel fip statistics
Fabric name: proxy2

Interface name: vlan.100
FIP Message type      Received     Sent     Rx errors   Dropped
MDS                   22236        0        0           17089
UDS                   0            0        0           0
FLOGI                 1257         0        8           0
FDISC                 0            0        0           0
LOGO                  0            0        0           0
ENODE KA              455          0        6           0
VN_Port KA            22           0        0           0
MDA                   0            243      0           0
UDA                   0            5147     0           0
FLOGI ACC             0            376      0           0
FLOGI RJT             0            881      0           0
FDISC ACC             0            0        0           0
FDISC RJT             0            0        0           0
LOGO ACC              0            0        0           0
LOGO RJT              0            0        0           0
CVL                   0            374      0           0
CVL ALL               0            380      0           0



General Statistics:

Number of frame recvd with invalid src-mac:           0
Number of frame recvd with invalid version:           0
Number of frame recvd with invalid opcode:            0
Number of frame recvd with invalid subcode:           0
Number of frame recvd on inactive FCF:                0
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel fip | 19938
show fibre-channel fip enode | 19946
show fibre-channel fip fabric | 19951
show fibre-channel fip fcf | 19956
show fibre-channel fip interface | 19961
clear fibre-channel fip statistics | 3599

show fibre-channel flogi fport

IN THIS SECTION

- Syntax | 19971
- Description | 19971
- Options | 19972
- Required Privilege Level | 19972
- Output Fields | 19972
- Sample Output | 19973
- Release Information | 19973

Syntax

```
show fibre-channel flogi fport
<fabric fabric-name>
```

Description

Display Fibre Channel fabric login (FLOGI) F_Port information.

Options

`fabric` *`fabric-name`*	(Optional) Display output only for the specified fabric.

Required Privilege Level

view

Output Fields

Table 557 on page 19972 lists the output fields for the `show fibre-channel flogi fport` command. Output fields are listed in the approximate order in which they appear.

Table 557: show fibre-channel flogi fport Output Fields

Field Name	Field Description
`Fabric`	Name of the fabric.
`Interface`	Name of the switch VF_Port interface.
`Mac-Address`	Media access control (MAC) address of the ENode.
`State`	Interface physical state: `up` or `down`.
`Logins`	Number of logins to the VF_Port.
`NPIV`	N_Port ID virtualization (NPIV) state: `Yes` or `No`.
`FLOGI-Port-WWN`	Unique worldwide name (WWN) of the VN_Port performing fabric login (FLOGI) to the switch VF_Port.

Sample Output

show fibre-channel flogi fport

```
user@switch> show fibre-channel flogi fport
Fabric: proxy2
Interface       Mac-Address         State  Logins   NPIV  FLOGI-Port-WWN
vlan.100        00:10:94:00:00:02   Up     2        Yes   20:00:10:94:00:01:00:01
vlan.100        00:10:94:00:00:03   Up     2        Yes   20:00:10:94:00:02:00:01
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel flogi nport | 19973

show fibre-channel flogi statistics | 19977

show fibre-channel flogi nport

IN THIS SECTION

- Syntax | 19974
- Description | 19974
- Options | 19974
- Required Privilege Level | 19974
- Output Fields | 19974
- Sample Output | 19975
- Release Information | 19976

Syntax

```
show fibre-channel flogi nport
<brief | detail>
<fabric fabric-name>
```

Description

Display Fibre Channel fabric login (FLOGI) VN_Port information.

Options

brief | detail	(Optional) Display the specified level of output.
fabric *fabric-name*	(Optional) Display output only for the specified fabric.

Required Privilege Level

view

Output Fields

Table 558 on page 19974 lists the output fields for the `show fibre-channel flogi nport` command. Output fields are listed in the approximate order in which they appear.

Table 558: show fibre-channel flogi nport Output Fields

Field Name	Field Description	Level of Output
`Fabric`	Name of the fabric.	All
`Virtual-switch`	Name of the fabric.	`detail`
`Interface`	Name of the VF_Port interface.	All

Table 558: show fibre-channel flogi nport Output Fields *(Continued)*

Field Name	Field Description	Level of Output
FCID	VN_Port Fibre Channel identifier provided by the Fibre Channel over Ethernet Forwarder (FCoE forwarder) or the Fibre Channel switch.	All
Port-WWN	Unique worldwide name (WWN) of the VN_Port.	All
Node-WWN	Unique WWN of the node hosting the VN_Port.	All
State or Flogi-state	Login state internal to Junos OS.	All
FLOGI-Port-WWN	Unique worldwide name (WWN) of the VN_Port performing fabric login (FLOGI) to the switch VF_Port.	detail

Sample Output

show fibre-channel flogi nport

```
user@switch> show fibre-channel flogi nport
Fabric: proxy2
Interface     FCID      Port-WWN                 Node-WWN                State
vlan.100      0x030001  20:00:10:94:00:01:00:01  10:00:10:94:00:00:00:01 online
vlan.100      0x030002  20:00:10:94:00:01:00:05  10:00:10:94:00:00:00:01 online
vlan.100      0x030003  20:00:10:94:00:02:00:01  10:00:10:94:00:00:00:02 online
vlan.100      0x030004  20:00:10:94:00:02:00:05  10:00:10:94:00:00:00:02 online
```

show fibre-channel flogi nport detail

```
user@switch> show fibre-channel flogi nport detail
Fabric: proxy2
    Virtual-switch: proxy2
```

```
Interface: vlan.100
Flogi-state: online
FCID: 0x030001
Port-WWN: 20:00:10:94:00:01:00:01
Node-WWN: 10:00:10:94:00:00:00:01
FLOGI-Port-WWN: 20:00:10:94:00:01:00:01

Interface: vlan.100
Flogi-state: online
FCID: 0x030002
Port-WWN: 20:00:10:94:00:01:00:05
Node-WWN: 10:00:10:94:00:00:00:01
FLOGI-Port-WWN: 20:00:10:94:00:01:00:01

Interface: vlan.100
Flogi-state: online
FCID: 0x030003
Port-WWN: 20:00:10:94:00:02:00:01
Node-WWN: 10:00:10:94:00:00:00:02
FLOGI-Port-WWN: 20:00:10:94:00:02:00:01

Interface: vlan.100
Flogi-state: online
FCID: 0x030004
Port-WWN: 20:00:10:94:00:02:00:05
Node-WWN: 10:00:10:94:00:00:00:02
FLOGI-Port-WWN: 20:00:10:94:00:02:00:01
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel flogi fport | **19971**

show fibre-channel flogi statistics | **19977**

show fibre-channel flogi statistics

IN THIS SECTION

- Syntax | 19977
- Description | 19977
- Options | 19977
- Required Privilege Level | 19977
- Output Fields | 19977
- Sample Output | 19979
- Release Information | 19980

Syntax

```
show fibre-channel flogi statistics
<fabric fabric-name>
```

Description

Display Fibre Channel fabric login (FLOGI) statistics.

Options

fabric ***fabric-name***	(Optional) Display output only for the specified fabric.

Required Privilege Level

view

Output Fields

Table 559 on page 19978 lists the output fields for the `show fibre-channel flogi statistics` command. Output fields are listed in the approximate order in which they appear.

Table 559: show fibre-channel flogi statistics Output Fields

Field Name		Field Description
`Fabric`		Name of the fabric.
`FLOGI-Server Message type`		Type of message: • `FLOGI`—Fabric login (FLOGI) messages. • `FDISC`—Fabric discovery (FDISC) messages. • `FLOGO`—Fabric logout messages. • `FLOGO-LS-ACC`—Fabric logout link service accept messages. • `LS-Accept`—Link service accept messages. • `LS-Reject`—Link service reject messages. • `invalid`—Invalid messages.
`Received`		Number of messages received for a given message type.
`Sent`		Number of messages sent for a given message type.
`Fabric`		Name of the fabric.
`Rx errors`		Number of receive errors for a given type of message.
`General Statistics`	• `Number of FC2 Header Parse Errors`	Number of errors parsing the FC-2 header.
	• `Number of FLOGI Parse Errors`	Number of errors parsing fabric login requests.
	• `Number of FDISC Parse Errors`	Number of errors parsing fabric discovery requests.

Table 559: show fibre-channel flogi statistics Output Fields *(Continued)*

Field Name		Field Description
	• Number of FLOGO Parse Errors	Number of errors parsing fabric logout requests.
	• Number of Logins Discarded as Domain-ID not available	Number of discarded logins due to unavailability of a domain ID.
	• Number of Logins Discarded as FCID not available	Number of discarded logins due to the unavailability of a Fibre Channel ID.
	• Number of FCID requests deferred	Number of deferred FCID requests.
	• Number of deferred FCID requests failed	Number of deferred FCID requests that failed.

Sample Output

show fibre-channel flogi statistics

```
user@switch> show fibre-channel flogi statistics
Fabric: proxy2

FLOGI-Server Message type   Received       Sent   Rx errors
FLOGI                       2              0      0
FDISC                       2              0      0
FLOGO                       0              0      0
FLOGO-LS-ACC                0              0      0
```

```
LS-Accept                    0              4      0
LS-Reject                    0              0      0
invalid                      0              0      0

General Statistics:

Number of FC2 Header Parse Errors:                          0
Number of FLOGI Parse Errors:                               0
Number of FDISC Parse Errors:                               0
Number of FLOGO Parse Errors:                               0
Number of Logins Discarded as Domain-ID not available:      0
Number of Logins Discarded as FCID not available:           0
Number of FCID requests deferred:                           0
Number of deferred FCID requests failed:                    0
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- show fibre-channel flogi fport | 19971
- show fibre-channel flogi nport | 19973
- clear fibre-channel flogi statistics | 3602

show fibre-channel interfaces

IN THIS SECTION

- Syntax | 19981
- Description | 19981
- Options | 19981
- Required Privilege Level | 19981
- Output Fields | 19981
- Sample Output | 19983

Release Information | 19984

Syntax

```
<brief | detail>
<fabric fabric-name>
show fibre-channel interfaces interface-name
```

Description

Display information about Fibre Channel (FC) interfaces.

Options

brief | detail	(Optional) Display the specified level of output.
`fabric` *`fabric-name`*	(Optional) Display output only for the specified fabric.
`interface-name`	Display output for the specified interface.

Required Privilege Level

view

Output Fields

Table 560 on page 19981 lists the output fields for the `show fibre-channel interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 560: show fibre-channel interfaces Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the FC interface.	All

Table 560: show fibre-channel interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Idx` or `Index`	Interface index internal to Junos OS.	All
`Type`	Type of interface: • `FC`—Native FC interface • `FCOE`—Fibre Channel over Ethernet interface	All
`Native Fabric-id`	Identification number of the QFX Series fabric.	All
`NPIV`	N_Port ID virtualization (NPIV) state: `Yes` or `No`.	All
`Config-Mode`	User-configured port mode: • `F`—The port is configured as a VF_Port, an FCoE port connected to FCoE devices. • `NP`—The port is configured as a proxy N_Port (NP_Port), a native FC port connected to an FC switch.	All
`Oper-Mode`	Operational port mode: • `F`—The port is operating as a VF_Port, an FCoE port connected to FCoE devices. • `NP`—The port is operating as an NP_Port, a native FC port connected to an FC switch or an FCoE forwarder (FCF).	All
`State`	Interface state: `up` or `down`.	All
`WWN`	Unique worldwide name (WWN) of the port.	`detail`
`FSM-State`	Finite state machine state, internal to Junos OS.	`detail`

Table 560: show fibre-channel interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Class ID	Fibre Channel interface class ID, internal to Junos OS.	detail
BB_SC_N	Buffer-to-buffer state change number.	detail
Tx B2B credits	Number of buffer-to-buffer credits advertised by the neighbor switch that is connected to the FC interface.	detail
Fabric	Name of the fabric.	detail
Remote-MAC	Media access control (MAC) address of the remotely connected FCoE device VN_Port interface.	detail
Tagging	Not used. Value is shown as untagged.	detail
Mode	Logical interface (LIF) mode of operation.	detail
H/W token	Unique identifier for the FCoE VLAN interface, internal to Junos OS.	detail

Sample Output

show fibre-channel interfaces

```
user@switch> show fibre-channel interfaces
                                Native                Config  Oper
Interface          Idx    Type  Fabric-id    NPIV     Mode    Mode    State
fc-0/0/1.0         70     FC    200          YES      NP      NP      up
vlan.100           84     FCOE  200          YES      F       F       up
```

show fibre-channel interfaces detail

```
user@switch> show fibre-channel interfaces detail
Interface: fc-0/0/1.0, Index: 70, Type: FC, Native Fabric-id: 200
NPIV: YES, Config-Mode: NP, Oper-Mode: NP, State: up
WWN: 10:00:00:15:17:a9:98:64, FSM-State: up, Class ID: 1, BB_SC_N: 0
Tx B2B credits: 32
  Fabric                Remote-MAC         Tagging    Mode   Oper state
  proxy2                -                  untagged   NP     up

Interface: vlan.100, Index: 84, Type: FCOE, Native Fabric-id: 200
NPIV: YES, Config-Mode: F, Oper-Mode: F, State: up
WWN: 10:00:00:30:48:b0:ee:d2, FSM-State: up
H/W token: 13
  Fabric                Remote-MAC         Tagging    Mode   Oper state
  proxy2                00:10:94:00:00:02  untagged   VF     up
  proxy2                00:10:94:00:00:03  untagged   VF     up
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric

Configuring a Physical Fibre Channel Interface

Configuring an FCoE VLAN Interface on an FCoE-FC Gateway

Converting an Ethernet Interface To a Fibre Channel Interface

Assigning Interfaces to a Fibre Channel Fabric

show fibre-channel next-hops

IN THIS SECTION

- Syntax | 19985
- Description | 19985
- Required Privilege Level | 19985
- Output Fields | 19985
- Sample Output | 19986
- Release Information | 19987

Syntax

```
show fibre-channel next-hops
```

Description

Display Fibre Channel next-hop route information.

Required Privilege Level

view

Output Fields

Table 561 on page 19985 lists the output fields for the `show fibre-channel next-hops` command. Output fields are listed in the approximate order in which they appear.

Table 561: show fibre-channel next-hops Output Fields

Field Name	Field Description
`Type`	Type of next hop internal to Junos OS.

Table 561: show fibre-channel next-hops Output Fields *(Continued)*

Field Name	Field Description
State	State of the NP_Port interface: • Active—The interface is online. • Deleted—The interface is deleted.
Interface	Name of the interface.
Mac-Address	Media access control (MAC) address of the interface.
Index	Next-hop index identifier.
Ref-count	Reference count internal to Junos OS.
Flags	Flags internal to Junos OS.

Sample Output

show fibre-channel next-hops

```
user@switch> show fibre-channel next-hops
```

```
Type    State  Interface    Mac-Address         Index  Ref-count Flags
intf    Active fc-0/0/0.0                       0       1
ucast   Active vlan.100     00:15:17:a9:98:64   674     1        kernel, self
ucast   Active vlan.100     0e:fc:00:03:00:01   675     1        kernel, self
ucast   Active vlan.100     0e:fc:00:03:00:02   676     1        kernel, self
ucast   Active vlan.100     0e:fc:00:03:00:03   677     1        kernel, self
ucast   Active vlan.100     0e:fc:00:03:00:04   678     1        kernel, self
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel routes | 20006

show route forwarding-table family fibre-channel | 23673

show fibre-channel proxy fabric-state

IN THIS SECTION

- Syntax | 19987
- Description | 19987
- Options | 19988
- Required Privilege Level | 19988
- Output Fields | 19988
- Sample Output | 19991
- Release Information | 19992

Syntax

```
show fibre-channel proxy fabric-state
<fabric fabric-name>
```

Description

Display Fibre Channel (FC) proxy fabric state information.

Options

`fabric` *`fabric-name`*	(Optional) Display output only for the specified fabric.

Required Privilege Level

view

Output Fields

Table 562 on page 19988 lists the output fields for the `show fibre-channel proxy fabric-state` command. Output fields are listed in the approximate order in which they appear.

Table 562: show fibre-channel proxy fabric-state Output Fields

Field Name	Field Description
`Fabric`	Name of the fabric.
`Fabric-id`	Fabric ID number.

Table 562: show fibre-channel proxy fabric-state Output Fields *(Continued)*

Field Name	Field Description
`Proxy load balance algorithm`	Load-balancing algorithm used on the FCoE-FC gateway FC fabric: • Simple—Load balancing is based on the weighted utilization (load) of the NP_Ports connected to an FC fabric. Each new FLOGI or FDISC is assigned to the least-loaded link. On a link load rebalance, only the sessions that need to be moved to another link are logged out. When those sessions log in again, they are placed on active NP_Port interfaces in a balanced manner. • ENode-based—Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. On a link load rebalance, all sessions are logged out. When the sessions log in again, they are placed on active NP_Port interfaces in a balanced manner. • FLOGI-based—Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. On a link load rebalance, only the sessions that need to be moved to another link are logged out. When those sessions log in again, they are placed on active NP_Port interfaces in a balanced manner.
`Fabric WWN verification`	Fabric worldwide name (WWN) verification check state on the FCoE-FC gateway fabric: • Yes—Fabric WWN verification check is enabled. • No—Fabric WWN verification check is disabled.
`Auto load rebalance enabled`	Automated link load rebalancing configuration for the FCoE-FC gateway fabric: • No—Automated load balancing is disabled (default state). • Yes—Automated load balancing is enabled.

Table 562: show fibre-channel proxy fabric-state Output Fields *(Continued)*

Field Name	Field Description
`Last rebalance start-time`	Time that the last link load rebalance began on the FCoE-FC gateway fabric: • Never—The link load has never been rebalanced. • Timestamp value—Time the last link load rebalancing started.
`Last rebalance end-time`	Time that the last link load rebalance ended on the FCoE-FC gateway fabric: • Never—The link load has never been rebalanced. • Timestamp value—Time the last link load rebalancing ended.
`Last rebalance trigger`	Event that triggered the last link load rebalance on the FCoE-FC gateway fabric: • None—The link load has never been rebalanced. • Config-CLI—Configure (enable) automated load balancing. • Request-CLI—Rebalance requested from the CLI using the `request fibre-channel proxy load-rebalance fabric` *`fabric-name`* operational command. • Preview-CLI—Rebalancing *dry run* requested from the CLI using the `request fibre-channel proxy load-rebalance dry-run fabric` *`fabric-name`* operational command. Indicates that the switch completed the dry run. A dry run simulates a link load rebalance and displays a list of sessions that might be affected if you request an actual rebalance. • Link-up—New FC link (NP_Port) up on the FCoE-FC gateway fabric, which causes a rebalance to distribute sessions to the new link. • Restore-complete—If the FC process on the switch restarts, the switch attempts to restore the session state that existed before the restart. When automated rebalance is enabled, restore-complete indicates that the sessions have been restored and rebalanced.
`Last rebalance trigger-time`	Time that the last link load rebalance was triggered on the FCoE-FC gateway fabric: • Never—Link load rebalancing has never been triggered. • Timestamp value—Time the last link load rebalancing was triggered.

Table 562: show fibre-channel proxy fabric-state Output Fields *(Continued)*

Field Name	Field Description
`Last rebalance trigger-result`	Result of the last trigger event on the FCoE-FC gateway fabric: • Never—Link load rebalancing has never been triggered. • Not-configured—Automated rebalancing is not configured on the FCoE-FC gateway fabric. • Not-required—Last rebalance trigger did not require rebalancing the link load (the link load was already balanced across the active NP_Port links). • In-progress—Link load rebalancing is in progress and has not finished yet. • Restore-in-progress—The switch is recovering from an FC process restart and is in the process of restoring the sessions to the active NP_Port links. • Success—Link load rebalancing was successful. • Logged-out-all—All sessions have been logged out. • Preview-complete—The switch has finished simulating a dry run rebalancing request from the CLI (`request fibre-channel proxy load-rebalance dry-run fabric` *`fabric-name`* operational command) and reported the sessions that might be affected if you request an actual link load rebalance. • Fabric-deletion-in-progress—FCoE-FC gateway fabric is in the process of being deleted. **NOTE**: A trigger event does not necessarily result in a rebalance action. Link load rebalancing only occurs if the NP_Port interface session load is not balanced at the time of the trigger event.

Sample Output

show fibre-channel proxy fabric-state

```
user@switch> show fibre-channel proxy fabric-state
```

```
Fabric: san_fab1, Fabric-id: 10
Proxy load balance algorithm: Simple, Fabric WWN verification: Yes
```

```
Auto load rebalance enabled  : No
Last rebalance start-time    : Never
Last rebalance end-time      : Never
Last rebalance trigger       : Link-up
Last rebalance trigger-time  : Mon Sep 10 21:42:30 2012 usec: 814602
Last rebalance trigger-result: Not-configured

Fabric: san_fab2, Fabric-id: 20
Proxy load balance algorithm: ENode based, Fabric WWN verification: Yes
Auto load rebalance enabled  : No
Last rebalance start-time    : Never
Last rebalance end-time      : Never
Last rebalance trigger       : Link-up
Last rebalance trigger-time  : Mon Sep 17 17:23:35 2012 usec: 619684
Last rebalance trigger-result: Not-configured
```

show fibre-channel proxy fabric-state fabric

```
user@switch> show fibre-channel proxy fabric-state fabric fc_fabric_100
Fabric: fc_fabric_100, Fabric-id: 100
Proxy load balance algorithm: FLOGI based, Fabric WWN verification: No
Auto load rebalance enabled  : Yes
Last rebalance start-time    : Never
Last rebalance end-time      : Never
Last rebalance trigger       : Config-CLI
Last rebalance trigger-time  : Fri Nov  2 08:56:16 2012 usec: 004487
Last rebalance trigger-result: Not-required
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Monitoring Fibre Channel Interface Load Balancing

show fibre-channel proxy login-table | **19993**

show fibre-channel proxy np-port | **19997**

show fibre-channel proxy statistics | **20003**

show fibre-channel proxy login-table

IN THIS SECTION

- Syntax | 19993
- Description | 19993
- Options | 19993
- Required Privilege Level | 19993
- Output Fields | 19994
- Sample Output | 19995
- Release Information | 19996

Syntax

```
show fibre-channel proxy login-table
<brief | detail>
<fabric fabric-name>
<interface interface-name>
```

Description

Display Fibre Channel (FC) proxy fabric login table information.

Options

brief | detail	(Optional) Display the specified level of output.
fabric *fabric-name*	(Optional) Display output only for the specified fabric.
interface *interface-name*	(Optional) Display output only for the specified interface.

Required Privilege Level

view

Output Fields

Table 563 on page 19994 lists the output fields for the `show fibre-channel proxy login-table` command. Output fields are listed in the approximate order in which they appear.

Table 563: show fibre-channel proxy login-table Output Fields

Field Name	Field Description	Level of Output
`Fabric`	Name of the fabric.	All
`Fabric-id`	Fabric ID number.	All
`F-Port`	One of the following two values: • VF_Port interface connected to the Fibre Channel over Ethernet (FCoE) host, shown as the FCoE VLAN interface. • QFX Series FC port that is logged in to the FC switch, shown by a hyphen (-) to indicate that it is not the FCoE device VN_Port.	All
`FCID`	VN_Port Fibre Channel identifier provided by the Fibre Channel over Ethernet (FCoE) forwarder (FCF) or the Fibre Channel switch.	All
`Port-WWN`	Unique worldwide name (WWN) of the VN_Port.	All
`Node-WWN`	Unique WWN of the node hosting the VN_Ports.	`detail`
`NP-Port`	NP_Port interface connected to the FCoE forwarder (FCF) or the Fibre Channel switch.	All
`Class`	FLOGI service class.	`detail`

Table 563: show fibre-channel proxy login-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Fabric port WWN	Unique WWN of the fabric port (VF_Port).	detail
Fabric WWN	Unique WWN of the fabric generated by the FCF.	detail

Sample Output

show fibre-channel proxy login-table

```
user@switch> show fibre-channel proxy login-table
```

```
Fabric: proxy2, Fabric-id: 200
F-Port                 FCID     Port-WWN                NP-Port
-                      0x030000 10:00:00:15:17:a9:98:64 fc-0/0/0.0
vlan.100               0x030001 20:00:10:94:00:01:00:01 fc-0/0/0.0
vlan.100               0x030002 20:00:10:94:00:01:00:05 fc-0/0/0.0
vlan.100               0x030003 20:00:10:94:00:02:00:01 fc-0/0/0.0
vlan.100               0x030004 20:00:10:94:00:02:00:05 fc-0/0/0.0
```

show fibre-channel proxy login-table detail

```
user@switch> show fibre-channel proxy login-table detail
Fabric: proxy2, Fabric-id: 200

    FCID:             0x030000
    F-Port:           -
    NP-Port:          fc-0/0/0.0
    Port WWN:         10:00:00:15:17:a9:98:64
    Node WWN:         20:c8:11:22:33:44:55:66
    Class:            3
    Fabric port WWN:  10:00:00:15:17:a9:99:48
    Fabric WWN:       00:0a:df:ff:0b:11:22:34
```

```
FCID:             0x030001
F-Port:           vlan.100
NP-Port:          fc-0/0/0.0
Port WWN:         20:00:10:94:00:01:00:01
Node WWN:         10:00:10:94:00:00:00:01
Class:            3
Fabric port WWN:  10:00:00:15:17:a9:99:48
Fabric WWN:       00:0a:df:ff:0b:11:22:34

FCID:             0x030002
F-Port:           vlan.100
NP-Port:          fc-0/0/0.0
Port WWN:         20:00:10:94:00:01:00:05
Node WWN:         10:00:10:94:00:00:00:01
Class:            3
Fabric port WWN:  10:00:00:15:17:a9:99:48
Fabric WWN:       00:0a:df:ff:0b:11:22:34

FCID:             0x030003
F-Port:           vlan.100
NP-Port:          fc-0/0/0.0
Port WWN:         20:00:10:94:00:02:00:01
Node WWN:         10:00:10:94:00:00:00:02
Class:            3
Fabric port WWN:  10:00:00:15:17:a9:99:48
Fabric WWN:       00:0a:df:ff:0b:11:22:34

FCID:             0x030004
F-Port:           vlan.100
NP-Port:          fc-0/0/0.0
Port WWN:         20:00:10:94:00:02:00:05
Node WWN:         10:00:10:94:00:00:00:02
Class:            3
Fabric port WWN:  10:00:00:15:17:a9:99:48
Fabric WWN:       00:0a:df:ff:0b:11:22:34
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring an FCoE-FC Gateway Fibre Channel Fabric

show fibre-channel proxy fabric-state | **19987**

show fibre-channel proxy np-port | **19997**

show fibre-channel proxy statistics | **20003**

show fibre-channel proxy np-port

IN THIS SECTION

- Syntax | **19997**
- Description | **19997**
- Options | **19998**
- Required Privilege Level | **19998**
- Output Fields | **19998**
- Sample Output | **20000**
- Release Information | **20002**

Syntax

```
show fibre-channel proxy np-port
<brief | detail>
<fabric fabric-name>
<interface interface-name>
```

Description

Display Fibre Channel gateway fabric proxy Node Port (NP_Port) information.

Options

brief | detail	(Optional) Display the specified level of output.
fabric *fabric-name*	(Optional) Display output only for the specified fabric.
interface *interface-name*	(Optional) Display output only for the specified interface.

Required Privilege Level

view

Output Fields

Table 564 on page 19998 lists the output fields for the `show fibre-channel proxy np-port` command. Output fields are listed in the approximate order in which they appear.

Table 564: show fibre-channel proxy np-port Output Fields

Field Name	Field Description	Level of Output
Fabric	Name of the fabric.	All
Fabric-id	Fabric ID number.	All
NP-Port	NP_Port interface connected to the FCoE forwarder (FCF) or the Fibre Channel switch.	All
State	FCID state of the NP_Port interface.	All
Sessions	Number of active sessions on the NP_Port interface. A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions.	All

Table 564: show fibre-channel proxy np-port Output Fields *(Continued)*

Field Name	Field Description	Level of Output

(Junos OS Release 18.1R1 and beyond)
`Supported max login sessions`
 	
(Junos OS Releases before 18.1R1)
`Configured max login sessions`
	Maximum number of FIP login sessions supported on the NP_Port interface.
	
	NOTE: If the fabric is configured to have a limit on the maximum number of FIP login sessions that is less than the configured maximum for the individual NP_Port displayed by this field, then the actual number of sessions permitted on the NP_Port will be constrained by the fabric limit. See "max-login-sessions" on page 10442 for more information.
		`detail`

`Enodes`	Number of ENodes with sessions on the NP_Port.	`detail`
`LB state`	Load-balancing state: • `On`—Load balancing is on • `Off`—Load balancing is off.	All
`LB weight`	Load balance weight, which reflects the port speed: • `2`—Port speed is 2 Gbps. • `4`—Port speed is 4 Gbps. • `8`—Port speed is 8 Gbps.	All
`Ref-count`	Reference count internal to Junos OS.	`detail`

Table 564: show fibre-channel proxy np-port Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Flags internal to Junos OS. **NOTE**: When an NP_Port interface reaches its configured maximum number of FIP sessions, the `Flags` field displays the flag `MAX-LOGINS-REACHED`.	`detail`

Sample Output

show fibre-channel proxy np-port

```
user@switch> show fibre-channel proxy np-port
```

```
Fabric: proxy1, Fabric-id: 10
NP-Port     State              Sessions       LB state   LB weight
fc-0/0/0.0  online             3              ON         4
fc-0/0/1.0  online             3              ON         4
fc-0/0/2.0  online             3              ON         4
root@junos1> show fibre-channel proxy np-port detail
```

show fibre-channel proxy np-port detail (Junos OS Release 18.1R1 and beyond)

```
user@switch> show fibre-channel proxy np-port detail
Fabric: my-fabric, Fabric-id: 100

    NP-Port:                        fc-0/2/0.0
    State:                          online
    Sessions:                       2
    Supported max login sessions:   2500
    Enodes:                         1
    LB state:                       ON
```

```
    LB weight:                      10
    Ref-count:                      1
    Flags:                          UP LB C3

    NP-Port:                        fc-0/2/1.0
    State:                          online
    Sessions:                       2
    Supported max login sessions:  2500
    Enodes:                         1
    LB state:                       ON
    LB weight:                      10
    Ref-count:                      1
    Flags:                          UP LB C3

    NP-Port:                        fc-0/2/2.0
    State:                          online
    Sessions:                       2
    Supported max login sessions:  2500
    Enodes:                         1
    LB state:                       ON
    LB weight:                      10
    Ref-count:                      1
    Flags:                          UP LB C3
```

show fibre-channel proxy np-port detail (Junos OS Releases before 18.1R1)

```
user@switch> show fibre-channel proxy np-port detail
Fabric: proxy1, Fabric-id: 10

    NP-Port:                         fc-0/0/0.0
    State:                           online
    Sessions:                        3
    Configured max login sessions:  130
    Enodes                           1
    LB state:                        ON
    LB weight:                       4
    Ref-count:                       4
    Flags:                           UP LB

    NP-Port:                         fc-0/0/1.0
```

```
    State:                          online
    Sessions:                       3
    Configured max login sessions:  130
    Enodes                          2
    LB state:                       ON
    LB weight:                      4
    Ref-count:                      4
    Flags:                          UP LB

    NP-Port:                        fc-0/0/2.0
    State:                          online
    Sessions:                       130
    Configured max login sessions:  130
    Enodes                          17
    LB state:                       OFF
    LB weight:                      4
    Ref-count:                      131
    Flags:                          UP MAX-LOGINS-REACHED
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring an FCoE-FC Gateway Fibre Channel Fabric

Monitoring Fibre Channel Interface Load Balancing

show fibre-channel proxy fabric-state | **19987**

show fibre-channel proxy login-table | **19993**

show fibre-channel proxy statistics | **20003**

show fibre-channel proxy statistics

IN THIS SECTION

- Syntax | 20003
- Description | 20003
- Options | 20003
- Required Privilege Level | 20003
- Output Fields | 20003
- Sample Output | 20005
- Release Information | 20005

Syntax

```
show fibre-channel proxy statistics
<fabric fabric-name>
```

Description

Display Fibre Channel proxy fabric statistics.

Options

fabric ***fabric-name***	(Optional) Display output only for the specified fabric.

Required Privilege Level

view

Output Fields

Table 565 on page 20004 lists the output fields for the `show fibre-channel proxy statistics` command. Output fields are listed in the approximate order in which they appear.

Table 565: show fibre-channel proxy statistics Output Fields

Field Name		Field Description
Fabric		Name of the fabric.
Fabric-id		Fabric ID number.
NP-Port Transmit Command Statistics		Transmitted command statistics for the NP_Port.
	• Command	Type of command issued on the NP_Port: • FLOGI—Fabric login commands issued. • FDISC—Fabric discovery commands issued. • LOGO—Logout commands issued. • Others—Other commands issued.
	• Tx	Number of times the command type was transmitted.
	• Rx-ACC	Number of times the NP_Port transmitted a receive accept message for the command type.
	• Rx-RJT	Number of times the NP_Port transmitted a receive reject message for the command type.
	• Abort	Number of times the NP_Port transmitted a termination message for the command type.
NP-Port Receive Command Statistics		Received command statistics for the NP_Port.
	• Command	The type of command received on the NP_Port: • LOGO—Logout commands issued. • Others—Other commands issued.

Table 565: show fibre-channel proxy statistics Output Fields *(Continued)*

Field Name		Field Description
	• Rx	Number of times the command type was received.
	• Tx-ACC	Number of times the NP_Port received a transmit accept message for the command type.
	• Tx-RJT	Number of times the NP_Port received a transmit reject message for the command type.
	• Abort	Number of times the NP_Port received a termination message for the command type.

Sample Output

show fibre-channel proxy statistics

```
user@switch> show fibre-channel proxy statistics
Fabric: proxy1, Fabric-id: 10

NP-Port Transmit Command Statistics:
Command         Tx        Rx-ACC    Rx-RJT    Abort
FLOGI           3         3         0         0
FDISC           3         3         0         0
LOGO            0         0         0         0
Others          0         0         0         0

NP-Port Receive Command Statistics:
Command         Rx        Tx-ACC    Tx-RJT    Abort
LOGO            0         0         0         0
Others          0         0         0         0
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring an FCoE-FC Gateway Fibre Channel Fabric

show fibre-channel proxy fabric-state | **19987**

show fibre-channel proxy login-table | **19993**

show fibre-channel proxy np-port | **19997**

clear fibre-channel proxy statistics | **3603**

show fibre-channel routes

IN THIS SECTION

- Syntax | **20006**
- Description | **20006**
- Options | **20006**
- Required Privilege Level | **20007**
- Output Fields | **20007**
- Sample Output | **20007**
- Release Information | **20008**

Syntax

```
show fibre-channel routes
<fabric fabric-name>
```

Description

Display Fibre Channel route information.

Options

fabric ***fabric-name***	(Optional) Display output only for the specified fabric.

Required Privilege Level

view

Output Fields

Table 566 on page 20007 lists the output fields for the `show fibre-channel routes` command. Output fields are listed in the approximate order in which they appear.

Table 566: show fibre-channel routes Output Fields

Field Name	Field Description
`Fabric`	Name of the fabric.
`Route-prefix`	Route destination.
`State`	State of the NP_Port interface: • `Active`—The interface is online. • `Deleted`—The interface is deleted.
`Interface`	Name of the interface.
`Mac-Address`	Media access control (MAC) address of the interface.
`Index`	Next-hop index identifier.
`Flags`	Flags internal to Junos OS.

Sample Output

show fibre-channel routes

```
user@switch> show fibre-channel routes
```

```
Fabric: proxy2
Route-prefix       State     Interface     Mac-Address         Index   Flags
0x030000/24        Active    fc-0/0/0.0    00:15:17:a9:98:64   674     kernel
0x030001/24        Active    vlan.100      0e:fc:00:03:00:01   675     kernel
0x030002/24        Active    vlan.100      0e:fc:00:03:00:02   676     kernel
0x030003/24        Active    vlan.100      0e:fc:00:03:00:03   677     kernel
0x030004/24        Active    vlan.100      0e:fc:00:03:00:04   678     kernel
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel next-hops | 19985

show route forwarding-table family fibre-channel | 23673

show fip snooping

IN THIS SECTION

- Syntax | 20009
- Description | 20009
- Options | 20009
- Required Privilege Level | 20009
- Output Fields | 20009
- Sample Output | 20012
- Release Information | 20015

Syntax

```
show fip snooping
<brief | detail>
```

Description

Display FIP snooping information.

Options

none	Display FIP snooping information.
brief | detail	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 567 on page 20009 lists the output fields for the `show fip snooping` command. Output fields are listed in the approximate order in which they appear.

Table 567: show fip snooping Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All

Table 567: show fip snooping Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Mode`	(QFX Series only) Snooping mode enabled on the FCoE VLAN: • VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port. • VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports.	All
`FC-MAP`	FCoE mapped address prefix of the FCoE forwarder for the VLAN.	All
`FCF` or `FCF-MAC`	MAC address of the FCF.	All
`Session Count` or `Active Sessions`	Current number of virtual link sessions with VN_Ports.	All
`VN_Port Count`	(QFX Series only) Number of VN_Ports active on an ENode.	`brief`
`Configured FKA-ADV`	FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.	`detail`

Table 567: show fip snooping Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Running FKA-ADV`	Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV. For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.	`detail`
`Beacon Period`	(QFX Series only) Beacon period interval in milliseconds.	`detail`
`VN2VN Mode`	(QFX Series only) Mode of VN2VN_Port snooping: • Multi-Point—Multiple ENodes are connected to the network and form multiple virtual links. Each virtual link is created between one pair of VN_Ports. This is analogous to the loop mode in traditional FC networks. • Point-to-Point—Two ENodes are connected to the network and form a single VN_Port to VN_Port virtual link. This is analogous to the point-to-point FC link between an FC initiator and an FC target.	`detail`
`ENode-MAC`	MAC address of the connected FCoE node (ENode).	All
`Interface`	Interface connected to the ENode. (QFabric System or Junos Fusion `satellite` command output only) When an FCoE LAG has been configured, this field displays both the LAG interface and the LAG member interface connected to the ENode.	`detail`

Table 567: show fip snooping Output Fields *(Continued)*

Field Name	Field Description	Level of Output
VN-Port MAC	MAC address of a VN_Port on the ENode.	All
FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.	detail
Active VN_Ports	(QFX Series only) Number of VN_Ports active on an ENode.	detail
Vlink far-end VN-Port-MAC	(QFX Series only) Media access control (MAC) address of the VN_Port at the other end of the virtual link.	detail

Sample Output

show fip snooping

```
user@switch> show fip snooping
VLAN : fcoevlan1     FC-MAP : 0e:fc:00
  FCF : 00:10:94:00:00:01   Session Count : 2
    Enode-MAC : 00:10:94:00:00:02
      VN-Port-MAC : 0E:FC:00:01:00:05
      VN-Port-MAC : 0E:FC:00:01:00:01
```

show fip snooping brief (QFX Series)

```
user@switch> show fip snooping brief
VLAN: vlan100,    Mode: VN2VF Snooping
  FC-MAP: 0e:fc:00
  FCF: 30:10:94:01:00:00    Session Count: 2
    Enode-MAC: 10:10:94:01:00:01
```

```
      VN-Port-MAC: 0e:fc:00:01:0d:01
      VN-Port-MAC: 0e:fc:00:01:0e:01
VLAN: vlan101,    Mode: VN2VN Snooping
  FC-MAP: 0e:fc:00
    Enode-MAC: 10:10:94:01:00:02 VN_Port count: 1
      VN-Port-MAC: 0e:fc:00:01:0a:01  Session Count: 2
    Enode-MAC: 10:10:94:01:00:03 VN_Port count: 0
```

show fip snooping detail (QFX Series Switches)

```
user@switch> show fip snooping detail
root@sw-pa02v> show fip snooping detail
VLAN: vlan100,  Mode: VN2VF Snooping
  FC-MAP: 0e:fc:00
  FCF Information
  FCF-MAC            : 30:10:94:01:00:00
  Active Sessions    : 2
  Configured FKA-ADV : 258
  Running FKA-ADV    : 188
    Enode Information
    Enode-MAC: 10:10:94:01:00:01,       Interface: xe-0/0/10
    Configured FKA-ADV : 258
    Running FKA-ADV    : 230
      Session Information
      VN-Port MAC: 0e:fc:00:01:0d:01,   FKA-ADV : 230
      VN-Port MAC: 0e:fc:00:01:0e:01,   FKA-ADV : 245

VLAN: vlan101,  Mode: VN2VN Snooping
  FC-MAP: 0e:fd:00
  Beacon_Period:  90000
  VN2VN Mode: Multi-Point
    Enode Information
    Enode-MAC: 10:10:94:01:00:02,       Interface: xe-0/0/10
      Active VN_Ports : 1
      VN_Port Information
      VN-Port MAC: 0e:fd:00:01:0a:01
         Active Sessions    : 2
         Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:01:0b:01
        Vlink far-end VN-Port-MAC: 0e:fd:00:01:0c:01
```

```
    Enode-MAC: 10:10:94:01:00:02,       Interface: xe-0/0/11
      Active VN_Ports : 0
```

show fip snooping detail (QFabric System FCoE with LAG Configured)

```
admin@qfabric> show fip snooping detail
VLAN: vlan_100, Mode: VN2VF Snooping
FC-MAP: 0e:fc:00
  FCF Information
  FCF-MAC            : 84:18:88:d1:f5:cc
  Active Sessions    : 2
  Configured FKA-ADV : 8000
  Running FKA-ADV    : 23962
    Enode Information
    Enode-MAC: 00:c0:dd:14:ae:6d,       Interface: P4546-C:ae0  P4546-C:xe-0/0/39
    Configured FKA-ADV : 8000
    Running FKA-ADV    : 16622
      Session Information
      VN-Port MAC: 0e:fc:00:6c:06:a5,   FKA-ADV : 246303
    Enode Information
    Enode-MAC: 00:c0:dd:14:ae:6f,       Interface: P4546-C:ae0  P4546-C:xe-0/0/38
    Configured FKA-ADV : 8000
    Running FKA-ADV    : 16512
      Session Information
      VN-Port MAC: 0e:fc:00:6c:06:a4,   FKA-ADV : 238150
```

show fip snooping detail (EX Series Switches)

```
user@switch> show fip snooping detail
VLAN : fcoevlan1    FC-MAP : 0e:fc:00
  FCF Information
  FCF-MAC            : 00:10:94:00:00:01
  Active Sessions    : 2
  Configured FKA-ADV : 258
  Running FKA-ADV    : 244
    Enode Information
    Enode-MAC : 00:10:94:00:00:02       Interface : xe-0/0/1
    Configured FKA-ADV : 258
    Running FKA-ADV    : 248
      Session Information
```

```
        VN-Port MAC : 0E:FC:00:01:00:05   FKA-ADV : 264
        VN-Port MAC : 0E:FC:00:01:00:01   FKA-ADV : 260
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

Configuring an FCoE LAG

Example: Configuring an FCoE LAG on a Redundant Server Node Group

show fip snooping enode | **20015**

show fip snooping fcf | **20021**

show fip snooping statistics | **20029**

show fip snooping vlan | **20034**

show fip snooping interface | **20024**

show fip snooping enode

IN THIS SECTION

- Syntax | **20016**
- Description | **20016**
- Options | **20016**
- Required Privilege Level | **20016**
- Output Fields | **20016**
- Sample Output | **20019**
- Release Information | **20020**

Syntax

```
show fip snooping enode enode-mac
<brief | detail>
<vlan vlan-name>
```

Description

Display FIP snooping FCoE node (ENode) information.

Options

brief | detail (Optional) Display the specified level of output.

enode-mac Display information for the ENode specified by the MAC address.

vlan *vlan-name* (Optional) Display FIP snooping information for the ENode on only the specified VLAN.

Required Privilege Level

view

Output Fields

Table 568 on page 20016 lists the output fields for the `show fip snooping enode` command. Output fields are listed in the approximate order in which they appear.

Table 568: show fip snooping enode Output Fields

Field Name	Field Description	Level of Output
ENode and ENode MAC	MAC address of the ENode.	All
VLAN	Name of the VLAN.	All
Interface	Interface connected to the ENode.	All

Table 568: show fip snooping enode Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Mode`	(QFX Series only) Snooping mode enabled on the FCoE VLAN: • VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port. • VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports.	All
`VN_Port Count`	(QFX Series only) Number of VN_Ports active on an ENode.	`brief`
`Session Count`	Current number of virtual link sessions with VN_Ports.	All
`Configured FKA-ADV`	FIP keepalive interval in seconds configured on the FCoE forwarder (FCF) multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant. For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.	`detail`

Table 568: show fip snooping enode Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Running FKA-ADV`	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF. For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.	`detail`
`VN-Port` or `VN-Port-MAC`	MAC address of a VN_Port on the ENode.	All
`FKA-ADV`	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.	`detail`
`FCF` or `FCF-MAC`	MAC address of the FCF to which the VN_Port is connected.	All
`Beacon Period`	(QFX Series only) Beacon period interval in milliseconds.	`detail`

Table 568: show fip snooping enode Output Fields *(Continued)*

Field Name	Field Description	Level of Output
VN2VN Mode	(QFX Series only) Mode of VN2VN_Port snooping: • Multi-Point—Multiple ENodes are connected to the network and form multiple virtual links. Each virtual link is created between one pair of VN_Ports. This is analogous to the loop mode in traditional FC networks. • Point-to-Point—Two ENodes are connected to the network and form a single VN_Port to VN_Port virtual link. This is analogous to the point-to-point FC link between an FC initiator and an FC target.	detail
Vlink far-end VN-Port-MAC	(QFX Series only) Media access control (MAC) address of the VN_Port at the other end of the virtual link.	detail

Sample Output

show fip snooping enode

```
user@switch> show fip snooping enode 00:10:94:00:00:02
Enode : 00:10:94:00:00:02   VLAN : vlan1   Interface : xe-0/0/1
        VN-Port-MAC           FCF-MAC
        0E:FC:00:00:00:05     00:10:94:00:00:01
        0E:FC:00:00:00:01     00:10:94:00:00:01
```

show fip snooping enode brief (QFX Series)

```
user@switch> show fip snooping enode 10:10:94:01:00:02 brief
Enode: 10:10:94:01:00:02 ,  VLAN: vlan101,  Interface: xe-0/0/10
  Mode: VN2VF Snooping      VN_Port Count: 1
        VN_Port Information
        VN_Port Mac: 0e:fc:00:01:0a:01     Session Count: 2
```

show fip snooping enode detail (QFX Series)

```
user@switch> show fip snooping enode 10:10:94:01:00:02 detail
Enode MAC: 10:10:94:01:00:02,   VLAN: vlan101,  Interface: xe-0/0/10
  Mode: VN2VF Snooping       VN_Port Count: 1
  Beacon_Period: 90000       VN2VN Mode: Multi-Point
        VN_Port Information
        VN_Port Mac: 0e:fc:00:01:0a:01     Session Count: 2
        Vlink far-end VN-Port-MAC: 0e:fc:00:01:0b:01
        Vlink far-end VN-Port-MAC: 0e:fc:00:01:0c:01
```

show fip snooping enode detail

```
user@switch> show fip snooping enode 00:10:94:00:00:02 detail
Enode MAC : 00:10:94:00:00:02   VLAN : vlan1    Interface : xe-0/0/1
Configured FKA-ADV : 258        Running FKA-ADV : 213
  Session Information
  VN-Port : 0E:FC:00:00:00:05   FKA-ADV : 229   FCF : 00:10:94:00:00:01
  VN-Port : 0E:FC:00:00:00:01   FKA-ADV : 225   FCF : 00:10:94:00:00:01
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

Example: Configuring an FCoE Transit Switch

show fip snooping | **20008**

show fip snooping fcf | **20021**

show fip snooping statistics | **20029**

show fip snooping vlan | **20034**

show fip snooping interface | **20024**

show fip snooping fcf

IN THIS SECTION

- Syntax | 20021
- Description | 20021
- Options | 20021
- Required Privilege Level | 20021
- Output Fields | 20022
- Sample Output | 20023
- Release Information | 20024

Syntax

```
show fip snooping fcf fcf-mac
<brief | detail>
<vlan vlan-name>
```

Description

Display FIP snooping FCoE forwarder (FCF) information.

Options

brief | detail	(Optional) Display the specified level of output.
fcf-mac	Display information for the FCF specified by the MAC address.
vlan-name	(Optional) Display FIP snooping information for the FCF on only the specified VLAN.

Required Privilege Level

view

Output Fields

Table 569 on page 20022 lists the output fields for the `show fip snooping fcf` command. Output fields are listed in the approximate order in which they appear.

Table 569: show fip snooping fcf Output Fields

Field Name		Field Description	Level of Output
FCF or FCF-MAC		MAC address of the FCoE forwarder.	All
VLAN		Name of the VLAN.	All
Session Count		Current number of virtual link sessions with VN_Ports.	None
Configured FKA-ADV		FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258.	detail
Running FKA-ADV		Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV.	detail
ENode-MAC		MAC address of the connected ENode.	All
	• Interface	Interface connected to the ENode.	detail
	• Configured FKA-ADV	FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant.	detail

Table 569: show fip snooping fcf Output Fields *(Continued)*

Field Name		Field Description	Level of Output
	• Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF.	detail
	• VN-Port MAC	MAC address of a VN_Port on the ENode.	All
	• FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.	detail

Sample Output

show fip snooping fcf

```
user@switch> show fip snooping fcf 00:10:94:00:00:01
FCF : 00:10:94:00:00:01   VLAN : vlan1   Session Count : 2
  Enode-MAC : 00:10:94:00:00:02
    VN-Port-MAC : 0E:FC:00:00:00:05
    VN-Port-MAC : 0E:FC:00:00:00:01
```

show fip snooping fcf detail

```
user@switch> show fip snooping fcf 00:10:94:00:00:01 detail
FCF-MAC : 00:10:94:00:00:01     VLAN : vlan1
Configured FKA-ADV : 258        Running FKA-ADV : 222
  Enode Information
  Enode-MAC : 00:10:94:00:00:02 Interface: xe-0/0/1
  Configured FKA-ADV : 258
  Running FKA-ADV    : 226
```

```
    Session Information
    VN-Port MAC : 0E:FC:00:00:00:05     FKA-ADV : 242
    VN-Port MAC : 0E:FC:00:00:00:01     FKA-ADV : 238
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- *Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch*
- Example: Configuring an FCoE Transit Switch
- show fip snooping | 20008
- show fip snooping enode | 20015
- show fip snooping statistics | 20029
- show fip snooping vlan | 20034
- show fip snooping interface | 20024

show fip snooping interface

IN THIS SECTION

- Syntax | 20025
- Description | 20025
- Options | 20025
- Required Privilege Level | 20025
- Output Fields | 20025
- Sample Output | 20028
- Release Information | 20028

Syntax

```
show fip snooping interface interface-name
<brief | detail>
```

Description

Display FIP snooping information for the specified interface.

Options

brief | detail	(Optional) Display the specified level of output.
interface-name	Display information for the specified interface.

Required Privilege Level

view

Output Fields

Table 570 on page 20025 lists the output fields for the `show fip snooping interface` *`interface-name`* command. Output fields are listed in the approximate order in which they appear.

Table 570: show fip snooping interface Output Fields

Field Name	Field Description	Level of Output
`VLAN`	Name of the VLAN.	All
`FC-MAP`	FCoE mapped address prefix of the FCoE forwarder for the VLAN.	All
`FCF` or `FCF-MAC`	MAC address of the FCF.	All
`Session Count` or `Active Sessions`	Current number of virtual link sessions with VN_Ports.	All

Table 570: show fip snooping interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Configured FKA-ADV`	FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.	`detail`
`Running FKA-ADV`	Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV. For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.	`detail`
`ENode-MAC`	MAC address of the connected FCoE node (ENode).	All
`Interface`	Interface connected to the ENode.	`detail`

Table 570: show fip snooping interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Configured FKA-ADV`	FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant. For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.	`detail`
`Running FKA-ADV`	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF. For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.	`detail`
`VN-Port MAC`	MAC address of a VN_Port on the ENode.	All
`FKA-ADV`	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.	`detail`

Sample Output

show fip snooping interface

```
user@switch> show fip snooping interface xe-0/0/9.0
VLAN: vlan_100,    FC-MAP: 0e:fc:00
  FCF: 30:10:94:01:00:00    Session Count: 1
    Enode-MAC: 10:10:94:01:00:01
      VN-Port-MAC: 0e:fc:00:01:0a:01
```

show fip snooping interface detail

```
user@switch> show fip snooping interface xe-0/0/9.0 detail
VLAN: vlan_100, FC-MAP: 0e:fc:00
  FCF Information
  FCF-MAC            : 30:10:94:01:00:00
  Active Sessions    : 1
  Configured FKA-ADV : 368640000
  Running FKA-ADV    : 0
    Enode Information
    Enode-MAC: 10:10:94:01:00:01,       Interface: xe-0/0/9
    Configured FKA-ADV : 368640000
    Running FKA-ADV    : 0
      Session Information
      VN-Port MAC: 0e:fc:00:01:0a:01,   FKA-ADV : 0
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

show fip snooping | **20008**

show fip snooping enode | **20015**

show fip snooping fcf | **20021**

show fip snooping statistics | 20029
show fip snooping vlan | 20034

show fip snooping statistics

IN THIS SECTION

- Syntax | 20029
- Syntax (Junos Fusion) | 20029
- Description | 20029
- Options | 20030
- Required Privilege Level | 20030
- Output Fields | 20030
- Sample Output | 20032
- Release Information | 20033

Syntax

```
show fip snooping statistics
<vlan vlan-name>
```

Syntax (Junos Fusion)

```
show fip snooping satellite statistics
<vlan vlan-name>
```

Description

Display FIP snooping statistics.

The command syntax in a Junos Fusion environment includes the `satellite` keyword and displays FIP snooping statistics for satellite device FCoE VLANs, which have FCoE and FIP functions and status

information distributed between the aggregation devices and satellite devices. The command validates that a specified VLAN is a satellite FCoE VLAN, and displays an error message if the satellite syntax is not used for a satellite FCoE VLAN, or if the satellite syntax is used with a VLAN that is not a satellite FCoE VLAN.

Options

vlan *vlan-name* (Optional) Display FIP snooping statistics for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 571 on page 20030 lists the output fields for the `show fip snooping statistics` command. Output fields are listed in the approximate order in which they appear.

Table 571: show fip snooping statistics Output Fields

Field Name	Field Description
VLAN	Name of the VLAN for which a set of statistics is displayed.
Mode	(QFX Series only) Snooping mode enabled on the FCoE VLAN: • VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port. • VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports.
Number of MDS	Number of multicast discovery solicitation messages sent on the VLAN.
Number of UDS	Number of unicast discovery solicitation messages sent on the VLAN.
Number of FLOGI	Number of fabric logins on the VLAN.

Table 571: show fip snooping statistics Output Fields *(Continued)*

Field Name	Field Description
`Number of FDISC`	Number of fabric discovery logins on the VLAN.
`Number of LOGO`	Number of fabric logouts on the VLAN.
`Number of ENode-keep-alive`	Number of ENode keepalive messages sent on the VLAN.
`Number of VNPort-keep-alive`	Number of VN_Port keepalive messages sent on the VLAN.
`Number of MDA`	Number of multicast discovery advertisement messages sent on the VLAN.
`Number of UDA`	Number of unicast discovery advertisement messages sent on the VLAN.
`Number of FLOGI_ACC`	Number of fabric logins accepted on the VLAN.
`Number of FLOGI_RJT`	Number of fabric logins rejected on the VLAN.
`Number of FDISC_ACC`	Number of fabric discoveries accepted on the VLAN.
`Number of FDISC_RJT`	Number of fabric discoveries rejected on the VLAN.
`Number of LOGO_ACC`	Number of fabric logouts accepted on the VLAN.
`Number of LOGO_RJT`	Number of fabric logouts rejected on the VLAN.
`Number of CVL`	Number of clear virtual links (CVL) actions on the VLAN.
`Number of VN_Port Probes Req`	(QFX Series only) Number of multicast N_Port_ID probes sent to the ALL-VN2VN-ENode-MACs multicast address on the VLAN.

Table 571: show fip snooping statistics Output Fields *(Continued)*

Field Name	Field Description
`Number of VN_Port Claim Notif`	(QFX Series only) Number of multicast N_Port_ID claim notifications sent on the VLAN.
`Number of VN_Port Beacons`	(QFX Series only) Number of multicast beacons sent on the VLAN.
`Number of VN_Port Probes Reply`	(QFX Series only) Number of replies to N_Port_ID probes sent on the VLAN. Replies are unicast to the ENode MAC address of the probe requester.
`Number of VN_Port Claim Reply`	(QFX Series only) Number of replies to N_Port_ID claim notifications sent on the VLAN. Replies are unicast to the ENode MAC address of the claim notifier.

Sample Output

show fip snooping statistics (FIP Snooping)

```
user@switch> show fip snooping statistics
VLAN: fcoevlan1      Mode: VN2VF Snooping

    Number of MDS:               2
    Number of UDS:               2
    Number of FLOGI:             2
    Number of FDISC:             2
    Number of LOGO:              0
    Number of Enode-keep-alive:  200
    Number of VNPort-keep-alive: 200

    Number of MDA:               25
    Number of UDA:               2
    Number of FLOGI_ACC:         2
    Number of FLOGI_RJT:         0
    Number of FDISC_ACC:         2
    Number of FDISC_RJT:         0
```

```
    Number of LOGO_ACC:          0
    Number of LOGO_RJT:          0
    Number of CVL:               0
```

show fip snooping statistics (VN2VN_Port Snooping)

```
user@switch> show fip snooping statistics
VLAN: vlan101   Mode: VN2VN Snooping

Number of VN_Port Probes Req:        3
Number of VN_Port Claim Notif:       3
Number of VN_Port Beacons:           0

Number of VN_Port Probes Reply:      3
Number of VN_Port Claim Reply:       3
Number of FLOGI:                     0
Number of FLOGI_ACC:                 0
Number of FLOGI_RJT:                 0
Number of FDISC:                     0
Number of FDISC_ACC:                 0
Number of FDISC_RJT:                 0
Number of LOGO:                      0
Number of LOGO_ACC:                  0
Number of LOGO_RJT:                  0
```

Release Information

Command introduced in Junos OS Release 10.4.

Command introduced with satellite keyword in Junos OS Release 17.4R1 for Junos Fusion Data Center.

RELATED DOCUMENTATION

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

Example: Configuring an FCoE Transit Switch

show fip snooping | **20008**

show fip snooping enode | **20015**

show fip snooping fcf | **20021**

show fip snooping vlan | 20034
show fip snooping interface | 20024

show fip snooping vlan

IN THIS SECTION

- Syntax | 20034
- Description | 20034
- Options | 20034
- Required Privilege Level | 20035
- Output Fields | 20035
- Sample Output | 20037
- Release Information | 20039

Syntax

```
show fip snooping vlan vlan-name
<brief | detail>
```

Description

Display FIP snooping VLAN information.

Options

brief | detail	(Optional) Display the specified level of output.
vlan-name	Display information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 572 on page 20035 lists the output fields for the `show fip snooping vlan` command. Output fields are listed in the approximate order in which they appear.

Table 572: show fip snooping vlan Output Fields

Field Name	Field Description	Level of Output
`VLAN`	Name of the VLAN.	All
`Mode`	(QFX Series only) Snooping mode enabled on the FCoE VLAN: • VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port. • VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports.	All
`VN_Port count`	(QFX Series only) Number of VN_Ports active on an ENode when the mode is VN2VN_Port FIP snooping.	
`FC-MAP`	FCoE mapped address prefix of the FCoE forwarder for the VLAN.	All
`Beacon_Period`	(QFX Series only) Beacon period interval in milliseconds.	`detail`

Table 572: show fip snooping vlan Output Fields *(Continued)*

Field Name		Field Description	Level of Output
`VN2VN Mode`		(QFX Series only) Mode of VN2VN_Port snooping: • Multi-Point—Multiple ENodes are connected to the network and form multiple virtual links. Each virtual link is created between one pair of VN_Ports. This is analogous to the loop mode in traditional FC networks. • Point-to-Point—Two ENodes are connected to the network and form a single VN_Port to VN_Port virtual link. This is analogous to the point-to-point FC link between an FC initiator and an FC target.	`detail`
`FCF` or `FCF-MAC`		MAC address of the FCF.	All
`Session Count` or `Active Sessions`		Current number of virtual link sessions with VN_Ports.	All
`Configured FKA-ADV`		FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258.	`detail`
`Running FKA-ADV`		Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV.	`detail`
`ENode-MAC`		MAC address of the connected ENode.	All
	• `Interface`	Interface connected to the ENode.	`detail`

Table 572: show fip snooping vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
• Configured FKA-ADV	FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant.	detail
• Running FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF.	detail
• VN-Port MAC	MAC address of a VN_Port on the ENode.	All
• FKA-ADV	Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.	detail
• Active VN_Ports	(QFX Series only) Number of VN_Ports active on an ENode.	detail
• Vlink far-end VN-Port-MAC	(QFX Series only) Media access control (MAC) address of the VN_Port at the other end of the virtual link.	detail

Sample Output

show fip snooping vlan

```
user@switch> show fip snooping vlan fcoevlan1
VLAN : fcoevlan1     FC-MAP : 0e:fc:00
  FCF : 00:10:94:00:00:01   Session Count : 2
    Enode-MAC : 00:10:94:00:00:02
```

```
      VN-Port-MAC : 0E:FC:00:00:00:05
      VN-Port-MAC : 0E:FC:00:00:00:01
```

show fip snooping vlan (QFX Series, VN2VF_Port FIP Snooping)

```
user@switch> show fip snooping vlan fcoevlan1
VLAN : fcoevlan1     Mode: VN2VF Snooping
  FC-MAP : 0e:fc:00
  FCF : 00:10:94:00:00:01   Session Count : 2
    Enode-MAC : 00:10:94:00:00:02
      VN-Port-MAC : 0E:FC:00:00:00:05
      VN-Port-MAC : 0E:FC:00:00:00:01
```

show fip snooping vlan (QFX Series, VN2VN_Port FIP Snooping)

```
user@switch> show fip snooping vlan vlan101
VLAN: vlan101,    Mode: VN2VN Snooping
  FC-MAP: 0e:fd:00
    Enode-MAC: 10:10:94:01:00:02 VN_Port count: 1
      VN-Port-MAC: 0e:fd:00:00:0a:01  Session Count: 2
    Enode-MAC: 10:10:94:01:00:03 VN_Port count: 0
```

show fip snooping vlan detail (QFX Series, VN2VN_Port FIP Snooping)

```
user@switch> show fip snooping vlan vlan101 detail
VLAN: vlan101,  Mode: VN2VN Snooping
  FC-MAP: 0e:fd:00
  Beacon_Period:  90000
  VN2VN Mode: Multi-Point
    Enode Information
    Enode-MAC: 10:10:94:01:00:02,       Interface: xe-0/0/10
      Active VN_Ports : 1
      VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0a:01
         Active Sessions    : 2
         Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0c:01
```

```
    Enode-MAC: 10:10:94:01:00:02,       Interface: xe-0/0/11
      Active VN_Ports : 0
```

show fip snooping vlan detail

```
user@switch> show fip snooping vlan fcoevlan1 detail
VLAN : fcoevlan1    FC-MAP : 0e:fc:00
  FCF Information
  FCF-MAC            : 00:10:94:00:00:01
  Active Sessions    : 2
  Configured FKA-ADV : 258
  Running FKA-ADV    : 235
    Enode Information
    Enode-MAC : 00:10:94:00:00:02       Interface : xe-0/0/1
    Configured FKA-ADV : 258
    Running FKA-ADV    : 239
      Session Information
      VN-Port MAC : 0E:FC:00:00:00:05   FKA-ADV : 255
      VN-Port MAC : 0E:FC:00:00:00:01   FKA-ADV : 251
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

Example: Configuring an FCoE Transit Switch

show fip snooping | **20008**

show fip snooping enode | **20015**

show fip snooping fcf | **20021**

show fip snooping statistics | **20029**

show fip snooping interface | **20024**

show fip vlan-discovery

IN THIS SECTION

- Syntax | 20040
- Description | 20040
- Options | 20040
- Required Privilege Level | 20040
- Output Fields | 20040
- Sample Output | 20041
- Release Information | 20042

Syntax

```
show fip vlan-discovery (enodes | statistics)
```

Description

Display FCoE VLAN information from the Fibre Channel switch or FCoE forwarder (FCF).

Options

enodes	Display VLAN discovery information for each ENode.
statistics	Display VLAN discovery information statistics.

Required Privilege Level

view

Output Fields

Table 573 on page 20041 lists the output fields for the `show fip vlan-discovery` command. Output fields are listed in the approximate order in which they appear.

Table 573: show fip vlan-discovery Output Fields

Field Name	Field Description	Level of Output
Enode-MAC	Media access control (MAC) address of the ENode.	enodes
Interface	Name of the interface.	enodes
Unsolicited notification count	Number of unsolicited VLAN discovery notifications.	All
Solicited notification count	Number of solicited VLAN discovery notifications.	statistics
Node Group Name	Displays the name of the Node group on QFabric systems.	statistics
Request count	Number of VLAN discovery requests sent by the ENode. This number should match the Solicited notification count number.	statistics
VLAN tags	Tags of the FIP-enabled VLANs.	enodes

Sample Output

show fip vlan-discovery enodes

```
user@switch> show fip vlan-discovery enodes

Enode-MAC            Interface      Unsolicited     Vlan Tags
                                    Notification
                                    Count
00:10:94:00:00:02    xe-0/0/9.0        0            400
```

show fip vlan-discovery statistics (QFX3500)

```
user@switch> show fip vlan-discovery statistics
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

show fip vlan-discovery statistics (QFabric Systems)

```
user@switch> show fip vlan-discovery statistics
NW-NG-0:
--------------------------------------------
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1

BBAK0399:
--------------------------------------------
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1

FCG001:
--------------------------------------------
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

clear fip vlan-discovery statistics | 3610

show firewall (QFX)

IN THIS SECTION

- Syntax | 20043
- Description | 20043
- Options | 20043
- Required Privilege Level | 20044
- Output Fields | 20044
- Sample Output | 20045
- Release Information | 20048

Syntax

```
show firewall
<application (CFM | eswd | RMPS)>>
<counter counter-name>
<filter filter-name>
<log <detail | interface interface-name>>
<terse>
```

Description

Display statistics about configured firewall filters.

Options

application (CFM | eswd | RMPS)

(Optional) Show firewall elements owned by the selected software component:

- Connectivity Fault Management (CFM)
- Ethernet switching daemon (eswd)—Shows only on devices that support it.
- Resource Management and Packet Steering (RMPS)

counter ***counter-name***	(Optional) Display statistics about a particular firewall filter counter.
filter ***filter-name***	(Optional) Display statistics about a particular firewall filter.
log	(Optional) Display log entries for all firewall filter activity.
terse	(Optional) Display firewall filter names only.

Required Privilege Level

view

Output Fields

Table 574 on page 20044 lists the output fields for the `show firewall` command. Output fields are listed in the approximate order in which they appear.

Table 574: show firewall Output Fields

Field Name	Field Description	Level of Output
Filter	Name of the filter that is configured at the `[edit firewall family` *`family-name`* `filter]` hierarchy level.	All levels
Counters	Display filter counter information: • Name—Name of a filter counter that has been configured with the **count** firewall filter action modifier. • Bytes—Number of bytes that match the filter term where the **count** action modifier was specified. • Packets—Number of packets that matched the filter term where the **count** action modifier was specified.	All levels

Table 574: show firewall Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Policers	Display policer information: • Name—Name of the policer that is configured at the `[edit firewall policer]` hierarchy level. • Packets—Number of packets that matched the filter term where the **policer** action modifier was specified. This is the number of packets that exceeded the rate limits that the policer specifies.	All levels
Action	Filter action: • **A**—Accept • **D**—Discard	All levels
Interface	Interface on which the firewall filter is applied.	All levels
Protocol	Name of the packet protocol.	All levels
Packet Length	Length of the packet.	All levels
Src Addr	Source address of the packet.	All levels
Dest Addr	Destination address of the packet.	All levels

Sample Output

show firewall

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                                                Bytes              Packets
counter-employee-web                                    0                    0
Filter: ingress-port-limit-tcp-icmp
```

```
Counters:
Name                                                Bytes              Packets
icmp-counter                                          560                 10
Policers:
Name                                               Packets
icmp-connection-policer                                10
tcp-connection-policer                                 0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

show firewall filter filter-name

```
user@switch> show firewall filter ingress-port-limit-tcp-icmp
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                                                Bytes              Packets
icmp-counter                                          560                 10
Policers:
Name                                               Packets
icmp-connection-policer                                10
tcp-connection-policer                                  0
```

show firewall counter counter-name

```
user@switch> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                                                Bytes              Packets
icmp-counter                                          560                 10
```

show firewall log

```
user@switch> show firewall log
Log :

Time      Filter    Action Interface     Protocol        Src Addr                          Dest
Addr
08:00:53  pfe       R      ge-1/0/6.0    ICMP            192.168.3.5
```

```
192.168.3.4
08:00:52  pfe       R      ge-1/0/6.0    ICMP            192.168.3.5
192.168.3.4
08:00:51  pfe       R      ge-1/0/6.0    ICMP            192.168.3.5
192.168.3.4
08:00:50  pfe       R      ge-1/0/6.0    ICMP            192.168.3.5
192.168.3.4
08:00:49  pfe       R      ge-1/0/6.0    ICMP            192.168.3.5
192.168.3.4
08:00:48  pfe       R      ge-1/0/6.0    ICMP            192.168.3.5
192.168.3.4
08:00:47  pfe       R      ge-1/0/6.0    ICMP            192.168.3.5
192.168.3.4
```

show firewall log detail

```
user@switch> show firewall log detail
Log :

Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Verifying That Two-Color Policers Are Operational

Overview of Firewall Filters (QFX Series)

Overview of Policers

show firewall

IN THIS SECTION

- Syntax | 20048
- Description | 20049
- Options | 20049
- Required Privilege Level | 20050
- Output Fields | 20050
- Sample Output | 20054
- Release Information | 20059

Syntax

```
show firewall
<application (cfm | eswd | rmps)>>
<counter counter-name>
<filter (filter-name| regex regular-expression |  version filter (filter-name)>>
|<log <(detail | interface regex  | version interface-namefilter-name)>>
<log <(detail | interface interface-name )>>
<prefix-action-stats filter filter-name prefix-action (prefix-action-name | prefix-action-name-
term-name) <from number to number filter >>
<logicalfilter-name prefix-action (prefix-action-nameprefix-action-name-term-name | all) <from
```

```
number to number>>
<logical-system (logical-system-name | all)>
< detail | terse>
```

Description

Display statistics about configured firewall filters.

If you query for options on the `show firewall filter` command, on Junos OS systems, you will see this output, which includes the configured Flowspec filters:

```
show firewall filter ?
Possible completions:
  <filtername>         Filter name
  __flowspec_default_inet__                                # Flowspec filter name
  application          Owner application
  counter              Counter name
  logical-system       Name of logical system, or 'all'
  regex                Show filter using regular expression
  version              Show filter version installed
```

However, on Junos OS Evolved systems, the Flowspec filters names are not shown here. To view Flowspec filters, use the `show firewall application routing` command.

Options

none

(Optional) Display statistics and counters for all configured firewall filters and counters. For EX Series switches, this command also displays statistics about all configured policers.

application (cfm | eswd | rmps

(Optional) Show firewall elements owned by the selected software component:

- Connectivity Fault Management (cfm)
- Ethernet switching daemon (eswd)—Shows only on devices that support it.
- Resource Management and Packet Steering (rmps)

counter *counter-name*	(Optional) Name of a filter counter.
detail	(EX Series switches and MX Series routers only) (Optional) Display firewall filter statistics and enhanced policer statistics and counters.
filter *filter-name*	(Optional) Name of a configured filter.
filter regex *regular-expression*	(Optional) Regular expression that matches the names of a subset of filters.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system.
log	(Optional) Display log entries for firewall filters.
log <(detail | interface *interface-name***)>**	(EX Series switches only) (Optional) Display detailed log entries of firewall activity or log information about a specific interface.
policer counters <(detail | counter-id *counter-index* **<detail>)>**	(EX8200 switches only) (Optional) Display enhanced policer counter statistics in brief or in detail.
prefix-action-stats	(Optional) Display prefix action statistics for firewall filters.
terse	(Optional) Display firewall filter names only.

Required Privilege Level

view

Output Fields

Table 575 on page 20051 lists the output fields for the `show firewall` command. Output fields are listed in the approximate order in which they appear.

Table 575: show firewall Output Fields

Field Name	Field Description
`Filter`	Name of a filter that has been configured with the `filter` statement at the `[edit firewall]` hierarchy level. Except on EX Series switches: • When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either `-i` for an input filter or `-o` for an output filter. • When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either `-in` for an input filter or `-out` for an output filter. When a logical system–specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, `__ls1/filter1`). • When a service filter is displayed that uses a service set, the separator between the service-set name and the service-filter name is a semicolon (:). **NOTE**: For `bridge family filter`, the `ip-protocol` match criteria is supported only for IPv4 and not for IPv6. This is applicable for line cards that support the Junos Trio chipset, such as the MX 3D MPC line cards.
`Counters`	Display filter counter information: • `Name`—Name of a filter counter that has been configured with the `counter` firewall filter action. • `Bytes`—Number of bytes that match the filter term under which the `counter` action is specified. • `Packets`—Number of packets that matched the filter term under which the `counter` action is specified. **NOTE**: On M and T Series routers, firewall filters cannot count `ip-options` packets on a per option type and per interface basis. A limited work around is to use the `show pfe statistics ip options` command to see `ip-options` statistics on a per Packet Forwarding Engine (PFE) basis. See show pfe statistics ip for sample output.

Table 575: show firewall Output Fields *(Continued)*

Field Name	Field Description
`Policers`	Display policer information: • `Name`—Name of policer. • `Bytes`—(For two-color policers on MX Series routers and EX Series switches, and for hierarchical policers on interfaces hosted on MICs and MPCs in MX Series routers) Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer. For other combinations of policer type, device, and line card type, this field is blank. • `Packets`—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.
`Policer Counter Index`	(EX8200 switch only) Global management counter ID. The counter ID value (*counter-index*) can be 0, 1, or 2.
`Green`	(EX8200 switch only) Number of packets within the limits. The number of packets is smaller than the committed information rate (CIR).
`Yellow`	(EX8200 switch only) Number of packets partially within the limits. The number of packets is greater than the CIR, but the burst size is within the excess burst size (EBS) limit.
`Discard`	(EX8200 switch only) Number of discarded packets.
`Bytes`	(EX8200 switch only) Number of green, yellow, red, or discarded packets in bytes.
`Packets`	(EX8200 switch only) Number of green, yellow, red, or discarded packets.
`Filter name`	(EX8200 switch only) Name of the filter with a term associated to a policer.
`Term name`	(EX8200 switch only) Name of the term associated with a policer.

Table 575: show firewall Output Fields *(Continued)*

Field Name	Field Description
`Policer name`	(EX8200 switch only) Name of the policer that is associated with a global management counter.
P1-t1	• OOS packet statistics for packets that are marked out-of-specification (out-of-spec) by the policer. Changes to all packets that have out-of-spec actions, such as discard, color marking, or forwarding-class, are included in this counter. • Offered packet statistics for traffic subjected to policing. • Transmitted packet statistics for traffic that is not discarded by the policer. When the policer action is discard, the statistics are the same as the in-spec statistics; when the policer action is non-discard (loss-priority or forwarding-class), the statistics are included in this counter.
`Action`	Filter action: • **A**—Accept • **D**—Discard
`Interface`	Interface on which the firewall filter is applied.
`Protocol`	Name of the packet protocol.
`Packet Length`	Length of the packet.
`Src Addr`	Source address of the packet.
`Dest Addr`	Destination address of the packet.

Sample Output

show firewall

```
user@host> show firewall
Filter: ef_path
Counters:
Name                                        Bytes              Packets
def-count                                   0                    0
video-count                                 0                    0
voice-count                                 0                    0

Filter: __default_bpdu_filter__

Filter: deep
Counters:
Name                                        Bytes              Packets
deep2                                       302076               5031

Filter: deep-flood
Counters:
Name                                        Bytes              Packets
deep_flood_def                              302136               5032
deep1                                       0                    0
Policers:
Name                                                                        Packets
deep-pol-op-first                                                  0
```

show firewall filter (MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name                              Bytes             Packets
Counter-1                             0                   0
Counter-2                             0                   0
Policers:
```

```
Name                            Bytes             Packets
Policer-1                        2770                  70
```

show firewall filter (non MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name                            Bytes             Packets
Counter-1                           0                   0
Counter-2                           0                   0
Policers:
Name                            Bytes             Packets
Policer-1                                              70
```

command-name

show firewall filter (Dynamic Input Filter)

```
user@host> show firewall filter dfwd-ge-5/0/0.1-in
Filter: dfwd-ge-5/0/0.1-in
Counters:
Name                                                Bytes              Packets
c1-ge-5/0/0.1-in                                        0                    0
```

show firewall (counter counter-name)

```
user@host> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                                                Bytes              Packets
icmp-counter                                            0                    0
```

show firewall log

```
user@host> show firewall log
Log :

Time      Filter    Action Interface     Protocol        Src Addr                        Dest
Addr
08:00:53  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
192.168.3.4
08:00:52  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
192.168.3.4
08:00:51  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
192.168.3.4
08:00:50  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
192.168.3.4
08:00:49  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
192.168.3.4
08:00:48  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
192.168.3.4
08:00:47  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
192.168.3.4
```

show firewall policer counters (EX8200 Switch)

```
user@switch> show firewall policer counters
Policer Counter Index 0:
                          Bytes            Packets
Green:                       73              15914
Yellow:                       9               1962
Discard:                    119              25942

Policer Counter Index 1:
                           Bytes            Packets
Green:                         0                  0
Yellow:                        0                  0
Discard:                       0                  0

Policer Counter Index 2:
                           Bytes            Packets
```

```
Green:                          0                  0
Yellow:                         0                  0
Discard:                        0                  0
```

show firewall policer counters (detail) (EX8200 Switch)

```
user@switch> show firewall policer counters detail
Policer Counter Index 0:
                          Bytes            Packets
Green:                       73              15914
Yellow:                       9               1962
Discard:                    119              25942

Filter name                Term name           Policer name
myfilter                   polcr-term-1        myfilter-polcr-1
inet-filter-ae             ae-snmp             policer-1
inet-filter-ae             ae-ssh              policer-2

Policer Counter Index 1:
                          Bytes            Packets
Green:                        0                  0
Yellow:                       0                  0
Discard:                      0                  0

Filter name                Term name           Policer name

Policer Counter Index 2:
                          Bytes            Packets
Green:                        0                  0
Yellow:                       0                  0
Discard:                      0                  0

Filter name                Term name           Policer name
```

show firewall policer counters (counter-id counter-index) (EX8200 Switch)

```
user@switch> show firewall policer counters counter-id 0
Policer Counter Index 0:
                          Bytes            Packets
Green:                       73              15914
```

```
Yellow:                       9               1962
Discard:                    119              25942
```

show firewall policer counters (counter-id counter-index detail) (EX8200 Switch)

```
user@switch> show firewall policer counters counter-id 0 detail
Policer Counter Index 0:
                          Bytes            Packets
Green:                       73              15914
Yellow:                       9               1962
Discard:                    119              25942

Filter name                Term name           Policer name
myfilter                   polcr-term-1        myfilter-polcr-1
inet-filter-ae             ae-snmp             policer-1
inet-filter-ae             ae-ssh              policer-2
```

show firewall detail

```
user@host> show firewall detail
Filter: __default_bpdu_filter__

Filter: foo
Counters:
Name                                                Bytes              Packets
c1                                               17652140               160474
Policers:
Name                                                Bytes              Packets
P1-t1
    OOS                                                 0                18286
    Offered                                             0 18446744073709376546
    Transmitted                                         0 18446744073709358260
```

show firewall application cfm (Junos OS Evolved)

```
user@host> show firewall application cfm
Filter: __cfm_filter_et-0/0/0__
```

```
Counters:
Name                                                Bytes              Packets
__cfm_cc_term_lvl_0__                                   0                    0
__cfm_cc_term_lvl_1__                                   0                    0
__cfm_cc_term_lvl_2__                                   0                    0
__cfm_cc_term_lvl_3__                                   0                    0
__cfm_cc_term_lvl_4__                                   0                    0
__cfm_cc_term_lvl_5__                                   0                    0
__cfm_cc_term_lvl_6__                                   0                    0
__cfm_cc_term_lvl_7__                                   0                    0
__cfm_ethtype_term__                                    0                    0
__cfm_lt_term_lvl_0__                                   0                    0
__cfm_lt_term_lvl_1__                                   0                    0
__cfm_lt_term_lvl_2__                                   0                    0
__cfm_lt_term_lvl_3__                                   0                    0
__cfm_lt_term_lvl_4__                                   0                    0
__cfm_lt_term_lvl_5__                                   0                    0
__cfm_lt_term_lvl_6__                                   0                    0
__cfm_lt_term_lvl_7__                                   0                    0
__cfm_ucast_term_536__                                  0                    0
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `logical-system` introduced in Junos OS Release 9.3.

Option `terse` introduced in Junos OS Release 9.4.

Option `policer counters` introduced in Junos OS Release 12.2 for EX Series switches.

Option `detail` introduced in Junos OS Release 12.3 for EX Series switches.

Option `detail` introduced in Junos OS Release 14.1 for MX Series routers.

Option `regex` *`regular-expression`* introduced in Junos OS Release 14.2.

Option `lsp` introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- Verifying That Firewall Filters Are Operational
- Verifying That Policers Are Operational

Overview of Policers

Enhanced Policer Statistics Overview

show firewall filter version | **20060**

show firewall prefix-action-stats | **20068**

show firewall filter version

IN THIS SECTION

- Syntax | **20060**
- Description | **20060**
- Options | **20060**
- Additional Information | **20061**
- Required Privilege Level | **20061**
- Output Fields | **20061**
- Sample Output | **20061**
- Release Information | **20061**

Syntax

```
show firewall filter version <filter-name>
```

Description

Display the version number of the installed firewall filter in the Routing Engine.

Options

none—(Optional) Display the version number of all installed firewall filters.

filter-name—(Optional) Name of a configured filter. If you specify the name of a filter, only the version number of that filter is displayed.

Additional Information

The initial version number is 1. This number increments by one when you modify the firewall filter settings or an associated prefix action. The maximum version number is 4,294,967,295. When the version number reaches 4,294,967,295, this number is reset to 1.

Required Privilege Level

view

Output Fields

Table 576 on page 20061 lists the output fields for the `show firewall filter version` command. Output fields are listed in the approximate order in which they appear.

Table 576: show firewall filter version Output Fields

Field Name	Field Description
`Filter`	Name of a filter that has been configured with the `filter` statement at the `[edit firewall]` hierarchy level.
`Version`	Display the version number of the firewall filter.

Sample Output

show firewall filter version

```
user@host> show firewall filter version
Filter version information :
Filter                                                Version
test                                                       10
```

Release Information

Command introduced in Junos OS Release 10.2R2.

show firewall log

IN THIS SECTION

- Syntax | 20062
- Syntax (EX Series Switches) | 20062
- Description | 20062
- Options | 20063
- Required Privilege Level | 20063
- Output Fields | 20063
- Sample Output | 20064
- Release Information | 20066

Syntax

```
show firewall log
<detail>
<extensive>
<interface interface-name>
<logical-system (logical-system-name | all)>
```

Syntax (EX Series Switches)

```
show firewall log
<detail>
<interface interface-name>
```

Description

Display log information about firewall filters.

Options

none	Display log information about firewall filters.
detail	(Optional) Display detailed information.
extensive	(Optional) Display hex dump of packet captured by log action.
interface *interface-name*	(Optional) Display log information about a specific interface.
logical-system (*logical-system-name* | all)	(Optional) Perform this operation on all logical systems or on a particular system.

Required Privilege Level

view

Output Fields

Table 577 on page 20063 lists the output fields for the `show firewall log` command. Output fields are listed in the approximate order in which they appear.

Table 577: show firewall log Output Fields

Field Name	Field Description
`Time of Log`	Time that the event occurred.
`Filter`	• Displays the name of a configured firewall filter or service filter only if the packet hit the filter's `log` action in a kernel filter (in the control plane). For any traffic that reaches the Routing Engine, the packets hit the `log` action in the kernel. • For all other logged packets (packet hit the filter's `log` action in the Packet Forwarding Engine), this field displays **pfe** instead of a configured filter name.

Table 577: show firewall log Output Fields *(Continued)*

Field Name	Field Description
`Filter Action`	Filter action: • `A`—Accept • `D`—Discard • `R`—Reject
`Name of Interface`	• Displays a physical interface name if the packet arrived at a port on a line card. • Displays **local** if the packet was generated by the device's internal Ethernet interface, `em1` or **fxp1**, which connects the Routing Engine with the router's packet-forwarding components.
`Name of protocol`	Packet's protocol name: `egp`, `gre`, `icmp`, `ipip`, `ospf`, `pim`, `rsvp`, `tcp`, or `udp`.
`Packet length`	Length of the packet.
`Source address`	Packet's source address.
`Destination address`	Packet's destination address and port.

Sample Output

show firewall log

```
user@host>show firewall log
Time      Filter    Action Interface     Protocol  Src Addr      Dest Addr
13:10:12  pfe       D      rlsq0.902     ICMP      192.0.2.2   192.0.2.1
13:10:11  pfe       D      rlsq0.902     ICMP      192.0.2.2   192.0.2.1
```

show firewall log detail

```
user@host> show firewall log detail
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
....
```

show firewall log extensive

```
user@host> show firewall log extensive
Time of Log: 2016-01-17 22:16:21 PST, Filter: pfe, Filter action: accept, Name of interface:
xe-0/0/1.0
Name of protocol: UDP, Packet Length: 98, Source address: 203.0.113.1, Destination address:
203.0.113.1
 :  00-0F: 00 01 03 ee ee ff 00 01 - 09 22 55 ee 81 00 02 58
 :  10-1F: 08 00 45 00 00 62 00 00 - 00 00 40 11 77 8a 01 00
 :  20-2F: 00 01 02 00 00 01 1c 00 - 1c 00 00 4e 19 83 00 01
 :  30-3F: 02 03 04 05 06 07 08 09 - 0a 0b 0c 0d 0e 0f 10 11
 :  40-4F: 12 13 14 15 16 17 18 19 - 1a 1b 1c 1d 1e 1f 20 21
 :  50-5F: 22 23 24 25 26 27 28 29 - 2a 2b 00 00 00 00 00 00
 :  60-6F: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00
 :  70-7F: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00
```

Release Information

Command introduced before Junos OS Release 7.4.

`extensive` option introduced in Junos OS Release 16.1.

`logical-system` option introduced in Junos OS Release 9.3.

show firewall policer

IN THIS SECTION

- Syntax | 20066
- Description | 20066
- Options | 20066
- Required Privilege Level | 20067
- Output Fields | 20067
- Sample Output | 20067
- Release Information | 20068

Syntax

```
show firewall policer
<policer-name>
```

Description

Display statistics about configured policers.

Options

none	Display the count of policed packets for all configured policers.
policer-name	(Optional) Display the count of policed packets for the specified policer.

Required Privilege Level

view

Output Fields

Table 578 on page 20067 lists the output fields for the `show firewall policer` command. Output fields are listed in the approximate order in which they appear.

Table 578: show firewall policer Output Fields

Field Name	Field Description	Level of Output
Filter	Name of the filter that is configured at the `[edit firewall family` *`family-name`* `filter]` hierarchy level.	All levels
Policers	Display policer information: • Filter—Name of filter that specifies the **policer** action modifier. • Name—Name of policer. • Packets—Number of packets that matched the filter term in which the **policer** action modifier is specified. This is the number of packets that exceed the rate limits that the policer specifies.	All levels

Sample Output

show firewall policer

```
user@switch> show firewall policer
Filter: egress-vlan-filter
Filter: ingress-port-filter
Policers:
Name                                                Packets
icmp-connection-policer                                   0
tcp-connection-policer                                    0
Filter: ingress-vlan-rogue-block
```

show firewall policer policer-name

```
user@switch> show firewall policer tcp-connection-policer
Filter: ingress-port-filter
Policers:
Name                                                Packets
tcp-connection-policer                                    0
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Verifying That Two-Color Policers Are Operational

Overview of Firewall Filters (QFX Series)

Overview of Policers

show firewall prefix-action-stats

IN THIS SECTION

- Syntax (filter-specific mode) | 20069
- Syntax (term-specific mode) | 20069
- Description | 20069
- Options | 20069
- Required Privilege Level | 20069
- Output Fields | 20070
- Sample Output | 20070
- Release Information | 20071

Syntax (filter-specific mode)

```
show firewall prefix-action-stats filter filter-name prefix-action prefix-action-name
<from number to number>
<logical-system (logical-system-name | all)>
```

Syntax (term-specific mode)

```
show firewall prefix-action-stats filter filter-name prefix-action prefix-action-name-term-name
<from number to number>
<logical-system (logical-system-name | all)>
```

Description

Display prefix action statistics about configured firewall filters.

If you clear statistics for firewall filters that are applied to MPCs and that also use the `prefix-action` action on matched packets, wait at least 5 seconds before you enter the `show firewall prefix-action-stats` command. A 5-second pause between issuing the `clear firewall` and `show firewall prefix-action-stats` commands avoids a possible timeout of the `show firewall prefix-action-stats` command.

By default, policers operate in *term-specific* mode.

See *Filter-Specific Policer Overview* for information about how to configure policers in *filter-specific* mode.

Options

filter *filter-name*	Name of a filter.
prefix-action *prefix-action-name*	Name of a prefix action.
from *number* to *number*	(Optional) Starting and ending counter or policer.
logical-system (*logical-system-name* | all)	(Optional) Perform this operation on all logical systems or on a particular system.

Required Privilege Level

view

Output Fields

Table 579 on page 20070 lists the output fields for the `show firewall prefix-action-stats` command. Output fields are listed in the approximate order in which they appear.

Table 579: show firewall prefix-action-stats Output Fields

Field Name	Field Description
`Filter`	Filter name. Filters configured for logical systems include the name of the filter prefixed with the two underscore characters (__) and the name of the logical system (for example, `__ls1/filter1`).

Sample Output

The following sample output assumes that the policer *act1* is in term mode and that there is a term named *term1* configured in the firewall filter *test.*

show firewall prefix-action-stats

```
user@host> show firewall prefix-action-stats filter test prefix-action act1-term1 from 0 to 9
Filter: test
Counters:
Name                        Bytes              Packets
act1-0                         0                    0
act1-1                         0                    0
act1-2                         0                    0
act1-3                         0                    0
act1-4                         0                    0
act1-5                         0                    0
act1-6                         0                    0
act1-7                         0                    0
act1-8                         0                    0
act1-9                         0                    0
Policers:
Name                        Bytes              Packets
act1-0                         0                    0
act1-1                         0                    0
act1-2                         0                    0
```

```
act1-3                      0                   0
act1-4                      0                   0
act1-5                      0                   0
act1-6                      0                   0
act1-7                      0                   0
act1-8                      0                   0
act1-9                      0                   0
```

Release Information

Command introduced before Junos OS Release 7.4.

`logical-system` option introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

clear firewall | 3613

show firewall templates-in-use

IN THIS SECTION

- Syntax | 20071
- Description | 20072
- Required Privilege Level | 20072
- Output Fields | 20072
- Sample Output | 20072
- Release Information | 20073

Syntax

```
show firewall templates-in-use
```

Description

Display the names of configured filter templates that are currently in use by dynamic subscribers and the number of times each template is referenced.

Required Privilege Level

view

Output Fields

Table 580 on page 20072 lists the output fields for the `show firewall templates-in-use` command. Output fields are listed in the approximate order in which they appear.

Table 580: show firewall templates-in-use Output Fields

Field Name	Field Description
`Filter Template`	Name of a filter that has been configured using the `filter` statement at either the `[edit firewall]` or `[edit dynamic-profiles` *`profile-name`* `firewall]` hierarchy and is being used as a template for dynamic subscriber filtering.
`Reference Count`	Number of times the filter has been referenced by subscribers accessing the network.

Sample Output

show firewall templates-in-use

```
user@host> show firewall templates-in-use
                                Dynamic Subscribers Reference Counts
Filter
Template
Reference Count
---------------
  ---------------
egressFilter
        10
ingressFilter
        10
```

```
dfilter
         5
dfilter-
pol
 5
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

clear firewall | 3613

show firewall log | 20062

show firewall (View)

IN THIS SECTION

- Syntax | 20073
- Description | 20074
- Options | 20074
- Required Privilege Level | 20074
- Output Fields | 20074
- Sample Output | 20076
- Release Information | 20076

Syntax

```
show firewall
<filter filter-name>
<counter counter-name>
```

```
<log>
<prefix-action-stats>
<terse>
```

Description

Display statistics about configured firewall filters.

Options

none	Display statistics about configured firewall filters.
filter ***filter-name***	Name of a configured filter.
counter ***counter-name***	Name of a filter counter.
log	Display log entries for firewall filters.
prefix-action-stats	Display prefix action statistics for firewall filters.
terse	Display firewall filter names only.

Required Privilege Level

view

Output Fields

Table 581 on page 20075 lists the output fields for the `show firewall` command. Output fields are listed in the approximate order in which they appear.

Table 581: show firewall Output Fields

Field Name	Field Description
`Filter`	Name of a filter that has been configured with the `filter` at the `[edit firewall]` hierarchy level. When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either `-i` for an input filter or `-o` for an output filter. When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either `-in` for an input filter or `-out` for an output filter. When a logical system–specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, `__ls1/filter1`).
`Counters`	Display filter counter information: • `Name`—Name of a filter counter that has been configured with the `counter` firewall filter action. • `Bytes`—Number of bytes that match the filter term under which the `counter` action is specified. • `Packets`—Number of packets that matched the filter term under which the `counter` action is specified.
`Policers`	Display policer information: • `Name`—Name of policer. • `Bytes`—Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer. • `Packets`—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.

Sample Output

show firewall

```
user@host> show firewall
Filter: ef_path
Counters:
Name                                        Bytes              Packets
def-count                                   0                    0
video-count                                 0                    0
voice-count                                 0                    0

Filter: __default_bpdu_filter__

Filter: deep
Counters:
Name                                        Bytes              Packets
deep2                                       302076               5031

Filter: deep-flood
Counters:
Name                                        Bytes              Packets
deep_flood_def                              302136               5032
deep1                                       0                    0
Policers:
Name                                                                        Packets
deep-pol-op-first                                        0
```

Release Information

Command introduced before Junos OS Release 10.0 .

RELATED DOCUMENTATION

firewall

show flexible-tunnels profiles

IN THIS SECTION

- Syntax | 20077
- Description | 20077
- Options | 20077
- Required Privilege Level | 20078
- Output Fields | 20078
- Sample Output | 20079
- Release Information | 20082

Syntax

```
show flexible-tunnels profiles
<brief | detail>
<tunnel-profile-name>
<route ip-address>
```

Description

Display information about flexible Virtual Extensible LAN (VXLAN) tunnel de-encapsulation profiles, which are programmed by controllers through the Juniper Extension Toolkit (JET) APIs.

Options

none	(Same as `brief`) Display information about all tunnel de-encapsulation profiles.
brief | detail	(Optional) Display the specified level of output.
tunnel-profile-name	Display information about the specified tunnel de-encapsulation profile.
route *ip-address*	Display information about the specified internal route in the **__flexible_tunnel_profiles__.inet.0** routing table.

Required Privilege Level

admin

Output Fields

Table 582 on page 20078 lists the output fields for the `show flexible-tunnels profiles` command. Output fields are listed in the approximate order in which they appear.

Table 582: show flexible-tunnels profiles Output Fields

Field Name	Field Description
tunnel-profile-name	Tunnel de-encapsulation profile name, if one was specified.
Client ID	An identifier of the client that injected the route specified in the tunnel profile.
Route prefix	An internal route that is associated with a particular tunnel de-encapsulation profile.
Table	Name of routing table, which is **__flexible_tunnel_profiles__.inet.0**.
Action	Action is `decapsulate`. If the contents of an incoming encapsulated packet matches the attributes specified in a particular tunnel de-capsulation profile, the packet is de-encapsulated as specified in the profile.
Tunnel type	Tunnel type is either `VXLAN-IPv4` or `VXLAN-IPv6`.
Interface	Name of the flexible tunnel interface for use with statistics, policing, filtering, and default encapsulation parameters.
VNI	VXLAN network identifier (VNI) specified in the tunnel de-encapsulation profile
Source prefix	IPv4 or IPv6 network address and length of the source that originated the packet.

Table 582: show flexible-tunnels profiles Output Fields *(Continued)*

Field Name	Field Description
Source UDP port range	A range of UDP port numbers, one of which is configured on the port that originated the packet.
Destination address	IPv4 or IPv6 network address of the host to which the packet is destined.
Destination UDP port	UDP port number that is configured on the host port to which the packet is destined.
VXLAN flags	An 8-bit VXLAN flag.

Sample Output

show flexible-tunnels profiles (none and brief)

```
user@host> show flexible-tunnels profiles
user@host> show flexible-tunnels profiles brief
CUSTOMER_0001-decap-mvnet1_1
  Client ID: ABCD, Route Prefix: 1.0.0.0/32, Table: __flexible_tunnel_profiles__.inet.0
  Action: Decapsulate, Tunnel type: VXLAN-IPv6
CUSTOMER_0001-decap-mvnet1_10
  Client ID: ABCD, Route Prefix: 1.0.0.9/32, Table: __flexible_tunnel_profiles__.inet.0
  Action: Decapsulate, Tunnel type: VXLAN-IPv6
CUSTOMER_0001-decap-mvnet1_11
  Client ID: ABCD, Route Prefix: 1.0.0.10/32, Table: __flexible_tunnel_profiles__.inet.0
  Action: Decapsulate, Tunnel type: VXLAN-IPv6
...
CUSTOMER_0001-decap-mvnet1_50001
  Client ID: ABCD, Route Prefix: 1.0.195.80/32, Table: __flexible_tunnel_profiles__.inet.0
  Action: Decapsulate, Tunnel type: VXLAN-IPv6
...
```

show flexible-tunnels profiles detail

```
user@host> show flexible-tunnels profiles detail
CUSTOMER_0001-decap-mvnet1_1
  Client ID: ABCD, Route Prefix: 1.0.0.0/32, Table: __flexible_tunnel_profiles__.inet.0
  Flexible IPv6 VXLAN tunnel profile
      Action: Decapsulate
      Interface: fti0.6 (Index: 10921)
      VNI: 16777213
      Source Prefix: 2a01:13:80:1:1:1::/96
      Source UDP Port Range: 49152 - 65535
      Destination Address: 2a01:10:255::2
      Destination UDP Port: 4789
      VXLAN Flags: 0x08
CUSTOMER_0001-decap-mvnet1_10
  Client ID: ABCD, Route Prefix: 1.0.0.9/32, Table: __flexible_tunnel_profiles__.inet.0
  Flexible IPv6 VXLAN tunnel profile
      Action: Decapsulate
      Interface: fti0.9 (Index: 10918)
      VNI: 16777213
      Source Prefix: 2a01:13:80:1:1:10::/96
      Source UDP Port Range: 49152 - 65535
      Destination Address: 2a01:10:255::2
      Destination UDP Port: 4789
      VXLAN Flags: 0x08
CUSTOMER_0001-decap-mvnet1_11
  Client ID: ABCD, Route Prefix: 1.0.0.10/32, Table: __flexible_tunnel_profiles__.inet.0
  Flexible IPv6 VXLAN tunnel profile
      Action: Decapsulate
      Interface: fti0.10 (Index: 10917)
      VNI: 16777213
      Source Prefix: 2a01:13:80:1:1:11::/96
      Source UDP Port Range: 49152 - 65535
      Destination Address: 2a01:10:255::2
      Destination UDP Port: 4789
      VXLAN Flags: 0x08
...
CUSTOMER_0001-decap-mvnet1_50001
  Client ID: ABCD, Route Prefix: 1.0.195.80/32, Table: __flexible_tunnel_profiles__.inet.0
  Flexible IPv6 VXLAN tunnel profile
      Action: Decapsulate
      Interface: fti0.6201 (Index: 10915)
```

```
      VNI: 16777214
      Source Prefix: 2a01:13:80:1:1:1::/96
      Source UDP Port Range: 49152 - 65535
      Destination Address: 2a01:10:255:1::2
      Destination UDP Port: 4789
      VXLAN Flags: 0x08
...
```

show flexible-tunnels profiles detail (to display GRE tunnel information)

```
user@host> show flexible-tunnels profiles detail
user@host> show flexible-tunnels profiles detail
decap-001
  Client ID: app1, Route Prefix: 10.0.0.1/32, Table: __flexible_tunnel_profiles__.inet.0
  Flexible IPv4 GRE tunnel profile
      Action: Decapsulate
      Source Prefix: 10.1.1.1
      Destination Address: 20.2.2.2
      GRE Key : 1
```

show flexible-tunnels profiles (Specific Profile, none and brief)

```
user@host> show flexible-tunnels profiles CUSTOMER_0001-decap-mvnet1_50001
user@host> show flexible-tunnels profiles CUSTOMER_0001-decap-mvnet1_50001
brief
CUSTOMER_0001-decap-mvnet1_50001
  Client ID: ABCD, Route Prefix: 1.0.195.80/32, Table: __flexible_tunnel_profiles__.inet.0
  Action: Decapsulate, Tunnel type: VXLAN-IPv6
```

show flexible-tunnels profiles detail (Specific Profile)

```
user@host> show flexible-tunnels profiles CUSTOMER_0001-decap-mvnet1_50001
detail
CUSTOMER_0001-decap-mvnet1_50001
  Client ID: ABCD, Route Prefix: 1.0.195.80/32, Table: __flexible_tunnel_profiles__.inet.0
  Flexible IPv6 VXLAN tunnel profile
      Action: Decapsulate
      Interface: fti0.6201 (Index: 10915)
```

```
      VNI: 16777214
      Source Prefix: 2a01:13:80:1:1:1::/96
      Source UDP Port Range: 49152 - 65535
      Destination Address: 2a01:10:255:1::2
      Destination UDP Port: 4789
      VXLAN Flags: 0x08
```

show flexible-tunnels profiles (Specific Route, none and brief)

```
user@host> show flexible-tunnels profiles route
1.0.195.80
user@host> show flexible-tunnels profiles route
1.0.195.80 brief
CUSTOMER_0001-decap-mvnet1_50001
  Client ID: ABCD, Route Prefix: 1.0.195.80/32, Table: __flexible_tunnel_profiles__.inet.0
  Action: Decapsulate, Tunnel type: VXLAN-IPv6
```

show flexible-tunnels profiles detail (Specific Route)

```
user@host> show flexible-tunnels profiles route
10.0.195.80 detail
CUSTOMER_0001-decap-mvnet1_50001
  Client ID: ABCD, Route Prefix: 1.0.195.80/32, Table: __flexible_tunnel_profiles__.inet.0
  Flexible IPv6 VXLAN tunnel profile
      Action: Decapsulate
      Interface: fti0.6201 (Index: 10915)
      VNI: 16777214
      Source Prefix: 2a01:13:80:1:1:1::/96
      Source UDP Port Range: 49152 - 65535
      Destination Address: 2a01:10:255:1::2
      Destination UDP Port: 4789
      VXLAN Flags: 0x08
```

Release Information

Command introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

Understanding Programmable Flexible VXLAN Tunnels

show flight-recorder status

IN THIS SECTION

- Syntax | 20083
- Description | 20083
- Options | 20083
- Required Privilege Level | 20083
- Output Fields | 20084
- Sample Output | 20085
- Release Information | 20085

Syntax

```
show flight-recorder status
```

Description

Display the current status of the flight recorder tool and associated parameters, such as the running status of the tool, and the current data snapshot list.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 583 on page 20084 lists the output fields for the `show flight-recorder status` command. Output fields are listed in the approximate order in which they appear.

Table 583: show flight-recorder status Output Fields

Field Name	Field Description
`Flight-recorder status`	State of the flight recorder tool: • `Running`—The flight recorder tool is enabled using the `request flight-recorder set high-cpu` command. • `Not Running`—The flight recorder tool is not enabled. By default, the flight recorder tool is disabled.
`Recent Parameter Data`	Information about configured parameters for the flight recorder tool: • `Cpu-threshold`—Specify the maximum value of CPU utilization in percentage, beyond which the collection of data is triggered. • `Polling-frequency`—Specify the time in seconds for polling for high CPU utilization. • `Backoff-duration`—Specify the time interval in seconds between two snapshots of data. • `Num-snapshots`—Specify the number of snapshots of data to be collected before quitting the collection process.
Flags set	Information about additional flags configured for the flight recorder tool: • `Collect-core`—Perform snapshot collection of the running core with every snapshot of data taken. • `Logical System`—Enable data collection on logical systems.
Snapshot Directory	Log file that is recorded and saved in the flight recorder directory. The recorded snapshots and core log files are saved in a folder under the **/var/log/flight_recorder/** directory. The folder format is `Flr_MONTH_DD_YYYY_HH:MM:SS`; for example, `Flr_May_09_2018_02:20:50`.

Table 583: show flight-recorder status Output Fields *(Continued)*

Field Name	Field Description
List of snapshots	List of log files recorded and saved under the flight recorder directory.

Sample Output

show flight-recorder status

```
user@host> show flight-recorder status
Flight-recorder status: Not Running!

Recent Parameter Data:
Cpu-threshold           10
Polling-frequency       5
Backoff-duration        10
Num-snapshots           3

Flags set:
Collect-core flag is set
Logical System flag is Not set (default)

Snapshot Directory : Flr_Feb_22_2018_13:26:41

List of snapshots:
 flr_2018-02-22_13:26:41.txt
 flr_2018-02-22_13:27:04.txt
 flr_2018-02-22_13:27:28.txt
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

request flight-recorder set high-cpu | **14633**

show forwarding-options enhanced-hash-key

IN THIS SECTION

- Syntax | 20086
- Description | 20086
- Required Privilege Level | 20086
- Output Fields | 20086
- Sample Output | 20089
- Release Information | 20094

Syntax

```
show forwarding-options enhanced-hash-key
```

Description

Display information about which packet fields are used by the hashing algorithm to make hashing decisions.

You can configure the fields that are inspected by the hashing algorithm to make hashing decisions for traffic entering a LAG bundle using the `forwarding-options enhanced-hash-key` statement.

Required Privilege Level

view

Output Fields

Table 584 on page 20087 lists the output fields for the `show forwarding-options enhanced-hash-key` command. Output fields are listed in the approximate order in which they first appear. Output fields vary by platform.

Table 584: show forwarding-options enhanced-hash-key Output Fields

Field Name	Field Description
Hash-Mode	Current hash mode: Layer 2 header or Layer 2 payload.
Protocol	Indicates whether the Protocol field is or is not used by the hashing algorithm: Yes or No.
Destination L4 Port	Indicates whether the Destination L4 Port field is or is not used by the hashing algorithm: Yes or No.
Source L4 Port	Indicates whether the Source L4 Port field is or is not used by the hashing algorithm: Yes or No.
Destination IPv4 Addr	Indicates whether the Destination IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.
Source IPv4 Addr	Indicates whether the Source IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.
Incoming port	Indicates whether the incoming port number (interface) is or is not used by the hashing algorithm. Yes or No **NOTE**: When passive monitoring is enabled on a QFX10000 Series switch interface, the `inet`, `inet6` and `L2` fields are all set to No.
Vlan id	Indicates whether the Vlan ID field is or is not used by the hashing algorithm: Yes or No.
Inner-Vlan ID	Indicates whether the inner Vlan field is or is not used by the hashing algorithm: Yes or No.
Next Hdr	Indicates whether the Next Hdr field is or is not used by the hashing algorithm: Yes or No.
Destination IPv6 Addr	Indicates whether the Destination IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.

Table 584: show forwarding-options enhanced-hash-key Output Fields *(Continued)*

Field Name	Field Description
Source IPv6 Addr	Indicates whether the Source IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.
Ether Type	Indicates whether the Ether Type field is or is not used by the hashing algorithm: Yes or No.
Destination MAC Address	Indicates whether the Destination MAC Address field is or is not used by the hashing algorithm: Yes or No.
Source MAC Address	Indicates whether the Source MAC Address field is or is not used by the hashing algorithm: Yes or No.
Symmetric-hashing	Indicates whether symmetric hashing is set for `inet` and `inet6`: Yes or No.
Load Balancing Method for ALB	Indicates the load balancing method for adaptive load balancing (ALB): flowlet or per-packet. The load balancing method is flowlet by default, and can be configured using the **fabric-load-balance** statement.
Load Balancing Method for DLB (QFX5120-32C, QFX5120-48Y, and QFX5220 switches)	Indicates various Dynamic Load Balancing (DLB) modes: • Flowlet • Assigned flow • Per-packet Refer Dynamic Load Balancing for more details.
Fabric Link Scale	Indicates the fabric link scale, in mbps.
Inactivity Interval	Indicates the fabric load balance inactivity interval, in microseconds (us). The inactivity interval is 16 microseconds by default, and can be configured using the **inactivity-interval** statement.

Table 584: show forwarding-options enhanced-hash-key Output Fields *(Continued)*

Field Name	Field Description
Hash Region Size/ Trunk	Indicates the hash region size, in buckets per fabric trunk.
Seed	A hash seed value, between 0 and 4294967295. If a hash-seed value is not configured it is automatically assigned on the QFX10000 Series switches. A hash-seed prevents traffic polarization to same links on the next hop QFX switch when two are connected with LAG/ECMP.
Key	Indicates whether the GRE key field is or is not used by the hashing algorithm: Yes or No.
Protocol	Indicates if a Generic Router Encapsulation (GRE) endpoint over routes was dynamically learned by a routing protocol such as RIP or OSPF.
MPLS Enabled	Indicates if MPLS is enabled under L2 switching.
VXLAN VNID	A 24-bit virtual network identifier (VNID) that uniquely identifies the Virtual Extensible Local Area Networks (VXLAN) segment.

Sample Output

show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0


Current Hash Settings
----------------------
Hash-Mode                              :layer2-payload

  inet Hash settings-
  --------------------
    inet packet fields
       Protocol                        : Yes
```

```
        Destination L4 Port             : Yes
        Source L4 Port                  : Yes
        Destination IPv4 Addr           : Yes
        Source IPv4 Addr                : Yes
        Vlan id                         : No


  inet6 Hash settings-
  ---------------------
    inet6 packet fields
        Next Hdr                        : Yes
        Destination L4 Port             : Yes
        Source L4 Port                  : Yes
        Destination IPv6 Addr           : Yes
        Source IPv6 Addr                : Yes
        Vlan id                         : No
```

show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0


Current Hash Settings
----------------------

Hash-Mode                               : layer2-header

  layer2 Hash settings-
  ----------------------
     layer2 packet fields
        Ether Type                      : Yes
        Destination MAC Address         : Yes
        Source MAC Address              : Yes
        VLAN ID                         : No
```

show forwarding-options enhanced-hash-key (Symmetric Hashing)

```
user@switch> show forwarding-options enhanced-hash-key
<some output removed for brevity>
```

```
  Symmetric-hashing SETTINGS
  ------------------------------------------------------------
     inet  : Yes
     inet6 : Yes
```

show forwarding-options enhanced-hash-key (Fabric Load Balancing Options)

```
user@switch> show forwarding-options enhanced-hash-key
<some output removed for brevity>

  Fabric Load Balancing Options
  ------------------------------------------------------------
      Load Balancing Method  : Flowlet
      Fabric Link Scale      : 40960 (mbps)
      Inactivity Interval    : 16 (us)
      Hash Region Size/Trunk : 1024 (buckets)
```

show forwarding-options enhanced-hash-key (Dynamic Load Balancing Options)

```
user@switch> show forwarding-options enhanced-hash-key ecmp-dlb
```

```
Slot 0
Current RTAG7 Settings
---------------------
Hash-Mode                              : layer2-payload
  inet RTAG7 settings-
  --------------------
    inet packet fields
       Protocol                        : Yes
       Destination L4 Port             : Yes
       Source L4 Port                  : Yes
       Destination IPv4 Addr           : Yes
       Source IPv4 Addr                : Yes
       Vlan id                         : No

  inet6 RTAG7 settings-
  ---------------------
```

```
    inet6 packet fields
       Next Hdr                        : Yes
       Destination L4 Port             : Yes
       Source L4 Port                  : Yes
       Destination IPv6 Addr           : Yes
       Source IPv6 Addr                : Yes
       Vlan id                         : No

  ECMP Load Balancing Options
   ------------------------------------------------------------
      Load Balancing Method   : Flowlet
      Inactivity Interval    : 64 (us)
```

show forwarding-options enhanced-hash-key (QFX10000 Series Switches)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0


Seed value for Hash function           0: 2301323130
Seed value for Hash function           1: 2301323130
Seed value for Hash function           2: 2301323130
Seed value for Hash function           3: 2301323130

    Inet settings:
    --------------
        IPV4 dest address:    Yes
        IPV4 source address:  Yes
        L4 Dest Port:         Yes
        L4 Source Port:       Yes
        Incoming port:        No
    Inet6 settings:
    --------------
        IPV6 dest address:    Yes
        IPV6 source address:  Yes
        L4 Dest Port:         Yes
        L4 Source Port:       Yes
        Incoming port:        No
    L2 settings:
    ------------
        Dest Mac address:    No
```

```
        Source Mac address:  No
        Vlan Id:             Yes
        Inner-vlan Id:       No
        Incoming port:       No
    GRE settings:
    -------------
        Key:                 No
        Protocol:            No
    MPLS settings:
    --------------
    MPLS Enabled:              Yes

    VXLAN settings:
    ---------------
        VXLAN VNID:          No
```

show forwarding-options enhanced-hash-key (Junos OS Evolved)

```
user@host>show forwarding-options enhanced-hash-key

    Hash Seed
    -------------------------
        Hash Seed                       : 2227128662
        Resilient Hash Seed             : 2227128662

    Inet RTAG7 Settings
    --------------------
    Inet packet fields

        Protocol                        : enable
        Source L4 Port                  : tcp, udp, sctp
        Destination L4 Port             : tcp, udp, sctp
        IPV4 Source Address             : enable
        IPV4 Destination Address        : enable

    GTP HASHING INFORMATION:
    ---------------------------
        Inet Tunnel Endpoint ID         : disable
        Inet6 Tunnel Endpoint ID        : disable

    Inet6 RTAG7 settings-
```

```
--------------------
Inet6 packet fields

    Next Header                     : enable
    Source L4 Port                  : tcp, udp, sctp
    Destination L4 Port             : tcp, udp
    IPV6 Source Address             : enable
    IPV6 Destination Address        : enable
    Flow label                      : enable

MPLS RTAG7 settings-
--------------------
    label                           : enable

VXLAN Settings
-------------------------
    Vxlan Id                        : disable

Layer 2 Settings
-------------------------
    Source Mac Address              : disable
    Destination Mac Address         : disable
    Vlan Id                         : enable
    Inner Vlan Id                   : enable
    Incoming Interface              : disable
```

Release Information

Command introduced in Junos OS Release 13.2X51-D15.

`Fabric Load Balancing Options` output fields introduced in Junos OS Release 14.1X53-D10.

`Incoming port` output field introduced in Junos OS Release 18.4R1 for QFX10000 Series switches.

Command introduced in Junos OS Evolved Release 21,3R1.

The `ecmp-dlb` statement introduced in Junos OS Evolved Release 19.4R2 for QFX5220 switches.

RELATED DOCUMENTATION

- Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)
- Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic

show forwarding-options hash-key

IN THIS SECTION

- Syntax (PTX10003-80C, PTX10003-160C (PTX10003), and ACX7000 Series) | 20095
- Description | 20095
- Required Privilege Level | 20095
- Output Fields | 20095
- Sample Output | 20098
- Release Information | 20100

Syntax (PTX10003-80C, PTX10003-160C (PTX10003), and ACX7000 Series)

```
show forwarding-options hash-key
```

Description

Display information about which packet fields are used by the hashing algorithm to make hashing decisions.

Required Privilege Level

View

Output Fields

Table 585: show forwarding-options hash-key Output Fields

Field Name	Field Description
Hash-Mode	Current hash mode: Layer 4 header or Layer 4 payload.
Protocol	Indicates whether the Protocol field is or is not used by the hashing algorithm: Yes or No.

Table 585: show forwarding-options hash-key Output Fields *(Continued)*

Field Name	Field Description
`Destination L4 Port`	Indicates whether the Destination L4 Port field is or is not used by the hashing algorithm: Enable or Disable.
`Source L4 Port`	Indicates whether the Source L4 Port field is or is not used by the hashing algorithm: Yes or No.
`Destination IPv4 Addr`	Indicates whether the Destination IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.
`Source IPv4 Addr`	Indicates whether the Source IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.
`Vlan id`	Indicates whether the Vlan ID field is or is not used by the hashing algorithm: Yes or No.
`Inner-Vlan ID`	Indicates whether the inner Vlan field is or is not used by the hashing algorithm: Yes or No.
`Inet Tunnel Endpoint ID`	Indicates whether the GTP hashing, Inet Tunnel Endpoint ID is or is not used by the hashing algorithm: Yes or No.
`Inet6 Tunnel Endpoint ID`	Indicates whether the GTP hashing, Inet6 Tunnel Endpoint ID is or is not used by the hashing algorithm: Yes or No.
`Incoming Interface`	Indicates whether the interfaceis or is not used by the hashing algorithm: Yes or No.
`Next Hdr`	Indicates whether the Next Hdr field is or is not used by the hashing algorithm: Yes or No.
`Destination IPv6 Addr`	Indicates whether the Destination IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.
`Source IPv6 Addr`	Indicates whether the Source IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.

Table 585: show forwarding-options hash-key Output Fields *(Continued)*

Field Name	Field Description
`Destination MAC Address`	Indicates whether the Destination MAC Address field is or is not used by the hashing algorithm: Yes or No.
`Source MAC Address`	Indicates whether the Source MAC Address field is or is not used by the hashing algorithm: Yes or No.
`Hash Seed`	A hash seed value, between 0 and 2882365321. A hash-seed prevents traffic polarization to same links on the next hop QFX switch when two are connected with LAG/ECMP.
Resilient Hash Seed	A resilient hash seed value, between 0 and 2557914061.
`MPLS Enabled`	Indicates if MPLS is enabled under L4 routing.
`VXLAN ID`	A identifier (ID) that uniquely identifies the Virtual Extensible Local Area Networks (VXLAN) segment.

Table 586: show forwarding-options hash-key Output Fields on ACX7000 Series

Field Name	Field Description
`Multiservice Hash Settings`	Enables source and destination mac fields in the hashkey.
`Inet Hash Settings`	• Layer3 Proto-Incorporate IPv4 Layer 3 data in the hashkey • Layer4 Proto-Incorporate Layer 4 data in the hashkey
`Inet6 Hash Settings`	• Layer3 Proto-Incorporate IPv4 Layer 3 data in the hashkey • Layer4 Proto-Incorporate Layer 4 data in the hashkey

Table 586: show forwarding-options hash-key Output Fields on ACX7000 Series *(Continued)*

Field Name	Field Description
Mpls Hash Settings	• Label-Incorporate MPLS label in the hashkey • IpPayload-Incorporate layer3 Payload in the hashkey • EtherPayload-Incorporate ether payload in the hashkey • ZeroCW-Incorporate preceding ether packet payload in the hashkey

Sample Output

show forwarding-options hash-key

```
user@host> show forwarding-options hash-key
```

```
Hash Seed
    -------------------------
        Hash Seed                       : 2882365321
        Resilient Hash Seed             : 2557914061

    Inet RTAG7 Settings
    --------------------
    Inet packet fields

        Protocol                        : enable
        Source L4 Port                  : disabled
        Destination L4 Port             : disabled
        IPV4 Source Address             : enable
        IPV4 Destination Address        : enable

    GTP HASHING INFORMATION:
    ---------------------------
        Inet Tunnel Endpoint ID         : disable
        Inet6 Tunnel Endpoint ID        : disable

    Inet6 RTAG7 settings-
```

```
    --------------------
    Inet6 packet fields

        Next Header                     : enable
        Source L4 Port                  : tcp, udp
        Destination L4 Port             : tcp
        IPV6 Source Address             : enable
        IPV6 Destination Address        : enable
        Flow label                      : disable

    MPLS RTAG7 settings-
    --------------------
        label                           : enable

    VXLAN Settings
    -------------------------
        Vxlan Id                        : disable

    Layer 2 Settings
    -------------------------
        Source Mac Address              : enable
        Destination Mac Address         : enable
        Vlan Id                         : enable
        Inner Vlan Id                   : enable
        Incoming Interface              : enable
```

show forwarding-options hash-key (ACX7000 Series)

```
user@host> show forwarding-options hash-key
```

```
SENT: Ukern command: show evo-pfemand hashkey
Multiservice Hash Settings
----------------------------
    Source MAC: disabled
    Destination MAC: disabled

Inet Hash Settings
--------------------
    Layer3 Proto: disabled
    Layer4 Proto: disabled
```

```
Inet6 Hash Settings
----------------------
    Layer3 Proto: disabled
    Layer4 Proto: disabled

Mpls Hash Settings
--------------------
    Label: disabled
    IpPayload: disabled
    EtherPayload: disabled
    ZeroCW: disabled
```

Release Information

Command introduced in Junos OS Evolved Release 21.3R1.

RELATED DOCUMENTATION

- Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

show forwarding-options hyper-mode

IN THIS SECTION

- Syntax | 20101
- Description | 20101
- Required Privilege Level | 20101
- Output Fields | 20101
- Sample Output | 20102
- Release Information | 20102

Syntax

```
show forwarding-options hyper-mode
```

Description

Display information about the hyper mode feature. After you configure the hyper mode feature, you must reboot the system for the network device (a router or a switch) to reflect the change. For instance, if you have configured hyper mode but not rebooted the system, the `Current mode` field displays `normal` while the `Configured mode` field displays `hyper mode`.

> **NOTE**: Since hyper mode is a chassis-specific feature, the `show forwarding-options hyper-mode` command on a guest network function (GNF) displays the hyper-mode status of the chassis (BSYS), not of the GNF. To know more about GNFs, see Components of Junos Node Slicing.

Required Privilege Level

view

Output Fields

Table 587 on page 20101 lists the output fields for the `show forwarding-options hyper-mode` command. Output fields are listed in the order in which they appear.

Table 587: show forwarding-options hyper-mode Output Fields

Field Name	Field Description
`Current mode`	Displays the current mode, either `normal` or `hyper mode`.
`Configured mode`	Displays the configured mode, either `normal` or `hyper mode`.

Sample Output

show forwarding-options hyper-mode

```
user@host> show forwarding-options hyper-mode
Current mode: hyper mode
Configured mode: hyper mode
```

Release Information

Command introduced in Junos OS Release 13.3R4.

RELATED DOCUMENTATION

Configuring Hyper Mode on Enhanced MPCs to Speed Up Packet Processing

Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches

hyper-mode (forwarding-options) | 8221

show forwarding-options load-balance

IN THIS SECTION

- Syntax | 20103
- Description | 20103
- Options | 20104
- Required Privilege Level | 20104
- Output Fields | 20104
- Sample Output | 20105
- Sample Output | 20105
- Sample Output | 20106
- Sample Output | 20106
- Release Information | 20106

Syntax

```
show forwarding-options load-balance
```

Description

Displays the load-balancing hash result. You can view this information for one, two, or three levels of load balancing. This command can be used in two ways to get different load-balancing information based on the parameters:

- To get the load-balancing decision result for routed IPv4, IPv6, and other L3 traffic, use the following:

  ```
  show forwarding-options load-balance ingress-interface <interface-name> family <family-type> source-address <src-IP> destination-address <dest-IP> source-mac-address <source-mac-address> dest-mac-address <destination-mac-address> transport-protocol <transport-protocol> source-port <src-port> destination-port <dest-port> tos <TOS>
  ```

  > **NOTE**: source-mac-address and dest-mac-address are only supported on BT based line-cards on PTX10001-36MR, LC1201 line card, LC1202 line card, PTX10004, PTX10008 and PTX10016 ,PTX10003

  You can use this command when all the packet header details are available.

- To get the load-balancing decision result for raw packet dump files, use the following:

  ```
  show forwarding-options load-balance ingress-interface <interface-name> family <family-type> packet-dump <pkt-dump>
  ```

  You can use this command when the packets in the traffic are complex to be described by the 5-tuple.

> **NOTE**:
>
> - This feature is not supported for multicast flows and L2 packets.
> - The maximum size of the raw packets entered on the CLI is 256 KB.
> - When injecting raw packets, ensure 24 bytes of trailing space, which is accounted for in the packet header length. This space will be used for inserting metadata to the injected probe packet.

Options

interface-name	Ingress logical interface.
family-type	Layer 3 family “inet/inet6”.
src-IP	Source IP address.
dest-IP	Destination IP address.
source-mac-address	source MAC address.
dest-mac-address	Destination MAC address.
transport-protocol	Transport layer protocol “tcp/udp”.
src-port	Source port (0 – 65535).
dest-port	Destination port (0 – 65535).
TOS	Type of service field (0 – 255).
pkt-dump	Raw packet dump in hexadecimal without '0x'. The hexadecimal code of packet information must start from the L3 header (IPv4/IPv6) to the end of the packet. It should be less than 256,000 characters in length.

Required Privilege Level

view

Output Fields

Table 588 on page 20104 lists the output fields for the `show forwarding-options load-balance` command. Output fields are listed in the order in which they appear.

Table 588: show forwarding-options hyper-mode Output Fields

Field Name	Field Description
`Outgoing logical aggregate interface`	Egress aggregated Ethernet interface for current parameters

Table 588: show forwarding-options hyper-mode Output Fields *(Continued)*

Field Name	Field Description
`Outgoing member physical interface`	Egress physical interface for current parameters
`Outgoing next hop address`	Egress next hop chosen for current parameters
`Outgoing next hop id`	Egress next-hop ID chosen for current parameters

Sample Output

AE Egress-IPv6 Egress

```
user@host> show forwarding-options load-balance family inet6 ingress-interface xe-5/0/3
transport-protocol tcp source-address 2201::2 destination-address 2202::2 source-port 1617
destination-port 1640 tos 224
================ fpc5 ================

Outgoing logical aggregate interface: ae0.0
Outgoing member physical interface  : xe-4/2/1
Outgoing next hop address           : fe80::2e21:72ff:fe71:1d
Outgoing next hop id                : 700
```

Sample Output

MPLS Egress IPv4

```
user@host> show forwarding-options load-balance ingress-interface xe-5/0/3 family inet source-
address 102.1.1.2 destination-address 202.1.1.2 tos 4 transport-protocol udp source-port 1431
destination-port 6531
================ fpc5 ================
Outgoing logical interface          : xe-5/0/2.0
Outgoing next hop address           : 40.40.40.2
Outgoing next hop id                : 747
```

Sample Output

Packet hash:IPv4 (MPLS egress)

```
user@host> show forwarding-options load-balance ingress-interface xe-5/0/3 family inet packet-
dump
4580004e000000003d114b1966010102ca01010205971983003a31abb9b7c07aa5bafa4349786960000000001101112134
bdfa97e00161a1b1c1d1e1f202122232425262728292a2b2c2d2e2fa073
================ fpc5 ================
Outgoing logical interface          : xe-4/1/1.0
Outgoing next hop address           : -
Outgoing next hop id                : 708
```

Note that for the MPLS egress, the outgoing next-hop address will not be displayed.

Sample Output

Packet hash:IPv6 (MPLS egress)

```
user@host> show forwarding-options load-balance ingress-interface xe-5/0/3 family inet6 packet-
dump
6c0000000026113d2101000000000000000000000000000022201000000000000000000000000000206b6078000265a60b
9b7c07aa5bafa4349786960000000021011121336 1b3e6300021a1bd63a
================ fpc5 ================

Outgoing logical interface          : xe-4/1/2.0
Outgoing next hop address           : -
Outgoing next hop id                : 705
```

Release Information

To learn more about platform support for `load-balance`, refer Getting load-balancing hash result information.

RELATED DOCUMENTATION

- show forwarding-options analyzer
- show forwarding-options port-mirroring | **20116**

show forwarding-options hyper-mode | 20100

show forwarding-options load-balance ecmp|trunk

IN THIS SECTION

- Syntax | 20107
- Syntax | 20107
- Description | 20108
- Options | 20109
- Required Privilege Level | 20110
- Output Fields | 20110
- Sample Output | 20110
- Release Information | 20111

Syntax

```
show forwarding-options load-balance ecmp

     (ipv4 | ipv6 | l2)
     ◆ ipv4 source-ip source-ip  destination-ip destination-ip vlan-id vlan-id l4-source-port l4-source-port l4-destination-port l4-destination-port protocol protocol ingress-interface interface
    ◆ ipv6 source-ip source-ip  destination-ip destination-ip vlan-id vlan-id l4-source-port l4-source-port l4-destination-port l4-destination-port  next-header next-header ingress-interface interface
◆ l2 source-mac source-mac destination-mac destination-mac ether-type ethertype  vlan-id vlan-id ingress-interface interface (destination-ipv4-address destination-ipv4-address  | destination-ipv6-address destination-ipv6-address)
```

Syntax

```
show forwarding-options load-balance trunk
```

```
    (ipv4 | ipv6 | l2)
    ◆ ipv4 source-ip source-ip destination-ip destination-ip vlan-id vlan-id l4-source-port l4-source-port l4-destination-port l4-dest-port protocol protocol ingress-interface interface egress-ae-interface interface packet-type (broadcast | l2-multicast | l3-multicast|unicast | unknown-unicast)
    ◆ ipv6 source-ip source-ip destination-ip destination-ip vlan-id vlan-id l4-source-port l4-source-port l4-destination-port l4-destination-port next-header next-header ingress-interface interface egress-ae-interface interface packet-type (broadcast | l2-multicast | l3-multicast | unicast | unknown-unicast)
◆ l2 source-mac source-mac destination-mac destination-mac ether-type ethertype vlan-id vlan-id ingress-interface interface egress-ae-interface interface packet-type (broadcast | l2-multicast | l3-multicast| unicast | unknown-unicast)
```

Description

Display hash parameters used by the hashing algorithm and the final egress interface for the traffic you◆re interested in. Use the command for ECMP (Layer 3 load balancing) or for aggregated Ethernet interfaces (Layer 2 load balancing). Use these commands to troubleshoot issues for which you need to know the packet path.

NOTE: The `show forwarding-options load-balance` option is not available in Junos EVO.

NOTE: Some options are required. If you do not specify those required options when you enter the command, you are provided with an error message that tells you that you must provide a value for the field. The required options are shown in **boldface** text in the preceding Syntax sections and are marked as ◆(Required)◆ in the following Options section.

To use this command, specify *one* of the `ipv4`, `ipv6`, or `l2` required options, with the choice of that option based on the hash mode that is configured on the switch. Check the `hash-mode` value by using the `show forwarding-options enhanced-hash-key` command. By default, the hash mode is `layer2-payload`. When the hash mode is `layer2-payload`, specify the `l2` option for Layer 2 packets (ARP, STP, and so on), the `ipv4` option for IPv4 packets, and the `ipv6` option for IPv6 packets. When the hash mode is `layer2-header`, specify the `l2` option for all types of packets. Next, provide valid inputs for `Hash settings` that are marked `Yes` in the `show forwarding-options enhanced-hash-key` command output. A value of 0 (zero) is automatically supplied for input for `Hash settings` fields that are marked as `No` in the same command output.

> **NOTE**: You do not need to provide input for fields such as `ECMP Next-hop` and `Next-hop ID/Type` (see these in the sample later in this document). Values for those fields are provided automatically.

Options

destination-ip ***destination-ip***	(Required) Destination IP address.
destination-ipv4-address ***destination-ipv4-address***	(Required) Destination IPv4 address.
destination-ipv6-address ***destination-ipv6-address***	(Required) Destination IPv6 address.
destination-mac ***destination-mac***	(Required) Destination MAC address.
egress-ae-interface ***interface***	Egress aggregated Ethernet LAG interface﹪the LAG bundle to which the traffic is destined.
ether-type ***ethertype***	Ethertype value.
ingress-interface ***interface***	(Required) Ingress interface﹪the network-facing interface where packets enter the switch.
ipv4	(Required) IPv4 packets.
ipv6	(Required) IPv6 packets.
l2	(Required) Layer 2 packets (such as ARP or STP) when the hash mode is `layer2-payload`. All packet types when the hash mode is `layer2-header`.
l4-destination-port ***l4-destination-port***	Layer 4 destination port number.
l4-source-port ***l4-source-port***	Layer 4 source port number.
next-header ***next-header***	Next header, as an integer value. Specifies the type of the next header for an IPv6 packet. This field usually specifies the transport layer protocol used by a packet's payload. When extension headers are present in the packet, this field indicates which extension header follows.
packet-type (broadcast | l2-multicast | l3-multicast|unicast | unknown-unicast)	Packet type: *One* of `broadcast`, `l2-multicast`, `l3-multicast`, `unicast`, or `unknown-unicast`.

protocol *protocol*	Protocol, as an integer value. The Protocol field in the IPv4 header contains a number indicating the type of data found in the payload portion of the datagram.
source-ip *source-ip*	(Required) Source IP address.
source-mac *source-mac*	(Required) Source MAC address.
vlan-id *vlan-id*	VLAN ID.

Required Privilege Level

view

Output Fields

Table 589 on page 20110 lists the output fields for the `show forwarding-options load-balance ecmp` and `show forwarding-options load-balance trunk` commands. Output fields are listed in the order in which they appear.

Table 589: show forwarding-options load-balance ecmp or show forwarding-options load-balance trunk Output Fields

Field Name	Field Description
`Hash parameters`	Hash parameters for the selected hash mode `layer2-payload` or `layer2-header`.
`Hash result`	`Egress interface:` *`interface-name`*: ECMP member or trunk member for the packet.

Sample Output

show forwarding-options load-balance ecmp ipv4

```
user@switch> show forwarding-options load-balance ecmp ipv4 source-ip 192.168.8.193 destination-
ip 203.0.113.22 vlan-id 0 l4-source-port 0 l4-destination-port 0 protocol 6 ingress-interface
ge-0/0/12
Hash parameters
---------------
        Hash Mode            : layer2-payload
```

```
        Ingress interface    : 14
        VLAN ID              : 0
        Ether Type           : 0x800
        Source IP Address    : 192.168.8.193
        Dest IP Address      : 203.0.113.22
        Protocol             : 6
        Source L4 port       : 0
        Dest L4 port         : 0
        ECMP Next-hop        : 200001
        Next-hop ID/ Type    : 131071/Unilist

Hash result
-----------
  Egress interface     : ge-0/0/34
```

Release Information

Commands introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show forwarding-options enhanced-hash-key | 20086

show forwarding-options load-balance | 20102

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

show forwarding-options next-hop-group

IN THIS SECTION

- Syntax | 20112
- Description | 20112
- Options | 20112
- Required Privilege Level | 20112
- Output Fields | 20112
- Sample Output | 20113

- Release Information | **20115**

Syntax

```
show forwarding-options next-hop-group
<terse | brief | detail>
<group-name>
```

Description

Display current state of next-hop groups.

Options

terse | brief | detail	(Optional) Display the specified level of output.
group-name	(Optional) Display a single next-hop group.

Required Privilege Level

view

Output Fields

Table 590 on page 20112 lists the output fields for the `show forwarding-options next-hop-group` command. Output fields are listed in the approximate order in which they appear.

Table 590: show forwarding-options next-hop-group Output Fields

Field Name	Field Description	Level of Output
Next-hop-group	Name of next-hop group.	All levels
Type	Next-hop group type, such as **inet**, **inet6** or **layer-2**.	All levels

Table 590: show forwarding-options next-hop-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Next-hop group state, either **up** or **down**.	All levels
Members Interfaces	Names of interfaces to which next-hop group members belong.	**brief detail**
Member Subgroup	Names of subgroups to which next-hop group members belong.	**brief detail**
Number of members configured	Number of next-hop group members configured.	**detail**
Number of members that are up	Number of next-hop group members that are up.	**detail**
Number of subgroups configured	Number of subgroups configured.	**detail**
Number of subgroups that are up	Number of subgroups that are up.	**detail**

Sample Output

show forwarding-options next-hop-group terse

```
user@host> show forwarding-options next-hop-group terse
Next-hop-group                 Type             State
  nhg                            inet             up
  nhg6                           inet6            up
  vpls_nhg_2                           layer-2             down
```

show forwarding-options next-hop-group brief

```
user@host> show forwarding-options next-hop-group brief

Next-hop-group: nhg
      Type:  inet
      State: up
      Members Interfaces:
        ge-0/2/8.0             next-hop  192.0.2.10
        ge-5/1/8.0             next-hop  198.51.100.10
        ge-5/1/9.0             next-hop  203.0.113.10

Next-hop-group: nhg6
      Type:  inet6
      State: up
      Members Interfaces:
        ge-5/1/5.0             next-hop  2001:db8::1:10
        ge-5/1/6.0             next-hop  2001:db8::20:10        Member Subgroup: nhsg6
          Members Interfaces:
            ge-5/0/4.0             next-hop  2001:db8::3:1
            ge-5/1/4.0             next-hop  2001:db8::4:1

Next-hop-group: vpls_nhg_2
Type: layer-2        State: down
```

show forwarding-options next-hop-group detail

```
user@host> show forwarding-options next-hop-group detail

Next-hop-group: nhg
Type: inet
State: up
Number of members configured    : 3
Number of members that are up   : 3
Number of subgroups configured  : 0
Number of subgroups that are up : 0
Members Interfaces:                                 State
  ge-0/2/8.0             next-hop  192.0.2.10        up
  ge-5/1/8.0             next-hop  203.0.113.10        up
  ge-5/1/9.0             next-hop  198.51.100.10.10        up
```

```
Next-hop-group: nhg6
Type: inet6
State: up
Number of members configured    : 2
Number of members that are up   : 2
Number of subgroups configured  : 1
Number of subgroups that are up : 1
Members Interfaces:                                 State
ge-5/1/5.0             next-hop  2001:db8::1:10       up
ge-5/1/6.0             next-hop  2001:db8::20:10      up
Member Subgroup: nhsg6                            up
       Number of members configured    : 2
       Number of members that are up   : 2
       Members Interfaces:                                 State
           ge-5/0/4.0             next-hop  2001:db8::3:1   up
           ge-5/1/4.0             next-hop  2001:db8::4:1   up

Next-hop-group: vpls_nhg_2
Number of members configured    : 2
Number of members that are up   : 0
Number of subgroups configured  : 0
Number of subgroups that are up : 0
Type: layer-2        State: down
Members Interfaces:             State
   ge-2/2/1.100                 down
   ge-2/3/9.0                 down
```

Release Information

Command introduced in Junos OS Release 9.6.

Support for IPv6 introduced in Junos OS Release 14.2 for the MX Series routers.

RELATED DOCUMENTATION

show forwarding-options port-mirroring | **20116**

show forwarding-options port-mirroring

IN THIS SECTION

- Syntax | 20116
- Description | 20116
- Options | 20116
- Required Privilege Level | 20116
- Output Fields | 20117
- Sample Output | 20118
- Release Information | 20120

Syntax

```
show forwarding-options port-mirroring
<terse | detail>
<instance-name>
<self-mirror (none | start |statistics |stop)>
```

Description

Display current state of port-mirroring instances.

Options

terse | detail	(Optional) Display the specified level of output.
instance-name	(Optional) Display a single port-mirroring instance.

Required Privilege Level

view

Output Fields

Table 591 on page 20117 lists the output fields for the `show forwarding-options port-mirroring` command. Output fields are listed in the approximate order in which they appear.

Table 591: show forwarding-options port-mirroring Output Fields

Field Name	Field Description	Level of Output
Instance Name	Name of port-mirroring instance.	All levels
Instance Id	Instance identification number.	All levels
State	Instance state, either `up` or `down`.	All levels
Input parameters		
Rate	Rate (ratio of packets sampled).	`detail`
Run-length	Run length (number of consecutive packets sampled).	`detail`
Maximum-packet-length	Maximum packet length.	`detail`
Output parameters		
Family	Protocol family.	`detail`
State	Instance state, either `up` or `down`.	`detail`
Destination	Destination (next-hop group name).	`detail`
Next-hop	IP address of the next hop to the destination.	`detail`

Table 591: show forwarding-options port-mirroring Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Parameters (self-mirror)	Parameters for self-mirror : • filename of PCAP file • family : *family-name* • size : *file-size* • packet count : *number of mirrored packets saved in file* • total bytes written : *count* • start record time : *time* • stop record time : *time* • duration (seconds) : *seconds* • rate : *number* • run-length : *number* • mode : *write mode* (circular | linear) • self mirror state : *recording state* (inactive | active | recording) •	

Sample Output

show forwarding-options port-mirroring terse

```
user@host> show forwarding-options port-mirroring terse
Instance Name                 Instance Id    State
&global_instance             1            up
inst1                     2            up
```

show forwarding-options port-mirroring detail

```
user@host> show forwarding-options port-mirroring detail
Instance Name: pm1
  Instance Id: 2
  Input parameters:
    Rate                  : 2
    Run-length            : 0
    Maximum-packet-length : 0
  Output parameters:
    Family              State     Destination          Next-hop
    inet                up        ge-0/0/0.0           10.1.1.2
    inet6               up        ge-0/0/0.0           2001:db8::2
    any                 up        ge-0/0/1.0           NA
```

show forwarding-options port-mirroring self-mirror

```
user@host> show forwarding-options port-mirroring self-mirror
Instance Name: I1
 Instance Id: 1
 Input parameters:
   Rate                  : 1
   Run-length            : 0
   file name             :/var/tmp/bt_capture_inet_egress
   family                :inet
   size                  :5000000
   mode                  :circular
   self-mirror state     : inactive <recording state inactive/active/recording>
```

show forwarding-options port-mirroring self-mirror statistics

```
user@host> show forwarding-options port-mirroring self-mirror statistics
Instance Name    :I1
  Instance Id    :1
Parameters
    file name           :/var/tmp/mirror1.pcap
    family              :inet
    size                :10485760
```

```
    packet count        :3021
    total bytes written :362130
```

show forwarding-options port-mirroring self-mirror start

```
user@host> show forwarding-options port-mirroring self-mirror start
Instance Name    :I1
  Instance Id    :1
 Parameters
    family                :inet
    start record time     :2024-06-29 22:12:35
    duration(seconds)     :300
```

show forwarding-options port-mirroring self-mirror stop

```
user@host> show forwarding-options port-mirroring self-mirror stop
Instance Name    :I1
  Instance Id    :1
 Parameters
    family                 :inet
    stop record time       :2024-06-29 22:15:35
    duration(seconds)      :300
```

Release Information

Command introduced in Junos OS Release 9.6.

show groups junos-defaults

IN THIS SECTION

- Syntax | 20121
- Description | 20121
- Required Privilege Level | 20122

- Release Information | **20122**

Syntax

```
show groups junos-defaults
```

Description

Display the full set of available preset statements from the defaults group.

```
user@host# show groups junos-defaults
    groups {
            junos-defaults {
                    applications {
                            # File Transfer Protocol
                            application junos-ftp {
                                    application-protocol ftp;
                                    protocol tcp;
                                    destination-port 21;
                }
                            # Trivial File Transfer Protocol
                            application junos-tftp {
                                    application-protocol tftp;
                                    protocol udp;
                                    destination-port 69;
                }
                            # RPC port mapper on TCP
                            application junos-rpc-portmap-tcp {
                                    application-protocol rpc-portmap;
                                    protocol tcp;
                                    destination-port 111;
                }
                            # RPC port mapper on UDP
            }
        }
    }
```

Required Privilege Level

view

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Using Configuration Groups to Quickly Configure Devices

show gtp tunnels

IN THIS SECTION

- Syntax | 20122
- Description | 20122
- Options | 20123
- Required Privilege Level | 20123
- Output Fields | 20123
- Sample Output | 20123
- Sample Output | 20125
- Release Information | 20127

Syntax

```
show security gtp tunnels (brief | summary | detail)
```

Description

Use this command to display all existing GTP tunnels.

Options

- brief—Displays a short listing of all GTP tunnels.
- summary—Displays a summary of all GTP tunnels.
- detail–Displays detailed information about all the GTP tunnels.

Required Privilege Level

view

Output Fields

Sample Output

show security gtp tunnels

```
user@host> show security gtp tunnels
```

Refer the GTP tunnel output for IPv4.

```
FPC 7 PIC 0:

 Index: 72000002, EBI/LBI: 5/5(V2)to sgw, Timeout: 1440m
 User: 61.0.0.102, 12345678 --> 62.0.0.102, 00000021
 Control: 61.0.0.101, 00325ac1 --> 62.0.0.101, 00000001

FPC 8 PIC 0:

 Index: 0x02000040 Tunnel ID: 0x50502410121507f5(V0), Timeout: 59m
 User: 20.1.0.1, 00000001 -> 20.0.2.1, 00000001
 Ctrl: 20.1.0.1, 00000001 -> 20.0.2.1, 00000001

 3 tunnels active in total
```

show security gtp tunnels summary

```
user@host> show security gtp tunnels summary
```

Refer the GTP tunnel summary output for IPv4 address.

```
FPC 1 PIC 0:

FPC 1 PIC 1:

FPC 2 PIC 0:

FPC 2 PIC 1:

2 tunnels active in total
```

show security gtp tunnels detail

```
user@host> show security gtp tunnels detail
```

Refer the GTP tunnel detail output for IPv4 address.

```
node0:
--------------------------------------------------------------------------
 FPC 0 PIC 0:

 FPC 0 PIC 1:

 FPC 0 PIC 2:

 FPC 0 PIC 3:

Index: 0x02000040 Tunnel ID: 0x50502410121507f5(V0), Timeout: 59m
User: 20.1.0.1, 00000001 -> 20.0.2.1, 00000001
Ctrl: 20.1.0.1, 00000001 -> 20.0.2.1, 00000001

 1 tunnels active in total
```

command-name

Sample Output

Starting with Junos OS 18.4R1, the GTP tunnel is supported on both IPv4 and IPv6 address.

show security gtp tunnels

```
user@host> show security gtp tunnels
```

Refer the GTP tunnel output for IPv4 and IPv6 address.

```
node0:
--------------------------------------------------------------------------
FPC 0 PIC 0:

Index: 0x00000002, V2, Control tunnel, IMSI: 110469790910178, MSISDN: 123456780912345
   IPv6: 2001:10::108, 0000042a -> 2001:13::104, 0000042a
   IPv4: 201.10.0.101, 0000042a -> 201.13.0.103, 0000042a
    User tunnel, EBI/LBI:5/5, to pgw, Timeout:50
      IPv6: 2001:10::102, 00000822 -> 2001:13::109, 00000832
      IPv4: 201.10.0.104, 00000822 -> 201.13.0.107, 00000832
    User tunnel, EBI/LBI:6/5, to pgw, Timeout:51
      IPv6: ::, 00000823 -> ::, 00000833
      IPv4: 201.10.0.105, 00000823 -> 201.13.0.108, 00000833
2 tunnels active in total

node1:
--------------------------------------------------------------------------
FPC 0 PIC 0:

Index: 0x00000002, V2, Control tunnel, IMSI: 110469790910178, MSISDN: 123456780912345
   IPv6: 2001:10::108, 0000042a -> 2001:13::104, 0000042a
   IPv4: 201.10.0.101, 0000042a -> 201.13.0.103, 0000042a
    User tunnel, EBI/LBI:5/5, to pgw, Timeout:50
      IPv6: 2001:10::102, 00000822 -> 2001:13::109, 00000832
      IPv4: 201.10.0.104, 00000822 -> 201.13.0.107, 00000832
    User tunnel, EBI/LBI:6/5, to pgw, Timeout:50
      IPv6: ::, 00000823 -> ::, 00000833
```

```
      IPv4: 201.10.0.105, 00000823 -> 201.13.0.108, 00000833
2 tunnels active in total
```

show security gtp tunnels summary

```
user@host> show security gtp tunnels summary
```

Refer the GTP tunnel summary output for IPv4 and IPv6 address.

```
node0:
--------------------------------------------------------------------------
FPC 0 PIC 0: 2 tunnels active

2 tunnels active in total


node1:
--------------------------------------------------------------------------
FPC 0 PIC 0: 2 tunnels active

2 tunnels active in total
```

show security gtp tunnels detail

```
user@host> show security gtp tunnels detail
```

Refer the GTP tunnel detail output for IPv4 and IPv6 address.

```
node0:
--------------------------------------------------------------------------
FPC 0 PIC 0:

Index: 0x00000002, V2, Control tunnel, IMSI: 110469790910178, MSISDN: 123456780912345
   IPv6: 2001:10::108, 0000042a -> 2001:13::104, 0000042a
   IPv4: 201.10.0.101, 0000042a -> 201.13.0.103, 0000042a
    User tunnel, EBI/LBI:5/5, to pgw, Timeout:51, alive time: 8
    Uplink: Packets 0, Bytes 0, Downlink: Packets 0, Bytes 0
```

```
      IPv6: 2001:10::102, 00000822 -> 2001:13::109, 00000832
      IPv4: 201.10.0.104, 00000822 -> 201.13.0.107, 00000832
    User tunnel, EBI/LBI:6/5, to pgw, Timeout:51, alive time: 8
    Uplink: Packets 0, Bytes 0, Downlink: Packets 0, Bytes 0
      IPv6: ::, 00000823 -> ::, 00000833
      IPv4: 201.10.0.105, 00000823 -> 201.13.0.108, 00000833
2 tunnels active in total

node1:
--------------------------------------------------------------------------
FPC 0 PIC 0:

Index: 0x00000002, V2, Control tunnel, IMSI: 110469790910178, MSISDN: 123456780912345
   IPv6: 2001:10::108, 0000042a -> 2001:13::104, 0000042a
   IPv4: 201.10.0.101, 0000042a -> 201.13.0.103, 0000042a
    User tunnel, EBI/LBI:5/5, to pgw, Timeout:51, alive time: 8
    Uplink: Packets 0, Bytes 0, Downlink: Packets 0, Bytes 0
      IPv6: 2001:10::102, 00000822 -> 2001:13::109, 00000832
      IPv4: 201.10.0.104, 00000822 -> 201.13.0.107, 00000832
    User tunnel, EBI/LBI:6/5, to pgw, Timeout:51, alive time: 8
    Uplink: Packets 0, Bytes 0, Downlink: Packets 0, Bytes 0
      IPv6: ::, 00000823 -> ::, 00000833
      IPv4: 201.10.0.105, 00000823 -> 201.13.0.108, 00000833
2 tunnels active in total
```

Release Information

Command introduced in Junos OS Release 10.0.

Support for GPRS tunneling protocol version 2 (GTPv2) added in Junos OS Release 11.4.

Command output updated in Junos OS Release 15.1X49-D40.

Starting in Junos OS Release 18.4R1, GPRS tunneling protocol (GTP) traffic is supported on both IPv4 and IPv6 through tunnel-based session distribution for security inspection.

RELATED DOCUMENTATION

clear gtp tunnels | 3618

show hfrr profiles

IN THIS SECTION

- Syntax | 20128
- Description | 20128
- Options | 20128
- Required Privilege Level | 20128
- Output Fields | 20129
- Sample Output | 20130
- Release Information | 20130

Syntax

```
show hfrr profiles
<brief |extensive>
```

Description

Display host fast reroute (HFRR) profile information.

HFRR adds a precomputed protection path into the Packet Forwarding Engine, such that if a link between a provider edge device and a server farm becomes unusable for forwarding, the Packet Forwarding Engine can use another path without having to wait for the router or the protocols to provide updated forwarding information.

Options

none	Display information about HRFF profiles.
brief | extensive	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 592 on page 20129 describes the output fields for the `show hfrr profiles` command. Output fields are listed in the approximate order in which they appear.

Table 592: show hfrr profiles Output Fields

Field Name	Field Description
HFRR	
HFRR current state	Status of the HFRR profile: HFRR_ACTIVE, HFRR_INACTIVE, HFRR_IFLH-NOT-CONF, and so on.
HFRR Prefix limit blackout timer expiry (in secs)	Time interval between an HFRR profile becoming inactive on exceeding the ARP prefix limit, and the profile starting the SYNC process.
HFRR prefix limit hit count	Number of times that an HFRR profile becomes inactive on exceeding the ARP prefix limit.
HFRR protected IFL name	Interface configured for the HFRR feature.
HFRR protected IFL handle	
HFRR routing instance name	The routing instance in which the HFRR interface is configured.
HFRR routing instance handle	
HFRR sync BG scheduled	
HFRR RTS filter on	
HFRR delete BG scheduled	
HFRR ARP prefix limit	Configured ARP prefix limit.

Table 592: show hfrr profiles Output Fields *(Continued)*

Field Name	Field Description
HFRR ARP supplementary blackout timeout (in mins)	Supplementary time-out value configured for profile to be inactive when it hits ARP prefix limit.
HFRR number of ARP routes learned	Number of ARP routes learned on the configured interface.
HFRR number of FRR routes created	Number of ARP routes created on the configured interface.

Sample Output

show hfrr profiles

```
user@host> show hfrr profiles
HFRR pointer: 0x9254000
HFRR current state: HFRR_ACTIVE
HFRR Prefix limit blackout timer expiry (in secs): 0
HFRR prefix limit hit count: 0
HFRR protected IFL name: ge-4/1/0.0
HFRR protected IFL handle: 0x9248738
HFRR routing instance name: test
HFRR routing instance handle: 0x9145740
HFRR sync BG scheduled: NO
HFRR RTS filter on: YES
HFRR delete BG scheduled: NO
HFRR ARP prefix limit: 0
HFRR ARP supplementary blackout timeout (in mins): 1
HFRR number of ARP routes learned: 4
HFRR number of FRR routes created: 2
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring Link Protection with Host Fast Reroute

show host

IN THIS SECTION

- Syntax | 20131
- Syntax (Junos OS Evolved) | 20131
- Syntax | 20131
- Description | 20132
- Options | 20132
- Additional Information | 20132
- Required Privilege Level | 20132
- Sample Output | 20132
- Release Information | 20132

Syntax

```
show host hostname
<routing-instance mgmt_junos>
<server server-name>
```

Syntax (Junos OS Evolved)

```
show host hostname
<routing-instance mgmt_junos>
<server server-name>
```

Syntax

Description

Display Domain Name System (DNS) hostname information.

Options

hostname	Hostname or address.
routing-instance mgmt_junos	(Optional) Side host server that is running.
server *server-name*	(Optional) Name server to use.

Additional Information

The `show host` command displays the raw data received from the name server.

Required Privilege Level

view

Sample Output

show host

```
user@host> show host device
device.example.net has address 192.0.2.0

user@host> show host 192.0.2.0
Name: device.example.net
Address: 192.0.2.0
Aliases:
```

Release Information

Command introduced before Junos OS Release 7.4.

`routing-instance mgmt_junos` option introduced in Junos OS Evolved Release 18.3R1.

`routing-instance mgmt_junos` option introduced in Junos OS Release 19.2R1.

show iccp

IN THIS SECTION

- Syntax | 20133
- Description | 20133
- Options | 20133
- Required Privilege Level | 20134
- Output Fields | 20134
- Sample Output | 20134
- Release Information | 20135

Syntax

```
show iccp <brief | detail>
logical-system [system-name | all]
```

Description

Display Inter-Chassis Control Protocol (ICCP) information about the multichassis link aggregation group (MC-LAG) peers, including the state of the TCP connection, Bidirectional Forwarding Detection (BFD) protocol, backup liveness peer status, and MCSNOOPD, LACPD, and ESWD applications.

Options

none—Display ICCP information about the MC-LAG peers, including the state of the TCP connection and BFD protocol, and MCSNOOPD, LACP, and ESWD applications.

brief—Display brief ICCP information about the MC-LAG peers, including the state of the TCP connection and BFD protocol, and MCSNOOPD, LACP, and ESWD applications.

detail—Display detailed ICCP information about the MC-LAG peers, including the state of the TCP connection and BFD protocol, and MCSNOOPD, LACP, and ESWD applications.

logical-system [*system-name* | all]—(Optional) Display information for a specified logical system or all systems.

Required Privilege Level

view

Output Fields

Table 593 on page 20134 lists the output fields for the `show iccp` command. Output fields are listed in the approximate order in which they appear.

Table 593: show iccp Output Fields

Field Name	Field Description
Redundancy Group Information for peer	Aggregated Ethernet interface name.
TCP Connection	Specifies if the TCP connection between the peers hosting the MC-LAG is up or down.
Liveness Detection	Specifies if liveness detection, also known as Bidirectional Forwarding Detection (BFD), is up or down.
Client Application	Specifies information regarding the state of the MCSNOOPD and client applications.

Sample Output

show iccp (QFX Series)

```
user@switch> show iccp
Redundancy Group Information for peer 10.3.3.2
  TCP Connection       : Established
  Liveliness Detection : Up

Client Application: MCSNOOPD

Client Application: eswd
```

show iccp (MX Series)

```
user@host> show iccp
Logical system    :LS1
    Redundancy Group Information for peer 10.1.1.1
       TCP Connection       : Established
       Liveliness Detection : Up
       Redundancy Group ID          Status
         1                           Up
         2                           Up

Client Application: lacpd
Redundancy Group IDs Joined:  1
Redundancy Group IDs Joined:  2

Client Application: l2ald_iccpd_client
 Redundancy Group IDs Joined:  1
 Redundancy Group IDs Joined:  2
```

Release Information

Command introduced in Junos OS Release 10.0.

Support for logical systems added in Junos OS Release 14.1 for MX Series routers.

RELATED DOCUMENTATION

iccp | 8250

show igmp group

IN THIS SECTION

- Syntax | 20136
- Syntax (EX Series Switch and the QFX Series) | 20136

- Description | **20136**
- Options | **20136**
- Required Privilege Level | **20137**
- Output Fields | **20137**
- Sample Output | **20138**
- show igmp group (User Plane output for ACX7100) | **20141**
- Release Information | **20141**

Syntax

```
show igmp group
<brief | detail>
<group-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show igmp group
<brief | detail>
<group-name>
```

Description

Display Internet Group Management Protocol (IGMP) group membership information.

Options

none	Display standard information about membership for all IGMP groups.
brief | detail	(Optional) Display the specified level of output.
group-name	(Optional) Display group membership for the specified IP address only.

logical-system (all | *logical-system-name)*	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 594 on page 20137 describes the output fields for the `show igmp group` command. Output fields are listed in the approximate order in which they appear.

Table 594: show igmp group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the IGMP membership report. A name of **local** indicates that the local routing device joined the group itself.	All levels
Group	Group address.	All levels
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.	All levels
Source	Source address.	All levels
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	**detail**
Last reported by	Address of the host that last reported membership in this group.	All levels
Timeout	Time remaining until the group membership is removed.	**brief none**
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	**detail**

Table 594: show igmp group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Type	Type of group membership: • **Dynamic**—Host reported the membership. • **Static**—Membership is configured.	All levels

Sample Output

show igmp group (Include Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.2
        Last reported by: 203.0.113.52
        Timeout:      24 Type: Dynamic
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.3
        Last reported by: 203.0.113.52
        Timeout:      24 Type: Dynamic
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.4
        Last reported by: 203.0.113.52
        Timeout:      24 Type: Dynamic
    Group: 198.51.100.2
        Group mode: Include
        Source: 203.0.113.4
        Last reported by: 203.0.113.52
        Timeout:      24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
```

```
    Group: 198.51.100.12
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 198.51.100.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show igmp group (Exclude Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 198.51.100.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 198.51.100.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show igmp group brief

The output for the `show igmp group brief` command is identical to that for the `show igmp group` command.

show igmp group detail

```
user@host> show igmp group detail
Interface: t1-0/1/0.0
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.2
        Source timeout: 12
        Last reported by: 203.0.113.52
```

```
        Group timeout:       0 Type: Dynamic
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.3
        Source timeout: 12
        Last reported by: 203.0.113.52
        Group timeout:       0 Type: Dynamic
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.4
        Source timeout: 12
        Last reported by: 203.0.113.52
        Group timeout:       0 Type: Dynamic
    Group: 198.51.100.2
        Group mode: Include
        Source: 203.0.113.4
        Source timeout: 12
        Last reported by: 203.0.113.52
        Group timeout:       0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 198.51.100.12
        Group mode: Exclude
        Source: 0.0.0.0
        Source timeout: 0
        Last reported by: Local
        Group timeout:       0 Type: Dynamic
    Group: 198.51.100.22
        Group mode: Exclude
        Source: 0.0.0.0
        Source timeout: 0
        Last reported by: Local
        Group timeout:       0 Type: Dynamic
```

show igmp group (User Plane output for ACX7100)

Displays the following output in the user-plane for customer premise equipment (CPE).

```
user@up1> show igmp group
Interface: pp0.3221225481, Groups: 1
    Group: 225.0.0.1
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: 100.1.1.2
        Timeout:     232 Type: Dynamic
Interface: demux0.2147483652, Groups: 1
    Group: 225.0.0.1
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: ROUTE
Interface: local, Groups: 2
    Group: 224.0.0.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 224.0.0.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

Release Information

Command introduced before Junos OS Release 7.4.

Command updated in Junos OS Release 24.4R1 to include multicast support monitoring for ACX7100 routers.

RELATED DOCUMENTATION

clear igmp membership | 3619

show igmp interface

IN THIS SECTION

- Syntax | 20142
- Syntax (EX Series Switches and the QFX Series) | 20142
- Description | 20142
- Options | 20143
- Required Privilege Level | 20143
- Output Fields | 20143
- Sample Output | 20146
- show igmp interface (User Plane output for ACX7100) | 20147
- Release Information | 20148

Syntax

```
show igmp interface
<brief | detail>
<interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and the QFX Series)

```
show igmp interface
<brief | detail>
<interface-name>
```

Description

Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.

Options

none	Display standard information about all IGMP-enabled interfaces.
brief | detail	(Optional) Display the specified level of output.
interface-name	(Optional) Display information about the specified IGMP-enabled interface only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 595 on page 20143 describes the output fields for the `show igmp interface` command. Output fields are listed in the approximate order in which they appear.

Table 595: show igmp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the routing device that has been elected to send membership queries.	All levels
State	State of the interface: **Up** or **Down**.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.	All levels
Timeout	How long until the IGMP querier is declared to be unreachable, in seconds.	All levels
Version	IGMP version being used on the interface: **1** , **2** , or **3**.	All levels

Table 595: show igmp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Groups	Number of groups on the interface.	All levels
Group limit	Maximum number of groups allowed on the interface. Any joins requested after the limit is reached are rejected.	All levels
Group threshold	Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.	All levels
Group log-interval	Time (in seconds) between consecutive log messages.	All levels
Immediate Leave	State of the immediate leave option: • **On**—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface. • **Off**—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.	All levels
Promiscuous Mode	State of the promiscuous mode option: • **On**—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces. • **Off**—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces.	All levels
Distributed	State of IGMP, which, by default, takes place on the Routing Engine for MX Series routers but can be distributed to the Packet Forwarding Engine to provide faster processing of join and leave events. • **On**—distributed IGMP is enabled.	All levels

Table 595: show igmp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Passive	State of the passive mode option: • **On**—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves. • **Off**—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves. The `passive` statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the **on** state declaration: • **send-general-query**—The interface sends general queries. • **send-group-query**—The interface sends group-specific and group-source-specific queries. • **allow-receive**—The interface receives control traffic.	All levels
OIF map	Name of the OIF map (if configured) associated with the interface.	All levels
SSM map	Name of the source-specific multicast (SSM) map (if configured) used on the interface.	All levels
Configured Parameters	Information configured by the user: • **IGMP Query Interval**—Interval (in seconds) at which this router sends membership queries when it is the querier. • **IGMP Query Response Interval**—Time (in seconds) that the router waits for a report in response to a general query. • **IGMP Last Member Query Interval**—Time (in seconds) that the router waits for a report in response to a group-specific query. • **IGMP Robustness Count**—Number of times the router retries a query.	All levels

Table 595: show igmp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Derived Parameters	Derived information: • **IGMP Membership Timeout**—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed. • **IGMP Other Querier Present Timeout**—Time (in seconds) that the router waits for the IGMP querier to send a query.	All levels

Sample Output

show igmp interface

```
user@host> show igmp interface
Interface: at-0/3/1.0
    Querier: 203.0.3.113.31
    State:         Up Timeout:    None Version:  2 Groups:      4
    SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
    Querier: 203.0.113.11
    State:         Up Timeout:    None Version:  2 Groups:      2
    SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
    Querier: 203.0.113.21
    State:         Up Timeout:    None Version:  2 Groups:      4
    SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off
Passive: Off
Distributed: OnConfigured Parameters:

IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
```

```
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

show igmp interface brief

The output for the `show igmp interface brief` command is identical to that for the `show igmp interface` command. For sample output, see "show igmp interface" on page 20146.

show igmp interface detail

The output for the `show igmp interface detail` command is identical to that for the `show igmp interface` command. For sample output, see "show igmp interface" on page 20146.

show igmp interface <interface-name>

```
user@host# show igmp interface ge-3/2/0.0
Interface: ge-3/2/0.0
  Querier: 203.0.113.111
  State: Up Timeout:    None
  Version:  3
  Groups:      1
  Group limit: 8
  Group threshold: 60
  Group log-interval: 10
  Immediate leave:  Off
  Promiscuous mode:  Off
  Distributed: On
```

show igmp interface (User Plane output for ACX7100)

Displays the following output in the user-plane for customer premise equipment (CPE).

```
•    show igmp interface

user@up1> show igmp interface
Interface: pp0.3221225481
    Querier: 128.220.14.35
    State:         Up Timeout:    1000 Version:  3 Groups:      1
    Group threshold:   0
```

```
    Group log-interval:  10000
    Immediate leave: Off
    Promiscuous mode: On
    Passive: Off
Interface: demux0.2147483652
    Querier: 0.0.0.0
    State:         Up Timeout:    None Version:  3 Groups:      1
    Immediate leave: On
    Promiscuous mode: On
    Passive: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

Release Information

Command introduced before Junos OS Release 7.4.

Command updated in Junos OS Release 24.4R1 to include multicast support monitoring for ACX7100 router interfaces.

RELATED DOCUMENTATION

clear igmp membership | 3619

show igmp snooping data-forwarding

IN THIS SECTION

Syntax | 20149

- Description | 20149
- Options | 20149
- Required Privilege Level | 20149
- Output Fields | 20149
- Sample Output | 20151
- Release Information | 20152

Syntax

```
show igmp snooping data-forwarding
<vlan vlan-name>
```

Description

Display multicast source VLAN (MVLAN) and data-forwarding receiver VLAN associations and related information when you configure multicast VLAN registration (MVR) in a routing instance.

Options

vlan *vlan-name* (Optional) Display configured MVR information about a particular VLAN only.

Required Privilege Level

view

Output Fields

Table 596 on page 20150 lists the output fields for the `show igmp snooping data-forwarding` command. Output fields are listed in the approximate order in which they appear.

Table 596: show igmp snooping data-forwarding Output Fields

Field Name	Field Description
`Instance`	Routing instance in which MVR is configured.
`Vlan`	VLAN names of the multicast source and receiver VLANs configured in the routing instance.
`Learning Domain`	Learning domain for snooping and MVR data forwarding.
`Type`	MVR VLAN type configured for the listed VLAN, either `MVR Receiver Vlan` or `MVR Source Vlan`.
`Group subnet`	Group subnet address for the multicast source VLAN in the MVR configuration (the MVLAN).
`Receiver vlans`	Multicast receiver VLANs associated with the MVLAN. When you configure a source MVLAN, you associate one or more MVR receiver VLANs with it.
`Mode`	MVR operating mode configured for the listed receiver VLAN: • `PROXY`—MVR receiver VLAN is in proxy mode. • `TRANSPARENT`—MVR receiver VLAN is in transparent mode. See "mode (Multicast VLAN Registration)" on page 10809.
`Egress translate`	VLAN tag translation setting for an MVR receiver VLAN: • `TRUE`—The `translate` option at the `[edit protocols igmp-snooping vlans` *`vlan-name`* `data-forwarding receiver]` hierarchy level is configured for the MVR receiver VLAN. With this option enabled, the device translates MVLAN tags into the MVR receiver VLAN tag when forwarding multicast traffic on the MVLAN. • `FALSE`—The `translate` option for VLAN tag translation is not configured for the MVR receiver VLAN. MVLAN traffic is forwarded with the MVLAN tag for receivers on trunk ports or untagged for hosts on access ports.
`Install route`	If `TRUE`, the device installs forwarding entries for the MVR receiver VLAN as well as for the MVLAN. If `FALSE`, only MVLAN forwarding entries are stored.

Table 596: show igmp snooping data-forwarding Output Fields *(Continued)*

Field Name	Field Description
Source vlans	One or more source MVLANs associated with the listed MVR receiver VLAN.

Sample Output

show igmp snooping data-forwarding

```
user@host> show igmp snooping data-forwarding
Instance: default-switch

Vlan: v2

Learning-Domain     : default
Type                : MVR Source Vlan
Group subnet        : 225.0.0.0/24
Receiver vlans:
    vlan: v1
    vlan: v3

Vlan: v1

Learning-Domain     : default
Type                : MVR Receiver Vlan
Mode                : PROXY
Egress translate    : FALSE
Install route       : FALSE
Source vlans:
    vlan: v2

Vlan: v3

Learning-Domain     : default
Type                : MVR Receiver Vlan
Mode                : TRANSPARENT
Egress translate    : FALSE
Install route       : TRUE
Source vlans:
```

```
    vlan: v2
```

show igmp snooping data-forwarding (vlan)

```
user@host> show igmp snooping data-forwarding vlan v1
Instance: default-switch

Vlan: v1

Learning-Domain     : default
Type                : MVR Receiver Vlan
Mode                : PROXY
Egress translate    : FALSE
Install route       : FALSE
Source vlans:
    vlan: v2
```

Release Information

Command introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 18.4R1 on EX2300 and EX3400 switches.

Support added in Junos OS Release 19.4R1 on EX4300 multigigabit switches.

RELATED DOCUMENTATION

Configuring Multicast VLAN Registration on EX Series Switches

show igmp snooping interface | **20157**

show igmp snooping membership | **20165**

show igmp snooping evpn status

IN THIS SECTION

- Syntax | 20153
- Description | 20153
- Options | 20153
- Required Privilege Level | 20154
- Output Fields | 20154
- Sample Output | 20156
- Release Information | 20157

Syntax

```
show igmp snooping evpn status
<detail>
<vlan name>
```

Description

Display Layer 2 multicast IGMP snooping information for the EVPN instances on devices in an EVPN network.

By default, this command displays brief output for all EVPN instances with IGMP snooping enabled for one or more VLANs.

Options

(brief | detail) Display brief information (the default output level) or more detailed information.

When you enable the optimized intersubnet multicast (OISM) feature, the `detail` option provides more about the different bridge domains (VLANs) OISM uses in the EVPN network.

vlan *name* Display IGMP snooping information only for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 597: show igmp snooping evpn status Output Fields

Field Name	Field Description	Output Level
Bridge-Domain	Bridge domain or VLAN with IGMP snooping enabled	all
External VLAN	OISM external VLAN: • Yes—OISM is enabled and uses this VLAN to send multicast traffic to an external multicast receiver or receive multicast traffic from an external multicast source. • No—OISM is not enabled, or OISM is enabled but doesn't use this VLAN to exchange traffic with an external multicast domain. An OISM external VLAN must be different from the SBDs or OISM revenue VLANs in the EVPN fabric. See *OISM Bridge Domains (VLANs)* for details.	detail
Instance	EVPN instance (EVI) name	all

Table 597: show igmp snooping evpn status Output Fields *(Continued)*

Field Name	Field Description	Output Level
Mode	Configured multicast replication mode: • Enhanced OISM—Enhanced OISM (the asymmetric bridge domains OISM model, also called the bridge domains not everywhere [BDNE] version) • Default—Any other mode besides enhanced OISM, including regular OISM (the symmetric bridge domains OISM model, also called the bridge domains everywhere [BDE] version). To set the multicast mode in EVPN networks, you use the `irb (local-only | local-remote | oism | enhanced-oism)` statement at the `[edit forwarding-options multicast-replication evpn]` hierarchy level. See "multicast-replication" on page 11080 for details.	default (brief)
OISM	OISM status—Enabled or not enabled	all
SMET enabled	Selective multicast Ethernet tag (SMET) forwarding using EVPN Type 6 routes—Yes (enabled) or No (not enabled)	default (brief)
Supplementary BD	OISM supplemental bridge domain (SBD): • Yes—OISM is enabled and uses this bridge domain or VLAN as an OISM SBD. • No—OISM is not enabled, or OISM is enabled but doesn't use this bridge domain or VLAN as an OISM SBD. An OISM SBD must be a VLAN that is different from any OISM revenue VLANs or external VLANs in the EVPN network. See *OISM Bridge Domains (VLANs)* for details.	detail
VN Identifier	Virtual network identifier (VNI) corresponding to this VLAN	all

Sample Output

show igmp snooping evpn status (Regular OISM enabled)

```
user@SL1> show igmp snooping evpn status
Mode: Default
SMET enabled: Yes

Instance: default-switch
    Bridge-Domain: VLAN-600, VN Identifier: 600
        OISM: Enabled
    Bridge-Domain: VLAN-601, VN Identifier: 601
        OISM: Enabled
    Bridge-Domain: VLAN-602, VN Identifier: 602
        OISM: Enabled
    Bridge-Domain: VLAN-603, VN Identifier: 603
        OISM: Enabled
    Bridge-Domain: VLAN-901, VN Identifier: 901
        OISM: Enabled
```

show igmp snooping evpn status (MAC-VRF EVPN instance and enhanced OISM enabled)

```
user@SL1> show igmp snooping evpn status
Mode: Enhanced OISM
SMET enabled: Yes

Instance: v100
    Bridge-Domain: vlan100, VN Identifier: 10100
        OISM: Enabled
    Bridge-Domain: vlan200, VN Identifier: 10200
        OISM: Enabled
    Bridge-Domain: vlan999, VN Identifier: 10999
        OISM: Enabled
```

show igmp snooping evpn status vlan VLAN-300 detail (default-switch EVPN instance, regular OISM enabled, and specified VLAN is the OISM SBD)

```
user@Leaf-1> show igmp snooping evpn status vlan VLAN-900 detail
Instance: default-switch
  Bridge-Domain: VLAN-900, VN Identifier: 11900
        OISM            : Enabled
        Supplementary BD: Yes
        External VLAN   : No
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Optimized Inter-Subnet Multicast in EVPN Networks

show igmp snooping interface

IN THIS SECTION

- Syntax | 20158
- Description | 20158
- Options | 20158
- Required Privilege Level | 20158
- Output Fields | 20158
- Sample Output | 20160
- Release Information | 20165

Syntax

```
show igmp snooping interface interface-name
<brief | detail>
<bridge-domain bridge-domain-name>
<logical-system logical-system-name>
<virtual-switch virtual-switch-name>
<vlan-id vlan-identifier>
```

Description

Display IGMP snooping interface information.

Options

none	Display detailed information.
brief | detail	(Optional) When applicable, this option lets you choose the how much detail to display.
bridge-domain *bridge-domain-name*	(Optional) Display information about a particular bridge domain.
logical-system *logical-system-name*	(Optional) Display information about a particular logical system, or type 'all'.
virtual-switch *virtual-switch-name*	(Optional) Display information about a particular virtual switch.
vlan-id *vlan-identifier*	(Optional) Display information about a particular VLAN.

Required Privilege Level

view

Output Fields

Table 598 on page 20159 lists the output fields for the `show igmp snooping interface` command. Output fields are listed in the approximate order in which they appear.

Table 598: show igmp snooping interface Output Fields

Field Name	Field Description	Level of Output
Routing-instance or **Instance**	Routing instance for IGMP snooping.	All levels
Bridge Domain or **Vlan**	Bridge domain or VLAN for which IGMP snooping is enabled.	All levels
Learning Domain	Learning domain for snooping.	All levels
interface	Interfaces that are being snooped in this learning domain.	All levels
Groups	Number of groups on the interface.	All levels
State	State of the interface: **Up** or **Down**.	All levels
Up Groups	Number of active multicast groups attached to the logical interface.	All levels
immediate-leave	State of immediate leave: **On** or **Off**.	All levels
router-interface	Router interfaces that are part of this learning domain.	All levels
Group limit	Maximum number of (source,group) pairs allowed per interface. When a group limit is not configured, this field is not shown.	All levels
Data-forwarding receiver: yes	VLAN associated with the interface is configured as a data-forwarding multicast receiver VLAN using multicast VLAN registration (MVR) on EX Series switches with Enhanced Layer 2 Software (ELS).	All levels
IGMP Query Interval	Frequency (in seconds) with which this router sends membership queries when it is the querier.	All levels

Table 598: show igmp snooping interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IGMP Query Response Interval	Time (in seconds) that the router waits for a response to a general query.	All levels
IGMP Last Member Query Interval	Time (in seconds) that the router waits for a report in response to a group-specific query.	All levels
IGMP Robustness Count	Number of times the router retries a query.	All levels
IGMP Membeship Timeout	Timeout for group membership. If no report is received for these groups before the timeout expires, the group membership is removed.	All levels
IGMP Other Querier Present Timeout	Time that the router waits for the IGMP querier to send a query.	All levels

Sample Output

show igmp snooping interface

```
user@host> show igmp snooping interface ge-0/1/4
Instance: default-switch

Bridge-Domain: sample

Learning-Domain: default
Interface: ge-0/1/4.0
State: Up Groups: 0
Immediate leave: Off
Router interface: no
```

```
Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

show igmp snooping interface (logical systems)

```
user@host> show igmp snooping interface logical-system all
logical-system: default
Instance: VPLS-6
Learning-Domain: default
Interface: ge-0/2/2.601
    State:         Up Groups:     10
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Instance: VS-4
Bridge-Domain: VS-4-BD-1
Learning-Domain: vlan-id 1041
Interface: ae2.3
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no
Interface: ge-0/2/2.1041
    State:         Up Groups:     20
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
```

```
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Instance: default-switch
Bridge-Domain: bd-200
Learning-Domain: default
Interface: ge-0/2/2.100
    State:         Up Groups:     20
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Bridge-Domain: bd0
Learning-Domain: default
Interface: ae0.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: yes
Interface: ae1.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no
Interface: ge-0/2/2.0
    State:         Up Groups:     32
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Instance: VPLS-1
Learning-Domain: default
Interface: ge-0/2/2.502
    State:         Up Groups:     11
```

```
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Instance: VS-1
Bridge-Domain: VS-BD-1
Learning-Domain: default
Interface: ae2.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no
Interface: ge-0/2/2.1010
    State:         Up Groups:     20
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Bridge-Domain: VS-BD-2
Learning-Domain: default
Interface: ae2.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no
Interface: ge-0/2/2.1011
    State:         Up Groups:     20
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

```
Instance: VPLS-p2mp
Learning-Domain: default
Interface: ge-0/2/2.3001
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

show igmp snooping interface (Group Limit Configured)

```
user@host> show igmp snooping interface instance vpls1
Instance: vpls1

Learning-Domain: default
Interface: ge-1/3/9.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: yes
Interface: ge-1/3/8.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: yes
    Group limit:   1000

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

show igmp snooping interface (ELS EX Series switches with MVR configured)

```
user@host> show igmp snooping interface instance inst1
Instance: inst1
```

```
Vlan: v2

Learning-Domain: default
Interface: ge-0/0/0.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no
    Group limit:      3
    Data-forwarding receiver: yes
```

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

show igmp snooping membership | 20165

show igmp snooping statistics | 20187

show igmp snooping membership

IN THIS SECTION

- Syntax | 20166
- Description | 20166
- Options | 20166
- Required Privilege Level | 20167
- Output Fields | 20167
- Sample Output | 20169
- Release Information | 20173

Syntax

```
show igmp snooping membership
<brief | detail>
<instance routing-instance-name>
<interface interface-name>
<vlan (vlan-id | vlan-name)>
<bridge-domain bridge-domain-name>
<group group-name>
<logical-system logical-system-name>
<virtual-switch virtual-switch-name>
<vlan-id vlan-identifier>
```

Description

Display the multicast group membership information maintained by IGMP snooping.

Options

none

Display the multicast group membership information about all VLANs on which IGMP snooping is enabled.

brief | detail

(Optional) Display the specified level of output. The default is `brief`.

> **NOTE**: On QFX Series switches, the output is the same for either`brief` or `detail` levels.

instance *routing-instance-name*

(Optional) Display the multicast group membership information about the specified routing instance.

interface *interface-name*

(Optional) Display the multicast group membership information about the specified interface.

vlan (*vlan-id* | *vlan-name*)

(Optional) Display the multicast group membership for the specified VLAN.

bridge-domain *bridge-domain-name*

(Optional) Display information about a particular bridge domain.

group *group-name*

(Optional) Display information about this group address.

logical-system *logical-system-name*

(Optional) Display information about a particular logical system, or type 'all'.

virtual-switch ***virtual-switch-name***	(Optional) Display information about a particular virtual switch. You can also use this option to display IGMP snooping information about a particular routing instance of type "mac-vrf" on page 10276.
vlan-id ***vlan-identifier***	(Optional) Display information about a particular VLAN.

Required Privilege Level

view

Output Fields

Table 599 on page 20167 lists the output fields for the `show igmp snooping membership` command. Output fields are listed in the approximate order in which they appear.

Table 599: show igmp snooping membership Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All
Instance	Routing instance for IGMP snooping.	All levels
Learning Domain	Learning domain for snooping.	All levels
Interface	Interface on which this router is a proxy.	**detail**
Data-forwarding receiver: yes	(EX Series switches with Enhanced Layer 2 Software (ELS) only) VLAN associated with the interface is configured as a data-forwarding multicast receiver VLAN using multicast VLAN registration (MVR). **NOTE**: Interfaces configured on MVR receiver VLANs are listed under the associated MVR *source* VLAN (MVLAN) for which the interface forwards multicast streams.	All levels
Up Groups or **Groups**	Number of active multicast groups attached to the logical interface.	All levels

Table 599: show igmp snooping membership Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Group	(Not displayed on QFX Series switches) IP multicast address of the multicast group. The following information is provided for the multicast group: • `Last reporter`—Last host to report membership for the multicast group. • `Receiver count`—Number of hosts on the interface that are members of the multicast group (field appears only if `immediate-leave` is configured on the VLAN), or number of interfaces that have membership in a multicast group. • `Uptime`—Length of time (in hours, minutes, and seconds) a multicast group has been active on the interface. • `timeout`—Time (in seconds) left until the entry for the multicast group is removed from the multicast group if no membership reports are received on the interface. This counter is reset to its maximum value when a membership report is received. • `Flags`—The lowest IGMP version in use by a host that is a member of the group on the interface. • `Include source`—Source addresses from which multicast streams are allowed based on IGMPv3 reports.	detail
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.	All levels
Source	Source address used on queries.	All levels
Last reported by	Address of source last replying to the query.	All levels
Group Timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	All levels
Timeout	Length of time (in seconds) left until the entry is purged.	**detail**

Table 599: show igmp snooping membership Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Type	Way that the group membership information was learned: • **Dynamic**—Group membership was learned by the IGMP protocol. • **Static**—Group membership was learned by configuration.	All levels
Include receiver	Source address of receiver included in membership with timeout (in seconds).	**detail**

Sample Output

show igmp snooping membership

```
user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
Up Groups:      0
Interface: ge-3/1/0.2
Up Groups:      0
Interface: ge-3/1/5.2
Up Groups:      0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
    Group: 233.252.0.99
        Group mode: Exclude
        Source: 0.0.0.0
```

```
        Last reported by: 233.252.0.87
        Group timeout:     173 Type: Dynamic
```

show igmp snooping membership (SRX1500)

```
user@host> show igmp snooping membership
Instance: default-switch

Vlan: v1

Learning-Domain: default
Interface: ge-0/0/3.0, Groups: 1
Group: 233.252.0.100
Group mode: Exclude
Source: 0.0.0.0
Last reported by: Local
Group timeout: 0 Type: Static
```

show igmp snooping membership detail (SRX1500)

```
user@host> show igmp snooping membership detail

VLAN: vlan2 Tag: 2 (Index: 3)
    Router interfaces:
        ge-1/0/0.0 dynamic Uptime: 00:14:24 timeout: 253
    Group: 233.252.0.99
      ge-1/0/17.0 259 Last reporter: 10.0.0.90 Receiver count: 1
      Uptime: 00:00:19 timeout: 259 Flags: <V3-hosts>
        Include source: 10.2.11.5, 10.2.11.12
```

show igmp snooping membership (Exclude Mode)

```
user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
```

```
Up Groups:      0
Interface: ge-3/1/0.2
Up Groups:      0
Interface: ge-3/1/5.2
Up Groups:      0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
    Group: 233.252.0.99
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: 233.252.0.87
        Group timeout:     173 Type: Dynamic
```

show igmp snooping membership interface ge-0/1/2.200

```
user@host> show igmp snooping membership interface ge-0/1/2.200
Instance: bridge-domain bar

Learning-Domain: default
Interface: ge-0/1/2.200
    Group: 233.252.0.1
        Source: 0.0.0.0
        Timeout:     391 Type: Static
    Group: 232.1.1.1
        Source: 192.128.1.1
        Timeout:       0 Type: Static
```

show igmp snooping membership vlan-id 1

```
user@host> show igmp snooping membership vlan-id 1
Instance: vpls2
```

```
Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
    Group: 233.252.0.1
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: 233.252.0.82
        Group timeout:     209 Type: Dynamic
```

show igmp snooping membership (ELS EX Series switches with MVR)

```
user@host> show igmp snooping membership
Instance: default-switch

Vlan: v2

Learning-Domain: default
Interface: ge-0/0/0.0, Groups: 0
Data-forwarding receiver: yes

Learning-Domain: default
Interface: ge-0/0/12.0, Groups: 1
    Group: 233.252.0.1
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: Local
        Group timeout:       0 Type: Static
```

show igmp snooping membership <detail> (QFX5100 switches—same output with or without detail option)

```
user@host> show igmp snooping membership detail
Instance: default-switch

Vlan: v100

Learning-Domain: default
Interface: xe-0/0/51:0.0, Groups: 1
    Group: 233.252.0.1
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: 233.252.0.82
        Group timeout:     251 Type: Dynamic
```

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

- show igmp snooping interface | 20157
- show igmp snooping statistics | 20187
- clear igmp snooping membership | 3623

show igmp-snooping membership

IN THIS SECTION

- Syntax | 20174
- Description | 20174
- Options | 20174

- Required Privilege Level | 20174
- Output Fields | 20175
- Sample Output | 20177
- Release Information | 20179

Syntax

```
show igmp-snooping membership
<brief | detail>
<interface interface-name>
<vlan vlan-id | vlan-name>
```

Description

Display the multicast group membership information maintained by IGMP snooping.

> **NOTE**: To display similar information on routing devices or switches that support the Enhanced Layer 2 Software (ELS) configuration style, use the equivalent command "show igmp snooping membership" on page 20165.

Options

none	Display general parameters.
brief | detail	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display IGMP snooping information for the specified interface.
vlan *vlan-id* | *vlan-name*	(Optional) Display IGMP snooping information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 600 on page 20175 lists the output fields for the `show igmp-snooping membership` command. Output fields are listed in the approximate order in which they appear.

Table 600: show igmp-snooping membership Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All
Interfaces	Interfaces that are members of the listed multicast group.	All
Tag	Numerical identifier of the VLAN.	detail
Router interfaces	List of information about multicast router interfaces: • Name of the multicast router interface. • `static` or `dynamic`—Whether the multicast router interface is statically or dynamically assigned. • `Uptime`—For static interfaces, amount of time since the interface was configured as a multicast-router interface or since the interface last flapped. For dynamic interfaces, amount of time since the first query was received on the interface or since the interface last flapped. • `timeout`—Query timeout in seconds.	detail

Table 600: show igmp-snooping membership Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Group	IP multicast address of the multicast group. The following information is provided for the multicast group: • `Last reporter`—Last host to report membership for the multicast group. • `Receiver count`—Number of hosts on the interface that are members of the multicast group (field appears only if `immediate-leave` is configured on the VLAN), or number of interfaces that have membership in a multicast group. • `Uptime`—Length of time (in hours, minutes, and seconds) a multicast group has been active on the interface. • `timeout`—Time (in seconds) left until the entry for the multicast group is removed from the multicast group if no membership reports are received on the interface. This counter is reset to its maximum value when a membership report is received.	detail

Table 600: show igmp-snooping membership Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Flags`—The lowest IGMP version in use by a host that is a member of the group on the interface. • `Include source`—Source addresses from which multicast streams are allowed based on IGMPv3 reports.	

Sample Output

show igmp-snooping membership

```
user@switch> show igmp-snooping membership
VLAN: v1
    224.1.1.1      *               258 secs
        Interfaces: ge-0/0/0.0
    224.1.1.3      *               258 secs
        Interfaces: ge-0/0/0.0
    224.1.1.5      *               258 secs
        Interfaces: ge-0/0/0.0
    224.1.1.7      *               258 secs
        Interfaces: ge-0/0/0.0
    224.1.1.9      *               258 secs
        Interfaces: ge-0/0/0.0
    224.1.1.11     *               258 secs
        Interfaces: ge-0/0/0.0
```

show igmp-snooping membership detail (QFX Series)

```
user@switch> show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
   Group: 225.0.0.2
```

```
   Receiver count: 1, Flags: <V3-hosts>
     ge-0/0/15.0 Uptime: 00:00:11 timeout: 248 Last reporter: 10.2.10.16
       Include source: 1.2.1.1, 1.3.1.1
VLAN: v44 Tag: 44 (Index: 5)
   Group: 225.0.0.1
   Receiver count: 1, Flags: <V2-hosts>
     ge-0/0/21.0 Uptime: 00:00:02 timeout: 257
VLAN: v110 Tag: 110 (Index: 4)
 Router interfaces:
         ge-0/0/3.0 static Uptime: 00:08:45
         ge-0/0/2.0 static Uptime: 00:08:45
         ge-0/0/4.0 dynamic Uptime: 00:16:41 timeout: 254
     Group: 225.0.0.3
     Receiver count: 1, Flags: <V3-hosts>
         ge-0/0/5.0 Uptime: 00:00:19 timeout: 259
     Group: 225.1.1.1
     Receiver count: 1, Flags: <V2-hosts>
         ge-0/0/5.0 Uptime: 00:22:43 timeout: 96
     Group: 225.2.2.2
     Receiver count: 1, Flags: <V2-hosts Static>
         ge-0/0/5.0 Uptime: 00:23:13
```

show igmp-snooping membership detail (EX Series)

```
user@switch> show igmp-snooping membership detail
VLAN: vlan2 Tag: 2 (Index: 3)
    Router interfaces:
        ge-1/0/0.0 dynamic Uptime: 00:14:24 timeout: 253
    Group: 233.252.0.99
      ge-1/0/17.0 259 Last reporter: 10.0.0.90 Receiver count: 1
      Uptime: 00:00:19 timeout: 259 Flags: <V3-hosts>
        Include source: 10.2.11.5, 10.2.11.12
```

show igmp-snooping membership vlan detail (EX Series)

```
user@switch> show igmp-snooping membership vlan vlan700 detail
VLAN: vlan700 Tag: 700 (Index: 52)
    Router interfaces:
        ae2.0 dynamic Uptime: 16:53:13 timeout: 245
```

```
    Group: 233.252.0.1
50      ge-0/0/1.0 Last reporter: 10.2.188.201
      Uptime: 17:00:52 timeout: 237 Flags: <V2-hosts>
        ge-0/0/0.0 Last reporter: 10.2.188.202
      Uptime: 17:00:50 timeout: 243 Flags: <V2-hosts>
```

Release Information

Command introduced in Junos OS Release 9.1.

IGMPv3 output introduced in Junos OS Release 12.1 for the QFX Series.

RELATED DOCUMENTATION

Monitoring IGMP Snooping

Configuring IGMP Snooping on Switches

show igmp-snooping route | **20181**

show igmp-snooping statistics | **20196**

show igmp-snooping vlans | **20198**

show igmp snooping options

IN THIS SECTION

- Syntax | **20180**
- Description | **20180**
- Options | **20180**
- Required Privilege Level | **20180**
- Sample Output | **20180**
- Release Information | **20181**

Syntax

```
show igmp snooping options
<brief | detail>
instance <instance-name>
<logical-system logical-system-name>
```

Description

Show the operational status of point-to-multipoint LSP for IGMP snooping routes.

Options

brief | detail	Display the specified level of output per routing instance. The default is brief.
instance-name	(Optional) Output for the specified routing instance only.
logical-system *logical-system-name*	(Optional) Display information about a particular logical system, or type 'all'.

Required Privilege Level

view

Sample Output

show igmp snooping options

```
user@host> show igmp snooping options

Instance: master
    P2MP LSP in use: no
Instance: default-switch
    P2MP LSP in use: no
Instance: name
    P2MP LSP in use: yes
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Configuring Point-to-Multipoint LSP with IGMP Snooping

use-p2mp-lsp | **30530**

multicast-snooping-options

show igmp-snooping route

IN THIS SECTION

- Syntax | **20181**
- Description | **20182**
- Options | **20182**
- Required Privilege Level | **20182**
- Output Fields | **20182**
- Sample Output | **20183**
- Release Information | **20184**

Syntax

```
show igmp-snooping route
<brief | detail>
<ethernet-switching <brief | detail | vlan (vlan-id | vlan-name)>>
<inet <brief | detail | vlan vlan-name>>
<vlan vlan-name>
```

Description

Display IGMP snooping route information.

> **NOTE**: This command is only available on switches that do not support the Enhanced Layer 2 Software (ELS) configuration style.

Options

none	Display general route information for all VLANs on which IGMP snooping is enabled.
brief | detail	(Optional) Display the specified level of output. The default is `brief`.
ethernet-switching	(Optional) Display information on Layer 2 multicast routes. This is the default.
inet	(Optional) Display information for Layer 3 multicast routes.
vlan *vlan-name*	(Optional) Display route information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 601 on page 20182 lists the output fields for the `show igmp-snooping route` command. Output fields are listed in the approximate order in which they appear. Some output fields are not displayed by this command on some devices.

Table 601: show igmp-snooping route Output Fields

Field Name	Field Description
Table	Routing table ID for virtual routing instances, or 0 on devices where this is not used.
Routing Table	Routing table ID for virtual routing instances.

Table 601: show igmp-snooping route Output Fields *(Continued)*

Field Name	Field Description
VLAN	Name of the VLAN for which IGMP snooping is enabled.
Group	Multicast IPv4 group address.
Interface or Interfaces	Name of the interface or interfaces in the VLAN associated with the multicast group.
Next-hop	ID associated with the next-hop device.
Layer 2 next-hop	ID associated with the Layer 2 next-hop device.
Routing next-hop	ID associated with the Layer 3 next-hop device.

Sample Output

show igmp-snooping route

```
user@switch> show igmp-snooping route
VLAN             Group          Next-hop
V11              224.1.1.1, *        533
        Interfaces: ge-0/0/13.0, ge-0/0/1.0
VLAN             Group          Next-hop
v12              224.1.1.3, *        534
        Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

show igmp-snooping route vlan v1

```
user@switch> show igmp-snooping route vlan v1
Table: 0
VLAN             Group          Next-hop
v1               224.1.1.1, *       1266
        Interfaces: ge-0/0/0.0
```

```
v1              224.1.1.3, *       1266
        Interfaces: ge-0/0/0.0
v1              224.1.1.5, *       1266
        Interfaces: ge-0/0/0.0
v1              224.1.1.7, *       1266
        Interfaces: ge-0/0/0.0
v1              224.1.1.9, *       1266
        Interfaces: ge-0/0/0.0
v1              224.1.1.11, *      1266
        Interfaces: ge-0/0/0.0
```

show igmp-snooping route detail

```
user@switch> show igmp-snooping route detail
 VLAN            Group          Next-hop
 default         233.252.0.0, *
 vlan100         233.252.0.0, *       1332
         Interfaces: ge-1/0/1.0
 VLAN            Group          Next-hop
 vlan100         233.252.0.1, *       1334
         Interfaces: ge-1/0/1.0, ge-5/0/30.0
```

show igmp-snooping route inet detail

```
user@switch> show igmp-snooping route inet detail
Routing table: 0
Group: 233.252.0.1, 192.168.60.100
    Routing next-hop: 3448
        vlan.100
    Interface: vlan.100, VLAN: vlan100, Layer 2 next-hop: 3343
```

Release Information

Command introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Monitoring IGMP Snooping

Configuring IGMP Snooping on Switches

show igmp-snooping statistics | **20196**

show igmp-snooping vlans | **20198**

show igmp-snooping route (View)

IN THIS SECTION

- Syntax | **20185**
- Description | **20185**
- Options | **20185**
- Required Privilege Level | **20186**
- Output Fields | **20186**
- Sample Output | **20186**
- Release Information | **20187**

Syntax

```
show igmp-snooping route ( brief | detail | ethernet-switching | inet | vlan)
```

Description

Display IGMP snooping route information.

Options

- `none`—Display general parameters.
- `brief` | `detail`—(Optional) Display the specified level of output.
- `ethernet-switching`—(Optional) Display Ethernet switching information.
- `inet`—(Optional) Display inet information.

- `vlan` *vlan-id* | *vlan-name*—(Optional) Display route information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 602 on page 20186 lists the output fields for the `show igmp-snooping route` command. Output fields are listed in the approximate order in which they appear.

Table 602: show igmp-snooping route Output Fields

Field Name	Field Description
`VLAN`	Name of the VLAN.
`Group`	Multicast group address.
`Next-hop`	ID associated with the next-hop device.

Sample Output

show igmp-snooping route

```
user@host> show igmp-snooping route
VLAN       Group          Next-hop
v11         203.0.113.0, *   533
Interfaces: ge-0/0/13.0, ge-0/0/1.0
v12         203.0.113.1, *   534
Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

show igmp-snooping route vlan v1

```
user@host> show igmp-snooping route vlan v1
Table: 0
VLAN      Group              Next-hop
```

```
v1        203.0.113.2, *       1266
Interfaces: ge-0/0/0.0
v1        203.0.113.3, *       1266
Interfaces: ge-0/0/0.0
v1        203.0.113.4, *       1266
Interfaces: ge-0/0/0.0
v1        203.0.113.5, *       1266
Interfaces: ge-0/0/0.0
v1        203.0.113.6, *       1266
Interfaces: ge-0/0/0.0
v1        203.0.113.6, *      1266
Interfaces: ge-0/0/0.0
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Interfaces

show igmp snooping statistics

IN THIS SECTION

- Syntax | 20188
- Description | 20188
- Options | 20188
- Required Privilege Level | 20188
- Output Fields | 20188
- Sample Output | 20191
- Release Information | 20195

Syntax

```
show igmp snooping statistics
<brief | detail>
<bridge-domain bridge-domain-name>
<logical-system logical-system-name>
<virtual-switch virtual-switch-name>
<vlan-id vlan-identifier>
```

Description

Display IGMP snooping statistics.

Options

none	(Optional) Display detailed information.
brief | detail	(Optional) Display the specified level of output.
bridge-domain ***bridge-domain-name***	(Optional) Display information about a particular bridge domain.
logical-system ***logical-system-name***	(Optional) Display information about a particular logical system, or type 'all'.
virtual-switch ***virtual-switch-name***	(Optional) Display information about a particular virtual switch.
vlan-id ***vlan-identifier***	(Optional) Display information about a particular VLAN.

Required Privilege Level

view

Output Fields

Table 603 on page 20189 lists the output fields for the `show igmp snooping statistics` command. Output fields are listed in the approximate order in which they appear.

Table 603: show igmp snooping statistics Output Fields

Field Name	Field Description	Level of Output
Routing-instance	Routing instance for IGMP snooping.	All levels
IGMP packet statistics	Heading for IGMP snooping statistics for all interfaces or for the specified interface.	All levels
learning-domain	Appears at end of “IGMP packets statistics” line.	All levels

Table 603: show igmp snooping statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IGMP Message type	Summary of IGMP statistics: • **Membership Query**—Number of membership queries sent and received. • **V1 Membership Report**—Number of version 1 membership reports sent and received. • **DVMRP**—Number of DVMRP messages sent or received. • **PIM V1**—Number of PIM version 1 messages sent or received. • **Cisco Trace**—Number of Cisco trace messages sent or received. • **V2 Membership Report**—Number of version 2 membership reports sent or received. • **Group Leave**—Number of group leave messages sent or received. • **Domain Wide Report**—Number of domain-wide reports sent or received. • **V3 Membership Report**—Number of version 3 membership reports sent or received. • **Other Unknown types**—Number of unknown message types received. • **IGMP v3 unsupported type**—Number of messages received with unknown and unsupported IGMP version 3 message types. • **IGMP v3 source required for SSM**—Number of IGMP version 3 messages received that contained no source. • **IGMP v3 mode not applicable for SSM**—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).	All levels
Received	Number of messages received.	All levels
Sent	Number of messages sent.	All levels

Table 603: show igmp snooping statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Rx errors	Number of received packets that contained errors.	All levels
IGMP Global Statistics	Summary of IGMP snooping statistics for all interfaces. • **Bad Length**—Number of messages received with length errors so severe that further classification could not occur. • **Bad Checksum**—Number of messages received with a bad IP checksum. No further classification was performed. • **Rx non-local**—Number of messages received from senders that are not local.	All levels

Sample Output

show igmp snooping statistics

```
user@host> show igmp snooping statistics
Routing-instance foo

 IGMP packet statistics for all interfaces in learning-domain vlan-100

 IGMP Message type      Received       Sent  Rx errors
 Membership Query             89         51          0
 V1 Membership Report          0          0          0
 DVMRP                         0          0          0
 PIM V1                        0          0          0
 Cisco Trace                   0          0          0
 V2 Membership Report        139          0          0
 Group Leave                   0          0          0
 Domain Wide Report            0          0          0
 V3 Membership Report        136          0          0
 Other Unknown types                                 0
 IGMP v3 unsupported type                            0
 IGMP v3 source required for SSM                    23
 IGMP v3 mode not applicable for SSM                 0
```

```
 IGMP Global Statistics
 Bad Length                    0
 Bad Checksum                  0
 Rx non-local                  0

Routing-instance bar

 IGMP packet statistics for all interfaces in learning-domain vlan-100

 IGMP Message type      Received       Sent  Rx errors
 Membership Query             89         51          0
 V1 Membership Report          0          0          0
 DVMRP                         0          0          0
 PIM V1                        0          0          0
 Cisco Trace                   0          0          0
 V2 Membership Report        139          0          0
 Group Leave                   0          0          0
 Domain Wide Report            0          0          0
 V3 Membership Report        136          0          0
 Other Unknown types                                 0
 IGMP v3 unsupported type                            0
 IGMP v3 source required for SSM                    23
 IGMP v3 mode not applicable for SSM                 0

 IGMP Global Statistics
 Bad Length                    0
 Bad Checksum                  0
 Rx non-local                  0
```

show igmp snooping statistics (SRX1500)

```
user@host> show igmp snooping statistics
Vlan: v1
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          0          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
```

```
Mtrace Response              0          0          0
Mtrace Request               0          0          0
Domain Wide Report           0          0          0
V3 Membership Report         0          0          0
Other Unknown types                                 0
```

show igmp snooping statistics logical-systems all

```
user@host> show igmp snooping statistics logical-systems all

logical-system: default
Bridge: VPLS-6
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          4          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                 0

Learning-Domain: vlan-id 1041 bridge-domain VS-4-BD-1
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          4          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                 0
```

```
Bridge: VPLS-p2mp
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          2          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                  0

Bridge: VS-BD-1
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          6          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                  0
```

show igmp snooping statistics interface (Bridge Domains Configured)

```
user@host> show igmp snooping statistics interface

 Bridge: bridge-domain1
 IGMP interface packet statistics for ge-2/0/8.0
 IGMP Message type      Received       Sent  Rx errors
 Membership Query              0          2          0
 V1 Membership Report          0          0          0
 DVMRP                         0          0          0
```

```
PIM V1                          0          0          0
Cisco Trace                     0          0          0
V2 Membership Report            0          0          0
Group Leave                     0          0          0
Mtrace Response                 0          0          0
Mtrace Request                  0          0          0
Domain Wide Report              0          0          0
V3 Membership Report            0          0          0
Other Unknown types                                   0

Bridge: bridge-domain2
IGMP interface packet statistics for ge-2/0/8.0
IGMP Message type      Received       Sent  Rx errors
Membership Query                0          2          0
V1 Membership Report            0          0          0
DVMRP                           0          0          0
PIM V1                          0          0          0
Cisco Trace                     0          0          0
V2 Membership Report            0          0          0
Group Leave                     0          0          0
Mtrace Response                 0          0          0
Mtrace Request                  0          0          0
Domain Wide Report              0          0          0
V3 Membership Report            0          0          0
Other Unknown types                                   0
```

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

- show igmp snooping interface | 20157
- show igmp snooping membership | 20165
- clear igmp snooping statistics | 3625

show igmp-snooping statistics

IN THIS SECTION

- Syntax | 20196
- Description | 20196
- Required Privilege Level | 20196
- Output Fields | 20196
- Sample Output | 20197
- Release Information | 20198

Syntax

```
show igmp-snooping statistics
```

Description

Display IGMP snooping statistics information.

NOTE: To display similar information on routing devices or switches that support the Enhanced Layer 2 Software (ELS) configuration style, use the equivalent command "show igmp snooping statistics" on page 20187.

Required Privilege Level

view

Output Fields

Table 604 on page 20197 lists the output fields for the `show igmp-snooping statistics` command. Output fields are listed in the approximate order in which they appear.

Table 604: show igmp-snooping statistics Output Fields

Field Name	Field Description
Bad length	IGMP packet has illegal or bad length.
Bad checksum	IGMP or IP checksum is incorrect.
Invalid interface	Packet was received through an invalid interface.
Not local	Number of packets received from senders that are not local, or 0 if not used (on some devices).
Receive unknown	Unknown IGMP type.
Timed out	Number of timeouts for all multicast groups, or 0 if not used (on some devices).
IGMP Type	Type of IGMP message (`Queries`, `Reports`, `Leaves`, or `Other`).
Received	Number of IGMP packets received.
Transmitted	Number of IGMP packets transmitted.
Recv Errors	Number of general receive errors, for packets received that did not conform to IGMP version 1 (IGMPv1), IGMPv2, or IGMPv3 standards.

Sample Output

show igmp-snooping statistics

```
user@switch> show igmp-snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 58

    IGMP Type           Received      Transmitted    Recv Errors
```

```
    Queries:                74295               0               0
    Reports:             18148423               0        16333523
    Leaves:                     0               0               0
    Other:                      0               0               0
```

Release Information

Command introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Monitoring IGMP Snooping

Configuring IGMP Snooping on Switches

show igmp-snooping route | **20181**

show igmp-snooping vlans | **20198**

show igmp-snooping vlans

IN THIS SECTION

- Syntax | **20199**
- Description | **20199**
- Options | **20199**
- Required Privilege Level | **20199**
- Output Fields | **20199**
- Sample Output | **20201**
- Release Information | **20202**

Syntax

```
show igmp-snooping vlans
<brief | detail>
<vlan vlan-id | vlan-name>
```

Description

Display IGMP snooping VLAN information.

> **NOTE**: To display similar information on routing devices or switches that support the Enhanced Layer 2 Software (ELS) configuration style, use equivalent commands such as "show igmp snooping interface" on page 20157.

Options

none	Display general IGMP snooping information for all VLANs on which IGMP snooping is enabled.
brief | detail	(Optional) Display the specified level of output. The default is brief.
vlan *vlan-id* | vlan *vlan-number*	(Optional) Display VLAN information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 605 on page 20200 lists the output fields for the `show igmp-snooping vlans` command. Output fields are listed in the approximate order in which they appear. Some output fields are not displayed by this command on some devices.

Table 605: show igmp-snooping vlans Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All levels
IGMP-L2-Querier	Source address for IGMP snooping queries (if switch is an IGMP querier)	All levels
Interfaces	Number of interfaces in the VLAN.	All levels
Groups	Number of groups in the VLAN to which the interface belongs.	All levels
MRouters	Number of multicast routers associated with the VLAN.	All levels
Receivers	Number of host receivers in the VLAN. Indicates how many VLAN interfaces would receive data because of IGMP membership.	All levels
RxVlans	Number of multicast VLAN registration (MVR) receiver VLANs configured for that MVR source VLAN.	
Tag	Numerical identifier of the VLAN (VLAN tag).	detail
tagged | untagged	Interface accepts tagged (802.1Q) packets for trunk mode and tagged-access mode ports, or untagged (native VLAN) packets for access mode ports.	detail
vlan-interface	Internal VLAN interface identifier or Layer 3 interface associated with the VLAN.	detail
Membership timeout	Membership timeout value.	detail
Querier timeout	Maximum length of time the switch waits to take over as IGMP querier if no query is received.	detail
Interface	Name of the interface.	detail

Table 605: show igmp-snooping vlans Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Reporters	Number of hosts on the interface that are current members of multicast groups. This field appears only when immediate-leave is configured on the VLAN.	detail
Router	Interface is a multicast router interface.	

Sample Output

show igmp-snooping vlans

```
user@switch> show igmp-snooping vlans
VLAN     Interfaces Groups MRouters Receivers
default     0           0      0         0
v1          11          50     0         0
v10         1           0      0         0
v11         1           0      0         0
v180        3           0      1         0
v181        3           0      0         0
v182        3           0      0         0
```

show igmp-snooping vlans vlan

```
user@switch> show igmp-snooping vlans vlan v10
user@switch> show igmp-snooping vlans vlan v10
VLAN       Interfaces Groups MRouters Receivers
v10        1          0      0        0
```

show igmp-snooping vlans detail

```
user@switch> show igmp-snooping vlans detail
VLAN: default, Tag: 0
Membership timeout: 54, Querier timeout: 52
```

```
VLAN: v2146-API, Tag: 2146, vlan-interface: vlan.2146
Membership timeout: 54, Querier timeout: 52
    Interface: ae0.0, tagged, Groups: 0, Reporters: 0,
    Interface: ge-7/0/21.0, untagged, Groups: 0, Reporters: 0
    Interface: ge-1/0/24.0, untagged, Groups: 0, Reporters: 0
    Interface: ge-1/0/25.0, untagged, Groups: 0, Reporters: 0
    Interface: ge-1/0/26.0, untagged, Groups: 0, Reporters: 0
    Interface: ge-1/0/36.0, untagged, Groups: 0, Reporters: 0
    Interface: ge-1/0/37.0, untagged, Groups: 0, Reporters: 0
    Interface: ge-1/0/38.0, untagged, Groups: 0, Reporters: 0
```

show igmp-snooping vlans vlan detail

```
user@switch> show igmp-snooping vlans vlan v10 detail
VLAN: v10, Tag: 10, vlan-interface: vlan.10
        Interface: ge-0/0/10.0, tagged, Groups: 0
IGMP-L2-Querier: Stopped, SourceAddress: 10.10.1.2
```

Release Information

Command introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Monitoring IGMP Snooping

Configuring IGMP Snooping on Switches

Verifying IGMP Snooping on EX Series Switches

show igmp-snooping route | **20181**

show igmp-snooping statistics | **20196**

show igmp statistics

IN THIS SECTION

- Syntax | 20203
- Syntax (ACX Series) | 20203
- Syntax (EX Series and QFX Series) | 20204
- Description | 20204
- Options | 20205
- Required Privilege Level | 20205
- Output Fields | 20205
- Sample Output | 20207
- show igmp statistics (User Plane output for ACX7100) | 20209
- Release Information | 20210

Syntax

```
show igmp statistics
<brief | detail>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (ACX Series)

```
show igmp statistics
<brief | detail>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series and QFX Series)

```
show igmp statistics
<brief | detail>
<interface interface-name>
```

Description

Display Internet Group Management Protocol (IGMP) statistics.

By default, Junos OS multicast devices collect statistics of received and transmitted IGMP control messages that reflect currently active multicast group subscribers.

MX series devices also automatically maintain *continuous* IGMP statistics globally on the device in addition to the default active subscriber statistics—these are persistent, continuous statistics of received and transmitted IGMP control packets that account for both past and current multicast group subscriptions processed on the device.

With continuous statistics, you can see the total count of IGMP control packets the device processed since the last device reboot or `clear igmp statistics continuous` command only for MX series devices. The device collects and displays continuous statistics only for the fields shown in the `IGMP packet statistics` output section of this command, and does not display the `IGMP Global statistics` section.

Devices that support continuous statistics maintain this information in a shared database and copy it to the backup Routing Engine at a configurable interval to avoid too much processing overhead on the Routing Engine. These actions preserve statistics counts across the following events or operations (which doesn't happen for the default active subscriber statistics):

- Routing daemon restart
- Graceful Routing Engine switchover (GRES)
- In-service software upgrade (ISSU)
- Line card reboot

You can change the default interval (300 seconds) using the `cont-stats-collection-interval` configuration statement at the `[edit routing-options multicast]` hierarchy level.

You can display either the default currently active subscriber statistics or continuous subscriber statistics (if supported), but not both at the same time. Include the `continuous` option to display continuous statistics, otherwise the command displays the statistics only for active subscribers.

Run the `clear igmp statistics` command to clear the currently active subscriber statistics. On devices that support continuous statistics, run the clear command with the `continuous` option to clear all continuous

statistics. You must run these commands separately to clear both types of statistics because the device maintains and clears the two types of statistics separately.

Options

none	Display IGMP statistics for all interfaces. These statistics represent currently active subscribers.
brief | detail	(Optional) Display the specified level of output.
continuous	(Optional) Display continuous IGMP statistics that account for both past and current multicast group subscribers instead of the default statistics that only reflect currently active subscribers. This option is not available with the `interface` option for interface-specific statistics.
interface ***interface-name***	(Optional) Display IGMP statistics about the specified interface only. This option is not available with the `continuous` option.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 606 on page 20205 describes the output fields for the `show igmp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 606: show igmp statistics Output Fields

Field Name	Field Description
`IGMP packet statistics`	Heading for IGMP packet statistics for all interfaces or for the specified interface name. **NOTE**: Shows currently active subscriber statistics in this section by default, or when the command includes the `continuous` option, shows continuous, persistent statistics that account for all IGMP control packets processed on the device.

Table 606: show igmp statistics Output Fields *(Continued)*

Field Name	Field Description
`IGMP Message type`	Summary of IGMP statistics: • `Membership Query`—Cumulative sum of membership queries sent and received. • `V1 Membership Query`—Number of V1 membership queries sent and received. • `V2 Membership Query`—Number of V2 membership queries sent and received. • `V3 Membership Query`—Number of V3 membership queries sent and received. • `V1 Membership Report`—Number of version 1 membership reports sent and received. • `DVMRP`—Number of DVMRP messages sent or received. • `PIM V1`—Number of PIM version 1 messages sent or received. • `Cisco Trace`—Number of Cisco trace messages sent or received. • `V2 Membership Report`—Number of version 2 membership reports sent or received. • `Group Leave`—Number of group leave messages sent or received. • `Mtrace Response`—Number of Mtrace response messages sent or received. • `Mtrace Request`—Number of Mtrace request messages sent or received. • `Domain Wide Report`—Number of domain-wide reports sent or received. • `V3 Membership Report`—Number of version 3 membership reports sent or received. • `Other Unknown types`—Number of unknown message types received. • `IGMP v3 unsupported type`—Number of messages received with unknown and unsupported IGMP version 3 message types. • `IGMP v3 source required for SSM`—Number of IGMP version 3 messages received that contained no source. • `IGMP v3 mode not applicable for SSM`—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM). Beginning with certain releases, this type includes records received for groups in the SSM range of addresses and in which the mode is MODE_IS_EXCLUDE or CHANGE_TO_EXCLUDE_MODE. This includes records with a non-empty source list.

Table 606: show igmp statistics Output Fields *(Continued)*

Field Name	Field Description
`Received`	Number of messages received.
`Sent`	Number of messages sent.
`Rx errors`	Number of received packets that contained errors.
`Max Rx rate (pps)`	Maximum number of IGMP packets received during 1 second interval.
`IGMP Global Statistics`	Summary of IGMP statistics for all interfaces. **NOTE**: These statistics are not supported or displayed with the `continuous` option. • `Bad Length`—Number of messages received with length errors so severe that further classification could not occur. • `Bad Checksum`—Number of messages received with a bad IP checksum. No further classification was performed. • `Bad Receive If`—Number of messages received on an interface not enabled for IGMP. • `Rx non-local`—Number of messages received from senders that are not local. • `Timed out`—Number of groups that timed out as a result of not receiving an explicit leave message. • `Rejected Report`—Number of reports dropped because of the IGMP group policy. • `Total Interfaces`—Number of interfaces configured to support IGMP.

Sample Output

show igmp statistics

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query           8883        459          0
```

```
V1 Membership Query        3333        159          0
V2 Membership Query        2500        100          0
V3 Membership Query        3050        200          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                  0
IGMP v3 unsupported type                             0
IGMP v3 source required for SSM                      0
IGMP v3 mode not applicable for SSM                  0

IGMP Global Statistics
Bad Length                    0
Bad Checksum                  0
Bad Receive If                0
Rx non-local               1227
Timed out                     0
Rejected Report               0
Total Interfaces              2
Max Rx rate (pps)          1536
```

show igmp statistics interface

```
user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type      Received       Sent  Rx errors
Membership Query              0        230          0
V1 Membership Report          0          0          0
```

show igmp statistics continuous

```
user@host> show igmp statistics continuous
IGMP packet statistics for all interfaces
```

```
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          9          0
V1 Membership Report          3          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          3          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          3          0          0
Other Unknown types                                 0
IGMP v3 unsupported type                            0
IGMP v3 source required for SSM                     0
IGMP v3 mode not applicable for SSM                 0
```

show igmp statistics (User Plane output for ACX7100)

Displays the following output in the user-plane for customer premise equipment (CPE).

```
user@up1> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          2          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          2          0          0
Other Unknown types                                 0
IGMP v3 unsupported type                            0
IGMP v3 source required for SSM                     0
IGMP v3 mode not applicable for SSM                 0

IGMP Global Statistics
```

```
Bad Length                    2
Bad Checksum                  0
Bad Receive If             4878
Rx non-local                  6
Timed out                     6
Rejected Report               0
Total Interfaces              2
Max Rx rate (pps)            58
```

Release Information

Command introduced before Junos OS Release 7.4.

`continuous` option added in Junos OS Release 19.4R1 for MX Series routers.

Command updated in Junos OS Release 24.4R1 to include multicast support monitoring for ACX7100 routers.

RELATED DOCUMENTATION

clear igmp statistics | 3627

show ike security-associations

IN THIS SECTION

- Syntax | 20211
- Description | 20211
- Options | 20211
- Required Privilege Level | 20211
- Output Fields | 20211
- Sample Output | 20215
- Release Information | 20216

Syntax

```
show ike security-associations
<brief | detail>
<peer-address>
```

Description

(Encryption interface on M Series and T Series routers only) Display information about Internet Key Exchange (IKE) security associations.

Options

none	Display standard information about all IKE security associations.
brief | detail	(Optional) Display the specified level of output.
peer-address	(Optional) Display IKE security associations for the specified peer address.

Required Privilege Level

view

Output Fields

Table 607 on page 20211 lists the output fields for the `show ike security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 607: show ike security-associations Output Fields

Field Name	Field Description	Level of Output
IKE peer	Remote end of the IKE negotiation.	**detail**
Role	Part played in the IKE session. The router triggering the IKE negotiation is the initiator, and the router accepting the first IKE exchange packets is the responder.	**detail**

Table 607: show ike security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Remote Address	Responder's address.	none specified
State	State of the IKE security association: • **Matured**—The IKE security association is established. • **Not matured**—The IKE security association is in the process of negotiation.	none specified
Initiator cookie	When the IKE negotiation is triggered, a random number is sent to the remote node.	All levels
Responder cookie	The remote node generates its own random number and sends it back to the initiator as a verification that the packets were received. Of the numerous security services available, protection against denial of service (DoS) is one of the most difficult to address. A “cookie” or anticlogging token (ACT) is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity. An exchange prior to CPU-intensive public key operations can thwart some DoS attempts (such as simple flooding with invalid IP source addresses).	All levels
Exchange type	Specifies the number of messages in an IKE exchange, and the payload types that are contained in each message. Each exchange type provides a particular set of security services, such as anonymity of the participants, perfect forward secrecy of the keying material, and authentication of the participants. Junos OS supports two types of exchanges: • **Main**—The exchange is done with six messages. **Main** encrypts the payload, protecting the identity of the neighbor. • **Aggressive**—The exchange is done with three messages. **Aggressive** does not encrypt the payload, leaving the identity of the neighbor unprotected.	All Levels

Table 607: show ike security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Authentication method	Type of authentication determines which payloads are exchanged and when they are exchanged. The Junos OS supports only **pre-shared keys**.	**detail**
Local	Prefix and port number of the local end.	**detail**
Remote	Prefix and port number of the remote end.	**detail**
Lifetime	Number of seconds remaining until the IKE security association expires.	**detail**
Algorithms	Header for the IKE algorithms output. • **Authentication**—Type of authentication algorithm used:**md5** or **sha1**. • **Encryption**—Type of encryption algorithm used: **des-cbc**, **3des-cbc**, or **None**. • **Pseudo random function**—Function that generates highly unpredictable random numbers:**hmac-md5** or**hmac-sha1**.	**detail**
Traffic statistics	Number of bytes and packets received and transmitted on the IKE security association. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the IKE security association. • **Input packets, Output packets**—Number of packets received and transmitted on the IKE security association.	**detail**

Table 607: show ike security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Notification to the key management process of the status of the IKE negotiation: • **caller notification sent**—Caller program notified about the completion of the IKE negotiation. • **waiting for done**—Negotiation is done. The library is waiting for the remote end retransmission timers to expire. • **waiting for remove**—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation. • **waiting for policy manager**—Negotiation is waiting for a response from the policy manager.	**detail**
IPsec security associates	Number of IPsec security associations created and deleted with this IKE security association.	**detail**

Table 607: show ike security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Phase 2 negotiations in progress	Number of phase 2 IKE negotiations in progress and status information:
	• **Negotiation type**—Type of phase 2 negotiation. The Junos OS currently supports **quick mode**.
	• **Message ID**—Unique identifier for a phase 2 negotiation.
	• **Local identity**—Identity of the local phase 2 negotiation. The format is *id-type-name* (*proto-name*:*port-number*,[0..*id-data-len*] = *iddata-presentation*)
	• **Remote identity**—Identity of the remote phase 2 negotiation. The format is *id-type-name* (*proto-name*:*port-number*,[0..*id-data-len*] = *iddata-presentation*)
	• **Flags**—Notification to the key management process of the status of the IKE negotiation:
	  • **caller notification sent**—Caller program notified about the completion of the IKE negotiation.
	  • **waiting for done**—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.
	  • **waiting for remove**—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.
	  • **waiting for policy manager**—Negotiation is waiting for a response from the policy manager.
		detail

Sample Output

show ike security-associations

```
user@host> show ike security-associations
Remote Address  State         Initiator cookie  Responder cookie  Exchange type
192.0.2.4         Matured       93870456fa000011  723a20713700003e  Main
```

show ike security-associations detail

```
user@host> show ike security-associations detail
IKE peer 192.0.2.4
  Role: Initiator, State: Matured
  Initiator cookie: cf22bd81a7000001, Responder cookie: fe83795c2800002e
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 192.0.2.5:500, Remote: 192.0.2.4:500
  Lifetime: Expires in 187 seconds
  Algorithms:
   Authentication        : md5
   Encryption            : 3des-cbc
   Pseudo random function: hmac-md5
  Traffic statistics:
   Input  bytes  :                 1000
   Output bytes  :                 1280
   Input  packets:                    5
   Output packets:                    9
  Flags: Caller notification sent
  IPsec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

Negotiation type: Quick mode, Role: Initiator, Message ID: 3582889153
    Local: 192.0.2.5:500, Remote: 192.0.2.4:500
    Local identity: ipv4_subnet(tcp:80,[0..7]=10.1.1.0/24)
    Remote identity: ipv4_subnet(tcp:100,[0..7]=10.1.2.0/24)
    Flags: Caller notification sent, Waiting for done
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear ike security-associations | 3632

show ilmi statistics

IN THIS SECTION

- Syntax | 20217
- Description | 20217
- Options | 20217
- Required Privilege Level | 20217
- Output Fields | 20217
- Sample Output | 20220
- Release Information | 20221

Syntax

```
show ilmi statistics
```

Description

Display input and output Integrated Local Management Interface (ILMI) statistics.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 608 on page 20218 lists the output fields for the `show ilmi statistics` command. Output fields are listed in the approximate order in which they appear.

Table 608: show ilmi statistics Output Fields

Field Name	Field Description
Input	Information about received ILMI packets: • **Packets**—Total number of messages delivered to the ILMI entity from the transport service. • **Bad versions**—Total number of messages delivered to the ILMI entity that were for an unsupported ILMI version. • **Bad community names**—Total number of messages delivered to the ILMI entity that did not use an ILMI community name. • **Bad community uses**—Total number of messages delivered to the ILMI entity that represented an ILMI operation that was not allowed by the ILMI community named in the message. • **ASN parse errors**—Total number of ASN.1 or BER errors encountered by the ILMI entity when decoding received ILMI messages. • **Too bigs**—Total number of ILMI packets delivered to the ILMI entity with an error status field of **tooBig**. • **No such names**—Total number of ILMI packets delivered to the ILMI entity with an error status field of **noSuchName**. • **Bad values**—Total number of ILMI packets delivered to the ILMI entity with an error status field of **badValue.** • **Read onlys**—Total number of valid ILMI packets delivered to the ILMI entity with an error status field of **readOnly**. Only incorrect implementations of ILMI generate this error. • **General errors**—Total number of ILMI packets delivered to the ILMI entity with an error status field of **genErr**. • **Total request varbinds**—Total number of objects retrieved successfully by the ILMI entity as a result of receiving valid ILMI **GetRequest** and **GetNext** packets. • **Total set varbinds**—Total number of objects modified successfully by the ILMI entity as a result of receiving valid ILMI **SetRequest** packets. • **Get requests**—Total number of ILMI **GetRequest** packets that have been accepted and processed by the ILMI entity. • **Get nexts**—Total number of ILMI **GetNext** packets that have been accepted and processed by the ILMI entity.

Table 608: show ilmi statistics Output Fields *(Continued)*

Field Name	Field Description
	• **Set requests**—Total number of ILMI **SetRequest** packets that have been accepted and processed by the ILMI entity. • **Get responses**—Total number of ILMI **GetResponse** packets that have been accepted and processed by the ILMI entity. • **Traps**—Total number of ILMI traps received by the ILMI entity. • **Silent drops**—Total number of **GetRequest**, **GetNextRequest**, **GetBulkRequest**, **SetRequest**, and **InformRequest** packets delivered to the ILMI entity that were silently dropped because the size of a reply containing an alternate response packet with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests. • **Proxy drops**—Total number of **GetRequest**, **GetNextRequest**, **GetBulkRequest**, **SetRequest**, and **InformRequest** packets delivered to the ILMI entity that were silently dropped because the transmission of the (possibly translated) message to a proxy target failed in such a way (other than a timeout) that no response packet could be returned.

Table 608: show ilmi statistics Output Fields *(Continued)*

Field Name	Field Description
Output	Information about transmitted ILMI packets: • **Packets**—Total number of messages passed from the ILMI entity to the transport service. • **Too bigs**—Total number of ILMI packets generated by the ILMI entity with an error status field of **tooBig**. • **No such names**—Total number of ILMI packets generated by the ILMI entity with an error status field of **noSuchName**. • **Bad values**—Total number of ILMI packets generated by the ILMI entity with an error status field of **badValue**. • **General errors**—Total number of ILMI packets generated by the ILMI entity with an error status field of **genErr**. • **Get requests**—Total number of ILMI **GetRequest** packets that have been generated by the ILMI entity. • **Get nexts**—Total number of ILMI **GetNext** packets that have been generated by the ILMI entity. • **Set requests**—Total number of ILMI **SetRequest** packets that have been generated by the ILMI entity. • **Get responses**—Total number of ILMI **GetResponse** packets that have been generated by the ILMI entity. • **Traps**—Total number of ILMI traps generated by the ILMI entity.

Sample Output

show ilmi statistics

```
user@host> show ilmi statistics
   ILMI statistics:
    Input:
      Packets: 0, Bad versions: 0, Bad community names: 0,
      Bad community uses: 0, ASN parse errors: 0,
      Too bigs: 0, No such names: 0, Bad values: 0,
```

```
      Read onlys: 0, General errors: 0,
      Total request varbinds: 0, Total set varbinds: 0,
      Get requests: 0, Get nexts: 0, Set requests: 0,
      Get responses: 0, Traps: 0,
      Silent drops: 0, Proxy drops 0
    Output:
      Packets: 0, Too bigs: 0, No such names: 0,
      Bad values: 0, General errors: 0,
      Get requests: 0, Get nexts: 0, Set requests: 0,
      Get responses: 0, Traps: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear ilmi statistics | 3634

show ingress-replication mvpn

IN THIS SECTION

- Syntax | 20221
- Description | 20222
- Required Privilege Level | 20222
- Output Fields | 20222
- Sample Output | 20223
- Release Information | 20223

Syntax

show ingress-replication mvpn

Description

Display the state and configuration of the ingress replication tunnels created for the MVPN application when using the **mpls-internet-multicast** routing instance type.

Required Privilege Level

View

Output Fields

Table 609 on page 20222 lists the output fields for the `show ingress-replication mvpn` command. Output fields are listed in the approximate order in which they appear.

Table 609: show ingress-replication mvpn Output Fields

Field Name	Field Description
Ingress tunnel	Identifies the MVPN ingress replication tunnel.
Application	Identifies the application (MVPN).
Unicast tunnels	List of unicast tunnels in use.
Leaf address	Address of the tunnel.
Tunnel type	Identifies the unicast tunnel type.
Mode	Indicates whether the tunnel was created as a new tunnel for the ingress replication, or if an existing tunnel was used.
State	Indicates whether the tunnel is Up or Down.

Sample Output

show ingress-replication mvpn

```
user@host> show ingress-replication mvpn
Ingress Tunnel: mvpn:1
  Application: MVPN
  Unicast tunnels
    Leaf Address       Tunnel-type       Mode       State
    10.255.245.2       P2P LSP           New        Up
    10.255.245.4       P2P LSP           New        Up
Ingress Tunnel: mvpn:2
  Application: MVPN
  Unicast tunnels
    Leaf Address       Tunnel-type       Mode       State
    10.255.245.2       P2P LSP           Existing   Up
```

Release Information

Command introduced in Junos OS Release 10.4.

show interfaces (NFX150)

IN THIS SECTION

- Syntax | 20224
- Description | 20224
- Required Privilege Level | 20224
- Sample Output | 20224
- Release Information | 20226

Syntax

```
show interfaces extensive interface-name
show interfaces interface-name
```

Description

Display details of the interfaces configured on NFX150.

Required Privilege Level

view

Sample Output

show interfaces extensive

```
user@host> show interfaces extensive heth-0-1
Physical interface: heth-0-1, Enabled, Physical link is Up
  Link-level type: Ethernet, Media type: Copper, MTU: 9192, Speed: 1Gbps, Duplex: Full-duplex,
Auto-negotiation: Enabled
  Device flags   : Present Running
  Current address: e8:b6:c2:cb:b8:89, Hardware address: e8:b6:c2:cb:b8:89
  Traffic statistics:
   Input  bytes  :                    0                 N/A bps
   Output bytes  :                    0                 N/A bps
   Input  packets:                    0                 N/A pps
   Output packets:                    0                 N/A pps
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0              N/A
    FIFO errors                              0                0
    DMA errors                               0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
```

```
    Oversized frames                        0
  VF statistics:
    VF Number: 0, PCI Address: 0000:06:10:0, Mapped to: ge-1/0/1
      Input bytes: 0, Input packets: 0, Output bytes: 0, Output packets: 0, Multicast packets: 0
    VF Number: 1, PCI Address: 0000:06:10:4, Mapped to: ge-1/0/2
      Input bytes: 0, Input packets: 0, Output bytes: 0, Output packets: 0, Multicast packets: 0
    VF Number: 2, PCI Address: 0000:06:11:0, Mapped to: None
      Input bytes: 0, Input packets: 0, Output bytes: 0, Output packets: 0, Multicast packets: 0
    VF Number: 3, PCI Address: 0000:06:11:4, Mapped to: None
      Input bytes: 0, Input packets: 0, Output bytes: 0, Output packets: 0, Multicast packets: 0
```

show interfaces <interface1>

```
user@host> show interfaces ge-1/0/1
Physical interface: ge-1/0/1, Enabled, Physical link is Up
  Interface index: 184, SNMP ifIndex: 551
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Link-mode: Full-duplex, Speed: 1000mbps,
BPDU Error: None, Loop Detect PDU Error: None,
  Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: e8:b6:c2:cb:b8:5c, Hardware address: e8:b6:c2:cb:b8:5c
  Last flapped   : 2019-09-12 15:43:49 UTC (00:00:03 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Ethernet FEC statistics              Errors
    FEC Corrected Errors                    0
    FEC Uncorrected Errors                  0
    FEC Corrected Errors Rate               0
    FEC Uncorrected Errors Rate             0
```

```
  PRBS Statistics : Disabled
  Interface transmit statistics: Disabled
```

show interfaces <interface1>

```
user@host> show interfaces ge-1/0/2
Physical interface: ge-1/0/2, Enabled, Physical link is Up
  Interface index: 185, SNMP ifIndex: 552
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Link-mode: Full-duplex, Speed: 1000mbps,
BPDU Error: None, Loop Detect PDU Error: None,
  Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: e8:b6:c2:cb:b8:5d, Hardware address: e8:b6:c2:cb:b8:5d
  Last flapped   : 2019-09-12 15:43:49 UTC (00:00:07 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Ethernet FEC statistics              Errors
    FEC Corrected Errors                    0
    FEC Uncorrected Errors                  0
    FEC Corrected Errors Rate               0
    FEC Uncorrected Errors Rate             0
  PRBS Statistics : Disabled
  Interface transmit statistics: Disabled
```

Release Information

Command introduced in Junos OS Release 18.1R1.

show interfaces

IN THIS SECTION

- Syntax (Gigabit Ethernet) | 20227
- Syntax (10 Gigabit Ethernet) | 20227
- Syntax (ACX5448, ACX5448-D, ACX710) | 20228
- Syntax (QFX5130-32CD) | 20228
- Syntax (SRX Series Firewalls and (vSRX Virtual Firewall and vSRX Virtual Firewall 3.0 platforms)) | 20228
- Description | 20229
- Options | 20230
- Additional Information | 20231
- Required Privilege Level | 20231
- Output Fields | 20232
- Sample output for G.fast and Annex J support | 20294
- Sample Output Gigabit Ethernet | 20311
- Sample Output | 20327
- Release Information | 20369

Syntax (Gigabit Ethernet)

```
show interfaces ge-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (10 Gigabit Ethernet)

```
show interfaces xe-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
```

```
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (ACX5448, ACX5448-D, ACX710)

```
show interfaces et-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (QFX5130-32CD)

```
show interfaces et-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (SRX Series Firewalls and (vSRX Virtual Firewall and vSRX Virtual Firewall 3.0 platforms))

```
show interfaces (
    <interface-name>
    <brief | detail | extensive | terse>
    <controller interface-name>|
    <descriptions interface-name>|
    <destination-class (all / destination-class-name logical-interface-name)>|
    <diagnostics optics interface-name>|
    <far-end-interval interface-fpc/pic/port>|
    <filters interface-name>|
    <flow-statistics interface-name>|
    <interval interface-name>|
    <load-balancing  (detail | interface-name)>|
    <mac-database mac-address mac-address>|
```

```
    <mc-ae id identifier unit number revertive-info>|
    <media interface-name>|
    <policers interface-name>|
    <queue both-ingress-egress egress forwarding-class forwarding-class ingress l2-statistics>|
    <redundancy  (detail | interface-name)>|
    <routing brief detail summary interface-name>|
    <routing-instance (all | instance-name)>|
    <snmp-index snmp-index>|
    <source-class (all / destination-class-name logical-interface-name)>|
    <statistics interface-name>|
    <switch-port switch-port number>|
    <transport pm (all | optics | otn) (all | current | currentday | interval | previousday)
(all | interface-name)>|
    <zone interface-name>|
    <dsl-sfp-options (adsl-options | gfast-options | vdsl-options) >
)
```

Description

Display status information about the specified Gigabit Ethernet interface.

(M320, M120, MX Series, and T Series routers only) Display status information about the specified 10-Gigabit Ethernet interface.

Display the IPv6 interface traffic statistics about the specified Gigabit Ethernet interface for MX series routers. The input and output bytes (bps) and packets (pps) rates are not displayed for IFD and local traffic.

Display status information and statistics about interfaces on SRX Series, vSRX Virtual Firewall, and vSRX Virtual Firewall 3.0 platforms running Junos OS.

SRX4600 supports 40-Gigabit Ethernet breakouts only in PIC mode. Use the `show interfaces extensive` command to view the speed configured for the interface on SRX4600. Reboot the device for the changed configuration to take effect.

On SRX Series appliances, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

Starting in Junos OS Release 18.4R1, Output fields `Next-hop` and `vpls-status` is displayed in the `show interfaces` *`interface name`* `detail` command, only for Layer 2 protocols on MX480 routers.

In Junos OS Releases 19.2R3, 19.3R3, 19.4R3, 20.1R2, and 20.2R1, on QFX5120-48Y switch, the `show interfaces` *`interface-name`*`<media><extensive>` command displays the autonegotiation status only for the interface that supports autonegotiation. This is applicable when the switch operates at 1-Gbps speed. In

the earlier Junos OS releases, incorrect autonegotiation status was displayed even when autonegotiation was disabled.

QFX5130-32CD switches does not display the `Filters statistics` when the `show interfaces extensive` command is executed due to interface-level filter statistics related hardware limitations. See "show interfaces extensive (QFX5130-32CD)" on page 20358.

Starting in Junos OS Release 20.4R1, we support G.fast and Annex J specification with SFP xDSL for ADSL2/ADSL2+ and all VDSL2 profiles on SRX380, SRX300, SRX320, SRX340, and SRX345 devices.

Options

For Gigabit interfaces:

ge-*fpc*/*pic*/*port*	Display standard information about the specified Gigabit Ethernet interface.

> **NOTE**: Interfaces with different speeds are named uniformly with ge-0/0/x for backward compatibility. Use the `show interfaces` command to view the interface speeds.

brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

For 10 Gigabit interfaces:

xe-*fpc*/*pic*/*port*	Display standard information about the specified 10-Gigabit Ethernet interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index* (Optional) Display information for the specified SNMP index of the interface.

statistics (Optional) Display static interface statistics.

For SRX interfaces:

interface-name (Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.

- `at-pim/0/port`—ATM-over-ADSL or ATM-over-SHDSL interface.
- `ce1-pim/0/port`—Channelized E1 interface.
- `cl-0/0/8`—3G wireless modem interface for SRX320 devices.
- `ct1-pim/0/port`—Channelized T1 interface.
- `dl0`—Dialer Interface for initiating ISDN and USB modem connections.
- `e1-pim/0/port`—E1 interface.
- `e3-pim/0/port`—E3 interface.
- `fe-pim/0/port`—Fast Ethernet interface.
- `ge-pim/0/port`—Gigabit Ethernet interface.
- `se-pim/0/port`—Serial interface.
- `t1-pim/0/port`—T1 (also called DS1) interface.
- `t3-pim/0/port`—T3 (also called DS3) interface.
- `wx-slot/0/0`—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

Additional Information

In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.

Required Privilege Level

view

Output Fields

Table 610 on page 20232 describes the output fields for the `show interfaces` (Gigabit Ethernet) command. Output fields are listed in the approximate order in which they appear. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see Table 611 on page 20280.

Table 610: show interfaces (Gigabit Ethernet) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical interface`	Name of the physical interface.	All levels
`Enabled`	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
`Interface index`	Index number of the physical interface, which reflects its initialization sequence.	`detail` `extensive` none
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail` `extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`
`Link-level type`	Encapsulation being used on the physical interface.	All levels
`MTU`	Maximum transmission unit size on the physical interface.	All levels
`Speed`	Speed at which the interface is running.	All levels
`Loopback`	Loopback status: `Enabled` or `Disabled`. If loopback is enabled, type of loopback: `Local` or `Remote`.	All levels
`Source filtering`	Source filtering status: `Enabled` or `Disabled`.	All levels

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`LAN-PHY mode`	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
`WAN-PHY mode`	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
`Unidirectional`	Unidirectional link mode status for 10-Gigabit Ethernet interface: `Enabled` or `Disabled` for parent interface; `Rx-only` or `Tx-only` for child interfaces.	All levels
`Flow control`	Flow control status: `Enabled` or `Disabled`.	All levels
`Auto-negotiation`	(Gigabit Ethernet interfaces) Autonegotiation status: `Enabled` or `Disabled`.	All levels
`Remote-fault`	(Gigabit Ethernet interfaces) Remote fault status: • `Online`—Autonegotiation is manually configured as online. • `Offline`—Autonegotiation is manually configured as offline.	All levels
`Device flags`	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
`Interface flags`	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
`Link flags`	Information about the link. Possible values are described in the “Links Flags” section under *Common Output Fields Description*.	All levels
`Wavelength`	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Frequency`	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
`CoS queues`	Number of CoS queues configured.	`detail extensive` none
`Schedulers`	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	`extensive`
`Hold-times`	Current interface hold-time up and hold-time down, in milliseconds (ms).	`detail extensive`
`Current address`	Configured MAC address.	`detail extensive` none
`Hardware address`	Hardware MAC address.	`detail extensive` none
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second`*`:`*`timezone`* (*`hour`*`:`*`minute`*`:`*`second`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive` none
`Input Rate`	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
`Output Rate`	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`
`Egress account overhead`	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	`detail extensive`
`Ingress account overhead`	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	`detail extensive`
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • `Output bytes`—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the "show interfaces " on page 20227 command.	`detail extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input errors`	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Errors`—Sum of the incoming frame terminated and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • `L3 incompletes`—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` statement. • `L2 channel errors`—Number of times the software did not find a valid logical interface for an incoming frame. • `L2 mismatch timeouts`—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • `FIFO errors`—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `Resource errors`—Sum of transmit drops.	`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • `Errors`—Sum of the outgoing frame terminated and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. **NOTE**: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the `Drops` field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs. • `Collisions`—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number must always be 0. If it is nonzero, there is a software bug. • `Aged packets`—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field must never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • `FIFO errors`—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `HS link CRC errors`—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • `MTU errors`—Number of packets whose size exceeded the MTU of the interface.	`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Resource errors`—Sum of transmit drops.	
`Egress queues`	Total number of egress queues supported on the specified interface. **NOTE**: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the `show interfaces` command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	`detail extensive`
`Queue counters (Egress)`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism. **NOTE**: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the `Dropped packets` field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.	`detail extensive`
`Ingress queues`	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Queue counters (Ingress)`	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	`extensive`
`Active alarms` and `Active defects`	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value `None` or `Link`. • `None`—There are no active defects or alarms. • `Link`—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	`detail` `extensive` none
Interface transmit statistics	(On MX Series devices) Status of the `interface-transmit-statistics` configuration: Enabled or Disabled. • `Enabled`—When the `interface-transmit-statistics` statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • `Disabled`—When the `interface-transmit-statistics` statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface.	`detail` `extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`OTN FEC statistics`	The forward error correction (FEC) counters provide the following statistics: • `Corrected Errors`—Count of corrected errors in the last second. • `Corrected Error Ratio`—Corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	`detail extensive`
`PCS statistics`	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • `Bit errors`—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • `Errored blocks`—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	`detail extensive`
`PRBS Statistics`	Displays the Pseudo Random Binary Sequence (PRBS) statistics. The PRBS Statistics are displayed in extensive, detailed, media and normal output except terse. The output is displayed per serdes lane. The output consists of the total number of iterations in error, the total number of iterations and the number of monitored seconds. An error iteration is one in which at least one bit error is seen.	`detail extensive`
`PRBS Pattern`	Specifies the pattern type, that is in the range from 7 to 58.	`detail extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Link Degrade`	Shows the link degrade status of the physical link and the estimated bit error rates (BERs). This field is available only for the PICs supporting the physical link monitoring feature.
	• `Link Monitoring`—Indicates if physical link degrade monitoring is enabled on the interface.
	  • `Enable`—Indicates that link degrade monitoring has been enabled (using the `link-degrade-monitor` statement) on the interface.
	  • `Disable`—Indicates that link degrade monitoring has not been enabled on the interface. If link degrade monitoring has not been enabled, the output does not show any related information, such as BER values and thresholds.
	• `Link Degrade Set Threshold`—The BER threshold value at which the link is considered degraded and a corrective action is triggered.
	• `Link Degrade Clear Threshold`—The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted.
	• `Estimated BER`—The estimated bit error rate.
	• `Link-degrade event`—Shows link degrade event information.
	  • `Seconds`—Time (in seconds) elapsed after a link degrade event occurred.
	  • `Count`—The number of link degrade events recorded.
	  • `State`—Shows the link degrade status (example: `Defect Active`).
		`detail extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`MAC statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC subsystem, including the following:
	• `Total octets` and `total packets`—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the "show interfaces " on page 20227 command.
	• `Unicast packets`, `Broadcast packets`, and `Multicast packets`—Number of unicast, broadcast, and multicast packets.
	• `CRC/Align errors`—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
	• `FIFO error`—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
	• `MAC control frames`—Number of MAC control frames.
	• `MAC pause frames`—Number of MAC control frames with `pause` operational code.
	• `Oversized frames`—There are two possible conditions regarding the number of oversized frames:
	  • Packet length exceeds interface MTU, or
	  • Packet length exceeds MRU
	• `Jabber frames`—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
		`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Fragment frames`—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • `VLAN tagged frames`—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. **NOTE**: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the `VLAN tagged frames` field displays 0 when the `show interfaces` command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC. • `Code violations`—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	
`OTN Received Overhead Bytes`	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	`extensive`
`OTN Transmitted Overhead Bytes`	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Filter statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet may enter the system or be rejected. • `Input packet count`—Number of packets received from the MAC hardware that the filter processed. • `Input packet rejects`—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • `Input DA rejects`—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). **NOTE**: On PTX Series routers, ignore any values displayed for Input DA rejects. Input DA rejects is not supported on PTX Series routers. • `Input SA rejects`—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field must increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • `Output packet count`—Number of packets that the filter has given to the MAC hardware. • `Output packet pad count`—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured.	`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Output packet error count`—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field must not increment. • `CAM destination filters, CAM source filters`—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields must be 0.	
`PMA PHY`	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `PHY Lock`—Phase-locked loop • `PHY Light`—Loss of optical signal	`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`WIS section`	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `BIP-B1`—Bit interleaved parity for SONET section overhead • `SEF`—Severely errored framing • `LOL`—Loss of light • `LOF`—Loss of frame • `ES-S`—Errored seconds (section) • `SES-S`—Severely errored seconds (section) • `SEFS-S`—Severely errored framing seconds (section)	`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`WIS line`	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information: • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `BIP-B2`—Bit interleaved parity for SONET line overhead • `REI-L`—Remote error indication (near-end line) • `RDI-L`—Remote defect indication (near-end line) • `AIS-L`—Alarm indication signal (near-end line) • `BERR-SF`—Bit error rate fault (signal failure) • `BERR-SD`—Bit error rate defect (signal degradation) • `ES-L`—Errored seconds (near-end line) • `SES-L`—Severely errored seconds (near-end line) • `UAS-L`—Unavailable seconds (near-end line) • `ES-LFE`—Errored seconds (far-end line) • `SES-LFE`—Severely errored seconds (far-end line) • `UAS-LFE`—Unavailable seconds (far-end line)	`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`WIS path`	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information: • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `BIP-B3`—Bit interleaved parity for SONET section overhead • `REI-P`—Remote error indication • `LOP-P`—Loss of pointer (path) • `AIS-P`—Path alarm indication signal • `RDI-P`—Path remote defect indication • `UNEQ-P`—Path unequipped • `PLM-P`—Path payload (signal) label mismatch • `ES-P`—Errored seconds (near-end STS path) • `SES-P`—Severely errored seconds (near-end STS path) • `UAS-P`—Unavailable seconds (near-end STS path) • `SES-PFE`—Severely errored seconds (far-end STS path) • `UAS-PFE`—Unavailable seconds (far-end STS path)	`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Autonegotiation information`	Information about link autonegotiation.
	• `Negotiation status:`
	  • `Incomplete`—Ethernet interface has the speed or link mode configured.
	  • `No autonegotiation`—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • `Complete`—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• `Link partner status`—`OK` when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• `Link partner`—Information from the remote Ethernet device:
	  • `Link mode`—Depending on the capability of the link partner, either `Full-duplex` or `Half-duplex`.
	  • `Flow control`—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are `Symmetric` (link partner supports `PAUSE` on receive and transmit), `Asymmetric` (link partner supports `PAUSE` on transmit), `Symmetric/Asymmetric` (link partner supports `PAUSE` on receive and transmit or only `PAUSE` on transmit), and `None` (link partner does not support flow control).
	  • `Remote fault`—Remote fault information from the link partner—`Failure` indicates a receive link error. `OK` indicates that the link partner is receiving. `Negotiation error` indicates a negotiation error. `Offline` indicates that the link partner is going offline.
	• `Local resolution`—Information from the local Ethernet device:
	  • `Flow control`—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are `Symmetric/Asymmetric` (local device supports `PAUSE` on receive
		`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	and transmit or only `PAUSE` on receive) and `None` (local device does not support flow control). Depending on the result of the negotiation with the link partner, local resolution flow control type will display `Symmetric` (local device supports `PAUSE` on receive and transmit), `Asymmetric` (local device supports `PAUSE` on receive), and `None` (local device does not support flow control). • `Remote fault`—Remote fault information. `Link OK` (no error detected on receive), `Offline` (local interface is offline), and `Link Failure` (link error detected on receive).	
`Received path trace`, `Transmitted path trace`	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.	`extensive`
`Packet Forwarding Engine configuration`	Information about the configuration of the Packet Forwarding Engine: • `Destination slot`—FPC slot number.	`extensive`

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`CoS information`	Information about the CoS queue for the physical interface. • `CoS transmit queue`—Queue number and its associated user-configured forwarding class name. • `Bandwidth %`—Percentage of bandwidth allocated to the queue. • `Bandwidth bps`—Bandwidth allocated to the queue (in bps). • `Buffer %`—Percentage of buffer space allocated to the queue. • `Buffer usec`—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • `Priority`—Queue priority: `low` or `high`. • `Limit`—Displayed if rate limiting is configured for the queue. Possible values are `none` and `exact`. If `exact` is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If `none` is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	`extensive`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Index number of the logical interface, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP interface index number for the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Flags`	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	All levels

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed.	brief detail extensive none
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`ACI VLAN`	Information displayed for agent circuit identifier (ACI) interface set configured with the `agent-circuit-id` autoconfiguration stanza. `Dynamic Profile`—Name of the dynamic profile that defines the ACI interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information. **NOTE**: The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the `line-identity` autoconfiguration stanza.	`brief detail extensive` none
Line Identity	Information displayed for access-line-identifier (ALI) interface sets configured with the `line-identity` autoconfiguration stanza. • `Dynamic Profile`—Name of the dynamic profile that defines the ALI interface set. • Trusted option used to create the ALI interface set: `Circuit-id`, `Remote-id`, or `Accept-no-ids`. More than one option can be configured. If configured, the ALI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ALI information. **NOTE**: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the `agent-circuit-id` autoconfiguration stanza.	`detail`
`Protocol`	Protocol family. Possible values are described in the “Protocol Field” section under *Common Output Fields Description*.	`detail extensive` none
`MTU`	Maximum transmission unit size on the logical interface.	`detail extensive` none

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Neighbor Discovery Protocol (NDP)Queue Statistics`	NDP statistics for protocol `inet6` under logical interface statistics. • `Max nh cache`—Maximum interface neighbor discovery nexthop cache size. • `New hold nh limit`—Maximum number of new unresolved nexthops. • `Curr nh cnt`—Current number of resolved nexthops in the NDP queue. • `Curr new hold cnt`—Current number of unresolved nexthops in the NDP queue. • `NH drop cnt`—Number of NDP requests not serviced.	All levels
`Dynamic Profile`	Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	`detail extensive` none
`Service Name Table`	Name of the service name table for the interface configured with a PPPoE family.	`detail extensive` none
`Max Sessions`	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	`detail extensive` none
`Duplicate Protection`	State of PPPoE duplicate protection: `On` or `Off`. When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	`detail extensive` none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: `On` or `Off`. When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	`detail extensive` none
`AC Name`	Name of the access concentrator.	`detail extensive` none

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Maximum labels`	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	`detail extensive` none
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the specified interface set. • `Input bytes, Output bytes`—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • `Input packets, Output packets`—Number of packets received and transmitted on the interface set.	`detail extensive`
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	`extensive`
`Local statistics`	Number and rate of bytes and packets destined to the router.	`extensive`
`Transit statistics`	Number and rate of bytes and packets transiting the switch. **NOTE**: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the `Output bytes` and `Output packets` interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	`extensive`
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Route Table`	Route table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail extensive` none

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Information about protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	`detail extensive`
`Donor interface`	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	`detail extensive` none
`Preferred source address`	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	`detail extensive` none
`Input Filters`	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	`detail extensive`
`Output Filters`	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	`detail extensive`
`Mac-Validate Failures`	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	`detail extensive` none
`Addresses, Flags`	Information about the address flags. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	`detail extensive` none
`protocol-family`	Protocol family configured on the logical interface. If the protocol is `inet`, the IP address of the interface is also displayed.	`brief`
`Flags`	Information about the address flag. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	`detail extensive` none
`Destination`	IP address of the remote side of the connection.	`detail extensive` none

Table 610: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Local`	IP address of the logical interface.	`detail` `extensive` none
`Broadcast`	Broadcast address of the logical interface.	`detail` `extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`

The following table describes the output fields for the `show interfaces` (10–Gigabit Ethernet) command.

Field Name	Field Description	Level of Output
`Physical interface`	Name of the physical interface.	All levels
`Enabled`	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
`Interface index`	Index number of the physical interface, which reflects its initialization sequence.	`detail` `extensive` none
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail` `extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`
`Link-level type`	Encapsulation being used on the physical interface.	All levels
`MTU`	Maximum transmission unit size on the physical interface.	All levels

`Speed`	Speed at which the interface is running.	All levels
`Loopback`	Loopback status: `Enabled` or `Disabled`. If loopback is enabled, type of loopback: `Local` or `Remote`.	All levels
`Source filtering`	Source filtering status: `Enabled` or `Disabled`.	All levels
`LAN-PHY mode`	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
`WAN-PHY mode`	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
`Unidirectional`	Unidirectional link mode status for 10-Gigabit Ethernet interface: `Enabled` or `Disabled` for parent interface; `Rx-only` or `Tx-only` for child interfaces.	All levels
`Flow control`	Flow control status: `Enabled` or `Disabled`.	All levels
`Auto-negotiation`	(Gigabit Ethernet interfaces) Autonegotiation status: `Enabled` or `Disabled`.	All levels
`Remote-fault`	(Gigabit Ethernet interfaces) Remote fault status: • `Online`—Autonegotiation is manually configured as online. • `Offline`—Autonegotiation is manually configured as offline.	All levels
`Device flags`	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
`Interface flags`	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
`Link flags`	Information about the link. Possible values are described in the “Links Flags” section under *Common Output Fields Description*.	All levels

Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: *year-month-day hour*:*minute*:*second*:*timezone* (*hour*:*minute*:*second* ago). For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified

`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail` `extensive`
`Egress account overhead`	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	`detail` `extensive`
`Ingress account overhead`	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	`detail` `extensive`
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • `Output bytes`—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail` `extensive`

`Input errors`	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Errors`—Sum of the incoming frame terminated and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • `L3 incompletes`—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` statement. • `L2 channel errors`—Number of times the software did not find a valid logical interface for an incoming frame. • `L2 mismatch timeouts`—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • `FIFO errors`—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `Resource errors`—Sum of transmit drops.	`extensive`

`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • `Errors`—Sum of the outgoing frame terminated and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Collisions`—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • `Aged packets`—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • `FIFO errors`—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `HS link CRC errors`—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • `MTU errors`—Number of packets whose size exceeded the MTU of the interface. • `Resource errors`—Sum of transmit drops.	`extensive`

`Egress queues`	Total number of egress queues supported on the specified interface. **NOTE**: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the `show interfaces` command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenorio, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	`detail` `extensive`
`Queue counters (Egress)`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	`detail` `extensive`
`Ingress queues`	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	`extensive`
`Queue counters (Ingress)`	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	`extensive`

`Active alarms and Active defects`	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value `None` or `Link`.
	• `None`—There are no active defects or alarms.
	• `Link`—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.
		`detail`
		`extensive`
		none

`OTN alarms`	Active OTN alarms identified on the interface.	`detail` `extensive`
`OTN defects`	OTN defects received on the interface.	`detail` `extensive`
`OTN FEC Mode`	The FECmode configured on the interface. • `efec`—Enhanced forward error correction (EFEC) is configured to defect and correct bit errors. • `gfec`—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors. • `none`—FEC mode is not configured.	`detail` `extensive`
`OTN Rate`	OTN mode. • `fixed-stuff-bytes`—Fixed stuff bytes 11.0957 Gbps. • `no-fixed-stuff-bytes`—No fixed stuff bytes 11.0491 Gbps. • `pass-through`—Enable OTN passthrough mode. • `no-pass-through`—Do not enable OTN passthrough mode.	`detail` `extensive`
`OTN Line Loopback`	Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: `enabled` or `disabled`.	`detail` `extensive`

OTN FEC statistics	The forward error correction (FEC) counters for the DWDM OTN PIC. • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	detail extensive
OTN FEC alarms	OTN FEC excessive or degraded error alarms triggered on the interface. • FEC Degrade—OTU FEC Degrade defect. • FEC Excessive—OTU FEC Excessive Error defect.	detail extensive
OTN OC	OTN OC defects triggered on the interface. • LOS—OC Loss of Signal defect. • LOF—OC Loss of Frame defect. • LOM—OC Loss of Multiframe defect. • Wavelength Lock—OC Wavelength Lock defect.	detail extensive

OTN OTU	OTN OTU defects detected on the interface • AIS—OTN AIS alarm. • BDI—OTN OTU BDI alarm. • IAE—OTN OTU IAE alarm. • TTIM—OTN OTU TTIM alarm. • SF—OTN ODU bit error rate fault alarm. • SD—OTN ODU bit error rate defect alarm. • TCA-ES—OTN ODU ES threshold alarm. • TCA-SES—OTN ODU SES threshold alarm. • TCA-UAS—OTN ODU UAS threshold alarm. • TCA-BBE—OTN ODU BBE threshold alarm. • BIP—OTN ODU BIP threshold alarm. • BBE—OTN OTU BBE threshold alarm. • ES—OTN OTU ES threshold alarm. • SES—OTN OTU SES threshold alarm. • UAS—OTN OTU UAS threshold alarm.	detail extensive
Received DAPI	Destination Access Port Interface (DAPI) from which the packets were received.	detail extensive
Received SAPI	Source Access Port Interface (SAPI) from which the packets were received.	detail extensive
Transmitted DAPI	Destination Access Port Interface (DAPI) to which the packets were transmitted.	detail extensive
Transmitted SAPI	Source Access Port Interface (SAPI) to which the packets were transmitted.	detail extensive

`PCS statistics`	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • `Bit errors`—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • `Errored blocks`—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	`detail` `extensive`

MAC statistics	Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following: • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	extensive

OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

`Filter statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • `Input packet count`—Number of packets received from the MAC hardware that the filter processed. • `Input packet rejects`—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • `Input DA rejects`—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • `Input SA rejects`—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • `Output packet count`—Number of packets that the filter has given to the MAC hardware. • `Output packet pad count`—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • `Output packet error count`—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment.	`extensive`

	• CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	
PMA PHY	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem.	extensive
WIS section	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. Subfields are: • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section)	extensive

WIS line	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. Subfields are: • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line)	extensive

WIS path	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. Subfields are: • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path)	extensive

`Autonegotiation information`	Information about link autonegotiation.
	• `Negotiation status:`
	  • `Incomplete`—Ethernet interface has the speed or link mode configured.
	  • `No autonegotiation`—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • `Complete`—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• `Link partner status`—`OK` when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• `Link partner:`
	  • `Link mode`—Depending on the capability of the attached Ethernet device, either `Full-duplex` or `Half-duplex`.
	  • `Flow control`—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is `None`. For Gigabit Ethernet interfaces, types are `Symmetric` (link partner supports `PAUSE` on receive and transmit), `Asymmetric` (link partner supports `PAUSE` on transmit), and `Symmetric/Asymmetric` (link partner supports both `PAUSE` on receive and transmit or only `PAUSE` receive).
	  • `Remote fault`—Remote fault information from the link partner—`Failure` indicates a receive link error. `OK` indicates that the link partner is receiving. `Negotiation error` indicates a negotiation error. `Offline` indicates that the link partner is going offline.
	• `Local resolution`—Information from the link partner:
	  • `Flow control`—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are `Symmetric` (link partner supports `PAUSE` on receive and transmit), `Asymmetric` (link partner supports `PAUSE` on transmit), and `Symmetric/Asymmetric` (link partner supports both `PAUSE` on receive and transmit or only `PAUSE` receive).
	  • `Remote fault`—Remote fault information. `Link OK` (no error detected on receive), `Offline` (local interface is offline), and `Link Failure` (link error detected on receive).
		`extensive`

`Received path trace`, `Transmitted path trace`	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.	`extensive`
`Packet Forwarding Engine configuration`	Information about the configuration of the Packet Forwarding Engine: • `Destination slot`—FPC slot number.	`extensive`
`CoS information`	Information about the CoS queue for the physical interface. • `CoS transmit queue`—Queue number and its associated user-configured forwarding class name. • `Bandwidth %`—Percentage of bandwidth allocated to the queue. • `Bandwidth bps`—Bandwidth allocated to the queue (in bps). • `Buffer %`—Percentage of buffer space allocated to the queue. • `Buffer usec`—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • `Priority`—Queue priority: `low` or `high`. • `Limit`—Displayed if rate limiting is configured for the queue. Possible values are `none` and `exact`. If `exact` is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If `none` is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	`extensive`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels

Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed.	brief detail extensive none

`Demux:`	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:
	• Source Family Inet
	• Destination Family Inet
		`detail`
		`extensive`
		none

`Encapsulation`	Encapsulation on the logical interface.	All levels

`Protocol`	Protocol family. Possible values are described in the “Protocol Field” section under *Common Output Fields Description*.	`detail`
		`extensive`
		none

`MTU`	Maximum transmission unit size on the logical interface.	`detail`
		`extensive`
		none

`Maximum labels`	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	`detail`
		`extensive`
		none

`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the specified interface set. • `Input bytes, Output bytes`—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • `Input packets, Output packets`—Number of packets received and transmitted on the interface set.	`detail` `extensive`
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	`extensive`
`Local statistics`	Number and rate of bytes and packets destined to the routing device.	`extensive`

`Transit statistics`	Number and rate of bytes and packets transiting the switch. **NOTE**: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the `Output bytes` and `Output packets` interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	`extensive`
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`

`Route Table`	Route table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail`
		`extensive`
		none

`Flags`	Information about protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	`detail` `extensive`

`Donor interface`	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	`detail`
		`extensive`
		none

`Preferred source address`	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	`detail`
		`extensive`
		none

`Input Filters`	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	`detail` `extensive`
`Output Filters`	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	`detail` `extensive`

`Mac-Validate Failures`	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	`detail` `extensive` none
`Addresses, Flags`	Information about the address flags. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	`detail` `extensive` none
`protocol-family`	Protocol family configured on the logical interface. If the protocol is `inet`, the IP address of the interface is also displayed.	`brief`
`Flags`	Information about address flag (possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	`detail` `extensive` none
`Destination`	IP address of the remote side of the connection.	`detail` `extensive` none
`Local`	IP address of the logical interface.	`detail` `extensive` none
`Broadcast`	Broadcast address of the logical interlace.	`detail` `extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. The following table describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500 bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). The `ge-0/3/0` interface is the inbound physical interface, and the `ge-0/0/0` interface is the outbound physical interface. On both interfaces, traffic is carried on logical unit `.50` (VLAN 50).

Table 611: Gigabit and 10 Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	`show interfaces ge-0/3/0 extensive`	Traffic statistics: Input bytes: 496 bytes per packet, representing the Layer 2 packet MAC statistics: Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	`show interfaces ge-0/3/0.50 extensive`	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	`show interfaces ge-0/0/0 extensive`	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes include 6 bytes for the destination MAC address plus 4 bytes for VLAN plus 2 bytes for the Ethernet type.
Outbound logical interface	`show interfaces ge-0/0/0.50 extensive`	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Table 612 on page 20281 lists the output fields for the `show interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 612: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical interface`	Name of the physical interface.	All levels
`Enabled`	State of the interface.	All levels
`Interface index`	Index number of the physical interface, which reflects its initialization sequence.	`detail` `extensive` `none`
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail` `extensive` `none`
`Link-level type`	Encapsulation being used on the physical interface.	All levels
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`
`MTU`	Maximum transmission unit size on the physical interface.	All levels
`Link mode`	Link mode: Full-duplex or Half-duplex.	
`Speed`	Speed at which the interface is running.	All levels
`BPDU error`	Bridge protocol data unit (BPDU) error: Detected or None	
`Loopback`	Loopback status: `Enabled` or `Disabled`. If loopback is enabled, type of loopback: `Local` or `Remote`.	All levels
`Source filtering`	Source filtering status: `Enabled` or `Disabled`.	All levels
`Flow control`	Flow control status: `Enabled` or `Disabled`.	All levels

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Auto-negotiation`	(Gigabit Ethernet interfaces) Autonegotiation status: `Enabled` or `Disabled`.	All levels
`Remote-fault`	(Gigabit Ethernet interfaces) Remote fault status: • `Online`—Autonegotiation is manually configured as online. • `Offline`—Autonegotiation is manually configured as offline.	All levels
`Device flags`	Information about the physical device.	All levels
`Interface flags`	Information about the interface.	All levels
`Link flags`	Information about the physical link.	All levels
`CoS queues`	Number of CoS queues configured.	`detail extensive` none
`Current address`	Configured MAC address.	`detail extensive` none
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second`*`:`*`timezone`* `(`*`hour`*`:`*`minute`*`:`*`second`* `ago)`. For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive` none
`Input Rate`	Input rate in bits per second (bps) and packets per second (pps).	None
`Output Rate`	Output rate in bps and pps.	None

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Active alarms and Active defects`	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value `None` or `Link`. • `None`—There are no active defects or alarms. • `Link`—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	`detail extensive` none
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input errors`	Input errors on the interface. • `Errors`—Sum of the incoming frame terminated and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • `L3 incompletes`—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` . • `L2 channel errors`—Number of times the software did not find a valid logical interface for an incoming frame. • `L2 mismatch timeouts`—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • `FIFO errors`—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `Resource errors`—Sum of transmit drops.	`extensive`

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • `Errors`—Sum of the outgoing frame terminated and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Collisions`—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation; therefore, for Gigabit Ethernet PICs, this number must always remain 0. If it is nonzero, there is a software bug. • `Aged packets`—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field must never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • `FIFO errors`—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `HS link CRC errors`—Number of errors on the high-speed links between the ASICs responsible for handling the interfaces. • `MTU errors`—Number of packets whose size exceeded the MTU of the interface. • `Resource errors`—Sum of transmit drops.	`extensive`
`Ingress queues`	Total number of ingress queues supported on the specified interface.	`extensive`

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Queue counters and queue number`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	`detail extensive`
`ECN CE marked packets`	Number of packets marked with CE (congestion experienced) due to congestion on a port.	`extensive`

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`MAC statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC subsystem, including the following:
	• `Total octets` and `total packets`—Total number of octets and packets.
	• `Unicast packets`, `Broadcast packets`, and `Multicast packets`—Number of unicast, broadcast, and multicast packets.
	• `CRC/Align errors`—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
	• `FIFO error`—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
	• `MAC control frames`—Number of MAC control frames.
	• `MAC pause frames`—Number of MAC control frames with `pause` operational code.
	• `Oversized frames`—There are two possible conditions regarding the number of oversized frames:
	  • Packet length exceeds 1518 octets, or
	  • Packet length exceeds MRU
	• `Jabber frames`—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
	• `Fragment frames`—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames
		`extensive`

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • `VLAN tagged frames`—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • `Code violations`—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Filter statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • `Input packet count`—Number of packets received from the MAC hardware that the filter processed. • `Input packet rejects`—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • `Input DA rejects`—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local device (which the router is rejecting). • `Input SA rejects`—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • `Output packet count`—Number of packets that the filter has given to the MAC hardware. • `Output packet pad count`—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • `Output packet error count`—Number of packets with an indicated error that the filter was given to transmit. These packets are usually	`extensive`

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • `CAM destination filters, CAM source filters`—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields must be 0.	

`Autonegotiation information`	Information about link autonegotiation.
	• `Negotiation status:`
	  • `Incomplete`—Ethernet interface has the speed or link mode configured.
	  • `No autonegotiation`—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • `Complete`—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
		`extensive`

`Packet Forwarding Engine configuration`	Information about the configuration of the Packet Forwarding Engine: • `Destination slot`—FPC slot number.	`extensive`

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`CoS information`	Information about the CoS queue for the physical interface. • `CoS transmit queue`—Queue number and its associated user-configured forwarding class name. • `Bandwidth %`—Percentage of bandwidth allocated to the queue. • `Bandwidth bps`—Bandwidth allocated to the queue (in bps). • `Buffer %`—Percentage of buffer space allocated to the queue. • `Buffer usec`—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • `Priority`—Queue priority: `low` or `high`. • `Limit`—Displayed if rate limiting is configured for the queue. Possible values are `none` and `exact`. If `exact` is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If `none` is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	`extensive`
`Interface transmit statistics`	Status of the `interface-transmit-statistics` configuration: Enabled or Disabled.	`detail extensive`
`Queue counters (Egress)`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	`detail extensive`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Index`	Index number of the logical interface, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP interface index number for the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Flags`	Information about the logical interface.	All levels
`Encapsulation`	Encapsulation on the logical interface.	All levels
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the specified interface set. • `Input bytes, Output bytes`—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • `Input packets, Output packets`—Number of packets received and transmitted on the interface set.	`detail extensive`
`Local statistics`	Number and rate of bytes and packets destined to the device.	`extensive`

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Transit statistics`	Number and rate of bytes and packets transiting the switch. **NOTE**: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the `Output bytes` and `Output packets` interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	`extensive`
`Security`	Security zones that interface belongs to.	`extensive`
`Flow Input statistics`	Statistics on packets received by flow module.	`extensive`
`Flow Output statistics`	Statistics on packets sent by flow module.	`extensive`
`Flow error statistics (Packets dropped due to)`	Statistics on errors in the flow module.	`extensive`
`Protocol`	Protocol family.	`detail extensive` none
`MTU`	Maximum transmission unit size on the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Route Table`	Route table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail extensive` none

Table 612: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Information about protocol family flags. .	`detail extensive`
`Addresses, Flags`	Information about the address flags..	`detail extensive` none
`Destination`	IP address of the remote side of the connection.	`detail extensive` none
`Local`	IP address of the logical interface.	`detail extensive` none
`Broadcast`	Broadcast address of the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Sample output for G.fast and Annex J support

show interfaces (SRX380, SRX300, SRX320, SRX340, and SRX345)

```
user@host>  show interfaces ge-0/0/8
Physical interface: ge-0/0/8, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 520
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Link-mode: Full-duplex, Speed: 1000mbps,
BPDU Error: None, Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled, Auto-negotiation: Enabled, Remote fault: Online
  DSL SFP Status:
    Chip Type             : xDSL
    Chip Firmware Version : 1_62_8463
    Training Status       : Showtime
    Training Mode         : ADSL2PLUS
    Annex Type            : Annex J
```

```
  Profile Type          : NA
  Carrier Set           : NA
  Line Status           : No Defect
DSL SFP Statistics:                     XTU-R (DS)           XTU-C (US)
  Packet Count             :                    0                    0
  CRC Error Count          :                    0                    0
  Electrical Length (dB)   :                    0                    0
  Net Data Rate (Kbps)     :                25737                 3143
  SNR Margin (dB)          :                  100                   -2
  CV Count                 :                    0                    0
  ES Count                 :                    0                    0
  SES Count                :                    0                    0
  UAS Count                :                    0                    0
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Current address: 4c:16:fc:de:30:89, Hardware address: 4c:16:fc:de:30:89
Last flapped   : 2020-10-28 19:56:29 PDT (3d 23:07 ago)
Input rate     : 20544 bps (42 pps)
Output rate    : 20544 bps (42 pps)
Active alarms  : None
Active defects : None
PCS statistics                      Seconds
  Bit errors                             0
  Errored blocks                         0
Ethernet FEC statistics              Errors
  FEC Corrected Errors                    0
  FEC Uncorrected Errors                  0
  FEC Corrected Errors Rate               0
  FEC Uncorrected Errors Rate             0
Interface transmit statistics: Disabled

Logical interface ge-0/0/8.0 (Index 77) (SNMP ifIndex 538)
  Flags: Up SNMP-Traps 0x0 VLAN-Tag [ 0x8100.10 ]  Encapsulation: ENET2
  Input packets : 0
  Output packets: 105040
  Security: Zone: trust
  Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp
  ospf ospf3 pgm pim rip ripng router-discovery rsvp sap vrrp dhcp finger ftp
  tftp ident-reset http https ike netconf ping reverse-telnet reverse-ssh
  rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text xnm-ssl
  lsping ntp sip dhcpv6 r2cp webapi-clear-text webapi-ssl tcp-encap
```

```
    sdwan-appqoe l3-ha
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 0,
    Curr new hold cnt: 0, NH drop cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.3/24, Local: 10.1.3.2, Broadcast: 10.1.3.255
```

show interfaces (G.fast related information on SRX380, SRX300, SRX320, SRX340, and SRX345)

```
user@host>  show interfaces ge-0/0/8
Physical interface: ge-0/0/8, Enabled, Physical link is Up
G.fast mode, DS Speed: 400Mbps, US Speed: 400Mbps
Cont.........................................
.............................................
```

show interfaces terse (ACX5448, ACX5448-D, ACX710 channelized interface)

```
user@host>  show interfaces terse et-0/1/2
Interface        Admin Link Proto      Local Remote
       et-0/1/2:0             up                 down
       et-0/1/2:1             up                 down
       et-0/1/2:2             up                 down
       et-0/1/2:3             up                 down
```

show interfaces (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
```

```
  Last flapped   : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push     0x8100.512
0x8100.513)
    Encapsulation: VLAN-CCC
    Egress account overhead: 100
    Ingress account overhead: 90
    Input packets : 0
    Output packets: 0
    Protocol ccc, MTU: 1522
      Flags: Is-Primary
```

show interfaces (Gigabit Ethernet on MX Series Routers)

```
user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None, Loopback:
Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote fault:
Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Current address: 00:00:5e:00:53:c0, Hardware address: 00:00:5e:00:53:76
  Last flapped   : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
```

```
        Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 10232
    Output packets: 10294
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255    Protocol
inet6, MTU: 1500
    Max nh cache: 4, New hold nh limit: 100000, Curr nh cnt: 4, Curr new hold cnt: 4, NH drop
cnt: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 2001:db8:/32, Local: 2001:db8::5
      Addresses, Flags: Is-Preferred
        Destination: 2001:db8:1::/32, Local: 2001:db8:223:9cff:fe9f:3e78
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

show interfaces (link degrade status)

```
user@host> show interfaces et-3/0/0
Physical interface: et-3/0/0, Enabled, Physical link is Down
  Interface index: 157, SNMP ifIndex: 537
  Link-level type: Ethernet, MTU: 1514, MRU: 0, Speed: 100Gbps, BPDU Error: None, Loopback:
Disabled, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 54:e0:32:23:9d:38, Hardware address: 54:e0:32:23:9d:38
  Last flapped   : 2014-06-18 02:36:38 PDT (02:50:50 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Link Degrade* :
  Link Monitoring                   :  Enable
```

```
  Link Degrade Set Threshold:       :  1E-7
  Link Degrade Clear Threshold:     :  1E-12
  Estimated BER                     :  1E-7
  Link-degrade event                :  Seconds     Count    State
                                          782         1     Defect Active
```

show interfaces et-0/0/0 (25-Gigabit Ethernet interfaces on PTX Series for default FEC) (Junos OS Evolved Release)

```
user@host> show interfaces et-0/0/0
Physical interface: et-0/0/0, Enabled, Physical link is Up
  Interface index: 1007, SNMP ifIndex: 503
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 25Gbps, BPDU Error: None, Loop
Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled, Media type: Fiber
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 84:03:28:eb:d4:44, Hardware address: 84:03:28:eb:d4:44
  Last flapped   : 2021-05-03 13:23:03 PDT (01:05:00 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Ethernet FEC Mode  :                  FEC91  <<< Default FEC setting starting 21.1R1-EVO
  Ethernet FEC statistics              Errors
    FEC Corrected Errors                    0
    FEC Uncorrected Errors                  0
    FEC Corrected Errors Rate               0
    FEC Uncorrected Errors Rate             0
  Interface transmit statistics: Disabled
  Link Degrade :
    Link Monitoring                   :  Disable
  Logical interface et-0/0/0.16386 (Index 1003) (SNMP ifIndex 611)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
```

```
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

show interfaces extensive (link degrade status) (PTX10001-36MR)

```
user@host> show interfaces et-0/0/1 extensive
Physical interface: et-0/0/1, Enabled, Physical link is Down
  Interface index: 1017, SNMP ifIndex: 519, Generation: 712964572820
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Media type: Fiber
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x80
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 40:de:ad:28:7a:0a, Hardware address: 40:de:ad:28:7a:0a
  Last flapped   : 2020-08-27 12:05:18 IST (00:50:56 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :     2239638274139000         392152080440 bps
   Output bytes  :                    0                    0 bps
   Input  packets:        2239638274049             49019008 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 1, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                    0                    0
    1                                0                    0                    0
    2                                0                    0                    0
    3                                0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
```

```
  2                   assured-forwarding
  3                   network-control
Active alarms  : LINK
Active defects : LINK, LOCAL-DEGRADE
PCS statistics                      Seconds
  Bit errors                             0
  Errored blocks                         0
Ethernet FEC Mode  :                 FEC119
Ethernet FEC statistics              Errors
  FEC Corrected Errors         166615427699
  FEC Uncorrected Errors                 12
  FEC Corrected Errors Rate         3687323
  FEC Uncorrected Errors Rate             0
MAC statistics:                      Receive         Transmit
  Total octets              2239648398609518                0
  Total packets                2239648398691                0
  Unicast packets              2239648398036                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
  Total errors                             1                0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot:  ( )
CoS information:
  Direction :
Interface transmit statistics: Disabled
```

```
  Link Degrade :
    Link Monitoring                   :  Enable
    Link Degrade Set Threshold        :  1E-5
    Link Degrade Clear Threshold      :  1E-10
    Link Degrade War Set Threshold    :  1E-9
    Link Degrade War Clear Threshold  :  1E-11
    Estimated BER                     :  1E-5
    Link-degrade event                :  Seconds              Count                State
                                         3054                 1                    Defect Active
```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```
user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
  Interface index: 151, SNMP ifIndex: 530, Generation: 154
  Interface flags: SNMP-Traps Internal: 0x4000
   Output bytes  :         240614363944            772721536 bps
   Output packets:           3538446506              1420444 pps
    Direction : Output
  Interface transmit statistics: Enabled

  Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
     Output bytes  :         195560312716            522726272 bps
     Output packets:           4251311146              1420451 pps

user@host> show interfaces ge-5/2/0.0 statistics detail
Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :               271524
     Output bytes  :             37769598
     Input  packets:                 3664
     Output packets:               885790
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :            16681118
      Input  packets:                   0
      Output packets:              362633
```

```
    Local statistics:
     Input  bytes  :               271524
     Output bytes  :               308560
     Input  packets:                 3664
     Output packets:                 3659
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :             37461038                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:               882131                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                   0                 0 bps
      Output bytes  :            16681118                 0 bps
      Input  packets:                   0                 0 pps
      Output packets:              362633                 0 pps
```

show interfaces extensive (MX960 Router with MPC10E-10C-MRATE, MPC10E-15C-MRATE line cards)

```
user@host> show interfaces et-x/y/4 extensive

Physical interface: et-1/0/4, Enabled, Physical link is Up
  Interface index: 240, SNMP ifIndex: 773, Generation: 271
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 400Gbps, BPDU Error: None, Loop Detect
PDU Error: None,
  Ethernet-Switching Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control:
Enabled
Wavelength     : 1552.52 nm, Frequency: 193.10 THz
  Pad to minimum frame size: Disabled
```

show interfaces brief (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
```

```
  Link flags     : None

  Logical interface ge-3/0/2.0
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push     0x8100.512
0x8100.513)
    Encapsulation: VLAN-CCC
    ccc

  Logical interface ge-3/0/2.32767
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
```

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 detail
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  Last flapped   : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Drop   bytes  :                    0                    0 bps
   Drop   packets:                    0                    0 pps
  Ingress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
```

```
    1 expedited-fo                  0                    0                    0
    2 assured-forw                  0                    0                    0
    3 network-cont                  0                    0                    0
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                   0                    0                    0
    1 expedited-fo                  0                    0                    0
    2 assured-forw                  0                    0                    0
    3 network-cont                  0                    0                    0
  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push 0x8100.512
0x8100.513)
    Encapsulation: VLAN-CCC
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                  0
     Output bytes  :                  0
     Input  packets:                  0
     Output packets:                  0
    Local statistics:
     Input  bytes  :                  0
     Output bytes  :                  0
     Input  packets:                  0
     Output packets:                  0
    Transit statistics:
     Input  bytes  :                  0                    0 bps
     Output bytes  :                  0                    0 bps
     Input  packets:                  0                    0 pps
     Output packets:                  0                    0 pps
    Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
      Flags: Is-Primary

  Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
  (Generation 139)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                  0
     Output bytes  :                  0
```

```
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
```

show interfaces extensive (Gigabit Ethernet IQ2)

```
user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70, Generation: 171
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4004000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 256
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:74, Hardware address: 00:00:5e:00:53:74
  Last flapped   : 2007-11-07 21:31:41 PST (02:03:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :          38910844056                 7952 bps
   Output bytes  :              7174605                 8464 bps
   Input  packets:            418398473                   11 pps
   Output packets:                78903                   12 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Ingress traffic statistics at Packet Forwarding Engine:
```

```
 Input  bytes  :          38910799145                 7952 bps
 Input  packets:            418397956                   11 pps
 Drop   bytes  :                    0                    0 bps
 Drop   packets:                    0                    0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort            418390823            418390823                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont                7133                 7133                    0
Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                 1031                 1031                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont               77872                77872                    0
Active alarms  : None
Active defects : None
MAC statistics:                      Receive         Transmit
  Total octets                   38910844056          7174605
  Total packets                    418398473            78903
  Unicast packets               408021893366             1026
  Broadcast packets                       10               12
  Multicast packets                418398217            77865
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0  OTN Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
  Payload Type: 0x08
OTN Transmitted Overhead Bytes:
```

```
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x08
Filter statistics:
  Input packet count               418398473
  Input packet rejects                   479
  Input DA rejects                       479
  Input SA rejects                         0
  Output packet count                                    78903
  Output packet pad count                                    0
  Output packet error count                                  0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
      Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
      Remote fault: OK
  Local resolution:
      Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                            %            bps     %           usec
  0 best-effort            95      950000000    95              0      low    none
  3 network-control         5       50000000     5              0      low    none
  Direction : Input
  CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                            %            bps     %           usec
  0 best-effort            95      950000000    95              0      low    none
  3 network-control         5       50000000     5              0      low    none

Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
  Flags: SNMP-Traps Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :               812400
   Output bytes  :              1349206
   Input  packets:                 9429
   Output packets:                 9449
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
```

```
      Output packets:                    0
    Local statistics:
     Input  bytes  :                 812400
     Output bytes  :                1349206
     Input  packets:                   9429
     Output packets:                   9449
    Transit statistics:
     Input  bytes  :                    0                  7440 bps
     Output bytes  :                    0                  7888 bps
     Input  packets:                    0                    10 pps
     Output packets:                    0                    11 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: 1500, Generation: 169, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
      Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
      Output Filters: F2-ge-3/0/1.0-out (53)
      Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
        Generation: 196
    Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__
```

NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the `show interfaces` command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed information, see the description of the logical interface `Transit statistics` fields in Table 610 on page 20232.

show interfaces (Gigabit Ethernet Unnumbered Interface)

```
user@host> show interfaces ge-3/2/0
Physical interface: ge-3/2/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 50
```

```
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:f8, Hardware address: 00:00:5e:00:53:f8
  Last flapped   : 2006-10-27 04:42:23 PDT (08:01:52 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 624 bps (1 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 0
    Output packets: 6
    Protocol inet, MTU: 1500
      Flags: Unnumbered
      Donor interface: lo0.0 (Index 64)
      Preferred source address: 203.0.113.22
```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
  Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ]  Encapsulation: PPP-over-
    Ethernet
    ACI VLAN:
      Dynamic Profile: aci-vlan-set-profile
    PPPoE:
      Dynamic Profile: aci-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
    Input packets : 9
```

```
    Output packets: 8
    Protocol multiservice, MTU: Unlimited
```

show interfaces (ALI Interface Set)

```
user@host> show interfaces ge-1/0/0.10
  Logical interface ge-1/0/0.10 (Index 346) (SNMP ifIndex 554) (Generation 155)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.10 ]  Encapsulation: ENET2
    Line Identity:
      Dynamic Profile: ali-set-profile
      Circuit-id Remote-id Accept-no-ids
    PPPoE:
      Dynamic Profile: ali-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
    Input packets : 9
    Output packets: 8
    Protocol multiservice, MTU: Unlimited
```

Sample Output Gigabit Ethernet

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```
user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 630, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback: None, Source
filtering: Enabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 1024
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:f6, Hardware address: 00:00:5e:00:53:f6
```

```
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :           6970332384                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:             81050506                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :           6970299398                    0 bps
   Input  packets:             81049992                    0 pps
   Drop   bytes  :                    0                    0 bps
   Drop   packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0,
    MTU errors: 0, Resource errors: 0
  Ingress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort             81049992             81049992                    0
    1 expedited-fo                    0                    0                    0
    2 assured-forw                    0                    0                    0
    3 network-cont                    0                    0                    0
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                     0                    0                    0
    1 expedited-fo                    0                    0                    0
    2 assured-forw                    0                    0                    0
    3 network-cont                    0                    0                    0
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  MAC statistics:                      Receive         Transmit
```

```
  Total octets                    6970332384                0
  Total packets                     81050506                0
  Unicast packets                   81050000                0
  Broadcast packets                      506                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
Filter statistics:
  Input packet count                81050506
  Input packet rejects                   506
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 5
CoS information:
  Direction : Output
  CoS transmit queue          Bandwidth               Buffer Priority   Limit
                       %            bps     %         usec
  0 best-effort       95      950000000    95         0          low     none
  3 network-control    5       50000000     5         0          low     none
  Direction : Input
  CoS transmit queue          Bandwidth               Buffer Priority   Limit
                       %            bps     %         usec
  0 best-effort       95      950000000    95         0          low     none
  3 network-control    5       50000000     5         0          low     none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]  Encapsulation: ENET2
  Egress account overhead: 100
  Ingress account overhead: 90
  Traffic statistics:
   Input  bytes  :                    0
```

```
     Output bytes  :                   46
     Input  packets:                    0
     Output packets:                    1
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                   46
     Input  packets:                    0
     Output packets:                    1
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: 1500, Generation: 253, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255, Generation: 265
    Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```
user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 630, Generation: 47
  Link-level type: Ethernet, MTU: 1514, Speed: 9.294GbpsGbps, Loopback: Disabled
  WAN-PHY mode
  Source filtering: Disabled, Flow control: Enabled Speed Configuration: Auto
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Link flags     : None
```

```
CoS queues     : 4 supported
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:9d, Hardware address: 00:00:5e:00:53:9d
Last flapped   : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :                    0                    0 bps
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  HS Link CRC errors: 0, HS Link FIFO overflows: 0,
  Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
  Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
  Resource errors: 0
Queue counters:       Queued packets  Transmitted packets     Dropped packets
  0 best-effort                   0                    0                   0
  1 expedited-fo                  0                    0                   0
  2 assured-forw                  0                    0                   0
  3 network-cont                  0                    0                   0
Active alarms : LOL, LOS, LBL
Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
PCS statistics             Seconds       Count
  Bit errors                     0           0
  Errored blocks                 0           0
MAC statistics:                      Receive         Transmit
  Total octets                             0                0
  Total packets                            0                0
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
```

```
    Code violations                          0
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                    0
    Output packet pad count                                0
    Output packet error count                              0
    CAM destination filters: 0, CAM source filters: 0
  PMA PHY:              Seconds        Count  State
    PLL lock                  0            0  OK
    PHY light             63159            1  Light Missing
WIS section:
  BIP-B1                      0            0
  SEF                    434430       434438  Defect Active
  LOS                    434430            1  Defect Active
  LOF                    434430            1  Defect Active
  ES-S                   434430
  SES-S                  434430
  SEFS-S                 434430
WIS line:
  BIP-B2                      0            0
  REI-L                       0            0
  RDI-L                       0            0  OK
  AIS-L                  434430            1  Defect Active
  BERR-SF                     0            0  OK
  BERR-SD                     0            0  OK
  ES-L                   434430
  SES-L                  434430
  UAS-L                  434420
  ES-LFE                      0
  SES-LFE                     0
  UAS-LFE                     0
WIS path:
  BIP-B3                      0            0
  REI-P                       0            0
  LOP-P                       0            0  OK
  AIS-P                  434430            1  Defect Active
  RDI-P                       0            0  OK
  UNEQ-P                      0            0  OK
  PLM-P                       0            0  OK
  ES-P                   434430
```

```
    SES-P                434430
    UAS-P                434420
    ES-PFE                    0
    SES-PFE                   0
    UAS-PFE                   0
  Received path trace:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
  Transmitted path trace: orissa so-1/0/0
    6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00   orissa so-1/0/0.
    Packet Forwarding Engine configuration:
      Destination slot: 1
    CoS information:
      CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                           %            bps     %    bytes
      0 best-effort       95      950000000    95        0          low     none
      3 network-control    5       50000000     5        0          low     none
```

show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```
user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
  Interface index: 143, SNMP ifIndex: 508, Generation: 208
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  Wavelength     : 1550.12 nm, Frequency: 193.40 THz
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:72, Hardware address: 00:00:5e:00:53:72
  Last flapped   : 2011-04-20 15:48:54 PDT (18:39:49 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
```

```
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
 Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                    0                    0                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:                      Receive         Transmit
  Total octets                             0                0
  Total packets                            0                0
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
  Total octets                             0                0
  Total packets                            0                0
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
```

```
  CRC/Align errors                        0                0
  FIFO errors                             0                0
  MAC control frames                      0                0
  MAC pause frames                        0                0
  Oversized frames                        0
  Jabber frames                           0
  Fragment frames                         0
  VLAN tagged frames                      0
  Code violations                         0
OTN alarms            :   None
OTN defects           :   None
OTN FEC Mode       : GFEC
OTN Rate               : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback : Enabled
OTN FEC statistics :
  Corrected Errors                                        0
  Corrected Error Ratio (          0 sec average)     0e-0
OTN FEC alarms:          Seconds        Count  State
  FEC Degrade                  0            0  OK
  FEC Excessive                0            0  OK
OTN OC:                  Seconds        Count  State
  LOS                          2            1  OK
  LOF                      67164            2  Defect Active
  LOM                      67164           71  Defect Active
  Wavelength Lock              0            0  OK
OTN OTU:
  AIS                          0            0  OK
  BDI                      65919         4814  Defect Active
  IAE                      67158            1  Defect Active
  TTIM                         7            1  OK
  SF                       67164            2  Defect Active
  SD                       67164            3  Defect Active
  TCA-ES                       0            0  OK
  TCA-SES                      0            0  OK
  TCA-UAS                     80           40  OK
  TCA-BBE                      0            0  OK
  BIP                          0            0  OK
  BBE                          0            0  OK
  ES                           0            0  OK
  SES                          0            0  OK
  UAS                        587            0  OK
  Received DAPI:
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
```

```
    Received SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  OTN Received Overhead Bytes:
    APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
    Payload Type: 0x03
  OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x03
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                      0
    Output packet pad count                                  0
    Output packet error count                                0
    CAM destination filters: 0, CAM source filters: 0
  Packet Forwarding Engine configuration:
    Destination slot: 7
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 best-effort            95     9500000000    95              0      low    none
    3 network-control         5      500000000     5              0      low    none
    ...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```
user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 212, Generation: 174
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional: Enabled,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```
user@host> show interfaces xe-7/0/0-tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :      322891152287160           9627472888 bps
   Input  packets:                    0                    0 pps
   Output packets:         328809727380              1225492 pps

...

  Filter statistics:
    Output packet count           328810554250
    Output packet pad count                  0
    Output packet error count                0
...

  Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)
    Flags: SNMP-Traps Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :      322891152287160
     Input  packets:                    0
     Output packets:         328809727380
     IPv6 transit statistics:
      Input  bytes  :                    0
      Output bytes  :                    0
      Input  packets:                    0
```

```
      Output packets:                  0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :      322891152287160           9627472888 bps
     Input  packets:                   0                    0 pps
     Output packets:         328809727380              1225492 pps
     IPv6 transit statistics:
      Input  bytes  :                  0
      Output bytes  :                  0
      Input  packets:                  0
      Output packets:                  0
    Protocol inet, MTU: 1500, Generation: 147, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255, Generation: 141
    Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```
user@host> show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional: Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped   : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :      322857456303482           9627496104 bps
   Output bytes  :                    0                    0 bps
```

```
   Input  packets:         328775413751              1225495 pps
   Output packets:                    0                    0 pps

...

  Filter statistics:
    Input packet count            328775015056
    Input packet rejects                     1
    Input DA rejects                         0

...

  Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :      322857456303482
     Output bytes  :                    0
     Input  packets:         328775413751
     Output packets:                    0
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :      322857456303482           9627496104 bps
     Output bytes  :                    0                    0 bps
     Input  packets:         328775413751              1225495 pps
     Output packets:                    0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: 1500, Generation: 145, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255, Generation: 139
    Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
```

```
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces media (Gigabit Ethernet, 100GE)

```
user@host show interfaces et-8/0/6 media
Physical interface: et-8/0/6, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 821
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, BPDU Error: None,
  Loop Detect PDU Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: 64:64:9b:b0:6e:60, Hardware address: 64:64:9b:b0:6e:60
  Last flapped   : 2020-07-02 06:35:32 IST (00:12:10 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Ethernet FEC Mode  :                  FEC91
  Ethernet FEC statistics              Errors
    FEC Corrected Errors                    0
    FEC Uncorrected Errors                  0
    FEC Corrected Errors Rate               0
    FEC Uncorrected Errors Rate             0
  MAC statistics:
    Input bytes: 0, Input packets: 0, Output bytes: 0, Output packets: 0
  Filter statistics:
    Filtered packets: 0, Padded packets: 0, Output packet errors: 0
  Interface transmit statistics: Disabled
```

show interfaces extensive (Gigabit Ethernet, 100GE)

```
user@host show interfaces et-0/0/0 extensive
Physical interface: et-0/0/0, Enabled, Physical link is Up
  Interface index: 1067, SNMP ifIndex: 548, Generation: 962072676083
  Link-level type: Ethernet, MTU: 1560, LAN-PHY mode, Speed: 100Gbps, BPDU Error: None,
  Loop Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Media type: Fiber
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 1c:9c:8c:a0:e2:f4, Hardware address: 1c:9c:8c:a0:e2:f4
  Last flapped   : 2021-04-21 14:29:36 JST (1d 00:37 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              5675584                 512 bps
   Output bytes  :              5702400                 512 bps
   Input  packets:                88681                   1 pps
   Output packets:                89100                   1 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0,
    L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
    HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                           89098                89098                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
```

```
  Bit errors                              0
  Errored blocks                          0
Ethernet FEC Mode  :                   FEC91
Ethernet FEC statistics               Errors
  FEC Corrected Errors                    11
  FEC Uncorrected Errors                   1
  FEC Corrected Errors Rate                0
  FEC Uncorrected Errors Rate              0
MAC statistics:                       Receive         Transmit
  Total octets                        5675648          5702464
  Total packets                         88682            89101
  Unicast packets                           0                0
  Broadcast packets                         0                0
  Multicast packets                     88682            89101
  CRC/Align errors                          0                0
  FIFO errors                               0                0
  MAC control frames                        0                0
  MAC pause frames                          0                0
  Oversized frames                          0
  Jabber frames                             0
  Fragment frames                           0
  VLAN tagged frames                        0
  Code violations                           0
  Total errors                              0                0
Priority Flow Control Watchdog Statistics:
  Queue          Detected        Recovered  LastPacketDropCount  TotalPacketDropCount
    0                   0                0                    0                     0
    1                   0                0                    0                     0
    2                   0                0                    0                     0
    3                   0                0                    0                     0
    4                   0                0                    0                     0
    5                   0                0                    0                     0
    6                   0                0                    0                     0
    7                   0                0                    0                     0
Filter statistics:
  Input packet count                        0
  Input packet rejects                      0
  Input DA rejects                          0
  Input SA rejects                          0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
```

```
  Packet Forwarding Engine configuration:
    Destination slot:  ( )
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 best-effort            95    95000000000    95              0      low    none
    3 network-control         5     5000000000     5              0      low    none
  PRBS Mode : Disabled
  Interface transmit statistics: Disabled
  Link Degrade :
    Link Monitoring                   :  Disable
```

Sample Output

Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
```

```
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255
```

Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255
```

show interfaces (Gigabit Ethernet for vSRX Virtual Firewall and vSRX Virtual Firewall 3.0)

```
user@host> show interfaces ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Up
```

```
  Interface index: 136, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Link-mode: Half-duplex, Speed: 1000mbps,
BPDU Error: None, Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:50:56:93:ef:25, Hardware address: 00:50:56:93:ef:25
  Last flapped   : 2019-03-29 01:57:45 UTC (00:00:41 ago)
  Input rate     : 1120 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
```

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1 detail
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps, BPDU Error:
None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                   0                   0                    0
    1 expedited-fo                  0                   0                    0
    2 assured-forw                  0                   0                    0
    3 network-cont                  0                   0                    0
```

```
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : LINK
Active defects : LINK
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Local statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Transit statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Security: Zone: public
  Flow Statistics :
  Flow Input statistics :
    Self packets :                     0
    ICMP packets :                     0
    VPN packets :                      0
    Multicast packets :                0
    Bytes permitted by policy :        0
    Connections established :          0
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        0
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
```

```
      Multiple user authentications:     0
      Multiple incoming NAT:             0
      No parent for a gate:              0
      No one interested in self packets: 0
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
      No route present:                  0
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 1500, Generation: 150, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255, Generation: 150
```

show interfaces statistics st0.0 detail

```
user@host> show interfaces statistics st0.0 detail
  Logical interface st0.0 (Index 71) (SNMP ifIndex 609) (Generation 136)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
    Traffic statistics:
     Input  bytes  :         528152756774
     Output bytes  :         575950643520
     Input  packets:          11481581669
     Output packets:          12520666095
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0            121859888 bps
     Output bytes  :                    0            128104112 bps
     Input  packets:                    0               331141 pps
```

```
     Output packets:                    0               348108 pps
    Security: Zone: untrust
    Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp ospf ospf3 pgm
pim rip ripng router-discovery rsvp
    sap vrrp
    Flow Statistics :
    Flow Input statistics :
      Self packets :                     0
      ICMP packets :                     0
      VPN packets :                      0
      Multicast packets :                0
      Bytes permitted by policy :        525984295844
      Connections established :          7
    Flow Output statistics:
      Multicast packets :                0
      Bytes permitted by policy :        576003290222
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
      Authentication failed:             0
      Incoming NAT errors:               0
      Invalid zone received packet:      0
      Multiple user authentications:     0
      Multiple incoming NAT:             0
      No parent for a gate:              0
      No one interested in self packets: 0
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
      No route present:                  2000280
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 9192
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
    Generation: 155, Route table: 0
```

```
      Flags: Sendbcast-pkt-to-re
```

show interfaces extensive (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
```

```
  2                   assured-forwarding
  3                   network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:                       Receive         Transmit
  Total octets                              0                0
  Total packets                             0                0
  Unicast packets                           0                0
  Broadcast packets                         0                0
  Multicast packets                         0                0
  CRC/Align errors                          0                0
  FIFO errors                               0                0
  MAC control frames                        0                0
  MAC pause frames                          0                0
  Oversized frames                          0
  Jabber frames                             0
  Fragment frames                           0
  VLAN tagged frames                        0
  Code violations                           0
Filter statistics:
  Input packet count                        0
  Input packet rejects                      0
  Input DA rejects                          0
  Input SA rejects                          0
  Output packet count                                        0
  Output packet pad count                                    0
  Output packet error count                                  0
  CAM destination filters: 2, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth                Buffer Priority   Limit
                            %            bps     %           usec
  0 best-effort            95      950000000    95              0      low    none
  3 network-control         5       50000000     5              0      low    none
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
```

```
 Input  bytes  :                    0
 Output bytes  :                    0
 Input  packets:                    0
 Output packets:                    0
Local statistics:
 Input  bytes  :                    0
 Output bytes  :                    0
 Input  packets:                    0
 Output packets:                    0
Transit statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :                    0                    0 bps
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
Security: Zone: public
Flow Statistics :
Flow Input statistics :
  Self packets :                     0
  ICMP packets :                     0
  VPN packets :                      0
  Multicast packets :                0
  Bytes permitted by policy :        0
  Connections established :          0
Flow Output statistics:
  Multicast packets :                0
  Bytes permitted by policy :        0
Flow error statistics (Packets dropped due to):
  Address spoofing:                  0
  Authentication failed:             0
  Incoming NAT errors:               0
  Invalid zone received packet:      0
  Multiple user authentications:     0
  Multiple incoming NAT:             0
  No parent for a gate:              0
  No one interested in self packets: 0
  No minor session:                  0
  No more sessions:                  0
  No NAT gate:                       0
  No route present:                  0
  No SA for incoming SPI:            0
  No tunnel found:                   0
  No session for a gate:             0
  No zone or NULL zone binding       0
```

```
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 1500, Generation: 150, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
        Generation: 150
```

show interfaces terse

```
user@host> show interfaces terse
    Interface               Admin Link Proto    Local                 Remote
    ge-0/0/0                up    up
    ge-0/0/0.0              up    up   inet     10.209.4.61/18
    gr-0/0/0                up    up
    ip-0/0/0                up    up
    st0                     up    up
    st0.1                   up    ready inet
    ls-0/0/0                up    up
    lt-0/0/0                up    up
    mt-0/0/0                up    up
    pd-0/0/0                up    up
    pe-0/0/0                up    up
    e3-1/0/0                up    up
    t3-2/0/0                up    up
    e1-3/0/0                up    up
    se-4/0/0                up    down
    t1-5/0/0                up    up
    br-6/0/0                up    up
    dc-6/0/0                up    up
    dc-6/0/0.32767          up    up
    bc-6/0/0:1              down  up
    bc-6/0/0:1.0            up    down
    dl0                     up    up
    dl0.0                   up    up   inet
    dsc                     up    up
```

```
    gre                     up    up
    ipip                    up    up
    lo0                     up    up
    lo0.16385               up    up   inet     10.0.0.1            --> 0/0
                                                 10.0.0.16           --> 0/0
    lsi                     up    up
    mtun                    up    up
    pimd                    up    up
    pime                    up    up
    pp0                     up    up
```

show interfaces terse (vSRX Virtual Firewall and vSRX Virtual Firewall 3.0)

```
user@host> show interfaces terse
Interface            Admin Link Proto    Local                 Remote
ge-0/0/0              up    up
ge-0/0/0.0            up    up               inet     10.1.65.1/24
ge-0/0/1              up    up
ge-0/0/2              up    up
e-0/0/3               up    up
ge-0/0/4              up    up
```

show interfaces controller (Channelized E1 IQ with Logical E1)

```
user@host> show interfaces controller ce1-1/2/6

Controller                                                        Admin Link
ce1-1/2/6                                                         up    up
e1-1/2/6                                                          up    up
```

show interfaces controller (Channelized E1 IQ with Logical DS0)

```
user@host> show interfaces controller ce1-1/2/3

Controller                                                        Admin Link
ce1-1/2/3                                                         up    up
```

```
ds-1/2/3:1                                                          up    up
ds-1/2/3:2                                                          up    up
```

show interfaces descriptions

```
user@host> show interfaces descriptions
 Interface       Admin Link Description
 so-1/0/0        up    up   M20-3#1
 so-2/0/0        up    up   GSR-12#1
 ge-3/0/0        up    up   SMB-OSPF_Area300
 so-3/3/0        up    up   GSR-13#1
 so-3/3/1        up    up   GSR-13#2
 ge-4/0/0        up    up   T320-7#1
 ge-5/0/0        up    up   T320-7#2
 so-7/1/0        up    up   M160-6#1
 ge-8/0/0        up    up   T320-7#3
 ge-9/0/0        up    up   T320-7#4
 so-10/0/0       up    up   M160-6#2
 so-13/0/0       up    up   M20-3#2
 so-14/0/0       up    up   GSR-12#2
 ge-15/0/0       up    up   SMB-OSPF_Area100
 ge-15/0/1       up    up   GSR-13#3
```

show interfaces destination-class all

```
user@host> show interfaces destination-class all
Logical interface so-4/0/0.0
                                              Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                        gold                        0                    0
                                 (                  0) (                 0)
                      silver                        0                    0
                                 (                  0) (                 0)
  Logical interface so-0/1/3.0
                                              Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                        gold                        0                    0
                                 (                  0) (                 0)
```

```
                        silver                        0                    0
                                   (                  0) (                 0)
```

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                        :  7.408 mA
    Laser output power                        :  0.3500 mW / -4.56 dBm
    Module temperature                        :  23 degrees C / 73 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
```

```
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:00-04:15:
...
```

show interfaces far-end-interval coc1-5/2/1:1

```
user@host> run show interfaces far-end-interval coc1-5/2/1:1
 Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
```

```
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:00-04:15:
```

show interfaces filters

```
user@host> show interfaces filters
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
ge-5/0/0        up    up
ge-5/0/0.0      up    up   any                        f-any
                           inet                       f-inet
                           multiservice
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
vt-0/3/0        up    up
at-1/0/0        up    up
at-1/0/0.0      up    up   inet
                           iso
at-1/1/0        up    down
at-1/1/0.0      up    down inet
                           iso
....
```

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
```

```
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp  ldp msdp nhrp ospf pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http https ike
    netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                     0
      ICMP packets :                     0
      VPN packets :                      2564
      Bytes permitted by policy :        3478
      Connections established :          1
    Flow Output statistics:
      Multicast packets :                0
      Bytes permitted by policy :        16994
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
      Authentication failed:             0
      Incoming NAT errors:               0
      Invalid zone received packet:      0
      Multiple user authentications:     0
      Multiple incoming NAT:             0
      No parent for a gate:              0
      No one interested in self packets: 0
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
      No route present:                  0
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
```

```
        Destination:     203.0.113.1/24, Local: 203.0.113.2, Broadcast: 203.0.113.255
```

show interfaces interval (Channelized OC12)

```
user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (E3)

```
user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
```

```
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH) (SRX Series Firewalls)

```
user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
  20:02-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:47-20:02:
    ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
    ES-P: 267, SES-P: 267, UAS-P: 267
  19:32-19:47:
    ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,     SES-P: 56,
UAS-P: 46
  19:17-19:32:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:02-19:17:
.....
```

show interfaces load-balancing (SRX Series Firewalls)

```
user@host> show interfaces load-balancing
Interface  State            Last change  Member count
ams0       Up               1d 00:50     2
ams1       Up               00:00:59     2
```

show interfaces load-balancing detail (SRX Series Firewalls)

```
user@host>show interfaces load-balancing detail
Load-balancing interfaces detail
Interface        : ams0
  State          : Up
  Last change    : 1d 00:51
  Member count   : 2
  Members        :
      Interface    Weight   State
      mams-2/0/0    10        Active
      mams-2/1/0    10        Active
```

show interfaces mac-database (All MAC Addresses on a Port SRX Series Firewalls)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback: None, Source
filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
```

```
  00:00:c8:01:01:02      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:03      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:04      30424716     1399536936       37448523     1722632058
  00:00:c8:01:01:05      30424789     1399540294       37448598     1722635508
  00:00:c8:01:01:06      30424788     1399540248       37448597     1722635462
  00:00:c8:01:01:07      30424783     1399540018       37448597     1722635462
  00:00:c8:01:01:08      30424783     1399540018       37448596     1722635416
  00:00:c8:01:01:09       8836796      406492616        8836795      406492570
  00:00:c8:01:01:0a      30424712     1399536752       37448521     1722631966
  00:00:c8:01:01:0b      30424715     1399536890       37448523     1722632058
  Number of MAC addresses : 21
```

show interfaces mac-database (All MAC Addresses on a Service SRX Series Firewalls)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      31016568     1426762128       38040381     1749857526
  00:00:c8:01:01:03      31016568     1426762128       38040382     1749857572
  00:00:c8:01:01:04      31016499     1426758954       38040306     1749854076
  00:00:c8:01:01:05      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:06      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:07      31016567     1426762082       38040380     1749857480
  00:00:c8:01:01:08      31016567     1426762082       38040379     1749857434
  00:00:c8:01:01:09       9428580      433714680        9428580      433714680
  00:00:c8:01:01:0a      31016496     1426758816       38040304     1749853984
  00:00:c8:01:01:0b      31016498     1426758908       38040307     1749854122
```

show interfaces mac-database mac-address

```
user@host> show interfaces mac-database xe-0/3/3 mac-address  (SRX devices)
00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback: None, Source
filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes     : 202324652
    Output bytes    : 202324560
    Input frames    : 4398362
    Output frames   : 4398360
  Policer statistics:
  Policer type      Discarded frames    Discarded bytes
  Output aggregate           3992386          183649756
```

show interfaces mc-ae (SRX Series Firewalls)

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links   : ae0
  Local Status   : active
  Peer Status    : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL         : Label Ethernet Interface
```

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the `show interfaces media` command for a SONET interface (with no level of output specified):

```
user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,   Loopback:
None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,   mpls: Not-
configured
  CHAP state: Not-configured
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
  SONET errors:
     BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

show interfaces policers (SRX Series Firewalls)

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
```

```
pd-0/3/0        up    up
pe-0/3/0        up    up
...
so-2/0/0        up    up
so-2/0/0.0      up    up   inet  so-2/0/0.0-in-policer so-2/0/0.0-out-policer
                           iso
so-2/1/0        up    down
...
```

show interfaces policers interface-name (SRX Series Firewalls)

```
user@host> show interfaces policers so-2/1/0
Interface       Admin Link Proto Input Policer         Output Policer
so-2/1/0        up    down
so-2/1/0.0      up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                           iso
                           inet6
```

show interfaces queue (SRX Series Firewalls)

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```
user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: class0
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
```

```
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
  Queue Buffer Usage:
    Reserved buffer      :             118750000 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
Queue: 1, Forwarding classes: class1
..
..
 Queue Buffer Usage:
    Reserved buffer      :                  9192 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
Queue: 3, Forwarding classes: class3
  Queued:
..
..
Queue Buffer Usage:
    Reserved buffer      :               6250000 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
```

show interfaces redundancy (SRX Series Firewalls)

```
user@host> show interfaces redundancy
Interface  State         Last change  Primary    Secondary  Current status
rsp0       Not present                sp-1/0/0   sp-0/2/0   both down
rsp1       On secondary  1d 23:56     sp-1/2/0   sp-0/3/0   primary down
```

```
rsp2       On primary    10:10:27     sp-1/3/0   sp-0/2/0   secondary down
rlsq0      On primary    00:06:24     lsq-0/3/0  lsq-1/0/0  both up
```

show interfaces redundancy (Aggregated Ethernet SRX Series Firewalls)

```
user@host> show interfaces redundancy
Interface  State            Last change  Primary      Secondary    Current status
rlsq0      On secondary     00:56:12     lsq-4/0/0    lsq-3/0/0    both up
ae0
ae1
ae2
ae3
ae4
```

show interfaces redundancy detail (SRX Series Firewalls)

```
user@host> show interfaces redundancy detail
Interface        : rlsq0
  State          : On primary
  Last change    : 00:45:47
  Primary        : lsq-0/2/0
  Secondary      : lsq-1/2/0
  Current status : both up
  Mode           : hot-standby

Interface        : rlsq0:0
  State          : On primary
  Last change    : 00:45:46
  Primary        : lsq-0/2/0:0
  Secondary      : lsq-1/2/0:0
  Current status : both up
  Mode           : warm-standby
```

show interfaces routing brief (SRX Series Firewalls)

```
user@host> show interfaces routing brief
Interface        State Addresses
so-5/0/3.0       Down  ISO   enabled
```

```
so-5/0/2.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.120
                       INET  enabled
so-5/0/1.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.130
                       INET  enabled
at-1/0/0.3       Up    CCC   enabled
at-1/0/0.2       Up    CCC   enabled
at-1/0/0.0       Up    ISO   enabled
                       INET  192.168.90.10
                       INET  enabled
lo0.0            Up    ISO   47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
                       ISO   enabled
                       INET  127.0.0.1
fxp1.0           Up
fxp0.0           Up    INET  192.168.6.90
```

show interfaces routing detail (SRX Series Firewalls)

```
user@host> show interfaces routing detail
so-5/0/3.0
  Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>
  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  ISO address (null)
    State: <Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>
  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  MPLS address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
  ISO address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  INET address 192.168.2.120
    State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
```

```
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
    Local address: 192.168.2.120
    Destination: 192.168.2.110/32
  INET address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...
```

show interfaces routing-instance all (SRX Series Firewalls)

```
user@host> show interfaces terse routing-instance all
Interface   Admin   Link   Proto    Local             Remote Instance
at-0/0/1       up      up     inet        10.0.0.1/24
ge-0/0/0.0  up      up     inet      192.168.4.28/24          sample-a
at-0/1/0.0  up      up     inet6     fe80::a:0:0:4/64         sample-b
so-0/0/0.0     up      up     inet        10.0.0.1/32
```

show interfaces snmp-index (SRX Series Firewalls)

```
user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
  Interface index: 149, SNMP ifIndex: 33
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,   Loopback:
None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : LOL, PLL, LOS
  SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P
```

show interfaces source-class all (SRX Series Firewalls)

```
user@host> show interfaces source-class all
Logical interface so-0/1/0.0
```

```
                                                Packets                Bytes
    Source class                  (packet-per-second)    (bits-per-second)
                            gold               1928095            161959980
                                 (                889) (            597762)
                          bronze                     0                   0
                                 (                  0) (                 0)
                          silver                     0                   0
                                 (                  0) (                 0)
  Logical interface so-0/1/3.0
                                                Packets                Bytes
    Source class                  (packet-per-second)    (bits-per-second)
                            gold                     0                   0
                                 (                  0) (                 0)
                          bronze                     0                   0
                                 (                  0) (                 0)
                          silver                116113             9753492
                                 (                939) (            631616)
```

show interfaces statistics (Fast Ethernet SRX Series Firewalls)

```
user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : None
  Active defects : None
  Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary, DCU, SCU-in
```

```
                                                Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                          silver1                    0                    0
                                  (                  0) (                 0)
                          silver2                    0                    0
                                  (                  0) (                 0)
                          silver3                    0                    0
                                  (                  0) (                 0)
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10.27.245/24, Local: 10.27.245.2,
        Broadcast: 10.27.245.255
    Protocol iso, MTU: 1497
      Flags: Is-Primary
```

show interfaces switch-port (SRX Series Firewalls)

```
user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
           Speed: 100mbps, Auto-negotiation: Enabled
            Statistics:                          Receive         Transmit
             Total bytes                       28437086         21792250
             Total packets                       409145            88008
             Unicast packets                       9987            83817
             Multicast packets                   145002                0
             Broadcast packets                   254156             4191
             Multiple collisions                     23               10
             FIFO/CRC/Align errors                    0                0
             MAC pause frames                         0                0
             Oversized frames                         0
             Runt frames                              0
             Jabber frames                            0
             Fragment frames                          0
             Discarded frames                         0
           Autonegotiation information:
             Negotiation status: Complete
             Link partner:
                 Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link partner
Speed: 100 Mbps
             Local resolution:
                 Flow control: None, Remote fault: Link OK
```

show interfaces transport pm (SRX Series Firewalls)

```
user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current          Elapse time:900 Seconds
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  OTU-BBE                 0                800              No              No
  OTU-ES                  0                135              No              No
  OTU-SES                 0                90               No              No
  OTU-UAS                 427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  OTU-BBE                 0                800              No              No
  OTU-ES                  0                135              No              No
  OTU-SES                 0                90               No              No
  OTU-UAS                 0                90               No              No
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  ODU-BBE                 0                800              No              No
  ODU-ES                  0                135              No              No
  ODU-SES                 0                90               No              No
  ODU-UAS                 427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  ODU-BBE                 0                800              No              No
  ODU-ES                  0                135              No              No
  ODU-SES                 0                90               No              No
  ODU-UAS                 0                90               No              No
  FEC            Suspect Flag:False             Reason:None
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  FEC-CorrectedErr        2008544300       0                NA              NA
  FEC-UncorrectedWords    0                0                NA              NA
  BER            Suspect Flag:False             Reason:None
    PM                 MIN       MAX       AVG       THRESHOLD      TCA-ENABLED      TCA-RAISED
  BER                  3.6e-5    5.8e-5    3.6e-5     10.0e-3        No                Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current
  Suspect Flag:True             Reason:Object Disabled
    PM                            CURRENT   MIN       MAX      AVG       THRESHOLD      TCA-
ENABLED       TCA-RAISED
                                                                           (MIN) (MAX)    (MIN)
```

```
(MAX)       (MIN) (MAX)
  Lane chromatic dispersion         0        0         0         0         0     0        NA
NA         NA     NA
  Lane differential group delay     0        0         0         0         0     0        NA
NA         NA     NA
    q Value                           120      120       120       120       0     0
NA    NA         NA     NA
  SNR                               28       28        29        28        0     0        NA
NA         NA     NA
  Tx output power(0.01dBm)          -5000    -5000     -5000     -5000     -300  -100     No
No         No     No
  Rx input power(0.01dBm)           -3642    -3665     -3626     -3637     -1800 -500     No
No         No     No
  Module temperature(Celsius)       46       46        46        46        -5    75       No
No         No     No
  Tx laser bias current(0.1mA)      0        0         0         0         0     0        NA
NA         NA     NA
  Rx laser bias current(0.1mA)      1270     1270      1270      1270      0     0        NA
NA         NA     NA
  Carrier frequency offset(MHz)     -186     -186      -186      -186      -5000 5000     No
No         No     No
```

show security zones (SRX Series Firewalls)

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
```

```
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

show interfaces extensive (QFX5130-32CD)

```
user@host> show interfaces et-0/0/29 extensive | no-more
Physical interface: et-0/0/29, Enabled, Physical link is Up
  Interface index: 1086, SNMP ifIndex: 549, Generation: 618475300929
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps,
  BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Media type: Copper
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 0c:59:9c:81:fb:12, Hardware address: 0c:59:9c:81:fb:12
  Last flapped   : 2020-09-14 06:27:45 PDT (1d 01:34 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :      130245061850190                    0 bps
   Output bytes  :      132765627331264                    0 bps
   Input  packets:          86830042098                    0 pps
   Output packets:          88510419103                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 4, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
```

```
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 5 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0                               0                    0                    0
  3                               0                    0                    0
  4                               0                    0                    0
  7                      88510418098          88510418098                    0
  8                               0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  3                   fcoe
  4                   no-loss
  7                   network-control
  8                   mcast
Active alarms  : None
Active defects : None
PCS statistics                      Seconds
  Bit errors                             0
  Errored blocks                         0
Ethernet FEC Mode  :                 FEC119
Ethernet FEC statistics              Errors
  FEC Corrected Errors           5796771110
  FEC Uncorrected Errors                  0
  FEC Corrected Errors Rate               0
  FEC Uncorrected Errors Rate             0
MAC statistics:                      Receive         Transmit
  Total octets               130245061850190  132765627331264
  Total packets                  86830042098      88510419103
  Unicast packets                86830041265      88510418256
  Broadcast packets                      179              190
  Multicast packets                      654              657
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
  Total errors                             0                0
MAC Priority Flow Control Statistics:
  Priority :  0                             0                0
```

```
    Priority :  1                            0                0
    Priority :  2                            0                0
    Priority :  3                            0                0
    Priority :  4                            0                0
    Priority :  5                            0                0
    Priority :  6                            0                0
    Priority :  7                            0                0
  Filter statistics:
    Input packet count                      0
    Input packet rejects                    0
    Input DA rejects                        0
    Input SA rejects                        0
    Output packet count                                     0
    Output packet pad count                                 0
    Output packet error count                               0
    CAM destination filters: 0, CAM source filters: 0
  Packet Forwarding Engine configuration:
    Destination slot:  ( )
  CoS information:
    Direction :
  Interface transmit statistics: Disabled
  Link Degrade :
    Link Monitoring                   :  Disable
```

show interfaces (PTX10001-36MR) (400G ZR and 400G ZR-M optics)

```
user@host> show interfaces et-0/0/10
Physical interface: et-0/0/10, Enabled, Physical link is Up
  Interface index: 1016, SNMP ifIndex: 519
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps,
  BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.40 THz
  Optic-loopback : Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 40:de:ad:8c:ca:22, Hardware address: 40:de:ad:8c:ca:22
  Last flapped   : 2022-02-28 07:32:26 UTC (03:58:05 ago)
  Input rate     : 0 bps (0 pps)
```

```
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Ethernet FEC Mode  :                 FEC119
  Ethernet FEC statistics              Errors
    FEC Corrected Errors               179957
    FEC Uncorrected Errors                  0
    FEC Corrected Errors Rate              18
    FEC Uncorrected Errors Rate             0
  Optic FEC Mode  :                      OFEC
  Optic FEC statistics:
    Corrected Errors                           2695089428427
    Uncorrected Words                                      0
    Corrected Error rate                           194566696
    Uncorrected Error rate                                 0
    Corrected Error Ratio (      14407 seconds average)  4.08e-04
  PRBS Mode : Disabled
  Interface transmit statistics: Disabled
  Link Degrade :
    Link Monitoring                   :  Disable
```

show interfaces (MX304)

```
user@host>show interfaces ge-0/2/12
Physical interface: ge-0/2/12, Enabled, Physical link is Up
  Interface index: 209, SNMP ifIndex: 746
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
None, Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Disabled, Remote fault:
Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
```

```
  Current address: 68:f3:8e:7a:1c:67, Hardware address: 68:f3:8e:7a:1c:67
  Last flapped   : 2023-03-09 16:04:18 IST (00:04:17 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             1
    Errored blocks                         1
  Ethernet FEC Mode  :                   NONE
    FEC Codeword size                       0
    FEC Codeword rate                   0.000
  Ethernet FEC statistics              Errors
    FEC Corrected Errors                    0
    FEC Uncorrected Errors                  0
    FEC Corrected Errors Rate               0
    FEC Uncorrected Errors Rate             0
  PRBS Mode : Disabled
  Interface transmit statistics: Disabled

  Logical interface ge-0/2/12.0 (Index 356) (SNMP ifIndex 748)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 21.1.1/24, Local: 21.1.1.2, Broadcast: 21.1.1.255
    Protocol multiservice, MTU: Unlimited
```

show interfaces (PTX10001-36MR) (Junos OS Evolved Release)

```
user@host> show interfaces et-0/1/0 media
Physical interface: et-0/1/0, Enabled, Physical link is Down
  Interface index: 1017, SNMP ifIndex: 554
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled, Media type: Copper
  Device flags   : Present Running Down
```

```
  Interface flags: Hardware-Down SNMP-Traps
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 88:d9:8f:08:46:57, Hardware address: 88:d9:8f:08:46:57
  Last flapped   : 2021-08-30 02:45:11 PDT (00:00:34 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK, REMOTE-FAULT
  Ethernet FEC Mode  :                 FEC119
  Ethernet FEC statistics              Errors
    FEC Corrected Errors            545479262
    FEC Uncorrected Errors                 24
    FEC Corrected Errors Rate            7942
    FEC Uncorrected Errors Rate             0
  MAC statistics:
    Input bytes: 0, Input packets: 0, Output bytes: 0, Output packets: 0
  Filter statistics:
    Filtered packets: 0, Padded packets: 0, Output packet errors: 0
  PRBS Mode : Enabled
  PRBS Pattern : 7
  PRBS Statistics
    Lane 0 : Number of iterations in error :   4538027 Total iterations :     4538027 Monitored
Seconds: 2
    Lane 1 : Number of iterations in error :   2269036 Total iterations :     2269036 Monitored
Seconds: 1
    Lane 2 : Number of iterations in error :   2269247 Total iterations :     2269247 Monitored
Seconds: 1
    Lane 3 : Number of iterations in error :   4538551 Total iterations :     4538551 Monitored
Seconds: 2
    Lane 4 : Number of iterations in error :   4537975 Total iterations :     4537975 Monitored
Seconds: 2
    Lane 5 : Number of iterations in error :   4538053 Total iterations :     4538053 Monitored
Seconds: 2
    Lane 6 : Number of iterations in error :   4538555 Total iterations :     4538555 Monitored
Seconds: 2
    Lane 7 : Number of iterations in error :   4538433 Total iterations :     4538433 Monitored
Seconds: 2
  Interface transmit statistics: Disabled
  Link Degrade :
    Link Monitoring                   :  Disable

  Logical interface et-0/1/0.16386 (Index 7003) (SNMP ifIndex 560)
    Flags: Up SNMP-Traps Encapsulation: ENET2
```

```
    Input packets : 0
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

show interfaces (SRX devices with IPv4 source-interface configured for the IP tunnel interface)

```
user@host> show interfaces ip-0/0/0
Physical interface: ip-0/0/0, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 533
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Link flags     : Keepalives DTE
  Device flags   : Present Running
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ip-0/0/0.0 (Index 113) (SNMP ifIndex 556)
    Flags: Up Point-To-Point SNMP-Traps 0x0 IP-Header 2.0.0.1:2.0.0.2:4:df:64:00000000
Encapsulation: IPIP-NULL
    Input packets : 0
    Output packets: 0
    Security: Zone: untrust
    Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp ospf ospf3 pgm pim
rip ripng router-discovery rsvp sap vrrp dhcp finger ftp tftp
    ident-reset http https ike netconf ping reverse-telnet reverse-ssh rlogin rpm rsh snmp snmp-
trap ssh telnet traceroute xnm-clear-text xnm-ssl lsping
    lsselfping ntp sip dhcpv6 r2cp webapi-clear-text webapi-ssl tcp-encap sdwan-appqoe high-
availability
    Protocol inet, MTU: 1480
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Flags: Sendbcast-pkt-to-re
    Protocol inet6, MTU: 1480
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Flags: None
      Addresses, Flags: Is-Preferred 0x0
        Destination: fe80::/64, Local: fe80::3a4f:4900:94:a601
```

show interfaces (SRX devices with IPv4 source-hostname and detsination-hostname configured for the IP tunnel interface)

```
user@host> show interfaces ip-0/0/0 Physical interface: ip-0/0/0, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 533
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Link flags     : Keepalives DTE
  Device flags   : Present Running
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ip-0/0/0.0 (Index 115) (SNMP ifIndex 556)
    Flags: Up Point-To-Point SNMP-Traps 0x0 IP-Header 2.0.0.1:2.0.0.2:4:df:64:00000000
Encapsulation: IPIP-NULL
    Input packets : 0
    Output packets: 0
    Security: Zone: untrust
    Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp ospf ospf3 pgm pim
rip ripng router-discovery rsvp sap vrrp dhcp finger ftp tftp
    ident-reset http https ike netconf ping reverse-telnet reverse-ssh rlogin rpm rsh snmp snmp-
trap ssh telnet traceroute xnm-clear-text xnm-ssl lsping
    lsselfping ntp sip dhcpv6 r2cp webapi-clear-text webapi-ssl tcp-encap sdwan-appqoe high-
availability
    Protocol inet, MTU: 1480
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Flags: Sendbcast-pkt-to-re
    Protocol inet6, MTU: 1480
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Flags: None
      Addresses, Flags: Is-Preferred 0x0
        Destination: fe80::/64, Local: fe80::3a4f:4900:94:a601
```

show interfaces (SRX devices with IPv6 source-interface configured for the IP tunnel interface)

```
user@host> show interfaces ip-0/0/0 Physical interface: ip-0/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 519
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Link flags     : Keepalives DTE
```

```
  Device flags   : Present Running
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ip-0/0/0.0 (Index 71) (SNMP ifIndex 550)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 IP-Header 2000::2-2000::1-41-64-0-0-00000000
Encapsulation: IPIP-NULL
    Input packets : 0
    Output packets: 0
    Security: Zone: untrust
    Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp ospf ospf3 pgm pim
rip ripng router-discovery rsvp sap vrrp dhcp finger ftp tftp
    ident-reset http https ike netconf ping reverse-telnet reverse-ssh rlogin rpm rsh snmp snmp-
trap ssh telnet traceroute xnm-clear-text xnm-ssl lsping
    lsselfping ntp sip dhcpv6 r2cp webapi-clear-text webapi-ssl tcp-encap sdwan-appqoe high-
availability
    Protocol inet, MTU: 1460
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Flags: Sendbcast-pkt-to-re
    Protocol inet6, MTU: 1460
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Flags: None
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::250:5600:93:3ea9
```

show interfaces (SRX devices with IPv6 source-hostname and detsination-hostname configured for the IP tunnel interface)

```
user@host> show interfaces ip-0/0/0 Physical interface: ip-0/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 519
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Link flags     : Keepalives DTE
  Device flags   : Present Running
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ip-0/0/0.0 (Index 67) (SNMP ifIndex 550)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 IP-Header 2000::2-2000::1-41-64-0-0-00000000
Encapsulation: IPIP-NULL
    Input packets : 0
```

```
    Output packets: 0
    Security: Zone: untrust
    Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp ospf ospf3 pgm pim
rip ripng router-discovery rsvp sap vrrp dhcp finger ftp tftp
    ident-reset http https ike netconf ping reverse-telnet reverse-ssh rlogin rpm rsh snmp snmp-
trap ssh telnet traceroute xnm-clear-text xnm-ssl lsping
    lsselfping ntp sip dhcpv6 r2cp webapi-clear-text webapi-ssl tcp-encap sdwan-appqoe high-
availability
    Protocol inet, MTU: 1460
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Flags: Sendbcast-pkt-to-re
    Protocol inet6, MTU: 1460
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Flags: None
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::250:5600:93:3ea9
```

show interfaces extensive (ACX5448, ACX5448-D, ACX5448-M and ACX710)

```
user@host> show interfaces xe-0/0/18 extensive
Physical interface: xe-0/0/18, Enabled, Physical link is Up
Interface index: 154, SNMP ifIndex: 525, Generation: 185
Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 1000mbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
Auto-negotiation: Enabled, Remote fault: Offline,
Media type: Fiber
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Schedulers : 0
Hold-times : Up 0 ms, Down 0 ms
Damping : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state: unsuppressed
Current address: 98:a4:04:79:50:14, Hardware address: 98:a4:04:79:50:14
Last flapped : 2021-06-29 09:33:22 PDT (13:43:07 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
```

```
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0, HS
link CRC errors: 0, MTU errors: 0,
Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
0 0 0 0
1 0 0 0
2 0 0 0
3 0 0 0
Queue number: Mapped forwarding classes
0 best-effort
1 expedited-forwarding
2 assured-forwarding
3 network-control
Active alarms : None
Active defects : None
PCS statistics Seconds
Bit errors 0
Errored blocks 0
Ethernet FEC statistics Errors
FEC Corrected Errors 0
FEC Uncorrected Errors 0
FEC Corrected Errors Rate 0
FEC Uncorrected Errors Rate 0
MAC statistics: Receive Transmit
Total octets 0 0
Total packets 0 0
Unicast packets 0 0
Broadcast packets 0 0
Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
```

```
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
PRBS Mode : Disabled
Filter statistics:
Input packet count 0
Input packet rejects 0
Input DA rejects 0
Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault: OK, Link partner
Speed: 1000 Mbps
Local resolution:
Flow control: Symmetric, Remote fault: Link OK, Local link Speed: 1000 Mbps, Link mode: Full-
duplex
Packet Forwarding Engine configuration:
Destination slot: 0 (0x00)
CoS information:
Direction : Output
CoS transmit queue Bandwidth Buffer Priority Limit
% bps % usec
0 best-effort 95 950000000 95 0 low none
3 network-control 5 50000000 5 0 low none
Interface transmit statistics: Disabled
```

Release Information

Command introduced before Junos OS Release 7.4.

Command modified in Junos OS Release 9.5 for SRX Series Firewalls.

Command modified in Junos OS Release 19.3R1 for MX Series Routers.

Starting in Junos OS Release 19.3R1, Output fields `Ifindex` and `speed` is modified in the `show interfaces` *`interface name`* `extensive` command, on all MX Series routers.

- The default behavior of WAN-PHY interface remains the same.The new `precise-bandwidth` option reflects the new speed (9.294-Gbps) configured on the supported line cards.
- The WAN-PHY framing mode is supported only on MPC5E and MPC6E line cards.

Starting in Junos OS Release 19.3R1, class of service (CoS) features can be configured on the physical interface with speed rates of 1-Gbps, 10-Gbps, 40-Gbps, and 100-Gbps to provide better bandwidth for processing traffic during congestion using variant speeds.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
20.4R1	Starting in Junos OS Release 20.4R1, we support G.fast and Annex J specification with SFP xDSL for ADSL2/ADSL2+ and all VDSL2 profiles on SRX380, SRX300, SRX320, SRX340, and SRX345 devices.
19.2R3	In Junos OS Releases 19.2R3, 19.3R3, 19.4R3, 20.1R2, and 20.2R1, on QFX5120-48Y switch, the `show interfaces` *`interface-name`*`<media><extensive>` command displays the autonegotiation status only for the interface that supports autonegotiation.
18.4R1	Starting in Junos OS Release 18.4R1, Output fields `Next-hop` and `vpls-status` is displayed in the `show interfaces` *`interface name`* `detail` command, only for Layer 2 protocols on MX480 routers.

RELATED DOCUMENTATION

- Understanding Layer 2 Interfaces on Security Devices
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- *Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*
- dsl-sfp-options

show interfaces (Abstracted Fabric)

IN THIS SECTION

- Syntax | 20371
- Description | 20371
- Options | 20372
- Required Privilege Level | 20372
- Output Fields | 20372
- Sample Output | 20378
- Release Information | 20400

Syntax

```
show interfaces af-interface-name
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information for the specified abstracted fabric (af) interface.

NOTE:

- Starting in Junos OS Release 18.3R1, the `show interfaces af-interface-name` output provides transmit statistics of each Packet Forwarding Engine peer list on a given af interface, in addition to the physical interface statistics.
- Starting from Junos OS Release 21.4R1, the `show interface af-name` command does not display the information about the load balancing of transmit traffic on each peer GNF

Forwarding Element (FE). For this information, use the show interfaces lb-stats *af-name* command.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 613 on page 20372describes the output fields for the `show interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 613: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name and status of the physical interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none

Table 613: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Type`	Interface type.	
`Link-level type`	Encapsulation being used on the physical interface.	All levels
`MTU`	Maximum transmission unit size on the physical interface.	All levels
`Speed`	Speed at which the interface is running.	All levels
`Device flags`	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
`Interface flags`	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
`Link type`	Link type. For example, Full-Duplex.	All levels
`Damping`	Damping information.	`detail extensive`
`Alternate link address`	Address of the alternative link.	`detail extensive`
`Hold-times`	Current interface hold-time up and hold-time down, in milliseconds (ms).	`detail extensive`
`Current address`	Configured MAC address.	`detail extensive` none

Table 613: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Hardware address`	Hardware MAC address.	`detail extensive` none
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second`*`:`*`timezone`* `(`*`hour`*`:`*`minute`*`:`*`second`* `ago)`. For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive` none
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the interface if IPv6 statistics tracking is enabled.	`extensive`

Table 613: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input errors`	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Errors`—Sum of the incoming frame terminates and FCS errors. • `Drops`—Number of packets dropped by the input queue. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • `Resource errors`—Sum of transmit drops.	`extensive`
`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. • `Errors`—Sum of the outgoing frame terminates and FCS errors. • `Drops`—Number of packets dropped by the output queue. **NOTE**: • `MTU errors`—Number of packets whose size exceeded the MTU of the interface. • `Resource errors`—Sum of transmit drops.	`extensive`
`Bandwidth`	Shows the interface bandwidth.	`detail extensive` none
Peer GNF id	The GNF peer connected using the abstracted fabric interface.	`detail extensive` none

Table 613: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Peer GNF Forwarding element(FE) view	Shows forwarding element (FE) number and the FPC slot, FE bandwidth, and FE status (up/down).	`detail extensive` none
Collapsed Forwarding Optimized Statistics	This section has the following subsections to display fabric optimization information. • Optimal Packets—shows the packets that are already forwarded to optimal Packet Forwarding Engine based on the abstracted fabric interface load balancing. • Non-optimal Packets—By default, these packets are forwarded to the non-optimal Packet Forwarding Engine by abstracted fabric interface. If you have configured the `optimize` mode using the command `set chassis network-slices guest-network-functions gnf` *`id`* `collapsed-forward optimize`, this section shows the packets that are optimized because of this configuration. If you have configured the `monitor` mode using the command `set chassis network-slices guest-network-functions gnf` *`id`* `collapsed-forward monitor`, this section shows non-optimal packets.	All levels
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Index number of the logical interface, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP interface index number for the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Flags`	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	All levels

Table 613: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Encapsulation`	Encapsulation on the logical interface.	All levels
`Protocol`	Protocol family. Possible values are described in the “Protocol Field” section under *Common Output Fields Description*.	`detail extensive` none
`MTU`	Maximum transmission unit size on the logical interface.	`detail extensive` none
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the specified interface set. • `Input bytes, Output bytes`—Number of bytes received and transmitted on the interface set. • `Input packets, Output packets`—Number of packets received and transmitted on the interface set.	`detail extensive`
`Transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	`extensive`
`Local statistics`	Number and rate of bytes and packets destined to the router.	`extensive`
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Route Table`	Route table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail extensive` none
`Flags`	Information about protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	`detail extensive`
`Addresses, Flags`	Information about the address flags. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	`detail extensive` none

Table 613: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
protocol-family	Protocol family configured on the logical interface. If the protocol is `inet`, the IP address of the interface is also displayed.	`brief`
`Flags`	Information about the address flag. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	`detail extensive` none
`Destination`	IP address of the remote side of the connection.	`detail extensive` none
`Local`	IP address of the logical interface.	`detail extensive` none
`Broadcast`	Broadcast address of the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Sample Output

show interfaces af-interface-name brief

This output is applicable to Junos OS Releases 21.4R1 and later, and versions prior to 18.3R1.

```
user@router> show interfaces af3 brief
Physical interface: af3, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 9212, Clocking: Unspecified, Speed: 1040000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000

  Logical interface af3.0
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 0x8100.200 ]  Encapsulation: ENET2
    inet  34.34.34.1/24
    mpls
    multiservice
```

```
  Logical interface af3.1
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.300 0x8100.400 ]  Encapsulation: ENET2
    inet6 34::1/64
          fe80::290:6901:2c2c:df00/64
    mpls
    multiservice

  Logical interface af3.32767
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    multiservice
```

show interfaces af-interface-name

This output is applicable to Junos OS Releases 21.4R1 and later, and versions prior to 18.3R1.

```
user@router> show interfaces af3
Physical interface: af3, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 931
  Type: Ethernet, Link-level type: Ethernet, MTU: 9212, Speed: 1040000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Current address: 00:90:69:2c:c0:63, Hardware address: 00:90:69:2c:c0:63
  Last flapped   : 2021-10-17 17:24:10 PDT (1w2d 08:14 ago)
  Input rate     : 616 bps (1 pps)
  Output rate    : 1088 bps (0 pps)
  Bandwidth      : 1040 Gbps
  Peer GNF id    : 3
  Collapsed Forwarding : Disabled
  Peer GNF Forwarding element(FE) view :
  FPC slot:FE num  FE Bandwidth(Gbps) Status
       7:0                   130         Up
       7:1                   130         Up
       7:2                   130         Up
       7:3                   130         Up
      17:0                   130         Up
      17:1                   130         Up
      17:2                   130         Up
      17:3                   130         Up
  Collapsed Forwarding Statistics :
```

```
          Optimal Packets            Non-optimal Packets
                        0                                0
          Optimal Bytes              Non-optimal Bytes
                      0                              0

Logical interface af3.0 (Index 375) (SNMP ifIndex 1123)
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 0x8100.200 ]
  Encapsulation: ENET2
  Input packets : 8371
  Output packets: 1683613372
  Protocol inet, MTU: 9190
  Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1,
  Curr new hold cnt: 0, NH drop cnt: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 192.0.2.0/24, Local: 192.0.2.1, Broadcast: 192.0.2.255
  Protocol mpls, MTU: 9178, Maximum labels: 3
    Flags: Is-Primary
  Protocol multiservice, MTU: Unlimited
    Flags: Is-Primary

Logical interface af3.1 (Index 376) (SNMP ifIndex 1125)
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.300 0x8100.400 ]
  Encapsulation: ENET2
  Input packets : 1528
  Output packets: 1524
  Protocol inet6, MTU: 9190
  Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1,
  Curr new hold cnt: 0, NH drop cnt: 0
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 34::/64, Local: 34::1
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::290:6901:2c2c:df00
  Protocol mpls, MTU: 9178, Maximum labels: 3
  Protocol multiservice, MTU: Unlimited

Logical interface af3.32767 (Index 377) (SNMP ifIndex 1124)
  Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
```

```
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

show interfaces af-interface-name detail

This output is applicable to Junos OS Releases 21.4R1 and later, and versions prior to 18.3R1.

```
user@router> show interfaces af3 detail
Physical interface: af3, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 931, Generation: 148
  Type: Ethernet, Link-level type: Ethernet, MTU: 9212, Clocking: Unspecified, Speed: 1040000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 00:90:69:2c:c0:63, Hardware address: 00:90:69:2c:c0:63
  Alternate link address: Unspecified
  Last flapped   : 2021-10-17 17:24:10 PDT (1w2d 08:40 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :     1713156145106484                  280 bps
   Output bytes  :     1704691343752013                 1008 bps
   Input  packets:        3467927904602                    0 pps
   Output packets:        3466096797054                    0 pps
   IPv6 transit statistics:
   Input  bytes  :                 1792
   Output bytes  :                    0
   Input  packets:                   30
   Output packets:                    0
  Bandwidth      : 1040 Gbps
  Peer GNF id    : 3
  Collapsed Forwarding : Disabled
  Peer GNF Forwarding element(FE) view :
  FPC slot:FE num  FE Bandwidth(Gbps) Status
       7:0                   130         Up
       7:1                   130         Up
       7:2                   130         Up
       7:3                   130         Up
```

```
      17:0                  130         Up
      17:1                  130         Up
      17:2                  130         Up
      17:3                  130         Up
  Collapsed Forwarding Statistics :
            Optimal Packets             Non-optimal Packets
                          0                               0
            Optimal Bytes               Non-optimal Bytes
                        0                               0

  Logical interface af3.0 (Index 375) (SNMP ifIndex 1123) (Generation 448)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 0x8100.200 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               673655
     Output bytes  :         824966843167
     Input  packets:                 9811
     Output packets:           1683614943
    Local statistics:
     Input  bytes  :               443858
     Output bytes  :               616248
     Input  packets:                 6217
     Output packets:                 7233
    Transit statistics:
     Input  bytes  :               229797                    8 bps
     Output bytes  :         824966226919                  280 bps
     Input  packets:                 3594                    0 pps
     Output packets:           1683607710                    0 pps
    Protocol inet, MTU: 9190
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
    Generation: 416, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2.0/24, Local: 192.0.2.1, Broadcast: 192.0.2.255, Generation: 302
    Protocol mpls, MTU: 9178, Maximum labels: 3, Generation: 417, Route table: 0
      Flags: Is-Primary
    Protocol multiservice, MTU: Unlimited, Generation: 418, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__

  Logical interface af3.1 (Index 376) (SNMP ifIndex 1125) (Generation 449)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.300 0x8100.400 ]  Encapsulation: ENET2
    Traffic statistics:
```

```
     Input  bytes  :               149864
     Output bytes  :               181586
     Input  packets:                 1780
     Output packets:                 1773
     IPv6 transit statistics:
     Input  bytes  :               149864
     Output bytes  :               181586
     Input  packets:                 1780
     Output packets:                 1773
    Local statistics:
     Input  bytes  :               149672
     Output bytes  :               181586
     Input  packets:                 1777
     Output packets:                 1773
    Transit statistics:
     Input  bytes  :                  192                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    3                    0 pps
     Output packets:                    0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                 192                    0 bps
      Output bytes  :                   0                    0 bps
      Input  packets:                   3                    0 pps
      Output packets:                   0                    0 pps
    Protocol inet6, MTU: 9190
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
    Generation: 419, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 34::/64, Local: 34::1
    Generation: 304
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::290:6901:2c2c:df00
    Protocol mpls, MTU: 9178, Maximum labels: 3, Generation: 306
    Protocol multiservice, MTU: Unlimited, Generation: 420, Route table: 0
    Generation: 421, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface af3.32767 (Index 377) (SNMP ifIndex 1124) (Generation 450)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
```

```
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol multiservice, MTU: Unlimited, Generation: 422, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces af-interface-name extensive

This output is applicable to Junos OS Releases 21.4R1 and later, and versions prior to 18.3R1.

```
user@router> show interfaces af3 extensive
Physical interface: af3, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 931, Generation: 148
  Type: Ethernet, Link-level type: Ethernet, MTU: 9212, Clocking: Unspecified, Speed: 1040000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 00:90:69:2c:c0:63, Hardware address: 00:90:69:2c:c0:63
  Alternate link address: Unspecified
  Last flapped   : 2021-10-17 17:24:10 PDT (1w2d 08:42 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :     1713156145118374                  768 bps
   Output bytes  :     1704691343766497                    0 bps
   Input  packets:        3467927904777                    1 pps
   Output packets:        3466096797232                    0 pps
```

```
    IPv6 transit statistics:
    Input  bytes  :                 1792
    Output bytes  :                    0
    Input  packets:                   30
    Output packets:                    0
   Input errors:
     Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
   Output errors:
     Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
   Bandwidth      : 1040 Gbps
   Peer GNF id    : 3
   Collapsed Forwarding : Disabled
   Peer GNF Forwarding element(FE) view :
   FPC slot:FE num  FE Bandwidth(Gbps) Status
        7:0                   130         Up
        7:1                   130         Up
        7:2                   130         Up
        7:3                   130         Up
       17:0                   130         Up
       17:1                   130         Up
       17:2                   130         Up
       17:3                   130         Up
   Collapsed Forwarding Statistics :
             Optimal Packets             Non-optimal Packets
                           0                                0
             Optimal Bytes               Non-optimal Bytes
                         0                                0

   Logical interface af3.0 (Index 375) (SNMP ifIndex 1123) (Generation 448)
     Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 0x8100.200 ]  Encapsulation: ENET2
     Traffic statistics:
      Input  bytes  :               683507
      Output bytes  :         824966855195
      Input  packets:                 9962
      Output packets:           1683615097
     Local statistics:
      Input  bytes  :               450058
      Output bytes  :               624868
      Input  packets:                 6310
      Output packets:                 7333
     Transit statistics:
      Input  bytes  :               233449                  768 bps
```

```
     Output bytes  :         824966230327                    0 bps
     Input  packets:                 3652                    1 pps
     Output packets:           1683607764                    0 pps
    Protocol inet, MTU: 9190
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
    Generation: 416, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2.0/24, Local: 192.0.2.1, Broadcast: 192.0.2.255, Generation: 302
    Protocol mpls, MTU: 9178, Maximum labels: 3, Generation: 417, Route table: 0
      Flags: Is-Primary
    Protocol multiservice, MTU: Unlimited, Generation: 418, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__

  Logical interface af3.1 (Index 376) (SNMP ifIndex 1125) (Generation 449)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.300 0x8100.400 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               152024
     Output bytes  :               184042
     Input  packets:                 1806
     Output packets:                 1797
     IPv6 transit statistics:
     Input  bytes  :               152024
     Output bytes  :               184042
     Input  packets:                 1806
     Output packets:                 1797
    Local statistics:
     Input  bytes  :               151832
     Output bytes  :               184042
     Input  packets:                 1803
     Output packets:                 1797
    Transit statistics:
     Input  bytes  :                  192                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    3                    0 pps
     Output packets:                    0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                 192                    0 bps
      Output bytes  :                   0                    0 bps
      Input  packets:                   3                    0 pps
      Output packets:                   0                    0 pps
```

```
    Protocol inet6, MTU: 9190
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
    Generation: 419, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 34::/64, Local: 34::1
    Generation: 304
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::290:6901:2c2c:df00
    Protocol mpls, MTU: 9178, Maximum labels: 3, Generation: 306
    Protocol multiservice, MTU: Unlimited, Generation: 420, Route table: 0
    Generation: 421, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface af3.32767 (Index 377) (SNMP ifIndex 1124) (Generation 450)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol multiservice, MTU: Unlimited, Generation: 422, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces af-interface-name

This output is applicable to Junos OS Releases starting from 18.3R1 to 21.3R1.

```
user@router> show interfaces af9
Physical interface: af9, Enabled, Physical link is Up
  Interface index: 209, SNMP ifIndex: 527
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 370000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 00:90:69:2b:00:4c, Hardware address: 00:90:69:2b:00:4c
  Last flapped   : 2018-09-12 01:44:01 PDT (00:01:02 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Bandwidth      : 370 Gbps
  Peer GNF id    : 9
  Peer GNF Forwarding element(FE) view :
  FPC slot:FE num  FE Bandwidth(Gbps) Status      Transmit Packets         Transmit Bytes
       6:0                   130         Up                      0                      0
      12:0                   120         Up                      0                      0
      12:1                   120         Up                      0                      0

  Residual Transmit Statistics :
  Packets :                    0 Bytes :                    0

  Fabric Queue Statistics :
  FPC slot:FE num    High priority(pkts)        Low priority(pkts)
       6:0                            0                          0
      12:0                            0                          0
      12:1                            0                          0
  FPC slot:FE num    High priority(bytes)      Low priority(bytes)
       6:0                              0                        0
      12:0                              0                        0
      12:1                              0                        0
  Residual Queue Statistics :
      High priority(pkts)       Low priority(pkts)
                       0                        0
      High priority(bytes)      Low priority(bytes)
                        0                        0
```

```
  Logical interface af9.0 (Index 332) (SNMP ifIndex 528)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 0
    Output packets: 13
    Protocol inet, MTU: 1500
```

show interfaces af-interface-name detail

This output is applicable to Junos OS Releases starting from 18.3R1 to 21.3R1.

```
user@router> show interfaces af2 detail

Physical interface: af2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 674, Generation: 349
  Type: Ethernet, Link-level type: Ethernet, MTU: 1518, Clocking: Unspecified, Speed: 130000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 00:90:69:fd:85:a4, Hardware address: 00:90:69:fd:85:a4
  Alternate link address: Unspecified
  Last flapped   : 2018-11-01 20:44:26 PDT (6d 02:57 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :             23381827                 280 bps
   Output bytes  :            652664393                 416 bps
   Input  packets:               341618                   0 pps
   Output packets:              5986312                   0 pps
   IPv6 transit statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Bandwidth      : 130 Gbps
  Peer GNF id    : 2
  Peer GNF Forwarding element(FE) view :
  FPC slot:FE num  FE Bandwidth(Gbps) Status      Transmit Packets         Transmit Bytes
       2:0                   130         Up                5656951              622264610
```

```
  Residual Transmit Statistics :
  Packets :                    0 Bytes :                    0

  Fabric Queue Statistics :
  FPC slot:FE num    High priority(pkts)        Low priority(pkts)
       2:0                      5656951                         0
  FPC slot:FE num    High priority(bytes)      Low priority(bytes)
       2:0                      622264610                        0
  Residual Queue Statistics :
      High priority(pkts)       Low priority(pkts)
                        0                        0
      High priority(bytes)      Low priority(bytes)
                         0                        0

  Logical interface af2.0 (Index 334) (SNMP ifIndex 647) (Generation 234)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             23382038
     Output bytes  :            650688227
     Input  packets:               341621
     Output packets:              5986312
    Local statistics:
     Input  bytes  :             23381827
     Output bytes  :             28423617
     Input  packets:               341618
     Output packets:               329361
    Transit statistics:
     Input  bytes  :                  211                     0 bps
     Output bytes  :            622264610                     0 bps
     Input  packets:                    3                     0 pps
     Output packets:              5656951                     0 pps
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
    Generation: 314, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2.0/24, Local: 192.0.2.1, Broadcast: 192.0.2.255, Generation: 224
    Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 315, Route table: 0
      Flags: Is-Primary
    Protocol multiservice, MTU: Unlimited, Generation: 316, Route table: 0
      Flags: Is-Primary
```

```
      Policer: Input: __default_arp_policer__

  Logical interface af2.1 (Index 336) (SNMP ifIndex 649) (Generation 235)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop
cnt: 0
    Generation: 317, Route table: 0
      Flags: Sendbcast-pkt-to-re
    Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 318, Route table: 0
    Protocol multiservice, MTU: Unlimited, Generation: 319, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface af2.32767 (Index 337) (SNMP ifIndex 675) (Generation 236)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
```

```
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol multiservice, MTU: Unlimited, Generation: 320, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces af-interface-name extensive

This output is applicable to Junos OS Releases starting from 18.3R1 to 21.3R1.

```
user@router> show interfaces af2 extensive

Physical interface: af2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 674, Generation: 349
  Type: Ethernet, Link-level type: Ethernet, MTU: 1518, Clocking: Unspecified, Speed: 130000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 00:90:69:fd:85:a4, Hardware address: 00:90:69:fd:85:a4
  Alternate link address: Unspecified
  Last flapped   : 2018-11-01 20:44:26 PDT (6d 02:57 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :             23382976                    0 bps
   Output bytes  :            652665950                    0 bps
   Input  packets:               341635                    0 pps
   Output packets:              5986329                    0 pps
   IPv6 transit statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
```

```
  Bandwidth      : 130 Gbps
  Peer GNF id    : 2
  Peer GNF Forwarding element(FE) view :
  FPC slot:FE num  FE Bandwidth(Gbps) Status      Transmit Packets         Transmit Bytes
       2:0                   130         Up                5656951              622264610

  Residual Transmit Statistics :
  Packets :                    0 Bytes :                    0

  Fabric Queue Statistics :
  FPC slot:FE num    High priority(pkts)        Low priority(pkts)
       2:0                      5656951                        0
  FPC slot:FE num    High priority(bytes)      Low priority(bytes)
       2:0                      622264610                       0
  Residual Queue Statistics :
      High priority(pkts)       Low priority(pkts)
                        0                        0
      High priority(bytes)      Low priority(bytes)
                         0                        0

  Logical interface af2.0 (Index 334) (SNMP ifIndex 647) (Generation 234)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             23383187
     Output bytes  :            650689682
     Input  packets:               341638
     Output packets:              5986329
    Local statistics:
     Input  bytes  :             23382976
     Output bytes  :             28425072
     Input  packets:               341635
     Output packets:               329378
    Transit statistics:
     Input  bytes  :                  211                    0 bps
     Output bytes  :            622264610                    0 bps
     Input  packets:                    3                    0 pps
     Output packets:              5656951                    0 pps
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
    Generation: 314, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
```

```
        Destination: 192.0.2.0/24, Local: 192.0.2.1, Broadcast: 192.0.2.255, Generation: 224
    Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 315, Route table: 0
      Flags: Is-Primary
    Protocol multiservice, MTU: Unlimited, Generation: 316, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__

  Logical interface af2.1 (Index 336) (SNMP ifIndex 649) (Generation 235)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop
cnt: 0
    Generation: 317, Route table: 0
      Flags: Sendbcast-pkt-to-re
    Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 318, Route table: 0
    Protocol multiservice, MTU: Unlimited, Generation: 319, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface af2.32767 (Index 337) (SNMP ifIndex 675) (Generation 236)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
```

```
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol multiservice, MTU: Unlimited, Generation: 320, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces af-interface-name statistics

This output is applicable to Junos OS Releases starting from 18.3R1 to 21.3R1.

```
user@router> show interfaces af4 statistics

Physical interface: af4, Enabled, Physical link is Up
  Interface index: 165, SNMP ifIndex: 958
  Type: Ethernet, Link-level type: Flexible-Ethernet, MTU: 1522, Speed: 240000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Current address: 00:90:69:c7:31:62, Hardware address: 00:90:69:c7:31:62
  Last flapped   : 2018-08-07 21:47:10 PDT (00:58:48 ago)
  Statistics last cleared: Never
  Input rate     : 1523826080 bps (2976222 pps)
  Output rate    : 112168 bps (232 pps)
  Input errors: 0, Output errors: 0
  Bandwidth      : 240 Gbps
  Peer GNF id    : 4
  Peer GNF Forwarding element(FE) view :
  FPC slot:FE num  FE Bandwidth(Gbps) Status      Transmit Packets         Transmit Bytes
      12:0                   120         Up                   2802                 168120
      12:1                   120         Up                   2848                 170880

  Residual Transmit Statistics :
  Packets :                    0 Bytes :                    0

  Fabric Queue Statistics :
  FPC slot:FE num    High priority(pkts)        Low priority(pkts)
```

```
      12:0                        2802                        0
      12:1                        2848                        0
  FPC slot:FE num    High priority(bytes)      Low priority(bytes)
      12:0                        168120                       0
      12:1                        170880                       0
  Residual Queue Statistics :
      High priority(pkts)       Low priority(pkts)
                       0                        0
      High priority(bytes)      Low priority(bytes)
                        0                        0

  Logical interface af4.1 (Index 335) (SNMP ifIndex 7354)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]  Encapsulation: ENET2
    Input packets : 71572652
    Output packets: 5740
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2.0/24, Local: 192.0.2.1, Broadcast: 192.0.2.255
    Protocol mpls, MTU: 1488, Maximum labels: 3
    Protocol multiservice, MTU: Unlimited

  Logical interface af4.32767 (Index 336) (SNMP ifIndex 7355)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

show interfaces af2 (with collapsed forwarding configured)

```
user@router> show interfaces af2

Physical interface: af2, Enabled, Physical link is Up
  Interface index: 152, SNMP ifIndex: 626
  Type: Ethernet, Link-level type: Ethernet, MTU: 1518, Speed: 1600000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
```

```
Current address: 00:90:69:00:c0:4a, Hardware address: 00:90:69:00:c0:4a
Last flapped   : 2019-10-01 06:19:50 PDT (2d 00:54 ago)
Input rate     : 312 bps (0 pps)
Output rate    : 144327758752 bps (18371659 pps)
Bandwidth      : 1600 Gbps
Peer GNF id    : 7
Collapsed Forwarding : Optimized
Peer GNF Forwarding element(FE) view :
FPC slot:FE num  FE Bandwidth(Gbps) Status      Transmit Packets         Transmit Bytes
     1:0                   400         Up           111437126843        110991378335628
     1:1                   400         Up                      0                      0
     1:2                   400         Up                      0                      0
     1:3                   400         Up                      0                      0

Residual Transmit Statistics :
Packets :                    0 Bytes :                    0

Fabric Queue Statistics :
FPC slot:FE num    High priority(pkts)        Low priority(pkts)
     1:0                            0              111437126843
     1:1                            0                         0
     1:2                            0                         0
     1:3                            0                         0
FPC slot:FE num    High priority(bytes)      Low priority(bytes)
     1:0                              0          110991378335628
     1:1                              0                        0
     1:2                              0                        0
     1:3                              0                        0
Residual Queue Statistics :
    High priority(pkts)       Low priority(pkts)
                      0                        0
    High priority(bytes)      Low priority(bytes)
                       0                        0
Collapsed Forwarding Optimized Statistics :
          Optimal Packets             Non-optimal Packets
                        0                    111437126843
          Optimal Bytes               Non-optimal Bytes
                       0                  110991378335628

Logical interface af2.0 (Index 342) (SNMP ifIndex 689)
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
  Input packets : 532
  Output packets: 41761133660
```

```
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2.0/24, Local: 192.0.2.1, Broadcast: 192.0.2.255
    Protocol multiservice, MTU: Unlimited

  Logical interface af2.1 (Index 341) (SNMP ifIndex 690)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    Input packets : 536
    Output packets: 41764698890
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2.0/24, Local: 192.0.2.1, Broadcast: 192.0.2.1
    Protocol multiservice, MTU: Unlimited

  Logical interface af2.32767 (Index 343) (SNMP ifIndex 692)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

show interfaces af2 (with sub line cards or SLCs configured on MX2K-MPC11E)

```
user@router> show interfaces af2

Physical interface: af2, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 501
  Type: Ethernet, Link-level type: Ethernet, MTU: 1518, Speed: 3600000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Current address: 00:90:69:1d:20:47, Hardware address: 00:90:69:1d:20:47
  Last flapped   : 2021-03-02 22:40:47 PST (4d 21:31 ago)
  Input rate     : 4128 bps (9 pps)
  Output rate    : 5248 bps (12 pps)
```

```
Bandwidth      : 3600 Gbps
Peer GNF id    : 2
Collapsed Forwarding : Optimized
Peer GNF Forwarding element(FE) view :
FPC slot:FE num  FE Bandwidth(Gbps) Status      Transmit Packets         Transmit Bytes
     0:0                   400         Up                3927296              204218108
     0:1                   400         Up                 911595               47346520
     0:2                   400         Up                    176                   8416
     0:3                   400         Up                  48627                4771730
     2:4                   500         Up                      0                      0
     2:5                   500         Up                      0                      0
     2:6                   500         Up                      0                      0
     2:7                   500         Up                      0                      0

Residual Transmit Statistics :
Packets :              4894889 Bytes :            256741776

Fabric Queue Statistics :
FPC slot:FE num    High priority(pkts)        Low priority(pkts)
     0:0                            0                   3927296
     0:1                            0                    911595
     0:2                            0                       176
     0:3                            0                     48627
     2:4                            0                         0
     2:5                            0                         0
     2:6                            0                         0
     2:7                            0                         0
FPC slot:FE num    High priority(bytes)      Low priority(bytes)
     0:0                              0                204218108
     0:1                              0                 47346520
     0:2                              0                     8416
     0:3                              0                 4771730
     2:4                              0                       0
     2:5                              0                       0
     2:6                              0                       0
     2:7                              0                       0
Residual Queue Statistics :
    High priority(pkts)       Low priority(pkts)
                      0                  4894889
    High priority(bytes)      Low priority(bytes)
                       0                256741776
Collapsed Forwarding Optimized Statistics :
          Optimal Packets             Non-optimal Packets
```

```
                     960612                        3927082
            Optimal Bytes               Non-optimal Bytes
                 52136510                      204208264

  Logical interface af2.0 (Index 361) (SNMP ifIndex 1032)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Input packets : 8659876
    Output packets: 9897785
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2.0/24, Local: 192.0.2.1, Broadcast: 192.0.2.255
    Protocol iso, MTU: 1497
      Flags: Is-Primary
      Addresses, Flags: Is-Primary
        Local: 49.0002.0192.0168.0001
    Protocol mpls, MTU: 1488, Maximum labels: 3
      Flags: Is-Primary
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary

  Logical interface af2.32767 (Index 364) (SNMP ifIndex 1033)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Setting Up Junos Node Slicing

show interfaces (Adaptive Services)

IN THIS SECTION

- Syntax | 20401
- Description | 20401
- Options | 20401
- Required Privilege Level | 20402
- Output Fields | 20402
- Sample Output | 20408
- Release Information | 20411

Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified adaptive services interface.

Options

interface-type	On M Series and T Series routers, the interface type is **sp-** *fpc/pic/port*.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.

snmp-index ***snmp-index***	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 614 on page 20402 lists the output fields for the `show interfaces` (adaptive services and redundant adaptive services) command. Output fields are listed in the approximate order in which they appear.

Table 614: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields

Field Name	Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Fields Description*.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the physical interface.
Generation	Unique number for use by Juniper Networks technical support only.
Type	Encapsulation being used on the interface.
Link-level type	Encapsulation being used on the physical interface.

Table 614: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields *(Continued)*

Field Name	Field Description
MTU	MTU size on the physical interface.
Clocking	Reference clock source: can be **Internal** or **External**.
Speed	Speed at which the interface is running.
Device flags	Information about the physical device. Possible values are described in the “Device Flags *Common Output Fields Description*.
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” se *Common Output Fields Description*.
Link type	Physical interface link type: **Full-Duplex** or **Half-Duplex**.
Link flags	Information about the link. Possible values are described in the “Link Flags” section unde *Output Fields Description*.
Physical info	Information about the physical interface.
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
Current address	Configured MAC address.
Hardware address	MAC address of the hardware.
Alternate link address	Backup address of the link.

Table 614: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields *(Continued)*

Field Name	Field Description
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last fla** ***month-day hour.minute.second timezone*** **(*hour.minute.second* ago)**. For example, **Last fl** **2002-04-26 10:52:40 PDT (04:33:20 ago)**.
Input Rate	Input rate in bits per second (bps) and packets per second (pps).
Output Rate	Output rate in bps and pps.
Statistics last cleared	Time when the statistics for the interface were last set to zero.
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interfac **NOTE**: With static NAT configured as basic NAT44 or destination NAT44 on MX Seri MS-MICs and MS-MPCs, the Input bytes field might show 16 more bytes than the Ou This is caused by the accounting of 16 bytes of the Juniper Forwarding Module cooki • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.

Table 614: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields *(Continued)*

Field Name	Field Description
Input errors	Input errors on the interface. The following paragraphs explain the counters whose mea obvious: • **Errors**—Sum of the incoming frame terminates and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If saturated, this number increments once for every packet that is dropped by the ASIC mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Frames received smaller than the runt threshold. • **Giants**—Frames received larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded were not recognized or not of interest. Usually, this field reports protocols that the Ju handle. • **Resource errors**—Sum of transmit drops.
Output errors	Output errors on the interface. The following paragraphs explain the counters whose me be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This n normally increment quickly, increasing only when the cable is unplugged, the far-end powered down and then up, or another problem occurs. If the number of carrier tran quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is • **Errors**—Sum of the outgoing frame terminates and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. I saturated, this number increments once for every packet that is dropped by the ASIC mechanism. • **MTU errors**—Number of packets larger than the MTU threshold. • **Resource errors**—Sum of transmit drops.
Logical Interface	

Table 614: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields *(Continued)*

Field Name	Field Description
Logical interface	Name of the logical interface.
Index	Logical interface index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP interface index number.
Generation	Unique number for use by Juniper Networks technical support only.
Flags	Information about the logical interface. Possible values are described in the "Logical Inte section under *Common Output Fields Description*.
Encapsulation	Encapsulation on the logical interface.
Input packets	Number of packets received on the logical interface.
Output packets	Number of packets transmitted on the logical interface.
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst received, the value in the output packet rate field might briefly exceed the peak cell rate (generally, less than 1 second) for this counter to stabilize.

Table 614: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields *(Continued)*

Field Name	Field Description
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in rate field might briefly exceed the peak cell rate. It takes generally less than 1 second for stabilize.
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address also displayed.
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, **mpls**.
MTU	MTU size on the logical interface.
Generation	Unique number for use by Juniper Networks technical support only.
Route table	Routing table in which the logical interface address is located. For example, **0** refers to th inet.0.
Flags	Information about the protocol family flags. Possible values are described in the "Family under *Common Output Fields Description*.
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flag *Common Output Fields Description*.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

Table 614: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields *(Continued)*

Field Name	Field Description
Generation	Unique number for use by Juniper Networks technical support only.

Sample Output

show interfaces (Adaptive Services)

```
user@host> show interfaces sp-1/2/0
Physical interface: sp-1/2/0, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 72
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : 2006-03-06 11:37:18 PST (00:57:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
    Input packets : 3057
    Output packets: 3044
    Protocol inet, MTU: 9192
      Flags: Receive-options, Receive-TTL-Exceeded
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.0.34, Local: 10.0.0.1
```

show interfaces brief (Adaptive Services)

```
user@host> show interfaces sp-1/2/0 brief
Physical interface: sp-1/2/0, Enabled, Physical link is Up
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
```

```
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000

  Logical interface sp-1/2/0.16383
    Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
    inet  10.0.0.1         --> 10.0.0.34
```

show interfaces detail (Adaptive Services)

```
user@host> show interfaces sp-1/2/0 detail
Physical interface: sp-1/2/0, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 72, Generation: 30
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2006-03-06 11:37:18 PST (00:57:56 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               125147                    0 bps
   Output bytes  :              1483113                    0 bps
   Input  packets:                 3061                    0 pps
   Output packets:                 3048                    0 pps

  Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73) (Generation 7)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
    Traffic statistics:
     Input  bytes  :               125147
     Output bytes  :              1483113
     Input  packets:                 3061
     Output packets:                 3048
    Local statistics:
     Input  bytes  :               125147
     Output bytes  :              1483113
     Input  packets:                 3061
```

```
     Output packets:                 3048
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 9192, Generation: 20, Route table: 1
      Flags: Receive-options, Receive-TTL-Exceeded
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.0.34, Local: 10.0.0.1, Broadcast: Unspecified,
        Generation: 22
```

show interfaces extensive (Adaptive Services)

```
user@host> show interfaces sp-1/2/0 extensive
Physical interface: sp-1/2/0, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 72, Generation: 30
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2006-03-06 11:37:18 PST (00:58:40 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               125547                    0 bps
   Output bytes  :              1483353                    0 bps
   Input  packets:                 3065                    0 pps
   Output packets:                 3052                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0
```

```
  Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73) (Generation 7)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
    Traffic statistics:
     Input  bytes  :               125547
     Output bytes  :              1483353
     Input  packets:                 3065
     Output packets:                 3052
    Local statistics:
     Input  bytes  :               125547
     Output bytes  :              1483353
     Input  packets:                 3065
     Output packets:                 3052
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 9192, Generation: 20, Route table: 1
      Flags: Receive-options, Receive-TTL-Exceeded
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.0.34, Local: 10.0.0.1, Broadcast: Unspecified,
        Generation: 22
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces (Aggregated Ethernet)

IN THIS SECTION

- Syntax | 20412
- Description | 20412
- Options | 20412
- Required Privilege Level | 20413
- Output Fields | 20413

Sample Output | 20423
Release Information | 20437

Syntax

```
show interfaces ae number
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M Series, T Series, MX Series, and PTX Series routers) Display status information about the specified aggregated Ethernet interfaces.

Options

ae*number*	Display standard information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

NOTE: On Junos OS Evolved, in untagged aggregated ethernet (ae) interfaces with no logical interface configuration, the ae interface will not be shown as "down" and the

speed will not be shown as "unspecified." The speed will be the aggregate speed of all the child member interfaces which are "up." In Junos OS, the speed is shown as "unspecified" in this case.

Required Privilege Level

view

Output Fields

Table 615 on page 20413 lists the output fields for the `show interfaces` (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 615: Aggregated Ethernet show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical interface`	Name of the physical interface and state of the interface.	All levels
`Enabled`	State of the physical interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
`Interface index`	Index number of the physical interface, which reflects its initialization sequence.	All levels
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Link-level type`	Encapsulation being used on the physical interface.	All levels
`MTU`	Maximum transmission unit size on the physical interface.	All levels

Table 615: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Speed`	Speed at which the interface is running.	All levels
`Loopback`	Loopback status: `Enabled` or `Disabled`. If loopback is enabled, type of loopback: `Local` or `Remote`.	All levels
`Source filtering`	Source filtering status: `Enabled` or `Disabled`.	All levels
`Flow control`	Flow control status: `Enabled` or `Disabled`.	All levels
`Minimum links needed`	Number of child links that must be operational for the aggregate interface to be operational.	All levels
`Device flags`	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
`Interface flags`	Information about the interface. Possible values are described in the "Interfaces Flags" section under *Common Output Fields Description*.	All levels
`Current address`	Configured MAC address.	`detail extensive`
`Hardware address`	Hardware MAC address.	`detail extensive`
`Last flapped`	Date, time, and how long ago the interface went from down to up or from up to down. The format is `Last flapped:` *`year-month-day hours:minutes:seconds timezone`* (*`hours:minutes:seconds`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive`
`Input Rate`	Input rate in bits per second (bps) and packets per second (pps).	None specified
`Output Rate`	Output rate in bps and pps.	None specified

Table 615: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes and rate, in bps, at which bytes are received on the interface. • `Output bytes`—Number of bytes and rate, in bps, at which bytes are transmitted on the interface. • `Input packets`—Number of packets and rate, in pps, at which packets are received on the interface. • `Output packets`—Number of packets and rate, in pps, at which packets are transmitted on the interface.	`detail extensive`

Table 615: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input errors`	Input errors on the interface: • `Errors`—Sum of incoming frame terminates and frame check sequence (FCS) errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's random early detection (RED) mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Giants`—Number of frames received that are larger than the giant threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos OS does not handle. • `Resource errors`—Sum of transmit drops.	`detail extensive`

Table 615: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions` —Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), then the cable, the far-end system, or the PIC is malfunctioning. • `Errors`—Sum of the outgoing frame terminates and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `MTU errors`—Number of packets whose size exceeded the MTU of the interface. • `Resource errors`—Sum of transmit drops.	`detail extensive`
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`
`MAC statistics`	Display media access control (MAC) statistics. • `Broadcast packets`—Number of broadcast packets received and transmitted. • `Multicast packets`—Number of multicast packets received and transmitted.	`detail extensive`

Table 615: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Queue counters`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism. **NOTE**: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the `show interfaces` command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	`detail extensive`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Index number of the logical interface (which reflects its initialization sequence).	`detail extensive` none
`SNMP ifIndex`	SNMP interface index number of the logical interface.	`detail extensive` none

Table 615: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Flags`	Information about the logical interface. Possible values are described in the “Logical Interface Flags Field” section under *Common Output Fields Description*.	All levels
`VLAN-Tag`	Tag Protocol Identifier (TPID) and VLAN identifier.	All levels
`Demux`	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • `Source Family Inet` • `Destination Family Inet`	`detail extensive` none
`Encapsulation`	Encapsulation on the logical interface.	All levels

Table 615: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Statistics`	Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface.
	• `Bundle`—Information about input and output bundle rates. For Junos OS Evolved, LACP packets and LAG links on the members of an aggregated Ethernet interface are not counted as part of the aggregated Ethernet bundle input or output statistics.
	• `Link`—(`detail` and `extensive` only) Information about specific links in the aggregate, including link state and input and output rates.
	• `Adaptive Statistics`—(`extensive` only) Information about adaptive load balancing counter statistics.
	  • `Adaptive Adjusts`—Number of times traffic flow imbalance was corrected by implementation of adaptive load balancing.
	  • `Adaptive Scans`—Number of times the link utilization on each member link of the AE bundle was scanned by for adaptive load balancing
	  • `Adaptive Tolerance`—Tolerance level, in percentage, for load imbalance on link utilization on each member link of the AE bundle.
	  • `Adaptive Updates`—Number of times traffic flow loads have been updated on an AE bundle.
	• `Marker Statistics`—(`detail` and `extensive` only) Information about 802.3ad marker protocol statistics on the specified links.
	  • `Marker Rx`—Number of valid marker protocol data units (PDUs) received on this aggregation port.
	  • `Resp Tx`—Number of marker response PDUs transmitted on this aggregation port.
	  • `Unknown Rx`—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown PDU, or are addressed to the slow protocols group
		`detail extensive`
		none

Table 615: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	MAC address (43B.3) but do not carry the slow protocols Ethernet type. • `Illegal Rx`—Number of frames received that carry the slow protocols Ethernet type value (43B.4) but contain a badly formed PDU or an illegal value of protocol subtype (43B.4).	

LACP info	Link Aggregation Control Protocol (LACP) information for each aggregated interface.
	• `Role` can be one of the following:
	  • `Actor`—Local device participating in LACP negotiation.
	  • `Partner`—Remote device participating in LACP negotiation.
	• `System priority`—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer.
	• `System identifier`—Actor or partner system ID, encoded as a MAC address.
	• `Port priority`—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer.
	• `Unknown Rx`—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown protocol data unit (PDU), or are addressed to the slow protocols group MAC address (43B.3) but do not carry the slow protocols Ethernet type.
	• `Port key`—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer.	

Table 615: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
LACP Statistics	LACP statistics for each aggregated interface. • `LACP Rx`—LACP received counter that increments for each normal hello. • `LACP Tx`—Number of LACP transmit packet errors logged. • `Unknown Rx`—Number of unrecognized packet errors logged. • `Illegal Rx`—Number of invalid packets received. **NOTE**: For `LACP Rx` and `LACP Tx`, Packet count is updated only on snmp timer expiry (30 secs).	`detail extensive` none
`protocol-family`	Protocol family configured on the logical interface. Possible values are described in the “Protocol Field” section under *Common Output Fields Description*.	`brief`
`Protocol`	Protocol family configured on the logical interface. Possible values are described in the “Protocol Field” section under *Common Output Fields Description*.	`detail extensive` none
`MTU`	Maximum transmission unit size on the logical interface.	`detail extensive` none
`Maximum labels`	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Route Table`	Routing table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail extensive`
`Flags`	Information about protocol family flags. Possible values are described in the “Family Flags Field” section under *Common Output Fields Description*.	`detail extensive` none

Table 615: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Mac-Validate Failures`	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	`detail extensive` none
`Addresses, Flags`	Information about address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	`detail extensive` none
`Destination`	IP address of the remote side of the connection.	`detail extensive` none
`Local`	IP address of the logical interface.	`detail extensive` none
`Broadcast`	Broadcast address of the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Sample Output

show interfaces (Aggregated Ethernet)

```
user@host> show interfaces ae0
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 59
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Current address: 00:00:5e:00:53:f0, Hardware address: 00:00:5e:00:53:f0
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
```

```
  Logical interface ae0.0 (Index 72) (SNMP ifIndex 60)
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 10.100.1.255
```

show interfaces brief (Aggregated Ethernet)

```
user@host> show interfaces ae0 brief
Physical interface: ae0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384

  Logical interface ae0.0
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    inet  203.0.113.2/24
```

show interfaces detail (Aggregated Ethernet)

```
user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 59, Generation: 36
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Current address: 00:00:5e:00:53:f0, Hardware address: 00:00:5e:00:53:f0
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
```

```
   Output packets:                    0                    0 pps
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                 7375                 7375                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                2268                 2268                    0

  Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Link:
      fe-0/1/0.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      fe-0/1/2.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      fe-0/1/3.0
        Input :             0          0             0            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      fe-0/1/0.0                 0           0            0            0
      fe-0/1/2.0                 0           0            0            0
      fe-0/1/3.0                 0           0            0            0
    Protocol inet, MTU: 1500, Generation: 37, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
        Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
        Generation: 49
```

show interfaces extensive (Aggregated Ethernet)

```
user@host> show interfaces ae0 extensive
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 304, SNMP ifIndex: 830, Generation: 342
  Link-level type: Flexible-Ethernet, MTU: 16000, Speed: 200Gbps, BPDU Error: None, MAC-REWRITE
Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
```

```
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:00:ab:cd:e1, Hardware address: f6:cc:55:34:af:8f
  Last flapped   : 2019-11-04 14:09:06 IST (00:21:46 ago)
  Statistics last cleared: 2019-11-04 14:30:46 IST (00:00:06 ago)
  Traffic statistics:
   Input  bytes  :          12955595850          16067340264 bps
   Output bytes  :          12955839144          16067696488 bps
   Input  packets:             12904045              2000425 pps
   Output packets:             12904281              2000469 pps
   IPv6 transit statistics:
   Input  bytes  :          12955586722
   Output bytes  :          12955830836
   Input  packets:             12903971
   Output packets:             12904214
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  MAC statistics:                      Receive         Transmit
    Broadcast packets                        0                0
    Multicast packets                       11               11
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                        12623539             12623539                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                              67                   67                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control

  Logical interface ae1.1501 (Index 405) (SNMP ifIndex 982) (Generation 12278)
```

```
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1501 ]  Encapsulation: ENET2
Statistics        Packets        pps         Bytes          bps
Bundle:
    Input :      12903972    2000413   12955586798  16067363080
    Output:      12904214    2000456   12955830836  16067715624
Adaptive Statistics:
    Adaptive Adjusts:          0
    Adaptive Scans  :          0
    Adaptive Updates:          0
Link:
  et-8/7/4.1501
    Input :       6453558     999544    6479365950   8028338176
    Output:       5419689     840489    5441367750   6750815312
  xe-8/0/0:0.1501
    Input :        630823      97367     633346374    782052320
    Output:        946439     146652     950224754   1177910328
  xe-8/0/0:1.1501
    Input :        681646     105969     684374260    851151032
    Output:        860269     133412     863710070   1071571528
  xe-8/0/0:2.1501
    Input :        629697      97558     632215406    783587288
    Output:        774250     120067     777347000    964383592
  xe-8/0/0:3.1501
    Input :        630012      98027     632532324    787360080
    Output:       1032345     160054    1036474378   1285560688
  xe-8/6/0:0.1501
    Input :        654743     101651     657362010    816468784
    Output:        473122      73326     475100508    589061512
  xe-8/6/0:1.1501
    Input :        630360      97905     632883672    786375112
    Output:        860240     133284     863680960   1070540592
  xe-8/6/0:2.1501
    Input :        654850     101450     657469796    814853936
    Output:        688226     106642     690978904    856554888
  xe-8/7/0:1.1501
    Input :        654370     101442     656986660    814784400
    Output:        516173      79971     518237690    642332976
  xe-8/7/0:2.1501
    Input :        655771     101940     658394074    818781864
    Output:       1075371     166587    1079586458   1337921528
  xe-8/7/0:3.1501
    Input :        628142      97560     630656272    783610088
    Output:        258090      39972     259122364    321062680
```

```
    Aggregate member links: 11

    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-8/7/4.1501              0           0            0            0
      xe-8/0/0:0.1501            0           0            0            0
      xe-8/0/0:1.1501            0           0            0            0
      xe-8/0/0:2.1501            0           0            0            0
      xe-8/0/0:3.1501            0           0            0            0
      xe-8/6/0:0.1501            0           0            0            0
      xe-8/6/0:1.1501            0           0            0            0
      xe-8/6/0:2.1501            0           0            0            0
      xe-8/7/0:1.1501            0           0            0            0
      xe-8/7/0:2.1501            0           0            0            0
      xe-8/7/0:3.1501            0           0            0            0
    Protocol inet, MTU: 15978
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
    Generation: 36373, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 172.161.1/24, Local: 172.161.1.1, Broadcast: 172.161.1.255, Generation:
72255
    Protocol inet6, MTU: 15978
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 2, Curr new hold cnt: 0, NH drop
cnt: 0
    Generation: 36374, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2005:1:161:1::/64, Local: 2005:1:161:1::1
    Generation: 72257
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::f6cc:5505:dd34:af8f
    Protocol multiservice, MTU: Unlimited, Generation: 72259
    Generation: 36375, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae1.32767 (Index 406) (SNMP ifIndex 985) (Generation 12279)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :            73          1          9052        10488
        Output:            67          0          8308        10032
```

```
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      et-8/7/4.32767
        Input :             7          1           868         1368
        Output:             7          0           868          912
      xe-8/0/0:0.32767
        Input :             6          0           744          912
        Output:             6          0           744          912
      xe-8/0/0:1.32767
        Input :             6          0           744          912
        Output:             6          0           744          912
      xe-8/0/0:2.32767
        Input :             7          0           868          912
        Output:             6          0           744          912
      xe-8/0/0:3.32767
        Input :             7          0           868          912
        Output:             6          0           744          912
      xe-8/6/0:0.32767
        Input :             7          0           868          912
        Output:             6          0           744          912
      xe-8/6/0:1.32767
        Input :             7          0           868          912
        Output:             6          0           744          912
      xe-8/6/0:2.32767
        Input :             7          0           868          912
        Output:             6          0           744          912
      xe-8/7/0:1.32767
        Input :             7          0           868          912
        Output:             6          0           744          912
      xe-8/7/0:2.32767
        Input :             6          0           744          912
        Output:             6          0           744          912
      xe-8/7/0:3.32767
        Input :             6          0           744          912
        Output:             6          0           744          912
    LACP info:        Role     System             System       Port     Port    Port
                             priority         identifier   priority   number     key
      et-8/7/4.32767    Actor      127 f4:cc:55:34:b7:40        127       26       2
      et-8/7/4.32767  Partner      127 d8:b1:22:f2:57:c0        127        8       2
      xe-8/0/0:0.32767    Actor      127 f4:cc:55:34:b7:40      127        3       2
```

```
    xe-8/0/0:0.32767  Partner      127 d8:b1:22:f2:57:c0      127       16        2
    xe-8/0/0:1.32767    Actor      127 f4:cc:55:34:b7:40      127        4        2
    xe-8/0/0:1.32767  Partner      127 d8:b1:22:f2:57:c0      127       17        2
    xe-8/0/0:2.32767    Actor      127 f4:cc:55:34:b7:40      127        5        2
    xe-8/0/0:2.32767  Partner      127 d8:b1:22:f2:57:c0      127       18        2
    xe-8/0/0:3.32767    Actor      127 f4:cc:55:34:b7:40      127        6        2
    xe-8/0/0:3.32767  Partner      127 d8:b1:22:f2:57:c0      127       19        2
    xe-8/6/0:0.32767    Actor      127 f4:cc:55:34:b7:40      127       17        2
    xe-8/6/0:0.32767  Partner      127 d8:b1:22:f2:57:c0      127       12        2
    xe-8/6/0:1.32767    Actor      127 f4:cc:55:34:b7:40      127       18        2
    xe-8/6/0:1.32767  Partner      127 d8:b1:22:f2:57:c0      127       13        2
    xe-8/6/0:2.32767    Actor      127 f4:cc:55:34:b7:40      127       19        2
    xe-8/6/0:2.32767  Partner      127 d8:b1:22:f2:57:c0      127       14        2
    xe-8/7/0:1.32767    Actor      127 f4:cc:55:34:b7:40      127       22        2
    xe-8/7/0:1.32767  Partner      127 d8:b1:22:f2:57:c0      127       20        2
    xe-8/7/0:2.32767    Actor      127 f4:cc:55:34:b7:40      127       23        2
    xe-8/7/0:2.32767  Partner      127 d8:b1:22:f2:57:c0      127       21        2
    xe-8/7/0:3.32767    Actor      127 f4:cc:55:34:b7:40      127       24        2
    xe-8/7/0:3.32767  Partner      127 d8:b1:22:f2:57:c0      127       22        2
  LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
    et-8/7/4.32767             0           0            0            0
    xe-8/0/0:0.32767           0           0            0            0
    xe-8/0/0:1.32767           0           0            0            0
    xe-8/0/0:2.32767           0           0            0            0
    xe-8/0/0:3.32767           0           0            0            0
    xe-8/6/0:0.32767           0           0            0            0
    xe-8/6/0:1.32767           0           0            0            0
    xe-8/6/0:2.32767           0           0            0            0
    xe-8/7/0:1.32767           0           0            0            0
    xe-8/7/0:2.32767           0           0            0            0
    xe-8/7/0:3.32767           0           0            0            0
  Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
    et-8/7/4.32767             0           0            0            0
    xe-8/0/0:0.32767           0           0            0            0
    xe-8/0/0:1.32767           0           0            0            0
    xe-8/0/0:2.32767           0           0            0            0
    xe-8/0/0:3.32767           0           0            0            0
    xe-8/6/0:0.32767           0           0            0            0
    xe-8/6/0:1.32767           0           0            0            0
    xe-8/6/0:2.32767           0           0            0            0
    xe-8/7/0:1.32767           0           0            0            0
    xe-8/7/0:2.32767           0           0            0            0
    xe-8/7/0:3.32767           0           0            0            0
```

```
    Protocol multiservice, MTU: Unlimited, Generation: 36376, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (Aggregated Ethernet with VLAN Stacking)

```
user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 48, Generation: 186
  Link-level type: 52, MTU: 1518, Speed: 2000mbps, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5e:00:53:3f, Hardware address: 00:00:5e:00:53:3f
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              2406875                40152 bps
   Output bytes  :              1124470                22056 bps
   Input  packets:                 5307                    5 pps
   Output packets:                13295                   21 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0               859777                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0              1897615                    0
```

```
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0               662505                    0

  Logical interface ae0.451 (Index 69) (SNMP ifIndex 167) (Generation 601)
    Flags: SNMP-Traps VLAN-Tag [ 0x8100.451 ]  Encapsulation: VLAN-VPLS
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :           289          0         25685          376
        Output:          1698          4        130375         3096
    Link:
      ge-1/2/0.451
        Input :           289          0         25685          376
        Output:             0          0             0            0
      ge-1/2/1.451
        Input :             0          0             0            0
        Output:          1698          4        130375         3096
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-1/2/0.451               0           0            0            0
      ge-1/2/1.451               0           0            0            0
    Protocol vpls, MTU: 1518, Generation: 849, Route table: 3
      Flags: Is-Primary

  Logical interface ae0.452 (Index 70) (SNMP ifIndex 170) (Generation 602)
    Flags: SNMP-Traps VLAN-Tag [ 0x8100.452 ]  Encapsulation: VLAN-VPLS
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :           293          1         26003         1072
        Output:          1694          3        130057         2400
    Link:
      ge-1/2/0.452
        Input :           293          1         26003         1072
        Output:          1694          3        130057         2400
      ge-1/2/1.452
        Input :             0          0             0            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-1/2/0.452               0           0            0            0
      ge-1/2/1.452               0           0            0            0
    Protocol vpls, MTU: 1518, Generation: 850, Route table: 3
      Flags: None
...
```

show interfaces extensive (Aggregated Ethernet for Junos OS)

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 561, Generation: 131
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled, Minimum links
needed: 1, Minimum bandwidth needed: 1bps,

  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps
  Current address: 3c:61:04:76:9e:d6, Hardware address: 3c:61:04:76:9e:d6
  Last flapped   : 2020-12-08 16:51:00 IST (00:03:35 ago)
  Statistics last cleared: 2020-12-12 01:13:41 IST (00:00:26 ago)
  Traffic statistics:
    Input  bytes  :        86956823       28608000 bps
    Output bytes  :        86763213       28418520 bps
    Input  packets:         1086965          44698 pps
    Output packets:          834942          34334 pps
    IPv6 transit statistics:
    Input  bytes  :               0
    Output bytes  :               0
    Input  packets:               0
    Output packets:               0
  MAC statistics:                 Receive       Transmit
    Broadcast packets                   0              0
    Multicast packets                   9              9
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Police discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:     Queued packets     Transmitted packets     Dropped packets
    0                         816857                  816857                   0
    1                              0                       0                   0
    2                              0                       0                   0
    3                             20                      20                   0
  Queue number:        Mapped forwarding classes
    0                  best-effort
    1                  expedited-forwarding
    2                  assured-forwarding
```

```
    3                  network-control

  Logical interface ae0.0 (Index 70) (SNMP ifIndex 571) (Generation 135)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :       1177122      44785      94169363     28663824
        Output:        903806      34206      94000321     28760848
    Adaptive Statistics:
        Adaptive Adjusts:          30
        Adaptive Scans  :          10
        Adaptive Updates:          30
    Link:
      et-4/0/0:0.0
        Input :        286853      10917      22947986      6986824
        Output:        230478       8717      23666988      7234448
      et-4/0/0:1.0
        Input :        293489      11165      23478977      7146184
        Output:        221212       8369      23272661      7121096
      et-4/0/0:2.0
        Input :        304490      11584      24359200      7414200
        Output:        228303       8648      23840112      7300128
      et-4/0/0:3.0
        Input :        292290      11119      23383200      7116616
        Output:        223813       8472      23220560      7105176

    Aggregate member links: 4

    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-4/0/0:0.0               0           0            0            0
      et-4/0/0:1.0               0           0            0            0
      et-4/0/0:2.0               0           0            0            0
      et-4/0/0:3.0               0           0            0            0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 157, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.0/24, Local: 192.168.0.1, Broadcast: 192.168.0.255, Generation: 150
    Protocol multiservice, MTU: Unlimited, Generation: 159, Route table: 0
```

```
      Flags: Is-Primary
      Policer: Input: default arp policer
```

show interfaces extensive (Aggregated Ethernet for Junos OS Evolved)

```
user@host> show interfaces extensive ae0
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 1014, SNMP ifIndex: 514, Generation: 661424965257
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled, Minimum links
needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Current address: 2c:21:31:4c:fc:7a, Hardware address: 2c:21:31:4c:fc:7a
  Last flapped   : 2020-12-11 21:10:37 IST (03:44:07 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes  :        57035046844       4377472 bps
    Output bytes  :        57999987279      23236888 bps
    Input  packets:          581991978          5582 pps
    Output packets:          591838667         29637 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Police discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Members:
     Statistics         Packets        pps          Bytes          bps
     et-0/0/23:0
        Input :      29413023           0     2882475760      1092112
        Output:      27678003        7607     2712307244      5964120
      et-0/0/23:1
        Input :      29415832           0     2882751626      1267400
        Output:      34012078        7155     3333183616      5610008
      et-0/0/23:2
        Input :      29805451           0     2920923968      1046232
        Output:      31564271        7820     3093298478      6131176
      et-0/0/23:3
        Input :      26248079           0     2572311742       971728
        Output:      27145040        7055     2660203820      5531584
  Egress queues: 8 supported, 4 in use
```

```
  Queue counters:     Queued packets     Transmitted packets     Dropped packets
    0                      120388041               120388041                   0
    1                              0                       0                   0
    2                              0                       0                   0
    3                          11271                   11271                   0
  Queue number:        Mapped forwarding classes
    0                  best-effort
    1                  expedited-forwarding
    2                  assured-forwarding
    3                  network-control
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0

  Logical interface ae0.0 (Index 70) (SNMP ifIndex 571) (Generation 135)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Statistics        Packets        pps           Bytes        bps
    Bundle:
        Input :     581991745       5635     46559172600    3607616
        Output:     591833599      29890     47346516111   19129848

    Aggregate member links: 4

    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-4/0/0:0.0               0           0            0            0
      et-4/0/0:1.0               0           0            0            0
      et-4/0/0:2.0               0           0            0            0
      et-4/0/0:3.0               0           0            0            0
    Protocol inet, MTU: 1500, Generation: 661424965263, Route table: 0
      Flags: Sendbcast-pkt-to-re, Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.0/24, Local: 192.168.0.1, Broadcast: 192.168.0.255, Generation:
661424965266
    Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 661424965264, Route table: 0
      Flags: Is-Primary
    Protocol multiservice, MTU: Unlimited, Generation: 661424965275, Route table: 0
      Flags: Is-Primary
```

Starting in Junos OS Evolved Release 20.3R1, when you configure adaptive load balancing, the show interfaces ae0 extensive command displays adaptive statistics under the physical interface, not the logical interface for the PTX10008 Series routers.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Ethernet Interfaces User Guide for Routing Devices

show interfaces anchor-group (Aggregated Packet Forwarding Engine)

IN THIS SECTION

- Syntax | 20437
- Description | 20437
- Options | 20438
- Required Privilege Level | 20438
- Output Fields | 20438
- Sample Output | 20440
- Release Information | 20441

Syntax

```
show interfaces anchor-group
<brief | detail>
interface-name
```

Description

Display interface information for the aggregated Packet Forwarding Engine group.

Options

none (Same as brief) Display a summary of the aggregated Packet Forwarding Engine interface information.

brief | detail (Optional) Display the specified level of output.

interface-name Name of the interface within the anchor Packet Forwarding Engine group.

> **NOTE**: The interface must be an aggregated Packet Forwarding Engine interface (apfe-).

Required Privilege Level

view

Output Fields

Table 616 on page 20438 lists the output fields for the `show interfaces anchor-group` command. Output fields are listed in the approximate order in which they appear.

Table 616: show interfaces anchor-group Output Fields

Field Name	Field Description	Level of Output
`Active`	Anchor Packet Forwarding Engine is operational.	All levels
`Inactive`	Anchor Packet Forwarding Engine is not operational.	All levels
`PF`	Primary Packet Forwarding Engine anchor has failed.	All levels
`MS`	Primary Packet Forwarding Engine is protected by a secondary Packet Forwarding Engine in manually switched mode for primary role change.	All levels
`HS`	Primary Packet Forwarding Engine is protected by a secondary Packet Forwarding Engine in hot standby mode.	All levels

Table 616: show interfaces anchor-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
WS	Primary Packet Forwarding Engine is protected by a secondary Packet Forwarding Engine in warm standby mode.	All levels
Group	Name of the aggregated Packet Forwarding Engine group.	brief none
Mode	Redundancy mode in which the aggregated Packet Forwarding Engine group operates. Currently, only warm standby mode is supported.	brief none
Sub-group ID	Redundancy subgroups within the anchor Packet Forwarding Engine group configuration that has FPCs as members. This is derived out of the Packet Forwarding Engines on a given FPC. For example, if the first Packet Forwarding Engine is assigned the number 0, then all the other Packet Forwarding Engines with sub-group ID 0 form the N:1 redundancy group.	brief none
Interface	Anchor Packet Forwarding Engine interface (pfe-).	All levels
Configured State	State in which the anchor Packet Forwarding Engine was configured. • Primary—Anchor Packet Forwarding Engine is in the pool of primary members. • Secondary—Anchor Packet Forwarding Engine is a backup to all the primary members.	All levels
Operational State	Indication whether the anchor Packet Forwarding Engine is operational (Active) or not operational (Inactive).	All levels
Redundancy State	Redundancy state (primary or secondary) in which the anchor Packet Forwarding Engine was configured.	All levels
Group Name	Name of the aggregated Packet Forwarding Engine group.	detail
Group Mode	Redundancy mode in which the aggregated Packet Forwarding Engine group operates. Currently, only warm standby mode is supported.	detail

Table 616: show interfaces anchor-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Group Id	Internal ID generated for the group.	detail
Switchover information	Switchover details, if any.	detail
Subgroup identifier	Number of redundancy subgroups within the anchor Packet Forwarding Engine group configuration that has FPCs as members. This is derived out of the Packet Forwarding Engines on a given FPC. For example, if the first Packet Forwarding Engine is assigned the number 0, then all the other Packet Forwarding Engines with subgroup ID 0 form the N:1 redundancy group.	detail

Sample Output

show interfaces anchor-group brief

```
user@host> show interfaces anchor-group brief
Redundancy Status Legend:

Active: Operational                Inactive: Non-operational
MS: Manually switched              PF: Primary failed
HS: Hot standby                    WS: Warm standby

Group     Mode  Sub-group  Interface  Configured  Operational    Redundancy
                   ID                    State       State          State

apfe0     WS       0       pfe-4/0/0  Primary     Active         Primary
                          pfe-5/0/0  Secondary   Active         Secondary
                   2       pfe-4/2/0  Primary     Active         Primary
                          pfe-5/2/0  Secondary   Active         Secondary
```

show interfaces anchor-group detail

```
user@host> show interfaces anchor-group detail
Active: Operational                Inactive: Non-operational
MS: Manually switched              PF: Primary failed
HS: Hot standby                    WS: Warm standby

Group Name: apfe0
  Group Mode: WS                       Group Id: 65
  Switchover information: None
  Interface pfe-4/2/0
    Configured state: Primary          Operational state: Active
    Redundancy state: Primary
    Subgroup identifier: 2
  Interface pfe-4/0/0
    Configured state: Primary          Operational state: Active
    Redundancy state: Primary
    Subgroup identifier: 0
  Interface pfe-5/0/0
    Configured state: Secondary        Operational state: Active
    Redundancy state: Secondary
    Subgroup identifier: 0
  Interface pfe-5/2/0
    Configured state: Secondary        Operational state: Active
    Redundancy state: Secondary
    Subgroup identifier: 2
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf system interfaces | 27532

show interfaces (ATM)

IN THIS SECTION

- Syntax | 20442
- Description | 20442
- Options | 20442
- Required Privilege Level | 20443
- Output Fields | 20443
- Sample Output | 20470
- Release Information | 20497

Syntax

```
show interfaces at-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M Series and T Series routers only) Display status information about the specified ATM interface.

Options

at-*fpc/pic/port*	Display standard information about the specified ATM interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.

snmp-index ***snmp-index*** (Optional) Display the SNMP index of the interface.

statistics (Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 617 on page 20443 lists the output fields for the `show interfaces` (ATM) command. Output fields are listed in the approximate order in which they appear.

Table 617: ATM show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical interface`	Name of the physical interface.	All levels
`Enabled`	State of the interface. Possible values are described in the “Enabled Field” section under Common Output Fields Description.	All levels
`Description`	Configured interface description.	All levels
`Interface index`	Physical interface's index number, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Link-level type`	Encapsulation being used on the physical interface: • `ATM-CCC-CELL-RELAY`—ATM cell relay for CCC. • `ATM-CCC-VC-MUX`—ATM virtual circuit (VC) for CCC. • `ATM-CISCO-NLPID`—Cisco-compatible ATM NLPID encapsulation. • `ATM-MIPP-LLC`—ATM MLPPP over ATM Adaptation Layer 5 (AAL5)/ logical link control (LLC). • `ATM-NLPID`—ATM NLPID encapsulation. • `ATM-PPP-LLC`—ATM PPP over AAL5/LLC. • `ATM-PPP-VC-MUX`—ATM PPP over raw AAL5. • `ATM-PVC`—ATM permanent virtual circuits. • `ATM-SNAP`—ATM LLC/SNAP encapsulation. • `ATM-TCC-SNAP`—ATM LLC/SNAP for translational cross-connection. • `ATM-TCC-VC-MUX`—ATM VC for translational cross-connection. • `ATM-VC-MUX`—ATM VC multiplexing. • `ETHER-OVER-ATM-LLC`—Ethernet over ATM (LLC/SNAP) encapsulation. • `ETHER-VPLS-OVER-ATM-LLC`—Ethernet VPLS over ATM (bridging) encapsulation.	All levels
`MTU`	MTU size on the physical interface.	All levels
`Clocking`	Reference clock source: `Internal` or `External`.	All levels
`framing` `Mode`	Framing mode: `SONET` or `SDH`.	All levels
`Speed`	Speed at which the interface is running as represented by the interface type (for example, `OC3`, `ADSL2+`, and `SHDSL(2-wire)`.	All levels

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Loopback`	Whether loopback is enabled and the type of loopback (`local` or `remote`).	All levels
`Payload scrambler`	Whether payload scrambling is enabled.	All levels
`Device flags`	Information about the physical device. Possible values are described in the "Device Flags" section under Common Output Fields Description.	All levels
`Link flags`	Information about the link. Possible values are described in the "Link Flags" section under Common Output Fields Description.	All levels
`CoS queues`	Number of CoS queues configured.	`detail extensive` none
`Hold-times`	Current interface hold-time up and hold-time down, in milliseconds.	`detail extensive`
`Current address`	Ethernet MAC address for this interface for Ethernet over ATM encapsulation.	`detail extensive` none
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second timezone`* (*`hour`*`:`*`minute`*`:`*`second`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive` none
`Input Rate`	Input rate in bits per second (bps) and packets per second (pps).	None specified
`Output Rate`	Output rate in bps and pps.	None specified
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Traffic statistics`	Statistics for traffic on the interface. • `Input bytes`—Number of bytes received on the interface • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input errors`	Input errors on the interface whose definitions are as follows: • `Errors`—Sum of the incoming frame terminations and frame check sequence (FCS) errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's random early detection (RED) mechanism. • `Invalid VCs`—Number of cells that arrived for a nonexistent VC. • `Framing errors`—Sum of AAL5 packets that have FCS errors, reassembly timeout errors, and length errors. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • `L3 incompletes`—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. • `L2 channel errors`—Number of times the software did not find a valid logical interface for an incoming frame. • `L2 mismatch timeouts`—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • `Resource errors`—Sum of transmit drops.	`extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If it increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • `Errors`—Sum of the outgoing frame terminations and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Aged packets`—Number of packets that remained so long in shared packet SDRAM that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • `MTU errors`—Number of packets larger than the MTU threshold. • `Resource errors`—Sum of transmit drops.	`extensive`
`Egress queues`	Total number of egress queues supported on the specified interface.	`detail extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Queue counters`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism. **NOTE**: Physical interface queue counters of ATM2 PICs displayed by the `show interfaces at-`*`fpc`*`/`*`pic`*`/`*`port`* `detail` command show the packet forwarding stream statistics associated with the ATM2 ports. Since multiple ports of the ATM2 PICs (except for the ATM2 dual-port OC12) share one packet forwarding stream, the physical interface queue counters reflect the aggregate of ATM2 port statistics.	`detail extensive`
`SONET alarms` `SONET defects`	SONET media-specific defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: `SONET PHY`, `SONET section`, `SONET line`, and `SONET path`.	`detail extensive` none
`SONET PHY`	Counts of specific SONET errors with detailed information. • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. State other than `OK` indicates a problem. Subfields are: • `PLL Lock`—Phase-locked loop • `PHY Light`—Loss of optical signal	`extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`SONET section`	Counts of specific SONET errors with detailed information. • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. State other than `OK` indicates a problem. Subfields are: • `BIP-B1`—Bit interleaved parity for SONET section overhead • `SEF`—Severely errored framing • `LOL`—Loss of light • `LOF`—Loss of frame • `ES-S`—Errored seconds (section) • `SES-S`—Severely errored seconds (section) • `SEFS-S`—Severely errored framing seconds (section)	`extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`SONET line`	Active alarms and defects, plus counts of specific SONET errors with detailed information. • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. State other than `OK` indicates a problem. Subfields are: • `BIP-B2`—Bit interleaved parity for SONET line overhead • `REI-L`—Remote error indication (near-end line) • `RDI-L`—Remote defect indication (near-end line) • `AIS-L`—Alarm indication signal (near-end line) • `BERR-SF`—Bit error rate fault signal failure • `BERR-SD`—Bit error rate defect signal degradation • `ES-L`—Errored seconds (near-end line) • `SES-L`—Severely errored seconds (near-end line) • `UAS-L`—Unavailable seconds (near-end line) • `ES-LFE`—Errored seconds (far-end line) • `SES-LFE`—Severely errored seconds (far-end line) • `UAS-LFE`—Unavailable seconds (far-end line)	`extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
SONET path	Active alarms and defects, plus counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. Subfields are: • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • ES-PFE—Errored seconds (far-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path)	extensive

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Received SONET overhead`
 	
`Transmitted SONET overhead`
	Values of the received and transmitted SONET overhead:
	
	• `C2`—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P.
	
	• `F1`—Section user channel byte. This byte is set aside for the purposes of users.
	
	• `K1` and `K2`—These bytes are allocated for APS signaling for the protection of the multiplex section.
	
	• `J0`—Section trace. This byte is defined for STS-1 number 1 of an STS-*N* signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter.
	
	• `S1`—Synchronization status. The S1 byte is located in the first STS-1 of an STS-*N*.
	
	• `Z3` and `Z4`—Allocated for future use.
		`extensive`

`SDH alarms`
 	
`SDH defects`
	SDH media-specific defects that can prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: `SDH PHY`, `SDH regenerator section`, `SDH multiplex section`, and `SDH path`.	All levels

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
SDH PHY	Active alarms and defects, plus counts of specific SDH errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. Subfields are: • PLL Lock—Phase-locked loop • PHY Light—Loss of optical signal	extensive
SDH regenerator section	Active alarms and defects, plus counts of specific SDH errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. Subfields are: • RS-BIP8—24-bit BIP for multiplex section overhead (B2 bytes) • OOF—Out of frame • LOS—Loss of signal • LOF—Loss of frame • RS-ES—Errored seconds (near-end regenerator section) • RS-SES—Severely errored seconds (near-end regenerator section) • RS-SEFS—Severely errored framing seconds (regenerator section)	extensive

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`SDH multiplex section`	Active alarms and defects, plus counts of specific SDH errors with detailed information. • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. State other than `OK` indicates a problem. Subfields are: • `MS-BIP24`—8-bit BIP for high-order path overhead (B3 byte) • `MS-FEBE`—Far-end block error (multiplex section) • `MS-FERF`—Far-end remote fail (multiplex section) • `MS-AIS`—Alarm indication signal (multiplex section) • `BERR-SF`—Bit error rate fault (signal failure) • `BERR-SD`—Bit error rate defect (signal degradation) • `MS-ES`—Errored seconds (near-end multiplex section) • `MS-SES`—Severely errored seconds (near-end multiplex section) • `MS-UAS`—Unavailable seconds (near-end multiplex section) • `MS-ES-FE`—Errored seconds (far-end multiplex section) • `MS-SES-FE`—Severely errored seconds (far-end multiplex section) • `MS-UAS-FE`—Unavailable seconds (far-end multiplex section)	`extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
SDH path	Active alarms and defects, plus counts of specific SDH errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. Subfields are: • HP-BIP8—8-bit BIP for regenerator section overhead (B1 byte) • HP-FEBE—Far-end block error (high-order path) • HP-LOP—Loss of pointer (high-order path) • HP-AIS—High-order-path alarm indication signal • HP-FERF—Far-end remote fail (high-order path) • HP-UNEQ—Unequipped (high-order path) • HP-PLM—Payload label mismatch (high-order path) • HP-ES—Errored seconds (near-end high-order path) • HP-SES—Severely errored seconds (near-end high-order path) • HP-UAS—Unavailable seconds (near-end high-order path) • HP-ES-FE—Errored seconds (far-end high-order path) • HP-SES-FE—Severely errored seconds (far-end high-order path) • HP-UAS-FE—Unavailable seconds (far-end high-order path)	extensive

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Received SDH overhead`
 	
`Transmitted SDH overhead`
	Values of the received and transmitted SONET overhead:
	
	• `C2`—Signal label. This byte is allocated to identify the construction and content of the STS-level SPE and for PDI-P.
	
	• `F1`—Section user channel byte. This byte is set aside for the purposes of users.
	
	• `K1` and `K2`—These bytes are allocated for APS signaling for the protection of the multiplex section.
	
	• `J0`—Section trace. This byte is defined for STS-1 number 1 of an STS-*N* signal. This bye is used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter.
	
	• `S1`—Synchronization status. The S1 byte is located in the first STS-1 of an STS-*N*.
	
	• `Z3` and `Z4`—These bytes are allocated for future use.
		`extensive`

`Received path trace`
 	
`Transmitted path trace`
	SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.	`extensive`

`ATM Status`	ATM state information: • `HCS State`—Status of the header check sequence. ATM uses the HCS field in the cell header in the cell delineation process to frame ATM cell boundaries. The HCS is an FCS-8 calculation over the first four octets of the ATM cell header. • `LOC`—Current loss of cell (LOC) delineation state. `OK` means that no LOC is currently asserted.	`extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`ATM Statistics`	ATM statistics for the interface: • `Uncorrectable HCS errors`—Number of cells dropped because the cell delineation failed. These errors most likely indicate that a SONET/SDH layer problem has occurred. • `Correctable HCS errors`—Number of correctable HCS errors that occurred. The cell delineation process can recover from these errors and locate the ATM cell boundary, although the framing process is not quite stable. The ATM cell is not dropped. This counter increases when the cell delineation process changes its state from `present` to `sync` (for example, when a cable is plugged into the interface). The following error statistics are from the framer: • `Tx cell FIFO overruns`—Number of overruns in the transmit FIFO. • `Rx cell FIFO overruns`—Number of overruns in the receive FIFO. • `Rx cell FIFO underruns`—Number of underruns in the receive FIFO. • `Input cell count`—Number of ATM cells received by the interface (not including idle cells). • `Output cell count`—Number of ATM cells transmitted by the interface (including idle cells). • `Output idle cell count`—Number of idle cells sent by the port. When ATM has nothing to send, it sends idle cells to fill the time slot. • `Output VC queue drops`—Number of packets dropped by a port on the PIC. Packets are dropped because of queue limits on the VCs. The following error statistics are from the SAR: • `Input no buffers`—Number of AAL5 packets dropped because no channel blocks or buffers were available to handle them. • `Input length errors`—Number of AAL5 packets dropped because their length was incorrect. Usually, these errors occur because a cell has been corrupted or lost, or because the length field was corrupted. They can also mean the AAL5 length field was zero.	`extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Input timeouts`—Number of AAL5 packets dropped because of a reassembly timeout. • `Input invalid VCs`—Number of AAL5 packets dropped because the header was unrecognized (because the VC was not correct or not configured). • `Input bad CRCs`—Number of AAL5 packets dropped because of frame check sequence errors. • `Input OAM cell no buffers`—Number of received OAM cells or raw cells dropped because no buffers were available to handle them. • `L2 circuit out-of-sequence packets`—(Layer 2 AAL5 mode) Number of AAL5 packets that are out of sequential order. • `Denied packets count`—The number of packets dropped due to VLAN priority deny packets or due to an error forwarding configuration that might cause a negative frame length, that is, the stripping size is larger than the packet size.	
`Packet Forwarding Engine configuration`	Information about the configuration of the Packet Forwarding Engine: • `Destination slot`—FPC slot number.	`extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`CoS information`	Information about the CoS queue for the physical interface. • `CoS transmit queue`—Queue number and its associated user-configured forwarding class name. • `Bandwidth %`—Percentage of bandwidth allocated to the queue. • `Bandwidth bps`—Bandwidth allocated to the queue (in bps). • `Buffer %`—Percentage of buffer space allocated to the queue. • `Buffer usec`—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • `Priority`—Queue priority: `low` or `high`. • `Limit`—Displayed if rate limiting is configured for the queue. Possible values are `none` and `exact`. If `exact` is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If `none` is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	`extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`VPI`	(ATM2) Virtual path identifier information:
	• `Flags`—VPI flags can be one or more of the following:
	  • `Active` (virtual path is up)
	  • `OAM` (operation and maintenance is enabled)
	  • `Shaping` (shaping is configured)
	• `CBR, Peak`
	• `OAM, Period`—Interval at which OAM F4 loopback cells are sent.
	• `Up count`—Number of F4 OAM cells required to consider the virtual path up; the range is `1` through `255`.
	• `Down count`—Number of F4 OAM cells required to consider the virtual path down; the range is `1` through `255`.
	• `Total down time`—Total number of seconds the VPI has been down since it was opened, using the format `Total down time:` *`hh`*`:`*`mm`*`:`*`ss`* or `Never`.
	• `Last down`—Time of last `Down` transition, using the format `Last down:` *`hh`*`:`*`mm`*`:`*`ss`* `ago` or `Never`.
	• `OAM F4 cell statistics`—(Nonpromiscuous mode) OAM F4 statistics:
	  • `Total received`—Number of OAM F4 cells received.
	  • `Total sent`—Number of OAM F4 cells sent.
	  • `Loopback received`—Number of OAM F4 loopback cells received.
	  • `Loopback sent`—Number of OAM F4 loopback cells sent.
	  • `Last received`—Time at which the last OAM F4 cell was received.
	  • `Last sent`—Time at which the last OAM F4 cell was sent.
		`detail extensive`
		none

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `RDI received`—Number of OAM F4 cells received with the remote defect indication bit set. • `RDI sent`—Number of OAM F4 cells sent with the RDI bit set. • `AIS received`—Number of OAM F4 cells received with the alarm indication signal bit set. • `AIS sent`—Number of OAM F4 cells sent with the AIS bit set. `Traffic statistics:` • `Input bytes`—Number of bytes received on the VPI. • `Output bytes`—Number of bytes transmitted on the VPI. • `Input packets`—Number of packets received on the VPI. • `Output packets`—Number of packets transmitted on the VPI.	
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Logical interface index number, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	Logical interface SNMP interface index number.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Flags`	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under Common Output Fields Description.	All levels
`Input packets`	Number of packets received on the logical interface.	None specified

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output packets`	Number of packets transmitted on the logical interface.	None specified
`Encapsulation`	Encapsulation on the logical interface.	All levels
`Traffic statistics`	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.	`detail extensive`
`Local statistics`	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.	`detail extensive`
`Transit statistics`	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.	`detail extensive`
`Input packets`	Number of packets received on the logical interface.	None specified
`Output packets`	Number of packets transmitted on the logical interface.	None specified
`protocol-family`	Protocol family configured on the logical interface. If the protocol is `inet`, the IP address of the interface is also displayed.	`brief`
`Protocol`	Protocol family configured on the logical interface.	`detail extensive` none
`MTU`	MTU size on the logical interface.	`detail extensive` none

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Route table`	Routing table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail extensive`
`Flags`	Information about the protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	`detail extensive` none
`Addresses, Flags`	Information about the address flags. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	`detail extensive` none
`Destination`	IP address of the remote side of the connection.	`detail extensive` none
`Local`	IP address of the logical interface.	`detail extensive` none
`Broadcast`	Broadcast address.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

VCI	Virtual circuit identifier number and information:
	• Flags—VCI flags:
	  • Active—VCI is up and in working condition.
	  • CCC down—VCI CCC is not in working condition.
	  • Closed—VCI is closed because the user disabled the logical or physical interface from the CLI.
	  • Configured—VCI is configured.
	  • Down—VCI is not in working condition. The VCI might have alarms, defects, F5 AIS/RDI, or no response to OAM loopback cells.
	  • ILMI—VCI is up and in working condition.
	  • OAM—OAM loopback is enabled.
	  • Multicast—VCI is a multicast VCI or DLCI.
	  • Multipoint destination—VCI is configured as a multipoint destination.
	  • None—No VCI flags.
	  • Passive-OAM—Passive OAM is enabled.
	  • Shaping—Shaping is enabled.
	  • Sustained—Shaping rate is set to Sustained.
	  • Unconfigured—VCI is not configured.
	• Total down time—Total number of seconds the VCI has been down, using the format Total down time: *hh*:*mm*:*ss* orNever.
	• Last down—Time of last Down transition, using the format Last down: *hh*:*mm*:*ss*.
		All levels

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `EPD threshold`—(ATM2 only) Threshold at which a packet is dropped when the queue size (in number of cells) exceeds the early packet-discard (EPD) value.	

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

VCI (continued)	• Transmit weight cells—(ATM2 only) Amount of bandwidth assigned to this queue.
	• ATM per-VC transmit statistics:
	  • Tail queue packet drops—Number of packets dropped because of bandwidth constraints. This value indicates that packets are queued to send out at a rate faster than allowed.
	• OAM F4 cell statistics—(Nonpromiscuous mode) OAM F4 statistics:
	  • Total received—Number of OAM F4 cells received.
	  • Total sent—Number of OAM F4 cells sent.
	  • Loopback received—Number of OAM F4 loopback cells received.
	  • Loopback sent—Number of OAM F4 loopback cells sent.
	  • Last received—Time at which the last OAM F4 cell was received.
	  • Last sent—Time at which the last OAM F4 cell was sent.
	  • RDI received—Number of OAM F4 cells received with the remote defect indication bit set.
	  • RDI sent—Number of OAM F4 cells sent with the RDI bit set.
	  • AIS received—Number of OAM F4 cells received with the alarm indication signal bit set.
	  • AIS sent—Number of OAM F4 cells sent with the AIS bit set.
	• Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.
	  • Input bytes—Number of bytes received on the interface.
	  • Output bytes—Number of bytes transmitted on the interface.
	  • Input packets—Number of packets received on the interface
	  • Output packets—Number of packets transmitted on the interface.
		All levels

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`IMA group properties`	• `Version`—The specified IMA specification version, either IMA 1.0 or IMA 1.1.
	• `Frame length`—The specified frame size, which can be 32, 64, 128, or 256.
	• `Differential delay`—Maximum differential delay among links in milliseconds.
	• `Symmetry`—Either Common Transmit Clock or Independent Transmit Clock timing mode.
	• `Transmit clock`—The specified IMA clock mode, either common or independent.
	• `Minimum links`—The number of minimum active links specified in both transmit and receive directions.
	  • `Transmit`—The per-PIC limit on the number of minimum active links in the transmit direction.
	  • `Receive`—The per-PIC limit on the number of minimum active links in the receive direction.
	• `Frame synchronization`—The specified IMA frame synchronization state transition variables (Alpha, Beta, and Gamma) and their specified values.
	  • `Alpha`—The number of consecutive invalid ICP cells for IFSM.
	  • `Beta`—The number of consecutive errored ICP cells for IFSM.
	  • `Gamma`—The number of consecutive valid ICP cells for IFSM.
	• `Links`—The number of IMA links assigned to the IMA group.
		`detail extensive`
		none

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`IMA group alarms`	• `Start-up-FE`—Far-end group alarm status • `Config-Aborted`—Near-end configuration terminated group alarm status • `Config-Aborted-FE`—Far-end configuration force-stopped group alarm status • `Insufficient-Links`—Near-end insufficient links group alarm status • `Insufficient-Links-FE`—Far-end insufficient links group alarm status • `Blocked-FE`—Far-end blocked group alarm status • `GR-Timing-Mismatch`—Group timing mismatch alarm status	`detail extensive` none
`IMA group defects`	• `Start-up-FE`—Far-end group defect status • `Config-Aborted`—Near-end configuration terminated group defect status • `Config-Aborted-FE`—Far-end configuration terminated group defect status • `Insufficient-Links`—Near-end insufficient links group defect status • `Insufficient-Links-FE`—Far-end insufficient links group defect status • `Blocked-FE`—Far-end blocked group defect status • `GR-Timing-Mismatch`—Group timing mismatch defect status	`detail extensive` none
`IMA Group state`	Near-end and far-end group status	`detail extensive` none

Table 617: ATM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`IMA group media`	IMA group media status, including seconds, count and state for the following media parameters: • `FC` • `FC-FE` • `Addr-Mismatch` • `Running` • `UAS`	`detail extensive` none

Sample Output

show interfaces (ATM, IMA Group)

```
user@host> show interfaces at-1/0/0
Physical interface: at-1/0/0, Enabled, Physical link is Up
  IMA group properties:
    Version              : 1.1
    Frame length         : 128
    Differential delay   : 25 milliseconds
    Symmetry             : Symmetrical Configuration and Operation
    Transmit clock       : Common
    Minimum links        : Transmit: 1, Receive: 1
    Frame synchronization: Alpha: 2, Beta: 2, Gamma: 1
    Links                : None
  IMA group alarms  : Start-up-FE Config-Aborted Config-Aborted-FE Insufficient-Links
Insufficient-Links-FE Blocked-FE GR-Timing-Mismatch
  IMA group defects : Start-up-FE Config-Aborted Config-Aborted-FE Insufficient-Links
Insufficient-Links-FE Blocked-FE GR-Timing-Mismatch
  IMA Group state:
    Near end : Start up
    Far  end : Start up
  IMA group media:         Seconds        Count  State
    FC                                        0
```

```
    FC-FE                                    0
    Addr-Mismatch                            0
    Running                      0
    UAS                          0
```

show interfaces extensive (ATM IMA Group)

```
user@host> show interfaces at-0/0/10 extensive
Physical interface: at-0/0/10, Enabled, Physical link is Up
   Interface index: 178, SNMP ifIndex: 540, Generation: 531
   Link-level type: ATM-PVC, MTU: 2048, Speed: Unspecified, Loopback: None, Payload scrambler:
Enabled
   Device flags   : Present Running
   Link flags     : None
   CoS queues     : 8 supported, 4 maximum usable queues
   Hold-times     : Up 0 ms, Down 0 ms
   Current address: 00:00:5e:00:53:0a
   Last flapped   : 2012-03-16 16:49:15 PDT (2d 07:12 ago)
   Statistics last cleared: 2012-03-16 16:56:58 PDT (2d 07:05 ago)
   Traffic statistics:
    Input  bytes  :                    0                    0 bps
    Output bytes  :                    0                    0 bps
    Input  packets:                    0                    0 pps
    Output packets:                    0                    0 pps
    IPv6 transit statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
   Input errors:
     Errors: 0, Drops: 0, Invalid VCs: 0, Framing errors: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
     L2 mismatch timeouts: 0, Resource errors: 0
   Output errors:
     Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0, Resource
errors: 0
   IMA group properties:
     Version              : 1.1
     Frame length         : 128
     Differential delay   : 25 milliseconds
     Symmetry             : Symmetrical Configuration and Operation
```

```
    Transmit clock       : Common
    Minimum links        : Transmit: 1, Receive: 1
    Frame synchronization: Alpha: 2, Beta: 2, Gamma: 1
    Link #1              : t1-0/0/4             up
  IMA Group alarms   : None
  IMA Group defects  : None

  IMA Group state:
    Near end : Operational
    Far  end : Operational
  IMA group media:         Seconds        Count  State
    FC                                        0
    FC-FE                                     0
    Addr-Mismatch                             0
    Running                 198306
    UAS                          0
  ATM status:
    HCS state:     Sync
    LOC      :       OK
  ATM Statistics:
    Uncorrectable HCS errors: 0, Correctable HCS errors: 0, Tx cell FIFO overruns: 0, Rx cell
FIFO overruns: 0,
    Rx cell FIFO underruns: 0, Input cell count: 0, Output cell count: 0, Output idle cell
count: 0,
    Output VC queue drops: 0, Input no buffers: 0, Input length errors: 0, Input timeouts: 0,
Input invalid VCs: 0,
    Input bad CRCs: 0, Input OAM cell no buffers: 0
  Packet Forwarding Engine configuration:
    Destination slot: 0
    VPI 2
      Flags: Active
      Total down time: 0 sec, Last down: Never
        Traffic statistics:
         Input  bytes  :                    0
         Output bytes  :                    0
         Input  packets:                    0
         Output packets:                    0

  Logical interface at-0/0/10.602 (Index 71) (SNMP ifIndex 1057) (Generation 17226)
    Flags: Point-To-Point SNMP-Traps CCC-Down 0x0 Encapsulation: ATM-CCC-Cell-Relay
    L2 circuit cell bundle size: 1, bundle timeout: 125 usec, timeout count: 0
    L2 circuit out-of-sequence count: 0, denied packets count: 0
```

show interfaces (ATM1, SONET Mode)

```
user@host> show interfaces at-1/0/0
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 300, SNMP ifIndex: 194
  Description: to allspice at-1/0/0
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:fe
  Last flapped   : 2006-02-24 14:28:12 PST (6d 01:51 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None

  Logical interface at-1/0/0.0 (Index 64) (SNMP ifIndex 204)
    Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-SNAP
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 4470
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.220.24/30, Local: 192.168.220.26,
        Broadcast: 192.168.220.27
    Protocol iso, MTU: 4470
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
        Input packets : 0
        Output packets: 0
```

show interfaces brief (ATM1, SONET Mode)

```
user@host> show interfaces at-1/0/0 brief
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Description: to allspice at-1/0/0
```

```
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None

  Logical interface at-1/0/0.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-SNAP
    inet  192.168.220.26/30
    iso
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
```

show interfaces detail (ATM1, SONET Mode)

```
user@host> show interfaces at-1/0/0 detail
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 300, SNMP ifIndex: 194, Generation: 183
  Description: to allspice at-1/0/0
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:fe
  Last flapped   : 2006-02-24 14:28:12 PST (6d 01:55 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  SONET alarms   : None
  SONET defects  : None
```

```
  Logical interface at-1/0/0.0 (Index 64) (SNMP ifIndex 204) (Generation 5)
    Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-SNAP
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 4470, Generation: 13, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.220.24/30, Local: 192.168.220.26,
        Broadcast: 192.168.220.27, Generation: 14
    Protocol iso, MTU: 4470, Generation: 14, Route table: 0
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
```

show interfaces extensive (ATM1, SONET Mode)

```
user@host> show interfaces at-1/0/0 extensive
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 300, SNMP ifIndex: 194, Generation: 183
```

```
Description: to allspice at-1/0/0
Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
Speed: OC3, Loopback: None, Payload scrambler: Enabled
Device flags   : Present Running
Link flags     : None
CoS queues     : 4 supported, 4 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:fe
Last flapped   : 2006-02-24 14:28:12 PST (6d 01:56 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :                    0                    0 bps
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
Input errors:
  Errors: 0, Drops: 0, Invalid VCs: 0, Framing errors: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,
  Resource errors: 0
Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                    0                    0                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont                   0                    0                    0
SONET alarms   : None
SONET defects  : None
SONET PHY:            Seconds        Count  State
  PLL Lock                  0            0  OK
  PHY Light                 0            0  OK
SONET section:
  BIP-B1                    0            0
  SEF                       0            0  OK
  LOS                       0            0  OK
  LOF                       0            0  OK
  ES-S                      0
  SES-S                     0
  SEFS-S                    0
SONET line:
  BIP-B2                    0            0
```

```
  REI-L                    0            0
  RDI-L                    0            0  OK
  AIS-L                    0            0  OK
  BERR-SF                  0            0  OK
  BERR-SD                  0            0  OK
  ES-L                     0
  SES-L                    0
  UAS-L                    0
  ES-LFE                   0
  SES-LFE                  0
  UAS-LFE                  0
SONET path:
  BIP-B3                   0            0
  REI-P                    0            0
  LOP-P                    0            0  OK
  AIS-P                    0            0  OK
  RDI-P                    0            0  OK
  UNEQ-P                   1            1  OK
  PLM-P                    0            0  OK
  ES-P                     1
  SES-P                    1
  UAS-P                    0
  ES-PFE                   0
  SES-PFE                  0
  UAS-PFE                  0
Received SONET overhead:
  F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, C2(cmp) : 0x13, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00
Transmitted SONET overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, F2      : 0x00, Z3      : 0x00
  Z4      : 0x00
ATM status:
  HCS state:     Sync
  LOC      :       OK
ATM Statistics:
  Uncorrectable HCS errors: 0, Correctable HCS errors: 0,
  Tx cell FIFO overruns: 0, Rx cell FIFO overruns: 0,
  Rx cell FIFO underruns: 0, Input cell count: 0, Output cell count: 0,
  Output idle cell count: 0, Output VC queue drops: 0, Input no buffers: 0,
  Input length errors: 0, Input timeouts: 0, Input invalid VCs: 0,
  Input bad CRCs: 0, Input OAM cell no buffers: 0
```

```
Packet Forwarding Engine configuration:
  Destination slot: 1
CoS information:
  CoS transmit queue         Bandwidth          Buffer      Priority    Limit
                      %            bps     %     usec
  0 best-effort      95      147744000    95        0           low    none
  3 network-control   5        7776000     5        0           low    none

Logical interface at-1/0/0.0 (Index 64) (SNMP ifIndex 204) (Generation 5)
  Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-SNAP
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Local statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Transit statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Protocol inet, MTU: 4470, Generation: 13, Route table: 0
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 192.168.220.24/30, Local: 192.168.220.26,
      Broadcast: 192.168.220.27, Generation: 14
  Protocol iso, MTU: 4470, Generation: 14, Route table: 0
    Flags: None
  VCI 0.128
    Flags: Active
    Total down time: 0 sec, Last down: Never
    ATM per-VC transmit statistics:
     Tail queue packet drops: 0
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
```

show interfaces (ATM2, SDH Mode)

```
user@host> show interfaces at-0/2/1
Physical interface: at-0/2/1, Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 42
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SDH mode, Speed: OC3,
  Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:3f
  Last flapped   : 2006-03-24 13:29:58 PST (00:04:48 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SDH   alarms   : None
  SDH   defects  : None
    VPI 0
      Flags: Active
      Total down time: 0 sec, Last down: Never
  Traffic statistics:
        Input  packets:                    0
        Output packets:                    0

  Logical interface at-0/2/1.0 (Index 75) (SNMP ifIndex 51)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: ATM-SNAP
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 4470
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.12.6, Local: 10.0.12.5
    Protocol iso, MTU: 4470
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0
        Input packets : 0
        Output packets: 0

  Logical interface at-0/2/1.32767 (Index 76) (SNMP ifIndex 50)
    Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x4000
```

```
    Encapsulation: ATM-VCMUX
    Input packets : 0
    Output packets: 0
    VCI 0.4
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
        Input packets : 0
        Output packets: 0
```

show interfaces brief (ATM2, SDH Mode)

```
user@host> show interfaces at-0/2/1 brief
Physical interface: at-0/2/1, Enabled, Physical link is Up
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SDH mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  Logical interface at-0/2/1.0
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: ATM-SNAP
    inet  10.0.12.5        --> 10.0.12.6
    iso
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0

  Logical interface at-0/2/1.32767
    Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x4000
    Encapsulation: ATM-VCMUX
    VCI 0.4
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
```

show interfaces detail (ATM2, SDH Mode)

```
user@host> show interfaces at-0/2/1 detail
Physical interface: at-0/2/1, Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 42, Generation: 40
```

```
Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SDH mode, Speed: OC3,
Loopback: None, Payload scrambler: Enabled
Device flags   : Present Running
Link flags     : None
CoS queues     : 4 supported, 4 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:3f
Last flapped   : 2006-03-24 13:29:58 PST (00:05:10 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :                    0                    0 bps
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                   0                    0                    0
  1 expedited-fo                  0                    0                    0
  2 assured-forw                  0                    0                    0
  3 network-cont                  0                    0                    0
SDH   alarms   : None
SDH   defects  : None
  VPI 0
    Flags: Active
    Total down time: 0 sec, Last down: Never
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0

Logical interface at-0/2/1.0 (Index 75) (SNMP ifIndex 51) (Generation 25)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: ATM-SNAP
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Local statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
```

```
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 4470, Generation: 62, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.12.6, Local: 10.0.12.5, Broadcast: Unspecified,
        Generation: 58
    Protocol iso, MTU: 4470, Generation: 63, Route table: 0
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
  Logical interface at-0/2/1.32767 (Index 76) (SNMP ifIndex 50) (Generation 26)
    Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x4000
    Encapsulation: ATM-VCMUX
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    VCI 0.4
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
```

```
        Input  bytes  :                    0
        Output bytes  :                    0
        Input  packets:                    0
        Output packets:                    0
```

show interfaces extensive (ATM2, SDH Mode)

```
user@host> show interfaces at-0/2/1 extensive
Physical interface: at-0/2/1, Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 42, Generation: 40
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SDH mode, Speed: OC3,
  Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:3f
  Last flapped   : 2006-03-24 13:29:58 PST (00:06:49 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Invalid VCs: 0, Framing errors: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,
    Resource errors: 0
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  SDH   alarms   : None
  SDH   defects  : None
  SDH PHY:              Seconds        Count  State
    PLL Lock                  0            0  OK
```

```
  PHY Light                 1            1  OK
SDH regenerator section:
  RS-BIP8                   2         8828
  OOF                       2            2  OK
  LOS                       2            1  OK
  LOF                       2            1  OK
  RS-ES                     4
  RS-SES                    3
  RS-SEFS                   2
SDH multiplex section:
  MS-BIP24                  2          771
  MS-FEBE                   1        17476
  MS-FERF                   2            1  OK
  MS-AIS                    2            1  OK
  BERR-SF                   0            0  OK
  BERR-SD                   0            0  OK
  MS-ES                     4
  MS-SES                    2
  MS-UAS                    0
  MS-ES-FE                  3
  MS-SES-FE                 2
  MS-UAS-FE                 0
SDH path:
  HP-BIP8                   1            6
  HP-FEBE                   1          251
  HP-LOP                    0            0  OK
  HP-AIS                    2            1  OK
  HP-FERF                   3            2  OK
  HP-UNEQ                   1            1  OK
  HP-PLM                    2            1  OK
  HP-ES                     4
  HP-SES                    3
  HP-UAS                    0
  HP-ES-FE                  3
  HP-SES-FE                 3
  HP-UAS-FE                 0
Received SDH overhead:
  F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, C2(cmp) : 0x13, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00
Transmitted SDH overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, F2      : 0x00, Z3      : 0x00
```

```
  Z4      : 0x00
ATM status:
  HCS state:     Sync
  LOC      :       OK
ATM Statistics:
  Uncorrectable HCS errors: 0, Correctable HCS errors: 0,
  Tx cell FIFO overruns: 0, Rx cell FIFO overruns: 0,
  Rx cell FIFO underruns: 0, Input cell count: 0, Output cell count: 0,
  Output idle cell count: 0, Output VC queue drops: 0, Input no buffers: 0,
  Input length errors: 0, Input timeouts: 0, Input invalid VCs: 0,
  Input bad CRCs: 0, Input OAM cell no buffers: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
  VPI 0
    Flags: Active
    Total down time: 0 sec, Last down: Never
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0

Logical interface at-0/2/1.0 (Index 75) (SNMP ifIndex 51) (Generation 25)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: ATM-SNAP
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Local statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Transit statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Protocol inet, MTU: 4470, Generation: 62, Route table: 0
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.0.12.6, Local: 10.0.12.5, Broadcast: Unspecified,
```

```
        Generation: 58
    Protocol iso, MTU: 4470, Generation: 63, Route table: 0
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
  Logical interface at-0/2/1.32767 (Index 76) (SNMP ifIndex 50) (Generation 26)
    Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x4000
    Encapsulation: ATM-VCMUX
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    VCI 0.4
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
```

show interfaces (ATM2, SONET Mode)

```
user@host> show interfaces at-0/3/1
Physical interface: at-0/3/1, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 67
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:5e
  Last flapped   : 2006-03-13 17:46:36 PST (16:01:12 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
    VPI 0
      Flags: Active, OAM, Shaping
      CBR, Peak: 50kbps
      OAM, Period 30 sec, Up count: 10, Down count: 10
      Total down time: 0 sec, Last down: Never
     OAM F4 cell statistics:
      Total received: 4, Total sent: 4
      Loopback received: 4, Loopback sent: 4
      RDI received: 0, RDI sent: 0
      AIS received: 0
  Traffic statistics:
        Input  packets:                    4
        Output packets:                   30
    VPI 10
      Flags: Active
      Total down time: 0 sec, Last down: Never
  Traffic statistics:
        Input  packets:                    0
        Output packets:                    0
  Logical interface at-0/3/1.0 (Index 78) (SNMP ifIndex 77)
    Flags: Point-To-Point Copy-PLP-To-CLP SNMP-Traps 0x4000
    Encapsulation: ATM-SNAP
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 4470
      Flags: None
```

```
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.59.5, Local: 10.0.59.6
    Protocol iso, MTU: 4470
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 10
        Input packets : 0
        Output packets: 0

  Logical interface at-0/3/1.32767 (Index 79) (SNMP ifIndex 76)
    Flags: Point-To-Multipoint Copy-PLP-To-CLP No-Multicast SNMP-Traps 0x4000
    Encapsulation: ATM-VCMUX
    Input packets : 4
    Output packets: 30
    VCI 0.16
      Flags: Active, ILMI
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
        Input packets : 0
        Output packets: 26
    VCI 0.4
      Flags: Active, OAM
      OAM, Period 30 sec, Up count: 10, Down count: 10
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0
        Input packets : 4
        Output packets: 4
     OAM F4 cell statistics:
      Total received: 4, Total sent: 4
      Loopback received: 4, Loopback sent: 4
      RDI received: 0, RDI sent: 0
      AIS received: 0, AIS sent: 0
```

show interfaces brief (ATM2, SONET Mode)

```
user@host> show interfaces at-0/3/1 brief
Physical interface: at-0/3/1, Enabled, Physical link is Up
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
```

```
  Device flags   : Present Running
  Link flags     : None

  Logical interface at-0/3/1.0
    Flags: Point-To-Point Copy-PLP-To-CLP SNMP-Traps 0x4000
    Encapsulation: ATM-SNAP
    inet  10.0.59.6        --> 10.0.59.5
    iso
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 10

  Logical interface at-0/3/1.32767
    Flags: Point-To-Multipoint Copy-PLP-To-CLP No-Multicast SNMP-Traps 0x4000
    Encapsulation: ATM-VCMUX
    VCI 0.16
      Flags: Active, ILMI
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
    VCI 0.4
      Flags: Active, OAM
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0
```

show interfaces detail (ATM2, SONET Mode)

```
user@host> show interfaces at-0/3/1 detail
Physical interface: at-0/3/1, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 67, Generation: 22
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:5e
  Last flapped   : 2006-03-13 17:46:36 PST (16:02:39 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  312                    0 bps
```

```
 Output bytes  :                 2952                    0 bps
 Input  packets:                    6                    0 pps
 Output packets:                   50                    0 pps
Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                   44                   44                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont                   6                    6                    0
SONET alarms   : None
SONET defects  : None
  VPI 0
    Flags: Active, OAM, Shaping
    CBR, Peak: 50kbps
    OAM, Period 30 sec, Up count: 10, Down count: 10
    Total down time: 0 sec, Last down: Never
   OAM F4 cell statistics:
    Total received: 6, Total sent: 6
    Loopback received: 6, Loopback sent: 6
    Last received: 00:00:29, Last sent: 00:00:29
    RDI received: 0, RDI sent: 0
    AIS received: 0
      Traffic statistics:
       Input  bytes  :                  312
       Output bytes  :                 2952
       Input  packets:                    6
       Output packets:                   50
  VPI 10
    Flags: Active
    Total down time: 0 sec, Last down: Never
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0

Logical interface at-0/3/1.0 (Index 78) (SNMP ifIndex 77) (Generation 20)
  Flags: Point-To-Point Copy-PLP-To-CLP SNMP-Traps 0x4000
  Encapsulation: ATM-SNAP
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
```

```
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 4470, Generation: 38, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.59.5, Local: 10.0.59.6, Broadcast: Unspecified,
        Generation: 44
    Protocol iso, MTU: 4470, Generation: 39, Route table: 0
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 10
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
  Logical interface at-0/3/1.32767 (Index 79) (SNMP ifIndex 76) (Generation 21)
    Flags: Point-To-Multipoint Copy-PLP-To-CLP No-Multicast SNMP-Traps 0x4000
    Encapsulation: ATM-VCMUX
    Traffic statistics:
     Input  bytes  :                  360
     Output bytes  :                 3302
     Input  packets:                    6
     Output packets:                   50
    Local statistics:
     Input  bytes  :                  360
     Output bytes  :                 3302
     Input  packets:                    6
     Output packets:                   50
    VCI 0.16
```

```
      Flags: Active, ILMI
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                 2640
       Input  packets:                    0
       Output packets:                   44
    VCI 0.4
      Flags: Active, OAM
      OAM, Period 30 sec, Up count: 10, Down count: 10
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                  312
       Output bytes  :                  312
       Input  packets:                    6
       Output packets:                    6
     OAM F4 cell statistics:
      Total received: 6, Total sent: 6
      Loopback received: 6, Loopback sent: 6
      Last received: 00:00:29, Last sent: 00:00:29
      RDI received: 0, RDI sent: 0
      AIS received: 0, AIS sent: 0
```

show interfaces extensive (ATM2, SONET Mode)

```
user@host> show interfaces at-0/3/1 extensive
Physical interface: at-0/3/1, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 67, Generation: 22
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:5e
```

```
Last flapped   : 2006-03-13 17:46:36 PST (16:04:12 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                  520                    0 bps
 Output bytes  :                 4240                    0 bps
 Input  packets:                   10                    0 pps
 Output packets:                   72                    0 pps
Input errors:
  Errors: 0, Drops: 0, Invalid VCs: 0, Framing errors: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,
  Resource errors: 0
Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                   62                   62                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont                  10                   10                    0
SONET alarms   : None
SONET defects  : None
SONET PHY:            Seconds        Count  State
  PLL Lock                  0            0  OK
  PHY Light                 0            0  OK
SONET section:
  BIP-B1                    0            0
  SEF                       0            0  OK
  LOS                       0            0  OK
  LOF                       0            0  OK
  ES-S                      0
  SES-S                     0
  SEFS-S                    0
SONET line:
  BIP-B2                    0            0
  REI-L                     0            0
  RDI-L                     0            0  OK
  AIS-L                     0            0  OK
  BERR-SF                   0            0  OK
  BERR-SD                   0            0  OK
  ES-L                      0
  SES-L                     0
  UAS-L                     0
```

```
  ES-LFE                    0
  SES-LFE                   0
  UAS-LFE                   0
SONET path:
  BIP-B3                    0            0
  REI-P                     0            0
  LOP-P                     0            0  OK
  AIS-P                     0            0  OK
  RDI-P                     0            0  OK
  UNEQ-P                    1            1  OK
  PLM-P                     0            0  OK
  ES-P                      1
  SES-P                     1
  UAS-P                     0
  ES-PFE                    0
  SES-PFE                   0
  UAS-PFE                   0
Received SONET overhead:
  F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, C2(cmp) : 0x13, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00
Transmitted SONET overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, F2      : 0x00, Z3      : 0x00
  Z4      : 0x00
ATM status:
  HCS state:     Sync
  LOC      :       OK
ATM Statistics:
  Uncorrectable HCS errors: 0, Correctable HCS errors: 0,
  Tx cell FIFO overruns: 0, Rx cell FIFO overruns: 0,
  Rx cell FIFO underruns: 0, Input cell count: 0, Output cell count: 0,
  Output idle cell count: 0, Output VC queue drops: 0, Input no buffers: 0,
  Input length errors: 0, Input timeouts: 0, Input invalid VCs: 0,
  Input bad CRCs: 0, Input OAM cell no buffers: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
  VPI 0
    Flags: Active, OAM, Shaping
    CBR, Peak: 50kbps
    OAM, Period 30 sec, Up count: 10, Down count: 10
    Total down time: 0 sec, Last down: Never
   OAM F4 cell statistics:
```

```
      Total received: 10, Total sent: 10
      Loopback received: 10, Loopback sent: 10
      Last received: 00:00:02, Last sent: 00:00:02
      RDI received: 0, RDI sent: 0
      AIS received: 0
        Traffic statistics:
         Input  bytes  :                  520
         Output bytes  :                 4240
         Input  packets:                   10
         Output packets:                   72
    VPI 10
      Flags: Active
      Total down time: 0 sec, Last down: Never
        Traffic statistics:
         Input  bytes  :                    0
         Output bytes  :                    0
         Input  packets:                    0
         Output packets:                    0

  Logical interface at-0/3/1.0 (Index 78) (SNMP ifIndex 77) (Generation 20)
    Flags: Point-To-Point Copy-PLP-To-CLP SNMP-Traps 0x4000
    Encapsulation: ATM-SNAP
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 4470, Generation: 38, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.59.5, Local: 10.0.59.6, Broadcast: Unspecified,
        Generation: 44
    Protocol iso, MTU: 4470, Generation: 39, Route table: 0
```

```
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 10
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0

  Logical interface at-0/3/1.32767 (Index 79) (SNMP ifIndex 76) (Generation 21)
    Flags: Point-To-Multipoint Copy-PLP-To-CLP No-Multicast SNMP-Traps 0x4000
    Encapsulation: ATM-VCMUX
    Traffic statistics:
     Input  bytes  :                  660
     Output bytes  :                 5473
     Input  packets:                   11
     Output packets:                   83
    Local statistics:
     Input  bytes  :                  660
     Output bytes  :                 5473
     Input  packets:                   11
     Output packets:                   83
    VCI 0.16
      Flags: Active, ILMI
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                 4320
       Input  packets:                    0
       Output packets:                   72
    VCI 0.4
      Flags: Active, OAM
      OAM, Period 30 sec, Up count: 10, Down count: 10
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0
      ATM per-VC transmit statistics:
```

```
        Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                  572
       Output bytes  :                  572
       Input  packets:                   11
       Output packets:                   11
     OAM F4 cell statistics:
      Total received: 11, Total sent: 11
      Loopback received: 11, Loopback sent: 11
      Last received: 00:00:18, Last sent: 00:00:18
      RDI received: 0, RDI sent: 0
      AIS received: 0, AIS sent: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces demux0 (Demux Interfaces)

IN THIS SECTION

- Syntax | 20497
- Description | 20498
- Options | 20498
- Required Privilege Level | 20498
- Output Fields | 20499
- Sample Output | 20511
- Release Information | 20516

Syntax

```
show interfaces demux0.logical-interface-number
<brief | detail | extensive | terse>
<descriptions>
```

```
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(MX Series and M Series routers only) Display status information about the specified demux interface.

Display status information about the specified demux MPLS subscriber logical interface.

Options

none	Display standard information about the specified demux interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

MPLS Subscriber Logical Interfaces	Display standard information about the specified demux interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 618 on page 20499 lists the output fields for the `show interfaces` demux0 (Demux Interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 618: show interfaces demux0 (Demux Interfaces) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical interface`	Name of the physical interface.	`brief detail extensive` none
`Interface index`	Index number of the physical interface, which reflects its initialization sequence.	`brief detail extensive` none
`Enabled`	State of the interface. Possible values are described in the "Enabled Field" section under Common Output Fields Description.	`brief detail extensive` none
`Physical link`	Status of the physical link (`Up` or `Down`).	`detail extensive` none
`Admin`	Administrative state of the interface (`Up` or `Down`).	`terse`
`Interface index`	Index number of the physical interface, which reflects its initialization sequence.	`detail extensive` none
`Link`	Status of the physical link (`Up` or `Down`).	`terse`
`Targeting summary`	Status of aggregated Ethernet links that are configured with targeted distribution (`primary` or `backup`)	`extensive`
Bandwidth	Bandwidth allocated to the aggregated Ethernet links that are configured with targeted distribution.	`extensive`
`Proto`	Protocol family configured on the interface.	`terse`

Table 618: show interfaces demux0 (Demux Interfaces) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Type	Type of interface. Software-Pseudo indicates a standard software interface with no associated hardware device.	brief detail extensive none
Link-level type	Encapsulation being used on the physical interface.	brief detail extensive
MTU	Maximum transmission unit size on the physical interface.	brief detail extensive
Clocking	Reference clock source: Internal (1) or External (2).	brief detail extensive
Speed	Speed at which the interface is running.	brief detail extensive
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under Common Output Fields Description.	brief detail extensive none
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under Common Output Fields Description.	brief detail extensive none
Link type	Data transmission type.	detail extensive none
Link flags	Information about the link. Possible values are described in the “Link Flags” section under Common Output Fields Description.	detail extensive none

Table 618: show interfaces demux0 (Demux Interfaces) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Physical info`	Information about the physical interface.	`detail extensive`
`Hold-times`	Current interface hold-time up and hold-time down, in milliseconds.	`detail extensive`
`Current address`	Configured MAC address.	`detail extensive`
`Hardware address`	Hardware MAC address.	`detail extensive`
`Alternate link address`	Backup address of the link.	`detail extensive`
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second`*`:`*`timezone`* (*`hour`*`:`*`minute`*`:`*`second`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive` none
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`

Table 618: show interfaces demux0 (Demux Interfaces) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface. • `IPv6 transit statistics`—Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. **NOTE**: These fields include dropped traffic and exception traffic, as those fields are not separately defined. • `Input bytes`—Number of bytes received on the interface • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`

Table 618: show interfaces demux0 (Demux Interfaces) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input errors`	Input errors on the interface whose definitions are as follows: • `Errors`—Sum of the incoming frame terminations and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Giants`—Number of frames received that are larger than the giant packet threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • `Resource errors`—Sum of transmit drops.	`extensive`
`Input Rate`	Input rate in bits per second (bps) and packets per second (pps).	none

Table 618: show interfaces demux0 (Demux Interfaces) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • `Errors`—Sum of the outgoing frame terminations and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `MTU errors`—Number of packets whose size exceeded the MTU of the interface. • `Resource errors`—Sum of transmit drops.	`extensive`
`Output Rate`	Output rate in bps and pps.	none
Logical Interface		
`Logical interface`	Name of the logical interface.	`brief detail extensive` none
`Index`	Index number of the logical interface, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP interface index number for the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail`

Table 618: show interfaces demux0 (Demux Interfaces) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under Common Output Fields Description.	`brief detail extensive` none
`Encapsulation`	Encapsulation on the logical interface.	`brief extensive` none
`ACI VLAN: Dynamic Profile`	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying demux interface to create dynamic VLAN subscriber interfaces based on ACI information.	`brief detail extensive` none
`Demux`	Specific IP demultiplexing (demux) values: • `Underlying interface`—The underlying interface that the demux interface uses. • `Index`—Index number of the logical interface. • `Family`—Protocol family configured on the logical interface. • `Source prefixes, total`—Total number of source prefixes for the underlying interface. • `Destination prefixes, total`—Total number of destination prefixes for the underlying interface. • `Prefix`—`inet` family prefix.	`detail extensive` none
`protocol-family`	Protocol family configured on the logical interface.	`brief`

Table 618: show interfaces demux0 (Demux Interfaces) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the specified interface set. • `Input bytes, Output bytes`—Number of bytes received and transmitted on the interface set. • `Input packets, Output packets`—Number of packets received and transmitted on the interface set. • `IPv6 transit statistics`—Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. **NOTE**: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`
`Local statistics`	Number of transit bytes and packets received and transmitted on the local interface. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`

Table 618: show interfaces demux0 (Demux Interfaces) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Transit statistics`	Number and rate of bytes and packets transiting the switch. **NOTE**: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`
`IPv6 Transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. **NOTE**: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`
`Input packets`	Number of packets received on the interface.	none
`Output packets`	Number of packets transmitted on the interface.	none
`Protocol`	Protocol family. Possible values are described in the “Protocol Field” section under Common Output Fields Description.	`detail extensive` none
`MTU`	Maximum transmission unit size on the logical interface.	`detail extensive` none

Table 618: show interfaces demux0 (Demux Interfaces) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Maximum labels`	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	`detail` `extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`
`Route table`	Route table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail` `extensive`
`Flags`	Information about protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	`detail` `extensive` none
`Mac-Validate Failures`	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	`detail` `extensive` none
`Addresses, Flags`	Information about the address flags. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	`detail` `extensive` none
`Destination`	IP address of the remote side of the connection.	`detail` `extensive` `statistics` none
`Local`	IP address of the logical interface.	`detail` `extensive` `terse` none
`Remote`	IP address of the remote interface.	`terse`
`Broadcast`	Broadcast address of the logical interlace.	`detail` `extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`

Table 618: show interfaces demux0 (Demux Interfaces) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Link`	Name of the physical interfaces for member links in an aggregated Ethernet bundle for a PPPoE over aggregated Ethernet configuration. PPPoE traffic goes out on these interfaces.	`detail extensive` none
`Dynamic-profile`	Name of the PPPoE dynamic profile assigned to the underlying interface.	`detail extensive` none
`Service Name Table`	Name of the PPPoE service name table assigned to the PPPoE underlying interface.	`detail extensive` none
`Max Sessions`	Maximum number of dynamic PPPoE logical interfaces that the router can activate on the underlying interface.	`detail extensive` none
`Duplicate Protection`	State of duplicate protection: `On` or `Off`. Duplicate protection prevents the activation of another dynamic PPPoE logical interface on the same underlying interface when a dynamic PPPoE logical interface for a client with the same MAC address is already active on that interface.	`detail extensive` none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: `On` or `Off`. When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	`detail extensive` none
`AC Name`	Name of the access concentrator.	`detail extensive` none

The **show interfaces (MPLS Interfaces) Output Fields** table lists the output fields for the `show interfaces` *`ifl-name`* command. Output fields are listed in the approximate order in which they appear.

Table 619: show interfaces (MPLS Subscriber Logical Interfaces) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		

Table 619: show interfaces (MPLS Subscriber Logical Interfaces) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Link`	Status of the physical link (Up or Down).	`terse` none
`Targeting summary`	Status of LT links that are configured with targeted distribution (primary or backup)	`extensive` none
Logical Interface		
`Logical interface`	Name of the logical interface.	`brief detail extensive` none
`Index`	Index number of the logical interface, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP interface index number for the logical interface.	`detail extensive` none
`Flags`	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under Common Output Fields Description.	`brief detail extensive` none
`Encapsulation`	Type of encapsulation configured on the logical interface.	`brief extensive` none
`Demux`	Specific IP demultiplexing (demux) values: • Underlying interface—The underlying interface that the demux interface uses. • Index—Index number of the logical interface. • Prefix—inet family prefix. • Family—Protocol family configured on the logical interface. • Source prefixes, total—Total number of source prefixes for the underlying interface.	`detail extensive` none

Sample Output

show interfaces demux0 (Demux)

```
user@host> show interfaces demux0
Physical interface: demux0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 79, Generation: 129
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Not Applicable, Clocking: 1,
  Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    Not Applicable
   Output packets:                    Not Applicable
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   Not Applicable
    Output packets:                   Not Applicable
  Input errors:
    Errors: Not Applicable, Drops: Not Applicable, Framing errors: Not Applicable,
    Runts: Not Applicable, Giants: Not Applicable,
    Policed discards: Not Applicable, Resource errors: Not Applicable
  Output errors:
    Carrier transitions: Not Applicable, Errors: Not Applicable, Drops: Not Applicable,
    MTU errors: Not Applicable,
    Resource errors: Not Applicable
Logical interface demux0.0 (Index 87) (SNMP ifIndex 84) (Generation 312)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Demux:
     Underlying interface: ge-2/0/1.0 (Index 74)
      Family Inet Source prefixes, total 1
```

```
      Prefix: 203.0.113/24
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                 1554
     Input  packets:                    0
     Output packets:                   37
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                 1554
     Input  packets:                    0
     Output packets:                   37
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: 1500, Generation: 395, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.13, Broadcast: 203.0.113.255,
        Generation: 434
```

show interfaces demux0 statistics

```
user@host> show interfaces demux0 statistics
Physical interface: demux0, Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 501
  Type: Software-Pseudo, MTU: Not Applicable
  Device flags   : Present Running
  Interface Specific flags: Internal: 0x200
```

```
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
  Statistics last cleared: 2024-02-06 16:49:13 PST (00:08:48 ago)
    Input packets : Not Applicable
    Output packets: Not Applicable
  Input errors: Not Applicable, Output errors: Not Applicable

  Logical interface demux0.1 (Index 337) (SNMP ifIndex 693)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.512 0x8100.1 ] In(pop-swap .1) Out(swap-push
0x0000.512 .1)  Encapsulation: VLAN-CCC
    Demux:
      Underlying interface: ae1 (Index 151)
    Link:
      xe-3/0/4
      xe-3/1/4
      xe-5/0/11
      xe-5/0/19
      xe-5/1/5
      xe-5/1/7
    Bandwidth: 20Gbps
    Input packets : 0
    Output packets: 0
    Protocol ccc, MTU: 2008
      Flags: Is-Primary, User-MTU

  Logical interface demux0.2 (Index 339) (SNMP ifIndex 6306)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.512 0x8100.2 ] In(pop-swap .2) Out(swap-push
0x0000.512 .2)  Encapsulation: VLAN-CCC
    Demux:
      Underlying interface: ae2 (Index 152)
    Link:
      xe-3/0/0
      xe-3/1/5
      xe-5/0/17
      xe-5/1/1
    Bandwidth: 20Gbps
    Input packets : 0
    Output packets: 0
    Protocol ccc, MTU: 2008
      Flags: User-MTU
```

```
  Logical interface demux0.3 (Index 323) (SNMP ifIndex 6307)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.512 0x8100.3 ] In(pop-swap .3) Out(swap-push
0x0000.512 .3)  Encapsulation: VLAN-CCC
    Demux:
      Underlying interface: ae1 (Index 151)
    Link:
      xe-3/0/4
      xe-3/1/4
      xe-5/0/11
      xe-5/0/19
      xe-5/1/5
      xe-5/1/7
    Bandwidth: 0
    Input packets : 5281718
    Output packets: 5281718
    Protocol ccc, MTU: 2008
      Flags: User-MTU
```

show interfaces demux0 (PPPoE over Aggregated Ethernet)

```
user@host> show interfaces demux0.100
Logical interface demux0.100 (Index 76) (SNMP ifIndex 61160)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]
    Encapsulation: ENET2
    Demux:
      Underlying interface: ae0 (Index 199)
    Link:
      ge-1/0/0
      ge-1/1/0
    Input packets : 0
    Output packets: 0
    Protocol pppoe
      Dynamic Profile: pppoe-profile,
      Service Name Table: service-table1,
      Max Sessions: 100, Duplicate Protection: On,
      Direct Connect: Off,
      AC Name: pppoe-server-1
```

show interfaces demux0 extensive (Targeted Distribution for Aggregated Ethernet Links)

```
user@host> show interfaces demux0.1073741824 extensive

  Logical interface demux0.1073741824 (Index 75) (SNMP ifIndex 558) (Generation 346)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Demux:
      Underlying interface: ae0 (Index 201)
    Link:
      ge-1/0/0
      ge-1/1/0
      ge-2/0/7
      ge-2/0/8
    Targeting summary:
      ge-1/1/0, primary, Physical link is Up
      ge-2/0/8, backup, Physical link is Up
    Bandwidth: 1000mbps
```

show interfaces demux0 (ACI Interface Set Configured)

```
user@host> show interfaces demux0.1073741827
  Logical interface demux0.1073741827 (Index 346) (SNMP ifIndex 527)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1802 0x8100.302 ]  Encapsulation: ENET2
    Demux: Source Family Inet
    ACI VLAN:
      Dynamic Profile: aci-vlan-set-profile
    Demux:
      Underlying interface: ge-1/0/0 (Index 138)
    Input packets : 18
    Output packets: 16
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re, Unnumbered
      Donor interface: lo0.0 (Index 322)
      Preferred source address: 203.0.113.202
      Addresses, Flags: Primary Is-Default Is-Primary
        Local: 203.0.113.119
    Protocol pppoe
      Dynamic Profile: aci-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
```

```
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
```

show interfaces demux0 *ifl-name*

```
 user@host> show interfaces demux0.3221225482
    Logical interface demux0.3221225482 (Index 536871015) (SNMP ifIndex 200000103)
    Flags: Up Encapsulation: ENET2
    Interface set: iflset100100
    Demux:
      Underlying interface: ps1.3221225480 (Index 536871011)
      Family Inet Source prefixes, total 1
          Prefix: 192.168.0.5/32
    Link:
      lt-1/0/0.32767
      lt-4/1/0.32767
      lt-1/1/0.32767
      lt-4/0/0.32767
    Targeting summary:
      lt-1/0/0, primary, Physical link is Up
      lt-4/1/0, backup, Physical link is Up
user@host> show interfaces demux0
Physical interface: demux0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 502
  Type: Software-Pseudo, MTU: Not Applicable, Clocking: 1
  Device flags   : Present Running
  Interface Specific flags: Internal: 0x200
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
    Input packets : Not Applicable
    Output packets: Not Applicable
```

Release Information

Command introduced in Junos OS Release 9.0.

Command introduced for MPLS subscriber logical interfaces in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

show interfaces destination-class

IN THIS SECTION

- Syntax | 20517
- Description | 20517
- Options | 20517
- Additional Information | 20518
- Required Privilege Level | 20518
- Output Fields | 20518
- Sample Output | 20518
- Release Information | 20520

Syntax

```
show interfaces destination-class
(all / destination-class-name logical-interface-name)
```

Description

Display information about interfaces grouped by destination class.

Options

all	Display information about all configured destination classes.
destination-class-name	Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

logical interface-name Name of a logical interface.

Additional Information

For interfaces that carry IPv4, IPv6, or Multiprotocol Label Switching (MPLS) traffic, you can maintain packet counts based on the entry and exit points for traffic passing through your network. Entry and exit points are identified by source and destination prefixes grouped into sets defined as source classes and destination classes. For more information, see the Junos OS Network Interfaces Library for Routing Devices.

Required Privilege Level

view

Output Fields

Table 620 on page 20518 lists the output fields for the `show interfaces destination-class` command. Output fields are listed in the approximate order in which they appear.

Table 620: show interfaces destination-class Output Fields

Field Name	Field Description
Logical interface	Name of the logical interface.
Destination class	Name of destination class usage (DCU) counters per class for this interface.
Packets	Packets going to designated user-selected prefixes.
Bytes	Bytes going to designated user-selected prefixes.

Sample Output

show interfaces destination-class all

```
user@host> show interfaces destination-class all
  Logical interface .local..1
```

```
Logical interface .local..2

Logical interface fxp0.0

Logical interface fxp1.0

Logical interface lo0.16384

Logical interface lo0.16385

Logical interface lo0.0

Logical interface .local..3

Logical interface .local..4

Logical interface .local..5

Logical interface .local..6

Logical interface .local..7

Logical interface .local..8

Logical interface .local..9

Logical interface .local..10

Logical interface lo0.3

Logical interface pfh-0/0/0.16383

Logical interface fe-0/1/0.0
  Protocol inet
                                            Packets                Bytes
    Destination class           (packet-per-second)    (bits-per-second)

                        SILVER1                   0                    0
                                (                 0) (                 0)
                        SILVER2                   0                    0
                                (                 0) (                 0)
                        SILVER3                   0                    0
```

```
                                (                  0) (                 0)
                        v4-dest                    0                   0
                                (                  0) (                 0)
  Protocol inet6
                                             Packets               Bytes
    Destination class            (packet-per-second)    (bits-per-second)

                        SILVER1                    0                   0
                                (                  0) (                 0)
                        SILVER2                    0                   0
                                (                  0) (                 0)
                        SILVER3                    0                   0
                                (                  0) (                 0)
                        v4-dest                    0                   0
                                (                  0) (                 0)

Logical interface fe-0/1/1.0

Logical interface fe-0/1/2.0
  Description: CE1-to-PE2

Logical interface ge-0/3/0.0
  Description: CE1-to-PE1

Logical interface ge-0/3/2.0
  Description: CE2-to-PE1

Logical interface pc-0/3/0.16383

Logical interface lt-1/2/0.3
  Description: LS3->LS2

Logical interface lt-1/2/0.5
  Description: LS3->LS1

Logical interface sp-1/2/0.16383
```

Release Information

Command introduced before Junos OS Release 7.4.

all option introduced in Junos OS Release 8.0.

show interfaces detail

IN THIS SECTION

- Syntax | 20521
- Description | 20521
- Options | 20521
- Additional Information | 20521
- Required Privilege Level | 20521
- Output Fields | 20522
- Sample Output | 20522
- Release Information | 20525

Syntax

```
show interfaces detail
```

Description

Display detailed information about all interfaces configured on the router.

Options

This command has no options.

Additional Information

In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.

Required Privilege Level

view

Output Fields

For more information, see the output fields table for the particular interface type in which you are interested. For information about destination class and source class statistics, see the “Destination Class Field” section and the “Source Class Field” section under *Common Output Fields Description*. For sample output for specific interfaces, see the other chapters in this manual.

Sample Output

show interfaces detail (SONET)

```
user@host> show interfaces so-1/1/0 detail
Physical interface: so-1/1/0, Enabled, Physical link is Up
  Interface index: 142, SNMP ifIndex: 47, Generation: 143
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC12,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 1934 (last seen 17:35:39 ago)
    Output: 1927 (last sent 17:35:48 ago)
  LCP state: Down
  NCP state: inet: Down, inet6: Not-configured, iso: Down, mpls: Not-configured
  CHAP state: Closed
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2006-04-19 15:22:33 PDT (05:25:55 ago)
  Statistics last cleared: 2006-04-18 03:58:02 PDT (1d 16:50 ago)
  Traffic statistics:
   Input  bytes  :              7910882                    0 bps
   Output bytes  :              5632131                    0 bps
   Input  packets:                89460                    0 pps
   Output packets:               116043                    0 pps
  SONET alarms   : None
  SONET defects  : None
  Logical interface so-1/1/0.0 (Index 69) (SNMP ifIndex 61) (Generation 138)
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
      Protocol inet, MTU: 4470, Generation: 156, Route table: 2
      Flags: Protocol-Down, SCU-out
```

```
                                                Packets                Bytes
      Source class                 (packet-per-second)    (bits-per-second)
                            gold1                    0                    0
                                  (                  0) (                 0)
                            gold2                    0                    0
                                  (                  0) (                 0)
                            gold3                    0                    0
                                  (                  0) (                 0)
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.27.248/24, Local: 10.27.248.1, Broadcast: 10.27.248.255,
        Generation: 152
    Protocol iso, MTU: 4470, Generation: 157, Route table: 2
      Flags: Protocol-Down, Is-Primary
```

show interfaces detail (MX Series Routers)

```
user@host> show interfaces xe-0/0/0 detail
Physical interface: xe-0/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 592, Generation: 148
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Enabled
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 08:81:f4:82:a3:f0, Hardware address: 08:81:f4:82:a3:f0
  Last flapped   : 2013-10-26 03:20:40 test (1w6d 04:47 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Egress queues: 8 supported, 4 in use
```

```
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                               0                    0                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                            79
    Errored blocks                        79
  Interface transmit statistics: Disabled
```

show interfaces detail (Junos OS Evolved)

```
user@host> show interfaces et-0/0/2:2 detail
Physical interface: et-0/0/2:2, Enabled, Physical link is Up
  Interface index: 1164, SNMP ifIndex: 592, Generation: 275
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present
  Interface flags: SNMP-Traps
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 28:8a:1c:68:84:ca, Hardware address: 28:8a:1c:68:84:ca
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
```

```
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Input errors: 3, Output errors: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 3, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Active alarms  : None
Active defects : None
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0                               0                    0                    0
  1                               0                    0                    0
  2                               0                    0                    0
  3                               0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
PCS statistics                      Seconds
  Bit errors                             0
  Errored blocks                         0
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces diagnostics optics

IN THIS SECTION

- Syntax | 20526
- Description | 20526
- Options | 20526
- Required Privilege Level | 20526
- Output Fields | 20527
- Sample Output | 20532
- Release Information | 20535

Syntax

```
show interfaces diagnostics optics interface-name
```

Description

Display diagnostics data and alarms for Gigabit Ethernet, 10-Gibabit Ethernet, and QSFP+ optical transceivers installed in a QFX Series product. The information provided by this command is known as digital optical monitoring (DOM) information.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.

Options

interface-name Name of the interface associated with the port in which the transceiver is installed.

Required Privilege Level

view

Output Fields

lists the output fields for the `show interfaces diagnostics optics` command. Output fields are listed in the approximate order in which they appear.

Table 621: show interfaces diagnostics optics Output Fields

Field Name	Field Description
Physical interface	Displays the name of the physical interface.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage (Not available for XFP transceivers)	Displays the voltage, in volts.
Laser rx power (Not available for SFP and SFP+ transceivers)	Displays the laser received optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Receiver signal average optical power (Not available for XFP transceivers)	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is **On** or **Off**.
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is **On** or **Off**.

Table 621: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Laser bias current high warning	Displays whether the laser bias power setting high warning is **On** or **Off**.
Laser bias current low warning	Displays whether the laser bias power setting low warning is **On** or **Off**.
Laser output power high alarm	Displays whether the laser output power high alarm is **On** or **Off**.
Laser output power low alarm	Displays whether the laser output power low alarm is **On** or **Off**.
Laser output power high warning	Displays whether the laser output power high warning is **On** or **Off**.
Laser output power low warning	Displays whether the laser output power low warning is **On** or **Off**.
Module temperature high alarm	Displays whether the module temperature high alarm is **On** or **Off**.
Module temperature low alarm	Displays whether the module temperature low alarm is **On** or **Off**.
Module temperature high warning	Displays whether the module temperature high warning is **On** or **Off**.
Module temperature low warning	Displays whether the module temperature low warning is **On** or **Off**.
Module voltage high alarm (Not available for XFP transceivers)	Displays whether the module voltage high alarm is **On** or **Off**.
Module voltage low alarm (Not available for XFP transceivers)	Displays whether the module voltage low alarm is **On** or **Off**.
Module voltage high warning (Not available for XFP transceivers)	Displays whether the module voltage high warning is **On** or **Off**.

Table 621: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Module voltage low warning (Not available for XFP transceivers)	Displays whether the module voltage low warning is **On** or **Off**.
Laser rx power high alarm	Displays whether the receive laser power high alarm is **On** or **Off**.
Laser rx power low alarm	Displays whether the receive laser power low alarm is **On** or **Off**.
Laser rx power high warning	Displays whether the receive laser power high warning is **On** or **Off**.
Laser rx power low warning	Displays whether the receive laser power low warning is **On** or **Off**.
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Module not ready alarm (Not available for SFP and SFP+ transceivers)	Displays whether the module not ready alarm is **On** or **Off**. When the output is **On**, the module has an operational fault.
Module power down alarm (Not available for SFP and SFP+ transceivers)	Displays whether the module power down alarm is **On** or **Off**. When the output is **On**, the module is in a limited power mode, low for normal operation.
Tx data not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx data not ready alarm is **On** or **Off**.
Tx not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx not ready alarm is **On** or **Off**.

Table 621: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Tx laser fault alarm (Not available for SFP and SFP+ transceivers)	Laser fault condition. Displays whether the Tx laser fault alarm is **On** or **Off**.
Tx CDR loss of lock alarm (Not available for SFP and SFP+ transceivers)	Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays whether the Tx CDR loss of lock alarm is **On** or **Off**.
Rx not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the receive path. Displays whether the Rx not ready alarm is **On** or **Off**.
Rx loss of signal alarm (Not available for SFP and SFP+ transceivers)	Receive loss of signal alarm. When **on**, indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is **On** or **Off**.
Rx CDR loss of lock alarm (Not available for SFP and SFP+ transceivers)	Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays whether the Rx CDR loss of lock alarm is **On** or **Off**.
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.

Table 621: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.
Module voltage high alarm threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage high warning.

Table 621: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Module voltage low warning threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser Rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser Rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser Rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser Rx power low warning.

Sample Output

show interfaces diagnostics optics xe-0/0/1 (SFP+ Transceiver)

```
user@host> show interfaces diagnostics optics
xe-0/0/1
Physical interface: xe-0/0/1
    Laser bias current                        :  4.968 mA
    Laser output power                        :  0.4940 mW / -3.06 dBm
    Module temperature                        :  27 degrees C / 81 degrees F
    Module voltage                            :  3.2310 V
    Receiver signal average optical power     :  0.0000
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
```

```
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
```

show interfaces diagnostics optics node1:xe-0/0/1 (SFP+ Transceiver)

```
user@host> show interfaces diagnostics optics
node1:xe-0/0/1
Physical interface: node1:xe-0/0/1
    Laser bias current                        :  4.968 mA
    Laser output power                        :  0.4940 mW / -3.06 dBm
```

```
Module temperature                        :  27 degrees C / 81 degrees F
Module voltage                            :  3.2310 V
Receiver signal average optical power     :  0.0000
Laser bias current high alarm             :  Off
Laser bias current low alarm              :  Off
Laser bias current high warning           :  Off
Laser bias current low warning            :  Off
Laser output power high alarm             :  Off
Laser output power low alarm              :  Off
Laser output power high warning           :  Off
Laser output power low warning            :  Off
Module temperature high alarm             :  Off
Module temperature low alarm              :  Off
Module temperature high warning           :  Off
Module temperature low warning            :  Off
Module voltage high alarm                 :  Off
Module voltage low alarm                  :  Off
Module voltage high warning               :  Off
Module voltage low warning                :  Off
Laser rx power high alarm                 :  Off
Laser rx power low alarm                  :  On
Laser rx power high warning               :  Off
Laser rx power low warning                :  On
Laser bias current high alarm threshold   :  10.500 mA
Laser bias current low alarm threshold    :  2.000 mA
Laser bias current high warning threshold :  9.000 mA
Laser bias current low warning threshold  :  2.500 mA
Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
Module temperature high warning threshold :  70 degrees C / 158 degrees F
Module temperature low warning threshold  :  0 degrees C / 32 degrees F
Module voltage high alarm threshold       :  3.630 V
Module voltage low alarm threshold        :  2.970 V
Module voltage high warning threshold     :  3.465 V
Module voltage low warning threshold      :  3.135 V
Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
```

```
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Monitoring Interface Status and Traffic

Installing a Transceiver in a QFX Series Device

Removing a Transceiver from a QFX Series Device

Junos OS Network Interfaces Library for Routing Devices

show interfaces diagnostics optics

IN THIS SECTION

- Syntax | 20536
- Description | 20536
- Options | 20536
- Required Privilege Level | 20536
- Output Fields | 20536
- Sample Output | 20559
- Sample Output | 20561
- Sample Output | 20562
- Sample Output | 20563
- Sample Output | 20565
- Release Information | 20569

Syntax

```
show interfaces diagnostics optics interface-name
```

Description

Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP, SFP+, XFP, QSFP+, or CFP) installed in EX Series Switches or QFX Series Switches. The information provided by this command is known as digital optical monitoring (DOM) information.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.

Options

interface-name	Name of the interface associated with the port in which the transceiver is installed: ge-*fpc/pic/port*, xe-*fpc/pic/port*, or et-*fpc/pic/port*.

Required Privilege Level

View

Output Fields

Table 622 on page 20537 lists the output fields for the `show interfaces diagnostics optics` command. Output fields are listed in the approximate order in which they appear.

Transceivers generate monitoring data such as power supply voltage, transceiver temperature, TEC current, receive optical power, laser output power, laser temperature, and laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.

Warning flags associated with power supply voltage, transceiver temperature, TEC current, receive optical power, laser output power, laser temperature, and laser bias current indicate conditions outside the normal limits; however, that need not cause immediate link failures.

Alarm flags associated with power supply voltage, transceiver temperature, TEC current, receive optical power, laser output power, laser temperature, and laser bias current indicate conditions likely, but not

necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.

Table 622: show interfaces diagnostics optics Output Fields

Field Name	Field Description
`Physical interface`	Displays the name of the physical interface.
`Lane 'x'`	Displays an individual Tx/Rx data transmission channel "x" associated with the given physical interface device.
`Laser bias current`	Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents.
`Laser output power` (Availability is dependent on SFP vendor)	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
`Laser temperature` (Availability is dependent on SFP vendor)	Displays the laser temperature, in Celsius and Fahrenheit.
`Module temperature`	Displays the temperature, in Celsius and Fahrenheit.
`Module voltage` (Availability is dependent on SFP vendor)	Displays the voltage in Volts.
`Laser rx power` (Availability is dependent on SFP vendor)	Displays the laser received optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Receiver signal average optical power` (Availability is dependent on SFP vendor)	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
`Laser bias current high alarm`	Displays whether the laser bias power setting high alarm is `On` or `Off`. Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser bias current indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser bias current low alarm`	Displays whether the laser bias power setting low alarm is `On` or `Off`. Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser bias current indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser bias current high warning`	Displays whether the laser bias power setting high warning is `On` or `Off`. Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser bias current indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Laser bias current low warning`	Displays whether the laser bias power setting low warning is `On` or `Off`. Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser bias current indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Laser output power high alarm` (Availability is dependent on SFP vendor)	Displays whether the laser output power high alarm is `On` or `Off`. Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser output power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser output power low alarm` (Availability is dependent on SFP vendor)	Displays whether the laser output power low alarm is `On` or `Off`. Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser output power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser output power high warning` (Availability is dependent on SFP vendor)	Displays whether the laser output power high warning is `On` or `Off`. Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser output power indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Laser output power low warning` (Availability is dependent on SFP vendor)	Displays whether the laser output power low warning is `On` or `Off`. Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser output power indicate conditions outside the normal limits; however, that need not cause immediate link failures.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser temperature high alarm` (Availability is dependent on SFP vendor)	Displays whether the laser temperature high alarm is `On` or `Off`. Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser temperature low alarm` (Availability is dependent on SFP vendor)	Displays whether the laser temperature low alarm is `On` or `Off`. Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser temperature high warning` (Availability is dependent on SFP vendor)	Displays whether the laser temperature high warning is `On` or `Off`. Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser temperature low warning` (Availability is dependent on SFP vendor)	Displays whether the laser temperature low warning is `On` or `Off`. Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Module temperature high alarm` (Availability is dependent on SFP vendor)	Displays whether the module temperature high alarm is `On` or `Off`. Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with transceiver temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Module temperature low alarm` (Availability is dependent on SFP vendor)	Displays whether the module temperature low alarm is `On` or `Off`. Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with transceiver temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Module temperature high warning` (Availability is dependent on SFP vendor)	Displays whether the module temperature high warning is `On` or `Off`. Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with transceiver temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Module temperature low warning` (Availability is dependent on SFP vendor)	Displays whether the module temperature low warning is `On` or `Off`. Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with transceiver temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Module voltage high alarm` (Availability is dependent on SFP vendor)	Displays whether the module voltage high alarm is `On` or `Off`. Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with power supply voltage indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Module voltage low alarm` (Availability is dependent on SFP vendor)	Displays whether the module voltage low alarm is `On` or `Off`. Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with power supply voltage indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Module voltage high warning` (Availability is dependent on SFP vendor)	Displays whether the module voltage high warning is `On` or `Off`. Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with power supply voltage indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Module voltage low warning` (Availability is dependent on SFP vendor)	Displays whether the module voltage low warning is `On` or `Off`. Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with power supply voltage indicate conditions outside the normal limits; however, that need not cause immediate link failures.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser rx power high alarm` (Availability is dependent on SFP vendor)	Displays whether the receive laser power high alarm is `On` or `Off`. Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with receive optical power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser rx power low alarm` (Availability is dependent on SFP vendor)	Displays whether the receive laser power low alarm is `On` or `Off`. Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with receive optical power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser rx power high warning` (Availability is dependent on SFP vendor)	Displays whether the receive laser power high warning is `On` or `Off`. Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with receive optical power indicate conditions outside the normal limits; however, that need not cause immediate link failures.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser rx power low warning` (Availability is dependent on SFP vendor)	Displays whether the receive laser power low warning is `On` or `Off`. Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with receive optical power indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Laser bias current high alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser bias current high alarm. Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser bias current indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Module not ready alarm` (Availability is dependent on SFP vendor)	Displays whether the module not ready alarm is `On` or `Off`. When the output is `On`, the module has an operational fault.
`Module low power alarm` (Availability is dependent on SFP vendor)	Displays whether the module low power alarm is `On` or `Off`. The host might use the module power mode to limit module power consumption to one of two thresholds—low power mode or high power mode. When operating in low power mode, the maximum power consumption of the module will remain below a defined maximum. Active module functionality is limited to host-to-module interactions over the management interface. Passive copper cable assemblies might operate in low power mode at all times and are fully functional.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Module initialization incomplete alarm` (Availability is dependent on SFP vendor)	Displays whether the module initialization incomplete alarm is `On` or `Off`. This alarm is generated when the transceiver module initialization is not complete.
`Module fault alarm` (Availability is dependent on SFP vendor)	Displays whether the module fault alarm is `On` or `Off`. This alarm is generated when the transceiver module detects a condition (such as TEC runaway or flash corruption) that could cause damage to the module. This alarm is provided as a notification to the host that a module fault has occurred. The implementation of the module fault is form factor dependent.
`PLD Flash initialization fault alarm` (Availability is dependent on SFP vendor)	Displays whether the PLD Flash initialization fault alarm is `On` or `Off`. This alarm is generated if the initialization of transceiver module internal programmable logic device (PLD) or flash device fails.
`Power supply fault alarm` (Availability is dependent on SFP vendor)	Displays whether the power supply fault alarm is `On` or `Off`. This alarm is generated if the power supply is out of range.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Checksum fault alarm` (Availability is dependent on SFP vendor)	Displays whether the checksum fault alarm is `On` or `Off`. To protect against degradation in the non-volatile storage and potential errors in transfer, the module management interface defines a few registers to let module manufacturers save checksums or check codes. Each checksum or check code is assigned for a small-sized datum derived from a block of digital data. The algorithms of calculation for checksums or check codes are defined by management interface specifications of form factors. For example, the checksum or check code shall be the low order 8 bits of the sum of the contents of that small-sized datum. Prior to using the stored information, module users calculate the checksums or check codes again using the same algorithm. If the results are different from those saved in checksum bytes by module manufacturers, checksum fault alarm is triggered. A Checksum fault means that the content of the small-sized datum covered by this checksum is corrupted or changed.
`Tx laser disabled alarm` (Availability is dependent on SFP vendor)	Displays whether the Tx laser disabled alarm is `On` or `Off`. This alarm is generated when the transmitter optical output is shut down.
`Module power down alarm` (Availability is dependent on SFP vendor)	Displays whether the module power down alarm is `On` or `Off`. When the output is `On`, module is in a limited power mode, low for normal operation.
`Tx data not ready alarm` (Availability is dependent on SFP vendor)	Any condition leading to invalid data on the transmit path. Displays whether the Tx data not ready alarm is `On` or `Off`.
`Tx not ready alarm` (Availability is dependent on SFP vendor)	Any condition leading to invalid data on the transmit path. Displays whether the Tx not ready alarm is `On` or `Off`.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Tx laser fault alarm` (Availability is dependent on SFP vendor)	Laser fault condition. Displays whether the Tx laser fault alarm is `On` or `Off`. This alarm is generated when the transmitter/laser is operating in an abnormal state.
`Tx CDR loss of lock alarm` (Availability is dependent on SFP vendor)	Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays whether the Tx CDR loss of lock alarm is `On` or `Off`. Some transceivers have Clock and Data Recovery (CDR) circuits at the transmit side. The CDR has its own oscillator and uses a phase locked loop (PLL) to calibrate the oscillator and match it to the phase and frequency of the incoming signal. The feedback circuit of PLL eventually locks the local oscillator on to the phase and frequency of the incoming signal. The PLL maintains that lock by continuously monitoring the received signal. If the lock cannot be maintained, the Loss of Lock (LOL) alarm is generated.
`Rx not ready alarm` (Availability is dependent on SFP vendor)	Any condition leading to invalid data on the receive path. Displays whether the Rx not ready alarm is `On` or `Off`.
`Rx loss of signal alarm`	Receive loss of signal alarm. When the output is `On`, indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is `On` or `Off`.
`Rx CDR loss of lock alarm` (Availability is dependent on SFP vendor)	Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays whether the Rx CDR loss of lock alarm is `On` or `Off`. Some transceivers have Clock and Data Recovery (CDR) circuits at the receive side. The CDR has its own oscillator and uses a phase locked loop (PLL) to calibrate the oscillator and match it to the phase and frequency of the incoming signal. The feedback circuit of PLL eventually locks the local oscillator on to the phase and frequency of the incoming signal. The PLL maintains that lock by continuously monitoring the received signal. If the lock cannot be maintained, the Loss of Lock (LOL) alarm is generated.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser bias current low alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser bias current low alarm. Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser bias current indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser bias current high warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser bias current high warning. Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser bias current indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Laser bias current low warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser bias current low warning. Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser bias current indicate conditions outside the normal limits; however, that need not cause immediate link failures.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser output power high alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser output power high alarm. Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser output power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser output power low alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser output power low alarm. Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser output power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser output power high warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser output power high warning. Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser output power indicate conditions outside the normal limits; however, that need not cause immediate link failures.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser output power low warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser output power low warning. Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser output power indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Module temperature high alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the module temperature high alarm. Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Module temperature low alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the module temperature low alarm. Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with laser temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Module temperature high warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the module temperature high warning. Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Module temperature low warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the module temperature low warning. Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with laser temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Module voltage high alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the module voltage high alarm. Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with power supply voltage indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Module voltage low alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the module voltage low alarm. Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with power supply voltage indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Module voltage high warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the module voltage high warning. Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with power supply voltage indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Module voltage low warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the module voltage low warning. Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with power supply voltage indicate conditions outside the normal limits; however, that need not cause immediate link failures.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser rx power high alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser Rx power high alarm. Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with receive optical power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser rx power low alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser Rx power low alarm. Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with receive optical power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser rx power high warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser Rx power high warning. Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with receive optical power indicate conditions outside the normal limits; however, that need not cause immediate link failures.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser rx power low warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser Rx power low warning. Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with receive optical power indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Laser temperature high alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser temperature high alarm, in Celsius and Fahrenheit. Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with transceiver temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.
`Laser temperature low alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser temperature low alarm, in Celsius and Fahrenheit. Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Alarm flags associated with transceiver temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser temperature high warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser temperature high warning, in Celsius and Fahrenheit. Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with transceiver temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`Laser temperature low warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for the laser temperature low warning, in Celsius and Fahrenheit. Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits. Warning flags associated with transceiver temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.
`SOA bias current high alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for SOA bias current high alarm.
`SOA bias current low alarm threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for SOA bias current low alarm.

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`SOA bias current high warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for SOA bias current high warning.
`SOA bias current low warning threshold` (Availability is dependent on SFP vendor)	Displays the vendor-specified threshold for SOA bias current low warning.
`Laser receiver power high alarm` (Availability is dependent on SFP vendor)	Displays whether the laser receiver power high alarm is `On` or `Off`. This alarm is triggered when the receive optical power is higher than the high alarm threshold.
`Laser receiver power low alarm` (Availability is dependent on SFP vendor)	Displays whether the laser receiver power low alarm is `On` or `Off`. This alarm is triggered when the receive optical power is lower than the high alarm threshold.
`Laser receiver power high warning` (Availability is dependent on SFP vendor)	Displays whether the laser receiver power high warning is `On` or `Off`. This warning is triggered when the receive optical power is higher than the high warning threshold.
`Laser receiver power low warning` (Availability is dependent on SFP vendor)	Displays whether the laser receiver power low warning is `On` or `Off`. This warning is triggered when the receive optical power is lower than the high warning threshold.
`Laser receiver power` (Availability is dependent on SFP vendor)	Displays the laser receiver power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).

Table 622: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Tx loss of signal functionality alarm` (Availability is dependent on SFP vendor)	Displays whether the Tx loss of signal functionality alarm is `On` or `Off`. This alarm is generated if the high-speed electrical signal amplitude of Tx input is below a threshold.
`APD supply fault alarm` (Availability is dependent on SFP vendor)	Displays whether the APD supply fault alarm is `On` or `Off`. Some transceivers have Avalanche Photodiode (APD) receiver to improve receiver sensitivity. APD requires a bias voltage. If there is a disruption in the bias voltage supply, the APD supply fault alarm is generated.
`TEC fault alarm` (Availability is dependent on SFP vendor)	Displays whether the TEC fault alarm is `On` or `Off`. Some transceivers have an internal temperature controller—Thermoelectric Cooler (TEC). If the TEC is faulty, the TEC fault alarm is generated.
`Wavelength unlocked alarm` (Availability is dependent on SFP vendor)	Displays whether the Wavelength unlocked alarm is `On` or `Off`. The transmitter wavelengths of some transceivers must be maintained at greater accuracy over operational lifetimes. Wavelength locker is used to stabilize the wavelength of the laser diode. If the wavelength is out of the specified range, the wavelength unlocked alarm is generated.

Sample Output

show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver)

```
user@switch> show interfaces diagnostics optics ge-0/1/0
Physical interface: ge-0/1/0
    Laser bias current                        :  5.444 mA
    Laser output power                        :  0.3130 mW / -5.04 dBm
    Module temperature                        :  36 degrees C / 97 degrees F
    Module voltage                            :  3.2120 V
    Receiver signal average optical power     :  0.3840 mW / -4.16 dBm
    Laser bias current high alarm             :  Off
```

```
Laser bias current low alarm              :  Off
Laser bias current high warning           :  Off
Laser bias current low warning            :  Off
Laser output power high alarm             :  Off
Laser output power low alarm              :  Off
Laser output power high warning           :  Off
Laser output power low warning            :  Off
Module temperature high alarm             :  Off
Module temperature low alarm              :  Off
Module temperature high warning           :  Off
Module temperature low warning            :  Off
Module voltage high alarm                 :  Off
Module voltage low alarm                  :  Off
Module voltage high warning               :  Off
Module voltage low warning                :  Off
Laser rx power high alarm                 :  Off
Laser rx power low alarm                  :  Off
Laser rx power high warning               :  Off
Laser rx power low warning                :  Off
Laser bias current high alarm threshold   :  15.000 mA
Laser bias current low alarm threshold    :  1.000 mA
Laser bias current high warning threshold :  12.000 mA
Laser bias current low warning threshold  :  2.000 mA
Laser output power high alarm threshold   :  0.6300 mW / -2.01 dBm
Laser output power low alarm threshold    :  0.0660 mW / -11.80 dBm
Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
Laser output power low warning threshold  :  0.0780 mW / -11.08 dBm
Module temperature high alarm threshold   :  109 degrees C / 228 degrees F
Module temperature low alarm threshold    :  -29 degrees C / -20 degrees F
Module temperature high warning threshold :  103 degrees C / 217 degrees F
Module temperature low warning threshold  :  -13 degrees C / 9 degrees F
Module voltage high alarm threshold       :  3.900 V
Module voltage low alarm threshold        :  2.700 V
Module voltage high warning threshold     :  3.700 V
Module voltage low warning threshold      :  2.900 V
Laser rx power high alarm threshold       :  1.2589 mW / 1.00 dBm
Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
Laser rx power high warning threshold     :  0.7939 mW / -1.00 dBm
Laser rx power low warning threshold      :  0.0157 mW / -18.04 dBm
```

Sample Output

show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver)

```
user@switch> show interfaces diagnostics optics xe-0/1/0
Physical interface: xe-0/1/0
    Laser bias current                        :  4.968 mA
    Laser output power                        :  0.4940 mW / -3.06 dBm
    Module temperature                        :  27 degrees C / 81 degrees F
    Module voltage                            :  3.2310 V
    Receiver signal average optical power     :  0.0000
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
```

```
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
```

Sample Output

show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver)

```
user@switch> show interfaces diagnostics optics xe-0/1/0
Physical interface: xe-0/1/0
    Laser bias current                        :  8.029 mA
    Laser output power                        :  0.6430 mW / -1.92 dBm
    Module temperature                        :  4 degrees C / 39 degrees F
    Laser rx power                            :  0.0012 mW / -29.21 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Module not ready alarm                    :  On
    Module power down alarm                   :  Off
    Tx data not ready alarm                   :  Off
    Tx not ready alarm                        :  Off
    Tx laser fault alarm                      :  Off
```

```
    Tx CDR loss of lock alarm                 :  Off
    Rx not ready alarm                        :  On
    Rx loss of signal alarm                   :  On
    Rx CDR loss of lock alarm                 :  On
    Laser bias current high alarm threshold   :  13.000 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  12.000 mA
    Laser bias current low warning threshold  :  3.000 mA
    Laser output power high alarm threshold   :  0.8310 mW / -0.80 dBm
    Laser output power low alarm threshold    :  0.1650 mW / -7.83 dBm
    Laser output power high warning threshold :  0.7410 mW / -1.30 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  90 degrees C / 194 degrees F
    Module temperature low alarm threshold    :  0 degrees C / 32 degrees F
    Module temperature high warning threshold :  85 degrees C / 185 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Laser rx power high alarm threshold       :  0.8912 mW / -0.50 dBm
    Laser rx power low alarm threshold        :  0.0912 mW / -10.40 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
```

Sample Output

show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver)

```
user@switch> show interfaces diagnostics optics et-3/0/0
Physical interface: et-3/0/0
    Module temperature                        :  33 degrees C / 92 degrees F
    Module voltage                            :  3.3060 V
  Lane 0
    Laser bias current                        :  7.182 mA
    Laser receiver power                      :  0.743 mW / -1.29 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
```

```
    Rx loss of signal alarm                  :  Off
  Lane 1
    Laser bias current                       :  7.326 mA
    Laser receiver power                     :  0.752 mW / -1.24 dBm
    Laser bias current high alarm            :  Off
    Laser bias current low alarm             :  Off
    Laser bias current high warning          :  Off
    Laser bias current low warning           :  Off
    Laser receiver power high alarm          :  Off
    Laser receiver power low alarm           :  Off
    Laser receiver power high warning        :  Off
    Laser receiver power low warning         :  Off
    Tx loss of signal functionality alarm    :  Off
    Rx loss of signal alarm                  :  Off
  Lane 2
    Laser bias current                       :  7.447 mA
    Laser receiver power                     :  0.790 mW / -1.03 dBm
    Laser bias current high alarm            :  Off
    Laser bias current low alarm             :  Off
    Laser bias current high warning          :  Off
    Laser bias current low warning           :  Off
    Laser receiver power high alarm          :  Off
    Laser receiver power low alarm           :  Off
    Laser receiver power high warning        :  Off
    Laser receiver power low warning         :  Off
    Tx loss of signal functionality alarm    :  Off
    Rx loss of signal alarm                  :  Off
  Lane 3
    Laser bias current                       :  7.734 mA
    Laser receiver power                     :  0.768 mW / -1.15 dBm
    Laser bias current high alarm            :  Off
    Laser bias current low alarm             :  Off
    Laser bias current high warning          :  Off
    Laser bias current low warning           :  Off
    Laser receiver power high alarm          :  Off
    Laser receiver power low alarm           :  Off
    Laser receiver power high warning        :  Off
    Laser receiver power low warning         :  Off
    Tx loss of signal functionality alarm    :  Off
    Rx loss of signal alarm                  :  Off
```

Sample Output

show interfaces diagnostics optics et-4/1/0 (CFP Transceiver)

```
user@switch> show interfaces diagnostics optics et-4/1/0
Physical interface: et-4/1/0
    Module temperature                        :  38 degrees C / 101 degrees F
    Module voltage                            :  3.2500 V
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Module not ready alarm                    :  Off
    Module low power alarm                    :  Off
    Module initialization incomplete alarm    :  Off
    Module fault alarm                        :  Off
    PLD Flash initialization fault alarm      :  Off
    Power supply fault alarm                  :  Off
    Checksum fault alarm                      :  Off
    Tx laser disabled alarm                   :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.5000 V
    Module voltage low alarm threshold        :  3.0990 V
    Module voltage high warning threshold     :  3.4000 V
    Module voltage low warning threshold      :  3.2000 V
    Laser bias current high alarm threshold   :  250.000 mA
    Laser bias current low alarm threshold    :  37.500 mA
    Laser bias current high warning threshold :  225.000 mA
    Laser bias current low warning threshold  :  50.000 mA
    Laser output power high alarm threshold   :  3.9800 mW / 6.00 dBm
    Laser output power low alarm threshold    :  0.4670 mW / -3.31 dBm
```

```
    Laser output power high warning threshold :  3.5480 mW / 5.50 dBm
    Laser output power low warning threshold  :  0.5240 mW / -2.81 dBm
    Laser rx power high alarm threshold       :  3.5481 mW / 5.50 dBm
    Laser rx power low alarm threshold        :  0.0616 mW / -12.10 dBm
    Laser rx power high warning threshold     :  3.1622 mW / 5.00 dBm
    Laser rx power low warning threshold      :  0.0691 mW / -11.61 dBm
    Laser temperature high alarm threshold    :  67 degrees C / 153 degrees F
    Laser temperature low alarm threshold     :  35 degrees C / 95 degrees F
    Laser temperature high warning threshold  :  62 degrees C / 144 degrees F
    Laser temperature low warning threshold   :  40 degrees C / 104 degrees F
    SOA bias current high alarm threshold     :  0.000 mA
    SOA bias current low alarm threshold      :  0.000 mA
    SOA bias current high warning threshold   :  0.000 mA
    SOA bias current low warning threshold    :  0.000 mA
  Lane 0
    Laser bias current                        :  131.684 mA
    Laser output power                        :  1.002 mW / 0.01 dBm
    Laser temperature                         :  54 degrees C / 128 degrees F
    Laser receiver power                      :  0.497 mW / -3.03 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    APD supply fault alarm                    :  Off
    TEC fault alarm                           :  Off
    Wavelength unlocked alarm                 :  Off
  Lane 1
```

```
    Laser bias current                        :  122.345 mA
    Laser output power                        :  1.002 mW / 0.01 dBm
    Laser temperature                         :  51 degrees C / 124 degrees F
    Laser receiver power                      :  0.611 mW / -2.14 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    APD supply fault alarm                    :  Off
    TEC fault alarm                           :  Off
    Wavelength unlocked alarm                 :  Off
  Lane 2
    Laser bias current                        :  112.819 mA
    Laser output power                        :  1.000 mW / 0.00 dBm
    Laser temperature                         :  50 degrees C / 122 degrees F
    Laser receiver power                      :  0.540 mW / -2.67 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
```

```
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    APD supply fault alarm                    :  Off
    TEC fault alarm                           :  Off
    Wavelength unlocked alarm                 :  Off
  Lane 3
    Laser bias current                        :  100.735 mA
    Laser output power                        :  1.002 mW / 0.01 dBm
    Laser temperature                         :  50 degrees C / 122 degrees F
    Laser receiver power                      :  0.637 mW / -1.96 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    APD supply fault alarm                    :  Off
    TEC fault alarm                           :  Off
    Wavelength unlocked alarm                 :  Off
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- Monitoring Interface Status and Traffic
- *Monitoring Interface Status and Traffic*
- *Install a Transceiver*
- Installing a Transceiver in a QFX Series Device
- *Remove a Transceiver*
- Removing a Transceiver from a QFX Series Device
- Junos OS Ethernet Interfaces Configuration Guide

show interfaces diagnostics optics (Security)

IN THIS SECTION

- Syntax | 20569
- Description | 20570
- Options | 20570
- Required Privilege Level | 20570
- Output Fields | 20570
- Sample Output | 20574
- Release Information | 20575

Syntax

```
show interfaces diagnostics optics interface-name
```

Description

Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP) installed in SRX Series Firewalls. The information provided by this command is known as digital optical monitoring (DOM) information.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.

NOTE: In a chassis cluster, the `show interfaces diagnostics optics` command works only on the node that is primary in redundancy group 0 (RG0).

Options

interface-name Name of the interface associated with the port in which the transceiver is installed: `ge-`*`fpc/pic/port`* .

Required Privilege Level

view

Output Fields

Table 623 on page 20570 lists the output fields for the show interfaces diagnostics optics command. Output fields are listed in the general order in which they appear.

Table 623: show interfaces diagnostics optics Output Fields

Field Name	Field Description
`Physical interface`	Displays the name of the physical interface.
`Laser bias current`	Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents.

Table 623: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser output power`	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
`Module temperature`	Displays the temperature, in Celsius and Fahrenheit.
`Module voltage`	Displays the voltage, in Volts.
`Receiver signal average optical power`	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
`Laser bias current high alarm`	Displays whether the laser bias power setting high alarm is `On` or `Off`.
`Laser bias current low alarm`	Displays whether the laser bias power setting low alarm is `On` or `Off`.
`Laser bias current high warning`	Displays whether the laser bias power setting high warning is `On` or `Off`.
`Laser bias current low warning`	Displays whether the laser bias power setting low warning is `On` or `Off`.
`Laser output power high alarm`	Displays whether the laser output power high alarm is `On` or `Off`.
`Laser output power low alarm`	Displays whether the laser output power low alarm is `On` or `Off`.
`Laser output power high warning`	Displays whether the laser output power high warning is `On` or `Off`.
`Laser output power low warning`	Displays whether the laser output power low warning is `On` or `Off`.
`Module temperature high alarm`	Displays whether the module temperature high alarm is `On` or `Off`.
`Module temperature low alarm`	Displays whether the module temperature low alarm is `On` or `Off`.

Table 623: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Module temperature high warning`	Displays whether the module temperature high warning is `On` or `Off`.
`Module temperature low warning`	Displays whether the module temperature low warning is `On` or `Off`.
`Module voltage high alarm`	Displays whether the module voltage high alarm is `On` or `Off`.
`Module voltage low alarm`	Displays whether the module voltage low alarm is `On` or `Off`.
`Module voltage high warning`	Displays whether the module voltage high warning is `On` or `Off`.
`Module voltage low warning`	Displays whether the module voltage low warning is `On` or `Off`.
`Laser rx power high alarm`	Displays whether the receive laser power high alarm is `On` or `Off`.
`Laser rx power low alarm`	Displays whether the receive laser power low alarm is `On` or `Off`.
`Laser rx power high warning`	Displays whether the receive laser power high warning is `On` or `Off`.
`Laser rx power low warning`	Displays whether the receive laser power low warning is `On` or `Off`.
`Laser bias current high alarm threshold`	Displays the vendor-specified threshold for the laser bias current high alarm.
`Laser bias current low alarm threshold`	Displays the vendor-specified threshold for the laser bias current low alarm.
`Laser bias current high warning threshold`	Displays the vendor-specified threshold for the laser bias current high warning.

Table 623: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser bias current low warning threshold`	Displays the vendor-specified threshold for the laser bias current low warning.
`Laser output power high alarm threshold`	Displays the vendor-specified threshold for the laser output power high alarm.
`Laser output power low alarm threshold`	Displays the vendor-specified threshold for the laser output power low alarm.
`Laser output power high warning threshold`	Displays the vendor-specified threshold for the laser output power high warning.
`Laser output power low warning threshold`	Displays the vendor-specified threshold for the laser output power low warning.
`Module temperature high alarm threshold`	Displays the vendor-specified threshold for the module temperature high alarm.
`Module temperature low alarm threshold`	Displays the vendor-specified threshold for the module temperature low alarm.
`Module temperature high warning threshold`	Displays the vendor-specified threshold for the module temperature high warning.
`Module temperature low warning threshold`	Displays the vendor-specified threshold for the module temperature low warning.
`Module voltage high alarm threshold`	Displays the vendor-specified threshold for the module voltage high alarm.
`Module voltage low alarm threshold`	Displays the vendor-specified threshold for the module voltage low alarm.

Table 623: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Module voltage high warning threshold	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser rx power low warning.

Sample Output

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                        :  7.408 mA
    Laser output power                        :  0.3500 mW / -4.56 dBm
    Module temperature                        :  23 degrees C / 73 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
```

```
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Understanding Interfaces

show interfaces diagnostics optics-applications

IN THIS SECTION

- Syntax | 20576
- Hierarchy Level | 20576
- Description | 20576
- Required Privilege Level | 20576
- Output Fields | 20577
- Sample Output | 20578
- Release Information | 20580

Syntax

```
show interfaces diagnostics optics-applications
```

Hierarchy Level

```
[edit interfaces diagnostics]
```

Description

Displays application selection, host interface code, media interface code, host lanes, media lanes, host assign, and media assign. The output shows the switchable modes.

Required Privilege Level

view

Output Fields

Table 624 on page 20577 lists the output fields for the `show interfaces diagnostics optics-applications` command.

Table 624: show interfaces diagnostics optics-applications Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Interface index	Index number of the physical interface, which reflects its initialization sequence.
Wavelength	Wavelength of the optics in nm.
Media Code	Describes the line side configuration.
Host Code	Describes the system side configuration.
Device Flags	Provides information about the physical interface.
Interface Flags	Provides information about the physical interface.
CoS queues	Number of CoS queues configured.
Current address	Configured MAC address.
Last Flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour :minute :second :timezone (hour :minute :second ago).
Input rate	Input rate in bits per second (bps) and packets per second (pps). Thevalue in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.

Table 624: show interfaces diagnostics optics-applications Output Fields *(Continued)*

Field Name	Field Description
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link. None—There are no active defects or alarms. Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.
PCS statistics	Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.
Last Flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour :minute :second :timezone (hour :minute :second ago).

Sample Output

show interfaces diagnostics optics-applications

To display the 3 different modes for 1x400G speed, use the following command:

```
user@router> show interfaces diagnostics optics-applications et-0/0/10
```

```
root@ardbeg-01# ...gnostics optics-applications et-0/0/10
Physical interface: et-0/0/10
  Interface Name               : et-0/0/10
  Current Speed                : 1x400G
Ap Sel    Host Intf Code                                    Media Intf
Code                                     Host Lanes       Media Lanes    Host Assign     Media
Assign
1         400GAUI-8 C2M (Annex 120E)(17)                    400ZR, DWDM,
amplified(62)                          8                1              1               1
```

```
2         400GAUI-8 C2M (Annex 120E)(17)                    400ZR, Single
Wavelength,Unamplified(63)            8                1              1               1
3         100GAUI-2 C2M (Annex 135G)(13)                    400ZR, DWDM,
amplified(62)                          2                1              85              1
```

To display the media code and host code, use the following command:

```
root@router> show interfaces et-0/0/10
Physical interface: et-0/0/10, Enabled, Physical link is Up
  Interface index: 1045, SNMP ifIndex: 561
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control:
Enabled, Auto-negotiation: Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.40 THz
  Optic-loopback : Enabled , Optic-loopbacktype : media-side-output
  Media Code     : ZR-400-OFEC-16QAM
  Host Code      : 400GAUI-8 C2M (Annex 120E)
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 88:d9:8f:0d:b1:54, Hardware address: 88:d9:8f:0d:b1:54
  Last flapped   : 2023-02-09 01:40:32 PST (04:06:07 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Ethernet FEC Mode  :                 FEC119
  Ethernet FEC statistics              Errors
    FEC Corrected Errors               376995
    FEC Uncorrected Errors               2408
    FEC Corrected Errors Rate               9
    FEC Uncorrected Errors Rate             0
  Optic FEC Mode  :                      OFEC
  Optic FEC statistics:
    Corrected Errors                             678954129269
    Uncorrected Words                                     443
    Corrected Error rate                             23813914
    Uncorrected Error rate                                  0
```

```
    Corrected Error Ratio (      26857 seconds average)  5.73e-05
  PRBS Mode : Disabled
  Interface transmit statistics: Disabled
  Link Degrade :
    Link Monitoring                    :  Disable
  Logical interface et-0/0/10.16386 (Index 1038) (SNMP ifIndex 578)
    Flags: Up SNMP-Traps Encapsulation: ENET2 DF
    Input packets : 0
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

Release Information

Command introduced in Junos OS Evolved Release 23.1R1.

show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)

IN THIS SECTION

- Syntax | 20581
- Description | 20581
- Options | 20581
- Additional Information | 20581
- Required Privilege Level | 20582
- Output Fields | 20582
- Sample Output | 20611
- Release Information | 20644

Syntax

```
show interfaces diagnostics optics interface-name
```

Description

Display diagnostics data, warnings, and alarms for Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, or Virtual Chassis port interfaces.

Options

interface-name	Interface name. For example: ge-*fpc*/*pic*/*port* et-*fpc*/*pic*/*port* et-*fpc*/*pic*/*port*:*channel* xe-*fpc*/*pic*/*port* vcp-*fpc*/*pic*/*port*

Additional Information

The transceivers are polled in 1-second intervals for diagnostics data, warnings, and alarms. The alarms do not cause the links to go down or the LEDs to change color, nor generate SNMP traps. Changes in alarm and warning status generate system log messages.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transceiver vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a device is not working.

> **NOTE**: Some transceivers do not support all optical diagnostics features described in the output fields.
>
> If optics measures transmit or receive power as zero, then, the measured power is displayed as `0.000 mW / - Inf dBm`

You can configure the P2-10G-40G-QSFPP PIC to operate either in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode. When the PIC is in 40-Gigabit Ethernet mode, you must execute the `show interfaces diagnostics optics et-`*fpc*/*pic*/*port* command. The output of this command displays the diagnostic

optics information about the corresponding 40-Gigabit Ethernet port of the PIC. However, when the PIC is in 10-Gigabit Ethernet mode, you must execute the `show interfaces diagnostics optics et-fpc/pic/port:channel` command. The output of this command displays the diagnostic optics information about the corresponding 10-Gigabit Ethernet port of the PIC. For information about the P2-10G-40G-QSFPP PIC, see *P2-10G-40G-QSFPP PIC Overview*.

Required Privilege Level

view

Output Fields

Table 625 on page 20582 lists the output fields for the `show interfaces diagnostics optics` command for DWDM and DWDM OTN PICs. Output fields are listed in the approximate order in which they appear.

Table 625: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs

Field Name	Field Description
`Physical interface`	Name of the physical interface.
`Laser bias current`	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
`Laser output power`	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm). This is a software equivalent to the `LsPOWMON` pin in hardware.
`Receiver signal average optical power`	Average received optical power, in mW and dBm. This indicator is a software equivalent to the `RxPOWMON` pin in hardware. Average optical power is vendor-specific.
`Laser end-of-life alarm`	Laser end-of-life alarm: `On` or `Off`.
`Laser wavelength alarm`	Laser wavelength alarm: `On` or `Off`.

Table 625: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs *(Continued)*

Field Name	Field Description
`Laser bias current alarm`	Laser bias current alarm: `On` or `Off`.
`Laser temperature alarm`	Laser temperature alarm: `On` or `Off`.
`Laser power alarm`	Laser power alarm: `On` or `Off`.
`Modulator temperature alarm`	Modulator temperature alarm: `On` or `Off`. Transceivers from some vendors do not support this field.
`Modulator bias alarm`	Modulator bias alarm: `On` or `Off`.
`Tx multiplexer FIFO error alarm`	Transmit multiplexer first in, first out (FIFO) error alarm: `On` or `Off`.
`Tx loss of PLL lock alarm`	Transmit loss of phase-locked loop (PLL) lock alarm: `On` or `Off`.
`Rx loss of average optical power alarm`	Receive loss of average optical power alarm: `On` or `Off`.
`Rx loss of AC power alarm`	Receive loss of AC power alarm: `On` or `Off`. Transceivers from some vendors do not support this field.
`Rx loss of PLL lock alarm`	Receive loss of phase-locked loop (PLL) lock alarm: `On` or `Off`.

Table 626 on page 20584 lists the output fields for the `show interfaces diagnostics optics` command when the router is operating with bidirectional SFP optics. Output fields are listed in the approximate order in which they appear.

Table 626: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics

Field Name	Field Description
`Physical interface`	Name of the physical interface.
`Laser bias current`	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
`Laser output power`	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm).
`Module temperature`	Temperature of the optics module, in Celsius and Fahrenheit.
`Module voltage`	Internally measured module voltage.
`Receiver signal average optical power`	Average received optical power, in mW and dBm.
`Wavelength Channel number`	Wavelength channel number set in the optics module.
`Wavelength setpoint`	Wavelength set in the optics module.
`Tx Dither`	Transmit dither status. Displays whether transmit dither is enabled or disabled.
`Frequency Error`	Frequency error reported from optics module.
`Wavelength Error`	Wavelength error reported from optics module.
`Laser bias current high alarm`	Laser bias power setting high alarm. Displays `on` or `off`.

Table 626: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(Continued)*

Field Name	Field Description
`Laser bias current low alarm`	Laser bias power setting low alarm. Displays `on` or `off`.
`Laser bias current high warning`	Laser bias power setting high warning. Displays `on` or `off`.
`Laser bias current low warning`	Laser bias power setting low warning. Displays `on` or `off`.
`Laser output power high alarm`	Laser output power high alarm. Displays `on` or `off`.
`Laser output power low alarm`	Laser output power low alarm. Displays `on` or `off`.
`Laser output power high warning`	Laser output power high warning. Displays `on` or `off`.
`Laser output power low warning`	Laser output power low warning. Displays `on` or `off`.
`Module temperature high alarm`	Module temperature high alarm. Displays `on` or `off`.
`Module temperature low alarm`	Module temperature low alarm. Displays `on` or `off`.
`Module temperature high warning`	Module temperature high warning. Displays `on` or `off`.

Table 626: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(Continued)*

Field Name	Field Description
`Module temperature low warning`	Module temperature low warning. Displays `on` or `off`.
`Module voltage high alarm`	Module voltage high alarm. Displays `on` or `off`.
`Module voltage low alarm`	Module voltage low alarm. Displays `on` or `off`.
`Module voltage high warning`	Module voltage high warning. Displays `on` or `off`.
`Module voltage low warning`	Module voltage high warning. Displays `on` or `off`.
`Laser rx power high alarm`	Receive laser power high alarm. Displays `on` or `off`.
`Laser rx power low alarm`	Receive laser power low alarm. Displays `on` or `off`.
`Laser rx power high warning`	Receive laser power high warning. Displays `on` or `off`.
`Laser rx power low warning`	Receive laser power low warning. Displays `on` or `off`.
`TEC fault alarm`	TEC fault alarm. Displays `on` or `off`.
`Wavelength unlocked alarm`	Wavelength unlocked alarm. Displays `on` or `off`.

Table 626: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(Continued)*

Field Name	Field Description
`TxTune`	Optical transmit side status. Displays whether optical transmit side is not ready due to tuning.
`Laser bias current high alarm threshold`	Vendor-specified threshold for the laser bias current high alarm: `70.000 mA`.
`Laser bias current low alarm threshold`	Vendor-specified threshold for the laser bias current low alarm: `0.0002 mA`.
`Laser bias current high warning threshold`	Vendor-specified threshold for the laser bias current high warning: `65.000 mA`.
`Laser bias current low warning threshold`	Vendor-specified threshold for the laser bias current low warning: `0.0002 mA`.
`Laser output power high alarm threshold`	Vendor-specified threshold for the laser output power high alarm: `1.0000 mW` or `0.00 dBm`.
`Laser output power low alarm threshold`	Vendor-specified threshold for the laser output power low alarm: `0.0560 mW` or `-12.52 dBm`.
`Laser output power high warning threshold`	Vendor-specified threshold for the laser output power high warning: `0.6300 mW` or `-2.01 dBm`.
`Laser output power low warning threshold`	Vendor-specified threshold for the laser output power low warning: `0.0890 mW` or `-10.51 dBm`.

Table 626: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(Continued)*

Field Name	Field Description
`Module temperature high alarm threshold`	Vendor-specified threshold for the module temperature high alarm: `100°` C or `212°` F.
`Module temperature low alarm threshold`	Vendor-specified threshold for the module temperature low alarm: `-50°` C or `-58°` F.
`Module temperature high warning threshold`	Vendor-specified threshold for the module temperature high warning: `95 °` C or `203 °` F.
`Module temperature low warning threshold`	Vendor-specified threshold for the module temperature low warning: `-48°` C or `-54°` F.
`Module voltage high alarm threshold`	Module voltage high alarm threshold: `3.700 v`.
`Module voltage low alarm threshold`	Module voltage low alarm threshold: `2.900 v`.
`Module voltage high warning threshold`	Module voltage high warning threshold: `3.7600 v`.
`Module voltage low warning threshold`	Module voltage low warning threshold: `3.000 v`.
`Laser rx power high alarm threshold`	Vendor-specified threshold for the laser Rx power high alarm: `1.9953 mW` or `3.00 dBm`.
`Laser rx power low alarm threshold`	Vendor-specified threshold for the laser Rx power low alarm: `0.0001 mW` or `-40.00 dBm`.

Table 626: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(Continued)*

Field Name	Field Description
`Laser rx power high warning threshold`	Vendor-specified threshold for the laser Rx power high warning: `1.0000 mW` or `0.00 dBm`.
`Laser rx power low warning threshold`	Vendor-specified threshold for the laser Rx power low warning: `0.0010 mW` or `-30.00 dBm`.

Table 627 on page 20589 lists the output fields for the `show interfaces diagnostics optics` command for SFP transceivers. Output fields are listed in the approximate order in which they appear.

Table 627: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers

Field Name	Field Description
`Physical interface`	Name of the physical interface.
`Laser bias current`	Measured laser bias current in uA.
`Laser output power`	Measured laser output power in mW.
`Module temperature`	Internally measured module temperature.
`Module voltage`	Internally measured module voltage.
`Laser rx power`	Measured receive optical power in mW.
`Laser bias current high alarm`	Laser bias current high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser bias current low alarm`	Laser bias current low alarm: `On` or `Off`. Alarm ranges are vendor-specific.

Table 627: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers *(Continued)*

Field Name	Field Description
`Laser output power high alarm`	Laser output power high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser output power low alarm`	Laser output power low alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Module temp high alarm`	Module temperature high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Module temp low alarm`	Module temperature low alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser rx power high alarm`	Laser receive power high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser rx power low alarm`	Laser receive power low alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser bias current high warning`	Laser bias current high warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser bias current low warning`	Laser bias current low warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser output power high warning`	Laser output power high warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser output power low warning`	Laser output power low warning: `On` or `Off`. Warning ranges are vendor-specific.
`Module temperature high warning`	Module temperature high warning: `On` or `Off`. Warning ranges are vendor-specific.

Table 627: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers *(Continued)*

Field Name	Field Description
`Module temperature low warning`	Module temperature low warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser rx power high warning`	Laser receive power high warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser rx power low warning`	Laser receive power low warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser bias current high alarm threshold`	Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser bias current low alarm threshold`	Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser bias current high warning threshold`	Laser bias current high warning threshold. Warning ranges are vendor-specific.
`Laser bias current low warning threshold`	Laser bias current low warning threshold. Warning ranges are vendor-specific.
`Laser output power high alarm threshold`	Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser output power low alarm threshold`	Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser output power high warning threshold`	Laser output power high warning threshold. Warning ranges are vendor-specific.

Table 627: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers *(Continued)*

Field Name	Field Description
`Laser output power low warning threshold`	Laser output power low warning threshold. Warning ranges are vendor-specific.
`Module temperature high alarm threshold`	Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.
`Module temperature low alarm threshold`	Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.
`Module temperature high warning threshold`	Module temperature high warning threshold. Warning ranges are vendor-specific.
`Module temperature low warning threshold`	Module temperature low warning threshold. Warning ranges are vendor-specific.
`Module voltage high alarm threshold`	Module voltage high alarm threshold. Alarm ranges are vendor-specific.
`Module voltage low alarm threshold`	Module voltage low alarm threshold. Alarm ranges are vendor-specific.
`Module voltage high warning threshold`	Module voltage high warning threshold. Warning ranges are vendor-specific.
`Module voltage low warning threshold`	Module voltage low warning threshold. Warning ranges are vendor-specific.
`Laser rx power high alarm threshold`	Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.

Table 627: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers *(Continued)*

Field Name	Field Description
`Laser rx power low alarm threshold`	Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser rx power high warning threshold`	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.
`Laser rx power high low threshold`	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.

Table 628 on page 20593 lists the output fields for the `show interfaces diagnostics optics` command for 10-Gigabit Ethernet transceivers. Output fields are listed in the approximate order in which they appear.

Table 628: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers

Field Name	Field Description
`Physical interface`	Name of the physical interface.
`Laser bias current`	Measured laser bias current in mA.
`Laser output power`	Measured laser output power in mW.
`Module temperature`	Internally measured module temperature.
`Laser rx power`	Measured receive optical power in mW.
`Laser bias current high alarm`	Laser bias current high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser bias current low alarm`	Laser bias current low alarm: `On` or `Off`. Alarm ranges are vendor-specific.

Table 628: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers *(Continued)*

Field Name	Field Description
`Laser output power high alarm`	Laser output power high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser output power low alarm`	Laser output power low alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Module temp high alarm`	Module temperature high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Module temp low alarm`	Module temperature low alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser rx power high alarm`	Laser receive power high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser rx power low alarm`	Laser receive power low alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser bias current high warning`	Laser bias current high warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser bias current low warning`	Laser bias current low warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser output power high warning`	Laser output power high warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser output power low warning`	Laser output power low warning: `On` or `Off`. Warning ranges are vendor-specific.

Table 628: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers *(Continued)*

Field Name	Field Description
`Module temperature high warning`	Module temperature high warning: `On` or `Off`. Warning ranges are vendor-specific.
`Module temperature low warning`	Module temperature low warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser rx power high warning`	Laser receive power high warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser rx power low warning`	Laser receive power low warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser bias current high alarm threshold`	Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser bias current low alarm threshold`	Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser output power high alarm threshold`	Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser output power low alarm threshold`	Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.
`Module temperature high alarm threshold`	Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.
`Module temperature low alarm threshold`	Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.

Table 628: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers *(Continued)*

Field Name	Field Description
`Laser rx power high alarm threshold`	Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser rx power low alarm threshold`	Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser bias current high warning threshold`	Laser bias current high warning threshold. Warning ranges are vendor-specific.
`Laser bias current low warning threshold`	Laser bias current low warning threshold. Warning ranges are vendor-specific.
`Laser output power high warning threshold`	Laser output power high warning threshold. Warning ranges are vendor-specific.
`Laser output power low warning threshold`	Laser output power low warning threshold. Warning ranges are vendor-specific.
`Module temperature high warning threshold`	Module temperature high warning threshold. Warning ranges are vendor-specific.
`Module temperature low warning threshold`	Module temperature low warning threshold. Warning ranges are vendor-specific.
`Laser rx power high warning threshold`	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.

Table 628: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers *(Continued)*

Field Name	Field Description
`Laser rx power low warning threshold`	Laser receive power low warning threshold. Warning threshold ranges are vendor-specific.

Table 629 on page 20597 lists the output fields for the `show interfaces diagnostics optics` command for 100-Gigabit Ethernet transceivers. Output fields are listed in the approximate order in which they appear.

Table 629: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers

Field Name	Field Description
`Physical interface`	Name of the physical interface.
Grid Channel Number	Set of wavelengths are divided into 40 grids. At a time transceiver will be programmed in one of this set (grid). Mapping of center wavelength to grid number is presented by this parameter.
Corrected Error Ratio	Indicates accumulated Bit Error Ratio.
Uncorrected Words Ratio	Monitors the error rate for either the full 100G link or on a channel-by-channel basis. Indicates the frame errors.
`Laser bias current`	Measured laser bias current in mA.
`Laser output power`	Measured laser output power in mW.
`Module temperature`	Internally measured module temperature.
`Laser rx power`	Measured receive optical power in mW.

Table 629: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers *(Continued)*

Field Name	Field Description
`Laser bias current high alarm`	Laser bias current high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser bias current low alarm`	Laser bias current low alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser output power high alarm`	Laser output power high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser output power low alarm`	Laser output power low alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Module temp high alarm`	Module temperature high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Module temp low alarm`	Module temperature low alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser rx power high alarm`	Laser receive power high alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser rx power low alarm`	Laser receive power low alarm: `On` or `Off`. Alarm ranges are vendor-specific.
`Laser bias current high warning`	Laser bias current high warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser bias current low warning`	Laser bias current low warning: `On` or `Off`. Warning ranges are vendor-specific.

Table 629: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers *(Continued)*

Field Name	Field Description
`Laser output power high warning`	Laser output power high warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser output power low warning`	Laser output power low warning: `On` or `Off`. Warning ranges are vendor-specific.
`Module temperature high warning`	Module temperature high warning: `On` or `Off`. Warning ranges are vendor-specific.
`Module temperature low warning`	Module temperature low warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser rx power high warning`	Laser receive power high warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser rx power low warning`	Laser receive power low warning: `On` or `Off`. Warning ranges are vendor-specific.
`Laser bias current high alarm threshold`	Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser bias current low alarm threshold`	Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser output power high alarm threshold`	Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser output power low alarm threshold`	Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.

Table 629: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers *(Continued)*

Field Name	Field Description
`Module temperature high alarm threshold`	Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.
`Module temperature low alarm threshold`	Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser rx power high alarm threshold`	Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser rx power low alarm threshold`	Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.
`Laser bias current high warning threshold`	Laser bias current high warning threshold. Warning ranges are vendor-specific.
`Laser bias current low warning threshold`	Laser bias current low warning threshold. Warning ranges are vendor-specific.
`Laser output power high warning threshold`	Laser output power high warning threshold. Warning ranges are vendor-specific.
`Laser output power low warning threshold`	Laser output power low warning threshold. Warning ranges are vendor-specific.
`Module temperature high warning threshold`	Module temperature high warning threshold. Warning ranges are vendor-specific.

Table 629: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers *(Continued)*

Field Name	Field Description
`Module temperature low warning threshold`	Module temperature low warning threshold. Warning ranges are vendor-specific.
`Laser rx power high warning threshold`	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.
`Laser rx power low warning threshold`	Laser receive power low warning threshold. Warning threshold ranges are vendor-specific.
Lane carrier frequency offset	Difference (in frequency units) between the target frequency and the actual frequency.
Lane SNR	Signal-to-noise ratio of the electrical data present on the channel.
TEC Current	Monitors the amount of current flowing to the TEC of a cooled laser. It is a 16-bit signed 2s complement value with a LSB unit of 0.1 mA. Thus the total range is from -3.2768 A to +3.2767 A.
Residual ISI	Measures the amount of correction being done by the module to account for residual inter-symbol interference (ISI). The usual cause for this is optical dispersion so this measurement is a proxy for residual (uncorrected) optical dispersion that is being corrected by the module. The parameter is unitless and the threshold alarm and warning values will give an indication of the severity of the uncorrected dispersion.
PAM Histogram	Provides the rate of measured signal on the line that has an analog level near the cutoff for a PAM bit transition (for example, 0 <-> 1, 1<->2, 2<->3).

Table 630 on page 20602 lists the output fields for the `show interfaces diagnostics optics` command for XFP transceivers. Output fields are listed in the approximate order in which they appear.

Table 630: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers

Field Name	Field Description
`Physical interface`	Name of the physical interface.
`Laser bias current`	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
`Laser output power`	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm). This is a software equivalent to the `LsPOWMON` pin in hardware.
`Module temperature`	Temperature of the XFP optics module, in Celsius and Fahrenheit.
`Laser rx power`	Laser received optical power, in mW and dBm.
`Laser bias current high alarm`	Laser bias power setting high alarm. Displays `on` or `off`.
`Laser bias current low alarm`	Laser bias power setting low alarm. Displays `on` or `off`.
`Laser bias current high warning`	Laser bias power setting high warning. Displays `on` or `off`.
`Laser bias current low warning`	Laser bias power setting low warning. Displays `on` or `off`.
`Laser output power high alarm`	Laser output power high alarm. Displays `on` or `off`.
`Laser output power low alarm`	Laser output power low alarm. Displays `on` or `off`.
`Laser output power high warning`	Laser output power high warning. Displays `on` or `off`.

Table 630: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers *(Continued)*

Field Name	Field Description
`Laser output power low warning`	Laser output power low warning. Displays `on` or `off`.
`Module temperature high alarm`	Module temperature high alarm. Displays `on` or `off`.
`Module temperature low alarm`	Module temperature low alarm. Displays `on` or `off`.
`Module temperature high warning`	Module temperature high warning. Displays `on` or `off`.
`Module temperature low warning`	Module temperature low warning. Displays `on` or `off`.
`Laser rx power high alarm`	Receive laser power high alarm. Displays `on` or `off`.
`Laser rx power low alarm`	Receive laser power low alarm. Displays `on` or `off`.
`Laser rx power high warning`	Receive laser power high warning. Displays `on` or `off`.
`Laser rx power low warning`	Receive laser power low warning. Displays `on` or `off`.
`Module not ready alarm`	Module not ready alarm. When `on`, indicates the module has an operational fault. Displays `on` or `off`.

Table 630: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers *(Continued)*

Field Name	Field Description
`Module power down alarm`	Module power down alarm. When `on`, module is in a limited power mode, low for normal operation. Displays on or off.
`Tx data not ready alarm`	Any condition leading to invalid data on the transmit path. Displays `on` or `off`.
`Tx not ready alarm`	Any condition leading to invalid data on the transmit path. Displays `on` or `off`.
`Tx laser fault alarm`	Laser fault condition. Displays `on` or `off`.
`Tx CDR loss of lock alarm`	Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays `on` or `off`.
`Rx not ready alarm`	Any condition leading to invalid data on the receive path. Displays `on` or `off`.
`Rx loss of signal alarm`	Receive Loss of Signal alarm. When `on`, indicates insufficient optical input power to the module. Displays `on` or `off`.
`Rx CDR loss of lock alarm`	Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays `on` or `off`.
`Laser bias current high alarm threshold`	Vendor-specified threshold for the laser bias current high alarm: `130.000 mA`.
`Laser bias current low alarm threshold`	Vendor-specified threshold for the laser bias current low alarm: `10.000 mA`.
`Laser bias current high warning threshold`	Vendor-specified threshold for the laser bias current high warning: `120.000 mA`.

Table 630: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers *(Continued)*

Field Name	Field Description
`Laser bias current low warning threshold`	Vendor-specified threshold for the laser bias current low warning: `12.000 mA`.
`Laser output power high alarm threshold`	Vendor-specified threshold for the laser output power high alarm: `0.8910 mW` or `-0.50 dBm`.
`Laser output power low alarm threshold`	Vendor-specified threshold for the laser output power low alarm: `0.2230 mW` or `-6.52 dBm`.
`Laser output power high warning threshold`	Vendor-specified threshold for the laser output power high warning: `0.7940 mW` or `-100 dBm`.
`Laser output power low warning threshold`	Vendor-specified threshold for the laser output power low warning: `0.2510 mW` or `-600 dBm`.
`Module temperature high alarm threshold`	Vendor-specified threshold for the module temperature high alarm: `90°` C or `194°` F.
`Module temperature low alarm threshold`	Vendor-specified threshold for the module temperature low alarm: `-5°` C or `23°` F.
`Module temperature high warning threshold`	Vendor-specified threshold for the module temperature high warning: `85 °` C or `185 °` F.
`Module temperature low warning threshold`	Vendor-specified threshold for the module temperature low warning: `0°` C or `32°` F.

Table 630: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers *(Continued)*

Field Name	Field Description
`Laser rx power high alarm threshold`	Vendor-specified threshold for the laser Rx power high alarm: `1.2589 mW` or `1.00 dBm`.
`Laser rx power low alarm threshold`	Vendor-specified threshold for the laser Rx power low alarm: `0.0323 mW` or `-14.91 dBm`.
`Laser rx power high warning threshold`	Vendor-specified threshold for the laser Rx power high warning: `1.1220 mW` or `0.50 dBm`.
`Laser rx power low warning threshold`	Vendor-specified threshold for the laser Rx power low warning: `0.0363 mW` or `-14.40 dBm`.

Table 631 on page 20606 lists the output fields for the `show interfaces diagnostics optics` command for VCP. Output fields are listed in the approximate order in which they appear.

Table 631: show interfaces diagnostics optics Output for Virtual Chassis Ports

Field Name	Field Description
`Physical interface`	Name of the physical interface.
`Laser bias current`	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
`Laser output power`	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm).
`Module temperature`	Temperature of the optics module, in Celsius and Fahrenheit.
`Module voltage`	Internally measured module voltage.

Table 631: show interfaces diagnostics optics Output for Virtual Chassis Ports *(Continued)*

Field Name	Field Description
`Receiver signal average optical power`	Average received optical power, in mW and dBm.
`Laser bias current high alarm`	Laser bias power setting high alarm. Displays `on` or `off`.
`Laser bias current low alarm`	Laser bias power setting low alarm. Displays `on` or `off`.
`Laser bias current high warning`	Laser bias power setting high warning. Displays `on` or `off`.
`Laser bias current low warning`	Laser bias power setting low warning. Displays `on` or `off`.
`Laser output power high alarm`	Laser output power high alarm. Displays `on` or `off`.
`Laser output power low alarm`	Laser output power low alarm. Displays `on` or `off`.
`Laser output power high warning`	Laser output power high warning. Displays `on` or `off`.
`Laser output power low warning`	Laser output power low warning. Displays `on` or `off`.
`Module temperature high alarm`	Module temperature high alarm. Displays `on` or `off`.

Table 631: show interfaces diagnostics optics Output for Virtual Chassis Ports *(Continued)*

Field Name	Field Description
`Module temperature low alarm`	Module temperature low alarm. Displays `on` or `off`.
`Module temperature high warning`	Module temperature high warning. Displays `on` or `off`.
`Module temperature low warning`	Module temperature low warning. Displays `on` or `off`.
`Module voltage high alarm`	Module voltage high alarm. Displays `on` or `off`.
`Module voltage low alarm`	Module voltage low alarm. Displays `on` or `off`.
`Module voltage high warning`	Module voltage high warning. Displays `on` or `off`.
`Module voltage low warning`	Module voltage high warning. Displays `on` or `off`.
`Laser rx power high alarm`	Receive laser power high alarm. Displays `on` or `off`.
`Laser rx power low alarm`	Receive laser power low alarm. Displays `on` or `off`.
`Laser rx power high warning`	Receive laser power high warning. Displays `on` or `off`.
`Laser rx power low warning`	Receive laser power low warning. Displays `on` or `off`.

Table 631: show interfaces diagnostics optics Output for Virtual Chassis Ports *(Continued)*

Field Name	Field Description
`Laser bias current high alarm threshold`	Vendor-specified threshold for the laser bias current high alarm.
`Laser bias current low alarm threshold`	Vendor-specified threshold for the laser bias current low alarm.
`Laser bias current high warning threshold`	Vendor-specified threshold for the laser bias current high warning.
`Laser bias current low warning threshold`	Vendor-specified threshold for the laser bias current low warning.
`Laser output power high alarm threshold`	Vendor-specified threshold for the laser output power high alarm.
`Laser output power low alarm threshold`	Vendor-specified threshold for the laser output power low alarm.
`Laser output power high warning threshold`	Vendor-specified threshold for the laser output power high warning.
`Laser output power low warning threshold`	Vendor-specified threshold for the laser output power low warning.
`Module temperature high alarm threshold`	Vendor-specified threshold for the module temperature high alarm.
`Module temperature low alarm threshold`	Vendor-specified threshold for the module temperature low alarm.

Table 631: show interfaces diagnostics optics Output for Virtual Chassis Ports *(Continued)*

Field Name	Field Description
`Module temperature high warning threshold`	Vendor-specified threshold for the module temperature high warning.
`Module temperature low warning threshold`	Vendor-specified threshold for the module temperature low warning.
`Module voltage high alarm threshold`	Module voltage high alarm threshold.
`Module voltage low alarm threshold`	Module voltage low alarm threshold.
`Module voltage high warning threshold`	Module voltage high warning threshold.
`Module voltage low warning threshold`	Module voltage low warning threshold.
`Laser rx power high alarm threshold`	Vendor-specified threshold for the laser Rx power high alarm.
`Laser rx power low alarm threshold`	Vendor-specified threshold for the laser Rx power low alarm.
`Laser rx power high warning threshold`	Vendor-specified threshold for the laser Rx power high warning.
`Laser rx power low warning threshold`	Vendor-specified threshold for the laser Rx power low warning.

Sample Output

show interfaces diagnostics optics for 10-Gigabit Ethernet on JNP10K-LC1201 line card (PTX10008)

```
user@host> show interfaces diagnostics optics et-4/0/20
Physical interface: et-4/0/20
    Module temperature                      :  34 degrees C / 93 degrees F
    Module voltage                          :  3.315 V
    Wavelength channel number               :  42
    Wavelength setpoint                     :  1550.10 nm
    Tx dither                               :  Enabled
    Frequency error                         :  0.20 GHz
    Wavelength error                        :  0 nm
    TEC fault alarm                         :  False
    Wavelength unlocked alarm               :  False
    Tx tune alarm                           :  False
    Module temperature high alarm           :  Off
    Module temperature low alarm            :  Off
    Module temperature high warning         :  Off
    Module temperature low warning          :  Off
    Module voltage high alarm               :  Off
Cont.............................................
```

show interfaces diagnostics optics (DWDM and DWDM OTN)

```
user@host> show interfaces diagnostics optics ge-5/0/0
Physical interface: ge-5/0/0
  Laser bias current                    :  79.938 mA
  Laser output power                    :  1.592 mW / 2.02 dBm
  Receiver signal average optical power :  1.3854 mW / 1.42 dBm
  Laser end-of-life alarm               :  Off
  Laser wavelength alarm                :  Off
  Laser bias current alarm              :  Off
  Laser temperature alarm               :  Off
  Laser power alarm                     :  Off
  Modulator temperature alarm           :  Off
  Modulator bias alarm                  :  Off
  Tx multiplexer FIFO error alarm       :  Off
  Tx loss of PLL lock alarm             :  Off
```

```
  Rx loss of average optical power alarm:  Off
  Rx loss of AC power alarm             :  Off
  Rx loss of PLL lock alarm             :  Off
```

show interfaces diagnostics optics (MPC6E with OTN MIC)

```
user@host> show interfaces diagnostics optics xe-3/0/0
Physical interface: xe-3/0/0
    Laser bias current                        :  7.806 mA
    Laser output power                        :  0.5660 mW / -2.47 dBm
    Module temperature                        :  32 degrees C / 89 degrees F
    Module voltage                            :  3.3560 V
    Receiver signal average optical power     :  0.5501 mW / -2.60 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  11.800 mA
    Laser bias current low alarm threshold    :  4.000 mA
    Laser bias current high warning threshold :  10.800 mA
    Laser bias current low warning threshold  :  5.000 mA
    Laser output power high alarm threshold   :  0.8310 mW / -0.80 dBm
    Laser output power low alarm threshold    :  0.2510 mW / -6.00 dBm
    Laser output power high warning threshold :  0.6600 mW / -1.80 dBm
    Laser output power low warning threshold  :  0.3160 mW / -5.00 dBm
```

```
    Module temperature high alarm threshold   :  78 degrees C / 172 degrees F
    Module temperature low alarm threshold    :  -13 degrees C / 9 degrees F
    Module temperature high warning threshold :  73 degrees C / 163 degrees F
    Module temperature low warning threshold  :  -8 degrees C / 18 degrees F
    Module voltage high alarm threshold       :  3.700 V
    Module voltage low alarm threshold        :  2.900 V
    Module voltage high warning threshold     :  3.600 V
    Module voltage low warning threshold      :  3.000 V
    Laser rx power high alarm threshold       :  1.0000 mW / 0.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces diagnostics optics (Bidirectional SFP)

```
user@host> show interfaces diagnostics optics ge-3/0/6
Physical interface: ge-3/0/6
    Laser bias current                        :  13.356 mA
    Laser output power                        :  0.2210 mW / -6.56 dBm
    Module temperature                        :  36 degrees C / 96 degrees F
    Module voltage                            :  3.2180 V
    Receiver signal average optical power     :  0.2429 mW / -6.15 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
```

```
    Laser rx power high warning                :  Off
    Laser rx power low warning                 :  Off
    Laser bias current high alarm threshold    :  70.000 mA
    Laser bias current low alarm threshold     :  0.002 mA
    Laser bias current high warning threshold :  65.000 mA
    Laser bias current low warning threshold   :  0.002 mA
    Laser output power high alarm threshold    :  1.0000 mW / 0.00 dBm
    Laser output power low alarm threshold     :  0.0560 mW / -12.52 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold   :  0.0890 mW / -10.51 dBm
    Module temperature high alarm threshold    :  100 degrees C / 212 degrees F
    Module temperature low alarm threshold     :  -50 degrees C / -58 degrees F
    Module temperature high warning threshold :  95 degrees C / 203 degrees F
    Module temperature low warning threshold   :  -48 degrees C / -54 degrees F
    Module voltage high alarm threshold        :  3.700 V
    Module voltage low alarm threshold         :  2.900 V
    Module voltage high warning threshold      :  3.600 V
    Module voltage low warning threshold       :  3.000 V
    Laser rx power high alarm threshold        :  1.9953 mW / 3.00 dBm
    Laser rx power low alarm threshold         :  0.0001 mW / -40.00 dBm
    Laser rx power high warning threshold      :  1.0000 mW / 0.00 dBm
    Laser rx power low warning threshold       :  0.0010 mW / -30.00 dBm
```

show interfaces diagnostics optics (SFP)

```
user@host> show interfaces diagnostics optics ge-0/3/0
Physical interface: ge-0/3/0
    Laser bias current                         :  23.408 mA
    Laser output power                         :  1.479 mW / 1.70 dBm
    Module temperature                         :  37 degrees C / 99 degrees F
    Laser rx power                             :  0.121 mW / -9.16 dBm
    Laser bias current high alarm              :  Off
    Laser bias current low alarm               :  Off
    Laser output power high alarm              :  Off
    Laser output power low alarm               :  Off
    Module temperature high alarm              :  Off
    Module temperature low alarm               :  Off
    Laser rx power high alarm                  :  Off
    Laser rx power low alarm                   :  Off
    Laser bias current high warning            :  Off
    Laser bias current low warning             :  Off
```

```
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  31.000 mA
    Laser bias current low alarm threshold    :  10.000 mA
    Laser output power high alarm threshold   :  6.000 mW / 7.78 dBm
    Laser output power low alarm threshold    :  0.100 mW / -10.00 dBm
    Module temperature high alarm threshold   :  85 degrees C / 185 degrees F
    Module temperature low alarm threshold    :  0 degrees C / 32 degrees F
    Laser rx power high alarm threshold       :  1.000 mW / 0.00 dBm
    Laser rx power low alarm threshold        :  0.001 mW / -30.00 dBm
    Laser bias current high warning threshold :  28.000 mA
    Laser bias current low warning threshold  :  11.000 mA
    Laser output power high warning threshold :  5.000 mW / 6.99 dBm
    Laser output power low warning threshold  :  0.500 mW / -3.01 dBm
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  10 degrees C / 50 degrees F
    Laser rx power high warning threshold     :  0.501 mW / -3.00 dBm
    Laser rx power low warning threshold      :  0.001 mW / -28.86 dBm
```

show interfaces diagnostics optics (SFP)

```
user@host> show interfaces diagnostics optics ge-1/0/0
Physical interface: ge-1/0/0
        Laser bias current                        :  49.010 mA
        Laser output power                        :  1.263 mW / 1.01 dBm
        Module temperature                        :  17 degrees C / 62 degrees F
        Module voltage                            :  4.21 V
        Laser rx power                            :  0.060 mW / -12.21 dBm
        Laser bias current high alarm             :  Off
        Laser bias current low alarm              :  Off
        Laser output power high alarm             :  Off
        Laser output power low alarm              :  Off
        Module temperature high alarm             :  Off
        Module temperature low alarm              :  Off
        Module voltage high alarm                 :  Off
        Module voltage low alarm                  :  Off
        Laser rx power high alarm                 :  Off
```

```
        Laser rx power low alarm                  :  Off
        Laser bias current high warning           :  Off
        Laser bias current low warning            :  Off
        Laser output power high warning           :  Off
        Laser output power low warning            :  Off
        Module temperature high warning           :  Off
        Module temperature low warning            :  Off
        Module voltage high warning               :  Off
        Module voltage low warning                :  Off
        Laser rx power high warning               :  Off
        Laser rx power low warning                :  Off
        Laser bias current high alarm threshold   :  70.000 mA
        Laser bias current low alarm threshold    :  20.000 mA
        Laser bias current high warning threshold :  65.000 mA
        Laser bias current low warning threshold  :  25.000 mA
        Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
        Laser output power low alarm threshold    :  0.1990 mW / -7.01 dBm
        Laser output power high warning threshold :  1.2580 mW / 1.00 dBm
        Laser output power low warning threshold  :  0.2230 mW / -6.52 dBm
        Module temperature high alarm threshold   :  78 degrees C /172 degrees F
        Module temperature low alarm threshold    :  13 degrees C / 9 degrees F
        Module temperature high warning threshold :  75 degrees C /167 degrees F
        Module temperature low warning threshold  :  10 degrees C / 14 degrees F
        Module voltage high alarm threshold       :  5.71 V
        Module voltage low alarm threshold        :  2.05 V
        Module voltage high warning threshold     :  5.20 V
        Module voltage low warning threshold      :  3.11 V
        Laser rx power high alarm threshold       :  1.7783 mW / 2.50 dBm
        Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
        Laser rx power high warning threshold     :  1.5849 mW / 2.00 dBm
        Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces diagnostics optics (XFP and CFP Optics)

```
user@host> show interfaces diagnostics optics xe-2/1/0
Physical interface: xe-2/1/0
    Laser bias current                        :  52.060 mA
    Laser output power                        :  0.5640 mW / -2.49 dBm
    Module temperature                        :  31 degrees C / 88 degrees F
    Laser rx power                            :  0.0844 mW / -10.74 dBm
    Laser bias current high alarm             :  Off
```

```
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Module not ready alarm                    :  Off
    Module power down alarm                   :  Off
    Tx data not ready alarm                   :  Off
    Tx not ready alarm                        :  Off
    Tx laser fault alarm                      :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx not ready alarm                        :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    Laser bias current high alarm threshold   :  130.000 mA
    Laser bias current low alarm threshold    :  10.000 mA
    Laser bias current high warning threshold :  120.000 mA
    Laser bias current low warning threshold  :  12.000 mA
    Laser output power high alarm threshold   :  0.8910 mW / -0.50 dBm
    Laser output power low alarm threshold    :  0.2230 mW / -6.52 dBm
    Laser output power high warning threshold :  0.7940 mW / -1.00 dBm
    Laser output power low warning threshold  :  0.2510 mW / -6.00 dBm
    Module temperature high alarm threshold   :  90 degrees C / 194 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  85 degrees C / 185 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Laser rx power high alarm threshold       :  1.2589 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0323 mW / -14.91 dBm
    Laser rx power high warning threshold     :  1.1220 mW / 0.50 dBm
    Laser rx power low warning threshold      :  0.0363 mW / -14.40 dBm
```

show interfaces diagnostics optics for 10-Gigabit Ethernet (PTX 24-10GE-SFPP)

```
user@host> show interfaces diagnostics optics et-2/0/23
Physical interface: et-2/0/23
    Laser bias current                        :  8.482 mA
    Laser output power                        :  0.5890 mW / -2.30 dBm
    Module temperature                        :  51 degrees C / 123 degrees F
    Module voltage                            :  3.2970 V
    Receiver signal average optical power     :  0.5574 mW / -2.54 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  11.800 mA
    Laser bias current low alarm threshold    :  4.000 mA
    Laser bias current high warning threshold :  10.800 mA
    Laser bias current low warning threshold  :  5.000 mA
    Laser output power high alarm threshold   :  0.8310 mW / -0.80 dBm
    Laser output power low alarm threshold    :  0.2510 mW / -6.00 dBm
    Laser output power high warning threshold :  0.6600 mW / -1.80 dBm
    Laser output power low warning threshold  :  0.3160 mW / -5.00 dBm
    Module temperature high alarm threshold   :  93 degrees C / 199 degrees F
    Module temperature low alarm threshold    :  -13 degrees C / 9 degrees F
    Module temperature high warning threshold :  88 degrees C / 190 degrees F
    Module temperature low warning threshold  :  -8 degrees C / 18 degrees F
    Module voltage high alarm threshold       :  3.700 V
```

```
    Module voltage low alarm threshold        :  2.900 V
    Module voltage high warning threshold     :  3.600 V
    Module voltage low warning threshold      :  3.000 V
    Laser rx power high alarm threshold       :  1.0000 mW / 0.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces diagnostics optics for 40-Gigabit Ethernet

```
user@host> show interfaces diagnostics optics et-7/1/0
Physical interface: et-7/1/0
    Module temperature                        :  34 degrees C / 94 degrees F
    Module voltage                            :  3.4720 V
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Module not ready alarm                    :  Off
    Module low power alarm                    :  Off
    Module initialization incomplete alarm    :  Off
    Module fault alarm                        :  Off
    PLD Flash initialization fault alarm      :  Off
    Power supply fault alarm                  :  Off
    Checksum fault alarm                      :  Off
    Tx laser disabled alarm                   :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    Module temperature high alarm threshold   :  80 degrees C / 176 degrees F
    Module temperature low alarm threshold    :  -10 degrees C / 14 degrees F
    Module temperature high warning threshold :  75 degrees C / 167 degrees F
    Module temperature low warning threshold  :  -5 degrees C / 23 degrees F
    Module voltage high alarm threshold       :  3.5990 V
    Module voltage low alarm threshold        :  3.0000 V
```

```
    Module voltage high warning threshold     :  3.5000 V
    Module voltage low warning threshold      :  3.0990 V
    Laser bias current high alarm threshold   :  100.000 mA
    Laser bias current low alarm threshold    :  10.000 mA
    Laser bias current high warning threshold :  80.000 mA
    Laser bias current low warning threshold  :  15.000 mA
    Laser output power high alarm threshold   :  2.8180 mW / 4.50 dBm
    Laser output power low alarm threshold    :  0.2390 mW / -6.22 dBm
    Laser output power high warning threshold :  2.2380 mW / 3.50 dBm
    Laser output power low warning threshold  :  0.3010 mW / -5.21 dBm
    Laser rx power high alarm threshold       :  2.5119 mW / 4.00 dBm
    Laser rx power low alarm threshold        :  0.0316 mW / -15.00 dBm
    Laser rx power high warning threshold     :  1.9953 mW / 3.00 dBm
    Laser rx power low warning threshold      :  0.0631 mW / -12.00 dBm
    Laser temperature high alarm threshold    :  80 degrees C / 176 degrees F
    Laser temperature low alarm threshold     :  -10 degrees C / 14 degrees F
    Laser temperature high warning threshold  :  75 degrees C / 167 degrees F
    Laser temperature low warning threshold   :  -5 degrees C / 23 degrees F
  Lane 0
    Laser bias current                        :  27.829 mA
    Laser output power                        :  0.851 mW / -0.70 dBm
    Laser temperature                         :  34 degrees C / 94 degrees F
    Laser receiver power                      :  0.894 mW / -0.49 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
```

```
    APD supply fault alarm                    :  Off
    TEC fault alarm                           :  Off
    Wavelength unlocked alarm                 :  Off
  Lane 1
    Laser bias current                        :  35.374 mA
    Laser output power                        :  0.896 mW / -0.48 dBm
    Laser temperature                         :  34 degrees C / 94 degrees F
    Laser receiver power                      :  0.707 mW / -1.50 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    APD supply fault alarm                    :  Off
    TEC fault alarm                           :  Off
    Wavelength unlocked alarm                 :  Off
  Lane 2
    Laser bias current                        :  29.173 mA
    Laser output power                        :  0.890 mW / -0.51 dBm
    Laser temperature                         :  34 degrees C / 94 degrees F
    Laser receiver power                      :  0.704 mW / -1.52 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
```

```
    Laser output power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    APD supply fault alarm                    :  Off
    TEC fault alarm                           :  Off
    Wavelength unlocked alarm                 :  Off
  Lane 3
    Laser bias current                        :  36.164 mA
    Laser output power                        :  0.899 mW / -0.46 dBm
    Laser temperature                         :  34 degrees C / 94 degrees F
    Laser receiver power                      :  0.892 mW / -0.50 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    APD supply fault alarm                    :  Off
```

```
    TEC fault alarm                          :  Off
    Wavelength unlocked alarm                :  Off
```

show interfaces diagnostics optics (P1-PTX-2-100G-WDM)

```
user@host> show interfaces diagnostics optics et-1/0/0
Physical interface: et-1/0/0
    Module temperature                       :  37 degrees C / 98 degrees F
    Module voltage                           :  3.3370 V
    Module temperature high alarm            :  Off
    Module temperature low alarm             :  Off
    Module temperature high warning          :  Off
    Module temperature low warning           :  Off
    Module voltage high alarm                :  Off
    Module voltage low alarm                 :  Off
    Module voltage high warning              :  Off
    Module voltage low warning               :  Off
    Module not ready alarm                   :  Off
    Module low power alarm                   :  Off
    Module initialization incomplete alarm   :  Off
    Module fault alarm                       :  Off
    PLD Flash initialization fault alarm     :  Off
    Power supply fault alarm                 :  Off
    Checksum fault alarm                     :  Off
    Tx laser disabled alarm                  :  Off
    Tx loss of signal functionality alarm    :  Off
    Tx CDR loss of lock alarm                :  Off
    Rx loss of signal alarm                  :  Off
    Rx CDR loss of lock alarm                :  Off
    Module temperature high alarm threshold  :  70 degrees C / 158 degrees F
    Module temperature low alarm threshold   :  0 degrees C / 32 degrees F
    Module temperature high warning threshold :  68 degrees C / 154 degrees F
    Module temperature low warning threshold  :  2 degrees C / 36 degrees F
    Module voltage high alarm threshold      :  3.4640 V
Module voltage low alarm threshold        :  3.1340 V
    Module voltage high warning threshold    :  3.4310 V
    Module voltage low warning threshold     :  3.1670 V
    Laser bias current high alarm threshold  :  300.000 mA
    Laser bias current low alarm threshold   :  75.000 mA
    Laser bias current high warning threshold :  287.500 mA
    Laser bias current low warning threshold  :  87.500 mA
```

```
Rx power high alarm threshold       :  2.8184 mW / 4.50 dBm
    Rx power low alarm threshold        :  0.0251 mW / -16.00 dBm
    Rx power high warning threshold     :  2.5119 mW / 4.00 dBm
    Rx power low warning threshold      :  0.0501 mW / -13.00 dBm
    LOS alarm threshold                   : 0.0158mW/ -18.01 dBm
    LOS warning threshold               : 0.0251mW/ -16.00 DBm
    Laser temperature high alarm threshold    :  57 degrees C / 135 degrees F
    Laser temperature low alarm threshold     :  25 degrees C / 77 degrees F
    Laser temperature high warning threshold  :  55 degrees C / 131 degrees F
    Laser temperature low warning threshold   :  27 degrees C / 81 degrees F
  Lane 0
    Laser bias current                        :  164.384 mA
Tx power                        :  1.181 mW / 0.72 dBm
    Laser temperature                         :  41 degrees C / 106 degrees F
    Rx power                      :  0.632 mW / -1.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Tx power high alarm             :  Off
    Tx power low alarm              :  Off
    Tx power high warning           :  Off
    Tx power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Rx  power high alarm                    :  Off
    Rx power low alarm                    :  Off
    Rx power high warning                 :  Off
    Rx power low warning                  :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    APD supply fault alarm                    :  Off
    TEC fault alarm                           :  Off
    Wavelength unlocked alarm                 :  Off
```

show interfaces diagnostics optics (P1-PTX-24-10G-W-SFPP)

```
user@host> show interfaces diagnostics optics ge-3/0/6
Physical interface: ge-3/0/6
    Laser bias current                        :  13.356 mA
    Laser output power                        :  0.2210 mW / -6.56 dBm
    Module temperature                        :  36 degrees C / 96 degrees F
    Module voltage                            :  3.2180 V
    Receiver signal average optical power     :  0.2429 mW / -6.15 dBm
    Wavelength Channel number                 :  1
    Wavelength setpoint                       :  1568.80 nm
    Tx Dither                                 :  Disabled
    Frequency Error                           :  0.00 GHz
    Wavelength Error                          :  0.00 nm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    TEC fault alarm                           :  Off
    Wavelength unlocked alarm                 :  Off
    Tx Tune                                   :  Off
    Laser bias current high alarm threshold   :  70.000 mA
    Laser bias current low alarm threshold    :  0.002 mA
    Laser bias current high warning threshold :  65.000 mA
    Laser bias current low warning threshold  :  0.002 mA
    Laser output power high alarm threshold   :  1.0000 mW / 0.00 dBm
```

```
    Laser output power low alarm threshold    :  0.0560 mW / -12.52 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold  :  0.0890 mW / -10.51 dBm
    Module temperature high alarm threshold   :  100 degrees C / 212 degrees F
    Module temperature low alarm threshold    :  -50 degrees C / -58 degrees F
    Module temperature high warning threshold :  95 degrees C / 203 degrees F
    Module temperature low warning threshold  :  -48 degrees C / -54 degrees F
    Module voltage high alarm threshold       :  3.700 V
    Module voltage low alarm threshold        :  2.900 V
    Module voltage high warning threshold     :  3.600 V
    Module voltage low warning threshold      :  3.000 V
    Laser rx power high alarm threshold       :  1.9953 mW / 3.00 dBm
    Laser rx power low alarm threshold        :  0.0001 mW / -40.00 dBm
    Laser rx power high warning threshold     :  1.0000 mW / 0.00 dBm
    Laser rx power low warning threshold      :  0.0010 mW / -30.00 dBm
```

show interfaces diagnostics optics (P2-10G-40G-QSFPP PIC in 40-Gigabit Ethernet mode)

```
user@host> show interfaces diagnostics optics et-0/1/5
Physical interface: et-0/1/5
    Module temperature                        :  30 degrees C / 85 degrees F
    Module voltage                            :  3.2760 V
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  5 degrees C / 41 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.6300 V
    Module voltage low alarm threshold        :  2.9700 V
    Module voltage high warning threshold     :  3.4640 V
    Module voltage low warning threshold      :  3.1340 V
    Laser bias current high alarm threshold   :  10.000 mA
    Laser bias current low alarm threshold    :  0.500 mA
    Laser bias current high warning threshold :  9.500 mA
```

```
    Laser bias current low warning threshold  :  1.000 mA
    Laser output power high alarm threshold   :  0.0000 mW / - Inf dBm
    Laser output power low alarm threshold    :  0.0000 mW / - Inf dBm
    Laser output power high warning threshold :  0.0000 mW / - Inf dBm
    Laser output power low warning threshold  :  0.0000 mW / - Inf dBm
    Laser rx power high alarm threshold       :  2.1878 mW / 3.40 dBm
    Laser rx power low alarm threshold        :  0.0446 mW / -13.51 dBm
    Laser rx power high warning threshold     :  1.7378 mW / 2.40 dBm
    Laser rx power low warning threshold      :  0.1122 mW / -9.50 dBm
  Lane 0
    Laser bias current                        :  7.065 mA
    Laser output power                        :  0.710 mW / -1.49 dBm
    Laser receiver power                      :  0.472 mW / -3.26 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
  Lane 1
    Laser bias current                        :  6.978 mA
    Laser output power                        :  0.771 mW / -1.13 dBm
    Laser receiver power                      :  0.450 mW / -3.47 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
  Lane 2
    Laser bias current                        :  6.955 mA
    Laser output power                        :  0.760 mW / -1.19 dBm
    Laser receiver power                      :  0.556 mW / -2.55 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
```

```
    Laser bias current high warning          :  Off
    Laser bias current low warning           :  Off
    Laser receiver power high alarm          :  Off
    Laser receiver power low alarm           :  Off
    Laser receiver power high warning        :  Off
    Laser receiver power low warning         :  Off
    Tx loss of signal functionality alarm    :  Off
    Rx loss of signal alarm                  :  Off
  Lane 3
    Laser bias current                       :  6.981 mA
    Laser output power                       :  0.736 mW / -1.33 dBm
    Laser receiver power                     :  0.537 mW / -2.70 dBm
    Laser bias current high alarm            :  Off
    Laser bias current low alarm             :  Off
    Laser bias current high warning          :  Off
    Laser bias current low warning           :  Off
    Laser receiver power high alarm          :  Off
    Laser receiver power low alarm           :  Off
    Laser receiver power high warning        :  Off
    Laser receiver power low warning         :  Off
    Tx loss of signal functionality alarm    :  Off
    Rx loss of signal alarm                  :  Off
```

show interfaces diagnostics optics (P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet mode)

```
user@host> show interfaces diagnostics optics et-0/1/5:3
Physical interface: et-0/1/5:3
    Module temperature                       :  30 degrees C / 85 degrees F
    Module voltage                           :  3.2760 V
    Module temperature high alarm            :  Off
    Module temperature low alarm             :  Off
    Module temperature high warning          :  Off
    Module temperature low warning           :  Off
    Module voltage high alarm                :  Off
    Module voltage low alarm                 :  Off
    Module voltage high warning              :  Off
    Module voltage low warning               :  Off
    Module temperature high alarm threshold  :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold   :  5 degrees C / 41 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
```

```
    Module voltage high alarm threshold       :  3.6300 V
    Module voltage low alarm threshold        :  2.9700 V
    Module voltage high warning threshold     :  3.4640 V
    Module voltage low warning threshold      :  3.1340 V
    Laser bias current high alarm threshold   :  10.000 mA
    Laser bias current low alarm threshold    :  0.500 mA
    Laser bias current high warning threshold :  9.500 mA
    Laser bias current low warning threshold  :  1.000 mA
    Laser output power high alarm threshold   :  0.0000 mW / - Inf dBm
    Laser output power low alarm threshold    :  0.0000 mW / - Inf dBm
    Laser output power high warning threshold :  0.0000 mW / - Inf dBm
    Laser output power low warning threshold  :  0.0000 mW / - Inf dBm
    Laser rx power high alarm threshold       :  2.1878 mW / 3.40 dBm
    Laser rx power low alarm threshold        :  0.0446 mW / -13.51 dBm
    Laser rx power high warning threshold     :  1.7378 mW / 2.40 dBm
    Laser rx power low warning threshold      :  0.1122 mW / -9.50 dBm
  Lane 3
    Laser bias current                        :  6.981 mA
    Laser output power                        :  0.736 mW / -1.33 dBm
    Laser receiver power                      :  0.537 mW / -2.70 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
```

show interfaces diagnostics optics (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show interfaces diagnostics optics et-2/0/0
Physical interface: et-2/0/0
    Module temperature                        :  39 degrees C / 102 degrees F
    Module voltage                            :  3.2300 V
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
```

```
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Module not ready alarm                    :  Off
    Module low power alarm                    :  Off
    Module initialization incomplete alarm    :  Off
    Module fault alarm                        :  Off
    Tx laser disabled alarm                   :  Off
    Rx loss of signal alarm                   :  Off
    Module temperature high alarm threshold   :  90 degrees C / 194 degrees F
    Module temperature low alarm threshold    :  -20 degrees C / -4 degrees F
    Module temperature high warning threshold :  0 degrees C / 32 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.6300 V
    Module voltage low alarm threshold        :  2.9700 V
    Module voltage high warning threshold     :  0.0000 V
    Module voltage low warning threshold      :  0.0000 V
    Rx power high alarm threshold             :  6.5535 mW / 8.16 dBm
    Rx power low alarm threshold              :  0.0028 mW / -25.53 dBm
    Rx power high warning threshold           :  6.5535 mW / 8.16 dBm
    Rx power low warning threshold            :  0.0028 mW / -25.53 dBm
    LOS alarm threshold                       :  0.0028 mW / -25.53 dBm
    LOS warning threshold                     :  0.0028 mW / -25.53 dBm
    Modem lock state                          :  OK
  Lane 0
    Tx power                                  :  1.000 mW / 0.00 dBm
    Module temperature                        :  51 degrees C / 124 degrees F
    Rx power (total)                          :  0.644 mW / -1.91 dBm
    Rx power (signal)                         :  0.618 mW / -2.09 dBm
    Lane chromatic dispersion                 :  -22 ps/nm
    Lane differential group delay             :  5 ps
    Lane Q2 factor                            :  14.20 dB
    Lane carrier frequency offset             :  -534 Mz
    Lane electrical SNR                       :  9.20 dB
    Tx power high alarm                       :  Off
    Tx power low alarm                        :  Off
    Tx power high warning                     :  Off
    Tx power low warning                      :  Off
    Rx power high alarm                       :  Off
    Rx power low alarm                        :  Off
    Rx power high warning                     :  Off
```

```
    Rx power low warning                      :  Off
    Rx loss of signal alarm                   :  Off
    Wavelength unlocked alarm                 :  Off
```

show interfaces diagnostics optics (MX960 Router with MPC10E-10C-MRATE, MPC10E-15C-MRATE line cards)

```
user@host> show interfaces diagnostics optics et-0/2/4

Physical interface: et-0/2/4
    Module temperature                        :  58 degrees C / 136 degrees F
    Module voltage                            :  3.2440 V
    Wavelength channel number                 :  28
    Wavelength setpoint                       :  1550.12 nm
    Tx dither                                 :  Disabled
    Frequency error                           :  0.00 GHz
    Wavelength error                          :  0.00 nm
    TEC fault alarm                           :  False
    Wavelength unlocked alarm                 :  False
    Tx tune alarm                             :  False
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Module temperature high alarm threshold   :  76 degrees C / 169 degrees F
    Module temperature low alarm threshold    :  -6 degrees C / 21 degrees F
    Module temperature high warning threshold :  73 degrees C / 163 degrees F
    Module temperature low warning threshold  :  -3 degrees C / 27 degrees F
    Module voltage high alarm threshold       :  3.7000 V
    Module voltage low alarm threshold        :  2.9990 V
    Module voltage high warning threshold     :  3.5990 V
    Module voltage low warning threshold      :  3.0490 V
    Laser bias current high alarm threshold   :  Not supported
    Laser bias current low alarm threshold    :  Not supported
    Laser bias current high warning threshold :  Not supported
    Laser bias current low warning threshold  :  Not supported
    Laser output power high alarm threshold   :  0.1995 mW / -7.00 dBm
```

```
    Laser output power low alarm threshold    :  0.0158 mW / -18.01 dBm
    Laser output power high warning threshold :  0.1585 mW / -8.00 dBm
    Laser output power low warning threshold  :  0.0200 mW / -16.99 dBm
    Laser rx power high alarm threshold       :  2.5119 mW / 4.00 dBm
    Laser rx power low alarm threshold        :  0.0050 mW / -23.01 dBm
    Laser rx power high warning threshold     :  1.9953 mW / 3.00 dBm
    Laser rx power low warning threshold      :  0.0100 mW / -20.00 dBm
    Laser temperature high alarm threshold    :  0 degrees C / 32 degrees F
    Laser temperature low alarm threshold     :  0 degrees C / 32 degrees F
    Laser temperature high warning threshold  :  0 degrees C / 32 degrees F
    Laser temperature low warning threshold   :  0 degrees C / 32 degrees F
  Lane 0
    Laser bias current                        :  242.600 mA
    Laser output power                        :  0.131 mW / -8.82 dBm
    Laser receiver power                      :  0.119 mW / -9.23 dBm
    Laser temperature                         :  58 degrees C / 136 degrees F
    Lane chromatic dispersion                 :  0.0 ps/nm
    Lane differential group delay             :  39.1 ps
    Lane carrier frequency offset             :  -257.0 MHz
    Lane polarization dependent loss          :  0.5 dB
    Lane snr                                  :  34.50 dB
    Lane Optical signal-to-noise ratio        :  37.5 dB
    Lane sopmd                                :  0.0
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
    Tx laser disabled alarm                   :  Off
```

show interfaces diagnostics optics (PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC

```
user@host> show interfaces diagnostics optics et-4/0/0
Physical interface: et-4/0/0
    Laser output power                        :  54 degrees C / 129 degrees F
```

```
    Tx module temperature                     :  0.0000
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Module not ready alarm                    :  Off
    Module low power alarm                    :  Off
    Module initialization incomplete alarm    :  Off
    Module fault alarm                        :  Off
    PLD Flash initialization fault alarm      :  Off
    Power supply fault alarm                  :  Off
    Checksum fault alarm                      :  Off
    Tx laser disabled alarm                   :  Off
    Rx loss of signal alarm                   :  Off
    Module temperature high alarm threshold   :  80 degrees C / 176 degrees F
    Module temperature low alarm threshold    :  0 degrees C / 32 degrees F
    Module temperature high warning threshold :  65 degrees C / 149 degrees F
    Module temperature low warning threshold  :  5 degrees C / 41 degrees F
    Module voltage high alarm threshold       :  0.0000 V
    Module voltage low alarm threshold        :  0.0000 V
    Module voltage high warning threshold     :  0.0000 V
    Module voltage low warning threshold      :  0.0000 V
    Rx power high alarm threshold             :  0.0000 mW / - Inf dBm
    Rx power low alarm threshold              :  0.0000 mW / - Inf dBm
    Rx power high warning threshold           :  0.0000 mW / - Inf dBm
    Rx power low warning threshold            :  0.0000 mW / - Inf dBm
    LOS alarm threshold                       :  0.0158 mW / -18.01 dBm
    LOS warning threshold                     :  0.0251 mW / -16.00 dBm
    Modem lock state                          :  OK
  Lane 0
    Tx power                                  :  1.000 mW / 0.00 dBm
    Module temperature                        :  0 degrees C / 32 degrees F
    Rx power (total)                          :  0.000 mW / - Inf dBm
    Rx power (signal)                         :  0.999 mW / -0.00 dBm
    Lane chromatic dispersion                 :  6 ps/nm
    Lane differential group delay             :  3 ps
    Lane Q2 factor                            :  15.40 dB
    Lane carrier frequency offset             :  0 MHz
    Lane electrical SNR                       :  16.60 dB
```

```
    Tx power high alarm                        :  Off
    Tx power low alarm                         :  Off
    Tx power high warning                      :  Off
    Tx power low warning                       :  Off
    Rx power high alarm                        :  Off
    Rx power low alarm                         :  Off
    Rx power high warning                      :  Off
Rx power low warning                       :  Off
    Rx loss of signal alarm                    :  Off
    Wavelength unlocked alarm                  :  Off
    Laser end-of-life alarm                    :  Off
  Lane 1
    Tx power                                   :  1.000 mW / 0.00 dBm
    Module temperature                         :  0 degrees C / 32 degrees F
    Rx power (total)                           :  0.000 mW / - Inf dBm
    Rx power (signal)                          :  0.999 mW / -0.00 dBm
    Tx power high alarm                        :  Off
    Tx power low alarm                         :  Off
    Tx power high warning                      :  Off
    Tx power low warning                       :  Off
    Rx power high alarm                        :  Off
    Rx power low alarm                         :  Off
    Rx power high warning                      :  Off
    Rx power low warning                       :  Off
    Rx loss of signal alarm                    :  Off
    Wavelength unlocked alarm                  :  Off
    Laser end-of-life alarm                    :  Off
  Lane 2
    Tx power                                   :  1.000 mW / 0.00 dBm
    Module temperature                         :  0 degrees C / 32 degrees F
    Rx power (total)                           :  0.000 mW / - Inf dBm
    Rx power (signal)                          :  0.999 mW / -0.00 dBm
    Tx power high alarm                        :  Off
    Tx power low alarm                         :  Off
    Tx power high warning                      :  Off
    Tx power low warning                       :  Off
    Rx power high alarm                        :  Off
    Rx power low alarm                         :  Off
    Rx power high warning                      :  Off
    Rx power low warning                       :  Off
    Rx loss of signal alarm                    :  Off
    Wavelength unlocked alarm                  :  Off
    Laser end-of-life alarm                    :  Off
```

```
  Lane 3
    Tx power                                  :  1.000 mW / 0.00 dBm
    Module temperature                        :  0 degrees C / 32 degrees F
    Rx power (total)                          :  0.000 mW / - Inf dBm
    Rx power (signal)                         :  0.999 mW / -0.00 dBm
    Tx power high alarm                       :  Off
    Tx power low alarm                        :  Off
    Tx power high warning                     :  Off
    Tx power low warning                      :  Off
Rx power high alarm                       :  Off
    Rx power low alarm                        :  Off
    Rx power high warning                     :  Off
    Rx power low warning                      :  Off
    Rx loss of signal alarm                   :  Off
    Wavelength unlocked alarm                 :  Off
    Laser end-of-life alarm                   :  Off
```

show interfaces diagnostics optics (for VCP)

```
user@host> show interfaces diagnostics optics vcp-2/0/1
Physical interface: vcp-2/0/1
    Laser bias current                        :  5.494 mA
    Laser output power                        :  0.2960 mW / -5.29 dBm
    Module temperature                        :  22 degrees C / 71 degrees F
    Module voltage                            :  3.2810 V
    Receiver signal average optical power     :  0.2426 mW / -6.15 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
```

```
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces diagnostics optics (MPC7 with interfaces disabled)

```
user@host> show interfaces diagnostics optics et-3/0/0
Physical interface: et-3/0/0
    Module temperature                        :  34 degrees C / 93 degrees F
    Module voltage                            :  3.2660 V
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
```

```
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.6300 V
    Module voltage low alarm threshold        :  2.9700 V
    Module voltage high warning threshold     :  3.4640 V
    Module voltage low warning threshold      :  3.1340 V
    Laser bias current high alarm threshold   :  9.999 mA
    Laser bias current low alarm threshold    :  0.499 mA
    Laser bias current high warning threshold :  9.499 mA
    Laser bias current low warning threshold  :  0.999 mA
    Laser output power high alarm threshold   :  0.0000 mW / - Inf dBm
    Laser output power low alarm threshold    :  0.0000 mW / - Inf dBm
    Laser output power high warning threshold :  0.0000 mW / - Inf dBm
    Laser output power low warning threshold  :  0.0000 mW / - Inf dBm
    Laser rx power high alarm threshold       :  2.1878 mW / 3.40 dBm
    Laser rx power low alarm threshold        :  0.0446 mW / -13.51 dBm
    Laser rx power high warning threshold     :  1.7378 mW / 2.40 dBm
    Laser rx power low warning threshold      :  0.1122 mW / -9.50 dBm
  Lane 0
    Laser bias current                        :  6.697 mA
    Laser output power                        :  0.738 mW / -1.32 dBm
    Laser receiver power                      :  0.790 mW / -1.02 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
    Tx laser disabled alarm                   :  Off
  Lane 1
    Laser bias current                        :  6.961 mA
    Laser output power                        :  0.908 mW / -0.42 dBm
    Laser receiver power                      :  0.827 mW / -0.83 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
```

```
    Laser receiver power high alarm          :  Off
    Laser receiver power low alarm           :  Off
    Laser receiver power high warning        :  Off
    Laser receiver power low warning         :  Off
    Tx loss of signal functionality alarm    :  Off
    Rx loss of signal alarm                  :  Off
    Tx laser disabled alarm                  :  Off
  Lane 2
    Laser bias current                       :  6.926 mA
    Laser output power                       :  0.888 mW / -0.51 dBm
    Laser receiver power                     :  0.820 mW / -0.86 dBm
    Laser bias current high alarm            :  Off
    Laser bias current low alarm             :  Off
    Laser bias current high warning          :  Off
    Laser bias current low warning           :  Off
    Laser receiver power high alarm          :  Off
    Laser receiver power low alarm           :  Off
    Laser receiver power high warning        :  Off
    Laser receiver power low warning         :  Off
    Tx loss of signal functionality alarm    :  Off
    Rx loss of signal alarm                  :  Off
    Tx laser disabled alarm                  :  Off
  Lane 3
    Laser bias current                       :  6.817 mA
    Laser output power                       :  0.846 mW / -0.73 dBm
    Laser receiver power                     :  0.827 mW / -0.82 dBm
    Laser bias current high alarm            :  Off
    Laser bias current low alarm             :  Off
    Laser bias current high warning          :  Off
    Laser bias current low warning           :  Off
    Laser receiver power high alarm          :  Off
    Laser receiver power low alarm           :  Off
    Laser receiver power high warning        :  Off
    Laser receiver power low warning         :  Off
    Tx loss of signal functionality alarm    :  Off
    Rx loss of signal alarm                  :  Off
    Tx laser disabled alarm                  :  Off
```

show interfaces diagnostics optics for 100-Gigabit Ethernet (QSFP-100GE-DWDM2)

```
user@host> show interfaces diagnostics optics et-18/0/2
Physical interface: et-18/0/2
    Module temperature                        :  37 degrees C / 98 degrees F
    Module voltage                            :  3.2770 V
    Grid Channel Number                       :  1 (191.40 THz)
    Corrected Error Ratio                     :  (6379 sec average) 3.29e-05
    Uncorrected Words Ratio                   :  (6379 sec average) 0.0e-05
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Module temperature high alarm threshold   :  79 degrees C / 174 degrees F
    Module temperature low alarm threshold    :  -4 degrees C / 25 degrees F
    Module temperature high warning threshold :  75 degrees C / 167 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.6300 V
    Module voltage low alarm threshold        :  2.9700 V
    Module voltage high warning threshold     :  3.4640 V
    Module voltage low warning threshold      :  3.1340 V
    Laser bias current high alarm threshold   :  109.999 mA
    Laser bias current low alarm threshold    :  19.999 mA
    Laser bias current high warning threshold :  99.999 mA
    Laser bias current low warning threshold  :  29.999 mA
    Laser output power high alarm threshold   :  3.5481 mW / 5.50 dBm
    Laser output power low alarm threshold    :  0.2344 mW / -6.30 dBm
    Laser output power high warning threshold :  2.8184 mW / 4.50 dBm
    Laser output power low warning threshold  :  0.2951 mW / -5.30 dBm
    Laser rx power high alarm threshold       :  3.5481 mW / 5.50 dBm
    Laser rx power low alarm threshold        :  0.0436 mW / -13.61 dBm
Laser rx power high warning threshold     :  2.8183 mW / 4.50 dBm
    Laser rx power low warning threshold      :  0.0871 mW / -10.60 dBm
  Lane 0
    Laser bias current                        :  73.804 mA
    Laser output power                        :  0.948 mW / -0.23 dBm
    Laser receiver power                      :  0.003 mW / -25.23 dBm
    Lane carrier frequency offset             :  362 MHz
```

```
    Lane SNR                                  :  9.60 dB
    TEC Current                               :  0.1 mA
    Residual ISI                              :  125
    PAM Histogram                             :  235
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  On
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  On
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  On
    Tx laser disabled alarm                   :  Off
  Lane 1
    Laser bias current                        :  83.363 mA
    Laser output power                        :  1.042 mW / 0.18 dBm
    Laser receiver power                      :  0.000 mW / - Inf dBm
    Lane carrier frequency offset             :  362 MHz
    Lane SNR                                  :  9.60 dB
    TEC Current                               :  0.1 mA
    Residual ISI                              :  125
    PAM Histogram                             :  235
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  On
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  On
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  On
    Tx laser disabled alarm                   :  Off
```

show interfaces diagnostics optics (PTX10001-36MR) (400G ZR optics)

```
user@host> show interfaces diagnostics optics et-0/2/1
Physical interface: et-0/2/1
     Module temperature                        :  59 degrees C / 138 degrees F
```

```
    Module voltage                            :  3.247 V
    Wavelength channel number                 :  28
    Wavelength setpoint                       :  1550.12 nm
    Tx dither                                 :  Disabled
    Frequency error                           :  0.00 GHz
    Wavelength error                          :  0.0 nm
    TEC fault alarm                           :  False
    Wavelength unlocked alarm                 :  False
    Tx tune alarm                             :  False
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Module temperature high alarm threshold   :  76 degrees C / 168 degrees F
    Module temperature low alarm threshold    :  -6 degrees C / 21 degrees F
    Module temperature high warning threshold :  73 degrees C / 163 degrees F
    Module temperature low warning threshold  :  -3 degrees C / 26 degrees F
    Module voltage high alarm threshold       :  3.7 V
    Module voltage low alarm threshold        :  3.0 V
    Module voltage high warning threshold     :  3.599 V
    Module voltage low warning threshold      :  3.049 V
    Laser bias current high alarm threshold   :  Not supported
    Laser bias current low alarm threshold    :  Not supported
    Laser bias current high warning threshold :  Not supported
    Laser bias current low warning threshold  :  Not supported
    Laser output power high alarm threshold   :  1.0 mW / 0.00 dBm
    Laser output power low alarm threshold    :  0.015 mW / -18.01 dBm
    Laser output power high warning threshold :  0.63 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.025 mW / -16.00 dBm
    Laser rx power high alarm threshold       :  1.584 mW / 1.99 dBm
    Laser rx power low alarm threshold        :  0.004 mW / -23.01 dBm
    Laser rx power high warning threshold     :  1.0 mW / 0.00 dBm
    Laser rx power low warning threshold      :  0.009 mW / -20.00 dBm
    Laser temperature high alarm threshold    :  80 degrees C / 176 degrees F
    Laser temperature low alarm threshold     :  -80 degrees C / -112 degrees F
    Laser temperature high warning threshold  :  75 degrees C / 167 degrees F
    Laser temperature low warning threshold   :  -40 degrees C / -40 degrees F
  Lane 0
    Laser bias current                        :  240.199 mA
```

```
    Laser output power                        :  0.153 mW / -8.14 dBm
    Laser temperature                         :  52 degrees C / 125 degrees F
    Laser receiver power                      :  0.134 mW / -8.69 dBm
    Lane chromatic dispersion                 :  1.0 ps/nm
    Lane differential group delay             :  4.0 ps
    Lane carrier frequency offset             :  0.0 MHz
    Lane polarization dependent loss          :  0.8 dB
    Lane snr                                  :  17.9 dB
    Lane Optical signal-to-noise ratio        :  36.3 dB
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
    Tx laser disabled alarm                   :  Off
```

show interfaces diagnostics optics (PTX10001-36MR) (400G ZR-M optics)

```
user@host> show interfaces diagnostics optics et-0/0/10
Physical interface: et-0/0/10
    Module temperature                        :  56 degrees C / 132 degrees F
    Module voltage                            :  3.202 V
    Wavelength channel number                 :  28
    Wavelength setpoint                       :  1550.12 nm
    Tx dither                                 :  Disabled
    Frequency error                           :  0.00 GHz
    Wavelength error                          :  0.0 nm
    TEC fault alarm                           :  False
```

```
    Wavelength unlocked alarm                 :  False
    Tx tune alarm                             :  False
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Module temperature high alarm threshold   :  76 degrees C / 168 degrees F
    Module temperature low alarm threshold    :  -6 degrees C / 21 degrees F
    Module temperature high warning threshold :  73 degrees C / 163 degrees F
    Module temperature low warning threshold  :  -3 degrees C / 26 degrees F
    Module voltage high alarm threshold       :  3.7 V
    Module voltage low alarm threshold        :  3.0 V
    Module voltage high warning threshold     :  3.599 V
    Module voltage low warning threshold      :  3.049 V
    Laser bias current high alarm threshold   :  Not supported
    Laser bias current low alarm threshold    :  Not supported
    Laser bias current high warning threshold :  Not supported
    Laser bias current low warning threshold  :  Not supported
    Laser output power high alarm threshold   :  1.0 mW / 0.00 dBm
    Laser output power low alarm threshold    :  0.015 mW / -18.01 dBm
    Laser output power high warning threshold :  0.63 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.025 mW / -16.00 dBm
    Laser rx power high alarm threshold       :  1.584 mW / 1.99 dBm
    Laser rx power low alarm threshold        :  0.001 mW / -28.23 dBm
    Laser rx power high warning threshold     :  1.0 mW / 0.00 dBm
    Laser rx power low warning threshold      :  0.004 mW / -23.01 dBm
    Laser temperature high alarm threshold    :  80 degrees C / 176 degrees F
    Laser temperature low alarm threshold     :  -80 degrees C / -112 degrees F
    Laser temperature high warning threshold  :  75 degrees C / 167 degrees F
    Laser temperature low warning threshold   :  -40 degrees C / -40 degrees F
  Lane 0
    Laser bias current                        :  235.695 mA
    Laser output power                        :  0.086 mW / -10.62 dBm
    Laser temperature                         :  52 degrees C / 125 degrees F
    Laser receiver power                      :  0.073 mW / -11.33 dBm
    Lane chromatic dispersion                 :  2.0 ps/nm
    Lane differential group delay             :  3.0 ps
    Lane carrier frequency offset             :  -66.0 MHz
    Lane polarization dependent loss          :  0.6 dB
```

```
    Lane snr                                  :  19.0 dB
    Lane Optical signal-to-noise ratio        :  35.8 dB
    Lane sopmd                                :  34.0
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
    Tx laser disabled alarm                   :  Off
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 12.1 for PTX Series routers.

Command introduced in Junos OS Release 19.2R1 for QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 routers.

RELATED DOCUMENTATION

Determining Transceiver Support and Specifications

show interfaces diagnostics optics-profile

IN THIS SECTION

- Syntax | 20645
- Description | 20645
- Options | 20645
- Additional Information | 20645
- Required Privilege Level | 20645
- Output Fields | 20646
- Sample Output | 20648
- Release Information | 20650

Syntax

```
show interfaces diagnostics optics-profile
```

Description

Display diagnostic data for the supported optics (QSFP+, QSFP28, QSFP-DD). No output is displayed for unsupported optics (SFP, CFP, etc.) and for QSFP DAC (direct attach cables) modules.

Options

interface-name	Interface name. For example: et-*fpc*/*pic*/*port*

Additional Information

Required Privilege Level

view

Output Fields

Table 632 on page 20646 lists the output fields for the `show interfaces diagnostics optics-profile` command.

Table 632: show interfaces diagnostics optics-profile Output Fields

Field Name	Field Description
Optics type	Type of optics supported. Supported optics are QSFP+, QSFP28, and QSFP-DD.
Number of lanes	Number of lanes depending on the optics connected. Number of lanes supported is 0 through 7.
Power class	Specifies the power class. Possible values: 1 through 8.
Power Mode	Specifies the power mode. Possible values: High or Low.
Host Side FEC	Specifies the FEC on host side optics. Possible values: Off, On, or Unsupported.
Media Side FEC	Specifies the FEC on media side optics. Possible values: Off, On, or Unsupported.
Rx output manual control support	Specifies if the optics can receive output. Possible values: Yes or No.
Tx CDR Implemented	Specifies if the transmit clock and data recovery is implemented. Possible values: Yes or No.
Tx CDR manual control support	Specifies if the transmit clock and data recovery can be configured. Possible values: Yes or No.
Rx CDR Implemented	Specifies if the receive clock and data recovery status is implemented. Possible values: Yes or No.
Rx CDR manual control support	Specifies if the receive clock and data recovery can be configured. Possible values: Yes or No.
Tx input equalization type	Specifies the Input Equalization type. Also called Continous Time Linear Equalization (CTLE). Possible values: Manual or Adaptive.
Tx input equalization manual control	Specifies if the input equalization is manual. Possible values: Yes or No.

Table 632: show interfaces diagnostics optics-profile Output Fields *(Continued)*

Field Name	Field Description
Tx input equalization max value	Specifies the maximum equalization value. Possible values: 0 through 10 dB.
Tx input equalization max settling time	Specifies the time required for the maximum equalization value. Possible values: 100 through 25500 ms.
Rx output emphasis manual control support	Specifies if the output emphasis can be configured. Possible values: Yes or No.
Rx output amplitude 100-400 mV no equalization	Specifies that there is no equalization for 100mV to 400mV amplitude. Possible values: Off or On.
Rx output amplitude 300-600 mV no equalization	Specifies that there is no equalization for 300mV to 600mV amplitude. Possible values: Off or On.
Rx output amplitude 400-800 mV no equalization	Specifies that there is no equalization for 400mV to 800mV amplitude. Possible values: Off or On.
Rx output amplitude 600-1200 mV no equalization	Specifies that there is no equalization for 600mV to 1200mV amplitude. Possible values: Off or On.
Tx squelch manual control support	Specifies if the transmit squelch status is configured. Possible values: Off or On.
Rx squelch manual control support	Specifies if the receive squelch status is configured. Possible values: Off or On.
Tx rate select version	Version of the transmit rate select option. Possible values: Version 1, Version 2, and not supported.
Rx rate select version	Version of the receive rate select option. Possible values: Version 1, Version 2, and not supported.
Lane	Lane number. Possible values: 0 through 7.

Table 632: show interfaces diagnostics optics-profile Output Fields *(Continued)*

Field Name	Field Description
Tx Laser	Specifies the status of the transmit Laser. Possible values: Off or On.
Rx Output	Specifies if the optics can receive output. Possible values: Off or On.
Tx CDR	Specifies if the transmit clock and data recovery status of the optics. Possible values: Off or On.
Rx CDR	Specifies if the receive clock and data recovery status of the optics. Possible values: Off or On.
Tx input equalization (CTLE)	Specifies if input equalization is configured. Possible values: 0 through 10 dB.
Rx output emphasis	Specifies if the optics can receive output emphasis. Possible values : 0 through 10 dB.
Rx output amplitude	Specifies if the optics can receive output amplitude setting. Possible values: • 1200mV—No equalization for 600mV - 1200mV amplitude. • 400mV—No equalization for 100mV - 400mV amplitude. • 600mV—No equalization for 300mV - 600mV amplitude. • 800mV—No equalization for 400mV - 800mV amplitude.
Tx squelch	Specifies the transmit squelch status of the optics. Possible values: Off or On.
Rx squelch	Specifies the receive squelch status of the optics. Possible values: Off or On.

Sample Output

show interfaces diagnostics optics-profile

```
user@host> show interfaces diagnostics optics-profile et-0/0/0
```

```
Physical interface: et-0/0/0
    Optics type                                       :  QSFP28
```

```
  Number of lanes                                    :  4
  Power class                                        :  class 4 (3.5 W max)
  Power mode                                         :  high
  Host side FEC                                      :  not supported
  Media side FEC                                     :  not supported
  Rx output manual control support                   :  yes
  Tx CDR implemented                                 :  yes
  Tx CDR manual control support                      :  yes
  Rx CDR implemented                                 :  yes
  Rx CDR manual control support                      :  yes
  Tx input equalization type                         :  manual
  Tx input equalization manual control support       :  yes
  Tx input equalization max value                    :  10 dB
  Tx input equalization max settling time            :  100 ms
  Rx output emphasis manual control support          :  yes
  Rx output emphasis max value                       :  6 dB
  Rx output amplitude 100-400 mV no equalization     :  yes
  Rx output amplitude 300-600 mV no equalization     :  yes
  Rx output amplitude 400-800 mV no equalization     :  yes
  Rx output amplitude 600-1200 mV no equalization    :  yes
  Tx squelch manual control support                  :  no
  Rx squelch manual control support                  :  yes
  Tx rate select version                             :  not supported
  Rx rate select version                             :  not supported
Lane 0
    Tx laser                                         :  enabled
    Rx output                                        :  enabled
    Tx CDR                                           :  enabled
    Rx CDR                                           :  enabled
    Tx input equalization (CTLE)                     :  1 dB
    Rx output emphasis                               :  0 dB
    Rx output amplitude                              :  400-800 mV
    Tx squelch                                       :  enabled
    Rx squelch                                       :  enabled
Lane 1
    Tx laser                                         :  enabled
    Rx output                                        :  enabled
    Tx CDR                                           :  enabled
    Rx CDR                                           :  enabled
    Tx input equalization (CTLE)                     :  1 dB
    Rx output emphasis                               :  0 dB
    Rx output amplitude                              :  400-800 mV
    Tx squelch                                       :  enabled
```

```
    Rx squelch                                      :  enabled
  Lane 2
    Tx laser                                        :  enabled
    Rx output                                       :  enabled
    Tx CDR                                          :  enabled
    Rx CDR                                          :  enabled
    Tx input equalization (CTLE)                    :  1 dB
    Rx output emphasis                              :  0 dB
    Rx output amplitude                             :  400-800 mV
    Tx squelch                                      :  enabled
    Rx squelch                                      :  enabled
  Lane 3
    Tx laser                                        :  enabled
    Rx output                                       :  enabled
    Tx CDR                                          :  enabled
    Rx CDR                                          :  enabled
    Tx input equalization (CTLE)                    :  1 dB
    Rx output emphasis                              :  0 dB
    Rx output amplitude                             :  400-800 mV
    Tx squelch                                      :  enabled
    Rx squelch                                      :  enabled
```

show interfaces diagnostics optics-profile (for SFP, CFP, QSFP DAC)

```
user@host> show interfaces diagnostics optics-profile et-0/0/0
```

```
Physical interface: et-0/0/0
    Optical profile diagnostics                     :  Not supported
```

Release Information

Command introduced in Junos OS Evolved Release 21.1R1

RELATED DOCUMENTATION

lane | 9623

lane-all | 9626

optics-options | 11847

show interfaces (Discard)

IN THIS SECTION

- Syntax | 20651
- Description | 20651
- Options | 20651
- Required Privilege Level | 20652
- Output Fields | 20652
- Sample Output | 20657
- Release Information | 20659

Syntax

```
show interfaces dsc
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified discard interface.

Options

dsc	Display standard information about the specified discard interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.

descriptions	(Optional) Display interface description strings.
media	This option is not relevant for the discard interface and always shows a value of 0.
snmp-index ***snmp-index***	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) This option is not relevant for the discard interface and always shows a value of 0.

Required Privilege Level

view

Output Fields

Table 633 on page 20652 lists the output fields for the `show interfaces` (discard) command. Output fields are listed in the approximate order in which they appear.

Table 633: Discard show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface, whether the interface is enabled, and the state of the physical interface: **Up** or **Down**.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Type	Type of interface. **Software-Pseudo** indicates a standard software interface with no associated hardware device.	All levels

Table 633: Discard show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Clocking	Reference clock source. It can be **Internal** or **External**.	**brief detail extensive**
Speed	Speed at which the interface is running.	**brief detail extensive**
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
Link type	Encapsulation being used on the physical interface.	**detail extensive**
Link flags	Information about the link. Possible values are described in the “Link Flags” section under *Common Output Fields Description*.	**detail extensive**
Physical info	Information about the physical interface.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	**detail extensive**
Current address, Hardware address	Configured MAC address and hardware MAC address.	**detail extensive**
Alternate link address	Backup address of the link.	**detail extensive**

Table 633: Discard show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour.minute.second timezone*** **(*hour.minute.second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets, Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**

Table 633: Discard show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface: • **Errors**—Sum of incoming frame terminated and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**detail extensive**

Table 633: Discard show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	(Extensive only) Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • **Errors**—Sum of the outgoing frame terminated and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive**
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	All levels

Table 633: Discard show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, or **mpls**.	**All levels**
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**

Sample Output

show interfaces dsc

```
user@host> show interfaces dsc
Physical interface: dsc, Enabled, Physical link is Up
  Interface index: 5, SNMP ifIndex: 5
  Type: Software-Pseudo, MTU: Unlimited
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface dsc.0 (Index 66) (SNMP ifIndex 235)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
    Protocol inet, MTU: Unlimited
      Flags: None
```

show interfaces dsc brief

```
user@host> show interfaces dsc brief
Physical interface: dsc, Enabled, Physical link is Up
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking: Unspecified,
Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

  Logical interface dsc.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
    inet
```

show interfaces dsc detail

```
user@host> show interfaces dsc detail
Physical interface: dsc, Enabled, Physical link is Up
  Interface index: 5, SNMP ifIndex: 5, Generation: 9
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking: Unspecified,
Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Unspecified
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0

  Logical interface dsc.0 (Index 66) (SNMP ifIndex 235) (Generation 6)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
    Protocol inet, MTU: Unlimited, Generation: 14, Route table: 0
      Flags: None
```

show interfaces dsc extensive

```
user@host> show interfaces dsc extensive
Physical interface: dsc, Enabled, Physical link is Up
  Interface index: 5, SNMP ifIndex: 5, Generation: 9
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking: Unspecified,
Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Unspecified
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0
  Logical interface dsc.0 (Index 66) (SNMP ifIndex 235) (Generation 6)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
    Protocol inet, MTU: Unlimited, Generation: 14, Route table: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show interfaces (ATM) | 20442

show interfaces routing | 21316

show interfaces (Dynamic Flow Capture)

IN THIS SECTION

- Syntax | 20660
- Description | 20660
- Options | 20661
- Required Privilege Level | 20661
- Output Fields | 20661
- Sample Output | 20666
- Release Information | 20667

Syntax

```
show interfaces dfc-fpc/pic/port:channel
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M320 and M120 Series routers and T Series routers only) Display status information about the specified dynamic flow capture interface.

Options

dfc-*fpc/pic/port:channel*	Display standard status information about the specified dynamic flow capture interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 634 on page 20661 lists the output fields for the `show interfaces` (Dynamic Flow Capture) command. Output fields are listed in the approximate order in which they appear.

Table 634: Dynamic Flow Capture show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	Sate of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none

Table 634: Dynamic Flow Capture show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Type	Type of interface.	All levels
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum Transmit Unit (MTU). Size of the largest packet to be transmitted.	All levels
Speed	Network speed on the interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
Link type	Data transmission type.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under *Common Output Fields Description*.	All levels
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour.minute.second timezone*** **(*hour.minute.second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive**
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified

Table 634: Dynamic Flow Capture show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input rate**, **Output rate**—Number of bits per second (packets per second) received and transmitted on the interface. • **Input packets**, **Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**
Input errors	• **Errors**—Input errors on the interface. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Frames received smaller than the runt threshold. • **Giants**—Frames received larger than the giant threshold. • **Policed Discards**—Frames that the incoming packet match code discarded because the frames did not recognize them or were not of interest. Usually, this field reports protocols that the Junos OS does not support. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 634: Dynamic Flow Capture show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	• **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • **Errors**—Sum of outgoing frame terminates and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC RED mechanism. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Flags	Information about the logical interface; values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified

Table 634: Dynamic Flow Capture show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets, Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**
Protocol	Protocol family configured on the logical interface (such as **iso** or **inet6**).	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive** none
Addresses, Flags	Addresses associated with the logical interface and information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none

Sample Output

show interfaces (Dynamic Flow Capture)

```
user@host> show interfaces dfc-0/0/0
Physical interface: dfc-0/0/0, Enabled, Physical link is Up
 Interface index: 146, SNMP ifIndex: 36
 Type: Adaptive-Services, Link-level type: Dynamic-Flow-Capture, MTU: 9192, Speed: 2488320kbps
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps 16384
 Link type : Full-Duplex
 Link flags : None
 Last flapped : 2005-08-26 15:08:36 PDT (01:18:42 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 44800440 bps (100000 pps)

Logical interface dfc-0/0/0.0 (Index 67) (SNMP ifIndex 43)
 Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
Input packets : 74
Output packets: 132
 Protocol inet, MTU: 9192
  Flags: Receive-options, Receive-TTL-Exceeded
  Addresses, Flags: Is-Preferred Is-Primary
   Destination: 10.36.100.1, Local: 10.36.100.2

Logical interface dfc-0/0/0.1 (Index 68) (SNMP ifIndex 49)
 Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
Input packets : 0
Output packets: 402927263
 Protocol inet, MTU: 9192
  Flags: Receive-options, Receive-TTL-Exceeded

Logical interface dfc-0/0/0.2 (Index 69) (SNMP ifIndex 50)
 Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
Input packets : 0
Output packets: 0
 Protocol inet, MTU: 9192
  Flags: Receive-options, Receive-TTL-Exceeded

Logical interface dfc-0/0/0.16383 (Index 70) (SNMP ifIndex 44)
 Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
Input packets : 1427
```

```
Output packets: 98
 Protocol inet, MTU: 9192
  Flags: Receive-options, Receive-TTL-Exceeded
  Addresses, Flags: Is-Preferred Is-Primary
   Destination: 10.0.0.16, Local: 10.0.0.1
```

Release Information

Command introduced in Junos OS Release 7.4.

show interfaces (Encryption)

IN THIS SECTION

- Syntax | 20667
- Description | 20668
- Options | 20668
- Required Privilege Level | 20668
- Output Fields | 20668
- Sample Output | 20672
- Release Information | 20676

Syntax

```
show interfaces es-fpc/pic/port:channel
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M Series and T Series routers only) Display status information about the specified encryption interface.

Options

es-*fpc/pic/port:channel*	Display standard status information about the specified encryption interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 635 on page 20668 lists the output fields for the `show interfaces` (Encryption) command. Output fields are listed in the approximate order in which they appear.

Table 635: Encryption show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels

Table 635: Encryption show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Encapsulation being used on the interface.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Device flags	Information about the physical device. Possible values are described in the “Link Flags” section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 635: Encryption show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets, Output packets**—Number of packets received and transmitted on the interface. • **Anti-replay failures**—Total number of antireplay failures seen on all tunnels configured on the ES PIC. • **Authentication**—Total number of authentication failures seen on all tunnels configured on the ES PIC.	**detail extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**
Queue counters	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Table 635: Encryption show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
IP-Header	IP header of the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, **mpls**.	**detail extensive** none

Table 635: Encryption show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.Address	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces (Encryption)

```
user@host> show interfaces es-0/3/0
Physical interface: es-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 71
```

```
  Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface es-0/3/0.0 (Index 70) (SNMP ifIndex 45)
    Flags: Hardware-Down Point-To-Point SNMP-Traps
    IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 3800
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.10.0.2, Local: 10.10.0.1
```

show interfaces brief (Encryption)

```
user@host> show interfaces es-0/3/0 brief
Physical interface: es-0/3/0, Enabled, Physical link is Up
  Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

  Logical interface es-0/3/0.0
    Flags: Hardware-Down Point-To-Point SNMP-Traps
    IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC
    inet  10.10.0.1        --> 10.10.0.2s
```

show interfaces detail (Encryption)

```
user@host> show interfaces es-0/3/0 detail
Physical interface: es-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 71, Generation: 21
  Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
```

```
  Input  bytes  :                    0                    0 bps
  Output bytes  :                    0                    0 bps
  Input  packets:                    0                    0 pps
  Output packets:                    0                    0 pps
  Anti-replay failures       : 0
  Authentication failures    : 0
 Egress queues: 4 supported, 4 in use
 Queue counters:       Queued packets  Transmitted packets      Dropped packets
   0 best-effort                   0                    0                    0
   1 expedited-fo                  0                    0                    0
   2 assured-forw                  0                    0                    0
   3 network-cont                  0                    0                    0

 Logical interface es-0/3/0.0 (Index 70) (SNMP ifIndex 45) (Generation 9)
   Flags: Hardware-Down Point-To-Point SNMP-Traps
   IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC
   Traffic statistics:
    Input  bytes  :                    0
    Output bytes  :                    0
    Input  packets:                    0
    Output packets:                    0
   Local statistics:
    Input  bytes  :                    0
    Output bytes  :                    0
    Input  packets:                    0
    Output packets:                    0
   Transit statistics:
    Input  bytes  :                    0                    0 bps
    Output bytes  :                    0                    0 bps
    Input  packets:                    0                    0 pps
    Output packets:                    0                    0 pps
   Protocol inet, MTU: 3800, Generation: 22, Route table: 0
     Flags: None
     Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
       Destination: 10.10.0.2, Local: 10.10.0.1, Broadcast: Unspecified,
       Generation: 26
```

show interfaces extensive (Encryption)

```
user@host> show interfaces es-0/3/0 extensive
Physical interface: es-0/3/0, Enabled, Physical link is Up
```

```
Interface index: 138, SNMP ifIndex: 71, Generation: 21
Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps
Hold-times     : Up 0 ms, Down 0 ms
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :                    0                    0 bps
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
 Anti-replay failures       : 0
 Authentication failures    : 0
Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                    0                    0                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont                   0                    0                    0

Logical interface es-0/3/0.0 (Index 70) (SNMP ifIndex 45) (Generation 9)
  Flags: Hardware-Down Point-To-Point SNMP-Traps
  IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Local statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Transit statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Protocol inet, MTU: 3800, Generation: 22, Route table: 0
    Flags: None
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
```

```
        Destination: 10.10.0.2, Local: 10.10.0.1, Broadcast: Unspecified,
        Generation: 26
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces extensive

IN THIS SECTION

- Syntax | 20676
- Description | 20676
- Options | 20678
- Required Privilege Level | 20678
- Output Fields | 20678
- Sample Output | 20678
- Sample Output | 20679
- Release Information | 20747

Syntax

```
show interfaces extensive
```

Description

Display extensive information about all interfaces configured on the router.

NOTE:

- For EX4300 switches, you can expect a temporary spike of CPU usage for *ifinfo* process. The temporary spike is not an error with the device. There can be thousands of logical interfaces (*ifls*) to process, depending on the hardware and software configuration of the virtual chassis. In order to produce the required output through the `show interfaces extensive | no-more` command, the process runs at the available CPU power. When the configuration statement is loaded in the device, the high CPU usage can last up to 10-15 seconds.

- At some times, the cumulative byte counters displayed with the `show interfaces extensive` command on the 10-Gigabit Ethernet MPC with SFP+ is not always increasing and cumulative and does not give the correct results. There is a time lag in collecting these statistics, during which the display might decrease or go from a nonzero number to zero. Eventually, the counter will display the correct result.

- When the `show interfaces extensive` command is executed on a router with an MPC or a T4000 Type 5 FPC, the `Input packet rejects` counter of the `Filter statistics` field also displays statistics related to the following packet errors:

  - Invalid VLAN range

  - Tagged packet received on an untagged interface

- When the `show interfaces extensive` command is executed on an interface that is configured on a T4000 Type 5 FPC, the `IPv6 transit statistics` field displays:

  - Total statistics (sum of transit and local statistics) at the physical interface level

  - Transit statistics at the logical interface level

- When the `show interfaces extensive` command is executed on an aggregate interface in a T1600 Core Router, the IPv6 `Input bytes` is displayed for an aggregate interface. However, the IPv6 `Input bytes` is always zero on a member link of an aggregated bundle even when there is IPv6 transit traffic on the member link. This is because the logical interface index of the aggregate logical interface is updated but not the logical interface of the member links in the channel lookup table.

- The `Output packets` field under the `Traffic statistics` section in the output of the `show interfaces extensive` command includes both IPv4 and IPv6 packets. For example, in a scenario in which both IPv4 and IPv6 packets are being mirrored on the same interface and when you deactivate an IPv4 port mirroring instance on the chassis, the output of the `show interfaces extensive` command shows a value in the Output packets

field of the Traffic statistics section, which is the value of IPv6 packets that are mirrored and not of the IPv4 packets. This behavior is expected.

- For IQ2 PIC interfaces, the output of the `show interfaces extensive` command displays byte statistics that includes L2 headers.
- If there are active OTN defects when an ISSU is performed, and the defect persists after the upgrade completes, the OTN alarm count is incremented by 1. For example, if an OTN alarm is active with a count of 1 and the defect remains after ISSU, the alarm count is incremented to 2. This behavior is expected.
- On Junos OS Evolved, the `show interfaces extensive` output displays zero for all loopback interface (lo0) statistics.

Options

This command has no options.

Required Privilege Level

view

Output Fields

For more information, see the output fields table for the particular interface type in which you are interested. For information about destination class and source class statistics, see the “Destination Class Field” section and the “Source Class Field” section under *Common Output Fields Description*. For sample output for specific interfaces, see the other topics in this collection.

Sample Output

show interfaces extensive (Circuit Emulation)

If a Circuit Emulation (CE) PIC is configured for SAToP pseudowire, then pseudowire statistics are displayed in the CE information section of the `show interface extensive` output. If SAToP pseudowire is not configured on the CE PIC, then all CE information counters display 0 (zero).

```
user@host> show interface t1-0/0/0 extensive
Physical interface :t1-0/0/0, Enabled, Physical Link : Up
        Interface index:61441
        Speed : 1.54 Mbps, Loopback: Disabled
```

```
        Operational state : Enabled,    Encapsulation : Trans
        Encoding : b8zs,        Framing    : unframe,   Build-out : 0-30
        Inversion    : enable,  Clock source : master
        Description :
        Traffic statistics:
        T1 media:       Seconds
        ES              1643
        SES             1643

        CE Info         Packets              Bytes
        CE Rx     :     2395529              306627712
        CE Tx     :     2396259              306721152
        CE Rx Drop:     0                    0
        CE Tx Drop:     0                    0

        CE Overrun  Events: 0
        CE Underrun Events: 0
```

Sample Output

show interfaces extensive (MX240, MX480, MX960, Routers with MPC7E-10G)

```
user@host> show interfaces xe-8/1/0 extensive
Physical interface: xe-8/0/4, Enabled, Physical link is Up
  Interface index: 203, SNMP ifIndex: 964, Generation: 206
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, Loop Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: None,
  Source filtering: Disabled, Flow control: Enabled, Speed Configuration: 1G
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 28:8a:1c:c9:25:34, Hardware address: 28:8a:1c:c9:25:34
  Last flapped   : 2020-03-31 09:25:32 PDT (00:00:38 ago)
  Statistics last cleared: Never
  Traffic statistics:
```

```
   Input  bytes  :                    0                    0 bps
   Output bytes  :                  672                    0 bps
   Input  packets:                    0                    0 pps
Output packets:                   16                    0 pps
```

show interfaces extensive (Fast Ethernet)

```
user@host> show interfaces fe-0/2/1 extensive
Physical interface: fe-0/2/0, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2006-04-16 23:00:41 PDT (02:08:05 ago)
  Statistics last cleared: 2006-04-16 21:42:00 PDT (03:26:46 ago)
  Traffic statistics:
   Input  bytes  :                17539                  152 bps
   Output bytes  :                92968                  224 bps
   Input  packets:                  348                    0 pps
   Output packets:                 1349                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                   66                   66                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                1283                 1283                    0
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                         24721           105982
```

```
  Total packets                          348             1349
  Unicast packets                        347              430
  Broadcast packets                        1               37
  Multicast packets                        0              882
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
Filter statistics:
  Input packet count                     348
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                     1349
  Output packet pad count                                    0
  Output packet error count                                  0
  CAM destination filters: 3, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
      Link mode: Full-duplex, Flow control: None, Remote fault: OK
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  CoS transmit queue                Bandwidth      Buffer    Priority    Limit
                            %            bps   %   usec
  0 best-effort            95       95000000  95    0            low    none
  3 network-control         5        5000000   5    0            low    none
Logical interface fe-0/2/0.0 (Index 66) (SNMP ifIndex 46) (Generation 133)
  Flags: SNMP-Traps Encapsulation: ENET2
  Protocol inet, MTU: 1500, Generation: 142, Route table: 0
    Flags: DCU, SCU-out
                                            Packets                Bytes
    Destination class            (packet-per-second)    (bits-per-second)
                      silv1_new                    0                    0
                               (                   0) (                 0)
                      silv2_new                    0                    0
                               (                   0) (                 0)
```

```
                        silv_misc                   0                    0
                                 (                  0) (                 0)
                          silver0                   0                    0
                                 (                  0) (                 0)
                          silver2                   0                    0
                                 (                  0) (                 0)
                          silver3                   0                    0
                                 (                  0) (                 0)
                          silver4                   0                    0
                                 (                  0) (                 0)
                          silver5                   0                    0
                                 (                  0) (                 0)
                          silver6                   0                    0
                                 (                  0) (                 0)
                          silver7                   0                    0
                                 (                  0) (                 0)
                          silver9                   0                    0
                                 (                  0) (                 0)
                                                Packets                Bytes
      Source class                (packet-per-second)    (bits-per-second)
                            gold1                   0                    0
                                 (                  0) (                 0)
                            gold2               16600              1062400
                                 (                  0) (                 0)
                            gold3                   0                    0
                                 (                  0) (                 0)
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.220.24/30, Local: 192.168.220.26, Broadcast: 192.168.220.27,
        Generation: 150
```

show interfaces extensive (Gigabit Ethernet)

```
user@host> show interfaces ge-5/0/0.0 extensive

  Logical interface ge-5/0/0.0 (Index 71) (SNMP ifIndex 1930) (Generation 139)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                   42
     Input  packets:                    0
     Output packets:                    1
```

```
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                  42
     Input  packets:                   0
     Output packets:                   1
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
    Output Filters: f-any
    Protocol inet, MTU: 1500, Generation: 155, Route table: 0
      Output Filters: f-inet,
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.220.24/30, Local: 192.168.220.26, Broadcast: 192.168.220.27,
        Generation: 170
    Protocol multiservice, MTU: Unlimited, Generation: 156, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (10-Gigabit Ethernet)

```
user@host> show interfaces xe-2/1/0 extensive

Physical interface: xe-2/1/0, Enabled, Physical link is Up
  Interface index: 258, SNMP ifIndex: 762, Generation: 2046
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None, Loopback:
None, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2011-12-17 00:19:02 PST (07:36:37 ago)
  Statistics last cleared: 2011-12-17 07:55:24 PST (00:00:15 ago)
  Traffic statistics:
   Input  bytes  :               110000                    0 bps
   Output bytes  :                    0                    0 bps
```

```
   Input  packets:                 1000                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :              110000
    Output bytes  :                   0
    Input  packets:                1000
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0,
    MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  MAC statistics:                      Receive         Transmit
    Total octets                        128000                0
    Total packets                         1000                0
    Unicast packets                       1000                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
```

```
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
Filter statistics:
  Input packet count                    1000
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                        0
  Output packet pad count                                    0
  Output packet error count                                  0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 2
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth               Buffer Priority   Limit
                            %            bps     %           usec
  0 best-effort            95     9500000000    95              0      low    none
  3 network-control         5      500000000     5              0      low    none
Interface transmit statistics: Disabled

Logical interface xe-2/1/0.0 (Index 83) (SNMP ifIndex 1677) (Generation 10082)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :               110000
   Output bytes  :                    0
   Input  packets:                 1000
   Output packets:                    0
   IPv6 transit statistics:
    Input  bytes  :               55000
    Output bytes  :                   0
    Input  packets:                 500
    Output packets:                   0
  Local statistics:
   Input  bytes  :                55000
   Output bytes  :                    0
   Input  packets:                  500
   Output packets:                    0
  Transit statistics:
   Input  bytes  :                55000                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                  500                    0 pps
```

```
     Output packets:                    0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :               55000
      Output bytes  :                   0
      Input  packets:                 500
      Output packets:                   0
    Protocol inet6, MTU: 1500, Generation: 23739, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2001:0db8:0a0b:12f0:0000:0000:0000/112, Local:
2001:0db8:0a0b:12f0:0000:0000:0000:0001
    Generation: 506
      Addresses, Flags: Is-Preferred
        Destination: Odb8::/64, Local: Odb8::21d:b5ff:fef8:6deb
    Protocol multiservice, MTU: Unlimited, Generation: 508
    Generation: 23740, Route table: 0
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (IQ2 and IQ2E)

```
user@host> show interfaces ge-3/2/2 extensive
    Physical interface: ge-3/2/2, Enabled, Physical link is Up
      Interface index: 156, SNMP ifIndex: 548, Generation: 159
      Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled,
      Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
      Device flags   : Present Running
      Interface flags: SNMP-Traps Internal: 0x4000
      CoS queues     : 8 supported, 8 maximum usable queues
      Schedulers     : 128
      Hold-times     : Up 0 ms, Down 0 ms
      Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
      Last flapped   : 2010-03-17 04:03:11 PDT (00:45:30 ago)
      Statistics last cleared: Never
      Traffic statistics:
       Input  bytes  :              1716096                    0 bps
       Output bytes  :              1716448                    0 bps
       Input  packets:                13407                    0 pps
       Output packets:                13411                    0 pps
       IPv6 total statistics:
        Input  bytes  :             1716096
        Output bytes  :             1716096
```

```
        Input  packets:               13407
        Output packets:               13407
      Ingress traffic statistics at Packet Forwarding Engine:
       Input  bytes  :              1716096                    0 bps
       Input  packets:                13407                    0 pps
       Drop   bytes  :                    0                    0 bps
       Drop   packets:                    0                    0 pps
      Input errors:
        Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes:
0, L2 channel errors: 1, L2 mismatch timeouts: 0, FIFO errors: 0,
        Resource errors: 0
      Output errors:
        Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO
errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
      Ingress queues: 8 supported, 4 in use
      Queue counters:       Queued packets  Transmitted packets      Dropped packets
        0 best-effort                13407                13407                    0
        1 expedited-fo                   0                    0                    0
        2 assured-forw                   0                    0                    0
        3 network-cont                   0                    0                    0
      Egress queues: 8 supported, 4 in use
      Queue counters:       Queued packets  Transmitted packets      Dropped packets
        0 best-effort                13407                13407                    0
        1 expedited-fo                   0                    0                    0
        2 assured-forw                   0                    0                    0
        3 network-cont                   4                    4                    0
      Active alarms  : None
      Active defects : None
      MAC statistics:                      Receive         Transmit
        Total octets                       1716096          1716448
        Total packets                        13407            13411
        Unicast packets                      13407            13407
        Broadcast packets                        0                0
        Multicast packets                        0                4
        CRC/Align errors                         0                0
        FIFO errors                              0                0
        MAC control frames                       0                0
        MAC pause frames                         0                0
        Oversized frames                         0
        Jabber frames                            0
        Fragment frames                          0
        VLAN tagged frames                       0
        Code violations                          0
```

```
    Filter statistics:
      Input packet count                  13407
      Input packet rejects                    0
      Input DA rejects                        0
      Input SA rejects                        0
      Output packet count                                  13411
      Output packet pad count                                  0
      Output packet error count                                0
      CAM destination filters: 0, CAM source filters: 0
    Autonegotiation information:
      Negotiation status: Complete
      Link partner:
          Link mode: Full-duplex, Flow control: None, Remote fault: OK
      Local resolution:
          Flow control: Symmetric, Remote fault: Link OK
    Packet Forwarding Engine configuration:
      Destination slot: 3
    CoS information:
      Direction : Output
      CoS transmit queue               Bandwidth               Buffer Priority   Limit
                                 %            bps     %           usec
      0 best-effort             95      950000000    95              0      low    none
      3 network-control          5       50000000     5              0      low    none
      Direction : Input
      CoS transmit queue               Bandwidth               Buffer Priority   Limit
                                 %            bps     %           usec
      0 best-effort             95      950000000    95              0      low    none
      3 network-control          5       50000000     5              0      low    none

    Logical interface ge-3/2/2.0 (Index 83) (SNMP ifIndex 6080) (Generation 148)
      Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]  Encapsulation: ENET2
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                  336
       Input  packets:                    0
       Output packets:                    4
       IPv6 total statistics:
        Input  bytes  :             1716096
        Output bytes  :             1716096
        Input  packets:               13407
        Output packets:               13407
      Local statistics:
       Input  bytes  :                    0
```

```
     Output bytes  :                  336
     Input  packets:                    0
     Output packets:                    4
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
     IPv6 total statistics:
      Input  bytes  :             1716096
      Output bytes  :             1716096
      Input  packets:               13407
      Output packets:               13407
    Protocol inet6, MTU: 1500, Generation: 159, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Primary
        Destination: Unspecified, Local: 2000::2
    Generation: 146
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::214:f600:6412:86fa
    Protocol multiservice, MTU: Unlimited, Generation: 148
    Generation: 160, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ge-3/2/2.32767 (Index 84) (SNMP ifIndex 6081) (Generation 149)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol multiservice, MTU: Unlimited, Generation: 161, Route table: 0
```

```
          Flags: None
          Policer: Input: __default_arp_policer__
```

show interfaces extensive (100-Gigabit Ethernet Type 4 PIC with CFP)

```
user@host> show interfaces et-0/0/0:0 extensive
Physical interface: et-0/0/0:0, Enabled, Physical link is Down
  Interface index: 156, SNMP ifIndex: 516, Generation: 163
  Link-level type: Ethernet, MTU: 9192, Speed: 50000mbps, BPDU Error: None,
MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping      : half-life: 5 sec, max-suppress: 20 sec, reuse 1000, suppress: 2000, state:
enabled
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2010-01-07 16:36:49 PST (18:02:35 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0,
    L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
    HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 8 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 DEFAULT, NC-                   0                    0                    0
    1 REALTIME                       0                    0                    0
```

```
  2 PRIVATE, NC-                    0                    0                    0
  3 CONTROL                      1253                 1253                    0
  4 BC-H, CLASS_                    0                    0                    0
  5 BC-M, CLASS_                    0                    0                    0
  6 IA, CLASS_V_                    0                    0                    0
  7 CLASS_S_OUTP                    0                    0                    0
Queue                          Mapped Forwarding Class
   0                                 DEFAULT, NC-Q0
  1                              REALTIME
  2                              PRIVATE, NC-Q1
  3                              CONTROL
  4                              BC-H, CLASS-Q4
  5                              BC-M, CLASS-Q5
  6                              IA, CLASS_V_OUTPUT
  7                              CLASS_S_OUTPUT
Active alarms  : None
Active defects : None
MAC statistics:                       Receive         Transmit
  Total octets                              0                0
  Total packets                             0                0
  Unicast packets                           0                0
  Broadcast packets                         0                0
  Multicast packets                         0                0
  CRC/Align errors                          0                0
  FIFO errors                               0                0
  MAC control frames                        0                0
  MAC pause frames                          0                0
  Oversized frames                          0
  Jabber frames                             0
  Fragment frames                           0
  VLAN tagged frames                        0
  Code violations                           0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth               Buffer Priority Limit
                            %            bps     %           usec
  0 best-effort            95    47500000000    95              0      low none
  3 network-control         5     2500000000     5              0      low none

Logical interface et-0/0/0:0.0 (Index 68) (SNMP ifIndex 546) (Generation 161)
  Flags: Deviet-Down SNMP-Traps Encapsulation: ENET2
```

```
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
    Protocol inet, MTU: 9178, Generation: 220, Route table: 0
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 192.168.220.24/30, Local: 192.168.220.26, Broadcast: 192.168.220.27,
Generation: 192
    Protocol mpls, MTU: 9166, Maximum labels: 3, Generation: 221, Route table: 0
    Protocol multiservice, MTU: Unlimited, Generation: 222, Route table: 0
      Policer: Input: __default_arp_policer
```

show interfaces extensive (PTX5000 Packet Transport Router)

```
user@host> show interfaces et-0/0/6 extensive
Physical interface: et-0/0/6, Enabled, Physical link is Up
  Interface index: 347, SNMP ifIndex: 531, Generation: 350
  Link-level type: Ethernet, MTU: 1514, Speed: 40Gbps, BPDU Error: None, Loop Detect PDU Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 30:b6:4f:02:29:06, Hardware address: 30:b6:4f:02:29:06
  Last flapped   : 2017-02-15 21:40:06 PST (22:55:13 ago)
  Statistics last cleared: 2017-02-16 20:33:02 PST (00:02:17 ago)
  Traffic statistics:
```

```
   Input  bytes  :              1760000                    0 bps
   Output bytes  :              1540000                    0 bps
   Input  packets:                16000                    0 pps
   Output packets:                14000                    0 pps
   IPv6 transit statistics:
    Input  bytes  :              880000
    Output bytes  :              770000
    Input  packets:                8000
    Output packets:                7000
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                            7000                 7000                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                            7000                 7000                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  MAC statistics:                      Receive         Transmit
    Total octets                       2048000          1792000
    Total packets                        16000            14000
    Unicast packets                      16000            14000
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
```

```
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                        0
  Output packet pad count                                    0
  Output packet error count                                  0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth                Buffer Priority   Limit
                            %            bps     %           usec
  0 best-effort            95    38000000000    95              0      low    none
  3 network-control         5     2000000000     5              0      low    none
Preclassifier statistics:
  Traffic Class        Received Packets   Transmitted  Packets      Dropped Packets
  best-effort                         0                      0                    0
  best-effort                         0                      0                    0
  best-effort                         0                      0                    0
  best-effort                         0                      0                    0
  best-effort                         0                      0                    0
  best-effort                         0                      0                    0
  best-effort                         0                      0                    0
  best-effort                         0                      0                    0
Link Degrade :
  Link Monitoring                   :  Disable
Interface transmit statistics: Disabled

Logical interface et-0/0/6.0 (Index 93) (SNMP ifIndex 841) (Generation 158)
  Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :              1760000
   Output bytes  :              1540000
   Input  packets:                16000
   Output packets:                14000
   IPv6 transit statistics:
    Input  bytes  :              880000
```

```
      Output bytes  :              770000
      Input  packets:                8000
      Output packets:                7000
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :              1760000                    0 bps
     Output bytes  :              1540000                    0 bps
     Input  packets:                16000                    0 pps
     Output packets:                14000                    0 pps
     IPv6 transit statistics:
      Input  bytes  :              880000
      Output bytes  :              770000
      Input  packets:                8000
      Output packets:                7000
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 206, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 11.0.0/24, Local: 11.0.0.2, Broadcast: 11.0.0.255, Generation: 228
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 207, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1100::/120, Local: 1100::2
    Generation: 230
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4fff:fe02:2906
    Protocol multiservice, MTU: Unlimited, Generation: 232
    Generation: 208, Route table: 0
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (PTX10008 Routers)

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 917, SNMP ifIndex: 817, Generation: 4436
  Link-level type: Ethernet, MTU: 1518, Speed: 20Gbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled, Minimum links
needed: 1,
  Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 30:b6:4f:e9:7c:05, Hardware address: 30:b6:4f:e9:7c:05
  Last flapped   : 2017-04-10 05:20:29 PDT (00:03:52 ago)
  Statistics last cleared: 2017-04-10 05:21:52 PDT (00:02:29 ago)
  Traffic statistics:
   Input  bytes  :          36463816334                    0 bps
   Output bytes  :          36463816334                    0 bps
   Input  packets:             24671053                    0 pps
   Output packets:             24671053                    0 pps
   IPv6 transit statistics:
    Input  bytes  :         18231905950
    Output bytes  :         18231905950
    Input  packets:            12335525
    Output packets:            12335525
  MAC statistics:                      Receive         Transmit
    Broadcast packets                        0                0
    Multicast packets                        0                0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                         24671053             24671053                    0
    1                                0                    0                    0
    2                                0                    0                    0
    3                                0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
```

```
    3                   network-control

  Logical interface ae0.0 (Index 99) (SNMP ifIndex 832) (Generation 43813)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :       4934211          0    7292763858            0
        Output:       4934211          0    7292763858            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      et-0/0/28:0.0
        Input :       4934211          0    7292763858            0
        Output:       4934211          0    7292763858            0
      et-0/0/28:3.0
        Input :             0          0             0            0
        Output:             0          0             0            0


    Aggregate member links: 2

    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/28:0.0              0           0            0            0
      et-0/0/28:3.0              0           0            0            0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 89219, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 21.0.0.0/30, Local: 21.0.0.1, Broadcast: 21.0.0.3, Generation: 62420
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 89220, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 3001::1500:0/126, Local: 3001::1500:1
    Generation: 62422
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4f00:2e9:7c05
    Protocol multiservice, MTU: Unlimited, Generation: 62424
```

```
    Generation: 89221, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae0.1 (Index 100) (SNMP ifIndex 833) (Generation 43814)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.3 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :       4934211          0    7292763858            0
        Output:       4934211          0    7292763858            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      et-0/0/28:0.1
        Input :             0          0             0            0
        Output:       4934211          0    7292763858            0
      et-0/0/28:3.1
        Input :       4934211          0    7292763858            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/28:0.1              0           0            0            0
      et-0/0/28:3.1              0           0            0            0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 89222, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 21.0.0.4/30, Local: 21.0.0.5, Broadcast: 21.0.0.7, Generation: 62426
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 89223, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 3001::1500:4/126, Local: 3001::1500:5
    Generation: 62428
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4f00:3e9:7c05
    Protocol multiservice, MTU: Unlimited, Generation: 62430
    Generation: 89224, Route table: 0
      Policer: Input: __default_arp_policer__
```

```
  Logical interface ae0.2 (Index 101) (SNMP ifIndex 834) (Generation 43815)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :       4934211          0    7292763858            0
        Output:       4934211          0    7292763858            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      et-0/0/28:0.2
        Input :       2467106          0    3646382668            0
        Output:       4934211          0    7292763858            0
      et-0/0/28:3.2
        Input :       2467105          0    3646381190            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/28:0.2              0           0            0            0
      et-0/0/28:3.2              0           0            0            0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 89225, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 21.0.0.8/30, Local: 21.0.0.9, Broadcast: 21.0.0.11, Generation: 62432
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 89226, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 3001::1500:8/126, Local: 3001::1500:9
    Generation: 62434
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4f00:4e9:7c05
    Protocol multiservice, MTU: Unlimited, Generation: 62436
    Generation: 89227, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae0.3 (Index 102) (SNMP ifIndex 835) (Generation 43816)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.5 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
```

```
    Bundle:
        Input :       4934210          0    7292762380            0
        Output:       4934210          0    7292762380            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      et-0/0/28:0.3
        Input :       4934210          0    7292762380            0
        Output:             0          0             0            0
      et-0/0/28:3.3
        Input :             0          0             0            0
        Output:       4934210          0    7292762380            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/28:0.3              0           0            0            0
      et-0/0/28:3.3              0           0            0            0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 89228, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 21.0.0.12/30, Local: 21.0.0.13, Broadcast: 21.0.0.15, Generation: 62438
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 89229, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 3001::1500:c/126, Local: 3001::1500:d
    Generation: 62440
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4f00:5e9:7c05
    Protocol multiservice, MTU: Unlimited, Generation: 62442
    Generation: 89230, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae0.4 (Index 103) (SNMP ifIndex 836) (Generation 43817)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.6 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :       4934210          0    7292762380            0
        Output:       4934210          0    7292762380            0
```

```
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      et-0/0/28:0.4
        Input :       2467105          0    3646381190            0
        Output:       2467105          0    3646381190            0
      et-0/0/28:3.4
        Input :       2467105          0    3646381190            0
        Output:       2467105          0    3646381190            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/28:0.4              0           0            0            0
      et-0/0/28:3.4              0           0            0            0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 89231, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 21.0.0.16/30, Local: 21.0.0.17, Broadcast: 21.0.0.19, Generation: 62444
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new hold cnt: 0, NH
drop cnt: 0
    Generation: 89232, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 3001::1500:10/126, Local: 3001::1500:11
    Generation: 62446
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4f00:6e9:7c05
    Protocol multiservice, MTU: Unlimited, Generation: 62448
    Generation: 89233, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae0.32767 (Index 104) (SNMP ifIndex 5645) (Generation 43818)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
```

```
        Adaptive Updates:          0
    Link:
      et-0/0/28:0.32767
        Input :             0          0             0            0
        Output:             0          0             0            0
      et-0/0/28:3.32767
        Input :             0          0             0            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/28:0.32767          0           0            0            0
      et-0/0/28:3.32767          0           0            0            0
    Protocol multiservice, MTU: Unlimited, Generation: 89234, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces extensive for 10-Gigabit Ethernet on JNP10K-LC1201 line card (PTX10008)

```
user@host> show interfaces et-4/0/20 extensive
Physical interface: et-4/0/20, Enabled, Physical link is Up
  Interface index: 1090, SNMP ifIndex: 537, Generation: 579820586853
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None, Loop
Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled
  Wavelength     : 1550.10 nm, Frequency: 193.40 THz - Output
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: e4:5d:37:38:cc:d8, Hardware address: e4:5d:37:38:cc:d8
  Last flapped   : 2020-10-20 05:42:54 UTC (00:00:29 ago)
  Statistics last cleared: Never
Cont.....................................
```

show interfaces extensive (PTX1000 Routers)

```
user@host> show interfaces et-0/0/48:1 extensive
Physical interface: et-0/0/48:1, Enabled, Physical link is Up
  Interface index: 306, SNMP ifIndex: 697, Generation: 311
```

```
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None, Loop
Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: None, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: ec:13:db:62:4a:f6, Hardware address: ec:13:db:62:4a:f6
  Last flapped   : 2017-05-08 11:07:59 PDT (12:08:13 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                               0                    0                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
```

```
  Bit errors                              3
  Errored blocks                          3
MAC statistics:                      Receive         Transmit
  Total octets                             0                0
  Total packets                            0                0
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth               Buffer Priority   Limit
                            %            bps     %           usec
  0 best-effort            95     9500000000    95              0      low    none
  3 network-control         5      500000000     5              0      low    none
Link Degrade :
  Link Monitoring                    :  Disable
Interface transmit statistics: Disabled
```

show interfaces extensive (MX Series Routers)

```
user@host> show interfaces xe-0/0/0 extensive
Physical interface: xe-0/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 592, Generation: 148
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2013-10-26 03:20:40 test (2w3d 03:15 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                    0                    0
    1                                0                    0                    0
```

```
  2                                0                    0                    0
  3                                0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : LINK
Active defects : LINK
PCS statistics                      Seconds
  Bit errors                           109
  Errored blocks                       109
MAC statistics:                      Receive         Transmit
  1                                0                    0                    0
  2                                0                    0                    0
  3                                0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : LINK
Active defects : LINK
PCS statistics                      Seconds
  Bit errors                           109
  Errored blocks                       109
MAC statistics:                      Receive         Transmit
  Total octets                           0                0
  Total packets                          0                0
  Unicast packets                        0                0
  Broadcast packets                      0                0
  Multicast packets                      0                0
  CRC/Align errors                       0                0
  FIFO errors                            0                0
  MAC control frames                     0                0
  MAC pause frames                       0                0
  Oversized frames                       0
  Jabber frames                          0
  Fragment frames                        0
  VLAN tagged frames                     0
  Code violations                        0
  Total errors                           0                0
Filter statistics:
```

```
    Input packet count                      0
    Input packet rejects                    0
    Input DA rejects                        0
    Input SA rejects                        0
    Output packet count                                     0
    Output packet pad count                                 0
    Output packet error count                               0
    CAM destination filters: 0, CAM source filters: 0
  Packet Forwarding Engine configuration:
    Destination slot: 0
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 best-effort            95     9500000000    95              0      low    none
    3 network-control         5      500000000     5              0      low    none
  Interface transmit statistics: Disabled
```

When an ASIC is wedged, the interfaces are brought down along with the physical interface. The reason for the link down is displayed as `ASIC-Error` in the `Device flags`.

```
user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Administratively down, Physical link is Down
  Interface index: 147, SNMP ifIndex: 563, Generation: 150
  Link-level type: Ethernet, MTU: 1514, MRU: 0, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None,
Loop Detect PDU Error: None,
  MAC-REWRITE Error: None, Loopback: None, Source filtering: Disabled, Flow control: Disabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running Down ASIC-Error
  Interface flags: Hardware-Down Down Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: cc:e1:7f:a8:05:4a, Hardware address: cc:e1:7f:a8:05:4a
  Last flapped   : 2017-06-05 17:20:54 PDT (00:03:51 ago)
  Statistics last cleared: Never
```

show interfaces extensive (MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface)

```
user@host> show interfaces xe-0/0/3 extensive
Physical interface: xe-0/0/3, Enabled, Physical link is Up
  Interface index: 200, SNMP ifIndex: 577, Generation: 203
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, MAC-REWRITE Error: None, Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2014-06-26 18:16:50 PDT (04:58:35 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                 0                    0                    0
```

```
  1                                 0                    0                    0
  2                                 0                    0                    0
  3                                 0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : None
Active defects : None
PCS statistics                      Seconds
  Bit errors                             0
  Errored blocks                         4
MAC statistics:                      Receive         Transmit
  Total octets                             0                0
  Total packets                            0                0
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
  Total errors                             0                0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth               Buffer Priority   Limit
```

```
                              %            bps     %           usec
    0 best-effort            95     9500000000    95              0      low    none
    3 network-control         5      500000000     5              0      low    none
  Interface transmit statistics: Disabled
```

show interfaces extensive (MX480 Router with MPC5E and 100-Gigabit Ethernet OTN Interface)

```
user@host> show interfaces et-2/1/0 extensive
Physical interface: et-2/1/0, Enabled, Physical link is Up
  Interface index: 215, SNMP ifIndex: 872, Generation: 218
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, BPDU Error: None, Loopback:
Disabled, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2014-06-26 18:42:04 PDT (04:36:58 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
```

```
  Output errors:
    Carrier transitions: 263, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors:
0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                               0                    0                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                       754
  MAC statistics:                      Receive         Transmit
    Total octets                         14960                0
    Total packets                          104                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
 MAC control frames                        0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            2
    Fragment frames                          6
    VLAN tagged frames                       0
    Code violations                          0
    Total errors                            98                0
  Filter statistics:
    Input packet count                     104
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                      0
    Output packet pad count                                  0
    Output packet error count                                0
```

```
    CAM destination filters: 0, CAM source filters: 0
  OTN alarms        : None
  OTN defects       : None
  OTN FEC Mode      : GFEC
  OTN Rate          :     OTU4 100Gbps
  OTN Line Loopback : None
  OTN Local Loopback: None
  OTN Payload PRBS   : None
  OTN FEC statistics:
    Corrected Errors                            169828399453
    Uncorrected Words                            28939961456
    Corrected Error Ratio (      17963 sec average)  8.46e-05
  OTN FEC alarms:          Seconds        Count  State
    FEC Degrade               1180            3  OK
    FEC Excessive             1160            5  OK
  OTN OC:                  Seconds        Count  State
    LOS                        129            1  OK
    LOF                          2            1  OK
    LOM                          0            0  OK
    Wavelength Lock              0            0  OK
  OTN OTU:
    AIS                          0            0  OK
    BDI                          7            1  OK
    IAE                          0            0  OK
    TTIM                       168           45  OK
    BIAE                         0            0  OK
    TSF                          0            0  OK
    SSF                          0            0  OK
    Received DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Received SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  OTN ODU:
    AIS                        130            1  OK
    OCI                          0            0  OK
    LCK                          0            0  OK
BDI                            7            1  OK
    TTIM                       133            1  OK
    IAE                          0            0  OK
```

```
    LTC                          0            0  OK
    CSF                          8            4  OK
    TSF                          0            0  OK
    SSF                          0            0  OK
    PTIM                       130            1  OK
    Received DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Received SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  OTN Received Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x00
  ODU Delay Management :
    Result  : 0x00
  PRBS:
    Result: Test not enabled
  OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x00
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 best-effort            95    95000000000    95              0      low    none
    3 network-control         5     5000000000     5              0      low    none
  Interface transmit statistics: Disabled
```

show interfaces extensive ((MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show interfaces et-3/0/0 extensive
Physical interface: et-3/0/0, Enabled, Physical link is Up
  Interface index: 163, SNMP ifIndex: 564, Generation: 166
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, BPDU Error: None, Loopback:
Disabled, Source filtering:
Disabled,
```

```
  Flow control: Enabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  Wavelength     : 1550.12 nm, Frequency: 193.40 THz
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2016-02-17 14:26:31 PST (09:04:28 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU
errors: 0,
    Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                    0                    0
    1                                0                    0                    0
    2                                0                    0                    0
```

```
  3                                0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : None
Active defects : None
PCS statistics                      Seconds
  Bit errors                             8
  Errored blocks                        10
MAC statistics:                      Receive         Transmit
  Total octets                             0                0
  Total packets                            0                0
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
  Total errors                             0                0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
OTN alarms        : None
OTN defects       : None
OTN FEC Mode      : SDFEC
OTN Rate          :     OTU4 (120.5Gbps)
OTN Line Loopback : None
OTN Local Loopback: None
OTN Payload PRBS   : None
```

```
  OTN Laser Enable   : On
  OTN FEC statistics:
    Corrected Errors                              7065332638
    Uncorrected Words                                3412572
    Corrected Error Ratio (      32785 sec average)  1.79e-06 (INVALID)
 OTN FEC alarms:          Seconds        Count  State
    FEC Degrade                 0            0  OK
    FEC Excessive               3            1  OK
  OTN OC:                  Seconds        Count  State
    LOS                         3            1  OK
    LOF                        50            3  OK
    LOM                         3            3  OK
    Wavelength Lock             0            0  OK
  OTN OTU:
    AIS                         0            0  OK
    BDI                         4            4  OK
    IAE                         4            4  OK
    TTIM                        1            1  OK
    BIAE                        3            3  OK
    TSF                        50            3  OK
    SSF                        50            3  OK
    Received DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Received SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  OTN ODU:
    AIS                        20            2  OK
    OCI                         4            4  OK
    LCK                         4            4  OK
    BDI                         2            2  OK
    TTIM                       20            2  OK
    IAE                         0            0  OK
    LTC                         0            0  OK
    CSF                        18            2  OK
    TSF                        66            2  OK
    SSF                        66            2  OK
    PTIM                       43            2  OK
    Received DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
```

```
    Received SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  OTN Received Overhead Bytes:
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x07
  ODU Delay Management :
    Result  : 0ms
  PRBS:
    Result: Test not enabled
  OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x07
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 best-effort            95    95000000000    95              0      low    none
    3 network-control         5     5000000000     5              0      low    none
  Interface transmit statistics: Disabled
```

show interfaces extensive (PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host > show interfaces extensive et-4/0/0
Physical interface: et-4/0/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 544, Generation: 161
  Link-level type: Ethernet, MTU: 1514, Speed: 100Gbps, BPDU Error: None, Loopback: Disabled,
Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  Wavelength     : 1550.12 nm, Frequency: 193.40 THz
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
```

```
unsuppressed
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2016-06-04 21:42:42 PDT (1d 05:09 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                    0                    0
    1                                0                    0                    0
    2                                0                    0                    0
    3                                0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             7
 Errored blocks                         10
  MAC statistics:                       Receive         Transmit
    Total octets                              0                0
    Total packets                             0                0
    Unicast packets                           0                0
    Broadcast packets                         0                0
```

```
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                       0
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 0, CAM source filters: 0
  OTN alarms        : None
  OTN defects       : None
  OTN FEC Mode      : SDFEC
  OTN Rate          :     OTU4 (120.5Gbps)
  OTN Line Loopback : None
  OTN Local Loopback: None
  OTN Payload PRBS   : None
  OTN Laser Enable   : On
  OTN FEC statistics:
    Corrected Errors                                 19637746
    Uncorrected Words                                       0
    Corrected Error Ratio (     104923 sec average)  1.55e-09
  OTN FEC alarms:          Seconds        Count  State
    FEC Degrade                  0            0  OK
    FEC Excessive                0            0  OK
  OTN OC:                  Seconds        Count  State
    LOS                          0            0  OK
    LOF                          2            1  OK
    LOM                          2            1  OK
    Wavelength Lock              0            0  OK
 OTN OTU:
    AIS                          0            0  OK
    BDI                          2            1  OK
    IAE                          0            0  OK
```

```
    TTIM                          0            0  OK
    BIAE                          0            0  OK
    TSF                           2            1  OK
    SSF                           0            0  OK
    Received DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Received SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  OTN ODU:
    AIS                           0            0  OK
    OCI                           0            0  OK
    LCK                           0            0  OK
    BDI                           2            1  OK
    TTIM                          0            0  OK
    IAE                           0            0  OK
    LTC                           0            0  OK
    CSF                           0            0  OK
    TSF                           2            1  OK
    SSF                           0            0  OK
    PTIM                          2            1  OK
    Received DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Received SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  OTN Received Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x07
  ODU Delay Management :
    Result  : 0ms
  PRBS:
    Result: Test not enabled
  OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
 Payload Type: 0x07
  Packet Forwarding Engine configuration:
```

```
    Destination slot: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 best-effort            95    95000000000    95              0      low    none
    3 network-control         5     5000000000     5              0      low    none
  Interface transmit statistics: Disabled
```

show interfaces extensive (MX2020 Router with MPC6E and OTN MIC)

```
user@host> show interfaces xe-3/0/0 extensive
Physical interface: xe-3/0/0, Enabled, Physical link is Up
  Interface index: 166, SNMP ifIndex: 516, Generation: 169
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, MAC-REWRITE Error: None, Loopback: None, Source filtering: Disabled, Flow control:
Enabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2014-05-28 17:53:12 PDT (05:56:24 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
```

```
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes:
0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO
errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-forwarding                   0                    0                    0
    2 assured-forwarding                   0                    0                    0
    3 network-control                   0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             2
    Errored blocks                         2
  MAC statistics:                      Receive         Transmit
    Total octets                           0                0
    Total packets                          0                0
    Unicast packets                        0                0
    Broadcast packets                      0                0
    Multicast packets                      0                0
    CRC/Align errors                       0                0
 FIFO errors                             0                0
    MAC control frames                     0                0
    MAC pause frames                       0                0
    Oversized frames                       0
    Jabber frames                          0
    Fragment frames                        0
    VLAN tagged frames                     0
    Code violations                        0
    Total errors                           0                0
  Filter statistics:
    Input packet count                     0
    Input packet rejects                   0
    Input DA rejects                       0
    Input SA rejects                       0
```

```
    Output packet count                                        0
    Output packet pad count                                    0
    Output packet error count                                  0
    CAM destination filters: 0, CAM source filters: 0
  OTN alarms        : None
  OTN defects       : None
  OTN FEC Mode      : GFEC
  OTN Rate          : Fixed Stuff Bytes 11.0957Gbps
  OTN Line Loopback : None
  OTN Local Loopback: None
  OTN Payload PRBS   : None
  OTN FEC statistics:
    Corrected Errors                                        0
    Uncorrected Words                                       0
    Corrected Error Ratio (      21387 sec average)  0.00e+00
  OTN FEC alarms:          Seconds        Count  State
    FEC Degrade                  0            0  OK
    FEC Excessive                0            0  OK
  OTN OC:                  Seconds        Count  State
    LOS                          0            0  OK
    LOF                          0            0  OK
    LOM                          0            0  OK
    Wavelength Lock              0            0  OK
  OTN OTU:
    AIS                          0            0  OK
    BDI                          0            0  OK
    IAE                          0            0  OK
    TTIM                         0            0  OK
    BIAE                         0            0  OK
    TSF                          0            0  OK
    SSF                          0            0  OK
    Received DAPI:
    00 53 4d 2d 54 52 43 20 44 41 50 49 2d 53 45 43   .SM-TRC DAPI-SEC
    Received SAPI:
    00 53 4d 2d 54 52 43 20 53 41 50 49 2d 53 45 43   .SM-TRC SAPI-SEC
    Transmitted DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  OTN ODU:
    AIS                          0            0  OK
    OCI                          0            0  OK
 LCK                           0            0  OK
```

```
  BDI                          0            0  OK
  TTIM                         0            0  OK
  IAE                          0            0  OK
  LTC                          0            0  OK
  CSF                          0            0  OK
  TSF                          0            0  OK
  SSF                          0            0  OK
  PTIM                         0            0  OK
  Received DAPI:
  00 50 4d 2d 54 52 43 20 44 41 50 49 2d 53 45 43   .PM-TRC DAPI-SEC
  Received SAPI:
  00 50 4d 2d 54 52 43 20 53 41 50 49 2d 53 45 43   .PM-TRC SAPI-SEC
  Transmitted DAPI:
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  Transmitted SAPI:
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
OTN Received Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x00
ODU Delay Management :
  Result  : 0x00
PRBS:
  Result: Test not enabled
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x00
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth               Buffer Priority   Limit
                            %            bps     %           usec
  0 best-effort            95     9500000000    95              0      low    none
  3 network-control         5      500000000     5              0      low    none
Interface transmit statistics: Disabled
```

show interfaces extensive (MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface)

```
user@host> show interfaces et-9/0/0 extensive
Physical interface: et-9/0/0, Enabled, Physical link is Up
  Interface index: 196, SNMP ifIndex: 623, Generation: 199
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, BPDU Error: None, Loopback:
Disabled, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2014-06-26 18:18:34 PDT (04:17:07 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                 0                    0                    0
    1                                 0                    0                    0
```

```
  2                                0                    0                    0
  3                                0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : None
Active defects : None
PCS statistics                      Seconds
  Bit errors                             0
  Errored blocks                         0
MAC statistics:                      Receive         Transmit
  Total octets                             0                0
  Total packets                            0                0
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
  Total errors                             0                0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth                Buffer Priority   Limit
                            %            bps     %            usec
```

```
    0 best-effort            95    95000000000    95              0      low    none
    3 network-control         5     5000000000     5              0      low    none
  Interface transmit statistics: Disabled
```

show interfaces extensive (MX2010 Router with MPC6E and 10-Gigabit Ethernet Interface)

```
user@host> show interfaces xe-6/1/0 extensive
Physical interface: xe-6/1/0, Enabled, Physical link is Up
  Interface index: 159, SNMP ifIndex: 603, Generation: 162
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, MAC-REWRITE Error: None, Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2014-06-26 18:16:50 PDT (04:21:04 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
```

```
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                               0                    0                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         1
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
 MAC control frames                        0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
    Total errors                             0                0
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                       0
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 0, CAM source filters: 0
```

```
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 best-effort            95     9500000000    95              0      low    none
    3 network-control         5      500000000     5              0      low    none
  Interface transmit statistics: Disabled
```

show interfaces extensive (T4000 Routers with Type 5 FPCs)

The output fields for the `show interfaces` *`interface`* `extensive` command remain the same for 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP), 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP), and 100-Gigabit Ethernet Type 5 PIC with CFP (PF-1CGE-CFP).

```
user@host> show interfaces xe-4/0/0 extensive
Physical interface: xe-4/0/0, Enabled, Physical link is Up
  Interface index: 200, SNMP ifIndex: 592, Generation: 203
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping      : half-life: 5 sec, max-suppress: 20 sec, reuse 1000, suppress: 2000, state:
enabled
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2013-06-03 16:01:56 PDT (06:04:07 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
```

```
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                       0
```

```
    Output packet pad count                                    0
    Output packet error count                                  0
    CAM destination filters: 0, CAM source filters: 0
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue             Bandwidth             Buffer Priority  Limit
                            %          bps     %         usec
    0 best-effort          95   9500000000    95            0      low    none
    3 network-control       5    500000000     5            0      low    none
  Preclassifier statistics:
   Traffic Class       Received Packets   Transmitted Packets      Dropped Packets
   real-time                         0                     0                    0
   network-control                   0                     0                    0
   best-effort                       0                     0                    0
  Interface transmit statistics: Disabled
```

show interfaces extensive (Aggregated Ethernet)

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 199, SNMP ifIndex: 570, Generation: 202
  Link-level type: Ethernet, MTU: 1514, Speed: 2Gbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2012-06-06 23:33:03 PDT (00:00:58 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                18532                 1984 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                  158                    2 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
```

```
Dropped traffic statistics due to STP State:
 Input  bytes  :                    0
 Output bytes  :                    0
 Input  packets:                    0
 Output packets:                    0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0,
  Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                    0                    0                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont                   0                    0                    0
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                   57                   57                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont               63605                63605                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control

Logical interface ae0.0 (Index 331) (SNMP ifIndex 583) (Generation 142)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  Statistics        Packets        pps         Bytes          bps
  Bundle:
      Input :           149          2         17416         1984
      Output:             0          0             0            0
  Link:
    ge-3/2/5.0
      Input :            90          1         10100          992
      Output:             0          0             0            0
    ge-3/3/9.0
      Input :            59          1          7316          992
      Output:             0          0             0            0
  LACP info:        Role             System             System         Port        Port  Port
                           priority          identifier  priority        number      key
```

```
      ge-3/2/5.0     Actor           100   00:00:00:00:00:01        127       1     1
      ge-3/2/5.0   Partner        127  00:24:dc:98:67:c0       127       1     1
      ge-3/3/9.0     Actor          100  00:00:00:00:00:01       127       2     1
      ge-3/3/9.0   Partner        127  00:24:dc:98:67:c0       127       2     1
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-3/2/5.0                38         137            0            0
      ge-3/3/9.0                36         139            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-3/2/5.0                 0           0            0            0
      ge-3/3/9.0                 0           0            0            0
    Protocol inet, MTU: 1500, Generation: 169, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.2, Broadcast: 1.1.1.255, Generation: 153
    Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (Aggregated Ethernet Interfaces on Junos OS Evolved)

On Junos OS Evolved Release 23.1R1 and later, when you disable a member link of an aggregated Ethernet Interface, ae0, the `show interfaces terse ae0` command and the `show interfaces extensive ae0` command display the status of the member link as down. In previous releases, the status of the member link was not displayed as down when you use the `show interfaces extensive ae0` command.

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 1005, SNMP ifIndex: 760, Generation: 837518623731
  Link-level type: Ethernet, MTU: 1514, Speed: 20Gbps, BPDU Error: None, Ethernet-Switching
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
control: Enabled,
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Current address: ec:7c:5c:cd:73:69, Hardware address: ec:7c:5c:cd:73:69
  Last flapped   : 2023-03-15 20:52:13 PDT (00:07:54 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               225545                 1880 bps
   Output bytes  :               247303                 2800 bps
   Input  packets:                 1869                    1 pps
```

```
 Output packets:                 2077                    2 pps
Members:
    Statistics        Packets       pps          Bytes         bps
    et-3/0/10:2
      Input :           620          1         73387         1184
      Output:          1068          1        119831         1400
    et-3/0/10:3
      Input :           511          0         66089          696
      Output:           482          1         61696         1400
    et-3/1/14:2 <-- down
      Input :           738          0         86069            0
      Output:           527          0         65776            0
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0                               0                    0                    0
  1                               0                    0                    0
  2                               0                    0                    0
  3                            2073                 2073                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control

Logical interface ae0.0 (Index 1003) (SNMP ifIndex 761) (Generation 837518623738)
  Flags: Up SNMP-Traps Encapsulation: ENET2 DF
  Statistics        Packets        pps         Bytes          bps
  Bundle:
      Input :          1865          4        225137         4176
      Output:           627          0         59242            0

  Aggregate member links: 3

  LACP info:        Role     System             System       Port     Port    Port
                           priority         identifier   priority   number     key
    et-3/0/10:2.0    Actor      127  ec:7c:5c:cd:73:69        127        3       1
    et-3/0/10:2.0  Partner      127  00:24:dc:bd:17:c0        127        7       1
    et-3/0/10:3.0    Actor      127  ec:7c:5c:cd:73:69        127        4       1
    et-3/0/10:3.0  Partner      127  00:24:dc:bd:17:c0        127        1       1
    et-3/1/14:2.0    Actor      127  ec:7c:5c:cd:73:69        127        5       1
    et-3/1/14:2.0  Partner        1  00:00:00:00:00:00          1        5       1
  LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
    et-3/0/10:2.0            474         478            0            0
```

```
      et-3/0/10:3.0            473         477            0            0
      et-3/1/14:2.0            466         484            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-3/0/10:2.0              0           0            0            0
      et-3/0/10:3.0              0           0            0            0
      et-3/1/14:2.0              0           0            0            0
    Protocol inet
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
    MTU: 1500, Generation: 837518623740, Route table: 0
      Flags: Sendbcast-pkt-to-re, Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1.0.0.0/30, Local: 1.0.0.1, Broadcast: 1.0.0.3, Generation: 837518623751
    Protocol iso, MTU: 1497, Generation: 837518623741, Route table: 0
      Flags: Is-Primary
    Protocol inet6
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
    MTU: 1500, Generation: 837518623742, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::ee7c:5cff:fecd:7369
    Protocol mpls, MTU: 1476, Maximum labels: 6, Generation: 837518623774
    Generation: 837518623743, Route table: 0
      Flags: Is-Primary
    Protocol multiservice, MTU: Unlimited, Generation: 837518623750, Route table: 0
      Flags: Is-Primary
```

show interfaces extensive

Starting in Junos OS Release 21.1R1, the 10 Gbps, 1 Gbps, or 100 Mbps port switches to low power idle (LPI) mode based on the following conditions:

- When a port operates at 1 Gbps or 10 Gbps speed and no traffic is either received or transmitted, then the port enters LPI mode. If the 1Gbps port transfers unidirectional or bidirectional traffic, then the port does not enter LPI mode.
- When a port operates at 100 Mbps speed, the port switches to LPI mode based on the direction of the traffic. The `show interfaces interface-name extensive` command displays `RX LPI` when no RX traffic and `TX LPI` when no TX traffic.
- When you enable EEE and the port has not entered LPI, the `show interfaces interface-name extensive` command displays `No LPI`.

Use Feature Explorer to confirm platform and release support for time domain reflectometry (TDR).

You can view the interface that is in low power idle mode, by executing the `show interfaces interface-name extensive` command. The output field, `IEEE 802.3az Energy Efficient Ethernet`, displays the status of the LPI mode.

```
user@host> show interfaces ge-0/0/0 extensive
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 648, SNMP ifIndex: 514
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Link-mode: Full-duplex, Speed: Auto, BPDU
Error: None, Loop Detect PDU Error: None, Ethernet-Switching Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control:
Disabled, Auto-negotiation: Enabled, Remote fault: Online, Media type: Copper,
  IEEE 802.3az Energy Efficient Ethernet: Enabled, TX LPI, Auto-MDIX: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 40:de:ad:f8:73:0e, Hardware address: 40:de:ad:f8:73:0e
  Last flapped   : 2020-10-08 06:18:24 PDT (02:37:26 ago)
  Input rate     : 217203784 bps (212117 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Ethernet FEC statistics              Errors
    FEC Corrected Errors                    0
    FEC Uncorrected Errors                  0
    FEC Corrected Errors Rate               0
    FEC Uncorrected Errors Rate             0
  PRBS Statistics : Disabled
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/0.0 (Index 555) (SNMP ifIndex 667)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 17236976
    Output packets: 175620525
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
```

```
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2.2.2/24, Local: 2.2.2.2, Broadcast: 2.2.2.255
```

show interfaces extensive (ACX7509)

```
user@host>show interfaces et-0/0/0 extensive
Physical interface: et-0/0/0, Enabled, Physical link is Down
  Interface index: 1035, SNMP ifIndex: 505, Generation: 588410520745
  Link-level type: Ethernet-Bridge, MTU: 9996, LAN-PHY mode, Speed: 100Gbps,
  BPDU Error: None, Loop Detect PDU Error: None, Ethernet-Switching Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Disabled, Media type: Fiber
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 74:e7:98:60:58:00, Hardware address: 74:e7:98:60:58:00
  Last flapped   : 2023-02-11 05:36:31 PST (01:22:17 ago)
```

show interfaces extensive (QFX5130-32CD and QFX5220-32CD)

```
user@host> show interfaces extensive et-0/0/30:0
Physical interface: et-0/0/30:0, Enabled, Physical link is Up
  Interface index: 1043, SNMP ifIndex: 601, Generation: 618475292371
  Link-level type: Ethernet, MTU: 9216, LAN-PHY mode, Speed: 25Gbps,
  BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled,
  Media type: Fiber
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 0c:59:9c:82:7d:16, Hardware address: 0c:59:9c:82:7d:16
  Last flapped   : 2021-04-12 01:46:47 PDT (00:06:13 ago)
```

```
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :         574772409296                    0 bps
 Output bytes  :         547697053354                    0 bps
 Input  packets:            559393055                    0 pps
 Output packets:            533041598                    0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 5 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0                        533041545            533041545                    0
  3                                0                    0                    0
  4                                0                    0                    0
  7                                0                    0                    0
  8                               53                   53                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  3                   fcoe
  4                   no-loss
  7                   network-control
  8                   mcast
Active alarms  : None
Active defects : None
PCS statistics                      Seconds
  Bit errors                             0
  Errored blocks                         0
Ethernet FEC Mode  :                  FEC91
Ethernet FEC statistics              Errors
  FEC Corrected Errors                    0
  FEC Uncorrected Errors                  6
  FEC Corrected Errors Rate               0
  FEC Uncorrected Errors Rate             0
MAC statistics:                      Receive         Transmit
  Total octets                  574772409296     547697053354
  Total packets                    559393055        533041598
  Unicast packets                  559393055        533041598
  Broadcast packets                        0                0
  Multicast packets                        0                0
```

```
  CRC/Align errors                          0                0
  FIFO errors                               0                0
  MAC control frames                        0                0
  MAC pause frames                          0                0
  Oversized frames                          0
  Jabber frames                             0
  Fragment frames                           0
  VLAN tagged frames                        0
  Code violations                           0
  Total errors                              0                0
MAC Priority Flow Control Statistics:
  Priority :  0                              0                0
  Priority :  1                              0                0
  Priority :  2                              0                0
  Priority :  3                              0                0
  Priority :  4                              0                0
  Priority :  5                              0                0
  Priority :  6                              0                0
  Priority :  7                              0                0
Filter statistics:
  Input packet count                        0
  Input packet rejects                      0
  Input DA rejects                          0
  Input SA rejects                          0
  Output packet count                                        0
  Output packet pad count                                    0
  Output packet error count                                  0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot:  ( )
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth               Buffer Priority   Limit
                            %            bps     %           usec
  0 best-effort             5     1250000000     5              0      low    none
  3 fcoe                   35     8750000000    35              0      low    none
  4 no-loss                35     8750000000    35              0      low    none
  7 network-control         5     1250000000     5              0      low    none
  8 mcast                  20     5000000000    20              0      low    none
Interface transmit statistics: Disabled
Link Degrade :
  Link Monitoring                   :  Disable
```

```
  Logical interface et-0/0/30:0.0 (Index 1017) (SNMP ifIndex 594)
   (Generation 618475292376)
    Flags: Up SNMP-Traps Encapsulation: Ethernet-Bridge
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol ethernet-switching, MTU: 9216, Generation: 618475292381,
    Route table: 51, Mesh Group: __all_ces__, Next-hop: 51023
      Flags: None
```

show interfaces extensive (QFX5210-64C)

```
show interfaces et-0/0/9 extensivePhysical interface: et-0/0/16, Enabled, Physical link is Up
  Interface index: 676, SNMP ifIndex: 657
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 100Gbps, BPDU Error: None, Loop
Detect PDU Error: None,
  Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled, Flow control: Disabled,
  Media type: Fiber
  Device flags   : Present Running
  Interface Specific flags: Internal: 0x100000
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 10 supported, 10 maximum usable queues
  Current address: 64:c3:d6:60:49:b4, Hardware address: 64:c3:d6:60:49:b4
  Last flapped   : 2023-06-22 11:37:44 PDT (00:18:03 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
```

```
  Active defects : None
  Pre-emphasis(reg 236/237)
    Register-value                    0
  Ctle(reg 234/235)
    Register-value                    0
  Differential-amplitude(reg 238/239)
    Register-value                    0
  PCS statistics                     Seconds
    Bit errors                            0
    Errored blocks                        0
  Ethernet FEC Mode  :                 FEC91
    FEC Codeword size                    528
    FEC Codeword rate                  0.973
  Ethernet FEC statistics             Errors
    FEC Corrected Errors                   0
    FEC Uncorrected Errors                 0
    FEC Corrected Errors Rate              0
    FEC Uncorrected Errors Rate            0
  PRBS Mode : Disabled
  Interface transmit statistics: Disabled

  Logical interface et-0/0/16.16386 (Index 839) (SNMP ifIndex 658)
    Flags: Up 0x4004000 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
```

show interfaces extensive (QFX5120-48Y)

```
show interfaces ge-0/0/44 extensive
Physical interface: ge-0/0/44, Enabled, Physical link is Down
  Interface index: 671, SNMP ifIndex: 581, Generation: 162
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Link-mode: Unknown, Speed: 1000mbps,
Source filtering: Disabled
  Link-mode: Half-duplex, BPDU Error: None, Loop Detect PDU Error: None, Ethernet-Switching
Error: None, Remote Bounce: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Flow control: Disabled, Auto-negotiation: Enabled, Remote fault: Online,
Media type: Copper, IEEE 802.3az Energy Efficient Ethernet: Disabled,
  Auto-MDIX: Enabled
  Device flags   : Present Running Down
  Interface Specific flags: Internal: 0x100000
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
```

```
  Link flags     : None
  CoS queues     : 10 supported, 10 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 54:4b:8c:19:e7:b0, Hardware address: 54:4b:8c:19:e7:b0
  Last flapped   : 2024-11-27 20:45:40 PST (00:00:00 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 10 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                    0                    0
    3                                0                    0                    0
    4                                0                    0                    0
    7                                0                    0                    0
    8                                0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    3                   fcoe
    4                   no-loss
    7                   network-control
    8                   mcast
  Active alarms  : LINK
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
```

```
  Total packets                          0                0
  Unicast packets                        0                0
  Broadcast packets                      0                0
  Multicast packets                      0                0
  CRC/Align errors                       0                0
  FIFO errors                            0                0
  MAC control frames                     0                0
  MAC pause frames                       0                0
  Oversized frames                       0
  Jabber frames                          0
  Fragment frames                        0
  VLAN tagged frames                     0
  Code violations                        0
MAC Priority Flow Control Statistics:
  Priority :  0                           0                0
  Priority :  1                           0                0
  Priority :  2                           0                0
  Priority :  3                           0                0
  Priority :  4                           0                0
  Priority :  5                           0                0
  Priority :  6                           0                0
  Priority :  7                           0                0
PRBS Mode : Disabled
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth               Buffer Priority   Limit
                            %            bps     %           usec
  0 best-effort             5       50000000     5              0      low    none
  3 fcoe                   35      350000000    35              0      low    none
  4 no-loss                35      350000000    35              0      low    none
  7 network-control         5       50000000     5              0      low    none
  8 mcast                  20      200000000    20              0      low    none
Interface transmit statistics: Disabled
MACSec statistics:
  Output
      Secure Channel Transmitted
      Protected Packets               : 0
      Encrypted Packets               : 0
      Protected Bytes                 : 0
```

```
        Encrypted Bytes                 : 0
     Input
        Secure Channel Received
        Accepted Packets                : 0
        Validated Bytes                 : 0
        Decrypted Bytes                 : 0

  Logical interface ge-0/0/44.16386 (Index 855) (SNMP ifIndex 0) (Generation 184)
    Flags: Device-Down 0x4004000 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
   Local statistics:
    Input  bytes  :                    0
    Output bytes  :                    0
    Input  packets:                    0
    Output packets:                    0
   Transit statistics:
    Input  bytes  :                    0                    0 bps
    Output bytes  :                    0                    0 bps
    Input  packets:                    0                    0 pps
    Output packets:                    0                    0 pps

{master:0}
root@bng-sw-onyx-p1a>
```

show interfaces extensive (IPv4 and IPv6 traffic statistics)

```
show interfaces ge-1/2/9.0 extensive
Logical interface ge-1/2/9.0 (Index 333) (SNMP ifIndex 585) (Generation 169)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
     IPv6 transit statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
```

```
 Input  packets:                    0
 Output packets:                    0
Local statistics:
 Input  bytes  :                    0
 Output bytes  :                    0
 Input  packets:                    0
 Output packets:                    0
Transit statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :                    0                    0 bps
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
 IPv6 transit statistics:
  Input  bytes  :                   0                    0 bps
  Output bytes  :                   0                    0 bps
  Input  packets:                   0                    0 pps
  Output packets:                   0                    0 pps
Layer3 Transit Statistics
 IPv4 Unicast
   Input  bytes  :                  0                    0 bps
   Output bytes  :                  0                    0 bps
   Input  packets:                  0                    0 pps
   Output packets:                  0                    0 pps
IPv4 Multicast
   Input  bytes  :                  0                    0 bps
   Output bytes  :                  0                    0 bps
   Input  packets:                  0                    0 pps
   Output packets:                  0                    0 pps
 IPv6 Unicast
   Input  bytes  :                  0                    0 bps
   Output bytes  :                  0                    0 bps
   Input  packets:                  0                    0 pps
   Output packets:                  0                    0 pps
 IPv6 Multicast
    Input  bytes  :                 0                    0 bps
    Output bytes  :                 0                    0 bps
    Input  packets:                 0                    0 pps
    Output packets:                 0                    0 pps
```

show interfaces extensive (IPv4 sub-interface)

```
user@host>show interfaces ge-0/0/0 extensive
   Logical interface ge-0/0/0.0 (Index 334) (SNMP ifIndex 630) (Generation 143)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Traffic statistics:
    Input  bytes  :         618
    Output bytes  :         972
    Input  packets:         12
    Output packets:         16
   IPv4 transit statistics:
    Input  bytes  :         72
    Output bytes  :         398
    Input  packets:          1
    Output packets:          5
    Local statistics:
    Input  bytes  :         618
    Output bytes  :         972
    Input  packets:         12
    Output packets:         16
    Transit statistics:
    Input  bytes  :          0          0 bps
    Output bytes  :          0          0 bps
    Input  packets:          0          0 pps
    Output packets:          0          0 pps
   IPv4 transit statistics:
     Input  bytes  :         0          0 bps
     Output bytes  :         0          0 bps
     Input  packets:         0          0 pps
     Output packets:         0          0 pps
```

show interfaces extensive (pseudowire interface)

```
user@host> show interfaces ps0.1 extensive
   Logical interface ps0.1 (index 402) (SNMP ifindex 68800) (Generation 8416)
   Egress queues: 8 supported, 4 in use
   Queue counters:    Queued packets    Transmitted packets    Dropped packets
        0                  0                  0                      0
        1                  0                  0                      0
        2                  0                  0                      0
```

```
     3                  0                  0                    0
Flags: Up 0x4000 VLAN-Tag [ 0x9100.512 ] Encapsulation: ENET2
Link:
   lt-2/0/10.32767
   lt-4/0/10/32767
Traffic statistics:
   Input bytes :    251000
   Output bytes :   247242
   Input packets:   2998
   Output packets:  2934
Local Statistics:
   Input bytes :    246104      6648 bps
   Output bytes :   247242      6936 bps
   Input packets:   2936        9 pps
   Output packets:  2934        10 pps
IPv4 transit statistics:
   Input bytes :     0      0 bps
   Output bytes :    0      0 bps
   Input packets:    0      0 bps
   Output packets:   0      0 bps
IPv6 transit statistics:
   Input bytes :     0      0 bps
   Output bytes :    0      0 bps
   Input packets:    0      0 bps
   Output packets:   0      0 bps
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces extensive satellite-device

IN THIS SECTION

- Syntax | 20748
- Description | 20748
- Options | 20748

- Required Privilege Level | 20748
- Output Fields | 20748
- Sample Output | 20758
- Release Information | 20773

Syntax

```
show interfaces extensive satellite-device (device-alias | all)
```

Description

Display the satellite device extended ports in a Junos Fusion.

Options

device-alias *device-alias*	Display extended port information for the satellite device using the specified device alias only.
all	Display information for all extended ports and aggregated Ethernet interfaces with extended ports as members configured on all of the satellite devices.

Required Privilege Level

view

Output Fields

Table 636 on page 20749 lists the output fields for the `show interfaces extensive satellite-device` command. Output fields are listed in the approximate order in which they appear.

Table 636: show interfaces extensive satellite-device Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
Device flags	Information about the physical device.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**. **NOTE**: This field is only displayed if asymmetric flow control is not configured.	All levels
`Pad to minimum frame size`	Pad Tx VLAN-tagged frame to minimum of 68 bytes.	
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Current address	Configured MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour: :minute.second.timezone*** **(*hour.minute.second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).**	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
`Extended port information`	Satellite device port ID	

Table 636: show interfaces extensive satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	**detail extensive**
IPv6 transit statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	**detail extensive**

Table 636: show interfaces extensive satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame terminates and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are greater than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 636: show interfaces extensive satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame terminates and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**
Queue counters	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Queue Number	The CoS queue number and the forwarding classes mapped to the queue number. The **Mapped forwarding class** column lists the forwarding classes mapped to each CoS queue.	**detail extensive**

Table 636: show interfaces extensive satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none

Table 636: show interfaces extensive satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem. • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of packets that exceeds the configured MTU. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.	**extensive**

Table 636: show interfaces extensive satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels

Table 636: show interfaces extensive satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface.	All levels
`Statistics`	• `Packets` • `pps` • `Bytes` • `bps`	All levels

`Bundle`	Provide information for each active bundle link.
	• `Input`
	  • `Packets—`
	  • `pps`
	  • `Bytes`
	  • `bps`
	  • `Output`
	  • `Packets—`
	  • `pps`
	  • `Bytes`
	  • `bps`
		All levels

Table 636: show interfaces extensive satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Adaptive Statistics`	• `Adaptive Adjusts` • `Adaptive Scans` • `Adaptive Updates`	All levels
`Link`	Link state: up or down.	All levels

LACP info	LACP state information for each aggregated interface:
	• **Role priority**—Role played by the interface. It can be one of the following:
	  • **Actor**—Local device participating in LACP negotiation.
	  • **Partner**—Remote device participating in LACP negotiation.
	  • `System identifier`—48-bit (6-byte) globally unique field.
	  • `System priority`—LACP system priority at the aggregated Ethernet interface level. This system priority value takes precedence over a system priority value configured at the global [edit chassis] hierarchy level.
	  • `Port number`
	  • `Port key`
	  • `Port`
		All levels

LACP Statistics	LACP statistics are returned when the **extensive** option is used and provides the following information: • **LACP Rx**—LACP received counter that increments for each normal hello. • **LACP Tx**—Number of LACP transmit packet errors logged. • **Unknown Rx**—Number of unrecognized packet errors logged. • **Illegal Rx**—Number of invalid packets received.	All levels

Table 636: show interfaces extensive satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Marker statistics`	Marker statistics are returned when the **extensive** option is used and provides the following information: • **Marker Rx**—Marker received counter that increments for each normal hello. • **Resp Tx**—Number of RESP transmit packet errors logged. • **Unknown Rx**—Number of unrecognized packet errors logged. • **Illegal Rx**—Number of invalid packets received.	All levels
`Protocol`	Protocol family configured on the logical interface.	All levels
`MTU`	MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked Adjusted.	All levels
`Generation`	Unique number for use by Juniper Networks technical support only.	All levels
`Route table`	Routing table in which this address exists. For example, Route table:0 refers to inet.0.	All levels
`Mesh table`	Information regarding mesh topology.	All levels

Sample Output

show interfaces extensive satellite-device all

```
user@aggregation-device> show interfaces extensive satellite-device all
Physical interface: ae0 (Extended Port, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 574, Generation: 131
  Link-level type: Ethernet, MTU: 1514, Speed: 2Gbps, BPDU Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
```

```
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: f4:b5:2f:f7:3f:c0, Hardware address: f4:b5:2f:f7:3f:c0
  Last flapped   : 2015-03-31 18:36:43 PDT (07:05:56 ago)
  Statistics last cleared: Never
  Extended port information:
    Satellite device port id  : 415
  Traffic statistics:
   Input  bytes  :             13515908                 2032 bps
   Output bytes  :             12289920                 2032 bps
   Input  packets:                99514                    2 pps
   Output packets:                96015                    2 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 7 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                               0                95867                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
    4                               0                    0                    0
    5                               0                    0                    0
    7                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   FC0
    1                   FC1
    2                   FC2
    3                   FC3
    4                   FC4
    5                   FC5, be-3
    7                   be-2

  Logical interface ae0.0 (Index 337) (SNMP ifIndex 575) (Generation 1194)
```

```
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          1729          0        601692            0
        Output:             0          0             0            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      ge-104/0/24.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      ge-103/0/0.0
        Input :          1729          0        601692            0
        Output:             0          0             0            0
    LACP info:        Role     System             System      Port    Port  Port
                             priority          identifier  priority  number   key
      ge-104/0/24.0    Actor      127  f4:b5:2f:f7:3f:c0       127      31     1
      ge-104/0/24.0  Partner      127  f4:b5:2f:41:0a:40       127      24     1
      ge-103/0/0.0    Actor       127  f4:b5:2f:f7:3f:c0       127       7     1
      ge-103/0/0.0  Partner       127  f4:b5:2f:41:0a:40       127       1     1
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-104/0/24.0          25470       25495            0            0
      ge-103/0/0.0           25469       25512            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-104/0/24.0              0           0            0            0
      ge-103/0/0.0               0           0            0            0
    Protocol bridge, MTU: 1514, Generation: 1229, Route table: 3, Mesh Group: __all_ces__
Physical interface: ae1 (Extended Port, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 790, Generation: 132
  Link-level type: Ethernet, MTU: 1514, Speed: 200mbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: f4:b5:2f:f7:3f:c1, Hardware address: f4:b5:2f:f7:3f:c1
  Last flapped   : 2015-03-31 18:36:44 PDT (07:05:55 ago)
  Statistics last cleared: Never
  Extended port information:
    Satellite device port id  : 431
```

```
  Traffic statistics:
   Input  bytes  :             13285288                 2032 bps
   Output bytes  :             12166400                 2032 bps
   Input  packets:                98447                    2 pps
   Output packets:                95050                    2 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 7 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                               0                94909                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
    4                               0                    0                    0
    5                               0                    0                    0
    7                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   FC0
    1                   FC1
    2                   FC2
    3                   FC3
    4                   FC4
    5                   FC5, be-3
    7                   be-2

  Logical interface ae1.0 (Index 338) (SNMP ifIndex 1216) (Generation 1195)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          2785          0        688380            0
        Output:             0          0             0            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
```

```
    Link:
      ge-104/0/25.0
        Input :            10          0           600            0
        Output:             0          0             0            0
      ge-103/0/1.0
        Input :          2775          0        687780            0
        Output:             0          0             0            0
    LACP info:        Role     System             System      Port    Port  Port
                               priority          identifier  priority  number   key
      ge-104/0/25.0    Actor      127  f4:b5:2f:f7:3f:c0       127      32     2
      ge-104/0/25.0  Partner      127  f4:b5:2f:41:0a:40       127      25     2
      ge-103/0/1.0    Actor       127  f4:b5:2f:f7:3f:c0       127       8     2
      ge-103/0/1.0  Partner       127  f4:b5:2f:41:0a:40       127       2     2
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-104/0/25.0          25470       25494            0            0
      ge-103/0/1.0           25469       25513            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-104/0/25.0              0           0            0            0
      ge-103/0/1.0               0           0            0            0
    Protocol bridge, MTU: 1514, Generation: 1230, Route table: 3, Mesh Group: __all_ces__
Physical interface: ae0 (Extended Port, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 574, Generation: 131
  Link-level type: Ethernet, MTU: 1514, Speed: 2Gbps, BPDU Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: f4:b5:2f:f7:3f:c0, Hardware address: f4:b5:2f:f7:3f:c0
  Last flapped   : 2015-03-31 18:36:43 PDT (07:05:56 ago)
  Statistics last cleared: Never
  Extended port information:
    Satellite device port id  : 415
  Traffic statistics:
   Input  bytes  :             13515908                 2032 bps
   Output bytes  :             12289920                 2032 bps
   Input  packets:                99514                    2 pps
   Output packets:                96015                    2 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
```

```
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 7 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                               0                95867                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
    4                               0                    0                    0
    5                               0                    0                    0
    7                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   FC0
    1                   FC1
    2                   FC2
    3                   FC3
    4                   FC4
    5                   FC5, be-3
    7                   be-2

  Logical interface ae0.0 (Index 337) (SNMP ifIndex 575) (Generation 1194)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          1729          0        601692            0
        Output:             0          0             0            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      ge-104/0/24.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      ge-103/0/0.0
        Input :          1729          0        601692            0
        Output:             0          0             0            0
    LACP info:        Role     System             System      Port    Port  Port
                             priority          identifier  priority  number   key
```

```
      ge-104/0/24.0    Actor      127  f4:b5:2f:f7:3f:c0       127      31     1
      ge-104/0/24.0  Partner      127  f4:b5:2f:41:0a:40       127      24     1
      ge-103/0/0.0    Actor       127  f4:b5:2f:f7:3f:c0       127       7     1
      ge-103/0/0.0  Partner       127  f4:b5:2f:41:0a:40       127       1     1
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-104/0/24.0          25470       25495            0            0
      ge-103/0/0.0           25469       25512            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-104/0/24.0              0           0            0            0
      ge-103/0/0.0               0           0            0            0
    Protocol bridge, MTU: 1514, Generation: 1229, Route table: 3, Mesh Group: __all_ces__
Physical interface: ae1 (Extended Port, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 790, Generation: 132
  Link-level type: Ethernet, MTU: 1514, Speed: 200mbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: f4:b5:2f:f7:3f:c1, Hardware address: f4:b5:2f:f7:3f:c1
  Last flapped   : 2015-03-31 18:36:44 PDT (07:05:55 ago)
  Statistics last cleared: Never
  Extended port information:
    Satellite device port id  : 431
  Traffic statistics:
   Input  bytes  :             13285288                 2032 bps
   Output bytes  :             12166400                 2032 bps
   Input  packets:                98447                    2 pps
   Output packets:                95050                    2 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 7 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                94909                    0
```

```
  1                                 0                    0                    0
  2                                 0                    0                    0
  3                                 0                    0                    0
  4                                 0                    0                    0
  5                                 0                    0                    0
  7                                 0                    0                    0
Queue number:         Mapped forwarding classes
  0                   FC0
  1                   FC1
  2                   FC2
  3                   FC3
  4                   FC4
  5                   FC5, be-3
  7                   be-2

Logical interface ae1.0 (Index 338) (SNMP ifIndex 1216) (Generation 1195)
  Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
  Statistics        Packets        pps         Bytes          bps
  Bundle:
      Input :          2785          0        688380            0
      Output:             0          0             0            0
  Adaptive Statistics:
      Adaptive Adjusts:          0
      Adaptive Scans  :          0
      Adaptive Updates:          0
  Link:
    ge-104/0/25.0
      Input :            10          0           600            0
      Output:             0          0             0            0
    ge-103/0/1.0
      Input :          2775          0        687780            0
      Output:             0          0             0            0
  LACP info:        Role     System             System      Port    Port  Port
                           priority          identifier  priority  number   key
    ge-104/0/25.0    Actor      127  f4:b5:2f:f7:3f:c0       127      32     2
    ge-104/0/25.0  Partner      127  f4:b5:2f:41:0a:40       127      25     2
    ge-103/0/1.0    Actor       127  f4:b5:2f:f7:3f:c0       127       8     2
    ge-103/0/1.0  Partner       127  f4:b5:2f:41:0a:40       127       2     2
  LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
    ge-104/0/25.0          25470       25494            0            0
    ge-103/0/1.0           25469       25513            0            0
  Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
    ge-104/0/25.0              0           0            0            0
```

```
      ge-103/0/1.0               0           0            0            0
    Protocol bridge, MTU: 1514, Generation: 1230, Route table: 3, Mesh Group: __all_ces__

Physical interface: ae0 (Extended Port, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 574, Generation: 131
  Link-level type: Ethernet, MTU: 1514, Speed: 2Gbps, BPDU Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: f4:b5:2f:f7:3f:c0, Hardware address: f4:b5:2f:f7:3f:c0
  Last flapped   : 2015-03-31 18:36:43 PDT (07:05:56 ago)
  Statistics last cleared: Never
  Extended port information:
    Satellite device port id  : 415
  Traffic statistics:
   Input  bytes  :             13515908                 2032 bps
   Output bytes  :             12289920                 2032 bps
   Input  packets:                99514                    2 pps
   Output packets:                96015                    2 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 7 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                               0                95867                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
    4                               0                    0                    0
    5                               0                    0                    0
    7                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   FC0
```

```
    1                   FC1
    2                   FC2
    3                   FC3
    4                   FC4
    5                   FC5, be-3
    7                   be-2

  Logical interface ae0.0 (Index 337) (SNMP ifIndex 575) (Generation 1194)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          1729          0        601692            0
        Output:             0          0             0            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      ge-104/0/24.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      ge-103/0/0.0
        Input :          1729          0        601692            0
        Output:             0          0             0            0
    LACP info:        Role     System             System      Port    Port  Port
                             priority          identifier  priority  number   key
      ge-104/0/24.0    Actor      127  f4:b5:2f:f7:3f:c0       127      31     1
      ge-104/0/24.0  Partner      127  f4:b5:2f:41:0a:40       127      24     1
      ge-103/0/0.0    Actor       127  f4:b5:2f:f7:3f:c0       127       7     1
      ge-103/0/0.0  Partner       127  f4:b5:2f:41:0a:40       127       1     1
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-104/0/24.0          25470       25495            0            0
      ge-103/0/0.0           25469       25512            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-104/0/24.0              0           0            0            0
      ge-103/0/0.0               0           0            0            0
    Protocol bridge, MTU: 1514, Generation: 1229, Route table: 3, Mesh Group: __all_ces__
Physical interface: ae1 (Extended Port, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 790, Generation: 132
  Link-level type: Ethernet, MTU: 1514, Speed: 200mbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Pad to minimum frame size: Disabled
```

```
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: f4:b5:2f:f7:3f:c1, Hardware address: f4:b5:2f:f7:3f:c1
  Last flapped   : 2015-03-31 18:36:44 PDT (07:05:55 ago)
  Statistics last cleared: Never
  Extended port information:
    Satellite device port id  : 431
  Traffic statistics:
   Input  bytes  :             13285288                 2032 bps
   Output bytes  :             12166400                 2032 bps
   Input  packets:                98447                    2 pps
   Output packets:                95050                    2 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 7 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                               0                94909                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
    4                               0                    0                    0
    5                               0                    0                    0
    7                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   FC0
    1                   FC1
    2                   FC2
    3                   FC3
    4                   FC4
    5                   FC5, be-3
    7                   be-2

  Logical interface ae1.0 (Index 338) (SNMP ifIndex 1216) (Generation 1195)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
```

```
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          2785          0        688380            0
        Output:             0          0             0            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      ge-104/0/25.0
        Input :            10          0           600            0
        Output:             0          0             0            0
      ge-103/0/1.0
        Input :          2775          0        687780            0
        Output:             0          0             0            0
    LACP info:        Role     System             System      Port    Port  Port
                               priority          identifier  priority  number   key
      ge-104/0/25.0    Actor      127  f4:b5:2f:f7:3f:c0       127      32     2
      ge-104/0/25.0  Partner      127  f4:b5:2f:41:0a:40       127      25     2
      ge-103/0/1.0    Actor       127  f4:b5:2f:f7:3f:c0       127       8     2
      ge-103/0/1.0  Partner       127  f4:b5:2f:41:0a:40       127       2     2
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-104/0/25.0          25470       25494            0            0
      ge-103/0/1.0           25469       25513            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-104/0/25.0              0           0            0            0
      ge-103/0/1.0               0           0            0            0
    Protocol bridge, MTU: 1514, Generation: 1230, Route table: 3, Mesh Group: __all_ces__
Physical interface: ge-101/0/7 (Extended Port, Enabled, Physical link is Down
  Interface index: 328, SNMP ifIndex: 1587, Generation: 331
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 10:0e:7e:bf:2d:0c, Hardware address: 10:0e:7e:bf:2d:0c
  Last flapped   : Never
  Statistics last cleared: Never
  Extended port information:
    Satellite device port id  : 143
```

```
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 7 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                               0                    0                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
    4                               0                    0                    0
    5                               0                    0                    0
    7                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   FC0
    1                   FC1
    2                   FC2
    3                   FC3
    4                   FC4
    5                   FC5, be-3
    7                   be-2
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
```

```
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
    Total errors                             0                0
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                       0
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 0, CAM source filters: 0
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 FC0                    95      950000000    95              0      low    none
    3 FC3                     5       50000000     5              0      low    none
  Interface transmit statistics: Disabled

Physical interface: ge-101/0/8 (Extended Port, Enabled, Physical link is Down
  Interface index: 329, SNMP ifIndex: 1586, Generation: 332
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 10:0e:7e:bf:2d:0d, Hardware address: 10:0e:7e:bf:2d:0d
  Last flapped   : Never
  Statistics last cleared: Never
  Extended port information:
    Satellite device port id  : 159
```

```
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 7 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                               0                    0                    0
    1                               0                    0                    0
    2                               0                    0                    0
    3                               0                    0                    0
    4                               0                    0                    0
    5                               0                    0                    0
    7                               0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   FC0
    1                   FC1
    2                   FC2
    3                   FC3
    4                   FC4
    5                   FC5, be-3
    7                   be-2
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
```

```
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
    Total errors                             0                0
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                       0
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 0, CAM source filters: 0
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 FC0                    95      950000000    95              0      low    none
    3 FC3                     5       50000000     5              0      low    none
  Interface transmit statistics: Disabled
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

- Configuring or Expanding a Junos Fusion Enterprise

show interfaces fabric

IN THIS SECTION

- Syntax | 20774
- Description | 20774
- Options | 20774
- Required Privilege Level | 20775
- Output Fields | 20775
- Sample Output | 20786
- Release Information | 20803

Syntax

```
show interfaces fabric
<interface-name>
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified fabric interface.

Options

interface-name (QFabric systems only) Either the serial number or the alias of the QFabric switch component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and not contain any colons.

brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
routing-instance (all | *instance-name*)	(Optional) Display all routing instances or the name of an individual routing instance.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 637 on page 20775 lists the output fields for the `show interfaces fabric` command. Output fields are listed in the approximate order in which they appear.

Table 637: show interfaces fabric Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Type	Physical interface type; for example, Ethernet.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none

Table 637: show interfaces fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Clocking	Reference clock source.	detail
Speed	Speed at which the interface is running.	All levels
Duplex	Duplex mode of the interface, either Full-Duplex or Half-Duplex.	All levels
MAC-REWRITE Error	Specifies if the encapsulation of the packet has been changed.	none
BPDU Error	Specifies if a BPDU has been received on a blocked interface.	none
Loopback	Loopback status: Enabled or Disabled. If loopback is enabled, type of loopback: Local or Remote.	All levels
Source filtering	Source filtering status: Enabled or Disabled.	All levels
Flow control	Flow control status: Enabled or Disabled. This field is only displayed if asymmetric flow control is not configured.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
CoS queues	Number of CoS queues configured.	detail extensive none

Table 637: show interfaces fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Hold-Times	Current interface hold-time up and hold-time down, in milliseconds.	detail
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped: *year-month-day hour.minute.second.timezone* (*hour.minute.second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).**	detail extensive none
Statistics last cleared	Date, time, and how long ago the statistics for the interface were cleared. The format is **Statistics last cleared: *year-month-day hour.minute.second.timezone* (*hour.minute.second* ago)**. For example, 2010-05-17 07:51:28 PDT (00:04:33 ago).	detail extensive
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail extensive

Table 637: show interfaces fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IPv6 transit statistics	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface: • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail extensive

Table 637: show interfaces fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • Errors—Sum of the incoming frame terminated and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the `ignore-l3-incompletes` statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops.	extensive

Table 637: show interfaces fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame terminated and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the fabric interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops.	extensive
Egress queues	Total number of egress queues supported on the specified interface.	detail extensive

Table 637: show interfaces fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Queue counters	CoS queue number and its associated user-configured forwarding class name. • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism.	detail extensive
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	detail extensive none

Table 637: show interfaces fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
MAC statistics	Receive and Transmit statistics reported by the PIC's MAC subsystem. • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of packets that exceed the configured MTU. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.	extensive

Table 637: show interfaces fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0. • Code violations—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	
Packet Forwarding Engine Configuration	Information about the configuration of the Packet Forwarding Engine: • Destination slot—FPC slot number. • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	extensive
Logical Interface		
Item	Type of QFabric system component being viewed. Possible values include Node group, Interconnect device, Fabric control, Fabric manager, Diagnostic routing engine, and Ungrouped Node device.	none
Identifier	Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.	none

Table 637: show interfaces fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Connection	Status of a QFabric system component: either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for the listed component.	none
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured, Failed (unsuccessful), Pending (in the process of being written or retried), or Unknown.	none
Node group	Name of the Node groups associated with the QFabric system, and the Node devices assigned to each Node group. The group can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for the devices in the group. This field also displays the serial ID for the Node group and the status for the Node group.	none
Fabric control	Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition. The fabric control Routing Engine can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric control Routing Engine.	none
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none

Table 637: show interfaces fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Information about the logical interface. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.	All levels
Encapsulation	Encapsulation method used on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail extensive
Local statistics	• Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail extensive

Table 637: show interfaces fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Transit statistics	• Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail extensive
protocol-family	Protocol family configured on the logical interface. If the protocol is inet, the IP address of the interface is also displayed.	brief
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none

Sample Output

show interfaces fabric

```
user@switch> show interfaces fabric
Item                    Identifier              Connection      Configuration
Node group
  BBAK3775                                      Connected       Configured
  NW-NG-0                                       Connected       Configured
  P2659-C                                       Connected       Configured
  ptor-0                                        Connected       Configured
Fabric control
  FC-0                                          Connected       Configured
  FC-1                                          Connected       Configured
```

show interfaces fabric brief

```
user@switch> show interfaces fabric brief
Physical interface: BBAK0372:fte-0/1/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface BBAK0372:fte-0/1/0.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: BBAK0372:fte-0/1/2, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface BBAK0372:fte-0/1/2.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: BBAK0394:fte-0/1/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000

  Logical interface BBAK0394:fte-0/1/0.32768
    Flags: SNMP-Traps Encapsulation: ENET2
    eth-switch

Physical interface: BBAK0394:fte-0/1/2, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000

  Logical interface BBAK0394:fte-0/1/2.32768
    Flags: SNMP-Traps Encapsulation: ENET2
    eth-switch
```

```
Physical interface: BBAK3775:bme0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running

  Logical interface BBAK3775:bme0.0
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.1/2
          128.0.0.16/2
          128.0.32.0/2
    tnp   0x10

  Logical interface BBAK3775:bme0.1
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.13/2
          128.0.130.0/2

  Logical interface BBAK3775:bme0.2
    Flags: Encapsulation: ENET2
    inet  128.0.0.13/8
          128.0.130.0/8
          169.254.128.13/16
          169.254.193.0/16

Physical interface: BBAK3775:qfabric, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface BBAK3775:qfabric.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet
    mpls
    eth-switch

Physical interface: BBAK3775:vcp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface BBAK3775:vcp0.32769
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: BBAK3775:vcp1, Enabled, Physical link is Up
```

```
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface BBAK3775:vcp1.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: BBAK3775:vcp2, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface BBAK3775:vcp2.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: BBAK3775:fte-0/1/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface BBAK3775:fte-0/1/0.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: EE3093:fte-0/1/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface EE3093:fte-0/1/0.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: EE3093:fte-0/1/2, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface EE3093:fte-0/1/2.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch
```

```
Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-0/0/0.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-0/0/4, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-0/0/4.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-0/0/6.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-0/0/13.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
```

```
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-0/0/15.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-1/0/2.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-1/0/7.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-1/0/10.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:bme0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running

  Logical interface IC-WS001:bme0.0
    Flags: LinkAddress 0-0 Encapsulation: ENET2
```

```
    inet  128.0.32.0       --> 0/0

  Logical interface IC-WS001:bme0.1
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.7/2
          128.0.130.2/2

  Logical interface IC-WS001:bme0.2
    Flags: Encapsulation: ENET2
    inet  128.0.0.7/8
          128.0.130.2/8
          169.254.128.7/16
          169.254.193.1/16

Physical interface: IC-WS001:bme1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface IC-WS001:bme1.0
    Flags: Encapsulation: ENET2
    inet  128.0.0.1/2
          128.0.0.4/2
          128.0.0.16/2
          128.0.0.17/2
          128.0.0.24/2
          128.0.0.25/2
          128.0.0.26/2
          128.0.0.28/2
          128.0.0.29/2
          128.0.0.31/2
    tnp   0x4

Physical interface: IC-WS001:qfabric, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface IC-WS001:qfabric.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet
    mpls
    eth-switch
```

```
Physical interface: IC-WS001:pme0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme2, Enabled, Physical link is Down
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme3, Enabled, Physical link is Down
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

Physical interface: IC-WS001:vcp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface IC-WS001:vcp0.32769
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface IC-WS001:vcp1.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp2, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface IC-WS001:vcp2.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp3, Enabled, Physical link is Down
```

```
  Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface IC-WS001:vcp3.32768
    Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp4, Enabled, Physical link is Down
  Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface IC-WS001:vcp4.32768
    Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2

Physical interface: NW-NG-0:bme0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running

  Logical interface NW-NG-0:bme0.0
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.1/2
          128.0.0.5/2
          128.0.32.0/2
    tnp   0x5

  Logical interface NW-NG-0:bme0.1
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.9/2
          128.0.128.4/2

  Logical interface NW-NG-0:bme0.2
    Flags: Encapsulation: ENET2
    inet  128.0.0.9/8
          128.0.128.68/8
          169.254.128.9/16
          169.254.192.34/16

Physical interface: NW-NG-0:qfabric, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface NW-NG-0:qfabric.0
    Flags: SNMP-Traps Encapsulation: ENET2
```

```
    inet
    mpls
    eth-switch

Physical interface: NW-NG-0:vcp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface NW-NG-0:vcp0.32769
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: NW-NG-0:vcp1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface NW-NG-0:vcp1.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: P2659-C:bme0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running

  Logical interface P2659-C:bme0.0
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.1/2
          128.0.0.16/2
          128.0.32.0/2
    tnp   0x10

  Logical interface P2659-C:bme0.1
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.8/2
          128.0.130.4/2

  Logical interface P2659-C:bme0.2
    Flags: Encapsulation: ENET2
    inet  128.0.0.8/8
          128.0.130.4/8
          169.254.128.8/16
          169.254.193.2/16

Physical interface: P2659-C:qfabric, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified, Speed: Unspecified
```

```
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface P2659-C:qfabric.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet
    mpls
    eth-switch

Physical interface: P2659-C:vcp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface P2659-C:vcp0.32769
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: P2659-C:vcp1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface P2659-C:vcp1.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: P2659-C:vcp2, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface P2659-C:vcp2.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: P2659-C:fte-0/1/2, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface P2659-C:fte-0/1/2.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: ptor-0:bme0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running
```

```
  Logical interface ptor-0:bme0.0
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.1/2
          128.0.0.17/2
          128.0.32.0/2
    tnp   0x11

  Logical interface ptor-0:bme0.1
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.16/2
          128.0.130.18/2

  Logical interface ptor-0:bme0.2
    Flags: Encapsulation: ENET2
    inet  128.0.0.16/8
          128.0.130.18/8
          169.254.128.16/16
          169.254.193.9/16

Physical interface: ptor-0:qfabric, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface ptor-0:qfabric.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet
    mpls
    eth-switch

Physical interface: ptor-0:vcp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface ptor-0:vcp0.32769
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: ptor-0:vcp1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface ptor-0:vcp1.32768
```

```
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: ptor-0:vcp2, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface ptor-0:vcp2.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2
```

show interfaces fabric detail

```
user@switch> show interfaces fabric detail
Physical interface: BBAK0372:fte-0/1/0, Enabled, Physical link is Up
  Interface index: 49165, SNMP ifIndex: 1212678666, Generation: 140
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 84:18:88:d1:fa:1f, Hardware address: 84:18:88:d1:fa:1f
  Last flapped   : 2012-11-09 21:36:41 UTC (4d 00:23 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :             14256654                    0 bps
   Output bytes  :              9618986                    0 bps
   Input  packets:                90511                    0 pps
   Output packets:                60101                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Egress queues: 12 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    3 fcoe                           0                    0                    0
    4 no-loss                        0                    0                    0
    7 network-cont                   0                    0                    0
    8 mcast                          0                    0                    0
  Active alarms  : None
```

```
  Active defects : None

  Logical interface BBAK0372:fte-0/1/0.32768 (Index 71) (SNMP ifIndex 1212678667) (Generation
136)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             12450372
     Output bytes  :             11986557
     Input  packets:                90510
     Output packets:                62750
    Local statistics:
     Input  bytes  :             12450372
     Output bytes  :             11986557
     Input  packets:                90510
     Output packets:                62750
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol eth-switch, MTU: 0, Generation: 163, Route table: 0
```

show interfaces fabric extensive

```
user@switch> show interfaces fabric extensive

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Up
  Interface index: 49176, SNMP ifIndex: 1209008655, Generation: 155
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:00:00:00:06, Hardware address: 00:00:00:00:00:06
  Last flapped   : 2012-11-13 23:53:30 UTC (00:53:20 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                91179                    0 bps
   Output bytes  :         361268221791            952985992 bps
```

```
   Input  packets:                  590                    0 pps
   Output packets:           2580487185               850880 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 12 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 fabric_fcset                   0                    0                    0
    1 fabric_fcset                   0                    0                    0
    2 fabric_fcset                   0                    0                    0
    3 fabric_fcset                   0                    0                    0
    4 fabric_fcset                   0                    0                    0
    5 fabric_fcset                   0                    0                    0
    6 fabric_fcset                   0                    0                    0
    7 fabric_fcset                   0                    0                    0
    8 fabric_fcset                   0           2582632925                    0
    9 fabric_fcset                   0                    0                    0
    10 fabric_fcset                   0                    0                    0
    11 fabric_fcset                   0                    0                    0
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                         91179     361268221791
    Total packets                          590       2580487185
    Unicast packets                        590       2580487185
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
```

```
    Code violations                          0
  MAC Priority Flow Control Statistics:
    Priority :  0                             0                0
    Priority :  1                             0                0
    Priority :  2                             0                0
    Priority :  3                             0                0
    Priority :  4                             0                0
    Priority :  5                             0                0
    Priority :  6                             0                0
    Priority :  7                             0                0
  Packet Forwarding Engine configuration:
    Destination slot: 0
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 best-effort             5     2000000000     5              0      low    none
    3 fcoe                   35    14000000000    35              0      low    none
    4 no-loss                35    14000000000    35              0      low    none
    7 network-control         5     2000000000     5              0      low    none
    8 mcast                  20     8000000000    20              0      low    none

  Logical interface IC-WS001:fte-0/0/6.32768 (Index 85) (SNMP ifIndex 1209008656) (Generation
150)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                79496
     Output bytes  :               179860
     Input  packets:                 590
     Output packets:                 948
    Local statistics:
     Input  bytes  :                79496
     Output bytes  :               179860
     Input  packets:                 590
     Output packets:                 948
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
    Protocol eth-switch, MTU: 0, Generation: 178, Route table: 0
```

show interfaces fabric terse

```
user@switch> show interfaces fabric terse
Item                    Identifier              Connection      Configuration
Node group
  BBAK3775                                      Connected       Configured
  NW-NG-0                                       Connected       Configured
  P2659-C                                       Connected       Configured
  ptor-0                                        Connected       Configured
Fabric control
  FC-0                                          Connected       Configured
  FC-1                                          Connected       Configured
```

show interfaces fabric device-name

```
user@switch> show interfaces fabric IC-WS001:fte-0/0/13
Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Up
  Interface index: 49177, SNMP ifIndex: 1209008767
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:00:00:00:00:0d, Hardware address: 00:00:00:00:00:0d
  Last flapped   : 2012-11-13 23:55:15 UTC (00:55:38 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface IC-WS001:fte-0/0/13.32768 (Index 86) (SNMP ifIndex 1209008768)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 748
    Output packets: 954
    Protocol eth-switch, MTU: 0
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Monitoring Interface Status and Traffic

Troubleshooting Network Interfaces

Troubleshooting an Aggregated Ethernet Interface

Junos OS Network Interfaces Library for Routing Devices

show interfaces (far-end-interval)

IN THIS SECTION

- Syntax | 20803
- Description | 20803
- Required Privilege Level | 20804
- Output Fields | 20804
- Sample Output | 20804
- Release Information | 20805

Syntax

```
show interfaces far-end-interval interface-fpc/pic/port
```

Description

On channelized interfaces, display the far end interval data for the specified interface.

Required Privilege Level

view

Output Fields

Table 638 on page 20804 lists the output fields for the `show interfaces far-end-interval` command. Output fields are listed in the approximate order in which they appear.

Table 638: show interfaces far-end-interval Output Fields

Field Name	Field Description
Physical interface	Interface FPC/PIC/port values.
SNMP ifIndex	SNMP interface index value.
ES-L/P	Error detection—Errored seconds.
SES-L/P	Error detection—Severely errored seconds.
UAS-L/P	Error detection—Unavailable seconds.

Sample Output

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:30-04:45:
```

```
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:00-04:15:
...
```

show interfaces far-end-interval coc1-5/2/1:1

```
user@host> run show interfaces far-end-interval coc1-5/2/1:1
 Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:00-04:15:
```

Release Information

Command introduced in Junos OS Release 9.4.

show interfaces (Fast Ethernet)

IN THIS SECTION

- Syntax | 20806
- Description | 20806
- Options | 20806

- Required Privilege Level | 20806
- Output Fields | 20807
- Sample Output | 20828
- Release Information | 20831

Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified Fast Ethernet interface.

Options

interface-type	On M Series and T Series routers, the interface type is `fe-`*`fpc/pic/port`*.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 639 on page 20807 lists the output fields for the `show interfaces` (Fast Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 639: show interfaces Fast Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical interface`	Name of the physical interface.	All levels
`Enabled`	State of the interface. Possible values are described in the "Enabled Field" section under Common Output Fields Description.	All levels
`Interface index`	Index number of the physical interface, which reflects its initialization sequence.	`detail` `extensive` none
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail` `extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`
`Link-level type`	Encapsulation being used on the physical interface.	All levels
`MTU`	Maximum transmission unit size on the physical interface.	All levels
`Link-mode`	Type of link connection configured for the physical interface: `Full-duplex` or `Half-duplex`	`extensive`
`Speed`	Speed at which the interface is running.	All levels
`Loopback`	Loopback status: `Enabled` or `Disabled`. If loopback is enabled, type of loopback: `Local` or `Remote`.	All levels

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Source filtering`	Source filtering status: `Enabled` or `Disabled`.	All levels
`LAN-PHY mode`	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
`WAN-PHY mode`	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
`Unidirectional`	Unidirectional link mode status for 10-Gigabit Ethernet interface: `Enabled` or `Disabled` for parent interface; `Rx-only` or `Tx-only` for child interfaces.	All levels
`Flow control`	Flow control status: `Enabled` or `Disabled`.	All levels
`Auto-negotiation`	(Gigabit Ethernet interfaces) Autonegotiation status: `Enabled` or `Disabled`.	All levels
`Remote-fault`	(Gigabit Ethernet interfaces) Remote fault status: • `Online`—Autonegotiation is manually configured as online. • `Offline`—Autonegotiation is manually configured as offline.	All levels
`Device flags`	Information about the physical device. Possible values are described in the “Device Flags” section under Common Output Fields Description.	All levels
`Interface flags`	Information about the interface. Possible values are described in the “Interface Flags” section under Common Output Fields Description.	All levels
`Link flags`	Information about the link. Possible values are described in the “Links Flags” section under Common Output Fields Description.	All levels

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Wavelength`	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
`Frequency`	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
`CoS queues`	Number of CoS queues configured.	`detail extensive` none
`Schedulers`	(GigabitEthernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	`extensive`
`Hold-times`	Current interface hold-time up and hold-time down, in milliseconds.	`detail extensive`
`Current address`	Configured MAC address.	`detail extensive` none
`Hardware address`	Hardware MAC address.	`detail extensive` none
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second`*`:`*`timezone`* (*`hour`*`:`*`minute`*`:`*`second`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive` none
`Input Rate`	Input rate in bits per second (bps) and packets per second (pps).	None specified
`Output Rate`	Output rate in bps and pps.	None specified
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces command.	`detail extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input errors`	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Errors`—Sum of the incoming frame terminations and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • `L3 incompletes`—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` statement. • `L2 channel errors`—Number of times the software did not find a valid logical interface for an incoming frame. • `L2 mismatch timeouts`—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • `FIFO errors`—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `Resource errors`—Sum of transmit drops.	`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • `Errors`—Sum of the outgoing frame terminations and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Collisions`—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • `Aged packets`—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • `FIFO errors`—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `HS link CRC errors`—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • `MTU errors`—Number of packets whose size exceeded the MTU of the interface. • `Resource errors`—Sum of transmit drops.	`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Egress queues`	Total number of egress queues supported on the specified interface. **NOTE**: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the `show interfaces` command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenorio, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	`detail extensive`
`Queue counters (Egress)`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	`detail extensive`
`Ingress queues`	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Queue counters (Ingress)`	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	`extensive`
`Active alarms` and `Active defects`	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value `None` or `Link`. • `None`—There are no active defects or alarms. • `Link`—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	`detail extensive` none
`OTN FEC statistics`	The forward error correction (FEC) counters provide the following statistics:. • `Corrected Errors`—The count of corrected errors in the last second. • `Corrected Error Ratio`—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	
`PCS statistics`	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • `Bit errors`—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • `Errored blocks`—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	`detail extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`MAC statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC subsystem, including the following: • `Total octets` and `total packets`—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces command. • `Unicast packets`, `Broadcast packets`, and `Multicast packets`—Number of unicast, broadcast, and multicast packets. • `CRC/Align errors`—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • `FIFO error`—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • `MAC control frames`—Number of MAC control frames. • `MAC pause frames`—Number of MAC control frames with `pause` operational code. • `Oversized frames`—Number of frames that exceed 1518 octets. • `Jabber frames`—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • `Fragment frames`—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.	`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `VLAN tagged frames`—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • `Code violations`—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	
`OTN Received Overhead Bytes`	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	`extensive`
`OTN Transmitted Overhead Bytes`	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Filter statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • `Input packet count`—Number of packets received from the MAC hardware that the filter processed. • `Input packet rejects`—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • `Input DA rejects`—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • `Input SA rejects`—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • `Output packet count`—Number of packets that the filter has given to the MAC hardware. • `Output packet pad count`—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured.	`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Output packet error count`—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • `CAM destination filters`, `CAM source filters`—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	
`PMA PHY`	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `PHY Lock`—Phase-locked loop • `PHY Light`—Loss of optical signal	`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`WIS section`	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `BIP-B1`—Bit interleaved parity for SONET section overhead • `SEF`—Severely errored framing • `LOL`—Loss of light • `LOF`—Loss of frame • `ES-S`—Errored seconds (section) • `SES-S`—Severely errored seconds (section) • `SEFS-S`—Severely errored framing seconds (section)	`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`WIS line`	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. State other than `OK` indicates a problem. Subfields are: • `BIP-B2`—Bit interleaved parity for SONET line overhead • `REI-L`—Remote error indication (near-end line) • `RDI-L`—Remote defect indication (near-end line) • `AIS-L`—Alarm indication signal (near-end line) • `BERR-SF`—Bit error rate fault (signal failure) • `BERR-SD`—Bit error rate defect (signal degradation) • `ES-L`—Errored seconds (near-end line) • `SES-L`—Severely errored seconds (near-end line) • `UAS-L`—Unavailable seconds (near-end line) • `ES-LFE`—Errored seconds (far-end line) • `SES-LFE`—Severely errored seconds (far-end line) • `UAS-LFE`—Unavailable seconds (far-end line)	`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`WIS path`	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `BIP-B3`—Bit interleaved parity for SONET section overhead • `REI-P`—Remote error indication • `LOP-P`—Loss of pointer (path) • `AIS-P`—Path alarm indication signal • `RDI-P`—Path remote defect indication • `UNEQ-P`—Path unequipped • `PLM-P`—Path payload (signal) label mismatch • `ES-P`—Errored seconds (near-end STS path) • `SES-P`—Severely errored seconds (near-end STS path) • `UAS-P`—Unavailable seconds (near-end STS path) • `SES-PFE`—Severely errored seconds (far-end STS path) • `UAS-PFE`—Unavailable seconds (far-end STS path)	`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Autonegotiation information`	Information about link autonegotiation.
	• `Negotiation status:`
	  • `Incomplete`—Ethernet interface has the speed or link mode configured.
	  • `No autonegotiation`—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • `Complete`—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• `Link partner status`—`OK` when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• `Link partner:`
	  • `Link mode`—Depending on the capability of the attached Ethernet device, either `Full-duplex` or `Half-duplex`.
	  • `Flow control`—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is `None`. For Gigabit Ethernet interfaces, types are `Symmetric` (link partner supports `PAUSE` on receive and transmit), `Asymmetric` (link partner supports `PAUSE` on transmit), and `Symmetric/Asymmetric` (link partner supports both `PAUSE` on receive and transmit or only `PAUSE` receive).
	  • `Remote fault`—Remote fault information from the link partner—`Failure` indicates a receive link error. `OK` indicates that the link partner is receiving. `Negotiation error` indicates a negotiation error. `Offline` indicates that the link partner is going offline.
	• `Local resolution`—Information from the link partner:
	  • `Flow control`—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are
		`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	`Symmetric` (link partner supports `PAUSE` on receive and transmit), `Asymmetric` (link partner supports `PAUSE` on transmit), and `Symmetric/Asymmetric` (link partner supports both `PAUSE` on receive and transmit or only `PAUSE` receive). • `Remote fault`—Remote fault information. `Link OK` (no error detected on receive), `Offline` (local interface is offline), and `Link Failure` (link error detected on receive).	
`Received path trace`, `Transmitted path trace`	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other routing device manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.	`extensive`
`Packet Forwarding Engine configuration`	Information about the configuration of the Packet Forwarding Engine: • `Destination slot`—FPC slot number.	`extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`CoS information`	Information about the CoS queue for the physical interface. • `CoS transmit queue`—Queue number and its associated user-configured forwarding class name. • `Bandwidth %`—Percentage of bandwidth allocated to the queue. • `Bandwidth bps`—Bandwidth allocated to the queue (in bps). • `Buffer %`—Percentage of buffer space allocated to the queue. • `Buffer usec`—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • `Priority`—Queue priority: `low` or `high`. • `Limit`—Displayed if rate limiting is configured for the queue. Possible values are `none` and `exact`. If `exact` is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If `none` is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	`extensive`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Index number of the logical interface, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP interface index number for the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Flags`	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under Common Output Fields Description.	All levels

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`VLAN-Tag`	Rewrite profile applied to incoming or outgoing frames on the outer (`Out`) VLAN tag or for both the outer and inner (`In`) VLAN tags. • `push`—An outer VLAN tag is pushed in front of the existing VLAN tag. • `pop`—The outer VLAN tag of the incoming frame is removed. • `swap`—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • `push`—An outer VLAN tag is pushed in front of the existing VLAN tag. • `push-push`—Two VLAN tags are pushed in from the incoming frame. • `swap-push`—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • `swap-swap`—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • `pop-swap`—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • `pop-pop`—Both the outer and inner VLAN tags of the incoming frame are removed.	`brief` `detail` `extensive` none
`Demux:`	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	`detail` `extensive` none
`Encapsulation`	Encapsulation on the logical interface.	All levels

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Protocol`	Protocol family. Possible values are described in the "Protocol Field" section under Common Output Fields Description.	`detail extensive` none
`MTU`	Maximum transmission unit size on the logical interface.	`detail extensive` none
`Maximum labels`	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	`detail extensive` none
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the specified interface set. • `Input bytes`, `Output bytes`—Number of bytes received and transmitted on the interface set • `Input packets`, `Output packets`—Number of packets received and transmitted on the interface set.	`detail extensive`
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	`extensive`
`Local statistics`	Number and rate of bytes and packets destined to the routing device.	`extensive`
`Transit statistics`	Number and rate of bytes and packets transiting the switch. **NOTE**: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the `Output bytes` and `Output packets` interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	`extensive`
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Route Table`	Route table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail extensive` none
`Flags`	Information about protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	`detail extensive`
`Donor interface`	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	`detail extensive` none
`Preferred source address`	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	`detail extensive` none
`Input Filters`	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	`detail extensive`
`Output Filters`	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	`detail extensive`
`Mac-Validate Failures`	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	`detail extensive` none
`Addresses, Flags`	Information about the address flags. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	`detail extensive` none
`protocol-family`	Protocol family configured on the logical interface. If the protocol is `inet`, the IP address of the interface is also displayed.	`brief`
`Flags`	Information about address flag (possible values are described in the “Addresses Flags” section under Common Output Fields Description.	`detail extensive` none

Table 639: show interfaces Fast Ethernet Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:44 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: None
```

```
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255
```

show interfaces brief (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 brief
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Logical interface fe-0/0/0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  203.0.113.1/24
```

show interfaces detail (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 detail
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22, Generation: 5391
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:3f:38
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:45 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                   42                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    1                    0 pps
  Active alarms  : None
  Active defects : None
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 105, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
```

```
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
      Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,
        Generation: 136
```

show interfaces extensive (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 extensive
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22, Generation: 5391
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed:
  100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:46 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                   42                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    1                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0               64
    Total packets                            0                1
    Unicast packets                          0                0
    Broadcast packets                        0                1
    Multicast packets                        0                0
    CRC/Align errors                         0                0
```

```
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                       1
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 1, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link partner: Full-duplex, Flow control: None, Remote fault: Ok
    Local resolution:
Packet Forwarding Engine configuration:
    Destination slot: 0
CoS information:
                 Bandwidth           Buffer Priority   Limit
                              %          bps   %         usec
    0 best-effort            95    950000000  95            0      low    none
    3 network-control         5     50000000   5            0      low    none
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 105, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,
        Generation: 136
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces filters

IN THIS SECTION

- Syntax | 20832
- Description | 20832
- Options | 20832
- Additional Information | 20832
- Required Privilege Level | 20833
- Output Fields | 20833
- Sample Output | 20833
- Release Information | 20835

Syntax

```
show interfaces filters
<interface-name>
```

Description

Display all firewall filters that are installed on each interface in a system.

Options

none	Display filter information about all interfaces.
interface-name	(Optional) Display filter information about a particular interface.

Additional Information

For information about how to configure firewall filters, see the Routing Policies, Firewall Filters, and Traffic Policers User Guide. For related operational mode commands, see the CLI Explorer.

Required Privilege Level

view

Output Fields

Table 640 on page 20833 lists the output fields for the `show interfaces filters` command. Output fields are listed in the approximate order in which they appear.

Table 640: show interfaces filters Output Fields

Field Name	Field Description
`Interface`	Name of the interface.
`Admin`	Interface state: `up` or `down`.
`Link`	Link state: `up` or `down`.
`Proto`	Protocol configured on the interface.
`Input Filter`	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
`Output Filter`	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.

Sample Output

show interfaces filters

```
user@host> show interfaces filters
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
```

```
ge-5/0/0        up    up
ge-5/0/0.0      up    up   any                        f-any
                           inet                       f-inet
                           multiservice
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
vt-0/3/0        up    up
at-1/0/0        up    up
at-1/0/0.0      up    up   inet
                           iso
at-1/1/0        up    down
at-1/1/0.0      up    down inet
                           iso
....
```

show interfaces filters (Interface-Name)

```
user@host> show interfaces filters so-2/1/0
Interface       Admin Link Proto Input Filter         Output Filter
so-2/1/0        up    down
so-2/1/0.0      up    down inet  goop                 outfilter
                           iso
                           inet6 v6in                 v6out
```

```
user@host > show interfaces filters ge-3/0/1
Interface       Admin Link Proto Input Filter         Output Filter
ge-3/0/1        up    up
ge-3/0/1.0      up    up   inet  F1-ge-3/0/1.0-in     F2-ge-3/0/1.0-out
                           inet  F3-ge-3/0/1.0-in
```

show interfaces filters (PTX Series Packet Transport Routers)

```
user@host > show interfaces filters em0
Interface       Admin Link Proto Input Filter         Output Filter
```

```
            em0             up    up
            em0.0           up    up   inet
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

show interfaces (Flow Collector)

IN THIS SECTION

- Syntax | 20835
- Description | 20836
- Options | 20836
- Required Privilege Level | 20836
- Output Fields | 20836
- Sample Output | 20842
- Release Information | 20844

Syntax

```
show interfaces cp-fpc/pic/port:channel
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M Series and T Series routers only) Display status information about the specified flow collector interface.

Options

cp-*fpc/pic/port:channel*	Display standard status information about the specified flow collector interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 641 on page 20836 lists the output fields for the `show interfaces` (Flow Collector) command. Output fields are listed in the approximate order in which they appear.

Table 641: Flow Collector Show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical Interface	Name of the physical interface type.	All levels
Link	Status of the link: **up** or **down**.	All levels

Table 641: Flow Collector Show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Enabled	State of the interface type. Possible values are described in the “Enabled Devices” section under *Common Output Fields Description*.	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface.	All levels
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum Transmit Unit (MTU). Size of the largest packet to be transmitted.	All levels
Clocking	Reference clock source of the interface.	All levels
Speed	Network speed on the interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
Link type	Data transmission type.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under *Common Output Fields Description*.	All levels

Table 641: Flow Collector Show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Physical info	Information about the physical interface.	All levels
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	**detail extensive** none
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Media access control (MAC) address of the interface.	**detail extensive** none
Alternate link address	Backup link address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*****hour:minute:second*** **ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive**
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**, **Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets**, **Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**

Table 641: Flow Collector Show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	• **Errors**—Input errors on the interface. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Frames received smaller than the runt threshold. • **Giants**—Frames received larger than the giant threshold. • **Policed Discards**—Frames that the incoming packet match code discarded because the frames did not recognize them or were not of interest. Usually, this field reports protocols that Junos does not support. • **Resource errors**—Sum of transmit drops.	**extensive**
Output errors	• **Carrier transitions** —Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • **Errors**—Sum of outgoing frame terminates and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC RED mechanism. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface	All levels

Table 641: Flow Collector Show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface; values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets, Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**

Table 641: Flow Collector Show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Protocol	Protocol family configured on the logical interface (such as **iso** or **inet6**).	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Route table in which this address exists; for example, **Route table:0** refers to inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces extensive (Flow Collector)

```
user@host> show interfaces extensive cp-5/0/0
Physical interface: cp-5/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 52, Generation: 29
  Type: Flow-collector, Link-level type: Flow-collection, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2005-05-24 16:48:11 PDT (00:12:04 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :           2041661287                    0 bps
   Output bytes  :           3795049544             43816664 bps
   Input  packets:              1365534                    0 pps
   Output packets:              3865644                 3670 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface cp-5/0/0.0 (Index 74) (SNMP ifIndex 53) (Generation 28)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Flow-collection
    Traffic statistics:
     Input  bytes  :           1064651568
     Output bytes  :             37144290
     Input  packets:               711324
     Output packets:               713672
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
```

```
    Output packets:                    0
  Transit statistics:
   Input  bytes  :           1064651568                    0 bps
   Output bytes  :             37144290                    0 bps
   Input  packets:               711324                    0 pps
   Output packets:               713672                    0 pps
  Protocol inet, MTU: 9192, Generation: 39, Route table: 0
    Flags: Receive-options, Receive-TTL-Exceeded
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 192.0.2.2, Local: 192.0.2.1, Broadcast: Unspecified,
      Generation: 40

Logical interface cp-5/0/0.1 (Index 75) (SNMP ifIndex 54) (Generation 29)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Flow-collection
  Traffic statistics:
   Input  bytes  :            976793823
   Output bytes  :             34099481
   Input  packets:               652729
   Output packets:               655127
  Local statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Transit statistics:
   Input  bytes  :            976793823                    0 bps
   Output bytes  :             34099481                    0 bps
   Input  packets:               652729                    0 pps
   Output packets:               655127                    0 pps
  Protocol inet, MTU: 9192, Generation: 40, Route table: 0
    Flags: Receive-options, Receive-TTL-Exceeded
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 198.51.100.2, Local: 198.51.100.1, Broadcast: Unspecified,
      Generation: 42

Logical interface cp-5/0/0.2 (Index 80) (SNMP ifIndex 55) (Generation 30)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Flow-collection
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :           3723079376
   Input  packets:                    0
   Output packets:              2495372
  Local statistics:
```

```
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :           3723079376             43816664 bps
     Input  packets:                    0                    0 pps
     Output packets:              2495372                 3670 pps
    Protocol inet, MTU: 9192, Generation: 41, Route table: 0
      Flags: Receive-options, Receive-TTL-Exceeded
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113.2, Local: 203.0.113.1, Broadcast: Unspecified,
        Generation: 44

  Logical interface cp-5/0/0.16383 (Index 81) (SNMP ifIndex 56) (Generation 31)
...
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces (Flow Monitoring)

IN THIS SECTION

- Syntax | 20845
- Description | 20845
- Options | 20845
- Required Privilege Level | 20845
- Output Fields | 20845
- Sample Output | 20851
- Release Information | 20852

Syntax

```
show interfaces mo-fpc/pic/port:channel
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M Series and T Series routers only) Display status information about the specified flow monitoring interface.

Options

mo-*fpc/pic/port:channel*	Display standard status information about the specified flow monitoring interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 642 on page 20846 lists the output fields for the `show interfaces` (Flow Monitoring) command. Output fields are listed in the approximate order in which they appear.

Table 642: show interfaces Output Fields (Flow Monitoring)

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Link	Status of the link: **up** or **down**.	All levels
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Description	Description and name of the interface.	All levels
Type	Type of interface.	All levels
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum Transmit Unit (MTU). Size of the largest packet to be transmitted.	All levels
Clocking	Reference clock source of the interface.	All levels
Speed	Network speed on the interface.	All levels

Table 642: show interfaces Output Fields (Flow Monitoring) *(Continued)*

Field Name	Field Description	Level of Output
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels
Link type	Data transmission type.	All levels
Link flags	Information about the link. Possible values are described in the "Link Flags" section under *Common Output Fields Description*.	All levels
Physical info	Information about the physical interface.	All levels
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Media access control (MAC) address of the interface.	**detail extensive** none
Alternate link address	Backup link address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour.minute.second timezone*** **(*hour.minute.second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**	**detail extensive**
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 642: show interfaces Output Fields (Flow Monitoring) ***(Continued)***

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**, **Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets**, **Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**
Input errors	• **Errors**—Input errors on the interface. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Frames received smaller than the runt threshold. • **Giants**—Frames received larger than the giant threshold. • **Policed Discards**—Frames that the incoming packet match code discarded because the frames did not recognize them or were not of interest. Usually, this field reports protocols that Junos does not support. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 642: show interfaces Output Fields (Flow Monitoring) *(Continued)*

Field Name	Field Description	Level of Output
Output errors	• **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • **Errors**—Sum of outgoing frame terminates and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC Red mechanism. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface; values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels

Table 642: show interfaces Output Fields (Flow Monitoring) ***(Continued)***

Field Name	Field Description	Level of Output
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets, Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Protocol	Protocol family configured on the logical interface (such as **iso** or **inet6**).	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Route table in which this address exists; for example, **Route table:0** refers to inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	**detail extensive** none

Sample Output

show interfaces extensive (Flow Monitoring)

```
user@host> show interfaces mo-4/0/0     extensive
Physical interface: mo-4/0/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 42, Generation: 28
  Description: monitor pic 2
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: Unlimited,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2005-05-24 16:43:12 PDT (00:17:46 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :            756824218              8328536 bps
   Output bytes  :            872916185              8400160 bps
   Input  packets:               508452                  697 pps
   Output packets:             15577196                18750 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface mo-4/0/0.0 (Index 83) (SNMP ifIndex 43) (Generation 26)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
    Traffic statistics:
     Input  bytes  :            756781796
     Output bytes  :            872255328
     Input  packets:               507233
     Output packets:             15575988
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
```

```
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :            756781796              8328536 bps
     Output bytes  :            872255328              8400160 bps
     Input  packets:               507233                  697 pps
     Output packets:             15575988                18750 pps
    Protocol inet, MTU: Unlimited, Generation: 38, Route table: 0
      Flags: None

  Logical interface mo-4/0/0.16383 (Index 84) (SNMP ifIndex 58) (Generation 27)
...
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces flow-statistics

IN THIS SECTION

- Syntax | 20852
- Description | 20853
- Options | 20853
- Required Privilege Level | 20853
- Output Fields | 20853
- Sample Output | 20857
- Release Information | 20858

Syntax

```
show interfaces flow-statistics                    <interface-name>
```

Description

Display interfaces flow statistics.

Options

Interface-name —(Optional) Display flow statistics about the specified interface. Following is a list of typical interface names. Replace *pim* with the PIM slot and *port* with the port number. For a complete list, see the *Interface Naming Conventions*.

- at-*pim*/0/*port*—ATM-over-ADSL or ATM-over-SHDSL interface.
- br-*pim*/0/*port*—Basic Rate Interface for establishing ISDN connections.
- ce1-*pim*/0/*port*—Channelized E1 interface.
- ct1-*pim*/0/*port*—Channelized T1 interface.
- dl0—Dialer Interface for initiating ISDN and USB modem connections.
- e1-*pim*/0/*port*—E1 interface.
- e3-*pim*/0/*port*—E3 interface.
- fe-*pim*/0/ *port*—Fast Ethernet interface.
- ge-*pim*/0/*port*—Gigabit Ethernet interface.
- se-*pim*/0/*port*—Serial interface.
- t1-*pim*/0/*port*—T1 (also called DS1) interface.
- t3-*pim*/0/ *port*—T3 (also called DS3) interface.
- wx-*slot*/0/0—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

Required Privilege Level

view

Output Fields

Table 643 on page 20854 lists the output fields for the `show interfaces flow-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 643: show interfaces flow-statistics Output Fields

Field Name	Field Description
`Traffic statistics`	Number of packets and bytes transmitted and received on the physical interface.
`Local statistics`	Number of packets and bytes transmitted and received on the physical interface.
`Transit statistics`	Number of packets and bytes transiting the physical interface.
`Flow input statistics`	Statistics on packets received by flow module.
`Flow output statistics`	Statistics on packets sent by flow module.
`Flow error statistics`	Packet drop statistics for the flow module. For further details, see Table 644 on page 20854.

Table 644: Flow Error Statistics (Packet Drop Statistics for the Flow Module)

Error	Error Description
Screen:	
Address spoofing	The packet was dropped when the screen module detected address spoofing.
Syn-attack protection	The packet was dropped because of SYN attack protection or SYN cookie protection.
VPN:	
Authentication failed	The packet was dropped because the IPsec Encapsulating Security Payload (ESP) or Authentication Header (AH) authentication failed.

No SA for incoming SPI	The packet was dropped because the incoming IPsec packet's security parameter index (SPI) does not match any known SPI.
Security association not active	The packet was dropped because an IPsec packet was received for an inactive SA.
NAT:	
Incoming NAT errors	The source NAT rule search failed, an invalid source NAT binding was found, or the NAT allocation failed.
Multiple incoming NAT	Sometimes packets are looped through the system more than once; if source NAT is specified more than once, the packet will be dropped.
Auth:	
Multiple user authentications	Sometimes packets are looped through the system more than once. Each time a packet passes through the system, that packet must be permitted by a policy. If the packet matches more than one policy that specifies user authentication, then it will be dropped.
User authentication errors	Packet was dropped because policy requires authentication; however: • Only Telnet, FTP, and HTTP traffic can be authenticated. • The corresponding authentication entry could not be found, if web-auth is specified. • The maximum number of authenticated sessions per user was exceeded.
Flow:	
No one interested in self packets	This counter is incremented for one of the following reasons: • The outbound interface is a self interface, but the packet is not marked as a to-self packet and the destination address is in a source NAT pool. • No service is interested in the to-self packet • When a zone has ident-reset service enabled, the TCP RST to IDENT request for port 113 is sent back and this counter is incremented.

No minor session	The packet was dropped because no minor sessions are available and a minor session was requested. Minor sessions are allocated for storing additional TCP state information.
No more sessions	The packet was dropped because there were no more free sessions available.
No route present	The packet was dropped because a valid route was not available to forward the packet. For new sessions, the counter is incremented for one of the following reasons: • No valid route was found to forward the packet. • A discard or reject route was found. • The route could not be added due to lack of memory. • The reverse path forwarding check failed for an incoming multicast packet. For existing sessions, the prior route was changed or deleted, or a more specific route was added. The session is rerouted, and this reroute could fail because: • A new route could not be found; either the previous route was removed, or the route was changed to discard or reject. • Multiple packets may concurrently force rerouting to occur, and only one packet can successfully complete the rerouting process. Other packets will be dropped. • The route table was locked for updates by the Routing Engine. Packets that match a new session are retried, whereas packets that match an existing session are not.
No tunnel found	The packet was dropped because a valid tunnel could not be found
No session for a gate	This counter is incremented when a packet is destined for an ALG, and the ALG decides to drop this packet.
No zone or NULL zone binding	The packet was dropped because its incoming interface was not bound to any zone.

Policy denied	The error counter is incremented for one of the following reasons: • Source and/or destination NAT has occurred and policy says to drop the packet. • Policy specifies user authentication, which failed. • Policy was configured to deny this packet.
TCP sequence number out of window	A TCP packet with a sequence number failed the TCP sequence number check that was received.
Counters Not Currently in Use	
No parent for a gate	-
Invalid zone received packet	-
No NAT gate	-

Sample Output

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp igmp ldp msdp nhrp ospf pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http https ike
    netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                    0
      ICMP packets :                    0
      VPN packets :                     2564
```

```
    Bytes permitted by policy :        3478
    Connections established :          1
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        16994
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
    No minor session:                  0
    No more sessions:                  0
    No NAT gate:                       0
    No route present:                  0
    No SA for incoming SPI:            0
    No tunnel found:                   0
    No session for a gate:             0
    No zone or NULL zone binding       0
    Policy denied:                     0
    Security association not active:   0
    TCP sequence number out of window: 0
    Syn-attack protection:             0
    User authentication errors:        0
  Protocol inet, MTU: 1500
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination:     203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```

Release Information

Command introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

Understanding Interfaces

show interfaces forwarding-class-counters

IN THIS SECTION

- Syntax | 20859
- Description | 20859
- Options | 20859
- Additional Information | 20859
- Required Privilege Level | 20860
- Output Fields | 20860
- Sample Output | 20861
- Release Information | 20865

Syntax

```
show interfaces forwarding-class-counters interface-name <comprehensive>
```

Description

Display interface accounting information by forwarding class for IPv4, IPv6, MPLS, Layer 2, and Other traffic.

Options

comprehensive (Optional) Display forwarding-class-counters per traffic family for all logical interfaces under the physical interface along with other quality-of-service information.

Additional Information

For physical interface-level statistics, if none of the logical interfaces have any of the traffic families configured on them, the forwarding class statistics for that family are still displayed with a value of 0.

For physical interface-level statistics, in case of Layer 2 families such as **ccc**, **tcc**, or **vpls**, the **Layer2** keyword is displayed because it is possible that different Layer 2 families are configured on the logical interface.

For logical interface-level statistics, the output displays statistics only for families that are configured on that logical interface. The statistics under **Other** family are still displayed because these are packets that are not classified as belonging to any family.

In the case of Layer 2 families such as **ccc**, **tcc**, or **vpls** configured on the logical interface, the actual family name is displayed in the output.

Statistics include input and output byte and packets and corresponding rates.

Required Privilege Level

view

Output Fields

Table 645 on page 20860 lists the output fields for the `show interfaces forwarding-class-counters` command. Output fields are listed in the approximate order in which they appear.

Table 645: show interfaces forwarding-class-counters Output Fields

Field Name	Field Description
`Input bytes`	A count of received bytes that match the forwarding class.
`Output bytes`	A count of transmitted bytes that match the forwarding class.
`Input packets`	A count of received packets that match the forwarding class.
`Output packets`	A count of transmitted packets that match the forwarding class.

Sample Output

show interfaces forwarding-class-counters interface-name

```
user@host> show interfaces forwarding-class-counters ge-4/2/1

user@host> show interfaces forwarding-class-counters ge-4/2/1
Physical interface ge-4/2/1 (Index 228) (SNMP ifIndex 870)
             Aggregate Forwarding-class statistics :
       Forwarding-class statistics : best-effort
        Input   bytes   :                     0        0 bps
        output  bytes   :                     0        0 bps
        Input   packets :                     0        0 pps
        output  packets :                     0     0 pps
       Forwarding-class statistics : network-control
        Input   bytes   :                     0        0 bps
        output  bytes   :                     0        0 bps
        Input   packets :                     0        0 pps
        output  packets :                     0     0 pps

    IPv4 Forwarding-class statistics :
       Forwarding-class statistics : best-effort
        Input   bytes   :                     0        0 bps
        output  bytes   :                     0        0 bps
        Input   packets :                     0        0 pps
        output  packets :                     0     0 pps
       Forwarding-class statistics : expedited-forwarding
        Input   bytes   :                     0        0 bps
        output  bytes   :                     0        0 bps
        Input   packets :                     0        0 pps
        output  packets :                     0     0 pps
       Forwarding-class statistics : assured-forwarding
        Input   bytes   :                     0        0 bps
        output  bytes   :                     0        0 bps
        Input   packets :                     0        0 pps
        output  packets :                     0     0 pps
       Forwarding-class statistics : network-control
        Input   bytes   :                     0        0 bps
        output  bytes   :                     0        0 bps
        Input   packets :                     0        0 pps
        output  packets :                     0     0 pps
```

```
    IPv6 Forwarding-class statistics :
        Forwarding-class statistics : best-effort
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps
        Forwarding-class statistics : expedited-forwarding
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps
        Forwarding-class statistics : assured-forwarding
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps
        Forwarding-class statistics : network-control
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps

    MPLS Forwarding-class statistics :
        Forwarding-class statistics : best-effort
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps
        Forwarding-class statistics : expedited-forwarding
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps
        Forwarding-class statistics : assured-forwarding
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps
        Forwarding-class statistics : network-control
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps
```

```
Layer2 Forwarding-class statistics
    Forwarding-class statistics : best-effort
     Input   bytes   :                    0        0 bps
     output  bytes   :                    0        0 bps
     Input   packets :                    0        0 pps
     output  packets :                    0     0 pps
    Forwarding-class statistics : expedited-forwarding
     Input   bytes   :                    0        0 bps
     output  bytes   :                    0        0 bps
     Input   packets :                    0        0 pps
     output  packets :                    0     0 pps
    Forwarding-class statistics : assured-forwarding
     Input   bytes   :                    0        0 bps
     output  bytes   :                    0        0 bps
     Input   packets :                    0        0 pps
     output  packets :                    0     0 pps
    Forwarding-class statistics : network-control
     Input   bytes   :                    0        0 bps
     output  bytes   :                    0        0 bps
     Input   packets :                    0        0 pps
     output  packets :                    0     0 pps

Other Forwarding-class statistics :
    Forwarding-class statistics : best-effort
     Input   bytes   :                    0        0 bps
     output  bytes   :                    0        0 bps
     Input   packets :                    0        0 pps
     output  packets :                    0     0 pps
    Forwarding-class statistics : expedited-forwarding
     Input   bytes   :                    0        0 bps
     output  bytes   :                    0        0 bps
     Input   packets :                    0        0 pps
     output  packets :                    0     0 pps
    Forwarding-class statistics : assured-forwarding
     Input   bytes   :                    0        0 bps
     output  bytes   :                    0        0 bps
     Input   packets :                    0        0 pps
     output  packets :                    0     0 pps
    Forwarding-class statistics : network-control
     Input   bytes   :                    0        0 bps
     output  bytes   :                    0        0 bps
     Input   packets :                    0        0 pps
```

```
        output  packets :                     0     0 pps


Logical interface ge-4/2/1.0 (Index 347) (SNMP ifIndex 1032)
   Forwarding-class accounting parameters :
 Aggregate Forwarding-class statistics :
            Forwarding-class statistics : best-effort
      Input   bytes   :                     0        0 bps
      output  bytes   :                     0        0 bps
      Input   packets :                     0        0 pps
      output  packets :                     0     0 pps
     Forwarding-class statistics : expedited-forwarding
      Input   bytes   :                     0        0 bps
      output  bytes   :                     0        0 bps
      Input   packets :                     0        0 pps
      output  packets :                     0     0 pps
     Forwarding-class statistics : assured-forwarding
      Input   bytes   :                     0        0 bps
      output  bytes   :                     0        0 bps
      Input   packets :                     0        0 pps
      output  packets :                     0     0 pps
     Forwarding-class statistics : network-control
      Input   bytes   :                     0        0 bps
      output  bytes   :                     0        0 bps
      Input   packets :                     0        0 pps
      output  packets :                     0     0 pps

 ccc Forwarding-class statistics :
     Forwarding-class statistics : best-effort
      Input   bytes   :                     0        0 bps
      output  bytes   :                     0        0 bps
      Input   packets :                     0        0 pps
      output  packets :                     0     0 pps
     Forwarding-class statistics : expedited-forwarding
      Input   bytes   :                     0        0 bps
      output  bytes   :                     0        0 bps
      Input   packets :                     0        0 pps
      output  packets :                     0     0 pps
     Forwarding-class statistics : assured-forwarding
      Input   bytes   :                     0        0 bps
      output  bytes   :                     0        0 bps
      Input   packets :                     0        0 pps
      output  packets :                     0     0 pps
```

```
        Forwarding-class statistics : network-control
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps

    Other Forwarding-class statistics :
        Forwarding-class statistics : best-effort
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps
        Forwarding-class statistics : expedited-forwarding
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps
        Forwarding-class statistics : assured-forwarding
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps
        Forwarding-class statistics : network-control
         Input   bytes   :                    0        0 bps
         output  bytes   :                    0        0 bps
         Input   packets :                    0        0 pps
         output  packets :                    0     0 pps
```

command-name

Release Information

Command introduced in Junos OS 14.1.

RELATED DOCUMENTATION

forwarding-class-accounting | **7413**

Class of Service User Guide (Routers and EX9200 Switches)

show interfaces fti

IN THIS SECTION

- Syntax | 20866
- Description | 20866
- Options | 20866
- Required Privilege Level | 20866
- Output Fields | 20867
- Sample Output | 20874
- Release Information | 20880

Syntax

```
show interfaces fti number
<detail | extensive | terse | statistics >
```

Description

Display status information about the specified FTIs.

Options

fti *number*	Display standard information about the specified FTIs.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
statistics	(Optional) Display FTI statistics.

Required Privilege Level

view

Output Fields

Table 646 on page 20867 lists the output fields for the `show interfaces` (Flexible Tunnel Interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 646: FTI show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical interface`	Name of the physical interface and state of the interface.	All levels
`Enabled`	State of the physical interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
`Interface index`	Index number of the physical interface, which reflects its initialization sequence.	All levels
`SNMP ifIndex`	SNMP index number for the physical interface.	All levels
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Link-level type`	Encapsulation being used on the physical interface. Possible values: Flexible- tunnel-Interfaces	All levels
`MTU`	Maximum transmission unit size on the physical interface.	All levels
`Clocking`	Reference clock source of the interface.	All levels
`Speed`	Speed at which the interface is running.	All levels
`Device flags`	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
`Interface flags`	Information about the interface. Possible values are described in the "Interfaces Flags" section under *Common Output Fields Description*.	All levels

Table 646: FTI show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Link-type`	Data transmission type.	All levels
`Link-flags`	Information about the link. Possible values are described in the “Link Flags” section under *Common Output Fields Description*.	All levels
`Last flapped`	Date, time, and how long ago the interface went from down to up or from up to down. The format is `Last flapped:` *`year-month-day hours:minutes:seconds timezone`* (*`hours:minutes:seconds`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	All levels
`Input packets`	Number of packets received on the physical interface.	All levels
`Output packets`	Number of packets transmitted on the physical interface.	All levels
`Physical-info`	Information about the physical interface.	`detail extensive`
`Hold-times`	Current interface hold-time up and hold-time down. Value is in milliseconds.	`detail extensive`

Table 646: FTI show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Damping`	Provides a smoothing of the up and down transitions (flapping) on an interface. • `halflife`— Decay half-life. The number of seconds after which the accumulated interface penalty counter is reduced by half if the interface remains stable. • `max-suppress`— Maximum hold-down time. The maximum number of seconds that an interface can be suppressed irrespective of how unstable the interface has been. • `reuse`— Reuse threshold. When the accumulated interface penalty counter falls below this number, the interface is no longer suppressed. • `suppress`— Cutoff (suppression) threshold. When the accumulated interface penalty counter exceeds this number, the interface is suppressed. • `state`— Interface damping state. If damping is enabled on an interface, it is suppressed during interface flaps that match the configured damping parameters.	`detail extensive`
`Current address`	Configured MAC address.	`detail extensive`
`Hardware address`	Hardware MAC address.	`detail extensive`
`Alternate-link-address`	Backup link address.	`detail extensive`
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`, `statistics`

Table 646: FTI show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes and rate, in bps, at which bytes are received on the interface. • `Output bytes`—Number of bytes and rate, in bps, at which bytes are transmitted on the interface. • `Input packets`—Number of packets and rate, in pps, at which packets are received on the interface. • `Output packets`—Number of packets and rate, in pps, at which packets are transmitted on the interface.	`detail extensive`
`Input errors`	Input errors on the interface: • `Errors`—Sum of incoming frame terminates and frame checksum (FCS) errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's random early detection (RED) mechanism. • `Framing errors`—Number of packets received with an invalid FCS. • `Runts`—Number of frames received that are smaller than the runt threshold. • `Giants`—Number of frames received that are larger than the giant threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos OS does not handle. • `Resource errors`—Sum of transmit drops.	`detail extensive`

Table 646: FTI show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions` —Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), then the cable, the far-end system, or the PIC is malfunctioning. • `Errors`—Sum of the outgoing frame terminates and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `MTU errors`—Number of packets whose size exceeded the MTU of the interface. • `Resource errors`—Sum of transmit drops.	`detail extensive`
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Index number of the logical interface (which reflects its initialization sequence).	`detail extensive` none

Table 646: FTI show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`SNMP ifIndex`	SNMP interface index number of the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Flags`	Information about the logical interface. Possible values are described in the "Logical Interface Flags Field" section under *Common Output Fields Description*.	All levels
`Encapsulation`	Encapsulation on the logical interface.	All levels
`Destination UDP Port`	Value to write to the destination-udp-port field.	All levels
`Source UDP Port Range`	Value to write to the source-udp-port-range field. • **min** ***min-port-number***: 1 through 65,535 • **max** ***max-port-number***: 1 through 65,535	All levels
`VNI`	Value to write to the `vni` attribute	All levels
`Source address`	IP address of the source side of the connection.	All levels
`Transit Statistics`	Statistics for traffic transiting the router. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`
`Protocol inet, MTU`	Protocol family configured on the logical interface. If the protocol is `inet`, the IP address of the interface is also displayed. MTU is the maximum transmission unit size on the logical interface	All levels

Table 646: FTI show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Max nh cache`	Maximum interface neighbor discovery nexthop cache size.	All levels
`New hold nh limit`	Maximum number of new unresolved nexthops.	All levels
`Curr nh cnts`	Current number of resolved next-hop addresses.	All levels
`Curr new hold cnt`	Current number of unresolved next-hop addresses.	All levels
`NH drop cnt`	Number of neighbor discovery requests not serviced	All levels
`Flags`	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	
`Route Table`	Routing table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail extensive`
`Flags`	Information about protocol family flags. Possible values are described in the “Family Flags Field” section under *Common Output Fields Description*.	All levels
`Addresses, Flags`	Information about address flags. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	All levels
`Destination address`	IP address of the remote side of the connection.	All levels
`Local`	IP address of the logical interface.	All levels
`Broadcast`	Broadcast address of the logical interface.	`detail extensive` none

Sample Output

show interfaces (Flexible Tunnel Interfaces on MX Series)

```
user@host> show interfaces fti0
Physical interface: fti0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 504
  Type: FTI, Link-level type: Flexible-tunnel-Interface, MTU: Unlimited, Speed: Unlimited
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface fti0.0 (Index 340) (SNMP ifIndex 581)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: VXLAN-GPEv4
    Destination UDP port: 4789, VNI: 1000, Source address: 105.5.5.5, Destination address:
106.6.6.6
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: Unlimited
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 103.3.3/24, Local: 103.3.3.5, Broadcast: 103.3.3.255
```

show interfaces (When source UDP port range for VXLAN encapsulation is configured at the Flexible Tunnel Interfaces on MX Series)

```
user@host> show interfaces fti0  Physical interface: fti0, Enabled, Physical link is Up
  Interface index: 137, SNMP ifIndex: 586
  Type: FTI, Link-level type: Flexible-tunnel-Interface, MTU: Unlimited, Speed: Unlimited
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
```

```
    Output packets: 0

  Logical interface fti0.0 (Index 329) (SNMP ifIndex 1678)
    Flags: Hardware-Down Up Point-To-Point SNMP-Traps Encapsulation: VXLAN-GPEv4
    Destination UDP port: 1000, Source UDP port range: [1000 - 2000],  VNI: 1000, Source
address: 13.13.13.13,     Destination address: 14.14.14.14
    Input packets : 0
    Output packets: 0
```

show interfaces detail (Flexible Tunnel Interfaces on MX Series)

```
user@host> show interfaces fti0 detail
Physical interface: fti0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 504, Generation: 139
  Type: FTI, Link-level type: Flexible-tunnel-Interface, MTU: Unlimited, Clocking: Unspecified,
  Speed: Unlimited
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
   IPv6 transit statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0

  Logical interface fti0.0 (Index 340) (SNMP ifIndex 581) (Generation 149)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: VXLAN-GPEv4
```

```
    Destination UDP port: 4789, VNI: 1000, Source address: 10.5.5.5, Destination address:
10.6.6.6
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
    Protocol inet, MTU: Unlimited
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
    Generation: 176, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.3.3/24, Local: 10.3.3.5, Broadcast: 10.3.3.255, Generation: 149

  Logical interface fti0.1 (Index 341) (SNMP ifIndex 582) (Generation 150)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: VXLAN-GPEv4
    Destination UDP port: 4789, VNI: 1000, Source address: 10.8.8.8, Destination address:
10.7.7.7
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
```

```
    Protocol inet, MTU: Unlimited
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
    Generation: 177, Route table: 9
      Flags: Sendbcast-pkt-to-re, Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255, Generation: 151
```

show interfaces extensive (Flexible Tunnel Interfaces on MX Series)

```
user@host> show interfaces fti0 extensive
Physical interface: fti0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 504, Generation: 139
  Type: FTI, Link-level type: Flexible-tunnel-Interface, MTU: Unlimited, Clocking: Unspecified,
  Speed: Unlimited
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
   IPv6 transit statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
```

```
  Logical interface fti0.0 (Index 340) (SNMP ifIndex 581) (Generation 149)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: VXLAN-GPEv4
    Destination UDP port: 4789, VNI: 1000, Source address: 10.5.5.5, Destination address:
10.6.6.6
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
    Protocol inet, MTU: Unlimited
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
    Generation: 176, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.3.3/24, Local: 10.3.3.5, Broadcast: 10.3.3.255, Generation: 149
```

show interfaces terse (Flexible Tunnel Interfaces on MX Series)

```
user@host> show interfaces fti0 terse
Interface               Admin    Link      Proto    Local         Remote
        fti0                    up        up
        fti0.0                  up        up       inet     10.3.3.5/24
```

show interfaces statistics (Flexible Tunnel Interfaces on MX Series)

```
user@host> show interfaces fti0 statistics
Physical interface: fti0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 504
```

```
  Type: FTI, Link-level type: Flexible-tunnel-Interface, MTU: Unlimited, Speed: Unlimited
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 0
  Input errors: 0, Output errors: 0

  Logical interface fti0.0 (Index 340) (SNMP ifIndex 581)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: VXLAN-GPEv4
    Destination UDP port: 4789, VNI: 1000, Source address: 10.5.5.5, Destination address:
10.6.6.6
    Input packets : 3
    Output packets: 3
    Protocol inet, MTU: Unlimited
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.3.3/24, Local: 10.3.3.5, Broadcast: 10.3.3.255
```

show interfaces (Flexible Tunnel Interfaces on PTX and QFX Series)

```
user@host> show interfaces fti0 detail
Physical interface: ft0, Enabled, Physical link is Up
  Interface index: 193, SNMP ifIndex: 576, Generation: 196
  Type: FT, Link-level type: FT, MTU: Unlimited, Speed: Unlimited
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
```

```
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Logical interface ft0.0 (Index 337) (SNMP ifIndex 583) (Generation 146)
    Flags: Hardware-Down Up Point-To-Point SNMP-Traps 0x4000
    Encapsulation: UDP
    Source: 10.1.1.1 Destination: 10.2.2.2 Destination-Port: 4789
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Flexible Tunnel Interfaces Overview

Configuring Flexible Tunnel Interfaces

Example: Configuring Flexible Tunnel Interfaces on MX Series Routers

vni(Interfaces) | **30911**

show interfaces ge

IN THIS SECTION

- Syntax (EX Series) | 20881
- Syntax (QFX Series) | 20881
- Description | 20882
- Options | 20882
- Required Privilege Level | 20883
- Output Fields | 20883
- Sample Output | 20894
- Release Information | 20900

Syntax (EX Series)

```
show interfaces ge-fpc/pic/port
<brief | detail | extensive | terse>
<media>
<statistics>
```

Syntax (QFX Series)

```
show interfaces device-name:type-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.

> **NOTE**: You must have a transceiver plugged into an SFP or SFP+ port before information about the interface can be displayed.

Options

ge-*fpc/pic/port*	(EX Series) Display standard information about the specified Gigabit Ethernet interface.
brief | detail | extensive | terse	(Optional) (EX Series) Display the specified level of output.
media	(Optional) (EX Series) Display media-specific information about network interfaces.
statistics	(Optional) (EX Series) Display static interface statistics.
brief | detail | extensive | terse	(Optional) (QFX Series) Display the specified level of output.
device-name*:*type-fpc/pic/port	(QFX Series) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.
descriptions	(Optional) (QFX Series) Display interface description strings.
media	(Optional) (QFX Series) Display media-specific information about network interfaces.
routing instance (all | *instance-name*)	(Optional) (QFX Series) Display the name of an individual routing-instance or display all routing-instances.
snmp-index *snmp-index*	(Optional) (QFX Series) Display information for the specified SNMP index of the interface.
statistics	(Optional) (QFX Series) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 647 on page 20883 lists the output fields for the **show interfaces ge-** command. Output fields are listed in the approximate order in which they appear.

Table 647: show interfaces ge- Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface: **Enabled** or **Disabled**.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Description	Optional user-specified description.	**brief detail extensive**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface. Default is 1514.	All levels
Speed	Speed of the interface: Auto if autonegotiation of speed is enabled; speed in megabits per second if the interface speed is explicitly configured.	All levels

Table 647: show interfaces ge- Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Duplex	Link mode of the interface: Auto if autonegotiation of link mode is enabled; Full-Duplex or Half-Duplex if the link mode is explicitly configured.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
IEEE 802.3az Energy Efficient Ethernet	IEEE 802.3az Energy Efficient Ethernet status: **Enabled** or **Disabled** (appears only for EEE-capable Base-T copper Ethernet interfaces).	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the link.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**

Table 647: show interfaces ge- Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Current address	Configured MAC address.	**detail extensive** none
Hardware address	MAC address of the hardware.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).**	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled on the switch.	**detail extensive**

Table 647: show interfaces ge- Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame terminated and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 sanity checks of the headers. For example, a frame with less than 20 bytes of available IP header is discarded. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 647: show interfaces ge- Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame terminated and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**

Table 647: show interfaces ge- Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain time, it is promoted to an alarm. Based on the switch configuration, a defect can activate the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none

Table 647: show interfaces ge- Output Fields *(Continued)*

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem. • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed 1518 octets. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**

Table 647: show interfaces ge- Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Filter Statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem.	**extensive**

Table 647: show interfaces ge- Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Autonegotiation information	Information about link autonegotiation:
	• **Negotiation status:**
	  • **Complete**—The autonegotiation process between the local and remote Ethernet interfaces was successful.
	  • **Incomplete**—Remote Ethernet interface has the speed or link mode configured or does not perform autonegotiation.
	  • **No autonegotiation**—Local Ethernet interface has autonegotiation disabled and the link mode and speed are manually configured.
	• **Link partner**—Information from the link partner:
	  • **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**. If the link mode of the remote device cannot be determined, the value is **Unknown**.
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, the types are: **Symmetric** (link partner supports **PAUSE** on receive and transmit); **Asymmetric** (link partner supports **PAUSE** on transmit); and **Symmetric/Asymmetric** (link partner supports **PAUSE** on both receive and transmit or **PAUSE** only on receive).
	  • **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	  • **Link partner speed**—Speed of the link partner.
	• **Local resolution**—Resolution of the autonegotiation process on the local interface:
	  • **Flow control**—Type of flow control that is used by the local interface. For Gigabit Ethernet interfaces, the types are: **Symmetric** (link partner supports **PAUSE** on receive and
		extensive

Table 647: show interfaces ge- Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	transmit); **Asymmetric** (link partner supports **PAUSE** on transmit); and **Symmetric/Asymmetric** (link partner supports **PAUSE** on both receive and transmit or **PAUSE** only on receive). • **Link mode**—Link mode of local interface: either **Full-duplex** or **Half-duplex**. Displayed when **Negotiation status** is **Incomplete**. • **Local link speed**—Speed of the local interface. Displayed when **Negotiation status** is **Incomplete**. • **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).	

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine:
	• **Destination slot**—FPC slot number:
	  • On standalone switches with built-in interfaces, the slot number refers to the switch itself and is always 0.
	  • On Virtual Chassis composed of switches with built-in interfaces, the slot number refers to the member ID of the switch.
	  • On switches with line cards or on Virtual Chassis composed of switches with line cards, the slot number refers to the line card slot number on the switch or Virtual Chassis.
		extensive

Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Table 647: show interfaces ge- Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface. **NOTE**: For logical interfaces on EX Series switches, the traffic statistics fields in **show interfaces** commands show only control traffic; the traffic statistics do not include data traffic.	**detail extensive**
IPv6 transit statistics	EX Series switches do not support the collection and reporting of IPv6 transit statistics.	**extensive**
Local statistics	Number and rate of bytes and packets destined to and from the switch.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch.	**extensive**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table **inet.0**.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface.	**detail extensive**
Output Filters	Names of any output filters applied to this interface.	**detail extensive**

Table 647: show interfaces ge- Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Information about protocol family flags. If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled.	**detail extensive**
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about the address flags.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces ge-0/0/0

```
user@switch> show interfaces ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Interface index: 129, SNMP ifIndex: 21
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
```

```
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:19:e2:50:3f:41, Hardware address: 00:19:e2:50:3f:41
  Last flapped   : 2008-01-16 11:40:53 UTC (4d 02:30 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 22)
    Flags: SNMP-Traps
    Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch
      Flags: None
```

show interfaces ge-0/0/0 brief

```
user@switch> show interfaces ge-0/0/0 brief
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Description: voice priority and tcp and icmp traffic rate-limiting filter at i
ngress port
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None

  Logical interface ge-0/0/0.0
    Flags: Device-Down SNMP-Traps Encapsulation: ENET2
    eth-switch
```

show interfaces ge-0/0/0 brief (with EEE Enabled on the EEE-capable Base-T copper Ethernet interfaces)

```
user@switch> show interfaces ge-0/0/0 brief
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online,
  IEEE 802.3az Energy Efficient Ethernet: Enabled, NO LPI
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
```

show interfaces ge-0/0/0 detail

```
user@switch> show interfaces ge-0/0/0 detail

Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 193, SNMP ifIndex: 206, Generation: 196
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
  Last flapped   : 2009-05-05 06:03:05 UTC (00:22:13 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
```

```
    Output packets:                  0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 assured-forw                   0                    0                    0
    5 expedited-fo                   0                    0                    0
    7 network-cont                   0                    0                    0
  Active alarms  : None
  Active defects : None

  Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
    Flags: SNMP-Traps Encapsulation: ENET2
    Bandwidth: 0
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol eth-switch, Generation: 146, Route table: 0
      Flags: Is-Primary
      Input Filters: f1,
      Output Filters: f2,,,,
```

show interfaces ge-0/0/4 extensive

```
user@switch> show interfaces ge-0/0/4 extensive

Physical interface: ge-0/0/4, Enabled, Physical link is Up
  Interface index: 165, SNMP ifIndex: 152, Generation: 168
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
```

```
Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:1f:12:33:65:44, Hardware address: 00:1f:12:33:65:44
Last flapped   : 2008-09-17 11:02:25 UTC (16:32:54 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :              2989761                  984 bps
 Input  packets:                    0                    0 pps
 Output packets:                24307                    1 pps
 IPv6 transit statistics:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                    0                    0                    0
  1 assured-forw                   0                    0                    0
  5 expedited-fo                   0                    0                    0
  7 network-cont                   0                24307                    0
Active alarms  : None
Active defects : None
MAC statistics:                      Receive         Transmit
  Total octets                             0          2989761
  Total packets                            0            24307
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0            24307
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
```

```
  MAC pause frames                        0                0
  Oversized frames                        0
  Jabber frames                           0
  Fragment frames                         0
  Code violations                         0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
      Link mode: Full-duplex, Flow control: None, Remote fault: OK,
      Link partner Speed: 1000 Mbps
  Local resolution:
      Flow control: None, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0
  Direction : Output
  CoS transmit queue               Bandwidth               Buffer Priority   Limit
                            %            bps     %           usec
  0 best-effort            95      950000000    95             NA      low    none
  7 network-control         5       50000000     5             NA      low    none

Logical interface ge-0/0/4.0 (Index 82) (SNMP ifIndex 184) (Generation 147)
  Flags: SNMP-Traps Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :              4107883
   Input  packets:                    0
   Output packets:                24307
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Local statistics:
   Input  bytes  :                   0
   Output bytes  :              4107883
   Input  packets:                    0
   Output packets:                24307
  Transit statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
   IPv6 transit statistics:
```

```
    Input  bytes  :                    0
    Output bytes  :                    0
    Input  packets:                    0
    Output packets:                    0
  Protocol eth-switch, Generation: 159, Route table: 0
    Flags: None
    Input Filters: f2,
    Output Filters: f1,,,,
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Monitoring Interface Status and Traffic
- Troubleshooting Network Interfaces on EX3200 Switches
- *Troubleshooting Network Interfaces on EX4200 Switches*
- *Troubleshooting an Aggregated Ethernet Interface*
- Junos OS Ethernet Interfaces Configuration Guide
- *Monitoring Interface Status and Traffic*
- *Troubleshooting Network Interfaces*
- *Troubleshooting an Aggregated Ethernet Interface*
- Junos OS Network Interfaces Library for Routing Devices

show interfaces (GRE)

IN THIS SECTION

- Syntax | **20901**
- Description | **20901**
- Options | **20901**
- Required Privilege Level | **20902**

- Output Fields | 20902
- Sample Output | 20907
- Release Information | 20913

Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified generic routing encapsulation (GRE) interface.

Options

interface-type	On M Series and T Series routers and EX Series switches, the interface type is **gr-*fpc/pic/port***.
brief | detail | extensive | terse	(Optional) Display the specified output level of interface information.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

NOTE: You can configure generic routing encapsulation (GRE) interfaces (gre-x/y/z) only for GMPLS control channels. GRE interfaces are not supported or configurable for other

applications. For more information about GMPLS, see the MPLS Applications User Guide.

Required Privilege Level

view

Output Fields

Table 648 on page 20902 lists the output fields for the `show interfaces` (GRE) command. Output fields are listed in the approximate order in which they appear.

Table 648: GRE show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface.	All levels
Link-level type	Encapsulation used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels

Table 648: GRE show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Speed	Speed at which the interface is running.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Device Flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Interface Flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	The number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none

Table 648: GRE show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
Flags	Information about the logical interface. Possible values listed in the "Logical Interface Flags" section under *Common Output Fields Description*. describe general information about the logical interface. GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field: • **Reassemble-Pkts**—If the **Flags** field includes this string, the GRE tunnel is configured to reassemble tunnel packets that were fragmented after tunnel encapsulation.	All levels
IP-Header	IP header of the logical interface. If the `tunnel key` statement is configured, this information is included in the **IP Header** entry. GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field: • **df**—If the **IP-Header** field includes this string immediately following the 16 bits of identification information (that is, if **:df:** displays after the twelfth byte), the GRE tunnel is configured to allow fragmentation of GRE packets after encapsulation.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
`L2 Routing Instance`	Name of the Layer 2 routing instance associated with the GRE interface.	All levels
`L3 Routing Instance`	Name of the Layer 3 routing instance associated with the GRE interface.	All levels

Table 648: GRE show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Copy-tos-to-outer-ip-header`	Status of type of service (ToS) bits in the GRE packet header: • `On`—ToS bits were copied from the payload packet header into the header of the IP packet sent through the GRE tunnel. • `Off`—ToS bits were not copied from the payload packet header and are set to 0 in the GRE packet header. **NOTE**: EX Series switches do not support copying ToS bits to the encapsulated packet, so the value of this field is always `Off` in switch output.	`detail extensive`
`Gre keepalives configured`	Indicates whether a GRE keepalive time and hold time are configured for the GRE tunnel. **NOTE**: EX Series switches do not support configuration of GRE tunnel keepalive times and hold times, so the value of this field is always `Off` in switch output.	`detail extensive`
`Gre keepalives adjacency state`	Status of the other end of the GRE tunnel: `Up` or `Down`. If keepalive messages are not received by either end of the GRE tunnel within the hold-time period, the GRE keepalive adjacency state is down even when the GRE tunnel is up.	`detail extensive`
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. • **Input rate**—Rate of bits and packets received on the interface. • **Output rate**—Rate of bits and packets transmitted on the interface.	**detail extensive**

Table 648: GRE show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive** none
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, or **mpls**.	**detail extensive** **none**
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
MTU	MTU size on the logical interface.	**detail extensive** **none**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none

Table 648: GRE show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces (GRE)

```
user@host> show interfaces gr-1/2/0
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
    Flags: Point-To-Point SNMP-Traps 16384
    IP-Header 192.0.2.2:192.0.2.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Input packets : 0
  Output packets: 0
    Protocol inet, MTU: 1476
      Flags: None
      Addresses, Flags: Is-Primary
        Local: 198.51.100.1
```

show interfaces brief (GRE)

```
user@host> show interfaces gr-1/2/0 brief
Physical interface: gr-1/2/0, Enabled, Physical link is Up
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

  Logical interface gr-1/2/0.0
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000
    IP-Header 10.10.0.2:10.10.0.1:47:df:64:0000000000000000
    Encapsulation: GRE-NULL
    inet  10.100.0.1/30
    mpls
```

show interfaces detail (GRE)

```
user@host> show interfaces gr-1/2/0 detail
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26, Generation: 13
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps

  Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47) (Generation 8)
    Flags: Point-To-Point SNMP-Traps 16384
    IP-Header 192.0.2.2:192.0.2.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
```

```
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 1476, Generation: 12, Route table: 0
      Flags: None
      Addresses, Flags: Is-Primary
        Destination: Unspecified, Local: 198.51.100.1, Broadcast: Unspecified,
        Generation: 15
```

show interfaces (Layer 2 Services Over GRE Interfaces)

```
user@host> show interfaces gr-2/2/10
show interfaces gr-2/2/10
Physical interface: gr-2/2/10, Enabled, Physical link is Up
  Interface index: 214, SNMP ifIndex: 690
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface gr-2/2/10.0 (Index 342) (SNMP ifIndex 10834)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 IP-Header
203.0.113.1:203.0.113.254:47:df:64:0000000000000000 Encapsulation: GRE-NULL
    L2 Routing Instance: vs1, L3 Routing Instance: default
    Copy-tos-to-outer-ip-header: Off
    Gre keepalives configured: Off, Gre keepalives adjacency state: down
    Input packets : 2
    Output packets: 0
    Protocol bridge, MTU: 1476
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 6/8, Local: 6.0.0.1, Broadcast: 6.255.255.255
```

show interfaces extensive (Layer 2 Services Over GRE Interfaces)

```
user@host> show interfaces gr-2/2/10.0 extensive

     Flags: SNMP-Traps Encapsulation: ENET2
     L2 Routing Instance: vs1, L3 Routing Instance: default
     Traffic statistics:
      Input  bytes  :             58851250
      Output bytes  :                    0
      Input  packets:              1279375
      Output packets:                    0
     Local statistics:
      Input  bytes  :                    0
      Output bytes  :                    0
      Input  packets:                    0
      Output packets:                    0
     Transit statistics:
      Input  bytes  :             58851250                75136 bps
      Output bytes  :                    0                    0 bps
      Input  packets:              1279375                  204 pps
      Output packets:                    0                    0 pps
     Protocol bridge, MTU: 1476, Generation: 175, Route table: 7
       Flags: Access-Mode
```

show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch

```
user@host> show interfaces gr-2/0/15 detail
Physical interface: gr-2/0/15, Enabled, Physical link is Up
  Interface index: 195, SNMP ifIndex: 846, Generation: 198
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:d2, Hardware address: 00:00:5e:00:53:d2
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: 2011-09-14 17:43:15 UTC (00:00:18 ago)
  Traffic statistics:
   Input  bytes  :              5600636                    0 bps
   Output bytes  :              5600636                    0 bps
   Input  packets:                20007                    0 pps
   Output packets:                20007                    0 pps
```

```
 IPv6 transit statistics:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0

Logical interface gr-2/0/15.0 (Index 75) (SNMP ifIndex 847) (HW Token 4093)
(Generation 140)
  Flags: Point-To-Point SNMP-Traps 0x0
  IP-Header 192.168.30.2:192.168.20.3:47:df:64:0000000000000000
  Encapsulation: GRE-NULL
  Copy-tos-to-outer-ip-header: Off
  Gre keepalives configured: Off, Gre keepalives adjacency state: down
  Traffic statistics:
   Input  bytes  :              5600886
   Output bytes  :              2881784
   Input  packets:                20010
   Output packets:                10018
  Local statistics:
   Input  bytes  :                  398
   Output bytes  :                  264
   Input  packets:                    5
   Output packets:                    3
  Transit statistics:
   Input  bytes  :              5600488                    0 bps
   Output bytes  :              2881520                    0 bps
   Input  packets:                20005                    0 pps
   Output packets:                10015                    0 pps
  Protocol inet, Generation: 159, Route table: 0
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.10.10/8, Local: 10.10.10.10, Broadcast: 10.10.10.255,
      Generation: 144

Logical interface gr-2/0/15.1 (Index 80) (SNMP ifIndex 848) (HW Token 4088)
(Generation 150)
  Flags: Point-To-Point SNMP-Traps 0x0
  IP-Header 192.168.40.2:192.168.30.1:47:df:64:0000000000000000
  Encapsulation: GRE-NULL
  Copy-tos-to-outer-ip-header: Off
  Gre keepalives configured: Off, Gre keepalives adjacency state: down
  Traffic statistics:
   Input  bytes  :                  260
```

```
     Output bytes  :              2880148
     Input  packets:                    4
     Output packets:                10002
    Local statistics:
     Input  bytes  :                  112
     Output bytes  :                    0
     Input  packets:                    2
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                  148                    0 bps
     Output bytes  :              2880148                    0 bps
     Input  packets:                    2                    0 pps
     Output packets:                10002                    0 pps
    Protocol inet, Generation: 171, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.10.10/8, Local: 10.10.10.10, Broadcast: 10.10.10.255,
        Generation: 160
```

show interfaces extensive (GRE)

The output for the `show interfaces extensive` command is identical to that for the `show interfaces detail` command. For sample output, see "show interfaces detail (GRE)" on page 20908 and "show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch" on page 20910.

show interfaces gr-2/0/10 for GRE IPv6 tunnel

```
user@host> show interfaces gr-2/0/10
show interfaces gr-2/0/10
Physical interface: gr-2/0/10, Enabled, Physical link is Up
  Interface index: 140, SNMP ifIndex: 559
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 4952 bps (3 pps)
  Output rate    : 200 bps (0 pps)

  Logical interface gr-2/0/10.0 (Index 355) (SNMP ifIndex 857)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 IP-Header
1000::11:0:11:1-1000::11:2:13:2-47-64-0-0-0000000000000000 Encapsulation: GRE-NULL
    Copy-tos-to-outer-ip-header: Off, Copy-tos-to-outer-ip-header-transit: Off
```

```
    Gre keepalives configured: Off, Gre keepalives adjacency state: down
    Input packets : 60
    Output packets: 83
    Protocol inet, MTU: 9082
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 14.0.13/24, Local: 14.0.13.1, Broadcast: 14.0.13.255
    Protocol iso, MTU: 9082
    Protocol inet6, MTU: 9082
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1400::14:0:13:0/120, Local: 1400::14:0:13:1
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a520:2875:4992
    Protocol mpls, MTU: 9070, Maximum labels: 3
      Flags: Is-Primary
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced before Junos OS Release 17.3R1.

show interfaces interface-set (Ethernet Interface Set)

IN THIS SECTION

- Syntax | 20914
- Description | 20914
- Options | 20914
- Required Privilege Level | 20914
- Output Fields | 20914
- Sample Output | 20917
- Release Information | 20920

Syntax

```
show interfaces interface-set interface-set-name
<detail | terse>
```

Description

Display information about the specified gigabit or 10-Gigabit Ethernet interface set.

You can also use the `show interfaces interface-set` command to display information about agent circuit identifier (ACI) interface sets.

Options

interface-set *interface-set-name*	Display information about the specified Gigabit Ethernet, 10-Gigabit Ethernet, ACI, or ALI interface set.
detail | terse	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 649 on page 20914 describes the information for the `show interfaces interface-set` command. Output fields are listed in the approximate order in which they appear.

Table 649: Ethernet show interfaces interface-set Output Fields

Field Name	Field Description	Level of Output
Physical Interface		

Table 649: Ethernet show interfaces interface-set Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface set`	Name of the interface set or sets. For ACI interface sets, the set name is prefixed with `aci-`. For ALI interface sets, the set name is prefixed with the trusted option that the interface set is based on: • `aci-`— The ACI is configured as the trusted option. • `ari-`— The ARI is configured as the trusted option. • `aci+ari-`— Both ACI and ARI are configured as the trusted option. • `noids-`— Neither the ACI nor the ARI is configured as the trusted option and neither ACI nor ARI is received.	All levels
`Interface set index`	Index number of the interface set.	`detail` none
`ACI VLAN`	For ACI interface sets, the string received in DHCP or PPPoE control packets that uniquely identifies the subscriber's access node and the DSL line on the access node. Only the `Agent Circuit ID` can be used to create the interface set. **NOTE**: The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the `line-identity` autoconfiguration stanza.	`detail` none

Table 649: Ethernet show interfaces interface-set Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Line Identity`	For ALI interface sets, the trusted option received in DHCP or PPPoE control packets that uniquely identifies the subscriber's access node and the DSL line on the access node. The trusted option can be either or both of the following: • `Agent Circuit ID`—The ACI value • `Agent Remote ID`—The ARI value. **NOTE**: When only `accept-no-ids` is configured as the trusted option, this field is not displayed. **NOTE**: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the `agent-circuit-id` autoconfiguration stanza.	`detail` none
`PPPoE`	Dynamic PPPoE subscriber interface that the router creates using the ACI or ALI interface set.	`detail` none
`Max Sessions`	For dynamic PPPoE subscriber interfaces, maximum number of PPPoE logical interfaces that that can be activated on the underlying interface.	`detail` none
`Max Sessions VSA Ignore`	For dynamic PPPoE subscriber interfaces, whether the router is configured to ignore (clear) the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks VSA [26-143] and restore the PPPoE maximum session value on the underlying interface to the value configured with the `max-sessions` statement: `Off` (default) or `On`.	`detail` none
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the specified interface set. • `Input bytes, Output bytes`—Number of bytes and number of bytes per second received and transmitted on the interface set • `Input packets, Output packets`—Number of packets and number of packets per second received and transmitted on the interface set.	`detail`

Table 649: Ethernet show interfaces interface-set Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Egress queues supported`	Total number of egress queues supported on the specified interface set.	`detail`
`Egress queues in use`	Total number of egress queues used on the specified interface set.	`detail`
`Queue counters`	`Queued packets`, `Transmitted packets`, and `Dropped packets` statistics for the four forwarding classes.	`detail`
`Members`	List of all interface sets or, for ACI interface sets, list of all subscriber interfaces belonging to the specified ACI interface set.	`detail` none

Sample Output

show interfaces interface-set terse

```
user@host> show interfaces interface-set terse
Interface set:
   iflset-xe-11/3/0-0
   ge-1/0/1-0
   ge-1/0/1-2
```

show interfaces interface-set detail

```
user@host> show interfaces interface-set iflset-xe-11/3/0-0 detail
Interface set: iflset-xe-11/3/0-0
  Interface set index: 19
  Traffic statistics:
   Output bytes  :            751017840            401673504 bps
   Output packets:             11044380               738377 pps
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets    Dropped packets
    0                        211091327             11044380          199995746
```

```
    1                                 0                    0                    0
    2                                 0                    0                    0
    3                                 0                    0                    0
  Members:
   xe-11/3/0.0
```

show interfaces interface-set (ACI Interface Set based on ACI)

```
user@host> show interfaces interface-set
Interface set: aci-1001-ge-5/2/0.10
  Interface set index: 1
  Interface set snmp index: 67108865
    ACI VLAN:
      Agent Circuit ID: circuit0
    PPPoE:
      Max Sessions: 32000, Max Sessions VSA Ignore: Off
  Members:
   demux0.3221225472
```

show interfaces interface-set (ACI Interface Set based on ACI Trusted Option)

```
user@host> show interfaces interface-set
Interface set: ari-1002-demux0.3221225473
  Interface set index: 2
  Interface set snmp index: 67108866
    Line Identity:
      Agent Circuit ID: remote20
    PPPoE:
      Max Sessions: 32000, Max Sessions VSA Ignore: Off
  Members:
   demux0.3221225474
```

show interfaces interface-set (ACI Interface Set based on ARI Trusted Option)

```
user@host> show interfaces interface-set
Interface set: aci-1002-demux0.3221225473
  Interface set index: 2
  Interface set snmp index: 67108866
```

```
    Line Identity:
      Agent Remote ID: remote20
    PPPoE:
      Max Sessions: 32000, Max Sessions VSA Ignore: Off
  Members:
   demux0.3221225474
```

show interfaces interface-set (ACI Interface Set based on ARI Trusted Option when both ACI and ARI are received)

```
user@host> show interfaces interface-set
Interface set: ari-1002-demux0.3221225473
  Interface set index: 2
  Interface set snmp index: 67108866
    Line Identity:
      Agent Remote ID: remote20
    PPPoE:
      Max Sessions: 32000, Max Sessions VSA Ignore: Off
  Members:
   demux0.3221225474
```

show interfaces interface-set (ACI Interface Set based on Accept-No-IDs Trusted Option when neither ACI nor ARI is received)

```
user@host> show interfaces interface-set
Interface set: noids-1002-demux0.3221225473
  Interface set index: 2
  Interface set snmp index: 67108866
  Members:
   demux0.3221225474
```

show interfaces interface-set (L2BSA and PPPoE Subscribers)

```
user@host> show interfaces interface-set
Interface set: ge-1/0/4
  Interface set index: 6
  Members:
```

```
ge-1/0/4.1073741908
pp0.1073741907
```

Release Information

Command introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers

show interfaces interface-set (PS Interface Set)

IN THIS SECTION

- Syntax | 20920
- Description | 20920
- Options | 20921
- Required Privilege Level | 20921
- Output Fields | 20921
- Sample Output | 20922
- Release Information | 20922

Syntax

```
show interfaces interface-set iflset-name
<detail | terse>
```

Description

Display information about the specified demux interface set.

Options

none Display standard information about the specified iflset.

detail | terse (Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 650 on page 20921 lists the output fields for the `show interfaces interface-set` *`iflset-name`* command. Output fields are listed in the approximate order in which they appear.

Table 650: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Members`	Status of the physical link (Up or Down).	`detail` none
`Targeting summary`	Status of LT links that are configured with targeted distribution (primary or backup)	none
Logical Interface		
`Interface set`	Name of the interface set or sets. • Interface set index—Index number of the interface set. • Interface set snmp index—SNMP interface index number for the interface set.	none

Sample Output

show interfaces interface-set *iflset-name*

```
user@host> show interfaces interface-set iflset100101
  Interface set: iflset100101
  Interface set index: 536870916
  Interface set snmp index: 603979780
  Members:
   demux0.3221225483
  Targeting summary:
    lt-4/1/0, Primary, Physical link is Up
    lt-1/0/0, Backup, Physical link is Up
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

- MPLS Pseudowire Subscriber Logical Interfaces

show interfaces interface-set queue

IN THIS SECTION

- Syntax | 20923
- Description | 20923
- Options | 20923
- Required Privilege Level | 20923
- Output Fields | 20924
- Sample Output | 20925
- Release Information | 20933

Syntax

```
show interfaces interface-set queue interface-set-name
<aggregate | remaining-traffic>
<both-ingress-egress>
<egress>
<forwarding-class class-name>
<ingress>
```

Description

Display information about the interface set queue.

Options

Following are all possible options for the `show interfaces interface-set queue` command. The options that appear may vary depending on the platform and software release.

interface-set-name	(Optional) Display information about the specified interface set. Wildcard values can be used in the interface set name.
aggregate	(Optional) Display the aggregated queuing statistics of all member logical interfaces for interface sets that have traffic-control profiles configured.
both-ingress-egress	(Optional) Display both ingress and egress queue statistics.
egress	(Optional) Display egress queue statistics.
forwarding-class *class-name*	(Optional) Display queuing statistics for the specified forwarding class.
ingress	(Optional) Display ingress queue statistics.
remaining-traffic	(Optional) Display the queuing statistics of all member logical interfaces for interface sets that do not have traffic-control profiles configured.

Required Privilege Level

view

Output Fields

Table 651 on page 20924 lists all possible fields for the `show interfaces interface-set queue` command. The fields that appear may vary depending on the platform and software release.

Table 651: Ethernet show interfaces interface-set queue Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Interface set	Name of the interface set.	All levels
Interface set index	Index number of the interface set.	All levels
Forwarding classes supported	Total number of forwarding classes supported on the specified interface set.	All levels
Forwarding classes in use	Total number of forwarding classes used on the specified interface set.	All levels
Egress queues supported	Total number of egress queues supported on the specified interface set.	All levels
Egress queues in use	Total number of egress queues used on the specified interface set.	All levels
Ingress queues supported	Total number of ingress queues supported on the specified interface set.	All levels
Ingress queues in use	Total number of ingress queues used on the specified interface set.	All levels
Queue	Egress or ingress queue number for the statistics being displayed.	All levels

Table 651: Ethernet show interfaces interface-set queue Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Forwarding classes	Forwarding class name for the statistics being displayed.	All levels
Queued	**Packet** and **Byte** statistics for the specified queue. • **Packets**—Number of packets queued and input rate in packets per second. • **Bytes**—Number of bytes queued and input rate in bytes per second.	All levels
Transmitted	**Packet** and **Byte** statistics for the specified forwarding class. • **Packets**—Number of packets transmitted and transmit rate in packets per second. • **Bytes**—Number of bytes transmitted and transmit rate in bytes per second. • **Tail-dropped packets**—Number of packets tail dropped. • **RED-dropped packets**—Number of RED-dropped packets for the **low**, **medium-low**, **medium-high**, and **high** loss priorities. • **RED-dropped bytes**—Number of RED-dropped bytes for the **low**, **medium-low**, **medium-high**, and **high** loss priorities.	All levels

Sample Output

show interfaces interface-set queue (Gigabit Ethernet)

```
user@host> show interfaces interface-set queue ge-2/2/0-0
Interface set: ge-2/2/0-0
  Interface set index: 3
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :               3998482                    1 pps
```

```
    Bytes                :             271896884                   688 bps
  Transmitted:
    Packets              :               1077474                     1 pps
    Bytes                :              73268340                   688 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :               2921008                     0 pps
     Low                 :               2921008                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :             198628544                     0 bps
     Low                 :             198628544                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
...
```

show interfaces interface-set queue both-ingress-egress (Enhanced DPC)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 both-ingress-egress
Interface set: ge-2/2/0-0
  Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :             185968478                473161 pps
    Bytes                :           10042313520             204441336 bps
  Transmitted:
    Packets              :               5441673                 13780 pps
    Bytes                :             293850342               5952960 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :             180526772                459372 pps
    RED-dropped bytes    :            9748446282             198451512 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
```

```
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :            522021472               473602 pps
    Bytes                :          28190332480            204599944 bps
  Transmitted:
    Packets              :              5791772                 4055 pps
    Bytes                :            312755688              1751976 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :            516227139               469546 pps
    RED-dropped bytes    :          27876265560            202843872 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              5417304                13797 pps
    Bytes                :            368429508              7506096 bps
  Transmitted:
    Packets              :              5014996                12769 pps
    Bytes                :            341019728              6946560 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :               402189                 1028 pps
     Low                 :               402189                 1028 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
```

```
    RED-dropped bytes    :              27348852                559536 bps
     Low                 :              27348852                559536 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
     Low                 :                     0                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low                 :                     0                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :               5770534                  3963 pps
    Bytes                :             396943252               2156144 bps
  Transmitted:
    Packets              :               3945152                  1457 pps
    Bytes                :             268270336                792608 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :               1815141                  2506 pps
     Low                 :               1815141                  2506 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :             123429524               1363536 bps
     Low                 :             123429524               1363536 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
```

```
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
```

show interfaces interface-set queue egress (Enhanced DPC)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 egress
Interface set: ge-2/2/0-0
  Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :               3958253                13822 pps
    Bytes                :             269217592              7519712 bps
  Transmitted:
    Packets              :               3665035                12729 pps
    Bytes                :             249222380              6924848 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                293091                 1093 pps
     Low                 :                293091                 1093 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :              19930188               594864 bps
     Low                 :              19930188               594864 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
```

```
     High                :                    0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :              5350989                 3904 pps
    Bytes                :            368412924              2124048 bps
  Transmitted:
    Packets              :              3790469                 1465 pps
    Bytes                :            257751892               796960 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :              1550282                 2439 pps
     Low                 :              1550282                 2439 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :            105419176              1327088 bps
     Low                 :            105419176              1327088 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
```

```
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
```

show interfaces interface-set queue forwarding-class (Gigabit Ethernet)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 forwarding-class best-effort
Interface set: ge-2/2/0-0
  Interface set index: 3
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :             101857694               1420083 pps
    Bytes                :            6927234456             772532320 bps
  Transmitted:
    Packets              :               3984693                 55500 pps
    Bytes                :             270959592              30192512 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :              97870952               1364583 pps
     Low                 :              97870952               1364583 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :            6655225776             742339808 bps
     Low                 :            6655225776             742339808 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
```

show interfaces interface-set queue (Enhanced DPC)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 ingress
Interface set: foo
  Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :             149036817                473711 pps
    Bytes                :            8048003934             204642936 bps
  Transmitted:
    Packets              :               4360749                 13891 pps
    Bytes                :             235480446               6000912 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :             144676035                459820 pps
    RED-dropped bytes    :            7812506592             198642024 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :             485089207                473605 pps
    Bytes                :           26195987476             204597576 bps
  Transmitted:
    Packets              :               5480799                  3959 pps
    Bytes                :             295963146               1710504 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :             479605853                469646 pps
    RED-dropped bytes    :           25898716170             202887072 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
```

```
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
```

show interfaces interface-set queue remaining-traffic (Gigabit Ethernet)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 remaining-traffic
Interface set: ge-2/2/0-0
  Interface set index: 12
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :               2201552                    0 pps
    Bytes                :             149705536                    0 bps
  Transmitted:
    Packets              :                609765                    0 pps
    Bytes                :              41464020                    0 bps
    Tail-dropped packets :                     0                    0 pps
    RED-dropped packets  :               1591787                    0 pps
     Low                 :               1591787                    0 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :             108241516                    0 bps
     Low                 :             108241516                    0 bps
     Medium-low          :                     0                    0 bps
     Medium-high         :                     0                    0 bps
     High                :                     0                    0 bps
```

Release Information

Command introduced in Junos OS Release 8.5.

show interfaces interval

IN THIS SECTION

- Syntax | 20934
- Description | 20934
- Options | 20934
- Required Privilege Level | 20934
- Output Fields | 20935
- Sample Output | 20936
- Release Information | 20938

Syntax

```
show interfaces interval
<interface-name>
```

Description

Display the channel service unit (CSU) interface alarm and error count in 15-minute intervals for the past 24 hours. If the system has been operational for less than 24 hours, the maximum number of intervals available is displayed.

Options

interface-name	(Optional) Name of a particular interface.

Required Privilege Level

view

Output Fields

Table 652 on page 20935 lists the output fields for the `show interfaces interval` command. Output fields are listed in the approximate order in which they appear.

Table 652: show interfaces interval Output Fields

Field Name	Field Description
`Physical interface`	Name of the interface.
`SNMP ifIndex`	SNMP index number for the physical interface.
`hh:mm``-current`	Time of day (in hours and minutes) at the beginning of the latest counter interval. The value of the latest counter interval is always less than 15 minutes.
`hh:mm-hh:mm`	Time of day (in hours and minutes) at the beginning and end of each 15-minute interval.
`alarm` or `event:` *`n`*	Count of alarms and events within each 15-minute interval. The specific alarm or event depends on the interface media type. For a description of the alarm or event listed, see the *`interface-type`* `media` field (for example, `T1 media`) under the `show interfaces` command for the particular interface type in which you are interested.
`Interval Total`	Sum of all the alarm and defect counters for the last 24-hour period.
Interval Total	Sum of all the alarm and defect counters for the last 24-hour period.
`Current Day Interval Total`	Sum of all the alarm and defect counters in the current day. **NOTE**: The `Current Day Interval` output field is reset after 24 hours.
`Previous Day Interval Total`	Sum of all the alarm and defect counters in the previous day.

Sample Output

show interfaces interval (Channelized OC12)

```
user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (E3)

```
user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
```

```
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH)

```
user@host> show interfaces interval so-2/2/0
Physical interface: so-2/2/0, SNMP ifIndex: 553
  02:53-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  02:38-02:53:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  02:23-02:38:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  02:08-02:23:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  01:53-02:08:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  01:38-01:53:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  01:23-01:38:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  01:08-01:23:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  00:53-01:08:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  00:38-00:53:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
....
  Current Day Interval Total:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  Previous Day Interval Total (Last updated at 02:23):
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear interfaces interval | 3641

show interfaces (IP-over-IP)

IN THIS SECTION

- Syntax | 20938
- Description | 20938
- Options | 20939
- Required Privilege Level | 20939
- Output Fields | 20939
- Sample Output | 20943
- Release Information | 20945

Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified IP-over-IP interface.

Options

interface-type	On M Series and T Series routers, the interface type is **ip-***fpc/pic/port*.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 653 on page 20939 lists the output fields for the `show interfaces` (IP-over-IP) command. Output fields are listed in the approximate order in which they appear.

Table 653: IP-over-IP show interfaces Output Fields

Field	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none

Table 653: IP-over-IP show interfaces Output Fields *(Continued)*

Field	Field Description	Level of Output
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface.	All levels
Link-level type	Encapsulation used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 653: IP-over-IP show interfaces Output Fields *(Continued)*

Field	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	All levels
IP Header	IP header of the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified

Table 653: IP-over-IP show interfaces Output Fields *(Continued)*

Field	Field Description	Level of Output
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. • **Input rate**—Rate of bits and packets received on the interface. • **Output rate**—Rate of bits and packets transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, or **mpls**.	**detail extensive** **none**
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**

Table 653: IP-over-IP show interfaces Output Fields *(Continued)*

Field	Field Description	Level of Output
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive** none

Sample Output

show interfaces (IP-over-IP)

```
user@host> show interfaces ip-0/0/0
Physical interface: ip-0/0/0, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 27
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ip-0/0/0.0 (Index 69) (SNMP ifIndex 49)
    Flags: Point-To-Point SNMP-Traps 16384
    IP-Header 192.0.2.1:192.0.2.2:4:df:64:00000000 Encapsulation: IPv4-NULL
  Input packets : 0
  Output packets: 0
    Protocol inet, MTU: 1480
      Flags: None
```

show interfaces brief (IP-over-IP)

```
user@host> show interfaces ip-0/0/0 brief
Physical interface: ip-0/0/0, Enabled, Physical link is Up
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface ip-0/0/0.0
```

```
    Flags: Point-To-Point SNMP-Traps 16384
    IP-Header 192.0.2.1:192.0.2.2:4:df:64:00000000 Encapsulation: IPv4-NULL
    inet
```

show interfaces detail (IP-over-iP)

```
user@host> show interfaces ip-0/0/0 detail
Physical interface: ip-0/0/0, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 27, Generation: 14
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps

  Logical interface ip-0/0/0.0 (Index 69) (SNMP ifIndex 49) (Generation 9)
    Flags: Point-To-Point SNMP-Traps 16384
    IP-Header 192.0.2.1:192.0.2.2:4:df:64:00000000 Encapsulation: IPv4-NULL
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 1480, Generation: 13, Route table: 0
      Flags: None
```

show interfaces extensive (IP-over-IP)

The output for the show interfaces extensive command is identical to that for the show interfaces detail command. For sample output, see "show interfaces detail (IP-over-iP)" on page 20944.

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces irb

IN THIS SECTION

- Syntax | 20945
- Description | 20946
- Options | 20946
- Additional Information | 20946
- Required Privilege Level | 20946
- Output Fields | 20946
- Sample Output | 20953
- Release Information | 20957

Syntax

```
show interfaces irb
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display integrated routing and bridging interfaces information.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
mac	Display hardware MAC address
media	(Optional) Display media-specific information about network interfaces.
snmp-index ***snmp-index***	(Optional) Display information for the interface with the specified SNMP index.
statistics	(Optional) Display static interface statistics.

Additional Information

Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another bridging domain that has a Layer 3 protocol configured.

Required Privilege Level

view

Output Fields

Table 654 on page 20946 lists the output fields for the `show interfaces irb` command. Output fields are listed in the approximate order in which they appear.

Table 654: show interfaces irb Output Fields

Field Name	Field Description	Level of Output
Physical Interface		

Table 654: show interfaces irb Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Physical interface	Name of the physical interface.	All levels
Enabled	State of the physical interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
Proto	Protocol configured on the interface.	**terse**
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Type	Physical interface type.	**detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	**detail extensive brief** none
MTU	MTU size on the physical interface.	**detail extensive brief** none
Clocking	Reference clock source: **Internal** or **External**. Always unspecified on IRB interfaces.	**detail extensive brief**
Speed	Speed at which the interface is running. Always unspecified on IRB interfaces.	**detail extensive brief**
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	**detail extensive brief** none
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	**detail extensive brief** none

Table 654: show interfaces irb Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Link type	Physical interface link type: **full duplex** or **half duplex**.	**detail extensive** none
Link flags	Information about the link. Possible values are described in the "Links Flags" section under *Common Output Fields Description*.	**detail extensive** none
Physical Info	Physical interface information.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	MAC address of the hardware.	**detail extensive** none
Alternate link address	Backup address of the link.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 654: show interfaces irb Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Table 654: show interfaces irb Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame terminates and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**detail extensive**

Table 654: show interfaces irb Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the DPC is malfunctioning. • **Errors**—Sum of the outgoing frame terminates and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface (which reflects its initialization sequence).	**detail extensive** none
SNMP ifIndex	SNMP interface index number of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	**detail extensive**

Table 654: show interfaces irb Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Encapsulation	Encapsulation on the logical interface.	**detail extensive**
Bandwidth	Dummy value that is ignored by an IRB interface. IRB interfaces are pseudo interfaces and do not have physical bandwidth associated with them.	**detail extensive**
Routing Instance	Routing instance IRB is configured under.	**detail extensive**
Bridging Domain	Bridging domain IRB is participating in.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router.	**detail extensive**

Table 654: show interfaces irb Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Protocol	Protocol family configured on the local interface. Possible values are described in the "Protocol Field" section under *Common Output Fields Description*.	**detail extensive**
MTU	Maximum transmission unit size on the logical interface.	**detail extensive**
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Addresses, Flags	Information about address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive**
Policer	The policer that is to be evaluated when packets are received or transmitted on the interface.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	**detail extensive**

Sample Output

show interfaces irb extensive

```
user@host> show interfaces irb extensive
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running
```

```
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0

  Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
    Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
    Bandwidth: 1000mbps
    Routing Instance: customer_0 Bridging Domain: bd0
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
     IPv6 transit statistics:
      Input  bytes  :                  0
      Output bytes  :                  0
      Input  packets:                  0
      Output packets:                  0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
```

```
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: 1500, Generation: 154, Route table: 0
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255,
        Generation: 155
    Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer
```

show interfaces irb snmp-index

```
user@host> show interfaces snmp-index 25
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 25
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
    Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
    Bandwidth: 1000mbps
    Routing Instance: customer_0 Bridging Domain: bd0
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1500
```

```
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
    Protocol multiservice, MTU: 1500
      Flags: Is-Primary
```

show interfaces irb (SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM)

```
user@host> show interfaces irb
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 503
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 00:90:69:7f:e7:f0, Hardware address: 00:90:69:7f:e7:f0
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface irb.0 (Index 327) (SNMP ifIndex 507)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
    MAC: 00:11:22:33:44:55
    Bandwidth: 1000mbps
    Routing Instance: default-switch Bridging Domain: bd512+512
    Input packets : 0
    Output packets: 1
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
    Protocol multiservice, MTU: 1500
```

show interfaces irb.0 (SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM)

```
user@host> show interfaces irb.0
regress@bloodlust# run show interfaces irb.0
  Logical interface irb.0 (Index 327) (SNMP ifIndex 507)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
    MAC: 00:11:22:33:44:55
```

```
    Bandwidth: 1000mbps
    Routing Instance: default-switch Bridging Domain: bd512+512
    Input packets : 0
    Output packets: 1
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
    Protocol multiservice, MTU: 1500
```

Release Information

Command introduced in Junos OS Release 8.4.

show interfaces l2-routing-instance

IN THIS SECTION

- Syntax | 20957
- Description | 20958
- Options | 20958
- Required Privilege Level | 20958
- Output Fields | 20958
- Sample Output | 20959
- Release Information | 20959

Syntax

```
show interfaces l2-routing-instance routing-instance-name
```

Description

Display information about core-facing physical interfaces in the specified routing instance. The routing instance must be configured for Layer 2 wholesale operations or the command displays no output.

Options

routing-instance-name	Name of the routing instance.

Required Privilege Level

view

Output Fields

Table 655 on page 20958 lists the output fields for the `show interfaces l2-routing-instance` command. Output fields are listed in the approximate order in which they appear.

Table 655: show interfaces l2-routing-instance Output Fields

Field Name	Field Description
`Interface`	Name of .the core-facing physical interface.
`VLAN Id`	Trunk VLAN ID assigned to the core-facing physical interface within its input VLAN map.
`Inner VLAN Id total`	Aggregate number of VLAN IDs assigned to this physical interface based on the inner VLAN ID-swap-ranges configured for the interface.
`Inner VLAN range`	Inner VLAN ID swap range; numbers that can be used for swapping inner VLAN IDs.
`Inner VLAN range total`	Size of the inner VLAN ID swap range.
`Inner VLAN id use count`	Number of inner VLAN ID tags in use.
`Inner VLAN id free count`	Number of inner VLAN ID tags not in use.

Sample Output

show interfaces l2-routing-instance

```
user@host> show interfaces l2-routing-instance NSP1
Interface: ge-1/1/1.0
VLAN Id: 100
  Inner VLAN Id total: 6
  Inner VLAN range: 15-20
    Inner VLAN range total   : 6
    Inner VLAN id use count  : 1
    Inner VLAN id free count : 5
```

Release Information

Command introduced in Junos OS Release 16.1R4.

RELATED DOCUMENTATION

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

show interfaces lb-stats af

IN THIS SECTION

- Syntax | 20960
- Description | 20960
- Required Privilege Level | 20960
- Output Fields | 20960
- Sample Output | 20963
- Release Information | 20964

Syntax

```
show interfaces lb-stats af-interface-name
```

Description

Displays information about the load balancing of transmit traffic (from the local FPC to the remote FPC) on each peer Packet Forwarding Engine of the abstracted fabric (af) interface. This command also displays the statistics of the transmit traffic on the fabric queues (high and low queues) for each peer Packet Forwarding Engine on the af interface.

NOTE:

- This command does not display information on traffic that is not load-balanced over af interface.

Required Privilege Level

View

Output Fields

Table 656 on page 20960lists the output fields for the show interfaces lb-stats *af-interface* command.

Table 656: show interfaces lb-stats *af-interface-name* Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name and status of the physical interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none

Table 656: show interfaces lb-stats *af-interface-name* Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Type`	Interface type.	
`Link-level type`	Encapsulation being used on the physical interface.	All levels
`MTU`	Maximum transmission unit size on the physical interface.	All levels
`Speed`	Speed at which the interface is running.	All levels
`Device flags`	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
`Interface flags`	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
`Link type`	Link type. For example, Full-Duplex.	All levels
`Current address`	Configured MAC address.	`detail extensive` none
`Hardware address`	Hardware MAC address.	`detail extensive` none
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second`*`:`*`timezone`* (*`hour`*`:`*`minute`*`:`*`second`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive` none

Table 656: show interfaces lb-stats *af-interface-name* Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Collapsed Forwarding	Displays the fabric optimization status. • Optimized • Disabled • Monitored See *Understanding Junos Node Slicing*	All levels
Peer GNF id	The GNF peer connected using the abstracted fabric interface.	`detail` `extensive` none
`Peer GNF Forwarding element(FE) view`	Shows Packet Forwarding Engine (also called forwarding element or FE) number and the FPC slot, FE bandwidth, FE status (up/down), FE transmit packets, and FE transmit bytes. The FE transmit statistics (in packets and bytes) provide insights into the load balancing of transmit traffic to each peer Packet Forwarding Engine of the abstracted fabric (`af`) interface.	`detail` `extensive` none
Residual Transmit Statistics	Displays the historical transmit statistics for the peer `af` interface from the local or remote FPC events (FPC offline/online/removal/power off on the local or remote side). The transmit values (in packets and bytes) indicates the statistics of the transmit traffic on the fabric queues (high and low queues) for each peer Packet Forwarding Engine on the `af` interface. **NOTE**: The data displayed in the Residual Transmit Statistics field is based on the last user triggered query, using the `show interfaces lb-stats` *`af-name`* command before the last FPC reboot.	
Fabric Queue Statistics	Displays the statistics of the transmit traffic on the fabric queues (high and low queues) for each peer PFE on the `af` interface.	
Residual Queue Statistics	Displays the historical fabric queue statistics for the peer `af` interface from the FPC events (FPC offline/online/removal/power off on the local or remote side).	

Sample Output

show interfaces lb-stats af2

```
user@host> show interfaces lb-stats af2
```

```
Physical interface: af2, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 501
  Type: Ethernet, Link-level type: Ethernet, MTU: 1518, Speed: 3000000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Current address: 00:90:69:15:20:47, Hardware address: 00:90:69:15:20:47
  Last flapped   : 2021-09-01 00:40:40 PDT (2w6d 01:56 ago)

  Peer GNF id    : 2
  Collapsed Forwarding : Disabled
  Peer GNF Forwarding element(FE) view :
  FPC slot:FE num  FE Bandwidth(Gbps) Status      Transmit Packets
Trans                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                
mit
Bytes
       2:2                   500         Up
0
             0
       2:3                   500         Up
0
             0
       2:4                   500         Up
0
             0
       2:5                   500         Up
0
             0
       2:6                   500         Up
203674
           19873016
       2:7                   500         Up
834
           38364
```

```
Residual Transmit Statistics :
Packets :                    0 Bytes :                    0

Fabric Queue Statistics :
FPC slot:FE num    High priority(pkts)        Low priority(pkts)
     2:2                            0                         0
     2:3                            0                         0
     2:4                            0                         0
     2:5                            0                         0
     2:6                            0                    203674
     2:7                            0                       834
FPC slot:FE num    High priority(bytes)      Low priority(bytes)
     2:2                              0                        0
     2:3                              0                        0
     2:4                              0                        0
     2:5                              0                        0
     2:6                              0                 19873016
     2:7                              0                    38364
Residual Queue Statistics :
    High priority(pkts)       Low priority(pkts)
                     0                        0
    High priority(bytes)      Low priority(bytes)
                      0                        0
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

clear interfaces lb-stats af | 3643

show interfaces (Link Services)

IN THIS SECTION

- Syntax | 20965
- Description | 20965
- Options | 20966
- Required Privilege Level | 20966
- Output Fields | 20966
- Sample Output | 20981
- Release Information | 20986

Syntax

For Multilink Frame Relay user-to-user network-to-network interface (UNI NNI):

```
show interfaces interface-type :channel
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

For Multilink Frame Relay end-to-end:

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified link services interface.

Options

interface-type	On M Series and T Series routers, the interface type is **ls-*fpc/pic/port***.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
channel	Channel number of the link services interface.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 657 on page 20966 lists the output fields for the `show interfaces` (link services) command. Output fields are listed in the approximate order in which they appear.

Table 657: Link Services show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface: **Multilink-Frame-Relay-UNI-NNI** (default), **LinkService**, **Frame-relay**, **Frame-relay-ccc**, or **Frame-relay-tcc**.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour.minute.second timezone*** **(*hour.minute.second* ago)** . For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Link flags	Information about the link. Possible values are described in the "Link Flags" section under *Common Output Fields Description*.	All levels
Hold-times	Current interface hold time up and hold time down, in milliseconds, in the format **Up *n* ms**, **Down *n* ms**.	**detail extensive**

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Multilink Frame Relay UNI NNI bundle options	Multilink Frame Relay UNI NNI only) Configured information about Multilink Frame Relay bundle options.
	• **Device type**—**DCE** (Data Communication Equipment) or **DTE** (Data Terminal Equipment).
	• **MRRU**—Configured size of the maximum received reconstructed unit (MRRU): **1500** to **4500** bytes. The default is **1524** bytes.
	• **Fragmentation threshold**—Configured fragmentation threshold: **128** through **16,320** bytes, in integer multiples of **64** bytes. The default setting is **0**, which disables fragmentation.
	• **Red differential delay limit**—Red differential delay limit among bundle links has been reached, indicating an action will occur.
	• **Yellow differential delay limit**—Yellow differential delay among bundle links has been reached, indicating a warning will occur.
	• **Red differential delay action**—Type of actions taken when the red differential delay exceeds the red limit: **Disable link transmit** or **Remove link from service**.
	• **Reassembly drop timer**—Drop timeout value to provide a recovery mechanism if individual links in the link services bundle drop one or more packets: **1** through **127** milliseconds. By default, the drop timeout parameter is **0** (disabled). A value that is under **5** ms is not recommended.
	• **Links needed to sustain bundle**—Minimum number of links to sustain the bundle: **1** through **8**.
	• **LIP Hello timer**—Link Interleaving Protocol hello timer: **1** through **180** seconds.
	  • **Acknowledgement timer**—Maximum period to wait for an add link acknowledgement, hello acknowledgement, or remove link acknowledgement: **1** through **10** seconds.
	  • **Acknowledgement retries**—Number of retransmission attempts to be made for consecutive hello or remove link messages after the expiration of the acknowledgement timer: **1** through **5.**
		detail extensive
		none

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Multilink Frame Relay UNI NNI bundle options (continued)	• **Bundle class**—Bundle class ID.
	• **LMI type**—Multilink Frame Relay UNI NNI LMI type:**ANSI** or **Q.933 ANNEX A**.
	  • **T391 LIV polling timer**—Multilink Frame Relay UNI NNI Full status polling counter: **1** through **255**, with a default value of **6**.
	  • **T392 polling verification timer**—Multilink Frame Relay UNI NNI LMI error threshold. The number of errors required to bring down the link, within the event count specified by **N393**. The range is **1** through **10**, with a default value of **3**.
	  • **N391 full status polling count**—Multilink Frame Relay UNI NNI Full status polling counter: **1** through **255**.
	  • **N392 error threshold**—Multilink Frame Relay UNI NNI LMI error threshold: **1** through **10**.
	• **N393 monitored event count**—Multilink Frame Relay UNI NNI LMI monitored event count: **1** through **10**, with a default value of **4.**
		detail extensive
		none

Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.	**detail extensive**
Multilink Frame Relay UNI NNI bundle errors	Information about Multilink Frame Relay bundle errors. • **Packet drops**—Number of packets dropped. • **Fragment drops**—Number of fragments dropped. • **MRRU exceeded**—Number of times a packet was dropped because the configured MRRU value was exceeded. • **Exception events**—Exception events counter.	**detail extensive**

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Multilink Frame Relay UNI NNI bundle statistics	Information about Multilink Frame Relay bundles.
	• **Fragments**—Bundle fragment information.
	  • **Input**—Total number and rate of frames and packets received, in **Frames**, **fps** (frames per second), **Bytes**, and **bps** (bits per second).
	  • **Output**—Total number and rate of frames and packets transmitted, in **Frames**, **fps**, **Bytes**, and **bps**.
	• **Packets**—Bundle packet information.
	  • **Input**—Total number and rate of frames and packets received, in **Frames**, **fps** (frames per second), **Bytes**, and **bps** (bits per second).
	  • **Output**—Total number and rate of frames and packets transmitted, in **Frames**, **fps**, **Bytes**, and **bps**.
		detail extensive

Multilink Frame Relay UNI NNI bundle links information	• **Active bundle links**—Number of bundle links that are currently active. • **Removed bundle links**—Number of bundle links that have been removed (RED differential delay action). • **Disabled bundle links**—Number of bundle links that have been disabled (RED differential delay action).	**detail extensive** none

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Multilink Frame Relay UNI NNI active bundle links statistics	(Multilink Frame Relay UNI NNI only) Display information for each active bundle link. • **Frames**—Number of multilink control frames received on this bundle link. • **fps**—Rate of multilink control frames received on this bundle link (in frames per second). • **Bytes**—Number of bytes received on this bundle link. • **bps**—Number of bits per second received on this bundle link. • ***interface-name***—Name of the bundle link interface. • **Input**—Total number and rate of frames and packets received. • **Output**—Total number and rate of frames and packets transmitted. • **Current differential delay**—Compare this bundle link's round trip time to the average of all bundle links' round trip times in **ms** (milliseconds). • **Recent high differential delay**—Highest differential delay value from the latest 10 intervals, in milliseconds. • **Times over red diff delay**—Number of times this bundle link exceeded the configured red differential delay limit. • **Times over yellow diff delay**—Number of times this bundle link exceeded the configured yellow differential delay limit.	**detail extensive**

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Multilink Frame Relay UNI NNI active bundle links statistics (continued)	• **LIP**—Link Interleaving Protocol information. • **Rcv**—Number of messages received. • **Xmt**—Number of messages transmitted. • **add_lnk**—ADD__LINK message notifies the peer endpoint that the local endpoint supports frame processing. It is generated on both ends of a bundle link when a bundle link endpoint is ready to become operational. • **lnk_ack**—ADD_LINK_ACK message notifies the peer that the local router has received a valid ADD_LINK message. • **lnk_rej**—ADD_LINK_REJ message notifies the peer that the local router has received an invalid ADD_LINK message. • **hello**—HELLO message notifies the peer that the local router is up. Both ends of a link bundle generate this message. • **hel_ack**—HELLO_ACK message notifies the peer that the local router has received a valid HELLO message. • **lnk_rem**—REMOVE_LINK message notifies the peer that the local router has received a REMOVE_LINK message. • **rem_ack**—REMOVE_LINK_ACK message notifies the peer that the local router has received a valid ADD_LINK message.	**detail extensive**

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Frame exceptions	For Multilink Frame Relay end-to-end only. Information about framing exceptions. Includes events recorded under **Exception Events** for each logical interface. • **Oversized frames**—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits). • **Errored input frames**—Number of input frame errors. • **Input on disabled link/bundle**—Number of frames received on disabled links. These frames can result either from an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it. • **Output for disabled link/bundle**—Number of frames sent for a disabled or unavailable link. These frames can result either from an inconsistent configuration, or from a bundle being brought up or down while traffic is flowing through it. • **Queuing drops**—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed.	**detail extensive**
Buffering exceptions	For Multilink Frame Relay end-to-end only. Information about buffering exceptions. Includes events recorded under **Exception Events** for each logical interface: • **Packet data buffer overflow**—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services interface capacity. Check the logical interface exception event counters to determine which bundle is responsible.	**detail extensive**

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Assembly exceptions	For Multilink Frame Relay end-to-end only. Information about assembly exceptions. Includes events recorded under **Exception Events** for each logical interface. An assembly exception does not necessarily indicate an operational problem with the physical link services interface itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the assembling multilink interface records one or more exception events. The multilink interface can drop multilink-encapsulated fragments itself as a result. Any multilink packets or fragments dropped by the physical link services interface itself result in packet or fragment drop counts on individual logical interfaces. If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, excess differential delay or traffic losses on individual interfaces can be the root cause.	**detail extensive**

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Assembly exceptions (continued)	• **Fragment timeout**—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the physical link services interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • **Missing sequence number**—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the physical link services interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers occurred within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible. • **Out-of-range sequence number**—Frame was received with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible.	**detail extensive**

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Hardware errors	For Multilink Frame Relay end-to-end only. Information about hardware errors: • **Data memory error**—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. • **Control memory error**—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	All levels
Encapsulation	Encapsulation being used: PPP, Multilink - FR or Multilink - PPP	All levels
Bandwidth	Speed at which the interface is running.	All levels

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bundle options	For Multilink Frame Relay end-to-end interfaces only: • **MRRU**—Configured size of the maximum received reconstructed unit (MRRU): **1500** to **4500** bytes. The default is **1524** bytes. • **Drop timer period**—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: **1** through **127** milliseconds. Values under **5** milliseconds are not recommended. The default setting is **0**, which disables the timer. • **Sequence number format**—(MLPPP) Short sequence number header format. • **Fragmentation threshold**—Configured fragmentation threshold: **128** through **16,320** bytes, in integer multiples of **64** bytes. The default setting is **0**, which disables fragmentation. • **Links needed to sustain bundle**—Minimum number of links to sustain the bundle: **1** through **8**. • **Interleave fragments**—State of the process that interleaves long packets with high-priority ones. Only **Disabled** is currently supported. • **Remote MRRU**—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed.	**detail extensive** none

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bundle status (MLPPP) or Multilink class status (MC-MLPPP)	Information about bundle status: • **Remote MRRU**—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. • **Received sequence number**—Sequence number for received packets. • **Transmit sequence number**—Sequence number for transmitted packets. • **Packet drops**—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • **Fragment drops**—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers.	**detail extensive** none

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bundle status (MLPPP) or Multilink class status (MC-MLPPP) (continued)	• **MRRU exceeded**—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • **Fragment timeout**—Drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • **Missing sequence number**—Gap detected in the sequence numbers of fragments on a bundle. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • **Out-of-range sequence number**—Frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • **Packet data buffer overflow**—Packet buffer memory full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity.	**detail extensive** none

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bundle errors	Information about bundle errors. • **Packet drops**—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. • **Fragment drops**—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. • **MRRU exceeded**—Number of reassembled packets exceeding the MRRU. • **Exception events**—Number of exceptional events encountered other than MRRU exceeded errors. These events are categorized under the physical interface: **Frame exceptions**, **Buffering exceptions**, and **Fragment exceptions**. Exception events do not necessarily indicate that the multilink interface is not operating properly. Individual link failures can produce exceptional events.	**detail extensive**
Statistics	Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. • **Bundle**—Information about bundles. • **Link**—Information about links used in the multilink operation.	**detail extensive**
Protocol	Protocol family configured on the logical interface.	**detail extensive** none
MTU	MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked **Adjusted**.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Table 657: Link Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Route table	Routing table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**

Sample Output

show interfaces extensive (MFR UNI NNI)

```
user@host> show interfaces ls-1/3/0:0 extensive
Physical interface: ls-1/3/0:0, Enabled, Physical link is Up
  Interface index: 25, SNMP ifIndex: 35, Generation: 124
  Link-level type: Multilink-FR-UNI-NNI, MTU: 1524
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2002-11-01 15:26:25 PST (00:34:49 ago)
  Statistics last cleared: Never
  Link flags     : None
```

```
Hold-times     : Up 0 ms, Down 0 ms
Multilink Frame Relay UNI NNI bundle options:
  Device type                       DTE
  MRRU                              1524
  Fragmentation threshold           1500
  Red differential delay limit      10
  Yellow differential delay limit   6
  Red differential delay action     Disable link transmit
  Reassembly drop timer             0
  Links needed to sustain bundle    1
  LIP Hello timer                   10
      Acknowledgement timer         4
      Acknowledgement retries       2
  Bundle class                      A
  LMI type                          Q.933 Annex A
    T391 LIV polling timer          10
    T392 polling verification timer 15
    N391 full status polling count  6
    N392 error threshold            3
    N393 monitored event count      4
Traffic statistics:
 Input  bytes  :                    0
 Output bytes  :                    0
 Input  packets:                    0
 Output packets:                    0
Multilink Frame Relay UNI NNI bundle errors:
  Packet drops                    0 (0 bytes)
  Fragment drops                  0 (0 bytes)
  MRRU exceeded                   0
  Exception events                0
Multilink Frame Relay UNI NNI bundle statistics
                   Frames        fps         Bytes          bps
  Fragments:
    Input :             0          0             0            0
    Output:           824          0         17304          320
  Packets:
    Input :             0          0             0            0
    Output:           824          0         17304          320
Multilink Frame Relay UNI NNI bundle links info:
    Active bundle links       4
    Removed bundle links      0
    Disabled bundle links     0
Multilink Frame Relay UNI NNI active bundle links statistics:
```

```
                    Frames        fps         Bytes          bps
  t1-0/2/0:0.0
    Input :             0          0             0            0
    Output:           206          0          4326           80
    Current differential delay       0.2 ms
    Recent high differential delay   3.8 ms
    Times over red diff delay        0
    Times over yellow diff delay     0
    LIP:add_lnk lnk_ack lnk_rej    hello hel_ack lnk_rem rem_ack
    Rcv:      2       2       0      206     207       0       0
    Xmt:      2       1       0      207     206       0       0
  t1-0/2/0:1.0
    Input :             0          0             0            0
    Output:           206          0          4326           80
    Current differential delay       0.2 ms
    Recent high differential delay   3.7 ms
    Times over red diff delay        0
    Times over yellow diff delay     0
    LIP:add_lnk lnk_ack lnk_rej    hello hel_ack lnk_rem rem_ack
    Rcv:      2       2       0      206     207       0       0
    Xmt:      2       1       0      207     206       0       0
  t1-0/2/0:2.0
    Input :             0          0             0            0
    Output:           206          0          4326           80
    Current differential delay       0.4 ms
    Recent high differential delay   3.8 ms
    Times over red diff delay        0
    Times over yellow diff delay     0
    LIP:add_lnk lnk_ack lnk_rej    hello hel_ack lnk_rem rem_ack
    Rcv:      2       2       0      206     207       0       0
    Xmt:      2       1       0      207     206       0       0
  t1-0/2/0:3.0
    Input :             0          0             0            0
    Output:           206          0          4326           80
    Current differential delay       0.3 ms
    Recent high differential delay   3.8 ms
    Times over red diff delay        0
    Times over yellow diff delay     0
    LIP:add_lnk lnk_ack lnk_rej    hello hel_ack lnk_rem rem_ack
    Rcv:      2       2       0      206     207       0       0
    Xmt:      2       1       0      207     206       0       0
Logical interface ls-1/3/0:0.0 (Index 5) (SNMP ifIndex 28) (Generation 10)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI
```

```
    Bandwidth: 622080kbps
    Bundle errors:
      Packet drops                    0 (0 bytes)
      Fragment drops                  0 (0 bytes)
      MRRU exceeded                   0
      Exception events                0
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Fragments:
        Input :             0          0             0            0
        Output:           824          0         17304          320
      Packets:
        Input :             0          0             0            0
        Output:           824          0         17304          320
    Link:
      t1-0/2/0:0.0
        Input :             0          0             0            0
        Output:           206          0          4326           80
      t1-0/2/0:1.0
        Input :             0          0             0            0
        Output:           206          0          4326           80
      t1-0/2/0:2.0
        Input :             0          0             0            0
        Output:           206          0          4326           80
      t1-0/2/0:3.0
        Input :             0          0             0            0
        Output:           206          0          4326           80
    Protocol inet, MTU: 1500 [Adjusted], Generation: 15, Route table: 0
      Flags: User-MTU, MTU-Protocol-Adjusted
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2.1, Local: 192.0.2.2, Broadcast: Unspecified,
        Generation: 10
```

show interfaces extensive (MFR End-to-End)

```
user@host> show interfaces ls-0/3/0 extensive
Physical interface: ls-0/3/0, Enabled, Physical link is Up
  Interface index: 264, SNMP ifIndex: 104, Generation: 525
  Link-level type: LinkService, MTU: 1524
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
```

```
  Last flapped   : 2002-10-16 17:53:49 PDT (00:22:00 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                73471                  264 bps
   Output bytes  :                80335                    0 bps
   Input  packets:                  822                    0 pps
   Output packets:                  819                    0 pps
  Frame exceptions:
    Oversized frames                 0
    Errored input frames             0
    Input on disabled link/bundle    0
    Output for disabled link/bundle  4
    Queuing drops                    3
  Buffering exceptions:
    Packet data buffer overflow      0
    Fragment data buffer overflow    0
  Assembly exceptions:
    Fragment timeout                 0
    Missing sequence number          0
    Out-of-order sequence number     0
    Out-of-range sequence number     0
  Hardware errors (sticky):
    Data memory error                0
    Control memory error             0
  Logical interface ls-0/3/0.0 (Index 5) (SNMP ifIndex 527) (Generation 47)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
    Bandwidth: 1536kbps
    Bundle options:
      MRRU                           1524
      Drop timer period              0
      Sequence number format         long (24 bits)
      Fragmentation threshold        0
      Links needed to sustain bundle 1
      Interleave fragments           Disabled
    Bundle status:
      Remote MRRU                    1500
      Received sequence number       0x19ec14
      Transmit sequence number       0x38cfa8
      Packet drops                   0 (0 bytes)
      Fragment drops                 0 (0 bytes)
      MRRU exceeded                  0
      Fragment timeout               0
      Missing sequence number        0
```

```
    Out-of-order sequence number    0
    Out-of-range sequence number    0
    Packet data buffer overflow     0
    Fragment data buffer overflow   0
  Bundle errors:
    Packet drops                    2 (68 bytes)
    Fragment drops                  0 (0 bytes)
    MRRU exceeded                   0
    Exception events                0
  Statistics         Frames        fps         Bytes          bps
  Bundle:
    Fragments:
      Input :           172          0         15544          288
      Output:           165          0         16645            0
    Packets:
      Input :           143          0         12885          288
      Output:           134          0         12276            0
  Link:
    t1-0/0/0.0
      Input :           143          0         12885          288
      Output:           134          0         12276            0
  Protocol inet, MTU: 1500, Generation: 76, Route table: 0
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.16.1.2, Local: 10.16.1.1, Broadcast:
      Unspecified, Generation: 81
  Protocol iso, MTU: 1500 [Adjusted], Generation: 77, Route table: 0
    Flags: Is-Primary
  Protocol inet6, MTU: 1500, Generation: 78, Route table: 0
    Flags: Is-Primary
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: 2001:db8::1:0/126, Local: 2001:db8::1:1,
      Broadcast: Unspecified, Generation: 83
    Addresses, Flags: Is-Preferred
      Destination: 2001:db8::/64, Local: 2001:db8::2a0:a5ff:fe12:4777,
      Broadcast: Unspecified,
      Generation: 85
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces (Link Services IQ)

IN THIS SECTION

- Syntax | 20987
- Description | 20987
- Options | 20987
- Additional Information | 20988
- Required Privilege Level | 20988
- Output Fields | 20988
- Sample Output | 21011
- Release Information | 21023

Syntax

```
show interfaces lsq-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<l2-statistics>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M Series, MX Series, and T Series routers only) Display status information about the specified link services intelligent queuing (IQ) interface.

Options

lsq-*fpc/pic/port*	Display standard status information about the specified link services IQ interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.

descriptions	(Optional) Display interface description strings.
l2-statistics	(Optional) Display Layer 2 queue statistics for Multilink Point-to-Point Protocol (MLPPP), FRF.15, and FRF.16 bundles.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Additional Information

Link services IQ interfaces are similar to link services interfaces. The important difference is that link services IQ interfaces fully support Junos OS class-of-service (CoS) components.

Required Privilege Level

view

Output Fields

Table 658 on page 20988 lists the output fields for the `show interfaces` (link services IQ) command. Output fields are listed in the approximate order in which they appear.

Table 658: show interfaces (Link Services IQ) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical interface`	Name of the physical interface.	All levels
`Enabled`	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
`Interface index`	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`SNMP ifIndex`	SNMP index number for the physical interface.	**detail extensive** none
`Generation`	Unique number for use by Juniper Networks technical support.	**detail extensive**
`Link-level type`	Encapsulation being used on the physical interface: `Multilink-Frame-Relay-UNI-NNI` `Multilink-Frame-Relay-UNI-NNI` (default), `LinkService`, `Frame-relay`, `Frame-relay-ccc`, or `Frame-relay-tcc`.	All levels
`MTU`	Maximum transmission unit size on the physical interface.	All levels
`Device flags`	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
`Interface flags`	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second timezone`* (*`hour`*`:`*`minute`*`:`*`second`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Multilink Frame Relay UNI NNI bundle options`	(Multilink Frame Relay UNI NNI only) Configured information about Multilink Frame Relay bundle options. • `Device type`—DCE (data communication equipment) or DTE (data terminal equipment). • `MRRU`—Configured size of the maximum received reconstructed unit (MRRU): 1500 to 4500 bytes. The default is 1524 bytes. • `Bandwidth`—Speed at which the interface is running. • `Fragmentation threshold`—Configured fragmentation threshold: 128 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. • `Red differential delay limit`—Red differential delay limit among bundle links has been reached, indicating an action will occur. • `Yellow differential delay limit`—Yellow differential delay among bundle links has been reached, indicating a warning will occur. • `Red differential delay action`—Type of actions taken when the red differential delay exceeds the red limit: *Disable link transmit* or *Remove link from service*. • `Link layer overhead`—Percentage of bundle bandwidth to be set aside for link layer overhead. • `Reassembly drop timer`—Drop timeout value to provide a recovery mechanism if individual links in the link services bundle drop one or more packets: 1 through 127 milliseconds. By default, the drop timeout parameter is 0 (disabled). A value under 5 ms is not recommended. • `Links needed to sustain bundle`—Minimum number of links to sustain the bundle: 1 through 8. • `LIP Hello timer`—Link Interleaving Protocol hello timer: 1 through 180 seconds.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Acknowledgement timer`—Maximum period to wait for an add link acknowledgement, hello acknowledgement, or remove link acknowledgement: 1 through 10 seconds. • `Acknowledgement retries`—Number of retransmission attempts to be made for consecutive hello or remove link messages after the expiration of the acknowledgement timer: 1through 5.	

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Multilink Frame Relay UNI NNI bundle options` (continued)	• `Bundle class`—Bundle class ID.
	• `LMI type`—Multilink Frame Relay UNI NNI LMI type: `ANSI`, `Q.933 ANNEX A`, or `Consortium`.
	  • `T391 LIV polling timer`—Multilink Frame Relay UNI NNI Full status polling counter: 1 through 255, with a default value of 6.
	  • `T392 polling verification timer`—Multilink Frame Relay UNI NNI LMI error threshold. The number of errors required to bring down the link, within the event count specified by *N393*. The range is 1 through 10, with a default value of 3.
	  • `N391 full status polling count`—Multilink Frame Relay UNI NNI Full status polling counter: 1 through 255.
	  • `N392 error threshold`—Multilink Frame Relay UNI NNI LMI error threshold: 1 through 10.
	• `N393 monitored event count`—Multilink Frame Relay UNI NNI LMI monitored event count: 1 through 10, with a default value of 4.
	• `Consortium LMI Settings`
	  • `n391dte`—DTE full status polling interval in seconds: 1 through 255.
	  • `n392dce`—DCE error threshold: 1 through 10.
	  • `n392dte`—DTE error threshold: 1 through 10.
	  • `n393dce`—DCE monitored event count: 1 through 10.
	  `n393dte`—DTE monitored event count: 1 through 10.
	  • `t391dte`—DTE polling verification timer (in seconds): 5 through 30.
	  • `t392dce`—DCE polling verification timer (in seconds): 5 through 30.
		detail extensive
		none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`LMI`	Local Managment Interface packet statistics:
	• `Input`—Number of packets arriving on the interface (nn) and timestamp of the most recent packet arrival, in the format:
	`Input: nn (last seen hh:mm:ss ago)`
	`Output`—Number of packets sent out on the interface (nn) and how much time has passed since the last packet was sent, in the format:
	`Output: nn (last seen hh:mm:ss ago)`
		detail extensive
		none

`DTE Statistics`	Statistics about information transferred from the data terminal equipment (DTE) to the data communications equipment (DCE). • `Enquiries sent`—Number of link status enquiries sent from the DTE to the DCE. • `Full enquiries sent`—Number of full enquiries sent from the DTE to the DCE. • `Enquiry responses received`—Number of enquiry responses received by the DCE from the DTE. • `Full enquiry responses received`—Number of full enquiry responses received by DCE from the DTE.	**detail extensive** none
`DCE Statistics`	Statistics about information transferred from the DCE to the DTE. • `Enquiries received`—Number of enquiries received by the DCE from the DTE. • `Full enquiries received`—Number of full enquiries received by the DCE from the DTE. • `Enquiry responses sent`—Number of enquiry responses sent from the DCE to the DTE. • `Full enquiry responses sent`—Number of full enquiry responses sent from the DCE to the DTE.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Common Statistics`	Statistics about messages snet between the DTE and the DCE. • `Unknown messages received`—Number of received packets that do not fall into any other category. • `Asynchronouts updates received`—Number of link status peer changes received. • `Out-of-sequence packets received`—Number of packets for which the sequence of the packets received is different from the expected sequence. • `Keepalive responses timed out`—Number of keepalive responses that time out when no Local Management Interface (LMI) packet was reported for `n392dte` or `n393dce` intervals. (See *LMI settings*.)	
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the Packet Forwarding Engine (PFE). Input traffic refers to the fragments received by the ingress PFE, which get assembled into Layer 3 input packets. Output packets refer to the IP packets transmitted out of the ingress PFE to the LSQ, which get segmented into output fragments.	**detail extensive**

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`DLCI`*nn*	Data-link connection identifier (DLCI) number of the logicial interview. The following information is displayed.
	• `Flags`—Values are:
	  • `Active`—Set when the link is active and the DTE and DCE are exchanging information.
	  • `Down`—Set when the link is active, but no information is received from the DTE.
	  • `DCE unconfigured`—Set when the corresponding DCLI in the DCE is not configured.
	  • `Configured`—Set when the correspondening DLCCI is configured.
	  • `DCE-Configured`—Displayed when the command is issued from the DTE.	

`DLCI Statistics`	(Frame Relay) Data-link connection identifier (DLCI) statistics. • `Active DLCI`—Number of active DLCIs. • `Inactive DLCI`—Number of inactive DLCIs.	
`Input rate`	(Redundant LSQ) Rate of bits and packets received on the interface.	None specified
`Output rate`	(Redundant LSQ) Rate of bits and packets transmitted on the interface.	None specified
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	**detail extensive**
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.	**detail extensive**

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Frame exceptions`	Information about framing exceptions. Includes events recorded under `Exception Events` for each logical interface. • `Oversized frames`—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits). • `Errored input frames`—Number of input frame errors. • `Input on disabled link/bundle`—Number of frames received on disabled links. These frames can result either from an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it. • `Output for disabled link/bundle`—Number of frames sent for a disabled or unavailable link. These frames can result either from an inconsistent configuration, or from a bundle being brought up or down while traffic is flowing through it. • `Queuing drops`—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed.	**extensive**
`Buffering exceptions`	Information about buffering exceptions. Includes events recorded under `Exception Events` for each logical interface: • `Packet data buffer overflow`—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • `Fragment data buffer overflow`—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. Check the logical interface exception event counters to determine which bundle is responsible.	**extensive**

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Assembly exceptions`	(Multilink Frame Relay end-to-end only) Information about assembly exceptions. Includes events recorded under `Exception Events` for each logical interface. An assembly exception does not necessarily indicate an operational problem with the physical link services IQ interface itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the interface records one or more exception events. The physical interface can drop multilink-encapsulated fragments itself as a result. Any multilink packets or fragments dropped by the interface itself result in packet or fragment drop counts on individual logical interfaces. If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, excess differential delay or traffic losses on individual interfaces can be the root cause. • `Fragment timeout`—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • `Missing sequence number`—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • `Out-of-order sequence number`—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the link services IQ interface unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible.	**extensive**

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Out-of-range sequence number`—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible.	
`Hardware errors (sticky)`	(Multilink Frame Relay end-to-end only) Information about hardware errors: • `Data memory error`—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. • `Control memory error`—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support.	**extensive**
`Egress queues`	Total number of egress queues supported on the specified interface.	**detail extensive** none
`Queue counters`	Queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive** none
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`SNMP ifIndex`	Logical interface SNMP interface index number.	**detail extensive** none
`Generation`	Unique number for use by Juniper Networks technical support.	**detail extensive**
`Flags`	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
`Encapsulation`	Encapsulation being used: PPP or Multilink PPP.	All levels
`Bandwidth`	Speed at which the interface is running.	All levels
`Bundle options`	(Multilink Frame Relay end-to-end interfaces only) • `MRRU`—Configured size of the maximum received reconstructed unit (MRRU): 1500 through 4500 bytes. The default is 1504 bytes. • `Drop timer period`—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: 0 though 2000 milliseconds. Values under 5 ms are not recommended. The default setting is 0, which disables the timer. • `Sequence number format`—Short sequence number header format (MLPPP only). • `Fragmentation threshold`—Configured fragmentation threshold: 64 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. • `Links needed to sustain bundle`—Minimum number of links to sustain the bundle: 1 through 8. • `Multilink classes`—Number of multilink classes negotiated. • `Link layer overhead`—Percentage of bundle bandwidth to be set aside for link-layer overhead.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Bundle status` (MLPPP) or `Multilink class status` (Multiclass MLPPP)	Information about bundle status: • `Remote MRRU`—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. • `Received sequence number`—Sequence number for received packets. • `Transmitted sequence number`—Sequence number for transmitted packets. • `Packet drops`—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • `Fragment drops`—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully, but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • `MRRU exceeded`—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • `Fragment timeout`—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • `Missing sequence number`—A gap was detected in the sequence numbers of fragments on a bundle. • `Out-of-order sequence number`—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Out-of-range sequence number`—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • `Packet data buffer overflow`—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • `Fragment data buffer overflow`—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity.	

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Statistics`	Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).
	• `Bundle`—Information for each active bundle link.
	  • `Fragments: Input and Output`—Total number and rate of fragments received and transmitted.
	  • `Packets: Input and Output`—Total number and rate of packets received and transmitted.
	  • `Multilink class`—(Multiclass MLPPP only) Information about multiclass links used in the multilink operation.
	• `Link`—Information about links used in the multilink operation.
	  • `Link name`—Interface name of the link services IQ channel and state information (physical link `up` or `down`).
	  • `Input and Output`—Total number and rate of fragments and packets received and transmitted.
		detail extensive

`NCP state`	(PPP) Network Control Protocol state. • `Conf-ack-received`—Acknowledgement was received. • `Conf-ack-sent`—Acknowledgement was sent. • `Conf-req-sent`—Request was sent. • `Down`—NCP negotiation is incomplete (not yet completed or has failed). • `Not-configured`—NCP is not configured on the interface. • `Opened`—NCP negotiation is successful.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Protocol`	Protocol family configured on the logical interface.	**detail extensive** none
`MTU`	MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked `Adjusted`.	**detail extensive** none
`Generation`	Unique number for use by Juniper Networks technical support only.	**detail extensive**
`Route Table`	Routing table in which this address exists. For example, `Route table:0` refers to inet.0.	**detail extensive**
`Flags`	Information about the protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	**detail extensive** none
`Addresses, Flags`	Information about the addresses configured on the logical interface. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	**detail extensive** none
`Destination`	IP address of the remote side of the connection.	**detail extensive** none
`Local`	IP address of the logical interface.	**detail extensive** none
`Broadcast`	Broadcast address on the logical interface.	**detail extensive** none
`Generation`	Unique number for use by Juniper Networks technical support.	**detail extensive**
MLPPP Bundle Interface		

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Logical interface`	Name of the logical interface.	All levels
`Index`	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
`SNMP ifIndex`	Logical interface SNMP interface index number.	**detail extensive** none
`Generation`	Unique number for use by Juniper Networks technical support.	**detail extensive**
`Flags`	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
`SNMP-Traps`	SNMP trap notifications are enabled.	All levels
`Encapsulation`	Encapsulation being used: PPP, Multilink PPP, or Multilink-FR.	All levels
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second timezone`* (*`hour`*`:`*`minute`*`:`*`second`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	**detail extensive** none
`Bandwidth`	Speed at which the interface is running.	All levels
`Bundle links Iinformation`	Information about the bundled links. • `Active bundle links`—Number of active links. • `Removed bundle links`—Information about links used in the multilink operation. • `Disabled bundle links`—Number of disabled links.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Bundle options`	(Multilink Frame Relay end-to-end interfaces only) • `MRRU`—Configured size of the maximum received reconstructed unit (MRRU): 1500 through 4500 bytes. The default is 1504 bytes. • `Drop timer period`—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: 0 though 2000 milliseconds. Values under 5 ms are not recommended. The default setting is 0, which disables the timer. • `Inner PPP Protocol field compression`—Inner PPP protocol compression is enabled or disabled. • `Sequence number format`—Short sequence number header format (MLPPP only). • `Fragmentation threshold`—Configured fragmentation threshold: 64 through 16,320 bytes, in integer multiples of 64 bytes. The default setting is 0, which disables fragmentation. • `Links needed to sustain bundle`—Minimum number of links to sustain the bundle: 1 through 8. • `Multilink classes`—Number of multilink classes negotiated. • `Link layer overhead`—Percentage of bundle bandwidth to be set aside for link-layer overhead.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Bundle status` (MLPPP)	Information about bundle status: • `Received sequence number`—Sequence number for received packets. • `Transmit sequence number`—Sequence number for transmitted packets. • `Packet drops`—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • `Fragment drops`—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • `MRRU exceeded`—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • `Fragment timeout`—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • `Missing sequence number`—A gap was detected in the sequence numbers of fragments on a bundle. • `Out-of-order sequence number`—Two frames with out-of-order sequence numbers occurred within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Out-of-range sequence number`—A frame was received with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • `Packet data buffer overflow`—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • `Fragment data buffer overflow`—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity.	

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Statistics`	Information about frames, bytes, and bits per second received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).
	
	The bundle, multilink, and network statistics are reported by the Packet Forwarding Engine (PFE). The Multi Link Detail statistics like fragments, non-fragments and LFI are reported by the PIC.
	
	However, the PFE reports an extra overhead of 2 bytes in the output when compared with the Multilink Detail Statistics. This is due to the service-cookie in the PFE which does the link demux for the ML header.
	
	The difference in the bytes received and transmitted from Network and Multilink interfaces and Multilink statistics for each member link is divided between the ML and the PPP headers. For example the header counter for a long sequence configuration would be as follows.
	
	• Input side - Total overhead = 6 bytes.
	  • ML: 4 bytes of ML header = 1 byte of Flag + 3 bytes of long sequence number.
	  • PPP: 2 bytes of protocol field.
	• Output side - Total overhead = 11 bytes.
	  • ML: 4 bytes of ML Header = 1 byte of Flag + 3 bytes of Long sequence number.
	  • PPP: 5 bytes = 4 bytes of header + 1 byte of Idle flag.
	  • 2 bytes of Service Cookie.
	• `Bundle`—Information for each active bundle link.
	  • `Multilink: Input and Ouput`—Total number and rate of multilink frames, bytes, and bits per second received and transmitted. It is a module connecting LSQ PIC and its member link. Multilink
		extensive

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	Input displays L2 fragments received from the member link to the LSQ PIC. Multilink Output displays the L2 fragments transmitted from LSQ PIC to the member links. • `Network: Input and Output`—Total number of network frames, bytes, and bits per second received and transmitted. It refers to the packets transmitted from an ingress interface to the PFE and then to the LSQ PIC. Network Input displays the L3 packets received from the LSQ PIC to the PFE. Network Output displays the L3 packets transmitted from PFE to LSQ PIC. • `Link`—Information about links used in the multilink operation. • `Link name`—The interface name of the link services IQ channel and state information (physical link *up* or *down*) and up time. • `Input and Output`—Total number and rate of frames, bytes, and bits per second received and transmitted.	

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Multilink detail statistics`	Frames, bytes, and bits per second received and sent by the bundle. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).
	
	The difference in the bytes received and transmitted from the bundle is divided between the ML and the PPP headers. For example the header counter for a long sequence configuration would be as follows:
	
	• Input side - Total overhead = 6 bytes.
	  • ML: 4 bytes of ML header = 1 byte of Flag + 3 bytes of long sequence number.
	  • PPP: 2 bytes of protocol field.
	• Output side - Total overhead = 9 bytes.
	  • ML: 4 bytes of ML Header = 1 byte of Flag + 3 bytes of Long sequence number.
	  • PPP: 5 bytes = 4 bytes of header + 1 byte of Idle flag.
	• `Bundle`—Information for the bundle link.
	  • `Fragments: Input and Output`—Total number and rate of multilink fragments received and transmitted.
	  • `Non-fragments: Input and Output`—Total number and rate of nonfragmented multilink frames received and transmitted.
	  • `LFI: Input and Output`—Total number and rate of link fragmented and interleaved frames and bytes.
		extensive

`Protocol`	Protocol family configured on the logical interface.	**detail extensive** none
`MTU`	MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked *Adjusted*.	**detail extensive** none

Table 658: show interfaces (Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Generation`	Unique number for use by Juniper Networks technical support only.	**detail extensive**
`Route Table`	Routing table in which this address exists. For example, `Route table:0` refers to inet.0.	**detail extensive**
`Addresses, Flags`	Information about the addresses configured on the logical interface. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	**detail extensive** none
`Destination`	IP address of the remote side of the connection.	**detail extensive** none
`Local`	IP address of the logical interface.	**detail extensive** none
`Broadcast`	Broadcast address on the logical interface.	**detail extensive** none
`Generation`	Unique number for use by Juniper Networks technical support.	**detail extensive**

Sample Output

show interfaces extensive (MLPPP on Link Services IQ)

```
user@host> show interfaces lsq-0/2/0 extensive
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
  Interface index: 140, SNMP ifIndex: 25, Generation: 23
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2005-06-02 08:54:36 PDT (00:05:45 ago)
  Statistics last cleared: Never
  Traffic statistics:
```

```
 Input  bytes  :              8872424               229080 bps
 Output bytes  :              9856960               234448 bps
 Input  packets:                38202                  117 pps
 Output packets:                39453                  117 pps
Frame exceptions:
  Oversized frames                  0
  Errored input frames              0
  Input on disabled link/bundle     0
  Output for disabled link/bundle   0
  Queuing drops                     0
Buffering exceptions:
  Packet data buffer overflow       0
  Fragment data buffer overflow     0
Assembly exceptions:
  Fragment timeout                  0
  Missing sequence number           0
  Out-of-order sequence number      0
  Out-of-range sequence number      0
Hardware errors (sticky):
  Data memory error                 0
  Control memory error              0
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 be                             0                    0                    0
  1 ef                             0                    0                    0
  2 af                             0                    0                    0
  3 nc                             0                    0                    0
Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
  Bandwidth: 256kbps
  Bundle options:
    MRRU                            1504
    Drop timer period               2000
    Sequence number format          long (24 bits)
    Fragmentation threshold         0
    Links needed to sustain bundle  1
    Multilink classes               0
    Link layer overhead             4.0 %
  Bundle status:
    Remote MRRU                     1500
    Received sequence number        0x0
    Transmit sequence number        0x0
    Packet drops                    0 (0 bytes)
    Fragment drops                  9 (1401 bytes)
```

```
      MRRU exceeded                   0
      Fragment timeout                0
      Missing sequence number         0
      Out-of-order sequence number    4
      Out-of-range sequence number    0
      Packet data buffer overflow     0
      Fragment data buffer overflow   0
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Multilink:
        Input :         79827        239       9593009       232288
        Output:         77533        234       9811743       238056
      Network:
        Input :         38202        117       8872424       229080
        Output:         39453        117       9856960       234448
    Link:
      ds-1/0/2:1:1.0 <-- up
        Input :          1114         87        180183       113608
        Output:          1577        118        199215       119064
      ds-1/0/2:1:2.0 <-- down
        Input :          1941        152        187948       118680
        Output:          1574        116        199494       118992
    Protocol inet, MTU: 1500 [Adjusted]
      Flags: User-MTU, MTU-Protocol-Adjusted
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.74.11/24, Local: 10.74.11.10
    Protocol iso, MTU: 1500 [Adjusted]
      Flags: User-MTU, MTU-Protocol-Adjusted
    Protocol mpls, MTU: 1488 [Adjusted], Maximum labels: 3
      Flags: User-MTU, MTU-Protocol-Adjusted
```

show interfaces extensive (Multiclass MLPPP on Link Services IQ)

```
user@host> show interfaces extensive lsq-0/2/0
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
  Interface index: 140, SNMP ifIndex: 25, Generation: 23
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2005-06-02 08:54:36 PDT (00:02:25 ago)
  Statistics last cleared: Never
```

```
Traffic statistics:
 Input  bytes  :              3474024                223704 bps
 Output bytes  :              4193992                233888 bps
 Input  packets:                15809                   116 pps
 Output packets:                16788                   117 pps
Frame exceptions:
  Oversized frames                  0
  Errored input frames              0
  Input on disabled link/bundle     0
  Output for disabled link/bundle   0
  Queuing drops                     0
Buffering exceptions:
  Packet data buffer overflow       0
  Fragment data buffer overflow     0
Assembly exceptions:
  Fragment timeout                  0
  Missing sequence number           0
  Out-of-order sequence number      0
  Out-of-range sequence number      0
Hardware errors (sticky):
  Data memory error                 0
  Control memory error              0
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 be                             0                    0                    0
  1 ef                             0                    0                    0
  2 af                             0                    0                    0
  3 nc                             0                    0                    0
Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
  Bandwidth: 256kbps
  Bundle options:
    MRRU                            1504
    Drop timer period               2000
    Sequence number format          long (24 bits)
    Fragmentation threshold         0
    Links needed to sustain bundle  1
    Multilink classes               2
    Link layer overhead             4.0 %
  Multilink class 0 status:
    Received sequence number        0x4c38
    Transmit sequence number        0x4890
    Packet drops                    0 (0 bytes)
    Fragment drops                  2551 (397084 bytes)
```

```
    MRRU exceeded                   0
    Fragment timeout                52
    Missing sequence number         0
    Out-of-order sequence number    953
    Out-of-range sequence number    0
    Packet data buffer overflow     0
    Fragment data buffer overflow   0
  Multilink class 1 status:
    Received sequence number        0xffffff
    Transmit sequence number        0x3710
    Packet drops                    0 (0 bytes)
    Fragment drops                  0 (0 bytes)
    MRRU exceeded                   0
    Fragment timeout                0
    Missing sequence number         0
    Out-of-order sequence number    0
    Out-of-range sequence number    0
    Packet data buffer overflow     0
    Fragment data buffer overflow   0
  Statistics         Frames        fps         Bytes          bps
  Bundle:
    Fragments:
      Input :         33719        239       4041763       231632
      Output:         32371        234       4096545       237488
    Packets:
      Input :         15809        116       3474024       223704
      Output:         16788        117       4193992       233888
  Multilink class 0:
    Fragments:
      Input :         19331          0             0            0
      Output:             0          0             0            0
    Packets:
      Input :          2064          0             0            0
      Output:          1864          0             0            0
  Multilink class 1:
    Fragments:
      Input :             0          0             0            0
      Output:         14096          0             0            0
    Packets:
      Input :         14096          0             0            0
      Output:             0          0             0            0
  Link:
    ds-1/0/2:1:1.0, Enabled, Physical link is Up
```

```
        Input :         20972        151       2030595       118080
        Output:         16184        116       2048468       118488
      ds-1/0/2:1:2.0, Enabled, Physical link is Up
        Input :         12747         88       2011168       113552
        Output:         16187        118       2048077       119000
    Protocol inet, MTU: 1500 [Adjusted], Generation: 14, Route table: 0
      Flags: User-MTU, MTU-Protocol-Adjusted
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast: Unspecified,
        Generation: 18
```

show interfaces extensive (MLPPP on Link Services IQ Bundle)

```
user@host> show interfaces lsq-7/1/0.0 extensive
Logical interface lsq-7/1/0.0 (Index 88) (SNMP ifIndex 114) (Generation 188)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR
    Last flapped: Never
    Bandwidth: 256kbps
    Bundle links information:
      Active bundle links       2
      Removed bundle links      0
      Disabled bundle links     0
    Bundle options:
      MRRU                                  1504
      Drop timer period                     1500
      Inner PPP Protocol field compression enabled
      Sequence number format                short (12 bits)
      Fragmentation threshold               0
      Links needed to sustain bundle        1
      Multilink classes                     0
      Link layer overhead                   4.0 %
    Bundle status:
      Received sequence number        0xb74
      Transmit sequence number        0xb74
      Packet drops                    0 (0 bytes)
      Fragment drops                  0 (0 bytes)
      MRRU exceeded                   0
      Fragment timeout                0
      Missing sequence number         0
      Out-of-order sequence number    0
      Out-of-range sequence number    0
```

```
      Packet data buffer overflow     0
      Fragment data buffer overflow   0
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Multilink:
        Input :        315381          0      42757818            0
        Output:        315381          0      43388580            0
      Network:
        Input :        315381          0      40952064            0
        Output:        315381          0      40952064            0
    Link:
      ds-6/0/0:1:1.0
        Up time: Up since boot
        Input :         63794          0      25146728            0
        Output:         63778          0      25273164            0
      ds-6/0/0:1:2.0
        Up time: Up since boot
        Input :        251587          0      17611090            0
        Output:        251603          0      18115416            0
    Multilink detail statistics:
    Bundle:
      Fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Non-fragments:
        Input :        293748          0      19387368            0
        Output:        293748          0      20562360            0
      LFI:
        Input :         21633          0      22152192            0
        Output:         21633          0      22325256            0
    Protocol inet, MTU: 1500, Generation: 204, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast:
Unspecified, Generation: 214
```

show interfaces extensive (MFR on Link Services IQ Bundle)

```
user@host> show interfaces lsq-1/0/0:0 extensive
Physical interface: lsq-1/0/0:0, Enabled, Physical link is Up
  Interface index: 179, SNMP ifIndex: 746, Generation: 182
```

```
Link-level type: Multilink-FR-UNI-NNI, MTU: 1508
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Last flapped   : 2010-11-15 01:11:00 PST (00:31:58 ago)
Statistics last cleared: Never
Hold-times     : Up 0 ms, Down 0 ms
Multilink Frame Relay UNI NNI bundle options:
  Device type                       DCE
  MRRU                              1508
  Bandwidth                         1536kbps
  Fragmentation threshold           0
  Red differential delay limit      120
  Yellow differential delay limit   72
  Red differential delay action     Remove link
  Reassembly drop timer             65535
  Links needed to sustain bundle    1
  Link layer overhead               4.0 %
  LIP Hello timer                   10
      Acknowledgement timer         4
      Acknowledgement retries       2
  Bundle class                      A
  LMI type                          Consortium
    T391 LIV polling timer          10
    T392 polling verification timer 15
    N391 full status polling count  6
    N392 error threshold            3
    N393 monitored event count      4
Consortium LMI settings: n392dce 3, n393dce 4, t392dce 15 seconds
LMI statistics:
  Input : 188 (last seen 00:00:01 ago)
  Output: 189 (last sent 00:00:01 ago)
DTE statistics:
  Enquiries sent                      : 0
  Full enquiries sent                 : 0
  Enquiry responses received          : 0
  Full enquiry responses received     : 0
DCE statistics:
  Enquiries received                  : 157
  Full enquiries received             : 31
  Enquiry responses sent              : 158
  Full enquiry responses sent         : 31
Common statistics:
  Unknown messages received           : 0
```

```
  Asynchronous updates received        : 0
  Out-of-sequence packets received     : 0
  Keepalive responses timedout         : 0
Traffic statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :                    0                    0 bps
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
 IPv6 transit statistics:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Multilink Frame Relay UNI NNI bundle errors:
  Packet drops                    0 (0 bytes)
  Fragment drops                  0 (0 bytes)
  MRRU exceeded                   0
  Exception events                0
Multilink Frame Relay UNI NNI bundle statistics:
                    Frames        fps         Bytes          bps


  Multilink:
    Input :             0          0             0            0
    Output:             0          0             0            0
  Network:
    Input :             0          0             0            0
    Output:             0          0             0            0
Multilink Frame Relay UNI NNI bundle links information:
    Active bundle links       1
    Removed bundle links      0
    Disabled bundle links     0
Multilink Frame Relay UNI NNI active bundle links statistics:
                    Frames        fps         Bytes          bps
  t1-7/0/0:1:3.0
    Up time: 00:31:24
    Input :             0          0             0            0
    Output:             0          0             0            0
    Current differential delay       0.0 ms
    Recent high differential delay   0.0 ms
    Times over red diff delay        0
    Times over yellow diff delay     0
    LIP:add_lnk lnk_ack lnk_rej    hello hel_ack lnk_rem rem_ack
    Rcv:      2       2       0        0     189       0       0
```

```
    Xmt:      2       1       0      189       0       0       0

Logical interface lsq-1/0/0:2.0 (Index 77) (SNMP ifIndex 751) (Generation 142)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI
  Last flapped: 2010-11-15 01:11:40 PST (00:31:18 ago)
  Bundle status:
    Received sequence number        0xfff
    Transmit sequence number        0x0
    Packet drops                    0 (0 bytes)
    Fragment drops                  0 (0 bytes)
    MRRU exceeded                   0
    Fragment timeout                0
    Missing sequence number         0
    Out-of-order sequence number    0
    Out-of-range sequence number    0
    Packet data buffer overflow     0
    Fragment data buffer overflow   0
  Statistics         Frames        fps         Bytes          bps
  Bundle:
    Multilink:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Network:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Link:
    t1-7/0/0:1:3.0
      Up time: 00:31:24
      Input :             0          0             0            0
      Output:             0          0             0            0
  Multilink detail statistics:
  Bundle:
    Fragments:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Non-fragments:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Protocol inet, MTU: 1500, Generation: 153, Route table: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.0.1.8/30, Local: 10.0.1.9, Broadcast: Unspecified, Generation: 154
  DLCI 12
```

```
      Flags: Active
      Total down time: 00:00:32 sec, Last down: 00:31:50 ago
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
  DLCI statistics:
    Active DLCI  :1  Inactive DLCI  :0
```

show interfaces extensive (Multiclass MLPPP on Link Services IQ)

```
user@host> show interfaces extensive lsq-0/2/0
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
  Interface index: 140, SNMP ifIndex: 25, Generation: 23
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2005-06-02 08:54:36 PDT (00:02:25 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              3474024               223704 bps
   Output bytes  :              4193992               233888 bps
   Input  packets:                15809                  116 pps
   Output packets:                16788                  117 pps
  Frame exceptions:
    Oversized frames                  0
    Errored input frames              0
    Input on disabled link/bundle     0
    Output for disabled link/bundle   0
    Queuing drops                     0
  Buffering exceptions:
    Packet data buffer overflow       0
    Fragment data buffer overflow     0
  Assembly exceptions:
    Fragment timeout                  0
    Missing sequence number           0
    Out-of-order sequence number      0
    Out-of-range sequence number      0
  Hardware errors (sticky):
    Data memory error                 0
```

```
    Control memory error              0
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 be                             0                    0                    0
    1 ef                             0                    0                    0
    2 af                             0                    0                    0
    3 nc                             0                    0                    0
  Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
    Bandwidth: 256kbps
    Bundle options:
      MRRU                            1504
      Drop timer period               2000
      Sequence number format          long (24 bits)
      Fragmentation threshold         0
      Links needed to sustain bundle  1
      Multilink classes               2
      Link layer overhead             4.0 %
    Multilink class 0 status:
      Received sequence number        0x4c38
      Transmit sequence number        0x4890
      Packet drops                    0 (0 bytes)
      Fragment drops                  2551 (397084 bytes)
      MRRU exceeded                   0
      Fragment timeout                52
      Missing sequence number         0
      Out-of-order sequence number    953
      Out-of-range sequence number    0
      Packet data buffer overflow     0
      Fragment data buffer overflow   0
    Multilink class 1 status:
      Received sequence number        0xffffff
      Transmit sequence number        0x3710
      Packet drops                    0 (0 bytes)
      Fragment drops                  0 (0 bytes)
      MRRU exceeded                   0
      Fragment timeout                0
      Missing sequence number         0
      Out-of-order sequence number    0
      Out-of-range sequence number    0
      Packet data buffer overflow     0
      Fragment data buffer overflow   0
    Statistics         Frames        fps         Bytes          bps
    Bundle:
```

```
    Fragments:
      Input :         33719        239       4041763       231632
      Output:         32371        234       4096545       237488
    Packets:
      Input :         15809        116       3474024       223704
      Output:         16788        117       4193992       233888
  Multilink class 0:
    Fragments:
      Input :         19331          0             0            0
      Output:             0          0             0            0
    Packets:
      Input :          2064          0             0            0
      Output:          1864          0             0            0
  Multilink class 1:
    Fragments:
      Input :             0          0             0            0
      Output:         14096          0             0            0
    Packets:
      Input :         14096          0             0            0
      Output:             0          0             0            0
  Link:
    ds-1/0/2:1:1.0, Enabled, Physical link is Up
      Input :         20972        151       2030595       118080
      Output:         16184        116       2048468       118488
    ds-1/0/2:1:2.0, Enabled, Physical link is Up
      Input :         12747         88       2011168       113552
      Output:         16187        118       2048077       119000
  Protocol inet, MTU: 1500 [Adjusted], Generation: 14, Route table: 0
    Flags: User-MTU, MTU-Protocol-Adjusted
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast: Unspecified,
      Generation: 18
```

Release Information

Command introduced before Junos OS Release 7.4.

`l2-statistics` option introduced with Junos OS Release 12.1.

RELATED DOCUMENTATION

Link and Multilink Services Overview

Multilink Interfaces on Channelized MICs Overview

show interfaces load-balancing (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 21024
- Description | 21024
- Options | 21024
- Required Privilege Level | 21025
- Output Fields | 21025
- Sample Output | 21027
- Release Information | 21028

Syntax

```
show interfaces load-balancing
<detail>
<interface-name>
```

Description

Display information about the aggregated multiservices interface (AMS) as well as its individual member interfaces and the status of the replication state.

Options

none	Display standard information about status of all AMS interfaces.
detail	(Optional) Display detailed status of all AMS interfaces.

interface-name (Optional) Name of the aggregated multiservices interface (`ams`). If this is omitted, then the information for all the aggregated multiservices interfaces, including those used in control plane redundancy and high availability (HA) for service applications, is displayed.

Required Privilege Level

view

Output Fields

Table 659 on page 21025 lists the output fields for the `show interfaces load-balancing` (aggregated multiservices interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 659: Aggregated Multiservices show interfaces load-balancing Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the aggregated multiservices (AMS) interface.	`detail` none
`State`	Status of AMS interfaces: • `Coming Up`—Interface is becoming operational. • `Members Seen`—Member interfaces (mams) are available. • `Up`—Interface is configured and operational. • `Wait for Members`—Member interfaces (mams) are not available. • `Wait Timer`—Interface is waiting for member interfaces (mams) to come online.	`detail` none
`Last change`	Time (in *hh:mm:ss* [*hours:minutes:seconds*] format) when the state last changed.	`detail` none
`Members`	Number of member interfaces (mams-).	none specified

Table 659: Aggregated Multiservices show interfaces load-balancing Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Member count`	Number of member PICs (mams) that are part of the aggregated interface.	`detail` none
`HA Model`	High availability (HA) model supported on the interface. • `Many-to-One`—The preferred backup Multiservices PIC, in hot standby mode, backs up one or more (N) active Multiservices PICs. • `One-to-One`—–The preferred backup Multiservices PIC, in hot standby mode, backs up only one active Multiservices PIC. **NOTE**: `One-to-One` is not supported on MX-SPC3 cards.	`detail` none

`Members`	Information about the member interfaces:
	• `Interface`—Name of the member interface.
	• `Weight`—Not applicable for the current release.
	• `State`—State of the member interface (mams-).
	  • `Active`—Member is an active member.
	  • `Backup`—Member is a backup.
	  • `Discard`—Member has not yet rejoined the ams interface after failure.
	  • `Down`—Member has not yet powered on.
	  • `Inactive`—Member has failed to rejoin the ams interface within the configured `rejoin-timeout`.
	  • `Invalid`—Multiservices PIC corresponding to the member interface has been configured but is not physically present in the chassis.
		`detail`

Table 659: Aggregated Multiservices show interfaces load-balancing Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Sync-state`	Synchronization (sync) status of the control plane redundancy. The sync state is displayed only when the ams interface is `Up`.
	• `Interface`—Name of the member interface.
	• `Status`—Synchronization status of the member interfaces.
	  • `In progress`—The active member is currently synchronizing its state information with the backup member.
	  • `In sync`—The active member has finished synchronizing its state information with the backup and the backup is ready to take over if the active member fails.
	  • `NA (Not applicable)`—The backup member is not yet ready to synchronize with the active (primary) member. This condition may occur if the backup is still powered off or still booting.
	  • `Unknown`—The daemons are still initializing and the state information is unavailable.
		`detail`

Sample Output

show interfaces load-balancing

```
user@host> show interfaces load-balancing
Interface  State       Last change  Members  HA Model
ams0       Up          00:10:02     4        Many-to-One
```

show interfaces load-balancing detail

```
user@host> show interfaces load-balancing detail
Load-balancing interfaces detail
Interface        : ams0
  State          : Up
  Last change    : 00:10:23
```

```
  Member count   : 4
  HA Model       : Many-to-One
  Members        :
      Interface    Weight   State
      mams-4/0/0   10       Active
      mams-4/1/0   10       Active
      mams-5/0/0   10       Active
      mams-5/1/0   10       Backup
  Sync-state     :
      Interface    Status
      mams-4/0/0   Unknown
      mams-4/1/0   Unknown
      mams-5/0/0   Unknown
```

show interfaces load-balancing detail (Specific Interface)

```
user@host> show interfaces load-balancing ams0 detail
Load-balancing interfaces detail
Interface        : ams0
  State          : Up
  Last change    : 00:11:28
  Member count   : 4
  HA Model       : Many-to-One
  Members        :
      Interface    Weight   State
      mams-4/0/0   10       Active
      mams-4/1/0   10       Active
      mams-5/0/0   10       Active
      mams-5/1/0   10       Backup
  Sync-state     :
      Interface    Status
      mams-4/0/0   Unknown
      mams-4/1/0   Unknown
      mams-5/0/0   Unknown
```

Release Information

Command introduced in Junos OS Release 11.4.

interface-name option added in Junos OS Release 16.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Understanding Aggregated Multiservices Interfaces

Understanding Aggregated Multiservices Interfaces for Next Gen Services

Example: Configuring an Aggregated Multiservices Interface (AMS)

show interfaces (Logical Tunnel)

IN THIS SECTION

- Syntax | **21029**
- Description | **21030**
- Options | **21030**
- Required Privilege Level | **21030**
- Output Fields | **21030**
- Sample Output | **21036**
- Release Information | **21039**

Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified logical tunnel interface.

Options

interface-type	On M Series and T Series routers, the interface type is **lt-***fpc/pic/port*. On ACX7K series routers, the interface type is **lt-***fpc/pfe/core/channel*. or **lt-***feb/pfe/core/channel*
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 660 on page 21030 lists the output fields for the `show interfaces` (logical tunnel) command. Output fields are listed in the approximate order in which they appear.

Table 660: Logical Tunnel show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels

Table 660: Logical Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface. **Software-Pseudo** indicates a standard software interface with no associated hardware device.	All levels
Link-level type	Encapsulation used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Clocking	Reference clock source: **Internal or External** when configured. Otherwise, **Unspecified**.	All levels
Speed	Speed at which the interface is running.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels
Link type	Type of link.	All levels
Link flags	Information about the link. Possible values are described in the "Link Flags" section under *Common Output Fields Description*.	All levels

Table 660: Logical Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Physical info	Information about the physical interface.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Alternate link address	Backup link address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets, Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**

Table 660: Logical Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame terminates and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 660: Logical Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • **Errors**—Sum of the outgoing frame terminates and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets larger than the MTU threshold. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	All levels

Table 660: Logical Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. • **Input bytes**—Rate of bytes received on the interface. • **Output bytes**—Rate of bytes transmitted on the interface. • **Input packets**—Rate of packets received on the interface. • **Output packets**—Rate of packets transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, **mpls**.	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Route table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**

Table 660: Logical Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces extensive (Logical Tunnel)

```
user@host> show interfaces lt-1/0/0 extensive
Physical interface: lt-1/0/0, Enabled, Physical link is Up
  Interface index: 143, SNMP ifIndex: 70, Generation: 26
  Type: Logical-tunnel, Link-level type: Logical-tunnel, MTU: 0,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Unspecified
  Link flags     : None
  Physical info  : 13
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:7e, Hardware address: Unspecified
```

```
  Alternate link address: Unspecified
  Last flapped   : 2004-03-03 15:53:52 PST (22:08:46 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0

  Logical interface lt-1/0/0.0 (Index 66) (SNMP ifIndex 467) (Generation 3024)
    Flags: Point-To-Point SNMP-Traps 16384 DLCI 100 Encapsulation: FR-NLPID
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 4470, Generation: 7034, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: Unspecified,
        Generation: 2054
```

show interfaces extensive (Logical Tunnel) on ACX7K Series Routers

```
user@host> show interfaces lt-0/0/0:0
Physical interface: lt-0/0/0:0, Enabled, Physical link is Up
```

```
  Interface index: 1271, SNMP ifIndex: 8065
  Type: Logical Tunnel, Link-level type: Ethernet, MTU: Unlimited, Speed: 10Gbps, Auto-
negotiation: Disabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
```

show interfaces extensive (Logical Tunnel) on ACX7K Series Routers

```
user@host> show interfaces lt-0/0/0:0 terse
Interface               Admin Link Proto    Local                 Remote
lt-0/0/0:0              up    up
```

show interfaces extensive (Logical Tunnel) on ACX7K Series Routers

```
user@host> show interfaces lt-0/0/0:0 detail
Physical interface: lt-0/0/0:0, Enabled, Physical link is Up
  Interface index: 1271, SNMP ifIndex: 8065, Generation: 605590390119
  Type: Logical Tunnel, Link-level type: Ethernet, MTU: Unlimited, Clocking: Unspecified, Speed:
10Gbps, Auto-negotiation: Disabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Unspecified
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
```

show interfaces extensive (Logical Tunnel) on ACX7K Series Routers

```
user@host> show interfaces lt-0/0/0:0 extensive
Physical interface: lt-0/0/0:0, Enabled, Physical link is Up
  Interface index: 1271, SNMP ifIndex: 8065, Generation: 605590390119
  Type: Logical Tunnel, Link-level type: Ethernet, MTU: Unlimited, Clocking: Unspecified, Speed:
10Gbps, Auto-negotiation: Disabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Unspecified
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0, ECN Marked packets: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces (Loopback)

IN THIS SECTION

- Syntax | 21040
- Description | 21040
- Options | 21040
- Required Privilege Level | 21041
- Output Fields | 21041
- Sample Output | 21046
- Release Information | 21051

Syntax

```
show interfaces lo0
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the local loopback interface.

NOTE: Logical interface `lo0.16385` is the loopback interface for the internal routing instance. Created by the internal routing service process, this interface facilitates internal traffic. It prevents any filter created on loopback `lo0.0` from blocking internal traffic.

Options

lo0 Display standard status information about the local loopback interface.

brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information.
snmp-index ***snmp-index***	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 661 on page 21041 lists the output fields for the `show interfaces` (loopback) command. Output fields are listed in the approximate order in which they appear.

Table 661: Loopback show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical Interface`	Name of the physical interface.	All levels
`Enabled`	State of the interface. Possible values are described in the “Enabled Field” section under Common Output Fields Description.	All levels
`Interface index`	Physical interface index number, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Table 661: Loopback show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Type`	Type of interface.	All levels
`Link-level type`	Encapsulation type used on the physical interface.	All levels
`MTU`	Size of the largest packet to be transmitted.	All levels
`Clocking`	Reference clock source of the interface.	All levels
`Speed`	Network speed on the interface.	All levels
`Device flags`	Information about the physical device. Possible values are described in the “Device Flags” section under Common Output Fields Description.	All levels
`Interface flags`	Information about the interface. Possible values are described in the “Interface Flags” section under Common Output Fields Description.	All levels
`Link type`	Data transmission type.	`detail extensive`
`Link flags`	Information about the link. Possible values are described in the “Link Flags” section under Common Output Fields Description.	`detail extensive` none
`Physical info`	Information about the physical interface.	`detail extensive`
`Hold-times`	Current interface hold-time up and hold-time down. Value is in milliseconds.	`detail extensive`
`Current address`	Configured MAC address.	`detail extensive`
`Hardware address`	Media access control (MAC) address of the interface.	`detail extensive`

Table 661: Loopback show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Alternate link address`	Backup link address.	`detail extensive`
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second timezone`* (*`hour`*`:`*`minute`*`:`*`second`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive`
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`, `Output bytes`—Number of bytes received and transmitted on the interface. • `Input packets`, `Output packets`—Number of packets received and transmitted on the interface.	`detail extensive`
`Input errors`	• `Errors`—Input errors on the interface. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Frames received smaller than the runt threshold. • `Giants`—Frames received larger than the giant threshold. • `Policed Discards`—Frames that the incoming packet match code discarded because the frames were not recognized or were not of interest. Usually, this field reports protocols that Junos does not support. • `Resource errors`—Sum of transmit drops.	`extensive`

Table 661: Loopback show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	• `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • `Errors`—Sum of outgoing frame terminations and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC RED mechanism. • `MTU errors`—Number of packets larger than the MTU threshold. • `Resource errors`—Sum of transmit drops.	`extensive`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Logical interface index number, which reflects its initialization sequence.	`detail extensive`
`SNMP ifIndex`	Logical interface SNMP interface index number.	`detail extensive`
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Flags`	Information about the logical interface; values are described in the "Logical Interface Flags" section under Common Output Fields Description.	`brief detail extensive`
`Encapsulation`	Encapsulation on the logical interface.	`brief detail extensive`

Table 661: Loopback show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input packets`	Number of packets received on the logical interface.	None specified
`Output packets`	Number of packets transmitted on the logical interface.	None specified
`Traffic statistics`	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	`detail extensive`
`Local statistics`	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	`detail extensive`
`Protocol`	Protocol family configured on the logical interface (such as `iso` or `inet6`).	`detail extensive` none
`MTU`	MTU size on the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Route Table`	Route table in which this address exists; for example, `Route table:0` refers to inet.0.	`detail extensive`
`Flags`	Information about the protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	`detail extensive` none
`Addresses, Flags`	Information about the address flags. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	`detail extensive`

Table 661: Loopback show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Destination`	IP address of the remote side of the connection.	`detail extensive` none
`Local`	IP address of the logical interface.	`detail extensive` none
`Broadcast`	Broadcast address on the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Sample Output

show interfaces (Loopback)

```
user@host> show interfaces lo0
Physical interface: lo0, Enabled, Physical link is Up
  Interface index: 6, SNMP ifIndex: 6
  Type: Loopback, MTU: Unlimited
  Device flags   : Present Running Loopback
  Interface flags: SNMP-Traps
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface lo0.0 (Index 64) (SNMP ifIndex 16)
    Flags: SNMP-Traps Encapsulation: Unspecified
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: Unlimited
      Flags: None
      Addresses, Flags: Is-Default Is-Primary
        Local: 203.0.113.1
```

```
      Addresses
        Local: 127.0.0.1
    Protocol iso, MTU: Unlimited
      Flags: None
      Addresses, Flags: Is-Default Is-Primary
        Local: 49.0004.1000.0000.0001

  Logical interface lo0.16385 (Index 65) (SNMP ifIndex 76)
    Flags: SNMP-Traps Encapsulation: Unspecified
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: Unlimited
      Flags: None
```

show interfaces brief (Loopback)

```
user@host> show interfaces lo0 brief
Physical interface: lo0, Enabled, Physical link is Up
  Type: Loopback, Link-level type: Unspecified, MTU: Unlimited,
  Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running Loopback
  Interface flags: SNMP-Traps

  Logical interface lo0.0
    Flags: SNMP-Traps Encapsulation: Unspecified
    inet  203.0.113.1         --> 0/0
          127.0.0.1        --> 0/0
    iso   49.0004.1000.0000.0001

  Logical interface lo0.16385
    Flags: SNMP-Traps Encapsulation: Unspecified
    inet
```

show interfaces detail (Loopback)

```
user@host> show interfaces lo0 detail
Physical interface: lo0, Enabled, Physical link is Up
  Interface index: 6, SNMP ifIndex: 6, Generation: 4
  Type: Loopback, Link-level type: Unspecified, MTU: Unlimited,
  Clocking: Unspecified, Speed: Unspecified
```

```
  Device flags   : Present Running Loopback
  Interface flags: SNMP-Traps
  Link type      : Unspecified
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Logical interface lo0.0 (Index 64) (SNMP ifIndex 16) (Generation 3)
    Flags: SNMP-Traps Encapsulation: Unspecified
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0

    Protocol inet, MTU: Unlimited, Generation: 10, Route table: 0
      Flags: None
      Addresses, Flags: Is-Default Is-Primary
        Destination: Unspecified, Local: 203.0.113.1, Broadcast: Unspecified,
        Generation: 10
      Addresses, Flags: None
        Destination: Unspecified, Local: 127.0.0.1, Broadcast: Unspecified,
        Generation: 12
    Protocol iso, MTU: Unlimited, Generation: 11, Route table: 0
      Flags: None
      Addresses, Flags: Is-Default Is-Primary
        Destination: Unspecified, Local: 49.0004.1000.0000.0001,
        Broadcast: Unspecified, Generation: 14

  Logical interface lo0.16385 (Index 65) (SNMP ifIndex 76) (Generation 4)
```

```
    Flags: SNMP-Traps Encapsulation: Unspecified
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Protocol inet, MTU: Unlimited, Generation: 12, Route table: 1
      Flags: None
```

show interfaces extensive (Loopback)

```
user@host> show interfaces lo0 extensive
Physical interface: lo0, Enabled, Physical link is Up
  Interface index: 6, SNMP ifIndex: 6, Generation: 4
  Type: Loopback, Link-level type: Unspecified, MTU: Unlimited,
  Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running Loopback
  Interface flags: SNMP-Traps
  Link type      : Unspecified
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
```

```
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

Logical interface lo0.0 (Index 64) (SNMP ifIndex 16) (Generation 3)
    Flags: SNMP-Traps Encapsulation: Unspecified
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Protocol inet, MTU: Unlimited, Generation: 10, Route table: 0
      Flags: None
      Addresses, Flags: Is-Default Is-Primary
        Destination: Unspecified, Local: 203.0.113.1, Broadcast: Unspecified,
        Generation: 10
      Addresses, Flags: None
        Destination: Unspecified, Local: 127.0.0.1, Broadcast: Unspecified,
        Generation: 12
    Protocol iso, MTU: Unlimited, Generation: 11, Route table: 0
      Flags: None
      Addresses, Flags: Is-Default Is-Primary
        Destination: Unspecified, Local: 49.0004.1000.0000.0001,
        Broadcast: Unspecified, Generation: 14

  Logical interface lo0.16385 (Index 65) (SNMP ifIndex 76) (Generation 4)
    Flags: SNMP-Traps Encapsulation: Unspecified
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Protocol inet, MTU: Unlimited, Generation: 12, Route table: 1
```

```
    Flags: None
```

NOTE: On Junos OS Evolved, the `show interfaces extensive` output displays zero for all loopback interface (lo0) statistics.

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces lsi (Label-Switched Interface)

IN THIS SECTION

- Syntax | 21051
- Description | 21052
- Options | 21052
- Required Privilege Level | 21052
- Output Fields | 21052
- Sample Output | 21054
- Release Information | 21055

Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance instance-name>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified label-switched interface (LSI).

Options

interface-type	On most routers, the interface type is `lt-`*`fpc`*/*`pic`*/*`port`*.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
routing-instance *instance-name*	(Optional) Display information for the specified routing instance.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 662 on page 21052 lists the output fields for the `show interfaces` (logical tunnel) command. Output fields are listed in the approximate order in which they appear.

Table 662: Logical Tunnel show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels

Table 662: Logical Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	`detail` `extensive` none
SNMP ifIndex	SNMP interface index number.	`detail` `extensive` none
Generation	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under Common Output Fields Description.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. • `Input bytes`—Rate of bytes received on the interface. • `Output bytes`—Rate of bytes transmitted on the interface. • `Input packets`—Rate of packets received on the interface. • `Output packets`—Rate of packets transmitted on the interface.	`detail` `extensive`
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	`detail` `extensive`

Table 662: Logical Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	`detail extensive`
Protocol	Protocol family configured on the logical interface, such as `iso`, `inet6`, `mpls`.	`detail extensive` none
MTU	MTU size on the logical interface.	`detail extensive` none
Generation	Unique number for use by Juniper Networks technical support only.	`detail extensive`
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	`detail extensive` none

Sample Output

show interfaces lsi extensive

```
user@host> show interfaces lsi extensive
Physical interface: lsi

  Logical interface lsi.84934656 (Index 363) (SNMP ifIndex 586) (Generation 194)
    Flags: Up Point-To-Point SNMP-Traps 0x4000000 Encapsulation: LSI-NULL
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
```

```
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol vpls, MTU: Unlimited, Generation: 279, Route table: 10

  Logical interface lsi.84934657 (Index 366) (SNMP ifIndex 589) (Generation 197)
    Flags: Up Point-To-Point SNMP-Traps 0x4000000 Encapsulation: LSI-NULL
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol vpls, MTU: Unlimited, Generation: 282, Route table: 10
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet)

IN THIS SECTION

- Syntax (M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface) | 21056
- Syntax (M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface) | 21056
- Description | 21057
- Options | 21057
- Required Privilege Level | 21057
- Output Fields | 21057
- Sample Output | 21063
- Release Information | 21080

Syntax (M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface)

```
show interfaces em0 | fxp0 | mgmtre0
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface)

```
show interfaces bcm0  | em0 | em1| fxp1 | fxp2 | ixgbe0 | ixgbe1
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M Series, T Series, TX Matrix Plus, and PTX Series devices only) Display status information about the management Ethernet and internal Ethernet interfaces.

Options

em0 | fxp0 | mgmtre0
(M Series, MX Series, T Series, and PTX Series) Display standard information about the management Ethernet interface. For supported Ethernet interface by chassis and Routing Engine, see Supported Routing Engines by Router.

bcm0 | em0 | em1 | fxp1 | fxp2 | ixgbe0 | ixgbe1
(M Series, MX Series, T Series, and PTX Series) Display standard information about the internal Ethernet interfaces. See Supported Routing Engines by Router for the internal Ethernet interface names for each Routing Engine by hardware platform.

> **NOTE**: On Junos OS Evolved, the ixgbe0 and ixgbe1 internal interfaces are deprecated.

brief | detail | extensive | terse
(Optional) Display the specified level of output.

descriptions
(Optional) Display interface description strings.

media
(Optional) Display media-specific information.

snmp-index *snmp-index*
(Optional) Display information for the specified SNMP index of the interface.

statistics
(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 663 on page 21058 lists the output fields for the `show interfaces` (management) command on the M Series routers, T Series routers, TX Matrix Plus routers, and PTX Series. Output fields are listed in the approximate order in which they appear.

Table 663: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface.	All levels
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum transmission unit (MTU)—Size of the largest packet to be transmitted.	All levels
Clocking	Reference clock source of the interface.	All levels
Speed	Network speed on the interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels

Table 663: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface *(Continued)*

Field Name	Field Description	Level of Output
Link type	Data transmission type.	**detail extensive** none
Link flags	Information about the link. Possible values are described in the "Link Flags" section under *Common Output Fields Description*.	**detail extensive**
Physical info	Information about the physical interface.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Media access control (MAC) address of the interface.	**detail extensive** none
Alternate link address	Backup link address.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour.minute.second timezone*** **(*hour.minute.second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Input packets	Number of packets received on the physical interface.	None specified
Output packets	Number of packets transmitted on the physical interface.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 663: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical and physical interface. • **Input bytes**, **Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets**, **Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**
Input errors	• **Errors**—Input errors on the interface. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Frames received smaller than the runt threshold. • **Giants**—Frames received larger than the giant threshold. • **Policed Discards**—Frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos does not support. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 663: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface *(Continued)*

Field Name	Field Description	Level of Output
Output errors	• **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • **Errors**—Sum of outgoing frame terminates and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC RED mechanism. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface; values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Encapsulation	Encapsulation on the logical interface.	**detail extensive** none
inet	IP address of the logical interface.	**brief**

Table 663: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface *(Continued)*

Field Name	Field Description	Level of Output
Protocol	Protocol family configured on the logical interface (such as **iso** or **inet6**).	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Route table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive** none
Addresses, Flags	Information about address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces brief (Management Ethernet)

```
user@host> show interfaces fxp0 brief
Physical interface: fxp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface fxp0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  192.168.70.143/21
```

show interfaces (Management Ethernet)

```
user@host> show interfaces fxp0
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Last flapped   : Never
    Input packets : 80804
    Output packets: 1105

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255
```

show interfaces (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Last flapped   : Never
    Input packets : 1424
    Output packets: 5282

  Logical interface em0.0 (Index 3) (SNMP ifIndex 18)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 1424
    Output packets: 5282
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast: 192.168.178.127
```

show interfaces (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Last flapped   : Never
    Input packets : 212581
    Output packets: 71

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 212551
```

```
    Output packets: 71
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 192.168.3/24, Local: 192.168.3.30,
        Broadcast: 192.168.3.255
```

show interfaces detail (Management Ethernet)

```
user@host> show interfaces fxp0 detail
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1, Generation: 0
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              6484031
   Output bytes  :               167503
   Input  packets:                81008
   Output packets:                 1110

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 6, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255, Generation: 1
```

show interfaces detail (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 detail
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17, Generation: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               124351
   Output bytes  :              1353212
   Input  packets:                 1804
   Output packets:                 5344
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0

  Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               117135
     Output bytes  :              1331647
     Input  packets:                 1804
     Output packets:                 5344
    Local statistics:
     Input  bytes  :               117135
     Output bytes  :              1331647
     Input  packets:                 1804
     Output packets:                 5344
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
```

```
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast: 192.168.178.127,
Generation: 1
```

show interfaces detail (Management Ethernet [PTX Packet Transport Routers])

```
user@host> show interfaces detail em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0, Generation: 3
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
      Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Physical info  : Unspecified
      Hold-times     : Up 0 ms, Down 0 ms
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Alternate link address: Unspecified
      Last flapped   : Never
      Statistics last cleared: Never
      Traffic statistics:
       Input  bytes  :             15255909
       Output bytes  :                 4608
       Input  packets:               214753
       Output packets:                   72
       IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             14394630
     Output bytes  :                 3024
     Input  packets:               214723
     Output packets:                   72
    Local statistics:
     Input  bytes  :             14394630
     Output bytes  :                 3024
     Input  packets:               214723
```

```
     Output packets:                    72
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 192.168.3/24, Local: 192.168.3.30,
        Broadcast: 192.168.3.255, Generation: 1
```

show interfaces extensive (Management Ethernet)

```
user@host> show interfaces fxp0 extensive
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1, Generation: 0
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              6678904
   Output bytes  :               169657
   Input  packets:                83946
   Output packets:                 1127
  Input errors:
    Errors: 12, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 6, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
```

```
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255, Generation: 1
```

show interfaces extensive (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 extensive

Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17, Generation: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               127120
   Output bytes  :              1357414
   Input  packets:                 1843
   Output packets:                 5372
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0

  Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               119748
     Output bytes  :              1335719
     Input  packets:                 1843
```

```
     Output packets:                 5372
    Local statistics:
     Input  bytes  :               119748
     Output bytes  :              1335719
     Input  packets:                 1843
     Output packets:                 5372
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast: 192.168.178.127,
Generation: 1
```

show interfaces extensive (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces extensive em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0, Generation: 3
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
      Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Physical info  : Unspecified
      Hold-times     : Up 0 ms, Down 0 ms
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Alternate link address: Unspecified
      Last flapped   : Never
      Statistics last cleared: Never
      Traffic statistics:
       Input  bytes  :             15236459
       Output bytes  :                 4608
       Input  packets:               214482
       Output packets:                   72
       IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
      Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
```

```
      Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             14376264
     Output bytes  :                 3024
     Input  packets:               214452
     Output packets:                   72
    Local statistics:
     Input  bytes  :             14376264
     Output bytes  :                 3024
     Input  packets:               214452
     Output packets:                   72
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 192.168.3/24, Local: 192.168.3.30,
        Broadcast: 192.168.3.255, Generation: 1
```

show interfaces mgmtre0 (Management Ethernet [PTX5000 Router])

```
user@host> show interfaces mgmtre0 extensive
Physical interface: mgmtre0, Enabled, Physical link is Up
  Interface index: 1001, SNMP ifIndex: 501
  Link-level type: Ethernet, MTU: 1500
  Device flags   : Present
  Interface flags: None
  Link flags     : None
  Current address: ec:9e:cd:06:30:da, Hardware address: ec:9e:cd:06:30:da
  Last flapped   : Never

  Logical interface mgmtre0.0 (Index 1001) (SNMP ifIndex 503)
    Flags: Encapsulation: ENET2
    Protocol inet, MTU: 1486
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.92.248/23, Local: 10.92.248.22,
        Broadcast: 10.92.249.255
```

```
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

show interfaces brief (Management Ethernet)

```
user@host> show interfaces fxp1 brief
Physical interface: fxp1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface fxp1.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  10.0.0.4/8
    inet6 fe80::200:ff:fe00:4/64
          fec0::10:0:0:4/64
    tnp   4
```

show interfaces brief (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface em0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  192.168.178.11/25
```

show interfaces brief (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
      Speed: 1000mbps
      Device flags   : Present Running
```

```
      Interface flags: SNMP-Traps

      Logical interface em0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  192.168.3.30/24

    root@aboslutely> show interfaces em0 terse
    Interface               Admin Link Proto    Local                 Remote
    em0                     up    up
    em0.0                   up    up   inet     192.168.3.30/24
```

show interfaces (Internal Ethernet)

```
user@host> show interfaces fxp1
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Last flapped   : Never
    Input packets : 30655
    Output packets: 33323

  Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255
    Protocol inet6, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe00:4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4
    Protocol tnp, MTU: 1500
      Flags: Primary, Is-Primary
```

```
      Addresses
        Local: 4
```

show interfaces (Internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Last flapped   : Never
    Input packets : 2301738
    Output packets: 3951155

  Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 2301595
    Output packets: 3951155
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255
      Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
        Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 192.168.0.4
    Protocol inet6, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe22:4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::a:22:0:4
    Protocol tnp, MTU: 1500
      Flags: Primary, Is-Primary
      Addresses
        Local: 0x22000004
```

show interfaces detail (Internal Ethernet)

```
user@host> show interfaces fxp1 detail
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              2339969
   Output bytes  :             15880707
   Input  packets:                30758
   Output packets:                33443

  Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 7, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
        Generation: 3
    Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
        Broadcast: Unspecified, Generation: 5
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
        Generation: 7
    Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
```

```
        Destination: Unspecified, Local: 4, Broadcast: Unspecified,
        Generation: 8
```

show interfaces detail (Internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0 detail
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :            238172825
   Output bytes  :           1338948955
   Input  packets:              2360984
   Output packets:              4061512
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0

  Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :            228720309
     Output bytes  :           1261387447
     Input  packets:              2360841
     Output packets:              4061512
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
```

```
    Local statistics:
     Input  bytes  :            228720309
     Output bytes  :           1261387447
     Input  packets:              2360841
     Output packets:              4061512
    Protocol inet, MTU: 1500, Generation: 2, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255, Generation: 2
      Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
        Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 191.255.255.255, Generation: 3
    Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe22:4
    Generation: 4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::a:22:0:4
    Protocol tnp, MTU: 1500, Generation: 5
    Generation: 4, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
        Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified, Generation: 6
```

show interfaces extensive (internal Ethernet)

```
user@host> show interfaces fxp1 extensive
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
```

```
   Input  bytes  :              2349897
   Output bytes  :             15888605
   Input  packets:                30896
   Output packets:                33607
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 7, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
        Generation: 3
    Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
        Broadcast: Unspecified, Generation: 5
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
        Generation: 7
    Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
        Destination: Unspecified, Local: 4, Broadcast: Unspecified,
        Generation: 8
```

show interfaces extensive (internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0 extensive
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
```

```
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :            242730780
   Output bytes  :           1348312269
   Input  packets:              2398737
   Output packets:              4133510
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0

  Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :            233127252
     Output bytes  :           1269350897
     Input  packets:              2398594
     Output packets:              4133510
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :            233127252
     Output bytes  :           1269350897
     Input  packets:              2398594
     Output packets:              4133510
    Protocol inet, MTU: 1500, Generation: 2, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255, Generation: 2
```

```
    Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
      Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 191.255.255.255, Generation: 3
  Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::200:ff:fe22:4
  Generation: 4
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: fec0::/64, Local: fec0::a:22:0:4
  Protocol tnp, MTU: 1500, Generation: 5
  Generation: 4, Route table: 1
    Flags: Primary, Is-Primary
    Addresses, Flags: None
      Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified, Generation: 6
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

show interfaces mac-database

IN THIS SECTION

- Syntax | 21081
- Description | 21081
- Options | 21081
- Additional Information | 21081
- Required Privilege Level | 21082
- Output Fields | 21082
- Sample Output | 21085
- Release Information | 21090

Syntax

```
show interfaces mac-database (ge-fpc/pic/port | ge-fpc/pic/port.n | aex| xe-fpc/pic/port | xe-
fpc/pic/port.n | et-fpc/pic/port | et-fpc/pic/port.n) <mac-address mac-address>
```

Description

(M Series, T Series, MX Series routers, and PTX Series Packet Transport Routers only) Display media access control (MAC) address information for the specified interface.

Options

ge-*fpc/pic/port*	Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.
ge-*fpc/pic/port.n*	Display MAC addresses that have been learned on a particular logical interface.
ae*x*	Display MAC addresses that have been learned on a particular aggregated Ethernet interface.
xe-*fpc/pic/port*	Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.
xe-*fpc/pic/port.n*	Display MAC addresses that have been learned on a particular logical interface.
et-*fpc/pic/port*	Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.
et-*fpc/pic/port.n*	Display MAC addresses that have been learned on a particular logical interface.
mac-address *mac-address*	(Optional) Display detailed MAC address statistics, including policer information for `ge`, `xe`, and `et` interfaces.

Additional Information

On IQ2 PIC interfaces, the default value for maximum retention of entries in the MAC address table has changed, for cases in which the table is not full. The new holding time is 12 hours. The previous retention time of 3 minutes is still in effect when the table is full.

Required Privilege Level

view

Output Fields

Table 664 on page 21082 lists the output fields for the `show interfaces mac-database` command. Output fields are listed in the approximate order in which they appear.

Table 664: show interfaces mac-database Output Fields

Field Name	Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.
Interface index	Physical interface index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the physical interface.
Description	Description and name of the interface.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback: **local** or **remote**.
Source filtering	Whether source filtering is configured.
Flow control	Whether flow control is enabled or disabled.

Table 664: show interfaces mac-database Output Fields *(Continued)*

Field Name	Field Description
`Minimum links needed`	(Aggregated Ethernet interfaces only) Number of child links that must be operational for the aggregated interface to be operational.
`Minimum bandwidth needed`	(Aggregated Ethernet interfaces only) Minimum amount of bandwidth of child links that must be operational for the aggregated interface to be operational.
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.
`Currrent address`	(Aggregated Ethernet interfaces only) Configured MAC address.
`Hardware address`	(Aggregated Ethernet interfaces only) Hardware MAC address.
`Last flapped`	(Aggregated Ethernet interfaces only) Date, time, and how long ago the interface went from down to up or from up to down. The format is `Last flapped:` *`year-month-day hours:minutes:seconds timezone`* (`w`*weeks*`d`*days hours:minutes* `ago`). For example, `Last flapped: 2013-12-18 04:33:22 PST (1w5d 22:23 ago)`.
`Input Rate`	(Aggregated Ethernet interfaces only) Input rate in bits per second (bps) and packets per second (pps).
`Output Rate`	(Aggregated Ethernet interfaces only) Output rate in bps and pps.
Interface flags	Information about the interface. Possible values are described in the “Links Flags” section under *Common Output Fields Description*.
Link flags	Information about the link. Possible v4alues are described in the “Device Flags” section under *Common Output Fields Description*.
Logical Interface	
Logical interface	Name of the logical interface.

Table 664: show interfaces mac-database Output Fields *(Continued)*

Field Name	Field Description
Index	Logical interface index number, which reflects its initialization sequence.
SNMP ifIndex	Logical interface SNMP interface index number.
Flags	Information about the logical interface (possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.
Encapsulation	Encapsulation on the logical interface.
MAC address, Input frames, Input bytes, Output frames, Output bytes	MAC address and corresponding number of input frames, input bytes, output frames, and output bytes.
Number of MAC addresses	Number of MAC addresses configured.

Policer Statistics	(Displayed for **mac-address** option for `ge`, `xe`, and `et` interface types only) Display information about policers applied to a logical interface-MAC pair.
	• **Policer type**—Type of policer that is out of spec with respect to the configuration. It can be one or more of the following:
	  • **Input premium**—Number of high-priority rating out-of-spec frames or bytes received.
	  • **Output premium**—Number of high-priority rating out-of-spec frames or bytes sent.
	  • **Input aggregate**—Total number of out-of-spec frames or bytes received.
	  • **Output aggregate**—Total number of out-of-spec frames or bytes sent.
	• **Discarded Frames**—Number of discarded frames.
	• **Discarded Bytes**—Number of discarded bytes.

Sample Output

show interfaces mac-database (All MAC Addresses on a Port)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback: None, Source
filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:03      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:04      30424716     1399536936       37448523     1722632058
  00:00:c8:01:01:05      30424789     1399540294       37448598     1722635508
  00:00:c8:01:01:06      30424788     1399540248       37448597     1722635462
  00:00:c8:01:01:07      30424783     1399540018       37448597     1722635462
  00:00:c8:01:01:08      30424783     1399540018       37448596     1722635416
  00:00:c8:01:01:09       8836796      406492616        8836795      406492570
  00:00:c8:01:01:0a      30424712     1399536752       37448521     1722631966
  00:00:c8:01:01:0b      30424715     1399536890       37448523     1722632058
  Number of MAC addresses : 21
```

show interfaces mac-database (All MAC Addresses on an Aggregated Ethernet Interface)

```
user@host> show interfaces mac-database ae4
Physical interface: ae4, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 588
  Description: Member links xe-0/2/0
  Link-level type: Ethernet, MTU: 9188, Speed: Unspecified, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth
needed: 0
  Device flags   : Present Running
  Interface flags: Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:22:83:76:ff:c4, Hardware address: 00:22:83:76:ff:c4
  Last flapped   : 2013-12-18 04:33:22 PST (1w5d 22:23 ago)
  Input rate     : 62756384 bps (85266 pps)
  Output rate    : 62759472 bps (85272 pps)


  Logical interface ae4.0 (Index 334) (SNMP ifIndex 647)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
    MAC address        Input frames    Input bytes  Output frames   Output bytes
    00:00:00:aa:00:02      23888711     2627758118            300          22200
    00:00:00:aa:00:03             0              0              0              0
    00:00:00:aa:00:04             0              0              0              0
    Number of MAC addresses : 3
```

show interfaces mac-database (All MAC Addresses on a Service)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
```

```
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      31016568     1426762128       38040381     1749857526
  00:00:c8:01:01:03      31016568     1426762128       38040382     1749857572
  00:00:c8:01:01:04      31016499     1426758954       38040306     1749854076
  00:00:c8:01:01:05      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:06      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:07      31016567     1426762082       38040380     1749857480
  00:00:c8:01:01:08      31016567     1426762082       38040379     1749857434
  00:00:c8:01:01:09       9428580      433714680        9428580      433714680
  00:00:c8:01:01:0a      31016496     1426758816       38040304     1749853984
  00:00:c8:01:01:0b      31016498     1426758908       38040307     1749854122
```

show interfaces mac-database mac-address

```
user@host> show interfaces mac-database xe-0/3/3 mac-address       00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback: None, Source
filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes     : 202324652
    Output bytes    : 202324560
    Input frames    : 4398362
    Output frames   : 4398360
  Policer statistics:
  Policer type      Discarded frames    Discarded bytes
  Output aggregate           3992386          183649756
```

show interfaces mac-database (Physical interface)

```
user@host> show interfaces mac-database et-1/0/1
Physical interface: et-1/0/1, Enabled, Physical link is Up
  Interface index: 1022, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled, Auto-negotiation: Disabled, Media type: Fiber
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface et-1/0/1.0 (Index 1005) (SNMP ifIndex 545)
    Flags: Up SNMP-Traps Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:03      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:04      30424716     1399536936       37448523     1722632058
  00:00:c8:01:01:05      30424789     1399540294       37448598     1722635508
  00:00:c8:01:01:06      30424788     1399540248       37448597     1722635462
  00:00:c8:01:01:07      30424783     1399540018       37448597     1722635462
  00:00:c8:01:01:08      30424783     1399540018       37448596     1722635416
  00:00:c8:01:01:09       8836796      406492616        8836795      406492570
  00:00:c8:01:01:0a      30424712     1399536752       37448521     1722631966
  00:00:c8:01:01:0b      30424715     1399536890       37448523     1722632058
  Number of MAC addresses : 21
```

show interfaces mac-database (Logical interface)

```
user@host> show interfaces mac-database et-1/0/1.0
  Logical interface et-1/0/1.0 (Index 1005) (SNMP ifIndex 545)
    Flags: SNMP-Traps Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:03      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:04      30424716     1399536936       37448523     1722632058
  00:00:c8:01:01:05      30424789     1399540294       37448598     1722635508
  00:00:c8:01:01:06      30424788     1399540248       37448597     1722635462
  00:00:c8:01:01:07      30424783     1399540018       37448597     1722635462
  00:00:c8:01:01:08      30424783     1399540018       37448596     1722635416
  00:00:c8:01:01:09       8836796      406492616        8836795      406492570
  00:00:c8:01:01:0a      30424712     1399536752       37448521     1722631966
  00:00:c8:01:01:0b      30424715     1399536890       37448523     1722632058
```

show interfaces mac-database (Aggregated Ethernet)

```
user@host> show interfaces mac-database ae0
Physical interface: ae0, up, Physical link is Up
  Interface index: 6000, SNMP ifIndex: 541
  Link-level type: Ethernet, MTU: 1514, Speed: 800Gbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Minimum links
needed: 1,
  Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Current address: da:5e:26:3b:03:57, Hardware address: da:5e:26:3b:03:57
```

```
  Last flapped   : 2021-12-16 12:08:26 IST (00:04:00 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ae0.0 (Index 9700) (SNMP ifIndex 546)
    Flags: Up SNMP-Traps Encapsulation: Ethernet-Bridge
    MAC address        Input frames    Input bytes  Output frames   Output bytes
    00:00:00:aa:00:02      23888711     2627758118            300          22200
    00:00:00:aa:00:03             0              0              0              0
    00:00:00:aa:00:04             0              0              0              0
    Number of MAC addresses : 3
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

Support for statement with the `ae`*x* option introduced in Junos OS Release 15.1.

show interfaces mc-ae (MX Series)

IN THIS SECTION

- Syntax | **21091**
- Description | **21091**
- Options | **21091**
- Required Privilege Level | **21091**
- Output Fields | **21091**
- Sample Output | **21094**
- Release Information | **21097**

Syntax

```
show interfaces mc-ae
extensive
revertive-info
<id identifier unit number>
```

Description

On MX Series routers with multichassis aggregated Ethernet (**ae*X***) interfaces, displays information about the **ae*X*** interfaces.

Options

extensive	(Optional) Display extensive information for multichassis aggregated Ethernet interface.
revertive-info	(Optional) Display revertive mode information for multichassis aggregated Ethernet interface.
identifier	(Optional) Identifier of the multichassis aggregated Ethernet interface.
number	(Optional) Specify the logical interface by unit number.

Required Privilege Level

view

Output Fields

Table 665 on page 21091 lists the output fields for the `show interfaces mc-ae` command. Output fields are listed in the approximate order in which they appear.

Table 665: show interfaces mc-ae Output Fields

Output Field Name	Field Description
`Member Link`	Identifiers of the configured multichassis link aggregate interfaces configured interfaces.

Table 665: show interfaces mc-ae Output Fields *(Continued)*

Output Field Name	Field Description
`Current State Machine's State`	Current state of the MCLAG state machine. The MCLAG state machine is responsible for synchronization with the peer MCLAG node.
`Local Status`	Status of the local link: **active** or **standby**.
`Peer Status`	Status of the peer link: **active** or **standby**.
`Local State`	Up or down state of the local device.
`Peer State`	Status of the local and peer links in an **active/active** bridge or VRRP over integrated routing and bridging (IRB) configuration on MX Series routers, including: **Logical Interface** — Aggregated Ethernet (AE) aggregate number and unit number. **Topology Type** — The bridge or VRRP topology type configured on the AE. **Local State** — Up or down state of the local device. **Peer State** — Up or down state of the peer device. **Peer Ip/ICL-PL/State** — Address, interface and state of the peer device.
Core Facing Interface	Label: **pseudowire interface** or **Ethernet interface**.
ICL-PL	Label: **pseudowire interface** or **Ethernet interface**.
`switchover mode`	The configured switchover mode for the multichassis aggregated Ethernet interface: `revertive` or `non-revertive`.

Table 665: show interfaces mc-ae Output Fields *(Continued)*

Output Field Name	Field Description
`switchover status`	Status of the switchover if the `revert-time` statement is configured at the `[edit interfaces ae`*x* `mc-ae]` hierarchy level.
`revert time`	Revert time configured for the multichassis aggregated Ethernet interface.
`switchover time remaining`	Seconds left to trigger the switchover if the switchover is in progress.
`Configuration Error Status`	Reason for the configuration error.
MCAE Configuration	
`Redundancy Groups`	Identification number of the redundancy group. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions. Possible values: 1 through 4,294,967,294.
`MCAE ID`	Identification number of the MC-LAG device. The two MC-LAG network devices that manage a given MC-LAG must have the same identification number. Possible values:1 through 65,535.
`MCAE Mode`	Specifies whether the MC-LAG is in active-active or active-standby mode. Possible values: active-active or active-standby.
`Status Control`	Specifies whether the chassis becomes active or remains in standby mode when an interchassis link failure occurs. Possible values: active or standby.
`Chassis ID`	Chassis ID for Link Aggregation Control Protocol (LACP) to calculate the port number of MC-LAG physical member links. Possible values: 0 or 1.
LACP Configuration	

Table 665: show interfaces mc-ae Output Fields *(Continued)*

Output Field Name	Field Description
System ID	System id of the local system.
Admin Key	LACP administrative key of the node. Value should be the same on both MCLAG peers.
LACP Information	
Local Partner System ID	LACP partner system ID as seen by the local node.
Peer Partner System ID	LACP partner system ID as seen by MC-AE peer node.

Sample Output

show interfaces mc-ae

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links   : ae0
  Local Status   : active
  Peer Status    : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL         : Label Ethernet Interface
```

show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series Routers)

```
user@host# show interfaces mc-ae ge-0/0/0.0
 Member Link                  : ae0
 Current State Machine's State: active
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
```

```
     Logical Interface        : ae0.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/ICL-PL/State     : 192.168.100.10 ge-0/0/0.0 up
```

show interfaces mc-ae revertive-info

```
user@host> show interfaces mc-ae revertive-info id 2
Member Link           : ae1
Current State Machine's State: mcae active state
Local Status                 : active
Local State                  : up
Peer Status                  : standby
Peer State                   : up
Switchover Mode              : Non Revertive
Switchover Status            : N/A
Revert Time                  : 1 Minutes
Switchover Remaining Time    : N/A
     Logical Interface        : ae1.1024
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : N/A
```

show interfaces mc-ae extensive

```
user@host> show interfaces mc-ae extensive
 Member Link                  : ae2
Current State Machine's State: mcae active state
Local Status                 : active
Local State                  : up
Peer Status                  : active
Peer State                   : up
     Logical Interface        : ae2.1
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 192.168.143.17 ae0.1 up
```

```
MCAE Configuration
     Redundancy Group         : 1
     MCAE ID                  : 2
     MCAE Mode                : active_active
     Status Control           : active
     Chassis ID               : 0
LACP Configuration
     System ID                : 00:00:00:00:00:02
     Admin Key                : 10
```

show interfaces mc-ae extensive (MX Series Router after a configuration exchange error)

```
user@host> show interfaces mc-ae extensive
 Member Link                  : ae2
Current State Machine's State: mcae config exchange error
Configuration Error Status   : same chassis-id
Local Status                 : active
Local State                  : up
Peer Status                  : Unknown
Peer State                   : Unknown
     Logical Interface        : ae2.1
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 192.168.143.17 ae0.1 up

MCAE Configuration
     Redundancy Group         : 1
     MCAE ID                  : 2
     MCAE Mode                : active_active
     Status Control           : active
     Chassis ID               : 1
LACP Configuration
    System ID                : 00:00:00:00:00:02
     Admin Key                : 10
```

show interfaces mc-ae extensive

```
user@host> show interfaces mc-ae extensive
Member Link                  : ae0
```

```
Current State Machine's State: mcae active state
Local Status                 : active
Local State                  : up
Peer Status                  : active
Peer State                   : up
    Logical Interface        : ae0.1
    Topology Type            : bridge
    Local State              : up
    Peer State               : up
    Peer Ip/MCP/State        : 192.168.143.17 ge-0/0/2.1 up

MCAE Configuration
    Redundancy Group         : 1
    MCAE ID                  : 1
    MCAE Mode                : active_active
    Status Control           : active
    Chassis ID               : 0
LACP Configuration
    System ID                : 00:00:00:00:00:02
    Admin Key                : 10
LACP Information
    Local Partner System ID  : 00:00:00:00:00:01
    Peer Partner System ID   : 00:00:00:00:00:01
```

Release Information

Command introduced in Junos OS Release 9.6.

`revertive-info` statement introduced in Junos OS Release 13.3

`extensive` statement introduced in Junos OS Release 16.1R1

RELATED DOCUMENTATION

- Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up

show interfaces mc-ae

IN THIS SECTION

- Syntax | 21098
- Description | 21098
- Options | 21098
- Required Privilege Level | 21098
- Output Fields | 21099
- Sample Output | 21100
- Release Information | 21101

Syntax

```
show interfaces mc-ae id identifier unit number
```

Description

On peers with multichassis aggregated Ethernet (**mc-ae*X***) interfaces, use this command to display information about the multichassis aggregated Ethernet interfaces.

NOTE: In Junos OS Release 17.4R1, this command is not supported on EX4300, EX9200, PTX10000, QFX10002, and QFX10008 devices.

Options

id *identifier* (Optional) Specify the name of the multichassis aggregated Ethernet interface.

unit *number* (Optional) Specify the logical interface by unit number.

Required Privilege Level

view

Output Fields

Table 666 on page 21099 lists the output fields for the `show interfaces mc-ae` command. Output fields are listed in the approximate order in which they appear.

Table 666: show interfaces mc-ae Output Fields

Output Field Name	Field Description
Current State Machine's State	Specifies the state of the MC-LAG initialization state machine.
`Configuration Consistency Check`	Specifies the status of the MC-LAG configuration consistency check feature. The status is either `Passed` or `Failed`. If the status is `Failed`, the system will display the name of the parameter that failed consistency check. If there are multiple inconsistencies, only the first inconsistency is shown. If the enforcement level for the MC-LAG parameter was mandatory, and you did not configure that parameter correctly, the command will show that the MC-LAG interface is down.
Member Link	Specifies the identifiers of the configured multichassis link aggregated interface members.
Local Status	Specifies the status of the local link: **active** or **standby**.
Peer Status	Specifies the status of the peer link: **active** or **standby**.
Peer State	Specifies the status of the local and peer links in an **active/active** MC-LAG configuration.
Logical Interface	Specifies the identifier and unit of the AE interface.
Topology Type	Specifies the bridge configured on the AE.
Local State	Specifies if the local device is up or down.
Peer State	Specifies if the peer device is up or down.

Table 666: show interfaces mc-ae Output Fields *(Continued)*

Output Field Name	Field Description
Peer Ip/MCP/State	Specifies the multichassis protection (MCP) link or the interchassis link-protection link (ICL-PL) for all of the multichassis aggregated Ethernet interfaces that are part of the peer.

Sample Output

show interfaces mc-ae (EX Series)

```
user@switch> show interfaces mc-ae ae1 512
 Member Link                  : ae1
  Current State Machine's State: mcae active state
  Configuration Consistency Check    : Failed (redundancy group id mismatch)
  Local Status                 : active
  Local State                  : up
  Peer Status                  : standby
  Peer State                   : up
      Logical Interface        : ae1.0
      Topology Type            : bridge
      Local State              : up
      Peer State               : up
      Peer Ip/MCP/State        : 10.1.1.1 ae0.0 up
```

show interfaces mc-ae (MX Series)

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links   : ae0
  Local Status   : active
  Peer Status    : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL         : Label Ethernet Interface
```

show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series)

```
user@host# show interfaces mc-ae ge-0/0/0.0
 Member Link                  : ae0
 Current State Machine's State: active
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae0.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/ICL-PL/State     : 192.168.100.10 ge-0/0/0.0 up
```

Release Information

Command introduced in Junos OS Release 9.6.

`Configuration Consistency Check` output field added in Junos OS Release 15.1X53-D60 for the QFX Series.

show interfaces me0

IN THIS SECTION

- Syntax | 21102
- Description | 21102
- Options | 21102
- Required Privilege Level | 21102
- Output Fields | 21102
- Sample Output | 21109
- Release Information | 21113

Syntax

```
show interfaces me0
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance>
<statistics>
```

Description

Display status information about the management Ethernet interface.

Options

none	Display standard information about the management Ethernet interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
routing-instance	(Optional) Display the name of the routing instance.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 667 on page 21103 lists the output fields for the `show interfaces me0` command. Output fields are listed in the approximate order in which they appear.

Table 667: show interfaces me0 Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface: **Enabled** or **Disabled**.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Description	Optional user-specified description.	**brief detail extensive**
Type	Information about the type of functional interface.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface. The default is 1514.	All levels
Clocking	Interface that acts as a clock source. This field is not supported on EX Series switches and the default value is always **Unspecified**.	**detail extensive**
Speed	Speed at which the interface is running.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels

Table 667: show interfaces me0 Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Link type	Information about whether the link is duplex and whether the negotiation is manual or automatic.	**detail extensive** none
Physical info	Information about the device dependent physical interface selector. This field is applied only when a clocking option is specified. This field is not supported on EX Series switches and the default value is always **Unspecified**.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	MAC address of the hardware.	**detail extensive** none
Alternate link address	Information about alternate hardware address.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*****weeks*****w:*****days*****d*****hour:minute:second*** **ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3w:3d 22:58 ago).**	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface was last set to zero. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*****weeks*****w:*****days*****d*****hour:minute:second*** **ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3w:3d 22:58 ago).**	**detail extensive**

Table 667: show interfaces me0 Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. Following are fields in **Traffic statistics**: • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 transit statistics	Number and rate of bytes and IPv6 packets received and transmitted on the physical interface. Following are fields in **IPv6 transit statistics**: • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Table 667: show interfaces me0 Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame terminated and frame checksum (FCS) errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. • **Framing errors**—Number of packets received with an invalid FCS. • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**— Number of packets that exceed the size for the medium. For example, if the medium is Ethernet, the **Giant** field shows the count of packets with size greater than 1518 bytes. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 667: show interfaces me0 Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly. It increases only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increment quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame terminated and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface.	All levels

Table 667: show interfaces me0 Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	**detail extensive**
IPv6 transit statistics	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	**detail extensive**
Local statistics	Number and rate of bytes and packets destined to and exiting from the switch.	**extensive**
Protocol	Protocol family.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table **inet.0**.	**detail extensive**
Flags	Information about protocol family flags.	**detail extensive**
Input Filter	Ingress filter name.	**extensive**
Output Filter	Egress filter name.	**extensive**
Addresses	Information about the management interface addresses.	**detail extensive** none
Flags	Information about the address flags.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none

Table 667: show interfaces me0 Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces me0

```
user@switch> show interfaces me0
Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:1f:12:35:3c:bf, Hardware address: 00:1f:12:35:3c:bf
  Last flapped   : 2010-07-31 23:45:50 PDT (5d 00:32 ago)
    Input packets : 1661830
    Output packets: 3200

  Logical interface me0.0 (Index 3) (SNMP ifIndex 34)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 1661830
    Output packets: 3200
    Protocol inet
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.204.32/20, Local: 10.204.33.103,
        Broadcast: 10.204.47.255
    Protocol inet6
      Flags: Is-Primary
```

```
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21f:12ff:fe35:3cbf
```

show interfaces me0 brief

```
user@switch> show interfaces me0 brief
Physical interface: me0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface me0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  10.204.33.103/20
    inet6 fe80::21f:12ff:fe35:3cbf/64
```

show interfaces me0 detail

```
user@switch> show interfaces me0 detail
Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:35:3c:bf, Hardware address: 00:1f:12:35:3c:bf
  Alternate link address: Unspecified
  Last flapped   : 2010-07-31 23:45:50 PDT (5d 00:37 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :            366663167
   Output bytes  :               498590
   Input  packets:              1664031
   Output packets:                 3259
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
```

```
    Input  packets:                  0
    Output packets:                  0

  Logical interface me0.0 (Index 3) (SNMP ifIndex 34) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :            366665637
     Output bytes  :               500569
     Input  packets:              1664048
     Output packets:                 3275
     IPv6 transit statistics:
      Input  bytes  :                  0
      Output bytes  :                  0
      Input  packets:                  0
      Output packets:                  0
    Local statistics:
     Input  bytes  :            366665637
     Output bytes  :               500569
     Input  packets:              1664048
     Output packets:                 3275
    Protocol inet, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.204.32/20, Local: 10.204.33.103, Broadcast: 10.204.47.255, Generation: 1
    Protocol inet6, Generation: 2, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21f:12ff:fe35:3cbf
    Generation: 2
```

show interfaces me0 extensive

```
user@switch> show interfaces me0 extensive
Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
```

```
Physical info  : Unspecified
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:1f:12:38:58:bf, Hardware address: 00:1f:12:38:58:bf
Alternate link address: Unspecified
Last flapped   : 2010-08-15 06:27:33 UTC (03:06:22 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :             82310392
 Output bytes  :              1966952
 Input  packets:               110453
 Output packets:                17747
 IPv6 transit statistics:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0

Logical interface me0.0 (Index 3) (SNMP ifIndex 34) (Generation 1)
  Flags: SNMP-Traps Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :             82310392
   Output bytes  :              1966952
   Input  packets:               110453
   Output packets:                17747
  Local statistics:
   Input  bytes  :             82310392
   Output bytes  :              1966952
   Input  packets:               110453
   Output packets:                17747
  Protocol inet, Generation: 1, Route table: 0
    Flags: Is-Primary
    Input Filters: mgmt_filter,
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: 10.204.96/20, Local: 10.204.96.234,
      Broadcast: 10.204.111.255, Generation: 1
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch

Configuring Firewall Filters (CLI Procedure)

show interfaces media

IN THIS SECTION

- Syntax | 21113
- Description | 21113
- Options | 21114
- Additional Information | 21114
- Required Privilege Level | 21114
- Output Fields | 21114
- Sample Output | 21114
- Release Information | 21117

Syntax

```
show interfaces media
```

Description

Display media-specific information about all configured network interfaces.

> **NOTE**: `show interfaces media` lists details for all interfaces, whereas `show interfaces media` *`interface-name`* lists details only for the specified interface.

Options

This command has no options.

Additional Information

Output from both the `show interfaces` *`interface-name`* `detail` and the `show interfaces` *`interface-name`* `extensive` commands includes all the information displayed in the output from the `show interfaces media` command.

Required Privilege Level

view

Output Fields

The output from the `show interfaces media` command includes fields that display interface media-specific information. These fields are also included in the `show interfaces` *`interface-name`* command for each particular interface type, and the information provided in the fields is unique to each interface type.

One field unique to the `show interfaces media` command is interface-type errors (for example, SONET errors). This field appears for channelized E3, channelized T3, channelized OC, E1, E3, SONET, T1, and T3 interfaces. The information provided in this output field is also provided in the output from the `show interfaces` *`interface-name`* command. (For example, for SONET interfaces, these fields are SONET section, SONET line, and SONET path). For a description of errors, see the chapter with the particular interface type in which you are interested.

Sample Output

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the `show interfaces media` command for a SONET interface (with no level of output specified):

```
user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
```

```
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,   Loopback:
None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,   mpls: Not-
configured
  CHAP state: Not-configured
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
  SONET errors:
     BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

show interfaces media (MX Series Routers)

```
user@host>show interfaces media xe-0/0/0
Physical interface: xe-0/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 592
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Enabled
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 08:81:f4:82:a3:f0, Hardware address: 08:81:f4:82:a3:f0
  Last flapped   : 2013-10-26 03:20:40 test (1w6d 00:19 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
```

```
    Bit errors                            78
    Errored blocks                        78
  MAC statistics:
    Input bytes: 0, Input packets: 0, Output bytes: 0, Output packets: 0
  Filter statistics:
    Filtered packets: 0, Padded packets: 0, Output packet errors: 0
  Interface transmit statistics: Disabled
```

show interfaces media (PTX Series Packet Transport Routers)

```
user@host> show interfaces media em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Current address: 00:80:f9:25:00:1b, Hardware address: 00:80:f9:25:00:1b
      Last flapped   : Never
          Input packets : 215151
          Output packets: 72
```

show interfaces media (EX4100 Series Switches)

```
user@host> show interfaces ge-0/0/10 media
Physical interface: ge-0/0/10, Enabled, Physical link is Up
  Interface index: 658, SNMP ifIndex: 529
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Link-mode: Full-duplex, Speed: 1000mbps,
Duplex: Full-Duplex, BPDU Error: None, Loop Detect PDU Error: None,
  Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online,
  Media type: Copper, IEEE 802.3az Energy Efficient Ethernet: Disabled, Auto-MDIX: Enabled
  Device flags   : Present Running
  Interface Specific flags: Internal: 0x100000
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: ac:a0:9d:8f:dd:84, Hardware address: ac:a0:9d:8f:dd:84
  Last flapped   : 2024-04-01 02:39:22 PDT (01:20:29 ago)
  Input rate     : 0 bps (0 pps)
```

```
  Output rate     : 0 bps (0 pps)
  Active alarms   : None
  Active defects  : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  MAC statistics:
    Input bytes: 60200, Input packets: 175, Output bytes: 0, Output packets: 0
  PRBS Mode : Disabled
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault: OK, Link
partner Speed: 1000 Mbps
    Local resolution:
        Flow control: Symmetric, Flow control tx: Enabled, Flow control rx: Enabled, Remote
fault: Link OK, Local link Speed: 1000 Mbps, Link mode: Full-duplex
  Interface transmit statistics: Disabled
```

NOTE: For EX4100 1G copper ports, flow control Tx state is `Enabled`. The expected Tx state is `Disabled`. You see this behavior because of an internal block. The actual functionality remains the same.

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

show interfaces (Multicast Tunnel)

IN THIS SECTION

- Syntax | 21118
- Description | 21118
- Options | 21118

- Additional Information | 21119
- Required Privilege Level | 21119
- Output Fields | 21119
- Sample Output | 21121
- Release Information | 21124

Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified multicast tunnel interface and its logical encapsulation and de-encapsulation interfaces.

Options

interface-type	On M Series and T Series routers, the interface type is **mt-*fpc/pic/port***.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Additional Information

The multicast tunnel interface has two logical interfaces: encapsulation and de-encapsulation. These interfaces are automatically created by the Junos OS for every multicast-enabled VPN routing and forwarding (VRF) instance. The encapsulation interface carries multicast traffic traveling from the edge interface to the core interface. The de-encapsulation interface carries traffic coming from the core interface to the edge interface.

Required Privilege Level

view

Output Fields

Table 668 on page 21119 lists the output fields for the `show interfaces` (Multicast Tunnel) command. Output fields are listed in the approximate order in which they appear.

Table 668: Multicast Tunnel show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under Common Output Fields Description.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface.	All levels

Table 668: Multicast Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Link-level type	Encapsulation used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under Common Output Fields Description.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under Common Output Fields Description.	All levels
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	All levels

Sample Output

show interfaces (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0
Physical interface: mt-1/2/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 41
  Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
```

show interfaces brief (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0 brief
Physical interface: mt-1/2/0, Enabled, Physical link is Up
  Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
```

show interfaces detail (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0 detail
Physical interface: mt-1/2/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 41, Generation: 28
  Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
   Traffic statistics:
    Input  bytes  :            170664562               560000 bps
    Output bytes  :            112345376               368176 bps
    Input  packets:              2439107                 1000 pps
    Output packets:              2439120                 1000 pps
```

show interfaces extensive (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0 extensive
Physical interface: mt-1/2/0, Enabled, Physical link is Up
   Interface index: 141, SNMP ifIndex: 529, Generation: 144
   Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
   Hold-times     : Up 0 ms, Down 0 ms
   Device flags   : Present Running
   Interface flags: SNMP-Traps
   Statistics last cleared: Never
   Traffic statistics:
    Input  bytes  :            170664562               560000 bps
    Output bytes  :            112345376               368176 bps
    Input  packets:              2439107                 1000 pps
    Output packets:              2439120                 1000 pps
    IPv6 transit statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0

   Logical interface mt-1/2/0.32768 (Index 83) (SNMP ifIndex 556) (Generation 148)
     Flags: Point-To-Point SNMP-Traps 0x4000 IP-Header
192.0.2.1:10.0.0.6:47:df:64:0000000800000000 Encapsulation: GRE-NULL
     Traffic statistics:
      Input  bytes  :            170418430
      Output bytes  :            112070294
      Input  packets:              2434549
      Output packets:              2435593
      IPv6 transit statistics:
       Input  bytes  :                   0
       Output bytes  :                   0
       Input  packets:                   0
       Output packets:                   0
     Local statistics:
      Input  bytes  :                   0
      Output bytes  :                80442
      Input  packets:                   0
      Output packets:                1031
     Transit statistics:
      Input  bytes  :            170418430               560000 bps
      Output bytes  :            111989852               368176 bps
```

```
     Input  packets:              2434549                 1000 pps
     Output packets:              2434562                 1000 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: 1572, Generation: 182, Route table: 4
      Flags: None
    Protocol inet6, MTU: 1572, Generation: 183, Route table: 4
      Flags: None

  Logical interface mt-1/2/0.1081344 (Index 84) (SNMP ifIndex 560) (Generation 149)
    Flags: Point-To-Point SNMP-Traps 0x6000 Encapsulation: GRE-NULL
    Traffic statistics:
     Input  bytes  :               246132
     Output bytes  :               355524
     Input  packets:                 4558
     Output packets:                 4558
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :               246132
     Output bytes  :                    0
     Input  packets:                 4558
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :               355524                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                 4558                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: Unlimited, Generation: 184, Route table: 4
      Flags: None
```

```
     Protocol inet6, MTU: Unlimited, Generation: 185, Route table: 4
       Flags: None
```

show interfaces (Multicast Tunnel Encapsulation)

```
user@host> show interfaces mt-3/1/0.32768
Logical interface mt-3/1/0.32768 (Index 67) (SNMP ifIndex 0)
    Flags: Point-To-Point SNMP-Traps 0x4000
    IP-Header 198.51.100.1:10.255.70.15:47:df:64:0000000800000000
    Encapsulation: GRE-NULL
  Input packets : 0
  Output packets: 2
    Protocol inet, MTU: Unlimited
      Flags: None
```

show interfaces (Multicast Tunnel De-Encapsulation)

```
user@host> show interfaces mt-3/1/0.49152
Logical interface mt-3/1/0.49152 (Index 74) (SNMP ifIndex 0)
    Flags: Point-To-Point SNMP-Traps 0x6000 Encapsulation: GRE-NULL
  Input packets : 0
  Output packets: 2
    Protocol inet, MTU: Unlimited
      Flags: None
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces (Multilink Services)

IN THIS SECTION

Syntax | 21125

- Description | 21125
- Options | 21125
- Required Privilege Level | 21126
- Output Fields | 21126
- Sample Output | 21135
- Release Information | 21137

Syntax

```
show interfaces ml-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index>
<statistics>
```

Description

(M Series and T Series routers only) Display status information about the specified multilink services interface.

Options

ml-*fpc/pic/port*	Display standard status information about the specified multilink services interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index	(Optional) Display the SNMP index of interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 669 on page 21126 lists the output fields for the `show interfaces` (Multilink Services) command. Output fields are listed in the approximate order in which they appear.

Table 669: Multilink Services show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface: **Multilink**.	All levels
MTU	MTU size on the physical interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels

Table 669: Multilink Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour.minute.second timezone*** **(*hour.minute.second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.	**detail extensive**
Frame exceptions	Information about framing exceptions. Includes events recorded under **Exception Events** for each logical interface: • **Oversized frames**—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits). • **Errored input frames**—Number of input frame errors. • **Input on disabled link/bundle**—Number of frames received on disabled links. These can result either from an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it. • **Output for disabled link/bundle**—Number of frames sent for a disabled or unavailable link. These can result either from an inconsistent configuration, or from a bundle being brought up or down with traffic actively flowing through it. • **Queuing drops**—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed.	**extensive**

Table 669: Multilink Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Buffering exceptions	Information about buffering exceptions. Includes events recorded under **Exception Events** for each logical interface: • **Packet data buffer overflow**—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical multilink services interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical multilink services interface capacity. Check the logical interface exception event counters to determine which bundle is responsible.	**extensive**

Table 669: Multilink Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Assembly exceptions	Information about assembly exceptions. Includes events recorded under **Exception Events** for each logical interface. An assembly exception does not necessarily indicate an operational problem with the Multilink PIC itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the assembling multilink interface records one or more exception events. The multilink interface can drop multilink-encapsulated fragments itself as a result. Any multilink packets or fragments dropped by the Multilink PIC itself result in packet or fragment drop counts on individual logical interfaces. If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, excess differential delay or traffic losses on individual interfaces can be the root cause. • **Fragment timeout**—Drop-timer expired while a fragment sequence number was outstanding. Occurs only if drop-timer is enabled. This can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the Multilink PIC itself. If one or more individual links drop traffic, these events can occur. Check the logical interface exception event counters to determine which bundle is responsible. • **Missing sequence number**—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the Multilink PIC itself. If one or more individual links drop traffic, these events can occur. Check the logical interface exception event counters to determine which bundle is responsible. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers occurred within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible.	**extensive**

Table 669: Multilink Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **Out-of-range sequence number**—Frame was received with out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost, or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible.	
Hardware errors	Information about hardware errors: • **Data memory error**—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. • **Control memory error**—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support.	**extensive**
Logical Interface		
Logical interface	Logical interface name.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number.	**detail extensive** none
Encapsulation	Encapsulation being used: PPP or Multilink PPP.	All levels
Bandwidth	Speed at which the interface is running.	All levels
Flags	Logical interface flags. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	**detail extensive** none

Table 669: Multilink Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bundle options	Information about configured bundle options: • **MRRU**—Configured size of the MRRU (maximum received reconstructed unit). It can be **1500** to **4500** bytes. • **Drop timer period**—Configured drop timeout period. It can be **0** through **127** ms. A value of **0** disables the timer. The default setting is **0**. • **Sequence number format**—Configured size of the sequence header: **12** or **24** bits. The default is **24** bits. • **Fragmentation threshold**—Configured fragmentation threshold. A value of 0 results in no fragmentation. Nonzero values can be **128** through **16,320** bytes, in integer multiples of **64** bytes. The default is **0**. • **Links needed to sustain bundle**—Minimum number of links to sustain the bundle: **1** through **8**.	**detail extensive** none

Table 669: Multilink Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bundle status (MLPPP) or **Multilink class status (MC-MLPPP)**	Information about bundle status: • **Remote MRRU**—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. • **Received sequence number**—Sequence number for received packets. • **Transmitted sequence number**—Sequence number for transmitted packets. • **Packet drops**—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • **Fragment drops**—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • **MRRU exceeded**—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • **Fragment timeout**—Drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • **Missing sequence number**—Gap detected in the sequence numbers of fragments on a bundle. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream.	**detail extensive**

Table 669: Multilink Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **Out-of-range sequence number**—Frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • **Packet data buffer overflow**—Packet buffer memory full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity.	
Remote MRRU	MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed.	**detail extensive** none
Bundle errors	Information about bundle errors: • **Packet drops**—Number and byte count of output packets dropped, rather than being encapsulated and sent out of the router as fragments. • **Fragment drops**—Number and byte count of input fragments dropped, rather than being reassembled and handled by the router as packets. • **MRRU exceeded**—Number of reassembled packets exceeding the MRRU. • **Exception events**—Number of exceptional events encountered while handling traffic on the bundle, other than MRRU exceeded errors. These events are categorized under the physical interface: **Frame exceptions**, **Buffering exceptions**, and **Fragment exceptions**. Exception events do not necessarily indicate that the multilink interface is not operating properly. Individual link failures can produce exceptional events.	**detail extensive**

Table 669: Multilink Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Statistics	Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. • **Bundle**—Information about bundles. • **Link**—Information about links used in the multilink operation.	**detail extensive**
Protocol	Protocol family configured on the logical interface.	**detail extensive** none
MTU	MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked **Adjusted**.	**detail extensive** none
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Route table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none

Table 669: Multilink Services show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Broadcast	Broadcast address on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces extensive (Multilink Services)

```
user@host> show interfaces ml-0/3/0 extensive
Physical interface: ml-0/3/0, Enabled, Physical link is Up
  Interface index: 273, SNMP ifIndex: 196, Generation: 535
  Link-level type: Multilink, MTU: 4474
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2002-04-25 14:21:34 PDT (21:06:59 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 3535                    0 bps
   Output bytes  :                 4135                    0 bps
   Input  packets:                   87                    0 pps
   Output packets:                  103                    0 pps
  Frame exceptions:
    Oversized frames                 0
    Errored input frames             0
    Input on disabled link/bundle    0
    Output for disabled link/bundle  0
    Queuing drops                    0
  Buffering exceptions:
    Packet data buffer overflow      0
    Fragment data buffer overflow    0
  Assembly exceptions:
    Fragment timeout                 0
    Missing sequence number          0
    Out-of-order sequence number     0
```

```
    Out-of-range sequence number      0
  Hardware errors (sticky):
    Data memory error                 0
    Control memory error              0

  Logical interface ml-0/3/0.1 (Index 110) (SNMP ifIndex 674)
(Generation 402)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
    Bandwidth: 12288kbps
    Bundle options:
      MRRU                            1524
      Drop timer period               0
      Sequence number format          long (24 bits)
      Fragmentation threshold         0
      Links needed to sustain bundle  1
    Bundle status:
      Remote MRRU                     1500
      Received sequence number        0x19ec14
      Transmit sequence number        0x38cfa8
      Packet drops                    0 (0 bytes)
      Fragment drops                  0 (0 bytes)
      MRRU exceeded                   0
      Fragment timeout                0
      Missing sequence number         0
      Out-of-order sequence number    0
      Out-of-range sequence number    0
      Packet data buffer overflow     0
      Fragment data buffer overflow   0
    Bundle errors:
      Packet drops                    0 (0 bytes)
      Fragment drops                  0 (0 bytes)
      MRRU exceeded                   0
      Exception events                0
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Fragments:
        Input :             5          0           450            0
        Output:             6          0           499            0
      Packets:
        Input :             5          0           450            0
        Output:            12          0          1202            0
    Link:
      t1-0/1/0:11.0
```

```
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:12.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:10.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:14.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:13.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:8.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      t1-0/1/0:9.0
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet, MTU: 1500 [Adjusted], Flags: Generation: 752 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary, MTU-Protocol-Adjusted
        Destination: 198.51.100.2, Local: 198.51.100.1, Broadcast: Unspecified,
        Generation: 1090
    Protocol iso, MTU: 1500 [Adjusted], Flags: Is-Primary,
    Generation: 753 Route table: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces (PIM)

IN THIS SECTION

- Syntax | 21138
- Description | 21138

- Options | 21138
- Required Privilege Level | 21139
- Output Fields | 21139
- Sample Output | 21140
- Release Information | 21143

Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified Protocol Independent Multicast (PIM) de-encapsulation or PIM encapsulation interface, respectively.

Options

interface-type	On M Series and T Series routers, the PIM de-encapsulation interface type is **pd-*fpc/pic/port*** and the PIM encapsulation interface type is **pe-*fpc/pic/port***.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 670 on page 21139 lists the output fields for the `show interfaces` (PIM de-encapsulation or encapsulation) command. Output fields are listed in the approximate order in which they appear.

Table 670: PIM show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface.	All levels
Link-level type	Encapsulation used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels

Table 670: PIM show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Sample Output

show interfaces (PIM De-Encapsulation)

```
user@host> show interfaces pd-0/0/0
Physical interface: pd-0/0/0, Enabled, Physical link is Up
  Interface index: 130, SNMP ifIndex: 25
  Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
```

```
  Interface flags: SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
```

show interfaces brief (PIM De-Encapsulation)

```
user@host> show interfaces pd-0/0/0 brief
Physical interface: pd-0/0/0, Enabled, Physical link is Up
  Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
```

show interfaces detail (PIM De-Encapsulation)

```
user@host> show interfaces pd-0/0/0 detail
Physical interface: pd-0/0/0, Enabled, Physical link is Up
  Interface index: 130, SNMP ifIndex: 25, Generation: 11
  Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
```

show interfaces extensive (PIM Encapsulation)

```
user@host> show interfaces pd-0/0/0 extensive
Physical interface: pd-0/0/0, Enabled, Physical link is Up
  Interface index: 130, SNMP ifIndex: 25, Generation: 11
  Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
```

```
  Traffic statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
```

show interfaces (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0
Physical interface: pe-0/0/0, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 26
  Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
```

show interfaces brief (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0 brief
Physical interface: pe-0/0/0, Enabled, Physical link is Up
  Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
```

show interfaces detail (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0 detail
Physical interface: pe-0/0/0, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 26, Generation: 12
  Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
```

```
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
```

show interfaces extensive (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0 extensive
Physical interface: pe-0/0/0, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 26, Generation: 12
  Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces policers

IN THIS SECTION

- Syntax | 21144
- Description | 21144
- Options | 21144
- Additional Information | 21144
- Required Privilege Level | 21144
- Output Fields | 21144
- Sample Output | 21145

Release Information | 21146

Syntax

```
show interfaces policers
<interface-name>
```

Description

Display all policers that are installed on each interface in a system.

Options

none	Display policer information about all interfaces.
interface-name	(Optional) Display filter information about a particular interface.

Additional Information

For information about how to configure policers, see the *Junos Policy Framework Configuration Guide*. For related operational mode commands, see the *Junos Routing Protocols and Policies Command Reference*.

Required Privilege Level

view

Output Fields

Table 671 on page 21145 lists the output fields for the `show interfaces policers` command. Output fields are listed in the approximate order in which they appear.

Table 671: show interfaces policers Output Fields

Field Name	Field Description
Interface	Name of the interface.
Admin	Interface state: **up** or **down**.
Link	Link state: **up** or **down**.
Proto	Protocol configured on the interface.
Input Policer	Policer to be evaluated when packets are received on the interface. It has the format ***interface-name*-in-policer**.
Output Policer	Policer to be evaluated when packets are transmitted on the interface. It has the format ***interface-name*-out-policer.**

Sample Output

show interfaces policers

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
...
so-2/0/0        up    up
so-2/0/0.0      up    up   inet  so-2/0/0.0-in-policer so-2/0/0.0-out-policer
                           iso
```

```
so-2/1/0        up    down
...
```

show interfaces policers interface-name

```
user@host> show interfaces policers so-2/1/0
Interface       Admin Link Proto Input Policer         Output Policer
so-2/1/0        up    down
so-2/1/0.0      up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                           iso
                           inet6
```

show interfaces policers (PTX Series Packet Transport Routers)

```
user@host> show interfaces policers em0
Interface       Admin Link Proto Input Policer         Output Policer
    em0             up    up
    em0.0           up    up
                                                 inet
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

show interfaces port-profile

IN THIS SECTION

- Syntax | 21147
- Description | 21147
- Options | 21147
- Required Privilege Level | 21147

- Output Fields | 21147
- Sample Output | 21148
- Release Information | 21154

Syntax

```
show interfaces port-profile
```

Description

Display information about the operating port speed summary for the line card. You can configure channelized and non-channelized interfaces on each physical port on a PIC. The different channelized and non-channelized interfaces can operate at different speeds. Use the command to view the port speeds for the interfaces (channelized and non-channelized) configured on the line card.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 672 on page 21147lists the output fields for the `show interfaces port-profile` commands. Output fields are listed in the approximate order in which they appear.

Table 672: show interfaces port-profile Output fields

Field Name	Description
Interface	Name of the channelized or non-channelized interface.

Table 672: show interfaces port-profile Output fields *(Continued)*

Field Name	Description
Port Speed	Speed configured on the port for the interface (channelized or non-channelized).

Sample Output

show interfaces port-profile (PTX10001-36MR)

```
user@host> show interfaces port-profile
Interface       Port Speed
      et-0/0/0          400Gbps
      et-0/0/1:0        25Gbps
      et-0/0/1:1        25Gbps
      et-0/0/1:2        25Gbps
      et-0/0/1:3        25Gbps
      et-0/0/1:4        25Gbps
      et-0/0/1:5        25Gbps
      et-0/0/1:6        25Gbps
      et-0/0/1:7        25Gbps
      et-0/0/2          40Gbps
      et-0/0/3:0        10Gbps
      et-0/0/3:1        10Gbps
      et-0/0/3:2        10Gbps
      et-0/0/3:3        10Gbps
      et-0/0/4          100Gbps
      et-0/0/5          100Gbps
      et-0/0/6          100Gbps
      et-0/0/7          100Gbps
      et-0/0/8:0        100Gbps
      et-0/0/8:1        100Gbps
      et-0/0/9          400Gbps
      et-0/0/10:0       50Gbps
      et-0/0/10:1       50Gbps
      et-0/0/10:2       50Gbps
      et-0/0/10:3       50Gbps
      et-0/0/10:4       50Gbps
      et-0/0/10:5       50Gbps
      et-0/0/10:6       50Gbps
```

```
    et-0/0/10:7       50Gbps
    et-0/0/11         400Gbps
    et-0/1/0          400Gbps
    et-0/1/1:0        100Gbps
    et-0/1/1:1        100Gbps
    et-0/1/2          400Gbps
    et-0/1/3          100Gbps
    et-0/1/4          100Gbps
```

show interfaces port-profile (PTX10008 Series routers with JNP10K-LC1201)

```
user@host> show interfaces port-profile
Interface       Port Speed
et-7/0/0:0         100Gbps
et-7/0/0:1         100Gbps
et-7/0/0:2         100Gbps
et-7/0/0:3         100Gbps
et-7/0/1           100Gbps
et-7/0/2           400Gbps
et-7/0/3           400Gbps
et-7/0/4           400Gbps
et-7/0/5           400Gbps
et-7/0/6           400Gbps
et-7/0/7           400Gbps
et-7/0/8           400Gbps
et-7/0/9           400Gbps
et-7/0/10          400Gbps
et-7/0/11          400Gbps
et-7/0/12          400Gbps
et-7/0/13          400Gbps
et-7/0/14          400Gbps
et-7/0/15          400Gbps
et-7/0/16          400Gbps
et-7/0/17          400Gbps
et-7/0/18          400Gbps
et-7/0/19          400Gbps
et-7/0/20          400Gbps
et-7/0/21          400Gbps
et-7/0/22          400Gbps
et-7/0/23          400Gbps
et-7/0/24          400Gbps
```

```
et-7/0/25          400Gbps
et-7/0/26          400Gbps
et-7/0/27          400Gbps
et-7/0/28          400Gbps
et-7/0/29          400Gbps
et-7/0/30          400Gbps
et-7/0/31          400Gbps
et-7/0/32          400Gbps
et-7/0/33          400Gbps
et-7/0/34          400Gbps
et-7/0/35          400Gbps
```

show interfaces port-profile (PTX10008 Series routers with JNP10K-LC1201)

```
user@host> show interfaces port-profile
Interface       Port Speed
et-0/0/1           400Gbps
et-0/0/2           400Gbps
et-0/0/3           400Gbps
et-0/0/4           400Gbps
et-0/0/5           400Gbps
et-0/0/6           400Gbps
et-0/0/7           400Gbps
et-0/0/8           400Gbps
et-0/0/9           400Gbps
et-0/0/10          400Gbps
et-0/0/11          400Gbps
et-0/0/12          400Gbps
et-0/0/13          400Gbps
et-0/0/14          400Gbps
et-0/0/15          400Gbps
et-0/0/16          400Gbps
et-0/0/17          10Gbps
et-0/0/18          400Gbps
et-0/0/19          400Gbps
et-0/0/20          400Gbps
et-0/0/21          400Gbps
et-0/0/22          400Gbps
et-0/0/23          400Gbps
et-0/0/24          400Gbps
et-0/0/25          400Gbps
```

```
et-0/0/26          400Gbps
et-0/0/27          400Gbps
et-0/0/28          400Gbps
et-0/0/29          400Gbps
et-0/0/30          400Gbps
et-0/0/31          400Gbps
et-0/0/32          400Gbps
et-0/0/33          400Gbps
et-0/0/34          400Gbps
et-0/0/35          400Gbps
et-3/0/0           400Gbps
et-3/0/1           400Gbps
et-3/0/2           400Gbps
et-3/0/3           400Gbps
et-3/0/4           400Gbps
et-3/0/5           100Gbps
et-3/0/6           400Gbps
et-3/0/7           400Gbps
et-3/0/8           400Gbps
et-3/0/9           400Gbps
et-3/0/10          400Gbps
et-3/0/11          400Gbps
et-3/0/12          400Gbps
et-3/0/13          400Gbps
et-3/0/14          400Gbps
et-3/0/15          400Gbps
et-3/0/16          400Gbps
et-3/0/17          400Gbps
et-3/0/18          400Gbps
et-3/0/19          400Gbps
et-3/0/20          400Gbps
et-3/0/21          400Gbps
et-3/0/22          400Gbps
et-3/0/23          400Gbps
et-3/0/24          400Gbps
et-3/0/25          400Gbps
et-3/0/26          400Gbps
et-3/0/27          400Gbps
et-3/0/28          400Gbps
et-3/0/29          400Gbps
et-3/0/30          400Gbps
et-3/0/31          10Gbps
et-3/0/32          10Gbps
```

```
et-3/0/33          400Gbps
et-3/0/34          400Gbps
et-3/0/35          400Gbps
et-4/0/0           400Gbps
et-4/0/1           100Gbps
et-4/0/2           100Gbps
et-4/0/3           100Gbps
et-4/0/4           400Gbps
et-4/0/5           400Gbps
et-4/0/6           100Gbps
et-4/0/7           100Gbps
et-4/0/8:0         100Gbps
et-4/0/8:1         100Gbps
et-4/0/9           400Gbps
et-4/0/10          100Gbps
et-4/0/11          40Gbps
et-4/0/12          100Gbps
et-4/0/13          400Gbps
et-4/0/14          400Gbps
et-4/0/15          400Gbps
et-4/0/16:0        100Gbps
et-4/0/16:1        100Gbps
et-4/0/17          400Gbps
et-4/0/18          400Gbps
et-4/0/19          400Gbps
et-4/0/20          400Gbps
et-4/0/21          400Gbps
et-4/0/22          400Gbps
et-4/0/23          100Gbps
et-4/0/24          400Gbps
et-4/0/25          400Gbps
et-4/0/26          400Gbps
et-4/0/27          400Gbps
et-4/0/28          400Gbps
et-4/0/29          400Gbps
et-4/0/30          400Gbps
et-4/0/31          100Gbps
et-4/0/32          100Gbps
et-4/0/33          400Gbps
et-4/0/34          400Gbps
et-4/0/35          400Gbps
et-5/0/0           400Gbps
et-5/0/1:0         100Gbps
```

```
et-5/0/1:1          100Gbps
et-5/0/2            100Gbps
et-5/0/3            100Gbps
et-5/0/4            400Gbps
et-5/0/5            40Gbps
et-5/0/6:0          25Gbps
et-5/0/6:1          25Gbps
et-5/0/6:2          25Gbps
et-5/0/6:3          25Gbps
et-5/0/6:4          25Gbps
et-5/0/6:5          25Gbps
et-5/0/6:6          25Gbps
et-5/0/6:7          25Gbps
et-5/0/7:0          10Gbps
et-5/0/7:1          10Gbps
et-5/0/7:2          10Gbps
et-5/0/7:3          10Gbps
et-5/0/8:0          100Gbps
et-5/0/8:1          100Gbps
et-5/0/9            40Gbps
et-5/0/10           100Gbps
et-5/0/11           400Gbps
et-5/0/12           100Gbps
et-5/0/13:0         10Gbps
et-5/0/13:1         10Gbps
et-5/0/13:2         10Gbps
et-5/0/13:3         10Gbps
et-5/0/14           400Gbps
et-5/0/15           40Gbps
et-5/0/16           400Gbps
et-5/0/17           400Gbps
et-5/0/18           400Gbps
et-5/0/19           400Gbps
et-5/0/20           400Gbps
et-5/0/21           400Gbps
et-5/0/22           400Gbps
et-5/0/23           400Gbps
et-5/0/24           400Gbps
et-5/0/25           400Gbps
et-5/0/26           400Gbps
et-5/0/27           400Gbps
et-5/0/28           400Gbps
et-5/0/29           400Gbps
```

```
et-5/0/30          400Gbps
et-5/0/31          400Gbps
et-5/0/32          100Gbps
et-5/0/33          100Gbps
et-5/0/34:0        10Gbps
et-5/0/34:1        10Gbps
et-5/0/34:2        10Gbps
et-5/0/34:3        10Gbps
et-5/0/35          400Gbps
```

Release Information

Command introduced in Junos OS Evolved Release 20.1R2.

show interfaces (PPPoE)

IN THIS SECTION

- Syntax | 21154
- Description | 21155
- Options | 21155
- Required Privilege Level | 21155
- Output Fields | 21155
- Sample Output | 21165
- Release Information | 21170

Syntax

```
show interfaces pp0.logical
<brief | detail | extensive | terse>
<descriptions>
<media>
```

```
<snmp-index snmp-index>
<statistics>
```

Description

(M120 routers, M320 routers, and MX Series routers only). Display status information about the PPPoE interface.

Options

pp0.*logical*	Display standard status information about the PPPoE interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about PPPoE interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display PPPoE interface statistics.

Required Privilege Level

view

Output Fields

Table 673 on page 21155 lists the output fields for the `show interfaces` (PPPoE) command. Output fields are listed in the approximate order in which they appear.

Table 673: show interfaces (PPPoE) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical interface`	Name of the physical interface.	All levels

Table 673: show interfaces (PPPoE) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Enabled`	State of the interface. Possible values are described in the “Enabled Field” section under Common Output Fields Description.	All levels
`Interface index`	Physical interface index number, which reflects its initialization sequence.	`detail` `extensive` none
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail` `extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`
`Type`	Physical interface type (PPPoE).	All levels
`Link-level type`	Encapsulation on the physical interface (PPPoE).	All levels
`MTU`	MTU size on the physical interface.	All levels
`Clocking`	Reference clock source. It can be `Internal` or `External`.	All levels
`Speed`	Speed at which the interface is running.	All levels
`Device flags`	Information about the physical device. Possible values are described in the “Device Flags” section under Common Output Fields Description.	All levels
`Interface flags`	Information about the interface. Possible values are described in the “Interface Flags” section under Common Output Fields Description.	All levels
`Link type`	Physical interface link type: `full duplex` or `half duplex`.	All levels
`Link flags`	Information about the interface. Possible values are described in the “Link Flags” section under Common Output Fields Description.	All levels

Table 673: show interfaces (PPPoE) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input rate`	Input rate in bits per second (bps) and packets per second (pps).	None specified
`Output rate`	Output rate in bps and pps.	None specified
`Physical Info`	Physical interface information.	All levels
`Hold-times`	Current interface hold-time up and hold-time down, in milliseconds.	`detail extensive`
`Current address`	Configured MAC address.	`detail extensive`
`Hardware address`	MAC address of the hardware.	`detail extensive`
`Alternate link address`	Backup address of the link.	`detail extensive`
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`

Table 673: show interfaces (PPPoE) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. **NOTE**: These fields include dropped traffic and exception traffic, as those fields are not separately defined. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`
`Input errors`	Input errors on the interface: • `Errors`—Sum of incoming frame terminations and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Giants`—Number of frames received that are larger than the giant threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • `Resource errors`—Sum of B chip Tx drops and IXP Tx net transmit drops.	`extensive`

Table 673: show interfaces (PPPoE) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions` —Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), then the cable, the far-end system, or the PIM is malfunctioning. • `Errors`—Sum of the outgoing frame terminations and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `MTU errors`—Number of packets whose size exceeded the MTU of the interface. • `Resource errors`—Sum of B chip Tx drops and IXP Tx net transmit drops.	`extensive`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Logical interface index number (which reflects its initialization sequence).	`detail extensive` none
`SNMP ifIndex`	Logical interface SNMP interface index number.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Table 673: show interfaces (PPPoE) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under Common Output Fields Description.	All levels
`Encapsulation`	Type of encapsulation configured on the logical interface.	All levels
`PPP parameters`	PPP status: • LCP restart timer—Length of time (in milliseconds) between successive Link Control Protocol (LCP) configuration requests. • NCP restart timer—Length of time (in milliseconds) between successive Network Control Protocol (NCP) configuration requests.	`detail`
`PPPoE`	PPPoE status: • `State`—State of the logical interface (`up` or `down`). • `Session ID`—PPPoE session ID. • `Service name`—Type of service required. Can be used to indicate an Internet service provider (ISP) name or a class or quality of service. • `Configured AC name`—Configured access concentrator name. • `Auto-reconnect timeout`—Time after which to try to reconnect after a PPPoE session is terminated, in seconds. • `Idle Timeout`—Length of time (in seconds) that a connection can be idle before disconnecting. • `Underlying interface`—Interface on which PPPoE is running.	All levels
`Link`	Name of the physical interfaces for member links in an aggregated Ethernet bundle for a PPPoE over aggregated Ethernet configuration. PPPoE traffic goes out on these interfaces.	All levels

Table 673: show interfaces (PPPoE) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Traffic statistics`	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.	`detail extensive`
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. **NOTE**: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`
`Local statistics`	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.	`detail extensive`
`Transit statistics`	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize. **NOTE**: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.	`detail extensive`

Table 673: show interfaces (PPPoE) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Keepalive settings`	(PPP and HDLC) Configured settings for keepalives. • `interval` *seconds*—The time in seconds between successive keepalive requests. The range is `10` seconds through `32,767` seconds, with a default of `10` seconds. • `down-count` *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is `1` through `255`, with a default of `3`. • `up-count` *number*—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is `1` through `255`, with a default of `1`.	`detail extensive`

`Keepalive statistics`	(PPP and HDLC) Information about keepalive packets.
	• `Input`—Number of keepalive packets received by PPP.
	  • `(last seen 00:00:00 ago)`—Time the last keepalive packet was received, in the format *hh*:*mm*:*ss*.
	• `Output`—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  • `(last seen 00:00:00 ago)`—Time the last keepalive packet was sent, in the format *hh*:*mm*:*ss*.
	(MX Series routers with MPCs/MICs) When an MX Series router with MPCs/MICs is using PPP fast keepalive for a PPP link, the display does not include the number of keepalive packets received or sent, or the amount of time since the router received or sent the last keepalive packet.
		`detail extensive`

`Input packets`	Number of packets received on the logical interface.	None specified
`Output packets`	Number of packets transmitted on the logical interface.	None specified

Table 673: show interfaces (PPPoE) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`LCP state`	(PPP) Link Control Protocol state. • `Conf-ack-received`—Acknowledgement was received. • `Conf-ack-sent`—Acknowledgement was sent. • `Conf-req-sent`—Request was sent. • `Down`—LCP negotiation is incomplete (not yet completed or has failed). • `Not-configured`—LCP is not configured on the interface. • `Opened`—LCP negotiation is successful.	none `detail` `extensive`
`NCP state`	(PPP) Network Control Protocol state. • `Conf-ack-received`—Acknowledgement was received. • `Conf-ack-sent`—Acknowledgement was sent. • `Conf-req-sent`—Request was sent. • `Down`—NCP negotiation is incomplete (not yet completed or has failed). • `Not-configured`—NCP is not configured on the interface. • `Opened`—NCP negotiation is successful.	`detail extensive` none

Table 673: show interfaces (PPPoE) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`CHAP state`	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • `Chap-Chal-received`—Challenge was received but response not yet sent. • `Chap-Chal-sent`—Challenge was sent. • `Chap-Resp-received`—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • `Chap-Resp-sent`—Response was sent for the challenge received. • `Closed`—CHAP authentication is incomplete. • `Failure`—CHAP authentication failed. • `Not-configured`—CHAP is not configured on the interface. • `Success`—CHAP authentication was successful.	none `detail` `extensive`
`Protocol`	Protocol family configured on the logical interface.	`detail extensive` none
`protocol-family`	Protocol family configured on the logical interface. If the protocol is `inet`, the IP address of the interface is also displayed.	`brief`
`MTU`	MTU size on the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Route table`	Routing table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail extensive` none

Table 673: show interfaces (PPPoE) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Information about the protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	`detail extensive` none
`Addresses, Flags`	Information about the addresses configured for the protocol family. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	`detail extensive` none
`Destination`	IP address of the remote side of the connection.	`detail extensive` none
`Local`	IP address of the logical interface.	`detail extensive` none
`Broadcast`	Broadcast address.	`detail extensive` none

Sample Output

show interfaces (PPPoE)

```
user@host> show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 24
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface pp0.0 (Index 72) (SNMP ifIndex 72)
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
```

```
    PPPoE:
      State: SessionDown, Session ID: None,
      Service name: None, Configured AC name: sapphire,
      Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
      Underlying interface: at-5/0/0.0 (Index 70)
  Input packets : 0
  Output packets: 0
  LCP state: Not-configured
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,   mpls: Not-
configured
  CHAP state: Closed
    Protocol inet, MTU: 100
      Flags: User-MTU, Negotiate-Address
```

show interfaces (PPPoE over Aggregated Ethernet)

```
user@host> show interfaces pp0.1073773821
Logical interface pp0.1073773821 (Index 80) (SNMP ifIndex 32584)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 1,
      Session AC name: alcor, Remote MAC address: 00:00:5e:00:53:01,
      Underlying interface: demux0.100 (Index 88)
    Link:
      ge-1/0/0.32767
      ge-1/0/1.32767
    Input packets : 6
    Output packets: 6
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured,  mpls: Not-configured
  CHAP state: Closed
  PAP state: Success
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Primary
        Local: 203.0.113.1
```

show interfaces brief (PPPoE)

```
user@host> show interfaces pp0 brief
Physical interface: pp0, Enabled, Physical link is Up
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

  Logical interface pp0.0
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionDown, Session ID: None,
      Service name: None, Configured AC name: sapphire,
      Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
      Underlying interface: at-5/0/0.0 (Index 70)
    inet
```

show interfaces detail (PPPoE)

```
user@host> show interfaces pp0 detail
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 24, Generation: 9
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Logical interface pp0.0 (Index 72) (SNMP ifIndex 72) (Generation 14)
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
```

```
      State: SessionDown, Session ID: None,
      Service name: None, Configured AC name: sapphire,
      Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
      Underlying interface: at-5/0/0.0 (Index 70)
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
  LCP state: Not-configured
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,   mpls: Not-
configured
  CHAP state: Closed
    Protocol inet, MTU: 100, Generation: 14, Route table: 0
      Flags: User-MTU, Negotiate-Address
```

show interfaces extensive (PPPoE on M120 and M320 Routers)

```
user@host> show interfaces pp0 extensive
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 93, Generation: 129
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Statistics last cleared: Never
```

```
  Traffic statistics:
   Input  bytes  :               972192                    0 bps
   Output bytes  :               975010                    0 bps
   Input  packets:                 1338                    0 pps
   Output packets:                 1473                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0

  Logical interface pp0.0 (Index 69) (SNMP ifIndex 96) (Generation 194)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 26,
      Session AC name: None, AC MAC address: 00:00:5e:00:53:12,
      Service name: None, Configured AC name: None,
      Auto-reconnect timeout: Never, Idle timeout: Never,
      Underlying interface: ge-3/0/1.0 (Index 67)
    Traffic statistics:
     Input  bytes  :                  252
     Output bytes  :                  296
     Input  packets:                    7
     Output packets:                    8
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :                  252
     Output bytes  :                  296
     Input  packets:                    7
     Output packets:                    8
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
```

```
    Output packets:                    0                    0 pps
    IPv6 transit statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 1 (last seen 00:00:00 ago)
    Output: 1 (last sent 00:00:03 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-configured
  CHAP state: Closed
  PAP state: Closed
    Protocol inet, MTU: 1492, Generation: 171, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113.2, Local: 203.0.113.1, Broadcast: Unspecified, Generation: 206
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces prbs-stats

IN THIS SECTION

- Syntax (MX10003 and MX204) | **21171**
- Description | **21171**
- Required Privilege Level | **21171**
- Sample Output | **21171**
- Release Information | **21172**

Syntax (MX10003 and MX204)

```
show interfaces interface-name prbs-stats
```

Description

Displays the Pseudo Random Binary Sequence (PRBS) statistics and the status of the test (PASS/FAIL) along with error counters.

Use the "prbs-test-start" on page 12989 and "prbs-test-stop" on page 12991 commands to run and stop the PRBS statistics collection respectively.

For the step-by-step procedure on how to collect and view the PRBS statistics, refer *Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test*.

A 10-Gigabit ethernet interface contains a single lane of data transmission and reception, while a 40-Gigabit ethernet and 100-Gigabit ethernet interface comprises of four lanes of 10G and 25G for data transmission and reception respectively. The PRBS tests are executed per lane of an interface. Hence, The PRBS test status displays the status per lane of the interface.

Required Privilege Level

view

Sample Output

show interfaces et-0/1/2 prbs-stats (MX10003 and MX204 routers)

```
user@host> show interfaces et-0/1/2 prbs-stats
PRBS Statistics : Enabled

Lane 0 : State : Pass, Error count : 0
Lane 1 : State : Pass, Error count : 0
Lane 2 : State : Pass, Error count : 0
Lane 3 : State : Pass, Error count : 0
```

The PRBS tests are executed per lane of an interface. Hence, The PRBS test status displays the status per lane of the interface.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

- prbs-test-start | 12989
- prbs-test-stop | 12991
- clear interfaces statistics | 3651
- *Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test*

show interfaces priority-group

IN THIS SECTION

- Syntax | 21172
- Description | 21173
- Options | 21173
- Additional Information | 21173
- Required Privilege Level | 21173
- Output Fields | 21173
- Sample Output | 21174
- Release Information | 21175

Syntax

```
show interfaces priority-group
<interface-name>
<buffer-occupancy>
```

Description

Display class-of-service (CoS) priority group information for physical interfaces.

Options

interface-name (Optional) Show detailed CoS priority group statistics for the specified interface.

buffer-occupancy Displays the peak interface ingress priority-group buffer occupancy while `buffer-monitor-enable` is enabled at the `[edit chassis fpc slot-number traffic-manager]` hierarchy level.

Additional Information

For related CoS operational mode commands, see the CLI Explorer.

Required Privilege Level

view

Output Fields

Table 674 on page 21173 lists the output fields for the `show interfaces priority-group` command. Output fields are listed in the approximate order in which they appear.

Table 674: show interfaces queue Output Fields

Field Name	Field Description
`Physical interface`	Name of the physical interface.
`Enabled`	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.
`Interface index`	Physical interface's index number, which reflects its initialization sequence.
`SNMP ifIndex`	SNMP index number for the interface.

Table 674: show interfaces queue Output Fields *(Continued)*

Field Name	Field Description
PG	Priority group number.
PG-Utilization bytes	Priority group utilization in bytes.
Peak	(QFX5000 Series switches only) Diplays the peak buffer occupancy for the priority group while buffer-monitor-enable is enabled at the [edit chassis fpc *slot-number* traffic-manager] hierarchy level.

Sample Output

show interfaces priority-group buffer-occupancy

```
user@host> show interfaces priority-group xe-0/0/6 buffer-occupancy
Physical interface: xe-0/0/6, Enabled, Physical link is Down
  Interface index: 650, SNMP ifIndex: 580
            PG: 0
            PG-Utilization bytes  :
            Peak                : 88192
            PG: 1
            PG-Utilization bytes  :
            Peak                : 87984
            PG: 2
            PG-Utilization bytes  :
            Peak                : 87984
            PG: 3
            PG-Utilization bytes  :
            Peak                : 88192
            PG: 4
            PG-Utilization bytes  :
            Peak                : 88608
            PG: 5
            PG-Utilization bytes  :
            Peak                : 87776
            PG: 6
            PG-Utilization bytes  :
```

```
            Peak                : 0
            PG: 7
            PG-Utilization bytes  :
            Peak                : 0
```

Release Information

Command introduced before Junos OS Release 7.4.

`buffer-occupancy` statement introduced in Junos OS Release 19.1R1 for QFX5000 Series switches.

RELATED DOCUMENTATION

Understanding CoS Priority Group Scheduling

Defining CoS Traffic Control Profiles (Priority Group Scheduling)

show interfaces ps0 (Pseudowire Subscriber Interfaces)

IN THIS SECTION

- Syntax | 21175
- Description | 21176
- Options | 21176
- Required Privilege Level | 21176
- Output Fields | 21176
- Sample Output | 21180
- Release Information | 21182

Syntax

```
show interfaces ps0
<brief | detail | extensive | terse>
```

Description

Display status information about the pseudowire subscriber interface.

Options

brief | detail | extensive | terse (Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 675 on page 21176 lists the output fields for the `show interfaces ps0` command. Output fields are listed in the approximate order in which they appear.

Table 675: show interfaces ps0 Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical interface`	Name of the physical interface.	`brief detail extensive` none
`Enabled`	State of the interface. Possible values are described in the “Enabled Field” section under Common Output Fields Description.	`brief detail extensive` none
`Interface index`	Physical interface index number, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail extensive` none
`Type`	Physical interface type (Software-Pseudo).	`brief detail extensive` none

Table 675: show interfaces ps0 Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Link-level type`	Encapsulation being used on the physical interface.	`brief detail extensive`
`MTU`	MTU size on the physical interface.	`brief detail extensive`
`Clocking`	Reference clock source. It can be `Internal` or `External`.	`brief detail extensive`
`Speed`	Speed at which the interface is running.	`brief detail extensive`
`Device flags`	Information about the physical device. Possible values are described in the “Device Flags” section under Common Output Fields Description.	`brief detail extensive` none
`Interface flags`	Information about the interface. Possible values are described in the “Interface Flags” section under Common Output Fields Description.	`brief detail extensive` none
`Current address`	Configured MAC address.	`detail extensive` none
`Hardware address`	MAC address of the hardware.	`detail extensive` none
`Last flapped`	Date, time, and how long ago the interface went from down to up or up to down. The format is Last flapped: *year-month-day hours:minutes:seconds*: *timezone* (*hours:minutes:seconds* ago). or Never. For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).	`detail extensive` none
`input packets`	Number of packets received on the logical interface.	`detail extensive` none

Table 675: show interfaces ps0 Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`output packets`	Number of packets transmitted on the logical interface.	`detail extensive` none
Logical Interface		
`Logical interface`	Name of the logical interface.	`brief detail extensive` none
`Index`	Logical interface index number (which reflects its initialization sequence).	`detail extensive` none
`SNMP ifIndex`	Logical interface SNMP interface index number.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Flags`	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under Common Output Fields Description.	`brief detail extensive` none
`Encapsulation`	Type of encapsulation configured on the logical interface.	`brief extensive` none
`Traffic statistics`	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.	`detail extensive`

Table 675: show interfaces ps0 Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. **NOTE**: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • `Input bytes`——Number of bytes received on the interface. • `Output bytes`——Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`
`Local statistics`	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.	`detail extensive`
`Transit statistics`	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize. **NOTE**: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.	`detail extensive`
`Protocol`	Protocol family configured on the logical interface.	`detail extensive` none
`MTU`	MTU size on the logical interface.	`detail extensive` none
`Flags`	Information about the protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	`detail extensive` none

Table 675: show interfaces ps0 Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Donor interface`	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	`detail extensive` none
`Addresses, Flags`	Information about the addresses configured for the protocol family. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	`detail extensive` none
`Destination`	IP address of the remote side of the connection.	`detail extensive` none
`Local`	IP address of the logical interface.	`detail extensive` `terse` none
`Broadcast`	Broadcast address.	`detail extensive` none

Sample Output

show interfaces ps0

```
user@host> show interfaces ps0
Physical interface: ps0, Enabled, Physical link is Up
 Interface index: 166, SNMP ifIndex: 658
 Type: Software-Pseudo, Link-level type: 90, MTU: 1518, Clocking: 1, Speed: 800mbps
 Device flags : Present Running
 Interface flags: Point-To-Point Internal: 0x4000
 Current address: 00:00:5E:00:53:4a, Hardware address: 00:00:5E:00:53:4a
 Last flapped : Never
  Input packets : 0
  Output packets: 0

 Logical interface ps0.0 (Index 74) (SNMP ifIndex 656)
  Flags: Point-To-Point 0x4000 Encapsulation: Ethernet-CCC
  Input packets : 482
```

```
  Output packets: 0
  Protocol ccc, MTU: 1518
   Flags: Is-Primary

 Logical interface ps0.1 (Index 78) (SNMP ifIndex 665)
  Flags: Point-To-Point 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
  Input packets : 0
  Output packets: 482
  Protocol inet, MTU: 1500
   Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
     Destination: 203.0.113.0/24, Local: 203.0.113.1, Broadcast: 203.0.113.255

 Logical interface ps0.32767 (Index 75) (SNMP ifIndex 692)
  Flags: Point-To-Point 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
```

show interfaces ps0 extensive

```
user@host> show interfaces ps0.1 extensive
  Logical interface ps0.1 (Index 389) (SNMP ifIndex 0) (Generation 199)
    Flags: Up 0x4000 VLAN-Tag [ 0x8100.100 ]  Encapsulation: ENET2
    Traffic statistics:
      Input  bytes  :                    0
      Output bytes  :                    0
      Input  packets:                    0
      Output packets:                    0
    IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
      Input  bytes  :                    0
      Output bytes  :                    0
      Input  packets:                    0
      Output packets:                    0
    Transit statistics:
      Input  bytes  :                    0                    0 bps
      Output bytes  :                    0                    0 bps
```

```
    Input  packets:                    0                    0 pps
    Output packets:                    0                    0 pps
  IPv4 transit statistics:
    Input  bytes  :                   0                    0 bps
    Output bytes  :                   0                    0 bps
    Input  packets:                   0                    0 bps
    Output packets:                   0                    0 bps
  IPv6 transit statistics:
    Input  bytes  :                   0                    0 bps
    Output bytes  :                   0                    0 bps
    Input  packets:                   0                    0 bps
    Output packets:                   0                    0 bps
  Protocol inet, MTU: 1500, Generation: 194, Route table: 0
    Flags: Sendbcast-pkt-to-re, Unnumbered
    Donor interface: lo0.0 (Index 322)
    Addresses, Flags: Primary Is-Default Is-Primary
      Destination: Unspecified, Local: 203.0.113.144, Broadcast: Unspecified, Generation: 138
  Protocol inet6, MTU: 1500, Generation: 198, Route table: 0
    Flags: Unnumbered
    Donor interface: lo0.0 (Index 322)
      Destination: Unspecified, Local: 2001:db8::e187
  Generation: 157
```

Release Information

Command introduced at Junos OS Release 13.1.

RELATED DOCUMENTATION

MPLS Pseudowire Subscriber Logical Interfaces

show interfaces (PTX Series Packet Transport Routers)

IN THIS SECTION

Syntax | 21183

- Description | 21183
- Options | 21183
- Required Privilege Level | 21184
- Output Fields | 21184
- Sample Output | 21201
- Release Information | 21208

Syntax

```
show interfaces et-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(PTX Series Packet Transport Routers only) Display status information about the specified Ethernet interface.

Options

et-*fpc/pic/port*	Display standard information about the specified Ethernet interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

See Table 676 on page 21184 for the output fields for the `show interfaces` (PTX Series Packet Transport Routers) command.

Table 676: show interfaces PTX Series Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
BPDU Error	Bridge protocol data unit (BPDU) errors (if any).	All levels

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
MAC-Rewrite	MAC Rewrite errors (if any).	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels
Link flags	Information about the link. Possible values are described in the "Links Flags" section under *Common Output Fields Description*.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. **NOTE**: **Input bytes** and **output bytes** are counted as Layer 3 packet length.	**detail extensive**

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame terminates and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` statement. **NOTE**: The L3 incompletes field is *not* supported on PTX Series Packet Transport Routers. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame terminates and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Ingress queues	Total number of ingress queues supported on the specified interface.	**extensive**
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. • `LOCAL-FAULT`—Link fault signaling operates between the remote PHY RS (Reconciliation sub-layer) and the local RS. A Local Fault is used to signal a detected fault between the remote RS and the local RS to the local Ethernet interface. • `REMOTE-FAULT`—When the Local Fault status reaches an RS, the RS stops sending MAC data and continuously generates the Remote Fault status on the transmit data path .	**detail extensive** none

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed configured MTU. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.	**extensive**

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • **Output packet count**—Number of packets that the filter has given to the MAC hardware. • **Output packet pad count**—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • **Output packet error count**—Number of packets with an indicated error that the filter was given to transmit. These packets are usually	**extensive**

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  • **Incomplete**—Ethernet interface has the speed or link mode configured.
	  • **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner:**
	  • **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is **None**. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/ Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  • **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**—Information from the link partner:
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and
		extensive

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive). • **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).	
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	All levels

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Protocol	Protocol family. Possible values are described in the "Protocol Field" section under *Common Output Fields Description*.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set • **Input packets, Output packets**—Number of packets received and transmitted on the interface set. **NOTE**: **Input bytes** and **output bytes** are counted as Layer 3 packet length.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	**extensive**
Local statistics	Number and rate of bytes and packets destined to the router.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch.	**extensive**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive**

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	**detail extensive** none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	**detail extensive**
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	**detail extensive**
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about flags (possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none

Table 676: show interfaces PTX Series Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces brief (PTX5000 Packet Transport Router)

```
user@host> show interfaces brief et-7/0/0
Physical interface: et-7/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled, Source filtering:
Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
```

show interfaces extensive (PTX5000 Packet Transport Router)

```
user@host> show interfaces et-7/0/0 extensive
Physical interface: et-7/0/0, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 501, Generation: 171
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 88:e0:f3:3b:de:43, Hardware address: 88:e0:f3:3b:de:43
  Last flapped   : 2012-01-18 11:48:24 PST (01:51:00 ago)
  Statistics last cleared: 2012-01-18 13:38:54 PST (00:00:30 ago)
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
```

```
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
  Filter statistics:
```

```
    Input packet count                      0
    Input packet rejects                    0
    Input DA rejects                        0
    Input SA rejects                        0
    Output packet count                                    0
    Output packet pad count                                0
    Output packet error count                              0
    CAM destination filters: 0, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Incomplete
  Packet Forwarding Engine configuration:
    Destination slot: 7
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 best-effort            95     9500000000    95              0      low    none
    3 network-control         5      500000000     5              0      low    none
  Interface transmit statistics: Disabled
```

show interfaces terse (PTX5000 Packet Transport Router)

```
user@host> show interfaces terse
Interface               Admin Link Proto    Local                 Remote
et-2/0/0                up    up
et-2/0/1                up    up
et-2/0/2                up    up
et-2/0/3                up    up
et-2/0/4                up    up
et-2/0/5                up    down
et-2/0/6                up    up
et-2/0/7                up    up
et-2/0/8                up    up
et-2/0/9                up    down
et-2/0/10               up    up
et-2/0/11               up    up
et-2/0/12               up    up
et-2/0/13               up    down
et-2/0/14               up    up
et-2/0/15               up    up
et-2/0/16               up    up
```

```
et-2/0/17               up    down
et-2/0/18               up    down
et-2/0/19               up    up
et-2/0/20               up    down
et-2/0/21               up    up
et-2/0/22               up    down
et-2/0/23               up    up
et-2/1/0                up    up
et-2/1/1                up    up
et-2/1/2                up    up
et-2/1/3                up    up
et-2/1/4                up    up
et-2/1/5                up    up
et-2/1/6                up    up
et-2/1/7                up    up
et-2/1/8                up    up
et-2/1/9                up    up
et-2/1/10               up    up
et-2/1/11               up    up
et-2/1/12               up    up
et-2/1/13               up    up
et-2/1/14               up    up
et-2/1/15               up    up
et-2/1/16               up    up
et-2/1/17               up    up
et-2/1/18               up    up
et-2/1/19               up    up
et-2/1/20               up    up
et-2/1/21               up    up
et-2/1/22               up    up
et-2/1/23               up    up
et-5/0/0                up    up
et-5/0/0.0              up    up   ccc
et-5/0/0.32767          up    up   multiservice
et-5/0/1                up    up
et-5/0/2                up    up
et-5/0/3                up    down
et-5/0/4                up    down
et-5/0/5                up    up
et-5/0/5.0              up    up   ccc
et-5/0/5.32767          up    up   multiservice
et-5/0/6                up    up
et-5/0/7                up    up
```

```
et-5/0/8                up    down
et-5/0/9                up    up
et-5/0/10               up    up
et-5/0/11               up    up
et-5/0/12               up    up
et-5/0/13               up    down
et-5/0/14               up    down
et-5/0/15               up    up
et-5/0/16               up    up
et-5/0/17               up    up
et-5/0/18               up    up
et-5/0/19               up    up
et-5/0/20               up    down
et-5/0/21               up    down
et-5/0/22               up    up
et-5/0/23               up    up
et-5/1/0                up    up
et-5/1/1                up    up
et-7/0/0                up    up
et-7/0/1                up    up
et-7/0/2                up    up
et-7/0/3                up    up
et-7/0/4                up    up
et-7/0/5                up    up
et-7/0/6                up    up
et-7/0/7                up    up
et-7/0/8                up    up
et-7/0/9                up    up
et-7/0/10               up    down
et-7/0/11               up    down
et-7/0/12               up    down
et-7/0/13               up    down
et-7/0/14               up    down
et-7/0/15               up    down
et-7/0/16               up    down
et-7/0/17               up    down
et-7/0/18               up    down
et-7/0/19               up    down
et-7/0/20               up    down
et-7/0/21               up    down
et-7/0/22               up    down
et-7/0/23               up    down
dsc                     up    up
```

```
em0                     up    up
em0.0                   up    up   inet     192.168.177.61/25
gre                     up    up
ipip                    up    up
ixgbe0                  up    up
ixgbe0.0                up    up   inet     10.0.0.4/8
                                           128.0.0.1/2
                                           128.0.0.4/2
                                     inet6    fe80::200:ff:fe00:4/64
                                           fec0::a:0:0:4/64
                                     tnp      0x4
ixgbe1                  up    up
ixgbe1.0                up    up   inet     10.0.0.4/8
                                           128.0.0.1/2
                                           128.0.0.4/2
                                     inet6    fe80::200:1ff:fe00:4/64
                                           fec0::a:0:0:4/64
                                     tnp      0x4
lo0                     up    up
lo0.0                   up    up   inet     10.255.177.61       --> 0/0
                                           127.0.0.1           --> 0/0
                                     iso      47.0005.80ff.f800.0000.0108.0001.0102.5517.7061
                                     inet6    abcd::10:255:177:61
                                           fe80::ee9e:cd0f:fc02:b01e
lo0.16384               up    up   inet     127.0.0.1           --> 0/0
lo0.16385               up    up   inet
lsi                     up    up
mtun                    up    up
pimd                    up    up
pime                    up    up
tap                     up    up
```

show interfaces extensive (Junos OS Evolved)

```
user@host> show interfaces et-0/0/0 extensive
Physical interface: et-0/0/0, Enabled, Physical link is Up
  Interface index: 1002, SNMP ifIndex: 505, Generation: 113
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 100Gbps, BPDU Error: None, MAC-
REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running
```

```
  Interface flags: SNMP-Traps
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 88:e0:f3:3b:de:43, Hardware address: 88:e0:f3:3b:de:43
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets     Dropped packets
    0              16045690984503098046                    0                    0
    1              16045690984503098046                    0                    0
    2              16045690984503098046                    0                    0
    3              16045690984503098046                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  MAC statistics:                      Receive         Transmit
```

```
    Total octets                              0                0
    Total packets                             0                0
    Unicast packets                           0                0
    Broadcast packets                         0                0
    Multicast packets                         0                0
    CRC/Align errors                          0                0
    FIFO errors                               0                0
    MAC control frames                        0                0
    MAC pause frames                          0                0
    Oversized frames                          0
    Jabber frames                             0
    Fragment frames                           0
    VLAN tagged frames                        0
    Code violations                           0
    Total errors                              0                0
  Filter statistics:
    Input packet count                        0
    Input packet rejects                      0
    Input DA rejects                          0
    Input SA rejects                          0
    Output packet count                                        0
    Output packet pad count                                    0
    Output packet error count                                  0
    CAM destination filters: 0, CAM source filters: 0
```

Release Information

Command introduced in Junos OS Release 8.0.

show interfaces queue

IN THIS SECTION

- Syntax | 21209
- Description | 21209
- Options | 21209
- Overhead for Layer 2 Statistics | 21210

- Additional Information | 21213
- Required Privilege Level | 21213
- Output Fields | 21214
- Sample Output | 21222
- Release Information | 21256

Syntax

```
show interfaces queue
<aggregate>
<both-ingress-egress>
<buffer-occupancy>
<egress>
<forwarding-class forwarding-class>
<ingress>
<interface-name>
<l2-statistics>
<remaining-traffic>
<slice slice-name | detail | summary>
```

Description

Display class-of-service (CoS) queue information for physical interfaces.

Options

NOTE: This topic lists all possible options for the `show interfaces queue` command. The options that appear may vary depending on the platform and software release.

none	Show detailed CoS queue statistics for all physical interfaces.
aggregate	Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured.
both-ingress-egress	Display both ingress and egress queue statistics.

buffer-occupancy Displays the peak buffer occupancy for each queue while `buffer-monitor-enable` is enabled at the `[edit chassis fpc slot-number traffic-manager]` hierarchy level.

egress Display egress queue statistics.

forwarding-class *forwarding-class* Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.

ingress Display ingress queue statistics.

interface-name Show detailed CoS queue statistics for the specified interface. The other listed options are available for the specified interface.

l2-statistics Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles

remaining-traffic Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured.

slice *slice-name* | detail | summary *`slice-name`* displays the detailed hierarchical CoS statistics on the interface for the specified slice. `detail` displays the same statistics for all slices on the interface. `summary` displays a summary of the statistics for each slice on the interface. Both the `detail` and `summary` provide additional options, such as `forwarding-class` *`class-name`*, to filter the command results.

Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the Layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in Table 677 on page 21210. Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the Layer 3 level. In the case of link fragmentation and interleaving (LFI) for which fragmentation is not applied, corresponding Layer 2 overheads are added, as shown in Table 677 on page 21210.

Table 677: Layer 2 Overhead and Transmitted Packets or Byte Counts

Protocol	Fragmentation		LFI
	First fragmentation	Second to *n* fragmentations	
	Bytes	Bytes	
MLPPP (Long)	13	12	8

Table 677: Layer 2 Overhead and Transmitted Packets or Byte Counts *(Continued)*

Protocol	Fragmentation		LFI
	First fragmentation	Second to *n* fragmentations	
	Bytes	Bytes	
MLPPP (short)	11	10	8
MLFR (FRF15)	12	10	8
MFR (FRF16)	10	8	-
MCMLPPP(Long)	13	12	-
MCMLPPP(Short)	11	10	-

Layer 2 Statistics—Fragmentation Overhead Calculation

```
MLPPP/MC-MLPPP Overhead details:
   ==============================
   Fragment 1:

    Outer PPP header                        : 4 bytes
    Long or short sequence MLPPP header     : 4 bytes or 2 bytes
    Inner PPP header                        : 1 byte
    HDLC flag and FCS bytes                 : 4 bytes

   Fragments 2 .. n :

    Outer PPP header                        : 4 bytes
    Long or short sequence MLPPP header     : 4 bytes or 2 bytes
    HDLC flag and FCS bytes                 : 4 bytes


   MLFR (FRF15) Overhead details:
   ============================
   Fragment 1:
```

```
        Framerelay header     :  2 bytes
        Control,NLPID         :  2 bytes
        Fragmentaion header   :  2 bytes
        Inner proto           :  2 bytes
        HDLC flag and  FCS    :  4 bytes

    Fragments 2 ...n :

        Framerelay header     :  2 bytes
        Control,NLPID         :  2 bytes
        Fragmentaion header   :  2 bytes
        HDLC flag and  FCS    :  4 bytes


    MFR (FRF16) Overhead details:
    ==============================
    Fragment 1:
       Fragmentaion header   :  2 bytes
       Framerelay header     :  2 bytes
       Inner proto           :  2 bytes
       HDLC flag and  FCS    :  4 bytes

    Fragments 2 ...n :
       Fragmentaion header   :  2 bytes
       Framerelay header     :  2 bytes
       HDLC flag and  FCS    :  4 bytes
```

Overhead with LFI

```
MLPPP(Long & short sequence):
   ============================
       Outer PPP header      :  4 bytes
       HDLC flag and  FCS    :  4 bytes

  MLFR (FRF15):
  =============
       Framerelay header     :  2 bytes
       Control,NLPID         :  2 bytes
       HDLC flag and  FCS    :  4 bytes
```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the Layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the Layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the Layer 2 level, bytes transmitted is 1008 in 1 packet.

Additional Information

For rate-limited interfaces hosted on Modular Interface Cards (MICs), Modular Port Concentrators (MPCs), or Enhanced Queuing DPCs, rate-limit packet-drop operations occur *before* packets are queued for transmission scheduling. For such interfaces, the statistics for queued traffic do not include the packets that have already been dropped due to rate limiting, and consequently the displayed statistics for queued traffic are the same as the displayed statistics for transmitted traffic.

> **NOTE**: For rate-limited interfaces hosted on other types of hardware, rate-limit packet-drop operations occur *after* packets are queued for transmission scheduling. For these other interface types, the statistics for queued traffic include the packets that are later dropped due to rate limiting, and consequently the displayed statistics for queued traffic equals the sum of the statistics for transmitted and rate-limited traffic.

Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the `show interfaces statistics` command for that child interface.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the `Packet Forwarding Engine Chassis Queues` field) shows the prefragmentation values.

The behavior of the `egress` queues for the `Routing Engine-Generated Traffic` is not same as the configured queue for MLPPP and MFR configurations.

For related CoS operational mode commands, see the CLI Explorer.

Required Privilege Level

view

Output Fields

The following table lists all possible output fields for the `show interfaces queue` command. The output fields that appear may vary depending on the platform and software release.

Table 678: show interfaces queue Output Fields

Field Name	Field Description
`Physical/Anchor interface`	Name of the physical interface.
`Enabled`	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.
`Interface index`	Physical interface's index number, which reflects its initialization sequence.
`SNMP ifIndex`	SNMP index number for the interface.
`Slice`	Name of the slice.
`Slice index`	Slice's index number, which reflects its initialization sequence.
`Forwarding classes supported`	Total number of forwarding classes supported on the specified interface.
`Forwarding classes in use`	Total number of forwarding classes in use on the specified interface.
`Ingress queues supported`	Total number of ingress queues supported on the specified interface.
`Ingress queues in use`	Total number of ingress queues in use on the specified interface.

Table 678: show interfaces queue Output Fields *(Continued)*

Field Name	Field Description
`Output queues supported`	Total number of output queues supported on the specified interface.
`Output queues in use`	Total number of output queues in use on the specified interface.
`Egress queues supported`	Total number of egress queues supported on the specified interface.
`Egress queues in use`	Total number of egress queues in use on the specified interface.
`Queue counters (Ingress)`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.
`Burst size`	Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.

The following output fields are applicable to both interface component and Packet Forwarding component in the `show interfaces queue` command.

For the various queued and transmitted packets and bytes rows, the first column shows the total number of packets/bytes queued or transmitted. The second column shows the rate at which packets/bytes are queued or transmitted, in packets/bytes per second. You can reset the total packets/bytes to zero by running the `clear interfaces statistics` *`interface-name`* command.

`Queue`	Queue number.
`Forwarding classes`	Forwarding class name.

Table 678: show interfaces queue Output Fields *(Continued)*

Field Name	Field Description
`Queued Packets`	Number of packets queued to this queue. For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see "Additional Information" on page 21213.
`Queued Bytes`	Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see Table 679 on page 21218. For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see "Additional Information" on page 21213.
`Transmitted Packets`	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the `Packet Forwarding Engine Chassis Queues` field) shows the prefragmentation values. **NOTE**: For Layer 2 statistics, see "Overhead for Layer 2 Statistics" on page 21210
`Transmitted Bytes`	Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see Table 679 on page 21218. **NOTE**: On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation. **NOTE**: For Layer 2 statistics, see "Overhead for Layer 2 Statistics" on page 21210
`Tail-dropped packets`	Number of packets dropped because of tail drop.
`RL-dropped packets`	Number of packets dropped due to rate limiting. For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see "Additional Information" on page 21213.

Table 678: show interfaces queue Output Fields *(Continued)*

Field Name	Field Description
`RL-dropped bytes`	Number of bytes dropped due to rate limiting. For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see "Additional Information" on page 21213.

`RED-dropped packets`	Number of packets dropped because of random early detection (RED).
	
	• The output classifies dropped packets into the following categories:
	  • `Low`—Number of low-loss priority packets dropped because of RED.
	  • `Medium-low`—Number of medium-low loss priority packets dropped because of RED.
	  • `Medium-high`—Number of medium-high loss priority packets dropped because of RED.
	  • `High`—Number of high-loss priority packets dropped because of RED.

`RED-dropped bytes`	Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see Table 679 on page 21218.
	
	• The output classifies dropped bytes into the following categories:
	  • `Low, non-TCP`—Number of low-loss priority non-TCP bytes dropped because of RED.
	  • `Low, TCP`—Number of low-loss priority TCP bytes dropped because of RED.
	  • `High, non-TCP`—Number of high-loss priority non-TCP bytes dropped because of RED.
	  • `High, TCP`—Number of high-loss priority TCP bytes dropped because of RED.

`ECN-CE packets`	The number of packets that experienced congestion, either on a specific queue or on another device prior to being received.
`ECN-CE bytes`	The number of bytes that experienced congestion, either on a specific queue or on another device prior to being received.

Table 678: show interfaces queue Output Fields *(Continued)*

Field Name	Field Description
`Queue Buffer Usage:`	• `Reserved buffer`—The size of the memory buffer that is allocated for storing packets • `Current`—The amount of buffer memory that is currently in use on this queue.
`Queue-depth bytes`	Displays the amount of queue buffer that is in occupation at this instance. This is an indicator of the amount of data that is present in a queue at that point in time. The amount of data present is in the units of bytes.
`Peak`	Diplays the peak buffer occupancy for the queue while `buffer-monitor-enable` is enabled at the `[edit chassis fpc` *`slot-number`* `traffic-manager]` hierarchy level.
`Last-packet enqueued`	If `packet-timestamp` is enabled for an FPC, shows the day, date, time, and year in the format *day-of-the-week month day-date hh:mm:ss yyyy* when a packet was enqueued in the CoS queue. When the timestamp is aggregated across all active Packet Forwarding Engines, the latest timestamp for each CoS queue is reported.

Byte counts vary by interface hardware. Table 679 on page 21218 shows how the byte counts on the outbound interfaces vary depending on the interface hardware. Table 679 on page 21218 is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 679: Byte Count by Interface Hardware

Interface Hardware	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes	The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type. For RED dropped, 6 bytes are added for the source MAC address.

Table 679: Byte Count by Interface Hardware *(Continued)*

Interface Hardware	Output Level	Byte Count Includes	Comments
	Packet forwarding component	Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet	–

Table 679: Byte Count by Interface Hardware *(Continued)*

Interface Hardware	Output Level	Byte Count Includes	Comments
Non-IQ PIC	Interface	T Series, TX Series, T1600, and MX Series routers: • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet. T4000 routers with Type 5 FPCs : • Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap. • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap. M Series routers: • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead. PTX Series Packet Transport Routers: • Queued: The sum of the transmitted bytes and the RED dropped bytes. • Transmitted: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble. • RED dropped: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble (does not include the VLAN header or MPLS pushed bytes).	The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.

Table 679: Byte Count by Interface Hardware *(Continued)*

Interface Hardware	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes	For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.
Non-IQ PIC with a SONET/SDH interface	Interface	T Series, TX Series, T1600, and MX Series routers: • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet. M Series routers: • Queued: 478 bytes of Layer 3 packet. • Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes • RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet	For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP).
Interfaces configured with Frame Relay Encapsulation	Interface	The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.	

Table 679: Byte Count by Interface Hardware *(Continued)*

Interface Hardware	Output Level	Byte Count Includes	Comments
1-port 10-Gigabit Ethernet IQ2 and IQ2–E PICs 4-port 1G IQ2 and IQ2-E PICs	Interface	Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC. Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.	The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.
8-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet.	–

Sample Output

show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC)

The following example shows queue information for the rate-limited interface ge-4/2/0 on a Gigabit Ethernet MIC in an MPC. For rate-limited queues for interfaces hosted on MICs or MPCs, rate-limit packet drops occur prior to packet output queuing. In the command output, the nonzero statistics displayed in the `RL-dropped packets` and `RL-dropped bytes` fields quantify the traffic dropped to rate-limit queue 0 output to 10 percent of 1 gigabyte (100 megabits) per second. Because the RL-dropped traffic is not included in the `Queued` statistics, the statistics displayed for queued traffic are the same as the statistics for transmitted traffic.

```
user@host> show interfaces queue ge-4/2/0
Physical interface: ge-4/2/0, Enabled, Physical link is Up
  Interface index: 203, SNMP ifIndex: 1054
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :            131300649                141751 pps
    Bytes                :          11287964840              99793248 bps
  Transmitted:
    Packets              :            131300649                141751 pps
    Bytes                :          11287964840              99793248 bps
    Tail-dropped packets :                    0                     0 pps
    RL-dropped packets   :            205050862                602295 pps
```

```
    RL-dropped bytes     :           13595326612             327648832 bps
    RED-dropped packets  :                     0                     0 pps
     Low                 :                     0                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low                 :                     0                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
```

show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```
user@host> show interfaces queue ge-2/2/9 aggregate
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :             148450735                947295 pps
    Bytes                :            8016344944             409228848 bps
  Transmitted:
    Packets              :              76397439                487512 pps
    Bytes                :            4125461868             210602376 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :              72053285                459783 pps
     Low                 :              72053285                459783 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :            3890877444             198626472 bps
     Low                 :            3890877444             198626472 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
```

```
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :             410278257               473940 pps
    Bytes                :           22156199518            204742296 bps
  Transmitted:
    Packets              :               4850003                 4033 pps
    Bytes                :             261900162              1742256 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :             405425693               469907 pps
     Low                 :             405425693               469907 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :           21892988124            203000040 bps
     Low                 :           21892988124            203000040 bps
     Medium-low          :                     0                    0 bps
     Medium-high         :                     0                    0 bps
     High                :                     0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
```

```
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     Medium-low          :                     0                    0 bps
     Medium-high         :                     0                    0 bps
     High                :                     0                    0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              76605230               485376 pps
    Bytes                :            5209211400            264044560 bps
  Transmitted:
    Packets              :              76444631               484336 pps
    Bytes                :            5198235612            263478800 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                160475                 1040 pps
     Low                 :                160475                 1040 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :              10912300               565760 bps
     Low                 :              10912300               565760 bps
     Medium-low          :                     0                    0 bps
     Medium-high         :                     0                    0 bps
     High                :                     0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
  Transmitted:
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
```

```
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :              4836136                 3912 pps
    Bytes                :            333402032              2139056 bps
  Transmitted:
    Packets              :              3600866                 1459 pps
    Bytes                :            244858888               793696 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :              1225034                 2450 pps
     Low                 :              1225034                 2450 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :             83302312              1333072 bps
     Low                 :             83302312              1333072 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
```

```
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              77059796                486384 pps
    Bytes                :            3544750624             178989576 bps
  Transmitted:
    Packets              :              77059797                486381 pps
    Bytes                :            3544750670             178988248 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
     Low                 :                     0                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low                 :                     0                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
     Low                 :                     0                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low                 :                     0                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :               4846580                  3934 pps
    Bytes                :             222942680               1447768 bps
  Transmitted:
```

```
    Packets              :               4846580                  3934 pps
    Bytes                :             222942680               1447768 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
     Low                 :                     0                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low                 :                     0                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
     Low                 :                     0                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low                 :                     0                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
```

show interfaces queue (Gigabit Ethernet IQ2 PIC)

```
user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in use Ingress
queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :             418390039                    10 pps
```

```
    Bytes                :           38910269752                  7440 bps
  Transmitted:
    Packets              :             418390039                    10 pps
    Bytes                :           38910269752                  7440 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                  7055                     1 pps
    Bytes                :                451552                   512 bps
  Transmitted:
    Packets              :                  7055                     1 pps
    Bytes                :                451552                   512 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :                  1031                     0 pps
    Bytes                :                143292                     0 bps
  Transmitted:
```

```
    Packets              :                  1031                     0 pps
    Bytes                :                143292                     0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                 77009                    11 pps
    Bytes                :               6894286                  7888 bps
  Transmitted:
    Packets              :                 77009                    11 pps
    Bytes                :               6894286                  7888 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
```

```
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :                  1031                     0 pps
    Bytes                :                147328                     0 bps
  Transmitted:
    Packets              :                  1031                     0 pps
    Bytes                :                147328                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
     Low, non-TCP        :                     0                     0 pps
     Low, TCP            :                     0                     0 pps
     High, non-TCP       :                     0                     0 pps
     High, TCP           :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low, non-TCP        :                     0                     0 bps
     Low, TCP            :                     0                     0 bps
     High, non-TCP       :                     0                     0 bps
     High, TCP           :                     0                     0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
     Low, non-TCP        :                     0                     0 pps
     Low, TCP            :                     0                     0 pps
     High, non-TCP       :                     0                     0 pps
     High, TCP           :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low, non-TCP        :                     0                     0 bps
     Low, TCP            :                     0                     0 bps
     High, non-TCP       :                     0                     0 bps
     High, TCP           :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
```

```
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
     Low, non-TCP        :                    0                    0 pps
     Low, TCP            :                    0                    0 pps
     High, non-TCP       :                    0                    0 pps
     High, TCP           :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low, non-TCP        :                    0                    0 bps
     Low, TCP            :                    0                    0 bps
     High, non-TCP       :                    0                    0 bps
     High, TCP           :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                94386                   12 pps
    Bytes                :             13756799                 9568 bps
  Transmitted:
    Packets              :                94386                   12 pps
    Bytes                :             13756799                 9568 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
     Low, non-TCP        :                    0                    0 pps
     Low, TCP            :                    0                    0 pps
     High, non-TCP       :                    0                    0 pps
     High, TCP           :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low, non-TCP        :                    0                    0 bps
     Low, TCP            :                    0                    0 bps
     High, non-TCP       :                    0                    0 bps
     High, TCP           :                    0                    0 bps
```

show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```
user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
```

```
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                  254                    0 pps
    Bytes                :                16274                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
```

```
    Packets              : Not Available
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     3                     0 pps
    Bytes                :                   126                     0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              : Not Available
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              80564692                     0 pps
```

```
    Bytes                :            3383717100                     0 bps
  Transmitted:
    Packets              :              80564692                     0 pps
    Bytes                :            3383717100                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :              80564685                     0 pps
    Bytes                :            3383716770                     0 bps
  Transmitted:
    Packets              :              80564685                     0 pps
    Bytes                :            3383716770                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                  9397                     0 pps
    Bytes                :               3809052                   232 bps
  Transmitted:
    Packets              :                  9397                     0 pps
    Bytes                :               3809052                   232 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
```

show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```
user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                  288                    0 pps
    Bytes                :                18450                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
```

```
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
```

show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```
user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    3                    0 pps
    Bytes                :                  126                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
```

```
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              : Not Available
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              80564692                     0 pps
    Bytes                :            3383717100                     0 bps
  Transmitted:
    Packets              :              80564692                     0 pps
    Bytes                :            3383717100                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :              80564685                     0 pps
    Bytes                :            3383716770                     0 bps
  Transmitted:
    Packets              :              80564685                     0 pps
    Bytes                :            3383716770                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
```

```
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                 9538                    0 pps
    Bytes                :              3819840                    0 bps
  Transmitted:
    Packets              :                 9538                    0 pps
    Bytes                :              3819840                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
```

show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```
user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :            110208969               472875 pps
    Bytes                :           5951284434            204282000 bps
  Transmitted:
    Packets              :            110208969               472875 pps
    Bytes                :           5951284434            204282000 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
```

```
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
```

```
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :            109355853               471736 pps
    Bytes                :           7436199152            256627968 bps
  Transmitted:
    Packets              :            109355852               471736 pps
    Bytes                :           7436198640            256627968 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
```

```
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
```

show interfaces queue (QFX Series)

```
user@switch> show interfaces queue xe-0/0/15
   Physical interface: xe-0/0/15, Enabled, Physical link is Up
     Interface index: 49165, SNMP ifIndex: 539
   Forwarding classes: 12 supported, 8 in use
   Egress queues: 12 supported, 8 in use
   Queue: 0, Forwarding classes: best-effort
     Queued:
       Packets              :                     0                     0 pps
       Bytes                :                     0                     0 bps
     Transmitted:
       Packets              :                     0                     0 pps
       Bytes                :                     0                     0 bps
       Tail-dropped packets :                     0                     0 pps
       Tail-dropped bytes   :                     0                     0 bps
       RED-dropped packets  :                     0                     0 pps
       RED-dropped bytes    :                     0                     0 bps
       ECN-CE packets       :                     0                     0 pps
       ECN-CE bytes         :                     0                     0 bps
       Total-dropped packets:                     0                     0 pps
       Total-dropped bytes  :                     0                     0 bps
   Queue: 3, Forwarding classes: fcoe
     Queued:
       Packets              :                     0                     0 pps
       Bytes                :                     0                     0 bps
     Transmitted:
       Packets              :                     0                     0 pps
       Bytes                :                     0                     0 bps
       Tail-dropped packets : Not Available
       Total-dropped packets:                     0                     0 pps
       Total-dropped bytes  :                     0                     0 bps
   Queue: 4, Forwarding classes: no-loss
     Queued:
       Packets              :                     0                     0 pps
       Bytes                :                     0                     0 bps
     Transmitted:
       Packets              :                     0                     0 pps
       Bytes                :                     0                     0 bps
       Tail-dropped packets : Not Available
       Total-dropped packets:                     0                     0 pps
       Total-dropped bytes  :                     0                     0 bps
```

```
  Queue: 7, Forwarding classes: network-control
    Queued:
      Packets              :                    0                    0 pps
      Bytes                :                    0                    0 bps
    Transmitted:
      Packets              :                    0                    0 pps
      Bytes                :                    0                    0 bps
      Tail-dropped packets : Not Available
      Total-dropped packets:                    0                    0 pps
      Total-dropped bytes  :                    0                    0 bps
  Queue: 8, Forwarding classes: mcast
    Queued:
      Packets              :                    0                    0 pps
      Bytes                :                    0                    0 bps
    Transmitted:
      Packets              :                    0                    0 pps
      Bytes                :                    0                    0 bps
      Tail-dropped packets : Not Available
      Total-dropped packets:                    0                    0 pps
      Total-dropped bytes  :                    0                    0 bps
```

show interfaces queue l2-statistics (lsq interface)

```
user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets              :                    1                    0 pps
    Bytes                :                 1001                    0 bps
  Transmitted:
    Packets              :                    5                     0 pps
    Bytes                :                 1062                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets              :                    1                    0 pps
```

```
    Bytes                :                 1500                     0 bps
  Transmitted:
    Packets              :                    6                      0 pps
    Bytes                :                 1573                     0 bps
    Tail-dropped packets :                    0                     0 pps
    RED-dropped packets  :                    0                     0 pps
    RED-dropped bytes    :                    0                     0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets              :                    1                     0 pps
    Bytes                :                  512                     0 bps
  Transmitted:
    Packets              :                    3                      0 pps
    Bytes                :                  549                     0 bps
    Tail-dropped packets :                    0                     0 pps
    RED-dropped packets  :                    0                     0 pps
    RED-dropped bytes    :                    0                     0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                      0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets :                    0                     0 pps
    RED-dropped packets  :                    0                     0 pps
    RED-dropped bytes    :                    0                     0 bps
=========
```

show interfaces queue lsq (lsq-ifd)

```
user@switch> show interfaces queue lsq-1/0/0
Logical interface lsq-1/0/0 (Index 348) (SNMP ifIndex 660)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets              :                 55576                  1206 pps
    Bytes                :              29622008               5145472 bps
  Transmitted:
```

```
    Packets              :                 55576                  1206 pps
    Bytes                :              29622008               5145472 bps
    Tail-dropped packets :                     0                     0 pps
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    RED-dropped packets  :                     0                     0 pps
     Low                 :                     0                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low                 :                     0                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    RED-dropped packets  :                     0                     0 pps
     Low                 :                     0                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low                 :                     0                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
```

```
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets              :                22231                  482 pps
    Bytes                :             11849123              2057600 bps
  Transmitted:
    Packets              :                22231                  482 pps
    Bytes                :             11849123              2057600 bps
    Tail-dropped packets :                    0                    0 pps
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
```

show interfaces queue (Aggregated Ethernet on a MX series Router)

```
user@host> show interfaces queue ae0 remaining-traffic
Physical interface: ae0    , Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 543
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
```

```
  Queued:
    Packets              :                    16                    0 pps
    Bytes                :                  1896                    0 bps
  Transmitted:
    Packets              :                    16                    0 pps
    Bytes                :                  1896                    0 bps
    Tail-dropped packets :                     0                    0 pps
    RL-dropped packets   :                     0                    0 pps
    RL-dropped bytes     :                     0                    0 bps
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     Medium-low          :                     0                    0 bps
     Medium-high         :                     0                    0 bps
     High                :                     0                    0 bps
  Queue-depth bytes      :
    Average              :                     0
    Current              :                     0
    Peak                 :                     0
    Maximum              :             119013376
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
  Transmitted:
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
    Tail-dropped packets :                     0                    0 pps
    RL-dropped packets   :                     0                    0 pps
    RL-dropped bytes     :                     0                    0 bps
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     Medium-low          :                     0                    0 bps
     Medium-high         :                     0                    0 bps
```

```
     High                :                    0                    0 bps
  Queue-depth bytes      :
    Average              :                    0
    Current              :                    0
    Peak                 :                    0
    Maximum              :                32768
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
  Queue-depth bytes      :
    Average              :                    0
    Current              :                    0
    Peak                 :                    0
    Maximum              :                32768
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
```

```
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
  Queue-depth bytes      :
    Average              :                    0
    Current              :                    0
    Peak                 :                    0
    Maximum              :              6258688
```

show interfaces queue xe-0/0/2 buffer-occupancy (QFX5000 Series switch)

```
user@switch> show interfaces queue xe-0/0/2 buffer-occupancy
Physical interface: xe-0/0/2, Enabled, Physical link is Up
  Interface index: 689, SNMP ifIndex: 514
Forwarding classes: 12 supported, 5 in use
Egress queues: 12 supported, 5 in use
            Queue: 0, Forwarding classes: fc0
            Queue-depth bytes  :
            Peak               : 1786720
            Queue: 3, Forwarding classes: fcoe
            Queue-depth bytes  :
            Peak               : 0
            Queue: 4, Forwarding classes: no-loss
            Queue-depth bytes  :
            Peak               : 0
            Queue: 7, Forwarding classes: network-control
            Queue-depth bytes  :
            Peak               : 416
            Queue: 8, Forwarding classes: mcast
            Queue-depth bytes  :
            Peak               : 0
```

show interfaces queue xe-3/0/3 slice slice1

NOTE: `show interfaces queue xe-3/0/3 slice detail` shows the same level of information as below for all slices on the xe-3/0/3 interface.

```
user@router show interfaces queue xe-3/0/3 slice slice1
Physical interface: xe-3/0/3, Enabled, Physical link is Up
  Interface index: 275, SNMP ifIndex: 776
  Slice name: slice1, Slice index: 5
Forwarding classes: 16 supported, 7 in use
Egress queues: 8 supported, 7 in use
CoS scheduler resource information:
  Maximum units supported per MIC/PIC: 84
  Configured units per MIC/PIC: 1
  Maximum units allowed per port: 21
  Configured units on this port: 1
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              84374500                    0 pps
    Bytes                :            8099952000                    0 bps
  Transmitted:
    Packets              :               6594962                    0 pps
    Bytes                :             633116352                    0 bps
    Tail-dropped packets :              77779538                    0 pps
    RL-dropped packets   :                     0                    0 pps
    RL-dropped bytes     :                     0                    0 bps
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     Medium-low          :                     0                    0 bps
     Medium-high         :                     0                    0 bps
     High                :                     0                    0 bps
  Queue-depth bytes      :
    Average              :               5963760
    Current              :                     0
    Peak                 :               5963760
    Maximum              :               5996544
```

```
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
. . .
```

show interfaces queue xe-4/0/5 slice summary

This sample shows two network slices, `slice1` and `slice2`, which are anchored on `xe-4/0/5`. The counters here are the sum of corresponding counters across all queues of each slice on this physical interface. For example, the queued packets counter of `slice2` on `xe-4/0/5` here is the sum of queued packets across all the queues/forwarding-classes of `slice2` on `xe-4/0/5`.

```
user@router show interfaces queue xe-4/0/5 slice summary
Slice : slice1 (Index : 1)
  Anchor interface : xe-4/0/5 (Index : 191)
  Aggregated Egress Queue Counters:
    Queued:
      Packets       :              84374500                    0 pps
      Bytes         :            8099952000                    0 bps
    Transmitted:
      Packets       :               6594962                    0 pps
      Bytes         :             633116352                    0 bps
    Dropped:
      Packets       :              77779538                    0 pps
      Bytes         :            7466835648                    0 bps
Slice : slice2 (Index : 2)
  Anchor interface : xe-4/0/5 (Index : 191)
  Aggregated Egress Queue Counters:
    Queued:
      Packets       :             168748911                    0 pps
      Bytes         :           16199895456                    0 bps
    Transmitted:
      Packets       :              13190364                    0 pps
      Bytes         :            1266274944                    0 bps
    Dropped:
```

```
      Packets       :             155558547                    0 pps
      Bytes         :           14933620512                    0 bps
```

show interfaces queue xe-4/0/5 slice detail forwarding-class assured-forwarding

This sample shows two network slices, `slice1` and `slice2`, which are anchored on `xe-4/0/5`. The counters here are for the `assured-forwarding` forwarding class of each slice on this physical interface.

```
user@router show interfaces queue xe-4/0/5 slice detail forwarding-class assured-forwarding
Slice : slice1 (Index : 1)
  Anchor interface : xe-4/0/5 (Index : 191)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
  Queue-depth bytes      :
    Average              :                    0
    Current              :                    0
    Peak                 :                    0
    Maximum              :                33280
Slice : slice2 (Index : 2)
  Anchor interface : xe-4/0/5 (Index : 191)
Forwarding classes: 16 supported, 4 in use
```

```
Egress queues: 8 supported, 4 in use
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :              84374411                    0 pps
    Bytes                :            8099943456                    0 bps
  Transmitted:
    Packets              :                159290                    0 pps
    Bytes                :              15291840                    0 bps
    Tail-dropped packets :              84215121                    0 pps
    RL-dropped packets   :                     0                    0 pps
    RL-dropped bytes     :                     0                    0 bps
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     Medium-low          :                     0                    0 bps
     Medium-high         :                     0                    0 bps
     High                :                     0                    0 bps
  Queue-depth bytes      :
    Average              :                 32976
    Current              :                     0
    Peak                 :                 32976
    Maximum              :                 33280
```

show interfaces queue (vSRX Virtual Firewall)

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```
user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 135, SNMP ifIndex: 510
Forwarding classes: 8 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :                 14686                    0 pps
    Bytes                :                616812                    0 bps
```

```
  Transmitted:
    Packets              :                 14686                    0 pps
    Bytes                :                616812                    0 bps
    Tail-dropped packets :                     0                    0 pps
    RL-dropped packets   :                     0                    0 pps
    RL-dropped bytes     :                     0                    0 bps
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     Medium-low          :                     0                    0 bps
     Medium-high         :                     0                    0 bps
     High                :                     0                    0 bps
  Queue Buffer Usage:
    Reserved buffer      :             118750000 bytes
  Queue-depth bytes      :
    Current              :                     0
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
  Transmitted:
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
    Tail-dropped packets :                     0                    0 pps
    RL-dropped packets   :                     0                    0 pps
    RL-dropped bytes     :                     0                    0 bps
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     Medium-low          :                     0                    0 bps
     Medium-high         :                     0                    0 bps
     High                :                     0                    0 bps
  Queue Buffer Usage:
    Reserved buffer      :                  9192 bytes
  Queue-depth bytes      :
```

```
    Current              :                    0
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
  Queue Buffer Usage:
    Reserved buffer      :                 9192 bytes
  Queue-depth bytes      :
    Current              :                    0
...
```

Release Information

Command introduced before Junos OS Release 7.4.

`both-ingress-egress`, `egress`, and `ingress` options introduced in Junos OS Release 7.6.

`l2-statistics` option introduced in Junos OS Release 12.1.

`buffer-occupancy` statement introduced in Junos OS Release 19.1R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
22.3R1	Starting with Junos OS 24.2R1, the `slice summary` and `slice detail` options are supported on MX Series Routers.
22.3R1	Starting with Junos OS 22.3R1, the `slice` *`slice-name`* option is supported on MX Series Routers.
18.3R1	Starting with Junos OS 18.3R1, the `Tail-dropped packets` counter is supported on PTX Series Packet Transport Routers.
16.1	Starting with Junos OS Release 16.1, `Last-packet enqueued` output field is introduced.

RELATED DOCUMENTATION

- Monitoring Interface Status and Traffic
- *Monitoring Interfaces That Have CoS Components*
- *Defining CoS Schedulers and Scheduler Maps (CLI Procedure)*
- ATM Interfaces User Guide for Routing Devices
- Layer 2 Bridging, Address Learning, and Forwarding User Guide

show interfaces queue fabric

IN THIS SECTION

- Syntax | **21258**
- Description | **21258**
- Options | **21258**
- Required Privilege Level | **21258**
- Output Fields | **21258**
- Sample Output | **21260**
- Release Information | **21288**

Syntax

```
show interfaces queue fabric
<egress>
<forwarding-class forwarding-class>
<interface-name interface-name>
```

Description

Display class-of-service (CoS) queue information for the fabric interfaces that are configured between Node devices and Interconnect devices.

Options

none	Show detailed CoS queue statistics for all physical interfaces.
egress	(Optional) Display egress queue statistics.
forwarding-class ***forwarding-class***	(Optional) Forwarding class name for this queue. Show detailed CoS statistics for the queue associated with the specified forwarding class.
interface-name ***interface-name***	(Optional) Show detailed CoS queue statistics for the specified interface.

Required Privilege Level

view

Output Fields

Table 680 on page 21258 lists the output fields for the `show interfaces queue fabric` command. Output fields are listed in the approximate order in which they appear.

Table 680: show interfaces queue fabric Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.

Table 680: show interfaces queue fabric Output Fields *(Continued)*

Field Name	Field Description
Enabled	State of the interface. Possible values are: • Administratively down, Physical link is Down–The interface is turned off, and the physical link is inoperable. • Administratively down, Physical link is Up–The interface is turned off, but the physical link is operational and can pass packets when it is enabled. • Enabled, Physical link is Down–The interface is turned on, but the physical link is inoperable and cannot pass packets. • Enabled, Physical link is Up–The interface is turned on, and the physical link is operational and can pass packets.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the physical interface.
Forwarding classes	Number of forwarding classes supported and in use for the interface.
Egress queues	Number of output queues supported and in use on the specified interface.
Queue	CoS queue number.
Transmitted	Number of packets and bytes transmitted by this queue. Information on transmitted packets and bytes can include: • Packets–Number of packets transmitted. • Bytes–Number of bytes transmitted. • Tail-dropped packets–Number of arriving packets dropped because output queue buffers were full. • Total-dropped pkts–Number of transmitted packets dropped. • Total dropped bytes–Number of transmitted bytes dropped.

Table 680: show interfaces queue fabric Output Fields *(Continued)*

Field Name	Field Description
Queued	Number of packets and bytes queued to this queue. • Packets–Number of packets queued. • Bytes–Number of bytes queued.

Sample Output

show interfaces queue fabric

```
user@switch> show interfaces queue fabric
Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Up
  Interface index: 49178, SNMP ifIndex: 1208484475
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :             62665971                     0 pps
    Bytes                :           7770580404                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
```

```
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
```

```
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
```

```
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Up
  Interface index: 49211, SNMP ifIndex: 1208484377
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
```

```
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
```

```
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
```

```
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Up
  Interface index: 49212, SNMP ifIndex: 1208484365
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
```

```
    Total-dropped bytes  :                    0                    0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
```

```
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
```

```
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Up
  Interface index: 49213, SNMP ifIndex: 1208484625
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets              :                    0                    0 pps
```

```
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
```

```
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
```

```
    Total-dropped bytes  :                    0                    0 bps

Physical interface: P2659-C:fte-0/1/2, Enabled, Physical link is Up
  Interface index: 49161, SNMP ifIndex: 1209008630
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 7, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
```

```
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 8, Forwarding classes: mcast
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
```

show interfaces queue fabric egress

```
user@switch> show interfaces queue fabric egress
Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Up
  Interface index: 49178, SNMP ifIndex: 1208484475
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :             62665971                     0 pps
    Bytes                :           7770580404                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
```

```
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets              :                    0                    0 pps
```

```
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
```

```
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Up
  Interface index: 49211, SNMP ifIndex: 1208484377
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
```

```
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
```

```
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
```

```
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Up
  Interface index: 49212, SNMP ifIndex: 1208484365
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
```

```
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
```

```
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
```

```
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Up
  Interface index: 49213, SNMP ifIndex: 1208484625
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                     0 pps
    Total-dropped bytes  :                    0                     0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
```

```
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
```

```
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
```

```
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps

Physical interface: P2659-C:fte-0/1/2, Enabled, Physical link is Up
  Interface index: 49161, SNMP ifIndex: 1209008630
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 7, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
```

```
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 8, Forwarding classes: mcast
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
```

show interfaces queue fabric interface-name egress

```
user@switch> show interfaces queue fabric BBAK0394:fte-0/1/0 egress
Physical interface: BBAK0394:fte-0/1/0, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 1091568120 Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :           74777763341                 844587 pps
    Bytes                :         9272442654284             837830728 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
```

```
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 7, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
Queue: 8, Forwarding classes: mcast
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
```

show interfaces queue fabric interface-name egress forwarding-class forwarding-class-name

```
user@switch> show interfaces queue fabric BBAK0394:fte-0/1/0 egress forwarding-class best-effort
Physical interface: BBAK0394:fte-0/1/0, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 1091568120 Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :                    0                    0 pps
```

```
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :           74793424543                844612 pps
    Bytes                :         9274384643332            837855936 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                    0                    0 pps
    Total-dropped bytes  :                    0                    0 bps
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

show interfaces fabric | 20774

show interfaces redundancy

IN THIS SECTION

- Syntax | 21288
- Description | 21289
- Options | 21289
- Required Privilege Level | 21289
- Output Fields | 21289
- Sample Output | 21290
- Release Information | 21292

Syntax

```
show interfaces redundancy
<brief | detail>
```

Description

(M Series, T Series, and MX Series routers only) Display general information about redundancy for aggregated multiservices (AMS) interfaces configured for warm standby, adaptive services and link services intelligent queuing (IQ) interfaces, aggregated Ethernet interfaces redundancy, and LNS aggregated inline service interfaces.

> **NOTE**: When you run the `show interfaces redundancy` command on an MX80 router, it displays the error message, `error:the redundancy-interface-process subsystem is not running`. This is because an MX80 router does not have a redundant FPC and does not support link protection.

Options

brief | detail (Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 681 on page 21289 lists the output fields for the `show interfaces redundancy` command. Output fields are listed in the approximate order in which they appear.

Table 681: show interfaces redundancy Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the AMS interface, redundant adaptive services, link services IQ interfaces, aggregated Ethernet interfaces, or LNS aggregated inline service interfaces.	All levels
`State`	State of the redundant interface: `Not present`, `On primary`, `On secondary`, or `Waiting for primary MS PIC`.	All levels

Table 681: show interfaces redundancy Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Last Change`	Timestamp for the last change in status. This value resets after a primary Routing Engine switchover event if any of the following conditions is met: • GRES is not configured on the router. • The `rlsq` interface is configured without the `hot-standby` or `warm-standby` statements and the backup `lsq` interface was active before the switchover. • No logical interfaces are configured or all of the configured logical interfaces are down at the time of the switchover.	All levels
`Primary`	Name of the interface configured to be the primary interface.	All levels
`Secondary`	Name of the interface configured to be the backup interface.	All levels
`Current Status`	Physical status of the primary and secondary interfaces.	All levels
`Mode`	Standby mode.	`detail`

Sample Output

show interfaces redundancy

```
user@host> show interfaces redundancy
Interface  State         Last change  Primary    Secondary  Current status
rsp0       Not present                sp-1/0/0   sp-0/2/0   both down
rsp1       On secondary  1d 23:56     sp-1/2/0   sp-0/3/0   primary down
rsp2       On primary    10:10:27     sp-1/3/0   sp-0/2/0   secondary down
rlsq0      On primary    00:06:24     lsq-0/3/0  lsq-1/0/0  both up
ams0       On primary    00:39:51     mams-5/0/0 mams-5/1/0 both up
```

show interfaces redundancy (Aggregated Ethernet)

```
user@host> show interfaces redundancy
Interface  State            Last change  Primary      Secondary    Current status
rlsq0      On secondary     00:56:12     lsq-4/0/0    lsq-3/0/0    both up
ae0
ae1
ae2
ae3
ae4
```

show interfaces redundancy (Aggregated Inline Service Interface)

```
user@host> show interfaces redundancy asi0
Interface  State          Last change  Primary      Secondary    Current status
asi0       On primary     00:00:09     si-1/0/0     si-0/0/0     both up
```

show interfaces redundancy detail

```
user@host> show interfaces redundancy detail
Interface        : rlsq0
  State          : On primary
  Last change    : 00:45:47
  Primary        : lsq-0/2/0
  Secondary      : lsq-1/2/0
  Current status : both up
  Mode           : hot-standby

Interface        : rlsq0:0
  State          : On primary
  Last change    : 00:45:46
  Primary        : lsq-0/2/0:0
  Secondary      : lsq-1/2/0:0
  Current status : both up
  Mode           : warm-standby

Interface        : asi0
  State          : On primary
  Last change    : 00:03:42
```

```
  Primary        : si-1/0/0
  Secondary      : si-0/0/0
  Mode           : hot-standby
  Current status : both up

Interface        :ams0
  State : On primary
  Last change : 00:39:52
  Primary : mams-5/0/0
  Secondary : mams-5/1/0
  Mode : warm-standby
  Current status : both up
  Replication state : Disconnected
```

Release Information

Command introduced before Junos OS Release 7.4.

`detail` option added in Junos OS Release 10.0.

show interfaces (Redundant Adaptive Services)

IN THIS SECTION

- Syntax | 21293
- Description | 21293
- Options | 21293
- Required Privilege Level | 21293
- Output Fields | 21293
- Sample Output | 21294
- Release Information | 21295

Syntax

```
show interfaces rsp number
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M Series and T Series routers only) Display status information about the specified redundant adaptive services configuration.

Options

rsp*number*	Display standard status information about the specified redundant adaptive services configuration.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

See the output field table for the "show interfaces (Adaptive Services)" on page 20401 command.

Sample Output

show interfaces extensive (Redundant Adaptive Services)

```
user@host> show interfaces rsp0 extensive
Physical interface: rsp0, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 40, Generation: 44
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Redundancy-Device 16384
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2005-03-11 18:36:37 UTC (00:00:08 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface rsp0.0 (Index 68) (SNMP ifIndex 42) (Generation 30)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
```

```
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 9192, Generation: 37, Route table: 0
      Flags: Receive-options, Receive-TTL-Exceeded
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces (Redundant Link Services IQ)

IN THIS SECTION

- Syntax | 21295
- Description | 21296
- Options | 21296
- Required Privilege Level | 21296
- Output Fields | 21296
- Sample Output | 21312
- Release Information | 21315

Syntax

```
show interfaces rlsqnumber
<brief | detail | extensive | terse>
<descriptions>
<media>
<queue>
<routing>
```

```
<snmp-index snmp-index>
<statistics>
```

Description

(M Series and T Series routers only) Display status information about the specified redundant link services intelligent queuing (IQ) configuration.

Options

rlsq*number*	Redundant link services IQ interface name. The logical interface number range of values is **0** through **127**.
none	Display standard status information about the specified redundant link services IQ configuration.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
queue	(Optional) Display queue information about network interfaces.
routing	(Optional) Display routing information about network interfaces.
snmp-index ***snmp-index***	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 682 on page 21297 lists the output fields for the `show interfaces` (redundant link services IQ) command. Output fields are listed in the approximate order in which they appear.

Table 682: show interfaces (Redundant Link Services IQ) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface: **Multilink-Frame-Relay-UNI-NNI** (default), **LinkService**, **Frame-relay**, **Frame-relay-ccc**, or **Frame-relay-tcc**.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour.minute.second timezone*** **(*hour.minute.second* ago)** . For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Input rate	(Redundant LSQ) Rate of bits and packets received on the interface.	None specified

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output rate	(Redundant LSQ) Rate of bits and packets transmitted on the interface.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.	**detail extensive**
Frame exceptions	Information about framing exceptions. Includes events recorded under **Exception Events** for each logical interface. • **Oversized frames**—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits). • **Errored input frames**—Number of input frame errors. • **Input on disabled link/bundle**—Number of frames received on disabled links. These frames can result either from an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it. • **Output for disabled link/bundle**—Number of frames sent for a disabled or unavailable link. These frames can result either from an inconsistent configuration, or from a bundle being brought up or down while traffic is flowing through it. • **Queuing drops**—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed.	**extensive**

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Buffering exceptions	Information about buffering exceptions. Includes events recorded under **Exception Events** for each logical interface: • **Packet data buffer overflow**—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. Check the logical interface exception event counters to determine which bundle is responsible.	**extensive**

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Assembly exceptions	(Multilink Frame Relay end-to-end only) Information about assembly exceptions. Includes events recorded under **Exception Events** for each logical interface. An assembly exception does not necessarily indicate an operational problem with the physical link services IQ interface itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the interface records one or more exception events. The physical interface can drop multilink-encapsulated fragments itself as a result. Any multilink packets or fragments dropped by the interface itself result in packet or fragment drop counts on individual logical interfaces. If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, excess differential delay or traffic losses on individual interfaces can be the root cause. • **Fragment timeout**—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • **Missing sequence number**—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the link services IQ interface unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible.	**extensive**

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **Out-of-range sequence number**—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible.	
Hardware errors (sticky)	(Multilink Frame Relay end-to-end only) Information about hardware errors: • **Data memory error**—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. • **Control memory error**—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive** none
Queue counters	Queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive** none
Logical Interface		
Logical interface	Name of the logical interface	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
Encapsulation	Encapsulation being used: PPP or Multilink PPP.	All levels
Bandwidth	Speed at which the interface is running.	All levels
Bundle options	(Multilink Frame Relay end-to-end interfaces only) • **MRRU**—Configured size of the maximum received reconstructed unit (MRRU): **1500** through **4500** bytes. The default is **1504** bytes. • **Drop timer period**—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: **0** though **2000** milliseconds. Values under 5 ms are not recommended. The default setting is **0**, which disables the timer. • **Sequence number format**—Short sequence number header format (MLPPP only). • **Fragmentation threshold**—Configured fragmentation threshold: **64** through **16,320** bytes, in integer multiples of **64** bytes. The default setting is **0**, which disables fragmentation. • **Links needed to sustain bundle**—Minimum number of links to sustain the bundle: **1** through **8**. • **Multilink classes**—Number of multilink classes negotiated. • **Link layer overhead**—Percentage of bundle bandwidth to be set aside for link-layer overhead.	**detail extensive** none

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bundle status (MLPPP) or **Multilink class status** (MC-MLPPP)	Information about bundle status: • **Remote MRRU**—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. • **Received sequence number**—Sequence number for received packets. • **Transmitted sequence number**—Sequence number for transmitted packets. • **Packet drops**—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • **Fragment drops**—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • **MRRU exceeded**—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • **Fragment timeout**—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • **Missing sequence number**—A gap was detected in the sequence numbers of fragments on a bundle. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers within a single link. This event indicates that an	**detail extensive** none

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • **Out-of-range sequence number**—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • **Packet data buffer overflow**—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity.	

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Statistics	Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).
	• **Bundle**—Information for each active bundle link.
	  • **Fragments: Input and Output**—Total number and rate of fragments received and transmitted.
	  • **Packets: Input and Output**—Total number and rate of packets received and transmitted.
	  • **Multilink class**—(MC-MLPPP only) Information about multiclass links used in the multilink operation.
	• **Link**—Information about links used in the multilink operation.
	  • **Link name**—Interface name of the link services IQ channel and state information (physical link **up** or **down**).
	  • **Input and Output**—Total number and rate of fragments and packets received and transmitted.
		detail extensive

NCP state	(PPP) Network Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Protocol	Protocol family configured on the logical interface.	**detail extensive** none
MTU	MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked **Adjusted**.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive** none
Addresses, Flags	Information about the addresses configured on the logical interface. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
MLPPP Bundle Interface		
Logical interface	Name of the logical interface.	All levels

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
SNMP-Traps	SNMP trap notifications are enabled.	All levels
Encapsulation	Encapsulation being used: PPP, Multilink PPP or Multilink-FR.	All levels
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped: *year-month-day hour.minute.second timezone* (*hour.minute.second* ago)** . For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Bandwidth	Speed at which the interface is running.	All levels
Bundle links linformation	Information about the bundled links. • **Active bundle links**—Number of active links. • **Removed bundle links**—Information about links used in the multilink operation. • **Disabled bundle links**—Number of disabled links.	**detail extensive** none

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bundle options	(Multilink Frame Relay end-to-end interfaces only) • **MRRU**—Configured size of the maximum received reconstructed unit (MRRU): **1500** through **4500** bytes. The default is **1504** bytes. • **Drop timer period**—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: **0** though **2000** milliseconds. Values under 5 ms are not recommended. The default setting is **0**, which disables the timer. • **Inner PPP Protocol field compression**—Inner PPP protocol compression is enabled or disabled. • **Sequence number format**—Short sequence number header format (MLPPP only). • **Fragmentation threshold**—Configured fragmentation threshold: **64** through **16,320** bytes, in integer multiples of **64** bytes. The default setting is **0**, which disables fragmentation. • **Links needed to sustain bundle**—Minimum number of links to sustain the bundle: **1** through **8**. • **Multilink classes**—Number of multilink classes negotiated. • **Link layer overhead**—Percentage of bundle bandwidth to be set aside for link-layer overhead.	**detail extensive** none

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bundle status (MLPPP)	Information about bundle status: • **Received sequence number**—Sequence number for received packets. • **Transmit sequence number**—Sequence number for transmitted packets. • **Packet drops**—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • **Fragment drops**—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • **MRRU exceeded**—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • **Fragment timeout**—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • **Missing sequence number**—A gap was detected in the sequence numbers of fragments on a bundle. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers occurred within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream.	**detail extensive** none

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **Out-of-range sequence number**—A frame was received with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • **Packet data buffer overflow**—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity.	

Statistics	Information about frames, bytes, and bits per second received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).
	• **Bundle**—Information for each active bundle link.
	  • **Multilink: Input and Ouput**—Total number and rate of multilink frames, bytes, and bits per second received and transmitted.
	  • **Network: Input and Output**—Total number of multilink frames, bytes, and bits per second received and transmitted.
	• **Link**—Information about links used in the multilink operation.
	  • **Link name** is the interface name of the link services IQ channel and state information (physical link **up** or **down**) and up time.
	  • **Input and Output**—Total number and rate of frames, bytes, and bits per second received and transmitted.
		extensive

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Multilink detail statistics	Frames, bytes, and bits per second received and sent by the bundle. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).
	• **Bundle**—Information for the bundle link.
	  • **Fragments: Input and Output**—Total number and rate of multilink fragments received and transmitted.
	  • **Non-fragments: Input and Output**—Total number and rate of nonfragmented multilink frames received and transmitted.
	  • **LFI: Input and Output**—Total number and rate of link fragmented and interleaved frames and bytes.
		extensive

Protocol	Protocol family configured on the logical interface.	**detail extensive** none
MTU	MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked **Adjusted**.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**
Addresses, Flags	Information about the addresses configured on the logical interface. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none

Table 682: show interfaces (Redundant Link Services IQ) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**

Sample Output

show interfaces (Redundant Link Services IQ)

```
user@host> show interfaces rlsq0
Physical interface: rlsq0, Enabled, Physical link is Up
  Interface index: 196, SNMP ifIndex: 27
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface rlsq0.0 (Index 72) (SNMP ifIndex 88)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
    Bandwidth: 0
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Fragments:
        Input :             3          0           255            0
        Output:             3          0           264            0
      Packets:
        Input :             3          0           252            0
        Output:             0          0             0            0
    Link:
      t1-1/3/0:1.0
```

```
        Input :             3          0           255            0
        Output:             0          0             0            0
      t1-1/3/0:2.0
        Input :             0          0             0            0
        Output:             3          0           264            0
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2.0/30, Local: 192.0.2.1
```

show interfaces brief (Redundant Link Services IQ)

```
user@host> show interfaces rlsq0 brief
Physical interface: rlsq0, Enabled, Physical link is Up
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000

  Logical interface rlsq0.0
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
    inet  192.0.2.1/30
```

show interfaces detail (Redundant Link Services IQ)

```
user@host> show interfaces rlsq0 detail
Physical interface: rlsq0, Enabled, Physical link is Up
  Interface index: 196, SNMP ifIndex: 27, Generation: 144
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  252                    0 bps
   Output bytes  :                  276                    0 bps
   Input  packets:                    3                    0 pps
   Output packets:                    3                    0 pps
  Frame exceptions:
```

```
    Oversized frames                  0
    Errored input frames              0
    Input on disabled link/bundle     0
    Output for disabled link/bundle   0
    Queuing drops                     0
  Buffering exceptions:
    Packet data buffer overflow       0
    Fragment data buffer overflow     0
  Assembly exceptions:
    Fragment timeout                  0
    Missing sequence number           0
    Out-of-order sequence number      0
    Out-of-range sequence number      0
  Hardware errors (sticky):
    Data memory error                 0
    Control memory error              0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 be                             0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0

  Logical interface rlsq0.0 (Index 72) (SNMP ifIndex 88) (Generation 31)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
    Bandwidth: 0
    Bundle options:
      MRRU                            1504
      Remote MRRU                     N/A
      Drop timer period               2000
      Sequence number format          long (24 bits)
      Fragmentation threshold         0
      Links needed to sustain bundle  1
      Multilink classes               0
      Link layer overhead             4.0 %
    Bundle status:
      Received sequence number        0xffffff
      Transmit sequence number        0x0
      Packet drops                    0 (0 bytes)
      Fragment drops                  0 (0 bytes)
      MRRU exceeded                   0
      Fragment timeout                0
      Missing sequence number         0
```

```
      Out-of-order sequence number    0
      Out-of-range sequence number    0
      Packet data buffer overflow     0
      Fragment data buffer overflow   0
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Fragments:
        Input :             3          0           255            0
        Output:             3          0           264            0
      Packets:
        Input :             3          0           252            0
        Output:             0          0             0            0
    Link:
      t1-1/3/0:1.0
        Input :             3          0           255            0
        Output:             0          0             0            0
      t1-1/3/0:2.0
        Input :             0          0             0            0
        Output:             3          0           264            0
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
    Protocol inet, MTU: 1500, Generation: 43, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2.0/30, Local: 2.2.2.1, Broadcast: Unspecified,
        Generation: 45
```

show interfaces extensive (Redundant Link Services IQ)

The output for the `show interfaces rlsq extensive` command is identical to that for the `show interfaces rlsq detail` command. For sample output, see "show interfaces detail (Redundant Link Services IQ)" on page 21313.

Release Information

Command introduced in Junos OS Release 7.6.

show interfaces routing

IN THIS SECTION

- Syntax | 21316
- Description | 21316
- Options | 21316
- Additional Information | 21317
- Required Privilege Level | 21317
- Output Fields | 21317
- Sample Output | 21319
- Release Information | 21325

Syntax

```
show interfaces routing
<brief | summary | detail>
<interface-name>
<logical-system (all | logical-system-name)>
```

Description

Display the state of the router's interfaces. Use this command for performing router diagnostics only, when you are determining whether the routing protocols and the Junos OS differ about the state of an interface.

Options

none	Display standard information about the state of all router interfaces on all logical systems.
brief | summary | detail	(Optional) Display the specified level of output.
interface-name	(Optional) Name of a specific interface.

logical-system (all | *logical-system-name*) (Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

For information about related operational mode commands for routing instances and protocols, see the CLI Explorer.

Required Privilege Level

view

Output Fields

Table 683 on page 21317 lists the output fields for the `show interfaces routing` command. Output fields are listed in the approximate order in which they appear.

Table 683: show interfaces routing Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the physical interface.	none `brief`
`State`	State of the physical interface: `Up` or `Down`.	none `brief`
`Addresses`	Protocols and addresses configured on the interface.	none `brief`
`Index`	Interface index number, which reflects its initialization sequence.	`detail`
`Refcount`	Number of references to the interface in the routing software.	`detail`
`State`	State (`Up` or `Down`) and type of interface.	`detail`

Table 683: show interfaces routing Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Change`	Reflects one or more of the following recent changes to the interface: • `Add`—The interface was just added. • `Address`—The interface's link-layer address has changed. • `Delete`—The interface is being deleted. • `Encapsulation`—The type of encapsulation on the interface has changed. • `Metric`—The interface's metric value has changed. • `MTU`—The interface's maximim transmission unit size has changed. • `UpDown`—The interface has made an up or down transition.	`detail`
`Up/down transitions`	Number of times the interface has gone from `Down` to `Up`.	`detail`
`Link layer`	Describes the link layer of the interface.	`detail`
`Encapsulation`	Encapsulation on the interface.	`detail`
`Bandwidth`	Speed at which the interface is running.	`detail`

Table 683: show interfaces routing Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Protocol address	Information about the configuration of protocols on the interface:
	• Address—Address configured on the interface for the protocol type.
	• State—State (Up or down) and type of interface.
	• Change—Reflects one or more of the following recent changes to the interface:
	  • Add—The interface was just added.
	  • Address—The interface's address has changed.
	  • Broadcast—The interface's broadcast address has changed.
	  • Delete—The interface is being deleted.
	  • Netmask—The interface's netmask has changed.
	  • UpDown—The interface has made an up or down transition.
	• Preference—Preference value for the route for this address.
	• Metric—Metric value on the interface for the protocol type.
	• MTU—Maximim transmission unit value of the interface.
	• Local address—On a point-to-point link, the address of the local side of the link. Not used for multicast links.
	• Destination—For a point-to-point link, the address of the remote side of the link. For multicast links, the network address.
		detail

Sample Output

show interfaces routing brief

```
user@host> show interfaces routing brief
Interface        State Addresses
so-5/0/3.0       Down  ISO   enabled
so-5/0/2.0       Up    MPLS  enabled
```

```
                        ISO   enabled
                        INET  192.168.2.120
                        INET  enabled
so-5/0/1.0       Up     MPLS  enabled
                        ISO   enabled
                        INET  192.168.2.130
                        INET  enabled
at-1/0/0.3       Up     CCC   enabled
at-1/0/0.2       Up     CCC   enabled
at-1/0/0.0       Up     ISO   enabled
                        INET  192.168.90.10
                        INET  enabled
lo0.0            Up     ISO   47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
                        ISO   enabled
                        INET  127.0.0.1
fxp1.0           Up
fxp0.0           Up     INET  192.168.6.90
```

show interfaces routing brief (TX Matrix Plus Router)

```
user@host> show interfaces routing brief
Interface        State Addresses
...
ge-23/0/4.0      Up     INET  203.0.113.91
                        ISO   enabled
                        MPLS  enabled
ge-23/0/3.0      Up     INET  203.0.113.81
                        ISO   enabled
                        MPLS  enabled
ge-23/0/2.0      Up     INET  203.0.113.71
                        ISO   enabled
                        MPLS  enabled
ge-23/0/1.0      Up     INET  203.0.113.61
                        ISO   enabled
                        MPLS  enabled
ge-23/0/0.0      Up     INET  203.0.113.51
                        ISO   enabled
                        MPLS  enabled
ge-31/0/7.599    Up     INET  192.0.2.93
ge-31/0/7.598    Up     INET  192.0.2.89
ge-31/0/7.597    Up     INET  192.0.2.85
```

```
ge-31/0/7.596    Up    INET  192.0.2.81
ge-31/0/7.595    Up    INET  192.0.2.77
ge-31/0/7.594    Up    INET  192.0.2.73
...
ixgbe1.0         Up    INET  203.0.113.34
                       INET  198.51.100.4
                       INET6 fe80::200:1ff:fe22:4
                       INET6 fec0::a:22:0:4
ixgbe0.0         Up    INET  203.0.113.34
                       INET  198.51.100
                       INET6 fe80::200:ff:fe22:4
                       INET6 fec0::a:22:0:4
em0.0            Up    INET  192.168.178.11
```

show interfaces routing detail

```
user@host> show interfaces routing detail
so-5/0/3.0
  Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>
  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  ISO address (null)
    State: <Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>
  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  MPLS address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
  ISO address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  INET address 192.168.2.120
    State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
    Local address: 192.168.2.120
    Destination: 192.168.2.110/32
  INET address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
```

```
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...
```

show interfaces routing detail (TX Matrix Plus Router)

```
user@host> show interfaces routing detail
ge-23/0/4.0
  Index: 77, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 0.1d.b5.14.da.2d
  INET address 203.0.113.91
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 203.0.113.93
    Destination: 203.0.113.0/30
    System flags: <Is-Preferred Is-Primary>
  ISO address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1497 bytes
    System flags: <>
  MPLS address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1488 bytes
    System flags: <>
ge-23/0/3.0
  Index: 76, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 0.1d.b5.14.da.2c
  INET address 203.0.113.81
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 10.8.1.3
    Destination: 203.0.113.80/30
    System flags: <Is-Preferred Is-Primary>
  ISO address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1497 bytes
    System flags: <>
  MPLS address (null)
```

```
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1488 bytes
    System flags: <>
ge-23/0/2.0
  Index: 75, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 0.1d.b5.14.da.2b
  INET address 203.0.113.71
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 203.0.113.73
    Destination: 203.0.113.70/30
    System flags: <Is-Preferred Is-Primary>
  ISO address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1497 bytes
    System flags: <>
  MPLS address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1488 bytes
    System flags: <>
ge-23/0/1.0
  Index: 74, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 0.1d.b5.14.da.2a
  INET address 203.0.113.61
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 203.0.113.63
...
ixgbe1.0
  Index: 5, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 2.0.1.22.0.4
  INET address 203.0.113.34
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 203.0.113.255
    Destination: 203.0.113.0/8
    System flags: <Is-Preferred>
```

```
  INET address 198.51.100.4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 191.255.255.255
    Destination: 192.0.2.0/2
    System flags: <Primary Is-Preferred Is-Primary>
  INET6 address fe80::200:1ff:fe22:4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Destination: fe80::/64
    System flags: <Is-Preferred>
  INET6 address fec0::a:22:0:4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Destination: fec0::/64
    System flags: <Is-Preferred Is-Primary>
ixgbe0.0
  Index: 4, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 2.0.0.22.0.4
  INET address 203.0.113.34
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 203.0.113.255
    Destination: 203.0.113.0/8
    System flags: <Is-Preferred>
  INET address 198.51.100.4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 191.255.255.255
    Destination: 192.0.2.0/2
    System flags: <Primary Is-Default Is-Preferred Is-Primary>
  INET6 address fe80::200:ff:fe22:4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Destination: fe80::/64
    System flags: <Is-Preferred>
  INET6 address fec0::a:22:0:4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Destination: fec0::/64
    System flags: <Is-Default Is-Preferred Is-Primary>
```

```
em0.0
  Index: 3, Refcount: 2, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 100Mbps
  Link address #0 0.80.f9.26.0.c0
  INET address 192.168.178.11
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 192.168.178.127
    Destination: 192.168.178.0/25
    System flags: <Is-Preferred Is-Primary>
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces routing-instance

IN THIS SECTION

- Syntax | 21325
- Description | 21326
- Options | 21326
- Required Privilege Level | 21326
- Output Fields | 21326
- Sample Output | 21326
- Sample Output | 21327
- Release Information | 21328

Syntax

```
show interfaces routing-instance (instance-name | all)
<brief | detail | extensive | terse>
```

Description

Display information about the interfaces configured for either a specific routing instance or for all of the routing instances.

Options

all	Display information about all of the interfaces configured for all of the routing instances on the router.
instance-name	Display information about the interfaces configured for the specified routing instance.
brief | detail | extensive | terse	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

The output fields from the `show interfaces routing-instance` command are identical to those produced by the `show interfaces` *`interface-name`* command. For a description of output fields, see the other chapters in this manual.

Sample Output

show interfaces routing-instance terse

```
user@host> show interfaces routing-instance sample terse
Interface      Admin   Link      Proto    Local                Remote
ge-0/0/0.0     up      up        inet     192.168.4.28/24
```

Sample Output

show interfaces routing-instance all

```
user@host> show interfaces terse routing-instance all
Interface   Admin   Link   Proto    Local             Remote Instance
at-0/0/1       up      up     inet        203.0.113.1/24
ge-0/0/0.0  up      up     inet      192.168.4.28/24          sample-a
at-0/1/0.0  up      up     inet6     fe80::a:0:0:4/64         sample-b
so-0/0/0.0     up      up     inet        203.0.113.1/32
```

show interfaces routing-instance extensive

```
user@hostshow interfaces fe-0/1/3 routing-instance instance2 extensive
Logical interface fe-0/1/3.0 (Index 70) (SNMP ifIndex 53) (Generation 211)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes :                   0
     Output bytes :                  42
     Input packets:                   0
     Output packets:                  1
     IPv6 transit statistics:
      Input bytes :                   0
      Output bytes  :                 0
      Input packets:                  0
      Output packets:                 0
    Local statistics:
     Input bytes :                   0
     Output bytes :                 42
     Input packets:                  0
     Output packets:                 1
    Transit statistics:
     Input bytes :                   0                  0 bps
     Output bytes :                  0                  0 bps
     Input packets:                  0                  0 pps
     Output packets:                 0                  0 pps
     IPv6 transit statistics:
      Input bytes :                  0
      Output bytes :                 0
      Input packets:                 0
```

```
    Output packets:                0
  Protocol inet, MTU: 1500, Generation: 252, Route table: 4
    Flags: Is-Primary
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
    Destination: 192.0.2/24, Local: 192.0.2.51, Broadcast: 192.0.2.255, Generation: 263
```

Release Information

Command introduced in Junos OS Release 9.1.

show interfaces routing summary

IN THIS SECTION

- Syntax | 21328
- Description | 21329
- Options | 21329
- Additional Information | 21329
- Required Privilege Level | 21329
- Output Fields | 21329
- Sample Output | 21330
- Release Information | 21332

Syntax

```
show interfaces routing summary
<interface-name>
<logical-system (all | logical-system-name)>
```

Description

Display a summary of the state of the router interfaces. Use this command for performing router diagnostics only, when you are determining whether the routing protocols and the Junos OS differ about the state of an interface.

Options

none	Display summary information about the state of all router interfaces on all logical systems.
interface-name	(Optional) Name of a specific interface.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

For information about how to configure routing protocols, see the Junos OS Routing Protocols Library for Routing Devices. For information about related operational mode commands for routing instances and protocols, see the CLI Explorer.

Required Privilege Level

view

Output Fields

Table 684 on page 21329 lists the output fields for the `show interfaces routing summary` command. Output fields are listed in the approximate order in which they appear.

Table 684: show interfaces routing summary Output Fields

Field Name	Field Description
***n* physical interfaces**	Number of routing interfaces and number of interfaces in the **up** state.
***n protocol* protocol interfaces**	Type and number of routing protocols and number of related interfaces in the **up** state.

Table 684: show interfaces routing summary Output Fields *(Continued)*

Field Name	Field Description
Interface	Logical interface name.
Index	Logical interface index number, which reflects its initialization sequence.
Metric	Metric value for the interface.
Trans	Number of times the interface has transitioned from **Down** to **Up**.
Status	Interface status (**Up** or **Down**) and type.

Sample Output

show interfaces routing summary

```
user@host> show interfaces routing summary
14 physical interfaces (12 up)
  11 INET protocol addresses (11 up)
  6 ISO protocol addresses (4 up)
  3 MPLS protocol addresses (3 up)
  3 CCC protocol addresses (3 up)
Interface  Index   Metric   Trans. Status
so-5/0/3.0    15        0        0 Broadcast PointToPoint Multicast
so-5/0/2.0    14        0        0 Up Broadcast PointToPoint Multicast
so-5/0/1.0    13        0        5 Up Broadcast PointToPoint Multicast
so-5/0/0.0    12        0        2 Up Broadcast PointToPoint Multicast
so-1/2/0.0    11        0        0 Broadcast PointToPoint Multicast
so-1/1/0.0    10        0        5 Up Broadcast PointToPoint Multicast
at-1/0/0.6     9        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.5     8        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.4     7        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.3     6        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.2     5        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.0     4        0        0 Up Broadcast PointToPoint Multicast
lo0.0          3        0        0 Up Broadcast Loopback Multicast
```

```
fxp1.0         2        0        1 Up Broadcast Multicast
fxp0.0         1        0        0 Up Broadcast Multicast
```

show interfaces routing summary (TX Matrix Plus Router)

```
user@host> show interfaces routing summary
9 physical interfaces (9 up)
  11 INET protocol addresses (11 up)
  6 MPLS protocol addresses (6 up)
  4 INET6 protocol addresses (4 up)

Interface             Index   Metric   Trans. Status
ge-23/0/8.0              73        0        0 Up Broadcast Multicast
ge-23/0/7.0              72        0        0 Up Broadcast Multicast
ge-23/0/6.0              71        0        0 Up Broadcast Multicast
ge-7/0/9.0               69        0        0 Up Broadcast Multicast
ge-15/0/9.0              70        0        0 Up Broadcast Multicast
xe-6/1/1.0               68        0        0 Up Broadcast Multicast
lo0.16385                66        0        0 Up Broadcast Loopback Multicast
lo0.16384                65        0        0 Up Broadcast Loopback Multicast
lo0.0                    64        0        0 Up Broadcast Loopback Multicast
ixgbe1.0                  5        0        0 Up Broadcast Multicast
ixgbe0.0                  4        0        0 Up Broadcast Multicast
em0.0                     3        0        0 Up Broadcast Multicast
```

show interfaces routing summary (PTX5000 Packet Transport Routers)

```
user@host> show interfaces routing summary
7 physical interfaces (68 up)
  7 INET protocol addresses (7 up)
  2 CCC protocol addresses (2 up)
  4 INET6 protocol addresses (4 up)

Interface             Index   Metric   Trans. Status
lo0.16385                66        0        0 Up Broadcast Loopback Multicast
lo0.16384                64        0        0 Up Broadcast Loopback Multicast
lo0.0                    65        0        0 Up Broadcast Loopback Multicast
ixgbe1.0                  5        0        0 Up Broadcast Multicast
ixgbe0.0                  4        0        0 Up Broadcast Multicast
et-5/0/5.32767           72        0        0 Up Broadcast Multicast
```

```
et-5/0/5.0                68        0        0 Up Broadcast Multicast
et-5/0/0.32767            67        0        0 Up Broadcast Multicast
et-5/0/0.0                71        0        0 Up Broadcast Multicast
em0.0                      3        0        0 Up Broadcast Multicast
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces satellite-device

IN THIS SECTION

- Syntax | 21332
- Description | 21332
- Options | 21333
- Required Privilege Level | 21333
- Output Fields | 21333
- Sample Output | 21337
- Sample Output | 21337
- Release Information | 21339

Syntax

```
show interfaces satellite-device  (device-alias | all)
```

Description

Display the satellite device extended ports in a Junos Fusion.

Options

device-alias *device-alias* — Display extended port information for the satellite device using the specified device alias only.

all — Display information for all extended ports and aggregated Ethernet interfaces with extended ports as members configured on all of the satellite devices.

Required Privilege Level

view

Output Fields

Table 685 on page 21333 lists the output fields for the `show interfaces satellite-device` command. Output fields are listed in the approximate order in which they appear.

Table 685: show interfaces satellite-device Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
`Link-level type`	Encapsulation being used on the physical interface.	All levels
Device flags	Information about the physical device.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**. **NOTE**: This field is only displayed if asymmetric flow control is not configured.	All levels

Table 685: show interfaces satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Pad to minimum frame size`	Pad Tx VLAN-tagged frame to minimum of 68 bytes.	All levels
`Minimum links needed`	Minimum number of aggregated links.	All levels
`Minimum bandwidth needed`	Minimum bandwidth configured for aggregated bundle.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Current address	Configured MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour: :minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).**	**detail extensive** none
`Input rate`	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	All levels
`Output rate`	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	All levels
`Extended port information`	Satellite device port ID	All levels

Table 685: show interfaces satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
`Interface transmit statistics`	All levels	All levels
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface.	All levels

Table 685: show interfaces satellite-device Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Statistics	• Packets • pps • Bytes • bps	All levels

Bundle	Provides information for each active bundle link.
	• Input
	  • Packets
	  • pps
	  • Bytes
	  • bps
	  • Output
	  • Packets—
	  • pps
	  • Bytes
	  • bps
		All levels

Adaptive Statistics	• Adaptive Adjusts • Adaptive Scans • Adaptive Updates	All levels
Protocol	Protocol family configured on the logical interface.	All levels

Sample Output

Sample Output

show interfaces satellite-device all

```
user@aggregation-device> show interfaces satellite-device all
Physical interface: ae0 (Extended Port, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 574
  Link-level type: Ethernet, MTU: 1514, Speed: 2Gbps, BPDU Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: f4:b5:2f:f7:3f:c0, Hardware address: f4:b5:2f:f7:3f:c0
  Last flapped   : 2015-03-31 18:36:43 PDT (06:54:08 ago)
  Input rate     : 2032 bps (2 pps)
  Output rate    : 3048 bps (2 pps)

  Logical interface ae0.0 (Index 337) (SNMP ifIndex 575)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          1704          0        592992            0
        Output:             0          0             0            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Protocol bridge, MTU: 1514

Physical interface: ae1 (Extended Port, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 790
  Link-level type: Ethernet, MTU: 1514, Speed: 200mbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
```

```
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: f4:b5:2f:f7:3f:c1, Hardware address: f4:b5:2f:f7:3f:c1
  Last flapped   : 2015-03-31 18:36:44 PDT (06:54:07 ago)
  Input rate     : 2032 bps (2 pps)
  Output rate    : 2032 bps (2 pps)

  Logical interface ae1.0 (Index 338) (SNMP ifIndex 1216)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          2759          0        679982            0
        Output:             0          0             0            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Protocol bridge, MTU: 1514
Physical interface: xe-101/0/31 (Extended Port, Enabled, Physical link is Up
  Interface index: 336, SNMP ifIndex: 829
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 10:0e:7e:bf:2d:24, Hardware address: 10:0e:7e:bf:2d:24
  Last flapped   : 2015-03-31 08:28:23 PDT (17:02:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Interface transmit statistics: Disabled

  Logical interface xe-101/0/31.0 (Index 491) (SNMP ifIndex 926)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Input packets : 0
    Output packets: 0
    Protocol bridge, MTU: 1514

Physical interface: xe-101/0/32 (Extended Port, Enabled, Physical link is Up
  Interface index: 337, SNMP ifIndex: 836
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps
  Device flags   : Present Running
```

```
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Current address: 10:0e:7e:bf:2d:25, Hardware address: 10:0e:7e:bf:2d:25
Last flapped   : 2015-03-31 08:28:23 PDT (17:02:29 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Active alarms  : None
Active defects : None
Interface transmit statistics: Disabled

Logical interface xe-101/0/32.0 (Index 492) (SNMP ifIndex 935)
  Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
  Input packets : 0
  Output packets: 0
  Protocol bridge, MTU: 1514
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

show interfaces (Serial)

IN THIS SECTION

- Syntax | 21340
- Description | 21340
- Options | 21340
- Required Privilege Level | 21340
- Output Fields | 21341
- Sample Output | 21350

Release Information | 21359

Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about serial interfaces, including RS-232, RS-422/449, EIA-530, X.21, and V.35.

Options

interface-type	On M Series and T Series routers, the interface type is **se-*fpc/pic/port***.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 686 on page 21341 lists the output fields for the `show interfaces` (Serial) command. Output fields are listed in the approximate order in which they appear.

Table 686: show interfaces (Serial) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit (MTU) size on the physical interface.	All levels
Maximum speed	Maximum speed. The nonconfigurable value is 16,384 kbps.	**detail extensive** none
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels

Table 686: show interfaces (Serial) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels
Link flags	Information about the link. Possible values are described in the "Link Flags" section under *Common Output Fields Description*.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Keepalive settings	(PPP and HDLC) Configured settings for keepalive packets. • **Interval *seconds***—Time between successive keepalive requests. The range of values, in seconds, is 10 to 32,767. The default value is 10. • **Up-count *number***—Number of keepalive packets a destination must receive to change a link's status from down to up. The range of values is 1 to 255. The default value is 1. • **Down-count *number***—Number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 to 255. The default value is 3.	All levels
Keepalive	(PPP and HDLC) Information about keepalive packets. • **Input: *number* (*hh:mm:ss* ago)**—Number of keepalive packets received by PPP and the time since the last keepalive packet was received. • **Output: *number* (*hh:mm:ss* ago)**—Number of keepalive packets sent by PPP and the time since the last keepalive packet was sent.	**brief** none
Keepalive statistics	(PPP and HDLC) Information about keepalive packets. • **Input: *number* (last seen *hh:mm:ss*ago)**—Number of keepalive packets received by PPP and the time since the last keepalive packet was received. • **Output: *number*(last seen *hh:mm:ss* ago)**—Number of keepalive packets sent by PPP and the time since the last keepalive packet was sent.	**detail extensive**

Table 686: show interfaces (Serial) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
LCP state	(PPP) Link Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—LCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—LCP is not configured on the interface. • **Opened**—LCP negotiation is successful.	**detail extensive** none
NCP state	(PPP) Network Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none

Table 686: show interfaces (Serial) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • **Chap-Chal-received**—Challenge was received but response not yet sent. • **Chap-Chal-sent**—Challenge was sent. • **Chap-Resp-received**—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • **Chap-Resp-sent**—Response was sent for the challenge received. • **Closed**—CHAP authentication is incomplete. • **Failure**—CHAP authentication failed. • **Not-configured**—CHAP is not configured on the interface. • **Success**—CHAP authentication was successful.	**detail extensive** none
CoS queues	Number of CoS queues configured.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 686: show interfaces (Serial) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame terminated and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 686: show interfaces (Serial) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • **Errors**—Sum of the outgoing frame terminated and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets whose size exceeds the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues supported	Total number of egress queues supported on the specified interface. Displayed with the **statistics** option.	**detail extensive**
Egress queues in use	Total number of egress queues in use on the specified interface. Displayed with the **statistics** option.	**detail extensive**
Queue counters	CoS queue number and its associated user-configured forwarding class name. Displayed with the **statistics** option. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**

Table 686: show interfaces (Serial) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Serial media information	Information about the physical media:
	• **Line protocol**—**eia530**, **eia530a**, **rs232**, **rs449**, **v.35**, or **x.21.**.
	• **Resync history**—Information about resynchronization events:
	  • **Sync loss count**—Number of times the synchronization was lost.
	• **Data signal**—(X.21 and V.35) Information about the data signal:
	  • **Rx Clock**—Receive clock status: OK (DTE is receiving the receive clock signal) or Not detected (receive clock signal is not being received).
	• **Control signals**—Information about modem control signals:
	  • **Local mode:DCE** (data communication equipment) or **DTE** (data terminal equipment)
	  • **To DCE**—Control signals that the Serial PIC sent to the DCE: DTR (Data Terminal Ready:**up** or **down**) or RTS (Request To Send: **up** or **down**.)
	  • **From DC**—Control signals that the Serial PIC received from the DCE: CTS (Clear To Send: **up** or **down**), DCD (Data Carrier Detect: **up** or **down**), **DSR** (Data Set Ready: **up** or **down**), or TM (Test Mode: **up** or **down**).
	• **Clocking mode**—Clocking used for the transmit clock:
	  • **dte**—Transmit clock is generated by DTE.
	  • **dce**—Transmit clock is generated by the DCE and is looped back as the transmit clock.
	  • **loop-timed**—Receive clock from the DCE is looped back as the transmit clock.
	• **Clock rate**—Rate, in megahertz (MHz), at which the clock is configured.
	• **Loopback**—Configured loopback mode for the interface: **dce-remote**, **dce-local**, **liu**, **local**, or **none**.
		detail extensive

Table 686: show interfaces (Serial) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **Tx clock**—Clocking phase of the transmit clock: **invert** (transmit clock polarity is inverted) or **non-invert** (transmit clock polarity is not inverted). • **Line encoding**—Type of line encoding used: **nrz** (nonreturn to zero) or **nrzi** (return to zero inverted).	
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **PLP byte**—Packet Level Protocol byte.	**extensive**
CoS information	Information about the CoS queue for the physical interface: • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels

Table 686: show interfaces (Serial) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the source and destination address are also displayed.	**brief**
Protocol	Protocol family configured on the logical interface, such as iso, inet6, mpls.	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive**
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none

Table 686: show interfaces (Serial) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces (Serial, EIA-530)

```
user@host> show interfaces se-5/0/1
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 41
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 32 (00:00:10 ago), Output: 31 (00:00:07 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2006-04-26 15:10:18 PDT (00:05:22 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
```

```
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3
```

show interfaces brief (Serial, EIA-530)

```
user@host> show interfaces se-5/0/1 brief
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Type: Serial, Link-level type: PPP, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 235 (00:00:10 ago), Output: 234 (00:00:00 ago)

  Logical interface se-5/0/1.0
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    inet  12.0.0.1/30
```

show interfaces detail (Serial, EIA-530)

```
user@host> show interfaces se-5/0/1 detail
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 41, Generation: 25
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 37 (last seen 00:00:06 ago)
    Output: 35 (last sent 00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2006-04-26 15:10:18 PDT (00:06:02 ago)
```

```
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 928                   40 bps
   Output bytes  :                1023                   48 bps
   Input  packets:                  76                    0 pps
   Output packets:                  77                    0 pps
  Serial media information:
    Line protocol: eia530
    Resync history:
      Sync loss count: 0
    Data signal:
      Rx Clock: OK
    Control signals:
      Local mode: DTE
      To DCE: DTR: up, RTS: up
      From DCE: CTS: up, DCD: up, DSR: up
    Clocking mode: loop-timed
    Clock rate: 8.0 MHz
    Loopback: none
    Tx clock: non-invert
    Line encoding: nrz

  Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45) (Generation 9)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 15, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3,
        Generation: 23
```

show interfaces extensive (Serial, EIA-530)

```
user@host> show interfaces se-5/0/1 extensive
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 41, Generation: 25
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
```

```
Keepalive statistics:
  Input : 40 (last seen 00:00:00 ago)
  Output: 37 (last sent 00:00:09 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
CoS queues     : 8 supported, 8 maximum usable queues
Last flapped   : 2006-04-26 15:10:18 PDT (00:06:28 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                  988                   40 bps
 Output bytes  :                 1088                   48 bps
 Input  packets:                   81                    0 pps
 Output packets:                   82                    0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 2, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0
Serial media information:
  Line protocol: eia530
  Resync history:
    Sync loss count: 0
  Data signal:
    Rx Clock: OK
  Control signals:
    Local mode: DTE
    To DCE: DTR: up, RTS: up
    From DCE: CTS: up, DCD: up, DSR: up
  Clocking mode: loop-timed
  Clock rate: 8.0 MHz
  Loopback: none
  Tx clock: non-invert
  Line encoding: nrz
Packet Forwarding Engine configuration:
  Destination slot: 5, PLP byte: 1 (0x00)
CoS information:
  CoS transmit queue        Bandwidth          Buffer   Priority   Limit
                       %          bps     %     usec
  0 best-effort       95     15564800    95        0         low    none
  3 network-control    5       819200     5        0         low    none
```

```
  Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45) (Generation 9)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 15, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3,
        Generation: 23
```

show interfaces (Serial, V.35)

```
user@host> show interfaces se-5/0/0
Physical interface: se-5/0/0, Enabled, Physical link is Down
  Interface index: 150, SNMP ifIndex: 39
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
  Link flags     : Loose-NCP
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 0 (never)
  LCP state: Down
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,   mpls: Not-
configured
  CHAP state: Closed
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2006-04-26 14:51:27 PDT (01:02:23 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27)
    Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
    Encapsulation: PPP
    Protocol inet, MTU: 1500
      Flags: Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3
```

show interfaces brief (Serial, V.35)

```
user@host> show interfaces se-5/0/0 brief
Physical interface: se-5/0/0, Enabled, Physical link is Down
  Type: Serial, Link-level type: PPP, MTU: 1504
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
  Link flags     : Loose-NCP
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 0 (never)

  Logical interface se-5/0/0.0
    Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
    Encapsulation: PPP
    inet  13.0.0.2/30
```

show interfaces detail (Serial, V.35)

```
user@host> show interfaces se-5/0/0 detail
Physical interface: se-5/0/0, Enabled, Physical link is Down
  Interface index: 150, SNMP ifIndex: 39, Generation: 31
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
  Link flags     : Loose-NCP
  Hold-times     : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 0 (last seen: never)
    Output: 0 (last sent: never)
  LCP state: Down
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,   mpls: Not-
configured
  CHAP state: Closed
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2006-04-26 14:51:27 PDT (01:03:15 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
```

```
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Serial media information:
    Line protocol: v.35
    Resync history:
      Sync loss count: 0
    Data signal:
      Rx Clock: Not Detected
    Control signals:
      Local mode: DCE
      To DTE: CTS: down, DCD: down, DSR: up
      From DTE: DTR: down, RTS: down
    DCE loopback override: Off
    Clocking mode: internal
    Clock rate: 38.4 KHz
    Loopback: none
    Tx clock: non-invert
    Line encoding: nrz

  Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27) (Generation 12)
    Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
    Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 17, Route table: 0
      Flags: Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3,
        Generation: 23
```

show interfaces extensive (Serial, V.35)

```
user@host> show interfaces se-5/0/0 extensive
Physical interface: se-5/0/0, Enabled, Physical link is Down
  Interface index: 150, SNMP ifIndex: 39, Generation: 31
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
  Link flags     : Loose-NCP
  Hold-times     : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 0 (last seen: never)
```

```
    Output: 0 (last sent: never)
  LCP state: Down
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,   mpls: Not-
configured
  CHAP state: Closed
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2006-04-26 14:51:27 PDT (01:04:17 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0
  Serial media information:
    Line protocol: v.35
    Resync history:
      Sync loss count: 0
    Data signal:
      Rx Clock: Not Detected
    Control signals:
      Local mode: DCE
      To DTE: CTS: down, DCD: down, DSR: up
      From DTE: DTR: down, RTS: down
    DCE loopback override: Off
    Clocking mode: internal
    Clock rate: 38.4 KHz
    Loopback: none
    Tx clock: non-invert
    Line encoding: nrz
  Packet Forwarding Engine configuration:
    Destination slot: 5, PLP byte: 1 (0x00)
  CoS information:
    CoS transmit queue          Bandwidth         Buffer      Priority   Limit
                          %           bps     %    usec
    0 best-effort        95       15564800    95      0            low    none
    3 network-control     5         819200     5      0            low    none
```

```
  Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27) (Generation 12)
    Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
    Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 17, Route table: 0
      Flags: Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3,
        Generation: 23
```

show interfaces statistics detail (RS 449)

```
user@host> show interfaces se-6/0/0 statistics detail
 Interface index: 149, SNMP ifIndex: 59, Generation: 150
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 8mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : No-Keepalives Loose-NCP
  Hold-times     : Up 0 ms, Down 0 ms
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  PAP state: Closed
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2007-11-28 19:38:36 PST (00:14:06 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 744                    0 bps
   Output bytes  :                5978                    0 bps
   Input  packets:                  33                    0 pps
   Output packets:                 129                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 13, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                   24                   24                    0
    1 expedited-fo                   0                    0                    0
    2 bulk                           0                    0                    0
```

```
    3 assured-forw                 105                  105                    0
    4 voip                           0                    0                    0
  Serial media information:
    Line protocol: rs449
    Resync history:
      Sync loss count: 0
    Data signal:
      Rx Clock: OK
    Control signals:
      Local mode: DTE
      To DCE: DTR: up, RTS: up
      From DCE: CTS: up, DCD: up, DSR: up
    Clocking mode: internal
    Loopback: none
    Tx clock: non-invert
    Line encoding: nrz

  Logical interface se-6/0/0.0 (Index 75) (SNMP ifIndex 69) (Generation 141)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Protocol inet, MTU: 256, Generation: 145, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 11.11.11/24, Local: 11.11.11.2, Broadcast: 11.11.11.255,
        Generation: 157
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces smart-sfp-defects

IN THIS SECTION

- Syntax | 21360
- Description | 21360
- Options | 21360

- Required Privilege Level | 21360
- Output Fields | 21360
- Sample Output | 21369
- Release Information | 21372

Syntax

```
show interfaces smart-sfp-defects
<brief | detail | terse>
<interface-name>
```

Description

Display information about the defects on the smart SFP interface. The defects can be Smart SFP device defects or TDM legacy defects.

Options

brief | detail | terse	(Optional) Display the specified level of output.
interface-name	Name of the interface.

Required Privilege Level

view

Output Fields

Table 687 on page 21361 lists the output fields for the `show interfaces smart-sfp-defects` command. Output fields are listed in the approximate order in which they appear.

Table 687: show interfaces smart-sfp-defects Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail none
SNMP ifIndex	SNMP index number for the physical interface.	detail none
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Clocking	Reference clock source.	detail
Speed	Speed at which the interface is running.	All levels
Duplex	Duplex mode of the interface, either Full-Duplex or Half-Duplex.	All levels
MAC-REWRITE Error	Specifies if the encapsulation of the packet has been changed.	none
BPDU Error	Specifies if a BPDU has been received on a blocked interface.	none
Loopback	Loopback status: Enabled or Disabled. If loopback is enabled, type of loopback: Local or Remote.	All levels
Source filtering	Source filtering status: Enabled or Disabled.	All levels

Table 687: show interfaces smart-sfp-defects Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flow control	Flow control status: Enabled or Disabled. This field is only displayed if asymmetric flow control is not configured.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Smart Transceiver Type	Type of smart SFP transceiver. Possible values: E1, T1, DS3, STM1, STM4, or STM16.	All levels
Smart SFP	Configurable Type. You can configure these Smart SFP transceivers unlike other Smart SFP transceivers which you cannot configure.	All levels
Smart SFP Configurations		
Loopback	Specifies if loopback is configured.	All levels
Encapsulation Circuit Id	Emulation circuit id for encapulsation. Possible values: 0 through 1048575.	All levels
Decapsulation Circuit Id	Emulation circuit id for decapsulation. Possible values: 0 through 1048575.	All levels
Mode	Framing mode. Possible values: MEF8 or MPLS.	All levels
Dest MAC	Destination MAC Address	detail
Smart SFP Defects		
TDM defects	Legacy TDM defects. Defects specific to PDH and SDH legacy error codes and defects.	All

Table 687: show interfaces smart-sfp-defects Output Fields *(Continued)*

Field Name	Field Description	Level of Output
System defects	Smart SFP device defects. Defects are specific to the Smart SFP you use.	All
CoS queues	Number of CoS queues configured.	detail none
Hold-Times	Current interface hold-time up and hold-time down, in milliseconds.	detail
Current address	Configured MAC address.	detail none
Hardware address	Hardware MAC address.	detail none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour.minute.second.timezone*** **(*hour.minute.second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).**	detail none
Statistics last cleared	Date, time, and how long ago the statistics for the interface were cleared. The format is **Statistics last cleared:** ***year-month-day hour.minute.second.timezone*** **(*hour.minute.second* ago)**. For example, 2010-05-17 07:51:28 PDT (00:04:33 ago).	detail
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail

Table 687: show interfaces smart-sfp-defects Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IPv6 transit statistics	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface: • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail

Table 687: show interfaces smart-sfp-defects Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • Errors—Sum of the incoming frame terminates and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops.	detail none

Table 687: show interfaces smart-sfp-defects Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame terminates and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the fabric interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops.	detail none
Egress queues	Total number of egress queues supported on the specified interface.	detail

Table 687: show interfaces smart-sfp-defects Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Queue counters	CoS queue number and its associated user-configured forwarding class name. • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism.	detail
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail none
Flags	Information about the logical interface. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.	All levels
Input packets	Number of packets received on the interface.	detail none
Output packets	Number of packets transmitted on the interface.	detail none

Table 687: show interfaces smart-sfp-defects Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Encapsulation	Encapsulation method used on the logical interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail
Local statistics	• Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail
Transit statistics	• Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail

Table 687: show interfaces smart-sfp-defects Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Addresses, Flags	Information about the address flags.	detail none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	brief
MTU	Maximum transmission unit size on the physical interface.	All levels
Destination	IP address of the remote side of the connection.	detail none
Local	IP address of the logical interface.	detail none
Broadcast	Broadcast address of the logical interlace.	detail none
Generation	Unique number for use by Juniper Networks technical support only.	detail
Route table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail none

Sample Output

show interfaces smart-sfp-defects (T1)

```
user@host> show interfaces ge-4/1/0 smart-sfp-defects
Physical interface: ge-4/1/0, Enabled, Physical link is Up
  Interface index: 216, SNMP ifIndex: 742
  Link-level type: Ethernet-Bridge, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU
Error: None,
  Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote fault:
Online
  Pad to minimum frame size: Disabled
```

```
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x20004000
  Link flags     : None
  Smart Transceiver Type: T1
  Smart SFP: Configurable Type
  Smart SFP Configurations
    Loopback: None
    DMAC check Enabled: enabled
    Encapulation Circuit Id: 400
    Decapsulation Circuit Id: 16
    VLAN1: 400 VLAN2: 400
    Mode: MEF8
    Dest MAC: 00:19:3a:02:a4:b8
  Smart SFP Defects:
      TDM defects     : Local Packet Loss
      System defects  : None

  Logical interface ge-4/1/0.0 (Index 395) (SNMP ifIndex 615)
    Flags: Up SNMP-Traps 0x24004000 Encapsulation: Ethernet-Bridge
    Tenant Name: (null)
    Input packets : 2490261
    Output packets: 2486756
    Protocol bridge, MTU: 1514
```

show interfaces smart-sfp-defects (DS3)

```
user@host> show interfaces ge-4/0/5 smart-sfp-defects
Physical interface: ge-4/0/5, Enabled, Physical link is Up
  Interface index: 211, SNMP ifIndex: 735
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
None, Loop Detect PDU Error: None,
  Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  Smart Transceiver Type: DS3
  Smart SFP: Configurable Type
  Smart SFP Configurations
```

```
    Loopback: None
    Encapulation Circuit Id: 16
    Decapsulation Circuit Id: 16
    Mode: MPLS
  Smart SFP Defects:
      TDM defects     : Loss of Signal, Local Packet Loss, Out Of Frame
      System defects  : None

  Logical interface ge-4/0/5.0 (Index 355) (SNMP ifIndex 612)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 430
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
```

show interfaces smart-sfp-defects (STM1)

```
user@host > show interfaces ge-4/0/0 smart-sfp-defects
Physical interface: ge-4/0/0, Enabled, Physical link is Up
  Interface index: 192, SNMP ifIndex: 546, Generation: 269
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
None,
  Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Enabled,
  Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  Smart Transceiver Type: STM1
  Smart SFP: Configurable Type
  Smart SFP Configurations
    Loopback: None
    DMAC check Enabled: enabled
    Encapulation Circuit Id: 1000
    Decapsulation Circuit Id: 2000
    VLAN1: 100
    Mode: MEF8
    Dest MAC: 10:0e:7e:37:cd:29
  Smart SFP Defects:
      TDM defects     : Local Packet Loss
```

```
      System defects   : None

Logical interface ge-4/0/0.0 (Index 369) (SNMP ifIndex 799) (Generation 190)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
Input packets : 2400
    Output packets: 0
   Protocol multiservice, MTU: Unlimited
```

Release Information

Command introduced in Junos OS Release 19.4.

RELATED DOCUMENTATION

Smart SFPs To Transport Legacy Network Traffic over Packet Switched Networks

ces-psn-channel (tdm-options) | **3265**

iwf-params (tdm-options) | **9420**

tdm-options (Interfaces) | **29039**

show interfaces smart-sfp-statistics

IN THIS SECTION

- Syntax | **21373**
- Description | **21373**
- Options | **21373**
- Required Privilege Level | **21373**
- Output Fields | **21373**
- Sample Output | **21382**
- Release Information | **21388**

Syntax

```
show interfaces smart-sfp-statistics
<brief | detail | extensive | terse>
<descriptions>
<interface-name>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
```

Description

Display status information about the specified smart SFP interface.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
interface-name	Name of the interface.
media	(Optional) Display media-specific information about network interfaces.
routing-instance (all | *instance-name*)	(Optional) Display all routing instances or the name of an individual routing instance.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.

Required Privilege Level

view

Output Fields

Table 688 on page 21374 lists the output fields for the `show interfaces smart-sfp-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 688: show interfaces smart-sfp-statistics Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail none
SNMP ifIndex	SNMP index number for the physical interface.	detail none
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Clocking	Reference clock source.	detail
Speed	Speed at which the interface is running.	All levels
Duplex	Duplex mode of the interface, either Full-Duplex or Half-Duplex.	All levels
MAC-REWRITE Error	Specifies if the encapsulation of the packet has been changed.	none
BPDU Error	Specifies if a BPDU has been received on a blocked interface.	none
Loopback	Loopback status: Enabled or Disabled. If loopback is enabled, type of loopback: Local or Remote.	All levels
Source filtering	Source filtering status: Enabled or Disabled.	All levels

Table 688: show interfaces smart-sfp-statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flow control	Flow control status: Enabled or Disabled. This field is only displayed if asymmetric flow control is not configured.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Smart Transceiver Type	Type of smart SFP transceiver. Possible values: E1, T1, DS3, STM1, STM4, or STM16.	All levels
Smart SFP	Configurable Type. You can configure these Smart SFP transceivers unlike other Smart SFP transceivers which you cannot configure.	All levels
Smart SFP Configurations		
Loopback	Specifies if loopback is configured.	All levels
Encapsulation Circuit Id	Emulation circuit id for encapulsation. Possible values: 0 through 1048575.	all
Decapsulation Circuit Id	Emulation circuit id for decapsulation. Possible values: 0 through 1048575.	all
Mode	Framing mode. Possible values: MEF8 or MPLS.	all
Dest MAC	Destination MAC address.	all
Smart SFP Defects		
TDM defects	Legacy TDM defects. Defects specific to PDH and SDH legacy error codes and defects.	all

Table 688: show interfaces smart-sfp-statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
System defects	Smart SFP device defects. Defects are specific to the Smart SFP you use.	
CoS queues	Number of CoS queues configured.	detail none
Hold-Times	Current interface hold-time up and hold-time down, in milliseconds.	detail
Current address	Configured MAC address.	detail none
Hardware address	Hardware MAC address.	detail none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour.minute.second.timezone*** **(*hour.minute.second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).**	detail none
Statistics last cleared	Date, time, and how long ago the statistics for the interface were cleared. The format is **Statistics last cleared:** ***year-month-day hour.minute.second.timezone*** **(*hour.minute.second* ago)**. For example, 2010-05-17 07:51:28 PDT (00:04:33 ago).	detail
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail

Table 688: show interfaces smart-sfp-statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IPv6 transit statistics	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface: • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail

Table 688: show interfaces smart-sfp-statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • Errors—Sum of the incoming frame terminates and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops.	detail none

Table 688: show interfaces smart-sfp-statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame terminates and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the fabric interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops.	detail none
Egress queues	Total number of egress queues supported on the specified interface.	detail

Table 688: show interfaces smart-sfp-statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Queue counters	CoS queue number and its associated user-configured forwarding class name. • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism.	detail
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail none
Flags	Information about the logical interface. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.	All levels
Input packets	Number of packets received on the interface.	detail none
Output packets	Number of packets transmitted on the interface.	detail none

Table 688: show interfaces smart-sfp-statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Encapsulation	Encapsulation method used on the logical interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail
Local statistics	• Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail
Transit statistics	• Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail

Table 688: show interfaces smart-sfp-statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Addresses, Flags	Information about the address flags.	detail none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	brief
MTU	Maximum transmission unit size on the physical interface.	All levels
Destination	IP address of the remote side of the connection.	detail none
Local	IP address of the logical interface.	detail none
Broadcast	Broadcast address of the logical interlace.	detail none
Generation	Unique number for use by Juniper Networks technical support only.	detail
Route table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail none

Sample Output

show interfaces smart-sfp-statistics (STM1)

```
user@host> show interfaces ge-5/3/9 smart-sfp-statistics
Physical interface: ge-5/3/9, Enabled, Physical link is Up
  Interface index: 300, SNMP ifIndex: 562
  Link-level type: Ethernet, MTU: 1518, MRU: 1526, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
None, Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Enabled, Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
```

```
  Smart Transceiver Type: STM1/OC3
  Smart SFP: Configurable Type
  Smart SFP Configurations
    Loopback: None
    Mode: MEF8
  Smart SFP Ethernet port[P1] Statistics:            Counters
    Rx frames                                        64376
    Rx bytes                                       4128352
    Rx errored fcs frames                                0
    Rx length error frames                               0
    Tx frames                                 2552954627
    Tx bytes                                   120088140
  Smart SFP AV IWF Encap/Decap Statistics:           Counters
    Rx Packets                                           0
    Tx Packets                                2552812339
    Malformed  Packets                                   0
    Reordered  Packets                                   0
    Misordered Dropped Packets                           0
    Missing    Packets                                   0
    PlayedOut  Packets                                   0
    JitterBuffer Overrun Packets                         0
    JitterBuffer Underrun Packets                        0
  Smart SFP STM port[P0] statistics:             Counters
    Rx b1errors                                          0
    Tx b1errors                                          0

  Logical interface ge-5/3/9.32767 (Index 347) (SNMP ifIndex 659)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Input packets : 171346
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

show interfaces smart-sfp-statistics detail (STM1)

```
user@host> show interfaces ge-5/3/9 smart-sfp-statistics detail
Physical interface: ge-5/3/9, Enabled, Physical link is Up
  Interface index: 300, SNMP ifIndex: 562, Generation: 303
  Link-level type: Ethernet, MTU: 1518, MRU: 1526, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
None, Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
```

```
Enabled, Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Smart Transceiver Type: STM1/OC3
  Smart SFP: Configurable Type
  Smart SFP Configurations
    Loopback: None
    Mode: MEF8
  Smart SFP Ethernet port[P1] Statistics:            Counters
    Rx frames                                       105783
    Rx bytes                                       6778400
    Rx errored fcs frames                                0
    Rx length error frames                               0
    Tx frames                                 4209457628
    Tx bytes                                   301542720
  Smart SFP AV IWF Encap/Decap Statistics:           Counters
    Rx Packets                                           0
    Tx Packets                                4209223157
    Malformed  Packets                                   0
    Reordered  Packets                                   0
    Misordered Dropped Packets                           0
    Missing    Packets                                   0
    PlayedOut  Packets                                   0
    JitterBuffer Overrun Packets                         0
    JitterBuffer Underrun Packets                        0
  Smart SFP STM port[P0] statistics:             Counters
    Rx b1errors                                          0
    Tx b1errors                                          0

  Logical interface ge-5/3/9.32767 (Index 347) (SNMP ifIndex 659) (Generation 186)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             23232960
     Output bytes  :                    0
     Input  packets:               281769
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
```

```
     Input  bytes  :             23232960                 1008 bps
     Output bytes  :                    0                    0 bps
     Input  packets:               281769                    1 pps
     Output packets:                    0                    0 pps
    Protocol multiservice, MTU: Unlimited, Generation: 227, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__
```

show interfaces smart-sfp-statistics terse (STM1)

```
user@host> show interfaces ge-5/3/9 smart-sfp-statistics terse
Interface               Admin Link Proto    Local                 Remote
ge-5/3/9                up    up
ge-5/3/9.32767          up    up   multiservice
```

show interfaces smart-sfp-statistics (DS3)

```
user@host> show interfaces ge-5/3/0 smart-sfp-statistics
Physical interface: ge-5/3/0, Enabled, Physical link is Up
  Interface index: 291, SNMP ifIndex: 559
  Link-level type: Ethernet, MTU: 1518, MRU: 1526, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
None, Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Enabled, Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Smart Transceiver Type: DS3
  Smart SFP: Configurable Type
  Smart SFP Configurations
    Loopback: None
    Encapulation Circuit Id: 16
    Decapsulation Circuit Id: 16
    Mode: MEF8
    Dest MAC: 01:0f:25:00:00:03
  Smart SFP Ethernet port[P1] Statistics:            Counters
    Rx frames                                         3013
    Rx bytes                                        220548
    Rx errored fcs frames                                0
    Rx unicast frames                                 3013
```

```
    Rx multicast frames                                0
    Rx broadcast frames                                0
    Rx fragments                                       0
    Rx undersize frames                                0
    Rx oversize frames                                 0
    Rx invalid vlan mismatch frames                    0
    Tx frames                                  110842256
    Tx bytes                                  4202709764
    Tx unicast frames                              18256
    Tx multicast frames                        110824000
    Tx broadcast frames                                0
  Smart SFP AV IWF Encap/Decap Statistics:           Counters
    Rx Packets                                         0
    Tx Packets                                 110824000
    Malformed  Packets                                 0
    Reordered  Packets                                 0
    Misordered Dropped Packets                         0
    Missing    Packets                                 0
    PlayedOut  Packets                                 0
    JitterBuffer Overrun Packets                       0
    JitterBuffer Underrun Packets                      0
  Smart SFP DS3 port[P0] statistics:             Counters
    BiPolarVariations/Excessive zero errors                  0
    Tx B3 Errors                                             0
    Code Violation path errors                               0

  Logical interface ge-5/3/0.32767 (Index 358) (SNMP ifIndex 660)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Input packets : 18341
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

show interfaces smart-sfp-statistics detail (DS3)

```
Physical interface: ge-5/3/0, Enabled, Physical link is Up
  Interface index: 291, SNMP ifIndex: 559, Generation: 294
  Link-level type: Ethernet, MTU: 1518, MRU: 1526, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
None, Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Enabled, Remote fault: Online
```

```
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Smart Transceiver Type: DS3
  Smart SFP: Configurable Type
  Smart SFP Configurations
    Loopback: None
    Encapulation Circuit Id: 16
    Decapsulation Circuit Id: 16
    Mode: MEF8
    Dest MAC: 01:0f:25:00:00:03
  Smart SFP Ethernet port[P1] Statistics:            Counters
    Rx frames                                          4329
    Rx bytes                                         316992
    Rx errored fcs frames                                 0
    Rx unicast frames                                  4329
    Rx multicast frames                                   0
    Rx broadcast frames                                   0
    Rx fragments                                          0
    Rx undersize frames                                   0
    Rx oversize frames                                    0
    Rx invalid vlan mismatch frames                       0
    Tx frames                                     159426251
    Tx bytes                                      746264028
    Tx unicast frames                                 26251
    Tx multicast frames                           159400000
    Tx broadcast frames                                   0
  Smart SFP AV IWF Encap/Decap Statistics:           Counters
    Rx Packets                                            0
    Tx Packets                                    159400000
    Malformed  Packets                                    0
    Reordered  Packets                                    0
    Misordered Dropped Packets                            0
    Missing    Packets                                    0
    PlayedOut  Packets                                    0
    JitterBuffer Overrun Packets                          0
    JitterBuffer Underrun Packets                         0
  Smart SFP DS3 port[P0] statistics:             Counters
    BiPolarVariations/Excessive zero errors                  0
    Tx B3 Errors                                             0
    Code Violation path errors                               0

  Logical interface ge-5/3/0.32767 (Index 358) (SNMP ifIndex 660) (Generation 197)
```

```
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :              2985128
     Output bytes  :                    0
     Input  packets:                26336
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
  Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :              2985128                 1024 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                26336                    1 pps
     Output packets:                    0                    0 pps
    Protocol multiservice, MTU: Unlimited, Generation: 242, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces smart-sfp-statistics terse (DS3)

```
user@host> show interfaces ge-5/3/0 smart-sfp-statistics terse
Interface               Admin Link Proto    Local                 Remote
ge-5/3/0                up    up
ge-5/3/0.32767          up    up   multiservice
```

command-name

Release Information

Command introduced in Junos OS Release 19.4.

RELATED DOCUMENTATION

Smart SFPs To Transport Legacy Network Traffic over Packet Switched Networks

ces-psn-channel (tdm-options) | 3265

iwf-params (tdm-options) | 9420
tdm-options (Interfaces) | 29039

show interfaces snmp-index

IN THIS SECTION

- Syntax | 21389
- Description | 21389
- Options | 21389
- Additional Information | 21389
- Required Privilege Level | 21390
- Output Fields | 21390
- Sample Output | 21390
- Release Information | 21390

Syntax

```
show interfaces snmp-index snmp-index
```

Description

Display information for the interface with the specified SNMP index.

Options

This command has no options.

Additional Information

Output from both the `show interfaces` *`interface-name`* `detail` and the `show interfaces` *`interface-name`* `extensive` command includes all the information displayed in the output from the `show interfaces snmp-index` command.

Required Privilege Level

view

Output Fields

The output fields from the `show interfaces snmp-index` *`snmp-index`* command are identical to those produced by the `show interfaces` *`interface-name`* command. For a description of output fields, see the other chapters in this manual.

Sample Output

show interfaces snmp-index

```
user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
  Interface index: 149, SNMP ifIndex: 33
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,   Loopback:
None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : LOL, PLL, LOS
  SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces source-class

IN THIS SECTION

- Syntax | 21391
- Description | 21391
- Options | 21391
- Additional Information | 21392
- Required Privilege Level | 21392
- Output Fields | 21392
- Sample Output | 21392
- Release Information | 21394

Syntax

```
show interfaces source-class
(all / destination-class-name logical-interface-name)
```

Description

Display information about interfaces grouped by source class.

Options

all	Display information about all configured source classes.
source-class-name	Name of a logical grouping of prefixes that count packets having the source address matching those prefixes.
interface-name	Name of a logical interface.

Additional Information

For interfaces that carry IPv4, IPv6, or Multiprotocol Label Switching (MPLS) traffic, you can maintain packet counts based on the entry and exit points for traffic passing through your network. Entry and exit points are identified by source and destination prefixes grouped into sets defined as source classes and destination classes. For more information, see the Junos OS Network Interfaces Library for Routing Devices.

Required Privilege Level

view

Output Fields

Table 689 on page 21392 lists the output fields for the `show interfaces source-class` command. Output fields are listed in the approximate order in which they appear.

Table 689: show interfaces source-class Output Fields

Field Name	Field Description
Logical interface	Name of the logical interface.
Source class	Source class usage (SCU) counters per class for this interface.
Packets	Packets going to designated user-selected prefixes.
Bytes	Bytes going to designated user-selected prefixes.

Sample Output

show interfaces source-class all

```
user@host> show interfaces source-class all
 Logical interface .local..1

  Logical interface .local..2
```

```
Logical interface fxp0.0

Logical interface fxp1.0

Logical interface lo0.16384

Logical interface lo0.16385

Logical interface lo0.0

Logical interface .local..3

Logical interface .local..4

Logical interface .local..5

Logical interface .local..6

Logical interface .local..7

Logical interface .local..8

Logical interface .local..9

Logical interface .local..10

Logical interface lo0.3

Logical interface pfh-0/0/0.16383

Logical interface fe-0/1/0.0

Logical interface fe-0/1/1.0
  Protocol inet
                                              Packets                Bytes
    Source class                 (packet-per-second)    (bits-per-second)

                           GOLD1                    0                    0
                                 (                  0) (                 0)
                           GOLD2                    0                    0
                                 (                  0) (                 0)
                           GOLD3                    0                    0
                                 (                  0) (                 0)
```

```
                          v4-src                    0                    0
                                 (                  0) (                 0)
    Protocol inet6
                                                Packets                Bytes
      Source class                 (packet-per-second)    (bits-per-second)

                           GOLD1                    0                    0
                                 (                  0) (                 0)
                           GOLD2                    0                    0
                                 (                  0) (                 0)
                           GOLD3                    0                    0
                                 (                  0) (                 0)
                          v4-src                    0                    0
                                 (                  0) (                 0)

  Logical interface fe-0/1/2.0
    Description: CE1-to-PE2

  Logical interface ge-0/3/0.0
    Description: CE1-to-PE1

  Logical interface ge-0/3/2.0
    Description: CE2-to-PE1

  Logical interface pc-0/3/0.16383

  Logical interface lt-1/2/0.3
    Description: LS3->LS2

  Logical interface lt-1/2/0.5
    Description: LS3->LS1

  Logical interface sp-1/2/0.16383
```

Release Information

Command introduced before Junos OS Release 7.4.

all option introduced in Junos OS Release 8.0.

show interfaces (SRX Series)

IN THIS SECTION

- Syntax | 21396
- Description | 21397
- Options | 21397
- Required Privilege Level | 21398
- Output Fields | 21398
- Sample Output | 21412
- Sample Output | 21413
- Sample Output | 21414
- Sample Output | 21416
- Sample Output | 21419
- Sample Output | 21420
- Sample Output | 21421
- Sample Output | 21421
- Sample Output | 21422
- Sample Output | 21423
- Sample Output | 21424
- Sample Output | 21425
- Sample Output | 21426
- Sample Output | 21428
- Sample Output | 21428
- Sample Output | 21431
- Sample Output | 21432
- Sample Output | 21432
- Sample Output | 21434
- Sample Output | 21435
- Sample Output | 21436
- Sample Output | 21437
- Sample Output | 21437

- Sample Output | 21438
- Sample Output | 21439
- Sample Output | 21440
- Sample Output | 21441
- Release Information | 21442

Syntax

```
show interfaces {
    <brief | detail | extensive | terse>
    controller interface-name
    descriptions interface-name
    destination-class (all / destination-class-name logical-interface-name)
    diagnostics optics interface-name
    far-end-interval interface-fpc/pic/port
    filters interface-name
    flow-statistics interface-name
    interval interface-name
    load-balancing  (detail | interface-name)
    mac-database mac-address mac-address
    mc-ae id identifier unit number revertive-info
    media interface-name
    policers interface-name
    queue both-ingress-egress egress forwarding-class forwarding-class ingress l2-statistics
    redundancy  (detail | interface-name)
    routing brief detail summary interface-name
    routing-instance (all | instance-name)
    snmp-index snmp-index
    source-class (all / destination-class-name logical-interface-name)
    statistics interface-name
    switch-port switch-port number
    transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all
| interface-name)
    zone interface-name
}
```

Description

Display status information and statistics about interfaces on SRX Series appliance running Junos OS.

On SRX Series appliance, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

Options

- `interface-name`—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.
  - `at-` *pim*`/0/`*port*—ATM-over-ADSL or ATM-over-SHDSL interface.
  - `ce1-`*pim*`/0/` *port*—Channelized E1 interface.
  - `cl-0/0/8`—3G wireless modem interface for SRX320 devices.
  - `ct1-`*pim*`/0/`*port*—Channelized T1 interface.
  - `dl0`—Dialer Interface for initiating ISDN and USB modem connections.
  - `e1-`*pim*`/0/`*port*—E1 interface.
  - `e3-`*pim*`/0/`*port*—E3 interface.
  - `fe-`*pim*`/0/`*port*—Fast Ethernet interface.
  - `ge-`*pim*`/0/`*port*—Gigabit Ethernet interface.
  - `se-`*pim*`/0/`*port*—Serial interface.
  - `t1-`*pim*`/0/`*port*—T1 (also called DS1) interface.
  - `t3-`*pim*`/0/`*port*—T3 (also called DS3) interface.
  - `wx-`*slot*`/0/0`—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).
- `brief | detail | extensive | terse`—(Optional) Display the specified level of output.
- `controller`—(Optional) Show controller information.
- `descriptions`—(Optional) Display interface description strings.
- `destination-class`—(Optional) Show statistics for destination class.
- `diagnostics`—(Optional) Show interface diagnostics information.

- far-end-interval—(Optional) Show far end interval statistics.
- filters—(Optional) Show interface filters information.
- flow-statistics—(Optional) Show security flow counters and errors.
- interval—(Optional) Show interval statistics.
- load-balancing—(Optional) Show load-balancing status.
- mac-database—(Optional) Show media access control database information.
- mc-ae—(Optional) Show MC-AE configured interface information.
- media—(Optional) Display media information.
- policers—(Optional) Show interface policers information.
- queue—(Optional) Show queue statistics for this interface.
- redundancy—(Optional) Show redundancy status.
- routing—(Optional) Show routing status.
- routing-instance—(Optional) Name of routing instance.
- snmp-index—(Optional) SNMP index of interface.
- source-class—(Optional) Show statistics for source class.
- statistics—(Optional) Display statistics and detailed output.
- switch-port—(Optional) Front end port number (0..15).
- transport—(Optional) Show interface transport information.
- zone—(Optional) Interface's zone.

Required Privilege Level

view

Output Fields

Table 690 on page 21399 lists the output fields for the show interfaces command. Output fields are listed in the approximate order in which they appear.

Table 690: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical interface`	Name of the physical interface.	All levels
`Enabled`	State of the interface.	All levels
`Interface index`	Index number of the physical interface, which reflects its initialization sequence.	`detail` `extensive` none
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail` `extensive` none
`Link-level type`	Encapsulation being used on the physical interface.	All levels
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`
`MTU`	Maximum transmission unit size on the physical interface.	All levels
`Link mode`	Link mode: Full-duplex or Half-duplex.	
`Speed`	Speed at which the interface is running.	All levels
`BPDU error`	Bridge protocol data unit (BPDU) error: Detected or None	
`Loopback`	Loopback status: `Enabled` or `Disabled`. If loopback is enabled, type of loopback: `Local` or `Remote`.	All levels
`Source filtering`	Source filtering status: `Enabled` or `Disabled`.	All levels
`Flow control`	Flow control status: `Enabled` or `Disabled`.	All levels

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Auto-negotiation`	(Gigabit Ethernet interfaces) Autonegotiation status: `Enabled` or `Disabled`.	All levels
`Remote-fault`	(Gigabit Ethernet interfaces) Remote fault status: • `Online`—Autonegotiation is manually configured as online. • `Offline`—Autonegotiation is manually configured as offline.	All levels
`Device flags`	Information about the physical device.	All levels
`Interface flags`	Information about the interface.	All levels
`Link flags`	Information about the physical link.	All levels
`CoS queues`	Number of CoS queues configured.	`detail extensive` none
`Current address`	Configured MAC address.	`detail extensive` none
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second`*`:`*`timezone`* `(`*`hour`*`:`*`minute`*`:`*`second`* `ago)`. For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive` none
`Input Rate`	Input rate in bits per second (bps) and packets per second (pps).	None
`Output Rate`	Output rate in bps and pps.	None

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Active alarms and Active defects`	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value `None` or `Link`. • `None`—There are no active defects or alarms. • `Link`—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	`detail extensive` none
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input errors`	Input errors on the interface. • `Errors`—Sum of the incoming frame terminates and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • `L3 incompletes`—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` statement. • `L2 channel errors`—Number of times the software did not find a valid logical interface for an incoming frame. • `L2 mismatch timeouts`—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • `FIFO errors`—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `Resource errors`—Sum of transmit drops.	`extensive`

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • `Errors`—Sum of the outgoing frame terminates and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Collisions`—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • `Aged packets`—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • `FIFO errors`—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `HS link CRC errors`—Number of errors on the high-speed links between the ASICs responsible for handling the interfaces. • `MTU errors`—Number of packets whose size exceeded the MTU of the interface. • `Resource errors`—Sum of transmit drops.	`extensive`
`Ingress queues`	Total number of ingress queues supported on the specified interface.	`extensive`

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Queue counters and queue number`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	`detail extensive`

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`MAC statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC subsystem, including the following:
	• `Total octets` and `total packets`—Total number of octets and packets.
	• `Unicast packets`, `Broadcast packets`, and `Multicast packets`—Number of unicast, broadcast, and multicast packets.
	• `CRC/Align errors`—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
	• `FIFO error`—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
	• `MAC control frames`—Number of MAC control frames.
	• `MAC pause frames`—Number of MAC control frames with `pause` operational code.
	• `Oversized frames`—There are two possible conditions regarding the number of oversized frames:
	  • Packet length exceeds 1518 octets, or
	  • Packet length exceeds MRU
	• `Jabber frames`—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
	• `Fragment frames`—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames
		`extensive`

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • `VLAN tagged frames`—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • `Code violations`—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Filter statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • `Input packet count`—Number of packets received from the MAC hardware that the filter processed. • `Input packet rejects`—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • `Input DA rejects`—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local device (which the router is rejecting). • `Input SA rejects`—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • `Output packet count`—Number of packets that the filter has given to the MAC hardware. • `Output packet pad count`—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • `Output packet error count`—Number of packets with an indicated error that the filter was given to transmit. These packets are usually	`extensive`

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • `CAM destination filters, CAM source filters`—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	

`Autonegotiation information`	Information about link autonegotiation.
	• `Negotiation status:`
	  • `Incomplete`—Ethernet interface has the speed or link mode configured.
	  • `No autonegotiation`—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • `Complete`—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
		`extensive`

`Packet Forwarding Engine configuration`	Information about the configuration of the Packet Forwarding Engine: • `Destination slot`—FPC slot number.	`extensive`

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`CoS information`	Information about the CoS queue for the physical interface. • `CoS transmit queue`—Queue number and its associated user-configured forwarding class name. • `Bandwidth %`—Percentage of bandwidth allocated to the queue. • `Bandwidth bps`—Bandwidth allocated to the queue (in bps). • `Buffer %`—Percentage of buffer space allocated to the queue. • `Buffer usec`—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • `Priority`—Queue priority: `low` or `high`. • `Limit`—Displayed if rate limiting is configured for the queue. Possible values are `none` and `exact`. If `exact` is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If `none` is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	`extensive`
`Interface transmit statistics`	Status of the `interface-transmit-statistics` configuration: Enabled or Disabled.	`detail extensive`
`Queue counters (Egress)`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	`detail extensive`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Index`	Index number of the logical interface, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP interface index number for the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Flags`	Information about the logical interface.	All levels
`Encapsulation`	Encapsulation on the logical interface.	All levels
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the specified interface set. • `Input bytes, Output bytes`—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • `Input packets, Output packets`—Number of packets received and transmitted on the interface set.	`detail extensive`
`Local statistics`	Number and rate of bytes and packets destined to the device.	`extensive`

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Transit statistics`	Number and rate of bytes and packets transiting the switch. **NOTE**: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the `Output bytes` and `Output packets` interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	`extensive`
`Security`	Security zones that interface belongs to.	`extensive`
`Flow Input statistics`	Statistics on packets received by flow module.	`extensive`
`Flow Output statistics`	Statistics on packets sent by flow module.	`extensive`
`Flow error statistics (Packets dropped due to)`	Statistics on errors in the flow module.	`extensive`
`Protocol`	Protocol family.	`detail extensive` none
`MTU`	Maximum transmission unit size on the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Route Table`	Route table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail extensive` none

Table 690: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Information about protocol family flags. .	`detail extensive`
`Addresses, Flags`	Information about the address flags..	`detail extensive` none
`Destination`	IP address of the remote side of the connection.	`detail extensive` none
`Local`	IP address of the logical interface.	`detail extensive` none
`Broadcast`	Broadcast address of the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Sample Output

show interfaces Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:12:e4:b1:01, Hardware address: 00:00:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
```

```
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255
```

Sample Output

show interfaces brief (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface ge-3/0/2.0
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push     0x8100.512
0x8100.513)
    Encapsulation: VLAN-CCC
    ccc

  Logical interface ge-3/0/2.32767
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
```

Sample Output

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1 detail
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps, BPDU Error:
None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
```

```
 Output bytes  :                    0
 Input  packets:                    0
 Output packets:                    0
Local statistics:
 Input  bytes  :                    0
 Output bytes  :                    0
 Input  packets:                    0
 Output packets:                    0
Transit statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :                    0                    0 bps
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
Security: Zone: public
Flow Statistics :
Flow Input statistics :
  Self packets :                     0
  ICMP packets :                     0
  VPN packets :                      0
  Multicast packets :                0
  Bytes permitted by policy :        0
  Connections established :          0
Flow Output statistics:
  Multicast packets :                0
  Bytes permitted by policy :        0
Flow error statistics (Packets dropped due to):
  Address spoofing:                  0
  Authentication failed:             0
  Incoming NAT errors:               0
  Invalid zone received packet:      0
  Multiple user authentications:     0
  Multiple incoming NAT:             0
  No parent for a gate:              0
  No one interested in self packets: 0
  No minor session:                  0
  No more sessions:                  0
  No NAT gate:                       0
  No route present:                  0
  No SA for incoming SPI:            0
  No tunnel found:                   0
  No session for a gate:             0
  No zone or NULL zone binding       0
  Policy denied:                     0
```

```
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 1500, Generation: 150, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255, Generation: 150
```

Sample Output

show interfaces extensive (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
```

```
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                    0                    0                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont                   0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:                      Receive         Transmit
  Total octets                             0                0
  Total packets                            0                0
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 2, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth               Buffer Priority   Limit
```

```
                              %            bps     %           usec
  0 best-effort              95      950000000    95              0      low    none
  3 network-control           5       50000000     5              0      low    none
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Local statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Transit statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Security: Zone: public
  Flow Statistics :
  Flow Input statistics :
    Self packets :                     0
    ICMP packets :                     0
    VPN packets :                      0
    Multicast packets :                0
    Bytes permitted by policy :        0
    Connections established :          0
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        0
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
```

```
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
      No route present:                  0
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 1500, Generation: 150, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
        Generation: 150
```

Sample Output

show interfaces terse

```
user@host> show interfaces terse
    Interface               Admin Link Proto    Local                 Remote
    ge-0/0/0                up    up
    ge-0/0/0.0              up    up   inet     10.209.4.61/18
    gr-0/0/0                up    up
    ip-0/0/0                up    up
    st0                     up    up
    st0.1                   up    ready inet
    ls-0/0/0                up    up
    lt-0/0/0                up    up
    mt-0/0/0                up    up
    pd-0/0/0                up    up
    pe-0/0/0                up    up
    e3-1/0/0                up    up
    t3-2/0/0                up    up
    e1-3/0/0                up    up
```

```
    se-4/0/0                up    down
    t1-5/0/0                up    up
    br-6/0/0                up    up
    dc-6/0/0                up    up
    dc-6/0/0.32767          up    up
    bc-6/0/0:1              down  up
    bc-6/0/0:1.0            up    down
    dl0                     up    up
    dl0.0                   up    up   inet
    dsc                     up    up
    gre                     up    up
    ipip                    up    up
    lo0                     up    up
    lo0.16385               up    up   inet     10.0.0.1            --> 0/0
                                                 10.0.0.16           --> 0/0
    lsi                     up    up
    mtun                    up    up
    pimd                    up    up
    pime                    up    up
    pp0                     up    up
```

Sample Output

show interfaces controller (Channelized E1 IQ with Logical E1)

```
user@host> show interfaces controller ce1-1/2/6

Controller                                                        Admin Link
ce1-1/2/6                                                         up    up
e1-1/2/6                                                          up    up
```

show interfaces controller (Channelized E1 IQ with Logical DS0)

```
user@host> show interfaces controller ce1-1/2/3

Controller                                                        Admin Link
ce1-1/2/3                                                         up    up
ds-1/2/3:1                                                        up    up
ds-1/2/3:2                                                        up    up
```

Sample Output

show interfaces descriptions

```
user@host> show interfaces descriptions
 Interface       Admin Link Description
 so-1/0/0        up    up   M20-3#1
 so-2/0/0        up    up   GSR-12#1
 ge-3/0/0        up    up   SMB-OSPF_Area300
 so-3/3/0        up    up   GSR-13#1
 so-3/3/1        up    up   GSR-13#2
 ge-4/0/0        up    up   T320-7#1
 ge-5/0/0        up    up   T320-7#2
 so-7/1/0        up    up   M160-6#1
 ge-8/0/0        up    up   T320-7#3
 ge-9/0/0        up    up   T320-7#4
 so-10/0/0       up    up   M160-6#2
 so-13/0/0       up    up   M20-3#2
 so-14/0/0       up    up   GSR-12#2
 ge-15/0/0       up    up   SMB-OSPF_Area100
 ge-15/0/1       up    up   GSR-13#3
```

Sample Output

show interfaces destination-class all

```
user@host> show interfaces destination-class all
Logical interface so-4/0/0.0
                                              Packets                Bytes
      Destination class           (packet-per-second)    (bits-per-second)
                        gold                        0                    0
                                (                   0) (                 0)
                      silver                        0                    0
                                (                   0) (                 0)
  Logical interface so-0/1/3.0
                                              Packets                Bytes
      Destination class           (packet-per-second)    (bits-per-second)
                        gold                        0                    0
                                (                   0) (                 0)
```

```
                        silver                        0                    0
                                   (                  0) (                 0)
```

Sample Output

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                        :  7.408 mA
    Laser output power                        :  0.3500 mW / -4.56 dBm
    Module temperature                        :  23 degrees C / 73 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
```

```
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

Sample Output

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:00-04:15:
...
```

show interfaces far-end-interval coc1-5/2/1:1

```
user@host> run show interfaces far-end-interval coc1-5/2/1:1
 Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:00-04:15:
```

Sample Output

show interfaces filters

```
user@host> show interfaces filters
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
ge-5/0/0        up    up
ge-5/0/0.0      up    up   any                        f-any
                           inet                       f-inet
                           multiservice
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
vt-0/3/0        up    up
at-1/0/0        up    up
at-1/0/0.0      up    up   inet
                           iso
```

```
at-1/1/0        up    down
at-1/1/0.0      up    down inet
                           iso
....
```

Sample Output

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp  ldp msdp nhrp ospf pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http https ike
    netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                     0
      ICMP packets :                     0
      VPN packets :                      2564
      Bytes permitted by policy :        3478
      Connections established :          1
    Flow Output statistics:
      Multicast packets :                0
      Bytes permitted by policy :        16994
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
      Authentication failed:             0
      Incoming NAT errors:               0
      Invalid zone received packet:      0
      Multiple user authentications:     0
      Multiple incoming NAT:             0
      No parent for a gate:              0
      No one interested in self packets: 0
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
```

```
      No route present:                  0
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination:     203.0.113.1/24, Local: 203.0.113.2, Broadcast: 203.0.113.255
```

Sample Output

show interfaces interval (Channelized OC12)

```
user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (E3)

```
user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH)

```
user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
  20:02-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:47-20:02:
    ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
    ES-P: 267, SES-P: 267, UAS-P: 267
  19:32-19:47:
    ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,     SES-P: 56,
UAS-P: 46
  19:17-19:32:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
```

```
  19:02-19:17:
.....
```

Sample Output

show interfaces load-balancing

```
user@host> show interfaces load-balancing
Interface  State            Last change  Member count
ams0       Up               1d 00:50     2
ams1       Up               00:00:59     2
```

show interfaces load-balancing detail

```
user@host>show interfaces load-balancing detail
Load-balancing interfaces detail
Interface        : ams0
  State          : Up
  Last change    : 1d 00:51
  Member count   : 2
  Members        :
      Interface    Weight   State
      mams-2/0/0    10        Active
      mams-2/1/0    10        Active
```

Sample Output

show interfaces mac-database (All MAC Addresses on a Port)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback: None, Source
filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
```

```
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:03      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:04      30424716     1399536936       37448523     1722632058
  00:00:c8:01:01:05      30424789     1399540294       37448598     1722635508
  00:00:c8:01:01:06      30424788     1399540248       37448597     1722635462
  00:00:c8:01:01:07      30424783     1399540018       37448597     1722635462
  00:00:c8:01:01:08      30424783     1399540018       37448596     1722635416
  00:00:c8:01:01:09       8836796      406492616        8836795      406492570
  00:00:c8:01:01:0a      30424712     1399536752       37448521     1722631966
  00:00:c8:01:01:0b      30424715     1399536890       37448523     1722632058
  Number of MAC addresses : 21
```

show interfaces mac-database (All MAC Addresses on a Service)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
```

```
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      31016568     1426762128       38040381     1749857526
  00:00:c8:01:01:03      31016568     1426762128       38040382     1749857572
  00:00:c8:01:01:04      31016499     1426758954       38040306     1749854076
  00:00:c8:01:01:05      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:06      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:07      31016567     1426762082       38040380     1749857480
  00:00:c8:01:01:08      31016567     1426762082       38040379     1749857434
  00:00:c8:01:01:09       9428580      433714680        9428580      433714680
  00:00:c8:01:01:0a      31016496     1426758816       38040304     1749853984
  00:00:c8:01:01:0b      31016498     1426758908       38040307     1749854122
```

show interfaces mac-database mac-address

```
user@host> show interfaces mac-database xe-0/3/3 mac-address       00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback: None, Source
filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes     : 202324652
    Output bytes    : 202324560
    Input frames    : 4398362
    Output frames   : 4398360
  Policer statistics:
  Policer type      Discarded frames    Discarded bytes
  Output aggregate           3992386          183649756
```

Sample Output

show interfaces mc-ae

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links   : ae0
  Local Status   : active
  Peer Status    : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL         : Label Ethernet Interface
```

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the `show interfaces media` command for a SONET interface (with no level of output specified):

```
user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,   Loopback:
None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,   mpls: Not-
configured
  CHAP state: Not-configured
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
  SONET errors:
     BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
```

```
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

Sample Output

show interfaces policers

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
...
so-2/0/0        up    up
so-2/0/0.0      up    up   inet  so-2/0/0.0-in-policer so-2/0/0.0-out-policer
                           iso
so-2/1/0        up    down
...
```

show interfaces policers interface-name

```
user@host> show interfaces policers so-2/1/0
Interface       Admin Link Proto Input Policer         Output Policer
so-2/1/0        up    down
so-2/1/0.0      up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                           iso
                           inet6
```

Sample Output

show interfaces queue

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```
user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: class0
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
  Queue Buffer Usage:
    Reserved buffer      :            118750000 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
Queue: 1, Forwarding classes: class1
..
..
 Queue Buffer Usage:
    Reserved buffer      :                 9192 bytes
  Queue-depth bytes      :
    Current              :                    0
..
```

```
..
Queue: 3, Forwarding classes: class3
  Queued:
..
..
Queue Buffer Usage:
    Reserved buffer      :               6250000 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
```

Sample Output

show interfaces redundancy

```
user@host> show interfaces redundancy
Interface  State         Last change  Primary    Secondary  Current status
rsp0       Not present                sp-1/0/0   sp-0/2/0   both down
rsp1       On secondary  1d 23:56     sp-1/2/0   sp-0/3/0   primary down
rsp2       On primary    10:10:27     sp-1/3/0   sp-0/2/0   secondary down
rlsq0      On primary    00:06:24     lsq-0/3/0  lsq-1/0/0  both up
```

show interfaces redundancy (Aggregated Ethernet)

```
user@host> show interfaces redundancy
Interface  State            Last change  Primary      Secondary    Current status
rlsq0      On secondary     00:56:12     lsq-4/0/0    lsq-3/0/0    both up
ae0
ae1
ae2
ae3
ae4
```

show interfaces redundancy detail

```
user@host> show interfaces redundancy detail
Interface        : rlsq0
```

```
  State          : On primary
  Last change    : 00:45:47
  Primary        : lsq-0/2/0
  Secondary      : lsq-1/2/0
  Current status : both up
  Mode           : hot-standby

Interface        : rlsq0:0
  State          : On primary
  Last change    : 00:45:46
  Primary        : lsq-0/2/0:0
  Secondary      : lsq-1/2/0:0
  Current status : both up
  Mode           : warm-standby
```

Sample Output

show interfaces routing brief

```
user@host> show interfaces routing brief
Interface        State Addresses
so-5/0/3.0       Down  ISO   enabled
so-5/0/2.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.120
                       INET  enabled
so-5/0/1.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.130
                       INET  enabled
at-1/0/0.3       Up    CCC   enabled
at-1/0/0.2       Up    CCC   enabled
at-1/0/0.0       Up    ISO   enabled
                       INET  192.168.90.10
                       INET  enabled
lo0.0            Up    ISO   47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
                       ISO   enabled
                       INET  127.0.0.1
fxp1.0           Up
fxp0.0           Up    INET  192.168.6.90
```

show interfaces routing detail

```
user@host> show interfaces routing detail
so-5/0/3.0
  Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>
  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  ISO address (null)
    State: <Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>
  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  MPLS address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
  ISO address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  INET address 192.168.2.120
    State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
    Local address: 192.168.2.120
    Destination: 192.168.2.110/32
  INET address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...
```

Sample Output

show interfaces routing-instance all

```
user@host> show interfaces terse routing-instance all
Interface   Admin   Link   Proto    Local             Remote Instance
at-0/0/1       up      up     inet        10.0.0.1/24
ge-0/0/0.0  up      up     inet      192.168.4.28/24          sample-a
```

```
at-0/1/0.0  up      up     inet6     fe80::a:0:0:4/64         sample-b
so-0/0/0.0     up      up     inet        10.0.0.1/32
```

Sample Output

show interfaces snmp-index

```
user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
  Interface index: 149, SNMP ifIndex: 33
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,   Loopback:
None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : LOL, PLL, LOS
  SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P
```

Sample Output

show interfaces source-class all

```
user@host> show interfaces source-class all
Logical interface so-0/1/0.0
                                                Packets                Bytes
      Source class                 (packet-per-second)    (bits-per-second)
                             gold              1928095            161959980
                                  (                889) (            597762)
                           bronze                    0                    0
                                  (                  0) (                  0)
                           silver                    0                    0
                                  (                  0) (                  0)
  Logical interface so-0/1/3.0
                                                Packets                Bytes
      Source class                 (packet-per-second)    (bits-per-second)
```

```
                              gold                    0                    0
                                   (                  0) (                 0)
                            bronze                    0                    0
                                   (                  0) (                 0)
                            silver               116113              9753492
                                   (                939) (            631616)
```

Sample Output

show interfaces statistics (Fast Ethernet)

```
user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : None
  Active defects : None
  Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary, DCU, SCU-in
                                              Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                            silver1                   0                    0
                                   (                  0) (                 0)
                            silver2                   0                    0
                                   (                  0) (                 0)
                            silver3                   0                    0
                                   (                  0) (                 0)
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
```

```
        Destination: 10.27.245/24, Local: 10.27.245.2,
        Broadcast: 10.27.245.255
    Protocol iso, MTU: 1497
      Flags: Is-Primary
```

Sample Output

show interfaces switch-port

```
user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
           Speed: 100mbps, Auto-negotiation: Enabled
            Statistics:                          Receive         Transmit
             Total bytes                       28437086         21792250
             Total packets                       409145            88008
             Unicast packets                       9987            83817
             Multicast packets                   145002                0
             Broadcast packets                   254156             4191
             Multiple collisions                     23               10
             FIFO/CRC/Align errors                    0                0
             MAC pause frames                         0                0
             Oversized frames                         0
             Runt frames                              0
             Jabber frames                            0
             Fragment frames                          0
             Discarded frames                         0
           Autonegotiation information:
             Negotiation status: Complete
             Link partner:
                 Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link partner
Speed: 100 Mbps
             Local resolution:
                 Flow control: None, Remote fault: Link OK
```

Sample Output

show interfaces transport pm

```
user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current          Elapse time:900 Seconds
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  OTU-BBE                 0                800              No              No
  OTU-ES                  0                135              No              No
  OTU-SES                 0                90               No              No
  OTU-UAS                 427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  OTU-BBE                 0                800              No              No
  OTU-ES                  0                135              No              No
  OTU-SES                 0                90               No              No
  OTU-UAS                 0                90               No              No
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  ODU-BBE                 0                800              No              No
  ODU-ES                  0                135              No              No
  ODU-SES                 0                90               No              No
  ODU-UAS                 427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  ODU-BBE                 0                800              No              No
  ODU-ES                  0                135              No              No
  ODU-SES                 0                90               No              No
  ODU-UAS                 0                90               No              No
  FEC            Suspect Flag:False             Reason:None
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  FEC-CorrectedErr        2008544300       0                NA              NA
  FEC-UncorrectedWords    0                0                NA              NA
  BER            Suspect Flag:False             Reason:None
    PM                 MIN       MAX       AVG       THRESHOLD      TCA-ENABLED      TCA-RAISED
  BER                  3.6e-5    5.8e-5    3.6e-5     10.0e-3        No                Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current
  Suspect Flag:True             Reason:Object Disabled
    PM                            CURRENT   MIN       MAX      AVG       THRESHOLD      TCA-
```

```
ENABLED       TCA-RAISED
                                                                                  (MIN) (MAX)    (MIN)
(MAX)       (MIN) (MAX)
  Lane chromatic dispersion         0        0         0         0         0     0        NA
NA         NA     NA
  Lane differential group delay     0        0         0         0         0     0        NA
NA         NA     NA
    q Value                           120      120       120       120       0     0
NA    NA         NA     NA
  SNR                               28       28        29        28        0     0        NA
NA         NA     NA
  Tx output power(0.01dBm)          -5000    -5000     -5000     -5000     -300  -100     No
No         No     No
  Rx input power(0.01dBm)           -3642    -3665     -3626     -3637     -1800 -500     No
No         No     No
  Module temperature(Celsius)       46       46        46        46        -5    75       No
No         No     No
  Tx laser bias current(0.1mA)      0        0         0         0         0     0        NA
NA         NA     NA
  Rx laser bias current(0.1mA)      1270     1270      1270      1270      0     0        NA
NA         NA     NA
  Carrier frequency offset(MHz)     -186     -186      -186      -186      -5000 5000     No
No         No     No
```

Sample Output

show security zones

```
user@host>               show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
```

```
      fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

Release Information

Command modified in Junos OS Release 9.5.

show interfaces statistics

IN THIS SECTION

- Syntax | 21443
- Description | 21443
- Options | 21443
- Required Privilege Level | 21443
- Output Fields | 21443
- Sample Output | 21444
- Release Information | 21460

Syntax

```
show interfaces statistics interface-name
<satellite-device [device-alias-name |all ]>
<detail>
```

Description

Display static interface statistics, such as errors.

> **NOTE**: When the **show interfaces statistics** command is executed on an interface that is configured on T4000 Type 5 FPC, the *IPv6 transit statistics* field displays:
>
> - Total statistics (sum of transit and local statistics) at the physical interface level
> - Transit statistics at the logical interface level

Options

interface-name	Name of an interface.
satellite-device [*device-alias-name* | all]	(Junos Fusion only) (Optional) Display interface statistics for interfaces on the specified satellite device in the Junos Fusion, or on all satellite devices in the Junos Fusion.

> **NOTE**: In a Junos Fusion Enterprise, logical interface statistics are not synced across aggregation devices in a dual-aggregation device topology.

detail	(Optional) Display detailed output.

Required Privilege Level

view

Output Fields

Output from both the `show interfaces interface-name detail` and the `show interfaces interface-name extensive` commands include all the information displayed in the output from the `show interfaces statistics`

command. For more information, see the particular interface type in which you are interested. For information about destination class and source class statistics, see the “Destination Class Field” section and the “Source Class Field” section under Common Output Fields Description. For information about the input errors and output errors, see Fast Ethernet and Gigabit Ethernet Counters.

Sample Output

show interfaces statistics (Fast Ethernet)

```
user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5E:00:53:dc, Hardware address: 00:00:5E:00:53:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : None
  Active defects : None
  Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary, DCU, SCU-in
                                              Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                        silver1                    0                    0
                                (                  0) (                 0)
                        silver2                    0                    0
                                (                  0) (                 0)
                        silver3                    0                    0
                                (                  0) (                 0)
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10.27.245/24, Local: 10.27.245.2,
        Broadcast: 10.27.245.255
```

```
    Protocol iso, MTU: 1497
      Flags: Is-Primary
```

show interfaces statistics (Gigabit Ethernet PIC—Egress)

```
user@host> show interfaces ge-5/2/0 statistics detail
Physical interface: ge-5/2/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 519, Generation: 149
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote fault:
Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:74, Hardware address: 00:00:5E:00:53:74
  Last flapped   : 2009-11-11 11:24:00 PST (09:23:08 ago)
  Statistics last cleared: 2009-11-11 17:50:58 PST (02:56:10 ago)
  Traffic statistics:
   Input  bytes  :               271524                    0 bps
   Output bytes  :             37769598                  352 bps
   Input  packets:                 3664                    0 pps
   Output packets:               885790                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :            16681118
    Input  packets:                   0
    Output packets:              362633
   Multicast statistics:
    IPV4 multicast statistics:
     Input  bytes  :               112048                    0 bps
     Output bytes  :             20779920                    0 bps
     Input  packets:                 1801                    0 pps
     Output packets:               519498                    0 pps
    IPV6 multicast statistics:
     Input  bytes  :               156500                    0 bps
     Output bytes  :             16681118                    0 bps
     Input  packets:                 1818                    0 pps
     Output packets:               362633                    0 pps
```

```
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort               882558               882558                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                3232                 3232                    0
  Active alarms  : None
  Active defects : None

  Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :               271524
     Output bytes  :             37769598
     Input  packets:                 3664
     Output packets:               885790
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :            16681118
      Input  packets:                   0
      Output packets:              362633
    Local statistics:
     Input  bytes  :               271524
     Output bytes  :               308560
     Input  packets:                 3664
     Output packets:                 3659
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :             37461038                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:               882131                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
```

```
      Output bytes  :            16681118
      Input  packets:                   0
      Output packets:              362633
    Multicast statistics:
     IPV4 multicast statistics:
      Input  bytes  :               112048                    0 bps
      Output bytes  :             20779920                    0 bps
      Input  packets:                 1801                    0 pps
      Output packets:               519498                    0 pps
     IPV6 multicast statistics:
      Input  bytes  :               156500                    0 bps
      Output bytes  :             16681118                    0 bps
      Input  packets:                 1818                    0 pps
      Output packets:               362633                    0 pps
    Protocol inet, MTU: 1500, Generation: 151, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.40.40.0/30, Local: 10.40.40.2, Broadcast: 10.40.40.3, Generation: 167
    Protocol inet6, MTU: 1500, Generation: 152, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::10.40.40.0/126, Local: ::10.40.40.2
    Generation: 169
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21d:b5ff:fe61:d974
    Protocol multiservice, MTU: Unlimited, Generation: 171
    Generation: 153, Route table: 0
      Policer: Input: __default_arp_policer__
```

show interfaces statistics detail (Aggregated Ethernet)

```
user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 186, SNMP ifIndex: 111, Generation: 187
  Link-level type: Ethernet, MTU: 1514, Speed: 2000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5E:0053:f0, Hardware address: 00:00:5E:00:53:f0
  Last flapped   : Never
  Statistics last cleared: 2006-12-23 03:04:16 PST (01:16:24 ago)
```

```
  Traffic statistics:
   Input  bytes  :                28544                    0 bps
   Output bytes  :                39770                    0 bps
   Input  packets:                  508                    0 pps
   Output packets:                  509                    0 pps
   Input  bytes  :           IPv6 28544
   Output bytes  :               IPv6 0
   Input  packets:             IPv6 508
   Output packets:               IPv6 0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface ae0.0 (Index 67) (SNMP ifIndex 139) (Generation 145)
    Flags: SNMP-Traps Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :           508          0         28544            0
        Output:           509          0         35698            0
    Link:
      ge-3/3/8.0
        Input :           508          0         28544            0
        Output:             0          0             0            0
      ge-3/3/9.0
        Input :             0          0             0            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-3/3/8.0                0           0            0            0
      ge-3/3/9.0                0           0            0            0
  Egress queues: 8 supported, 8 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
    Protocol inet, MTU: 1500, Generation: 166, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
        Generation: 159
```

```
    Protocol inet6, MTU: 1500, Generation: 163, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::206:5bff:fe05:c321,
        Broadcast: Unspecified, Generation: 161
```

show interfaces statistics detail (Aggregated Ethernet—Ingress)

```
user@host> show interfaces statistics detail ae0 | no-more
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 504, Generation: 278
  Link-level type: Ethernet, MTU: 1514, Speed: 1Gbps, BPDU Error: None, MAC-REWRITE Error: None,
Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth
needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5E:00:53:f0, Hardware address: 00:00:5E:00:53:f0
  Last flapped   : 2009-11-09 03:30:23 PST (00:01:28 ago)
  Statistics last cleared: 2009-11-09 03:26:18 PST (00:05:33 ago)
  Traffic statistics:
   Input  bytes  :            544009602             54761856 bps
   Output bytes  :                 3396                    0 bps
   Input  packets:             11826292               148809 pps
   Output packets:                   42                    0 pps
   IPv6 transit statistics:
    Input  bytes  :           350818604
    Output bytes  :                   0
    Input  packets:             7626488
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
```

```
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                   21                   21                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                 451                  451                    0

  Logical interface ae0.0 (Index 70) (SNMP ifIndex 574) (Generation 177)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :      11826292     148809     544009602     54761856
        Output:            42          0          3396            0
    Link:
      ge-5/2/0.0
        Input :      11826292     148809     544009602     54761856
        Output:            42          0          3396            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-5/2/0.0                 0           0            0            0
    Protocol inet, MTU: 1500, Generation: 236, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.30.30.0/30, Local: 10.30.30.2, Broadcast: 10.30.30.3, Generation: 310
    Protocol inet6, MTU: 1500, Generation: 237, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::10.30.30.0/126, Local: ::10.30.30.2
    Generation: 312
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21d:b5ff:fe61:dbf0
    Protocol multiservice, MTU: Unlimited, Generation: 314
    Generation: 238, Route table: 0
      Policer: Input: __default_arp_policer__
```

show interfaces statistics detail (Aggregated Ethernet—Egress)

```
user@host> show interfaces statistics detail ae0 | no-more
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 501, Generation: 319
  Link-level type: Ethernet, MTU: 1514, Speed: 1Gbps, BPDU Error: None, MAC-REWRITE Error: None,
Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth
needed: 0
```

```
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5E:00:53:f0, Hardware address: 00:00:5E:00:53:f0
  Last flapped   : 2009-11-09 03:30:24 PST (00:02:42 ago)
  Statistics last cleared: 2009-11-09 03:26:42 PST (00:06:24 ago)
  Traffic statistics:
   Input  bytes  :                  440                    0 bps
   Output bytes  :           1047338120             54635848 bps
   Input  packets:                    7                    0 pps
   Output packets:             22768200               148466 pps
   IPv6 transit statistics:
    Input  bytes  :                 288
    Output bytes  :           723202616
    Input  packets:                   4
    Output packets:            15721796
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort            201985796            201985796                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                  65                   65                    0

  Logical interface ae0.0 (Index 72) (SNMP ifIndex 505) (Generation 204)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             7          0           440            0
        Output:      22768200     148466    1047338120     54635848
    Link:
      ge-2/1/6.0
        Input :             7          0           440            0
        Output:      22768200     148466    1047338120     54635848
```

```
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-2/1/6.0                 0           0            0            0
    Protocol inet, MTU: 1500, Generation: 291, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.30.30.0/30, Local: 10.30.30.1, Broadcast: 10.30.30.3, Generation: 420
    Protocol inet6, MTU: 1500, Generation: 292, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::/26, Local: ::10.30.30.1
    Generation: 422
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21f:12ff:fec2:37f0
    Protocol multiservice, MTU: Unlimited, Generation: 424
    Generation: 293, Route table: 0
      Policer: Input: __default_arp_policer__
```

show interfaces statistics (SONET/SDH)

```
user@host> show interfaces statistics detail so-3/0/0 | no-more
Physical interface: so-3/0/0, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 538, Generation: 283
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC192, Loopback: None,
FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 13 (last seen 00:00:04 ago)
    Output: 14 (last sent 00:00:02 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Not-configured, mpls: Not-configured
  CHAP state: Closed
  PAP state: Closed
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2009-11-09 02:52:34 PST (01:12:39 ago)
  Statistics last cleared: 2009-11-09 03:58:54 PST (00:06:19 ago)
  Traffic statistics:
   Input  bytes  :           2559160294             54761720 bps
   Output bytes  :                10640                   48 bps
   Input  packets:             55633975               148809 pps
```

```
   Output packets:                  216                     0 pps
   IPv6 transit statistics:
    Input  bytes  :           647922328
    Output bytes  :                   0
    Input  packets:            14085269
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Bucket drops: 0, Policed
discards: 0, L3 incompletes: 0,
    L2 channel errors: 0, L2 mismatch timeouts: 0, HS link CRC errors: 0, HS link FIFO
overflows: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, HS link FIFO underflows: 0,
MTU errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    4                    4                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                 213                  213                    0
  SONET alarms   : None
  SONET defects  : None

  Logical interface so-3/0/0.0 (Index 72) (SNMP ifIndex 578) (Generation 182)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Protocol inet, MTU: 4470, Generation: 244, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.30.30.0/30, Local: 10.30.30.2, Broadcast: 10.30.30.3, Generation: 322
    Protocol inet6, MTU: 4470, Generation: 245, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::10.30.30.0/126, Local: ::10.30.30.2
    Generation: 324
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe61:9264
    Generation: 326
```

show interfaces statistics (Aggregated SONET/SDH—Ingress)

```
user@host> show interfaces statistics detail as0 | no-more
Physical interface: as0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 534, Generation: 282
```

```
  Link-level type: PPP, MTU: 4474, Speed: OC192, Minimum links needed: 1, Minimum bandwidth
needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Last flapped   : 2009-11-09 03:45:53 PST (00:09:38 ago)
  Statistics last cleared: 2009-11-09 03:48:17 PST (00:07:14 ago)
  Traffic statistics:
   Input  bytes  :           2969786332             54761688 bps
   Output bytes  :                11601                    0 bps
   Input  packets:             64560636               148808 pps
   Output packets:                  225                    0 pps
   IPv6 transit statistics:
    Input  bytes  :          2086013152
    Output bytes  :                   0
    Input  packets:            45348114
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    3                    3                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                 222                  222                    0

  Logical interface as0.0 (Index 71) (SNMP ifIndex 576) (Generation 179)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :      64560550     148808    2969785300     54761688
        Output:           139          0         10344            0
    Link:
      so-3/0/0.0
        Input :      64560550     148808    2969785300     54761688
        Output:           139          0         10344            0
    Protocol inet, MTU: 4470, Generation: 240, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.30.30.0/30, Local: 10.30.30.2, Broadcast: 10.30.30.3, Generation: 316
```

```
    Protocol inet6, MTU: 4470, Generation: 241, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::10.30.30.0/126, Local: ::10.30.30.2
    Generation: 318
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe61:9264
    Generation: 320
```

show interfaces statistics (Aggregated SONET/SDH—Egress)

```
user@host> show interfaces statistics detail as0 | no-more
Physical interface: as0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 565, Generation: 323
  Link-level type: PPP, MTU: 4474, Speed: OC192, Minimum links needed: 1, Minimum bandwidth
needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Last flapped   : 2009-11-09 03:43:37 PST (00:12:48 ago)
  Statistics last cleared: 2009-11-09 03:48:54 PST (00:07:31 ago)
  Traffic statistics:
   Input  bytes  :                11198                  392 bps
   Output bytes  :           3101452132             54783448 bps
   Input  packets:                  234                    0 pps
   Output packets:             67422937               148868 pps
   IPv6 transit statistics:
    Input  bytes  :                5780
    Output bytes  :          2171015678
    Input  packets:                  72
    Output packets:            47195993
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort             67422830             67422830                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
```

```
    3 network-cont                  90                  90                    0

  Logical interface as0.0 (Index 71) (SNMP ifIndex 548) (Generation 206)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :           144          0         10118          392
        Output:      67422847     148868    3101450962     54783448
    Link:
      so-0/1/0.0
        Input :           144          0         10118          392
        Output:      67422847     148868    3101450962     54783448
    Protocol inet, MTU: 4470, Generation: 295, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.30.30.0/30, Local: 10.30.30.1, Broadcast: 10.30.30.3, Generation: 426
    Protocol inet6, MTU: 4470, Generation: 296, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::/26, Local: ::10.30.30.1
    Generation: 428
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe63:1d0a
    Generation: 429
```

show interfaces statistics (MX Series Routers)

```
user@host> show interfaces xe-0/0/0 statistics
Physical interface: xe-0/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 592
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5E:00:53:f0, Hardware address: 00:00:5E:00:53:f0
  Last flapped   : 2013-10-26 03:20:40 test (2w3d 03:29 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
```

```
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                           109
    Errored blocks                       109
  Interface transmit statistics: Disabled
```

show interfaces statistics (MX Series Routers; With RPF Check Detail of Ethernet Interface)

```
user@host> show interfaces ge-0/0/0 extensive
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 149, SNMP ifIndex: 527, Generation: 152
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
None, Loop Detect PDU Error: None,
  Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface Specific flags: RPF_statistics
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 56:68:ad:c0:1f:72, Hardware address: 56:68:ad:c0:1f:72
  Last flapped   : 2021-02-12 17:21:37 IST (3d 15:55 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
IPv6 transit statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
```

```
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
 RPF counters statistics:
   Rpf IPv4 bytes   :                     0
   Rpf IPv4 packets :                   0
   Rpf IPv6 bytes   :                     0
   Rpf IPv6 packets :                   0
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
```

show interfaces statistics (MX Series Routers: Dynamic Interfaces with RPF Check Detail)

```
user@host> show interfaces statistics pp0.3221225475 detail
  Logical interface pp0.3221225475(Index 536870921)(SNMP ifIndex 200000009) (Generation 6)
    Flags: Up Point-To-Point Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 1,
      Session AC name: B, Remote MAC address:00:00:5E:00:53:01,
      Underlying interface: xe-1/0/0.3221225474 (Index 536870919)
      Ignore End-Of-List tag: Disable
    Bandwidth: 0
    Traffic statistics:
     Input  bytes  :                   34
     Output bytes  :                    0
     Input  packets:                    1
     Output packets:                    1
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
```

```
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                   34                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    1                    0 pps
     Output packets:                    1                    0 pps
  Keepalive settings: Interval 30 seconds, Up-count 3, Down-count 3
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-configured
  CHAP state: Success
  PAP state: Closed
    Protocol inet, MTU: 1492
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
    Generation: 0, Route table: 0
      Flags: uRPF, Unnumbered
      RPF Failures: Packets: 0, Bytes: 0
      Donor interface: lo0.0 (Index 320)
      Input Filters: upstrm1-inet-pp0.3221225475-in
      Output Filters: dwnstrm1-inet-pp0.3221225475-out
      Addresses, Flags: Is-Primary
        Destination: Unspecified, Local: 10.255.96.19, Broadcast: Unspecified, Generation: 0
```

show interfaces statistics (PTX Series Packet Transport Routers)

```
user@host> show interfaces statistics em0
 Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Last flapped   : Never
      Statistics last cleared: Never
    Input packets : 212620
    Output packets: 71
      Input errors: 0, Output errors: 0

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 212590
```

```
    Output packets: 71
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 192.168.3/24, Local: 192.168.3.30,
        Broadcast: 192.168.3.255
```

show interfaces statistics (ACX Series routers)

```
user@host> show interfaces statistics ge-0/1/7
Physical interface: ge-0/1/7, Enabled, Physical link is Down
  Interface index: 151, SNMP ifIndex: 524
  Link-level type: Ethernet, Media type: Copper, MTU: 1514, Link-mode: Full-duplex, Speed:
1000mbps, BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote fault:
Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5E:00:53:a3, Hardware address: 00:00:5E:00:53:a3
  Last flapped   : 2012-05-11 04:25:28 PDT (2d 20:23 ago)
  Statistics last cleared: 2012-05-13 23:07:23 PDT (01:41:25 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 12.2 for ACX Series Routers.

`satellite-device` option introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

clear interfaces statistics

show interfaces statistics fabric

IN THIS SECTION

- Syntax | 21461
- Description | 21461
- Options | 21461
- Required Privilege Level | 21462
- Output Fields | 21462
- Sample Output | 21471
- Release Information | 21488

Syntax

```
show interfaces statistics fabric
<brief | detail | terse>
<descriptions>
<interface-name>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
```

Description

Display status information about the specified fabric interface.

Options

brief | detail | terse (Optional) Display the specified level of output.

descriptions	(Optional) Display interface description strings.
interface-name	(QFabric systems only) The interface name is either the serial number or the alias of the QFabric switch component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.
media	(Optional) Display media-specific information about network interfaces.
routing-instance (all | *instance-name*)	(Optional) Display all routing instances or the name of an individual routing instance.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.

Required Privilege Level

view

Output Fields

Table 691 on page 21462 lists the output fields for the `show interfaces statistics fabric` command. Output fields are listed in the approximate order in which they appear.

Table 691: show interfaces statistics fabric Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail none
SNMP ifIndex	SNMP index number for the physical interface.	detail none

Table 691: show interfaces statistics fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Clocking	Reference clock source.	detail
Speed	Speed at which the interface is running.	All levels
Duplex	Duplex mode of the interface, either Full-Duplex or Half-Duplex.	All levels
MAC-REWRITE Error	Specifies if the encapsulation of the packet has been changed.	none
BPDU Error	Specifies if a BPDU has been received on a blocked interface.	none
Loopback	Loopback status: Enabled or Disabled. If loopback is enabled, type of loopback: Local or Remote.	All levels
Source filtering	Source filtering status: Enabled or Disabled.	All levels
Flow control	Flow control status: Enabled or Disabled. This field is only displayed if asymmetric flow control is not configured.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
CoS queues	Number of CoS queues configured.	detail none
Hold-Times	Current interface hold-time up and hold-time down, in milliseconds.	detail

Table 691: show interfaces statistics fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Current address	Configured MAC address.	detail none
Hardware address	Hardware MAC address.	detail none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour.minute.second.timezone*** **(*hour.minute.second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).**	detail none
Statistics last cleared	Date, time, and how long ago the statistics for the interface were cleared. The format is **Statistics last cleared:** ***year-month-day hour.minute.second.timezone*** **(*hour.minute.second* ago)**. For example, 2010-05-17 07:51:28 PDT (00:04:33 ago).	detail
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail

Table 691: show interfaces statistics fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IPv6 transit statistics	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface: • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail

Table 691: show interfaces statistics fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • Errors—Sum of the incoming frame terminated and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops.	detail none

Table 691: show interfaces statistics fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame terminated and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the fabric interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops.	detail none
Egress queues	Total number of egress queues supported on the specified interface.	detail

Table 691: show interfaces statistics fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Queue counters	CoS queue number and its associated user-configured forwarding class name. • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism.	detail
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail none
Flags	Information about the logical interface. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.	All levels
Input packets	Number of packets received on the interface.	detail none
Output packets	Number of packets transmitted on the interface.	detail none

Table 691: show interfaces statistics fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input packets	Number of packets received on the interface.	detail none
Output packets	Number of packets transmitted on the interface.	detail none
Encapsulation	Encapsulation method used on the logical interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail
Local statistics	• Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail

Table 691: show interfaces statistics fabric Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Transit statistics	• Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	detail
Addresses, Flags	Information about the address flags.	detail none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	brief
MTU	Maximum transmission unit size on the physical interface.	All levels
Destination	IP address of the remote side of the connection.	detail none
Local	IP address of the logical interface.	detail none
Broadcast	Broadcast address of the logical interlace.	detail none
Generation	Unique number for use by Juniper Networks technical support only.	detail
Route table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail none

Sample Output

show interfaces statistics fabric

```
user@switch> show interfaces statistic fabric
Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Down
  Interface index: 49174, SNMP ifIndex: 1208484473
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:00:00:00:00:00, Hardware address: 00:00:00:00:00:00
  Last flapped   : 2012-11-27 20:30:30 UTC (01:55:19 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:fte-0/0/0.32768 (Index 83) (SNMP ifIndex 1208484474)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-0/0/4, Enabled, Physical link is Down
  Interface index: 49175, SNMP ifIndex: 1208484363
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:00:00:00:00:04, Hardware address: 00:00:00:00:00:04
  Last flapped   : 2012-11-27 20:30:30 UTC (01:55:20 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:fte-0/0/4.32768 (Index 84) (SNMP ifIndex 1208484364)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
```

```
    Output packets: 0
    Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Down
  Interface index: 49176, SNMP ifIndex: 1208484367
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:00:00:00:00:06, Hardware address: 00:00:00:00:00:06
  Last flapped   : 2012-11-27 20:30:30 UTC (01:55:20 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:fte-0/0/6.32768 (Index 85) (SNMP ifIndex 1208484368)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Down
  Interface index: 49177, SNMP ifIndex: 1208484479
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:00:00:00:00:0d, Hardware address: 00:00:00:00:00:0d
  Last flapped   : 2012-11-27 20:30:30 UTC (01:55:20 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:fte-0/0/13.32768 (Index 86) (SNMP ifIndex 1208484480)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch, MTU: 0
```

```
Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Down
  Interface index: 49178, SNMP ifIndex: 1208484475
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:00:00:00:00:0f, Hardware address: 00:00:00:00:00:0f
  Last flapped   : 2012-11-27 20:30:30 UTC (01:55:20 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:fte-0/0/15.32768 (Index 87) (SNMP ifIndex 1208484476)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Down
  Interface index: 49211, SNMP ifIndex: 1208484377
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:00:00:00:00:02, Hardware address: 00:00:00:00:00:02
  Last flapped   : 2012-11-27 20:30:47 UTC (01:55:03 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:fte-1/0/2.32768 (Index 120) (SNMP ifIndex 1208484378)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Down
  Interface index: 49212, SNMP ifIndex: 1208484365
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
```

```
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:00:00:00:00:07, Hardware address: 00:00:00:00:00:07
  Last flapped   : 2012-11-27 20:30:47 UTC (01:55:04 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:fte-1/0/7.32768 (Index 121) (SNMP ifIndex 1208484366)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Down
  Interface index: 49213, SNMP ifIndex: 1208484625
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:00:00:00:00:0a, Hardware address: 00:00:00:00:00:0a
  Last flapped   : 2012-11-27 20:30:47 UTC (01:55:04 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:fte-1/0/10.32768 (Index 122) (SNMP ifIndex 1208484626)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:bme0, Enabled, Physical link is Up
  Interface index: 64, SNMP ifIndex: 1208483877
  Type: Ethernet, Link-level type: Ethernet, MTU: 1500
  Device flags   : Present Running
  Current address: 02:00:00:00:40:06, Hardware address: 02:00:00:00:40:01
  Last flapped   : Never
```

```
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 26730
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:bme0.0 (Index 4) (SNMP ifIndex 1208484065)
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    Input packets : 2715
    Output packets: 18
    Protocol inet, MTU: 1482
        Local: 128.0.32.0

  Logical interface IC-WS001:bme0.1 (Index 5) (SNMP ifIndex 1208484091)
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 999
    Protocol inet, MTU: 1482
      Addresses
        Destination: 128/2, Local: 128.0.0.6, Broadcast: 191.255.255.255
        Destination: 128/2, Local: 128.0.130.2, Broadcast: 191.255.255.255

  Logical interface IC-WS001:bme0.2 (Index 6) (SNMP ifIndex 1208484092)
    Flags: Encapsulation: ENET2
    Input packets : 180408
    Output packets: 23051
    Protocol inet, MTU: 1486
        Destination: 128/8, Local: 128.0.0.6, Broadcast: 128.255.255.255
        Destination: 128/8, Local: 128.0.130.2, Broadcast: 128.255.255.255
        Destination: 169.254/16, Local: 169.254.128.6, Broadcast: 169.254.255.255
        Destination: 169.254/16, Local: 169.254.193.1, Broadcast: 169.254.255.255

Physical interface: IC-WS001:bme1, Enabled, Physical link is Up
  Interface index: 49156, SNMP ifIndex: 1208483949
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:0d:0c:0f:00:03, Hardware address: 00:0d:0c:0f:00:03
  Last flapped   : 1970-01-01 00:00:01 UTC (2238w5d 22:25 ago)
  Statistics last cleared: Never
    Input packets : 168885
    Output packets: 184712
  Input errors: 0, Output errors: 0
```

```
  Logical interface IC-WS001:bme1.0 (Index 3) (SNMP ifIndex 1208483950)
    Flags: Encapsulation: ENET2
    Input packets : 168885
    Output packets: 184712
    Protocol inet, MTU: 1500
        Destination: 128/2, Local: 128.0.0.1, Broadcast: 191.255.255.255
        Destination: 128/2, Local: 128.0.0.5, Broadcast: 191.255.255.255
        Destination: 128/2, Local: 128.0.0.16, Broadcast: 191.255.255.255
        Destination: 128/2, Local: 128.0.0.17, Broadcast: 191.255.255.255
        Destination: 128/2, Local: 128.0.0.24, Broadcast: 191.255.255.255
        Destination: 128/2, Local: 128.0.0.25, Broadcast: 191.255.255.255
        Destination: 128/2, Local: 128.0.0.26, Broadcast: 191.255.255.255
        Destination: 128/2, Local: 128.0.0.28, Broadcast: 191.255.255.255
        Destination: 128/2, Local: 128.0.0.29, Broadcast: 191.255.255.255
        Destination: 128/2, Local: 128.0.0.31, Broadcast: 191.255.255.255
    Protocol tnp, MTU: 1500
        Local: 0x5

Physical interface: IC-WS001:dcfabric, Enabled, Physical link is Up
  Interface index: 27, SNMP ifIndex: 1208484093
  Type: Ethernet, Link-level type: Ethernet, MTU: 1572
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Current address: 00:0b:ca:fe:00:01, Hardware address: 00:0b:ca:fe:00:01
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 0
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:dcfabric.0 (Index 64) (SNMP ifIndex 1208484094)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1558
    Protocol mpls, MTU: 1546, Maximum labels: 3
    Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:pme0, Enabled, Physical link is Up
  Interface index: 66, SNMP ifIndex: 1208484104
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps
  Device flags   : Present Running
```

```
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:23:9c:f1:a2:e6, Hardware address: 00:23:9c:f1:a2:e6
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 1007238
    Output packets: 63383
  Input errors: 0, Output errors: 0

Physical interface: IC-WS001:pme1, Enabled, Physical link is Up
  Interface index: 67, SNMP ifIndex: 1208484105
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:23:9c:f1:a2:e7, Hardware address: 00:23:9c:f1:a2:e7
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 1007118
    Output packets: 55381
  Input errors: 0, Output errors: 0

Physical interface: IC-WS001:pme2, Enabled, Physical link is Down
  Interface index: 68, SNMP ifIndex: 1208484106
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:23:9c:f1:a2:e8, Hardware address: 00:23:9c:f1:a2:e8
  Last flapped   : 2012-11-27 02:52:03 UTC (19:33:54 ago)
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 0
  Input errors: 0, Output errors: 0

Physical interface: IC-WS001:pme3, Enabled, Physical link is Down
  Interface index: 69, SNMP ifIndex: 1208484107
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:23:9c:f1:a2:e9, Hardware address: 00:23:9c:f1:a2:e9
  Last flapped   : 2012-11-27 02:52:03 UTC (19:33:54 ago)
```

```
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 0
  Input errors: 0, Output errors: 0

Physical interface: IC-WS001:vcp0, Enabled, Physical link is Up
  Interface index: 74, SNMP ifIndex: 1208484372
  Type: Ethernet, Link-level type: 70, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Link type      : Full-Duplex
  Current address: 00:23:9c:f1:a2:e3, Hardware address: 00:23:9c:f1:a2:e3
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 121842
    Output packets: 3548
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:vcp0.32769 (Index 11) (SNMP ifIndex 1208484376)
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    Input packets : 13044
    Output packets: 3548

Physical interface: IC-WS001:vcp1, Enabled, Physical link is Up
  Interface index: 70, SNMP ifIndex: 1208484108
  Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps
  Device flags   : Present Running
  Link type      : Full-Duplex
  Current address: 00:23:9c:f1:a2:e6, Hardware address: 00:23:9c:f1:a2:e6
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 767413
    Output packets: 46503
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:vcp1.32768 (Index 7) (SNMP ifIndex 1208484109)
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    Input packets : 735889
    Output packets: 46503

Physical interface: IC-WS001:vcp2, Enabled, Physical link is Up
  Interface index: 71, SNMP ifIndex: 1208484369
  Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps
  Device flags   : Present Running
```

```
  Link type      : Full-Duplex
  Current address: 00:23:9c:f1:a2:e7, Hardware address: 00:23:9c:f1:a2:e7
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 831710
    Output packets: 44548
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:vcp2.32768 (Index 8) (SNMP ifIndex 1208484373)
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    Input packets : 737844
    Output packets: 44548

Physical interface: IC-WS001:vcp3, Enabled, Physical link is Down
  Interface index: 72, SNMP ifIndex: 1208484370
  Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps
  Device flags   : Present Running
  Link type      : Full-Duplex
  Current address: 00:23:9c:f1:a2:e8, Hardware address: 00:23:9c:f1:a2:e8
  Last flapped   : 2012-11-27 20:31:36 UTC (01:54:21 ago)
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 0
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:vcp3.32768 (Index 9) (SNMP ifIndex 1208484374)
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0

Physical interface: IC-WS001:vcp4, Enabled, Physical link is Down
  Interface index: 73, SNMP ifIndex: 1208484371
  Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps
  Device flags   : Present Running
  Link type      : Full-Duplex
  Current address: 00:23:9c:f1:a2:e9, Hardware address: 00:23:9c:f1:a2:e9
  Last flapped   : 2012-11-27 20:31:36 UTC (01:54:21 ago)
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 0
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:vcp4.32768 (Index 10) (SNMP ifIndex 1208484375)
```

```
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
```

show interfaces statistics fabric brief

```
user@switch> show interfaces statistics fabric brief
Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-0/0/0.32768
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-0/0/4, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-0/0/4.32768
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-0/0/6.32768
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
```

```
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-0/0/13.32768
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-0/0/15.32768
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-1/0/2.32768
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-1/0/7.32768
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface IC-WS001:fte-1/0/10.32768
```

```
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:bme0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running

  Logical interface IC-WS001:bme0.0
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.32.0       --> 0/0

  Logical interface IC-WS001:bme0.1
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.6/2
          128.0.130.2/2

  Logical interface IC-WS001:bme0.2
    Flags: Encapsulation: ENET2
    inet  128.0.0.6/8
          128.0.130.2/8
          169.254.128.6/16
          169.254.193.1/16

Physical interface: IC-WS001:bme1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface IC-WS001:bme1.0
    Flags: Encapsulation: ENET2
    inet  128.0.0.1/2
          128.0.0.5/2
          128.0.0.16/2
          128.0.0.17/2
          128.0.0.24/2
          128.0.0.25/2
          128.0.0.26/2
          128.0.0.28/2
          128.0.0.29/2
          128.0.0.31/2
    tnp   0x5

Physical interface: IC-WS001:dcfabric, Enabled, Physical link is Up
```

```
  Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface IC-WS001:dcfabric.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet
    mpls
    eth-switch

Physical interface: IC-WS001:pme0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme2, Enabled, Physical link is Down
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme3, Enabled, Physical link is Down
  Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

Physical interface: IC-WS001:vcp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface IC-WS001:vcp0.32769
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface IC-WS001:vcp1.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2
```

```
Physical interface: IC-WS001:vcp2, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface IC-WS001:vcp2.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp3, Enabled, Physical link is Down
  Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface IC-WS001:vcp3.32768
    Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp4, Enabled, Physical link is Down
  Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running

  Logical interface IC-WS001:vcp4.32768
    Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2
```

show interfaces statistics fabric detail

```
user@switch> show interfaces statistics fabric detail
 show interfaces statistics fabric detail
Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Down
  Interface index: 49174, SNMP ifIndex: 1208484473, Generation: 153
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:00:00:00:00, Hardware address: 00:00:00:00:00:00
  Last flapped   : 2012-11-27 20:30:30 UTC (02:04:59 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
```

```
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 12 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 fabric_fcset                   0                    0                    0
    1 fabric_fcset                   0                    0                    0
    2 fabric_fcset                   0                    0                    0
    3 fabric_fcset                   0                    0                    0
    4 fabric_fcset                   0                    0                    0
    5 fabric_fcset                   0                    0                    0
    6 fabric_fcset                   0                    0                    0
    7 fabric_fcset                   0                    0                    0
    8 fabric_fcset                   0                    0                    0
    9 fabric_fcset                   0                    0                    0
    10 fabric_fcset                   0                    0                    0
    11 fabric_fcset                   0                    0                    0

  Logical interface IC-WS001:fte-0/0/0.32768 (Index 83) (SNMP ifIndex 1208484474) (Generation
148)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
```

```
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol eth-switch, MTU: 0, Generation: 176, Route table: 0
```

show interfaces statistics fabric terse

```
user@switch> show interfaces statistics fabric terse
Interface               Admin Link Proto    Local                 Remote
IC-WS001:fte-0/0/0      up    down
IC-WS001:fte-0/0/0.32768 up   down eth-switch
IC-WS001:fte-0/0/4      up    down
IC-WS001:fte-0/0/4.32768 up   down eth-switch
IC-WS001:fte-0/0/6      up    down
IC-WS001:fte-0/0/6.32768 up   down eth-switch
IC-WS001:fte-0/0/13     up    down
IC-WS001:fte-0/0/13.32768 up  down eth-switch
IC-WS001:fte-0/0/15     up    down
IC-WS001:fte-0/0/15.32768 up  down eth-switch
IC-WS001:fte-1/0/2      up    down
IC-WS001:fte-1/0/2.32768 up   down eth-switch
IC-WS001:fte-1/0/7      up    down
IC-WS001:fte-1/0/7.32768 up   down eth-switch
IC-WS001:fte-1/0/10     up    down
IC-WS001:fte-1/0/10.32768 up  down eth-switch
IC-WS001:bme0           up    up
IC-WS001:bme0.0         up    up   inet     128.0.32.0          --> 0/0
IC-WS001:bme0.1         up    up   inet     128.0.0.6/2
                                            128.0.130.2/2
IC-WS001:bme0.2         up    up   inet     128.0.0.6/8
                                            128.0.130.2/8
                                            169.254.128.6/16
                                            169.254.193.1/16
IC-WS001:bme1           up    up
IC-WS001:bme1.0         up    up   inet     128.0.0.1/2
                                            128.0.0.5/2
                                            128.0.0.16/2
                                            128.0.0.17/2
                                            128.0.0.24/2
                                            128.0.0.25/2
                                            128.0.0.26/2
                                            128.0.0.28/2
```

```
                                          128.0.0.29/2
                                          128.0.0.31/2
                                     tnp      0x5
IC-WS001:dcfabric       up    up
IC-WS001:dcfabric.0     up    up   inet
                                   mpls
                                   eth-switch
IC-WS001:pme0           up    up
IC-WS001:pme1           up    up
IC-WS001:pme2           up    down
IC-WS001:pme3           up    down
IC-WS001:vcp0           up    up
IC-WS001:vcp0.32769     up    up
IC-WS001:vcp1           up    up
IC-WS001:vcp1.32768     up    up
IC-WS001:vcp2           up    up
IC-WS001:vcp2.32768     up    up
IC-WS001:vcp3           up    down
IC-WS001:vcp3.32768     up    down
IC-WS001:vcp4           up    down
IC-WS001:vcp4.32768     up    down
```

show interfaces statistics fabric device-name

```
user@switch> show interfaces statistics fabric IC-WS001:fte-0/0/13
Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Down
  Interface index: 49177, SNMP ifIndex: 1208484479
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:00:00:00:00:0d, Hardware address: 00:00:00:00:00:0d
  Last flapped   : 2012-11-27 20:30:30 UTC (02:09:53 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0

  Logical interface IC-WS001:fte-0/0/13.32768 (Index 86) (SNMP ifIndex 1208484480)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
```

```
    Input packets : 0
    Output packets: 0
    Protocol eth-switch, MTU: 0
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

- *Monitoring Interface Status and Traffic*
- *Troubleshooting Network Interfaces*
- *Troubleshooting an Aggregated Ethernet Interface*
- Junos OS Network Interfaces Library for Routing Devices

show interfaces statistics st0

IN THIS SECTION

- Syntax | 21488
- Description | 21488
- Required Privilege Level | 21489
- Sample Output | 21489
- Release Information | 21489

Syntax

```
show interfaces statistics interface-name
```

Description

Displays the interface input and output statistics for physical and logical interface.

Required Privilege Level

view

Sample Output

show interfaces statistics

```
user@host> show interfaces statistics st0.1
Logical interface st0.1 (Index 91) (SNMP ifIndex 268)
     Flags: Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
     Input packets : 2743333
     Output packets: 6790470992
     Security: Zone: untrust
     Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp ospf pgm pim rip
router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http https ike netconf ping
reverse-telnet
     reverse-ssh rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text xnm-ssl
lsping ntp sip
     Protocol inet, MTU: 9192
       Addresses, Flags: Is-Preferred Is-Primary
         Destination: 192.167.1.0/30, Local: 192.167.1.1
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Understanding Interfaces

show interfaces summary

IN THIS SECTION

- Syntax | 21490
- Description | 21490
- Options | 21490
- Required Privilege Level | 21490
- Output Fields | 21490
- Sample Output | 21491
- Release Information | 21492

Syntax

```
show interfaces summary
```

Description

Display the status and statistics on logical interfaces configured on the device at the Flexible PIC Concentrator (FPC) level.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 692 on page 21491 describes the output fields for the **show interfaces summary** command. Output fields are listed in the approximate order in which they appear.

Table 692: show interfaces summary Output Fields

Field Name	Field Description
`System's maximum logical interfaces`	Total number of logical interfaces in the device.
`Logical interfaces allocated`	Number of allocated logical interfaces.
`Logical interfaces available`	Number of available logical interfaces.
`Logical interface type`	The type of logical interfaces. • `LSI`—Number of label-switched logical interfaces and their status. • `Ethernet Untagged`—Number of untagged logical interfaces and their status. • `Ethernet VLAN`—Number of tagged logical interfaces and their status. • `Others`—Number of dynamic and other logical interfaces, and their status.
`System`	Statistics on the global logical interfaces in the system.
`FPC x`	Statistics on the logical interfaces in a specific FPC.

Sample Output

show interfaces summary

```
user@host> show interfaces summary
Logical interfaces:
    System's maximum logical interfaces  : 262144
    Logical interfaces allocated         :     31
    Logical interfaces available         : 262113

    System:
    Logical interface type    Count         UP       DOWN
    Total                        28         28          0
```

```
    LSI                             0          0          0
    Ethernet Untagged              15         15          0
    Ethernet VLAN                   0          0          0
    Others                         13         13          0

    FPC1:
    Logical interface type    Count         UP       DOWN
    Total                           3          3          0
    LSI                             0          0          0
    Ethernet Untagged               3          3          0
    Ethernet VLAN                   0          0          0
    Others                          0          0          0

    FPC2:
    Logical interface type    Count         UP       DOWN
    Total                           0          0          0
    LSI                             0          0          0
    Ethernet Untagged               0          0          0
    Ethernet VLAN                   0          0          0
    Others                          0          0          0
```

Release Information

Command introduced in Junos OS Release 14.1R2.

show interfaces swfabx

IN THIS SECTION

- Syntax | 21493
- Description | 21493
- Required Privilege Level | 21493
- Output Fields | 21493
- Sample Output | 21493
- Release Information | 21494

Syntax

```
show interfaces (swfab0 | swfab1)
```

Description

Display the configured interfaces for each swfab interface. The swfab interface can contain one or more members because it is an aggregated interface.

Required Privilege Level

view

Output Fields

Table 693 on page 21493 lists the output fields for the `show interfaces <swfab0 | swfab1>` command. Output fields are listed in the approximate order in which they appear.

Table 693: show interfaces <swfab0 | swfab1> Output Fields

Field Name	Field Description
`fabric-options`	The fabric-options hierarchy is configured to be in sync with the fab interfaces.
`member-interfaces`	Interfaces specified under member-interfaces are single aggregate interfaces. This interface carries internode switching traffic.

Sample Output

show interfaces swfab0

```
user@host# show interfaces swfab0
fabric-options {
                member-interfaces {
                    ge-0/0/9;
                    ge-0/0/10;
```

```
                }
        }
```

show interfaces swfab1

```
user@host# show interfaces swfab1
fabric-options {
                member-interfaces {
                    ge-7/0/9;
                    ge-7/0/10;
                }
        }
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

clear interfaces statistics swfabx | 3653

show interfaces targeting (Aggregated Ethernet for Subscriber Management)

IN THIS SECTION

- Syntax | 21495
- Description | 21495
- Options | 21495
- Required Privilege Level | 21495
- Output Fields | 21495
- Sample Output | 21496
- Release Information | 21497

Syntax

```
show interfaces targeting aex
```

Description

(MX Series routers only) Display status information about the distribution of subscribers on different links in an aggregated Ethernet bundle.

Options

Required Privilege Level

view

Output Fields

Table 694 on page 21495 lists the output fields for the `show interfaces targeting` command. Output fields are listed in the approximate order in which they appear.

Table 694: show interfaces targeting Output Fields

Field Name	Field Description	Level of Output
Aggregated Ethernet Interface		
`Aggregated interface`	Name of the aggregated Ethernet bundle.	All levels
`Redundancy mode`	Redundancy mechanism on the interface: `Link Level Redundancy` or `FPC Redundancy`.	All levels
`Total number of distributed interfaces`	Number of distributed links in the bundle.	All levels
Physical Interface		

Table 694: show interfaces targeting Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Physical interface`	Name of the physical interface and state of the interface.	All levels
`Link status`	Status of the link on the physical interface: `up` or `down`.	
`Number of primary distributions`	Number of subscribers distributed on primary links.	All levels
`Number of backup distributions`	Number of subscribers distributed on backup links.	All levels

Sample Output

show interfaces targeting ae0

```
user@host> show interfaces targeting ae0
Aggregated interface: ae0
Redundancy mode: Link Level Redundancy
Total number of distributed interfaces: 3
Physical interface: ge-1/0/0, Link status: Up
Number of primary distributions: 200
Number of backup distributions: 200
Physical interface: ge-1/1/0, Link status: Up
Number of primary distributions: 200
Number of backup distributions: 199
Physical interface: ge-2/0/7, Link status: Up
Number of primary distributions: 200
Number of backup distributions: 200
Physical interface: ge-2/0/8, Link status: Up
Number of primary distributions: 199
Number of backup distributions: 200
```

Release Information

Command introduced in Junos OS Release 11.2.

show interfaces terse

IN THIS SECTION

- Syntax | 21497
- Description | 21497
- Options | 21497
- Additional Information | 21498
- Required Privilege Level | 21498
- Output Fields | 21498
- Sample Output | 21499
- Release Information | 21503

Syntax

```
show interfaces terse
```

Description

Display summary information about interfaces.

Options

This command has no options.

Additional Information

Interfaces are always displayed in numerical order, from the lowest to the highest FPC slot number. Within that slot, the lowest PIC slot is shown first. On an individual PIC, the lowest port number is always first.

Required Privilege Level

view

Output Fields

Table 695 on page 21498 lists the output fields for the `show interfaces terse` command. Output fields are listed in the approximate order in which they appear.

Table 695: show interfaces terse Output Fields

Field Name	Field Description
Interface	Interface name.
Admin	Whether the interface is turned on (up) or off (down).
Link	Link state: **up** or **down**.
Proto	Protocol family configured on the logical interface. A logical interface on a router that supports Ethernet OAM always shows the multiservice protocol.
Local	Local IP address of the logical interface.
Remote	Remote IP address of the logical interface.

Sample Output

show interfaces terse

```
user@host> show interfaces terse
Interface               Admin Link Proto    Local                 Remote
t1-0/1/0:0              up    up
t1-0/1/0:0.0            up    up   inet     192.168.220.18/30
t1-0/1/0:1              up    up
t1-0/1/0:2              up    up
t1-0/1/0:3              up    up
at-1/0/0                up    up
at-1/0/1                up    up
dsc                     up    up
fxp0                    up    up
fxp0.0                  up    up   inet     192.168.71.249/21
fxp1                    up    up
fxp1.0                  up    up   inet     10.0.0.4/8
                                   tnp      4
gre                     up    up
ipip                    up    up
lo0                     up    up
lo0.0                   up    up   inet     10.0.1.4            --> 0/0
                                            127.0.0.1           --> 0/0
lo0.16385               up    up   inet

lsi                     up    up
mtun                    up    up
```

show interfaces terse (SRX1600 firewall)

```
user@host> show interfaces terse

Interface               Admin Link Proto    Local                 Remote
ge-0/0/0                up    down
gr-0/0/0                up    up
ip-0/0/0                up    up
lt-0/0/0                up    up
ge-0/0/1                up    up
```

```
ge-0/0/1.0              up    up   aenet    --> reth0.0
ge-0/0/2                up    up
ge-0/0/3                up    up
ge-0/0/3.0              up    up   aenet    --> reth1.0
ge-0/0/4                up    up
ge-0/0/5                up    up
ge-0/0/5.0              up    up   aenet    --> fab0.0
ge-0/0/6                up    up
ge-0/0/6.0              up    up   aenet    --> fab0.0
ge-0/0/7                up    down
ge-0/0/8                up    up
ge-0/0/9                up    up
ge-0/0/10               up    up
ge-0/0/11               up    up
ge-0/0/12               up    up
ge-0/0/13               up    down
ge-0/0/14               up    down
ge-0/0/15               up    down
xe-0/2/0                up    up
xe-0/2/1                up    up
xe-0/2/2                up    up
ge-7/0/0                up    down
ge-7/0/1                up    up
ge-7/0/1.0              up    up   aenet    --> reth0.0
ge-7/0/2                up    up
ge-7/0/3                up    up
ge-7/0/3.0              up    up   aenet    --> reth1.0
ge-7/0/4                up    up
ge-7/0/5                up    up
ge-7/0/5.0              up    up   aenet    --> fab1.0
ge-7/0/6                up    up
ge-7/0/6.0              up    up   aenet    --> fab1.0
ge-7/0/7                up    down
ge-7/0/8                up    up
ge-7/0/9                up    up
ge-7/0/10               up    up
ge-7/0/11               up    up
ge-7/0/12               up    up
ge-7/0/13               up    down
ge-7/0/14               up    down
ge-7/0/15               up    down
xe-7/2/0                up    up
xe-7/2/1                up    up
```

```
xe-7/2/2                up    up
dsc                     up    up
em0                     up    up
em0.0                   up    up   inet     129.16.0.1/2
                                           143.16.0.1/2
                                     tnp    0x1100001
em1                     up    up
em1.0                   up    up   inet     129.16.0.1/2
                                           143.16.0.1/2
                                     tnp    0x1100001
em2                     up    up
em2.32768               up    up   inet     192.168.1.2/24
fab0                    up    up
fab0.0                  up    up   inet     30.17.0.200/24
                                     multiservice
fab1                    up    up
fab1.0                  up    up   inet     30.18.0.200/24
                                     multiservice
fti0                    up    up
fxp0                    up    up
fxp0.0                  up    up   inet     10.157.75.73/19
gre                     up    up
ipip                    up    up
irb                     up    up
lo0                     up    up
lo0.0                   up    up   inet     10.255.75.73        --> 0/0
                                           127.0.0.1           --> 0/0
                                     iso    47.0005.80ff.f800.0000.0108.0001.0102.5507.5073
                                     inet6  abcd::10:255:75:73
                                           fe80::baf0:150f:fc8a:5302
lo0.16384               up    up   inet     127.0.0.1           --> 0/0
lo0.16385               up    up   inet     10.0.0.1            --> 0/0
                                           10.0.0.16           --> 0/0
                                           128.0.0.1           --> 0/0
                                           128.0.0.4           --> 0/0
                                           128.0.1.16          --> 0/0
lsi                     up    up
mtun                    up    up
pimd                    up    up
pime                    up    up
pp0                     up    up
ppd0                    up    up
ppe0                    up    up
```

```
reth0                   up    up
reth0.0                 up    up   inet     111.1.1.1/8
                                   multiservice
reth1                   up    up
reth1.0                 up    up   inet     112.1.1.1/8
                                   multiservice
st0                     up    up
st0.16000               up    up
tap                     up    up
vtep                    up    up
```

show interfaces terse (TX Matrix Plus Router)

```
user@host> show interfaces terse

Interface               Admin Link Proto    Local                 Remote
xe-0/0/0                up    up
xe-0/0/1                up    up
xe-0/0/2                up    up
xe-0/0/3                up    up
xe-6/0/0                up    up
xe-6/0/1                up    up
xe-6/0/2                up    up
xe-6/0/3                up    up
xe-6/1/0                up    up
xe-6/1/1                up    up
xe-6/1/2                up    up
xe-6/1/3                up    up
so-0/0/0                up    up
so-0/0/0.0              up    up   inet     10.1.1.1/30
ge-1/3/0.0              up    up   inet     --> 0/0
ge-7/0/0                up    up
ge-7/0/0.0              up    up   inet     10.2.1.1/30
ge-7/0/0.1              up    up   inet     10.2.1.5/30
ge-7/0/0.2              up    up   inet     10.2.1.9/30
ge-7/0/0.3              up    up   inet     10.2.1.13/30
ge-7/0/0.4              up    up   inet     10.2.1.17/30
ge-7/0/0.5              up    up   inet     10.2.1.21/30
...
em0                     up    up
em0.0                   up    up   inet     192.168.178.11/25
```

```
gre                     up    up
ipip                    up    up
ixgbe0                  up    up
ixgbe0.0                up    up   inet     10.34.0.4/8
                                           162.0.0.4/2
                                     inet6  fe80::200:ff:fe22:4/64
                                            fec0::a:22:0:4/64
                                     tnp    0x22000004
ixgbe1                  up    up
ixgbe1.0                up    up   inet     10.34.0.4/8
                                           162.0.0.4/2
                                     inet6  fe80::200:1ff:fe22:4/64
                                            fec0::a:22:0:4/64
                                     tnp    0x22000004
```

show interfaces terse (PTX Series Packet Transport Routers)

```
user@host> show interfaces em0 terse
   Interface               Admin Link Proto    Local                 Remote
    em0                     up    up
    em0.0                   up    up   inet     192.168.3.30/24
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

RELATED DOCUMENTATION

Setting Up Logical Systems

show interfaces terse satellite-device

IN THIS SECTION

- Syntax | 21504
- Description | 21504
- Options | 21504
- Required Privilege Level | 21504
- Output Fields | 21505
- Sample Output | 21505
- Sample Output | 21506
- Release Information | 21508

Syntax

```
show interfaces terse satellite-device (device-alias | all)
```

Description

Display the satellite device extended ports in a Junos Fusion.

Options

device-alias *device-alias*
: Display extended port information for the satellite device using the specified device alias only.

all
: Display information for all extended ports and aggregated Ethernet interfaces with extended ports as members configured on all of the satellite devices.

Required Privilege Level

view

Output Fields

Table 696 on page 21505 lists the output fields for the `show interfaces terse satellite-device` command. Output fields are listed in the approximate order in which they appear.

Table 696: show interfaces terse satellite-device Output Fields

Field Name	Field Description
`Interface`	Interface name.
`Admin`	Whether the interface is turned on (up) or off (down).
`Link`	Link state: **up** or **down**.
`Proto`	Protocol family configured on the logical interface.
`Local`	Local IP address of the logical interface.
`Remote`	Remote IP address of the logical interface.

Sample Output

show interfaces terse satellite-device device-alias

```
user@aggregation-device> show interfaces terse satellite-device TOR1
Interface               Admin Link Proto    Local                 Remote
sd-101/0/0              up    up
sd-101/0/0.32770        up    up   bridge
xe-101/0/14             up    up
xe-101/0/15             up    up
xe-101/0/16             up    up
xe-101/0/17             up    up
xe-101/0/24             up    up
xe-101/0/25             up    up
xe-101/0/31             up    up
xe-101/0/31.0           up    up   bridge
```

```
xe-101/0/32             up    down
xe-101/0/32.0           up    down bridge
xe-101/0/33             up    down
xe-101/0/33.0           up    down bridge
ge-101/0/36             up    down
et-101/0/48             up    down
xe-101/0/50:0           up    up
xe-101/0/50:0.0         up    up   bridge
xe-101/0/50:1           up    up
xe-101/0/50:2           up    up
xe-101/0/50:2.0         up    up   bridge
xe-101/0/50:3           up    up
```

Sample Output

show interfaces terse satellite-device all

```
user@aggregation-device> show interfaces terse satellite-device all
Interface               Admin Link Proto    Local                 Remote
ae0                     up    up
ae0.0                   up    up   bridge
ae1                     up    up
ae1.0                   up    up   bridge
ae2                     up    up
ae2.0                   up    up   bridge
ae3                     up    up
ae3.0                   up    up   bridge
ae4                     up    up
ae4.0                   up    up   bridge
ae5                     up    up
ae5.0                   up    up   bridge
ae6                     up    up
ae6.0                   up    up   bridge
ae7                     up    up
ae7.0                   up    up   bridge
ae8                     up    up
ae8.0                   up    up   bridge
ae9                     up    up
ae9.0                   up    up   bridge
ae10                    up    down
ae10.0                  up    down bridge
```

```
xe-101/0/14             up    up
xe-101/0/15             up    up
xe-101/0/16             up    up
xe-101/0/17             up    up
xe-101/0/24             up    up
xe-101/0/25             up    up
xe-101/0/31             up    up
xe-101/0/31.0           up    up   bridge
xe-101/0/32             up    down
xe-101/0/32.0           up    down bridge
xe-101/0/33             up    down
xe-101/0/33.0           up    down bridge
ge-101/0/36             up    down
et-101/0/48             up    down
xe-101/0/50:0           up    up
xe-101/0/50:0.0         up    up   bridge
xe-101/0/50:1           up    up
xe-101/0/50:2           up    up
xe-101/0/50:2.0         up    up   bridge
xe-101/0/50:3           up    up
xe-102/0/10             up    up
xe-102/0/11             up    up
xe-102/0/12             up    down
xe-102/0/13             up    up
xe-102/0/14             up    up
xe-102/0/15             up    up
xe-102/0/16             up    up
xe-102/0/17             up    up
xe-102/0/24             up    up
xe-102/0/25             up    up
xe-102/0/31             up    up
xe-102/0/31.0           up    up   bridge
xe-102/0/32             up    up
xe-102/0/32.0           up    up   bridge
xe-102/0/33             up    up
xe-102/0/45             up    down
ge-102/0/46             up    down
xe-102/0/47             up    down
et-102/0/48             up    down
et-102/0/49             up    down
et-102/0/50             up    down
et-102/0/51             up    down
et-102/0/52             up    down
```

```
et-102/0/53             up    down
ge-103/0/0              up    up
ge-103/0/0.0            up    up   aenet    --> ae0.0
ge-103/0/1              up    down
ge-103/0/1.0            up    down aenet    --> ae1.0
ge-103/0/2              up    up
ge-103/0/2.0            up    up   aenet    --> ae2.0
ge-103/0/3              up    up
```

Release Information

Command introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

- Configuring or Expanding a Junos Fusion Enterprise

show interfaces terse zone

IN THIS SECTION

- Syntax | 21508
- Description | 21509
- Options | 21509
- Required Privilege Level | 21509
- Sample Output | 21509
- Release Information | 21509

Syntax

```
show interfaces terse zone
```

Description

Display summary information about zone interfaces.

Options

This command has no options.

Required Privilege Level

view

Sample Output

show interface terse zone

```
user@host> show interface terse zone
Interface     Admin    Link   Proto     Local           Remote     Zone
ge-0/0/0.0      up      up     inet     1.4.253.251/16             trust
```

Release Information

Command introduced in Junos OS Release 12.3X48-D20.

show interfaces transport pm

IN THIS SECTION

- Syntax | 21510
- Description | 21510
- Options | 21510
- Required Privilege Level | 21511
- Output Fields | 21511
- Sample Output (Junos OS) | 21516

- Release Information | **21524**

Syntax

```
show interfaces transport pm <category> <time period> <interface>)
```

Description

Display performance monitoring (PM) and threshold crossing alert (TCA) data for long-haul optics.

Display performance data of coherent optics monitored over a definite time period. The data includes the following:

- Optical power received and transmitted
- Module temperature
- Signal-to-noise ratio of the received data

You collect the data in 15-minute or 1-day time interval.

Options

<category>

all	Display information about both optics and OTN
optics	Display information only about optics
otn	Display information only about OTN

<time period>

all	Display information for the current 15-minute interval, the current day, the ninety-six 15-minute intervals, and the previous day.
current	Display information only for the current 15-minute interval.
currentday	Display information only for the current 24 hours
interval	Display information only for the ninety-six 15-minute intervals
previousday	Display information only for the previous day.

<interface>

all Display information for all interfaces or only for the specified interface (for example, *et-fpc*

interface-name Display information for all interfaces or only for the specified interface (for example, *et-fpc*

Required Privilege Level

view

Output Fields

The following tables lists all possible fields for the `show interfaces transport pm` command. The fields that appear may vary depending on the platform and software release.

Table 698 on page 21514 lists the output fields for the `show interfaces transport pm optics` command. Fields are listed in the approximate order in which they appear.

Table 697: show interfaces transport pm Output Fields for Coherent-ZR Optics (400G ZR, 400G ZR-M, 400G ZR-M-HP, and 100G ZR optics)

Field Name	Field Description
`Physical interface`	Name of the physical interface.
SNMP ifIndex	The SNMP index of the interface.
Current	Applicable only for output of current 15-minute PM data, indicates the time range for which t displayed data was accumulated.
Current day	Applicable only for output of current 1-day PM data, indicates the time range for which the displayed data was accumulated.
`Interval`	Applicable only for output of previous 15-minute intervals PM data, indicates the interval number and the time range for which the displayed data was accumulated.
Previous day	Applicable only for output of previous 1-day PM data, indicates the time range for which the displayed data was accumulated.

Table 697: show interfaces transport pm Output Fields for Coherent-ZR Optics (400G ZR, 400G ZR-M, 400G ZR-M-HP, 100G ZR optics) *(Continued)*

Field Name	Field Description
Suspect Flag	Set to *True* if performance monitoring could not be conducted for part or all of the interval duration, otherwise it is set to *False*.
Reason	If the suspect flag is set to True, this field provides a reason why the data is considered to be suspect.
PM	Performance monitoring parameter.
CURRENT	The current instantaneous value of the parameter.
MIN	The minimum observed value of the parameter over the interval duration.
MAX	The maximum observed value of the parameter over the interval duration.
AVG	The average observed value of the parameter over the interval duration.
THRESHOLD (MIN)	Threshold setting for low threshold crossing alert monitoring, if applicable.
THRESHOLD (MAX)	Threshold setting for high threshold crossing alert, if applicable
TCA-ENABLED (MIN)	Indicates if the low threshold crossing alert is enabled, if applicable.
TCA-ENABLED (MAX)	Indicates if the high threshold crossing alert is enabled, if applicable.
TCA-RAISED (MIN)	Indicates if the low threshold crossing has occurred during the interval.
TCA-RAISED (MAX)	Indicates if the high threshold crossing alert has occurred during the interval.
Module temperature (Celsius)	The temperature of the optics module.

Table 697: show interfaces transport pm Output Fields for Coherent-ZR Optics (400G ZR, 400G ZR-M, 400G ZR-M-HP, 100G ZR optics) *(Continued)*

Field Name	Field Description
`Supply voltage`	Supply voltage
`Pre-FEC BER`	The bit error ratio of the received signal based on the FEC corrected bits.
`Uncorrected FER`	The ratio of uncorrectable FEC frames.
`SNR`	Signal to noise ratio.
`Tx power`	The transmit laser output power.
`Rx total power`	The total received optical power on the port.
`Rx signal power`	The received signal optical power on the port.
`Tx bias current(mA)`	Bias current of the transmit laser in mA.
`Carrier frequency offset (MHz)`	Measured carrier frequency offset.
`Chromatic dispersion`	Residual chromatic dispersion measured.
`Differential group delay`	Measured differential group delay.
`SOPMD`	Second order polarization mode dispersion.
`PDL`	Polarization dispersion loss.
`SOP ROC`	State of polarization rate of change in krad/sec.
`OSNR`	Optical signal to noise ratio.

Table 697: show interfaces transport pm Output Fields for Coherent-ZR Optics (400G ZR, 400G ZR-M, 400G ZR-M-HP, 100G ZR optics) *(Continued)*

Field Name	Field Description
`Modulator bias X/I`	Modulator bias in %.
`Modulator bias X/Q`	Modulator bias in %.
`Modulator bias Y/I`	Modulator bias in %.
`Modulator bias Y/Q`	Modulator bias in %.
`Modulator bias X phase`	Modulator bias in %
`Modulator bias Y phase`	Modulator bias in %.
`Q value`	Q value in dB.
`Q margin`	Q value in dB.

Table 698: show interfaces transport pm Output Fields

Field Name	Field Description
`Physical interface`	Name of the physical interface.
`Interval`	The 15 minute interval for performance monitoring.
`Suspect Flag`	Set to *True* if performance monitoring could not be conducted for part or all of the interval duration, otherwise it is set to *False*.
`Reason`	If the suspect flag is set to True, this field provides a reason why the data is considered to be suspect.

Table 698: show interfaces transport pm Output Fields *(Continued)*

Field Name	Field Description
COUNT	Measured value.
THRESHOLD	Threshold value set.
TCA-ENABLED	Set to TRUE if enabled.
TCA-RAISED	TRUE if enabled and the value crosses the threshold.
Near End PM	Near end threshold crossing defect trigger. For more information about each trigger, see "tca" page 28973.
Far End PM	Far end threshold crossing defect trigger. For more information about each trigger, see "tca" o page 28973.
FEC PM	Forwarding Error Correction (FEC) threshold crossing defect trigger. For more information abo each trigger, see "tca" on page 28973.
BER PM	BEC threshold crossing defect trigger. For more information about each trigger, see "tca" on p 28973.
CURRENT	The current instantaneous value of the parameter.
PM	Performance monitor.
MIN	The minimum observed value of the parameter over the interval duration.
MAX	The maximum observed value of the parameter over the interval duration.
AVG	The average observed value of the parameter over the interval duration.
Lane Chromatic dispersion	Residual chromatic dispersion measured.

Table 698: show interfaces transport pm Output Fields *(Continued)*

Field Name	Field Description
Lane differential group delay	Measured differential group delay.
q Value	Measured Quality factor value.
SNR	Signal to noise ratio.
Tx output power	The transmit laser output power.
Rx input power	The received optical power on the port.
Module temperature (Celsius)	The temperature of the optics module.
Tx Laser bias current (0.1mA)	Magnitude of the laser bias power setting current. The laser bias provides direct modulation c laser diodes and modulates currents.
Rx Laser bias current (0.1mA)	Magnitude of the laser bias power setting current.
Carrier frequency offset (MHz)	Measured carrier frequency offset.

Sample Output (Junos OS)

show interfaces transport pm

```
user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current          Elapse time:900 Seconds
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  OTU-BBE                 0                800              No              No
  OTU-ES                  0                135              No              No
  OTU-SES                 0                90               No              No
  OTU-UAS                 427              90               No              No
```

```
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  OTU-BBE                 0                800              No              No
  OTU-ES                  0                135              No              No
  OTU-SES                 0                90               No              No
  OTU-UAS                 0                90               No              No
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  ODU-BBE                 0                800              No              No
  ODU-ES                  0                135              No              No
  ODU-SES                 0                90               No              No
  ODU-UAS                 427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  ODU-BBE                 0                800              No              No
  ODU-ES                  0                135              No              No
  ODU-SES                 0                90               No              No
  ODU-UAS                 0                90               No              No
  FEC            Suspect Flag:False             Reason:None
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  FEC-CorrectedErr        2008544300       0                NA              NA
  FEC-UncorrectedWords    0                0                NA              NA
  BER            Suspect Flag:False             Reason:None
    PM                 MIN       MAX       AVG       THRESHOLD      TCA-ENABLED      TCA-RAISED
  BER                  3.6e-5    5.8e-5    3.6e-5     10.0e-3        No                Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current
  Suspect Flag:True             Reason:Object Disabled
    PM                            CURRENT   MIN       MAX      AVG       THRESHOLD      TCA-
ENABLED       TCA-RAISED
                                                                                 (MIN) (MAX)    (MIN)
(MAX)       (MIN) (MAX)
  Lane chromatic dispersion         0        0         0         0         0     0        NA
NA         NA     NA
  Lane differential group delay     0        0         0         0         0     0        NA
NA         NA     NA
  q Value                           120      120       120       120       0     0        NA
NA         NA     NA
  SNR                               28       28        29        28        0     0        NA
NA         NA     NA
  Tx output power(0.01dBm)          -5000    -5000     -5000     -5000     -300  -100     No
No         No     No
  Rx output power(0.01dBm)          -3642    -3665     -3626     -3637     -1800 -500     No
```

```
No          No     No
  Module temperature(Celsius)       46       46        46        46        -5    75       No
No          No     No
  Tx laser bias current(0.1mA)      0        0         0         0         0     0        NA
NA          NA     NA
  Rx laser bias current(0.1mA)      1270     1270      1270      1270      0     0        NA
NA          NA     NA
  Carrier frequency offset(MHz)     -186     -186      -186      -186      -5000 5000     No
No          No     No
```

show interfaces transport pm optics (400G ZR, 400G ZR-M, 400G ZR-M-HP, and 100G ZR optics)

```
user@host show interfaces transport pm optics current et-2/2/4
Physical interface: et-2/2/4, SNMP ifIndex 650
  Current: 2024-07-04,11:30:00 - current
  Suspect Flag: False  Reason: Not Applicable
    PM                              CURRENT  MIN      MAX      AVG      THRESHOLD        TCA-
ENABLED  TCA-RAISED
                                                                       (MIN)   (MAX)    (MIN)
(MAX)  (MIN) (MAX)
  Module temperature(Celsius)       49       49       49       49       0       73        No
No     No    No
  Supply voltage(V)                 3.31     3.29     3.31     3.29     0       0         NA
NA     NA    NA
  Pre-FEC BER                       1.27e-3  1.20e-3  1.88e-3  1.40e-3  0       9.00e-3   NA
No     NA    No
  Uncorrected FER                   0.00e-0  0.00e-0  0.00e-0  0.00e-0  0       0.00e-0   NA
No     NA    No
  SNR(dB)                           17.4     16.7     17.4     16.7     21.5    0         Yes
NA     Yes   NA
  Tx power(dBm)                     -7.53    -7.54    0.00     -5.99    -16.00  0.00      No
No     No    No
  Rx total power(dBm)               -15.65   -15.68   0.00     -12.55   -21.00  2.00      No
No     No    No
  Rx signal power(dBm)              -15.22   -15.34   -15.22   -15.23   -13.50  0.00      No
No     No    No
  Tx bias current(mA)               244.20   244.20   244.20   244.20   0       0         NA
NA     NA    NA
  Carrier frequency offset(MHz)     0        0        0        0        -3600   3600      No
No     No    No
```

```
  Chromatic dispersion(ps/nm)       0        0        2        0        0        0         NA
NA     NA    NA
  Differential group delay(ps)      0.9      0.9      0.9      0.9      0        0         NA
NA     NA    NA
  SOPMD(ps^2)                       15.9     4.9      167.9    32.6     0        0         NA
NA     NA    NA
  PDL(dB)                           0.7      0.3      0.7      0.3      0        0         NA
NA     NA    NA
  SOP ROC(krad/sec)                 0        0        0        0        0        0         NA
NA     NA    NA
  OSNR(dB)                          36.2     36.2     36.2     36.2     0        0         NA
NA     NA    NA
  Modulator bias X/I(%)             31.89    31.76    32.03    31.86    0        0         NA
NA     NA    NA
  Modulator bias X/Q(%)             29.50    29.38    29.64    29.46    0        0         NA
NA     NA    NA
  Modulator bias Y/I(%)             35.15    35.03    35.31    35.11    0        0         NA
NA     NA    NA
  Modulator bias Y/Q(%)             34.58    34.58    34.85    34.64    0        0         NA
NA     NA    NA
  Modulator bias X phase(%)         35.58    35.48    35.68    35.54    0        0         NA
NA     NA    NA
  Modulator bias Y phase(%)         48.65    48.64    48.83    48.69    0        0         NA
NA     NA    NA
  Q value(dB)                       0.0      0.0      0.0      0.0      0        0         NA
NA     NA    NA
  Q margin(dB)                      0.0      0.0      0.0      0.0      0        0         NA
NA     NA    NA
```

show interfaces transport (MX960 router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host > show interfaces transport et-3/0/0
Administrative State:      In Service
Operational    State:      Normal
```

show interfaces transport pm (MX960 router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host > show interfaces transport pm otn current et-3/0/0
Physical interface: et-3/0/0, SNMP ifIndex 564
  23:30-current          Elapsed time:455 Seconds
  Near End       Suspect Flag:False             Reason:Not Applicable
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  OTU-BBE                 0                800              No              No
  OTU-ES                  0                135              No              No
  OTU-SES                 0                90               No              No
  OTU-UAS                 0                90               No              No
  Far End        Suspect Flag:False             Reason:Not Applicable
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  OTU-BBE                 0                800              No              No
  OTU-ES                  0                135              No              No
  OTU-SES                 0                90               No              No
  OTU-UAS                 0                90               No              No
  Near End       Suspect Flag:False             Reason:Not Applicable
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  ODU-BBE                 0                800              No              No
  ODU-ES                  0                135              No              No
  ODU-SES                 0                90               No              No
  ODU-UAS                 0                90               No              No
  Far End        Suspect Flag:False             Reason:Not Applicable
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  ODU-BBE                 0                800              No              No
  ODU-ES                  0                135              No              No
  ODU-SES                 0                90               No              No
  ODU-UAS                 0                90               No              No
  FEC            Suspect Flag:False             Reason:Not Applicable
    PM                   COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  FEC-CorrectedErr        30865849         0                NA              NA
  FEC-UncorrectedWords    0                0                NA              NA
  BER            Suspect Flag:False             Reason:Not Applicable
    PM                 MIN       MAX       AVG       THRESHOLD      TCA-ENABLED      TCA-RAISED
  BER                  4.0e-7    5.9e-7    5.1e-7     1.0e-2         No                No
```

show interfaces transport (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host > show interfaces transport et-8/0/0
Administrative State:      In Service
Operational    State:      Normal
```

show interfaces transport pm optics (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host > show interfaces transport pm optics current et-4/0/0
Physical interface: et-4/0/0, SNMP ifIndex 544
  02:45-current
  Suspect Flag:False            Reason:Not Applicable
    PM                              CURRENT   MIN       MAX       AVG       THRESHOLD      TCA-
ENABLED       TCA-RAISED
                                                                            (MIN) (MAX)    (MIN)
(MAX)       (MIN) (MAX)
  Lane chromatic dispersion(ps/nm)  -6       -32       45        -1        0     0        NA
NA         NA     NA
  Lane differential group delay(ps) 3        2         4         3         0     0        NA
NA         NA     NA
  Lane Q2 factor(0.1dB)             154      154       155       154       0     0        NA
NA         NA     NA
  SNR(0.1dB)                        167      164       171       165       0     0        NA
NA         NA     NA
  Carrier frequency offset(MHz)     0        0         0         0         -3600 3600     No
No         No     No
  Tx output power(0.01dBm)          0        0         0         0         -1100 300      No
No         No     No
  Rx input total power(0.01dBm)     0        0         0         0         -3000 300      No
No         No     No
  Module temperature(Celsius)       53       53        55        53        -5    75       No
No         No     No
```

show interfaces transport pm otn (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host> show interfaces transport pm otn previousday et-4/0/0
Physical interface: et-4/0/0, SNMP ifIndex 544
```

```
  02:45-current
  Suspect Flag:False            Reason:Not Applicable
    PM                              CURRENT   MIN       MAX      AVG       THRESHOLD      TCA-
ENABLED       TCA-RAISED
                                                                                 (MIN) (MAX)    (MIN)
(MAX)       (MIN) (MAX)
  Lane chromatic dispersion(ps/nm)  -6       -32       45        -1        0     0        NA
NA         NA     NA
  Lane differential group delay(ps) 3        2         4         3         0     0        NA
NA         NA     NA
  Lane Q2 factor(0.1dB)             154      154       155       154       0     0        NA
NA         NA     NA
  SNR(0.1dB)                        167      164       171       165       0     0        NA
NA         NA     NA
  Carrier frequency offset(MHz)     0        0         0         0         -3600 3600     No
No         No     No
  Tx output power(0.01dBm)          0        0         0         0         -1100 300      No
No         No     No
  Rx input total power(0.01dBm)     0        0         0         0         -3000 300      No
No         No     No
  Module temperature(Celsius)       53       53        55        53        -5    75       No
No         No     No
```

show interfaces transport pm optics (MX10003, MX10008, MX10016, and MX204 routers with QSFP-100GE-DWDM2 transceiver)

```
user@host> show interfaces transport pm optics current et-2/0/0

Physical interface: et-2/0/0, SNMP ifIndex 934
  18:00-current
  Suspect Flag:False        Reason:Not Applicable
    PM                              CURRENT   MIN       MAX      AVG       THRESHOLD      TCA-
ENABLED       TCA-RAISED
                                                                                 (MIN) (MAX)    (MIN)
(MAX)       (MIN) (MAX)
  Module temperature(Celsius)       32       32        32        32        -5    75       No
No         No     No
  Pre-FEC BER                       0.00e-0  0.00e-0   0.00e-0   0.00e-0   0     1.00e-3  NA
No         NA     No
  Uncorrected FER                   0.00e-0  0.00e-0   0.00e-0   0.00e-0   0     0.00e-0  NA
No         NA     No
```

```
  Lane 0
    PM                              CURRENT   MIN       MAX      AVG       THRESHOLD      TCA-
ENABLED       TCA-RAISED
                                                                           (MIN) (MAX)    (MIN)
(MAX)       (MIN) (MAX)
  SNR(0.1dB)                         0        0         0         0         150   0        No
NA         No     NA
  Laser tx power(0.01dBm)           -951     -966      -943      -943      -300  -100     No
No         No     No
  Laser rx power(0.01dBm)           8        6         16        6         -1800 -500     No
No         No     No
  Tx laser bias current(0.1mA)      691      675       713       675       0     0        NA
NA         NA     NA
  Laser frequency Error(MHz)        -389     -429      -339      -353      -3000 3000     No
No         No     No
  TEC Current(0.1mA)                492      423       509       435       -5000 5000     No
No         No     No
  Residual ISI(0.1ps/nm)            0        0         0         0         -800  800      No
No         No     No
  PAM Histogram                     0        0         0         0         0     16384    NA
No         NA     No
  Lane 1
    PM                              CURRENT   MIN       MAX      AVG       THRESHOLD      TCA-
ENABLED       TCA-RAISED
                                                                           (MIN) (MAX)    (MIN)
(MAX)       (MIN) (MAX)
  SNR(0.1dB)                         0        0         111       0         150   0        No
NA         No     NA
  Laser tx power(0.01dBm)           -954     -956      -943      -943      -300  -100     No
No         No     No
  Laser rx power(0.01dBm)           27       14        28        14        -1800 -500     No
No         No     No
  Tx laser bias current(0.1mA)      781      770       797       770       0     0        NA
NA         NA     NA
  Laser frequency Error(MHz)        0        -59       19        0         -3000 3000     No
No         No     No
  TEC Current(0.1mA)                486      452       534       475       -5000 5000     No
No         No     No
  Residual ISI(0.1ps/nm)            0        0         0         0         -800  800      No
No         No     No
  PAM Histogram                     0        0         0         0         0     16384    NA
```

```
No          NA      No
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

- clear interfaces transport pm | 3654
- tca | 28973
- transport-monitoring | 30126

show interfaces vcp

IN THIS SECTION

- Syntax | 21524
- Description | 21525
- Options | 21525
- Additional Information | 21525
- Required Privilege Level | 21525
- Output Fields | 21525
- Sample Output | 21549
- Release Information | 21552

Syntax

```
show interfaces vcp-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
```

```
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified Virtual Chassis Port.

> **NOTE**: This command only displays the statistics for the local interface of the router that you are directly logged in to. You cannot display interface status information of other member routers.

Options

`vcp-`*`fpc/pic/port`*	Display standard information about the specified Virtual Chassis Port.
`brief | detail | extensive | terse`	(Optional) Display the specified level of output.
`descriptions`	(Optional) Display interface description strings.
`media`	(Optional) Display media-specific information about network interfaces.
`snmp-index` *`snmp-index`*	(Optional) Display information for the specified SNMP index of the interface.
`statistics`	(Optional) Display static interface statistics.

Additional Information

In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.

Required Privilege Level

view

Output Fields

Table 699: show interfaces vcp Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Physical Interface`	Name of the physical interface.	All levels
`Enabled`	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
`Interface index`	Index number of the physical interface, which reflects its initialization sequence.	`detail` `extensive` none
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail` `extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`
`Link-level type`	Encapsulation being used on the physical interface.	All levels
`MTU`	Maximum transmission unit size on the physical interface.	All levels
`Speed`	Speed at which the interface is running.	All levels
`Loopback`	Loopback status: `Enabled` or `Disabled`. If loopback is enabled, type of loopback: `Local` or `Remote`.	All levels
`Source filtering`	Source filtering status: `Enabled` or `Disabled`.	All levels
`LAN-PHY mode`	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`WAN-PHY mode`	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
`Unidirectional`	Unidirectional link mode status for 10-Gigabit Ethernet interface: `Enabled` or `Disabled` for parent interface; `Rx-only` or `Tx-only` for child interfaces.	All levels
`Flow control`	Flow control status: `Enabled` or `Disabled`.	All levels
`Auto-negotiation`	(Gigabit Ethernet interfaces) Autonegotiation status: `Enabled` or `Disabled`.	All levels
`Remote-fault`	(Gigabit Ethernet interfaces) Remote fault status: • `Online`—Autonegotiation is manually configured as online. • `Offline`—Autonegotiation is manually configured as offline.	All levels
`Device flags`	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
`Interface flags`	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
`Link flags`	Information about the link. Possible values are described in the “Links Flags” section under *Common Output Fields Description*.	All levels
`Wavelength`	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
`Frequency`	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`CoS queues`	Number of CoS queues configured.	`detail extensive` none
`Schedulers`	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	`extensive`
`Hold-times`	Current interface hold-time up and hold-time down, in milliseconds (ms).	`detail extensive`
`Current address`	Configured MAC address.	`detail extensive` none
`Hardware address`	Hardware MAC address.	`detail extensive` none
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *`year-month-day hour`*`:`*`minute`*`:`*`second`*`:`*`timezone`* (*`hour`*`:`*`minute`*`:`*`second`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive` none
`Input Rate`	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
`Output Rate`	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Egress account overhead`	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	`detail extensive`
`Ingress account overhead`	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	`detail extensive`
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • `Output bytes`—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the "show interfaces " on page 20227 command.	`detail extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input errors`	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Errors`—Sum of the incoming frame terminates and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • `L3 incompletes`—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` statement. • `L2 channel errors`—Number of times the software did not find a valid logical interface for an incoming frame. • `L2 mismatch timeouts`—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • `FIFO errors`—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `Resource errors`—Sum of transmit drops.	`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • `Errors`—Sum of the outgoing frame terminates and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. **NOTE**: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the `Drops` field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs. • `Collisions`—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • `Aged packets`—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • `FIFO errors`—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `HS link CRC errors`—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • `MTU errors`—Number of packets whose size exceeded the MTU of the interface.	`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Resource errors`—Sum of transmit drops.	
`Egress queues`	Total number of egress queues supported on the specified interface. **NOTE**: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the `show interfaces` command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenorio, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	`detail extensive`
`Queue counters (Egress)`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism. **NOTE**: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the `Dropped packets` field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.	`detail extensive`
`Ingress queues`	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Queue counters (Ingress)`	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	`extensive`
`Active alarms` and `Active defects`	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value `None` or `Link`. • `None`—There are no active defects or alarms. • `Link`—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	`detail` `extensive` none
Interface transmit statistics	(On MX Series devices) Status of the `interface-transmit-statistics` configuration: Enabled or Disabled. • `Enabled`—When the `interface-transmit-statistics` statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • `Disabled`—When the `interface-transmit-statistics` statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface.	`detail` `extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`OTN FEC statistics`	The forward error correction (FEC) counters provide the following statistics: • `Corrected Errors`—The count of corrected errors in the last second. • `Corrected Error Ratio`—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	`detail extensive`
`PCS statistics`	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • `Bit errors`—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • `Errored blocks`—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	`detail extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`MAC statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC subsystem, including the following:
	• `Total octets` and `total packets`—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the "show interfaces " on page 20227 command.
	• `Unicast packets`, `Broadcast packets`, and `Multicast packets`—Number of unicast, broadcast, and multicast packets.
	• `CRC/Align errors`—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
	• `FIFO error`—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
	• `MAC control frames`—Number of MAC control frames.
	• `MAC pause frames`—Number of MAC control frames with `pause` operational code.
	• `Oversized frames`—There are two possible conditions regarding the number of oversized frames:
	  • Packet length exceeds 1518 octets, or
	  • Packet length exceeds MRU
	• `Jabber frames`—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
		`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Fragment frames`—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • `VLAN tagged frames`—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. **NOTE**: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the `VLAN tagged frames` field displays 0 when the `show interfaces` command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC. • `Code violations`—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	
`OTN Received Overhead Bytes`	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	`extensive`
`OTN Transmitted Overhead Bytes`	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Filter statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • `Input packet count`—Number of packets received from the MAC hardware that the filter processed. • `Input packet rejects`—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • `Input DA rejects`—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). • `Input SA rejects`—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • `Output packet count`—Number of packets that the filter has given to the MAC hardware. • `Output packet pad count`—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • `Output packet error count`—Number of packets with an indicated error that the filter was given to transmit. These packets are usually	`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • `CAM destination filters`, `CAM source filters`—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	
`PMA PHY`	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `PHY Lock`—Phase-locked loop • `PHY Light`—Loss of optical signal	`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`WIS section`	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `BIP-B1`—Bit interleaved parity for SONET section overhead • `SEF`—Severely errored framing • `LOL`—Loss of light • `LOF`—Loss of frame • `ES-S`—Errored seconds (section) • `SES-S`—Severely errored seconds (section) • `SEFS-S`—Severely errored framing seconds (section)	`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`WIS line`	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information: • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `BIP-B2`—Bit interleaved parity for SONET line overhead • `REI-L`—Remote error indication (near-end line) • `RDI-L`—Remote defect indication (near-end line) • `AIS-L`—Alarm indication signal (near-end line) • `BERR-SF`—Bit error rate fault (signal failure) • `BERR-SD`—Bit error rate defect (signal degradation) • `ES-L`—Errored seconds (near-end line) • `SES-L`—Severely errored seconds (near-end line) • `UAS-L`—Unavailable seconds (near-end line) • `ES-LFE`—Errored seconds (far-end line) • `SES-LFE`—Severely errored seconds (far-end line) • `UAS-LFE`—Unavailable seconds (far-end line)	`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`WIS path`	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information: • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `BIP-B3`—Bit interleaved parity for SONET section overhead • `REI-P`—Remote error indication • `LOP-P`—Loss of pointer (path) • `AIS-P`—Path alarm indication signal • `RDI-P`—Path remote defect indication • `UNEQ-P`—Path unequipped • `PLM-P`—Path payload (signal) label mismatch • `ES-P`—Errored seconds (near-end STS path) • `SES-P`—Severely errored seconds (near-end STS path) • `UAS-P`—Unavailable seconds (near-end STS path) • `SES-PFE`—Severely errored seconds (far-end STS path) • `UAS-PFE`—Unavailable seconds (far-end STS path)	`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Autonegotiation information`	Information about link autonegotiation.
	• `Negotiation status:`
	  • `Incomplete`—Ethernet interface has the speed or link mode configured.
	  • `No autonegotiation`—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • `Complete`—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• `Link partner status`—`OK` when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• `Link partner`—Information from the remote Ethernet device:
	  • `Link mode`—Depending on the capability of the link partner, either `Full-duplex` or `Half-duplex`.
	  • `Flow control`—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are `Symmetric` (link partner supports `PAUSE` on receive and transmit), `Asymmetric` (link partner supports `PAUSE` on transmit), `Symmetric/Asymmetric` (link partner supports `PAUSE` on receive and transmit or only `PAUSE` on transmit), and `None` (link partner does not support flow control).
	  • `Remote fault`—Remote fault information from the link partner—`Failure` indicates a receive link error. `OK` indicates that the link partner is receiving. `Negotiation error` indicates a negotiation error. `Offline` indicates that the link partner is going offline.
	• `Local resolution`—Information from the local Ethernet device:
	  • `Flow control`—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are `Symmetric/Asymmetric` (local device supports `PAUSE` on receive
		`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	and transmit or only PAUSE on receive) and None (local device does not support flow control). Depending on the result of the negotiation with the link partner, local resolution flow control type will display Symmetric (local device supports PAUSE on receive and transmit), Asymmetric (local device supports PAUSE on receive), and None (local device does not support flow control). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive).	
Received path trace, Transmitted path trace	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.	extensive
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • Destination slot—FPC slot number.	extensive

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`CoS information`	Information about the CoS queue for the physical interface. • `CoS transmit queue`—Queue number and its associated user-configured forwarding class name. • `Bandwidth %`—Percentage of bandwidth allocated to the queue. • `Bandwidth bps`—Bandwidth allocated to the queue (in bps). • `Buffer %`—Percentage of buffer space allocated to the queue. • `Buffer usec`—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • `Priority`—Queue priority: `low` or `high`. • `Limit`—Displayed if rate limiting is configured for the queue. Possible values are `none` and `exact`. If `exact` is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If `none` is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	`extensive`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Index number of the logical interface, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP interface index number for the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
`VLAN-Tag`	Rewrite profile applied to incoming or outgoing frames on the outer (`Out`) VLAN tag or for both the outer and inner (`In`) VLAN tags. • `push`—An outer VLAN tag is pushed in front of the existing VLAN tag. • `pop`—The outer VLAN tag of the incoming frame is removed. • `swap`—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • `push`—An outer VLAN tag is pushed in front of the existing VLAN tag. • `push-push`—Two VLAN tags are pushed in from the incoming frame. • `swap-push`—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • `swap-swap`—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • `pop-swap`—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • `pop-pop`—Both the outer and inner VLAN tags of the incoming frame are removed.	`brief` `detail` `extensive` none

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Demux`	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	`detail` `extensive` none
`Encapsulation`	Encapsulation on the logical interface.	All levels
`ACI VLAN: Dynamic Profile`	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.	`brief` `detail` `extensive` none
`Protocol`	Protocol family. Possible values are described in the “Protocol Field” section under *Common Output Fields Description*.	`detail` `extensive` none
`MTU`	Maximum transmission unit size on the logical interface.	`detail` `extensive` none
`Dynamic Profile`	(MX Series routers with Trio MPCs only) Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	`detail` `extensive` none
`Service Name Table`	(MX Series routers with Trio MPCs only) Name of the service name table for the interface configured with a PPPoE family.	`detail` `extensive` none
`Max Sessions`	(MX Series routers with Trio MPCs only) Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	`detail` `extensive` none
`Duplicate Protection`	(MX Series routers with Trio MPCs only) State of PPPoE duplicate protection: `On` or `Off`. When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	`detail` `extensive` none

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Direct Connect	State of the configuration to ignore DSL Forum VSAs: `On` or `Off`. When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	`detail extensive` none
`AC Name`	Name of the access concentrator.	`detail extensive` none
`Maximum labels`	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	`detail extensive` none
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the specified interface set. • `Input bytes, Output bytes`—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • `Input packets, Output packets`—Number of packets received and transmitted on the interface set.	`detail extensive`
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	`extensive`
`Local statistics`	Number and rate of bytes and packets destined to the router.	`extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Transit statistics`	Number and rate of bytes and packets transiting the switch. **NOTE**: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the `Output bytes` and `Output packets` interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	`extensive`
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Route Table`	Route table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail extensive` none
`Flags`	Information about protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	`detail extensive`
`Donor interface`	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	`detail extensive` none
`Preferred source address`	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	`detail extensive` none
`Input Filters`	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	`detail extensive`
`Output Filters`	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	`detail extensive`

Table 699: show interfaces vcp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Mac-Validate Failures`	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	`detail extensive` none
`Addresses, Flags`	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	`detail extensive` none
`protocol-family`	Protocol family configured on the logical interface. If the protocol is `inet`, the IP address of the interface is also displayed.	`brief`
`Flags`	Information about the address flag. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	`detail extensive` none
`Destination`	IP address of the remote side of the connection.	`detail extensive` none
`Local`	IP address of the logical interface.	`detail extensive` none
`Broadcast`	Broadcast address of the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`

Sample Output

show interfaces vcp (specific interface)

```
user@host> show interfaces vcp-4/0/0
Physical interface: vcp-4/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 821
  Link-level type: 70, MTU: 9148, Speed: 1000mbps
  Device flags   : Present Running
```

```
  Interface flags: Promiscuous Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: Unspecified, Hardware address: Unspecified
  Last flapped   : 2015-11-05 11:29:58 PST (10:39:25 ago)
  Input rate     : 18232 bps (26 pps)
  Output rate    : 17160 bps (25 pps)
  Active alarms  : None
  Active defects : None
  Interface transmit statistics: Disabled

  Logical interface vcp-4/0/0.32768 (Index 64) (SNMP ifIndex 822)
    Flags: Encapsulation: Unspecified
```

show interfaces vcp brief (specific interface)

```
user@host> show interfaces vcp-4/0/0 brief
Physical interface: vcp-4/0/0, Enabled, Physical link is Up
  Link-level type: 70, MTU: 9148, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: Promiscuous Internal: 0x4000
  Link flags     : None

  Logical interface vcp-4/0/0.32768
    Flags: Down Encapsulation: Unspecified
```

show interfaces vcp detail (specific interface)

```
user@host> show interfaces vcp-4/0/0 detail
Physical interface: vcp-4/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 821, Generation: 131
  Link-level type: 70, MTU: 9148, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: Promiscuous Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
```

```
Hold-times     : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Last flapped   : 2015-11-05 11:29:58 PST (10:40:23 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :             74190791                12624 bps
 Output bytes  :             72984813                12080 bps
 Input  packets:              1050663                   26 pps
 Output packets:              1041055                   26 pps
 IPv6 transit statistics:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Egress queues: 8 supported, 4 in use
Queue counters:        Queued packets  Transmitted packets      Dropped packets
  0 best-effort

                                 698                  698                    0
  1 expedited-forwarding

                                   0                    0                    0
  2 assured-forwarding

                                   0                    0                    0
  3 network-control

                              1040348              1040348                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : None
Active defects : None
Interface transmit statistics: Disabled

Logical interface vcp-4/0/0.32768 (Index 64) (SNMP ifIndex 822) (Generation 129)
  Flags: Down Encapsulation: Unspecified
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Guidelines for Configuring Virtual Chassis Ports

show virtual-chassis protocol interface (MX Series Virtual Chassis) | **27650**

show interfaces (View Aggregated Ethernet)

IN THIS SECTION

- Syntax | **21552**
- Description | **21553**
- Options | **21553**
- Required Privilege Level | **21553**
- Output Fields | **21553**
- Sample Output | **21564**
- Release Information | **21567**

Syntax

```
show interfaces <aenumber | rethnumber >
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, status information for all interfaces is displayed.

> **NOTE**: This command only provides interface statistics for a redundant ethernet interface (reth) when executed on the node which has the active members/links of the redundant ethernet interface.

Options

- ae*number* | reth*number*—(Optional) Display standard information about the specified aggregated Ethernet interface or redundant Ethernet interface.
- brief | detail | extensive | terse—(Optional) Display the specified level of output.
- descriptions—(Optional) Display interface description strings.
- media—(Optional) Display media-specific information.
- snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.
- statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 700 on page 21553 lists the output fields for the `show interfaces` (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 700: Aggregated Ethernet show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		

Table 700: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Physical interface`	Name of the physical interface and state of the interface.	All levels
`Enabled`	State of the physical interface.	All levels
`Interface index`	Index number of the physical interface, which reflects its initialization sequence.	All levels
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail` `extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`
`Link-level type`	Encapsulation being used on the physical interface.	All levels
`MTU`	Maximum transmission unit size on the physical interface.	All levels
`Speed`	Speed at which the interface is running.	All levels
`Loopback`	Loopback status: `Enabled` or `Disabled`. If loopback is enabled, type of loopback: `Local` or `Remote`.	All levels
`Source filtering`	Source filtering status: `Enabled` or `Disabled`.	All levels
`Flow control`	Flow control status: `Enabled` or `Disabled`.	All levels
`Minimum links needed`	Number of child links that must be operational for the aggregate interface to be operational.	All levels
`Device flags`	Information about the physical device.	All levels

Table 700: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface flags`	Information about the interface.	All levels
`Current address`	Configured MAC address.	`detail extensive`
`Hardware address`	Hardware MAC address.	`detail extensive`
`Last flapped`	Date, time, and how long ago the interface went from down to up or up to down. The format is `Last flapped:` *`year-month-day hours:minutes:seconds timezone`* (*`hours:minutes:seconds`* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.	`detail extensive`
`Input Rate`	Input rate in bits per second (bps) and packets per second (pps).	None specified
`Output Rate`	Output rate in bps and pps.	None specified
`Statistics last cleared`	Time when the statistics for the interface were last set to zero.	`detail extensive`
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`

Table 700: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.	`detail extensive`
`Input errors`	Input errors on the interface: • `Errors`—Sum of incoming frame terminates and frame check sequence (FCS) errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid FCS. • `Runts`—Number of frames received that are smaller than the runt threshold. • `Giants`—Number of frames received that are larger than the giant threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • `Resource errors`—Sum of transmit drops.	`detail extensive`

Table 700: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions` —Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • `Errors`—Sum of the outgoing frame terminates and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `MTU errors`—Number of packets whose size exceeded the MTU of the interface. • `Resource errors`—Sum of transmit drops.	`detail extensive`
`Egress queues`	Total number of egress queues supported on the specified interface	`detail extensive`
`Queue counters`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.	`detail extensive`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels

Table 700: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Index`	Index number of the logical interface (which reflects its initialization sequence).	`detail` `extensive` none
`SNMP ifIndex`	SNMP interface index number of the logical interface.	`detail` `extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail` `extensive`
`Flags`	Information about the logical interface.	All levels
`VLAN-Tag`	Tag Protocol Identifier (TPID) and VLAN identifier.	All levels
`Demux`	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	`detail` `extensive` none
`Encapsulation`	Encapsulation on the logical interface.	All levels
`Statistics`	Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface. • `Bundle`—Information about input and output bundle rates. • `Link`—(`detail` and `extensive` only) Information about specific links in the aggregate, including link state and input and output rates.	`detail` `extensive` none

Table 700: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`LACP info`	Link Aggregation Control Protocol (LACP) information for each aggregated interface.
	• `Role` can be one of the following:
	  • `Actor`—Local device participating in LACP negotiation.
	  • `Partner`—Remote device participating in LACP negotiation.
	• `System priority`—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer.
	• `System identifier`—Actor or partner system ID, encoded as a MAC address.
	• `Port priority`—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer.
	• `Port number`—Port number assigned to the port by the actor or partner, encoded as an unsigned integer.
	• `Port key`—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer.
		`detail extensive`
		none

`LACP Statistics`	LACP statistics for each aggregated interface. • `LACP Rx`—LACP received counter that increments for each normal hello. • `LACP Tx`—Number of LACP transmit packet errors logged. • `Unknown Rx`—Number of unrecognized packet errors logged. • `Illegal Rx`—Number of invalid packets received.	`detail extensive` none

Table 700: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Marker Statistic`	(`detail` and `extensive` only) Information about 802.3ad marker protocol statistics on the specified links. • `Marker Rx`—Number of valid marker PDUs received on this aggregation port. • `Resp Tx`—Number of marker response PDUs transmitted on this aggregation port. • `Unknown Rx`—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown protocol data unit (PDU), or are addressed to the slow protocols group MAC address (43B.3) but do not carry the slow protocols Ethernet type. • `Illegal Rx`—Number of frames received that carry the slow protocols Ethernet type value (43B.4) but contain a badly formed PDU or an illegal value of protocol subtype (43B.4).	`detail extensive` none
`Flow Statistics`	Flow statistics for each aggregated interface.	`detail extensive` none
`Flow Input statistics`	Statistics for packets received by the flow module.	`detail extensive` none
`Flow Output statistics`	Statistics for packets sent by the flow module.	`detail extensive` none

Table 700: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Flow error statistics (Packets dropped due to)`	Packet drop statistics for the flow module.
	• `Address spoofing`—Packet dropped when the screen module detected address spoofing.
	• `Authentication failed`—Packet dropped because the IPsec Encapsulating Security Payload (ESP) or Authentication Header (AH) authentication failed.
	• `Incoming NAT errors`—Packet dropped because the source NAT rule search failed, an invalid source NAT binding was found, or the NAT allocation failed.
	• `Invalid zone received packet`—This counter is not currently in use.
	• `Multiple user authentications`—Packet dropped if it matches more than one policy that specifies user authentication. (Sometimes packets are looped through the system more than once. Each time a packet passes through the system, that packet must be permitted by a policy.)
	• `Multiple incoming NAT`—Packet dropped if source NAT is specified more than once. (Sometimes packets are looped through the system more than once.)
	• `No parent for a gate`—This counter is not currently in use.
	• `No one interested in self packets`—This counter is incremented for one of the following reasons:
	  • The outbound interface is a self interface, but the packet is not marked as a to-self packet and the destination address is in a source NAT pool.
	  • No service is interested in the to-self packet
	  • When a zone has ident-reset service enabled, the TCP RST to IDENT request for port 113 is sent back and this counter is incremented.
		`detail extensive`
		none

Table 700: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Flow error statistics (Packets dropped due to)`	Packet drop statistics for the flow module (continued).
	• `No minor session`—Packet dropped because no minor sessions are available and a minor session was requested. Minor sessions are allocated for storing additional TCP state information.
	• `No more sessions`—Packet dropped because there were no more free sessions available.
	• `No NAT gate`—This counter is not currently in use.
	• `No route present`—Packet dropped because a valid route was not available to forward the packet.
	• `No SA for incoming SPI`—Packet dropped because the incoming IPsec packet's security parameter index (SPI) does not match any known SPI.
	• `No tunnel found`—Packet dropped because a valid tunnel could not be found.
	• `No session for a gate`—Packet dropped by an ALG.
	• `No zone or NULL zone binding`—Packet dropped because its incoming interface was not bound to any zone.
	• `Policy denied`—The error counter is incremented for one of the following reasons:
	  • Source or destination NAT (or both) has occurred and policy says to drop the packet.
	  • Policy specifies user authentication, which failed.
	  • Policy was configured to deny this packet.
	• `Security association not active`—Packet dropped because an IPsec packet was received for an inactive SA.
	• `TCP sequence number out of window`—TCP packet with a sequence number failed the TCP sequence number check that was received.
		`detail extensive`
		none

Table 700: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

	• `Syn-attack protection`—Packet dropped because of SYN attack protection or SYN cookie protection.
	• `User authentication errors`—Packet dropped because policy requires authentication; however:
	  • Only Telnet, FTP, and HTTP traffic can be authenticated.
	  • The corresponding authentication entry could not be found, if web-auth is specified.
	  • The maximum number of authenticated sessions per user was exceeded.	

`protocol-family`	Protocol family configured on the logical interface.	`brief`
`Protocol`	Protocol family configured on the logical interface.	`detail extensive` none
`MTU`	Maximum transmission unit size on the logical interface.	`detail extensive` none
`Maximum labels`	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`Route Table`	Routing table in which the logical interface address is located. For example, `0` refers to the routing table inet.0.	`detail extensive`
`Flags`	Information about protocol family flags.	`detail extensive` none
`Mac-Validate Failures`	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	`detail extensive` none

Table 700: Aggregated Ethernet show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Addresses, Flags`	Information about address flags.	`detail extensive` none
`Destination`	IP address of the remote side of the connection.	`detail extensive` none
`Local`	IP address of the logical interface.	`detail extensive` none
`Broadcast`	Broadcast address of the logical interface.	`detail extensive` none
`Policer`	Policer to be evaluated when packets are received or transmitted on the interface.	`detail extensive` none
`Generation`	Unique number for use by Juniper Networks technical support only.	`detail extensive`
`BFD View`	Number of child interfaces that bind with an Aggregated Ethernet and the details about each of the child interface.	`detail extensive`

Sample Output

show interfaces extensive (Aggregated Ethernet)

```
user@host> show interfaces ae0 extensive

Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 511, Generation: 132
  Description: peering_circuit_1
  Link-level type: Ethernet, MTU: 9192, Speed: 300Gbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
```

```
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: d6:04:ff:61:b8:34, Hardware address: d6:04:ff:61:b8:34
  Last flapped   : 2021-12-28 05:29:46 AST (12w0d 17:21 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :    11536008239086728          15907122832 bps
   Output bytes  :    79363162639161356          94333055640 bps
   Input  packets:       34243932723902              6000894 pps
   Output packets:       71090422616807             10394710 pps
   IPv6 transit statistics:
   Input  bytes  :            132607566
   Output bytes  :             64814847
   Input  packets:              1214245
   Output packets:               205780
  Label-switched interface (LSI) traffic statistics:
   Input  bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  MAC statistics:                      Receive         Transmit
    Broadcast packets                     3325             3336
    Multicast packets                507930812        508053261
  Input errors:
    Errors: 3, Drops: 0, Framing errors: 3, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 1, Drops: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 8 supported, 6 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                    0                    0
    1                                0                    0                    0
    2                                0                    0                    0
    3                                0                    0                    0
    4                                0                    0                    0
    5                                0                    0                    0
  Egress queues: 8 supported, 6 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                   71060094226903       71060094226903                    0
    1                         22965705             22965705                    0
```

```
  2                        685726357            685726357                    0
  3                        663839376            663839376                    0
  4                      29235543467          29235543467                    0
  5                        223155458            223155458                    0
Queue number:         Mapped forwarding classes
  0                   BEST-EFFORT
  1                   PREMIUM-DATA
  2                   BUSINESS
  3                   NC
  4                   VIDEO
  5                   VOICE

Logical interface ae0.0 (Index 333) (SNMP ifIndex 517) (Generation 145)
  Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
  Statistics        Packets        pps         Bytes          bps
  Bundle:
      Input : 34243935928183    6010160 11536009310170467  16116995960
      Output: 71090429597976   10451659 79363169937806454  94884697104
  Adaptive Statistics:
      Adaptive Adjusts:          0
      Adaptive Scans  :          0
      Adaptive Updates:          0
  Link:
    et-0/0/0.0
      Input : 11298708983442    2030928 3805108340032805   5555110248
      Output: 23308119937548    3424862 26023713920361802  30952365312
    et-3/0/1.0
      Input : 11419377486627    1980800 3848204937492876   5200758512
      Output: 23747173576030    3493846 26511514194652619  31799446616
    et-10/0/0.0
      Input : 11412365473613    1998432 3843844231672195   5361127200
      Output: 23736059007904    3532951 26496037962502537  32132885176

  Aggregate member links: 3

  BFD View :           Link        Usable
                 et-0/0/0.0         Yes
                 et-3/0/1.0         Yes
                et-10/0/0.0         Yes
  LACP info:        Role     System             System       Port     Port    Port
                           priority         identifier   priority   number     key
    et-0/0/0.0     Actor        127  d4:04:ff:61:c7:e0        127       41       1
    et-0/0/0.0   Partner        127  54:1e:56:54:0d:08        127       49      15
```

```
      et-3/0/1.0     Actor        127  d4:04:ff:61:c7:e0        127        6       1
      et-3/0/1.0   Partner        127  54:1e:56:54:0d:08        127        7      15
      et-10/0/0.0    Actor        127  d4:04:ff:61:c7:e0        127       15       1
      et-10/0/0.0  Partner        127  54:1e:56:54:0d:08        127       19      15
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      et-0/0/0.0           7238124     7237616            0            0
      et-3/0/1.0           7317685     7317124            0            0
      et-10/0/0.0          7317430     7317147            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/0.0                 0           0            0            0
      et-3/0/1.0                 0           0            0            0
      et-10/0/0.0                0           0            0            0
    Protocol inet, MTU: 9178
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
```

Release Information

Command modified in Junos OS Release 10.2.

show interfaces (Virtual Loopback Tunnel)

IN THIS SECTION

- Syntax | 21568
- Description | 21568
- Options | 21568
- Required Privilege Level | 21568
- Output Fields | 21568
- Sample Output | 21572
- Release Information | 21574

Syntax

```
show interfaces vt-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified virtual loopback tunnel interface.

Options

vt-*fpc/pic/port*	Display standard information about the specified virtual loopback tunnel interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 701 on page 21569 lists the output fields for the `show interfaces` (virtual loopback tunnel) command. Output fields are listed in the approximate order in which they appear.

Table 701: Virtual Loopback Tunnel show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface.	All levels
Link-level type	Encapsulation used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified

Table 701: Virtual Loopback Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets, Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified
Bandwidth	Bandwidth allotted to the logical interface, in kilobytes per second.	All levels

Table 701: Virtual Loopback Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
protocol-family	Protocol family configured on the logical interface. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**brief**
Protocol	Protocol family configured on the logical interface. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**

Table 701: Virtual Loopback Tunnel show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	**detail extensive** none

Sample Output

show interfaces (Virtual Loopback Tunnel)

```
user@host> show interfaces vt-1/2/0
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 40
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Device flags   : Present Running
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface vt-1/2/0.0 (Index 76) (SNMP ifIndex 57)
    Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel
  Input packets : 0
  Output packets: 0
    Protocol inet, MTU: Unlimited
      Flags: None
    Protocol mpls, MTU: Unlimited, Maximum labels: 3
      Flags: None
```

show interfaces brief (Virtual Loopback Tunnel)

```
user@host> show interfaces vt-1/2/0 brief
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Device flags   : Present Running

  Logical interface vt-1/2/0.0
```

```
    Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel
    inet
    mpls
```

show interfaces detail (Virtual Loopback Tunnel)

```
user@host> show interfaces vt-1/2/0 detail
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 40, Generation: 27
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps

  Logical interface vt-1/2/0.0 (Index 76) (SNMP ifIndex 57) (Generation 17)
    Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: Unlimited, Generation: 33, Route table: 0
      Flags: None
    Protocol mpls, MTU: Unlimited, Maximum labels: 3, Generation: 34, Route table: 0
      Flags: None
```

show interfaces extensive (Virtual Loopback Tunnel)

```
user@host> show interfaces vt-1/2/0 extensive
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 40, Generation: 27
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps

  Logical interface vt-1/2/0.0 (Index 76) (SNMP ifIndex 57) (Generation 17)
    Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: Unlimited, Generation: 33, Route table: 0
      Flags: None
    Protocol mpls, MTU: Unlimited, Maximum labels: 3, Generation: 34, Route table: 0
      Flags: None
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces vlan

IN THIS SECTION

- Syntax | 21575
- Description | 21575
- Options | 21575
- Required Privilege Level | 21576
- Output Fields | 21576
- Sample Output | 21588
- Release Information | 21593

Syntax

```
show interfaces (vlan | vlan.vlan-id)
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about routed VLAN interfaces (RVIs).

Options

vlan | vlan.*vlan-id*	Display status information for the specified RVI.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.

media	(Optional) Display media-specific information about network interfaces.
routing-instance (all | *instance-name*)	(Optional) Associate this RVI with the named routing instance.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 702 on page 21576 lists the output fields for the `show interfaces vlan` command. Output fields are listed in the approximate order in which they appear. The level of output *none* means the basic command with no optional options—that is, either just **show interfaces vlan** or **show interfaces vlan.*vlan-id***.

Table 702: show interfaces vlan Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface, which is always **vlan**.	All levels
Enabled	State of the interface: **Enabled** or **Disabled**, followed by the statement **Physical link is <Up/Down>**	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Type	Because this is routed VLAN interface information, this entry is always **VLAN**.	**detail extensive** none
Link-level type	Encapsulation (added control information) being used on the physical interface. Because this is routed VLAN interface information, this entry is always **VLAN**.	All levels
MTU	Maximum transmission unit (MTU) size on the physical interface. The default MTU size depends on the switch platform. Changing either the media MTU or protocol MTU causes an interface to be deleted and added again.	All levels
Clocking	Value is always **Unspecified**—not applicable on switches.	**detail extensive**
Speed	Speed of the interface, either **Auto** if autonegotiation of speed is enabled or a number representing the configured speed in megabits per second.	**detail extensive** none

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Device flags	Information about the physical device such as: **Dest-route-down**—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status. **Down**—Device has been administratively disabled. **Hear-Own-Xmit**—Device receives its own transmissions. **Is-Default**—This address is the default address of the switch. The default address is used as the source address by SNMP, ping, traceroute, and other network utilities. **Is-Preferred**—This address is the default local address for packets originating from the local switch and sent to destinations on the subnet. **Is-Primary**—This address is the default local address for broadcast and multicast packets originated locally and sent out the interface. **Link-Layer-Down**—The link-layer protocol has failed to connect with the remote endpoint. **Loopback**—Switch is in physical loopback. **Loop-Detected**—The link layer has received frames that it sent, thereby detecting a physical loopback. **No-Carrier**—On media that support carrier recognition, no carrier is currently detected. **No-Multicast**—Device does not support multicast traffic. **Preferred**—This address is a candidate to become the preferred address. **Present**—Device is physically present and recognized. **Promiscuous**—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media. **Primary**—This address is a candidate to become the primary address. **Quench**—Transmission on the device is quenched, because the output buffer is overflowing **Recv-All-Multicasts**—Device is in multicast promiscuous mode and therefore provides no multicast filtering. **Running**—Device is active and enabled.	**detail extensive** none
Link type	Link mode of the interface—**Auto** if autonegotiation is enabled, or the configured **Full-Duplex** or **Half-Duplex**.	**detail extensive** none
Link flags	Value is always **None**—not applicable on switches.	**detail extensive** none
Physical Info	Value is always **Unspecified**—not applicable on switches.	**detail extensive**

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	MAC address of the hardware.	**detail extensive** none
Hardware address	MAC address of the switch.	**detail extensive** none
Alternate link address	Value is always **Unspecified**—not applicable on switches.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).** The entry can also be **Never**.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive** none

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets transmitted or received on the physical interface for supported switches. • **Input bytes**—Number of bytes received on the interface for this switch. This value reflects the information gathered by the automatic ingress counter on EX3200 switches and EX4200 switches. EX8200 switches can also be configured to collect this information with the command l3-interface-ingress-counting. • **Output bytes**—Number of bytes sent on the interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches. • **Input packets**—Number of packets received on the interface for this switch. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command l3-interface-ingress-counting. • **Output packets**—Number of packets sent on the interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches.	**detail extensive**

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IPv6 transit statistics	Number and rate of bytes and packets transmitted and/or received on the IPv6 interface for supported switches. • **Input bytes**—Number of bytes received on the interface. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command l3-interface-ingress-counting. • **Output bytes**—Number of bytes sent on the IPv6 interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches. • **Input packets**—Number of packets received on the interface. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command l3-interface-ingress-counting. • **Output packets**—Number of packets sent on the IPv6 interface. This value reflects the information gathered by the automatic egress counter for and EX8200 switches.	**detail extensive**

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input Errors	Input errors on the interface. The following paragraphs explain some of the counters whose meaning may not be obvious. • **Errors**—Sum of the incoming frame terminated and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this value increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 sanity checks of the headers. For example, a frame with less than 20 bytes of available IP header is discarded. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC. If this value is ever nonzero, the interface is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This value does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the interface is malfunctioning. • **Errors**—Sum of the outgoing frame terminated and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. Both Gigabit Ethernet interfaces and 10 Gigabit Ethernet interfaces support only full-duplex operation, so for those two interfaces, this value should always be zero. If the value is nonzero for either Gigabit Ethernet or 10 Gigabit Ethernet, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the interface. If this value is ever nonzero, the interface is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
vlan.*vlan-id*, Index, SNMP ifIndex	VLAN ID, index, and SNMP index number for the logical interface. The logical interface index values reflect the item's initialization sequence.	**detail extensive** none
Generation	Unique number for Juniper Networks Technical support use only.	**detail extensive** none
Flags	Errors that have occurred on this interface, such as **Link Layer Down**. Other possible flags include: • **Device-down**—Device has been administratively disabled. • **Disabled**—Interface is administratively disabled. • **Down**—A hardware failure has occurred. • **Hardware-Down**—Interface protocol initialization failed to complete successfully. • **SNMP-Traps**—SNMP trap notifications are enabled. • **Up**—Interface is enabled and operational.	**detail extensive** none
SNMP-Traps	Each configured SNMP trap has a number that appears here—0x0 is always displayed for logical interface SNMP traps.	**detail extensive** none
Encapsulation	Encapsulation method, which is the process of adding control information. The value is always Ethernet 2 (**ENET2**) for logical encapsulation.	**detail extensive** none

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical interface of supported switches. Traffic statistics represent the sum of the next two fields, Local statistics and Transit statistics. Note that these are not the values for the RVI ingress or egress counters—for that value, see Transit statistics below. • **Input bytes**—Number of bytes received on the interface. Same value as the physical interface. • **Output bytes**—Number of bytes sent on the interface. Same value as the physical interface. • **Input packets**—Number of packets received on the interface. Same value as the physical interface. • **Output packets**—Number of bytes sent on the interface. Same value as the physical interface. **NOTE**: The bandwidth bps counter is not enabled on the switches.	**detail extensive**
Local statistics	Number and rate of bytes and packets received and transmitted locally by the Routing Engine on the logical interface of supported switches. All packets for protocols and process statistics are counted here. • **Input bytes**—Number of bytes received on the interface. Same value as for the physical interface. • **Output bytes**—Number of bytes sent on the interface. Same value as for the physical interface. • **Input packets**—Number of packets received on the interface. Same value as for the physical interface. • **Output packets**—Number of bytes sent on the interface. Same value as for the physical interface.	**detail extensive** none

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Transit statistics	Number and rate of bytes and packets received and transmitted on the RVI logical interface of supported switches. Look at this value to see the RVI ingress and egress count. • **Input bytes**—Number of bytes received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. • **Output bytes**—Number of bytes sent on the interface. This egress counter is automatic for EX8200. • **Input packets**—Number of packets received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. • **Output packets**—Number of packets sent on the interface. This egress counter is automatic for EX8200 switches.	**detail extensive**
IPv6 transit statistics	Number and rate of IPv6 bytes and packets received and transmitted on the RVI logical interface of supported switches. Transit values are unique to the logical interface and do not appear in physical interface output. Look at the values listed below to see the RVI ingress and egress count for IPv6 traffic. • **Input bytes**—Number of bytes received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. • **Output bytes**—Number of bytes sent by the interface. This egress counter is automatic for EX8200 switches. • **Input packets**—Number of packets received on the interface. This ingress counter is automatic for EX3200 and EX4200 and configurable for EX8200 switches. • **Output packets**—Number of packets sent by the interface. This egress counter is automatic for EX8200 switches. **NOTE**: The bandwidth bps counter is not enabled on the switches.	**detail extensive**

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Protocol	Protocol used for the logical interface—this value is **inet** for IPv4 traffic and **inet6** for IPv6 traffic.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Route table in which the logical interface address is located. For example, **0** refers to the routing table **inet.0**.	**detail extensive** none
Protocol flags	Information about the protocol such as **Targeted-broadcast**.	**detail extensive** none
Protocol addresses and Address flags	Protocol address values here can be: **Dest-route-down**—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status **Device-down**—Device has been administratively disabled. **Disabled**—Interface is administratively disabled. **Down**—A hardware failure has occurred. **Hardware-Down**—Interface protocol initialization failed to complete successfully. **Is-Default**—This address is the default address of the switch. The default address is used as the source address by SNMP, ping, traceroute, and other network utilities. **Is-Preferred**—This address is the default local address for packets originating from the local switch and sent to destinations on the subnet. **Is-Primary**—This address is the default local address for broadcast and multicast packets originated locally and sent out the interface. **Preferred**—This address is a candidate to become the preferred address. **Primary**—This address is a candidate to become the primary address. **SNMP-Traps**—SNMP trap notifications are enabled. **Up**—Interface is enabled and operational.	**detail extensive** none
Address destination	Logical destination's network address.	**detail extensive** none
Local address	IP address of the logical interface.	**detail extensive** none

Table 702: show interfaces vlan Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Broadcast address	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces vlan

```
user@switch> show interfaces vlan
Physical interface: vlan, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 556
  Type: VLAN, Link-level type: VLAN, MTU: 1518, Speed: 1000mbps
  Device flags   : Present Running
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface vlan.0 (Index 82) (SNMP ifIndex 557)
    Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 1
    Protocol inet
      Flags: Targeted-broadcast
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255

  Logical interface vlan.1 (Index 83) (SNMP ifIndex 558)
    Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 1
    Protocol inet
```

```
      Flags: Targeted-broadcast
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.2/24, Local: 10.1.2.1, Broadcast: 10.1.2.255
```

show interfaces vlan terse

```
user@switch> show interfaces vlan terse
Interface               Admin Link Proto    Local                 Remote
vlan                    up    up
vlan.0                  up    down inet     10.1.1.1/24
vlan.1                  up    down inet     10.1.2.1/24
```

show interfaces vlan extensive

```
user@switch> show interfaces vlan extensive
Physical interface: vlan, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 556, Generation: 153
  Type: VLAN, Link-level type: VLAN, MTU: 1518, Clocking: Unspecified,
  Speed: 1000mbps
  Device flags   : Present Running
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
```

```
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0

Logical interface vlan.0 (Index 82) (SNMP ifIndex 557) (Generation 147)
  Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                   42
   Input  packets:                    0
   Output packets:                    1
  Local statistics:
   Input  bytes  :                    0
   Output bytes  :                   42
   Input  packets:                    0
   Output packets:                    1
  Transit statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Protocol inet, Generation: 159, Route table: 0
    Flags: Targeted-broadcast
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
      Generation: 138

Logical interface vlan.1 (Index 83) (SNMP ifIndex 558) (Generation 148)
  Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                   42
   Input  packets:                    0
   Output packets:                    1
  Local statistics:
   Input  bytes  :                    0
   Output bytes  :                   42
   Input  packets:                    0
   Output packets:                    1
  Transit statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
```

```
     Input  packets:                    0                    0 pps
    Output packets:                    0                    0 pps
   Protocol inet, Generation: 160, Route table: 0
     Flags: Targeted-broadcast
     Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
       Destination: 10.1.2/24, Local: 10.1.2.1, Broadcast: 10.1.2.255,
       Generation: 140
```

show interfaces vlan detail

```
user@switch> show interfaces vlan detail
Physical interface: vlan, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 556, Generation: 153
  Type: VLAN, Link-level type: VLAN, MTU: 1518, Clocking: Unspecified,
  Speed: 1000mbps
  Device flags   : Present Running
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0

  Logical interface vlan.0 (Index 82) (SNMP ifIndex 557) (Generation 147)
    Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                   42
```

```
    Input  packets:                    0
    Output packets:                    1
  Local statistics:
    Input  bytes  :                    0
    Output bytes  :                   42
    Input  packets:                    0
    Output packets:                    1
  Transit statistics:
    Input  bytes  :                    0                    0 bps
    Output bytes  :                    0                    0 bps
    Input  packets:                    0                    0 pps
    Output packets:                    0                    0 pps
  Protocol inet, Generation: 159, Route table: 0
     Flags: Targeted-broadcast
     Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
       Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
       Generation: 138

Logical interface vlan.1 (Index 83) (SNMP ifIndex 558) (Generation 148)
  Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
    Input  bytes  :                    0
    Output bytes  :                   42
    Input  packets:                    0
    Output packets:                    1
  Local statistics:
    Input  bytes  :                    0
    Output bytes  :                   42
    Input  packets:                    0
    Output packets:                    1
  Transit statistics:
    Input  bytes  :                    0                    0 bps
    Output bytes  :                    0                    0 bps
    Input  packets:                    0                    0 pps
    Output packets:                    0                    0 pps
  Protocol inet, Generation: 160, Route table: 0
     Flags: Targeted-broadcast
     Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
       Destination: 10.1.2/24, Local: 10.1.2.1, Broadcast: 10.1.2.255,
       Generation: 140
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- show ethernet-switching table | **19700**
- show vlans
- Monitoring Interface Status and Traffic
- Troubleshooting Network Interfaces on EX3200 Switches
- *Troubleshooting Network Interfaces on EX4200 Switches*
- Verifying Routed VLAN Interface Status and Statistics on EX Series Switches

show interfaces voq

IN THIS SECTION

- Syntax | **21594**
- Syntax (Junos OS Evolved) | **21594**
- Description | **21594**
- Options | **21594**
- Additional Information | **21595**
- Required Privilege Level | **21595**
- Output Fields | **21595**
- Sample Output | **21596**
- Release Information | **21611**

Syntax

```
show interfaces voq interface-name
<forwarding-class forwarding-class-name>
<non-zero>
```

Syntax (Junos OS Evolved)

```
show interfaces voq interface-name
<forwarding-class forwarding-class-name>
<non-zero>
<source-fpc source-fpc-number>
```

Description

Display the random early detection (RED) drop statistics from all ingress Packet Forwarding Engines associated with the specified physical egress interface. In the VOQ architecture, egress output queues (shallow buffers) buffer data in virtual queues on ingress Packet Forwarding Engines. In cases of congestion, you can use this command to identify which ingress Packet Forwarding Engine is the source of RED-dropped packets contributing to congestion.

NOTE: On the PTX Series routers and QFX10000 switches, these statistics include tail-dropped packets.

Options

interface *interface-name*	Display the ingress VOQ RED drop statistics for the specified egress interface.
forwarding-class *forwarding-class-name*	Display VOQ RED drop statistics for a specified forwarding class.
non-zero	Display only non-zero VOQ RED drop statistics counters.
source-fpc *source-fpc-number*	Display VOQ RED drop statistics for the specified source FPC.

Additional Information

- On PTX Series routers, you can display VOQ statistics for only the WAN physical interface.
- VOQ statistics for aggregated physical interfaces are not supported. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can use the `show interfaces queue` command to identify the child link which is experiencing congestion and then view the VOQ statistics on the respective child link using the `show interfaces voq` command.

For information on virtual output queuing on PTX routers, see *Understanding Virtual Output Queues on PTX Series Packet Transport Routers*. For information on virtual output queueing on QFX10000 switches, see *Understanding CoS Virtual Output Queues (VOQs) on QFX10000 Switches*.

Required Privilege Level

view

Output Fields

Table 703 on page 21595 lists the output fields for the show interfaces queue command. Output fields are listed in the approximate order in which they appear.

Table 703: show interfaces voq Output Fields

Field Name	Field Description
`Physical interface`	Name of the physical interface.
`Enabled`	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.
`Interface index`	Physical interface's index number, which reflects its initialization sequence.
`SNMP ifIndex`	SNMP index number for the interface.
`Queue`	Egress queue number.

Table 703: show interfaces voq Output Fields *(Continued)*

Field Name	Field Description
`Forwarding classes`	Forwarding class name.
`FPC number`	Number of the Flexible PIC Concentrator (FPC) located on ingress.
`PFE`	Number of the Packet Forwarding Engine providing virtual output queues on the ingress.
`RED-dropped packets`	Number of packets per second (pps) dropped because of random early detection (RED). **NOTE**: On the PTX Series routers, these statistics include tail-dropped packets.
`RED-dropped bytes`	Number of bytes per second dropped because of RED. The byte counts vary by interface hardware. **NOTE**: On the PTX Series routers, these statistics include tail-dropped packets.

Sample Output

show interfaces voq (For a Specific Physical Interface) (PTX Series Routers)

The following example shows ingress RED-dropped statistics for the egress Ethernet interface configured on port 0 of Physical Interface Card (PIC) 0, located on the FPC in slot 7.

The sample output below shows that the cause of the congestion is ingress Packet Forwarding Engine PFE 0, which resides on FPC number 4, as denoted by the count of RED-dropped packets and RED-dropped bytes for egress queue 0, forwarding classes best-effort and egress queue 3, forwarding class network control.

```
user@host> show interfaces voq et-7/0/0
Physical interface: et-7/0/0, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 699
```

```
Queue: 0, Forwarding classes: best-effort

FPC number: 1
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 4
    PFE: 0
      RED-dropped packets  :             19969426              2323178 pps
      RED-dropped bytes    :           2196636860           2044397464 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 6
    PFE: 0
      RED-dropped packets  :             19969424              2321205 pps
      RED-dropped bytes    :           2196636640           2042660808 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
```

```
    PFE: 4
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 5
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 6
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 7
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 7
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

Queue: 1, Forwarding classes: expedited-forwarding

  FPC number: 1
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
```

```
FPC number: 4
  PFE: 0
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
  PFE: 1
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
  PFE: 2
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
  PFE: 3
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps

FPC number: 6
  PFE: 0
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
  PFE: 1
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
  PFE: 2
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
  PFE: 3
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
  PFE: 4
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
  PFE: 5
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
  PFE: 6
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
  PFE: 7
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps

FPC number: 7
  PFE: 0
    RED-dropped packets  :                    0                    0 pps
```

```
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

Queue: 2, Forwarding classes: assured-forwarding

  FPC number: 1
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 4
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 6
    PFE: 0
```

```
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 4
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 5
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 6
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 7
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 7
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

Queue: 3, Forwarding classes: network-control

  FPC number: 1
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
```

```
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 4
    PFE: 0
      RED-dropped packets  :             16338670              1900314 pps
      RED-dropped bytes    :           1797253700           1672276976 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 6
    PFE: 0
      RED-dropped packets  :             16338698              1899163 pps
      RED-dropped bytes    :           1797256780           1671263512 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 4
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 5
      RED-dropped packets  :                    0                    0 pps
```

```
      RED-dropped bytes    :                    0                    0 bps
    PFE: 6
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 7
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 7
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
```

show interfaces voq (For a Specific Physical Interface) (QFX10000 Switches)

The sample output below shows congestion on ingress PFE 1 on FPC number 0, and on ingress PFE 2 on FPC number 1, as denoted by the count of RED-dropped packets and RED-dropped bytes for best-effort egress queue 0.

```
user@host> show interfaces voq et-1/0/0
Physical interface: et-1/0/0, Enabled, Physical link is Up
  Interface index: 659, SNMP ifIndex: 539

Queue: 0, Forwarding classes: best-effort

  FPC number: 0
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :            411063248             16891870 pps
      RED-dropped bytes    :          52616095744          17297275600 bps
```

```
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 1
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :            411063012             16891870 pps
      RED-dropped bytes    :          52616065536          17297275376 bps

Queue: 3, Forwarding classes: fcoe

  FPC number: 0
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 1
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

Queue: 4, Forwarding classes: no-loss

  FPC number: 0
    PFE: 0
```

```
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 1
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

Queue: 7, Forwarding classes: network-control

  FPC number: 0
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 1
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
```

```
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
```

show interfaces voq et-7/0/0 (For a Specific Forwarding Class)

```
user@host> show interfaces voq et-7/0/0 forwarding-class best-effort
Physical interface: et-7/0/0, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 699

Queue: 0, Forwarding classes: best-effort

  FPC number: 1
    PFE: 0
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 4
    PFE: 0
      RED-dropped packets  :             66604786              2321519 pps
      RED-dropped bytes    :           7326526460           2042936776 bps
    PFE: 1
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 2
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps
    PFE: 3
      RED-dropped packets  :                    0                    0 pps
      RED-dropped bytes    :                    0                    0 bps

  FPC number: 6
    PFE: 0
```

```
      RED-dropped packets  :              66604794                371200 pps
      RED-dropped bytes    :            7326527340             326656000 bps
    PFE: 1
      RED-dropped packets  :                     0                     0 pps
      RED-dropped bytes    :                     0                     0 bps
    PFE: 2
      RED-dropped packets  :                     0                     0 pps
      RED-dropped bytes    :                     0                     0 bps
    PFE: 3
      RED-dropped packets  :                     0                     0 pps
      RED-dropped bytes    :                     0                     0 bps
    PFE: 4
      RED-dropped packets  :                     0                     0 pps
      RED-dropped bytes    :                     0                     0 bps
    PFE: 5
      RED-dropped packets  :                     0                     0 pps
      RED-dropped bytes    :                     0                     0 bps
    PFE: 6
      RED-dropped packets  :                     0                     0 pps
      RED-dropped bytes    :                     0                     0 bps
    PFE: 7
      RED-dropped packets  :                     0                     0 pps
      RED-dropped bytes    :                     0                     0 bps

  FPC number: 7
    PFE: 0
      RED-dropped packets  :                     0                     0 pps
      RED-dropped bytes    :                     0                     0 bps
    PFE: 1
      RED-dropped packets  :                     0                     0 pps
      RED-dropped bytes    :                     0                     0 bps
    PFE: 2
      RED-dropped packets  :                     0                     0 pps
      RED-dropped bytes    :                     0                     0 bps
    PFE: 3
      RED-dropped packets  :                     0                     0 pps
      RED-dropped bytes    :                     0                     0 bps
```

show interfaces voq et-5/0/12 (For a Specific Source FPC)

```
user@host> show interfaces voq et-5/0/12 source-fpc 0
Physical interface: et-5/0/12, Enabled, Physical link is Up
          Interface index: 166, SNMP ifIndex: 1104

        Queue: 0, Forwarding classes: best-effort

          FPC number: 0
            PFE: 0
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 1
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 2
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 3
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps

        Queue: 1, Forwarding classes: expedited-forwarding

          FPC number: 0
            PFE: 0
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 1
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 2
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 3
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps

        Queue: 2, Forwarding classes: assured-forwarding

          FPC number: 0
            PFE: 0
```

```
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 1
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 2
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 3
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps

        Queue: 3, Forwarding classes: network-control

          FPC number: 0
            PFE: 0
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 1
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 2
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 3
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
```

show interfaces voq et-5/0/12 (For a Specific Forwarding Class and Source FPC)

```
user@host> show interfaces voq et-5/0/12 forwarding-class best-effort source-fpc 5
Physical interface: et-5/0/12, Enabled, Physical link is Up
          Interface index: 166, SNMP ifIndex: 1104

        Queue: 0, Forwarding classes: best-effort

          FPC number: 5
            PFE: 0
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 1
```

```
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 2
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 3
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 4
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 5
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 6
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
            PFE: 7
              RED-dropped packets  :                    0                    0 pps
              RED-dropped bytes    :                    0                    0 bps
```

show interfaces voq et-7/0/0 (Non-Zero)

```
user@host> show interfaces voq et-7/0/0 non-zero

Physical interface: et-7/0/0, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 699

Queue: 0, Forwarding classes: best-effort

  FPC number: 4
    PFE: 0
      RED-dropped packets  :              95862238               2301586 pps
      RED-dropped bytes    :           10544846180            2025396264 bps

  FPC number: 6
    PFE: 0
      RED-dropped packets  :              95866639               2322569 pps
      RED-dropped bytes    :           10545330290            2043860728 bps
```

```
Queue: 3, Forwarding classes: network-control

  FPC number: 4
    PFE: 0
      RED-dropped packets  :              78433066               1899727 pps
      RED-dropped bytes    :            8627637260            1671760384 bps

  FPC number: 6
    PFE: 0
      RED-dropped packets  :              78436704               1900628 pps
      RED-dropped bytes    :            8628037440            1672553432 bps
```

show interfaces voq et-7/0/0 (For a Specific Forwarding Class and Non-Zero)

```
user@host show interfaces voq et-7/0/0 forwarding-class best-effort non-zero
Physical interface: et-7/0/0, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 699

Queue: 0, Forwarding classes: best-effort

  FPC number: 4
    PFE: 0
      RED-dropped packets  :             119540012               2322319 pps
      RED-dropped bytes    :           13149401320            2043640784 bps

  FPC number: 6
    PFE: 0
      RED-dropped packets  :             119540049               2322988 pps
      RED-dropped bytes    :           13149405390            2044229744 bps
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Virtual Output Queues on PTX Series Packet Transport Routers

Understanding CoS Virtual Output Queues (VOQs) on QFX10000 Switches

show interfaces xe

IN THIS SECTION

- Syntax (QFX Series) | 21612
- Syntax (EX Series) | 21612
- Description | 21613
- Options | 21613
- Required Privilege Level | 21614
- Output Fields | 21614
- Sample Output | 21627
- Release Information | 21639

Syntax (QFX Series)

```
show interfaces device-name:type-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
<statistics>
```

Syntax (EX Series)

```
show interfaces xe-fpc/pic/port
<brief | detail | extensive | terse>
<media>
<statistics>
```

Description

Display status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.

> **NOTE**: You must have a transceiver plugged into an SFP+ or an XFP port before information about the interface can be displayed.

> **NOTE**: On an EX Series switch, the traffic statistics for a LAG might vary slightly from the cumulative traffic statistics of the member interfaces of the LAG. This difference is more likely to be seen when the traffic is bursty in nature, and because the statistics are not fetched from the LAG and the members in the same instant. For accurate traffic statistics for a LAG, use the aggregated Ethernet counters.

Options

device-name:type-fpc/pic/port	(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.
brief | detail | extensive | terse	(Optional) (QFX Series) Display the specified level of output.
descriptions	(Optional) (QFX Series) Display interface description strings.
media	(Optional) (QFX Series) Display media-specific information about network interfaces.
routing-instance (all | *instance-name*)	(Optional) (QFX Series) Display the name of an individual routing instance or display all routing instances.
snmp-index *snmp-index*	(Optional) (QFX Series) Display information for the specified SNMP index of the interface.
statistics	(Optional) (QFX Series) Display static interface statistics.
xe-*fpc/pic/port*	(EX Series) Display standard information about the specified 10-Gigabit Ethernet interface.
brief | detail | extensive | terse	(Optional) (EX Series) Display the specified level of output.

media (Optional) (EX Series) Display media-specific information about network interfaces. For 10-Gigabit Ethernet interfaces, using the media option does not provide you with new or additional information. The output is the same as when the media option is not used.

statistics (Optional) (EX Series) Display static interface statistics. For 10-Gigabit Ethernet interfaces, using the statistics option does not provide you with new or additional information. The output is the same as when the statistics option is not used.

Required Privilege Level

view

Output Fields

Table 704 on page 21614 lists the output fields for the `show interfaces xe` command. Output fields are listed in the approximate order in which they appear.

Table 704: show interfaces xe Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface.	All levels

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Duplex	Duplex mode of the interface, either **Full-Duplex** or **Half-Duplex**.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: **Enabled** or **Disabled** for parent interface; **Rx-only** or **Tx-only** for child interfaces.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**. **NOTE**: This field is only displayed if asymmetric flow control is not configured.	All levels

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Configured-flow- control	Configured flow control for the interface transmit buffers (**tx-buffers**) and receive buffers (**rx-buffers**): • **tx-buffers**—**On** if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer. **Off** if the interface is not configured to respond to received PAUSE messages. • **rx-buffers**—**On** if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer. **Off** if the interface is not configured to generate and send PAUSE messages. **NOTE**: This field is only displayed if asymmetric flow control is configured.	All levels
Auto-negotiation	Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the link.	All levels
Wavelength	Configured wavelength, in nanometers (nm).	All levels
Frequency	Frequency associated with the configured wavelength, in terahertz (THz).	All levels

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
CoS queues	Number of CoS queues configured.	**detail extensive** none
Schedulers	Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour: :minute.second.timezone*** **(*hour.minute.second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	**detail extensive**

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame terminated and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the `ignore-l3-incompletes` statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame terminated and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Queue Number	The CoS queue number and the forwarding classes mapped to the queue number. The **Mapped forwarding class** column lists the forwarding classes mapped to each CoS queue.	**detail extensive**
Ingress queues	Total number of ingress queues supported on the specified interface.	**extensive**
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
PCS statistics	Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.	**detail extensive**

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem. • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of packets that exceeds the configured MTU. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.	**extensive**

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem.	**extensive**

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  • **Incomplete**—Ethernet interface has the speed or link mode configured.
	  • **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner:**
	  • **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is **None**. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  • **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**:
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and
		extensive

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive). For asymmetric **PAUSE**, shows if the PAUSE transmit and **PAUSE** receive states on the interface are **enable** or **disable**. • **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).	
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	**detail extensive**
IPv6 transit statistics	If IPv6 statics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	**extensive**
Local statistics	Number and rate of bytes and packets destined to and from the switch.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch.	**extensive**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface.	**detail extensive**
Output Filters	Names of any output filters applied to this interface.	**detail extensive**

Table 704: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Information about protocol family flags. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.	**detail extensive**
Addresses, Flags	Information about the address flags.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about the address flag.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces

```
user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
```

```
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch, MTU: 0
      Flags: Trunk-Mode
```

show interfaces (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Configured-flow-control  tx-buffers: off  rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
```

```
    Output packets: 0
    Protocol eth-switch, MTU: 0
      Flags: Trunk-Mode
```

show interfaces brief

```
user@switch> show interfaces xe-0/0/1 brief
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None

  Logical interface xe-0/0/1.0
    Flags: SNMP-Traps Encapsulation: ENET2
    eth-switch
```

show interfaces detail

```
user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:50 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
```

```
 IPv6 transit statistics:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Egress queues: 12 supported, 9 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                    0                    0                    0
  1 fc7                            0                    0                    0
  2 no-loss                        0                    0                    0
  3 fcoe                           0                    0                    0
  4 fc4                            0                    0                    0
  5 fc5                            0                    0                    0
  6 fc6                            0                    0                    0
  7 network-cont                   0                    0                    0
  8 mcast                          0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   fc7
  2                   no-loss
  3                   fcoe
  4                   fc4
  5                   fc5
  6                   fc6
  7                   network-control
  8                   mcast
Active alarms  : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Local statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Transit statistics:
   Input  bytes  :                   0                    0 bps
```

```
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
      Flags: Trunk-Mode
```

show interfaces detail (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Configured-flow-control  tx-buffers: off  rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:50 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Egress queues: 12 supported, 9 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 fc7                            0                    0                    0
    2 no-loss                        0                    0                    0
    3 fcoe                           0                    0                    0
    4 fc4                            0                    0                    0
    5 fc5                            0                    0                    0
    6 fc6                            0                    0                    0
```

```
    7 network-cont                  0                    0                    0
    8 mcast                         0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   fc7
    2                   no-loss
    3                   fcoe
    4                   fc4
    5                   fc5
    6                   fc6
    7                   network-control
    8                   mcast
  Active alarms  : None
  Active defects : None

  Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
      Flags: Trunk-Mode
```

show interfaces extensive

```
user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
```

```
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:03:08 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 12 supported, 9 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 fc7                            0                    0                    0
    2 no-loss                        0                    0                    0
    3 fcoe                           0                    0                    0
    4 fc4                            0                    0                    0
    5 fc5                            0                    0                    0
    6 fc6                            0                    0                    0
    7 network-cont                   0                    0                    0
    8 mcast                          0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   fc7
    2                   no-loss
    3                   fcoe
    4                   fc4
```

```
  5                   fc5
  6                   fc6
  7                   network-control
  8                   mcast
Active alarms  : None
Active defects : None
MAC statistics:                       Receive         Transmit
  Total octets                              0                0
  Total packets                             0                0
  Unicast packets                           0                0
  Broadcast packets                         0                0
  Multicast packets                         0                0
  CRC/Align errors                          0                0
  FIFO errors                               0                0
  MAC control frames                        0                0
  MAC pause frames                          0                0
  Oversized frames                          0
  Jabber frames                             0
  Fragment frames                           0
  VLAN tagged frames                        0
  Code violations                           0
MAC Priority Flow Control Statistics:
  Prioity :  0                              0                0
  Prioity :  1                              0                0
  Prioity :  2                              0                0
  Prioity :  3                              0                0
  Prioity :  4                              0                0
  Prioity :  5                              0                0
  Prioity :  6                              0                0
  Prioity :  7                              0                0
Filter statistics:
  Input packet count                        0
  Input packet rejects                      0
  Input DA rejects                          0
  Input SA rejects                          0
  Output packet count                                        0
  Output packet pad count                                    0
  Output packet error count                                  0
  CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
```

```
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 best-effort            75     7500000000    75              0      low    none
    7 network-control         5      500000000     5              0      low    none
    8 mcast                  20     2000000000    20              0      low    none

  Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
      Flags: Trunk-Mode
```

show interfaces extensive (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Configured-flow-control  tx-buffers: off  rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:03:08 ago)
```

```
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 12 supported, 9 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 fc7                            0                    0                    0
    2 no-loss                        0                    0                    0
    3 fcoe                           0                    0                    0
    4 fc4                            0                    0                    0
    5 fc5                            0                    0                    0
    6 fc6                            0                    0                    0
    7 network-cont                   0                    0                    0
    8 mcast                          0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   fc7
    2                   no-loss
    3                   fcoe
    4                   fc4
    5                   fc5
    6                   fc6
    7                   network-control
    8                   mcast
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
```

```
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
MAC Priority Flow Control Statistics:
  Prioity :  0                             0                0
  Prioity :  1                             0                0
  Prioity :  2                             0                0
  Prioity :  3                             0                0
  Prioity :  4                             0                0
  Prioity :  5                             0                0
  Prioity :  6                             0                0
  Prioity :  7                             0                0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue             Bandwidth           Buffer Priority   Limit
                          %            bps     %       usec
  0 best-effort          75     7500000000    75          0      low    none
  7 network-control       5      500000000     5          0      low    none
  8 mcast                20     2000000000    20          0      low    none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
```

```
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
      Flags: Trunk-Mode
```

show interfaces terse

```
user@switch> show interfaces xe-0/0/1 terse
Interface               Admin Link Proto    Local                 Remote

xe-0/0/1               up    up
xe-0/0/1.0             up    up   eth-switch
```

show interfaces (QFabric System)

```
user@switch> show interfaces node1:xe-0/0/0
Physical interface: node1:xe-0/0/0, Enabled, Physical link is Down
  Interface index: 129, SNMP ifIndex: 2884086
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
  Interface flags: Internal: 0x4000
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
```

```
Output rate     : 0 bps (0 pps)
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- *Monitoring Interface Status and Traffic*
- Monitoring Interface Status and Traffic
- *Troubleshooting Network Interfaces*
- *Troubleshooting an Aggregated Ethernet Interface*
- Junos OS Network Interfaces Library for Routing Devices
- Troubleshooting Network Interfaces on EX3200 Switches
- *Troubleshooting Network Interfaces on EX4200 Switches*
- *Troubleshooting an Aggregated Ethernet Interface*
- Junos OS Ethernet Interfaces Configuration Guide

show interfaces xle

IN THIS SECTION

- Syntax | 21640
- Description | 21640
- Options | 21640
- Required Privilege Level | 21640
- Output Fields | 21641
- Sample Output | 21654
- Release Information | 21666

Syntax

```
show interfaces device-name:type-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
<statistics>
```

Description

Display status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.

Options

device-name:type-fpc/pic/port	(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
routing-instance (all | *instance-name*)	(Optional) Display the name of an individual routing instance or display all routing instances.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 705 on page 21641 lists the output fields for the **show interfaces xe** command. Output fields are listed in the approximate order in which they appear.

Table 705: show interfaces xe Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Duplex	Duplex mode of the interface, either **Full-Duplex** or **Half-Duplex**.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: **Enabled** or **Disabled** for parent interface; **Rx-only** or **Tx-only** for child interfaces.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**. **NOTE**: This field is only displayed if asymmetric flow control is not configured.	All levels
Configured-flow- control	Configured flow control for the interface transmit buffers (**tx-buffers**) and receive buffers (**rx-buffers**): • **tx-buffers**—**On** if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer. **Off** if the interface is not configured to respond to received PAUSE messages. • **rx-buffers**—**On** if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer. **Off** if the interface is not configured to generate and send PAUSE messages. **NOTE**: This field is only displayed if asymmetric flow control is configured.	All levels
Auto-negotiation	Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the link.	All levels
Wavelength	Configured wavelength, in nanometers (nm).	All levels
Frequency	Frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Schedulers	Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour: :minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. **NOTE**: The bandwidth bps counter is not enabled.	**detail extensive**

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame terminated and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame terminated and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Queue Number	The CoS queue number and the forwarding classes mapped to the queue number. The **Mapped forwarding class** column lists the forwarding classes mapped to each CoS queue.	**detail extensive**
Ingress queues	Total number of ingress queues supported on the specified interface.	**extensive**
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
PCS statistics	Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.	**detail extensive**

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem. • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of packets that exceeds the configured MTU. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.	**extensive**

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem.	**extensive**

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  • **Incomplete**—Ethernet interface has the speed or link mode configured.
	  • **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner:**
	  • **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is **None**. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  • **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**:
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and
		extensive

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive). For asymmetric **PAUSE**, shows if the PAUSE transmit and **PAUSE** receive states on the interface are **enable** or **disable**. • **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).	
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	**detail extensive**
IPv6 transit statistics	If IPv6 statics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	**extensive**
Local statistics	Number and rate of bytes and packets destined to and from the switch.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch.	**extensive**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface.	**detail extensive**
Output Filters	Names of any output filters applied to this interface.	**detail extensive**

Table 705: show interfaces xe Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Information about protocol family flags. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.	**detail extensive**
Addresses, Flags	Information about the address flags.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about the address flag.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces

```
user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
```

```
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch, MTU: 0
      Flags: Trunk-Mode
```

show interfaces (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Configured-flow-control  tx-buffers: off  rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
```

```
    Output packets: 0
    Protocol eth-switch, MTU: 0
      Flags: Trunk-Mode
```

show interfaces brief

```
user@switch> show interfaces xe-0/0/1 brief
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None

  Logical interface xe-0/0/1.0
    Flags: SNMP-Traps Encapsulation: ENET2
    eth-switch
```

show interfaces detail

```
user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:50 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
```

```
 IPv6 transit statistics:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Egress queues: 12 supported, 9 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                    0                    0                    0
  1 fc7                            0                    0                    0
  2 no-loss                        0                    0                    0
  3 fcoe                           0                    0                    0
  4 fc4                            0                    0                    0
  5 fc5                            0                    0                    0
  6 fc6                            0                    0                    0
  7 network-cont                   0                    0                    0
  8 mcast                          0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   fc7
  2                   no-loss
  3                   fcoe
  4                   fc4
  5                   fc5
  6                   fc6
  7                   network-control
  8                   mcast
Active alarms  : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Local statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Transit statistics:
   Input  bytes  :                   0                    0 bps
```

```
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
      Flags: Trunk-Mode
```

show interfaces detail (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Configured-flow-control  tx-buffers: off  rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:50 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Egress queues: 12 supported, 9 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 fc7                            0                    0                    0
    2 no-loss                        0                    0                    0
    3 fcoe                           0                    0                    0
    4 fc4                            0                    0                    0
    5 fc5                            0                    0                    0
    6 fc6                            0                    0                    0
```

```
    7 network-cont                  0                    0                    0
    8 mcast                         0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   fc7
    2                   no-loss
    3                   fcoe
    4                   fc4
    5                   fc5
    6                   fc6
    7                   network-control
    8                   mcast
  Active alarms  : None
  Active defects : None

  Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
      Flags: Trunk-Mode
```

show interfaces extensive

```
user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
```

```
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:03:08 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 12 supported, 9 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 fc7                            0                    0                    0
    2 no-loss                        0                    0                    0
    3 fcoe                           0                    0                    0
    4 fc4                            0                    0                    0
    5 fc5                            0                    0                    0
    6 fc6                            0                    0                    0
    7 network-cont                   0                    0                    0
    8 mcast                          0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   fc7
    2                   no-loss
    3                   fcoe
    4                   fc4
```

```
  5                   fc5
  6                   fc6
  7                   network-control
  8                   mcast
Active alarms  : None
Active defects : None
MAC statistics:                       Receive         Transmit
  Total octets                              0                0
  Total packets                             0                0
  Unicast packets                           0                0
  Broadcast packets                         0                0
  Multicast packets                         0                0
  CRC/Align errors                          0                0
  FIFO errors                               0                0
  MAC control frames                        0                0
  MAC pause frames                          0                0
  Oversized frames                          0
  Jabber frames                             0
  Fragment frames                           0
  VLAN tagged frames                        0
  Code violations                           0
MAC Priority Flow Control Statistics:
  Prioity :  0                              0                0
  Prioity :  1                              0                0
  Prioity :  2                              0                0
  Prioity :  3                              0                0
  Prioity :  4                              0                0
  Prioity :  5                              0                0
  Prioity :  6                              0                0
  Prioity :  7                              0                0
Filter statistics:
  Input packet count                        0
  Input packet rejects                      0
  Input DA rejects                          0
  Input SA rejects                          0
  Output packet count                                        0
  Output packet pad count                                    0
  Output packet error count                                  0
  CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
```

```
    CoS transmit queue               Bandwidth               Buffer Priority   Limit
                              %            bps     %           usec
    0 best-effort            75     7500000000    75              0      low    none
    7 network-control         5      500000000     5              0      low    none
    8 mcast                  20     2000000000    20              0      low    none

  Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
      Flags: Trunk-Mode
```

show interfaces extensive (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Configured-flow-control  tx-buffers: off  rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:03:08 ago)
```

```
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 12 supported, 9 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 fc7                            0                    0                    0
    2 no-loss                        0                    0                    0
    3 fcoe                           0                    0                    0
    4 fc4                            0                    0                    0
    5 fc5                            0                    0                    0
    6 fc6                            0                    0                    0
    7 network-cont                   0                    0                    0
    8 mcast                          0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   fc7
    2                   no-loss
    3                   fcoe
    4                   fc4
    5                   fc5
    6                   fc6
    7                   network-control
    8                   mcast
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
```

```
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
MAC Priority Flow Control Statistics:
  Prioity :  0                             0                0
  Prioity :  1                             0                0
  Prioity :  2                             0                0
  Prioity :  3                             0                0
  Prioity :  4                             0                0
  Prioity :  5                             0                0
  Prioity :  6                             0                0
  Prioity :  7                             0                0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue             Bandwidth           Buffer Priority   Limit
                          %            bps     %       usec
  0 best-effort          75     7500000000    75          0      low    none
  7 network-control       5      500000000     5          0      low    none
  8 mcast                20     2000000000    20          0      low    none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
```

```
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
      Flags: Trunk-Mode
```

show interfaces terse

```
user@switch> show interfaces xe-0/0/1 terse
Interface               Admin Link Proto    Local                 Remote

xe-0/0/1               up    up
xe-0/0/1.0             up    up   eth-switch
```

show interfaces (QFabric System)

```
user@switch> show interfaces node1:xe-0/0/0
Physical interface: node1:xe-0/0/0, Enabled, Physical link is Down
  Interface index: 129, SNMP ifIndex: 2884086
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU Error: None,
MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
  Interface flags: Internal: 0x4000
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
```

```
Output rate     : 0 bps (0 pps)
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Monitoring Interface Status and Traffic

Troubleshooting Network Interfaces

Troubleshooting an Aggregated Ethernet Interface

Junos OS Network Interfaces Library for Routing Devices

show ip-source-guard

IN THIS SECTION

- Syntax | 21666
- Description | 21667
- Required Privilege Level | 21667
- Output Fields | 21667
- Sample Output | 21668
- Release Information | 21668

Syntax

```
show ip-source-guard
```

Description

Display IP source guard database information.

Required Privilege Level

view

Output Fields

Table 706 on page 21667 lists the output fields for the `show ip-source-guard` command. Output fields are listed in the approximate order in which they appear.

Table 706: show ip-source-guard Output Fields

Field Name	Field Description
VLAN	VLAN on which IP source guard is enabled.
Interface	Access interface associated with the VLAN in column 1.
Tag	VLAN ID for the VLAN in column 1. Possible values are: • 0, indicating the VLAN is not tagged. • 1 – 4093
IP Address	Source IP address for a device connected to the interface in column 2. A value of * (star, or asterisk) indicates that IP source guard is not enabled on this VLAN but the interface is shared with a VLAN that is enabled for IP source guard.
MAC Address	Source MAC address for a device connected to the interface in column 2. A value of * (star, or asterisk) indicates that IP source guard is not enabled on this VLAN but the interface is shared with a VLAN that is enabled for IP source guard.

Sample Output

show ip-source-guard

```
user@switch> show ip-source-guard
IP source guard information:
Interface    Tag  IP Address   MAC Address        VLAN

ge-0/0/12.0  0    10.10.10.7   00:30:48:92:A5:9D  vlan100

ge-0/0/13.0  0    10.10.10.9   00:30:48:8D:01:3D  vlan100

ge-0/0/13.0  100  *            *                  voice
```

Release Information

Command introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN

Example: Configuring IP Source Guard with Other EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces

Verifying That IP Source Guard Is Working Correctly

show ipsec certificates

IN THIS SECTION

- Syntax | 21669
- Description | 21669
- Options | 21669
- Required Privilege Level | 21669
- Output Fields | 21669

Sample Output | 21671
Release Information | 21672

Syntax

```
show ipsec certificates
<brief | detail>
<crl crl-name | serial-number>
```

Description

(Encryption interface on M Series and T Series routers only) Display information about the IPsec certificate database.

Options

none	Display standard information about all of the entries in the IPsec certificate database.
brief | detail	(Optional) Display the specified level of output.
crl ***crl-name* | *serial-number***	(Optional) Display information about the entries on the certificate revocation list (CRL) or for the specified serial number. A CRL is a timestamped list identifying revoked certificates. The CRL is signed by a certificate authority (CA) or CRL issuer and made freely available in a public repository. Each revoked certificate is identified in a CRL by its certificate serial number.

Required Privilege Level

view

Output Fields

Table 707 on page 21670 lists the output fields for the `show ipsec certificates` command. Output fields are listed in the approximate order in which they appear.

Table 707: show ipsec certificates Output Fields

Field Name	Field Description	Level of Output
Database	Display information about the IPsec certificate database. • **Total entries**—Number of database entries, including entries that are not trusted or that are in the process of being deleted. • **Active entries**—Number of database entries, excluding entries that are marked as deleted. • **Locked entries**—Number of statically configured database entries that cannot expire, such as CA certificates that are root or trusted.	All levels
Subject	Distinguished name for the certificate for **C**, **O**, **CN**, as described in RFC 3280, *Internet x.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile.*	All levels
ID	Identification number of the database entry. **ID** is generated by the internal certificate database.	All levels
References	Reference number the certificate manager has for the particular entry.	**detail**
Serial	Unique serial number assigned to each certificate by the CA.	All levels
Flags	State of the certificate. • **Trusted**—Passed validity checks. • **Not trusted**—Failed validity checks. • **Root**—Entry is locked and may have been learned through IKE or a locally configured CA certificate. • **Non-root**—Entry is not locked. • **Crl-issuer**—Entity issues CRLs. • **Non-crl-issuer**—Entity does not issue CRLs.	**detail**

Table 707: show ipsec certificates Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Validity period starts	Start time that the certificate is valid, in the format ***yyyy mon dd, hh:mm:ss*** **GMT**.	**detail**
Validity period ends	End time that the certificate is valid, in the format ***yyyy mon dd, hh:mm:ss*** **GMT**.	**detail**
Alternative name information	Auxiliary identity for the certificate: ***dns-name***, ***email-address***, ***ip-address***, or ***uri*** (uniform resource identifier).	**detail**
Issuer	Information about the entity that has signed and issued the CRL as described in RFC 2459, *Internet X.509 Public Key Infrastructure Certificate and CRL Profile.*	**detail**

Sample Output

show ipsec certificates detail

```
user@host> show ipsec certificates detail
Database: Total entries: 3 Active entries: 4 Locked entries: 1
Subject: C=us, O=x
  ID: 5, References: 0, Serial: 22314868
  Flags: Trusted Non-root Crl-issuer
  Validity period starts: 2003 Mar  1st, 01:20:42 GMT
  Validity period ends: 2003 Mar 31st, 01:50:42 GMT
  Alternative name information:
    IP address: 10.20.210.1
  Issuer: C=FI, O=Company-ABC, CN=Company ABC class 2

Subject: C=us, O=x
  ID: 4, References: 0, Serial: 22315496
  Flags: Trusted Non-root Crl-issuer
  Validity period starts: 2003 Mar  1st, 01:21:45 GMT
  Validity period ends: 2003 Mar 31st, 01:51:45 GMT
  Alternative name information:
    IP address: 10.20.210.20
```

```
  Issuer: C=FI, O=Company-ABC, CN=Company ABC class 2

Subject: C=FI, O=SSH Company-ABC, CN=Company ABC class 2
  ID: 1, References: 1, Serial: 1538512
  Flags: Trusted Root Non-crl-issuer
  Validity period starts: 2001 Aug  1st, 07:08:32 GMT
  Validity period ends: 2004 Aug  1st, 07:08:32 GMT
  Alternative name information:
    Email address: certifier-support@ssh.com
  Issuer: C=FI, O=Company-ABC, CN=Company ABC class 2
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear ipsec security-associations

show ipsec redundancy

IN THIS SECTION

- Syntax | 21673
- Description | 21673
- Options | 21673
- Required Privilege Level | 21673
- Output Fields | 21673
- Sample Output | 21674
- Release Information | 21675

Syntax

```
show ipsec redundancy (interface <es-fpc/pic/port> | security-associations <sa-name>)
```

Description

(Encryption interface on M Series and T Series routers only) Display information about IPsec redundancy.

Options

interface <es-*fpc/pic/port*>	Display information about all encryption interfaces, or optionally, about a particular encryption interface.
security-associations <*sa-name*>	Display information about all remote tunnels, or optionally, about a particular remote tunnel.

Required Privilege Level

view

Output Fields

Table 708 on page 21673 lists the output fields for the `show ipsec redundancy` command. Output fields are listed in the approximate order in which they appear.

Table 708: show ipsec redundancy Output Fields

Field Name	Field Description
Failure counter	Number of times a PIC switched between primary and backup interfaces, or the number of times the tunnel switched between the primary and remote peers since the software has been activated.
Primary interface '	Name of the interface configured to be the primary interface.
Backup interface	Name of the interface configured to be the backup interface.

Table 708: show ipsec redundancy Output Fields *(Continued)*

Field Name	Field Description
State	State of the primary or backup interface can be **Active**, **Offline**, or **Standby**. Both ES PICs are initialized to **Offline**. For primary and remote peers, **State** can be **Active** or **Standby**. Both peers are in a state of **Standby** by default (there is not yet a connection between the two peers).
Security association	Name of the security association.
Local IP	Local IP address.
Primary remote IP	IP address of the configured primary remote peer.
Backup remote IP	IP address of the configured backup remote peer.

Sample Output

show ipsec redundancy interface

```
user@host> show ipsec redundancy interface
    Failure counter: 0
    Primary interface: es-1/3/0, State: Active
    Backup interface : es-1/1/0, State: Standby
```

show ipsec redundancy security-associations

```
user@host> show ipsec redundancy security-associations sa-dynamic
   Security association: sa-dynamic, Failure counter: 0
   Local IP: 192.0.2.4
   Primary remote IP: 198.51.100.5, State: Standby
   Backup remote IP : 192.0.2.3, State: Standby
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

request ipsec switch | 14665

show ipsec security-associations

IN THIS SECTION

- Syntax | 21675
- Description | 21675
- Options | 21676
- Required Privilege Level | 21676
- Output Fields | 21676
- Sample Output | 21679
- Release Information | 21679

Syntax

```
show ipsec security-associations
<brief | detail>
<sa-name>
```

Description

Display information about the IPsec security associations applied to the local or transit traffic stream.

Options

none	Display standard information about all IPsec security associations.
brief | detail	(Optional) Display the specified level of output.
sa-name	(Optional) Display the specified IPsec security association.

Required Privilege Level

view

Output Fields

Table 709 on page 21676 lists the output fields for the **show ipsec security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 709: show ipsec security-associations Output Fields

Field Name	Field Description	Level of Output
Security association	Name of the security association.	All levels
Interface family	Status of the interface family of the security association. If the interface family field is absent, it is a transport mode security association. The interface family can have one of three options: • **Up**—The security association is referenced in the interface family and the interface family is up. • **Down**—The security association is referenced in the interface family and the interface family is down. • **No reference**—The security association is not referenced in the interface family.	All levels
Local gateway	Gateway address of the local system.	All levels
Remote gateway	Gateway address of the remote system.	All levels

Table 709: show ipsec security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Local identity	Prefix and port number of the local end	All levels
Remote identity	Prefix and port number of the remote end.	All levels
Direction	Direction of the security association: **inbound** or **outbound**.	All levels
SPI	Value of the security parameter index.	All levels
AUX-SPI	Value of the auxiliary security parameter index. • When the value is **AH** or **ESP**, **AUX-SPI** is always **0**. • When the value is**AH+ESP**, **AUX-SPI** is always a positive integer.	All levels
State	Status of the security association: • **Installed**—The security association is installed in the security association database. (For transport mode security associations, the value of **State** must always be **Installed**.) • **Not installed**—The security association is not installed in the security association database.	**detail**
Mode	Mode of the security association: • **transport**—Protects single host-to-host protections. • **tunnel**—Protects connections between security gateways.	All levels
Type	Type of security association:. • **manual**—Security parameters require no negotiation. They are static, and are configured by the user. • **dynamic**—Security parameters are negotiated by the IKE protocol. Dynamic security associations are not supported in transport mode.	All levels

Table 709: show ipsec security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Protocol	Protocol supported: • **transport** mode—Supports Encapsulation Security Protocol (**ESP**) or Authentication Header (**AH**). • **tunnel** mode—Supports **ESP** or **AH+ESP**.	All levels
Authentication	Type of authentication used: **hmac-md5-96**, **hmac-sha1-96**, or **None**.	**detail**
Encryption	Type of encryption used: **des-cbc**, **3des-csc**, or **None**.	**detail**
Soft lifetime **Hard lifetime**	(**dynamic** output only) Each lifetime of a security association has two display options, hard and soft, one of which must be present for a dynamic security association. The **hard lifetime** specifies the lifetime of the SA. The **soft lifetime**, which is derived from the hard lifetime, informs the IPsec key management system that the SA is about to expire. This allows the key management system to negotiate a new SA before the hard lifetime expires. • **Expires in** *seconds* **seconds**—Number of seconds left until the security association expires. • **Expires in** *kilobytes* **kilobytes**—Number of kilobytes left until the security association expires.	**detail**
Anti-replay service	State of the service that prevents packets from being replayed: **Enabled** or **Disabled**.	**detail**
Replay window size	Configured size, in packets, of the antireplay service window: **32** or **64**. The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets. If the replay window size is **0**, the antireplay service is **disabled**.	**detail**

Sample Output

show ipsec security-associations sa-name

```
user@host> show ipsec security-associations sa-cosmic brief
Security association: sa-cosmic, Interface family: Up
Local gateway: 192.0.2.1, Remote gateway: 198.51.100.1
Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Direction SPI         AUX-SPI     Mode       Type     Protocol
inbound   2908734119  0           tunnel     dynamic  AH
outbound  3494029335  0           tunnel     dynamic  AH
```

show ipsec security-associations sa-name detail

```
user@host> show ipsec security-associations sa-cosmic detail
Security association: sa-cosmic, Interface family: Up

Local gateway: 192.0.2.1, Remote gateway: 198.51.100.1
Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Direction: inbound, SPI: 2908734119, AUX-SPI: 0, State: Installed
Mode: tunnel, Type: dynamic
Protocol: AH, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expired
Hard lifetime: Expires in 120 seconds
Anti-replay service: Disabled

Direction: outbound, SPI: 3494029335, AUX-SPI: 0, State: Installed
Mode: tunnel, Type: dynamic
Protocol: AH, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expired
Hard lifetime: Expires in 120 seconds
Anti-replay service: Disabled
```

Release Information

Command introduced before Junos OS Release 7.4.

show ipv6 neighbors

IN THIS SECTION

- Syntax | 21680
- Description | 21680
- Options | 21681
- Required Privilege Level | 21681
- Output Fields | 21681
- Sample Output | 21682
- Release Information | 21683

Syntax

```
show ipv6 neighbors
<flags>
<hostname host-name>
<interface interface-name>
<logical-system logical-system-name>
<reference-count count>
<tenant name>
<vpn vpn-name>
```

Description

Display information about the IPv6 neighbor cache.

NOTE: Starting with Junos OS Release 16.1, `show ipv6 neighbors` command does not display the underlying ifl information if `enhanced-convergence` statement at [`edit irb unit` *`unit-number`*] hierarchy level and `enhanced-ip` statement at [`edit chassis network-services`] hierarchy level is configured for the destination interface IRB.

Options

none	Display the entries in the IPv6 table.
flags	(Optional) Display the flags set, if any.
hostname *host-name*	(Optional) Display the hostname.
interface *interface-name*	(Optional) Display information about IPv6 for the specified logical interface
logical-system *logical-system-name*	(Optional) Display the IPv6 entries for the specified logical system; only available on the main router context.
reference-count *count*	(Optional) Display the IPv6 next-hop reference count.
tenant *name*	(Optional) Displays the IPv6 entries for the specified tenant system; only available in the main router context.
vpn *vpn-name*	(Optional) Display entries in the IPv6 table for the specified virtual private network's (VPN) routing table.

Required Privilege Level

view

Output Fields

Table 710 on page 21681 describes the output fields for the `show ipv6 neighbors` command. Output fields are listed in the approximate order in which they appear.

Table 710: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: **up**, **down**, **incomplete**, **reachable**, **stale**, or **unreachable**.

Table 710: show ipv6 neighbors Output Fields *(Continued)*

Field Name	Field Description
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: **yes** or **no**.
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: **yes** or **no**.
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```
user@host> show ipv6 neighbors
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale       546 yes no      fe-1/2/0.1
fe80::2a0:a514:0:24c         00:05:85:8f:c8:bd  stale       258 yes no      fe-1/2/0.1
fe80::2a0:a514:0:64c         00:05:85:8f:c8:bd  stale       111 yes no      fe-1/2/1.5
fe80::2a0:a514:0:a4c         00:05:85:8f:c8:bd  stale       327 yes no      fe-1/2/2.9
```

show ipv6 neighbors

The command displaying the underlying l2 ifl information when `enhanced-convergence` statement and `enhanced-ip` statement is not configured.

```
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure Interface
23::23:0:0:2                 00:00:23:00:00:02  reachable   0   no  no      irb.0 [xe-2/2/0.0]
```

The command not displaying the underlying l2 ifl information when `enhanced-convergence` statement and `enhanced-ip` statement is configured.

```
IPv6 Address                  Linklayer Address  State       Exp Rtr Secure Interface
23::23:0:0:2                  00:00:23:00:00:02  reachable   0   no  no      irb.0
```

Release Information

Command introduced before Junos OS Release 7.4.

flags, **hostname**, **interface**, **logical-system**, **reference-count**, **tenant**, and **vpn** options added in Junos OS Release 18.3.

RELATED DOCUMENTATION

clear ipv6 neighbors | 3659

show ipv6 neighbors (SRX Series)

IN THIS SECTION

- Syntax | 21683
- Description | 21684
- Options | 21684
- Required Privilege Level | 21684
- Output Fields | 21684
- Sample Output | 21685
- Release Information | 21685

Syntax

```
show ipv6 neighbors
```

Description

Display information about the IPv6 neighbor cache.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 711 on page 21684 lists the output fields for the `show ipv6 neighbors` command. Output fields are listed in the approximate order in which they appear.

Table 711: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: up, down, incomplete, reachable, stale, or unreachable.
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: yes or no.
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: yes or no.
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```
user@host> show ipv6 neighbors
IPv6 Address     Linklayer Address      State       Exp Rtr Secure     Interface
            10:1::2          00:00:0a:00:00:00      reachable   17   yes     no     reth0.0
            11:11::2         00:19:e2:4b:61:83      stale       1197 yes     no     at-1/0/0.0
            12:12::2         00:19:e2:4b:61:83      stale       1188 yes     no     at-3/0/0.0
```

Release Information

Command introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

clear ipv6 neighbors (SRX Series) | 3661

show ipv6 router-advertisement

IN THIS SECTION

- Syntax | 21686
- Description | 21686
- Options | 21686
- Additional Information | 21686
- Required Privilege Level | 21687
- Output Fields | 21687
- Sample Output | 21689
- Release Information | 21692

Syntax

```
show ipv6 router-advertisement
<conflicts>
<interface interface>
<logical-system (all | logical-system-name)>
<prefix prefix/prefix length>
```

Description

Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers.

The router advertisement module does not function in the backup Routing Engine as the Routing Engine does not send an acknowledgment message after receiving the packets. Starting in Junos OS Release 22.2R1, you can view the router advertisement module information using the `show ipv6 router-advertisement` operational command.

Options

none	Display all IPv6 router advertisement information for all interfaces.
conflicts	(Optional) Display only the IPv6 router advertisement information that is conflicting.
interface ***interface***	(Optional) Display IPv6 router advertisement information for the specified interface.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
prefix ***prefix/prefix length***	(Optional) Display IPv6 router advertisement information for the specified prefix.

Additional Information

The display identifies conflicting information by enclosing the value the router is advertising in brackets.

Required Privilege Level

view

Output Fields

Table 712 on page 21687 describes the output fields for the `show ipv6 router-advertisement` command. Output fields are listed in the approximate order in which they appear.

Table 712: show ipv6 router-advertisement Output Fields

Field Name	Field Description
`Interface`	Name of the interface.
`Advertisements sent`	Number of router advertisements sent and the elapsed time since they were sent.
`Solicits received`	Number of solicitation messages received.
`Advertisements received`	Number of router advertisements received.
`Advertisements from`	Names of interfaces from which router advertisements have been received and the elapsed time since the last one was received.
`Managed`	Managed address configuration flag: `0` (stateless) or `1` (stateful).
`Other configuration`	Other stateful configuration flag: `0` (stateless) or `1` (stateful).
`Reachable time`	Time that a node identifies a neighbor as reachable after receiving a reachability confirmation, in milliseconds.
`Default lifetime`	Default lifetime, in seconds: from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router.
`Retransmit timer`	Time between retransmitted Neighbor Solicitation messages, in milliseconds.

Table 712: show ipv6 router-advertisement Output Fields *(Continued)*

Field Name	Field Description
`Current hop limit`	Configured current hop limit.
`Prefix`	Name and length of the prefix.
`Valid lifetime`	How long the prefix remains valid for onlink determination.
`Preferred lifetime`	How long the prefix generated by stateless autoconfiguration remains preferred.
`On link`	Onlink flag: `0` (not onlink) or `1` (onlink).
`Autonomous`	Autonomous address configuration flag: `0` (not autonomous) or `1` (autonomous).
`Upstream Mode`	Configured interface as upstream interface for RA proxy
`Downstream Mode`	Configured interface as downstream interface for RA proxy.
`Downstream`	Downstream interface for RA proxy.
`Passive Mode`	RA receive only mode is enabled.
`Proxy Blackout Timer`	Proxy blackout timer interval is the time interval for which the interface must not be used as a proxy interface. Proxy functionality is disabled on that interface.
`Parameter Preference`	Preference to select configured or proxied parameters for downstream interface
`error`	Displays the details of the error.

Sample Output

show ipv6 router-advertisement (with RA and NAT64 IPv6 address prefix) (cSRX Container Firewall, MX Series, SRX Series, vMX, and vSRX Virtual Firewall)

```
user@host> show ipv6 router-advertisement
Interface: ge-0/0/0.0
  Advertisements sent: 3, last sent 00:00:05 ago
  Solicits sent: 1, last sent 00:00:37 ago
  Solicits received: 0
  Advertisements received: 0
  Solicited router advertisement unicast: Disable
  IPv6 RA Preference: DEFAULT/MEDIUM
  Passive mode: Disable
  Upstream mode: Disable
  Downstream mode: Disable
  Proxy blackout timer: Not Running
  NAT prefix: 6000::/96 Lifetime: 8192 sec
```

show ipv6 router-advertisement

```
user@host> show ipv6 router-advertisement
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
    Managed: 0
    Other configuration: 0 [1]
      Reachable time: 0 ms
      Default lifetime: 1800 sec
      Retransmit timer: 0 ms
      Current hop limit: 64
```

show ipv6 router-advertisement (Without RA proxy) (SRX Series and vSRX Virtual Firewall 3.0)

(Without RA proxy)

```
user@host> show ipv6 router-advertisement
Interface: ge-0/0/1.0
  Advertisements sent: 7, last sent 00:00:11 ago
  Solicits sent: 1, last sent 00:00:41 ago
  Solicits received: 0
  Advertisements received: 0
  Solicited router advertisement unicast: Disable
  IPv6 RA Preference: DEFAULT/MEDIUM
  Passive mode: Enable
  Upstream mode: Disable
  Downstream mode: Disable
  Proxy blackout timer: Not Running
```

(With RA proxy)

```
Interface: ge-0/0/2.0
  Advertisements sent: 2, last sent 00:00:49 ago
  Solicits sent: 1, last sent 00:01:21 ago
  Solicits received: 0
  Advertisements received: 18
  Solicited router advertisement unicast: Disable
  IPv6 RA Preference: DEFAULT/MEDIUM
  Passive mode: Enable
  Upstream mode: Enable
  Downstream mode: Disable
  Proxy parameter preference: Proxied
  Proxy blackout timer: Not Running
  Downstream: ge-0/0/0.0
  Downstream: ge-0/0/1.0
  Advertisement from fe80::5668:adff:fed8:101b, heard 00:00:00 ago
  Managed: 0
  Other configuration: 1 [0]
  Link MTU: 1500 bytes
  Reachable time: 5555 ms
  Default lifetime: 1799 sec [1800 sec]
  Retransmit timer: 4444 ms
```

```
  Current hop limit: 50 [64]
  RDNSS address: abcd:1::1
  Lifetime: 3333 sec
  Prefix: 2002:2:0:2000::/64
  Valid lifetime: 3600 sec
  Preferred lifetime: 2400 sec
  On link: 1
  Autonomous: 1
  Route Information: 2002:2:0:2000::/64
  IPv6 RA Preference: LOW
  Route lifetime: 1111 sec
  DNSSL suffix: juniper.net
  Lifetime: 6666 sec
```

show ipv6 router-advertisement conflicts

```
user@host> show ipv6 router-advertisement conflicts
Interface: fxp0.0
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago
    Other configuration: 0 [1]
```

show ipv6 router-advertisement prefix

```
user@host> show ipv6 router-advertisement prefix 2001:db8:8040::/16
Interface: fe-0/1/3.0
  Advertisements sent: 3, last sent 00:04:11 ago
  Solicits received: 0
  Advertisements received: 3
  Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago
    Managed: 0
    Other configuration: 0
    Reachable time: 0 ms
    Default lifetime: 180 sec [1800 sec]
    Retransmit timer: 0 ms
    Current hop limit: 64
    Prefix: 2001:db8:8040:1::/64
      Valid lifetime: 2592000 sec
      Preferred lifetime: 604800 sec
```

```
        On link: 1
        Autonomous: 1
```

show ipv6 router-advertisement (Backup Routing Engine)

```
user@host> show ipv6 router-advertisement
error: Module not running (routing)
```

Release Information

Command introduced before Junos OS Release 7.4.

Starting in Junos OS Release 22.1, we support RA proxy on SRX Series and vSRX Virtual Firewall 3.0 . This command output is modified to display the configured upstream interfaces, downstream interfaces, the proxy flag, the proxy blackout timer, and the passive mode information.

RELATED DOCUMENTATION

clear ipv6 router-advertisement | 3663

show ipv6-source-guard

IN THIS SECTION

- Syntax | 21693
- Description | 21693
- Required Privilege Level | 21693
- Output Fields | 21693
- Sample Output | 21694
- Release Information | 21694

Syntax

```
show ipv6-source-guard
```

Description

(For non-ELS switches) Display IPv6 source guard database information.

Required Privilege Level

view

Output Fields

Table 713 on page 21693 lists the output fields for the `show ipv6-source-guard` command. Output fields are listed in the approximate order in which they appear.

Table 713: show ipv6-source-guard Output Fields

Field Name	Field Description
VLAN	VLAN on which IPv6 source guard is enabled.
Interface	Access interface associated with the VLAN desribed in row 1.
Tag	VLAN ID for the VLAN described in row 1. Possible values are: • 0, indicating the VLAN is not tagged. • 1 through 4093
IP Address	Source IP address for a device connected to the interface described in row 2. A * (asterisk) indicates that IPv6 source guard is not enabled on this VLAN, but the interface is shared with a VLAN that is enabled for IPv6 source guard.
MAC Address	Source MAC address for a device connected to the interface described in row 2. A * (asterisk) indicates that IPv6 source guard is not enabled on this VLAN but the interface is shared with a VLAN that is enabled for IPv6 source guard.

Sample Output

show ipv6-source-guard

```
user@switch> show ipv6-source-guard
IP source guard information:
Interface    Tag   IP Address                    MAC Address         VLAN
ge-0/0/6.0   0     2001:db8::10:0:15             00:10:94:10:00:01   vlan1
ge-0/0/6.0   0     fe80::210:94ff:fe10:1         00:10:94:10:00:01   vlan1
ge-0/0/7.0   0     2001:db8::10:0:14             00:10:94:10:00:02   vlan1
ge-0/0/7.0   0     fe80::210:94ff:fe10:2         00:10:94:10:00:02   vlan1
```

Release Information

Command introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN

Example: Configuring IP Source Guard with Other EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces

Verifying That IP Source Guard Is Working Correctly

show isis adjacency

IN THIS SECTION

- Syntax | 21695
- Syntax (EX Series Switches and QFX Series) | 21695
- Description | 21695
- Options | 21695
- Required Privilege Level | 21696
- Output Fields | 21696

Sample Output | 21700
Release Information | 21702

Syntax

```
show isis adjacency
<system-id>
<brief | detail | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show isis adjacency
<system-id>
<brief | detail | extensive>
<instance instance-name>
```

Description

Display information about IS-IS neighbors.

Options

none	Display standard information about IS-IS neighbors for all routing instances.
system id	(Optional) Display information about IS-IS neighbors for the specified intermediate system.
brief | detail | extensive	(Optional) Display standard information about IS-IS neighbors with the specified level of output.
instance *instance-name*	(Optional) Display information about IS-IS neighbors for the specified routing instance.

logical-system (all | *logical-system-name*) (Optional) Display information about IS-IS neighbors for all logical systems or for a particular logical system.

Required Privilege Level

view

Output Fields

Table 714 on page 21696 describes the output fields for the `show isis adjacency` command. Output fields are listed in the approximate order in which they appear.

Table 714: show isis adjacency Output Fields

Field Name	Field Description	Level of Output
`Interface`	Interface through which the neighbor is reachable.	All levels
`System`	System identifier (`sysid`), displayed as a name, if possible.	`brief`
`L or Level`	Level: • `1`—Level 1 only • `2`—Level 2 only • `3`—Level 1 and Level 2 An exclamation point (`!`) preceding the level number indicates that the adjacency is missing an IP address.	All levels
`State`	State of the adjacency: `Up`, `Down`, `New`, `One-way`, `Initializing`, or `Rejected`.	All levels
`Hold (secs)`	Remaining hold time of the adjacency.	`brief`
`SNPA`	Subnetwork point of attachment (MAC address of the next hop).	`brief`
`Expires in`	How long until the adjacency expires, in seconds.	`detail`

Table 714: show isis adjacency Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Priority	Priority to become the designated intermediate system.	detail extensive
Up/Down transitions	Count of adjacency status changes from Up to Down or from Down to Up.	detail
Last transition	Time of the last Up/Down transition.	detail
Circuit type	Bit mask of levels on this interface: 1=Level 1 router; 2=Level 2 router; 3=both Level 1 and Level 2 router.	detail
Speaks	Protocols supported by this neighbor.	detail extensive
MAC address	MAC address of the interface.	detail extensive
Topologies	Supported topologies.	detail extensive
Restart capable	Whether a neighbor is capable of graceful restart: Yes or No.	detail extensive
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their link-state PDUs.	detail extensive
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise the interface in the routing device's outbound link-state PDUs.	detail extensive
IP addresses	IP address of this neighbor.	detail extensive
IPV6 Address	IPv6 address of the neighbor.	detail extensive

Table 714: show isis adjacency Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Level 1 IPv4 Adj-SID	Level 1 IPv4 node-SID of the adjacent neighbor.	`detail extensive`
Level 1 IPv6 Adj-SID	Level 1 IPv6 node-SID of the adjacent neighbor.	
Level 2 IPv4 Adj-SID	Level 2 IPv4 node-SID of the adjacent neighbor.	`detail extensive`
Level 2 IPv6 Adj-SID	Level 2 IPv6 node-SID of the adjacent neighbor.	`detail extensive`
SRv6	Segment routing capability is enabled on the device.	extensive
End-X-SID	Function for SRv6 adjacency.	extensive
Flavor	End-X-SID function can have the following flavors: • PSP— Penultimate segment pop of the segment routing header (SRH). • USD— Ultimate segment pop of the SRH. • USP— Ultimate segment decapsulation of the SRH.	extensive

Table 714: show isis adjacency Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Transition log`	List of recent transitions, including:
	• `When`—Time at which an IS-IS adjacency transition occurred.
	• `State`—Current state of the IS-IS adjacency (`up`, `down`, or `rejected`).
	  • `Up`—Adjacency is up and operational.
	  • `Down`—Adjacency is down and not available.
	  • `Rejected`—Adjacency has been rejected.
	• `Event`—Type of transition that occurred.
	  • `Seenself`—Possible routing loop has been detected.
	  • `Interface down`—IS-IS interface has gone down and is no longer available.
	  • `Error`—Adjacency error.
	• `Down reason`—Reason that an IS-IS adjacency is down:
	  • `3-Way Handshake Failed`—Connection establishment failed.
	  • `Address Mismatch`—Address mismatch caused link failure.
	  • `Aged Out`—Link expired.
	  • `ISO Area Mismatch`—IS-IS area mismatch caused link failure.
	  • `Bad Hello`—Unacceptable hello message caused link failure.
	  • `BFD Session Down`—Bidirectional failure detection caused link failure.
	  • `Interface Disabled`—IS-IS interface is disabled.
	  • `Interface Down`—IS-IS interface is unavailable.
	  • `Interface Level Disabled`—IS-IS level is disabled.
		`extensive`

Table 714: show isis adjacency Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Level Changed`—IS-IS level has changed on the adjacency. • `Level Mismatch`—Levels on adjacency are not compatible. • `MPLS LSP Down`—Label-switched path (LSP) is unavailable. • `MT Topology Changed`—IS-IS topology has changed. • `MT Topology Mismatch`—IS-IS topology is mismatched. • `Remote System ID Changed`—Adjacency peer system ID changed. • `Protocol Shutdown`—IS-IS protocol is disabled. • `CLI Command`—Adjacency brought down by user. • `Unknown`—Unknown.	

Sample Output

show isis adjacency

```
user@host> show isis adjacency
Interface             System         L State        Hold (secs) SNPA
at-2/3/0.0            ranier         3  Up                   23
```

show isis adjacency brief

The output for the `show isis adjacency brief` command is identical to that for the `show isis adjacency` command. For sample output, see "show isis adjacency" on page 21700.

show isis adjacency detail

```
user@host> show isis adjacency detail
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 21 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:09 ago
```

```
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast, IPV6-Unicast  Restart capable: Yes, Adjacency advertisement: Advertise
  LAN id: pro-bng3-c-F.02, IP addresses: 11.1.1.2
  IPv6 addresses: fe80::2a0:a514:0:4745
  Level 1 IPv4 Adj-SID: 299808, IPv6 Adj-SID: 299824
```

show isis adjacency extensive

```
user@host> show isis adjacency extensive
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 22 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:16 ago
  Circuit type: 3, Speaks: IP, IPv6

  Topologies: Unicast, IPV6-Unicast
  Restart capable: Yes, Adjacency advertisement: Advertise
  LAN id: ro_re.02, IP addresses: 11.1.1.2
  IPv6 addresses: fe80::2a0:a514:0:3e45
  IPv6 Global address: ::40.1.1.3
  Level 1 IPv4 Adj-SID: 300112, IPv6 Adj-SID: 300304
  Level 2 IPv4 Adj-SID: 300320, IPv6 Adj-SID: 300336
  Transition log:
  When                  State        Event           Down reason
  Thu Mar 26 06:13:18   Up           Seenself
```

command-name

```
Interface: ge-0/0/0.0, Level: 1, State: Up, Expires in 7 secs
  Priority: 64, Up/Down transitions: 1, Last transition: 00:14:29 ago
  Circuit type: 1, Speaks: IP, IPv6, MAC address: 56:68:ac:d3:30:2f
  Topologies: Unicast
  Restart capable: Yes, Adjacency advertisement: Advertise
  LAN id: R5_re.02, IP addresses: 5.5.5.5
  IPv6 addresses: fe80::5668:acff:fed3:302f
  IPv6 Global Interface Address: 10:57:57::5
  Level 1 SRv6 unprotected END-X-SID: 7000::6
    Flavor: PSP, Flags: --P
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear isis adjacency | 3664

show isis adjacency holddown

IN THIS SECTION

- Syntax | 21702
- Description | 21702
- Options | 21703
- Required Privilege Level | 21703
- Output Fields | 21703
- Sample Output | 21704
- Release Information | 21705

Syntax

```
show isis adjacency holddown
<level>
<brief | detail | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display holddown status and time when IS-IS adjacencies are being formed. Adjacency holddown process takes place on an IS-IS level basis. When adjacency holdown is enabled, IS-IS adjacencies are

formed sequentially. There is a holddown time between each adjacency and the process is completed when all IS-IS adjacencies are formed. This holddown time might cause network instability.

This command is useful to verify whether the adjacency holddown is enabled and facilitates troubleshooting when there are adjacency issues due to IS-IS adjacency holddown .

Options

none	Display standard overview information about IS-IS adjacency holddown for all routing instances.
brief | detail | extensive	(Optional) Display standard information about IS-IS adjacency holddown with the specified level of output.
instance ***instance-name***	(Optional) Display IS-IS adjacency holddown information for the specified routing instance.
level	(Optional) Display information about IS-IS neighbors for the specified IS-IS level.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Field	Field Description
`Level`	IS-IS level: • 1—Level 1 information • 2—Level 2 information

(Continued)

Field	Field Description
status	IS-IS Adjacency holddown status: • active—IS-IS adjacency holddown is in process. • inactive—IS-IS adjacency holddown process is completed. • disabled—IS-IS adjacency holddown is disabled.
when	The time stamp for the IS-IS adjacency holddown.
Event	• seenself—IS-IS has changed the IS-IS adjacency holddown status. • status change—IS-IS adjacency holddown status has changed.
reason	The reason for change in the IS-IS adjacency holddown state.

Sample Output

show isis adjacency holddown

```
user@host> show isis adjacency holddown
Level           status
1               active
2               active
```

show isis adjacency holddown extensive

```
user@host> show isis adjacency holddown extensive
Level: 1
  Adjacency holddown is active, Adjacency holddown reset is not completed
  Holddown process has not started, Network might not be stable
  Holddown Started 0 secs ago
  First adjacency up since 0 secs
  Total adjacency up count: 0
```

```
  Transition log:
  When                  State        Event           reason
  Mon Jun 22 05:35:03   Active       Seenself
  Tue Jun 23 02:13:30   Active       Seenself
  Tue Jun 23 02:14:40   Inactive     Status change   holddown complete
  Tue Jun 23 02:15:52   Active       Seenself

Level: 2
  Adjacency holddown is active, Adjacency holddown reset is not completed
  Holddown process has not started, Network might not be stable
  Holddown Started 0 secs ago
  First adjacency up since 0 secs
  Total adjacency up count: 0
  Transition log:
  When                  State        Event           reason
  Mon Jun 22 05:35:03   Active       Seenself
  Tue Jun 23 02:13:30   Active       Seenself
  Tue Jun 23 02:14:40   Inactive     Status change   holddown complete
  Tue Jun 23 02:15:52   Active       Seenself
```

show isis adjacency holddown detail

```
user@host> show isis adjacency holddown detail
 Level: 1
  Adjacency holddown is active, Adjacency holddown reset is not completed
  Holddown process has not started, Network might not be stable
  Holddown Started 0 secs ago
  First adjacency up since 0 secs
  Total adjacency up count: 0

Level: 2
  Adjacency holddown is active, Adjacency holddown reset is not completed
  Holddown process has not started, Network might not be stable
  Holddown Started 0 secs ago
  First adjacency up since 0 secs
  Total adjacency up count: 0
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

show isis adjacency | 21694

show isis authentication

IN THIS SECTION

- Syntax | 21706
- Syntax (EX Series Switches and QFX Series) | 21706
- Description | 21706
- Options | 21707
- Required Privilege Level | 21707
- Output Fields | 21707
- Sample Output | 21708
- Release Information | 21708

Syntax

```
show isis authentication
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show isis authentication
<instance instance-name>
```

Description

Display information about Intermediate System-to-Intermediate System (IS-IS) authentication.

Options

none	Display information about IS-IS authentication.
instance ***instance-name***	(Optional) Display IS-IS authentication for the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 715 on page 21707 describes the output fields for the `show isis authentication` command. Output fields are listed in the approximate order in which they appear.

Table 715: show isis authentication Output Fields

Field Name	Field Description
`Interface`	Interface name.
`Level`	IS-IS level.
`IIH Auth`	IS-IS Hello (IIH) packet authentication type. Displays the name of the active keychain if hitless authentication key rollover is configured.
`CSN Auth`	Complete sequence number authentication type.
`PSN Auth`	Partial sequence number authentication type.
`L1 LSP Authentication`	Layer 1 link-state PDU authentication type.
`L2 LSP Authentication`	Layer 2 link-state PDU authentication type.

Sample Output

show isis authentication

```
user@host> show isis authentication
Interface             Level IIH Auth  CSN Auth  PSN Auth
at-2/3/0.0            1     Simple    Simple    Simple
                      2     MD5       MD5       MD5

L1 LSP Authentication: Simple
L2 LSP Authentication: MD5
```

show isis authentication (With Hitless Authentication Key Rollover Configured)

```
user@host> show isis authentication
Interface             Level IIH Auth  CSN Auth  PSN Auth
so-0/1/3.0            2     hakrhello MD5       MD5

L2 LSP Authentication: MD5
```

Release Information

Command introduced in Junos OS Release 7.5.

Support for hitless authentication key rollover introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

show security keychain | 24765

show isis backup coverage

IN THIS SECTION

- Syntax | 21709
- Syntax (EX Series Switches and QFX Series) | 21709
- Description | 21709
- Options | 21709
- Required Privilege Level | 21710
- Output Fields | 21710
- Sample Output | 21711
- Release Information | 21711

Syntax

```
show isis backup coverage
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show isis backup coverage
<instance instance-name>
```

Description

Display information about the level of backup coverage available.

Options

none Display information about the level of backup coverage available for all the nodes and prefixes in the network.

instance ***instance-name***	(Optional) Display information about the level of backup coverage for a specific IS-IS routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 716 on page 21710 lists the output fields for the `show isis backup coverage` command. Output fields are listed in the approximate order in which they appear.

Table 716: show isis backup coverage Output Fields

Field Name	Field Description
Topology	Type of topology or address family: `IPV4 Unicast` or `IPV6 Unicast`.
Level	IS-IS level: • `1`—Level 1 • `2`—Level 2
Node	By topology, the percentage of all routes configured on the node that are protected through backup coverage.
IPv4	Percentage of IPv4 unicast routes that are protected through backup coverage.
IPv6	Percentage of IPv6 unicast routes that are protected through backup coverage.
CLNS	Percentage of Connectionless Network Service (CLNS) routes that are protected through backup coverage.

Sample Output

show isis backup coverage

```
user@host> show isis backup coverage
Backup Coverage:
    Topology        Level   Node    IPv4    IPv6    CLNS
    IPV4 Unicast        2  28.57%  22.22%   0.00%   0.00%
    IPV6 Unicast        2   0.00%   0.00%   0.00%   0.00%
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Loop-Free Alternate Routes for IS-IS

show isis backup label-switched-path | 21711

show isis backup label-switched-path

IN THIS SECTION

- Syntax | 21712
- Syntax (EX Series Switches and QFX Series) | 21712
- Description | 21712
- Options | 21712
- Required Privilege Level | 21712
- Output Fields | 21712
- Sample Output | 21713
- Release Information | 21713

Syntax

```
show isis backup label-switched-path
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show isis backup label-switched-path
```

Description

Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes.

Options

none	Display information about MPLS LSPs designated as backup routes for IS-IS routes.
logical-system (all | *logical-system-name*	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 717 on page 21712 lists the output fields for the `show isis backup label-switched-path` command. Output fields are listed in the approximate order in which they appear.

Table 717: show isis backup label-switched-path Output Fields

Field Name	Field Description
Backup MPLS LSPs	List of MPLS LSPs designated as backup paths for IS-IS routes.

Table 717: show isis backup label-switched-path Output Fields *(Continued)*

Field Name	Field Description
Egress	IP address of the egress routing device for the LSP.

Status	State of the LSP:
	• `Up`—The routing device can detect RSVP hello messages from the neighbor.
	• `Down`—The routing device has received one of the following indications:
	  • Communication failure from the neighbor.
	  • Communication from IGP that the neighbor is unavailable.
	  • Change in the sequence numbers in the RSVP hello messages sent by the neighbor.
	• `Deleted`—LSP is no longer available as a backup path.

Last change	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is *hh*:*mm*:*ss*.
TE-metric	Configured traffic engineering metric.
Metric	Configured metric.

Sample Output

show isis backup label-switched-path

```
user@host> show isis backup label-switched-path
 Backup MPLS LSPs:
 f-to-g, Egress: 192.168.1.4, Status: up, Last change: 06:12:03
    TE-metric: 9, Metric: 0
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Loop-Free Alternate Routes for IS-IS

show isis backup coverage | 21709

show isis backup spf results

IN THIS SECTION

- Syntax | 21714
- Syntax (EX Series Switches) | 21714
- Description | 21715
- Options | 21715
- Required Privilege Level | 21715
- Output Fields | 21715
- Sample Output | 21716
- Release Information | 21720

Syntax

```
show isis backup spf results
<instance instance-name>
<level (1 | 2)>
<logical-system (all | logical-system-name)>
<no-coverage>
<topology (ipv4-unicast | ipv6-multicast | ipv6-unicast | unicast)>
```

Syntax (EX Series Switches)

```
show isis backup spf results
<instance instance-name>
<level (1 | 2)>
```

```
<no-coverage>
<topology (ipv4-unicast | unicast)>
```

Description

Display information about IS-IS shortest-path-first (SPF) calculations for backup paths.

Options

none	Display information about IS-IS SPF calculations for all backup paths for all destination nodes.
instance *instance-name*	(Optional) Display SPF calculations for backup paths for the specified routing instance.
level (1 | 2)	(Optional) Display SPF calculations for the backup paths for the specified IS-IS level.
logical-system *logical-system-name*	(Optional) Display SPF calculations for the backup paths for all logical systems or on a particular logical system.
no-coverage	(Optional) Display SPF calculations only for destinations that do not have backup coverage.
topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)	(Optional) Display SPF calculations for backup paths for the specified topology only.

Required Privilege Level

view

Output Fields

Table 718 on page 21716 lists the output fields for the `show isis backup spf results` command. Output fields are listed in the approximate order in which they appear.

Table 718: show isis backup spf results Output Fields

Field Name	Field Description
node-name	Name of the destination node.
Address	Address of the destination node.
Primary next-hop	Interface and name of the node of the primary next hop to reach the destination.
Root	Name of the next-hop neighbor.
Metric	Metric to the node.
Eligible	Indicates that the next-hop neighbor has been designated as a backup path to the destination node.
Backup next-hop	Name of the interface of the backup next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
LSP	Name of the MPLS label-switched path (LSP) designated as a backup path.
Not eligible	Indicates that the next-hop neighbor cannot function as a backup path to the destination.
Reason	Describes why the next-hop neighbor is designated as `Not eligible` as a backup path.

Sample Output

show isis backup spf results

```
user@host> show isis backup spf results
```

```
IS-IS level 1 SPF results:
  0 nodes

 IS-IS level 2 SPF results:
banff.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: ae0.0, IPV4, camaro, SNPA:  0:90:69:f:67:f0
  Primary next-hop: so-6/0/0.0, IPV6, olympic
  Primary next-hop: ae0.0, IPV6, camaro, SNPA:  0:90:69:f:67:f0
    Root: camaro, Root Metric: 10, Metric: 10
      Not eligible, Reason: Primary next-hop multipath
    Root: olympic, Root Metric: 10, Metric: 10
      Not eligible, Reason: Primary next-hop multipath
    Root: glacier, Root Metric: 10, Metric: 25
      Not eligible, Reason: Primary next-hop multipath
crater.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: so-6/0/0.0, IPV6, olympic
    Root: olympic, Root Metric: 10, Metric: 10
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: glacier, Root Metric: 10, Metric: 15
      Eligible, Backup next-hop: as0.0, IPV4, glacier
      Eligible, Backup next-hop: as0.0, IPV6, glacier
    Root: camaro, Root Metric: 10, Metric: 20
      Not eligible, Reason: Interface is already covered
olympic.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: so-6/0/0.0, IPV6, olympic
    Root: olympic, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: camaro, Root Metric: 10, Metric: 20
      track-item: olympic.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: glacier, Root Metric: 10, Metric: 20
      track-item: olympic.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
camaro.00
  Primary next-hop: ae0.0, IPV4, camaro, SNPA:  0:90:69:f:67:f0
  Primary next-hop: ae0.0, IPV6, camaro, SNPA:  0:90:69:f:67:f0
    Root: camaro, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
```

```
    Root: glacier, Root Metric: 10, Metric: 20
      track-item: camaro.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: olympic, Root Metric: 10, Metric: 20
      track-item: camaro.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
glacier.00
  Primary next-hop: as0.0, IPV4, glacier
  Primary next-hop: as0.0, IPV6, glacier
    Root: glacier, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: camaro, Root Metric: 10, Metric: 20
      track-item: glacier.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: olympic, Root Metric: 10, Metric: 20
      track-item: glacier.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
  5 nodes
```

show isis backup spf results no-coverage

```
user@host> show isis backup spf results no-coverage
IS-IS level 1 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA:  b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA:  b0:c6:9a:2c:f0:de
    Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
    Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
      track-item: pro-bng3-k.00-00
      track-item: pro-bng3-j.00-00
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA:  b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA:  b0:c6:9a:2a:f4:21
    Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
    Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
      track-item: pro-bng3-j.00-00
      track-item: pro-bng3-i.00-00
```

```
   2 nodes

 IS-IS level 2 SPF results:
olympic.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: so-6/0/0.0, IPV6, olympic
    Root: olympic, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: camaro, Root Metric: 10, Metric: 20
      track-item: olympic.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: glacier, Root Metric: 10, Metric: 20
      track-item: olympic.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
camaro.00
  Primary next-hop: ae0.0, IPV4, camaro, SNPA:  0:90:69:f:67:f0
  Primary next-hop: ae0.0, IPV6, camaro, SNPA:  0:90:69:f:67:f0
    Root: camaro, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: glacier, Root Metric: 10, Metric: 20
      track-item: camaro.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: olympic, Root Metric: 10, Metric: 20
      track-item: camaro.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
glacier.00
  Primary next-hop: as0.0, IPV4, glacier
  Primary next-hop: as0.0, IPV6, glacier
    Root: glacier, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: camaro, Root Metric: 10, Metric: 20
      track-item: glacier.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: olympic, Root Metric: 10, Metric: 20
      track-item: glacier.00-00
      track-item: kobuk.00-00
```

```
    Not eligible, Reason: Path loops
3 nodes
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

show isis backup coverage | 21709

Understanding Loop-Free Alternate Routes for IS-IS

show isis context-identifier

IN THIS SECTION

- Syntax | 21720
- Description | 21721
- Options | 21721
- Required Privilege Level | 21721
- Output Fields | 21721
- Sample Output | 21722
- Release Information | 21722

Syntax

```
show isis context-identifier
<brief | detail | extensive>
<identifier name>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display IS-IS context identifier information.

Options

`brief | detail | extensive`	(Optional) Display the specified level of output.
`identifier name`	(Optional) Display information about the specified context identifier.
`instance` *`instance-name`*	(Optional) Display entries for the specified routing instance.
`logical-system (all |` *`logical-system-name`*`)`	(Optional) Display the context identifier information for all logical systems or for a particular logical system.

Required Privilege Level

View

Output Fields

Table 719 on page 21721 lists the output fields for the `show isis context-identifier` command. Output fields are listed in the approximate order in which they appear.

Table 719: show isis context-identifier Output Fields

Field Name	Field Description	Level of Output
Context	IPv4 address that defines a protection pair. The context is manually configured on both primary and protector PEs.	detail
Owner	Protocol that requires the context.	detail
Role	Role of the PE, which is either primary or protector.	detail
Primary	Name of the primary PE.	detail
Metric	Advertised interior gateway protocol (IGP) metric.	detail

Sample Output

command-name

```
user@host> show isis context-identifier detail

IS-IS context database:
Context                Owner     Role       Primary      Metric
2.2.4.3               MPLS      Primary    pro3-e        1
  Advertiser pro3-e, Router ID 10.255.245.198, Metric 1, Level 1
  Advertiser pro3-e, Router ID 10.255.245.198, Metric 1, Level 2
  Advertiser pro3-c, Router ID 10.255.245.196, Metric 11, Level 2
```

Release Information

Command introduced in Junos OS Release 10.4.

show isis database

IN THIS SECTION

- Syntax | 21723
- Syntax (EX Series Switches and QFX Series) | 21723
- Description | 21723
- Options | 21723
- Required Privilege Level | 21724
- Output Fields | 21724
- Sample Output | 21728
- Release Information | 21732

Syntax

```
show isis database
<system-id>
<brief | detail | extensive>
<instance instance-name>
<level (1 | 2)>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show isis database
<system-id>
<brief | detail | extensive>
<level (1 | 2)>
<instance instance-name>
```

Description

Display the entries in the Intermediate System-to-Intermediate System (IS-IS) link-state database, which contains data about PDU packets.

Options

none	Display standard information about IS-IS link-state database entries for all routing instances.
system id	(Optional) Display IS-IS link-state database entries for the specified intermediate system.
brief | detail | extensive	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display IS-IS link-state database entries for the specified routing instance.
level (1 | 2)	(Optional) Display IS-IS link-state database entries for the specified IS-IS level.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 720 on page 21724 describes the output fields for the `show isis database` command. Output fields are listed in the approximate order in which they appear. Fields that contain internal IS-IS information useful only in troubleshooting obscure problems are not described in the table. For more details about these fields, contact your customer support representative.

Table 720: show isis database Output Fields

Field Name	Field Description	Level of Output
Interface name	Name of the interface on which the link-state PDU has been received; always `IS-IS` for this command.	All levels
level	Level of intermediate system: • `1`—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system. • `2`—Intermediate system routes between areas and toward other ASs.	All levels
LSP ID	Link-state PDU identifier.	All levels
Sequence	Sequence number of the link-state PDU.	All levels
Checksum	Checksum value of the link-state PDU.	All levels
Lifetime (secs)	Remaining lifetime of the link-state PDU, in seconds.	All levels
Attributes	Attributes of the specified database: `L1`, `L2`, `Overload`, or `Attached` (L1 only).	none `brief`

Table 720: show isis database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
#LSPs	Total number of link-state PDUs in the specified link-state database.	none `brief`
IP prefix	Prefix advertised by this link-state PDU.	`detail extensive`
IS neighbor	IS-IS neighbor of the advertising system.	`detail extensive`
IP prefix	IPv4 prefix advertised by this link-state PDU.	`detail extensive`
V6 prefix	IPv6 prefix advertised by this link-state PDU.	`detail extensive`
Metric	Metric of the prefix or neighbor.	`detail extensive`
Header	• `LSP ID`—Link state PDU identifier of the header. • `Length`—Header length. • `Allocated Length`—Amount of length available for the header. • `Router ID`—Address of the local routing device. • `Remaining Lifetime`—Remaining lifetime of the link-state PDU, in seconds.	`extensive`

Table 720: show isis database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Packet	• `LSP ID`—The identifier for the link-state PDU. • `Length`—Packet length. • `Lifetime`—Remaining lifetime, in seconds. • `Checksum`—The checksum of the link-state PDU. • `Sequence`—The sequence number of the link-state PDU. Every time the link-state PDU is updated, this number increments. • `Attributes`—Packet attributes. • `NLPID`—Network layer protocol identifier. • `Fixed length`—Specifies the set length for the packet.	`extensive`

Table 720: show isis database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
TLVs	• `Area Address`—Area addresses that the routing device can reach. • `Speaks`—Supported routing protocols. • `IP router id`—ID of the routing device (usually the IP address). • `IP address`—IPv4 address. • `Hostname`—Assigned name of the routing device. • `IP prefix`—IP prefix of the routing device. • `Metric`—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device. • `IP extended prefix`—Extended IP prefix of the routing device. • `IS neighbor`—Directly attached neighbor's name and metric. • `IS extended neighbor`—Directly attached neighbor's name, metric, IP address, local interface index, and remote interface index. The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205. • `Router Capability`—ID of the routing device and flag. **NOTE**: Router capability also specifies SPRING capability and SPRING algorithm when segment routing is enabled on the routing device.	`extensive`
P2P SRv6 END-X-SID	Point-to-Point function END-X-SID for SRv6 adjacency	extensive
SRv6 Locator	Locator address to which the END-X-SID is attached.	extensive
SRv6 Capability	ID and flag indicating the segment routing capability of the device.	extensive

Table 720: show isis database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flavor	Behavior of the END-X-SID function. This field can have the following values: • PSP—Penultimate segment pop of the segment routing header (SRH). • USP—Ultimate segment pop of the SRH. • USD—Ultimate segment decapsulation of the SRH.	extensive
SRv6 Maximum Segments Left	Maximum number of segments left for processing. If no value is advertised, then it is set to 0.	extensive

Sample Output

show isis database

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
kobuk.00-00                      0x3   0x3167     1057 L1 L2
camaro.00-00                     0x5   0x770e     1091 L1 L2
ranier.00-00                     0x4   0xaa95     1091 L1 L2
glacier.00-00                    0x4   0x206f     1089 L1 L2
glacier.02-00                    0x1   0xd141     1089 L1 L2
badlands.00-00                   0x3   0x87a2     1093 L1 L2
  6 LSPs

IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
kobuk.00-00                      0x6   0x8d6b     1096 L1 L2
camaro.00-00                     0x9   0x877b     1101 L1 L2
ranier.00-00                     0x8   0x855d     1103 L1 L2
glacier.00-00                    0x7   0xf892     1098 L1 L2
glacier.02-00                    0x1   0xd141     1089 L1 L2
```

```
badlands.00-00                    0x6    0x562     1105 L1 L2
  6 LSPs
```

show isis database brief

The output for the `show isis database brief` command is identical to that for the `show isis database` command. For sample output, see "show isis database" on page 21728.

show isis database detail

```
user@host> show isis database logical-system CE3 sisira.00-00 detail

IS-IS level 1 link-state database:

sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 975 secs
   IS neighbor: hemantha-CE3.02               Metric:       10
   ES neighbor: 0015.0015.0015                Metric:       10 Down
   ES neighbor: 0025.0025.0025                Metric:       10 Down
   ES neighbor: 0030.0030.0030                Metric:       10 Down
   ES neighbor: 0040.0040.0040                Metric:       10 Down
   ES neighbor: sisira                        Metric:        0
   IP prefix: 10.1.0.0/24                      Metric:       10 External Down
   IP prefix: 10.3.0.0/24                      Metric:       10 External Down
   IP prefix: 10.4.0.0/24                      Metric:       10 External Down
   IP prefix: 10.5.0.0/24                      Metric:       10 Internal Up
   IP prefix: 10.15.15.15/32                  Metric:       10 External Down
   IP prefix: 10.25.25.25/32                  Metric:       10 External Down
   IP prefix: 10.30.30.30/32                  Metric:       10 External Down
   IP prefix: 10.40.40.40/32                  Metric:       10 External Down
   IP prefix: 10.60.60.60/32                  Metric:        0 Internal Up

IS-IS level 2 link-state database:

sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 993 secs
   IS neighbor: hemantha-CE3.02               Metric:       10
   IP prefix: 10.1.0.0/24                      Metric:       10 External Down
   IP prefix: 10.3.0.0/24                      Metric:       10 External Down
   IP prefix: 10.4.0.0/24                      Metric:       10 External Down
   IP prefix: 10.5.0.0/24                      Metric:       10 Internal Up
   IP prefix: 10.15.15.15/32                  Metric:       10 External Down
   IP prefix: 10.25.25.25/32                  Metric:       10 External Down
```

```
   IP prefix: 10.30.30.30/32                  Metric:       10 External Down
   IP prefix: 10.40.40.40/32                  Metric:       10 External Down
   IP prefix: 10.50.50.50/32                  Metric:       10 Internal Up
   IP prefix: 10.60.60.60/32                  Metric:        0 Internal Up
   ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
                                              Metric:       10 External Down
   ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
                                              Metric:       10 External Down
   ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
                                              Metric:       10 External Down
   ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
                                              Metric:       10 External Down
   ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
                                              Metric:        0 Internal Up
```

show isis database extensive (TI-LFA for IS-IS v6)

```
user@host> show isis database extensive
TLVs:
    Area address: 47.0005.80ff.f800.0000.0108.0001 (13)
    LSP Buffer Size: 1492
    Speaks: IP
    Speaks: IPV6
    IP router id: 10.1.1.1
    IP address: 10.1.1.1
    Hostname: r1_re0
    IS neighbor: r0_re.00, Internal, Metric: default 10
    IS neighbor: r1_re0.02, Internal, Metric: default 10
    Extended IS Reachability TLV, Type: 22, Length: 154
    IS extended neighbor: r0_re.00, Metric: default 10 SubTLV len: 72
       IPv6 address: ::10.1.1.1
       Neighbor's IPv6 address: ::10.1.1.10
      Local interface index: 333, Remote interface index: 333
      P2P IPV6 Adj-SID - Flags:0xf0(F:1,B:1,V:1,L:1,S:0,P:0), Weight:0, Label: 24
      P2P IPv6 Adj-SID:      24, Weight:   0, Flags: FBVL—

IS-IS level 1 link-state database:
R2_re.00-00 Sequence: 0x6, Checksum: 0x9f29, Lifetime: 553 secs
   IS neighbor: R2_re.02                      Metric:       10
     Two-way fragment: R2_re.02-00, Two-way first fragment: R2_re.02-00
     LAN IPv6 Adj-SID:  299776, Weight:   0, Neighbor: R4_re, Flags: F-VL--
```

```
   IS neighbor: R1_re.02                      Metric:       10
     Two-way fragment: R1_re.02-00, Two-way first fragment: R1_re.02-00
     LAN IPv6 Adj-SID:  299792, Weight:   0, Neighbor: R1_re, Flags: F-VL--
   V6 prefix: 2::/24                          Metric:        0 Internal Up
   V6 prefix: 10:12:12::/96                   Metric:       10 Internal Up
   V6 prefix: 10:24:24::/96                   Metric:       10 Internal Up
…..
  TLVs:
    Area address: 47.0005.80ff.f800.0000.0108.0001 (13)
    LSP Buffer Size: 1492
    Speaks: IP
    Speaks: IPV6
    IP router id: 10.2.2.2
    IP address: 10.2.2.2
    Hostname: R2_re
    IS neighbor: r5_re0.00, Internal, Metric: default 10
    IS neighbor: r0_re0.00, Internal, Metric: default 10
    IS neighbor: r0_re0.00, Internal, Metric: default 10
    Extended IS Reachability TLV, Type: 22, Length: 120
    IS extended neighbor: r5_re0.00, Metric: default 10 SubTLV len: 29
      IP address: 10.40.1.3
      Neighbor's IP address: 10.40.1.5
      Local interface index: 334, Remote interface index: 334
      P2P IPV6 Adj-SID - Flags:0xb0(F:1,B:0,V:1,L:1,S:0,P:0), Weight:0, Label: 18
      P2P IPv6 Adj-SID:      18, Weight:   0, Flags: F-VL--
      +P2P SRV6 END-X-SID:2001:db8:7000::11 , Flags:--P, Weight:0
       + Flags:0x20(B:0,S:0,P:1), Flavor: USD
     ….
     ….
     SRv6 Locator: 2001:db8:2::/48, Metric: 0, MTID: 0, Flags: 0
      SRv6 SID: 2001:db8:2::2:2, Flavor: PSP, USP, USD
      SRv6 SID: 2001:db8:2::2:3, Flavor: USD
      SRv6 SID: 2001:db8:2::2:2, Flavor: None
    Router Capability:  Router ID 10.2.2.2, Flags: 0x00
      SPRING Algorithm - Algo: 0
      SRv6 Capability - Flags: 0
      Node MSD Advertisement Sub-TLV, Type: 23, Length: 2
        SRv6 Maximum Segments Left MSD Sub-TLV, Type:41, Value: 6
IPv6 prefix: 2001:db8:2::/48 Metric 0 Up
    IPv6 prefix: 10:24:24::/96 Metric 10 Up
    IPv6 prefix: 10:12:12::/96 Metric 10 Up
    IS neighbor: R2_re.02, Internal, Metric: default 10
```

```
    IS neighbor: R1_re.02, Internal, Metric: default 10
    Extended IS Reachability TLV, Type: 22, Length: 68
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear isis database | 3667

show isis flex-algorithm

IN THIS SECTION

- Syntax | 21732
- Description | 21732
- Options | 21733
- Required Privilege Level | 21733
- Output Fields | 21733
- Sample Output | 21735
- Release Information | 21736

Syntax

```
show isis (spring | spf | route) flex-algorithm flex-algorithm-id
logical-system (all |logical-system-name)
```

Description

Display the configured flexible algorithm and the defined parameters.

Options

flex-algorithm-id	Specify the ID of the flex algorithm to view its configured parameters.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 721 on page 21733 describes the output fields for the `show isis flex-algorithm` command. Output fields are listed in the approximate order in which they appear.

Table 721: show isis flex-algorithm Output Fields

Field Name	Field Description
flex algorithm	Name of the defined flexible algorithm.
color	Name of the BGP colored community mapped with the flexible algorithm.
Chosen Definition	Name of the flexible algorithm definition that the device has chosen to participate in.
calculation-type	The calculation type specified in the flexible algorithm definition (FAD).
metric-type	The metric type specified in the FAD
priority	The specified priority for the advertised flexible algorithm.
Include Any	The administrative groups included in the FAD.
Include All	All the administrative groups included in the FAD.

Table 721: show isis flex-algorithm Output Fields *(Continued)*

Field Name	Field Description
Exclude	The administrative groups to be excluded.
Prefix	Destination of the route.
L	IS-IS level: • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2
Version	Version of SPF that generated the route.
Metric	Metric value associated with the route.
Type	Metric type: `int` (internal) or `ext` (external).
Interface	Interface to the next hop.
NH	Next hop address type: • IPv4 • IPv6 • MPLS
Via	System identifier of the next hop, displayed as a name if possible.

Sample Output

show isis spring flex-algorithm flex-algorithm-id

```
user@host> show isis spring flex-algorithm flex-algorithm-id 129
Flex algortihm : 129, Paritcipating
Color: <color>
Chosen Definition: <hostname>
calculation-type: spf/strict-spf
    metric-type: igp-metric/te-metric
priority: <priority>
Include Any: <admin-groups>
Include All: <admin-groups>
Exclude: <admin-groups>
```

show isis spf log flex-algorithm id

```
user@host> show isis spf log flex-algorithm-id 129
IS-IS level 1 SPF log:

 IS-IS level 2 SPF log:
Start time          Elapsed (secs) Count Reason
Mon Jun 24 11:22:29        0.003264   15 Updated LSP r4.00-00
Mon Jun 24 11:22:36        0.003769    9 Flushing adjacencies
```

show isis route inet flex-algorithm-id

```
user@host> show isis route inet flex-algorithm-id 129
IS-IS routing table             Current version: L1: 0 L2: 857
IPv4/IPv6 Routes
----------------
Prefix             L Version   Metric Type Interface       NH   Via                 Backup Score
10.1.1.1/32        2     857       30 int  ge-0/0/4.1      IPV4 r3
10.2.2.2/32       2     857       20 int  ge-0/0/4.1      IPV4 r3
10.3.3.3/32       2     857       10 int  ge-0/0/4.1      IPV4 r3
10.4.4.4/32       2     857       30 int  ge-0/0/7.1      IPV4 r6
```

```
10.5.5.5/32       2     857       20 int  ge-0/0/7.1      IPV4 r6
10.6.6.6/32       2     857       10 int  ge-0/0/7.1      IPV4 r6
10.10.10.10/32     2     857       40 int  ge-0/0/4.1      IPV4 r3
```

Release Information

Command introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

definition | **5231**

flex-algorithm | **7221**

How to Configure Flexible Algorithms in IS-IS for Segment Routing Traffic Engineering

show isis hostname

IN THIS SECTION

- Syntax | **21736**
- Syntax (EX Series Switches and QFX Series) | **21737**
- Description | **21737**
- Options | **21737**
- Required Privilege Level | **21737**
- Output Fields | **21737**
- Sample Output | **21738**
- Release Information | **21738**

Syntax

```
show isis hostname
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show isis hostname
```

Description

Display IS-IS hostname database information.

This command displays the system ID-to-name cache. The output shows if the mapping has been learned by receipt of a Hostname TLV #137 (type dynamic) configured in Junos OS with the `set system host-name` command, or a static mapping defined in Junos OS with the `set system static-host-mapping` *`hostname`* `sysid` command (type static). The local router always has its type set to static even if `static-host-mapping` is not configured.

Options

none	Display IS-IS hostname database information.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 722 on page 21737 describes the output fields for the `show isis hostname` command. Output fields are listed in the approximate order in which they appear.

Table 722: show isis hostname Output Fields

Field Name	Field Description
System Id	System identifier mapped to the hostname.
Hostname	Hostname mapped to the system identifier.

Table 722: show isis hostname Output Fields *(Continued)*

Field Name	Field Description
Type	Type of mapping between system identifier and hostname. • Dynamic—Hostname mapping determined as described in RFC 2763, *Dynamic Hostname Exchange Mechanism for IS-IS*. • Static—Hostname mapping configured by user.

Sample Output

show isis hostname

```
user@host> show isis hostname
   IS-IS hostname database:
   System Id      Hostname                                        Type
   1921.6800.4201 isis1                                           Dynamic
   1921.6800.4202 isis2                                           Static
   1921.6800.4203 isis3                                           Dynamic
```

Release Information

Command introduced before Junos OS Release 7.4.

show isis interface

IN THIS SECTION

- Syntax | 21739
- Syntax (EX Series Switches and QFX Series) | 21739
- Description | 21739
- Options | 21740

- Required Privilege Level | 21740
- Output Fields | 21740
- Sample Output | 21744
- Release Information | 21745

Syntax

```
show isis interface
<brief | detail | extensive>
<interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show isis interface
<brief | detail | extensive>
<interface-name>
```

Description

Display status information about Intermediate System-to-Intermediate System (IS-IS)-enabled interfaces.

> **NOTE**: If the configured metric for an IS-IS level is above 63, and the `wide-metrics-only` statement is not configured, the `show isis interface detail` command and the `show isis interface extensive` command display `63` as the metric value for that level. Configure the `wide-metrics-only` statement to generate metric values greater than 63 on a per IS-IS level basis.
>
> The `show isis interface` command displays the configured metric value for an IS-IS level irrespective of whether is configured or not.

Options

none	Display standard information about all IS-IS-enabled interfaces.
brief | detail | extensive	(Optional) Display the specified level of output.
interface-name	(Optional) Display information about the specified interface only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 723 on page 21740 describes the output fields for the `show isis interface` command. Output fields are listed in the approximate order in which they appear.

Table 723: show isis interface Output Fields

Field Name	Field Description	Level of Output
interface-name	Name of the interface.	`detail`
Designated router	Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.	`detail`
Index	Interface index assigned by the Junos OS kernel.	`detail`
State	Internal implementation information.	`detail`

Table 723: show isis interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Circuit id	Circuit identifier. **NOTE**: Each IS-IS interface is assigned a circuit ID value to identify the interface within the linkstate database. All interfaces (loopback, broadcast, and so on) and all point-to-point links share the locally significant value of 0x01, and this value is not incremented.	`detail`
Circuit type	Circuit type: • `1`—Level 1 only • `2`—Level 2 only • `3`—Level 1 and Level 2	`detail`
LSP interval	Interval between link-state PDUs sent from the interface.	`detail`
CSNP interval	Interval between complete sequence number PDUs sent from the interface.	`detail extensive`
Sysid	System identifier.	`detail`
Interface	Interface through which the adjacency is made.	none `brief`
L or Level	Level: • `1`—Level 1 only • `2`—Level 2 only • `3`—Level 1 and Level 2 **NOTE**: The default IS-IS level on loopback interfaces are always same as the IS-IS level configured on other IS-IS interfaces in a router. You can also configure IS-IS level on loopback interfaces per your requirement.	All levels

Table 723: show isis interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
CirID	Circuit identifier.	none `brief`
Level 1 DR	Level 1 designated intermediate system.	none `brief`
Level 2 DR	Level 2 designated intermediate system.	none `brief`
L1/L2 Metric	Interface's metric for Level 1 and Level 2. If there is no information, the metric is `0`.	none `brief`
Flood-group Area-ID	Flood-group is configured on a specific IS-IS interface **NOTE**: Seen only when Flood-group is configured.	`detail extensive`
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their label-switched paths (LSPs).	`detail extensive`
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs.	`detail extensive`
Adjacencies	Number of adjacencies established on this interface.	`detail`
Priority	Priority value for this interface.	`detail`
Metric	Metric value for this interface.	`detail`
Hello(s) / Hello Interval	Interface's hello interval.	`detail extensive`
Hold(s) / Hold Time	Interface's hold time.	`detail extensive`

Table 723: show isis interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Post convergence Protection	Post convergence protection can have the following types when enabled,
	• `Fate Sharing` can have the following values
	  • **Yes**-You have configured fate-sharing protection.
	  • **No**-You have not configured fate-sharing protection.
	• `node protection` can have the following values:
	  • **Yes**–You have configured node protection.
	  • **No**–You have not configured node protection.
	• `srlg protection` can have the following values:
	  • **Yes**–You have configured Shared Risk Link Groutp (SRLG) protection.
	  • **No**–You have not configured SRLG protection.
	`Node cost` is the recalculated metric cost of the node.
		`extensive`

Designated Router	Router responsible for sending network link-state advertisements, which describe all the routing devices attached to the network.	`detail`
Hello padding	Type of hello padding: • `Adaptive`—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state TLV. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. • `Loose`—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. • `Strict`—Padding is performed on all interface types and for all adjacency states, and is continuous.	`extensive`

Table 723: show isis interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
LDP sync state	Current LDP synchronization state: `in sync`, `in holddown`, or `not supported`.	`extensive`
reason	Reason for being in the LDP sync state.	`extensive`
config holdtime	Configured value of the hold timer.	`extensive`
remaining	If the state is not in sync and the hold time is not infinity, then this field displays the remaining hold time in seconds.	`extensive`
IIH max size	Configured value of IS-IS hello packets	`extensive`

Sample Output

show isis interface

```
user@host> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
at-2/3/0.0            3   0x1 Point to Point    Point to Point         10/10
lo0.0                 3   0x1 Passive           Passive                 0/0
```

show isis interface brief

The output for the `show isis interface brief` command is identical to that for the `show isis interface` command. For sample output, see "show isis interface" on page 21744.

show isis interface detail

```
user@host> show isis interface detail
IS-IS interface database:
at-2/3/0.0
```

```
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1             1       64     10     9.000       27
    2             1       64     10     9.000       27
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Adjacency advertisement: Advertise
  Protection Type: Node Link, No eligible Backup
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1             0       64      0 Passive
    2             0       64      0 Passive
```

show isis interface extensive

```
user@host> show isis interface extensive
IS-IS interface database:
xe-6/1/0.0
  Index: 75, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding, IIH max size: 1505
  Adjacency advertisement: Advertise

  Flood-group Area-ID: 49.0001
  Level 1
    Adjacencies: 0, Priority: 64, Metric: 10
    Disabled
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 20.000 s, Hold Time: 60 s
    Post convergence Protection: Enabled, Fate sharing: Yes, Srlg: Yes, Node cost: 0
    Designated Router: nemean.03
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding Wide IS-IS Metrics for Traffic Engineering

Example: Enabling Wide IS-IS Metrics for Traffic Engineering

show isis interface-group

IN THIS SECTION

- Syntax | 21746
- Description | 21746
- Options | 21746
- Required Privilege Level | 21747
- Output Fields | 21747
- Sample Output | 21748
- Release Information | 21748

Syntax

```
show isis interface-group
<instance instance>
<logical-system (all| logical-system-name)>
```

Description

Display status information about the specified interface group.

Options

none	Display standard status information about the interface group.
instance	Display the status information about the interface group for the specified instance.

logical-system (all | logical-system-name)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 724 on page 21747 describes the output fields for the `show isis interface-group` command. Output fields are listed in the approximate order in which they appear.

Table 724: show layer2-map Output Fields

Field Name	Field Description
Interface-group	Name of the interface group followed by the interfaces that belong to the specified group and their status.
Total Nominal Bandwidth	Minimum bandwidth reserved for this interface group.
Total Actual Bandwidth	Total bandwidth that is actually available.
Level	IS-IS Level • 1—Level 1 • 2—Level 2
Address family	Address family can have one of the following values: • IPv6 Unicast • IPv6 Multicast
Metric	Metric value for this interface group.

Table 724: show layer2-map Output Fields *(Continued)*

Field Name	Field Description
Bandwidth Treshold	The maximum bandwidth allowed for this interface group.

Sample Output

show isis interface-group

```
user@host> show isis interface-group
Interface-group: R2-B
  et-8/0/0.0, 100Gbps, Up
  et-9/0/0.0, 100Gbps, Up
  Total Nominal Bandwidth: 200Gbps, Total Actual Bandwidth: 200Gbps
  Level 1
    IPV6 Unicast, Metric: 30
      Bandwidth Threshold: 100Gbps, Metric: 1000
      Bandwidth Threshold: 200Gbps, Metric: 30
    IPV6 Multicast, Metric: 30
      Bandwidth Threshold: 100Gbps, Metric: 1000
      Bandwidth Threshold: 200Gbps, Metric: 30
```

Release Information

Command introduced in Junos OS Release 16.1.

show isis layer2-map

IN THIS SECTION

- Syntax | 21749
- Description | 21749
- Options | 21749

- Required Privilege Level | 21749
- Output Fields | 21750
- Sample Output | 21750
- Release Information | 21751

Syntax

```
show isis layer2-map
<destination>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display the details of the Layer 2 ARP or neighbor discovery next hops and the mapped data link address in the kernel for the routing instances.

Options

none	Display all the Layer 2 mappings in the kernel for all supported address families for all routing instances.
destination	(Optional) Display the Layer 2 mapping for a specified destination address.
instance *instance-name*	(Optional) Display the Layer 2 mapping for the specified routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 725 on page 21750 describes the output fields for the `show isis layer2-map` command. Output fields are listed in the approximate order in which they appear.

Table 725: show isis layer2-map Output Fields

Field Name	Field Description
Interface	The interface name of the next hop.
IP address	IP address of the next hop.
IPv4 records	Total number of IPv4 records.
IPv6 records	Total number of IPv6 records.
SNPA	The subnetwork point of attachment (SNPA) indicates the data link address of the next hop.
State	State of the next hop. This field can have one of the following values: • Delete—The specified entry is deleted. • Disabled—The Layer 2 mapping is disabled. • Rexmit—This additional state is displayed along with other states for entries that need to be retransmitted • Sent—The entry was sent or added to the kernel.
Refresh	Time in seconds before the next hop is refreshed.

Sample Output

show isis layer2-map

```
user@host> show isis layer2-map logical-system all
logical-system: r2
```

```
Layer2 mapping database for instance master

IP Address                              Interface  SNPA              Refresh  State
12.0.0.1                                lt-1/2/0.21 b0:c6:9a:2c:d8:bc 00:14:34 Sent

  IPv4 records: 1
  IPv6 records: 0

-----

logical-system: r1

Layer2 mapping database for instance master

IP Address                              Interface  SNPA              Refresh  State
12.0.0.2                                lt-1/2/0.12 b0:c6:9a:2c:d8:bd 00:14:34 Sent

  IPv4 records: 1
  IPv6 records: 0
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Layer 2 Mapping for IS-IS

layer2-map | **9651**

show isis overview

IN THIS SECTION

- Syntax | **21752**
- Syntax (EX Series Switches and QFX Series) | **21752**

- Description | 21752
- Options | 21752
- Required Privilege Level | 21752
- Output Fields | 21753
- Sample Output | 21756
- Release Information | 21759

Syntax

```
show isis overview
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show isis overview
<instance instance-name>
```

Description

Display IS-IS overview information.

Options

none	Display standard overview information about IS-IS for all routing instances.
instance *instance-name*	(Optional) Display overview information for the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 726 on page 21753 lists the output fields for the `show isis overview` command. Output fields are listed in the approximate order in which they appear.

Table 726: show isis overview Output Fields

Field Name	Field Description
Hostname	Name of the router.
Sysid	Part of the ISO address of the routing device.
Areaid	The area number of the routing device.
Instance	IS-IS routing instance.
Router ID	Router ID of the routing device.
Adjacency holddown	Adjacency holddown capability: `enabled` or `disabled`.
Maximum Areas	Maximum number of IS-IS areas advertised by the routing device.
LSP life time	Lifetime of the link-state PDU, in seconds.
Filter low life time LSPs up to	LSPs with a lifetime lower than this value are filtered out.
Attached bit evaluation	Attached bit capability: `enabled` or `disabled`.
SPF delay	Delay before performing consecutive shortest-path-first (SPF) calculations.
SPF holddown	Delay before performing additional SPF calculations after the maximum number of consecutive SPF calculations is reached.

Table 726: show isis overview Output Fields *(Continued)*

Field Name	Field Description
SPF rapid runs	Maximum number of SPF calculations that can be performed in succession before the holddown timer begins.
Overload bit at startup is set	Overload bit capability is enabled.
Overload high metrics	Overload high metrics capability: `enabled` or `disabled`.
Allow internal prefix overloading	Allow internal prefixes to be advertised with high metric: `enabled` or `disabled`
Allow external prefix overloading	Allow external prefixes to be advertised with high metric: `enabled` or `disabled`
Overload timeout	Time period after which overload is reset and the time that remains before the timer is set to expire.
Traffic engineering	Traffic engineering capability: `enabled` or `disabled`.
Restart	Graceful restart capability: `enabled` or `disabled`.
Restart duration	Time period for complete reacquisition of IS-IS neighbors.
Helper mode	Graceful restart helper capability: `enabled` or `disabled`.
Level	IS-IS level: • `1`—Level 1 information • `2`—Level 2 information
IPv4 is enabled	IP Protocol version 4 capability is enabled.

Table 726: show isis overview Output Fields *(Continued)*

Field Name	Field Description
IPv6 is enabled	IP Protocol version 6 capability is enabled.
Micro-loop avoidance	Micro-loop avoidance is enabled. Generally adjacent nodes converge faster than neighboring nodes causing traffic to loop. A route convergence delay is configured to avoid such micro loops.
Internal route preference	Preference value of internal routes.
External route preference	Preference value of external routes.
Prefix export limit	Number of prefixes allowed to be exported, as configured by the `prefix-export-limit` statement.
Prefix export count	Number of prefixes exported.
Wide area metrics are enabled	Wide area metrics capability is enabled.
Narrow metrics are enabled	Narrow metrics capability is enabled.
Adjacency holddown is active	IS-IS adjacencies come up one after another when adjacency holddown is enabled.
SRv6	Indicates the SRv6 capability of the device.
Locator	END-SID is attached to this locator address.
END-SID	Function for the prefix segments.

Table 726: show isis overview Output Fields *(Continued)*

Field Name	Field Description
Flavor	Behavior of the END-SID function. This field can have the following values: • PSP—Penultimate segment pop of the segment routing header (SRH). • USP—Ultimate segment pop of the SRH. • USD—Ultimate segment decapsulation of the SRH.

Sample Output

show isis overview

```
user@host> show isis overview
Instance: master
  Router ID: 10.255.107.183
  Hostname: pro-bng3-a
  Sysid: 0192.0168.0001
  Areaid: 49.0002
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Filter low life time LSPs up to: 300
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  Overload bit at startup is set
    Overload high metrics: disabled
    Allow route leaking: disabled
    Allow internal prefix overloading: enabled
    Allow external prefix overloading: enabled
    Overload timeout: 60 sec
  IPv4 is enabled, IPv6 is enabled
Micro-loop avoidance: Enabled
    Method: Route Convergence Delay, Route convergence delay: 5000 msec
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
  Level 1
```

```
    Internal route preference: 15
    External route preference: 160
    Wide metrics are enabled, Narrow metrics are enabled
    Adjacency holddown is active
  Level 2
    Internal route preference: 18
    External route preference: 165
    Prefix export limit: 5, Prefix export count: 5
    Wide metrics are enabled
    Adjacency holddown is active
```

command-name

```
user@host> show isis overview (Segment Routing)
Source Packet Routing (SPRING): Enabled
    SRGB Config Range :
      SRGB Start-Label : 40000, SRGB Index-Range : 4000
    SRGB Block Allocation: Success
      SRGB Start Index : 40000, SRGB Size : 4000, Label-Range: [ 40000, 43999 ]
    Node Segments: Enabled
      Ipv4 Index : 11, Ipv6 Index : 12
  SRv6: Enabled
     Locator: 7000::/64, Algorithm: 128
       END-SID: 7000::1, Flavor: PSP, USP, USD
       END-SID: 7000::2, Flavor: None
       END-SID: 7000::3, Flavor: PSP
  Post Convergence Backup: Disabled
  Level 1
    Internal route preference: 15
    External route preference: 160
```

command-name

```
user@host> show isis overview logical-system R2
Instance: master
  Router ID: 192.168.0.2
  Hostname: pro-bng3-a-R2
  Sysid: 0192.0168.0002
  Areaid: 49.0002
  Adjacency holddown: enabled
```

```
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
  Level 1
    Internal route preference: 15
    External route preference: 160
    Prefix export count: 0
    Wide metrics are enabled, Narrow metrics are enabled
  Level 2
    Internal route preference: 18
    External route preference: 165
    Prefix export count: 0
    Wide metrics are enabled, Narrow metrics are enabled

user@host> show isis overview logical-system R3
Instance: master
  Router ID: 192.168.0.3
  Hostname: pro-bng3-a-R3
  Sysid: 0192.0168.0003
  Areaid: 49.0002
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
  Level 1
    Internal route preference: 15
    External route preference: 160
    Prefix export count: 0
    Wide metrics are enabled, Narrow metrics are enabled
  Level 2
    Internal route preference: 18
    External route preference: 165
```

```
    Prefix export count: 0
    Wide metrics are enabled, Narrow metrics are enabled
```

Release Information

Command introduced in Junos OS Release 8.5.

show isis purge log

IN THIS SECTION

- Syntax | 21759
- Description | 21759
- Options | 21760
- Required Privilege Level | 21760
- Output Fields | 21760
- Sample Output | 21761
- Release Information | 21761

Syntax

```
show isis purge log
<instance instance-name>
<level level>
<logical-system (all | logical-system-name)>
```

Description

Displays the purge log history received from adjacent devices. The purge log can include the type, length and value (TLV) details and the Intermediate System (IS) identification. Use this command to identify the purge originator and view other details of recent purges. You can display the last 50 entries in the IS-IS purge log and filter the log history based on the IS-IS level or by routing instance.

Options

none	Display purge log of all routing instances.
instance *instance-name*	(Optional) Display purge log of the specified routing instance.
level *level*	(Optional) Display purge log of IS-IS routers that belong to the specified level.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 727 on page 21760 describes the output fields for the `show isis purge log` command. Output fields are listed in the approximate order in which they appear.

Table 727: show isis purge log Output Fields

Field Name	Field Description
IS-IS Level	Level: • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address.
Time	Date and time in hours, minutes, and seconds when the purge occurred.
LSP Name	Name of the LSP where the purge occurred
Originator	Source of the purge

Table 727: show isis purge log Output Fields *(Continued)*

Field Name	Field Description
Received	Router that received the purge instruction

Sample Output

show isis purge log

```
user@host> show isis purge log
IS-IS level Level: 1 Purge log:
Time                   LSP Name             Originator         Received
Wed Aug  5 17:52:32    VMX_R5_re.00-00       VMX_R6_re        VMX_R4_re
Wed Aug  5 17:52:32    VMX_R5_re.02-00       VMX_R6_re        VMX_R4_re
Wed Aug  5 17:52:32    VMX_R2_re.00-00       VMX_R6_re        VMX_R4_re
Wed Aug  5 17:52:32    VMX_R3_re.00-00       VMX_R6_re        VMX_R4_re
Wed Aug  5 17:52:32    VMX_R1_re.00-00       VMX_R6_re        VMX_R4_re
Wed Aug  5 17:52:32    VMX_R6_re.00-00       VMX_R6_re   0000.0000.0000
Wed Aug  5 17:52:45    VMX_R6_re.00-00       VMX_R3_re        VMX_R3_re
Wed Aug  5 17:52:45    VMX_R6_re.00-00       VMX_R3_re        VMX_R3_re
Wed Aug  5 17:52:45    VMX_R5_re.00-00       VMX_R3_re        VMX_R3_re
Wed Aug  5 17:52:45    VMX_R5_re.02-00       VMX_R3_re        VMX_R3_re
Wed Aug  5 17:52:45    VMX_R2_re.00-00       VMX_R3_re        VMX_R3_re
Wed Aug  5 17:52:45    VMX_R2_re.02-00       VMX_R3_re        VMX_R3_re
```

Release Information

Command introduced in Junos OS Release 16.2R1.

show isis route

IN THIS SECTION

- Syntax | 21762
- Syntax (EX Series Switches and QFX Series) | 21762
- Description | 21762
- Options | 21763
- Required Privilege Level | 21763
- Output Fields | 21763
- Sample Output | 21764
- Release Information | 21768

Syntax

```
show isis route
<destination>
<inet | inet6>
<instance instance-name>
<logical-system (all | logical-system-name)>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>
```

Syntax (EX Series Switches and QFX Series)

```
show isis route
<destination>
<inet | inet6>
<instance instance-name>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>
```

Description

Display the routes in the IS-IS routing table.

Options

none	Display all routes in the IS-IS routing table for all supported address families for all routing instances.
destination	(Optional) Destination address for the route.
inet | inet6	(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.
instance *instance-name*	(Optional) Display routes for the specified routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)	(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.

Required Privilege Level

view

Output Fields

Table 728 on page 21763 describes the output fields for the `show isis route` command. Output fields are listed in the approximate order in which they appear.

Table 728: show isis route Output Fields

Field Name	Field Description
Current version	Number of the current version of the IS-IS routing table.
L1	Version of Level 1 SPF that was run.
L2	Version of Level 2 SPF that was run.
Prefix or Label	Destination of the route.

Table 728: show isis route Output Fields *(Continued)*

Field Name	Field Description
L	IS-IS level: • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2
Version	Version of SPF that generated the route.
Metric	Metric value associated with the route.
Type	Metric type: int (internal) or ext (external).
Interface	Interface to the next hop.
Via	System identifier of the next hop, displayed as a name if possible.
ISO Routes	ISO routing table entries.
snpa	MAC address.

Sample Output

show isis route logical-system

```
user@host> show isis route logical-system ls1
IS-IS routing table             Current version: L1: 8 L2: 11
Prefix             L Version Metric Type Interface    Via
10.9.7.0/30        2      11     20 int  gr-0/2/0.0   h
10.9.201.1/32      2      11     60 int  gr-0/2/0.0   h
IPV6 Unicast IS-IS routing table            Current version: L1: 9 L2: 11
```

```
Prefix             L Version Metric Type Interface    Via
2001:db8:8009:3::a09:3200/126 2      11     20 int gr-0/2/0.0  h
```

show isis route floodgroup

```
user@R2> show isis route floodgroup 49.0001
IS-IS routing table             Current version: L1: 14 L2: 27
IPv4/IPv6 Routes
----------------
Prefix             L Version   Metric Type Interface       NH   Via                 Backup Score
0.0.0.0/0          1      14       10 int  ge-0/0/8.0      IPV4 R1
10.81.3.3/32       1      14       10 int  ge-0/0/8.0      IPV4 R1
10.128.17.202/32   1      14       10 int  ge-0/0/8.0      IPV4 R1
```

show isis route (CLNS)

```
user@host> show isis route
IS-IS routing table             Current version: L1: 10 L2: 8
IPv4/IPv6 Routes
Prefix             L Version   Metric Type Interface    Via
0.0.0.0/0          1      10       10 int  fe-0/0/1.0   ISIS.0
ISO Routes
Prefix  L   Version   Metric Type Interface    Via    snpa
0/0
        1        10       10 int  fe-0/0/1.0   isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001/104
        1        10        0 int 47.0005.80ff.f800.0000.0108.0001.1921.6800.4001/152
        1        10       10 int  fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001.1921.6800.4002/152
        1        10       20 int  fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0002/104
        1        10        0 int
47.0005.80ff.f800.0000.0108.0002.1921.6800.4001/152
        1        10       10 int  fe-0/0/1.0 isis.0 0:12:0:34:0:56
```

show isis route

```
user@host> show isis route

IS-IS routing table            Current version: L1: 4 L2: 13
IPv4/IPv6 Routes
----------------
Prefix                   L   Version  Metric Type Interface       NH   Via
10.255.71.52/32          2      13       10   int  ae0.0           IPV4 camaro
10.255.71.238/32         2      13       20   int  so-6/0/0.0      IPV4 olympic
                                                   as0.0           IPV4 glacier
10.255.71.239/32         2      13       20   int  so-6/0/0.0      IPV4 olympic
                                                   ae0.0           IPV4 camaro
10.255.71.242/32         2      13       10   int  as0.0           IPV4 glacier
10.255.71.243/32         2      13       10   int  so-6/0/0.0      IPV4 olympic
10.12.13.0/30            2      13       20   int  so-6/0/0.0      IPV4 olympic
10.12.15.0/30            2      13       20   int  so-6/0/0.0      IPV4 olympic
10.13.15.0/30            2      13       30   int  ae0.0           IPV4 camaro
                                                   so-6/0/0.0      IPV4 olympic
                                                   as0.0           IPV4 glacier
10.13.16.0/30            2      13       25   int  as0.0           IPV4 glacier
10.14.15.0/30            2      13       20   int  ae0.0           IPV4 camaro
10.192.1.0/30            2      13       30   int  so-6/0/0.0      IPV4 olympic
                                                   as0.0           IPV4 glacier
2001:db8:1eee::/64       2      13       30   int  so-6/0/0.0      IPV6 olympic
                                                   as0.0           IPV6 glacier
2001:db8::10:255:71:52/128   2  13       10   int  ae0.0           IPV6 camaro
2001:db8::10:255:71:238/128  2  13       20   int  so-6/0/0.0      IPV6 olympic
                                                   as0.0           IPV6 glacier
2001:db8::10:255:71:239/128  2  13       20   int  so-6/0/0.0      IPV6 olympic
                                                   ae0.0           IPV6 camaro
2001:db8::10:255:71:242/128  2  13       10   int  as0.0           IPV6 glacier
2001:db8::10:255:71:243/128  2  13       10   int  so-6/0/0.0      IPV6 olympic


MPLS Routes
-----------
Label              L Version   Metric Type Interface       NH     Via
300032 /52         2      38        0 int  lt-1/2/0.13     MPLS   Direct forward to
10.0.7.60(pro-bng3-c-E)
300048 /52         1      27        0 int  lt-1/2/0.12     MPLS   Direct forward to
10.0.6.60(pro-bng3-c-E)
```

```
300064 /52         1      27        0 int  lt-1/2/0.12     MPLS   Direct forward to
10.0.6.60(pro-bng3-c-E)
300080 /52         2      38        0 int  lt-1/2/0.12     MPLS   Direct forward to
10.0.6.60(pro-bng3-c-E)
300096 /52         2      38        0 int  lt-1/2/0.12     MPLS   Direct forward to
10.0.6.60(pro-bng3-c-E)
299920 /52         1      27        0 int  lt-1/2/0.14     MPLS   Direct forward to
10.0.10.70(pro-bng3-c-F)
299936 /52         1      27        0 int  lt-1/2/0.14     MPLS   Direct forward to
10.0.10.70(pro-bng3-c-F)
299952 /52         2      38        0 int  lt-1/2/0.14     MPLS   Direct forward to
10.0.10.70(pro-bng3-c-F)
299968 /52         2      38        0 int  lt-1/2/0.14     MPLS   Direct forward to
10.0.10.70(pro-bng3-c-F)
299984 /52         1      27        0 int  lt-1/2/0.13     MPLS   Direct forward to
10.0.7.60(pro-bng3-c-E)
300000 /52         1      27        0 int  lt-1/2/0.13     MPLS   Direct forward to
10.0.7.60(pro-bng3-c-E)
300016 /52         2      38        0 int  lt-1/2/0.13     MPLS   Direct forward to
10.0.7.60(pro-bng3-c-E)
```

command-name

```
user@host>show isis route (SRv6)
IS-IS routing table             Current version: L1: 12 L2: 0
IPv4/IPv6 Routes
----------------
Prefix             L Version   Metric Type Interface       NH   Via                 Backup Score
::10.1.1.1/128      1      12       10 int  ge-0/0/1.0      IPV6 r1_re0
IPv4/IPv6->MPLS Routes
----------------------
Prefix             L Version   Metric Type Interface       NH   Via                 Backup Score
::10.1.1.1/128      1      12       10 int  ge-0/0/1.0      MPLS r1_re0(128.205.51.101)
                                           ge-0/0/0.0      MPLS r3_re0(128.205.51.103)

+SRV6 Routes
+-----------
+Prefix             L Version   Metric Type Interface       NH     Via
+2001:db8:1000::1/128        1      13        0 int
+2001:db8:1000::2/128        1      13        0 int
+2001:db8:1000::21/128       1      13        0 int  ge-0/0/0.0      IPV6   R1
```

```
+SRv6->INET6.3 Routes
+--------------------
+Prefix             L Version   Metric Type Interface       NH   Via                 +Backup
Score
+2001:db8:3000::1/128        1      13       10 int  ge-0/0/4.0      SRv6 TCNH
```

Release Information

Command introduced before Junos OS Release 7.4.

show isis route download priority

IN THIS SECTION

- Syntax | 21768
- Description | 21769
- Options | 21769
- Required Privilege Level | 21770
- Output Fields | 21770
- Sample Output | 21771
- Release Information | 21771

Syntax

```
show isis route download priority
<destination>
<floodgroup floodgroup>
<inet | inet6>
<instance instance-name>
<iso>
<logical-system (all | logical-system-name>
```

```
<nexthop interface | neighbor>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>
```

Description

Display the routes in the IS-IS routing table in the sequence of set priority. IS-IS routes set with a high priority are displayed first followed by medium and low priority prefixes. Internet Service Providers (ISP) can configure route priority to ensure faster convergence for important customers.

Configure a routing policy to assign a priority to IS-IS routes. These routes are updated in the routing table in the order of their priority. In the event of an IS-IS topology change, high priority prefixes are updated in the routing table first followed by medium and then low priority prefixes.

Options

none	Display all routes in the IS-IS routing table in order of their priority for all supported address families for all routing instances.
destination	(Optional) Destination address for the route for which you want to display the priority.
floodgroup *floodgroup*	(Optional) Display the priority of routes that belong to a specific floodgroup.
inet | inet6	(Optional) Display IS-IS route priority for inet (IPv4) or inet6 (IPv6) routes, respectively.
instance *instance-name*	(Optional) Display route priority for the specified routing instance only.
iso	(Optional) Display the priority of Connectionless Network Service (CLNS) routes.
nexthop *interface* | *neighbor*	Display the priority of routes reachable through a specific next hop.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)	(Optional) Display routes in order of the set priority for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.

Required Privilege Level

view

Output Fields

Table 729 on page 21770 describes the output fields for the `show isis route` command. Output fields are listed in the approximate order in which they appear.

Table 729: show isis route Output Fields

Field Name	Field Description
Current version	Number of the current version of the IS-IS routing table.
L1	Version of Level 1 SPF that was run.
L2	Version of Level 2 SPF that was run.
Prefix or Label	Destination of the route.
L	IS-IS level: • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2
Version	Version of SPF that generated the route.
Metric	Metric value associated with the route.
Type	Metric type: int (internal) or ext (external).
Interface	Interface to the next hop.
Via	System identifier of the next hop, displayed as a name if possible.

Table 729: show isis route Output Fields *(Continued)*

Field Name	Field Description
NH	Next hop address type: IPv4, IPv6, or MPLS.
Backup Score	Count of available backup paths.

Sample Output

show isis route download-priority

```
user@host> show isis route download-priority
IS-IS routing table             Current version: L1: 0 L2: 122
IPv4/IPv6 Routes
----------------
Prefix             L Version   Metric Type Interface       NH   Via                Backup Score
3.3.3.3/32         2     122       20 int  ge-1/0/1.0      IPV4 R2
                                           ge-5/0/9.0      IPV4 R2
2.2.2.2/32         2     122       10 int  ge-1/0/1.0      IPV4 R2
                                           ge-5/0/9.0      IPV4 R2
198.1.1.0/24        2     122       20 int  ge-1/0/1.0      IPV4 R2
                                           ge-5/0/9.0      IPV4 R2
198.1.2.0/24        2     122       20 int  ge-1/0/1.0      IPV4 R2
                                           ge-5/0/9.0      IPV4 R2
2001:db8:2:2::1/128  2     122       10 int  ge-1/0/1.0      IPV6 R2
                                           ge-5/0/9.0      IPV6 R2
2001:db8:3:3::3/128  2     122       20 int  ge-1/0/1.0      IPV6 R2
                                           ge-5/0/9.0      IPV6 R2
2001:db8:1:1::/64   2     122       20 int  ge-1/0/1.0      IPV6 R2
                                           ge-5/0/9.0      IPV6 R2
2001:db8:1:2::/64   2     122       20 int  ge-1/0/1.0      IPV6 R2
                                           ge-5/0/9.0      IPV6 R2
```

Release Information

Command introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show isis route | 21762

import (Protocols IS-IS) | 8476

show isis spf

IN THIS SECTION

- Syntax | 21772
- Syntax (EX Series Switches) | 21772
- Description | 21773
- Options | 21773
- Required Privilege Level | 21773
- Output Fields | 21773
- Sample Output | 21774
- Release Information | 21778

Syntax

```
show isis spf (brief | log | results)
<instance instance-name>
<level (1 | 2)>
<logical-system (all | logical-system-name)>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>
```

Syntax (EX Series Switches)

```
show isis spf (brief | log | results)
<instance instance-name>
<level (1 | 2)>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>
```

Description

Display information about IS-IS shortest-path-first (SPF) calculations.

Options

brief	Display an overview of SPF calculations.
log	Display the log of SPF calculations.
results	Display the results of SPF calculations.
instance *instance instance-name*	(Optional) Display SPF calculations for the specified routing instance.
level (1 | 2)	(Optional) Display SPF calculations for the specified IS-IS level.
log	Display the log of SPF calculations.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system.
results	Display the results of SPF calculations.
topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)	(Optional) Display SPF calculations for the specified topology only.

Required Privilege Level

view

Output Fields

Table 730 on page 21773 describes the output fields for the `show isis spf` command. Output fields are listed in the approximate order in which they appear.

Table 730: show isis spf Output Fields

Field Name	Field Description
Node	System ID of a node.

Table 730: show isis spf Output Fields *(Continued)*

Field Name	Field Description
Metric	Metric to the node.
Interface	Interface of the next hop.
Via	System ID of the next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
Start time	(`log` option only) Time that the SPF computation started.
Elapsed (secs)	(`log` option only) Length of time, in seconds, required to complete the SPF computation.
Count	(`log` option only) Number of times the SPF was triggered.
Reason	(`log` option only) Reason that the SPF computation was completed.

Sample Output

show isis spf log

```
user@host> show isis spf log logical-system lsl
 IS-IS level 1 SPF log:
Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000069    1 Reconfig
Fri Oct 31 12:41:18        0.000107    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000050    3 Address change on so-1/2/2.0
Fri Oct 31 12:41:23        0.000033    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28        0.000178    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59        0.000060    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30        0.000161    2 Multi area attachment change
Fri Oct 31 12:56:58        0.000198    1 Periodic SPF
```

```
Fri Oct 31 13:10:29        0.000209    1 Periodic SPF
 IS-IS level 2 SPF log:

Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000035    1 Reconfig
Fri Oct 31 12:41:18        0.000047    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000043    5 Address change on gr-0/2/0.0
Fri Oct 31 12:41:23        0.000022    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59        0.000144    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30        0.000257    3 New LSP skag.00-00
Fri Oct 31 12:54:37        0.000195    1 Periodic SPF
Fri Oct 31 12:55:50        0.000178    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55        0.000174    1 Updated LSP h.00-00
Fri Oct 31 12:55:58        0.000176    1 Updated LSP skag.00-00
Fri Oct 31 13:08:14        0.000198    1 Periodic SPF
IPV6 Unicast IS-IS level 1 SPF log:

Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000028    1 Reconfig
Fri Oct 31 12:41:18        0.000043    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000112    4 Updated LSP fix.00-00
Fri Oct 31 12:41:23        0.000059    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25        0.000041    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28        0.000103    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59        0.000040    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30        0.000118    2 Multi area attachment change
Fri Oct 31 12:56:08        0.000289    1 Periodic SPF
Fri Oct 31 13:11:07        0.000214    1 Periodic SPF
IPV6 Unicast IS-IS level 2 SPF log:

Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000027    1 Reconfig
Fri Oct 31 12:41:18        0.000039    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000049    6 Updated LSP fix.00-00
Fri Oct 31 12:41:23        0.000025    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25        0.000023    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59        0.000087    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30        0.000123    3 New LSP skag.00-00
Fri Oct 31 12:55:50        0.000121    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55        0.000121    1 Updated LSP h.00-00
Fri Oct 31 12:55:58        0.000121    1 Updated LSP skag.00-00
```

```
Fri Oct 31 13:09:46        0.000201    1 Periodic SPF
...
```

show isis spf results logical-system

```
user@host> show isis spf results logical-system ls1
IS-IS level 1 SPF results:
Node             Metric     Interface         Via             SNPA
scat.00          10         ge-1/1/0.0       scat            0:90:69:a6:48:9d
                 20         10.9.1.0/30
fix.02           10
fix.00           0
                 10         10.9.1.0/30
                 10         10.9.5.0/30
                 10         10.9.6.0/30
                 20         10.9.7.0/30
                 60         10.9.201.1/32
  3 nodes

 IS-IS level 2 SPF results:
Node             Metric     Interface         Via             SNPA
skag.00          20         gr-0/2/0.0       h
                 30         10.9.7.0/30
skag.02          20         gr-0/2/0.0       h
h.00             10         gr-0/2/0.0       h
                 20         10.9.6.0/30
                 20         10.9.7.0/30
                 60         10.9.201.1/32
fix.00           0
                 10         10.9.1.0/30
                 10         10.9.5.0/30
                 10         10.9.6.0/30
  4 nodes

IPV6 Unicast IS-IS level 1 SPF results:
Node             Metric     Interface         Via             SNPA
scat.00          10         ge-1/1/0.0       scat            0:90:69:a6:48:9d
                            ge-1/1/0.0       scat            0:90:69:a6:48:9d
                 20         8009:1::a09:1400/126
fix.02           10
fix.00           0
```

```
                 10         8009:1::a09:1400/126
                 10         8009:2::a09:1e00/126
                 20         8009:3::a09:3200/126
                 10         8009:4::a09:2800/126
  3 nodes

IPV6 Unicast IS-IS level 2 SPF results:
Node             Metric     Interface         Via             SNPA
skag.00          20         gr-0/2/0.0       h
                            gr-0/2/0.0       h
                 30         8009:3::a09:3200/126
skag.02          20         gr-0/2/0.0       h
                            gr-0/2/0.0       h
h.00             10         gr-0/2/0.0       h
                            gr-0/2/0.0       h
                 20         8009:3::a09:3200/126
                 20         8009:4::a09:2800/126
fix.00           0
                 10         8009:1::a09:1400/126
                 10         8009:2::a09:1e00/126
                 10         8009:4::a09:2800/126
  4 nodes

Multicast IS-IS level 1 SPF results:
Node             Metric     Interface         Via             SNPA
scat.00          10         ge-1/1/0.0       scat            0:90:69:a6:48:9d
fix.02           10
fix.00           0
  3 nodes

Multicast IS-IS level 2 SPF results:
Node             Metric     Interface         Via             SNPA
skag.00          20         gr-0/2/0.0       h
skag.02          20         gr-0/2/0.0       h
h.00             10         gr-0/2/0.0       h
fix.00           0
  4 nodes
...
```

show isis spf results (CLNS)

```
user@host> show isis spf results
IS-IS level 1 SPF results:
Node             Metric     Interface         Via             SNPA
skag.00 10                  fe-0/0/1.0       toothache       0:12:0:34:0:56
                            fe-0/0/1.0       toothache       0:12:0:34:0:56
                 20         192.168.37.64/29
                 10         1921.6800.4001
                 20         1921.6800.4002
pro1-a.02        10
pro1-a.00        0
                 0          10.255.245.1/32
                 10         192.168.37.64/29
                 0          1921.6800.4211
  3 nodes

IS-IS level 2 SPF results:
Node             Metric     Interface         Via             SNPA
skag.00 10                  fe-0/0/1.0       toothache       0:12:0:34:0:56
                            fe-0/0/1.0       toothache       0:12:0:34:0:56
                 20         10.255.245.1/32
                 20         192.168.37.64/29
                 20         47.0005.80ff.f800.0000.0109.0010/104
pro1-a.02        10
pro1-a.00        0
                 0          10.255.245.1/32
                 10         192.168.37.64/29
  3 nodes
```

Release Information

Command introduced before Junos OS Release 7.4.

show isis spring interface traffic-statistics

IN THIS SECTION

- Syntax | 21779
- Description | 21779
- Options | 21779
- Required Privilege Level | 21779
- Sample Output | 21780
- Release Information | 21780

Syntax

```
show isis spring interface traffic-statistics
logical-system (all |logical-system-name)
```

Description

Display the traffic statistics of an IS-IS SPRING link. SPRING traffic is measured to recalculate the actual available bandwidth to RSVP for traffic engineering.

Options

none	Display the traffic statistics of an IS-IS SPRING link.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Sample Output

show isis spring interface traffic-statistics

```
user@host> show isis spring interface traffic-statistics
Interface name     Bytes          Packets
ge-0/0/1.0         1287495        3000
ae0.0              2332435435     10233
```

Release Information

Command introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

- auto-bandwidth | 2682
- traffic-statistics (Protocols IS-IS) | 30048
- update-threshold-max-reservable | 30458
- clear isis spring traffic-statistics | 3673
- show auto-bandwidth | 17116

show isis spring sensor info

IN THIS SECTION

- Syntax | 21781
- Description | 21781
- Options | 21781
- Required Privilege Level | 21781
- Output Fields | 21781
- Sample Output | 21782
- Release Information | 21783

Syntax

```
show isis spring sensor info
logical-system (all |logical-system-name)
```

Description

Displays a list of sensors associated with the label IS-IS route and next hops for segment routing traffic. The command only displays the information related to the sensors and not the traffic statistics.

Options

none	Display the sensor information of an IS-IS SPRING route.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1 describes the output fields for the `show isis spring sensor info` command. Output fields are listed in the approximate order in which they appear.

Table 731: show isis spring sensor info Output Fields

Field	Field Description
Sensor-name	Represents the router or interface that the sensor is associated with.
Sensor-id	Unique number associated either with route or interface.

Sample Output

show isis spring sensor info

```
user@host> show isis spring sensor info
Per-interface-per-member-link Ingress Sensor:
---------------------------------------------
Sensor-name                        Sensor-id
aggr_ingress_intf_sensor           3221225484

Per-interface-per-member-link Egress Sensor:
---------------------------------------------
Sensor-name                        Sensor-id
ge-0/0/0.0                         3221225497
ge-0/0/1.0                         3221225498
ge-0/0/2.0                         3221225499

Per-sid Ingress Sensor:
----------------------
Sensor-name                        Sensor-id
16                                 3221225478
17                                 3221225479
18                                 3221225474
19                                 3221225475
20                                 3221225482
21                                 3221225483
22                                 3221225480
23                                 3221225481
24                                 3221225489
25                                 3221225490
400001                             3221225491
400002                             3221225492
400005                             3221225487
400006                             3221225488
400009                             3221225493
400010                             3221225494
400011                             3221225495
400012                             3221225496

IPv4/IPv6 Per-sid Egress Sensor:
--------------------------------
```

```
Sensor-name                    Sensor-id
L-ISIS-::10.10.10.1            3221225474
```

Release Information

Command introduced in 19.1R1.

RELATED DOCUMENTATION

- sensor-based-stats | 16555
- source-packet-routing (Protocols IS-IS) | 28266
- *Understanding Source Packet Routing in Networking (SPRING)*

show isis statistics

IN THIS SECTION

- Syntax | 21783
- Syntax (EX Series Switches and QFX Series) | 21784
- Description | 21784
- Options | 21784
- Required Privilege Level | 21784
- Output Fields | 21784
- Sample Output | 21786
- Release Information | 21787

Syntax

```
show isis statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show isis statistics
<instance instance-name>
```

Description

Display statistics about IS-IS traffic.

Options

none	Display IS-IS traffic statistics for all routing instances.
instance ***instance-name***	(Optional) Display statistics for the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 732 on page 21785 describes the output fields for the `show isis statistics` command. Output fields are listed in the approximate order in which they appear.

Table 732: show isis statistics Output Fields

Field Name	Field Description
PDU type	PDU type: • CSNP—Complete sequence number PDUs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU. • IIH—IS-IS hello packets are broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems. • LSP—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area. • PSNP—Partial sequence number PDUs are sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device. • Unknown—The PDU type is unknown.
Received	Number of PDUs received since IS-IS started or since the statistics were set to zero.
Processed	Number of PDUs received less the number dropped.
Drops	Number of PDUs dropped.
Sent	Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.
Rexmit	Number of PDUs retransmitted since IS-IS started or since the statistics were set to zero.
Total packets received/sent	Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.

Table 732: show isis statistics Output Fields *(Continued)*

Field Name	Field Description
SNP queue length	Number of CSPN and PSNP packets currently waiting in the queue for processing. This value is almost always `0`.
LSP queue length	Number of link-state PDUs waiting in the queue for processing. This value is almost always `0`.
SPF runs	Number of shortest-path-first (SPF) calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable.
Fragments rebuilt	Number of link-state PDU fragments that the local system has computed.
LSP regenerations	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software decides that a link-state PDU must be removed from the network.

Sample Output

show isis statistics

```
user@host> show isis statistics
IS-IS statistics for merino:

PDU type     Received  Processed      Drops       Sent     Rexmit
LSP             12227      12227          0       8184        683
IIH            113808     113808          0     115817          0
CSNP           198868     198868          0     198934          0
PSNP             6985       6979          6       8274          0
Unknown             0          0          0          0          0
Totals         331888     331882          6     331209        683

Total packets received: 331888 Sent: 331892
```

```
SNP queue length:           0 Drops:          0
LSP queue length:           0 Drops:          0

SPF runs:                1014
Fragments rebuilt:       1038
LSP regenerations:        425
Purges initiated:           0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear isis statistics | 3674

show l2-learning global-information

IN THIS SECTION

- Syntax | 21787
- Description | 21788
- Options | 21788
- Required Privilege Level | 21788
- Output Fields | 21788
- Sample Output | 21789
- Release Information | 21790

Syntax

```
show l2-learning global-information
```

Description

Display Layer 2 learning process-related information for the entire device.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 733 on page 21788 describes the output fields for the `show l2-learning global-information` command. Output fields are listed in the approximate order in which they appear.

Table 733: show l2-learning global-information Output Fields

Field Name	Field Description
MAC aging interval	Configured timeout interval, in seconds, for all MAC table entries.
MAC learning	Status of MAC learning: **Enabled** or **Disabled**.
MAC statistics	Status of MAC accounting: **Enabled** or **Disabled**.
MAC limit Count	Configured maximum limit on the number of MAC addresses that can be learned.
MAC limit hit flag	Status of the learned MAC limit hit flag: **Enabled** (the learned MAC exceeds the global MAC limit) or **Disabled** (the learned MAC does not exceed the global MAC limit).
MAC packet action drop	Status of action to drop packets after the configured MAC address limit is reached: **Enabled** (packets are dropped) or **Disabled** (packets are forwarded).
MAC+IP aging interval	Configured timeout interval, in seconds, for all MAC-IP table entries.

Table 733: show l2-learning global-information Output Fields *(Continued)*

Field Name	Field Description
MAC+IP limit Count	Configured maximum limit on the number of MAC-IP addresses that can be learned.
MAC+IP limit reached	Status of the learned MAC-IP limit reached flag: **Yes** (the learned MAC-IP exceeds the global MAC-IP limit) or **No** (the learned MAC-IP does not exceed the global MAC-IP limit).
RE state	State of routing engine: **Master** or **Backup**.
Fast Update	Status of evpn fast-mac-mobility update feature, **Enabled** or **Disabled**.
Host Pkts GBP src tag	GBP tag for host bound packets
Event Retry Time	Configured retry time interval, in seconds, for l2-learning events in case of failure.

Sample Output

show l2-learning global-information

```
user@host> show l2-learning global-information
Global Configuration:

MAC aging interval    : 300
MAC learning          : Enabled
MAC statistics        : Disabled
MAC limit Count       : 393215
MAC limit hit flag    : Disabled
MAC packet action drop: Disabled

MAC+IP aging interval : IPv4 - 1200 seconds
                        IPv6 - 1200 seconds
MAC+IP limit Count    : 393215
MAC+IP limit reached  : No
RE state              : Master
Fast Update           : Disabled
```

```
Host Pkts GBP src tag : 0
Event Retry Time      : 600
```

Release Information

Command introduced in Junos OS Release 8.4.

show l2-learning global-mac-count

IN THIS SECTION

- Syntax | 21790
- Description | 21790
- Options | 21790
- Required Privilege Level | 21791
- Output Fields | 21791
- Sample Output | 21791
- Release Information | 21791

Syntax

```
show l2-learning global-mac-count
```

Description

(MX Series routers only) Display the total number of dynamic and static MAC addresses learned for the entire router.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Displays the total number of dynamic and static MAC addresses learned for the entire router.

Sample Output

show l2-learning global-mac-count

```
user@host> show l2-learning global-mac-count
100 dynamic and static MAC addresses learned globally
```

Release Information

Command introduced in Junos OS Release 9.3.

show l2-learning instance

IN THIS SECTION

- Syntax | 21792
- Description | 21792
- Options | 21792
- Required Privilege Level | 21792
- Output Fields | 21792
- Sample Output | 21793
- Release Information | 21793

Syntax

```
show l2-learning instance
```

Description

Display Layer 2 learning properties for all the configured routing instances.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 734 on page 21792 describes the output fields for the `show l2-learning instance` command. Output fields are listed in the approximate order in which they appear.

Table 734: show l2-learning instance Output Fields

Field Name	Field Description
Routing Instance	Name of routing instance.
Bridging Domain	Name of bridging domain. On MX Series routers you can use the `show l2-learning instance <extensive>` command option to display the Bridge Service-id information which includes the Config Service ID and the Active Service ID.
Index	Number associated with the routing instance or bridging domain.
Logical System	Name of logical system or **Default** if no logical system is configured.

Table 734: show l2-learning instance Output Fields *(Continued)*

Field Name	Field Description
Routing instance flags	Status of Layer 2 learning properties for each routing instance: • **DL**—MAC learning is disabled. • **SE**—MAC accounting is enabled. • **AD**—Packets are dropped after MAC address limit is reached. • **LH**—The maximum number of MAC addresses has been learned on the routing instance. The routing instance is not able to learn any additional MAC addresses.
MAC limit	Maximum number of MAC addresses that can be learned from each interface in the routing instance or bridging domain.

Sample Output

show l2-learning instance

```
user@host> show l2-learning instance
Information for routing instance:

Routing Instance flags (DL -disable learning, SE -stats enabled,
                        AD -packet action drop, LH -mac limit hit)

Routing              Bridging     Index   Logical       Routing     MAC
Instance             Domain               System        flags       limit
__juniper_private1__                  1   Default                    5000
vs1                  vlan100          3   Default                    5120
vs1                  vlan200          4   Default                    5120
```

Release Information

(MX Series routers only) Command introduced in Junos OS Release 8.4.

show l2-learning interface

IN THIS SECTION

- Syntax | 21794
- Description | 21794
- Options | 21794
- Required Privilege Level | 21794
- Output Fields | 21794
- Sample Output | 21795
- Release Information | 21796

Syntax

```
show l2-learning interface
```

Description

(MX Series routers only) Display Layer 2 learning information for all the interfaces.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 735 on page 21795 describes the output fields for the `show l2-learning interface` command. Output fields are listed in the approximate order in which they appear.

Table 735: show l2-learning interfaceOutput Fields

Field Name	Field Description
Logical interface	Name of the logical interface.
Index	Index of the interface.
Routing Instance	Number of the routing instance to which the interface belongs.
Interface device	Value of the order in which the Junos OS finds and initializes the interface.
Logical interface flags	Status of Layer 2 learning properties for each interface: • **DL**—MAC learning is disabled. • **SE**—MAC accounting is enabled. • **AD**—Packets are dropped after the MAC interface limit is reached. • **MAC limit**—Maximum number of MAC addresses that can be learned from the interface. • `MP`—MAC Pinning enabled.

Sample Output

show l2-learning interface

```
user@host> showl2-learning interface
Information for interface family:

Logical Interface flags (DL -disable learning, SE -stats enabled,
                         AD -packet action drop, LH -mac limit hit)

Logical           Index  Routing   Interface   Logical              MAC
interface                instance  device      Interface flags      limit
ge-11/0/3.0          79         3        136                          1024
ge-11/1/4.100        84         3        150                          1024
ge-11/1/1.100        86         3        147                          1024
```

```
ge-11/1/0.100        87         3        146                        1024
xe-10/2/0.100        88         3        144                        1024
xe-10/0/0.100        89         3        129                        1024
ge-11/1/0.200        90         4        146                        1024
ge-11/1/1.200        91         4        147                        1024
ge-11/1/4.200        92         4        150                        1024
xe-10/0/0.200        93         4        129                        1024
xe-10/2/0.200        94         4        144                        1024
```

show l2 learning-interface

```
user@host> run show l2-learning interface
Routing Instance Name : default-switch
Logical Interface flags (DL -disable learning, AD -packet action drop,
                         LH - MAC limit hit, DN - Interface Down, MP - MAC Pinning enabled)
Logical             BD         MAC        STP          Logical
Interface           Name       Limit      State        Interface flags
ae0.0                          8192                      MP
```

Release Information

Command introduced in Junos OS Release 8.4.

Added sample output to indicate an EVPN MAC Pinned interface, introduced in Junos OS 16.2R1.

show l2-learning mac-move-buffer

IN THIS SECTION

- Syntax | 21797
- Description | 21797
- Options | 21797
- Required Privilege Level | 21797
- Output Fields | 21797
- Sample Output | 21797

Release Information | 21798

Syntax

```
show l2-learning mac-move-buffer
<brief | detail | extensive>
<active>
```

Description

(MX Series routers only) Display action as a result of configuring the MAC address move feature.

Options

none	Display action as a result of the MAC address move feature.
brief | detail | extensive	(Optional) Display the specified level of output.
active	(Optional) Display the set of interfaces blocked as a result of the MAC address move action.

Required Privilege Level

view

Output Fields

Display action as a result of the MAC address move feature.

Sample Output

show l2-learning mac-move-buffer active

```
user@host> show l2-learning mac-move-buffer active
MAC Address: 00:00:00:00:01:01, VLAN Id: 0
```

```
    Time Rec : 2012-06-25 06:23:41   Bridge Domain: bd10
    Prev IFL : ge-1/0/5.0            New IFL: ge-1/0/6.0
    IFBD     : ge-1/0/6.0:10         Blocked  : YES
```

show l2-learning mac-move-buffer extensive

```
user@host> show l2-learning mac-move-buffer extensive | display xml
<l2ald-mac-move-buffer>
<l2ald-mac-move-entry junos:style="extensive">
<l2ald-mac-address>aa:00:00:00:02:00</l2ald-mac-address>
<l2ald-learn-vlan-id>0<l2ald-learn-vlan-id>
<l2ald-mac-move-time-rec>11:27:57</l2ald-mac-move-time-rec>
<l2ald-mac-move-bridge-domain>bd</l2ald-mac-move-bridge-domain>
<l2ald-mac-move-from-ifl>ge-1/0/5.200</l2ald-mac-move-from-ifl>
<l2ald-mac-move-to-ifl>ge-1/0/6.200</l2ald-mac-move-to-ifl>
<l2ald-mac-move-to-ifbd>ge-1/0/6.200</l2ald-mac-move-to-ifbd>
<l2ald-mac-move-is-blocked>Yes</l2ald-mac-move-is-blocked>
</l2ald-mac-move-entry>
```

Release Information

Command introduced in Junos OS Release 9.4.

show l2-learning redundancy-groups

IN THIS SECTION

- Syntax | 21799
- Description | 21799
- Options | 21799
- Required Privilege Level | 21799
- Output Fields | 21799
- Sample Output | 21802
- Release Information | 21804

Syntax

```
show l2-learning redundancy-groups
logical-system [system-name | all]
<redundancy-group-id [0 to 4294967294]>
arp-statistics
nd-statistics
remote-macs
```

Description

(MX Series routers only) Display ARP statistics, Neighbor Discovery statistics, or remote MAC addresses for the Multi-Chassis Aggregated Ethernet (MC-AE) nodes for all or specified redundancy groups on a router or switch or logical systems on a router or switch. Note that the Redundancy Group ID is inherited by the bridging domain or VLAN from member AE interfaces.

Options

logical-system [*system-name* | all]	(Optional) Display information for a specified logical system or all systems.
redundancy-group-id	(Optional) The redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate the routing or switching devices contained in a redundancy group.
arp-statistics	(Optional) Count of ARP packets sent and received by the two MC-AE nodes.
nd-statistics	(Optional) Count of Neighbor Discovery packets sent and received by the two MC-AE nodes.
remote-macs	(Optional) List of remote MAC addresses in the “Installed” state, as learned from the remote MC-AE node.

Required Privilege Level

view

Output Fields

Output fields are listed in the approximate order in which they appear.

Table 736: show l2-learning redundancy-groups arp-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
MCLAG ARP Statistics Group ID	ARP statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.
ARP Rx Count From Line	Total number of ARPs received from the Line.
ARP Tx Count To Peer	Total number of ARPs sent to the peer.
ARP Rx Count From Peer	Total number of ARPs received from the peer.
ARP Drop Count received from line	Total number of ARPs sent by the peer that were received.
ARP Drop Count received from peer	Total number of ARPs sent by the peer that were dropped
Service-id	Service ID (configured at the routing instance level).

Table 737: show l2-learning redundancy-groups nd-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
MCLAG ND Statistics Group ID	Neighbor Discovery statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.

Table 737: show l2-learning redundancy-groups nd-statistics Output Fields *(Continued)*

Field Name	Field Description
ND Rx Count From Line	Total number of Neighbor Discovery packets received from the Line.
ND Tx Count To Peer	Total number of Neighbor Discovery packets sent to the peer.
ND Rx Count From Peer	Total number of Neighbor Discovery packets received from the peer.
ND Drop Count received from line	Total number of Neighbor Discovery packets sent by the peer that were received.
ND Drop Count received from peer	Total number of Neighbor Discovery packets sent by the peer that were dropped
Service-id	Service ID (configured at the routing instance level).

Table 738: show l2-learning redundancy-groups remote-macs Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
Peer-Addr	IP address of the remote peer.
VLAN	Virtual LAN identifier associated with the redundancy group.
MAC	Hardware media access control address associated with the redundancy group.
MCAE-ID	ID number of the MC-AE used by the redundancy group.

Table 738: show l2-learning redundancy-groups remote-macs Output Fields *(Continued)*

Field Name	Field Description
Flags	Connection state: local connect or Remote connect. If no flag is shown, the redundancy group may not be connected.
Status	Installation state: Installed or Not Installed.

Sample Output

show l2-learning redundancy-groups arp-statistics

```
user@host> show l2-learning redundancy-groups arp-statistics
Logical System : default
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ARP Statistics
Group ID                                : 1
ARP Rx Count From Line                  : 52
ARP Tx Count To Peer                    : 15
ARP Rx Count From Peer                  : 39
ARP Install Count                       : 34
ARP Drop Count received from line       : 37
ARP Drop Count received from peer       : 5
```

show l2-learning redundancy-groups nd-statistics

```
user@host> show l2-learning redundancy-groups nd-statistics
Logical System : default
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID                                : 1
ND Rx Count From Line                   : 52
ND Tx Count To Peer                     : 15
ND Rx Count From Peer                   : 39
```

```
ND Install Count                        : 34
ND Drop Count received from line        : 37
ND Drop Count received from peer        : 5
```

show l2-learning redundancy-groups remote-macs

```
user@host> show l2-learning redundancy-groups <redundancy-group-id> remote-macs
  Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

Service-id  Peer-Addr    VLAN        MAC            MCAE-ID  Subunit  Opcode   Flags   Status
10          10.1.1.2      100   64:87:88:6a:df:f0      1        0       1        0     Installed
```

show l2-learning redundancy-groups logical-system arp-statistics (for Logical Systems)

```
user@host> show l2-learning redundancy-groups logical-system LS1 arp-statistics

 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ARP Statistics
Group ID                                : 1
ARP Rx Count From Line                  : 52
ARP Tx Count To Peer                    : 15
ARP Rx Count From Peer                  : 39
ARP Install Count                       : 34
ARP Drop Count received from line       : 37
ARP Drop Count received from peer       : 5
```

show l2-learning redundancy-groups logical-system nd-statistics (for Logical Systems)

```
user@host> show l2-learning redundancy-groups logical-system LS1 nd-statistics

 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID                                : 1
ND Rx Count From Line                   : 52
ND Tx Count To Peer                     : 15
```

```
ND Rx Count From Peer                    : 39
ND Install Count                         : 34
ND Drop Count received from line         : 37
ND Drop Count received from peer         : 5
```

show l2-learning redundancy-groups group-id

```
user@host> show l2-learning redundancy-groups 1

Redundancy Group ID : 1       Flags : Local Connect,Remote Connect
```

show l2-learning redundancy-groups logical-system

```
user@host> show l2-learning redundancy-groups logical-system ls1

Redundancy Group ID : 2       Flags : Local Connect,Remote Connect
```

Release Information

Command introduced in Junos OS Release 13.2.

Support for logical systems added in Junos OS Release 14.1.

RELATED DOCUMENTATION

- Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up
- show interfaces mc-ae (MX Series) | **21090**
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers and QFX Series Switches
- Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up

show l2circuit connections

IN THIS SECTION

- Syntax | 21805
- Description | 21805
- Options | 21806
- Required Privilege Level | 21806
- Output Fields | 21806
- Sample Output | 21812
- Sample Output | 21812
- Sample Output | 21813
- Sample Output | 21814
- Release Information | 21815

Syntax

```
show l2circuit connections
<brief | extensive | summary>
<down | up | up-down>
<history>
<interface interface-name>
<logical-system (all | logical-system-name)>
<neighbor neighbor>
<status>
```

Description

Display status information about Layer 2 virtual circuits from the local provider edge (PE) router to its neighbors.

Options

none	Display standard information about Layer 2 virtual circuits on all interfaces for all neighbors.
brief | extensive | summary	(Optional) Display the specified level of output. Use history to display information about connection history. Use status to display information about the connection and interface status.
down | up | up-down	(Optional) Display nonoperational, operational, or both kinds of connections.
history	(Optional) Display information about connection history.
interface *interface-name*	(Optional) Show all Layer 2 virtual circuits on an interface.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS.
neighbor *neighbor*	(Optional) IP address of a specific neighbor.
status	(Optional) Display information about the connection and interface status.

Required Privilege Level

view

Output Fields

Table 739 on page 21806 lists the output fields for the `show l2circuit connections` command. Output fields are listed in the approximate order in which they appear.

Table 739: show l2circuit connections Output Fields

Field Name	Field Description
`Layer-2 Circuit Connections`	Displays the legends for connection and interface status.
`Neighbor`	Remote PE neighbor.

Table 739: show l2circuit connections Output Fields *(Continued)*

Field Name	Field Description
`Interface`	Logical PE-to-CE interface on which the virtual circuit is configured.
`Type`	VC type: `rmt` (remote) or `loc` (local).

Table 739: show l2circuit connections Output Fields *(Continued)*

Field Name	Field Description
`Legend for connection status (St)`	Status of the virtual circuit connection: • `EI`—The local virtual circuit interface is configured with an encapsulation that is not supported. • `MM`—The two routers do not agree on an MTU value, which causes an MTU mismatch. • `EM`—The encapsulation type received on this virtual circuit from the neighbor does not match the local virtual circuit interface encapsulation type. • `CM`—The two routers do not agree on a control word, which causes a control word mismatch. • `VM`—The remote and local VLAN IDs do not match across the Layer 2 circuit. • `OL`—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • `NC`—The interface is not configured as a CCC or TCC interface. • `BK`—The virtual circuit has switched to a backup connection. • `CB`—The remote PE router is advertising a different cell bundle from that configured on the local PE router. • `LD`—The connection to the local site is signaled down, because the CE-facing interface to the local site is down. • `RD`—The remote neighbor is down. It has signaled a problem using the pseudowire status code. • `NP`—The router detects that interface hardware is not present. The hardware may be offline, a PIC may not be of the desired type, or the interface may be configured in a different routing instance. • `Dn`—The virtual circuit is down. • `VC-Dn`—The virtual circuit is down because there is no tunnel LSP from the local PE router to the neighbor. • `UP`—The virtual circuit is operational.

Table 739: show l2circuit connections Output Fields *(Continued)*

Field Name	Field Description
	• CF—The router cannot find enough bandwidth to the remote router to satisfy the Layer 2 circuit bandwidth requirement. • IB—The bit rate is incompatible for Time Division Multiplexing (TDM). • TDM—TDM is not configured correctly. • ST—The virtual circuit has been switched to a standby connection. • SP—The virtual circuit connection is using a static pseudowire. • RS—The remote site is in a standby state. • XX—The virtual circuit is down for an unknown reason. This is a programming error.
Time last up	Date and time the virtual circuit was last operational.
# Up trans	Number of times the virtual circuit came up.
local-interface-name	Name of the local PE-to-CE interface.
Status	Status of the local interface.
Up	Interface is operational.
Dn	Interface is not operational.
NP	Not present. Interface does not exist.
DS	Disabled. Interface has been administratively disabled.
WE	Wrong encapsulation. The interface is not configured as CCC.
UN	Interface status is initialized.

Table 739: show l2circuit connections Output Fields *(Continued)*

Field Name	Field Description
`Encapsulation`	Encapsulation of the local interface.
`Flow Label Transmit`	Flow label transmit status.
`Flow Label Receive`	Flow label receive status.
`Remote PE`	Prefix of the remote PE router.
`Negotiated control-word`	Whether the use of the control word has been negotiated for this virtual circuit: `Yes (Null)` or `No`.
`Incoming label`	Label used by the remote side of the virtual circuit to send packets destined to the local side. This label is routed to the local virtual circuit interface.
`Outgoing label`	Label used by the local side of the virtual circuit to send packets to the remote side of the virtual circuit. Packets originated on the local virtual circuit interface are encapsulated with this label before being placed on the tunnel LSP to the neighbor for this virtual circuit. This label is allocated by the neighbor and is used in demultiplexing incoming packets destined for this virtual circuit.
`Negotiated PW status TLV`	Displays the pseudowire status type, length, and value (TLV). TLVs are a method of encoding variable-length or optional information. If the pseudowire status TLV is used, the corresponding local or neighbor PE router status code is also displayed.
`local PW status code`	If the pseudowire status TLV is used, displays the local PE router status code.
`Neighbor PW status code`	If the pseudowire status TLV is used, displays the neighbor PE router status code.
`Local interface`	Name of the local interface used for the Layer 2 circuit connection.
`Status`	Status of the local interface (`Up` or `Down`).

Table 739: show l2circuit connections Output Fields *(Continued)*

Field Name	Field Description
`Encapsulation`	Encapsulation configured for the local interface.
`APS-active`	Indicates that the interface belongs to the working circuit.
`APS-inactive`	Indicates that the interface belongs to the protect circuit.
`Connection protection`	Whether or not connection protection is configured for the Layer 2 circuit to the neighbor: `Yes` or `No`.
`VC bandwidth`	Bandwidth requirement of the Layer 2 circuit.
`Time`	Time at which the event occurred.
`Connection History`	Event types logged in history. • `loc intf up`—Local virtual circuit interface went up. • `loc intf down`—Local virtual circuit interface went down. • `In lbl Update`—Incoming label has been updated. • `Out lbl Update`—Outgoing label has been updated. • `PE route changed`—Route to PE router has been updated. • `PE route down`—Route to PE router is down. • `rmt side marked`—Remote side is marked. • `VC Dn`—Remote side indicated that its end of the virtual circuit is down (if the tunnel LSP from the remote side to the local side is down). • `status update timer`—Status update timer processing. It computes the state of the virtual circuit, and determines whether it should be advertised to or withdrawn from the remote side.

Sample Output

show l2circuit connections

```
user@host> show l2circuit  connections
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch               Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch           CF -- Call admission control failure
OL -- no outgoing label          IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- hot standby
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 10.255.245.51
    Interface                 Type  St     Time last up          # Up trans
    ge-2/0/2.600(vc 5)        rmt   Up     Dec  7 18:11:18 2009           1
      Remote PE: 10.255.245.51, Negotiated control-word: No
      Incoming label: 299856, Outgoing label: 299808
      Negotiated PW status TLV: No
      Local interface: ge-2/0/2.600, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
      Auto-sensed or Programmed by XYZ
```

Sample Output

show l2circuit connections interface

```
user@host> show l2circuit connections interface t1-2/0/0:1:1.0
Layer-2 Circuit Connections:
```

```
Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch               Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch           CF -- Call admission control failure
OL -- no outgoing label          IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- hot standby
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 10.1.1.1
    Interface                 Type  St     Time last up          # Up trans
    t1-2/0/0:1:1.0(vc 1)(SP)  rmt   Up     Apr 27 04:21:02 2011           1
      Remote PE: 10.1.1.1, Negotiated control-word: Yes (Non-null)
      Incoming label: 1010001, Outgoing label: 1000001
      Negotiated PW status TLV: No
      Local interface: t1-1/0/0:1:1.0, Status: Up, Encapsulation: SATOP-T1, APS-active
      Local interface: t1-2/0/0:1:1.0, Status: Up, Encapsulation: SATOP-T1, APS-inactive
```

Sample Output

show l2circuit connections extensive

```
user@host>show l2circuit connections extensive
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch               Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch           CF -- Call admission control failure
OL -- no outgoing label          IB -- TDM incompatible bitrate
```

```
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- hot standby
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 10.255.49.149
    Interface                 Type  St     Time last up          # Up trans
    ae0.0(vc 100)             rmt   Up     Aug 31 09:36:12 2009           1
      Remote PE: 10.255.49.149, Negotiated control-word: Yes (Null)
      Incoming label: 299824, Outgoing label: 299776
      Negotiated PW status TLV: Yes
      local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
      Local interface: ae0.0, Status: Up, Encapsulation: ETHERNET
      Connection protection: Yes
    Connection History:
        Aug 31 09:36:12 2009  status update timer
        Aug 31 09:36:12 2009  PE route changed
        Aug 31 09:36:12 2009  Out lbl Update                    299776
        Aug 31 09:36:12 2009  In lbl Update                     299824
        Aug 31 09:36:12 2009  loc intf up                        ae0.0
```

Sample Output

show l2circuit connections extensive (Pseudowire Redundancy with Hot Standby)

```
user@host>show l2circuit connections extensive
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch               Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch           CF -- Call admission control failure
OL -- no outgoing label          IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
```

```
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- Hot-standby Connection
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 192.0.2.101
    Interface                 Type  St     Time last up          # Up trans
    ge-1/3/2.600(vc 1)        rmt   Up     Jan 24 11:00:26 2013           1
      Remote PE: 192.0.2.101, Negotiated control-word: Yes (Null)
      Incoming label: 299776, Outgoing label: 299776
      Negotiated PW status TLV: Yes
      local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
      Local interface: ge-1/3/2.600, Status: Up, Encapsulation: VLAN
    Connection History:
        Jan 24 11:00:26 2013  status update timer
        Jan 24 11:00:26 2013  PE route changed
        Jan 24 11:00:26 2013  Out lbl Update                    299776
        Jan 24 11:00:26 2013  In lbl Update                     299776
        Jan 24 11:00:26 2013  loc intf up                ge-1/3/2.600
Neighbor: 192.0.2.102
    Interface                 Type  St     Time last up          # Up trans
    ge-1/3/2.600(vc 2)        rmt   HS     -----                       ----
      Remote PE: 192.0.2.102, Negotiated control-word: Yes (Null)
      Incoming label: 299792, Outgoing label: 299776
      Negotiated PW status TLV: Yes
      local PW status code: 0x00000020, Neighbor PW status code: 0x00000000
      Local interface: ge-1/3/2.600, Status: Up, Encapsulation: VLAN
```

Release Information

Command introduced before Junos OS Release 7.4.

Display enhancements in Junos OS Release 9.6.

Display enhancements in Junos OS Release 10.2.

Display enhancements in Junos OS Release 12.1.

Display enhancements in Junos OS Release 13.2.

show l2pt interface

IN THIS SECTION

- Syntax | 21816
- Description | 21816
- Options | 21817
- Required Privilege Level | 21817
- Output Fields | 21817
- Sample Output | 21818
- Release Information | 21819

Syntax

```
show l2pt interface
<brief | detail>
<interface-name>
```

Description

Display information about the interfaces and protocols for which the device will use Layer 2 (L2) protocol tunneling (L2PT) across destination encapsulation tunnels in an EVPN network. When the device receives protocol control frames on the configured interfaces, the device transparently transmits those frames across the encapsulation tunnels and doesn't trap those frames to the control plane.

Use this command when you configure L2PT over destination encapsulation tunnels with the "l2pt" on page 9537 statement at the `[edit protocols layer2-control]` hierarchy level. The output of this command includes the L2PT-enabled protocols and the destination encapsulation tunnel type.

> **NOTE**: Standard L2PT implements protocol tunneling using MAC rewrite operations, so instead of using the `l2pt` statement, you configure standard L2PT using the "mac-rewrite" on page 10259 statement at the `[edit protocols layer2-control]` hierarchy level. See *Layer 2 Protocol Tunneling* for details on the standard L2PT feature.

To see information about a standard L2PT configuration, use the "show mac-rewrite interface" on page 21993 command. That command displays the L2PT-enabled protocols and the destination MAC address the device uses for the MAC rewrite operation.

However, you can still use this `show l2pt interface` command to see information about interfaces where you enabled standard L2PT. In that case, instead of a destination encapsulation tunnel, the command shows the destination MAC address the device uses for the MAC rewrite operation.

Options

(brief | detail) (Optional) Display the specified level of output.

interface *interface-name* (Optional) Display information for L2PT over encapsulation tunnels for the specified interface. Without this option, the command displays L2PT information for all interfaces configured with L2PT over encapsulation tunnels.

Required Privilege Level

view

Output Fields

Table 740: show l2pt interface Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the access interface on which you have configured L2PT over encapsulation tunnels.	brief detail
`Destination`	Either of the following: • When you enable L2PT over encapsulation tunnels, the destination encapsulation tunnel type: • `vxlan-tunnel`—L2PT enabled over VXLAN tunnels on the device. • When you enable standard L2PT, the MAC address the device uses for the MAC rewrite operation	brief detail

Table 740: show l2pt interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Protocols	L2 protocol control packets that the device receives on the listed interface and transparently tunnels over the encapsulation tunnels on the device.	brief detail

Sample Output

show l2pt interface

```
user@device> show l2pt interface

Interface      Destination       Protocols
xe-0/0/1       VXLAN-TUNNEL      8021X 8023AH ELMI LACP LLDP MMRP MVRP
xe-0/1/1       VXLAN-TUNNEL      LLDP MVRP
```

show l2pt interface (L2PT over VXLAN and standard L2PT on different interfaces)

This sample output shows L2PT over VXLAN (Destination=VXLAN-TUNNEL) and standard L2PT (Destination=Protocol MAC rewrite address) enabled on different interfaces for same protocol.

```
user@device> show l2pt interface

Interface           Destination            Protocols
xe-0/0/2:3          VXLAN-TUNNEL            LLDP
xe-0/0/3:1          01:00:0C:CD:CD:D0       LLDP
```

show l2pt interface *name* detail

```
user@device> show l2pt interface xe-0/0/1 detail

Interface: xe-0/0/1
Destination: VXLAN-TUNNEL
Protocols: 8021X 8023AH ELMI LACP LLDP MMRP MVRP STP
```

Release Information

Command introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Layer 2 Protocol Tunneling over VXLAN Tunnels in EVPN-VXLAN Bridged Overlay Networks

l2pt (Destination Tunnels) | 9537

show l2vpn connections

IN THIS SECTION

- Syntax | 21819
- Description | 21820
- Options | 21820
- Required Privilege Level | 21820
- Output Fields | 21821
- Sample Output | 21825
- Release Information | 21829

Syntax

```
show l2vpn connections
<brief | extensive>
<down | up | up-down>
<history>
<instance instance>
<instance-history>
<local-site local-site>
<logical-system (all | logical-system-name)>
<remote-site remote-site>
```

```
<status>
<summary>
```

Description

Display Layer 2 virtual private network (VPN) connections.

Options

none	Display all Layer 2 VPN connections for all routing instances.
brief | extensive	(Optional) Display the specified level of output.
down | up | up-down	(Optional) Display nonoperational, operational, or both kinds of connections.
history	(Optional) Display information about connection history.
instance *instance*	(Optional) Display connections for the specified routing instance only.
instance-history	(Optional) Display information about connection history for a particular instance.
local-site *local-site*	(Optional) Display connections for the specified Layer 2 VPN local site name or ID only.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system.
remote-site *remote-site*	(Optional) Display connection for the specified Layer 2 VPN remote site ID only.
status	(Optional) Display information about the connection and interface status.
summary	(Optional) Display summary of all Layer 2 VPN connections information.

Required Privilege Level

view

Output Fields

Table 741 on page 21821 lists the output fields for the `show l2vpn connections` command. Output fields are listed in the approximate order in which they appear.

Table 741: show l2vpn connections Output Fields

Field Name	Field Description
`Instance`	Name of Layer 2 VPN instance.
`L2vpn-id`	For BGP autodiscovery, a globally unique Layer 2 VPN community identifier for the instance.
`Local-ID`	BGP `local-address` assigned to the local routing device.
`Local site`	Name of local site.
`Local source-attachment-id`	For FEC 129, the VPWS source attachment identifier. The point-to-point nature of VPWS requires that you specify the source access individual identifier (SAII) and the target access individual identifier (TAII). This SAII-TAII pair defines a unique pseudowire between two PE devices.
`Target-attachment-id`	For FEC 129, the VPWS target attachment identifier. If the configured target identifier matches a source identifier advertised by a remote PE device by way of a BGP auto-discovery message, the pseudowire between that source-target pair is signaled. If there is no match between an advertised source identifier and the configured target identifier, the pseudowire is not established.
`Interface name`	Name of interface.
`Remote Site ID`	Remote site ID.
`Label Offset`	Numbers within the label block that are skipped to find the next label base.
`Label-base`	Advertises the first label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router.

Table 741: show l2vpn connections Output Fields *(Continued)*

Field Name	Field Description
`Range`	Advertises the label block size.
`status-vector`	Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of `0` indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of `1` indicates either one or both are down.
`connection-site`	Name of the connection site.
`Type`	Type of connection: `loc` (local) or `rmt` (remote).
`St`	Status of the connection. (For a list of possible values, see the `Legend for connection status (St)` field.)
`Time last up`	Time that the connection was last in the `Up` condition.
`# Up trans`	Number of transitions from `Down` to `Up` condition.
`Local circuit`	Address and status of local circuit.
`Remote circuit`	Address and status of remote circuit.

Table 741: show l2vpn connections Output Fields *(Continued)*

Field Name	Field Description
St	Status of the Layer 2 VPN connection (corresponds with Legend for Connection Status): • EI—The local Layer 2 VPN interface is configured with an encapsulation that is not supported. • EM—The encapsulation type received on this Layer 2 VPN connection from the neighbor does not match the local Layer 2 VPN connection interface encapsulation type. • VC-Dn—The virtual circuit is currently down. • CM—The two routers do not agree on a control word, which causes a control word mismatch. • CN—The virtual circuit is not provisioned properly. • OR—The label associated with the virtual circuit is out of range. • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • LD—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established. • RD—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established. • LN—The local site has lost path selection to the remote site and therefore no pseduowires can be established from this local site. • RN—The remote site has lost path selection to a local site or other remote site and therefore no pseudowires are established to this remote site. • XX—The Layer 2 VPN connection is down for an unn reason. This is a programming error. • NC—The interface encapsulation is not configured as an appropriate CCC, TCC, or Layer 2 VPN encapsulation. • WE—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the Layer 2 VPN routing instance.

Table 741: show l2vpn connections Output Fields *(Continued)*

Field Name	Field Description
	• NP—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the desired type, or the interface might be configured in a different routing instance. • ->—Only the outbound connection is up. • <-—Only the inbound connection is up. • Up—The Layer 2 VPN connection is operational. • Dn—The Layer 2 VPN connection is down. • CF—The router cannot find enough bandwidth to the remote router to satisfy the Layer 2 VPN connection bandwidth requirement. • SC—The local site identifier matches the remote site identifier. No pseudowire can be established between these two sites. You should configure different values for the local and remote site identifiers. • LM—The local site identifier is not the minimum designated, meaning it is not the lowest. There is another local site with a lower site identifier. Pseudowires are not being established to this local site. and the associated local site identifier is not being used to distribute Layer 2 VPN label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. • RM—The remote site identifier is not the minimum designated, meaning it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot established a pseduowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic can continue to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. • IL—The incoming packets for the Layer 2 VPN connection have no MPLS label.
Remote PE	Address of the remote provider edge router.
Incoming label	Name of the incoming label.

Table 741: show l2vpn connections Output Fields *(Continued)*

Field Name	Field Description
Outgoing label	Name of the outgoing label.
Egress Protection	Whether the given PVC is protected by connection protection logic using egress protection for BGP signaled layer 2 services.
Flow Label Receive	Capability to pop the flow label in the receive direction to the remote provider edge (PE) router
Flow Label Transmit	Capability to push the flow label in the transmit direction to the provider edge (PE) router
Time	Date and time of Layer 2 VPN connection event.
Event	Type of event.
Interface/Lbl/PE	Interface, label, or PE router.

Sample Output

show l2vpn connections

```
user@host> show l2vpn connections
L2VPN Connections :
Instance : vpna
Edge protection: Not-Primary
Local site: 2 (ce-2)
offset: 1, range: 3, label-base: 32768
    connection-site           Type  St  Time last up          # Up trans
    3 (3)                     loc   Up  Jul 18 20:45:46 2001           1
      Local circuit: fe-0/0/0.1, Status: Up
      Remote circuit: fe-0/0/3.0, Status: Up
    1                         rmt   Up  Jul 18 21:47:25 2001           1
      Local circuit: fe-0/0/0.0, Status: Up
```

```
      Remote PE: 192.0.2.1
      Incoming label: 32768, Outgoing label: 32769
Local site: 3 (ce-3)
offset: 1, range: 2, label-base: 33792
    connection-site           Type  St  Time last up          # Up trans
    2 (ce-b)                  loc   Up  Jul 18 20:45:46 2001           1
      Local circuit: fe-0/0/0.1, Status: Up
      Remote circuit: fe-0/0/3.0, Status: Up
    1                         rmt   Up  Jul 18 21:47:25 2001           1
      Local circuit: fe-0/0/3.1, Status: Up
      Remote PE: 192.0.2.1
      Incoming label: 33792, Outgoing label: 32770
```

show l2vpn connections

```
user@host> show l2vpn connections

Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
```

```
Dn -- down

Instance: l2vpn-inst
Edge protection: Not-Primary
  Local site: pe2 (2)
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Jun 22 14:46:50 2015           1
      Remote PE: 10.255.255.1, Negotiated control-word: Yes (Null)
      Incoming label: 800002, Outgoing label: 800003
      Local interface: ge-0/0/1.300, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: Yes, Flow Label Receive: Yes
```

show l2vpn connections extensive

```
user@host> show l2vpn connections extensive
L2VPN Connections:
Instance: vpn-a
Edge protection: Not-Primary
Local site: ce-a (1)
  Interface name      Remote Site ID
  fe-0/0/0.0                       2
  Label Offset      Offset     Range
         32768           1         2
    connection-site           Type  St  Time last up          # Up trans
    2                         rmt   Up  Aug  3 00:08:14 2001           1
      Local circuit: fe-0/0/0.0, Status: Up
      Remote PE: 192.168.24.1
      Incoming label: 32769, Outgoing label: 32768
      Egress Protection: Yes
        Time                  Event            Interface/Lbl/PE
        Aug  3 00:08:14 2001  PE route up
        Aug  3 00:08:14 2001  Out lbl Update               32768
        Aug  3 00:08:14 2001  In lbl Update                32769
        Aug  3 00:08:14 2001  ckt0 up                 fe-0/0/0.0
```

show l2vpn connections extensive (VPWS)

```
user@host> show l2vpn connections
Layer-2 VPN connections:
```

```
Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: FEC129-VPWS
  L2vpn-id: 100:100
    Number of local interfaces: 1
    Number of local interfaces up: 1
    ge-2/0/5.0
  Local source-attachment-id: 1 (ONE)
    Target-attachment-id      Type  St     Time last up          # Up trans
    2                         rmt   Up     Nov 28 16:16:14 2012           1
      Remote PE: 198.51.100.2, Negotiated control-word: No
      Incoming label: 299792, Outgoing label: 299792
      Local interface: ge-2/0/5.0, Status: Up, Encapsulation: ETHERNET
    Connection History:
        Nov 28 16:16:14 2012  status update timer
        Nov 28 16:16:14 2012  PE route changed
        Nov 28 16:16:14 2012  Out lbl Update                    299792
        Nov 28 16:16:14 2012  In lbl Update                     299792
        Nov 28 16:16:14 2012  loc intf up                   ge-2/0/5.0
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance-history` option introduced in Junos OS Release 12.3R2.

show lacp interfaces

IN THIS SECTION

- Syntax | 21829
- Description | 21829
- Options | 21829
- Required Privilege Level | 21830
- Output Fields | 21830
- Sample Output | 21835
- Release Information | 21836

Syntax

```
show lacp interfaces
<interface-name>
extensive
```

Description

Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet, Fast Ethernet, or Gigabit Ethernet interface.

Options

none	Display LACP information for all interfaces.
interface-name	(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—**ae***number*
- Fast Ethernet—**fe-***fpc/pic/port*
- Gigabit Ethernet—**ge-***fpc/pic/port*
- 10 Gigabit Ethernet—**xe-***fpc/pic/port*

extensive Display LACP information for the interface in detail.

> **NOTE**: The `show lacp interfaces` command returns the following error message if your system is not configured in either active or passive LACP mode:
> "Warning: lacp subsystem not running – not needed by configuration"

Required Privilege Level

view

Output Fields

Table 742 on page 21830 lists the output fields for the `show lacp interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 742: show lacp interfaces Output Fields

Field Name	Field Description	Level of Outpu
LACP State	For a child interface configured with the force-up statement, LACP state displays FUP along with the interface name.	All Levels
Aggregated interface	Aggregated interface value.	All Levels

Table 742: show lacp interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Outpu

LACP State	LACP state information for each aggregated interface:
	• **Role**—Role played by the interface. It can be one of the following:
	  • **Actor**—Local device participating in LACP negotiation.
	  • **Partner**—Remote device participating in LACP negotiation.
	• **Exp**—Expired state. **Yes** indicates the actor or partner is in an expired state. **No** indicates the actor or partner is not in an expired state.
	• **Def**—Default. **Yes** indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. **No** indicates the operational partner information in use has been received in an LACP PDU.
	• **Dist**—Distribution of outgoing frames. **No** indicates distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is **Yes**.
	• **Col**—Collection of incoming frames. **Yes** indicates collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is **No**.
	• **Syn**—Synchronization. If the value is **Yes**, the link is considered synchronized. It has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is **No**, the link is not synchronized. It is currently not in the right aggregation.
	• **Aggr**—Ability of aggregation port to aggregate (**Yes**) or to operate only as an individual link (**No**).
	• `Timeout`—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (`Slow Timeout` or `Fast Timeout`). In a fast timeout, PDUs are sent every second and in a slow timeout, PDUs are sent every 30 seconds. LACP timeout occurs when 3 consecutive PDUs are missed. If LACP timeout is a fast timeout, the time taken when 3 consecutive PDUs are missed is 3 seconds (3x1 second). If LACP timeout is a slow timeout, the time taken is 90 seconds(3x30 seconds).
	• **Activity**—Actor or partner's port activity. **Passive** indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is **Active**. **Active** indicates the port's preference to participate in the protocol regardless of the partner's control value.
		All Levels

Table 742: show lacp interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Outpu
	• Core isolation state down (CDN)— LACP interface state. Down indicates the LACP interface is down because all the eBGP sessions for Ethernet VPN (EVPN) are down.	

Table 742: show lacp interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Outpu

LACP Protocol	LACP protocol information for each aggregated interface:
	• Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.
	• **Receive State**—One of the following values:
	  • **Current**—The state machine receives an LACP PDU and enters the **Current** state.
	  • **Defaulted**—If no LACP PDU is received before the timer for the **Current** state expires a second time, the state machine enters the **Defaulted** state.
	  • **Expired**—If no LACP PDU is received before the timer for the **Current** state expires once, the state machine enters the **Expired** state.
	  • **Initialize**—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the **Initialize** state.
	  • **LACP Disabled**—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to **LACP Disabled**. This state is similar to the **Defaulted** state, except that the port is forced to operate as an individual port.
	  • **Port Disabled**—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the **Port Disabled** state.
	• **Transmit State**—Transmit state of state machine. One of the following values:
	  • **Fast Periodic**—Periodic transmissions are enabled at a fast transmission rate.
	  • **No Periodic**—Periodic transmissions are disabled.
	  • **Periodic Timer**—Transitory state entered when the periodic timer expires.
	  • **Slow Periodic**—Periodic transmissions are enabled at a slow transmission rate.
	• **Mux State**—State of the multiplexer state machine for the aggregation port. The state is one of the following values:
	  • **Attached**—Multiplexer state machine initiates the process of attaching the port to the selected aggregator.
		All Levels

Table 742: show lacp interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Outpu
	• **Collecting**—**Yes** indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. **No** indicates the receive function of this link is not enabled. • **Collecting Distributing**—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution. • **Detached**—Process of detaching the port from the aggregator is in progress. • **Distributing**—**Yes** indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames may be passed down from the aggregator's distribution function for transmission. **No** indicates the transmit function of this link is not enabled. • **Waiting**—Multiplexer state machine is in a holding process, awaiting an outcome.	

LACP info	• `Role` can be one of the following:
	  • `Actor`—Local device participating in LACP negotiation.
	  • `Partner`—Remote device participating in LACP negotiation.
	• `System priority`—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer.
	• `System identifier`—Actor or partner system ID, encoded as a MAC address.
	• `Port priority`—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer.
	• `Port number`—Port number assigned to the port by the actor or partner, encoded as an unsigned integer.
	• `Port key`—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer.
		Extensive

Sample Output

show lacp interfaces (Aggregated Ethernet)

```
user@host> show lacp interfaces ae0 extensive
 LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-0/0/1                  Current   Fast periodic Collecting distributing
      ge-0/0/2                  Current   Fast periodic Collecting distributing
      ge-0/0/3                  Current   Fast periodic Collecting distributing
    LACP info:        Role     System             System       Port     Port    Port
                             priority         identifier   priority   number     key
      ge-0/0/1       Actor        127  00:05:86:4e:b6:c0        127        1       1
      ge-0/0/1     Partner        127  00:05:86:7e:d3:c0        127        1       1
      ge-0/0/2       Actor        127  00:05:86:4e:b6:c0        127        2       1
      ge-0/0/2     Partner        127  00:05:86:7e:d3:c0        127        2       1
      ge-0/0/3       Actor        127  00:05:86:4e:b6:c0        127        3       1
      ge-0/0/3     Partner        127  00:05:86:7e:d3:c0        127        3       1
```

show lacp interfaces (Gigabit Ethernet)

```
user@host> show lacp interfaces ge-0/3/0
Aggregated interface: ae0
    LACP State:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-0/3/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/3/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP Protocol:   Receive State    Transmit State           Mux State
      ge-0/3/0             Current     Fast periodic  Collecting distributing
```

show lacp interfaces (10 Gigabit Ethernet)

```
user@host> show lacp interfaces xe-1/0/2
Aggregated interface: ae0
    LACP State:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-1/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP Protocol:   Receive State    Transmit State          Mux State
      xe-1/0/2             Current     Fast periodic  Collecting distributing
```

Release Information

Command introduced in Junos OS Release 7.6.

`extensive` statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

- Configuring Aggregated Ethernet Links (CLI Procedure)
- *Configuring Link Aggregation*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*
- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*
- Junos OS Interfaces Fundamentals Configuration Guide

show lacp interfaces (View)

IN THIS SECTION

- Syntax | 21837
- Description | 21837
- Options | 21837
- Required Privilege Level | 21837

- Output Fields | 21838
- Sample Output | 21841
- Release Information | 21843

Syntax

```
show lacp interfaces interface-name
```

Description

Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet interface, redundant Ethernet interface, Gigabit Ethernet interface, or 10-Gigabit Ethernet interface. If you do not specify an interface name, LACP information for all interfaces is displayed.

Options

none Display LACP information for all interfaces.

interface-name (Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—`ae`*number*
- Redundant Ethernet—`reth`*number*
- Gigabit Ethernet—`ge-`*fpc*/*pic*/*port*
- 10-Gigabit Ethernet—`xe-`*fpc*/*pic*/*port*

> **NOTE**: The `show lacp interfaces` command returns the following error message if your system is not configured in either active or passive LACP mode:
> "Warning: lacp subsystem not running – not needed by configuration"

Required Privilege Level

view

Output Fields

Table 743 on page 21838 lists the output fields for the `show lacp interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 743: show lacp interfaces Output Fields

Field Name	Field Description
`Aggregated interface`	Aggregated interface value.

Table 743: show lacp interfaces Output Fields *(Continued)*

Field Name	Field Description

`LACP State`	LACP state information for each aggregated interface:
	• `Role`—Role played by the interface. It can be one of the following:
	  • `Actor`—Local device participating in LACP negotiation.
	  • `Partner`—Remote device participating in LACP negotiation.
	• `Exp`—Expired state. `Yes` indicates the actor or partner is in an expired state. `No` indicates the actor or partner is not in an expired state.
	• `Def`—Default. `Yes` indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. `No` indicates the operational partner information in use has been received in a link aggregation control protocol data unit (PDU).
	• `Dist`—Distribution of outgoing frames. `No` indicates distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is `Yes`.
	• `Col`—Collection of incoming frames. `Yes` indicates collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is `No`.
	• `Syn`—Synchronization. If the value is `Yes`, the link is considered synchronized. It has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is `No`, the link is not synchronized. It is currently not in the right aggregation.
	• `Aggr`—Ability of aggregation port to aggregate (`Yes`) or to operate only as an individual link (`No`).
	• `Timeout`—LACP timeout preference. Periodic transmissions of link aggregation control PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (`Long Timeout` or `Short Timeout`).
	• `Activity`—Actor or partner's port activity. `Passive` indicates the port's preference for not transmitting link aggregation control PDUs unless its partner's control value is `Active`. `Active` indicates the port's preference to participate in the protocol regardless of the partner's control value.

Table 743: show lacp interfaces Output Fields *(Continued)*

Field Name	Field Description

`LACP Protocol`	LACP protocol information for each aggregated interface:
	• Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.
	• `Receive State`—One of the following values:
	  • `Current`—The state machine receives a link aggregation control PDU and enters the `Current` state.
	  • `Defaulted`—If no link aggregation control PDU is received before the timer for the `Current` state expires a second time, the state machine enters the `Defaulted` state.
	  • `Expired`—If no link aggregation control PDU is received before the timer for the `Current` state expires once, the state machine enters the `Expired` state.
	  • `Initialize`—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the `Initialize` state.
	  • `LACP Disabled`—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to `LACP Disabled`. This state is similar to the `Defaulted` state, except that the port is forced to operate as an individual port.
	  • `Port Disabled`—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the `Port Disabled` state.
	• `Transmit State`—Transmit state of state machine. One of the following values:
	  • `Fast Periodic`—Periodic transmissions are enabled at a fast transmission rate.
	  • `No Periodic`—Periodic transmissions are disabled.
	  • `Periodic Timer`—Transitory state entered when the periodic timer expires.
	  • `Slow Periodic`—Periodic transmissions are enabled at a slow transmission rate.
	• `Mux State`—State of the multiplexer state machine for the aggregation port. The state is one of the following values:
	  • `Attached`—Multiplexer state machine initiates the process of attaching the port to the selected aggregator.

Table 743: show lacp interfaces Output Fields *(Continued)*

Field Name	Field Description
	• `Collecting Distributing`—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution. • `Detached`—Process of detaching the port from the aggregator is in progress. • `Waiting`—Multiplexer state machine is in a holding process, awaiting an outcome.

Sample Output

show lacp interfaces (Aggregated Ethernet)

```
user@host> show lacp interfaces ae0
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-2/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-2/0/0                  Current   Fast periodic Collecting distributing
      ge-2/0/1                  Current   Fast periodic Collecting distributing
      ge-2/2/0                  Current   Fast periodic Collecting distributing
      ge-2/2/1                  Current   Fast periodic Collecting distributing
```

show lacp interfaces (Redundant Ethernet)

```
user@host> show lacp interfaces reth0
Aggregated interface: reth0
```

```
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-11/0/0      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/0    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/1      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/1    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/2      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/2    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/3      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/3    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-11/0/0                 Current   Fast periodic Collecting distributing
      ge-11/0/1                 Current   Fast periodic Collecting distributing
      ge-11/0/2                 Current   Fast periodic Collecting distributing
      ge-11/0/3                 Current   Fast periodic Collecting distributing
      ge-3/0/0                  Current   Fast periodic Collecting distributing
      ge-3/0/1                  Current   Fast periodic Collecting distributing
      ge-3/0/2                  Current   Fast periodic Collecting distributing
      ge-3/0/3                  Current   Fast periodic Collecting distributing
{primary:node1}
```

show lacp interfaces (Gigabit Ethernet)

```
user@host> show lacp interfaces ge-0/3/0
Aggregated interface: ae0
    LACP State:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-0/3/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/3/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP Protocol:   Receive State    Transmit State           Mux State
      ge-0/3/0             Current     Fast periodic  Collecting distributing
```

Release Information

Command modified in Junos OS Release 10.2.

RELATED DOCUMENTATION

Verifying LACP on Redundant Ethernet Interfaces

show lacp statistics interfaces (View)

IN THIS SECTION

- Syntax | 21843
- Description | 21843
- Options | 21844
- Required Privilege Level | 21844
- Output Fields | 21844
- Sample Output | 21844
- Release Information | 21845

Syntax

```
show lacp statistics interfaces interface-name
```

Description

Display Link Aggregation Control Protocol (LACP) statistics about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, LACP statistics for all interfaces are displayed.

Options

interface-name (Optional) Name of an interface.

Required Privilege Level

view

Output Fields

Table 744 on page 21844 lists the output fields for the `show lacp statistics interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 744: show lacp statistics interfaces Output Fields

Field Name	Field Description
`Aggregated interface`	Aggregated interface value.
`LACP Statistics`	LACP statistics provide the following information: • `LACP Rx`—counter that increments for each received LACP packet. • `LACP Tx`—counter that increments for each transmitted LACP packet. • `Unknown Rx`—number of unrecognized packet errors logged. • `Illegal Rx`—number of invalid packets received. **NOTE**: Starting in Junos OS Evolved Release 18.3R1, the `clear interfaces statistics` command clears LACP statistics as well as the counters displayed in the `show lacp statistics interfaces` command.

Sample Output

show lacp statistics interfaces

```
user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
```

```
    ge-2/0/0              1352        2035          0           0
    ge-2/0/1              1352        2056          0           0
    ge-2/2/0              1352        2045          0           0
    ge-2/2/1              1352        2043          0           0
```

Release Information

Command modified in Release 10.2 of Junos OS.

Command introduced in Release 11.1 of Junos OS.

RELATED DOCUMENTATION

- Verifying LACP on Redundant Ethernet Interfaces
- Verifying the Status of a LAG Interface
- *Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets*
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*
- *Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch*

show lacp timeouts

IN THIS SECTION

- Syntax | 21846
- Description | 21846
- Options | 21846
- Output Fields | 21846
- Sample Output | 21847

Syntax

```
show lacp timeouts <interface-name>
```

Description

In a fast timeout, PDUs are sent every second and in a slow timeout, PDUs are sent every 30 seconds. LACP timeout occurs when 3 consecutive PDUs are missed. If LACP timeout is a fast timeout, the time taken when 3 consecutive PDUs are missed is 3 seconds (3x1 second). If LACP timeout is a slow timeout, the time taken is 90 seconds(3x30 seconds).

Options

interface-name (Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—**ae*number***
- Fast Ethernet—**fe-*fpc/pic/port***
- Gigabit Ethernet—**ge-*fpc/pic/port***
- 10 Gigabit Ethernet—**xe-*fpc/pic/port***

Output Fields

"Table 1" on page 21845 lists the output fields for the `show lacp timeouts` command. Output fields are listed in the approximate order in which they appear.

Table 745: show lacp timeouts Output Fields

Field Name	Field Description	Level of Outpu
Aggregated interface	Aggregated interface value.	All Levels
LACP Interfaces	It holds only lacp member interface name.	All Levels

Table 745: show lacp timeouts Output Fields *(Continued)*

Field Name	Field Description	Level of Outpu
Current State	It shows the current mux state of the member interface. **Mux State**—State of the multiplexer state machine for the aggregation port. The state is one of the following values: • **Attached**—Multiplexer state machine initiates the process of attaching the port to the selected aggregator. • **Collecting**—Yes indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. No indicates the receive function of this link is not enabled. • **Collecting Distributing**—Collecting and distributing states are merged to form a combined state (coupled control). Because independent control is impossible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling collection and distribution. • **Detached**—The process of detaching the port from the aggregator is in progress. • **Distributing**—Yes indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames may be passed down from the aggregator's distribution function for transmission. No indicates the transmit function of this link is not enabled. • **Waiting**—Multiplexer state machine is in a holding process, awaiting an outcome.	All Levels
Last Timeout	This value represents the timestamp of the most recent timeout generated on the member interface as a result of missing three consecutive PDUs on that port.	All Levels

Sample Output

show lacp timeouts (Gigabit Ethernet)

```
user@host> run show lacp timeouts
Aggregated interface: ae1
    LACP Interfaces:   Current state   Last timeout
      ge-0/0/0           Detached        2022-12-21 12:09:03 IST (06:59:56 ago)
```

show ldp database

IN THIS SECTION

- Syntax | 21848
- Description | 21848
- Options | 21848
- Required Privilege Level | 21849
- Output Fields | 21849
- Sample Output | 21853
- Release Information | 21861

Syntax

```
show ldp database
<brief | detail | extensive>
<inet | l2circuit>
<instance instance-name>
<logical-system (all | logical-system-name)>
<p2mp>
<session session>
<summary>
```

Description

Display entries in the LDP database.

Options

none	Display standard information about all entries in the LDP database for all routing instances.
brief | detail | extensive	(Optional) Display the specified level of output.

inet | l2circuit	(Optional) Display only IPv4 or Layer 2 circuit bindings.
instance *instance-name*	(Optional) Display routing instance information for the specified instance only.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS.
p2mp	(Optional) Display point-to-multipoint binding information. This option is only supported on Junos OS.
session *session*	(Optional) Display database for the specified session only. *session* is the destination address of the LDP session.
summary	(Optional)—Display summary output. This option displays the number of labels received and advertised for each LDP session.

Required Privilege Level

view

Output Fields

Table 746 on page 21849 describes the output fields for the `show ldp database` command. Output fields are listed in the approximate order in which they appear.

Table 746: show ldp database Output Fields

Field Name	Field Description	Level of Output
`Input label database`	Label received from the other router.	All levels
`Output label database`	Label advertised to the other router.	All levels
session-identifier	Session identifier, which includes the local and remote label space identifiers.	All levels
`Labels received`	Number of labels received from the other router.	All levels

Table 746: show ldp database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Labels advertised`	Number of labels advertised to the other router.	All levels.
`Label`	Label binding to a route prefix.	All levels

Table 746: show ldp database Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Prefix`	Route prefix.
	
	It can be one of the following values:
	
	• IP prefix.
	
	• Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
	
	• Layer 2 encapsulation type.
	
	Layer 2 encapsulation types are displayed in the format `L2CKT` *`control word status encapsulation-type vc-number`*, for example, `L2CKT CtflWord FRAME RELAY VC 2`
	
	• *`control-word-status`*—Displays whether the use of the control word has been negotiated for this virtual circuit:
	
	  • `NoCtrlWord`
	
	  • `CtrlWord`
	
	• *`encapsulation-type`*—Encapsulation type:
	
	  • `FRAME RELAY`
	
	  • `ATM AAL5`
	
	  • `ATM CELL`
	
	  • `VLAN`
	
	  • `ETHERNET`
	
	  • `CISCO_HDLC`
	
	  • `PPP`
	
	• `VC number`—Virtual circuit number. It can have any numeric value.
	
	• `(Stale)`—When you display the LDP database for the neighbor of a restarting router, the bindings leaned from the restarting neighbor
		All levels

Table 746: show ldp database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	are displayed as (Stale). Stale bindings are deleted if they are not refreshed within the recovery time.	
`MTU`	MTU of the Layer 2 circuit. MTU is displayed for all encapsulation types except ATM cell encapsulations.	`detail`
`VCCV Control Channel types`	Virtual Circuit Connection Verification (VCCV) control channel types. • `MPLS router alert label` • `MPLS PW label with TTL=1`	`extensive`
`VCCV Control Verification types`	The only valid VCCV control verification type is `LSP ping`.	`extensive`
`TDM payload size`	Size of the Time Division Multiplex (TDM) payload.	All levels
`TDM bitrate`	Bit rate for the TDM traffic.	All levels
`Requested VLAN ID`	(VLANs) VLAN identifier of the Layer 2 circuit.	`detail`
`Cell bundle size`	(ATM cell encapsulations) Maximum number of cells that the Layer 2 circuit can receive in a packet.	`detail`

Table 746: show ldp database Output Fields *(Continued)*

Field Name	Field Description	Level of Output

State	State of the label binding:
	• Active—Label binding has been installed and distributed appropriately. A label binding is almost always in this state.
	• New—New label that has not yet been distributed.
	  • MapRcv—Waiting to receive a label mapping message.
	  • MapSend—Waiting to send a label mapping message.
	  • RelRcv—Waiting to receive a label release message.
	  • RelRsnd—Waiting to receive a label release message before resending label mapping message.
	  • RelSend—Waiting to send a label release message.
	  • ReqSend—Waiting to send a label request message.
	  • W/dSend—Waiting to send a label withdrawal message.
		detail

Age	Time elapsed since the binding was created.	detail

Sample Output

show ldp database (primary)

```
user@host> show ldp database extensive
Input label database, 10.255.107.232:0--10.255.107.236:0
  Label     Prefix
 299840      10.255.107.232/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No
      3      10.255.107.236/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No
```

```
 299776      L2CKT CtrlWord VLAN VC 100
            MTU: 1500 Requested VLAN ID: 600 Flow Label T Bit: 1 Flow Label R Bit: 1
            State: Active
            Age: 9:35
            Entropy Label Capability: No
            VCCV Control Channel types:
                PWE3 control word
                MPLS router alert label
                MPLS PW label with TTL=1
            VCCV Control Verification types:
                LSP ping
                BFD with PW-ACH-encapsulation for Fault Detection
                BFD with IP/UDP-encapsulation for Fault Detection

Output label database, 10.255.107.232:0--10.255.107.236:0
  Label     Prefix
      3      10.255.107.232/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No
 299776      10.255.107.236/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No
```

show ldp database (standby)

```
user@host> show ldp database extensive

Input label database, 10.255.107.236:0--10.255.107.234:0
  Label     Prefix
 299808      10.255.107.230/32
            State: Active
            Age: 1d 2:46:36
            Standby binding state:
                Map messages: 1
                Release messages: 0
  Label     Prefix
 301136      10.255.107.232/32
            State: Active
            Age: 1d 2:46:36
```

```
            Standby binding state:
                Map messages: 1
                Release messages: 0
  Label     Prefix
      3      10.255.107.234/32
            State: Active
            Age: 1d 2:46:36
            Standby binding state:
                Map messages: 1
                Release messages: 0
  Label     Prefix
 302480      10.255.107.236/32
            State: Active
            Age: 1d 2:46:36
            Standby binding state:
                Map messages: 1
                Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.234:0
  Label     Prefix
 299904      10.255.107.230/32
            State: Active
            Age: 1d 2:46:36
 299936      10.255.107.232/32
            State: Active
            Age: 1d 2:46:36
 299872      10.255.107.234/32
            State: Active
            Age: 1d 2:46:36
      3      10.255.107.236/32
            State: Active
            Age: 1d 2:46:36
 299952      P2MP root-addr 10.255.107.230, lsp-id 16777217
            State: Active
            Age: 1d 2:46:36
```

show ldp database l2circuit detail

```
user@host> show ldp database l2circuit detail
 Input label database, 10.255.245.44:0--10.255.245.45:0
    Label     Prefix
```

```
   100176     L2CKT CtrlWord ATM CELL (VC Mode) VC 100
              Cell bundle size: 80
              State: Active
              Age: 9:48
   100256     L2CKT CtrlWord FRAME RELAY VC 101
              MTU: 4470
              State: Active
              Age: 9:48

 Output label database, 10.255.245.44:0--10.255.245.45:0
    Label     Prefix
   100048     L2CKT CtrlWord ATM CELL (VC Mode) VC 100
              Cell bundle size: 80
              State: Active
              Age: 9:48
   100112     L2CKT CtrlWord FRAME RELAY VC 101
              MTU: 4470
              State: Active
              Age: 9:48
```

show ldp database l2circuit extensive

```
user@host> show ldp database l2circuit extensive
Input label database, 10.255.245.198:0--10.255.245.194:0
  Label     Prefix
 299872     L2CKT CtrlWord PPP VC 100
            MTU: 4470
            VCCV Control Channel types:
                MPLS router alert label
                MPLS PW label with TTL=1
            VCCV Control Verification types:
                LSP ping
  Label     Prefix
            State: Active
            Age: 19:23:08
```

show ldp database p2mp (primary)

```
user@host> show ldp database p2mp extensive
```

This command is only supported on Junos OS.

```
Input label database, 10.255.107.232:0--10.255.107.236:0
  Label     Prefix
 569649      P2MP root-addr 10.255.107.232, lsp-id 16777217
            State: Active
            Age: 2d 6:41:46

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0
  Label     Prefix
 299888      P2MP root-addr 10.255.107.230, lsp-id 16777217
            State: Active
            Age: 2d 6:41:35
```

show ldp database p2mp (standby)

```
user@host> show ldp database p2mp extensive
```

This command is only supported on Junos OS.

```
Input label database, 10.255.107.236:0--10.255.107.232:0
  Label     Prefix
 299968      P2MP root-addr 10.255.107.230, lsp-id 16777217
            State: Active
            Age: 4d 22:21:57
          Standby binding state:
               Map messages: 1
               Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.232:0
  Label     Prefix
      3      P2MP root-addr 10.255.107.232, lsp-id 1
```

```
            State: Active
            Age: 4d 22:21:57
```

show ldp database session

```
user@host> show ldp database session 10.1.1.195
Input label database, 10.0.0.194:0--10.1.1.195:0
  Label     Prefix
 100002     10.255.245.197/32
 100003     10.255.245.196/32
 100004     10.0.0.194/32
      3     10.1.1.195/32
 100000     L2CKT NoCtrlWord FRAME RELAY VC 1
 100001     L2CKT CtrlWord FRAME RELAY VC 2
Output label database, 10.0.0.194:0--10.1.1.195:0
  Label     Prefix
 100003     10.255.245.197/32
 100004     10.1.1.195/32
 100002     10.255.245.196/32
      3     10.0.0.194/32
 100000     L2CKT CtrlWord FRAME RELAY VC 2
 100001     L2CKT NoCtrlWord FRAME RELAY VC 1
```

show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show ldp database
```

This command is only supported on Junos OS.

```
Input label database, 1.1.1.2:0--1.1.1.3:0
  Label     Prefix
 299808      1.1.1.2/32
      3      1.1.1.3/32
 299792      1.1.1.6/32
 299776      10.255.2.227/32
 299840      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
 299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
```

```
Output label database, 1.1.1.2:0--1.1.1.3:0
  Label     Prefix
      3      1.1.1.2/32
 299776      1.1.1.3/32
 299808      1.1.1.6/32
 299792      10.255.2.227/32

Input label database, 1.1.1.2:0--1.1.1.6:0
  Label     Prefix
 299856      1.1.1.2/32
 299792      1.1.1.3/32
      3      1.1.1.6/32
 299776      10.255.2.227/32
 299888      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
 299808      P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
 299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
 299840      P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
 299872      P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

Output label database, 1.1.1.2:0--1.1.1.6:0
  Label     Prefix
      3      1.1.1.2/32
 299776      1.1.1.3/32
 299808      1.1.1.6/32
 299792      10.255.2.227/32
```

show ldp database (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show ldp database
```

This command is only supported on Junos OS.

```
Input label database, 10.255.2.227:0--1.1.1.3:0
  Label     Prefix
 299808      1.1.1.2/32
      3      1.1.1.3/32
 299792      1.1.1.6/32
 299776      10.255.2.227/32
```

```
Output label database, 10.255.2.227:0--1.1.1.3:0
  Label     Prefix
 299856      1.1.1.2/32
 299776      1.1.1.3/32
 299792      1.1.1.6/32
      3      10.255.2.227/32

Input label database, 10.255.2.227:0--1.1.1.6:0
  Label     Prefix
 299856      1.1.1.2/32
 299792      1.1.1.3/32
      3      1.1.1.6/32
 299776      10.255.2.227/32

Output label database, 10.255.2.227:0--1.1.1.6:0
  Label     Prefix
 299856      1.1.1.2/32
 299776      1.1.1.3/32
 299792      1.1.1.6/32
      3      10.255.2.227/32
 299888      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
 299808      P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
 299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
 299840      P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
 299872      P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7
```

show ldp database summary

```
user@host> show ldp database summary
Session ID                      Labels received    Labels advertised
10.255.0.1:0--10.255.0.2:0                     4                    4
10.255.0.1:0--10.255.0.3:0                     4                    4
```

show ldp database (Multipoint LDP Recursive FEC)

```
user@host> show ldp database
```

```
Input label database, 192.168.1.2:0--192.168.1.3:0
Labels received: 4
  Label     Prefix
 299840      P2MP root-addr 192.168.1.2, lsp-id 16777217

Output label database, 192.168.1.2:0--192.168.1.3:0
Labels advertised: 4
  Label     Prefix
     16      P2MP root-addr 192.168.1.8, lsp-id 16777217 next-root:192.168.1.4
Input label database, 192.168.1.2:0--192.168.1.4:0
Labels received: 4

Output label database, 192.168.1.2:0--192.168.1.4:0
Labels advertised: 3
```

Release Information

Command introduced before Junos OS Release 7.4.

`summary` option introduced in Junos OS Release 14.2.

show ldp database-label-requests

IN THIS SECTION

- Syntax | 21862
- Description | 21862
- Options | 21862
- Required Privilege Level | 21863
- Output Fields | 21863

Sample Output | 21863
Release Information | 21864

Syntax

```
show ldp database-label-requests
<brief | detail | extensive>
<inet | l2circuit>
<instance instance-name>
<logical-system (all | logical-system-name)>
<p2mp>
<session session>
```

Description

Display entries in the LDP database label requests.

Options

none	Display standard information about all entries in the LDP database for all routing instances.
brief | detail | extensive	(Optional) Display the specified level of output.
inet | l2circuit	(Optional) Display only IPv4 or Layer 2 circuit bindings.
instance *instance-name*	(Optional) Display routing instance information for the specified instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS.
p2mp	(Optional) Display point-to-multipoint binding information. This option is only supported on Junos OS.
session *session*	(Optional) Display database for the specified session only. *session* is the destination address of the LDP session.

Required Privilege Level

view

Output Fields

Table 747 on page 21863 describes the output fields for the `show ldp database-label-requests` command. Output fields are listed in the approximate order in which they appear.

Table 747: show ldp database-label-requests Output Fields

Field Name	Field Description	Level of Output
`Input label database`	Label received from the other router.	All levels
`Output label database`	Label advertised to the other router.	All levels
`Prefix`	Route prefix. It can be one of the following values: • IP prefix.	All levels
`State`	State of the label binding: • `Active`—Label binding has been installed and distributed appropriately. A label binding is almost always in this state.	`detail`

Sample Output

show ldp database-label-requests

```
user@host> show ldp database-label-requests
Input label database, 198.51.100.0/24--203.0.113.0/24
  Prefix       State
203.0.113.0/24   Active
192.0.2.0/24   Active
```

```
Output label database, 198.51.100.0/24--203.0.113.0/24
  Prefix      State
198.51.100.0/24   Active
```

Release Information

Statement introduced in Junos OS Release 20.3R1.

show ldp fec-filters

IN THIS SECTION

- Syntax | 21864
- Description | 21864
- Options | 21865
- Required Privilege Level | 21865
- Output Fields | 21865
- Sample Output | 21865
- Release Information | 21865

Syntax

```
show ldp fec-filters
<fec>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display information about configured Label Distribution Protocol (LDP) forwarding equivalence class (FEC) filters.

Options

fec	(Optional) Display FEC filter information for the specified FEC.
instance *instance-name*	(Optional) Display FEC filter information for the specified instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 748 on page 21865 lists the output fields for the `show ldp fec-filters` command. Output fields are listed in the approximate order in which they appear.

Table 748: show ldp fec-filters Output Fields

Field Name	Field Description
Ingress	Names of the FEC filters on the ingress routers.
Transit	Names of the FEC filters on the transit routers.

Sample Output

show ldp fec-filters

```
user@host> show ldp fec-filters 10/8
10.22.1.2/32
     Ingress: f1-10.22.1.2/32 (index: 3)
     Transit: (null) (index: 0)
```

Release Information

Command introduced before Junos OS Release 7.4.

show ldp interface

IN THIS SECTION

- Syntax | 21866
- Description | 21866
- Options | 21866
- Required Privilege Level | 21867
- Output Fields | 21867
- Sample Output | 21868
- Release Information | 21868

Syntax

```
show ldp interface
<brief | detail | extensive>
<interface-name>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display the status of Label Distribution Protocol (LDP)-enabled interfaces.

Options

none	Display standard status information about all LDP-enabled interface for all routing instances.
interface-name	(Optional) Display information for the specified interface.
brief | detail | extensive	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display information for the specified routing instance.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS.

Required Privilege Level

view

Output Fields

Table 749 on page 21867 describes the output fields for the `show ldp interface` command. Output fields are listed in the approximate order in which they appear.

Table 749: show ldp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Interface name.	All levels
Label space ID	Label space identifier that the router is advertising on the interface.	All levels
Nbr count	Number of neighbors on the interface.	All levels
Next hello	How long until the next hello packet is sent on this interface, in seconds.	All levels
Hello interval	One-third of the negotiated hold time (in seconds). If the user-configured value for the hello interval is smaller than the computed value, the user-configured value is used.	**detail** **extensive**
Hold time	Configured hold time, in seconds.	**detail** **extensive**
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	**extensive**
Local hello interval	Locally configured hello interval.	**extensive**

Sample Output

show ldp interface extensive

```
user@host> show ldp interface extensive
Interface            Label space ID        Nbr count   Next hello
 fe-0/0/3.0           10.255.245.6:0           2           0
   Hello interval: 1, Hold time: 15, Transport address: 10.255.245.6
   Local hello interval: 2, Index: 69
```

Release Information

Command introduced before Junos OS Release 7.4.

show ldp neighbor

IN THIS SECTION

- Syntax | 21868
- Description | 21869
- Options | 21869
- Required Privilege Level | 21869
- Output Fields | 21869
- Sample Output | 21870
- Release Information | 21871

Syntax

```
show ldp neighbor
<brief | detail | extensive>
<auto-targeted>
<instance instance-name>
```

```
<logical-system (all | logical-system-name)>
<neighbor-address>
```

Description

Display Label Distribution Protocol (LDP) neighbor information.

Options

none	Display standard information about LDP neighbors for all routing instances.
brief | detail | extensive	(Optional) Display the specified level of output.
auto-targeted	(Optional) Display information about LDP neighbors that are automatically targeted using the loopback addresses.
instance *instance-name*	(Optional) Display information for the specified routing instance.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS.
neighbor-address	(Optional) Display information about the specified LDP neighbor.

Required Privilege Level

view

Output Fields

Table 750 on page 21869 describes the output fields for the `show ldp neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 750: show ldp neighbor Output Fields

Field Name	Field Description	Level of Output
Address	IP address of the neighbor.	All levels
Interface	Interface over which the neighbor was discovered.	All levels

Table 750: show ldp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Label space ID	Label space identifier advertised by the neighbor.	All levels
Hold time	Remaining hold time before the neighbor expires, in seconds.	All levels
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	**detail**
Configuration sequence	Counter that increments whenever the neighbor changes its configuration.	**detail**
Up for	Length of time the LDP neighbor has been in operation.	**detail extensive**
Reference count	Reference count for the LDP neighbor.	**extensive**
Hold time	Displays the neighbor's hold time. The hold time is the proposed hold times for the local and peer routers.	**extensive**
Proposed local/peer	Hold time value proposed by the local router and the peer router.	**extensive**

Sample Output

show ldp neighbor extensive

```
user@host> show ldp neighbor extensive
Address            Interface          Label space ID         Hold Time
 192.168.37.23      so-1/0/0.0         10.255.245.5:0           44
   Transport address: 10.255.245.5, Configuration sequence: 6
   Up for 00:03:37
   Reference count: 1
   Hold time: 45, Proposed local/peer: 15/45
```

show ldp neighbor auto-targeted extensive

```
user@host> show ldp neighbor auto-targeted extensive
Address            Interface          Label space ID         Hold time
10.255.107.236     lo0.0              10.255.107.236:0         41
  Transport address: 10.255.107.236, Configuration sequence: 14
  Up for 00:10:53
  Reference count: 2
  Hold time: 45, Proposed local/peer: 45/45
  Hello interval: 15
  Hello flags: targeted
  Neighbor types: Auto-targeted
```

Release Information

Command introduced before Junos OS Release 7.4.

`neighbor-address` option added in Junos OS Release 8.5.

RELATED DOCUMENTATION

clear ldp neighbor | 3680

show ldp overview

IN THIS SECTION

- Syntax | 21872
- Syntax (EX Series Switch and QFX Series) | 21872
- Description | 21872
- Options | 21872
- Required Privilege Level | 21872
- Output Fields | 21872
- Sample Output | 21876

Release Information | 21878

Syntax

```
show ldp overview
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and QFX Series)

```
show ldp overview
<instance instance-name>
```

Description

Display LDP overview information.

Options

none	Display standard overview information about LDP for all routing instances.
instance *instance-name*	(Optional) Display LDP overview information for the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Display LDP information from systems or a particular logical system on the devices. This option is only supported on Junos OS.

Required Privilege Level

view

Output Fields

Table 751 on page 21873 lists the output fields for the `show ldp overview` command. Output fields are listed in the approximate order in which they appear.

Table 751: show ldp overview Output Fields

Field Name	Field description	Level of Output
`Instance`	LDP routing instance.	All Levels
`Router ID`	Router ID of the routing device.	All Levels
`Message ID`	Unique identifier of message.	All Levels
`Configuration sequence`	Value of configuration sequence.	All Levels
`Deaggregate`	Status of control forwarding equivalence class (FEC) deaggregation on the router. By default it is `disabled` on the router.	All Levels
`Explicit null`	Advertise label 0 to the egress routing device of an LSP. Explicit null: `enabled` or `disabled`. **NOTE**: If you do not include the explicit-null statement in the configuration, label 3 (implicit null) is advertised.	All Levels
`IPv6 tunneling`	Internet Protocol version 6 tunneling: `enabled` or `disabled`.	All Levels
`Strict targeted hellos`	Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. Strict targeted hellos: `enabled` or `disabled`. **NOTE**: LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.	All Levels
`Loopback if added`	Loopback interface is added: `yes` or `no`.	All Levels
`Route preference`	Default preference value (also known as an administrative distance) assigned to each route that the routing table receives. LDP preference is: `9`	All Levels

Table 751: show ldp overview Output Fields *(Continued)*

Field Name	Field description	Level of Output
`Unicast transit LSP chaining`	Unicast transit LSP chaining: `enabled` or `disabled`.	All Levels
`P2MP transit LSP chaining`	P2MP transit LSP chaining: `enabled` or `disabled`.	All Levels
`Transit LSP statistics based on route statistics`	Transit LSP statistics based on route statistics: `enabled` or `disabled`.	All Levels
Longest match	Longest match for label mapping: `enabled` or `disabled`.	All Levels
`Capabilities enabled`	Enabled capabilities: `none`	All Levels
`Timers`	• `Keepalive interval`: Keepalive interval value. • `Keepalive timeout`: Time interval for which the neighbor LDP node waits before determining session failure. • `Link hello interval`: Specify how often the router sends Link Management Protocol (LMP) hello packets. • `Link hello hold time`: Time interval for which an LDP node waits for a hello message before declaring a neighbor is down. • `Targeted hello interval`: Specify how often LDP sends targeted hello messages. • `Targeted hello hold time`: Time interval for which a sending LSR maintains a record of targeted hello messages from the receiving LSR without receipt of another targeted hello message from that LSR. • `Label withdraw delay`: Time interval for withdrawing labels to reduce router workload during IGP convergence.	All Levels

Table 751: show ldp overview Output Fields *(Continued)*

Field Name	Field description	Level of Output
`Graceful restart`	Graceful restart attributes. • `Restart`– Graceful restart capability: `enabled` or `disabled`. • `Helper`– Standard graceful restart helper capability: `enabled` or `disabled`. • `Restart in process`– Graceful restart in process. • `Reconnect time`– Period of time that a restarting LSR (label switched router) designates to LDP neighbors to wait until the former reestablishes the session after restarting. • `Max neighbor reconnect time`– Maximum reconnect time. • `Recovery time`– Period of time that an LSR preserves its state across the restart. • `Max neighbor recovery time`– Maximum recovery time designated to LDP neighbors by a restarting LSR.	All Levels
`Traffic Engineering`	• `Bgp igp`– BGP and IGP destinations: `enabled` or `disabled`. When enabled, IGPs use MPLS paths for forwarding traffic. • `Both ribs`– BGP and IGP destinations with routes in both RIBs: `enabled` or `disabled`. • `Mpls forwarding`– MPLS routes used for forwarding: `enabled` or `disabled`.	All Levels
`IGP`	• `Tracking igp metric`–Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1). • `Sync session up delay`– Time interval to synchronize LDP session.	All Levels

Table 751: show ldp overview Output Fields *(Continued)*

Field Name	Field description	Level of Output
`Session protection`	• `Session protection`– Remote neighbor added to LDP configuration which enables protection for all sessions in the corresponding LDP instance: `enabled` or `disabled`. • `Session protection timeout`– Period of time until which the remote neighbor is connected to LSR in the absence of link neighbors.	All Levels
`Interface addresses advertising`	Advertises interface address.	All Levels
`Label allocation`	Label accounting information. • `Current number of LDP labels allocated`—Number of labels currently in use. • `Total number of LDP labels allocated`—Cumulative number of labels being allocated. • `Total number of LDP labels freed`—Cumulative number of labels being freed. • `Total number of LDP label allocation failures`—Cumulative number of failures for allocating a label. • `Current number of labels allocated by all protocols`—Number of labels currently being used by routing protocols.	All Levels
`Recursive fec`	Recursive opaque value per RFC 6512: `enabled` or `disabled`	All Levels

Sample Output

show ldp overview

```
user@host> show ldp overview
Instance: master
Reference count: 4
```

```
Router ID: 10.1.1.1
Inet LSR ID: 10.1.1.4
Inet6 LSR ID: 10.1.1.6
LDP inet: enabled
LDP inet6: enabled
Transport preference: Single-stack
Message id: 21
Configuration sequence: 1
Deaggregate: disabled
Explicit null: disabled
IPv6 tunneling: disabled
Strict targeted hellos: disabled
Loopback if added: yes
Route preference: 9
Unicast transit LSP chaining: disabled
P2MP transit LSP chaining: disabled
Transit LSP statistics based on route statistics: disabled
LDP route acknowledgement: enabled
BGP export: enabled
No TTL propagate: enabled
LDP mtu discovery: disabled
LDP SR Mapping Client: disabled
Capabilities enabled: none
Egress FEC capabilities enabled: entropy-label-capability
Downstream unsolicited Sessions:
Operational: 2
Retention: liberal
Control: ordered
Auto targeted sessions:
Auto targeted: disabled
Dynamic tunnel session count: 0
P2MP:
Recursive route: disabled
Recursive fec: enabled
No rsvp tunneling: disabled
Timers:
Keepalive interval: 10, Keepalive timeout: 30
Link hello interval: 5, Link hello hold time: 15
Targeted hello interval: 15, Targeted hello hold time: 45
Label withdraw delay: 60, Make before break timeout: 30
Make before break switchover delay: 3
Link protection timeout: 120
Graceful restart:
```

```
Restart: enabled, Helper: enabled, Restart in process: false
Reconnect time: 60000, Max neighbor reconnect time: 120000
Recovery time: 160000, Max neighbor recovery time: 240000
Traffic Engineering:
Bgp igp: disabled
Both ribs: disabled
Mpls forwarding: disabled
IGP:
Tracking igp metric: disabled
Sync session up delay: 10
Session protection:
Session protection: disabled
Session protection timeout: 0
Interface addresses advertising:
10.1.1.1
10.100.2.1
10.101.2.1
10.1.1.4
10.1.1.6
10.128.85.142
1000:1:2::1
1001:1:2::1
1111::1
2001:db8:abcd::10.128:85:142
fe80:1:2::1
fe80:1001:1002::1
LDP Job:
Read job time quantum: 1000, Write job time quantum: 1000
Read job loop quantum: 100, Write job loop quantum: 100
Backup inbound read job time quantum: 1000, Backup outbound read job time quantum: 1000
Backup inbound read job loop quantum: 100, Backup outbound read job loop quantum: 100
Label allocation:
Current number of LDP labels allocated: 4
Total number of LDP labels allocated: 7
Total number of LDP labels freed: 3
Total number of LDP label allocation failure: 0
Current number of labels allocated by all protocols: 4
```

Release Information

Command introduced in Junos OS Release 11.2.

show ldp p2mp fec

IN THIS SECTION

- Syntax | 21879
- Description | 21879
- Options | 21879
- Required Privilege Level | 21880
- Output Fields | 21880
- Sample Output | 21881
- Release Information | 21881

Syntax

```
show ldp p2mp fec
<brief | detail | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display Label Distribution Protocol (LDP) point-to-multipoint (P2MP) forwarding equivalence class (FEC) information.

Options

brief | detail | extensive	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display routing instance information for the specified instance only.
logical-system (all | *logical-system-name*)	(Optional) Display LDP point-to-multipoint tunnel table information of all logical systems or a particular logical system.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `name-of-command` command.

Table 752: name-of-command Output Fields

Field Name	Field Description	Level of Output
FEC	FEC for which LDP information is displayed. For P2MP LSPs, FEC appears as a combination of root address and the LSP ID (`root_addr:lsp_id`) or as a combination of root address multicast source address, and multicast group address (`root_addr:lsp_id/grp,src`) in case of LDP inband-signalling. For multipoint LDP P2MP LSPs, FEC appears as a combination of root address multicast source address, and multicast group address (`root_addr:lsp_id/grp,src`).	any
`Type`	Role of the router for this P2MP FEC, either `Ingress` (originating from this router), `Transit` (forwarded through this router), or `Egress`.	any
Label	Label in the output label database entry for the P2MP FEC; that is the label on which the incoming traffic for the P2MP FEC is received.	any
Reference count	Count of references for this FEC.	extensive

Sample Output

name-of-sample-output-command

```
user@host> show ldp p2mp fec extensive
```

```
LDP P2MP FECs:
 P2MP root-addr 192.168.1.1, lsp-id 16777217 next-root: 192.168.1.4
  Fec type: Ingress (Active)
  Label: 300288, Reference count:    1
P2MP root-addr 192.168.1.8, lsp-id 16777217
  Fec type: Transit (Active)
  Label: 300304, Reference count:    1
```

Release Information

Command introduced in Junos OS Release 15.1

RELATED DOCUMENTATION

No Link Title

show ldp p2mp path

IN THIS SECTION

- Syntax | 21882
- Description | 21882
- Options | 21882
- Required Privilege Level | 21882
- Output Fields | 21882
- Sample Output | 21883

Release Information | 21884

Syntax

```
show ldp p2mp path
<brief | detail | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display Label Distribution Protocol (LDP) point-to-multipoint (P2MP) label-switched paths (LSPs).

Options

brief | detail | extensive	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display routing instance information for the specified instance only.
logical-system (all | *logical-system-name*)	(Optional) Display LDP point-to-multipoint tunnel table information of all logical systems or a particular logical system.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `name-of-command` command.

Table 753: name-of-command Output Fields

Field Name	Field Description	Level of Output
`Output Session (label)`	Session ID and labels that this system has sent using LDP. These correspond to MPLS packets received.	any
`P2MP path type`	Type of P2MP traffic originating from a router, either `Egress` (leaving this router) or `Transit` (forwarded through this router).	any
`Input Session (label)`	Session ID and labels that this system has received using LDP. These correspond to MPLS packets transmitted.	any
`Attached FEC`	P2MP FECs corresponding to the paths (combination or input and output label).	any
`Address`	Memory address of the P2MP path.	extensive
`Reference Count`	Count of references for this FEC.	extensive
Transit route	P2MP traffic forwarded through this router.	extensive

Sample Output

name-of-sample-output-command

```
user@host> show ldp p2mp path extensive
```

```
P2MP path type: Transit/Egress
  Output Session (label): 192.168.1.3:0 (299984) (Primary)
  Input Session (label): 192.168.1.5:0 (300000)
  Attached FECs:  P2MP root-addr 192.168.1.2, lsp-id 16777217 (Active)
  Address: 0x76f94a0, Reference count: 2, Transit route
P2MP path type: Transit/Egress
```

```
Output Session (label): 192.168.1.8:0 (300304) (Primary)
Input Session (label): 192.168.1.6:0 (299953)
Attached FECs:  P2MP root-addr 192.168.1.8, lsp-id 16777217 (Active)
Address: 0x76faa60, Reference count: 2, Transit route
```

Release Information

Command introduced in Junos OS Release 15.1

RELATED DOCUMENTATION

No Link Title

show ldp p2mp tunnel

IN THIS SECTION

- Syntax | 21884
- Description | 21885
- Options | 21885
- Required Privilege Level | 21885
- Sample Output | 21885
- Release Information | 21885

Syntax

```
show ldp p2mp tunnel
<brief | detail | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display LDP point-to-multipoint tunnel table information.

Options

brief | detail | extensive	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display routing instance information for the specified instance only.
logical-system (all | *logical-system-name***)**	(Optional) Display LDP point-to-multipoint tunnel table information of all logical systems or a particular logical system.

Required Privilege Level

View

Sample Output

show ldp p2mp tunnel

```
user@host> show ldp p2mp tunnel extensive

Instance    Tunnel type          Tunnel name
0           Name              10.254.1.1:1:ldp-p2mp:mvpn:vpn-1
    P2MP root-addr 10.255.107.232, lsp-id 16777217
    Self id 805306372
    Reference count 2
```

Release Information

Command introduced in Junos OS Release 13.3.

show ldp path

IN THIS SECTION

- Syntax | 21886
- Description | 21886
- Options | 21886
- Required Privilege Level | 21887
- Output Fields | 21887
- Sample Output | 21888
- Release Information | 21888

Syntax

```
show ldp path
<brief | detail | extensive>
<destination>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display Label Distribution Protocol (LDP) label-switched paths (LSPs).

Options

none	Display standard information about all LDP LSPs for all routing instances.
brief | detail | extensive	(Optional) Display the specified level of output.
destination	(Optional) Restrict the output to entries that match the specified destination prefix.
instance *instance-name*	(Optional) Display information for the specified routing instance only.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS.

Required Privilege Level

view

Output Fields

Table 754 on page 21887 describes the output fields for the `show ldp path` command. Output fields are listed in the approximate order in which they appear.

Table 754: show ldp path Output Fields

Field Name	Field Description
Output Session (label)	Session ID and labels that this system has sent using LDP. These correspond to MPLS packets received.
Input Session (label)	Session ID and labels that this system has received using LDP. These correspond to MPLS packets transmitted.
route	MPLS route.
Attached route	Route corresponding to the LSP.
Ingress route	The router acts as the ingress for the LSP.
Reference count	Reference count for the LDP neighbor.
Transit route	Names of the forwarding equivalence class (FEC) filters on the transit routers.
Global label	MPLS label that is used globally.

Sample Output

show ldp path extensive

```
user@host> show ldp path extensive
Output Session (label)          Input Session (label)
10.255.14.220:0(3)              ( )
   Attached route:  10.255.14.221/32
   Reference count: 3, Global label: 3
10.255.14.220:0(100000)         10.255.14.220:0(3)
   Attached route:  10.255.14.220/32, Ingress route
   Reference count: 2, Transit route, Global label: 100000
10.255.14.220:0(100001)         10.255.14.220:0(100001)
   Attached route:  10.255.14.214/32, Ingress route
   Reference count: 2, Transit route, Global label: 100001
```

Release Information

Command introduced before Junos OS Release 7.4.

show ldp rib-groups

IN THIS SECTION

- Syntax | 21889
- Description | 21889
- Options | 21889
- Required Privilege Level | 21889
- Output Fields | 21889
- Sample Output | 21890
- Release Information | 21890

Syntax

```
show ldp rib-groups
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Displays the rib group names.

Options

instance *instance-name*	(Optional) Display routing instance information for the specified instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS.

Required Privilege Level

view

Output Fields

Table 755 on page 21889 describes the output fields for the `show ldp rib-groups` command. Output fields are listed in the approximate order in which they appear.

Table 755: show ldp rib-groups Output Fields

Field Name	Field Description
`Rib group name`	Displays the name of Rib group.

Sample Output

show ldp rib-groups

```
user@host> show ldp rib-groups
Rib group name: LDP
Rib group name: LDP
Rib group name: ldp-install master
Rib group name: ldp-install master
Rib group name: ldp-color master
Rib group name: ldp-color6 master
```

Release Information

Statement introduced in Junos OS Release 20.3R1.

show ldp route

IN THIS SECTION

- Syntax | 21890
- Description | 21891
- Options | 21891
- Required Privilege Level | 21891
- Output Fields | 21891
- Sample Output | 21892
- Release Information | 21901

Syntax

```
show ldp route
<brief | detail | extensive>
<destination>
```

```
<fec-and-route>
<fec-only>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display the entries in the Label Distribution Protocol (LDP) internal topology table. The internal topology table contains routes from inet.0 and inet.3 and is used when binding a label to a forwarding equivalence class (FEC).

Options

none	Display standard information about all entries in the LDP internal topology table for all routing instances.
brief | detail | extensive	(Optional) Display the specified level of output.
destination	(Optional) Restrict the output to entries that are longer than the specified destination prefix and prefix length.
fec-and-route	Display the show routes and the FECs.
fec-only	Display only LDP FECs.
instance *instance-name*	(Optional) Display entries for the specified routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS.

Required Privilege Level

view

Output Fields

Table 756 on page 21892 describes the output fields for the `show ldp route` command. Output fields are listed in the approximate order in which they appear.

Table 756: show ldp route Output Fields

Field Name	Field Description
`Destination`	Destination prefix.
`Next-hop intf/lsp/table`	Interface that is the next hop to the destination prefix.
`Next-hop address`	IP address of the next hop.
`Session ID`	LDP session ID.
`Route flags`	Information about the route. For example, the `Ingress TTL propagate` flag indicates that the time-to-live (TTL) value is being propagated with the route.
`Bound to outgoing label`	The route has been bound to LSPs with the label being distributed for that LSP.
`Topology entry`	The topology that the route is bound to.
`Ingress route status`	Status of the ingress route. For example, it could be `Active` or `Inactive`.
`Last modified`	The length of time since the ingress route status last changed.
`Last event(s)`	The last event that occurred.

Sample Output

show ldp route detail

```
user@host> show ldp route 10.255.8.5 detail
Destination         Next-hop intf/lsp                Next-hop address
 10.255.8.5/32      f1
   Session ID 10.255.170.84:0--10.255.170.92:0
                    fe-0/0/0.0                       192.168.100.2
   Session ID 10.255.170.84:0--10.255.8.5:0
```

```
                    so-0/2/1.0
   Session ID 10.255.170.84:0--10.255.8.5:0
                    so-0/2/2.0
   Session ID 10.255.170.84:0--10.255.8.3:0
   Bound to outgoing label 299776, Topology entry: 0x8c38a80
   BFD dest addr   BFD state LSP-ping Next-hop addr   Next-hop intf/lsp
   127.0.0.64      up        up       192.168.100.2   fe-0/0/0.0
   127.0.1.64      up        up                       so-0/2/1.0
   127.0.2.64      up        up                       so-0/2/2.0
   127.0.3.64      up        up                       f1
.....
```

show ldp route extensive

```
user@host> show ldp route extensive

Destination         Next-hop intf/lsp/table           Next-hop address
 10.0.0.0/30        ge-1/2/0.18                       10.0.0.17
   Session ID 192.168.0.6:0--192.168.0.5:0
   Route flags: None
Destination         Next-hop intf/lsp/table           Next-hop address
 10.0.0.4/30        ge-1/2/0.18                       10.0.0.17
   Session ID 192.168.0.6:0--192.168.0.5:0
   Route flags: None
Destination         Next-hop intf/lsp/table           Next-hop address
 10.0.0.8/30        ge-1/2/1.21                       10.0.0.22
   Session ID 192.168.0.6:0--192.168.0.4:0
   Route flags: None
Destination         Next-hop intf/lsp/table           Next-hop address
 10.0.0.12/30       ge-1/2/1.21                       10.0.0.22
   Session ID 192.168.0.6:0--192.168.0.4:0
   Route flags: None
Destination         Next-hop intf/lsp/table           Next-hop address
 10.0.0.16/30       ge-1/2/0.18
   Route flags: None
Destination         Next-hop intf/lsp/table           Next-hop address
 10.0.0.18/32
   Route flags: None
Destination         Next-hop intf/lsp/table           Next-hop address
 10.0.0.20/30       ge-1/2/1.21
   Route flags: None
```

```
Destination         Next-hop intf/lsp/table           Next-hop address
 10.0.0.21/32
   Route flags: None
Destination         Next-hop intf/lsp/table           Next-hop address
 192.168.0.1/32     ge-1/2/0.18                       10.0.0.17
   Session ID 192.168.0.6:0--192.168.0.5:0
   Route flags: None
Destination         Next-hop intf/lsp/table           Next-hop address
 192.168.0.2/32     ge-1/2/1.21                       10.0.0.22
   Session ID 192.168.0.6:0--192.168.0.4:0
                    ge-1/2/0.18                       10.0.0.17
   Session ID 192.168.0.6:0--192.168.0.5:0
   Route flags: None
Destination         Next-hop intf/lsp/table           Next-hop address
 192.168.0.3/32     ge-1/2/1.21                       10.0.0.22
   Session ID 192.168.0.6:0--192.168.0.4:0
   Route flags: None
Destination         Next-hop intf/lsp/table           Next-hop address
 192.168.0.4/32     ge-1/2/1.21                       10.0.0.22
   Session ID 192.168.0.6:0--192.168.0.4:0
   Bound to outgoing label 299808, Topology entry: 0x92a483c
   Ingress route status: Active, Last modified: 00:01:19 ago
   Route flags: Ingress TTL propagate, Transit TTL propagate
Destination         Next-hop intf/lsp/table           Next-hop address
 192.168.0.5/32     ge-1/2/0.18                       10.0.0.17
   Session ID 192.168.0.6:0--192.168.0.5:0
   Bound to outgoing label 299792, Topology entry: 0x92a47f8
   Ingress route status: Active, Last modified: 00:01:19 ago
   Route flags: Ingress TTL propagate, Transit TTL propagate
Destination         Next-hop intf/lsp/table           Next-hop address
 192.168.0.6/32     lo0.6
   Bound to outgoing label 3, Topology entry: 0x92a4a5c
   Ingress route status: Inactive
   Route type: Egress route
   Route flags: None
Destination         Next-hop intf/lsp/table           Next-hop address
    10.10.20.1/32       fe-1/0/0.0                        192.168.199.37
                        LSP LDP->10.255.107.230
```

show ldp route fec-and-route

```
user@host> show ldp route fec-and-route
Destination          Next-hop intf/lsp/table           Next-hop address
 10.4.0.0/16         fxp0.0                            10.92.31.254
 10.5.0.0/16         fxp0.0                            10.92.31.254
 10.6.128.0/17       fxp0.0                            10.92.31.254
 10.9.0.0/16         fxp0.0                            10.92.31.254
 10.10.0.0/16        fxp0.0                            10.92.31.254
 10.13.4.0/23        fxp0.0                            10.92.31.254
 10.13.10.0/23       fxp0.0                            10.92.31.254
 10.82.0.0/15        fxp0.0                            10.92.31.254
 10.84.0.0/16        fxp0.0                            10.92.31.254
 10.85.12.0/22       fxp0.0                            10.92.31.254
 10.92.0.0/16        fxp0.0                            10.92.31.254
 10.92.16.0/20       fxp0.0
 10.92.20.175/32
 10.94.0.0/16        fxp0.0                            10.92.31.254
 10.99.0.0/16        fxp0.0                            10.92.31.254
 10.102.0.0/16       fxp0.0                            10.92.31.254
 10.150.0.0/16       fxp0.0                            10.92.31.254
 10.155.0.0/16       fxp0.0                            10.92.31.254
 10.157.64.0/19      fxp0.0                            10.92.31.254
 10.160.0.0/16       fxp0.0                            10.92.31.254
 10.204.0.0/16       fxp0.0                            10.92.31.254
 10.205.0.0/16       fxp0.0                            10.92.31.254
 10.206.0.0/16       fxp0.0                            10.92.31.254
 10.207.0.0/16       fxp0.0                            10.92.31.254
 10.209.0.0/16       fxp0.0                            10.92.31.254
 10.212.0.0/16       fxp0.0                            10.92.31.254
 10.213.0.0/16       fxp0.0                            10.92.31.254
 10.214.0.0/16       fxp0.0                            10.92.31.254
 10.215.0.0/16       fxp0.0                            10.92.31.254
 10.216.0.0/16       fxp0.0                            10.92.31.254
 10.218.13.0/24      fxp0.0                            10.92.31.254
 10.218.14.0/24      fxp0.0                            10.92.31.254
 10.218.16.0/20      fxp0.0                            10.92.31.254
 10.218.32.0/20      fxp0.0                            10.92.31.254
 10.227.0.0/16       fxp0.0                            10.92.31.254
 10.255.111.0/24     ge-0/0/2.0                        11.11.11.2
 10.255.111.1/32     ge-0/0/2.0                        11.11.11.2
 10.255.111.2/32     ge-0/0/2.0                        11.11.11.2
```

```
 10.255.111.3/32    ge-0/0/2.0                        11.11.11.2
 10.255.111.4/32    ge-0/0/2.0                        11.11.11.2
 10.255.111.4/32    ge-0/0/2.0                        11.11.11.2
 10.255.112.1/32    lo0.0
 10.255.112.1/32    lo0.0
 10.255.112.2/32    ge-0/0/2.0                        11.11.11.2
 10.255.112.2/32    ge-0/0/2.0                        11.11.11.2
 11.11.11.0/24      ge-0/0/2.0
 11.11.11.1/32
 12.12.12.0/24      ge-0/0/2.0                        11.11.11.2
 15.15.15.0/24      ge-0/0/1.0
 15.15.15.1/32
 22.22.22.0/24      ge-0/0/0.0
 22.22.22.1/32
 23.23.23.0/24      ge-0/0/2.0                        11.11.11.2
 24.24.24.0/24      ge-0/0/2.0                        11.11.11.2
 25.25.25.0/24      ge-0/0/2.0                        11.11.11.2
 128.92.17.45/32    ge-0/0/2.0                        11.11.11.2
 128.92.20.175/32   lo0.0
 128.92.21.186/32   ge-0/0/2.0                        11.11.11.2
 128.92.25.135/32   ge-0/0/2.0                        11.11.11.2
 128.92.27.91/32    ge-0/0/2.0                        11.11.11.2
 128.92.28.70/32    ge-0/0/2.0                        11.11.11.2
 172.16.0.0/12      fxp0.0                            10.92.31.254
 192.168.0.0/16     fxp0.0                            10.92.31.254
 192.168.102.0/23   fxp0.0                            10.92.31.254
 207.17.136.0/24    fxp0.0                            10.92.31.254
 207.17.136.192/32  fxp0.0                            10.92.31.254
 207.17.137.0/24    fxp0.0                            10.92.31.254
 224.0.0.5/32
```

show ldp route fec-and-route

```
user@host> show ldp route fec-and-route
Destination         Next-hop intf/lsp/table           Next-hop address
 10.4.0.0/16        fxp0.0                            10.92.31.254
 10.5.0.0/16        fxp0.0                            10.92.31.254
 10.6.128.0/17      fxp0.0                            10.92.31.254
 10.9.0.0/16        fxp0.0                            10.92.31.254
 10.10.0.0/16       fxp0.0                            10.92.31.254
 10.13.4.0/23       fxp0.0                            10.92.31.254
```

```
  10.13.10.0/23      fxp0.0                            10.92.31.254
  10.82.0.0/15       fxp0.0                            10.92.31.254
  10.84.0.0/16       fxp0.0                            10.92.31.254
  10.85.12.0/22      fxp0.0                            10.92.31.254
  10.92.0.0/16       fxp0.0                            10.92.31.254
  10.92.16.0/20      fxp0.0
  10.92.20.175/32
  10.94.0.0/16       fxp0.0                            10.92.31.254
  10.99.0.0/16       fxp0.0                            10.92.31.254
  10.102.0.0/16      fxp0.0                            10.92.31.254
  10.150.0.0/16      fxp0.0                            10.92.31.254
  10.155.0.0/16      fxp0.0                            10.92.31.254
  10.157.64.0/19     fxp0.0                            10.92.31.254
  10.160.0.0/16      fxp0.0                            10.92.31.254
  10.204.0.0/16      fxp0.0                            10.92.31.254
  10.205.0.0/16      fxp0.0                            10.92.31.254
  10.206.0.0/16      fxp0.0                            10.92.31.254
  10.207.0.0/16      fxp0.0                            10.92.31.254
  10.209.0.0/16      fxp0.0                            10.92.31.254
  10.212.0.0/16      fxp0.0                            10.92.31.254
  10.213.0.0/16      fxp0.0                            10.92.31.254
  10.214.0.0/16      fxp0.0                            10.92.31.254
  10.215.0.0/16      fxp0.0                            10.92.31.254
  10.216.0.0/16      fxp0.0                            10.92.31.254
  10.218.13.0/24     fxp0.0                            10.92.31.254
  10.218.14.0/24     fxp0.0                            10.92.31.254
  10.218.16.0/20     fxp0.0                            10.92.31.254
  10.218.32.0/20     fxp0.0                            10.92.31.254
  10.227.0.0/16      fxp0.0                            10.92.31.254
  10.255.111.0/24    ge-0/0/2.0                        11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
  10.255.111.1/32    ge-0/0/2.0                        11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
    Bound to outgoing label 300192, Topology entry: 0xb5de1b0
    Ingress route status: Active, Last modified: 09:57:49 ago
    Last event(s): Rebind
    Route flags: Transit TTL propagate, Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
  10.255.111.2/32    ge-0/0/2.0                        11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
    Bound to outgoing label 300208, Topology entry: 0xb5de1f8
    Ingress route status: Active, Last modified: 09:57:49 ago
    Last event(s): Rebind
```

```
   Route flags: Transit TTL propagate, Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.111.3/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
   Bound to outgoing label 300224, Topology entry: 0xb5de240
   Ingress route status: Active, Last modified: 09:57:49 ago
   Last event(s): Rebind
   Route flags: Transit TTL propagate, Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.111.4/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
   Bound to outgoing label 300112, Topology entry: 0xb5de708
   Ingress route status: Active, Last modified: 10:10:56 ago
   Last event(s): Evaluate Update ingress route Update transit route
   Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.111.4/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
   Bound to outgoing label 300112, Topology entry: 0xb5de708
   Ingress route status: Active, Last modified: 10:10:56 ago
   Last event(s): Evaluate Update ingress route Update transit route
   Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.112.1/32    lo0.0
   Bound to outgoing label 3, Topology entry: 0xb5de120
   Ingress route status: Inactive
   Last event(s): Evaluate Update history
   Route type: Egress route
   Route flags: Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.112.1/32    lo0.0
   Bound to outgoing label 3, Topology entry: 0xb5de120
   Ingress route status: Inactive
   Last event(s): Evaluate Update history
   Route type: Egress route
   Route flags: Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.112.2/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
   Bound to outgoing label 300064, Topology entry: 0xb5de630
   Ingress route status: Active, Last modified: 10:11:04 ago
   Last event(s): Update ingress route
   Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match
```

```
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.112.2/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
   Bound to outgoing label 300064, Topology entry: 0xb5de630
   Ingress route status: Active, Last modified: 10:11:04 ago
   Last event(s): Update ingress route
   Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 11.11.11.0/24      ge-0/0/2.0
 11.11.11.1/32
 12.12.12.0/24      ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
 15.15.15.0/24      ge-0/0/1.0
 15.15.15.1/32
 22.22.22.0/24      ge-0/0/0.0
 22.22.22.1/32
 23.23.23.0/24      ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
 24.24.24.0/24      ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
 25.25.25.0/24      ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
 128.92.17.45/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
 128.92.20.175/32   lo0.0
 128.92.21.186/32   ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
 128.92.25.135/32   ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
 128.92.27.91/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
 128.92.28.70/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
 172.16.0.0/12      fxp0.0                            10.92.31.254
 192.168.0.0/16     fxp0.0                            10.92.31.254
 192.168.102.0/23   fxp0.0                            10.92.31.254
 207.17.136.0/24    fxp0.0                            10.92.31.254
 207.17.136.192/32  fxp0.0                            10.92.31.254
 207.17.137.0/24    fxp0.0                            10.92.31.254
 224.0.0.5/32
```

show ldp route fec-only

```
user@host> show ldp route fec-only
user@host_re0> show ldp route fec-only
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.111.1/32    ge-0/0/2.0                        11.11.11.2
 10.255.111.2/32    ge-0/0/2.0                        11.11.11.2
 10.255.111.3/32    ge-0/0/2.0                        11.11.11.2
 10.255.111.4/32    ge-0/0/2.0                        11.11.11.2
 10.255.112.1/32    lo0.0
 10.255.112.2/32    ge-0/0/2.0                        11.11.11.2
```

show ldp route fec-only detail

```
user@host> show ldp route fec-only detail
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.111.1/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
   Bound to outgoing label 300192, Topology entry: 0xb5de1b0
   Ingress route status: Active, Last modified: 09:55:10 ago
   Last event(s): Rebind
   Route flags: Transit TTL propagate, Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.111.2/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
   Bound to outgoing label 300208, Topology entry: 0xb5de1f8
   Ingress route status: Active, Last modified: 09:55:10 ago
   Last event(s): Rebind
   Route flags: Transit TTL propagate, Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.111.3/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
   Bound to outgoing label 300224, Topology entry: 0xb5de240
   Ingress route status: Active, Last modified: 09:55:10 ago
   Last event(s): Rebind
   Route flags: Transit TTL propagate, Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.111.4/32    ge-0/0/2.0                        11.11.11.2
   Session ID 10.255.112.1:0--10.255.112.2:0
   Bound to outgoing label 300112, Topology entry: 0xb5de708
```

```
    Ingress route status: Active, Last modified: 10:08:17 ago
    Last event(s): Evaluate Update ingress route Update transit route
    Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.112.1/32    lo0.0
    Bound to outgoing label 3, Topology entry: 0xb5de120
    Ingress route status: Inactive
    Last event(s): Evaluate Update history
    Route type: Egress route
    Route flags: Allow longest match
Destination         Next-hop intf/lsp/table           Next-hop address
 10.255.112.2/32    ge-0/0/2.0                        11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
    Bound to outgoing label 300064, Topology entry: 0xb5de630
    Ingress route status: Active, Last modified: 10:08:25 ago
    Last event(s): Update ingress route
    Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match
```

Release Information

Command introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 20.3R1.

show ldp session

IN THIS SECTION

- Syntax | 21902
- Description | 21902
- Options | 21902
- Required Privilege Level | 21902
- Output Fields | 21902
- Sample Output | 21909
- Release Information | 21911

Syntax

```
show ldp session
<brief | detail | extensive>
<auto-targeted>
<destination>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display information about Label Distribution Protocol (LDP) sessions.

Options

none	Display standard information about all LDP sessions for all routing instances.
brief | detail | extensive	(Optional) Display the specified level of output.
auto-targeted	(Optional) Display information about LDP sessions that are automatically targeted using loopback addresses.
destination	(Optional) Restrict LDP session display to the specified address.
instance *instance-name*	(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, information is displayed for the master instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS.

Required Privilege Level

view

Output Fields

Table 757 on page 21903 describes the output fields for the `show ldp session` command. Output fields are listed in the approximate order in which they appear.

Table 757: show ldp session Output Fields

Field Name	Field Description	Level of Output
`Address`	Transport address of the session.	any
`State`	State of the session: `Nonexistent`, `Connecting`, `Initialized`, `OpenRec`, `OpenSent`, `Operational`, or `Closing`. The states correspond to the state diagram specified in Internet Draft LDP Specification draft-ietf-mpls-rfc3036bis-01.txt.	any
`Connection`	TCP connection state: `Closed`, `Opening`, or `Open`.	any
`Hold time`	Time remaining until the session will be closed, in seconds.	any
`Session ID`	LDP identifiers of the peers of this session.	detail extensive
`Next keepalive`	Time until next keepalive is sent, in seconds.	`detail extensive`
`Active`	Whether the local router is playing the active role in the session and during session establishment.	`detail extensive`
`Passive`	Whether the local router is playing the passive role in the session and during session establishment.	`detail extensive`
`Maximum PDU`	Maximum protocol data unit (*PDU*) size (packet size) for the session.	`detail extensive`
`Hold time`	Time remaining until the session will be closed, in seconds. This value corresponds to the one configured using the `keepalive-timeout` statement configured at the `[edit protocols ldp]` hierarchy level.	`detail extensive`
`Neighbor count`	Number of neighbors that are contributing to the session.	`detail extensive`
`Neighbor types`	Category of LDP session: `discovered` or `auto-targeted`.	any

Table 757: show ldp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Keepalive interval`	Keepalive interval, in seconds.	`detail extensive`
`Connect retry interval`	TCP connection retry interval, in seconds.	`detail extensive`
`Local address`	Local transport address.	`detail extensive`
`Remote address`	Remote transport address.	`detail extensive`
`Up for`	Time that this session has been up.	`detail extensive`
`Last down`	Time since the session last went down.	`detail extensive`

Table 757: show ldp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Reason	Reason the session went down: • Aborted graceful restart • Authentication key was changed • Bad type length value (TLV) • Bad protocol data unit (PDU) packets • Command-line interface (CLI) command • Connect time expired • Connection error • Connection reset • Error during initialization • Hold time expired • No adjacency or all adjacencies down • Notification received • Received notification from peer • Unexpected End of File (EOF) • Unknown reason	detail extensive
Number of session flaps	Number of times the session changes from up to down.	detail extensive
Restarting	LDP is in the process of gracefully restarting.	detail extensive
Capabilities advertised	LDP capabilities advertised to a peer.	detail extensive

Table 757: show ldp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Capabilities received`	LDP capabilities received from a peer.	`detail extensive`
`Protection`	Information about the status of MPLS LDP session protection.	`detail extensive`
`restart complete in` *`nnn`* `msec`	Amount of time (in milliseconds) remaining until graceful restart is declared complete.	`detail extensive`
`Authentication type`	Shows the longest match MD5 authentication	`detail extensive`
`Local`	Information about graceful restart for the local end of an LDP session. Graceful restart and helper mode are independent. • `Restart`—Status of the grateful restart feature at the local end of the LDP session: `enabled` or `disabled`. • `Helper mode`—Status of the helper mode feature at the local end of the LDP session: `enabled` or `disabled`. When this feature is enabled, the local end of the LDP session can help the restarting router with its LDP restart procedures. • `Reconnect time`—Amount of time to wait from when a restart is initiated until the router can exchange LDP messages with its neighbors. The default is `60000 msec` and is not configurable. (`Reconnect timeout` refers to "FT Reconnect timeout" in draft-ietf-mpls-ldp-restart-06, *Internet Draft Graceful Restart Mechanism for LDP.)*	`detail extensive`

Table 757: show ldp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Remote`	Information about graceful restart at the remote end of an LDP session. Graceful restart and helper mode are independent. • `Restart`—Status of the grateful restart feature at the remote end of the LDP session: `enabled` or `disabled`. • `Helper mode`—Status of the helper mode feature at the remote end of the LDP session: `enabled` or `disabled`. When this feature is enabled, the remote end of the LDP session can help the restarting router with its LDP restart procedures. • `Reconnect time`—Amount of time in milliseconds from when a restart is initiated until the remote router can exchange LDP messages with its neighbors.	`detail extensive`
`Local maximum recovery time`	Amount of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds).	`detail extensive`
`Next-hop addresses received`	Next-hop addresses received on the session.	`detail extensive`
`Queue depth`	Number of messages that are queued for sending to the peers in the group.	`extensive`

Table 757: show ldp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Message type`	Type of message being sent: • • `Initialization`—Session initialization negotiation messages sent by an LSR to an LDP peer when the transport connection is established. • `Keeaplive`—Keepalive timer messages sent by an LSR to an LDP peer to keep the session active when there is no information or PDU exchanged between them. • `Notification`—Notification messages (such as state of the LDP session) or error information (such as bad PDU length) sent by an LSR to an LDP peer. • `Address`—Message sent by an LSR to an LDP peer to advertise interface addresses. • `Address withdraw`—Message sent by an LSR to an LDP peer to withdraw a previously advertised interface address. • `Label mapping`—Message sent by an LSR to an LDP peer to advertise label mapping for a forwarding equivalence class (FEC). • `Label request`—Message sent by an LSR to an LDP peer to request a label mapping for an FEC. • `Label withdraw`— Message sent by an LSR to an LDP peer to withdraw a previously advertised FEC-label mapping. • `Label release`—Message sent by an LSR to an LDP peer to notify the peer that a specific FEC-label mapping has been released. • `Label abort`—Message sent by an LSR to an LDP peer to abort a label request message. • `Total`—Messages sent and received during the lifetime of the session. • `Last 5 seconds`—Messages sent and received during the current session.	`extensive`

Sample Output

show ldp session brief

```
user@host> show ldp session brief
  Address           State        Connection     Hold time
10.255.72.160       Operational  Open             21
10.255.72.164       Operational  Open             20
10.255.72.172       Operational  Open             21
```

show ldp session detail

```
user@host> show ldp session detail
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 27
  Session ID: 192.168.0.2:0--192.168.0.3:0
  Next keepalive in 7 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 192.168.0.2, Remote address: 192.168.0.3
  Up for 00:00:02
  Capabilities advertised: none
  Capabilities received: none
  Protection: disabled
  Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
  Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
  Local maximum neighbor reconnect time: 120000 msec
  Local maximum neighbor recovery time: 240000 msec
  Local Label Advertisement mode: Downstream unsolicited
  Remote Label Advertisement mode: Downstream unsolicited
  Negotiated Label Advertisement mode: Downstream unsolicited
  Nonstop routing state: Not in sync
  Next-hop addresses received:
    10.0.0.5
    10.0.0.33
```

show ldp session extensive

```
user@host> show ldp session extensive
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 22
  Session ID: 192.168.0.2:0--192.168.0.3:0
  Next keepalive in 2 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 192.168.0.2, Remote address: 192.168.0.3
  Up for 00:05:37
  Capabilities advertised: none
  Capabilities received: none
  Protection: disabled
  Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
  Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
  Local maximum neighbor reconnect time: 120000 msec
  Local maximum neighbor recovery time: 240000 msec
  Local Label Advertisement mode: Downstream unsolicited
  Remote Label Advertisement mode: Downstream unsolicited
  Negotiated Label Advertisement mode: Downstream unsolicited
  Nonstop routing state: Not in sync
  Next-hop addresses received:
    10.0.0.5
    10.0.0.33
  Queue depth: 0
Message type               Total                     Last 5 seconds
                       Sent      Received          Sent      Received
Initialization            1             1             0             0
Keepalive                33            33             1             1
Notification              0             0             0             0
Address                   1             1             0             0
Address withdraw          0             0             0             0
Label mapping             7             5             0             0
Label request             0             0             0             0
Label withdraw            3             1             0             0
Label release             1             3             0             0
Label abort               0             0             0             0
```

show ldp session auto-targeted detail

```
user@host> show ldp session auto-generated detail
Address: 192.168.1.5, State: Operational, Connection: Open, Hold time: 25
    Session ID: 192.168.1.1:0--192.168.1.5:0
    Next keepalive in 5 seconds
    Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
    Neighbor types: discovered, Auto-targeted
                                ^^^^^^^^^^^^^
    Keepalive interval: 10, Connect retry interval: 1
    Local address: 192.168.1.1, Remote address: 192.168.1.5
    Up for 00:00:34
    Capabilities advertised: none
    Capabilities received: none
    Protection: disabled
    Local - Restart: disabled, Helper mode: enabled
    Remote - Restart: disabled, Helper mode: enabled
    Local maximum neighbor reconnect time: 120000 msec
    Local maximum neighbor recovery time: 240000 msec
    Local Label Advertisement mode: Downstream unsolicited
    Remote Label Advertisement mode: Downstream unsolicited
    Negotiated Label Advertisement mode: Downstream unsolicited
    Nonstop routing state: Not in sync
    Next-hop addresses received:
      192.168.1.2
      192.168.1.3
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear ldp session | 3682

show ldp statistics

IN THIS SECTION

- Syntax | 21912
- Description | 21912
- Options | 21912
- Required Privilege Level | 21912
- Output Fields | 21913
- Sample Output | 21916
- Release Information | 21917

Syntax

```
show ldp statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display Label Distribution Protocol (LDP) statistics.

Options

none	Display LDP statistics for all routing instances.
instance *instance-name*	(Optional) Display information for the specified routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS.

Required Privilege Level

view

Output Fields

Table 758 on page 21913 lists the output fields for the `show ldp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 758: show ldp statistics Output Fields

Field Name	Field Description
Total Sent, Received	Total number of each message type sent and received.
Last 5 seconds Sent, Received	Number of each message type sent and received in the last 5 seconds.

Table 758: show ldp statistics Output Fields *(Continued)*

Field Name	Field Description
Message type	LDP message types: • **Hello**—Messages that enable LDP nodes to discover one another and to detect the failure of a neighbor or of the link to the neighbor. • **Initialization**—Messages that indicate an LDP session has started. • **Keepalive**—Messages that ensure that the keepalive timeout is not exceeded. • **Notification**—Advisory information and signal error information. • **Address**—Messages with address information. • **Address withdrawal**—Messages regarding address withdrawal. • **Label mapping**—Messages with label mapping information. • **Label request**—Request for a label mapping from a neighboring router. • **Label withdrawal**—Withdrawal message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use. • **Label release**—Message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use. • **Label abort**—Messages about label interruptions. • **All UDP**—All hello messages sent by LSRs to the well-known UDP port, 646. • **All TCP**—All LDP session messages.

Table 758: show ldp statistics Output Fields *(Continued)*

Field Name	Field Description
Event type	LDP events and errors: • **Sessions opened**—Number of LDP sessions that have been opened. • **Sessions closed**—Number of LDP sessions that have been closed. • **Topology changes**—Number of changes to the known LDP topology. • **No interface**—Number of missing interface address messages. When a new LDP session is initialized and before sending label lapping or label request messages, the LSR advertises its interface addresses with one or more address messages. • **No session**—Number of missing session messages. Session messages are used to establish, maintain, and terminate sessions between LDP peers. • **No adjacency**—The exchange of hello adjacency messages results in the creation of an adjacency. The LDP identifier, together with the sender's LDP identifier in the PDU header, enables the receiver to match the initialization message with one of its hello adjacencies. If there is no matching hello adjacency, the LSR sends a session the initialization message is rejected. • **Unknown version**—The LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment. • **Malformed PDU**—An LDP PDU received on a TCP connection for an LDP session is malformed if the LDP identifier in the PDU header is unknown to the receiver, or if it is known but is not the LDP identifier associated by the receiver with the LDP peer for this LDP session. An LDP PDU is considered to be malformed if the LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment. An LDP PDU is considered malformed if the PDU length field is too small (less than 14) or too large (greater than maximum PDU length). • **Malformed message**—Malformed LDP messages that are part of the LDP discovery mechanism are handled by silently discarding them. An LDP message is malformed if the message type is unknown. If the message type is less than 0x8000 (high order bit = 0), it is an error signaled by the unknown message type status code.

Table 758: show ldp statistics Output Fields *(Continued)*

Field Name	Field Description
	An LDP message is considered to be malformed If the message length is too large, meaning that the message extends beyond the end of the containing LDP PDU. The LDP message is considered to be malformed if the message length is too small, meaning that it is smaller than the smallest possible value component. The LDP message is considered to be malformed if the message is missing one or more mandatory parameters. • **Unknown message type**—If the message type is less than 0x8000 (high order bit = 0) or greater than or equal to 0x8000 (high order bit = 1) it is considered to be an unknown message. • **Inappropriate message**—The message is not of the type that the receiver expects to receive. • **Malformed TLV**—The TLV lLength is too large or the receiver cannot decode the TLV value. This can indicate an issue in either the sending or receiving LSR. • **Bad TLV value**—The TLV Length is too large. • **Missing TLV**—The TLV is missing one or more mandatory parameters. • **PDU too large**—The PDF is greater than the maximum PDU length. Section "Initialization Message" in RFC 5036 describes how the maximum PDU length for a session is determined.
Total	Total number of each event or error.
Last 5 seconds	Number of each event or error in the last 5 seconds.

Sample Output

show ldp statistics

```
user@host> show ldp statistics
Message type              Total                     Last 5 seconds
                       Sent      Received          Sent      Received
Hello                   265           263             2             2
```

```
Initialization            2            2            0            0
Keepalive               112          111            1            0
Notification              0            0            0            0
Address                   2            2            0            0
Address withdraw          0            0            0            0
Label mapping             7            6            0            0
Label request             0            0            0            0
Label withdraw            2            0            0            0
Label release             0            2            0            0
Label abort               0            0            0            0
All UDP                 265          263            2            2
All TCP                 123          121            1            0

Event type                            Total       Last 5 seconds

Sessions opened                           2                 0
Sessions closed                           0                 0
Topology changes                         11                 0

No interface                              0                 0
No session                                0                 0
No adjacency                              0                 0
Unknown version                           0                 0
Malformed PDU                             0                 0
Malformed message                         0                 0
Unknown message type                      0                 0
Inappropriate message                     0                 0
Malformed TLV                             0                 0
Bad TLV value                             0                 0
Missing TLV                               0                 0
PDU too large                             0                 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear ldp statistics | 3683

show ldp traffic-statistics

IN THIS SECTION

- Syntax | 21918
- Description | 21918
- Options | 21918
- Additional Information | 21919
- Required Privilege Level | 21919
- Output Fields | 21919
- Sample Output | 21921
- Release Information | 21923

Syntax

```
show ldp traffic-statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
<p2mp>
```

Description

Display Label Distribution Protocol (LDP) traffic statistics.

NOTE: If nonstop active routing features is configured, `show ldp traffic-statistics` command is not supported on backup Routing Engines.

Options

none	Display LDP traffic statistics for all routing instances.
instance *instance-name*	(Optional) Display LDP traffic statistics for the specified routing instance only.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
p2mp	(Optional) Display only the data traffic statistics for a point-to-multipoint LSP.

Additional Information

To collect output from this command on a periodic basis, configure the `traffic-statistics` statement for the LDP protocol. For more information, see the *Junos MPLS Applications Configuration Guide*.

Required Privilege Level

view

Output Fields

Table 759 on page 21919 lists the output fields for the `show ldp traffic-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 759: show ldp traffic-statistics Output Fields

Field Name	Field Description
`Message type`	LDP message types.
`FEC`	Forwarding equivalence class (FEC) for which LDP traffic statistics are collected. For P2MP LSPs, FEC appears as a combination of root address and the LSP ID (`root_addr:lsp_id`). For multipoint LDP P2MP LSPs, FEC appears as a combination of root address multicast source address, and multicast group address (`root_addr:lsp_id/grp,src`).
`Type`	Type of traffic originating from a router, either `Ingress` (originating from this router) or `Transit` (forwarded through this router).
`Packets`	Number of packets passed by the FEC since its LSP came up.

Table 759: show ldp traffic-statistics Output Fields *(Continued)*

Field Name	Field Description
`Bytes`	Number of bytes of data passed by the FEC since its LSP came up.
`Shared`	Whether a label is shared by prefixes: `Yes` or `No`. A `Yes` value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
`Nexthop`	The next hop address for P2MP LSPs. (This is the downstream LDP Session ID.)
`Label`	For multipoint LDP with multicast-only fast reroute (MoFRR), the multipoint LDP node selects two separate upstream peers and sends two separate labels, one to each upstream peer. The same algorithm described in RFC 6388 is used to select the primary upstream path. The backup upstream path selection again uses the same algorithm but excludes the primary upstream LSR as a candidate. Two streams of MPLS traffic are sent to the egress node from the two different upstream peers. The MPLS traffic from only one of the upstream neighbors is selected as the primary path to accept the traffic, and the other becomes the backup path. The traffic on the backup path is dropped. When the primary upstream path fails, the traffic from the backup path is then accepted. The multipoint LDP node selects the two upstream paths based on the interior gateway protocol (IGP) root node next hop. Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
`Backup route`	For multipoint LDP with MoFRR, the route that is used if the primary route becomes unavailable.
Interface	Name of the point-to-multipoint interface for which statistics are displayed.
type	Type of LDP traffic on the interface.

Sample Output

show ldp traffic-statistics

```
user@host> show ldp traffic-statistics

 FEC                  Type                Packets              Bytes    Shared

  10.35.3.0/30        Transit                   0                  0    Yes

                      Ingress                   0                  0    No

 10.35.10.1/32       Transit                    0                  0    Yes

                      Ingress                   0                  0    No

  10.255.245.214/32   Transit                   0                  0    No

                      Ingress                  11                752    No

  192.168.37.36/30    Transit                    0                  0    Yes

                      Ingress                    0                  0    No



 FEC(root_addr:lsp_id)  Nexthop             Packets          Bytes     Shared

 10.255.72.160:16777217 192.168.8.81         152056          14597376   No

                         192.168.8.1         152056          14597376   No

                         192.168.8.65        152056          14597376   No
NET FEC Statistics:

  FEC                  Type                Packets              Bytes    Shared
  10.255.107.230/32   Transit               30858            2022345    No
                      Ingress                  20               5120    No
```

show ldp traffic-statistics p2mp (Ingress or transit router only, Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show ldp traffic-statistics p2mp
P2MP FEC Statistics:
FEC(root_addr:lsp_id/grp,src)     Nexthop              Packets             Bytes Shared
10.11.0.73:239.10.0.1,10.98.0.10  10.99.0.117           243408         121217184 No
                                  10.99.0.13            236286         117670428 No
10.11.0.73:239.10.0.2,10.98.0.10  10.99.0.117           248800         123902400 No
                                  10.99.0.13            240759         119897982 No
10.11.0.73:239.10.0.1,10.98.0.20  10.99.0.117           250286         124642428 No
                                  10.99.0.13            243741         121383018 No
10.11.0.73:239.10.0.2,10.98.0.20  10.99.0.117           252970         125979060 No
                                  10.99.0.13            245218         122118564 No
```

show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show ldp traffic-statistics p2mp

P2MP FEC Statistics:

FEC(root_addr:lsp_id/grp,src)     Nexthop              Packets             Bytes Shared
10.1.1.1:232.1.1.1,192.168.219.11, Label: 301568
                                  10.1.8.2                   0                 0 No
                                  10.1.4.2                   0                 0 No
10.1.1.1:232.1.1.1,192.168.219.11, Label: 301584, Backup route
                                  10.1.4.2                   0                 0 No
                                  10.1.8.2                   0                 0 No
10.1.1.1:232.1.1.2,192.168.219.11, Label: 301600
                                  10.1.8.2                   0                 0 No
                                  10.1.4.2                   0                 0 No
10.1.1.1:232.1.1.2,192.168.219.11, Label: 301616, Backup route
                                  10.1.4.2                   0                 0 No
                                  10.1.8.2                   0                 0 No
```

show ldp traffic-statistics interface p2mp (Multipoint LDP Interfaces)

```
user@host> show ldp traffic-statistics interface p2mp
```

```
                                Receive                Transmit
                Interface type  Packets       Bytes    Packets
lo0.0           p2mp                  0           0          0
ge-0/0/0.0      p2mp                 10         920          0
ge-0/0/1.0      p2mp                  0           0          0
ge-0/0/2.0      p2mp                  0           0          0
ge-0/0/3.0      p2mp                  0           0          5
ge-0/0/4.0      p2mp                  0           0         10
ge-0/0/5.0      p2mp                  0           0         15
```

Release Information

Command introduced before Junos OS Release 7.4.

`p2mp` option added in Junos OS Release 11.2.

RELATED DOCUMENTATION

clear ldp statistics | 3683

LDP Configuration

show link-management

IN THIS SECTION

- Syntax | 21924
- Description | 21924
- Options | 21924
- Required Privilege Level | 21924
- Output Fields | 21924
- Sample Output | 21927
- Release Information | 21927

Syntax

```
show link-management
```

Description

Display Multiprotocol Label Switching (MPLS) peer and traffic engineering link information.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 760 on page 21924 describes the output fields for the `show link-management` command. Output fields are listed in the approximate order in which they appear.

Table 760: show link-management Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
State	State of the peer: **Up** or **Down**.
Control address	Address to which a control channel is established.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is **1** through **4,294,967,296**.

Table 760: show link-management Output Fields *(Continued)*

Field Name	Field Description
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is **1** through **4,294,967,296**.
State	State of the control channel: **Up** or **Down**.
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is **1** through **4,294,967,295**.
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is **0** through **4,294,967,295.**
Flags	Code that provides information about the control channel. Currently supports only code value **R**, which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.
TE link name	Name of the traffic-engineered link.
State	State of the traffic-engineered link: **Up**, **Down**, or **Init**.
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET, Ethernet**, **Packet**, and **PDH**.

Table 760: show link-management Output Fields *(Continued)*

Field Name	Field Description
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are **PSC-1** and **Packet**.
Minimum bandwidth	Smallest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link (in bps).
Maximum bandwidth	Largest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bits per second (bps) and megabits per second (Mbps), of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidths of all interfaces that are members of the link and that are not yet allocated (in bps).
Name	Name of the interface.
State	State of the interface: **Up** or **Down**.
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: **Yes** or **No**.
LSP-name	LSP name.

Sample Output

show link-management

```
user@host> show link-management
Peer name: PEER-A, System identifier: 11973
 State: Up, Control address: 10.255.245.4
   CC local ID CC remote ID State       TxSeqNum  RcvSeqNum Flags
         24547        24547 Up              1027       1026
  TE links:
   pro4-ba

 TE link name: pro4-ba, State: Init
  Local identifier: 2662, Remote identifier: 0, Encoding: SDH/SONET, Switching:
  PSC-1,
  Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
  155.52Mbps,
  Available bandwidth: 155.52Mbps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    so-1/0/2    Up       21271          0     155.52Mbps   No
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show link-management peer | 21928
- show link-management routing | 21931
- show link-management statistics | 21935
- show link-management te-link | 21939

show link-management peer

IN THIS SECTION

- Syntax | 21928
- Description | 21928
- Options | 21928
- Required Privilege Level | 21928
- Output Fields | 21929
- Sample Output | 21930
- Release Information | 21930

Syntax

```
show link-management peer
<name peer-name>
```

Description

Display Multiprotocol Label Switching (MPLS) peer link information.

Options

none	Display all peer link information.
name *peer-name*	(Optional) Display information for the specified peer only.

Required Privilege Level

view

Output Fields

Table 761 on page 21929 describes the output fields for the `show link-management peer` command. Output fields are listed in the approximate order in which they appear.

Table 761: show link-management peer Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
State	State of the peer: **Up** or **Down**.
Control address	Address to which a control channel is established.
Hello interval	How often the routing device sends Link Management Protocol (LMP) hello packets.
Hello dead interval	How long LMP waits before declaring the control channel to be dead. This is an interval during which the routing device receives no LMP hello packets from the neighbor on a control that is active or up.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is **1** through **4,294,967,296**.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is **1** through **4,294,967,296**.
State	State of the control channel: **Up** or **Down**.
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is **1** through **4,294,967,295**.
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is **0** through **4,294,967,295.**

Table 761: show link-management peer Output Fields *(Continued)*

Field Name	Field Description
Flags	Code that provides information about the control channel. Currently supports only code value **R**, which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.

Sample Output

show link-management peer

```
user@host> show link-management peer
Peer name: sonet, System identifier: 41448
  State: Up, Control address: 70.70.70.70
  Hello interval: 10000, Hello dead interval: 30000
    CC local ID CC remote ID State       TxSeqNum  RcvSeqNum Flags
           3265            0 ConfSnd            1          0 R
   TE links:
    to-sonet
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show link-management | 21923

show link-management routing | 21931

show link-management statistics | 21935

show link-management te-link | 21939

show link-management routing

IN THIS SECTION

- Syntax | 21931
- Description | 21931
- Options | 21931
- Required Privilege Level | 21932
- Output Fields | 21932
- Sample Output | 21934
- Release Information | 21935

Syntax

```
show link-management routing
<peer <name name> | te-link <name name>>
<resource <name name>>
```

Description

Display Multiprotocol Label Switching (MPLS) peer or traffic engineering link information from the routing process.

Options

none	Display all peer and traffic-engineered link information.
peer <name *name*>	(Optional) Display information for all peers or for the specified peer only.
resource <name *name*>	(Optional) Display information for all resources or for the specified resource only.
te-link <name *name*>	(Optional) Display information for all traffic-engineered forwarding paths or for the specified path only.

Required Privilege Level

view

Output Fields

Table 762 on page 21932 describes the output fields for the `show link-management routing` command. Output fields are listed in the approximate order in which they appear.

Table 762: show link-management routing Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
State	State of the peer: **Up** or **Down**.
Control address	Address to which a control channel is established.
Control channel	Interface over which control packets are sent.
State	State of the control channel.
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: **Up** or **Down**.
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.

Table 762: show link-management routing Output Fields *(Continued)*

Field Name	Field Description
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET**, **Ethernet**, and **Packet**.
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Resource	Forwarding adjacency LSP information.
Type	Type of resource. The type is always a forwarding adjacency LSP.
State	State of the LSP: **Up** or **Down**.
System Identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
Total bandwidth	Bandwidth resource, in bps or Mbps, on the TE-link learned from the routing process.
Traffic parameters	• **Encoding**—Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET**, **Ethernet**, and **Packet**. • **Switching**—Type of switching that can be performed on the traffic-engineered link: **PSC-1** and **Packet**. • **Granularity**—Layer 2 data for switching Layer 2 LSPs for this resource. Not supported. This value is always **unknown**.

Sample Output

show link-management routing

```
user@host> show link-management routing
Peer name: __rpd:fe-0/1/0.0, System identifier: 2147483649
 State: Up, Control address: (null)
   Control-channel                  State
   fe-0/1/0.0                       Active

Peer name: __rpd:fe-0/1/2.0, System identifier: 2147483650
 State: Up, Control address: (null)
   Control-channel                  State
   fe-0/1/2.0                       Active

Peer name: __rpd:so-0/2/0.0, System identifier: 2147483651
 State: Down, Control address: (null)
   Control-channel                  State
   so-0/2/0.0

Peer name: __rpd:so-0/2/1.0, System identifier: 2147483652
 State: Down, Control address: (null)
   Control-channel                  State
   so-0/2/1.0

...

 TE link name: __rpd:fe-0/1/0.0, State: Up
  Local identifier: 2147483649, Remote identifier: 0,
  Local address: 192.168.37.66, Remote address: 192.168.37.66,
  Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
  Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

 TE link name: __rpd:fe-0/1/2.0, State: Up
  Local identifier: 2147483650, Remote identifier: 0,
  Local address: 192.168.37.73, Remote address: 192.168.37.73,
  Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
  Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

 TE link name: __rpd:so-0/2/0.0, State: Down
  Local identifier: 2147483651, Remote identifier: 0,
  Local address: 192.168.37.82, Remote address: 192.168.37.95,
```

```
  Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 155.52Mbps,
  Total bandwidth: 155.52Mbps, Available bandwidth: 155.52Mbps

...

Resource: falsp-bd, Type: LSP, State: Dn  System identifier: 2147483652,
Total bandwidth: 0bps,  Traffic parameters: Encoding: Packet, Switching: Packet, Granularity:
Unknown

Resource: falsp-be, Type: LSP, State: Up  System identifier: 2147483654,
Total bandwidth: bw[1]=10Mbps,  Traffic parameters: Encoding: Packet,
Switching: Packet, Granularity: Unknown
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show link-management | 21923
- show link-management peer | 21928
- show link-management statistics | 21935
- show link-management te-link | 21939

show link-management statistics

IN THIS SECTION

- Syntax | 21936
- Description | 21936
- Options | 21936
- Required Privilege Level | 21936
- Output Fields | 21936
- Sample Output | 21938

Release Information | 21938

Syntax

```
show link-management statistics
<peer <name name>>
```

Description

Display statistical information for Link Management Protocol (LMP) packets.

Options

none	Display information for all peers.
peer <name *name*>	(Optional) Display information for all peers or for the specified peer only.

Required Privilege Level

view

Output Fields

Table 763 on page 21936 describes the output fields for the `show link-management statistics` command. Output fields are listed in the approximate order in which they appear.

Table 763: show link-management statistics Output Fields

Field Name	Field Description
Received packets	Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.

Table 763: show link-management statistics Output Fields *(Continued)*

Field Name	Field Description
Received bad packets	Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Small packets	Number of packets that are too small.
Wrong protocol version	Number of packets specifying the wrong LMP version.
Messages for unknown peer	Number of packets destined for an unknown peer.
Messages for bad state	Number of packets indicating a state that does not match the recipient.
Stale acknowledgments	Number of **configAck** and **LinkSummaryAck** packets received that have a stale message ID.
Stale negative acknowledgments	Number of **configNack** and **LinkSummaryNack** packets received that have a stale message ID.
Sent packets	Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Retransmitted packets	Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Dropped packets	Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.

Sample Output

show link-management statistics

```
user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
    Received packets
        Config: 1
        Hello: 2572
    Small packets: 0
    Wrong protocol version: 0
    Messages for unknown peer: 0
    Messages for bad state: 0
    Stale acknowledgments: 0
    Stale negative acknowledgments: 0
    Sent packets
        Config: 2
        ConfigAck: 1
        Hello: 2572
    Retransmitted packets
        Config: 1
```

Release Information

Command introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

show link-management | 21923

show link-management peer | 21928

show link-management routing | 21931

show link-management te-link | 21939

show link-management te-link

IN THIS SECTION

- Syntax | 21939
- Description | 21939
- Options | 21939
- Required Privilege Level | 21939
- Output Fields | 21940
- Sample Output | 21941
- Release Information | 21942

Syntax

```
show link-management te-link
<brief | detail>
<name name>
```

Description

Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.

Options

none	Display information for all traffic-engineered links.
brief | detail	(Optional) Display the specified level of output.
name *name*	(Optional) Display information for the specified traffic-engineered link only.

Required Privilege Level

view

Output Fields

Table 764 on page 21940 describes the output fields for the `show link-management te-link` command. Output fields are listed in the approximate order in which they appear.

Table 764: show link-management te-link Output Fields

Field Name	Field Description
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: **Up** or **Down.**
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET**, **Ethernet**, **Packet**, and **PDH**.
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are **PSC-1** and **Packet**.
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).

Table 764: show link-management te-link Output Fields *(Continued)*

Field Name	Field Description
Available Bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Name	Name of the interface.
State	State of the interface: **Up** or **Down**.
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: **Yes** or **No**.
LSP-name	LSP name.

Sample Output

show link-management te-link

```
user@host> show link-management te-link
TE link name: FA-bd, State: Up
  Local identifier: 4144, Remote identifier: 0, Local address: 10.2.2.1,
  Remote address: 10.2.2.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,   Available
bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    falsp-bd    Dn       43077          0           0bps No

 TE link name: FA-be, State: Up
  Local identifier: 4145, Remote identifier: 0, Local address: 10.1.1.1,
   Remote address: 10.1.1.2, Encoding: Ethernet, Switching: Packet,
```

```
  Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,   Available
bandwidth: 8Mbps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    falsp-be    Up       43076          0         10Mbps  Yes  e2elsp-bf
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show link-management | 21923
- show link-management peer | 21928
- show link-management routing | 21931
- show link-management statistics | 21935

show lldp

IN THIS SECTION

- Syntax | 21943
- Description | 21943
- Options | 21943
- Required Privilege Level | 21943
- Output Fields | 21943
- Sample Output | 21949
- Release Information | 21952

Syntax

```
show lldp
<detail>
```

Description

`Port ID TLV`Display information about Link Layer Discovery Protocol (LLDP) and Link Level Discovery Protocol–Media Endpoint Discovery (LLDP-MED) configuration and capabilities on the switch. LLDP and LLDP-MED are used to learn about and to distribute device information on network links.

Options

none Display LLDP information for all interfaces.

detail (Optional) Display detailed LLDP information for all interfaces.

Required Privilege Level

view

Output Fields

Table 765 on page 21943 lists the output fields for the `show lldp` command. Output fields are listed in the approximate order in which they appear.

Table 765: show lldp Output Fields

Field Name	Field Description	Level of Output
LLDP	LLDP operating state. The state can be **enabled** or **disabled**. **NOTE**: If a VLAN that has been configured for untagged packets on an interface also has Layer 2 protocol tunneling (L2PT) enabled for LLDP, the LLDP operating state for that interface is displayed as **disabled**.	All levels

Table 765: show lldp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Advertisement interval	Frequency, in seconds, at which LLDP advertisements are sent. This value is set by the **advertisement-interval** configuration statement.	All levels
Transmit delay	Seconds of delay before advertisements are sent to neighbors following a change to a TLV (type, length, or value) element in the LLDP protocol or to the state of the local system, such as a change in hostname or management address. You can set this value to reduce the delay in notifying neighbors of a change in the local system. This value is set by the `transmit-delay` configuration statement.	All levels
Hold timer	On EX4300 switches, the hold timer shows the length of time LLDP information is held before it is discarded. The hold timer value is equal to the advertisement interval multiplied by the hold multiplier. On all other switches, the hold timer shows the value of the hold multiplier. The hold multipler value is set by the `hold-multiplier` configuration statement.	All levels
Notification interval	How often LLDP trap notifications are generated as a result of LLDP database changes. If the interval value is 0, LLDP trap notifications on database changes are disabled. This value is set by the **lldp-configuration-notification-interval** configuration statement.	All levels
Config Trap Interval	How often LLDP trap notifications are generated as a result of changes in topology—for example, when an endpoint connects or disconnects. If the interval value is 0, LLDP trap notifications on topology changes are disabled. This value is set by the **ptopo-configuration-trap-interval** configuration statement.	All levels

Table 765: show lldp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Connection Hold timer	Amount of time the system maintains dynamic topology entries. This value is set by the **ptopo-configuration-maximum-hold-time** configuration statement.	All levels
Port ID TLV subtype	• *interface-name*—Indicates the interface name as the port information for the local device. • `locally-assigned`—Indicates that the sub-type for port ID TLV generation is locally assigned value of SNMP index of the interface.	
Port Description TLV type	Following value used for port description TLV: • `interface-alias (ifAlias)`—Indicates that the *ifAlias* MIB object value is used to generate the port description TLV. • `interface-description (ifDescr)`—Indicates that the *ifDescr* MIB object value is used to generate the port description TLV.	
LLDP-MED	LLDP-MED operating state. The state can be `Enabled` or `Disabled`.	All levels
MED fast start count	Number of advertisements sent from a switch to a device, such as a VoIP telephone, when the device is first detected by the switch. These increased advertisements are temporary. After a device and a switch exchange information and can communicate, advertisements are reduced to one per second. This value is set by using the **fast-start** configuration statement.	All levels
Interface	Name of the interface for which LLDP configuration information is being reported.	All levels
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs.	All levels
LLDP	LLDP operating state. The state can be `Enabled` or `Disabled`.	All levels

Table 765: show lldp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Power Negotiation	LLDP power negotiation operating state. The state can be **Enabled** or **Disabled**.	All levels
Neighbor count	Total number of new LLDP neighbors detected since the last switch reboot.	**detail**
Interface	Name of the interface that is advertising VLAN information.	All levels
Vlan-id	VLAN tag associated with the interface sending LLDP frames. If the interface is not a member of a VLAN, the VLAN ID is advertised as 0.	**detail**
Vlan-name	VLAN name associated with the VLAN ID. For switches running Junos OS releases with Enhanced Layer 2 Software (ELS), this column displays the string `vlan-`*`vlan-id`* by default. Starting in Junos OS Release 15.1X53-D59 and 18.2R1, you can configure the `vlan-name-tlv-option name` option at the `[edit protocols lldp]` hierarchy level to transmit the VLAN name in the LLDP VLAN name TLV in place of the VLAN ID, and display the actual VLAN name in this output field instead.	**detail**

Table 765: show lldp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
LLDP basic TLVs supported	Basic TLVs supported on the switch: • **Chassis identifier**—TLV that advertises the MAC address associated with the local system. • **Port identifier**—TLV that advertises the port identification for the specified port in the local system. • **Port description**—Interface name for the port. • **System name**—TLV that advertises the user-configured name of the local system. • **System description**—TLV that advertises the system description containing information about the software and current image running on the system. This information is taken from the software and is not configurable. • **System capabilities**—TLV that advertises the primary functions performed by the system—for example, bridge or router. • **Management address**—TLV that advertises the IP management address of the local system.	**detail**

Table 765: show lldp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Supported LLDP 802 TLVs	802.3 TLVs supported on the switch: • **MAC/PHY configuration status**—TLV that advertises information about the physical interface, such as autonegotiation status and support and MAU type. The information is based on the physical interface structure and is not configurable. • **Power via MDI**—TLV that advertises MDI power support, PSE power pair, and power class information. • **Link aggregation**—TLV that advertises if the interface is aggregated and its aggregated interface ID. • **Maximum frame size**—TLV that advertises the maximum transmission unit (MTU) of the interface sending LLDP frames. • **Port VLAN tag**—TLV that advertises the VLAN tag configured on the interface. • **Port VLAN name**—TLV that advertises the VLAN name configured on the interface.	**detail**

Table 765: show lldp Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Supported LLDP MED TLVs	LLDP-MED TLVs supported on the switch:
	• **LLDP MED capabilities**—TLV that advertises the primary function of the port. The capabilities values range from 0 through 15:
	  • **0**—Capabilities
	  • **1**—Network Policy
	  • **2**—Location Identification
	  • **3**—Extended Power via MDI-PSE
	  • **4**—Inventory
	  • **5–15**—Reserved
	• **Network policy**—TLV that advertises the port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier, application types—such as voice or streaming video—802.1Q VLAN tagging, and 802.1p priority bits and DiffServ code points.
	• **Endpoint location**—TLV that advertises the physical location of the endpoint.
	• **Extended power Via MDI**—TLV that advertises the power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port.
		detail

Sample Output

show lldp (EX3200 switches)

```
user@switch> show lldp
LLDP                    : Enabled
Advertisement interval : 30 seconds
Transmit delay          : 2 seconds
Hold timer              : 4 seconds
```

```
Notification interval  : 0 Second(s)
Config Trap Interval   : 0 seconds
Connection Hold timer  : 300 seconds

LLDP MED               : Disabled
MED fast start count   : 3 Packets


Interface      Parent Interface    LLDP        LLDP-MED     Power Negotiation
all            -                   Enabled     Enabled      Enabled
```

show lldp (EX4300 switches)

```
user@switch> show lldp
LLDP                   : Enabled
Advertisement interval : 30 seconds
Transmit delay         : 2 seconds
Hold timer             : 120 seconds
Notification interval  : 0 Second(s)
Config Trap Interval   : 0 seconds
Connection Hold timer  : 300 seconds

LLDP MED               : Disabled
MED fast start count   : 3 Packets


Interface      Parent Interface    LLDP        LLDP-MED     Power Negotiation
all            -                   Enabled     Enabled      Enabled
```

show lldp detail (EX4300 switches)

```
user@switch> show lldp detail
LLDP                   : Enabled
Advertisement interval : 30 seconds
Transmit delay         : 2 seconds
Hold timer             : 120 seconds
Notification interval  : 0 Second(s)
Config Trap Interval   : 0 seconds
Connection Hold timer  : 300 seconds
```

```
LLDP MED                : Disabled
MED fast start count    : 3 Packets


Interface      Parent Interface    LLDP        LLDP-MED     Power Negotiation     Neighbor count
all            -                   Enabled     Enabled      Enabled               8

Interface      Parent Interface    Vlan-id     Vlan-name
xe-3/0/0.0     ae31.0              100         vlan-100
xe-3/0/0.0     ae31.0              101         vlan-101
xe-3/0/0.0     ae31.0              4000        vlan-4000
xe-3/0/1.0     ae31.0              100         vlan-100
xe-3/0/1.0     ae31.0              101         vlan-101
xe-3/0/1.0     ae31.0              4000        vlan-4000
xe-3/0/2.0     ae31.0              100         vlan-100
xe-3/0/2.0     ae31.0              101         vlan-101
xe-3/0/2.0     ae31.0              4000        vlan-4000

LLDP basic TLVs supported:
Chassis identifier, Port identifier, Port description, System name, System
description, System capabilities, Management address.

Supported LLDP 802 TLVs:
MAC/PHY configuration/status, Power via MDI, Link aggregation, Maximum frame
size, Port VLAN tag, Port VLAN name.

Supported LLDP MED TLVs:
LLDP MED capabilities, Network policy, Endpoint location, Extended power
Via MDI.
```

show lldp detail (EX3400 switches with VLAN name TLV)

```
user@switch> show lldp detail
LLDP                      : Enabled
Advertisement interval    : 30 seconds
Transmit delay            : 2 seconds
Hold timer                : 120 seconds
Notification interval     : 5 Second(s)
Config Trap Interval      : 0 seconds
Connection Hold timer     : 300 seconds
```

```
LLDP MED                  : Enabled
MED fast start count      : 3 Packets

Port ID TLV subtype       : locally-assigned
Port Description TLV type : interface-alias (ifAlias)


Interface   Parent Interface    LLDP      LLDP-MED    Power Negotiation   Neighbor count
all         -                   -         Enabled     -                   5

Interface   Parent Interface    Vlan-id   Vlan-name
ge-0/0/0    -                   2         dc-vlan
ge-0/0/1    -                   2         dc-vlan
ge-0/0/2    -                   2         dc-vlan
ge-0/0/3    -                   2         dc-vlan
ge-0/0/4    -                   2         dc-vlan
ge-0/0/5    -                   2         dc-vlan
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring LLDP (CLI Procedure)

Configuring LLDP-MED (CLI Procedure)

Understanding LLDP and LLDP-MED on EX Series Switches

Understanding LLDP

show lldp local-information

IN THIS SECTION

Syntax | **21953**

- Description | 21953
- Options | 21953
- Required Privilege Level | 21953
- Output Fields | 21953
- Sample Output | 21955
- Release Information | 21955

Syntax

```
show lldp local-information
```

Description

Display the information that the switch provides in Link Layer Discovery Protocol (LLDP) advertisements to its neighbors.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 766 on page 21954 lists the output fields for the `show lldp local-information` command. Output fields are listed in the approximate order in which they appear.

Table 766: show lldp local-information Output Fields

Field Name	Field Description
LLDP Local Information details	Information about the local system (the switch): • **Chassis ID**—MAC address associated with the switch. • **System name**—User-configured name of the switch. • **System descr**—System description containing information about the switch model and the current software image running on the switch. This information is taken from the software and is not configurable.
System Capabilities	Capabilities (such as **bridge** or **router**) that are supported or enabled on the system.
Management Information	Details of the management information: **Port Name**, **Port Address** (such as **10.204.34.35**), **Address Type** (such as **ipv4** or **ipv6**), **Port ID** (SNMP interface index), **Port ID Subtype**, and **Port Subtype**. The **Port Subtype** displays: • **ifIndex(2)**— IP address of the switch's management Ethernet interface (**me0**) or virtual management Ethernet (**VME**) interface address (for a virtual chassis) is used to manage the switch. • **unknown(1)**—IP management address has been configured with set **protocols lldp management-address**.
Interface name	Name of the local interface which is configured for either LLDP or LLDP-MED.
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the local interface belongs.
SNMP Index	SNMP interface index.
Interface description	User-configured port description.
Status	Administrative status of the interface: either **up** or **down**.
Tunneling	Status of tunneling on the interface: either **enabled** or **disabled**.

Sample Output

show lldp local-information (EX Series Switch)

```
user@switch> show lldp local-information

LLDP Local Information details

Chassis ID   : 00:1d:b5:aa:b9:f0
System name  : switch
System descr : Juniper Networks, Inc. ex8208 , version 10.4I0 [builder] Build
               date: 2010-11-17 12:38:30 UTC

System Capabilities
    Supported       : Bridge Router
    Enabled         : Bridge Router

Management Information
    Port Name       : -
    Port Address    : 10.93.54.6
    Address Type    : IPv4
    Port ID         : 34
    Port ID Subtype : local(7)
    Port Subtype    : ifIndex(2)

Interface name Parent Interface  SNMP Index Interface description Status Tunneling
me0.0          -                 34          -                     Down   Disabled
xe-3/0/0.0     ae31.0            769         xe-3/0/0.0            Up     Disabled
xe-3/0/1.0     ae31.0            770         xe-3/0/1.0            Up     Disabled
xe-3/0/2.0     ae31.0            771         xe-3/0/2.0            Up     Disabled
xe-3/0/3.0     ae31.0            772         xe-3/0/3.0            Up     Disabled
xe-3/0/4.0     ae31.0            577         xe-3/0/4.0            Up     Disabled
xe-3/0/5.0     ae31.0            578         xe-3/0/5.0            Up     Disabled
xe-3/0/6.0     ae31.0            579         xe-3/0/6.0            Up     Disabled
xe-3/0/7.0     ae31.0            581         xe-3/0/7.0            Up     Disabled
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring LLDP (CLI Procedure)

Understanding LLDP and LLDP-MED on EX Series Switches

Understanding LLDP

show lldp neighbors

IN THIS SECTION

- Syntax | 21956
- Description | 21956
- Options | 21957
- Required Privilege Level | 21957
- Output Fields | 21957
- Sample Output | 21961
- Release Information | 21972

Syntax

```
show lldp neighbors
<interface interface-name>
detail
```

Description

Display information about LLDP neighbors.

For information about extended port names in the Junos Fusion technology, see *Understanding Junos Fusion Ports*.

NOTE: When a port with DCBX enabled begins to exchange type, length, and value (TLV) entries, optional LLDP TLVs on that port are not advertised to neighbors in order to interoperate with a wider variety of converged network adapters (CNAs). As a result, information for those ports will not be listed in the output for this command.

NOTE: The Chassis ID TLV has a subtype for Network Address Family. The supported network address families are IPv4 and IPv6. LLDP frames are validated only if the Network Address subtype of the Chassis ID TLV has a value of 1 (IPv4) or 2 (IPv6). For any other value, the transmitting device is detected by LLDP as a neighbor and displayed in the output of the `show lldp neighbors` command, but it is not assigned to the VLAN.

Options

none	Display learned LLDP information on all neighboring interfaces and devices.
interface ***interface-name***	(Optional) Display the neighbor information about a particular physical interface.
detail	(Optional) Display detailed information about LLDP neighbors.

Required Privilege Level

view

Output Fields

Table 767 on page 21957 describes the output fields for the `show lldp neighbors` command. Output fields are listed in the approximate order in which they appear.

Table 767: show lldp neighbors Output Fields

Field Name	Field Description
`LLDP Remote Devices Information`	Information about remote devices.
`Local Interface`	List of local interfaces for which neighbor information is available.

Table 767: show lldp neighbors Output Fields *(Continued)*

Field Name	Field Description
`Parent Interface`	List of aggregated Ethernet interfaces, if any, to which the local interfaces belong.
`Chassis Id`	List of chassis identifiers for neighbors.
`Port info`	List of port information gathered from neighbors. This could be the port identifier or port description.
`System name`	List of system names gathered from neighbors, which includes the host name and the domain name.
`LLDP Neighbor Information`	Information about both local and neighbor systems on the interface (appears when the `interface` option is used).
`Local Information`	Information about the local system on the interface (appears when the `interface` option is used).
`Neighbor Information`	Information about the neighbor system on the interface (appears when the `interface` option is used).
`Index`	Local interface index (appears when the `interface` option is used).
`Time Mark`	Date and timestamp of information (appears when the `interface` option is used).
`Time To Live`	Number of seconds for which this information is valid (appears when the `interface` option is used).
`Age`	Age (in seconds) of the neighbor since the TLV is received from the neighbor (appears when the `interface` option is used).
`Local Interface`	Name of the local physical interface (appears when the `interface` option is used).

Table 767: show lldp neighbors Output Fields *(Continued)*

Field Name	Field Description
`Parent Interface`	Name of the aggregated Ethernet interface, if any, to which the interface belongs (appears when the **interface** option is used).
`Local Port ID`	Local interface SNMP index (appears when the interface option is used).
`Ageout Count`	Number of times the complete set of information advertised by the neighbor has been deleted from LLDP neighbor information maintained by the local system because the information timeliness interval has expired (appears when the `interface` option is used).
`Neighbor Information`	Information about a neighbor system on the interface (appears when the `interface` option is used).
`Chassis type`	Type of chassis identifier supplied, such as `MAC address` (appears when the `interface` option is used).
`Chassis ID`	Chassis identifier of the chassis type listed (appears when the `interface` option is used).
`Port type`	Type of port identifier supplied, such as `Locally assigned` (appears when the `interface` option is used).
`Port ID`	Port identifier of the port type listed (appears when the `interface` option is used).
`Port description`	Configured port description, port name, or the SNMP ifIndex (appears when the `interface` option is used).
`System name`	Name supplied by the system on the interface (appears when the `interface` option is used). When NetBIOS snooping is enabled, this is the NetBIOS name of the host.
`System Description`	Description supplied by the system on the interface (appears when the `interface` option is used).

Table 767: show lldp neighbors Output Fields *(Continued)*

Field Name	Field Description
System Capabilities	Capabilities (such as Bridge, Router, and Telephone) that are Supported or Enabled by the system on the interface (appears when the interface option is used).
Management address	Details of the management address: Address Type (such as IPv4 and IPv6), Address (such as 10.204.34.35, 1fd::1a10), Interface Number, Interface Subtype, and object identifier (OID) (appears when the interface option is used).
Management Info	Details of management information: Type (such as IPv4 or IPv6), Address (such as 10.204.34.35), Port ID, Subtype, Interface Subtype, and object identifier (OID) (appears when the interface option is used). The Interface Subtype displays: • ifIndex(2)— IP address of the neighbor's management Ethernet interface (me0) or virtual management Ethernet (VME) interface address (for a Virtual Chassis) that is used to manage the switch. • unknown(1)—Neighbor's IP management address has been configured with set protocols lldp management-address.
Media Info	Additional details about the endpoint device when a device that supports LLDP-MED is attached to the interface. The specific details depend upon the capabilities of the device. Details might include: Media endpoint class (such as Class 3 for communication devices such as IP phones), MED Hardware revision, MED Firmware revision, MED Software revision, MED Serial number, MED Manufacturer name, MED Model name.
Organization Info	One or more entries (indexed by the Index element) listing remote information by Organizationally Unique Identifier (OUI), Subtype, Index, and Info (appears when the interface option is used).

Sample Output

show lldp neighbors

```
user@switch> show lldp neighbors

Local Interface   Parent Interface   Chassis Id          Port info    System Name
xe-3/0/4.0        ae31.0             b0:c6:9a:63:80:40   xe-0/0/0.0   host.jnpr.net
xe-3/0/5.0        ae31.0             b0:c6:9a:63:80:40   xe-0/0/1.0   host.jnpr.net
xe-3/0/6.0        ae31.0             b0:c6:9a:63:80:40   xe-0/0/2.0   host.jnpr.net
xe-3/0/7.0        ae31.0             b0:c6:9a:63:80:40   xe-0/0/3.0   host.jnpr.net
xe-3/0/0.0        ae31.0             b0:c6:9a:63:80:40   xe-0/1/0.0   host.jnpr.net
xe-3/0/1.0        ae31.0             b0:c6:9a:63:80:40   xe-0/1/1.0   host.jnpr.net
xe-3/0/2.0        ae31.0             b0:c6:9a:63:80:40   xe-0/1/2.0   host.jnpr.net
xe-3/0/3.0        ae31.0             b0:c6:9a:63:80:40   xe-0/1/3.0   host.jnpr.net
```

show lldp neighbors (Management Ethernet Interfaces)

```
user@host> show lldp neighbors

Local Interface    Parent Interface    Chassis Id          Port info          System Name
fxp0               -                   78:fe:3d:ee:4e:00   151                x2-sw35
xe-0/0/0           -                   a8:d0:e5:50:26:c0   512                sitara
xe-0/0/1           -                   a8:d0:e5:50:26:c0   513                sitara
```

show lldp neighbors interface ge-0/0/8

```
user@switch> show lldp neighbors interface ge-0/0/8

LLDP Neighbor Information:
Local Information:
Index: 1 Time to live: 120 Time mark: Thu Nov 26 06:41:24 2015 Age: 1 secs
Local Interface    : ge-0/0/8
Parent Interface   : -
Local Port ID      : 518
Ageout Count       : 0
```

```
Neighbour Information:
Chassis type       : Mac address
Chassis ID         : 88:e0:f3:1f:14:e0
Port type          : Locally assigned
Port ID            : 880
Port description   : ge-0/0/8
System name        : bng-nw6moj.juniper.net

System Description : Juniper Networks, Inc. ex4300-24p Ethernet Switch, kernel JUNOS
14.1I20151125_0548_rajjs, Build date: 2015-11-25 06:06:58 UTC Copyright (c) 1996-2015 Juniper
Networks, Inc.

System capabilities
        Supported: Bridge Router
        Enabled  : Bridge Router

Management address
        Address Type      : IPv4(1)
        Address           : 10.204.39.232
        Interface Number  : 33
        Interface Subtype : ifIndex(2)
        OID               : 1.3.6.1.2.1.31.1.1.1.1.33.
Media endpoint class: Network Connectivity

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MAC/PHY Configuration/Status (1)
       Info     : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation Capability
(0x1), MAU Type (0x0)
       Index    : 1

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MDI Power (2)
       Info     : MDI Power Support [PSE bit set, supported, disabled, CONTROL bit not set
(0x3)], MDI Power Pair [signal], MDI Power Class [Unknown (7)]
       Index    : 2

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Link Aggregation (3)
       Info     : Aggregation Status [supported, disabled (0x1)], Aggregation Port ID (0)
       Index    : 3
```

```
Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Maximum Frame Size (4)
       Info     : MTU Size (1514)
       Index    : 4

Organization Info
       OUI      : Juniper Specific (0x009069)
       Subtype  : Chassis Serial Type (1)
       Info     : Juniper Slot Serial [MS3112240009]
       Index    : 5
```

show lldp neighbors interface ge-0/0/4 (Management Address is IPv4)

```
user@host> show lldp neighbors interface ge-0/0/4
LLDP Neighbor Information:
Local Information:
Index: 2 Time to live: 120 Time mark: Tue Dec 31 11:47:46 2013 Age: 15 secs
Local Interface    : ge-2/0/1
Parent Interface   : ae0
Local Port ID      : 1476
Ageout Count      : 0

Neighbour Information:
Chassis type       : Mac address
Chassis ID         : ac:4b:c8:92:67:c0
Port type          : Locally assigned
Port ID            : 529
Port description   : ge-1/2/3
System name        : apg-hp

System Description : Juniper Networks, Inc. mx240 , version 14.1-20131222.0 [builder] Build
date: 2013-12-22 09:13:26 UTC



System capabilities
        Supported: Bridge Router
        Enabled  : Bridge Router

Management address
```

```
         Address Type      : IPv4(1)
         Address           : 10.216.98.57
         Interface Number  : 1
         Interface Subtype : ifIndex(2)
         OID               : 1.3.6.1.2.1.31.1.1.1.1.1.

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MAC/PHY Configuration/Status (1)
       Info     : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation Capability
(0x1d), MAU Type (0x0)
       Index    : 1

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Link Aggregation (3)
       Info     : Aggregation Status (0x3), Aggregation Port ID (1694498816)
       Index    : 2

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Maximum Frame Size (4)
       Info     : MTU Size (1518)
       Index    : 3
```

show lldp neighbors interface ge-0/0/4 (Management Address is IPv6)

```
user@host> show lldp neighbors interface ge-0/0/4
LLDP Neighbor Information:
Local Information:
Index: 1 Time to live: 120 Time mark: Thu Dec 12 07:19:45 2013 Age: 28 secs
Local Interface    : ge-1/2/2
Parent Interface   : -
Local Port ID      : 528
Ageout Count      : 0

Neighbour Information:
Chassis type       : Mac address
Chassis ID         : 64:87:88:65:37:c0
Port type          : Locally assigned
```

```
Port ID            : 1475
Port description   : ge-2/0/0
System name        : apg-hp1

System Description : Juniper Networks, Inc. mx240 , version 11.4R10 Build date: 2013-10-24
10:10:02 UTC



System capabilities
        Supported: Bridge Router
        Enabled  : Bridge Router

Management address
        Address Type      : IPv6(2)
        Address           : 1fd::1a10
        Interface Number  : 1
        Interface Subtype : ifIndex(2)
        OID               : 1.3.6.1.2.1.31.1.1.1.1.1.

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MAC/PHY Configuration/Status (1)
       Info     : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation Capability
(0x5), MAU Type (0x0)
       Index    : 1

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Link Aggregation (3)
       Info     : Aggregation Status (0x1), Aggregation Port ID (0)
       Index    : 2

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Maximum Frame Size (4)
       Info     : MTU Size (1518)
       Index    : 3

Organization Info
       OUI      : Ethernet Bridged (0x0080c2)
       Subtype  : VLAN Name (3)
```

```
        Info     : VLAN ID (100), VLAN Name (vlan-100)
        Index    : 4
```

show lldp neighbors detail

```
user@host> show lldp neighbors detail
LLDP Neighbor Information:
Local Information:
Index: 1 Time to live: 120 Time mark: Wed Apr  3 08:25:57 2019 Age: 8 secs
Local Interface    : me0
Parent Interface   : -
Local Port ID      : 33
Ageout Count       : 0

Neighbour Information:
Chassis type       : Mac address
Chassis ID         : 1a:22:23:dc:d9:50
Port type          : Locally assigned
Port ID            : 517
Port description   : ge-0/0/7.0
System name        : test

System Description : Juniper Networks, Inc. ex3300-48t , version 12.3R Build date: 2014-03-13
07:02:54 UTC



System capabilities
        Supported: Bridge Router
        Enabled  : Bridge Router

Management address
        Address Type      : IPv4(1)
        Address           : 10.221.0.111
        Interface Number  : 34
        Interface Subtype : ifIndex(2)
        OID               : 1.3.6.1.2.1.31.1.1.1.1.34.
Media endpoint class: Network Connectivity

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MAC/PHY Configuration/Status (1)
```

```
       Info     : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation Capability
(0x6c11), MAU Type (0x0)
       Index    : 1

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Link Aggregation (3)
       Info     : Aggregation Status [supported, disabled (0x1)], Aggregation Port ID (0)
       Index    : 2

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Maximum Frame Size (4)
       Info     : MTU Size (1514)
       Index    : 3

Organization Info
       OUI      : Ethernet Bridged (0x0080c2)
       Subtype  : Port Vid (1)
       Info     : VLAN ID (52),
       Index    : 4

Organization Info
       OUI      : Juniper Specific (0x009069)
       Subtype  : Chassis Serial Type (1)
       Info     : Juniper Slot Serial [GA0215270535]
       Index    : 5

Organization Info
       OUI      : Ethernet Bridged (0x0080c2)
       Subtype  : VLAN Name (3)
       Info     : VLAN ID (52), VLAN Name (vlan52)
       Index    : 6
       Index    : 7

LLDP Neighbor Information:
Local Information:
Index: 2 Time to live: 120 Time mark: Wed Apr  3 08:25:56 2019 Age: 9 secs
Local Interface    : ge-0/0/2
Parent Interface   : -
Local Port ID      : 512
Ageout Count       : 0
```

```
Neighbour Information:
Chassis type       : Mac address
Chassis ID         : 1a:22:23:dc:d9:50
Port type          : Locally assigned
Port ID            : 511
Port description   : ge-0/0/2

System Description : Juniper Networks, Inc. ex4300-24p Ethernet Switch, kernel JUNOS 17.1, Build
date: 2019-03-26 08:12:13 UTC Copyright (c) 1996-2019 Juniper Networks, Inc.



System capabilities
        Supported: Bridge Router
        Enabled  : Bridge Router

Management address
        Address Type      : IPv4(1)
        Address           : 10.204.39.247
        Interface Number  : 33
        Interface Subtype : ifIndex(2)
        OID               : 1.3.6.1.2.1.31.1.1.1.1.33.
Media endpoint class: Network Connectivity

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MAC/PHY Configuration/Status (1)
       Info     : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation Capability
(0x1), MAU Type (0x0)
       Index    : 1

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MDI Power (2)
       Info     : MDI Power Support [PSE bit set, supported, disabled, CONTROL bit not set
(0x3)], MDI Power Pair  [signal], MDI Power Class [class1 (1)]
       Index    : 2

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Link Aggregation (3)
       Info     : Aggregation Status [supported, enabled (0x3)], Aggregation Port ID (556)
       Index    : 3
```

```
Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Maximum Frame Size (4)
       Info     : MTU Size (6500)
       Index    : 4

Organization Info
       OUI      : Juniper Specific (0x009069)
       Subtype  : Chassis Serial Type (1)
       Info     : Juniper Slot Serial [MS3112240002]
       Index    : 5

Organization Info
       OUI      : Ethernet Bridged (0x0080c2)
       Subtype  : VLAN Name (3)
       Info     : VLAN ID (100), VLAN Name (vlan-100)
       Index    : 6
       Index    : 7
       Index    : 8
```

show lldp neighbors interface ge-0/0/0.0 (for a VoIP AvayaTelephone with LLDP-MED Support)

```
user@switch>show lldp neighbors interface ge-0/0/0.0

LLDP Neighbor Information:
Local Information:
Index: 20 Time to live: 120 Time mark: Thu Apr 15 22:26:22 2010 Age: 16 secs
Local Interface    : ge-0/0/0.0
Parent Interface   : -
Local Port ID      : 517
Ageout Count       : 0

Neighbour Information:
Chassis type       : Network address
Chassis ID         : 0.0.0.0
Port type          : Mac address
Port ID            : 00:04:0d:fc:55:48
System name        : AVAFC5548.juniper.net

System capabilities
```

```
        Supported  : Bridge Telephone
        Enabled    : Bridge

Management Info
        Type              : IPv4
        Address           : 0.0.0.0
        Port ID           : 1
        Subtype           : 1
        Interface Subtype : ifIndex(2)
        OID               : 1.3.6.1.2.1.31.1.1.1.1.1
Media endpoint class: Class III Device

MED Hardware revision : 4610D01A
MED Firmware revision : b10d01b2_9.bin
MED Software revision : a10d01b2_9.bin
MED Serial number     : 07N510103424
MED Manufacturer name : Avaya
MED Model name        : 4610

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MAC/PHY Configuration/Status (1)
       Info     : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation Capability
(0x1d00), MAU Type (0x0)
       Index    : 1

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MDI Power (2)
       Info     : MDI Power Support [PSE bit set, supported, disabled, CONTROL bit not set
(0x3)], MDI Power Pair [signal], MDI Power Class [Unknown (7)]
       Index    : 2

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Link Aggregation (3)
       Info     : Aggregation Status [supported, disabled (0x1)], Aggregation Port ID (0)
       Index    : 3

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Maximum Frame Size (4)
       Info     : MTU Size (1514)
```

```
        Index    : 4

Organization Info
        OUI      : Ethernet Bridged (0x0080c2)
        Subtype  : Port Vid (1)
        Info     : VLAN ID (10),
        Index    : 5

Organization Info
        OUI      : Juniper Specific (0x009069)
        Subtype  : Chassis Serial Type (1)
        Info     : Juniper Slot Serial [BQ0208211462]
        Index    : 6

Organization Info
        OUI      : Ethernet Bridged (0x0080c2)
        Subtype  : VLAN Name (3)
        Info     : VLAN ID (10), VLAN Name (vtest)
        Index    : 7
```

show lldp neighbors interface ge-0/0/5.0 (with NetBIOS Snooping Enabled on the Switch)

```
user@switch> show lldp neighbors interface ge-0/0/5

Age: 299999 secs
Local Interface    : ge-0/0/5.0
Parent Interface   : -
Chassis ID         : 00:10:94:00:00:02
Port description   : 192.0.2.1
System name        : host.juniper.net
```

show lldp neighbors detail (l2ng and non-l2ng platform)

```
user@switch show lldp neighbors detail
LLDP Neighbor Information:
Local Information:
Index: 11 Time to live: 20 Time mark: Fri Feb  4 11:15:46 2022 Age: 3 secs
Local Interface    : ge-0/0/2
Parent Interface   : -
```

```
Local Port ID      : 521
Ageout Count       : 1

Neighbour Information:
Chassis type       : Mac address
Chassis ID         : 00:05:86:90:bd:c0
Port type          : Interface name
Port ID            : ge-0/0/2
Port description   : INTF2_DISECR_CE1
System name        : CE1_re0
```

show lldp neighbors interface (l2ng and non-l2ng platform)

```
user@switch show lldp neighbors interface ge-0/0/2
LLDP Neighbor Information:
Local Information:
Index: 3 Time to live: 20 Time mark: Fri Feb  4 11:16:50 2022 Age: 3 secs
Local Interface    : ge-0/0/2
Parent Interface   : -
Local Port ID      : 521
Ageout Count      : 0

Neighbour Information:
Chassis type       : Mac address
Chassis ID         : 2c:6b:f5:0d:3d:c0
Port type          : Interface name
Port ID            : ge-0/0/2
Port description   : FIRST_INTF_DESC_CE2
System name        : CE2_re0.example.com
```

Release Information

Command introduced in Junos OS Release 9.6.

`detail` option introduced in Junos OS Release 19.1R2.

RELATED DOCUMENTATION

clear lldp neighbors | 3685

Understanding LLDP

Configuring LLDP (CLI Procedure)

Understanding LLDP and LLDP-MED on EX Series Switches

show lldp neighbors-vlan-name-tlv-list interface

IN THIS SECTION

- Syntax | 21973
- Description | 21973
- Required Privilege Level | 21973
- Sample Output | 21973
- Release Information | 21974

Syntax

```
show lldp neighbors-vlan-name-tlv-list interface <name>
```

Description

Display the VLAN Name TLVs information about neighboring devices learned by the device by using the Link Layer Discovery Protocol (LLDP).

Required Privilege Level

view

Sample Output

```
user@switch> show lldp neighbors-vlan-name-tlv-list interface <name>
Organization Info
OUI      : Ethernet Bridged (0x0080c2)
```

```
Subtype   : VLAN Name (3)
 Vlan-id        Vlan-name      Index
 100            vlan-100       7

 101            vlan-101       8

 102            vlan-102       9

 103            vlan-103       10

 104            vlan-104       11

 105            vlan-105       12

 106            vlan-106       13

 107            vlan-107       14
```

Release Information

Command introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

- Configuring the Transmission of Maximum VLAN Name TLVs in LLDP
- *Configuring LLDP (CLI Procedure)*
- *Understanding LLDP*

show lldp remote-global-statistics

IN THIS SECTION

- Syntax | 21975
- Description | 21975
- Options | 21975

- Required Privilege Level | 21975
- Output Fields | 21975
- Sample Output | 21976
- Release Information | 21976

Syntax

```
show lldp remote-global-statistics
```

Description

Display remote Link Layer Discovery Protocol (LLDP) global statistics.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 768 on page 21975 describes the output fields for the `show lldp remote-global-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 768: show lldp remote-global-statistics Output Fields

Field Name	Field Description
LLDP Remote Database Table Counters	Information about remote database table counters.

Table 768: show lldp remote-global-statistics Output Fields *(Continued)*

Field Name	Field Description
LastchangeTime	Time elapsed between LLDP agent startup and the last change to the remote database table information.
Inserts	Number of insertions made in the remote database table.
Deletes	Number of deletions made in the remote database table.
Drops	Number of LLDP frames dropped from the remote database table because of errors.
Ageouts	Number of remote database table entries that have aged out of the table.

Sample Output

show lldp remote-global-statistics

```
user@host> show lldp remote-global-statistics
user@host> show lldp remote-global-statistics
   LLDP Remote Database Table Counters
   LastchangeTime          Inserts    Deletes    Drops    Ageouts
   00:00:76 (76 sec)       192        0          0        0
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring LLDP (CLI Procedure)

Understanding LLDP and LLDP-MED on EX Series Switches

show lldp statistics

IN THIS SECTION

- Syntax | 21977
- Description | 21977
- Options | 21977
- Required Privilege Level | 21977
- Output Fields | 21978
- Sample Output | 21978
- Release Information | 21980

Syntax

```
show lldp statistics
<interface interface-name>
```

Description

Display information about Link Layer Discovery Protocol (LLDP) statistics.

Options

none	Display LLDP statistics on all interfaces and devices.
interface *interface-name*	(Optional) Display LLDP statistics for the specified physical interface.

Required Privilege Level

view

Output Fields

Table 769 on page 21978 describes the output fields for the `show lldp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 769: show lldp statistics Output Fields

Field Name	Field Description
Interface	Interface name. For information about extended port names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs. **NOTE**: Because LLDP packets are transmitted and received on member interfaces only, statistics are available only for the member interfaces, not for the aggregated interface.
Received	Number of LLDP frames received on this interface.
Unknown TLVs	Number of LLDP frames with unsupported content received on this interface.
With Errors	Number of LLDP frames with errors received on this interface.
Discarded	Number of LLDP frames received and then discarded on this interface.
Transmitted	Number of LLDP frames sent on this interface.
Untransmitted	Number of LLDP frames that were untransmitted on this interface.

Sample Output

show lldp statistics

```
user@host> show lldp statistics
```

```
Interface    Parent Interface  Received  Unknown TLVs  With Errors
xe-3/0/0.0   ae31.0            1564      0             0
xe-3/0/1.0   ae31.0            1564      0             0
xe-3/0/2.0   ae31.0            1565      0             0
xe-3/0/3.0   ae31.0            1566      0             0
xe-3/0/4.0   ae31.0            1598      0             0
xe-3/0/5.0   ae31.0            1598      0             0
xe-3/0/6.0   ae31.0            1596      0             0
xe-3/0/7.0   ae31.0            1597      0             0
xe-5/0/6.0   -                 0         0             0
xe-5/0/7.0   -                 0         0             0


Discarded TLVs  Transmitted  Untransmitted
0               3044         1
0               3044         1
0               3044         1
0               3044         1
0               3075         1
0               3075         1
0               3075         1
0               3075         1
0               17312        0
0               17312        0
```

show lldp statistics interface ge-0/1/1

```
user@host> show lldp statistics interface ge-0/1/1
Interface Received Transmitted Unknown-TLVs With-Errors Discarded
--------- -------- ----------- ------------ ----------- ---------
ge-0/1/1  544      540         0            0           0
```

show lldp statistics interface xe-3/0/0.0

```
user@switch> show lldp statistics interface xe-3/0/0.0

Interface    Parent Interface  Received  Unknown TLVs  With Errors
xe-3/0/0.0   ae31.0            1566      0             0
```

```
Discarded TLVs  Transmitted  Untransmitted
0               3046         1
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

clear lldp statistics | 3686

Configuring LLDP (CLI Procedure)

Understanding LLDP and LLDP-MED on EX Series Switches

show log

IN THIS SECTION

- Syntax | 21980
- Syntax (QFX Series and OCX Series) | 21981
- Syntax (TX Matrix Router) | 21981
- Description | 21981
- Options | 21981
- Required Privilege Level | 21982
- Sample Output | 21982
- Release Information | 21986

Syntax

```
show log
<filename | user <username>>
```

Syntax (QFX Series and OCX Series)

```
show log filename
<device-type (device-id | device-alias)>
```

Syntax (TX Matrix Router)

```
show log
<all-lcc | lcc number | scc>
<filename / user <username>>
```

Description

List log files, display log file contents, or display information about users who have logged in to the router or switch.

> **NOTE**: On MX Series routers, modifying a configuration to replace a service interface with another service interface is treated as a catastrophic event. When you modify a configuration, the entire configuration associated with the service interface—including NAT pools, rules, and service sets—is deleted and then re-created for the newly specified service interface. If there are active sessions associated with the service interface that is being replaced, these sessions are deleted and the NAT pools are then released, which leads to the generation of the NAT_POOL_RELEASE system log messages. However, because NAT pools are already deleted as a result of the catastrophic configuration change and no longer exist, the NAT_POOL_RELEASE system log messages are not generated for the changed configuration.

Options

none — List all log files.

<all-lcc | lcc *number* | scc> — (Routing matrix only)(Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace *number* with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).

device-type — (QFabric system only) (Optional) Display log messages for only one of the following device types:

- `director-device`—Display logs for Director devices.
- `infrastructure-device`—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).
- `interconnect-device`—Display logs for Interconnect devices.
- `node-device`—Display logs for Node devices.

> **NOTE**: If you specify the `device-type` optional parameter, you must also specify either the `device-id` or `device-alias` optional parameter.

(*device-id* | *device-alias*)

If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename

(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.

> **NOTE**: The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of **messages**.

user <*username*>

(Optional) Display logging information about users who have recently logged in to the router or switch. If you include *username*, display logging information about the specified user.

Required Privilege Level

trace

Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r--  1 root  bin      211663 Oct  1 19:44 dcd
-rw-r--r--  1 root  bin      999947 Oct  1 19:41 dcd.0
```

```
-rw-r--r--  1 root  bin      999994 Oct  1 17:48 dcd.1
-rw-r--r--  1 root  bin      238815 Oct  1 19:44 rpd
-rw-r--r--  1 root  bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r--  1 root  bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r--  1 root  bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r--  1 root  bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r--  1 root  bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r--  1 root  bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r--  1 root  bin     1056350 Sep 30 07:04 rpd.6
-rw-r--r--  1 root  bin     1048876 Sep 30 03:21 rpd.7
-rw-rw-r--  1 root  bin       19656 Oct  1 19:37 wtmp
```

show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr 192.0.2.21 nhop type
local nhop 192.0.2.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr 192.0.2.22 nhop type
unicast nhop 192.0.2.22
Oct  1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex 43
Oct  1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

user@host:LSYS1> show log flow_lsys1.log
Nov  7 07:34:09 07:34:09.491800:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table lock

Nov  7 07:34:09 07:34:09.491809:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route table lock

Nov  7 07:34:09 07:34:09.491840:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table lock

Nov  7 07:34:09 07:34:09.491841:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route table lock

Nov  7 07:34:09 07:34:09.491854:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh 0x0

Nov  7 07:34:09 07:34:09.491868:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table lock
```

```
Nov  7 07:34:09 07:34:09.491869:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route table lock

Nov  7 07:34:09 07:34:09.491881:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh 0x0
user@host:TSYS1> show log flow_tsys1.log
Nov  7 13:21:47 13:21:47.217744:CID-0:THREAD_ID-05:LSYS_ID-32:RT:<192.0.2.0/0-
>198.51.100.0/9011;1,0x0> :

Nov  7 13:21:47 13:21:47.217747:CID-0:THREAD_ID-05:LSYS_ID-32:RT:packet [84] ipid = 39281,
@0x7f490ae56d52

Nov  7 13:21:47 13:21:47.217749:CID-0:THREAD_ID-05:LSYS_ID-32:RT:---- flow_process_pkt: (thd 5):
flow_ctxt type 0, common flag 0x0, mbuf 0x4882b600, rtbl7

Nov  7 13:21:47 13:21:47.217752:CID-0:THREAD_ID-05:LSYS_ID-32:RT: flow process pak fast ifl 88
in_ifp lt-0/0/0.101

Nov  7 13:21:47 13:21:47.217753:CID-0:THREAD_ID-05:LSYS_ID-32:RT:  lt-0/0/0.101:192.0.2.0-
>198.51.100.0, icmp, (0/0)

Nov  7 13:21:47 13:21:47.217756:CID-0:THREAD_ID-05:LSYS_ID-32:RT: find flow: table 0x11d0a2680,
hash 20069(0xffff), sa 192.0.2.0, da 198.51.100.0, sp 0, d0

Nov  7 13:21:47 13:21:47.217760:CID-0:THREAD_ID-05:LSYS_ID-32:RT:Found: session id 0x12. sess
tok 28685

Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT:  flow got session.

Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT:  flow session id 18

Nov  7 13:21:47 13:21:47.217763:CID-0:THREAD_ID-05:LSYS_ID-32:RT: vector bits 0x200 vector
0x84ae85f0

Nov  7 13:21:47 13:21:47.217764:CID-0:THREAD_ID-05:LSYS_ID-32:RT:set nat 0x11e463550(18) timeout
const to 2

Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT: set_nat_timeout 2 on session 18

Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT:refresh nat 0x11e463550(18)
timeout to 2

Nov  7 13:21:47 13:21:47.217767:CID-0:THREAD_ID-05:LSYS_ID-32:RT:insert usp tag for apps
```

```
Nov  7 13:21:47 13:21:47.217768:CID-0:THREAD_ID-05:LSYS_ID-32:RT:mbuf 0x4882b600, exit nh
0xfffb0006
```

show log filename (QFabric System)

```
user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486 chassisd:
CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 1,
jnxFruL2Index 1, jnxFruL3Index 0, jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11,
jnxFruSlot 0, jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486 chassisd:
CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 1,
jnxFruL2Index 2, jnxFruL3Index 0, jnxFruName PIC:  @ 0/1/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492 chassisd:
CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 1,
jnxFruL2Index 1, jnxFruL3Index 0, jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11,
jnxFruSlot 0, jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492 chassisd:
CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 1,
jnxFruL2Index 2, jnxFruL3Index 0, jnxFruName PIC:  @ 0/1/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486 file: UI_COMMIT:
User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486 file: UI_COMMIT:
User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50 _DCF_default___NW-
INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50 _DCF_default___NW-
INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492 file: UI_COMMIT:
User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492 file: UI_COMMIT:
User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491 chassisd:
CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 1,
jnxFruL2Index 1, jnxFruL3Index 0, jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11,
jnxFruSlot 0, jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)
```

show log user

```
user@host> show log user
usera    mg2546                  Thu Oct  1 19:37   still logged in
usera    mg2529                  Thu Oct  1 19:08 - 19:36  (00:28)
usera    mg2518                  Thu Oct  1 18:53 - 18:58  (00:04)
root     mg1575                  Wed Sep 30 18:39 - 18:41  (00:02)
root     ttyp2    aaa.bbbb.com   Wed Sep 30 18:39 - 18:41  (00:02)
userb    ttyp1    192.0.2.0      Wed Sep 30 01:03 - 01:22  (00:19)
```

show log accepted-traffic (SRX4600, SRX5400, SRX5600, and SRX5800)

```
user@host> show log accepted-traffic
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.5/2-
>4.4.4.2/63 0x0 None 3.3.3.5/2->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2617282058
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.4/4-
>4.4.4.2/63 0x0 None 3.3.3.4/4->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2550162754
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.4/1-
>4.4.4.2/63 0x0 None 3.3.3.4/1->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2550162755
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.3/0-
>4.4.4.2/63 0x0 None 3.3.3.3/0->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2550162752
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.5/5-
>4.4.4.2/63 0x0 None 3.3.3.5/5->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2550162751
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.3/3-
>4.4.4.2/63 0x0 None 3.3.3.3/3->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2550162753
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
```

Release Information

Command introduced before Junos OS Release 7.4.

Option *device-type (device-id | device-alias)* is introduced in Junos OS Release 13.1 for the QFX Series.

RELATED DOCUMENTATION

syslog (System)

show loop-detect enhanced interface

IN THIS SECTION

- Syntax | 21987
- Description | 21987
- Options | 21988
- Required Privilege Level | 21988
- Output Fields | 21988
- Sample Output | 21989
- Release Information | 21993

Syntax

```
show loop-detect enhanced interface
<interface-name>
```

Description

Display status information about leaf device logical interfaces with loop detection enabled in a Layer 2 network.

You can configure loop detection on leaf devices for server-facing Layer 2 logical interfaces, and specify an action that the leaf device performs on the interface when it detects a loop, such as bringing the interface down. You can also configure an interval after which the device will automatically revert the interface to its state prior to the loop detection action following repair of the loop condition.

After you’ve enabled loop detection, the interfaces periodically send multicast loop-detection protocol data units (PDUs). If the interface receives a loop-detection PDU, which indicates a loop exists, the interface triggers the configured action to break the loop. The network administrator might then need to repair the conditions that caused the loop before reverting the interface to its state prior to loop

detection. This command shows the configured loop detection parameters, current interface status, and detail about the remote interface that formed the loop.

Options

interface-name

Display loop detection status for the specified logical interface.

If you don't specify an interface name, this command displays loop detection status for all logical interfaces that have this feature enabled.

Required Privilege Level

view

Output Fields

Table 770: show loop-detect enhanced interface Output Fields

Field Name	Field Description
Interface	Local interface name with loop detection configured.
Vlan-id	Local interface VLAN ID for which loop detection is configured.
ESI	Ethernet Segment ID (ESI) associated with the local interface.
Current status	Current status of the interface—Loop-detected or Normal [Link Up] (no loop detected) This output section displays the following additional fields if the interface detected a loop. These fields describe the interface that routed loop detection PDUs back to the local interface: • Remote Host: Remote host name. • Remote Chassis: Remote host chassis ID. • Remote Interface: Remote interface name that is looping PDUs. • Remote ESI: ESI associated with the remote interface. **NOTE**: The interface causing the loop might be on the same local device or on a remote device.

Table 770: show loop-detect enhanced interface Output Fields *(Continued)*

Field Name	Field Description
Last loop-detect time	Timestamp when the device initially detected the loop and changed the interface status to Loop-detected. When the interface has no `Revert Interval` set, or the `Revert Interval` timer is counting down the revert time after a loop has been cleared, if the feature detects a new loop, it updates the `Last loop-detect time` output field to display that most recent time loop detection time.
Receive statistics	Number of loop-detection PDUs received since the current loop detection period timestamp.
Action configured	The configured action that the device triggers upon loop detection.
Action count	Cumulative count of loop detection actions triggered on the local interface.
Transmit Interval	Default or configured interval at which the interface sends loop detection PDUs.
Revert Interval	Configured interval after the loop condition is repaired in which the interface state reverts to the previous state before the loop detection action occurred.

Sample Output

show loop-detect enhanced interface *interface-name* (no loop detected)

```
user@host> show loop-detect enhanced interface xe-0/0/12.0
```

```
Interface                  :xe-0/0/12.0
Vlan-id                    :100
ESI                        :00:00:00:00:00:00:00:00:00:00
Current status             :Normal[Link Up]
Last loop-detect time      :-
Receive statistics         :0
Action configured          :Interface-down
Action count               :0
```

```
Transmit Interval          :1s
Revert Interval            :60s
```

show loop-detect enhanced interface *interface-name* (loop detected)

```
user@host> show loop-detect enhanced interface xe-0/0/12.0
```

```
Interface                  :xe-0/0/12.0
Vlan-id                    :100
ESI                        :00:00:00:00:00:00:00:00:00:00
Current status             :Loop-detected
                Remote Host       :leaf04
                Remote Chassis    :ab:cd:ef:12:34:56
                Remote Interface  :xe-0/0/12.0
                Remote ESI        :00:00:00:00:00:00:00:00:00:00
Last loop-detect time      :Tue Oct 20 04:36:37 2020
Receive statistics         :4
Action configured          :Interface-down
Action count               :1
Transmit Interval          :1s
Revert Interval            :60s
```

show loop-detect enhanced interface (EX4300-48MP, all interfaces configured with loop detection)

```
user@host> show loop-detect enhanced interface
```

```
Interface                  :ae60.0
Vlan-id                    :104
ESI                        :00:00:00:00:e1:06:06:06:06:00
Current status             :Loop-detected
                Remote Host        :host-A
                Remote Chassis     :aa:bb:cc:cc:ee:fa
                Remote Interface   :ae65.0
                Remote ESI         :00:00:00:00:e1:06:06:06:06:00
Last loop-detect time      :Wed Dec 23 23:20:26 2020
Receive statistics         :1624
```

```
Action configured          :Interface-down
Action count               :1
Transmit Interval          :1s
Revert Interval            :220s

Interface                  :mge-0/0/31.0
Vlan-id                    :101
ESI                        :00:00:00:00:00:00:00:00:00:00
Current status             :Loop-detected
                Remote Host        :host-B
                Remote Chassis     :aa:bb:cc:cc:ee:fb
                Remote Interface   :mge-0/0/32.0
                Remote ESI         :00:00:00:00:00:00:00:00:00:00
Last loop-detect time      :Wed Dec 23 23:20:25 2020
Receive statistics         :1625
Action configured          :Interface-down
Action count               :1
Transmit Interval          :1s
Revert Interval            :300s

Interface                  :mge-0/0/32.0
Vlan-id                    :101
ESI                        :00:00:00:00:00:00:00:00:00:00
Current status             :Loop-detected
                Remote Host        :host-B
                Remote Chassis     :aa:bb:cc:cc:ee:fb
                Remote Interface   :mge-0/0/31.0
                Remote ESI         :00:00:00:00:00:00:00:00:00:00
Last loop-detect time      :Wed Dec 23 23:20:25 2020
Receive statistics         :1625
Action configured          :Interface-down
Action count               :1
Transmit Interval          :1s
Revert Interval            :270s

Interface                  :mge-0/0/41.0
Vlan-id                    :105
ESI                        :00:00:00:00:00:00:00:00:00:00
Current status             :Loop-detected
                Remote Host        :host-A
                Remote Chassis     :aa:bb:cc:cc:ee:fa
                Remote Interface   :xe-0/0/6.0
                Remote ESI         :00:00:00:00:00:00:00:00:00:00
```

```
Last loop-detect time      :Wed Dec 23 23:20:24 2020
Receive statistics         :1626
Action configured          :Interface-down
Action count               :1
Transmit Interval          :1s
Revert Interval            :240s
```

show loop-detect enhanced interface (ACX7100-32C, ACX7100-48L, all interfaces configured with loop detection, no loop detected)

```
user@host> show loop-detect enhanced interface
Interface                  :xe-0/0/12.0
Vlan-id                    :100
ESI                        :00:00:00:00:00:00:00:00:00:00
Current status             :Normal[Link Up]
Last loop-detect time      :-
Receive statistics         :0
Action configured          :Interface-down
Action count               :0
Transmit Interval             :1s
Revert Interval          :60s
```

show loop-detect enhanced interface (ACX7100-32C, ACX7100-48L, all interfaces configured with loop detection, loop detected)

```
user@host> show loop-detect enhanced interface
Interface                  :xe-0/0/12.0
Vlan-id                    :100
ESI                        :00:00:00:00:00:00:00:00:00:00
Current status             :Loop-detected
    Remote Host        :leaf04
    Remote Chassis     :94:f7:ad:94:dd:40
    Remote Interface   :xe-0/0/12.0
    Remote ESI       :00:00:00:00:00:00:00:00:00:00
Last loop-detect time      :Tue May 26 04:36:37 2020
Receive statistics         :4
Action configured          :Interface-down
Action count               :1
```

```
Transmit Interval       :1s
Revert Interval           :60s
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

EVPN-VXLAN Lightweight Leaf to Server Loop Detection

loop-detect | **10113**

clear loop-detect enhanced interface | **3693**

show mac-rewrite interface

IN THIS SECTION

- Syntax | **21993**
- Description | **21994**
- Options | **21994**
- Required Privilege Level | **21994**
- Output Fields | **21994**
- Sample Output | **21995**
- Release Information | **21995**

Syntax

```
show mac-rewrite interface
<brief | detail>
<interface-name>
```

Description

Display Layer 2 protocol tunneling (L2PT) information.

> **NOTE**: The implementation of L2PT over VXLAN tunnels for EVPN-VXLAN networks doesn't do MAC rewrite operations for the protocol frame tunneling. To verify operation when you enable L2PT over VXLAN, use the "show l2pt interface" on page 21816 command instead of this command. See *Layer 2 Protocol Tunneling over VXLAN Tunnels in EVPN-VXLAN Bridged Overlay Networks* for details.

Options

brief | detail	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display L2PT information for the specified interface.

Required Privilege Level

view

Output Fields

Table 771 on page 21994 lists the output fields for the `show mac-rewrite interface` command. Output fields are listed in the approximate order in which they appear.

Table 771: show mac-rewrite interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface on which L2PT is configured.	brief detail
Tunnel destination MAC address	Name of the configured Layer 2 multicast MAC address.	brief detail

Table 771: show mac-rewrite interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Protocols	Layer 2 protocols being tunneled on this interface. All devices that support L2PT can tunnel the following protocols: Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP). The following Layer 2 protocols can also be tunneled on some devices that support L2PT: E-LMI, GVRP, IEEE 802.1X, IEEE 802.3AH, LACP, LLDP, MMRP, MVRP, PVSTP+, UDLD, or VSTP. See "protocol" on page 13424 for more information on the supported protocols for tunneling on different devices.	brief detail

Sample Output

show mac-rewrite interface

```
user@host> show mac-rewrite interface
Tunnel dest mac            : 01:00:0C:CD:CD:D0
Interface       Protocols

et-0/0/0       STP 8021X 8023AH ELMI LACP LLDP MMRP MVRP VTP CDP
```

show mac-rewrite interface (EX Series Switches)

```
user@switch> show mac-rewrite interface
Interface       Protocols

 ge-0/0/1       802.3AH LLDP STP
```

Release Information

Command introduced in Junos OS Release 9.1.

Command introduced in Junos OS Release 14.1X53-D10 and 17.3R1 for EX4300 switches.

Command introduced in Junos OS Release 15.1X53-D55 and 18.2R1 for EX2300 and EX3400 switches.

RELATED DOCUMENTATION

- layer2-control | 9641
- mac-rewrite | 10259
- protocol | 13424
- *Understanding Layer 2 Protocol Tunneling*
- *Configuring Layer 2 Protocol Tunneling*

show mac-vrf forwarding flood

IN THIS SECTION

- Syntax | 21996
- Description | 21997
- Options | 21997
- Additional Information | 21998
- Required Privilege Level | 21998
- Output Fields | 21998
- Sample Output | 21999
- Release Information | 22000

Syntax

```
show mac-vrf forwarding flood
<brief | detail | extensive>
<event-queue>
<instance instance-name>
<route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-flood)>
<vlan-name vlan-name>
```

Description

Show Ethernet-switching flooding information for the bridging or switching instances.

> **NOTE**: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

none
Display all Ethernet-switching flooding information for all VLANs.

brief | detail | extensive
(Optional) Display the specified level of output.

event-queue
(Optional) Display the queue of pending Ethernet-switching flood events.

instance *instance-name*
(Optional) Display Ethernet-switching flooding information for the specified routing instance.

route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-flood)
(Optional) Display the route for the following traffic types::

- `all-ce-flood`—Display the route for flooding traffic to all customer edge routers or switches if you have configured the `no-local-switching` option.
- `all-ve-flood`—Display the route for flooding traffic to all VPLS edge routers or switches if you have configured the `no-local-switching` option.
- `alt-root-flood`—Display the Spanning Tree Protocol (STP) alt-root flooding route used for the interface. The alt-root flooding route is used in spanning tree protocol (STP) when the root port is in discarding state.
- `bd-flood`—Display the route for flooding traffic of a VLAN if you have not configured the `no-local-switching` option.
- `mlp-flood`—Display the route for flooding traffic to MAC learning chips..
- `re-flood`—Display the route that the Routing Engine (RE) uses when flooding to all interfaces.

vlan-name *vlan-name*
(Optional) Display Ethernet-switching flooding information for the specified VLAN.

Additional Information

On MX Series routers and the EX9200 line of switches, this command hierarchy is an alias for the `show bridge flood` command hierarchy.

On QFX Series switches, this command hierarchy is an alias for the `show ethernet-switching flood` command hierarchy.

Required Privilege Level

view

Output Fields

"show mac-vrf forwarding flood" on page 21996 lists the output fields for the `show mac-vrf forwarding flood` command.

Table 772: show mac-vrf forwarding flood Output Fields

Field Name	Field Description
Name	Name of the switching instance
CEs	Number of customer edge routers or switches
VEs	Number of VPLS edge routers or switches
VLAN Name	Name of the associated VLAN in the switching instance (if any)
Flood Routes	Routes that get flooded to other destinations
Prefix	Prefix of the flooded route

Table 772: show mac-vrf forwarding flood Output Fields *(Continued)*

Field Name	Field Description
Owner	Owner of the flood route. The displayed values could be one of: • __all_ces__ • __all_ves__ • __alt_root__ • __bd_flood__ • __mlp_flood__ • __re_flood__
Type	Type of routes flooded
NhType	Type of next-hop receiver. The displayed values could be one of: • comp • ucast
NhIndex	Index number of next hop entry

Sample Output

show mac-vrf forwarding flood

```
user@host> show mac-vrf forwarding flood
```

```
Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
```

```
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  0x3057b/51 FLOOD_GRP_COMP_NH __all_ces__      comp            12866
  0x30004/51 FLOOD_GRP_COMP_NH __re_flood__     comp            12863
VLAN Name: VLAN102
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  0x3057c/51 FLOOD_GRP_COMP_NH __all_ces__      comp            12875
  0x30005/51 FLOOD_GRP_COMP_NH __re_flood__     comp            12872
VLAN Name: VLAN103
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  0x3057d/51 FLOOD_GRP_COMP_NH __all_ces__      comp            12884
  0x30006/51 FLOOD_GRP_COMP_NH __re_flood__     comp            12881
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show mac-vrf forwarding flood-group | 22000

show mac-vrf forwarding flood-group

IN THIS SECTION

- Syntax | 22001
- Description | 22001
- Options | 22001
- Additional Information | 22001
- Required Privilege Level | 22001
- Output Fields | 22002
- Sample Output | 22002
- Release Information | 22005

Syntax

```
show mac-vrf forwarding flood-group
<eamnh>
<instance instance name>
```

Description

Display flood-group information for all active MAC-VRF instances.

> **NOTE**: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

none Display flood-group information for the MAC-VRF instance.

eamnh	(Optional)
instance <*instance-name*> | bd-index <*bridge-domain-index-id*>	(Optional) Display flood-group information for the specified routing instance or for a specific bridge domain index ID.

Additional Information

On MX Series routers, this command is an alias for the `show l2 learning flood-group` command.

On QFX Series switches, this command is an alias for the `show ethernet-switching flood-group` command.

Required Privilege Level

view

Output Fields

Table 773: show mac-vrf forwarding flood-group Output Fields

Field Name	Field Description
Bd Name	Bridge domain name
Eamnh	Enhanced mesh group and next hop information
Fgrp Flag	Flood group flag
Flags	Number of flags
FldGrp Name	Flood group name
Flood Route	Route identification in hexadecimal
Flood Token	Token identification in hexadecimal
Fnh id	Forwarding next-hop id.
id	EAMNH id
MeshGrpName	Name of the mesh group

Sample Output

show mac-vrf forwarding flood-group

```
user@host> show mac-vrf forwarding flood-group
```

```
FldGrp Name : __ves__+++2 MeshGrp Name: __ves__ Bd Name: v150+150 Fgrp Flag: 0xc
Eamnh: 0x94aae88    id: 1801      Flags: 0
Flood Route : 0x94aae08 Flood Token : 0x30002/51 Fnh id      : 1803

FldGrp Name : __all_ces__+++2 MeshGrp Name: __all_ces__ Bd Name: v150+150 Fgrp Flag: 0xa
Eamnh: 0x94aaf08    id: 1781      Flags: 0
Flood Route : 0x94aaec8 Flood Token : 0x30003/51 Fnh id      : 1783
```

```
FldGrp Name : __mlp_flood__+++2 MeshGrp Name: __mlp_flood__ Bd Name: v150+150 Fgrp Flag: 0x40c
Eamnh: 0x956f348    id: 0         Flags: 1
Flood Route : 0x94aaf48 Flood Token : 0x20002/51 Fnh id      : 0

FldGrp Name : __re_flood__+++2 MeshGrp Name: __re_flood__ Bd Name: v150+150 Fgrp Flag: 0x4c
Eamnh: 0x956f608    id: 0         Flags: 1
Flood Route : 0x956f5c8 Flood Token : 0x30000/51 Fnh id      : 1788

FldGrp Name : __ves__+++3 MeshGrp Name: __ves__ Bd Name: v250+250 Fgrp Flag: 0xc
Eamnh: 0x956f808    id: 1800      Flags: 0
Flood Route : 0x956f7c8 Flood Token : 0x30004/51 Fnh id      : 1802

FldGrp Name : __all_ces__+++3 MeshGrp Name: __all_ces__ Bd Name: v250+250 Fgrp Flag: 0xa
Eamnh: 0x956f888    id: 1782      Flags: 0
Flood Route : 0x956f848 Flood Token : 0x30005/51 Fnh id      : 1789

FldGrp Name : __mlp_flood__+++3 MeshGrp Name: __mlp_flood__ Bd Name: v250+250 Fgrp Flag: 0x40c
Eamnh: 0x956f908    id: 0         Flags: 1
Flood Route : 0x956f8c8 Flood Token : 0x20003/51 Fnh id      : 0

FldGrp Name : __re_flood__+++3 MeshGrp Name: __re_flood__ Bd Name: v250+250 Fgrp Flag: 0x4c
Eamnh: 0x956f988    id: 0         Flags: 1
Flood Route : 0x956f948 Flood Token : 0x30001/51 Fnh id      : 1790
```

```
user@host> show mac-vrf forwarding flood-group eamnh
```

```
Routing Instance Name : macvrf_v150
Bridge Domain Name   : v150+150
Mesh Group Name   : __ves__
NH Index   : 1801
IF NH tree:

Component         Comp       Index     Gen      NH/Ucst   Flood
Name              Type                 Number   INDEX     Ready
vtep.32771        L2_IFF     576       180       1793     True
Routing Instance Name : macvrf_v150
Bridge Domain Name   : v150+150
Mesh Group Name   : __all_ces__
NH Index   : 1781
```

```
IF NH tree:

Component         Comp       Index     Gen      NH/Ucst   Flood
Name              Type                 Number   INDEX     Ready
ae0.0             L2_IFF     575       176       1787     True
Routing Instance Name : macvrf_v150
Bridge Domain Name   : v150+150
Mesh Group Name   : __mlp_flood__
NH Index   : 0
Routing Instance Name : macvrf_v150
Bridge Domain Name   : v150+150
Mesh Group Name   : __re_flood__
NH Index   : 0
Routing Instance Name : macvrf_v150
Bridge Domain Name   : v250+250
Mesh Group Name   : __ves__
NH Index   : 1800
IF NH tree:

Component         Comp       Index     Gen      NH/Ucst   Flood
Name              Type                 Number   INDEX     Ready
vtep.32771        L2_IFF     576       180       1793     True
Routing Instance Name : macvrf_v150
Bridge Domain Name   : v250+250
Mesh Group Name   : __all_ces__
NH Index   : 1782
IF NH tree:

Component         Comp       Index     Gen      NH/Ucst   Flood
Name              Type                 Number   INDEX     Ready
ae0.0             L2_IFF     575       176       1787     True
Routing Instance Name : macvrf_v150
Bridge Domain Name   : v250+250
Mesh Group Name   : __mlp_flood__
NH Index   : 0
Routing Instance Name : macvrf_v150
Bridge Domain Name   : v250+250
Mesh Group Name   : __re_flood__
NH Index   : 0
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show mac-vrf forwarding flood | 21996

clear mac-vrf forwarding mac-table | 3698

show mac-vrf forwarding global-information

IN THIS SECTION

- Syntax | 22005
- Description | 22005
- Options | 22006
- Additional Information | 22006
- Required Privilege Level | 22006
- Output Fields | 22006
- Sample Output | 22007
- Release Information | 22008

Syntax

```
show mac-vrf forwarding global-information
```

Description

Display global process and configuration information about the MAC-VRF subsystems.

> **NOTE**: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

This command has no options.

Additional Information

On MX Series routers and the EX9200 line of switches, this command is an alias for the `show l2-learning global-information` command.

On QFX Series switches, this command is an alias for the `show ethernet-switching global-information` command.

Required Privilege Level

view

Output Fields

Table 774: show mac-vrf forwarding global-information Output Fields

Field Name	Field Description
MAC aging interval	How fast learned MAC address entries age out of the table (in seconds)
MAC learning	Allow MAC learning—enabled or disabled
MAC statistics	Keep MAC statistics—enabled or disabled
MAC limit count	Limit MAC entries to this number
MAC limit hit	Warn when MAC limit is hit—enabled or disabled
MAC packet action drop	Drop MAC packets—enabled or disabled

Table 774: show mac-vrf forwarding global-information Output Fields *(Continued)*

Field Name	Field Description
MAC+IP aging interval	Time to keep old MAC+IP entries—in seconds. One timer for IPv4 entries and one timer for IPv6 entries.
MAC+IP limit count	Total MAC+IP addresses supported on the platform.
MAC+IP limit reached	Has your system reached the MAC+IP limit?—Yes or No
LE aging time	Time to keep LE entries.
LE VLAN aging time	Time to keep LE VLAN entries
RE state	State of the Routing Engine—primary or backup
VXLAN Overlay load balance	VXLAN overlay load balancing—enabled or disabled
VXLAN ECMP	VXLAN ECMP—enabled or disabled

Sample Output

show mac-vrf forwarding global-information

```
user@host> show mac-vrf forwarding global-information
```

```
Global Configuration:

MAC aging interval    : 300
MAC learning          : Enabled
MAC statistics        : Disabled
MAC limit Count       : 393215
MAC limit hit         : Disabled
MAC packet action drop: Disabled
MAC+IP aging interval : IPv4 - 1200 seconds
                        IPv6 - 1200 seconds
MAC+IP limit Count    : 393215
MAC+IP limit reached  : No
LE  aging time        : 1200
```

```
LE  VLAN aging time   : 1200
RE state              : Master
VXLAN Overlay load bal: Disabled
VXLAN ECMP            : Disabled
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show mac-vrf forwarding global-mac-count | 22008

show mac-vrf forwarding global-mac-ip-count | 22010

show mac-vrf forwarding global-mac-count

IN THIS SECTION

- Syntax | 22008
- Description | 22009
- Options | 22009
- Additional Information | 22009
- Required Privilege Level | 22009
- Sample Output | 22009
- Release Information | 22009

Syntax

```
show mac-vrf forwarding global-mac-count
```

Description

Display the number of globally-learned Layer 2 MAC addresses.

> **NOTE**: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

No options exist for this command.

Additional Information

On MX Series routers, this command is an alias for `show l2-learning global-mac-count`.

On QFX Series switches, this command is an alias for `show ethernet-switching global-mac-count`.

Required Privilege Level

view

Sample Output

show mac-vrf forwarding global-mac-count

```
user@host> show mac-vrf forwarding global-mac-count
```

```
1032 dynamic and static MAC addresses learned globally
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show mac-vrf forwarding flood | 21996

show mac-vrf forwarding global-information | 22005

show mac-vrf forwarding global-mac-ip-count | 22010

show mac-vrf forwarding global-mac-ip-count

IN THIS SECTION

- Syntax | 22010
- Description | 22010
- Options | 22010
- Additional Information | 22011
- Required Privilege Level | 22011
- Sample Output | 22011
- Release Information | 22011

Syntax

```
show mac-vrf forwarding global-mac-ip-count
```

Description

Display the number of globally-learned MAC address to IP address binding count. The first column of output shows the number of MAC addresses that the device can learn. The second column shows the number of IP addresses that are bound to learned MAC addresses.

Options

No options exist for this command.

Additional Information

On MX Series routers, this command is an alias for `show l2-learning global-mac-ip-count`.

On QFX Series switches, this command is an alias for `show ethernet-switching global-mac-ip-count`.

Required Privilege Level

view

Sample Output

show mac-vrf forwarding global-mac-ip-count

```
user@r1_dc1_spine1> show mac-vrf forwarding global-mac-ip-count
```

```
735 MAC+IP bindings learned globally
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show mac-vrf forwarding global-information | 22005

show mac-vrf forwarding global-mac-count | 22008

show mac-vrf forwarding instance

IN THIS SECTION

- Syntax | 22012
- Description | 22012

- Options | 22012
- Additional Information | 22013
- Required Privilege Level | 22013
- Output Fields | 22013
- Sample Output | 22014
- Release Information | 22017

Syntax

```
show mac-vrf forwarding instance
<brief|detail|extensive>
<vlan-name vlan name>
```

Description

Display the MAC-IP addresses for all IPv4 (ARP) and IPv6 (ND) bindings for VLANs in routing instances where the instance type is MAC-VRF.

> **NOTE**: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

none	Display the brief listing of MAC-IP addresses for all IPv4 (ARP) and IPv6 (ND) bindings for VLANs in routing instances where the instance type is MAC-VRF.
brief | detail | extensive	(Optional) Display the specified level of output.
vlan-name *vlan-name*	(Optional) Display information about the specified VLAN only.

Additional Information

On MX Series routers and the EX9200 line of switches, this command hierarchy is an alias for the `show l2-learning instance` command hierarchy.

On QFX Series switches, this command hierarchy is an alias for the `show ethernet-switching instance` command hierarchy.

Required Privilege Level

view

Output Fields

Table 775: show mac-vrf forwarding instance Output Fields

Field Name	Field Description
Instance type	Routing instance type
Logical system	Name of the logical system associated with the routing instance
Routing instance	Name of the routing instance
VLAN name	Name of the associated VLAN
Index	Routing instance index number
IRB index	Integrated routing and bridging (IRB) index number (if any)
Flags	Flags for the routing instance: **DL** disable learning **SE** statistics enabled **AD** packet action drop **LH** MAC limit hit **MI** MAC+IP limit hit
Tag	VLAN tags (if any)

Sample Output

show mac-vrf forwarding instance

```
user@host> show mac-vrf forwarding instance
Information for routing instance and VLAN:

Flags (DL - disable learning, SE - stats enabled,
       AD - packet action drop, LH - MAC limit hit,
       MI - mac+ip limit hit)

Inst Logical   Routing             VLAN name          Index IRB   Flags Tag
type system    instance                                      index
RTT  Default   __juniper_private1__                   1
RTT  Default   default-switch                         4
RTT  Default   default_internal_6                     5
RTT  Default   macvrf_v150                            6
vlan Default   macvrf_v150         v150               2     553         150
vlan Default   macvrf_v150         v250               3     554         250
```

show mac-vrf forwarding instance detail

```
user@host show mac-vrf forwarding instance detailInformation for routing instance and VLAN:
Routing instance : __juniper_private1__
   RTB index: 1
   MAC limit: 5000                 MACs learned: 0          Local Macs learned: 0
   Static MACs learned: 0          Non config static MACs learned: 0
   MAC+IP limit: 0                 MAC+IP bindings learned: 0        Local MAC+IP learned:0
   Sequence number: 0              Handle: 0x931a008
   VLAN id:
   Flags: Bridge instance,RTT is present in all PFE.
   Config service id: 0            Active service id: 0
   Config VLAN Id          :              Config operation: none
   Config params: mac tbl sz: 0, mac age: 300000, intf mac limit: 8192,
   Config MAC+IP params: mac+ip tbl sz: 0, intf mac+ip limit: 8192,
   Config static MAC count :
 Counters:
   Kernel write errors        : 0

Information for routing instance and VLAN:
```

```
Routing instance : default-switch
   RTB index: 4
   MAC limit: 0                    MACs learned: 0          Local Macs learned: 0
   Static MACs learned: 0          Non config static MACs learned: 0
   MAC+IP limit: 0                 MAC+IP bindings learned: 0        Local MAC+IP learned:0
   Sequence number: 0              Handle: 0x9317008
   VLAN id:
   Flags: Bridge instance,RTT defined,Default RTT,RTT is present in all PFE.
   Config service id: 0            Active service id: 0
   Config VLAN Id          :               Config operation: none
   Config params: mac tbl sz: 0, mac age: 300000, intf mac limit: 8192,
   Config MAC+IP params: mac+ip tbl sz: 0, intf mac+ip limit: 8192,
   Config static MAC count :
 Counters:
   Kernel write errors        : 0

Information for routing instance and VLAN:
Routing instance : default_internal_6
   RTB index: 5
   MAC limit: 5000                 MACs learned: 0          Local Macs learned: 0
   Static MACs learned: 0          Non config static MACs learned: 0
   MAC+IP limit: 0                 MAC+IP bindings learned: 0        Local MAC+IP learned:0
   Sequence number: 0              Handle: 0x9570008
   VLAN id:
   Flags: Bridge instance,RTT defined,RTT is present in all PFE.
   Config service id: 0            Active service id: 0
   Config VLAN Id          :               Config operation: none
   Config params: mac tbl sz: 0, mac age: 300000, intf mac limit: 8192,
   Config MAC+IP params: mac+ip tbl sz: 0, intf mac+ip limit: 8192,
   Config flags: mac-vrf fwd-inst
   Config static MAC count :
 Counters:
   Kernel write errors        : 0

Information for routing instance and VLAN:
Routing instance : macvrf_v150
   RTB index: 6
   MAC limit: 0                    MACs learned: 6          Local Macs learned: 0
   Static MACs learned: 0          Non config static MACs learned: 0
   MAC+IP limit: 0                 MAC+IP bindings learned: 7        Local MAC+IP learned:10
   Sequence number: 0              Handle: 0x956d008
   VLAN id:
   Flags: Bridge instance,RTT defined,vpls RTT,RTT is present in all PFE.,EVPN VXLAN
```

```
    Forwarding instance: default_internal_6    (Index: 5)
    Config service id: 0            Active service id: 0
    Config VLAN Id          :              Config operation: none
    Config params: mac tbl sz: 0, mac age: 300000, intf mac limit: 8192,
    Config MAC+IP params: mac+ip tbl sz: 0, intf mac+ip limit: 8192,
    Config flags: vpls,irb ifl,svtep ifl,evpn vxlan,mac-vrf
    Service type: vlan-aware
    Config static MAC count :
  Counters:
    Kernel write errors        : 0

Information for routing instance and VLAN:
Routing instance : macvrf_v150
  VLAN Name : v150
    RTB index: 6                    VLAN index: 2
    MAC limit: 5120                 MACs learned: 3          Local Macs learned: 0
    MAC+IP limit: 5120              MAC+IP bindings learned: 3
    Sequence number: 0              Handle: 0x957e008
    VLAN id: 150
    Flags: BD defined,VxLAN,Ageless CPMAC,Recd MAC SYNC,Sent MAC SYNC,ARP/NDP flood suppression
    Config VLAN Id          : 150          Config operation: none
    Config params: mac tbl sz: 5120, mac age: 300000, intf mac limit: 1024,
    Config MAC+IP params: mac+ip tbl sz: 5120, intf mac+ip limit: 1024,
    Config flags: vlan,VxLAN,Ageless CPMAC,no st mac cfg
    Config static MAC count : 0
    Config ownership flags: config
    Config RG Id: 0                        Active RG Id: 0
    Config service id: 0            Active service id: 0
    Config VXLAN   Id: 150                   Active VXLAN   Id: 150
    Config VXLAN MCIP: 0.0.0.0                Active VXLAN MCIP: 0.0.0.0
    Config VXLAN Decapsulate accept inner vlan: 0
    Config VXLAN Encapsulate inner vlan       : 0
    Config VXLAN Unreachable RVTEP age timer  : 300
    Config VXLAN OVSDB Enabled: No
    SVTEP L3-RTT ID:   0                      Source-VTEP IP: 10.255.1.1
    Kernel ownership flags: config
    MVRP ref count: 0
  Counters:
    Kernel write errors        : 0

Information for routing instance and VLAN:
Routing instance : macvrf_v150
  VLAN Name : v250
```

```
    RTB index: 6                    VLAN index: 3
    MAC limit: 5120                 MACs learned: 3          Local Macs learned: 0
    MAC+IP limit: 5120              MAC+IP bindings learned: 4
    Sequence number: 0              Handle: 0x9580008
    VLAN id: 250
    Flags: BD defined,VxLAN,Ageless CPMAC,Recd MAC SYNC,Sent MAC SYNC,ARP/NDP flood suppression
    Config VLAN Id          : 250          Config operation: none
    Config params: mac tbl sz: 5120, mac age: 300000, intf mac limit: 1024,
    Config MAC+IP params: mac+ip tbl sz: 5120, intf mac+ip limit: 1024,
    Config flags: vlan,VxLAN,Ageless CPMAC,no st mac cfg
    Config static MAC count : 0
    Config ownership flags: config
    Config RG Id: 0                        Active RG Id: 0
    Config service id: 0            Active service id: 0
    Config VXLAN   Id: 250                   Active VXLAN   Id: 250
    Config VXLAN MCIP: 0.0.0.0                Active VXLAN MCIP: 0.0.0.0
    Config VXLAN Decapsulate accept inner vlan: 0
    Config VXLAN Encapsulate inner vlan       : 0
    Config VXLAN Unreachable RVTEP age timer  : 300
    Config VXLAN OVSDB Enabled: No
    SVTEP L3-RTT ID:   0                      Source-VTEP IP: 10.255.1.1
    Kernel ownership flags: config
    MVRP ref count: 0
  Counters:
    Kernel write errors        : 0
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show mac-vrf forwarding global-information | 22005

show mac-vrf forwarding instance-mapping

IN THIS SECTION

- Syntax | 22018
- Description | 22018
- Options | 22018
- Required Privilege Level | 22018
- Output Fields | 22019
- Sample Output | 22019
- Release Information | 22020

Syntax

```
show mac-vrf forwarding instance-mapping
<brief|detail>
<forwarding-instance-id forwarding instance id number>
```

Description

Display information about the mapping of configured MAC-VRF instances to forwarding instances.

Options

none	Display brief information about the mapping of the global MAC-VRF instance to the forwarding instance.
brief | detail	(Optional) Display the specified level of detail.
forwarding-instance-id *forwarding instance id number*	(Optional) Display the instances mapped to a specific forwarding instance id.

Required Privilege Level

view

Output Fields

Table 776: show mac-vrf forwarding instance-mapping Output Fields

Field Name	Field Description
Forwarding-instance-id	Configuration ID for the MAC-VRF routing instance—Id 0 is reserved for the default-switch instance. The output displays other ID numbers only if you've configured the IDs using the `set routing-instances` *`mac-vrf-instance-name`* `forwarding-instance identifier` *`id-number`* command.
Forwarding-instance	Forwarding-instance name—The instance name can be either `default-switch` (if you have not configured a forwarding-instance identifier) or the name of the configured forwarding instance.
Table-id	Table ID for the forwarding instance—Unique, kernel-allocated table ID for the forwarding instance.
Mac-vrf instance	Name of the MAC-VRF instance that is mapped to the forwarding instance.
Table-id	Unique, kernel-allocated table ID for the MAC-VRF instance.

Sample Output

show mac-vrf forwarding instance-mapping

```
user@host> show mac-vrf forwarding instance-mapping
```

```
Forwarding-instance-id     Forwarding instance        Table-id     Mac-vrf instance      Table-id

    0                     default-switch              6
                                                                 macvrf_v500          13
    6                     default_internal_6          7
                                                                 macvrf_v100          8
                                                                 macvrf_v200          9
    60                    default_internal_60         10
                                                                 macvrf_v300          11
```

```
                        vs400                       14
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show mac-vrf forwarding instance | 22011

show mac-vrf forwarding mac-table | 22031

show mac-vrf forwarding interface

IN THIS SECTION

- Syntax | 22020
- Description | 22021
- Options | 22021
- Additional Information | 22021
- Required Privilege Level | 22021
- Output Fields | 22021
- Sample Output | 22023
- Release Information | 22025

Syntax

```
show mac-vrf forwarding interface
<brief | detail | extensive>
<interface-name>
```

Description

Show Ethernet switching information about all the interfaces within all active MAC-VRF routing instances.

Options

none	Show brief Ethernet switching information about all the interfaces within all MAC-VRF routing instances.
`brief` | `detail` | `extensive`	(Optional) Show the specified level of output.
`interface-name`	(Optional) Show the Ethernet switching information for the specified interface only.

> **NOTE**: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Additional Information

On MX Series routers and the EX9200 line of switches, this command hierarchy is an alias for the `show l2-learning instance` command hierarchy.

On QFX Series switches, this command hierarchy is an alias for the `show ethernet-switching interface` command hierarchy.

Required Privilege Level

view

Output Fields

Table 777: show mac-vrf forwarding interface Output Fields

Field Name	Field Description
Routing instance name	Name of the mac-vrf routing instance.

Table 777: show mac-vrf forwarding interface Output Fields *(Continued)*

Field Name	Field Description	
Logical interface	Logical interface name associated with the mac-vrf routing instance	
VLAN members	For aggregated Ethernet (ae) interfaces, this lists the member VLANs for the displayed interface.	
TAG	VLAN tag associated with VLAN member interfaces.	
MAC limit	Maximum number of MAC address entries allowed for each interface. The MAC limit for each member interface must add up to less than the total MAC limit for the aggregated Ethernet interface.	
MAC+IP limit	Maximum number of MAC+IP address entries allowed for each interface. The MAC+IP limit for each member interface must add up to less than the total MAC+IP limit for the aggregated Ethernet interface.	
STP state	The spanning tree protocol state for the displayed VLAN member interfaces.	
Logical interface flags	**DL**	disable learning
	AD	packet action drop
	LH	MAC limit hit
	DN	interface down
	MMAS	MAC move action shutdown
	AS	Autostate-exclude enabled
	SCTL	shutdown by storm control
	MI	MAC+IP limit hit
Tagging	Display shows tagged, or blank depending on whether the interface is tagged or not.	

Sample Output

show mac-vrf forwarding interface

```
user@host> show mac-vrf forwarding interface
```

```
Routing Instance Name : macvrf_v150
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown,  AS - Autostate-exclude enabled,
                         SCTL - shutdown by Storm-control, MI - MAC+IP limit hit)

Logical         Vlan                  TAG   MAC    MAC+IP STP         Logical          Tagging
interface       members                     limit  limit  state       interface flags
vtep.32770                                  0      0                                   tagged
Routing Instance Name : macvrf_v150
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown,  AS - Autostate-exclude enabled,
                         SCTL - shutdown by Storm-control, MI - MAC+IP limit hit)

Logical         Vlan                  TAG   MAC    MAC+IP STP         Logical          Tagging
interface       members                     limit  limit  state       interface flags
vtep.32769                                  0      0                                   tagged
Routing Instance Name : macvrf_v150
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown,  AS - Autostate-exclude enabled,
                         SCTL - shutdown by Storm-control, MI - MAC+IP limit hit)

Logical         Vlan                  TAG   MAC    MAC+IP STP         Logical          Tagging
interface       members                     limit  limit  state       interface flags
ae0.0                                       8192   8192                                tagged
                v150                  150   1024   1024   Forwarding                   tagged
                v250                  250   1024   1024   Forwarding                   tagged
Routing Instance Name : macvrf_v150
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown,  AS - Autostate-exclude enabled,
                         SCTL - shutdown by Storm-control, MI - MAC+IP limit hit)
```

```
Logical         Vlan                  TAG   MAC    MAC+IP STP         Logical          Tagging
interface       members                     limit  limit  state       interface flags
vtep.32771                                  0      0                                   tagged
                v250                  250   0      0      Forwarding                   tagged
                v150                  150   0      0      Forwarding                   tagged
Routing Instance Name : macvrf_v150
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown,  AS - Autostate-exclude enabled,
                         SCTL - shutdown by Storm-control, MI - MAC+IP limit hit)

Logical         Vlan                  TAG   MAC    MAC+IP STP         Logical          Tagging
interface       members                     limit  limit  state       interface flags
esi.1786                                    0      0                                   tagged
                v150                  150   0      0      Forwarding                   tagged
Routing Instance Name : macvrf_v150
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown,  AS - Autostate-exclude enabled,
                         SCTL - shutdown by Storm-control, MI - MAC+IP limit hit)

Logical         Vlan                  TAG   MAC    MAC+IP STP         Logical          Tagging
interface       members                     limit  limit  state       interface flags
esi.1794                                    0      0                                   tagged
                v250                  250   0      0      Forwarding                   tagged
Routing Instance Name : macvrf_v150
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown,  AS - Autostate-exclude enabled,
                         SCTL - shutdown by Storm-control, MI - MAC+IP limit hit)

Logical         Vlan                  TAG   MAC    MAC+IP STP         Logical          Tagging
interface       members                     limit  limit  state       interface flags
esi.1804                                    0      0                                   tagged
Routing Instance Name : macvrf_v150
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown,  AS - Autostate-exclude enabled,
                         SCTL - shutdown by Storm-control, MI - MAC+IP limit hit)

Logical         Vlan                  TAG   MAC    MAC+IP STP         Logical          Tagging
interface       members                     limit  limit  state       interface flags
```

```
esi.1805                                    0      0                                  tagged
Paste device command output here
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

- show mac-vrf forwarding instance | 22011
- show ethernet-switching mac-ip-table | 19660

show mac-vrf forwarding mac-ip-table

IN THIS SECTION

- Syntax | 22025
- Description | 22026
- Options | 22026
- Additional Information | 22026
- Required Privilege Level | 22027
- Output Fields | 22027
- Sample Output | 22029
- Release Information | 22031

Syntax

```
show mac-vrf forwarding mac-ip-table
<brief | count | detail | extensive>
<age>
<instance>
<ipv4>
```

```
<ipv6>
<kernel >
<vlan-name (all | vlan-name)>
```

Description

Display the MAC-IP address for all IPv4 (ARP) and IPv6 (ND) bindings for VLANs in MAC-VRF type routing instances.

> **NOTE**: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific subcommand or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

none	Display brief information about the MAC-IP table.
brief | count | detail | extensive	(Optional) Display the specified level of output.
age	(Optional) (For MX Series routers and the EX9200 line of switches only) Display the age of a single MAC address.
instance *instance-name*	(Optional) Display MAC-IP table information for the specified routing instance.
ipv4	(Optional) Display IPv4 address entries associated with MAC addresses learned by ARP.
ipv6	(Optional) Display IPv6 address entries associated with MAC addresses learned by NDP.
vlan-name (all | *vlan name*)	(Optional) Display MAC-IP table information for the specified VLAN.

Additional Information

On MX Series routers and the EX9200 line of switches, this command hierarchy is an alias for the `show bridge mac-ip-table` command hierarchy.

On QFX Series switches, this command hierarchy is an alias for the `show ethernet-switching mac-ip-table` command hierarchy.

Required Privilege Level

view

Output Fields

Table 778: show mac-vrf forwarding mac-ip-table Output Fields

Field Name	Field Description	Level of Output
IP address	IP address associated with the EVPN routing instance.	All levels
MAC address	MAC address	All levels

Table 778: show mac-vrf forwarding mac-ip-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	MAC IP flags—Identifies statically installed MAC-IP entries. • S—Statically installed MAC-IP address entries. • D—Dynamic installed MAC-IP address entries. • L—Locally learned MAC-IP address entries • R— MAC-IP address entries that are learned from a remote device through the control plane. • K—MAC-IP address entries that are installed in the kernel. • RT—Destination route that corresponds to an installed entry. • AD—Advertised to remote device. • RE—ARP/ND entry has expired and another ARP or network solicitation request has been sent. • R0—Router • OV—Override	Brief, detail
Logical Interface	The logical interface associated with the routing instance.	Brief, detail

Table 778: show mac-vrf forwarding mac-ip-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Active source	The source of the learned MAC-IP address entry. Displays the ESI or router ID.	All levels
Routing instance	Information about the routing instance where the MAC-IP address entry was learned. • Bridging domain • MAC-IP local mask • MAC-IP remote mask • MAC-IP RPD flags • MAC-IP flags	Extensive
BD ID	Bridge domain ID	Brief

Sample Output

show mac-vrf forwarding mac-ip-table vlan-name v100

```
user@host> show mac-vrf forwarding mac-ip-table vlan-name v100
```

```
MAC IP flags  (S - Static, D - Dynamic, L - Local , R - Remote, K - Kernel, RT - Dest Route,
               AD - Advt to remote, RE - Re-ARP/ND, RO - Router, OV - Override)
 Routing instance : default-
 Bridging domain : v100
   IP                          MAC                 Flags            Logical            Active
   address                     address                              Interface          source
   10.1.1.1                    00:11:00:00:00:01   DR,K,RT          vtep.32769
101.1.1.1
   2001::192:100:1:1           00:00:10:00:11:00   DR,K,RT          vtep.32769
```

```
101.1.1.1
   fe80::205:8600:6471:e600     00:00:10:00:11:00    DR,K,RT           vtep.32769
101.1.1.1
   10.0.1.2                     00:00:10:00:22:00    S,K
   2001::192:100:1:2            00:00:10:00:22:00    S,K
   fe80::205:8600:6471:d900     00:00:10:00:22:00    S,K
   10.0.1.3                     00:00:10:00:33:00    DR,K,RT           vtep.32771
103.1.1.1
   2001::192:100:1:3            00:00:10:00:33:00    DR,K,RT           vtep.32771
103.1.1.1
   fe80::205:8600:6471:7400     00:00:10:00:33:00    DR,K,RT           vtep.32771
103.1.1.1
   10.0.1.4                     00:00:10:00:44:00    DR,K,RT           vtep.32770
104.1.1.1
   2001::192:100:1:4            00:00:10:00:44:00    DR,K,RT           vtep.32770
104.1.1.1
   fe80::205:8600:6471:c000     00:00:10:00:44:00    DR,K,RT           vtep.32770
104.1.1.1
   10.0.1.254                   40:00:10:11:11:00    S,K               esi.1874
05:00:00:00:64:00:00:00:64:00
   2001::192:100:1:254          60:00:10:11:11:00    S,K               esi.1874
05:00:00:00:64:00:00:00:64:00
Paste device command output here
```

show mac-vrf forwarding mac-ip-table extensive

```
IP address: 10.1.1.1
  MAC address: 00:11:00:00:00:01
  Routing instance: default
   Bridging domain: v100
   MAC-IP local mask: 0x0000000000000000
   MAC-IP remote mask: 0x8000000000000000
   MAC-IP RPD flags: 0x00000081
   MAC-IP flags: remote,kernel
```

show mac-vrf forwarding mac-ip-table kernel

```
show mac-vrf forwarding mac-ip-table kernel
    BD       IP                            MAC                  Logical           Flags
```

```
id        address                    address            Interface
2         192.168.1.1                00:21:59:aa:77:f0  irb.1            0x00000209
2         192.168.1.100              00:00:5e:00:01:01  irb.1            0x00000609
3         192.168.2.1                00:21:59:aa:77:f0  irb.2            0x00000209
3         192.168.2.100              00:00:5e:00:01:01  irb.2            0x00000609
4         192.168.3.1                00:21:59:aa:77:f0  irb.3            0x00000209
4         192.168.3.100              00:00:5e:00:01:01  irb.3            0x00000609
5         192.168.4.1                00:21:59:aa:77:f0  irb.4            0x00000209
5         192.168.4.100              00:00:5e:00:01:01  irb.4            0x00000609
```

Release Information

Command introduced in Junos OS Release 20.4R1.

show mac-vrf forwarding mac-table

IN THIS SECTION

- Syntax | 22031
- Description | 22032
- Options | 22032
- Additional Information | 22033
- Required Privilege Level | 22033
- Output Fields | 22033
- Sample Output | 22035
- Release Information | 22035

Syntax

```
show mac-vrf forwarding mac-table
<brief | count | detail | extensive>
<mac address>
<age>
<instance vrf-instance-name>
```

```
<vlan-name vlan-name>
<vlan-id vlan-id>
<interface interface-name>
```

Description

Display information about the forwarding table in the `mac-vrf` instance. The commands work only within a `mac-vrf` type routing instance that is present on a supported platform.

> **NOTE**: We've implemented the `show mac-vrf` commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

none	(Same as brief) Display information about the MAC-VRF forwarding MAC table.
brief | count | detail | extensive	(Optional) Display the specified level of output.
mac-address	(Optional) Display forwarding table information for the specified MAC address.
age	(Optional) For MX Series routers only. Display the age of a single MAC address.
instance *instance-name*	(Optional) Display MAC table information for the specified routing instance.
interface *interface-name*	(Optional) Display MAC table information for the specified interface.
vlan-name *vlan name*	(Optional) Display MAC table information for the specified VLAN.
vlan-id *vlan id*	(Optional) Display MAC table information for the specified VLAN.

Additional Information

On MX Series routers and the EX9200 line of switches, this command hierarchy is an alias for the `show bridge mac-table` command hierarchy. All the command options available in this command hierarchy are aliases to the respective command options in the `show bridge mac-table` command hierarchy.

On QFX Series switches, this command hierarchy is an alias for the `show ethernet-switching table` command hierarchy. All the command options available in this command hierarchy are aliases to the respective command options in the `show ethernet-switching table` command hierarchy.

Required Privilege Level

view

Output Fields

Table 779 on page 22033 lists the output fields for the `show mac-vrf forwarding mac-table` command.

Table 779: show mac-vrf forwarding mac-table Output Fields

Field Name	Description
VLAN name	Name of the associated VLAN.
MAC address	Learned MAC address

Table 779: show mac-vrf forwarding mac-table Output Fields *(Continued)*

Field Name	Description
MAC flags	Flags that provide information about the MAC addresses in the MAC-VRF instance. • S—Statically installed MAC-IP address entries. • D—Dynamic installed MAC-IP address entries. • L—Locally learned MAC-IP address entries • R— MAC-IP address entries that are learned from a remote device through the control plane. • K—Kernel-based MAC-IP address entries. • RT—Destination route that corresponds to an installed entry. • AD—Advertised to remote device. • RE—ARP/ND entry has expired and another ARP or network solicitation request has been sent. • R0—Router • OV—Override
Logical interface	Interface on which the MAC address was learned.
SVLBNH/VENH index	Shared VXLAN load balancing next hop (SVLBNH)/ VXLAN encapsulated next hop (VENH) associated with the SVLBNH.
Active source	IP address or MAC address of the source interface.
Routing instance	Name of the routing instance with which this MAC table is associated.

Sample Output

show mac-vrf forwarding mac-table

```
root@vQFX-Leaf1> show mac-vrf forwarding mac-table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 6 entries, 6 learned
Routing instance : macvrf_v150
   Vlan                MAC                 MAC      Logical                SVLBNH/      Active
   name                address             flags    interface              VENH Index   source
   v150                02:05:86:71:bd:00   DR       vtep.32771
10.255.1.2
   v150                02:05:86:71:d1:00   DR       vtep.32769
10.255.1.1
   v150                40:00:10:11:11:00   DR       esi.1787
05:00:00:fc:4d:00:00:00:96:00
   v150                40:00:10:11:11:02   DR       esi.1787
05:00:00:fc:4d:00:00:00:96:00
   v150                50:00:10:11:11:00   DR       esi.1787
05:00:00:fc:4d:00:00:00:96:00
   v150                50:00:10:11:11:02   DR       esi.1787
05:00:00:fc:4d:00:00:00:96:00
```

Release Information

Command introduced in Junos OS Release 20.4R1.

show mac-vrf forwarding mgrp-policy

IN THIS SECTION

Syntax | 22036

- Description | 22036
- Options | 22037
- Additional Information | 22037
- Required Privilege Level | 22037
- Output Fields | 22037
- Sample Output | 22039
- Release Information | 22042

Syntax

```
show mac-vrf forwarding mgrp-policy
<instance instance-name>
<instance instance-name mgrp-id mgrp-id>
```

Description

Display information about logical mesh groups, which are part of a mechanism that Juniper Networks uses to control the flooding of broadcast, unknown unicast, and multicast (BUM) traffic.

You can assign all interfaces that you configure on a spine or super spine device to a particular mesh group. We support the following default mesh groups for a VLAN:

- CE mesh group, to which you assign all directly connected interfaces.
- VE mesh group, to which you assign all virtual tunnel endpoint (VTEP) interfaces for the local POD or data center.
- WAN mesh group, to which you assign all remote VTEP interfaces within standalone or interconnected PODs or data centers.
- RE mesh group, to which you typically do not assign interfaces.

When a BUM packet that must be flooded enters an interface on the spine or super spine device, the system identifies the interface's mesh group. The mesh group associated with a particular VLAN determines the outgoing groups to which the packet should flood. For example, if a BUM packet enters an interface assigned to the CE mesh group, the outgoing flood groups might be the following:

- All interfaces in the CE mesh group except the interface through which the packet entered.

- The VE and WAN mesh groups.
- Whether the spine or super spine device that receives a packet acts as a designated forwarder (DF) or a non-DF.
- The local bias feature, also known as split horizon, which eliminates duplication of traffic to its destination in another POD or data center.

> **NOTE**: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

none	Display the global mesh group policy.
instance *routing instance name*	(Optional) Display the mesh group policy for the specified routing instance.
mgrp-id *mesh group id*	(Optional) Display the mesh group policy for the specified mac-vrf routing instance and mesh group id number.

Additional Information

On MX Series routers and the EX92000 line of switches, this command hierarchy is an alias for the `show l2-learning mgrp-policy` command hierarchy.

On QFX Series switches, this command hierarchy is an alias for the `show ethernet-switching mgrp-policy` command hierarchy.

Required Privilege Level

admin

Output Fields

Table 780 on page 22038 lists the output fields for the `show mac-vrf forwarding mgrp-policy` command.

Table 781 on page 22038 lists the output fields for the `show mac-vrf forwarding mgrp-policy instance` command.

Table 782 on page 22038 lists the output fields for the `show mac-vrf forwarding mgrp-policy instance mgrp-id` command.

Table 780: show mac-vrf forwarding mgrp-policy Output Fields

Field Name	Field Description
Role	Roles include either DF (designated forwarder) or non-DF (non-designated forwarder). We organize the output by role. Mesh groups listed under each role provide a summary of the groups to which a BUM packet might get flooded.
Mesh Group	Name of a mesh group. For each mesh group, the output displays the outgoing flood groups to which a BUM packet gets flooded.

Table 781: show mac-vrf forwarding mgrp-policy instance Output Fields

Field Name	Field Description
Role	Roles include either DF (designated forwarder) or non-DF (non-designated forwarder). The output displays the role for each mesh group.
Mesh Group	Name of a mesh group. For each mesh group, the output displays: • Flags associated with the mesh group. • The outgoing flood groups to which a BUM packet gets flooded.

Table 782: show mac-vrf forwarding mgrp-policy instance mgrp-id Output Fields

Field Name	Field Description
Mesh Group	Name of a mesh group. For the mesh group, the output displays the associated flags.
Interfaces	Interfaces associated with the mesh group. When a spine or super spine device floods a BUM packet to the next hop, it does so through these interfaces.

Sample Output

show mac-vrf forwarding mgrp-policy

```
user@host> show mac-vrf forwarding mgrp-policy
```

```
 Role: DF

Mesh Group:    __all_ces__
    __all_ces__
    __ves__
    __wan_flood__
    __icl_flood__
    __etree_leaf__

Mesh Group:    __ar_flood__
    __all_ces__
    __ar_flood__
    __wan_flood__

Mesh Group:    __esi_flood__

Mesh Group:    __etree_leaf__
    __all_ces__
    __ves__

Mesh Group:    __icl_flood__
    __all_ces__

Mesh Group:    __mlp_flood__

Mesh Group:    __re_flood__
    __all_ces__
    __ves__
    __wan_flood__
    __icl_flood__
    __etree_leaf__

Mesh Group:    __ves__
```

```
    __all_ces__
    __wan_flood__
    __etree_leaf__

Mesh Group:    __wan_flood__
    __all_ces__
    __ves__

 Role: NDF

Mesh Group:    __all_ces__
    __all_ces__
    __ves__
    __wan_flood__
    __icl_flood__

Mesh Group:    __ar_flood__
    __all_ces__
    __ar_flood__
    __wan_flood__

Mesh Group:    __esi_flood__

Mesh Group:    __etree_leaf__

Mesh Group:    __icl_flood__

Mesh Group:    __mlp_flood__

Mesh Group:    __re_flood__
    __all_ces__
    __ves__
    __wan_flood__

Mesh Group:    __ves__
    __all_ces__

Mesh Group:    __wan_flood__
    __all_ces__
Paste device command output here
```

show mac-vrf forwarding mgrp-policy instance

```
user@host> show mac-vrf forwarding mgrp-policy
instance test
Role: DF
Mesh Group:    __all_ces__Flags: 0xa ( split-horizon, shared-RT)
    __all_ces__
    __ves__
    __wan_flood__

Role: DF
Mesh Group:    __mlp_flood__Flags: 0xc ( shared-RT, no-local-flooding)

Role: DF
Mesh Group:    __re_flood__Flags: 0x4c ( shared-RT, no-local-flooding, permanent-comp-nh)
    __all_ces__
    __ves__
    __wan_flood__

Role: DF
Mesh Group:    __ves__Flags: 0xc ( shared-RT, no-local-flooding)
    __all_ces__
    __wan_flood__

Role: DF
Mesh Group:    __wan_flood__Flags: 0xc ( shared-RT, no-local-flooding)
    __all_ces__
    __ves__

Role: Non-DF
Mesh Group:    __ves__Flags: 0xc ( shared-RT, no-local-flooding)

Role: Non-DF
Mesh Group:    __wan_flood__Flags: 0xc ( shared-RT, no-local-flooding)
    __all_ces__
```

show mac-vrf forwarding mgrp-policy instance mgrp-id

```
user@host> show mac-vfr forwarding mgrp-policy
instance evpn-vxlan mgrp-id 0
```

```
Mesh Group:    __ves__Flags: 0xc ( shared-RT, no-local-flooding)
Interfaces
     vtep.32769
     vtep.32770
     vtep.32771
     vtep.32772
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

- show ethernet-switching mac-ip-table | 19660
- show mac-vrf forwarding flood | 21996
- show mac-vrf forwarding mac-table | 22031

show mac-vrf forwarding statistics

IN THIS SECTION

- Syntax | 22043
- Description | 22043
- Options | 22043
- Additional Information | 22043
- Required Privilege Level | 22044
- Output Fields | 22044
- Sample Output | 22045
- Release Information | 22046

Syntax

```
show mac-vrf statistics
<instance instance-name>
<vlan-name vlan-name>
```

Description

Display bridge statistics for the device. The statistics include a count of local MAC addresses, a count of remote MAC addresses, a count of remote VXLAN tunnel endpoints (VTEPs), and the number of enabled VLANs on the device including those in the current `mac-vrf` routing instance. The command output displays traffic statistics by packet and byte.

> **NOTE**: `mac-statistics` must be enabled at the `edit routing-instances <instance-name> protocols evpn` hierarchy or no statistics are gathered.

> **NOTE**: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

none	Display bridge statistics for all VLANs in all routing instances.
instance *instance-name*	(Optional) Display statistics for the specified routing instance.
vlan-name *vlan-name*	(Optional) Display statistics for the specified VLAN.

Additional Information

On supported MX Series routers, this command is an alias for the `show bridge statistics` and `show evpn statistics` commands.

On QFX Series switches, this command is an alias for the `show ethernet-switching statistics` command.

Required Privilege Level

view

Output Fields

Table 783: Output Fields

Field Name	Description
Routing instance	The displayed output is separated by routing instance. The name of the routing instance serves as the section separator.
Index	For the routing instance dislpay: The index number associated with the routing instance. For the interface display: The index number of the local interface.
Sequence number	Sequence number assigned to this bridging domain. Used for debugging.
Non-config static MACs learned	
Bridging domain	The bridge domain associated with the specified routing instance.
MACs learned	Count of MAC addresses learned in the specified routing instance.
MAC limit	The maximum number of MAC addresses that can be associated with the specified routing instance.
Static MACs learned	Count of static MAC addresses learned in the specified routing instance.
Flags	Shows the type of routing instance configured.
Local interface	Shows the local interface associated with the specified routing instance.

Table 783: Output Fields *(Continued)*

Field Name	Description
Broadcast packets	Count of broadcast packets forwarded on the specified interface.
Broadcast bytes	Count of broadcast bytes forwarded on the specified interface
Multicast packets	Count of multicast packets forwarded on the specified interface.
Multicast bytes	Count of multicast bytes forwarded on the specified interface
Flooded packets	Count of packets flooded out on the specified interface
Flooded bytes	Count of bytes flooded out on the specified interface.
Unicast packets	Count of unicast packets forwarded on the specified interface
Unicast bytes	Count of unicast bytes forwarded on the specified interface.
Current MAC count	Count of the number of MAC addresses learned through the specified interface.

Sample Output

show mac-vrf forwarding statistics

```
user@host> show mac-vrf forwarding statistics
```

```
  Local interface: vtep.32769, Index: 574
    Multicast packets:                    0
    Multicast bytes  :                    0
```

```
    Flooded packets  :                    0
    Flooded bytes    :                    0
    Current MAC count:                    0
  Local interface: vtep.32771, Index: 576
    Multicast packets:                    0
    Multicast bytes  :                    0
    Flooded packets  :                    0
    Flooded bytes    :                    0
    Current MAC count:                    2
  Local interface: vtep.32770, Index: 564
    Multicast packets:                    0
    Multicast bytes  :                    0
    Flooded packets  :                    0
    Flooded bytes    :                    0
    Current MAC count:                    0
  Local interface: ae0.0, Index: 575
    Multicast packets:                    0
    Multicast bytes  :                    0
    Flooded packets  :                    0
    Flooded bytes    :                    0
    Current MAC count:                    0 (Limit 8192)
```

Release Information

Command introduced in Junos OS Release 20.4.

RELATED DOCUMENTATION

show mac-vrf forwarding global-information | 22005

show mac-vrf forwarding vlans

IN THIS SECTION

- Syntax | 22047
- Description | 22047

- Options | 22047
- Additional Information | 22048
- Required Privilege Level | 22048
- Output Fields | 22048
- Sample Output | 22049
- Release Information | 22050

Syntax

```
show mac-vrf forwarding vlans
<brief | detail | extensive>
<instance instance-name>
<operational>
<vlan-name>
<interface interface-name>
```

Description

Display information about VLANs configured on mac-vrf routing instances.

> **NOTE**: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

none	Display brief information for all VLANs.
brief | detail | extensive	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display information about the specified routing instance.
operational	(Optional) Display information about the operational routing instances.

`vlan` *`vlan-name`*	(Optional) Display information about the specified VLAN.
`interface` *`interface-name`*	(Optional) Display information about the specified logical interface.

Additional Information

On MX Series routers and the EX9200 line of switches, this command hierarchy is an alias for the `show bridge domain` command hierarchy.

On QFX Series switches, this command hierarchy is an alias for the `show vlans` command hierarchy.

Required Privilege Level

view

Output Fields

Table 784: show mac-vrf forwarding vlans Output Fields

Field Name	Field Description	Level of Output
VLAN name	Name of the VLAN.	none, **brief**
Tag	802.1Q tag applied to this VLAN. If **none** is displayed, no tag is applied.	All levels
Interfaces	Interface associated with learned MAC addresses or **All-members** option (flood entry). An asterisk (*****) beside the interface indicates that the interface is **UP**.	All levels
Routing instance	Name of the routing instance on which the VLAN is configured.	All levels
Number of interfaces	Number of tagged and untagged interfaces associated with a VLAN.	**detail**, **extensive**
Internal Index	VLAN index internal to Junos OS software.	**extensive**
Origin	Manner in which the VLAN was created: **static** or **learn**.	**extensive**

Table 784: show mac-vrf forwarding vlans Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Total MAC count	Total number of 802.1Q-tagged and untagged VLANs on the device.	**detail**, **extensive**
VXLAN Enabled	Shows whether VXLAN (Virtual Extensible LAN) is enabled within the mac-vrf routing instance.	**detail**, **extensive**
MAC aging time	Time, in seconds, for the MAC entry to age out.	**detail**, **extensive**

Sample Output

show mac-vrf forwarding vlans brief

```
user@router> show mac-vrf forwarding vlans brief

Routing instance        VLAN name             Tag          Interfaces
default-switch          c1                    20
                                                           ge-0/0/0.0*
                                                           ge-1/0/0.0*
                                                           ge-2/0/0.0*
default-switch          c2                    30
                                                           ge-0/0/0.0*
                                                           ge-2/0/0.0*
default-switch          default               1

default-switch          iso                   10
                                                           ge-0/0/1.0*
default-switch          iso1                  50
                                                           ge-0/0/0.0*
                                                           ge-2/0/0.0*
default-switch          pri                   100
                                                           ge-0/0/0.0*
                                                           ge-1/0/0.0*
                                                           ge-2/0/0.0*
test                    v100                  NA           ge-2/0/1.100*
```

Release Information

Command introduced in Junos OS Release 20.4R1.

show mac-vrf forwarding vxlan-tunnel-end-point esi

IN THIS SECTION

- Syntax | 22050
- Description | 22050
- Options | 22051
- Additional Information | 22051
- Required Privilege Level | 22051
- Output Fields | 22051
- Sample Output | 22053
- Release Information | 22053

Syntax

```
show mac-vrf forwarding vxlan-tunnel-end-point esi
<esi-identifier esi-identifier>
<instance instance>
<svlbnh>
```

Description

Displays the Ethernet segment identifier (ESI) information about VXLAN tunnel end-points (VTEPs).

NOTE: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only

on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

none Display information about the VTEP.

esi-identifier *esi-identifier* (Optional) Display information about the specified Ethernet segment identified by the ESI value.

instance *instance* (Optional) Display ESI information about the specified routing instance.

svlbnh (Optional) Display the shared VXLAN load-balancing next hop (SVLBNH) index for a given ESI and instance.

You must use this option in conjunction with both the `esi-identifier` and `instance` options. For detailed information about a particular SVLBNH index, see "show ethernet-switching vxlan-tunnel-end-point svlbnh" on page 19733.

Additional Information

On MX Series routers and the EX9200 line of switches, this command hierarchy is an alias for the `show l2-learning vxlan-tunnel-end-point` command hierarchy.

On QFX Series switches, this command hierarchy is an alias for the `show ethernet-switching vxlan-tunnel-end-point` command hierarchy.

Required Privilege Level

view

Output Fields

Table 785: show mac-vrf forwarding vxlan-tunnel-end-point esi Output Fields

Field Name	Field Description
ESI	Ethernet segment identifier

Table 785: show mac-vrf forwarding vxlan-tunnel-end-point esi Output Fields *(Continued)*

Field Name	Field Description
RTT	Routing instance in which you have configured the ESI
VLBNH	VXLAN load balancing next hop.
INH	Indirect next hop
ESI-IFL	ESI logical interface
LOC-IFL	Local logical interface for a given Ethernet segment
#RVTEPS	Number of remote VTEPs
RVTEP-IP	IP address of the remote VTEP
RVTEP-IFL	Index number of the remote VTEP.
VENH	VXLAN encapsulated composite next hop
MASK-ID	Position of the remote VTEP in the ESI mask. Each remote VTEP is assigned a unique position.
FLAGS	Internal flags used for debugging. The flags help to identify the state of the remote VTEP.
MAC-COUNT	Count of MACs learned on the remote VTEP for a given ESI.
SVLBNH	Index number of the SVLBNH.

Sample Output

show mac-vrf forwarding vxlan-tunnel-end-point esi

```
user@host> show mac-vrf forwarding vxlan-tunnel-end-point esi
```

```
PESI                            RTT                      VLNBH INH     ESI-IFL    LOC-IFL    #RVTEPs
00:10:11:12:13:14:15:16:17:11 default-switch           1758  131077  esi.1758             2
Aliasing
    RVTEP-IP             RVTEP-IFL      VENH     MASK-ID   FLAGS        MAC-COUNT
    111.1.1.1            vtep.32769     1753     0         2            0
    112.1.1.1            vtep.32773     1757     1         2            0
ESI                          RTT                      VLNBH INH     ESI-IFL    LOC-IFL    #RVTEPs
00:10:11:12:13:14:15:16:17:22 default-switch           1768  131080  esi.1768   ae0.0,     1
Aliasing
    RVTEP-IP             RVTEP-IFL      VENH     MASK-ID   FLAGS        MAC-COUNT
    113.1.1.1            vtep.32774     1767     0         2            0
ESI                          RTT                      VLNBH INH     ESI-IFL    LOC-IFL    #RVTEPs
05:00:00:00:64:00:00:00:64:00 default-switch           1762  131079  esi.1762             3
    RVTEP-IP             RVTEP-IFL      VENH     MASK-ID   FLAGS        MAC-COUNT
    101.1.1.1            vtep.32770     1754     0         2            2
    103.1.1.1            vtep.32771     1755     1         2            2
    102.1.1.1            vtep.32772     1756     2         2            2
ESI                          RTT                      VLNBH INH     ESI-IFL    LOC-IFL    #RVTEPs
05:00:00:00:64:00:00:00:c8:00 default-switch           1760  131078  esi.1760             3
    RVTEP-IP             RVTEP-IFL      VENH     MASK-ID   FLAGS        MAC-COUNT
    101.1.1.1            vtep.32770     1754     0         2            2
    103.1.1.1            vtep.32771     1755     1         2            2
    102.1.1.1            vtep.32772     1756     2         2            2
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show mac-vrf forwarding vxlan-tunnel-end-point svlbnh | 22063

show mac-vrf forwarding vxlan-tunnel-end-point remote

IN THIS SECTION

- Syntax | 22054
- Description | 22054
- Options | 22055
- Additional Information | 22055
- Required Privilege Level | 22055
- Output Fields | 22055
- Sample Output | 22061
- Release Information | 22063

Syntax

```
show mac-vrf forwarding vxlan-tunnel-end-point remote
<detail>
<instance mac-vrf routing-instance name>
<ip remote vtep ip address>
<mac-table>
<summary>
<vtep-source-interface local vtep source interface ip address>
```

Description

Display information about remote VXLAN tunnel endpoints (VTEPs).

By default, this command displays VTEP information for multiple MAC-VRF routing instances, if configured.

NOTE: We've implemented the `show mac-vrf` commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only

on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

none	Display information about remote VTEPs associated with all MAC-VRF routing instances.
detail	Display information about remote VTEPs associated with all MAC-VRF routing instances.
instance *mac-vrf routing-instance name*	(Optional) Display VTEP information about the specified MAC-VRF routing instance.
ip *remote vtep ip address*	(Optional) Display detailed information about the specified remote VTEP IP address.
mac-table	(Optional) Display MAC table information along with the remote VTEP information.
summary	(Optional) Display only summary information about all MAC-VRF routing instances.
vtep-source-interface *local vtep source interface ip address*	(Optional) Display detailed remote VTEP information associated with the specified local VTEP source IP address.

Additional Information

On MX Series routers and the EX9200 line of switches, this command hierarchy is an alias for the `show l2-learning vxlan-tunnel-end-point remote` command hierarchy.

On QFX Series switches, this command hierarchy is an alias for the `show ethernet-switching vxlan-tunnel-end-point remote` command hierarchy.

Required Privilege Level

view

Output Fields

Table 1 on page 22056 lists the output fields for the `show mac-vrf forwarding vxlan-tunnel-end-point remote` command.

Table 2 on page 22057 lists the output fields for the `show mac-vrf forwarding vxlan-tunnel-end-point remote instance` command.

Table 3 on page 22058 lists the output fields for the `show mac-vrf forwarding vxlan-tunnel-end-point remote ip` command.

Table 4 on page 22059 lists the output fields for the `show mac-vrf forwarding vxlan-tunnel-end-point remote mac-table` command.

Table 5 on page 22059 lists the output fields for the `show mac-vrf forwarding vxlan-tunnel-end-point remote summary` command.

Table 6 on page 22060 lists the output fields for the `show mac-vrf forwarding vxlan-tunnel-end-point remote vtep-source-interface` command.

Table 786: show mac-vrf forwarding vxlan-tunnel-end-point remote Output Fields

Field Name	Field Description
Logical System Name	Name of the logical system on which you have configured the MAC-VRF routing instance
Id	Logical system ID
SVTEP-IP	IP address of the (local) source VTEP
IFL	Logical interface name
L3-Idx	System index number of the associated Layer 3 interface
SVTEP-Mode	Source VTEP mode
RVTEP-IP	IP address of the remote VTEP
L2-RTT	Layer 2 routing instance name
IFL-Idx	System index number of the associated logical interface
Interface	Name of the VTEP interface
NH-Id	Next-hop ID
RVTEP-Mode	Mode of the remote VTEP interface
Flags	Flag
VNID	VLAN network ID

Table 786: show mac-vrf forwarding vxlan-tunnel-end-point remote Output Fields *(Continued)*

Field Name	Field Description
MC-Group-IP	Multicast group IP address. Defaults to 0.0.0.0 if not configured.

Table 787: show mac-vrf forwarding vxlan-tunnel-end-point remote instance *mac-vrf routing-instance name* Output Fields

Field Name	Field Description
Logical System Name	Name of the logical system on which you have configured the MAC-VRF routing instance
Id	Logical system ID
SVTEP-IP	IP address of the (local) source VTEP
IFL	Logical interface name
L3-Idx	System index number of the associated Layer 3 interface
SVTEP-Mode	Source VTEP mode
RVTEP-IP	IP address of the remote VTEP
L2-RTT	Layer 2 routing instance name
IFL-Idx	System index number of the associated logical interface
Interface	Name of the VTEP interface
NH-Id	Next-hop ID
RVTEP-Mode	Mode of the remote VTEP interface
Flags	Flag
VNID	VXLAN network ID
MC-Group-IP	Multicast group IP address. Defaults to 0.0.0.0 if not configured.

Table 788: show mac-vrf forwarding vxlan-tunnel-end-point remote ip *remote vtep ip address* Output Fields

Field Name	Field Description
Logical System	Name of the logical system on which you have configured the MAC-VRF routing instance
Routing instance	Name of the MAC-VRF routing instance whose tunnel connects with the remote IP address
Bridging domain	Name of the bridging domain
VLAN	VLAN ID of the tunnel
VXLAN ID	VXLAN ID
Multicast Group IP	Multicast group IP address. Defaults to 0.0.0.0 if not configured.
Remote VTEP	IP address of the remote VTEP
Nexthop Id	ID of the next hop
MAC address	MAC address associated with the remote VTEP IP address
MAC Flags	Flags: **S** Static MAC address **D** Dynamic MAC address **L** Locally learned MAC address **C** Control MAC address **SE** Statistics enabled **NM** Non-configured MAC address **R** Remote PE MAC address **P** Pinned MAC address
Logical interface	Logical interface associated with the remote MAC address and the remote IP address
Remote VTEP IP address	IP address of the remote VTEP

Table 789: show mac-vrf forwarding vxlan-tunnel-end-point remote mac-table Output Fields

Field Name	Field Description
Logical System	Name of the logical system on which you have configured the MAC-VRF routing instance
Routing instance	Name of the MAC-VRF routing instance whose tunnel connects with the remote IP address
Bridging domain	Name of the bridging domain
VLAN	VLAN ID of the tunnel
VNID	VXLAN network ID
MAC address	MAC address associated with the remote VTEP IP address
MAC Flags	Flags: **S** Static MAC address **D** Dynamic MAC address **L** Locally learned MAC address **C** Control MAC address **SE** Statistics enabled **NM** Non-configured MAC address **R** Remote PE MAC address **P** Pinned MAC address
Logical interface	Logical interface associated with the remote MAC address and the remote IP address
Remote VTEP IP address	IP address of the remote VTEP

Table 790: show mac-vrf forwarding vxlan-tunnel-end-point remote summary Output Fields

Field Name	Field Description
Logical System Name	Name of the logical system on which you have configured the MAC-VRF routing instance
Id	Logical system ID

Table 790: show mac-vrf forwarding vxlan-tunnel-end-point remote summary Output Fields *(Continued)*

Field Name	Field Description
SVTEP-IP	IP address of the (local) source VTEP
IFL	Logical interface name
L3-Idx	System index number of the associated Layer 3 interface
SVTEP-Mode	Source VTEP mode
RVTEP-IP	IP address of the remote VTEP
L2-RTT	Layer 2 routing instance name
IFL-Idx	System index number of the associated logical interface
Interface	Name of the VTEP interface
NH-Id	Next-hop ID
RVTEP-Mode	Mode of the remote VTEP interface.
Flags	Flag

Table 791: show mac-vrf forwarding vxlan-tunnel-end-point remote vtep-source-interface *local vtep source ip address* Output Fields

Field Name	Field Description
Logical System Name	Name of the logical system on which you have configured the MAC-VRF routing instance
Id	Logical system ID
SVTEP-IP	IP address of the (local) source VTEP
IFL	Logical interface name
L3-Idx	System index number of the associated Layer 3 interface

Table 791: show mac-vrf forwarding vxlan-tunnel-end-point remote vtep-source-interface *local vtep source ip address* Output Fields *(Continued)*

Field Name	Field Description
SVTEP-Mode	Source VTEP mode
RVTEP-IP	IP address of the remote VTEP
L2-RTT	Layer 2 routing instance name
IFL-Idx	System index number of the associated logical interface
Interface	Name of the VTEP interface
NH-Id	Next-hop ID
RVTEP-Mode	Mode of the remote VTEP interface
Flags	Flag
VNID	VXLAN network ID
MC-Group-IP	Multicast group IP address. Defaults to 0.0.0.0 if not configured.

Sample Output

show mac-vrf forwarding vxlan-tunnel-end-point remote

```
user@host> show mac-vrf forwarding vxlan-tunnel-end-point remote
```

```
Logical System Name       Id  SVTEP-IP         IFL   L3-Idx    SVTEP-Mode
<default>                 0   10.255.1.1       lo0.0    0
 RVTEP-IP         L2-RTT                   IFL-Idx   Interface    NH-Id   RVTEP-Mode Flags
 10.255.1.2       macvrf_v150              557       vtep.32770   1793    RNVE
    VNID          MC-Group-IP
    150           0.0.0.0
 RVTEP-IP         L2-RTT                   IFL-Idx   Interface    NH-Id   RVTEP-Mode Flags
```

```
 10.255.1.3       macvrf_v150              556       vtep.32769   1791    RNVE
 10.255.1.4       macvrf_v150              558       vtep.32771   1799    RNVE
```

show mac-vrf forwarding vxlan-tunnel-end-point remote ip *remote vtep ip address*

```
user@host> show mac-vrf forwarding vxlan-tunnel-end-point remote ip 10.255.1.2
```

```


MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC , P -Pinned MAC)

Logical system   : <default>
Routing instance : macvrf_v150
 Bridging domain : v150+150, VLAN : 150
 VXLAN ID : 150, Multicast Group IP : 0.0.0.0
 Remote VTEP : 10.255.1.2, Nexthop ID : 1793
   MAC                 MAC      Logical          Remote VTEP
   address             flags    interface        IP address
   02:05:86:71:14:00   DR       vtep.32770       10.255.1.2
   40:00:10:11:11:02   DR       esi.1792         10.255.1.2
   50:00:10:11:11:02   DR       esi.1792         10.255.1.2
```

show mac-vrf forwarding vxlan-tunnel-end-point remote mac-table

```
user@host> show mac-vrf forwarding vxlan-tunnel-end-point remote mac-table
```

```


MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC , P -Pinned MAC)

Logical system   : <default>
Routing instance : macvrf_v150
 Bridging domain : v150+150, VLAN : 150
 VXLAN ID : 150, Multicast Group IP : 0.0.0.0
 Remote VTEP : 10.255.1.2, Nexthop ID : 1793
   MAC                 MAC      Logical          Remote VTEP
   address             flags    interface        IP address
```

```
    02:05:86:71:14:00   DR       vtep.32770       10.255.1.2
    40:00:10:11:11:02   DR       esi.1792         10.255.1.2
    50:00:10:11:11:02   DR       esi.1792         10.255.1.2
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show mac-vrf forwarding vxlan-tunnel-end-point svlbnh | 22063

show mac-vrf forwarding vxlan-tunnel-end-point esi | 22050

show mac-vrf forwarding vxlan-tunnel-end-point svlbnh

IN THIS SECTION

- Syntax | 22063
- Description | 22064
- Options | 22064
- Additional Information | 22064
- Required Privilege Level | 22064
- Output Fields | 22064
- Sample Output | 22065
- Release Information | 22066

Syntax

```
show mac-vrf forwarding vxlan-tunnel-end-point svlbnh
<brief | detail | extensive>
<overflow>
<svlbnh-index index-number>
```

Description

Display information about a shared VXLAN load-balancing next hop (SVLBNH) and the virtual tunnel end-points (VTEPs) associated with the SVLBNH.

> **NOTE**: We've implemented the show mac-vrf commands on all supported platforms. If a particular platform doesn't support a specific sub-command or option, when you run the command, the output is empty. For example, `show mac-vrf forwarding age` is available only on MX Series platforms. If you run that command on any other supported platform, the output is empty.

Options

none	(Same as `brief`) Display information for all SVLBNHs.
`brief` | `detail` | `extensive`	(Optional) Display the specified level of output.
`overflow`	(Optional) Display the pending requests for SVLBNH allocation.
`svlbnh-index` *`index-number`*	(Optional) Display information about the specified SVLBNH index.

Additional Information

On MX Series routers and the EX9200 line of switches, this command hierarchy is an alias for the `show l2-learning vxlan-tunnel-end-point svlbnh` command hierarchy.

On QFX Series switches, this command hierarchy is an alias for the `show ethernet-switching vxlan-tunnel-end-point svlbnh` command hierarchy.

Required Privilege Level

view

Output Fields

Field Name	Field Description
SVLBNH *index-number*	Index number of the SVLBNH.

(Continued)

Field Name	Field Description
RVTEP count	Number of remote VTEPs associated with the SVLBNH.
ESI count	Number of Ethernet segment identifiers (ESIs) associated with the SVLBNH.
RVTEP-IP	IP address of the remote VTEP.
RVTEP-IFL	Index number of remote VTEP.
VENH	VXLAN encapsulated next hop (VENH) associated with the SVLBNH.
ESI	ESIs associated with the SVLBNH.
RTT	Routing instance in which you have configured the ESI.

Sample Output

show mac-vrf forwarding vxlan-tunnel-end-point svlbnh detail

```
user@host> show mac-vrf forwarding vxlan-tunnel-end-point svlbnh detail
```

```
SVLBNH: 1882        RVTEP count: 2           ESI count: 1
    RVTEP-IP             RVTEP-IFL            VENH
    10.0.5.1             vtep.32770           1859
    10.0.3.1             vtep.32769           1835
    ESI                           RTT
    00:00:00:00:00:00:00:00:01:03 default-switch

SVLBNH: 1879        RVTEP count: 2           ESI count: 1
    RVTEP-IP             RVTEP-IFL            VENH
    10.0.2.1             vtep.32771           1869
    10.0.3.1             vtep.32769           1835
```

```
    ESI                           RTT
    00:00:00:00:00:00:00:00:01:01 default-switch

SVLBNH: 1881        RVTEP count: 2           ESI count: 1
    RVTEP-IP             RVTEP-IFL            VENH
    10.0.2.1             vtep.32771           1869
    10.0.5.1             vtep.32770           1859
    ESI                           RTT
    00:00:00:00:00:00:00:00:01:02 default-switch

SVLBNH: 1870        RVTEP count: 3           ESI count: 5
    RVTEP-IP             RVTEP-IFL            VENH
    10.0.2.1             vtep.32771           1869
    10.0.5.1             vtep.32770           1859
    10.0.3.1             vtep.32769           1835
    ESI                           RTT
    00:00:00:00:00:00:00:00:01:04 default-switch
    05:00:00:02:9a:00:00:00:65:00 default-switch
    05:00:00:02:9a:00:00:00:66:00 default-switch
    05:00:00:02:9a:00:00:00:67:00 default-switch
    05:00:00:02:9a:00:00:00:68:00 default-switch
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show mac-vrf forwarding vxlan-tunnel-end-point svlbnh | 22063

show mld group

IN THIS SECTION

- Syntax | 22067
- Description | 22067

- Options | 22067
- Required Privilege Level | 22067
- Output Fields | 22068
- Sample Output | 22069
- show mld group (User Plane output for ACX7100) | 22073
- Release Information | 22074

Syntax

```
show mld group
<brief | detail>
<group-name>
<logical-system (all | logical-system-name)>
```

Description

Display information about Multicast Listener Discovery (MLD) group membership.

(ACX7100 Routers as CPE) Displays monitoring information for multicast interfaces on customer premises equipments (CPE) with user-plane configuration.

Options

none	Display standard information about all MLD groups.
brief | detail	(Optional) Display the specified level of output.
group-name	(Optional) Display MLD information about the specified group.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 792 on page 22068 describes the output fields for the `show mld group` command. Output fields are listed in the approximate order in which they appear.

Table 792: show mld group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the MLD membership report; **local** means that the local router joined the group itself.	All levels
Group	Group address.	All levels
Source	Source address.	All levels
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.	All levels
Last reported by	Address of the host that last reported membership in this group.	All levels
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	**detail**
Timeout	Time remaining until the group membership is removed.	**brief none**
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	**detail**
Type	Type of group membership: • **Dynamic**—Host reported the membership. • **Static**—Membership is configured.	All levels

Sample Output

show mld group (Include Mode)

```
user@host> show mld group
Interface: fe-0/1/2.0
    Group: ff02::1:ff05:1a67
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     245 Type: Dynamic
    Group: ff02::1:ffa8:c35e
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     241 Type: Dynamic
    Group: ff02::2:43e:d7f6
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     244 Type: Dynamic
    Group: ff05::2
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     244 Type: Dynamic
Interface: local
    Group: ff02::2
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show mld group (Exclude Mode)

```
user@host> show mld group
Interface: ge-0/2/2.0
```

```
Interface: ge-0/2/0.0
    Group: ff02::6
        Source: ::
        Last reported by: fe80::21f:12ff:feb6:4b3a
        Timeout:     245 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: fe80::21f:12ff:feb6:4b3a
        Timeout:      28 Type: Dynamic
Interface: local
    Group: ff02::2
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show mld group brief

The output for the `show mld group brief` command is identical to that for the `show mld group` command. For sample output, see "show mld group (Include Mode)" on page 22069 "show mld group (Exclude Mode)" on page 22069.

show mld group detail (Include Mode)

```
user@host> show mld group detail
Interface: fe-0/1/2.0
    Group: ff02::1:ff05:1a67
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     224 Type: Dynamic
    Group: ff02::1:ffa8:c35e
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     220 Type: Dynamic
    Group: ff02::2:43e:d7f6
        Group mode: Include
```

```
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     223 Type: Dynamic
    Group: ff05::2
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     223 Type: Dynamic
Interface: so-1/0/1.0
    Group: ff02::2
        Group mode: Include
        Source: ::
        Last reported by: fe80::280:42ff:fe15:f445
        Timeout:     258 Type: Dynamic
Interface: local
    Group: ff02::2
        Group mode: Include
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show mld group detail (Exclude Mode)

```
user@host> show mld group detail
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
    Group: ff02::6
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: fe80::21f:12ff:feb6:4b3a
        Group timeout:     226 Type: Dynamic
    Group: ff02::16
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: fe80::21f:12ff:feb6:4b3a
```

```
        Group timeout:     246 Type: Dynamic
Interface: local
    Group: ff02::2
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: Local
        Group timeout:       0 Type: Dynamic
    Group: ff02::16
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: Local
        Group timeout:       0 Type: Dynamicuser@up1> show mld group
Interface: pp0.3221225483, Groups: 2
    Group: ff1e::1
        Group mode: Exclude
        Source: ::
        Last reported by: fe80::e
        Timeout:     243 Type: Dynamic
    Group: ff1e::2
        Group mode: Exclude
        Source: ::
        Last reported by: fe80::e
        Timeout:     249 Type: Dynamic
Interface: demux0.2147483653, Groups: 2
    Group: ff1e::1
        Group mode: Exclude
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: ROUTE
    Group: ff1e::2
        Group mode: Exclude
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: ROUTE
Interface: local, Groups: 4
    Group: ff02::2
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
```

```
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::1:2
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff05::1:3
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show mld group (User Plane output for ACX7100)

Displays the following output in the user-plane for customer premise equipment (CPE).

```
user@UP1> show mld group
Interface: pp0.3221225483, Groups: 2
    Group: ff1e::1
        Group mode: Exclude
        Source: ::
        Last reported by: fe80::e
        Timeout:     243 Type: Dynamic
    Group: ff1e::2
        Group mode: Exclude
        Source: ::
        Last reported by: fe80::e
        Timeout:     249 Type: Dynamic
Interface: demux0.2147483653, Groups: 2
    Group: ff1e::1
        Group mode: Exclude
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: ROUTE
    Group: ff1e::2
        Group mode: Exclude
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: ROUTE
Interface: local, Groups: 4
    Group: ff02::2
        Source: ::
```

```
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::1:2
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff05::1:3
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

Release Information

Command introduced before Junos OS Release 7.4.

Command updated in Junos OS Release 24.4R1 to include multicast support monitoring for ACX7100 router interfaces.

RELATED DOCUMENTATION

clear mld membership | 3702

show mld interface

IN THIS SECTION

- Syntax | 22075
- show mld interface (User Plane output for ACX7100) | 22075
- Description | 22076
- Options | 22076
- Required Privilege Level | 22076

- Output Fields | 22076
- Sample Output | 22080
- Release Information | 22081

Syntax

```
show mld interface
<brief | detail>
<interface-name>
<logical-system (all | logical-system-name)>
```

show mld interface (User Plane output for ACX7100)

Displays the following output in the user-plane for customer premise equipment (CPE).

```
user@UP1> show mld interface
Interface: pp0.3221225483
    Querier: fe80::5668:adff:fec0:2ec4
    State:         Up Timeout:    1000 Version:  2 Groups:      2
    Group threshold:   0
    Group log-interval:  10000
    Immediate leave: On
    Promiscuous mode: Off
    Passive: Off
Interface: demux0.2147483653
    Querier: fe80::5668:adff:fec0:2ec4
    State:         Up Timeout:    None Version:  2 Groups:      2
    Immediate leave: On
    Promiscuous mode: Off
    Passive: Off

Configured Parameters:
MLD Query Interval: 125.0
MLD Query Response Interval: 10.0
MLD Last Member Query Interval: 1.0
MLD Robustness Count: 2
```

```
Derived Parameters:
MLD Membership Timeout: 260.0
MLD Other Querier Present Timeout: 255.0
```

Description

Display information about multipoint Listener Discovery (MLD)-enabled interfaces.

(ACX7100 Routers as CPE) Displays monitoring information for multicast interfaces on customer premises equipments (CPE) with user-plane configuration.

Options

none	Display standard information about all MLD-enabled interfaces.
brief | detail	(Optional) Display the specified level of output.
interface-name	(Optional) Display information about the specified interface.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 793 on page 22076 describes the output fields for the `show mld interface` command. Output fields are listed in the approximate order in which they appear.

Table 793: show mld interface Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the interface.	All levels

Table 793: show mld interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Querier	Address of the router that has been elected to send membership queries.	All levels
State	State of the interface: **Up** or **Down**.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the interface.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy at the MLD interface.	All levels
Timeout	How long until the MLD querier is declared to be unreachable, in seconds.	All levels
Version	MLD version being used on the interface: **1** or **2**.	All levels
Groups	Number of groups on the interface.	All levels

Table 793: show mld interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Passive	State of the passive mode option: • **On**—Indicates that the router can run IGMP or MLD on the interface but not send or receive control traffic such as IGMP or MLD reports, queries, and leaves. • **Off**—Indicates that the router can run IGMP or MLD on the interface and send or receive control traffic such as IGMP or MLD reports, queries, and leaves. The `passive` statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the **on** state declaration: • **send-general-query**—The interface sends general queries. • **send-group-query**—The interface sends group-specific and group-source-specific queries. • **allow-receive**—The interface receives control traffic	All levels
OIF map	Name of the OIF map associated to the interface.	All levels
SSM map	Name of the source-specific multicast (SSM) map used on the interface, if configured.	All levels
Group limit	Maximum number of groups allowed on the interface. Any memberships requested after the limit is reached are rejected.	All levels
Group threshold	Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.	All levels
Group log-interval	Time (in seconds) between consecutive log messages.	All levels

Table 793: show mld interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Immediate Leave	State of the immediate leave option: • **On**—Indicates that the router removes a host from the multicast group as soon as the router receives a multicast listener done message from a host associated with the interface. • **Off**—Indicates that after receiving a multicast listener done message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.	All levels
Distributed	State of MLD, which, by default, takes place on the Routing Engine for MX Series routers but can be distributed to the Packet Forwarding Engine to provide faster processing of join and leave events. • **On**—distributed MLD is enabled.	All levels
Configured Parameters	Information configured by the user. • **MLD Query Interval (.1 secs)**—Interval at which this router sends membership queries when it is the querier. • **MLD Query Response Interval (.1 secs)**—Time that the router waits for a report in response to a general query. • **MLD Last Member Query Interval (.1 secs)**—Time that the router waits for a report in response to a group-specific query. • **MLD Robustness Count**—Number of times the router retries a query.	All levels
Derived Parameters	Derived information. • **MLD Membership Timeout (.1 secs)**—Timeout period for group membership. If no report is received for these groups before the timeout expires, the group membership will be removed. • **MLD Other Querier Present Timeout (.1 secs)**—Time that the router waits for the IGMP querier to send a query.	All levels

Sample Output

show mld interface

```
user@host> show mld interface
Interface: fe-0/0/0
    Querier: None
    State: Up         Timeout:       0    Version:  1    Groups:      0
    SSM Map Policy: ssm-policy-A
Interface: at-0/3/1.0
    Querier:  8038::c0a8:c345
    State: Up         Timeout:    None    Version:  1    Groups:      0
    SSM Map Policy: ssm-policy-B
Interface: fe-1/0/1.0
    Querier:  ::192.168.195.73
    State: Up         Timeout:    None    Version:  1    Groups:      3
    SSM Map Policy: ssm-policy-C
    SSM map: ipv6map1
Immediate Leave: On

Promiscuous Mode: Off
Passive: Off
Distributed: OnConfigured Parameters:

Configured Parameters:
MLD Query Interval (.1 secs): 1250
MLD Query Response Interval (.1 secs): 100
MLD Last Member Query Interval (.1 secs): 10
MLD Robustness Count: 2

Derived Parameters:
MLD Membership Timeout (.1secs): 2600
MLD Other Querier Present Timeout (.1 secs): 2550
```

show mld interface brief

The output for the `show mld interface brief` command is identical to that for the `show mld interface` command. For sample output, see "show mld interface" on page 22080.

show mld interface detail

The output for the `show mld interface detail` command is identical to that for the `show mld interface` command. For sample output, see "show mld interface" on page 22080.

show mld interface <interface-name>

```
user@host# show mld interface ge-3/2/0.0
Interface: ge-3/2/0.0
  Querier: 203.0.113.111
  State: Up Timeout:    None Version:  3 Groups:      1
  Group limit: 8
  Group threshold: 60
  Group log-interval: 10
  Immediate leave:  Off
  Promiscuous mode:  Off   Distributed: On
```

Release Information

Command introduced before Junos OS Release 7.4.

Command updated in Junos OS Release 24.4R1 to include multicast support monitoring for ACX7100 router interfaces.

RELATED DOCUMENTATION

clear mld membership | 3702

show mld snooping evpn database

IN THIS SECTION

- Syntax | 22082
- Description | 22082
- Options | 22082

- Required Privilege Level | 22083
- Output Fields | 22083
- Sample Output | 22084
- Release Information | 22085

Syntax

```
show mld snooping evpn database
<brief | detail>
<group>
<bridge-domain bridge-domain-id>
<instance routing-instance-name>
<interface interface-name>
<logical-system (logical-system-name | all)>
<vlan vlan-id>
```

Description

Display MLD snooping information for IPv6 multicast forwarding in an EVPN network.

Options

brief | detail Level of output to display. The default output is the same as `brief`.

group Display output only for the specified IPv6 multicast group IP address range.

bridge-domain Display output only for the specified bridge domain name.

instance Display output for the specified EVPN instance (EVI).

interface Display output only for the specified interface.

logical-system Display output for a specific logical system name or `all` logical systems.

vlan Display output only for a specified VLAN ID (0 through 4094).

Required Privilege Level

routing

Output Fields

Table 794 on page 22083 lists the output fields for the `show mld snooping evpn database` command. Output fields are listed in the approximate order in which they appear.

Table 794: show mld snooping evpn database Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the EVPN routing instance for which information is displayed.	All levels
Bridge-Domain, VLAN ID	Bridge domain name and VLAN ID.	All levels
Group IP	Multicast IPv6 address of the multicast group.	All levels
Source IP	IPv6 address of the multicast data source for the group.	All levels
Access OIF Count	Number of outgoing interface (OIF) list entries on the access side.	All levels
Core NH	Next hop ID corresponding to the next hop interface.	All levels
Access OIF (more information)	More details about OIF entries: • `Interface`—Interface name • `Local`—Local receiver. • `Remote`—Remote receiver.	`detail`

Sample Output

show mld snooping evpn database

```
user@device> show mld snooping evpn database
Instance: EVPN-2
    Bridge-Domain: V100, VLAN ID: 100
        Group IP: ff02::db8:25:1, Source IP: ::, Access OIF Count: 2, Core NH: 1049231
        Group IP: ff02::db8:33:1, Source IP: ::, Access OIF Count: 2, Core NH: 1049231

Instance: EVPN-4
    Bridge-Domain: V200, VLAN ID: 200
        Group IP: ff02::db8:25:1, Source IP: ::, Access OIF Count: 1, Core NH: 1049232
        Group IP: ff02::db8:33:1, Source IP: ::, Access OIF Count: 1, Core NH: 1049232
```

show mld snooping evpn database detail

```
user@device> show mld snooping evpn database
detail
Instance: EVPN-2
  Bridge-Domain: V100, VLAN ID: 100
        Group IP: ff02::db8:25:1, Source IP: ::
            Core NH: 1049231
            Access OIF Count: 2
                  Interface    Local    Remote
                ae3.0            0        1
                ge-0/0/4.0       1        0
        Group IP: ff02::db8:33:1, Source IP: ::
            Core NH: 1049231
            Access OIF Count: 2
                  Interface    Local    Remote
                ge-0/0/4.0       1        0
                ae3.0            1        0

Instance: EVPN-4
  Bridge-Domain: V200, VLAN ID: 200
        Group IP: ff02::db8:25:1, Source IP: ::
            Core NH: 1049232
            Access OIF Count: 1
```

```
                Interface    Local    Remote
            ae4.0            0        1
     Group IP: ff02::db8:33:1, Source IP: ::
         Core NH: 1049232
         Access OIF Count: 1
                Interface    Local    Remote
            ae4.0            1        0
```

Release Information

Introduced in Junos OS Release 18.4R1.

Introduced in Junos OS Release 19.4R2 on ACX Series routers.

RELATED DOCUMENTATION

Configure Multicast Forwarding with MLD Snooping in an EVPN-MPLS Environment

Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment

show mld snooping interface

IN THIS SECTION

- Syntax | 22086
- Description | 22086
- Options | 22086
- Required Privilege Level | 22086
- Output Fields | 22086
- Sample Output | 22088
- Release Information | 22090

Syntax

```
show mld snooping interface
<brief | detail>
<instance routing-instance>
<interface-name>
<qualified-vlan vlan-name>
<vlan vlan-name>
```

Description

Display MLD snooping information for an interface.

Options

none	Display MLD snooping information for all interfaces on which MLD snooping is enabled.
brief | detail	(Optional) Display the specified level of output. The default is `brief`.
instance *routing-instance*	(Optional) Display MLD snooping information for the specified routing instance.
interface-name	(Optional) Display MLD snooping information for the specified interface.
qualified-vlan *vlan-name*	(Optional) Display MLD snooping information for the specified qualified VLAN.
vlan *vlan-name*	(Optional) Display MLD snooping information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 795 on page 22087 lists the output fields for the `show mld snooping interface` command. Output fields are listed in the approximate order in which they appear. Details may differ for EX switches and MX routers.

Table 795: show mld snooping interface Output Fields

Field Name	Field Description	Level of Output
`Instance`	Routing instance for MLD snooping.	All levels
`Learning Domain`	Learning domain for MLD snooping.	All levels
`Vlan`	Name of the VLAN for which MLD snooping is enabled.	All levels
`Interface`	Name of the interface.	All levels
`State`	State of the interface: `Up` or `Down`.	`detail`, none
`Groups`	Number of multicast groups on the interface.	`detail`, none
`Immediate leave`	State of the immediate leave option: • `On`—Indicates that the MLD querier removes a host from the multicast group as soon as it receives a leave report from a host associated with the interface. • `Off`—Indicates that after receiving a leave report, instead of removing a host from the multicast group immediately, the MLD querier sends a group query to determine if there are any other hosts on that interface still interested in the multicast group.	`detail`, none
`Router interface`	Indicates whether the interface is a multicast router interface: `Yes` or `No`.	`detail`

Table 795: show mld snooping interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Configured Parameters`	Information configured by the user. • `MLD Query Interval`—Interval (in seconds) at which the MLD querier sends membership queries. • `MLD Query Response Interval`—Time (in seconds) that the MLD querier waits for a report in response to a general query. • `MLD Last Member Query Interval`—Time (in seconds) that the MLD querier waits for a report in response to a group-specific query. • `MLD Robustness Count`—Number of times the MLD querier retries a query.	All levels

Sample Output

show mld snooping interface

```
user@switch> show mld snooping interface
Instance: default-switch

Vlan: v100

Learning-Domain: default
Interface: ge-0/0/1.0
    State:         Up Groups:      1
    Immediate leave: Off
    Router interface: no
Interface: ge-0/0/2.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no

Configured Parameters:
MLD Query Interval: 125.0
MLD Query Response Interval: 10.0
```

```
MLD Last Member Query Interval: 1.0
MLD Robustness Count: 2
```

show mld snooping interface ge-0/0/2.0

```
user@switch> show mld snooping interface ge-0/0/2.0
Instance: default-switch

Vlan: v100

Learning-Domain: default
Interface: ge-0/0/2.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no

Configured Parameters:
MLD Query Interval: 125.0
MLD Query Response Interval: 10.0
MLD Last Member Query Interval: 1.0
MLD Robustness Count: 2
```

show mld snooping interface brief

```
user@switch> show mld snooping interface brief
Instance: default-switch

Vlan: v1

Learning-Domain: default
Interface: ge-0/0/1.0
Interface: ge-0/0/2.0

Configured Parameters:
MLD Query Interval: 125.0
MLD Query Response Interval: 10.0
MLD Last Member Query Interval: 1.0
MLD Robustness Count: 2
```

show mld snooping interface detail

The output for the `show mld snooping interface detail` command is identical to that for the `show mld snooping interface` command. For sample output, see "show mld snooping interface" on page 22088.

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Verifying MLD Snooping on Switches

Configuring MLD Snooping on a Switch VLAN with ELS Support (CLI Procedure)

show mld snooping membership

IN THIS SECTION

- Syntax | 22090
- Description | 22091
- Options | 22091
- Required Privilege Level | 22091
- Output Fields | 22091
- Sample Output | 22093
- Release Information | 22093

Syntax

```
show mld snooping membership
<brief | detail>
<interface logical-interface-name>
```

```
<vlan (vlan-id | vlan-name) >
```

Description

Display the multicast group membership information maintained by MLD snooping.

Options

none	Display the multicast group membership information for all VLANs on which MLD snooping is enabled.
brief | detail	(Optional) Display the specified level of output. The default is `brief`.
interface *interface-name*	(Optional) Display the multicast group membership information for the specified interface.
vlan (*vlan-id* | *vlan-name*)	(Optional) Display the multicast group membership for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 796 on page 22091 lists the output fields for the `show mld snooping membership` command. Output fields are listed in the approximate order in which they appear.

Table 796: show mld snooping membership Output Fields

Field Name	Field Description	Level of Output
`VLAN`	Name of the VLAN.	All
`Interfaces`	Interfaces that are members of the listed multicast group.	`brief`
`Tag`	Numerical identifier of the VLAN.	`detail`

Table 796: show mld snooping membership Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Router interfaces`	List of information about multicast-router interfaces: • Name of the multicast-router interface. • `static` or `dynamic`—Whether the multicast-router interface has been statically configured or dynamically learned. • `Uptime`—For static interfaces, amount of time since the interface was configured as a multicast-router interface or since the interface last flapped. For dynamic interfaces, amount of time since the first query was received on the interface or since the interface last flapped. • `timeout`—Seconds remaining before a dynamic multicast-router interface times out.	`detail`
`Group`	IP multicast address of the multicast group. The following information is provided for the multicast group: • Name of the interface belonging to the multicast group. • `Timeout`—Time (in seconds) left until a dynamically learned interface is removed from the multicast group if no MLD membership reports are received on the interface. This counter is reset to its maximum value when a membership report is received. • `Flags`—The lowest MLD version in use by a host that is a member of the group on the interface. If the flag `static` is included, the interface has been configured as static member of the multicast group. • `Receiver count`—Number of hosts on the interface that are members of the multicast group. This field appears only if `immediate-leave` is configured on the VLAN. • `Last reporter`—Last host to report membership for the multicast group. • `Include source`—Multicast source addresses from all MLDv2 membership reports received for the group on the interface.	`detail`

Sample Output

show mld snooping membership

```
user@host> show mld snooping membership
VLAN: mld_vlan
    2001:db8:ff1e::2010
        Interfaces: ge-1/0/30.0
    2001:db8:ff1e::2011
        Interfaces: ge-1/0/30.0
    2001:db8:ff1e::2012
        Interfaces: ge-1/0/30.0
    2001:db8:ff1e::2013
        Interfaces: ge-1/0/30.0
    2001:db8:ff1e::2014
        Interfaces: ge-1/0/30.0
```

show mld snooping membership detail

```
user@host> show mld snooping membership detail
VLAN: mld-vlan Tag: 100 (Index: 3)
    Router interfaces:
        ge-1/0/0.0 static Uptime: 00:57:13
   Group: 2001:db8:ff1e::2010
     ge-1/0/30.0 Timeout: 180 Flags: <V2-hosts>
     Last reporter: 2001:db8:2020:1:1:3
       Include source: 2001:db8:1:1::2
VLAN: mld-vlan1 Tag: 200 (Index: 4)
    Router interfaces:
        ae200.0 dynamic Uptime: 00:14:24 timeout: 244
   Group: 2001:db8:ff1e::2010
     ge-12/0/31.0 Timeout: 224 Flags: <V1-hosts>
     Last reporter: 2001:db8:2020:1:1:4
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding MLD Snooping

Example: Configuring MLD Snooping on SRX Series Devices

mld-snooping | **10783**

clear mld snooping membership | **3704**

show mld snooping statistics | **22098**

Verifying MLD Snooping on EX Series Switches (CLI Procedure)

Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)

show mld-snooping route

IN THIS SECTION

- Syntax | **22094**
- Description | **22094**
- Options | **22095**
- Required Privilege Level | **22095**
- Output Fields | **22095**
- Sample Output | **22096**
- Release Information | **22097**

Syntax

```
show mld-snooping route
<brief | detail>
<ethernet-switching | inet6>
<vlan (vlan-id | vlan-name)>
```

Description

Display multicast route information maintained by MLD snooping.

Options

none	Display route information for all VLANs on which MLD snooping is enabled.
brief | detail	(Optional) Display the specified level of output. The default is **brief**.
ethernet-switching	(Optional) Display information on Layer 2 IPv6 multicast routes. This is the default.
inet6	(Optional) Display information on Layer 3 IPv6 multicast routes.
vlan (*vlan-id* | *vlan-name*)	(Optional) Display route information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 797 on page 22095 lists the output fields for the `show mld-snooping route` command. Output fields are listed in the approximate order in which they appear.

Table 797: show mld-snooping route Output Fields

Field Name	Field Description
Table	Routing table ID for virtual routing instances.
Routing Table	Routing table ID for virtual routing instances.
VLAN	Name of the VLAN on which MLD snooping is enabled.
Group	Multicast IPv6 group address. Only the last 32 bits of the address are shown. The switch uses only these bits in determining multicast routes.
Next-hop	ID associated with the next-hop device.
Routing next-hop	ID associated with the Layer 3 next-hop device.

Table 797: show mld-snooping route Output Fields *(Continued)*

Field Name	Field Description
Interface or **Interfaces**	Name of the interface or interfaces in the VLAN associated with the multicast group.
Layer 2 next-hop	ID associated with the Layer 2 next-hop device.

Sample Output

show mld-snooping route

```
user@switch> show mld-snooping route

VLAN             Group          Next-hop
vlan1            ::0000:0001        1464
vlan1            ff00::
vlan10           ::0000:0002        1599
vlan10           ff00::
vlan11           ::0000:0002        1513
vlan11           ff00::
vlan12           ff00::
vlan13           ff00::
vlan14           ff00::
vlan15           ff00::
vlan16           ff00::
vlan17           ff00::
vlan18           ff00::
vlan19           ff00::
vlan2            ff00::
vlan20           ::0000:0002        1602
vlan20           ff00::
vlan3            ff00::
vlan4            ff00::
vlan5            ff00::
vlan6            ff00::
vlan7            ff00::
vlan8            ff00::
```

```
vlan9            ff00::
default          ff00::
```

show mld-snooping route detail

```
user@switch> show mld-snooping route detail
VLAN             Group          Next-hop
mld-vlan         ::0000:2010        1323
        Interfaces: ge-1/0/30.0
VLAN             Group          Next-hop
mld-vlan         ff00::             1317
        Interfaces: ge-1/0/0.0
VLAN             Group          Next-hop
mld-vlan         ::0000:0000        1317
        Interfaces: ge-1/0/0.0
VLAN             Group          Next-hop
mld-vlan1        ::0000:2010        1324
        Interfaces: ge-12/0/31.0
VLAN             Group          Next-hop
mld-vlan1        ff00::             1318
        Interfaces: ae200.0
VLAN             Group          Next-hop
mld-vlan1        ::0000:0000        1318
        Interfaces: ae200.0
```

show mld-snooping route inet6 detail

```
user@switch> show mld-snooping route inet6 detail
Routing table: 0
Group: ff05::1, 4001::11
    Routing next-hop: 1352
        vlan.2
    Interface: vlan.2, VLAN: vlan2, Layer 2 next-hop: 1387
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

show mld snooping membership | 22090

show mld snooping statistics | 22098

show mld-snooping vlans | 22100

Verifying MLD Snooping on EX Series Switches (CLI Procedure)

Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)

show mld snooping statistics

IN THIS SECTION

- Syntax | 22098
- Description | 22098
- Required Privilege Level | 22098
- Output Fields | 22099
- Sample Output | 22099
- Release Information | 22100

Syntax

```
show mld snooping statistics
```

Description

Display MLD snooping statistics.

Required Privilege Level

view

Output Fields

Table 798 on page 22099 lists the output fields for the `show mld snooping statistics` command. Output fields are listed in the approximate order in which they appear.

Table 798: show mld snooping statistics Output Fields

Field Name	Field Description
Bad length	MLD packet has illegal or bad length.
Bad checksum	MLD or IP checksum is incorrect.
Invalid interface	Packet was received through an invalid interface.
Not Local	Not used—always 0.
Receive unknown	Unknown MLD message type.
Timed out	Not used—always 0.
MLD Type	Type of MLD message (Query, Report, Leaves, or Other).
Received	Number of MLD packets received.
Transmitted	Number of MLD packets transmitted.
Recv Errors	Number of packets received that did not conform to the MLD version 1 (MLDv1) or MLDv2 standards.

Sample Output

show mld snooping statistics

```
user@host> show mld snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
```

```
Not local: 0 Receive unknown: 0 Timed out: 0

    MLD Type           Received      Transmitted    Recv Errors
    Queries:               74295                0              0
    Reports:            18148423                0       16333523
    Leaves:                    0                0              0
    Other:                     0                0              0
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding MLD Snooping

Example: Configuring MLD Snooping on SRX Series Devices

mld-snooping | **10783**

clear mld snooping statistics | **3705**

show mld snooping membership | **22090**

Verifying MLD Snooping on EX Series Switches (CLI Procedure)

show mld-snooping vlans

IN THIS SECTION

- Syntax | **22101**
- Description | **22101**
- Options | **22101**
- Required Privilege Level | **22101**
- Output Fields | **22101**
- Sample Output | **22102**
- Release Information | **22103**

Syntax

```
show mld-snooping vlans
<brief | detail>
<vlan vlan-name>
```

Description

Display MLD snooping information for a VLAN or for all VLANs.

Options

none Display MLD snooping information for all VLANs on which MLD snooping is enabled.

brief | detail (Optional) Display the specified level of output. The default is `brief`.

vlan *vlan-name* (Optional) Display MLD snooping information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 799 on page 22101 lists the output fields for the `show mld-snooping vlans` command. Output fields are listed in the approximate order in which they appear.

Table 799: show mld-snooping vlans Output Fields

Field Name	Field Description	Level of Output
`VLAN`	Name of the VLAN.	All levels
`Interfaces`	Number of interfaces in the VLAN.	`brief`
`Groups`	Number of groups in the VLAN.	`brief`

Table 799: show mld-snooping vlans Output Fields *(Continued)*

Field Name	Field Description	Level of Output
MRouters	Number of multicast-router interfaces in the VLAN.	brief
Receivers	Number of interfaces in the VLAN with a receiver for any group. Indicates how many interfaces might receive data because of MLD group membership.	brief
Tag	VLAN tag.	detail
vlan-interface	The Layer 3 interface, if any, associated with the VLAN.	detail
Interface	Name of the interface. The following information is provided for each interface: • tagged or untagged—Whether the interface accepts tagged packets (trunk mode and tagged-access mode ports) or untagged packets (access mode ports) • Groups—The number of multicast groups the interface belongs to • Reporters—The number of hosts on the interface that are current members of multicast groups. This field appears only when you configure "immediate-leave" on page 8454 on the VLAN. • Router—Indicates the interface is a multicast-router interface	detail

Sample Output

show mld-snooping vlans

```
user@host> show mld-snooping vlans
VLAN           Interfaces Groups MRouters Receivers
default                 0      0        0         0
v1                     11     50        0         0
v10                     1      0        0         0
v11                     1      0        0         0
```

```
v180                    3      0        1         0
v181                    3      0        0         0
v182                    3      0        0         0
```

show mld-snooping vlans vlan v10

```
user@host> show mld-snooping vlans vlan v10
 VLAN             Interfaces Groups MRouters Receivers
v10                        3      1        1         0       0
```

show mld-snooping vlans vlan vlan2 detail

```
user@host> show mld-snooping vlans vlan vlan2 detail

VLAN: vlan2, Tag: 2, vlan-interface: vlan.2
    Interface: ge-0/0/2.0, untagged, Groups: 5
    Interface: ge-0/0/4.0, tagged, Groups: 3, Router
```

show mld-snooping vlans detail

```
user@host> show mld-snooping vlans detail
VLAN: mld-vlan, Tag: 100
    Interface: ge-1/0/0.0, untagged, Groups: 0, Router
    Interface: ge-1/0/30.0, untagged, Groups: 1
    Interface: ge-1/0/33.0, untagged, Groups: 0
    Interface: ge-12/0/30.0, untagged, Groups: 0
VLAN: mld-vlan1, Tag: 200
    Interface: ge-1/0/31.0, untagged, Groups: 0
    Interface: ge-12/0/31.0, untagged, Groups: 1
    Interface: ae200.0, untagged, Groups: 0, Router
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

mld-snooping | 10783

show mld snooping membership | 22090

show mld-snooping route | 22094

show mld snooping statistics | 22098

Verifying MLD Snooping on EX Series Switches (CLI Procedure)

Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)

show mld statistics

IN THIS SECTION

- Syntax | 22104
- Syntax (MX Series) | 22105
- Syntax (ACX Series) | 22105
- Description | 22105
- Options | 22106
- Required Privilege Level | 22106
- Output Fields | 22106
- Sample Output | 22108
- show mld statistics (User Plane output for ACX7100) | 22110
- Release Information | 22110

Syntax

```
show mld statistics
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (MX Series)

```
show mld statistics
(<continuous> | <interface interface-name>)
<logical-system (all | logical-system-name)>
```

Syntax (ACX Series)

```
show mld statistics
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Description

Display information about Multicast Listener Discovery (MLD) statistics.

By default, Junos OS multicast devices collect statistics of received and transmitted MLD control messages that reflect currently active multicast group subscribers.

MX series devices also automatically maintain *continuous* MLD statistics globally on the device in addition to the default active subscriber statistics—these are persistent, continuous statistics of received and transmitted MLD control packets that account for both past and current multicast group subscriptions processed on the device.

With continuous statistics, you can see the total count of MLD control packets the MX series device processed since the last device reboot or `clear mld statistics continuous` command. The device collects and displays continuous statistics only for the fields shown in the `MLD packet statistics...` output section of this command, and does not display the `MLD Global statistics` section.

MX series devices that support continuous statistics maintain this information in a shared database and copy it to the backup Routing Engine at a configurable interval to avoid too much processing overhead on the Routing Engine. These actions preserve statistics counts across the following events or operations (which doesn't happen for the default active subscriber statistics):

- Routing daemon restart
- Graceful Routing Engine switchover (GRES)
- In-service software upgrade (ISSU)
- Line card reboot

You can change the default interval (300 seconds) using the `cont-stats-collection-interval` configuration statement at the `[edit routing-options multicast]` hierarchy level.

You can display either the default currently active subscriber statistics or continuous subscriber statistics (if supported), but not both at the same time. Include the `continuous` option to display continuous statistics, otherwise the command displays the statistics only for currently active subscribers.

Run the `clear mld statistics` command to clear the currently active subscriber statistics. On devices that support continuous statistics, run the clear command with the `continuous` option to clear all continuous statistics. You must run these commands separately to clear both types of statistics because the device maintains and clears the two types of statistics separately.

Options

none	Display MLD statistics for all interfaces. These statistics represent currently active subscribers.
continuous	(Optional) Display continuous MLD statistics that account for both past and current multicast group subscribers instead of the default statistics that only reflect currently active subscribers. This option is not available with the `interface` option for interface-specific statistics.
interface *interface-name*	(Optional) Display statistics about the specified interface. This option is not available with the `continuous` option.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 800 on page 22107 describes the output fields for the `show mld statistics` command. Output fields are listed in the approximate order in which they appear.

Table 800: show mld statistics Output Fields

Field Name	Field Description
MLD Packet Statistics...	Heading for MLD packet statistics for all interfaces or for the specified interface name. **NOTE**: Shows currently active subscriber statistics in this section by default, or when the command includes the `continuous` option, shows continuous, persistent statistics that account for all MLD control packets processed on the device.
Received	Number of received packets.
Sent	Number of transmitted packets.
Rx errors	Number of received packets that contained errors.
MLD Message type	Summary of MLD statistics. • **Listener Query (v1/v2)**—Number of membership queries sent and received. • **Listener Report (v1)**—Number of version 1 membership reports sent and received. • **Listener Done (v1/v2)**—Number of Listener Done messages sent and received. • **Listener Report (v2)**—Number of version 2 membership reports sent and received. • **Other Unknown types**—Number of unknown message types received. • **MLD v2 source required for SSM**—Number of MLD version 2 messages received that contained no source. • **MLD v2 mode not applicable for SSM**—Number of MLD version 2 messages received that did not contain a mode applicable for source-specific multicast (SSM).

Table 800: show mld statistics Output Fields *(Continued)*

Field Name	Field Description
MLD Global Statistics	Summary of MLD statistics for all interfaces. **NOTE**: These statistics are not supported or displayed with the `continuous` option. • **Bad Length**—Number of messages received with length errors so severe that further classification could not occur. • **Bad Checksum**—Number of messages received with an invalid IP checksum. No further classification was performed. • **Bad Receive If**—Number of messages received on an interface not enabled for MLD. • **Rx non-local**—Number of messages received from nonlocal senders. • **Timed out**—Number of groups that timed out as a result of not receiving an explicit leave message. • **Rejected Report**—Number of reports dropped because of the MLD group policy. • **Total Interfaces**—Number of interfaces configured to support IGMP.

Sample Output

show mld statistics

```
user@host> show mld statistics
MLD packet statistics for all interfaces
MLD Message type            Received       Sent  Rx errors
Listener Query (v1/v2)             0          2          0
Listener Report (v1)               0          0          0
Listener Done (v1/v2)              0          0          0
Listener Report (v2)               0          0          0
Other Unknown types                                      0
MLD v2 source required for SSM     2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length                         0
```

```
Bad Checksum                        0
Bad Receive If                      0
Rx non-local                        0
Timed out                           0
Rejected Report                     0
Total Interfaces                    2
```

show mld statistics interface

```
user@host> show mld statistics interface fe-1/0/1.0
MLD interface packet statistics for fe-1/0/1.0
MLD Message type            Received       Sent  Rx errors
Listener Query (v1/v2)             0          2          0
Listener Report (v1)               0          0          0
Listener Done (v1/v2)              0          0          0
Listener Report (v2)               0          0          0
Other Unknown types                                      0
MLD v2 source required for SSM     2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length                          0
Bad Checksum                        0
Bad Receive If                      0
Rx non-local                        0
Timed out                           0
Rejected Report                     0
Total Interfaces                    2
```

show mld statistics continuous

```
user@host> show mld statistics continuous
MLD packet statistics for all interfaces
MLD Message type            Received       Sent  Rx errors
Listener Query (v1/v2)             0          3          0
Listener Report (v1)               1          0          0
Listener Done (v1/v2)              1          0          0
Listener Report (v2)               1          0          0
Other Unknown types                                      0
MLD v2 unsupported type                                  0
```

```
MLD v2 source required for SSM                          0
MLD v2 mode not applicable for SSM                      0
```

show mld statistics (User Plane output for ACX7100)

Displays the following output in the user-plane for customer premise equipment (CPE).

```
user@up1> show mld statistics
MLD packet statistics for all interfaces
MLD Message type            Received       Sent  Rx errors
Listener Query (v1/v2)             0          3          0
Listener Report (v1)               0          0          0
Listener Done (v1/v2)              0          0          0
Listener Report (v2)               7          0          0
Other Unknown types                                      0
MLD v2 unsupported type                                  0
MLD v2 source required for SSM                           0
MLD v2 mode not applicable for SSM                       0

MLD Global Statistics
Bad Length                    1
Bad Checksum                  0
Bad Receive If               26
Rx non-local                  0
Timed out                     4
Rejected Report               0
Max Rx rate (pps)             4
Total Interfaces              2
```

Release Information

Command introduced before Junos OS Release 7.4.

`continuous` option added in Junos OS Release 19.4R1 for MX Series routers.

Command updated in Junos OS Release 24.4R1 to include multicast support monitoring for ACX7100 router interfaces.

RELATED DOCUMENTATION

clear mld statistics | 3707

show modem wireless firmware

IN THIS SECTION

- Syntax | 22111
- Description | 22111
- Options | 22111
- Required Privilege Level | 22111
- Output Fields | 22112
- Sample Output | 22114
- Release Information | 22115

Syntax

```
show modem wireless firmware interface-name
```

Description

Display modem firmware details for the LTE Mini-PIM.

Options

- *interface-name*—The LTE interface is `cl-x/0/0`, where *x* is the slot number in which the LTE Mini-PIM is installed.

Required Privilege Level

view

Output Fields

Table 801 on page 22112 lists some of the output fields for the `show modem wireless firmware` command. Output fields are listed in the approximate order in which they appear.

Table 801: show modem wireless firmware Output Fields

Field Name	Description
`LTE mPIM firmware details`	Displays the details of the firmware installed on the LTE Mini-PIM.
`Wireless modem firmware details`	Displays the details of the modem firmware.
`OTA status`	Displays the status of over-the-air (OTA) upgrade. The OTA upgrade can be enabled or disabled on the LTE Mini-PIM. OTA upgrade is disabled by default.

Table 801: show modem wireless firmware Output Fields *(Continued)*

Field Name	Description

`Status of SIM`	• `Number of SIM`—Number of SIM cards installed.
	• `Slot of active`—The slot in which the active SIM card is installed.
	• `SIM state`—Indicates whether the SIM card is present in the slot.
	• `Modem PIN security status`—Indicates the security status of the SIM. If the SIM is locked by using the `request modem wireless sim-lock enable` command, then the security status is displayed as enabled.
	• `SIM status`—Status of the Subscriber Identity Module (SIM) in the LTE Mini-PIM. The status can be one of the following:
	  • `SIM Okay`
	  • `No status`—The device is being powered on or powered off, or the SIM card has been removed from the slot.
	  • `SIM init failure`—There is a problem with the SIM; the SIM might need to be replaced.
	  • `SIM locked`
	  • `PIN1 blocked`—Obtain a PIN unblocking key (PUK) to unblock the SIM.
	  • `PIN1 rejected`—The wrong PIN was entered.
	  • `PIN2 rejected`—The wrong PIN was entered.
	  • `Network rejected`
	• `SIM user operation needed`—Action required by the user. This can be one of the following:
	  • `No op`—No user operation required.
	  • `Enter PIN`—Enter the personal identification number (PIN) to unlock the SIM card.
	  • `Enter PUK`—Enter the PUK to unblock the SIM card.
	• `Retries remaining`—If the value of `SIM user operation needed` is `Enter PIN`, this is the number of PIN unlock attempts remaining before the modem is blocked. If the PIN is entered incorrectly three consecutive times, the SIM card is blocked.

Table 801: show modem wireless firmware Output Fields *(Continued)*

Field Name	Description
	If the value of `SIM user operation needed` is `Enter PUK`, this is the number of unblock attempts remaining before the modem is unusable. If the PUK is entered incorrectly ten times, the SIM card must be returned to the service provider for reactivation.

Sample Output

show modem wireless firmware

```
user@host> show modem wireless firmware cl-1/0/0
 LTE mPIM firmware details
  Product name: Junos LTE mPIM
  Serial number: AG50071852
  Hardware version: AcceleratedConcepts/sprite
  Firmware version: 17.4.3
  MAC: 00:00:5e:00:a0:61
  System uptime: 3430 seconds
Wireless modem firmware details
  Modem firmware version: 9999999_9904609_SWI9X30C_02.23.00.00_00_GENERIC_002.018_000
  Modem Firmware build date: 22/10/2016
  Card type: MC7430
  Modem manufacturer: Sierra Wireless, Inc
  Hardware version: 1.0
  Power & Temperature: Normal 3343 mV, Normal 30.00 C
OTA status
  State: Enabled
  New firmware available: No
Number of SIM: 2
Slot of active: 2
Status of SIM 1
  SIM state: SIM present
  Modem PIN security status: Disabled
  SIM status: SIM Okay
  SIM user operation needed: No Op
  Retries remaining: 3
Status of SIM 2
  SIM state: SIM present
```

```
Modem PIN security status: Disabled
SIM status: SIM Okay
SIM user operation needed: No Op
Retries remaining: 3
```

Release Information

Command introduced in Junos OS 15.1X49-D100

RELATED DOCUMENTATION

show modem wireless network | 22115

show modem wireless network

IN THIS SECTION

- Syntax | 22115
- Description | 22115
- Options | 22116
- Required Privilege Level | 22116
- Output Fields | 22116
- Sample Output | 22118
- Release Information | 22119

Syntax

```
show modem wireless network interface-name
```

Description

Display the status of the modem and the status of the network connection for the LTE Mini-PIM.

Options

- *interface-name*—The LTE interface is `cl-x/0/0`, where *x* is the slot number in which the LTE Mini-PIM is installed.

Required Privilege Level

view

Output Fields

Table 802 on page 22116 lists some of the output fields for the `show modem wireless network` command. Output fields are listed in the approximate order in which they appear.

Table 802: show modem wireless network Output Fields

Field Name	Field Description
`Current Modem Status`	Status of the modem on the Mini-PIM. The status can be one of the following states: • Disconnected • Calling • Connected

Table 802: show modem wireless network Output Fields *(Continued)*

Field Name	Field Description
`Current Service Status`	Status of the network connection. The status can be one of the following states: • Normal • Emergency Call Only • No Service Available • Unable To Register • Forbidden PLMN • Forbidden Area • Roaming Not Permitted • Account Not Permitted • Modem Not Permitted • Unknown IMSI • Authentication Failure
`Current Service Type`	One of the following: • Circuit switched (CS) • Packet switched (PS) • Combo (CS, PS) • Invalid

Table 802: show modem wireless network Output Fields *(Continued)*

Field Name	Field Description
`Current Service Mode`	One of the following: • Unknown • LTE • DC-HSPA+ • HSPA+ • HSPA • UMTS
`Current Band`	Current radio band in use.
`Mobile Country Code (MCC)`	Number that uniquely identifies the country.
`Mobile Network Code`	Number that uniquely identifies a network within a country.

Sample Output

show modem wireless network

```
user@host> show modem wireless network cl-1/0/0
LTE Connection details
Connected time: 147
IP: 172.16.52.4
Gateway: 172.16.52.5
DNS: 123.123.123.123
Input bps: 0
Output bps: 0
Bytes Received: 1308
Bytes Transferred: 1164
Packets Received: 10
Packets Transferred: 10
```

```
Wireless Modem Network Info
Current Modem Status: Connected
Current Service Status: Normal
Current Service Type: PS
Current Service Mode: LTE
Current Band: B3
Network: UNICOM
Mobile Country Code (MCC): 460
Mobile Network Code (MNC): 1
Location Area Code (LAC): 65534
Routing Area Code (RAC): 0
Cell Identification: 4865903
Access Point Name (APN): abcde
Public Land Mobile Network (PLMN): CHN-UNICOM
Physical Cell ID (PCI): 333
International Mobile Subscriber Identification (IMSI): ***************
International Mobile Equipment Identification (IMEI/MEID): ***************
Integrate Circuit Card Identity (ICCID): 89860114721100697502
Reference Signal Receiving Power (RSRP): -97
Reference Signal Receiving Quality (RSRQ): -16
Signal to Interference-plus-Noise Ratio (SiNR): 0
Signal Noise Ratio (SNR): 0
Energy per Chip to Interference (ECIO): 0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

show modem wireless profiles | 22120

show modem wireless firmware | 22111

show modem wireless profiles

IN THIS SECTION

- Syntax | 22120
- Description | 22120
- Options | 22120
- Required Privilege Level | 22120
- Output Fields | 22120
- Sample Output | 22121
- Release Information | 22122

Syntax

```
show modem wireless profiles interface-name slot slot-number
```

Description

Display the profiles configured on the LTE Mini-PIM.

Options

- *interface-name*—The LTE interface is `cl-x/0/0`, where *x* is the slot number in which the LTE Mini-PIM is installed.
- *slot-number*—The slot in which the SIM card is inserted. The value can be either 1 or 2.

Required Privilege Level

view

Output Fields

Table 803 on page 22121 lists some of the output fields for the `show modem wireless profiles` command. Output fields are listed in the approximate order in which they appear.

Table 803: show modem wireless profiles Output Fields

Field Name	Field Description
`Max profiles`	The maximum number of profiles available for each SIM card. This value is always 16. The LTE Mini-PIM supports two SIM cards and so you can configure a total of 32 profiles, although only one profile can be active at a time.
`Default profile Id`	The profile used to connect to the network when there is no profile selected. The default profile ID is always 1.
`Profile details`	• `Username`—The username provided by the service provider. • `Password`—The password provided by the service provider. • `Access point name (APN)`—The APN provided by the service provider. • `Authentication`—The protocol used for authentication.

Sample Output

show modem wireless profiles

```
user@host> show modem wireless profiles cl-1/0/0 slot 1
Profile details
  Max profiles: 16
  Default profile Id: 1

Profile 1: ACTIVE
  Valid: TRUE
  Access point name (APN): ctnet
  Authentication: None
Profile 2: Inactive
  Valid: TRUE
  Username: myuser
  Password: 123456
  Access point name (APN): testapn
  Authentication: PAP
```

```
Profile 3: Invalid
Profile 4: Invalid
Profile 5: Invalid
Profile 6: Invalid
Profile 7: Invalid
Profile 8: Invalid
Profile 9: Invalid
Profile 10: Invalid
Profile 11: Invalid
Profile 12: Invalid
Profile 13: Invalid
Profile 14: Invalid
Profile 15: Invalid
Profile 16: Invalid
```

Release Information

Command introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

show modem wireless firmware | 22111

show modem wireless network | 22115

show monitor security flow

IN THIS SECTION

- Syntax | 22123
- Description | 22123
- Required Privilege Level | 22123
- Output Fields | 22123
- Sample Output | 22124
- Release Information | 22125

Syntax

```
show monitor security flow
```

Description

Display information about the security flow session monitoring.

Required Privilege Level

view

Output Fields

Lists the output fields for the `show monitor security flow` command. Output fields are listed in the approximate order in which they appear.

Table 804: show monitor security flow Output Fields

Field Name	Field Description
`Monitor security flow session status`	State of the security flow session monitoring: `active` or `inactive`.
`Monitor security flow trace file`	Name of the file for monitoring output.

Table 804: show monitor security flow Output Fields *(Continued)*

Field Name	Field Description
`Monitor security flow filters`	• `conn-tag`—Tag that uniquely identifies a session. The session key is expanded to include this tuple. • `Destination Address`—Address of the destination to be matched. • `Destination Port`—Name of the destination port to be matched. • `Interface Name`—Interface name to be matched. • `Logical System Name`—Logical system name to be matched. • `Name`—Name of the security flow filter. • `Protocol`—Name of the protocol to be matched. • `Source Address`—Address of the source to be matched. • `Source Port`—Name of the source port to be matched. • `Status`—State of the security flow filter: `active` or `inactive`.

Sample Output

show monitor security flow

```
user@host>show monitor security flow

Monitor security flow session status: Active
Monitor security flow trace file: flow
Monitor security flow filters:
  Name: server-sql
    Status: Active
    source: 10.2.2.1 (port *), destination: 10.20.30.40 (port 1433)
    protocol: TCP
    conn-tag: 0
  Name: internet-access
    Status: Active
    source: * (port *), destination: * (port 80)
```

```
    protocol: TCP
    conn-tag: 0
```

Release Information

Command introduced in Junos OS Release 12.1X46-D10. This topic was updated to include the flow session conn-tag filter in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

Monitoring Security Flow Sessions Overview

monitor security flow filter | 10916

monitor security flow start | 10918

clear monitor security flow filter | 3709

show mpls abstract-hop-membership

IN THIS SECTION

- Syntax | 22125
- Description | 22126
- Options | 22126
- Required Privilege Level | 22126
- Output Fields | 22126
- Sample Output | 22127
- Release Information | 22127

Syntax

```
show mpls abstract-hop-membership
<abstract-hop-name>
```

```
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display MPLS abstract hop membership tables for each abstract hop configured on the device.

Options

none	(Optional) Display the MPLS abstract hop membership table for all the configured abstract hops on the router.
abstract-hop-name	(Optional) Display the MPLS abstract hop membership table for the specified abstract hop.
instance *instance-name*	(Optional) Display the MPLS abstract hop membership table for the specified instance. If *instance-name* is omitted, information is displayed for the master instance.
logical-system (all | *logical-system-name*)	(Optional) Display the MPLS abstract hop membership table for all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 805 on page 22126 describes the output fields for the `show mpls abstract-hop-membership` command. Output fields are listed in the approximate order in which they appear.

Table 805: show mpls abstract-hop-membership Output Fields

Field Name	Field Description
`Abstract hop`	Name of the abstract hop.
`Credibility`	Credibility value associated with the interior gateway protocol in use.

Table 805: show mpls abstract-hop-membership Output Fields *(Continued)*

Field Name	Field Description
Address	IP address of the abstract hop member nodes.

Sample Output

show mpls abstract-hop-membership

```
user@host> show mpls abstract-hop-membership
Abstract hop: ah1
   Credibility: 0
Address: 127.0.0.6
Address: 127.0.0.1
Address: 127.0.0.2
Address: 127.0.0.3

Abstract hop: ah2
   Credibility: 0
Address: 127.0.0.6
Address: 127.0.0.3
Address: 127.0.0.4

Abstract hop: ah3
   Credibility: 0
Address: 127.0.0.6
Address: 127.0.0.3
Address: 127.0.0.5
```

Release Information

Command introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Example: Configuring Abstract Hops for MPLS LSPs

abstract-hop | 1625
constituent-list | 4571
show mpls lsp abstract-computation | 22203

show mpls admin-groups

IN THIS SECTION

- Syntax | 22128
- Syntax (EX Series Switches) | 22128
- Description | 22128
- Options | 22129
- Required Privilege Level | 22129
- Output Fields | 22129
- Sample Output | 22129
- Release Information | 22130

Syntax

```
show mpls admin-groups
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show mpls admin-groups
```

Description

Display information about configured Multiprotocol Label Switching (MPLS) administrative groups.

Options

none	Display information about the configured MPLS administrative groups.
instance *instance-name*	(Optional) Display MPLS administrative group information for the specified instance. If *instance-name* is omitted, MPLS administrative group information for the master instance is displayed.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 806 on page 22129 describes the output fields for the `show mpls admin-groups` command. Output fields are listed in the approximate order in which they appear.

Table 806: show mpls admin-groups Output Fields

Field Name	Field Description
Group	Name of the administrative group.
Bit index	Value assigned to the administrative group.

Sample Output

show mpls admin-groups

```
user@host> show mpls admin-groups
Group        Bit index
black         3
blue          2
gold          1
green         0
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` *`instance-name`* option added in Junos OS Release 15.1.

show mpls association

IN THIS SECTION

- Syntax | 22130
- Description | 22130
- Options | 22131
- Required Privilege Level | 22131
- Output Fields | 22131
- Sample Output | 22132
- Release Information | 22132

Syntax

```
show mpls association (iif incoming-interface | oif outgoing-interface)
<logical-system (all | logical-system-name)>
```

Description

Display the Multiprotocol Label Switching (MPLS) label-switched paths (LSPs) based on the association with an incoming or outgoing LSP interface. The command output displays the list of RSVP-TE LSPs carrying traffic in and out of the same interface.

Options

iif *incoming-interface-name*	Display list of RSVP-TE LSPs that share the specified incoming interface to bring in traffic. This option works on transit label-switching routers (LSRs) and egress label edge routers (LERs).
oif *outgoing-interface-name*	Display list of RSVP-TE LSPs that share the specified outgoing interface to carry out traffic. This option works on ingress LERs and transit LSRs.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 807 on page 22131 describes the output fields for the `show mpls association` command. Output fields are listed in the approximate order in which they appear.

Table 807: show mpls association Output Fields

Field Name	Field Description
To	Destination IP address of the corresponding LSP.
From	Source IP address of the corresponding LSP.
State	State of the corresponding LSP handled by this RSVP session: `Up`, `Dn` (down), or `Restart`.
LSPname	Name of the LSP.

Sample Output

show mpls association iif

```
user@host> show mpls association iif ge-0/0/0.0
To              From            State       LSPname
128.102.174.121 128.102.180.21  Up          LSP-ABC
128.102.174.121 128.102.180.21  Up          LSP-ABC1
128.102.174.121 128.102.180.21  Up          LSP-ABC2
128.102.174.121 128.102.180.21  Up          LSP-ABC3
Total 4 displayed, Up 4, Down 0
```

show mpls association oif

```
user@host> show mpls association oif ge-0/0/0.0
To              From            State       LSPname
128.102.174.121 128.102.180.21  Up          LSP-ABC
128.102.174.121 128.102.180.21  Up          LSP-ABC1
128.102.174.121 128.102.180.21  Up          LSP-ABC2
128.102.174.121 128.102.180.21  Up          LSP-ABC3
```

Release Information

Command introduced in Junos OS Release 16.1 .

RELATED DOCUMENTATION

show mpls correlation nexthop-id | 22151

show mpls call-admission-control

IN THIS SECTION

Syntax | 22133

- Syntax (EX Series Switches) | 22133
- Description | 22133
- Options | 22133
- Additional Information | 22134
- Required Privilege Level | 22134
- Output Fields | 22134
- Sample Output | 22135
- Release Information | 22136

Syntax

```
show mpls call-admission-control
<instance instance-name>
<logical-system (all | logical-system-name)>
<lsp-name>
```

Syntax (EX Series Switches)

```
show mpls call-admission-control
<lsp-name>
```

Description

Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.

Options

none	Display CAC information for all LSPs.
instance *instance-name*	(Optional) Display MPLS LSP CAC information for the specified instance. If *instance-name* is omitted, MPLS LSP CAC information for the master instance is displayed.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
lsp-name	(Optional) Display CAC information for the specified LSP only.

Additional Information

The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.

Required Privilege Level

view

Output Fields

Table 808 on page 22134 describes the output fields for the `show mpls call-admission-control` command. Output fields are listed in the approximate order in which they appear.

Table 808: show mpls call-admission-control Output Fields

Field Name	Field Description
Available bandwidth	Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at **ct0**) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type.
Layer2 connections	Different Layer 2 connections that had some bandwidth requirement and were admitted into an LSP path.
LSP name	LSP pathname.
Neighbor address	Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits.
Circuit	Interface name and circuit information.

Table 808: show mpls call-admission-control Output Fields *(Continued)*

Field Name	Field Description
Primary	LSP's primary standby path.
Standby	LSP's secondary standby path.
VC bandwidth	Bandwidth constraints associated with a Layer 2 circuit route.

Sample Output

show mpls call-admission-control

```
user@host# show mpls call-admission-control

LSP name: pro1-be
 *Primary
  Available bandwidth: 0bps

LSP name: pro1-be-1
 *Primary
  Available bandwidth: 60kbps

LSP name: pro1-be-gold
 *Primary
  Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
    Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
    VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>

LSP name: pro1-be-gold-2
 *Primary
  Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: pro1-be-silver
 *Primary   prim1
  Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
    Layer2 connections:
```

```
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
  Standby   sec1
  Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
    Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` *`instance-name`* option added in Junos OS Release 15.1.

show mpls container-lsp

IN THIS SECTION

- Syntax | 22136
- Description | 22137
- Options | 22137
- Required Privilege Level | 22138
- Output Fields | 22138
- Sample Output | 22144
- Release Information | 22146

Syntax

```
show mpls container-lsp
<brief | detail | extensive | terse>
<count-active-routes>
<defaults>
<descriptions>
<down | up>
```

```
<egress>
<ingress>
<logical-system (all | logical-system-name)>
<name name>
<statistics>
<transit>
<unidirectional>
```

Description

Display information about configured and active Multiprotocol Label Switching (MPLS) container label-switched paths (LSPs).

Options

none Display standard information about all configured and active member LSPs of the container LSP.

brief | detail | extensive | terse (Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

count-active-routes (Optional) Show active routes for the container LSP.

defaults (Optional) Display the default settings of the container LSP.

descriptions (Optional) Display the container LSP descriptions. To view this information, you must configure the description statement at the `[edit protocol mpls container-lsp]` hierarchy level. Only the LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

down | up (Optional) Display only LSPs that are inactive or active, respectively.

egress (Optional) Display the LSPs ending at this device.

> **NOTE**: The `egress` option displays all the LSPs including regular LSPs, members of container LSPs, and transit LSPs. This is an expected behavior for all platforms.

ingress (Optional) Display the member LSPs originating from this device.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
name *name*	(Optional) Display information about the specified LSP or group of LSPs.
statistics	(Optional) Display accounting information about LSPs. Statistics are not available for LSPs on the egress routing device, because the penultimate routing device in the LSP sets the label to 0. Also, as the packet arrives at the egress routing device, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.
transit	(Optional) Display LSPs transiting this routing device.
unidirectional	(Optional) Display unidirectional LSP information.

Required Privilege Level

view

Output Fields

Table 809 on page 22138 describes the output fields for the `show mpls container-lsp` command. Output fields are listed in the approximate order in which they appear.

Table 809: show mpls container-lsp Output Fields

Field Name	Field Description	Level of Output
`Ingress LSP`	Information about the member LSPs on the ingress routing device. Each LSP has one line of output.	All levels
`Container LSP name`	Name of the container LSP.	All levels
`Member LSP count`	Number of member LSPs in the container LSP.	All levels
`To`	Destination (egress routing device) of the session.	brief
`From`	Source (ingress routing device) of the session.	brief detail

Table 809: show mpls container-lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the LSP handled by this RSVP session: • Up • Dn (down) • Restart	brief detail
Rt	Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	brief
P	Path. An asterisk (*) underneath this column indicates that the LSP is a primary path.	brief
ActivePath	(Ingress LSP) Name of the active path: Primary or Secondary.	detail extensive
LSPname	Name of the member LSP.	brief detail
Egress LSP	Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output.	All levels
Transit LSP	Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
Min LSPs	Minimum number of member LSPs. Default: 1	extensive
Max LSPs	Number of member LSPs that the container LSP can have at maximum. Default: 64 (due to ECMP limit)	extensive

Table 809: show mpls container-lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Aggregate bandwidth`	Sum of the bandwidths of all member LSPs.	extensive
`Sampled Aggregate bandwidth`	Sum of the Maximum Average Bandwidth utilization of all member LSPs	extensive
`NormalizeTimer`	Duration between two normalization events. When not configured, 21600 seconds (6 hours) is set as the default value.	extensive
`NormalizeThreshold`	Change in aggregate LSP utilization to trigger splitting or merging expressed in percentage.	extensive
`Max Signaling BW`	Maximum bandwidth used to signal LSPs after a normalization event. Default value is 0 bps. When not configured, the value is inherited from the splitting bandwidth configuration. **NOTE**: Between two normalization events, when auto-bandwidth adjustment happens, the per-LSP auto-bandwidth configuration and thresholds are used, instead of the maximum signaling bandwidth threshold.	extensive
`Min Signaling BW`	Minimum bandwidth used to signal LSPs after a normalization event. Default value is 0 bps. When not configured, the value is inherited from the merging bandwidth configuration. **NOTE**: Between two normalization events, when auto-bandwidth adjustment happens, the per-LSP auto-bandwidth configuration and thresholds are used, instead of the minimum signaling bandwidth threshold.	extensive
`Splitting BW`	Bandwidth used for LSP splitting and merging. Default value is 0 bps. When not configured, the value is inherited from the auto-bandwidth maximum bandwidth configuration.	extensive

Table 809: show mpls container-lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Merging BW`	Bandwidth used for LSP splitting and merging. Default value is 0 bps. When not configured, the value is inherited from the auto-bandwidth minimum bandwidth configuration.	extensive
`LSPtype`	Type of LSP.	extensive
`LoadBalance`	Loadbalanced bandwidth.	extensive
`MinBW`	Minimum LSP bandwidth in bps related to auto-bandwidth.	extensive
`AdjustTimer`	Total amount of time in seconds allowed before LSP bandwidth adjustment take place. Range: 300 through 315360000 seconds	extensive
`Max AvgBW util`	Current value of the actual maximum average bandwidth utilization in bps.	extensive
`Overflow limit`	Threshold overflow limit.	extensive
`Underflow limit`	Threshold underflow limit.	extensive
`Priorities`	Setup priority and hold priority values. For setup priority, 0 and 7 is the highest and lowest priority, respectively. When not explicitly configured, 7 and 0 are set as the default values for the setup priority and hold priority, respectively.	extensive
`SmartOptimizeTimer`	Time in seconds allowed before path reoptimization.	extensive

Table 809: show mpls container-lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Computed ERO`	Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (S) or loose (L).	extensive

Table 809: show mpls container-lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Received RRO`	Received record route. RRO is a series of hops, each with an address followed by a flag. In most cases, the received RRO is the same as the computed ERO. If the received RRO is different from the computed ERO, there is a topology change in the network, and the route is taking a detour. The following flags identify the protection capability and status of the downstream node: • `0x01`—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the local protection flag was set in the SESSION_ATTRIBUTE object of the corresponding path message. • `0x02`—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously). • `0x03`—Combination of `0x01` and `0x02`. • `0x04`—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section. • `0x08`—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the local protection available bit is set but the node protection bit is cleared. • 0x09—Detour is established. Combination of `0x01` and `0x08`. • `0x10`—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted. • `0x20`—Node ID. Indicates that the address specified in the RRO's IPv4 or IPv6 sub-object is a node ID address, which refers to the router address or router ID. Nodes must use the same address consistently.	extensive

Table 809: show mpls container-lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• 0xb—Detour is in use. Combination of 0x01, 0x02, and 0x08.	
Created	Date and time the LSP was created.	extensive

Sample Output

show mpls container-lsp

```
user@host> show mpls container-lsp
Ingress LSP: 1 sessions
Container LSP name                              Member LSP count
test                                                          2
To              From            State Rt P     ActivePath       LSPname
10.255.107.76   10.255.107.78   Up     0 *                      test-1
10.255.107.76   10.255.107.78   Up     0 *                      test-2
Total 2 displayed, Up 2, Down 0

naling Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls container-lsp extensive

```
user@host> show mpls container-lsp extensive
Ingress LSP: 1 sessions
Container LSP name: test, Member count: 2
 Normalization
  Min LSPs: 2, Max LSPs: 64, Aggregate bandwidth: 0bps
  NormalizeTimer: 1800 secs, NormalizeThreshold: 0%
  Max Signaling BW: 2kbps, Min Signaling BW: 2kbps, Splitting BW: 5Mbps, Merging BW: 2kbps
  Normalization in 989 second(s)
10.255.107.76
```

```
  From: 10.255.107.78, State: Up, ActiveRoute: 0, LSPname: test-1
  ActivePath:  (primary)
  LSPtype: Dynamic Configured, Penultimate hop popping
  LoadBalance: Random
  Autobandwidth
  MinBW: 1000bps
  AdjustTimer: 300 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 89 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up, No-decrement-ttl
    Priorities: 7 0
    Bandwidth: 1000bps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
 1.3.0.2 S 1.7.0.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          1.3.0.2 1.7.0.1
   11 Jul 13 20:08:26.613 Make-before-break: Switched to new instance
   10 Jul 13 20:08:04.360 Record Route:  1.3.0.2 1.7.0.1
    9 Jul 13 20:08:04.360 Up
    8 Jul 13 20:08:04.360 Automatic Autobw adjustment succeeded: BW changes from 2000 bps to
1000 bps
    7 Jul 13 20:08:04.314 Originate make-before-break call
    6 Jul 13 20:08:04.314 CSPF: computation result accepted  1.3.0.2 1.7.0.1
    5 Jul 13 20:05:02.423 Selected as active path
    4 Jul 13 20:05:02.422 Record Route:  1.3.0.2 1.7.0.1
    3 Jul 13 20:05:02.421 Up
    2 Jul 13 20:05:02.376 Originate Call
    1 Jul 13 20:05:02.376 CSPF: computation result accepted  1.3.0.2 1.7.0.1
  Created: Sat Jul 13 20:03:03 2013
10.255.107.76
From: 10.255.107.78, State: Up, ActiveRoute: 0, LSPname: test-2
  ActivePath:  (primary)
  LSPtype: Dynamic Configured, Penultimate hop popping
  LoadBalance: Random
  Autobandwidth
  MinBW: 1000bps
  AdjustTimer: 300 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 89 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
```

```
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                    State: Up, No-decrement-ttl
    Priorities: 7 0
    Bandwidth: 1000bps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
 1.2.0.2 S 1.4.0.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          1.2.0.2 1.4.0.2
   11 Jul 13 20:08:05.363 Make-before-break: Switched to new instance
   10 Jul 13 20:08:04.450 Record Route:  1.2.0.2 1.4.0.2
    9 Jul 13 20:08:04.449 Up
    8 Jul 13 20:08:04.449 Automatic Autobw adjustment succeeded: BW changes from 2000 bps to
1000 bps
    7 Jul 13 20:08:04.327 Originate make-before-break call
    6 Jul 13 20:08:04.327 CSPF: computation result accepted  1.2.0.2 1.4.0.2
    5 Jul 13 20:05:00.849 Selected as active path
    4 Jul 13 20:05:00.841 Record Route:  1.3.0.2 1.7.0.1
    3 Jul 13 20:05:00.831 Up
    2 Jul 13 20:05:00.513 Originate Call
    1 Jul 13 20:05:00.502 CSPF: computation result accepted  1.3.0.2 1.7.0.1
  Created: Sat Jul 13 20:03:03 2013
Total 2 displayed, Up 2, Down 0

Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Release Information

Command introduced in Junos OS Release 14.2.

Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.

RELATED DOCUMENTATION

request mpls container-lsp | 14689

clear mpls container-lsp | 3711

show mpls context-identifer

IN THIS SECTION

- Syntax | 22147
- Description | 22147
- Options | 22147
- Required Privilege Level | 22148
- Output Fields | 22148
- Sample Output | 22149
- Sample Output | 22149
- Release Information | 22149

Syntax

```
show mpls context-identifer
<brief | detail>
<logical-system (all | logical-system-name)>
<primary>;
<protector>;
```

Description

Display information about configured egress protection context identifiers.

Options

none	Display standard information about egress protection.
brief | detail	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
primary	(Optional) Perform this operation on the primary node.

protector (Optional) Perform this operation on the protector node.

Required Privilege Level

view

Output Fields

Table 810 on page 22148 describes the output fields for the `show mpls egress-protection detail` command. Output fields are listed in the approximate order in which they appear.

Table 810: show mpls lsp Output Fields

Field Name	Field Description	Level of Output
ID	Context identifier.	All levels
Type	Indicates node type: protector or primary	All levels
Metric	MPLS cost value of the context identifier route. This route appears in inet.0 on the protector and primary nodes. On the protector node, the metric is a larger number.	All levels
Mode	Indicates "advertise-mode" on page 2018: proxy or alias	`detail`
Context table	Name of the MPLS routing table created for egress protection.	All levels
Context LSPs	Names of the LSPs that have egress protection configured. Loopback interface addresses of the devices from which the LSPs are originated.	`detail`
Total	Total number of primary and protector nodes.	All levels

Table 810: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Primary	Number of primary nodes.	All levels
Protector	Number of protector nodes.	All levels

Sample Output

show mpls context-identifier detail (Protector)

```
user@host> show mpls context-identifier detail
ID: 166.1.3.1
  Type: protector, Metric: 16777215, Mode: alias
  Context table: __166.1.3.1__.mpls.0, Label out: 299968
```

Sample Output

show mpls context-identifier detail (Primary)

```
user@host> show mpls context-identifier detail

ID: 166.1.3.1
  Type: primary, Metric: 1, Mode: alias

Total 1, Primary 1, Protector 0
```

Release Information

Command introduced in Junos OS Release 11.4R3.

RELATED DOCUMENTATION

Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP

Example: Configuring MPLS Egress Protection for Layer 3 VPN Services

show mpls correlation label

IN THIS SECTION

- Syntax | 22150
- Description | 22150
- Options | 22150
- Required Privilege Level | 22151
- Release Information | 22151

Syntax

```
show mpls correlation label label-value
<brief | detail | extensive | terse>
<descriptions>
<logical-system (all | logical-system-name)>
```

Description

Display the correlation information for the Multiprotocol Label Switching (MPLS) label-switched path (LSP) with the owner of the label.

Options

label-value	Display information about the specified label.
brief | detail | extensive | terse	(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

descriptions (Optional) Display the LSP descriptions. To view this information, you must configure the description statement at the `[edit protocol mpls lsp]` hierarchy level. Only the LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

logical-system (all | *logical-system-name*) (Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 15.2 .

RELATED DOCUMENTATION

show mpls correlation nexthop-id | 22151

show mpls association | 22130

show mpls correlation nexthop-id

IN THIS SECTION

- Syntax | 22152
- Description | 22152
- Options | 22152
- Required Privilege Level | 22152
- Output Fields | 22152
- Sample Output | 22153
- Release Information | 22153

Syntax

```
show mpls correlation nexthop-id nexthop-id
<descriptions>
<logical-system (all | logical-system-name)>
```

Description

Display the correlation information for the Multiprotocol Label Switching (MPLS) label-switched path (LSP) with the owner of the next-hop ID.

Options

nexthop-id	Display information about the specified next-hop ID.
descriptions	(Optional) Display the LSP descriptions. To view this information, you must configure the description statement at the `[edit protocol mpls lsp]` hierarchy level. Only the LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 811 on page 22152 describes the output fields for the `show mpls correlation nexthop-id` command. Output fields are listed in the approximate order in which they appear.

Table 811: show mpls correlation nexthop-id Output Fields

Field Name	Field Description
LSP name	Name of the LSP associated with the specified next-hop ID.

Sample Output

show mpls correlation nexthop-id

```
user@host> show mpls correlation nexthop-id nexthop-id
LSP name: LSP-ABC
```

Release Information

Command introduced in Junos OS Release 16.1 .

RELATED DOCUMENTATION

show mpls association | 22130

show mpls cspf

IN THIS SECTION

- Syntax | 22154
- Syntax (EX Series Switches) | 22154
- Description | 22154
- Options | 22154
- Required Privilege Level | 22154
- Output Fields | 22154
- Sample Output | 22156
- Release Information | 22156

Syntax

```
show mpls cspf
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show mpls cspf
```

Description

Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.

Options

none	Display MPLS CSFP statistics.
instance *instance-name*	(Optional) Display MPLS CSPF information for the specified instance. If *instance-name* is omitted, MPLS CSPF information for the primary instance is displayed.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 812 on page 22155 describes the output fields for the `show mpls cspf` command. Output fields are listed in the approximate order in which they appear.

Table 812: show mpls cspf Output Fields

Field Name	Field Description
Queue length	Number of LSPs queued for automatic path computation.
current	Current queue length.
maximum	Maximum queue length (high-water mark).
dequeued	Number of terminated computation attempts.
Paths	Counters for label-switched path computations.
total	Sum of the next four fields.
successful	Number of path computations that were successfully completed.
no route	Number of path computations that failed because the destination is unreachable.
Sys Error	Number of path computations that failed because of lack of memory.
CSPFs	Total number of CSPF computations. A single path might require multiple CSPF computations.
Time	Time, in seconds, required to perform the label-switched path computation.
Total	Total amount of time consumed by the CSPF path computation algorithm.
CSPFs	Total number of CSPF computations.
Avg per CSPF	Average amount of time required for each CSPF computation.
% of rpd	Percentage of routing process CPU used in the CSPF computation.

Sample Output

show mpls cspf

```
user@host> show mpls cspf
CSPF statistics
Queue length   current       maximum      dequeued
                     0             0             0
Paths            total    successful      no route     sys error         CSPFs
                     0             0             0             0             0
Time (secs)      total         CSPFs  avg per CSPF      % of rpd
               0.000000      0.000000      0.000000        0.0000
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` *`instance-name`* option added in Junos OS Release 15.1.

show mpls diffserv-te

IN THIS SECTION

- Syntax | 22157
- Syntax (EX Series Switches) | 22157
- Description | 22157
- Options | 22157
- Required Privilege Level | 22157
- Output Fields | 22157
- Sample Output | 22158
- Release Information | 22158

Syntax

```
show mpls diffserve-te
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show mpls diffserve-te
```

Description

Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information.

Options

none	Display DiffServ classes and priorities used by MPLS LSPs.
instance *instance-name*	(Optional) Display DiffServ classes and priorities used by MPLS LSPs for the specified instance. If *instance-name* is omitted, DiffServ information for the master instance is displayed.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 813 on page 22158 describes the output fields for the `show mpls diffserv-te` command. Output fields are listed in the approximate order in which they appear.

Table 813: show mpls diffserv-te Output Fields

Field Name	Field Description
Bandwidth model	Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.
TE class	DiffServ traffic engineering class.
Traffic class	MPLS class type that corresponds to the DiffServ traffic engineering class: • **ct0**—Best effort • **ct1**—Assured forwarding • **ct2**—Expedited forwarding • **ct3**—Network control
Priority	MPLS preemption priority for this class type, a value from **0** through **7**. Interior gateway protocols (IGPs) distribute information about the available bandwidth for each traffic engineering class.

Sample Output

show mpls diffserv-te

```
user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class     Traffic class       Priority
te0          ct0                 3
te1          ct1                 2
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` *`instance-name`* option added in Junos OS Release 15.1.

show mpls egress-protection

IN THIS SECTION

- Syntax | 22159
- Description | 22159
- Options | 22159
- Required Privilege Level | 22160
- Output Fields | 22160
- Sample Output | 22160
- Sample Output | 22161
- Release Information | 22161

Syntax

```
show mpls egress-protection
<brief | detail>
<logical-system (all | logical-system-name)>
```

Description

Display information about egress protection.

NOTE: Use this command on the device configured as the protector PE router to display information about egress protection. If you use this command on the device configured as the primary PE router, no output is displayed.

Options

none	Display standard information about egress protection.
brief | detail	(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 814 on page 22160 describes the output fields for the `show mpls egress-protection detail` command. Output fields are listed in the approximate order in which they appear.

Table 814: show mpls lsp Output Fields

Field Name	Field Description
Instance	Indicates egress instance name
Type	Indicates type of the VRF. It can be either local-vrf or remote-vrf
RIB	Indicates the edge-protection created routing table
Context-Id	Indicates the context-ID associated with the RIB.
Interface/Enhanced-lookup	Show VT interfaces associated with the backup RIB. Shows **Enhanced-lookup** for MX Series 5G Universal Routing Platforms with the Enhanced IP Network Services mode configured using the `network-services enhanced-ip` statement at the `[edit chassis]` hierarchy level.

Sample Output

show mpls egress-protection detail (Centralized Protector)

```
user@host> show mpls egress-protection detail

Instance                Type      Protection-Type
```

```
rsite1                  remote-vrf  Protector
  RIB __99.99.1.4-rsite1__.inet.0, Context-Id 99.99.1.4, Enhanced-lookup
  Route Target 1:1
rsite24                 remote-vrf  Protector
  RIB __99.99.1.4-rsite24__.inet.0, Context-Id 99.99.1.4, Enhanced-lookup
  Route Target 100:1023
```

Sample Output

show mpls egress-protection detail (Collocated Protector)

```
user@host> show mpls egress-protection detail

Instance                 Type      Protection-Type
site2                   local-vrf   Protector
  RIB __66.6.6.6-site2__.inet.0, Context-Id 66.6.6.6, Interface vt-1/3/0.87031809
  Route Target 100:251
site12                  local-vrf   Protector
  RIB __66.6.6.6-site12__.inet.0, Context-Id 66.6.6.6, Interface vt-1/3/0.87031808
  Route Target 100:250
  Route Target 100:251

site2                   local-vrf   Protector
  RIB __66.6.6.6-site2__.inet.0, Context-Id 66.6.6.6, Interface vt-1/3/0.87031809
  Route Target 100:251
```

Release Information

Command introduced in Junos OS Release 11.4R3.

RELATED DOCUMENTATION

Example: Configuring MPLS Egress Protection for Layer 3 VPN Services

Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP

show mpls interface

IN THIS SECTION

- Syntax | 22162
- Syntax (EX Series Switches) | 22162
- Description | 22162
- Options | 22162
- Additional Information | 22163
- Required Privilege Level | 22163
- Output Fields | 22163
- Sample Output | 22164
- Release Information | 22165

Syntax

```
show mpls interface
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show mpls interface
```

Description

Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.

Options

none Display information about MPLS-enabled interfaces.

instance *instance-name*	(Optional) Display information about MPLS-enabled interfaces for the specified routing instance. If *instance-name* is omitted, information about MPLS-enabled interfaces is displayed for the master instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

MPLS is enabled on an interface when the interface is configured with both the `set protocol mpls interface` *interface-name* and `set interface` *interface-name* `unit 0 family mpls` statements.

Required Privilege Level

view

Output Fields

Table 815 on page 22163 describes the output fields for the `show mpls interface` command. Output fields are listed in the approximate order in which they appear.

Table 815: show mpls interface Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: **Up** or **Dn** (down).
Administrative groups	Administratively assigned colors of the link.
Maximum labels	Maximum number of MPLS labels upon which MPLS can operate on a logical interface. This is configured using the `maximum-labels` statement at the `[edit logical-systems` *logical-system-name* `interfaces` *interface-name* `unit` *logical-unit-number* `family mpls]` or the `[edit interfaces` *interface-name* `unit` *logical-unit-number* `family mpls]` hierarchy levels.

Table 815: show mpls interface Output Fields *(Continued)*

Field Name	Field Description
Static protection revert time	Time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path. This is configured using the `protection-revert-time` statement at the `[edit logical-systems` *`logical-system-name`* `protocols mpls interface` *`interface-name`* `static]` or the `[edit protocols mpls interface` *`interface-name`* `static]` hierarchy levels.
Always mark connection protection tlv	**Enabled** or **Disabled**: **Enabled** indicates that the `always-mark-connection-protection-tlv` statement is configured at the `[edit logical-systems` *`logical-system-name`* `protocols mpls interface` *`interface-name`* `static]` or the `[edit protocols mpls interface` *`interface-name`* `static]` hierarchy levels. When this statement is configured, it marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, the `switch-away-lsps` statement must be configured.
Switch away lsps	**Enabled** or **Disabled**: **Enabled** indicates that the `switch-away-lsps` statement is configured at the `[edit logical-systems` *`logical-system-name`* `protocols mpls interface` *`interface-name`* `static]` or the `[edit protocols mpls interface` *`interface-name`* `static]` hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.

Sample Output

show mpls interface

```
user@host> show mpls interface

Interface: ge-0/2/1.57
  State: Up
  Administrative group: <none>
  Maximum labels: 5
  Static protection revert time: 5 seconds
```

```
Always mark connection protection tlv: Disabled
Switch away lsps : Disabled
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` *`instance-name`* option added in Junos OS Release 15.1.

show mpls label usage

IN THIS SECTION

- Syntax | 22165
- Description | 22165
- Options | 22166
- Additional Information | 22166
- Required Privilege Level | 22166
- Output Fields | 22167
- Sample Output | 22167
- Release Information | 22168

Syntax

```
show mpls label usage
<label label value>
<label-range range-start range-end>
<logical-system  (all | logical-sytem-name)>
```

Description

Show the available label space resource in RPD and also the applications that use the label space in RPD. There are four different label spaces currently used in MPLS—namely LSI, dynamic, block, and static.

Each label space has a fixed number and cannot grow beyond the fixed value. Using this command, the administrator can monitor the available labels in each label space and the applications that are using the labels. Based on the availability of labels, the administrator can decide to stop any service and free some labels or use other service where the labels are available.

Starting in Junos OS Release 17.2, you can configure the `enhanced-ip` command, which is supported on platforms using Modular Port Concentrators (MPCs) equipped with Junos Trio chipsets. You can also separate the MPLS labels used for different label spaces which provides more flexibility and scalability.

When you set each member router's network services to `enhanced-ip`, only MPC or Modular Interface Cards (MICs) modules and Multiservices Dense Port Concentrator (MS-DPC) modules are powered on in the chassis. Non-service DPCs do not work with enhanced IP network services.

Options

none	Display the available labels in each label space and the applications using the labels.
label ***label value***	(Optional) Display the information about which *label value* is used by which protocol, if any.
label-range ***range-start range-end***	(Optional) Display the complete information about the `label-range` specified. WIth the `enhanced-ip` command enabled on the supported device, effective ranges and configured ranges along with details of different label spaces such as LSI, dynamic, block, and static types are displayed.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

Once the label space crosses the threshold, a new syslog message is added.

<label-space-name> label space usage crossed threshold limit of 90%.

For instance, LSI label space usage crossed threshold limit of 90%.

Required Privilege Level

view

Output Fields

Table 816 on page 22167 describes the output fields for the `show mpls label usage` command. Output fields are listed in the order in which they appear.

Table 816: show mpls label usage Fields

Field Name	Field Description
Label Space	Indicates the different types of labels currently used in MPLS.
Total	Indicates the total label space available.
Available	Indicates the number of freely available labels and also the percentage of the label space available.
Applications	Indicates the applications that use the MPLS label spaces.
Effective Ranges	Indicates actual ranges in use, which can be different from configured ranges, if conflicting with label already allocated.
Configured Ranges	Indicates the currently configured range assigned to different label spaces on the device.

Sample Output

show mpls label usage

```
user@host> show mpls label usage
Label space Total   Available        Applications
LSI         999984  999971 (100.00%) BGP/LDP VPLS with no-tunnel-services, BGP L3VPN with vrf-
table-label
Block       999984  999971 (100.00%) BGP/LDP VPLS with tunnel-services, BGP L2VPN
Dynamic     999984  999971 (100.00%) RSVP, LDP, PW, L3VPN, RSVP-P2MP, LDP-P2MP, MVPN, EVPN, BGP
Static      48576   48576  (100.00%) Static LSP, Static PW
```

With `enhanced-ip` enabled on the supported device, you get the following additional output.

```
user@host> show mpls label usage
Label space Total   Available        Applications
LSI         996983  996983 (100.00%) BGP/LDP VPLS with no-tunnel-services, BGP L3VPN with vrf-
table-label
Block       996983  996983 (100.00%) BGP/LDP VPLS with tunnel-services, BGP L2VPN
Dynamic     996983  996983 (100.00%) RSVP, LDP, PW, L3VPN, RSVP-P2MP, LDP-P2MP, MVPN, EVPN, BGP
Static      48576   48576  (100.00%) Static LSP, Static PW
Effective Ranges
Range name  Shared with Start   End
Dynamic     16      999
Dynamic     4001    999999
Static      1000000 1048575
SRGB        1000    2999     OSPF
SRGB        3000    4000     GLOBAL
Configured Ranges
Range name  Shared with Start   End
Dynamic     16      999
Dynamic     4001    999999
Static      1000000 1048575
SRGB        1000    2999     OSPF
SRGB        3000    4000     GLOBAL

user@host> show mpls label usage label 101
Label 101 is used by protocol BGP
user@host> show mpls label usage label 102
Label 102 is used by protocol LDP
user@host> show mpls label usage label 103
Label 103 is not allocated to any protocol
```

Release Information

Command introduced in Junos OS Release 15.1.

Support for the `label` statement added in Junos OS Release 17.2.

Support for the `label-range` statement added in Junos OS Release 17.2.

show mpls label usage label-range

IN THIS SECTION

- Syntax | 22169
- Description | 22169
- Options | 22170
- Required Privilege Level | 22170
- Output Fields | 22170
- Sample Output | 22171
- Release Information | 22172

Syntax

```
show mpls label usage label-range
<block-label-range range-start range-end>
<dynamic-label-range range-start range-end>
<label-limit label-limit value>
<lsi-label-range range-start range-end>
<static-label-range range-start range-end>
```

Description

There are four different label spaces currently used in MPLS—namely LSI, dynamic, block, and static. Each label space has a fixed number and cannot grow beyond the fixed value. Using `show mpls label usage label-range` command, the administrator can monitor the available labels in each label space and the applications that are using the labels. Based on the availability of labels, the administrator can decide to stop any service and free some labels or use other service where the labels are available.

Starting in Junos OS Release 17.2, you can configure the `enhanced-ip` command, which is supported on platforms using Modular Port Concentrators (MPCs) equipped with Junos Trio chipsets. You can also separate the MPLS labels used for different label spaces which provides more flexibility and scalability.

When you set each member router's network services to `enhanced-ip`, only MPC or Modular Interface Cards (MICs) modules and Multiservices Dense Port Concentrator (MS-DPC) modules are powered on in the chassis. Non-service DPCs do not work with enhanced IP network services.

Options

block-label-range *range-start range-end*	Display the details of block label type.
dynamic-label-range *range-start range-end*	Display the details of dynamic label type.
label-limit *label-limit value*	Limit for the number of concurrent active labels.
lsi-label-range *range-start range-end*	Display the details of LSI label type.
static-label-range *range-start range-end*	Display the details of static label type.

Required Privilege Level

view

Output Fields

Table 817 on page 22170 describes the output fields for the `show mpls label usage label-range` command. Output fields are listed in the order in which they appear.

Table 817: show mpls label usage label-range Fields

Field Name	Field Description
Label Space	Indicates the different types of labels currently used in MPLS.
Total	Indicates the total label space available.
Available	Indicates the number of freely available labels and also the percentage of the label space available.
Applications	Indicates the applications that use the MPLS label spaces.
Effective Ranges	Indicates actual ranges in use, which can be different from configured ranges, if conflicting with label already allocated.
Configured Ranges	Indicates the currently configured range assigned to different label spaces on the device.

Table 817: show mpls label usage label-range Fields *(Continued)*

Field Name	Field Description
Total	Indicates the currently used labels with label type and application type details.

Sample Output

show mpls label usage label-range

With the `enhanced-ip` command enabled on the supported device, you get the following output.

```
user@host> show mpls label usage label-range 16 600
Label space Total   Available        Applications
LSI         101     99     (98.02% ) BGP/LDP VPLS with no-tunnel-services, BGP L3VPN with vrf-
table-label
Block       101     101    (100.00%) BGP/LDP VPLS with tunnel-services, BGP L2VPN
Dynamic     101     98     (97.03% ) RSVP, LDP, PW, L3VPN, RSVP-P2MP, LDP-P2MP, MVPN, EVPN, BGP
Static      48576   48576  (100.00%) Static LSP, Static PW
Effective Ranges
Range name  Shared with Start   End
LSI         300     400
Block       500     600
Dynamic     100     200
Static      1000000 1048575
Configured Ranges
Range name  Shared with Start   End
LSI         300     400
Block       500     600
Dynamic     100     200
Static      1000000 1048575
Total(16 to 600) 5
Label type  Alloc count
LSI         2
Dynamic     3
App type    Alloc count
LDP         1
BGP         2
RT_INSTANCE 2
```

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

show mpls label usage | 22165

show mpls lsp

IN THIS SECTION

- Syntax | 22172
- Syntax (EX Series Switches Only) | 22173
- Description | 22173
- Options | 22173
- Required Privilege Level | 22177
- Output Fields | 22177
- Sample Output | 22189
- Release Information | 22202

Syntax

```
show mpls lsp
<brief | detail | extensive | terse>
<abstract-computation>
<autobandwidth>
<bidirectional | unidirectional>
<bypass>
<count-active-routes>
<defaults>
<descriptions>
<down | up>
```

```
<externally-controlled>
<externally-provisioned>
<instance routing-instance-name>
<locally-provisioned>
<logical-system (all | logical-system-name)>
<lsp-type>
<name name>
<p2mp>
<reverse-statistics>
<segment>
<statistics>
<transit>
```

Syntax (EX Series Switches Only)

```
show mpls lsp
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<descriptions>
<down | up>
<externally-controlled>
<externally-provisioned>
<lsp-type>
<name name>
<p2mp>
<statistics>
<transit>
```

Description

Display information about configured and active dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options

none Display standard information about all configured and active dynamic MPLS LSPs.

brief | detail | extensive | terse

(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

In the extensive command output, the duplicate back-to-back messages are recorded as aggregated messages. An additional timestamp is included for these aggregated messages, where if the aggregated messages are five or less, timestamp deltas are recorded for each message, and if the aggregated messages are greater than five, the first and last timestamp is recorded.

For example:

- All timestamps

```
9204 Jun 29 13:23:45.405 10.54.43.110: Explicit Route: bad strict
route [3 times - 13:21:00, 13:22:01, 13:23:10]
```

- Timestamp deltas

```
9204 Jun 29 13:23:45.405 510.4.43.110: Explicit Route: bad strict
route [3 times - 13:21:00, +1:01, +2:10]
```

- First and last timestamp

```
9204 Jun 29 13:23:45.405 10.4.43.110: Explicit Route: bad strict route
[6 times - 13:21:00, 13:23:10]
```

abstract-computation

(Optional) Display abstract computation preprocessing for LSPs.

See show mpls lsp abstract-computation for more details.

autobandwidth

(Optional) Display automatic bandwidth information. This option is explained separately (see `show mpls lsp autobandwidth`).

bidirectional | unidirectional

(Optional) Display bidirectional or unidirectional LSP information, respectively.

bypass

(Optional) Display LSPs used for protecting other LSPs.

count-active-routes

(Optional) Display active routes for LSPs.

defaults

(Optional) Display the MPLS LSP default settings.

descriptions (Optional) Display the MPLS label-switched path (LSP) descriptions. To view this information, you must configure the description statement at the `[edit protocol mpls lsp]` hierarchy level. Only LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

down | up (Optional) Display only LSPs that are inactive or active, respectively.

externally-controlled (Optional) Display the LSPs that are under the control of an external Path Computation Element (PCE).

externally-provisioned (Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).

instance *instance-name* (Optional) Display MPLS LSP information for the specified instance. If *instance-name* is omitted, MPLS LSP information is displayed for the master instance.

locally-provisioned (Optional) Display LSPs that have been provisioned locally by the Path Computation Client (PCC).

logical-system (all | *logical-system-name*) (Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type (Optional) Display information about a particular LSP type:

- `bypass`—Sessions for bypass LSPs.
- `egress`—Sessions that terminate on this routing device.
- `ingress`—Sessions that originate from this routing device.
- `pop-and-forward`—Sessions that originate from RSVP-TE pop-and-forward LSP tunnels.
- `transit`—Sessions that pass through this routing device.

name *name* (Optional) Display information about the specified LSP or group of LSPs.

p2mp (Optional) Display information about point-to-multipoint LSPs.

reverse-statistics (Optional) Display packet statistics for reverse direction of LSPs.

segment (Optional) Display segment identifier (SID) labels.

statistics (Optional) (Ingress and transit routers only) Display accounting information about LSPs. Statistics are generally not available for LSPs on the egress routing device,

because in most cases either penultimate hop popping or explicit null labels are used. These labels, 3 and 0, respectively, fall into the reserved range of 0-15, for which egress statistics are not available. For example, with explicit null operation (label 0), as the packet arrives at the egress routing device, the hardware removes the MPLS label and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.

NOTE: When a label from the 0-15 reserved range is used at the egress LSR, you expect to see egress LSP statistics as not applicable (NA):

```
user@ptx> show mpls lsp statistics
Ingress LSP: 1 sessions
To              From            State     Packets
Bytes LSPname
10.0.255.42     10.0.255.47     Up             10
920    PTX-to-MX_UHP
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To              From            State     Packets
Bytes LSPname
10.0.255.47     10.0.255.42     Up          NA       NA  MX-to-
PTX_PHP
Total 1 displayed, Up 1, Down 0
```

NOTE: If a bypass LSP is configured for the primary static LSP, display cumulative statistics of packets traversing through the protected LSP and bypass LSP when traffic is re-optimized when the protected LSP link is restored. (Bypass LSPs are not supported on QFX Series switches.)

When used with the `bypass` option (`show mpls lsp bypass statistics`), display statistics for the traffic that flows only through the bypass LSP.

transit (Optional) Display LSPs transiting this routing device.

Required Privilege Level

view

Output Fields

Table 818 on page 22177 describes the output fields for the `show mpls lsp` command. Output fields are listed in the approximate order in which they appear.

Table 818: show mpls lsp Output Fields

Field Name	Field Description	Level of Output
`Ingress LSP`	Information about LSPs on the ingress routing device. Each session has one line of output.	All levels
`Egress LSP`	Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output.	All levels
`Transit LSP`	Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
`P2MP name`	Name of the point-to-multipoint LSP. Dynamically generated P2MP LSPs used for VPLS flooding use dynamically generated P2MP LSP names. The name uses the format *identifier*:`vpls`:*router-id*:*routing-instance-name*. The *identifier* is automatically generated by Junos OS.	All levels
`P2MP branch count`	Number of destination LSPs the point-to-multipoint LSP is transmitting to.	All levels
`P`	An asterisk (`*`) under this heading indicates that the LSP is a primary path.	All levels
address	(`detail` and `extensive`) Destination (egress routing device) of the LSP.	`detail extensive`
`To`	Destination (egress routing device) of the session.	`brief`

Table 818: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`From`	Source (ingress routing device) of the session.	`brief detail`
`State`	State of the LSP handled by this RSVP session: `Up`, `Dn` (down), or `Restart`.	`brief detail`
`Active Route`	Number of active routes (prefixes) installed in the forwarding table. For ingress LSPs, the forwarding table is the primary IPv4 table (`inet.0`). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (`mpls.0`).	`detail extensive`
`Rt`	Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (`inet.0`). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	`brief`
`P`	Path. An asterisk (`*`) underneath this column indicates that the LSP is a primary path.	`brief`
`ActivePath`	(Ingress LSP) Name of the active path: `Primary` or `Secondary`.	`detail extensive`
`LSPname`	Name of the LSP.	`brief detail`
`Statistics`	Displays the number of packets and the number of bytes transmitted over the LSP. These counters are reset to zero whenever the LSP path is optimized (for example, during an automatic bandwidth allocation).	`extensive`
`Aggregate statistics`	Displays the number of packets and the number of bytes transmitted over the LSP. These counters continue to iterate even if the LSP path is optimized. You can reset these counters to zero using the `clear mpls lsp statistics` command.	`extensive`
`Packets`	Displays the number of packets transmitted over the LSP. `NA` is shown on egress LSP if incoming label is in the reserved range of 0 -15	`brief extensive`

Table 818: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bytes	Displays the number of bytes transmitted over the LSP. NA is shown on egress LSP if incoming label is in the reserved range of 0 -15	brief extensive
DiffServeInfo	Type of LSP: multiclass LSP (multiclass diffServ-TE LSP) or Differentiated-Services-aware traffic engineering LSP (diffServ-TE LSP).	detail
LSPtype	Type of LSP: • Static configured—Static • Dynamic configured—Dynamic • Externally controlled—External path computing entity Also indicates if the LSP is a Penultimate hop popping LSP or an Ultimate hop popping LSP.	detail extensive
Bypass	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. LSPpath can be either primary or secondary and can be displayed on the ingress, egress, and transit routing devices.	detail
Bidir	(GMPLS) The LSP allows data to travel in both directions between GMPLS devices.	All levels
Bidirectional	(GMPLS) The LSP allows data to travel both ways between GMPLS devices.	All levels
FastReroute desired	Fast reroute has been requested by the ingress routing device.	detail

Table 818: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Link protection desired`	Link protection has been requested by the ingress routing device.	`detail`
Node/Link protection desired	Link protection has been requested by the ingress routing device.	`detail`
`LSP Control Status`	(Ingress LSP) LSP control mode: • External—By default, all PCE-controlled LSPs are under external control. When an LSP is under external control, the PCC uses the PCE-provided parameters to set up the LSP. • Local—A PCE-controlled LSP can come under local control. When the LSP switches from external control to local control, path computation is done using the CLI-configured parameters and constraint-based routing. Such a switchover happens only when there is a trigger to re-signal the LSP. Until then, the PCC uses the PCE-provided parameters to signal the PCE-controlled LSP, although the LSP remains under local control. A PCE-controlled LSP switches to local control from its default external control mode in cases such as no connectivity to a PCE or when a PCE returns delegation of LSPs back to the PCC.	`extensive`
`External Path CSPF status`	(PCE-controlled LSPs) Status of the PCE-controlled LSP with per path attributes: • Local • External	`extensive`
`Externally Computed ERO`	(PCE-controlled LSPs) Externally computed explicit route when the route object is not null or empty. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (`S`) or loose (`L`).	`extensive`

Table 818: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
EXTCTRL_LSP	(PCE-controlled LSPs) Display path history including the bandwidth, priority, and metric values received from the external controller.	extensive
flap counter	Counts the number of times a LSP flaps down or up.	extensive
LoadBalance	(Ingress LSP) CSPF load-balancing rule that was configured to select the LSP's path among equal-cost paths: Most-fill, Least-fill, or Random.	detail extensive
Signal type	Signal type for GMPLS LSPs. The signal type determines the peak data rate for the LSP: DS0, DS3, STS-1, STM-1, or STM-4.	All levels
Encoding type	LSP encoding type: Packet, Ethernet, PDH, SDH/SONET, Lambda, or Fiber.	All levels
Switching type	Type of switching on the links needed for the LSP: Fiber, Lamda, Packet, TDM, or PSC-1.	All levels
GPID	Generalized Payload Identifier (identifier of the payload carried by an LSP): HDLC, Ethernet, IPv4, PPP, or Unknown.	All levels
Protection	Configured protection capability desired for the LSP: Extra, Enhanced, none, One plus one, One to one, or Shared.	All levels
Upstream label in	(Bidirectional LSPs) Incoming label for reverse direction traffic for this LSP.	All levels
Upstream label out	(Bidirectional LSPs) Outgoing label for reverse direction traffic for this LSP.	All levels
Suggested label received	(Bidirectional LSPs) Label the upstream interface suggests to use in the Resv message that is sent.	All levels

Table 818: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Suggested label sent`	(Bidirectional LSPs) Label the downstream node suggests to use in the Resv message that is returned.	All levels
`Autobandwidth`	(Ingress LSP) The LSP is performing autobandwidth allocation.	`detail extensive`
Mbb counter	Counts the number of times a LSP incurs MBB.	extensive
`MinBW`	(Ingress LSP) Configured minimum value of the LSP, in bps.	`detail extensive`
`MaxBW`	(Ingress LSP) Configured maximum value of the LSP, in bps.	`detail extensive`
`Dynamic MinBW`	(Ingress LSP) Displays the current dynamically specified minimum bandwidth allocation for the LSP, in bps.	`detail extensive`
`Dynamic MinBW`	(Ingress LSP) Displays the current dynamically specified minimum bandwidth allocation for the LSP, in bps.	`detail extensive`
`AdjustTimer`	(Ingress LSP) Configured value for the `adjust-timer` statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	`detail extensive`
`Adjustment Threshold`	(Ingress LSP) Configured value for the `adjust-threshold` statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	`detail extensive`
`Time for Next Adjustment`	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	`detail extensive`
`Time of Last Adjustment`	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	`detail extensive`

Table 818: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`MaxAvgBW util`	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.	`detail extensive`
`Overflow limit`	(Ingress LSP) Configured value of the threshold overflow limit.	`detail extensive`
`Overflow sample count`	(Ingress LSP) Current value for the overflow sample count.	`detail extensive`
`Bandwidth Adjustment in` *`nnn`* `second(s)`	(Ingress LSP) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds.	`detail extensive`
`In-place Update Count`	Current value of the in-place LSP bandwidth update counter indicating the number of times an LSP-ID is reused when LSP-ID re-use is enabled for an LSP.	`detail extensive`
`Underflow limit`	(Ingress LSP) Configured value of the threshold underflow limit.	`detail extensive`
`Underflow sample count`	(Ingress LSP) Current value for the underflow sample count.	`detail extensive`
`Underflow Max AvgBW`	(Ingress LSP) The highest sample bandwidth among the underflow samples recorded currently. This is the signaling bandwidth if an adjustment occurs because of an underflow.	`detail extensive`

`Active path indicator`	(Ingress LSP) A value of `*` indicates that the path is active. The absence of `*` indicates that the path is not active. In the following example, “long” is the active path.
	
	`*Primary long`
	`Standby short`
		`detail extensive`

`Primary`	(Ingress LSP) Name of the primary path.	`detail extensive`

Table 818: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Secondary	(Ingress LSP) Name of the secondary path.	detail extensive
Standby	(Ingress LSP) Name of the path in standby mode.	detail extensive
State	(Ingress LSP) State of the path: Up or Dn (down).	detail extensive
COS	(Ingress LSP) Class-of-service value.	detail extensive
Bandwidth per class	(Ingress LSP) Active bandwidth for the LSP path for each MPLS class type, in bps.	detail extensive
Priorities	(Ingress LSP) Configured value of the setup priority and the hold priority respecitively (the setup priority is displayed first), where 0 is the highest priority and 7 is the lowest priority. If you have not explicitly configured these values, the default values are displayed (7 for the setup priority and 0 for the hold priority).	detail extensive
OptimizeTimer	(Ingress LSP) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	detail extensive
SmartOptimizeTimer	(Ingress LSP) Configured value of the smart optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	detail extensive
Reoptimization in xxx seconds	(Ingress LSP) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds.	detail extensive
Computed ERO (S [L] denotes strict [loose] hops)	(Ingress LSP) Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (S) or loose (L).	detail extensive
CSPF metric	(Ingress LSP) Constrained Shortest Path First metric for this path.	detail extensive

Table 818: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Received RRO`	(Ingress LSP) Received record route. A series of hops, each with an address followed by a flag. (In most cases, the received record route is the same as the computed explicit route. If `Received RRO` is different from `Computed ERO`, there is a topology change in the network, and the route is taking a detour.) The following flags identify the protection capability and status of the downstream node: • `0x01`—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the Local protection flag was set in the `SESSION_ATTRIBUTE` object of the corresponding Path message. • `0x02`—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously). • `0x03`—Combination of `0x01` and `0x02`. • `0x04`—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section. • `0x08`—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the `Local protection available` bit is set but the `Node protection` bit is cleared. • `0x09`—Detour is established. Combination of `0x01` and `0x08`. • `0x10`—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted. • `0x20`—Node ID. Indicates that the address specified in the RRO's IPv4 or IPv6 sub-object is a node ID address, which refers to the router address or router ID. Nodes must use the same address consistently. • `0xb`—Detour is in use. Combination of `0x01`, `0x02`, and `0x08`.	`detail extensive`

Table 818: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Labels	Labels of pop-and-forward LSP tunnel: • P—Pop labels. • D—Delegation labels.	extensive
Index number	(Ingress LSP) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed.	extensive
Date	(Ingress LSP) Date of the LSP event.	extensive
Time	(Ingress LSP) Time of the LSP event.	extensive
Event	(Ingress LSP) Description of the LSP event.	extensive
Created	(Ingress LSP) Date and time the LSP was created.	extensive
Resv style	(Bypass) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	brief detail extensive
Labelin	Incoming label for this LSP.	brief detail
Labelout	Outgoing label for this LSP.	brief detail
LSPname	Name of the LSP.	brief detail
Time left	Number of seconds remaining in the lifetime of the reservation.	detail
Since	Date and time when the RSVP session was initiated.	detail

Table 818: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Tspec`	Sender's traffic specification, which describes the sender's traffic parameters.	`detail`
`Port number`	Protocol ID and sender or receiver port used in this RSVP session.	`detail`
`PATH rcvfrom`	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this router, and number of packets received from the upstream neighbor.	`detail`
`PATH sentto`	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor, and number of packets sent to the downstream routing device.	`detail`
`RESV rcvfrom`	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor. The output in this field, which is consistent with that in the `PATH rcvfrom` field, indicates that the RSVP negotiation is complete.	`detail`
`Record route`	Recorded route for the session, taken from the record route object.	`detail`
`Pop-and-forward`	Attributes of the pop-and-forward LSP tunnel.	`extensive`
`ETLD In`	Number of transport labels that the LSP-Hop can potentially receive from its upstream hop. It is recorded as Effective Transport Label Depth (ETLD) at the transit and egress devices.	`extensive`
`ETLD Out`	Number of transport labels the LSP-Hop can potentially send to its downstream hop. It is recorded as ETLD at the transit and ingress devices.	`extensive`

Table 818: show mpls lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Delegation hop`	Specifies if the transit hop is selected as a delegation label: • Yes • No	`extensive`
`Soft preempt`	Number of soft preemptions that occurred on a path and when the last soft preemption occurred. Only successful soft preemptions are counted (those that actually resulted in a new path being used).	`detail`
`Soft preemption pending`	Path is in the process of being soft preempted. This display is removed once the ingress router has calculated a new path.	`detail`
`MPLS-TE LSP Defaults`	Default settings for MPLS traffic engineered LSPs: • `LSP Holding Priority`—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully. • `LSP Setup Priority`—Determines whether a new LSP that preempts an existing LSP can be established. • `Hop Limit`—Specifies the maximum number of routers the LSP can traverse (including the ingress and egress). • `Bandwidth`—Specifies the bandwidth in bits per second for the LSP. • `LSP Retry Timer`—Length of time in seconds that the ingress router waits between attempts to establish the primary path.	`defaults`

The XML tag name of the `bandwidth` tag under the `auto-bandwidth` tag has been updated to `maximum-average-bandwidth` . You can see the new tag when you issue the `show mpls lsp extensive` command with the `| display xml` pipe option. If you have any scripts that use the `bandwidth` tag, ensure that they are updated to `maximum-average-bandwidth`.

Sample Output

show mpls lsp defaults

```
user@host> show mpls lsp defaults
   MPLS-TE LSP Defaults
      LSP Holding Priority      0
      LSP Setup Priority        7
      Hop Limit                 255
      Bandwidth                 0
      LSP Retry Timer           30 seconds
```

show mpls lsp descriptions

```
user@host> show mpls lsp descriptions
Ingress LSP: 3 sessions
To              LSP name                              Description
10.0.0.195      to-sanjose                            to-sanjose-desc
10.0.0.195      to-sanjose-other-desc                 other-desc
Total 2 displayed, Up 2, Down 0
```

show mpls lsp detail

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath:  (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
```

```
          10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  157, Since: Wed Jul 18 17:55:12 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp detail (When Egress Protection Is in Standby Mode)

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath:  (primary)
  LSPtype: Static Configured, Ultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
```

```
          10.0.0.18 10.0.0.22
   11 Sep 20 15:54:35.032 Make-before-break: Switched to new instance
   10 Sep 20 15:54:34.029 Record Route:  10.0.0.18 10.0.0.22
    9 Sep 20 15:54:34.029 Up
    8 Sep 20 15:54:20.271 Originate make-before-break call
    7 Sep 20 15:54:20.271 CSPF: computation result accepted  10.0.0.18 10.0.0.22
    6 Sep 20 15:52:10.247 Selected as active path
    5 Sep 20 15:52:10.246 Record Route:  10.0.0.18 10.0.0.22
    4 Sep 20 15:52:10.243 Up
    3 Sep 20 15:52:09.745 Originate Call
    2 Sep 20 15:52:09.745 CSPF: computation result accepted  10.0.0.18 10.0.0.22
    1 Sep 20 15:51:39.903 CSPF failed: no route toward 192.168.0.4
  Created: Thu Sep 20 15:51:08 2012
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  148, Since: Thu Sep 20 15:52:10 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 49601 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 27 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp detail (When Egress Protection Is in Effect During a Local Repair)

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions
```

```
192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath:  (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Down, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  157, Since: Wed Jul 18 17:55:12 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
Egress protection PLR as protector: In Use
PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp extensive

```
user@host> show mpls lsp extensive
Ingress LSP: 1 sessions
```

```
192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath:  (primary)
  LSPtype: Static Configured, Ultimate hop popping
  LSP Control Status: Externally controlled
  LoadBalance: Random
  Metric: 10
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    External Path CSPF status: local
    Bandwidth: 98.76kbps
    SmartOptimizeTimer: 180
    Include All: green
    Externally Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 0) 10.1.3.2 S
10.2.3.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          10.0.0.18 10.0.0.22
    9 May 17 16:55:06.574 EXTCTRL LSP: Sent Path computation request and LSP status
    8 May 17 16:55:06.574 EXTCTRL_LSP: Computation request/lsp status contains: signalled bw
98760 req BW 0 admin group(exclude 0 include any 0 include all 16) priority setup 5 hold 4 hops:
10.1..3.2 10.2..3.2
    7 May 17 16:55:06.574 Selected as active path
    6 May 17 16:55:06.558 EXTCTRL LSP: Sent Path computation request and LSP status
    8 May 17 16:55:06.574 EXTCTRL_LSP: Computation request/lsp status contains: signalled bw
98760 req BW 0 admin group(exclude 0 include any 0 include all 16) priority setup 5 hold 4 hops:
10.1.3.2 10.2..3.2
    7 May 17 16:55:06.574 Selected as active path
    6 May 17 16:55:06.558 EXTCTRL LSP: Sent Path computation request and LSP status
    5 May 17 16:55:06.558 EXTCTRL_LSP: Computation request/lsp status contains: signalled bw
98760 req BW 0 admin group(exclude 0 include any 0 include all 16) priority setup 5 hold 4 hops:
10.1..3.2 10.2..3.2
    4 May 17 16:55:06.557 Record Route: 10.1.3.2 10.2.3.2
    3 May 17 16:55:06.557 Up
    2 May 17 16:55:06.382 Originate Call
    1 May 17 16:55:06.382 EXTCTRL_LSP: Received setup parameters :: local_cspf, 10.1.3.2
10.2.3.2
  Created: Tue May 17 16:55:07 2016
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
```

```
192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  148, Since: Thu Sep 20 15:52:10 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 49601 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 27 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
```

show mpls lsp ingress extensive

```
user@host> show mpls lsp ingress extensive
Ingress LSP: 1 sessions

10.50.0.1
  From: 10.0.0.1, State: Up, ActiveRoute: 0, LSPname: test
  ActivePath:  (primary)
  LSPtype: Static Pop-and-forward Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    OptimizeTimer: 300
    SmartOptimizeTimer: 180
    Reoptimization in 240 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.1.1.2 S 10.4.4.1 S 10.5.5.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
                            (Labels: P=Pop D=Delegation)
          10.80.1.2(Label=18 P) 10.50.1.2(Label=17 P) 10.70.1.2(Label=16 P)
          10.92.1.1(Label=16 D) 10.93.1.2(Label=16 P) 10.99.1.1(Label=16 P)
          10.99.1.1(Label=16 P) 10.99.1.2(Label=3)
   17 Aug  3 13:17:33.601 CSPF: computation result ignored, new path less avail bw[3 times]
   16 Aug  3 13:02:51.283 CSPF: computation result ignored, new path no benefit[2 times]
```

```
   15 Aug  3 12:54:36.678 Selected as active path
   14 Aug  3 12:54:36.676 Record Route:  10.1.1.2 10.4.4.1 10.5.5.2
   13 Aug  3 12:54:36.676 Up
   12 Aug  3 12:54:33.924 Deselected as active
   11 Aug  3 12:54:33.924 Originate Call
   10 Aug  3 12:54:33.923 Clear Call
    9 Aug  3 12:54:33.923 CSPF: computation result accepted  10.1.1.2 10.4.4.1 10.5.5.2
    8 Aug  3 12:54:33.922 10.2.2.2: No Route toward dest
    7 Aug  3 12:54:28.177 CSPF: computation result ignored, new path no benefit[4 times]
    6 Aug  3 12:35:03.830 Selected as active path
    5 Aug  3 12:35:03.828 Record Route:  10.2.2.2 10.3.3.2
    4 Aug  3 12:35:03.827 Up
    3 Aug  3 12:35:03.814 Originate Call
    2 Aug  3 12:35:03.814 CSPF: computation result accepted  10.2.2.2 10.3.3.2
    1 Aug  3 12:34:34.921 CSPF failed: no route toward 10.50.0.1
  Created: Tue Aug  3 12:34:35 2010
Total 1 displayed, Up 1, Down 0
```

show mpls lsp extensive (automatic bandwidth adjustment enabled)

```
user@host> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath:  (primary)
  Node/Link protection desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Autobandwidth
  MinBW: 300bps, MaxBW: 1000bps, Dynamic MinBW: 1000bps
  Adjustment Timer: 300 secs AdjustThreshold: 25%
  Max AvgBW util: 963.739bps, Bandwidth Adjustment in 0 second(s).
  Min BW Adjust Interval: 1000, MinBW Adjust Threshold (in %): 50
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 9, Underflow Max AvgBW: 614.421bps
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    Bandwidth: 1000bps
    SmartOptimizeTimer: 180
```

```
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          192.168.0.6(flag=0x20) 10.0.0.18(Label=299792) 192.168.0.4(flag=0x20)
10.0.0.22(Label=3)
   12 Apr 30 10:25:17.024 Make-before-break: Switched to new instance
   11 Apr 30 10:25:16.023 Record Route:  192.168.0.6(flag=0x20) 10.0.0.18(Label=299792)
192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
   10 Apr 30 10:25:16.023 Up
    9 Apr 30 10:25:16.023 Automatic Autobw adjustment succeeded: BW changes from 300 bps to 1000
bps
    8 Apr 30 10:25:15.946 Originate make-before-break call
    7 Apr 30 10:25:15.946 CSPF: computation result accepted  10.0.0.18 10.0.0.22
    6 Apr 30 10:16:42.891 Selected as active path
    5 Apr 30 10:16:42.891 Record Route:  192.168.0.6(flag=0x20) 10.0.0.18(Label=299776)
192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
    4 Apr 30 10:16:42.890 Up
    3 Apr 30 10:16:42.828 Originate Call
    2 Apr 30 10:16:42.828 CSPF: computation result accepted  10.0.0.18 10.0.0.22
    1 Apr 30 10:16:14.064 CSPF: could not determine self[2 times]
  Created: Tue Apr 30 10:15:16 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp detail (in-place LSP bandwidth update enabled)

```
user@host >show mpls lsp detail
Ingress LSP: 1 sessions

10.2.5.2
  From: 192.168.255.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R4-1
  ActivePath: path-R2-R3 (primary)
  Link protection desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Follow destination IGP metric
```

```
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  LSP Self-ping Status : Enabled
*Primary   path-R2-R3       State: Up
    Priorities: 7 0
    Bandwidth: 100Mbps
    SmartOptimizeTimer: 180
    Flap Count: 1
    MBB Count: 4
    In-place Update Count: 2
```

show mpls lsp extensive (in-place LSP bandwidth update enabled)

```
user@host >show mpls lsp extensive
Ingress LSP: 1 sessions

10.2.5.2
  From: 192.168.255.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R4-1
  ActivePath: path-R2-R3 (primary)
  Link protection desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Follow destination IGP metric
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  LSP Self-ping Status : Enabled
*Primary   path-R2-R3       State: Up
    Priorities: 7 0
    Bandwidth: 100Mbps
    SmartOptimizeTimer: 180
    Flap Count: 1
    MBB Count: 4
    In-place Update Count: 2
  48 Mar  3 16:43:40.438 In-place LSP Update successful
   47 Mar  3 16:43:40.477 Record Route:  192.168.255.2(flag=0x21) 10.1.2.2(flag=1 Label=415072)
192.168.255.3(flag=0x21) 10.2.3.3(flag=1 Label=418192) 192.168.255.4(flag=0x20) 10.3.4.4(Label=
3)
   46 Mar  3 16:43:39.617 CSPF: ERO retrace was successful  10.1.2.2 10.2.3.3 10.3.4.4
  45 Mar  3 16:43:39.617 Originate In-place LSP Update call
  44 Mar  3 16:42:28.263 LSP-ID: 1 deleted
   43 Mar  3 16:42:28.263 Make-before-break: Cleaned up old instance: Hold dead expiry
   42 Mar  3 16:41:07.416 Link-protection Up
   41 Mar  3 16:41:06.169 Make-before-break: Switched to new instance
```

```
   49 Mar  3 16:41:06.167 Self-ping ended successfully
   39 Mar  3 16:41:05.839 Record Route:  192.168.255.2(flag=0x21) 10.1.2.2(flag=1 Label=415072)
192.168.255.3(flag=0x21) 10.2.3.3(flag=1 Label=418192) 192.168.255.4(flag=0x20) 10.3.4.4(Label=3)
   38 Mar  3 16:41:05.449 Record Route:  10.2.2.2(flag=0x20) 10.1.2.2(Label=415072)
192.168.255.3(flag=0x21) 10.2.3.3(flag=1 Label=418192) 192.168.255.4(flag=0x20) 10.3.4.4(Label=3)
   37 Mar  3 16:41:05.419 Up
   36 Mar  3 16:41:05.419 Self-ping started
   35 Mar  3 16:41:05.419 Self-ping enqueued
   34 Mar  3 16:41:05.419 Record Route:  192.168.255.2(flag=0x20) 10.1.2.2(Label=415072)
10.3.3.3(flag=0x20) 10.2.3.3(Label=418192) 192.168.255.4(flag=0x20) 10.3.4.4(Label=3)
   33 Mar  3 16:41:05.362 Originate make-before-break call
+ 32 Mar  3 16:41:05.362 LSP-ID: 2 created
   31 Mar  3 16:41:05.362 CSPF: computation result accepted  10.1.2.2 10.2.3.3 10.3.4.4
```

show mpls lsp bypass extensive

```
user@host # show mpls lsp bypass extensive

Ingress LSP: 1 sessions

10.2.2.2
  From: 10.1.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->10.1.2.2
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300032
  Resv style: 1 SE, Label in: -, Label out: 300032
  Time left:    -, Since: Tue Dec  3 15:19:49 2013
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 55750 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 1
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.1.5.2 (lt-1/2/0.15) 1221 pkts
  RESV rcvfrom: 10.1.5.2 (lt-1/2/0.15) 1221 pkts, Entropy label: No
  Explct route: 10.1.5.2 10.1.5.1
  Record route: <self>  10.1.5.2 10.1.5.1
+     4 Dec  3 15:19:49 Record Route:  10.1.5.2 10.1.5.1
```

```
+    3 Dec  3 15:19:49 Up
+     2 Dec  3 15:19:49 CSPF: computation result accepted
+    1 Dec  3 15:19:47 Originate Call
Total 1 displayed, Up 1, Down 0
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit LSP: 0 sessions
```

show mpls lsp p2mp

```
user@host> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To              From            State Rt P ActivePath       LSPname
10.255.245.51   10.255.245.50   Up     0 * path1            p2mp-branch-1
P2MP name: p2mp-lsp2, P2MP branch count: 1
To              From            State Rt P ActivePath       LSPname
10.255.245.51   10.255.245.50   Up     0 * path1            p2mp-st-br1
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp p2mp detail

```
user@host> show mpls lsp p2mp detail
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-branch-1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp1
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path1            State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
```

```
 192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
          192.168.208.17
P2MP name: p2mp-lsp2, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-st-br1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp2
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path1            State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
          192.168.208.17
Total 2 displayed, Up 2, Down 0
```

show mpls lsp detail count-active-routes

```
user@host> show mpls lsp detail count-active-routes
Ingress LSP: 1 sessions

10.213.192.2
  From: 10.156.162.128, State: Up, ActiveRoute: 1, LSPname: to-lahore
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Autobandwidth
  MinBW: 5Mbps MaxBW: 250Mbps
  AdjustTimer: 300 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 102 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    Bandwidth: 5Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 4)
 10.252.0.177 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
```

```
          10.252.0.177
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp statistics extensive

```
user@host> show mpls lsp statistics extensive
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  Statistics: Packets 302, Bytes 28992
  Aggregate statistics: Packets 302, Bytes 28992
  ActivePath:  (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          10.0.0.18 10.0.0.22
    6 Oct  3 11:18:28.281 Selected as active path
    5 Oct  3 11:18:28.281 Record Route:  10.0.0.18 10.0.0.22
    4 Oct  3 11:18:28.280 Up
    3 Oct  3 11:18:27.995 Originate Call
    2 Oct  3 11:18:27.995 CSPF: computation result accepted  10.0.0.18 10.0.0.22
    1 Oct  3 11:17:59.118 CSPF failed: no route toward 192.168.0.4[2 times]
  Created: Wed Oct  3 11:17:01 2012
Total 1 displayed, Up 1, Down 0
```

show mpls lsp extensive (conditional metric)

```
user@host> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.4.4
  From: 192.168.3.3, State: Up, ActiveRoute: 0, LSPname: PE1_to_PE2, LSPid: 2
  ActivePath:  (primary)
  Link protection desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Follow destination IGP metric conditionally
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  LSP Self-ping Status : Enabled
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Flap Count: 0
    MBB Count: 0
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2 )
 192.168.1.1 S 192.168.0.2 S
```

Release Information

Command introduced before Junos OS Release 7.4.

`defaults` option added in Junos OS Release 8.5.

`autobandwidth` option added in Junos OS Release 11.4.

`externally-controlled` option added in Junos OS Release 12.3.

`externally-provisioned` option added in Junos OS Release 13.3.

`instance` *`instance-name`* option added in Junos OS Release 15.1.

RELATED DOCUMENTATION

- clear mpls lsp
- show mpls lsp autobandwidth

show mpls lsp abstract-computation

IN THIS SECTION

- Syntax | 22203
- Description | 22203
- Options | 22203
- Required Privilege Level | 22204
- Output Fields | 22204
- Sample Output | 22205
- Release Information | 22205

Syntax

```
show mpls lsp abstract-computation
<brief | detail | extensive>;
<logical-system (all | logical-system-name)>;
<name lsp-name>;
```

Description

Display the ingress to egress abstract hop computation used by the constrained shortest path in the preprocessing for LSPs. The command output displays the various computation passes involved per LSP, and the qualifying exit devices for each pass. It also displays the affinity per pass, and the current start device chosen for the pass.

Options

brief | detail | extensive	(Optional) Display the desired level of output.
logical-system (all | *logical-system-name*)	(Optional) Display the abstract computation for abstract hop constraints on all logical systems or on a particular logical system.
lsp-name	(Optional) Name of the LSP for which the abstract hop computation is displayed.

Required Privilege Level

view

Output Fields

Table 819 on page 22204 describes the output fields for the `show mpls lsp abstract-computation` command. Output fields are listed in the approximate order in which they appear.

Table 819: show mpls lsp abstract-computation Output Fields

Field Name	Field Description
`Path computation using abstract hops for LSP`	Name of the LSP for which the abstract hop computation is performed.
`Path type`	The type of the path can be primary or secondary.
`Path name`	Name of the path.
`Credibility`	Credibility value associated with the interior gateway protocol in use.
`Total no of CSPF passes`	Number of constrained shortest path passes for the abstract hop.
`CSPF pass no`	Constrained shortest path pass number for the abstract hop computation.
`Start address of the pass`	IP address where the pass starts.
`Affinity`	Name of the abstract hop.
`Destination`	Destination IP address for a node in the pass.
`State`	State of the backtracking: • Valid • Disqualified

Sample Output

show mpls lsp abstract-computation

```
user@R0> show mpls lsp abstract-computation
Path computation using abstract hops for LSP: R0-R31
Path type: Primary, Path name: prim
Credibility: 0, Total no of CSPF passes: 2
CSPF pass no: 0
Start address of the pass: 127.0.0.6
Destination: 127.0.0.1, State: VALID
Destination: 127.0.0.2, State: VALID
Destination: 127.0.0.3, State: VALID
Affinity: ah1
CSPF pass no: 1
Start address of the pass: 127.0.0.1
Destination: 127.0.0.3, State: VALID
Path type: Secondary, Path name: nonstdby
Path type: Standby, Path name: stdby
Credibility: 0, Total no of CSPF passes: 2
CSPF pass no: 0
Start address of the pass: 127.0.0.6
Destination: 127.0.0.3, State: VALID
Destination: 127.0.0.4, State: VALID
Affinity: ah2
CSPF pass no: 1
Start address of the pass: 127.0.0.4
Destination: 127.0.0.3, State: VALID
```

Release Information

Command introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Example: Configuring Abstract Hops for MPLS LSPs

abstract-hop | **1625**

constituent-list | **4571**

show mpls abstract-hop-membership | **22125**

show mpls lsp autobandwidth

IN THIS SECTION

- Syntax | 22206
- Description | 22206
- Options | 22206
- Required Privilege Level | 22207
- Output Fields | 22207
- Sample Output | 22209
- Release Information | 22209

Syntax

```
show mpls lsp autobandwidth
<brief | detail | extensive>
<logical-system (all | logical-system-name)>
<name lsp-name>
```

Description

Display automatic bandwidth information for the LSP(s).

After a Routing Engine switchover, the output of the `show mpls autobandwidth` command might not be up-to-date, as the automatic bandwidth information for the LSP(s) is gathered by the new primary Routing Engine during the first adjustment interval.

Options

brief | detail | extensive (Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

logical-system (all | *logical-system-name*) (Optional) Perform this operation on all logical systems or on a particular logical system.

name *lsp-name* (Optional) Specify name of the LSP for which the automatic bandwidth information should be displayed.

Required Privilege Level

view

Output Fields

Table 820 on page 22207 describes the output fields for the `show mpls lsp autobandwidth` command. Output fields are listed in the approximate order in which they appear.

Table 820: show mpls lsp autobandwidth Output Fields

Field Name	Field Description	Level of Output
`To`	Destination (egress routing device) of the session.	All Levels
`From`	Source (ingress routing device) of the session.	All Levels
`LSPname`	Name of the LSP.	All Levels
`Min BW`	(Ingress LSP) Configured minimum value of the LSP, in bps.	`detail extensive`
`Max BW`	(Ingress LSP) Configured maximum value of the LSP, in bps.	`detail extensive`
`Max AvgBW util`	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps. **NOTE**: In calculating this value, the sample collected during make before break (MBB) is ignored to prevent inaccurate results. The first sample after a bandwidth adjustment, or after a change in the LSP ID (regardless of path change), is also ignored.	`detail extensive`
`Overflow limit`	(Ingress LSP) Configured value of the threshold overflow limit.	`detail extensive`

Table 820: show mpls lsp autobandwidth Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Overflow sample count`	(Ingress LSP) Current value for the overflow sample count.	`detail extensive`
`Underflow limit`	(Ingress LSP) Configured value of the threshold underflow limit.	`detail extensive`
`Underflow sample count`	(Ingress LSP) Current value for the underflow sample count.	`detail extensive`
`Adjustment Timer`	(Ingress LSP) Configured value for the adjust-timer statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	`detail extensive`
`Adjustment Threshold`	(Ingress LSP) Configured value for the adjust-threshold statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	`detail extensive`
`Time for Next Adjustment`	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	`detail extensive`
`Time of Last Adjustment`	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	`detail extensive`
`Last BW`	Previous active bandwidth of the LSP.	`detail extensive`
`Last Requested BW`	Bandwidth requested in the previous automatic bandwidth adjustment.	`detail extensive`
`Last Signaled BW`	Bandwidth signaled in the previous automatic bandwidth adjustment.	`detail extensive`
`Highest Watermark BW`	Maximum bandwidth used by the LSP.	`detail extensive`

Table 820: show mpls lsp autobandwidth Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Total AutoBw Adjustments`	Total number of attempts to adjust automatic bandwidth including failed and successful adjustments.	`detail extensive`
`Successful Adjustments`	Number of successful automatic bandwidth adjustments.	`detail extensive`
`Failed Adjustments`	Number of failed automatic bandwidth adjustments.	`detail extensive`

Sample Output

show mpls lsp autobandwidth

```
user@host> show mpls lsp autobandwidth extensive
To: 10.255.106.133,
From: 10.255.106.135, LSPname: r0-r1
Min BW: 100kbps, Max BW: 0bps, Max AvgBW util: 2.33249Mbps
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 0
Adjustment Timer: 300 sec, Adjustment Threshold: 0
Time for Next Adjustment: 23 sec, Time of Last Adjustment: Fri Jun  3 21:05:37
2011
Last BW: 100kbps, Last Requested BW: 2.2169Mbps, Last Signaled BW: 2.2169Mbps,
Highest Watermark BW: 2.33249Mbps
Total AutoBw Adjustments: 1, Successful Adjustments: 1, Failed Adjustments: 0
```

Release Information

Statement introduced in Junos OS Release 11.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

`name`*`lsp-name`* option introduced in Junos OS Release 18.1R1 for all platforms.

RELATED DOCUMENTATION

show mpls lsp | 22172

Achieving a Make-Before-Break, Hitless Switchover for LSPs

show mpls path

IN THIS SECTION

- Syntax | 22210
- Syntax (EX Series Switches) | 22210
- Description | 22211
- Options | 22211
- Required Privilege Level | 22211
- Output Fields | 22211
- Sample Output | 22212
- Release Information | 22212

Syntax

```
show mpls path
<instance instance-name>
<logical-system (all | logical-system-name)>
<path-name>
```

Syntax (EX Series Switches)

```
show mpls path
<path-name>
```

Description

Display dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options

none	Display standard information about all MPLS LSPs.
instance *instance-name*	(Optional) Display the dynamic MPLS LSP for the specified instance. If *instance-name* is omitted, dynamic MPLS LSP for the master instance is displayed.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
path-name	(Optional) Display information about the specified LSP only.

Required Privilege Level

view

Output Fields

Table 821 on page 22211 describes the output fields for the `show mpls path` command. Output fields are listed in the approximate order in which they appear.

Table 821: show mpls path Output Fields

Field Name	Field Description
Path name	Information about ingress LSPs. Each path has one line of output.
Address	Addresses of the routing devices that form the LSP.
Strict/loose address	Whether the address is a configured as a strict or loose address.

Sample Output

show mpls path

```
user@host> show mpls path
Path name        Address          Strict/loose address
p1               123.456.55.6     Strict
                 123.456.1.6      Loose
p2               191.456.1.4      Strict
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` *`instance-name`* option added in Junos OS Release 15.1.

show mpls srlg

IN THIS SECTION

- Syntax | 22212
- Description | 22213
- Options | 22213
- Required Privilege Level | 22213
- Output Fields | 22213
- Sample Output | 22214
- Release Information | 22214

Syntax

```
show mpls srlg
<logical-systems (all | logical-system-name)>
```

Description

Display Shared Risk Link Group (SRLG) cost and value configuration information.

> **NOTE**: If an SRLG is associated with a link that is used by an ingress LSP in the router, then on deleting the SRLG configuration from that router, the SRLG gets removed from the SRLG table only on the next reoptimization of the LSP. Until then, the output of the `run show mpls srlg` command displays `Unknown-XXX` instead of the SRLG name and a non zero srlg-cost for that SRLG.

Options

logical-system (all | *logical-system-name*)	(Optional) View SRLG configuration information for all logical systems or a particular logical system.

Required Privilege Level

view

Output Fields

Table 822 on page 22213 lists the output fields for the `show mpls srlg` command. Output fields are listed in the approximate order in which they appear.

Table 822: show mpls srlg Output Fields

Field Name	Field Description
SRLG	Name of the SRLG.
Value	A group ID for the SRLG ranging from **1** through **4294967295**.
Cost	A cost for the Shared Risk Link Group (SRLG) ranging from **1** through **65535**.

Sample Output

command-name

```
user@host> show mpls srlg

SRLG             Value      Cost
srlg-a             101        10
```

Release Information

Command introduced before Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring SRLG

show mpls static-lsp

IN THIS SECTION

- Syntax | 22215
- Description | 22215
- Options | 22215
- Required Privilege Level | 22216
- Output Fields | 22216
- Sample Output | 22218
- Release Information | 22219

Syntax

```
show mpls static-lsp
<brief | detail | extensive | terse>
<bypass>
<descriptions>
<down | up>
<ingress>
<instance instance-name>
<logical-system (all | logical-system-name)>
<lsp-type>
<name name>
<statistics>
<transit>
```

Description

Display information about configured and active static Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options

none	Display standard information about all configured and active static MPLS LSPs.
brief | detail | extensive | terse	(Optional) Display the specified level of output. The **extensive** option displays the same information as the **detail** option, but covers the most recent 50 events.
bypass	(Optional) Display LSPs used for protecting other static LSPs.
descriptions	(Optional) Display the MPLS static LSP descriptions. To view this information, you must configure the description statement at the `[edit protocols mpls static-label-switched-path path-name bypass]`, `[edit protocols mpls static-label-switched-path path-name ingress]`, or `[edit protocols mpls static-label-switched-path path-name transit incoming-label]` hierarchy levels. Only static LSPs with a description are displayed.
down | up	(Optional) Display only static LSPs that are inactive or active, respectively.
instance *instance-name*	(Optional) Display information about all configured and active static MPLS LSPs for the specified routing instance. If *instance-name* is omitted, information about all configured and active static MPLS LSPs for the master instance is displayed.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
lsp-type	(Optional) Display information about a particular LSP type: • **bypass**—Sessions for bypass LSPs. • **ingress**—Sessions that originate from this routing device. • **transit**—Sessions that pass through this routing device.
name *name*	(Optional) Display information about the specified static LSP or group of LSPs.
statistics	(Optional) Display accounting information about static LSPs.
transit	(Optional) Display static LSPs transiting this routing device.

Required Privilege Level

view

Output Fields

Table 823 on page 22216 describes the output fields for the `show mpls static-lsp` command. Output fields are listed in the approximate order in which they appear.

Table 823: show mpls static-lsp Output Fields

Field Name	Field Description	Level of Output
Ingress LSPs	Information about the static LSPs on the ingress routing device. Each session has one line of output.	All levels
Transit LSPs	Number of static LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
Bypass LSPs	Information about the bypass LSPs configured on the routing device. Each session has one line of output.	All levels
LSPname	Name of the static LSP.	All levels

Table 823: show mpls static-lsp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
To	Destination (egress routing device) of the session.	All levels
State	State of the static LSP handled by this RSVP session: **Up**, **Dn** (down), or **Restart**.	All levels
Packets	Number of packet transiting the static LSP (**statistics** option only).	All levels
Bytes	Number of bytes transiting the static LSP (**statistics** option only).	All levels
Nexthop	IP address for the next-hop router for the static LSP.	**detail**, **extensive**
Bypass	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels
Link protection desired	Link protection has been requested by the ingress routing device.	**detail**, **extensive**
LabelOperation	Label operation to perform: **Push**, **Pop**, **Swap**.	**detail**, **extensive**
Outgoing-label	Outgoing label to use for the MPLS packet in either push or swap label operations.	**detail**, **extensive**
Created	(Ingress LSP) Date and time the static LSP was created.	**extensive**
Bandwidth	Bandwidth configured for the static LSP.	**detail**, **extensive**
Resv style	(Bypass) RSVP reservation style. This field consists of two parts: the number of active reservations and the reservation style, which can be **FF** (fixed filter), **SE** (shared explicit), or **WF** (wildcard filter).	All levels

Sample Output

show mpls static-lsp extensive

```
user@host> show mpls static-lsp extensive
Ingress LSPs:
LSPname: alpha-to-beta, To: 192.168.14.1
  State: Dn
  Nexthop: 192.168.10.1
  LabelOperation: Push, Outgoing-label: 1000001
  Created: Thu Jan 14 16:44:43 2010
  Bandwidth: 0 bps
Total 1, displayed 1, Up 0, Down 1

Transit LSPs:
Total 0, displayed 0, Up 0, Down 0

Bypass LSPs:
Total 0, displayed 0, Up 0, Down 0
```

show mpls static-lsp statistics ingress

```
user@host> show mpls static-lsp statistics ingress
Ingress LSPs:
LSPname                      To              State     Packets            Bytes
alpha-to-beta                192.168.14.1    Dn             NA               NA
Total 1, displayed 1, Up 0, Down 1
```

show mpls static-lsp (when MPLS stitching is used)

The show mpls static-lsp command was extended in Junos release 14.1X53-D25 to accommodate the stitching feature of MPLS. This example shows the LSP state as 'InProgress' because the LSP is waiting for protocol next-hop resolution. For more information, see

```
user@host> show mpls static-lsp
Ingress LSPs:
Total 0, displayed 0, Up 0, Down 0
Transit LSPs: LSPname                            Incoming-label  State
```

```
to-165                                                          1000000
InProgress
```

Release Information

Command introduced in Junos OS Release 10.1.
Command updated in Junos OS Release 14.1X53-D25 to accommodate the stitching feature of MPLS.

show mpls tunnel manager statistics

IN THIS SECTION

- Syntax | 22219
- Description | 22219
- Options | 22220
- Required Privilege Level | 22220
- Output Fields | 22220
- Sample Output | 22221
- Release Information | 22221

Syntax

```
show mpls tunnel-manager-statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display a running counter of the number of Multiprotocol Label Switching (MPLS) label-switched paths (LSPs) that have been updated in-place without any change to LSP identifier (ID) through in-place LSP update and without any change to RSVP installed in the inet.3 routing table.

Options

none	Display the number of LSPs that have been updated in-place without any change to LSP-ID.
instance ***instance-name***	(Optional) Display the count of the LSPs that have been updated in-place LSPs for the specified instance. If `instance-name` is omitted, LSP statistics for the primary instance are displayed.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 824 on page 22220 describes the output fields for the `show mpls mpls tunnel manager statistics` command. Output fields are listed in the approximate order in which they appear.

Table 824: show mpls tunnel manager statistics Output Fields

Field Name	Field Description	Level of Output
`LSP make-before-break`	Indicates signalling a new LSP path before tearing down the old LSP path. Note that the old and the new LSP instances might traverse the same path.	All levels
`In-place LSP BW increase`	Indicates LSP bandwidth increase through in-place update of RSVP signaling states corresponding to the LSP.	All levels
`In-place LSP BW decrease`	Indicates LSP bandwidth decrease through in-place update of RSVP signaling states corresponding to the LSP.	All levels
`In-place LSP update CSPF failure`	Indicates the in-place LSP auto-bandwidth resizing failure at the CSPF path computation stage.	All levels

Table 824: show mpls tunnel manager statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`In-place LSP update signaling error`	Indicates the in-place LSP auto-bandwidth resizing failure when RSVP signaling error is received.	All levels
`In-place LSP update signaling timeout`	Indicates the in-place LSP auto-bandwidth resizing failure when RSVP signaling takes too long to complete.	All levels

Sample Output

show mpls tunnel-manager-statistics

```
user@host> show tunnel manager statistics
MPLS Tunnel Manager Statistics
  LSP make-before-break                    1
  In-place LSP BW increase                 1
  In-place LSP BW decrease                 1
  In-place LSP update CSPF failure         0
  In-place LSP update signaling error      0
  In-place LSP update signaling timeout    0
```

Release Information

Command introduced before Junos OS Release 20.4R1.

RELATED DOCUMENTATION

clear mpls lsp | 3713

show mpls lsp autobandwidth | 22206

show msdp

IN THIS SECTION

- Syntax | 22222
- Description | 22222
- Options | 22222
- Required Privilege Level | 22223
- Output Fields | 22223
- Sample Output | 22224
- Release Information | 22225

Syntax

```
show msdp
<brief | detail>
<instance instance-name>
<logical-system (all / logical-system-name)>
<peer peer-address>
```

Description

Display Multicast Source Discovery Protocol (MSDP) information.

Options

none	Display standard MSDP information for all routing instances.
brief | detail	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display information for the specified instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

peer ***peer-address*** (Optional) Display information about the specified peer only,

Required Privilege Level

view

Output Fields

Table 825 on page 22223 describes the output fields for the `show msdp` command. Output fields are listed in the approximate order in which they appear.

Table 825: show msdp Output Fields

Field Name	Field Description	Level of Output
Peer address	IP address of the peer.	All levels
Local address	Local address of the peer.	All levels
State	Status of the MSDP connection: **Listen**, **Established**, or **Inactive**.	All levels
Last up/down	Time at which the most recent peer-state change occurred.	All levels
Peer-Group	Peer group name.	All levels
SA Count	Number of source-active cache entries advertised by each peer that were accepted, compared to the number that were received, in the format ***number-accepted/ number-received***.	All levels
Peer Connect Retries	Number of peer connection retries.	**detail**
State timer expires	Number of seconds before another message is sent to a peer.	**detail**
Peer Times out	Number of seconds to wait for a response from the peer before the peer is declared unavailable.	**detail**

Table 825: show msdp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
SA accepted	Number of entries in the source-active cache accepted from the peer.	**detail**
SA received	Number of entries in the source-active cache received by the peer.	**detail**

Sample Output

show msdp

```
user@host> show msdp
Peer address    Local address   State       Last up/down Peer-Group SA Count
198.32.8.193    198.32.8.195    Established  5d 19:25:44 North23    120/150
198.32.8.194    198.32.8.195    Established  3d 19:27:27 North23    300/345
198.32.8.196    198.32.8.195    Established  5d 19:39:36 North23    10/13
198.32.8.197    198.32.8.195    Established  5d 19:32:27 North23    5/6
198.32.8.198    198.32.8.195    Established  3d 19:33:04 North23    2305/3000
```

show msdp brief

The output for the `show msdp brief` command is identical to that for the `show msdp` command. For sample output, see "show msdp" on page 22224.

show msdp detail

```
user@host> show msdp detail
Peer: 10.255.70.15
Local address: 10.255.70.19
State: Established
Peer Connect Retries: 0
State timer expires: 22
Peer Times out: 49
SA accepted: 0
SA received: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show msdp source | 22225
show msdp source-active | 22228
show msdp statistics | 22232

show msdp source

IN THIS SECTION

- Syntax | 22225
- Description | 22225
- Options | 22226
- Required Privilege Level | 22226
- Output Fields | 22226
- Sample Output | 22227
- Release Information | 22227

Syntax

```
show msdp source
<instance instance-name>
<logical-system (all / logical-system-name)>
<source-address>
```

Description

Display multicast sources learned from Multicast Source Discovery Protocol (MSDP).

Options

none	Display standard MSDP source information for all routing instances.
instance *instance-name*	(Optional) Display information for the specified instance only.
logical-system (all | *logical-system-name)*	(Optional) Perform this operation on all logical systems or on a particular logical system.
source-address	(Optional) IP address and optional prefix length. Display information for the specified source address only.

Required Privilege Level

view

Output Fields

Table 826 on page 22226 describes the output fields for the `show msdp source` command. Output fields are listed in the approximate order in which they appear.

Table 826: show msdp source Output Fields

Field Name	Field Description
Source address	IP address of the source.
/Len	Length of the prefix for this IP address.
Type	Discovery method for this multicast source: • **Configured**—Source-active limit explicitly configured for this source. • **Dynamic**—Source-active limit established when this source was discovered.
Maximum	Source-active limit applied to this source.
Threshold	Source-active threshold applied to this source.

Table 826: show msdp source Output Fields *(Continued)*

Field Name	Field Description
Exceeded	Number of source-active messages received from this source exceeding the established maximum.

Sample Output

show msdp source

```
user@host> show msdp source
Source address /Len   Type      Maximum   Threshold    Exceeded
0.0.0.0        /0     Configured        5      none       0
10.1.0.0       /16    Configured      500      none       0
10.1.1.1       /32    Configured    10000      none       0
10.1.1.2       /32    Dynamic        6936      none       0
10.1.5.5       /32    Dynamic         500      none         123
10.2.1.1       /32    Dynamic           2      none       0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show msdp | 22222

show msdp source-active | 22228

show msdp statistics | 22232

show msdp source-active

IN THIS SECTION

- Syntax | 22228
- Description | 22228
- Options | 22228
- Required Privilege Level | 22229
- Output Fields | 22229
- Sample Output | 22230
- Release Information | 22231

Syntax

```
show msdp source-active
<brief | detail>
<group group>
<instance instance-name>
<local>
<logical-system (all | logical-system-name)>
<originator originator>
<peer peer-address>
<source source-address>
```

Description

Display the Multicast Source Discovery Protocol (MSDP) source-active cache.

Options

none	Display standard MSDP source-active cache information for all routing instances.
brief | detail	(Optional) Display the specified level of output.

group *group*	(Optional) Display source-active cache information for the specified group.
instance *instance-name*	(Optional) Display information for the specified instance.
local	(Optional) Display all source-active caches originated by this router.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
originator *originator*	(Optional) Display information about the peer that originated the source-active cache entries.
peer *peer-address*	(Optional) Display the source-active cache of the specified peer.
source *source-address*	(Optional) Display the source-active cache of the specified source.

Required Privilege Level

view

Output Fields

Table 827 on page 22229 describes the output fields for the `show msdp source-active` command. Output fields are listed in the approximate order in which they appear.

Table 827: show msdp source-active Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.

Table 827: show msdp source-active Output Fields *(Continued)*

Field Name	Field Description
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.
Group address	Multicast address of the group.
Source address	IP address of the source.
Peer address	IP address of the peer.
Originator	Router ID configured on the source of the rendezvous point (RP) that originated the message, or the loopback address when the router ID is not configured.
Flags	Flags: **Accept**, **Reject**, or **Filtered**.

Sample Output

show msdp source-active

```
user@host> show msdp source-active
Group address   Source address  Peer address    Originator      Flags
230.0.0.0       192.168.195.46  local           10.255.14.30    Accept
230.0.0.1       192.168.195.46  local           10.255.14.30    Accept
230.0.0.2       192.168.195.46  local           10.255.14.30    Accept
230.0.0.3       192.168.195.46  local           10.255.14.30    Accept
230.0.0.4       192.168.195.46  local           10.255.14.30    Accept
```

show msdp source-active brief

The output for the `show msdp source-active brief` command is identical to that for the `show msdp source-active` command. For sample output, see "show msdp source-active" on page 22230.

show msdp source-active detail

The output for the `show msdp source-active detail` command is identical to that for the `show msdp source-active` command. For sample output, see "show msdp source-active" on page 22230.

show msdp source-active source

```
user@host> show msdp source-active source 192.168.215.246
Global active source limit exceeded: 0
Global active source limit maximum: 25000
Global active source limit threshold: 24000
Global active source limit log-warning: 100
Global active source limit log interval: 0

Group address   Source address  Peer address    Originator      Flags
226.2.2.1       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.3       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.4       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.5       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.7       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.10      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.11      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.13      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.14      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.15      192.168.215.246 10.255.182.140  10.255.182.140  Accept
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show msdp | 22222

show msdp source | 22225

show msdp statistics | 22232

show msdp statistics

IN THIS SECTION

- Syntax | 22232
- Description | 22232
- Options | 22232
- Required Privilege Level | 22233
- Output Fields | 22233
- Sample Output | 22235
- Release Information | 22236

Syntax

```
show msdp statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
<peer peer-address>
```

Description

Display statistics about Multicast Source Discovery Protocol (MSDP) peers.

Options

none	Display statistics about all MSDP peers for all routing instances.
instance *instance-name*	(Optional) Display statistics about a specific MSDP instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
peer *peer-address*	(Optional) Display statistics about a particular MSDP peer.

Required Privilege Level

view

Output Fields

Table 828 on page 22233 describes the output fields for the `show msdp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 828: show msdp statistics Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.
Peer	Address of peer.
Last State Change	How long ago the peer state changed.
Last message received from the peer	How long ago the last message was received from the peer.
RPF Failures	Number of reverse path forwarding (RPF) failures.

Table 828: show msdp statistics Output Fields *(Continued)*

Field Name	Field Description
Remote Closes	Number of times the remote peer closed.
Peer Timeouts	Number of peer timeouts.
SA messages sent	Number of source-active messages sent.
SA messages received	Number of source-active messages received.
SA request messages sent	Number of source-active request messages sent.
SA request messages received	Number of source-active request messages received.
SA response messages sent	Number of source-active response messages sent.
SA response messages received	Number of source-active response messages received.
SA messages with zero Entry Count received	Entry Count is a field within SA message that defines how many source/group tuples are present in the SA message. The counter is incremented each time an SA with an Entry Count of zero is received.
Active source exceeded	Number of times this peer has exceeded configured source-active limits.
Active source Maximum	Configured number of active source messages accepted by this peer.
Active source threshold	Configured threshold on this peer for applying random early discard (RED) to drop some but not all MSDP active source messages.
Active source log-warning	Configured threshold on this peer at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Active source log-interval	Time (in seconds) between consecutive log messages on this peer.

Table 828: show msdp statistics Output Fields *(Continued)*

Field Name	Field Description
Keepalive messages sent	Number of keepalive messages sent.
Keepalive messages received	Number of keepalive messages received.
Unknown messages received	Number of unknown messages received.
Error messages received	Number of error messages received.

Sample Output

show msdp statistics

```
user@host> show msdp statistics
Global active source limit exceeded: 0
Global active source limit maximum: 10
Global active source limit threshold: 8
Global active source limit log-warning: 60
Global active source limit log interval: 60

Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA messages with zero Entry Count received: 0
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 0
Active source Maximum: 10
```

```
Active source threshold: 8
Active source log-warning: 60
Active source log-interval 120
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0
```

show msdp statistics peer

```
user@host> show msdp statistics peer 10.255.182.140
Peer: 10.255.182.140
    Last State Change: 8:19:23 (00:01:08)
    Last message received from peer: 8:20:05 (00:00:26)
    RPF Failures: 0
    Remote Closes: 0
    Peer Timeouts: 0
    SA messages sent: 17
    SA messages received: 16
    SA request messages sent: 0
    SA request messages received: 0
    SA response messages sent: 0
    SA response messages received: 0
    Active source exceeded: 20
    Active source Maximum: 10
    Active source threshold: 8
    Active source log-warning: 60
    Active source log-interval: 120
    Keepalive messages sent: 0
    Keepalive messages received: 0
    Unknown messages received: 0
    Error messages received: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear msdp statistics | 3719

show multi-chassis mc-lag configuration-consistency

IN THIS SECTION

- Syntax | 22237
- Description | 22237
- Options | 22237
- Required Privilege Level | 22238
- Output Fields | 22238
- Sample Output | 22239
- Release Information | 22243

Syntax

```
show multi-chassis mc-lag configuration-consistency (brief | detail)
```

Description

Displays configuration consistency check status for various MC-LAG parameters .

NOTE: This command only displays MC-LAG parameters that are committed.

Options

none Display configuration consistency check status for various MC-LAG parameters.

Required Privilege Level

view

Output Fields

Table 829 on page 22238 lists the output fields for the `show multi-chassis mc-lag configuration-consistency` command. Output fields are listed in the approximate order in which they appear.

Table 829: show multi-chassis mc-lag configuration-consistency Output Fields

Output Field Name	Field Description
`Configuration Item`	Name of the committed MC-LAG parameter.
`Local Physical Interface`	Name of the physical interface configured on the local MC-LAG peer.
`Peer Physical Interface`	Name of the physical interface configured on the remote MC-LAG peer.
`Local Logical Interface`	Name of the logical interface configured on the local MC-LAG peer.
`Peer Logical Interface`	Name of the logical interface configured on the remote MC-LAG peer.
`Local IRB`	Name of the integrated routing and bridging (IRB) interface configured on the local MC-LAG peer.
`Peer IRB`	Name of the integrated routing and bridging (IRB) interface configured on the remote MC-LAG peer.
`Local VLAN`	Name of the VLAN configured on the local MC-LAG peer.
`Peer VLAN`	Name of the VLAN configured on the remote MC-LAG peer.
`Enforcement Level`	Enforcement level for the MC-LAG parameter is `Mandatory` or `Desirable`.

Table 829: show multi-chassis mc-lag configuration-consistency Output Fields *(Continued)*

Output Field Name	Field Description
`Local Value`	Value of the committed MC-LAG parameter for the local MC-LAG peer.
`Peer Value`	Value of the committed MC-LAG parameter for the remote MC-LAG peer.
`Result`	Result of the configuration consistency check of the MC-LAG parameter is `PASS` or `FAIL`.

Sample Output

show multi-chassis mc-lag configuration-consistency

```
user@host> show multi-chassis mc-lag configuration-consistency
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
service-id                          Mandatory           1
1                     PASS
session-establishment-hold-time     Mandatory           300
300                   PASS
local-ip-addr                       Mandatory           10.1.1.1
10.1.1.2              PASS
iccp/bfd multiplier                 Mandatory           3
3                     PASS
iccp/minimum-transmit-interval      Mandatory           60
60                    PASS
bfd minimum-interval                Mandatory           1000
1000                  PASS

Local Physical Interface:ae0
Peer Physical Interface :ae0
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
```

```
------------------                  -----------------   -----------
----------            -------
lacp admin-key                      Mandatory           21
21                    PASS
lacp system-id                      Mandatory           00:00:00:00:00:01
00:00:00:00:00:01     PASS
lacp mode                           Mandatory           0
0                     PASS
prefer-status-control-active        Desirable           TRUE
TRUE                  FAIL
mcae status-control                 Mandatory           standby
active                PASS
mcae deployment mode                Mandatory           active-active         active-
active         PASS
mcae chassis-id                     Mandatory           0
1                     PASS
mcae redundancy-group               Mandatory           101
101                   PASS
force-icl-down                      Mandatory           --
TRUE                  PASS

Local Logical Interface:ae0.0
Peer Logical Interface :ae0.0
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
vlan membership                     Mandatory           501-502
501-502               PASS
interface-mode                      Mandatory           trunk
trunk                 PASS

Local Physical Interface:ae1
Peer Physical Interface :ae1
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
lacp admin-key                      Mandatory           22
22                    PASS
lacp system-id                      Mandatory           00:00:00:00:00:03
00:00:00:00:00:03     PASS
lacp mode                           Mandatory           0
```

```
0                     PASS
prefer-status-control-active         Desirable           TRUE
TRUE                  FAIL
mcae status-control                  Mandatory           standby
active                PASS
mcae deployment mode                 Mandatory           active-active         active-
active         PASS
mcae chassis-id                      Mandatory           0
1                     PASS
mcae redundancy-group                Mandatory           101
101                   PASS
force-icl-down                       Mandatory           --
TRUE                  PASS

Local Logical Interface:ae1.0
Peer Logical Interface :ae1.0
Configuration Item                   Enforcement Level   Local Value           Peer
Value            Result
------------------                   -----------------   -----------
----------            -------
vlan membership                      Mandatory           601-602
601-602               PASS
interface-mode                       Mandatory           trunk
trunk                 PASS

Local Physical Interface:ae5
Peer Physical Interface :ae5

Local Logical Interface:ae5.0
Peer Logical Interface :ae5.0
Configuration Item                   Enforcement Level   Local Value           Peer
Value            Result
------------------                   -----------------   -----------
----------            -------
vlan membership                      Mandatory           501-502,601-602
501-502,601-602       PASS
interface-mode                       Mandatory           trunk
trunk                 PASS

Local VLAN:client-vlan-1
Peer VLAN :client-vlan-1
Configuration Item                   Enforcement Level   Local Value           Peer
Value            Result
```

```
------------------                  -----------------   -----------
----------            -------
mcae-mac-synchronize                Mandatory           TRUE
TRUE                  PASS

Local IRB:irb.501
Peer IRB :irb.501
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
IPv4 Addresses                      Mandatory           10.1.1.1/24
10.1.1.2/24           FAIL

Local VLAN:client-vlan-2
Peer VLAN :client-vlan-2
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
mcae-mac-synchronize                Mandatory           TRUE
TRUE                  PASS

Local IRB:irb.502
Peer IRB :irb.502
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
IPv4 Addresses                      Mandatory           10.1.1.1/24
10.1.1.2/24           FAIL

Local VLAN:server-vlan-1
Peer VLAN :server-vlan-1
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
mcae-mac-synchronize                Mandatory           TRUE
TRUE                  PASS

Local IRB:irb.601
Peer IRB :irb.601
```

```
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
IPv4 Addresses                      Mandatory           10.2.1.1/24
10.2.1.2/24           FAIL

Local VLAN:server-vlan-2
Peer VLAN :server-vlan-2
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
mcae-mac-synchronize                Mandatory           TRUE
TRUE                  PASS

Local IRB:irb.602
Peer IRB :irb.602
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
IPv4 Addresses                      Mandatory           10.3.1.1/24
10.3.1.2/24           FAIL
```

Release Information

Command introduced in Junos OS Release 16.1R1.

show multi-chassis mc-lag configuration-consistency global-config

IN THIS SECTION

- Syntax | 22244
- Description | 22244
- Options | 22244
- Required Privilege Level | 22245

- Output Fields | 22245
- Sample Output | 22245
- Release Information | 22246

Syntax

```
show multi-chassis mc-lag configuration-consistency global-config (brief | detail)
```

Description

View configuration consistency check status for all committed global configuration related to MC-LAG functionality, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.

This command shows only a subset of what is shown in the `show multi-chassis mc-lag configuration-consistency` command. The following parameters related to the global configuration are checked for consistency.

- ICL interface
- RSTP bridge priority
- service ID
- session establishment hold time
- local IP address of the ICCP interface
- backup liveness detection peer IP address
- ICCP/BFD multiplier

NOTE: This command only displays MC-LAG parameters that are committed.

Options

none Display configuration consistency check status for all global configuration related to MC-LAG functionality.

Required Privilege Level

view

Output Fields

Table 830 on page 22245 lists the output fields for the `show multi-chassis mc-lag configuration-consistency global-config` command. Output fields are listed in the approximate order in which they appear.

Table 830: show multi-chassis mc-lag configuration-consistency global-config Output Fields

Output Field Name	Field Description
`Configuration Item`	Name of the committed MC-LAG parameter.
`Enforcement Level`	Enforcement level for the MC-LAG parameter is `Mandatory` or `Desirable`.
`Local Value`	Value of the committed MC-LAG parameter on the local peer.
`Peer Value`	Value of the committed MC-LAG parameter on the remote peer.
`Result`	Result of the configuration consistency check of the MC-LAG parameter is `PASS` or `FAIL`.

Sample Output

show multi-chassis mc-lag configuration-consistency global-config

```
user@host> show multi-chassis mc-lag configuration-consistency global-config
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
service-id                          Mandatory           1
1                    PASS
session-establishment-hold-time     Mandatory           300
300                  PASS
```

```
local-ip-addr                       Mandatory           10.1.1.1
10.1.1.2              PASS
iccp/bfd multiplier                 Mandatory           3
3                     PASS
iccp/minimum-transmit-interval      Mandatory           60
60                    PASS
bfd minimum-interval                Mandatory           1000
1000                  PASS
```

Release Information

Command introduced in Junos OS Release 15.1X53-D60.

show multi-chassis mc-lag configuration-consistency icl-config

IN THIS SECTION

- Syntax | 22246
- Description | 22246
- Options | 22247
- Required Privilege Level | 22247
- Output Fields | 22247
- Sample Output | 22248
- Release Information | 22249

Syntax

```
show multi-chassis mc-lag configuration-consistency icl-config (brief | detail)
```

Description

View configuration consistency check status for parameters related to the ICL, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the

configuration consistency check. The results are either pass or fail. Some example of parameters related to the ICL interface are the interface mode and which VLAN the interface belongs to.

This command shows only a subset of what is shown in the `show multi-chassis mc-lag configuration-consistency` command. The following parameters related to the ICL configuration are checked for consistency check:

- VLAN membership
- interface mode

> **NOTE**: This command only displays MC-LAG parameters that are committed.

Options

none Display configuration consistency check status for MC-LAG parameters related to the interchassis control link.

Required Privilege Level

view

Output Fields

Table 831 on page 22247 lists the output fields for the `show multi-chassis mc-lag configuration-consistency icl-config` command. Output fields are listed in the approximate order in which they appear.

Table 831: show multi-chassis mc-lag configuration-consistency icl-config Output Fields

Output Field Name	Field Description
`Configuration Item`	Name of the committed MC-LAG parameter.
`Local Physical Interface`	Name of the physical interface configured on the local MC-LAG peer.
`Peer Physical Interface`	Name of the physical interface configured on the remote MC-LAG peer.

Table 831: show multi-chassis mc-lag configuration-consistency icl-config Output Fields *(Continued)*

Output Field Name	Field Description
`Local Logical Interface`	Name of the logical interface configured on the local MC-LAG peer.
`Peer Logical Interface`	Name of the logical interface configured on the remote MC-LAG peer.
`Enforcement Level`	Enforcement level for the MC-LAG parameter is `Mandatory` or `Desirable`.
`Local Value`	Value of the committed MC-LAG parameter on the local peer.
`Peer Value`	Value of the committed MC-LAG parameter on the remote peer.
`Result`	Result of the configuration consistency check of the MC-LAG parameter is `PASS` or `FAIL`.

Sample Output

show multi-chassis mc-lag configuration-consistency icl-config

```
user@host> show multi-chassis mc-lag configuration-consistency icl-config
Local Physical Interface:ae5
Peer Physical Interface :ae5

Local Logical Interface:ae5.0
Peer Logical Interface :ae5.0
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
vlan membership                     Mandatory           501-502,601-602
501-502,601-602       PASS
interface-mode                      Mandatory           trunk
trunk                 PASS
```

Release Information

Command introduced in Junos OS Release 15.1X53-D60.

show multi-chassis mc-lag configuration-consistency list-of-parameters

IN THIS SECTION

- Syntax | 22249
- Description | 22249
- Options | 22250
- Required Privilege Level | 22250
- Output Fields | 22250
- Sample Output | 22258
- Release Information | 22262

Syntax

```
show multichassis configuration-check list-of-parameters
```

Description

Displays the list of MC-LAG parameters (referred to as configuration knobs in the CLI) that are checked for consistency across MC-LAG peers. There are certain parameters that must identical and others that must be unique on both peers. Enforcement of the consistency check for the parameters is either mandatory or desired. If the enforcement is mandatory, and you have not configured the parameter correctly, the multichassis aggregated Ethernet interface (MC-AE) interface will not come up. If the enforcement is desired, and you have not configured the parameter correctly, the MC-AE interface will come up, but performance might be sub-optimal. In this situation, the system will issue a syslog message.

The following list provides the hierarchies in which the MC-LAG parameters are configured:

- `ICL ifd`

Specifies configuration parameters related to the interchassis control link

- `ICCP Peer`

  Specifies configuration parameters related to ICCP functionality

- `IRB Interface`

  Specifies configuration parameters related to the integrated routing and bridging interface

- `MCAE IFBD`

  Specifies configuration parameters related to the VLAN membership of a given MC-AE interface

- `MCAE ifd`

  Specifies configuration parameters related to a given MC-AE interface

- `MCAE ifl`

  Specifies configuration parameters related to a given MC-AE logical interface

- `VLAN`

  Specifies configuration parameters related to a given VLAN

- `VRRP Group`

  Specifies configuration parameters related to a VRRP session

Options

There are no options for this command.

Required Privilege Level

view

Output Fields

Table 832 on page 22251 lists the output fields for the `show multi-chassis mc-lag configuration-consistency list-of-parameters` command.

Table 832: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields

Configuration Knob	Hierarchy	Consistency R
`session-establishment-hold-time` Specify the time during which an Inter-Chassis Control Protocol (ICCP) connection must be established between peers.	`Global, ICCP Peer`	`Identical`
`mac-limit` Specify the maximum number of MAC addresses to be associated with a VLAN—the default is unlimited, which can leave the network vulnerable to flooding.	`Global`	`Identical`
`mac-aging-timer` Specify how long MAC addresses remain in the Ethernet switching table.	`Global`	`Identical`
`arp-aging-timer` Specify the ARP aging timer in minutes for a logical interface of `inet`.	`Global`	`Identical`
`rstp-system-identifier` Specify different bridge identifiers for different RSTP routing instances.	`Global`	`Identical`
`mstp-system-identifier` Specify different bridge identifiers for different MSTP routing instances.	`Global`	`Identical`
`rstp-bridge-priority` Determine which bridge is elected as the root bridge for RSTP. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.	`Global`	`Identical`

Table 832: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields *(Continued)*

Configuration Knob	Hierarchy	Consistency R
`mstp-bridge-priority` Determine which bridge is elected as the root bridge for MSTP. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.	`Global`	`Identical`
`rstp-bpdu-block-on-edge` Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch for RSTP.	`Global`	`Identical`
`vstp-bpdu-block-on-edge` Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch for VSTP.	`Global`	`Identical`
`mstp-bpdu-block-on-edge` Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch for MSTP.	`Global`	`Identical`
`service-id` Specify a service identifier for each multichassis aggregated Ethernet interface that belongs to a link aggregation group (LAG).	`Global`	`Identical`
`bfd minimum-interval` Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session.	`ICCP Peer`	`Identical`

Table 832: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields *(Continued)*

Configuration Knob	Hierarchy	Consistency R
`iccp/minimum-transmit-interval` Specify the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.	`ICCP Peer`	`Identical`
`iccp/minimum-receive-interval` Specify the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session.	`ICCP Peer`	`Identical`
`iccp/bfd multiplier` Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.	`ICCP Peer`	`Identical`
`iccp single-hop` Configure single hop BFD sessions.	`ICCP Peer`	`Identical`
`iccp/authentication-key` Specify the authentication key password to verify the authenticity of packets sent from the peers hosting an MC-LAG.	`ICCP Peer`	`Identical`
`redundancy-group-id-list` Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions.	`ICCP Peer`	`Identical`

Table 832: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields *(Continued)*

Configuration Knob	Hierarchy	Consistency R
`backup-liveness-detection` Determine whether a peer is up or down by exchanging keepalive messages over the management link between the two Inter-Chassis Control Protocol (ICCP) peers.	`ICCP Peer`	`Unique`
`mc-ae-id` Specify the identification number of the MC-LAG device.	`MCAE ifd`	`Identical`
`mcae redundancy-group` Used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other.	`MCAE ifd`	`Identical`
`mcae chassis-id` Used by LACP for calculating the port number of the MC-LAG's physical member links.	`MCAE ifd`	`Unique`
`mcae deployment mode` Indicates whether an MC-LAG is in active-standby mode or active-active mode.	`MCAE ifd`	`Identical`
`mcae status-control` Specify whether the chassis becomes active or remains in standby mode when an interchassis link failure occurs.	`MCAE ifd`	`Unique`
`force-icl-down` Specify that if the node's ICCP peer goes down, the system brings down the interchassis-link logical interface.	`MCAE ifd`	`Unique`

Table 832: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields *(Continued)*

Configuration Knob	Hierarchy	Consistency R
`prefer-status-control-active` Specify that the node configured as `status-control active` becomes the active node if the peer of this node goes down.	`MCAE ifd`	`Unique`
`lacp mode` Specify LACP is active or passive.	`MCAE ifd`	`Identical`
`lacp periodic` Specify the interval for periodic transmission of LACP packets.	`MCAE ifd`	`Identical`
`lacp system-id` Define the LACP system identifier at the aggregated Ethernet interface level.	`MCAE ifd`	`Identical`
`lacp admin-key` Specify an administrative key for the router or switch.	`MCAE ifd`	`Identical`
`native-vlan-id` Configure mixed tagging support for untagged packets on a port.	`MCAE ifd`	`Identical`
`mcae-mac-synchronize` Synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG.	`VLAN`	`Identical`

Table 832: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields *(Continued)*

Configuration Knob	Hierarchy	Consistency R
`Interface mac Limit` Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.	`VLAN`	`Identical`
`l3-interface` Associate a Layer 3 interface with the VLAN.	`VLAN`	`Identical`
`igmp-snooping` Enable IGMP snooping. A Layer 2 device monitors the IGMP join and leave messages sent from each connected host to a multicast router. This enables the Layer 2 device to keep track of the multicast groups and associated member ports. The Layer 2 device uses this information to make intelligent decisions and to forward multicast traffic to only the intended destination hosts.	`VLAN`	`Identical`
`family` Specify the protocol family configured on the logical interface.	`IRB Interface`	`Identical`
`ipv4 address` Specify an IPv4 address for the IRB interface.	`IRB Interface`	`Unique`
`ipv6 address` Specify an IPv6 address for the IRB interface.	`IRB Interface`	`Unique`
`vrrp-group id` Specify a VRRP group identifier.	`IRB Interface`	`Identical`

Table 832: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields *(Continued)*

Configuration Knob	Hierarchy	Consistency R
`proxy-arp-type` For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.	`IRB Interface`	`Identical`
`vrrp-group priority` Configure a Virtual Router Redundancy Protocol (VRRP) router's priority for becoming the primary default router. The router with the highest priority within the group becomes the primary.	`VRRP Group`	`Unique`
`vrrp-group authentication-key` Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the authentication-type statement.	`VRRP Group`	`Identical`
`vrrp-group authentication-type` Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group.	`VRRP Group`	`Identical`
`vrrp-group virtual-address` Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group.	`VRRP Group`	`Identical`
`encapsulation` Configure a logical link-layer encapsulation type.	`MCAE ifd`	`Identical`

Table 832: show multi-chassis mc-lag configuration-consistency list-of-parameters Output Fields *(Continued)*

Configuration Knob	Hierarchy	Consistency R
`flexible-vlan-tagging` Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.	`MCAE ifd`	`Identical`
`vlan-tagging` For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.	`MCAE ifd`	`Identical`
`mtu` Specify the maximum transmission unit (MTU) size for the media or protocol.	`MCAE ifd`, `ICL ifd`	`Identical`
`interface-mode` Determine whether the logical interface accepts or discards packets based on VLAN tags.	`MCAE ifl`	`Identical`
`vlan membership` Specify the name of the VLAN that belongs to an interface.	`MCAE ifl`	`Identical`

Sample Output

show multi-chassis mc-lag configuration-consistency list-of-parameters

```
user@host> show multi-chassis mc-lag configuration-consistency list-of-parameters
Possible completions:
  configuration-consistency  Show all configuration consistency information
```

```
regress@liki-pe2_re0> show multi-chassis mc-lag configuration-consistency list-of-parameters
#Item   Configuration Knob                   Hierarchy         Consistency Requirement
Enforcement
------  -----------------                   ---------         -----------------------
-----------
0       local-ip-addr                        Global                Unique
Mandatory
1       session-establishment-hold-time      Global                Identical
Mandatory
2       local-ip-addr                        ICCP Peer             Unique
Mandatory
3       session-establishment-hold-time      ICCP Peer             Identical
Mandatory
5       bfd minimum-interval                 ICCP Peer             Identical
Mandatory
6       iccp/minimum-transmit-interval       ICCP Peer             Identical
Mandatory
7       iccp/minimum-receive-interval        ICCP Peer             Identical
Mandatory
8       iccp/bfd multiplier                  ICCP Peer             Identical
Mandatory
9       iccp single-hop                      ICCP Peer             Identical
Mandatory
11      iccp/authentication-key              ICCP Peer             Identical
Mandatory
4       redundancy-group-id-list             ICCP Peer             Identical
Mandatory
12      backup-liveness-detection            ICCP Peer             Unique
Mandatory
13      service-id                           Global                Identical
Mandatory
14      mac-limit                            Global                Identical
Desired
15      mac-ageing-timer                     Global                Identical
Desired
16      arp-ageing-timer                     Global                Identical
Desired
17      rstp-bpdu-block-on-edge              Global                Identical
Desired
18      rstp-bridge-priority                 Global                Identical
Desired
19      rstp-system-identifier               Global                Identical
Desired
```

```
20      vstp-bpdu-block-on-edge              Global                Identical
Desired
21      mstp-bpdu-block-on-edge              Global                Identical
Desired
22      mstp-bridge-priority                 Global                Identical
Desired
23      mstp-system-identifier               Global                Identical
Desired
24      mc-ae-id                             MCAE ifd              Identical
Mandatory
25      mcae redundancy-group                MCAE ifd              Identical
Mandatory
26      mcae chassis-id                      MCAE ifd              Unique
Mandatory
27      mcae deployment mode                 MCAE ifd              Identical
Mandatory
28      mcae status-control                  MCAE ifd              Unique
Mandatory
29      force-icl-down                       MCAE ifd              Unique
Mandatory
30      prefer-status-control-active         MCAE ifd              Unique
Desired
31      lacp mode                            MCAE ifd              Identical
Mandatory
32      lacp periodic                        MCAE ifd              Identical
Mandatory
33      lacp system-id                       MCAE ifd              Identical
Mandatory
34      lacp admin-key                       MCAE ifd              Identical
Mandatory
59      vlan id list                         VLAN                  Identical
Mandatory
60      vlan-ids                             VLAN                  Identical
Mandatory
62      Interface mac Limit                  VLAN                  Identical
Desired
58      service-id                           VLAN                  Identical
Mandatory
64      igmp-snooping-enabled                VLAN                  Identical
Mandatory
61      mcae-mac-synchronize                 VLAN                  Identical
Mandatory
45      l3-interface                         VLAN                  Identical
```

```
Desired
47      ipv4 address                        IRB Interface         Unique
Mandatory
48      ipv6 address                        IRB Interface         Unique
Mandatory
49      vrrp-group id                       IRB Interface         Identical
Mandatory
53      vrrp-group priority                 VRRP Group            Unique
Mandatory
51      vrrp-group authentication-key       VRRP Group            Identical
Mandatory
52      vrrp-group authentication-type      VRRP Group            Identical
Mandatory
50      vrrp-group virtual-address          VRRP Group            Identical
Mandatory
54      proxy-arp-type                      IRB Interface         Identical
Mandatory
35      encapsulation                       MCAE ifd              Identical
Mandatory
36      flexible-vlan-tagging               MCAE ifd              Identical
Mandatory
37      vlan-tagging                        MCAE ifd              Identical
Mandatory
38      mtu                                 MCAE ifd              Identical
Mandatory
39      native-vlan-id                      MCAE ifd              Identical
Mandatory
40      family                              MCAE ifl              Identical
Mandatory
42      interface-mode                      MCAE ifl              Identical
Mandatory
43      vlans                               MCAE ifl              Identical
Mandatory
44      vlan membership                     MCAE ifl              Identical
Mandatory
65      ICL interface                       ICL ifd               Identical
Mandatory
65      ICL interface                       ICL ifd               Identical
Mandatory
65      ICL interface                       ICL ifd               Identical
Mandatory
68      encapsulation                       ICL ifd               Identical
Mandatory
```

```
69      flexible-vlan-tagging                ICL ifd              Identical
Mandatory
71      vlan-tagging                         ICL ifd              Identical
Mandatory
70      mtu                                  ICL ifd              Identical
Mandatory
72      native-vlan-id                       ICL ifd              Identical
Mandatory
73      family                               ICL ifl              Identical
Mandatory
75      interface-mode                       ICL ifl              Identical
Mandatory
76      vlans                                ICL ifl              Identical
Mandatory
77      vlan membership                      ICL ifl              Identical
Mandatory
```

Release Information

Command introduced in Junos OS Release 16.1R1.

show multi-chassis mc-lag configuration-consistency mcae-config

IN THIS SECTION

- Syntax | 22263
- Description | 22263
- Options | 22263
- Required Privilege Level | 22264
- Output Fields | 22264
- Sample Output | 22265
- Release Information | 22266

Syntax

```
show multi-chassis mc-lag configuration-consistency mcae-config (brief | detail)
```

Description

View configuration consistency check status for committed parameters related to the multichassis aggregated Ethernet interfaces, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.

This command shows only a subset of what is shown in the `show multi-chassis mc-lag configuration-consistency` command. The following parameters related to the MC-AE interfaces are checked for consistency:

- LACP administrative key
- LACP system ID
- LACP periodic interval
- prefer status control setting
- status control setting
- mode
- chassis ID
- redundancy group ID
- VLAN membership of the ICL
- interface mode of the ICL

NOTE: This command only displays MC-LAG parameters that are committed.

Options

none Display configuration consistency check status for MC-LAG parameters related to the multichassis aggregated Ethernet interface.

Required Privilege Level

view

Output Fields

Table 833 on page 22264 lists the output fields for the `show multi-chassis mc-lag configuration-consistency mcae-config` command. Output fields are listed in the approximate order in which they appear.

Table 833: show multi-chassis mc-lag configuration-consistency mcae-config Output Fields

Output Field Name	Field Description
`Configuration Item`	Name of the committed MC-LAG parameter.
`Local Physical Interface`	Name of the physical interface configured on the local MC-LAG peer.
`Peer Physical Interface`	Name of the physical interface configured on the remote MC-LAG peer.
`Local Logical Interface`	Name of the logical interface configured on the local MC-LAG peer.
`Peer Logical Interface`	Name of the logical interface configured on the remote MC-LAG peer.
`Enforcement Level`	Enforcement level for the MC-LAG parameter is `Mandatory` or `Desirable`.
`Local Value`	Value of the committed MC-LAG parameter on the local peer.
`Peer Value`	Value of the committed MC-LAG parameter on the remote peer.
`Result`	Result of the configuration consistency check of the MC-LAG parameter is `PASS` or `FAIL`.

Sample Output

show multi-chassis mc-lag configuration-consistency mcae-config

```
user@host> show multi-chassis mc-lag configuration-consistency mcae-config
Local Physical Interface:ae0
Peer Physical Interface :ae0
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
lacp admin-key                      Mandatory           21
21                    PASS
lacp system-id                      Mandatory           00:00:00:00:00:01
00:00:00:00:00:01     PASS
lacp mode                           Mandatory           0
0                     PASS
prefer-status-control-active        Desirable           TRUE
TRUE                  FAIL
mcae status-control                 Mandatory           standby
active                PASS
mcae deployment mode                Mandatory           active-active         active-
active         PASS
mcae chassis-id                     Mandatory           0
1                     PASS
mcae redundancy-group               Mandatory           101
101                   PASS
force-icl-down                      Mandatory           --
TRUE                  PASS

Local Logical Interface:ae0.0
Peer Logical Interface :ae0.0
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
vlan membership                     Mandatory           501-502
501-502               PASS
interface-mode                      Mandatory           trunk
trunk                 PASS

Local Physical Interface:ae1
```

```
Peer Physical Interface :ae1
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
lacp admin-key                      Mandatory           22
22                    PASS
lacp system-id                      Mandatory           00:00:00:00:00:03
00:00:00:00:00:03     PASS
lacp mode                           Mandatory           0
0                     PASS
prefer-status-control-active        Desirable           TRUE
TRUE                  FAIL
mcae status-control                 Mandatory           standby
active                PASS
mcae deployment mode                Mandatory           active-active         active-
active         PASS
mcae chassis-id                     Mandatory           0
1                     PASS
mcae redundancy-group               Mandatory           101
101                   PASS
force-icl-down                      Mandatory           --
TRUE                  PASS

Local Logical Interface:ae1.0
Peer Logical Interface :ae1.0
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
vlan membership                     Mandatory           601-602
601-602               PASS
interface-mode                      Mandatory           trunk
trunk                 PASS
```

Release Information

Command introduced in Junos OS Release 15.1X53-D60.

show multi-chassis mc-lag configuration-consistency vlan-config

IN THIS SECTION

- Syntax | 22267
- Description | 22267
- Options | 22268
- Required Privilege Level | 22268
- Output Fields | 22268
- Sample Output | 22269
- Release Information | 22270

Syntax

```
show multi-chassis mc-lag configuration-consistency vlan-config (brief | detail)
```

Description

View configuration consistency check status for committed parameters related to MC-LAG VLAN configuration, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.

This command shows only a subset of what is shown in the `show multi-chassis mc-lag configuration-consistency` command. The following parameters related to the VLAN and IRB configuration are checked for consistency:

- VRRP group ID
- IP address of IRB interface

NOTE: This command only displays MC-LAG parameters that are committed.

Options

none Display configuration consistency check status for MC-LAG parameters related to VLAN configuration.

Required Privilege Level

view

Output Fields

Table 834 on page 22268 lists the output fields for the `show multi-chassis mc-lag configuration-consistency vlan-config` command. Output fields are listed in the approximate order in which they appear.

Table 834: show multi-chassis mc-lag configuration-consistency vlan-config Output Fields

Output Field Name	Field Description
`Configuration Item`	Name of the committed MC-LAG parameter.
`Local Physical Interface`	Name of the physical interface configured on the local MC-LAG peer.
`Peer Physical Interface`	Name of the physical interface configured on the remote MC-LAG peer.
`Local Logical Interface`	Name of the logical interface configured on the local MC-LAG peer.
`Peer Logical Interface`	Name of the logical interface configured on the remote MC-LAG peer.
`Local IRB`	Name of the integrated routing and bridging (IRB) interface configured on the local MC-LAG peer.
`Peer IRB`	Name of the integrated routing and bridging (IRB) interface configured on the remote MC-LAG peer.
`Local VLAN`	Name of the VLAN configured on the local MC-LAG peer.

Table 834: show multi-chassis mc-lag configuration-consistency vlan-config Output Fields *(Continued)*

Output Field Name	Field Description
Peer VLAN	Name of the VLAN configured on the remote MC-LAG peer.
Enforcement Level	Enforcement level for the MC-LAG parameter is Mandatory or Desirable.
Local Value	Value of the committed MC-LAG parameter on the local peer.
Peer Value	Value of the committed MC-LAG parameter on the remote peer.
Result	Result of the configuration consistency check of the MC-LAG parameter is PASS or FAIL.

Sample Output

show multi-chassis mc-lag configuration-consistency vlan-config

```
user@host> show multi-chassis mc-lag configuration-consistency vlan-config
Local VLAN:client-vlan-1
Peer VLAN :client-vlan-1

Local IRB:irb.501
Peer IRB :irb.501
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
vrrp-group id                       Mandatory           11
11                    PASS
IPv4 Addresses                      Mandatory           10.1.1.2/8
10.1.1.1/8           PASS

Local VLAN:client-vlan-2
Peer VLAN :client-vlan-2
```

```
Local IRB:irb.502
Peer IRB :irb.502
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
vrrp-group id                       Mandatory           12
12                    PASS
IPv4 Addresses                      Mandatory           10.0.1.2/8
10.0.1.1/8           PASS
```

Release Information

Command introduced in Junos OS Release 15.1X53-D60.

show multi-chassis mc-lag configuration-consistency vrrp-config

IN THIS SECTION

- Syntax | 22270
- Description | 22271
- Options | 22271
- Required Privilege Level | 22271
- Output Fields | 22271
- Sample Output | 22272
- Release Information | 22273

Syntax

```
show multi-chassis mc-lag configuration-consistency vrrp-config
```

Description

View configuration consistency check status for committed parameters related to VRRP configuration, the consistency requirement (identical or unique), the enforcement level (mandatory or desired), and the result of the configuration consistency check. The results are either pass or fail.

This command shows only a subset of what is shown in the `show multi-chassis mc-lag configuration-consistency` command. The following parameters related to the VRRP configuration are checked for consistency: VRRP group virtual IP address and VRRP group priority value.

NOTE: This command only displays MC-LAG parameters that are committed.

Options

none Displays configuration consistency check status for MC-LAG parameters related to Virtual Router Redundancy Protocol (VRRP) configuration.

Required Privilege Level

view

Output Fields

Table 835 on page 22271 lists the output fields for the `show multi-chassis mc-lag configuration-consistency vrrp-config` command. Output fields are listed in the approximate order in which they appear.

Table 835: show multi-chassis mc-lag configuration-consistency vrrp-config Output Fields

Output Field Name	Field Description
`Configuration Item`	Name of the committed MC-LAG parameter.
`Local Physical Interface`	Name of the physical interface configured on the local MC-LAG peer.
`Peer Physical Interface`	Name of the physical interface configured on the remote MC-LAG peer.

Table 835: show multi-chassis mc-lag configuration-consistency vrrp-config Output Fields *(Continued)*

Output Field Name	Field Description
`Local Logical Interface`	Name of the logical interface configured on the local MC-LAG peer.
`Peer Logical Interface`	Name of the logical interface configured on the remote MC-LAG peer.
`Local IRB`	Name of the integrated routing and bridging (IRB) interface configured on the local MC-LAG peer.
`Peer IRB`	Name of the integrated routing and bridging (IRB) interface configured on the remote MC-LAG peer.
`Local VLAN`	Name of the VLAN configured on the local MC-LAG peer.
`Peer VLAN`	Name of the VLAN configured on the remote MC-LAG peer.
`Enforcement Level`	Enforcement level for the MC-LAG parameter is `Mandatory` or `Desirable`.
`Local Value`	Value of the committed MC-LAG parameter on the local peer.
`Peer Value`	Value of the committed MC-LAG parameter on the remote peer.
`Result`	Result of the configuration consistency check of the MC-LAG parameter is `PASS` or `FAIL`.

Sample Output

show multi-chassis mc-lag configuration-consistency vrrp-config

```
user@host> show multi-chassis mc-lag configuration-consistency vrrp-config
Lcal VRRP Group:11
```

```
Peer VRRP Group :11
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
vrrp-group virtual-address          Mandatory           010.001.001.010
010.001.001.010       PASS
vrrp-group priority                 Mandatory           202
201                   PASS

Local VRRP Group:12

Peer VRRP Group :12
Configuration Item                  Enforcement Level   Local Value           Peer
Value            Result
------------------                  -----------------   -----------
----------            -------
vrrp-group virtual-address          Mandatory           011.001.001.010
011.001.001.010       PASS
vrrp-group priority                 Mandatory           202
201                   PASS
```

Release Information

Command introduced in Junos OS Release 15.1X53-D60.

show multicast backup-pe-groups

IN THIS SECTION

- Syntax | 22274
- Description | 22274
- Options | 22274
- Required Privilege Level | 22274
- Output Fields | 22275
- Sample Output | 22275

Release Information | 22276

Syntax

```
show multicast backup-pe-groups
<address pe-address>
<group group-name>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display backup PE router group information when ingress PE redundancy is configured. Ingress PE redundancy provides a backup resource when point-to-multipoint LSPs are configured for multicast distribution.

Options

none	Display standard information about all backup PE groups.
address *pe-address*	(Optional) Display the groups that a PE address is associated with.
group *group*	(Optional) Display the backup PE group information for a particular group.
instance *instance-name*	(Optional) Display backup PE group information for a specific multicast instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 836 on page 22275 describes the output fields for the `show multicast backup-pe-groups` command. Output fields are listed in the approximate order in which they appear.

Table 836: show multicast backup-pe-groups Output Fields

Field Name	Field Description
Backup PE Group	Group name.
Designated PE	Primary PE router. Address of the PE router that is currently forwarding traffic on the static route.
Transitions	Number of times that the designated PE router has transitioned from the most eligible PE router to a backup PE router and back again to the most eligible PE router.
Last Transition	Time of the most recent transition.
Local Address	Address of the local PE router.
Backup PE List	List of PE routers that are configured to be backups for the group.

Sample Output

show multicast backup-pe-groups

```
user@host> show multicast backup-pe-groups
Instance: master

Backup PE group: b1
    Designated PE:   10.255.165.7
    Transitions:     1
    Last Transition: 03:15:01
    Local Address:   10.255.165.7
    Backup PE List:
                     10.255.165.8
```

```
Backup PE group: b2
    Designated PE:   10.255.165.7
    Transitions:     2
    Last Transition: 02:58:20
    Local Address:   10.255.165.7
    Backup PE List:
                     10.255.165.9
                     10.255.165.8
```

Release Information

Command introduced in Junos OS Release 9.0.

show multicast ecid-mapping satellite

IN THIS SECTION

- Syntax | 22276
- Description | 22277
- Options | 22277
- Required Privilege Level | 22277
- Output Fields | 22277
- Sample Output | 22278
- Release Information | 22280

Syntax

```
show multicast ecid-mapping satellite
    <brief | detail | extensive>
    <device-id device-id>
    <ecid ecid>
    <reference-id reference-id>
```

Description

Displays entries in the satellite multicast ECID mapping database.

ECID database entries map a group of extended ports to an ECID value for the satellite devices in a Junos Fusion. Each entry also records the next hop to reach the corresponding destination extended ports.

Options

brief | detail | extensive	(Optional) Display the specified level of output.
device-id *device-ID*	Display information from the ECID database for a specified satellite device ID.
ecid *ecid*	Display information from the ECID database for a specified ECID.
reference-id *reference ID*	Display information from the ECID database for a specified internally-assigned reference ID related to the ECID request messages exchanged during ECID allocation (used for troubleshooting issues with ECID allocation).

Required Privilege Level

view

Output Fields

Table 837 on page 22277 lists the output fields for the `show multicast ecid-mapping satellite` command. Output fields are listed in the approximate order in which they appear.

Table 837: show multicast ecid-mapping satellite Command Output Fields

Field Name	Field Description	Level of Output
`Satellite Device ID`	Satellite device ID.	All
`ECID`	ECID assigned to a grouping of extended ports on the satellite device. Satellite devices have multiple ECIDs assigned that represent different groups of extended ports.	All

Table 837: show multicast ecid-mapping satellite Command Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Status of requesting and allocating the ECID. Values include: • Init—An ECID mapping entry for a group of ports has been created. • Request Scheduled—An ECID mapping entry has been queued for an ECID value request. • Request Sent—The ECID request has been sent to the corresponding satellite device. • Response Received—An ECID value has been received from the satellite device. • Sync—An ECID mapping with an ECID value is to be synchronized with a satellite device. • Ready—An ECID next hop has been created and is ready to be referenced for forwarding.	All
Flags	Internal flag values for troubleshooting.	All
nhid	Satellite device next-hop ID associated with the ECID.	All
Reference ID	Internal reference ID assigned to an ECID request message, used for troubleshooting ECID communication with satellite devices.	All (with specified **device-id** or ecid)

Sample Output

show multicast ecid-mapping satellite

```
user@host> show multicast ecid-mapping satellite
Satellite   ECID        State                     Flags
Device ID
100         4129        Ready [nhid=612]          0x0
```

```
                xe-100/0/6.0
                xe-100/0/7.0
100         4097        Ready [nhid=1061]         0x0
                xe-100/0/0.0
                xe-100/0/1.0
100         4122        Ready [nhid=1190]         0x0
                xe-100/0/2.0
                xe-100/0/3.0
                xe-100/0/4.0
                xe-100/0/5.0
103         4103        Ready [nhid=1062]         0x0
                xe-103/0/5.0
                xe-103/0/6.0
103         4104        Ready [nhid=1068]         0x0
                xe-103/0/7.0
                xe-103/0/8.0
103         4105        Ready [nhid=1069]         0x0
                xe-103/0/10.0
                xe-103/0/9.0
103         4106        Ready [nhid=1070]         0x0
                xe-103/0/11.0
                xe-103/0/12.0
103         4107        Ready [nhid=1071]         0x0
                xe-103/0/13.0
                xe-103/0/14.0
103         4109        Ready [nhid=1097]         0x0
                xe-103/0/15.0
                xe-103/0/16.0
103         4110        Ready [nhid=1074]         0x0
                xe-103/0/17.0
                xe-103/0/18.0
103         4111        Ready [nhid=1075]         0x0
                xe-103/0/19.0
                xe-103/0/20.0
103         4112        Ready [nhid=1076]         0x0
                xe-103/0/21.0
                xe-103/0/22.0
```

show multicast ecid-mapping satellite (for specified satellite device-id and ECID)

```
user@host> show multicast ecid-mapping satellite device-id 100 ecid 4101
Satellite   ECID        State                     Reference ID
  Device ID
  100         4101        Ready [nhid=1845]         14
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding Multicast Replication in a Junos Fusion

Egress (Local) Replication on the Satellite Devices

show multicast flow-map

IN THIS SECTION

- Syntax | 22281
- Syntax (EX Series Switch and the QFX Series) | 22281
- Description | 22281
- Options | 22281
- Required Privilege Level | 22281
- Output Fields | 22281
- Sample Output | 22282
- Sample Output | 22283
- Release Information | 22283

Syntax

```
show multicast flow-map
<brief | detail>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show multicast flow-map
<brief | detail>
```

Description

Display configuration information about IP multicast flow maps.

Options

none	Display configuration information about IP multicast flow maps on all systems.
brief | detail	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 838 on page 22282 describes the output fields for the `show multicast flow-map` command. Output fields are listed in the approximate order in which they appear.

Table 838: show multicast flow-map Output Fields

Field Name	Field Description	Levels of Output
Name	Name of the flow map.	All levels
Policy	Name of the policy associated with the flow map.	All levels
Cache-timeout	Cache timeout value assigned to the flow map.	All levels
Bandwidth	Bandwidth setting associated with the flow map.	All levels
Adaptive	Whether or not adaptive mode is enabled for the flow map.	none
Flow-map	Name of the flow map.	**detail**
Adaptive Bandwidth	Whether or not adaptive mode is enabled for the flow map.	**detail**
Redundant Sources	Redundant sources defined for the same destination group.	**detail**

Sample Output

show multicast flow-map

```
user@host> show multicast flow-map
Instance: master
Name            Policy                  Cache timeout        Bandwidth Adaptive
map2            policy2                         never          2000000 no
map1            policy1                    60 seconds          2000000 no
```

Sample Output

show multicast flow-map detail

```
user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
    Policy:             policy1
    Cache Timeout:      600 seconds
    Bandwidth:          2000000
    Adaptive Bandwidth: yes
    Redundant Sources:  10.11.11.11
    Redundant Sources:  10.11.11.12
    Redundant Sources:  10.11.11.13
```

Release Information

Command introduced in Junos OS Release 8.2.

show multicast forwarding-cache statistics

IN THIS SECTION

- Syntax | 22284
- Description | 22284
- Options | 22284
- Required Privilege Level | 22284
- Output Fields | 22284
- Sample Output | 22285
- Release Information | 22286

Syntax

```
show multicast forwarding-cache statistics
<inet | inet6>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display IP multicast forwarding cache statistics.

Options

none	Display multicast forwarding cache statistics for all supported address families for all routing instances.
inet | inet6	(Optional) Display multicast forwarding cache statistics for IPv4 or IPv6 family addresses, respectively.
instance *instance-name*	(Optional) Display multicast forwarding cache statistics for a specific routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 839 on page 22285 describes the output fields for the `show multicast forwarding-cache statistics` command. Output fields are listed in the approximate order in which they appear.

Table 839: show multicast forwarding-cache statistics Output Fields

Field Name	Field Description
Instance	Name of the routing instance for which multicast forwarding cache statistics are displayed.
Family	Protocol family for which multicast forwarding cache statistics are displayed: **ALL**, **INET**, or **INET6**.
General (or MVPN RPT) Suppression Active	Indicates whether suppression is configured.
General (or MVPN RPT) Entries Used	Number of currently used multicast forwarding cache entries.
General (or MVPN RPT) Suppress Threshold	Maximum number of multicast forwarding cache entries that can be added to the cache. When the number of entries reaches the configured threshold, the device suspends adding new multicast forwarding cache entries.
General (or MVPN RPT) Reuse Value	Number of multicast forwarding cache entries that must be reached before the device creates new multicast forwarding cache entries. When the total number of multicast forwarding cache entries is below the reuse value, the device resumes adding new multicast forwarding cache entries.

Sample Output

show multicast forwarding cache statistics instance

```
user@host> show multicast forwarding-cache statistic instance mvpn1 intet6
  Instance: mvpn1 Family: INET6
General Suppression Active                     Yes
General Entries Used                             0
General Suppress Threshold                     200
General Reuse Value                            200
MVPN RPT Suppression Active                    Yes
MVPN RPT Entries Used                            0
```

```
MVPN RPT Suppress Threshold                    200
MVPN RPT Reuse Value                           200
```

show multicast forwarding cache statistics instance (Forwarding-cache suppression is disabled)

```
user@host> show multicast forwarding-cache statistic instance mvpn1
Instance: mvpn1 Family: ALL
Forwarding-cache suppression disabled Not enabled by configuration
```

Release Information

Command introduced in Junos OS Release 12.2.

Starting in Junos OS Release 16.1, output includes general and rendezvous-point tree (RPT) suppression states.

RELATED DOCUMENTATION

clear multicast forwarding-cache | 3723

threshold (Multicast Forwarding Cache) | 29329

show multicast interface

IN THIS SECTION

- Syntax | 22287
- Syntax (EX Series Switch and the QFX Series) | 22287
- Description | 22287
- Options | 22287
- Required Privilege Level | 22287
- Output Fields | 22287
- Sample Output | 22289

Release Information | 22289

Syntax

```
show multicast interface
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show multicast interface
```

Description

Display bandwidth information about IP multicast interfaces.

Options

none	Display all interfaces that have multicast configured.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 840 on page 22288 describes the output fields for the `show multicast interface` command. Output fields are listed in the approximate order in which they appear.

Table 840: show multicast interface Output Fields

Field Name	Field Description
Interface	Name of the multicast interface.
Maximum bandwidth (bps)	Maximum bandwidth setting, in bits per second, for this interface.
Remaining bandwidth (bps)	Amount of bandwidth, in bits per second, remaining on the interface.
Mapped bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface. **NOTE**: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. This field does not appear in the output when the no QoS adjustment feature is disabled.
Local bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface. **NOTE**: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. This field does not appear in the output when the no QoS adjustment feature is disabled.
Reverse OIF mapping	State of the reverse OIF mapping feature (**on** or **off**). **NOTE**: This field does not appear in the output when the no QoS adjustment feature is disabled.
Reverse OIF mapping no QoS adjustment	State of the no QoS adjustment feature (**on** or **off**) for interfaces that are using reverse OIF mapping. **NOTE**: This field does not appear in the output when the no QoS adjustment feature is disabled.

Table 840: show multicast interface Output Fields *(Continued)*

Field Name	Field Description
Leave timer	Amount of time a mapped interface remains active after the last mapping ends. **NOTE**: This field does not appear in the output when the no QoS adjustment feature is disabled.
No QoS adjustment	State (**on**) of the no QoS adjustment feature when this feature is enabled. **NOTE**: This field does not appear in the output when the no QoS adjustment feature is disabled.

Sample Output

show multicast interface

```
user@host> show multicast interface
Interface                Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3                                 10000000                        0
fe-0/0/3.210                             10000000                  -2000000
fe-0/0/3.220                            100000000                 100000000
fe-0/0/3.230                             20000000                  18000000
fe-0/0/2.200                            100000000                 100000000
```

Release Information

Command introduced in Junos OS Release 8.3.

show multicast mrinfo

IN THIS SECTION

- Syntax | 22290
- Description | 22290
- Options | 22290
- Required Privilege Level | 22290
- Output Fields | 22291
- Sample Output | 22292
- Release Information | 22292

Syntax

```
show multicast mrinfo
<host>
```

Description

Display configuration information about IP multicast networks, including neighboring multicast router addresses.

Options

none Display configuration information about all multicast networks.

host (Optional) Display configuration information about a particular host. Replace *host* with a hostname or IP address.

Required Privilege Level

view

Output Fields

Table 841 on page 22291 describes the output fields for the `show multicast mrinfo` command. Output fields are listed in the approximate order in which they appear.

Table 841: show multicast mrinfo Output Fields

Field Name	Field Description
source-address	Query address, hostname (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor.
ip-address-1--->ip-address-2	Queried router interface address and directly attached neighbor interface address, respectively.
(*name* or *ip-address*)	Name or IP address of neighbor.

[*metric/ threshold/ type/ flags*]	Neighbor's multicast profile:
	• *metric*—Always has a value of **1**, because **mrinfo** queries the directly connected interfaces of a device.
	• *threshold*—Multicast threshold time-to-live (TTL). The range of values is **0** through **255**.
	• *type*—Multicast connection type: **pim** or **tunnel**.
	• *flags*—Flags for this route:
	  • **querier** Queried router is the designated router for the neighboring session.
	  • **leaf** Link is a leaf in the multicast network.
	  • **down** Link status indicator.

Sample Output

show multicast mrinfo

```
user@host> show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
  192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
  0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```

Release Information

Command introduced before Junos OS Release 7.4.

show multicast next-hops

IN THIS SECTION

- Syntax | 22292
- Syntax (EX Series Switch and the QFX Series) | 22293
- Description | 22293
- Options | 22293
- Required Privilege Level | 22294
- Output Fields | 22294
- Sample Output | 22294
- Release Information | 22297

Syntax

```
show multicast next-hops
<brief | detail | terse>
```

```
<identifier-number>
<inet | inet6>
<logical-system (all | logical-system-name)>
<show multicast next-hops satellitesatellite>
<session>
```

Syntax (EX Series Switch and the QFX Series)

```
show multicast next-hops
<brief | detail>
<identifier-number>
<inet | inet6>
```

Description

Display the entries in the IP multicast next-hop table.

Options

none	Display standard information about all entries in the multicast next-hop table for all supported address families.
brief | detail | terse	(Optional) Display the specified level of output. Use `terse` to display the total number of outgoing interfaces (as opposed to listing them) When you include the **detail** option on M Series and T Series routers and EX Series switches, the downstream interface name includes the next-hop ID number in parentheses, in the form **fe-0/1/2.0-(1048574),** where **1048574** is the next-hop ID number. Starting in Junos OS release 16.1, the `show multicast next-hops` statement shows the hierarchical next hops contained in the top-level next hop.
identifier-number	(Optional) Show a particular next hop by ID number. The range of values is **1** through **65,535**.
inet | inet6	(Optional) Display entries for IPv4 or IPv6 family addresses, respectively.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
satellite	Show satellite multicast unicast next-hops

session Show all MCNH session IDs

Required Privilege Level

view

Output Fields

Table 842 on page 22294 describes the output fields for the `show multicast next-hops` command. Output fields are listed in the approximate order in which they appear.

Table 842: show multicast next-hops Output Fields

Field Name	Field Description
Family	Protocol family (such as **INET**).
ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine.
Refcount	Number of cache entries that are using this next hop.
KRefcount	Kernel reference count for the next hop.
Downstream interface	Interface names associated with each multicast next-hop ID.
Incoming interface list	List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.

Sample Output

show multicast next-hops

```
user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefcount Downstream interface
```

```
262142          4          2 so-1/0/0.0
262143          2          1 mt-1/1/0.49152
262148          2          1 mt-1/1/0.32769
```

show multicast next-hops (Ingress Router, Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show multicast next-hops
Family: INET
ID          Refcount KRefcount Downstream interface Addr
1048580            2         1 1048576
(0x600dc04)        1         0 1048584
(0x600ea04)        1         0 (0x600e924)
1048583            2         1 1048579
(0x600e144)        1         0 1048587
(0x600e844)        1         0 (0x600e764)
1048582            2         1 1048578
(0x600df84)        1         0 1048586
(0x600e684)        1         0 (0x600e5a4)
1048581            2         1 1048577
(0x600ddc4)        1         0 1048585
(0x600ebc4)        1         0 (0x600eae4)
```

show multicast next-hops (Egress Router, Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show multicast next-hops
Family: INET
ID          Refcount KRefcount Downstream interface Addr
(0x600e844)        8         0 1048575
1048575           16         0 distributed-gmp
```

show multicast next-hops (Bidirectional PIM)

```
user@host> show multicast next-hops
Family: INET
ID          Refcount KRefcount Downstream interface
2097151            8         4 ge-0/0/1.0
```

```
Family: INET6
ID          Refcount KRefcount Downstream interface
2097157            2         1 ge-0/0/1.0

Family: Incoming interface list
ID          Refcount KRefcount Downstream interface
513                5         2 lo0.0
                               ge-0/0/1.0
514                5         2 lo0.0
                               ge-0/0/1.0
                               xe-4/1/0.0
515                3         1 lo0.0
                               ge-0/0/1.0
                               xe-4/1/0.0
544                1         0 lo0.0
                               xe-4/1/0.0
```

show multicast next-hops brief

The output for the `show multicast next-hops brief` command is identical to that for the `show multicast next-hops` command. For sample output, see "show multicast next-hops" on page 22294.

show multicast next-hops detail

```
user@host> show multicast next-hops detail
Family: INET
ID          Refcount KRefcount Downstream interface Addr
1048584            2         1 1048581
                               1048580
                               Flags 0x208 type 0x18 members 0/0/2/0/0
                               Address 0xb1841c4
1048591            3         2 787
                               747
                               Flags 0x206 type 0x18 members 0/0/2/0/0
                               Address 0xb1847f4
1048580            4         1 ge-1/1/9.0-(1048579)
                               Flags 0x200 type 0x18 members 0/0/0/1/0
                               Address 0xb184134
1048581            2         0 736
                               765
```

```
                                Flags 0x3 type 0x18 members 0/0/2/0/0
                                Address 0xb183dd4
1048585           18          0 787
                                747
                                Flags 0x203 type 0x18 members 0/0/2/0/0
                                Address 0xb184404

Family: INET6
ID          Refcount KRefcount Downstream interface Addr
1048586            4          2 1048585
                                1048583
                                Flags 0x20c type 0x19 members 0/0/2/0/0
                                Address 0xb1842e4
1048583           14          4 ge-1/1/9.0-(1048582)
                                Flags 0x200 type 0x19 members 0/0/0/1/0
                                Address 0xb183ef4
1048592            4          2 1048583
                                1048591
                                Flags 0x20c type 0x19 members 0/0/2/0/0
                                Address 0xb184644
```

show multicast next-hops detail (PIM using point-to-multipoint mode)

```
user@host> show multicast next-hops detail
Family: INET
ID          Refcount KRefcount Downstream interface
262142             2          1 st0.0-192.0.2.0(573)
                                st0.0-198.51.100.0(572)
```

Release Information

Command introduced before Junos OS Release 7.4.

inet6 option introduced in Junos OS Release 10.0 for EX Series switches.

detail option display of next-hop ID number introduced in Junos OS Release 11.1 for M Series and T Series routers and EX Series switches.

Support for bidirectional PIM added in Junos OS Release 12.1.

`terse` option introduced in Junos OS Release 16.1 for the MX Series.

show multicast next-hops satellite

IN THIS SECTION

- Syntax | 22298
- Description | 22298
- Options | 22298
- Required Privilege Level | 22298
- Output Fields | 22299
- Sample Output | 22300
- Release Information | 22301

Syntax

```
show multicast next-hops satellite
```

Description

Display satellite multicast next-hop table information.

The output lists next-hops for all interfaces in the VPLS address family used in multicast replication.

Options

brief | detail (Optional) Output level is the same when either option or no option is specified.

Required Privilege Level

view

Output Fields

Table 843 on page 22299 lists the output fields for the `show multicast next-hops satellite` command. Output fields are listed in the approximate order in which they appear, although the display order varies between the different levels of output.

Table 843: show multicast next-hops satellite Command Output Fields

Field Name	Field Description	Level of Output
`Next-hop ID`	Next-hop ID in the database.	All
`Interface Index`	Internal ID for a satellite device virtual interface.	All
`Interface Name`	Interface name for configured interfaces in the Junos Fusion. A satellite device virtual interface named `sd-`*`fpc-id`*`/0/0` is used to send traffic to extended ports on satellite device *fpc-id*.	All
`Label`	A multicast ECID assigned to the satellite device interface next hop. If no ECIDs are assigned for the interface, the value displayed in this field is 0. Satellite device interfaces might have multiple ECIDs listed that represent different groups of multicast destination extended ports on that satellite device.	All
`Flags`	Flags giving additional information about a next-hop entry. Values include: • SAT—The entry is a satellite composite next hop for extended port destinations that has been resolved and updated with the corresponding satellite device interface, `sd-`*`fpc-id`*`/0/0`, and the associated ECID shown in the `Label` field.	All

Sample Output

show multicast next-hops satellite

```
user@host> show multicast next-hops satellite
Next-hop    Interface   Interface
ID          index       Name                  Label  Flags
186         1           unknown               0
530         323         unknown               0
574         331         lc-0/0/0.32769        0
578         335         ge-0/0/4.32770        0
583         341         ge-0/0/7.32770        0
584         343         ge-0/0/8.32770        0
585         345         ge-0/0/9.32770        0
586         337         ge-0/0/5.32770        0
587         339         ge-0/0/6.32770        0
610         347         xe-100/0/6.0          0
611         349         xe-100/0/7.0          0
612         348         sd-100/0/0.32770      4129   SAT
618         350         xe-100/0/0.0          0
619         351         xe-100/0/1.0          0
620         352         xe-100/0/2.0          0
621         353         xe-100/0/3.0          0
622         354         xe-100/0/4.0          0
…
1061        348         sd-100/0/0.32770      4097   SAT
1062        346         sd-103/0/0.32770      4103   SAT
1068        346         sd-103/0/0.32770      4104   SAT
1069        346         sd-103/0/0.32770      4105   SAT
1070        346         sd-103/0/0.32770      4106   SAT
1071        346         sd-103/0/0.32770      4107   SAT
1074        346         sd-103/0/0.32770      4110   SAT
1075        346         sd-103/0/0.32770      4111   SAT
1076        346         sd-103/0/0.32770      4112   SAT
1077        346         sd-103/0/0.32770      4113   SAT
1078        346         sd-103/0/0.32770      4114   SAT
1079        346         sd-103/0/0.32770      4115   SAT
1080        346         sd-103/0/0.32770      4116   SAT
1081        346         sd-103/0/0.32770      4117   SAT
1082        346         sd-103/0/0.32770      4118   SAT
1085        346         sd-103/0/0.32770      4121   SAT
1097        346         sd-103/0/0.32770      4109   SAT
```

```
1189        346         sd-103/0/0.32770       4120   SAT
1190        348         sd-100/0/0.32770       4122   SAT
1363        433         ge-0/0/1.0             0
1364        434         ge-0/0/2.0             0
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Understanding Multicast Replication in a Junos Fusion*
>
> show multicast snooping next-hops satellite | 22338

show multicast pim-to-igmp-proxy

IN THIS SECTION

- Syntax | 22301
- Syntax (EX Series Switch and the QFX Series) | 22302
- Description | 22302
- Options | 22302
- Required Privilege Level | 22302
- Output Fields | 22302
- Sample Output | 22303
- Release Information | 22303

Syntax

```
show multicast pim-to-igmp-proxy
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show multicast pim-to-igmp-proxy
<instance instance-name>
```

Description

Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.

Options

none	Display configuration information about PIM-to-IGMP message translation for all routing instances.
instance ***instance-name***	(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 844 on page 22302 describes the output fields for the `show multicast pim-to-igmp-proxy` command. Output fields are listed in the order in which they appear.

Table 844: show multicast pim-to-igmp-proxy Output Fields

Field Name	Field Description
Instance	Routing instance. Default instance is **master** (inet.0 routing table).
Proxy state	State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: **enabled** or **disabled**.

Table 844: show multicast pim-to-igmp-proxy Output Fields *(Continued)*

Field Name	Field Description
interface-name	Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured.

Sample Output

show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

show multicast pim-to-igmp-proxy instance

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```

Release Information

Command introduced in Junos OS Release 9.6.

instance option introduced in Junos OS Release 10.3.

instance option introduced in Junos OS Release 10.3 for EX Series switches.

RELATED DOCUMENTATION

Configuring PIM-to-IGMP and PIM-to-MLD Message Translation

show multicast pim-to-mld-proxy

IN THIS SECTION

- Syntax | 22304
- Syntax (EX Series Switch and the QFX Series) | 22304
- Description | 22304
- Options | 22305
- Required Privilege Level | 22305
- Output Fields | 22305
- Sample Output | 22305
- Release Information | 22306

Syntax

```
show multicast pim-to-mld-proxy
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show multicast pim-to-mld-proxy
<instance instance-name>
```

Description

Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.

Options

none	Display configuration information about PIM-to-MLD message translation for all routing instances.
instance *instance-name*	(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 845 on page 22305 describes the output fields for the `show multicast pim-to-mld-proxy` command. Output fields are listed in the order in which they appear.

Table 845: show multicast pim-to-mld-proxy Output Fields

Field Name	Field Description
Proxy state	State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: **enabled** or **disabled**.
interface-name	Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured.

Sample Output

show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

show multicast pim-to-mld-proxy instance

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```

Release Information

Command introduced in Junos OS Release 9.6.

instance option introduced in Junos OS Release 10.3.

instance option introduced in Junos OS Release 10.3 for EX Series switches.

show multicast route

IN THIS SECTION

- Syntax | 22306
- Syntax (EX Series Switch and the QFX Series) | 22307
- Description | 22307
- Options | 22307
- Required Privilege Level | 22308
- Output Fields | 22308
- Sample Output | 22311
- Release Information | 22322

Syntax

```
show multicast route
<brief | detail | extensive | summary>
<active | all | inactive>
<display-origin-protocol>
<display-tunnel-name>
```

```
<group group>
<inet | inet6>
<instance (all | instance name)>
<logical-system (all |  logical-system-name)>
<oif-count>
<regular-expression>
<source-prefix source-prefix>
```

Syntax (EX Series Switch and the QFX Series)

```
show multicast route
<brief | detail | extensive | summary>
<active | all | inactive>
<group group>
<inet | inet6>
<instance (all | instance name)>
<regular-expression>
<source-prefix source-prefix>
```

Description

Display the entries in the IP multicast forwarding table. You can display similar information with the `show route table inet.1` command.

> **NOTE**: On all SRX Series Firewalls, when a multicast route is not available, pending sessions are not torn down, and subsequent packets are queued. If no multicast route resolve comes back, then the traffic flow has to wait for the pending session to timed out. Then packets can trigger new pending session create and route resolve.

Options

none	Display standard information about all entries in the multicast forwarding table for all routing instances.
brief | detail | extensive | summary	(Optional) Display the specified level of output.
active | all | inactive	(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.

display-origin-protocol	(Optional) Display originating protocol instead of upstream protocol.
display-tunnel-name	(Optional) Display provider tunnel name instead of branch information.
group *group*	(Optional) Display the cache entries for a particular group.
inet | inet6	(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.
instance (all | *instance-name***)**	(Optional) Display entries in the multicast forwarding table for all multicast-enabled routing instances or a specific multicast instance.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system.
oif-count	(Optional) Display a count of outgoing interfaces rather than listing them.
regular-expression	(Optional) Display information about the multicast forwarding table entries that match a UNIX OS-style regular expression.
source-prefix *source-prefix*	(Optional) Display the cache entries for a particular source prefix.

Required Privilege Level

view

Output Fields

Table 846 on page 22308 describes the output fields for the `show multicast route` command. Output fields are listed in the approximate order in which they appear.

Table 846: show multicast route Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	`summary extensive`
family	IPv4 address family (`INET`) or IPv6 address family (`INET6`).	All levels

Table 846: show multicast route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Group	Group address. For any-source multicast routes, for example for bidirectional PIM, the group address includes the prefix length.	All levels
Source	Prefix and length of the source as it is in the multicast forwarding table.	All levels
Incoming interface list	List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	All levels
Upstream interface	Name of the interface on which the packet with this source prefix is expected to arrive.	All levels
Upstream rpf interface list	When multicast-only fast reroute (MoFRR) is enabled, a PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request.	All levels
Downstream interface list	List of interface names to which the packet with this source prefix is forwarded. **distributed-gmp** Added in Junos OS Release 17.4R1 to indicate that line cards with distributed IGMP interfaces are receiving multicast traffic for a given (s,g).	All levels
Number of outgoing interfaces	Total number of outgoing interfaces for each (S,G) entry.	`extensive`
Session description	Name of the multicast session.	`detail extensive`

Table 846: show multicast route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Statistics	Rate at which packets are being forwarded for this source and group entry (in KBps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobytes per second packet forwarding statistic queries fails or times out, the statistics field displays `Forwarding statistics are not available`. **NOTE**: On QFX Series switches and OCX Series switches, this field does not report valid statistics. EX4300 Series switches don't report valid multicast statistics.	`detail extensive`
Next-hop ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the `show multicast nexthops` command.	`detail extensive`
Incoming interface list ID	For bidirectional PIM, incoming interface list identifier. Identifiers for interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	`detail extensive`
Upstream protocol	The protocol that maintains the active multicast forwarding route for this group or source. When the `show multicast route extensive` command is used with the `display-origin-protocol` option, the field name is only `Protocol` and not `Upstream Protocol`. However, this field also displays the protocol that installed the active route.	`detail extensive`
Route type	Type of multicast route. Values can be (S,G) or (*,G).	`summary`
Route state	Whether the group is `Active` or `Inactive`.	`summary extensive`
Route count	Number of multicast routes.	`summary`
Forwarding state	Whether the prefix is pruned or forwarding.	`extensive`

Table 846: show multicast route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Cache lifetime/ timeout	Number of seconds until the prefix is removed from the multicast forwarding table. A value of `never` indicates a permanent forwarding entry. A value of `forever` indicates routes that do not have keepalive times.	`extensive`
Wrong incoming interface notifications	Number of times that the upstream interface was not available.	`extensive`
Uptime	Time since the creation of a multicast route.	`extensive`
Sensor ID	Sensor ID corresponding to multicast route.	`extensive`

Sample Output

Starting in Junos OS Release16.1, `show multicast route` displays the top-level hierarchical next hop.

show multicast route

```
user@host> show multicast route
Family: INET

Group: 233.252.0.0
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0

Group: 233.252.0.1
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0

Group: 233.252.0.1
```

```
    Source: 10.255.70.15/32
    Upstream interface: so-1/0/0.0
    Downstream interface list:
        mt-1/1/0.1081344

Family: INET6
```

show multicast route (Bidirectional PIM)

```
user@host> show multicast route
Family: INET

Group: 233.252.0.1/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0

Group: 233.252.0.3/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0

Group: 233.252.0.11/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0

Group: 233.252.0.13/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0
Family: INET6
```

show multicast route brief

The output for the `show multicast route brief` command is identical to that for the `show multicast route` command. For sample output, see "show multicast route" on page 22311 or "show multicast route (Bidirectional PIM)" on page 22312.

show multicast route summary

```
user@host>show multicast route summary
Instance: master Family: INET

Route type   Route state    Route count
(S,G)             Active             2
(S,G)             Inactive           3

Instance: master Family: INET6
```

show multicast route detail

```
user@host> show multicast route detail
Family: INET

Group: 233.252.0.0
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Unknown
    Statistics: 8 kBps, 100 pps, 45272 packets
    Next-hop ID: 262142
    Upstream protocol: PIM

Group: 233.252.0.1
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Administratively Scoped
    Statistics: 0 kBps, 0 pps, 13404 packets
    Next-hop ID: 262142
```

```
    Upstream protocol: PIM

Group: 233.252.0.1
    Source: 10.255.70.15/32
    Upstream interface: so-1/0/0.0
    Downstream interface list:
        mt-1/1/0.1081344
    Session description: Administratively Scoped
    Statistics: 46 kBps, 1000 pps, 921077 packets

    Next-hop ID: 262143
    Upstream protocol: PIM

Family: INET6
```

show multicast route extensive (BGP Multicast)

```
user@host> show multicast route extensive
Family: INET

Group: 225.1.1.1/32
    Source: 10.0.0.1
    Upstream interface: ge-2/1/0.0
    Downstream interface list:
        ge-2/1/1.1 ge-2/1/1.0
    Number of outgoing interfaces: 2
    Session description: Unknown
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 123456
    Upstream protocol: BGPMCAST
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
```

show multicast route extensive (sender based RPF MoFRR)

```
user@host> show multicast route extensive
Instance: BLACK Family: INET6
```

```
Group: 233.252.0.1
    Source: 101.1.1.2/32
    Upstream rpf interface list:
        vt-16/0/0.0 (P)
            Session Id: 0x16b Session Status: Up
            Min-rate: 4000 kbps Weight: 1
            Sender Id: Label 57
        vt-16/0/0.0 (B)
            Session Id: 0x167 Session Status: Up & Forwarding
            Min-rate: 4000 kbps Weight: 65533
            Sender Id: Label 53
    Downstream interface list:
        et-19/4/1.0
    Number of outgoing interfaces: 1
    Session description: NOB Cross media facilities
    Statistics: 1842 kBps, 16741 pps, 1250285 packets
    RPF Next-hop ID: 1039
    Next-hop ID: 1048588
    Upstream protocol: MVPN
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:01:15
```

NOTE: `Session Status: Up & Forwarding` indicates that the specific session is the current forwarding path.

show multicast route extensive (Bidirectional PIM)

```
user@host> show multicast route extensive
Family: INET

Group: 233.252.0.1/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0
    Number of outgoing interfaces: 1
```

```
    Session description: NOB Cross media facilities
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097153
    Incoming interface list ID: 585
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Group: 233.252.0.3/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0
    Number of outgoing interfaces: 1
    Session description: NOB Cross media facilities
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097153
    Incoming interface list ID: 589
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Family: INET6
```

show multicast route extensive (PIM using point-to-multipoint mode)

```
user@host> show multicast route extensive
Instance: master Family: INET

Group: 225.0.0.1
    Source: 192.0.2.0/24
    Upstream interface: st0.1
+   Upstream neighbor: 203.0.113.0/24
    Downstream interface list:
+        st0.0-198.51.100.0 st0.0-198.51.100.1
    Session description: Unknown
```

```
    Statistics: 0 kBps, 1 pps, 119 packets
    Next-hop ID: 262142
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 0
    Uptime: 00:02:00
```

show multicast route extensive (traffic counters)

```
user@host> show multicast route extensive
Instance: master Family: INET

Group: 225.0.0.1
    Source: 192.0.2.0/24
    Upstream interface: ge-3/0/12.0
    Downstream interface list:
        ge-0/0/18.0 ge-0/0/7.0 ge-2/0/11.0 ge-2/0/7.0 ge-3/0/20.0 ge-3/0/21.0
    Number of outgoing interfaces: 6
    Session description: Unknown
    Statistics: 102 kBps, 801 pps, 5735 packets
    Next-hop ID: 131076
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 0
    Uptime: 00:03:57
```

show multicast route instance <instance-name> extensive

```
user@host> show multicast route instance mvpn extensive
Family: INET
roup: 233.252.0.10
    Source: 10.0.0.2/32
    Upstream interface: xe-0/0/0.102
    Downstream interface list:
        xe-10/3/0.0 xe-0/3/0.0 xe-0/0/0.106 xe-0/0/0.105
        xe-0/0/0.103 xe-0/0/0.104 xe-0/0/0.107 xe-0/0/0.108
```

```
    Session description: Administratively Scoped
    Statistics: 256 kBps, 3998 pps, 670150 packets
    Next-hop ID: 1048579
    Upstream protocol: MVPN
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 58
    Uptime: 00:00:04

Instance: master Family: INET

Group: 225.0.0.1
    Source: 101.0.0.2/32
    Upstream interface: ge-2/2/0.101
    Downstream interface list:
        distributed-gmp
    Number of outgoing interfaces: 1
    Session description: Unknown
    Statistics: 105 kBps, 2500 pps, 4153361 packets
    Next-hop ID: 1048575
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 0
    Uptime: 00:31:46

Group: 225.0.0.1
    Source: 101.0.0.3/32
    Upstream interface: ge-2/2/0.101
    Downstream interface list:
        distributed-gmp
    Number of outgoing interfaces: 1
    Session description: Unknown
    Statistics: 105 kBps, 2500 pps, 4153289 packets
    Next-hop ID: 1048575
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
```

```
    Wrong incoming interface notifications: 0
    Uptime: 00:31:46
```

show multicast route extensive (PIM NSR support for VXLAN on primary Routing Engine)

```
user@host> show multicast route extensive
Instance: master Family: INET

Group: 233.252.0.1
    Source: 10.3.3.3/32
    Upstream interface: ge-3/1/2.0
    Downstream interface list:
        -(593)
    Number of outgoing interfaces: 1
    Session description: Organisational Local Scope
    Statistics: 0 kBps, 0 pps, 27 packets
    Next-hop ID: 1048576
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding (Forwarding state is set as 'Forwarding' in master RE.)
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:06:38

Group: 233.252.0.1
    Source: 10.2.1.4/32
    Upstream interface: local
    Downstream interface list:
        ge-3/1/2.0
    Number of outgoing interfaces: 1
    Session description: Organisational Local Scope
    Statistics: 0 kBps, 0 pps, 86 packets
    Next-hop ID: 1048575
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding (Forwarding state is set as 'Forwarding' in master RE.)
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:07:45
```

```
Instance: master Family: INET6
```

show multicast route extensive (PIM NSR support for VXLAN on backup Routing Engine)

```
user@host> show multicast route extensive
Instance: master Family: INET

Group: 233.252.0.1
    Source: 10.3.3.3/32
    Upstream interface: ge-3/1/2.0
    Number of outgoing interfaces: 0
    Session description: Organisational Local Scope
    Forwarding statistics are not available
    Next-hop ID: 0
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Pruned (Forwarding state is set as 'Pruned' in backup RE.)
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:06:46

Group: 233.252.0.1
    Source: 10.2.1.4/32
    Upstream interface: local
    Number of outgoing interfaces: 0
    Session description: Organisational Local Scope
    Forwarding statistics are not available
    Next-hop ID: 0
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Pruned  (Forwarding state is set as 'Pruned' in backup RE.)
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:07:54

Instance: master Family: INET6
```

show multicast route extensive (PIM NSR support for VXLAN on backup Routing Engine)

```
user@host> show multicast route extensive
Instance: master Family: INET

Group: 233.252.0.1
    Source: 10.3.3.3/32
    Upstream interface: ge-3/1/2.0
    Downstream interface list:
        -(593)
    Number of outgoing interfaces: 1
    Session description: Organisational Local Scope
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 1048576
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding (Forwarding state is set as 'Forwarding' in backup RE.)
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:06:38

Group: 233.252.0.1
    Source: 10.2.1.4/32
    Upstream interface: local
    Downstream interface list:
        ge-3/1/2.0
    Number of outgoing interfaces: 1
    Session description: Organisational Local Scope
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 1048575
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding (Forwarding state is set as 'Forwarding' in backup RE.)
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:07:45

Instance: master Family: INET6
```

show multicast route extensive (Junos OS Evolved)

```
user@host> show multicast route extensive
Instance: master Family: INET

Group: 232.255.255.100
    Source: 10.1.1.2/32
    Upstream interface: et-0/0/0:0.0
    Downstream interface list:
        et-0/0/2:1.0 et-0/0/1:0.0
    Number of outgoing interfaces: 2
    Session description: Source specific multicast
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 11066
    Upstream protocol: Multicast
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 14:58:34
    Sensor ID: 0xf0000002
```

Release Information

Command introduced before Junos OS Release 7.4.

inet6 and `instance` options introduced in Junos OS Release 10.0 for EX Series switches.

Support for bidirectional PIM added in Junos OS Release 12.1.

`oif-count` option introduced in Junos OS Release 16.1 for the MX Series.

Support for PIM NSR support for VXLAN added in Junos OS Release 16.2.

Support for multicast traffic counters added in Junos OS 19.2R1 for EX4300 switches.

RELATED DOCUMENTATION

- Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

show multicast rpf

IN THIS SECTION

- Syntax | 22323
- Syntax (EX Series Switch and the QFX Series) | 22323
- Description | 22323
- Options | 22324
- Required Privilege Level | 22324
- Output Fields | 22324
- Sample Output | 22325
- Release Information | 22328

Syntax

```
show multicast rpf
<inet | inet6>
<instance instance-name>
<logical-system (all |  logical-system-name)>
<prefix>
<summary>
```

Syntax (EX Series Switch and the QFX Series)

```
show multicast rpf
<inet | inet6>
<instance instance-name>
<prefix>
<summary>
```

Description

Display information about multicast reverse-path-forwarding (RPF) calculations.

Options

none	Display RPF calculation information for all supported address families.
inet | inet6	(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.
instance *instance-name*	(Optional) Display information about multicast RPF calculations for a specific multicast instance.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system.
prefix	(Optional) Display the RPF calculation information for the specified prefix.
summary	(Optional) Display a summary of all multicast RPF information.

Required Privilege Level

view

Output Fields

Table 847 on page 22324 describes the output fields for the `show multicast rpf` command. Output fields are listed in the approximate order in which they appear.

Table 847: show multicast rpf Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Source prefix	Prefix and length of the source as it exists in the multicast forwarding table.
Protocol	How the route was learned.

Table 847: show multicast rpf Output Fields *(Continued)*

Field Name	Field Description
Interface	Upstream RPF interface. **NOTE**: The displayed interface information does not apply to bidirectional PIM RP addresses. This is because the `show multicast rpf` command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF interface information, always use the `show pim join extensive` command when bidirectional PIM is configured.
Neighbor	Upstream RPF neighbor. **NOTE**: The displayed neighbor information does not apply to bidirectional PIM. This is because the `show multicast rpf` command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF neighbor information, always use the `show pim join extensive` command when bidirectional PIM is configured.

Sample Output

show multicast rpf

```
 user@host>  show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
    Protocol: Static

10.255.14.132/32
    Protocol: Direct
    Interface: lo0.0

10.255.245.91/32
    Protocol: IS-IS
    Interface: so-1/1/1.0
    Neighbor: 192.168.195.21

172.16.0.1/32
```

```
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

192.168.0.0/16
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

192.168.14.0/24
Protocol: Direct
Interface: fxp0.0

192.168.14.132/32
Protocol: Local

192.168.195.20/30
Protocol: Direct
Interface: so-1/1/1.0

192.168.195.22/32
Protocol: Local

192.168.195.36/30
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 192.168.195.21
```

show multicast rpf inet6

```
user@host>  show multicast rpf inet6

Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128
    Protocol: Direct
    Interface: lo0.0

::10.255.245.91/128
```

```
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 2001:db8::2a0:a5ff:fe28:2e8c

::192.168.195.20/126
Protocol: Direct
Interface: so-1/1/1.0

::192.168.195.22/128
Protocol: Local

::192.168.195.36/126
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 2001:db8::2a0:a5ff:fe28:2e8c

::192.168.195.76/126
Protocol: Direct
Interface: fe-2/2/0.0

::192.168.195.77/128
Protocol: Local

2001:db8::/64
Protocol: Direct
Interface: so-1/1/1.0

2001:db8::290:69ff:fe0c:993a/128
Protocol: Local

2001:db8::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0

2001:db8::2/128
Protocol: PIM

2001:db8::d/128
Protocol: PIM
```

show multicast rpf prefix

```
user@host>  show multicast rpf 2001:db8::/16

Multicast RPF table: inet6.0, 13 entries

2001:db8::2/128
    Protocol: PIM

2001:db8::d/128
    Protocol: PIM

...
```

show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```

Release Information

Command introduced before Junos OS Release 7.4.

inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.

show multicast scope

IN THIS SECTION

- Syntax | 22329
- Syntax (EX Series Switch and the QFX Series) | 22329
- Description | 22329
- Options | 22329

- Required Privilege Level | 22330
- Output Fields | 22330
- Sample Output | 22330
- Release Information | 22331

Syntax

```
show multicast scope
<inet | inet6>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show multicast scope
<inet | inet6>
<instance instance-name>
```

Description

Display administratively scoped IP multicast information.

Options

none	Display standard information about administratively scoped multicast information for all supported address families in all routing instances.
inet | inet6	(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.
instance *instance-name*	(Optional) Display administratively scoped information for a specific multicast instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 848 on page 22330 describes the output fields for the `show multicast scope` command. Output fields are listed in the approximate order in which they appear.

Table 848: show multicast scope Output Fields

Field Name	Field Description
Scope name	Name of the multicast scope.
Group Prefix	Range of multicast groups that are scoped.
Interface	Interface that is the boundary of the administrative scope.
Resolve Rejects	Number of kernel resolve rejects.

Sample Output

show multicast scope

```
user@host> show multicast scope
                                                                Resolve
Scope name           Group Prefix       Interface               Rejects
233-net              233.252.0.0/16     fe-0/0/0.1                    0
local                233.252.0.1/16     fe-0/0/0.1                    0
local                2001:db8::/16          fe-0/0/0.1                    0
larry                2001:db8::1234/128     fe-0/0/0.1                    0
```

show multicast scope inet

```
user@host> show multicast scope inet
                                                                Resolve
```

```
Scope name           Group Prefix       Interface                 Rejects
233-net              233.252.0.0/16     fe-0/0/0.1                      0
local                233.252.0.0/16     fe-0/0/0.1                      0
```

show multicast scope inet6

```
user@host> show multicast scope inet6
                                                                 Resolve
Scope name           Group Prefix       Interface                 Rejects
local                2001:db8::/16          fe-0/0/0.1                     0
larry                2001:db8::1234/128     fe-0/0/0.1                     0
```

Release Information

Command introduced before Junos OS Release 7.4.

inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.

show multicast sessions

IN THIS SECTION

- Syntax | 22332
- Syntax (EX Series Switch and the QFX Series) | 22332
- Description | 22332
- Options | 22332
- Required Privilege Level | 22332
- Output Fields | 22333
- Sample Output | 22333
- Release Information | 22334

Syntax

```
show multicast sessions
<brief | detail | extensive>
<logical-system (all | logical-system-name)>
<regular-expression>
```

Syntax (EX Series Switch and the QFX Series)

```
show multicast sessions
<brief | detail | extensive>
<regular-expression>
```

Description

Display information about announced IP multicast sessions.

> **NOTE**: On all SRX Series Firewalls, only 100 packets can be queued during pending (S, G) route. However, when multiple multicast sessions enter the route resolve process at the same time, buffer resources are not sufficient to queue 100 packets for each session.

Options

none	Display standard information about all multicast sessions for all routing instances.
brief | detail | extensive	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
regular-expression	(Optional) Display information about announced sessions that match a UNIX-style regular expression.

Required Privilege Level

view

Output Fields

Table 849 on page 22333 describes the output fields for the `show multicast sessions` command. Output fields are listed in the approximate order in which they appear.

Table 849: show multicast sessions Output Fields

Field Name	Field Description
session-name	Name of the known announced multicast sessions.

Sample Output

show multicast sessions

```
user@host> show multicast sessions
1-Department of Biological Sciences, LSU
  ...
  Monterey Bay - DockCam
  Monterey Bay - JettyCam
  Monterey Bay - StandCam
  Monterey DockCam
  Monterey DockCam / ROV cam
  ...
  NASA TV (MPEG-1)
  ...
  UO Broadcast - NASA Videos - 25 Years of Progress
  UO Broadcast - NASA Videos - Journey through the Solar System
  UO Broadcast - NASA Videos - Life in the Universe
  UO Broadcast - NASA Videos - Nasa and the Airplane
  UO Broadcasts OPB's Oregon Story
  UO DOD News Clips
  UO Medical Management of Biological Casualties (1)
  UO Medical Management of Biological Casualties (2)
  UO Medical Management of Biological Casualties (3)
...
  376 active sessions.
```

show multicast sessions regular-expression detail

```
user@host> show multicast sessions "NASA TV" detail
SDP Version: 0  Originated by: -@10.223.83.33
 Session: NASA TV (MPEG-1)
 Description: NASA television in MPEG-1 format, provided by Private University. Please contact
the UO if you have problems with this feed.
 Email: Your Name Here <multicast@lists.private.edu>
 Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 233.252.0.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 10.223.91.191 live
Attribute: fmtp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106 Connection
Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
Attribute: rtpmap:103 L16/22050/2
Attribute: rtpmap:104 L16/22050

  1 matching sessions.
```

Release Information

Command introduced before Junos OS Release 7.4.

show multicast snooping next-hops

IN THIS SECTION

- Syntax | 22335
- Description | 22335
- Options | 22335
- Required Privilege Level | 22336
- Output Fields | 22336
- Sample Output | 22336
- Release Information | 22338

Syntax

```
show multicast snooping next-hops
<brief | detail>
<identifier next-hop-ID>
<inet>
<inet6>
<logical-system logical-system-name>
```

Description

Display information about the IP multicast snooping next-hops.

Options

brief | detail (Optional) Display the specified level of output.

inet (Optional) Display information for IPv4 multicast next hops only. If a family is not specified, both IPv4 and IPv6 results will be shown.

inet6 (Optional) Display information for IPv6 multicast next hops only. If a family is not specified, both IPv4 and IPv6 results will be shown.

logical-system *logical-system-name*	(Optional) Display information about a particular logical system, or type 'all'.

Required Privilege Level

view

Output Fields

Table 850 on page 22336 describes the output fields for the `show multicast snooping next-hops` command. Output fields are listed in the approximate order in which they appear.

Table 850: show multicast snooping next-hops Output Fields

Field Name	Field Description
Family	Protocol family for which multicast snooping next hops are displayed: **INET** or **INET6**.
Refcount	Number of cache entries that are using this next hop.
KRefcount	Kernel reference count for the next hop.
Downstream interface	Interface names associated with each multicast next-hop ID.
Nexthop Id	Identifier for the next-hop. **NOTE**: To see the next-hop ID for a given PE mesh group, `igmp-snooping` must be enabled for the relevant VPLS routing instance. (Junos OS creates a default CE and VE mesh groups for each VPLS routing instance. The next hop of the VE mesh group is the set of VE mesh-group interfaces of the remaining PEs in the same VPLS routing instance.)

Sample Output

show multicast snooping next-hops

```
user@host> show multicast snooping next-hops
Family: INET
```

```
ID          Refcount KRefcount Downstream interface Nexthop Id
1048574            4         1 ge-0/1/0.1000
                               ge-0/1/2.1000
                               ge-0/1/3.1000

1048574            4         1 ge-0/1/0.1000-(2000)
                               1048575
                               1048576

1048575            2         0 ge-0/1/2.1000-(2001)
                               ge-0/1/3.1000-(2002)

1048576            2         0 lsi.1048578-(2003)
                               lsi.1048579-(2004)
```

show multicast snooping next-hops (IGMP snooping enabled on a VPLS)

In thIS example, ID 1048585 is the VE next-hop ID created for the VE next hop that is holding VE interfaces for the routing instance. It only appears if igmp snooping is enabled on the VPLS.

```
user@host> show multicast snooping next-hops
Family: INET
ID          Refcount KRefcount Downstream interface Addr
1048588            2         1 1048585
1048589            2         1 1048585
                               ge-0/0/5.100
0                  2         0 ge-0/0/0.100
                               ge-0/0/1.100
1048583            2         1 local
1048587            2         1 local
                               1048585
1048586            4         2 local
                               1048585
                               ge-0/0/5.100
1048584            2         1 local
                               ge-0/0/5.100
1048582            6         2 ge-0/0/5.100
0                  2         0 ge-0/0/0.200
                               ge-0/0/2.200
0                  2         0 ge-0/0/0.300
                               ge-0/0/2.300
```

```
0                   1          0 vt-0/0/10.17825792
                                 vt-0/0/10.17825793
0                   1          0 vt-0/0/10.1048576
                                 vt-0/0/10.1048578
1048585             5          0 vt-0/0/10.1048577
                                 vt-0/0/10.1048579
0                   1          0 vt-0/0/10.34603008
                                 vt-0/0/10.34603009
```

Release Information

Command introduced in Junos OS Release 11.2.

show multicast snooping next-hops satellite

IN THIS SECTION

- Syntax | 22338
- Description | 22339
- Options | 22339
- Required Privilege Level | 22339
- Output Fields | 22339
- Sample Output | 22341
- Release Information | 22343

Syntax

```
show multicast snooping next-hops satellite
<brief | detail | extensive>
<nexthop-id nexthop-id>
```

Description

Display detailed multicast next-hop information for satellite device destinations.

You can use this command to:

- View the current list of all multicast traffic next hops.
- See details about a specified multicast next hop.
- Follow aggregation device composite next-hop processing as the aggregation device resolves and updates multicast next-hop table entries for extended port destinations.

The aggregation device allocates ECID tags that represent multicast or broadcast destinations behind satellite device extended ports, associates them with the corresponding satellite device virtual interfaces (`sd-`*`fpc-id`*`/0/0`), and updates multicast next-hop table entries accordingly. More detailed output from this command shows events that result in next-hop table updates.

Options

brief | detail | extensive (Optional) Display the specified level of output. The default output level is `brief`.

nexthop-id *nexthop-id* Display more detailed multicast satellite next-hop information only for the specified next-hop ID.

Required Privilege Level

view

Output Fields

Table 851 on page 22339 lists the output fields for the `show ethernet-switching flood next-hops satellite` command. Output fields are listed in the approximate order in which they appear.

Table 851: show multicast snooping next-hops satellite Command Output Fields

Field Name	Field Description	Level of Output
`Next-hop ID`	Multicast next-hop ID (original multicast next hop used with ingress multicast replication).	All

Table 851: show multicast snooping next-hops satellite Command Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Forwarding Next-Hop Type`	Type of next-hop entry. Values include: • COMPOSITE—Composite next hop.	All
`Table`	Name of routing or forwarding table for the listed entry.	All
`Flags`	Flags giving additional information about a next-hop entry. Values include: • SAT—The entry is a satellite destination composite next hop for an extended port destination that has been resolved and updated with the corresponding satellite device interface, `sd-`*`fpc`*`/0/0`, and the associated ECID shown in the `label` field. • ST—The entry is stale and waiting to be refreshed. • RU—The entry is stale, but is marked to be reused.	All
`Mrouter`	Next-hop list for multicast routers connected to the bridge domain or VLAN.	detail extensive
`aggregation-device`	Next-hop IDs and corresponding interfaces for the composite next hop to reach local multicast destination ports on the aggregation device.	detail extensive
`satellite-device-id` *`id`*	Satellite device ID with the next-hop IDs and corresponding interface names for the composite next-hop chain to reach extended ports that are multicast destinations on that satellite device. When an ECID has been assigned to destination extended ports on a satellite device, this field lists the satellite next-hop ID, the corresponding virtual satellite device interface (`sd-`*`fpc-id`*`/0/0`), and the allocated ECID (`label` field).	detail extensive

Table 851: show multicast snooping next-hops satellite Command Output Fields *(Continued)*

Field Name	Field Description	Level of Output
AE	Link aggregation group interface next-hop IDs and interface names for multicast destination composite next hops.	detail extensive
label	ECID associated with one or more multicast destination extended ports on the satellite device in the satellite-device-id field.	detail extensive
Linked Next-hop	Multicast satellite next-hop ID (next hop used with egress multicast replication).	detail extensive
When	Elapsed time since an event related to a multicast next-hop entry change.	extensive
Event	Brief description of the event related to a multicast next-hop entry change.	extensive
Action	Brief description of actions that resulted from the event.	extensive

Sample Output

show multicast snooping next-hops satellite

```
user@host> show multicast snooping next-hops satellite

Next-hop    Forwarding          Table                   Flags
ID          Next-Hop
            Type
2097265     COMPOSITE           inet.0
2097267     COMPOSITE           inet.0
2097270     COMPOSITE           inet.0
```

```
2097271     COMPOSITE           inet.0
2094267     COMPOSITE           inet.0                  SAT
```

show multicast snooping next-hops satellite nexthop-id (detail view for a specified next-hop ID)

```
user@host> show multicast snooping next-hops satellite nexthop-id 524296 detail
  Next-hop    Forwarding          Table                   Flags
  ID          Next-Hop
              Type
  524296      COMPOSITE           inet.0                  SAT
      satellite-device-id 106:
          ->1839 sd-106/0/0.32770 label=4104
      MRouter:
        524293
          aggregation-device:
              ->1708 et-0/0/30.0
              ->1826 xe-0/0/28:1.0
          satellite-device-id 100:
              ->1845 sd-100/0/0.32770 label=4101
          satellite-device-id 106:
              ->1847 sd-106/0/0.32770 label=4103
      Linked Next-hop: 524295
```

show multicast snooping next-hops satellite nexthop-id (extensive view for a specified next-hop ID)

```
user@host> show multicast snooping next-hops satellite nexthop-id 1048576 extensive
Next-hop    Forwarding          Table                   Flags
ID          Next-Hop
            Type
1048576     COMPOSITE           default
    satellite-device-id 100:
        ->54297 xe-100/0/12.0
        ->54299 xe-100/0/13.0
    satellite-device-id 101:
        ->54303 xe-101/0/12.0
        ->54339 xe-101/0/13.0
```

```
    Linked Next-hop: 1048578

When            Event                          Action
1w0d 21:07:51.281 Next-hop linking             Linked with satellite nhid:1048578
1w0d 21:07:51.260 Adding satellite Indirect    DB add
1w0d 21:07:51.260 Processing translate Q       Translating Multicast composite INH
                                               sd-100/0/0.32770 ECID:4105 [nhid=627]
                                               sd-101/0/0.32770 ECID:4104 [nhid=626]
1w0d 21:07:51.260 inh dependency               [ECID 4105] nhid:627 --> inhid:1048576
1w0d 21:07:51.260 inh dependency               [ECID 4104] nhid:626 --> inhid:1048576
1w0d 21:07:51.254 inh dependency               [ECID 4105] nhid:627 --> inhid:1048576
1w0d 21:07:51.254 Add to translate Q           inhid:1048576 added
1w0d 21:07:51.254 inh dependency               [ECID 4104] nhid:626 --> inhid:1048576
1w0d 21:07:51.241 Translate Q Skip             inhid:1048576: [ECID=0] Next-hop not ready
1w0d 21:07:51.241 Translate Q Skip             inhid:1048576: [ECID=0] Next-hop not ready
1w0d 21:07:51.241 Processing translate Q       Translating Multicast composite INH
1w0d 21:07:51.241 Add to translate Q           inhid:1048576 added
1w0d 21:07:51.241 Add to translate Q           Original INH decoded - Adding to Translate Q
1w0d 21:07:51.241 ECID reference               Added ECID reference-ID:2
1w0d 21:07:51.241 ECID reference               Added ECID reference-ID:1
1w0d 21:07:51.241 Next-hop add                 Linked to fwd CNH extension: 0xa67a200
                                               list[4]: {54297 54299 54303 54339 }
                                               nhid:1048576 from kernel
                                               fwd nhid[619]:
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding Multicast Replication in a Junos Fusion

Egress (Local) Replication on the Satellite Devices

show multicast snooping route satellite | 22350

show multicast snooping route

IN THIS SECTION

- Syntax | 22344
- Description | 22345
- Options | 22345
- Required Privilege Level | 22346
- Output Fields | 22346
- Sample Output | 22347
- Release Information | 22349

Syntax

```
show multicast snooping route
<regexp>
<active>
<all>
<bridge-domain bridge-domain-name>
<brief >
<control>
<data>
<detail >
<extensive>
<group group>
<inactive>
<inet>
<inet6>
<instance instance-name>
<logical-system logical-system-name>
<mesh-group mesh-group-name>
<qualified-vlan vlan-id>
<source-prefix source-prefix>
<vlan vlan-id>
```

Description

Display the entries in the IP multicast snooping forwarding table. You can display some of this information with the `show route table inet.1` command.

Options

none	Display standard information about all entries in the multicast snooping table for all virtual switches and all bridge domains.
active | all | inactive	(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast snooping table.
bridge-domain *bridge-domain*	(Optional) Display the entries for a particular bridge domain.
brief | detail | extensive	(Optional) Display the specified level of output.
control	(Optional) Display control route entries.
data	(Optional) Display data route entries.
group *group*	(Optional) Display the entries for a particular group.
inet	(Optional) Display IPv4 information.
inet6	(Optional) Display IPv6 information.
instance *instance-name*	(Optional) Display the entries for a multicast instance.
logical-system *logical-system-name*	(Optional) Display information about a particular logical system, or type 'all'.
mesh-group *mesh-group-name*	(Optional) Display the entries for a particular mesh group.
qualified-vlan *vlan-id*	(Optional) Display the entries for a particular qualified VLAN.
regexp	(Optional) Display information about the multicast forwarding table entries that match a UNIX-style regular expression.
source-prefix *source-prefix*	(Optional) Display the entries for a particular source prefix.
vlan *vlan-id*	(Optional) Display the entries for a particular VLAN.

Required Privilege Level

view

Output Fields

Table 852 on page 22346 describes the output fields for the `show multicast snooping route` command. Output fields are listed in the approximate order in which they appear.

Table 852: show multicast snooping route Output Fields

Field Name	Field Description	Level of Output
Nexthop Bulking	Displays whether next-hop bulk updating is **ON** or **OFF** (only for routing-instance type of **virtual switch** or **vpls**).	All levels
Family	IPv4 address family (**INET**) or IPv6 address family (**INET6**).	All levels
Group	Group address.	All levels
Source	Prefix and length of the source as it is in the multicast forwarding table. For (*,G) entries, this field is set to "*".	All levels
Routing-instance	Name of the routing instance to which this routing information applies. (Displayed when multicast is configured within a routing instance.)	All levels
Learning Domain	Name of the learning domain to which this routing information applies.	**detail extensive**

Table 852: show multicast snooping route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Statistics	Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. **NOTE**: EX4600, EX4650, and the QFX5000 line of switches don't provide packet rates for multicast transit traffic at Layer 2, and any values displayed in this field for `kBps` and `pps` are not valid. Up until and including the following Junos OS Release versions, the same is true of the packet count (`packets` value in this field is also not a valid count): Junos OS Releases 18.4R2-S2, 19.1R2-S1, 19.2R1, 19.3R2, and 19.4R1. Starting *after* those Junos OS releases, EX4600, EX4650, and QFX5000 switches count packets forwarded to this prefix and display valid statistics for the `packets` value only.	**detail extensive**
Next-hop ID	Next-hop identifier of the prefix. The identifier is returned by the router's Packet Forwarding Engine and is also displayed in the output of the `show multicast nexthops` command.	**detail extensive**
Route state	Whether the group is **Active** or **Inactive**.	**extensive**
Forwarding state	Whether the prefix is **Pruned** or **Forwarding**.	**extensive**
Cache lifetime/ timeout	Number of seconds until the prefix is removed from the multicast forwarding table. A value of **never** indicates a permanent forwarding entry.	**extensive**

Sample Output

show multicast snooping route bridge-domain

```
 user@host>  show multicast snooping route bridge-domain br-dom-1 extensive
Family: INET

Group: 232.1.1.1
    Source: 192.168.3.100/32
    Downstream interface list:
```

```
        ge-0/1/0.200
    Statistics: 0 kBps, 0 pps, 1 packets
    Next-hop ID: 1048577
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 240 seconds
```

show multicast snooping route instance vs

```
 user@host>  show multicast snooping route instance vs
Nexthop Bulking: ON

Family: INET

Group: 224.0.0.0
    Bridge-domain: vsid500

Group: 225.1.0.1
    Bridge-domain: vsid500
    Downstream interface list: vsid500
       ge-0/3/8.500 ge-1/1/9.500 ge1/2/5.500
```

show multicast snooping route extensive

```
 user@host>  show multicast snooping route extensive inet6 group ff03::1
Nexthop Bulking: OFF

Family: INET6
Group: ff03::1/128
    Source: ::
    Bridge-domain: BD-1
    Mesh-group: __all_ces__
    Downstream interface list:
        ae0.1 -(562) 1048576
    Statistics: 2697 kBps, 3875 pps, 758819039 packets
    Next-hop ID: 1048605
    Route state: Active
    Forwarding state: Forwarding

Group: ff03::1/128
```

```
    Source: 6666::2/128
    Bridge-domain: BD-1
    Mesh-group: __all_ces__
    Downstream interface list:
        ae0.1 -(562) 1048576
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 1048605
    Route state: Active
    Forwarding state: Forwarding
```

show multicast snooping route extensive group

```
 user@host>  show multicast snooping route extensive iinstance evpn-vxlan group 233.252.0.1/
Group: 233.252.0.1/32
    Source: *
    Vlan: VLAN-100
    Mesh-group: __all_ces__
        Downstream interface list:
        ge-0/0/3.0 -(662)
        evpn-core-nh -(131076)
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 131070
    Route state: Active
    Forwarding state: Forwarding
```

Release Information

Command introduced in Junos OS Release 8.5.

Support for `control`, `data`, `qualified-vlan` and `vlan` options introduced in Junos OS Release 13.3 for EX Series switches.

show multicast snooping route satellite

IN THIS SECTION

- Syntax | 22350
- Description | 22350
- Options | 22351
- Required Privilege Level | 22351
- Output Fields | 22351
- Sample Output | 22353
- Release Information | 22356

Syntax

```
show multicast snooping route satellite
<brief | detail | extensive>
<bridge-domain-name bridge-domain-name | <vlan-name vlan-name>
<group group-address>
<source source-address>
<vswitch-name virtual-switch-name>
```

Description

Display Layer 2 multicast routing information (with IGMP snooping) for destination ports on satellite devices.

This command lists multicast routing entries by route prefix and bridge domain name (Junos Fusion Provider Edge). Each route entry lists the next-hop ID (`NH Index`) used when the aggregation device performs ingress multicast replication. When local replication is enabled and the VLAN has multicast destination extended ports on a satellite device, the aggregation device:

- Creates a satellite device next-hop chain to those multicast destinations through their corresponding satellite device.
- Updates the multicast route entry with a link to the satellite next-hop chain (`Linked NH Index`).

If a multicast route does not have a satellite next-hop chain, `Linked NH Index` is `0`. When the `Linked NH Index` value is non-zero, `NH Index` refers to the *original* next-hop information, but the aggregation device uses the satellite next-hop chain (`Linked NH Index`) for routing multicast traffic to satellite device extended ports. Use the `detail` or `extensive` option to expand the output to include details about the original next-hop and satellite next-hop chains.

Options

brief | detail | extensive (Optional) Display the specified level of output. The default output level is `brief`. The `detail` output level expands the original and satellite device next-hop chains for each route displayed. The `extensive` output level includes details about multicast next-hop entry update events used mainly for troubleshooting.

bridge-domain-name *bridge-domain-name* (Junos Fusion Provider Edge only) Filter output to display bridge domain flooding information only for the specified bridge domain name.

group *group-address* Filter output to display flooding information only for the specified multicast group.

source *source-address* Filter output to display flooding information only for the specified multicast source.

vswitch-name *virtual-switch-name* Filter output to display flooding information only for the specified Layer 2 virtual switch.

Required Privilege Level

view

Output Fields

Table 853 on page 22351 lists the output fields for the `show multicast snooping route satellite` command. Output fields are listed in the approximate order in which they appear.

Table 853: show multicast snooping route satellite Command Output Fields

Field Name	Field Description	Level of Output
`Route`	Multicast route information listed by route prefix under this heading.	All

Table 853: show multicast snooping route satellite Command Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`(S, G/m)`	Multicast entry state.	All
`Bridge Domain`	Bridge domain or VLAN name.	All
`NH Index`	Multicast next-hop ID (original multicast next hop used with ingress multicast replication).	brief
`Linked NH Index`	Multicast satellite next-hop ID (next hop used with egress multicast replication).	brief
`Next-hop information`	Detailed list of next-hop chain information for the original and satellite next-hop chains, listed by original next-hop index (`NH Index`) and satellite next-hop chain index (`Linked NH Index`) output field values.	detail extensive
`Mrouter`	Next-hop chain list for multicast routers connected to the bridge domain or VLAN.	detail extensive
`aggregation-device`	Next-hop IDs and corresponding interfaces for the composite next hop to reach local multicast destination ports on the aggregation device.	detail extensive
`satellite-device-id` *`id`*	Satellite device ID with the next-hop IDs and corresponding interface names for the composite next-hop chain to reach extended ports that are multicast destinations on that satellite device. When an ECID has been assigned to destination extended ports on a satellite device, this field lists the satellite next-hop ID, the corresponding virtual satellite device interface (`sd-`*`fpc-id`*`/0/0`), and the allocated ECID (`label` field).	detail extensive

Table 853: show multicast snooping route satellite Command Output Fields *(Continued)*

Field Name	Field Description	Level of Output
AE	Link aggregation group interface next-hop IDs and interface names multicast destination composite next hops.	detail extensive
label	ECID associated with one or more multicast destination extended ports on the satellite device in the satellite-device-id field.	
When	Elapsed time since an event related to a multicast next-hop entry addition or update.	extensive
Event	Brief description of the event related to a multicast next-hop entry addition or update.	extensive
Action	Brief description of actions that resulted from the event.	extensive

Sample Output

show multicast snooping route satellite

```
user@host> show multicast snooping route satellite

-----VSwitch Instance: default-switch----
Route: 00.04.00.01.00.00.233.252.0.0.0.0.0.0.00.00/72
  (S, G/m): (*, 233.252.0.0.0/24)
    Bridge Domain:   bd100
    NH Index:        1048582
    Linked NH Index: 1048620
Route: 00.05.00.01.00.00.233.252.0.0.0.0.0.0.00.00/72
  (S, G/m): (*, 233.252.0.0.0/24)
    Bridge Domain:   bd1
    NH Index:        1048582
    Linked NH Index: 1048620
Route: 00.06.00.01.00.00.233.252.0.0.0.0.0.0.00.00/72
```

```
  (S, G/m): (*, 233.252.0.0.0/24)
    Bridge Domain:   bd10
    NH Index:        1048582
    Linked NH Index: 1048620
```

show multicast snooping route satellite detail

```
user@host> show multicast snooping route satellite detail

-----VSwitch Instance: default-switch----
  Route: 00.03.00.01.00.00.233.252.0.0.0.0.0.0.00.00/52
      (S, G/m): (*, 233.252.0.0/4)
       Bridge Domain: VLAN800
        Next-hop information:
          524287
            MRouter:
              524286
                aggregation-device:
                    ->1708 et-0/0/30.0
                    ->1826 xe-0/0/28:1.0
                satellite-device-id 100:
                    ->1838 xe-100/0/49:3.0
                satellite-device-id 106:
                    ->1834 xe-106/0/11.0
          524294
            MRouter:
              524293
                aggregation-device:
                    ->1708 et-0/0/30.0
                    ->1826 xe-0/0/28:1.0
                satellite-device-id 100:
                    ->1845 sd-100/0/0.32770 label=4101
                satellite-device-id 106:
                    ->1847 sd-106/0/0.32770 label=4103
  Route: 00.03.00.01.00.00.233.252.0.1.0.0.0.0.00.00/80
      (S, G/m): (*, 233.252.0.1/32)
       Bridge Domain: VLAN800
        Next-hop information:
          524295
            satellite-device-id 106:
                ->1804 xe-106/0/13.0
```

```
            MRouter:
              524286
                aggregation-device:
                    ->1708 et-0/0/30.0
                    ->1826 xe-0/0/28:1.0
                satellite-device-id 100:
                    ->1838 xe-100/0/49:3.0
                satellite-device-id 106:
                    ->1834 xe-106/0/11.0
          524296
            satellite-device-id 106:
                ->1839 sd-106/0/0.32770 label=4104
            MRouter:
              524293
                aggregation-device:
                    ->1708 et-0/0/30.0
                    ->1826 xe-0/0/28:1.0
                satellite-device-id 100:
                    ->1845 sd-100/0/0.32770 label=4101
                satellite-device-id 106:
                    ->1847 sd-106/0/0.32770 label=4103
```

show multicast snooping route satellite group

```
user@host> show multicast snooping route satellite group 233.252.0.1
-----VSwitch Instance: default-switch----
Route: 00.04.00.01.00.00.233.252.0.1.0.0.0.0.00.00/80
  (S, G/m): (*, 233.252.0.1/32)
    Bridge Domain:   bd100
    NH Index:        1048585
    Linked NH Index: 1048621
```

show multicast snooping route satellite bridge-domain-name (for specific bridge domain with detail view on Junos Fusion Provider Edge)

```
user@host> show multicast snooping route satellite bridge-domain-name bd100 group 233.252.0.1
detail
```

```
-----VSwitch Instance: default-switch----
Route: 00.04.00.01.00.00.233.252.0.1.0.0.0.0.00.00/80
    (S, G/m): (*, 233.252.0.1/32)
     Bridge Domain: bd100
      Next-hop information:
        1048576
          satellite-device-id 100:
              ->54297 xe-100/0/12.0
              ->54299 xe-100/0/13.0
          satellite-device-id 101:
              ->54303 xe-101/0/12.0
              ->54339 xe-101/0/13.0
        1048578
          satellite-device-id 100:
              ->627 sd-100/0/0.32770 label=4105
          satellite-device-id 101:
              ->626 sd-101/0/0.32770 label=4104
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding Multicast Replication in a Junos Fusion

Egress (Local) Replication on the Satellite Devices

show multicast snooping next-hops satellite | 22338

show multicast statistics

IN THIS SECTION

- Syntax | 22357
- Description | 22357

- Options | 22357
- Additional Information | 22358
- Required Privilege Level | 22358
- Output Fields | 22358
- Sample Output | 22360
- Release Information | 22363

Syntax

```
show multicast statistics
<inet | inet6>
<instance (all | instance-name)>
<interface interface-name>
<logical-system (all | logical-system-name)>
<satelliteshow multicast statistics satellite >
```

Description

Display IP multicast statistics.

Options

none	Display multicast statistics for all supported address families for all routing instances.
inet | inet6	(Optional) Display multicast statistics for IPv4 or IPv6 family addresses, respectively.
instance (all | *instance-name*)	(Optional) Display statistics for all multicast-enabled routing instances or a specific multicast routing instance.
interface *interface-name*	(Optional) Display statistics for a specific interface.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

satellite (Optional) Show satellite multicast statistics.

Additional Information

The input and output interface multicast statistics are consistent, but not timely. They are constructed from the forwarding statistics, which are gathered at 30-second intervals. Therefore, the output from this command always lags the true count by up to 30 seconds.

Required Privilege Level

view

Output Fields

Table 854 on page 22358 describes the output fields for the `show multicast statistics` command. Output fields are listed in the approximate order in which they appear.

Table 854: show multicast statistics Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Family	Protocol family for which multicast statistics are displayed: **INET** or **INET6**.
Interface	Name of the interface for which statistics are being reported.
Routing Protocol	Primary multicast protocol on the interface: **PIM**, **DVMRP for INET**, or **PIM for INET6**.
Mismatch	Number of multicast packets that did not arrive on the correct upstream interface.
Kernel Resolve	Number of resolve requests processed by the primary multicast protocol on the interface.
Resolve No Route	Number of resolve requests that were ignored because there was no route to the source.
Resolve Filtered	Number of resolve requests filtered by policy if any policy is configured.

Table 854: show multicast statistics Output Fields *(Continued)*

Field Name	Field Description
In Kbytes	Total accumulated incoming packets (in KB) since the last time the `clear multicast statistics` command was issued.
Out Kbytes	Total accumulated outgoing packets (in KB) since the last time the `clear multicast statistics` command was issued.
Mismatch error	Number of mismatches that were ignored because of internal errors.
Mismatch No Route	Number of mismatches that were ignored because there was no route to the source.
Routing Notify	Number of times that the multicast routing system has been notified of a new multicast source by a multicast routing protocol .
Resolve Error	Number of resolve requests that were ignored because of internal errors.
In Packets	Total number of incoming packets since the last time the `clear multicast statistics` command was issued.
Out Packets	Total number of outgoing packets since the last time the `clear multicast statistics` command was issued.
Resolve requests on interfaces not enabled for multicast *n*	Number of resolve requests on interfaces that are not enabled for multicast that have accumulated since the `clear multicast statistics` command was last issued.
Resolve requests with no route to source *n*	Number of resolve requests with no route to the source that have accumulated since the `clear multicast statistics` command was last issued.

Table 854: show multicast statistics Output Fields *(Continued)*

Field Name	Field Description
Routing notifications on interfaces not enabled for multicast *n*	Number of routing notifications on interfaces not enabled for multicast that have accumulated since the `clear multicast statistics` command was last issued.
Routing notifications with no route to source *n*	Number of routing notifications with no route to the source that have accumulated since the `clear multicast statistics` command was last issued.
Interface Mismatches on interfaces not enabled for multicast *n*	Number of interface mismatches on interfaces not enabled for multicast that have accumulated since the `clear multicast statistics` command was last issued.
Group Membership on interfaces not enabled for multicast *n*	Number of group memberships on interfaces not enabled for multicast that have accumulated since the `clear multicast statistics` command was last issued.

Sample Output

show multicast statistics

```
user@host> show multicast statistics
    Address family: INET
    Interface: fe-0/0/0
        Routing Protocol:          PIM   Mismatch error:                0
        Mismatch:                    0   Mismatch No Route:             0
        Kernel Resolve:             10   Routing Notify:                0
        Resolve No Route:            0   Resolve Error:                 0
        In Kbytes:                4641   In Packets:                50454
        Out Kbytes:                  0   Out Packets:                   0
    Interface: so-0/1/1.0
        Routing Protocol:          PIM   Mismatch error:                0
        Mismatch:                    0   Mismatch No Route:             0
```

```
        Kernel Resolve:              0   Routing Notify:                0
        Resolve No Route:            0   Resolve Error:                 0
        In Kbytes:                   0   In Packets:                    0
        Out Kbytes:               4641   Out Packets:               50454

    Resolve requests on interfaces not enabled for multicast 0
    Resolve requests with no route to source 0
    Routing notifications on interfaces not enabled for multicast 0
    Routing notifications with no route to source 0
    Interface Mismatches on interfaces not enabled for multicast 0
    Group Membership on interfaces not enabled for multicast 25

    Address family: INET6
    Interface: fe-0/0/0.0
        Routing Protocol:          PIM   Mismatch error:                0
        Mismatch:                    0   Mismatch No Route:             0
        Kernel Resolve:              0   Routing Notify:                0
        Resolve No Route:            0   Resolve Error:                 0
        In Kbytes:                   0   In Packets:                    0
        Out Kbytes:                  0   Out Packets:                   0
    Interface: so-0/1/1.0
        Routing Protocol:          PIM   Mismatch error:                0
        Mismatch:                    0   Mismatch No Route:             0
        Kernel Resolve:              0   Routing Notify:                0
        Resolve No Route:            0   Resolve Error:                 0
        In Kbytes:                   0   In Packets:                    0
        Out Kbytes:                  0   Out Packets:                   0

    Resolve requests on interfaces not enabled for multicast 0
    Resolve requests with no route to source 0
    Routing notifications on interfaces not enabled for multicast 0
    Routing notifications with no route to source 0
    Interface Mismatches on interfaces not enabled for multicast 0
    Group Membership on interfaces not enabled for multicast 0
```

show multicast statistics (PIM using point-to-multipoint mode)

```
user@host> show multicast statistics
Interface: st0.0-192.0.2.0
    Routing protocol:        PIM   Mismatch error:            0
    Mismatch:                  0   Mismatch no route:         0
```

```
    Kernel resolve:            0   Routing notify:             0
    Resolve no route:          0   Resolve error:              0
    Resolve filtered:          0   Notify filtered:            0
    In kbytes:                 0   In packets:                 0
    Out kbytes:                0   Out packets:                0

Interface: st0.0-192.0.2.0
    Routing protocol:        PIM   Mismatch error:             0
    Mismatch:                  0   Mismatch no route:          0
    Kernel resolve:            0   Routing notify:             0
    Resolve no route:          0   Resolve error:              0
    Resolve filtered:          0   Notify filtered:            0
    In kbytes:                 0   In packets:                 0
    Out kbytes:                0   Out packets:                0

Interface: st0.1-198.51.100.0
    Routing protocol:        PIM   Mismatch error:             0
    Mismatch:                  0   Mismatch no route:          0
    Kernel resolve:            0   Routing notify:             0
    Resolve no route:          0   Resolve error:              0
    Resolve filtered:          0   Notify filtered:            0
    In kbytes:                 0   In packets:                 0
    Out kbytes:                0   Out packets:                0
```

show multicast statistics interface

```
user@host> show multicast statistics interface vt-3/0/10.2097152
Instance: master Family: INET
Interface: vt-3/0/10.2097152
    Routing protocol:          PIM   Mismatch error:               0
    Mismatch:                    0   Mismatch no route:            0
    Kernel resolve:              0   Routing notify:               0
    Resolve no route:            0   Resolve error:                0
    Resolve filtered:            0   Notify filtered:              0
    In kbytes:                   0   In packets:                   0
    Out kbytes:                  0   Out packets:                  0
```

Release Information

Command introduced before Junos OS Release 7.4.

`interface` option introduced in Junos OS Release 16.1 for the MX Series.

RELATED DOCUMENTATION

clear multicast statistics | 3730

show multicast statistics satellite

IN THIS SECTION

- Syntax | 22363
- Description | 22363
- Options | 22364
- Required Privilege Level | 22364
- Sample Output | 22364
- Release Information | 22364

Syntax

```
show multicast statistics satellite
<brief | detail >
```

Description

Display statistics about multicast satellite routing tables and ECID management.

Options

brief | detail (Optional) Display the specified level of output. The default output level is `brief`.

Required Privilege Level

view

Sample Output

show multicast statistics satellite

```
user@host> show multicast statistics satellite
Multicast Statistics:
    Number of flood route entries:            8000
    Number of satellite flood route entries:  8000
    Number of MCINET route entries:           44000
    Number of satellite MCINET route entries: 36000
    Unicast VPLS next-hops(non-satellite):    32
    Number of satellite ECID next-hops:       12
    Number of VPLS composite next-hops:       12000
    Number of satellite composite next-hops:  12000
    Number of Indirect next-hops:             28002
    Number of Satellite Indirect next-hops:   28001
    Number of ECIDs requested:                14
    Number of ECID responses received:        14
    Number of ECID delete messages:           2
    Number of ECID mapping entries in DB:     12
    Number of ECID mapping entries ready:     12
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding Multicast Replication in a Junos Fusion

show multicast summary satellite

IN THIS SECTION

- Syntax | 22365
- Description | 22365
- Options | 22365
- Required Privilege Level | 22365
- Sample Output | 22366
- Release Information | 22366

Syntax

```
show multicast summary satellite
```

Description

Display summary status of multicast replication features in a Junos Fusion.

This command displays whether or not egress multicast replication (also called local replication) is enabled. When local replication is configured, this command displays `Egress replication: Enabled`, and `Egress replication: Disabled` otherwise.

This command also displays the graceful restart state of the satellite management control plane processes for local replication when these processes are first activated or have been restarted. The `Restart phase` output field value indicates the phase where the restart process stalled or failed, or displays a `Complete` message if the restart process completed successfully.

Options

This command has no options.

Required Privilege Level

view

Sample Output

show multicast summary satellite

```
user@host>show multicast summary satellite
Multicast:
    Restart phase:      Complete (11/11)
    Egress replication: Enabled
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Understanding Multicast Replication in a Junos Fusion

Configuring Egress (Local) Replication on a Junos Fusion

show multicast statistics satellite | 22363

show multicast usage

IN THIS SECTION

- Syntax | 22367
- Syntax (EX Series Switch and the QFX Series) | 22367
- Description | 22367
- Options | 22367
- Required Privilege Level | 22367
- Output Fields | 22368
- Sample Output | 22368
- Release Information | 22370

Syntax

```
show multicast usage
<brief | detail>
<inet | inet6>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show multicast usage
<brief | detail>
<inet | inet6>
<instance instance-name>
```

Description

Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups.

Options

none	Display multicast usage information for all supported address families for all routing instances.
brief | detail	(Optional) Display the specified level of output.
inet | inet6	(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.
instance *instance-name*	(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 855 on page 22368 describes the output fields for the `show multicast usage` command. Output fields are listed in the approximate order in which they appear.

Table 855: show multicast usage Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Group	Group address.
Sources	Number of sources.
Packets	Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays **unavailable**.
Bytes	Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays **unavailable**.
Prefix	IP address.
/len	Prefix length.
Groups	Number of multicast groups.
Sensor ID	Sensor ID corresponding to multicast route.

Sample Output

show multicast usage

```
user@host> show multicast usage
Group           Sources Packets             Bytes
```

```
233.252.0.0       1       52847                4439148
233.252.0.1       2       13450                1125530


Prefix          /len Groups Packets              Bytes
10.255.14.144   /32  2      66254                5561304
10.255.70.15    /32  1      43                   3374...
```

show multicast usage brief

The output for the `show multicast usage brief` command is identical to that for the `show multicast usage` command. For sample output, see "show multicast usage" on page 22368.

show multicast usage instance

```
user@host> show multicast usage instance VPN-A
Group           Sources Packets              Bytes
233.252.0.254   1       5538                 509496
233.252.0.39    1       13                   624
233.252.0.40    1       13                   624

Prefix          /len Groups Packets              Bytes
192.168.195.34  /32  1      5538                 509496
10.255.14.30    /32  1      13                   624
10.255.245.91   /32  1      13                   624
...
```

show multicast usage detail

```
user@host> show multicast usage detail
Group           Sources Packets              Bytes
233.252.0.0       1       53159                4465356
    Source: 10.255.14.144   /32 Packets: 53159 Bytes: 4465356
233.252.0.1       2       13450                1125530
    Source: 10.255.14.144   /32 Packets: 13407 Bytes: 1122156
    Source: 10.255.70.15    /32 Packets: 43 Bytes: 3374


Prefix          /len Groups Packets              Bytes
```

```
10.255.14.144   /32  2      66566                5587512
    Group: 233.252.0.0       Packets: 53159 Bytes: 4465356
    Group: 233.252.0.1       Packets: 13407 Bytes: 1122156
10.255.70.15    /32  1      43                   3374
    Group: 233.252.0.1       Packets: 43 Bytes: 3374
```

Release Information

Command introduced before Junos OS Release 7.4.

inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.

show mvpn c-multicast

IN THIS SECTION

- Syntax | 22370
- Description | 22371
- Options | 22371
- Required Privilege Level | 22371
- Output Fields | 22371
- Sample Output | 22372
- Release Information | 22374

Syntax

```
show mvpn c-multicast
<extensive | summary>
<instance-name instance-name>
<source-pe>
```

Description

Display the multicast VPN customer multicast route information.

Options

extensive | summary	(Optional) Display the specified level of output.
instance-name *instance-name*	(Optional) Display output for the specified routing instance.
source-pe	(Optional) Display source-pe output for the specified c-multicast entries.

Required Privilege Level

view

Output Fields

Table 856 on page 22371 lists the output fields for the `show mvpn c-multicast` command. Output fields are listed in the approximate order in which they appear.

Table 856: show mvpn c-multicast Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the VPN routing instance.	summary extensive none
C-mcast IPv4 (S:G)	Customer router IPv4 multicast address.	extensive none
Ptnl	Provider tunnel attributes, *tunnel type:tunnel source, tunnel destination group*.	extensive none

Table 856: show mvpn c-multicast Output Fields *(Continued)*

Field Name	Field Description	Level of Output
St	State: • DS—Represents (S,G) and is created due to (*,G) • RM—Remote VPN route learned from the remote PE router • St display blank—SSM group join	extensive none
MVPN instance	Name of the multicast VPN routing instance	extensive none
C-multicast IPv4 route count	Number of customer multicast IPv4 routes associated with the multicast VPN routing instance.	summary
C-multicast IPv6 route count	Number of customer multicast IPv6 routes associated with the multicast VPN routing instance.	summary

Sample Output

show mvpn c-multicast

```
user@host> show mvpn c-multicast
Family : INET

Instance : VPN-A
  C-mcast IPv4 (S:G)            Provider Tunnel                          Label-In       St
FwdNh    Segment
  192.168.195.78/32:233.252.0.1/24 PIM-SM:10.255.14.144, 192.168.0.1                   RM  M-
mt-2/0/10.32771

Instance : VPN-B
  C-mcast IPv4 (S:G)            Provider Tunnel                          Label-In       St
FwdNh    Segment
  192.168.195.94/32:233.252.0.2/24 PIM-SM:10.255.14.144, 192.168.0.2                   RM  M-
mt-2/0/10.32770

Instance : VPN-C
```

```
  MVPN Mode : RPT-SPT

  C-mcast IPv4 (S:G)            Provider Tunnel                                     Label-In
St    FwdNh    Segment
  192.168.1.1/0:233.252.0.3/32        INGRESS-REPLICATION:MPLS Label 16:192.168.2.1
16        M-1048665
```

show mvpn c-multicast summary

```
user@host> show mvpn c-multicast summary
MVPN Summary:
Family: INET
Family: INET6

Instance: mvpn1
  C-multicast IPv6 route count: 1
```

show mvpn c-multicast extensive

```
user@host> show mvpn c-multicast extensive
Family : INET

Instance : VPN-A
  MVPN Mode : SPT-ONLY
  MVPN Forwarding Entry uses hierarchical nexthop : No
  MVPN Forwarding Entry share tunnel nexthop : Yes

  C-mcast IPv4 (S:G)            Provider Tunnel                          Label-In       St
FwdNh    Segment
  192.168.195.78/32:233.252.0.1/24 PIM-SM:10.255.14.144, 192.168.0.1                  RM  M-
mt-2/0/10.32771

Instance : VPN-B
  MVPN Mode : SPT-ONLY
  MVPN Forwarding Entry uses hierarchical nexthop : No
  MVPN Forwarding Entry share tunnel nexthop : Yes

  C-mcast IPv4 (S:G)            Provider Tunnel                          Label-In       St
FwdNh    Segment
  192.168.195.94/32:233.252.0.2/24 PIM-SM:10.255.14.144, 192.168.0.2                  RM  M-
```

```
mt-2/0/10.32770

Instance : VPN-C
  MVPN Mode : RPT-SPT
  MVPN Forwarding Entry uses hierarchical nexthop : No
  MVPN Forwarding Entry share tunnel nexthop : Yes

  C-mcast IPv4 (S:G)            Provider Tunnel                                     Label-In
St    FwdNh    Segment
  192.168.1.1/0:233.252.0.3/32  INGRESS-REPLICATION:MPLS Label 16:192.168.2.1
16         M-1048665
```

show mvpn c-multicast source-pe

```
user@host> show mvpn c-multicast  source-pe
Family : INET
Family : INET6

Instance : mvpn1
  MVPN Mode : RPT-SPT
    C-Multicast route address: ::/0:ff05::1/128
        MVPN Source-PE1:
                extended-community: no-advertise target:10.1.0.0:9
                Route Distinguisher: 10.1.0.0:1
                Autonomous system number: 1
                Interface: ge-0/0/9.1 Index: 343
        PIM Source-PE1:
                extended-community: target:10.1.0.0:9
                Route Distinguisher: 10.1.0.0:1
                Autonomous system number: 1
                Interface: ge-0/0/9.1 Index: 343
```

Release Information

Command introduced in Junos OS Release 8.4.

Option to show `source-pe` introduced in Junos OS Release 15.1.

show mvpn instance

IN THIS SECTION

- Syntax | 22375
- Description | 22375
- Options | 22375
- Required Privilege Level | 22376
- Output Fields | 22376
- Sample Output | 22377
- Sample Output | 22378
- Sample Output | 22379
- Release Information | 22381

Syntax

```
show mvpn instance
<instance-name>
<display-tunnel-name>
<extensive | summary>
<inet | inet6>
<logical-system>
```

Description

Display the multicast VPN routing instance information according to the options specified.

Options

instance-name	(Optional) Display statistics for the specified routing instance, or press Enter without specifying an instance name to show output for all instances.
display-tunnel-name	(Optional) Display the ingress provider tunnel name rather than the attribute.

extensive | summary (Optional) Display the specified level of output.

inet | inet6 (Optional) Display output for the specified IP type.

inet | inet6 (Optional) Display output for the specified IP type.

logical-system (Optional) Display details for the specified logical system, or type “all”.

Required Privilege Level

view

Output Fields

Table 857 on page 22376 lists the output fields for the `show mvpn instance` command. Output fields are listed in the approximate order in which they appear.

Table 857: show mvpn instance Output Fields

Field Name	Field Description	Level of Output
MVPN instance	Name of the multicast VPN routing instance	extensive none
Instance	Name of the VPN routing instance.	summary extensive none
Provider tunnel	Provider tunnel attributes, *tunnel type:tunnel source, tunnel destination group*.	extensive none
Neighbor	Address, type of provider tunnel (I-P-tnl, inclusive provider tunnel and S-P-tnl, selective provider tunnel) and provider tunnel for each neighbor.	extensive none
C-mcast IPv4 (S:G)	Customer IPv4 router multicast address.	extensive none
C-mcast IPv6 (S:G)	Customer IPv6 router multicast address.	extensive none

Table 857: show mvpn instance Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Ptnl	Provider tunnel attributes, *tunnel type:tunnel source, tunnel destination group*.	extensive none
St	State: • DS—Represents (S,G) and is created due to (*,G) • RM—Remote VPN route learned from the remote PE router • St display blank—SSM group join	extensive none
Neighbor count	Number of neighbors associated with the multicast VPN routing instance.	summary
C-multicast IPv4 route count	Number of customer multicast IPv4 routes associated with the multicast VPN routing instance.	summary
C-multicast IPv6 route count	Number of customer multicast IPv6 routes associated with the multicast VPN routing instance.	summary

Sample Output

show mvpn instance

```
user@host> show mvpn instance
Family : INET

Instance : VPN-A
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 192.168.0.1
  Neighbor                      Inclusive Provider Tunnel                          Label-In
St             Segment
  10.255.14.160                 PIM-SM:10.255.14.160, 192.168.0.1
  10.255.70.17                  PIM-SM:10.255.70.17, 192.168.0.1

  C-mcast IPv4 (S:G)            Provider Tunnel                                    Label-In
```

```
St    FwdNh    Segment
  192.168.195.78/32:233.252.0.1/24 PIM-SM:10.255.14.144, 192.168.0.1
RM  M-mt-2/0/10.32771

Instance : VPN-B
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 192.168.0.2
  Neighbor                      Inclusive Provider Tunnel                          Label-In
St             Segment
  10.255.14.160                 PIM-SM:10.255.14.160, 192.168.0.2
  10.255.70.17                  PIM-SM:10.255.70.17, 192.168.0.2

  C-mcast IPv4 (S:G)            Provider Tunnel                                    Label-In
St    FwdNh    Segment
  192.168.195.94/32:233.252.0.2/24 PIM-SM:10.255.14.144, 192.168.0.2
RM  M-mt-2/0/10.32770

Instance : VPN-C
  MVPN Mode : SPT-ONLY
  Provider tunnel: I-P-tnl:RSVP-TE P2MP:10.255.14.144, 29061,10.255.14.144
  Neighbor                      Inclusive Provider Tunnel                          Label-In
St             Segment
  10.255.14.160                 RSVP-TE P2MP:10.255.14.160, 41451,10.255.14.160           19
  10.255.70.17                  RSVP-TE P2MP:10.255.70.17, 5310,10.255.70.17              19

Instance : VPN-D
  MVPN Mode : SPT-ONLY
  Provider tunnel: I-P-tnl:LDP-P2MP:10.255.14.144, lsp-id 16777217
  Neighbor                      Inclusive Provider Tunnel                          Label-In
St             Segment
  10.255.14.160                 LDP-P2MP:10.255.14.160, lsp-id 16777217                   20
  10.255.70.17                  LDP-P2MP:10.255.70.17, lsp-id 16777217                    20
```

Sample Output

show mvpn instance summary

```
user@host> show mvpn instance summary
MVPN Summary:
Family: INET
Family: INET6
```

```
Instance: mvpn1
  Sender-Based RPF: Disabled. Reason: Not enabled by configuration.
  Hot Root Standby: Disabled. Reason: Not enabled by configuration.
  Neighbor count: 3
  C-multicast IPv6 route count: 1
```

Sample Output

show mvpn instance extensive

```
user@host> show mvpn instance extensive
Family : INET

Instance : VPN-A
  MVPN Mode : SPT-ONLY
  MVPN Forwarding Entry uses hierarchical nexthop : No
  MVPN Forwarding Entry share tunnel nexthop : Yes
  Sender-Based RPF: Disabled. Reason: Not enabled by configuration.
  Hot Root Standby: Disabled. Reason: Not enabled by configuration.
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 192.168.0.1
  Neighbor                      Inclusive Provider Tunnel                          Label-In
St             Segment
  10.255.14.160                 PIM-SM:10.255.14.160, 192.168.0.1
  10.255.70.17                  PIM-SM:10.255.70.17, 192.168.0.1

  C-mcast IPv4 (S:G)            Provider Tunnel                                    Label-In
St    FwdNh    Segment
  192.168.195.78/32:233.252.0.1/24 PIM-SM:10.255.14.144, 192.168.0.1
RM  M-mt-2/0/10.32771

Instance : VPN-B
  MVPN Mode : SPT-ONLY
  MVPN Forwarding Entry uses hierarchical nexthop : No
  MVPN Forwarding Entry share tunnel nexthop : Yes
  Sender-Based RPF: Disabled. Reason: Not enabled by configuration.
  Hot Root Standby: Disabled. Reason: Not enabled by configuration.
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 192.168.0.2
  Neighbor                      Inclusive Provider Tunnel                          Label-In
St             Segment
  10.255.14.160                 PIM-SM:10.255.14.160, 192.168.0.2
  10.255.70.17                  PIM-SM:10.255.70.17, 192.168.0.2
```

```
  C-mcast IPv4 (S:G)            Provider Tunnel                                     Label-In
St    FwdNh    Segment
  192.168.195.94/32:233.252.0.2/24 PIM-SM:10.255.14.144, 192.168.0.2
RM  M-mt-2/0/10.32770

Instance : VPN-C
  MVPN Mode : SPT-ONLY
  MVPN Forwarding Entry uses hierarchical nexthop : No
  MVPN Forwarding Entry share tunnel nexthop : Yes
  Sender-Based RPF: Disabled. Reason: Not enabled by configuration.
  Hot Root Standby: Disabled. Reason: Not enabled by configuration.
  Provider tunnel: I-P-tnl:RSVP-TE P2MP:10.255.14.144, 29061,10.255.14.144
  Neighbor                      Inclusive Provider Tunnel                           Label-In
St             Segment
  10.255.14.160                 RSVP-TE P2MP:10.255.14.160, 41451,10.255.14.160           19
  10.255.70.17                  RSVP-TE P2MP:10.255.70.17, 5310,10.255.70.17              19
  C-mcast IPv4 (S:G)            Provider Tunnel                                     Label-In
St    FwdNh    Segment
  192.168.195.66/0:233.252.0.3/32        Primary unbound

Instance : VPN-D
  MVPN Mode : SPT-ONLY
  MVPN Forwarding Entry uses hierarchical nexthop : No
  MVPN Forwarding Entry share tunnel nexthop : Yes
  Sender-Based RPF: Disabled. Reason: Not enabled by configuration.
  Hot Root Standby: Disabled. Reason: Not enabled by configuration.
  Provider tunnel: I-P-tnl:LDP-P2MP:10.255.14.144, lsp-id 16777217
  Neighbor                      Inclusive Provider Tunnel                           Label-In
St             Segment
  10.255.14.160                 LDP-P2MP:10.255.14.160, lsp-id 16777217                   20
  10.255.70.17                  LDP-P2MP:10.255.70.17, lsp-id 16777217                    20
  C-mcast IPv4 (S:G)            Provider Tunnel                                     Label-In
St    FwdNh    Segment
  192.168.195.67/0:233.252.0.4/32        Primary unbound
```

show mvpn instance summary (IPv6)

```
user@host> show mvpn instance summary
MVPN Summary:
Instance: VPN-A
```

```
  C-multicast IPv6 route count: 2
Instance: VPN-B
  C-multicast IPv6 route count: 2
```

Release Information

Command introduced in Junos OS Release 8.4.

Additional details in output for extensive option introduced in Junos OS Release 15.1.

show mvpn neighbor

IN THIS SECTION

- Syntax | 22382
- Description | 22382
- Options | 22382
- Required Privilege Level | 22382
- Output Fields | 22382
- Sample Output | 22383
- Sample Output | 22384
- Sample Output | 22385
- Sample Output | 22385
- Sample Output | 22386
- Sample Output | 22386
- Sample Output | 22387
- Sample Output | 22387
- Release Information | 22388

Syntax

```
show mvpn neighbor
<extensive | summary>
<inet | inet6>
<instance instance-name | neighbor-address address>
<logical-system logical-system-name>
```

Description

Display multicast VPN neighbor information.

Options

extensive | summary	(Optional) Display the specified level of output for all multicast VPN neighbors.
inet | inet6	(Optional) Display IPv4 or IPv6 information for all multicast VPN neighbors.
instance *instance-name* | neighbor-address *address*	(Optional) Display multicast VPN neighbor information for the specified instance or the specified neighbor.
logical-system *logical-system-name*	(Optional) Display multicast VPN neighbor information for the specified logical system.

Required Privilege Level

view

Output Fields

Table 858 on page 22383 lists the output fields for the `show mvpn neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 858: show mvpn neighbor Output Fields

Field Name	Field Description	Level of Output
MVPN instance	Name of the multicast VPN routing instance	**extensive** none
Instance	Name of the VPN routing instance.	**summary extensive** none
Neighbor	Address, type of provider tunnel (**I-P-tnl**, inclusive provider tunnel and **S-P-tnl**, selective provider tunnel) and provider tunnel for each neighbor.	**extensive** none
Provider tunnel	Provider tunnel attributes, ***tunnel type:tunnel source, tunnel destination group***.	**extensive** none

Sample Output

show mvpn neighbor

```
user@host> show mvpn neighbor
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.1
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
```

```
Instance: VPN-B
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.2
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.2
```

Sample Output

show mvpn neighbor extensive

```
user@host> show mvpn neighbor extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.1
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.2
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.2
```

show mvpn neighbor extensive

```
user@host> show mvpn neighbor extensive
MVPN instance:
```

```
Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: mvpn-a
  Neighbor                              I-P-tnl
  10.255.72.45
  10.255.72.50                          LDP P2MP:10.255.72.50, lsp-id 1
```

Sample Output

show mvpn neighbor instance-name

```
user@host> show mvpn neighbor instance-name VPN-A
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.1
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.1
```

Sample Output

show mvpn neighbor neighbor-address

```
user@host> show mvpn neighbor neighbor-address 10.255.14.160
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
```

```
Instance: VPN-A
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.2
```

Sample Output

show mvpn neighbor neighbor-address summary

```
user@host> show mvpn neighbor neighbor-address 10.255.70.17 summary
MVPN Summary:
Instance: VPN-A
Instance: VPN-B
```

Sample Output

show mvpn neighbor neighbor-address extensive

```
user@host> show mvpn neighbor neighbor-address 10.255.70.17 extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                              I-P-tnl
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.1
MVPN instance:
```

```
Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Neighbor                              I-P-tnl
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.2
```

Sample Output

show mvpn neighbor neighbor-address instance-name

```
user@host> show mvpn neighbor neighbor-address 10.255.70.17 instance-name VPN-A
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                              I-P-tnl
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.1
```

Sample Output

show mvpn neighbor summary

```
user@host> show mvpn neighbor summary
MVPN Summary:
Family: INET
Family: INET6

Instance: mvpn1
  Neighbor count: 3
```

Release Information

Command introduced in Junos OS Release 8.4.

show mvpn suppressed

IN THIS SECTION

- Syntax | 22388
- Description | 22388
- Options | 22389
- Required Privilege Level | 22389
- Output Fields | 22389
- Sample Output | 22389
- Sample Output | 22390
- Release Information | 22390

Syntax

```
show mvpn suppressed
<instance-name>
<general | mvpn-rpt>
<inet | inet6>
```

Description

MVPN maintains a list of suppressed customer-multicast states and the reason they were suppressed. Display it, for example, to help understand the enforcement of forwarding-cache limits

Options

instance-name	(Optional) Display statistics for the specified routing instance, or press Enter without specifying an instance name to show output for all instances.
general | mvpn-rpt	(Optional) Display suppressed multicast prefixes and reason they were suppressed.
inet | inet6	(Optional) Display output for the specified IP type.

Required Privilege Level

view

Output Fields

Table 859 on page 22389 lists the output fields for the `show mvpn suppressed` command. Output fields are listed in the approximate order in which they appear.

Table 859: show mvpn suppressed Output Fields

Field Name	Field Description
MVPN instance	Name of the multicast VPN routing instance.
Prefix	Shown as a single line per prefix, group followed by source.
reason	MVPN *,G entries are deleted either because they exceed either the general forwarding-cache limit or because they exceed the forwarding-cache limit set for MVPN RPT.

Sample Output

show mvpn suppressed

```
user@host> show mvpn suppressed instance name
Instance: mvpn1 Family: INET

  Prefix 0.0.0.0/0:239.1.1.1/32, Suppressed due to MVPN RPT forwarding-cache limit
```

```
  Instance: mvpn1 Family: INET6
  Prefix ::91.1.1.1/128:Ff05::1/128, Suppressed due to general forwarding-cache limit
  Prefix ::/0:ff05::2/128, Suppressed due to general forwarding-cache limit
  Prefix ::/0:ff05::3/128, Suppressed due to MVPN RPT forwarding-cache limit
```

Sample Output

show mvpn suppressed summary

```
user@host> show mvpn suppressed instance name summary
Instance: mvpn1 Family: INET

  General entries suppressed:    5
  MVPN RPT entries suppressed:   1

  Instance: mvpn1 Family: INET6
  General entries suppressed:    5
  MVPN RPT entries suppressed:   1
```

Release Information

Command introduced in Junos OS Release16.1.

show mvrp

IN THIS SECTION

- Syntax | 22391
- Description | 22391
- Required Privilege Level | 22391
- Output Fields | 22391
- Sample Output | 22392
- Sample Output | 22393

Sample Output | 22393
Release Information | 22394

Syntax

```
show mvrp
```

Description

Display Multiple VLAN Registration Protocol (MVRP) configuration information.

Required Privilege Level

view

Output Fields

Table 860 on page 22391 lists the output fields for the `show mvrp` command. Output fields are listed in the approximate order in which they appear.

Table 860: show mvrp Output Fields

Field Name	Field Description
`MVRP dynamic VLAN creation`	Displays whether global MVRP dynamic VLAN creation is `Enabled` or `Disabled`.
`Global MVRP configuration`	Displays global MVRP information: • **MVRP status**—Displays whether MVRP is **Enabled** or **Disabled**. • **MVRP dynamic vlan creation**—Displays whether global MVRP dynamic VLAN creation is **Dnabled** or **Disabled**.

Table 860: show mvrp Output Fields *(Continued)*

Field Name	Field Description
`MVRP BPDU MAC address`	Displays the multicast media access control (MAC) address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the customer MVRP multicast MAC address is used.
`MVRP timers (ms)`	Displays MVRP timer information: • `Interface`—The interface on which MVRP is configured. • `Join`—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements. • `Leave`—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message. • `LeaveAll`— The interval at which LeaveAll messages are sent on interfaces. LeaveAll messages maintain current MVRP VLAN membership information in the network.
Interface based configuration	Displays interface-specific MVRP information: • **Interface**—The interface on which MVRP is configured. • **Status**—Displays whether MVRP is **Enabled** or **Disabled**. • **Registration**—Displays whether registration for the interface is **Forbidden** or **Normal**. • **Dynamic VLAN Creation**—Displays whether interface dynamic VLAN creation is **Enabled** or **Disabled**.

Sample Output

show mvrp (EX Series switches and MX Series routers)

```
user@host> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface       Join   Leave  LeaveAll
```

```
  ge-11/2/8        200     800     10000
  ge-11/0/9        200     800     10000
  ge-11/3/0        200     800     10000
```

Sample Output

show mvrp (SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320)

```
user@host> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (00-00-5E-00-53-00)
MVRP timers (ms)
  Interface       Join   Leave  LeaveAll
  ge-0/0/1        200     800     60
```

Sample Output

show mvrp (EX Series switches)

```
user@switch> show mvrp

Global MVRP configuration
  MVRP status               : Enabled
  MVRP dynamic vlan creation: Enabled
  MVRP Timers (ms):
   Interface        Join   Leave   LeaveAll
   --------------   ----   -----   --------
   all              200    600     10000
   xe-0/1/1.0       200    600     10000

  Interface based configuration:
   Interface        Status     Registration   Dynamic VLAN Creation
   --------------   --------   ------------   ---------------------
   all                   Disabled   Normal             Enabled
   xe-0/1/1.0           Enabled    Normal         Enabled
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches

Verifying That MVRP Is Working Correctly on Switches

show mvrp statistics | **22406**

show mvrp applicant-state | **22394**

show mvrp dynamic-vlan-memberships | **22398**

show mvrp interface | **22401**

show mvrp registration-state | **22403**

show mvrp statistics | **22406**

show mvrp applicant-state | **22394**

show mvrp dynamic-vlan-memberships | **22398**

show mvrp interface | **22401**

show mvrp registration-state | **22403**

show mvrp applicant-state

IN THIS SECTION

- Syntax | **22395**
- Description | **22395**
- Required Privilege Level | **22395**
- Output Fields | **22395**
- Sample Output (EX Series and MX Series) | **22396**
- Sample Output (SRX1500, SRX300, SRX550M, SRX345, SRX340, and SRX320) | **22397**
- Release Information | **22397**

Syntax

```
show mvrp applicant-state
```

Description

For MX Series routers, EX Series switches, SRX1500, SRX300, SRX550M, SRX345, SRX340, and SRX320, display Multiple VLAN Registration Protocol (MVRP) applicant state information.

Required Privilege Level

view

Output Fields

Table 861 on page 22395 lists the output fields for the `show mvrp applicant-state` command. Output fields are listed in the approximate order in which they appear.

Table 861: show mvrp applicant-state Output Fields

Field Name	Field Description
`VLAN Id`	Displays the VLAN ID number.
`Interface`	Displays the interface number associated with the VLAN ID.

Table 861: show mvrp applicant-state Output Fields *(Continued)*

Field Name	Field Description
State	Displays one of the following MVRP registrar states: • VO— Very anxious observer. • VP —Very anxious passive. • VA —Very anxious new. • AN —Anxious new. • AA —Anxious active. • QA —Quiet active. • LA —Leaving active. • AO —Anxious observer. • QO —Quiet observer. • LO —Leaving observer. • AP —Anxious passive. • QA —Quiet passive.

Sample Output (EX Series and MX Series)

show mvrp applicant-state (EX Series and MX Series)

```
user@host> show mvrp applicant-state
MVRP applicant state for routing instance 'default-switch'
 (VO) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,
 (AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,
 (AO) Anxious observer, (QO) Quiet observer, (LO) Leaving observer,
 (AP) Anxious passive, (QP) Quiet passive

VLAN Id     Interface          State
    100     ge-11/3/0          Declaring (QA)
```

```
    200     ge-11/3/0          Declaring (QA)
    300     ge-11/3/0          Declaring (QA)
```

Sample Output (SRX1500, SRX300, SRX550M, SRX345, SRX340, and SRX320)

show mvrp applicant-state

```
user@host> show mvrp applicant-state
MVRP applicant state for routing instance 'default-switch'
 (VO) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,
 (AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,
 (AO) Anxious observer, (QO) Quiet observer, (LO) Leaving observer,
 (AP) Anxious passive, (QP) Quiet passive

VLAN Id     Interface          State
    1       ge-0/0/1          Idle (VO)
    30      ge-0/0/1          Idle (VO)
    40      ge-0/0/1          Idle (VO)
    50      ge-0/0/1          Idle (VO)
    100     ge-0/0/1          Idle (VO)
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

show mvrp | 22390

show mvrp interface | 22401

show mvrp registration-state | 22403

show mvrp statistics | 22406

show mvrp interface | 22401

show mvrp registration-state | 22403

show mvrp dynamic-vlan-memberships

IN THIS SECTION

- Syntax | 22398
- Description | 22398
- Required Privilege Level | 22398
- Output Fields | 22398
- Sample Output (MX Series Routers and EX Series Switches) | 22399
- Sample Output (EX Series Switches) | 22399
- Sample Output (SRX1500, SRX300, SRX550M, SRX345, SRX340, SRX320) | 22400
- Release Information | 22400

Syntax

```
show mvrp dynamic-vlan-memberships
```

Description

Display all VLANs that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the router, switch, or SRX Series Firewall.

Required Privilege Level

clear

Output Fields

Table 862 on page 22399 lists the output fields for the `show mvrp dynamic-vlan-memberships` command on MX Series routers and EX Series switches. Output fields are listed in the approximate order in which they appear.

Table 862: show mvrp dynamic-vlan-memberships Output Fields

Field Name	Field Description
VLAN Id	The VLAN ID of the dynamically created VLAN.
Interfaces	The interface or interfaces that are bound to the dynamically created VLAN.

Sample Output (MX Series Routers and EX Series Switches)

show mvrp dynamic-vlan- memberships (MX Series and EX Series)

```
user@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
 (s) static vlan, (f) fixed registration

VLAN Id            Interfaces
    100 (s)        ge-11/3/0
    200 (s)        ge-11/3/0
    300 (s)
```

Sample Output (EX Series Switches)

show mvrp dynamic-vlan-memberships (EX Series)

```
user@switch> show mvrp dynamic-vlan-memberships
VLAN Name             Interfaces
-------------------       ----------------
__mvrp_100__              xe-0/1/1.0
                          xe-0/1/0.0
__mvrp_200__              xe-0/1/1.0
                          xe-0/1/0.0
__mvrp_300__              xe-0/1/1.0
```

Sample Output (SRX1500, SRX300, SRX550M, SRX345, SRX340, SRX320)

show mvrp dynamic-vlan- memberships

```
user@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
 (s) static vlan, (f) fixed registration

VLAN Id            Interfaces
    1  (s)
    30 (s)
    40 (s)        ge-0/0/1
    50 (s)        ge-0/0/1
   100 (s)        ge-0/0/1 (f)
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches

Verifying That MVRP Is Working Correctly on Switches

show mvrp | **22390**

show mvrp applicant-state | **22394**

show mvrp interface | **22401**

show mvrp registration-state | **22403**

show mvrp registration-state | **22403**

show mvrp statistics | **22406**

show mvrp interface

IN THIS SECTION

- Syntax | 22401
- Description | 22401
- Required Privilege Level | 22401
- Output Fields | 22401
- Sample Output (MX Series Routers and SX Series Switches) | 22402
- Sample Output (SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320) | 22402
- Release Information | 22403

Syntax

```
show mvrp interface
```

Description

Display Multiple VLAN Registration Protocol (MVRP) interface-specific information.

Required Privilege Level

view

Output Fields

Table 863 on page 22402 lists the output fields for the `show mvrp interface` command. Output fields are listed in the approximate order in which they appear.

Table 863: show mvrp interface Output Fields

Field Name	Field Description
Interface	Interface on which MVRP is configured.
Status	Status of the MVRP: Enabled or Disabled.
Registration Mode	Registration for the interface: Fixed, Forbidden, or Normal.
Applicant Mode	Applicant mode.

Sample Output (MX Series Routers and SX Series Switches)

show mvrp interface

```
user@host> show mvrp interface
MVRP interface information for routing instance 'default-switch'

Interface       Status       Registration      Applicant
                             Mode              Mode
ge-11/2/8       Enabled      Normal            Normal
ge-11/0/9       Enabled      Normal            Normal
ge-11/3/0       Enabled      Normal            Normal
```

Sample Output (SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320)

show mvrp interface

```
user@host> show mvrp interface
MVRP interface information for routing instance 'default-switch'

Interface       Status       Registration      Applicant
                Mode         Mode              Mode
```

```
ge-0/0/1        Enabled      Normal            Normal
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

- show mvrp | 22390
- show mvrp applicant-state | 22394
- show mvrp dynamic-vlan-memberships | 22398
- show mvrp registration-state | 22403
- show mvrp registration-state | 22403
- show mvrp statistics | 22406

show mvrp registration-state

IN THIS SECTION

- Syntax | 22403
- Description | 22404
- Required Privilege Level | 22404
- Output Fields | 22404
- Sample Output | 22405
- Sample Output | 22405
- Release Information | 22405

Syntax

```
show mvrp registration-state
```

Description

For MX Series routers, EX Series switches and SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320, display Multiple VLAN Registration Protocol (MVRP) registration state information.

Required Privilege Level

view

Output Fields

Table 864 on page 22404 lists the output fields for the `show mvrp registration-state` command. Output fields are listed in the approximate order in which they appear.

Table 864: show mvrp registration-state Output Fields

Field Name	Field Description
`VLAN Id`	Displays the VLAN ID number.
`Interface`	Displays the interface number associated with the VLAN ID.
`Registrar State`	Displays whether the registrar state is Registered or Empty.
`Forced State`	Displays whether the forced state is Registered or Empty.
`Managed State`	Displays one of the following states: • `fixed`—VLANs always stay in a registered state and are declared as such on all other forwarding ports. • `normal` —VLANs participate in the MVRP protocol and honor incoming join requests normally. • `forbidden` —VLANs ignore the incoming join requests and always stay in an unregistered state.
`STP State`	Displays whether the Spanning Tree Protocol (STP) is Blocking or Forwarding.

Sample Output

show mvrp registration-state (EX Series and MX Series)

```
user@host> show mvrp registration-state
MVRP registration state for routing instance 'default-switch'

VLAN Id   Interface   Registrar   Forced      Managed    STP
                      State       State       State      State
    100   ge-11/2/8   Empty       Registered  Fixed      Forwarding
          ge-11/0/9   Empty       Empty       Normal     Forwarding
          ge-11/3/0   Registered  Registered  Normal     Forwarding
    101   ge-11/2/8   Empty       Registered  Fixed      Forwarding
          ge-11/0/9   Empty       Empty       Normal     Forwarding
          ge-11/3/0   Registered  Registered  Normal     Forwarding
```

Sample Output

show mvrp registration-state (SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320)

```
user@host> show mvrp registration-state
MVRP registration state for routing instance 'default-switch'

VLAN Id   Interface   Registrar   Forced      Managed    STP
                      State       State       State      State
    1     ge-0/0/1    Empty       Empty       Normal     Forwarding
    30    ge-0/0/1    Empty       Empty       Normal     Forwarding
    40    ge-0/0/1    Registered  Registered  Normal     Forwarding
    50    ge-0/0/1    Registered  Registered  Normal     Forwarding
   100    ge-0/0/1    Empty       Registered  Fixed      Forwarding
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

show mvrp | 22390
show mvrp dynamic-vlan-memberships | 22398
show mvrp interface | 22401
show mvrp statistics | 22406

show mvrp statistics

IN THIS SECTION

- Syntax (EX Series Switches) | 22406
- Syntax (Switches with ELS Support) | 22406
- Syntax (SRX Series Firewalls) | 22407
- Description | 22407
- Options | 22407
- Required Privilege Level | 22407
- Output Fields | 22407
- Sample Output | 22411
- Release Information | 22413

Syntax (EX Series Switches)

```
show mvrp statistics
<interface interface-name>
```

Syntax (Switches with ELS Support)

```
show mvrp statistics
<interface interface-name>
<routing-instance routing-instance-name>
```

Syntax (SRX Series Firewalls)

```
show mvrp statistics
```

Description

Display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages.

Options

none	Show MVRP statistics for all interfaces on the switch.
interface *interface-name*	(Optional) Show MVRP statistics for the specified interface.

Required Privilege Level

view

Output Fields

Table 865 on page 22407 lists the output fields for the `show mvrp statistics` command on EX Series switches. Output fields are listed in the approximate order in which they appear.

Table 865: show mvrp statistics Output Fields

Field Name	Field Description
MRPDU received	Number of MRPDU messages received on the switch.
Invalid PDU received	Number of invalid MRPDU messages received on the switch.
New received	Number of new messages received on the switch.

Table 865: show mvrp statistics Output Fields *(Continued)*

Field Name	Field Description
Join Empty received	Number of MRP JoinEmpty messages received on the switch. Either this value or the value for *JoinIn received* should increase when the value for *MRPDU received* increases. If this value is not incrementing when it should, you might have a Junos OS release version compatibility issue. To fix a version compatibility issue, see *Configuring Multiple VLAN Registration Protocol (MVRP) on Switches* .
Join In received	Number of MRP JoinIn messages received on the switch. Either this value or the value for *JoinEmpty received* should increase when the value for *MRPDU received* increases. If this value is not incrementing when it should, you might have a Junos OS release version compatibility issue. To fix a version compatibility issue, see *Configuring Multiple VLAN Registration Protocol (MVRP) on Switches* .
Empty received	Number of MRP Empty messages received on the switch.
In received	Number of MRP In messages received on the switch.
Leave received	Number of MRP Leave messages received on the switch.
LeaveAll received	Number of LeaveAll messages received on the switch.
MRPDU transmitted	Number of MRPDU messages transmitted from the switch.
MRPDU transmit failures	Number of MRPDU transmit failures from the switch.
New transmitted	Number of new messages transmitted from the switch.
Join Empty transmitted	Number of JoinEmpty messages sent from the switch.
Join In transmitted	Number of MRP JoinIn messages sent from the switch.

Table 865: show mvrp statistics Output Fields *(Continued)*

Field Name	Field Description
Empty transmitted	Number of MRP Empty messages sent from the switch.
In transmitted	Number of MRP In messages sent from the switch.
Leave transmitted	Number of MRP Leave Empty messages sent from the switch.
LeaveAll transmitted	Number of MRP LeaveAll messages sent from the switch.

Table 866 on page 22409 lists the output fields for the `show mvrp statistics` command on SRX Series Firewalls. Output fields are listed in the approximate order in which they appear.

Table 866: show mvrp statistics Output Fields

Field Name	Field Description
`Interface name`	Interface for which MVRP statistics are displayed.
`VLAN IDs registered`	Number of Virtual LAN (VLAN) IDs registered.
`Sent MVRP PDUs`	Number of MRPDU messages transmitted from the switch.
`Received MVRP PDUs without error`	Number of MRPDU messages received on the switch.
`Received MVRP PDUs with error`	Number of invalid MRPDU messages received on the switch.
`Transmitted Join Emply`	Number of JoinEmpty messages sent from the switch.

Table 866: show mvrp statistics Output Fields *(Continued)*

Field Name	Field Description
`Transmitted Leave All`	Number of MRP LeaveAll messages sent from the switch.
`Received Join In`	Number of MRP JoinIn messages received on the switch. Either this value or the value for `Received Join Empty` should increase when the value for `Received MVRP PDUs without error` increases. If this value is not incrementing when it should, you might have a Junos OS release compatibility issue. To resolve the issue, see *Configuring Multiple VLAN Registration Protocol (MVRP) on Switches* .
`Transmitted Join In`	Number of MRP JoinIn messages sent from the switch.
`Transmitted Empty`	Number of MRP Empty messages sent from the switch.
`Transmitted Leave`	Number of MRP LeaveEmpty messages sent from the switch.
`Transmitted In`	Number of MRP In messages sent from the switch.
`Transmitted New`	Number of New messages transmitted from the switch.
`Received Leave All`	Number of LeaveAll messages received on the switch.
`Received Leave`	Number of MRP Leave messages received on the switch.
`Received In`	Number of MRP In messages received on the switch.
`Received Empty`	Number of MRP Empty messages received on the switch.

Table 866: show mvrp statistics Output Fields *(Continued)*

Field Name	Field Description
`Received Join Empty`	Number of MRP JoinEmpty messages received on the switch. Either this value or the value for `Received Join In` should increase when the value for `Received MVRP PDUs without error` increases. If this value is not incrementing when it should, you might have a Junos OS release compatibility issue. To resolve the issue, see *Configuring Multiple VLAN Registration Protocol (MVRP) on Switches* .
`Received New`	Number of New messages received on the switch.

Sample Output

show mvrp statistics interface xe-0/1/1.0

```
user@switch> show mvrp statistics interface xe-0/1/1.0
MVRP statistics
  MRPDU received                : 3342
  Invalid PDU received          : 0
  New received                  : 2
  Join Empty received           : 1116
  Join In received              : 2219
  Empty received                : 2
  In received                   : 2
  Leave received                : 1
  LeaveAll received             : 1117
  MRPDU transmitted             : 3280
  MRPDU transmit failures       : 0
  New transmitted               : 0
  Join Empty transmitted        : 1114
  Join In transmitted           : 2163
  Empty transmitted             : 1
  In transmitted                : 1
  Leave transmitted             : 1
  LeaveAll transmitted          : 1111
```

show mvrp statistics

```
user@host> show mvrp statistics
MVRP statistics for routing instance 'default-switch'

Interface name                  : xe-0/1/1
VLAN IDs registered             : 117
Sent MVRP PDUs                  : 118824
Received MVRP PDUs without error: 118848
Received MVRP PDUs with error   : 0
Transmitted Join Empty          : 5229
Transmitted Leave All           : 2
Recieved Join In                : 11884924
Transmitted Join In             : 1835
Transmitted Empty               : 93606408
Transmitted Leave               : 888
Transmitted In                  : 13780024
Transmitted New                 : 2692
Received Leave All              : 118761
Received Leave                  : 97
Received In                     : 3869
Received Empty                  : 828
Received Join Empty             : 2020152
Received New                    : 224
...
```

show mvrp statistics (SRX Series Firewalls)

```
user@host> show mvrp statistics
MVRP statistics for routing instance 'default-switch'

Interface name                  : ge-0/0/1
VLAN IDs registered             : 2
Sent MVRP PDUs                  : 41
Received MVRP PDUs without error: 28
Received MVRP PDUs with error   : 0
Transmitted Join Empty          : 0
Transmitted Leave All           : 20
Received Join In                : 0
Transmitted Join In             : 0
```

```
Transmitted Empty              : 114
Transmitted Leave              : 0
Transmitted In                 : 10
Transmitted New                : 0
Received Leave All             : 1
Received Leave                 : 0
Received In                    : 0
Received Empty                 : 67
Received Join Empty            : 24
Received New                   : 0
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- show mvrp | 22390
- clear mvrp statistics | 3732
- *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches*
- *Verifying That MVRP Is Working Correctly on Switches*
- *Verifying That MVRP Is Working Correctly on EX Series Switches with ELS Support*

show neighbor-discovery-inspection statistics

IN THIS SECTION

- Syntax | 22414
- Description | 22414
- Required Privilege Level | 22414
- Output Fields | 22414
- Sample Output | 22415
- Release Information | 22415

Syntax

```
show neighbor-discovery-inspection statistics
```

Description

Display neighbor discovery inspection statistics.

Required Privilege Level

view

Output Fields

Table 867 on page 22414 lists the output fields for the `show neighbor-discovery-inspection statistics` command. Output fields are listed in the approximate order in which they appear.

Table 867: show neighbor-discovery-inspection statistics Output Fields

Field Name	Field Description	Level of Output
Interface	Interface on which neighbor discovery inspection has been applied.	All levels
Packets received	Total number of packets total that underwent neighbor discovery inspection.	All levels
ND inspection pass	Total number of packets that passed neighbor discovery inspection.	All levels
ND inspection failed	Total number of packets that failed neighbor discovery inspection.	All levels

Sample Output

show neighbor-discovery-inspection statistics

```
user@switch> show neighbor-discovery-inspection statistics
Interface     Packets received     ND inspection pass  ND inspection failed
  ge-0/0/0                    5                      1                     4
  ge-0/0/1                    0                      0                     0
```

Release Information

Command introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

clear arp inspection statistics | 3426

Example: Configuring Port Security (non-ELS)

Verifying That DAI Is Working Correctly

show network-access aaa accounting

IN THIS SECTION

- Syntax | 22416
- Description | 22416
- Required Privilege Level | 22416
- Output Fields | 22416
- Sample Output | 22417
- Release Information | 22417

Syntax

```
show network-access aaa accounting
```

Description

Display the state of the RADIUS Acct-On response sent from the RADIUS server.

Required Privilege Level

view

Output Fields

Table 868 on page 22416 lists the output fields for the `show network-access aaa accounting` command. Output fields are listed in the approximate order in which they appear.

Table 868: show network-access aaa accounting Output Fields

Field Name	Field Description
`Profile`	Name of the profile associated with the RADIUS server. A RADIUS server can be associated with more than one profile.
`Logical System`	Logical system associated with the access profile.
`Routing Instance`	Routing instance associated with the access profile.
`Acct-On-Response`	Status of the RADIUS Acct-On response. • ACK—ACK response for the Acct-On message is received from the RADIUS server. • ERROR—An error condition has occurred. • NONE— No Acct-On message is sent. • PENDING—Acct-On message is sent to RADIUS server, but no response has been received yet.

Sample Output

show network-access aaa accounting

```
user@host> show network-access aaa accounting
Profile        Logical System   Routing Instance   Acct-On-Response
ppp-profile    default          default            ACK
l2tp-profile   default          l2tp_RI            PENDING
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

RADIUS Accounting for Subscriber Access

show network-access aaa radius-servers

IN THIS SECTION

- Syntax | 22418
- Description | 22418
- Options | 22418
- Required Privilege Level | 22418
- Output Fields | 22418
- Sample Output | 22425
- Release Information | 22428

Syntax

```
show network-access aaa radius-servers
<detail>
```

Description

Display RADIUS server status and information.

Options

detail (Optional) Display detailed level of information.

Required Privilege Level

view

Output Fields

Table 869 on page 22418 lists the output fields for the `show network-access aaa radius-servers` command. Output fields are listed in the approximate order in which they appear.

Table 869: show network-access aaa radius-servers Output Fields

Field Name	Field Description	Level of Output
`Profile`	Name of the profile associated with the RADIUS server. A RADIUS server can be associated with more than one profile.	All levels
`Server address`	IPv4 or IPv6 address of the RADIUS server.	All levels
`Authentication port`	RADIUS server authentication port number.	All levels
`Preauthentication port`	RADIUS server preauthentication port number.	All levels

Table 869: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting port`	RADIUS server accounting port number.	All levels
`Accounting retry`	Number of times the router retransmits RADIUS accounting messages when no response is received from the server.	Detail
`Accounting timeout`	Period the local router waits to receive a response from a RADIUS accounting server before retransmitting the message.	Detail

Table 869: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Status`	RADIUS server status, `UP` (Alive), `UNREACHABLE`, or `DOWN` (DEAD). If status is `DOWN`, the Status field includes the number of seconds configured by the `revert-interval` statement. The router does not send requests to servers in the `DOWN` state, but does send requests to servers with a status of either `UP` or `UNREACHABLE`. This field also displays the status of AAA accounting suspension or resumption, and the status of baselining of accounting statistics if you suspended or resumed accounting operations or initiated the generation of a baseline. This information is applicable only for RADIUS servers that are in the `UP` state. **NOTE**: After requests to a server or set of servers time out after 10 seconds, the status of the servers changes. The following guidelines apply to server status: • For the purpose of marking a server as `Down` (DEAD), the request includes the original request and any retries that are configured. The 10-second timeout period starts after the initial request and all retries have expired without receiving a response from the server. The amount of the timeout period that elapses before the server is marked `Down` is not always exactly 10 seconds, and can vary depending on how frequently subscribers are logging in. When subscribers are continually and rapidly logging in, the server is marked as `Down` at 10 seconds. However, if subscribers are logging in less frequently and at a slower pace, then the server is not marked `Down` until a subsequent subscriber attempts to log in. • All servers cannot be marked as `DOWN`; instead, the unresponsive servers are marked as `UNREACHABLE`. For example, if only one RADIUS server is configured and that server is unresponsive, the server status is marked as `UNREACHABLE` rather than `DOWN`.	All levels

Table 869: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• If at least one server has a status of UP, the status of all unresponsive servers is set to DOWN for the remainder of the configured revert-interval setting. • If no server has a status of UP, then the status of the unresponsive servers is set to UNREACHABLE for the remainder of the revert-interval setting or for 30 seconds, whichever is less. • The status of unresponsive servers is returned to UP from DOWN or UNREACHABLE at the end of the revert-interval setting (or the 30-second interval). • If no requests are sent to a server, the server's status is always UP.	
RADIUS servers	Details for specific RADIUS server, identified by IP address.	Detail
Authentication requests	Number of authentication requests received by the authentication server.	Detail
Authentication rollover requests	Number of requests coming into the server as a result of the previous server timing out.	Detail
Authentication retransmissions	Number of retransmissions.	Detail
Accepts	Number of authentication requests accepted by the authentication server.	Detail
Rejects	Number of authentication requests rejected by the authentication server.	Detail
Challenges	Number of authentication requests challenged by the authentication server.	Detail

Table 869: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Authentication malformed responses`	Number of responses with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).	Detail
`Authentication bad authenticators`	Number of responses in which the authenticator is incorrect for the authentication request. This can occur if the RADIUS secrets for the client and server do not match.	Detail
`Authentication requests pending`	Number of authentication requests waiting for a response.	Detail
`Authentication request timeouts`	Number of times an authentication request to the server timed out.	Detail
`Authentication unknown responses`	Number of unknown responses. The RADIUS response type in the header is invalid or unsupported.	Detail
`Authentication packets dropped`	Number of packets dropped because they are too short or because the router receives a response for which there is no corresponding request.	Detail
`Preauthentication requests`	Number of preauthentication requests received by the preauthentication server.	Detail
`Preauthentication rollover requests`	Number of preauthentication requests coming into the server as a result of the previous server timing out.	Detail
`Preauthentication retransmissions`	Number of retransmissions of preauthentication requests.	Detail
`Preauthentication Accepts`	Number of preauthentication requests accepted by the preauthentication server.	Detail

Table 869: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Preauthentication Rejects`	Number of preauthentication requests rejected by the preauthentication server.	Detail
`Preauthentication Challenges`	Number of preauthentication requests challenged by the preauthentication server.	Detail
`Preauthentication malformed responses`	Number of responses to preauthentication requests with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).	Detail
`Preauthentication bad authenticators`	Number of responses in which the authenticator is incorrect for the preauthentication request. This can occur if the RADIUS secrets for the client and server do not match.	Detail
`Preauthentication requests pending`	Number of preauthentication requests waiting for a response.	Detail
`Preauthentication request timeouts`	Number of times a preauthentication request to the server timed out.	Detail
`Preuthentication unknown responses`	Number of unknown responses during the preauthentication phase. The RADIUS response type in the header is invalid or unsupported.	Detail
`Preauthentication packets dropped`	Number of preauthentication packets dropped because they are too short or because the router receives a response for which there is no corresponding request.	Detail
`Accounting start requests`	Number of accounting start requests received.	Detail
`Accounting interim requests`	Number of accounting interim requests received.	Detail

Table 869: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting stop requests`	Number of accounting stop requests received.	Detail
`Accounting rollover requests`	Number of requests coming into the server as a result of the previous server timing out.	Detail
`Accounting retransmissions`	Number of retransmissions.	Detail
`Accounting start responses`	Number of accounting start responses sent by the server.	Detail
`Accounting interim responses`	Number of accounting interim responses sent by the server.	Detail
`Accounting stop responses`	Number of accounting stop responses sent by the server.	Detail
`Accounting malformed responses`	Number of responses with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).	Detail
`Accounting bad authenticators`	Number of responses in which the authenticator is incorrect for the accounting request. This can occur if the RADIUS secrets for the client and server do not match.	Detail
`Accounting requests pending`	Number of accounting requests waiting for a response.	Detail
`Accounting request timeouts`	Number of accounting requests to the accounting server that timed out.	Detail

Table 869: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting unknown responses`	Number of unknown responses. The RADIUS response type in the header is invalid or unsupported.	Detail
`Accounting packets dropped`	Number of packets dropped because they are too short or because the router receives a response for which there is no corresponding request.	Detail

Sample Output

show network-access aaa radius-servers

```
user@host> show network-access aaa radius-servers
Profile: xyz-profile1
    Server address: 192.168.30.188
      Authentication port: 1645
      Preauthentication port: 1810
      Accounting port: 1646
      Status: UP
Profile: xyz-profile2
    Server address: 192.168.30.190
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Status: DOWN ( 60 seconds )
Profile: xyz-profile11
    Server address: 2001:DB8:0:f101::2
      Authentication port: 1645
      Preauthentication port: 1810
      Accounting port: 1646
      Status: UP
```

show network-access aaa radius-servers

```
user@host> show network-access aaa radius-servers
Profile: xyz-profile3
    Server address: 192.168.30.188
      Authentication port: 1645
      Preauthentication port: 1810
      Accounting port: 1646
      Status: UNREACHABLE
Profile: xyz-profile3
    Server address: 192.168.30.190
      Authentication port: 1812
      Accounting port: 1813
      Preauthentication port: 1810
      Status: UNREACHABLE
```

show network-access aaa radius-servers detail

```
user@host> show network-access aaa radius-servers detail
Profile: xyz_profile5
    Server address: 192.168.30.188
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Status: UP (accounting suspended, baseline in progress)
    Server address: 192.168.30.190
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Accounting retry: 5
      Accounting port: 60
      Status: UP (accounting suspended, baseline in progress)
    Server address: 192.168.30.192
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Status: UP
    Server address: 192.168.30.190
      Authentication port: 1812
      Accounting port: 1813
```

```
      Accounting retry: 5
      Accounting port: 60
      Status: UP
    Server address: 192.168.30.192
      Authentication port: 1812
      Accounting port: 1813
      Status: UP

RADIUS Servers
  192.168.30.188
    Authentication requests: 7658
    Authentication rollover requests: 0
    Authentication retransmissions: 3600
    Accepts: 6458
    Rejects: 0
    Challenges: 0
    Authentication malformed responses: 0
    Authentication bad authenticators: 0
    Authentication requests pending: 0
    Authentication request timeouts: 4800
    Authentication unknown responses: 0
    Authentication packets dropped: 0
    Preauthentication requests: 7658
    Preauthentication rollover requests: 0
    Preauthentication retransmissions: 3600
    Preauthentication Accepts: 6458
    Preauthentication Rejects: 0
    Preauthentication Challenges: 0
    Preauthentication malformed responses: 0
    Preauthentication bad authenticators: 0
    Preauthentication requests pending: 0
    Preauthentication request timeouts: 4800
    Preauthentication unknown responses: 0
    Preauthentication packets dropped: 0
    Accounting start requests: 1
    Accounting interim requests: 1
    Accounting stop requests: 0
    Accounting rollover requests: 0
    Accounting retransmissions: 0
    Accounting start responses: 1
    Accounting interim responses: 1
    Accounting stop responses: 0
    Accounting malformed responses: 0
```

```
    Accounting bad authenticators: 0
    Accounting requests pending: 0
    Accounting request timeouts: 0
    Accounting unknown responses: 0
    Accounting packets dropped: 0
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information

show network-access aaa statistics

IN THIS SECTION

- Syntax | 22428
- Description | 22429
- Options | 22429
- Required Privilege Level | 22429
- Output Fields | 22429
- Sample Output | 22442
- Release Information | 22446

Syntax

```
show network-access aaa statistics
<accounting (detail)>
<address-assignment (client | pool pool-name)>
<dynamic-requests>
<radius>
```

```
<saml>
<session-limit-per-username>
```

Description

Display AAA accounting, address-assignment, dynamic request statistics, RADIUS settings and statistics, Security Assertion Markup Language (SAML) authentication statistics, and subscriber session limit statistics.

Options

accounting (detail)	(Optional) Display AAA accounting statistics. The `detail` keyword displays additional accounting information
address-assignment (client | pool *pool-name*)	(Optional) Display AAA address-assignment client and pool statistics.
dynamic-requests	(Optional) Display AAA dynamic requests.
radius	(Optional) Display RADIUS settings and statistics.
saml	(Optional) Display SAML authentication statistics.
session-limit-per-username	Maximum number of sessions allowed for a username per access profile. Use the `brief` option to display only active users with blocked requests. Use the `detail` option to display all active users.

Required Privilege Level

view

Output Fields

Table 870 on page 22430 lists the output fields for the `show network-access aaa statistics` command. Output fields are listed in the approximate order in which they appear.

Table 870: show network-access aaa statistics Output Fields

Field Name	Field Description	Level of Output
`Requests received`	• Number of accounting requests generated by the AAA framework. • Number of dynamic requests received from the external server. Does not include requests sent from backup accounting.	All levels
`Accounting request failures`	Number of accounting requests that failed to be sent or queued from a client to a RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`
`Accounting request success`	Number of accounting requests successfully sent or queued from a client to a RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`
`Account on requests`	Number of accounting on requests sent from a client to a RADIUS accounting server.	`detail`
`Accounting start requests`	Number of accounting start requests sent from a client to a RADIUS accounting server.	`detail`
`Accounting interim requests`	Number of accounting interim requests sent from a client to a RADIUS accounting server.	`detail`

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting stop requests`	Number of accounting stop requests sent from a client to a RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`
`Accounting request timeouts`	Number of accounting requests to the accounting server that timed out. This field was named `Timed out requests` in releases before Junos OS Release 16.1. Does not include requests sent from backup accounting.	All levels
`Accounting Response failures`	Number of accounting requests not acknowledged (NAK) by the accounting server. Does not include requests sent from backup accounting.	All levels
`Accounting response success`	Number of accounting requests acknowledged by the accounting server. Does not include requests sent from backup accounting.	All levels
`Account on responses`	Number of accounting on requests acknowledged by the RADIUS accounting server.	`detail`
`Accounting start responses`	Number of accounting start requests acknowledged by the RADIUS accounting server.	`detail`
`Accounting interim responses`	Number of accounting interim requests acknowledged by the RADIUS accounting server.	`detail`

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting stop responses`	Number of accounting stop requests acknowledged by the RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`
`Accounting rollover requests`	Number of accounting requests coming to a RADIUS accounting server after a previous server timing out.	`detail`
`Accounting unknown requests`	Number of unknown accounting requests sent from a client to a RADIUS accounting server (for example, when the header has invalid or unsupported information).	`detail`
`Accounting radius pending requests`	Number of accounting requests sent from a client to a RADIUS accounting server that are waiting for a response from the server.	`detail`
`Accounting malformed responses`	Number of accounting responses from a RADIUS accounting server that have invalid or unexpected attributes.	`detail`
`Accounting retransmissions`	Number of accounting requests made by a client to the RADIUS sever that were retransmitted. Does not include requests sent from backup accounting.	`detail`
`Accounting bad authenticators`	Number of accounting responses from a RADIUS accounting server that have an incorrect authenticator (for example, the client and server RADIUS secret do not match).	`detail`

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting packets dropped`	Number of accounting responses from a RADIUS accounting server that are dropped by a client.	`detail`
`Accounting backup record creation requests`	Number of accounting stop requests from a client to a RADIUS accounting server that were forwarded to be backed up.	`detail`
`Accounting backup replay request success`	Number of backup accounting stop requests successfully created by clients after each timeout for replay to a RADIUS accounting server.	`detail`
`Accounting backup request failures`	Number of backup accounting requests that failed to be sent or queued from a client to a RADIUS accounting server.	`detail`
`Accounting backup request success`	Number of backup accounting requests successfully sent or queued from a client to a RADIUS accounting server.	`detail`
`Accounting backup timeouts`	Number of backup accounting requests that timed out after being sent to a RADIUS accounting server.	`detail`

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting backup in-flight requests`	Number of backup accounting requests that were successfully sent or queued to a RADIUS accounting server for which no response or error has been received yet. Backup requests are replayed only in the following circumstances: • When the request being replayed receives a positive response, the next request can be replayed. • When the request being replayed receives a timeout response, it can be replayed again. Consequently this intermediate timer displays 1 or 0. The value eventually drops to 0 as requests are responded to positively or fail due to error.	`detail`
`Accounting backup responses success`	Number of backup records that were successfully acknowledged with a positive response from a RADIUS accounting server.	`detail`
`Accounting backup radius requests`	Number of backup requests sent to UDP level. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. An observation that the value is increasing is more significant than the exact value of the counter.	`detail`

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting backup radius responses`	Number of responses received at the UDP level for backup requests. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	`detail`
`Accounting backup radius timeouts`	Number of backup requests that timed out after being sent to UDP. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	`detail`
`Accounting backup radius pending requests`	Number of backup requests sent to a RADIUS accounting server that are waiting for a response from the server. This is an intermediate state counter that eventually drops to zero as requests are responded to or failed due to error.	`detail`
`Accounting backup radius retransmissions`	Sum of backup request retransmissions for each RADIUS accounting server. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	`detail`

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting backup malformed responses`	Sum of malformed responses received for backup requests sent to each RADIUS accounting server at the UDP level.	`detail`
`Accounting backup bad authenticators`	Sum of responses received for backup accounting requests for each RADIUS accounting server where authenticators were mismatched.	`detail`
`Accounting backup responses dropped`	Sum of responses for backup accounting requests for each RADIUS accounting server that were dropped due to various sanity checks.	`detail`
`Accounting backup rollover requests`	Sum of backup accounting requests rolled over for each RADIUS accounting server.	`detail`
`Accounting backup unknown responses`	Sum of unknown responses for backup accounting requests for each RADIUS accounting server.	`detail`
`Client`	Client type; for example, DHCP, Mobile IP, PPP.	none specified
`Out of Memory`	Number of times an address was not given to the client due to memory issues.	none specified
`No Matches`	Number of times there were no network matches for the pool.	none specified

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Pool Name`	Name of the address-assignment pool for this client.	none specified
`Out of Addresses`	Number of times there were no available addresses in the pool.	none specified
`Address total`	Number of addresses in the pool.	none specified
`Addresses in use`	Number of addresses in use.	none specified
`Addresses excluded`	Number of addresses excluded from being allocated from the pool with the `excluded-address` or `excluded-range` statements.	none specified
`Address Usage (percent)`	Percentage of total addresses in use. This value does not take excluded addresses into account.	none specified
`Assertion attributes missing`	Number of SAML assertions where mandatory attributes like username configured at the identity-provider object are not present in the assertion.	none specified
`Assertion decryption fail`	Instances where the identity provider (IdP) could not decrypt the assertion it had encrypted in SAML.	none specified
`Assertion node missing`	Instances when the SAML assertion or the encrypted-assertion node is missing in the assertion response.	none specified

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Assertion parse fail	Instances when the firewall failed to parse SAML assertion.	none specified
`Assertion received`	Number of SAML assertion responses received from IdP.	none specified
Assertion sanity fail	Instances when the issuer node value failed the sanity checks because it does not match the configured IdP entity-id in SAML.	none specified
Assertion signature verify fail	Count of SAML assertion message signature verification failures.	none specified
`Assertion timeout`	Number of requests that did not receive a response from the IdP before the assertion wait time expired in SAML.	none specified
Assertion username mismatch	Instances when the username differs in the authentication request and the assertion response in SAML.	none specified
`Authentication request received`	Number of SAML-based authentication requests received.	none specified
`Authentication response sent`	Number of authentication responses sent in SAML.	none specified
`Logout parse fail`	Instances of failures in parsing in the logout request message in SAML.	none specified

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Logout request received`	Number of user logout requests received from the IdP in SAML.	none specified
`Logout response sent`	Number of user logout responses sent to the IdP in SAML.	none specified
`Logout sanity fail`	Count of logout request failed sanity checks such as issuer node value does not match configured IdP entity-id in SAML.	none specified
`Logout signature verify fail`	Count of logout request signature validation failures in SAML.	none specified
`Memory allocation fail`	Number of memory allocation failure while processing SAML request/responses.	none specified
`Request hit cache`	Count of SAML requests that hit the assertion cache.	none specified
`Request sent to IdP`	SAML requests miss assertion cache and sent to IdP.	none specified
`Pool drain configured`	Configuration state of active drain for the specified local address pool, `yes` or `no`.	none specified
`Pool Usage`	Percentage of allocated addresses in the specified address pool.	none specified

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`processed successfully`	Number of dynamic requests processed successfully by the AAA framework.	All levels
`errors during processing`	Number of dynamic requests that resulted in processing errors by the AAA framework.	All levels
`Link Name`	Name of the secondary address-assignment pool to which the primary pool is linked.	
`silently dropped`	Number of dynamic requests dropped by the AAA framework due to multiple back-to-back or duplicate requests.	All levels
`RADIUS Server`	IPv4 or IPv6 address of the RADIUS server to which the router is sending requests.	All levels
`Profile`	Name of the RADIUS profile associated with the RADIUS server. A RADIUS server can be associated with more than one RADIUS profile.	All levels
`Configured`	Configured maximum number of outstanding requests from the router to the RADIUS server for a specific profile. An outstanding request is a request to which the RADIUS server has not yet responded. The range of values is 0 through 2000 outstanding requests. The default value is 1000.	All levels
`Current`	Current number of outstanding requests from the router to the RADIUS server for a specific profile. An outstanding request is a request to which the RADIUS server has not yet responded.	All levels

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Peak`	Highest number of outstanding requests from the router to the RADIUS server for a specific profile at any point in time since the router was started or since the counter was last cleared. **NOTE**: If the value of this field is equal to the value of the `Configured` field, you may want to increase the value of the `Configured` field.	All levels
`Exceeded`	Number of times that the router attempted to send requests to the RADIUS server in excess of the configured maximum value for a specific profile. **NOTE**: If the value of this field is nonzero, you may want to increase the value of the `Configured` field.	All levels
`Username`	Username for a subscriber with one or more active sessions for an access profile.	`briefdetail`
`Access-profile`	Name of the access profile where the username is active.	`briefdetail`
`Blocked requests`	Number of session requests that have been blocked for the username for an access profile. A request is blocked when it exceeds the configured session limit.	`briefdetail`
`Session count`	Number of active sessions for the username for an access profile.	`briefdetail`
`Total usernames`	Number of active usernames for all access profiles.	none `summary`

Table 870: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Total usernames exceeding session limit`	Number of usernames that have attempted sessions greater than the limit configured for the username.	none `summary`
`Total blocked requests`	Number of session requests that have been blocked because the session limit is exceeded.	none `summary`

Sample Output

show network-access aaa statistics accounting

```
user@host> show network-access aaa statistics accounting
Accounting module statistics
Accounting module statistics
  Requests received: 5000
  Accounting request timeouts: 2000
  Accounting response failures: 0
  Accounting response success: 3000
```

show network-access aaa statistics accounting detail

```
user@host> show network-access aaa statistics accounting detail
Accounting module statistics
Accounting module statistics
  Requests received: 5000
  Accounting request failures: 0
  Accounting request success: 5000
    Account on requests: 0
    Accounting start requests: 3000
    Accounting interim requests: 0
    Accounting stop requests: 2000
```

```
  Accounting request timeouts: 2000
  Accounting response failures: 0
  Accounting response success: 3000
    Account on responses: 0
    Accounting start responses: 3000
    Accounting interim responses: 0
    Accounting stop responses: 0
  Accounting rollover requests: 0
  Accounting unknown responses: 0
  Accounting radius pending requests: 0
  Accounting malformed responses: 0
  Accounting retransmissions: 6000
  Accounting bad authenticators: 0
  Accounting packets dropped: 0

  Accounting backup record creation requests: 3000
  Accounting backup request replay success: 9808
  Accounting backup request failures: 0
  Accounting backup request success: 3006
  Accounting backup timeouts: 6
  Accounting backup in-flight requests: 0
  Accounting backup responses success: 3000
  Accounting backup radius requests: 3006
  Accounting backup radius responses: 3000
  Accounting backup radius timeouts: 99
  Accounting backup radius pending requests: 0
  Accounting backup radius retransmissions: 99
  Accounting backup malformed responses: 0
  Accounting backup bad authenticators: 0
  Accounting backup responses dropped: 0
  Accounting backup rollover requests: 0
  Accounting backup unknown responses: 0
```

show network-access aaa statistics address-assignment client

```
user@host> show network-access aaa statistics address-assignment client
Address-assignment statistics
 Client: jdhcpd
 Out of Memory: 0
 No Matches: 2
```

show network-access aaa statistics address-assignment pool

```
user@host> show network-access aaa statistics address-assignment pool isp_1
Address-assignment statistics
 Pool Name: isp_1
 Pool Name: (all pools in chain)
 Out of Memory: 0
 Out of Addresses: 9
 Address total: 47
 Addresses in use: 47
 Address Usage (percent): 100
 Pool drain configured: yes
```

show network-access aaa statistics address-assignment pool (Excluded Addresses)

```
user@host> show network-access aaa statistics address-assignment pool isp_1
Address-assignment statistics
 Pool Name: isp_1
 Pool Name: (all pools in chain)
 Out of Memory: 0
 Out of Addresses: 0
 Address total: 24000
 Addresses in use: 12000
 Addresses excluded: 1000
 Address Usage (percent): 50
 Pool drain configured: yes
```

show network-access aaa statistics dynamic-requests

```
user@host> show network-access aaa statistics dynamic-requests
requests received: 0
processed successfully: 0
errors during processing: 0
silently dropped: 0
```

show network-access aaa statistics radius

```
user@host> show network-access aaa statistics radius
Outstanding Requests
RADIUS Server            Profile        Configured   Current   Peak   Exceeded
198.51.100.239           prof1          1000         0         1000   14
                         prof2          500          17        432    0
198.51.100.211           myprof         200          0         200    27
203.0.113.254            pppoe-auth     111          0         1      0
2001:db8:0:f101::2       xyz-profile11  1000         10        135    0
```

show network-access aaa statistics saml

```
user@host> show network-access aaa statistics saml
SAML Authentication statistics
    Authentication request received        13
    Authentication response sent           13
    Request hit cache                      1
    Request sent to IdP                    3
    Assertion received                     3
    Assertion timeout                      0
    Assertion parse fail                   0
    Assertion sanity fail                  0
    Assertion signature verify fail        0
    Assertion decryption fail              1
    Assertion node missing                 0
    Assertion atributes missing            0
    Assertion username mismatch            0
    Logout request received                0
    Logout response sent                   0
    Logout parse fail                      0
    Logout sanity fail                     0
    Logout signature verify fail           0
    Memory allocation fail                 0
```

show network-access aaa statistics session-limit-per-username (Users with Blocked Requests)

```
user@host> show network-access aaa statistics session-limit-per-username brief
Username           Access-profile     Blocked requests     Session count
xyz@example.net    BNG1               3                    5
abc@example.net    BNG2               2                    5
```

show network-access aaa statistics session-limit-per-username (All Active Users)

```
user@host> show network-access aaa statistics session-limit-per-username detail
Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               0                    4
xyz@example.net    BNG1               3                    5
abc@example.net    BNG2               2                    5
pqr@example.net    BNG2               0                    1
```

show network-access aaa statistics session-limit-per-username

```
user@host> show network-access aaa statistics on-limit-per-username
Total usernames:  15
Total usernames exceeding session limit: 2
Total blocked requests:  5
```

Release Information

Command introduced in Junos OS Release 9.1.

`address-assignment` option introduced in Junos OS Release 10.0.

`radius` option introduced in Junos OS Release 11.4.

`detail` option introduced in Junos OS Release 13.3.

`session-limit-per-username` option introduced in Junos OS Release 18.4R1 on MX Series routers.

Support for `saml` option added in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information

Session Options for Subscriber Access

clear network-access aaa statistics | 3736

show network-access aaa statistics accounting

IN THIS SECTION

- Syntax | 22447
- Description | 22447
- Required Privilege Level | 22447
- Output Fields | 22447
- Sample Output | 22448
- Release Information | 22448

Syntax

```
show network-access aaa statistics accounting;
```

Description

Display authentication, authorization, and accounting (AAA) accounting statistics.

Required Privilege Level

view

Output Fields

Table 871 on page 22448 lists the output fields for the `show network-access aaa statistics accounting` command. Output fields are listed in the approximate order in which they appear.

Table 871: show network-access aaa statistics accounting Output Fields

Field Name	Field Description
Requests received	The number of accounting-request packets sent from a switch to a RADIUS accounting server.
Accounting Response failures	The number of accounting-response failure packets sent from the RADIUS accounting server to the switch.
Accounting Response Success	The number of accounting-response success packets sent from the RADIUS accounting server to the switch.
Requests timedout	The number of requests-timedout packets sent from the RADIUS accounting server to the switch.

Sample Output

show network-access aaa statistics accounting

```
user@switch> show network-access aaa statistics accounting
Accounting module statistics
  Requests received: 1
  Accounting Response failures: 0
  Accounting Response Success: 1
  Requests timedout: 0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

accounting-server | 1739

accounting-stop-on-access-deny

Configuring 802.1X RADIUS Accounting (CLI Procedure)

show network-access aaa statistics authentication

IN THIS SECTION

- Syntax | 22449
- Description | 22449
- Options | 22449
- Required Privilege Level | 22449
- Output Fields | 22449
- Sample Output | 22452
- Release Information | 22453

Syntax

```
show network-access aaa statistics authentication
<detail>
```

Description

Display AAA authentication statistics.

Options

detail (Optional) Displays detailed information about authentication.

Required Privilege Level

view

Output Fields

Table 872 on page 22450 lists the output fields for the `show network-access aaa statistics authentication` command. Output fields are listed in the approximate order in which they appear.

Table 872: show network-access aaa statistics authentication Output Fields

Field Name	Field Description	Level of Output
`Requests received`	Number of authentication requests received from clients.	All levels
`Accepts`	Number of authentication requests accepted by the authentication server.	All levels
`Rejects`	Number of authentication requests rejected by the authentication server.	All levels
`Challenges`	Number of authentication requests challenged by the authentication server.	All levels
`Timed out requests`	Number of authentication requests that timed out.	All levels
`RADIUS authentication failures`	Number of RADIUS authentication requests that have failed.	Detail
`Queue request deleted`	Number of queue requests that have been deleted.	Detail
`Malformed reply`	Number of malformed replies received from the RADIUS authentication server.	Detail
`No server configured`	Number of authentication requests that failed because no authentication server is configured.	Detail
`Access Profile configuration not found`	Number of authentication requests that failed because no access profile is configured.	Detail
`Unable to create client record`	Number of times that the router is unable to create the client record for the authentication request.	Detail
`Unable to create client request`	Number of times that the router is unable to create the client request for the authentication request.	Detail

Table 872: show network-access aaa statistics authentication Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Unable to build authentication request`	Number of times that the router is unable to build the authentication request.	Detail
`No server found`	Number of requests to the authentication server that have timed out; the server is then considered to be down.	Detail
`Unable to create handle`	Number of authentication requests that have failed because of an internal allocation failure.	Detail
`Unable to queue request`	Number of times the router was unable to queue the request to the authentication server.	Detail
`Invalid credentials`	Number of times the router did not have proper authorization to access the authentication server.	Detail
`Malformed request`	Number of times the router request to the authentication server is malformed.	Detail
`License unavailable`	Number of times the router did not have a license to access the authentication server.	Detail
`Redirect requested`	Number of authentication requests that have been redirected based on routing instance.	Detail
`Internal failure`	Number of internal failures.	Detail
`Local authentication failures`	Number of times local authentication failed.	Detail
`LDAP lookup failures`	Number of times the LDAP lookup operation failed.	Detail

Sample Output

show network-access aaa statistics authentication

```
user@host> show network-access aaa statistics authentication
Authentication module statistics
  Requests received: 2118
    Accepts: 261
  Rejects: 975
  Challenges: 0
  Timed out requests: 882
```

show network-access aaa statistics authentication detail

```
user@host> show network-access aaa statistics authentication detail
Authentication module statistics
  Requests received: 2118
    Accepts: 261
  Rejects: 975
    RADIUS authentication failures: 975
      Queue request deleted: 0
      Malformed reply: 0
      No server configured: 0
      Access Profile configuration not found: 0
      Unable to create client record: 0
      Unable to create client request: 0
      Unable to build authentication request: 0
      No server found: 975
      Unable to create handle: 0
      Unable to queue request: 0
      Invalid credentials: 0
      Malformed request: 0
      License unavailable: 0
      Redirect requested: 0
      Internal failure: 0
    Local authentication failures: 0
    LDAP lookup failures: 0
  Challenges: 0
  Timed out requests: 882
```

Release Information

Command introduced in Junos OS Release 9.1.

Option `detail` introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information

show network-access aaa statistics dynamic-requests

IN THIS SECTION

- Syntax | 22453
- Description | 22453
- Required Privilege Level | 22453
- Output Fields | 22454
- Sample Output | 22454
- Release Information | 22454

Syntax

```
show network-access aaa statistics dynamic-requests;
```

Description

Display authentication, authorization, and accounting (AAA) authentication statistics for disconnects.

Required Privilege Level

view

Output Fields

Table 873 on page 22454 lists the output fields for the `show network-access aaa statistics dynamic-requests` command. Output fields are listed in the approximate order in which they appear.

Table 873: show network-access aaa statistics dynamic-requests Output Fields

Field Name	Field Description
Requests received	The number of dynamic requests received by the RADIUS server.
Processed successfully	The number of dynamic requests successfully processed by the RADIUS server.
Errors during processing	The number of errors that occurred while the RADIUS server was processing the dynamic request.
Silently dropped	The number of silently dropped requests.

Sample Output

show network-access aaa statistics authentication

```
user@switch> show network-access aaa statistics dynamic-requests
Dynamic-requests module statistics
  Requests received: 0
  Processed successfully: 0
  Errors during processing: 0
  Silently dropped: 0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

authentication-server | 2652

Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch

show network-access aaa statistics pending-accounting-stops

IN THIS SECTION

- Syntax | 22455
- Description | 22455
- Options | 22455
- Required Privilege Level | 22455
- Output Fields | 22456
- Sample Output | 22456
- Release Information | 22456

Syntax

```
show network-access aaa statistics pending-accounting-stops
```

Description

Display the number of pending accounting stop requests.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 874 on page 22456 lists the output field for the `show network-access aaa statistics pending-accounting-stops` command.

Table 874: show network-access aaa statistics pending-accounting-stops Output Fields

Field Name	Field Description
`Pending accounting stops`	Total number of accounting stop messages queued.

Sample Output

show network-access aaa statistics pending-accounting-stops

```
user@host> show network-access aaa statistics pending-accounting-stops
Pending accounting stops: 10,000
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

request network-access aaa replay pending-accounting-stops | 14695

show accounting pending-accounting-stops | 17011

show network-access aaa statistics preauthentication

IN THIS SECTION

Syntax | 22457

- Description | 22457
- Options | 22457
- Required Privilege Level | 22457
- Output Fields | 22457
- Sample Output | 22458
- Release Information | 22458

Syntax

```
show network-access aaa statistics preauthentication
```

Description

Display AAA preauthentication statistics.

Options

detail (Optional) Displays detailed information about authentication.

Required Privilege Level

view

Output Fields

Table 875 on page 22458 lists the output fields for the `show network-access aaa statistics preauthentication` command. Output fields are listed in the approximate order in which they appear.

Table 875: show network-access aaa statistics preauthentication Output Fields

Field Name	Field Description	Level of Output
`Requests received`	Number of preauthentication requests received from clients.	All levels
`Multistack requests`	Number of preauthentication requests for dual-stack subscribers.	All levels
`Accepts`	Number of preauthentication requests accepted by the preauthentication server.	All levels
`Rejects`	Number of preauthentication requests rejected by the preauthentication server.	All levels
`Challenges`	Number of preauthentication requests challenged by the preauthentication server.	All levels
`Timed out requests`	Number of preauthentication requests that timed out.	All levels

Sample Output

show network-access aaa statistics preauthentication

```
user@host> show network-access aaa statistics preauthentication
Preauthentication module statistics
  Requests received: 2118
  Multistack requests: 0
  Accepts: 261
  Rejects: 975
  Challenges: 0
  Timed out requests: 882
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

RADIUS Logical Line Identification

RADIUS Logical Line Identification

show network-access aaa statistics re-authentication

IN THIS SECTION

- Syntax | 22459
- Description | 22459
- Required Privilege Level | 22459
- Output Fields | 22459
- show network-access aaa statistics re-authentication | 22460
- Release Information | 22460

Syntax

```
show network-access aaa statistics re-authentication
```

Description

Display statistics for RADIUS re-authentication, and starting in Junos OS Release 18.2R1, for local re-authentication.

Required Privilege Level

view

Output Fields

Table 876 on page 22460 lists the output fields for the show network-access aaa statistics re-authentication command. Output fields are listed in the approximate order in which they appear.

Table 876: show network-access aaa statistics re-authentication Output Fields

Field Name	Field Description
`Re-authentication statistics`	Displays re-authentication statistics.
`Requests received`	Total number of re-authentication requests that the device received from clients.
`Accepts`	Total number of accepted re-authentications.
`Challenges`	Total number of re-authentication challenges.
`Internal errors`	Total number of re-authentication internal errors.
`Rejects`	Total number of re-authentications rejected.
`Timed out requests`	Total number of re-authentication accounting timeouts.

show network-access aaa statistics re-authentication

command-name

```
user@host> show network-access aaa statistics re-authentication
Re-authentication statistics
  Requests received: 0
  Accepts: 0
  Rejects: 0
  Challenges: 0
  Timed out requests: 0
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

reauthenticate (DHCP Local Server) | 14103

show network-access aaa subscribers

IN THIS SECTION

- Syntax | 22461
- Description | 22461
- Options | 22461
- Required Privilege Level | 22462
- Output Fields | 22462
- Sample Output | 22465
- Release Information | 22468

Syntax

```
show network-access aaa subscribers
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<statistics>
<username>
<session-id session-id-number detail>
```

Description

Display subscriber-specific AAA statistics.

Options

logical-system *logical-system-name*	(Optional) List subscribers in the specific logical system.

routing-instance ***routing-instance-name***	(Optional) List subscribers for the specific routing instance. If you do not specify a routing instance name, the default routing instance is assumed.
statistics	(Optional) Display statistics for the subscriber events.
username	(Optional) Display information for the specified subscriber.
session-id ***session-id-number*** **detail**	(Optional) Display information for the specified session ID.

Required Privilege Level

view

Output Fields

Table 877 on page 22462 lists the output fields for the `show network-access aaa subscribers` command. Output fields are listed in the approximate order in which they appear.

Table 877: show network-access aaa subscribers Output Fields

Field Name	Field Description
`Challenge requests`	Number of authentication requests challenged by the authentication server for this subscriber.
`Challenge responses`	Number of challenge responses sent by the subscriber to the authentication server.
`START sent successfully`	Number of accounting start requests generated by the AAA framework for this subscriber.
`START send failures`	Number of accounting start requests that failed to make it to the accounting server for this subscriber.
`START ack received`	Number of accounting start requests acknowledged by the accounting server for this subscriber.
`INTERIM sent successfully`	Number of accounting interim requests generated by the AAA framework for this subscriber.

Table 877: show network-access aaa subscribers Output Fields *(Continued)*

Field Name	Field Description
`INTERIM send failures`	Number of accounting interim requests that failed to make it to the accounting server for this subscriber.
`INTERIM ack received`	Number of accounting interim requests acknowledged by the accounting server for this subscriber.
`Requests received`	Number of reauthentication requests received by the authentication server.
`Successful responses`	Number of successful reauthentication requests granted by the authentication server.
`Aborts handled`	Number of reauthentication requests terminated by the authentication server.
`Service name`	Name of the subscriber service.
`Creation requests`	Number of requests to create the service.
`Deletion requests`	Number of requests to delete the service.
`Request timeouts`	Number of times the service request was timed out.
`Client type`	Type of client; for example, DHCP, Mobile IP, PPP.
`Session-ID`	ID of the subscriber session.
`Session uptime`	How long the session has been up, in *HH:MM:SS.*
`Accounting`	Status of accounting, and type of accounting if accounting is on.
`Stripped username`	Username of the subscriber session.

Table 877: show network-access aaa subscribers Output Fields *(Continued)*

Field Name	Field Description
AAA Logical system/ Routing instance	AAA framework for the subscriber of logical system or routing instance.
Target Logical system/Routing instance	Target framework for the subscriber of logical system or routing instance.
Access-profile	Profile of the subscriber.
Accounting Session ID	ID of the subscriber session for accounting.
Multi Accounting Session ID	ID of the subscriber session for multiple accouting.
IP Address	IPv4 address of the subscriber.
IPv6 Address	IPv6 address of the subscriber.
IPv6 Prefix	IPv6 prefix of the subscriber.
Authentication State	State of subscriber session authentication.
Accounting State	State of subscriber session accounting.
Provisioning Type	Type of subscriber provisioning.

Sample Output

show network-access aaa subscribers logical-system

```
user@host> show network-access aaa subscribers logical-system
Username            Client type    Logical system/Routing instance
user61@example.net    ppp            default
00010e020304.1231   dhcp           isp-bos-metro-12:isp-cmbrg-12
user54@example.com     dhcp           default:isp-gtown-r3-00
0020df980102.2334   dhcp           isp-bos-metro-16:isp-cmbrg-12
```

show network-access aaa subscribers logical-system routing-instance

```
user@host> show network-access aaa subscribers logical-system isp-bos-metro-16 routing-instance
isp-cmbrg-12-32
Username            Client type   Logical system/Routing instance
00010e020304.1231   dhcp          isp-bos-metro-12:isp-cmbrg-12
user54@example.com     dhcp          default:isp-gtown-r3-00
0020df980102.2334   dhcp          isp-bos-metro-16:isp-cmbrg-12
```

show network-access aaa subscribers statistics username

```
user@host> show network-access aaa subscribers statistics username 00010e020304.1231
Authentication statistics
   Challenge requests: 0
   Challenge responses: 0
 Accounting statistics
   START sent successfully: 1
   START send failures: 0
   START ack received: 1
   INTERIM sent successfully: 0
   INTERIM send failures: 0
   INTERIM ack received: 0
 Re-authentication statistics
   Requests received: 0
   Sucessfull responses: 0
   Aborts handled: 0
 Service statistics
```

```
    Service name: filter-serv
    Creation requests: 1
    Deletion requests: 0
    Request timeouts: 0
    Service name: filter-serv2
    Creation requests: 144
    Deletion requests: 0
    Request timeouts: 144
```

show network-access aaa subscribers username

```
user@host> show network-access aaa subscribers username user80@example.net
Logical system/Routing instance   Client type  Session-ID  Session uptime  Accounting
isp-bos-metro-16:isp-cmbrg-12     dhcp         7           01:12:56        on/volume
Service name       Service type      Quota          Accounting
I-Cast             volume            1200 Mbps      on/volume+time
Voip                                                on/volume
GamingBurst        time              6000 secs      on/volume
```

show network-access aaa subscribers session-id 26 detail

The following command output is seen when only an IPv4 client is associated with the session:

```
user@host> show network-access aaa subscribers session-id 26 detail
Type: dhcp
Stripped username: my-customer
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.0.0.2
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
```

The following command output is seen when IPv6 client logs in (after IPv4 association) and is associated with the same session ID:

```
user@host> show network-access aaa subscribers session-id 26 detail
Type: dhcp
Stripped username: my-customer
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.0.0.2
IPv6 Address: 2001:db8:0:8003::2
IPv6 Prefix: 2001:db8:ffff:0:4::/64
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
```

show network-access aaa subscribers (Tenant systems)

```
user@host:TSYS1> show network-access aaa subscribers
Username                 Logical system/Routing instance   Client type    Session-ID
userX                    default:TSYS1-ri                    xauth          1
```

show network-access aaa subscribers (Tenant systems)

```
user@host> show network-access aaa subscribers session-id 1 detail
Type: xauth
Username: userX
Stripped username: userX
AAA Logical system/Routing instance: default:TSYS1-ri
Target Logical system/Routing instance: default:TSYS1-ri
Access-profile: ap1
+    Tenant: TSYS1
Session ID: 1
Multi Accounting Session ID: 0
IP Address: 192.0.2.0
Authentication State: AuthStateActive
```

```
Accounting State: Acc-Init
Converted to time accounting: no
Provisioning Type: None

user@host> show network-access aaa subscribers session-id 2 detail
Type: xauth
Username: userY
Stripped username: userY
AAA Logical system/Routing instance: default:TSYS2-ri
Target Logical system/Routing instance: default:TSYS2-ri
Access-profile: ap1
+    Tenant: TSYS2
Session ID: 2
Multi Accounting Session ID: 0
IP Address: 192.0.2.1
Authentication State: AuthStateActive
Accounting State: Acc-Init
Converted to time accounting: no
Provisioning Type: None
```

Release Information

Command introduced in Junos OS Release 9.1.

Command updated with `session-id` *`session-id-number`* `detail` in Junos OS Release 17.3.

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information

show network-access aaa subscribers session-id

IN THIS SECTION

- Syntax | 22469
- Description | 22469

- Options | 22469
- Required Privilege Level | 22469
- Output Fields | 22469
- Sample Output | 22475
- Release Information | 22479

Syntax

```
show network-access aaa subscribers session-id session-id
<brief | detail>
```

Description

Display information about the specified subscriber session.

Options

session-id	ID of the subscriber session.
brief | detail	(Optional) Display the specified level of information.

Required Privilege Level

view

Output Fields

Table 878 on page 22470 lists the output fields for the `show network-access aaa subscribers session-id` command. Output fields are listed in the approximate order in which they appear.

Table 878: show network-access aaa subscribers session-id Output Fields

Field Name	Field Description	Level of Output
`Type` and `Client type`	Type of client.	All levels
`Accounting`	Status of the accounting configuration for the service, `on` or `off`, and the type of accounting, `time`, `volume+time`, or `flat-file`. The `time` and `volume+time` types are configured in RADIUS Service-Statistics VSA [26-69].	brief none
`Service type`	Type of accounting: `volume`, `time`, `volume+time`, or `na`.	brief
`Quota`	Quota for service: volume (in Mbps) or time (seconds).	brief
`Username`	Name of the user logged in to the session.	detail
`Stripped username`	Username after the domain has been removed.	detail
`Logical system/ Routing instance` and `AAA Logical system/ Routing instance`	Name of the routing instance, logical system name, or both used for the session.	All levels
`Target Logical system/Routing instance`	Logical system/routing instance to which the session is mapped.	detail
`Access-profile`	Access profile used for AAA services for the session.	detail
`Session ID`	ID of the subscriber session. The session ID value displayed under `Service name` is the service session ID.	detail

Table 878: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting Session ID`	ID of the accounting session (RADIUS attribute 44). The ID appears in decimal or description format, as specified by the `accounting-session-id-format` statement.	detail
`Multi Accounting Session ID`	Bundle ID for MLPPP sessions. Acct-Multi-Session-Id (RADIUS attribute 50) uses the value of the session database bundle session ID to enable RADIUS to link together multiple related sessions. The value of this field is zero when no MLPPP sessions exist.	detail
`IP Address`	IP address of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
`IPv6 Address`	IPv6 address of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
`IPv6 Prefix`	IPv6 prefix of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
`Authentication State`	State of the subscriber authentication session: `AuthInit`, `AuthStart`, `AuthChallenge`, `AuthRedirect`, `AuthClntRespWait`, `AuthAcctVolStatsAckWait`, `AuthAcctStopAckWait`, `AuthServCreateRespWait`, `AuthLogoutStart`, `AuthStateActive`, `AuthClntLogoutRespWait`, `AuthProfileUpdateWait`, `AuthProvisionRespWait`, `AuthProvisionServiceCreationWait`	detail
`Gx-Plus Provisioning State`	State of Gx-Plus provisioning: • ignored—Subscriber has no IPv4 address or NAS-Port-ID. • in-progress—Provisioning is in progress. • logout—Subscriber logout is in progress. • logout-done—Logout response has been received. • response-received—Provisioning response has been received.	detail

Table 878: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Pcrf Provisioning State`	State of PCRF provisioning: • `active`—PCRF provisioning is active. • `ignored`—Subscriber has no IPv4 address or NAS-Port-ID. • `in-progress`—Provisioning is in progress. • `logout`—Subscriber logout is in progress. • `logout-done`—Logout response has been received. • `response-received`—Provisioning response has been received.	detail
`Pcrf Subscription-Id-Type`	Type of subscriber for a PCRF partition. You can define your own or use a predefined value: `0 (END_USER_E164), 1 (END_USER_IMSI), 2 (END_USER_SIP_URI), 3 (END_USER_NAI), 4 (END_USER_PRIVATE)`.	detail
`Pcrf Subscription-Id-Data`	Subscriber data string concatenated from a list of user-selected data options used to identify the subscriber type for a PCRF partition; for example, demux0	detail
`Ocs Subscription-Id-Type`	Type of subscriber for an OCS partition. You can define your own or use a predefined value: `0 (END_USER_E164), 1 (END_USER_IMSI), 2 (END_USER_SIP_URI), 3 (END_USER_NAI), 4 (END_USER_PRIVATE)`.	detail
`Ocs Subscription-Id-Data`	Subscriber data string concatenated from a list of user-selected data options used to identify the subscriber type for an OCS partition; for example: test-sid	detail
`Ocs Interrogation State`	State of the OCS interrogation: `first, intermediate, final`.	detail
`Ocs Data State`	State of the OCS data: `none`	detail

Table 878: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting State`	State of the subscriber accounting session: `Acc-Init`, `Acc-Start-Sent`, `Imm-Update-Stats-Pending`, `Acc-Interim-Sent`, `Acc-Stop-Stats-Pending`, `Acc-Stop-Sent`, `Acc-Stop-On-Fail-Deny-Sent`, `Acc-Stop-Ackd`	detail
`Provisioning-type`	Provisioning type for this session: • gx-plus—Subscriber service uses Gx-Plus provisioning. • jsrc—Subscriber service uses JSRC provisioning. • none—Provisioning is not enabled.	detail
`Service name`	Name of the attached service or policy. • For RADIUS-activated and CLI-activated services, displays the full activation string for the service. If the activation string includes service parameters, then both the service name and service parameters are displayed. • For JSRC-activated policies—displays the policy name.	All levels
`Service State`	State of the service provided in the subscriber session.	detail
`Service Family`	Network family of the service provided in the subscriber session.	detail
`Service Activation Source`	Source used to activate the service.	detail
`Session uptime`	How long the session has been up, in *HH:MM:SS*.	All levels
`Service CC-Service-Identifier`	Data identification element of the 3GPP Diameter credit control service charging system that uniquely defines the `CC-Service-Context`.	detail

Table 878: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Service Rating-Group`	Value associated with a charging rule and part of the accounting data stream for the PCRF.	detail
`Ocs Control`	Whether OCS controls the service: `yes` or `no`.	detail
`Accounting status`	Status of the accounting configuration for the service, `on` or `off`, and the type of accounting, `time`, `volume+time`, or `flat-file`. The `time` and `volume+time` types are configured in RADIUS Service-Statistics VSA [26-69].	detail
`Service accounting session ID`	ID of the service accounting session; RADIUS Acct-Session-Id attribute (44). The ID appears in decimal or description format, as specified by the `accounting-session-id-format` statement.	detail
`Service accounting state`	State of the service accounting session: `Acc-Init`, `Acc-Start-Sent`, `Imm-Update-Stats-Pending`, `Acc-Interim-Sent`, `Acc-Stop-Stats-Pending`, `Acc-Stop-Sent`, `Acc-Stop-On-Fail-Deny-Sent`, `Acc-Stop-Ackd`	detail
`Accounting interim interval`	Amount of time between interim accounting updates for this service, in seconds; RADIUS Service-Interim-Acct-Interval VSA [26-140] or Diameter Acct-Interim-Interval AVP (85).	detail
`Pcrf session-stamp`	Value that consists of the UTC time when the router creates the CCR-GX-I. The router appends the session stamp to the session ID when a CCR packet is sent to the PCRF. You configure the `use-session-stamp` option at the `[edit access pcrf partition` *`partition-name`* hierarchy level. If you do not configure this option, the field displays a value of zero. For local reinitialization, you must configure the `use-session-stamp` option.	detail

Sample Output

show network-access aaa subscribers session-id brief

```
user@host> show network-access aaa subscribers session-id 6 brief
Logical system/Routing instance   Client type   Session uptime    Accounting
default:default                   dhcp          00:01:29          on/time

Service name               Service type   Quota           Accounting
filter-service             -na-           -na-            off
filter-service-2           volume+time    77.00MB/120secs off
1337994190863204450        -na-           -na-            off
```

show network-access aaa subscribers session-id 2 (Flat File Accounting)

```
user@host> show network-access aaa subscribers session-id 2
Logical system/Routing instance   Client type    Session-ID     Session uptime    Accounting
default:default                   dhcp           2              00:00:48          on/volume+time

Service name               Service type   Quota           Accounting
filter-service             -na-           -na-            on/flat-file
```

show network-access aaa subscribers session-id detail

```
user@host> show network-access aaa subscribers session-id 5 detail
Type: dhcp
Username: user23@example.net
Stripped username: user23
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:retail-onlinecompany-ca
Access-profile:retailer-onlinecompany-sjc
Session ID: 5
Accounting Session ID: jnpr ge-1/0/0.101:1
Multi Accounting Session ID: 0
IP Address: 192.168.44.104
Authentication State: AuthStateActive
Pcrf session-stamp: 0
Pcrf Provisioning State: active
```

```
Pcrf Subscription-Id-Type: 4
Pcrf Subscription-Id-Data: demux0
Ocs Subscription-Id-Type: 15
Ocs Subscription-Id-Data: test-sid
Ocs Interrogation State: intermediate
Ocs Data State: none
Gx-Plus Provisioning State: response-received
Accounting State: Acc-Interim-Sent
Provisioning-type: jsrc
Service name: filter-service-1
 Service State: SvcActive
 Service Family: inet
 Service Activation Source: PCRF-LOGIN
 Session ID: 7
 Session uptime: 00:01:33
 Service CC-Service-Identifier: 777
 Service Rating Group: 10
 Ocs Control: yes
Service name: filter-service-2
 Service State: SvcActive
 Service Family: inet
 Service Activation Source: PCRF-LOGIN Session ID: 8
 Session uptime: 00:01:33
 Service CC-Service-Identifier: 778
 Service Rating Group: 11
 Ocs Control: no
Accounting status: on/volume+time
 Service accounting session ID: 1:2-1322506006
 Service accounting state: Acc-Interim-Sent
 Accounting interim interval: 600
```

show network-access aaa subscribers session-id detail (Service with Multiple Instances)

```
user@host> show network-access aaa subscribers session-id 6 detail
Type: dhcp
Stripped username: user-test-fms2
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: attr_test_profile1
Session ID: 6
Accounting Session ID: 6
```

```
Multi Accounting Session ID: 0
IP Address: 198.51.100.10
Authentication State: AuthStateActive
Pcrf session-stamp: 0
Accounting State: Acc-Interim-Sent
Provisioning Type: None
Service name: economy-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius at Reauth
  Session ID: 7
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:7-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
Service name: economy-service(upstrm-filter,dwnstrm-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 8
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:8-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
```

show network-access aaa subscribers session-id detail (Single Session Dual Stack with active V4 and V6 subscribers)

```
user@host> show network-access aaa subscribers session-id 26 detail
Type: dhcp
Stripped username: user-test-25
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.10.0.6
IPv6 Address: 00:00:5E:00:53:02
```

```
IPv6 Prefix: 00:00:5E:00:53:00/64
Authentication State: AuthStateActive
Pcrf session-stamp: 0
Accounting State: Acc-Interim-Sent
Provisioning Type: None
```

show network-access aaa subscribers session-id detail (PCRF with Session-Stamp)

```
user@host> show network-access aaa subscribers session-id 23 detail
Type: dhcp
Username: user45-28-abc@example.net
Stripped username: user45-28-abc
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: test-profile-abc-user
Session ID: 23
Accounting Session ID: 23
Multi Accounting Session ID: 0
IP Address: 198.51.100.5
Authentication State: AuthStateActive
Pcrf session-stamp: 1557788595
Pcrf Provisioning State: active
Pcrf Subscription-Id-Type: 4
Pcrf Subscription-Id-Data: user45-28-abc@example.net
Ocs Subscription-Id-Type: 15
Ocs Subscription-Id-Data: 987654321
Ocs Interrogation State: idle
Ocs Data State: none
Accounting State: Acc-Start-Sent
Provisioning Type: PCRF-LOGIN
Service name: test-filters(filter-123-xyz-in,filter-123-xyz-out)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: PCRF-LOGIN
  Session ID: 24
  Session uptime: 00:00:39
  Service CC-Service-Identifier: 12345
  Service Rating Group: 3300
  Ocs Control: yes
  Service session type: Service-Profile
Service name: test-cos(abc-video-100M)
```

```
    Service State: SvcActive
    Service Activation Source: PCRF-LOGIN
    Session ID: 25
    Session uptime: 00:00:39
    Service session type: Service-Profile
  Service name: test-multi
    Service State: SvcActive
    Service Family: inet
    Service Activation Source: PCRF-LOGIN
    Session ID: 26
    Session uptime: 00:00:39
    Service session type: Service-Profile
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- *Verifying and Managing Subscriber AAA Information*
- *Local and Remote Service Activation and Deactivation Using the CLI*
- *Deactivating a Single Instance of a Subscriber Service*
- *Deactivating All Instances of a Subscriber Service*
- *Verifying Subscriber Services with Multiple Instances*
- *Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers*
- *Understanding Gy Interactions Between the Router and the OCS*
- *Understanding Interactions Between the PCRF, PCEF, and OCS*

show network-access aaa terminate-code

IN THIS SECTION

- Syntax | 22480
- Description | 22480

- Options | 22480
- Required Privilege Level | 22481
- Output Fields | 22481
- Sample Output | 22482
- Release Information | 22485

Syntax

```
show network-access aaa terminate-code
<brief | detail | summary>
<reverse>
<(aaa | dhcp | l2tp | ppp)>
```

Description

Display the count for termination cause types and the current mapping between session termination cause types and code values.

Options

none	Display all mappings.
brief | detail | summary	(Optional) Display the specified level of output. The `summary` output is displayed by default and includes base count information about mappings. The `brief` output displays mappings with non-zero usage count and custom mappings. The `detail` output displays all mappings.
aaa	(Optional) Limit display to AAA mappings only.
dhcp	(Optional) Limit display to DHCP mappings only.
l2tp	(Optional) Limit display to L2TP mappings only.
ppp	(Optional) Limit display to PPP mappings only.
reverse	(Optional) Display mapping of the code value conveyed in the RADIUS Acct-Terminate-Cause attribute (49) to the termination cause type.

vlan (Optional) Limit display to VLAN mappings only.

Required Privilege Level

view

Output Fields

Table 879 on page 22481 lists the output fields for the `show network-access aaa terminate-code` command. Output fields are listed in the approximate order in which they appear.

Table 879: show network-access aaa terminate-code Output Fields

Field Name	Field Description	Level of Output
RADIUS	RFC-defined code value conveyed in the RADIUS Acct-Terminate-Cause attribute (49) or a nonstandard, customized value that you configure with the `terminate-code aaa` statement at the `[edit access]` hierarchy level.	`brief` `detail` None (with `reverse` option)
Custom	Whether or not the termination cause is a customized mapping or the default mapping.	All levels
Mapping-Count	Number of mappings that occurred for a specific terminate cause type or category (standard or summary output) or per termination cause (reverse output).	`summary` None
Usage-Count	Number of times the terminate code mapping was used.	All levels
Type	Termination cause type—null, aaa, dhcp, l2tp, ppp, or vlan. **NOTE**: The null termination cause type indicates that no termination reason was provided by the subscriber and the RADIUS Acct-Terminate-Cause attribute (49) was not included in the Acct-Stop request	All levels

Table 879: show network-access aaa terminate-code Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Code	Specific termination cause.	`brief detail`

Sample Output

show network-access aaa terminate-code

```
user@host> show network-access aaa terminate-code
Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     1             0           null
  no     12            0           aaa
  no     5             0           dhcp
  no     364           0           l2tp
  no     210           0           ppp
  no     13            10          vlan
```

show network-access aaa terminate-code reverse

```
user@host> show network-access aaa terminate-code reverse
Terminate-code:
  RADIUS    Custom Mapping-Count Usage-Count Type
  0         no     1             0           null
  1         no     1             0           aaa
  1         no     1             0           dhcp
  1         no     5             0           l2tp
  1         no     8             0           ppp
  1         no     2             10          vlan
  2         no     1             0           dhcp
  2         no     3             0           ppp
  2         no     2             0           vlan
  4         no     1             0           aaa
  4         no     1             0           dhcp
  4         no     1             0           l2tp
  4         no     1             0           ppp
```

```
  5         no     2             0           aaa
  5         no     1             0           l2tp
  5         no     1             0           ppp
  6         no     2             0           aaa
  6         no     13            0           l2tp
  6         no     3             0           ppp
  6         no     3             0           vlan
  8         no     3             0           l2tp
  8         no     5             0           ppp
  9         no     13            0           l2tp
  9         no     12            0           ppp
  9         no     4             0           vlan
  10        no     4             0           aaa
  10        no     1             0           dhcp
  10        no     128           0           l2tp
  10        no     171           0           ppp
  15        no     1             0           dhcp
  15        no     190           0           l2tp
  15        no     1             0           vlan
  16        no     1             0           vlan
  17        no     2             0           aaa
  17        no     10            0           l2tp
  17        no     6             0           ppp
```

show network-access aaa terminate-code dhcp

```
user@host> show network-access aaa terminate-code dhcp
Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     5             0           dhcp
```

show network-access aaa terminate-code detail

```
user@host> show network-access aaa terminate-code aaa detail

Terminate-code:
  RADIUS     Custom Usage-Count Type Code
  17         no     1           aaa  deny-authentication-denied
  10         no     1           aaa  deny-no-resources
  17         no     0           aaa  deny-server-request-timeout
```

```
  6          no     0           aaa  service-shutdown-network-logout
  10         no     0           aaa  service-shutdown-remote-reset
  1200       yes    5           aaa  service-shutdown-subscriber-logout
  5          no     0           aaa  service-shutdown-time-limit
  10         no     0           aaa  service-shutdown-volume-limit
  6          no     13          aaa  shutdown-administrative-reset
  4          no     0           aaa  shutdown-idle-timeout
  10         no     0           aaa  shutdown-remote-reset
  5          no     0           aaa  shutdown-session-timeout
```

show network-access aaa terminate-code brief

```
user@host> show network-access aaa terminate-code brief
Terminate-code:
  RADIUS     Custom Usage-Count Type Code
  17         no     1           aaa  deny-authentication-denied
  10         no     1           aaa  deny-no-resources
  1200       yes    5           aaa  service-shutdown-subscriber-logout
  6          no     13          aaa  shutdown-administrative-reset
  15         no     7           dhcp nak
  10         no     1           l2tp session-receive-cdn-avp-missing-secret
  10         no     1           ppp  bundle-fail-create
  1          no     1           ppp  lcp-peer-terminate-term-req
  10         no     1           ppp  lcp-tunnel-disconnected
  1          no     10          vlan out-of-band-ancp-port-down
```

show network-access aaa terminate-code summary

```
user@host> show network-access aaa terminate-code summary

Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     1             0           null
  no     12            0           aaa
  no     5             0           dhcp
  no     364           0           l2tp
  no     210           0           ppp
  no     13            10          vlan
```

show network-access aaa terminate-code vlan

```
user@host> show network-access aaa terminate-code vlan
Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     13            0           vlan
```

show network-access aaa terminate-code vlan detail

```
user@host> show network-access aaa terminate-code vlan detail
Terminate-code:
  RADIUS     Custom Usage-Count Type Code
  6          no     0           vlan admin-logout
  16         no     0           vlan admin-reconnect
  9          no     0           vlan other
  2          no     0           vlan out-of-band-access-interface-down
  6          no     0           vlan out-of-band-admin-access-interface-down
  6          no     0           vlan out-of-band-admin-core-interface-down
  1          no     0           vlan out-of-band-ancp-port-down
  1          no     0           vlan out-of-band-ancp-port-vlan-id-change
  2          no     0           vlan out-of-band-core-interface-down
  15         no     0           vlan out-of-band-l2-wholesale-no-free-vlans
  9          no     0           vlan profile-request-error
  9          no     0           vlan sdb-error
  9          no     0           vlan subscriber-activate-error
```

Release Information

Command introduced in Junos OS Release 11.4.

`vlan` option introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Session Termination Causes and RADIUS Termination Cause Codes

Session Termination Causes and RADIUS Termination Cause Codes

AAA Termination Causes and Code Values

DHCP Termination Causes and Code Values

L2TP Termination Causes and Code Values

PPP Termination Causes and Code Values

VLAN Termination Causes and Code Values

show network-access address-assignment pool

IN THIS SECTION

- Syntax | 22486
- Description | 22486
- Options | 22486
- Required Privilege Level | 22487
- Output Fields | 22487
- Sample Output | 22488
- Release Information | 22488

Syntax

```
show network-access address-assignment pool pool-name
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display state information for each address-assignment pool.

Options

none	Display information about clients that have obtained addresses from the address-assignment pool.
pool *pool-name*	Display information about the specified address-assignment pool.

logical-system ***logical-system-name***	(Optional) Perform this operation on the specified logical system.
routing-instance ***routing-instance-name***	(Optional) Perform this operation on the specified routing instance.

Required Privilege Level

view and system

Output Fields

Table 880 on page 22487 lists the output fields for the `show network-access address-assignment pool` command. Output fields are listed in the approximate order in which they appear.

Table 880: show network-access address-assignment pool Output Fields

Field Name	Field Description
`IP address/prefix`	IP address of the client.
`Hardware address`	MAC address of the client. Displays `NA` for addresses excluded from being allocated from the pool with the `excluded-address` or `excluded-range` statements.
Host/User	Hostname or username of the client. Displays `EXCLUDED` for addresses excluded from being allocated from the pool with the `excluded-address` or `excluded-range` statements.
`Type`	Type of client. Displays `unknown` for addresses excluded from being allocated from the pool with the `excluded-address` or `excluded-range` statements.

Sample Output

show network-access address-assignment pool

```
user@host> show network-access address-assignment pool sunnywest logical-system ls1 routing-
instance routinst2
IP address/prefix       Hardware address      Host/User       Type
192.168.2.1             00:00:5e:00:53:01      user1          DHCP
192.168.2.2             00:00:5e:00:53:02      user2          DHCP
192.168.2.3             00:00:5e:00:53:03      user3          DHCP
192.168.2.4             NA                     EXCLUDED       unknown
```

Release Information

Command introduced in Junos OS Release 9.0.

show network-access address-assignment pool (View)

IN THIS SECTION

- Syntax | 22488
- Description | 22489
- Required Privilege Level | 22489
- Output Fields | 22489
- Sample Output | 22489
- Release Information | 22490

Syntax

```
show network-access address-assignment pool name
```

Description

Display information summary about a specific pool.

Required Privilege Level

view

Output Fields

Table 881 on page 22489 lists the output fields for the `show network-access address-assignment pool` command. Output fields are listed in the approximate order in which they appear.

Table 881: show network-access address-assignment pool Output Fields

Field Name	Field Description
`IP address`	IP address assigned to a client.
`Hardware address`	MAC address of the client. For XAuth clients, the value is NA.
`Host/User`	For static IP address assignment, the user name and profile are displayed in the format *username@profile*. If the client is assigned an IP address from an address pool and a user name exists, the user name is displayed. For DHCP applications, if the host name is configured the host name is displayed; otherwise NA is displayed.
`Type`	Either XAuth or DHCP attributes are configured.

Sample Output

command-name

```
user@host> show network-access address-assignment pool xauth1
IP address       Hardware address        Host/User                 Type
192.0.2.1         NA                      jason@dvpn-auth           XAUTH
192.0.2.2         NA                      jacky                     XAUTH
```

```
192.0.2.3          00:00:5E:00:53:01       host1                      DHCP
192.0.2.4          00:00:5E:00:53:02       NA                         DHCP
```

Release Information

Command introduced in Junos OS Release 10.4.

show network-access domain-map

IN THIS SECTION

- Syntax | 22490
- Description | 22490
- Options | 22490
- Required Privilege Level | 22491
- Output Fields | 22491
- Sample Output | 22491
- Release Information | 22492

Syntax

```
show network-access domain-map
<statistics>
```

Description

Display domain map information.

Options

statistics (Optional) Display domain map statistics.

Required Privilege Level

view

Output Fields

Table 882 on page 22491 lists the output fields for the `show network-access domain-map statistics` command. Output fields are listed in the approximate order in which they appear.

Table 882: show network-access domain-map Output Fields

Field Name	Field Description
`Matched domains`	Number of usernames with domain names that are matched.
`Unmatched domains`	Number of usernames with domain names that are not matched.
`Missing domain names`	Number of usernames without a domain name.
`Stripped username`	Number of usernames from which the domain name has been stripped.
`Default used`	Number of times the default domain map is used.

Sample Output

show network-access domain-map statistics

```
user@host> show network-access domain-map statistics
General domain mapping statistics
  Matched domains: 7
  Unmatched domains: 1
  Missing domain names: 0
  Stripped username: 7
Domain statistics for domain-name: default
  Default used: 1
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Verifying Domain Maps

show network-access gx-plus

IN THIS SECTION

- Syntax | 22492
- Description | 22492
- Options | 22492
- Required Privilege Level | 22493
- Output Fields | 22493
- Sample Output | 22495
- Release Information | 22509

Syntax

```
show network-access gx-plus
<state | statistics | sync-state>
```

Description

Display Gx-Plus provisioning state, synchronization state, and statistics information.

Options

state (Optional) Display Gx-Plus provisioning state.

statistics (Optional) Display Gx-Plus statistics.

sync-state (Optional) Display Gx-Plus synchronization state.

Required Privilege Level

view

Output Fields

Table 883 on page 22493 lists the output fields for the `show network-access gx-plus` command. Output fields are listed in the approximate order in which they appear.

Table 883: show network-access gx-plus Output Fields

Field Name	Field Description
`Gx-plus state`	State of the Gx-Plus application, including the following: • `Engine created` • `Partition configured` • `Configuration active` • `Diameter interface configured` • `Total number of entries`—Number of provisioned, pending, and terminating subscribers. • `Number of pending entries`—Number of pending subscribers. • `Number of pending logouts`—Number of subscribers logging out.
`Sync-Event`	Type of synchronization event.
`Timeout`	Number of times notification sent without response.
`Gx-plus general counters`	Number and state of general events.

Table 883: show network-access gx-plus Output Fields *(Continued)*

Field Name	Field Description
`Gx-plus sync-event counters`	Number and state of synchronization events.
`Gx-plus sync-event diameter-events`	Events or errors that occur in jdiameterd on behalf of Gx-Plus.
`Gx-plus upstream diameter-events`	Errors encountered when the Gx-Plus component in authd is processing outbound requests.
`Gx-plus upstream counters`	Events or errors that occur with the Gx-Plus component of authd.
`Gx-plus downstream counters`	Events for requests that originated from jdiameterd and are intended for the Diameter layer on a PCRF server.
`cold-boot`	Event that takes place after a cold boot (system reboot).
`warm-boot`	Event that takes place after an authd or jdiameterd process restart.
`d-req`	Event that takes place after a discovery request from the router.
`ayt`	Event that takes place after an “Are You There” check. The check is triggered when the connection breaks between Diameter peers. It is the same as an AWD, except it is sent from the router side rather than the server side.
`awd`	Event that takes place after an application watchdog flag. The AWD is triggered when the connections breaks between Diameter peers. It is the same as an AYT, except it is sent from the server side rather than the router side.
`prov`	Events that take place as a result of the CCR-I provisioning request.
`logout`	Events that take place as a result of the CCR-T logout request.

Table 883: show network-access gx-plus Output Fields *(Continued)*

Field Name	Field Description
ruleReport	Events that take place as a result of the CCR-U rule report request.
threshold	Events that take place as a result of the CCR-U threshold request.
general	Events that cannot be attributed to a specific type of request.
audit	Events that take place as a result requests for an audit of a specific subscriber session; includes information about state and services.
d-exact	Events that take place as a result of discovery exchange for Gx-Plus. The request originates from the PCRF server for a specific subscriber.
d-bulk	Events that take place as a result of discovery requests for all active Gx-Plus subscribers. The request originates on the PCRF server.
d-done	Events that take place when the discovery exchange is completed.

Sample Output

show network-access gx-plus state

```
user@host> show network-access gx-plus state
Gx-plus state:
  Engine created                : yes
  Partition configured          : yes
  Configuration active          : yes
  Diameter interface configured : yes
  Total number of entries       : 0
  Number of pending entries     : 0
  Number of pending logouts     : 0
```

show network-access gx-plus statistics

```
user@host> show network-access gx-plus statistics
Gx-plus general counters:
  Counter                                                   Value
  engine created                                            1
  initial config: active                                    1
  recovery: process restart                                 1
  diameter-app initial config: success                      1

Gx-plus sync-event counters:
  Category      Counter                                     Value
  warm-boot     activated                                   1
  warm-boot     posted                                      1
  warm-boot     response                                    1
  awd           posted                                      12
  awd           response                                    12
```

show network-access gx-plus sync-state

```
user@host> show network-access gx-plus sync-state
Gx-plus sync-events:
  Sync-Event                Timeout
  cold-boot                 6100
```

show network-access gx-plus statistics extensive (Extensive)

```
user@host> show network-access gx-plus statistics extensive

Gx-plus general counters:
  Counter                                                   Value
  engine created                                            0
  engine created after recovery                             0
  engine destruction started                                0
  engine destroyed                                          0
  diameter-app config: failed                               0
  initial config: active                                    0
  initial config: inactive                                  0
  config changed: inactive to active                        0
```

```
  config changed: active to inactive                        0
  config changed: active to active                          0
  config changed: inactive to inactive                      0
  recovery: cold-boot                                       0
  recovery: re-failover                                     0
  recovery: process restart                                 0
  diameter-app initial config: failure                      0
  diameter-app initial config: success                      0
  diameter-app config retry: failure                        0
  diameter-app config retry: success                        0
  response when no active config                            0
  bad response                                              0
  response: unknown session                                 0
  request when on no origins                                0
  request when no config: response queue full               0
  request when no config: no memory                         0
  request when no config: bad request                       0
  request when no config: response build failure            0
  request when no config: response sent                     0
  request when no config: response post failure             0
  bad request                                               0
  diameter tx                                               124986
  diameter rx                                               7231
  diameter event                                            1253
  retrieving pending provisioning request                   0
  retrieving pending rule-report                            0
  retrieving pending rule-report                            0
  missed rule report request                                0
  tick                                                      1213965

Gx-plus sync-event counters:
  Sync-Event    Counter                                     Value
  cold-boot     activated                                   0
  cold-boot     re-activated                                0
  cold-boot     duplicate activation                        0
  cold-boot     activated to queue                          0
  cold-boot     message build failure                       0
  cold-boot     posted                                      0
  cold-boot     post failure                                0
  cold-boot     response                                    0
  warm-boot     activated                                   0
  warm-boot     re-activated                                0
  warm-boot     duplicate activation                        0
```

```
  warm-boot     activated to queue                          0
  warm-boot     message build failure                       0
  warm-boot     posted                                      0
  warm-boot     post failure                                0
  warm-boot     response                                    0
  d-req         activated                                   0
  d-req         re-activated                                0
  d-req         duplicate activation                        0
  d-req         activated to queue                          0
  d-req         message build failure                       0
  d-req         posted                                      0
  d-req         post failure                                0
  d-req         response                                    0
  ayt           activated                                   5
  ayt           re-activated                                0
  ayt           duplicate activation                        37
  ayt           activated to queue                          0
  ayt           message build failure                       0
  ayt           posted                                      0
  ayt           post failure                                0
  ayt           response                                    0
  awd           activated                                   0
  awd           re-activated                                0
  awd           duplicate activation                        0
  awd           activated to queue                          0
  awd           message build failure                       0
  awd           posted                                      0
  awd           post failure                                0
  awd           response                                    0

Gx-plus sync-event diameter-events:
  Category      Counter                                     Value
  cold-boot     bad flags                                   0
  cold-boot     bad fixed destination                       0
  cold-boot     bad routed destination                      0
  cold-boot     tx is over limit                            0
  cold-boot     bad end-to-end id                           0
  cold-boot     no peer for tx                              0
  cold-boot     peer down while waiting for answer          0
  cold-boot     timeout while waiting for answer            0
  cold-boot     tx timeout                                  0
  cold-boot     tx try limit                                0
  cold-boot     tx failure                                  0
```

```
cold-boot     discarded                                  0
cold-boot     received answer is over limit              0
cold-boot     tx failure: no memory                      0
cold-boot     base-app-tx-timeout                        0
cold-boot     base-app-rx-timeout                        0
cold-boot     base-app-tx-discard                        0
cold-boot     base-app-rx-discard                        0
warm-boot     bad flags                                  0
warm-boot     bad fixed destination                      0
warm-boot     bad routed destination                     0
warm-boot     tx is over limit                           0
warm-boot     bad end-to-end id                          0
warm-boot     no peer for tx                             0
warm-boot     peer down while waiting for answer         0
warm-boot     timeout while waiting for answer           0
warm-boot     tx timeout                                 0
warm-boot     tx try limit                               0
warm-boot     tx failure                                 0
warm-boot     discarded                                  0
warm-boot     received answer is over limit              0
warm-boot     tx failure: no memory                      0
warm-boot     base-app-tx-timeout                        0
warm-boot     base-app-rx-timeout                        0
warm-boot     base-app-tx-discard                        0
warm-boot     base-app-rx-discard                        0
d-req         bad flags                                  0
d-req         bad fixed destination                      0
d-req         bad routed destination                     0
d-req         tx is over limit                           0
d-req         bad end-to-end id                          0
d-req         no peer for tx                             0
d-req         peer down while waiting for answer         0
d-req         timeout while waiting for answer           0
d-req         tx timeout                                 0
d-req         tx try limit                               0
d-req         tx failure                                 0
d-req         discarded                                  0
d-req         received answer is over limit              0
d-req         tx failure: no memory                      0
d-req         base-app-tx-timeout                        0
d-req         base-app-rx-timeout                        0
d-req         base-app-tx-discard                        0
d-req         base-app-rx-discard                        0
```

```
  ayt           bad flags                                0
  ayt           bad fixed destination                    0
  ayt           bad routed destination                   0
  ayt           tx is over limit                         0
  ayt           bad end-to-end id                        0
  ayt           no peer for tx                           1248
  ayt           peer down while waiting for answer       0
  ayt           timeout while waiting for answer         0
  ayt           tx timeout                               0
  ayt           tx try limit                             0
  ayt           tx failure                               0
  ayt           discarded                                0
  ayt           received answer is over limit            0
  ayt           tx failure: no memory                    0
  ayt           base-app-tx-timeout                      0
  ayt           base-app-rx-timeout                      0
  ayt           base-app-tx-discard                      0
  ayt           base-app-rx-discard                      2
  awd           bad flags                                0
  awd           bad fixed destination                    0
  awd           bad routed destination                   0
  awd           tx is over limit                         0
  awd           bad end-to-end id                        0
  awd           no peer for tx                           42
  awd           peer down while waiting for answer       0
  awd           timeout while waiting for answer         0
  awd           tx timeout                               0
  awd           tx try limit                             0
  awd           tx failure                               0
  awd           discarded                                0
  awd           received answer is over limit            0
  awd           tx failure: no memory                    0
  awd           base-app-tx-timeout                      0
  awd           base-app-rx-timeout                      0
  awd           base-app-tx-discard                      0
  awd           base-app-rx-discard                      0

Gx-plus upstream counters:
  Category      Counter                                  Value
  prov          start                                    6
  prov          start: failure: bad data                 0
  prov          start: ignore                            0
  prov          start: failure: in recovery              0
```

```
prov          start: failure: spruious                    0
prov          start: prov canceled                        0
prov          start: threshold canceled                   0
prov          start: rule-report canceled                 0
prov          start: rule-report posted to CCR-T          0
prov          start: no proc required                     0
prov          start: success + logout dropped             0
prov          start: failure + logout dropped             0
prov          start: success: no-config                   0
prov          start: failure: no-config                   0
prov          start: success: over active limit           0
prov          start: failure: over active limit           0
prov          start: success: pending                     0
prov          start: failure: duplicate                   0
prov          start: success: no-memory                   0
prov          start: failure: no-memory                   0
prov          start: success: message posted              0
prov          start: success: message to ready-queue      0
prov          start: async: message posted                4
prov          start: async: message to ready queue        2
prov          message posted from ready queue             2
prov          response: late response                     0
prov          response: positive                          6
prov          response: negative                          0
prov          response: bad data                          0
prov          response timeout                            2
prov          timeout: message already in ready queue     12
prov          timeout: message posted                     0
prov          timeout: message to ready queue             0
prov          active refill: message posted               0
prov          active refill: message to ready queue       0
prov          message build failure                       0
prov          post failure                                0
prov          implied post                                5
prov          discard                                     0
prov          discard(no effect)                          0
prov          implied response                            0
prov          not supported                               0
prov          element allocated                           6
prov          element freed                               6
logout        start                                       6
logout        start: failure: bad data                    0
logout        start: ignore                               0
```

```
logout        start: failure: in recovery                0
logout        start: failure: spruious                   0
logout        start: prov canceled                       0
logout        start: threshold canceled                  0
logout        start: rule-report canceled                0
logout        start: rule-report posted to CCR-T         0
logout        start: no proc required                    0
logout        start: success + logout dropped            0
logout        start: failure + logout dropped            0
logout        start: success: no-config                  0
logout        start: failure: no-config                  0
logout        start: success: over active limit          0
logout        start: failure: over active limit          0
logout        start: success: pending                    0
logout        start: failure: duplicate                  0
logout        start: success: no-memory                  0
logout        start: failure: no-memory                  0
logout        start: success: message posted             4
logout        start: success: message to ready-queue     2
logout        start: async: message posted               0
logout        start: async: message to ready queue       0
logout        message posted from ready queue            2
logout        response: late response                    0
logout        response: positive                         6
logout        response: negative                         0
logout        response: bad data                         0
logout        response timeout                           0
logout        timeout: message already in ready queue    12
logout        timeout: message posted                    0
logout        timeout: message to ready queue            0
logout        active refill: message posted              0
logout        active refill: message to ready queue      0
logout        message build failure                      0
logout        post failure                               0
logout        implied post                               0
logout        discard                                    0
logout        discard(no effect)                         0
logout        implied response                           0
logout        not supported                              0
logout        element allocated                          6
logout        element freed                              6
ruleReport    start                                      12
ruleReport    start: failure: bad data                   0
```

```
ruleReport    start: ignore                                0
ruleReport    start: failure: in recovery                  0
ruleReport    start: failure: spruious                     0
ruleReport    start: prov canceled                         0
ruleReport    start: threshold canceled                    0
ruleReport    start: rule-report canceled                  0
ruleReport    start: rule-report posted to CCR-T           6
ruleReport    start: no proc required                      0
ruleReport    start: success + logout dropped              0
ruleReport    start: failure + logout dropped              0
ruleReport    start: success: no-config                    0
ruleReport    start: failure: no-config                    0
ruleReport    start: success: over active limit            0
ruleReport    start: failure: over active limit            0
ruleReport    start: success: pending                      0
ruleReport    start: failure: duplicate                    0
ruleReport    start: success: no-memory                    0
ruleReport    start: failure: no-memory                    0
ruleReport    start: success: message posted               6
ruleReport    start: success: message to ready-queue       0
ruleReport    start: async: message posted                 0
ruleReport    start: async: message to ready queue         0
ruleReport    message posted from ready queue              0
ruleReport    response: late response                      0
ruleReport    response: positive                           6
ruleReport    response: negative                           0
ruleReport    response: bad data                           0
ruleReport    response timeout                             0
ruleReport    timeout: message already in ready queue      0
ruleReport    timeout: message posted                      0
ruleReport    timeout: message to ready queue              0
ruleReport    active refill: message posted                0
ruleReport    active refill: message to ready queue        0
ruleReport    message build failure                        0
ruleReport    post failure                                 0
ruleReport    implied post                                 0
ruleReport    discard                                      0
ruleReport    discard(no effect)                           0
ruleReport    implied response                             0
ruleReport    not supported                                0
ruleReport    element allocated                            6
ruleReport    element freed                                6
threshold     start                                        113
```

```
threshold     start: failure: bad data                    0
threshold     start: ignore                               0
threshold     start: failure: in recovery                 0
threshold     start: failure: spruious                    0
threshold     start: prov canceled                        0
threshold     start: threshold canceled                   0
threshold     start: rule-report canceled                 0
threshold     start: rule-report posted to CCR-T          0
threshold     start: no proc required                     0
threshold     start: success + logout dropped             0
threshold     start: failure + logout dropped             0
threshold     start: success: no-config                   0
threshold     start: failure: no-config                   0
threshold     start: success: over active limit           0
threshold     start: failure: over active limit           0
threshold     start: success: pending                     0
threshold     start: failure: duplicate                   0
threshold     start: success: no-memory                   0
threshold     start: failure: no-memory                   0
threshold     start: success: message posted              92
threshold     start: success: message to ready-queue      21
threshold     start: async: message posted                0
threshold     start: async: message to ready queue        0
threshold     message posted from ready queue             0
threshold     response: late response                     0
threshold     response: positive                          92
threshold     response: negative                          0
threshold     response: bad data                          0
threshold     response timeout                            21
threshold     timeout: message already in ready queue     0
threshold     timeout: message posted                     0
threshold     timeout: message to ready queue             0
threshold     active refill: message posted               0
threshold     active refill: message to ready queue       0
threshold     message build failure                       0
threshold     post failure                                0
threshold     implied post                                0
threshold     discard                                     0
threshold     discard(no effect)                          0
threshold     implied response                            0
threshold     not supported                               0
threshold     element allocated                           113
threshold     element freed                               113
```

```
    general       start                                       0
    general       start: failure: bad data                    0
    general       start: ignore                               0
    general       start: failure: in recovery                 0
    general       start: failure: spruious                    0
    general       start: prov canceled                        0
    general       start: threshold canceled                   0
    general       start: rule-report canceled                 0
    general       start: rule-report posted to CCR-T          0
    general       start: no proc required                     0
    general       start: success + logout dropped             0
    general       start: failure + logout dropped             0
    general       start: success: no-config                   0
    general       start: failure: no-config                   0
    general       start: success: over active limit           0
    general       start: failure: over active limit           0
    general       start: success: pending                     0
    general       start: failure: duplicate                   0
    general       start: success: no-memory                   0
    general       start: failure: no-memory                   0
    general       start: success: message posted              0
    general       start: success: message to ready-queue      0
    general       start: async: message posted                0
    general       start: async: message to ready queue        0
    general       message posted from ready queue             0
    general       response: late response                     0
    general       response: positive                          0
    general       response: negative                          0
    general       response: bad data                          0
    general       response timeout                            0
    general       timeout: message already in ready queue     0
    general       timeout: message posted                     0
    general       timeout: message to ready queue             0
    general       active refill: message posted               0
    general       active refill: message to ready queue       0
    general       message build failure                       0
    general       post failure                                0
    general       implied post                                0
    general       discard                                     0
    general       discard(no effect)                          0
    general       implied response                            0
    general       not supported                               0
    general       element allocated                           0
```

```
  general       element freed                              0

Gx-plus upstream diameter-events:
  Category      Counter                                    Value
  prov          bad flags                                  0
  prov          bad fixed destination                      0
  prov          bad routed destination                     0
  prov          tx is over limit                           0
  prov          bad end-to-end id                          0
  prov          no peer for tx                             0
  prov          peer down while waiting for answer         0
  prov          timeout while waiting for answer           0
  prov          tx timeout                                 0
  prov          tx try limit                               0
  prov          tx failure                                 0
  prov          discarded                                  0
  prov          received answer is over limit              0
  prov          tx failure: no memory                      0
  prov          base-app-tx-timeout                        0
  prov          base-app-rx-timeout                        0
  prov          base-app-tx-discard                        0
  prov          base-app-rx-discard                        0
  logout        bad flags                                  0
  logout        bad fixed destination                      0
  logout        bad routed destination                     0
  logout        tx is over limit                           0
  logout        bad end-to-end id                          0
  logout        no peer for tx                             0
  logout        peer down while waiting for answer         0
  logout        timeout while waiting for answer           0
  logout        tx timeout                                 0
  logout        tx try limit                               0
  logout        tx failure                                 0
  logout        discarded                                  0
  logout        received answer is over limit              0
  logout        tx failure: no memory                      0
  logout        base-app-tx-timeout                        0
  logout        base-app-rx-timeout                        0
  logout        base-app-tx-discard                        0
  logout        base-app-rx-discard                        0
  ruleReport    bad flags                                  0
  ruleReport    bad fixed destination                      0
  ruleReport    bad routed destination                     0
```

```
ruleReport    tx is over limit                          0
ruleReport    bad end-to-end id                         0
ruleReport    no peer for tx                            0
ruleReport    peer down while waiting for answer        0
ruleReport    timeout while waiting for answer          0
ruleReport    tx timeout                                0
ruleReport    tx try limit                              0
ruleReport    tx failure                                0
ruleReport    discarded                                 0
ruleReport    received answer is over limit             0
ruleReport    tx failure: no memory                     0
ruleReport    base-app-tx-timeout                       0
ruleReport    base-app-rx-timeout                       0
ruleReport    base-app-tx-discard                       0
ruleReport    base-app-rx-discard                       0
threshold     bad flags                                 0
threshold     bad fixed destination                     0
threshold     bad routed destination                    0
threshold     tx is over limit                          0
threshold     bad end-to-end id                         0
threshold     no peer for tx                            0
threshold     peer down while waiting for answer        0
threshold     timeout while waiting for answer          0
threshold     tx timeout                                0
threshold     tx try limit                              0
threshold     tx failure                                0
threshold     discarded                                 0
threshold     received answer is over limit             0
threshold     tx failure: no memory                     0
threshold     base-app-tx-timeout                       0
threshold     base-app-rx-timeout                       0
threshold     base-app-tx-discard                       0
threshold     base-app-rx-discard                       0
general       bad flags                                 0
general       bad fixed destination                     0
general       bad routed destination                    0
general       tx is over limit                          0
general       bad end-to-end id                         0
general       no peer for tx                            0
general       peer down while waiting for answer        0
general       timeout while waiting for answer          0
general       tx timeout                                0
general       tx try limit                              0
```

```
  general       tx failure                                0
  general       discarded                                 0
  general       received answer is over limit             0
  general       tx failure: no memory                     0
  general       base-app-tx-timeout                       0
  general       base-app-rx-timeout                       0
  general       base-app-tx-discard                       0
  general       base-app-rx-discard                       0

Gx-plus downstream counters:
  Category      Counter                                   Value
  audit         new request                               0
  audit         retry request                             0
  audit         queue full                                0
  audit         no-memory                                 0
  audit         message posted                            0
  audit         message build failure                     0
  audit         post failure                              0
  d-exact       new request                               0
  d-exact       retry request                             0
  d-exact       queue full                                0
  d-exact       no-memory                                 0
  d-exact       message posted                            0
  d-exact       message build failure                     0
  d-exact       post failure                              0
  d-bulk        new request                               5
  d-bulk        retry request                             0
  d-bulk        queue full                                0
  d-bulk        no-memory                                 0
  d-bulk        message posted                            5
  d-bulk        message build failure                     0
  d-bulk        post failure                              0
  d-done        new request                               0
  d-done        retry request                             0
  d-done        queue full                                0
  d-done        no-memory                                 0
  d-done        message posted                            0
  d-done        message build failure                     0
  d-done        post failure                              0
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

clear network-access gx-plus replay | 3743

clear network-access gx-plus statistics | 3745

show network access nas-5g

IN THIS SECTION

- Syntax | 22509
- Description | 22509
- Options | 22510
- Required Privilege Level | 22510
- Output Fields | 22510
- Sample Output | 22513
- Release Information | 22514

Syntax

```
show network access nas-5g (state | statistics)
```

Description

Display statistics on the interactions between the Access Gateway Function (AGF) and other network functions during subscriber authentication.

Options

`state`	Display the status of the 5G connections
`statistics`	Display the statistics for the 5G core (5GC) connections.

Required Privilege Level

view

Output Fields

Table 884 on page 22510 lists the output fields for the `show network access nas-5g` command.

Table 884: show network access nas-5g

Field Name	Field Description
Connection up transitions	Number of established Non-Access Stratum (NAS) connections.
Connection down transitions	Number of NAS connections shut down.
Connection errors	Number of NAS connections that failed.
Total messages sent	Number of NAS signaling messages sent.
Total messages received	Number of NAS signaling messages received.
Total request retries	Number of retry attempts for NAS signaling messages.
Total request timeouts	Number of request messages that timed out.
Authentication request attempts	Number of attempted authentication requests.

Table 884: show network access nas-5g *(Continued)*

Field Name	Field Description
Authentication request messages sent	Number of authentication request messages sent successfully.
Authentication request messages authenticated	Number of authentication request messages authenticated.
Authentication request failures	Number of authentication request messages that failed.
Authentication response failures	Number of errors encountered during the processing of authentication response messages.
Authentication request timeouts	Number of timed out authentication request messages.
Maximum number of outstanding requests exceeded failures	Number of times the outstanding requests exceeded the failure count.
Session terminate request messages sent	Number of session termination request messages sent.
Session terminate response messages received	Number of session termination response messages received
Session terminate request message failures	Number of session termination request messages that failed.
Session terminate request timeouts	Number of session termination request messages that timed out.

Table 884: show network access nas-5g *(Continued)*

Field Name	Field Description
Session command logout received	Number of session log-out command messages received.
Session command logout receive failures	Number of received log-out messages that failed.
Session disconnect request sent	Number of session disconnect request messages sent.
Session disconnect request send failures	Number of errors encountered while sending the session disconnect request messages.
Session disconnect response received	Number of session disconnect response messages received.
Session disconnect response receive failures	Number of errors encountered during the processing of session disconnect response messages.
Number of outstanding requests	Number of pending sessions that have not been processed.
Total number of allocated messages	Total number of messages allocated in memory.

Sample Output

show network-access nas-5g state

```
user@host> show network-access nas-5g state
```

```
NAS-5G module state information
  Connection state: UP
  Total number of subscribers: 11
```

show network-access nas-5g statistics

```
user@host> show network-access nas-5g statistics
```

```
NAS-5G module
statistics

  Connection up transitions:
2
  Connection down transitions:
1
  Connection errors:
0

  Total messages sent:
11

  Total messages received:
11
  Total request retries:
0
  Total request timeouts:
0
  Authentication request attempts:
5
  Authentication request messages sent:
5
```

```
  Authentication request messages authenticated:
5
  Authentication request failures:
0
  Authentication response failures:
0
  Authentication request timeouts:
0
  Maximum number of outstanding requests exceeded failures:
0
  Session terminate request messages sent:
6
  Session terminate response messages received:
6
  Session terminate request message failures:
0
  Session terminate response failures:
0
  Session terminate request timeouts:
0
  Number of outstanding requests:
0
  Total number of allocated messages:
0
```

Release Information

Command introduced in Junos OS Release 22.3R1.

show network-access nasreq statistics

IN THIS SECTION

- Syntax | 22515
- Description | 22515
- Required Privilege Level | 22515
- Output Fields | 22515

Sample Output | 22519
Release Information | 22520

Syntax

```
show network-access nasreq statistics
```

Description

Display information about the NASREQ protocol exchanges that are used to authenticate and authorize subscribers when they log in.

Required Privilege Level

view

Output Fields

Table 885 on page 22515 lists the output fields for the `show network-access nasreq statistics` command. Output fields are listed in the approximate order in which they appear.

Table 885: show network-access nasreq statistics Output Fields

Field Name	Field Description
`Authentication request attempts`	Number of authentication requests forwarded to the NASREQ engine. Note that authentication includes authorization in one message exchange.
`Authentication request failures`	Number of authentication request failures at the NASREQ engine. This includes timeout failures of the request waiting in the transmit queue.
`Authentication request messages sent`	Number of authentication requests forwarded to the Diameter engine.

Table 885: show network-access nasreq statistics Output Fields *(Continued)*

Field Name	Field Description
`Authentication request message failures`	Number of authentication requests failures at the Diameter engine.
`Authentication request messages timeouts`	Number of authentication requests that timed out waiting for a response from the NASREQ server.
`Authentication denies failures`	Number of authentication responses with an unsuccessful result-code.
`Authentication grants received`	Number of authentication responses with a successful result-code.
`Authorization request attempts`	Number of authorization requests to the NASREQ engine.
`Authorization request failures`	Number of authorization request failures at the NASREQ engine. This includes timeout failures of the request waiting in the transmit queue.
`Authorization request messages sent`	Number of authorization requests forwarded to the Diameter engine.
`Authorization request message failures`	Number of authorization requests failures at the Diameter engine.
`Authorization request messages timeouts`	Number of authorization requests that timed out waiting for a response from the NASREQ server.
`Authorization denies failures`	Number of authorization responses with an unsuccessful result-code.
`Authorization grants received`	Number of authorization responses with a successful result-code.
`Session-terminate request attempts`	Number of session-terminate requests forwarded to the NASREQ engine.

Table 885: show network-access nasreq statistics Output Fields *(Continued)*

Field Name	Field Description
`Session-terminate request failures`	Number of session-terminate requests failures at the NASREQ engine. This includes timeout failures of the request waiting in the transmit queue.
`Session-terminate request messages sent`	Number of session-terminate requests forwarded to the Diameter engine.
`Session-terminate request messages failures`	Number of session-terminate requests failures at the Diameter engine.
`Session-terminate request messages timeouts`	Number of session-terminate requests that timed out waiting for a response from the NASREQ server.
`Session-terminate response message failures`	Number of session-terminate responses with an unsuccessful result-code.
`Session-terminate response messages received`	Number of session-terminate responses with a successful result-code.
`Abort-session requests received`	Number of Abort-Session-Request messages received by the NASREQ engine.
`Abort-session response-ack messages sent`	Number of Abort-Session-Request messages receiving an ACK response.
`Abort-session response-nack messages sent`	Number of Abort-Session-Request messages receiving a NACK response.
`Abort-session response message failures`	Number of Abort-Session-Request messages that failed transmission, either due to timeout in the transmit queue, or failed by the Diameter engine.
`Number of NASREQ subscribers`	Number of active subscribers that received NASREQ authorization.

Table 885: show network-access nasreq statistics Output Fields *(Continued)*

Field Name	Field Description
`Number of result-code-1xxx`	Number of response messages with an informational result-code.
`Number of result-code-2xxx`	Number of response messages with a success result-code.
`Number of result-code-3xxx`	Number of response messages with a protocol-error result-code.
`Number of result-code-4xxx`	Number of response messages with an transient-error result-code.
`Number of result-code-5xxx`	Number of response messages with a permanent-error result-code.
`Number of result-code-other`	Number of response messages with an unregistered result-code.
`Transmit queue time-outs`	Number of request/responses that timed out waiting for Diameter engine resources.
`Transmit message drops`	Number of request/responses dropped.
`Transmit message failures`	Number of request/response failures at the Diameter engine.
`Number of requests on the Tx queue`	Number of requests/responses currently on the transmit queue.
`Number of requests waiting for a response`	Number of requests currently in the response queue waiting for a response.
`Number of outstanding requests`	Number of outstanding requests, including AAR and STR.
`Total number of allocated messages`	Number of currently allocated request/response messages.

Sample Output

command-name

```
user@host> show network-access nasreq statistics
Authentication request attempts:4000
Authentication request failures:0
Authentication request messages sent:4000
Authentication request message failures:0
Authentication request messages timeouts:0
Authentication denies failures:0
Authentication grants received:4000
Authorization request attempts:4000
Authorization request failures:0
Authorization request messages sent:4000
Authorization request message failures:0
Authorization request messages timeouts:0
Authorization denies failures:0
Authorization grants received:4000
Session-terminate request attempts:1000
Session-terminate request failures:0
Session-terminate request messages sent:1000
Session-terminate request message failures:0
Session-terminate request messages timeouts:0
Session-terminate response messages failures:0
Session-terminate response messages received:1000
Abort-session requests received:1000
Abort-session response-ack messages sent:1000
Abort-session response-nack messages sent:0
Abort-session response message failures:0
Number of NASREQ subscribers:3000
Number of result-code-1xxx:0
Number of result-code-2xxx:5000
Number of result-code-3xxx:0
Number of result-code-4xxx:0
Number of result-code-5xxx:0
Number of result-code-other:0
Transmit queue time-outs:0
Transmit message drops:0
Transmit message failures:0
Number of requests on the Tx queue:0
Number of requests waiting for a response:0
```

```
Number of outstanding requests:0
Total number of allocated messages:0
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

NASREQ for Authentication and Authorization

NASREQ for Authentication and Authorization

Diameter Base Protocol

show network-access ocs

IN THIS SECTION

- Syntax | 22520
- Description | 22521
- Options | 22521
- Required Privilege Level | 22521
- Output Fields | 22521
- Sample Output | 22522
- Release Information | 22522

Syntax

```
show network-access ocs
<state>
<statistics [brief | detail | extensive]>
```

Description

Display Online Charging System (OCS) provisioning state and statistics information.

Options

brief | detail | extensive	(Optional) Display the specified level of output.
state	(Optional) Display OCS provisioning state.
statistics	(Optional) Display OCS statistics.

Required Privilege Level

view

Output Fields

Table 886 on page 22521 lists the output fields for the `show network-access ocs` command. Output fields are listed in the approximate order in which they appear.

Table 886: show network-access ocs Output Fields

Field Name	Field Description	Level of Output
Ocs state	State of the OCS components, including the following: • **created**—Just created. • **dead**—Ready to be deleted. • **final**—Final interrogation in progress. • **idle**—Idle. • **initial**—Initial interrogation in progress. • **intermediate**—Intermediate interrogation in progress.	All levels
Ocs general counters	Type and number of Diameter response and answer messages for OCS communication using the Gy protocol.	All levels

Sample Output

show network-access ocs state

```
user@host> show network-access ocs state
Ocs state:
  Component                     Value
  state                         active
  active-configuration          yes
  queue-state                   normal
  subscribers                   2000
  sub-in-idle                   2000
```

show network-access ocs statistics

```
user@host> show network-access ocs statistics
Ocs general counters:
  Counter                                                   Value
  ccr-gy-i                                                  4001
  cca-gy-i                                                  4001
  ccr-gy-u                                                  38001
  cca-gy-u                                                  38001
  ccr-gy-t                                                  2001
  cca-gy-t                                                  2001
  asr                                                       1
  asa-ack                                                   1
```

Release Information

Command introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

ocs (Diameter Applications) | **11771**

clear network-access ocs statistics | **3747**

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

show network-access pcrf

IN THIS SECTION

- Syntax | 22523
- Description | 22523
- Options | 22523
- Required Privilege Level | 22524
- Output Fields | 22524
- Sample Output | 22527
- Release Information | 22528

Syntax

```
show network-access pcrf
<state>
<statistics [brief | detail | extensive]>
<subscribers>
```

Description

Display Policy and Charging Rules Function (PCRF) provisioning state and statistics information.

Options

brief | detail | extensive	(Optional) Display the specified level of output.
state	(Optional) Display PCRF provisioning state.
statistics	(Optional) Display PCRF statistics.

subscribers (Optional) Display session ID and provisioning state for each PCRF subscriber.

Required Privilege Level

view

Output Fields

Table 887 on page 22525 lists the output fields for the `show network-access pcrf` command. Output fields are listed in the approximate order in which they appear.

Table 887: show network-access pcrf Output Fields

Field Name	Field Description	Level of Output

`Pcrf state`	State of the PCRF components:
	• **active**—Installed services from local grant.
	• **created**—Just created.
	• **dead**—Received logout response or timed out waiting for logout response.
	• **deny**—Received deny from remote and waiting for subscriber logout.
	• **failed**—GRES before active state or provisioning failed, waiting for subscriber logout.
	• **grant**—Received remote grant and installing services.
	• **local-active**—Services installed after local grant.
	• **local-deny**—Denied by local decision waiting for subscriber logout.
	• **local-grant**—Received local grant and installing services.
	• `local-reinit`—Router is waiting for the CCA-GX-T so it can reinitialize the PCRF session. Local reinitialization must be configured on the PCRF partition.
	This state results from one of the following:
	  • When a RAR is received while the PCRF is in local-active state.
	  • When a RAR is received while the PCRF is in started state and no default service is configured.
	  • When the provisioning-done event occurs while the PCRF is in local-reinit-early state.
	• `local-reinit-early`—The PCRF is engaged in provisioning services when a failure is detected or the PCRF server sends a RAR message of any type. The session transitions
		All levels

Table 887: show network-access pcrf Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	to local-reinit state when the provisioning completes. Local reinitialization must be configured on the PCRF partition. • **logout**—Waiting for logout response. • **ocs-final-wait**—Waiting for OCS to complete final interrogation. • **report**—Sending report. • **update**—Processing RAR-update. • **upd-from-local**—Received remote grant, in local-active state, and installing services. • **started**—Log in.	
`Pcrf general counters`	Type and number of Diameter response and answer messages for PCRF communication using the Gx protocol.	All levels
`Pcrf diameter event counters`	State and number of Diameter events.	All levels
`Pcrf rule install counters`	Result and number of rule installations.	All levels
`Name`	PCRF subscriber name.	All levels
`Session-ID`	Identifier for the PCRF session.	All levels

Sample Output

show network-access pcrf state

```
user@host> show network-access pcrf state
Pcrf state:
  Component                     Value
  state                         active
  active-configuration          yes
  queue-state                   normal
  subscribers                   2000
  sub-in-idle                   2000
```

show network-access pcrf statistics

```
user@host> show network-access pcrf statistics
Pcrf general counters:
  Counter                                                   Value
  ccr-gx-i                                                  4001
  ccr-gx-i-retry                                            4
  cca-gx-i                                                  4001
  ccr-gx-u                                                  4001
  cca-gx-u                                                  4001
  ccr-gx-t                                                  2001
  cca-gx-t                                                  2001

Pcrf diameter event counters:
  Diameter event                                            Value
  no peer for tx                                            4

Pcrf rule install counters:
  Result                                                    Value
  success                                                   16004
  family mismatch                                           5
```

show network-access pcrf subscribers

```
user@host> show network-access pcrf subscribers
Pcrf subscribers:
  Name                              Session-Id            Prov-State
  user45-example                    2761                  active
  user48-example                    12731                 local-active
  user51-example                    12732                 local-active
  user54-example                    16887                 started
  user57-example                    16895                 started
```

Release Information

Command introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

- pcrf (Diameter Applications) | **12325**
- clear network-access pcrf | **3749**
- *3GPP Policy and Charging Control for Wireline Provisioning and Accounting*
- *3GPP Policy and Charging Control for Wireline Provisioning and Accounting*
- *3GPP Policy and Charging Control for Wireline Provisioning and Accounting*

show network-access radsec local-certificate

IN THIS SECTION

- Syntax | **22529**
- Description | **22529**
- Options | **22529**
- Required Privilege Level | **22529**
- Output Fields | **22530**
- Sample Output | **22530**

- Sample Output | 22531
- Sample Output | 22531
- Release Information | 22531

Syntax

```
show network-access radsec local-certificate
<state>
<statistics [brief | detail | extensive]>
<certificate-name>
```

Description

Display the state and statistics of local certificate acquisition. RADSEC uses local certificates dynamically acquired from the public key infrastructure to establish a TLS connection.

If a certificate is not available, or if it was revoked, the RADSEC client will try to retrieve it every 300 seconds. Response timeout is 10 seconds, and failures are retried in 10 seconds.

Options

state	Display the state of acquisition for the local certificate.
statistics	Display acquisition statistics for the local certificate.
brief | detail | extensive	(Optional) Display the specified level of output.
certificate-name	(Optional) Display detailed information about the specified certificate.

Required Privilege Level

view

Output Fields

Table 888 on page 22530 lists the output fields for the `show network-access local-certificate` command. Output fields are listed in the approximate order in which they appear.

Table 888: show network-access radsec local-certificate Output Fields

Field Name	Field Description	Level of Output
Local certificate state	State of acquisition for the local certificate. • `active`—Local certificate is active. • `waiting`—Waiting to acquire local certificate.	all
Local certificate general counters	Statistics for RADSEC local certificate acquisition.	all **NOTE**: Default output level will list only non-zero counters. Use `detail` or `extensive` to view all counters.

Sample Output

show network-access radsec local-certificate state

```
user@host> show network-access radsec local-certificate state
Local certificate state:
  cert-2                                        active
  cert-4                                        waiting
  qqq                                           waiting
```

Sample Output

show network-access radsec local-certificate statistics

```
user@host> show network-access radsec local-certificate statistics
Local certificate general counters:
  total-requests                                36
```

Sample Output

show network-access radsec local-certificate statistics detail

```
user@host> show network-access radsec local-certificate statistics detail
Local certificate general counters:
  total-requests                                36
  failed-requests                               0
  total-responses                               0
  configured-responses                           0
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

RADIUS over TLS (RADSEC)

show network-access radsec state

IN THIS SECTION

- Syntax | 22532
- Description | 22532

- Options | 22532
- Required Privilege Level | 22532
- Output Fields | 22532
- Sample Output | 22534
- Release Information | 22534

Syntax

```
show network-access radsec state
<destination destination-id>
```

Description

Display the connection state of RADSEC destinations.

Options

destination *destination-id* (Optional) Display detailed information about the request specified by this RADSEC destination.

Required Privilege Level

view

Output Fields

Table 889 on page 22533 lists the output fields for the `show network-access state` command. Output fields are listed in the approximate order in which they appear.

Table 889: show network-access radsec state Output Fields

Field Name	Field Description	Level of Output
Radsec state	State information for the RADSEC destination.	all
state	State of the RADSEC connection. • **connecting**—Establishing TCP connection. • **ssl-handshake**—SSL negotiation in progress. • **open**—RADSEC session is established for exchange of RADIUS messages. • **pause**—Pause for restart of connection process. The length of the pause is determined by the reason for restarting. • **local-cert-wait**—Connection initiated but waiting for local certificate to complete negotiation.	all
secs-in-state	Length of time in seconds of the current state.	all
remaining-secs	Length of time in seconds remaining for the current state.	all
pause-reason	The reason for restarting the connection, which triggers the pause state. The pause reason determines the length of the pause until reattempting the connection.	all
acct-support	Shows whether the remote server supports accounting. • `Y`—Remote server supports accounting. This is the default value. • `N`—If the client receives a NAK response to an accounting request, all accounting requests will be dropped.	all
remote-failures	Number of consecutive failures on the remote side.	all
tx-requests	Number of RADIUS request messages transmitted.	all

Table 889: show network-access radsec state Output Fields *(Continued)*

Field Name	Field Description	Level of Output
tx-responses	Number of RADIUS response messages transmitted.	all

Sample Output

show network-access radsec state

```
user@host> show network-access radsec state
Radsec state:
    destination                                     895
    state                                           open
    secs-in-state                                   66
    remaining-secs                                  4294967295
    pause-reason                                    none
    acct-support                                    Y
    remote-failures                                 0
    tx-requests                                     0
    tx-responses                                    0
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

RADIUS over TLS (RADSEC)

show network-access radsec statistics

IN THIS SECTION

- Syntax | 22535
- Description | 22535
- Options | 22535
- Required Privilege Level | 22535
- Output Fields | 22536
- Sample Output | 22536
- Sample Output | 22537
- Release Information | 22538

Syntax

```
show network-access radsec statistics
<[brief | detail | extensive]>
<destination destination-id>
```

Description

Display the connection statistics for the RADSEC destinations.

Options

brief | detail | extensive	(Optional) Display the specified level of output. The default is `brief`, which will list only non-zero counters.
destination *destination-id*	(Optional) Display detailed information about the request specified by this RADSEC destination.

Required Privilege Level

view

Output Fields

Table 890 on page 22536 lists the output fields for the `show network-access statistics` command. Output fields are listed in the approximate order in which they appear.

Table 890: show network-access radsec statistics Output Fields

Field Name	Field Description	Level of Output
Radsec general counters	Statistics for RADSEC syslog event counters.	all **NOTE**: Default output level will list only non-zero counters. Use `detail` or `extensive` to show all counters.
Destination	ID number of the RADSEC destination.	all

Sample Output

show network-access radsec statistics

```
user@host> show network-access radsec statistics
Radsec general counters:

  destination                                    895
  start-events                                   1
  clear-events                                   1
  timeout-events                                 1
  loc-cert-acq-events                            1
  connected-events                               1
  ssl-ready-events                               1
```

Sample Output

show network-access radsec statistics detail

```
user@host> show network-access radsec statistics detail

Radsec general counters:

  destination                                    895
  start-events                                   1
  clear-events                                   1
  force-disconnect-events                        0
  timeout-events                                 1
  loc-cert-acq-events                            1
  loc-cert-lost-events                           0
  connected-events                               1
  conn-failed-events                             0
  ssl-disconnected-events                        0
  ssl-ready-events                               1
  in-auth-reqs                                   0
  in-acct-reqs                                   0
  in-dyn-req-resps                               0
  tx-auth-reqs                                   0
  tx-acct-reqs                                   0
  tx-wd-reqs                                     9
  tx-late-auth-reqs                              0
  tx-late-acct-reqs                              0
  tx-dyn-req-resps                               0
  rx-auth-resps                                  0
  rx-acct-resps                                  0
  rx-dyn-reqs                                    0
  rx-dyn-req-naks                                0
  rx-dyn-req-drops                               0
  rx-wd-resps                                    9
  rx-resps                                       0
  rx-late-resps                                  0
  rx-other-drops                                 0
  resp-disconnect-drops                          0
  id-disconnect-drops                            0
  id-timeout-drops                               0
  tx-req-no-acct-supports                        0
  tx-req-dest-downs                              0
```

```
tx-req-overflows                          0
tx-req-disconnects                        0
tx-req-bad-responses                      0
tx-req-id-reuse-timeouts                  0
tx-resp-dest-downs                        0
tx-wd-reqs                                9
rx-wd-resps                               9
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

RADIUS over TLS (RADSEC)

show network-access requests pending

IN THIS SECTION

- Syntax | 22539
- Description | 22539
- Options | 22539
- Required Privilege Level | 22539
- Output Fields | 22539
- Sample Output | 22540
- Sample Output | 22541
- Sample Output | 22541
- Release Information | 22541

Syntax

```
show network-access requests pending
<detail>
<index number >
```

Description

Display the status of pending authentication requests.

Options

- none—Show pending authentication requests.
- "show network-access requests pending" on page 22538"show network-access requests pending" on page 22538`detail`—Display detailed information about all pending requests.
- `index` *`number`*—(Optional) Display detailed information about the request specified by this index number. Use the command without options to obtain a list of requests and index numbers.

Required Privilege Level

view

Output Fields

Table 891 on page 22539 lists the output fields for the show network-access requests pending command. Output fields are listed in the approximate order in which they appear.

Table 891: show network-access requests pending Output Fields

Field Name	Field Description
`Index`	Internal number identifying the pending request. Use this number to obtain more information on the record.
`User`	Originator of authentication request.

Table 891: show network-access requests pending Output Fields *(Continued)*

Field Name	Field Description
`Status`	The pending requests are requests and responses that are not yet sent back to the respective clients. The pending requests can be in one of the following states: • `Processing`: This request is being processed by the device. The authentication process has started but is not complete. • `Waiting on Auth Server`: The request is sent to an external authentication server, and the device is waiting for the response. • `Processed`: This request has completed authentication (success or failure). The results are not yet forwarded back to the client. • `Request cancelled by Admin`: This request was cancelled by the Admin. The reply with cancel code is not yet sent back to the client.
`Profile`	The profile determines how the user is authenticated. Local clients defined with the statement `access profile client` are authenticated with the password authentication. Clients configured external to the device, on a RADIUS or LDAP server are authenticated with RADIUS or LDAP authentication.

Sample Output

show network-access requests pending

```
user@host> show network-access requests pending
Information about pending authentication entries
  Total pending authentication requests: 2
Index User                  Status
1     Sun                   Processing
2     Sam                   Processed
```

Sample Output

show network-access requests pending detail

```
user@host> show network-access requests pending detail
Information about pending authentication entries
  Total pending authentication requests: 2
Index: 1  User: Sun
  Status: Processing
  Profile: Sunnyvale-firewall-users
Index: 2  User: Sam
  Status: Processed
  Profile: Westford-profile
```

Sample Output

show network-access requests pending index 1

```
user@host> show network-access requests pending index 1
Index: 1  User: Sun
  Status: Processing
  Profile: Sunnyvale-firewall-users
```

Release Information

Command introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

clear network-access requests pending | 3753

show network-access requests statistics

IN THIS SECTION

- Syntax | 22542
- Description | 22542
- Required Privilege Level | 22542
- Output Fields | 22542
- show network-access requests statistics | 22543
- Release Information | 22544

Syntax

```
show network-access requests statistics
```

Description

Display authentication statistics for the configured authentication type.

Required Privilege Level

view

Output Fields

Table 892 on page 22542 lists the output fields for the network-access requests statistics command. Output fields are listed in the approximate order in which they appear.

Table 892: show network-access requests statistics Output Fields

Field Name	Field Description
`Total requests received`	Total number of authentication requests that the device received from clients.

Table 892: show network-access requests statistics Output Fields *(Continued)*

Field Name	Field Description
`Total responses sent`	Total number of authentication responses that the device sent to the clients.
`Success responses`	Total number of clients that authenticated successfully.
`Failure responses`	Total number of clients that failed to authenticate.

show network-access requests statistics

command-name

```
user@host> show network-access requests statistics
General authentication statistics
  Total requests received: 100
  Total responses sent: 70
Radius authentication statistics
  Total requests received: 40
  Success responses: 20
  Failure responses: 20
Radius reauthentication statistics
  Total requests received: 0
  Success responses: 0
  Failure responses: 0
LDAP authentication statistics
  Total requests received: 30
  Success responses: 15
  Failure responses: 15
Local authentication statistics
  Total requests received: 5
  Success responses: 2
  Failure responses: 3
Local re-authentication statistics
  Total requests received: 0
  Success responses: 0
  Failure responses: 0
Securid authentication statistics
```

```
Total requests received: 15
Success responses: 3
Failure responses: 12
```

Release Information

Command modified in Release 9.1 of Junos OS.

RELATED DOCUMENTATION

clear network-access requests statistics

show network-access s6a

IN THIS SECTION

- Syntax | 22544
- Description | 22545
- Options | 22545
- Required Privilege Level | 22545
- Sample Output | 22545
- Release Information | 22547

Syntax

```
show network-access s6a
<state>
<statistics>
<statistics extensive>
```

Description

Displays information about the s6a protocol exchanges that are used to authenticate subscribers when subscribers log in.

Options

state	Displays the state of information about authentication.
statistics	Displays information about all s6a general counters.
statistics extensive	Displays detail information about all s6a general counters.

Required Privilege Level

view

Sample Output

show network-access s6a state

```
user@host> show network-access s6a state
S6a state:
  Component                                       Value
  active-configuration                            yes
  queue-state                                     normal
  request-count                                   0
```

show network-access s6a statistics

```
user@host> show network-access s6a statistics
S6a general counters:
Counter ............Value
aia-grant ..........1
```

show network-access s6a statistics extensive

```
user@host# show network-access s6a statistics extensive
S6a general counters:
  Counter                                                   Value
  air                                                       0
  air-retry                                                 0
  air-failures                                              0
  aia                                                       0
  aia-grant                                                 0
  aia-deny                                                  0
  aia-timeout                                               0
  aia-failure                                               0
  aia-late-response                                         0
  aia-parse-errors                                          0
  aia-drops-no-session                                      0
  aia-drops-bad-orealm                                      0
  aia-drops-bad-ohost                                       0
  aia-drops-no-result                                       0
  aia-drops-other                                           0
  aia-bad-result                                            0
  aia-bad-data                                              0
  rx-unsupported-resp-cmd                                   0
  rx-bad-experimental-result                                0
  rx-bad-authentication-info                                0
  rx-bad-utran-vector                                       0
  rx-bad-eutran-vector                                      0
  rx-bad-geran-vector                                       0
  rx-parse-errors                                           0

S6a diameter event counters:
  Diameter event                                            Value
  bad data message                                          0
  good data message                                         0
  bad flags                                                 0
  bad fixed destination                                     0
  bad routed destination                                    0
  tx is over limit                                          0
  bad end-to-end id                                         0
  no peer for tx                                            0
  peer down while waiting for answer                        0
  timeout while waiting for answer                          0
```

```
tx timeout                                          0
tx try limit                                        0
tx failure                                          0
discarded                                           0
received answer is over limit                       0
tx failure: no memory                               0
base-app-tx-timeout                                 0
base-app-rx-timeout                                 0
base-app-tx-discard                                 0
base-app-rx-discard                                 0
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Example: Configure S6a Application

authentication-order | **2633**

partition (s6a) | **12190**

show network-access securid-node-secret-file

IN THIS SECTION

- Syntax | **22548**
- Description | **22548**
- Required Privilege Level | **22548**
- Output Fields | **22548**
- Sample Output | **22548**
- Release Information | **22549**

Syntax

```
show network-access securid-node-secret-file
```

Description

Display the path to the node secret file for the SecurID authentication type.

Required Privilege Level

view

Output Fields

Table 893 on page 22548 lists the output fields for the network-access securid-node-secret-file command. Output fields are listed in the approximate order in which they appear.

Table 893: show network-access securid-node-secret-file Output Fields

Field Name	Field Description
`SecurID Server`	Name of the SecurID authentication server.
`Node Secret File`	Path to the node secret file.

Sample Output

show network-access securid-node-secret-file

```
user@host> show network-access securid-node-secret-file
SecurID server node secret file:
SecurID Server       Node Secret File
ace-server1          /var/db/securid/ace-server1/node-secret
```

Release Information

Command introduced in Release 9.1 of Junos OS.

RELATED DOCUMENTATION

configuration-file | 4514
securid-server | 16447
clear network-access securid-node-secret-file | 3756

show network-agent statistics

IN THIS SECTION

- Syntax | 22549
- Description | 22550
- Options | 22550
- Required Privilege Level | 22550
- Output Fields | 22550
- Sample Output | 22552
- Release Information | 22560

Syntax

```
show network-agent statistics
<brief | detail>
<juniper>
<gnmi>
<all>
<subscription-id identifier>
```

Description

Display details of all or specific sensor subscriptions configured for Junos telemetry interface (JTI). Subscriptions can be made using wither a Juniper API or a gRPC Network Management Interface (gNMI) API.

Options

brief | detail	(Optional) Display brief or detail level of output.
juniper	(Optional) Displays subscriptions using the Juniper API.
gnmi	(Optional) Displays subscriptions using the gRPC Network Management Interface (gNMI) API.
all	(Optional) Display both Juniper and gNMI-based subscription statistics.
subscription-id ***identifier***	(Optional) Display subscription statistics specific to a subscription. Configure an identifier in the range from **0** through **4294967295**. By default, information is displayed for all subscriptions.

NOTE: It is recommended to use the `show network-agent statistics brief/detail` command to obtain a more accurate output.

Required Privilege Level

view

Output Fields

Table 894 on page 22550 lists the output fields for the `show network-agent statistics` command. Output fields are listed in the approximate order in which they appear.

Table 894: show network-agent statistics Output Fields

Field Name	Field Description
`Subscription ID`	Numerical identifier of the sensor.

Table 894: show network-agent statistics Output Fields *(Continued)*

Field Name	Field Description
`Type`	Type of API used for the subscription. Values are `juniper` or `gNMI`.
`Client IP`	IP address of the client that is collecting sensor data.
`Subscription Time`	Time of subscription, which helps to correlate statistics information from the provisioning logs that are taken over a period of multiple subscriptions.
`Sensor Path`	Resource path configured for the subscription; for example, `/components/`.
`Reporting Interval`	Interval at which statistics are streamed for a subscription.
`Component(s)`	Component for which statistics are provided.
`Average Latency`	Average latency values per sensor, which helps to check the latency of any given sensor on the device.
`Circular Buffer used`	Approximate circular buffer usage per sensor, which provides an early alert if drops are likely for any specific sensor.
`Bytes Sent`	Number of bytes transiting for this sensor.
`Packets Sent`	Number of packets transiting for this sensor.
`Drops`	Number of drops for this sensor.
`Initial Sync Bytes Sent`	Number of bytes sent for the subscription's initial synchronization.
`Initial Sync Packets Sent`	Number of packets sent during the subscription's initial synchronization.

Table 894: show network-agent statistics Output Fields *(Continued)*

Field Name	Field Description
`Initial Sync Drops`	Number of drops during subscription's initial synchronization.

Sample Output

show network-agent statistics (detailed output, no subscription present)

```
user@host> show network-agent statistics detail
No sensors subscribed/available
```

show network-agent statistics (detailed output, subscription creation in progress)

```
user@host> show network-agent statistics detail
Subscription not created completely, so not available for queries right now
```

show network-agent statistics (detailed output, subscription created, system operating in a steady state)

```
user@host> show network-agent statistics detail
Subscription Details :
    Subscription ID                     : 1
    Type                                : juniper
    Client IP                           : ipv6:::ffff:10.209.0.224:45888
    Subscription Time                   : 2020-02-24 04:08:11 UTC

    Sensor Statistics :
        Sensor Path                     : /components/
        Reporting Interval              : 2
        Component(s)                    : jkdsd,chassisd
        Average Latency                 : 5
        Circular Buffer used            : 12
        Bytes Sent                      : 176530
        Packets Sent                    : 21
```

```
        Drops                           : 0

        Child Sensor Statistics :
            Path                        : /components/
            Component                   : chassisd
            Bytes Sent                  : 172773
            Packets Sent                : 14
            Drops                       : 0

        Child Sensor Statistics :
            Path                        : /components/
            Component                   : jkdsd
            Bytes Sent                  : 3757
            Packets Sent                : 7
            Drops                       : 0
```

show network-agent statistics brief gnmi (shows subscription path)

```
user@host> show network-agent statistics brief gnmi

Subscription Details :
    Subscription ID                     : 1
    Type                                : gnmi
    Client IP                           : ipv4:::ffff:10.53.113.65:37934
    GNMI mode                           : STREAM
    Subscription Time (UTC)             : Mon May 13 06:38:39 2024

    Sensor Statistics :
        Sensor Path                     : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
        Reporting Interval              : 10
        Node name(s)                    : r0_re0,r0_re0,tt-ygadiraju-15341379-vm,tt-
ygadiraju-15341379-vm
        Component(s)                    : PFE,na-grpcd,rpd,mib2d,xmlproxyd_TM_Thread_1
        GNMI Sub Mode                   : SAMPLE

        Child Sensor Statistics :
            Path                        : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
            Node-name                   : r0_re0
```

```
            Component                   : na-grpcd
            GNMI-SubMode                : SAMPLE

        Child Sensor Statistics :
            Path                        : /junos/system/linecard/interface/logical/usage/
            Node-name                   : r0_re0
            Component                   : PFE
            Component-ID                : 0
            GNMI-SubMode                : SAMPLE

        Child Sensor Statistics :
            Path                        : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
            Node-name                   : tt-ygadiraju-15341379-vm
            Component                   : mib2d
            GNMI-SubMode                : SAMPLE

        Child Sensor Statistics :
            Path                        : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
            Node-name                   : tt-ygadiraju-15341379-vm
            Component                   : rpd
            GNMI-SubMode                : SAMPLE

        Child Sensor Statistics :
            Path                        : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
            Node-name                   : tt-ygadiraju-15341379-vm
            Component                   : xmlproxyd_TM_Thread_1
            GNMI-SubMode                : SAMPLE

        Child Sensor Statistics :
            Path                        : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
            Node-name                   : tt-ygadiraju-15341379-vm
```

show network-agent statistics brief gnmi (shows subscription path)

```
user@host> show network-agent statistics brief gnmi

Subscription Details :
```

```
    Subscription ID                     : 1
    Type                                : gnmi
    Client IP                           : ipv4:::ffff:10.53.113.65:37934
    GNMI mode                           : STREAM
    Subscription Time (UTC)             : Mon May 13 06:38:39 2024

    Sensor Statistics :
        Sensor Path                     : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
        Reporting Interval              : 10
        Node name(s)                    : r0_re0,r0_re0,tt-ygadiraju-15341379-vm,tt-
ygadiraju-15341379-vm
        Component(s)                    : PFE,na-grpcd,rpd,mib2d,xmlproxyd_TM_Thread_1
        GNMI Sub Mode                   : SAMPLE

        Child Sensor Statistics :
            Path                        : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
            Node-name                   : r0_re0
            Component                   : na-grpcd
            GNMI-SubMode                : SAMPLE

        Child Sensor Statistics :
            Path                        : /junos/system/linecard/interface/logical/usage/
            Node-name                   : r0_re0
            Component                   : PFE
            Component-ID                : 0
            GNMI-SubMode                : SAMPLE

        Child Sensor Statistics :
            Path                        : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
            Node-name                   : tt-ygadiraju-15341379-vm
            Component                   : mib2d
            GNMI-SubMode                : SAMPLE

        Child Sensor Statistics :
            Path                        : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
            Node-name                   : tt-ygadiraju-15341379-vm
            Component                   : rpd
            GNMI-SubMode                : SAMPLE
```

```
        Child Sensor Statistics :
            Path                        : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
            Node-name                   : tt-ygadiraju-15341379-vm
            Component                   : xmlproxyd_TM_Thread_1
            GNMI-SubMode                : SAMPLE

        Child Sensor Statistics :
            Path                        : /interfaces/interface[name='ge-0/0/2']/subinterfaces/
subinterface/
            Node-name                   : tt-ygadiraju-15341379-vm
```

show network-agent statistics brief gnmi (shows subscription details)

```
user@host> show network-agent statistics detail gnmi
Subscription Details :
    Subscription ID                     : 1
    Type                                : gnmi
    Client IP                           : ipv4:::ffff:10.53.113.65:57596
    GNMI mode                           : STREAM
    Subscription Time (UTC)             : Mon May 13 08:37:21 2024

    Sensor Statistics :
        Sensor Path                                     : /interfaces/interface[name='ge-0/0/2']/
subinterfaces/subinterface/
        Reporting Interval                              : 10
        Node-name                                       : r0_re0,r0_re0,tt-ygadiraju-15341379-
vm,tt-ygadiraju-15341379-vm
        Component(s)                                    : PFE,na-
grpcd,mib2d,rpd,xmlproxyd_TM_Thread_1
        GNMI Sub Mode                                   : SAMPLE
        Bytes Sent                                      : 3219
        Packets Sent                                    : 7
        Drops                                           : 0
        Average iLatency (ms)                           : 3
        Average Circular Buffer Used (%)                : 0
        Initial Sync Bytes Sent                         : 5642
        Initial Sync Packets Sent                       : 5
        Initial Sync Drops                              : 0
        Initial Sync Average iLatency (ms)              : 3
        Initial Sync Average Circular Buffer Used (%)   : 0
```

```
        Child Sensor Statistics :
            Path                                : /interfaces/interface[name='ge-0/0/2']/
subinterfaces/subinterface/
            Node-name                           : r0_re0
            Component                           : na-grpcd
            GNMI-SubMode                        : SAMPLE
            Bytes Sent                          : 0
            Packets Sent                        : 0
            Drops                               : 0
            Average iLatency (ms)               : 0
            Initial Sync Bytes Sent             : 1598
            Initial Sync Packets Sent           : 1
            Initial Sync Drops                  : 0
            Initial Sync Average iLatency (ms) : 2

        Child Sensor Statistics :
            Path                                : /junos/system/linecard/interface/logical/usage/
            Node-name                           : r0_re0
            Component                           : PFE
            Component-ID                        : 0
            GNMI-SubMode                        : SAMPLE
            Bytes Sent                          : 0
            Packets Sent                        : 1
            Drops                               : 0
            Average iLatency (ms)               : 3
            Initial Sync Bytes Sent             : 1598
            Initial Sync Packets Sent           : 1
            Initial Sync Drops                  : 0
            Initial Sync Average iLatency (ms) : 2

        Child Sensor Statistics :
            Path                                : /interfaces/interface[name='ge-0/0/2']/
subinterfaces/subinterface/
            Node-name                           : tt-ygadiraju-15341379-vm
            Component                           : mib2d
            GNMI-SubMode                        : SAMPLE
            Bytes Sent                          : 2727
            Packets Sent                        : 4
            Drops                               : 0
            Average iLatency (ms)               : 2
            Initial Sync Bytes Sent             : 0
            Initial Sync Packets Sent           : 0
```

```
            Initial Sync Drops                : 0
            Initial Sync Average iLatency (ms) : 0

        Child Sensor Statistics :
            Path                              : /interfaces/interface[name='ge-0/0/2']/
subinterfaces/subinterface/
            Node-name                         : tt-ygadiraju-15341379-vm
            Component                         : rpd
            GNMI-SubMode                      : SAMPLE
            Bytes Sent                        : 237
            Packets Sent                      : 1
            Drops                             : 0
            Average iLatency (ms)             : 7
            Initial Sync Bytes Sent           : 0
            Initial Sync Packets Sent         : 0
            Initial Sync Drops                : 0
            Initial Sync Average iLatency (ms) : 0

        Child Sensor Statistics :
            Path                              : /interfaces/interface[name='ge-0/0/2']/
subinterfaces/subinterface/
            Node-name                         : tt-ygadiraju-15341379-vm
            Component                         : xmlproxyd_TM_Thread_1
            GNMI-SubMode                      : SAMPLE
            Bytes Sent                        : 255
            Packets Sent                      : 1
            Packets Sent                      : 4
            Drops                             : 0
            Average iLatency (ms)             : 2
            Initial Sync Bytes Sent           : 0
            Initial Sync Packets Sent         : 0
            Initial Sync Drops                : 0
            Initial Sync Average iLatency (ms) : 0

        Child Sensor Statistics :
            Path                              : /interfaces/interface[name='ge-0/0/2']/
subinterfaces/subinterface/
            Node-name                         : tt-ygadiraju-15341379-vm
            Component                         : rpd
            GNMI-SubMode                      : SAMPLE
            Bytes Sent                        : 237
            Packets Sent                      : 1
            Drops                             : 0
```

```
            Average iLatency (ms)              : 7
            Initial Sync Bytes Sent            : 0
            Initial Sync Packets Sent          : 0
            Initial Sync Drops                 : 0
            Initial Sync Average iLatency (ms) : 0

        Child Sensor Statistics :
            Path                               : /interfaces/interface[name='ge-0/0/2']/
subinterfaces/subinterface/
            Node-name                          : tt-ygadiraju-15341379-vm
            Component                          : xmlproxyd_TM_Thread_1
            GNMI-SubMode                       : SAMPLE
            Bytes Sent                         : 255
            Packets Sent                       : 1
            Drops                              : 0
            Average iLatency (ms)              : 0
            Initial Sync Bytes Sent            : 0
            Initial Sync Packets Sent          : 0
            Initial Sync Drops                 : 0
            Initial Sync Average iLatency (ms) : 0

        Child Sensor Statistics :
            Path                               : /interfaces/interface[name='ge-0/0/2']/
subinterfaces/subinterface/
            Node-name                          : tt-ygadiraju-15341379-vm
            Component                          : mib2d
            GNMI-SubMode                       : SAMPLE
            Bytes Sent                         : 0
            Packets Sent                       : 0
            Drops                              : 0
            Average iLatency (ms)              : 0
            Initial Sync Bytes Sent            : 2446
            Initial Sync Packets Sent          : 3
            Initial Sync Drops                 : 0
            Initial Sync Average iLatency (ms) : 4
```

show network-agent statistics (detailed output, subscription deletion in progress)

```
user@host> show network-agent statistics detail
Unable to receive the subscription details
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show agent sensors | 17024

show network-access aaa saml assertion-cache

IN THIS SECTION

- Syntax | 22560
- Description | 22560
- Options | 22561
- Required Privilege Level | 22561
- Sample Output | 22561
- Release Information | 22562

Syntax

```
show network-access aaa saml assertion-cache
<brief | detail>
<domain-name domain-name>
<username username>
```

Description

Display Security Assertion Markup Language (SAML) authentication assertion cache information.

Options

brief | detail (Optional) Display the specified level of output.

- Default: brief

domain-name *domain-name* Specify the domain name such as example.net or any. This domain must match the domain configured in the profile at `set access profile demo saml identity-provider`.

username *username* Specify the SAML authenticated username. This username is the same as the user login ID.

Required Privilege Level

view

Sample Output

show network-access aaa saml assertion-cache

```
user@host> show network-access aaa saml assertion-cache
  Username            Remaining Validity
  test1@example.com   4m 47s
  user1@example.com    9h 20m 10s
  user2@example.net    1d 1h 5m 20s
  user@example.net     3h 40s
```

show network-access aaa saml assertion-cache detail

```
user@host> show network-access aaa saml assertion-cache detail
  Username: test1@juniper.net
  Remaining Validity: 4m 41s
  group: srx-dept1
  firstname: test1
  lastname: jnpr
```

show network-access aaa saml assertion-cache username

```
user@host> show network-access aaa saml assertion-cache username user2@example.net
  Username: user2@example.net
  Remaining Validity: 4m 41s
  group: srx-dept1
  firstname: user
  lastname: name
```

show network-access aaa saml assertion-cache domain

```
user@host> show network-access aaa saml assertion-cache domain example.net
  Username                     Remaining Validity
  user@example.net             3h 40s
```

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

saml | 16281

clear network-access aaa saml assertion-cache | 3757

show nlsd mode

IN THIS SECTION

- Syntax | 22563
- Description | 22563
- Additional Information | 22563
- Required Privilege Level | 22563
- Sample Output | 22563

Release Information | 22563

Syntax

```
show nlsd mode
```

Description

The command `show nlsd mode` shows you if the netlink asynchronous (async) API is activated or deactivated.

Additional Information

Netlink async mode is deactivated by default. You can enable netlink async mode with the command `set system netlink-async-mode`.

Required Privilege Level

view

Sample Output

show nlsd mode

```
user@host> show nlsd mode
Netlink async mode status: activated
```

Release Information

Command introduced in Junos OS Evolved Release 20.1R1.

show node reboot

IN THIS SECTION

- Syntax | 22564
- Description | 22564
- Sample Output | 22564
- Release Information | 22564

Syntax

```
show node reboot node-name
```

Description

Use this command to display any pending reboot or shutdown operations for the specified nodes in your network.

Sample Output

show node reboot

```
user@host> show node reboot re0
reboot of node is scheduled at Mon Mar 7 09:33:13 2022
```

Release Information

Command introduced in Junos OS Evolved 19.2R1.

show node statistics

IN THIS SECTION

- Syntax | 22565
- Description | 22565
- Options | 22565
- Sample Output | 22566
- Release Information | 22570

Syntax

```
show node statistics protocol-name node-name
```

Description

Displays protocol statistics for nodes in the network. Protocol statistics help you monitor and troubleshoot your network.

Options

Table 895 on page 22565 lists the protocol options that you can execute as part of the `show node statistics` command to display specific statistics.

Table 895: Protocol Options

Protocol	Description
`icmp` *node-name*	Shows ICMP statistics for the selected node.
`icmpmsg` *node-name*	Shows ICMP message types for the selected node.
`ip` *node-name*	Displays IPv4 statistics for the selected node.

Table 895: Protocol Options *(Continued)*

Protocol	Description
`ipext` *node-name*	Shows detailed IPv4 statistics for the selected node.
`tcp` *node-name*	Displays TCP statistics for the selected node.
`tcpext` *node-name*	Shows extended TCP statistics for the selected node.
`udp` *node-name*	Displays UDP statistics for the selected node.
`udplite` *node-name*	Displays UDP Lite statistics for the selected node.

Sample Output

show node statistics ICMP

```
user@host> show node statistics icmp re0
------------------------------
node: re0
------------------------------
Icmp:
    15 ICMP messages received
    0 input ICMP message failed.
    ICMP input histogram:
        destination unreachable: 4
        echo requests: 9
        echo replies: 2
    19 ICMP messages sent
    0 ICMP messages failed
    ICMP output histogram:
        destination unreachable: 8
        echo request: 2
        echo replies: 9
```

show node statistics ICMPmsg

```
user@host> show node statistics icmpmsg re0
-------------------------------
node: re0
-------------------------------
IcmpMsg:
        InType0: 2
        InType3: 4
        InType8: 9
        OutType0: 9
        OutType3: 8
        OutType8: 2
```

show node statistics IP

```
user@host> show node statistics ip re0
-------------------------------
node: re0
-------------------------------
Ip:
    16866054 total packets received
    0 forwarded
    0 incoming packets discarded
    16349557 incoming packets delivered
    16250537 requests sent out
    9 dropped because of missing route
    643 fragments dropped after timeout
    25707155 reassemblies required
    1379681 packets reassembled ok
    55031 packet reassembles failed
```

show node statistics IPext

```
user@host> show node statistics ipext re0
-------------------------------
node: re0
-------------------------------
IpExt:
```

```
    InNoRoutes: 188
    InMcastPkts: 247566
    OutMcastPkts: 130799
    InBcastPkts: 114928
    InOctets: 4566984898
    OutOctets: 9703802978
    InMcastOctets: 95373650
    OutMcastOctets: 37490728
    InBcastOctets: 33839476
    InNoECTPkts: 17111859
```

show node statistics TCP

```
user@host> show node statistics tcp re0
-------------------------------
node: re0
-------------------------------
Tcp:
    4165 active connections openings
    3807 passive connection openings
    182 failed connection attempts
    1333 connection resets received
    2201 connections established
    16132952 segments received
    17229523 segments send out
    21228 segments retransmited
    0 bad segments received.
    1730 resets sent
```

show node statistics TCPext

```
user@host> show node statistics tcpext re0
-------------------------------
node: re0
-------------------------------
TcpExt:
    50 resets received for embryonic SYN_RECV sockets
    80 packets pruned from receive queue because of socket buffer overrun
    1524 TCP sockets finished time wait in fast timer
    2 packets rejects in established connections because of timestamp
```

```
    473665 delayed acks sent
    67 delayed acks further delayed because of locked socket
    Quick ack mode was activated 3604 times
    447222 packets directly queued to recvmsg prequeue.
    386029 bytes directly in process context from backlog
    1311470377 bytes directly received in process context from prequeue
    4076007 packet headers predicted
    445042 packets header predicted and directly queued to user
    4530199 acknowledgments not containing data payload received
    4987056 predicted acknowledgments
    1535 times recovered from packet loss by selective acknowledgements
    Detected reordering 23 times using FACK
    Detected reordering 87 times using SACK
    Detected reordering 60 times using time stamp
    28 congestion windows fully recovered without slow start
    59 congestion windows partially recovered using Hoe heuristic
    456 congestion windows recovered without slow start by DSACK
    9 congestion windows recovered without slow start after partial ack
    TCPLostRetransmit: 33
    47 timeouts after SACK recovery
    6155 fast retransmits
    36 forward retransmits
    40 retransmits in slow start
    72 other TCP timeouts
    TCPLossProbes: 14928
    TCPLossProbeRecovery: 792
    3604 DSACKs sent for old packets
    81 DSACKs sent for out of order packets
    13590 DSACKs received
    2 DSACKs for out of order packets received
    78 connections reset due to unexpected data
    14 connections reset due to early user close
    59 connections aborted due to timeout
    TCPDSACKIgnoredOld: 5
    TCPDSACKIgnoredNoUndo: 8387
    TCPSpuriousRTOs: 2
    TCPSackShifted: 1437
    TCPSackMerged: 4524
    TCPSackShiftFallback: 36906
    TCPRcvCoalesce: 150976
    TCPOFOQueue: 36979
    TCPOFOMerge: 81
    TCPSpuriousRtxHostQueues: 193
```

```
    TCPAutoCorking: 3400209
    TCPWantZeroWindowAdv: 400
    TCPSynRetrans: 3
    TCPOrigDataSent: 9679629
    TCPHystartTrainDetect: 186
    TCPHystartTrainCwnd: 3979
    TCPHystartDelayDetect: 2
    TCPHystartDelayCwnd: 589
    TCPWinProbe: 3
    TCPKeepAlive: 4310999
```

show node statistics UDP

```
user@host> show node statistics udp re0
-------------------------------
node: re0
-------------------------------
Udp:
    322005 packets received
    8 packets to unknown port received.
    0 packet receive errors
    233256 packets sent
    IgnoredMulti: 115962
```

show node statistics UDPLite

```
user@host> show node statistics udplite re0
-------------------------------
node: re0
-------------------------------
UdpLite:
```

Release Information

Command introduced in Junos OS Evolved Release 19.2R1.

show nonstop-routing

IN THIS SECTION

- Syntax | 22571
- Description | 22571
- Required Privilege Level | 22571
- Output Fields | 22571
- Sample Output | 22574
- Release Information | 22575

Syntax

```
show nonstop-routing
```

Description

Display the status of nonstop active routing that includes the automerge statistics and state.

Required Privilege Level

View

Output Fields

Table 896 on page 22571 describes the output fields for the **show nonstop-routing** command. Output fields are listed in the approximate order in which they appear.

Table 896: show nonstop-routing Output Fields

Field Name	Field Description
Nonstop Routing	State of NSR.

Table 896: show nonstop-routing Output Fields *(Continued)*

Field Name	Field Description
Precision Timers state	State of precision timer feature in the kernel. • **Enabled** By default, autokeepalive precision timers are enabled on the kernel after switchover. • **Disabled** Autokeepalive precision timers are disabled. • **Inactive** Precision timer is inactive if it is disabled. • **Ready** Kernel precision timer is ready but is never activated. • **InProcess** Kernel precision timer is operational and is generating keepalives on behalf of the RPD after switchover. The / count indicates the number of sessions being serviced against the total sessions. • **Completed** Kernel has completed keepalive generation for all the sessions after switchover, and RPD has taken over all of them successfully. • **Error** Error while retrieving the precision timer status of the kernel.
Precision Timers max period	Maximum period, in seconds, after the switchover from standby to primary event for which the kernel autogenerates keepalives on behalf of BGP.
Automerge	Status of the automerge. • **Active** Automerge of sockets by the kernel after switchover is active. • **Inactive** Automerge of sockets by the kernel after switchover is inactive.

Table 896: show nonstop-routing Output Fields *(Continued)*

Field Name	Field Description
Batching	Status of Batching. • **Yes** Automerge of sockets by the kernel after a switchover. • **No** Automerge of sockets by the kernel after switchover is inactive.
Batch count	Number of sockets merged per batch.
Batch count adjust	Speed at which the batch count is adjusted. • **Slow** Number of sockets merged per batch is incremented additively. • **Exp** Number of sockets merged per batch is incremented exponentially. • **None** Number of sockets merged per batch remains constant.
Batch interval	Time interval between batches of automerge operation.
Batch interval adjust	Speed at which the batch interval is adjusted. • **Exp** Time interval between automerge of batches is increased exponentially. • **None** Time interval between automerge of batches is not adjusted.

Table 896: show nonstop-routing Output Fields *(Continued)*

Field Name	Field Description
Automerge State	State of the automerge • **Ready** Ready to automerge socket pairs from secondary to primary routing engine • **InProgress** Kernel is performing automerge after switchover • **Switchover Completed** Sessions merged after switchover
Sessions Processed	Count of sessions that are automerged.

Sample Output

show nonstop-routing (MX Series Router)

```
user@host show nonstop-routing
Nonstop Routing : Enabled
  Precision Timers state: Enabled: Completed - 0/0
  Precision Timers max period: 200
  Automerge : Active
  Batching: No
  Batch count: 200
  Batch count adjust: Exponential
  Batch interval: 20 msec
  Batch interval adjust: None
  Automerge State: Ready
  Sessions Processed: 0
```

show nonstop-routing (MX Series Router)

```
user@host> show nonstop-routing

Nonstop Routing : Enabled
          Automerge : Active
```

```
            Batching: Yes
            Batch count: 500
            Batch count adjust: Slow
            Batch interval: 50 msec
            Batch interval adjust: None
            Automerge State: Ready
            Sessions Processed: 0
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

nonstop-routing | 11671

show nsagentd

IN THIS SECTION

- Syntax | 22575
- Description | 22576
- Required Privilege Level | 22576
- Sample Output | 22576
- Release Information | 22577

Syntax

```
show nsagentd (statistics | status)
```

Description

The `nsagentd` options show you the statistics and status of the nsagentd app. This app handles and resolves TCP-specific functionality, updates, and notifications when a TCP flow becomes stale.

Required Privilege Level

view

Sample Output

show nsagentd statistics

```
user@host> show nsagentd statistics
Invalid upcalls:              0
Ingress SYN upcalls:
  Upcalls:                    0
  FIB lookups:
    Attempted:                0
    Failed:                   0
    Succeeded:                0
      Result code not EOK:    0
      Result code EOK:        0
  Ioctl invocations:          0
    Attempted:                0
    Failed:                   0
    Succeeded:                0
Stale TCP session upcalls:
  Upcalls:                    0
  FIB lookups:
    Attempted:                0
    Failed:                   0
    Succeeded:                0
      Result code not EOK:    0
      Result code EOK:        0
  Ioctl invocations:
    Attempted:                0
    Failed:                   0
    Succeeded:                0
```

show nsagentd status

```
user@host> show nsagentd status
Local address:     224.0.0.4
Start time:        2022-04-06 08:05:27.535
Last upcall time:  Never
```

Release Information

Command introduced in Junos OS Evolved Release 21.3R1.

show ntp associations

IN THIS SECTION

- Syntax | 22577
- Description | 22577
- Options | 22578
- Required Privilege Level | 22578
- Output Fields | 22578
- Sample Output | 22579
- Release Information | 22580

Syntax

```
show ntp associations
<no-resolve>
```

Description

Display Network Time Protocol (NTP) peers and their state.

Options

none Display NTP peers and their state.

no-resolve (Optional) Suppress symbolic addressing.

Required Privilege Level

view

Output Fields

Table 897 on page 22578 describes the output fields for the `show ntp associations` command. Output fields are listed in the approximate order in which they appear.

Table 897: show ntp associations Output Fields

Field Name	Field Description
`remote`	Address or name of the remote NTP peer. Status of the peer in the clock selection process: • space—Discarded because of a high stratum value or failed sanity checks. • `x`—Designated "falseticker" by the intersection algorithm. • `.`—Culled from the end of the candidate list. • `-` —Discarded by the clustering algorithm. • `+`—Included in the final selection set. • `#`—Selected for synchronization, but the distance exceeds the maximum. • `*`—Selected for synchronization. • `o`—Selected for synchronization, but the packets-per-second (pps) signal is in use.
`refid`	Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of `0.0.0.0`.
`auth`	Display the authentication method used for the peers.

Table 897: show ntp associations Output Fields *(Continued)*

Field Name	Field Description
st	Stratum of the remote peer.
t	Type of peer: b (broadcast), l (local), m (multicast), or u (unicast).
when	When the last packet from the peer was received.
poll	Polling interval, in seconds.
reach	Reachability register, in octal.
delay	Current estimated delay of the peer, in milliseconds.
offset	Current estimated offset of the peer, in milliseconds.
disp	Current estimated dispersion of the peer, in milliseconds.
rootdelay	Unsigned fixed-point number indicating the total round-trip delay to the reference clock.
rootdisp	Unsigned fixed-point number indicating the maximum error relative to the reference clock.

Sample Output

show ntp associations

```
root@R1> show ntp associations
     remote               refid           auth  st  t  when  poll reach  delay     offset
jitter rootdelay rootdisp
=================================================================================================
====================
*anydns.jnpr.net        172.16.0.1           -   3  u    55    64  377  258.017    -2.852
0.909   23.010    5.112
```

show ntp associations | display xml (Starting in Junos OS and Junos OS Evolved Release 22.4R1)

```
<user@host> show ntp associations | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/22.4I0/junos">
    <ntp-associations>
        <association>
            <remote>*anyntp.juniper.net</remote>
            <refid>172.16.0.1</refid>
            <auth>-</auth>
            <st>3</st>
            <association-type>u</association-type>
            <when>55</when>
            <poll>64</poll>
            <reach>377</reach>
            <delay>258.017</delay>
            <offset>-2.852</offset>
            <jitter>0.909</jitter>
            <rootdelay>23.010</rootdelay>
            <rootdisp>5.112</rootdisp>
        </association>
    </ntp-associations>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ntp status | 22581

show ntp status

IN THIS SECTION

- Syntax | 22581
- Description | 22581
- Options | 22581
- Required Privilege Level | 22581
- Output Fields | 22582
- Sample Output | 22584
- Release Information | 22585

Syntax

```
show ntp status
<no-resolve>
```

Description

Display the values of internal variables returned by Network Time Protocol (NTP) peers.

Options

none Display the values of internal variables returned by NTP peers.

no-resolve (Optional) Suppress symbolic addressing.

Required Privilege Level

view

Output Fields

Table 898 on page 22582 describes the output fields for the `show ntp status` command. Output fields are listed in the approximate order in which they appear.

Table 898: show ntp status Output Fields

Field Name	Field Description
`status`	System status word, a code representing the status items listed.
`leap_none`	Indicates a normal synchronized state with no leap seconds imminent. Other options could be `leap_add_sec`, `leap_del_sec`, or `leap_alarm`, indicating a leap second will be added, deleted, or a leap second requirement is upcoming.
`sync_ntp`	Indicates the current synchronization source, in this case, an NTP server. Other options include `sync_alarm` and `sync_unspec`, both indicating that the router has not been synched.
`x events`	Indicates the number of events that have occurred since that last code change. An event is often the receipt of an NTP polling message.
`event_peer/strat_chg`	Describes the most recent event, in this case, the stratum of the peer server changed.
`version`	A detailed description of the version of NTP being used.
`processor`	Indicates the current hardware platform and version of the processor.
`system`	Detailed description of the name and version of the operating system in use.
`leap`	The number of leap seconds in use.
`stratum`	The stratum of the peer server. Anything greater than 1 is a secondary reference source, and the number roughly represents the number of hops away from the stratum 1 server.. Stratum 1 is a primary reference, such as an atomic clock.
`precision`	The precision of the peer clock, how precisely the frequency and time can be maintained with this particular timekeeping system.

Table 898: show ntp status Output Fields *(Continued)*

Field Name	Field Description
`rootdelay`	The total roundtrip delay to the primary reference source, in seconds.
`rootdispersion`	The maximum error relative to the primary reference source, in seconds.
`peer`	An identification number of the peer in use.
`refid`	Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of `0.0.0.0`.
`reftime`	The local time, in timestamp format, when the local clock was last updated. If the local clock has never been synchronized, the value is zero.
`poll`	The NTP broadcast message polling interval, in seconds.
`clock`	The current time on the local router clock.
`state`	The current mode of NTP operation, where 1 is symmetric active, 2 is symmetric passive, 3 is client, 4 is server, and 5 is broadcast.
`offset`	Current estimated offset of the peer, in milliseconds. Indicates the time difference between the reference clock and the local clock.
`frequency`	The frequency of the clock.
`jitter`	Indicates the magnitude of jitter, in milliseconds, between several time queries.
`stability`	A measure of how well this clock can maintain a constant frequency.

Sample Output

show ntp status

```
user@host> show ntp status
associd=0 status=0615 leap_none, sync_ntp, 1 event, clock_sync,
version="ntpd 4.2.8p15-a Sun Oct  2 12:41:44 PDT 2022 (1)",
processor="amd64", system="FreeBSDJNPR-12.1-20220929.32ed862a0f7_",
leap=00, stratum=4, precision=-21, rootdelay=25.091, rootdisp=29.386,
refid=172.16.0.1,
reftime=e7290692.310f9fda  Wed, Nov 23 2022 12:52:34.191,
clock=e7290819.e3d61bef  Wed, Nov 23 2022 12:59:05.889, peer=45960,
tc=10, mintc=3, offset=-0.294143, frequency=-14.096, sys_jitter=0.000000,
clk_jitter=0.169, clk_wander=0.020
```

show ntp status | display xml (Starting in Junos OS and Junos OS Evolved Release 22.4R1)

```
<user@host> show ntp status | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/22.4I0/junos">
    <ntp-status>
        <associd>0</associd>
        <statusinfo>0615 leap_none, sync_ntp, 1 event, clock_sync</statusinfo>
        <version>"ntpd 4.2.8p15-a Sun Oct  2 12:41:44 PDT 2022 (1)"</version>
        <processor>"amd64"</processor>
        <system>"FreeBSDJNPR-12.1-20220929.32ed862a0f7_"</system>
        <leap>00</leap>
        <stratum>4</stratum>
        <precision>-21</precision>
        <rootdelay>25.091</rootdelay>
        <rootdisp>27.346</rootdisp>
        <refid>172.16.0.1</refid>
        <reftime>e7290692.310f9fda  Wed, Nov 23 2022 12:52:34.191</reftime>
        <clock>e7290791.f3de0d45  Wed, Nov 23 2022 12:56:49.952</clock>
        <peer>45960</peer>
        <tc>10</tc>
        <mintc>3</mintc>
        <offset>-0.294143</offset>
        <frequency>-14.096</frequency>
        <sys-jitter>0.000000</sys-jitter>
        <clk-jitter>0.169</clk-jitter>
```

```
        <clk-wander>0.020</clk-wander>
    </ntp-status>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show ntp associations | 22577

show oam ethernet connectivity-fault-management delay-statistics

IN THIS SECTION

- Syntax | 22585
- Description | 22586
- Options | 22586
- Required Privilege Level | 22586
- Output Fields | 22586
- Sample Output | 22588
- Release Information | 22590

Syntax

```
show oam ethernet connectivity-fault-management delay-statistics
<count entry-count>
<local-mep mep-id>
<maintenance-association ma-name>
```

```
<maintenance-domain md-name>
<remote-mep remote-mep-id>
```

Description

On MX Series routers with Ethernet interfaces on Dense Port Concentrators (DPCs), display ETH-DM delay statistics.

On EX Series switches, display delay measurement results.

Options

count ***entry-count***	(Optional) Number of entries to display from the statistics table. The range of values is **1** through **100**. The default value is **100** entries.
local-mep ***mep-id***	(Optional) Numeric identifier of the local MEP. On MX Series routers, the range of values is **1** through **8192**. On EX Series switches, the range of values is **1** through **8191**.
maintenance-association ***ma-name***	Name of an existing CFM maintenance association.
maintenance-domain ***md-name***	Name of an existing connectivity fault management (CFM) maintenance domain.
remote-mep ***remote-mep-id***	(Optional) Numeric identifier of the remote MEP. On MX Series routers, the range of values is **1** through **8192**. On EX Series switches, the range of values is **1** through **8191**.

Required Privilege Level

view

Output Fields

Table 899 on page 22587 lists the output fields for the `show oam ethernet connectivity-fault-management delay-statistics` command and the `show oam ethernet connectivity-fault-management mep-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 899: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields

Output Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) numeric identifier.
MAC address	Unicast MAC address configured for the MEP.
Remote MEP count	Number of remote MEPs (unless you specify the **remote-mep** option).
Remote MEP identifier	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Index	Index number that corresponds to the ETH-DM entry in the CFM database.
One-way delay (usec)	For a one-way ETH-DM session, the frame delay time, in microseconds, measured at the receiver MEP. For a detailed description of one-way Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the Junos OS Network Interfaces Library for Routing Devices.
Two-way delay (usec)	For a two-way ETH-DM session, the frame delay time, in microseconds, measured at the initiator MEP. For a detailed description of two-way Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the Junos OS Network Interfaces Library for Routing Devices.
Average one-way delay	Average one-way frame delay for the statistics displayed.
Average one-way delay variation	Average one-way “frame jitter” for the statistics displayed.
Best-case one-way delay	Lowest one-way frame delay for the statistics displayed.

Table 899: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields *(Continued)*

Output Field Name	Field Description
Worst-case one-way delay	Highest one-way frame delay for the statistics displayed.
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.
Best-case two-way delay	Lowest two-way frame delay for the statistics displayed.
Worst-case two-way delay	Highest two-way frame delay calculated in this session.

Sample Output

show oam ethernet connectivity-fault- management delay-statistics

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics maintenance-domain
md6 maintenance-association ma6
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 2
  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        259            519
      2        273            550
      3        287            571
      4        299            610
      5        313            650
    Average one-way delay          : 286 usec
    Average one-way delay variation: 62 usec
    Best case one-way delay        : 259 usec
    Worst case one-way delay       : 313 usec
    Average two-way delay          : 580 usec
```

```
    Average two-way delay variation: 26 usec
    Best case two-way delay        : 519 usec
    Worst case two-way delay       : 650 usec

  Remote MEP identifier: 102
  Remote MAC address: 00:04:55:63:39:5a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        29            58
      2        23            59
      3        27            56
      4        29            62
      5        33            68
    Average one-way delay          : 28 usec
    Average one-way delay variation: 3 usec
    Best case one-way delay        : 23 usec
    Worst case one-way delay       : 33 usec
    Average two-way delay          : 60 usec
    Average two-way delay variation: 3 usec
    Best case two-way delay        : 56 usec
    Worst case two-way delay       : 68 usec
```

show oam ethernet connectivity-fault- management delay-statistics remote-mep

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics maintenance-domain
md6 maintenance-association ma6 remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39

  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        259            519
      2        273            550
      3        287            571
      4        299            610
      5        313            650
    Average one-way delay          : 286 usec
    Average one-way delay variation: 62 usec
```

```
    Best case one-way delay        : 259 usec
    Worst case one-way delay       : 313 usec
    Average two-way delay          : 580 usec
    Average two-way delay variation: 26 usec
    Best case two-way delay        : 519 usec
    Worst case two-way delay       : 650 usec
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

- clear oam ethernet connectivity-fault-management statistics | 3775
- clear oam ethernet connectivity-fault-management delay-statistics | 3763
- show oam ethernet connectivity-fault-management interfaces | 22596
- show oam ethernet connectivity-fault-management mep-database | 22620
- show oam ethernet connectivity-fault-management mep-statistics | 22637

show oam ethernet connectivity-fault-management forwarding-state

IN THIS SECTION

- Syntax | 22591
- Description | 22591
- Options | 22591
- Required Privilege Level | 22591
- Output Fields | 22591
- Sample Output | 22592
- Release Information | 22596

Syntax

```
show oam ethernet connectivity-fault-management forwarding-state
interface interface-name | instance instance-name
brief | detail | extensive
```

Description

Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management forwarding state information for Ethernet interfaces.

Options

interface *interface-name* Display forwarding state information for the specified Ethernet interface only.

instance *instance-name* Display forwarding state information for the specified forwarding instance only.

brief | detail | extensive (Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 900 on page 22591 lists the output fields for the `show oam ethernet connectivity-fault-management forwarding-state` command. Output fields are listed in the approximate order in which they appear.

Table 900: show oam ethernet connectivity-fault-management forwarding-state Output Fields

Field Name	Field Description	Level of Output
Interface name	Interface identifier.	All levels
Link (Status)	Local link status.	All levels

Table 900: show oam ethernet connectivity-fault-management forwarding-state Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Filter action	Filter action for messages at the level.	All levels
Next hop type	Next-hop type.	All levels
Next index	Next-hop index number.	**brief**
Level	Maintenance domain (MD) level.	**detail**
Direction	MEP direction configured.	none
Instance name	Forwarding instance name.	All levels
CEs	Number of customer edge (CE) interfaces.	All levels
VEs	Number of VPN endpoint (VE) interfaces.	All levels

Sample Output

show oam ethernet connectivity-fault- management forwarding- state

```
user@host> show oam ethernet connectivity-fault-management forwarding-state
CEs: 3

  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
  3                       Drop              none
  4                       Drop              none
```

```
  5                       Drop              none
  6                       Drop              none
  7                       Drop              none
```

show oam ethernet connectivity-fault- management forwarding- state instance

```
user@host> show oam ethernet connectivity-fault-management forwarding-state instance
Instance name: __+bd1__
CEs: 3
VEs: 0
  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
  3                       Drop              none
  4                       Drop              none
  5                       Drop              none
  6                       Drop              none
  7                       Drop              none
```

show oam ethernet connectivity-fault- management forwarding- state interface

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface
Interface name: ge-3/0/0.0
Instance name: __+bd1__
  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
  3                       Drop              none
  4                       Drop              none
  5                       Drop              none
  6                       Drop              none
  7        down           Receive           none
```

```
Interface name: xe-0/0/0.0
Instance name: __+bd1__
  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
  3                       Drop              none
  4                       Drop              none
  5                       Drop              none
  6                       Drop              none
  7        down           Receive           none
```

show oam ethernet connectivity-fault- management forwarding- state interface detail

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface detail
Interface name: ge-3/0/0.0
Instance name: __+bd1__

  Level: 0
  Filter action: Drop
  Nexthop type: none

  Level: 1
  Filter action: Drop
  Nexthop type: none

  Level: 2
  Filter action: Drop
  Nexthop type: none

  Level: 3
  Filter action: Drop
  Nexthop type: none

  Level: 4
  Filter action: Drop
  Nexthop type: none
```

```
  Level: 5
  Filter action: Drop
  Nexthop type: none

  Level: 6
  Filter action: Drop
  Nexthop type: none

  Level: 7
  Direction: down
  Filter action: Receive
  Nexthop type: none

Interface name: xe-0/0/0.0
Instance name: __+bd1__

  Level: 0
  Filter action: Drop
  Nexthop type: none

  Level: 1
  Filter action: Drop
  Nexthop type: none

  ...
```

show oam ethernet connectivity-fault- management forwarding- state interface interface-name

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface interface-
name ge-3/0/0/0.0
Interface name: ge-3/0/0.0
Instance name: __+bd1__
  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
```

```
3                       Drop              none
4                       Drop              none
5                       Drop              none
6                       Drop              none
7         down          Receive           none
```

Release Information

Command introduced in Junos OS Release 8.4.

show oam ethernet connectivity-fault-management interfaces

IN THIS SECTION

- Syntax | 22596
- Description | 22597
- Options | 22597
- Required Privilege Level | 22597
- Output Fields | 22597
- Sample Output | 22604
- Release Information | 22611

Syntax

```
show oam ethernet connectivity-fault-management interfaces
ethernet-interface-name
level md-level
brief | detail | extensive
```

Description

On M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M320, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for Ethernet interfaces.

In addition, for Ethernet interfaces on MX Series routers , also display any ITU-T Y.1731 frame delay measurement (ETH-DM) frame counts when **detail** or **extensive** mode is specified.

For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.

Options

brief | detail | extensive (Optional) Specified level of output.

ethernet-interface-name (Optional) CFM information only for CFM entities attached to the specified Ethernet interface.

level *md-level* (Optional) CFM information for CFM identities enclosed within a maintenance domain of the specified level.

Required Privilege Level

view

Output Fields

Table 901 on page 22597 lists the output fields for the `show oam ethernet connectivity-fault-management interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 901: show oam ethernet connectivity-fault-management interfaces Output Fields

Field Name	Field Description	Level of Output
`Interface`	Interface identifier.	All levels
`Interface status`	Local interface status.	All levels
`Link status`	Local link status. **Up**, **down**, or **oam-down**.	All levels

Table 901: show oam ethernet connectivity-fault-management interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Maintenance domain name`	Maintenance domain name.	**detail extensive**
`Format (Maintenance domain)`	Maintenance domain name format configured.	**detail extensive**
`Level`	Maintenance domain level configured.	All levels
`Maintenance association name`	Maintenance association name.	**detail extensive**
Format (Maintenance association)	Maintenance association name format configured.	**detail extensive**
Continuity-check status	Continuity-check status.	**detail extensive**
Ethernet-ais status	Status of alarm indication signal (AIS). `active` or `in-active`.	**detail extensive**
Interval	Continuity-check message interval.	**detail extensive**
Loss-threshold	Lost continuity-check message threshold.	**detail extensive**
Interface status TLV	Status of the interface status TLV, if configured on the MEP interface: **none**, **up**, **down**, **testing**, **unknown**, **dormant**, **notPresent**, **lowerLayerDown**	**detail extensive**
Port status TLV	Status of the port status TLV, if configured on the MEP interface: **none**, **no**, **yes**	**detail extensive**

Table 901: show oam ethernet connectivity-fault-management interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Connection Protection TLV	Status of the connection protection TLV if configured on the MEP interface: **no**, **yes** If **yes**, then the transmitted connection protection TLV is decoded and the following three fields are displayed: **Prefer me**, **Protection in use**, **FRR Flag**	**detail extensive**
Prefer me	If set to **yes**, the path through which CCM was transmitted is preferred (unless the path fails). It is used for signaling a manual-switch command to the remote side. Its value can be **yes** or **no**.	**detail extensive**
Protection in use	Used for protection decision coordination. Its value is set to **yes** if the endpoint transmitting the CCM is currently transmitting the user traffic to protection path. Its value can be **yes** or **no**.	**detail extensive**
FRR Flag	LSR/LER forwarding the CCM Frame into a bypass tunnel is set. Its value can be **yes** or **no**.	**detail extensive**
MEP identifier	Maintenance association end point (MEP) identifier.	All levels
Neighbors	Number of MEP neighbors.	All levels
Direction	MEP direction configured.	**detail extensive**
MAC address	MAC address configured for the MEP.	**detail extensive**
MEP status	Indicates the status of the connectivity fault management (CFM) protocol running on the MEP: **Running**, **inactive**, **disabled**, or **unsupported**.	**detail extensive**

Table 901: show oam ethernet connectivity-fault-management interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Remote MEP not receiving CCM	Whether the remote MEP is not receiving connectivity check messages (CCMs).	**detail extensive**
Erroneous CCM received	Whether erroneous CCMs have been received.	**detail extensive**
Cross-connect CCM received	Whether cross-connect CCMs have been received.	**detail extensive**
RDI sent by some MEP	Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.	**detail extensive**
Some remote MEP's MAC in error state	Indicates whether the remote MEP's MAC is in error state.	**detail extensive**
Alarm Indication Signal	Indicates whether the AIS is triggered or is cleared.	**detail extensive**
CCMs sent	Number of CCMs transmitted.	**detail extensive**
CCMs received out of sequence	Number of CCMs received out of sequence.	**detail extensive**
LBMs sent	Number of loopback request messages (LBMs) sent.	**detail extensive**
Valid in-order LBRs received	Number of loopback response messages (LBRs) received that were valid messages and in sequence.	**detail extensive**
Valid out-of-order LBRs received	Number of LBRs received that were valid messages and not in sequence.	**detail extensive**

Table 901: show oam ethernet connectivity-fault-management interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
LBRs received with corrupted data	Number of LBRs received that were corrupted.	**detail extensive**
LBRs sent	Number of LBRs transmitted.	**detail extensive**
LTMs sent	Linktrace messages (LTMs) transmitted.	**detail extensive**
LTMs received	Linktrace messages received.	**detail extensive**
LTRs sent	Linktrace responses (LTRs) transmitted.	**detail extensive**
LTRs received	Linktrace responses received.	**detail extensive**
Sequence number of next LTM request	Sequence number of next LTM request to be transmitted.	**detail extensive**
1DMs sent	If the interface is attached to an initiator MEP for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.	**detail extensive**
Valid 1DMs received	If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of valid 1DM frames received. For all other cases, this field displays **0**.	**detail extensive**
Invalid 1DMs received	If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of invalid 1DM frames received. For all other cases, this field displays **0**.	**detail extensive**

Table 901: show oam ethernet connectivity-fault-management interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Out of sync 1DMs received	If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of out-of-sync one-way delay measurement request packets received.	**detail extensive**
DMMs sent	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.	**detail extensive**
Valid DMMs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of valid two-way delay measurement request packets received.	**detail extensive**
Invalid DMMs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of invalid two-way delay measurement request packets received.	**detail extensive**
DMRs sent	If the interface is attached to a responder MEP for a two-way ETH-DM session: Number of delay measurement reply (DMR) frames sent. For all other cases, this field displays **0**.	**detail extensive**
Valid DMRs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of valid DMRs received. For all other cases, this field displays **0**.	**detail extensive**
Invalid DMRs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of invalid DMRs received. For all other cases, this field displays **0**.	**detail extensive**
LMM sent	If the interface is attached to an initiator MEP for a ETH-LM session: Number of loss measurement message (LMM) PDU frames sent to the peer MEP in this session.	**detail extensive**

Table 901: show oam ethernet connectivity-fault-management interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Valid LMM received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid loss measurement request packets received.	**detail extensive**
Invalid LMM received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid loss measurement request packets received.	**detail extensive**
LMR sent	If the interface is attached to a responder MEP for a ETH-LM session: Number of loss measurement reply (LMR) frames sent.	**detail extensive**
Valid LMR received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid LMR frames received.	**detail extensive**
Invalid LMR received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid LMR frames received.	**detail extensive**
Valid AIS frames transmitted	Number of valid AIS frames transmitted to the peer MEPs.	**detail extensive**
Valid AIS frames received	Number of valid AIS frames received from the peer MEPs.	**detail extensive**
`SLM sent`	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of synthetic loss measurement (SLM) request packets transmitted from the source MEP to the remote or destination MEP in this session.	`detail extensive`
`Valid SLM received`	If the interface is attached to a responder MEP for a ETH-SLM session: Number of valid SLM PDUs transmitted from the source MEP to the remote or destination MEP.	`detail extensive`
`Invalid SLM received`	If the interface is attached to a responder MEP for a ETH-SLM session: Number of invalid SLM PDUs transmitted from the source MEP to the remote or destination MEP.	`detail extensive`

Table 901: show oam ethernet connectivity-fault-management interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`SLR sent`	If the interface is attached to a responder MEP for a ETH-SLM session: Number detail extensive of synthetic loss reply (SLR) frames sent.	`detail extensive`
`Valid SLR received`	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of valid SLR PDUs that the source MEP received from the remote or destination MEP.	`detail extensive`
`Invalid SLR received`	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of invalid SLR PDUs that the source MEP received from the remote or destination MEP.	`detail extensive`
Remote MEP count	Number of remote MEPs.	**extensive**
Identifier (remote MEP)	MEP identifier of the remote MEP.	**extensive**
MAC address (remote MEP)	MAC address of the remote MEP.	**extensive**
State (remote MEP)	State of the remote MEP.	**extensive**
Interface (remote MEP)	Interface of the remote MEP.	**extensive**

Sample Output

show oam ethernet connectivity-fault- management interfaces

```
user@host> show oam ethernet connectivity-fault-management interfaces
Interface       Link       Status           Level       MEP          Neighbors
                                                        Identifier
ge-1/1/0.0      Up         Active           0           2            1
```

```
ge-1/1/0.1      Up         Active           0           2            1
ge-1/1/0.10     Up         Active           0           2            1
ge-1/1/0.100    Up         Active           0           2            1
ge-1/1/0.101    Up         Active           0           2            1
ge-1/1/0.102    Up         Active           0           2            1
ge-1/1/0.103    Up         Active           0           2            1
ge-1/1/0.104    Up         Active           0           2            1
ge-1/1/0.105    Up         Active           0           2            1
ge-1/1/0.106    Up         Active           0           2            1

...
```

show oam ethernet connectivity-fault- management interfaces detail

```
user@host> show oam ethernet connectivity-fault-management interfaces detail
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
  Maintenance domain name: md0, Format: string, Level: 5
  Maintenance association name: ma1, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
      MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
    MEP status: running
    Defects:
      Remote MEP not receiving CCM                  : no
      Erroneous CCM received                        : yes
      Cross-connect CCM received                    : no
      RDI sent by some MEP                          : yes
      Alarm Indication Signal                       : yes
    Statistics:
      CCMs sent                                     : 76
      CCMs received out of sequence                 : 0
      LBMs sent                                     : 0
      Valid in-order LBRs received                  : 0
      Valid out-of-order LBRs received              : 0
      LBRs received with corrupted data             : 0
      LBRs sent                                     : 0
      LTMs sent                                     : 0
      LTMs received                                 : 0
      LTRs sent                                     : 0
      LTRs received                                 : 0
      Sequence number of next LTM request           : 1320235363
```

```
      1DMs sent                                    : 0
      Valid 1DMs received                          : 0
      Invalid 1DMs received                        : 0
      DMMs sent                                    : 0
      DMRs sent                                    : 0
      Valid DMRs received                          : 0
      Invalid DMRs received                        : 0
      LMM sent                                     : 10
      Valid LMM received                           : 20
      Invalid LMM received                         : 0
      LMR sent                                     : 20
      Valid LMR received                           : 10
      Invalid LMR received                         : 0
      Valid AIS frames transmitted                 : 0
      Valid AIS frames received                        : 0
      SLM sent                                     : 10
      Valid SLM received                           : 20
      Invalid SLM received                         : 0
      SLR sent                                     : 20
      Valid SLR received                           : 10
      Invalid SLR received                         : 0
    Remote MEP count: 2
      Identifier    MAC address        State    Interface
       2001     00:90:69:0b:7f:71       ok    ge-5/2/9.0
       4001     00:90:69:0b:09:c5       ok    ge-5/2/9.0
```

show oam ethernet connectivity-fault- management interfaces detail (One-Way ETH-DM)

```
user@host show oam ethernet connectivity-fault-management interfaces detail
Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
  Maintenance domain name: md6, Format: string, Level: 6
    Maintenance association name: ma6, Format: string
    Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
        MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
    MEP status: running
    Defects:
      Remote MEP not receiving CCM                 : no
      Erroneous CCM received                       : no
      Cross-connect CCM received                   : no
      RDI sent by some MEP                         : no
    Statistics:
```

```
      CCMs sent                                       : 1590
      CCMs received out of sequence                   : 0
      LBMs sent                                       : 0
      Valid in-order LBRs received                    : 0
      Valid out-of-order LBRs received                : 0
      LBRs received with corrupted data               : 0
      LBRs sent                                       : 0
      LTMs sent                                       : 0
      LTMs received                                   : 0
      LTRs sent                                       : 0
      LTRs received                                   : 0
      Sequence number of next LTM request             : 1542035464

      1DMs sent                                       : 10
      Valid 1DMs received                             : 0
      Invalid 1DMs received                           : 0
      DMMs sent                                       : 0
      DMRs sent                                       : 0
      Valid DMRs received                             : 0
      Invalid DMRs received                           : 0
    Remote MEP count: 1
      Identifier    MAC address        State    Interface
        201     00:90:69:0a:43:94       ok    ge-0/2/5.0
```

show oam ethernet connectivity-fault- management interfaces detail (Connection Protection TLV Configured)

```
user@hostshow oam ethernet connectivity-fault-management interfaces detail

Interface name: xe-6/2/0.0 , Interface status: Active, Link status: Up
  Maintenance domain name: md6, Format: string, Level: 6
    Maintenance association name: ma6, Format: string
    Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
    Interface status TLV: none, Port status TLV: none
    Connection Protection TLV: yes
      Prefer me: no, Protection in use: no, FRR Flag: no
    MEP identifier: 1, Direction: down, MAC address: 00:19:e2:b1:14:30
    MEP status: running
    Defects:
      Remote MEP not receiving CCM                    : no
      Erroneous CCM received                          : no
```

```
      Cross-connect CCM received                    : no
      RDI sent by some MEP                          : no
      Some remote MEP's MAC in error state          : no
    Statistics:
      CCMs sent                                     : 225
      CCMs received out of sequence                 : 0
      LBMs sent                                     : 0
      Valid in-order LBRs received                  : 0
      Valid out-of-order LBRs received              : 0
      LBRs received with corrupted data             : 0
      LBRs sent                                     : 0
      LTMs sent                                     : 0
      LTMs received                                 : 0
      LTRs sent                                     : 0
      LTRs received                                 : 0
      Sequence number of next LTM request           : 1244305646

      1DMs sent                                     : 0
      Valid 1DMs received                           : 0
      Invalid 1DMs received                         : 0
      Out of sync 1DMs received                     : 0
      DMMs sent                                     : 0
      Valid DMMs received                           : 0
      Invalid DMMs received                         : 0
      DMRs sent                                     : 0
      Valid DMRs received                           : 0
      Invalid DMRs received                         : 0
      LMMs sent                                     : 0
      Valid LMMs received                           : 0
      Invalid LMMs received                         : 0
      LMRs sent                                     : 0
      Valid LMRs received                           : 0
      Invalid LMRs received                         : 0
    Remote MEP count: 1
      Identifier    MAC address        State    Interface
          2     00:90:69:7f:e4:30
```

show oam ethernet connectivity-fault- management interfaces extensive

```
user@host> show oam ethernet connectivity-fault-management interfaces extensive
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
```

```
 Maintenance domain name: md0, Format: string, Level: 5
   Maintenance association name: ma1, Format: string
   Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
   Interface status TLV: none, Port status TLV: none
   Connection Protection TLV: no
MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
   MEP status: running
   Defects:
     Remote MEP not receiving CCM                  : no
     Erroneous CCM received                        : yes
     Cross-connect CCM received                    : no
     RDI sent by some MEP                          : yes

     Alarm Indication Signal                       : yes
   Statistics:
     CCMs sent                                     : 76
     CCMs received out of sequence                 : 0
     LBMs sent                                     : 0
     Valid in-order LBRs received                  : 0
     Valid out-of-order LBRs received              : 0
     LBRs received with corrupted data             : 0
     LBRs sent                                     : 0
     LTMs sent                                     : 0
     LTMs received                                 : 0
     LTRs sent                                     : 0
     LTRs received                                 : 0
     Sequence number of next LTM request           : 1645032434

     1DMs sent                                     : 0
     Valid 1DMs received                           : 0
     Invalid 1DMs received                         : 0
     DMMs sent                                     : 0
     DMRs sent                                     : 0
     Valid DMRs received                           : 0
     Invalid DMRs received                         : 0

     Valid AIS frames transmitted                  : 0
     Valid AIS frames received                         : 0
     SLM sent                                      : 10
     Valid SLM received                            : 20
     Invalid SLM received                          : 0
     SLR sent                                      : 20
     Valid SLR received                            : 10
```

```
      Invalid SLR received                          : 0
    Remote MEP count: 2
      Identifier    MAC address        State    Interface
       2001     00:90:69:0b:7f:71       ok    ge-5/2/9.0
       4001     00:90:69:0b:09:c5       ok    ge-5/2/9.0
```

show oam ethernet connectivity-fault- management interfaces level

```
user@host> show oam ethernet connectivity-fault-management interfaces level 7
Interface       Link       Status           Level       MEP          Neighbors
                                                       Identifier
ge-3/0/0.0      Up         Active           7           201          0
xe-0/0/0.0      Up         Active           7           203          1
```

show oam ethernet connectivity-fault- management interfaces (trunk ports)

```
user@host> show oam ethernet connectivity-fault-management interfaces

Interface                   Link      Status       Level  MEP         Neighbors
                                                          Identifier
ge-4/0/1.0, vlan 100        Up        Active       5      100         0
ge-10/3/10.4091, vlan 4091  Down      Inactive     4      400         0
ge-4/0/0.0                  Up        Active       6      200         0


user@host> show oam ethernet connectivity-fault-management interfaces  ge-4/0/0.0

Interface                   Link      Status       Level  MEP         Neighbors
                                                          Identifier
ge-4/0/0.0                  Up        Active       6      200         0


user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/1.0 vlan 100

Interface                   Link      Status       Level  MEP         Neighbors
                                                          Identifier
ge-4/0/1.0, vlan 100        Up        Active       5      100         0
```

```
user@host> show oam ethernet connectivity-fault-management interfaces ge-10/3/10.4091 vlan 4091

Interface                   Link      Status       Level  MEP         Neighbors
                                                          Identifier
ge-10/3/10.4091, vlan 4091  Down      Inactive     4      400         0
```

Release Information

Command introduced in Junos OS Release 8.4.

Support for ITU-T Y.1731 frame delay measurement added in Junos OS Release 9.5.

Support for ITU-T Y.1731 Ethernet synthetic frame loss measurement (ETH-SLM) added in Junos OS Release 13.2 for ACX Series and MX Series routers.

RELATED DOCUMENTATION

clear oam ethernet connectivity-fault-management statistics | 3775

show oam ethernet connectivity-fault-management delay-statistics | 22585

show oam ethernet connectivity-fault-management mep-database | 22620

show oam ethernet connectivity-fault-management mep-statistics | 22637

Ethernet Alarm Indication

show oam ethernet connectivity-fault-management linktrace path-database

IN THIS SECTION

- Syntax | 22612
- Description | 22612
- Options | 22612
- Required Privilege Level | 22612
- Output Fields | 22612
- Sample Output | 22614

Release Information | 22615

Syntax

```
show oam ethernet connectivity-fault-management linktrace path-database mac-address maintenance-association
ma-name maintenance-domain md-name
```

Description

On M320, MX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management maintenance linktrace database information.

Options

mac-address	Display connectivity fault management path database information for the specified MAC address of the remote host.
maintenance-association *ma-name*	Display connectivity fault management path database information for the specified maintenance association.
maintenance-domain *md-name*	Display connectivity fault management path database information for the specified maintenance domain.

Required Privilege Level

view

Output Fields

Table 902 on page 22613 lists the output fields for the `show oam ethernet connectivity-fault-management linktrace path-database` command. Output fields are listed in the approximate order in which they appear.

Table 902: show oam ethernet connectivity-fault-management linktrace path-database Output Fields

Field Name	Field Description
`Linktrace to`	MAC address of the 802.1ag node to which the linktrace message is targeted.
`Interface`	Interface used by the local MEP to send the linktrace message (LTM).
`Maintenance Domain`	Maintenance domain identifier specified in the traceroute command.
`Maintenance Association`	Maintenance association identifier specified in the traceroute command.
`Level`	Maintenance domain level configured for the maintenance domain.
`Local Mep`	MEP identifier of the local MEP originating the linktrace.
`Hop`	Sequential hop count of the linktrace path.
`TTL`	Number of hops remaining in the linktrace message (LTM). The time to live (TTL) is decremented at each hop.
`Source MAC address`	MAC address of the 802.1ag node responding to the LTM or the source MAC address of the LTR.
`Next hop MAC address`	MAC address of the egress interface of the node to which the LTM is forwarded or the next-hop MAC address derived from the next egress identifier in the Egress-ID TLV of the LTR PDU.
`Transaction Identifier`	4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all maintenance domains. Use the transaction identifier to match an incoming linktrace responses (LTR), with a previously sent LTM.

Sample Output

show oam ethernet connectivity-fault- management linktrace path-database

```
user@host> show oam ethernet connectivity-fault-management linktrace path-database
maintenance-domain MD1 maintenance-association MA1 00:01:02:03:04:05
Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
       Maintenance Domain: MD1, Level: 7
       Maintenance Association: MA1, Local Mep: 1

        Hop     TTL      Source MAC address        Next hop MAC address
       Transaction Identifier:100001
       1       63       00:00:aa:aa:aa:aa         00:00:ab:ab:ab:ab
       2       62       00:00:bb:bb:bb:bb         00:00:bc:bc:bc:bc
       3       61       00:00:cc:cc:cc:cc         00:00:cd:cd:cd:cd
       4       60       00:01:02:03:04:05         00:00:00:00:00:00
```

show oam ethernet connectivity-fault- management linktrace path-database (Two traceroute Commands)

```
user@host> traceroute ethernet maintenance-domain md1 maintenance-association ma1
00:01:02:03:04:05
 Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
       Maintenance Domain: MD1, Level: 7
       Maintenance Association: MA1, Local Mep: 1

       Hop     TTL      Source MAC address        Next hop MAC address
       Transaction Identifier:100002
       1       63       00:00:aa:aa:aa:aa         00:00:ab:ab:ab:ab
       2       62       00:00:bb:bb:bb:bb         00:00:bc:bc:bc:bc
       3       61       00:00:cc:cc:cc:cc         00:00:cd:cd:cd:cd
       4       60       00:01:02:03:04:05         00:00:00:00:00:00

       Transaction Identifier:100003
       1       63       00:00:aa:aa:aa:aa         00:00:ab:ab:ab:ab
       2       62       00:00:bb:bb:bb:bb         00:00:bc:bc:bc:bc
       3       61       00:00:cc:cc:cc:cc         00:00:cd:cd:cd:cd
       4       60       00:01:02:03:04:05         00:00:00:00:00:00
```

Release Information

Command introduced in Junos OS Release 9.0.

show oam ethernet connectivity-fault-management loss-statistics

IN THIS SECTION

- Syntax | 22615
- Description | 22615
- Options | 22616
- Required Privilege Level | 22616
- Output Fields | 22616
- show oam ethernet connectivity fault management loss statistics | 22618
- Release Information | 22620

Syntax

```
show oam ethernet connectivity-fault-management loss-statistics
maintenance-domain md-name
maintenance-association ma-name
<count entry-count>
<local-mep local-mep-id>
<remote-mep remote-mep-id>
```

Description

On MX Series and ACX series routers with Ethernet interfaces, display ETH-LM statistics for on-demand mode only.

Options

maintenance-domain ***md-name***	Name of an existing CFM maintenance domain.
maintenance-association ***ma-name***	Name of an existing CFM maintenance association.
count ***entry-count***	(Optional) Number of entries to display from the statistics table. The range of values is from 1 through 100. The default value is 100.
local-mep ***local-mep-id***	(Optional) Numeric identifier of the local MEP. The range of values is from 1 through 8191.
remote-mep ***remote-mep-id***	(Optional) Numeric identifier of the remote MEP. The range of values is from 1 through 8191.

Required Privilege Level

view

Output Fields

Table 903 on page 22616 lists the output fields for the `show oam ethernet connectivity-fault-management loss-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 903: show oam ethernet connectivity-fault-management loss-statistics Output Fields

Output Field Name	Field Description
`MEP identifier`	Maintenance association end point (MEP) numeric identifier.
`MAC address`	Unicast MAC address configured for the MEP.
`Remote MEP count`	Number of remote MEPs (unless you specify the `remote-mep` option).
`Remote MEP identifier`	Numeric identifier of the remote MEP.
`Remote MAC address`	Unicast MAC address of the remote MEP.

Table 903: show oam ethernet connectivity-fault-management loss-statistics Output Fields *(Continued)*

Output Field Name	Field Description
`Index`	Index number that corresponds to the ETH-LM entry in the CFM database.
`Near-end frame loss`	Count of frame loss associated with ingress data frames.
`Far-end frame loss`	Count of frame loss associated with egress data frames.
`Near-end loss ratio`	Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the ingress interface.
`Far-end loss ratio`	Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the egress interface.
`Average near-end frame loss`	Average frame loss measured in this session associated with ingress data frames.
`Average near-end loss ratio`	Average frame loss ratio measured in this session associated with ingress data frames.
`Average far-end frame loss`	Average frame loss measured in this session associated with egress data frames.
`Average far-end loss ratio`	Average frame loss ratio measured in this session associated with egress data frames.
`Near-end best case loss`	Lowest frame loss measured in this session associated with ingress data frames.
`Near-end best case loss ratio`	Lowest frame loss ratio measured in this session associated with ingress data frames.

Table 903: show oam ethernet connectivity-fault-management loss-statistics Output Fields *(Continued)*

Output Field Name	Field Description
Near-end worst case loss	Highest frame loss measured in this session associated with ingress data frames.
Near-end worst case loss ratio	Highest frame loss ratio measured in this session associated with ingress data frames.
Far-end best case frame loss	Lowest frame loss measured in this session associated with egress data frames.
Far-end best case loss ratio	Lowest frame loss ratio measured in this session associated with egress data frames.
Far-end worst case loss	Highest frame loss measured in this session associated with egress data frames.
Far-end worst case loss ratio	Highest frame loss ratio measured in this session associated with egress data frames.

show oam ethernet connectivity fault management loss statistics

command-name

```
user@host>. show oam ethernet connectivity fault management loss statistics maintenance-domain
md maintenance-association ma
MEP identifier: 1, MAC address: 64:87:88:f9:7d:1b
Remote MEP count: 1

  Remote MAC address: 64:87:88:6a:da:94

    LM client session-id:4843
    CIR Loss measurement statistics:
    Index  Near-end      Far-end      Near-end       Far-end       Near-end     Far-end
           Frame loss    Total tx     Total rx       Frame loss    Total tx     Total rx
```

```
          (CIR)         (CIR)         (CIR)          (CIR)         (CIR)        (CIR)
  1        0             245           245           0             244           244
  2        0             488           488           0             489           489
  3        0             732           732           0             733           733
  4        0             977           977           0             976           976
EIR Loss measurement statistics:
Index  Near-end      Far-end      Near-end       Far-end      Near-end      Far-end
       Frame loss    Total tx     Total rx       Frame loss    Total tx     Total rx
        (EIR)         (EIR)         (EIR)         (EIR)         (EIR)        (EIR)
  1        0             272           272           0             273           273
  2        0             546           546           0             545           545
  3        0             820           820           0             819           819
  4        0             1092          1092          0             1093          1093
Total far-end Tx (CIR)             : 977
Total near-end Rx (CIR)            : 977
Total near-end loss(CIR)           : 0
Total near-end loss ratio(CIR)     : 0.00000%
Total near-end Tx (CIR)            : 976
Total far-end Rx (CIR)             : 976
Total far-end loss(CIR)            : 0
Total far-end loss ratio(CIR)      : 0.00000%
Average near-end loss(CIR)         : 0.00000
Average near-end loss ratio(CIR)   : 0.00000%
Average far-end loss(CIR)          : 0.00000
Average far-end loss ratio(CIR)    : 0.00000%
Near-end best case loss(CIR)       : 0
Near-end best case loss ratio(CIR) : 0.00000%
Near-end worst case loss(CIR)      : 0
Near-end worst case loss ratio(CIR): 0.00000%
Far-end best case loss(CIR)        : 0
Far-end best case loss ratio(CIR)  : 0.00000%
Far-end worst case loss(CIR)       : 0
Far-end worst case loss ratio(CIR) : 0.00000%
Total far-end Tx (EIR)             : 1092
Total near-end Rx (EIR)            : 1092
Total near-end loss(EIR)           : 0
Total near-end loss ratio(EIR)     : 0.00000%
Total near-end Tx (EIR)            : 1093
Total far-end Rx (EIR)             : 1093
Total far-end loss(EIR)            : 0
Total far-end loss ratio(EIR)      : 0.00000%
Average near-end loss(EIR)         : 0.00000
Average near-end loss ratio(EIR)   : 0.00000%
```

```
    Average far-end loss(EIR)          : 0.00000
    Average far-end loss ratio(EIR)    : 0.00000%
    Near-end best case loss(EIR)       : 0
    Near-end best case loss ratio(EIR) : 0.00000%
    Near-end worst case loss(EIR)      : 0
    Near-end worst case loss ratio(EIR): 0.00000%
    Far-end best case loss(EIR)        : 0
    Far-end best case loss ratio(EIR)  : 0.00000%
    Far-end worst case loss(EIR)       : 0
    Far-end worst case loss ratio(EIR) : 0.00000%
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show oam ethernet connectivity-fault-management mep-statistics | 22637

show oam ethernet connectivity-fault-management mep-database

IN THIS SECTION

- Syntax | 22621
- Description | 22621
- Options | 22621
- Required Privilege Level | 22622
- Output Fields | 22622
- Sample Output | 22629
- Release Information | 22637

Syntax

```
show oam ethernet connectivity-fault-management mep-database
maintenance-domain domain-name
maintenance-association ma-name
<local-mep local-mep-id>
<remote-mep remote-mep-id>
extensive
```

Description

On M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M320, M120, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

In addition, on M120, M320, and MX series routers, also display port status TLV, interface status TLV, and action profile information.

In addition, for Ethernet interfaces on MX Series routers , also display any ITU-T Y.1731 frame delay measurement (ETH-DM) frame counts.

For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.

Options

maintenance-association *ma-name*	Display connectivity fault management information for the specified maintenance association.
maintenance-domain *domain-name*	Display connectivity fault management information for the specified maintenance domain.
local-mep *local-mep-id*	(Optional) Display connectivity fault management information for the specified local MEP only.
remote-mep *remote-mep-id*	(Optional) Display connectivity fault management information for the specified remote MEP only.

Required Privilege Level

view

Output Fields

Table 904 on page 22622 lists the output fields for the `show oam ethernet connectivity-fault-management mep-database` command. Output fields are listed in the approximate order in which they appear.

Table 904: show oam ethernet connectivity-fault-management mep-database Output Fields

Field Name	Field Description
Maintenance domain name	Maintenance domain name.
Format (Maintenance domain)	Maintenance domain name format configured.
Level	Maintenance domain level configured.
Maintenance association name	Maintenance association name.
Format (Maintenance association)	Maintenance association name format configured.
Continuity-check status	Continuity-check status.
Interval	Continuity-check message interval.
Loss-threshold	Lost continuity-check message threshold.

Table 904: show oam ethernet connectivity-fault-management mep-database Output Fields *(Continued)*

Field Name	Field Description
Connection Protection TLV	Status of the connection protection TLV, if configured on the MEP interface: **no**, **yes** If **yes**, then the transmitted connection protection TLV is decoded and the following three fields are displayed: **Prefer me**, **Protection in use**, **FRR Flag**
Prefer me	If set to **yes**, the path through which CCM was transmitted is preferred (unless the path fails). It is used for signaling a manual-switch command to remote side. Its value can be **yes** or **no**.
Protection in use	Used for protection decision coordination. Its value is set to **yes** if the endpoint transmitting the CCM is currently transmitting the user traffic to protection path. Its value can be **yes** or **no**.
FRR Flag	LSR/LER forwarding the CCM Frame into a bypass tunnel is set. Its value can be **yes** or **no**.
MEP identifier	Maintenance association end point (MEP) identifier.
Direction	MEP direction configured.
MAC address	MAC address configured for the MEP.
Auto-discovery	Whether automatic discovery is enabled or disabled.
Priority	Priority used for CCMs and linktrace messages transmitted by the MEP.
Interface name	Interface identifier.
Interface status	Local interface status.
Link status	Local link status.

Table 904: show oam ethernet connectivity-fault-management mep-database Output Fields *(Continued)*

Field Name	Field Description
Remote MEP not receiving CCM	Whether the remote MEP is not receiving CCMs.
Erroneous CCM received	Whether erroneous CCMs have been received.
Cross-connect CCM received	Whether cross-connect CCMs have been received.
RDI sent by some MEP	Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.
CCMs sent	Number of CCMs transmitted.
CCMs received out of sequence	Number of CCMs received out of sequence.
LBMs sent	Number of loopback messages (LBMs) sent.
Valid in-order LBRs received	Number of loopback response messages (LBRs) received that were valid messages and in sequence.
1DMs sent	If the MEP is an initiator for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.
Valid 1DMs received	If the MEP is a receiver for a one-way ETH-DM session: Number of valid 1DM frames received. For all other cases, this field displays **0**.

Table 904: show oam ethernet connectivity-fault-management mep-database Output Fields *(Continued)*

Field Name	Field Description
Invalid 1DMs received	If the MEP is a receiver for a one-way ETH-DM session: Number of invalid 1DM frames received. For all other cases, this field displays **0**.
Out of sync 1DMs received	If the MEP is a receiver for a one-way ETH-DM session: Number of out-of-sync one-way delay measurement request packets received.
DMMs sent	If the MEP is an initiator for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.
Valid DMMs received	If the MEP is an initiator for a two-way ETH-DM session: Number of valid two-way delay measurement packets received.
Invalid DMMs received	If the MEP is an initiator for a two-way ETH-DM session: Number of invalid two-way delay measurement packets received.
DMRs sent	If the MEP is a responder for a ETH-DM session: Number of Delay Measurement Reply (DMR) frames sent. For all other cases, this field displays **0**.
Valid DMRs received	If the MEP is an initiator for a two-way ETH-DM session: Number of valid DMRs received. For all other cases, this field displays **0**.
Invalid DMRs received	If the MEP is an initiator for a two-way ETH-DM session: Number of invalid DMRs received. For all other cases, this field displays **0**.
Valid out-of-order LBRs received	Number of LBRs received that were valid messages and not in sequence.

Table 904: show oam ethernet connectivity-fault-management mep-database Output Fields *(Continued)*

Field Name	Field Description
LBRs received with corrupted data	Number of LBRs received that were corrupted.
LBRs sent	Number of LBRs transmitted.
LTMs sent	Linktrace messages (LTMs) transmitted.
LTMs received	Linktrace messages received.
LTRs sent	Linktrace responses (LTRs) transmitted.
LTRs received	Linktrace responses received.
Sequence number of next LTM request	Sequence number of the next linktrace message request to be transmitted.
LMM sent	If the interface is attached to an initiator MEP for a ETH-LM session: Number of loss measurement message (LMM) PDU frames sent to the peer MEP in this session.
Valid LMM received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid loss measurement request packets received.
Invalid LMM received	If the interface is attached to an initiator MEP for a ETH LM session: Number of invalid loss measurement request packets received.
LMR sent	If the interface is attached to a responder MEP for a ETH-LM session: Number of loss measurement reply (LMR) frames sent.
Valid LMR received	If the interface is attached to an initiator MEP for a ETH LM session: Number of valid LMR frames received.
Invalid LMR received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid LMR frames received.

Table 904: show oam ethernet connectivity-fault-management mep-database Output Fields *(Continued)*

Field Name	Field Description
`SLM sent`	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of synthetic loss measurement (SLM) request packets transmitted from the source MEP to the remote or destination MEP in this session.
`Valid SLM received`	If the interface is attached to a responder MEP for a ETH-SLM session: Number of valid SLM PDUs transmitted from the source MEP to the remote or destination MEP.
`Invalid SLM received`	If the interface is attached to a responder MEP for a ETH-SLM session: Number of invalid SLM PDUs transmitted from the source MEP to the remote or destination MEP.
`SLR sent`	If the interface is attached to a responder MEP for a ETH-SLM session: Number detail extensive of synthetic loss reply (SLR) frames sent.
`Valid SLR received`	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of valid SLR PDUs that the source MEP received from the remote or destination MEP.
`Invalid SLR received`	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of invalid SLR PDUs that the source MEP received from the remote or destination MEP.
Remote MEP identifier	MEP identifier of the remote MEP.
State (remote MEP)	State of the remote MEP: **idle**, **start**, **ok**, or **failed**.
MAC address	MAC address of the remote MEP.
Type	Whether the remote MEP MAC address was learned using automatic discovery or configured.
Interface	Interface of the remote MEP. A seven-digit number is appended if CFM is configured to run on a routing instance of type VPLS.

Table 904: show oam ethernet connectivity-fault-management mep-database Output Fields *(Continued)*

Field Name	Field Description
Last flapped	Date, time, and how long ago the remote MEP interface went from down to up. The format is **Last flapped:** ***year-month-day hours.minutes.seconds timezone*** **(*hours.minutes.seconds* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.
Remote defect indication	Whether the remote defect indication (RDI) bit is set in messages that have been received or transmitted.
Port status TLV	• In the Maintenance domain section, displays the last transmitted port status TLV value. • In the Remote MEP section, displays the last value of port status TLV received from the remote MEP. In the Action profile section, displays, the last occurred event **port-status-tlv blocked** event. This event occurred due to the reception of **blocked** value in the port status TLV from remote MEP.
Interface status TLV	• In the Maintenance domain section, displays the last transmitted interface status TLV value. • In the Remote MEP section, displays the last value of interface status TLV received from the remote MEP. In the Action profile section, if displays, the last occurred event interface-status-tlv event (either **lower-layer-down** or **down**). This event occurred due to the reception of either lower or **down** value in the interface status TLV from remote MEP.
Action profile	Name of the action profile occurrence associated with a remote MEP.
Last event	When an action profile occurs, displays the last event that triggered it.
Last event cleared	When all the configured and occurred events (under action profile) are cleared, then the action taken gets reverted (such as down interface is made up) and the corresponding time is noted and displayed.
Action	Action taken and the corresponding time of the action occurrence.

Sample Output

show oam ethernet connectivity-fault- management mep-database

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain vpls-
vlan2000 maintenance-association vpls-vlan200
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
  Maintenance association name: vpls-vlan200, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
  Auto-discovery: enabled, Priority: 0
 Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: no  Interface name: ge-0/0/1.0, Interface status: Active, Link
status: Up
  Defects:
    Remote MEP not receiving CCM                  : no
    Erroneous CCM received                        : no
    Cross-connect CCM received                    : no
    RDI sent by some MEP                          : no
  Statistics:
    CCMs sent                                     : 1476
    CCMs received out of sequence                 : 0
    LBMs sent                                     : 85
    Valid in-order LBRs received                  : 78
    Valid out-of-order LBRs received              : 0
    LBRs received with corrupted data             : 0
    LBRs sent                                     : 0
    LTMs sent                                     : 1
    LTMs received                                 : 0
    LTRs sent                                     : 0
    LTRs received                                 : 1
    Sequence number of next LTM request           : 1
    1DMs sent                                     : 0
    Valid 1DMs received                           : 0
    Invalid 1DMs received                         : 0
    DMMs sent                                     : 0
    DMRs sent                                     : 0
    Valid DMRs received                           : 0
    Invalid DMRs received                         : 0
  Remote MEP count: 1
```

```
    Identifier    MAC address        State    Interface
      100     00:19:e2:b2:81:4b       ok  vt-0/1/10.1049088
```

show oam ethernet connectivity-fault- management mep-database (One-Way ETH-DM)

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain md6
maintenance-domain ma6
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                  : no
    Erroneous CCM received                        : no
    Cross-connect CCM received                    : no
    RDI sent by some MEP                          : no
  Statistics:
    CCMs sent                                     : 1590
    CCMs received out of sequence                 : 0
    LBMs sent                                     : 0
    Valid in-order LBRs received                  : 0
    Valid out-of-order LBRs received              : 0
    LBRs received with corrupted data             : 0
    LBRs sent                                     : 0
    LTMs sent                                     : 0
    LTMs received                                 : 0
    LTRs sent                                     : 0
    LTRs received                                 : 0
    Sequence number of next LTM request           : 0
    1DMs sent                                     : 10
    Valid 1DMs received                           : 0
    Invalid 1DMs received                         : 0
    DMMs sent                                     : 0
    DMRs sent                                     : 0
    Valid DMRs received                           : 0
    Invalid DMRs received                         : 0
  Remote MEP count: 1
```

```
    Identifier    MAC address        State    Interface
      201     00:90:69:0a:43:94       ok    ge-0/2/5.0
```

show oam ethernet connectivity-fault- management mep-database local-mep remote-mep

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain vpls-
vlan2000 maintenance-association vpls-vlan200 local-mep 200 remote-mep 100
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
  Maintenance association name: vpls-vlan200, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 100, State: ok
    MAC address: 00:19:e2:b2:81:4b, Type: Learned
    Interface: vt-0/1/10.1049088
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
```

show oam ethernet connectivity-fault- management mep-database remote-mep (Action Profile Event)

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain md5
maintenance-association ma5 remote-mep 200
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 200, State: ok
    MAC address: 00:05:85:73:96:1f, Type: Configured
    Interface: ge-1/0/8.0
    Last flapped: Never
```

```
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: lower-layer-down
    Action profile: juniper
      Last event: Interface-status-tlv lower-layer-down
      Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)
```

show oam ethernet connectivity-fault- management mep-database (Connection Protection TLV Configured)

```
user@host>show oam ethernet connectivity-fault-management mep-database maintenance-domain md5
maintenance-association ma5
```

If connection-protection is not enabled on down MEPs, but connection-protection TLV is used, MX always sets the protection-in-use flag in connection-protection tlv, while CCMs are sent out. During reversion, this is an indicator to the receiver that protect-path is in use, otherwise the peer (receiver) assumes working is active and reversion does not work as expected. Setting this bit does not affect protection-switching/traffic-loss.

```
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 1, Direction: down, MAC address: 00:19:e2:b1:14:30
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: yes
    Prefer me: no, Protection in use: no, FRR Flag: no
  Interface name: xe-6/2/0.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                  : no
    Erroneous CCM received                        : no
    Cross-connect CCM received                    : no
    RDI sent by some MEP                          : no
    Some remote MEP's MAC in error state          : no
  Statistics:
    CCMs sent                                     : 251
    CCMs received out of sequence                 : 0
    LBMs sent                                     : 0
    Valid in-order LBRs received                  : 0
    Valid out-of-order LBRs received              : 0
```

```
    LBRs received with corrupted data             : 0
    LBRs sent                                     : 0
    LTMs sent                                     : 0
    LTMs received                                 : 0
    LTRs sent                                     : 0
    LTRs received                                 : 0
    Sequence number of next LTM request           : 0
    1DMs sent                                     : 0
    Valid 1DMs received                           : 0
    Invalid 1DMs received                         : 0
    Out of sync 1DMs received                     : 0
    DMMs sent                                     : 0
    Valid DMMs received                           : 0
    Invalid DMMs received                         : 0
    DMRs sent                                     : 0
    Valid DMRs received                           : 0
    Invalid DMRs received                         : 0
    LMMs sent                                     : 0
    Valid LMMs received                           : 0
    Invalid LMMs received                         : 0
    LMRs sent                                     : 0
    Valid LMRs received                           : 0
    Invalid LMRs received                         : 0
  Remote MEP count: 1
   Identifier    MAC address        State    Interface
        2     00:90:69:7f:e4:30
```

show oam ethernet connectivity-fault- management mep-database

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain md5
maintenance-association ma5
Maintenance association name: ma1, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:14:f6:b6:01:fe
Auto-discovery: enabled, Priority: 0
Interface name: ge-1/0/0.0, Interface status: Active, Link status: Up

Defects:
Remote MEP not receiving CCM         : no
Erroneous CCM received               : no
Cross-connect CCM received           : no
```

```
RDI sent by some MEP                 : no

Statistics:
CCMs sent                            : 328703
CCMs received out of sequence        : 0
LBMs sent                            : 85
Valid in-order LBRs received         : 78
Valid out-of-order LBRs received     : 0
LBRs received with corrupted data    : 0
LBRs sent                            : 0
LTMs sent                            : 0
LTMs received                        : 0
LTRs sent                            : 0
LTRs received                        : 0
Sequence number of next LTM request  : 0
1DMs sent                            : 10
Valid 1DMs received                  : 10
Invalid 1DMs received                : 0
DMMs sent                            : 20
DMRs sent                            : 0
Valid DMRs received                  : 10
Invalid DMRs received                : 0
LMM sent                             : 10
Valid LMM received                   : 20
Invalid LMM received                 : 0
LMR sent                             : 20
Valid LMR received                   : 10
Invalid LMR received                 : 0
SLM sent                                      : 10
Valid SLM received                            : 20
Invalid SLM received                          : 0
SLR sent                                      : 20
Valid SLR received                            : 10
Invalid SLR received                          : 0

Remote MEP count                     : 1

Identifier     MAC address         State     Interface
    2       00:12:1e:fb:ea:7d       ok        ge-1/0/0.0
```

show oam ethernet connectivity-fault- management mep-database (enhanced continuity measurement)

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain md5
maintenance-association ma5 local-mep 2001 remote-mep 1001
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address  : 00:19:e2:b0:74:00, Type: Learned
Interface    : ge-2/0/0.0
Last flapped  : Never
+ Continuity  : 91%, Admin-enable duration: 2100sec, Oper-down duration: 100sec
Remote defect indication: false
Port status TLV: none
Interface status TLV: none
```

show oam ethernet connectivity-fault- management mep-database local-mep remote-mep

When the timer is not running:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain md6
maintenance-association ma6 local-mep 201 remote-mep 101 extensive
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 201, Direction: up, MAC address: 00:24:dc:48:46:cd
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: no
  Chassis ID: 00:24:dc:48:47:f2
  Management Address: 10.216.34.30
  Sendid TLV: no
  Interface name: ge-5/1/9.0, Interface status: Active, Link status: Up
```

```
  Remote MEP identifier: 101, State: ok
    MAC address: 00:22:83:db:4c:5b, Type: Configured
    Interface: ge-5/3/0.0
    Last flapped: 2019-02-05 00:45:57 PST (00:54:51 ago)
    Continuity: 99%, Admin-enable duration: 3367sec, Oper-down duration: 43sec
    Effective loss threshold: 3 frames
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
    Remote Management Address: 0.0.0.0
    Sendid TLV: no
    Sender ID Remote Chassis Subtype: 0
    EDM time remaining: N/A
```

When the timer is running:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain md6
maintenance-association ma6 local-mep 201 remote-mep 101 extensive
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 201, Direction: up, MAC address: 00:24:dc:48:46:cd
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: no
  Chassis ID: 00:24:dc:48:47:f2
  Management Address: 10.216.34.30
  Sendid TLV: no
  Interface name: ge-5/1/9.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 101, State: ok
    MAC address: 00:22:83:db:4c:5b, Type: Configured
    Interface: ge-5/3/0.0
    Last flapped: 2019-02-05 00:45:57 PST (00:56:16 ago)
    Continuity: 99%, Admin-enable duration: 3451sec, Oper-down duration: 43sec
    Effective loss threshold: 3 frames
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
    Remote Management Address: 0.0.0.0
    Sendid TLV: no
```

```
    Sender ID Remote Chassis Subtype: 0
    EDM time remaining: 894 secs
```

Release Information

Command introduced in Junos OS Release 8.4.

Support for ITU-T Y.1731 frame delay measurement added in Junos OS Release 9.5.

Support for ITU-T Y.1731 synthetic frame loss measurement added in Junos OS Release 13.2 for MX Series routers.

Support for ITU-T Y.1731 ethernet expected defect function (ETH-ED) added in Junos OS Release 19.1 for MX Series routers

NOTE: Ethernet expected defect function (ETH-ED) will be supported on all MX series routers only when `enhanced-ip` mode is enabled. ETH-ED will not be supported if CFM is running in centralized mode.

RELATED DOCUMENTATION

- clear oam ethernet connectivity-fault-management statistics | 3775
- show oam ethernet connectivity-fault-management delay-statistics | 22585
- show oam ethernet connectivity-fault-management interfaces | 22596
- show oam ethernet connectivity-fault-management mep-statistics | 22637

show oam ethernet connectivity-fault-management mep-statistics

IN THIS SECTION

- Syntax | 22638
- Description | 22638
- Options | 22638
- Required Privilege Level | 22639

- Output Fields | 22639
- Sample Output | 22642
- Release Information | 22654

Syntax

```
show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md-name
maintenance-association ma-name
<mep mep-id>
<remote-mep remote-mep-id>
<count entry-count>
```

Description

On MX Series and ACX Series routers and EX Series switches with Ethernet interfaces, display ETH-DM statistics and ETH-DM frame counts.

For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.

Options

maintenance-domain *md-name*	Name of an existing CFM maintenance domain.
maintenance-association *ma-name*	Name of an existing CFM maintenance association.
mep *mep-id*	(Optional) Numeric identifier of the local MEP. The range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.
remote-mep *remote-mep-id*	(Optional) Numeric identifier of the remote MEP. The range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.
count *entry-count*	(Optional) Number of entries to display from the statistics table. The range of values is 1 through 100. The default value is 100 entries.

Required Privilege Level

view

Output Fields

Table 905 on page 22639 lists the output fields for the `show oam ethernet connectivity-fault-management mep-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 905: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields

Output Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) numeric identifier.
MAC address	Unicast MAC address configured for the MEP.
Remote MEP count	Number of remote MEPs (unless you specify the **remote-mep** option).
`CCMs sent`	Number of continuity check messages (CCMs) sent.
`CCMs received`	Number of continuity check messages (CCMs) received for a specific remote MEP and maintenance association.
`CCMs received out of sequence`	Number of continuity check messages (CCMs) received that were not in sequence.
Remote MEP identifier	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Index	Index number that corresponds to the ETH-DM entry in the CFM database.

Table 905: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields *(Continued)*

Output Field Name	Field Description
One-way delay (usec)	For a one-way ETH-DM session, the frame delay time, in microseconds, measured at the receiver MEP. For a detailed description of one-way Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the Junos OS Network Interfaces Library for Routing Devices.
Two-way delay (usec)	For a two-way ETH-DM session, the frame delay time, in microseconds, measured at the initiator MEP. For a detailed description of two-way Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the Junos OS Network Interfaces Library for Routing Devices.
Average one-way delay	Average one-way frame delay for the statistics displayed.
Average one-way delay variation	Average one-way “frame jitter” for the statistics displayed.
Best-case one-way delay	Lowest one-way frame delay for the statistics displayed.
Worst-case one-way delay	Highest one-way frame delay for the statistics displayed.
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.
Best-case two-way delay	Lowest two-way frame delay for the statistics displayed.
Worst-case two-way delay	Highest two-way frame delay calculated in this session.
`SLM packets sent`	Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.

Table 905: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields *(Continued)*

Output Field Name	Field Description
`SLM packets received`	Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.
`SLR packets sent`	Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.
`SLR packets received`	Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.
`Local TXFC1 value`	Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.
`Local RXFC1 value`	Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet.
`Last Received SLR frame TXFCf(tc)`	Value of the local counter TxFCl at the time of SLM frame transmission.
`Last Received SLR frame TXFCb(t`	Value of the local counter RxFCl at the time of SLR frame transmission.
`Frame loss (near-end)`	Count of frame loss associated with ingress data frames.
`Frame loss (far-end)`	Count of frame loss associated with egress data frames.

Sample Output

show oam ethernet connectivity-fault- management mep-statistics (CIR counters only)

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1
maintenance-association ma-1 local-mep 3 remote-mep 103 count 3
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count                    : 1
CCMs sent                           : 6550
CCMs received out of sequence       : 0
LBMs sent                           : 0
Valid in-order LBRs received        : 0
Valid out-of-order LBRs received    : 0
LBRs received with corrupted data   : 0
LBRs sent                           : 0
LTMs sent                           : 0
LTMs received                       : 0
LTRs sent                           : 0
LTRs received                       : 0
Sequence number of next LTM request : 0
1DMs sent                           : 5
Valid 1DMs received                 : 0
Invalid 1DMs received               : 0
DMMs sent                           : 5
DMRs sent                           : 0
Valid DMRs received                 : 5
Invalid DMRs received               : 0
LMM sent                            : 5
Valid LMM received                  : 5
Invalid LMM received                : 0
LMR sent                            : 0
Valid LMR received                  : 5
Invalid LMR received                : 0
Remote MEP identifier               : 101
Remote MAC address                  : 00:05:85:73:39:4a

Delay measurement statistics:
Index       One-way delay         Two-way delay
               (usec)                (usec)
  1             259                   519
  2             273                   550
  3             287                   571
```

```
  4             299                   610
  5             313                   650

Average one-way delay                 : 286 usec
Average one-way delay variation       : 62 usec
Best case one-way delay               : 259 usec
Average two-way delay                 : 580 usec
Average two-way delay variation       : 26 usec
Best case two-way delay               : 519 usec
Worst case two-way delay              : 650 usec

Loss measurement statistics:
Index     Near-end     Far-end       Near-end     Far-end
          Frame loss   Frame loss    Frame loss   Frame loss
            (CIR)       (CIR)          (EIR)        (EIR)
  1           9           9
  2           3           5
  3           7           5
  4           9           6
  5           3           6

Average near-end loss (CIR)           : 6.2
Average near-end loss ratio (CIR)     : 6.2%
Average far-end loss (CIR)            : 6.2
Average far-end loss ratio (CIR)      : 6.2%
Near-end best case loss (CIR)         : 3
Near-end best case loss ratio (CIR)   : 3%
Near-end worst case loss (CIR)        : 9
Near-end worst case loss ratio (CIR)  : 9%
Far-end best case loss (CIR)          : 5
Far-end best case loss ratio (CIR)    : 5%
Far-end worst case loss (CIR)         : 9
Far-end worst case loss ratio (CIR)   : 9%
```

show oam ethernet connectivity-fault- management mep-statistics (CIR and EIR counters enabled)

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1
maintenance-association ma-1 local-mep 3 remote-mep 103 count 3
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count                      : 1
```

```
CCMs sent                             : 6550
CCMs received out of sequence         : 0
LBMs sent                             : 0
Valid in-order LBRs received          : 0
Valid out-of-order LBRs received      : 0
LBRs received with corrupted data     : 0
LBRs sent                             : 0
LTMs sent                             : 0
LTMs received                         : 0
LTRs sent                             : 0
LTRs received                         : 0
Sequence number of next LTM request   : 0
1DMs sent                             : 5
Valid 1DMs received                   : 0
Invalid 1DMs received                 : 0
DMMs sent                             : 5
DMRs sent                             : 0
Valid DMRs received                   : 5
Invalid DMRs received                 : 0
LMM sent                              : 5
alid LMM received                     : 5
Invalid LMM received                  : 0
LMR sent                              : 0
Valid LMR received                    : 5
Invalid LMR received                  : 0
Remote MEP identifier                 : 101
Remote MAC address                    : 00:05:85:73:39:4a

Delay measurement statistics:
Index       One-way delay         Two-way delay
              (usec)                (usec)
  1            259                   519
  2            273                  550
  3            287                  571
  4            299                  610
  5            313                  650

Average one-way delay                   : 286 usec
Average one-way delay variation         : 62 usec
Best case one-way delay                 : 259 usec
Average two-way delay                   : 580 usec
Average two-way delay variation         : 26 usec
Best case two-way delay                 : 519 usec
```

```
Worst case two-way delay               : 650 usec

Loss measurement statistics:
Index     Near-end     Far-end       Near-end     Far-end
          Frame loss   Frame loss    Frame loss   Frame loss
           (CIR)         (CIR)        (EIR)         (EIR)
  1          9             9             2            4
  2          3             5             4            6
  3          7             5             0            2
  4          9             6             8            2
  5          3             6             6            4

Average near-end loss (CIR)               : 6.2
Average near-end loss ratio (CIR)         : 6.2%
Average far-end loss (CIR)                : 6.2
Average far-end loss ratio (CIR)          : 6.2%
Near-end best case loss (CIR)             : 3
Near-end best case loss ratio (CIR)       : 3%
Near-end worst case loss (CIR)            : 9
Near-end worst case loss ratio (CIR)      : 9%
Far-end best case loss (CIR)              : 5
Far-end best case loss ratio (CIR)        : 5%
Far-end worst case loss (CIR)             : 9
Far-end worst case loss ratio (CIR)       : 9%
Average near-end loss (EIR)               : 4
Average near-end loss ratio (EIR)         : 4%
Average far-end loss (EIR)                : 3.4
Average far-end loss ratio (EIR)          : 3.4%
Near-end best case loss (EIR)             : 0
Near-end best case loss ratio (EIR)       : 0%
Near-end worst case loss (EIR)            : 8
Near-end worst case loss ratio (EIR)      : 8%
Far-end best case loss (EIR)              : 2
Far-end best case loss ratio (EIR)        : 2%
Far-end worst case loss (EIR)             : 6
Far-end worst case loss ratio (EIR)       : 6%
```

show oam ethernet connectivity-fault- management mep-statistics remote-mep (CIR counters only)

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1
maintenance-association ma-1 local-mep 3 remote-mep 103 count 3 remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
CCMs sent                              : 7762
CCMs received out of sequence          : 0
LBMs sent                              : 0
Valid in-order LBRs received           : 0
Valid out-of-order LBRs received       : 0
LBRs received with corrupted data      : 0
LBRs sent                              : 0
LTMs sent                              : 0
LTMs received                          : 0
LTRs sent                              : 0
LTRs received                          : 0
Sequence number of next LTM request    : 0
1DMs sent                              : 5
Valid 1DMs received                    : 0
Invalid 1DMs received                  : 0
DMMs sent                              : 5
DMRs sent                              : 0
Valid DMRs received                    : 5
Invalid DMRs received                  : 0
LMM sent                               : 5
Valid LMM received                     : 5
Invalid LMM received                   : 0
LMR sent                               : 0
Valid LMR received                     : 5
Invalid LMR received                   : 0
Remote MEP identifier                  : 101
Remote MAC address                     : 00:05:85:73:39:4a

Delay measurement statistics:
Index       One-way delay         Two-way delay
               (usec)                (usec)
  1             259                   519
  2             273                   550
  3             287                   571
  4             299                   610
  5             313                   650
```

```
Average one-way delay                  : 286 usec
Average one-way delay variation        : 62 usec
Best case one-way delay                : 259 usec
Average two-way delay                  : 580 usec
Average two-way delay variation        : 26 usec
Best case two-way delay                : 519 usec
Worst case two-way delay               : 650 usec

Loss measurement statistics:
Index     Near-end     Far-end       Near-end     Far-end
          Frame loss   Frame loss    Frame loss   Frame loss
            (CIR)        (CIR)         (EIR)       (EIR)
  1           9            9
  2           3            5
  3           7            5
  4           9            6
  5           3            6
Average near-end loss (CIR)               : 6.2
Average near-end loss ratio (CIR)         : 6.2%
Average far-end loss (CIR)                : 6.2
Average far-end loss ratio (CIR)          : 6.2%
Near-end best case loss (CIR)             : 3
Near-end best case loss ratio (CIR)       : 3%
Near-end worst case loss (CIR)            : 9
Near-end worst case loss ratio (CIR)      : 9%
Far-end best case loss (CIR)              : 5
Far-end best case loss ratio (CIR)        : 5%
Far-end worst case loss (CIR)             : 9
Far-end worst case loss ratio (CIR)       : 9%
Average near-end loss (EIR)               : 4
Average near-end loss ratio (EIR)         : 4%
Average far-end loss (EIR)                : 3.4
Average far-end loss ratio (EIR)          : 3.4%
Near-end best case loss (EIR)             : 0
Near-end best case loss ratio (EIR)       : 0%
Near-end worst case loss (EIR)            : 8
Near-end worst case loss ratio (EIR)      : 8%
Far-end best case loss (EIR)              : 2
Far-end best case loss ratio (EIR)        : 2%
Far-end worst case loss (EIR)             : 6
Far-end worst case loss ratio (EIR)       : 6%
```

show oam ethernet connectivity-fault- management mep-statistics remote-mep (CIR and EIR counters enabled)

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1
maintenance-association ma-1 local-mep 3 remote-mep 103 count 3 remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
CCMs sent                            :7762
CCMs received out of sequence        : 0
LBMs sent                            : 0
Valid in-order LBRs received         : 0
Valid out-of-order LBRs received     : 0
LBRs received with corrupted data    : 0
LBRs sent                            : 0
LTMs sent                            : 0
LTMs received                        : 0
LTRs sent                            : 0
LTRs received                        : 0
Sequence number of next LTM request  : 0
1DMs sent                            : 5
Valid 1DMs received                  : 0
Invalid 1DMs received                : 0
DMMs sent                            : 5
DMRs sent                            : 0
Valid DMRs received                  : 5
Invalid DMRs received                : 0
LMM sent                             : 5
Valid LMM received                   : 5
Invalid LMM received                 : 0
LMR sent                             : 0
Valid LMR received                   : 5
Invalid LMR received                 : 0
Remote MEP identifier                : 101
Remote MAC address                   : 00:05:85:73:39:4a

Delay measurement statistics:
Index       One-way delay         Two-way delay
               (usec)                (usec)
  1             259                   519
  2             273                   550
  3             287                   571
  4             299                   610
  5             313                   650
```

```
Average one-way delay                 : 286 usec
Average one-way delay variation       : 62 usec
Best case one-way delay               : 259 usec
Average two-way delay                 : 580 usec
Average two-way delay variation       : 26 usec
Best case two-way delay               : 519 usec
Worst case two-way delay              : 650 usec

Loss measurement statistics:
Index     Near-end     Far-end       Near-end     Far-end
          Frame loss   Frame loss    Frame loss   Frame loss
            (CIR)       (CIR)          (EIR)        (EIR)
  1          10           8              5            12
  2          12           7              6            16
  3          7            5              0            2
  4          9            6              8            2
  5          3            6              6            4

Average near-end loss (CIR)            : 6.2
Average near-end loss ratio (CIR)      : 6.2%
Average far-end loss (CIR)             : 6.2
Average far-end loss ratio (CIR)       : 6.2%
Near-end best case loss (CIR)          : 3
Near-end best case loss ratio (CIR)    : 3%
Near-end worst case loss (CIR)         : 9
Near-end worst case loss ratio (CIR)   : 9%
Far-end best case loss (CIR)           : 5
Far-end best case loss ratio (CIR)     : 5%
Far-end worst case loss (CIR)          : 9
Far-end worst case loss ratio (CIR)    : 9%
Average near-end loss (EIR)            : 4
Average near-end loss ratio (EIR)      : 4%
Average far-end loss (EIR)             : 3.4
Average far-end loss ratio (EIR)       : 3.4%
Near-end best case loss (EIR)          : 0
Near-end best case loss ratio (EIR)    : 0%
Near-end worst case loss (EIR)         : 8
Near-end worst case loss ratio (EIR)   : 8%
Far-end best case loss (EIR)           : 2
Far-end best case loss ratio (EIR)     : 2%
```

```
Far-end worst case loss (EIR)          : 6
Far-end worst case loss ratio (EIR)    : 6%
```

show oam ethernet connectivity-fault- management mep-statistics

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1
maintenance-association ma-1
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 1
  CCMs sent                                      : 6550
  CCMs received out of sequence                  : 0
  LBMs sent                                      : 0
  Valid in-order LBRs received                   : 0
  Valid out-of-order LBRs received               : 0
  LBRs received with corrupted data              : 0
  LBRs sent                                      : 0
  LTMs sent                                      : 0
  LTMs received                                  : 0
  LTRs sent                                      : 0
  LTRs received                                  : 0
  Sequence number of next LTM request            : 0
  1DMs sent                                      : 5
  Valid 1DMs received                            : 0
  Invalid 1DMs received                          : 0
  DMMs sent                                      : 5
  DMRs sent                                      : 0
  Valid DMRs received                            : 5
  Invalid DMRs received                          : 0
  SLM sent                                       : 10
  Valid SLM received                             : 20
  Invalid SLM received                           : 0
  SLR sent                                       : 20
  Valid SLR received                             : 10
  Invalid SLR received                           : 0

  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        259            519
```

```
      2        273            550
      3        287            571
      4        299            610
      5        313            650
    Average one-way delay          : 286 usec
    Average one-way delay variation: 62 usec
    Best case one-way delay        : 259 usec
    Worst case one-way delay       : 313 usec
    Average two-way delay          : 580 usec
    Average two-way delay variation: 26 usec
    Best case two-way delay        : 519 usec
    Worst case two-way delay       : 650 usec    Synthetic Loss measurement statistics:
     SLM packets sent                             : 100
     SLM packets received                         : 0
     SLR packets sent                             : 100
     SLR packets received                         : 0
     Accumulated SLM statistics:
     Local TXFC1 value                            : 100
     Local RXFC1 value                            : 100
     Last Received SLR frame TXFCftc              : 100
     Last Received SLR frame TXFCbtc              : 100
     SLM Frame Loss:
     Frame Loss (far-end)                         : 0 (0.00 %)
     Frame Loss (near-end)                        : 0 (0.00 %)
```

show oam ethernet connectivity-fault-management mep-statistics remote-mep

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1
maintenance-association ma1 remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
  CCMs sent                                       : 7762
  CCMs received out of sequence                   : 0
  LBMs sent                                       : 0
  Valid in-order LBRs received                    : 0
  Valid out-of-order LBRs received                : 0
  LBRs received with corrupted data               : 0
  LBRs sent                                       : 0
  LTMs sent                                       : 0
  LTMs received                                   : 0
  LTRs sent                                       : 0
  LTRs received                                   : 0
```

```
Sequence number of next LTM request           : 0
1DMs sent                                     : 5
Valid 1DMs received                           : 0
Invalid 1DMs received                         : 0
DMMs sent                                     : 5
DMRs sent                                     : 0
Valid DMRs received                           : 5
Invalid DMRs received                         : 0
SLM sent                                      : 10
Valid SLM received                            : 20
Invalid SLM received                          : 0
SLR sent                                      : 20
Valid SLR received                            : 10
Invalid SLR received                          : 0

Remote MEP identifier: 101
Remote MAC address: 00:05:85:73:39:4a
  Delay measurement statistics:
  Index  One-way delay  Two-way delay
            (usec)         (usec)
    1        259            519
    2        273            550
    3        287            571
    4        299            610
    5        313            650
  Average one-way delay          : 286 usec
  Average one-way delay variation: 62 usec
  Best case one-way delay        : 259 usec
  Worst case one-way delay       : 313 usec
  Average two-way delay          : 580 usec
  Average two-way delay variation: 26 usec
  Best case two-way delay        : 519 usec
  Worst case two-way delay       : 650 usec    Synthetic Loss measurement statistics:
   SLM packets sent                             : 100
   SLM packets received                         : 0
   SLR packets sent                             : 100
   SLR packets received                         : 0
   Accumulated SLM statistics:
   Local TXFC1 value                            : 100
   Local RXFC1 value                            : 100
   Last Received SLR frame TXFCftc              : 100
   Last Received SLR frame TXFCbtc              : 100
   SLM Frame Loss:
```

```
    Frame Loss (far-end)                        : 0 (0.00 %)
    Frame Loss (near-end)                       : 0 (0.00 %)
```

show oam ethernet connectivity-fault-management mep-statistics local-mep remote-mep

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1
maintenance-association ma1 local-mep 121 remote-mep 101
MEP identifier: 1, MAC address: 3c:61:04:25:fa:95
Remote MEP count: 1
  CCMs sent                                     : 875
  CCMs Received                                 : 300
  CCMs received out of sequence                 : 0
  LBMs sent                                     : 0
  Valid in-order LBRs received                  : 0
  Valid out-of-order LBRs received              : 0
  LBRs received with corrupted data             : 0
  LBRs sent                                     : 0
  LTMs sent                                     : 0
  LTMs received                                 : 0
  LTRs sent                                     : 0
  LTRs received                                 : 0
  Sequence number of next LTM request           :  12
  1DMs sent                                     : 0
  Valid 1DMs received                           : 0
  Invalid 1DMs received                         : 0
  Out of sync 1DMs received                     : 0
  DMMs sent                                     : 0
  Valid DMMs received                           : 0
  Invalid DMMs received                         : 0
  DMRs sent                                     : 0
  Valid DMRs received                           : 0
  Invalid DMRs received                         : 0
  LMMs sent                                     : 0
  Valid LMMs received                           : 0
  Invalid LMMs received                         : 0
  LMRs sent                                     : 0
  Valid LMRs received                           : 0
  Invalid LMRs received                         : 0
  SLMs sent                                     : 0
  Valid SLMs received                           : 0
  Invalid SLMs received                         : 0
```

```
SLRs sent                                    : 0
Valid SLRs received                          : 0
Invalid SLRs received                        : 0
Valid AISs generated                         : 0
Valid AISs received                          : 0
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for ITU-T Y.1731 Ethernet synthetic frame loss measurement (ETH-SLM) added in Junos OS Release 13.2 for MX Series routers.

RELATED DOCUMENTATION

clear oam ethernet connectivity-fault-management statistics | 3775

show oam ethernet connectivity-fault-management delay-statistics | 22585

show oam ethernet connectivity-fault-management interfaces | 22596

show oam ethernet connectivity-fault-management mep-database | 22620

show oam ethernet connectivity-fault-management mip

IN THIS SECTION

- Syntax | 22655
- Description | 22655
- Options | 22655
- Required Privilege Level | 22655
- Output Fields | 22655
- Sample Output | 22656
- Release Information | 22656

Syntax

```
show oam ethernet connectivity-fault-management mip
<qualifier>
```

Description

Display all the maintenance association intermediate points (MIPs) created in the system. Specify qualifiers to display specific MIPs.

Options

qualifier—(Optional) Display the specified MIP.

Required Privilege Level

view

Output Fields

Table 906 on page 22655 lists the output fields for the `show oam ethernet connectivity-fault-management mip` command. Output fields are listed in the approximate order in which they appear.

Table 906: show oam ethernet connectivity-fault-management mip Output Fields

Field Name	Field Description
MIP information for instance	Header for the MIP information showing the MIP name.
Interface	Interface type-dpc/pic/port.unit-number.
Level	MIP level configured.

Sample Output

show oam ethernet connectivity-fault -management mip

```
user@host> show oam ethernet connectivity-fault-management mip
MIP information for  __mip_name__

MIP information for  default-switch bd1

  Interface       Level
  ge-3/0/0.0          7
  ge-3/0/1.0          6
  ge-3/0/3.0          6
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

No Link Title

No Link Title

show oam ethernet connectivity-fault-management path-database

IN THIS SECTION

- Syntax | 22657
- Description | 22657
- Options | 22657
- Required Privilege Level | 22657
- Output Fields | 22657
- Sample Output | 22659

Release Information | 22660

Syntax

```
show oam ethernet connectivity-fault-management path-database
<host-mac-address>
<maintenance-association ma-name>
<maintenance-domain domain-name>
```

Description

On M7i and M10i with Enhanced CFEB (CFEB-E), M320, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management path database information for a host configured with an MEP.

Options

host-mac-address	(Optional) Display connectivity fault management path database information for a specified Ethernet host.
maintenance-association ***ma-name***	(Optional) Display connectivity fault management path database information for the specified maintenance association.
maintenance-domain ***domain-name***	(Optional) Display connectivity fault management path database information for the specified maintenance domain.

Required Privilege Level

view

Output Fields

No Link Title lists the output fields for the `show oam ethernet connectivity- fault-management path-database` command. Output fields are listed in the approximate order in which they appear.

Table 907: show oam ethernet connectivity-fault-management linktrace path-database Output Fields

Field Name	Field Description
Linktrace to	MAC address of the 802.1ag node to which the linktrace message is targeted.
Interface	Interface used by the local MEP to send the linktrace message (LTM).
Maintenance Domain	Maintenance domain identifier specified in the traceroute command.
Maintenance Association	Maintenance association identifier specified in the traceroute command.
Level	Maintenance domain level configured for the maintenance domain.
Local Mep	MEP identifier of the local MEP originating the linktrace.
Hop	Sequential hop count of the linktrace path.
TTL	Number of hops remaining in the linktrace message (LTM). The time to live (TTL) is decremented at each hop.
Source MAC address	MAC address of the 802.1ag maintenance intermediate point (MIP) that is forwarding the LTM.
Next hop MAC address	MAC address of the 802.1ag node that is the next hop in the LTM path.
Transaction Identifier	4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all maintenance domains. Use the transaction identifier to match an incoming linktrace responses (LTR), with a previously sent LTM.

Sample Output

show oam ethernet connectivity-fault- management path-database

```
user@host> show oam ethernet connectivity-fault-management path-database maintenance-domain MD1
maintenance-association MA1 00:01:02:03:04:05
Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
       Maintenance Domain: MD1, Level: 7
       Maintenance Association: MA1, Local Mep: 1

       Hop     TTL      Source MAC address        Next hop MAC address
       Transaction Identifier:100001
       1       63       00:00:aa:aa:aa:aa         00:00:bb:bb:bb:bb
       2       62       00:00:bb:bb:bb:bb         00:00:cc:cc:cc:cc
       3       61       00:00:cc:cc:cc:cc         00:01:02:03:04:05
       4       60       00:01:02:03:04:05         00:00:00:00:00:00
```

show oam ethernet connectivity-fault- management linktrace path-database (Two traceroute commands)

```
user@host> show oam ethernet connectivity-fault-management path-database maintenance-domain MD2
maintenance-association MA2 00:06:07:08:09:0A
Linktrace to 00:06:07:08:09:0A, Interface : ge-5/0/1.0
       Maintenance Domain: MD2, Level: 6
       Maintenance Association: MA2, Local Mep: 10

       Hop     TTL      Source MAC address        Next hop MAC address
       Transaction Identifier:100002
       1       63       00:00:aa:aa:aa:aa         00:00:bb:bb:bb:bb
       2       62       00:00:bb:bb:bb:bb         00:00:cc:cc:cc:cc
       3       61       00:00:cc:cc:cc:cc         00:06:07:08:09:0A
       4       60       00:06:07:08:09:0A         00:00:00:00:00:00
       Transaction Identifier:100003
       1       63       00:00:aa:aa:aa:aa         00:00:bb:bb:bb:bb
       2       62       00:00:bb:bb:bb:bb         00:00:cc:cc:cc:cc
       3       61       00:00:cc:cc:cc:cc         00:06:07:08:09:0A
       4       60       00:06:07:08:09:0A         00:00:00:00:00:00
```

Release Information

Command introduced in Junos OS Release 8.4.

show oam ethernet connectivity-fault-management policer

IN THIS SECTION

- Syntax | 22660
- Description | 22660
- Options | 22660
- Required Privilege Level | 22661
- Output Fields | 22661
- Sample Output | 22662
- Release Information | 22663

Syntax

```
show oam ethernet connectivity-fault-management policer <maintenance-domain md-name> <maintenance-
association ma-name>
```

Description

On M120, M320, MX Series, T320, and T640 routers displays connectivity-fault-management policer statistics.

Options

This command has the following options:

maintenance-domain *md-name* — Name of an existing CFM maintenance domain. If this option is not specified, policer statistics are displayed for all maintenance associations for all maintenance domains.

maintenance-association *ma-name*

Name of an existing CFM maintenance association. If this option is not specified, policer statistics are displayed for all maintenance associations for given maintenance domain. This option cannot be specified without specifying maintenance-domain name.

Required Privilege Level

view

Output Fields

Table 908 on page 22661 lists the output fields for the `show oam ethernet connectivity-fault-management policer` command. Output fields are listed in the approximate order in which they appear.

Table 908: show oam ethernet connectivity-fault-management policer Output Fields

Field Name	Field Description
Legend for Policer	Describes the symbols used under the **Scope** and **Type** headings: • G - Global scope • S - Service scope • cc - Continuity check (Type)
Maintenance Domain	Displays the maintenance domain name.
Level	Displays the maintenance domain level configured.
Maintenance association	Displays the maintenance association name.
Policer	Displays the policer name.
Type	Policer type. Value **cc** means this policer is used only to police continuity check CFM messages. Value **other** means this policer is used only to police non-continuity check CFM messages. Value **all** means this policer is used to police all CFM messages.

Table 908: show oam ethernet connectivity-fault-management policer Output Fields *(Continued)*

Field Name	Field Description
Scope	Policer scope. Displays whether the *global* (G) policer configuration is applicable or the session (S) specific policer config is applicable.
Drop count	Displays the number of packets dropped by the indicated policer.

Sample Output

show oam ethernet connectivity-fault -management policer

Displays the policer information for all maintenance associations and their maintenance domains.

```
show oam ethernet connectivity-fault-management policer
Legend for Policer
G - Global scope
S - Service scope
cc - Continuity check

Maintenance Domain: md1 Level: 1
Maintenance association Policer      Type    Scope Drop count
  ma1                   cfm-policer1 all     G            300
  ma1-2                 cfm-policer1 cc      S            259
  ma1-2                 cfm-policer1 other   G            300
Maintenance Domain: md2 Level: 2
Maintenance association Policer      Type    Scope Drop count
  ma2                   cfm-policer1 cc      G            300
  ma2                   cfm-policer2 other   S            223
```

show oam ethernet connectivity-fault -management policer maintenance-domain md-name

Displays the policer information for the specified maintenance domain and its maintenance associations.

```
show oam ethernet connectivity-fault-management policer maintenance-domain md1
Legend for Policer
```

```
G - Global scope
S - Service scope
cc - Continuity check

Maintenance Domain: md1 Level: 1
Maintenance association Policer      Type    Scope Drop count
  ma1                   cfm-policer1 all     G            300
  ma1-2                 cfm-policer1 cc      S            259
  ma1-2                 cfm-policer1 other   G            300
```

show oam ethernet connectivity-fault -management policer maintenance-domain md-name maintenance-association ma-name

Displays the policer information for the specified **maintenance-domain** ***md-name*** and **maintenance-association** ***ma-name***.

```
show oam ethernet connectivity-fault-management policer maintenance-domain md5 maintenance-
association ma5
Legend for Policer
G - Global scope
S - Service scope
cc - Continuity check

Maintenance Domain: md5 Level: 5
Maintenance association Policer       Type    Scope Drop count
  ma5                   cfm-policer   cc      S            187
  ma5                   cfm-policer-2 other   S            234
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

clear oam ethernet connectivity-fault-management policer | 3772

show oam ethernet connectivity-fault-management sla-iterator-statistics

IN THIS SECTION

- Syntax | 22664
- Description | 22664
- Options | 22664
- Required Privilege Level | 22665
- Output Fields | 22665
- Sample Output | 22668
- Sample Output | 22669
- Sample Output | 22670
- Release Information | 22677

Syntax

```
show oam ethernet connectivity-fault-management sla-iterator-statistics
maintenance-domain md-name
maintenance-association ma-name
sla-iterator sla-iterator
<local-mep local-mep-id>
<remote-mep remote-mep-id>
```

Description

Display the Ethernet Operation, Administration, and Maintenance (OAM) service-level agreement (SLA) iterator statistics.

Options

maintenance-domain *md-name*	Name of an existing connectivity fault management (CFM) maintenance domain.
maintenance-association *ma-name*	Name of an existing CFM maintenance association.

sla-iterator ***sla-iterator***	Name of the iterator profile.
local-mep ***local-mep-id***	(Optional) Numeric identifier of the local MEP. The range of values is **1** through **8191**.
remote-mep ***remote-mep-id***	(Optional) Numeric identifier of the remote MEP. The range of values is **1** through **8192**.

Required Privilege Level

view

Output Fields

Table 909 on page 22665 lists the output fields for the `show oam ethernet connectivity-fault-management sla-iterator-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 909: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields

Output Field Name	Output Field Description
Maintenance domain	Name of the maintenance domain.
Level	Level of the maintenance domain level configured.
Maintenance association	Name of the maintenance association.
Local MEP id	Numeric identifier of the local MEP.
Remote MEP id	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Iterator name	Name of iterator.
Iterator Id	Numeric identifier of the iterator.

Table 909: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields *(Continued)*

Output Field Name	Output Field Description
Iterator cycle time	Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection
Iteration period	Maximum number of cycles per iteration
Iterator status	Current status of iterator whether running or stopped.
Infinite iterations	Status of iteration as infinite or finite.
Counter reset time	Date and time when the counter was reset.
Reset reason	Reason to reset counter.
Delay weight	Calculation weight of delay.
Delay variation weight	Calculation weight of delay variation.
DMM sent	Delay measurement message (DMM) PDU frames sent to the peer MEP in this session.
DMM skipped for threshold hit	Number of DMM frames sent to the peer MEP in this session skipped during threshold hit.
DMM skipped for threshold hit window	Number of DMM frames sent to the peer MEP in this session skipped during the last threshold hit window.
DMR received	Number of delay measurement reply (DMR) frames received.
DMR out of sequence	Total number of DMR out of sequence packets received.
DMR received with invalid time stamps	Total number of DMR frames received with invalid timestamps.

Table 909: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields *(Continued)*

Output Field Name	Output Field Description
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.
Average one-way forward delay variation	Average one-way forward delay variation for the statistics displayed in microseconds.
Average one-way backward delay variation	Average one-way backward delay variation for the statistics displayed in microseconds.
Weighted average two-way delay	Weighted average two-way delay for the statistics displayed in microseconds.
Weighted average two-way delay variation	Weighted average two-way delay variation for the statistics displayed in microseconds.
Weighted average one-way backward delay variation	Weighted average one-way backward delay variation for the statistics displayed in microseconds.
Weighted average one-way forward delay variation	Weighted average one-way forward delay variation for the statistics displayed in microseconds.
`SLM packets sent`	Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.
`SLM packets received`	Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.

Table 909: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields *(Continued)*

Output Field Name	Output Field Description
SLR packets sent	Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.
SLR packets received	Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.
Local TXFC1 value	Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.
Local RXFC1 value	Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet.
Last Received SLR frame TXFCf(tc)	Value of the local counter TxFCl at the time of SLM frame transmission.
Last Received SLR frame TXFCb(t	Value of the local counter RxFCl at the time of SLR frame transmission.
Frame loss (near-end)	Count of frame loss associated with ingress data frames.
Frame loss (far-end)	Count of frame loss associated with egress data frames.

Sample Output

show oam ethernet connectivity-fault -management sla-iterator-statistics

```
user@switch> show oam ethernet connectivity-fault-management sla-iterator-statistics sla-
iterator i1 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 2
```

```
Iterator statistics:
Maintenance domain: md6, Level: 6
Maintenance association: ma6, Local MEP id: 1000
Remote MEP id: 103, Remote MAC address: 00:90:69:0a:43:92
Iterator name: i1, Iterator Id: 1
Iterator cycle time: 10ms, Iteration period: 1 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2010-03-19 20:42:39 PDT (2d 18:24 ago)
Reset reason: Adjacency flap

Iterator delay measurement statistics:
Delay weight: 1, Delay variation weight: 1
DMM sent                                          : 23898520
DMM skipped for threshold hit                     : 11000
DMM skipped for threshold hit window              : 0
DMR received                                      : 23851165
DMR out of sequence                               : 1142
DMR received with invalid time stamps             : 36540
Average two-way delay                             : 129 usec
Average two-way delay variation                   : 15 usec
Average one-way forward delay variation           : 22 usec
Average one-way backward delay variation          : 22 usec
Weighted average two-way delay                    : 134 usec
Weighted average two-way delay variation          : 8 usec
Weighted average one-way forward delay variation  : 6 usec
Weighted average one-way backward delay variation : 2 usec
```

Sample Output

show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers)

```
user@router> show oam ethernet connectivity-fault-management sla-iterator-statistics maintenance-
domain mdu maintenance-association mau local-mep 4 remote-mep 3 sla-iterator lm
Iterator statistics:
  Maintenance domain: 2, Level: 2
  Maintenance association: W-160432000-001, Local MEP id: 2
  Remote MEP id: 1, Remote MAC address: 00:90:69:0a:43:39
  Iterator name: iter1, Iterator Id: 1
  Iterator cycle time: 100ms, Iteration period: 10 cycles
  Iterator status: running, Infinite iterations: true
  Counter reset time: 2012-09-25 02:15:31 PDT (00:00:45 ago)
```

```
  Reset reason: Adjacency flap
  Iterator loss measurement statistics:
    LMM sent                            : 444
    LMM skipped for threshold hit       : 0
    LMM skipped for threshold hit window: 0
    LMR received                        : 444
    LMR out of sequence                 : 0
    LMR forwarding-class mismatch       : 0
  Accumulated transmit statistics:
    Near-end (CIR)                      : 0
    Far-end (CIR)                       : 0
    Near-end (EIR)                      : 0
    Far-end (EIR)                       : 0
  Accumulated receive statistics:
    Near-end (CIR)                      : 0
    Far-end (CIR)                       : 0
    Near-end (EIR)                      : 0
    Far-end (EIR)                       : 0
  Accumulated loss statistics:
    Near-end loss (CIR)                 : 0
    Near-end loss-ratio (CIR)           : 0 (0.00000%)
    Far-end loss (CIR)                  : 0
    Far-end loss-ratio (CIR)            : 0 (0.00000%)
    Near-end loss (CIR)                 : 0
    Near-end loss-ratio (EIR)           : 0 (0.00000%)
    Far-end loss (CIR)                  : 0
    Far-end loss-ratio (EIR)            : 0 (0.00000%)
  Last loss measurement statistics:
    Near-end (CIR)                      : 0
    Far-end (CIR)                       : 0
    Near-end (EIR)                      : 0
    Far-end (EIR)
```

Sample Output

show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) - Delay Measurement (DM)in Metro Ethernet Forum (MEF) mode

```
user@router> show oam ethernet connectivity-fault-management sla-iterator-statistics maintenance-
domain md6-1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator DM
Iterator statistics:
```

```
  Maintenance domain: md6-1, Level: 6
  Maintenance association: ma1, Local MEP id: 1
  Remote MEP id: 2, Remote MAC address: 00:23:9c:db:0d:7a
  Iterator name: Test_DM, Iterator Id: 1
  Iterator cycle time: 1000ms, Iteration period: 2000 cycles
  Iterator status: running, Infinite iterations: true
  Counter reset time: 2018-12-05 19:48:23 PST (00:00:54 ago)
  Reset reason: Adjacency flap
  Current delay measurement statistics:
    Measurement Interval Index                : 2 (Suspect status : 1)
    Measurement Interval Start Time           : 2018-12-05 19:48:23 PST (Elapsed time : 53011
msec)
    Frame Delay two way (min, max, avg)       : 251 , 295 , 262  (usec)
    Frame Delay forward (min, max, avg)       : 125 , 147 , 131  (usec)
    Frame Delay backward (min, max, avg)      : 125 , 147 , 131  (usec)
    Inter Frame Delay two way (min, max, avg)  : 0 , 42 , 7  (usec)
    Inter Frame Delay forward (min, max, avg)  : 0 , 21 , 3  (usec)
    Inter Frame Delay backward (min, max, avg) : 0 , 21 , 3  (usec)
    Frame Delay Range two way (max, avg)      : 0 , 0  (usec)
    Frame Delay Range forward (max, avg)      : 0 , 0  (usec)
    Frame Delay Range backward (max, avg)     : 0 , 0  (usec)
    SOAM TXed                                 : 52
    SOAM RXed                                 : 52

  Delay measurement bin statistics:
    Measurement Interval Index     : 2
    Two Way Frame Delay
       [0           - 4999       ] (usec)   : 52
       [5000        - 9999       ] (usec)   : 0
       [10000       - Infinity   ] (usec)   : 0
    Forward Frame Delay
       [0           - 4999       ] (usec)   : 52
       [5000        - 9999       ] (usec)   : 0
       [10000       - Infinity   ] (usec)   : 0
    Backward Frame Delay
       [0           - 4999       ] (usec)   : 52
       [5000        - 9999       ] (usec)   : 0
       [10000       - Infinity   ] (usec)   : 0
    Two Way Inter Frame Delay Variation
       [0           - 4999       ] (usec)   : 51
       [5000        - 9999       ] (usec)   : 0
       [10000       - Infinity   ] (usec)   : 0
    Forward Inter Frame Delay Variation
```

```
        [0           - 4999       ] (usec)   : 51
        [5000        - 9999       ] (usec)   : 0
        [10000       - Infinity   ] (usec)   : 0
     Backward Inter Frame Delay Variation
        [0           - 4999       ] (usec)   : 51
        [5000        - 9999       ] (usec)   : 0
        [10000       - Infinity   ] (usec)   : 0
     Two Way Frame Delay Range
        [0           - 4999       ] (usec)   : 0
        [5000        - 9999       ] (usec)   : 0
        [10000       - Infinity   ] (usec)   : 0
     Forward Frame Delay Range
        [0           - 4999       ] (usec)   : 0
        [5000        - 9999       ] (usec)   : 0
        [10000       - Infinity   ] (usec)   : 0
     Backward Frame Delay Range
        [0           - 4999       ] (usec)   : 0
        [5000        - 9999       ] (usec)   : 0
        [10000       - Infinity   ] (usec)   : 0
```

show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) - Synthetic loss measurement (SLM) in Metro Ethernet Forum (MEF) mode

```
user@router> show oam ethernet connectivity-fault-management sla-iterator-statistics maintenance-
domain md6-1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator SLM
Iterator statistics:
  Maintenance domain: md6-1, Level: 6
  Maintenance association: ma1, Local MEP id: 1
  Remote MEP id: 2, Remote MAC address: 00:23:9c:db:0d:7a
  Iterator name: Test_SLM, Iterator Id: 2
  Iterator cycle time: 1000ms, Iteration period: 2000 cycles
  Iterator status: running, Infinite iterations: true
  Counter reset time: 2018-12-05 19:48:23 PST (00:00:11 ago)
  Reset reason: Adjacency flap
  Iterator synthetic loss measurement statistics:
    SLM skipped for threshold hit        : 0
    SLM skipped for threshold hit window : 0
    SLR out of sequence                  : 0
    SLM Sample Size                      : 180 SLMs
    Local packet Stats (TxFCl, RxFCl)    : 11, 11
```

```
    Last SLR packet stats (TxFCf, TxFCb) : 11, 11
    Last measured FLR (fwd, bkwd)        : 0.000%, 0.000% (Sample #NA)


  Current Measurement Interval loss statistics:
    Measurement Interval Index           : 2 (Suspect Status: 1)
    Measurement Interval Start Time      : 2018-12-05 19:48:23 PST (Elapsed time: 10042 msec)
    SOAM Frames (Tx, Rx)                 : 10, 10
    Forward Frame Stats (Tx, Rx)         : 10, 10
    Backward Frame Stats (Tx, Rx)        : 10, 10
    Frame Loss (fwd, bkwd)               : 0, 0
    Forward FLR minimum                  : 0.000%
    Forward FLR maximum                  : 0.000%
    Forward FLR average                  : 0.000%
    Backward FLR minimum                 : 0.000%
    Backward FLR maximum                 : 0.000%
    Backward FLR average                 : 0.000%

  Current Measurement Interval availability statistics:
    Measurement Interval Index           : 2 (Suspect Status: 1)
    Measurement Interval Start Time      : 2018-12-05 19:48:23 PST (Elapsed time: 10042 msec)
    High loss (fwd, bkwd)                : 0 , 0
    Consecutive high loss (fwd, bkwd)    : 0 , 0
    Available (fwd, bkwd)                : 0 , 0
    Unavailable (fwd, bkwd)              : 0 , 0
    Forward FLR minimum                  : 0.000%
    Forward FLR maximum                  : 0.000%
    Forward FLR average                  : 0.000%
    Backward FLR minimum                 : 0.000%
    Backward FLR maximum                 : 0.000%
    Backward FLR average                 : 0.000%

    Last known available status  (fwd, bkwd)    : unknown, unknown
    Last known forward availability transition  : NA
    Last known backward availability transition : NA
```

show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) - Delay measurement (DM) statistics in non-Metro Ethernet Forum (MEF) mode

```
user@router> show oam ethernet connectivity-fault-management sla-iterator-statistics maintenance-
domain md6-1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator DM
```

```
Iterator statistics:
  Maintenance domain: md6-1, Level: 6
  Maintenance association: ma1, Local MEP id: 1
  Remote MEP id: 2, Remote MAC address: 00:23:9c:db:0d:7a
  Iterator name: Test_DM, Iterator Id: 1
  Iterator cycle time: 1000ms, Iteration period: 2000 cycles
  Iterator status: running, Infinite iterations: true
  Counter reset time: 2018-12-05 19:43:47 PST (00:02:44 ago)
  Reset reason: Adjacency flap
  Iterator delay measurement statistics:
    Calculation weight: Delay: 1, Delay variation: 1
    DMM sent                                        : 164
    DMM skipped for threshold hit                   : 0
    DMM skipped for threshold hit window            : 0
    DMR received                                    : 164
    DMR out of sequence                             : 0
    DMR forwarding-class mismatch                   : 0
    DMR received with invalid time stamps           : 0
    Average two-way delay                           : 234 usec
    Average two-way delay variation                 : 9 usec
    Average one-way forward delay variation         : 346 usec
    Average one-way backward delay variation        : 346 usec
    Weighted average two-way delay                  : 221 usec
    Weighted average two-way delay variation        : 2 usec
    Weighted average one-way forward delay variation : 357 usec
    Weighted average one-way backward delay variation: 355 usec
    Bestcase two-way delay                          : 210 usec
    Worstcase two-way delay                         : 283 usec
    Weighted Bestcase two-way delay                 : 210 usec
    Weighted Worstcase two-way delay                : 283 usec
```

show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) - Synthetic loss measurement (SLM) statistics in non-Metro Ethernet Forum (MEF) mode

```
user@router> show oam ethernet connectivity-fault-management sla-iterator-statistics maintenance-
domain md6-1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator SLM

Iterator statistics:
  Maintenance domain: md6-1, Level: 6
  Maintenance association: ma1, Local MEP id: 1
  Remote MEP id: 2, Remote MAC address: 00:23:9c:db:0d:7a
```

```
  Iterator name: Test_SLM, Iterator Id: 2
  Iterator cycle time: 1000ms, Iteration period: 2000 cycles
  Iterator status: running, Infinite iterations: true
  Counter reset time: 2018-12-05 19:43:47 PST (00:03:23 ago)
  Reset reason: Adjacency flap
  Iterator synthetic loss measurement statistics:
    SLM sent                            : 202
    SLM skipped for threshold hit       : 0
    SLM skipped for threshold hit window: 0
    SLR received                        : 202
    SLR out of sequence                 : 0
  SLM  transmit statistics:
    SLM TXed                            : 202
    SLM RXed                            : 202
    SLM Last packet Tx count            : 202
    SLM Last packet Rx count            : 202
  Last loss measurement statistics:
    Near-end loss                       : 0 (0.00%)
    Far-end  loss                       : 0 (0.00%)
```

show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) - Delay Measurement (DM) with “legacy-pm-display” option in enhanced-cfm mode

```
user@router> show oam ethernet connectivity-fault-management sla-iterator-statistics maintenance-
domain md6-1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator Legacy
Iterator statistics:
  Maintenance domain: md6-1, Level: 6
  Maintenance association: ma1, Local MEP id: 1
  Remote MEP id: 2, Remote MAC address: 00:23:9c:db:0d:7a
  Iterator name: Test_DM_Legacy, Iterator Id: 3
  Iterator cycle time: 1000ms, Iteration period: 2000 cycles
  Iterator status: running, Infinite iterations: true
  Counter reset time: 2018-12-05 19:51:33 PST (00:00:15 ago)
  Reset reason: Adjacency flap
  Iterator delay measurement statistics:
    Calculation weight: Delay: 1, Delay variation: 1
    DMM sent                                        : 14
    DMM skipped for threshold hit                   : 0
    DMM skipped for threshold hit window            : 0
    DMR received                                    : 14
```

```
    DMR out of sequence                              : 0
    DMR forwarding-class mismatch                    : 0
    DMR received with invalid time stamps            : 0
    Average two-way delay                            : 324 usec
    Average two-way delay variation                  : 7 usec
    Average one-way forward delay variation          : 240 usec
    Average one-way backward delay variation         : 240 usec
    Weighted average two-way delay                   : 324 usec
    Weighted average two-way delay variation         : 6 usec
    Weighted average one-way forward delay variation : 222 usec
    Weighted average one-way backward delay variation: 222 usec
    Bestcase two-way delay                           : 312 usec
    Worstcase two-way delay                          : 333 usec
    Weighted Bestcase two-way delay                  : 312 usec
    Weighted Worstcase two-way delay                 : 333 usec
```

show oam ethernet connectivity-fault-management state - To verify the Connectivity Fault Management (CFM) state

The following command is to verify whether the CFM state is in enhanced-cfm mode or not.

```
user@router> show oam ethernet connectivity-fault-management state
Connectivity fault management state:
    CFM Mode Of Operation: Enhanced
    Enhanced IP Mode: Enabled
    CFM Config State: Ok
    CFM Cleanup State: Ok
    CFM Restart Timer State: Cleanup Timer State stopped Rebooting in 0 sec
    CFM CFMMAN Job State: Not Pending
    Number of sessions: 1
    Number of sessions created: 1
    Number of sessions deleted: 0
    Number of sessions freed: 0
    Number of sessions enqueued: 1
    Number of sessions dequeued: 1
    VPLS feature: enabled
    Token based forwarding feature: enabled
    Forwarding table filtering simulation feature: disabled
    Hardware assisted flooding feature: enabled
    Flood resynchronization for GRES feature: enabled
    Shared interface filter feature: disabled
```

```
    Hardware timestamping feature: disabled
    Marking of connection protection TLV feature: disabled
    CFMD config memory resource  limit(in bytes): 3221225472
    CFMD max resident set (peak) size (in bytes):  24158208

Packet processing state:
    State of the connection to packet processing daemon: down
    State of the flow to packet processing daemon: ready
    State of the packet processing job: ready
    Number of times the connection to packet processing daemon was blocked: 0
    State of the connection to cfmman: slots: 2 3 4 5

Filter state:
    State of the connection to firewall daemon: Connected
    Number of reconnects made to firewall daemon: 0
    Number of requests sent to firewall daemon: 13
    Number of requests accepted by firewall daemon: 13
    Number of requests rejected by firewall daemon: 0
    Number of requests lost due to disconnection: 0
```

Release Information

Command introduced in Junos OS Release 11.4 .

RELATED DOCUMENTATION

Configure an Iterator Profile on a Switch (CLI Procedure)

clear oam ethernet connectivity-fault-management sla-iterator-statistics | 3774

show oam ethernet connectivity-fault-management synthetic-loss-statistics

IN THIS SECTION

Syntax | 22678

- Description | 22678
- Options | 22678
- Required Privilege Level | 22679
- Output Fields | 22679
- Sample Output | 22680
- Release Information | 22681

Syntax

```
show oam ethernet connectivity-fault-management synthetic-loss-statistics
<local-mep local-mep-id>
maintenance-association ma-name
<count entry-count>
maintenance-domain md-name
<remote-mep remote-mep-id>
```

Description

On MX Series routers with Modular Port Concentrators (MPCs) with Ethernet interfaces, display the on-demand ETH-SLM statistics.

Options

count *entry-count*	(Optional) Number of entries to display from the statistics table. The range of values is from 1 through 100. The default value is 100.
local-mep *local-mep-id*	(Optional) Numeric identifier of the local MEP. The range of values is from 1 through 8192.
maintenance-association *ma-name*	Name of an existing CFM maintenance association.
maintenance-domain *md-name*	Name of an existing connectivity fault management (CFM) maintenance domain.
remote-mep *remote-mep-id*	(Optional) Numeric identifier of the remote MEP. The range of values is from 1 through 8192.

Required Privilege Level

view

Output Fields

Table 910 on page 22679 lists the output fields for the `show oam ethernet connectivity-fault-management synthetic-loss-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 910: show oam ethernet connectivity-fault-management synthetic-loss-statistics Output Fields

Output Field Name	Field Description
`MEP identifier`	Maintenance association end point (MEP) numeric identifier.
`MAC address`	Unicast MAC address configured for the MEP.
`Remote MEP count`	Number of remote MEPs (unless you specify the `remote-mep` option).
`Remote MEP identifier`	Numeric identifier of the remote MEP.
`Remote MAC address`	Unicast MAC address of the remote MEP.
`SLM packets sent`	Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.
`SLM packets received`	Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.
`SLR packets sent`	Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.
`SLR packets received`	Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.

Table 910: show oam ethernet connectivity-fault-management synthetic-loss-statistics Output Fields *(Continued)*

Output Field Name	Field Description
`Local TXFC1 value`	Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.
`Local RXFC1 value`	Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TxFC1 is sent in the packet.
`Last Received SLR frame TXFCf(tc)`	Value of the local counter TxFCl at the time of SLM frame transmission.
`Last Received SLR frame TXFCb(t`	Value of the local counter RxFCl at the time of SLR frame transmission.
`Frame loss (near-end)`	Count of frame loss associated with ingress data frames.
`Frame loss (far-end)`	Count of frame loss associated with egress data frames.

Sample Output

show oam ethernet connectivity-fault- management synthetic-loss-statistics

```
user@switch> show oam ethernet connectivity-fault-management synthetic-loss-statistics
maintenance-domain md6 maintenance-association ma6
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 2
  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
  Synthetic Loss measurement statistics:
   SLM packets sent                            : 100
```

```
    SLM packets received                        : 0
    SLR packets sent                            : 100
    SLR packets received                        : 0
   Accumulated SLM statistics:
    Local TXFC1 value                           : 100
    Local RXFC1 value                           : 100
    Last Received SLR frame TXFCftc             : 100
    Last Received SLR frame TXFCbtc             : 100
   SLM Frame Loss:
    Frame Loss (far-end)                        : 0 (0.00 %)
    Frame Loss (near-end)                       : 0 (0.00 %)
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

- clear oam ethernet connectivity-fault-management statistics | 3775
- clear oam ethernet connectivity-fault-management synthetic-loss-measurement | 3777
- show oam ethernet connectivity-fault-management interfaces | 22596
- show oam ethernet connectivity-fault-management mep-database | 22620
- show oam ethernet connectivity-fault-management mep-statistics | 22637

show oam ethernet evc

IN THIS SECTION

- Syntax | 22682
- Description | 22682
- Options | 22682
- Required Privilege Level | 22682
- Output Fields | 22682
- Sample Output | 22683

Release Information | 22683

Syntax

```
show oam ethernet evc <evc-id>
```

Description

On MX Series routers with OAM Ethernet Virtual Connection (EVC) configurations, displays the EVC configuration and status information.

Options

This command has no options.

Required Privilege Level

View

Output Fields

Table 911 on page 22682 lists the output fields for the `show oam ethernet evc` command. Output fields are listed in the approximate order in which they appear.

Table 911: show oam ethernet evc Output Fields

Field Name	Field Description
EVC identifier	Header for the EVC information showing the EVC name, configuration, and active/inactive status.
UNI count	Number of configured and active UNIs.
Protocol	Protocol configured between the UNIs.

Table 911: show oam ethernet evc Output Fields *(Continued)*

Field Name	Field Description
Local UNIs	Heading for the list of local UNIs
UNI Identifier	Name of the UNI.
Interface	Interface type-dpc/pic/port.unit-number.
Status	Status operational or not operational.

Sample Output

show oam ethernet evc

```
user@host> show oam ethernet evc
EVC identifier: evc1, Point-to-Point, Active
  UNI count: Configured(2), Active(2)
  Protocol: cfm, Management domain: md, Management association: ma
  Local UNIs:
    UNI Identifier      Interface     Status
    uni1                ge-1/1/1      Operational
    uni2                ge-1/1/1      Not Operational
```

Release Information

Command introduced in Junos OS Release 9.5.

show oam ethernet fnp interface

IN THIS SECTION

- Syntax | 22684
- Description | 22684
- Options | 22684
- Required Privilege Level | 22684
- Output Fields | 22685
- Sample Output | 22685
- Release Information | 22685

Syntax

```
show oam ethernet fnp interface
<ethernet-interface-name>
<routing-instance routing-instance-name>
```

Description

On MX Series routers with Gigabit Ethernet or aggregated Ethernet PICs, displays OAM Ethernet Failure Notification Protocol (FNP) information for Ethernet interfaces.

Options

interface-name	(Optional) Display Ethernet FNP information for the specified Ethernet interface only.
routing-instance-name	(Optional) Display FNP for the specified routing instance.

Required Privilege Level

view

Output Fields

Table 912 on page 22685 lists the output fields for the `show oam ethernet fnp interface` command. Output fields are listed in the approximate order in which they appear.

Table 912: show oam ethernet fnp interface Output Fields

Field Name	Field Description
Interface	Name of the interface for the displayed information.
VLAN	Name of the VLAN.
State	Displays state of the interface.
FNP Message Interface	Displays the message interface type.
FNP Message Source MAC	Displays the source MAC address.

Sample Output

show oam ethernet fnp interface

```
user@host> show oam ethernet fnp interface
  The FNP controlled interfaces are:

  Interface    VLAN   State   FNP message   FNP message
                              Interface     Source MAC

  ge-0/0/0.30  30     down    lsi.1054976   a0:aa:aa:aa:aa:aa
  ge-0/0/0.20  20     down    lsi.1054976   a0:aa:aa:aa:aa:aa
```

Release Information

Command introduced in Junos OS Release 11.4.

show oam ethernet fnp messages

IN THIS SECTION

- Syntax | 22686
- Description | 22686
- Options | 22686
- Required Privilege Level | 22686
- Output Fields | 22687
- Sample Output | 22688
- Release Information | 22688

Syntax

```
show oam ethernet fnp messages
<interface interface-name>
<routing instance routing-instance-name>
```

Description

On MX Series routers with Gigabit Ethernet or aggregated Ethernet PICs, displays OAM Ethernet Failure Notification Protocol (FNP) messages.

Options

interface-name	(Optional) Display Ethernet FNP messages for the specified Ethernet interface only.
routing-instance-name	(Optional) Display FNP messages for the specified routing instance.

Required Privilege Level

view

Output Fields

Table 913 on page 22687 lists the output fields for the `show oam ethernet fnp messages` command. Output fields are listed in the approximate order in which they appear.

Table 913: show oam ethernet fnp messages Output Fields

Field Name	Field Description
Message from source MAC address	The source MAC address of the message.
Originating port number	Port number of the original message.
Time since last message	Elapsed time in hours, minutes, and seconds since the last message was received.
Time since last message update	Elapsed time in hours, minutes, and seconds since the last message was updated.
Total messages received	Number of messages received.
Domain ID	Domain ID of the message.
STP Root ID	The spanning tree Root ID of the message.
Trigger Reason	The reason why the message was triggered.
Effected VLANs	Number of VLANs that are affected.
Disabled interfaces	Name of the interfaces that are disabled.

Sample Output

show oam ethernet fnp messages

```
user@host> show oam ethernet fnp messages
  Active FNP messages on interface lsi.1054465
  Message source MAC: a0:aa:aa:aa:aa:aa
    Originating port number: 141077
    Time since last message: 00:00:00
    Time since last message update: 00:00:00
    Total messages received: 1
    Domain ID: 0
    STP Root ID: 0.f0:ff:ff:ff:ff:ff
    Trigger reason: todo
    Effected VLANs: 10
    Disabled interfaces:
      Interface VLAN
      ge-0/0/0.10 10
```

Release Information

Command introduced in Junos OS Release 11.4

show oam ethernet fnp status

IN THIS SECTION

- Syntax | 22689
- Description | 22689
- Options | 22689
- Required Privilege Level | 22689
- Output Fields | 22689
- Sample Output | 22690

Release Information | 22690

Syntax

```
show oam ethernet fnp status
<interface interface-name>
<routing instance routing-instance-name>
```

Description

On MX Series routers with Gigabit Ethernet or aggregated Ethernet PICs, displays OAM Ethernet Failure Notification Protocol (FNP) status.

Options

interface-name	(Optional) Display Ethernet FNP information for the specified Ethernet interface only.
routing-instance-name	(Optional) Display FNP for the specified routing instance.

Required Privilege Level

view

Output Fields

Table 914 on page 22689 lists the output fields for the `show oam ethernet fnp status` command. Output fields are listed in the approximate order in which they appear.

Table 914: show oam ethernet fnp status Output Fields

Field Name	Field Description
FNP interval	The time interval between messages.

Table 914: show oam ethernet fnp status Output Fields *(Continued)*

Field Name	Field Description
Loss threshold	The number of messages that can be lost before FNP is marked as down.
FNP enabled interfaces	Displays interfaces that are enabled.
Interface	The name of the interface.
Domain ID	Domain ID of the message.
STP Root ID	The spanning tree Root ID of the message.
FNP Messages	The total number of messages received.

Sample Output

show oam ethernet fnp status

```
user@host> show oam ethernet status
  FNP interval:
  Loss threshold
  FNP enabled interfaces
  Interface      Domain ID      STP Root ID              FNP Messages
  ge-0/0/0.1278        100      0.f0:ff:ff:ff:ff:ff                 0
```

Release Information

Command introduced in Junos OS Release 11.4

show oam ethernet link-fault-management

IN THIS SECTION

- Syntax | 22691
- Description | 22691
- Options | 22691
- Required Privilege Level | 22691
- Output Fields | 22692
- Sample Output | 22698
- Release Information | 22701

Syntax

```
show oam ethernet link-fault-management
<brief | detail>
<interface-name>
```

Description

Display Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.

Options

brief | detail (Optional) Display the specified level of output.

interface-name (Optional) Display link fault management information for the specified Ethernet interface only.

Required Privilege Level

view

Output Fields

Table 915 on page 22692 lists the output fields for the `show oam ethernet link-fault-management` command. Output fields are listed in the approximate order in which they appear.

Table 915: show oam ethernet link-fault-management Output Fields

Field Name	Field Description	Level of Output
`Status`	Status of the established link. • `Fail`—A link fault condition exists. • `Running`—A link fault condition does not exist.	All levels
`Discovery state`	State of the discovery mechanism: • `Passive Wait` • `Send Any` • `Send Local Remote` • `Send Local Remote Ok`	All levels
`Peer address`	Address of the OAM peer.	All levels

Table 915: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Information about the interface. • `Remote-Stable`—Indicates remote OAM client acknowledgment of, and satisfaction with, local OAM state information. `False` indicates that remote DTE has either not seen or is unsatisfied with local state information. `True` indicates that remote DTE has seen and is satisfied with local state information. • `Local-Stable`—Indicates local OAM client acknowledgment of, and satisfaction with, remote OAM state information. `False` indicates that local DTE either has not seen or is unsatisfied with remote state information. `True` indicates that local DTE has seen and is satisfied with remote state information. • `Remote-State-Valid`—Indicates the OAM client has received remote state information found within local information TLVs (type, length, values) of received Information OAM PDUs. `False` indicates that the OAM client has not seen remote state information. `True` indicates that the OAM client has seen remote state information.	All levels
`Remote loopback status`	An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).	All levels

Table 915: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Remote entity information`	Remote entity information. • `Remote MUX action`—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs. • `Remote parser action`—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to the higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs. • `Discovery mode`—Indicates whether discovery mode is active or inactive. • `Unidirectional mode`—Indicates the ability to operate a link in unidirectional mode for diagnostic purposes. • `Remote loopback mode`—Indicates whether remote loopback is supported or not supported. • `Link events`—Indicates whether interpreting link events is supported or not supported on the remote peer. • `Variable requests`—Indicates whether variable requests are supported or not supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer.	All levels
OAM Receive Statistics		
`Information`	Number of information PDUs received.	`detail`
`Event`	Number of loopback control PDUs received.	`detail`
`Variable request`	Number of variable request PDUs received.	`detail`
`Variable response`	Number of variable response PDUs received.	`detail`
`Loopback control`	Number of loopback control PDUs received.	`detail`

Table 915: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Organization specific`	Number of vendor organization specific PDUs received.	`detail`
OAM Transmit Statistics		
`Information`	Number of information PDUs transmitted.	`detail`
`Event`	Number of event notification PDUs transmitted.	`detail`
`Variable request`	Number of variable request PDUs transmitted.	`detail`
`Variable response`	Number of variable response PDUs transmitted.	`detail`
`Loopback control`	Number of loopback control PDUs transmitted.	`detail`
`Organization specific`	Number of vendor organization specific PDUs transmitted.	`detail`
OAM Received Symbol Error Event information		
`Events`	Number of symbol error event TLVs that have been received after the OAM sublayer was reset.	`detail`
`Window`	Symbol error event window in the received PDU. The protocol default value is the number of symbols that can be received in one second on the underlying physical layer.	`detail`
`Threshold`	Number of errored symbols in the period required for the event to be generated.	`detail`

Table 915: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Errors in period`	Number of symbol errors in the period reported in the received event PDU.	`detail`
`Total errors`	Number of errored symbols that have been reported in received event TLVs after the OAM sublayer was reset. Symbol errors are coding symbol errors.	`detail`
OAM Received Frame Error Event Information		
`Events`	Number of errored frame event TLVs that have been received after the OAM sublayer was reset.	`detail`
`Window`	Duration of the window in terms of the number of 100 ms period intervals.	`detail`
`Threshold`	Number of detected errored frames required for the event to be generated.	`detail`
`Errors in period`	Number of detected errored frames in the period.	`detail`
`Total errors`	Number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset. A frame error is any frame error on the underlying physical layer.	`detail`
OAM Received Frame Period Error Event Information		
`Events`	Number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset.	`detail`
`Window`	Duration of the frame seconds window.	`detail`
`Threshold`	Number of frame seconds errors in the period.	`detail`

Table 915: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Errors in period`	Number of frame seconds errors in the period.	`detail`
`Total errors`	Number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset.	`detail`
OAM Transmitted Symbol Error Event Information		
`Events`	Number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset.	`detail`
`Window`	The symbol error event window in the transmitted PDU.	`detail`
`Threshold`	Number of errored symbols in the period required for the event to be generated.	`detail`
`Errors in period`	Number of symbol errors in the period reported in the transmitted event PDU.	`detail`
`Total errors`	Number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset.	`detail`
OAM Transmitted Frame Error Event Information		
`Events`	Number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset.	`detail`
`Window`	Duration of the window in terms of the number of 100-ms period intervals.	`detail`
`Threshold`	Number of detected errored frames required for the event to be generated.	`detail`
`Errors in period`	Number of detected errored frames in the period.	`detail`

Table 915: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Total errors`	Number of errored frames that have been detected after the OAM sublayer was reset.	`detail`

Sample Output

show oam ethernet link-fault-management brief

```
user@host> show oam ethernet link-fault-management brief
Interface: ge-3/1/3
  Status: Running, Discovery state: Send Any, ISSU
  Peer address: 00:90:69:72:2c:83
  Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50  Remote loopback status: Disabled on
local port, Enabled on peer port
  Remote entity information:
    Remote MUX action: discarding, Remote parser action: loopback
    Discovery mode: active, Unidirectional mode: unsupported
    Remote loopback mode: supported, Link events: supported
    Variable requests: unsupported, Remote in ISSU
```

show oam ethernet link-fault-management brief (Loopback tracking)

```
user@host> show oam ethernet link-fault-management
Interface: ge-3/1/3
  Status: Running, Discovery state: Active Send Local
  Peer address: 00:00:00:00:00:00
  Flags:0x8
     Loopback tracking: Enabled,      Loop Status: Found
```

show oam ethernet link-fault-management detail

```
user@host> show oam ethernet link-fault-management detail
  Interface: ge-6/1/0
    Status: Running, Discovery state: Send Any, ISSU
```

```
    Peer address: 00:90:69:0a:07:14
    Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
    OAM receive statistics:
      Information: 186365, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM transmit statistics:
      Information: 186347, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM received symbol error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame period error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM transmitted symbol error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    OAM current symbol error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
          OAM transmitted frame error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    OAM current frame error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
          Remote entity information:
      Remote MUX action: forwarding, Remote parser action: forwarding
      Discovery mode: active, Unidirectional mode: unsupported
      Remote loopback mode: supported, Link events: supported
      Variable requests: unsupported, Remote in ISSU
```

show oam ethernet link-fault-management detail (backup Routing Engine)

```
user@host> show oam ethernet link-fault-management ge-0/2/0 detail
Interface: ge-0/2/0
    Status: Running, Discovery state: Send Any
    Transmit interval: 100ms, PDU threshold: 3 frames, Hold time: 300ms
```

```
    Peer address: ac:4b:c8:81:90:a4
    Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
    OAM receive statistics:
      Information: 0, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM flags receive statistics:
      Critical event: 0, Dying gasp: 0, Link fault: 0
    OAM transmit statistics:
      Information: 0, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 786, Organization specific: 0
    OAM received symbol error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame period error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame seconds error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM transmitted symbol error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    OAM current symbol error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    OAM transmitted frame error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    OAM current frame error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    Loopback tracking: Enabled, Loop status: Not Found
    Detect LOC: Enabled, LOC status: Not Found
    Remote entity information:
      Remote MUX action: forwarding, Remote parser action: forwarding
      Discovery mode: active, Unidirectional mode: unsupported
      Remote loopback mode: unsupported, Link events: supported
      Variable requests: unsupported
Application profile statistics:
  Profile Name                   Invoked     Executed
```

```
LK_ADJ_LOSS100_1                        1        1
LK_ADJ_LOSS100_2                        1        0
LK_ADJ_LOSS100_3                        1        0
LK_ADJ_LOSS101_1                        1        1
LK_ADJ_LOSS101_2                        1        0
LK_ADJ_LOSS101_3                        1        0
LK_ADJ_LOSS106_1                        0        0
LK_ADJ_LOSS106_2                        0        0
LK_ADJ_LOSS106_3                        0        0
LK_ADJ_LOSS107_1                        0        0
LK_ADJ_LOSS107_2                        0        0
LK_ADJ_LOSS107_3                        0        0
```

Release Information

Statement for SRX Series Firewalls introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

clear oam ethernet connectivity-fault-management path-database | 3770

clear oam ethernet connectivity-fault-management statistics

Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways

Example: Configuring Ethernet OAM Link Fault Management on a Security Device

show oam ethernet lmi

IN THIS SECTION

- Syntax | 22702
- Description | 22702
- Options | 22702
- Required Privilege Level | 22702
- Output Fields | 22702
- Sample Output | 22704

Release Information | 22704

Syntax

```
show oam ethernet lmi (interface <interface-name>)
```

Description

On routers with Gigabit Ethernet or aggregated Ethernet, and OAM Ethernet Local Management Interface (E-LMI) configuration, display the LMI information for the configured interfaces or optionally for a specified interface.

NOTE: On MX Series routers, E-LMI is supported on Gigabit Ethernet (ge), 10-Gigabit Ethernet (xe), and Aggregated Ethernet (ae) interfaces configured on MX Series routers with DPC only.

Options

interface	(Optional) Display LMI information for a specified interface.
interface-name	(Optional) Display Ethernet LMI information for the specified interface only.

Required Privilege Level

View

Output Fields

Table 916 on page 22703 lists the output fields for the `show oam ethernet lmi` command. Output fields are listed in the approximate order in which they appear.

Table 916: show oam ethernet lmi Output Fields

Field Name	Field Description
Physical Interface	Header for the EVC information showing the Ethernet virtual circuit (EVC) name, configuration, and active/inactive status.
UNI Identifier	Name of the UNI.
EVC map type	EVC configuration.
Polling verification timer	Polling verification timer status.
E-LMI state	Operational status of the E-LMI configuration in the interfaces or specified interface.
Priority/Untagged VLAN ID	To be provided.
Default EVC	The EVC set as the default EVC.
Associated EVCs	Heading for the list of configured EVCs.
EVC Indentifier	EVC name.
Reference ID	To be provided.
Status	Status active or not active.
CE VLAN IDs	Customer edge VLAN ID numbers.

Sample Output

show oam ethernet lmi interface

```
user@host> show oam ethernet lmi interface ge-1/1/1
Physical interface: ge-1/1/1, Physical link is Up
  UNI identifier: uni-ce1, EVC map type: Bundling
  Polling verification timer: Enabled, E-LMI state: Operational
  Priority/Untagged VLAN ID: 20, Default EVC: evc1
  Associated EVCs:
    EVC                Reference        Status              CE VLAN IDs
    Identifier     ID
    evc1                   1           Active (New)           1-2048
    evc2                    2        Not Active             2049-4096
```

Release Information

Command introduced in Junos OS Release 9.5.

show oam ethernet lmi statistics

IN THIS SECTION

- Syntax | 22705
- Description | 22705
- Options | 22705
- Required Privilege Level | 22705
- Output Fields | 22705
- Sample Output | 22706
- Release Information | 22706

Syntax

```
show oam ethernet lmi statistics <interface interface-name>
```

Description

On MX Series routers with Gigabit Ethernet or aggregated Ethernet PICs, displays OAM Ethernet Local Management Interface (LMI) statistics.

Options

interface (Optional) Display LMI statistics for a specified interface.

interface-name (Optional) Display Ethernet LMI information for the specified Ethernet interface only.

Required Privilege Level

view

Output Fields

Table 917 on page 22705 lists the output fields for the `show oam ethernet lmi statistics` command. Output fields are listed in the approximate order in which they appear.

Table 917: show oam ethernet lmi statistics Output Fields

Field Name	Field Description
Physical interface	Name of the interface for the displayed statistics.
Reliability errors	Number of E-LMI reliability errors logged.
Protocol errors	Number of E-LMI protocol errorse.
Status check received	Number of E-LMI status check receive errors.

Table 917: show oam ethernet lmi statistics Output Fields *(Continued)*

Field Name	Field Description
Status check sent	Number of E-LMI status check sent errors.
Full status received	Number of E-LMI full status receive errors.
Full status sent	Number of E-LMI full status sent errors.
Full status continued received	Number of E-LMI status continued received errors.
Full status continued sent	Number of E-LMI full status continued sent errors.
Asynchronous status sent	Number of E-LMI asynchronous status sent errors.

Sample Output

show oam ethernet lmi statistics

```
user@host> show oam ethernet lmi statistics interface ge-1/1/1
Physical interface: ge-1/1/1
  Reliability errors                     4  Protocol errors                    0
  Status check received                  0  Status check sent                  0
  Full status received                 694  Full status sent                 694
  Full status continued received         0  Full status continued sent         0
  Asynchronous status sent               0
```

Release Information

Command introduced in Junos OS Release 9.5.

show oam fabric flow-specification

IN THIS SECTION

- Syntax | 22707
- Description | 22707
- Options | 22707
- Required Privilege Level | 22707
- Output Fields | 22708
- Sample Output | 22709
- Release Information | 22710

Syntax

```
show oam fabric flow-specification
<flow-specification-name>
```

Description

Display the fabric flow specifications that are configured in a QFabric system Node group.

Options

flow-specification-name	(Optional) Name of a flow specification.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Output Fields

Table 918 on page 22708 lists the output fields for the `show oam fabric flow-specification` command. Output fields are listed in the approximate order in which they appear.

Table 918: show oam fabric flow-specification Output Fields

Field Name	Field Description
`Flow specification name`	Name of the flow specification.
`Type`	Flow specification type: • `Ethernet` • `Ethernet Unicast IPV4` • `Multicast IPV4` • `Multicast VLAN flood`
`Ethernet frame size`	Ethernet frame size in bytes.
`Ether type`	EtherType protocol.
`Source-MAC`	MAC address of the source FMEP interface.
`Source-MAC Mask`	MAC address mask of the source FMEP interface.
`Destination-MAC`	MAC address of the destination FMEP interface.
`Destination-MAC Mask`	MAC address mask of the destination FMEP interface.
`Source-IP`	IP address of the source FMEP interface.
`Source-IP Mask`	IP address mask of the source FMEP interface.

Table 918: show oam fabric flow-specification Output Fields *(Continued)*

Field Name	Field Description
Destination-IP	IP address of the destination FMEP interface.
Destination-IP mask	IP address mask of the destination FMEP interface.
IP-Protocol	IPv4 protocol configured.
Source-L4-Port	Source TCP or UDP port.
Destination-L4-Port	Destination TCP or UDP port.
Destination multicast group IP	IP address of the destination multicast group.

Sample Output

show oam fabric flow-specification (All)

```
user@host> show oam fabric flow-specification

 Flow specification name  : fspec1 Type  : Ethernet
     Ethernet frame size  : Unspecified
     Ether type  : Unspecified
     Source-MAC  : Unspecified
     Source-MAC Mask : Unspecified
     Destination-MAC  : Unspecified
     Destination-MAC Mask : Unspecified
 Flow specification name  : fspec2 Type  : Ethernet
     Ethernet frame size  : Unspecified
     Ether type  : 1792
     Source-MAC  : 0:0:a0:f0:cc:22
     Source-MAC Mask : Unspecified
     Destination-MAC  : 0:11:22:33:a0:e2
     Destination-MAC Mask : Unspecified
 Flow specification name  : fspec3 Type : Ethernet Unicast IPV4
```

```
    Ethernet frame size  : Unspecified
    Source-IP  : 121.0.0.1
    Source-IP Mask : Unspecified
    Destination-IP  : 132.1.0.1
    Destination-IP Mask : Unspecified
    Destination-mac : Unspecified
    IP-protocol : TCP
    Source-L4-port : 270
    Destination-L4-port : Unspecified
Flow specification name  : fspec4 Type  : Multicast IPV4
    Ethernet frame size  : Unspecified
    Source IP  : 100.0.0.23
    Destination multicast group IP:             225.0.0.1
Flow specification name  : fspec5 Type  : Multicast VLAN flood
    Ethernet frame size  : Unspecified
```

show oam fabric flow-specification (Flow Specification Specified)

```
user@host> show oam fabric flow-specification fspec3

Flow specification name  : fspec3 Type : Ethernet Unicast IPV4
    Ethernet frame size  : Unspecified
    Source-IP  : 121.0.0.1
    Source-IP Mask : Unspecified
    Destination-IP  : 132.1.0.1
    Destination-IP Mask : Unspecified
    Destination-mac : Unspecified
    IP-protocol : TCP
    Source-L4-port : 270
    Destination-L4-port : Unspecified
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | **29924**

show oam fabric interfaces

IN THIS SECTION

- Syntax | **22711**
- Description | **22711**
- Options | **22712**
- Required Privilege Level | **22712**
- Output Fields | **22712**
- Sample Output | **22713**
- Release Information | **22713**

Syntax

```
show oam fabric interfaces
<interface-name>
<fabric-maintenance-association fma-name>
<brief>
```

Description

Display the fabric maintenance associations (FMAs) and fabric maintenance endpoints (FMEPs) that are configured in a QFabric system Node group.

If you do not specify an interface, display information for an FMA (if one is specified) or all FMAs configured in the Node group.

Options

brief	(Optional) Summarized version of the output.
fabric-maintenance-association ***fma-name***	(Optional) Name of a specific FMA.
interface-name	(Optional) Name of a specific interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Output Fields

Table 919 on page 22712 lists the output fields for the `show oam fabric interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 919: show oam fabric interfaces Output Fields

Field Name	Field Description
`Interface-name`	Name of the interface.
`Fabric-Maintenance Association`	FMA name.
`VLAN`	VLAN name.
`Interface state`	State of the interface: Up or down.
`MEP Identifier`	FMEP identifier.
`MEP Name`	FMEP name.

Sample Output

show oam fabric interfaces

```
user@host> show oam fabric interfaces

Interface-name    Fabric-Maintenance    VLAN    Interface    MEP           MEP
                  Association                   state        Identifier    Name
ED1479:NULL        fma-default           *       up          32774        fmep-default-ED1479
ED1494:NULL        fma-default           *       up          32773        fmep-default-ED1494
ED1494:xe-0/0/7.0  fma1                  v100    up          1
ED1497:NULL        fma-default           *       up          32775        fmep-default-ED1497
ED1497:xe-0/0/8.0  fma1                  v100    up          2
P3613-C:NULL       fma-default           *       up          32777        fmep-default-P3613-C
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | **29924**

show openflow capability

IN THIS SECTION

Syntax | **22714**

- Description | 22714
- Options | 22714
- Required Privilege Level | 22714
- Output Fields | 22714
- Sample Output | 22718
- Release Information | 22722

Syntax

```
show openflow capability
<action | feature | match-condition>
```

Description

Display support information for OpenFlow features, actions, and match conditions on the Juniper Networks device.

Options

none	Display support information for all OpenFlow capabilities.
action	(Optional) Display support information for OpenFlow actions.
feature	(Optional) Display support information for OpenFlow features.
match-condition	(Optional) Display support information for OpenFlow match conditions.

Required Privilege Level

admin

Output Fields

Table 920 on page 22715 lists the output fields for the `show openflow capability` command. Output fields are listed in the approximate order in which they appear.

Table 920: show openflow capability Output Fields

Field Name	Field Description
Supported Features—Indicates Whether the Juniper Networks Device Supports the Following OpenFlow Features	
Flow statistics	Indicates whether the device supports OpenFlow flow statistics.
Table statistics	Indicates whether the device supports OpenFlow flow table statistics.
Port statistics	Indicates whether the device supports OpenFlow port statistics.
Group statistics	Indicates whether the device supports OpenFlow group statistics. **NOTE**: This field appears only if the device supports OpenFlow v1.3.1 or later.
802.1d spanning tree	Indicates whether the device supports the 802.1D Spanning Tree Protocol.
Reassemble IP fragments	Indicates whether the device supports reassembling IP fragments.
Queue statistics	Indicates whether the device supports OpenFlow queue statistics.
Match IP addresses in ARP pkts	Indicates whether the device supports matching on IP addresses in ARP packets.
Supported Match Conditions—Indicates Whether the Juniper Networks Device Supports the Following OpenFlow Match Conditions	
Switch input port	Indicates whether the device supports matching against the ingress switch port.
VLAN vid	Indicates whether the device supports matching against the VLAN identifier in the outermost VLAN tag.
Ethernet source address	Indicates whether the device supports matching against the Ethernet source address.

Table 920: show openflow capability Output Fields *(Continued)*

Field Name	Field Description
Ethernet destination address	Indicates whether the device supports matching against the Ethernet destination address.
Ethernet frame type	Indicates whether the device supports matching against the Ethernet frame type.
IP protocol	Indicates whether the device supports matching against the IP protocol or lower 8 bits of the ARP opcode.
TCP/UDP source port	Indicates whether the device supports matching against the TCP or UDP source port.
TCP/UDP destination port	Indicates whether the device supports matching against the TCP or UDP destination port.
IP(v4) source address	Indicates whether the device supports matching against the IPv4 source address. **NOTE**: `IPv4` appears only if the device supports OpenFlow v1.3.1 or later.
IP(v4) destination address	Indicates whether the device supports matching against the IPv4 destination address. **NOTE**: `IPv4` appears only if the device supports OpenFlow v1.3.1 or later.
IPv6 source address	Indicates whether the device supports matching against the IPv6 source address. **NOTE**: This field appears only if the device supports OpenFlow v1.3.1 or later.
IPv6 destination address	Indicates whether the device supports matching against the IPv6 destination address. **NOTE**: This field appears only if the device supports OpenFlow v1.3.1 or later.

Table 920: show openflow capability Output Fields *(Continued)*

Field Name	Field Description
VLAN priority	Indicates whether the device supports matching against the VLAN priority in the outermost VLAN tag.
IP ToS (DSCP field)	Indicates whether the device supports matching against the IPv4 ToS bits.
Supported Actions—Indicates Whether the Juniper Networks Device Supports the Following OpenFlow Actions	
Output to switch port	Indicates whether the device supports forwarding the packet to a specified port.
Set the 802.1q VLAN id	Indicates whether the device supports the optional Modify-Field action to modify the existing 802.1Q VLAN ID of the outermost VLAN tag in the frame header or to add a new header with the VLAN ID if none exists.
Set the 802.1q priority	Indicates whether the device supports the optional Modify-Field action to modify the existing 802.1Q VLAN priority of the outermost VLAN tag in the frame header or to add a new header with the VLAN priority if none exists.
Strip the 802.1q header	Indicates whether the device supports the optional Modify-Field action to remove the outermost VLAN header in the frame.
Ethernet source address	Indicates whether the device supports the optional Modify-Field action to modify the Ethernet source address field in the frame header.
Ethernet destination address	Indicates whether the device supports the optional Modify-Field action to modify the Ethernet destination address field in the frame header.
IP source address	Indicates whether the device supports the optional Modify-Field action to modify the IP source address field and update the checksum in the packet header.
IP destination address	Indicates whether the device supports the optional Modify-Field action to modify the IP destination address field and update the checksum in the packet header.

Table 920: show openflow capability Output Fields *(Continued)*

Field Name	Field Description
IP ToS (DSCP)	Indicates whether the device supports the optional Modify-Field action to modify the IPv4 ToS field in the packet header.
TCP/UDP source port	Indicates whether the device supports the optional Modify-Field action to modify the TCP or UDP source port field and update the checksum in the packet header.
TCP/UDP destination port	Indicates whether the device supports the optional Modify-Field action to modify the TCP or UDP destination port field and update the checksum in the packet header.
Output to queue	Indicates whether the device supports the optional Enqueue action to set the queue ID for the packet.
Execute group	Indicates whether the device supports a group action to be executed. **NOTE**: This field appears only if the device supports OpenFlow v1.3.1 or later.

Sample Output

show openflow capability

```
user@host> show openflow capability
Openflowd platform feature support:
Flow statistics:    Yes
Table statistics:   Yes
Port statistics:    Yes
802.1d spanning tree:    No
Reassemble IP fragments:    No
Queue statistics:    Yes
Match IP addresses in ARP pkts:    No

Openflowd platform match condition support:
Switch input port:    Yes
```

```
VLAN vid:    Yes
Ethernet source address:    Yes
Ethernet destination address:  Yes
Ethernet frame type: Yes
IP protocol:    Yes
TCP/UDP source port:    Yes
TCP/UDP destination port:   Yes
IP source address:    Yes
IP destination address: Yes
VLAN priority:    Yes
IP ToS (DSCP field):    Yes

Openflowd platform action support:
Output to switch port:    Yes
Set the 802.1q VLAN id    Yes
Set the 802.1q priority:    No
Strip the 802.1q header:    Yes
Ethernet source address:    No
Ethernet destination address:    No
IP source address:    No
IP destination address:    No
IP ToS (DSCP):    No
TCP/UDP source port:    No
TCP/UDP destination port:    No
Output to queue:    No
```

show openflow capability (OpenFlow 1.3.1)

```
user@host> show openflow capability
Openflowd platform feature support:
Flow statistics:    Yes
Table statistics:   Yes
Port statistics:    Yes
Group statistics:    Yes
802.1d spanning tree:    No
Reassemble IP fragments:    No
Queue statistics:    Yes
Match IP addresses in ARP pkts:    No

Openflowd platform match condition support:
Switch input port:    Yes
```

```
VLAN vid:    Yes
Ethernet source address:    Yes
Ethernet destination address:  Yes
Ethernet frame type: Yes
IP protocol:    Yes
TCP/UDP source port:    Yes
TCP/UDP destination port:   Yes
IPv4 source address:    Yes
IPv4 destination address: Yes
IPv6 source address:    Yes
IPv6 destination address: Yes
VLAN priority:    Yes
IP ToS (DSCP field):    Yes

Openflowd platform action support:
Output to switch port:    Yes
Set the 802.1q VLAN id    Yes
Set the 802.1q priority:    No
Strip the 802.1q header:    Yes
Ethernet source address:    No
Ethernet destination address:    No
IP source address:    No
IP destination address:    No
IP ToS (DSCP):    No
TCP/UDP source port:    No
TCP/UDP destination port:    No
Output to queue:    No
Execute Group:    Yes
```

show openflow capability action

```
user@host> show openflow capability action
Openflowd platform action support:
Output to switch port:    Yes
Set the 802.1q VLAN id    Yes
Set the 802.1q priority:    No
Strip the 802.1q header:    Yes
Ethernet source address:    No
Ethernet destination address:    No
IP source address:    No
IP destination address:    No
```

```
IP ToS (DSCP):    No
TCP/UDP source port:    No
TCP/UDP destination port:    No
Output to queue:    No
```

show openflow capability feature

```
user@host> show openflow capability feature
Openflowd platform feature support:
Flow statistics:    Yes
Table statistics:   Yes
Port statistics:    Yes
802.1d spanning tree:    No
Reassemble IP fragments:    No
Queue statistics:    Yes
Match IP addresses in ARP pkts:    No
```

show openflow capability match-condition

```
user@host> show openflow capability match-condition
Openflowd platform match condition support:
Switch input port:    Yes
VLAN vid:    Yes
Ethernet source address:    Yes
Ethernet destination address:  Yes
Ethernet frame type: Yes
IP protocol:    Yes
TCP/UDP source port:    Yes
TCP/UDP destination port:   Yes
IP source address:    Yes
IP destination address: Yes
VLAN priority:    Yes
IP ToS (DSCP field):    Yes
```

show openflow capability match-condition (OpenFlow 1.3.1)

```
user@host> show openflow capability match-condition
Openflowd platform match condition support:
```

```
Switch input port:    Yes
VLAN vid:    Yes
Ethernet source address:    Yes
Ethernet destination address:  Yes
Ethernet frame type: Yes
IP protocol:    Yes
TCP/UDP source port:    Yes
TCP/UDP destination port:   Yes
IPv4 source address:    Yes
IPv4 destination address: Yes
IPv6 source address:    Yes
IPv6 destination address: Yes
VLAN priority:    Yes
IP ToS (DSCP field):    Yes
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

OpenFlow Support on Juniper Networks Devices

show openflow controller

IN THIS SECTION

- Syntax | 22723
- Description | 22723
- Options | 22723
- Required Privilege Level | 22723
- Output Fields | 22723
- Sample Output | 22724
- Release Information | 22725

Syntax

```
show openflow controller
<address address>
<switch switch-name>
```

Description

Display OpenFlow controller information and connection status. OpenFlow controllers are configured at the [edit protocols openflow switch *switch-name*] hierarchy level.

Options

none	Display information about all configured controllers.
address *address*	(Optional) Display information about the controller at the specified IP address.
switch *switch-name*	(Optional) Display information about controllers associated with the specified virtual switch.

Required Privilege Level

admin

Output Fields

Table 921 on page 22723 lists the output fields for the `show openflow controller` command. Output fields are listed in the approximate order in which they appear.

Table 921: show openflow controller Output Fields

Field Name	Field Description
`Controller socket`	Socket on the controller to which the OpenFlow virtual switch connects.
`Controller IP address`	IP address of the OpenFlow controller.

Table 921: show openflow controller Output Fields *(Continued)*

Field Name	Field Description
`Controller protocol`	Protocol used by the switch to initiate a connection with the controller.
`Controller port`	Port on the controller to which the OpenFlow virtual switch connects.
`Controller connection state`	Status of the connection between the OpenFlow virtual switch and the controller.
`Number of connection attempt`	Number of connection attempts made by the virtual switch to the controller.
`Controller role`	User-configured role of the controller.
`Negotiated version`	A numerical value that represents the OpenFlow version that is negotiated between the Junos OS device and the OpenFlow controller during the initial connection. **NOTE**: This field appears only if the Junos OS device supports OpenFlow version 1.3.1 or later. .
`Negotiated version`	A numerical value that represents the OpenFlow version that is negotiated between the Junos OS device and the OpenFlow controller during the initial connection. **NOTE**: This field appears only if the Junos OS device supports OpenFlow version 1.3.1 or later.

Sample Output

show openflow controller (OpenFlow 1.3.1)

```
user@host> show openflow controller
Openflowd controller information:
Controller socket: 15
Controller IP address: 198.51.100.174
Controller protocol: tcp
Controller port: 6633
```

```
Controller connection state: up
Number of connection attempt: 5
Controller role: equal
Negotiated version: 4
```

show openflow controller address (OpenFlow 1.3.1)

```
user@host> show openflow controller address 198.51.100.174
Openflowd controller information:
Controller socket: 15
Controller IP address: 198.51.100.174
Controller protocol: tcp
Controller port: 6633
Controller connection state: up
Number of connection attempt: 5
Controller role: equal
Negotiated version: 4
```

show openflow controller switch (OpenFlow 1.3.1)

```
user@host> show openflow controller switch OFswitch1
Openflowd controller information:
Controller socket: 15
Controller IP address: 198.51.100.174
Controller protocol: tcp
Controller port: 6633
Controller connection state: up
Number of connection attempt: 5
Controller role: equal
Negotiated version: 4
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

controller (Protocols OpenFlow) | 4654

OpenFlow Support on Juniper Networks Devices

Understanding the OpenFlow Version Negotiation Between the Controller and Devices Running Junos OS

show openflow filters

IN THIS SECTION

- Syntax | 22726
- Description | 22726
- Options | 22726
- Required Privilege Level | 22727
- Output Fields | 22727
- Sample Output | 22728
- Release Information | 22729

Syntax

```
show openflow filters
<interface interface-name>
<switch switch-name>
```

Description

Display information for filters bound to OpenFlow interfaces.

Options

none	Display information for all filters that are bound to OpenFlow interfaces.
interface *interface-name*	(Optional) Display information for the filter bound to the specified OpenFlow interface. The interface name must include the logical unit number.

switch *switch-name* (Optional) Display information for filters bound to the interfaces configured under the specified OpenFlow virtual switch.

Required Privilege Level

admin

Output Fields

Table 922 on page 22727 lists the output fields for the `show openflow filters` command. Output fields are listed in the approximate order in which they appear.

Table 922: show openflow filters Output Fields

Field Name	Field Description
Switch Name	User-configured identifier for the OpenFlow virtual switch associated with the interface to which the filter is bound.
Number of filters	Number of filters bound to OpenFlow interfaces on the virtual switch.
Default action	Default action executed for packets that do not match any existing flow entries. Values are `PACKET IN` or `DROP`.
Filter name	Filter identifier consisting of the concatenation of the interface name (including the logical unit number) and an internally assigned switch ID.
Filter index	Auto-generated string that identifies the filter.
Number of terms	Number of terms in the filter. Each term consists of match conditions and actions.
Number of priorities	Number of unique active flow priorities in the filter.
Term name	Filter term identifier, which consists of the filter name (interface name and switch ID), the flow priority, and a sequence number.
Priority ID	Flow entry priority. Higher priority terms are installed above lower priority terms.

Table 922: show openflow filters Output Fields *(Continued)*

Field Name	Field Description
Flow ID	Flow identifier associated with that flow entry.
Number of packets	Number of packets that have matched a filter term. A filter term is equivalent to a flow entry.
Number of bytes	Number of bytes that have matched a filter term. A filter term is equivalent to a flow entry.

Sample Output

show openflow filters

```
user@host> show openflow filters

Switch                  Filter      Number of  Number of  Number of
Name                    Index       terms      priorities packets
OFswitch1               96468992    0          0          0
                        96468993    0          0          0
                        96468994    0          0          0
                        96468995    0          0          0
                        96468996    1          1          7928017621
```

show openflow filters interface

```
user@host> show openflow filters interface ge-1/1/7.0
Switch Name: OFswitch1
Filter name: ge-1/1/7.0_0
Filter index: 96468996
Number of terms: 1            Number of priorities: 1

Term name: ge-1/1/7.0_0:32766^OF:1
Priority ID: 32766                        Flow ID:        16842752
Number of packets:7941332819              Number of bytes:476479969140
```

show openflow filters switch

```
user@host> show openflow filters switch OFswitch1
Switch Name: OFswitch1
Number of filters: 5            Default action: PACKET IN

Filter name: ge-1/1/0.0_0
Filter index: 96468992
Number of terms: 0            Number of priorities: 0

Filter name: ge-1/1/1.0_0
Filter index: 96468993
Number of terms: 0            Number of priorities: 0

Filter name: ge-1/1/2.0_0
Filter index: 96468994
Number of terms: 0            Number of priorities: 0

Filter name: ge-1/1/3.0_0
Filter index: 96468995
Number of terms: 0            Number of priorities: 0

Filter name: ge-1/1/7.0_0
Filter index: 96468996
Number of terms: 1            Number of priorities: 1
Priority  Flow                   Number of              Number of
ID        ID                     packets                bytes
32768     16842752               7941332819             476479969140
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

Understanding OpenFlow Flows and Filters on Devices Running Junos OS

show openflow flows

IN THIS SECTION

- Syntax | 22730
- Description | 22730
- Options | 22730
- Required Privilege Level | 22731
- Output Fields | 22731
- Sample Output | 22732
- Release Information | 22735

Syntax

```
show openflow flows
<brief | detail | summary>
<flow-id>
<switch switch-name>
```

Description

Display information about traffic flows that match filters on OpenFlow interfaces.

Options

none	(Same as `brief`) Display information for all flows.
brief | detail | summary	(Optional) Display the specified level of output.
flow-id	(Optional) Display information only for the specified flow.
switch *switch-name*	(Optional) Display information only for the flows on the specified OpenFlow virtual switch.

Required Privilege Level

admin

Output Fields

Table 923 on page 22731 lists the output fields for the `show openflow flows` command. Output fields are listed in the approximate order in which they appear.

Table 923: show openflow flows Output Fields

Field Name	Field Description
`Switch Name`	User-configured identifier for the OpenFlow virtual switch on which the flow resides.
`Number of flows`	Number of active flow entries associated with that OpenFlow virtual switch.
`Flow name`	Flow descriptor.
`Table ID`	Identifier for the flow table from which the flow originated.
`Flow ID`	Flow identifier associated with that flow entry.
`Number of packets`	Number of packets that have matched the flow entry.
`Priority`	Flow entry priority. Packets match higher priority entries before matching lower priority entries.
`Idle timeout`	Number of seconds after which the flow entry is removed from the flow table provided there are no matching packets.
`Hard timeout`	Number of seconds after which the flow entry is removed from the flow table regardless of the number of matching packets.

Table 923: show openflow flows Output Fields *(Continued)*

Field Name	Field Description
Cookie	An identifier specified by the OpenFlow controller when the flow is installed in the flow table. Cookies are used to filter flows for flow modification and delete operations. **NOTE**: This field appears only if the device running Junos OS supports OpenFlow version 1.3.1 or later.
Match	Configured match conditions against which the incoming packet is compared. **NOTE**: In the OpenFlow 1.3.1 sample output, the match fields that are populated for an IPv6 flow apply to MX Series routers only.
Action	Set of actions (for OpenFlow v1.0) or flow instructions (for OpenFlow v1.3.1) applied to a packet when it matches the flow entry.
Number of match	Number of match conditions against which the incoming packet is compared.
Number of action	Number of actions (for OpenFlow v1.0) or flow instructions (for OpenFlow v1.3.1) that are applied to a packet when it matches the flow entry.

Sample Output

show openflow flows switch brief

```
user@host> show openflow flows switch OFswitch1 brief

Switch                  Flow    Number of packets    Priority Number of Number of
Name                    ID                                    match    action
OFswitch1               16842752 8075372509          32768    1         1
```

show openflow flows switch detail

```
user@host> show openflow flows switch OFswitch1 detail
Flow name: flow-16842752
Table ID: 1     Flow ID: 16842752
```

```
Priority: 32768   Idle timeout(in sec):0        Hard timeout(in sec): 0
Cookie: 0
Match: Input port: 45549
       Ethernet src addr: wildcard
       Ethernet dst addr: wildcard
       Input vlan id: wildcard          Input VLAN priority: wildcard
       Ether type: wildcard
       IP ToS: wildcard                 IP protocol: wildcard
       IP src addr: wildcard            IP dst addr: wildcard
       Source port: wildcard            Destination port: wildcard
Action: Output port 41350,
```

show openflow flows switch detail (OpenFlow 1.3.1, IPv6 Flow)

```
user@host> show openflow flows switch OFswitch1 detail
Flow name: flow-67174400
Table ID: 1     Flow ID: 67174400
Priority: 8       Idle timeout(in sec):1800     Hard timeout(in sec): 1800
Cookie: 3333
Match: Input port: 2
       Ethernet src addr: aa:bb:cc:11:22:33
       Ethernet dst addr: aa:bb:cc:11:22:34
       Input vlan id: 50                Input VLAN priority: 3
       Ether type: 0x86dd
       IP ToS: 0x2                      IP protocol: 0x6
       IPv4 src addr: NA
       IPv4 dst addr: NA
       IPv6 src addr: 2001::/64
       IPv6 dst addr: 1001::1/128
       Source port: 8                   Destination port: 9
Action: Set_Field[VLAN_ID] 601, Output port 4, Set_Field[VLAN_ID] 701, Output port 6,
```

show openflow flows switch summary

```
user@host> show openflow flows switch OFswitch1 summary

Switch Name                                                    Number of Flows
OFswitch1                                                      1
```

show openflow flows brief (Specific Flow)

```
user@host> show openflow flows 16842752 brief

Switch                  Flow    Number of packets    Priority Number of Number of
Name                    ID                                    match    action
OFswitch1               16842752 8056139439          32768    1         1
```

show openflow flows detail (Specific Flow)

```
user@host> show openflow flows 16842752 detail
Flow name: flow-16842752
Table ID: 1     Flow ID: 16842752
Priority: 32768   Idle timeout(in sec):0        Hard timeout(in sec): 0
Cookie: 0
Match: Input port: 45549
       Ethernet src addr: wildcard
       Ethernet dst addr: wildcard
       Input vlan id: wildcard          Input VLAN priority: wildcard
       Ether type: wildcard
       IP ToS: wildcard                 IP protocol: wildcard
       IP src addr: wildcard            IP dst addr: wildcard
       Source port: wildcard            Destination port: wildcard
Action: Output port 41350,
```

show openflow flows detail (Specific Flow, OpenFlow 1.3.1, IPv6 Flow)

```
user@host> show openflow flows 67174400 detail
Flow name: flow-67174400
Table ID: 1     Flow ID: 67174400
Priority: 8       Idle timeout(in sec):1800     Hard timeout(in sec): 1800
Cookie: 3333
Match: Input port: 2
       Ethernet src addr: aa:bb:cc:11:22:33
       Ethernet dst addr: aa:bb:cc:11:22:34
       Input vlan id: 50                Input VLAN priority: 3
       Ether type: 0x86dd
       IP ToS: 0x2                      IP protocol: 0x6
       IPv4 src addr: NA
```

```
        IPv4 dst addr: NA
        IPv6 src addr: 2001::/64
        IPv6 dst addr: 1001::1/128
        Source port: 8                  Destination port: 9
Action: Set_Field[VLAN_ID] 601, Output port 4, Set_Field[VLAN_ID] 701, Output port 6,
```

show openflow flows detail (Specific Flow, OpenFlow 1.3.1, Destination MAC Address Rewrite Action)

```
user@host> show openflow flows 16842752 detail

Flow name: flow-16842752
Table ID: 1     Flow ID: 16842752
Priority: 1       Idle timeout(in sec):0        Hard timeout(in sec): 0
Cookie: 3333
Match: Input port: 1
       Ethernet src addr: wildcard
       Ethernet dst addr: wildcard
       Input vlan id: wildcard          Input VLAN priority: wildcard
       Ether type: wildcard
       IP ToS: wildcard                 IP protocol: wildcard
       IPv4 src addr: wildcard
       IPv4 dst addr: wildcard
       IPv6 src addr: wildcard
       IPv6 dst addr: wildcard
       Source port: wildcard            Destination port: wildcard
Action: Dst MAC ab:cd:ef:12:34:56, Output port 2,
```

show openflow flows summary (Specific Flow)

```
user@host> show openflow flows 16842752 summary
Flow name: flow-16842752
Number of packets: 8066495711
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

OpenFlow Support on Juniper Networks Devices

show openflow groups

IN THIS SECTION

- Syntax | 22736
- Description | 22736
- Options | 22736
- Required Privilege Level | 22737
- Output Fields | 22737
- Sample Output | 22738
- Release Information | 22740

Syntax

```
show openflow groups
<brief | details | summary>
<group-id>
```

Description

Display information about OpenFlow groups. Groups are supported only on Juniper Networks devices running OpenFlow v1.3.1 or later.

Options

none	Display information for all groups.
brief | detail | summary	(Optional) Display the specified level of output.

group-id (Optional) Display information about the specified group only.

Required Privilege Level

admin

Output Fields

Table 924 on page 22737 describes the output fields for the `show openflow groups` command. Table 924 on page 22737 lists the output fields in the approximate order in which they are displayed in the sample output.

Table 924: show openflow groups Output Fields

Field Name	Field Description
`Group ID`	Unique identifier assigned to a group by the OpenFlow controller.
`Type`	Group type, which can be either All or Indirect.
`Number of Buckets`	Number of buckets for a particular group. A group can have 0 to 32 buckets.
`Number of Flow Reference`	Number of flow entries that point to a particular group.
`Bucket`	Information about each bucket for a particular group.
`Actions`	Set of action(s) applied to a packet when it matches the flow entry.
`Switch Name`	User-configured identifier for the OpenFlow virtual switch on which the flow resides.
`Number of Groups`	Number of groups that currently exist in the OpenFlow virtual switch.
`Flow`	Identifier associated with a particular flow entry.

Sample Output

show openflow groups

```
user@host> show openflow groups

Group                   Type          Number of           Number of Flow
ID                                    Buckets             Reference
50                      All           2                   1
51                      All           2                   1
60                      Indirect      1                   0
```

show openflow groups brief

```
user@host> show openflow groups brief

Group                   Type          Number of           Number of Flow
ID                                    Buckets             Reference
50                      All           2                   1
51                      All           2                   1
60                      Indirect      1                   0
```

show openflow groups detail

```
user@host> show openflow groups detail

Group Id: 50                    Type: All
Bucket Bucket 1
      Actions: VLAN ID 2022, Output port 2,
Bucket Bucket 2
      Actions: VLAN ID 3022, Output port 4,

Group Id: 51                    Type: All
Bucket Bucket 3
      Actions: VLAN ID 2001, Output port 1,
Bucket Bucket 4
      Actions: VLAN ID 3001, Output port 3,
```

```
Group Id: 60                    Type: Indirect
Bucket Bucket 5
      Actions: VLAN ID 2060, Output port 3,
```

show openflow groups summary

```
user@host> show openflow groups summary

Switch Name                                                    Number of Groups
OF-ex92k                                                       3
```

show openflow groups (Specific Group)

```
user@host> show openflow groups 50

Group                   Type          Number of           Number of Flow
ID                                    Buckets             Reference
50                      All           2                   1
```

show openflow groups brief (Specific Group)

```
user@host> show openflow groups 50 brief

Group                   Type          Number of           Number of Flow
ID                                    Buckets             Reference
50                      All           2                   1
```

show openflow groups detail (Specific Group)

```
user@host> show openflow groups 50 detail

Group Id: 50                    Type: All
Bucket 1
      Actions: VLAN ID 2022, Output port 2,
Bucket 2
```

```
      Actions: VLAN ID 3022, Output port 4,
Flow 570710622208
```

show openflow groups summary (Specific Group)

```
user@host> show openflow groups 50 summary

Group                   Type          Number of           Number of Flow
ID                                    Buckets             Reference
50                      All           2                   1
```

Release Information

Command introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Understanding How the OpenFlow Group Action Works

OpenFlow Operational Mode Commands

show openflow statistics groups | 22750

show openflow interfaces

IN THIS SECTION

- Syntax | 22741
- Description | 22741
- Options | 22741
- Required Privilege Level | 22741
- Output Fields | 22741
- Sample Output | 22742
- Release Information | 22745

Syntax

```
show openflow interfaces
<interface-name>
<switch switch-name>
```

Description

Display physical characteristics and status information for interfaces participating in OpenFlow.

Options

none	Display information for all interfaces participating in OpenFlow.
interface-name	(Optional) Display information only for the specified interface. Specify the interface name including the logical unit number—for example, ge-1/1/0.0.
switch *switch-name*	(Optional) Display information only for those interfaces configured under the specified OpenFlow virtual switch.

Required Privilege Level

admin

Output Fields

Table 925 on page 22741 lists the output fields for the `show openflow interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 925: show openflow interfaces Output Fields

Field Name	Field Description
Switch name	User-configured identifier for the OpenFlow virtual switch to which the interface is bound.
Interface Name	Name of the logical interface.

Table 925: show openflow interfaces Output Fields *(Continued)*

Field Name	Field Description
Interface port number	Port identifier associated with the OpenFlow interface.
Interface Hardware Address	Media access control (MAC) address of the interface.
Interface speed	Speed and duplex mode of the interface.
Interface Auto-Negotiation	Autonegotiation status: `Enabled` or `Disabled`.
Interface media type	Media type of the interface. For example, copper or fiber.
Interface state	Current state of the interface.

Sample Output

show openflow interfaces

```
user@host> show openflow interfaces
Switch name: OFswitch1
Interface Name: ge-1/1/2.0
Interface port number: 41507
Interface Hardware Address: 00:00:5e:00:53:b4
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch1
Interface Name: ge-1/1/3.0
Interface port number: 44383
Interface Hardware Address: 00:00:5e:00:53:b5
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
```

```
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch1
Interface Name: ge-1/1/1.0
Interface port number: 41350
Interface Hardware Address: 00:00:5e:00:53:b7
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch1
Interface Name: ge-1/1/7.0
Interface port number: 45549
Interface Hardware Address: 00:00:5e:00:53:b6
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch1
Interface Name: ge-1/1/0.0
Interface port number: 44538
Interface Hardware Address: 00:00:5e:00:53:b2
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up
```

show openflow interfaces (Specific Interface)

```
user@host> show openflow interfaces ge-1/1/0.0
Switch name: OFswitch1
Interface Name: ge-1/1/0.0
Interface port number: 44538
Interface Hardware Address: 00:00:5e:00:53:b2
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
```

```
Interface media type: Fiber
Interface state: Up
```

show openflow interfaces switch

```
user@host> show openflow interfaces switch OFswitch1
Switch name: OFswitch1
Interface Name: ge-1/1/2.0
Interface port number: 41507
Interface Hardware Address: 00:00:5e:00:53:b4
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch1
Interface Name: ge-1/1/3.0
Interface port number: 44383
Interface Hardware Address: 00:00:5e:00:53:b5
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch1
Interface Name: ge-1/1/1.0
Interface port number: 41350
Interface Hardware Address: 00:00:5e:00:53:b7
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch1
Interface Name: ge-1/1/7.0
Interface port number: 45549
Interface Hardware Address: 00:00:5e:00:53:b6
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up
```

```
Switch name: OFswitch1
Interface Name: ge-1/1/0.0
Interface port number: 44538
Interface Hardware Address: 00:00:5e:00:53:b2
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

show openflow filters | 22726

show openflow flows | 22730

show openflow statistics interfaces | 22753

show openflow statistics flows

IN THIS SECTION

- Syntax | 22746
- Description | 22746
- Options | 22746
- Required Privilege Level | 22746
- Output Fields | 22746
- Sample Output | 22747
- Release Information | 22749

Syntax

```
show openflow statistics flows
<flow-id>
<switch switch-name>
```

Description

Display statistics for OpenFlow flows.

Options

none	Display flow statistics for all flows for all OpenFlow virtual switches.
flow-id	(Optional) Display flow statistics only for the specified flow.
switch *switch-name*	(Optional) Display flow statistics only for the specified OpenFlow virtual switch.

Required Privilege Level

admin

Output Fields

Table 926 on page 22746 lists the output fields for the `show openflow statistics flows` command. Output fields are listed in the approximate order in which they appear.

Table 926: show openflow statistics flows Output Fields

Field Name	Field Description
`Switch Name`	User-configured identifier for the OpenFlow virtual switch on which the flow resides.
`Table ID`	Identifier for the flow table from which the flow originated.
`Flow ID`	OpenFlow flow entry identifier.

Table 926: show openflow statistics flows Output Fields *(Continued)*

Field Name	Field Description
`Duration(in sec)`	Number of seconds the flow has been active.
`Duration(in nsec)`	Number of nanoseconds the flow has been active beyond the value of `Duration(in sec)`.
`Priority`	Flow entry priority. Packets match higher priority entries before matching lower priority entries.
`Idle timeout`	Number of seconds after which the flow entry is removed from the flow table provided there are no matching packets.
`Hard timeout`	Number of seconds after which the flow entry is removed from the flow table regardless of the number of matching packets.
`Number of packets`	Number of packets that have matched the flow entry.
`Number of bytes`	Number of bytes that have matched the flow entry.
`Match`	Fields against which the incoming packet is compared.
`Action`	Set of actions applied to a packet when it matches the flow entry.

Sample Output

show openflow statistics flows

```
user@host> show openflow statistics flows
Switch Name: OFswitch1
Table ID: 1     Flow ID: 16842752
Duration(in sec): 58772          Duration(in nsec): 215702000
Priority: 32768   Idle timeout(in sec):0        Hard timeout(in sec): 0
Number of packets: 8745275026
Number of bytes:   524716501560
```

```
Match: IN_PORT,
Action: OUTPUT,
```

show openflow statistics flows (OpenFlow 1.3.1)

```
user@host> show openflow statistics flows
Switch Name: 100
Table ID: 1     Flow ID: 67174400
Duration(in sec): 51539          Duration(in nsec): 3961483296
Priority: 1000    Idle timeout(in sec):0        Hard timeout(in sec): 0
Number of packets: 179965839
Number of bytes:   14757198798
Match: IN_PORT, DL_VLAN, DL_TYPE, IPv6_SRC, IPv6_DST,
Action: SET_FIELD[VLAN_ID], OUTPUT,
```

show openflow statistics flows (Specific Flow)

```
user@host> show openflow statistics flows 16842752
Switch Name: OFswitch1
Table ID: 1     Flow ID: 16842752
Duration(in sec): 58803          Duration(in nsec): 4127548296
Priority: 32768   Idle timeout(in sec):0        Hard timeout(in sec): 0
Number of packets: 8749713419
Number of bytes:   524982805140
Match: IN_PORT,
Action: OUTPUT,
```

show openflow statistics flows (Specific Flow, OpenFlow 1.3.1)

```
user@host> show openflow statistics flows 67174400
Switch Name: 100
Table ID: 1     Flow ID: 67174400
Duration(in sec): 51539          Duration(in nsec): 3961483296
Priority: 1000    Idle timeout(in sec):0        Hard timeout(in sec): 0
Number of packets: 179965839
Number of bytes:   14757198798
```

```
Match: IN_PORT, DL_VLAN, DL_TYPE, IPv6_SRC, IPv6_DST,
Action: SET_FIELD[VLAN_ID], OUTPUT,
```

show openflow statistics flows switch

```
user@host> show openflow statistics flows switch OFswitch1
Switch Name: OFswitch1
Table ID: 1     Flow ID: 16842752
Duration(in sec): 58829          Duration(in nsec): 4124448296
Priority: 32768   Idle timeout(in sec):0        Hard timeout(in sec): 0
Number of packets: 8752672358
Number of bytes:   525160341480
Match: IN_PORT,
Action: OUTPUT,
```

show openflow statistics flows switch (OpenFlow 1.3.1)

```
user@host> show openflow statistics flows switch 100
Switch Name: 100
Table ID: 1     Flow ID: 67174400
Duration(in sec): 51539          Duration(in nsec): 3961483296
Priority: 1000    Idle timeout(in sec):0        Hard timeout(in sec): 0
Number of packets: 179965839
Number of bytes:   14757198798
Match: IN_PORT, DL_VLAN, DL_TYPE, IPv6_SRC, IPv6_DST,
Action: SET_FIELD[VLAN_ID], OUTPUT,
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

show openflow flows | 22730

show openflow statistics interfaces | 22753

show openflow statistics packet | 22758

show openflow statistics tables | 22766

show openflow statistics groups

IN THIS SECTION

- Syntax | 22750
- Description | 22750
- Options | 22750
- Required Privilege Level | 22750
- Output Fields | 22751
- Sample Output | 22752
- Release Information | 22753

Syntax

```
show openflow statistics groups
<group-id>
```

Description

Display statistics for OpenFlow groups. Groups are supported only on Juniper Networks devices running OpenFlow v1.3.1 or later.

Options

none Display statistics for all groups defined on all OpenFlow virtual switches.

group-id (Optional) Display statistics only for the specified group.

Required Privilege Level

admin

Output Fields

Table 927 on page 22751 lists the output fields for the `show openflow statistics groups` command. Output fields are listed in the approximate order in which they appear.

Table 927: show openflow statistics groups Output Fields

Field	Description
`Switch Name`	User-configured identifier for the OpenFlow virtual switch in which the groups resides.
`Group ID`	Unique identifier assigned to the group by the OpenFlow controller.
`Ref Count`	Number of flow entries that reference the group.
`Number of packets` (group)	Number of packets handled by the group.
`Number of bytes` (group)	Number of bytes handled by the group.
`Duration(in sec)`	Number of seconds the group has been active.
`Duration(in nsec)`	Number of nanoseconds the group has been active beyond the `Duration(in sec)`
`Bucket` *`number`*	Statistics for a particular bucket.
`Number of packets` (bucket)	Number of packets handled by a bucket in the group.
`Number of bytes` (bucket)	Number of bytes handled by a bucket in the group.

For a group with the group type of all, the values specified in the `Number of packets` (group) and `Number of bytes` (group) fields are usually the same as those specified in the `Number of packets` (bucket) and `Number of bytes` (bucket) fields because all buckets in the group are executed.

Sample Output

show openflow statistics group

```
user@host> show openflow statistics groups

Switch Name: OF-ex92k
Group ID: 50                    Ref Count: 1
Number of packets: 62161
Number of bytes:   7956608
Duration(in sec): 22687          Duration(in nsec): 4255381296
    Bucket 0
      Number of packets: 62161
      Number of bytes:   7956608
    Bucket 1
      Number of packets: 62161
      Number of bytes:   7956608

Switch Name: OF-ex92k
Group ID: 51                    Ref Count: 1
Number of packets: 0
Number of bytes:   0
Duration(in sec): 22673          Duration(in nsec): 8549000
    Bucket 0
      Number of packets: 0
      Number of bytes:   0
    Bucket 1
      Number of packets: 0
      Number of bytes:   0
...
```

show openflow statistics group (Specific Group)

```
user@host> show openflow statistics groups 50

Switch Name: OF-ex92k
Group ID: 50                    Ref Count: 1
Number of packets: 64886
Number of bytes:   8305408
Duration(in sec): 22789          Duration(in nsec): 586200000
```

```
    Bucket 0
      Number of packets: 64886
      Number of bytes:   8305408
    Bucket 1
      Number of packets: 64886
      Number of bytes:   8305408
```

Release Information

Command introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Understanding How the OpenFlow Group Action Works

OpenFlow Operational Mode Commands

show openflow groups | 22736

show openflow statistics interfaces

IN THIS SECTION

- Syntax | 22754
- Description | 22754
- Options | 22754
- Required Privilege Level | 22754
- Output Fields | 22754
- Sample Output | 22756
- Release Information | 22757

Syntax

```
show openflow statistics interfaces
<switch switch-name>
```

Description

Display statistics for interfaces participating in OpenFlow.

Options

none	Display statistics for all interfaces participating in OpenFlow for all configured OpenFlow virtual switches.
switch *switch-name*	(Optional) Display statistics only for those interfaces on the specified OpenFlow virtual switch.

Required Privilege Level

admin

Output Fields

Table 928 on page 22754 lists the output fields for the `show openflow statistics interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 928: show openflow statistics interfaces Output Fields

Field Name	Field Description
Switch Name	User-configured identifier for the OpenFlow virtual switch to which the interface is bound.
Interface Name	Name of the logical interface.
Port Number	Port identifier associated with the OpenFlow interface.

Table 928: show openflow statistics interfaces Output Fields *(Continued)*

Field Name	Field Description
Num of rx pkts	Number of packets received on the OpenFlow interface.
Num of tx pkts	Number of packets transmitted on the OpenFlow interface.
Num of rx bytes	Number of bytes received on the OpenFlow interface.
Num of tx bytes	Number of bytes transmitted on the OpenFlow interface.
Num of rx error	Number of receive errors.
Num of tx error	Number of transmit errors.
Number of packets dropped by RX	Number of packets dropped by the ingress interface.
Number of packets dropped by TX	Number of packets dropped by the egress interface.
Number of rx frame error	Number of packets with frame alignment errors.
Number of rx overrun error	Number of packets with RX overrun.
Number of CRC error	Number of CRC errors.
Number of collisions	Number of Ethernet collisions.

Sample Output

show openflow statistics interfaces

```
user@host> show openflow statistics interfaces
Switch Name: OFswitch1
Interface Name: ge-1/1/2.0       Port Number: 41507
Num of rx pkts: 0                          Num of tx pkts: 1372301
Num of rx bytes: 0                         Num of tx bytes: 88665532
Num of rx error: 0                         Num of tx error:0
Number of packets dropped by RX: 0
Number of packets dropped by TX: 0
Number of rx frame error:        0
Number of rx overrun error:      0
Number of CRC error:             0
Number of collisions:            0

Switch Name: OFswitch1
Interface Name: ge-1/1/3.0       Port Number: 44383
Num of rx pkts: 0                          Num of tx pkts: 1372285
Num of rx bytes: 0                         Num of tx bytes: 88664476
Num of rx error: 0                         Num of tx error:0
Number of packets dropped by RX: 0
Number of packets dropped by TX: 0
Number of rx frame error:        0
Number of rx overrun error:      0
Number of CRC error:             0
Number of collisions:            0

Switch Name: OFswitch1
Interface Name: ge-1/1/1.0       Port Number: 41350
Num of rx pkts: 0                          Num of tx pkts: 8776241344
Num of rx bytes: 0                         Num of tx bytes: 526580807026
Num of rx error: 0                         Num of tx error:0
Number of packets dropped by RX: 0
Number of packets dropped by TX: 0
Number of rx frame error:        0
Number of rx overrun error:      0
Number of CRC error:             0
Number of collisions:            0

Switch Name: OFswitch1
```

```
Interface Name: ge-1/1/7.0       Port Number: 45549
Num of rx pkts: 8840952127                Num of tx pkts: 1047701
Num of rx bytes: 530457127620             Num of tx bytes: 69187816
Num of rx error: 0                        Num of tx error:0
Number of packets dropped by RX: 0
Number of packets dropped by TX: 0
Number of rx frame error:        0
Number of rx overrun error:      0
Number of CRC error:             0
Number of collisions:            0

Switch Name: OFswitch1
Interface Name: ge-1/1/0.0       Port Number: 44538
Num of rx pkts: 0                         Num of tx pkts: 1372031
Num of rx bytes: 0                        Num of tx bytes: 88647712
Num of rx error: 0                        Num of tx error:0
Number of packets dropped by RX: 0
Number of packets dropped by TX: 0
Number of rx frame error:        0
Number of rx overrun error:      0
Number of CRC error:             0
Number of collisions:            0
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

show openflow interfaces | 22740

show openflow statistics flows | 22745

show openflow statistics tables | 22766

show openflow statistics packet

IN THIS SECTION

- Syntax | 22758
- Description | 22758
- Options | 22758
- Required Privilege Level | 22758
- Output Fields | 22759
- Sample Output | 22759
- Release Information | 22760

Syntax

```
show openflow statistics packet (in | out)
<switch switch-name>
```

Description

Display statistics for packet-in and packet-out (send-packet) actions.

Options

none	Display statistics for all OpenFlow virtual switches.
switch *switch-name*	(Optional) Display statistics only for the specified OpenFlow virtual switch.

Required Privilege Level

admin

Output Fields

Table 929 on page 22759 lists the output fields for the `show openflow statistics packet` command. Output fields are listed in the approximate order in which they appear.

Table 929: show openflow statistics packet Output Fields

Field Name	Field Description
Switch Name	User-configured identifier for the OpenFlow virtual switch.
Rx packets	Number of packets received by the OpenFlow virtual switch that have been sent to the OpenFlow controller. The switch includes the packet in the data portion of an OFPT_PACKET_IN message.
Tx packets	Number of packets sent by the OpenFlow controller to an egress interface. The controller includes the packet in the data portion of an OFPT_PACKET_OUT message.
Drop packets	Number of dropped packets.

Sample Output

show openflow statistics packet in

```
user@host> show openflow statistics packet in
Openflow packet-in statistics information:
Switch Name                                          Rx packets  Drop packets
OFswitch1                                            1044137     0
```

show openflow statistics packet out

```
user@host> show openflow statistics packet out
Openflow packet-out statistics information:
Switch Name                                          Tx packets  Drop packets
OFswitch1                                            5260759     0
```

show openflow statistics packet out switch

```
user@host> show openflow statistics packet out switch OFswitch1
Openflow packet-out statistics information:
Switch Name                                          Tx packets  Drop packets
OFswitch1                                            5260759     0
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

show openflow statistics queue

IN THIS SECTION

- Syntax | 22760
- Description | 22761
- Options | 22761
- Required Privilege Level | 22761
- Output Fields | 22761
- Sample Output | 22762
- Release Information | 22763

Syntax

```
show openflow statistics queue
<interface interface-name>
```

Description

Display statistics for hardware queues for interfaces participating in OpenFlow.

Options

none Display queue statistics for all interfaces participating in OpenFlow.

interface *interface-name* (Optional) Display queue statistics only for the specified interface. Specify the interface name including the logical unit number—for example, ge-1/1/0.0

Required Privilege Level

admin

Output Fields

Table 930 on page 22761 lists the output fields for the `show openflow statistics queue` command. Output fields are listed in the approximate order in which they appear.

Table 930: show openflow statistics queue Output Fields

Field Name	Field Description
Switch Name	User-configured identifier for the OpenFlow virtual switch.
Port No	Port identifier associated with the OpenFlow interface.
Queue Id	Priority queue identifier.
TX bytes	Number of bytes transmitted through the queue.
TX packets	Number of packets transmitted through the queue.
Tx errors	Number of packets dropped by the queue due to overrun.

Sample Output

show openflow statistics queue

```
user@host> show openflow statistics queue
Openflow queue statistics information:
Switch Name            Port No Queue Id   TX bytes     TX packets   Tx errors
OFswitch1              41507  0           115327076    1372459      0
OFswitch1              41507  1           0            0            0
OFswitch1              41507  2           0            0            0
OFswitch1              41507  3           0            0            0
OFswitch1              41507  4           0            0            0
OFswitch1              41507  5           0            0            0
OFswitch1              41507  6           0            0            0
OFswitch1              41507  7           0            0            0
OFswitch1              44383  0           115325732    1372443      0
OFswitch1              44383  1           0            0            0
OFswitch1              44383  2           0            0            0
OFswitch1              44383  3           0            0            0
OFswitch1              44383  4           0            0            0
OFswitch1              44383  5           0            0            0
OFswitch1              44383  6           0            0            0
OFswitch1              44383  7           0            0            0
OFswitch1              41350  0           752072717540 8953246155   0
OFswitch1              41350  1           0            0            0
OFswitch1              41350  2           0            0            0
OFswitch1              41350  3           0            0            0
OFswitch1              41350  4           0            0            0
OFswitch1              41350  5           0            0            0
OFswitch1              41350  6           0            0            0
OFswitch1              41350  7           0            0            0
OFswitch1              45549  0           88060496     1047859      0
OFswitch1              45549  1           0            0            0
OFswitch1              45549  2           0            0            0
OFswitch1              45549  3           0            0            0
OFswitch1              45549  4           0            0            0
OFswitch1              45549  5           0            0            0
OFswitch1              45549  6           0            0            0
OFswitch1              45549  7           0            0            0
OFswitch1              44538  0           115304396    1372189      0
OFswitch1              44538  1           0            0            0
OFswitch1              44538  2           0            0            0
```

```
OFswitch1              44538  3           0            0            0
OFswitch1              44538  4           0            0            0
OFswitch1              44538  5           0            0            0
OFswitch1              44538  6           0            0            0
OFswitch1              44538  7           0            0            0
```

show openflow statistics queue interface

```
user@host> show openflow statistics queue interface ge-1/1/2.0
Openflow queue statistics information:
Switch Name            Port No Queue Id   TX bytes     TX packets   Tx errors
OFswitch1              41507  0           115327076    1372459      0
OFswitch1              41507  1           0            0            0
OFswitch1              41507  2           0            0            0
OFswitch1              41507  3           0            0            0
OFswitch1              41507  4           0            0            0
OFswitch1              41507  5           0            0            0
OFswitch1              41507  6           0            0            0
OFswitch1              41507  7           0            0            0
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

show openflow statistics flows | 22745

show openflow statistics tables | 22766

show openflow statistics summary

IN THIS SECTION

- Syntax | 22764
- Description | 22764
- Options | 22764
- Required Privilege Level | 22764
- Output Fields | 22764
- Sample Output | 22766
- Release Information | 22766

Syntax

```
show openflow statistics summary
```

Description

Display summary statistics for all installed OpenFlow flow entries for all OpenFlow virtual switches.

Options

This command has no options.

Required Privilege Level

admin

Output Fields

Table 931 on page 22765 lists the output fields for the `show openflow statistics summary` command. Output fields are listed in the approximate order in which they appear.

Table 931: show openflow statistics summary Output Fields

Field Name	Field Description
Switch Name	User-configured identifier for the OpenFlow virtual switch.
Port Number	Port identifier associated with the OpenFlow interface.
Number of RX packets	Number of packets received on the OpenFlow interface.
Number of TX packets	Number of packets transmitted on the OpenFlow interface.
Num of packets dropped by RX	Number of packets dropped by the ingress interface.
Flow ID	Flow identifier associated with that flow entry.
Number of packets	Number of packets that have matched the flow entry.
Duration (in sec)	Number of seconds the flow has been active.
Priority	Flow entry priority. Packets match higher priority entries before matching lower priority entries.
Idle Timeout	Number of seconds after which the flow entry is removed from the flow table provided there are no matching packets.
Hard Timeout	Number of seconds after which the flow entry is removed from the flow table regardless of the number of matching packets.

Sample Output

show openflow statistics summary

```
user@host> show openflow statistics summary

Switch    Port   Number of RX packets    Number of TX packets    Num of packets
Name      Number                                                 dropped by RX
OFswitch1    41507 0                     1372609                 0
OFswitch1    44383 0                     1372593                 0
OFswitch1    41350 0                     9119477900              0
OFswitch1    45549 9184188377            1048009                 0
OFswitch1    44538 0                     1372339                 0
Switch           Flow   Number of packets     Duration  Priority Idle    Hard
Name             ID                           (in sec)           Timeout Timeout
OFswitch1        16842752 9117212928          61278       32768  0       0
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

show openflow statistics flows | 22745

show openflow statistics tables | 22766

show openflow summary | 22771

show openflow statistics tables

IN THIS SECTION

- Syntax | 22767
- Description | 22767

- Options | 22767
- Required Privilege Level | 22767
- Output Fields | 22767
- Sample Output | 22769
- Release Information | 22770

Syntax

```
show openflow statistics tables
<switch switch-name>
```

Description

Display statistics for OpenFlow flow tables.

Options

none	Display statistics for flow tables on all OpenFlow virtual switches.
switch *switch-name*	(Optional) Display statistics only for flow tables on the specified OpenFlow virtual switch.

Required Privilege Level

admin

Output Fields

Table 932 on page 22768 lists the output fields for the `show openflow statistics tables` command. Output fields are listed in the approximate order in which they appear.

Table 932: show openflow statistics tables Output Fields

Field Name	Field Description
`Table Name`	String identifier for the OpenFlow flow table.
`Table id`	Numeric identifier for the OpenFlow flow table.
`Supported wildcards`	Wildcards supported by the flow table.
`Max number of entries`	Maximum number of entries supported in the flow table.
`Number of active entries`	Number of active entries in the flow table.
`Number of idle timeout entries`	Number of entries in the flow table that have been removed because the idle timeout expired and no packets matched those entries.
`Number of hard timeout entries`	Number of entries in the flow table that have been removed because the hard timeout expired.
`Number of flow delete entries`	Number of entries in the flow table that have been removed in response to controller requests.
`Number of flow add entries`	Number of entries in the flow table that have been added in response to controller requests.
`Number of flow modify entries`	Number of entries in the flow table that have been modified in response to controller requests.
`Number of total delete entries`	Number of entries in the flow table that have been removed for any reason.

Sample Output

show openflow statistics tables

```
user@host> show openflow statistics tables
Table name: Default flow table                  Table id:1
Supported wildcards: IN_PORT, DL_VLAN, DL_SRC, DL_DST, DL_TYPE, NW_PROTO, TP_SRC, TP_DST,
NW_SRC, NW_DST, DL_VLAN_PCP, NW_TOS,
Max number of entries: 65535              Number of active entries: 1
Number of idle timeout entries: 0
Number of hard timeout entries: 0
Number of flow delete entries:  0
Number of flow add entries:     1
Number of flow modify entries:  0
Number of total delete entries: 0
```

show openflow statistics tables (OpenFlow 1.3.1)

```
user@host> show openflow statistics tables
Table name: Default flow table                  Table id:1
Supported wildcards: IN_PORT, DL_VLAN, DL_SRC, DL_DST, DL_TYPE, NW_PROTO, TP_SRC, TP_DST,
IPv4_SRC, IPv4_DST, IPv6_SRC, IPv6_DST, DL_VLAN_PCP, NW_TOS,
Number of active entries: 7
Number of entries used: 0
Number of idle timeout entries: 0
Number of hard timeout entries: 0
Number of flow delete entries:  0
Number of flow add entries:     7
Number of flow modify entries:  0
Number of total delete entries: 0
```

show openflow statistics tables switch

```
user@host> show openflow statistics tables switch OFswitch1
Table name: Default flow table                  Table id:1
Supported wildcards: IN_PORT, DL_VLAN, DL_SRC, DL_DST, DL_TYPE, NW_PROTO, TP_SRC, TP_DST,
NW_SRC, NW_DST, DL_VLAN_PCP, NW_TOS,
Max number of entries: 65535              Number of active entries: 1
```

```
Number of idle timeout entries: 0
Number of hard timeout entries: 0
Number of flow delete entries:  0
Number of flow add entries:     1
Number of flow modify entries:  0
Number of total delete entries: 0
```

show openflow statistics tables switch (OpenFlow 1.3.1)

```
user@host> show openflow statistics tables switch 100
Table name: Default flow table                  Table id:1
Supported wildcards: IN_PORT, DL_VLAN, DL_SRC, DL_DST, DL_TYPE, NW_PROTO, TP_SRC, TP_DST,
IPv4_SRC, IPv4_DST, IPv6_SRC, IPv6_DST, DL_VLAN_PCP, NW_TOS,
Number of active entries: 7
Number of entries used: 0
Number of idle timeout entries: 0
Number of hard timeout entries: 0
Number of flow delete entries:  0
Number of flow add entries:     7
Number of flow modify entries:  0
Number of total delete entries: 0
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

show openflow statistics flows | 22745

show openflow statistics interfaces | 22753

show openflow statistics summary | 22764

show openflow summary

IN THIS SECTION

- Syntax | 22771
- Description | 22771
- Options | 22771
- Required Privilege Level | 22771
- Output Fields | 22771
- Sample Output | 22772
- Release Information | 22772

Syntax

```
show openflow summary
```

Description

Display summary information for OpenFlow including the number of configured virtual switches, controllers, interfaces, and flows.

Options

This command has no options.

Required Privilege Level

admin

Output Fields

Table 933 on page 22772 lists the output fields for the `show openflow summary` command. Output fields are listed in the approximate order in which they appear.

Table 933: show openflow summary Output Fields

Field Name	Field Description
Number of switches	Total number of configured OpenFlow virtual switches.
Number of controllers	Total number of configured OpenFlow controllers.
Number of interfaces	Number of logical interfaces participating in OpenFlow.
Number of active flow entries	Number of active entries in the flow table.

Sample Output

show openflow summary

```
user@host> show openflow summary
Number of switches:            1
Number of controllers:         1
Number of interfaces:          5
Number of active flow entries: 1
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

show openflow statistics summary | 22764

show openflow switch | 22773

show openflow switch

IN THIS SECTION

- Syntax | 22773
- Description | 22773
- Options | 22773
- Required Privilege Level | 22773
- Output Fields | 22774
- Sample Output | 22775
- Release Information | 22776

Syntax

```
show openflow switch
<switch-name>
```

Description

Display OpenFlow message statistics for OpenFlow virtual switches.

Options

none	Display information for all OpenFlow virtual switches.
switch *switch-name*	(Optional) Display information only for the specified OpenFlow virtual switch.

Required Privilege Level

admin

Output Fields

Table 934 on page 22774 lists the output fields for the `show openflow switch` command. Output fields are listed in the approximate order in which they appear.

Table 934: show openflow switch Output Fields

Field Name	Field Description
Switch Name	User-configured identifier for the OpenFlow virtual switch.
Switch ID	Device identifier for the OpenFlow virtual switch.
Switch DPID	Data path ID uniquely identifying the OpenFlow instance. This value is a concatenation of the switch ID for the virtual switch and the management port MAC address.
Flow mod received	Number of Modify Flow Entry messages (OFPT_FLOW_MOD) received from the controller.
Vendor received	Number of messages with vendor-specific extensions.
Packets sent	Number of packets sent to the controller.
Packets received	Number of packets received from the controller.
Echo req sent	Number of Echo Request messages (OFPT_ECHO_REQUEST) sent to the controller.
Echo req received	Number of Echo Request messages (OFPT_ECHO_REQUEST) received from the controller.
Echo reply sent	Number of Echo Reply messages (OFPT_ECHO_REPLY) sent to the controller.
Echo reply received	Number of Echo Reply messages (OFPT_ECHO_REPLY) received from the controller.
Port Status sent	Number of Port Status messages (OFPT_PORT_STATUS) sent to the controller.

Table 934: show openflow switch Output Fields *(Continued)*

Field Name	Field Description
Port mod received	Number of Port Modification messages (OFPT_PORT_MOD) received from the controller.
Barrier request	Number of Barrier Request messages (OFPT_BARRIER_REQUEST) received from the controller.
Barrier reply	Number of Barrier Reply messages (OFPT_BARRIER_REPLY) sent to the controller.
Error msg sent	Number of error messages (OFPT_ERROR) sent to the controller.
Error msg received	Number of error messages (OFPT_ERROR) received from the controller.

Sample Output

show openflow switch

```
user@host> show openflow switch
Switch Name:        OFswitch1
Switch ID:          0                   Switch DPID:    00:00:00:00:5e:00:53:d0
Flow mod received:  4                   Vendor received:      0
Packets sent:       1048258             Packets received:     1089664
Echo req sent:      4115                Echo req received:    0
Echo reply sent:    0                   Echo reply received:  4115
Port Status sent:   1                   Port mod received:    0
Barrier request:    0                   Barrier reply:        0
Error msg sent:     1                   Error msg received:   0
```

show openflow switch (Specific OpenFlow Virtual Switch)

```
user@host> show openflow switch OFswitch1
Switch Name:        OFswitch1
Switch ID:          0                   Switch DPID:    00:00:00:00:5e:00:53:d0
Flow mod received:  4                   Vendor received:      0
```

```
Packets sent:       1048259            Packets received:     1089675
Echo req sent:      4116               Echo req received:    0
Echo reply sent:    0                  Echo reply received:  4116
Port Status sent:   1                  Port mod received:    0
Barrier request:    0                  Barrier reply:        0
Error msg sent:     1                  Error msg received:   0
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

OpenFlow Operational Mode Commands

show openflow statistics tables | 22766

show openflow summary | 22771

show ospf context-identifier

IN THIS SECTION

- Syntax | 22777
- Syntax (EX Series Switches and QFX Series) | 22777
- Description | 22777
- Options | 22777
- Required Privilege Level | 22778
- Output Fields | 22778
- Sample Output | 22779
- Release Information | 22779

Syntax

```
show ospf context-identifier
<brief | detail>
<area area-id>
<context-id>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show ospf context-identifier
<brief | detail>
<area area-id>
<context-id>
<instance instance-name>
```

Description

Display the context identifier information processed and advertised by Open Shortest Path First (OSPF) for egress protection.

Options

none	Display information about all context identifiers.
brief | detail	(Optional) Display the specified level of output.
area *area-id*	(Optional) Display information about the context identifier for the specified area.
context-id	(Optional) Display information about the specified context identifier.
instance *instance-name*	(Optional) Display information about the context identifier for the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 935 on page 22778 lists the output fields for the `show ospf context-identifier` command. Output fields are listed in the approximate order in which they appear.

Table 935: show ospf context-identifier Output Fields

Field Name	Field Description	Level of Output
Context	IPv4 address that defines a protection pair. The context is manually configured on both primary and protector provider edge (PE) devices.	All levels
Status	State of the path: active or inactive.	All levels
Metric	Advertised OSPF metric.	All levels
Area	OSPF area number.	All levels
Other Advertisements	Other advertisements received by the OSPF node: • Advertising router– Address of the device that sent the advertisement. • Type–Type of OSPF path: inter-area and stub. • Metric–Advertised OSPF metric. • None–No additional advertisements were received by the OSPF node.	`detail`

Sample Output

show ospf context-identifier

```
user@host> show ospf context-identifier
Context-id: 2.2.4.3
  Status: active, Metric: 65534, PE role: protector, Area: 0.0.0.0
```

show ospf context-identifier detail

```
user@host> show ospf context-identifier detail
Context-id: 88.24.13.1
                Status: inactive, Metric: 0, PE role: protector, Area: 0.0.0.13
                Other Advertisements:
                  Advertising router: 8.8.8.103
                    Type: stub link
                    Metric: 65534
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- egress-protection (Layer 2 circuit) | 6218
- Junos OS VPNs Library for Routing Devices
- egress-protection (MPLS) | 6220
- Junos OS VPNs Library for Routing Devices

show ospf database

IN THIS SECTION

- Syntax | 22780
- Syntax (EX Series Switches and QFX Series) | 22781
- Description | 22781
- Options | 22781
- Required Privilege Level | 22782
- Output Fields | 22782
- Sample Output | 22785
- Release Information | 22791

Syntax

```
show ospf database
<brief | detail | extensive | summary>
<advertising-router (address | self)>
<area area-id>
<asbrsummary>
<external>
<instance instance-name>
<link-local>
<logical-system (all | logical-system-name)>
<lsa-id lsa-id>
<netsummary>
<network>
<nssa>
<opaque-area  (extended-link link-addr link-address)>
<opaque-area   (extended-prefix prefix prefix length)><router>
```

Syntax (EX Series Switches and QFX Series)

```
show ospf database
<brief | detail | extensive | summary>
<advertising-router (address | self)>
<area area-id>
<asbrsummary>
<external>
<instance instance-name>
<link-local>
<lsa-id lsa-id>
<netsummary>
<network>
<nssa>
<opaque-area>
<router>
```

Description

Display the entries in the OSPF version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets.

Options

none	Display standard information about entries in the OSPFv2 link-state database for all routing instances.
brief | detail | extensive | summary	(Optional) Display the specified level of output.
advertising-router (*address* | self)	(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.
area *area-id*	(Optional) Display the LSAs in a particular area.
asbrsummary	(Optional) Display summary AS boundary router LSA entries.
external	(Optional) Display external LSAs.
instance *instance-name*	(Optional) Display all OSPF database information under the named routing instance.

link-local	(Optional) Display information about link-local LSAs.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
lsa-id *lsa-id*	(Optional) Display the LSA with the specified LSA identifier.
netsummary	(Optional) Display summary network LSAs.
network	(Optional) Display information about network LSAs.
nssa	(Optional) Display information about not-so-stubby area (NSSA) LSAs.
opaque-area	(Optional) Display opaque area-scope LSAs.
extended-link link-addr *link-address*	(Optional) Displays the specific extended-link LSAs based on the link-address.
extended-prefix prefix *prefix/len*	(Optional) Displays the specific extended-prefix LSAs based on the prefix.
router	(Optional) Display information about router LSAs.

Required Privilege Level

view

Output Fields

Table 936 on page 22782 describes the output fields for the `show ospf database` command. Output fields are listed in the approximate order in which they appear.

Table 936: show ospf database Output Fields

Field Name	Field Description	Level of Output
area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Type	Type of link advertisement: ASBRSum, Extern, Network, NSSA, OpaqArea, Router, or Summary.	All levels

Table 936: show ospf database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
ID	LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device.	All levels
Adv Rtr	Address of the routing device that sent the advertisement.	All levels
Seq	Link sequence number of the advertisement.	All levels
Age	Time elapsed since the LSA was originated, in seconds.	All levels
Opt	Optional OSPF capabilities associated with the LSA.	All levels
Cksum	Checksum value of the LSA.	All levels
Len	Length of the advertisement, in bytes.	All levels
Router	Router link-state advertisement information: • bits—Flags describing the routing device that generated the LSP. • link count—Number of links in the advertisement. • id—ID of a routing device or subnet on the link. • data—For stub networks, the subnet mask. Otherwise, the IP address of the routing device that generated the LSP. • type—Type of link. It can be PointToPoint, Transit, Stub, or Virtual. • TOS count—Number of type-of-service (ToS) entries in the advertisement. • TOS 0 metric—Metric for ToS 0. • TOS—Type-of-service (ToS) value. • metric—Metric for the ToS.	`detail extensive`

Table 936: show ospf database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Network	Network link-state advertisement information: • mask—Network mask. • attached router—ID of the attached neighbor.	`detail extensive`
Summary	Summary link-state advertisement information: • mask—Network mask. • TOS—Type-of-service (ToS) value. • metric—Metric for the ToS.	`detail extensive`
Gen timer	How long until the LSA is regenerated.	`extensive`
Aging timer	How long until the LSA expires.	`extensive`
Installed *hh:mm:ss* ago	How long ago the route was installed.	`extensive`
expires in *hh:mm:ss*	How long until the route expires.	`extensive`
sent *hh:mm:ss* ago	How long ago the LSA was sent.	`extensive`
Last changed *hh:mm:ss* ago	How long ago the route was changed.	`extensive`
Change count	Number of times the route has changed.	`extensive`
Ours	Indicates that this is a local advertisement.	`extensive`
Router LSAs	Number of router link-state advertisements in the link-state database.	`summary`

Table 936: show ospf database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Network LSAs	Number of network link-state advertisements in the link-state database.	summary
Summary LSAs	Number of summary link-state advertisements in the link-state database.	summary
NSSA LSAs	Number of not-so-stubby area link-state advertisements in the link-state database.	summary

Sample Output

show ospf database

```
user@host> show ospf database
OSPF link state database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.70.103    10.255.70.103    0x80000002   215  0x20 0x4112  48
Router  *10.255.71.242    10.255.71.242    0x80000002   214  0x20 0x11b1  48
Summary *23.1.1.0         10.255.71.242    0x80000002   172  0x20 0x6d72  28
Summary *24.1.1.0         10.255.71.242    0x80000002   177  0x20 0x607e  28
NSSA    *33.1.1.1         10.255.71.242    0x80000002   217  0x28 0x73bd  36

    OSPF link state database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.52     10.255.71.52     0x80000004   174  0x20 0xd021  36
Router  *10.255.71.242    10.255.71.242    0x80000003   173  0x20 0xe191  36
Network *23.1.1.1         10.255.71.242    0x80000002   173  0x20 0x9c76  32
Summary *12.1.1.0         10.255.71.242    0x80000001   217  0x20 0xfeec  28
Summary *24.1.1.0         10.255.71.242    0x80000002   177  0x20 0x607e  28
NSSA    *33.1.1.1         10.255.71.242    0x80000001   222  0x28 0xe047  36

    OSPF link state database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.238    10.255.71.238    0x80000003   179  0x20 0x3942  36
Router  *10.255.71.242    10.255.71.242    0x80000003   177  0x20 0xf37d  36
Network *24.1.1.1         10.255.71.242    0x80000002   177  0x20 0xc591  32
```

```
Summary *12.1.1.0         10.255.71.242    0x80000001   217  0x20 0xfeec  28
Summary *23.1.1.0         10.255.71.242    0x80000002   172  0x20 0x6d72  28
NSSA    *33.1.1.1         10.255.71.242    0x80000001   222  0x28 0xeb3b  36
```

show ospf database

The output for `show ospf databse nssa` with `nssa-only` configuration statement enabled at `[edit policy-options policy-statement` *`policy-name`* `term` *`term name`* `then external]`, which clears P-bit on type 7 LSA.

```
user@host> show ospf database
OSPF link state database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.70.103    10.255.70.103    0x80000002   215  0x20 0x4112  48
Router  *10.255.71.242    10.255.71.242    0x80000002   214  0x20 0x11b1  48
Summary *23.1.1.0         10.255.71.242    0x80000002   172  0x20 0x6d72  28
Summary *24.1.1.0         10.255.71.242    0x80000002   177  0x20 0x607e  28
NSSA    *33.1.1.1         10.255.71.242    0x80000002   217  0x20 0x73bd  36

    OSPF link state database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.52     10.255.71.52     0x80000004   174  0x20 0xd021  36
Router  *10.255.71.242    10.255.71.242    0x80000003   173  0x20 0xe191  36
Network *23.1.1.1         10.255.71.242    0x80000002   173  0x20 0x9c76  32
Summary *12.1.1.0         10.255.71.242    0x80000001   217  0x20 0xfeec  28
Summary *24.1.1.0         10.255.71.242    0x80000002   177  0x20 0x607e  28
NSSA    *33.1.1.1         10.255.71.242    0x80000001   222  0x28 0xe047  36

    OSPF link state database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.238    10.255.71.238    0x80000003   179  0x20 0x3942  36
Router  *10.255.71.242    10.255.71.242    0x80000003   177  0x20 0xf37d  36
Network *24.1.1.1         10.255.71.242    0x80000002   177  0x20 0xc591  32
Summary *12.1.1.0         10.255.71.242    0x80000001   217  0x20 0xfeec  28
Summary *23.1.1.0         10.255.71.242    0x80000002   172  0x20 0x6d72  28
NSSA    *33.1.1.1         10.255.71.242    0x80000001   222  0x20 0xeb3b  36
```

show ospf database brief

The output for the `show ospf database brief` command is identical to that for the `show ospf database` command. For sample output, see "show ospf database " on page 22785.

show ospf database detail

```
user@host> show ospf database detail
    OSPF link state database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.70.103    10.255.70.103    0x80000002   261  0x20 0x4112  48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Router  *10.255.71.242    10.255.71.242    0x80000002   260  0x20 0x11b1  48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Summary *23.1.1.0         10.255.71.242    0x80000002   218  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0         10.255.71.242    0x80000002   223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA    *33.1.1.1         10.255.71.242    0x80000002   263  0x28 0x73bd  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0

    OSPF link state database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.52     10.255.71.52     0x80000004   220  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router  *10.255.71.242    10.255.71.242    0x80000003   219  0x20 0xe191  36
  bits 0x3, link count 1
  id 23.1.1.1, data 23.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *23.1.1.1         10.255.71.242    0x80000002   219  0x20 0x9c76  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.52
Summary *12.1.1.0         10.255.71.242    0x80000001   263  0x20 0xfeec  28
```

```
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0         10.255.71.242    0x80000002   223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA    *33.1.1.1         10.255.71.242    0x80000001   268  0x28 0xe047  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0

    OSPF link state database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.238    10.255.71.238    0x80000003   225  0x20 0x3942  36
  bits 0x0, link count 1
  id 24.1.1.1, data 24.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router  *10.255.71.242    10.255.71.242    0x80000003   223  0x20 0xf37d  36
  bits 0x3, link count 1
  id 24.1.1.1, data 24.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *24.1.1.1         10.255.71.242    0x80000002   223  0x20 0xc591  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.238
Summary *12.1.1.0         10.255.71.242    0x80000001   263  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *23.1.1.0         10.255.71.242    0x80000002   218  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA    *33.1.1.1         10.255.71.242    0x80000001   268  0x28 0xeb3b  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
```

show ospf database extensive

```
user@host>show ospf database extensive advertising-router self lsa-id 4.0.0.0
Flex-Algorithm Definition (16), length 28:
      Flex-Algo (1), length 1:
        128
      Metric-Type (2), length 1:
        2
```

```
      Calc-Type (3), length 1:
        1
      Priority (4),  length 1:
        10
     FAD AG Exclude (1), length 4:
        Exclude AG (1), length 4:
        0x4
     FAD AG Include Any (2), length 4:
        Include Any AG (1), length 4:
        0x2
     FAD AG Include All (3), length 4:
        Include All AG (1), length 4:
        0x2
     FAD Flags (4), length 4:
        Flags (1), length 4:
                     0X80000000
```

show ospf database extensive

```
user@host> show ospf database advertising-router self lsa-id 7.0.0.1 extensive
Extended Prefix (1), length 20:
      Route Type (1), length 1:
        1
      Prefix Length (2), length 1:
        32
      AF (3), length 1:
        0
      Flags (4),  length 1:
        0x40
      Prefix (5), length 32:
        2.2.2.2
      Prefix Sid (2), length 8:
       Flags (1), length 1:
        0x00
       MT ID (2), length 1:
        0
       Algorithm (3), length 1:
        0
       SID (4), length 4:
        1000
      FAPM (3), length 8:
```

```
        Algorithm (1), length 1:
         128
        Flags (2), length 1:
         0x00
        Metric (3), length 4:
         1000
```

show ospf database summary

```
user@host> show ospf database summary
Area 0.0.0.1:
   2 Router LSAs
   2 Summary LSAs
   1 NSSA LSAs
Area 0.0.0.2:
   2 Router LSAs
   1 Network LSAs
   2 Summary LSAs
   1 NSSA LSAs
Area 0.0.0.3:
   2 Router LSAs
   1 Network LSAs
   2 Summary LSAs
   1 NSSA LSAs
Externals:
Interface fe-2/2/1.0:
Interface ge-0/3/2.0:
Interface so-0/1/2.0:
Interface so-0/1/2.0:
```

show ospf database opaque-area extended-link link-addr ***link-address***

```
user@host> show ospf database extended-link link-addr 21.5.7.2

OSPF database, Area 0.0.0.0
Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
OpaqArea 8.0.0.1          7.7.7.7          0x8000021e   301  0x22 0x20ab  48
```

show ospf database opaque-area extended-prefix prefix ***prefix/len***

```
user@host>show ospf database extended-prefix  prefix 1.1.1.1/32

OSPF database, Area 0.0.0.0
Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
OpaqArea 7.0.0.1          1.1.1.1          0x80000213  1833  0x22 0x88ee  44
```

Release Information

Command introduced before Junos OS Release 7.4.

`advertising-router self` (*address* | `self`) option introduced in Junos OS Release 9.5.

`advertising-router self` (*address* | `self`) option introduced in Junos OS Release 9.5 for EX Series switches.

`show ospf database opaque-area extended-link link-addr` (*link-address*) option introduced in Junos OS Release 23.2R1 and Junos OS Evoved 23.2R1.

`show ospf database opaque-area extended-prefix prefix` (*prefix length*) option introduced in Junos OS Release 23.2R1 and Junos OS Evoved 23.2R1.

RELATED DOCUMENTATION

clear (ospf | ospf3) database | 3784

show ospf flex-algorithm

IN THIS SECTION

- Syntax | 22792
- Description | 22792
- Options | 22792
- Required Privilege Level | 22792
- Output Fields | 22792
- Sample Output | 22794

Release Information | 22799

Syntax

```
show ospf (spring | spf | route) flex-algorithm flex-algorithm-id
logical-system (all |logical-system-name)
```

Description

Display the configured flexible algorithm and the defined parameters.

Options

flex-algorithm-id	Specify the ID of the flex algorithm to view its configured parameters.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 937 on page 22792 describes the output fields for the `show ospf flex-algorithm` command. Output fields are listed in the approximate order in which they appear.

Table 937: show ospf flex-algorithm Output Fields

Field Name	Field Description
flex algorithm	Name of the defined flexible algorithm.
color	Name of the BGP colored community mapped with the flexible algorithm.

Table 937: show ospf flex-algorithm Output Fields *(Continued)*

Field Name	Field Description
Chosen Definition	Name of the flexible algorithm definition that the device has chosen to participate in.
calculation-type	The calculation type specified in the flexible algorithm definition (FAD).
metric-type	The metric type specified in the FAD
priority	The specified priority for the advertised flexible algorithm.
Include Any	The administrative groups included in the FAD.
Include All	All the administrative groups included in the FAD.
Exclude	The administrative groups to be excluded.
Prefix	Destination of the route.
Area	Area number. Area 0.0.0.0 is the backbone area.
Version	Version of SPF that generated the route.
Metric	Metric value associated with the route.
Type	Metric type: `int` (internal) or `ext` (external).
Interface	Interface to the next hop.

Table 937: show ospf flex-algorithm Output Fields *(Continued)*

Field Name	Field Description
NH	Next hop address type: • IPv4 • IPv6 • MPLS
Via	System identifier of the next hop, displayed as a name if possible.

Sample Output

show isis spring flex-algorithm flex-algorithm-id

```
user@host> show ospf spring flex-algorithm flex-algorithm-id 128
Flex algortihm : 129, Paritcipating
Color: <color>
Chosen Definition: <hostname>
calculation-type: spf/strict-spf
    metric-type: igp-metric/te-metric
priority: <priority>
Include Any: <admin-groups>
Include All: <admin-groups>
Exclude: <admin-groups>
```

```
user@host> show ospf spring flex-algorithm flex-algorithm-id 128
Flex Algo: 128, Area: 0.0.0.0
Color: 128, Participating, FAD supported
Winner: 192.168.255.10, Metric: 0, Calc: 0, Prio: 0, inc-any: 0x1, FAD supported
Include-Any:  0x1 RED
SPF Version: 296
Participation toggles: 1
Topo refresh count: 0
Full SPFs: 296, Partial SPFs: 0
```

show ospf spring flex-algorithm

```
user@host> show ospf spring flex-algorithm
Flex Algo: 128, Area: 0.0.0.0
  Color: 128, Participating, FAD supported
      Winner: 100.100.100.100, Metric: 2, Calc: 1, Prio: 10, inc-any: 0x2, inc-all: 0x2, Excl:
0x4, Flags: 0X80000000, FAD supported
        Include-Any:  0x2 red
        Include-All:  0x2 red
        Exclude:  0x4 green
  SPF Version: 14
  Participation toggles: 1
  Topo refresh count: 0
  Full SPFs: 14, Partial SPFs: 0
```

show ospf log

```
user@host> show ospf log
Topology default SPF log:

   Last instance of each event type
When            Type            Elapsed
1w2d 14:16:01   SPF             0.000316
1w2d 14:16:01   Stub            0.000233
1w2d 14:16:01   Interarea       0.000002
1w2d 14:16:01   External        0.000004
1w2d 14:16:01   NSSA            0.000001
1w2d 14:16:01   Cleanup         0.000551

   Maximum length of each event type
When            Type            Elapsed
1w2d 14:51:10   SPF             0.000997
1w2d 16:16:18   Stub            0.000675
1w3d 07:25:10   Interarea       0.000010
1w3d 07:45:50   External        0.000013
1w3d 07:32:04   NSSA            0.000008
1w3d 08:54:48   Cleanup         0.001044

   Last 100 events
When            Type            Elapsed
```

```
1w2d 14:25:19   FlexAlgo SPF    0.000680
1w2d 14:25:19   SPF             0.000204
1w2d 14:25:19   Stub            0.000025
1w2d 14:25:19   Interarea       0.000003
1w2d 14:25:19   External        0.000002
1w2d 14:25:19   NSSA            0.000001
1w2d 14:25:19   Prefix SID      0.000222
1w2d 14:25:19   Adj SID         0.000074
1w2d 14:25:19   Cleanup         0.000607
1w2d 14:25:19     Total         0.001209
1w2d 14:25:14   SPF             0.000188
1w2d 14:25:14   Stub            0.000054
1w2d 14:25:14   Interarea       0.000002
1w2d 14:25:14   External        0.000001
1w2d 14:25:14   NSSA            0.000001
1w2d 14:25:14   Prefix SID      0.000181
1w2d 14:25:14   Adj SID         0.000178
1w2d 14:25:14   Cleanup         0.000413
1w2d 14:25:14     Total         0.001656
1w2d 14:23:37   FlexAlgo SPF    0.001914
1w2d 14:23:37   FlexAlgo SPF    0.000081
1w2d 14:23:37   SPF             0.000215
```

show ospf route flex-algorithm-id

```
user@host> show ospf route flex-algorithm-id 128
Topology default Route Table:

Prefix             Path  Route      NH       Metric NextHop       Nexthop
                   Type  Type       Type            Interface     Address/LSP
192.168.255.11     Intra Router     IP            1 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
192.168.255.12     Intra Router     IP            2 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
                                                     ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
192.168.255.13     Intra Router     IP            1 ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
192.168.255.14     Intra Router     IP            2 ge-0/0/0.0    10.10.1.2
```

```
                                                ge-0/0/1.0    10.10.1.6
                                                ge-0/0/2.0    10.10.3.2
                                                ge-0/0/3.0    10.10.3.6
192.168.255.15     Intra Router     IP        3 ge-0/0/0.0    10.10.1.2
                                                ge-0/0/1.0    10.10.1.6
                                                ge-0/0/2.0    10.10.3.2
                                                ge-0/0/3.0    10.10.3.6
10.10.1.0/30       Intra Network    IP        1 ge-0/0/0.0
10.10.1.4/30       Intra Network    IP        1 ge-0/0/1.0
10.10.3.0/30       Intra Network    IP        1 ge-0/0/2.0
10.10.3.4/30       Intra Network    IP        1 ge-0/0/3.0
10.10.12.0/30      Intra Network    IP        2 ge-0/0/0.0    10.10.1.2
                                                ge-0/0/1.0    10.10.1.6
10.10.12.4/30      Intra Network    IP        2 ge-0/0/0.0    10.10.1.2
                                                ge-0/0/1.0    10.10.1.6
10.10.13.0/30      Intra Network    IP        2 ge-0/0/0.0    10.10.1.2
                                                ge-0/0/1.0    10.10.1.6
                                                ge-0/0/2.0    10.10.3.2
                                                ge-0/0/3.0    10.10.3.6
10.10.14.0/30      Intra Network    IP        2 ge-0/0/0.0    10.10.1.2
                                                ge-0/0/1.0    10.10.1.6
10.10.23.0/30      Intra Network    IP        2 ge-0/0/2.0    10.10.3.2
                                                ge-0/0/3.0    10.10.3.6
10.10.24.0/30      Intra Network    IP        3 ge-0/0/0.0    10.10.1.2
                                                ge-0/0/1.0    10.10.1.6
                                                ge-0/0/2.0    10.10.3.2
                                                ge-0/0/3.0    10.10.3.6
10.10.25.0/30      Intra Network    IP        3 ge-0/0/0.0    10.10.1.2
                                                ge-0/0/1.0    10.10.1.6
                                                ge-0/0/2.0    10.10.3.2
                                                ge-0/0/3.0    10.10.3.6
10.10.25.4/30      Intra Network    IP        3 ge-0/0/0.0    10.10.1.2
                                                ge-0/0/1.0    10.10.1.6
                                                ge-0/0/2.0    10.10.3.2
                                                ge-0/0/3.0    10.10.3.6
10.10.34.0/30      Intra Network    IP        2 ge-0/0/2.0    10.10.3.2
                                                ge-0/0/3.0    10.10.3.6
10.10.34.4/30      Intra Network    IP        2 ge-0/0/2.0    10.10.3.2
                                                ge-0/0/3.0    10.10.3.6
10.10.45.0/30      Intra Network    IP        3 ge-0/0/0.0    10.10.1.2
                                                ge-0/0/1.0    10.10.1.6
                                                ge-0/0/2.0    10.10.3.2
                                                ge-0/0/3.0    10.10.3.6
```

```
10.10.45.4/30      Intra Network    IP            3 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
                                                     ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
128.49.106.245/32  Intra Network    IP            1 ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
128.49.107.40/32   Intra Network    IP            2 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
                                                     ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
192.168.255.10/32  Intra Network    IP            0 lo0.0
192.168.255.10/32  Intra Network    Spring        0 lo0.0
192.168.255.11/32  Intra Network    IP            1 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
192.168.255.11/32  Intra Network    Spring        1 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
192.168.255.12/32  Intra Network    IP            2 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
                                                     ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
192.168.255.12/32  Intra Network    Spring        2 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
                                                     ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
192.168.255.13/32  Intra Network    IP            1 ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
192.168.255.13/32  Intra Network    Spring        1 ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
192.168.255.14/32  Intra Network    IP            2 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
                                                     ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
192.168.255.14/32  Intra Network    Spring        2 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
                                                     ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
192.168.255.15/32  Intra Network    IP            3 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
                                                     ge-0/0/2.0    10.10.3.2
                                                     ge-0/0/3.0    10.10.3.6
192.168.255.15/32  Intra Network    Spring        3 ge-0/0/0.0    10.10.1.2
                                                     ge-0/0/1.0    10.10.1.6
                                                     ge-0/0/2.0    10.10.3.2
```

```
                                                   ge-0/0/3.0    10.10.3.6
16                Intra Network    Mpls          0 ge-0/0/0.0    10.10.1.2
                                   Bkup MPLS       ge-0/0/1.0    10.10.1.6
16 (S=0)          Intra Network    Mpls          0 ge-0/0/0.0    10.10.1.2
                                   Bkup MPLS       ge-0/0/1.0    10.10.1.6
17                Intra Network    Mpls          0 ge-0/0/1.0    10.10.1.6
                                   Bkup MPLS       ge-0/0/0.0    10.10.1.2
17 (S=0)          Intra Network    Mpls          0 ge-0/0/1.0    10.10.1.6
                                   Bkup MPLS       ge-0/0/0.0    10.10.1.2
20                Intra Network    Mpls          0 ge-0/0/2.0    10.10.3.2
                                   Bkup MPLS       ge-0/0/3.0    10.10.3.6
20 (S=0)          Intra Network    Mpls          0 ge-0/0/2.0    10.10.3.2
                                   Bkup MPLS       ge-0/0/3.0    10.10.3.6
21                Intra Network    Mpls          0 ge-0/0/3.0    10.10.3.6
                                   Bkup MPLS       ge-0/0/2.0    10.10.3.2
21 (S=0)          Intra Network    Mpls          0 ge-0/0/3.0    10.10.3.6
                                   Bkup MPLS       ge-0/0/2.0    10.10.3.2
81001             Intra Network    Mpls          1 ge-0/0/0.0    10.10.1.2
                                                   ge-0/0/1.0    10.10.1.6
81001 (S=0)       Intra Network    Mpls          1 ge-0/0/0.0    10.10.1.2
                                                   ge-0/0/1.0    10.10.1.6
81003             Intra Network    Mpls          2 ge-0/0/0.0    10.10.1.2
                                                   ge-0/0/1.0    10.10.1.6
                                                   ge-0/0/2.0    10.10.3.2
                                                   ge-0/0/3.0    10.10.3.6
81004             Intra Network    Mpls          1 ge-0/0/2.0    10.10.3.2
                                                   ge-0/0/3.0    10.10.3.6
81004 (S=0)       Intra Network    Mpls          1 ge-0/0/2.0    10.10.3.2
                                                   ge-0/0/3.0    10.10.3.6
81006             Intra Network    Mpls          2 ge-0/0/0.0    10.10.1.2
                                                   ge-0/0/1.0    10.10.1.6
                                                   ge-0/0/2.0    10.10.3.2
                                                   ge-0/0/3.0    10.10.3.6
81007             Intra Network    Mpls          3 ge-0/0/0.0    10.10.1.2
                                                   ge-0/0/1.0    10.10.1.6
                                                   ge-0/0/2.0    10.10.3.2
                                                   ge-0/0/3.0    10.10.3.6
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Understanding OSPF Flexible Algorithm for Segment Routing

Example: OSPF Flexible Algorithm

flex-algorithm (Protocols OSPF)

definition (Protocols OSPF)

show (ospf | ospf3) backup coverage

IN THIS SECTION

- Syntax | **22800**
- Syntax (QFX Series) | **22801**
- Description | **22801**
- Options | **22801**
- Required Privilege Level | **22801**
- Output Fields | **22801**
- Sample Output | **22802**
- Release Information | **22803**

Syntax

```
show (ospf | ospf3) backup coverage
<instance instance-name>
< logical-system (all | logical-system-name)>
<realm (ipv4-unicast | ipv46-unicast>
<topology topology-name>
```

Syntax (QFX Series)

```
show (ospf | ospf3) backup coverage
<instance instance-name>
<topology topology-name>
```

Description

Display information about the level of backup coverage available for all the nodes and prefixes in the network.

Options

none	Display information about the level backup coverage for all OSPF routing instances in all logical systems.
logical-system (all | *logical-system-name*)	(Optional) Display information about the level of backup coverage for all logical systems or for a specific logical system.
instance *instance-name*	(Optional) Display information about the level of backup coverage for a specific OSPF routing instance.
realm (ipv4-unicast | ipv6-unicast)	(Optional) (OSPFv3 only) Display information about the level of backup coverage for the specific OSPFv3 realm, or address family.
topology (default | *topology-name*)	(Optional) (OSPFv2 only) Display information about the level of backup coverage for the specific OSPF topology.

Required Privilege Level

view

Output Fields

Table 938 on page 22802 lists the output fields for the `show (ospf | ospf3) backup coverage` command. Output fields are listed in the approximate order in which they appear.

Table 938: show (ospf | ospf3) backup coverage Output Fields

Field Name	Field Description
rmorn, June 2020: Uplift project, general cleanup	Information about backup coverage for each OSPF node.
Area	Area number. Area 0.0.0.0 is the backbone.
Covered Nodes	Number of nodes for which backup coverage is available.
Total Nodes	Total number of OSPF nodes.
Route Coverage	Information about backup coverage for each type of OSPF route.
Path Type	Type of OSPF path: Intra, Inter, Ext1, Ext2, and All.
Covered Routes	For each path type, the number of routes for which backup coverage is available.
Total Routes	For each path type, the total number of configured routes.
Percent Covered	For all nodes and for each path type, the percentage for which backup coverage is available.

Sample Output

show ospf backup coverage

```
user@host> show ospf backup coverage
Topology default coverage:

Node Coverage:

Area             Covered  Total  Percent
                   Nodes  Nodes  Covered
```

```
0.0.0.0                4      5   80.00%

Route Coverage:

Path Type  Covered   Total  Percent
            Routes  Routes  Covered
Intra            8      14   57.14%
Inter            0       0  100.00%
Ext1             0       0  100.00%
Ext2             1       1  100.00%
All              9      15   60.00%
```

show ospf3 backup coverage

```
user @host > show ospf3 backup coverage
show ospf3 backup coverage
Node Coverage:

Area             Covered  Total  Percent
                   Nodes  Nodes  Covered
0.0.0.0                4      5   80.00%

Route Coverage:

Path Type  Covered   Total  Percent
            Routes  Routes  Covered
Intra            4       6   66.67%
Inter            0       0  100.00%
Ext1             0       0  100.00%
Ext2             1       1  100.00%
All              5       7   71.43%
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show (ospf | ospf3) backup lsp | 22804

show (ospf | ospf3) backup lsp

IN THIS SECTION

- Syntax | 22804
- Description | 22804
- Options | 22804
- Required Privilege Level | 22805
- Output Fields | 22805
- Sample Output | 22806
- Sample Output | 22806
- Release Information | 22806

Syntax

```
show (ospf | ospf3) backup lsp
<logical-system (all | logical-system-name>
<realm (ipv4-unicast | ipv6-unicast)>
```

Description

Display information about MPLS label-switched-paths (LSPs) designated as backup routes for OSPF routes.

> **NOTE**: MPLS LSPs can be used as backup routes only for routes in the default OSPFv2 topology and not for any configured topology. Additionally, MPLS LSPs cannot be used as backup routes for nondefault instances either for OSPFv2 or OSPFv3.

Options

none	Display information all MPLS LSPs designated as backup routes.

logical-system (all | *logical-system-name*) (Optional) Display information about MPLS LSPs designated as backup routes for all logical systems or a specific logical system.

realm (ipv4-unicast | ipv6-unicast) (Optional) (OSPFv3 only) Display information about MPLS LSPs designated as backup routes for a specific realm, or address family.

Required Privilege Level

view

Output Fields

Table 939 on page 22805 lists the output fields for the `show (ospf | ospf3) backup lsp` command. Output fields are listed in the approximate order in which they appear.

Table 939: show (ospf | ospf3) backup lsp Output Fields

Field Name	Field Description
MPLS LSP name	Name of each MPLS LSP designated as a backup path.
Egress	IP address of the egress router for the LSP.

Status	State of the LSP:
	• Up—The router can detect RSVP hello messages from the neighbor.
	• Down—The router has received one of the following indications:
	  • Communication failure from the neighbor.
	  • Communication from IGP that the neighbor is unavailable.
	  • Change in the sequence numbers in the RSVP hello messages sent by the neighbor.
	• Deleted—The LSP is no longer available as a backup path.

Last change	Time elapsed since the neighbor state changed either from up or down or from down to up. The format is *hh:mm:ss*.

Table 939: show (ospf | ospf3) backup lsp Output Fields *(Continued)*

Field Name	Field Description
TE-metric	Configured traffic engineering metric.
Metric	Configured metric.

Sample Output

show ospf backup lsp

```
user@host> show ospf backup lsp
tobanff
  Egress: 10.255.71.239, Status: up, Last change: 00:00:23
  TE-metric: 0, Metric: 0
```

Sample Output

show ospf3 backup lsp

```
user@host> show ospf3 backup lsp
tobanff
  Egress: 10.255.71.239, Status: up, Last change: 00:00:45
  TE-metric: 0, Metric: 0
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show (ospf | ospf3) backup coverage | 22800

show (ospf | ospf3) backup neighbor

IN THIS SECTION

- Syntax | 22807
- Syntax (QFX Series) | 22807
- Description | 22807
- Options | 22808
- Required Privilege Level | 22808
- Output Fields | 22808
- Sample Output | 22809
- Release Information | 22809

Syntax

```
show (ospf | ospf3) backup neighbor
<area area-id>
<instance (default | instance-name)>
<logical-system (default | ipv4-multicast | logical-system-name)>
<topology (default | ipv4-multicast | topology-name)>
```

Syntax (QFX Series)

```
show (ospf | ospf3) backup neighbor
<area area-id>
<instance instance-name>
<topology (default | ipv4-multicast | topology-name)>
```

Description

Display the neighbors through which direct next hops for the backup paths are available.

Options

none	Display all neighbors that have direct next hops for backup paths.
area ***area-id***	(Optional) Display the area information.
instance (default | *instance-name*)	(Optional) Display information about the default routing instance or a particular routing instance.
logical-system (default | ipv4-multicast | *logical-system-name*)	(Optional) Display information about the default logical system, IPv4 multicast logical system, or a particular logical system.
topology (default | ipv4-multicast | *topology-name*)	(OSPFv2 only) (Optional) Display information about the default topology, IPv4 multicast topology, or a particular topology.

Required Privilege Level

view

Output Fields

Table 940 on page 22808 lists the output fields for the `show (ospf |ospf3) backup neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 940: show (ospf |ospf3) backup neighbor Output Fields

Field Name	Field Description	Level of Output
Neighbor to Self Metric	Metric from the backup neighbor to the OSPF node.	All levels
Self to Neighbor Metric	Metric from the OSPF node to the backup neighbor.	All levels
Direct next-hop	Interface and address of the direct next hop.	All levels

Sample Output

show ospf backup neighbor

```
user@host> show ospf backup neighbor
Topology default backup neighbors:

Area 0.0.0.5 backup neighbors:

10.0.0.5
  Neighbor to Self Metric: 5
  Self to Neighbor Metric: 5
  Direct next-hop: ge-4/0/0.111 via 10.0.175.5

10.0.0.6
  Neighbor to Self Metric: 5
  Self to Neighbor Metric: 5
  Direct next-hop: ge-4/1/0.110 via 10.0.176.6
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show (ospf | ospf3) backup spf | 22809

show (ospf | ospf3) backup spf

IN THIS SECTION

- Syntax | 22810
- Syntax (QFX Series) | 22810
- Description | 22810

- Options | 22811
- Required Privilege Level | 22811
- Output Fields | 22811
- Sample Output | 22812
- Release Information | 22813

Syntax

```
show (ospf | ospf3) backup spf
<brief | detail>
<area area-id>
<instance instance-name>
<logical-system (all | logical-system-name>
<no-coverage>
<node-id>
<realm (ipv4-unicast | ipv6-unicast)>
<topology (default | ipv4-multicast | topology-name)>
```

Syntax (QFX Series)

```
show (ospf | ospf3) backup spf
<brief | detail>
<area area-id>
<instance instance-name>
<no-coverage>
<node-id>
<topology (default | ipv4-multicast | topology-name)>
```

Description

Display information about OSPF shortest-path-first calculations for backup paths.

Options

none	Display information about OSPF shortest-path-first (SPF) calculations for all backup paths for all destination nodes.
brief | detail	(Optional) Display the specified level of output.
area *area-id*	(Optional) Display the area information.
instance *instance-name*	(Optional) Display information about the routing instance.
logical-system (all | *logical-system-name*)	(Optional) Display information about all logical systems or a specific logical system.
no-coverage	(Optional) Display information if there is no backup coverage.
node-id	(Optional) Display information about the node specified.
realm (ipv4–unicast | ipv6–unicast)	(Optional) Display information about the ipv4 or ipv6 realm.
topology (default | ipv4-multicast | *topology-name*)	(Optional) (OSPFv2 only) Display information about the default topology, IPv4 multicast topology, or a specifc topology.

Required Privilege Level

view

Output Fields

Table 941 on page 22811 lists the output fields for the `show (ospf |ospf3) backup spf` command. Output fields are listed in the approximate order in which they appear.

Table 941: show (ospf |ospf3) backup spf Output Fields

Field Name	Field Description	Level of Output
Area *area-id* results	Area for which the results are displayed. Area 0.0.0.0 is the backbone area.	All levels
address	Address of the node for which the results are displayed.	All levels

Table 941: show (ospf |ospf3) backup spf Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Self to Destination Metric	Metric from the node to the destination.	All levels
Parent Node	Address of the parent node.	All levels
Primary next-hop	Address of the next hop.	All levels
Backup Neighbor	Address of the backup neighbor or LSP endpoint and the following information: • Neighbor to Destination Metric • Neighbor to Self Metric • Self to Neighbor Metric • Status (Eligible, Not Eligible, Not Evaluated) and the reason for the status. **NOTE**: If the backup neighbor is an LSP endpoint, it is indicated as such after the neighbor address.	All levels

Sample Output

show ospf backup spf

```
user@host> show ospf backup spf
Topology default results:

Area 0.0.0.0 results:

pro16-d-lo0.xxx.yyyy.net
  Self to Destination Metric: 1
  Parent Node: pro16-b-lo0.xxx.yyyy.net
  Primary next-hop: at-1/0/1.0
```

```
  Backup Neighbor: pro16-c-lo0.xxx.yyyy.net (LSP endpoint)
    Neighbor to Destination Metric: 4, Neighbor to Self Metric: 3
    Self to Neighbor Metric: 3
    Not eligible, Reason: Path loops
  Backup Neighbor: pro16-d-lo0.xxx.yyyy.net
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  ...
```

Release Information

Command introduced in Junos OS Release 10.0.

show (ospf | ospf3) bgp-orr

IN THIS SECTION

- Syntax | 22813
- Description | 22814
- Options | 22814
- Required Privilege Level | 22815
- Output Fields | 22815
- Sample Output | 22816
- Release Information | 22816

Syntax

```
show (ospf | ospf3) bgp-orr
<abr>
<asbr>
<brief | detail | extensive>
<extern>
```

```
<group group-name>
<instance instance-name>
<inter>
<intra>
<logical-system (all | logical-system-name)>
<network>
<router>
<topology>
```

Description

Display information about OSPF BGP-ORR metric (RIB).

Options

none	Display group information about all OSPF BGP groups.
abr	(Optional) Display OSPF routes to area border routers.
asbr	(Optional) Display OSPF routes to autonomous systems border routers.
brief | detail | extensive	(Optional) Display the specified level of output.
extern	(Optional) Display external OSPF routes.
group *group-name*	(Optional) Display group information for the specified group.
instance *instance-name*	(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, **cust1**, **cust11**, and **cust111** are all displayed when you run the `show bgp group instance cust1` command). The instance name can be primary for the main instance, or any valid configured instance name or its prefix.
inter	(Optional) Display inter-area OSPF routes.
intra	(Optional) Display intra-area OSPF routes.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
network	(Optional) Display routes to networks.
router	(Optional) Display routes to all routers.

topology (Optional) Name of topology.

Required Privilege Level

view

Output Fields

Table 942 on page 22815 describes the output fields for the `show ospf bgp-orr` command. Output fields are listed in the approximate order in which they appear.

Table 942: show bgp group Output Fields

Field Name	Field Description	Level of Output
BGP ORR Peer Group	Name of the BGP ORR peer group.	All levels
Primary	Primary node (igp-primary) in a BGP peer group.	All levels
Backup	Backup node (igp-backup) in a BGP peer group, which is used when the primary node (igp-primary) goes down or becomes unreachable.	All levels
Prefix	Destination of the route.	All levels
Path Type	Display the route learned path (inter-area route or intra-area route).	All levels
Route Type	Display the type of router from which the route was learned (Router or Transit).	All levels
Metric	IGP metric value.	All levels

Sample Output

show ospf bgp-orr

```
user@host> show ospf bgp-orr
Topology default Route Table:

BGP ORR Peer Group: toClients
  Primary: 10.1.1.1, active
  Backup: 10.5.5.5
Prefix             Path  Route      Metric
                   Type  Type
10.1.1.2            Intra Router        100
10.1.1.3            Intra Router        130
10.1.1.1/32         Intra Network         0
10.1.1.2/32         Intra Network       100
10.1.1.0/30        Intra Network       100
10.1.1.4/30        Intra Network       130
```

show ospf3 bgp-orr

```
user@host> show ospf3 bgp-orr
BGP ORR Peer Group: toClients
  Primary: 10.1.1.1, active
Prefix                                        Path  Route      Metric
                                              Type  Type
10.1.1.2                                       Intra Router     100
10.1.1.2;0.0.0.4                               Intra Transit    100
10.1.1.3                                       Intra Router     130
10.1.1.3;0.0.0.2                               Intra Transit    130
::10.1.1.1/128                                 Intra Network    0
::10.1.1.2/128                                 Intra Network    100
::10.1.1.3/128                                 Intra Network    130
```

Release Information

Command introduced in Junos OS Release 16.2.

show (ospf | ospf3) interface

IN THIS SECTION

- Syntax | 22817
- Syntax (EX Series Switches and QFX Series) | 22817
- Description | 22818
- Options | 22818
- Required Privilege Level | 22818
- Output Fields | 22818
- Sample Output | 22822
- show ospf interface extensive (SRLG Protection Enabled) | 22824
- show ospf interface detail (HMAC-SHA1 Authentication Algorithm Enabled) | 22825
- show ospf interface detail (Delete-if-not-inuse Option Enabled) | 22826
- show ospf interface detail (HMAC-SHA2 Keychain Enabled | 22826
- show ospf interface detail (HMAC-SHA1 Authentication Algorithm Enabled for Virtual Link) | 22827
- show ospf interface detail (HMAC-SHA1 Authentication Algorithm Enabled for Sham Link) | 22828
- Release Information | 22828

Syntax

```
show (ospf | ospf3) interface
<brief | detail | extensive>
<area area-id>
<interface-name>
<instance instance-name>
<logical-system (all | logical-system-name)>
<realm (ip4-multicast | ipv4-unicast | ipv6-multicast)>
```

Syntax (EX Series Switches and QFX Series)

```
show (ospf | ospf3) interface
<brief | detail | extensive>
```

```
<area area-id>
<interface-name>
<instance instance-name>
```

Description

Display the status of OSPF interfaces.

Options

none	Display standard information about the status of all OSPF interfaces for all routing instances
brief | detail | extensive	(Optional) Display the specified level of output.
area *area-id*	(Optional) Display information about the interfaces that belong to the specified area.
interface-name	(Optional) Display information for the specified interface.
instance *instance-name*	(Optional) Display all OSPF interfaces under the named routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)	(OSPFv3 only) (Optional) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the `realm` option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

Output Fields

Table 943 on page 22819 lists the output fields for the `show (ospf | ospf3) interface` command. Output fields are listed in the approximate order in which they appear.

Table 943: show (ospf | ospf3) interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface running OSPF version 2 or OSPF version 3.	All levels
State	State of the interface: BDR, Down, DR, DRother, Loop, PtToPt, or Waiting.	All levels
Area	Number of the area that the interface is in.	All levels
DR ID	Address of the area's designated router.	All levels
BDR ID	Backup designated router for a particular subnet.	All levels
Nbrs	Number of neighbors on this interface.	All levels
Type	Type of interface: LAN, NBMA, P2MP, P2P, or Virtual.	`detail extensive`
Address	IP address of the neighbor.	`detail extensive`
Mask	Netmask of the neighbor.	`detail extensive`
Prefix-length	(OSPFv3) IPv6 prefix length, in bits.	`detail extensive`
OSPF3-Intf-Index	(OSPFv3) OSPF version 3 interface index.	`detail extensive`
MTU	Interface maximum transmission unit (MTU).	`detail extensive`
Cost	Interface cost (metric).	`detail extensive`
DR addr	Address of the designated router.	`detail extensive`
BDR addr	Address of the backup designated router.	`detail extensive`

Table 943: show (ospf | ospf3) interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Adj count	Number of adjacent neighbors.	`detail extensive`
Secondary	Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface for only one area.	`detail extensive`
Flood Reduction	Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the `DoNotAge` bit set. As a result, LSAs are refreshed only when a change occurs.	`extensive`
Priority	Router priority used in designated router (DR) election on this interface.	`detail extensive`
Flood list	List of link-state advertisements (LSAs) that might be about to flood this interface.	`extensive`
Ack list	Acknowledgment list. List of pending acknowledgments on this interface.	`extensive`
Descriptor list	List of packet descriptors.	`extensive`
Hello	Configured value for the hello timer.	`detail extensive`
Dead	Configured value for the dead timer.	`detail extensive`
Auth type	(OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: • MD5—The MD5 mechanism is configured in accordance with RFC 2328. • None—No authentication method is configured. • Password—A simple password (RFC 2328) is configured.	`detail extensive`

Table 943: show (ospf | ospf3) interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Topology	(Multiarea adjacency) Name of topology: default or *name*.	`detail extensive`
LDP sync state	(OSPFv2 and LDP synchronization) Current state of LDP synchronization: in sync, in holddown, and not supported.	`extensive`
reason	(OSPFv2 and LDP synchronization) Reason for the current state of LDP synchronization. The LDP session might be up or down, or adjacency might be up or down.	`extensive`
config holdtime	(OSPFv2 and LDP synchronization) Configured value of the hold timer. If the state is not synchronized, and the hold time is not infinity, the `remaining` field displays the number of seconds that remain until the configured hold timer expires.	`extensive`
IPSec SA name	(OSPFv2) Name of the IPSec security association name.	`detail extensive`
Active key ID	(OSPFv2 and MD5) Number from 0 to 255 that uniquely identifies an MD5 key.	`detail extensive`
Start time	(OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as Start time 1970 Jan 01 00:00:00 PST.	`detail extensive`
ReXmit	Configured value for the Retransmit timer.	`detail extensive`
Stub, Not Stub, or Stub NSSA	Type of area.	`detail extensive`

Table 943: show (ospf | ospf3) interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Post convergence Protection	Post convergence protection can have the following types when enabled • `Fate Sharing` can have the following values • `Yes`-You have configured fate-sharing protection. • `No`-You have not configured fate-sharing protection. • `node protection` can have the following values: • `Yes`-You have configured node protection. • `No`-You have not configured node protection. • `srlg protection` can have the following values: • `Yes`–You have configured Shared Risk Link Groutp (SRLG) protection. • `No`–You have not configured SRLG protection. `Node cost` is the recalculated metric cost of the node.	`extensive`

Sample Output

show ospf interface brief

```
user@host> show ospf interface brief
Intf                State     Area            DR ID           BDR ID       Nbrs
at-5/1/0.0          PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
ge-2/3/0.0          DR       0.0.0.0         192.168.4.16    192.168.4.15    1
lo0.0               DR       0.0.0.0         192.168.4.16    0.0.0.0         0
so-0/0/0.0          Down     0.0.0.0         0.0.0.0         0.0.0.0         0
so-6/0/1.0          PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-6/0/2.0          Down     0.0.0.0         0.0.0.0         0.0.0.0         0
so-6/0/3.0          PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
```

show ospf interface detail

```
user@host> show ospf interface detail
Interface              State     Area            DR ID           BDR ID Nbrs
fe-0/0/1.0             BDR      0.0.0.0         192.168.37.12   10.255.245.215 1
Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40
DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
t1-0/2/1.0             PtToPt   0.0.0.0         0.0.0.0         0.0.0.0 0
Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604
  Adj count 0
Hello 10, Dead 40, ReXmit 5, Not Stub
Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST
IPsec SA Name: sa
```

show ospf interface detail

```
user@hostshow ospf interface <> detail
Interface           State   Area            DR ID           BDR ID          Nbrs
ge-0/0/5.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 14.14.14.14, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 1
  , Active Key ID: 0
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

show ospf3 interface detail

```
user@host> show ospf3 interface so-0/0/3.0 detail
Interface              State     Area           DR-ID         BDR-ID     Nbrs
so-0/0/3.0             PtToPt   0.0.0.0        0.0.0.0       0.0.0.0        1
Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64
OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1
Hello 10, Dead 40, ReXmit 5, Not Stub
```

show ospf interface extensive (SRLG Protection Enabled)

command-name

```
user@host> show ospf interface extensive
Interface           State   Area            DR ID           BDR ID          Nbrs
ge-0/0/0.0          DR      0.0.0.0         10.205.172.20   10.205.171.195     1
  Type: LAN, Address: 81.1.2.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  DR addr: 81.1.2.1, BDR addr: 81.1.2.2, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: Post Convergence
  Post convergence protection: Enabled, Fate sharing: No, SRLG: Yes, Node cost: 65535
  Topology default (ID 0) -> Cost: 1
•    Checking backup route in rib:
root@R0# run show route 6.6.6.6
inet.0: 61 destinations, 61 routes (61 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

6.6.6.6/32         *[OSPF/10] 00:08:52, metric 20
                    > to 41.41.41.2 via ge-0/0/1.0
                    > to 31.31.31.2 via ge-0/0/2.0, Push 800030

inet.3: 6 destinations, 10 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

6.6.6.6/32         *[LDP/9] 00:07:33, metric 1
                    > to 41.41.41.2 via ge-0/0/1.0, Push 299808
                    [L-OSPF/10/5] 00:07:33, metric 20
                    > to 41.41.41.2 via ge-0/0/1.0, Push 800060
                    > to 31.31.31.2 via ge-0/0/2.0, Push 800030, Push 800060
```

command-name

```
user@host> show ospf interface extensive (Fate-Sharing Protection Enabled)
Interface           State   Area            DR ID           BDR ID          Nbrs
ge-0/0/0.0          DR      0.0.0.0         10.205.172.20   10.205.171.195     1
  Type: LAN, Address: 81.1.2.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  DR addr: 81.1.2.1, BDR addr: 81.1.2.2, Priority: 128
```

```
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: Post Convergence
  Post convergence protection: Enabled, Fate sharing: Yes, SRLG: No, Node cost: 65535
  Topology default (ID 0) -> Cost: 1
•    Checking backup route in rib:
root@R0# run show route 6.6.6.6
inet.0: 61 destinations, 61 routes (61 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

6.6.6.6/32         *[OSPF/10] 00:08:52, metric 20
                    > to 41.41.41.2 via ge-0/0/1.0
                    > to 31.31.31.2 via ge-0/0/2.0, Push 800030

inet.3: 6 destinations, 10 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

6.6.6.6/32         *[LDP/9] 00:07:33, metric 1
                    > to 41.41.41.2 via ge-0/0/1.0, Push 299808
                    [L-OSPF/10/5] 00:07:33, metric 20
                    > to 41.41.41.2 via ge-0/0/1.0, Push 800060
                    > to 31.31.31.2 via ge-0/0/2.0, Push 800030, Push 800060
```

show ospf interface detail (HMAC-SHA1 Authentication Algorithm Enabled)

```
user@host> show ospf interface detail
Interface           State   Area            DR ID           BDR ID          Nbrs
ge-0/0/5.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 14.14.14.14, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: Cryptographic, Active key ID: 0,
  Auth algo: HMAC-SHA1
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

show ospf interface detail (Delete-if-not-inuse Option Enabled)

```
user@host> show ospf interface detail
Interface           State   Area            DR ID           BDR ID          Nbrs
ge-0/0/5.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 14.14.14.14, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 0,
  Active key ID:1
  Delete-if-not-inuse:Yes
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

show ospf interface detail (HMAC-SHA2 Keychain Enabled

HMAC-SHA2-224 Keychain

```
user@host> show ospf interface detail
Interface           State   Area            DR ID           BDR ID          Nbrs
ge-0/0/5.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 14.14.14.14, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: Cryptographic, Active key ID: 0,
  Auth algo: HMAC-SHA2-224
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

HMAC-SHA2-256 Keychain

```
user@host> show ospf interface detail
Interface           State   Area            DR ID           BDR ID          Nbrs
ge-0/0/5.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 14.14.14.14, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: Cryptographic, Active key ID: 0,
```

```
  Auth algo: HMAC-SHA2-256
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

HMAC-SHA2-384 Keychain

```
user@host> show ospf interface detail
Interface           State   Area            DR ID           BDR ID          Nbrs
ge-0/0/5.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 14.14.14.14, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: Cryptographic, Active key ID: 0,
  Auth algo: HMAC-SHA2-384
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

HMAC-SHA2-512 keychain

```
user@host> show ospf interface detail
Interface           State   Area            DR ID           BDR ID          Nbrs
ge-0/0/5.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 14.14.14.14, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: Cryptographic, Active key ID: 0,
  Auth algo: HMAC-SHA2-512
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

show ospf interface detail (HMAC-SHA1 Authentication Algorithm Enabled for Virtual Link)

```
user@host> show ospf interface detail
vl-2.2.2.2          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: Virtual, Address: 15.15.15.11, Mask: 0.0.0.0, MTU: 0, Cost: 1
  Transit Area: 0.0.0.1, Destination: 15.15.15.12
```

```
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth Keychain name: red
  Auth type: Cryptographic, Active key ID: 3
  Auth algo: hmac-sha1
  Protection type: None, No eligible backup, No eligible remote backup
  Topology default (ID 0) -> Cost: 1  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

show ospf interface detail (HMAC-SHA1 Authentication Algorithm Enabled for Sham Link)

```
user@host> show ospf interface detail
shamlink.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 0, Cost: 10
  Local: 198.51.100.2, Remote: 198.51.100.4
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth Keychain name: red
  Auth type: Cryptographic, Active key ID: 3
  Auth algo: hmac-sha1
  Protection type: None, No eligible backup, No eligible remote backup
  Topology default (ID 0) -> Cost: 10
```

Release Information

Command introduced before Junos OS Release 7.4.

`area` option introduced in Junos OS Release 9.2.

`area` option introduced in Junos OS Release 9.2 for EX Series switches.

`realm` option introduced in Junos OS Release 9.2.

show (ospf | ospf3) io-statistics

IN THIS SECTION

- Syntax | 22829
- Syntax (EX Series Switch and QFX Series) | 22829
- Description | 22829
- Options | 22829
- Required Privilege Level | 22830
- Output Fields | 22830
- Sample Output | 22830
- Release Information | 22830

Syntax

```
show (ospf | ospf3) io-statistics
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and QFX Series)

```
show (ospf | ospf3) io-statistics
```

Description

Display Open Shortest Path First (OSPF) input and output statistics.

Options

none	Display OSPF input and output statistics.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 944 on page 22830 lists the output fields for the `show ospf io-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 944: show (ospf | ospf3) io-statistics Output Fields

Field Name	Field Description
Packets read	Number of OSPF packets read since the last time the routing protocol was started.
average per run	Total number of packets divided by the total number of times the OSPF read operation is scheduled to run.
max run	Maximum number of packets for a given run among all scheduled runs.
Receive errors	Number of faulty packets received with errors.

Sample Output

show ospf io-statistics

```
user@host> show ospf io-statistics

Packets read: 7361, average per run: 1.00, max run: 1
Receive errors:
  None
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear (ospf | ospf3) statistics | 3795

show (ospf | ospf3) log

IN THIS SECTION

- Syntax | 22831
- Syntax (EX Series Switch and QFX Series) | 22831
- Description | 22832
- Options | 22832
- Required Privilege Level | 22832
- Output Fields | 22832
- Sample Output | 22833
- Release Information | 22834

Syntax

```
show (ospf | osfp3) log
<instance instance-name>
<logical-system (all | logical-system-name)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<topology topology-name>
```

Syntax (EX Series Switch and QFX Series)

```
show (ospf | osfp3) log
<instance instance-name>
<topology topology-name>
```

Description

Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations.

Options

none	Display entries in the OSPF log of SPF calculations for all routing instances.
instance *instance-name*	(Optional) Display entries for the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
topology *topology-name*	(Optional) (OSPFv2 only) Display entries for the specified topology.
realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)	(OSPFv3 only) (Optional) Display entries for the specified OSPFv3 realm, or address family. Use the `realm` option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

Output Fields

Table 945 on page 22832 lists the output fields for the `show (ospf | ospf3) log` command. Output fields are listed in the approximate order in which they appear.

Table 945: show (ospf | ospf3) log Output Fields

Field Name	Field Description
When	Time, in weeks (`w`) and days (`d`), since the SPF calculation was made.
Type	Type of calculation: Cleanup, External, Interarea, NSSA, Redist, SPF, Stub, Total, or Virtuallink.

Table 945: show (ospf | ospf3) log Output Fields *(Continued)*

Field Name	Field Description
Elapsed	Amount of time, in seconds, that elapsed during the operation, or the time required to complete the SPF calculation. The start time is the time displayed in the `When` field.

Sample Output

show ospf log

```
user@host> show ospf log
When            Type            Elapsed

1w4d 17:25:58   Stub            0.000017
1w4d 17:25:58   SPF             0.000070
1w4d 17:25:58   Stub            0.000019
1w4d 17:25:58   Interarea       0.000054
1w4d 17:25:58   External        0.000005
1w4d 17:25:58   Cleanup         0.000203
1w4d 17:25:58     Total         0.000537
1w4d 17:24:48   SPF             0.000125
1w4d 17:24:48   Stub            0.000017
1w4d 17:24:48   SPF             0.000100
1w4d 17:24:48   Stub            0.000016
1w4d 17:24:48   Interarea       0.000056
1w4d 17:24:48   External        0.000005
1w4d 17:24:48   Cleanup         0.000238
1w4d 17:24:48     Total         0.000600
...
```

show ospf log topology voice

```
user@host> show ospf log topology voice
Topology voice SPF log:

   Last instance of each event type
```

```
When            Type            Elapsed
00:06:11        SPF             0.000116
00:06:11        Stub            0.000114
00:06:11        Interarea       0.000126
00:06:11        External        0.000067
00:06:11        NSSA            0.000037
00:06:11        Cleanup         0.000186

   Maximum length of each event type
When            Type            Elapsed
00:13:43        SPF             0.000140
00:13:33        Stub            0.000116
00:13:43        Interarea       0.000128
00:13:33        External        0.000075
00:13:38        NSSA            0.000039
00:13:53        Cleanup         0.000657

   Last 100 events
When            Type            Elapsed

00:13:53        SPF             0.000090
00:13:53        Stub            0.000041
00:13:53        Interarea       0.000123
00:13:53        External        0.000040
00:13:53        NSSA            0.000038
00:13:53        Cleanup         0.000657
00:13:53          Total         0.001252
.
.
00:06:11        SPF             0.000116
00:06:11        Stub            0.000114
00:06:11        Interarea       0.000126
00:06:11        External        0.000067
00:06:11        NSSA            0.000037
00:06:11        Cleanup         0.000186
00:06:11          Total         0.000818
```

Release Information

Command introduced before Junos OS Release 7.4.

`topology` option introduced in Junos OS Release 9.0.

`topology` option introduced in Junos OS Release 9.0 for EX Series switches.

`realm` option introduced in Junos OS Release 9.2.

show (ospf | ospf3) neighbor

IN THIS SECTION

- Syntax | 22835
- Syntax (EX Series Switches and QFX Series) | 22835
- Description | 22836
- Options | 22836
- Required Privilege Level | 22837
- Output Fields | 22837
- Sample Output | 22840
- Release Information | 22843

Syntax

```
show (ospf | ospf3) neighbor
<brief | detail | extensive>
<area area-id>
<instance (all |  instance-name)>
<interface interface-name>
<logical-system (all | logical-system-name)>
<neighbor>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
```

Syntax (EX Series Switches and QFX Series)

```
show (ospf | ospf3) neighbor
<brief | detail | extensive>
<area area-id>
```

```
<instance (all |  instance-name)>
<interface interface-name>
<neighbor>
```

Description

Display information about OSPF neighbors.

CPU utilization might increase while the device learns its OSPF neighbors. We recommend that you use the `show (ospf | ospf3) neighbor` command after the device learns and establishes OSPF neighbor adjacencies. Depending on the size of your network, this might take several minutes. If you receive a "timeout communicating with routing daemon" error when using the `show (ospf | ospf3) neighbor` command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the CLI.

Options

none	Display standard information about all OSPF neighbors for all routing instances.
brief | detail | extensive	(Optional) Display the specified level of output.
area *area-id*	(Optional) Display information about the OSPF neighbors for the specified area.
instance (all | *instance-name*)	(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.
interface *interface-name*	(Optional) Display information about OSPF neighbors for the specified logical interface.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
neighbor	(Optional) Display information about the specified OSPF neighbor.
realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)	(OSPFv3 only) (Optional) Display information about the OSPF neighbors for the specified OSPFv3 realm, or address family. Use the `realm` option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

Output Fields

Table 946 on page 22837 lists the output fields for the `show (ospf | ospf3) neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 946: show (ospf | ospf3) neighbor Output Fields

Field Name	Field Description	Level of Output
Address	Address of the neighbor.	All levels
Interface	Interface through which the neighbor is reachable.	All levels

Table 946: show (ospf | ospf3) neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the neighbor: • Attempt—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor. • Down—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the Down state, although at a reduced frequency. • Exchange—Routing device is describing its entire link-state database by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged. • ExStart—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the primary, and to determine the initial sequence number. • Full—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements. • Init—A hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state may occur, for example, because the routing device itself did not appear in the neighbor's hello packet. • Loading—Link-state request packets are sent to the neighbor to acquire more recent advertisements that have been discovered (but not yet received) in the Exchange state. • 2Way—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state 2Way or greater.	All levels
ID	Router ID of the neighbor.	All levels

Table 946: show (ospf | ospf3) neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Pri	Priority of the neighbor to become the designated router.	All levels
Dead	Number of seconds until the neighbor becomes unreachable.	All levels
Link state acknowledgment list	Number of link-state acknowledgments received.	`extensive`
Link state retransmission list	Total number of link-state advertisements retransmitted. For extensive output only, the following information is also displayed: • Type—Type of link advertisement: ASBR, Sum, Extern, Network, NSSA, OpaqArea, Router, or Summary. • LSA ID—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device. • Adv rtr—Address of the routing device that sent the advertisement. • Seq—Link sequence number of the advertisement.	`detail extensive`
Neighbor-address	(OSPFv3 only) If the neighbor uses virtual links, the Neighbor-address is the site-local, local, or global address. If the neighbor uses a physical interface, the Neighbor-address is an IPv6 link-local address.	`detail extensive`
area	Area that the neighbor is in.	`detail extensive`
OSPF3-Intf-Index	(OSPFv3 only) Displays the OSPFv3 interface index.	`detail extensive`
opt	Option bits received in the hello packets from the neighbor.	`detail extensive`
DR or DR-ID	Address of the designated router.	`detail extensive`

Table 946: show (ospf | ospf3) neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
BDR or BDR-ID	Address of the backup designated router.	`detail extensive`
Up	Length of time since the neighbor came up.	`detail extensive`
adjacent	Length of time since the adjacency with the neighbor was established.	`detail extensive`
SPRING Adjacency Labels	Segment routing in networking adjacency labels. **NOTE**: Displayed only when segment routing is enabled	`detail extensive`
Label	Segment routing label.	`detail extensive`
Flags	Segment routing flags. Flags VL indicate *value* and *local*.	`detail extensive`

Sample Output

show ospf neighbor brief

```
user@host> show ospf neighbor brief
  Address          Intf               State      ID              Pri  Dead
192.168.254.225  fxp3.0              2Way      10.250.240.32    128   36
192.168.254.230  fxp3.0              Full      10.250.240.8     128   38
192.168.254.229  fxp3.0              Full      10.250.240.35    128   33
10.1.1.129       fxp2.0              Full      10.250.240.12    128   37
10.1.1.131       fxp2.0              Full      10.250.240.11    128   38
10.1.2.1         fxp1.0              Full      10.250.240.9     128   32
10.1.2.81        fxp0.0              Full      10.250.240.10    128   33
```

show ospf neighbor detail

```
user@host> show ospf neighbor detail
Address          Interface              State     ID               Pri  Dead
```

```
10.0.6.60        lt-1/2/0.12            Full      1.1.1.60         128    38
  Area 0.0.0.0, opt 0x52, DR 0.0.0.0, BDR 0.0.0.0
  Up 23:53:47, adjacent 23:53:34
  SPRING Adjacency Labels:

    Label       Flags

    299968      VL

10.0.10.70       lt-1/2/0.14            Full      1.1.1.70         128    37
  Area 0.0.0.0, opt 0x52, DR 0.0.0.0, BDR 0.0.0.0
  Up 23:53:47, adjacent 23:53:47
  SPRING Adjacency Labels:

     Label       Flags

     299952      VL
```

show ospf neighbor extensive

```
user@host> show ospf neighbor extensive
  Address          Interface             State      ID               Pri  Dead
10.5.1.2         ge-1/2/0.1             Full      10.5.1.2         128   33
  area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
  Up 06:09:42, adjacent 05:17:50
    Link state retransmission list:

      Type      LSA ID           Adv rtr          Seq

     Summary  10.8.56.0        172.25.27.82     0x8000004d

     Router   10.5.1.94        10.5.1.94        0x8000005c

     Network  10.5.24.2        10.5.1.94        0x80000036

     Summary  10.8.57.0        172.25.27.82     0x80000024

     Extern   1.10.90.0        10.8.1.2         0x80000041

     Extern   1.4.109.0        10.6.1.2         0x80000041
```

```
     Router   10.5.1.190       10.5.1.190       0x8000005f

     Network  10.5.48.2        10.5.1.190       0x8000003d

     Summary  10.8.58.0        172.25.27.82     0x8000004d

     Extern   1.10.91.0        10.8.1.2         0x80000041

     Extern   1.4.110.0        10.6.1.2         0x80000041

     Router   10.5.1.18        10.5.1.18        0x8000005f

     Network  10.5.5.2         10.5.1.18        0x80000033

     Summary  10.8.59.0        172.25.27.82     0x8000003a

     Summary  10.8.62.0        172.25.27.82     0x80000025

10.5.10.2        ge-1/2/0.10            ExStart   10.5.1.38        128   38
  area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
  Up 06:09:42
    master, seq 0xac1530f8, rexmit DBD in 2 sec
    rexmit LSREQ in 0 sec
10.5.11.2        ge-1/2/0.11            Full      10.5.1.42        128   33
  area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
  Up 06:09:42, adjacent 05:27:00
    Link state retransmission list:

      Type      LSA ID           Adv rtr          Seq

     Summary  10.8.58.0        172.25.27.82     0x8000004d

     Extern   1.10.91.0        10.8.1.2         0x80000041

     Extern   1.1.247.0        10.5.1.2         0x8000003f

     Extern   1.4.110.0        10.6.1.2         0x80000041

     Router   10.5.1.18        10.5.1.18        0x8000005f

     Network  10.5.5.2         10.5.1.18        0x80000033
```

```
    Summary  10.8.59.0        172.25.27.82     0x8000003a
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance all` option introduced in Junos OS Release 9.1.

`instance all` option introduced in Junos OS Release 9.1 for EX Series switches.

`area`, `interface`, and `realm` options introduced in Junos OS Release 9.2.

`area` and `interface` options introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

clear (ospf | ospf3) neighbor | 3791

show (ospf | ospf3) overview

IN THIS SECTION

- Syntax | 22844
- Syntax (EX Series Switch and QFX Series) | 22844
- Description | 22844
- Options | 22844
- Required Privilege Level | 22844
- Output Fields | 22845
- Sample Output | 22848
- Release Information | 22851

Syntax

```
show (ospf | ospf3) overview
<brief | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
```

Syntax (EX Series Switch and QFX Series)

```
show (ospf | ospf3) overview
<brief | extensive>
<instance instance-name>
```

Description

Display Open Shortest Path First (OSPF) overview information.

Options

none	Display standard information about all OSPF neighbors for all routing instances.
brief | extensive	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display all OSPF interfaces under the named routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)	(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the `realm` option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

Output Fields

Table 947 on page 22845 lists the output fields for the `show ospf overview` command. Output fields are listed in the approximate order in which they appear.

Table 947: show ospf overview Output Fields

Field name	Field Description	Level of Output
Instance	OSPF routing instance.	All levels
Router ID	Router ID of the routing device.	All levels
Route table index	Route table index.	All levels
Configured overload	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
Topology	Topology identifier.	All levels
Prefix export count	Number of prefixes exported into OSPF.	All levels
Full SPF runs	Number of complete Shortest Path First calculations.	All levels
SPF delay	Delay before performing consecutive Shortest Path First calculations.	All levels
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.	All levels
LSA refresh time	Refresh period for link-state advertisement (in minutes).	All levels

Table 947: show ospf overview Output Fields *(Continued)*

Field name	Field Description	Level of Output
SPRING	Source protocol routing in networking: enable or disable.	All levels
Node Segments	Nodes of source protocol routing in networking:enable or disable.	All levels
Ipv4 Index	Ipv4 Index.	All levels
Index Range	Ipv4 Index range.	All levels
Node Segment Blocks Allocated	Details about node segment blocks.	All levels
Database protection state	Current state of database protection.	All levels
Warning threshold	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
Non self-generated LSAs	Number of LSAs whose router ID is not equal to the local router ID: Current, Warning (threshold), and Allowed.	All levels
Ignore time	How long the database has been in the ignore state.	All levels
Reset time	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
Ignore count	Number of times the database has been in the ignore state: Current and Allowed.	All levels
Restart	Graceful restart capability: enabled or disabled.	All levels
Restart duration	Time period for complete reacquisition of OSPF neighbors.	All levels

Table 947: show ospf overview Output Fields *(Continued)*

Field name	Field Description	Level of Output
Restart grace period	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels
Graceful restart helper mode	(OSPFv2) Standard graceful restart helper capability (based on RFC 3623): enabled or disabled.	All levels
Restart-signaling helper mode	(OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): enabled or disabled.	All levels
Helper mode	(OSPFv3) Graceful restart helper capability: enabled or disabled.	All levels
Trace options	OSPF-specific trace options.	`extensive`
Trace file	Name of the file to receive the output of the tracing operation.	`extensive`
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: Normal Stub, Not Stub, or Not so Stubby Stub.	All levels
Authentication Type	Type of authentication: None, Password, or MD5. **NOTE**: The Authentication Type field refers to the authentication configured at the [`edit protocols ospf area` *`area-id`*] level. Any authentication configured for an interface in this area will not affect the value of this field.	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

Sample Output

show ospf overview (without SRGB)

```
user@host> show ospf overview
Instance: master
    Router ID: 10.255.245.6
    Route table index: 0
      Configured overload, expires in 118 seconds
    LSA refresh time: 50 minutes
    SPRING: Enabled
    Node Segments: Enabled
    Ipv4 Index : 10, Index Range: 2048
    Node Segment Blocks Allocated:
    Start Index : 0, Size : 256, Label-Range: [ 802048, 802303 ]
    Start Index : 256, Size : 256, Label-Range: [ 802304, 802559 ]
    Start Index : 512, Size : 256, Label-Range: [ 802560, 802815 ]
    Start Index : 768, Size : 256, Label-Range: [ 802816, 803071 ]
    Start Index : 1024, Size : 256, Label-Range: [ 803072, 803327 ]
    Start Index : 1280, Size : 256, Label-Range: [ 803328, 803583 ]
    Start Index : 1536, Size : 256, Label-Range: [ 803584, 803839 ]
    Start Index : 1792, Size : 256, Label-Range: [ 803840, 804095 ]
     Restart: Enabled
        Restart duration: 20 sec
        Restart grace period: 40 sec
        Helper mode: enabled
     Area: 0.0.0.0
      Stub type: Not Stub
      Authentication Type: None
      Area border routers: 0, AS boundary routers: 0
      Neighbors
        Up (in full state): 0
    Topology: default (ID 0)
      Prefix export count: 0
      Full SPF runs: 1
      SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
```

show ospf overview (with SRGB)

```
user@host> show ospf overview
Instance: master
  Router ID: 10.10.10.10
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  SPRING: Enabled
    SRGB Config Range :
      SRGB Start-Label : 1000, SRGB Index-Range : 2000
    SRGB Block Allocation: Success
      SRGB Start Index : 1000, SRGB Size : 2000, Label-Range: [ 1000, 2999 ]
    Node Segments: Enabled
    Ipv4 Index : 1000
  Post Convergence Backup: Disabled
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 3
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 5
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Enabled, Remote Backup calculation enabled
```

show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
  Router ID: 10.255.112.218
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Restart: Enabled
    Restart duration: 180 sec
    Restart grace period: 210 sec
    Graceful restart helper mode: Enabled
```

```
    Restart-signaling helper mode: Enabled
  Database protection state: Normal
    Warning threshold: 70 percent
    Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 1
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 160
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 70
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
  Router ID: 10.255.112.128
  Route table index: 0
  LSA refresh time: 50 minutes
  Database protection state: Normal
    Warning threshold: 80 percent
    Non self-generated LSAs: Current 3, Warning 8, Allowed 10
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 2
  Area: 0.0.0.0
    Stub type: Not Stub
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 1
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 7
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
  Router ID: 1.1.1.103
  Route table index: 0
  Full SPF runs: 13, SPF delay: 0.200000 sec
  LSA refresh time: 50 minutes
  Restart: Disabled
  Trace options: lsa
  Trace file: /var/log/ospf size 131072 files 10
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 1
```

Release Information

Command introduced in Junos OS Release 7.4.

`realm` option introduced in Junos OS Release 9.2.

Database protection introduced in Junos 10.2.

show (ospf | ospf3) route

IN THIS SECTION

- Syntax | 22852
- Syntax (EX Series Switch and QFX Series) | 22852
- Description | 22852
- Options | 22853
- Required Privilege Level | 22854
- Output Fields | 22854

Sample Output | 22856
Release Information | 22859

Syntax

```
show (ospf | ospf3) route
<brief | detail | extensive>
<abr | asbr | extern | inter | intra>
<destination>
<instance (default | ipv4-multicast |  instance-name)>
<logical-system (default | ipv4-multicast | logical-system-name)>
<network>
<no-backup-coverage>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>
<topology  (default | ipv4-multicast | topology-name)>
<transit>
```

Syntax (EX Series Switch and QFX Series)

```
show (ospf | ospf3) route
<brief | detail | extensive>
<abr | asbr | extern | inter | intra>
<destination>
<instance instance-name
<network>
<no-backup-coverage>
<router>
<topology  (default | ipv4-multicast | topology-name)>
<transit>
```

Description

Display the entries in the Open Shortest Path First (OSPF) routing table.

Options

none	Display standard information about all entries in the OSPF routing table for all routing instances and all topologies.
destination	Display routes to the specified IP address (with optional destination prefix length).
brief | detail | extensive	(Optional) Display the specified level of output.
abr	(Optional) Display routes to area border routers.
asbr	(Optional) Display routes to autonomous system border routers.
extern	(Optional) Display external routes.
inter	(Optional) Display interarea routes.
intra	(Optional) Display intra-area routes.
instance (default | ipv4-multicast | *instance-name*)	(Optional) Display entries for the default routing instance, the IPv4 multicast routing instance, or for the specified routing instance.
logical-system (default | ipv4-multicast | *logical-system-name*)	(Optional) Perform this operation on the default logical system, the IPv4 multicast logical system, or on a particular logical system.
network	(Optional) Display routes to networks.
no-backup-coverage	(Optional) Display routes with no backup coverage.
realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)	(OSPFv3 only) (Optional) Display entries in the routing table for the specified OSPFv3 realm, or address family. Use the `realm` option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
router	(Optional) Display routes to all routers.
topology (default | ipv4-multicast | *topology-name*)	(OSPFv2 only) (Optional) Display routes for the default OSPF topology, IPv4 multicast topology, or for a particular topology.
transit	(Optional) (OSPFv3 only) Display OSPFv3 routes to pseudonodes.

Required Privilege Level

view

Output Fields

Table 948 on page 22854 list the output fields for the `show (ospf | ospf3) route` command. Output fields are listed in the approximate order in which they appear.

Table 948: show (ospf | ospf3) route Output Fields

Field Name	Field Description	Output Level
Topology	Name of the topology.	All levels
Prefix	Destination of the route.	All levels
Path type	How the route was learned: • Inter—Interarea route • Ext1—External type 1 route • Ext2—External type 2 route • Intra—Intra-area route	All levels
Route type	The type of routing device from which the route was learned: • AS BR—Route to AS border router. • Area BR—Route to area border router. • Area/AS BR—Route to router that is both an Area BR and AS BR. • Network—Network router. • Router—Route to a router that is neither an Area BR nor an AS BR. • Transit—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link. • Discard—Route to a summary discard.	All levels

Table 948: show (ospf | ospf3) route Output Fields *(Continued)*

Field Name	Field Description	Output Level
NH Type	Next-hop type: LSP or IP.	All levels
Metric	Route's metric value.	All levels
NH-interface	(OSPFv3 only) Interface through which the route's next hop is reachable.	All levels
NH-addr	(OSPFv3 only) IPv6 address of the next hop.	All levels
NextHop Interface	(OSPFv2 only) Interface through which the route's next hop is reachable.	All levels
Nexthop addr/label	(OSPFv2 only) If the NH Type is IP, then it is the address of the next hop. If the NH Type is LSP, then it is the name of the label-switched path.	All levels
Area	Area ID of the route.	`detail`
Origin	Router from which the route was learned.	`detail`
Type 7	Route was learned through a not-so-stubby area (NSSA) link-state advertisement (LSA).	`detail`
P-bit	Route was learned through NSSA LSA and the propagate bit was set.	`detail`
Fwd NZ	Forwarding address is nonzero. Fwd NZ is only displayed if the route is learned through an NSSA LSA.	`detail`

Table 948: show (ospf | ospf3) route Output Fields *(Continued)*

Field Name	Field Description	Output Level
optional-capability	Optional capabilities propagated in the router LSA. This field is in the output for intra-area router routes only (when Route Type is Area BR, AS BR, Area/AS BR, or Router), not for interarea router routes or network routes. Three bits in this field are defined as follows: • 0x4 (V)—Routing device is at the end of a virtual active link. • 0x2 (E)—Routing device is an autonomous system boundary router. • 0x1 (B)—Routing device is an area border router.	`detail`
priority	The priority assigned to the prefix: • high • medium • low **NOTE**: The `priority` field applies only to routes of type Network.	`detail`
BGP-ORR Generation-ID	Display the BGP-ORR generation identifier of the main OSPF route. This field is shown only for non-zero values.	`extensive`

Sample Output

show ospf route

```
user@host> show ospf route
Topology default Route Table:

Prefix             Path  Route      NH       Metric NextHop       Nexthop
                   Type  Type       Type            Interface     Address/LSP
10.1.1.60/32        Intra Network    Spring        6 lt-1/2/0.14   10.0.10.70
                                    Bkup SPRING     lt-1/2/0.12   10.0.6.60
10.1.1.70/32        Intra Network    IP            1 lt-1/2/0.14   10.0.10.70
                                    Bkup LSP                      (null)
10.1.1.70/32        Intra Network    Spring        1 lt-1/2/0.14   10.0.10.70
```

```
                                    Bkup SPRING     lt-1/2/0.12   10.0.6.60
10.1.1.80/32        Intra Network    IP            6 lt-1/2/0.14   10.0.10.70
                                    Bkup IP         lt-1/2/0.12   10.0.6.60
10.1.1.80/32        Intra Network    Spring        6 lt-1/2/0.14   10.0.10.70

802068 (S=0)       Intra Network    Mpls          0 lt-1/2/0.14   10.0.10.70
                                    Bkup MPLS       lt-1/2/0.12   10.0.6.60
802078 (S=0)       Intra Network    Mpls          0 lt-1/2/0.14   10.0.10.70
                                    Bkup MPLS       lt-1/2/0.12   10.0.6.60
802088 (S=0)       Intra Network    Mpls          0 lt-1/2/0.14   10.0.10.70
                                    Bkup MPLS       lt-1/2/0.12   10.0.6.60
802098 (S=0)       Intra Network    Mpls          0 lt-1/2/0.14   10.0.10.70
                                    Bkup MPLS       lt-1/2/0.12   10.0.6.60
802108 (S=0)       Intra Network    Mpls          0 lt-1/2/0.14   10.0.10.70
                                    Bkup MPLS       lt-1/2/0.12   10.0.6.60
802118 (S=0)       Intra Network    Mpls          0 lt-1/2/0.14   10.0.10.70
                                    Bkup MPLS       lt-1/2/0.12   10.0.6.60
802118             Intra Network    Mpls          0 lt-1/2/0.14   10.0.10.70
                                    Bkup MPLS       lt-1/2/0.12   10.0.6.60
```

show ospf route extensive

```
user@host> show ospf route extensive
Topology default Route Table:

Prefix             Path  Route      NH       Metric NextHop       Nexthop
                   Type  Type       Type            Interface     Address/LSP
10.1.1.1            Intra Router     IP          100 ge-0/0/2.0    10.1.1.1
  area 0.0.0.0, origin 10.1.1.1, optional-capability 0x0
10.1.1.1/32         Intra Network    IP          100 ge-0/0/2.0    10.1.1.1
  area 0.0.0.0, origin 10.1.1.1, priority medium
BGP-ORR generation-id: 1
```

show ospf3 route

```
user@host> show ospf3 route
Prefix               Path   Route       NH   Metric  NextHop       Nexthop
                     Type   Type        Type         Interface     addr/label
10.255.71.13         Intra  Router      IP        1
  NH-interface fe-0/0/2.0, NH-addr fe80::290:69ff:fe9b:e002
```

```
10.255.71.13;0.0.0.2
10.255.245.1         Intra  Router      IP       40  fxp1.1        192.168.36.17
  area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3         Intra  AS BR       IP        1  fxp2.3        192.168.36.34
  area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32      Intra  Network     IP       40  fxp1.1        192.168.36.17
  area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32      Intra  Network     IP        0  lo0.0
  area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32      Intra  Network     IP        1  fxp2.3        192.168.36.34
  area 0.0.0.0, origin 10245.3, priority low
                     Intra  Transit     IP        1
  NH-interface fe-0/0/2.0
192::168:222:84/126  Intra  Network     IP        1
  NH-interface fe-0/0/2.0
abcd::71:12/128      Intra  Network     IP        0
  NH-interface lo0.0
abcd::71:13/128      Intra  Network     LSP       1
  NH-interface fe-0/0/2.0, NH-addr lsp-cd
```

show ospf route topology voice

```
user@host show ospf route topology voice
Topology voice Route Table:
Prefix             Path   Route       NH   Metric  NextHop       Nexthop
                   Type   Type        Type         Interface     addr/label
10.255.8.2         Intra  Router      IP        1  so-0/2/0.0
10.255.8.3         Intra  Router      IP        2  so-0/2/0.0
10.255.8.1/32      Intra  Network     IP        0  lo0.0
10.255.8.2/32      Intra  Network     IP        1  so-0/2/0.0
10.255.8.3/32      Intra  Network     IP        2  so-0/2/0.0
192.168.8.0/29     Intra  Network     IP        2  so-0/2/0.0
192.168.8.44/30    Intra  Network     IP        2  so-0/2/0.0
192.168.8.46/32    Intra  Network     IP        1  so-0/2/0.0
192.168.8.48/30    Intra  Network     IP        1  so-0/2/1.0
192.168.8.52/30    Intra  Network     IP        2  so-0/2/0.0
192.168.9.44/30    Intra  Network     IP        1  so-0/2/0.0
192.168.9.45/32    Intra  Network     IP        2  so-0/2/0.0
```

Release Information

Command introduced before Junos OS Release 7.4.

`topology` option introduced in Junos OS Release 9.0.

`realm` option introduced in Junos OS Release 9.2.

show (ospf | ospf3) statistics

IN THIS SECTION

- Syntax | 22859
- Syntax (EX Series Switch and QFX Series) | 22859
- Description | 22860
- Options | 22860
- Required Privilege Level | 22860
- Output Fields | 22860
- Sample Output | 22862
- Release Information | 22863

Syntax

```
show (ospf | ospf3) statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
```

Syntax (EX Series Switch and QFX Series)

```
show (ospf | ospf3) statistics
<instance instance-name>
```

Description

Display OSPF statistics.

Options

none	Display OSPF statistics for all routing instances.
instance *instance-name*	(Optional) Display all statistics for the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)	(Optional) (OSPFv3 only) Display all statistics for the specified OSPFv3 realm, or address family. Use the `realm` option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

Output Fields

Table 949 on page 22860 lists the output fields for the `show (ospf | ospf3) statistics` command. Output fields are listed in the approximate order in which they appear.

Table 949: show (ospf | ospf3) statistics Output Fields

Field Name	Field Description
Packet type	Type of OSPF packet.
Total Sent/Total Received	Total number of packets sent and received.
Last 5 seconds Sent/Last 5 seconds Received	Total number of packets sent and received in the last 5 seconds.
DBDs retransmitted	Total number of database description packets retransmitted, and number retransmitted in the last 5 seconds.

Table 949: show (ospf | ospf3) statistics Output Fields *(Continued)*

Field Name	Field Description
LSAs flooded	Total number of link-state advertisements flooded, and number flooded in the last 5 seconds.
LSAs flooded high-prio	Total number of high priority link-state advertisements flooded, and number flooded in the last 5 seconds. A link-state advertisement is deemed a high priority if it has changed since it was last sent.
LSAs retransmitted	Total number of link-state advertisements retransmitted, and number retransmitted in the last 5 seconds.
LSAs transmitted to nbr	Total number of link-state advertisements transmitted to a neighbor, and number transmitted in the last 5 seconds.
LSAs requested	Total number of link-state advertisements requested by neighboring devices, and number requested in the last 5 seconds.
LSAs acknowledged	Total number of link-state advertisements acknowledged, and number acknowledged in the last 5 seconds.
Flood queue depth	Total number of entries in the extended queue.
Total rexmit entries	Total number of retransmission entries waiting to be sent from the OSPF routing instance.
db summaries	Total number of database description summaries waiting to be sent from the OSPF routing instance.
lsreq entries	Total number of link-state request entries waiting to be sent from the OSPF routing instance.

Table 949: show (ospf | ospf3) statistics Output Fields *(Continued)*

Field Name	Field Description
Receive errors	Number and type of receive errors. Some sample receive errors include: • mtu mismatches • no interface found • no virtual link found • nssa mismatches • stub area mismatches • subnet mismatches If there are no receive errors, the output displays `none`.

Sample Output

show ospf statistics

```
user@host> show ospf statistics
Packet type             Total                  Last 5 seconds
                   Sent      Received        Sent      Received
   Hello             31            14           2             2
     DbD              9            10           0             0
   LSReq              2             2           0             0
LSUpdate              8            16           0             0
   LSAck              9             9           0             0

DBDs retransmitted     :                    3, last 5 seconds :          0
LSAs flooded           :                   12, last 5 seconds :          0
LSAs flooded high-prio :                    0, last 5 seconds :          0
LSAs retransmitted     :                    0, last 5 seconds :          0
LSAs transmitted to nbr:                    3, last 5 seconds :          0
LSAs requested         :                    5, last 5 seconds :          0
LSAs acknowledged      :                   19, last 5 seconds :          0

Flood queue depth      :               0
```

```
Total rexmit entries   :               0
db summaries           :               0
lsreq entries          :               0

Receive errors:
  862 no interface found
  115923 no virtual link found
```

show ospf3 statistics

```
user@host> show ospf3 statistics
Packet type             Total                  Last 5 seconds
                   Sent      Received        Sent      Received
   Hello              0             0           0             0
     DbD              0             0           0             0
   LSReq              0             0           0             0
LSUpdate              0             0           0             0
   LSAck              0             0           0             0

DBDs retransmitted     :                    0, last 5 seconds :          0
LSAs flooded           :                    0, last 5 seconds :          0
LSAs flooded high-prio :                    0, last 5 seconds :          0
LSAs retransmitted     :                    0, last 5 seconds :          0
LSAs transmitted to nbr:                    0, last 5 seconds :          0
LSAs requested         :                    0, last 5 seconds :          0
LSAs acknowledged      :                    0, last 5 seconds :          0

Flood queue depth      :               0
Total rexmit entries   :               0
db summaries           :               0
lsreq entries          :               0

Receive errors:
  None
```

Release Information

Command introduced before Junos OS Release 7.4.

`realm` option introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

clear (ospf | ospf3) statistics | 3795

show ospf spring prefix sid map

IN THIS SECTION

- Syntax | 22864
- Description | 22864
- Options | 22864
- Additional Information | 22865
- Required Privilege Level | 22865
- Output Fields | 22865
- Sample Output | 22866
- Release Information | 22866

Syntax

```
show ospf spring prefix-sid-map
```

Description

Displays segment routing mapping server (SRMS) advertisements.

Options

none Displays the SRMS advertisements and status of the mapping entry.

Additional Information

The output displays all the prefixes and SIDs. If in case two prefixes have the same SID or single prefix advertised by two multiple mapping servers have different SIDs, it also displays the mapping entry is active or not.

Required Privilege Level

view

Output Fields

Table 1 describes the output fields for the `show ospf spring prefix-sid-map` command.

Table 950: show ospf spring prefix-sid-map Output Fields

Field Name	Field Description	Level of Output
SID	segment identifier	All levels
Prefix	Destination of the route.	All levels
Mapping server	Name of the mapping server.	All levels
Range	Number of prefixes covered by the advertisement.	All levels
Active	Status of mapping entry.	All levels

Sample Output

show ospf spring prefix-sid-map

```
user@host> show ospf spring prefix-sid-map
```

```
OSPF database, Area 0.0.0.0

SID     Prefix            Mapping-server   Range     Active
50      5.5.5.5/32        5.5.5.5          1         Yes
71      7.7.7.7/32        7.7.7.7          1         Yes
```

Release Information

Command introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/ref/command/show-ospf-database.html

show ospf spring sid database

IN THIS SECTION

- Syntax | 22867
- Description | 22867
- Options | 22867
- Additional Information | 22867
- Required Privilege Level | 22867
- Output Fields | 22867
- Sample Output | 22868

Release Information | 22869

Syntax

```
show ospf spring sid-database
```

Description

Displays the segment identifier (SID) database with the prefix and index of native segment routing nodes.

Options

none Displays the SID database.

Additional Information

The output displays all the prefixes and all segment identifiers (SIDs). If in any case two prefixes have the same SID or single prefix advertised by two different nodes has different SIDs, it will also display the route-type of the route.

Required Privilege Level

view

Output Fields

Table 1 describes the output fields for the `show ospf spring sid-database` command. Output fields are listed in the approximate order in which they appear.

Table 951: show ospf spring sid-database Output Fields

Field Name	Field Description	Level of Output
SID	segment identifier	All levels

Table 951: show ospf spring sid-database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Prefix	Destination of the route.	All levels
Advertised-by	Names of the nodes advertising the prefixes.	All levels
Route-type	The type of routing device from which the route was learned: • AS BR—Route to AS border router. • Area BR—Route to area border router. • Area/AS BR—Route to router that is both an Area BR and AS BR. • Network—Network router. • Router—Route to a router that is neither an Area BR nor an AS BR. • Transit—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link. • Discard—Route to a summary discard.	All levels

Sample Output

show ospf spring sid-database

```
user@host> show ospf spring sid-database
```

```
OSPF database, Area 0.0.0.0

SID     Prefix            Advertised-by    Route-type
10      1.1.1.1/32        1.1.1.1          Intra-Area
20      2.2.2.2/32        2.2.2.2          Intra-Area
```

```
30      3.3.3.3/32         3.3.3.3          Intra-Area
40      4.4.4.4/32         4.4.4.4          Intra-Area
50      5.5.5.5/32         5.5.5.5          Intra-Area
100     100.100.100.100/32 100.100.100.100  Intra-Area
```

Release Information

Command introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/ref/command/show-ospf-database.html

show ospf3 database

IN THIS SECTION

- Syntax | 22869
- Syntax (EX Series Switches and QFX Series) | 22870
- Description | 22870
- Options | 22871
- Required Privilege Level | 22872
- Output Fields | 22872
- Sample Output | 22879
- Release Information | 22880

Syntax

```
show ospf3 database
<brief | detail | extensive | summary>
<advertising-router (address | self)>
<area area-id>
```

```
<external>
<instance instance-name>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link>
<link-local>
<logical-system (all | logical-system-name)>
<lsa-id lsa-id>
<network>
<nssa>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>
```

Syntax (EX Series Switches and QFX Series)

```
show ospf3 database
<brief | detail | extensive | summary>
<advertising-router (address | self)>
<area area-id>
<external>
<instance instance-name>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link>
<link-local>
<lsa-id lsa-id>
<network>
<nssa>
<router>
```

Description

Display the entries in the OSPF version 3 (OSPFv3) link-state database, which contains data about link-state advertisement (LSA) packets.

Options

none	Display standard information about all entries in the OSPFv3 link-state database.
brief | detail | extensive | summary	(Optional) Display the specified level of output.
advertising-router (*address* | self)	(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.
area *area-id*	(Optional) Display the LSAs in a particular area.
external	(Optional) Display external LSAs.
instance *instance-name*	(Optional) Display all OSPF database information under the named routing instance.
inter-area-prefix	(Optional) Display information about interarea-prefix LSAs.
inter-area-router	(Optional) Display information about interarea-router LSAs.
intra-area-prefix	(Optional) Display information about intra-area-prefix LSAs.
link	(Optional) Display information about link LSAs.
link-local	(Optional) Display information about link-local LSAs.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
lsa-id *lsa-id*	(Optional) Display the LSA with the specified LSA identifier.
network	(Optional) Display information about network LSAs.
nssa	(Optional) Display information about not-so-stubby area (NSSA) LSAs.
realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)	(Optional) Display information about the specified OSPFv3 realm, or address family. Use the `realm` option to specify an address family other than IPv6 unicast, which is the default.
router	(Optional) Display information about router LSAs.

Required Privilege Level

view

Output Fields

Table 952 on page 22872 lists the output fields for the `show ospf3 database` command. Output fields are listed in the approximate order in which they appear.

Table 952: show ospf3 database Output Fields

Field Name	Field Description	Level of Output
OSPF link state database, area *area-number*	Entries in the link-state database for this area.	`brief detail extensive`
OSPF AS SCOPE link state database	Entries in the AS scope link-state database.	`brief detail extensive`
OSPF Link-Local link state database, interface *interface-name*	Entries in the link-local link-state database for this interface.	`brief detail extensive`
area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Type	Type of link advertisement: Extern, InterArPfx, InterArRtr, IntraArPrx , Link, Network, NSSA, or Router.	`brief detail extensive`
ID	Link identifier included in the advertisement. An asterisk (*) preceding the identifier marks database entries that originated from the local routing device.	`brief detail extensive`
Adv Rtr	Address of the routing device that sent the advertisement.	`brief detail extensive`
Seq	Link sequence number of the advertisement.	`brief detail extensive`

Table 952: show ospf3 database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Age	Time elapsed since the LSA was originated, in seconds.	`brief detail extensive`
Cksum	Checksum value of the LSA.	`brief detail extensive`
Len	Length of the advertisement, in bytes.	`brief detail extensive`
Router (Router Link-State Advertisements)		
bits	Flags describing the routing device that generated the LSP.	`detail extensive`
Options	Option bits carried in the router LSA.	`detail extensive`
For Each Router Link		
Type	Type of interface. The value of all other output fields describing a routing device interface depends on the interface's type: • PointToPoint (1)—Point-to-point connection to another routing device. • Transit (2)—Connection to a transit network. • Virtual (4)—Virtual link.	`detail extensive`
Loc-if-id	Local interface ID assigned to the interface that uniquely identifies the interface with the routing device.	`detail extensive`
Nbr-if-id	Interface ID of the neighbor's interface for this routing device link.	`detail extensive`
Nbr-rtr-id	Router ID of the neighbor routing device (for type 2 interfaces, the attached link's designated router).	`detail extensive`

Table 952: show ospf3 database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Metric	Cost of the router link.	`detail extensive`
Gen timer	How long until the LSA is regenerated, in the format *hours*:*minutes*:*seconds*.	`extensive`
Aging timer	How long until the LSA expires, in the format *hours*:*minutes*:*seconds*.	`extensive`
Installed *nn*:*nn*:*nn* ago	How long ago the route was installed, in the format *hours*:*minutes*:*seconds*.	`extensive`
expires in *nn*:*nn*:*nn*	How long until the route expires, in the format *hours*:*minutes*:*seconds*.	`extensive`
sent *nn*:*nn*:*nn* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format *hours*:*minutes*:*seconds*.	`extensive`
Ours	Indicates that this is a local advertisement.	`extensive`
Network (Network Link-State Advertisements)		
Options	Option bits carried in the network LSA.	`detail extensive`
Attached Router	Router IDs of each of the routing devices attached to the link. Only routing devices that are fully adjacent to the designated router are listed. The designated router includes itself in this list.	`detail extensive`
InterArPfx (Interarea-Prefix Link-State Advertisements)		
Prefix	IPv6 address prefix.	`detail extensive`
Prefix-options	Option bit associated with the prefix.	`detail extensive`

Table 952: show ospf3 database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	`detail extensive`
Gen timer	How long until the LSA is regenerated, in the format *hours*:*minutes*:*seconds*.	`extensive`
Aging timer	How long until the LSA expires, in the format *hours*:*minutes*:*seconds*.	`extensive`
Installed *nn*:*nn*:*nn* ago	How long ago the route was installed, in the format *hours*:*minutes*:*seconds*.	`extensive`
expires in *nn*:*nn*:*nn*	How long until the route expires, in the format *hours*:*minutes*:*seconds*.	`extensive`
sent *nn*:*nn*:*nn* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format *hours*:*minutes*:*seconds*.	`extensive`
Ours	Indicates that this is a local advertisement.	`extensive`
InterArRtr (Interarea-Router Link-State Advertisements)		
Dest-router-id	Router ID of the routing device described by the LSA.	`detail extensive`
options	Optional capabilities supported by the routing device.	`detail extensive`
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	`detail extensive`
Prefix	IPv6 address prefix.	`extensive`

Table 952: show ospf3 database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Prefix-options	Option bit associated with the prefix.	`extensive`
Extern (External Link-State Advertisements)		
Prefix	IPv6 address prefix.	`detail extensive`
Prefix-options	Option bit associated with the prefix.	`detail extensive`
Metric	Cost of the route, which depends on the value of Type.	`detail extensive`
Type *n*	Type of external metric: Type 1 or Type 2.	`detail extensive`
Aging timer	How long until the LSA expires, in the format *hours*:*minutes*:*seconds*.	`extensive`
Installed *nn:nn:nn* ago	How long ago the route was installed, in the format *hours*:*minutes*:*seconds*.	`extensive`
expires in *nn:nn:nn*	How long until the route expires, in the format *hours*:*minutes*:*seconds*.	`extensive`
sent *nn:nn:nn* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format *hours*:*minutes*:*seconds*.	`extensive`
Link (Link-State Advertisements)		
IPv6-Address	IPv6 link-local address on the link for which this link LSA originated.	`detail extensive`
Options	Option bits carried in the link LSA.	`detail extensive`
priority	Router priority of the interface attaching the originating routing device to the link.	`detail extensive`

Table 952: show ospf3 database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	`detail extensive`
Prefix	IPv6 address prefix.	`detail extensive`
Prefix-options	Option bit associated with the prefix.	`detail extensive`
Gen timer	How long until the LSA is regenerated, in the format *hours*:*minutes*:*seconds*.	`extensive`
Aging timer	How long until the LSA expires, in the format *hours*:*minutes*:*seconds*.	`extensive`
Installed *nn:nn:nn* ago	How long ago the route was installed, in the format *hours*:*minutes*:*seconds*.	`extensive`
expires in *nn:nn:nn*	How long until the route expires, in the format *hours*:*minutes*:*seconds*.	`extensive`
sent *nn:nn:nn* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format *hours*:*minutes*:*seconds*.	`extensive`
Ours	Indicates that this is a local advertisement.	`extensive`
IntraArPfx (Intra-Area-Prefix Link-State Advertisements)		
Ref-lsa-type	LSA type of the referenced LSA. • Router—Address prefixes are associated with a router LSA. • Network—Address prefixes are associated with a network LSA.	`detail extensive`
Ref-lsa-id	Link-state ID of the referenced LSA.	`detail extensive`

Table 952: show ospf3 database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Ref-router-id	Advertising router ID of the referenced LSA.	`detail extensive`
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	`detail extensive`
Prefix	IPv6 address prefix.	`detail extensive`
Prefix-options	Option bit associated with the prefix.	`detail extensive`
Metric	Cost of this prefix. Expressed in the same units as the interface costs in the router LSAs.	`detail extensive`
Gen timer	How long until the LSA is regenerated, in the format *hours*:*minutes*:*seconds*.	`extensive`
Aging timer	How long until the LSA expires, in the format *hours*:*minutes*:*seconds*.	`extensive`
Installed *hh*:*mm*:*ss* ago	How long ago the route was installed, in the format *hours*:*minutes*:*seconds*.	`extensive`
expires in *hh*:*mm*:*ss*	How long until the route expires, in the format *hours*:*minutes*:*seconds*.	`extensive`
sent *hh*:*mm*:*ss* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format *hours*:*minutes*:*seconds*.	`extensive`
n Router LSAs	Number of router LSAs in the link-state database.	`summary`
n Network LSAs	Number of network LSAs in the link-state database.	`summary`
n InterArPfx LSAs	Number of interarea-prefix LSAs in the link-state database.	`summary`

Table 952: show ospf3 database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
n InterArRtr LSAs	Number of interarea-router LSAs in the link-state database.	summary
n IntraArPfx LSAs	Number of intra-area-prefix LSAs in the link-state database.	summary
Externals	Display of the external LSA database.	summary
n Extern LSAs	Number of external LSAs in the link-state database.	summary
Interface *interface-name*	Name of the interface for which link-local LSA information is displayed.	summary
n Link LSAs	Number of link LSAs in the link-state database.	summary

Sample Output

show ospf3 database brief

```
user@host> show ospf3 database brief
    OSPF3 link state database, area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router      0.0.0.1          10.255.4.85      0x80000003   885  0xa697  40
Router     *0.0.0.1          10.255.4.93      0x80000002   953  0xc677  40
InterArPfx *0.0.0.2          10.255.4.93      0x80000001   910  0xb96f  44
InterArRtr *0.0.0.1          10.255.4.93      0x80000001   910  0xe159  32
IntraArPfx *0.0.0.1          10.255.4.93      0x80000002   432  0x788f  72

    OSPF3 link state database, area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router     *0.0.0.1          10.255.4.93      0x80000003   916  0xea40  40
Router      0.0.0.1          10.255.4.97      0x80000006   851  0xc95b  40
Network     0.0.0.2          10.255.4.97      0x80000002   916  0x4598  32
InterArPfx *0.0.0.1          10.255.4.93      0x80000002   117  0xa980  44
InterArPfx *0.0.0.2          10.255.4.93      0x80000002    62  0xd47e  44
NSSA        0.0.0.1          10.255.4.97      0x80000002   362  0x45ee  44
```

```
IntraArPfx  0.0.0.1          10.255.4.97      0x80000006   851  0x2f77  52

    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern      0.0.0.1          10.255.4.85      0x80000002    63  0x9b86  44
Extern     *0.0.0.1          10.255.4.93      0x80000001   910  0x59c9  44

    OSPF3 Link-Local link state database, interface ge-1/3/0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.2          10.255.4.93      0x80000003   916  0x4dab  64
```

show ospf3 database summary

```
user@host> show ospf3 database summary
Area 0.0.0.0:
   2 Router LSAs
   1 InterArPfx LSAs
   1 InterArRtr LSAs
   1 IntraArPfx LSAs
Area 0.0.0.1:
   2 Router LSAs
   1 Network LSAs
   2 InterArPfx LSAs
   1 NSSA LSAs
   1 IntraArPfx LSAs
Externals:
   2 Extern LSAs
Interface ge-1/3/0.0:
   1 Link LSAs
Interface lo0.0:
Interface so-2/2/0.0:
   1 Link LSAs
```

Release Information

Command introduced before Junos OS Release 7.4.

`realm` option introduced in Junos OS Release 9.2.

`advertising-router` (*address* | `self`) option introduced in Junos Relase 9.5.

`advertising-router` (*address* | `self`) option introduced in Junos OS Release 9.5 for EX Series switches.

RELATED DOCUMENTATION

clear (ospf | ospf3) database | 3784

show ovsdb commit failures

IN THIS SECTION

- Syntax | 22881
- Description | 22881
- Options | 22882
- Required Privilege Level | 22882
- Output Fields | 22882
- Sample Output | 22883
- Release Information | 22884

Syntax

```
show ovsdb commit failures
<transaction-id>
```

Description

Display configurations of *Open vSwitch Database (OVSDB)*-managed Virtual Extensible LANs (VXLANs) and associated logical interfaces that the Juniper Networks switch dynamically configured but was unable to commit.

For each OVSDB-managed VXLAN and associated logical interface that you plan to implement in a Junos OS environment, you must configure equivalent entities in NSX Manager or in the NSX API for an NSX environment, or in the Contrail Web user interface for a Contrail environment. The software-defined networking (SDN) controller pushes these configurations to the connected Juniper Networks

switch by way of the OVSDB schema for physical devices. After the Juniper Networks switch receives these configurations, it dynamically configures a Junos OS-equivalent VXLAN and associated logical interface, and attempts to commit the configurations.

During the commitment of the dynamic configurations, If there is an issue with one or more of the configurations, all configurations in the transaction, even the ones that are correctly configured, remain uncommitted and are saved in a queue. All configurations in the transaction remain uncommitted and in the queue until you troubleshoot and resolve the configuration issues. After you resolve the configuration issues, you must use the `clear ovsdb commit failures` command to remove the transaction from the queue and retry committing the configurations.

> **NOTE**: While an erroneous transaction exists in the queue, the Juniper Networks switch cannot commit the dynamic configurations of additional VXLANs and their associated logical interfaces. The commitment of these VXLANs and logical interfaces remain in a pending state until all VXLAN and logical interface configurations in the erroneous transaction are resolved and successfully committed.

Issues that can cause commitment errors include but are not limited to the detection of the same VXLAN name or VXLAN network identifier (VNI) in a dynamically configured VXLAN and in a VXLAN that was previously configured using the Junos OS CLI.

To monitor for issues with dynamically configured OVSDB-managed VXLANs and their associated interfaces, we recommend checking for system log messages and traceoptions files for OVSDB.

Options

none	Display information about an erroneous transaction.
transaction-id	Display information about the transaction with the specified numerical ID.

Required Privilege Level

admin

Output Fields

Table 953 on page 22883 lists the output fields for the `show ovsdb commit failures` command. Output fields are listed in the approximate order in which they appear.

Table 953: show ovsdb commit failures Output Fields

Field Name	Field Description
Txn ID	ID assigned to a transaction by the Juniper Networks switch.
Logical-switch	Name of the VXLAN that the Juniper Networks switch dynamically configured but was unable to commit the configuration of.
Port	Name of an OVSDB-managed physical interface that is associated with the VXLAN.
VLAN ID	ID that is assigned to the VXLAN.

Sample Output

show ovsdb commit failures

```
user@host> show ovsdb commit failures
Txn ID      Logical-switch                          Port             VLAN ID
1           28805c1d-0122-495d-85df-19abd647d772
1                                                   xe-0/0/5:0       1016
1           9acc24b3-7b0a-4c2e-b572-3370c3e1acff
1                                                   xe-0/0/5:0       1017
...
```

show ovsdb commit failure (Specific Transaction)

```
user@host> how ovsdb commit failures 1
Txn ID      Logical-switch                          Port             VLAN ID
1           28805c1d-0122-495d-85df-19abd647d772
1                                                   xe-0/0/5:0       1016
1           9acc24b3-7b0a-4c2e-b572-3370c3e1acff
1                                                   xe-0/0/5:0       1017
...
```

Release Information

Command introduced in Junos OS Release 14.1X53-D26.

RELATED DOCUMENTATION

Understanding Dynamically Configured VXLANs in an OVSDB Environment

traceoptions (OVSDB) | 29630

show ovsdb controller

IN THIS SECTION

- Syntax | 22884
- Description | 22884
- Options | 22885
- Required Privilege Level | 22885
- Output Fields | 22885
- Sample Output | 22886
- Release Information | 22887

Syntax

```
show ovsdb controller
<address ip-address>
```

Description

Display information and connection status for software-defined networking (SDN) controllers to which the Juniper Networks device is connected.

Options

none	Display information about all SDN controllers to which the Juniper Networks device is connected.
address ***ip-address***	Display information about the SDN controller at the specified IP address.

Required Privilege Level

admin

Output Fields

Table 954 on page 22885 lists the output fields for the `show ovsdb controller` command. Output fields are listed in the approximate order in which they appear.

Table 954: show ovsdb controller Output Fields

Field Name	Field Description
Controller IP address	IP address of the SDN controller to which the Juniper Networks device is connected.
Controller protocol	Protocol used by the Juniper Networks device to initiate the connection.
Controller port	Port to which the Juniper Networks device is connected.
Controller connection	State of the connection with the SDN controller.
Controller seconds-since-connect	Number of seconds since the connection with the SDN controller was established.
Controller seconds-since-disconnect	Number of seconds since the connection with the SDN controller was dropped.
Controller connection status	Status of the connection with the SDN controller.

Sample Output

show ovsdb controller

```
user@host> show ovsdb controller
VTEP controller information:
Controller IP address: 10.168.66.189
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56290
Controller seconds-since-disconnect: 0
Controller connection status: active

Controller IP address: 10.168.181.54
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56292
Controller seconds-since-disconnect: 0
Controller connection status: active

Controller IP address: 10.168.182.45
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56292
Controller seconds-since-disconnect: 0
Controller connection status: active
```

show ovsdb controller address

```
user@host> show ovsdb controller address 10.168.182.45
VTEP controller information:
Controller IP address: 10.168.182.45
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56347
```

```
Controller seconds-since-disconnect: 0
Controller connection status: active
```

Release Information

Command introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Setting Up the OVSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs

Setting Up OVSDB on Juniper Networks Devices That Support the Dynamic Configuration of VXLANs

Understanding How to Set Up OVSDB Connections on a Juniper Networks Device

show ovsdb interface

IN THIS SECTION

- Syntax | 22887
- Description | 22888
- Options | 22888
- Required Privilege Level | 22888
- Output Fields | 22888
- Sample Output | 22889
- Release Information | 22889

Syntax

```
show ovsdb interface
<interface-name>
```

Description

Display information about *Open vSwitch Database (OVSDB)*-managed interfaces configured by using the `interfaces` *`interface-name`* statement in the [`edit protocols ovsdb`] hierarchy.

Options

none Display information about all OVSDB-managed interfaces.

interface-name Display information about the specified OVSDB-managed interface.

Required Privilege Level

admin

Output Fields

Table 955 on page 22888 lists the output fields for the `show ovsdb interface` command. Output fields are listed in the approximate order in which they appear.

Table 955: show ovsdb interface Output Fields

Field Name	Field Description
Interface	Name of interface.
VLAN ID	ID of Virtual Extensible LAN (VXLAN) with which the interface is associated. **NOTE**: This field is not supported by MX Series routers or EX9200 switches.
Bridge domain or VLAN	Bridge domain or VLAN under which the VXLAN is created. **NOTE**: This field is not supported by MX Series routers or EX9200 switches.

Sample Output

show ovsdb interface

```
user@host> show ovsdb interface
Interface              VLAN ID         Bridge-domain
ge-7/0/9.0
ge-7/0/9.1
irb.11
irb.12
irb.2
irb.3
xe-10/3/0.0
xe-10/3/0.1
```

show ovsdb (Specific Interface)

```
user@host> show ovsdb interface ge-7/0/9.0
Interface              VLAN ID         Bridge-domain
ge-7/0/9.0
```

Release Information

Command introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Setting Up the OVSDB Protocol on Juniper Networks Devices that Support Manual Configuration of VXLANs

Setting Up OVSDB on Juniper Networks Devices That Support the Dynamic Configuration of VXLANs

show ovsdb logical-switch

IN THIS SECTION

- Syntax | 22890
- Description | 22890
- Options | 22891
- Required Privilege Level | 22891
- Output Fields | 22891
- Sample Output | 22892
- Release Information | 22893

Syntax

```
show ovsdb logical-switch
<logical-switch-name>
```

Description

NOTE: In the Open vSwitch Database (OVSDB) schema for physical devices, the logical switch table stores information about the Layer 2 broadcast domain that you configured in a VMware NSX or Contrail environment. In the NSX environment, the Layer 2 broadcast domain is known as a *logical switch*, while in the Contrail environment, the domain is known as a *virtual network*.

In the context of the `show ovsdb logical-switch` command, the term *logical switch* refers to the logical switch or virtual network that was configured in the NSX or Contrail environments, respectively, and the corresponding configuration that was pushed to the OVSDB schema.

Display information about logical switches and the corresponding Virtual Extensible LANs (VXLANs), which were configured on the Juniper Networks device.

In the command output, each logical switch is identified by a universally unique identifier (UUID), which in the context of this command, is also known as a logical switch name.

The `show ovsdb logical-switch` command displays the state of the logical switch (`Flags`), which can be one of the following:

`Created by Controller`	A logical switch is configured. However, a corresponding VXLAN is not yet configured. In this state, the logical switch and corresponding VXLAN are not yet operational.
`Created by L2ALD`	A VXLAN is configured. However, a corresponding logical switch is not yet configured. In this state, the logical switch and corresponding VXLAN are not yet operational.
`Created by both`	A logical switch and a corresponding VXLAN are configured. In this state, the logical switch and corresponding VXLAN are operational.
`Tunnel key mismatch`	The VNIs specified in the logical switch and corresponding VXLAN configurations do not match. In this state, the logical switch and corresponding VXLAN are not yet operational.

Options

none	Display information about all logical switches that are present in the OVSDB schema for physical devices.
logical-switch-name	Display information about the specified logical switch.

Required Privilege Level

admin

Output Fields

Table 956 on page 22892 lists the output fields for the `show ovsdb logical-switch` command. Output fields are listed in the approximate order in which they appear.

Table 956: show ovsdb logical-switch Output Fields

Field Name	Field Description
Logical Switch Name	UUID that is automatically generated and assigned to the logical switch. When you configure the corresponding VXLAN in the Junos OS CLI, you must specify the same UUID as the VXLAN name.
Flags	State of the logical switch. For possible states, see the Description section of this topic.
VNI	VNI that is configured for the logical switch and corresponding VXLAN.
Num of Remote MAC	The total number of remote media access control (MAC) addresses associated with the logical switch. These addresses are learned by software and hardware *virtual tunnel endpoints (VTEPs)*.
Num of Local MAC	The total number of local MAC addresses associated with the logical switch. *Local MAC addresses* are addresses learned on the local physical ports.

Sample Output

show ovsdb logical-switch

```
user@host> show ovsdb logical-switch
Logical switch information:
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
Flags: Created by both
VNI: 3
Num of Remote MAC: 13
Num of Local MAC: 12
Logical Switch Name: 9b4f880e-dac8-4612-a832-97ad9dec270f
Flags: Created by Controller
VNI: 50
Num of Remote MAC: 0
Num of Local MAC: 0
Logical Switch Name: bc0da2da-6c16-44bf-b655-442484294ded
Flags: Created by Controller
VNI: 51
```

```
Num of Remote MAC: 0
Num of Local MAC: 0
```

show ovsdb logical-switch (Specific Logical Switch)

```
user@host> show ovsdb logical-switch 24a76aff-7e61-4520-a78d-3eca26ad7510
Logical switch information:
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
Flags: Created by both
VNI: 3
Num of Remote MAC: 13
Num of Local MAC: 12
```

Release Information

Command introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

OVSDB Schema for Physical Devices

Troubleshooting a Nonoperational Logical Switch and Corresponding Junos OS OVSDB-Managed VXLAN

show ovsdb mac

IN THIS SECTION

- Syntax | 22894
- Description | 22894
- Options | 22894
- Required Privilege Level | 22895
- Output Fields | 22895
- Sample Output | 22896

Release Information | 22899

Syntax

```
show ovsdb mac
<address mac-address>
<local>
<logical-switch logical-switch-uuid>
<multicast>
<remote>
<unicast>
```

Description

Display media access control (MAC) addresses, as well as information about the MAC addresses, learned by a Juniper Networks device that functions as a hardware *virtual tunnel endpoint (VTEP)*. Using the *Open vSwitch Database (OVSDB)* management protocol, this hardware VTEP can learn about MAC addresses directly or from other software or hardware VTEPs. The MAC addresses learned directly by the hardware VTEP are known as *local addresses*, while the addresses learned from other software or hardware VTEPs are known as *remote addresses*.

Options

Use one or more of the following options to display a more specific list of MAC addresses and information about the MAC addresses. For example, to display a list of local unicast MAC addresses, you can issue the `show ovsdb mac local unicast` command.

none Display all MAC addresses, which includes all local, remote, unicast, and multicast addresses associated with all logical switches.

address *mac-address* Display the specified MAC address.

count (All Juniper Networks devices that support OVSDB except EX9200 switches) Display the number of MAC addresses learned by the Juniper Networks device. Using this option alone, the number includes all local, remote, unicast, and multicast MAC addresses associated with all logical switches in the logical switch table of the OVSDB schema for physical devices. You can use this option with one or more of the other

options to display a more specific count of MAC addresses. For example, to display the number of local and remote unicast MAC addresses, you can issue the `show ovsdb mac count local remote unicast` command.

local Display all local MAC addresses.

logical-switch *logical-switch-uuid* Display all MAC addresses associated with the specified logical switch in the logical switch table of the OVSDB schema for physical devices.

multicast Display all multicast MAC addresses.

remote Display all remote MAC addresses.

unicast Display all unicast MAC addresses.

Required Privilege Level

admin

Output Fields

Table 957 on page 22895 lists the output fields for the `show ovsdb mac` command. Output fields are listed in the approximate order in which they appear.

Table 957: show ovsdb mac Output Fields

Field Name	Field Description
Logical Switch Name	Universally unique identifier (UUID) of the logical switch.
MAC Address	MAC addresses of *virtual machines (VMs)*.
IP Address	IP address of VMs. **NOTE**: If the IP addresses of VMs are not published by the SDN controller, this field displays `0.0.0.0`.
Encapsulation	Encapsulation type.

Table 957: show ovsdb mac Output Fields *(Continued)*

Field Name	Field Description
VTEP Address	IP address of the hardware or software VTEP from which the MAC address was learned. Further, this VTEP can forward VM traffic to the associated host.
MAC Count	**NOTE**: This field is supported by all Juniper Networks devices that support OVSDB except EX9200 switches. Number of all or specified MAC addresses learned by the Juniper Networks device.

Sample Output

show ovsdb mac

```
user@host> show ovsdb mac
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
  Mac                    IP                 Encapsulation      Vtep
  Address                Address                               Address
  02:00:00:00:03:01      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  02:00:00:00:03:02      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  02:00:00:00:03:03      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  02:00:00:00:03:04      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  02:00:00:00:03:05      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  04:00:00:00:03:05      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  06:00:00:00:03:01      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  06:00:00:00:03:02      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  06:00:00:00:03:03      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  06:00:00:00:03:04      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  06:00:00:00:03:05      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  40:b4:f0:06:6f:f0      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  ff:ff:ff:ff:ff:ff      0.0.0.0            Vxlan over Ipv4    10.100.100.1

Logical Switch Name: bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab
  Mac                    IP                 Encapsulation      Vtep
  Address                Address                               Address
  02:00:00:00:11:01      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  02:00:00:00:11:02      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  02:00:00:00:11:03      0.0.0.0            Vxlan over Ipv4    10.1.1.29
```

```
  02:00:00:00:11:04      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  02:00:00:00:11:05      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  04:00:00:00:11:05      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  06:00:00:00:11:01      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  06:00:00:00:11:02      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  06:00:00:00:11:03      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  06:00:00:00:11:04      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  06:00:00:00:11:05      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  40:b4:f0:06:6f:f0      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  00:23:9c:5e:a7:f0      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  08:00:00:00:11:01      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  08:00:00:00:11:02      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  08:00:00:00:11:03      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  08:00:00:00:11:04      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  08:00:00:00:11:05      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  ff:ff:ff:ff:ff:ff      0.0.0.0            Vxlan over Ipv4    10.110.110.1
...
```

show ovsdb mac address

```
user@host> show ovsdb mac address 02:00:00:00:03:01

  Mac                    IP                 Encapsulation      Vtep
  Address                Address                               Address
  02:00:00:00:03:01      0.0.0.0            Vxlan over Ipv4    10.255.18.22
```

show ovsdb mac logical-switch

```
user@host> show ovsdb mac logical-switch bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab
Logical Switch Name: bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab
  Mac                    IP                 Encapsulation      Vtep
  Address                Address                               Address
  02:00:00:00:11:01      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  02:00:00:00:11:02      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  02:00:00:00:11:03      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  02:00:00:00:11:04      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  02:00:00:00:11:05      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  04:00:00:00:11:05      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  06:00:00:00:11:01      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  06:00:00:00:11:02      0.0.0.0            Vxlan over Ipv4    10.1.1.29
```

```
  06:00:00:00:11:03      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  06:00:00:00:11:04      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  06:00:00:00:11:05      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  40:b4:f0:06:6f:f0      0.0.0.0            Vxlan over Ipv4    10.1.1.29
  00:23:9c:5e:a7:f0      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  08:00:00:00:11:01      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  08:00:00:00:11:02      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  08:00:00:00:11:03      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  08:00:00:00:11:04      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  08:00:00:00:11:05      0.0.0.0            Vxlan over Ipv4    10.255.18.22
  ff:ff:ff:ff:ff:ff      0.0.0.0            Vxlan over Ipv4    10.110.110.1
```

show ovsdb mac local unicast

```
user@host> show ovsdb mac local unicast
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
  Mac                    IP                 Encapsulation      Vtep
  Address                Address                               Address
  02:00:00:00:03:01      0.0.0.0            Vxlan over Ipv4    10.255.181.72
  02:00:00:00:03:02      0.0.0.0            Vxlan over Ipv4    10.255.181.72
  02:00:00:00:03:03      0.0.0.0            Vxlan over Ipv4    10.255.181.72
  02:00:00:00:03:04      0.0.0.0            Vxlan over Ipv4    10.255.181.72
  02:00:00:00:03:05      0.0.0.0            Vxlan over Ipv4    10.255.181.72
  04:00:00:00:03:05      0.0.0.0            Vxlan over Ipv4    10.255.181.72
  06:00:00:00:03:01      0.0.0.0            Vxlan over Ipv4    10.255.181.72
  06:00:00:00:03:02      0.0.0.0            Vxlan over Ipv4    10.255.181.72
  06:00:00:00:03:03      0.0.0.0            Vxlan over Ipv4    10.255.181.72
  06:00:00:00:03:04      0.0.0.0            Vxlan over Ipv4    10.255.181.72
  06:00:00:00:03:05      0.0.0.0            Vxlan over Ipv4    10.255.181.72
  40:b4:f0:06:6f:f0      0.0.0.0            Vxlan over Ipv4    10.255.181.72
...
```

show ovsdb mac (Count of All Local, Remote, Unicast, and Multicast MAC Addresses for All Logical Switches)

```
user@host> show ovsdb mac count
MAC count: 6877
```

Release Information

Command introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

OVSDB Schema for Physical Devices

show ovsdb statistics interface

IN THIS SECTION

- Syntax | 22899
- Description | 22899
- Options | 22900
- Required Privilege Level | 22900
- Output Fields | 22900
- Sample Output | 22900
- Release Information | 22901

Syntax

```
show ovsdb statistics interface
<interface-name>
```

Description

Display statistics for *Open vSwitch Database (OVSDB)*-managed interfaces configured by using the `interfaces` *`interface-name`* statement in the [`edit protocols ovsdb`] hierarchy.

When an interface is configured as OVSDB-managed, the collection of statistics for that interface begins, and the statistics displayed at any given time reflects the data collected up to that point.

Options

none	Display statistics for all configured OVSDB-managed interfaces.
interface-name	Display statistics for the specified interface.

Required Privilege Level

admin

Output Fields

Table 958 on page 22900 lists the output fields for the `show ovsdb statistics interface` command. Output fields are listed in the approximate order in which they appear.

Table 958: show ovsdb statistics interface Output Fields

Field Name	Field Description
Num of rx pkts	Number of packets received by the interface.
Num of tx pkts	Number of packets sent by the interface.
Num of rx bytes	Number of bytes received by the interface.
Num of tx bytes	Number of bytes sent by the interface.

Sample Output

show ovsdb statistics interface

```
user@host> show ovsdb statistics interface
Interface Name: ge-7/0/9.0
Num of rx pkts: 945                    Num of tx pkts: 113280890
Num of rx bytes: 56700                  Num of tx bytes: 57531319540
Interface Name: ge-7/0/10.0
Num of rx pkts: 459                    Num of tx pkts: 473840856
```

```
Num of rx bytes: 84747                 Num of tx bytes: 45830738532
Interface Name: ge-7/0/11.0
Num of rx pkts: 305                   Num of tx pkts: 367483456
Num of rx bytes: 98974                 Num of tx bytes: 33495468092
```

show ovsdb statistics interface (Specific Interface)

```
user@host> show ovsdb statistics interface ge-7/0/9.0
Interface Name: ge-7/0/9.0
Num of rx pkts: 945                   Num of tx pkts: 113280890
Num of rx bytes: 56700                 Num of tx bytes: 57531319540
```

Release Information

Command introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

interfaces (OVSDB) | 9151

show ovsdb tunnels

IN THIS SECTION

- Syntax | 22902
- Description | 22902
- Options | 22903
- Required Privilege Level | 22903
- Output Fields | 22903
- Sample Output | 22904
- Release Information | 22905

Syntax

```
show ovsdb tunnels
<destination-address ip-address>
<source-address ip-address>
```

Description

Display the status of tunnels established between Juniper Networks devices that act as hardware virtual tunnel endpoints (VTEPs), and software VTEPs and service nodes that replicate Layer 2 broadcast, unknown unicast, and multicast (BUM) packets within a Virtual Extensible LAN (VXLAN) managed by the Open vSwitch Database (OVSDB) protocol. (In this topic, the software VTEPs and service nodes are known collectively as *replicators*.)

In this topology with VMware NSX controllers deployed, the hardware VTEP uses the Bidirectional Forwarding Detection (BFD) protocol to monitor replicators learned by the NSX controllers. In this situation, BFD is used to decrease the possibility that BUM packets are forwarded to non-functional replicators and eventually dropped.

To monitor the status of the replicators, the hardware VTEP, replicators, and NSX controllers must all be BFD-capable. (Juniper Networks switches that support OVSDB and VXLAN are BFD-capable.) In addition, the NSX controller must enable BFD on the hardware VTEP and replicators. When the NSX controller has enabled BFD on the hardware VTEP and replicators, their BFD status is considered to be enabled. If the NSX controller cannot enable BFD on these entities (for example, the entity is not BFD-capable), their BFD status is considered to be disabled,

With the BFD protocol enabled on a hardware VTEP, upon receipt of a BUM packet on an OVSDB-managed interface, the hardware VTEP can choose a functional replicator to handle the packet.

In the output of the `show ovsdb tunnels` command, each tunnel is identified by the IP addresses of the source and destination entities at the ends of the tunnel. Each tunnel that appears is from the perspective of the hardware VTEP. The hardware VTEP provides information about its end of a particular tunnel, while the NSX controller collects the same information from the replicator at the other end of the tunnel. Both hardware VTEP and NSX controller push the information to the tunnel table in the OVSDB schema for physical devices.

Some settings of the `BFD Enabled` and `BFD Status` fields in the `show ovsdb tunnels` output can indicate issues:

- A tunnel with a `BFD Enabled` setting of `Enabled` and a `BFD Status` of `Down` usually indicates that there is an issue with the replicator.
- A tunnel with a `BFD Enabled` setting of `Disabled` and a `BFD Status` of `AdminDown` indicates that either the hardware VTEP, the NSX controller, or the replicator is not BFD-capable or is having an issue with BFD. As a result, a BFD session between the hardware VTEP and the replicator cannot be enabled.

Options

Use one or more of the following options to display a more specific list of tunnels.

none	Display all tunnels.
destination-address *ip-address*(Optional)	Display tunnels that have the specified destination address.
source-address *ip-address*(Optional)	Display tunnels that have the specified source address.

Required Privilege Level

admin

Output Fields

Table 959 on page 22903 lists the output fields for the `show ovsdb tunnels` command. Output fields are listed in the approximate order in which they appear.

Table 959: show ovsdb tunnels Output Fields

Field Name	Field Descriptions
SRC IP Address	IP address of the OVSDB-managed physical interface from which the BFD control packet is received.
DST IP Address	IP address of the replicator to which the BFD control packet is sent.
BFD Enabled	Status of the BFD protocol on the hardware VTEP and replicator at either ends of the tunnel (`Enabled` or `Disabled`).

Table 959: show ovsdb tunnels Output Fields *(Continued)*

Field Name	Field Descriptions
BFD Status	Status of a BFD session between a hardware VTEP and a replicator. Possible values include: • `Up`—The BFD session is enabled and up. • `Down`—The BFD session is considered to be down if a replicator does not respond to three BFD control messages. • `AdminDown`—The BFD session is disabled by the NSX controller.

Sample Output

show ovsdb tunnels

```
user@host>  show ovsdb tunnels
SRC IP Address       DST IP Address       BFD Enabled     BFD Status
192.168.110.110      192.168.100.1        Enabled         Up
192.168.110.110      192.168.100.2        Disabled        AdminDown
192.168.110.120      192.168.100.20       Enabled         Down
```

show ovsdb tunnels (Specific Destination)

```
user@host>  show ovsdb tunnels destination-address 192.168.100.1
SRC IP Address       DST IP Address       BFD Enabled     BFD Status
192.168.110.110      192.168.100.1        Enabled         Up
```

show ovsdb tunnels (Specific Source)

```
user@host> show ovsdb tunnels source-address 192.168.110.110
SRC IP Address       DST IP Address       BFD Enabled     BFD Status
```

```
192.168.110.110      192.168.100.1        Enabled         Up
192.168.110.110      192.168.100.2        Disabled        AdminDown
```

Release Information

Command introduced in Junos OS Release 14.1X53-D40.

RELATED DOCUMENTATION

Understanding BFD in a VMware NSX Environment with OVSDB and VXLAN

show ovsdb virtual-tunnel-end-point

IN THIS SECTION

- Syntax | 22905
- Description | 22905
- Options | 22906
- Required Privilege Level | 22906
- Output Fields | 22906
- Sample Output | 22907
- Release Information | 22908

Syntax

```
show ovsdb virtual-tunnel-end-point
address <ip-address>
encapsulation <encapsulation-type>
```

Description

Display information about the following entities that the Juniper Networks device has learned:

- Other Juniper Networks devices that function as hardware *virtual tunnel endpoints (VTEPs)*
- Software VTEPs
- Service nodes
- Top-of-rack service nodes (TSNs)

Options

none	Display information about all VTEPs, service nodes, and TSNs that the Juniper Networks device has learned.
address *ip-address*	Display information about the entity with the specified IP address.
encapsulation *encapsulation-type*	Display information about all entities with the specified encapsulation type.

Required Privilege Level

admin

Output Fields

Table 960 on page 22906 lists the output fields for the `show ovsdb virtual-tunnel-end-point` command. Output fields are listed in the approximate order in which they appear.

Table 960: show ovsdb virtual-tunnel-end-point Output Fields

Field Name	Field Description
Encapsulation	Encapsulation type of entity.
IP Address	IP address of entity.
Num of MACs	Number of media access control (MAC) addresses learned by the entity.

Sample Output

show ovsdb virtual-tunnel-end-point

```
user@host> show ovsdb virtual-tunnel-end-point

Encapsulation               Ip Address                          Num of MAC's
VXLAN over IPv4             10.255.181.43                       24
VXLAN over IPv4             10.255.181.50                       12
VXLAN over IPv4             10.255.181.72                       24
```

show ovsdb virtual-tunnel-end-point address (Specific Address)

```
user@host> show ovsdb virtual-tunnel-end-point address 10.255.181.43

Encapsulation               Ip Address                          Num of MAC's
VXLAN over IPv4             10.255.181.43                       24
```

show ovsdb virtual-tunnel-end-point encapsulation (Specific Encapsulation)

```
user@host> show ovsdb virtual-tunnel-end-point encapsulation vxlan-over-ipv4

Encapsulation               Ip Address                          Num of MAC's
VXLAN over IPv4             10.255.181.43                       24
VXLAN over IPv4             10.255.181.50                       12
VXLAN over IPv4             10.255.181.72                       24
```

show ovsdb virtual-tunnel-end-point address (Specific Address) encapsulation (Specific Encapsulation)

```
user@host> show ovsdb virtual-tunnel-end-point address 10.255.181.43 encapsulation vxlan-over-
ipv4

Encapsulation               Ip Address                          Num of MAC's
VXLAN over IPv4             10.255.181.43                       24
```

Release Information

Command introduced in Junos OS Release 14.1R2.

show passive-monitoring error

IN THIS SECTION

- Syntax | 22908
- Description | 22908
- Options | 22908
- Required Privilege Level | 22909
- Output Fields | 22909
- Sample Output | 22910
- Release Information | 22911

Syntax

```
show passive-monitoring error (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring error statistics.

Options

*** | all | mo-*fpc/pic/port*** Display error statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

Table 961 on page 22909 lists the output fields for the `show passive-monitoring error` command. Output fields are listed in the approximate order in which they appear.

Table 961: show passive-monitoring error Output Fields

Field Name	Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.
Interface state	State of the passive monitoring interface: • **Monitoring**—Specified interface is actively monitoring. • **Disabled**—Specified interface has been disabled from the CLI. • **Not monitoring**—The interface is operational, but not monitoring. This condition occurs when an interface first comes online, or when the interface is operational, but no logical unit has been configured under the physical interface. • **Unknown**—Unknown state. • **Error**—An error occurred during the process of determining the state of the interface.
Error information	
Packets dropped (no memory)	Number of packets dropped because of memory shortage.
Packets dropped (not IP)	Number of non-IP packets dropped.

Table 961: show passive-monitoring error Output Fields *(Continued)*

Field Name	Field Description
Packets dropped (not IPv4)	Number of packets dropped because they failed the IPv4 version check.
Packets dropped (header too small)	Number of packets dropped because the packet length or IP header length was too small.
Memory allocation failures	Number of flow record memory allocation failures. A small number reflects failures to replenish the free list. A large number indicates the monitoring station is almost out of memory space.
Memory free failures	Number of flow record memory free failures.
Memory free list failures	Number of flow records received from free list that failed. Memory is nearly exhausted or too many new flows greater than 128 KB are being created per second.
Memory warning	Whether the flows have exceeded 1 million packets per second (Mpps) on a Monitoring Services PIC or 2 Mpps on a Monitoring Services II PIC. The response can be **Yes** or **No**.
Memory overload	Whether the memory has been overloaded. The response can be **Yes** or **No.**
PPS overload	Whether the PIC is receiving more packets per second than the configured threshold. The response can be **Yes** or **No**.
BPS overload	Whether the PIC is receiving more bits per second than the configured threshold. The response can be **Yes** or **No**.

Sample Output

show passive-monitoring error all

```
user@host> show passive-monitoring error all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
```

```
  Error information
    Packets dropped (no memory): 0, Packets dropped (not IP): 0
    Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
    Memory allocation failures: 0, Memory free failures: 0
    Memory free list failures: 0
    Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Not monitoring
  Error information
    Packets dropped (no memory): 0, Packets dropped (not IP): 0
    Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
    Memory allocation failures: 0, Memory free failures: 0
    Memory free list failures: 0
    Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No
```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring flow

IN THIS SECTION

- Syntax | 22912
- Description | 22912
- Options | 22912
- Required Privilege Level | 22912
- Output Fields | 22912
- Sample Output | 22914
- Release Information | 22914

Syntax

```
show passive-monitoring flow (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive flow statistics.

Options

*** | all | mo-*fpc/pic/port*** Display passive flow statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

Table 962 on page 22912 lists the output fields for the `show passive-monitoring flow` command. Output fields are listed in the approximate order in which they appear.

Table 962: show passive-monitoring flow Output Fields

Field Name	Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.

Table 962: show passive-monitoring flow Output Fields *(Continued)*

Field Name	Field Description
Interface state	State of the passive monitoring interface: • **Monitoring**—Specified interface is actively monitoring. • **Disabled**—Specified interface has been disabled from the CLI. • **Not monitoring**—The interface is operational, but not monitoring. This condition occurs when an interface first comes online, or when the interface is operational, but no logical unit has been configured under the physical interface. • **Unknown**—Unknown state. • **Error**—An error occurred during the process of determining the state of the interface.
Flow information	
Flow packets	Number of packets received by an operational PIC.
Flow bytes	Number of bytes received by an operational PIC.
Flow packets 10-second rate	Number of packets per second handled by the PIC and displayed as a 10-second average.
Flow bytes 10-second rate	Number of bytes per second handled by the PIC and displayed as a 10-second average.
Active flows	Number of currently active flows tracked by the PIC.
Total flows	Total number of flows received by an operational PIC.
Flows exported	Total number of flows exported by an operational PIC.
Flows packets exported	Total number of cflowd packets exported by an operational PIC.

Table 962: show passive-monitoring flow Output Fields *(Continued)*

Field Name	Field Description
Flows inactive timed out	Total number of flows that are exported because of inactivity.
Flows active timed out	Total number of long-lived flows that are exported because of an active timeout.

Sample Output

show passive-monitoring flow all

```
user@host> show passive-monitoring flow all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
  Flow information
    Flow packets: 6533434, Flow bytes: 653343400
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 1599
    Flows exported: 1599, Flows packets exported: 55
    Flows inactive timed out: 1599, Flows active timed out: 0

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Monitoring
  Flow information
    Flow packets: 6537780, Flow bytes: 653778000
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 1601
    Flows exported: 1601, Flows packets exported: 55
    Flows inactive timed out: 1601, Flows active timed out: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring memory

IN THIS SECTION

- Syntax | 22915
- Description | 22915
- Options | 22915
- Required Privilege Level | 22915
- Output Fields | 22915
- Sample Output | 22916
- Release Information | 22917

Syntax

```
show passive-monitoring memory (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring memory and flow record statistics

Options

*** | all | mo-*fpc/pic/port*** Display memory and flow record statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

Table 963 on page 22916 lists the output fields for the `show passive-monitoring memory` command. Output fields are listed in the approximate order in which they appear.

Table 963: show passive-monitoring memory Output Fields

Field Name	Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.
Memory utilization	
Allocation count	Number of flow records allocated.
Free count	Number of flow records freed.
Maximum allocated	Maximum number of flow records allocated since the monitoring station booted. This number represents the peak number of flow records allocated at a time.
Allocations per second	Flow records allocated per second during the last statistics interval on the PIC.
Frees per second	Flow records freed per second during the last statistics interval on the PIC.
Total memory used, Total memory free	Total memory currently used and total amount of memory currently free (in bytes).

Sample Output

show passive-monitoring memory all

```
user@host> show passive-monitoring memory all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
  Memory utilization
    Allocation count: 1600, Free count: 1599, Maximum allocated: 1600
    Allocations per second: 3200, Frees per second: 1438
    Total memory used (in bytes): 103579176, Total memory free (in bytes):     163914184
```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring status

IN THIS SECTION

- Syntax | 22917
- Description | 22917
- Options | 22917
- Required Privilege Level | 22917
- Output Fields | 22918
- Sample Output | 22919
- Release Information | 22919

Syntax

```
show passive-monitoring status (*| all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring status.

Options

*** | all | mo-*fpc/pic/port*** Display status for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

Table 964 on page 22918 lists the output fields for the `show passive-monitoring status` command. Output fields are listed in the approximate order in which they appear.

Table 964: show passive-monitoring status Output Fields

Output Field	Output Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.
Interface state	Monitoring state of the passive monitoring interface. • **Monitoring**—PIC is actively monitoring. • **Disabled**—PIC has been disabled using the CLI. • **Not monitoring**—PIC is operational, but not monitoring. This condition can happen while the PIC is coming online, or when the PIC is operational but has no logical unit configured under the physical interface. • **Unknown**
Group index	Integer that represents the monitoring group of which the PIC is a member. **Group index** is a mapping from the group name to an index. It is not related to the number of monitoring groups.
Export interval	Configured export interval for cflowd records, in seconds.
Export format	Configured export format (only cflowd version 5 is supported).
Protocol	Protocol the PIC is configured to monitor (only IPv4 is supported).
Engine type	Configured engine type that is inserted in output cflowd packets.
Engine ID	Configured engine ID that is inserted in output cflowd packets.

Sample Output

show passive-monitoring status all

```
user@host> show passive-monitoring status all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
  Group index: 0
  Export interval: 15 secs, Export format: cflowd v5
  Protocol: IPv4, Engine type: 1, Engine ID: 1

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Disabled

Passive monitoring interface: mo-4/2/0, Local interface index: 46
Interface state: Not monitoring
```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring usage

IN THIS SECTION

- Syntax | 22920
- Description | 22920
- Options | 22920
- Required Privilege Level | 22920
- Output Fields | 22920
- Sample Output | 22921
- Release Information | 22921

Syntax

```
show passive-monitoring usage (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring usage statistics.

Options

*** | all | mo-*fpc/pic/port*** Display usage statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

Table 965 on page 22920 lists the output fields for the `show passive-monitoring usage` command. Output fields are listed in the approximate order in which they appear.

Table 965: show passive-monitoring usage Output Fields

Output Field	Output Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.
CPU utilization	
Uptime	Time, in milliseconds, that the PIC has been operational.
Interrupt time	Total time that the PIC has spent processing packets since the last PIC reset.

Table 965: show passive-monitoring usage Output Fields *(Continued)*

Output Field	Output Field Description
Load (5 second)	CPU load on the PIC, averaged more than 5 seconds. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.
Load (1 minute)	CPU load on the PIC, averaged more than 1 minute. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.

Sample Output

show passive-monitoring usage all

```
user@host> show passive-monitoring usage
Passive monitoring interface: mo-4/0/0, Local interface index: 44
  CPU utilization
    Uptime: 653155 milliseconds, Interrupt time: 40213754 microseconds
    Load (5 second): 20%, Load (1 minute): 17%

Passive monitoring interface: mo-4/1/0, Local interface index: 45
  CPU utilization
    Uptime: 652292 milliseconds, Interrupt time: 40223178 microseconds
    Load (5 second): 22%, Load (1 minute): 15%

Passive monitoring interface: mo-4/2/0, Local interface index: 46
  CPU utilization
    Uptime: 649491 milliseconds, Interrupt time: 40173645 microseconds
    Load (5 second): 22%, Load (1 minute): 10098862%
```

Release Information

Command introduced before Junos OS Release 7.4.

show path-computation-client active-pce

IN THIS SECTION

- Syntax | 22922
- Description | 22922
- Options | 22922
- Required Privilege Level | 22922
- Output Fields | 22923
- Sample Output | 22925
- Release Information | 22927

Syntax

```
show path-computation-client active-pce
<brief | detail>
```

Description

Displays information about the current active Path Computation Element (PCE).

Options

none	Display brief information about the current active PCE.
brief | detail	(Optional) Display the specific level of output.

Required Privilege Level

view

Output Fields

Table 966 on page 22923 describes the output fields for the `show path-computation-client active-pce` command. Output fields are listed in the approximate order in which they appear.

Table 966: show path-computation-client active-pce Output Fields

Field Name	Field Description	Level of Output
IP address	IP address of the current active PCE.	All levels
Priority	Active PCE priority.	All levels
PCE status	Active PCE state: • PCE_STATE_NEW— Initial PCEP session state. • PCE_STATE_RECONNECT—Trying to re-establish TCP connection with the PCEP peer. • PCE_STATE_CONNECTING—Establishing TCP connection with the PCEP peer. • PCE_STATE_CONNECTED—TCP connection established with the PCEP peer. • PCE_STATE_SYNC—Open messages exchanged with the PCEP peer and entering SYNC state. • PCE_STATE_UP—PCEP session established.	All levels
Session type	Active PCE type: • PCE_TYPE_STATELESS—Does not learn LSP state information form PCC. • PCE_TYPE_STATEFUL—Uses LSP state information learned from PCCs to optimize path computations, but does not actively update LSP state. A PCC maintains synchronization with the PCE. • PCE_TYPE_STATEFULACTIVE—Uses LSP state information learned from PCCs to optimize path computations, and actively updates LSP parameters in those PCCs that delegate control of their LSPs to the PCE.	All levels

Table 966: show path-computation-client active-pce Output Fields *(Continued)*

Field Name	Field Description	Level of Output
PCE-primary role	PCE primary role state: • main—Current active PCE. • backup—Backup PCE.	All levels
PCRpts	Number of PC report (PCRpt) messages sent by PCC to a stateful PCE to report current state of LSP(s).	All levels
PCUpdates	Number of PC update (PCUpd) messages sent by a PCE to a PCC to update LSP parameters.	All levels
Local Keepalive timer	Keepalive timer used by or for the PCC.	All levels
Local Dead timer	Dead timer used by or for the PCC.	All levels
Remote Keepalive timer	Keepalive timer used by or for the PCE.	All levels
Remote Dead timer	Dead timer used by or for the PCE.	All levels
PCErr-recv	Information about type, value, and number of PC Error messages received.	All levels
Max unknown messages	Maximum number of unknown messages received for a PCEP session. Recommended value is 5. If the number of unknown messages received by a PCC or PCE is greater than or equal to the maximum number, the PCEP session is closed.	detail
Keepalives received	Number of Keepalive messages received by a PCC from a PCE.	detail
Keepalives sent	Number of Keepalive messages sent by a PCC to a PCE.	detail

Table 966: show path-computation-client active-pce Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Dead timer	Dead timer used by the current active PCE.	detail
Elapsed as main current	Time (in seconds) the PCE is in the main primary role state.	detail
Elapsed as main total	Time (in seconds) the PCE became main from the last PCCD restart.	detail
Unknown msgs/min rate	Number of unknown messages received per minute.	detail
Session failures	Number of PCEP session failures with the PCE.	detail
Delegation timeout in	Time (in seconds) left for LSP delegation to timeout.	detail
Delegation failures	Number of LSP delegation failures.	detail
Connection down	Time (in seconds) since the PCEP session is down.	detail
PCErr-sent	Information about type, value, and number of PC Error messages sent.	All levels

Sample Output

show path-computation-client active-pce

```
user@host> show path-computation-client active-pce
PCE pce1
General
    IP address                : 10.209.57.166
    Priority                  : 2
```

```
    PCE status                : PCE_STATE_NEW
    Session type              : PCE_TYPE_STATEFULACTIVE
    PCE-mastership            : main
    PCE Multipath Capable     : On
    PCC Multipath Capable     : On

Counters
    PCReqs              Total: 0            last 5min: 0            last hour: 0
    PCReps              Total: 0            last 5min: 0            last hour: 0
    PCRpts              Total: 0            last 5min: 0            last hour: 0
    PCUpdates           Total: 0            last 5min: 0            last hour: 0


Timers
    Local             Keepalive timer:         0 [s]   Dead timer:         0 [s]
    Remote            Keepalive timer:         0 [s]   Dead timer:         0 [s]

Errors
    PCErr-recv
    PCErr-sent
            Type: 19            Value: 3            Count: 1
    PCE-PCC-NTFS
    PCC-PCE-NTFS
```

show path-computation-client active-pce detail

```
user@host> show path-computation-client active-pce detail
PCE pce1
General
    IP address              : 172.22.25.223
    Priority                : 1
    PCE status              : PCE_STATE_RECONNECT
    Session type            : PCE_TYPE_STATEFULACTIVE
    PCE-mastership          : main
    Max unknown messages    : 5
    Keepalives received     : 0
    Keepalives sent         : 0
    Dead timer              : 0 [s]
    Elapsed as main current : 1 [s]
    Elapsed as main total   : 2542 [s]
    Unknown msgs/min rate   : 0
```

```
    Session failures        : 575
    Delegation timeout in   : 14 [s]
    Delegation failures     : 21928
    Connection down         : 16 [s]

Counters
    PCReqs              Total: 0            last 5min: 0            last hour: 0
    PCReps              Total: 0            last 5min: 0            last hour: 0
    PCRpts              Total: 31512        last 5min: 7243         last hour: 7243
    PCUpdates           Total: 80           last 5min: 40           last hour: 40
Timers
    Local             Keepalive timer:        30 [s]   Dead timer:       120 [s]
    Remote            Keepalive timer:        30 [s]   Dead timer:       120 [s]

Errors
    PCErr-recv
    PCErr-sent
                    Type: 1             Value: 2            Count: 12
    PCE-PCC-NTFS
    PCC-PCE-NTFS
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

request path-computation-client active-pce | 14728

show path-computation-client lsp

IN THIS SECTION

- Syntax | 22928
- Description | 22928
- Options | 22928

- Required Privilege Level | 22928
- Output Fields | 22929
- Sample Output | 22932
- Release Information | 22935

Syntax

```
show path-computation-client lsp
<extensive (name lsp name)>
<name>
<p2mp>
```

Description

Display the state of label-switched paths (LSPs) known to the Path Computation Client (PCC).

Options

none Display information about LSPs known to the PCC.

extensive (Optional) Display extensive level of output about each known LSP - point-to-point and point-to-multipoint LSPs.

name (Optional) Enter regular expression for LSP names to match. You can filter a particular LSP from a large number of LSPs by entering specific name of the LSP.

p2mp (Optional) Display information about known point-to-multipoint Path Computation Element (PCE)-initiated LSPs.

Required Privilege Level

view

Output Fields

Table 967 on page 22929 describes the output fields for the `show path-computation-client lsp` command. Output fields are listed in the approximate order in which they appear.

Table 967: show path-computation-client lsp Output Fields

Field Name	Field Description
LSP Name	Name of the LSP.
Status	LSP status: • Primary(Act)—Primary LSP. • Bypass—PCE-initiated bypass LSP.
PLSP-Id	PCEP-specific unique identifier for each LSP. The ID is created by the PCC for the lifetime of a PCEP session.
LSP-Type	Type of LSP: • External provisioned • Local
Controller	Name of the external path computing entity.
Path-Setup-Type	Protocol used to set up the LSP: • RSVP-TE • SPRING-TE
Template	Name of template used.
P2MP name	Name of the point-to-multipoint tree that includes the sub-LSPs of the PCE-initiated LSP.

Table 967: show path-computation-client lsp Output Fields *(Continued)*

Field Name	Field Description
P2MP Branch Name	Name of the branch sub-LSP that makes up the point-to-multipoint tree of the PCE-initiated LSP.
PathName	Name of LSP path.
From	Ingress IP address of LSP.
To	Egress IP address of LSP.
State	LSP state: up, down.
Active Path	Name of active path.
Link Protection	LSP Link protection.
P2mp tree	Name of the point-to-multipoint tree that includes the sub-LSPs of the PCE-initiated LSP.
Path cspf status	CSPF computation done by an external controller or local router.
LSP-ID	LSP identifier.
RSVP Error	RSVP error ID.
Priorities	Setup and hold priorities.
Bandwidth	LSP bandwidth.
Requested AutoBw	Requested bandwidth to controller for auto-bandwidth.
Controller	Name of external controller.

Table 967: show path-computation-client lsp Output Fields *(Continued)*

Field Name	Field Description
Record Route	LSP record route object.
From PCE ERO (received)	Explicit Route Object (ERO) received by controller from the routing protocol process.
From RPD ERO (reported)	ERO reported from the routing protocol process.
Configured ERO on PCC	ERO configured on router.
Association	Grouping of LSPs. Association output details: • Association-type—Type of association supported by PCCD. Currently, only `path protection` association is supported. **NOTE**: The `path protection` association is displayed as `PROTECTION` in the CLI output. • Association ID—Association ID. • Association Source—Address of the PCEP speaker that creates the association. For the `path protection` association, association source is always the ingress IP address of the LSP.
PCE Traffic Steering	MVPN flow specification capability.
FS-ID	Flow specification ID.
Route Distinguisher	Route Distinguisher of MVPN instance.
Source Prefix	Matching source address of MVPN flow specification.
Multicast Group Prefix	Matching group address of MVPN flow specification.

Table 967: show path-computation-client lsp Output Fields *(Continued)*

Field Name	Field Description
State	Flow specification state of specific flow specification ID.
Last Rpt/Pcreqest received from RPD at	Time the last PC report or PE request message was received.
Last Update sent to PCE at	Time the last update was sent to PCE.
Last PcUpdate/PcCreate received from PCE at	Time the last PC update message was received from the controller.
Last error sent to PCE at	Time the last PC error was sent to the controller.
Last 5 reasons to send Report/ Pcrequest	• Reconfig • Down • Get_info • Up • Active

Sample Output

show path-computation-client lsp

```
user@host> show path-computation-client lsp

Name        Status            PLSP-Id    LSP-Type         Controller       Path-Setup-Type
Template        p2mp-name
lsp1        Primary(Act)      1          ext-provised     pce1                rsvp
default_pvc     -
lsp1        Bypass            2          ext-provised     pce1                rsvp
default_pvc     -
```

show path-computation-client lsp p2mp

```
user@host> show path-computation-client lsp p2mp
P2MP name: p2mp_tree1

P2MP Branch Name: p2mp_tree1_leaf1

P2MP Branch Name: p2mp_tree1_leaf2

P2MP name: p2mp_tree2

P2MP Branch Name: p2mp_tree2_leaf1

P2MP Branch Name: p2mp_tree2_leaf2
```

show path-computation-client lsp extensive (PCE-Initiated Point-to-Multipoint LSP Mapping to MVPN)

```
user@host> show path-computation-client lsp extensive
LSP Name                  : lsp_test
PathName                  : primary
From : 128.220.10.163 To : 128.220.10.151 State : Up
Active Path               : primary
Link Protection none
LSP Type                  : ext-provised
P2mp tree                 : p2mp_test
Path cspf status          : external_cspf
Template                  : default_pvc
PLSP-ID                   : 2
LSP-ID                    : 1
RSVP Error                : 0x0
Priorities                : 0
Bandwidth                 : 98760
Requested AutoBw          : 0
Controller                : pce1
Record Route              : 1.4.0.2(S) 1.1.0.4(S) 4.7.0.2(S) 1.1.0.4(S) 7.8.0.2(S) 1.1.0.0(S)
From PCE ERO (received)   : 128.220.10.15(L)
From RPD ERO (reported)   : 128.220.10.15(L)
Configured ERO on PCC     : 128.220.10.15(L)
LSP Association:
```

```
     Association-type                Association ID    Association Source
      PROTECTION                      2                 128.220.10.163
PCE Traffic Steering      :
  FS-ID: 1
    Route Distinguisher    : 12345112:123
    Source Prefix.         : 10.220.10.15/24
    Multicast Group Prefix : 10.220.10.11/24
    State                  : Active(MVPN instance configured)
  FS-ID: 2
    Route Distinguisher.   : 12345112:123
    Source Prefix.         : 10.220.10.15/24
    Multicast Group Prefix : 10.220.10.12/24
    State                  : Inactive(MVPN S,G already exist)
Last Rpt/Pcreqest received from RPD at        : 23:46:12.000
Last Update sent to PCE at                    : 05:30:00.000
Last PcUpdate/PcCreate received from PCE at   : 23:46:12.000
Last error sent to PCE at                     : 05:30:00.000
Last 5 reasons to send Report/Pcrequest       : Reconfig, Down, Get_info, Up, Active,
```

show path-computation-client lsp extensive name <lsp name>

```
user@host> show path-computation-client lsp extensive name R0-R4-1/via-R1
LSP Name                  : R0-R4-1/via-R1
PathName                  : via-R1
From : 192.0.2.0       To : 198.51.100.0       State : Up
Active Path               : -
Link Protection           : none
LSP Type                  : local
P2mp tree                 : NULL
Path cspf status          : local_cspf
Controller                : -
Template                  : NULL
PLSP-ID                   : 4
LSP-ID                    : 0
RSVP Error                : 0x0
Requested AutoBw          : 0bps
Record Route              : (Label=27) (Label=289) (Label=158) [Type: Primary, path_id: 128,
Weight: 1, state: Active]
Record Route              : (Label=27) (Label=173) (Label=79) (Label=158) [Type: Primary,
path_id: 129, Weight: 1, state: Active]
From PCE ERO (received)   : NULL
```

```
From RPD ERO (reported)   : (Label=27) (Label=289) (Label=158) [Type: Primary, path_id: 128,
Weight: 1, state: Active]
From RPD ERO (reported)   : (Label=27) (Label=173) (Label=79) (Label=158) [Type: Primary,
path_id: 129, Weight: 1, state: Active]
Configured ERO on PCC     : Not Supported
Bandwidth:
        Intended          : 0bps
        Actual            : 0bps
Metric:
        Intended          : 0
        Actual            : 0
Mapped Flowspec (FS-Ids)  : -
LSP Attributes:
       Exclude-Any: 0, Include-Any: 0, Include-All: 0
       Setup Priority: 0, Hold-Priority: 0
       Local Protection Bit: FALSE
Last Rpt/Pcreqest received from RPD at        : 22:04:27.000
Last Update sent to PCE at                    : 16:00:00.000
Last PcUpdate/PcCreate received from PCE at   : 16:00:00.000
Last error sent to PCE at                     : 16:00:00.000
Last 5 reasons to send Report/Pcrequest
```

Release Information

Command introduced in Junos OS Release 17.2R1.

`extensive` and `p2mp` options introduced in Junos OS Release 18.3R1 on MX Series routers.

`name` option introduced in Junos OS Release 20.1R1 and Junos OS Evolved Release 20.1R1 on PTX10003-160C and PTX10008 routers.

RELATED DOCUMENTATION

- show path-computation-client active-pce | 22922
- show path-computation-client statistics | 22936
- show path-computation-client status | 22944

show path-computation-client statistics

IN THIS SECTION

- Syntax | 22936
- Description | 22936
- Options | 22936
- Required Privilege Level | 22936
- Output Fields | 22937
- Sample Output | 22940
- Release Information | 22944

Syntax

```
show path-computation-client statistics
<brief | detail>
<all>
```

Description

Display statistics about the Path Computation Element (PCE).

Options

none	Display statistics about the primary PCE.
brief | detail	(Optional) Display the specific level of output.
all	(Optional) Display the statistics about all PCEs configured on the PCC.

Required Privilege Level

view

Output Fields

Table 968 on page 22937 describes the output fields for the `show path-computation-client statistics` command. Output fields are listed in the approximate order in which they appear.

Table 968: show path-computation-client statistics Output Fields

Field Name	Field Description	Level of Output
IP address	IP address of the PCE.	All levels
Priority	PCE priority.	All levels
PCE status	PCE state: • PCE_STATE_NEW— Initial PCEP session state. • PCE_STATE_RECONNECT—Trying to re-establish TCP connection with the PCEP peer. • PCE_STATE_CONNECTING—Establishing TCP connection with the PCEP peer. • PCE_STATE_CONNECTED—TCP connection established with the PCEP peer. • PCE_STATE_SYNC—Open messages exchanged with the PCEP peer and entering SYNC state. • PCE_STATE_UP—PCEP session established.	All levels

Table 968: show path-computation-client statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Session type	Active PCE type: • PCE_TYPE_STATELESS—Does not learn LSP state information form PCC. • PCE_TYPE_STATEFUL—Uses LSP state information learned from PCCs to optimize path computations, but does not actively update LSP state. A PCC maintains synchronization with the PCE. • PCE_TYPE_STATEFULACTIVE—Uses LSP state information learned from PCCs to optimize path computations, and actively updates LSP parameters in those PCCs that delegate control of their LSPs to the PCE.	All levels
PCE-primary role	PCE primary role state: • main • primary • backup	All levels
PCRpts	Number of PC report (PCRpt) messages sent by PCC to a stateful PCE to report current state of LSP(s).	All levels
PCUpdates	Number of PC update (PCUpd) messages sent by a PCE to a PCC to update LSP parameters.	All levels
Local Keepalive timer	Keepalive timer used by or for the PCC.	All levels
Local Dead timer	Dead timer used by or for the PCC.	All levels
Remote Keepalive timer	Keepalive timer used by or for the PCE.	All levels

Table 968: show path-computation-client statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Remote Dead timer	Dead timer used by or for the PCE.	All levels
PCErr-recv	Information about type, value, and number of PC Error messages received.	All levels
PCErr-sent	Information about type, value, and number of PC Error messages sent.	All levels
Max unknown messages	Maximum number of unknown messages received for a PCEP session. Recommended value is **5**. If the number of unknown messages received by a PCC or PCE is greater than or equal to the maximum number, the PCEP session is closed.	detail
Keepalives received	Number of Keepalive messages received by a PCC from a PCE.	detail
Keepalives sent	Number of Keepalive messages sent by a PCC to a PCE.	detail
Elapsed as main current	Time (in seconds) the PCE is in the main primary role state.	detail
Elapsed as main total	Time (in seconds) the PCE became main from the last PCCD restart.	detail
Unknown msgs/min rate	Number of unknown messages received per minute.	detail
Session failures	Number of PCEP session failures with the PCE.	detail
Delegation timeout in	Time (in seconds) left for LSP delegation to timeout.	detail
Delegation failures	Number of LSP delegation failures.	detail

Table 968: show path-computation-client statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Connection down	Time (in seconds) since the PCEP session is down.	detail
Local IP address	IP address of the PCC.	detail
LSP provisioning allowed	LSP provisioning capability.	detail
P2MP LSP report allowed	Report capability of point-to-multipoint LSP.	detail
P2MP LSP update allowed	Update capability of point-to-multipoint LSP.	detail
P2MP LSP init allowed	Initiate capability of point-to-multipoint LSP.	detail
PCE Traffic Steering	Traffic steering capability of the PCE.	detail

Sample Output

show path-computation-client statistics all

```
user@host> show path-computation-client statistics all
PCE pce1

General
    IP address              : 10.209.57.166
    Priority                : 2
    PCE status              : PCE_STATE_NEW
    Session type            : PCE_TYPE_STATEFULACTIVE
    PCE-mastership          : main

Counters
    PCReqs              Total: 0            last 5min: 0            last hour: 0
```

```
    PCReps              Total: 0            last 5min: 0            last hour: 0
    PCRpts              Total: 0            last 5min: 0            last hour: 0
    PCUpdates           Total: 0            last 5min: 0            last hour: 0

Timers
    Local             Keepalive timer:         0 [s]   Dead timer:         0 [s]
    Remote            Keepalive timer:         0 [s]   Dead timer:         0 [s]

Errors
    PCErr-recv
    PCErr-sent
    PCE-PCC-NTFS
    PCC-PCE-NTFS

PCE pce2

General
    IP address               : 10.31.32.1
    Priority                 : 10
    PCE status               : PCE_STATE_NEW
    Session type             : PCE_TYPE_STATEFULACTIVE
    PCE-mastership           : backup

Counters
    PCReqs              Total: 0            last 5min: 0            last hour: 0
    PCReps              Total: 0            last 5min: 0            last hour: 0
    PCRpts              Total: 0            last 5min: 0            last hour: 0
    PCUpdates           Total: 0            last 5min: 0            last hour: 0

Timers
    Local             Keepalive timer:         0 [s]   Dead timer:         0 [s]
    Remote            Keepalive timer:         0 [s]   Dead timer:         0 [s]

Errors
    PCErr-recv
    PCErr-sent
    PCE-PCC-NTFS
    PCC-PCE-NTFS
```

show path-computation-client statistics detail

```
user@host> show path-computation-client statistics detail
PCE pce1
General
    IP address               : 10.209.57.166
    Priority                 : 2
    PCE status               : PCE_STATE_NEW
    Session type             : PCE_TYPE_STATEFULACTIVE
    PCE-mastership           : main
    Max unknown messages     : 5
    Keepalives received      : 0
    Keepalives sent          : 0
    Dead timer               : 0 [s]
    Elapsed as main current  : 294 [s]
    Elapsed as main total    : 294 [s]
    Unknown msgs/min rate    : 0
    Session failures         : 0
    Replies timedout         : 0
    Delegation timeout in    : 26 [s]
    Delegation failures      : 0
    Connection down          : 4 [s]

Counters
    PCReqs              Total: 0            last 5min: 0            last hour: 0
    PCReps              Total: 0            last 5min: 0            last hour: 0
    PCRpts              Total: 0            last 5min: 0            last hour: 0
    PCUpdates           Total: 0            last 5min: 0            last hour: 0

Timers
    Local             Keepalive timer:         0 [s]   Dead timer:         0 [s]
    Remote            Keepalive timer:         0 [s]   Dead timer:         0 [s]

Errors
    PCErr-recv
    PCErr-sent
    PCE-PCC-NTFS
    PCC-PCE-NTFS
```

show path-computation-client statistics (PCE-Initiated Point-to-Multipoint LSP Mapping to MVPN)

```
user@host> show path-computation-client statistics
PCE pce1
--------------------------------------------
General
    PCE IP address           : 10.220.11.59
    Local IP address         : 128.220.11.56
    Priority                 : 0
    PCE status               : PCE_STATE_UP
    Session type             : PCE_TYPE_STATEFULACTIVE
    LSP provisioning allowed : On
    P2MP LSP report allowed  : On
    P2MP LSP update allowed  : On
    P2MP LSP init allowed    : Off
    PCE-mastership           : main
    PCE Traffic Steering     : On

Counters
    PCReqs              Total: 0            last 5min: 0            last hour: 0
    PCReps              Total: 0            last 5min: 0            last hour: 0
    PCRpts              Total: 4            last 5min: 0            last hour: 0
    PCUpdates           Total: 2            last 5min: 0            last hour: 0
    PCCreates           Total: 1            last 5min: 0            last hour: 0

Timers
    Local  Keepalive timer:   30 [s]  Dead timer:  120 [s]  LSP cleanup timer:  500 [s]
    Remote Keepalive timer:   30 [s]  Dead timer:  120 [s]  LSP cleanup timer:    0 [s]

Errors
    PCErr-recv
    PCErr-sent
    PCE-PCC-NTFS
    PCC-PCE-NTFS

Pcupdate empty ero action counters
    Send-err               : 0
    Tear down path         : 0
    Routing decision       : 0
    Routing decision failed: 0
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

clear path-computation-client statistics | 3802

show path-computation-client status

IN THIS SECTION

- Syntax | 22944
- Description | 22944
- Options | 22944
- Required Privilege Level | 22945
- Output Fields | 22945
- Sample Output | 22946
- Release Information | 22947

Syntax

```
show path-computation-client status
<extensive>
```

Description

Display the status of the Path Computation Client (PCC).

Options

none Display the status of the PCC.

extensive (Optional) Display extensive information about the PCC including point-to-point and point-to-multipoint PCE-initiated LSPs.

For point-to-multipoint PCE-initiated LSPs, the extensive output displays the point-to-multipoint LSP tree and branches separately for a PCEP session.

Required Privilege Level

view

Output Fields

Table 969 on page 22945 describes the output fields for the `show path-computation-client status` command. Output fields are listed in the approximate order in which they appear.

Table 969: show path-computation-client status Output Fields

Field Name	Field Description
Session	Name of the PCE with which the PCEP session is established.
Type	Type of PCE.
Provisioning	Provisioning status of the PCE.
Status	PCEP session status.
Total number of LSPs	Number of LSPs in total.
Static LSPs	Status of point-to-point and point-to-multipoint static LSPs.
Externally controlled LSPs	Status of point-to-point and point-to-multipoint LSPs that are controlled by a PCE.
Externally provisioned LSPs	Status of point-to-point and point-to-multipoint LSPs that are provisioned by a PCE.

Table 969: show path-computation-client status Output Fields *(Continued)*

Field Name	Field Description
Orphaned LSPs	Status of LSPs that are in the orphaned state because of PCEP session failure.
Delegated	Status of point-to-point and point-to-multipoint LSPs that are delegated to the PCE by the PCC.

Sample Output

show path-computation-client status

```
user@host> show path-computation-client status extensive
Session              Type                        Provisioning   Status
 pce1               Stateful Active             On             Up

LSP Summary
 Total number of LSPs               : 0
 Static LSPs                        : 0
         P2P                        : 0
                                             : 0/0 (primary/bypass)
         P2MP                       : 0/0 (branches/trees)
 Externally controlled LSPs         : 0
         P2P                        : 0
                                                 : 0/0 (primary/bypass)
         P2MP                       : 0/0 (branches/trees)
 Externally provisioned LSPs        : 0/16000 (current/limit)
         P2P                        : 0
                                                 : 0/0 (primary/ bypass)
         P2MP                       : 0/0 (branches/trees)
 Orphaned LSPs                      : 0

pce1 (main)
 Delegated             : 0
        P2P            : 0
                               : 0/0 (primary/bypass)
        P2MP           : 0/0 (branches/trees)
 Externally provisioned: 0
```

```
        P2P            : 0
                               : 0/0 (primary/bypass)
        P2MP           : 0/0 (branches/trees)
```

Release Information

Command introduced in Junos OS Release 17.2R1.

`extensive` option introduced in Junos OS Release 18.3R1 on MX Series routers.

RELATED DOCUMENTATION

- show path-computation-client active-pce | 22922
- show path-computation-client statistics | 22936
- show path-computation-client lsp | 22927

show performance-monitoring mpls lsp

IN THIS SECTION

- Syntax | 22948
- Description | 22948
- Options | 22948
- Required Privilege Level | 22948
- Output Fields | 22948
- Sample Output | 22951
- Release Information | 22954

Syntax

```
show performance-monitoring mpls lsp
<brief | detail | extensive>
<name lsp name>
```

Description

Display the following performance monitoring data:

- Packet loss measurement
- Packet throughput measurement
- Two-way channel delay
- Round-trip delay
- Inter-packet delay variation (IPDV)

Options

none	Display standard information performance monitoring data.
brief | detail | extensive	(Optional) Display the specified level of output. **NOTE**: The extensive option displays the same information as the detail option.
name *lsp name*	(Optional) Display information about the specified LSP.

Required Privilege Level

View

Output Fields

Table 970 on page 22949 describes the output fields for the `show performance-monitoring mpls lsp` command. Output fields are listed in the approximate order in which they appear.

Table 970: show performance-monitoring mpls lsp Output Fields

Field Name	Display Data		Field Description	Level of Output
Session	Total		Total number of performance monitoring sessions created.	All Levels
	Up		Number of performance monitoring sessions that are up and running.	All Levels
	Down		Number of performance monitoring sessions that are down.	All Levels
LSP name			Name of the LSP.	All Levels
Loss measurement Data	Traffic-class		Traffic class for which loss measurement is performed.	All Levels
	Queries sent		Total number of queries sent for loss measurement.	All Levels
	Responses received		Total number of responses received for loss measurement queries.	All Levels
	Responses dropped due to errors		Total number of loss measurement responses dropped due to errors.	All Levels
	Queries timeout		Number of timed out queries sent for loss measurement.	All Levels
	Forward loss measurement	Average packet loss	Average packet loss (total loss of packets divided by the total number of samples used since the session is up).	All Levels

Table 970: show performance-monitoring mpls lsp Output Fields *(Continued)*

Field Name	Display Data		Field Description	Level of Output
		Average packet throughput	Total number of packets sent divided by the time considered for measurement.	All Levels
	Reverse loss measurement	Average packet loss	Average packet loss (total loss of packets divided by the total number of samples used since the session is up).	All Levels
		Average packet throughput	Total number of packets sent divided by the time considered for measurement.	All Levels
Delay measurement Data	Traffic-class		Traffic class for which delay measurement is performed.	All Levels
	Queries sent		Total number of queries sent for delay measurement.	All Levels
	Responses received		Total number of responses received for delay measurement queries.	All Levels
	Responses dropped due to errors		Total number of delay measurement responses dropped due to errors.	All Levels
	Queries timeout		Number of timed out queries sent for delay measurement.	All Levels
	Best 2-way channel delay		Best available two-way channel delay.	All Levels
	Worst 2-way channel delay		Worst available two-way channel delay.	All Levels

Table 970: show performance-monitoring mpls lsp Output Fields *(Continued)*

Field Name	Display Data	Field Description	Level of Output
	Best round trip time	Best available round-trip time.	All Levels
	Worst round trip time	Worst available round-trip time.	All Levels
	Avg absolute fw delay variation	Average of the variation in forward delay.	All Levels
	Avg absolute rv delay variation	Average of the variation in reverse delay.	All Levels
	Two-way channel delay	Sum of packet delays, excluding the processing time of the remote provider edge (PE) router.	`detail`, `extensive`
	Two-way round trip delay	Total time taken for completing round-trip of packet.	`detail`, `extensive`

Sample Output

show performance-monitoring mpls lsp

```
user@host> show performance-monitoring mpls lsp
Session Total: 3 Up: 3 Down: 0
  LSP name:to_bad, PM State:Up
    Loss measurement Data:
      Duration: 00:04:43
      Traffic-class: None
      Queries sent: 282
      Responses received: 282
      Responses dropped due to errors: 0
      Queries timeout: 0
      Forward loss measurement:
        Average packet loss: 0
```

```
        Average packet throughput: 554338
      Reverse loss measurement:
        Average packet loss: 0
        Average packet throughput: 1352077
LSP name:to_bad, PM State:Up
  Delay measurement Data:
    Duration: 00:04:43
    Traffic-class: 0
    Queries sent: 282
    Responses received: 282
    Responses dropped due to errors: 0
    Queries timeout: 0
    Best 2-way channel delay: 72 usecs
    Worst 2-way channel delay: 365 usecs
    Best round trip time: 843 usecs
    Worst round trip time: 105523 usecs
    Avg absolute fw delay variation: 1619 usecs
    Avg absolute rv delay variation: 1619 usecs
LSP name:to_bad, PM State:Up
  Loss measurement Data:
    Duration: 00:04:43
    Traffic-class: None
    Queries sent: 282
    Responses received: 282
    Responses dropped due to errors: 0
    Queries timeout: 0
    Forward loss measurement:
      Average packet loss: 0
      Average packet throughput: 553927
    Reverse loss measurement:
      Average packet loss: 0
      Average packet throughput: 1351531
  Delay measurement Data:
    Best 2-way channel delay: 76 usecs
    Worst 2-way channel delay: 368 usecs
    Best round trip time: 1082 usecs
    Worst round trip time: 126146 usecs
    Avg absolute fw delay variation: 1618 usecs
    Avg absolute rv delay variation: 1619 usecs
```

show performance-monitoring mpls lsp detail

```
user@host> show performance-monitoring mpls lsp detail
Session Total: 3 Up: 3 Down: 0
  LSP name:to_bad, PM State:Up
    Loss measurement Data:
      Duration: 00:04:53
      Traffic-class: None
      Queries sent: 292
      Responses received: 292
      Responses dropped due to errors: 0
      Queries timeout: 0
      Forward loss measurement:
        Average packet loss: 0
        Average packet throughput: 554486
        Packet loss samples:
          00000000 00000000 00000000 00000000 00000000
        Packet throughput samples:
          00554002 00557550 00557717 00558822 00557107
      Reverse loss measurement:
        Average packet loss: 0
        Average packet throughput: 1352406
        Packet loss samples:
          00000000 00000000 00000000 00000000 00000000
        Packet throughput samples:
          01351088 01365948 01353926 01362976 01358788
  LSP name:to_bad, PM State:Up
    Delay measurement Data:
      Duration: 00:04:53
      Traffic-class: 0
      Queries sent: 292
      Responses received: 292
      Responses dropped due to errors: 0
      Queries timeout: 0
      Best 2-way channel delay: 72 usecs
      Worst 2-way channel delay: 365 usecs
      Best round trip time: 843 usecs
      Worst round trip time: 105523 usecs
      Avg absolute fw delay variation: 1683 usecs
      Avg absolute rv delay variation: 1684 usecs
        Two-way channel delay:
          73 usecs 73 usecs 73 usecs 73 usecs 72 usecs
```

```
        Two-way round trip delay:
          922 usecs 2234 usecs 884 usecs 1121 usecs 1169 usecs
  LSP name:to_bad, PM State:Up
    Loss measurement Data:
      Duration: 00:04:53
      Traffic-class: None
      Queries sent: 292
      Responses received: 292
      Responses dropped due to errors: 0
      Queries timeout: 0
      Forward loss measurement:
        Average packet loss: 0
        Average packet throughput: 554089
        Packet loss samples:
          00000000 00000000 00000000 00000000 00000000
        Packet throughput samples:
          00554007 00557548 00557713 00558547 00557385
      Reverse loss measurement:
        Average packet loss: 0
        Average packet throughput: 1351914
        Packet loss samples:
          00000000 00000000 00000000 00000000 00000000
        Packet throughput samples:
          01358923 01352980 01362436 01223841 01496977
    Delay measurement Data:
      Best 2-way channel delay: 76 usecs
      Worst 2-way channel delay: 368 usecs
      Best round trip time: 1082 usecs
      Worst round trip time: 126146 usecs
      Avg absolute fw delay variation: 1682 usecs
      Avg absolute rv delay variation: 1683 usecs
        Two-way channel delay:
          76 usecs 76 usecs 76 usecs 77 usecs 77 usecs
        Two-way round trip delay:
          107496 usecs 102369 usecs 104048 usecs 1433 usecs 103306 usecs
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

clear performance-monitoring mpls lsp | 3803

performance-monitoring (Protocols MPLS) | 12434

show pfe cfeb

IN THIS SECTION

- Syntax | 22955
- Description | 22955
- Options | 22955
- Required Privilege Level | 22955
- Output Fields | 22956
- Sample Output | 22957
- Release Information | 22960

Syntax

```
show pfe cfeb
```

Description

(M7i routers only) Display Packet Forwarding Engine Compact Forwarding Engine Board (CFEB) status and statistics information.

Options

This command has no options.

Required Privilege Level

admin

Output Fields

Table 971 on page 22956 lists the output fields for the `show pfe cfeb` command. Output fields are listed in the approximate order in which they appear.

Table 971: show pfe cfeb Output Fields

Field Name	Field Description

CFEB status	Status of CFEB:
	• Slot—CFEB slot number.
	• State—Status of the CFEB:
	  • Online—CFEB is online and running.
	  • Offline—CFEB is powered down.
	• Last State Change—Date and time the CFEB state last changed.
	• Uptime (total)—How long the Routing Engine has been connected to the CFEB and, therefore, how long the Flexible PIC Concentrator (FPC) has been up and running.
	• Failures—Number of PFE Peer detach failures.
	• Pending—Number of messages waiting to be sent.
	• Policer Drop Probability—Current policer drop probability. The default is high, and can be configured using the `policer-drop-probability-low` statement.

Peer message type receive qualifiers	• Message Type—IPC Message Type. For example, interface and nexthop.
	• Receive Qualifier – Message receive qualifier for a peer (non-None only):
	  • All
	  • Only this slot
	  • Selective slot

Table 971: show pfe cfeb Output Fields *(Continued)*

Field Name	Field Description
PFE listener statistics	PFE listener statistics: • Open—Number of times a peer was opened. • Close—Number of times a peer was closed. • Sleep—Number of times a thread slept. • Wakeup—Number of times wakeup was issued. • Resync Request—Number of resync requests. • Resync Done—Number of successful resyncs. • Resync Fail—Number of failed resyncs. • Resync Time—Time the resync last happened.
PFE IPC statistics	• type—IPC Message Type. • TX Messages—Number of Tx messages. • RX Messages—Number of Rx messages.
PFE socket-buffer mbuf depth	• bucket—Bucket number. • count—Number of messages in the bucket.
PFE socket-buffer bytes pending transmit	• bucket—Bucket number. • count—Number of bytes pending transmit.

Sample Output

show pfe cfeb

```
user@host> show pfe cfeb
```

```
CFEB status:
  Slot:                 Present
  State:                Online
  Last State Change:    2005-03-10 09:01:25 PST
  Uptime (total):       2d 00:44
  Failures:             0
  Pending:              0
..Policer Drop Probability: HIGH

Peer message type receive qualifiers:
 Message Type      Receive Qualifier
 ----------------  -----------------
              TTP  All
              IFD  All
              IFL  All
          Nexthop  All
              COS  All
            Route  All
      SW Firewall  All
      HW Firewall  All
   PFE Statistics  All
   PIC Statistics  All
         Sampling  All
       Monitoring  None
              ASP  None
             L2TP  None
        Collector  None
 PIC Configuration  All
 Queue Statistics  All
           (null)  None

PFE listener statistics:
  Open:                 1
  Close:                0
  Sleep:                0
  Wakeup:               0
  Resync Request:       0
  Resync Done:          1
  Resync Fail:          0
  Resync Time:          0

PFE IPC statistics:
 type        TX Messages  RX messages
```

```
 ----------- -----------  -----------
      Header           0            0
        Test           0            0
   Interface         562        14582
     Chassis           0            0
        Boot           0            0
    Next-hop         104            0
       Jtree           0            0
       Cprod           0            0
       Route         103            1
         Pfe        3770         2925
         Dfw          10            0
  Mastership           0            0
    Sampling           0            0
        GUCP           0            0
         CoS          50            0
        GCCP           0            0
        GHCP           0            0
        IRSD           0            0
  Monitoring           0            0
          RE           0            0
         PIC           0            0
     ASP cfg           0            0
     ASP cmd           0            0
    L2TP cfg           0            0
   Collector           0            0
   PIC state           0            0
  Aggregator           0            0
       Empty           0            0

PFE socket-buffer mbuf depth:
 bucket          count
 ------          -----
     0               0
     1               0
     2               0
     3               0
     4               0
     5               0
     6               0
     7               0
     8               0
     9               0
```

```
    10                0
    11                0
    12                0
    13                0
    14                0
    15                0
    16                0
    17                0
    18                0
    19                0
    20                0
    21                0

PFE socket-buffer bytes pending transmit:
 bucket          count
 ------          -----
     0               0
     1               0
     2               0
     3               0
     4               0
     5               0
     6               0
     7               0
     8               0
     9               0
    10               0
    11               0
    12               0
    13               0
    14               0
    15               0
    16               0
    17               0
    18               0
    19               0
    20               0
    21               0
```

Release Information

Command introduced before Junos OS Release 7.4.

show pfe data

IN THIS SECTION

- Syntax | 22961
- Description | 22961
- Options | 22962
- Required Privilege Level | 22962
- Output Fields | 22962
- Sample Output | 22969
- Sample Output | 22972
- Release Information | 22974

Syntax

```
show pfe data
<”usp flow config”>
<”usp flow stats”>
```

Description

Displays the Packet Forwarding Engine data. The Packet Forwarding Engine is the central processing element of the forwarding plane, systematically moving the packets in and out of the device. In Junos OS, the Packet Forwarding Engine has a locally stored forwarding table.

The forwarding table is a synchronized copy of all the information from the Route Engine that the forwarding plane needs to handle each packet, including outgoing interfaces, addresses, and so on. Storing a local copy of this information allows the Packet Forwarding Engine to get its job done without going to the control plane every time that it needs to process a packet.

The locally stored forwarding table allows the Packet Forwarding Engine to continue forwarding packets, even when a disruption occurs to the control plane, such as when a routing or other process issue happens.

Options

"usp flow config" Displays current flow configuration in logical system root-logical-system.

"usp flow stats" Displays current flow statistics.

Required Privilege Level

view

Output Fields

Table 972 on page 22963 lists the output fields for the `show pfe data` command. Output fields are listed in the approximate order in which they appear.

Table 972: show pfe data Output Fields

Field Name	Field Description
`Flow main Parameters`	Information about flow main parameters: • `allow-dns-reply`—Displays whether `allow-dns-reply` is enabled or disabled. • `allow-embedded-icmp`—Displays whether `allow-embedded-icmp` is enabled or disabled. • `mcast-buffer-enhance`—Displays whether `mcast-buffer-enhance` is enabled or disabled. • `allow-reverse-ecmp`—Displays whether `allow-reverse-ecmp` is enabled or disabled. • `no-local-favor-ecmp`—Displays whether `no-local-favor-ecmp` is enabled or disabled. • `Reordering-mixed-traffic`—Displays whether `Reordering-mixed-traffic` is enabled or disabled. • `Route-change-timeout`—Displays the route change timeout value. • `NH-resolve-timeout`—Displays next-hop resolve timeout value. • `Mcast-nh-resolve-retry`—Displays multicast next-hop resolve value. • `Pending-sess-queue-length`—Displays the pending session queue length. • `Syn-flood-protection tcp-syn-cookie`—Displays the synchronizing flood protection TCP synchronizing cookies. • `Syn-flood-dynamic-whitelist-time`—Displays the syncronizing floow dynamic allowlist time. • `no-inter-node-forwarding`—Displays whether `no-inter-node-forwarding` is enabled or disabled. • `no-inter-node-forwarding-ipv6`—Displays whether `no-inter-node-forwarding-ipv6` is enabled or disabled. • `enable-reroute-nat-check`—Displays whether `enable-reroute-nat-check` is enabled or disabled. • `force_ip_reassembly`—Displays whether `force_ip_reassembly` is enabled or disabled. • `fragment_limit`—Displays the fragment limit. • `frag_reassembly_timeout`—Displays the fragment reassembly timeout value. • `sync_icmp_session`—Displays whether `sync_icmp_session` is enabled or disabled. • `log dropped-icmp-packet`—Displays whether log `dropped-icmp-packet` is enabled or disabled.

Table 972: show pfe data Output Fields *(Continued)*

Field Name	Field Description
	• `log dropped-illegal-packet`—Displays whether log `dropped-illegal-packet` is enabled or disabled. • `preserve incoming frag size`—Displays whether preserve incoming fragment size is enabled or disabled. • `log-extension`—Displays whether `log-extension` is enabled or disabled. • `advanced options drop_match_reserved_ip_address`—Displays whether advanced options `drop_match_reserved_ip_address` is enabled or disabled. • `advanced options drop_link_local_address`—Displays whether advanced options `drop_link_local_address` is enabled or disabled. • `advanced options reverse route packet mode vr option`—Displays whether advanced options reverse route packet mode vr option is enabled or disabled. • `advanced options close icmp session immediately`—Displays whether advanced options close Internet Control Message Protocol session immediately is enabled or disabled. • `load-distribution session-affinity ipsec`—Displays whether `load-distribution session-affinity ipsec` is enabled or disabled.
`Aging Parameters`	Information about aging parameters: • `early-ageout`—Displays the number of early ageout. • `low-watermark`—Displays the low watermark value. • `high-watermark`—Displays the high watermark value.
`TCP MSS Parameters`	Information about TCP MSS Parameters: • `all-tcp-mss`—Displays whether `all-tcp-mss` is enabled or disabled. • `ipsec-vpn-tcp-mss`—Displays whether `ipsec-vpn-tcp-mss` is enabled or disabled. • `gre-in-tcp-mss`—Displays whether `gre-in-tcp-mss` is enabled or disabled. • `gre-out-tcp-mss`—Displays whether `gre-out-tcp-mss` is enabled or disabled.

Table 972: show pfe data Output Fields *(Continued)*

Field Name	Field Description
`TCP Session Parameters`	Information about TCP Session Parameters: • `rst-invalid-session`—Displays whether `rst-invalid-session` is enabled or disabled. • `fin-invalid-session`—Displays whether `fin-invalid-session` is enabled or disabled. • `rst-check-sequence`—Displays whether `rst-check-sequence` is enabled or disabled. • `syn-check`—Displays whether `syn-check` is enabled or disabled. • `syn-check-in-tunnel`—Displays whether `syn-check-in-tunnel` is enabled or disabled. • `sequence-check`—Displays whether `sequence-check` is enabled or disabled. • `strict-syn-check`—Displays whether `strict-syn-check` is enabled or disabled. • `tcp-initial-timeout`—Displays TCP initial timeout value. • `maximum-scale-window`—Displays the value of maximum scale window. • `time-wait-state-session-timeout`—Displays the wait time state and session timeout value.
`Flow traceoptions`	Information about flow traceoptions: • `trace flags`—Trace the flag option. • `trace level`—Trace level of the trace message. • `trace present`—Trace the present option. • `trace dst id`—Trace the specified destination ID. • `trace module buffer`—Trace the module buffer value.

Table 972: show pfe data Output Fields *(Continued)*

Field Name	Field Description
`Flow Modes`	Information about flow modes: • `Flow mode for inet`—Flow mode packets to an IPv4 destination. Displays configuration and running mode. • `Flow mode for inet6`—Displays the IPv6 flow mode informations such as the running mode, configuration mode, and the status of IPv6 security knob. • `Flow mode for mpls`—Flow mode packets to MPLS. Displays configuration and running mode. • `Flow mode for iso`—Flow mode to ISO. Displays configuration and running mode.
`Flow action flag`	Information about the flow action flag.
`Session distribution`	• `Config mode`—Displays configuration mode. • `Running mode`—Displays running mode. • `GTP-U distribution`—Displays whether the GTP-U distribution is enabled or disabled.
`Ipsec performance acceleration`	Accelerate the IPSec traffic performance. Displays whether IPSec performance acceleration is enabled or disabled.
`Packet ordering`	Information about packet ordering: • `Config mode`—Displays configuration mode. • `Running mode`—Displays running mode.

Table 972: show pfe data Output Fields *(Continued)*

Field Name	Field Description
`Hash Bucket Size`	Information about hash bucket size: • `Ipv4`—Display Packet Forwarding Engine IPv4 routes. • `Ipv6`—Display Packet Forwarding Engine IPv6 routes. • `Flow packet log enable`—Displays flow packet log enable. • `Flow packet log index`—Displays flow packet log index. • `Flow packet log throttle-interval`—Displays flow packet log throttle-interval. • `Flow tap mode`—Displays whether flow tap mode is enabled or disabled. • `Flow power mode IPsec`—Displays whether flow power mode IPsec is enabled or disabled. • `Fat core group global switch`—Displays whether fat core group global switch is on or off. • `Fat core group size`—Displays the fat core group size.
`Ager Ring Statistics`	Information about the ager ring statistics: • `IPv4 and IPv6 Common Ager Ring`—Displays whether batching is enabled or disabled. • `current tick`—Displays the current session timeout value. • `insert`—Displays the insert value. • `update`—Displays the update value. • `callback invoked`—Displays the callback invoked value. • `remove`—Displays the remove value. • `update skipped`—Displays the skipped update value. • `max watchdog time`—Displays maximum watchdog time. • `timeout value larger`—Displays larger timeout value. • `ager update`—Displays update ager value.

Table 972: show pfe data Output Fields *(Continued)*

Field Name	Field Description
`Session Scan stats`	Information about session scan statistics: • `Requests recvd`—Displays the number of requests received. • `Requests Processing Done`—Displays the number of processing requests completed. • `Request Done Time`—Displays the number of requests done time. • `Skip session scan`—Displays the number of skip session scan. • `Total sessions scanned`—Displays the total number of sessions scanned. • `Batched (active sess limit)`—Displays the number of active sessions limit. • `Pending Que added`—Displays the number of pending queue added. • `Pending Que add fail`—Displays the number of pending queue add failures. • `Response Done`—Displays the number of response completed. • `Response Failed`—Displays the number of responses failed. • `Response Done Time`—Displays the response time completed. • `Scan Handler Invoked`—Displays the number of scan handler invoked. • `PendQ Handler Invoked`—Displays the number of pending handler invoked. • `Session scan running time`—Displays the timer of session scan running. • `Session scan handler yield`—Displays the number of session scan handler yield. • `Scan post even failed`—Displays the number of scan post even failed.

Table 972: show pfe data Output Fields *(Continued)*

Field Name	Field Description
`Flow Session Statistics`	Information about flow session statistics: • `Current flows`—Displays the number of current flows. • `IPv4 flow alloc`—Displays the number of IPv4 flow allocation. • `IPv4 flow free`—Displays the number of IPv4 flow free. • `IPv4 flow alloc failure`—Displays the number of IPv4 flow allocation failures. • `IPv4 flow route skip`—Displays the number of IPv4 flow route skipped. • `IPv4 flow fast reroute`—Displays the number of IPv4 flow fast reroute. • `IPv4 flow tunnel route skip`—Displays the number of IPv4 flow tunnel route skipped. • `ager phase1`—Displays the number of ager phase1. • `ager phase2`—Displays the number of ager phase2. • `cp sess del rexmitted`—Displays the cp session delete rexmitted.
`Flow Table Statistics`	Information about the flow table statistics: • `flow add`—Displays number of flow add entries. • `flow remove`—Displays number of flow delete entries. • `flow find`—Displays number of flow find entries.

Sample Output

show pfe data "usp flow config"

```
user@host> show pfe data "usp flow config"
================ tnp_0x10000080 ================


Current FLOW configuration in logical system root-logical-system:
==========================
```

```
Flow main Parameters::
 allow-dns-reply: disabled (default),
 allow-embedded-icmp: disabled (default),
 mcast-buffer-enhance: disabled (default),
 allow-reverse-ecmp: disabled (default),
 Reordering-mixed-traffic enabled (default),
 Route-change-timeout: 6 (default),
 Pending-sess-queue-length: 5 (default),
 Syn-flood-protection tcp-syn-cookie (default),
 Syn-flood-dynamic-whitelist-time 300,
 no-inter-node-forwarding: disabled (default),
 no-inter-node-forwarding-ipv6: disabled,
 enable-reroute-nat-check: disabled (default),
 force_ip_reassembly: disabled (default),
 sync_icmp_session: disabled (default),
 log dropped-icmp-packet: disabled (default),
 log dropped-illegal-packet: disabled (default),
  preserve incoming frag size : disabled
  log-extension : disabled (default)
 advanced options drop_match_reserved_ip_address: disabled (default),
 advanced options drop_link_local_address: disabled (default),
 advanced options reverse route packet mode vr option: disabled (default),
 advanced options close icmp session immediately: disabled (default),
 load-distribution session-affinity ipsec: disabled (default),
Aging Parameters::
 early-ageout: 20 (default),
 low-watermark: 100 (default),
 high-watermark: 100 (default),
TCP MSS Parameters::
 all-tcp-mss: disabled (default),
 ipsec-vpn-tcp-mss: disabled (default),
 gre-in-tcp-mss: disabled (default),
 gre-out-tcp-mss: disabled (default),
TCP Session Parameters::
 rst-invalid-session: disabled (default)
 fin-invalid-session: disabled (default)
 rst-check-sequence: disabled (default)
 syn-check: enabled (default)
 syn-check-in-tunnel: enabled (default)
 sequence-check: enabled (default)
 strict-syn-check: disabled (default)
 tcp-initial-timeout: 0x14 (default)
```

```
 maximum-scale-window: 0x3ebc0 (default)
 time-wait-state-session-timeout: 150 (default)
Flow traceoptions:
  trace flags: 0x0
  trace level: 0x0
  trace present: 0
  trace dst id: 0
  trace module buffer: 0
FLow Modes:
 Flow mode for inet:
  Running mode: FLOW_MODE_FLOW_BASED
  Config mode: FLOW_MODE_FLOW_BASED
 Flow mode for inet6:
  Running mode: FLOW_MODE_FLOW_BASED
  Config mode: FLOW_MODE_FLOW_BASED
  IPv6 security knob: enabled
 Flow mode for mpls:
  Running mode: FLOW_MODE_DROP
  Config mode: FLOW_MODE_DROP
 Flow mode for iso:
  Running mode: FLOW_MODE_DROP
  Config mode: FLOW_MODE_DROP
Flow action flag: 0x700080
Session distribution:
  Config mode: RR-based
  Running mode: RR-based
  GTP-U distribution: disabled
Ipsec performance acceleration: disabled
Packet ordering:
  Config mode: Hardware
  Running mode: Hardware
Hash Bucket Size:
Ipv4: 65536
Ipv6: 65536
Flow packet log enable: 0
Flow packet log index: 20
Flow packet log throttle-interval: 1024
Flow tap mode: disabled(default)
```

Sample Output

show pfe data "usp flow stats"

```
user@host> show pfe data "usp flow stats"
================ tnp_0x10000080 ================

Ager Ring Statistics
        IPv4 and IPv6 Common Ager Ring: (Batching Disabled)
        current tick                        24929
        insert                              53082
        update                           10618908
        callback invoked                    53082
        remove                                  0
        update skipped                          0
        max watchdog time                    5000
        timeout value larger                    0
        ager update                             0


Session Scan stats
        Requests recvd                          0
        Requests Processing Done                0
        Request Done Time                       0
        Skip session scan                       0
        Total sessions scanned                  0
        Batched (active sess limit)             0
        Pending Que added                       0
        Pending Que add fail                    0
        Response Done                           0
        Response Failed                         0
        Response Done Time                      0
        Scan Handler Invoked                    0
        PendQ Handler Invoked                   0
        Session scan running time              40
        Session scan handler yield              0
        Scan post even failed                   0

Flow Session Statistics
        Current flows                           0
        IPv4 flow alloc                     26541
        IPv4 flow free                      26541
```

```
        IPv4 flow alloc failure                 0
        IPv4 flow route skip                    0
        IPv4 flow fast reroute                  0
        IPv4 flow tunnel route skip             0
        ager phase1                         26541
        ager phase2                         26541
        cp sess del rexmitted                   0
Flow Table Statistics
        flow add                                0
        flow remove                             0
        flow find                       114370163
Ager Ring Statistics
        IPv4 and IPv6 Common Ager Ring: (Batching Disabled)
        current tick                        24929
        insert                              53082
        update                           10618908
        callback invoked                    53082
        remove                                  0
        update skipped                          0
        max watchdog time                    5000
        timeout value larger                    0
        ager update                             0



Flow Session Statistics
        Current flows                           0
        IPv6 flow alloc                         0
        IPv6 flow free                          0
        IPv6 flow alloc failure                 0
        IPv6 flow route skip                    0
        IPv6 flow fast reroute                  0
        IPv6 flow tunnel route skip             0
        ager phase1                             0
        ager phase2                             0
        cp sess del rexmitted                   0
Flow Table Statistics
        flow add                                0
        flow remove                             0
        flow find                               0
```

Release Information

Command introduced in Junos OS Release 18.2R3-S3.

Starting in Junos OS Release 18.2R3-S3, 19.1R3, 19.2R3, 19.3R3, 19.4R2, 20.1R2, you can now view the Packet Forwarding Engine data using the `usp flow config` and `usp flow stats` options for the `show pfe data` command.

RELATED DOCUMENTATION

Junos OS Architecture Overview

show pfe feb

IN THIS SECTION

- Syntax | 22974
- Description | 22974
- Options | 22975
- Required Privilege Level | 22975
- Output Fields | 22975
- Sample Output | 22977
- Release Information | 22983

Syntax

```
show pfe feb
```

Description

(M5 and M10 routers only) Display Packet Forwarding Engine Forwarding Engine Board (FEB) status and statistics information.

Options

This command has no options.

Required Privilege Level

admin

Output Fields

Table 973 on page 22975 lists the output fields for the `show pfe feb` command. Output fields are listed in the approximate order in which they appear.

Table 973: show pfe feb Output Fields

Field Name	Field Description

FEB status	Status of FEB:
	• Slot—FEB slot number.
	• State—State of the FEB:
	  • Offline—FEB is powered down.
	  • Online—FEB is operational and running.
	  • Check—FEB is in alarmed state where the Switch Interface Board (SIB) plane is partially operational for the following reasons:
	  • FEB is not inserted properly.
	  • Two or more links between the FEB and Packet Forwarding Engine fail.
	• Last State Change—Date and time the CFEB state last changed.
	• Uptime (total)—How long the Routing Engine has been connected to the FEB and, therefore, how long the Flexible PIC Concentrator (FPC) has been up and running.
	• Failures—Number of PFE Peer detach failures.
	• Pending—Number of messages waiting to be sent.
	• Policer Drop Probability—Current policer drop probability. The default is high, and can be configured using the `policer-drop-probability-low` statement.

Table 973: show pfe feb Output Fields *(Continued)*

Field Name	Field Description

Peer message type receive qualifiers	• Message Type—IPC Message Type. For example, interface and nexthop.
	• Receive Qualifier – Message receive qualifier for a peer (non-None only):
	  • All
	  • Only this slot
	  • Selective slot

PFE listener statistics	PFE listener statistics: • Open—Number of times a peer was opened. • Close—Number of times a peer was closed. • Sleep—Number of times a thread slept. • Wakeup—Number of times wakeup was issued. • Resync Request—Number of resync requests. • Resync Done—Number of successful resyncs. • Resync Fail—Number of failed resyncs. • Resync Time—Time the resync last happened.
PFE IPC statistics	• type—IPC Message Type. • TX Messages—Number of Tx messages. • RX Messages—Number of Rx messages.
PFE socket-buffer mbuf depth	• bucket—Bucket number. • count—Number of messages in the bucket.

Table 973: show pfe feb Output Fields *(Continued)*

Field Name	Field Description
PFE socket-buffer bytes pending transmit	• bucket—Bucket number. • count—Number of bytes pending transmit.

Sample Output

show pfe feb

```
user@host> show pfe feb

FEB status:
  Slot:                 Present
  State:                Online
  Last State Change:    2005-03-11 00:33:57 PST
  Uptime (total):       1d 09:14
  Failures:             0
  Pending:              0
..Policer Drop Probability: HIGH

Peer message type receive qualifiers:
 Message Type      Receive Qualifier
 ----------------  -----------------
              TTP  All
              IFD  All
              IFL  All
          Nexthop  All
              COS  All
            Route  All
      SW Firewall  All
      HW Firewall  All
   PFE Statistics  All
   PIC Statistics  All
         Sampling  All
       Monitoring  None
              ASP  None
             L2TP  None
```

```
        Collector  None
 PIC Configuration  All
 Queue Statistics  All
           (null)  None

PFE listener statistics:
  Open:                 1
  Close:                0
  Sleep:                0
  Wakeup:               0
  Resync Request:       0
  Resync Done:          1
  Resync Fail:          0
  Resync Time:          0

PFE IPC statistics:
 type        TX Messages  RX messages
 ----------- -----------  -----------
      Header           0            0
        Test           0            0
   Interface         639        11889
     Chassis           0            0
        Boot           0            0
    Next-hop         104            0
       Jtree           0            0

       Cprod           0            0
       Route         940            0
         Pfe        3008         1995
         Dfw           9            0
  Mastership           0            0
    Sampling           0            0
        GUCP           0            0
         CoS          35            0
        GCCP           0            0
        GHCP           0            0
        IRSD           0            0
  Monitoring           0            0
          RE           0            0
         PIC           0            0
     ASP cfg           0            0
     ASP cmd           0            0
    L2TP cfg           0            0
```

```
   Collector           0            0
   PIC state           0            0
  Aggregator           0            0
       Empty           0            0

PFE socket-buffer mbuf depth:
 bucket          count
 ------          -----
     0               0
     1               0
     2               0
     3               0
     4               0
     5               0
     6               0
     7               0
     8               0
     9               0
    10               0
    11               0
    12               0
    13               0
    14               0
    15               0
    16               0
    17               0
    18               0
    19               0
    20               0
    21               0

PFE socket-buffer bytes pending transmit:
 bucket          count
 ------          -----
     0               0
     1               0
     2               0
     3               0
     4               0
     5               0
     6               0
     7               0
     8               0
```

```
     9               0
    10               0
    11               0
    12               0
    13               0
    14               0
    15               0
    16               0
    17               0
    18               0
    19               0
    20               0
    21               0
```

show pfe feb

```
user@host> show pfe feb

FEB status:
  Slot:                 Present
  State:                Online
  Last State Change:    2005-03-11 00:33:57 PST
  Uptime (total):       1d 09:14
  Failures:             0
  Pending:              0

Peer message type receive qualifiers:
 Message Type      Receive Qualifier
 ----------------  -----------------
              TTP  All
              IFD  All
              IFL  All
          Nexthop  All
              COS  All
            Route  All
      SW Firewall  All
      HW Firewall  All
   PFE Statistics  All
   PIC Statistics  All
         Sampling  All
       Monitoring  None
```

```
               ASP  None
              L2TP  None
         Collector  None
 PIC Configuration  All
 Queue Statistics  All
            (null)  None

PFE listener statistics:
  Open:                 1
  Close:                0
  Sleep:                0
  Wakeup:               0
  Resync Request:       0
  Resync Done:          1
  Resync Fail:          0
  Resync Time:          0

PFE IPC statistics:
 type        TX Messages  RX messages
 ----------- -----------  -----------
      Header           0            0
        Test           0            0
   Interface         639        11889
     Chassis           0            0
        Boot           0            0
    Next-hop         104            0
       Jtree           0            0

       Cprod           0            0
       Route         940            0
         Pfe        3008         1995
         Dfw           9            0
  Mastership           0            0
    Sampling           0            0
        GUCP           0            0
         CoS          35            0
        GCCP           0            0
        GHCP           0            0
        IRSD           0            0
  Monitoring           0            0
          RE           0            0
         PIC           0            0
     ASP cfg           0            0
```

```
     ASP cmd           0            0
    L2TP cfg           0            0
   Collector           0            0
   PIC state           0            0
  Aggregator           0            0
       Empty           0            0

PFE socket-buffer mbuf depth:
 bucket          count
 ------          -----
     0               0
     1               0
     2               0
     3               0
     4               0
     5               0
     6               0
     7               0
     8               0
     9               0
    10               0
    11               0
    12               0
    13               0
    14               0
    15               0
    16               0
    17               0
    18               0
    19               0
    20               0
    21               0

PFE socket-buffer bytes pending transmit:
 bucket          count
 ------          -----
     0               0
     1               0
     2               0
     3               0
     4               0
     5               0
     6               0
```

```
     7              0
     8              0
     9              0
    10              0
    11              0
    12              0
    13              0
    14              0
    15              0
    16              0
    17              0
    18              0
    19              0
    20              0
    21              0
```

Release Information

Command introduced before Junos OS Release 7.4.

show pfe filter hw summary

IN THIS SECTION

- Syntax | 22984
- Description | 22984
- Required Privilege Level | 22984
- Output Fields | 22984
- Sample Output | 22985
- Release Information | 22985

Syntax

```
show pfe filter hw summary
```

Description

Display a summary of the access control list (ACL; also known as firewall filter) ternary content-addressable memory (TCAM) hardware utilization to show the allocated, used, and free TCAM entry space.

Command supported on standalone QFX Series switches, QFX5100-only (pure QFX5100) Virtual Chassis Fabric (VCF), QFX5100-only (pure QFX5100) Virtual Chassis (VC), and QFX3500-only (pure QFX3500) VC.

Required Privilege Level

view

Output Fields

Table 974 on page 22984 lists the output fields for the `show pfe filter hw summary` command. Output fields are listed in the approximate order in which they appear.

Table 974: show pfe filter hw summary Output Fields

Field Name	Field Description
`Group`	ACL ingress and egress filter groups: • iRACL group—ingress routing ACL filter group • iVACL group—ingress VLAN ACL filter group • iPACL group—ingress port ACL filter group • ePACL group—egress port ACL filter group • eVACL group—egress VLAN ACL filter group • eRACL group—egress routing ACL filter group • eRACL IPv6 group—egress IPv6 routing ACL filter group

Table 974: show pfe filter hw summary Output Fields *(Continued)*

Field Name	Field Description
Group-ID	Internal identification number of the filter group.
Allocated	Number of TCAM filter entries allocated to the filter group.
Used	Number of TCAM filter entries used by the filter group.
Free	Number of TCAM filter entries available for use by the filter group.

Sample Output

show pfe filter hw summary

```
user@switch> show pfe filter hw summary

Group                     Group-ID      Allocated      Used           Free
---------------------------------------------------------------------------
> Ingress filter groups:
   iRACL group            14            512            4              508
   iVACL group            13            512            2              510
   iPACL group            12            256            2              254
> Egress filter groups:
   ePACL group            20            256            3              253
   eVACL group            21            256            4              252
   eRACL group            22            256            245            11
   eRACL IPV6 group       24            256            3              253
```

Release Information

Command introduced in Junos OS Release 14.1X53-D10.

Command supported on EX Series in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Planning the Number of Firewall Filters to Create

show pfe fpc

IN THIS SECTION

- Syntax | 22986
- Syntax (TX Matrix and TX Matrix Plus Router) | 22986
- Syntax (MX Series Router) | 22987
- Description | 22987
- Options | 22987
- Required Privilege Level | 22988
- Output Fields | 22988
- Sample Output | 22993
- Release Information | 23013

Syntax

```
show pfe fpc slot
<detail | extensive>
```

Syntax (TX Matrix and TX Matrix Plus Router)

```
show pfe fpc
<lcc number>
```

Syntax (MX Series Router)

```
show pfe fpc slot
<detail | extensive>
<all-members>
<local>
<member member-id>
```

Description

Display Packet Forwarding Engine statistics for the specified Flexible PIC Concentrator (FPC).

> **NOTE**: Starting in Junos OS Evolved Releases 23.4R2 and 24.2R1, this command is disabled across all Junos OS Evolved platforms. Previously, the following error message was returned when executing the command: `error: the statistics-service subsystem is not running`

Options

slot FPC slot number. Replace ***slot*** with a value from **0** through **2**.

detail | extensive (Optional) Display the specified level of detail.

all-members (MX Series routers only) (Optional) Display Packet Forwarding Engine statistics for the specified FPC in all members of the Virtual Chassis configuration.

lcc *number* (TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, the slot number of the T640 router (or line-card chassis) that houses the FPC. On a TX Matrix Plus router, **lcc *number*** represents the slot number of the router (or line-card chassis) that houses the FPC.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Display Packet Forwarding Engine statistics for the specified FPC in the local Virtual Chassis member.

member ***member-id*** (MX Series routers only) (Optional) Display Packet Forwarding Engine statistics for the specified FPC in the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

Required Privilege Level

admin

Output Fields

Table 975 on page 22989 lists the output fields for the `show pfe fpc` command. Output fields are listed in the approximate order in which they appear.

Table 975: show pfe fpc Output Fields

Field Name	Field Description

FPC 1 status	Status of FPC 1:
	• Slot—FPC slot number – 1.
	• State—State of FPC1:
	  • Dead—Held in reset because of errors.
	  • Diag—Slot is being ignored while the FPC is running diagnostics.
	  • Dormant—Held in reset.
	  • Empty—No FPC is present.
	  • Online—FPC is online and running.
	  • Present—FPC is detected by the chassis daemon but either is not supported by the current version of Junos OS or is inserted in the wrong slot. The output also states either Hardware Not Supported or Hardware Not in Right Slot. The FPC is coming up but not yet online.
	  • Probed—Probe is complete; awaiting restart of the Packet Forwarding Engine (PFE).
	  • Probe-wait—Waiting to be probed.
	• Last State Change—Date and time the FPC state last changed.
	• Uptime—How long the Routing Engine has been connected to the FEB and, therefore, how long the Flexible PIC Concentrator (FPC) has been up and running.
	• Failures—Number of PFE Peer detach failures.
	• Pending—Number of messages waiting to be sent.
	• Route Memory Enhanced—Reallocation of the jtree memory on the Packet Forwarding Engine to allocate more memory for routing tables. Can be configured with the `memory-enhanced` statement.

Table 975: show pfe fpc Output Fields *(Continued)*

Field Name	Field Description
PFE listener statistics	PFE listener statistics: • Open—Number of times a peer was opened. • Close—Number of times a peer was closed. • Sleep—Number of times a thread slept. • Wakeup—Number of times wakeup was issued. • Resync Request—Number of resync requests. • Resync Done—Number of successful resyncs. • Resync Fail—Number of failed resyncs. • Resync Time—Time the resync last happened.
PFE IPC statistics	• type—IPC Message Type. • TX Messages—Number of Tx messages. • RX Messages—Number of Rx messages.
GFPC 0 status	Status of GFPC 0: • Slot—GFPC slot number – 0. • State—State of GFPC. • Last State Change—Date and time the GFPC state last changed.
Peer message type receive qualifiers [non-NONE(s) only]	• IPC Msg Type—IPC Message Type. For example, interface, nexthop. • Receive Qualifier—Message receive qualifier for a peer (non-NONE(s) only):
IFSTATE BITS SET	IFSTATE clients that have registered to receive the message types this slot is listening to.

Table 975: show pfe fpc Output Fields *(Continued)*

Field Name	Field Description
PFE listener statistics	PFE listener statistics: • Open—Number of times a peer was opened. • Close—Number of times a peer was closed. • Sleep—Number of times a thread slept. • Wakeup—Number of times wakeup was issued. • Resync Request—Number of resync requests. • Resync Done—Number of successful resyncs. • Resync Fail—Number of failed resyncs. • Resync Time—Time the resync last happened.
PFE IPC statistics	• type—IPC Message Type. • TX Messages—Number of Tx messages. • RX Messages—Number of Rx messages.
PFE socket-buffer mbuf depth	• bucket—Bucket number. • count—Number of messages in the bucket.
PFE socket-buffer bytes pending transmit	• bucket—Bucket number. • count—Number of bytes pending transmit.

Table 975: show pfe fpc Output Fields *(Continued)*

Field Name	Field Description
GFPC 2 status	Status of GFPC 2: • Slot—GFPC slot number – 2. • State—State of GFPC. • Last State Change—Date and time the GFPC state last changed. • Route Memory Enhanced—Reallocation of the jtree memory on the Packet Forwarding Engine to allocate more memory for routing tables. Can be configured with the `memory-enhanced` statement. • Filter Memory Enhanced—Reallocation of the jtree memory on the Packet Forwarding Engine to allocate more memory for firewall filters. Can be configured with the `memory-enhanced` statement.
XDPC status	XDPC status: • Slot—Present or empty. • State—Online or offline. • Last State Change—Date and time the DPC state last changed. • Uptime (total)—Length of time the DPC has been online. • Failures—Number of DPC failures. • Pending—Number of messages waiting to be sent. • Route Memory Enhanced—Reallocation of the jtree memory on the Packet Forwarding Engine to allocate more memory for routing tables. Can be configured with the `memory-enhanced` statement. • Policer Drop Probability—Current policer drop probability. The default is high, and can be configured using the `policer-drop-probability-low` statement.

Sample Output

show pfe fpc

```
user@host> show pfe fpc 1

FPC 1 status:
  Slot:                  Present
  State:                 Online
  Last State Change:     2000-01-10 18:12:27 UTC
  Uptime:                1d 03:31
  Failures:              0
  Pending:               0
  Route Memory Enhanced: 0
PFE listener statistics:
  Open:                 1
  Close:                0
  Sleep:                0
  Wakeup:               0
  Resync Request:       0
  Resync Done:          0
  Resync Fail:          0
  Resync Time:          0

PFE IPC statistics:
 type        TX Messages  RX messages
 ----------- -----------  -----------
      Header           0            0
        Test           0            0
   Interface        2251         2219
     Chassis           0            0
        Boot           0            0
    Next-hop           0            0
       Jtree           0            0
       Cprod           0            0
       Route           0            0
         Pfe           0            1
         Dfw
```

show pfe fpc lcc

```
user@host> show pfe fpc 0 lcc 0
lcc0-re0:
--------------------------------------------------------------------------
GFPC 0 status:
  Slot:                 Present
  State:                Online
  Last State Change:    2009-06-17 21:00:35 PDT
  Uptime (total):       02:31:45
  Failures:             0
  Pending:              0

Peer message type receive qualifiers [ non-NONE(s) only ]:
 IPC Msg Type (subtype)     Receive Qualifier
 ----------------  -----------------
 Interface        (0)           All
 Interface        (1)           All
 Interface        (2)           All
 Interface        (3)           All
 Interface        (4)           All
 Interface        (5)           All
 Interface        (6)           All
 Interface        (7)           All
 Interface        (8)           All
 Interface        (9)           All
 Interface        (10)          All
 Interface        (11)          All
 Interface        (12)          All
 Interface        (13)          All
 Interface        (14)          All
 Interface        (15)          All
 Interface        (16)          All
 Interface        (17)          All
 Interface        (18)          All
 Interface        (19)          All
 Interface        (20)          Slot only
 Interface        (21)          All
...
 Next-hop         (0)           All
 Next-hop         (1)           All
 Next-hop         (2)           All
```

```
 Next-hop         (3)           All
 Next-hop         (4)           All
 Next-hop         (5)           Always TRUE
...
 Route            (0)           All
 Route            (1)           All
 Route            (2)           All
 Route            (3)           All
 Route            (4)           All
 Route            (5)           All
 Route            (6)           All
 Route            (7)           All
 Route            (8)           All

...
 Pfe              (1)           Always TRUE
 Pfe              (3)           Always TRUE
 Pfe              (5)           Always TRUE
...
 Dfw              (0)           All
 Dfw              (1)           All
 Dfw              (2)           All
 Dfw              (3)           All

...
 Sampling         (1)           All
 Sampling         (2)           All
 Sampling         (3)           All
 CoS              (0)           All
 CoS              (1)           All
 CoS              (2)           All
 CoS              (3)           All
...
 PIC              (1)           Always TRUE
 PIC              (3)           Always TRUE
...
 GenCfg           (8)           All
 GenCfg           (15)          All
...
 IFSTATE BITS SET:
 ----------------
        IFD
        IFL
```

```
        IFF
        IFA
        RTTABLE
        ROUTE
        NEXTHOP
        FIREWALL
        NAME TABLE
        COS_FABRIC
        COS_POLICY
        COS_RED
        COS_REWRT_TABLE
        COS_REWRT_IFLMAP
        COS_CLASS_TABLE
        COS_CLASS_IFLMAP
        COS_POLICER
        COS_SHAPER
        SAMPLE
        RTCOS
        SYSCONF
        IFVP
        SADB
        IFVC
        COS_FC_QUEUE
        COS_FRAGMAP_TABLE
        COS_FRAGMAP_IFLMAP
        Generic config
        Mesh group

PFE listener statistics:
  Open:                 1
  Close:                0
  Sleep:                0
  Wakeup:               0
  Resync Request:       0
  Resync Done:          1
  Resync Fail:          0
  Resync Time:          0

PFE IPC statistics:
 Type (subtype)             TX Messages    RX messages
 --------------             -----------   -----------
 Interface  ( 3)                165            0
 Interface  ( 4)                 81            0
```

```
 Interface  ( 5)                 0          190
 Interface  ( 8)               145            0
 Interface  ( 9)               425            0
 Interface  (10)                24            0
...
PFE socket-buffer mbuf depth:
 bucket          count
 ------          -----
     0               0
     1               0
     2               0

PFE socket-buffer bytes pending transmit:
 bucket          count
 ------          -----
     0               0
     1               0
...
```

show pfe fpc 0 detail

```
user@host> show pfe fpc 0 detail

GFPC 2 status:
  Slot:                 Present
  State:                Online
  Last State Change:    2010-11-16 03:55:25 PST
  Uptime (total):       00:11:06
  Failures:             1
  Pending:              0
  Route Memory Enhanced: 0
 Filter Memory Enhanced: 1

Peer message type receive qualifiers [ non-NONE(s) only ]:
 IPC Msg Type (subtype)     Receive Qualifier
 ----------------  -----------------
 Interface        (0)           All
 Interface        (1)           All
 Interface        (2)           All
 Interface        (3)           All
 Interface        (4)           All
```

```
 Interface        (5)          All
 Interface        (6)          All
 Interface        (7)          All
 Interface        (8)          All
 Interface        (9)          All
 Interface        (10)         All
 Interface        (11)         All
 ...
 Next-hop         (0)          All
 Next-hop         (1)          All
 Next-hop         (2)          All
 Next-hop         (3)          All
 Next-hop         (4)          All
 Next-hop         (5)          All
...
 Route            (0)          All
 Route            (1)          All
 Route            (2)          All
 Route            (3)          All
 Route            (4)          All
 Route            (5)          All
...
 Pfe              (1)          Always TRUE
 Pfe              (3)          Always TRUE
 Pfe              (5)          Always TRUE
...
 Dfw              (0)          All
 Dfw              (1)          All
 Dfw              (2)          All
 Dfw              (3)          All
...
 Sampling         (1)          All
 Sampling         (2)          All
 Sampling         (3)          All
 CoS              (0)          All
 CoS              (1)          All
 CoS              (2)          All
 CoS              (3)          All
 CoS              (4)          All
...
 PIC              (1)          Always TRUE
 PIC              (3)          Always TRUE
...
```

```
 GenCfg           (8)           All
 GenCfg           (15)          All
...
 IFSTATE BITS SET:
 ----------------
        IFD
        IFL
        IFF
        IFA
        RTTABLE
        ROUTE
        NEXTHOP
        FIREWALL
        NAME TABLE
        COS_FABRIC
        COS_POLICY
        COS_RED
        COS_REWRT_TABLE
        COS_REWRT_IFLMAP
        COS_CLASS_TABLE
        COS_CLASS_IFLMAP
        COS_POLICER
        COS_SHAPER
        SAMPLE
        RTCOS
        SYSCONF
        IFVP
        SADB
        IFVC
        COS_FC_QUEUE
        COS_FRAGMAP_TABLE
        COS_FRAGMAP_IFLMAP
        Generic config
        Mesh group

PFE listener statistics:
  Open:                 2
  Close:                1
  Sleep:                0
  Wakeup:               0
  Resync Request:       0
  Resync Done:          2
  Resync Fail:          0
```

```
  Resync Time:          0

PFE IPC statistics:
 Type (subtype)             TX Messages    RX messages
 --------------             -----------   -----------
 Interface  ( 3)                104            0
 Interface  ( 5)                  0            8
 Interface  ( 8)                 85            0
 Interface  ( 9)                 67            0
 Interface  (10)                  4            0
...
 Next-hop   ( 1)                364            0
 Next-hop   ( 3)                 12            0
 Next-hop   (11)                 33            0
 Next-hop   (23)                 39            0
 Route      ( 1)                331            0
 Route      ( 2)                 34            0
 Route      ( 3)                  1            0
 Route      ( 6)                  1            0
 Route      ( 9)                 48            0
 Pfe        ( 1)                  0            1
 Pfe        ( 3)                  1            0
 Pfe        ( 4)                  0            1
 Pfe        ( 5)                  1            0
...
 Dfw        ( 1)                 20            0
 Dfw        (18)                  1            0
 GenCfg     ( 8)                 45            0
 GenCfg     (15)                  1            0
```

show pfe fpc 0 (MX 960 with DPC)

```
user@host> show pfe fpc 0

XDPC 0 status:
  Slot:                 Present
  State:                Online
  Last State Change:    2012-08-07 13:13:01 PDT
  Uptime (total):       21:01:41
  Failures:             0
```

```
  Pending:              0
  Route Memory Enhanced: 0
  Policer Drop Probability: HIGH

Peer message type receive qualifiers [ non-NONE(s) only ]:
 IPC Msg Type (subtype)     Receive Qualifier
 ----------------  -----------------
 Interface        (0)           All
 Interface        (1)           All
 Interface        (2)           All
 Interface        (3)           All
 Interface        (4)           All
 Interface        (5)           All
 Interface        (6)           All
 Interface        (7)           All
 Interface        (8)           All
 Interface        (9)           All
 Interface        (10)          All
 Interface        (11)          All
 Interface        (12)          All
 Interface        (13)          All
 Interface        (14)          All
 Interface        (15)          All
 Interface        (16)          All
 Interface        (17)          All
 Interface        (18)          All
 Interface        (19)          All
 Interface        (20)          Slot only
 Interface        (21)          All
 Interface        (22)          Slot only
 Interface        (23)          All
 Interface        (24)          All
 Interface        (25)          All
 Interface        (26)          All
 Interface        (27)          All
 Interface        (28)          All
 Interface        (29)          All
 Interface        (30)          All
 Interface        (31)          All
 Interface        (32)          All
 Interface        (33)          All
 Interface        (34)          All
 Interface        (35)          All
```

```
Interface        (36)          All
Interface        (37)          All
Interface        (38)          All
Interface        (39)          All
Interface        (40)          All
Interface        (41)          All
Interface        (42)          Slot only
Interface        (43)          Slot only
Interface        (44)          Slot only
Interface        (45)          All
Interface        (46)          All
Interface        (47)          All
Interface        (48)          Slot only
Interface        (49)          Slot only
Interface        (50)          Slot only
Interface        (51)          Slot only
Interface        (52)          All
Interface        (53)          All
Interface        (54)          All
Interface        (55)          All
Interface        (56)          Slot only
Interface        (57)          All
Interface        (58)          All
Interface        (59)          All
Interface        (60)          All
Interface        (61)          All
Interface        (62)          All
Interface        (63)          All
Interface        (64)          Slot only
Interface        (65)          All
Interface        (66)          All
Interface        (67)          All
Interface        (68)          All
Interface        (69)          All
Interface        (70)          All
Interface        (71)          All
Interface        (72)          All
Interface        (73)          All
Interface        (74)          All
Interface        (75)          All
Interface        (76)          Slot only
Interface        (77)          Slot only
Interface        (78)          Slot only
```

```
Interface        (79)          All
Interface        (80)          All
Interface        (81)          All
Interface        (82)          All
Interface        (83)          Slot only
Interface        (84)          All
Interface        (85)          All
Interface        (86)          All
Interface        (87)          All
Interface        (88)          All
Interface        (89)          All
Interface        (90)          All
Interface        (91)          All
Interface        (92)          All
Interface        (93)          Slot only
Interface        (94)          Slot only
Interface        (95)          Slot only
Interface        (96)          All
Interface        (97)          All
Interface        (98)          All
Interface        (99)          All
Interface        (100)         All
Interface        (101)         All
Interface        (102)         All
Interface        (103)         All
Interface        (104)         All
Interface        (105)         Slot only
Interface        (106)         Slot only
Interface        (107)         All
Interface        (108)         All
Interface        (109)         All
Interface        (110)         All
Interface        (111)         All
Interface        (112)         All
Interface        (113)         All
Interface        (114)         All
Interface        (115)         All
Interface        (116)         All
Interface        (117)         All
Interface        (118)         All
Interface        (119)         All
Interface        (120)         All
Interface        (121)         Slot only
```

```
Interface        (122)        All
Interface        (123)        All
Interface        (124)        All
Interface        (125)        Slot only
Interface        (126)        Slot only
Interface        (127)        Slot only
Interface        (128)        All
Interface        (129)        All
Interface        (130)        All
Interface        (131)        All
Interface        (132)        All
Interface        (133)        All
Interface        (134)        All
Interface        (135)        All
Interface        (138)        All
Interface        (139)        All
Interface        (142)        All
Interface        (145)        All
Interface        (146)        All
Interface        (147)        All
Interface        (148)        All
Interface        (149)        All
Interface        (150)        Slot only
Interface        (151)        All
Interface        (152)        Slot only
Interface        (153)        All
Interface        (154)        All
Interface        (155)        All
Interface        (156)        All
Interface        (157)        All
Interface        (158)        All
Interface        (159)        Slot only
Interface        (160)        All
Interface        (161)        All
Interface        (163)        All
Interface        (164)        Slot only
Interface        (165)        Slot only
Interface        (167)        All
Interface        (168)        All
Interface        (169)        All
Interface        (170)        Slot only
Interface        (171)        Slot only
Interface        (172)        All
```

```
Interface        (173)        All
Interface        (174)        All
Interface        (175)        All
Interface        (176)        All
Interface        (177)        All
Interface        (178)        All
Interface        (179)        All
Interface        (180)        All
Interface        (181)        All
Interface        (182)        All
Interface        (183)        All
Interface        (184)        All
Interface        (185)        All
Interface        (186)        All
Interface        (187)        All
Interface        (188)        All
Interface        (189)        All
Interface        (190)        All
Interface        (191)        All
Interface        (192)        All
Interface        (193)        All
Interface        (194)        All
Interface        (195)        All
Interface        (196)        All
Interface        (197)        All
Interface        (198)        All
Interface        (199)        All
Interface        (200)        All
Interface        (201)        All
Interface        (202)        All
Interface        (204)        All
Interface        (205)        All
Interface        (206)        All
Interface        (207)        All
Interface        (208)        All
Interface        (209)        All
Interface        (210)        All
Interface        (211)        All
Interface        (212)        All
Interface        (213)        All
Interface        (214)        All
Interface        (215)        All
Interface        (216)        All
```

```
Interface        (217)        All
Interface        (218)        All
Interface        (219)        All
Interface        (220)        All
Interface        (221)        All
Interface        (222)        All
Interface        (223)        All
Interface        (224)        All
Interface        (225)        All
Interface        (226)        All
Interface        (227)        All
Interface        (229)        All
Interface        (230)        All
Interface        (231)        All
Interface        (232)        All
Interface        (233)        All
Interface        (234)        All
Interface        (235)        All
Interface        (236)        All
Interface        (237)        All
Interface        (238)        All
Interface        (239)        All
Next-hop         (0)          All
Next-hop         (1)          All
Next-hop         (2)          All
Next-hop         (3)          All
Next-hop         (4)          All
Next-hop         (5)          All
Next-hop         (6)          All
Next-hop         (7)          All
Next-hop         (8)          All
Next-hop         (9)          All
Next-hop         (10)         All
Next-hop         (11)         All
Next-hop         (12)         All
Next-hop         (13)         All
Next-hop         (14)         All
Next-hop         (15)         All
Next-hop         (16)         All
Next-hop         (17)         All
Next-hop         (18)         All
Next-hop         (19)         All
Next-hop         (20)         All
```

```
Next-hop        (21)        All
Next-hop        (22)        All
Next-hop        (23)        All
Next-hop        (24)        All
Next-hop        (25)        All
Next-hop        (26)        All
Next-hop        (27)        All
Next-hop        (28)        All
Next-hop        (29)        All
Next-hop        (30)        All
Next-hop        (31)        All
Next-hop        (32)        All
Next-hop        (33)        All
Next-hop        (34)        All
Next-hop        (35)        All
Next-hop        (36)        All
Next-hop        (37)        All
Next-hop        (39)        Always TRUE
Next-hop        (40)        All
Next-hop        (41)        All
Next-hop        (42)        All
Next-hop        (43)        All
Route           (0)         All
Route           (1)         All
Route           (2)         All
Route           (3)         All
Route           (4)         All
Route           (5)         All
Route           (6)         All
Route           (7)         All
Route           (8)         All
Route           (9)         All
Route           (10)        All
Route           (11)        All
Route           (12)        All
Route           (13)        All
Route           (14)        All
Route           (15)        All
Route           (16)        All
Route           (17)        All
Route           (18)        All
Route           (19)        All
Route           (20)        All
```

```
Route           (22)        All
Route           (23)        All
Route           (24)        All
Route           (25)        All
Route           (26)        All
Route           (27)        All
Route           (28)        All
Route           (29)        Always TRUE
Route           (30)        Always TRUE
Pfe             (1)         Always TRUE
Pfe             (3)         Always TRUE
Pfe             (5)         Always TRUE
Pfe             (7)         Always TRUE
Pfe             (10)        Always TRUE
Pfe             (11)        Always TRUE
Pfe             (12)        Always TRUE
Pfe             (13)        Always TRUE
Pfe             (14)        Always TRUE
Pfe             (15)        Always TRUE
Pfe             (35)        Always TRUE
Dfw             (0)         All
Dfw             (1)         All
Dfw             (2)         All
Dfw             (3)         All
Dfw             (4)         All
Dfw             (5)         All
Dfw             (6)         All
Dfw             (7)         All
Dfw             (8)         All
Dfw             (9)         All
Dfw             (10)        All
Dfw             (11)        All
Dfw             (12)        All
Dfw             (13)        All
Dfw             (14)        All
Dfw             (18)        All
Dfw             (19)        All
Sampling        (1)         All
Sampling        (2)         All
Sampling        (3)         All
CoS             (0)         All
CoS             (1)         All
CoS             (2)         All
```

```
CoS             (3)          All
CoS             (4)          All
CoS             (5)          All
CoS             (6)          All
CoS             (7)          All
CoS             (8)          All
CoS             (9)          All
CoS             (10)         All
CoS             (11)         All
CoS             (12)         All
CoS             (13)         All
CoS             (14)         All
CoS             (15)         All
CoS             (16)         All
CoS             (17)         All
CoS             (18)         All
CoS             (19)         All
CoS             (20)         All
CoS             (21)         All
CoS             (22)         All
CoS             (23)         All
CoS             (27)         All
CoS             (29)         All
CoS             (31)         All
CoS             (32)         All
PIC             (1)          Always TRUE
PIC             (3)          Always TRUE
PIC             (5)          Always TRUE
PIC             (7)          Always TRUE
PIC             (10)         Always TRUE
PIC             (11)         Always TRUE
PIC             (12)         Always TRUE
PIC             (13)         Always TRUE
PIC             (14)         Always TRUE
PIC             (15)         Always TRUE
GenCfg          (2)          All
GenCfg          (4)          All
GenCfg          (5)          All
GenCfg          (6)          All
GenCfg          (8)          All
GenCfg          (9)          All
GenCfg          (10)         All
GenCfg          (15)         All
```

```
GenCfg          (17)          All
GenCfg          (24)          All
GenCfg          (27)          All
GenCfg          (29)          All
GenCfg          (31)          All
STP             (1)           All
BD              (0)           All
BD              (1)           All
BD              (2)           All
IFSTATE BITS SET:
----------------
       IFD
       IFL
       IFF
       IFA
       RTTABLE
       ROUTE
       NEXTHOP
       FIREWALL
       NAME TABLE
       COS_FABRIC
       COS_POLICY
       COS_RED
       COS_REWRT_TABLE
       COS_REWRT_IFLMAP
       COS_CLASS_TABLE
       COS_CLASS_IFLMAP
       COS_POLICER
       COS_SHAPER
       SAMPLE
       RTCOS
       SYSCONF
       IFVP
       SADB
       IFVC
       COS_FC_QUEUE
       COS_FRAGMAP_TABLE
       COS_FRAGMAP_IFLMAP
       Generic config
       STP
       Mesh group
       Bridge Domain
       IFBD
```

```
PFE listener statistics:
  Open:                 1
  Close:                0
  Sleep:                0
  Wakeup:               0
  Resync Request:       0
  Resync Done:          1
  Resync Fail:          0
  Resync Time:          0

PFE IPC statistics:
 Type (subtype)             TX Messages    RX messages
 --------------             -----------   -----------
 Interface  ( 3)                131            0
 Interface  ( 5)                  0          379
 Interface  ( 9)                 48            0
 Interface  (10)                102            0
 Interface  (11)                  1            0
 Interface  (12)                204            0
 Interface  (13)                177            0
 Interface  (15)                 90            0
 Interface  (23)                 49            0
 Interface  (24)                  8            0
 Interface  (29)                 27            0
 Interface  (30)                 11            0
 Interface  (33)                101            0
 Interface  (34)                101            0
 Interface  (35)                 84            0
 Interface  (36)                 18            0
 Interface  (37)                 38            0
 Interface  (39)                  0            1
 Interface  (53)                  0          379
 Interface  (54)                620            0
 Interface  (55)               2064            0
 Interface  (56)                  0          379
 Interface  (57)                 57            0
 Interface  (58)                  1            0
 Interface  (90)                  0           21
 Interface  (91)                  0           13
 Interface  (92)                  0           12
 Interface  (117)                 0         1516
 Interface  (138)                 0          758
```

```
Interface  (151)          244          0
Interface  (163)          124          0
Interface  (201)          101          0
Interface  (226)           91          0
Interface  (229)          124          0
Interface  (238)          205          0
Next-hop   ( 1)           159          0
Next-hop   ( 2)             5          0
Next-hop   ( 3)            16          0
Next-hop   (11)            51          0
Next-hop   (23)            12          0
Next-hop   (40)             3          0
Route      ( 1)           164          0
Route      ( 2)            70          0
Route      ( 3)            11          0
Route      ( 6)             1          0
Route      ( 9)            14          0
Route      (12)             2          0
Route      (13)             1          0
Route      (22)             4          0
Pfe        ( 1)             0          1
Pfe        ( 3)           157          0
Pfe        ( 4)             0        157
Pfe        ( 5)           158          0
Pfe        ( 6)             0        158
Pfe        ( 7)           158          0
Pfe        ( 8)             0        158
Pfe        ( 9)             0          1
Pfe        (10)             1          0
Pfe        (11)             1          0
Pfe        (12)          2772          0
Pfe        (13)           108        108
Pfe        (15)           158          0
Pfe        (16)             0        158
Pfe        (47)             0          1
Dfw        ( 1)            23          0
Dfw        ( 2)             1          0
Dfw        ( 6)             0          6
Dfw        (18)           175          0
GenCfg     ( 5)             1          0
GenCfg     ( 8)           157          0
GenCfg     ( 9)            21          0
GenCfg     (15)            57          0
```

```
STP        ( 1)                112            0
STP        ( 2)                  0           98
STP        ( 5)                  0           97
```

Release Information

Command introduced before Junos OS Release 7.4.

show pfe next-hop

IN THIS SECTION

- Syntax | 23013
- Syntax (TX Matrix and TX Matrix Plus Routers) | 23013
- Description | 23014
- Options | 23014
- Required Privilege Level | 23015
- Output Fields | 23015
- Sample Output | 23016
- Release Information | 23019

Syntax

```
show pfe next-hop <interface interface-name>
```

Syntax (TX Matrix and TX Matrix Plus Routers)

```
show pfe next-hop <fpc slot> <interface interface-name> <lcc number>
```

Description

Display Packet Forwarding Engine next-hop information.

Options

none

Display all Packet Forwarding Engine next-hop information.

fpc *slot*

(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.

- On a TX Matrix router, if you specify the number of a T640 router by using the `lcc`*number* option (the recommended method), replace *slot* with a value from `0` through `7`. Otherwise, replace *slot* with a value from `0` through `31`.
- On a TX Matrix Plus router, if you specify the number of a T1600 router by using the `lcc`*number* option (the recommended method), replace *slot* with a value from `0` through `7`. Otherwise, replace *slot* with a value from `0` through `31`.
- On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the `lcc`*number* option (the recommended method), replace *slot* with a value from `0` through `7`. Otherwise, replace *slot* with a value from `0` through `63`.

For example, the following commands have the same result:

```
user@host> show pfe next-hop fpc 1 lcc 1
user@host> show pfe next-hop fpc 9
```

interface *interface-name*

(Optional) Display the Packet Forwarding Engine next-hop interface.

lcc *number*

(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine next-hop interface for a specific T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine next-hop interface for the router (or line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.

- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

Required Privilege Level

admin

Output Fields

Table 976 on page 23015 lists the output fields for the `show pfe next hop` command. Output fields are listed in the approximate order in which they appear.

Table 976: show pfe next-hop Output Fields

Field Name	Field Description
`ID`	The next-hop ID for the entry.
`Type`	The next-hop type for the entry.
`Interface`	The interface to which the next-hop entry is assigned.
`Protocol`	The protocol type for the next-hop entry.
`Encap`	Encapsulation type for the next-hop entry.
`Next Hop Addr`	Next-hop address for the next-hop entry.
`MTU`	MTU value for the nexthop entry.

Sample Output

show pfe next-hop

```
user@host> show pfe next-hop
Nexthop Info:
   ID      Type      Interface   Protocol         Encap  Next Hop Addr    MTU
-----  --------  -------------  ---------  ------------  ---------------  -----
    4     Mcast  -                   IPv4             -  0.0.0.0             0
    5     Bcast  -                   IPv4             -  -                   0
    7   Discard  -                   IPv4             -  -                   0
    8  MDiscard  -                   IPv4             -  -                   0
    9    Reject  -                   IPv4             -  -                   0
   13     Local  -                   IPv4             -  192.168.4.60        0
   14   Resolve  fxp0.0              IPv4   Unspecified  -                   0
   17     Local  -                   IPv4             -  127.0.0.1           0
   18   Unicast  fxp0.0              IPv4   Unspecified  192.168.4.254       0
   21     Local  -                   IPv4             -  11.1.0.1            0
   22   Unicast  at-0/1/0.0          IPv4      ATM SNAP  11.1.0.2          4482
 ...
```

show pfe next-hop fpc (TX Matrix Router)

```
user@host> show pfe next-hop fpc 1
Slot 1
Nexthop Info:
   ID      Type      Interface    Next Hop Addr    Protocol       Encap     MTU
-----  --------  -------------  ---------------  ----------  ------------  ----
    5     Mcast  -              default                IPv4             -     0
    6     Bcast  -              -                      IPv4             -     0
    8   Discard  -              -                      IPv4             -     0
    9  MDiscard  -              -                      IPv4             -     0
   13     Mcast  -              default                IPV6             -     0
   17  MDiscard  -              -                      IPV6             -     0
   18    Reject  -              -                      IPV6             -     0
   24   Discard  -              -                      None             -     0
   68     Local  -              192.168.66.113         IPv4             -     0
   69   Resolve  fxp0.0         -                      IPv4   Unspecified     0
   70   Unicast  fxp0.0         192.168.71.254         IPv4   Unspecified     0
  256     Local  -              10.71.71.1             IPv4             -     0
```

```
  257     Local  -              127.0.0.1              IPv4             -     0
  258     Mcast.local..1      default                IPv4   Unspecified     0
  259     Bcast.local..1      -                      IPv4   Unspecified     0
  261   Discard.local..1      -                      IPv4   Unspecified     0
  262  MDiscard.local..1      -                      IPv4   Unspecified     0
  269     Mcast.local..1      default                IPV6   Unspecified     0
  271   Discard.local..1      -                      IPV6   Unspecified     0
...
```

show pfe next-hop fpc (TX Matrix Plus Router)

```
user@host> show pfe next-hop fpc 0

Slot 0

   ID      Type      Interface    Next Hop Addr    Protocol       Encap     MTU
-----  --------  -------------  ---------------  ----------  ------------  ----
   31     Mcast  -              default                IPv4             -     0
   32     Bcast  -              -                      IPv4             -     0
   34   Discard  -              -                      IPv4             -     0
   35  MDiscard  -              -                      IPv4             -     0
   36    Reject  -              -                      IPv4             -     0
   39     Mcast  -              default                IPv6             -     0
   42   Discard  -              -                      IPv6             -     0
   43  MDiscard  -              -                      IPv6             -     0
   44    Reject  -              -                      IPv6             -     0
   49   Receive  -              -                      MPLS             -     0
   50   Discard  -              -                      MPLS             -     0
  111     Mcast  .local..1      default                IPv4   Unspecified     0
  112     Bcast  .local..1      -                      IPv4   Unspecified     0
  114   Discard  .local..1      -                      IPv4   Unspecified     0
  115  MDiscard  .local..1      -                      IPv4   Unspecified     0
  116    Reject  .local..1      -                      IPv4   Unspecified     0
  119     Mcast  .local..1      default                IPv6   Unspecified     0
  122   Discard  .local..1      -                      IPv6   Unspecified     0
  123  MDiscard  .local..1      -                      IPv6   Unspecified     0
  124    Reject  .local..1      -                      IPv6   Unspecified     0
  191     Mcast  .local..2      default                IPv4   Unspecified     0
  192     Bcast  .local..2      -                      IPv4   Unspecified     0
  194   Discard  .local..2      -                      IPv4   Unspecified     0
  195  MDiscard  .local..2      -                      IPv4   Unspecified     0
```

```
196    Reject  .local..2      -                      IPv4   Unspecified     0
322     Local  -              10.1.0.5               IPv4             -     0
323   Resolve  bcm0.0         -                      IPv4   Unspecified     0
326     Local  -              129.0.0.5              IPv4             -     0
327   Resolve  bcm0.0         -                      IPv4   Unspecified     0
328     Local  -              fe80::201:ff:fe01:5
                                                     IPv6             -     0
329   Receive  bcm0.0         ff02::1:ff01:5         IPv6   Unspecified     0
330   Receive  bcm0.0         fe80::                 IPv6   Unspecified     0
331   Resolve  bcm0.0         -                      IPv6   Unspecified     0
332     Local  -              fec0::a:1:0:5          IPv6             -     0
333   Receive  bcm0.0         ff02::1:ff00:5         IPv6   Unspecified     0
334   Receive  bcm0.0         fec0::                 IPv6   Unspecified     0
335   Resolve  bcm0.0         -                      IPv6   Unspecified     0
348     Local  -              192.168.178.4          IPv4             -     0
349   Resolve  em0.0          -                      IPv4   Unspecified     0
350   Unicast  em0.0          192.168.178.126        IPv4   Unspecified     0
357     Local  -              fe80::201:1ff:fe01:5
                                                     IPv6             -     0
512     Local  -              10.255.178.11          IPv4             -     0
513     Local  -              127.0.0.1              IPv4             -     0
515     Local  -              abcd::10:255:178:11
                                                     IPv6             -     0
516     Local  -              fe80::200:ff:fe00:0
                                                     IPv6             -     0
517     Local  -              127.0.0.1              IPv4             -     0
518     Mcast  .local..3      default                IPv4   Unspecified     0
519     Bcast  .local..3      -                      IPv4   Unspecified     0
521   Discard  .local..3      -                      IPv4   Unspecified     0
522  MDiscard  .local..3      -                      IPv4   Unspecified     0
523    Reject  .local..3      -                      IPv4   Unspecified     0
531     Mcast  .local..3      default                IPv6   Unspecified     0
533   Discard  .local..3      -                      IPv6   Unspecified     0
534  MDiscard  .local..3      -                      IPv6   Unspecified     0
535    Reject  .local..3      -                      IPv6   Unspecified     0
539    Mgroup  -              -                      IPv4             -     0
540     Bcast  ge-15/0/3.0    -                      IPv4      Ethernet     0
541   Receive  ge-15/0/3.0    14.2.1.0               IPv4      Ethernet     0
542     Local  -              14.2.1.1               IPv4             -     0
543   Resolve  ge-15/0/3.0    -                      IPv4      Ethernet     0
544     Bcast  ge-31/0/4.0    -                      IPv4      Ethernet     0
545   Receive  ge-31/0/4.0    14.1.1.0               IPv4      Ethernet     0
546     Local  -              14.1.1.1               IPv4             -     0
```

```
547   Resolve  ge-31/0/4.0    -                    IPv4      Ethernet     0
548   Unicast  ge-31/0/4.0    14.1.1.2             IPv4      Ethernet     0
549   Unicast  ge-15/0/3.0    14.2.1.2             IPv4      Ethernet     0
550     Bcast  ae1.0          -                    IPv4      Ethernet     0
551   Receive  ae1.0          11.1.1.0             IPv4      Ethernet     0
552     Local  -              11.1.1.1             IPv4             -     0
553   Resolve  ae1.0          -                    IPv4      Ethernet     0
554   Aggreg.  ae1.0          -                    IPv4      Ethernet     0
555   Unicast  ge-23/0/8.0    11.1.1.2             IPv4      Ethernet     0
556   Unicast  ge-7/0/9.0     11.1.1.2             IPv4      Ethernet     0
557   Aggreg.  ae1.0          -                    MPLS      Ethernet     0
558   Unicast  ge-23/0/8.0    -                    MPLS      Ethernet     0
559   Unicast  ge-7/0/9.0     -                    MPLS      Ethernet     0
560   Aggreg.  ae1.0          -                    MPLS      Ethernet     0
561   Unicast  ge-23/0/8.0    -                    MPLS      Ethernet     0
562   Unicast  ge-7/0/9.0     -                    MPLS      Ethernet     0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Routing Matrix with TXP-T1600 Configuration
- Routing Matrix with TXP-T1600-3D Configuration
- Routing Matrix with TXP-T4000-3D Configuration
- Routing Matrix with a TXP-Mixed-LCC-3D Configuration

show pfe route

IN THIS SECTION

- Syntax | 23020
- Syntax (EX Series Switches) | 23020
- Syntax (QFX Series) | 23020

- Syntax (MX Series, PTX10001-36MR, PTX10004, PTX10008, PTX10016 Devices) | **23021**
- Syntax (TX Matrix and TX Matrix Plus Routers) | **23021**
- Description | **23021**
- Options | **23021**
- Required Privilege Level | **23023**
- Output Fields | **23023**
- Sample Output | **23026**
- Release Information | **23033**

Syntax

```
show pfe route
<<inet6 | ip | iso> <prefix prefix> | <table <table-name>   <index index> <prefix prefix>>>
<mpls>
<summary>
```

Syntax (EX Series Switches)

```
show pfe route
<<inet6 | ip> <prefix prefix> | <table <table-name>   <index index> <prefix prefix>>>
<mpls>
<summary>
```

Syntax (QFX Series)

```
show pfe route
<<inet6 | ip> <prefix prefix> | <table <table-name>   <index index> <prefix prefix>> <hw (host |
lpm | multicast)>>
<<clnp> <prefix prefix> | <table <table-name>   <index index> <prefix prefix>>>
<mpls>
<summary> <hw>
```

Syntax (MX Series, PTX10001-36MR, PTX10004, PTX10008, PTX10016 Devices)

```
show pfe route
<<inet6 | ip> <prefix prefix> | <table <table-name>   <index index> <prefix prefix>>>
<dhcp>
<mpls>
<summary>
```

Syntax (TX Matrix and TX Matrix Plus Routers)

```
show pfe route
<fpc slot>
<<inet6 | ip | iso> <prefix prefix> | <table <table-name>   <index index> <prefix prefix>>>
<lcc number>
<mpls>
<summary>
```

Description

Display the routes in the Packet Forwarding Engine forwarding table. The Packet Forwarding Engine forwards packets between input and output interfaces.

> **NOTE**: The Routing Engine maintains a master copy of the forwarding table. It copies the forwarding table to the Packet Forwarding Engine, which is the part of the router or switch responsible for forwarding packets. To display the routes in the Routing Engine forwarding table, use the `show route forwarding` table command. For more information, see the CLI Explorer.

Options

none Display all Packet Forwarding Engine forwarding table information.

clnp (Optional) Show International Standards Organization (ISO) connectionless-mode network protocol (CLNP) route table information.

dhcp (Optional) Display Packet Forwarding Engine DHCP-Snooping route table information.

fpc *slot*

(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.

- On a TX Matrix router, if you specify the number of a T640 router by using the `lcc` *number* option (the recommended method), replace *slot* with a value from `0` through `7`. Otherwise, replace *slot* with a value from `0` through `31`.
- On a TX Matrix Plus router, if you specify the number of a T1600 router by using the `lcc` *number* option (the recommended method), replace *slot* with a value from `0` through `7`. Otherwise, replace *slot* with a value from `0` through `31`
- On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the `lcc` *number* option (the recommended method), replace *slot* with a value from `0` through `7`. Otherwise, replace *slot* with a value from `0` through `63`.

For example, the following commands have the same result:

```
user@host> show pfe route fpc 1 lcc 1
user@host> show pfe route fpc 9
```

host

(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display host routes installed in the on-chip hardware table.

hw

(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display routes installed in the on-chip hardware table (as opposed to displaying routes from the routing table and the PFE forwarding table before they are installed in the hardware).

index *index*

(Optional) Display table index.

inet6

(Optional) Display Packet Forwarding Engine IPv6 routes.

ip

(Optional) Display Packet Forwarding Engine IPv4 routes.

iso

(Optional) Display ISO version routing tables.

lcc *number*

(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, the slot number of the T640 router (or line-card chassis) that houses the FPC. On a TX Matrix Plus router, the slot number of the router (line-card chassis) that houses the FPC.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

lpm (QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display longest prefix match (LPM) routes installed in the on-chip hardware table.

mpls (Optional) Display Packet Forwarding Engine MPLS information.

multicast (QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display multicast routes installed in the on-chip hardware table.

prefix *prefix* (Optional) IPv4 or IPv6 prefix for which to show table entries.

summary (Optional) Display summary of Packet Forwarding Engine information. Additionally, displays total number of errors per route table per FPC slot on PTX series devices (PTX10001-36MR, PTX10004, PTX10008, PTX10016).

table <*table-name*> (Optional) Display table information.

Required Privilege Level

admin

Output Fields

Table 977 on page 23024 lists the output fields for the `show pfe route` command. Output fields are listed in the approximate order in which they appear.

Table 977: show pfe route Output Fields

Field Name	Field Description
`Destination`	Destination address for the entry.
`NH IP Addr`	Next-hop IP address for the entry.
`Type`	Next-hop type for the entry
`NH ID`	Next-hop ID for the entry
`Encap`	Encapsulation type for the next-hop entry.
`Interface`	Interface to which the next-hop entry is assigned.

Table 978 on page 23024 lists the output fields for the QFX Series `show pfe route` hardware table (`hw`) commands. Output fields are listed in the approximate order in which they appear.

Table 978: QFX Series, EX4600 switches, and OCX Series show pfe route Hardware Table Output Fields

Field Name	Field Description
`Max`	Maximum routing entries per route type.
`Used`	Number of routing entries consumed per route type.
`Free`	Number of unused routing entries per route type.
`% Free`	Percentage of unused routing entries per route type.
`Rtt`	Internal routing engine index number of the route table.
`VRF`	Internal hardware index number for the corresponding route table.

Table 978: QFX Series, EX4600 switches, and OCX Series show pfe route Hardware Table Output Fields *(Continued)*

Field Name	Field Description
`Destination`	Destination address for the entry.
`Type`	(`show pfe route summary hw`)—Route type for the entry: IPv4 or IPv6 route, and host, LPM, or multicast route. (`show pfe route (ip | inet6) hw`)—Next-hop type for the entry.
`NH ID`	Next-hop ID for the entry
`Interface`	Interface to which the next-hop entry is assigned.
`HW NH-ID`	Internal hardware index number of the next-hop.
`Src-MAC-Address`	Source MAC address.
`Port`	Port number.
`Dst-MAC-Address`	Destination MAC address.
`VLAN`	ID of the multicast group VLAN.
`GROUP`	Internal hardware index number of the multicast group next-hop.
`CLASS`	Internal class number of the multicast group.

Sample Output

show pfe route ip

```
user@host> show pfe route ip

IPv4 Route Table 0, default.0, 0x0:
Destination                       NH IP Addr      Type     NH ID Interface
--------------------------------- --------------- -------- ----- ---------
default                                            Discard     8
127.0.0.1                         127.0.0.1          Local   256
172.16/12                         192.168.71.254   Unicast    68 fxp0.0
192.168.0/18                      192.168.71.254   Unicast    68 fxp0.0
192.168.40/22                     192.168.71.254   Unicast    68 fxp0.0
192.168.64/18                     192.168.71.254   Unicast    68 fxp0.0
192.168.64/21                                      Resolve    67 fxp0.0
192.168.71.249                    192.168.71.249     Local    66
192.168.220.0/30                                   Resolve   303 fe-0/0/0.0
192.168.220.0                     192.168.220.0    Receive   301 fe-0/0/0.0
224.0.0.1                                            Mcast     5
255.255.255.255                                      Bcast     6

...
```

show pfe route iso

```
user@host# show pfe route iso

CLNS Route Table 0, CLNP.0, 0x0:
Destination                              Type     NH ID Interface
---------------------------------------- -------- ----- ---------
default                                    Reject    60
47.0005.80ff.f800.0000.0108.0001.0102.5508.2159/152     Local   514
49.0001.00a0.c96b.c491/72                   Local   536
```

show pfe route lcc summary (TX Matrix Router)

```
user@host> show pfe route lcc 2 summary

Slot 0

IPv4 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default           43        3081
1                  4         281

MPLS Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            1          68

IPV6 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            9         717
1                  5         389

Slot 1


IPv4 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default           43        3081
1                  4         281

MPLS Route Tables:

Index         Routes     Size(b)
--------  ----------  ----------
Default            1          68

IPV6 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            9         717
```

```
1                   5         389

Slot 16

IPv4 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default           41        2938
1                  4         281

MPLS Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            1          68

IPV6 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            9         717
1                  5         389

Slot 17

IPv4 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default           41        2938
1                  4         281


MPLS Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            1          68

IPV6 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            9         717
1                  5         389
```

show pfe route lcc summary (TX Matrix Plus Router)

```
user@host> show pfe route lcc 2 summary

Slot 0



IPv4 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default           25        2266
1                  9         815
2                  6         545
3                  5         453
4                 15        1371
5                  5         453
6                 13        1187

MPLS Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            1          88
4                  5         452

IPv6 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            7         697
1                 13        1305
3                  4         385
4                  4         385
5                  4         385
6                 18        1833

Slot 6



IPv4 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default           25        2266
1                  9         815
```

```
2                   6         545
3                   5         453
4                  15        1371
5                   5         453
6                  13        1187

MPLS Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            1          88
4                  5         452

IPv6 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            7         697
1                 13        1305
3                  4         385
4                  4         385
5                  4         385
6                 18        1833
...
```

show pfe route summary (MX Series Router)

```
user@host> show pfe route summary

Slot 0


DHCP-Snooping Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            1         144

IPv4 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default           25        2266
1                  9         815
2                  6         545
```

```
3                   5         453
4                  15        1371
5                   5         453
6                  13        1187

MPLS Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            1          88
4                  5         452

IPv6 Route Tables:
Index         Routes     Size(b)
--------  ----------  ----------
Default            7         697
1                  13        1305
3                   4         385
4                   4         385
5                   4         385
6                  18        1833

...
```

show pfe route summary hw (QFX Series, EX4600 Switches, OCX Series)

```
user@switch> show pfe route summary hw
Slot 0
Unit: 0
Profile active: l2-profile-three
Type            Max       Used      Free      % free
---------------------------------------------------
IPv4 Host       8192      103       8073      98.55
IPv4 LPM        16384     9         16369     99.91
IPv4 Mcast      4096      2         4037      98.56

IPv6 Host       4096      6         4037      98.56
IPv6 LPM(< 64)  8192      3         8185      99.91
IPv6 LPM(> 64)  256          1       255          99.61
IPv6 Mcast      2048      0         2019      98.58
```

show pfe route ip hw host (QFX Series)

```
user@switch> show pfe route ip hw host
Slot 0
Unit: 0
IPv4 Host entries present: 103
Rtt  VRF  Destination                      Type     NH-ID     Interface                HW NH-
ID  Src-MAC-Address   Port Dst-MAC-Address
-----------------------------------------------------------------------------------------------
-----------------------------------------
4    3    255.255.255.255                  Bcast    1695      .local.    .4   ifl 550
100003    00:00:00:01:02:03 127  00:00:00:01:02:03
0    1    200.1.1.42                       Unicast  1743      et-0/1/1   .0   ifl 559
100268    84:18:88:de:96:fd 53   00:00:00:21:12:23
0    1    200.1.1.56                       Unicast  1743      et-0/1/1   .0   ifl 559
100268    84:18:88:de:96:fd 53   00:00:00:21:12:23
0    1    200.1.1.61                       Unicast  1743      et-0/1/1   .0   ifl 559
100268    84:18:88:de:96:fd 53   00:00:00:21:12:23
0    1    11.1.1.2                         Unicast  1743      et-0/1/1   .0   ifl 559
100268    84:18:88:de:96:fd 53   00:00:00:21:12:23
0    1    200.1.1.73                       Unicast  1743      et-0/1/1   .0   ifl 559
100268    84:18:88:de:96:fd 53   00:00:00:21:12:23
0    1    200.1.1.76                       Unicast  1743      et-0/1/1   .0   ifl 559
100268    84:18:88:de:96:fd 53   00:00:00:21:12:23
0    1    200.1.1.18                       Unicast  1743      et-0/1/1   .0   ifl 559
100268    84:18:88:de:96:fd 53   00:00:00:21:12:23
0    1    200.1.1.5                        Unicast  1743      et-0/1/1   .0   ifl 559
100268    84:18:88:de:96:fd 53   00:00:00:21:12:23
0    1    200.1.1.23                       Unicast  1743      et-0/1/1   .0   ifl 559
100268    84:18:88:de:96:fd 53   00:00:00:21:12:23
0    1    101.1.1.255                      Bcast    1664      ae0        .0   ifl 544
100003    00:00:00:01:02:03 127  00:00:00:01:02:03
0    1    200.1.1.40                       Unicast  1743      et-0/1/1   .0   ifl 559
100268    84:18:88:de:96:fd 53   00:00:00:21:12:23
0    1    200.1.1.58                       Unicast  1743      et-0/1/1   .0   ifl 559
100268    84:18:88:de:96:fd 53   00:00:00:21:12:23. . .
. . .
```

show pfe route summary (PTX10001-36MR, PTX10004, PTX10008, PTX10016 Devices)

```
user@host> show pfe route summary

Slot 0

IPv4 Route Tables:
Index         Routes     Size(b)  Prefixes     Aggr     Installed   Comp(%)   Errors
--------  ----------  ----------  ---------  ---------  ----------  ------   --------
Default           25        2600         22          0          22      0         2
1                  0           0          0          0           0      0         0
50                 5         520          5          0           5      0         1
36738              5         520          5          0           5      0         1
```

Release Information

Command introduced before Junos OS Release 7.4.

Command option `hw` introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Command support introduced in Junos OS Evolved Release 22.4R1 for the PTX Series.

RELATED DOCUMENTATION

- Routing Matrix with TXP-T1600 Configuration
- Routing Matrix with TXP-T1600-3D Configuration
- Routing Matrix with TXP-T4000-3D Configuration
- Routing Matrix with a TXP-Mixed-LCC-3D Configuration

show pfe sFlow egress-mcast-stats

IN THIS SECTION

- Syntax | 23034
- Description | 23034

- Required Privilege Level | 23034
- Output Fields | 23034
- Sample Output | 23035
- Release Information | 23036

Syntax

```
show pfe sFlow egress-mcast-stats
```

Description

Displays sFlow egress multicast sampling related statistics from the packet forwarding Engine. Dropped samples are the samples exceeding the maximum replication rate. When the router reaches 1990 replications of packets per second with the maximum being 2000 replications along with one packet for the multicast group with 20 replications (20 access side ports part of the multicast group), then since the total is greater than 2000 replications, this packet from the router will be dropped and the dropped samples will be incremented by 1.

When the router samples both unicast and multicast traffic with 1990 unicast and 10 multicast packets, in this case, we will replicate 10 samples with 5 (5 access side ports are part of that group), and the collector will receive 2040 samples.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show pfe sFlow egress-mcast-stats` command.

Table 979: show pfe sFlow egress-mcast-stats Output Fields

Field Name	Field Description
`Slot`	FPC slot.

Table 979: show pfe sFlow egress-mcast-stats Output Fields *(Continued)*

Field Name	Field Description
Rate per sec	Maximum replication rate per FPC. The default value is 2000 packets allowed to be replicated per second.
Dropped samples	Total samples exceeding maximum replication rate and dropped.
Replicated samples	Total number of samples replicated per FPC.
Received samples	Total number of samples received for replication.

Sample Output

show pfe sflow egress-mcast-stats

```
user@host> show pfe sflow egress-mcast-stats
```

```
Slot 0 
 
 
```

Release Information

Command introduced in Junos OS Release 21.2R1.

show pfe ssb

IN THIS SECTION

- Syntax | 23036
- Description | 23036
- Options | 23036
- Required Privilege Level | 23036
- Output Fields | 23037
- Sample Output | 23041
- Release Information | 23044

Syntax

```
show pfe ssb
```

Description

(M20 routers only) Display Packet Forwarding Engine System and Switch Board (SSB) status and statistics information.

Options

This command has no options.

Required Privilege Level

admin

Output Fields

Table 980 on page 23037 lists the output fields for the `show pfe ssb` command. Output fields are listed in the approximate order in which they appear.

Table 980: show pfe ssb Output Fields

Field Name	Field Description
`Uptime (total)`	SSB uptime.
`Failures`	Number of failures .
`Pending`	Number of pending.
`Peer message type receive qualifiers`	Information about Peer message type receive qualifiers.
`Message Type`	Peer message type.
`Receive Qualifier`	Peer receive qualifier.
`TTP`	Peer message type TTP.
`IFD`	Peer message type IFD.
`IFL`	Peer message type IFL.
`Nexthop`	Peer message type Nexthop.
`COS`	Peer message type COS.
`Route`	Peer message type Route.
`SW Firewall`	Peer message type SW Firewall.

Table 980: show pfe ssb Output Fields *(Continued)*

Field Name	Field Description
`HW Firewall`	Peer message type HW Firewall.
`PFE Statistics`	Peer message type PFE Statistics.
`PIC Statistics`	Peer message type PIC Statistics.
`Sampling`	Peer message type Sampling .
`Monitoring`	Peer message type Monitoring.
`ASP`	Peer message type ASP.
`L2TP`	Peer message type L2TP.
`Collector`	Peer message type Collector.
`PIC Configuration`	Peer message type PIC Configuration.
`Queue Statistics`	Peer message type Queue Statistics.

Table 980: show pfe ssb Output Fields *(Continued)*

Field Name	Field Description
`PFE Listener statistics`	Information about Packet Forwarding Engine listener statistics: • `Open`—Number of PFE listeners in the “open” state. • `Close`—Number of PFE listeners in the “close” state. • `Sleep`—Number of PFE listeners in the “sleep” state. • `Wakeup`—Number of PFE listeners in the “wakeup” state. • `Resync Request`—Number of PFE listeners in the “resync request” state. • `Resync Done`—Number of PFE listeners in the “resync done” state. • `Resync Fail`—Number of PFE listeners in the “resync fail” state • `Resync Time`—Number of PFE listeners in the resync time state.

Table 980: show pfe ssb Output Fields *(Continued)*

Field Name	Field Description

`PFE IPC statistics`	Information about Packet Forwarding Engine IPC statistics.
	• `type`—Type of IPC message.
	  • `Header`—IPC message type Header.
	  • `Test`—IPC message type Test.
	  • `Interface`—IPC message type Interface.
	  • `Chassis`—IPC message type Chassis.
	  • `Boot`—IPC message type Boot
	  • `Next-hop`—IPC message type Next-hop.
	  • `Jtree`—IPC message type Jtree.
	  • `Cprod`—IPC message type Cprod.
	  • `Route`—IPC message type Route.
	  • `Pfe`—IPC message type PFE.
	  • `Dfw`—IPC message type Dfw.
	  • `Mastership`—IPC message type Primary Role.
	  • `Sampling`—IPC message type Sampling.
	  • `GUCP`—IPC message type GUCP.
	  • `CoS`—IPC message type CoS.
	  • `GCCP`—IPC message type GCCP.
	  • `GHCP`—IPC message type GHCP.
	  • `IRSD`—IPC message type IRSD.
	  • `Monitoring`—IPC message type Monitoring.

Table 980: show pfe ssb Output Fields *(Continued)*

Field Name	Field Description
	• RE—IPC message type RE. • PIC—IPC message type PIC. • ASP cfg—IPC message type ASP configuration. • ASP cmd—IPC message type ASP command.. • L2TP cfg—IPC message type L2TP configuration. • Collector—IPC message type Collector. • PIC state—IPC message type PIC state. • Aggregator—IPC message type Aggregate. • Empty—IPC message type Empty. • PFE socket-buffer mbuf depth—Information about Packet Forwarding Engine socket-buffer depth • bucket—mbuf bucket value. • count—mbuf count value. •
PFE socket-buffer bytes pending transmit—	Information about Packet Forwarding Engine socket-buffer bytes pending for transmit. • TX Messages—Number of transmitted messages. • RX messages—Number of received messages.

Sample Output

show pfe ssb

```
user@host> show pfe ssb
SSB status:
  Slot:                 Present
```

```
  State:                Online
  Last State Change:    2005-03-06 03:10:28 PST
  Uptime (total):       11:23:27
  Failures:             0
  Pending:              0

Peer message type receive qualifiers:
 Message Type      Receive Qualifier
 ----------------  -----------------
              TTP  Slot only
              IFD  All
              IFL  All
          Nexthop  All
              COS  All
            Route  All
      SW Firewall  All
      HW Firewall  All
   PFE Statistics  All
   PIC Statistics  None
         Sampling  All
       Monitoring  None
              ASP  None
             L2TP  None
        Collector  None
 PIC Configuration  None
 Queue Statistics  None
           (null)  None


PFE listener statistics:
  Open:                 1
  Close:                0
  Sleep:                0
  Wakeup:               0
  Resync Request:       0
  Resync Done:          1
  Resync Fail:          0
  Resync Time:          0


PFE IPC statistics:
 type        TX Messages  RX messages
 ----------- -----------  -----------
```

```
       Header           0            0
         Test           0            0
    Interface         737         9911
      Chassis           0            0
         Boot           0            0
     Next-hop          48            0
        Jtree           0            0
        Cprod           0            0
        Route          94            0
          Pfe        2034          683
          Dfw           8            0
  Mastership            0            0
    Sampling            0            0
        GUCP            0            0
         CoS           73            0
        GCCP            0            0
        GHCP            0            0
        IRSD            0            0
  Monitoring            0            0
          RE            0            0
         PIC            0            0
     ASP cfg            0            0
     ASP cmd            0            0
    L2TP cfg            0            0
   Collector            0            0
   PIC state            0            0
  Aggregator            0            0
       Empty            0            0

PFE socket-buffer mbuf depth:
 bucket          count
 ------          -----
     0               0
     1               0
     2               0
     3               0
     4               0
     5               0
     6               0
     7               0
     8               0
     9               0
    10               0
```

```
    11                0
    12                0
    13                0
    14                0
    15                0
    16                0
    17                0
    18                0
    19                0
    20                0
    21                0

PFE socket-buffer bytes pending transmit:
 bucket          count
 ------          -----
     0               0
     1               0
     2               0
     3               0
     4               0
     5               0
     6               0
     7               0
     8               0
     9               0
    10               0
    11               0
    12               0
    13               0
    14               0
    15               0
    16               0
    17               0
    18               0
    19               0
    20               0
    21               0
```

Release Information

Command introduced before Junos OS Release 7.4.

show pfe statistics bridge

IN THIS SECTION

- Syntax | 23045
- Description | 23045
- Options | 23045
- Required Privilege Level | 23045
- Output Fields | 23046
- Sample Output | 23047
- Release Information | 23051

Syntax

```
show pfe statistics bridge
<fpc slot>
```

Description

Display information about the number of packets discarded in the ingress pipeline of the Packet Forwarding Engine, packets discarded because of egress filtering or congestion filtering, number of control packets, and general counters for dropped packets. You can use this information to inform troubleshooting investigations.

Options

none Display bridge counter statistics for all Flexible PIC Concentrator (FPC) slots.

fpc *slot* (Optional) Display bridge counter statistics for a specific FPC slot.

Required Privilege Level

view

Output Fields

Table 981 on page 23046 lists the output fields for the `show pfe statistics bridge` command. Output fields are listed in the approximate order in which they appear.

Table 981: show pfe statistics bridge Output Fields

Field Name	Field Description
Ingress Counters	Information about ingress counters: • **Received**—Number of packets received by the bridge. • **VLAN Filtered**—Number of packets discarded because of VLAN filtering. • **Security Filtered**—Number of packets discarded because of security filtering. • **Other Discards**—Number of packets dropped by the bridge for reasons other than VLAN or security filtering.
Egress Counters	Information about egress counters: • **Unicast**—Number of unicast packets transmitted. • **Multicast**—Number of multicast packets transmitted. • **Broadcast**—Number of broadcast packets transmitted. • **Egress Filtered**—Number of egress-filtered packets (regardless of port, priority, or mode). • **TailDrop**—Number of packets filtered because of egress queue congestion. • **Forward Restrict**—Number of packets filtered because of egress forward restrictions. • **Congestion Filtered**—Number of packets filtered because of transmit queue (TxQ) congestion. • **Control Packets**—Number of control packets (sent to CPU, received from CPU, and sent to analyzer).

Table 981: show pfe statistics bridge Output Fields *(Continued)*

Field Name	Field Description
Drop Counters	Information about drop counters: • **Drop Mode**—Count mode of the counter. • **Drop Counter**—Counter value.
General Counters	Information about general counters: • **Drop Mode**—Count mode of the counter. • **Drop Counter**—Counter value. • **Source Not Learnt**—Number of source addresses that were not learnt because of internal congestion.
MUX PFE	Information about multiplexer PFE for oversubscribed cards: • **Drop Mode**—Count mode of the counter. • **Drop Count**—Counter value.

Sample Output

show pfe statistics bridge (EX3200 and EX4200 Switches)

```
user@switch> show pfe statistics bridge
Slot 0

PFE:                        0           1           2
----------------------------------------------------------------
----  Ingress Counters ----
Received:                       0          52           0
VLAN Filtered:                  0           0           0
Security Filtered:              0           0           0
Other Discards:                 0           0           0
----  Egress Counters ----
Unicast:                        0         104          52
```

```
Multicast:                      0           0           0
Broadcast:                      0           0           0
Egress Filtered:                0           0           0
Congestion Filtered:            0           0           0
Control Packets:                5           0           0
----  General Counters ----
Drop Mode:                  0           0           0
Drop Counter:           34217       36080        6367
Source Not Learnt:             0           0           0
```

show pfe statistics bridge (EX8200 Switches and EX8200 Virtual Chassis)

```
user@switch> show pfe statistics bridge
Slot 0
PFE:                        0             1
--------------------------------------------------------------------
---- Ingress Counters ----
Received:                     946           48
VLAN Filtered:                  0            0
Security Filtered:              0            0
Other Discards:                 0            0
---- Egress Counters ----
Unicast:                        0            0
Multicast:                      0            0
Broadcast:                      0            0
Egress Filtered:                0            0
TailDrop:                       0            0
Forward Restrict:               0            0
Congestion Filtered:            0            0
Control Packets:             4103          896
---- Drop Counters ----
Drop Mode:                0            0
Drop Counter:         12528            2

Slot 1

PFE:                            0            1
--------------------------------------------------------------------
---- Ingress Counters ----
Received:                       0            0
VLAN Filtered:                  0            0
```

```
Security Filtered:              0            0
Other Discards:                 0            0
---- Egress Counters ----
Unicast:                        0            0
Multicast:                      0            0
Broadcast:                      0            0
Egress Filtered:                0            0
TailDrop:                       0            0
Forward Restrict:               0            0
Congestion Filtered:            0            0
Control Packets:                0            0
---- Drop Counters ----
Drop Mode:                      0            0
Drop Counter:                   0            0

Slot 2

PFE:                            0             1
---------------------------------------------------------------------
---- Ingress Counters ----
Received:                       0             0
VLAN Filtered:                  0             0
Security Filtered:              0             0
Other Filtered:                 0             0
---- Egress Counters ----
Unicast:                        0             0
Multicast:                      0             0
Broadcast:                      0             0
Egress Filtered:                0             0
TailDrop:                       0             0
Forward Restrict:               0             0
Congestion Filtered:            0             0
Control Packets:                0             0
---- Drop Counters ----
Drop Mode:                      0             0
Drop Counter:                   0             0
```

show pfe statistics bridge fpc (EX8200 Switches and EX8200 Virtual Chassis)

```
user@switch> show pfe statistics bridge fpc 2
Slot 2
```

```
PFE:                        0             1
---------------------------------------------------------------------
---- Ingress Counters ----
Received:                0             0
VLAN Filtered:           0             0
Security Filtered:       0             0
Other Discards:          0             0
---- Egress Counters ----
Unicast:                 0             0
Multicast:               0             0
Broadcast:               0             0
Egress Filtered:         0             0
TailDrop:                0             0
Forward Restrict:        0             0
Congestion Filtered:     0             0
Control Packets:         0             0
---- Drop Counters ----
Drop Mode:               0             0
Drop Counter:            0             0
```

show pfe statistics bridge fpc (EX8200-40XS (40-port SFP+) Line Card)

```
user@switch> show pfe statistics bridge fpc 8
Slot 8

PFE:                          0           1           2           3
---------------------------------------------------------------------
----  Ingress Counters ----
Received:                0           3           0           0
VLAN Filtered:           0           0           0           0
Security Filtered:       0           0           0           0
Other Discards:          0           1           0           0
----  Egress Counters ----
Unicast:                 0           0           0           0
Multicast:               0           0           0           0
Broadcast:               0           0           0           0
Egress Filtered:         0           0           0           0
TailDrop:                0           0           0           0
Forward Restrict:        0           0           0           0
Congestion Filtered:     0           2           0           0
```

```
Control Packets:         4           0           0           0
----  Drop Counters ----
Drop Mode:               0           0           0           0
Drop Counter:            0           1           0           0

MUX PFE:                    4           5
--------------------------------------------------------------------
Drop Mode:                  0           0
Drop Count:                 0           0
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Monitoring Switch Control Traffic

show pfe terse

IN THIS SECTION

- Syntax | 23052
- Syntax (TX Matrix and TX Matrix Plus Router) | 23052
- Syntax (MX Series Router) | 23052
- Description | 23052
- Options | 23052
- Required Privilege Level | 23053
- Sample Output | 23053
- Release Information | 23054

Syntax

```
show pfe terse
```

Syntax (TX Matrix and TX Matrix Plus Router)

```
show pfe terse
<lcc number | scc>
<sfc number>
```

Syntax (MX Series Router)

```
show pfe terse
<all-members>
<local>
<member member-id>
```

Description

Display Packet Forwarding Engine status information.

Options

none Display brief information about the Packet Forwarding Engine.

all-members (MX Series routers only) (Optional) Display Packet Forwarding Engine status information for all members in the Virtual Chassis configuration.

lcc *number* (TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine information for a T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine information for the router (or line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

scc (TX Matrix routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix router (or switch-card chassis).

sfc (TX Matrix Plus routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level

admin

Sample Output

show pfe terse (TX Matrix Router)

```
user@host> show pfe terse
Slot Type Slot    State   Flags Uptime
  0  SFM  Present Online  0x0bf 01:25:42
  2  SFM  Present Online  0x0bf 01:25:40
  0  FPC  Present Online  0x102 01:25:57
  1  FPC  Present Online  0x102 01:25:55
  2  FPC  Present Online  0x102 01:25:53
```

show pfe terse (TX Matrix Plus Router)

```
user@host> show pfe terse
sfc0-re0:
--------------------------------------------------------------------------
Slot Type Slot      State     Uptime
 0   LCC  Present   Online    2d 05:26

lcc0-re0:
--------------------------------------------------------------------------
Slot Type Slot      State     Uptime
 0   GFPC Present   Online    2d 05:25
 1   GFPC Present   Online    2d 05:25
```

show pfe terse sfc (TX Matrix Plus Router)

```
user@host> show pfe terse sfc 0
sfc0-re0:
--------------------------------------------------------------------------
Slot Type Slot      State     Uptime
 0   LCC  Present   Online    2d 05:25
```

Release Information

Command introduced before Junos OS Release 7.4.

show pfe version

IN THIS SECTION

- Syntax | 23055
- Description | 23055
- Options | 23055
- Required Privilege Level | 23055

Sample Output | 23055
Release Information | 23056

Syntax

```
show pfe version <brief | detail>
```

Description

Display Packet Forwarding Engine version information.

Options

brief | detail	Display the specified level of output.

Required Privilege Level

admin

Sample Output

show pfe version brief

```
user@host> show pfe version brief
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC
```

show pfe version detail

```
user@host> show pfe version detail
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC
    device01.example.com:/volume/build/junos/rpd_feb11/11.1/development/20101111.0/obj-i386/
junos/usr.sbin/pfed
```

Release Information

Command introduced before Junos OS Release 7.4.

show pim bidirectional df-election

IN THIS SECTION

- Syntax | 23056
- Description | 23056
- Options | 23056
- Required Privilege Level | 23057
- Output Fields | 23057
- Sample Output | 23058
- Release Information | 23059

Syntax

```
show pim bidirectional df-election
<brief | detail >
<inet | inet6>
<instance instance name>
<logical-system (all |  logical-system-name)>
<rpa address>
```

Description

For bidirectional PIM, display the designated forwarder (DF) election results for each interface grouped by the rendezvous point addresses (RPAs).

Options

none	Display standard information about all interfaces.

brief | detail	(Optional) Display the specified level of output.
inet | inet6	(Optional) Display DF election results for IPv4 or IPv6 family addresses, respectively.
instance *instance-name*	(Optional) Display DF election results for a specific routing instance.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system.
rpa *address*	(Optional) Display the DF election results for an RP address.

Required Privilege Level

view

Output Fields

Table 982 on page 23057 describes the output fields for the `show pim bidirectional df-election` command. Output fields are listed in the approximate order in which they appear.

Table 982: show pim bidirectional df-election Output Fields

Field Name	Field Description	Level of Output
Family	IPv4 address family (**INET**) or IPv6 address family (**INET6**).	All levels
Instance	Name of the routing instance.	All levels
RPA	RP address.	All levels
Group ranges	Address ranges of the multicast groups mapped to this RP address.	All levels
Interfaces	Bidirectional PIM interfaces on this routing device. An interface can win the DF election (**Win**), lose the DF election (**Lose**), or be the RP link (**RPL**). The RP link is the interface directly connected to a subnet that contains a phantom RP address. A phantom RP address is an RP address that is not assigned to a routing device interface.	All levels **brief** displays the DF election winner only.

Table 982: show pim bidirectional df-election Output Fields *(Continued)*

Field Name	Field Description	Level of Output
DF	IP address of the designated forwarder.	All levels

Sample Output

show pim bidirectional df-election

```
user@host> show pim bidirectional df-election
Instance: PIM.master Family: INET

RPA: 10.10.1.3
Group ranges: 224.1.3.0/24, 225.1.3.0/24
Interfaces:
    ge-0/0/1.0    (RPL)      DF: none
    lo0.0         (Win)      DF: 10.255.179.246
    xe-4/1/0.0    (Win)      DF: 10.10.2.1

RPA: 10.10.13.2
Group ranges: 224.1.1.0/24, 225.1.1.0/24
Interfaces:
    ge-0/0/1.0    (Lose)     DF: 10.10.1.2
    lo0.0         (Win)      DF: 10.255.179.246
    xe-4/1/0.0    (Lose)     DF: 10.10.2.2

Instance: PIM.master Family: INET6

RPA: fec0::10:10:1:3
Group ranges: ff00::/8
Interfaces:
    ge-0/0/1.0    (Lose)     DF: fe80::b2c6:9aff:fe95:86fa
    lo0.0         (Win)      DF: fe80::2a0:a50f:fc64:e661
    xe-4/1/0.0    (Win)      DF: fe80::226:88ff:fec5:3c37

RPA: fec0::10:10:13:2
Group ranges: ff00::/8
Interfaces:
    ge-0/0/1.0    (Lose)     DF: fe80::b2c6:9aff:fe95:86fa
```

```
    lo0.0         (Win)      DF: fe80::2a0:a50f:fc64:e661
    xe-4/1/0.0    (Win)      DF: fe80::226:88ff:fec5:3c37
```

show pim bidirectional df-election brief

```
user@host> show pim bidirectional df-election brief
Instance: PIM.master Family: INET

RPA: 10.10.1.3
Group ranges: 224.1.3.0/24, 225.1.3.0/24
Interfaces:
    lo0.0         (Win)      DF: 10.255.179.246
    xe-4/1/0.0    (Win)      DF: 10.10.2.1

RPA: 10.10.13.2
Group ranges: 224.1.1.0/24, 225.1.1.0/24
Interfaces:
    lo0.0         (Win)      DF: 10.255.179.246

Instance: PIM.master Family: INET6

RPA: fec0::10:10:1:3
Group ranges: ff00::/8
Interfaces:
    lo0.0         (Win)      DF: fe80::2a0:a50f:fc64:e661
    xe-4/1/0.0    (Win)      DF: fe80::226:88ff:fec5:3c37

RPA: fec0::10:10:13:2
Group ranges: ff00::/8
Interfaces:
    lo0.0         (Win)      DF: fe80::2a0:a50f:fc64:e661
    xe-4/1/0.0    (Win)      DF: fe80::226:88ff:fec5:3c37
```

Release Information

Command introduced in Junos OS Release 12.1.

show pim bidirectional df-election interface

IN THIS SECTION

- Syntax | 23060
- Description | 23060
- Options | 23060
- Required Privilege Level | 23061
- Output Fields | 23061
- Sample Output | 23062
- Release Information | 23063

Syntax

```
show pim bidirectional df-election interface
<inet | inet6>
<instance instance name>
<interface-name>
<logical-system (all |  logical-system-name)>
```

Description

For bidirectional PIM, display the default and the configured designated forwarder (DF) election parameters for each interface.

Options

none	Display standard information about all interfaces.
inet | inet6	(Optional) Display DF election parameters for IPv4 or IPv6 family addresses, respectively.
instance *instance-name*	(Optional) Display DF election parameters for a specific routing instance.

interface-name	(Optional) Display DF election parameters for a specific interface.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 983 on page 23061 describes the output fields for the `show pim bidirectional df-election interface` command. Output fields are listed in the approximate order in which they appear.

Table 983: show pim bidirectional df-election interface Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Family	IPv4 address family (**INET**) or IPv6 address family (**INET6**).
Interface	Name of the bidirectional PIM interface.
Robustnes Count	Minimum number of DF election messages that must fail to be received for DF election to fail.
Offer Period	Interval between repeated DF election messages.
Backoff Period	Period that the acting DF waits between receiving a better DF Offer and sending the Pass message to transfer DF responsibility.
RPA	RP address.
State	For each RP address, state of each interface with respect to the DF election: **Offer** (when the election is in progress), **Win**, or **Lose**.

Table 983: show pim bidirectional df-election interface Output Fields *(Continued)*

Field Name	Field Description
DF	IP address of the designated forwarder.

Sample Output

show pim bidirectional df-election interface

```
user@host> show pim bidirectional df-election interface
Instance: PIM.master Family: INET

Interface: ge-0/0/1.0
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                            State  DF
  10.10.1.3                      Offer  none
  10.10.13.2                     Lose   10.10.1.2

Interface: lo0.0
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                            State  DF
  10.10.1.3                      Win    10.255.179.246
  10.10.13.2                     Win    10.255.179.246

Interface: xe-4/1/0.0
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                            State  DF
  10.10.1.3                      Win    10.10.2.1
  10.10.13.2                     Lose   10.10.2.2
```

```
Instance: PIM.master Family: INET6

Interface: ge-0/0/1.0
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                            State  DF
  fec0::10:10:1:3                Lose   fe80::b2c6:9aff:fe95:86fa
  fec0::10:10:13:2               Lose   fe80::b2c6:9aff:fe95:86fa

Interface: lo0.0
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                            State  DF
  fec0::10:10:1:3                Win    fe80::2a0:a50f:fc64:e661
  fec0::10:10:13:2               Win    fe80::2a0:a50f:fc64:e661

Interface: xe-4/1/0.0
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                            State  DF
  fec0::10:10:1:3                Win    fe80::226:88ff:fec5:3c37
  fec0::10:10:13:2               Win    fe80::226:88ff:fec5:3c37
```

Release Information

Command introduced in Junos OS Release 12.1.

show pim bootstrap

IN THIS SECTION

- Syntax | 23064
- Syntax (EX Series Switch and the QFX Series) | 23064
- Description | 23064
- Options | 23064
- Required Privilege Level | 23065
- Output Fields | 23065
- Sample Output | 23066
- Release Information | 23066

Syntax

```
show pim bootstrap
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show pim bootstrap
<instance instance-name>
```

Description

For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers.

Options

none	Display PIM bootstrap router information for all routing instances.

instance *instance-name*	(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 984 on page 23065 describes the output fields for the `show pim bootstrap` command. Output fields are listed in the approximate order in which they appear.

Table 984: show pim bootstrap Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
BSR	Bootstrap router.
Pri	Priority of the routing device as elected to be the bootstrap router.
Local address	Local routing device address.
Pri	Local routing device address priority to be elected as the bootstrap router.
State	Local routing device election state: **Candidate**, **Elected**, or **Ineligible**.
Timeout	How long until the local routing device declares the bootstrap router to be unreachable, in seconds.

Sample Output

show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master

BSR                     Pri Local address           Pri State      Timeout
None                      0 10.255.71.46              0 InEligible       0
2001:db8:1:1:1:0:aff:785c 34 2001:db8:1:1:1:0:aff:7c12   0 InEligible       0
```

show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A

BSR             Pri Local address   Pri State      Timeout
None              0 192.168.196.105   0 InEligible       0
```

Release Information

Command introduced before Junos OS Release 7.4.

instance option introduced in Junos OS Release 10.0 for EX Series switches.

show pim interfaces

IN THIS SECTION

- Syntax | 23067
- Syntax (EX Series Switch and the QFX Series) | 23067
- Description | 23067
- Options | 23067
- Required Privilege Level | 23068

- Output Fields | 23068
- Sample Output | 23069
- Release Information | 23070

Syntax

```
show pim interfaces
<inet | inet6>
<instance (instance-name | all)>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show pim interfaces
<inet | inet6>
<instance (instance-name | all)>
```

Description

Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured.

Options

none	Display interface information for all family addresses for the main instance.
inet | inet6	(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.
instance (*instance-name* | all)	(Optional) Display information about interfaces for a specific PIM-enabled routing instance or for all routing instances.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 985 on page 23068 describes the output fields for the `show pim interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 985: show pim interfaces Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Name	Interface name.
State	State of the interface. The state also is displayed in the `show interfaces` command.
Mode	PIM mode running on the interface: • **B**—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers. • **S**—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic. • **Dense**—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.) • **Sparse-Dense**—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as **dense** is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as **sparse** is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules. When sparse-dense mode is configured, the output includes both **S** and **D**. When bidirectional-sparse mode is configured, the output includes **S** and **B**. When bidirectional-sparse-dense mode is configured, the output includes **B**, **S**, and **D**.

Table 985: show pim interfaces Output Fields *(Continued)*

Field Name	Field Description
IP	Version number of the address family on the interface: **4** (IPv4) or **6** (IPv6).
V	PIM version running on the interface: **1** or **2**.
State	State of PIM on the interface: • **Active**—Bidirectional mode is enabled on the interface and on all PIM neighbors. • **DR**—Designated router. • **NotCap**—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol. • **NotDR**—Not the designated router. • **P2P**—Point to point.
NbrCnt	Number of neighbors that have been seen on the interface.
JoinCnt(sg)	Number of (s,g) join messages that have been seen on the interface.
JointCnt(*g)	Number of (*,g) join messages that have been seen on the interface.
DR address	Address of the designated router.

Sample Output

show pim interfaces

```
user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
```

```
Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name                Stat Mode IP V State        NbrCnt JoinCnt(sg/*g) DR address
ge-0/3/0.0          Up   S     4 2 NotDR,NotCap      1 0/0            40.0.0.3
ge-0/3/3.50         Up   S     4 2 DR,NotCap         1 9901/100       50.0.0.2
ge-0/3/3.51         Up   S     4 2 DR,NotCap         1 0/0            51.0.0.2
pe-1/2/0.32769      Up   S     4 2 P2P,NotCap        0 0/0
```

show pim interfaces (PIM using point-to-multipoint mode)

```
user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable, P2MP = Point-to-multipoint
link

Name      Stat Mode IP V State        NbrCnt JoinCnt(sg/*g) DR address
st0.0     Up   S     4 2 DR,P2MP      0      10/0   192.0.2.0
```

Release Information

Command introduced before Junos OS Release 7.4.

inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.

Commmand introduced in Junos OS Release 11.3 for the QFX Series.

Support for bidirectional PIM added in Junos OS Release 12.1.

Support for the **instance all** option added in Junos OS Release 12.1.

show pim join

IN THIS SECTION

Syntax | 23071

- Syntax (EX Series Switch and the QFX Series) | 23071
- Description | 23072
- Options | 23072
- Required Privilege Level | 23072
- Output Fields | 23073
- Sample Output | 23078
- Release Information | 23092

Syntax

```
show pim join
<brief | detail | extensive | summary>
<bidirectional | dense | sparse>
<downstream-count>
<exact>
<inet | inet6>
<instance (all | instance-name)>
<logical-system (all | logical-system-name)>
<range>
<resolve-mvpn-neighbor>
<rp ip-address/prefix | source ip-address/prefix>
<sg | star-g>
```

Syntax (EX Series Switch and the QFX Series)

```
show pim join
<brief | detail | extensive | summary>
<dense | sparse>
<exact>
<inet | inet6>
<instance (all | instance-name)>
<range>
<rp ip-address/prefix | source ip-address/prefix>
<sg | star-g>
```

Description

Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.

For bidirectional PIM, display information about PIM group ranges (*,G-range) for each active bidirectional RP group range, in addition to each of the joined (*,G) routes.

Options

none	Display the standard information about PIM groups for all supported family addresses for all routing instances.
brief | detail | extensive | summary	(Optional) Display the specified level of output.
bidirectional | dense | sparse	(Optional) Display information about PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.
downstream-count	(Optional) Display the downstream count instead of a list.
exact	(Optional) Display information about only the group that exactly matches the specified group address.
inet | inet6	(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.
instance (all | *instance-name*)	(Optional) Display information about groups for all PIM-enabled routing instances or a specific PIM-enabled routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
range	(Optional) Address range of the group, specified as ***prefix/ prefix-length***.
resolve-mvpn-neighbor	(Optional) Display BGP-MVPN neighbor.
rp *ip-address/prefix* | source *ip-address/prefix*	(Optional) Display information about the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.
sg | star-g	(Optional) Display information about PIM (S,G) or (*,G) entries.

Required Privilege Level

view

Output Fields

Table 986 on page 23073 describes the output fields for the `show pim join` command. Output fields are listed in the approximate order in which they appear.

Table 986: show pim join Output Fields

Field Name	Field Description	Level of Output
`Instance`	Name of the routing instance.	`brief` `detail` `extensive` `summary` none
`Family`	Name of the address family: `inet` (IPv4) or `inet6` (IPv6).	`brief` `detail` `extensive` `summary` none
`Route type`	Type of multicast route: (S,G) or (*,G).	`summary`
`Route count`	Number of (S,G) routes and number of (*,G) routes.	`summary`
`R`	Rendezvous Point Tree.	`brief` `detail` `extensive` none
`S`	Sparse.	`brief` `detail` `extensive` none
`W`	Wildcard.	`brief` `detail` `extensive` none
`Group`	Group address.	`brief` `detail` `extensive` none
`Bidirectional group prefix length`	For bidirectional PIM, length of the IP prefix for RP group ranges.	All levels
`Source`	Multicast source: • `*` (wildcard value) • *`ipv4-address`* • *`ipv6-address`*	`brief` `detail` `extensive` none

Table 986: show pim join Output Fields *(Continued)*

Field Name	Field Description	Level of Output
RP	Rendezvous point for the PIM group.	brief detail extensive none
Flags	PIM flags: • bidirectional—Bidirectional mode entry. • dense—Dense mode entry. • rptree—Entry is on the rendezvous point tree. • sparse—Sparse mode entry. • spt—Entry is on the shortest-path tree for the source. • wildcard—Entry is on the shared tree.	brief detail extensive none
Upstream interface	RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*,G). For bidirectional PIM, RP Link means that the interface is directly connected to a subnet that contains a phantom RP address. A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.	brief detail extensive none
Upstream neighbor	Information about the upstream neighbor: Direct, Local, Unknown, or a specific IP address. For bidirectional PIM, Direct means that the interface is directly connected to a subnet that contains a phantom RP address. The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.	extensive
Upstream rpf-vector	Information about the upstream Reverse Path Forwarding (RPF) vector; appears in conjunction with the rpf-vector command.	extensive

Table 986: show pim join Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Active upstream interface	When multicast-only fast reroute (MoFRR) is configured in a PIM domain, the upstream interface for the active path. A PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. Preference is given to two paths that do not converge to the same immediate upstream router. PIM installs appropriate multicast routes with upstream neighbors as RPF next hops with two (primary and backup) interfaces.	`extensive`
Active upstream neighbor	On the MoFRR primary path, the IP address of the neighbor that is directly connected to the active upstream interface.	`extensive`
MoFRR Backup upstream interface	The MoFRR upstream interface that is used when the primary path fails. When the primary path fails, the backup path is upgraded to primary, and traffic is forwarded accordingly. If there are alternate paths available, a new backup path is calculated and the appropriate multicast route is updated or installed.	`extensive`
MoFRR Backup upstream neighbor	IP address of the MoFRR upstream neighbor.	extensive

Table 986: show pim join Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Upstream state`	Information about the upstream interface: • `Join to RP`—Sending a join to the rendezvous point. • `Join to Source`—Sending a join to the source. • `Local RP`—Sending neither join messages nor prune messages toward the RP, because this routing device is the rendezvous point. • `Local Source`—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device. • `No Prune to RP`—Automatically sent to RP when SPT and RPT are on the same path. • `Prune to RP`—Sending a prune to the rendezvous point. • `Prune to Source`—Sending a prune to the source. **NOTE**: RP group range entries have `None` in the `Upstream state` field because RP group ranges do not trigger actual PIM join messages between routing devices.	`extensive`

Table 986: show pim join Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Downstream neighbors`	Information about downstream interfaces: • `Interface`—Interface name for the downstream neighbor. A pseudo PIM-SM interface appears for all IGMP-only interfaces. A pseudo multipoint LDP (Pseudo-MLDP) interface appears on ingress root nodes in M-LDP point-to-multipoint LSPs with inband signaling. • `Interface address`—Address of the downstream neighbor. • `State`—Information about the downstream neighbor: `join` or `prune`. • `Flags`—PIM join flags: `R (RPtree)`, `S (Sparse)`, `W (Wildcard)`, or `zero`. • `Uptime`—Time since the downstream interface joined the group. • `Time since last Join`—Time since the last join message was received from the downstream interface. • `Time since last Prune`—Time since the last prune message was received from the downstream interface. • `rpf-vector`—IP address of the RPF vector TLV .	`extensive`
`Number of downstream interfaces`	Total number of outgoing interfaces for each (S,G) entry.	`extensive`
`Assert Timeout`	Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null.	`extensive`
`Keepalive timeout`	Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, `Keepalive timeout` is `Infinity`.	`extensive`

Table 986: show pim join Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Uptime`	Time since the creation of (S,G) or (*,G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*,G) state.	`extensive`
`Bidirectional accepting interfaces`	Interfaces on the routing device that forward bidirectional PIM traffic. The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (`DF Winner`), or the interface is the reverse path forwarding (RPF) interface toward the RP (`RPF`).	`extensive`

Sample Output

show pim join summary

```
user@host> show pim join summary
Instance: PIM.master Family: INET

Route type            Route count
(s,g)                 2
(*,g)                 1

Instance: PIM.master Family: INET6
```

show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
```

```
    Upstream interface: Local

Group: 233.252.0.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 233.252.0.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0

Group: 233.252.0.2
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)

Group: 233.252.0.3
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
```

```
Group: 233.252.0.4
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join inet6

```
user@host> show pim join inet6
Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

 Group: 2001:db8::e000:101
     Source: *
     RP: ::10.46.0.13
     Flags: sparse,rptree,wildcard
     Upstream interface: Local

 Group: 2001:db8::e000:101
     Source: ::10.1.1.1
     Flags: sparse
     Upstream interface: unknown (no neighbor)

 Group: 2001:db8::e800:101
     Source: ::10.1.1.1
     Flags: sparse
     Upstream interface: unknown (no neighbor)

 Group: 2001:db8::e800:101
     Source: ::10.1.1.2
     Flags: sparse
     Upstream interface: unknown (no neighbor)
```

show pim join inet6 star-g

```
user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

 Group: 2001:db8::e000:101
     Source: *
     RP: ::10.46.0.13
     Flags: sparse,rptree,wildcard
     Upstream interface: Local
```

show pim join instance <instance-name>

```
user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.2
    Source: *
    RP: 10.10.47.100
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 233.252.0.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0

Group: 233.252.0.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join instance <instance-name> downstream-count

```
user@host> show pim join instance VPN-A downstream-count
Instance: PIM.SML_VRF_4 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: *
    RP: 10.11.11.6
    Flags: sparse,rptree,wildcard
    Upstream interface: mt-1/2/10.32813
    Number of downstream interfaces: 4

Group: 233.252.0.1
    Source: 10.1.1.1
    Flags: sparse,spt
    Upstream interface: ge-0/0/3.5
    Number of downstream interfaces: 5
```

show pim join instance <instance-name> downstream-count extensive

```
user@host> show pim join instance VPN-A downstream-count extensive
Instance: PIM.SML_VRF_4 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: *
    RP: 10.11.11.6
    Flags: sparse,rptree,wildcard
    Upstream interface: mt-1/2/10.32813
    Upstream neighbor: 10.2.2.7 (assert winner)
    Upstream state: Join to RP
    Uptime: 02:51:41
    Number of downstream interfaces: 4
    Number of downstream neighbors: 4

Group: 233.252.0.1
    Source: 10.1.1.1
    Flags: sparse,spt
    Upstream interface: ge-0/0/3.5
```

```
    Upstream neighbor: 10.1.1.17
    Upstream state: Join to Source, Prune to RP
    Keepalive timeout: 0
    Uptime: 02:51:42
    Number of downstream interfaces: 5
    Number of downstream neighbors: 7
```

show pim join detail

```
user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 233.252.0.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 233.252.0.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (PIM Resolve TLV for Multicast in Seamless MPLS)

```
user@host> show pim join extensive
Group: 228.26.1.5
    Source: 10.60.0.101
    Flags: sparse,spt
    Upstream interface: ge-5/0/0.1
    Upstream neighbor: 10.100.1.13
```

```
    Upstream state: Join to Source
Upstream rpf-vector: 10.100.20.1
    Keepalive timeout: 178
    Uptime: 17:44:38
    Downstream neighbors:
        Interface: xe-2/0/3.1
            10.203.2.190 State: Join Flags: S Timeout: 156
            Uptime: 17:44:38 Time since last Join: 00:00:54
            rpf-vector: 10.100.20.1
        Interface: xe-2/0/2.1
            10.203.1.190 State: Join Flags: S Timeout: 156
            Uptime: 17:44:38 Time since last Join: 00:00:54
            rpf-vector: 10.100.20.2
    Number of downstream interfaces: 2
    Number of downstream neighbors: 2
```

show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: so-1/0/0.0
            10.111.10.2 State: Join Flags: SRW Timeout: 174
            Uptime: 00:03:49 Time since last Join: 00:01:49
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: SRW  Timeout: Infinity
            Uptime: 00:03:49 Time since last Join: 00:01:49
    Number of downstream interfaces: 2

Group: 233.252.0.1
    Source: 10.255.14.144
```

```
    Flags: sparse,spt
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local Source, Local RP
    Keepalive timeout: 344
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: so-1/0/0.0
            10.111.10.2 State: Join Flags: S Timeout: 174
            Uptime: 00:03:49 Time since last Prune: 00:01:49
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: S   Timeout: Infinity
            Uptime: 00:03:49 Time since last Prune: 00:01:49
    Number of downstream interfaces: 2

Group: 233.252.0.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0
    Upstream neighbor: 10.111.10.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 344
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: Pseudo-GMP
            fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
        Interface: so-1/0/0.0 (pruned)
            10.111.10.2 State: Prune Flags: SR Timeout: 174
            Uptime: 00:03:49 Time since last Prune: 00:01:49
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: S   Timeout: Infinity
            Uptime: 00:03:49 Time since last Prune: 00:01:49
    Number of downstream interfaces: 3

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
```

```
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.10.1.2
    Upstream state: None
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
    Number of downstream interfaces: 0

Group: 233.252.0.1
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.10.1.2
    Upstream state: None
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
     Downstream neighbors:
         Interface: lt-1/0/10.24
             10.0.24.4 State: Join   RW  Timeout: 185
         Interface: lt-1/0/10.23
             10.0.23.3 State: Join   RW  Timeout: 184
     Number of downstream interfaces: 2

Group: 233.252.0.2
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)
    Upstream neighbor: Direct
    Upstream state: Local RP
```

```
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
        Interface: xe-4/1/0.0    (DF Winner)
    Number of downstream interfaces: 0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)
    Upstream neighbor: Direct
    Upstream state: Local RP
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
        Interface: xe-4/1/0.0    (DF Winner)
    Number of downstream interfaces: 0
```

show pim join instance <instance-name> extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.2
    Source: *
    RP: 10.10.47.100
```

```
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: mt-1/1/0.32768
            10.10.47.101 State: Join Flags: SRW Timeout: 156
            Uptime: 00:03:49 Time since last Join: 00:01:49
    Number of downstream interfaces: 1

Group: 233.252.0.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0
    Upstream neighbor: 10.111.30.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156
    Uptime: 00:14:52

Group: 233.252.0.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0
    Upstream neighbor: 10.111.20.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156
    Uptime: 00:14:52
```

show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
```

```
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:55
    Downstream neighbors:
        Interface: Pseudo-MLDP
        Interface: lt-1/2/0.25
            10.2.5.2 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 233.252.0.2
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:41
    Downstream neighbors:
        Interface: Pseudo-MLDP

Group: 233.252.0.3
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:41
    Downstream neighbors:
        Interface: Pseudo-MLDP

Group: 233.252.0.22
    Source: 10.2.7.7
    Flags: sparse,spt
    Upstream interface: lt-1/2/0.27
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:25
    Downstream neighbors:
        Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
```

```
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 2001:db8::1:2
    Source: 2001:db8::1:2:7:7
    Flags: sparse,spt
    Upstream interface: lt-1/2/0.27
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:26
    Downstream neighbors:
        Interface: Pseudo-MLDP
```

show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.0
    Source: *
    RP: 10.1.1.1
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 11:31:33
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: SRW  Timeout: Infinity
            Uptime: 11:31:33 Time since last Join: 11:31:32

Group: 233.252.0.1
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <10.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
```

```
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 233.252.0.2
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <10.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30
    Downstream neighbors:
        Interface: lt-1/2/0.14
            10.1.4.4 State: Join Flags: S Timeout: 177
            Uptime: 11:30:33 Time since last Join: 00:00:33
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 233.252.0.3
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <10.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: fe-1/3/0.0
```

```
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 233.252.0.22
    Source: 10.2.7.7
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <10.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:30
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 2001:db8::1:2
    Source: 2001:db8::1:2:7:7
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <10.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: fe-1/3/0.0
            2001:db8::21f:12ff:fea5:c4db State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32
```

Release Information

Command introduced before Junos OS Release 7.4.

summary option introduced in Junos OS Release 9.6.

inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.

Support for bidirectional PIM added in Junos OS Release 12.1.

Multiple new filter options introduced in Junos OS Release 13.2.

`downstream-count` option introduced in Junos OS Release 16.1.

Support for PIM NSR support for VXLAN added in Junos OS Release 16.2

Support for RFC 5496 (via `rpf-vector`) added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

- clear pim join
- Example: Configuring Multicast-Only Fast Reroute in a PIM Domain
- Example: Configuring Bidirectional PIM
- *Example: Configuring PIM State Limits*

show pim mdt

IN THIS SECTION

- Syntax | **23093**
- Description | **23094**
- Options | **23094**
- Required Privilege Level | **23094**
- Output Fields | **23094**
- Sample Output | **23096**
- Release Information | **23099**

Syntax

```
show pim mdt instance instance-name
<brief | detail | extensive>
data-mdt-joins
data-mdt-limit
inet
```

```
inet6
<incoming | outgoing>
<logical-system (all | logical-system-name)>
<range>
```

Description

Display information about Protocol Independent Multicast (PIM) default multicast distribution tree (MDT) and the data MDTs in a Layer 3 VPN environment for a routing instance.

Options

instance *instance-name*	Display information about data-MDTs for a specific PIM-enabled routing instance.
brief | detail | extensive	(Optional) Display the specified level of output.
data-mdt-joins	Show received PIM data-mdt-joins.
data-mdt-limits	Show received PIM data-mdt-limits.
incoming | outgoing	(Optional) Display incoming or outgoing multicast data tunnels, respectively.
inet | inet6	Display IPv4 or IPv6 multicast data tunnels.
logical-system (all | *logical-system-name***)**	(Optional) Perform this operation on all logical systems or on a particular logical system.
range	(Optional) Display information about an IP address with optional prefix length representing a particular multicast group.

Required Privilege Level

view

Output Fields

Table 987 on page 23095 describes the output fields for the `show pim mdt` command. Output fields are listed in the approximate order in which they appear.

Table 987: show pim mdt Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Tunnel direction	Direction the tunnel faces, from the router's perspective: **Outgoing** or **Incoming**.	All levels
Tunnel mode	Mode the tunnel is operating in: **PIM-SSM** or **PIM-ASM**.	All levels
Default group address	Default multicast group address using this tunnel.	All levels
Default source address	Default multicast source address using this tunnel.	All levels
Default tunnel interface	Default multicast tunnel interface.	All levels
Default tunnel source	Address used as the source address for outgoing PIM control messages.	All levels
C-Group	Customer-facing multicast group address using this tunnel. If you enable dynamic reuse of data MDT group addresses, more than one group address can use the same data MDT.	**detail**
C-Source	IP address of the multicast source in the customer's address space. If you enable dynamic reuse of data MDT group addresses, more than one source address can use the same data MDT.	**detail**
P-Group	Service provider-facing multicast group address using this tunnel.	**detail**
Data tunnel interface	Multicast data tunnel interface that set up the data-MDT tunnel.	**detail**

Table 987: show pim mdt Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Last known forwarding rate	Last known rate, in kilobits per second, at which the tunnel was forwarding traffic.	**detail**
Configured threshold rate	Rate, in kilobits per second, above which a data-MDT tunnel is created and below which it is deleted.	**detail**
Tunnel uptime	Time that this data-MDT tunnel has existed. The format is *hours:minutes:seconds*.	**detail**

Sample Output

show pim mdt <variables> instance

Use this command to display MDT information for default MDT and data-MDT for IPv4 and/or IPv6 traffic.)

```
user@host> show pim mdt inet | inet6 instance VPN-A
Instance: PIM.VPN-A  Family: INET
Tunnel direction: Outgoing
Tunnel mode: PIM-SM
Default group address: 224.1.1.1
Default source address: 0.0.0.0
Default tunnel interface: mt-0/0/0.32768
Default tunnel source: 0.0.0.0

C-group address   C-source address   P-group address    Data tunnel interface
227.1.1.1         18.1.1.2           228.1.1.1          mt-0/0/0.32769

Instance: PIM.VPN-A
Tunnel direction: Incoming
Tunnel mode: PIM-SM
Default group address: 224.1.1.1
Default source address: 0.0.0.0
Default tunnel interface: mt-0/0/0.1081344
Default tunnel source: 0.0.0.0
```

```
Instance: PIM.VPN-A Family: INET6
```

show pim mdt instance detail

```
user@host> show pim mdt instance VPN-A detail
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

C-Group: 235.1.1.2
    C-Source: 192.168.195.74
    P-Group : 228.0.0.0
    Data tunnel interface      : mt-1/1/0.32769
    Last known forwarding rate : 48 kbps (6 kBps)
    Configured threshold rate  : 10 kbps
    Tunnel uptime              : 00:00:34

Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.1081344
```

show pim mdt instance extensive

```
user@host> show pim mdt instance VPN-A extensive
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

C-Group: 235.1.1.2
    C-Source: 192.168.195.74
    P-Group : 228.0.0.0
    Data tunnel interface      : mt-1/1/0.32769
    Last known forwarding rate : 48 kbps (6 kBps)
    Configured threshold rate  : 10 kbps
```

```
    Tunnel uptime              : 00:00:41

Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.1081344
```

show pim mdt instance incoming

```
user@host> show pim mdt instance VPN-A incoming
Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.1081344
```

show pim mdt instance outgoing

```
user@host> show pim mdt instance VPN-A outgoing
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

C-group address   C-source address   P-group address    Data tunnel interface
235.1.1.2         192.168.195.74     228.0.0.0          mt-1/1/0.32769
```

show pim mdt instance (SSM Mode)

```
user@host> show pim mdt instance vpn-a
Instance: PIM.vpn-a
Tunnel direction: Outgoing
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.216
Default tunnel interface: mt-1/3/0.32769
Default tunnel source: 192.168.7.1
```

```
Instance: PIM.vpn-a
Tunnel direction: Incoming
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.217
Default tunnel interface: mt-1/3/0.1081345

Instance: PIM.vpn-a
Tunnel direction: Incoming
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.218
Default tunnel interface: mt-1/3/0.1081345
```

Release Information

Command introduced before Junos OS Release 7.4.

Support for IPv6 added in Junos OS Release 17.3R1.

show pim mdt data-mdt-joins

IN THIS SECTION

- Syntax | 23100
- Description | 23100
- Options | 23100
- Required Privilege Level | 23100
- Output Fields | 23100
- Sample Output | 23101
- Release Information | 23101

Syntax

```
show pim mdt data-mdt-joins
<logical-system (all | logical-system-name)> instance instance-name
```

Description

In a draft-rosen Layer 3 multicast virtual private network (MVPN) configured with service provider tunnels, display the advertisements of new multicast distribution tree (MDT) group addresses cached by the provider edge (PE) routers in the specified VPN routing and forwarding (VRF) instance that is configured to use the Protocol Independent Multicast (PIM) protocol.

Options

instance *instance-name*	Display data MDT join packets cached by PE routers in a specific PIM instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

> **NOTE**: Draft-rosen multicast VPNs are not supported in a logical system environment even though the configuration statements can be configured under the logical-systems hierarchy.

Required Privilege Level

view

Output Fields

Table 988 on page 23101 describes the output fields for the `show pim mdt data-mdt-joins` command. Output fields are listed in the approximate order in which they appear.

Table 988: show pim mdt data-mdt-joins Output Fields

Field Name	Field Description
C-Group	IPv4 group address in the address space of the customer's VPN-specific PIM-enabled routing instance of the multicast traffic destination. This 32-bit value is carried in the C-group field of the MDT join TLV packet.
C-Source	IPv4 address in the address space of the customer's VPN-specific PIM-enabled routing instance of the multicast traffic source. This 32-bit value is carried in the C-source field of the MDT join TLV packet.
P-Group	IPv4 group address in the service provider's address space of the new data MDT that the PE router will use to encapsulate the VPN multicast traffic flow (C-Source, C-Group). This 32-bit value is carried in the P-group field of the MDT join TLV packet.
P-Source	IPv4 address of the PE router.
Timeout	Timeout, in seconds, remaining for this cache entry. When the cache entry is created, this field is set to 180 seconds. After an entry times out, the PE router deletes the entry from its cache and prunes itself off the data MDT.

Sample Output

show pim mdt data-mdt-joins

```
user@host show pim mdt data-mdt-joins instance VPN-A
C-Source         C-Group          P-Source         P-Group          Timeout
20.2.15.9        225.1.1.2        20.0.0.5         239.10.10.0      172
20.2.15.9        225.1.1.3        20.0.0.5         239.10.10.1      172
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Data MDTs

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

show pim mdt data-mdt-limit

IN THIS SECTION

- Syntax | 23102
- Description | 23102
- Options | 23102
- Required Privilege Level | 23103
- Output Fields | 23103
- Sample Output | 23103
- Release Information | 23104

Syntax

```
show pim mdt data-mdt-limit instance instance-name
<logical-system (all | logical-system-name)>
```

Description

Display the maximum number configured and the currently active data multicast distribution trees (MDTs) for a specific VPN routing and forwarding (VRF) instance.

Options

instance *instance-name*	Display data MDT information for the specified VRF instance.

logical-system (all | *logical-system-name*) (Optional) Perform this operation on all logical systems or on a particular logical system.

> **NOTE**: Draft-rosen multicast VPNs are not supported in a logical system environment even though the configuration statements can be configured under the logical-systems hierarchy.

Required Privilege Level

view

Output Fields

Table 989 on page 23103 describes the output fields for the `show pim mdt data-mdt-limit` command. Output fields are listed in the approximate order in which they appear.

Table 989: show pim mdt data-mdt-limit Output Fields

Field Name	Field Description
Maximum Data Tunnels	Maximum number of data MDTs created in this VRF instance. If the number is 0, no data MDTs are created for this VRF instance.
Active Data Tunnels	Active number of data MDTs in this VRF instance.

Sample Output

show pim mdt data-mdt-limit

```
user@host show pim mdt data-mdt-limit instance VPN-A
Maximum Data Tunnels                 10
Active Data Tunnels                   2
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Understanding Data MDTs

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

show pim mvpn

IN THIS SECTION

- Syntax | 23104
- Description | 23104
- Options | 23105
- Required Privilege Level | 23105
- Output Fields | 23105
- Sample Output | 23106
- Release Information | 23106

Syntax

```
show pim mvpn
<logical-system (all | logical-system-name)>
```

Description

Display information about multicast virtual private network (MVPN) instances.

Options

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 990 on page 23105 describes the output fields for the **show pim mvpn** command. Output fields are listed in the approximate order in which they appear.

Table 990: show pim mvpn Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
VPN-Group	Multicast group address configured for the default multicast distribution tree.	All levels
Mode	Mode the tunnel is operating in: **PIM-MVPN**, **NGEN-MVPN**, **NGEN-TRANSITION** or **None**.	All levels
Tunnel	Type of tunnel: **PIM-SSM**, **PIM-SM**, **NGEN PMSI**, or **None** (VRF-only). If **NGEN-PMSI** is displayed, enter the **show mvpn instance** command for more information.	All levels

Sample Output

show pim mvpn

```
user@host>  show pim mvpn
Instance              VPN-Group          Mode                Tunnel
PIM.ce1               232.1.1.1          PIM-MVPN            PIM-SSM
```

Release Information

Command introduced in Junos OS Release 9.4.

show pim neighbors

IN THIS SECTION

- Syntax | 23106
- Syntax (EX Series Switch and the QFX Series) | 23107
- Description | 23107
- Options | 23107
- Required Privilege Level | 23107
- Output Fields | 23107
- Sample Output | 23110
- Release Information | 23112

Syntax

```
show pim neighbors
<brief | detail>
<inet | inet6>
<instance (instance-name | all)>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show pim neighbors
<brief | detail>
<inet | inet6>
<instance (instance-name | all)>
```

Description

Display information about Protocol Independent Multicast (PIM) neighbors.

Options

none	(Same as `brief`) Display standard information about PIM neighbors for all supported family addresses for the main instance.
brief | detail	(Optional) Display the specified level of output.
inet | inet6	(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.
instance (*instance-name* | all)	(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 991 on page 23108 describes the output fields for the `show pim neighbors` command. Output fields are listed in the approximate order in which they appear.

Table 991: show pim neighbors Output Fields

Field Name	Field Description	Level of Output
`Instance`	Name of the routing instance.	All levels
`Interface`	Interface through which the neighbor is reachable.	All levels
`Neighbor addr`	Address of the neighboring PIM routing device.	All levels
`IP`	IP version: `4` or `6`.	All levels
`V`	PIM version running on the neighbor: `1` or `2`.	All levels
`Mode`	PIM mode of the neighbor: `Sparse`, `Dense`, `SparseDense`, or `Unknown`. When the neighbor is running PIM version 2, this mode is always `Unknown`.	All levels
`Option`	Can be one or more of the following: • `B`—Bidirectional Capable. • `G`—Generation Identifier. • `H`—Hello Option Holdtime. • `L`—Hello Option LAN Prune Delay. • `P`—Hello Option DR Priority. • `T`—Tracking bit. • `A`—Join attribute; used in conjunction with `pim rpf-vector`.	`brief` none
`Uptime`	Time the neighbor has been operational since the PIM process was last initialized. Starting in Junos OS release 17.3R1, `uptime` is not reset during ISSU.The time format is as follows: *dd*:*hh*:*mm*:*ss* `ago` for less than a week and *nwnd*:*hh*:*mm*:*ss* `ago` for more than a week.	All levels
`Address`	Address of the neighboring PIM routing device.	`detail`

Table 991: show pim neighbors Output Fields *(Continued)*

Field Name	Field Description	Level of Output
BFD	Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: Enabled, Operational state is up, or Disabled.	detail
Hello Option Holdtime	Time for which the neighbor is available, in seconds. The range of values is 0 through 65,535.	detail
Hello Default Holdtime	Default holdtime and the time remaining if the holdtime option is not in the received hello message.	detail
Hello Option DR Priority	Designated router election priority. The range of values is 0 through 255.	detail
Hello Option Join Attribute	Appears in conjunction with the rpf-vector command. The Join attribute is included in the PIM join messages of PIM routers that can receive type 1 Encoded-Source Address.	detail
Hello Option Generation ID	9-digit or 10-digit number used to tag hello messages.	detail
Hello Option Bi-Directional PIM supported	Neighbor can process bidirectional PIM messages.	detail
Hello Option LAN Prune Delay	Time to wait before the neighbor receives prune messages, in the format delay *nnn* ms override *nnnn* ms.	detail
Join Suppression supported	Neighbor is capable of join suppression.	detail

Table 991: show pim neighbors Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Rx Join`	Information about joins received from the neighbor. • `Group`—Group addresses in the join message. • `Source`—Address of the source in the join message. • `Timeout`—Time for which the join is valid.	`detail`

Sample Output

show pim neighbors

```
user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking bit
A = Hello Option Join Attribute

Instance: PIM.master
Interface  IP V Mode        Option       Uptime Neighbor addr
ae0.0      4  2             HPLGTA       19:01:24 20.0.0.13
ae1.0      4  2             HPLGTA       19:01:24 20.0.0.149
```

show pim neighbors instance

```
user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking bit

Interface           IP V Mode        Option      Uptime Neighbor addr
at-0/3/1.0           4 2             HPLG      00:07:54 10.111.30.2
```

```
mt-1/1/0.32768       4 2             HPLG      00:07:22 10.10.47.101
so-1/0/1.0           4 2             HPLG      00:07:50 10.111.20.2
```

show pim neighbors detail

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: ae1.0

Address: 20.0.0.149, IPv4, PIM v2, sg Join Count: 0, tsg Join Count: 332
        BFD: Disabled
        Hello Option Holdtime: 105 seconds 86 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 853386212
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Join Suppression supported
        Hello Option Join Attribute supported

    Address: 20.0.0.150,    IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, tsg Join Count: 0
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 358917871
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Join Suppression supported
             Hello Option Join Attribute supported

Interface: lo0.0

    Address: 10.255.179.246,    IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, tsg Join
Count: 0
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1997462267
        Hello Option Bi-Directional PIM supported
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                      Join Suppression supported
```

show pim neighbors detail (With BFD)

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: fe-1/0/0.0
    Address: 192.168.11.1,        IPv4, PIM v2, Mode: Sparse
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 836607909
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

    Address: 192.168.11.2,        IPv4, PIM v2
        BFD: Enabled, Operational state is up
        Hello Default Holdtime: 105 seconds 104 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1907549685
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Interface: fe-1/0/1.0
    Address: 192.168.12.1,        IPv4, PIM v2
        BFD: Disabled
        Hello Default Holdtime: 105 seconds 80 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1971554705
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Release Information

Command introduced before Junos OS Release 7.4.

`inet6` and `instance` options introduced in Junos OS Release 10.0 for EX Series switches.

Support for bidirectional PIM added in Junos OS Release 12.1.

Support for the `instance all` option added in Junos OS Release 12.1.

Support for RFC 5496 (via `rpf-vector`) added in Junos OS Release 17.3R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.3R1	Starting in Junos OS release 17.3R1, `uptime` is not reset during ISSU.

show pim rps

IN THIS SECTION

- Syntax | 23113
- Syntax (EX Series Switch and the QFX Series) | 23113
- Description | 23114
- Options | 23114
- Required Privilege Level | 23114
- Output Fields | 23114
- Sample Output | 23118
- Release Information | 23123

Syntax

```
show pim rps
<brief | detail | extensive>
<group-address>
<inet | inet6>
<instance (all | instance-name)>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show pim rps
<brief | detail | extensive>
```

```
<group-address>
<inet | inet6>
<instance (all | instance-name)>
```

Description

Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs).

Options

none	Display standard information about PIM RPs for all groups and family addresses for all routing instances.
brief | detail | extensive	(Optional) Display the specified level of output.
group-address	(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.
inet | inet6	(Optional) Display information for IPv4 or IPv6 family addresses, respectively.
instance (all | *instance-name*)	(Optional) Display information about RPs for all PIM-enabled routing instances or a specific PIM-enabled routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 992 on page 23114 describes the output fields for the `show pim rps` command. Output fields are listed in the approximate order in which they appear.

Table 992: show pim rps Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels

Table 992: show pim rps Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Family or Address family	Name of the address family: **inet** (IPv4) or **inet6** (IPv6).	All levels
RP address	Address of the rendezvous point.	All levels
Type	Type of RP: • **auto-rp**—Address of the RP known through the Auto-RP protocol. • **bootstrap**—Address of the RP known through the bootstrap router protocol (BSR). • **embedded**—Address of the RP known through an embedded RP (IPv6). • **static**—Address of RP known through static configuration.	**brief** none
Holdtime	How long to keep the RP active, with time remaining, in seconds.	All levels
Timeout	How long until the local routing device determines the RP to be unreachable, in seconds.	All levels
Groups	Number of groups currently using this RP.	All levels
Group prefixes	Addresses of groups that this RP can span.	**brief** none
Learned via	Address and method by which the RP was learned.	**detail extensive**
Mode	The PIM mode of the RP: bidirectional or sparse. If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats.	All levels
Time Active	How long the RP has been active, in the format ***hh:mm:ss***.	**detail extensive**

Table 992: show pim rps Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Device Index	Index value of the order in which Junos OS finds and initializes the interface. For bidirectional RPs, the **Device Index** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Subunit	Logical unit number of the interface. For bidirectional RPs, the **Subunit** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Interface	Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively. For bidirectional RPs, the **Interface** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Group Ranges	Addresses of groups that this RP spans.	**detail extensive** *group-address*
Active groups using RP	Number of groups currently using this RP.	**detail extensive**
total	Total number of active groups for this RP.	**detail extensive**

Table 992: show pim rps Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Register State for RP	Current register state for each group:
	• **Group**—Multicast group address.
	• **Source**—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively:
	• **First Hop**—PIM-designated routing device that sent the Register message (the source address in the IP header).
	• **RP Address**—RP to which the Register message was sent (the destination address in the IP header).
	• **State**:
	On the designated router:
	  • **Send**—Sending Register messages.
	  • **Probe**—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages.
	  • **Suppress**—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to **Probe** state.
	• On the RP:
	  • **Receive**—Receiving Register messages.
		extensive

Anycast-PIM rpset	If anycast RP is configured, the addresses of the RPs in the set.	**extensive**
Anycast-PIM local address used	If anycast RP is configured, the local address used by the RP.	**extensive**

Table 992: show pim rps Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Anycast-PIM Register State	If anycast RP is configured, the current register state for each group:
	• **Group**—Multicast group address.
	• **Source**—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively.
	• **Origin**—How the information was obtained:
	  • **DIRECT**—From a local attachment
	  • **MSDP**—From the Multicast Source Discovery Protocol (MSDP)
	  • **DR**—From the designated router
		extensive

RP selected	For sparse mode and bidirectional mode, the identity of the RP for the specified group address.	*group-address*

Sample Output

show pim rps

```
user@host> show pim rps
Instance: PIM.master

Address-family INET
RP address      Type        Mode   Holdtime Timeout Groups Group prefixes
10.100.100.100 auto-rp     sparse      150     146      0 233.252.0.0/8
                                                           233.252.0.1/24
10.200.200.200 auto-rp     sparse      150     146      0 233.252.0.2/4

address-family INET6
```

show pim rps brief

The output for the `show pim rps brief` command is identical to that for the `show pim rps` command. For sample output, see "show pim rps" on page 23118.

show pim rps <group-address>

```
user@host> show pim rps 233.252.0.0
Instance: PIM.master
Instance: PIM.master

RP selected: 10.100.100.100
```

show pim rps <group-address>

```
user@host> show pim rps 233.252.0.0
Instance: PIM.master
Instance: PIM.master

RP selected: 10.100.100.100
```

show pim rps <group-address> (Bidirectional PIM)

```
user@host> show pim rps 233.252.0.1
Instance: PIM.master

233.252.0.0/16
        10.4.12.75 (Bidirectional)

RP selected: 10.4.12.75
```

show pim rps <group-address> (PIM Dense Mode)

```
user@host> show pim rps 233.252.0.1
Instance: PIM.master
```

```
Dense Mode active for group 233.252.0.1
```

show pim rps <group-address> (SSM Range Without asm-override-ssm Configured)

```
user@host> show pim rps 233.252.0.1
Instance: PIM.master

Source-specific Mode (SSM) active for group 233.252.0.1
```

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP)

```
user@host> show pim rps 233.252.0.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 233.252.0.1

233.252.0.0/16
        10.4.12.75

RP selected: 10.4.12.75
```

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP)

```
user@host> show pim rps 233.252.0.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 233.252.0.1

233.252.0.0/16
        10.4.12.75 (Bidirectional)

RP selected: (null)
```

show pim rps instance

```
user@host> show pim rps instance VPN-A
Instance: PIM.VPN-A
Address family INET
RP address              Type        Holdtime Timeout Groups Group prefixes
10.10.47.100            static             0    None      1 233.252.0.0/4

Address family INET6
```

show pim rps extensive (PIM Sparse Mode)

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.255.245.91
Learned via: static configuration
Time Active: 00:05:48
Holdtime: 45 with 36 remaining
Device Index: 122
Subunit: 32768
Interface: pd-6/0/0.32768
Group Ranges:
    233.252.0.0/4, 36s remaining
Active groups using RP:
    233.252.0.1

    total 1 groups active

Register State for RP:
Group           Source          FirstHop        RP Address      State    Timeout
233.252.0.1       192.168.195.78  10.255.14.132   10.255.245.91   Receive        0
```

show pim rps extensive (Bidirectional PIM)

```
user@host> show pim rps extensive
Instance: PIM.master
Address family INET
```

```
RP: 10.10.1.3
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
        233.252.0.0/24
        233.252.0.01/24

RP: 10.10.13.2
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
       233.252.0.3/24
       233.252.0.4/24
```

show pim rps extensive (PIM Anycast RP in Use)

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.10.10.2
Learned via: static configuration
Time Active: 00:54:52
Holdtime: 0
Device Index: 130
Subunit: 32769
Interface: pimd.32769
Group Ranges:
    233.252.0.0/4
Active groups using RP:
    233.252.0.10

    total 1 groups active

Anycast-PIM rpset:
            10.100.111.34
```

```
            10.100.111.17
            10.100.111.55

Anycast-PIM local address used: 10.100.111.1
Anycast-PIM Register State:
Group                       Source                              Origin
233.252.0.1                 10.10.95.2                          DIRECT
233.252.0.2                 10.10.95.2                          DIRECT
233.252.0.3                 10.10.70.1                          MSDP
233.252.0.4                 10.10.70.1                          MSDP
233.252.0.5                 10.10.71.1                          DR

Address family INET6

Anycast-PIM rpset:
            ab::1
            ab::2
Anycast-PIM local address used: cd::1

Anycast-PIM Register State:
Group                     Source                              Origin
::224.1.1.1               ::10.10.95.2                        DIRECT
::224.1.1.2               ::10.10.95.2                        DIRECT
::224.20.20.1             ::10.10.71.1                        DR
```

Release Information

Command introduced before Junos OS Release 7.4.

inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.

Support for bidirectional PIM added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Bidirectional PIM

show pim snooping interfaces

IN THIS SECTION

- Syntax | 23124
- Description | 23124
- Options | 23124
- Required Privilege Level | 23125
- Output Fields | 23125
- Sample Output | 23126
- Release Information | 23127

Syntax

```
show pim snooping interfaces
<brief | detail>
<instance instance-name>
<interface interface-name>
<logical-system logical-system-name>
<vlan-id vlan-identifier>
```

Description

Display information about PIM snooping interfaces.

Options

none	Display detailed information.
brief | detail	(Optional) Display the specified level of output.
instance *<instance-name>*	(Optional) Display PIM snooping interface information for the specified routing instance.

interface *<interface-name>*	(Optional) Display PIM snooping information for the specified interface only.
logical-system *logical-system-name*	(Optional) Display information about a particular logical system, or type 'all'.
vlan-id *<vlan-identifier>*	(Optional) Display PIM snooping interface information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 993 on page 23125 lists the output fields for the `show pim snooping interface` command. Output fields are listed in the approximate order in which they appear.

Table 993: show pim snooping interface Output Fields

Field Name	Field Description	Level of Output
`Instance`	Routing instance for PIM snooping.	All levels
`Learning-Domain`	Learning domain for snooping.	All levels
`Name`	Router interfaces that are part of this learning domain.	All levels
`State`	State of the interface: `Up`, or `Down`.	All levels
`IP-Version`	Version of IP used: `4` for IPv4, or `6` for IPv6.	All levels
`NbrCnt`	Number of neighboring routers connected through the specified interface.	All levels
`DR address`	IP address of the designated router.	All levels

Sample Output

show pim snooping interfaces

```
user@host> show pim snooping interfaces
Instance: vpls1
Learning-Domain: vlan-id 10
Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
ge-1/3/3.10 Up 4 1
ge-1/3/5.10 Up 4 1
ge-1/3/7.10 Up 4 1
DR address: 192.0.2.5
DR flooding is ON

Learning-Domain: vlan-id 20
Name State IP-Version NbrCnt
ge-1/3/1.20 Up 4 1
ge-1/3/3.20 Up 4 1
ge-1/3/5.20 Up 4 1
ge-1/3/7.20 Up 4 1
DR address: 192.0.2.6
DR flooding is ON
```

show pim snooping interfaces instance vpls1

```
user@host> show pim snooping interfaces instance vpls1
Instance: vpls1

Learning-Domain: vlan-id 10
Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
ge-1/3/3.10 Up 4 1
ge-1/3/5.10 Up 4 1
ge-1/3/7.10 Up 4 1
DR address: 192.0.2.5
DR flooding is ON

Learning-Domain: vlan-id 20
Name State IP-Version NbrCnt
```

```
ge-1/3/1.20 Up 4 1
ge-1/3/3.20 Up 4 1
ge-1/3/5.20 Up 4 1
ge-1/3/7.20 Up 4 1
DR address: 192.0.2.6
DR flooding is ON
```

show pim snooping interfaces interface <interface-name>

```
user@host> show pim snooping interfaces interface ge-1/3/1.10
Instance: vpls1
Learning-Domain: vlan-id 10

Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
DR address: 192.0.2.5
DR flooding is ON

Learning-Domain: vlan-id 20
DR address: 192.0.2.6
DR flooding is ON
```

show pim snooping interfaces vlan-id <vlan-id>

```
user@host> show pim snooping interfaces vlan-id 10
Instance: vpls1
Learning-Domain: vlan-id 10

Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
ge-1/3/3.10 Up 4 1
ge-1/3/5.10 Up 4 1
ge-1/3/7.10 Up 4 1
DR address: 192.0.2.5
DR flooding is ON
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

PIM Snooping for VPLS

show pim snooping join

IN THIS SECTION

- Syntax | 23128
- Description | 23128
- Options | 23128
- Required Privilege Level | 23129
- Output Fields | 23129
- Sample Output | 23131
- Release Information | 23133

Syntax

```
show pim snooping join
<brief | detail | extensive>
<instance instance-name>
<logical-system logical-system-name>
<vlan-id vlan-id>
```

Description

Display information about Protocol Independent Multicast (PIM) snooping joins.

Options

none	Display detailed information.
brief | detail | extensive	(Optional) Display the specified level of output.

instance *instance-name*	(Optional) Display PIM snooping join information for the specified routing instance.
logical-system *logical-system-name*	(Optional) Display information about a particular logical system, or type 'all'.
vlan-id *vlan-identifier*	(Optional) Display PIM snooping join information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 994 on page 23129 lists the output fields for the `show pim snooping join` command. Output fields are listed in the approximate order in which they appear.

Table 994: show pim snooping join Output Fields

Field Name	Field Description	Level of Output
`Instance`	Routing instance for PIM snooping.	All levels
`Learning-Domain`	Learning domain for PIM snooping.	All levels
`Group`	Multicast group address.	All levels
`Source`	Multicast source address: • * (wildcard value) • *<ipv4-address>* • *<ipv6-address>*	All levels

Table 994: show pim snooping join Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	PIM flags: • `bidirectional`—Bidirectional mode entry. • `dense`—Dense mode entry. • `rptree`—Entry is on the rendezvous point tree. • `sparse`—Sparse mode entry. • `spt`—Entry is on the shortest-path tree for the source. • `wildcard`—Entry is on the shared tree.	All levels
`Upstream state`	Information about the upstream interface: • `Join to RP`—Sending a join to the rendezvous point. • `Join to Source`—Sending a join to the source. • `Local RP`—Sending neither join messages nor prune messages toward the RP, because this router is the rendezvous point. • `Local Source`—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device. • `Prune to RP`—Sending a prune to the rendezvous point. • `Prune to Source`—Sending a prune to the source. **NOTE**: RP group range entries have `None` in the `Upstream state` field because RP group ranges do not trigger actual PIM join messages between routers.	All levels
`Upstream neighbor`	Information about the upstream neighbor: `Direct`, `Local`, `Unknown`, or a specific IP address. For bidirectional PIM, `Direct` means that the interface is directly connected to a subnet that contains a phantom RP address.	All levels

Table 994: show pim snooping join Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Upstream port`	RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*,G). For bidirectional PIM, `RP Link` means that the interface is directly connected to a subnet that contains a phantom RP address.	All levels
`Downstream port`	Information about downstream interfaces.	`extensive`
`Downstream neighbors`	Address of the downstream neighbor.	`extensive`
`Timeout`	Time remaining until the downstream join state is updated (in seconds).	`extensive`

Sample Output

show pim snooping join

```
user@host> show pim snooping join
Instance: vpls1

Learning-Domain: vlan-id 10
Group: 198.51.100.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.4, port: ge-1/3/5.10


Learning-Domain: vlan-id 20
Group: 198.51.100.3
Source: *
Flags: sparse,rptree,wildcard
```

```
Upstream state: None
Upstream neighbor: 203.0.113.4, port: ge-1/3/5.20
```

show pim snooping join extensive

```
user@host> show pim snooping join extensive
Instance: vpls1
Learning-Domain: vlan-id 10

Group: 198.51.100.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.4, port: ge-1/3/5.10
Downstream port: ge-1/3/1.10
Downstream neighbors:
192.0.2.2 State: Join Flags: SRW Timeout: 166

Learning-Domain: vlan-id 20
Group: 198.51.100.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 203.0.113.4, port: ge-1/3/5.20
Downstream port: ge-1/3/3.20
Downstream neighbors:
203.0.113.3 State: Join Flags: SRW Timeout: 168
```

show pim snooping join instance

```
user@host> show pim snooping join instance vpls1
Instance: vpls1

Learning-Domain: vlan-id 10
Group: 198.51.100.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.4, port: ge-1/3/5.10
```

```
Learning-Domain: vlan-id 20
Group: 198.51.100.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 203.0.113.4, port: ge-1/3/5.20
```

show pim snooping join vlan-id

```
user@host> show pim snooping join vlan-id 10
Instance: vpls1
Learning-Domain: vlan-id 10
Group: 198.51.100.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.4, port: ge-1/3/5.10
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

PIM Snooping for VPLS

show pim snooping neighbors

IN THIS SECTION

- Syntax | 23134
- Description | 23134
- Options | 23134
- Required Privilege Level | 23135

- Output Fields | 23135
- Sample Output | 23136
- Release Information | 23140

Syntax

```
show pim snooping neighbors
<brief | detail>
<instance instance-name>
<interface interface-name>
<logical-system logical-system-name>
<vlan-id vlan-identifier>
```

Description

Display information about Protocol Independent Multicast (PIM) snooping neighbors.

Options

none	Display detailed information.
brief | detail	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display PIM snooping neighbor information for the specified routing instance.
interface *interface-name*	(Optional) Display information for the specified PIM snooping neighbor interface.
logical-system *logical-system-name*	(Optional) Display information about a particular logical system, or type 'all'.
vlan-id *vlan-identifier*	(Optional) Display PIM snooping neighbor information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 995 on page 23135 lists the output fields for the `show pim snooping neighbors` command. Output fields are listed in the approximate order in which they appear.

Table 995: show pim snooping neighbors Output Fields

Field Name	Field Description	Level of Output
`Instance`	Routing instance for PIM snooping.	All levels
`Learning-Domain`	Learning domain for PIM snooping.	All levels
`Interface`	Router interface for which PIM snooping neighbor details are displayed.	All levels
`Option`	PIM snooping options available on the specified interface: • H = Hello Option Holdtime • P = Hello Option DR Priority • L = Hello Option LAN Prune Delay • G = Generation Identifier • T = Tracking Bit	All levels
`Uptime`	Time the neighbor has been operational since the PIM process was last initialized, in the format *dd*:*hh*:*mm*:*ss* `ago` for less than a week and *nwnd*:*hh*:*mm*:*ss* `ago` for more than a week.	All levels
`Neighbor addr`	IP address of the PIM snooping neighbor connected through the specified interface.	All levels
`Address`	IP address of the specified router interface.	All levels

Table 995: show pim snooping neighbors Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Hello Option Holdtime`	Time for which the neighbor is available, in seconds. The range of values is `0` through `65,535`.	`detail`
`Hello Option DR Priority`	Designated router election priority. The range of values is `0` through `4294967295`. **NOTE**: By default, every PIM interface has an equal probability (priority 1) of being selected as the DR.	`detail`
`Hello Option Generation ID`	9-digit or 10-digit number used to tag hello messages.	`detail`
`Hello Option LAN Prune Delay`	Time to wait before the neighbor receives prune messages, in the format `delay` *`nnn`* `ms override` *`nnnn`* `ms`.	`detail`

Sample Output

show pim snooping neighbors

```
user@host> show pim snooping neighbors
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Instance: vpls1
Learning-Domain: vlan-id 10

Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:43:33 192.0.2.2
ge-1/3/3.10 HPLGT 00:43:33 192.0.2.3
ge-1/3/5.10 HPLGT 00:43:33 192.0.2.4
ge-1/3/7.10 HPLGT 00:43:33 192.0.2.5

Learning-Domain: vlan-id 20
```

```
Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:43:33 192.0.2.12
ge-1/3/3.20 HPLGT 00:43:33 192.0.2.13
ge-1/3/5.20 HPLGT 00:43:33 192.0.2.14
ge-1/3/7.20 HPLGT 00:43:33 192.0.2.15
```

show pim snooping neighbors detail

```
user@host> show pim snooping neighbors detail
Instance: vpls1
Learning-Domain: vlan-id 10

Interface: ge-1/3/1.10
Address: 192.0.2.2
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 83 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 830908833
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/3.10
Address: 192.0.2.3
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 97 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 2056520742
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/5.10
Address: 192.0.2.4
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 81 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1152066227
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/7.10
Address: 192.0.2.5
```

```
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 96 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1113200338
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
Learning-Domain: vlan-id 20

Interface: ge-1/3/1.20
Address: 192.0.2.12
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 81 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 963205167
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/3.20
Address: 192.0.2.13
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 104 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 166921538
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/5.20
Address: 192.0.2.14
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 88 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 789422835
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/7.20
Address: 192.0.2.15
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 88 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1563649680
```

```
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
```

show pim snooping neighbors instance

```
user@host> show pim snooping neighbors instance vpls1
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Instance: vpls1
Learning-Domain: vlan-id 10

Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:46:03 192.0.2.2
ge-1/3/3.10 HPLGT 00:46:03 192.0.2.3
ge-1/3/5.10 HPLGT 00:46:03 192.0.2.4
ge-1/3/7.10 HPLGT 00:46:03 192.0.2.5


Learning-Domain: vlan-id 20

Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:46:03 192.0.2.12
ge-1/3/3.20 HPLGT 00:46:03 192.0.2.13
ge-1/3/5.20 HPLGT 00:46:03 192.0.2.14
ge-1/3/7.20 HPLGT 00:46:03 192.0.2.15
```

show pim snooping neighbors interface

```
user@host> show pim snooping neighbors interface ge-1/3/1.20
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20
```

```
Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:48:04 192.0.2.12
```

show pim snooping neighbors vlan-id

```
user@host> show pim snooping neighbors vlan-id 10
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Instance: vpls1
Learning-Domain: vlan-id 10

Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:49:12 192.0.2.2
ge-1/3/3.10 HPLGT 00:49:12 192.0.2.3
ge-1/3/5.10 HPLGT 00:49:12 192.0.2.4
ge-1/3/7.10 HPLGT 00:49:12 192.0.2.5
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

- Configuring Interface Priority for PIM Designated Router Selection
- Modifying the PIM Hello Interval
- PIM Snooping for VPLS
- show pim neighbors

show pim snooping statistics

IN THIS SECTION

- Syntax | 23141
- Description | 23141
- Options | 23141
- Required Privilege Level | 23142
- Output Fields | 23142
- Sample Output | 23143
- Release Information | 23147

Syntax

```
show pim snooping statistics
<instance instance-name>
<interface interface-name>
<logical-system logical-system-name>
<vlan-id vlan-id>
```

Description

Display Protocol Independent Multicast (PIM) snooping statistics.

Options

none	Display PIM statistics.
instance *instance-name*	(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM) snooping.
interface *interface-name*	(Optional) Display statistics about the specified interface for PIM snooping.

logical-system *logical-system-name*	(Optional) Display information about a particular logical system, or type 'all'.
vlan-id *vlan-identifier*	(Optional) Display PIM snooping statistics information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 996 on page 23142 lists the output fields for the `show pim snooping statistics` command. Output fields are listed in the approximate order in which they appear.

Table 996: show pim snooping statistics Output Fields

Field Name	Field Description	Level of Output
`Instance`	Routing instance for PIM snooping.	All levels
`Learning-Domain`	Learning domain for PIM snooping.	All levels
`Tx J/P messages`	Total number of transmitted join/prune packets.	All levels
`RX J/P messages`	Total number of received join/prune packets.	All levels
`Rx J/P messages -- seen`	Number of join/prune packets seen but not received on the upstream interface.	All levels
`Rx J/P messages -- received`	Number of join/prune packets received on the downstream interface.	All levels
`Rx Hello messages`	Total number of received hello packets.	All levels
`Rx Version Unknown`	Number of packets received with an unknown version number.	All levels

Table 996: show pim snooping statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Rx Neighbor Unknown`	Number of packets received from an unknown neighbor.	All levels
`Rx Upstream Neighbor Unknown`	Number of packets received with unknown upstream neighbor information.	All levels
`Rx Bad Length`	Number of packets received containing incorrect length information.	All levels
`Rx J/P Busy Drop`	Number of join/prune packets dropped while the router is busy.	All levels
`Rx J/P Group Aggregate 0`	Number of join/prune packets received containing the aggregate group information.	All levels
`Rx Malformed Packet`	Number of malformed packets received.	All levels
`Rx No PIM Interface`	Number of packets received without the interface information.	All levels
`Rx No Upstream Neighbor`	Number of packets received without upstream neighbor information.	All levels
`Rx Unknown Hello Option`	Number of hello packets received with unknown options.	All levels

Sample Output

show pim snooping statistics

```
user@host> show pim snooping statistics
Instance: vpls1
Learning-Domain: vlan-id 10
```

```
Tx J/P messages 0
RX J/P messages 8
Rx J/P messages -- seen 0
Rx J/P messages -- received 8
Rx Hello messages 37
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0


Learning-Domain: vlan-id 20

Tx J/P messages 0
RX J/P messages 2
Rx J/P messages -- seen 0
Rx J/P messages -- received 2
Rx Hello messages 39
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0
```

show pim snooping statistics instance

```
user@host> show pim snooping statistics instance vpls1
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 9
Rx J/P messages -- seen 0
Rx J/P messages -- received 9
Rx Hello messages 45
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0


Learning-Domain: vlan-id 20

Tx J/P messages 0
RX J/P messages 3
Rx J/P messages -- seen 0
Rx J/P messages -- received 3
Rx Hello messages 47
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
```

```
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0
```

show pim snooping statistics interface

```
user@host> show pim snooping statistics interface ge-1/3/1.20
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20

PIM Interface statistics for ge-1/3/1.20
Tx J/P messages 0
RX J/P messages 0
Rx J/P messages -- seen 0
Rx J/P messages -- received 0
Rx Hello messages 13
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
```

show pim snooping statistics vlan-id

```
user@host> show pim snooping statistics vlan-id 10
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 11
Rx J/P messages -- seen 0
Rx J/P messages -- received 11
Rx Hello messages 64
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
```

```
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

PIM Snooping for VPLS

clear pim snooping statistics | 3817

show pim source

IN THIS SECTION

- Syntax | 23148
- Syntax (EX Series Switch and the QFX Series) | 23148
- Description | 23148
- Options | 23148
- Required Privilege Level | 23149
- Output Fields | 23149
- Sample Output | 23149
- Release Information | 23151

Syntax

```
show pim source
<brief | detail>
<inet | inet6>
<instance instance-name>
<logical-system (all | logical-system-name)>
<source-prefix>
```

Syntax (EX Series Switch and the QFX Series)

```
show pim source
<brief | detail>
<inet | inet6>
<instance instance-name>
<source-prefix>
```

Description

Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state.

Options

none	Display standard information about the PIM RPF state for all supported family addresses for all routing instances.
brief | detail	(Optional) Display the specified level of output.
inet | inet6	(Optional) Display information for IPv4 or IPv6 family addresses, respectively.
instance *instance-name*	(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
source-prefix	(Optional) Display the state for source RPF states in the given range.

Required Privilege Level

view

Output Fields

Table 997 on page 23149 describes the output fields for the `show pim source` command. Output fields are listed in the approximate order in which they appear.

Table 997: show pim source Output Fields

Field Name	Field Description
`Instance`	Name of the routing instance.
`Source`	Address of the source or reverse path.
`Prefix/length`	Prefix and prefix length for the route used to reach the RPF address.
`Upstream Protocol`	Protocol toward the source address.
`Upstream interface`	RPF interface toward the source address. A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.
`Upstream Neighbor`	Address of the RPF neighbor used to reach the source address. The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.

Sample Output

show pim source

```
user@host> show pim source
Instance: PIM.master Family: INET
```

```
Source 10.255.14.144
    Prefix 10.255.14.144/32
    Upstream interface Local
    Upstream neighbor Local

Source 10.255.70.15
    Prefix 10.255.70.15/32
    Upstream interface so-1/0/0.0
    Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6
```

show pim source brief

The output for the `show pim source brief` command is identical to that for the `show pim source` command. For sample output, see "show pim source" on page 23149.

show pim source detail

```
user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
    Prefix 10.255.14.144/32
    Upstream interface Local
    Upstream neighbor Local
    Active groups:233.252.0.0
        233.252.0.1
        233.252.0.1

Source 10.255.70.15
    Prefix 10.255.70.15/32
    Upstream interface so-1/0/0.0
    Upstream neighbor 10.111.10.2
    Active groups:233.252.0.1

Instance: PIM.master Family: INET6
```

show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim source
Instance: PIM.master Family: INET

Source 10.1.1.1
    Prefix 10.1.1.1/32
    Upstream interface Local
    Upstream neighbor Local

Source 10.2.7.7
    Prefix 10.2.7.0/24
    Upstream protocol MLDP
    Upstream interface Pseudo MLDP
    Upstream neighbor MLDP LSP root <10.1.1.2>

Source 192.168.219.11
    Prefix 192.168.219.0/28
    Upstream protocol MLDP
    Upstream interface Pseudo MLDP
       Upstream neighbor via MLDP-inband
      Upstream interface fe-1/3/0.0
    Upstream neighbor 192.168.140.1
    Upstream neighbor MLDP LSP root <10.1.1.2>

Instance: PIM.master Family: INET6
Source 2001:db8::1:2:7:7
    Prefix 2001:db8::1:2:7:0/120
    Upstream protocol MLDP
    Upstream interface Pseudo MLDP
       Upstream neighbor via MLDP-inband
      Upstream interface fe-1/3/0.0
    Upstream neighbor 192.168.140.1
Upstream neighbor MLDP LSP root <10.1.1.2>
```

Release Information

Command introduced before Junos OS Release 7.4.

`inet6` and `instance` options introduced in Junos OS Release 10.0 for EX Series switches.

show pim statistics

IN THIS SECTION

- Syntax | 23152
- Syntax (EX Series Switch and the QFX Series) | 23152
- Description | 23152
- Options | 23153
- Required Privilege Level | 23153
- Output Fields | 23153
- Sample Output | 23163
- Release Information | 23170

Syntax

```
show pim statistics
<inet |inet6)>
<instance (all | instance-name)>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show pim statistics
<inet |inet6)>
<instance (all | instance-name)>
<interface interface-name>
```

Description

Display Protocol Independent Multicast (PIM) statistics.

Options

none	Display PIM statistics.
inet | inet6	(Optional) Display IPv4 or IPv6 PIM statistics, respectively.
instance (all | *instance-name*)	(Optional) Display statistics for all routing instances or a specific routing instance enabled by Protocol Independent Multicast (PIM).
interface *interface-name*	(Optional) Display statistics about the specified interface.
logical-system (all | *logical-system-name)*	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 998 on page 23153 describes the output fields for the `show pim statistics` command. Output fields are listed in the approximate order in which they appear.

Table 998: show pim statistics Output Fields

Field Name	Field Description
Instance	Name of the routing instance. This field only appears if you specify an interface, for example: • `inet interface` *`interface-name`* • `inet6 interface` *`interface-name`* • `interface` *`interface-name`*

Table 998: show pim statistics Output Fields *(Continued)*

Field Name	Field Description
Family	Output is for IPv4 or IPv6 PIM statistics. `INET` indicates IPv4 statistics, and `INET6` indicates IPv6 statistics. This field only appears if you specify an interface, for example: • `inet interface` *`interface-name`* • `inet6 interface` *`interface-name`* • `interface` *`interface-name`*
PIM statistics	PIM statistics for all interfaces or for the specified interface.
PIM message type	Message type for which statistics are displayed.
Received	Number of received statistics.
Sent	Number of messages sent of a certain type.
Rx errors	Number of received packets that contained errors.
V2 Hello	PIM version 2 hello packets.
V2 Register	PIM version 2 register packets.
V2 Register Stop	PIM version 2 register stop packets.
V2 Join Prune	PIM version 2 join and prune packets.
V2 Bootstrap	PIM version 2 bootstrap packets.
V2 Assert	PIM version 2 assert packets.

Table 998: show pim statistics Output Fields *(Continued)*

Field Name	Field Description
V2 Graft	PIM version 2 graft packets.
V2 Graft Ack	PIM version 2 graft acknowledgment packets.
V2 Candidate RP	PIM version 2 candidate RP packets.
V2 State Refresh	PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh. State refresh is an extension to PIM-DM. It not supported in Junos OS.
V2 DF Election	PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election.
V2 Sparse Join	Count of PIM version 2 sparse join packets per (S,G) or (*,G) multicast route.
V2 Sparse Prune	Count of PIM version 2 sparse prune packets per (S,G) or (*,G) multicast route.
V1 Query	PIM version 1 query packets.
V1 Register	PIM version 1 register packets.
V1 Register Stop	PIM version 1 register stop packets.
V1 Join Prune	PIM version 1 join and prune packets.
V1 RP Reachability	PIM version 1 RP reachability packets.
V1 Assert	PIM version 1 assert packets.
V1 Graft	PIM version 1 graft packets.

Table 998: show pim statistics Output Fields *(Continued)*

Field Name	Field Description
V1 Graft Ack	PIM version 1 graft acknowledgment packets.
AutoRP Announce	Auto-RP announce packets.
AutoRP Mapping	Auto-RP mapping packets.
AutoRP Unknown type	Auto-RP packets with an unknown type.
Anycast Register	Auto-RP announce packets.
Anycast Register Stop	Auto-RP announce packets.
Global Statistics	Summary of PIM statistics for all interfaces.
Hello dropped on neighbor policy	Number of hello packets dropped because of a configured neighbor policy.
Unknown type	Number of PIM control packets received with an unknown type.
V1 Unknown type	Number of PIM version 1 control packets received with an unknown type.
Unknown Version	Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.
Neighbor unknown	Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.
Bad Length	Number of PIM control packets received for which the packet size does not match the PIM length field in the packet.
Bad Checksum	Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet.

Table 998: show pim statistics Output Fields *(Continued)*

Field Name	Field Description
Bad Receive If	Number of PIM control packets received on an interface that does not have PIM configured.
Rx Bad Data	Number of PIM control packets received that contain data for TCP Bad register packets.
Rx Intf disabled	Number of PIM control packets received on an interface that has PIM disabled.
Rx V1 Require V2	Number of PIM version 1 control packets received on an interface configured for PIM version 2.
Rx V2 Require V1	Number of PIM version 2 control packets received on an interface configured for PIM version 1.
Rx Register not RP	Number of PIM register packets received when the routing device is not the RP for the group.
Rx Register no route	Number of PIM register packets received when the RP does not have a unicast route back to the source.
Rx Register no decap if	Number of PIM register packets received when the RP does not have a de-encapsulation interface.
Null Register Timeout	Number of NULL register timeout packets.
RP Filtered Source	Number of PIM packets received when the routing device has a source address filter configured for the RP.
Rx Unknown Reg Stop	Number of register stop messages received with an unknown type.
Rx Join/Prune no state	Number of join and prune messages received for which the routing device has no state.

Table 998: show pim statistics Output Fields *(Continued)*

Field Name	Field Description
Rx Join/Prune on upstream if	Number of join and prune messages received on the interface used to reach the upstream routing device, toward the RP.
Rx Join/Prune for invalid group	Number of join or prune messages received for invalid multicast group addresses.
Rx Join/Prune messages dropped	Number of join and prune messages received and dropped.
Rx sparse join for dense group	Number of PIM sparse mode join messages received for a group that is configured for dense mode.
Rx Graft/Graft Ack no state	Number of graft and graft acknowledgment messages received for which the router or switch has no state.
Rx Graft on upstream if	Number of graft messages received on the interface used to reach the upstream routing device, toward the RP.
Rx CRP not BSR	Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap.
Rx BSR when BSR	Number of BSR messages received in which the PIM message type is Bootstrap.
Rx BSR not RPF if	Number of BSR messages received on an interface that is not the RPF interface.
Rx unknown hello opt	Number of PIM hello packets received with options that Junos OS does not support.
Rx data no state	Number of PIM control packets received for which the routing device has no state for the data type.

Table 998: show pim statistics Output Fields *(Continued)*

Field Name	Field Description
Rx RP no state	Number of PIM control packets received for which the routing device has no state for the RP.
Rx aggregate	Number of PIM aggregate MDT packets received.
Rx malformed packet	Number of PIM control packets received with a malformed IP unicast or multicast address family.
No RP	Number of PIM control packets received with no RP address.
No register encap if	Number of PIM register packets received when the first-hop routing device does not have an encapsulation interface.
No route upstream	Number of PIM control packets received when the routing device does not have a unicast route to the the interface used to reach the upstream routing device, toward the RP.
Nexthop Unusable	Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down.
RP mismatch	Number of PIM control packets received for which the routing device has an RP mismatch.
RP mode mismatch	RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages.
RPF neighbor unknown	Number of PIM control packets received for which the routing device has an unknown RPF neighbor for the source.
Rx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.

Table 998: show pim statistics Output Fields *(Continued)*

Field Name	Field Description
Tx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Embedded-RP invalid addr	Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.
Embedded-RP limit exceed	Number of times the limit configured with the `maximum-rps` statement is exceeded. The `maximum-rps` statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.
Embedded-RP added	Number of packets in which the embedded RP for IPv6 is added. The following receive events trigger extraction of an IPv6 embedded RP address on the routing device: • Multicast Listener Discovery (MLD) report for an embedded RP multicast group address • PIM join message with an embedded RP multicast group address • Static embedded RP multicast group address associated with an interface • Packets sent to an embedded RP multicast group address received on the DR An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.
Embedded-RP removed	Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.
Rx Register msgs filtering drop	Number of received register messages dropped because of a filter configured for PIM register messages.
Tx Register msgs filtering drop	Number of register messages dropped because of a filter configured for PIM register messages.

Table 998: show pim statistics Output Fields *(Continued)*

Field Name	Field Description
Rx Bidir Join/Prune on non-Bidir if	Error counter for join and prune messages received on non-bidirectional PIM interfaces.
Rx Bidir Join/Prune on non-DF if	Error counter for join and prune messages received on non-designated forwarder interfaces.
V4 (S,G) Maximum	Maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted.
V4 (S,G) Accepted	Number of accepted (S,G) IPv4 multicast routes.
V4 (S,G) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device).
V4 (S,G) Log Interval	Time (in seconds) between consecutive log messages.
V6 (S,G) Maximum	Maximum number of (S,G) IPv6 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted.
V6 (S,G) Accepted	Number of accepted (S,G) IPv6 multicast routes.
V6 (S,G) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv6 multicast routes accepted by the device).
V6 (S,G) Log Interval	Time (in seconds) between consecutive log messages.
V4 (grp-prefix, RP) Maximum	Maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.
V4 (grp-prefix, RP) Accepted	Number of accepted group-to-RP IPv4 multicast mappings.

Table 998: show pim statistics Output Fields *(Continued)*

Field Name	Field Description
V4 (grp-prefix, RP) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device).
V4 (grp-prefix, RP) Log Interval	Time (in seconds) between consecutive log messages.
V6 (grp-prefix, RP) Maximum	Maximum number of group-to RP IPv6 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.
V6 (grp-prefix, RP) Accepted	Number of accepted group-to-RP IPv6 multicast mappings.
V6 (grp-prefix, RP) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv6 multicast mappings accepted by the device).
V6 (grp-prefix, RP) Log Interval	Time (in seconds) between consecutive log messages.
V4 Register Maximum	Maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP.
V4 Register Accepted	Number of accepted IPv4 PIM registers.
V4 Register Threshold	Threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device).
V4 Register Log Interval	Time (in seconds) between consecutive log messages.
V6 Register Maximum	Maximum number of IPv6 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP.
V6 Register Accepted	Number of accepted IPv6 PIM registers.

Table 998: show pim statistics Output Fields *(Continued)*

Field Name	Field Description
V6 Register Threshold	Threshold at which a warning message is logged (percentage of the maximum number of IPv6 PIM registers accepted by the device).
V6 Register Log Interval	Time (in seconds) between consecutive log messages.
(*,G) Join drop due to SSM range check	PIM join messages that are dropped because the multicast addresses are outside of the SSM address range of 232.0.0.0 through 232.255.255.255. You can extend the accepted SSM address range by configuring the `ssm-groups` statement.

Sample Output

show pim statistics

```
user@host> show pim statistics
PIM Message type        Received       Sent  Rx errors
V2 Hello                      15         32          0
V2 Register                    0        362          0
V2 Register Stop             483          0          0
V2 Join Prune                 18        518          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V2 State Refresh               0          0          0
V2 DF Election                 0          0          0
V2 Sparse Join                 0          0          0
V2 Sparse Prune                0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
```

```
V1 Graft                        0          0          0
V1 Graft Ack                    0          0          0
AutoRP Announce                 0          0          0
AutoRP Mapping                  0          0          0
AutoRP Unknown type             0
Anycast Register                0          0          0
Anycast Register Stop           0          0          0

Global Statistics

Hello dropped on neighbor policy              0
Unknown type                                  0
V1 Unknown type                               0
Unknown Version                               0
ipv4 BSR pkt drop due to excessive rate       0
ipv6 BSR pkt drop due to excessive rate       0
Neighbor unknown                              0
Bad Length                                    0
Bad Checksum                                  0
Bad Receive If                                0
Rx Bad Data                                   0
Rx Intf disabled                              0
Rx V1 Require V2                              0
Rx V2 Require V1                              0
Rx Register not RP                            0
Rx Register no route                          0
Rx Register no decap if                       0
Null Register Timeout                         0
RP Filtered Source                            0
Rx Unknown Reg Stop                           0
Rx Join/Prune no state                        0
Rx Join/Prune on upstream if                  0
Rx Join/Prune for invalid group               5
Rx Join/Prune messages dropped                0
Rx sparse join for dense group                0
Rx Graft/Graft Ack no state                   0
Rx Graft on upstream if                       0
Rx CRP not BSR                                0
Rx BSR when BSR                               0
Rx BSR not RPF if                             0
Rx unknown hello opt                          0
Rx data no state                              0
Rx RP no state                                0
```

```
Rx aggregate                                0
Rx malformed packet                         0
Rx illegal TTL                              0
Rx illegal destination address              0
No RP                                       0
No register encap if                        0
No route upstream                           0
Nexthop Unusable                            0
RP mismatch                                 0
RP mode mismatch                            0
RPF neighbor unknown                        0
Rx Joins/Prunes filtered                    0
Tx Joins/Prunes filtered                    0
Embedded-RP invalid addr                    0
Embedded-RP limit exceed                    0
Embedded-RP added                           0
Embedded-RP removed                         0
Rx Register msgs filtering drop             0
Tx Register msgs filtering drop             0
Rx Bidir Join/Prune on non-Bidir if         0
Rx Bidir Join/Prune on non-DF if            0
(*,G) Join drop due to SSM range check      0
```

show pim statistics inet interface <interface-name>

```
user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          4          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V1 Query                       0          0          0
```

```
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
AutoRP Announce                0          0          0
AutoRP Mapping                 0          0          0
AutoRP Unknown type            0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0
```

show pim statistics inet6 interface <interface-name>

```
user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          4          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0
```

show pim statistics instance <instance-name>

```
user@host> show pim statistics instance VPN-A
PIM Message type        Received       Sent  Rx errors
V2 Hello                      31         37          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
```

```
V2 Join Prune                  0         16          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V2 State Refresh               0          0          0
V2 DF Election                 0          0          0
V2 Sparse Join                 0          0          0
V2 Sparse Prune                0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
AutoRP Announce                0          0          0
AutoRP Mapping                 0          0          0
AutoRP Unknown type            0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0

Global Statistics

Hello dropped on neighbor policy              0
Unknown type                                  0
V1 Unknown type                               0
Unknown Version                               0
Neighbor unknown                              0
Bad Length                                    0
Bad Checksum                                  0
Bad Receive If                                0
Rx Bad Data                                   0
Rx Intf disabled                              0
Rx V1 Require V2                              0
Rx V2 Require V1                              0
Rx Register not RP                            0
Rx Register no route                          0
Rx Register no decap if                       0
Null Register Timeout                         0
RP Filtered Source                            0
```

```
Rx Unknown Reg Stop                          0
Rx Join/Prune no state                       0
Rx Join/Prune on upstream if                 0
Rx Join/Prune for invalid group              0
Rx Join/Prune messages dropped               0
Rx sparse join for dense group               0
Rx Graft/Graft Ack no state                  0
Rx Graft on upstream if                      0
Rx CRP not BSR                               0
Rx BSR when BSR                              0
Rx BSR not RPF if                            0
Rx unknown hello opt                         0
Rx data no state                             0
Rx RP no state                               0
Rx aggregate                                 0
Rx malformed packet                          0
Rx illegal TTL                               0
Rx illegal destination address               0
No RP                                        0
No register encap if                         0
No route upstream                           28
Nexthop Unusable                             0
RP mismatch                                  0
RP mode mismatch                             0
RPF neighbor unknown                         0
Rx Joins/Prunes filtered                     0
Tx Joins/Prunes filtered                     0
Embedded-RP invalid addr                     0
Embedded-RP limit exceed                     0
Embedded-RP added                            0
Embedded-RP removed                          0
Rx Register msgs filtering drop              0
Tx Register msgs filtering drop              0
Rx Bidir Join/Prune on non-Bidir if          0
Rx Bidir Join/Prune on non-DF if             0
V4 (S,G) Maximum                            10
V4 (S,G) Accepted                            9
V4 (S,G) Threshold                          80
V4 (S,G) Log Interval                       80
V6 (S,G) Maximum                             8
V6 (S,G) Accepted                            8
V6 (S,G) Threshold                          50
V6 (S,G) Log Interval                      100
```

```
V4 (grp-prefix, RP) Maximum                100
V4 (grp-prefix, RP) Accepted                 5
V4 (grp-prefix, RP) Threshold               80
V4 (grp-prefix, RP) Log Interval            10
V6 (grp-prefix, RP) Maximum                 20
V6 (grp-prefix, RP) Accepted                 0
V6 (grp-prefix, RP) Threshold               90
V6 (grp-prefix, RP) Log Interval            20
V4 Register Maximum                        100
V4 Register Accepted                        10
V4 Register Threshold                       80
V4 Register Log Interval                    10
V6 Register Maximum                         20
V6 Register Accepted                         0
V6 Register Threshold                       90
V6 Register Log Interval                    20
(*,G) Join drop due to SSM range check       0
```

show pim statistics interface <interface-name>

```
user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          3          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
```

```
V1 Graft                        0          0          0
V1 Graft Ack                    0          0          0
AutoRP Announce                 0          0          0
AutoRP Mapping                  0          0          0
AutoRP Unknown type             0
Anycast Register                0          0          0
Anycast Register Stop           0          0          0

Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                        0          3          0
V2 Register                     0          0          0
V2 Register Stop                0          0          0
V2 Join Prune                   0          0          0
V2 Bootstrap                    0          0          0
V2 Assert                       0          0          0
V2 Graft                        0          0          0
V2 Graft Ack                    0          0          0
V2 Candidate RP                 0          0          0
Anycast Register                0          0          0
Anycast Register Stop           0          0          0
```

Release Information

Command introduced before Junos OS Release 7.4.

`inet6` and `instance` options introduced in Junos OS Release 10.0 for EX Series switches.

Support for bidirectional PIM added in Junos OS Release 12.1.

RELATED DOCUMENTATION

clear pim statistics | 3819

show poe controller

IN THIS SECTION

- Syntax | 23171
- Description | 23171
- Required Privilege Level | 23171
- Output Fields | 23171
- Sample Output | 23173
- Release Information | 23175

Syntax

```
show poe controller
```

Description

Display configuration and status of the PoE controllers.

Required Privilege Level

view

Output Fields

Table 999 on page 23172 lists the output fields for the `show poe controller` command. Output fields are listed in the approximate order in which they appear.

Table 999: show poe controller Output Fields

Field Name	Field Description
Controller index	PoE controller number: • 0 for EX2200, EX3200, standalone EX3300, standalone EX4200 switches, standalone EX4300 switches, and ACX2000 routers. • Member ID for switches in an EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4300 Virtual Chassis, a mixed EX4200 and EX4500 Virtual Chassis. • Slot number for line cards with a PoE controller in an EX6200 or EX8200 switch.
Maximum power	The maximum PoE power consumption for the switch or line card. This is the total amount of power available to the PoE controller to allocate to the PoE ports.
Power reserved	Total amount of power reserved by the PoE ports at the time the command is executed. This value, which represents actual power reserved, cannot exceed the value for `Maximum power`.
Power consumption	Total amount of power being consumed by the PoE ports at the time the command is executed. This value, which represents actual power consumption, cannot exceed the value for `Maximum power`.
Guard Band	Amount of power that has been placed in reserve for power demand spikes and that cannot be allocated to a PoE interface.
Management	Power management mode: **class** or **static** or **high-power**. **NOTE**: The mode **high-power** is available on only ACX2000 routers.

Table 999: show poe controller Output Fields *(Continued)*

Field Name	Field Description
Status	Status of the PoE controller: • **AF_ENHANCE**—Controller supports enhanced PoE. The maximum power per PoE port is 18.6 W in static mode (15.4 W in class mode). • **DEVICE FAIL**—Software download to the controller has failed or the PoE controller is not initialized because of a hardware failure. • **DOWNLOAD_INIT**—Software download to the controller is in the initial phase. • **AF_MODE**—Controller supports standard IEEE 802.3af. The maximum power per PoE port is 15.4 W. • **AT/AF COMBO**—Controller supports a mix of standard IEEE 802.3af and IEEE 802.3at (PoE+) ports. The maximum power per port is 30 W for IEEE 802.3at (PoE+) ports and 15.4 W for the IEEE 802.3af ports. • **AT_MODE**—Controller supports IEEE 802.3at (PoE+). The maximum power per PoE port is 30 W. • **BT_MODE**—Controller supports standard IEEE 802.3bt (PoE). The maximum power per PoE port is 90 W. • **SW_DOWNLOAD (*n*%)**—Software download to the controller is in progress.
Lldp Priority	Link Layer Discovery Protocol (LLDP) priority operating state. The state can be **Enabled** or **Disabled**. LLDP priority enables the PoE controller to assign interfaces the power priority provided by the connected powered device by using LLDP power negotiation rather than the power priority configured on the switch interface.

Sample Output

show poe controller (EX3200 Switch)

```
user@switch> show poe controller
Controller  Maximum   Power         Guard    Management  Status      Lldp
index       power     consumption   band                             Priority
   0        130.00W   81.20W         10W     Static      AF_ENHANCE  Disabled
```

```
   0       1800.00W    0.00W          0W     Class       DEVICE FAIL Disabled
   1       1800.00W    0.00W          0W     Class       BT_MODE     Disabled
```

show poe controller (EX4400 Switch)

```
user@switch> show poe controller
Controller  Maximum   Power      Power         Guard    Management   Status        Lldp
index       power     reserved   consumption   band                               Priority
   0        1650W     43.00W     8.00W           0W     Class        BT_MODE       Disabled
```

show poe controller (EX8200 Switch)

```
user@switch> show poe controller
Controller  Maximum   Power         Guard    Management  Status       Lldp
index       power     consumption   band                              Priority
   0        792.00W   603.50W         0W     Class       AT/AF COMBO  Disabled
   4        915.00W   781.00W         0W     Class       AT/AF COMBO  Disabled
   7        915.00W   0.00W           0W     Class       AT/AF COMBO  Disabled
```

show poe controller (Controller Software Upgrade in Progress)

```
user@switch> show poe controller

Controller  Maximum   Power         Guard    Management  Status            Lldp
index       power     consumption   band                                   Priority
   0        130.00W   0.00W           0W     Static      AF_ENHANCE        Disabled
   8**      130.00W   0.00W           0W     Static      SW_DOWNLOAD(10%)  Disabled
  **New PoE software upgrade available.
Use 'request system firmware upgrade poe fpc-slot <slot>'
 This procedure will take around 10 minutes (recommended to be performed during maintenance)
```

show poe controller (ACX2000 Router)

```
user@host> show poe controller

Controller  Maximum   Power         Guard    Management   Status        Lldp
```

```
index       power     consumption   band                                Priority
   0        130.0 W   14.2 W         0 W      high-power   UP
```

Release Information

Command introduced in Junos OS Release 9.0.

Command introduced in Junos OS Release 12.2 for ACX2000 routers.
Statement introduced in Junos OS Release 17.2R1 for a Junos Fusion Provider Edge.

RELATED DOCUMENTATION

- show poe interface | 23177
- request system firmware upgrade poe
- Verifying PoE Configuration and Status (CLI Procedure)
- *Verifying PoE Configuration and Status for a Junos Fusion (CLI Procedure)*
- Monitoring PoE Power Consumption (CLI Procedure)
- Upgrading the PoE Controller Software

show poe controller (View)

IN THIS SECTION

- Syntax | 23176
- Description | 23176
- Options | 23176
- Required Privilege Level | 23176
- Output Fields | 23176
- Sample Output | 23177
- Release Information | 23177

Syntax

```
show poe controller
```

Description

Display the status of the Power over Ethernet (PoE) controller.

Options

`none`—Display general parameters of the PoE software module controller.

Required Privilege Level

View

Output Fields

Table 1000 on page 23176 lists the output fields for the `show poe controller` command. Output fields are listed in the approximate order in which they appear.

Table 1000: show poe controller Output Fields

Field name	Field Description
Controller-index	Identifies the controller.
Maximum-power	Specifies the maximum power that can be provided by the SRX Series Firewall to PoE ports.
Power-consumption	Specifies the total amount of power allocated to the PoE ports.
Guard-band	Shows the guard band configured on the controller.
Management	Shows the power management mode.

Sample Output

show poe controller

user@host>**show poe controller**

```
Controller  Maximum   Power              Guard band  Management
index       power     consumption
   0        150.0 W    0.0 W               0 W         Static
```

Release Information

Command introduced in Junos OS Release 9.5.

show poe interface

IN THIS SECTION

- Syntax | 23177
- Description | 23178
- Options | 23178
- Required Privilege Level | 23178
- Output Fields | 23178
- Sample Output | 23183
- Release Information | 23187

Syntax

```
show poe interface
<fpc-slot number>
<interface-name>
```

Description

Display the status of PoE interfaces.

Options

none—Display status of all PoE interfaces on the switch or router.

`fpc-slot` *`number`*—(Optional) (EX6200 or EX8200 switches only) Display the status of the PoE interfaces on the specified line card.

`interface-name`—(Optional) Display the status of a specific PoE interface on the switch.

Required Privilege Level

view

Output Fields

Table 1001 on page 23178 lists the output fields for the `show poe interface` command. Output fields are listed in the approximate order in which they appear.

Table 1001: show poe interface Output Fields

Field Name (All Interfaces Output)	Field Name (Single Interface Output)	Field Description
Interface	**PoE Interface**	Interface name.
Admin status	**Administrative status**	Administrative state of the PoE interface: **Enabled** or **Disabled**. If the PoE interface is disabled, it can provide network connectivity, but it cannot provide power to connected devices.

Table 1001: show poe interface Output Fields *(Continued)*

Field Name (All Interfaces Output)	Field Name (Single Interface Output)	Field Description
Oper status	**Operational status**	Operational state of the PoE interface: • `ON`—The interface is currently supplying power to a powered device. • `OFF`—PoE is enabled on the interface, but the interface is not currently supplying power to a powered device. • `FAULT`—PoE interface is in the `OFF` state due to a fault condition. • `Disabled`—PoE is disabled on the interface.

Table 1001: show poe interface Output Fields *(Continued)*

Field Name (All Interfaces Output)	Field Name (Single Interface Output)	Field Description
	`Operational status detail`	Additional information for troubleshooting the operational state of the PoE interface: • `Admin up but disabled on hardware`—The interface is disabled due to power budget unavailability. • `Overload`—Interface is in the fault condition. • `IEEE PD Detected`—The interface is providing power to the powered device. • `Detection In Progress`—Detection of the powered device is ongoing. • `4P Port that deliver only 2 Pair non IEEE`—Signature failure on Alt-B, allowing power only on Alt-A (Non IEEE or Legacy PD). • `4P Port delivering 2P non IEEE`—Non-IEEE PD was detected using 4P matrix in BT mode and power as 2Pair . • `4P Port delivering 4P non IEEE`—Non-IEEE PD was detected using 4P matrix in BT mode and power as 4Pair. • `4P Port delivering 2P IEEE SSPD`—802.3BT- SSPD was detected using 4P matrix and operate as 2P if requested class =< 4. • `4P Port delivering 4P IEEE SSPD`—802.3BT- SSPD was detected using 4P matrix and operate as 4P if requested class > 4. • `4P Port delivering 2P IEEE DSPD in 1st phase`—802.3BT- DSPD was detected using 4P matrix and operate as 2P due to 4pair candidate validation in two cycles. • `4P Port delivering 2P IEEE DSPD`—802.3BT- DSPD was detected using 4P matrix and operate as 2P. • `4P Port delivering 4P IEEE DSPD`—802.3BT- DSPD was detected using 4P matrix and operate as 4P. • `Force Power BT 4P`—Port matrix 4P and delivers power on both pair sets due to force power command.

Table 1001: show poe interface Output Fields *(Continued)*

Field Name (All Interfaces Output)	Field Name (Single Interface Output)	Field Description
		• `Force Power BT Error`—Force power command was set, one of the port pair sets stop delivering power, from at least one reason out of various reasons (System related, Device related, port related or Pair set related). • `Connection Check error`—This error will be reported only in 4 pair port when invalid connection check signature was detected. In such case detection fail counter will be incremented. • `Open`—Port is not connected (Equivalent to Detection in Progress).
	`FourPair status`	Status of four-pair PoE (PoE-4P), a Juniper Networks extension to the IEEE 802.3at standard, which can be enabled for high or ultra-high power delivery: • `Enabled`—High or ultra-high power mode is enabled. • `Disabled`—High or ultra-high power mode is disabled. See Enabling PoE on EX Series Switches (CLI Procedure) for information on how to enable four-pair PoE.
`Pair/Mode status`		Shows the mode of power delivery configured on the interface. • `4P/AT`—Interface is configured for high power mode. • `4P/POH`—Interface is configured for ultra-high power mode. • `DS/BT`—Interface is configured for dual-signature powered devices. • `SS/BT`—Interface is configured for single-signature powered devices.

Table 1001: show poe interface Output Fields *(Continued)*

Field Name (All Interfaces Output)	Field Name (Single Interface Output)	Field Description
`Max power`	`Power limit on the interface`	Maximum power that can be provided by the interface. This is determined by the class of the powered device. For dual-signature devices, each pair set has its own class. The maximum power will be the total of the maximum power for the two classes: `max power = class x + class y` for dual signature `x/y`. For example, a dual-signature device with a class value displayed as 5/5 has a maximum power of 90W, because class 5 has a maximum power of 45W. An `(L)` next to the value indicates that the value on the port was negotiated by LLDP.
Priority	**Priority**	Interface power priority: either **High** or **Low**. An **(L)** next to the value indicates that the value on the port was negotiated by LLDP.
Power consumption	**Power consumed**	Amount of power being used by the interface at the time the command is executed.
Class	**Class of power device**	IEEE PoE class of the powered device. `Class 0` is the default class and is used when the class of the powered device is unknown. If no powered device is connected, this field contains `not applicable`. PoE-bt supports power devices with dual signatures. For dual-signature devices, the output value contains both the class values, e.g. `5/5`. For single-signature devices, the output value is formatted as `5/-`.
	PoE Mode	IEEE PoE standard supported by the interface—either `802.3af`, or `802.3at`, `ultra-poe`, or `802.3bt`.

Sample Output

show poe interface

```
user@switch> show poe interface
Interface Admin        Oper        Max       Priority Power             Class
          status       status      power               consumption
 ge-0/0/0 Enabled      ON          15.4W     Low      7.9W                0
 ge-0/0/1 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/2 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/3 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/4 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/5 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/6 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/7 Enabled      ON          15.4W     Low      3.2W                2
```

show poe interface (with LLDP negotiation)

```
user@switch> show poe interface
Interface    Admin       Oper    Max        Priority       Power          Class
             status      status  power                      consumption
 ge-0/0/0    Enabled      ON     17.5W(L)   Low(L)         16.2W           4
 ge-0/0/1    Enabled      ON     17.5W(L)   Low(L)         16.0W           4
 ge-0/0/2    Enabled      ON     17.5W(L)   High(L)        16.0W           4
 ge-0/0/3    Enabled      ON     17.5W(L)   Low(L)         16.0W           4
 ge-0/0/4    Enabled      ON     10.1W(L)   Low(L)         10.0W           3
 ge-0/0/5    Enabled      ON     3.5W(L)    High(L)        3.0W            2
(L) LLDP-negotiated value on the port.
```

show poe interface (specific interface)

```
user@switch> show poe interface ge-0/0/3
PoE interface status:
PoE interface                :  ge-0/0/3
Administrative status        : Enabled
Operational status           :  ON
Operational status detail    :  IEEE PD Detected
Power limit on the interface : 7.0W
```

```
Priority                      : Low
Power consumed                : 5.3W
Class of power device         :   2
PoE Mode                      : 802.3af
```

show poe interface (specific FPC slot)

```
user@switch> show poe interface fpc-slot 3
Interface Admin        Oper        Max       Priority Power             Class
          status       status      power               consumption
 ge-3/0/0 Enabled       ON         30.0W     Low      20.3W                 4
 ge-3/0/1 Enabled       ON         30.0W     Low      17.8W                 4
 ge-3/0/2 Enabled       ON         30.0W     High     16.3W                 4
 ge-3/0/3 Enabled       ON         30.0W     High     16.2W                 4
 ge-3/0/4 Enabled       ON         30.0W     Low      25.9W                 4
 ge-3/0/5 Enabled       ON         30.0W     Low      10.1W                 4
 ge-3/0/6 Enabled       ON         30.0W     Low      16.2W                 4
 ge-3/0/7 Enabled       ON         30.0W     Low      6.4W                  4
 ge-3/0/8 Enabled       ON         30.0W     Low      5.2W                  4
 ge-3/0/9 Enabled       ON         30.0W     Low      5.2W                  4
ge-3/0/10 Enabled       ON         30.0W     Low      21.5W                 4
ge-3/0/11 Enabled       ON         30.0W     Low      21.7W                 4
ge-3/0/12 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/13 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/14 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/15 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/16 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/17 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/18 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/19 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/20 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/21 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/22 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/23 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/24 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/25 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/26 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/27 Enabled       ON         15.4W     Low      9.4W                  0
ge-3/0/28 Enabled       ON         15.4W     Low      7.0W                  0
ge-3/0/29 Enabled       ON         15.4W     Low      2.2W                  1
ge-3/0/30 Enabled       ON         15.4W     Low      2.2W                  1
```

```
ge-3/0/31 Enabled       ON         15.4W     Low      2.2W                  1
ge-3/0/32 Enabled       ON         15.4W     Low      2.0W                  1
ge-3/0/33 Enabled       ON         15.4W     Low      2.0W                  1
ge-3/0/34 Enabled       ON         15.4W     Low      2.2W                  1
ge-3/0/35 Enabled       ON         15.4W     Low      2.2W                  1
ge-3/0/36 Enabled       ON         15.4W     Low      2.2W                  1
ge-3/0/37 Enabled       ON         15.4W     Low      2.2W                  1
ge-3/0/38 Enabled       ON         15.4W     Low      2.2W                  1
ge-3/0/39 Enabled       ON         15.4W     Low      2.2W                  1
```

show poe interface (specific interface on ACX2000 Universal Metro Routers)

```
user@host> show poe interface ge-0/1/7
PoE interface status:
PoE interface                : ge-0/1/7
Administrative status        : Enabled
Operational status           : Powered-up
Power limit on the interface : 9.0 W
Priority                     : Low
Power consumed               : 14.2 W
Class of power device        :  4
```

show poe interface (PoE-bt mode)

```
user@switch> show poe interface
Interface   Admin      Oper    Pair/Mode  Max     Priority  Power         Class
            status     status  status     power             consumption
 ge-0/0/0   Enabled    ON      SS/BT      15.4W   Low       12.0W (L)      3/-
 ge-0/0/1   Enabled    ON      4P/BT      60.0W   Low       55.0W          4/-
 ge-0/0/2   Enabled    ON      4P/POH     90.0W   Low       81.0W          4/-
 ge-0/0/3   Enabled    ON      SS/BT      90.0W   High      80.0W          8/-
 ge-0/0/4   Enabled    ON      SS/BT      75.0W   Low       65.0W(L)       7/-
 ge-0/0/5   Enabled    ON      SS/BT      30.0W   Low       27.0W          4/-
 ge-0/0/6   Enabled    ON      SS/BT      15.4W   Low       13.0W          3/-
 ge-0/0/7   Enabled    ON      SS/BT      60.0W   Low       49.0W          6/-
 ge-0/0/8   Enabled    ON      DS/BT      90.0W   Low       78.0W          5/5
 ge-0/0/9   Enabled    ON      DS/BT      75.0W   Low       68.0W(L)       5/4
 ge-0/0/10  Enabled    ON      DS/BT      60.0W   Low       55.0W          4/4
```

show poe interface (PoE-bt mode: specific interface)

```
user@switch> show poe interface ge-0/0/3
PoE interface status:
PoE interface                :  ge-0/0/3
Administrative status        : Enabled
Operational status           :  ON
Four-pair status             :  NA
Power limit on the interface : 90.0W
Priority                     : Low
Power consumed               : 78.0W
Class of power device        :   5/5
PoE Mode                     : 802.3bt
```

show poe interface (PoE bounce disabled)

```
user@switch> show poe interface mge-0/0/29
PoE interface status:
PoE interface                : mge-0/0/29
Administrative status        : Disabled
Operational status           : Disabled
Operational status detail    : Admin off
FourPair status              : Disabled
Power limit on the interface : 0.0W
Priority                     : Low
Power consumed               : 0.0W
Class of power device        : not-applicable
PoE Mode                     :   802.3bt
```

show poe interface (PoE bounce enabled)

```
user@switch> show poe interface mge-0/0/29
PoE interface status:
PoE interface                : mge-0/0/29
Administrative status        : Enabled
Operational status           :  OFF
Operational status detail    : Disabled
FourPair status              : Disabled
Power limit on the interface : 15.4W
```

```
Priority                      : Low
Power consumed                : 0.0W
Class of power device         : not-applicable
PoE Mode                      :   802.3bt
```

```
user@switch> show poe interface mge-0/0/29
PoE interface status:
PoE interface                 : mge-0/0/29
Administrative status         : Enabled
Operational status            :   ON
Operational status detail     : 4P Port delivering 2P IEEE SSPD
FourPair status               : Disabled
Power limit on the interface : 19.4W (L)
Priority                      : Low
Power consumed                : 7.8W
Class of power device         :      4/-
PoE Mode                      :   802.3bt
  (L) LLDP-negotiated value on the port.
```

Release Information

Command introduced in Junos OS Release 9.0.

Command introduced in Junos OS Release 12.2 for ACX2000 routers.

RELATED DOCUMENTATION

show poe controller | 23171

Enabling PoE on EX Series Switches (CLI Procedure)

show poe notification-control

IN THIS SECTION

Syntax | 23188

- Description | 23188
- Required Privilege Level | 23188
- Output Fields | 23188
- Sample Output | 23189
- Release Information | 23189

Syntax

```
show poe notification-control
```

Description

Display the state of the PoE **notification-control** option, which enables or disables PoE SNMP traps.

Required Privilege Level

view

Output Fields

Table 1002 on page 23188 lists the output fields for the `show poe notification-control` command. Output fields are listed in the approximate order in which they appear.

Table 1002: show poe notification-control Output Fields

Field Name	Field Description
FPC slot	FPC slot number: • 0 for a standalone switch • Member ID for a Virtual Chassis

Table 1002: show poe notification-control Output Fields *(Continued)*

Field Name	Field Description
Notification-control-status	Status of notification control: • **ON**—PoE traps are enabled. • **OFF**—PoE traps are disabled.

Sample Output

show poe notification-control

```
user@switch> show poe notification-control
FPC slot             Notification-control-status
    0                    OFF
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show poe controller | 23171

show poe interface | 23177

Verifying PoE Configuration and Status (CLI Procedure)

show poe telemetries

IN THIS SECTION

Syntax | 23190

- Description | 23190
- Options | 23190
- Required Privilege Level | 23190
- Output Fields | 23190
- Sample Output | 23191
- Release Information | 23192

Syntax

```
show poe telemetries
<interface interface-name count number>
<count number interface interface-name>
```

Description

Display a history of power consumption on the specified interface. Telemetries must be enabled on the interface before you can display a history of power consumption.

Options

- `Interface` *`interface-name`*—Display telemetries for the specified PoE interface.
- `count` *`number`*—Display the specified number of telemetries records for the specified PoE interface.

Required Privilege Level

View

Output Fields

Table 1003 on page 23191 lists the output fields for the `show poe telemetries interface` command. Output fields are listed in the approximate order in which they appear.

Table 1003: show poe telemetries interface Output Fields

Field name	Field Description
S1 No	Number of the record for the specified port. The last record is the most is the most recent.
Timestamp	Time that the power-consumption data was gathered.
Power	Amount of power provided by the specified port at the time the data was gathered.
Voltage	Voltage on the specified port at the time the data was gathered.

Sample Output

show poe telemetries interface

`user@host>`**show poe telemetries interface ge-0/0/1 count 8**

```
Sl No    Timestamp                Power    Voltage
   1     Fri Jan 04 11:41:15 2009 6.6 W    47.2 V
   2     Fri Jan 04 11:40:15 2009 6.6 W    47.2 V
   3     Fri Jan 04 11:39:15 2009 6.6 W    47.2 V
   4     Fri Jan 04 11:38:15 2009 6.6 W    47.2 V
   5     Fri Jan 04 11:37:15 2009 6.6 W    47.2 V
   6     Fri Jan 04 11:36:15 2009 6.6 W    47.2 V
   7     Fri Jan 04 11:35:15 2009 6.6 W    47.2 V
   8     Fri Jan 04 11:34:15 2009 6.6 W    47.2 V
```

`user@host>`**show poe telemetries count 5 interface ge-0/0/1**

```
Sl No    Timestamp                Power    Voltage
   1     Fri Jan 04 11:47:15 2009 6.6 W    47.2 V
   2     Fri Jan 04 11:38:15 2009 6.6 W    47.2 V
   3     Fri Jan 04 11:29:15 2009 6.6 W    47.2 V
   4     Fri Jan 04 11:11:15 2009 6.6 W    47.2 V
```

```
    5     Fri Jan 04 11:10:15 2009 6.6 W    47.2 V
```

Release Information

Command modified in Junos OS Release 12.3X48-D10.

show policer (EX)

IN THIS SECTION

- Syntax | 23192
- Description | 23192
- Options | 23192
- Required Privilege Level | 23193
- Output Fields | 23193
- Sample Output | 23193
- Release Information | 23194

Syntax

```
show policer
<policer-name>
```

Description

Display statistics about configured policers.

Options

none Display the count of policed packets for all configured policers in the system.

policer-name (Optional) Display the count of policed packets for the specified policer.

Required Privilege Level

view

Output Fields

Table 1004 on page 23193 lists the output fields for the `show policer` command. Output fields are listed in the approximate order in which they appear.

Table 1004: show policer Output Fields

Field Name	Field Description	Level of Output
Filter	Name of filter that is configured with the `filter` statement at the `[edit firewall]` hierarchy level.	All levels
Policers	Display policer information: • Filter—Name of filter that specifies the policer action. • Name—Name of policer. • Packets—Number of packets that matched the filter term where the policer action is specified. This is the number of packets that exceed the rate limits that the policer specifies.	All levels

Sample Output

show policer

```
user@host> show policer
Filter: egress-vlan-filter
Filter: ingress-port-filter
Policers:
Name                                                Packets
icmp-connection-policer                                   0
```

```
tcp-connection-policer                                  0
Filter: ingress-vlan-rogue-block
```

show policer (policer-name)

```
user@host> show policer tcp-connection-policer
Filter: ingress-port-filter
Policers:
Name                                                Packets
tcp-connection-policer                                    0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- *Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches*
- *Verifying That Firewall Filters Are Operational*
- *Verifying That Policers Are Operational*
- *Firewall Filters for EX Series Switches Overview*
- *Understanding the Use of Policers in Firewall Filters*

show policer

IN THIS SECTION

- Syntax | 23195
- Description | 23195
- Options | 23195
- Required Privilege Level | 23195
- Output Fields | 23196

Sample Output | 23197
Release Information | 23199

Syntax

```
show policer
<detail>
<policer-name>
```

Description

Display the number of policed packets for a given policer or an aggregate policer. An aggregate policer is an aggregate of different policers on the same logical interface.

Options

detail (Optional) Display additional statistics and counters. Requires the "enhanced-policer" on page 6403 statement to be enabled at the **[edit chassis]** hierarchy level.

> **NOTE**: `show policer detail` is not available for the following counters: three color policers, hierarchical policers, prefix-specific actions, LSP policers and other nexthop policers, and fast-update-filters on devices running MPC1 or MPC2 line cards.

policer-name (Optional) Display the number of policed packets for the specified policer.

<null> Displays the number of policed packets for all configured policers.

Required Privilege Level

view

Output Fields

Table 1005 on page 23196 lists the output fields for the `show policer` command. Output fields are listed in the approximate order in which they appear.

Table 1005: show policer Output Fields

Field Name	Field Description
`Name`	Name of the policer. Policier detail also includes the following statistics: **`OOS`** Packet statistics for packets that are marked out-of-specification by the policer. Changes to all packets that have out-of-specification actions, such as discard, color marking, or forwarding-class, are included in this counter. **`Offered`** Packet statistics for traffic subjected to policing. **`Transmitted`** Packet statistics for traffic that is not discarded by the policer. When the policer action is discard, the statistics are the same as the within-specification statistics; when the policer action is non-discard (loss-priority or forwarding-class), the statistics are included in this counter.
`Bytes`	• (For two-color policers on MX Series routers, and for hierarchical policers on MS-DPC, MIC, and MPC interfaces on MX Series routers)—Total number of bytes policed by the specified policer. For other combinations of policer type, device, and line card type, this field is blank. • (T Series and M10i)—Not applicable. The Bytes information is not displayed.
`Packets`	Total number of packets policed by the specified policer.

Sample Output

show policer (Junos OS Evolved)

In Junos OS Evolved, the output of this command doesn't display the default Address Resolution Protocol (ARP) policer because it isn't needed. Distributed denial of service (DDoS) protection replaces the functionality of the default ARP policer.

```
user@host> show policer
Policers:
Name                                                Bytes              Packets
__default_arp_mgmt0__                                   0                    0
__default_arp_mgmt1__                                   0                    0
__default_icmpv4_mgmt__                                 0                    0
__default_icmpv6_mgmt__                                 0                    0
__default_mcast_mgmt0__                                 0                    0
__default_mcast_mgmt1__                                 0                    0
```

show policer (MX Series)

```
user@host> show policer
Policers:
Name                                                Bytes              Packets
__default_arp_policer__                            314520                 5242
pol-2M-ge-1/2/0.1-inet-i                         10372300               103723
pol-2M-ge-1/2/0.1-inet6-i                         7727800                77278
pol-2M-ge-1/2/0.1-mpls-i                          7070336                67984
pol-2M-ge-1/2/0.1001-vpls-i                      65153700               651537
pol-2M-ge-1/2/0.2001-vpls-i                      65180900               651809
pol-2M-ge-1/2/0.3001-ccc-i                       62202144               647939
```

show policer (non MX Series Router)

```
user@host> show policer
Policers:

Name                                                Bytes              Packets
__default_arp_policer__                                NA                 5242
pol-2M-ge-1/2/0.1-inet-i                               NA               103723
```

```
pol-2M-ge-1/2/0.1-inet6-i                          NA                77278
pol-2M-ge-1/2/0.1-mpls-i                           NA                67984
pol-2M-ge-1/2/0.1001-vpls-i                        NA               651537
pol-2M-ge-1/2/0.2001-vpls-i                        NA               651809
pol-2M-ge-1/2/0.3001-ccc-i                         NA               647939
```

show policer (Aggregate Policer, non MX Series Router)

```
user@host> show policer
Policers:
Name                                            Bytes              Packets
__default_arp_policer__                            NA                    0
P1-ae0.0-log_int-o                                 NA                    0
P2-ge-7/0/2.0-inet-o                               NA                    0
P2-ge-7/0/2.0-inet6-o                              NA                    0
__policer_tmpl__-term                              NA                    0
__policer_tmpl__-fc0                               NA                    0
__policer_tmpl__-fc0                               NA                    0
__policer_tmpl__-fc1                               NA                    0
__policer_tmpl__-fc0                               NA                    0
__policer_tmpl__-fc1                               NA                    0
__policer_tmpl__-fc2                               NA                    0
__policer_tmpl__-fc0                               NA                    0
__policer_tmpl__-fc1                               NA                    0
__policer_tmpl__-fc2                               NA                    0
__policer_tmpl__-fc3                               NA                    0
```

show policer detail (MX Series Router running MPC 1 or MPC 2 line cards)

```
user@host> show policer detail
Policers:
Name                                            Bytes              Packets
__default_arp_policer__
    OOS                                            0                    0
    Offered                                        0                  496
    Transmitted                                    0                  496
P1-xe-1/0/0.0-inet-i
    OOS                                            0                11329
```

```
    Offered                                          0              111188
    Transmitted                                      0               99859
```

Release Information

Command introduced before Junos OS Release 7.4.

The command `show policer detail` was introduced in Junos OS Release 12.3.

show policy

IN THIS SECTION

- Syntax | 23199
- Syntax (EX Series Switches) | 23200
- Description | 23200
- Options | 23200
- Required Privilege Level | 23200
- Output Fields | 23200
- Sample Output | 23201
- Release Information | 23202

Syntax

```
show policy
<logical-system (all | logical-system-name)>
<policy-name>
<statistics >
```

Syntax (EX Series Switches)

```
show policy
<policy-name>
```

Description

Display information about configured routing policies.

Options

none	List the names of all configured routing policies.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
policy-name	(Optional) Show the contents of the specified policy.
statistics	(Optional) Use in conjunction with the `test policy` command to show the length of time (in microseconds) required to evaluate a given policy and the number of times it has been executed. This information can be used, for example, to help structure a policy so it is evaluated efficiently. Timers shown are per route; times are not cumulative. Statistics are incremented even when the router is learning (and thus evaluating) routes from peering routers.

Required Privilege Level

view

Output Fields

Table 1006 on page 23201 lists the output fields for the `show policy` command. Output fields are listed in the approximate order in which they appear.

Table 1006: show policy Output Fields

Field Name	Field Description
policy-name	Name of the policy listed.
term	Name of the user-defined policy term. The term name `unnamed` is used for policy elements that occur outside of user defined terms
from	Match condition for the policy.
then	Action for the policy.

Sample Output

show policy

```
user@host> show policy
Configured policies:
__vrf-export-red-internal__
__vrf-import-red-internal__
red-export
rf-test-policy
multicast-scoping
```

show policy policy-name

```
user@host> show policy vrf-import-red-internal
Policy vrf-import-red-internal:
    from
         203.0.113.0/28  accept
         203.0.113.32/28  accept
    then reject
```

show policy statistics policy-name

```
user@host> show policy statistics iBGP-v4-RR-Import
Policy iBGP-v4-RR-Import:
    [1243328] Term Lab-Infra:
        from [1243328 0]  proto BGP
             [28 0] route filter:
                 10.11.0.0/8 orlonger
                 10.13.0.0/8 orlonger
        then [28 0] accept
    [1243300] Term External:
        from [1243300 1]  proto BGP
         [1243296 0]  community Ext-Com1 [64496:1515 ]
         [1243296 0]  prefix-list-filter Customer-Routes
         [1243296 0]  aspath AS6221
             [1243296 1] route filter:
                 172.16.49.0/12 orlonger
                 172.16.50.0/12 orlonger
                 172.16.51.0/12 orlonger
                 172.16.52.0/12 orlonger
                 172.16.56.0/12 orlonger
                 172.16.60.0/12 orlonger
        then [1243296 2] community + Ext-Com2 [64496:2000 ] [1243296 0] accept
    [4] Term Final:
        then [4 0] reject
```

Release Information

Command introduced before Junos OS Release 7.4.

`statistics` option introduced in Junos OS Release 16.1 for MX Series routers.

RELATED DOCUMENTATION

show policy damping

test policy | **29209**

show policy conditions

IN THIS SECTION

- Syntax | 23203
- Syntax (EX Series Switches) | 23203
- Description | 23203
- Options | 23204
- Required Privilege Level | 23204
- Output Fields | 23204
- Sample Output | 23205
- Release Information | 23205

Syntax

```
show policy conditions
<condition-name>
<detail>
<dynamic>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show policy conditions
<condition-name>
<detail>
<dynamic>
```

Description

Display all the configured conditions as well as the routing tables with which the configuration manager is interacting. If the `detail` keyword is included, the output also displays dependent routes for each condition.

Options

none	Display all configured conditions and associated routing tables.
condition-name	(Optional) Display information about the specified condition only.
detail	(Optional) Display the specified level of output.
dynamic	(Optional) Display information about the conditions in the dynamic database.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1007 on page 23204 lists the output fields for the **show policy conditions** command. Output fields are listed in the approximate order in which they appear.

Table 1007: show policy conditions Output Fields

Field Name	Field Description	Level of Output
Condition	Name of configured condition.	All levels
event	Condition type. If the `if-route-exists` option is configured, the event type is: Existence of a route in a specific routing table.	All levels
Dependent routes	List of routes dependent on the condition, along with the latest generation number.	`detail`
Condition tables	List of routing tables associated with the condition, along with the latest generation number and number of dependencies.	All levels

Table 1007: show policy conditions Output Fields *(Continued)*

Field Name	Field Description	Level of Output
If-route-exists conditions	List of conditions configured to look for a route in the specified table.	All levels

Sample Output

show policy conditions detail

```
user@host> show policy conditions detail
Configured conditions:
Condition cond1, event: Existence of a route in a specific routing table
Dependent routes:
 172.16.4.4/32, generation 3
 6.6.6.6/32, generation 3
 10.10.10.10/32, generation 3

Condition cond2, event: Existence of a route in a specific routing table
Dependent routes:
None

Condition tables:
Table inet.0, generation 4, dependencies 3, If-route-exists conditions: cond1 (static) cond2
(static)
```

Release Information

Command introduced in Junos OS Release 9.0.

show policy damping

IN THIS SECTION

- Syntax | 23206
- Syntax (EX Series Switch and QFX Series) | 23206
- Description | 23206
- Options | 23206
- Additional Information | 23207
- Required Privilege Level | 23207
- Output Fields | 23207
- Sample Output | 23208
- Release Information | 23208

Syntax

```
show policy damping
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and QFX Series)

```
show policy damping
```

Description

Display information about BGP route flap damping parameters.

Options

none	Display information about BGP route flap damping parameters.

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

In the output from this command, figure-of-merit values correlate with the probability of future instability of a routing device. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time. A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes.

Required Privilege Level

view

Output Fields

Table 1008 on page 23207 describes the output fields for the `show policy damping` command. Output fields are listed in the approximate order in which they appear.

Table 1008: show policy damping Output Fields

Field Name	Field Description
`Halflife`	Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes.
`Reuse merit`	Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.
`Suppress/cutoff merit`	Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols.

Table 1008: show policy damping Output Fields *(Continued)*

Field Name	Field Description
`Maximum suppress time`	Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability.
`Computed values`	• `Merit ceiling`—Maximum merit that a flapping route can collect. • `Maximum decay`—Maximum decay half-life, in minutes.

Sample Output

show policy damping

```
user@host> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "standard-damping":
  Halflife: 10 minutes
  Reuse merit: 4000 Suppress/cutoff merit: 8000
  Maximum suppress time: 30 minutes
  Computed values:
    Merit ceiling: 32120
    Maximum decay: 12453
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Routing Policies, Firewall Filters, and Traffic Policers User Guide
- clear bgp damping
- show route damping | 23564

show ppp address-pool

IN THIS SECTION

- Syntax | 23209
- Description | 23209
- Options | 23209
- Required Privilege Level | 23210
- Output Fields | 23210
- Sample Output | 23210
- Release Information | 23211

Syntax

```
show ppp address-pool pool-name
<detail>
```

Description

Display PPP address pool information.

Options

pool-name	Address pool name.
detail	(Optional) Display detailed address pool information.

Required Privilege Level

view

Output Fields

Table 1009 on page 23210 lists the output fields for the `show ppp address-pool` command. Output fields are listed in the approximate order in which they appear.

Table 1009: show ppp address-pool Output Fields

Field Name	Field Description	Level of Output
`Address pool`	Trace address pool code.	All levels
`Address range`	Range of sequentially ordered IP addresses contained in the address pool.	`detail`
`Number of assigned addresses`	Fixed IP address that is to be given to remote users when they dial in. This is a host-only IP address (subnet mask is 255.255.255.255) and is only for single connection receiver profiles.	All levels
`Number of addresses configured`	Number of IP addresses that are available for allocation and used by PPP sessions.	All levels
`Assigned addresses`	Addresses assigned to PPP sessions from the address pool.	`detail`

Sample Output

show ppp address-pool

```
user@host> show ppp address-pool
Address pool ppp1
  Address range: 203.0.113.221 - 203.0.113.230
```

```
  Number of assigned addresses: 0
  Number of addresses configured: 10
```

show ppp address-pool detail

```
user@host> show ppp address-pool ppp1 detail
Address pool ppp1
  Address range: 203.0.113.221 - 203.0.113.230
  Number of assigned addresses: 2
  Number of addresses configured: 10
  Assigned addresses:
      203.0.113.221
      203.0.113.222
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Verifying and Managing PPP Configuration for Subscriber Management

show ppp interface

IN THIS SECTION

- Syntax | 23212
- Description | 23212
- Options | 23212
- Required Privilege Level | 23212
- Output Fields | 23212
- Sample Output | 23226
- Release Information | 23229

Syntax

```
show ppp interface interface-name
<extensive | terse>
```

Description

Display information about PPP interfaces.

Options

interface-name Name of a logical interface.

Starting in Junos OS Release 17.3, the * (asterisk) wildcard character is supported for the interface name for debugging purpose. With this support, you can match any string of characters in that position in the interface name. For example, so* matches all SONET/SDH interfaces.

extensive | terse (Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1010 on page 23212 lists the output fields for the `show ppp interface` command. Output fields are listed in the approximate order in which they appear.

Table 1010: show ppp interface Output Fields

Field Name	Field Description	Level of Output
`Session`	Name of the logical interface on which the session is running.	All levels
`Type`	Session type: PPP.	All levels

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Phase`	PPP process phase: `Authenticate`, `Pending`, `Establish`, `LCP`, `Network`, `Disabled`, and `Tunneled`.	All levels
`Session flags`	Special conditions present in the session: `Bundled`, `TCC`, `No-keepalives`, `Looped`, `Monitored`, and `NCP-only`.	All levels
`protocol` `State`	Protocol state information. See specific protocol state fields for information.	None specified
`AUTHENTICATION`	Challenge-Handshake Authentication Protocol (CHAP) authentication state information or Password Authentication Protocol (PAP) state information. See the `Authentication` field description for further information.	None specified
`Keepalive settings`	Keepalive settings for the PPP sessions on the L2TP network server (LNS). LNS-based PPP sessions are supported only on service interfaces (si). • `Interval`—Time in seconds between successive keepalive requests. Keepalive aging timeout is calculated as a product of the `interval` and `Down-count` values. If the keepalive aging timeout is greater than 180 seconds, the keepalive packets are handled by the Routing Engine. If the aging timeout is less than or equal to 180 seconds, the packets are handled by the Packet Forwarding Engine. • `Up-count`—The number of keepalive packets a destination must receive to change a link's status from down to up. • `Down-count`—The number of keepalive packets a destination must fail to receive before the network takes down a link.	`extensive`

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Magic-Number validation`	Indicates whether the local peer is configured to ignore mismatches between peer magic numbers when the numbers are validated during PPP keepalive (Echo-Request/Echo-Reply) exchanges. • Enable–Mismatch detection sends failed Echo-Reply packets to the Routing Engine. If a valid magic number is not received within the configurable keepalive interval, PPP treats this as a keepalive failure and tears down the PPP sessions. • Disable–The Packet Forwarding Engine does not perform a validation check for magic numbers received from remote peers. A mismatch cannot be detected, so receipt of its own magic number or an unexpected value does not trigger notification to the Routing Engine.	`extensive`
`RE Keepalive statistics`	Keepalive statistics for the packets handled by the Routing Engine. • `LCP echo req Tx`—LCP echo requests sent from the Routing Engine. • `LCP echo req Rx`—LCP echo requests received at the Routing Engine. • `LCP echo rep Tx`—LCP echo responses sent from the Routing Engine. • `LCP echo rep Rx`—LCP echo responses received at the Routing Engine. • `LCP echo req timeout`—Number of keepalive packets where the keepalive aging timer has expired. • `LCP Rx echo req Magic Num Failures`—LCP echo requests where the magic numbers shared between the PPP peers during LCP negotiation did not match. • `LCP Rx echo rep Magic Num Failures`—LCP echo responses where the magic numbers shared between the PPP peers during LCP negotiation did not match.	`extensive`

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output

LCP	LCP information:
	• State—LCP protocol state (all platforms except M120 and M320 routers):
	  • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received.
	  • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • Closed—Link is not available for traffic.
	  • Opened—Link is administratively available for traffic.
	  • Req-sent—An attempt has been made to configure the connection.
	• State—LCP protocol state (M120 and M320 routers):
	  • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received.
	  • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • Closed—Link is available (up), but no Open has occurred.
	  • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
	  • Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.
		extensive

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Stopping`—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • `Last started`—LCP state start time. • `Last completed`—LCP state completion time. • `Last updated`—Reports the timestamp of the last successful connection update exchange. 1. When LCP negotiation completes, this field has the same value as the `Last completed` field. 2. The field then reports the timestamp of any subsequent successful exchange of Connection-Update-Request and Connection-Update-Ack messages with the peer (such as a home gateway). This field is displayed only when the Connection-Status-Message option is successfully negotiated.	

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output

	• `Negotiated options:`
	  • `ACFC`—Address and-Control Field Compression. A configuration option that provides a method to negotiate the compression of the Data Link Layer Address and Control fields.
	  • `Asynchronous map`—Asynchronous control character map. A configuration option used on asynchronous links such as telephone lines to identify control characters that must be replaced by a two-character sequence to prevent them from being interpreted by equipment used to establish the link.
	  • `Authentication protocol`—Protocol used for authentication. This option provides a method to negotiate the use of a specific protocol for authentication. It requires a peer to authenticate itself before allowing network-layer protocol packets to be exchanged. By default, authentication is not required.
	  • `Authentication algorithm`—Type of authentication algorithm. The Message Digest algorithm (MD5) is the only algorithm supported.
	  • `Connection Update Requests`—Number of connection update requests sent by PPP to the remote peer (such as a home gateway). This value does not include retries.
	  This field is displayed even when negotiation fails for the Connection-Status-Message option. This enables you to confirm that an update request was sent. The absence of the Juniper Connection Status Message field indicates the peer does not support the updates.
	  • `Endpoint discriminator class`—For multilink PPP (MLPPP), a configuration option that identifies the system transmitting the packet. This option advises a system that the peer on this link could be the same as the peer on another existing link.
	  • `Juniper Connection Status Message`—The content of the Connection-Status-Message VSA (26-4874–218) most recently received from RADIUS.	

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	This field is displayed only when the Connection-Status-Message option is successfully negotiated. • `Magic number`—A configuration option that provides a method to detect looped-back links and other data-link layer anomalies. By default, the magic number is not negotiated. • `MRU`—Maximum receive unit. A configuration option that may be sent to inform the peer that the implementation can receive larger packets, or to request that the peer send smaller packets. The default value is 1500 octets. • `MRRU`—For multilink PPP, the maximum receive reconstructed unit. A configuration option that specifies the maximum number of octets in the Information fields of reassembled packets. • `Multilink header suspendable classes`—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number, with the maximum number of suspendable classes given. • `Multilink header format classes`—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number. • `PFC`—Protocol-Field-Compression. A configuration option that provides a method to negotiate the compression of the PPP Protocol field. • `short sequence`—For MLPPP, an option that advises the peer that the implementation wishes to receive fragments with short, 12-bit sequence numbers.	

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Authentication`	CHAP or PAP authentication state information. For CHAP authentication: • `Chap-ans-rcvd`—Packet was sent from the peer, indicating that the peer received the `Chap-resp-sent` packet. • `Chap-ans-sent`—Packet was sent from the authenticator, indicating that the authenticator received the peer's `Chap-resp-rcvd` packet. • `Chap-chal-rcvd`—Challenge packet has been received by the peer. • `Chap-chal-sent`—Challenge packet has been sent by the authenticator to begin the CHAP protocol or has been transmitted at any time during the Network-Layer Protocol (NCP) phase to ensure that the connection has not been altered. • `Chap-resp-rcvd`—CHAP response packet has been received by the authenticator. • `Chap-resp-sent`—CHAP response packet has been sent to the authenticator. • `Closed`—Link is not available for authentication. • `Failure`—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. • `Success`—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful. For PAP authentication: • `Pap-resp-sent`—PAP response sent to peer (ACK/NACK)t. • `Pap-req-rcvd`—PAP request packet received from peer. • `Pap-resp-rcvd`—PAP response received from the peer (ACK/NACK). • `Pap-req-sent`—PAP request packet sent to the peer. • `Closed`—Link is not available for authentication.	None specified

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Failure`—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. • `Success`—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful.	

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`IPCP`	Internet Protocol Control Protocol (IPCP) information.
	• `State`—(All platforms except M120 and M320 routers) One of the following values:
	  • `Ack-rcvd`—A Configure-Request has been sent and a Configure-Ack has been received.
	  • `Ack-sent`—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • `Closed`—Link is not available for traffic.
	  • `Opened`—Link is administratively available for traffic.
	  • `Req-sent`—An attempt has been made to configure the connection.
	• `State`—(M120 and M320 routers) One of the following values:
	  • `Ack-rcvd`—A Configure-Request has been sent and a Configure-Ack has been received.
	  • `Ack-sent`—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • `Closed`—Link is available (up), but no Open has occurred.
	  • `Closing`—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • `Opened`—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • `Req-sent`—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • `Starting`—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
		`extensive`

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output

	• `Stopped`—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.
	• `Stopping`—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	• `Last started`—IPCP state start time.
	• `Last completed`—IPCP state authentication completion time.
	• `Negotiated options:`
	  • `compression protocol`—Negotiate the use of a specific compression protocol. By default, compression is not enabled.
	  • `local address`—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address.
	  • `primary DNS server`—Negotiate with the remote peer to select the address of the primary DNS server to be used on the local end of the link.
	  • `primary WINS server`—Negotiate with the remote peer to select the address of the primary WINS server to be used on the local end of the link.
	  • `remote address`—IP address of the remote end of the link in dotted quad notation.
	  • `secondary DNS server`—Negotiate with the remote peer to select the address of the secondary DNS server to be used on the local end of the link.
	  • `secondary WINS server`—Negotiate with the remote peer to select the address of the secondary WINS server to be used on the local end of the link.
	• `Negotiation mode`—PPP Network Control Protocol (NCP) negotiation mode configured for IPCP: `Active` or `Passive`	

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output

IPV6CP	Internet Protocol version 6 Control Protocol (IPv6CP) information.
	• State—(All platforms except M120 and M320 routers) One of the following values:
	  • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received.
	  • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • Closed—Link is not available for traffic.
	  • Opened—Link is administratively available for traffic.
	  • Req-sent—An attempt has been made to configure the connection.
	• State—(M120 and M320 routers) One of the following values:
	  • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received.
	  • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • Closed—Link is available (up), but no Open has occurred.
	  • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
		extensive

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Stopped`—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. • `Stopping`—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • `Last started`—IPV6CP state start time. • `Last completed`—IPV6CP state authentication completion time. • `Negotiated options`: • `local interface identifier`—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address. • `remote interface identifier`—IP address of the remote end of the link in dotted quad notation. • `Negotiation mode`—PPP Network Control Protocol (NCP) negotiation mode configured for IPv6CP: `Active` or `Passive`	
`OSINLCP State`	OSI Network Layer Control Protocol (OSINLCP) protocol state information (all platforms except M120 and M320 routers): • `State`: • `Ack-rcvd`—Configure-Request has been sent and Configure-Ack has been received. • `Ack-sent`—Configure-Request and Configure-Ack have both been sent, but Configure-Ack has not yet been received. • `Closed`—Link is not available for traffic. • `Opened`—Link is administratively available for traffic. • `Req-sent`—Attempt has been made to configure the connection. • `Last started`—OSINLCP state start time. • `Last completed`—OSINCLP state completion time.	`extensive`

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`TAGCP`	TAGCP information.
	• `State`—(All platforms except M120 and M320 routers) One of the following values:
	  • `Ack-rcvd`—A Configure-Request has been sent and a Configure-Ack has been received.
	  • `Ack-sent`—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • `Closed`—Link is not available for traffic.
	  • `Opened`—Link is administratively available for traffic.
	  • `Req-sent`—An attempt has been made to configure the connection.
	• `State`—(M120 and M320 routers) One of the following values:
	  • `Ack-rcvd`—A Configure-Request has been sent and a Configure-Ack has been received.
	  • `Ack-sent`—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • `Closed`—Link is available (up), but no Open has occurred.
	  • `Closing`—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • `Opened`—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • `Req-sent`—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • `Starting`—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
		`extensive`
		none

Table 1010: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Stopped`—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. • `Stopping`—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • `Last started`—TAGCP state start time. • `Last completed`—TAGCP state authentication completion time.	

Sample Output

show ppp interface

```
user@host> show ppp interface si-1/3/0.0
Session si-1/3/0.0, Type: PPP, Phase: Authenticate
  Session flags: Monitored
    LCP State: Opened
    AUTHENTICATION: CHAP State: Chap-resp-sent, Chap-ans-sent
    IPCP State: Closed, OSINLCP State: Closed
```

show ppp interface extensive (LCP Connection Update Negotiation Successful)

```
user@host> show ppp interface extensive pp0.3221225489
  Session pp0.3221225489, Type: PPP, Phase: Network
   Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
               Magic-Number validation: enable
    LCP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
      Last completed: 2020-02-11 15:06:00 PDT
      Last updated: 2020-02-11 15:06:10 PDT
      Negotiated options:
        Magic number: 906403799, Initial Advertised MRU: 1492, Local MRU: 1492, Peer MRU: 149
        Juniper Connection Status Message: 10m:xxxx
```

```
        Connection-Update-Requests: 1
     Authentication: Off
    IPCP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
      Last completed: 2020-02-11 15:06:00 PDT
      Negotiated options:
        Local address: 198.51.100.30, Remote address: 203.0.113.9
      Negotiation mode: Passive
    IPV6CP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
      Last completed: 2020-02-11 15:06:00 PDT
      Negotiated options:
        Local interface identifier: 2001:db8::fc73:cba, Remote interface identifier: 2001:db8::3a
      Negotiation mode: Passive
```

show ppp interface extensive (LCP Connection Update Negotiation Failed)

```
user@host> show ppp interface extensive pp0.3221225489
  Session pp0.3221225489, Type: PPP, Phase: Network
   Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
               Magic-Number validation: enable
    LCP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
      Last completed: 2020-02-11 15:06:00 PDT
      Negotiated options:
        Magic number: 906403799, Initial Advertised MRU: 1492, Local MRU: 1492, Peer MRU: 149
        Connection-Update-Requests: 1
     Authentication: Off
    IPCP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
      Last completed: 2020-02-11 15:06:00 PDT
      Negotiated options:
        Local address: 198.51.100.30, Remote address: 203.0.113.9
      Negotiation mode: Passive
    IPV6CP
      State: Opened
      Last started: 2020-02-11 15:06:00 PDT
```

```
      Last completed: 2020-02-11 15:06:00 PDT
      Negotiated options:
        Local interface identifier: 2001:db8::fc73:cba, Remote interface identifier: 2001:db8::3a
      Negotiation mode: Passive
```

show ppp interface extensive (Inline Service Interface)

```
user@host> show ppp interface si-0/0/3.0 extensive
Session si-0/0/3.0, Type: PPP, Phase: Network
Keepalive settings: Interval 30 seconds, Up-count 1, Down-count 3
                    Magic-Number validation: disable
RE Keepalive statistics:
        LCP echo req Tx        :  657 (last sent 00:50:10 ago)
        LCP echo req Rx        :    0 (last seen: never)
        LCP echo rep Tx        :    0
        LCP echo rep Rx        :  657
        LCP echo req timout    :    0
        LCP Rx echo req Magic Num Failures :    0
        LCP Rx echo rep Magic Num Failures :    0
LCP
  State: Opened
  Last started: 2007-01-29 10:43:50 PST
  Last completed: 2007-01-29 10:43:50 PST
  Negotiated options:
    Authentication protocol: PAP, Magic number: 2341124815, MRU: 4470
Authentication: PAP
  State: Success
  Last started: 2007-01-29 10:43:50 PST
  Last completed: 2007-01-29 10:43:50 PST
IPCP
  State: Opened
  Last started: 2007-01-29 10:43:50 PST
  Last completed: 2007-01-29 10:43:50 PST
  Negotiated options:
    Local address: 203.0.113.21, Remote address: 203.0.113.22
  Negotiation mode: Active
IPV6CP
  State: Opened
  Last started: 2007-01-29 10:43:50 PST
  Last completed: 2007-01-29 10:43:50 PST
  Negotiated options:
```

```
    Local interface identifier: 2a0:a522:64:d319, Remote interface identifier: 0:0:0:c
  Negotiation mode: Passive
```

show ppp interface terse

```
user@host> show ppp interface si-1/3/0 terse
Session name    Session type   Session phase      Session flags
si-1/3/0.0      PPP            Authenticate       Monitored
```

Release Information

Command introduced in Junos OS Release 7.5.

show ppp statistics

IN THIS SECTION

- Syntax | 23229
- Description | 23230
- Options | 23230
- Required Privilege Level | 23230
- Output Fields | 23230
- Sample Output | 23236
- Release Information | 23238

Syntax

```
show ppp statistics
<detail>
<memory>
<recovery>
```

Description

Display PPP interface statistics information.

Options

detail (Optional) Display the detailed statistics.

memory (Optional) Display PPP process memory statistics.

recovery (Optional) Display recovery state of PPP after a GRES or restart. It is safe to force another GRES or restart only when the recovery state indicates the recovery is done.

> **NOTE**: When you issue this command option during the recovery process, the command may time out or fail silently rather than display output. Recovery is not complete until the command displays `Recovery state: recovery done`.

Required Privilege Level

view

Output Fields

Table 1011 on page 23230 lists the output fields for the `show ppp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1011: show ppp statistics Output Fields

Field Name	Field Description	Level of Output
`Total sessions`	Number of PPP sessions on an interface.	none `detail`
`Sessions in disabled phase`	Number of PPP sessions disabled. Number of sessions where the link is either administratively or physically down. Once the PPP process learns from the kernel that Layer 2 is ready to send and receive traffic, it will do a phase transition from disabled to established. When LCP and NCP transitions through states, links transition to the establish phase when terminate packets are exchanged or some other failure, such as authentication or expiration of a timer occurs.	none `detail`

Table 1011: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Sessions in establish phase`	Number of PPP sessions in establish phase. In order to establish communications over a point-to-point link, each end of the PPP link must first send LCP packets to configure and test the data link.	none `detail`
`Sessions in authenticate phase`	Number of PPP sessions in authenticate phase. Each end of the PPP link must first send LCP packets to configure the data link during the link establishment phase. After the link has been established, PPP provides for an optional authentication phase before proceeding to the Network-Layer Protocol (NLP) phase.	none `detail`
`Sessions in network phase`	Number of PPP sessions in the network phase. After a link has been established and optional facilities have been negotiated as needed by the LCP, PPP must send Network Control Protocol (NCP) packets to choose and configure one or more network-layer protocols, such as IP, IPX, or AppleTalk. Once each of the chosen network-layer protocols has been configured, datagrams from each network-layer protocol can be sent over the link.	none `detail`
`Bundles in pending phase`	Number of unique bundles to which PPP links are referring.	none `detail`

Table 1011: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Type`	Type of structure for which memory is allocated. • `Queued rtsock msgs`—Queued route socket messages. When a PPP process is unable to send a route socket message to the kernel (typically because of congestion of the route socket interface), the message is queued for deferred processing. • `PPP session`—Active PPP session. Stores all the information for a PPP session, such as authentication, sequence number, LCP session, and NCP session information. • `Interface address`—Interface address associated with a PPP connection. Stores the information about the interface address that PPP obtains from the kernel. • `Destination profile`—Stores the destination profile information associated with an interface address. • `ML link settings`—Stores information about an MLPPP link, such as the bundle name and compressed real-time transport protocol (CRTP) settings. • `IPCP blocked address`—When addresses are blocked in an address pool (for example, when the interface address is within the range of an address pool, it will be implicitly blocked), this structure is used to store the address in the pool. • `PPP session trace`—A PPP session trace is allocated for record keeping for each session listed at the [`set protocols ppp monitor-session`] hierarchy level. • `IFL redundancy state`—Stores redundancy state information needed for high availability (HA) operation. • `Protocol family`—Stores the information about the protocol family that PPP obtains from the kernel.	`detail`

Table 1011: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Type (continued)`	• `ML bundle settings`—Multilink bundle settings. Stores the context information for a MLPPP bundle. • `PPP LCP session`—PPP Link Control Protocol session, used for establishing, configuring, and testing the data-link connection. Stores the information for an LCP session, such as negotiated options, current state, and statistics. • `PPP NCP session`—PPP Network Control Protocol (NCP) phase in the PPP link connection process. Stores the information for an NCP session, such as negotiated options, current state, address family, and statistics. • `Physical interface`—Stores the information about the physical interface that PPP obtains from the kernel. • `Access profile`—Stores the information found at the [`edit access profile`] hierarchy level for each profile. • `ML wait entry`—Created when there are MLPPP links joining a bundle. before its addition to the PPP process. Links are saved here, and when the bundle is added, are properly assigned to the bundle. • `Group profile`—Stores information set in the PPP stanza of a group profile, such as the primary and secondary Domain Name System (DNS), primary and secondary NDNS, and address pool name. • `Profile client`—Stores the per-client information of the access profile (information obtained from the [`set access profile` *`name`* `client` *`client-name`*] hierarchy level. • `PPP Auth session`—PPP authentication session. Stores all the session-specific authentication protocol parameters. • `Logical interface`—Stores the information about the logical interface that PPP obtains from the kernel. • `Non-tagged`—Generic catch-all for allocations not of a particular structure type.	`detail`

Table 1011: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Type`	If you specify the `memory` keyword, the following memory statistics are displayed for Ethernet interfaces on M120 and M320 routers. • `authenticate`—Stores information common to all PPP authentication protocols. • `linkInterface`—Stores information about PPP link interfaces. • `pap`—Stores information about PPP PAP authentication protocol. Includes authenticator and authenticate state machines. • `lcp`—PPP Link Control Protocol session. Used for establishing, configuring and testing the data-link connection. Stores information for LCP session, such as negotiated options, state, and statistics. • `chap`—Stores information about PPP CHAP authentication protocol. Includes authenticator and authenticate state machines. • `eapBuffer`—Stores runtime authentication information for EAP. • `eap`—Stores information about PPP EAP authentication protocol. Includes authenticator and authenticate state machines. • `authNone`—Stores information about no PPP authentication. Includes the authenticator state machine. • `networkInterface`—Stores information about NCP portions of PPP protocol. • `ipNcp`—PPP IPCP session information. Used for configuring, negotiating, and establishing IPCP protocol. Stores the current state, and configured and negotiated options. • `ipv6Ncp`—PPP IPv6CP session information. Used for configuring, negotiating, and establishing IPv6CP protocol. Stores the current state, and configured and negotiated options. • `osiNcp`—PPP OSICP session information. Used for configuring, negotiating, and establishing OSICP protocol. Stores the current state, and configured and negotiated options.	`memory`

Table 1011: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `mplsNcp`—PPP MPLSCP session information. Used for configuring, negotiating, and establishing MPLSCP protocol. Stores the current state. • `trace`—Stores information for PPP debugging.	
`Total`	Total memory allocations.	`detail`
`Size`	Size of the structure.	`detail`
`Active`	Number of instances of the structure that are used.	`detail`
`Free`	Number of instances of the structure that are on the free list. Types with a number in the `Free` column are pooled structures, and are typically types that are often used.	`detail`
`Limit`	Maximum number of instances that can be on the free list. Types with a number in the `Limit` column are pooled structures, and are typically types that are often used.	`detail`
`Total size`	Total amount of memory being used by a type of structure (includes active and free instances).	`detail`
`Requests`	Number of allocation requests made by a type of structure.	`detail`
`Failures`	Number of failed allocations.	`detail`
`Recovery state`	State of PPP recovery after a GRES or restart: • recovery done—All sessions have recovered; it is safe to force another GRES or restart. • recovery cleanup pending—Not all PPP sessions have recovered; it is not safe to force another GRES or restart.	none

Table 1011: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Subscriber sessions pending retention`	Number of PPP subscriber sessions that are in the process of being recovered.	none
`Subscriber sessions recovered OK`	Number of PPP subscriber sessions that have recovered after a GRES or restart.	none
`Subscriber sessions recovery failed`	Number of PPP subscriber sessions that have failed to recover after a GRES or restart.	none

Sample Output

show ppp statistics

```
user@host> show ppp statistics
Session statistics from PPP process
  Total sessions: 0
    Sessions in disabled phase    : 0
    Sessions in establish phase   : 0
    Sessions in authenticate phase: 0
    Sessions in network phase     : 0
    Bundles in pending phase      : 0

Session statistics from PPP universal edge process
  Total subscriber sessions: 32
    Subscriber sessions in disabled phase    : 32
    Subscriber sessions in establish phase   : 0
    Subscriber sessions in authenticate phase: 0
    Subscriber sessions in network phase     : 0
```

show ppp statistics detail

```
user@host> show ppp statistics detail
Session statistics from PPP process
  Total sessions: 0
    Sessions in disabled phase    : 0
    Sessions in establish phase   : 0
    Sessions in authenticate phase: 0
    Sessions in network phase     : 0
    Bundles in pending phase      : 0
Type                   Size  Active  Free  Limit  Total size  Requests Failures
Queued rtsock msgs       28      0      0  65535           0         0
PPP session              60      0                          0         0
Interface address        64      0      0  65535           0         0
Destination profile      65      0                          0         0
ML link settings         68      0                          0         0
IPCP blocked address     68      0                          0         0
PPP session trace        76      0                          0         0
IFL redundancy state     76      0                          0         0
Protocol family          84      0      0  65535           0         0
ML bundle settings      108      0                          0         0
PPP LCP session         120      0                          0         0
PPP NCP session         124      0                          0         0
Physical interface      124    170      0  65535       21080       170
Access profile          132      0                          0         0
ML wait entry           144      0      0     20           0         0
Group profile           164      0                          0         0
Profile client          272      0                          0         0
PPP Auth session        356      0                          0         0
Logical interface       524      0      0  65535           0         0
Non-tagged                                                  8         2
Total                                                   21088       172       0

Session statistics from PPP universal edge process
  Total subscriber sessions: 32
    Subscriber sessions in disabled phase    : 32
    Subscriber sessions in establish phase   : 0
    Subscriber sessions in authenticate phase: 0
    Subscriber sessions in network phase     : 0
Type                   Size  Active  Free  Limit  Total size Requests Failures
authenticate            224       1    99  16384         224        0        0
linkInterface           152       1    99  16384         152        0        0
```

```
pap                     256       1    99  16384         256        0        0
lcp                     272       1    99  16384         272        0        0
chap                    284       0     0  16384           0        0        0
eapBuffer              1464       0     0  16384           0        0        0
eap                     276       0     0  16384           0        0        0
authNone
networkInterface        220       1    99  16384         220        0        0
ipNcp                   256       1    99  16384         256        0        0
ipv6Ncp                 204       0     0  16384           0        0        0
osiNcp                  192       0     0  16384           0        0        0
mplsNcp                 188       0     0  16384           0        0        0
trace                  2052       0    16     16           0        0        0
Total                                                    1380        0        0
```

show ppp statistics recovery (Safe to Restart)

```
user@host> show ppp statistics recovery
Recovery statistics from PPP universal edge process
  Recovery state: recovery done
    Subscriber sessions recovered OK      : 32001
    Subscriber sessions recovery failed   : 0
```

show ppp statistics recovery (Unsafe to Restart)

```
user@host> show ppp statistics recovery
Recovery statistics from PPP universal edge process
  Recovery state: recovery cleanup pending
    Subscriber sessions pending retention : 32001
    Subscriber sessions recovered OK      : 0
    Subscriber sessions recovery failed   : 0
```

Release Information

Command introduced in Junos OS Release 7.5.

show ppp summary

IN THIS SECTION

- Syntax | 23239
- Description | 23239
- Options | 23239
- Required Privilege Level | 23239
- Output Fields | 23239
- Sample Output | 23240
- Release Information | 23240

Syntax

```
show ppp summary
```

Description

Display PPP session summary information.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1012 on page 23240 lists the output fields for the `show ppp summary` command. Output fields are listed in the approximate order in which they appear.

Table 1012: show ppp summary Output Fields

Field Name	Field Description
Interface	Interface on which the PPP session is running. An interface type of pp0 indicates an Ethernet interface type on a M120 or M320 router.
Session type	Type of session: PPP or Cisco-HDLC.
Session phase	PPP process phases: Authenticate, Pending, Establish, Network, Disabled.
Session flags	Special conditions present in the session, such as Bundled, TCC, No-keepalives, Looped, Monitored, and NCP-only.

Sample Output

show ppp summary

```
user@host> show ppp summary
Interface       Session type  Session phase      Session flags
at-4/0/0.456    PPP           Network
lsq-0/3/0.0     PPP           Disabled
lsq-1/0/0.0     PPP           Disabled
rlsq0.0         PPP           Network            NCP-only
so-1/0/0.0      PPP           Authenticate
so-1/0/1.0      PPP           Disabled           Looped
so-2/0/0.0      Cisco-HDLC    Establish
so-4/0/0.0      PPP           Establish          Monitored
t1-1/3/0:1.0    PPP           Network            Bundled
t1-1/3/0:2.0    PPP           Network            Bundled
pp0.12          PPP           Network
```

Release Information

Command introduced in Junos OS Release 7.5.

show pppoe interfaces (Security)

IN THIS SECTION

- Syntax | 23241
- Description | 23241
- Options | 23241
- Required Privilege Level | 23242
- Output Fields | 23242
- Sample Output | 23244
- Release Information | 23246

Syntax

```
show pppoe interfaces
<brief | detail | extensive>
<pp0.logical>
```

Description

Display session-specific information about PPPoE interfaces.

Options

none Display interface information for all PPPoE interfaces.

brief | detail (Optional) Display the specified level of output.

extensive (Optional) Display information about the number of packets sent and received and the number of timeouts during a PPPoE session.

pp0.*logical* (Optional) Name of an interface. The logical unit number for static interfaces can be a value from 0 through 16,385. The logical unit number for dynamic interfaces can be a value from 1,073,741,824 through the maximum number of logical interfaces supported on your SRX300, SRX320, and SRX340, and SRX550M devices.

Required Privilege Level

view

Output Fields

Table 1013 on page 23242 lists the output fields for the `show pppoe interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 1013: show pppoe interfaces Output Fields

Field Name	Field Description
`Index`	Index number of the logical interface, which reflects its initialization sequence.
`State`	State of the logical interface: `up` or `down`.
`Session ID`	Session ID.
`Service name`	Type of service required (can be used to indicate an ISP name, a class, or quality of service).
`Configured AC name`	Configured access concentrator name.
`Session AC name`	Name of the access concentrator.
`Remote MAC address` or `Remote MAC`	MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.
`Auto-reconnect timeout`	Timeout value for reconnecting after a PPPoE session is terminated (in seconds).
`Idle timeout`	Length of time (in seconds) that a connection can be idle before disconnecting.
`Session uptime`	Length of time the session has been up, in *hh:mm:ss*.

Table 1013: show pppoe interfaces Output Fields *(Continued)*

Field Name	Field Description
`Ignore End-Of-List tag`	Disables the `End-of-List` tag to continue processing of other tags after the `End-of-List` tag in a PPPoE Active Discovery Offer (PADO) packet.
`Underlying interface`	Interface on which PPPoE is running.
`Packet Type`	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • `PADI`—PPPoE Active Discovery Initiation packets. • `PADO`—PPPoE Active Discovery Offer packets. • `PADR`—PPPoE Active Discovery Request packets. • `PADS`—PPPoE Active Discovery Session-Confirmation packets. • `PADT`—PPPoE Active Discovery Termination packets. • `Service name error`—Packets for which the Service-Name request could not be honored. • `AC system error`—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • `Generic error`—Packets that indicate an unrecoverable error occurred. • `Malformed packets`—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • `Unknown packets`—Unrecognized packets.
`Timeout`	Timeouts that occur during the PPPoE session: • `PADI`—No PADI packets received within the timeout period. • `PADO`—No PADO packets received within the timeout period. (This value is always zero and is not supported.) • `PADR`—No PADR packets received within the timeout period.

Table 1013: show pppoe interfaces Output Fields *(Continued)*

Field Name	Field Description
`Receive Error Counters`	Error counters received during the PPPoE session: • `PADI`—No PADI error counters received during the session. • `PADO`—No PADO error counters received during the session. • `PADR`—No PADR error counters received during the session. • `PADS`—No PADS error counters received during the session.

Sample Output

show pppoe interfaces

```
user@host> show pppoe interfaces
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:21 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
```

show pppoe interfaces brief

```
user@host> show pppoe interfaces brief
Interface       Underlying          State       Session    Remote
                interface                       ID         MAC
pp0.0           ge-0/0/1.0          Session up  4          b0:c6:9a:74:5e:c1
```

show pppoe interfaces detail

```
user@host> show pppoe interfaces detail
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:21 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
  Ignore End-Of-List tag: Enable
```

show pppoe interfaces extensive

```
user@host> show pppoe interfaces extensive
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:22 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
  PacketType                       Sent         Received
    PADI                              1                0
    PADO                              0                1
    PADR                              1                0
    PADS                              0                1
    PADT                              0                0
    Service name error                0                0
    AC system error                   0                0
    Generic error                     0                0
    Malformed packets                 0                0
    Unknown packets                   0                0
  Timeout
    PADI                              0
    PADO                              0
    PADR                              0
  Receive Error Counters
```

```
    PADI                            0
    PADO                            0
    PADR                            0
    PADS                            0
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

show pppoe interfaces

IN THIS SECTION

- Syntax | 23246
- Description | 23247
- Options | 23247
- Required Privilege Level | 23247
- Output Fields | 23247
- Sample Output | 23250
- Release Information | 23251

Syntax

```
show pppoe interfaces
<brief | detail
<pp0.logical>
```

Description

Display session-specific information about PPPoE interfaces.

Options

none Display interface information for all PPPoE interfaces.

brief | detail (Optional) Display the specified level of output.

pp0.*logical* (Optional) Name of an interface. The logical unit number for static interfaces can be a value from 0 through 16385. The logical unit number for dynamic interfaces can be a value from 1073741824 through the maximum number of logical interfaces supported on your router.

Required Privilege Level

view

Output Fields

Table 1014 on page 23247 lists the output fields for the `show pppoe interfaces` command. Output fields are listed in the approximate order in which they appear. Not all fields are displayed for PPPoE interfaces on M120 and M320 routers in server mode.

Table 1014: show pppoe interfaces Output Fields

Field Name	Field Description	Level of Output
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Index number of the logical interface, which reflects its initialization sequence.	`detail` `extensive` none
`State`	State of the logical interface: `up` or `down`.	All levels

Table 1014: show pppoe interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Session ID`	Session ID.	All levels
`Type`	Origin of the logical interface: `Static` or `Dynamic`. Indicates whether the interface was statically or dynamically created.	`detail extensive` `none`
`Service name`	Type of service required (can be used to indicate an ISP name or a class or quality of service).	`detail extensive` `none`
`Configured AC name`	Configured access concentrator name.	`detail extensive` `none`
`Session AC name`	Name of the access concentrator.	`detail extensive` `none`
`Remote MAC address` or `Remote MAC`	MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.	All levels
`Session uptime`	Length of time the session has been up, in *hh:mm:ss*.	`detail extensive` `none`
`Dynamic Profile`	Name of the dynamic profile that was used to create this interface. If the interface was statically created, this field is not displayed.	`detail extensive` `none`
`Underlying interface`	Interface on which PPPoE is running.	All levels
`Agent Circuit ID`	Agent circuit identifier (ACI) that corresponds to the DSLAM interface that initiated the client service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both the beginning and end of the string. If the agent circuit ID is not configured, this field is not displayed.	`detail extensive` `none`

Table 1014: show pppoe interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Agent Remote ID`	Agent remote identifier that corresponds to the subscriber associated with the DSLAM interface that initiated the service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both at the beginning and end of the string. If the agent remote ID is not configured, this field is not displayed.	`detail` `extensive` none
`ACI Interface Set`	Internally-generated name of the dynamic ACI interface set, if configured, and the set index number of the ACI entry in the session database.	`detail` `extensive` none
`Packet Type`	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • `PADI`—PPPoE Active Discovery Initiation packets. • `PADO`—PPPoE Active Discovery Offer packets. • `PADR`—PPPoE Active Discovery Request packets. • `PADS`—PPPoE Active Discovery Session-Confirmation packets. • `PADT`—PPPoE Active Discovery Termination packets. • `Service name error`—Packets for which the Service-Name request could not be honored. • `AC system error`—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • `Generic error`—Packets that indicate an unrecoverable error occurred. • `Malformed packets`—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • `Unknown packets`—Unrecognized packets.	`extensive`

Sample Output

show pppoe interfaces

```
user@host> show pppoe interfaces
pp0.0 Index 66
  State: Down, Session ID: None,
  Service name: None, Configured AC name: sapphire,
  Session AC name: None, Remote MAC address: 00:00:5e:00:53:00,
  Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
  Underlying interface: at-5/0/0.0 Index 71
```

show pppoe interfaces (Status for the Specified Interface)

```
user@host> show pppoe interfaces pp0.1073741827
pp0.1073741827 Index 70
  State: Session Up, Session ID: 30, Type: Dynamic,
  Session AC name: velorum,
  Remote MAC address: 00:00:5e:00:53:c1,
  Session uptime: 16:45:46 ago,
  Underlying interface: ge-2/0/3.1 Index 73
  Service name: premium
  Dynamic Profile: PppoeProfile
  Agent Circuit ID: velorum-ge-2/0/3
  Agent Remote ID: westford
```

show pppoe interfaces brief

```
user@host> show pppoe interfaces brief
Interface       Underlying            State       Session    Remote
                interface                         ID         MAC
pp0.0           ge-2/0/3.2            Session Up  27         00:00:5e:00:53:c1
pp0.1           ge-2/0/3.2            Session Up  28         00:00:5e:00:53:c1
pp0.1073741824  ge-2/0/3.1            Session Up  29         00:00:5e:00:53:c1
pp0.1073741825  ge-2/0/3.1            Session Up  30         00:00:5e:00:53:c1
pp0.1073741826  ge-2/0/3.1            Session Up  31         00:00:5e:00:53:c1
```

show pppoe interfaces detail

```
user@host> show pppoe interfaces detail
pp0.0 Index 66
  State: Down, Session ID: None, Type: Static,
  Service name: None, Configured AC name: sapphire,
  Session AC name: None, Remote MAC address: 00:00:5e:00:53:00,
  Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
  Underlying interface: at-5/0/0.0 Index 71
```

show pppoe interfaces (PPPoE Subscriber Interface with ACI Interface Set)

```
user@host> show pppoe interfaces pp0.1073741827
pp0.1073741827 Index 346
  State: Session Up, Session ID: 4, Type: Dynamic,
  Service name: AGILENT, Remote MAC address: 00:00:5e:00:53:62,
  Session AC name: nbc,
  Session uptime: 6d 02:22 ago,
  Dynamic Profile: aci-vlan-pppoe-profile,
  Underlying interface: demux0.1073741826 Index 345
  Agent Circuit ID: aci-ppp-dhcp-dvlan-50
  ACI Interface Set: aci-1002-demux0.1073741826 Index 2
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

show pppoe lockout

IN THIS SECTION

- Syntax | 23252
- Description | 23252
- Options | 23252
- Required Privilege Level | 23253
- Output Fields | 23253
- Sample Output | 23255
- Release Information | 23256

Syntax

```
show pppoe lockout
<underlying-interface-name>
```

Description

Display summary information about PPPoE clients currently undergoing lockout or currently in a lockout grace period on all PPPoE underlying logical interfaces or on a specified PPPoE underlying logical interface. You can configure PPPoE subscriber session lockout, also known as short-cycle protection, for VLAN, VLAN demux, and PPPoE-over-ATM dynamic subscriber interfaces.

Options

none	Display information about the lockout condition and the lockout grace period for PPPoE clients on all PPPoE underlying logical interfaces.
underlying-interface-name	(Optional) Name of the PPPoE underlying logical interface. If you do not specify an underlying interface, the router iteratively displays output for all existing clients undergoing lockout per PPPoE underlying logical interface.

Required Privilege Level

view

Output Fields

Table 1015 on page 23253 lists the output fields for the `show pppoe lockout` command. Output fields are listed in the approximate order in which they appear.

Table 1015: show pppoe lockout Output Fields

Field Name	Field Description
underlying-interface-name	Name of the PPPoE underlying logical interface.
Index	Index number of the logical interface, which reflects its initialization sequence.
Device	Name of the physical interface or aggregated Ethernet bundle.
SVLAN	Stacked VLAN ID, also known as the *outer tag*.
VLAN	VLAN ID, also know as the *inner tag*.
VPI	Virtual path identifier value for the PPPoE client.
VCI	Virtual circuit identifier value for the PPPoE client.

Table 1015: show pppoe lockout Output Fields *(Continued)*

Field Name	Field Description
`Short-Cycle Protection`	State of PPPoE short-cycle protection, also known as PPPoE subscriber session lockout, on the underlying interface: • `circuit-id`—Filters PPPoE client sessions by their agent circuit identifier (ACI) value when configured for short-cycle protection • `mac-address`—Filters PPPoE client sessions by their unique media access control (MAC) address when configured for short-cycle protection • `off`—Short-cycle protection not configured for PPPoE client sessions Enabling short-cycle protection temporarily prevents (locks out) a failed or short-lived (short-cycle) PPPoE subscriber session from reconnecting to the router for a default or configurable period of time.
`Lockout Time (seconds)`	Displays the PPPoE lockout time range, the number of PPPoE clients in lockout condition, and the number of PPPoE clients in a lockout grace period: • `Min`—Minimum lockout time, in seconds, configured on the PPPoE underlying interface. • `Max`—Maximum lockout time, in seconds, configured on the PPPoE underlying interface. • `Total clients in lockout`—Number of PPPoE clients currently undergoing lockout. • `Total clients in lockout grace period`—Number of PPPoE clients currently in a lockout grace period. A *lockout grace period* occurs when the time between lockout events is greater than either 15 minutes or the maximum lockout time.
`Client Address`	MAC source address or agent circuit idenfiier (ACI) value of the PPPoE client.
`Current`	Current lockout time, in seconds; displays `0` (zero) if the PPPoE client is not undergoing lockout.
`Elapsed`	Time elapsed into the lockout period, in seconds; displays `0` (zero) if the PPPoE client is not undergoing lockout
`Next`	Lockout time, in seconds, that the router uses for the next lockout event; displays a nonzero value if the PPPoE client is currently in a lockout grace period.

Sample Output

show pppoe lockout (ACI-Based Short-Cycle Protection)

```
user@host> show pppoe lockout at-1.0.0.30
at-1/0/0.30 Index 10305
Device: at-1/0/0, VPI: 1, VCI: 30
Short Cycle Protection: circuit-id,
  Lockout Time (seconds): Min: 1, Max: 300
    Total clients in lockout: 1
    Total clients in lockout grace period: 1

  Client Address                      Current   Elapsed    Next
    Relay-identifier atm 3/0:100.33        64        22     128
      00:00:5e:00:53:ab
      00:00:5e:00:53:21
```

show pppoe lockout (MAC Address-Based Short-Cycle Protection)

```
user@host> show pppoe lockout demux0.100
demux0.100 Index 10305
Device: xe-1/0/0, SVLAN: 100, VLAN: 100,
  Short-Cycle Protection: mac-address,
  Lockout Time (seconds): Min: 1, Max: 300
    Total clients in lockout: 3
    Total clients in lockout grace period: 1

  Client Address                      Current   Elapsed    Next
    00:00:5e:00:53:15                      16        10      32
    00:00:5e:00:53:ab                     256       168     300
    00:00:5e:00:53:23                       0         0       8
```

show pppoe lockout (Short-Cycle Protection Not Configured)

```
user@host> show pppoe lockout xe-1/0/0.1
xe-1/0/0.0 Index 10305
Device: xe-1/0/0,
  Short-Cycle Protection: Off,
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Verifying and Managing Dynamic PPPoE Configuration

Configuring Lockout of PPPoE Subscriber Sessions

show pppoe lockout atm-identifier

IN THIS SECTION

- Syntax | 23256
- Description | 23256
- Options | 23257
- Required Privilege Level | 23257
- Output Fields | 23257
- Sample Output | 23259
- Release Information | 23260

Syntax

```
show pppoe lockout atm-identifier device-name device-name vpi vpi-identifier vci vci-identifier
```

Description

Display information about the lockout condition or lockout grace period for all PPPoE subscriber sessions associated with the specified ATM encapsulation type identifiers. Because the lockout condition persists even in the absence of an underlying interface or after automatic removal of the VLAN or VLAN demux interface, using the `show pppoe lockout atm-identifier` command enables you to

display the lockout condition for PPPoE clients by specifying ATM identifying characteristics instead of the ATM interface name.

The following characteristics comprise the ATM encapsulation type identifier:

- Device name (physical interface or aggregated Ethernet bundle)
- Virtual path identifier (VPI)
- Virtual circuit identifier (VCI)

You can configure PPPoE subscriber session lockout, also known as PPPoE short-cycle protection, for VLAN, VLAN demux, and PPPoE-over-ATM dynamic subscriber interfaces.

Options

device-name Name of the ATM physical interface or aggregated Ethernet bundle associated with the PPPoE client for which you want to display lockout information.

vci-identifier ATM VCI value associated with the PPPoE client for which you want to display lockout information.

- **Range:** 0 through 65535

vpi-identifier ATM VPI value associated with the PPPoE client for which you want to display lockout information.

- **Range:** 0 through 255

Required Privilege Level

view

Output Fields

Table 1016 on page 23258 lists the output fields for the `show pppoe lockout atm-identifier` command. Output fields are listed in the approximate order in which they appear.

Table 1016: show pppoe lockout atm-identifier Output Fields

Field Name	Field Description
underlying-interface-name	Name of the PPPoE underlying logical interface. If no associated underlying interface exists, the underlying interface name is not displayed. Instead, the command output displays only the encapsulation type identifier.
Index	Index number of the logical interface, which reflects its initialization sequence. If no associated underlying interface exists, the index number is not displayed. Instead, the command output displays only the encapsulation type identifier.
Device	Name of the ATM physical interface or aggregated Ethernet bundle.
VPI	Virtual path identifier value for the PPPoE client.
VCI	Virtual circuit identifier value for the PPPoE client.
Short Cycle Protection	State of PPPoE short-cycle protection, also known as PPPoE subscriber session lockout, on the underlying interface: • circuit-id—Filters PPPoE client sessions by their agent circuit identifier (ACI) value when configured for short-cycle protection. • mac-address—Filters PPPoE client sessions by their unique media access control (MAC) address when configured for short cycle-protection . • off—Short-cycle protection not configured for PPPoE client sessions.

Table 1016: show pppoe lockout atm-identifier Output Fields *(Continued)*

Field Name	Field Description
`Lockout Time (seconds)`	Displays the PPPoE lockout time range, the number of PPPoE clients in lockout condition, and the number of PPPoE clients in a lockout grace period: • `Min`—Minimum lockout time, in seconds, configured on the PPPoE underlying interface. • `Max`—Maximum lockout time, in seconds, configured on the PPPoE underlying interface. • `Total clients in lockout`—Number of PPPoE clients currently undergoing lockout. • `Total clients in lockout grace period`—Number of PPPoE clients currently in a lockout grace period. A *lockout grace period* occurs when the time between lockout events is greater than either 15 minutes or the maximum lockout time.
`Client Address`	MAC source address or agent circuit identifier (ACI) value of the PPPoE client.
`Current`	Current lockout time, in seconds; displays `0` (zero) if the PPPoE client is not undergoing lockout.
`Elapsed`	Time elapsed into the lockout period, in seconds; displays `0` (zero) if the PPPoE client is not undergoing lockout.
`Next`	Lockout time, in seconds, that the router uses for the next lockout event; displays a nonzero value if the PPPoE client is currently in a lockout grace period.

Sample Output

show pppoe lockout atm-identifier device-name vpi vci (PPPoE Client with Specified VPI and VCI on ATM Physical Interface)

```
user@host> show pppoe-lockout atm-identifier device-name at-1/0/0 vpi 1 vci 30
at-1/0/0.30 Index 10305
Device: at-1/0/0, VPI: 1, VCI: 30
Short Cycle Protection: circuit-id,
  Lockout Time (seconds): Min: 1, Max: 300
    Total clients in lockout: 1
    Total clients in lockout grace period: 1
```

```
Client Address                      Current   Elapsed    Next
  Relay-identifier atm 3/0:100.33        64        22     128
    00:00:5e:00:53:ab
    00:00:5e:00:53:21
```

Release Information

Command introduced in Junos OS Release 15.2.

RELATED DOCUMENTATION

Verifying and Managing Dynamic PPPoE Configuration

Configuring Lockout of PPPoE Subscriber Sessions

show pppoe lockout vlan-identifier

IN THIS SECTION

- Syntax | 23260
- Description | 23261
- Options | 23261
- Required Privilege Level | 23261
- Output Fields | 23261
- Sample Output | 23263
- Release Information | 23264

Syntax

```
show pppoe lockout vlan-identifier device-name device-name
<svlan-id svlan-identifier>
<vlan-id vlan-identifier>
```

Description

Display information about the lockout condition or lockout grace period for all PPPoE subscriber sessions associated with the specified VLAN encapsulation type identifiers. Because the lockout condition persists even in the absence of an underlying interface or after automatic removal of the VLAN or VLAN demux interface, using the `show pppoe lockout vlan-identifier` command enables you to display the lockout condition for PPPoE clients by specifying VLAN identifying characteristics instead of the underlying interface name.

The following characteristics comprise the VLAN encapsulation type identifier:

- Device name (physical interface or aggregated Ethernet bundle)
- Stacked VLAN (S-VLAN) ID (also known as the *outer tag*)
- VLAN ID (also known as the *inner tag*)

You can configure PPPoE subscriber session lockout, also known as PPPoE short-cycle protection, for VLAN, VLAN demux, and PPPoE-over-ATM dynamic subscriber interfaces.

Options

device-name Name of the Ethernet physical interface or aggregated Ethernet bundle associated with the PPPoE client for which you want to display lockout information.

svlan-identifier (Optional) A valid S-VLAN identifier associated with the PPPoE client for which you want to display lockout information.

- **Range:** 1 through 4094

vlan-identifier (Optional) A valid VLAN identifier associated with the PPPoE client for which you want to display lockout information.

- **Range:** 1 through 4094

Required Privilege Level

view

Output Fields

Table 1017 on page 23262 lists the output fields for the `show pppoe lockout vlan-identifier` command. Output fields are listed in the approximate order in which they appear.

Table 1017: show pppoe lockout vlan-identifier Output Fields

Field Name	Field Description
underlying-interface-name	Name of the PPPoE underlying logical interface. If no associated underlying interface exists, the underlying interface name is not displayed. Instead, the command output displays only the encapsulation type identifier.
Index	Index number of the logical interface, which reflects its initialization sequence. If no associated underlying interface exists, the index number is not displayed. Instead, the command output displays only the encapsulation type identifier.
Device	Name of the Ethernet physical interface or aggregated Ethernet bundle.
SVLAN	Stacked VLAN ID, also known as the *outer tag*.
VLAN	VLAN ID, also know as the *inner tag*.
Short Cycle Protection	State of PPPoE short-cycle protection, also known as PPPoE subscriber session lockout, on the underlying interface: • circuit-id—Filters PPPoE client sessions by their agent circuit identifier (ACI) value when configured for short-cycle protection. • mac-address—Filters PPPoE client sessions by their unique media access control (MAC) address when configured for short-cycle protection . • off—Short-cycle protection not configured for PPPoE client sessions.

Table 1017: show pppoe lockout vlan-identifier Output Fields *(Continued)*

Field Name	Field Description
`Lockout Time (seconds)`	PPPoE lockout time range, the number of PPPoE clients in lockout condition, and the number of PPPoE clients in a lockout grace period: • `Min`—Minimum lockout time, in seconds, configured on the PPPoE underlying interface. • `Max`—Maximum lockout time, in seconds, configured on the PPPoE underlying interface. • `Total clients in lockout`—Number of PPPoE clients currently undergoing lockout. • `Total clients in lockout grace period`—Number of PPPoE clients currently in a lockout grace period. A *lockout grace period* occurs when the time between lockout events is greater than either 15 minutes or the maximum lockout time.
`Client Address`	MAC source address or agent circuit identifier (ACI) value of the PPPoE client.
`Current`	Current lockout time, in seconds; displays `0` (zero) if the PPPoE client is not undergoing lockout.
`Elapsed`	Time elapsed into the lockout period, in seconds; displays `0` (zero) if the PPPoE client is not undergoing lockout.
`Next`	Lockout time, in seconds, that the router uses for the next lockout event; displays a nonzero value if the PPPoE client is currently in a lockout grace period.

Sample Output

show pppoe lockout vlan-identifier device-name vlan-id (Single-Tagged VLAN on Aggregated Ethernet Bundle)

```
user@host> show pppoe lockout vlan-identifier device-name ae0 vlan-id 100
Device: ae0, VLAN: 100
  Short-Cycle Protection level: mac-address,
  Lockout Time (seconds): Min: 1, Max: 300
    Total clients in lockout: 3
    Total clients in lockout grace period: 1
```

```
  Client Address                      Current   Elapsed    Next
    00:00:5e:00:53:15                      16        10      32
    00:00:5e:00:53:ab                     256       168     300
    00:00:5e:00:53:23                       0         0       8
```

show pppoe lockout vlan-identifier device-name svlan-id vlan-id (Dual-Tagged VLAN on Gigabit Ethernet Interface)

```
user@host> show pppoe lockout vlan-identifier device-name ge-1/1/0 svlan-id 100 vlan-id 1
Device: ge-1/1/0, SVLAN: 100, VLAN: 1
  Short Cycle Protection: mac-address,
  Lockout Time (sec):  Min: 30, Max: 90
    Total clients in lockout: 0
    Total clients in lockout grace period: 1
  Client Address                     Current    Elapsed    Next
    00:00:5e:00:53:22                      0          0      60
```

show pppoe lockout vlan-identifier device-name (Untagged VLAN on Aggregated Ethernet Bundle)

```
user@host> show pppoe lockout vlan-identifier device-name ae2
Device: ae3
  Short Cycle Protection: mac-address,
  Lockout Time (sec):  Min: 30, Max: 90
    Total clients in lockout: 0
    Total clients in lockout grace period: 1
  Client Address                     Current    Elapsed    Next
    00:00:5e:00:53:22                      0          0      60
```

Release Information

Command introduced in Junos OS Release 15.2.

RELATED DOCUMENTATION

Verifying and Managing Dynamic PPPoE Configuration

Configuring Lockout of PPPoE Subscriber Sessions

show pppoe service-name-tables

IN THIS SECTION

- Syntax | 23265
- Description | 23265
- Options | 23265
- Required Privilege Level | 23265
- Output Fields | 23266
- Sample Output | 23267
- Release Information | 23268

Syntax

```
show pppoe service-name-tables
<table-name>
```

Description

Display configuration information about PPPoE service name tables.

Options

none	Display the names of configured PPPoE service name tables.
table-name	(Optional) Name of a configured PPPoE service name table.

Required Privilege Level

view

Output Fields

Table 1018 on page 23266 lists the output fields for the `show pppoe service-name-tables` command. Output fields are listed in the approximate order in which they appear.

Table 1018: show pppoe service-name-tables Output Fields

Field Name	Field Description	Level of Output
`Service Name Table`	Name of the PPPoE service name table.	none
`Service Name`	Name of a configured service in the PPPoE service name table: • `<empty>`—Service of zero length that represents an unspecified service • `<any>`—Default service for non-empty service entries that do not match the configured empty or named service entries • *`service-name`*—Named service entry	none
`Action`	Action taken when the PPPoE underlying interface interface receives a PPPoE Active Discovery Initiation (PADI) packet with the specified named service, `empty` service, `any` service, or ACI/ARI pair: • `Delay` *`seconds`*—Number of seconds that the interface delays before responding with a PPPoE Active Discovery Offer (PADO) packet • `Drop`—Interface drops (ignores) the packet. • `Terminate`—Interface responds immediately with a PADO packet	none
`Dynamic Profile`	Name of the dynamic profile with which the router creates a dynamic PPPoE subscriber interface. A dynamic profile can be assigned to a named service, `empty` service, `any` service, or ACI/ARI pair.	none
`Routing Instance`	Name of the routing instance in which to instantiate the dynamic PPPoE subscriber interface. A routing instance can be assigned to a named service, `empty` service, `any` service, or ACI/ARI pair.	none

Table 1018: show pppoe service-name-tables Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Max Sessions`	Maximum number of active PPPoE sessions that the router can establish with the specified named service, `empty` service, or `any` service.	none
`Active Sessions`	Current count of active PPPoE sessions created using the specified named service, `empty` service, or `any` service. The Active Sessions value cannot exceed the Max Sessions value.	none
`ACI`	Agent circuit identifier (ACI) that corresponds to the DSLAM interface that initiated the client service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both the beginning and end of the string. An ACI can be configured as part of an ACI/ARI pair for a named service, `empty` service, or `any` service.	none
`ARI`	Agent remote identifier (ARI) that corresponds to the subscriber associated with the DSLAM interface that initiated the service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both at the beginning and end of the string. An ARI can be configured as part of an ACI/ARI pair for a named service, `empty` service, or `any` service.	none
`Static Interface`	Name of the static PPPoE interface reserved for exclusive use by the PPPoE client with matching ACI/ARI information. A static interface can be configured only for an ACI/ARI pair.	none

Sample Output

show pppoe service-name-tables

```
user@host> show pppoe service-name-tables
Service Name Table: test1
Service Name Table: test2
Service Name Table: test3
```

show pppoe service-name-tables (For the Specified Table Name)

```
user@host> show pppoe service-name-tables Table1
Service Name Table: Table1
  Service Name: <empty>
    Action: Terminate
    Dynamic Profile: BasicPppoeProfile
    Max Sessions: 100
    Active Sessions: 3
  Service Name: <any>
    Action: Drop
    ACI: velorum-ge-2/0/3
    ARI: westford
      Action: Terminate
      Static Interface: pp0.100
    ACI: volantis-ge-5/0/5
    ARI: sunnyvale
       Action: Terminate
       Static Interface: pp0.101
  Service Name: Wholesale
    Action: Terminate
    Dynamic Profile: WholesalePppoeProfile
    Routing Instance: WholesaleRI
    Max Sessions: 16000
    Active Sessions: 4
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Verifying a PPPoE Configuration

Verifying and Managing Dynamic PPPoE Configuration

show pppoe sessions

IN THIS SECTION

- Syntax | 23269
- Description | 23269
- Options | 23269
- Required Privilege Level | 23270
- Output Fields | 23270
- Sample Output | 23270
- Release Information | 23271

Syntax

```
show pppoe sessions
<aci circuit-id-string>
<ari remote-id-string>
<service service-name>
```

Description

Display information about all active PPPoE sessions on the router, or about the active PPPoE sessions established for a specified service name, agent circuit identifier (ACI), or agent remote identifier (ARI).

Options

none	Display information for all active PPPoE sessions on the router.
aci ***circuit-id-string***	(Optional) Display information only for active PPPoE sessions established with the specified agent circuit identifier. The agent circuit identifier corresponds to the DSLAM interface that initiated the service request.
ari ***remote-id-string***	(Optional) Display information only for active PPPoE sessions established with the specified agent remote identifier. The agent remote identifier corresponds to the subscriber associated with the DSLAM interface that initiated the service request.

service *service-name* (Optional) Display information only for active PPPoE sessions established with the specified service, where *service-name* can be `empty`, `any`, or a named service.

Required Privilege Level

view

Output Fields

Table 1019 on page 23270 lists the output fields for the `show pppoe sessions` command. Output fields are listed in the approximate order in which they appear.

Table 1019: show pppoe sessions Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the statically-created or dynamically-created PPPoE interface for the active PPPoE session.	none
`Underlying interface`	Interface on which PPPoE is running.	none
`State`	State of the PPPoE session; displays `Session Up` for active PPPoE sessions.	none
`Session ID`	PPPoE session identifier.	none
`Remote MAC`	MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.	none

Sample Output

show pppoe sessions (For All Active Sessions)

```
user@host> show pppoe sessions
Interface       Underlying            State       Session    Remote
                interface                         ID         MAC
```

```
pp0.0           ge-2/0/3.2            Session Up  27         00:00:5e:00:53:c1
pp0.1           ge-2/0/3.2            Session Up  28         00:00:5e:00:53:c1
pp0.1073741824  ge-2/0/3.1            Session Up  29         00:00:5e:00:53:c1
pp0.1073741825  ge-2/0/3.1            Session Up  30         00:00:5e:00:53:c1
pp0.1073741826  ge-2/0/3.1            Session Up  31         00:00:5e:00:53:c1
```

show pppoe sessions (For All Active Sessions Matching the Agent Circuit Identifier)

```
user@host> show pppoe sessions aci “velorum-ge-2/0/3”
Interface       Underlying            State       Session    Remote
                interface                         ID         MAC
pp0.0           ge-2/0/3.2            Session Up  27         00:00:5e:00:53:c1
pp0.1           ge-2/0/3.2            Session Up  28         00:00:5e:00:53:c1
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Verifying a PPPoE Configuration

Verifying and Managing Dynamic PPPoE Configuration

show pppoe statistics

IN THIS SECTION

- Syntax | 23272
- Description | 23272
- Options | 23272
- Required Privilege Level | 23272
- Output Fields | 23272
- Sample Output | 23274

- Release Information | 23275

Syntax

```
show pppoe statistics
<logical-interface-name>
```

Description

Display statistics information about PPPoE interfaces.

Options

none	Display PPPoE statistics for all interfaces.
logical-interface-name	(Optional) Name of a PPPoE underlying logical interface.

Required Privilege Level

view

Output Fields

Table 1020 on page 23273 lists the output fields for the `show pppoe statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1020: show pppoe statistics Output Fields

Field Name	Field Description
`Active PPPoE sessions`	Total number of active PPPoE sessions and the number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • `PADI`—PPPoE Active Discovery Initiation packets. • `PADO`—PPPoE Active Discovery Offer packets. • `PADR`—PPPoE Active Discovery Request packets. • `PADS`—PPPoE Active Discovery Session-Confirmation packets. • `PADT`—PPPoE Active Discovery Termination packets. • `Service name error`—Packets for which the Service-Name request could not be honored. • `AC system error`—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • `Generic error`—Packets that indicate an unrecoverable error occurred. • `Malformed packets`—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • `Unknown packets`—Unrecognized packets.
`Timeouts`	Information about timeouts that occurred during the PPPoE session (not displayed for M120, M320, and MX Series routers): • `PADI`—No PADR packet has been received within the timeout period. (This value is always zero and is not supported.) • `PADO`—No PPPoE Active Discovery Offer packet has been received within the timeout period. • `PADR`—No PADS packet has been received within the timeout period.

Sample Output

show pppoe statistics

```
user@host> show pppoe statistics
Active PPPoE sessions: 1
  PacketType          Sent      Received
  PADI                   0             0
  PADO                   0             0
  PADR                   0             0
  PADS                   0             0
  PADT                   0             0
  Service name error     0             0
  AC system error        0             0
  Generic error          0             0
  Malformed packets      0             0
  Unknown packets        0             0
  Timeouts
  PADI                   0
  PADO                   0
  PADR                   0
```

show pppoe statistics (For the Specified Underlying Interface Only)

```
user@host> show pppoe statistics ge-4/0/3.2
Active PPPoE sessions: 4
  PacketType                        Sent         Received
    PADI                               0                5
    PADO                               5                0
    PADR                               0                5
    PADS                               4                0
    PADT                               0                1
    Service name error                 0                0
    AC system error                    0                0
    Generic error                      0                0
    Malformed packets                  0                0
    Unknown packets                    0                0
```

Release Information

Command introduced before Junos OS Release 7.4.

logical-interface-name option introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

show ppp address-pool | 23209

show pppoe underlying-interfaces | 23275

show pppoe underlying-interfaces

IN THIS SECTION

- Syntax | 23275
- Description | 23275
- Options | 23276
- Required Privilege Level | 23276
- Output Fields | 23276
- Sample Output | 23280
- Release Information | 23284

Syntax

```
show pppoe underlying-interfaces
<brief | detail | extensive>
<lockout>
<logical-interface-name>
```

Description

Display information about PPPoE underlying interfaces.

Options

brief | detail | extensive (Optional) Display the specified level of output.

lockout (Optional) Display summary information about the lockout condition and the lockout grace period for PPPoE clients on the PPPoE underlying interface.

logical-interface-name (Optional) Name of a PPPoE underlying logical interface.

Required Privilege Level

view

Output Fields

Table 1021 on page 23276 lists the output fields for the `show pppoe underlying-interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 1021: show pppoe underlying-interfaces Output Fields

Field Name	Field Description	Level of Output
Underlying Interface	Name of the PPPoE underlying logical interface.	All levels
Service Name Table	Name of the service name table.	All levels
Dynamic Profile	Name of the dynamic profile that was used to create this interface. If the interface was statically created, then the value is **none**.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	`detail extensive`
State	Origin of the logical interface: **Static** or **Dynamic**. Indicates whether the interface was statically or dynamically created.	`detail extensive`

Table 1021: show pppoe underlying-interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Operational States	Fields in this block are actual operational values rather than simply the configured values. The operational values can be the result of RADIUS-initiated changes.	`detail extensive`
Max Sessions	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface. When this number of logical interfaces has been established, all subsequent PPPoE Active Discovery Initiation (PADI) packets are dropped and all subsequent PPPoE Active Discovery Request (PADR) packets trigger PPPoE Active Discovery Session (PADS) error responses.	`detail extensive`
Max Sessions VSA Ignore	Whether the router is configured to ignore (clear) the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks VSA [26-143] and restore the PPPoE maximum session value on the underlying interface to the value configure with the `max-sessions` statement: **Off** (default) or **On**.	`detail extensive` none
Active Sessions	Number of active PPPoE sessions on the underlying interface. If a dynamic profile is listed, then it is the number of active PPPoE sessions on the underlying interface that are using this profile. The Active Sessions value must not exceed the Max Sessions value.	`detail extensive`
Agent Circuit Identifier	Whether the underlying interface is configured with the `agent-circuit-identifier` statement to enable creation of autosensed dynamic VLAN subscriber interfaces based on agent circuit identifier (ACI) information. `Autosensing` indicates that creation of ACI-based dynamic VLAN interfaces is enabled on the underlying interface. If creation of ACI-based dynamic VLANs is not configured on the underlying interface, this field does not appear. **NOTE**: The Agent Circuit Identifier field is replaced with the Line Identity field when an ALI interface set is configured with the `line-identity` autoconfiguration stanza.	`detail extensive` none

Table 1021: show pppoe underlying-interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Line Identity`	Whether the underlying interface is configured with the `line-identity` statement to enable creation of autosensed dynamic VLAN subscriber interfaces based on the specified trusted option: ACI, ARI, both, or neither. `Autosensing` indicates that creation of ALI-based dynamic VLAN interfaces is enabled on the underlying interface. If creation of ALI dynamic VLANs based on trusted options is not configured on the underlying interface, this field does not appear. **NOTE**: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the `agent-circuit-id` autoconfiguration stanza.	`detail extensive` none
Duplicate Protection	State of PPPoE duplicate protection: **On** or **Off**. When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client. The uniqueness of the PPPoE client is determined by the client's MAC address.	`detail extensive`
Short Cycle Protection	State of PPPoE short cycle protection: **mac-address**, **circuit-id**, or **Off**. Enabling short cycle protection, also known as PPPoE lockout, on the PPPoE underlying interface temporarily prevents (locks out) a failed or short-lived (short-cycle) PPPoE subscriber session from reconnecting to the router for a default or configurable period of time. PPPoE client sessions are identified by their unique media access control (MAC) source address or agent circuit identifier (ACI) value.	`detail extensive`
Direct Connect	State of the configuration to ignore DSL Forum VSAs: **On** or **Off**. When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	`detail extensive` none
AC Name	Name of the access concentrator.	`detail extensive`

Table 1021: show pppoe underlying-interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
PacketType	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • **PADI**—PPPoE Active Discovery Initiation packets. • **PADO**—PPPoE Active Discovery Offer packets. • **PADR**—PPPoE Active Discovery Request packets. • **PADS**—PPPoE Active Discovery Session-Confirmation packets. • **PADT**—PPPoE Active Discovery Termination packets. • **Service name error**—Packets for which the Service-Name request could not be honored. • **AC system error**—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • **Generic error**—Packets that indicate an unrecoverable error occurred. • **Malformed packets**—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • **Unknown packets**—Unrecognized packets.	`detail extensive`

Table 1021: show pppoe underlying-interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Lockout Time (sec)	The PPPoE lockout time range, the number of PPPoE clients in lockout condition, and the number of PPPoE clients in a lockout grace period if **Short Cycle Protection** is enabled (**On**): • **Min**—Minimum lockout time, in seconds, configured on the PPPoE underlying interface. • **Max**—Maximum lockout time, in seconds, configured on the PPPoE underlying interface. • **Total clients in lockout**—Number of PPPoE clients currently undergoing lockout. • **Total clients in lockout grace period**—Number of PPPoE clients currently in a lockout grace period. A *lockout grace period* occurs when the time between lockout events is greater than either 15 minutes or the maximum lockout time.	`extensive`
Client Address	MAC source address of the PPPoE client.	`extensive`
Current	Current lockout time, in seconds; displays **0** (zero) if the PPPoE client is not undergoing lockout.	`extensive`
Elapsed	Time elapsed into the lockout period, in seconds; displays `0` if the PPPoE client is not undergoing lockout	`extensive`
Next	Lockout time, in seconds, that the router uses for the next lockout event; displays a nonzero value if the PPPoE client is currently in a lockout grace period.	`extensive`

Sample Output

show pppoe underlying-interfaces brief

```
user@host> show pppoe underlying-interfaces brief
Underlying Interface    Service Name Table    Dynamic Profile
```

```
ge-4/0/3.1              Premium                 None
ge-4/0/3.2              None                    PppoeProfile
```

show pppoe underlying-interfaces detail

```
user@host> show pppoe underlying-interfaces detail
ge-4/0/3.1 Index 73
  Operational States:
  State: Static, Dynamic Profile: None,
  Max Sessions: 4000, Max Sessions VSA Ignore: Off,
  Active Sessions: 0,
  Service Name Table: Premium,
  Direct Connect: Off,
  AC Name: velorum, Duplicate Protection: On,
  Short Cycle Protection: Off

ge-4/0/3.2 Index 78
  Operational States:
  State: Dynamic, Dynamic Profile: PppoeProfile,
  Max Sessions: 500, Max Sessions VSA Ignore: Off,
  Active Sessions: 3,
  Service Name Table: None,
  Direct Connect: Off,
  AC Name: velorum, Duplicate Protection: On,
  Short Cycle Protection: Off
```

show pppoe underlying-interfaces extensive

```
user@host> show pppoe underlying-interfaces extensive
ge-4/0/3.1 Index 73
  Operational States:
  State: Static, Dynamic Profile: None,
  Max Sessions: 4000, Max Sessions VSA Ignore Off,
  Active Sessions: 0,
  Service Name Table: None,
  Direct Connect: Off,
  AC Name: velorum, Duplicate Protection: Off,
  Short Cycle Protection: Off
```

```
  PacketType                        Sent         Received

    PADI                               0                0
    PADO                               0                0
    PADR                               0                0
    PADS                               0                0
    PADT                               0                0
    Service name error                 0                0
    AC system error                    0                0
    Generic error                      0                0
    Malformed packets                  0                0
    Unknown packets                    0                0

ge-4/0/3.2 Index 78
  Operational States:
  State: Dynamic, Dynamic Profile: PppoeProfile,
  Max Sessions: 4000, Max Sessions VSA Ignore: Off
  Active Sessions: 3,
  Service Name Table: None,
  Direct Connect: Off,
  AC Name: velorum, Duplicate Protection: Off,
  Short Cycle Protection: Off

  PacketType                        Sent         Received
    PADI                               0                5
    PADO                               5                0
    PADR                               0                5
    PADS                               4                0
    PADT                               0                1
    Service name error                 0                0
    AC system error                    0                0
    Generic error                      0                0
    Malformed packets                  0                0
    Unknown packets                    0                0
```

show pppoe underlying-interfaces extensive (PPPoE client in lockout condition)

```
user@host> show pppoe underlying-interfaces ge-1/0/0/.0 extensive
ge-1/0/0.0 Index 71
  State: Static, Dynamic Profile: None,
  Max Sessions: 32000, Max Sessions VSA Ignore: Off,
```

```
  Active Sessions: 0,
  Service Name Table: None,
  Direct Connect: Off,
  AC name: winona, Duplicate Protection: On,
  Short Cycle Protection: Off

  PacketType                        Sent         Received
    PADI                               0                7
    PADO                               3                0
    PADR                               0                3
    PADS                               3                0
    PADT                               2                1
    Service name error                 0                0
    AC system error                    0                0
    Generic error                      0                0
    Malformed packets                  0                0
    Unknown packets                    0                0

  Lockout Time (sec):  Min: 1, Max: 30
    Total clients in lockout: 1
    Total clients in lockout grace period: 0

  Client Address             Current     Elapsed    Next
    00:00:5e:00:53:11              4           3       8
```

show pppoe underlying-interfaces lockout

```
user@host> show pppoe underlying-interfaces ge-1/0/0.0 lockout
ge-1/0/0.0 Index 71
  Short Cycle Protection: Off,
  Lockout Time (sec):  Min: 10, Max: 60
    Total clients in lockout: 0
    Total clients in lockout grace period: 0
```

show pppoe underlying-interfaces detail (Autosensing Configured for ACI-based Dynamic VLANs)

```
user@host> show pppoe underlying-interfaces demux0.1073741826 detail
demux0.1073741826 Index 345
```

```
  State: Dynamic, Dynamic Profile: aci-vlan-pppoe-profile,
  Max Sessions: 32000, Max Sessions VSA Ignore: Off,
  Active Sessions: 1,
  Agent Circuit Identifier: Autosensing,
  Service Name Table: None,
  Duplicate Protection: On, Short Cycle Protection: Off,
  Direct Connect: Off,
  AC Name: nbc,
  Short Cycle Protection: circuit-id,
```

show pppoe underlying-interfaces detail (Autosensing Configured for ALI-based Dynamic VLANs)

```
user@host> show pppoe underlying-interfaces demux0.1073741826 detail
demux0.1073741826 Index 345
  State: Dynamic, Dynamic Profile: aci-vlan-pppoe-profile,
  Max Sessions: 32000, Max Sessions VSA Ignore: Off,
  Active Sessions: 1,
  Line Identity: Autosensing,
  Service Name Table: None,
  Duplicate Protection: On, Short Cycle Protection: Off,
  Direct Connect: Off,
  AC Name: nbc,
  Short Cycle Protection: circuit-id,
```

Release Information

Command introduced in Junos OS Release 10.0.

`lockout` option added in Junos OS Release 11.4.

RELATED DOCUMENTATION

Verifying and Managing Dynamic PPPoE Configuration

Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces

Configuring the PPPoE Family for an Underlying Interface

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers

show pppoe version

IN THIS SECTION

- Syntax | 23285
- Description | 23285
- Options | 23285
- Required Privilege Level | 23285
- Output Fields | 23285
- Sample Output | 23286
- Release Information | 23287

Syntax

```
show pppoe version
```

Description

(M120 routers and M320 routers only) Display version information about PPPoE.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1022 on page 23286 lists the output fields for the `show pppoe version` command. Output fields are listed in the approximate order in which they appear.

Table 1022: show pppoe version Output Fields

Field Name	Field Description
version *n*	PPPoE version number and RFC. For example, **version 1, rfc 2516**.
PPPoE protocol	State of the PPPoE protocol: **enabled** or **disabled**.
Maximum Sessions	Maximum active sessions supported per router. The default is 256 sessions.
PADI resend timeout	Initial time, in seconds, that the router waits to receive a PPoE Active Discovery Offer (PADO) packet for the PPoE Active Discovery Initiation (PADI) packet sent. This timeout doubles for each successive PADI packet sent. Not displayed for M120 and M320 routers.
PADR resend timeout	Initial time, in seconds, that the router waits to receive a PPoE Active Discovery Session Confirmation (PADS) packet for the PPoE Active Discovery Request (PADR) packet sent. This timeout doubles for each successive PADR packet sent. Not displayed for M120 and M320 routers.
Max resend timeout	Maximum value, in seconds, that the PADI or PADR resend timer can accept. The maximum value is 64. Not displayed for M120 and M320 routers.
Max Configured AC timeout	Time, in seconds, during which the configured access concentrator must respond. Not displayed for M120 and M320 routers.

Sample Output

show pppoe version

```
user@host> show pppoe version
Point-to-Point Protocol Over Ethernet, version 1. rfc2516
   PPPoE protocol               = Enabled
   Maximum Sessions             = 256
   PADI resend timeout          = 2 seconds
   PADR resend timeout          = 16 seconds
   Max resend timeout           = 64 seconds
   Max Configured AC timeout    = 4 seconds
```

Release Information

Command introduced before Junos OS Release 7.4.

show programmable-rpd clients

IN THIS SECTION

- Syntax | 23287
- Description | 23287
- Options | 23288
- Required Privilege Level | 23288
- Sample Output | 23288
- Release Information | 23288

Syntax

```
show programmable-rpd clients
<client-id>
```

Description

Lists clients connected to the programmable routing protocol process (prpd) server. The prpd provides public APIs to program routing systems, making it possible for users to directly access the APIs to customize, create and modify behavior of their network.

The command output shows client specific details and statistics such as client ID, protocol and corresponding gateway handle, purge timer, the client up/down status, the route-summary of the number of routes installed, the per-route-table breakup of the route count and, if the client is disconnected, the time remaining before the client state is purged. If the client has not registered any protocol, the gateway handle is 0.

Options

none Display the information for all prpd clients.

client-id (Optional) Display information only about the specified prpd client.

Required Privilege Level

view

Sample Output

show programmable-rpd clients

```
user@host> show programmable-rpd clients
RPD global purge timeout: 120
RPD Server connected client details:
ClientIdentifier   PurgeTimer  Status   Timeout  Protocol       Gateway
3                   150         Up               BGP-Static      578
2                    75         Up               NoGwProtocol      0
1                   120         Down     117     BGP-Static       577
```

Release Information

Command introduced in Junos OS Release 16.2.

client-id option introduced in Junos OS Release 22.1R1.

RELATED DOCUMENTATION

purge-timeout (programmable-rpd) | **13870**

show route | **23512**

show programmable-rpd clients route-summary | **23289**

show programmable-rpd clients route-summary

IN THIS SECTION

- Syntax | 23289
- Description | 23289
- Options | 23289
- Required Privilege Level | 23290
- Sample Output | 23290
- Release Information | 23291

Syntax

```
show programmable-rpd clients route-summary
<client-id>
<display-vni-data>
<instance instance-name>
<table table-name>
```

Description

Displays the number of programmed routes for clients connected to the programmable routing protocol process (prpd) server. The prpd provides public APIs to program routing systems, making it possible for users to directly access the APIs to customize, create and modify behavior of their network.

You can filter the route count by client ID, routing instance, or routing table. The number of routes this command displays includes hidden routes.

Options

none	(Optional) Display the total number of programmed routes for all prpd clients.
client-id	(Optional) Display only the number of programmed routes for the specified prpd client.

display-vni-data (Optional) Display the number of programmed routes for each virtual routing and forwarding (VRF) instance and each tunnel that has a virtual network identifier (VNI) number.

instance *instance-name* (Optional) Display only the number of programmed routes for the specified routing instance.

table *table-name* (Optional) Display only the number of programmed routes for the specified routing table.

Required Privilege Level

view

Sample Output

show programmable-rpd clients route-summary

```
user@host> show programmable-rpd clients route-summary

Total Programmed Route Count: 25
                              primary         secondary
        inet.0                  15                 0
        vpnA.inet.0              0                10

client-id: client-1
      Static                     5
        inet.0                   5                 0
        vpnA.inet.0              0                 5
      Bgp-Static                 5
        inet.0                   5                 0
        vpnA.inet.0              0                 5
client-id: client-2
    Static                       5
        inet.0                   5                 0
```

show programmable-rpd clients route-summary display-vni-data

```
user@host> show programmable-rpd clients route-summary display-vni-data | no-more
Total programmed route count: 20
                                            Primary     Secondary
All Clients:                                  20            0
        vrf1.inet.0                            5            0
        vrf2.inet.0                            5            0
        vrf1.inet6.0                           5            0
        vrf2.inet6.0                           5            0
client-id: clientr0
    Static                                    20
        vrf1.inet.0                            5            0
            VNI 1                              1
            VNI 2                              1
            VNI 3                              1
            VNI 4                              1
            VNI 5                              1
        vrf2.inet.0                            5            0
            VNI 11                             1
            VNI 12                             1
            VNI 13                             1
            VNI 14                             1
            VNI 15                             1
        vrf1.inet6.0                           5            0
            VNI 6                              1
            VNI 7                              1
            VNI 8                              1
            VNI 9                              1
            VNI 10                             1
        vrf2.inet6.0                           5            0
            VNI 16                             1
            VNI 17                             1
            VNI 18                             1
            VNI 19                             1
            VNI 20                             1
```

Release Information

Command introduced in Junos OS Release 22.1R1.

`display-vni-data` option introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

show programmable-rpd clients | 23287

show programmable-rpd statistics

IN THIS SECTION

- Syntax | 23292
- Description | 23292
- Required Privilege Level | 23292
- Sample Output | 23293
- Release Information | 23294

Syntax

```
show programmable-rpd statistics
    flex-routes
        <destination>
        table
```

Description

View detailed traffic statistics for programmable RPD (PRPD) flex routes.

Required Privilege Level

view

Sample Output

show programmable-rpd statistics flex-routes

```
user@host> show programmable-rpd statistics flex-routes
Table: 002f10c440414b5c9bf03bfee82b2915.inet.0
         192.168.1.1/32
                 Packets 2570 (0 pps), Bytes 133960 (26 bps)
         192.168.1.2/32
                 Packets 34 (0 pps), Bytes 2856 (0 bps)
         192.168.1.3/32
                 Packets 0 (0 pps), Bytes 0 (0 bps)

Table: vrf1.inet.0
         192.168.1.1/32
                    Packets 2570 (0 pps), Bytes 133960 (26 bps)
         192.168.1.2/32
                    Packets 34 (0 pps), Bytes 2856 (0 bps)
         192.168.1.4/32
                    Packets 0 (0 pps), Bytes 0 (0 bps)
```

show programmable-rpd statistics flex-routes table

```
user@host> show programmable-rpd statistics flex-routes table
002f10c440414b5c9bf03bfee82b2915.inet.0
Table: 002f10c440414b5c9bf03bfee82b2915.inet.0
         192.168.1.1/32
                 Packets 2570 (0 pps), Bytes 133960 (26 bps)
         192.168.1.2/32
                 Packets 34 (0 pps), Bytes 2856 (0 bps)
         192.168.1.3/32
                 Packets 0 (0 pps), Bytes 0 (0 bps)
```

show programmable-rpd statistics flex-routes <destination>

```
user@host> show programmable-rpd statistics flex-routes 192.168.1.1
Table: 002f10c440414b5c9bf03bfee82b2915.inet.0
         192.168.1.1/32
```

```
                  Packets 2570 (0 pps), Bytes 133960 (26 bps)

Table: vrf1.inet.0
         192.168.1.1/32
                    Packets 2570 (0 pps), Bytes 133960 (26 bps)
```

Release Information

Command introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Routing Protocol Process Overview

programmable-rpd | 13369

show protection-group ethernet-ring aps

IN THIS SECTION

- Syntax | 23295
- Description | 23295
- Options | 23295
- Required Privilege Level | 23295
- Output Fields | 23295
- Sample Output | 23296
- Sample Output | 23297
- Sample Output | 23297
- Release Information | 23299

Syntax

```
show protection-group ethernet-ring aps
```

Description

Display the status of the Automatic Protection Switching (APS) and Ring APS (RAPS) messages on an Ethernet ring.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1023 on page 23295 lists the output fields for the `show protection-group ethernet-ring aps` command. Output fields are listed in the approximate order in which they appear.

Table 1023: show protection-group ethernet-ring aps Output Fields

Field Name	Field Description
Ethernet Ring	Name configured for the Ethernet ring.
Request/State	Status of the Ethernet ring RAPS messages. • **NR**—Indicates that there is no request for APS on the ring. • **SF**—Indicates that there is a signal failure on the ring. • `FS`—Indicates that there are active forced-switch requests in the ring. • `MS`—Indicates that there are active manual-switch requests in the ring. **NOTE**: Both `FS` and `MS` values are valid only when G.8032v2 is supported.

Table 1023: show protection-group ethernet-ring aps Output Fields *(Continued)*

Field Name	Field Description
Ring Protection Link Blocked	Blocking on the ring protection link: **Yes** or **No**.
No Flush	Indicates the value of the Do Not Flush (DNF) flag in the received RAPS PDU. If the value is Yes, then FDB flush is not triggered as part of processing of the received RAPS PDU.
`Blocked Port Reference`	This parameter is the reference to the blocked ring port. If the east ring port is blocked, the Blocked Port Reference (BPR) value is 0. If the west ring port is blocked, the BPR value is 1. If both ring ports are blocked, this parameter can take any value. If both east and west ports are blocked or not blocked, the value would be 0. This field is valid only when G.8032v2 is supported.
Blocked Port Reference	Reference of the ring port on which traffic is blocked.
Originator	Indicates whether the node is the originator of the RAPS messages.
Remote Node ID	Identifier (in MAC address format) of the remote node.

Sample Output

show protection-group ethernet-ring aps (EX Switches)

```
user@switch>show protection-group ethernet-ring aps
Ring Name    Request/state  No Flush  RPL Blocked  Originator  Remote Node ID erp1
NR             No        Yes          No          00:1F:12:30:B8:81
```

Sample Output

show protection-group ethernet-ring aps (Owner Node, Normal Operation on ACX and MX Routers)

```
user@host> show protection-group ethernet-ring aps
Ethernet Ring   Request/state  RPL Blocked  No Flush  BPR  Originator  Remote Node ID
Erp_1           NR             Yes          No        1    No          00:00:00:02:00:01
```

Sample Output

show protection-group ethernet-ring aps detail (Owner Node, Normal Operation on ACX and MX Routers)

```
user@host> show protection-group ethernet-ring aps detail
Ethernet-Ring name             : Erp_1
Request/State                  : NR
Ring Protection Link blocked   : Yes
No Flush Flag                  : No
Blocked Port Reference         : 1
Originator                     : No
Remote Node ID                 : 00:00:00:02:00:01
```

show protection-group ethernet-ring aps (MX RPL Owner Ring Node, Failure condition on non-RPL link of the ring)

```
user@host>show protection-group ethernet-ring aps
Ethernet Ring   Request/state   RPL Blocked  No Flush
pg101           SF              No           No

Originator  Remote Node ID
No          00:01:02:00:00:01
```

show protection-group ethernet-ring aps (MX Interconnection Ring Node, Failure condition in major ring on non-RPL link of the ring)

```
user@host>show protection-group ethernet-ring aps
Ethernet Ring   Request/state  RPL Blocked  No Flush  BPR
pg_major        SF             No           No           0
pg_subring      NR                Yes                    Yes                0

Originator  Remote Node ID
No          00:01:00:00:00:01
No                   00:02:00:00:00:02
```

show protection-group ethernet-ring aps (MX Series router)

```
user@host>show protection-group ethernet-ring aps
Ethernet Ring   Request/state  RPL Blocked  No Flush  BPR  Originator  Remote Node ID
Inst_Vlans_1-15 NR             Yes          Yes       1    Yes         NA
Inst_Vlans_16-30 NR            Yes          Yes       0    No          00:00:00:03:00:02
```

show protection-group ethernet-ring aps detail (MX Series router)

```
user@host>show protection-group ethernet-ring aps
Ethernet-Ring name             : Inst_Vlans_1-15
Request/State                  : NR
Ring Protection Link blocked   : Yes
No Flush Flag                  : Yes
Blocked Port Reference         : 1
Originator                     : Yes
Remote Node ID                 : NA

Ethernet-Ring name             : Inst_Vlans_16-30
Request/State                  : NR
Ring Protection Link blocked   : Yes
No Flush Flag                  : Yes
Blocked Port Reference         : 0
Originator                     : No
Remote Node ID                 : 00:00:00:03:00:02
```

show protection-group ethernet-ring aps (MX Interconnection Ring Node as RPL owner of major ring, rings in IDLE state)

```
user@host>show protection-group ethernet-ring aps detail

Ethernet-Ring name             : pg_major
Request/State                  : NR
Ring Protection Link blocked   : Yes
No Flush Flag                  : Yes
Blocked Port Reference         : 0
Originator                     : Yes
Remote Node ID                 : NA

Ethernet-Ring name             : pg_subring
Request/State                  : NR
Ring Protection Link blocked   : Yes
No Flush Flag                  : Yes
Blocked Port Reference         : 0
Originator                     : No
Remote Node ID                 : 00:00:03:00:00:03
```

show protection-group ethernet-ring aps detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring aps detail
Ethernet-Ring name             : pg1001
Request/State                  : NR
Ring Protection Link blocked   : Yes
No Flush Flag                  : Yes
Blocked Port Reference         : 0
Originator                     : Yes
Remote Node ID                 : NA
```

Release Information

Command introduced in Junos OS Release 9.4.

Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

RELATED DOCUMENTATION

show protection-group ethernet-ring data-channel | 23309
show protection-group ethernet-ring interface | 23315
show protection-group ethernet-ring node-state | 23320
show protection-group ethernet-ring statistics | 23328
show protection-group ethernet-ring vlan | 23335

show protection-group ethernet-ring configuration

IN THIS SECTION

- Syntax | 23300
- Description | 23300
- Required Privilege Level | 23300
- Output Fields | 23301
- Sample Output | 23304
- Release Information | 23308

Syntax

```
show protection-group ethernet-ring configuration
```

Description

Display the configuration of Ethernet ring protection group on EX Switches and MX Series routers.

Required Privilege Level

view

Output Fields

Table 1024 on page 23301 lists the output fields for the `show protection-group ethernet-ring configuration` command. Output fields are listed in the approximate order in which they appear.

Table 1024: show protection-group ethernet-ring configuration Output Fields

Output Fields	Field Description
`G8032 Compatability Version`	This is the compatibility version mode of ERP. This parameter always takes the value 1 in the case of G8032v1. This parameter is valid only for MX Series routers.
East Interface	One of the two switch interfaces that participates in a ring link. When Junos supports G8032v2, this interface is treated as interface 0.
West Interface	One of the two interfaces in a switch that participates in a ring link. When Junos supports G8032v2, this interface is treated as interface 1.
Restore Interval	Configured interval of wait time after a link is restored. When a link goes down, the RPL link is activated. When the down link becomes active again, the RPL owner receives a notification. The RPL owner waits for the restore interval before issuing a block on the RPL link. The configured restore interval can be 5 through 12 minutes for ERPv1 and 1 through 12 minutes for ERPv2. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value. **NOTE**: Wait to Restore (WTR) configuration values on EX2300 and EX3400 switches must be 5-12 minutes.

Table 1024: show protection-group ethernet-ring configuration Output Fields *(Continued)*

Output Fields	Field Description
`Wait to Block Interval`	Configured interval of wait time for link restoration when a manual command (manual switch or force switch) is cleared. On clearing the manual command, the RPL owner receives NR messages, which starts a timer with interval 'Wait to Block' to restore the RPL link after its expiration. This delay timer is set to be 5 seconds longer than the guard timer. The configured number can be from 5 seconds through 10 seconds. The parameter is valid only for G.8032v2. **NOTE**: The Wait To Block Timer (WTB) is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect.
Guard Interval	Configured number of milliseconds (in 10 millisecond intervals, 10 milliseconds through 2000 milliseconds) that the node does not process any Ethernet ring protection protocol data units (PDUs). This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.
`Hold off interval`	This is the interval at which the link is held down even before declaring that the link is down. Because the parameter is not supported at present, its value is always considered 0. This parameter is valid only for MX Series routers.
Node ID	Node ID for the switch or router. If the node ID is not configured, it is assigned by default. For EX Series switches, the Node ID value cannot be configured, whereas for MX Series routers, it can be configured.
`Ring ID`	In G8032v2, the ring ID can be within the range 1–239. All the nodes in a ring should have the same ring ID. In the case of G8032v1, the value of the ring ID is always 1. This parameter is valid only for MX Series routers.
`Node Role`	Indicates whether the ring node is operating as a normal ring-node or RPL-owner or RPL-neighbor. For G8032v1 RPL-neighbor role is not supported. This parameter is valid only for MX Series routers.

Table 1024: show protection-group ethernet-ring configuration Output Fields *(Continued)*

Output Fields	Field Description
`Revertive Mode of Operation`	This parameter indicates whether the ring is operating in revertive mode or nonrevertive mode. In nonrevertive mode of operation, when all links in the ring and Ethernet Ring Nodes have recovered and no external requests are active, the Ethernet Ring does not automatically revert. G8032v1 supports only revertive mode of operation. This parameter is valid only for MX Series routers.
`RAPS Tx Dot1p priority`	The RAPS Tx Dot1p priority is a parameter with which the RAPS is transmitted from the ring node. For G8032v1, the value of this parameter is always 0. For G8032v2, the value of this parameter can be within the range 0–7. This parameter is valid only for MX Series routers.
`Node type`	Indicates whether ring node is a normal ring node having two ring-links or a open ring-node having only a single ring-link or a interconnection ring-node. An interconnection ring node can be connected to major ring in non virtual-channel mode or in virtual channel mode. Ring interconnection is not supported for G8032v1. This parameter is valid only for MX Series routers.
`Major ring name`	If the node type is interconnection in the ring, this parameter takes the name of the major ring to which the sub-ring node is connected. This parameter is valid only for MX Series routers.
`Interconnection mode`	Indicates the interconnection mode if the type of the node is interconnection. An interconnection ring node can be connected to major ring in non-virtual channel mode or in virtual channel mode. This parameter is valid only for MX Series routers.
`Propagate Topology Change event`	When Propagate Topology Change event is set to 1, the change in the topology of sub-ring is propagated to the major ring, enabling the transmission of EVENT FLUSH RAPS PDU in the major ring. When the parameter is set to 0, the topology change in the sub-ring is not propagated to the major ring blocking EVENT FLUSH RAPS PDU transmission in the major ring. This parameter is valid only for MX Series routers.
Control Vlan	The VLAN that transfers ERP PDUs from one node to another.

Table 1024: show protection-group ethernet-ring configuration Output Fields *(Continued)*

Output Fields	Field Description
Physical Ring	Physical ring if the east and west interfaces are nontrunk ports. For MX Series routers, the ring is termed a physical ring if no data channels are defined for the ring and the entire physical port forwarding is controlled by ERP.
`Data Channel VLAN(s)`	Data VLANs for which forwarding behavior is controlled by the ring instance.

Sample Output

show protection-group ethernet-ring configuration (EX Switch)

```
user@switch>show protection-group ethernet-ring configuration
Ethernet ring configuration parameters for protection group erp1
East Interface   : ge-0/0/3.0
West Interface   : ge-0/0/9.0
Restore Interval : 5 minutes
Guard Interval   : 500 ms
Node Id          : 00:1F:12:30:B8:81
Control Vlan     : 101
Physical Ring    : yes
```

show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version                  : 2
East interface (interface 0)                 : xe-2/3/0.1
West interface (interface 1)                 : xe-2/2/1.1
Restore interval                             : 5 minutes
Wait to Block interval                       : 5 seconds
Guard interval                               : 500 ms
Hold off interval                            : 0 ms
Node ID                                      : 64:87:88:65:37:D0
Ring ID (1 ... 239)                          : 1
Node role (normal/rpl-owner/rpl-neighbour)   : normal
```

```
Revertive mode of operation                  : 1
RAPS Tx Dot1p priority (0 .. 7)              : 0
Node type (normal/open/interconnection)      : Normal
Control Vlan                                 : 100
Physical Ring                                : No
Data Channel Vlan(s)                         : 200,300
```

show protection-group ethernet-ring configuration (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration
Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version                  : 2
East interface (interface 0)                 : xe-2/3/0.1
West interface (interface 1)                 : xe-2/2/1.1
Restore interval                             : 5 minutes
Wait to Block interval                       : 5 seconds
Guard interval                               : 500 ms
Hold off interval                            : 0 ms
Node ID                                      : 64:87:88:65:37:D0
Ring ID (1 ... 239)                          : 1
Node role (normal/rpl-owner/rpl-neighbour)   : rpl-neighbour
Node RPL end                                 : east-port
Revertive mode of operation                  : 1
RAPS Tx Dot1p priority (0 .. 7)              : 0
Node type (normal/open/interconnection)      : Normal
Control Vlan                                 : 100
Physical Ring                                : No
Data Channel Vlan(s)                         : 200,300
```

show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version                  : 2
East interface (interface 0)                 : xe-2/3/0.1
West interface (interface 1)                 : xe-2/2/1.1
Restore interval                             : 5 minutes
Wait to Block interval                       : 5 seconds
Guard interval                               : 500 ms
Hold off interval                            : 0 ms
```

```
Node ID                                      : 64:87:88:65:37:D0
Ring ID (1 ... 239)                          : 1
Node role (normal/rpl-owner/rpl-neighbour)   : rpl-owner
Node RPL end                                 : east-port
Revertive mode of operation                  : 1
RAPS Tx Dot1p priority (0 .. 7)              : 0
Node type (normal/open/interconnection)      : Normal
Control Vlan                                 : 100
Physical Ring                                : No
Data Channel Vlan(s)                         : 200,300
```

show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version                  : 2
East interface (interface 0)                 : xe-2/3/0.1
West interface (interface 1)                 : (no erp)
Restore interval                             : 5 minutes
Wait to Block interval                       : 5 seconds
Guard interval                               : 500 ms
Hold off interval                            : 0 ms
Node ID                                      : 64:87:88:65:37:D0
Ring ID (1 ... 239)                          : 1
Node role (normal/rpl-owner/rpl-neighbour)   : rpl-owner
Node RPL end                                 : east-port
Revertive mode of operation                  : 1
RAPS Tx Dot1p priority (0 .. 7)              : 0
Node type (normal/open/interconnection)      : Open
Control Vlan                                 : 100
Physical Ring                                : No
Data Channel Vlan(s)                         : 200,300
```

show protection-group ethernet-ring configuration (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration
Ethernet Ring configuration information for protection group pg_major
G8032 Compatibility Version                  : 2
East interface (interface 0)                 : xe-2/3/0.1
West interface (interface 1)                 : xe-2/2/1.1
```

```
Restore interval                              : 5 minutes
Wait to Block interval                        : 5 seconds
Guard interval                                : 500 ms
Hold off interval                             : 0 ms
Node ID                                       : 64:87:88:65:37:D0
Ring ID (1 ... 239)                           : 1
Node role (normal/rpl-owner/rpl-neighbour)    : rpl-owner
Node RPL end                                  : east-port
Revertive mode of operation                   : 1
RAPS Tx Dot1p priority (0 .. 7)               : 0
Node type (normal/open/interconnection)       : Normal
Control Vlan                                  : 100
Physical Ring                                 : No
Data Channel Vlan(s)                          : 200,300

Ethernet Ring configuration information for protection group pg_subring
G8032 Compatibility Version                   : 2
East interface (interface 0)                  : ge-2/0/0.1
West interface (interface 1)                  : (no erp)
Restore interval                              : 5 minutes
Wait to Block interval                        : 5 seconds
Guard interval                                : 500 ms
Hold off interval                             : 0 ms
Node ID                                       : 64:87:88:65:37:D0
Ring ID (1 ... 239)                           : 2
Node role (normal/rpl-owner/rpl-neighbour)    : normal
Revertive mode of operation                   : 1
RAPS Tx Dot1p priority (0 .. 7)               : 0
Node type (normal/open/interconnection)       : Non-VC-Interconnection
Major ring name                               : pg_major
Interconnection mode (VC/Non-VC)              : Non-VC mode
Propagate Topology Change event               : 0
Control Vlan                                  : 101
Physical Ring                                 : No
Data Channel Vlan(s)                          : 200,300
```

show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
Ethernet Ring configuration information for protection group pg_major
G8032 Compatibility Version                   : 2
```

```
East interface (interface 0)                : xe-2/3/0.1
West interface (interface 1)                : xe-2/2/1.1
Restore interval                            : 5 minutes
Wait to Block interval                      : 5 seconds
Guard interval                              : 500 ms
Hold off interval                           : 0 ms
Node ID                                     : 64:87:88:65:37:D0
Ring ID (1 ... 239)                         : 1
Node role (normal/rpl-owner/rpl-neighbour)  : rpl-owner
Node RPL end                                : east-port
Revertive mode of operation                 : 1
RAPS Tx Dot1p priority (0 .. 7)             : 0
Node type (normal/open/interconnection)     : Normal
Control Vlan                                : 100
Physical Ring                               : No
Data Channel Vlan(s)                        : 200,300

Ethernet Ring configuration information for protection group pg_subring
G8032 Compatibility Version                 : 2
East interface (interface 0)                : ge-2/0/0.1
West interface (interface 1)                : (no erp)
Restore interval                            : 5 minutes
Wait to Block interval                      : 5 seconds
Guard interval                              : 500 ms
Hold off interval                           : 0 ms
Node ID                                     : 64:87:88:65:37:D0
Ring ID (1 ... 239)                         : 2
Node role (normal/rpl-owner/rpl-neighbour)  : normal
Revertive mode of operation                 : 1
RAPS Tx Dot1p priority (0 .. 7)             : 0
Node type (normal/open/interconnection)     : Non-VC-Interconnection
Major ring name                             : pg_major
Interconnection mode (VC/Non-VC)            : Non-VC mode
Propagate Topology Change event             : 0
Control Vlan                                : 101
Physical Ring                               : No
Data Channel Vlan(s)                        : 200,300
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

show protection-group ethernet-ring aps | 23294
show protection-group ethernet-ring data-channel | 23309
show protection-group ethernet-ring interface | 23315
show protection-group ethernet-ring node-state | 23320
show protection-group ethernet-ring statistics | 23328
show protection-group ethernet-ring vlan | 23335

show protection-group ethernet-ring data-channel

IN THIS SECTION

- Syntax | 23309
- Description | 23309
- Options | 23310
- Required Privilege Level | 23310
- Output Fields | 23310
- Sample Output | 23311
- Release Information | 23312

Syntax

```
show protection-group ethernet-ring data-channel
<brief | detail>
<group-name group-name>
```

Description

Display the configuration of Ethernet ring protection group on EX Switches and MX Series routers.

Options

brief | detail (Optional) Display the specified level of output.

group-name (Optional) Protection group for which to display statistics. If you omit this optional field, all protection group statistics for configured groups will be displayed.

Required Privilege Level

view

Output Fields

Table 1025 on page 23310 lists the output fields for the `show protection-group ethernet-ring data-channel` command. Output fields are listed in the approximate order in which they appear.

Table 1025: show protection-group ethernet-ring data-channel Output Fields

Field Name	Field Description
Interface	Name of the interface configured for the Ethernet ring.
STP index	The Spanning Tree Protocol (STP) index number used by each interface in an Ethernet ring. The STP index controls the forwarding behavior for a set of VLANs on a data channel on an Ethernet ring port. For multiple Ethernet ring instances on an physical ring port, there are multiple STP index numbers. Different ring instances will have different STP index numbers and may have different forwarding behavior.
Forward State	Forwarding state on the Ethernet ring. • **fowarding**—Indicates packets are being forwarded. • **discarding**—Indicates packets are being discarded.

Sample Output

show protection-group ethernet-ring data-channel

```
user@host> show protection-group ethernet-ring data-channel
Ethernet ring data channel information for protection group pg301

Interface    STP index  Forward State
xe-5/0/2     78         forwarding
xe-2/2/0     79         discarding

Ethernet ring data channel parameters for protection group pg302

Interface    STP index  Forward State
xe-5/0/2     80         forwarding
xe-2/2/0     81         forwarding
```

show protection-group ethernet-ring data-channel detail

```
user@host> show protection-group ethernet-ring data-channel detail
Ethernet ring data channel parameters for protection group pg301

Interface name                 : xe-5/0/2
STP index                      : 78
Forward State                  : forwarding

Interface name                 : xe-2/2/0
STP index                      : 79
Forward State                  : discarding


Ethernet ring data channel parameters for protection group pg302

Interface name                 : xe-5/0/2
STP index                      : 80
Forward State                  : forwarding

Interface name                 : xe-2/2/0
```

```
STP index                       : 81
Forward State                   : forwarding
```

show protection-group ethernet-ring data-channel detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring data-channel detail
Ethernet ring data channel parameters for protection group pg1001

Interface name                  : ge-0/0/42
STP index                       : 52
Forward State                   : discarding

Interface name                  : ge-0/0/38
STP index                       : 53
Forward State                   : forwarding
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

- show protection-group ethernet-ring aps | 23294
- show protection-group ethernet-ring interface | 23315
- show protection-group ethernet-ring node-state | 23320
- show protection-group ethernet-ring statistics | 23328
- show protection-group ethernet-ring vlan | 23335

show protection-group ethernet-ring flush-info

IN THIS SECTION

- Syntax | 23313

- Description | 23313
- Options | 23313
- Required Privilege Level | 23313
- Output Fields | 23313
- Sample Output | 23314
- Release Information | 23314

Syntax

```
show protection-group ethernet-ring flush-info
```

Description

Display information about flush ports in an Ethernet ring.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1026 on page 23313 lists the output fields for the `show protection-group ethernet-ring flush-info` command. Output fields are listed in the approximate order in which they appear.

Table 1026: show protection-group ethernet-ring flush-info Output Fields

Field Name	Field Description
`Interface`	Physical interface configured for the Ethernet ring. This can be an aggregated Ethernet link also.

Table 1026: show protection-group ethernet-ring flush-info Output Fields *(Continued)*

Field Name	Field Description
`Originating Node`	Node from which RAPS protocol data units originates on the Ethernet Ring.
`Blocked Port Reference`	Reference of the ring port on which traffic is blocked.

Sample Output

show protection-group ethernet-ring flush-info (ACX and MX Series Routers)

```
user@host> show protection-group ethernet-ring flush-info
Ethernet ring flush port information for protection group pg100

Interface     Originating Node   Blocked Port Reference
xe-5/0/2.4001 00:00:00:00:00:00   0
xe-2/2/0.4001 00:00:00:00:00:00   0
```

show protection-group ethernet-ring flush-info detail (ACX and MX Series Routers)

```
user@host> show protection-group ethernet-ring flush-info detail
Ethernet ring flush port information for protection group pg100

Interface name                 : xe-5/0/2.4001
Originating Node               : 00:00:00:00:00:00
Blocked Port Reference         : 0

Interface name                 : xe-2/2/0.4001
Originating Node               : 00:00:00:00:00:00
Blocked Port Reference         : 0
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

show protection-group ethernet-ring data-channel | 23309
show protection-group ethernet-ring aps | 23294
show protection-group ethernet-ring node-state | 23320
show protection-group ethernet-ring statistics | 23328
show protection-group ethernet-ring vlan | 23335

show protection-group ethernet-ring interface

IN THIS SECTION

- Syntax | 23315
- Description | 23315
- Options | 23315
- Required Privilege Level | 23316
- Output Fields | 23316
- Sample Output | 23317
- Release Information | 23320

Syntax

```
show protection-group ethernet-ring interface
```

Description

Displays the status of the Automatic Protection Switching (APS) interfaces on an Ethernet ring.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1027 on page 23316 lists the output fields for both the EX Series switch, and the ACX Series and MX Series router `show protection-group ethernet-ring interface` commands. Output fields are listed in the approximate order in which they appear.

Table 1027: MX Series Routers show protection-group ethernet-ring interface Output Fields

Field Name	Field Description
Ethernet ring port parameters for protection group *group-name*	Output is organized by configured protection group.
Interface	Physical interfaces configured for the Ethernet ring. This can be an aggregated Ethernet link also.
Control Channel	(MX Series router only) Logical unit configured on the physical interface.
Direction	Direction of the traffic.
Forward State	State of the ring forwarding on the interface: **discarding** or **forwarding**.
Ring Protection Link End	Whether this interface is the end of the ring: **Yes** or **No**.
Signal Failure	Whether there a signal failure exists on the link: **Clear** or **Set**.
Admin State	State of the interface: For EX switches, **ready**, **ifl ready**, or **waiting**. For MX routers, **IFF ready** or **IFF disabled**.

Sample Output

show protection-group ethernet-ring interface (EX Series Switch Owner Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface      Forward State  RPL End  Signal Failure  Admin State

ge-0/0/3.0     discarding     Yes        Clear          ready
ge-0/0/9.0     forwarding     No         Clear          ready
```

show protection-group ethernet-ring interface (Owner Node MX Series Router)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface  Control Channel  Direction  Forward State  RPL End  SF     Admin State
ge-1/2/0   ge-1/2/0.100     east       forwarding     No       Clear  IFF ready
ge-1/2/2   ge-1/2/2.100     west       forwarding     No       Clear  IFF ready
```

show protection-group ethernet-ring interface detail (Owner Node MX Series Router)

```
user@host> show protection-group ethernet-ring interface detail
Ethernet ring port parameters for protection group pg101

Interface name                 : ge-1/2/0
Control channel name           : ge-1/2/0.100
Interface direction            : east
Ring Protection Link End       : No
Signal Failure                 : Clear
Forward State                  : forwarding
Interface Admin State          : IFF ready

Interface name                 : ge-1/2/2
Control channel name           : ge-1/2/2.100
Interface direction            : west
```

```
Ring Protection Link End       : No
Signal Failure                 : Clear
Forward State                  : forwarding
Interface Admin State          : IFF ready
```

show protection-group ethernet-ring interface (EX Series Switch Ring Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Ethernet ring port parameters for protection group pg101

Interface      Forward State  RPL End  Signal Failure  Admin State

ge-0/0/3.0     discarding     Yes        Clear          ready
ge-0/0/9.0     forwarding     No         Clear          ready
```

show protection-group ethernet-ring interface detail (ACX Series and MX Series)

```
user@host> show protection-group ethernet-ring interface detail
Ethernet ring port parameters for protection group Erp_1

Interface name                 : xe-0/0/0
Control channel name           : xe-0/0/0.1
Interface direction            : east
Ring Protection Link End       : No
Signal Failure                 : Clear
Forward State                  : forwarding
Interface Admin State          : IFF ready

Interface name                 : et-0/0/48
Control channel name           : et-0/0/48.1
Interface direction            : west
Ring Protection Link End       : No
Signal Failure                 : Clear
Forward State                  : forwarding
Interface Admin State          : IFF ready
```

show protection-group ethernet-ring interface detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring interface detail
Ethernet ring port parameters for protection group pg1001

Interface name                 : ge-0/0/14
Control channel name           : ge-0/0/14.0
Interface direction            : east
Ring Protection Link End       : No
Signal Failure                 : Clear
Forward State                  : forwarding
Interface Admin State          : IFF ready

Interface name                 : ge-0/0/18
Control channel name           : ge-0/0/18.0
Interface direction            : west
Ring Protection Link End       : No
Signal Failure                 : Clear
Forward State                  : forwarding
Interface Admin State          : IFF ready
```

show protection-group ethernet-ring interface detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring interface detail
Ethernet ring port parameters for protection group pg1001

Interface name                 : ge-0/0/42
Control channel name           : ge-0/0/42.0
Interface direction            : east
Ring Protection Link End       : Yes
Signal Failure                 : Clear
Forward State                  : discarding
Interface Admin State          : IFF ready

Interface name                 : ge-0/0/38
Control channel name           : ge-0/0/38.0
Interface direction            : west
Ring Protection Link End       : No
Signal Failure                 : Clear
```

```
Forward State                  : forwarding
Interface Admin State          : IFF ready
```

Release Information

Command introduced in Junos OS Release 9.4.

Command introduced in Junos OS Release 12.3X54 for ACX Series routers.

RELATED DOCUMENTATION

- show protection-group ethernet-ring data-channel | 23309
- show protection-group ethernet-ring aps | 23294
- show protection-group ethernet-ring node-state | 23320
- show protection-group ethernet-ring statistics | 23328
- show protection-group ethernet-ring vlan | 23335

show protection-group ethernet-ring node-state

IN THIS SECTION

- Syntax | 23321
- Description | 23321
- Options | 23321
- Required Privilege Level | 23321
- Output Fields | 23321
- Sample Output | 23324
- Release Information | 23327

Syntax

```
show protection-group ethernet-ring node-state
```

Description

Display the status of the Automatic Protection Switching (APS) nodes on an Ethernet ring.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1028 on page 23321 lists the output fields for the `show protection-group ethernet-ring node-state` command. Output fields are listed in the approximate order in which they appear.

Table 1028: show protection-group ethernet-ring node-state Output Fields

Field Name	Field Description
Ring Name/Ethernet Ring	Name configured for the Ethernet ring.

Table 1028: show protection-group ethernet-ring node-state Output Fields *(Continued)*

Field Name	Field Description
APS State	State of the Ethernet ring APS. • **idle**—Indicates that the ring is working in normal condition and there is no active or pending protection-switching request in the ring. When the ring is in idle state, it is blocked at the RPL link. • **protected**—Indicates that there is a protection switch on the ring because of a signal failure condition on the ring link. • **MS**—Indicates that the manual switch command is active in the ring. • **FS**—Indicates that the forced switch command is active in the ring. • **pending**—Indicates that the ring is in pending state.

Table 1028: show protection-group ethernet-ring node-state Output Fields *(Continued)*

Field Name	Field Description
Event	Events on the ring. • **NR-RB**—Indicates that there is no APS request and the ring link is blocked on the ring owner node. • **NR**—Indicates that there is no APS request pending in the ring. • **local SF**—Indicates that there is signal failure on one or both of the ring links of the node. • **remote SF**—Indicates that there is signal failure on one or more ring links of any other node of the ring. • **local FS**—Indicates that there is a forced switched command active on one or both of the ring links of the node. • **remote FS**—Indicates that there is a forced switch command active on one or more ring links of any other node of the ring. • **local MS**—Indicates that there is a manual switch command active on one of the ring links of the node. • **remote MS**—Indicates that there is a manual switch command active on one or more ring links of any other node of the ring. • **WTR running**—Indicates that the wait to restore timer is running on the RPL owner. • **WTB running**—Indicates that the wait to block timer is running on the RPL owner.
RPL Owner / Ring Protection Link Owner	Whether this node is the ring owner: **Yes** or **No**.
WTR Timer / Restore Timer	Restoration timer: **running** or **disabled**.

Table 1028: show protection-group ethernet-ring node-state Output Fields *(Continued)*

Field Name	Field Description
WTB Timer / Wait to block timer	Wait to block timer: **running** or **disabled**. **NOTE**: The Wait To Block Timer (WTB) is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect.
Wait to block timer (WTB Timer)	Wait to block interval. **NOTE**: The Wait To Block Timer (WTB) is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect.
Guard Timer	Guard timer: **running** or **disabled**.
Op State / Operational State	State of the node: **Operational** or **any internal wait state.**.

Sample Output

show protection-group ethernet-ring node-state (MX Series Router - RPL Owner Node, Normal Operation)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring   APS State  Event       RPL Owner  WTR Timer  WTB Timer  Guard Timer  Operation
state
pg101           idle       NR-RB       Yes        disabled   disabled   disabled
operational
pg102            idle      NR-RB       No         disabled   disabled   disabled     operational
```

show protection-group ethernet-ring node-state (MX Series Router - Normal Ring Node, Normal Operation)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event         RPL Owner
pg102            idle         NR-RB         No

WTR Timer        WTB Timer   Guard Timer  Operation state
disabled   disabled             disabled            operational
```

show protection-group ethernet-ring node-state (MX Series Router - RPL Owner Node, Remote Failure Condition)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event         RPL Owner
pg101            protected    remote SF         Yes

WTR Timer      WTB Timer    Guard Timer  Operation state
disabled       disabled         disabled        operational
```

show protection-group ethernet-ring node-state detail (ACX Series and MX Series Router)

```
user@host> show protection-group ethernet-ring node-state detail
Ethernet-Ring name             : Erp_1
APS State                      : idle
Event                          : NR-RB
Ring Protection Link Owner     : No
Wait to Restore Timer          : disabled
Wait to Block Timer            : disabled
Guard Timer                    : disabled
Operation state                : operational
```

show protection-group ethernet-ring node-state detail (MX Series Router - RPL Owner Node, Normal Operation)

```
user@host> show protection-group ethernet-ring node-state detail
Ethernet-Ring name             : pg101
```

```
APS State                       : idle
Event                           : NR-RB
Ring Protection Link Owner      : Yes
Wait to Restore Timer           : disabled
Wait to Block Timer             : disabled
Guard Timer                     : disabled
Operation state                 : operational

Ethernet-Ring name              : pg102
APS State                       : idle
Event                           : NR-RB
Ring Protection Link Owner      : No
Wait to Restore Timer           : disabled
Wait to Block Timer             : disabled
Guard Timer                     : disabled
Operation state                 : operational
```

show protection-group ethernet-ring node-state detail (MX Series Router with WTR Timer)

```
user@host> show protection-group ethernet-ring node-state detail
Ethernet-Ring name              : pg_major
APS State                       : pending
Event                           : WTR running
Ring Protection Link Owner      : Yes
Wait to Restore Timer           : running (time to expire: 269 sec)
Wait to Block Timer             : disabled
Guard Timer                     : disabled
Operation state                 : operational

Ethernet-Ring name              : pg_subring
APS State                       : pending
Event                           : NR
Ring Protection Link Owner      : No
Wait to Restore Timer           : disabled
Wait to Block Timer             : disabled
Guard Timer                     : disabled
Operation state                 : operational
```

show protection-group ethernet-ring node-state detail (MX Series Router with WTB Timer)

```
user@host> show protection-group ethernet-ring node-state detail
Ethernet-Ring name              : Pg-2
APS State                       : pending
Event                           : WTB running
Ring Protection Link Owner      : Yes
Wait to Restore Timer           : disabled
Wait to Block Timer             : running (time to expire: 2 sec)
Guard Timer                     : disabled
Operation state                 : operational
```

show protection-group ethernet-ring node-state detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring node-state detail
Ethernet-Ring name              : pg1001
APS State                       : idle
Event                           : NR-RB
Ring Protection Link Owner      : Yes
Wait to Restore Timer           : disabled
Wait to Block Timer             : disabled    <-field not supported. Always disabled.
Guard Timer                     : disabled
Operation state                 : operational
```

Release Information

Command introduced in Junos OS Release 9.4.

Command introduced in Junos OS Release 12.3X54 for ACX Series routers.

RELATED DOCUMENTATION

show protection-group ethernet-ring data-channel | 23309

show protection-group ethernet-ring aps | 23294

show protection-group ethernet-ring interface | 23315

show protection-group ethernet-ring statistics | 23328

show protection-group ethernet-ring vlan | 23335

show protection-group ethernet-ring statistics

IN THIS SECTION

- Syntax | 23328
- Description | 23328
- Options | 23328
- Required Privilege Level | 23328
- Output Fields | 23329
- Sample Output | 23331
- Release Information | 23335

Syntax

```
show protection-group ethernet-ring statistics group-name group-name
```

<brief | detail>

Description

Display statistics regarding Automatic Protection Switching (APS) protection groups on an Ethernet ring.

Options

group-name Display statistics for the protection group. If you omit this option, protection group statistics for all configured groups are displayed.

`brief`—Display brief statistics for the protection group.

`detail`—Display detailed statistics for the protection group.

Required Privilege Level

view

Output Fields

Table 1029 on page 23329 lists the output fields for the `show protection-group ethernet-ring statistics` command.

Table 1029: show protection-group ethernet-ring statistics Output Fields

Field Name	Field Description
Ethernet Ring Statistics for PG	Name of the protection group for which statistics are displayed.
RAPS event sent	Number of times Ring Automatic Protection Switching (RAPS) message transmission event occurred locally. This field is applicable only to MX Series routers.
RAPS event received	Number of RAPS messages received and processed by ERP state-machine and which resulted in state transition. This field is applicable only to MX Series routers.
Local SF	Number of times a signal failure has occurred locally.
Remote SF	Number of times a signal failure has occurred anywhere else on the ring.
NR event	Number of times a No Request event has occurred on the ring. This field is applicable only to EX Series switches.
`NR event sent`	Number of times a No Request event has occurred locally. This field is applicable only to MX Series routers.
`NR event received`	Number of times a No Request event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.
NR-RB event	Number of times a No Request, Ring Blocked event has occurred on the ring. This field is applicable only to EX Series switches.
`NR-RB event sent`	Number of times a No Request, Ring Blocked event has occurred locally. This field is applicable only to MX Series routers.

Table 1029: show protection-group ethernet-ring statistics Output Fields *(Continued)*

Field Name	Field Description
`NR-RB event received`	Number of times a No Request, Ring Blocked event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.
`Flush event sent`	Number of times flush-event RAPS message transmission event occurred locally. This field is applicable only to MX Series routers.
`Flush event received`	Number of flush-event RAPS messages received and processed by the ring instance control process. This field is applicable only to MX Series routers.
`Local FS event sent`	Number of times a forced switch event has occurred locally. This field is applicable only to MX Series routers.
`Remote FS event received`	Number of times a forced switch event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.
`Local MS event sent`	Number of times a manual switch event has occurred locally. This field is applicable only to MX Series routers.
`Remote MS event received`	Number of times a manual switch event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.

Table 1030 on page 23330 lists the output fields for the `show protection-group ethernet-ring statistics` command when the `detail` option is used. These fields are valid only for MX Series routers.

Table 1030: show protection-group ethernet-ring statistics detail Output Fields (for MX Series Routers)

Field Name	Field Description
`Total number of FDB flush`	Number of times forwarding database (FDB) flush has happened for the ring instance.
`Flush-logic triggered flush`	Number of times FDB flush has happened because of flush-logic based on node ID and Blocked Port Reference (BPR).

Table 1030: show protection-group ethernet-ring statistics detail Output Fields (for MX Series Routers) *(Continued)*

Field Name	Field Description
`Remote RAPS PDU received`	Number of valid RAPS PDU messages received. This counter counts only RAPS messages generated by other devices on the ring.
`Remote RAPS dropped due to guard-timer`	Number of RAPS messages dropped by the device because the guard timer is running.
`Invalid remote RAPS PDU dropped`	Number of RAPS messages dropped by the device because the messages are invalid.
`RAPS dropped due to miscellaneous errors`	Number of RAPS messages dropped because of any other reason. For example, messages dropped because of unsupported functionality.
`Local received RAPS PDU dropped`	Number of self-generated RAPS messages received and dropped.

Sample Output

show protection-group ethernet-ring statistics (EX Series Switch)

```
user@switch> show protection-group ethernet-ring statistics
Ring Name Local SF Remote SF NR Event NR-RB Event
erp1       2       1       2        3
```

show protection-group ethernet-ring statistics (MX Series Router)

```
user@host> show protection-group ethernet-ring statistics
Ethernet Ring statistics for PG Pg-1
RAPS event sent                  : 1
RAPS event received              : 1152
Local SF happened:               : 0
Remote SF happened:              : 428
```

```
NR event sent:                    : 1
NR event received:                : 133
NR-RB event sent:                 : 0
NR-RB event received:             : 591
Flush event sent                  : 0
Flush event received:             : 0
Local FS event sent:              : 0
Remote FS event received:         : 0
Local MS event sent:              : 0
Remote MS event received:         : 0
```

show protection-group ethernet-ring statistics detail (Specific Group)(MX Series Router)

```
user@host> show protection-group ethernet-ring statistics detail
Ethernet Ring statistics for PG Pg-1
RAPS event sent                          : 1
RAPS event received                      : 0
Local SF happened                        : 0
Remote SF happened                       : 0
NR event sent                            : 1
NR event received                        : 0
NR-RB event sent                         : 0
NR-RB event received                     : 0
Flush event sent                         : 0
Flush event received                     : 0
Local FS event sent                      : 0
Remote FS event received                 : 0
Local MS event sent                      : 0
Remote MS event received                 : 0
Total number of FDB flush                : 0
Flush-logic triggered flush              : 0
Remote raps PDU received                 : 0
Remote raps dropped due to guard-timer   : 0
Invalid remote raps PDU dropped          : 0
Raps dropped due to miscellaneous errors : 0
Local received raps PDU dropped          : 0
```

show protection-group ethernet-ring statistics (Owner Node, Failure Condition on ACX and MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                        : 1
RAPS received                    : 0
Local SF happened:               : 0
Remote SF happened:              : 0
NR event happened:               : 0
NR-RB event happened:            : 1
NR event sent:                   : 0
NR event received:               : 0
NR-RB event sent:                : 1
NR-RB event received:            : 0
Flush event sent                 : 0
Flush event received:            : 0
Local FS event sent:             : 0
Remote FS event received:        : 0
Local MS event sent:             : 0
Remote MS event received:        : 0
```

show protection-group ethernet-ring statisitics (Ring Node, Failure Condition on ACX and MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent                        : 1
RAPS received                    : 0
Local SF happened:               : 0
Remote SF happened:              : 0
NR event happened:               : 0
NR-RB event happened:            : 1
NR event sent:                   : 0
NR event received:               : 0
NR-RB event sent:                : 1
NR-RB event received:            : 0
Flush event sent                 : 0
Flush event received:            : 0
Local FS event sent:             : 0
```

```
Remote FS event received:        : 0
Local MS event sent:             : 0
Remote MS event received:        : 0
```

show protection-group ethernet-ring statistics detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring statistics detail
Ethernet Ring statistics for PG pg1001
RAPS event sent                          : 1
RAPS event received                      : 1
Local SF happened                        : 0
Remote SF happened                       : 0
NR event sent                            : 1
NR event received                        : 0
NR-RB event sent                         : 0
NR-RB event received                     : 1
Flush event sent                         : 0
Flush event received                     : 0
Local FS event sent                      : 0
Remote FS event received                 : 0
Local MS event sent                      : 0
Remote MS event received                 : 0
Total number of FDB flush                : 0
Flush-logic triggered flush              : 0
Remote raps PDU received                 : 145
Remote raps dropped due to guard-timer   : 0
Invalid remote raps PDU dropped          : 0
Raps dropped due to miscellaneous errors : 0
Local received raps PDU dropped          : 0
```

show protection-group ethernet-ring statistics detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring statistics detail
Ethernet Ring statistics for PG pg1001
RAPS event sent                          : 2
RAPS event received                      : 0
Local SF happened                        : 0
Remote SF happened                       : 0
NR event sent                            : 1
NR event received                        : 0
```

```
NR-RB event sent                        : 1
NR-RB event received                    : 0
Flush event sent                        : 0
Flush event received                    : 0
Total number of FDB flush               : 0
Remote raps PDU received                : 211
Remote raps dropped due to guard-timer  : 0
Invalid remote raps PDU dropped         : 0
Raps dropped due to miscellaneous errors : 0
Local received raps PDU dropped         : 91
```

Release Information

Command introduced in Junos OS Release 9.4.

Command introduced in Junos OS Release 12.3X54 for ACX Series routers.

RELATED DOCUMENTATION

- show protection-group ethernet-ring data-channel | 23309
- show protection-group ethernet-ring aps | 23294
- show protection-group ethernet-ring node-state | 23320
- show protection-group ethernet-ring interface | 23315
- show protection-group ethernet-ring vlan | 23335

show protection-group ethernet-ring vlan

IN THIS SECTION

- Syntax | 23336
- Description | 23336
- Options | 23336
- Required Privilege Level | 23336
- Output Fields | 23336

Sample Output | 23337
Release Information | 23341

Syntax

```
show protection-group ethernet-ring vlan
<brief | detail>
<group-name group-name>
```

Description

On MX Series routers, display all data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.

Options

brief | detail (Optional) Display the specified level of output.

group-name (Optional) Protection group for which to display details such as data channel interfaces, vlan, and bridge-domain. If you omit this optional field, details for all configured protection groups will be displayed.

Required Privilege Level

view

Output Fields

Table 1031 on page 23337 lists the output fields for the `show protection-group ethernet-ring vlan` command. Output fields are listed in the approximate order in which they appear.

Table 1031: show protection-group ethernet-ring vlan Output Fields

Field Name	Field Description
Interface	Name of the interface configured for the Ethernet protection ring.
Vlan	Name of the VLAN associated with the interface configured for the Ethernet protection ring.
STP index	The Spanning Tree Protocol (STP) index number used by each interface in an Ethernet ring. The STP index controls the forwarding behavior for a set of VLANs on a data channel on an Ethernet ring port. For multiple Ethernet ring instances on an physical ring port, there are multiple STP index numbers. Different ring instances will have different STP index numbers and may have different forwarding behavior.
Bridge Domain	Name of the bridge domain that is associated with the VLAN configured for the Ethernet protection ring.

Sample Output

show protection-group ethernet-ring vlan

```
user@host> show protection-group ethernet-ring vlan
Ethernet ring IFBD parameters for protection group vkm01

Interface    Vlan    STP Index   Bridge Domain
xe-5/0/2     1       78          default-switch/bd1
xe-2/2/0     1       79          default-switch/bd1
xe-5/0/2     2       78          default-switch/bd2
xe-2/2/0     2       79          default-switch/bd2
xe-5/0/2     3       78          default-switch/bd3
xe-2/2/0     3       79          default-switch/bd3
xe-5/0/2     4       78          default-switch/bd4
xe-2/2/0     4       79          default-switch/bd4
xe-5/0/2     5       78          default-switch/bd5
xe-2/2/0     5       79          default-switch/bd5
xe-5/0/2     6       78          default-switch/bd6
xe-2/2/0     6       79          default-switch/bd6
```

```
xe-5/0/2     7       78          default-switch/bd7
xe-2/2/0     7       79          default-switch/bd7
xe-5/0/2     8       78          default-switch/bd8
xe-2/2/0     8       79          default-switch/bd8
xe-5/0/2     9       78          default-switch/bd9
xe-2/2/0     9       79          default-switch/bd9
xe-5/0/2     10      78          default-switch/bd10
xe-2/2/0     10      79          default-switch/bd10
xe-5/0/2     11      78          default-switch/bd11
xe-2/2/0     11      79          default-switch/bd11
xe-5/0/2     12      78          default-switch/bd12
xe-2/2/0     12      79          default-switch/bd12
xe-5/0/2     13      78          default-switch/bd13
xe-2/2/0     13      79          default-switch/bd13
xe-5/0/2     14      78          default-switch/bd14
xe-2/2/0     14      79          default-switch/bd14
xe-5/0/2     15      78          default-switch/bd15
xe-2/2/0     15      79          default-switch/bd15
```

show protection-group ethernet-ring vlan brief

```
user@host> show protection-group ethernet-ring vlan brief
Ethernet ring IFBD parameters for protection group vkm01

Interface    Vlan    STP Index   Bridge Domain
xe-5/0/2     1       78          default-switch/bd1
xe-2/2/0     1       79          default-switch/bd1
xe-5/0/2     2       78          default-switch/bd2
xe-2/2/0     2       79          default-switch/bd2
xe-5/0/2     3       78          default-switch/bd3
xe-2/2/0     3       79          default-switch/bd3
xe-5/0/2     4       78          default-switch/bd4
xe-2/2/0     4       79          default-switch/bd4
xe-5/0/2     5       78          default-switch/bd5
xe-2/2/0     5       79          default-switch/bd5
xe-5/0/2     6       78          default-switch/bd6
xe-2/2/0     6       79          default-switch/bd6
xe-5/0/2     7       78          default-switch/bd7
xe-2/2/0     7       79          default-switch/bd7
xe-5/0/2     8       78          default-switch/bd8
xe-2/2/0     8       79          default-switch/bd8
```

```
xe-5/0/2     9       78          default-switch/bd9
xe-2/2/0     9       79          default-switch/bd9
xe-5/0/2     10      78          default-switch/bd10
xe-2/2/0     10      79          default-switch/bd10
xe-5/0/2     11      78          default-switch/bd11
xe-2/2/0     11      79          default-switch/bd11
xe-5/0/2     12      78          default-switch/bd12
xe-2/2/0     12      79          default-switch/bd12
xe-5/0/2     13      78          default-switch/bd13
xe-2/2/0     13      79          default-switch/bd13
xe-5/0/2     14      78          default-switch/bd14
xe-2/2/0     14      79          default-switch/bd14
xe-5/0/2     15      78          default-switch/bd15
xe-2/2/0     15      79          default-switch/bd15
```

show protection-group ethernet-ring vlan detail

```
user@host> show protection-group ethernet-ring vlan detail
Ethernet ring IFBD parameters for protection group vkm01

Interface name                 : xe-5/0/2
Vlan                           : 1
STP index                      : 78
Bridge Domain                  : default-switch/bd1

Interface name                 : xe-2/2/0
Vlan                           : 1
STP index                      : 79
Bridge Domain                  : default-switch/bd1

Interface name                 : xe-5/0/2
Vlan                           : 2
STP index                      : 78
Bridge Domain                  : default-switch/bd2

Interface name                 : xe-2/2/0
Vlan                           : 2
STP index                      : 79
Bridge Domain                  : default-switch/bd2

Interface name                 : xe-5/0/2
```

```
Vlan                            : 3
STP index                       : 78
Bridge Domain                   : default-switch/bd3
```

show protection-group ethernet-ring vlan group-name vkm01

```
user@host> show protection-group ethernet-ring vlan vkm01

Ethernet ring IFBD parameters for protection group vkm01

Interface    Vlan    STP Index   Bridge Domain
xe-5/0/2     16      80          default-switch/bd16
xe-2/2/0     16      81          default-switch/bd16
xe-5/0/2     17      80          default-switch/bd17
xe-2/2/0     17      81          default-switch/bd17
xe-5/0/2     18      80          default-switch/bd18
xe-2/2/0     18      81          default-switch/bd18
xe-5/0/2     19      80          default-switch/bd19
xe-2/2/0     19      81          default-switch/bd19
xe-5/0/2     20      80          default-switch/bd20
xe-2/2/0     20      81          default-switch/bd20
xe-5/0/2     21      80          default-switch/bd21
xe-2/2/0     21      81          default-switch/bd21
xe-5/0/2     22      80          default-switch/bd22
xe-2/2/0     22      81          default-switch/bd22
xe-5/0/2     23      80          default-switch/bd23
xe-2/2/0     23      81          default-switch/bd23
xe-5/0/2     24      80          default-switch/bd24
xe-2/2/0     24      81          default-switch/bd24
xe-5/0/2     25      80          default-switch/bd25
xe-2/2/0     25      81          default-switch/bd25
xe-5/0/2     26      80          default-switch/bd26
xe-2/2/0     26      81          default-switch/bd26
xe-5/0/2     27      80          default-switch/bd27
xe-2/2/0     27      81          default-switch/bd27
xe-5/0/2     28      80          default-switch/bd28
xe-2/2/0     28      81          default-switch/bd28
xe-5/0/2     29      80          default-switch/bd29
xe-2/2/0     29      81          default-switch/bd29
```

```
xe-5/0/2     30      80          default-switch/bd30
xe-2/2/0     30      81          default-switch/bd30
```

show protection-group ethernet-ring vlan detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring vlan detail
Ethernet ring IFBD parameters for protection group pg1001

Interface name                  : ge-0/0/42
Vlan                            : 2001
STP index                       : 52
Bridge Domain                   : default-switch/vlan2001

Interface name                  : ge-0/0/38
Vlan                            : 2001
STP index                       : 53
Bridge Domain                   : default-switch/vlan2001
```

Release Information

Command introduced in Junos OS Release 10.2.

Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

RELATED DOCUMENTATION

- show protection-group ethernet-ring aps | 23294
- show protection-group ethernet-ring data-channel | 23309
- show protection-group ethernet-ring interface | 23315
- show protection-group ethernet-ring node-state | 23320
- show protection-group ethernet-ring statistics | 23328

show ptp aggregated-ethernet interfaces

IN THIS SECTION

- Syntax | 23342
- Description | 23342
- Options | 23342
- Required Privilege Level | 23342
- Output Fields | 23342
- Sample Output | 23343
- Release Information | 23344

Syntax

```
show ptp aggregated-ethernet interfaces
```

Description

Display the information related to aggregated Ethernet bundles.

Options

This command has no options.

Required Privilege Level

View

Output Fields

Table 1032 on page 23343 lists the output fields for the `show ptp aggregated-ethernet interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 1032: show ptp aggregated-ethernet interfaces

Field Name	Field Description
`Bundle`	Name of the aggregated Ethernet bundle.
`Primary`	Primary interface in an aggregated Ethernet bundle. • `active`—Indicates that the link is actively receiving ptp packets. • `down`—Indicates that the link is down.
`Secondary`	Secondary interface in an aggregated Ethernet bundle. • `active`—Indicates that the link is actively receiving ptp packets. • `ready`—Indicates that the link is ready to take over if the primary link fails.
`Role`	Role of the aggregated Ethernet bundle: • `Master`, `Slave`, or `Stateful`

Sample Output

show ptp aggregated-ethernet interfaces

```
user@host> show ptp aggregated-ethernet interfaces

bundle       primary                    secondary                  role
ae0.0        ge-1/0/1 (active)        ge-5/0/5 (ready)          slave
ae2.0        ge-5/0/1 (down)          ge-6/1/1 (active)   master
```

show ptp aggregated-ethernet interfaces (stateful ports)

```
user@host> show ptp aggregated-ethernet interfaces

AE Bundle         Ifstate  Primary                Secondary              Role
ae0.0               Up       et-11/0/5(active)      et-11/1/2(secondary)   Stateful
```

Release Information

Command introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Precision Time Protocol Overview

show ptp all master clock

IN THIS SECTION

- Syntax | 23344
- Description | 23344
- Options | 23344
- Required Privilege Level | 23345
- Output Fields | 23345
- Sample Output | 23348

Syntax

```
show ptp all-master-clock
```

Description

(MX240, MX480, MX960, MX2010, MX2020, MX2008, MX304, PTX10004, PTX10008, ACX7100-32, ACX7100-48, ACX7024, ACX7509 routers) Display the details of the timeTransmitter clock configured on the node.

Options

The command has no options.

Required Privilege Level

view

Output Fields

The Output Fields on page 23345 lists the output fields for the `show ptp all-master-clock` command. Output fields are listed in the approximate order in which they appear.

Table 1033: show ptp all master clock Output Fields

Field Name	Field Description
`Slot`	Number of the FPC or MIC slot. If underlying interface is not available momentarily or is unavailable for a virtual interface, the slot number is 255.
`Parent Id`	EUI-64 clock identifier of the immediate upstream timeTransmitter clock.
`GMC Id`	EUI-64 clock identifier of the grandmaster clock.
`GMC Class`	Denotes the grandmaster clock's traceability of the distributed time or frequency.
`GMC Accuracy`	Indicates the expected accuracy of the grandmaster clock.
`GMC Variance`	Represents an estimate of the variations of the grandmaster clock.
`GMC Priority 1`	Priority1-value of the grandmaster clock.
`GMC Priority 2`	Priority2-value of the grandmaster clock.

Table 1033: show ptp all master clock Output Fields *(Continued)*

Field Name	Field Description
UTC Offset	Offset between International Atomic Time (TAI) and Coordinated Universal Time (UTC) times. The value is 34 seconds as of January 2012.
Leap-59	When TRUE, the last minute of the current UTC day has only 59 seconds (instead of the 60 SI seconds).
Leap-61	When TRUE, the last minute of the current UTC day has 61 seconds (instead of the 60 SI seconds).
Time Traceable	When TRUE, the timescale and the UTC offset are traceable to a primary reference.
Freq Traceable	When TRUE, frequency determining the timescale is traceable to a primary reference.
Time Scale	The PTP timeTransmitter clock's best estimate of the international standard time scales.
Time Source	Time source external to the Precision Time Protocol (PTP), which provides time and/or frequency as appropriate. The time source is traceable to the international standards laboratories maintaining clocks that form the basis for the International Atomic Time (TAI) and Universal Coordinated Time (UTC) timescales. Examples of these are Global Positioning System (GPS), NTP, and National Institute of Standards and Technology (NIST) timeservers.

Table 1033: show ptp all master clock Output Fields *(Continued)*

Field Name	Field Description
`Unicast`	Configure upstream unicast PTP timeTransmitter clock sources.
`UTC Valid`	TRUE if the UTC offset is known to be correct.
`Two Step`	Whether the clock provides time information, which is a combination of an event message and a subsequent general message: True or False.
`Alternate Master`	Allows alternate timeTransmitters that are not the current best timeTransmitter to exchange PTP timing information with timeReceiver ports.
`Steps Removed`	Number of boundary clocks between the local clock and the foreign timeTransmitter clock.
`Parent SourcePort Id`	Source port Id of a timeTransmitter clock.
`ANN msg Interval`	Determines the mean time interval between successive announce messages. For example, announce Interval = 2 ^ Log Interval.
Interface	Interface on which announce message is received.
PTP Stream ID	Unique Local identifier index of PTP stream created.

Sample Output

show ptp all-master-clock

```
user@host> show ptp all-master-clock
Upstream Masters:

Slot                  : 1                          Parent Id        : 74:e7:98:ff:fe:60:8b:dd
GMC Id                : 74:e7:98:ff:fe:60:8b:dd    GMC Class        : 6
GMC Accuracy          : 33                         GMC Variance     : 20061
GMC Priority1         : 128                        GMC Priority2    : 128
UTC Offset            : 37                         Leap-59          : FALSE
Leap-61               : FALSE                      Time traceable   : TRUE
Freq Traceable        : TRUE                       Time Scale       : TRUE
Time source           : 32                         Unicast          : FALSE
UTC Valid             : TRUE                       Two Step         : FALSE
Alternate Master      : FALSE                      Steps Removed    : 0
Parent SourcePort Id  : 2                          ANN msg Interval : -3
Interface             : et-1/0/0:0.0               PTP Stream ID    : 4
```

show ptp clock

IN THIS SECTION

- Syntax | 23349
- Description | 23349
- Options | 23349
- Required Privilege Level | 23349
- Output Fields | 23349
- Sample Output | 23351
- Release Information | 23354

Syntax

```
show ptp clock
```

Description

(ACX Series, MX80, MX240, MX480, MX960, PTX10008 routers, and QFX Series switches) Display the details of the clock configured on the node.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1034 on page 23349 lists the output fields for the `show ptp clock` command. Output fields are listed in the approximate order in which they appear.

Table 1034: show ptp clock Output Fields

Field Name	Field Description
`Slot Number`	Number of the FPC or MIC slot.
`Two-step Clock`	Whether the clock provides time information which is a combination of an event message and a subsequent general message: `True` or `False`.
`Clock Identity`	Clock identity of the timeReceiver or timeReceiver as defined in IEEE 1588.
`Total Ports on Device`	Total number of PTP ports on the router.

Table 1034: show ptp clock Output Fields *(Continued)*

Field Name	Field Description
`Clock Class`	Attribute of an ordinary or boundary clock that denotes the traceability of the time or frequency distributed by the grandmaster clock.
`Clock Accuracy`	Indicates the expected accuracy of a clock when it is the grandmaster clock, or in the event it becomes the grandmaster.
`Log Variance`	Represents an estimate of the variations of the local clock when it is not synchronized via PTP to another clock.
`Clock Priority1`	Priority value of the clock. Lower value takes precedence.
`Clock Priority2`	Prioritize the primaries to avoid confusion when the `Clock Priority1` value is the same for different primaries in a network.
`UTC Offset`	Offset between International Atomic Time (TAI) and Coordinated Universal Time (UTC) times. The value is 34 seconds as of January 2012.
`Leap59`	When `TRUE`, the last minute of the current UTC day has only 59 seconds (instead of the 60 SI seconds).
`Leap61`	When `TRUE`, the last minute of the current UTC day has 61 seconds (instead of the 60 SI seconds).
`Time Traceable`	When `TRUE`, the timescale and the UTC offset are traceable to a timeTransmitter reference.
`Frequency Traceable`	When `TRUE`, frequency determining the timescale is traceable to a timeTransmitter reference.
`Time Source`	Time source external to the Precision Time Protocol (PTP), which provides time and/or frequency as appropriate. The time source is traceable to the international standards laboratories maintaining clocks that form the basis for the International Atomic Time (TAI) and Universal Coordinated Time (UTC) timescales. Examples of these are Global Positioning System (GPS), NTP, and National Institute of Standards and Technology (NIST) timeservers.

Table 1034: show ptp clock Output Fields *(Continued)*

Field Name	Field Description
Delay Req Sending Time	Interval in seconds between the delay-request messages sent by the timeReceiver to the timeTransmitter.
Steps Removed	Number of boundary clocks between the local clock and the foreign timeTransmitter clock.
Slave-only	Set to TRUE, when the system is used in ordinary timeReceiver clock mode; otherwise, FALSE.
Parent Id	EUI-64 clock identifier of the immediate upstream timeTransmitter clock.
GMC Id	EUI-64 clock identifier of the grandmaster clock.
GMC Class	Denotes the grandmaster clock's traceability of the distributed time or frequency.
GMC Accuracy	Indicates the expected accuracy of the grandmaster clock.
GMC Variance	Represents an estimate of the variations of the grandmaster clock.
GMC Priority1	Priority1-value of the grandmaster clock.
GMC Priority2	Priority2-value of the grandmaster clock.

Sample Output

show ptp clock

```
user@host> show ptp clock
Clock Details:

Slot Number           : 7
Default Data:
Two-step Clock        : FALSE                     Clock Identity : 00:05:85:ff:fe:73:ef:d0
```

```
Total Ports on Device : 0                          Clock Class    : 255
Clock Accuracy        : 49                         Log Variance   : -12944
Clock Priority1       : 128                        Clock Priority2: 128
UTC Offset            : 33                         Leap59         : FALSE
Leap61                : FALSE                      Time Traceable  : FALSE
Frequency Traceable   : FALSE                      Time master    : 0
Delay Req Sending Time: 0                          Steps Removed  : 1
Slave-only            : NA
Parent Data:
Parent Id             : 00:18:0b:ff:ff:20:01:62
GMC Id                : 00:18:0b:ff:ff:20:01:62   GMC Class        : 52
GMC Accuracy          : 254                       GMC Variance     : 11952
GMC Priority1         : 0                         GMC Priority2    : 0
Global Data:
UTC Offset            : 34                        Leap-59          : FALSE
Leap-61               : FALSE                     Time traceable    : FALSE
Freq Traceable        : FALSE                     Time Scale       : FALSE
Time master           : 160
```

show ptp clock (ACX Series Routers)

```
user@host> show ptp clock
Clock Details:
Slot Number           : 0
Default Data:
Two-step Clock        : FALSE                      Clock Identity : 84:18:88:ff:fe:c0:7a:00
Total Ports on Device : 0                          Clock Class    : 255
Clock Accuracy        : 34                         Log Variance   : 15353
Clock Priority1       : 128                        Clock Priority2: 128
UTC Offset            : 0                          Leap59         : FALSE
Leap61                : FALSE                      Time Traceable : FALSE
Frequency Traceable   : FALSE                      Time Source    : 0
Delay Req Sending Time: 0                          Steps Removed  : 0
Slave-only            : NA
Parent Data:
Parent Id             : 00:00:64:ff:fe:01:01:02
GMC Id                : 00:00:64:ff:fe:01:01:02   GMC Class        : 80
GMC Accuracy          : 35                        GMC Variance     : 0
GMC Priority1         : 128                       GMC Priority2    : 128
Global Data:
UTC Offset            : 0                         Leap-59          : FALSE
```

```
Leap-61                : FALSE                      Time tracable    : FALSE
Freq Traceable         : FALSE                      Time Scale       : FALSE
Time source            : 16
```

show ptp clock (QFX5120-48T)

```
user@host> show ptp clock
Clock Details:
Slot Number            : 0
Default Data:
Two-step Clock         : FALSE                       Clock Identity : 3c:8c:93:ff:fe:b4:c4:4d
Total Ports on Device : 0                            Clock Class    : 255
Clock Accuracy         : 254                         Log Variance   : 65535
Clock Priority1        : 128                         Clock Priority2: 255
UTC Offset             : 37                          Leap59         : FALSE
Leap61                 : FALSE                       Time Traceable : FALSE
Frequency Traceable    : FALSE                       Time Source    : 160
Delay Req Sending Time: 0                            Steps Removed  : 2
Slave-only             : TRUE
Parent Id              : 3c:8c:93:ff:fe:b5:4c:4c
GMC Id                 : 7c:e2:ca:ff:fe:a2:2b:d0   GMC Class        : 248
GMC Accuracy           : 254                        GMC Variance     : 65535
GMC Priority1          : 128                        GMC Priority2    : 40
Global Data:
UTC Offset             : 37                         Leap-59          : FALSE
Leap-61                : FALSE                      Time tracable    : FALSE
Freq Traceable         : FALSE                      Time Scale       : TRUE
Time source            : 160                        Path Trace count : 0
```

show ptp clock detail (PTX10008)

```
user@host> show ptp clock detail
Clock Details:
Slot Number            : 0 (CB)
Default Data:
Two-step Clock         : FALSE                       Clock Identity : e4:5d:37:ff:fe:4e:d7:76
Total Ports on Device : 128                          Clock Class    : 248
Clock Accuracy         : 254                         Log Variance   : 65535
Clock Priority1        : 128                         Clock Priority2: 128
UTC Offset             : 37                          Leap59         : FALSE
```

```
Leap61                : FALSE                      Time Traceable : FALSE
Frequency Traceable   : FALSE                      Time Source    : 160
Delay Req Sending Time: 0                          Steps Removed  : 0
Slave-only            : FALSE
Parent Data:
Parent Id             : 2c:21:72:ff:fe:d8:07:d0
GMC Id                : 00:00:00:00:00:00:00:01   GMC Class        : 6 (Tx: 248)
GMC Accuracy          : 33 (Tx: 254)              GMC Variance     : 65535 (Tx: 65535)
GMC Priority1         : 128                       GMC Priority2    : 128 (Tx: 128)
Global Data:
UTC Offset            : 37                        Leap-59          : FALSE
Leap-61               : FALSE                     Time traceable   : TRUE
Freq Traceable        : TRUE                      Time Scale       : TRUE
Time source           : 160                       Path Trace count : 0
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

IEEE 1588v2 PTP Boundary Clock Overview

IEEE 1588v2 Precision Timing Protocol (PTP)

Precision Time Protocol Overview

show ptp global-information

IN THIS SECTION

- Syntax | 23355
- Description | 23355
- Options | 23355
- Required Privilege Level | 23355
- Output Fields | 23355
- Sample Output | 23359

Release Information | 23368

Syntax

```
show ptp global-information
```

Description

Show Precision Time Protocol (PTP)–related global information.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1035 on page 23355 lists the output fields for the `show ptp global-information` command. Output fields are listed in the approximate order in which they appear.

Table 1035: show ptp global-information Output Fields

Field Name	Field Description
`PTP Global Configuration`	Displays if PTP is configured globally.
`Domain Number`	PTP domain with values from 0 through 127. The default value is 0. Only one PTP domain is supported at any given point in time.
`Clock mode`	Clock mode is either boundary or ordinary.

Table 1035: show ptp global-information Output Fields *(Continued)*

Field Name	Field Description
`Profile Type`	IEEE-2008, Enterprise, G.8275.1, G.8275.2.enh, SMPTE-ST 2059-2, AES67, and Combined AES67 and SMPTE-ST 2059-2.
`Servo Type`	Indicates whether the servo type applied on the platform is the default or the enhanced version.
`Priority Level1`	Priority value of the clock: 0 through 255. The default is 128. The lower value takes precedence.
`Priority Level2`	Priority value of the clock: 0 through 255. The default is 128. This value is used to differentiate and prioritize the timeTransmitter clocks when the *`priority1-value`* is the same for different timeTransmitter clocks in a network. The lower value takes precedence.
`Unicast Negotiation`	Method by which the announce, synchronization, and delay-response packet rates are negotiated between the timeTransmitter and the timeReceiver before a PTP session is established. Unicast negotiation is enabled or disabled.
`ESMC QL From Clock Class`	Denotes whether the conversion from clock class to QL is enabled or disabled.
`Clock Class/ESMC QLI`	Denotes the user defined clock class to QL conversion.
`SNMP Trap Status`	Denotes the SNMP trap generation status (Enabled or Disabled).
`Master Parameters`	`Delay Request Timeout` The default value is 30 seconds. The range is from 30 to 300 seconds.
`Transparent-clock-config`	Displays if transparent clock mode is enabled or disabled.

Table 1035: show ptp global-information Output Fields *(Continued)*

Field Name	Field Description
`Transparent-clock-status`	Display the status of the transparent clock operation. The following status is as follows: • **N/A**—Transparent clock is not configured. • **Active**—Transparent clock is configured and working properly. • **sync-in-progress**—This is a temporary state. During startup, all the PHYs are synchronized with each other. This status also occurs when a new PIC is plugged in to the switch and all the PHYs go through the synchronization cycle again. **NOTE**: Transparent clock operation is disabled during the synchronization of PHYs. • **Inactive**—Transparent clock is configured but not working properly. This indicates a hardware error in PHY timestamping logic.
`APTS`	Displays the status of assisted partial timing support (APTS) on the device.
`APTS Domain`	Indicates the domain on which APTS is enabled on.
`APTS Unicast Negotiation`	Displays the status of assisted partial timing support (APTS) unicast negotiation on the device.
`Holdover-in-spec Duration`	Indicates the time duration in minutes wherein the servo state will be holdover-in-specification after which the holdover time expires and the device moves to holdover-out-specification.
`Syntonized transparent-clock-config`	Displays if syntonized transparent clock mode is enabled or disabled.
`PPM Status`	Displays the Passive Port Monitoring (PPM) status (Enabled or Disabled) (only on PTX10004, PTX10008 devices).
`PPM Delay Request Interval`	PPM delay request interval supporting G.8275.1 specifications. The default value is -4, or (only on PTX10004, PTX10008 devices).

Table 1035: show ptp global-information Output Fields *(Continued)*

Field Name	Field Description
`Syntonized transparent-clock-status`	Display the status of the syntonized transparent clock operation. The following status is as follows: • **N/A**—Syntonized transparent clock is not configured. • **Active**—Syntonized transparent clock is configured and working properly. • **sync-in-progress**—This is a temporary state. During startup, all the PHYs are synchronized with each other. This status also occurs when a new PIC is plugged into the switch and all the PHYs go through the synchronization cycle again. **NOTE**: Syntonized transparent clock operation is disabled during the synchronization of PHYs. • **Inactive**—Syntonized transparent clock is configured but not working properly. This indicates a hardware error in PHY timestamping logic.
`UTC Leap Seconds`	The number of UTC leap seconds is 37 seconds by default. However, you can configure a different value.
`Phase adjust threshold`	Displays the maximum phase offset in nanoseconds that is adjusted by the G.8275.2 enhanced servo in a phase-aligned state. A higher phase adjust threshold permits a clock to tolerate phase offset changes because of higher network packet delay variation (PDV). **NOTE**: This field is displayed for supported platforms.
`Frequency lock threshold`	Displays the frequency offset in parts per billion (ppb) that the servo locks with. Locking with a higher frequency lock threshold allows a quicker frequency lock, although at the expense of a higher phase offset per window in the phase-aligned state. **NOTE**: This field is displayed for supported platforms.
`Phase lock threshold`	Displays he maximum phase offset in nanoseconds with which G.8275.2 enhanced servo enters the phase-aligned state. Phase locking with a higher phase lock threshold leads to the servo and time receiver getting into the phase-aligned state quickly. **NOTE**: This field is displayed for supported platforms.

Sample Output

show ptp global-information (Transparent Clock Configured)

```
user@switch> show ptp global-information
PTP Global Configuration:
Transparent-clock-config : ENABLED
Transparent-clock-status : ACTIVE
```

show ptp global-information (Syntonized Transparent Clock Configured)

```
user@switch> show ptp global-information
PTP Global Configuration:
Syntonized transparent-clock-config : ENABLED
Syntonized transparent-clock-status : ACTIVE
Syntonized clock-status : LOCKED
```

show ptp global-information (Default Profile)

```
user@switch> show ptp global-information
PTP Global Configuration:
Domain number           : 0
Clock mode              : Ordinary
Profile type            : IEEE-2008
Priority Level1         : 128
Priority Level2         : 128
Path Trace              : Disabled
Unicast Negotiation     : Disabled
ESMC QL From Clock Class: Disabled
Clock Class/ESMC QL     :  -
SNMP Trap Status        : Disabled
PHY Time Stamping       : Disabled
UTC Leap Seconds        : 37
```

show ptp global-information (Enterprise Profile)

```
user@switch> show ptp global-information
PTP Global Configuration:
Domain number           : 0
Clock mode              : Boundary
Profile type            : Enterprise
Priority Level1         : 128
Priority Level2         : 128
…

…
```

show ptp global-information (G.8275.1 Profile)

```
user@host> show ptp global-information
PTP Global Configuration:
Domain number           : 0
Clock mode              : Boundary
Profile type            : G.8275.1
Priority Level1         : 128
Priority Level2         : 128
Local Priority          : 128
Path Trace              : Disabled
Unicast Negotiation     : Disabled
ESMC QL From Clock Class: Disabled
Clock Class/ESMC QL     : -
SNMP Trap Status        : Disabled
PHY Time Stamping       : Enabled
UTC Leap Seconds        : 37
Transparent-clock-config : Disabled
Transparent-clock-status : N/A
  Sync Interval         : -
  Delay Request Interval: -4 (16 packets per second)
  Announce Interval     : -
  Announce Timeout      : 3
Grant Duration          : -
Master Parameters:
  Sync Interval         : -4 (16 packets per second)
  Delay Request Interval: -
```

```
  Announce Interval     : -3 (8 packets per second)
  Delay Request Timeout : -
  Clock Step            : one-step
  Arbitrary Mode        : FALSE
Number of Slaves        : 1
Number of Masters       : 0
Number of Stateful      : 0
```

show ptp global-information (G.8275.2.enh Profile)

```
user@host> show ptp global-information
PTP Global Configuration:
Domain number           : 44
Clock mode              : Ordinary
Profile type            : G.8275.2.ENH
Servo type              : Default
Priority Level1         : 128
Priority Level2         : 255
Local Priority          : 128
Path Trace              : Disabled
Unicast Negotiation     : Enabled
ESMC QL From Clock Class: Disabled
Clock Class/ESMC QL     :  -
SNMP Trap Status        : Disabled
PHY Time Stamping       : Disabled
UTC Leap Seconds        : 37
Transparent-clock-config : DISABLED
Transparent-clock-status : N/A
Slave Parameters:
  Sync Interval         : -6 (64 packets per second - unicast request)
  Delay Request Interval: -6 (64 packets per second - unicast request)
  Announce Interval     : -3 (8 packets per second - unicast request)
  Announce Timeout      : 3
  Grant Duration        : 300
Master Parameters:
  Sync Interval               : -6 (64 packets per second)
  Announce Interval           : 0 (1 packet per second)
  Min Announce Interval       : -3 (8 packets per second)
  Max Announce Interval       : 0 (1 packet per second)
  Min Sync Interval           : -7 (128 packets per second)
  Max Sync Interval           : 0 (1 packet per second)
```

```
  Min Delay Response Interval : -7 (128 packets per second)
  Max Delay Response Interval : 0 (1 packet per second)
  Delay Request Timeout       : <not applicable>
  Clock Step                  : one-step
  Arbitrary Mode              : FALSE

Number of Slaves       : 1
Number of Masters      : 0
Number of Stateful     : 0
```

show ptp global-information (APTS Configured)

```
user@host> show ptp global-informationPTP Global Configuration:
Domain number           : 24
Clock mode              : Boundary
Profile: G.8275.1
Priority Level1         : 128
Priority Level2         : 128
Local Priority          : 128
Path Trace              : Disabled
ESMC QL From Clock Class: Enabled
Clock Class/ESMC QL     : 6 / (QL PRC/SSM 0x2)
SNMP Trap Status        : Disabled
PHY Time Stamping       : Enabled
UTC Leap Seconds        : 37
Transparent-clock-config : DISABLED
Transparent-clock-status : N/A
APTS                     : Enabled
APTS Domain              : 44
APTS Unicast Negotiation : Enabled
Holdover-in-spec Duration: 240 Minutes

Slave Parameters:
  Sync Interval         : -6 (64 packets per second - unicast request)
  Delay Request Interval: -6 (64 packets per second - unicast request)
  Announce Interval     : 0 (1 packet per second - unicast request)
  Announce Timeout      : 3
  Grant Duration        : 300
Master Parameters:
  Sync Interval               : -4 (16 packets per second)
  Announce Interval           : -3 (8 packets per second)
```

```
  Delay Request Timeout       : <not applicable>
  Clock Step                  : one-step
  Arbitrary Mode              : FALSE
Master Parameters (G.8275.2.enh):
  Min Announce Interval       : -3 (8 packets per second)
  Max Announce Interval       : 0 (1 packets per second)
  Min Sync Interval           : -7 (128 packets per second)
  Max Sync Interval           : 0 (1 packet per second)
  Min Delay Response Interval : -7 (128 packets per second)
  Max Delay Response Interval : 0 (1 packet per second)
Number of Slaves       : 1
Number of Masters      : 3
Number of Stateful     : 0
```

show ptp global-information (MX10008, ACX7100-32C and ACX7100-48L)

```
user@host> show ptp global-information
PTP Global Configuration:
Domain number           : 44
Clock mode              : Boundary
Profile type            : G.8275.2.ENH
Servo type              : Enhanced
Priority Level1         : 128
Priority Level2         : 128
Local Priority          : 128
Path Trace              : Disabled
Unicast Negotiation     : Enabled
ESMC QL From Clock Class: Disabled
Clock Class/ESMC QL     :  -
SNMP Trap Status        : Disabled
PHY Time Stamping       : Enabled
UTC Leap Seconds        : 37
Transparent-clock-config : DISABLED
Transparent-clock-status : N/A
Phase adjust threshold   : 1100
Frequency lock threshold : 1
Phase lock threshold     : 200
Slave Parameters:
  Sync Interval         : -6 (64 packets per second - unicast request)
  Delay Request Interval: -6 (64 packets per second - unicast request)
  Announce Interval     : 0 (1 packet per second - unicast request)
```

```
  Announce Timeout      : 3
  Grant Duration        : 300
Master Parameters:
  Min Announce Interval       : -3 (8 packets per second)
  Max Announce Interval       : 0 (1 packet per second)
  Min Sync Interval           : -7 (128 packets per second)
  Max Sync Interval           : 0 (1 packet per second)
  Min Delay Response Interval : -7 (128 packets per second)
  Max Delay Response Interval : 0 (1 packet per second)
  Delay Request Timeout       : <not applicable>
  Clock Step                  : one-step
  Arbitrary Mode              : FALSE

Number of Slaves       : 1
Number of Masters      : 0
Number of Stateful     : 0
```

show ptp global-information (SMPTE ST-2059-2 Profile)

```
user@host> show ptp global-information
PTP Global Configuration:
Domain number           : 127
Clock mode              : Boundary
Profile type            : SMPTE ST-2059-2
Priority Level1         : 128
Priority Level2         : 128
Local Priority          : -
Path Trace              : Disabled
Unicast Negotiation     : Disabled
ESMC QL From Clock Class: Disabled
Clock Class/ESMC QL     : -
SNMP Trap Status        : Disabled
PHY Time Stamping       : Disabled
UTC Leap Seconds        : 37
Transparent-clock-config : Disabled
Transparent-clock-status : N/A
  Sync Interval         : -
  Delay Request Interval: -3 (8 packets per second)
  Announce Interval     : -
  Announce Timeout      : 3
  Grant Duration        : -
```

```
Master Parameters:
  Sync Interval         : -3 (8 packets per second)
  Announce Interval     : -2 (4 packets per second)
  Min Announce Interval : -3 (8 packets per second)
  Max Announce Interval : 1 (1 packet every 2 seconds)
  Min Sync Interval     : -7 (128 packets per second)
  Max Sync Interval     : -1 (2 packets per second)
  Min Delay Response Interval : -7 (128 packets per second)
  Max Delay Response Interval : 4 (1 packet every 16 seconds)
  Delay Request Timeout : 30
  Clock Step            : one-step
  Arbitrary Mode        : FALSE
Number of Slaves       : 1
Number of Masters      : 1
Number of Stateful     : 0
```

show ptp global-information (AES67 Profile)

```
user@host> show ptp global-information
PTP Global Configuration:
Domain number           : 0
Clock mode              : Boundary
Profile type            : AES67
Priority Level1         : 128
Priority Level2         : 128
Local Priority          : -
Path Trace              : Disabled
Unicast Negotiation     : Disabled
ESMC QL From Clock Class: Disabled
Clock Class/ESMC QL     : -
SNMP Trap Status        : Disabled
PHY Time Stamping       : Disabled
UTC Leap Seconds        : 37
Transparent-clock-config : Disabled
Transparent-clock-status : N/A
  Sync Interval         : -
  Delay Request Interval: -3 (8 packets per second)
  Announce Interval     : -
  Announce Timeout      : 3
  Grant Duration        : -
Master Parameters:
```

```
  Sync Interval         : -3 (8 packets per second)
  Announce Interval     : 1 (1 packet every 2 seconds)
  Min Announce Interval : 0 (1 packet per second)
  Max Announce Interval : 4 (1 packet every 16 seconds)
  Min Sync Interval     : -4 (16 packets per second)
  Max Sync Interval     : 1 (1 packet every 2 seconds)
  Min Delay Response Interval : -7 (128 packets per second)
  Max Delay Response Interval : 4 (1 packet every 16 seconds)
  Delay Request Timeout : 30
  Clock Step            : one-step
  Arbitrary Mode        : FALSE
Number of Slaves       : 1
Number of Masters      : 1
Number of Stateful     : 0
```

show ptp global-information (Combined AES67 and SMPTE ST-2059-2 Profile)

```
user@host> show ptp global-information
PTP Global Configuration:
Domain number           : 0
Clock mode              : Boundary
Profile type            : Combined AES67 and SMPTE ST-2059-2
Priority Level1         : 128
Priority Level2         : 128
Local Priority          : -
Path Trace              : Disabled
Unicast Negotiation     : Disabled
ESMC QL From Clock Class: Disabled
Clock Class/ESMC QL     : -
SNMP Trap Status        : Disabled
PHY Time Stamping       : Disabled
UTC Leap Seconds        : 37
Transparent-clock-config : Disabled
Transparent-clock-status : N/A
  Sync Interval         : -
  Delay Request Interval: -3 (8 packets per second)
  Announce Interval     : -
  Announce Timeout      : 3
  Grant Duration        : -
Master Parameters:
  Sync Interval         : -3 (8 packets per second)
```

```
  Announce Interval     : 0 (1 packet per second)
  Min Announce Interval : 0 (1 packet per second)
  Max Announce Interval : 1 (1 packet every 2 seconds)
  Min Sync Interval     : -4 (16 packets per second)
  Max Sync Interval     : -1 (2 packet per second)
  Min Delay Response Interval : -7 (128 packets per second)
  Max Delay Response Interval : 4 (1 packet every 16 seconds)
  Delay Request Timeout : 30
  Clock Step            : one-step
  Arbitrary Mode        : FALSE
Number of Slaves       : 1
Number of Masters      : 1
Number of Stateful     : 0
```

show ptp global-information (PTX10004, PTX10008 Devices)

```
user@host> show ptp global-information
PTP Global Configuration:
Domain number           : 24
Clock mode              : Boundary
Profile type            : G.8275.1
Priority Level1         : 128
Priority Level2         : 128
Local Priority          : 128
Path Trace              : Disabled
Unicast Negotiation     : Disabled
ESMC QL From Clock Class: Disabled
Clock Class/ESMC QL     :  -
SNMP Trap Status        : Disabled
PHY Time Stamping       : Enabled
UTC Leap Seconds        : 37
Transparent-clock-config : DISABLED
Transparent-clock-status : N/A
PPM Status              : ENABLED
PPM Delay Request Interval : -4 (16 packets per second)

Slave Parameters:
  Sync Interval         : <not applicable>
  Delay Request Interval: -4 (16 packets per second)
  Announce Interval     : <not applicable>
  Announce Timeout      : 3
```

```
  Grant Duration        : <not applicable>
Master Parameters:
  Sync Interval               : -4 (16 packets per second)
  Announce Interval           : -3 (8 packets per second)
  Delay Request Timeout       : <not applicable>
  Clock Step                  : one-step
  Arbitrary Mode              : FALSE

Number of Slaves       : 0
Number of Masters      : 1
Number of Stateful     : 1
```

Release Information

Command introduced in Junos OS Release 14.1X53-D25.

Output showing Syntonized Transparent Clock added in Junos OS Evolved Release 21.1R1.

SMPTE-ST 2059-2, AES67, and Combined AES67, and SMPTE-ST 2059-2 profile types are added in Junos OS Release 21.3 R1.

Frequency lock threshold, phase adjust threshold, and phase lock threshold output fields added in Junos OS Release 23.4R1 for MX10008.

Frequency lock threshold, phase adjust threshold, and phase lock threshold output fields supported for ACX7100-32C, ACX7100-48L, and ACX7024 in Junos OS Evolved Release 24.2R1.

APTS, APTS Domain, APTS Unicast Negotiation, Holdover-in-spec Duration output fields supported for ACX7024 and ACX7024X in Junos OS Evolved Release 24.2R2.

RELATED DOCUMENTATION

Understanding Transparent Clocks in Precision Time Protocol

Configuring Transparent Clock Mode for Precision Time Protocol

Configuring the Precision Time Protocol G.8275.2 Enhanced Profile (Telecom Profile)

show ptp hybrid

IN THIS SECTION

- Syntax | 23369
- Description | 23369
- Options | 23369
- Required Privilege Level | 23369
- Output Fields | 23370
- Sample Output | 23371
- Release Information | 23372

Syntax

```
show ptp hybrid
<config | status>
```

Description

Display the current configuration and current operation mode of the timeReceiver.

Options

config Display the PTP source to Synchronous Ethernet interface mappings.

status Display the current hybrid mode operational status.

Required Privilege Level

View

Output Fields

Table 1036 on page 23370 lists the output fields for the `show ptp hybrid` command. Output fields are listed in the approximate order in which they appear.

Table 1036: show ptp hybrid Output Fields

Field Name	Field Description
`ptp source`	Displays the IP address of the PTP source.
`synce source`	Displays the interface name of the Synchronous Ethernet source through which the PTP source is traceable.
`Configured Mode`	Displays the current configured mode of the router as `Hybrid`.
`Operating Mode`	Displays the current operation mode: `Hybrid` or `None`.
`PTP Reference`	Displays the IP address and the interface name of the PTP grandmaster clock.
`Synchronous Ethernet Reference`	Displays the interface name of the Synchronous Ethernet grandmaster clock.
`Lock state`	Displays the current lock state of the router: `Locked`, `Initializing`, or `Acquiring`.
`Lock state description`	Displays the description for the current lock state of the router: • `Initializing`—Hybrid mode is being initialized. • `Acquiring Frequency`—Synchronous Ethernet source identified for frequency synchronization, acquiring frequency-related data from timeTransmitter clock. • `Frequency Locked, Acquiring Phase`—Frequency locked from the Synchronous Ethernet source, acquiring phase-related data from timeTransmitter clock. • `Frequency and Phase Locked`—TimeReceiver clock is frequency and phase synchronized with timeTransmitter clock.

Sample Output

show ptp hybrid config

```
user@host> show ptp hybrid config
ptp source          synce source
100.1.1.2              ge-1/1/2
```

show ptp hybrid status

```
user@host> show ptp hybrid status
Hybrid Mode Status:
Configured Mode                : Hybrid
Operating  Mode                : Hybrid
PTP Reference                  : 100.1.1.2, ge-1/0/0.0
Synchronous Ethernet Reference : ge-1/1/2
Lock state                     : Locked
Lock state description         : Frequency and Phase Locked
```

show ptp hybrid status (PTX10008)

```
user@host> show ptp hybrid status
Hybrid Mode Status:
Configured Mode                : Hybrid
Operating  Mode                : Hybrid
PTP Reference                  : e4:5d:37:ff:fe:4e:d7:76, et-6/0/16:2.0
Synchronous Ethernet Reference : et-6/0/16:2
Lock state                     : Locked
Lock state description         : Frequency Locked Phase Locked
```

show ptp hybrid status (MX10008 with MX10k-LC2101 and MX10K-LC9600 line cards)

```
user@host> show ptp hybrid status
Hybrid Mode Status:
Configured Mode                : Hybrid
Operating  Mode                : Hybrid
PTP Reference                  : d4:04:ff:ff:fe:7e:5d:a2, xe-1/1/0:0
```

```
Frequency Reference            : external-0 (bits)
Lock state                     : Locked
Lock state description         : Frequency Locked Phase Locked
```

Release Information

Command introduced in Junos OS Release 12.2R2.

RELATED DOCUMENTATION

Understanding Hybrid Mode

show ptp lock-status

IN THIS SECTION

- Syntax | 23372
- Description | 23373
- Options | 23373
- Required Privilege Level | 23373
- Output Fields | 23373
- Sample Output | 23374
- Release Information | 23377

Syntax

```
show ptp lock-status
```

Description

(ACX Series, MX80, MX240, MX480, MX960, PTX10008 routers, and QFX Series switches) Display information about the lock status of the timeReceiver. The output verifies whether the timeReceiver is aligned to the reference (timeTransmitter clock) or not.

Options

detail Display detailed information about the lock status of the timeReceiver.

Required Privilege Level

view

Output Fields

Table 1037 on page 23373 lists the output fields for the `show ptp lock-status` command. Output fields are listed in the approximate order in which they appear.

Table 1037: show ptp lock-status Output Fields

Field Name	Field Description
`Lock State`	State of the timeReceiver clock with respect to its timeTransmitter clock: • Freerun • Holdover • Phase Aligned • Acquiring • Initializing • Freq locked
`Phase offset`	Time offset information of a timeReceiver clock with respect to its timeTransmitter clock. Precision of this time offset is 1 nanosecond.

Table 1037: show ptp lock-status Output Fields *(Continued)*

Field Name	Field Description
`State since`	Date, time, and how long ago the lock status of the PTP timeReceiver or timeReceiver clock changed. The format is `State since: year-month-day hour:minute:second:timezone (hour:minute:second ago)`. For example, `State since: 2002-04-26 10:52:40 PDT (04:33:20 ago)`. On ACX Series routers, this field is displayed in Junos OS Release 15.1 and later.
`Selected Master Details`	Details include the following: • `Upstream Master address`—The address of the remote timeTransmitter from which the timeReceiver acquires the clock. • `Slave interface`—The timeReceiver interface on this router corresponding to the timeTransmitter above. **NOTE**: On ACX Series router, if the PTP lock state is **FREERUN**, then the **Selected Master Details** field is not shown.

Sample Output

show ptp lock-status

```
user@host> run show ptp lock-status
Lock Status:

Lock State    : 5 (PHASE ALIGNED)
Phase offset  : 0.000000001 sec
```

show ptp lock-status

```
user@host> show ptp lock-status
Lock Status:

Lock State    : 1 (FREERUN)
Phase offset  : 0.000000869 sec
```

show ptp lock-status detail

```
user@host> show ptp lock-status detail
Lock Status:

Lock State    : 5 (PHASE ALIGNED)
State since   : 2014-09-10 11:24:11 PDT (00:02:51 ago)

Phase offset  : 0.000000030 sec

Selected Master Details:
Upstream Master address    : 13.13.13.1
Slave interface            : ge-0/1/5.0
```

show ptp lock-status detail

```
user@host> show ptp lock-status detail
Lock Status:
Lock State    : 5 (PHASE ALIGNED)
Phase offset  : -0.000000010 sec

Selected Master Details:
Upstream Master address    : 2001:cdba:0000:0000:0000:0000:3257:9652
Slave interface            : ge-0/2/0.0
Parent Id                  : 84:18:88:ff:fe:c0:34:00
GMC Id                     : 00:18:0b:ff:fe:20:03:14
```

show ptp lock-status detail

```
user@host> show ptp lock-status detail
Lock Status:

Lock State    : 5 (PHASE ALIGNED)
Phase offset  : -0.000000016 sec
State since   : 2018-10-17 01:53:29 UTC (01:09:37 ago)

Selected Master Details:
Upstream Master address    : 2003::a
Slave interface            : ae0.0
```

```
Parent Id                  : 12:34:56:78:9a:bc:de:01
GMC Id                     : 12:34:56:78:9a:bc:de:01
```

show ptp lock-status detail

```
user@host> show ptp lock-status detail
Lock Status:

Lock State    : 5 (PHASE ALIGNED)
Phase offset  : -0.000000016 sec
State since   : 2018-10-17 01:53:29 UTC (01:09:37 ago)

Selected Master Details:
Upstream Master address    : 2003::a
Slave interface            : lo0.0
Parent Id                  : 12:34:56:78:9a:bc:de:01
GMC Id                     : 12:34:56:78:9a:bc:de:01
```

show ptp lock-status detail (QFX5120-48T)

```
user@host> show ptp lock-status detail
Lock Status:

Lock State    : 5 (PHASE ALIGNED)
Phase offset  : -0.000000013 sec
State since   : 2021-05-20 23:15:01 PDT (1d 00:26 ago)

Selected Master Details:
Upstream Master address    : 51.0.0.1
Slave interface            : xe-0/0/50:0.0
Parent Id                  : 00:00:00:00:00:00:00:01
GMC Id                     : 00:00:00:00:00:00:00:01
```

show ptp lock-status detail (PTX10008)

```
user@host> show ptp lock-status detail
Lock Status:
```

```
Lock State    : 5 (PHASE ALIGNED)
Phase offset  : -0.000002880 sec
State since   : 2021-05-06 18:55:13 UTC (00:01:20 ago)

Selected Master Details:
Upstream Master address    : 00:00:00:00:00:00:00:01
Slave interface            : et-6/0/16:2.0
Clock reference state      : Clock locked
1pps reference state       : Clock qualified
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

IEEE 1588v2 PTP Boundary Clock Overview

IEEE 1588v2 Precision Timing Protocol (PTP)

Precision Time Protocol Overview

show ptp master

IN THIS SECTION

- Syntax | 23378
- Description | 23378
- Options | 23378
- Required Privilege Level | 23378
- Output Fields | 23378
- Sample Output | 23380
- Release Information | 23382

Syntax

```
show ptp master
<brief |detail | interface>
```

Description

(MX80, MX240, MX480, MX960, PTX10008 routers, and QFX Series switches) Display information about the configured timeTransmitter and the status of the timeTransmitter.

Options

brief Display information about the timeTransmitter in brief.

detail Display information about the timeTransmitter in detail.

interface Display information about the configured interface of the timeTransmitter.

Required Privilege Level

View

Output Fields

Table 1038 on page 23378 lists the output fields for the `show ptp master` command. Output fields are listed in the approximate order in which they appear.

Table 1038: show ptp primary Output Fields

Field Name	Field Description
`Interface`	Name of the interface configured for Precision Time Protocol (PTP) on the timeTransmitter.

Table 1038: show ptp primary Output Fields *(Continued)*

Field Name	Field Description
`Status`	Status of the Precision Time Protocol timeTransmitter: • `Master` or `Slave` • `Active` or `Inactive` • `Initializing` or `Down`
`Local Adddress`	IP or MAC address of the configured timeTransmitter clock.
`Status` (Local address Status)	Status of the local address of the interface: • `Configured` or `Not configured` • `Master` or `Slave` • `Active` or `Inactive`
`Status` (Remote address Status)	Status of the remote address of the interface on the QFX Series: • `Configured` or `Not configured` • `Master` or `Slave` • `Active` or `Inactive`
`Total Remote Slaves`	Number of remote timeReceiver clocks.
`Slave Address`	IP or MAC addresses of the timeReceivers.
`Status` (Member Address Status)	Status of the address of the timeReceiver: • `Configured` or `Not configured` • `Master` or `Slave` • `Active` or `Inactive` or `Ready`

Sample Output

show ptp primary

```
user@host> run show ptp master brief
PTP Master Interface Configured:

Master Interface           Status
ge-7/0/2.0                 Master, Active
```

show ptp primary detail

```
user@host> run show ptp master detail
PTP Master Interface Details:
Interface  : xe-0/0/6:1.0
Status     : Master, Active
Clock Info :
    Local Address: 50.50.50.1         Status: Configured, Master, Active
        Remote Address: 224.0.1.129
    Total Remote Slaves: 1
```

show ptp primary detail

```
user@host> run show ptp slave detail
PTP Master Interface Details:

Interface  : xe-0/0/31:3.0
Status     : Master, Active
Clock Info :
    Local Address: 10.10.10.2         Status: Configured, Master, Active
        Remote Address: 10.10.10.1         Status: Learned, Slave, Active
        Remote Address: 224.0.1.129        Status: Configured, Slave, Active
    Total Remote Slaves: 2

Interface  : xe-0/0/35:3.0
Status     : Master, Active
Clock Info :
    Local Address: 10.2.2.1           Status: Configured, Master, Active
```

```
        Remote Address: 10.2.2.24          Status: Learned, Slave, Active
        Remote Address: 10.2.2.29          Status: Learned, Slave, Active
        Remote Address: 10.2.2.30          Status: Learned, Slave, Active
        Remote Address: 10.2.2.32          Status: Learned, Slave, Active
        Remote Address: 10.2.2.35          Status: Learned, Slave, Active
        Remote Address: 10.2.2.52          Status: Learned, Slave, Active
        Remote Address: 10.2.2.61          Status: Learned, Slave, Active
        Remote Address: 224.0.1.129        Status: Configured, Slave, Active
    Total Remote Slaves: 8
```

show ptp primary detail

```
user@host> run show ptp master detail
PTP Master Interface Details:
Interface  : ge-7/0/2.0
Status     : Master, Active
Clock Info :
    Local Address: 10.0.0.1               Status: Configured, Master, Active
    Total Remote Slaves: 0
        Slave IP: 10.0.0.2           Status: Configured, Slave, Active
```

show ptp primary detail

```
user@host> run show ptp master detail
 PTP Master Interface Details:
Interface  : ge-0/1/5.0
Status     : Master, Active
Clock Info :
    Local Address: 84:18:88:c0:60:a1  Status: Configured, Master, Active
        Remote Address: [Slave Mac]  Status: Configured, Slave, Active
    Total Remote Slaves: 1
```

show ptp primary detail

```
user@host> run show ptp master detail
PTP Master Interface Details:
Interface  : lo0.0
Status     : Master, Active
```

```
Clock Info :
    Local Address: 2001::1:132        Status: Configured, Master, Active
        Remote Address: 2002::1            Status: SECURE AUTO SLAVE,
                                                   Sig State: Ann + Sync
        Remote Address: 3001::1:133        Status: SECURE AUTO SLAVE,
                                                   Sig State: Ann + Sync
    Total Remote Slaves: 2
```

show ptp primary detail (PTX10008)

```
user@host> run show ptp master detail
PTP Master Interface Details:
Interface  : et-6/0/16:0.0
Status     : Master, Active
Clock Info :
    Local Address: e4:5d:37:4e:d0:94   Status: Configured, Master, Active
        Remote Address: e4:5d:37:ff:fe:4e:d7:76  Status: Configured, Slave, Active
```

show ptp interface ge-7/0/2.0

```
user@host> run show ptp master interface ge-7/0/2.0
PTP Master Interface Configured:

Master Interface           Status
ge-7/0/2.0                 Master, Active
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Precision Time Protocol Overview

Configuring the Precision Time Protocol G.8275.2 Enhanced Profile (Telecom Profile)

show ptp passive-port-monitor-status

IN THIS SECTION

- Syntax | 23383
- Description | 23383
- Options | 23383
- Required Privilege Level | 23383
- Output Fields | 23383
- Sample Output | 23385
- Release Information | 23388

Syntax

```
show ptp passive-port-monitor-status
```

Description

(PTX10004 and PTX10008) Displays the details of the PTP passive port interfaces.

Options

The command has no options.

Required Privilege Level

view

Output Fields

The Output Fields lists the output fields for the `show ptp passive-port-monitor-status` command. Output fields are listed in the approximate order in which they appear.

Table 1039: show ptp passive-port-monitor-status Output Fields

Field Name	Field Description
PPM Status	Status of the PTP passive port monitoring feature. (Enabled or Disabled)
PPM Delay Request Interval	The configured delay request rate for port monitoring. The value of delay request interval ranges from -4 to +4. This represents the packet rates as 2^(-4) that is (16 packets per second) to 2^4 (1 packets per 16 second). Default value is -4.
Actively Monitored-PTP ports data	Displays data from actively monitored PTP ports.
Local Interface	Displays details about local interface. For example, ae0.4094 (et-4/0/0, primary): Specifies that the local interface has an aggregated Ethernet ID 0 with port number 4094, et-4/0/0 is the Ethernet IP address that is configured as primary.
Clock Stream	Priority value of the clock stream: 0 through 255. The default is 128. The lower value takes precedence.
Source-Port-ID	Displays source port ID: MAC address and port number.
Destination-Port-ID	Displays destination port ID: MAC address and port number.
Port State	State of the PTP port (Master, Slave or Passive)
Operating Mode	Denotes the configured role (Stateful/Slave/Master) of the port.
PhaseOffsetThreshold(ns)	Time phase offset threshold value for the local interface defined in nano-seconds. The threshold value is a range from 0 to 100000000. Raises PTP performance alarm when the measured phase error exceeds the threshold over 15 minutes interval.
Configured role	The PTP stateful port configuration for passive port monitoring. Can be not-master, measure-only or both.

Table 1039: show ptp passive-port-monitor-status Output Fields *(Continued)*

Field Name	Field Description
`Measurement-status`	State of the PTP performance measurement. ctive for passive ports with monitoring enabled, not-applicable when port state is not passive but PTP monitoring is enabled and not- • Active: For passive ports with monitoring enabled. • Not-Applicable: When monitoring is enabled but port state is not passive. • Inactive: When monitoring is enabled but PTP is other than phase aligned state.
`Latest measurement data`	Details of the latest measurement data.
`Timestamp`	Timestamp of the measurement data.
`PhaseOffsetFromMaster(ns)`	Time phase offset of non-master device with master. The range is from 0 to 100000000 (100ms). The default value is 100ms.
`Master->slave delay(ns)`	Master to non-master synchronization interval; minimum, mean and maximum values (in nano seconds).
`Slave->master delay(ns)`	Non-master to Master synchronization interval; minimum, mean and maximum values (in nano seconds).
`Mean-path-delay(ns)`	Delay of mean path minimum, mean and maximum values.

Sample Output

show ptp passive-port-monitor-status (with passive port monitoring enabled and configured as not-master)

```
user@root> show ptp passive-port-monitor-status
PPM Status              : ENABLED
PPM Delay Request Interval : -4 (16 packets per second)
```

```
Actively Monitored-PTP ports data:

Local Interface          : ae0.4094 (et-4/0/0, primary)
Clock Stream             : 5
Source-Port-ID           : 00:cc:34:ff:fe:77:c4:57  Port : 5
Destination-Port-ID      : 00:cc:34:ff:fe:77:c8:de  Port : 1
Port State               : Passive
Operating Mode           : Stateful
PhaseOffsetThreshold(ns) : 100000
Configured role          : Not-master
Measurement-status       : Valid
Latest measurement data  :
Timestamp                : Aug 20 22:26:16
PhaseOffsetFromMaster(ns): 15       / 17       / 791       (min/mean/max)
Master->slave delay(ns)  : 0        / 1575     / 1584      (min/mean/max)
Slave->master delay(ns)  : 0        / 1540     / 1582      (min/mean/max)
Mean-path-delay(ns)      : 1557      (mean)
```

show ptp passive-port-monitor-status (with passive port monitoring enabled and configured as measure-only)

```
user@root> show ptp passive-port-monitor-status
PPM Status               : ENABLED
PPM Delay Request Interval : -4 (16 packets per second)

Actively Monitored-PTP ports data:

Local Interface          : et-0/0/34
Clock Stream             : 6
Source-Port-ID           : 00:cc:34:ff:fe:77:c4:57  Port : 6
Destination-Port-ID      : 00:00:00:00:00:00:00:01  Port : 1
Port State               : Passive
Operating Mode           : Stateful
PhaseOffsetThreshold(ns) : 100000
Configured role          : Measure-only
Measurement-status       : Valid
Latest measurement data  :
Timestamp                : Aug 20 22:26:16
PhaseOffsetfromMaster(ns): 4078     /4087     /4115     (min/mean/max)
Master->slave delay(ns)  : 4078     /4087     /4115     (min/mean/max)
```

```
Slave->Master delay(ns)  : 4078     /4087     /4115     (min/mean/max)
Mean-path-delay(ns)      : 4078     (mean)
```

show ptp passive-port-monitor-status (with passive port monitoring enabled and configured as not-master and measure-only)

```
user@root> show ptp passive-port-monitor-status
PPM Status              : ENABLED
PPM Delay Request Interval : -4 (16 packets per second)

Actively Monitored-PTP ports data:

Local Interface          : et-0/0/16:2
Clock Stream             : 6
Source-Port-ID           : 00:cc:34:ff:fe:77:c4:57  Port : 7
Destination-Port-ID      : 00:00:00:00:00:00:00:01  Port : 2
Port State               : Passive
Operating Mode           : Stateful
PhaseOffsetThreshold(ns) : 100000
Configured role          : Measure-only, Not-master
Measurement-status       : Valid
Latest measurement data  :
Timestamp                : Aug 20 22:26:16
PhaseOffsetfromMaster(ns): 4078     /4087     /4115     (min/mean/max)
Master->slave delay(ns)  : 4078     /4087     /4115     (min/mean/max)
Slave->Master delay(ns)  : 4078     /4087     /4115     (min/mean/max)
Mean-path-delay(ns)      : 4078     (mean)
```

show ptp passive-port-monitor-status (with passive port monitoring enabled, but port state is not passive)

```
user@root> show ptp passive-port-monitor-status
PTP Passive Port Interface Details:

PPM-monitor-status       : ENABLED
PPM-delay-request-rate   : -4 (16 packets per second)

Actively Monitored-PTP ports data   :
```

```
Local Interface          : et-0/0/3.0
Clock stream             : 5
Source-Port-ID           : e4:5d:37:ff:fe:4e:d7:76 port-num: 2
Port State               : Master
Monitor mode             : Enabled
PhaseOffsetThreshold(ns) : 100000000
Configured role          : None
Measurement-status       : Not-Applicable
```

show ptp passive-port-monitor-status (with passive port monitoring enabled, but port state not phase-aligned)

```
user@root> show ptp passive-port-monitor-status
PTP Passive Port Interface Details:

PPM-monitor-status       : ENABLED
PPM-delay-request-rate   : -4 (16 packets per second)

Actively Monitored-PTP ports data   :

Local Interface          : et-0/0/16:1.0
Clock stream             : 7
Source-Port-ID           : e4:5d:37:ff:fe:4e:d7:76 port-num: 4
Port State               : Passive
Monitor mode             : Enabled
PhaseOffsetThreshold(ns) : 100000000
Configured role          : Not-master
Measurement-status       : Inactive
```

show ptp passive-port-monitor-status (passive port monitoring disabled)

```
user@root> show ptp passive-port-monitor-status
PTP Passive Port Interface Details:

PPM-monitor-status       : DISABLED
```

Release Information

Command introduced in Junos OS Evolved Release 23.4R1.

show ptp path-trace detail

IN THIS SECTION

- Syntax | 23389
- Description | 23389
- Options | 23389
- Required Privilege Level | 23389
- Output Fields | 23389
- Sample Output | 23390
- Release Information | 23390

Syntax

```
show ptp path-trace detail
```

Description

(MX80, MX240, MX480, and MX960 routers) Display the details of the path an announce message takes in a PTP ring topology.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1040 on page 23390 lists the output fields for the `show ptp path-trace detail` command. Output fields are listed in the approximate order in which they appear.

Table 1040: show ptp path-trace detail Output Fields

Field Name	Field Description
Hop Count	The count of the next router in a network trail where the announce message is received.
Member Clock Identity	Clock identity of the timeReceiver or timeReceiver as defined in IEEE 1588.

Sample Output

show ptp path-trace detail

```
user@host> show ptp path-trace detail
Hop count          Member Clock Identity

     1               00:05:85:ff:fe:74:1f:d0
     2               00:05:85:ff:fe:73:ef:d0
```

Release Information

Command introduced in Junos OS Release 13.3R4.

RELATED DOCUMENTATION

IEEE 1588v2 PTP Boundary Clock Overview

Precision Time Protocol Overview

show ptp phy-timestamping-interfaces

IN THIS SECTION

Syntax | 23391

- Description | 23391
- Options | 23391
- Required Privilege Level | 23391
- Output Fields | 23391
- Sample Output | 23392
- Release Information | 23392

Syntax

```
show ptp phy-timestamping-interfaces
```

Description

Display information about the interfaces that have PHY timestamping enabled. PHY timestamping is the timestamping of the IEEE 1588 event packets at the physical layer. Timestamping the packet at the physical layer eliminates the noise or the packet delay variation (PDV) that is introduced by the Packet Forwarding Engine.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1041 on page 23392 lists the output fields for the `show ptp phy-timestamping-interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 1041: show-ptp-phy-timestamping-interfaces

Field Name	Field Description
`PHY time-stamp interfaces`	Name of the interface configured for phy-timestamping.

Sample Output

show ptp phy-timestamping-interfaces

```
user@host> show ptp phy-timestamping-interfaces
PHY time-stamp interfaces:
    xe-0/2/1
    xe-4/1/3
```

Release Information

Command introduced in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

Precision Time Protocol Overview

show ptp port

IN THIS SECTION

- Syntax | 23393
- Description | 23393
- Options | 23393
- Required Privilege Level | 23393

- Output Fields | 23393
- Sample Output | 23395
- Release Information | 23399

Syntax

```
show ptp port
<brief | detail>
ifl logical-interface-name
```

Description

(MX80, MX240, MX480, MX960, PTX10004, PTX10008, and ACX710 routers) Display information about the number of ports created according to the configuration. For each unique local IP address, one Precision Time Protocol port is created.

Options

brief	Display information about the PTP port in brief.
detail	Display information about the PTP port in detail.
logical-interface-name	Display PTP port information for a specific logical interface.

Required Privilege Level

View

Output Fields

Table 1042 on page 23394 lists the output fields for the `show ptp port` command. Output fields are listed in the approximate order in which they appear.

Table 1042: show ptp port Output Fields

Field Name	Field Description
Local Interface	Local logical interface.
`Local IP`	IP address of the interface acting as the timeReceiver.
`Remote IP`	IP address of the remote node.
`Clock Stream`	Unique index for each session created.
`State Changes`	The total number of times the port state has changed.
`Last Change`	The time elapsed after the last change was made to the port state.
`Clock Identity`	IP address of the timeReceiver.
`Port State`	Status of the port: `PTP listening`, `PTP initializing` or `Passive`.
`Delay Req Interval`	Interval in seconds between the delay request messages sent by the timeReceiver to the timeTransmitter.
`Announce Interval`	Logarithmic mean interval for the announce messages to be sent by the timeTransmitter.
`Announce Timeout`	Number of times the announce interval message has to pass between the timeReceiver and the timeTransmitter without receipt of an announce message.
`Sync Interval`	Logarithmic mean interval for sync interval messages to be sent by the timeTransmitter.
`Delay Mechanism`	Type of delay mechanism used.
`Port Number`	PTP port number.

Table 1042: show ptp port Output Fields *(Continued)*

Field Name	Field Description
Operating Mode	Clock mode of the node.
Master Clock ID	Unique clock-identity of the timeTransmitter.
Previous Announce Messages	Previous announce messages.
Current Announce Message	Current announce messages.

Sample Output

show ptp port brief

```
user@host> show ptp port brief
PTP port-data:
Local Interface   : xe-0/0/18.0
Local Address     : 98:a4:04:7f:19:15
Remote Address    : 01:1b:19:00:00:00
Clock Stream      : 0              Clock Identity    : 98:a4:04:ff:fe:7f:19:7f
State Changes     : 2              Last Change       : 00:03:08 ago
Port State        : Master         Delay Req Interval: -4
Announce Interval : -3             Announce Timeout  : 3
Sync Interval     : -4             Delay Mechanism   : End-to-end
Port Number       : 1              Operating Mode    : Stateful

Local Interface   : xe-0/0/18.0
Local Address     : 98:a4:04:7f:19:15
Remote Address    : 01:1b:19:00:00:00
Clock Stream      : 0              Clock Identity    : 98:a4:04:ff:fe:7f:19:7f
State Changes     : 2              Last Change       : 00:03:08 ago
Port State        : Master         Delay Req Interval: -4
Announce Interval : -3             Announce Timeout  : 3
```

```
Sync Interval     : -4             Delay Mechanism   : End-to-end
Port Number       : 1              Operating Mode    : Stateful
```

show ptp port detail

```
user@host> show ptp port detail
PTP port-data:
Local Interface   : xe-0/0/18.0
Local Address     : 98:a4:04:7f:19:15
Remote Address    : 01:1b:19:00:00:00
Clock Stream      : 0              Clock Identity    : 98:a4:04:ff:fe:7f:19:7f
State Changes     : 2              Last Change       : 00:03:08 ago
Port State        : Master         Delay Req Interval: -4
Announce Interval : -3             Announce Timeout  : 3
Sync Interval     : -4             Delay Mechanism   : End-to-end
Port Number       : 1              Operating Mode    : Stateful
Local Priority    : 128            Asymmetry         : 0
Vlan-id           : <null>
Priority          : <null>
Foreign Master Clock Details:
Master Clock Id               : 98:a4:04:ff:fe:7f:19:7f
Previous Announce Messages    : 0
Current Announce Messages     : 0 Local Interface   : xe-0/0/7.0

Local Address     : 98:a4:04:7f:19:0a
Remote Address    : 01:1b:19:00:00:00
Clock Stream      : 1              Clock Identity    : 98:a4:04:ff:fe:7f:19:7f
State Changes     : 2              Last Change       : 00:03:08 ago
Port State        : Slave          Delay Req Interval: -4
Announce Interval : -3             Announce Timeout  : 3
Sync Interval     : -4             Delay Mechanism   : End-to-end
Port Number       : 2              Operating Mode    : Stateful
Local Priority    : 128            Asymmetry         : 0
Vlan-id           : <null>
Priority          : <null>
Foreign Master Clock Details:
Master Clock Id               : 98:a4:04:ff:fe:80:3a:7f
Previous Announce Messages    : 0
Current Announce Messages     : 0
```

show ptp port ifl

```
user@host> show ptp port ifl ge-1/0/5.0
PTP port-data:
Local Interface   : ge-1/0/5.0
Local Address     : 2001:db8:00:05:85:73:b0:aa
Remote Address    : 01:1b:19:00:00:00
Clock Stream      : 1              Clock Identity    : 98:a4:04:ff:fe:7f:19:7f
State Changes     : 2              Last Change       : 00:03:08 ago
Port State        : Slave          Delay Req Interval: -4
Announce Interval : -3             Announce Timeout  : 3
Sync Interval     : -4             Delay Mechanism   : End-to-end
Port Number       : 2              Operating Mode    : Stateful
Local Priority    : 128            Asymmetry         : 0
Vlan-id           : <null>
Priority          : <null>
Foreign Master Clock Details:
Master Clock Id               : 98:a4:04:ff:fe:80:3a:7f
Previous Announce Messages    : 0
Current Announce Messages     : 0
```

show ptp port detail (PTX10008)

```
user@host> show ptp port detail
PTP port-data:
Local Interface   : et-6/0/16:2.0
Local Address     : e4:5d:37:4e:d0:96
Remote Address    : 01:1b:19:00:00:00
Clock Stream      : 1              Clock Identity    : 98:a4:04:ff:fe:7f:19:7f
State Changes     : 2              Last Change       : 00:03:08 ago
Port State        : Slave          Delay Req Interval: -4
Announce Interval : -3             Announce Timeout  : 3
Sync Interval     : -4             Delay Mechanism   : End-to-end
Port Number       : 2              Operating Mode    : Stateful
Local Priority    : 128            Asymmetry         : 0
Vlan-id           : <null>
Priority          : <null>
Foreign Master Clock Details:
Master Clock Id               : 98:a4:04:ff:fe:80:3a:7f
```

```
Previous Announce Messages    : 0
Current Announce Messages     : 0
```

show ptp port status (PTX10008)

```
user@host> show ptp port detail
PTP port-data:
Local Interface   : et-6/0/16:2.0
Local Address     : e4:5d:37:4e:d0:96
Remote Address    : 01:1b:19:00:00:00
Clock Stream      : 1              Clock Identity    : 98:a4:04:ff:fe:7f:19:7f
State Changes     : 2              Last Change       : 00:03:08 ago
Port State        : Slave          Delay Req Interval: -4
Announce Interval : -3             Announce Timeout  : 3
Sync Interval     : -4             Delay Mechanism   : End-to-end
Port Number       : 2              Operating Mode    : Stateful
Local Priority    : 128            Asymmetry         : 0
Vlan-id           : <null>
Priority          : <null>
Foreign Master Clock Details:
Master Clock Id               : 98:a4:04:ff:fe:80:3a:7f
Previous Announce Messages    : 0
Current Announce Messages     : 0
```

show ptp port (PTX10004, PTX10008) (PTP Passive Port Monitoring enabled and configured as not-master and measure-only)

```
user@host> show ptp port
PTP port-data:

Local Interface   : et-0/0/16:1.0
Local Address     : e4:5d:37:4e:c8:19
Remote Address    : 01:1b:19:00:00:00
Clock Stream      : 5              Clock Identity    : e4:5d:37:ff:fe:4e:d7:76
Port State        : Passive        Delay Req Interval: -4
Announce Interval : -3             Announce Timeout  : 3
Sync Interval     : -4             Delay Mechanism   : End-to-end
Port Number       : 2              Operating Mode    : Stateful, Not-master, Measure-only
Local Priority    : 128
```

show ptp port (PTX10004, PTX10008) (PTP Passive Port Monitoring enabled and configured as not-master)

```
user@host> show ptp port
PTP port-data:

Local Interface   : et-0/0/16:1.0
Local Address     : e4:5d:37:4e:c8:19
Remote Address    : 01:1b:19:00:00:00
Clock Stream      : 5              Clock Identity    : e4:5d:37:ff:fe:4e:d7:76
Port State        : Passive        Delay Req Interval: -4
Announce Interval : -3             Announce Timeout  : 3
Sync Interval     : -4             Delay Mechanism   : End-to-end
Port Number       : 2              Operating Mode    : Stateful, Not-master
Local Priority    : 128
```

show ptp port (PTX10004, PTX10008) (PTP Passive Port Monitoring Enabled and configured as measure-only)

```
user@host> show ptp port
PTP port-data:

Local Interface   : et-0/0/16:1.0
Local Address     : e4:5d:37:4e:c8:19
Remote Address    : 01:1b:19:00:00:00
Clock Stream      : 5              Clock Identity    : e4:5d:37:ff:fe:4e:d7:76
Port State        : Passive        Delay Req Interval: -4
Announce Interval : -3             Announce Timeout  : 3
Sync Interval     : -4             Delay Mechanism   : End-to-end
Port Number       : 2              Operating Mode    : Stateful, Measure-only
Local Priority    : 128
```

Release Information

Command introduced in Junos OS Release 12.2.

`ifl` *`logical-interface-name`* `added in Junos OS Release 16.1.`

RELATED DOCUMENTATION

Precision Time Protocol Overview

show ptp quality-level-mapping

IN THIS SECTION

- Syntax | 23400
- Description | 23400
- Required Privilege Level | 23400
- Output Fields | 23400
- show ptp quality-level-mapping (ACX710) | 23401
- Release Information | 23401

Syntax

```
show ptp quality-level-mapping
```

Description

Display Precision Time Protocol (PTP) clock class to Ethernet Synchronization Message Channel (ESMC) quality mapping.

Required Privilege Level

maintenance

Output Fields

Table 1043 on page 23401 lists the output fields for the `show ptp quality-level-mapping` command. Output fields are listed in the approximate order in which they appear.

Table 1043: show ptp quality-level-mapping Output Fields

Field Name	Field Description
`quality level`	ESMC quality level
`ptp clock class`	PTP clock class

show ptp quality-level-mapping (ACX710)

command-name

```
user@host> show ptp quality-level-mapping

quality level         ptp clock class
PRC                      140
SSU-A                    150
SSU-B                    160
SEC                      255
```

Release Information

Command introduced in Junos OS Release 12.2R2.

RELATED DOCUMENTATION

Synchronous Ethernet Overview

show ptp slave

IN THIS SECTION

Syntax | 23402

- Description | 23402
- Options | 23402
- Required Privilege Level | 23402
- Output Fields | 23402
- Sample Output | 23404
- Release Information | 23405

Syntax

```
show ptp slave
<brief | detail | interface>
```

Description

(MX80, MX240, MX480, MX960, PTX10008 routers, and the QFX Series) Display information about the configured timeReceiver and the status of the timeReceiver.

Options

brief Display information about the timeReceiver in detail.

detail Display information about the timeReceiver in detail.

interface Display information about the configured interface of the timeReceiver.

Required Privilege Level

View

Output Fields

Table 1044 on page 23403 lists the output fields for the `show ptp slave` command. Output fields are listed in the approximate order in which they appear.

Table 1044: show ptp client Output Fields

Field Name	Field Description
`Interface`	Name of the interface configured for Precision Time Protocol.
`Status`	Status of the Precision Time Protocol timeReceiver: • `Master` or `Slave` • `Active` or `Inactive` • `Initializing` or `Down`
`Interface`	Interface configured on the timeReceiver.
`Local Address`	IP or MAC address of the local interface.
`Status` (Local address Status)	Status of the local address of the interface acting as the timeReceiver: • `Configured` or `Unconfigured` • `Master` or `Slave` • `Active` or `Inactive` or `Ready`
`Status` (Remote address Status)	Status of the remote address of the interface on the QFX Series: • `Configured` or `Not configured` • `Master` or `Slave` • `Active` or `Inactive`
`Total Remote Masters`	Number of remote primaries.
`Remote Master`	IP or MAC address of the remote node.

Table 1044: show ptp client Output Fields *(Continued)*

Field Name	Field Description
`Status` (Member IP Address Status)	Status of the address of the timeTransmitter: • `Configured` or `Unconfigured` • `Master` or `Slave` • `Active` or `Inactive`

Sample Output

show ptp slave

```
user@host> run show ptp slave
PTP Slave Interfaces Configured:

Slave Interface             Status
ge-7/0/0.0                  Slave, Active
```

show ptp slave detail

```
user@host> run show ptp slave detail
PTP Slave Interface Details:

Interface        : ge-7/0/0.0
Status           : Slave, Active
Clock Info
    Local address : 10.10.1.10              Status: Configured, Slave, Active
    Total Remote Masters: 0
        Remote Master: 10.10.1.2        Status: Configured, Master, Active
```

show ptp slave detail

```
user@host> run show ptp slave detail
PTP Slave Interface Details:
```

```
Interface        : ge-0/1/5.0
Status           : Slave, Active
Clock Info
    Local Address 2001:cdba:0000:0000:0000:0000:3257:9653        Status: Configured, Slave,
Active
        Remote Master:: 2001:cdba:0000:0000:0000:0000:3257:9652 Status: Configured, Master,
Active
    Total Remote Masters: 1
```

show ptp slave detail

```
user@host> run show ptp slave detail
PTP Slave Interface Details:

Interface        : ae0.0
Status           : Slave, Active
Clock Info
    Local Address : 2003::b                  Status: Configured, Slave, Active
        Remote Master: 2003::a          Status: Configured, Master, Active
    Total Remote Masters: 1
```

show ptp slave detail (PTX10008)

```
user@host> run show ptp slave detail
PTP Slave Interface Details:

Interface        : et-6/0/9:1, primary
Status           : Slave, Active
Clock Info
    Local Address : e4:5d:37:4e:d2:f1        Status: Configured, Slave, Active
        Remote Master: 2c:21:72:ff:fe:d8:07:d0 Status: Configured, Master, Active
    Total Remote Masters: 1
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Precision Time Protocol Overview

Configuring the Precision Time Protocol G.8275.2 Enhanced Profile (Telecom Profile)

show ptp stateful detail

IN THIS SECTION

- Syntax | 23406
- Description | 23406
- Options | 23406
- Required Privilege Level | 23406
- Output Fields | 23407
- Sample Output | 23407
- Release Information | 23408

Syntax

```
show ptp stateful detail
```

Description

Display information about the stateful interfaces.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1045 on page 23407 lists the output fields for the `show ptp stateful detail` command. Output fields are listed in the approximate order in which they appear.

Table 1045: show ptp stateful detail

Field Name	Field Description
`Interface`	Name of the configured stateful interface.
`Status`	Status of the stateful interface: stateful and active or inactive.
`Local Address`	MAC address of the configured stateful interfaces.
`Status (of local address)`	Status of the local address of the interface: • Configured or not configured • Active or inactive
`Port Status`	Status of the port configured: • TimeTransmitter or timeReceiver
`Remote Address`	MAC address of the remote stateful interface.

Sample Output

show ptp stateful detail

```
user@host> show ptp stateful detail
Interface  : ge-7/0/1.0
Status     : Stateful, Active
Clock Info :
    Local Address: 01:00:5E:90:10:00  Status: Configured, Active
    Port Status: Slave
        Remote Address: 01:00:5E:90:10:FF
```

show ptp stateful detail (PTX10004 and PTX10008)

```
user@host> show ptp stateful detail
PTP Stateful Interface Details:

Interface  : et-6/0/16:0.0
Status     : Stateful, Active
Clock Info :
    Local Address: e4:5d:37:4e:d0:94  Status: Configured, Active
    Port Status: Master
        Remote Address: 01:1b:19:00:00:00

Interface  : et-6/0/16:2.0
Status     : Stateful, Active
Clock Info :
    Local Address: e4:5d:37:4e:d0:96  Status: Configured, Active
    Port Status: Slave
        Remote Address: 01:1b:19:00:00:00
```

Release Information

Command introduced in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

Precision Time Protocol Overview

show ptp statistics

IN THIS SECTION

- Syntax | 23409
- Description | 23409
- Options | 23409
- Required Privilege Level | 23409

- Output Fields | 23409
- Sample Output | 23411
- Release Information | 23414

Syntax

```
show ptp statistics
<brief | detail>
```

Description

Display information about Precision Time Protocol (PTP) statistics.

Options

brief Display brief statistics about the operation of configured PTP clocks.

detail Display detailed statistics about the operation of configured PTP clocks.

Required Privilege Level

view

Output Fields

Table 1046 on page 23409 lists the output fields for the `show ptp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1046: show ptp statistics Output Fields

Field Name	Field Description
`Local Address`	IP address of the local PTP timeTransmitter and timeReceiver interfaces.

Table 1046: show ptp statistics Output Fields *(Continued)*

Field Name	Field Description
`Remote Address`	IP address of the remote PTP timeTransmitter and timeReceiver interfaces.
`Role`	Function performed by an Ethernet interface configured as a timeReceiver or timeTransmitter.
`Stream`	Stream ID uniquely identifies the connection between one timeTransmitter and one timeReceiver.
`Received`	1588v2 packets received by the timeTransmitter or timeReceiver interface. For the QFX Series, all packets transmitted by the timeTransmitter or timeReceiver interface.
`Transmitted`	1588v2 packets transmitted by the timeTransmitter or timeReceiver interface. For the QFX Series, all packets transmitted by the timeTransmitter or timeReceiver interface.
`Signalling`	Packet count for signalling messages: • Rx—Number of packets received. • Tx—Number of packets transmitted.
`Announce`	Packet count for announce messages: • Rx—Number of packets received. • Tx—Number of packets transmitted.
`Sync`	Packet count for synchronization messages: • Rx—Number of packets received. • Tx—Number of packets transmitted.

Table 1046: show ptp statistics Output Fields *(Continued)*

Field Name	Field Description
Delay	Packet count for delay request or response messages: • Rx—Number of packets received. • Tx—Number of packets transmitted.
Error	Packet count for signal loss errors: • Rx—Number of packets received with errors. • Tx—Number of packets transmitted with errors.

Sample Output

show ptp statistics

```
user@host> show ptp statistics
Local Address    Remote Address     Role  Stream      Received  Transmitted
2.2.2.2          10.10.20.50       Slave       0         45716        22826
6.6.6.2          6.6.6.1          Master       4         24960        74880
```

show ptp statistics

```
user@host> show ptp statistics
Local Clock      Remote Clock                              Role  Stream      Received
Transmitted
        ge-0/1/5.0       2001:cdba:0000:0000:0000:0000:3257:9653 Master      4
727205       345698
        ge-0/2/0.0       2001:cdba:0000:0000:0000:0000:3256:0101 Master      5
4493776      2222524
```

show ptp statistics detail

```
user@host> show ptp statistics detail
Local Address    Remote Address     Role  Stream      Received  Transmitted
2.2.2.2          10.10.20.50       Slave       0         47009        23470
         Signalling    Announce        Sync       Delay       Error
   Rx:            5         184       23399       23426           0
   Tx:           45           0           0       23426           0

6.6.6.2          6.6.6.1          Master       4         25600        76800
         Signalling    Announce        Sync       Delay       Error
   Rx:            0           0           0       25600           0
   Tx:            0       25600       25600       25600           0
```

show ptp statistics detail

```
user@host> show ptp statistics detail
Local Clock      Remote Clock       Role  Stream      Received  Transmitted
ge-0/1/5.0       2001:cdba:0000:0000:0000:0000:3257:9653  Master      4        727205
345698
         Signalling    Announce        Sync       Delay       Error
   Rx:            0       33978      347535      345692           6
   Tx:            0           0           0      345698           0
Local Clock      Remote Clock       Role  Stream      Received  Transmitted
ge-0/2/0.0       2001:cdba:0000:0000:0000:0000:3256:0101 Master      5       4493776      2222524
         Signalling    Announce        Sync       Delay       Error
   Rx:            0       36819     2234472     2222485          14
   Tx:            0           0           0     2222524           0
```

show ptp statistics)

```
user@host> show ptp statistics
Local Clock      Remote Clock       Role  Stream      Received  Transmitted
ae0.0            2003::a           Slave       0        666694       331379
lo0.0            2002::1          Master       4        696526      1405758
lo0.0            3001::1:133      Master       5        680312      1373318
```

show ptp statistics detail

```
user@host> show ptp statistics detail
Local Clock      Remote Clock       Role  Stream      Received  Transmitted
xe-0/0/6:0.0     224.0.1.129       Slave       0             0            0
         Signalling    Announce        Sync       Delay       Error
   Rx:            0        1688      217905      129833         131
   Tx:            0           0           0      130608           0


Local Clock      Remote Clock       Role  Stream      Received  Transmitted
xe-0/0/6:1.0     224.0.1.129      Master       0             0        80517
         Signalling    Announce        Sync       Delay       Error
   Rx:            0           0           0           0           0
   Tx:            0         625       79892           0           0

Local Clock      Remote Clock       Role  Stream      Received  Transmitted
xe-0/0/6:1.0     50.50.50.2       Master       0        249980       249980
         Signalling    Announce        Sync       Delay       Error
   Rx:            0           0           0      249980           0
   Tx:            0           0           0      249980           0

Local Clock      Remote Clock       Role  Stream      Received  Transmitted
xe-0/0/6:1.0     50.50.50.3       Master       0         18520        18600
         Signalling    Announce        Sync       Delay       Error
   Rx:            0           0           0       18520           0
   Tx:            0           0           0       18600           0

Local Clock      Remote Clock       Role  Stream      Received  Transmitted
xe-0/0/6:1.0     50.50.50.10       Master       0             0            0
         Signalling    Announce        Sync       Delay       Error
   Rx:            0           0           0           0           0
   Tx:            0           0           0           0           0
```

show ptp statistics detail

```
user@host> show ptp statistics detail
Local Clock      Remote Clock       Role  Stream      Received  Transmitted
ae0.0            2003::a           Slave       0        167235        82773
```

```
          Signalling    Announce        Sync       Delay       Error
   Rx:            15        1303       83174       82758           0
   Tx:            17           0           0       82761           0
Local Clock      Remote Clock       Role  Stream      Received  Transmitted
lo0.0            2002::1          Master       4         82152       167879
          Signalling    Announce        Sync       Delay       Error
   Rx:           132           0           0       82020           0
   Tx:           132        1320       84407       82020           0
Local Clock      Remote Clock       Role  Stream      Received  Transmitted
lo0.0            3001::1:133      Master       5           131         1450
          Signalling    Announce        Sync       Delay       Error
   Rx:           131           0           0           0           0
   Tx:           131        1319           0           0           0
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

IEEE 1588v2 PTP Boundary Clock Overview

IEEE 1588v2 Precision Timing Protocol (PTP)

show redundant-power-system led

IN THIS SECTION

- Syntax | 23415
- Description | 23415
- Required Privilege Level | 23415
- Output Fields | 23415
- Sample Output | 23416
- Release Information | 23417

Syntax

```
show redundant-power-system led
```

Description

Display information about fan status, Redundant Power System (RPS) status, and the switch connectors as displayed by the corresponding LEDs on the RPS.

Required Privilege Level

view

Output Fields

Table 1047 on page 23415 lists the output fields for the `show redundant-power-system led` command. Output fields are listed in the approximate order in which they appear.

Table 1047: show redundant-power-system led Output Fields

Field Name	Field Description	Level of Output
RPS	The serial number of the RPS.	
RPS Fan	Status of the RPS power supply fans as displayed by the LED: • Green—All RPS power supply fans are operating fine. • Amber—A fan has failed in at least one RPS power supply.	All levels
RPS System Status	Status of the RPS system as displayed by the LED: • Green—The RPS is active. • Blinking green—The RPS is booting. • Amber—An RPS power supply has failed. • Off—The RPS is off.	All levels

Table 1047: show redundant-power-system led Output Fields *(Continued)*

Field Name	Field Description	Level of Output
RPS Port LED Status	Status of the RPS switch connectors as displayed by the LEDs. These LEDs indicate whether the redundant power source is being used. • Green—The RPS connector is enabled and connected to a switch but the RPS is not actively backing up the switch. • Blinking green—The RPS is backing up the switch connected to the port. • Off—The RPS connector is not connected to a switch. • Amber—The RPS is oversubscribed and the backup power to the switch has failed.	All levels
Port	Number of one of the six switch connectors on the RPS.	All levels
Status	Status of each switch connector on the RPS.	All levels

Sample Output

show redundant-power-system led (Standalone Switch)

```
user@switch> show redundant-power-system led

Gathering requested information.
RPS-CG0209121807
    RPS Fan: GREEN
    RPS System Status: GREEN
RPS Port LED Status
Port  Status
   0  GREEN
   1  OFF
   2  OFF
   3  OFF
   4  OFF
   5  OFF
```

show redundant-power-system led (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system led

Gathering requested information.
RPS-CG0209121814
    RPS Fan: GREEN
    RPS System Status: GREEN
RPS Port LED Status
Port  Status
   0  OFF
   1  OFF
   2  OFF
   3  OFF
   4  OFF
   5  GREEN
RPS-CG0209121815
    RPS Fan: GREEN
    RPS System Status: GREEN
RPS Port LED Status
Port  Status
   0  OFF
   1  OFF
   2  OFF
   3  OFF
   4  GREEN
   5  OFF
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

LEDs on an EX Series Redundant Power System

show redundant-power-system multi-backup

IN THIS SECTION

- Syntax | 23418
- Description | 23418
- Required Privilege Level | 23418
- Sample Output | 23418
- Release Information | 23419

Syntax

```
show redundant-power-system multi-backup
```

```
show redundant-power-system multi-backup member member-number
```

Description

Display the current status of the Redundant Power System's (RPS's) ability to back up two switches per power supply when enough power to support Power over Ethernet (PoE) is not needed. This ability is referred to as the RPS's multi-backup ability.

Required Privilege Level

view

Sample Output

show redundant-power-system multi-backup

```
User@switch> show redundant-power-system multi-backup
Requesting information from redundant-power-system..     Multi-Backup: enabled
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

request redundant-power-system multi-backup | 14739

show redundant-power-system network

IN THIS SECTION

- Syntax | 23419
- Description | 23419
- Required Privilege Level | 23419
- Sample Output | 23420
- Release Information | 23420

Syntax

```
show redundant-power-system network
```

Description

Display the Redundant Power Supply (RPS) IP address, netmask address, and gateway address required for firmware backup.

Required Privilege Level

view

Sample Output

show redundant-power-system network

```
user@switch> show redundant-power-system network
 Requesting information from redundant-power-system..
    IP Address:   10.93.2.38
    Netmask: 255.255.254.0
    Gateway:  10.93.3.254
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Upgrading Firmware on an EX Series Redundant Power System

show redundant-power-system power-supply

IN THIS SECTION

- Syntax | 23421
- Description | 23421
- Required Privilege Level | 23421
- Output Fields | 23421
- Sample Output | 23422
- Release Information | 23422

Syntax

```
show redundant-power-system power-supply
```

Description

Display information about the power supplies installed in the Redundant Power System (RPS). After installing a power supply, we recommend that you use this command to be sure that the power supply installed correctly.

Required Privilege Level

view

Output Fields

Table 1048 on page 23421 lists the output fields for the `show redundant-power-system power-supply` command. Output fields are listed in the approximate order in which they appear.

Table 1048: show redundant-power-system power-supply Output Fields

Field Name	Field Description	Level of Output
RPS	Serial number of the RPS.	All levels
PSU Slot	Number of the power supply slot. Slots are numbered 1 through 3.	All levels
Status	Status of the power supply slots: • Present—The slot contains an RPS power supply. • Empty—The slot is empty.	All levels
Description	Description of the RPS power supply installed in the slot.	All levels

Sample Output

show redundant-power-system power-supply (Standalone Switch)

```
user@switch> show redundant-power-system power-supply

Gathering requested information.
RPS-CG0209121807
PSU Slot Status    Description
       1 Online    930W AC
       2 offline     ---
       3 Online    930W AC
```

show redundant-power-system power-supply (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system power-supply

Gathering requested information.
RPS-CG0209121814
PSU Slot Status    Description
       1 Online    930W AC
       2 offline     ---
       3 Online    930W AC
RPS-CG0209121815
PSU Slot Status    Description
       1 Online    930W AC
       2 Online    930W AC
       3 Online    930W AC
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Installing a Power Supply in the EX Series Redundant Power System

show redundant-power-system status

IN THIS SECTION

- Syntax | 23423
- Description | 23423
- Required Privilege Level | 23423
- Output Fields | 23423
- Sample Output | 23424
- Release Information | 23425

Syntax

```
show redundant-power-system status
```

Description

Display the status information for the switch connectors on the Redundant Power System (RPS).

Required Privilege Level

view

Output Fields

Table 1049 on page 23423 lists the output fields for the `show redundant-power-system status` command. Output fields are listed in the approximate order in which they appear.

Table 1049: show redundant-power-system status Output Fields

Field Name	Field Description	Level of Output
RPS	Serial number of the RPS.	All levels

Table 1049: show redundant-power-system status Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Port	Number of the switch connector.	All levels
Status	Status of the switch connector: • ARMED—The switch is ready to get backup power from RPS if power supply fails on the switch. • OFF—The switch has zero and is not configured to receive backup power from RPS. • BACKED-UP—The switch is receiving power backup from RPS. • OVER-SUBSCRIBED—The switch cannot receive backup power from RPS even if you set the .	All levels
Priority	Priority value of the switch connector.	All levels
Power-Requested	Power requested by the switch on the corresponding switch connector.	All levels

Sample Output

show redundant-power-system status (Standalone Switch)

```
user@switch> show redundant-power-system status

Gathering requested information.
RPS-CG0209121807
Port Status        Power-requested
  0  Armed        3         930W
  1  Off          1         ---
  2  Off          1         ---
  3  Off          1         ---
  4  Off          1         ---
  5  Off          1         ---
```

show redundant-power-system status (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system status

Gathering requested information.
RPS-CG0209121814
Port Status        Power-requested
  0  OFF          1         ---
  1  OFF          1         ---
  2  OFF          1         ---
  3  OFF          1         ---
  4  OFF          1         ---
  5  Armed        5         930W
RPS-CG0209121815
Port Status        Power-requested
  0  OFF          1         ---
  1  OFF          1         ---
  2  OFF          1         ---
  3  OFF          1         ---
  4  Armed        4         930W
  5  OFF          1         ---
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Determining and Setting Priority for Switches Connected to an EX Series RPS

Installing a Power Supply in the EX Series Redundant Power System

show redundant-power-system upgrade

IN THIS SECTION

- Syntax | 23426
- Description | 23426
- Required Privilege Level | 23426
- Output Fields | 23426
- Sample Output | 23427
- Release Information | 23427

Syntax

```
show redundant-power-system upgrade
```

Description

Display RPS firmware upgrade status (pass or fail), previous RPS firmware version, and current RPS firmware version.

Required Privilege Level

view

Output Fields

Table 1050 on page 23426 lists the output fields for the `show redundant-power-system status` command. Output fields are listed in the approximate order in which they appear.

Table 1050: show redundant-power-system upgrade Output Fields

Field Name	Field Description	Level of Output
Firmware Upgrade Status	Indicates whether the upgrade passed or failed	All levels

Table 1050: show redundant-power-system upgrade Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Previous Firmware Version	Firmware version before the upgrade	All levels
Current Firmware Version	Firmware version after the upgrade	

Sample Output

show redundant-power-system upgrade

```
user@switch> show redundant-power-system upgrade
Requesting information from redundant-power-system..
Firmware Upgrade Status:   Pass
Previous Firmware Version: 1.0
Current Firmware Version:  1.0
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

request redundant-power-system multi-backup | **14739**

show redundant-power-system version

IN THIS SECTION

- Syntax | **23428**
- Description | **23428**

- Required Privilege Level | 23428
- Output Fields | 23428
- Sample Output | 23429
- Release Information | 23429

Syntax

```
show redundant-power-system version
```

Description

Display version information about the Redundant Power System (RPS).

Required Privilege Level

view

Output Fields

Table 1051 on page 23428 lists the output fields for the `show redundant-power-system version` command. Output fields are listed in the approximate order in which they appear.

Table 1051: show redundant-power-system version Output Fields

Field Name	Field Description	Level of Output
RPS	Serial number of the RPS.	All levels
Model	Model name of the RPS.	All levels
RPS Firmware Version	Version number of the firmware installed on the RPS.	All levels
RPS U-Boot Version	Version of the bootup software installed on the RPS.	All levels

Sample Output

show redundant-power-system version (Standalone Switch)

```
user@switch> show redundant-power-system version

      RPS-CG0209121807
     Model: EX-PWR_RPS200
     RPS Firmware Version [1.0]
     RPS U-Boot   Version [1.1.6]
```

show redundant-power-system version (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system version

      RPS-CG0209121814
     Model: EX-PWR_RPS200
     RPS Firmware Version [1.0]
     RPS U-Boot   Version [1.1.6]
      RPS-CG0209121815
     Model: EX-PWR_RPS200
     RPS Firmware Version [1.0]
     RPS U-Boot   Version [1.1.6]
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Installing a Power Supply in the EX Series Redundant Power System
- Packing an EX Series Redundant Power System or Redundant Power System Components for Shipping

show redundant-trunk-group

IN THIS SECTION

- Syntax | 23430
- Description | 23430
- Options | 23430
- Required Privilege Level | 23430
- Output Fields | 23430
- Sample Output | 23431
- Release Information | 23431

Syntax

```
show redundant-trunk-group <group-name group-name>
```

Description

Display information about redundant trunk groups.

Options

group-name *group-name* Display information about the specified redundant trunk group.

Required Privilege Level

view

Output Fields

Table 1052 on page 23431 lists the output fields for the `show redundant-trunk-group` command. Output fields are listed in the approximate order in which they appear.

Table 1052: show redundant-trunk-group Output Fields

Field Name	Field Description
Group name	Name of the redundant trunk port group.
Interface	Name of an interface belonging to the trunk port group.
State	Operating state of the interface. • **Up** denotes the interface is up. • **Down** denotes the interface is down. • **Pri** denotes a primary interface. • **Act** denotes an active interface.
Time of last flap	Date and time at which the advertised link became unavailable, and then, available again.
Flap count	Total number of flaps since the last switch reboot.

Sample Output

show redundant-trunk-group group-name Group1

```
user@switch> show redundant-trunk-group group-name Group1

Group      Interface       State          Time of last flap        Flap
name                                                               Count

Group1     ge-0/0/45.0     UP/Pri/Act     Never                    0
           ge-0/0/47.0     UP             Never                    0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Example: Configuring Redundant Trunk Links for Faster Recovery on EX Series Switches
- Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support
- Understanding Redundant Trunk Links (Legacy RTG Configuration)

show remote-device-management service-devices

IN THIS SECTION

- Syntax | 23432
- Description | 23432
- Options | 23432
- Required Privilege Level | 23433
- Output Fields | 23433
- Sample Output | 23438
- Release Information | 23441

Syntax

```
show remote-device-management service-devices
name
<detail>
```

Description

Display information about all remote service devices or a specific remote service device.

Options

none Display summary information about all remote service devices.

detail Display detailed information about remote service devices.

name Name of the remote service device.

Required Privilege Level

view

Output Fields

Table 1053 on page 23433 lists the output fields for the `show remote-device-management service-devices` command. Output fields are listed in the approximate order in which they appear.

Table 1053: show remote-device-management service-devices Output Fields

Field Name	Field Description	Level of Output
`Device Name`	Name of the remote service device.	All levels
`State`	`State`—State of the remote device service management (RDSM) connection to the service device, `Up` or `Down`.	All levels
`Eligible Services`	Number of remote subscriber services that are eligible to be provisioned on the device. When a remote device is up without interruption, the `Eligible Services` and `Provisioned Services` counters should be equal. If the `Provisioned Services` counter is less than the `Eligible Services` counter, you may need to reconfigure the remote device to provision the outstanding eligible services.	All levels
`Provisioned Services`	Number of remote subscriber services that are provisioned on the device.	All levels
`Address`	IP address of the remote device used to configure the subscriber service; unique across all routing instances.	`detailextensive`
`Last State Change Time`	Timestamp when the RDSM connection to the remote device last changed state.	`detailextensive`

Table 1053: show remote-device-management service-devices Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Vlan Id List`	List of VLAN ranges and IDs that are served by the remote device and make up the access domain. The access domain corresponds to the set of subscriber-facing Layer 2 locations that map to the device.	`detailextensive`
`Dictionary File`	The absolute file path on the router for the vendor-specific dictionary that defines the set of NETCONF XML protocol commands required to provision, deprovision, and roll back a subscriber service on the remote device.	`detailextensive`
`Provisioning Method`	Attributes configured for the NETCONF XML Protocol method for provisioning and deprovisioning services on the remote device: • `User Name`—Name used to access the remote device during service management. • `Connection Retry Interval`—The interval between successive attempts to establish a NETCONF session with the remote device. • `Response Timeout`—Period during which the device must respond to an attempt to provision or deprovision a service. • `Response Timeout Count`—Number of consecutive response timeouts that occur before the BNG takes action. • `Bulk Interval`—Interval during which multiple services are provisioned or deprovisioned based on the assigned dictionary before the configuration is committed to the service device. • `Bulk Limit`—Maximum number of services provisioned or deprovisioned based on the assigned dictionary during the bulk interval—before the configuration is committed to the service device • `Reconfigure Bulk Limit`—Maximum number of services provisioned on the service device for the access domain when the device is reconfigured; number is based on the assigned dictionary before the configuration is committed to the service device • `Port`—TCP port number for the NETCONF protocol session.	`detailextensive`

Table 1053: show remote-device-management service-devices Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Round Trip Time (millisec)`	Aggregate duration for all dictionary RPCs required to provision or deprovision a service: `Minimum`, `Maximum`, `Average`, and `Last`. When the bulk limit is set to 1, these statistics are reported per remote service. When the bulk limit is greater than 1, the duration includes all services that are part of the bulk configuration.	`detailextensive`
`Bulk Count`	Number of services included in a bulk configuration: `Minimum`, `Maximum`, `Average`, and `Last`. The number is determined by the `bulk-limit` and `bulk-interval` configuration for the remote device at the `[edit system services remote-device-management service-device` *`device-name`* `provisioning-method netconf]` hierarchy level. When the bulk limit is set to 1, the count is 1 for all fields.	`detailextensive`
`Service Session ID`	ID number for the remote service session.	`extensive`
`Subscriber Session ID`	ID number for the subscriber session associated with the service session.	`extensive`

Table 1053: show remote-device-management service-devices Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`State`	Status of the service session; represents the RDSM processing state: • `Active`—Remote service session is up and fully provisioned. • `Deprovisioning`—Remote service session is in the process of being deprovisioned. • `Init`—Transient state that indicates that the state is transitioning to the provisioning state. • `Provisioning`—Remote service session is in the process of being provisioned. • `Provisioning-Complete`—Transient state that indicates service provisioning has completed on the device; transitions to Active or Provisioning-Rollback. • `Provisioning-Rollback`—Rollback of the service provisioning on the device is in progress. Provisioning is rolled back when provisioning fails for one or more other eligible service devices.	`extensive`
`Service Name`	Name of the remote service.	`extensive`

Table 1053: show remote-device-management service-devices Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Service Provisioning`	Status counts for all provisioning actions attempted for the remote device.
	• `Attempted`—Number of service provisioning attempts for a service device in the `Up` state.
	• `Succeeded`—Number of services successfully provisioned.
	• `Failed`—Number of services that failed provisioning for all reasons
	  • `Reject`—Number of failures due to an explicit error response during provisioning.
	  • `Timeout`—Number of failures due to no response during provisioning.
	• `In Progress`—Number of services being provisioned but not yet completed.
	• `Reconfigure Pending`—Number of services that are not configured but are pending reconfiguration. A nonzero value implies that you must initiate a reconfiguration request to provision these services.
	• `Queued`—Number of services for which provisioning has not yet been attempted because another service is in the process of being provisioned or deprovisioned.
	The typical case is when remote device reconfiguration is followed immediately by a new subscriber service provisioning action. In this instance, reconfiguration is allowed to complete, which delays action on the new subscriber service.
		`extensive`

Table 1053: show remote-device-management service-devices Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Service Deprovisioning	Status counts for all deprovisioning actions attempted for the remote device.
	• Attempted—Number of service deprovisioning attempts for a service device in the Up state.
	• Succeeded—Number of services successfully deprovisioned.
	• Failed—Number of services that failed deprovisioning for all reasons
	  • Reject—Number of failures due to an explicit error response during deprovisioning.
	  • Timeout—Number of failures due to no response during deprovisioning.
	• In Progress—Number of services being deprovisioned but not yet completed.
	• Queued—Number of services for which deprovisioning has not yet been attempted because another service is in the process of being provisioned or deprovisioned.
		extensive

Reconfiguration Requests	Status counts of reconfiguration requests: Received, Succeeded, Failed, and Pending. A reconfiguration request initiated while one is in-progress is rejected.	extensive

Sample Output

show remote-device-management service-devices

```
user@host> show remote-device-management service-devices
Device Name     State   Eligible Services     Provisioned Services
Olt-xyz         Up       2                    2
Olt-1           Up       3                    1
Olt-2           Down     2                    0
```

show remote-device-management service-devices (Device Name)

```
user@host> show remote-device-management service-devices olt-xyz
Device Name     State   Eligible Services     Provisioned Services
Olt-xyz         Up      2                     2
```

show remote-device-management service-devices (Device Name Detail)

```
user@host> show remote-device-management service-devices olt-xyz detail
Device Name: olt-xyz
State: Up
Address: 10.2.3.1
Routing Instance: default
Last State Change Time: Wed Apr 11 07:24:02 2018
Vlan Id List: 1-40 71-80
Dictionary File: /var/home/dict/dictionary-1.xml
Eligible Services: 1
Provisioned Services: 1

Provisioning Method: netconf
    User Name: regress
    Connection Retry Interval: 3
    Response Timeout: 10
    Response Timeout Count: 3
    Bulk Interval: 1000
    Bulk Limit: 1
    Reconfigure Bulk Limit: 100
    Port: 830

Round Trip Time (millisec)
Minimum      Maximum      Average      Last
1681         1681         1681         1681

Bulk Count
Minimum      Maximum      Average      Last
1            1            1            1
```

show remote-device-management service-devices (Device Name Extensive)

```
user@host> show remote-device-management service-devices olt-xyz extensive
Device Name: olt-xyz
State: Up
Address: 10.2.3.1
Routing Instance: default
Last State Change Time: Wed Apr 11 07:24:02 2018
Vlan Id List: 1-40 71-80
Dictionary File: /var/home/dict/dictionary-1.xml
Eligible Services: 1
Provisioned Services: 1

Provisioning Method: netconf
    User Name: regress
    Connection Retry Interval: 3
    Response Timeout: 10
    Response Timeout Count: 3
    Bulk Interval: 1000
    Bulk Limit: 1
    Reconfigure Bulk Limit: 100
    Port: 830

Round Trip Time (millisec)
Minimum      Maximum      Average      Last
1681         1681         1681         1681

Bulk Count
Minimum      Maximum      Average      Last
1            1            1            1

Service Sessions
Service Session ID    Subscriber Session ID State                    Service Name
3                     2                      Active                  s1

Service Provisioning
   Attempted: 1
   Succeeded: 1
   Failed: 0
      Reject: 0
      Timeout: 0
   In Progress: 0
```

```
    Reconfigure Pending: 0
    Queued: 0

 Service De-provisioning
    Attempted: 0
    Succeeded: 0
    Failed: 0
       Reject: 0
       Timeout: 0
    In Progress: 0
    Queued: 0

 Reconfiguration Requests
    Received: 0
    Succeeded: 0
    Failed: 0
    Pending: 0
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

- show remote-device-management services | 23441
- show remote-device-management statistics | 23445
- show remote-device-management subscribers | 23451
- show remote-device-management summary | 23455

show remote-device-management services

IN THIS SECTION

- Syntax | 23442
- Description | 23442

- Options | 23442
- Required Privilege Level | 23442
- Output Fields | 23442
- Sample Output | 23444
- Release Information | 23445

Syntax

```
show remote-device-management services
session-id id-number
<detail>
```

Description

Display information about all service sessions or a specific service session on remote service devices.

Options

none	Display summary information about all remote service sessions.
detail	Display detailed information about remote service sessions.
session-id *id-number*	Identification number for the service session.

Required Privilege Level

view

Output Fields

Table 1054 on page 23443 lists the output fields for the `show remote-device-management services` command. Output fields are listed in the approximate order in which they appear.

Table 1054: show remote-device-management services Output Fields

Field Name	Field Description	Level of Output
`Service Session ID`	ID number for the remote service session.	All levels
`Service State`	Status of the service session; represents the RDSM processing state: • `Active`—Remote service session is up and fully provisioned. • `Deprovisioning`—Remote service session is in the process of being deprovisioned. • `Init`—Transient state that indicates that the state is transitioning to the provisioning state. • `Provisioning`—Remote service session is in the process of being provisioned. • `Provisioning-Complete`—Transient state that indicates service provisioning has completed on the device; transitions to Active or Provisioning-Rollback. • `Provisioning-Rollback`—Rollback of the service provisioning on the device is in progress. Provisioning is rolled back when provisioning fails for one or more other eligible service devices.	All levels
`Service Name`	Name of the remote service.	All levels
`Subscriber Session ID`	ID number for the subscriber session associated with the service session.	All levels

Table 1054: show remote-device-management services Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Service Devices`	Detailed status information about each remote device eligible to be provisioned with the named remote service:
	• `Name`—Name of the service device.
	• `State`—State of the remote device service management (RDSM) connection to the service device, `Up` or `Down`.
	• `Provisioned`—State of the remote service session provisioning:
	  • `Yes`—Service has been successfully provisioned for the service device.
	  • `No`—Service provisioning is in progress or has not been performed. Provisioning might not have been performed because the remote device transitioned from Down to Up and requires reconfiguration. When a remote device is declared Down, it is no longer provisioned and must be reconfigured.
		`detail`

Sample Output

show remote-device-management services

```
user@host> show remote-device-management services
Service Session ID   Subscriber Session ID State                   Service Name
1234                 111                   Active                  s1
222                  45                    Provisioning            s1
555                  100                   Deprovisioning          s2
```

show remote-device-management services (Service Session)

```
user@host> show remote-device-management services session-id 1234
Service Session ID   Subscriber Session ID State                   Service Name
1234                 111                   Active                  s1
```

show remote-device-management services (Service Session Detail)

```
user@host> show remote-device-management services session-id 1234 detail
Service Session
Service Session ID: 1234
Service State: Active
Service Name: SUBSCRIBER-UPSTREAM-PROFILE
Subscriber Session ID: 111

Service Devices
Device Name                 State      Provisioned
olt-xyz                     Up         Yes
olt-1                       Up         No
olt-3                       Down       No
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

- show remote-device-management service-devices | 23432
- show remote-device-management statistics | 23445
- show remote-device-management subscribers | 23451
- show remote-device-management summary | 23455

show remote-device-management statistics

IN THIS SECTION

- Syntax | 23446
- Description | 23446
- Options | 23446
- Required Privilege Level | 23446

- Output Fields | 23446
- Sample Output | 23449
- Release Information | 23450

Syntax

```
show remote-device-management statistics
(summary | service-devices device-name)
```

Description

Display a global summary of service statistics for all remote devices or detailed statistics for a specific remote service device.

Options

service-devices *device-name*	Display statistics for the specified service device.
summary	Display a summary of service statistics for all remote devices.

Required Privilege Level

view

Output Fields

Table 1055 on page 23447 lists the output fields for the `show remote-device-management statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1055: show remote-device-management statistics Output Fields

Field Name	Field Description
`Service Activations`	Status of service profile instantiation requests from authd to provision remote services. Each service activation results in a service provisioning request to each remote device. The values displayed are the totals for all remote devices. • `Received`—Number of service provisioning requests received by RDSM. • `Acked`—Number of service provisioning requests acknowledged by RDSM. • `Nacked`—Number of service provisioning requests not acknowledged by RDSM. • `In progress`—Number of service provisioning requests currently in progress on remote devices.
`Service Deactivations`	Status of service profile deinstantiation request from authd to provision remote services. Each service deactivation results in a service deprovisioning request to each remote device. The values displayed are the totals for all remote devices. • `Received`—Number of service deprovisioning requests received by RDSM. • `Acked`—Number of service deprovisioning requests acknowledged by RDSM. • `Nacked`—Number of service deprovisioning requests not acknowledged by RDSM. • `In progress`—Number of service deprovisioning requests currently in progress on remote devices.
`Round Trip Time (millisec)`	Aggregate duration for all dictionary RPCs required to provision or deprovision a service: `Minimum`, `Maximum`, `Average`, and `Last`. When the bulk limit is set to 1, these statistics are reported per remote service. When the bulk limit is greater than 1, the duration includes all services that are part of the bulk configuration.
`Bulk Count`	Number of services included in a bulk configuration: `Minimum`, `Maximum`, `Average`, and `Last`. The number is determined by the `bulk-limit` and `bulk-interval` configuration for the remote device at the `[edit system services remote-device-management service-device` *`device-name`* `provisioning-method netconf]` hierarchy level. When the bulk limit is set to 1, the count is 1 for all fields.

Table 1055: show remote-device-management statistics Output Fields *(Continued)*

Field Name	Field Description

`Service Provisioning`	Status counts for all provisioning actions attempted for the remote device.
	• `Attempted`—Number of service provisioning attempts for a service device in the `Up` state.
	• `Succeeded`—Number of services successfully provisioned.
	• `Failed`—Number of services that failed provisioning for all reasons
	  • `Reject`—Number of failures due to an explicit error response during provisioning.
	  • `Timeout`—Number of failures due to no response during provisioning.
	• `In Progress`—Number of services being provisioned but not yet completed.
	• `Reconfigure Pending`—Number of services that are not configured but are pending reconfiguration. A nonzero value implies that you must initiate a reconfiguration request to provision these services.
	• `Queued`—Number of services for which provisioning has not yet been attempted because another service is in the process of being provisioned or deprovisioned.
	The typical case is when remote device reconfiguration is followed immediately by a new subscriber service provisioning action. In this instance, reconfiguration is allowed to complete, which delays action on the new subscriber service.

`Service Deprovisioning`	Status counts for all deprovisioning actions attempted for the remote device.
	• `Attempted`—Number of service deprovisioning attempts for a service device in the `Up` state.
	• `Succeeded`—Number of services successfully deprovisioned.
	• `Failed`—Number of services that failed deprovisioning for all reasons
	  • `Reject`—Number of failures due to an explicit error response during deprovisioning.
	  • `Timeout`—Number of failures due to no response during deprovisioning.
	• `In Progress`—Number of services being deprovisioned but not yet completed.
	• `Queued`—Number of services for which deprovisioning has not yet been attempted because another service is in the process of being provisioned or deprovisioned.

Table 1055: show remote-device-management statistics Output Fields *(Continued)*

Field Name	Field Description
Reconfiguration Requests	Status counts of reconfiguration requests: Received, Succeeded, Failed, and Pending. A reconfiguration request initiated while one is in-progress is rejected.

Sample Output

show remote-device-management statistics (Summary)

```
user@host> show remote-device-management statistics summary
Service Activations
   Received: 10
   Acked: 7
   Nacked: 2
   In Progress: 1

Service Deactivations
   Received: 4
   Acked: 3
   Nacked: 1
   In Progress: 0
```

show remote-device-management statistics (Service Devices)

```
user@host> show remote-device-management statistics service-devices olt-xyz

Round Trip Time (millisec)
    Minimum     Maximum     Average     Last
    1000        2000        1300        1500

Bulk Count
    Minimum     Maximum     Average     Last
    1           5           2           3

Service Provisioning
    Attempted: 2
```

```
    Succeeded: 1
    Failed: 0
        Reject: 0
        Timeout: 0
    In Progress: 1
    Reconfigure Pending: 1
    Queued: 0

Service Deprovisioning
    Attempted: 0
    Succeeded: 0
    Failed: 0
        Reject: 0
        Timeout: 0
    In Progress: 0
    Queued: 0

Reconfiguration Requests
    Received: 2
    Succeeded: 1
    Failed: 0
    Pending: 1
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

- show remote-device-management service-devices | 23432
- show remote-device-management services | 23441
- show remote-device-management subscribers | 23451
- show remote-device-management summary | 23455

show remote-device-management subscribers

IN THIS SECTION

- Syntax | 23451
- Description | 23451
- Options | 23451
- Required Privilege Level | 23451
- Output Fields | 23452
- Sample Output | 23453
- Release Information | 23454

Syntax

```
show remote-device-management subscribers
session-id id-number
<detail>
```

Description

Display information about service sessions for all subscriber sessions or about all service sessions for a specific subscriber session on remote service devices.

Options

none	Display summary information about remote service sessions for all subscribers.
detail	Display detailed information about remote service sessions for all subscribers.
session-id *id-number*	Identification number for the subscriber session.

Required Privilege Level

view

Output Fields

Table 1056 on page 23452 lists the output fields for the `show remote-device-management subscribers` command. Output fields are listed in the approximate order in which they appear.

Table 1056: show remote-device-management subscribers Output Fields

Field Name	Field Description	Level of Output
`Subscriber Session ID`	ID number for the subscriber session associated with the service session.	All levels
`Service Session ID`	ID number for the remote service session.	All levels
`State`	Status of the service session; represents the RDSM processing state: • `Active`—Remote service session is up and fully provisioned. • `Deprovisioning`—Remote service session is in the process of being deprovisioned. • `Init`—Transient state that indicates that the state is transitioning to the provisioning state. • `Provisioning`—Remote service session is in the process of being provisioned. • `Provisioning-Complete`—Transient state that indicates service provisioning has completed on the device; transitions to Active or Provisioning-Rollback. • `Provisioning-Rollback`—Rollback of the service provisioning on the device is in progress. Provisioning is rolled back when provisioning fails for one or more other eligible service devices.	All levels
`Service Name`	Name of the remote service.	All levels

Table 1056: show remote-device-management subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Service Devices	Detailed status information about each remote device eligible to be provisioned with the named remote service:
	• Name—Name of the service device.
	• State—State of the remote device service management (RDSM) connection to the service device, Up or Down.
	• Provisioned—State of the remote service session provisioning:
	  • Yes—Service has been successfully provisioned for the service device.
	  • No—Service provisioning is in progress or has not been performed. Provisioning might not have been performed because the remote device transitioned from Down to Up and requires reconfiguration. When a remote device is declared Down, it is no longer provisioned and must be reconfigured.
		detail

Sample Output

show remote-device-management subscribers

```
user@host> show remote-device-management subscribers
Subscriber Session ID   Service Session ID State                   Service Name
111                     1234               Active                  s1
1238                    111                Active                  s112
222                     45                 Provisioning            s1
555                     100                Deprovisioning          s2
```

show remote-device-management subscribers (Subscriber Session)

```
user@host> show remote-device-management subscribers session-id 111
Subscriber Session ID   Service Session ID State                   Service Name
111                     1234               Active                  s1
111                     1238               Active                  s112
```

show remote-device-management subscribers (Subscriber Session Detail)

```
user@host> show remote-device-management subscribers session-id 111 detail
Subscriber Session ID: 111

Service Session
Service Session ID: 1234
Service State: Active
Service Name: s1
Subscriber Session ID: 111

Service Devices
Device Name                 State      Provisioned
olt-xyz                     Up         Yes
olt-1                       Up         No
olt-3                       Down       No


Service Session
Service Session ID: 1238
Service State: Active
Service Name: s112
Subscriber Session ID: 111

Service Devices
Device Name                 State      Provisioned
olt-xyz                     Up         Yes
olt-1                       Up         No
olt-3                       Down       No
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

clear remote-device-management statistics | 3837

show remote-device-management service-devices | 23432

show remote-device-management services | 23441

show remote-device-management statistics | 23445
show remote-device-management summary | 23455

show remote-device-management summary

IN THIS SECTION

- Syntax | 23455
- Description | 23455
- Required Privilege Level | 23455
- Output Fields | 23455
- Sample Output | 23457
- Release Information | 23458

Syntax

```
show remote-device-management summary;
```

Description

Display summary information about the remote service devices, such as session state and service state.

Required Privilege Level

view

Output Fields

Table 1057 on page 23456 lists the output fields for the `show remote-device-management summary` command. Output fields are listed in the approximate order in which they appear.

Table 1057: show remote-device-management summary Output Fields

Field Name	Field Description
`Service Devices by Connection State`	Number of services devices where the remote device service management (RDSM) connection to the remote device is in the `Up` or `Down` state.
`Service Sessions by State`	Number of service sessions in each of the following states, which represent the RDSM processing state: • `Active`—Remote service session is up and fully provisioned. • `Deprovisioning`—Remote service session is in the process of being deprovisioned. • `Init`—Transient state that indicates that the state is transitioning to the provisioning state. • `Provisioning`—Remote service session is in the process of being provisioned. • `Provisioning-Complete`—Transient state that indicates service provisioning has completed on the device; transitions to Active or Provisioning-Rollback. • `Provisioning-Rollback`—Rollback of the service provisioning on the device is in progress. Provisioning is rolled back when provisioning fails for one or more other eligible service devices.
`Service Activations`	Status of service profile instantiation requests from authd to provision remote services. Each service activation results in a service provisioning request to each remote device. The values displayed are the totals for all remote devices. • `Received`—Number of service provisioning requests received by remote devices. • `Acked`—Number of service provisioning requests acknowledged by remote devices. • `Nacked`—Number of service provisioning requests not acknowledged by remote devices. • `In progress`—Number of service provisioning requests currently in progress on remote devices.

Table 1057: show remote-device-management summary Output Fields *(Continued)*

Field Name	Field Description
`Service Deactivations`	Status of service profile deinstantiation request from authd to provision remote services. Each service deactivation results in a service deprovisioning request to each remote device. The values displayed are the totals for all remote devices. • `Received`—Number of service deprovisioning requests received by remote devices. • `Acked`—Number of service deprovisioning requests acknowledged by remote devices. • `Nacked`—Number of service deprovisioning requests not acknowledged by remote devices. • `In progress`—Number of service deprovisioning requests currently in progress on remote devices.

Sample Output

show remote-device-management summary

```
user@host> show remote-device-management summary
Service Devices by Connection State
   Up: 2
   Down: 1

Service Sessions by State
   Active: 2
   Provisioning:  1
   Deprovisioning: 0

Service Activations
   Received: 10
   Acked: 7
   Nacked: 2
   In progress: 1

Service De-activations
```

```
    Received: 4
    Acked: 3
    Nacked: 1
    In progress: 0
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

- show remote-device-management service-devices | 23432
- show remote-device-management services | 23441
- show remote-device-management statistics | 23445
- show remote-device-management subscribers | 23451

show resource-manager

IN THIS SECTION

- Syntax | 23458
- Description | 23459
- Required Privilege Level | 23459
- Sample Output | 23459
- Release Information | 23459

Syntax

```
show resource-manager
```

Description

Display the CPU and memory utilization.

Required Privilege Level

view

Sample Output

show resource-manager

```
user@host> show resource-manager
FPC 0 PIC 0

Services Memory: 0 %

Packet Memory: 0 %

Session created: 631829

THREAD 0

Last Minute CPU Average %: 0
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Traffic Processing on SRX Series Devices Overview

show resource-manager cpu

IN THIS SECTION

- Syntax | 23460
- Description | 23460
- Required Privilege Level | 23460
- Sample Output | 23460
- Release Information | 23461

Syntax

```
show resource-manager cpu
```

Description

This command displays the CPU utilization.

Required Privilege Level

view

Sample Output

show resource-manager cpu

```
user@host> show resource-manager cpu
fpc 0 pic 0

CPU Threshold exceeded :0

CPU threshold Clear :0

Drop flows :0
```

```
Throttled Sessions :0

Packets Dropped :0

Alarm Sent :0

Messages logged :0
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Traffic Processing on SRX Series Devices Overview

show resource-manager memory

IN THIS SECTION

- Syntax | 23461
- Description | 23462
- Required Privilege Level | 23462
- Sample Output | 23462
- Release Information | 23462

Syntax

```
show resource-manager memory
```

Description

Display the CPU and memory utilization.

Required Privilege Level

view

Sample Output

show resource-manager memory

```
user@host> show resource-manager memory
fpc 0 pic 0

Session Low Watermark Hit :0

Session High Watermark Hit :0

Memory Threshold exceeded :15

Memory threshold Clear :14

Early Session Aged Out :0

Drop flows :0

Fail Open :0

Fail Close :167814

Alarm Sent :15

Messages logged :29
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Traffic Processing on SRX Series Devices Overview

show rift database

IN THIS SECTION

- Syntax | 23463
- Description | 23463
- Options | 23463
- Required Privilege Level | 23464
- Output Fields | 23464
- Sample Output | 23465
- Sample Output | 23466
- Release Information | 23466

Syntax

```
show rift database
content
statistics
```

Description

Display Routing in Fat Trees (RIFT) protocol link-state database information. You can use this information for debugging purpose.

Options

content Display RIFT link-state database topology information element (TIE) headers information.

statistics Display RIFT link-state database information.

Required Privilege Level

view

Output Fields

Table 1058 on page 23464 lists the output fields for the `show rift database content` command. Output fields are listed in the approximate order in which they appear.

Table 1058: show rift database content Output Fields

Field Name	Field Description
Dir	Direction of Topology information element (TIE).
Originator	Originating node.
Type	Type of TIE.
ID	TIE identifier.
SeqNr	TIE sequence number.
Lifetime	Remaining TIE lifetime in seconds.
Origin Creation Time	Time of TIE creation at the origin in seconds.
Origin Lifetime	Lifetime of TIE when created at the origin.
Content Size	Size of TIE as exchanged on the link.
Key ID	Displays whether the TIE has an inner security key, that is, carries originator's authentication.

Table 1059 on page 23465 lists the output fields for the `show rift database statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1059: show rift database statistics Output Fields

Field Name	Field Description
Peers	Configured peers.
in 3-WAY	Number of peers in three way.
Last UP/DOWN	Date and time peer was last UP or DOWN.
Last New TIE	Identification of last new TIE.
on	Date and time of the last new TIE.
TIE Version Collisions	Number of TIEs database received of same or older version.
Last Southbound Routes	Last time southbound routes were computed .
Dir	Direction of TIE.
Type	Type of TIE.
#TIES	Number of TIEs.

Sample Output

show rift database content

```
user@host> show rift database content
Dir Originator Type ID SeqNr Lifetime Origin Creation Time
 Origin Content Key ID
---+----------------+---------+--------+-------+--------+----------------------+-Lifetime-+--
Size-+--------
S 0000000000000001 Node 10000000 2829 601200 2019/11/08 22:28:42
 604800 805 0
S 0000000000000001 Node 10000001 21020 601200 2019/11/08 22:28:42
```

```
 604800 998 0
S 0000000000000001 Node 10000002 2340 601200 2019/11/08 22:28:42
 604800 862 0
S 0000000000000001 Node 10000003 12074 601200 2019/11/08 22:28:42
 604800 862 0
```

Sample Output

show rift database statistics

```
user@host> show rift database statistics
Peers: Configured 4, in 3-WAY 4, Last UP/DOWN 2019/06/06 16:04:36.354
Last New TIE: 00002c6bf5586fc0/S/Positive/30000008, on 2019/06/06 16:08:56.187
TIE Version Collisions: 28
Last Southbound Routes: 2019/06/06 16:02:31.407

Dir   Type      #TIES
-----+---------+-----+
South
      External      3
      Node          6
      Prefix        2
      Positive      1
```

Release Information

Command introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | **15794**

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

show rift flood-reduction

IN THIS SECTION

- Syntax | 23467
- Description | 23467
- Options | 23467
- Required Privilege Level | 23467
- Output Fields | 23467
- Sample Output | 23468
- Release Information | 23468

Syntax

```
show rift flood-reduction
statistics
```

Description

Display Routing in Fat Trees (RIFT) protocol flood-reduction information.

Options

statistics	Display RIFT flood-reduction statistics.

Required Privilege Level

view

Output Fields

Table 1060 on page 23468 lists the output fields for the `show rift flood-reduction statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1060: show rift flood-reduction statistics Output Fields

Field Name	Field Description
Runs	Number of runs.
Holds	Number of computations held down.
Leaders	Chosen flood leaders.
Last On	Date and time of last election.

Sample Output

show rift flood-reduction statistics

```
user@host> show rift flood-reduction statistics
Runs  Holds Leaders                                                          Last On
-----+-----+-----------------------------------------------------------------
+-------------------------
  110   123                                                                    2019/06/06
16:08:01.189
```

Release Information

Command introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

show rift interface

IN THIS SECTION

- Syntax | 23469
- Description | 23469
- Options | 23469
- Required Privilege Level | 23469
- Output Fields | 23470
- Sample Output | 23475
- Sample Output | 23475
- Release Information | 23476

Syntax

```
show rift interface
statistics
status
```

Description

Display Routing in Fat Trees (RIFT) protocol interface information.

Options

statistics	Display RIFT interface statistics.
status	Display status of RIFT interfaces.

Required Privilege Level

view

Output Fields

Table 1061 on page 23470 lists the output fields for the `show rift interface statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1061: show rift interface statistics Output Fields

Field Name	Field Description
Started	Date and time when the statistics was started.
TIES RX	Number of topology information elements (TIEs) received.
TX	Number of TIEs transmitted.
REQ	Number of TIEs requested.
Neighbor REQ	Number of TIEs requested by neighbor.
TIDES TX	Number of topology information description elements (TIDEs) transmitted.
TIRES TX	Number of topology information request elements (TIREs) transmitted.
TIRES RX	Number of TIREs received.
Pkt Rate/100msecs	Packet rate per 100 milliseconds: • Highest—Highest achieved flooding rate. • Current—Current flooding rate n packets per 100 milliseconds. • Packet Sequence Losses—Sequence loses in flooding leading to flood rate adaptations.
Last TIE RX	Last received TIE.
Last TIE RX on	Time at which last TIE was received.

Table 1061: show rift interface statistics Output Fields *(Continued)*

Field Name	Field Description
Last Newer TIE RX	Last newer received TIE.
Last Newer TIE RX on	Time at which last newer TIE was received.
Last TIE TX	Last transmitted TIE.
Last TIE TX on	Time at which last TIE was transmitted.
Last Newer TIE RX	Last newer received TIE.
Last Newer TIE RX on	Time at which last newer TIE was received.
Last TIE TX	Last transmitted TIE.
Last TIE TX on	Time at which last TIE was transmitted.
TIE TX Queue Len	TIE transmission queue length.
Last TIE TX Queued	Last TIE queued for transmission.
on	Event leading to queuing of last TIE.
Largest TX'ed - TIE/TIDE/TIRE	Largest transmitted size per packet type.
Three-Way UP	Number of times adjacency came up.
DOWN	Number of times adjacency came down.
Last UP	Date and time the last adjacency came up.

Table 1061: show rift interface statistics Output Fields *(Continued)*

Field Name	Field Description
Last DOWN	Date and time the last adjacency went down.
Last Reason DOWN	Reason why the last adjacency went down.
LIE TX	Number of link information elements (LIEs) transmitted.
RX	Number of LIEs received.
Corrupt	Number of corrupt LIEs received.
Last LIE TX	Date and time the last LIE was transmitted.
Largest TX'ed	Size of largest transmitted TIE, TIDE, and TIRE.
Reject Reason	Reason for rejecting the last LIE.
Current Level Self/Neighbor	Current level of device (self) or the neighboring device.
Level Changes Self/Neighbor	Number of level changes for the device (self) or the neighboring device.
Flood Leader	Role as flood leader as designated by the neighboring device.
Changes	Number of changes in the flood relationship.
Last Change	Date and time the flood leadership changed last.

Table 1062 on page 23473 lists the output fields for the `show rift interface status` command. Output fields are listed in the approximate order in which they appear.

Table 1062: show rift interface status Output Fields

Field Name	Field Description
Link ID	Link identifier of the local RIFT interface.
Interface	Interface name.
Status Admin	Administrative status of the interface. • True—Status is up. • False—Status is down.
Platform	Carrier status of interface. • True—Status is up. • False—Status is down.
BFD	BFD status of interface if BFD is negotiated. • True—Status is up. • False—Status is down.
State	State of interface: • one-way • two-way • three-way
Uptime	Duration interface was in three-way state.
LIE TX V4	Number of transmitted IPv4 LIEs.
LIE TX V6	Number of transmitted IPv6 LIEs.

Table 1062: show rift interface status Output Fields *(Continued)*

Field Name	Field Description
LIE TX Port	Port on which LIEs are transmitted.
TIE RX Port	Port on which LIEs are received.
PoD	Point of delivery.
Nonce	Local security nonce.
Neighbor	If in three-way state, describes neighbor with HoldTime and Outer security key.
Link ID	Link ID of neighbor.
Name	Link name of neighbor.
Level	Level of neighbor.
TIE V4	IPv4 addresses on which TIEs are received.
TIE V6	IPv6 addresses on which TIEs are received.
TIE Port	Port on which TIEs are received.
BW	Bandwidth of the interface.
Outer Key	Outer security key ID.
Holdtime	Adjacency hold time in 3-way.

Sample Output

show rift interface

```
user@host> show rift interface statistics
Link ID: 257, Interface: ge-0/0/1.0, Started: 2019/11/09 11:48:25.760
TIES RX: 65, TX: 6198, REQ: 1, Neighbor REQ: 2445
Same TIE RX: 5, TIES RE-TX: 714, Own TIES RX'ed: 6
TIDES TX: 104827, RX: 5442, TIRES TX: 33, TIRES RX: 1601
Pkt Rate/100msecs Highest: 50, Current: 50, Packet Sequence Losses: 22
Last TIE RX: 0000000000000002/S/Node____/10000003, Last TIE RX on: 2019/11/09
15:17:18.225
Last Newer TIE RX: 0000000000000002/S/Node____/10000000, Last Newer TIE RX on:
2019/11/09 15:06:54.486
Last TIE TX: 0000000000000001/S/PosExt__/700000b2, Last TIE TX on: 2019/11/09
15:17:23.747
TIE TX Queue Len: 0, Last TIE TX Queued: 0000000000000001/S/PosExt__/700001ff, on:
 NewerOnTIDE
Last TIE REQ'ed: 0000000000000001/S/PosExt__/700000c5, Last TIE REQ on: 2019/11/09
 15:07:00.033
Largest TX'ed: TIE: 1018, (0000000000000001/S/Node____/10000001), TIDE: 1192, TIRE:
 272
Three-Way UP 4, DOWN 3, Last UP 2019/11/09 13:54:22.950, Last DOWN 2019/11/09
13:54:19.777
Last Reason DOWN: HoldtimeExpired
LIE TX 6346, RX 6178
Last LIE TX 2019/11/09 15:19:54.824, RX 2019/11/09 15:19:53.715, Reject Reason:
None
Current Level Self 24, Neighbor 23, Level Changes Self 1, Neighbor 1
Flood Leader: False, Changes: 2, Last Change: 2019/11/09 11:48:26.049
```

Sample Output

show rift interface status

```
user@host> show rift interface status
Link ID: 258, Interface: ge-0/0/0.1
Status Admin True, Platform True, BFD True, State: ThreeWay, 3-Way Uptime: 3 hours,
 34 minutes, 18 seconds
LIE TX V4: 224.0.0.120, LIE TX V6: ff02::a1f7, LIE TX Port: 914, TIE RX Port: 915
```

```
PoD 0, Nonce 11589
Neighbor: ID 000000000000000a, Link ID 258, Name: rift 10:ge-0/0/0.1, Level: 23
TIE V4: 1.1.10.2, TIE V6: fe80::5668:a300:114:2212, TIE Port: 915, BW 1000 MBits/s
PoD: None, Nonce: 17251, Outer Key: 0, Holdtime: 3 secs
```

Release Information

Command introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

show rift node

IN THIS SECTION

- Syntax | 23476
- Description | 23477
- Options | 23477
- Required Privilege Level | 23477
- Output Fields | 23477
- Sample Output | 23479
- Sample Output | 23479
- Release Information | 23479

Syntax

```
show rift node
statistics
status
```

Description

Display Routing in Fat Trees (RIFT) protocol node information.

Options

statistics Display RIFT node statistics.

status Display RIFT node status.

Required Privilege Level

view

Output Fields

Table 1063 on page 23477 lists the output fields for the `show rift node statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1063: show rift node statistics Output Fields

Field Name	Field Description
Starttime	Time the statistics started.
Servicce Requests	Number of service requests.
Failed Requests	Number of failed requests.

Table 1064 on page 23477 lists the output fields for the `show rift node status` command. Output fields are listed in the approximate order in which they appear.

Table 1064: show rift node status Output Fields

Field Name	Field Description
System Name	Name of the system.

Table 1064: show rift node status Output Fields *(Continued)*

Field Name	Field Description
System ID	Identification of the system.
Level	Level of the system.
RIFT Encoding Major	Major version of the RIFT protocol.
Minor	Minor version of the RIFT protocol.
Flags	Flags of the RIFT protocol.
Capabilities	Configured capabilities for the RIFT protocol: • flood-reduction—True if enabled. False if disabled. • Hierarchy Indications—Indicates the hierarchy level.
LIE v4 Receive	IPv4 receive multicast address of the link information element (LIE).
LIE v6 Receive	IPv4 receive multicast address of the LIE.
Re-Connections	Number of tries to connect to the Redis server.
Peers	Number of peers configured.
3-way	Number of peers in three way state.
South	Number of peers in 3-way southbound.
North	Number of peers in 3-way northbound.

Sample Output

show rift node statistics

```
user@host> show rift node statistics
Starttime: 2019/06/06 16:01:00.786
Service Requests: 46, Failed Requests: 0
```

Sample Output

show rift node status

```
user@host> show rift node status
System Name: rift00, System ID: 00002c6bf5586fc0
Level: 24, RIFT Encoding Major: 29, Minor: 0
Flags: overload=False
Capabilities: flood-reduction: True, Hierarchy Indications: top_of_fabric
LIE v4 Receive: 224.0.0.120, LIE v6 Receive: ff02::a1f7
 Re-Connections: 2278
Peers: 4, 3-way: 4, South: 4, North: 0
```

Release Information

Command introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | **15794**

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

show rift path-computation

IN THIS SECTION

- Syntax | 23480
- Description | 23480
- Options | 23480
- Required Privilege Level | 23480
- Output Fields | 23480
- Sample Output | 23481
- Release Information | 23482

Syntax

```
show rift path-computation
statistics
```

Description

Display Routing in Fat Trees (RIFT) protocol path-computation information.

Options

statistics Display RIFT path computation statistics.

Required Privilege Level

view

Output Fields

Table 1065 on page 23481 lists the output fields for the `show rift path-computation statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1065: show rift path-computation statistics Output Fields

Field Name	Field Description
Dir	Direction of path computations.
Runs	Number of path computations.
Nodes#	Number of nodes after current path computation.
Total	• Nodes#—Total number of nodes in all path computations. • Prfxs#—Total prefixes in all path computations. • Deltas—Total generated prefix deltas in all path computations. • Holds—Number of times path computation was held down.
Last Run	Last path computation.
Last Trigger	Last triggering TIE.
On	Last triggering TIE received on.

Sample Output

show rift path-computation statistics

```
user@host> show rift path-computation statistics
          +----------- Total --------+
Dir   Runs Nodes#|Nodes# Prfxs# Deltas Holds|     Last Run Last
Trigger                                          On
-----+----+------+------+------+------+-----+-------------
+---------------------------------------------+-------------
South  144      9   1227      0      0     0  16:08:56.187            00002c6bf5586fc0/S/
Positive/30000008 16:08:56.187
```

```
North  144      1    142   2860    415   195  16:08:56.188           00002c6bf5c948c0/N/
External/60000157 16:08:01.158
```

Release Information

Command introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

show rift routes

IN THIS SECTION

- Syntax | 23482
- Description | 23483
- Options | 23483
- Required Privilege Level | 23483
- Output Fields | 23483
- Sample Output | 23485
- show rift routes statistics | 23485
- Release Information | 23486

Syntax

```
show rift routes
content
next-hops
statistics
```

Description

Display the routing table information of the Routing in Fat Trees (RIFT) protocol.

Options

content Display all route information of the RIFT protocol.

next-hops Display next-hop information of the RIFT protocol.

statistics Display the routing table statistics of the RIFT protocol.

Required Privilege Level

view

Output Fields

Table 1066 on page 23483 lists the output fields for the `show rift routes content` command. Output fields are listed in the approximate order in which they appear.

Table 1066: show rift routes content Output Fields

Field Name	Field Description
Prefix	IPv4 or IPv6 prefix address.
Active	Type of active RIFT route.
Metric	Metric of next hop.
N-Hop	Next-hop ID.
All Present	Types of all present RIFT routes.

Table 1067 on page 23484 lists the output fields for the `show rift routes next-hops` command. Output fields are listed in the approximate order in which they appear.

Table 1067: show rift routes next-hops Output Fields

Field Name	Field Description
Nexthop	ID of RIFT next hop.
SystemID	System ID of the adjacent node.
Links	Link IDs leading to the adjacent node.

Table 1068 on page 23484 lists the output fields for the `show rift routes statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1068: show rift routes statistics Output Fields

Field Name	Field Description
Nhops	Current number of next hops: • Deletes—Number of next hops deleted. • Adds/Changes—Number of next hops added or changed.
Last Transaction	Last transaction with: • Nhop Diff#—Number of next hop differentials downloaded. • Route Diff#—Number of route differentials downloaded. • Route Type—Number of route types involved.
AF	Numbers per address family for: • Prefixes—Number of prefixes in the table. • Deletes—Number of deletes performed over the lifetime. • Adds/Changes—Number of additions or changes performed over the lifetime.

Sample Output

show rift routes content

```
user@host> show rift routes content
Prefix                Active Metric N-Hop    All Present
---------------------+------+------+--------+-------------------
1.24.66.0/30          N           2 80004f04 N
2.2.2.11/32           NExt        2 80004f19 NExt
2.2.2.12/32           NExt        2 80004f13 NExt
2.2.2.13/32           NExt        2 80004f03 NExt
```

show rift routes next-hops

```
user@host> show rift routes next-hops
Nexthop   SystemID          Links
---------+-----------------+-----------------
80004f00  00002c6bf5a021c0  274 (ge-0/0/0.19)
80004f01  00002c6bf58703c0  270 (ge-0/0/0.15)
80004f02  00002c6bf5692fc0  278 (ge-0/0/0.23)
          00002c6bf55952c0  277 (ge-0/0/0.22)
```

show rift routes statistics

command-name

```
user@host> show rift routes statistics
Nhops Deletes  Adds/Changes
-----+--------+------------
   11        0          695

Last Transaction          Nhop Diff# Route Diff# Route Type
-------------------------+----------+-----------+-----------------------
2019/05/31 10:44:43.594            0           4                    NExt

AF   Prefixes Deletes Adds/Changes
----+--------+-------+------------
```

```
IPv4        12      170          256
IPv6         9      423          496
```

Release Information

Command introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

show rift tie

IN THIS SECTION

- Syntax | 23486
- Description | 23486
- Options | 23487
- Required Privilege Level | 23487
- Output Fields | 23487
- Sample Output | 23488
- Release Information | 23488

Syntax

```
show rift tie tie
```

Description

Display Routing in Fat Tress (RIFT) topology information element (TIE) information.

Options

tie Display information of the RIFT TIE in the following format:

```
(node-hex | node-name)/(North | South)/(node | prefix | positive | negative | key-value |
external)/TIE-number-hex
```

Where:

- (node-hex | node-name)—System ID in hexadecimal number of node name.

  To specify the value in hexadecimal form, include 0x as a prefix.

- (North | South)—Choice of direction. You can use one letter abbreviation.

- (node | prefix | positive | negative | key-value | external)—Choice of type. You can use two letter abbreviations.

- TIE-number-hex—TIE number in hexadecimal.

  To specify the value in hexadecimal form, include 0x as a prefix.

For example, *00002c6bf50f51c0/N/node/10000000*.

Required Privilege Level

view

Output Fields

Table 1069 on page 23487 lists the output fields for the `show rift tie` command. Output fields are listed in the approximate order in which they appear.

Table 1069: show rift tie Output Fields

Field Name	Field Description
TIE ID	Identification of TIE.
Content	Content of TIE as python parsable structure conforming to RIFT encoding schemas.

Sample Output

show rift tie

```
user@host> show rift tie 00002c6bf50f51c0/N/node/10000000
TIE ID: 00002c6bf50f51c0/N/Node____/10000000
Content: TIEElement(node=NodeTIEElement(neighbors={48842194470336L:
NodeNeighborsTIEElement(bandwidth=1000, cost=1,
...
```

Release Information

Command introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

show rift topology

IN THIS SECTION

- Syntax | 23489
- Description | 23489
- Options | 23489
- Required Privilege Level | 23489
- Output Fields | 23489
- Sample Output | 23491
- Release Information | 23491

Syntax

```
show rift topology
nodes
```

Description

Display the Routing in Fat Trees (RIFT) topology information.

Options

nodes Display information for all visible RIFT nodes.

Required Privilege Level

view

Output Fields

Table 1070 on page 23489 lists the output fields for the `show rift topology nodes` command. Output fields are listed in the approximate order in which they appear.

Table 1070: show rift topology nodes Output Fields

Field Name	Field Description
Lvl	RIFT node level.
Name	RIFT node name.
Originator	Node system ID.
Ovrld	Node overloaded: • Yes • No

Table 1070: show rift topology nodes Output Fields *(Continued)*

Field Name	Field Description
Dir	Direction of reachability. • N—North direction. • S—South direction.
3-way	Number of adjacencies in three way.
Miscbl	Number of miscabled links.
Secure	Number of secured links.
Auth	Number of authenticated links.
Non	Number of links without authentication.
V4	Number of originated IPv4 prefixes.
V6	Number of originated IPv6 prefixes.
Latest TIE Origination	Date of the newest TIE a node originated, that is, the last change it underwent.
Links	Number of secured (that is, authenticated through outer key) and unsecured links.
TIEs	Number of TIEs a node originated, both non-authenticated and secured by inner key.
Prefixs	Number of prefixes a node originated, both IPv4 and IPv6.

Sample Output

show rift topology nodes

```
user@host> show rift topology nodes
 +-------- Links-----+--- TIEs ----+-
Prefixs -+
Lvl Name Originator Ovrld Dir|3-way|Miscbl|Secure| Auth | Non | V4
 | V6 |Latest TIE Origination
---+----------+----------------+-----+---+-----+------+------+------+------+-----+-----
+------------------------
24 rift 01 0000000000000001 N 51 445 2
 1001 2019/11/09 15:07:01
```

Release Information

Command introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

show rift versions

IN THIS SECTION

- Syntax | 23492
- Description | 23492
- Options | 23492
- Required Privilege Level | 23492
- Output Fields | 23492
- Sample Output | 23493

Release Information | 23493

Syntax

```
show rift versions
info
```

Description

Display various package version of the Routing in Fat Trees (RIFT) protocol.

Options

info Display RIFT package information.

Required Privilege Level

view

Output Fields

Table 1071 on page 23492 lists the output fields for the `show rift versions info` command. Output fields are listed in the approximate order in which they appear.

Table 1071: show rift versions info Output Fields

Field Name	Field Description
Package	RIFT package version.
Build Date	Date and time of the RIFT package.
Encoding Version	RIFT protocol encoding version.

Table 1071: show rift versions info Output Fields *(Continued)*

Field Name	Field Description
Statistics Version	Statistics schema version.
Services Version	Services schema version.

Sample Output

show rift versions info

```
user@host> show rift versions info
Package: 1.0.0.1064751
Built On: 2019-11-29T20:04:18.141027255+00:00
Encoding Version: 2.0
Statistics Version: 3.0
Services Version: 17.1
```

Release Information

Command introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

rift | 15794

Understanding Junos Implementation of Routing in Fat Tree (RIFT) Protocol

show rip general-statistics

IN THIS SECTION

Syntax | 23494

- Syntax (EX Series Switches and QFX Series) | 23494
- Description | 23494
- Options | 23494
- Required Privilege Level | 23494
- Output Fields | 23495
- Sample Output | 23495
- Release Information | 23495

Syntax

```
show rip general-statistics
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show rip general-statistics
```

Description

Display brief RIP statistics.

Options

none	Display brief RIP statistics.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1072 on page 23495 lists the output fields for the `show rip general-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1072: show rip general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no recv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIP.
max memory	Most memory used by RIP.

Sample Output

show rip general-statistics

```
user@host> show rip general-statistics
RIPv2 I/O info:
     bad msgs       :         0
     no recv intf   :         0
     curr memory    :         0
     max memory     :         0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear rip general-statistics | 3838

show rip neighbor

IN THIS SECTION

- Syntax | 23496
- Syntax (EX Series Switches and QFX Series) | 23496
- Description | 23496
- Options | 23497
- Required Privilege Level | 23497
- Output Fields | 23497
- Sample Output | 23498
- Release Information | 23499

Syntax

```
show rip neighbor
<instance (all | instance-name)>
<logical-system (all | logical-system-name)>
<name>
```

Syntax (EX Series Switches and QFX Series)

```
show rip neighbor
<instance (all | instance-name)>
<name>
```

Description

Display information about RIP neighbors.

Options

none	Display information about all RIP neighbors for all instances.
instance (all | *instance-name*)	(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
name	(Optional) Display detailed information about only the specified RIP neighbor.

Required Privilege Level

view

Output Fields

Table 1073 on page 23497 lists the output fields for the `show rip neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 1073: show rip neighbor Output Fields

Field Name	Field Description
Neighbor	Name of the RIP neighbor. **NOTE**: Beginning with Junos OS Release 11.1, when you configure demand circuits, the output displays a demand circuit (DC) flag next to neighbor interfaces configured for demand circuits. If you configure demand circuits at the `[edit protocols rip group group-name neighbor neighbor-name]` hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the `[edit protocols rip group group-name]` hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits.
State	State of the connection: **Up** or **Dn** (Down).

Table 1073: show rip neighbor Output Fields *(Continued)*

Field Name	Field Description
Source Address	Address of the port on the local router.
Destination Address	Address of the port on the remote router.
Send Mode	Send options: **broadcast**, **multicast**, **none**, or **version 1**.
Receive Mode	Type of packets to accept: **both**, **none**, **version 1**, or **version 2**.
In Met	Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.

Sample Output

show rip neighbor

```
user@host> show rip neighbor
                  Local  Source          Destination     Send    Receive  In
Neighbor          State  Address         Address         Mode    Mode     Met
--------          -----  -------         -----------     ----    -------  ---
ge-2/3/0.0           Up  192.168.9.105   192.168.9.107   bcast   both       1
at-5/1/1.42          Dn  (null)          (null)          mcast   v2 only    3
at-5/1/0.42          Dn  (null)          (null)          mcast   both       3
at-5/1/0.0           Up  198.51.100.0    224.0.0.9       mcast   both       3
so-0/0/0.0           Up  192.168.9.97    224.0.0.9       mcast   both       3
```

show rip neighbor (With Demand Circuits Configured)

```
user@host> show rip neighbor
                  Local  Source          Destination     Send   Receive  In
Neighbor          State  Address         Address         Mode   Mode     Met
--------          -----  -------         -----------     ----   -------  ---
```

```
so-0/1/0.0(DC)       Up  10.10.10.2      224.0.0.9       mcast  both       1
so-0/2/0.0(DC)       Up  192.0.2.2       224.0.0.9       mcast  both       1
```

Release Information

Command introduced before Junos OS Release 7.4.

show rip statistics

IN THIS SECTION

- Syntax | 23499
- Syntax (EX Series Switches and QFX Series) | 23500
- Description | 23500
- Options | 23500
- Required Privilege Level | 23500
- Output Fields | 23500
- Sample Output | 23503
- Release Information | 23503

Syntax

```
show rip statistics
<instance (all | instance-name)>
<logical-system (all | logical-system-name)>
<name>
<peer (all | address)>
```

Syntax (EX Series Switches and QFX Series)

```
show rip statistics
<instance (all | instance-name)>
<name>
```

Description

Display RIP statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.

Options

none	Display RIP statistics for all routing instances.
instance (all | *instance-name*)	(Optional) Display RIP statistics for all instances or for only the specified routing instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
name	(Optional) Display detailed information about only the specified RIP neighbor.
peer (all | *address*)	(Optional) Display RIP statistics for a single peer or all peers.

Required Privilege Level

view

Output Fields

Table 1074 on page 23501 lists the output fields for the `show rip statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1074: show rip statistics Output Fields

Field Name	Field Description
RIP info	Information about RIP on the specified interface: • **port**—UDP port number used for RIP. • **update interval**—Interval between routing table updates, in seconds. • **holddown**—Hold-down interval, in seconds. • **timeout**—Timeout interval, in seconds. • **restart in progress**—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart. • **restart time**—Estimated time for the graceful restart to finish, in seconds. • **restart will complete in**—Remaining time for the graceful restart to finish, in seconds. • **rts learned**—Number of routes learned through RIP. • **rts held down**—Number of routes held down by RIP. • **rqsts dropped**—Number of received request packets that were dropped. • **resps dropped**—Number of received response packets that were dropped.
logical-interface	Name of the logical interface and its statistics: • **routes learned**—Number of routes learned on the logical interface. • **routes advertised**—Number of routes advertised by the logical interface.

Table 1074: show rip statistics Output Fields *(Continued)*

Field Name	Field Description
Counter	List of counter types: • **Updates Sent**—Number of update messages sent. • **Triggered Updates Sent**—Number of triggered update messages sent. • **Responses Sent**—Number of response messages sent. • **Bad Messages**—Number of invalid messages received. • **RIPv1 Updates Received**—Number of RIPv1 update messages received. • **RIPv1 Bad Route Entries**—Number of RIPv1 invalid route entry messages received. • **RIPv1 Updates Ignored**—Number of RIPv1 update messages ignored. • **RIPv2 Updates Received**—Number of RIPv2 update messages received. • **RIPv2 Bad Route Entries**—Number of RIPv2 invalid route entry messages received. • **RIPv2 Updates Ignored**—Number of RIPv2 update messages that were ignored. • **Authentication Failures**—Number of received update messages that failed authentication. • **RIP Requests Received**—Number of RIP request messages received. • **RIP Requests Ignored**—Number of RIP request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

Sample Output

show rip statistics

```
user@host> show rip statistics so-0/0/0.0
RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s
restart in progress: restart time 60s; restart will complete in 55s
    rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0
so-0/0/0.0:  0 routes learned; 501 routes advertised
Counter                         Total   Last 5 min  Last minute
-------                   -----------  -----------  -----------
Updates Sent                        0            0            0
Triggered Updates Sent              0            0            0
Responses Sent                      0            0            0
Bad Messages                        0            0            0
RIPv1 Updates Received              0            0            0
RIPv1 Bad Route Entries             0            0            0
RIPv1 Updates Ignored               0            0            0
RIPv2 Updates Received              0            0            0
RIPv2 Bad Route Entries             0            0            0
RIPv2 Updates Ignored               0            0            0
Authentication Failures             0            0            0
RIP Requests Received               0            0            0
RIP Requests Ignored                0            0            0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear rip statistics | 3840

show ripng general-statistics

IN THIS SECTION

- Syntax | 23504
- Syntax (EX Series Switch) | 23504
- Description | 23504
- Options | 23504
- Required Privilege Level | 23505
- Output Fields | 23505
- Sample Output | 23505
- Release Information | 23505

Syntax

```
show ripng general-statistics
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch)

```
show ripng general-statistics
```

Description

Display general RIP next-generation (RIPng) statistics.

Options

none	Display general RIPng statistics.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1075 on page 23505 lists the output fields for the `show ripng general-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1075: show ripng general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no recv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIPng.
max memory	Most memory used by RIPng.

Sample Output

show ripng general-statistics

```
user@host> show ripng general-statistics
RIPng I/O info:
     bad msgs       :         0
     no recv intf   :         0
     curr memory    :         0
     max memory     :         0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear ripng general-statistics | 3842

show ripng neighbor

IN THIS SECTION

- Syntax | 23506
- Syntax (EX Series Switch) | 23506
- Description | 23506
- Options | 23507
- Required Privilege Level | 23507
- Output Fields | 23507
- Sample Output | 23508
- Release Information | 23508

Syntax

```
show ripng neighbor
<logical-system (all | logical-system-name)>
<name>
```

Syntax (EX Series Switch)

```
show ripng neighbor
<name>
```

Description

Display information about RIP next-generation (RIPng) neighbors.

Options

none	Display information about all RIPng neighbors.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
name	(Optional) Display detailed information about a specific RIPng neighbor.

Required Privilege Level

view

Output Fields

Table 1076 on page 23507 lists the output fields for the `show ripng neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 1076: show ripng neighbor Output Fields

Field Name	Field Description
Neighbor	Name of RIPng neighbor.
State	State of the connection: **Up** or **Dn** (Down).
Source Address	Source address.
Destination Address	Destination address.
Send	Send options: **broadcast**, **multicast**, **none**, **version 1**, or **yes.**
Recv	Type of packets to accept: **both**, **none**, **version 1**, or **yes.**
In Met	Metric added to incoming routes when advertising into RIPng routes that were learned from other protocols.

Sample Output

show ripng neighbor

```
user@host> show ripng neighbor
                    Source                          Dest               In
Neighbor     State  Address                         Address  Send Recv Met
--------     -----  -------                         -------  ---- ---- ---
fe-0/0/2.0      Up  fe80::290:69ff:fe68:b002        ff02::9  yes  yes    1
```

Release Information

Command introduced before Junos OS Release 7.4.

show ripng statistics

IN THIS SECTION

- Syntax | 23508
- Syntax (EX Series Switch) | 23509
- Description | 23509
- Options | 23509
- Required Privilege Level | 23509
- Output Fields | 23509
- Sample Output | 23511
- Release Information | 23512

Syntax

```
show ripng statistics
<logical-system (all | logical-system-name)>
<name>
```

Syntax (EX Series Switch)

```
show ripng statistics
<name>
```

Description

Display RIP next generation (RIPng) statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.

Options

none	Display RIPng statistics for all neighbors.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
name	(Optional) Display detailed information about a specific RIPng neighbor.

Required Privilege Level

view

Output Fields

Table 1077 on page 23510 lists the output fields for the `show ripng statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1077: show ripng statistics Output Fields

Field Name	Field Description
RIPng info	Information about RIPng on the specified interface: • **port**—UDP port number used for RIPng. • **holddown**—Hold-down interval, in seconds. • **rts learned**—Number of routes learned through RIPng. • **rts held down**—Number of routes held down by RIPng. • **rqsts dropped**—Number of received request packets that were dropped. • **resps dropped**—Number of received response packets that were dropped. • **restart**—Graceful restart status. Displayed when RIPng is or has been in the process of graceful restart.
logical-interface	Name of the logical interface and its statistics: • **routes learned**—Number of routes learned on the logical interface. • **routes advertised**—Number of routes advertised by the logical interface. • **timeout**—Timeout interval, in seconds. • **update interval**—Interval between routing table updates, in seconds.

Table 1077: show ripng statistics Output Fields *(Continued)*

Field Name	Field Description
Counter	List of counter types: • **Updates Sent**—Number of update messages sent. • **Triggered Updates Sent**—Number of triggered update messages sent. • **Responses Sent**—Number of response messages sent. • **Bad Messages**—Number of invalid messages received. • **Updates Received**—Number of RIPng update messages received. • **Bad Route Entries**—Number of RIPng invalid route entry messages received. • **Updates Ignored**—Number of RIPng update messages ignored. • **RIPng Requests Received**—Number of RIPng request messages received. • **RIPng Requests Ignored**—Number of RIPng request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

Sample Output

show ripng statistics

```
user@host> show ripng statistics
RIPng info: port 521; holddown 120s;
    rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0

so-0/1/3.0:  0 routes learned; 1 routes advertised; timeout 180s; update interval 20s
Counter                          Total   Last 5 min  Last minute
```

```
-------                    -----------  -----------  -----------
Updates Sent                       934           16            4
Triggered Updates Sent               1            0            0
Responses Sent                       0            0            0
Bad Messages                         0            0            0
Updates Received                     0            0            0
Bad Route Entries                    0            0            0
Updates Ignored                      0            0            0
RIPng Requests Received              0            0            0
RIPng Requests Ignored               0            0            0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear ripng statistics | 3844

show route

IN THIS SECTION

- Syntax | 23513
- Syntax (EX Series Switches) | 23513
- Description | 23513
- Options | 23513
- Required Privilege Level | 23514
- Output Fields | 23514
- Sample Output | 23520
- Release Information | 23526

Syntax

```
show route
<all>
<destination-prefix>
<logical-system (all | logical-system-name)>
<private>
<te-ipv4-prefix-ip  te-ipv4-prefix-ip>
<te-ipv4-prefix-node-ip te-ipv4-prefix-node-ip>
<te-ipv4-prefix-node-iso te-ipv4-prefix-node-iso>
<te-ipv6-prefix-ipv6-addr te-ipv6-prefix-ipv6-addr>
<te-ipv6-prefix-node-iso te-ipv6-prefix-node-iso>
<rib-sharding (main | rib-shard-name)>
<validation-statistics>
```

Syntax (EX Series Switches)

```
show route
<all>
<destination-prefix>
<private>
```

Description

Display the active entries in the routing tables.

Options

none	Display brief information about all active entries in the routing tables.
all	(Optional) Display information about all routing tables, including private, or internal, routing tables.
destination-prefix	(Optional) Display active entries for the specified address or range of addresses.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

private	(Optional) Display information only about all private, or internal, routing tables.
programmed detail	(Optional) Display API-programmed routes.
display-client-data	(Optional) Display client id and cookie information for routes installed by the routing protocol process client applications.
te-ipv4-prefix-ip *te-ipv4-prefix-ip*	(Optional) Display IPv4 address of the traffic-engineering prefix, without the mask length if present in the routing table.
te-ipv4-prefix-node-ip *te-ipv4-prefix-node-ip*	(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 node addresses from the traffic-engineered routes in the `lsdist.0` table.
te-ipv4-prefix-node-iso *te-ipv4-prefix-node-iso*	(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 routes with the specified ISO circuit ID from the `lsdist.0` table.
te-ipv6-prefix-ipv6-addr *te-ipv6-prefix-ipv6-addr*	(Optional) Filter IPv6 node addresses from the traffic-engineering IPv6 prefix.
te-ipv6-prefix-node-iso *te-ipv6-prefix-node-iso*	(Optional) Filter IPv6 routes with the specified ISO circuit ID in the traffic-engineering IPv6 prefix.
rib-sharding (main | *rib-shard-name*)	(Optional) Display the rib shard name.
validation-statistics	(Optional) Display per-RIB specific statistics information.

Required Privilege Level

view

Output Fields

Table 1078 on page 23515 describes the output fields for the `show route` command. Output fields are listed in the approximate order in which they appear.

Table 1078: show route Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number destinations	Number of destinations for which there are routes in the routing table.
number routes	Number of routes in the routing table and total number of routes in the following states: • `active` (routes that are active). • `holddown` (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. However, if you have configured advertisement of multiple routes (with the `add-path` or `advertise-inactive` statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong. If you have configured `uRPF-loose` mode, the holddown bit is most likely set because Kernel Routing Table (KRT) is using inactive route to build valid incoming interfaces. In this case, you can ignore the holddown state because nothing is wrong. • `hidden` (routes that are not used because of a routing policy).

Table 1078: show route Output Fields *(Continued)*

Field Name	Field Description

destination-prefix	Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:
	• *MPLS-label* (for example, 80001).
	• *interface-name* (for example, ge-1/0/2).
	• *neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source* (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96):
	  • *neighbor-address*—Address of the neighbor.
	  • *control-word-status*—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord.
	  • *encapsulation type*—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  • *vc-id*—Virtual circuit identifier.
	  • *source*—Source of the advertisement: Local or Remote.

[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a `to` line indicates the best subpath to the route. In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.

Table 1078: show route Output Fields *(Continued)*

Field Name	Field Description
weeks:days hours:minutes:seconds	How long the route been known (for example, 2w4d 13:11:14, or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
localpref	Local preference value included in the route.
from	Interface from which the route was received.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • `I`—IGP. • `E`—EGP. • `?`—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • `[ ]`—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • `{ }`—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • `( )`—Parentheses enclose a confederation. • `( [ ] )`—Parentheses and brackets enclose a confederation set. **NOTE**: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.

Table 1078: show route Output Fields *(Continued)*

Field Name	Field Description
encapsulated	Extended next-hop encoding capability enabled for the specified BGP community for routing IPv4 traffic over IPv6 tunnels. When BGP receives routes without the tunnel community, IPv4-0ver IPv6 tunnels are not created and BGP routes are resolved without encapsulation.
Route Labels	Stack of labels carried in the BGP route update.
validation-state	(BGP-learned routes) Validation status of the route: • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Unverified—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the destination is `Discard`, traffic is dropped.

Table 1078: show route Output Fields *(Continued)*

Field Name	Field Description
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • *lsp-path-name*—Name of the LSP used to reach the next hop. • *label-action*—MPLS label and operation occurring at the next hop. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label). For VPNs, expect to see multiple push operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).
Private unicast	(Enhanced subscriber management for MX Series routers) Indicates that an access-internal route is managed by enhanced subscriber management. By contrast, access-internal routes not managed by enhanced subscriber management are displayed with associated next-hop and media access control (MAC) address information.
balance	Distribution of the load based on the underlying operational interface bandwidth for equal-cost multipaths (ECMP) across the nexthop gateways in percentages.
Valid Routes	Count of valid routes in the local routing information base (RIB).
Routes invalid due to max-length	Count of invalid routes due to `origin-as` mismatch in the local RIB.
Routes invalid due to origin-as	Count of invalid routes due to maximum-length mismatch in the local RIB.

Table 1078: show route Output Fields *(Continued)*

Field Name	Field Description
Unknown routes	Count of unknown routes in the local RIB.

Sample Output

show route

```
user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
                   *[MVPN/70] 19:53:41, metric2 1
                      Indirect
1:65500:1:10.0.0.40/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
                    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
                    [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
```

show route

The following sample output shows route hierarchy for translation route.

```
user@host> show route 10.1.1.1
```

```
C1.inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.1/32        *[PRPD/10] 00:16:50, metric 2
                    >  to 192.0.2.2 via ge-0/0/1.0
```

show route forwarding-table matching 10.1.1.1

```
user@host> show route forwarding-table matching 10.1.1.1
Routing table: C1.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif

10.1.1.1/32        user     0                    indr  1048574     4

                                                comp      624     2
```

show route 10.1.1.1 extensive expanded-nh

```
user@host> show route 10.1.1.1extensive expanded-nh
C1.inet
C1.inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
10.1.1.1/32 (1 entry, 1 announced)
Installed-nexthop:
Indr (0xc5c207c) ::44.0.0.1
  Krt_inh (0xc6fd004) Index:1048574 PNH: ::44.0.0.1
    Translate-comp (0xc5c2144) Index:624 v4tov6 src ::22.0.0.1 dest ::44.0.0.1
```

show route te-ipv6-prefix-ipv6-addr

```
user@host> show route te-ipv6-prefix-ipv6-addr 10::10

inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
mpls.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)

inet6.0: 10 destinations, 11 routes (10 active, 0 holddown, 0 hidden)

inet6.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

lsdist.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:10::10/128 } ISIS-L1:0 }/1216
                   *[IS-IS/15] 00:07:58
                       Fictitious
```

show route te-ipv6-prefix-node-iso

```
user@host> show route te-ipv6-prefix-node-iso 0100.0100.0100.00

inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)

inet6.0: 10 destinations, 11 routes (10 active, 0 holddown, 0 hidden)

inet6.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

lsdist.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:10::10/128 } ISIS-L1:0 }/
1216
                   *[IS-IS/15] 00:08:46
                       Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:21:0:1::1/128 } ISIS-L1:0 }/
1216
                   *[IS-IS/15] 00:08:46
                       Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:abcd::128:207:200:16/128 } ISIS-
```

```
L1:0 }/1216
                   *[IS-IS/15] 00:08:46
                       Fictitious
```

show route (VPN)

The following sample output shows a VPN route with composite next hops enabled. The first `Push` operation corresponds to the outer label. The second `Push` operation corresponds to the inner label.

```
user@host> show route 192.0.2.0

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.0/24        [BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                    [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                   #[Multipath/255] 00:28:28, metric2 102
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                      to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
```

show route (with Destination Prefix)

```
user@host> show route 192.168.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/12      *[Static/5] 2w4d 12:54:27
                    > to 192.168.167.254 via fxp0.0
```

show route destination-prefix detail

```
user@host> show route 198.51.100.0 detail
```

```
inet.0: 15 destinations, 20 routes (15 active, 0 holddown, 0 hidden)
198.51.100.0/24 (2 entries, 2 announced)
        *BGP    Preference: 170/-101
         ...
         BGP-Static Preference: 4294967292
         Next hop type: Discard
         Address: 0x9041ae4
         Next-hop reference count: 2
         State: <NoReadvrt Int Ext AlwaysFlash>
         Inactive reason: Route Preference
         Local AS:   200
         Age: 4d 1:40:40
         Validation State: unverified
         Task: RT
         Announcement bits (1): 2-BGP_RT_Background
         AS path: 4 5 6 I
```

show route extensive

```
user@host> show route extensive
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
    *BGP   Preference: 170/-101
           PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group 203.0.113.1
           Next hop type: Indirect
           Address: 0x92455b8
           Next-hop reference count: 2
           Source: 10.0.0.30
           Protocol next hop: 10.0.0.40
           Indirect next hop: 2 no-forward
           State: <Active Int Ext>
           Local AS: 64510 Peer AS: 64511
           Age: 3  Metric2: 1
           Validation State: unverified
           Task: BGP_64510.10.0.0.30+179
           Announcement bits (2): 0-PIM.v1 1-mvpn global task
           AS path: I (Originator) Cluster list:  10.0.0.30
           AS path:  Originator ID: 10.0.0.40
           Communities: target:64502:100 encapsulation:0L:14
           Import Accepted
           Localpref: 100
```

```
            Router ID: 10.0.0.30
            Primary Routing Table bgp.mvpn.0
            Indirect next hops: 1
                 Protocol next hop: 10.0.0.40 Metric: 1
                 Indirect next hop: 2 no-forward
                 Indirect path forwarding next hops: 1
                       Next hop type: Router
                       Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
                 10.0.0.40/32 Originating RIB: inet.3
                       Metric: 1                       Node path count: 1
                       Forwarding nexthops: 1
                          Nexthop: 10.0.24.4 via lt-0/3/0.24
```

show route programmed detail

```
user@host> show route programmed detail
inet.0: 36 destinations, 37 routes (36 active, 0 holddown, 0 hidden)
100.75.1.0/27 (2 entries, 1 announced)
        *Static Preference: 5/100
                Next hop type: Router, Next hop index: 0
                Address: 0xcc38a10
                Next-hop reference count: 1
                Next hop: 100.30.1.2 via ge-0/0/2.0 weight 0x1, selected
                Session Id: 0x0
                Next hop: via fti0.1001 weight 0x8001
                Session Id: 0x0
               State: <Active Int NSR-incapable Programmed>
                Age: 37
                Validation State: unverified
                Announcement bits (1): 0-KRT
                AS path: I
```

show route validation-statistics

```
user@host> show route validation-statistics
        Valid Routes: 432
        Routes invalid due to max-length: 646
        Routes invalid due to origin-as: 21
        Unknown routes: 5353
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `private` introduced in Junos OS Release 9.5.

Option `private` introduced in Junos OS Release 9.5 for EX Series switches.

Option `display-client-data` introduced in Junos OS Release 16.2R1 on MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series routers.

Options `te-ipv4-prefix-ip`, `te-ipv4-prefix-node-ip`, and `te-ipv4-prefix-node-iso` introduced in Junos OS Release 17.2R1 on MX Series and PTX Series.

`rib-sharding` option introduced in cRPD Release 20.1R1.

RELATED DOCUMENTATION

- Understanding IS-IS Configuration
- Verifying and Managing Junos OS Enhanced Subscriber Management

show route active-path

IN THIS SECTION

- Syntax | 23527
- Syntax (EX Series Switches) | 23527
- Description | 23527
- Options | 23527
- Required Privilege Level | 23527
- Output Fields | 23527
- Sample Output | 23528
- Release Information | 23533

Syntax

```
show route active-path
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route active-path
<brief | detail | extensive | terse>
```

Description

Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.

Options

none	Display all active routes.
brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to `brief`.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the `show route` command, the `show route detail` command, the `show route extensive` command, or the `show route terse` command.

Sample Output

show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
                    > via lo0.0
10.255.71.50/32    *[IS-IS/15] 00:18:13, metric 10
                    > to 172.16.100.1 via so-2/1/3.0
172.16.100.1/24      *[Direct/0] 00:18:36
                    > via so-2/1/3.0
172.16.100.1/32      *[Local/0] 00:18:41
                      Local via so-2/1/3.0
192.168.64.0/21    *[Direct/0] 21:33:52
                    > via fxp0.0
192.168.70.19/32   *[Local/0] 21:33:52
                      Local via fxp0.0
```

show route active-path brief

The output for the `show route active-path brief` command is identical to that for the `show route active-path` command. For sample output, see `show route active-path`.

show route active-path detail

```
user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: ‹Active Int›
                Local AS:   200
```

```
                Age: 21:37:10
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 172.16.100.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:31      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

172.16.100.0/24 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:54
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

172.16.100.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 21:59
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
```

```
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I
```

show route active-path extensive

```
user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
```

```
                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {172.16.100.1}
IS-IS level 2, LSP fragment 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 172.16.100.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:08      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

172.16.100.1/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:31
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

172.16.100.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 24:36
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
```

```
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I
```

show route active-path terse

```
user@host> show route active-path terse

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.19/32    D   0                        >lo0.0
* 10.255.71.50/32    I  15         10             >172.16.100.1.
* 172.16.100.0/24       D   0                        >so-2/1/3.0
* 172.16.100.2/32       L   0                         Local
* 192.168.64.0/21    D   0                        >fxp0.0
```

```
* 192.168.70.19/32   L   0                      Local
```

Release Information

Command introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

- show route | 23512
- show route detail | 23573
- show route extensive | 23619
- show route terse | 23824

show route advertising-protocol

IN THIS SECTION

- Syntax | 23534
- Description | 23534
- Options | 23534
- Additional Information | 23534
- Required Privilege Level | 23535
- Output Fields | 23535
- Sample Output | 23538
- Release Information | 23542

Syntax

```
show route advertising-protocol protocol neighbor-address
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Description

Display the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
neighbor-address	Address of the neighboring router to which the route entry is being transmitted.
protocol	Protocol transmitting the route: • `bgp`—Border Gateway Protocol • `dvmrp`—Distance Vector Multicast Routing Protocol • `msdp`—Multicast Source Discovery Protocol • `pim`—Protocol Independent Multicast • `rip`—Routing Information Protocol • `ripng`—Routing Information Protocol next generation

Additional Information

Routes displayed are routes that the routing table has exported into the routing protocol and that have been filtered by the associated protocol's `export` routing policy statements. Starting with Junos OS Release 13.3, you can display the routing instance table `foo` for any address family, on a VPN route reflector, or a VPN AS boundary router that is advertising local VPN routes. However, If you do not specify the `table` in the command, the output displays each VRF prefix twice.

Required Privilege Level

view

Output Fields

Table 1079 on page 23535 lists the output fields for the `show route advertising-protocol` command. Output fields are listed in the approximate order in which they appear.

Table 1079: show route advertising-protocol Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, inet.0.	All levels
number destinations	Number of destinations for which there are routes in the routing table.	All levels
number routes	Number of routes in the routing table and total number of routes in the following states: • `active` (routes that are active) • `holddown` (routes that are in the pending state before being declared inactive) • `hidden` (routes that are not used because of a routing policy)	All levels
Prefix	Destination prefix.	`brief` none
destination-prefix (entry , announced)	Destination prefix. The `entry` value is the number of routes for this destination, and the `announced` value is the number of routes being announced for this destination.	`detail extensive`
BGP group and type	BGP group name and type (Internal or External).	`detail extensive`
Route Distinguisher	Unique 64-bit prefix augmenting each IP subnet.	`detail extensive`

Table 1079: show route advertising-protocol Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Advertised Label	Incoming label advertised by the Label Distribution Protocol (LDP). When an IP packet enters a label-switched path (LSP), the ingress router examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label.	`detail extensive`
Label-Base, range	First label in a block of labels and label block size. A remote PE router uses this first label when sending traffic toward the advertising PE router.	`detail extensive`
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routers by advertising VPN labels. VPN labels transit over either a Resource Reservation Protocol (RSVP) or a Label Distribution Protocol (LDP) label-switched path (LSP) tunnel.	`detail extensive`
Nexthop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the next-hop advertisement to the peer is Self, and the RIB-out next hop is a specific IP address, the RIB-out IP address is included in the extensive output.	All levels
MED	Multiple exit discriminator value included in the route.	`brief`
Lclpref or Localpref	Local preference value included in the route.	All levels
Queued	When BGP route prioritization is enabled and a route is present in a priority queue, this shows which priority queue the route is in.	All levels except `brief`

Table 1079: show route advertising-protocol Output Fields *(Continued)*

Field Name	Field Description	Level of Output
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • `I`—IGP. • `E`—EGP. • `?`—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • `[ ]`—Brackets enclose the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured. • `{ }`—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • `( )`—Parentheses enclose a confederation. • `( [ ] )`—Parentheses and brackets enclose a confederation set. **NOTE**: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
Route Labels	Stack of labels carried in the BGP route update.	`detail extensive`
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	`detail extensive`
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	`detail extensive`

Table 1079: show route advertising-protocol Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Communities`	Community path attribute for the route. See the output field table for the "show route detail" on page 23573 command for all possible values for this field.	`detail extensive`
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	`detail extensive`
Attrset AS	Number, local preference, and path of the autonomous system (AS) that originated the route. These values are stored in the Attrset attribute at the originating router.	`detail extensive`
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	`detail extensive`
control flags	Control flags: `none` or `Site Down`.	`detail extensive`
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	`detail extensive`

Sample Output

show route advertising-protocol bgp (Layer 3 VPN)

```
user@host> show route advertising-protocol bgp 10.255.14.171
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.172/32   Self                     1        100 I
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.181/32   Self                     2        100 I
```

show route advertising-protocol bgp (IPv6 unicast address community)

```
user@host> show route advertising-protocol bgp 10.255.165.220 extensive
       vpn1.inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
    * 128.205.172.129/32 (1 entry, 1 announced)
    BGP group internal type Internal
    Route Distinguisher: 10.255.168.42:8
    VPN Label: 299808
    Nexthop: Self
    Flags: Nexthop Change
    MED: 1
    Localpref: 100
       AS path: [203] I
       Communities: rte-type:0.0.0.0:1:0 ipv6-origin:<2001:db8:9999::9>:89
       ipv6-target:<2001:db8:9999::9>:89 ipv6-extended:203:<2001:db8:9999::9>:89
       ipv6-origin:<2001:db8:9999::9>:137 ipv6-target:<2001:db8:9999::9>:137
       ipv6-extended:515:<2001:db8:9999::9>:137
```

show route advertising-protocol bgp detail

```
user@host> show route advertising-protocol bgp 10.222.1.3 detail
bgp20.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
10.222.1.11/32 (1 entry, 1 announced)
 BGP group pe-pe type Internal
     Route Distinguisher: 10.255.14.11:69
     Advertised Label: 100000
     next hop: Self
     Localpref: 100
     AS path: 2 I
     Communities: target:69:20
     AIGP 210
10.8.0.0/16 (1 entry, 1 announced)
 BGP group pe-pe type Internal
     Route Distinguisher: 10.255.14.11:69
     Advertised Label: 100000
     Next hop: Self
     Localpref: 100
     AS path: 2 I
     Communities: target:69:20
     AIGP 210
```

show route advertising-protocol bgp detail (Aggregate Extended Community Bandwidth)

```
user@host> show route advertising-protocol bgp 10.0.4.2 10.0.2.0/30 detail
inet.0: 20 destinations, 26 routes (20 active, 0 holddown, 0 hidden)
* 10.0.2.0/30 (2 entries, 1 announced)
 BGP group external2 type External
     Nexthop: Self
     AS path: [65000] 65001 I
     Communities: bandwidth:65000:80000000
```

show route advertising-protocol bgp detail (BGP Multicast)

```
user@host>show route advertising-protocol bgp 10.4.6.1 detail
bgpmcast.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 4:3:0:0:32:10.2.20.20:32:225.1.1.1:10.1.1.1:10.20.20.20/240 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Nexthop: 10.4.6.6
     Localpref: 100
     AS path: [65000] 65001 I
     Communities: target:65000:100
     Tunnel type: AnyEncap, RPF tunnel:, Remote end point: 10.1.2.1
     Tunnel type: AnyEncap, Remote end point: 10.1.4.1
```

show route advertising-protocol bgp detail (Labeled Unicast)

```
user@host>show route advertising bgp 10.1.1.3 detail
inet.0: 69 destinations, 70 routes (69 active, 0 holddown, 0 hidden)
* 10.1.1.8/32 (2 entries, 2 announced)
BGP group ibgp type Internal
Route Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: 10.1.1.4
MED: 7
Localpref: 100
AS path: [65005] I
Cluster ID: 10.3.3.3
Originator ID: 10.1.1.1
Entropy label capable
inet6.0: 26 destinations, 28 routes (26 active, 0 holddown, 0 hidden)
* 2001:db8:100::1/128 (2 entries, 1 announced)
```

```
BGP group ibgp type Internal
Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: 2001:db8:0:ffff:10.1.1.4
Localpref: 100
AS path: [65005] I
Cluster ID: 10.3.3.3
Originator ID: 10.1.1.1
```

show route advertising-protocol bgp detail (Layer 2 VPN)

```
user@host> show route advertising-protocol bgp 192.168.24.1 detail
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
192.168.16.10:1:1:1/96 (1 entry, 1 announced)
 BGP group int type Internal
     Route Distinguisher: 192.168.16.1:1
     Label-base : 32768, range : 3
     Nexthop: Self
     Localpref: 100
     AS path: I
     Communities: target:65412:100
     AIGP 210
     Layer2-info: encaps:VLAN, control flags:, mtu:
```

show route advertising-protocol bgp detail (Layer 3 VPN)

```
user@host> show route advertising-protocol bgp 10.255.14.176 detail
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101264
     Nexthop: Self
     Localpref: 100
     AS path: I
     Communities: target:10:100
     AIGP 210
     AttrSet AS: 65100
         Localpref: 100
```

```
        AS path: I
...
```

show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address)

```
user@host> show route advertising-protocol bgp 10.200.0.2 10.170.1.0/24 extensive all
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 6 hidden)
  10.170.1.0/24 (2 entries, 1 announced)
 BGP group eBGP-INTEROP type External
     Nexthop: Self (rib-out 10.100.3.2)
     AS path: [64713] 65200 I
...
```

Release Information

Command introduced before Junos OS Release 7.4.

`ipv6-origin`, `ipv6-target`, and `ipv6-extended Community` output fields supported in Junos OS Release and Junos OS Evolved Release 23.1.

RELATED DOCUMENTATION

Example: Configuring the MED Attribute That Determines the Exit Point in an AS

show route all

IN THIS SECTION

- Syntax | 23543
- Syntax (EX Series Switches) | 23543
- Description | 23543
- Options | 23543
- Required Privilege Level | 23543
- Output Fields | 23543

Sample Output | 23544
Release Information | 23545

Syntax

```
show route all
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route all
```

Description

Display information about all routes in all routing tables, including private, or internal, tables.

Options

none	Display information about all routes in all routing tables, including private, or internal, tables.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

In Junos OS Release 9.5 and later, only the output fields for the `show route all` command display all routing tables, including private, or hidden, routing tables. The output field table of the "show route" on page 23512 command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later.

Sample Output

show route all

The following example displays a snippet of output from the `show route` command and then displays the same snippet of output from the `show route all` command:

```
user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
1                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
2                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
800017             *[VPLS/7] 1d 14:00:16
                    > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 14:00:26
                    > via vt-3/2/0.32772, Pop

user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
1                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
2                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
800017             *[VPLS/7] 1d 13:54:49
                    > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 13:54:59
                    > via vt-3/2/0.32772, Pop
vt-3/2/0.32769      [VPLS/7] 1d 13:54:49
                      Unusable
vt-3/2/0.32772      [VPLS/7] 1d 13:54:59
                      Unusable
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show route brief | 23553

show route detail | 23573

show route aspath-regex

IN THIS SECTION

- Syntax | 23545
- Syntax (EX Series Switches) | 23546
- Description | 23546
- Options | 23546
- Additional Information | 23546
- Required Privilege Level | 23547
- Output Fields | 23547
- Sample Output | 23547
- Release Information | 23548

Syntax

```
show route aspath-regex regular-expression
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route aspath-regex regular-expression
```

Description

Display the entries in the routing table that match the specified autonomous system (AS) path regular expression.

Options

regular-expression	Regular expression that matches an entire AS path.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

You can specify a regular expression as:

- An individual AS number
- A period wildcard used in place of an AS number
- An AS path regular expression that is enclosed in parentheses

You also can include the operators described in the table of AS path regular expression operators in the *Junos Policy Framework Configuration Guide*. The following list summarizes these operators:

- {*m*,*n*}—At least *m* and at most *n* repetitions of the AS path term.
- {*m*}—Exactly *m* repetitions of the AS path term.
- {*m*,}—*m* or more repetitions of the AS path term.
- *—Zero or more repetitions of an AS path term.
- +—One or more repetitions of an AS path term.
- ?—Zero or one repetition of an AS path term.
- *aspath_term* | *aspath_term*—Match one of the two AS path terms.

When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:

```
show route aspath-regex ".* 234 .*"
```

Required Privilege Level

view

Output Fields

For information about output fields, see the output field table for the "show route" on page 23512 command.

Sample Output

show route aspath-regex (Matching a Specific AS Number)

```
user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25     *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65548 65536}) IGP
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25   *[IS-IS/15] 09:15:37, metric 37, tag 1
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
                    [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65548 65536}) IGP
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
...
```

show route aspath-regex (Matching Any Path with Two AS Numbers)

```
user@host> show route aspath-regex ".* 234 3561 .*"

inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
```

```
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 2685 2686 Incomplete
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 5696 7369 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 6327 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Using AS Path Regular Expressions

show route best

IN THIS SECTION

- Syntax | 23549
- Syntax (EX Series Switches) | 23549
- Description | 23549
- Options | 23549
- Required Privilege Level | 23549
- Output Fields | 23549
- Sample Output | 23550
- Release Information | 23552

Syntax

```
show route best destination-prefix
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route best destination-prefix
<brief | detail | extensive | terse>
```

Description

Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to `brief`.
destination-prefix	Address or range of addresses.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the `show route` command, the `show route detail` command, the `show route extensive` command, or the `show route terse` command.

Sample Output

show route best

```
user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[OSPF/10] 1d 13:19:20, metric 2
                    > to 10.31.1.6 via ge-3/1/0.0
                      via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[RSVP/7] 1d 13:20:13, metric 2
                    > via so-0/3/0.0, label-switched-path green-r1-r3

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8         *[Direct/0] 2d 01:43:34
                    > via fxp2.0
                    [Direct/0] 2d 01:43:34
                    > via fxp1.0
```

show route best detail

```
user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                Next hop: via so-0/3/0.0
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:06        Metric: 2
                Area: 0.0.0.0
                Task: OSPF
```

```
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 5
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100016
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:59        Metric: 2
                Task: RSVP
                Announcement bits (1): 1-Resolve tree 2
                AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp2.0, selected
                State: <Active Int>
                Age: 2d 1:44:20
                Task: IF
                AS path: I
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
                Inactive reason: No difference
                Age: 2d 1:44:20
                Task: IF
                AS path: I
```

show route best extensive

The output for the `show route best extensive` command is identical to that for the `show route best detail` command. For sample output, see "show route best detail" on page 23550.

show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   O  10          2            >10.31.1.6
                                                   so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   R   7          2            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.0.0.0/8         D   0                        >fxp2.0
                     D   0                        >fxp1.0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show route brief | 23553
- show route detail | 23573
- show route extensive | 23619

show route terse | 23824

show route brief

IN THIS SECTION

- Syntax | 23553
- Syntax (EX Series Switches) | 23553
- Description | 23553
- Options | 23554
- Required Privilege Level | 23554
- Output Fields | 23554
- Sample Output | 23554
- Release Information | 23555

Syntax

```
show route brief
<destination-prefix>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route brief
<destination-prefix>
```

Description

Display brief information about the active entries in the routing tables.

Options

none	Display all active entries in the routing table.
destination-prefix	(Optional) Display active entries for the specified address or range of addresses.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the Output Field table of the `show route` command.

Sample Output

show route brief

```
user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                      Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                    > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.0.0/18         *[Static/5] 1w5d 20:30:29
                    > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                    > via fxp0.0
```

```
192.168.164.51/32  *[Local/0] 2w4d 13:11:14
                      Local via fxp0.0
207.17.136.192/32  *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16     *[Direct/0] 1w5d 20:30:28
                    > via fe-0/0/3.0
100.101.2.3/32     *[Local/0] 1w5d 20:30:28
                      Local via fe-0/0/3.0
172.16.233.5/32       *[OSPF/10] 1w5d 20:30:29, metric 1
                      MultiRecv
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show route | 23512
- show route all | 23542
- show route best | 23548

show route ccc

IN THIS SECTION

- Syntax | 23556
- Description | 23556
- Options | 23556
- Required Privilege Level | 23556
- Output Fields | 23556
- Sample Output | 23556
- Release Information | 23557

Syntax

```
show route ccc ccc
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Description

Display circuit cross-connect (CCC) entries in the Multiprotocol Link Switching (MPLS) routing table.

Options

ccc	Name of an entry with a circuit cross-connect interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the "show route" on page 23512 command, the "show route detail" on page 23573 command, the "show route extensive" on page 23619 command, or the "show route terse" on page 23824 command.

Sample Output

show route ccc extensive

```
user@host> show route ccc fe-0/1/0.600 extensive
mpls.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
fe-0/1/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel fe-0/1/2.600.0      /16 -> {0.0.0.0}
```

```
        *CCC    Preference: 7
                Next-hop reference count: 2
                Next hop: via so-0/0/3.0 weight 0x1, selected
                Label operation: Push 101424
                State: <Active Int>
                Local AS:   100
                Age: 28:13   Metric: 3
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show connections | 19176

show route community

IN THIS SECTION

- Syntax | 23558
- Syntax (EX Series Switches) | 23558
- Description | 23558
- Options | 23558
- Additional Information | 23559
- Required Privilege Level | 23559
- Output Fields | 23559
- Sample Output | 23559
- Release Information | 23559

Syntax

```
show route community as-number:community-value
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route community as-number:community-value
<brief | detail | extensive | terse>
```

Description

Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community.

Options

as-number:community-value

One or more community identifiers. *as-number* is the AS number, and *community-value* is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks. Community identifiers can include wildcards.

For example:

```
user@host> show route table inet.0 protocol bgp community
"12083:6015" community "12083:65551"
```

or

```
user@host> show route table inet.0 protocol bgp community
[12083:6014 12083:65551]
```

brief | detail | extensive | terse

(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)

(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

Specifying the community option displays all routes matching the community found within the routing table. The community option does not limit the output to only the routes being advertised to the neighbor after any egress routing policy.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the "show route" on page 23512 command, the "show route detail" on page 23573 command, the "show route extensive" on page 23619 command, or the "show route terse" on page 23824 command.

Sample Output

show route community

```
user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.4.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                    AS Path: {666} 234 2548 1 IGP
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
172.16.6.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                    AS Path: {666} 234 2548 568 721 Incomplete
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
172.16.92.0/16         *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
                    AS Path: {666} 234 2548 1673 1675 1747 IGP
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show route detail | 23573

show route community-name

IN THIS SECTION

- Syntax | 23560
- Syntax (EX Series Switches) | 23560
- Description | 23560
- Options | 23561
- Required Privilege Level | 23561
- Output Fields | 23561
- Sample Output | 23561
- Release Information | 23562

Syntax

```
show route community-name community-name
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route community-name community-name
<brief | detail | extensive | terse>
```

Description

Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community, specified by a community name.

Options

community-name	Name of the community.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the "show route" on page 23512 command, the "show route detail" on page 23573 command, the "show route extensive" on page 23619 command, or the "show route terse" on page 23824 command.

Sample Output

show route community-name

```
user@host> show route community-name red-com
inet.0: 17 destinations, 17 routes (16 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

instance1.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.212/32  *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                      AS path: 300 I
                    > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
172.16.20.20/32     *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                      AS path: I
                    > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
172.16.100.0/24       *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
```

```
                      AS path: I
                    > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.204:10:10.255.245.212/32
                   *[BGP/170] 00:06:40, localpref 100, from 10.255.245.204
                      AS path: 300 I
                    > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:172.16.20.20/32
                   *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
                      AS path: I
                    > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:100.1.4.0/24
                   *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
                      AS path: I
                    > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

instance1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Release Information

Command introduced before Junos OS Release 7.4.

show route cumulative

IN THIS SECTION

- Syntax | 23563
- Description | 23563

- Options | 23563
- Required Privilege Level | 23563
- Output Fields | 23563
- Sample Output | 23564
- Release Information | 23564

Syntax

```
show route cumulative
<fabric>
<logical-system (all | logical-system-name)>
<vpn-family (inet.0 | inet6.0)>
```

Description

Shows the cumulative number of either IPv4 or IPv6 routes in the VRF table.

Options

fabric	Internal fabric state.
logical-system (all | *logical-system-name*)	(Optional) Show cumulative routes on all logical systems or on a particular logical system.
vpn-family (inet.0 | inet6.0)	Enter `inet.0` for IPv4 routes or `inet6.0` for IPv6 routes.

Required Privilege Level

view

Output Fields

Field Name	Field Description
destinations	Number of destinations for which there are VRF routes in the routing table.
routes	Number of VRF routes in the routing table: • active—Number of routes that are active. • holddown—Number of VRF routes that are in the hold-down state before being declared inactive. • hidden—Number of VRF routes that are not used because of routing policy.

Sample Output

show route cumulative

```
user@host> show route cumulative vpn-family inet.0
Total VRF Routes: 720 destinations, 722 routes (720 active, 0 holddown, 0 hidden)
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

show route summary | 23779

show route damping

IN THIS SECTION

Syntax | 23565

- Syntax (EX Series Switch and QFX Series) | 23565
- Description | 23565
- Options | 23565
- Required Privilege Level | 23566
- Output Fields | 23566
- Sample Output | 23571
- Release Information | 23572

Syntax

```
show route damping (decayed | history | suppressed)
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and QFX Series)

```
show route damping (decayed | history | suppressed)
<brief | detail | extensive | terse>
```

Description

Display the BGP routes for which updates might have been reduced because of route flap damping.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.
decayed	Display route damping entries that might no longer be valid, but are not suppressed.
history	Display entries that have already been withdrawn, but have been logged.

logical-system (all | *logical-system-name*) (Optional) Perform this operation on all logical systems or on a particular logical system.

suppressed Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.

Required Privilege Level

view

Output Fields

Table 1080 on page 23566 lists the output fields for the `show route damping` command. Output fields are listed in the approximate order in which they appear.

Table 1080: show route damping Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, `inet.0`.	All levels
`destinations`	Number of destinations for which there are routes in the routing table.	All levels
number `routes`	Number of routes in the routing table and total number of routes in the following states: • `active` • `holddown` (routes that are in a pending state before being declared inactive) • `hidden` (the routes are not used because of a routing policy)	All levels
destination-prefix (`entry`, `announced`)	Destination prefix. The `entry` value is the number of routes for this destination, and the `announced` value is the number of routes being announced for this destination.	`detail extensive`

Table 1080: show route damping Output Fields *(Continued)*

Field Name	Field Description	Level of Output
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a `to` line indicates the best subpath to the route. In every routing metric except for the BGP `LocalPref` attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the `LocalPref` value in the `Preference2` field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the `LocalPref` value for Route 2 is 155, the `Preference2` value is -156. Route 2 is preferred because it has a higher `LocalPref` value and a lower `Preference2` value.	All levels
`Next-hop reference count`	Number of references made to the next hop.	`detail extensive`
`Source`	IP address of the route source.	`detail extensive`
`Next hop`	Network layer address of the directly reachable neighboring system.	`detail extensive`
`via`	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word `Selected`.	`detail extensive`
`Protocol next hop`	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.	`detail extensive`
`Indirect next hop`	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.	`detail extensive`

Table 1080: show route damping Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Flags for this route. For a description of possible values for this field, see the output field table for the "show route detail" on page 23573 command.	detail extensive
Local AS	AS number of the local routing device.	detail extensive
Peer AS	AS number of the peer routing device.	detail extensive
Age	How long the route has been known.	detail extensive
Metric	Metric for the route.	detail extensive
Task	Name of the protocol that has added the route.	detail extensive
Announcement bits	List of protocols that announce this route. *n*-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. *n* is an index used by Juniper Networks customer support only.	detail extensive

Table 1080: show route damping Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`AS path`	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • `I`—IGP. • `E`—EGP. • `?`—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • `[ ]`—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured. • `{ }`—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • `( )`—Parentheses enclose a confederation. • `( [ ] )`—Parentheses and brackets enclose a confederation set. **NOTE**: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
`to`	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	`brief` none
`via`	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word `Selected`.	`brief` none
`Communities`	Community path attribute for the route. See the output field table for the "show route detail" on page 23573 command.	`detail extensive`

Table 1080: show route damping Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Localpref`	Local preference value included in the route.	All levels
`Router ID`	BGP router ID as advertised by the neighbor in the open message.	`detail extensive`
`Merit (last update/now)`	Last updated and current figure-of-merit value.	`detail extensive`
`damping-parameters`	Name that identifies the damping parameters used, which is defined in the damping statement at the `[edit policy-options]` hierarchy level.	`detail extensive`
`Last update`	Time of most recent change in path attributes.	`detail extensive`
`First update`	Time of first change in path attributes, which started the route damping process.	`detail extensive`
`Flaps`	Number of times the route has gone up or down or its path attributes have changed.	`detail extensive`
`Suppressed`	(`suppressed` keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it.	All levels
`Reusable in`	(`suppressed` keyword only) Time when a suppressed route will again be available.	All levels
`Preference will be`	(`suppressed` keyword only) Preference value that will be applied to the route when it is again active.	All levels

Sample Output

show route damping decayed detail

```
user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083 hidden)
10.0.111.0/24 (7 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 151973
                Source: 172.23.2.129
                Next hop: via so-1/2/0.0
                Next hop: via so-5/1/0.0, selected
                Next hop: via so-6/0/0.0
                Protocol next hop: 172.23.2.129
                Indirect next hop: 89a1a00 264185
                State: <Active Ext>
                Local AS: 64500 Peer AS:   64490
                Age: 3:28       Metric2: 0
                Task: BGP_64490.172.23.2.129+179
                Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179
         6-Resolve tree 2 7-Resolve tree 3
                AS path: 64499 64510 645511 645511 645511 645511 I ()
                Communities: 65551:390 65551:2000 65551:3000 65550:701
                Localpref: 100
                Router ID: 172.23.2.129
                Merit (last update/now): 1934/1790
                damping-parameters: damping-high
                Last update:       00:03:28 First update:       00:06:40
                Flaps: 2
```

show route damping history

```
user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122 hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15       [BGP ] 2d 22:47:58, localpref 100
                      AS path: 64220 65541 65542 I
                    > to 192.168.60.85 via so-3/1/0.0
```

show route damping history detail

```
user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105 hidden)
10.108.0.0/15 (3 entries, 1 announced)
         BGP                 /-101
                Next-hop reference count: 69058
                Source: 192.168.60.85
                Next hop: 192.168.60.85 via so-3/1/0.0, selected
                State: <Hidden Ext>
                Inactive reason: Unusable path
                Local AS: 64500 Peer AS:  64220
                Age: 2d 22:48:10
                Task: BGP_64220.192.168.60.85+179
                AS path: 64220 65541 65542 I ()
                Communities: 65541:390 65541:2000 65541:3000 65504:3561
                Localpref: 100
                Router ID: 192.168.80.25
                Merit (last update/now): 1000/932
                damping-parameters: set-normal
                Last update:       00:01:05 First update:       00:01:05
                Flaps: 1
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear bgp damping

show policy damping | 23206

show route detail

IN THIS SECTION

- Syntax | 23573
- Syntax (EX Series Switches) | 23573
- Description | 23573
- Options | 23573
- Required Privilege Level | 23574
- Output Fields | 23574
- Sample Output | 23591
- Release Information | 23608

Syntax

```
show route detail
<destination-prefix>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route detail
<destination-prefix>
```

Description

Display detailed information about the active entries in the routing tables.

Options

none Display all active entries in the routing table on all systems.

destination-prefix	(Optional) Display active entries for the specified address or range of addresses.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1081 on page 23574 describes the output fields for the `show route detail` command. Output fields are listed in the approximate order in which they appear.

Table 1082 on page 23584 describes all possible values for the Next-hop Types output field.

Table 1083 on page 23586 describes all possible values for the State output field. A route can be in more than one state (for example, `<Active NoReadvrt Int Ext>`).

Table 1084 on page 23589 describes the possible values for the Communities output field.

Table 1081: show route detail Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number destinations	Number of destinations for which there are routes in the routing table.
number routes	Number of routes in the routing table and total number of routes in the following states: • `active` (routes that are active) • `holddown` (routes that are in the pending state before being declared inactive) • `hidden` (routes that are not used because of a routing policy)

Table 1081: show route detail Output Fields *(Continued)*

Field Name	Field Description

route-destination (entry, announced)	Route destination (for example:10.0.0.1/24). The `entry` value is the number of routes for this destination, and the `announced` value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• *`MPLS-label`* (for example, 80001).
	• *`interface-name`* (for example, ge-1/0/2).
	• *`neighbor-address`*:*`control-word-status`*:*`encapsulation type`*:*`vc-id`*:*`source`* (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  • *`neighbor-address`*—Address of the neighbor.
	  • *`control-word-status`*—Whether the use of the control word has been negotiated for this virtual circuit: `NoCtrlWord` or `CtrlWord`.
	  • *`encapsulation type`*—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  • *`vc-id`*—Virtual circuit identifier.
	  • *`source`*—Source of the advertisement: `Local` or `Remote`.
	  • *`source`*—Source of the advertisement: `Local` or `Remote`.

label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • `S=0 route` indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no `S=` information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).

Table 1081: show route detail Output Fields *(Continued)*

Field Name	Field Description

[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route.
	• +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.
	• - —A hyphen indicates the last active route.
	• *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a `to` line indicates the best subpath to the route.
	
	In every routing metric except for the BGP `LocalPref` attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the `LocalPref` value in the `Preference2` field. For example, if the `LocalPref` value for Route 1 is 100, the `Preference2` value is -101. If the `LocalPref` value for Route 2 is 155, the `Preference2` value is -156. Route 2 is preferred because it has a higher `LocalPref` value.
	
	`Preference2` values are signed integers, that is, `Preference2` values can be either positive or negative values. However, Junos OS evaluates `Preference2` values as unsigned integers that are represented by positive values. Based on the `Preference2` values, Junos OS evaluates a preferred route differently in the following scenarios:
	• Both Signed `Preference2` values
	  • Route A = -101
	  • Route B = -156
	  Where both the `Preference2` values are signed, Junos OS evaluates only the unsigned value of `Preference2` and Route A, which has a lower `Preference2` value is preferred.
	• Unsigned `Preference2` values
	  Now consider both unsigned `Preference2` values:
	  • Route A = 4294967096
	  • Route B = 200
	  Here, Junos OS considers the lesser `Preference2` value and Route B with a `Preference2` value of 200 is preferred because it is less than 4294967096.
	• Combination of signed and unsigned `Preference2` values

Table 1081: show route detail Output Fields *(Continued)*

Field Name	Field Description
	When `Preference2` values of two routes are compared, and for one route the `Preference2` is a signed value, and for the other route it is an unsigned value, Junos OS prefers the route with the positive `Preference2` value over the negative `Preference2` value. For example, consider the following signed and unsigned `Preference2` values: • Route A = -200 • Route B = 200 In this case, Route B with a `Preference2` value of 200 is preferred although this value is greater than -200, because Junos OS evaluates only the unsigned value of the `Preference2` value.
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 1082 on page 23584.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.

Table 1081: show route detail Output Fields *(Continued)*

Field Name	Field Description
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word `Selected`. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be `pop` (where a label is removed from the top of the stack), `push` (where another label is added to the label stack), or `swap` (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 1083 on page 23586.
Local AS	AS number of the local routing device.

Table 1081: show route detail Output Fields *(Continued)*

Field Name	Field Description
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see `show route table` .
Task	Name of the protocol that has added the route.
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the KRT for installing the route into the Packet Forwarding Engine, to a resolve tree, a L2 VC, or even a VPN. For example, *n*-`Resolve inet` indicates that the specified route is used for route resolution for next hops found in the routing table. • *n*—An index used by Juniper Networks customer support only.

Table 1081: show route detail Output Fields *(Continued)*

Field Name	Field Description
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • `I`—IGP. • `E`—EGP. • `Recorded`—The AS path is recorded by the sample process (sampled). • `?`—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • `[ ]`—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • `[ ]`—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • `{ }`—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • `( )`—Parentheses enclose a confederation. • `( [ ] )`—Parentheses and brackets enclose a confederation set. **NOTE**: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.

Table 1081: show route detail Output Fields *(Continued)*

Field Name	Field Description
validation-state	(BGP-learned routes) Validation status of the route: • `Invalid`—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • `Unknown`—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • `Unverified`—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • `Valid`—Indicates that the prefix and autonomous system pair are found in the database.
ORR Generation-ID	Displays the optimal route reflection (ORR) generation identifier. ISIS and OSPF interior gateway protocol (IGP) updates filed whenever any of the corresponding ORR route has its metric valued changed, or if the ORR route is added or deleted.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.

Table 1081: show route detail Output Fields *(Continued)*

Field Name	Field Description
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 1084 on page 23589 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: `none` or `Site Down`.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.

Table 1081: show route detail Output Fields *(Continued)*

Field Name	Field Description
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table
ImportAccepted LongLivedStaleImport DeletePending	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. The DeletePending flag indicates that a BGP route needs to be processed due to a BGP peer down event.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Statistics ID Group	Indicates the Kernel ID number and Statistics ID number.
Statistics	Indicates the number of packets and data transferred.

Table 1082: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Dynamic List	Dynamic list next hop
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.

Table 1082: Next-hop Types Output Field Values *(Continued)*

Next-Hop Type	Description
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	A specific node or set of nodes to which the routing device forwards packets that match the route prefix. To qualify as next-hop type router, the route must meet the following criteria: • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Software	Next hop added to the Routing Engine forwarding table for remote IP addresses with prefix /32 for Junos OS Evolved only.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 1083: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.

Table 1083: State Output Field Values *(Continued)*

Value	Description
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.

Table 1083: State Output Field Values *(Continued)*

Value	Description
NSR-incapable	Route added by non-NSR supported protocols.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Programmed	Route installed programatically by on-box or off-box applications using API.
ProtectionCand	Indicates paths requesting protection.
ProtectionPath	Indicates the route entry that can be used as a protection path.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.

Table 1083: State Output Field Values *(Continued)*

Value	Description
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 1084: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: *local AS number:link-bandwidth-number*	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
link-bandwidth-number	Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).
local AS number	Local AS number: from 1 through 65,535.
options	1 byte. Currently this is only used if the route type is 5 or 7. Setting the least significant bit in the field indicates that the route carries a type 2 metric.

Table 1084: Communities Output Field Values *(Continued)*

Value	Description
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be0); 7 for NSSA routes; or 129 for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000. The format is *area-number:ospf-route-type:options*.
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306. The format is *area-number:ospf-route-type:options*.
target	Defines which VPN the route participates in; target has the format *32-bit IP address:16-bit number*. For example, 10.19.0.0:100.
ipv6-origin	Defines the source of the IPv6 unicast address in a policy match condition.
ipv6-target	Defines the VPN IPv6 target unicast address used in a policy match condition. The target has the *128-bit IP address:16-bit number* format. For example, 2001:db8:9999::9.
ipv6-extended	Defines the extended format of the IPv6 unicast address in a policy match. For example, `ipv6-extended:203:<2001:db8::7>:67`
unknown IANA	Incoming IANA codes with a value between 0x1 and 0x7fff. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above 0x8000. This code of the BGP extended community attribute is accepted, but it is not recognized.

Table 1084: Communities Output Field Values *(Continued)*

Value	Description
Validation State	Displays the origin validation information for each route entry and a reason code if the route is marked invalid. The possible reasons for an invalid route entry include mismatch due to `origin-as` or due to `maximum-length`.

Sample Output

show route detail

```
user@host> show route detail

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    65069
                Age: 1:31:43
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

10.31.1.0/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    65069
                Age: 1:30:17
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
```

```
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    65069
                Age: 1:30:17    Metric: 1
                ORR Generation-ID: 1
                        Area: 0.0.0.0
                Task: OSPF
                AS path: I

10.31.1.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    65069
                Age: 1:30:20
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    65069
                Age: 1:29:56    Metric: 2
                Area: 0.0.0.0
                ORR Generation-ID: 1
                        Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

172.16.233.2/32 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
```

```
                Local AS:    65069
                Age: 1:31:45
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

172.16.233.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    65069
                Age: 1:31:43
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    65069
                Age: 1:25:49    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

10.255.71.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    65069
                Age: 1:25:49    Metric: 1
```

```
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    65069
                Age: 1:31:44
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0   (1 entry, 1 announced)
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    65069
                Age: 1:31:45    Metric: 1
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
```

```
                  Address: 0x9174c28
                  Next-hop reference count: 2
                  Next hop: 172.16.0.2 via lt-1/2/0.9 weight 0x1
                  Label-switched-path R2-to-R4-2p2mp
                  Label operation: Pop
                  Next hop type: Router, Next hop index: 1048574
                  Address: 0x92544f0
                  Next-hop reference count: 2
                  Next hop: 172.16.0.2 via lt-1/2/0.7 weight 0x1
                  Label-switched-path R2-to-R200-p2mp
                  Label operation: Pop
                  Next hop: 172.16.0.2 via lt-1/2/0.5 weight 0x8001
                  Label operation: Pop
                  State: <Active Int>
                  Age: 1:29       Metric: 1
                  Task: RSVP
                  Announcement bits (1): 0-KRT
                  AS path: I...

800010 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:29:30
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 10.255.70.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:29:30    Metric2: 2
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
```

```
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    65069
                Age: 1:31:44
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    65069
                Age: 1:31:44
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    65069
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    65069
```

```
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    65069
                Age: 1:31:43
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:31:44
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    65069 Peer AS:    65069
                Age: 1:25:49    Metric2: 1
                AIGP 210
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
```

```
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
```

```
                State: <Active Int>
                Local AS:    65099
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512

inet.0: 45 destinations, 47 routes (44 active, 0 holddown, 1 hidden)
10.1.1.3/32 (1 entry, 1 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Router, Next hop index: 580
                Address: 0x9db6ed0
                Next-hop reference count: 8
                Next hop: 10.1.1.6 via lt-1/0/10.5, selected
                Session Id: 0x18a
                State: <Active Int>
                Local AS:     2
                Age: 1:32       Metric: 10
                Validation State: unverified
                ORR Generation-ID: 1
                Task: IS-IS
                Announcement bits (3): 0-KRT 5-Resolve tree 4 6-Resolve_IGP_FRR task
                AS path: I

inet.0: 61 destinations, 77 routes (61 active, 1 holddown, 0 hidden)
10.1.1.1/32 (2 entries, 1 announced)
        *OSPF   Preference: 10
                Next hop type: Router, Next hop index: 673
                Address: 0xc008830
                Next-hop reference count: 3
                Next hop: 10.1.1.8 via ge-0/0/2.0, selected
                Session Id: 0x1b7
                State: <Active Int>
                Local AS:     1
                Age: 3:06:59    Metric: 100
                Validation State: unverified
                ORR Generation-ID: 1
                Area: 0.0.0.0
                Task: OSPF
```

```
                Announcement bits (2): 1-KRT 9-Resolve tree 2
                AS path: I
```

show route programmed detail (with statistics information)

```
user@host> show route programmed detail

inet.0: 104 destinations, 104 routes (103 active, 0 holddown, 1 hidden)

inet.3: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
192.0.2.0/24 (1 entry, 1 announced)
        *Static Preference: 5/100
                Next hop type: Router, Next hop index: 0
                Address: 0xa667694
                Next-hop reference count: 15
                Next hop: 198.51.100.0 via ae0.0
                Label operation: Push 50008, Push 3400(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Load balance label: Label 50008: None; Label 3400: None;
                Label element ptr: 0x84e7e90
                Label parent element ptr: 0x0
                Label element references: 3
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x0
                Statistics ID Group: Kernel ID = 10, Stats IDs = { 536870923 }
                Statistics: Packets 1380, Bytes 681720
                Next hop: 198.51.100.1 via ae0.1
                Label operation: Push 50008, Push 3401(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Load balance label: Label 50008: None; Label 3401: None;
                Label element ptr: 0x84e7ee0
                Label parent element ptr: 0x0
                Label element references: 3
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x0
                Statistics ID Group: Kernel ID = 11, Stats IDs = { 536870924 }
                Statistics: Packets 1444, Bytes 713336
                Next hop: 198.51.100.2 via ae0.2
                Label operation: Push 50008, Push 3410(top)
```

```
                    Label TTL action: prop-ttl, prop-ttl(top)
                    Load balance label: Label 50008: None; Label 3410: None;
                    Label element ptr: 0xa1926a0
                    Label parent element ptr: 0x0
                    Label element references: 3
                    Label element child references: 0
                    Label element lsp id: 0
                    Session Id: 0x0
                    Statistics ID Group: Kernel ID = 12, Stats IDs = { 536870925 }
                    Statistics: Packets 1420, Bytes 701480
                    Next hop: 198.51.100.3 via ge-0/0/2.0
                    Label operation: Push 50008, Push 3411(top)
                    Label TTL action: prop-ttl, prop-ttl(top)
                    Load balance label: Label 50008: None; Label 3411: None;
                    Label element ptr: 0xa1926c8
                    Label parent element ptr: 0x0
                    Label element references: 3
                    Label element child references: 0
                    Label element lsp id: 0
                    Session Id: 0x0
                    Statistics ID Group: Kernel ID = 13, Stats IDs = { 536870926 }
                    Statistics: Packets 1486, Bytes 734084
                    Next hop: 198.51.100.4 via ge-0/0/2.1, selected
                    Label operation: Push 50008, Push 3420(top)
                    Label TTL action: prop-ttl, prop-ttl(top)
                    Load balance label: Label 50008: None; Label 3420: None;
                    Label element ptr: 0xa1926f0
                    Label parent element ptr: 0x0
                    Label element references: 3
                    Label element child references: 0
                    Label element lsp id: 0
                    Session Id: 0x0
                    Statistics ID Group: Kernel ID = 14, Stats IDs = { 536870927 }
                    Statistics: Packets 1399, Bytes 691106
                    Next hop: 198.51.100.5 via ge-0/0/2.2
                    Label operation: Push 50008, Push 3421(top)
                    Label TTL action: prop-ttl, prop-ttl(top)
                    Load balance label: Label 50008: None; Label 3421: None;
                    Label element ptr: 0xa192718
                    Label parent element ptr: 0x0
                    Label element references: 3
                    Label element child references: 0
                    Label element lsp id: 0
```

```
                Session Id: 0x0
                Statistics ID Group: Kernel ID = 15, Stats IDs = { 536870928 }
                Statistics: Packets 1455, Bytes 718770
                Next hop: 198.51.100.6 via ge-0/0/4.0
                Label operation: Push 50008, Push 3450(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Load balance label: Label 50008: None; Label 3450: None;
                Label element ptr: 0xa192740
                Label parent element ptr: 0x0
                Label element references: 3
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x0
                Statistics ID Group: Kernel ID = 16, Stats IDs = { 536870929 }
                Statistics: Packets 1407, Bytes 695058
                Next hop: 198.51.100.7 via ge-0/0/4.1
                Label operation: Push 50008, Push 3460(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Load balance label: Label 50008: None; Label 3460: None;
                Label element ptr: 0xa192768
                Label parent element ptr: 0x0
                Label element references: 3
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x0
                Statistics ID Group: Kernel ID = 17, Stats IDs = { 536870930 }
                Statistics: Packets 1418, Bytes 700492
                Next hop: 198.51.100.8 via ge-0/0/4.2
                Label operation: Push 50008, Push 3470(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Load balance label: Label 50008: None; Label 3470: None;
                Label element ptr: 0xa192790
                Label parent element ptr: 0x0
                Label element references: 3
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x0
                Statistics ID Group: Kernel ID = 18, Stats IDs = { 536870931 }
                Statistics: Packets 0, Bytes 0
                State: <Active Int NSR-incapable Programmed>
                Age: 12:46
                Validation State: unverified
                Announcement bits (3): 0-Resolve tree 6 1-Resolve tree 1 3-Resolve_IGP_FRR task
```

```
                AS path: I
                Session-IDs associated:
                Session-id: 370 Version: 0
                Thread: junos-main
```

show route detail (with IPv6 unicast address community)

```
user@host> show route 203.0.112.1 detail
inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
203.0.112.1/32 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 559
                Address: 0x7b5de84
                Next-hop reference count: 12, key opaque handle: 0x0, non-key opaque handle: 0x0
                Source: 13.1.1.1
                Next hop: 13.1.1.1 via ge-0/0/0.0, selected
                Session Id: 320
                State: <Active Ext>
                Local AS:   200 Peer AS:   100
                Age: 8:41       Metric: 1
                Validation State: unverified
                Task: BGP_100.13.1.1.1
                Announcement bits (1): 0-KRT
                AS path: 100 I
                Communities: target:100:65 target:172.16.7.7:2 ipv6-target:<2001:db8::5>:65
                ipv6-target:<2001:db8::5>:101 ipv6-origin:<2001:db8::6>:66
                ipv6-origin:<2001:db8::6>:102 ipv6-extended:100:<2001:db8::7>:67
                ipv6-extended:256:<2001:db8::7>:103
                Accepted
                Localpref: 100
                Router ID: 2.2.2.2
                Thread: junos-main
```

show route detail (with BGP Multipath)

```
user@host> show route detail

10.1.1.8/30 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 262142
```

```
                Address: 0x901a010
                Next-hop reference count: 2
                Source: 10.1.1.2
                Next hop: 10.1.1.2 via ge-0/3/0.1, selected
                Next hop: 10.1.1.6 via ge-0/3/0.5
                State: <Active Ext>
                Local AS:     1 Peer AS:     2
                Age: 5:04:43
                  Validation State: unverified
                Task: BGP_2.10.1.1.2+59955
                Announcement bits (1): 0-KRT
                AS path: 65002 I
                Accepted Multipath
                Localpref: 100
                Router ID: 172.16.1.2
         BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 678
                Address: 0x8f97520
                Next-hop reference count: 9
                Source: 10.1.1.6
                Next hop: 10.1.1.6 via ge-0/3/0.5, selected
                State: <NotBest Ext>
                Inactive reason: Not Best in its group - Active preferred
                Local AS:     65001 Peer AS:     65002
                Age: 5:04:43
                   Validation State: unverified
                Task: BGP_2.10.1.1.6+58198
                AS path: 65002 I
                Accepted MultipathContrib
                Localpref: 100
                Router ID: 172.16.1.3
```

show route detail (with BGP, DeletePending)

```
user@host> show route detail
2001:db8:2:1:10.1.1.12/30 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 65002:1
                Next hop type: Indirect
                Address: 0x95c4ee8
                Next-hop reference count: 6
```

```
                Source: 10.1.1.4
                Next hop type: Router, Next hop index: 809
                Next hop: 10.1.1.6 via lt-1/0/10.5, selected
                Label operation: Push 299888, Push 299792(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Load balance label: Label 299888: None; Label 299792: None;
                Session Id: 0x142
                Protocol next hop: 10.1.1.4
                Label operation: Push 299888
                Label TTL action: prop-ttl
                Load balance label: Label 299888: None;
                Indirect next hop: 0x96f0110 1048574 INH Session ID: 0x14e
                State: <Active Int Ext ProtectionPath ProtectionCand>
                Local AS:     65002 Peer AS:     65002
                Age: 2w1d 17:42:45      Metric2: 1
                Validation State: unverified
                Task: BGP_10.2.1.1.4+55190
                AS path: I
                Communities: target:65002:1
                Import Accepted DeletePending
                VPN Label: 299888
                Localpref: 100
                Router ID: 10.1.1.4
                Secondary Tables: red.inet.0
```

show route detail (with Validate State)

```
user@host>show route detail

inet.0: 48 destinations, 51 routes (48 active, 0 holddown, 0 hidden)
1.1.1.1/32 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 599
                Address: 0x7106094
                Next-hop reference count: 7
                Source: 10.1.1.1
                Next hop: 10.1.1.1 via ge-0/0/0.0, selected
                Session Id: 0x140
                State: <Active Ext>
                Local AS:     2 Peer AS:     1
                Age: 2w5d 4:29:09
```

```
                Validation State: [valid | invalid | unknown | unverified ]  +
                Task: RT
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
                State: **Active Int>
                Local AS:  1001
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1, src
192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
```

```
                Next-hop reference count: 3
                Next hop: 1.3.8.2 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 1.3.4.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    10
                Age: 54:05      Metric: 1
                Validation State: unverified
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                FECs bound to route: P2MP root-addr 172.16.1.1, grp: 232.1.1.1, src:
192.168.219.11
                Primary Upstream : 172.16.1.3:0--172.16.1.2:0
                   RPF Nexthops  :
                        ge-1/2/15.0, 10.1.94.1, Label: 301568, weight: 0x1
                        ge-1/2/14.0, 10.1.3.1, Label: 301568, weight: 0x1
                Backup Upstream : 172.16.1.3:0--172.16.1.6:0
                   RPF Nexthops :
                        ge-1/2/20.0, 10.1.96.1, Label: 301584, weight: 0xfffe
                        ge-1/2/19.0, 10.1.6.1, Label: 301584, weight: 0xfffe
```

show route detail (Flexible VXLAN Tunnel Profile)

```
user@host> show route 192.168.0.2 detail
...

CUSTOMER_0001.inet.0: 5618 destinations, 6018 routes (5618 active, 0 holddown, 0 hidden)

192.168.0.2/32 (1 entry, 1 announced)
        *Static Preference: 5/100
                Next hop type: Router, Next hop index: 74781
                Address: 0x5d9b03cc
                Next-hop reference count: 363
                Next hop: via fti0.6, selected
```

```
                Session Id: 0x24c8
                State: <Active Int NSR-incapable OpaqueData Programmed>
                Age: 1:25:53
                Validation State: unverified
                        Tag: 10000001   Tag2: 1
                Announcement bits (2): 1-KRT 3-Resolve tree 30
                AS path: I
                Flexible IPv6 VXLAN tunnel profile
                    Action: Encapsulate
                    Interface: fti0.6 (Index: 10921)
                    VNI: 10000001
                    Source Prefix: 2001:db8:255::2/128
                    Source UDP Port Range: 54614 - 60074
                    Source MAC Address: 00:00:5e:00:52:01
                    Destination Address: 2001:db8:10:1:1:1:0:1
                    Destination UDP Port: 4790
                    VXLAN Flags: 0x08
...
```

Release Information

Command introduced before Junos OS Release 7.4.

Comand introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

`DeletePending` flag added to the command output in Junos OS Release 19.4R1.

`ipv6-origin`, `ipv6-target`, and `ipv6-extended Community` output fields supported in Junos OS Release and Junos OS Evolved Release 23.1.

show route exact

IN THIS SECTION

- Syntax | 23609
- Syntax (EX Series Switches) | 23609
- Description | 23609
- Options | 23609

- Required Privilege Level | 23609
- Output Fields | 23610
- Sample Output | 23610
- Release Information | 23611

Syntax

```
show route exact destination-prefix
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route exact destination-prefix
<brief | detail | extensive | terse>
```

Description

Display only the routes that exactly match the specified address or range of addresses.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to `brief`.
destination-prefix	Address or range of addresses.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the `show route` command, the `show route detail` command, the `show route extensive` command, or the `show route terse` command.

Sample Output

show route exact

```
user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                    > to 192.168.71.254 via fxp0.0
```

show route exact detail

```
user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2d 3:30:26
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact extensive

```
user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
```

```
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:25:18
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact terse

```
user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 207.17.136.0/24    S   5                        >192.168.71.254
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show route | 23512
- show route detail | 23573
- show route extensive | 23619
- show route terse | 23824

show route export

IN THIS SECTION

- Syntax | 23612
- Syntax (EX Series Switches) | 23612
- Description | 23612
- Options | 23613
- Required Privilege Level | 23613
- Output Fields | 23613
- Sample Output | 23614
- Release Information | 23615

Syntax

```
show route export
<brief | detail>
<instance <instance-name> | routing-table-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route export
<brief | detail>
<instance <instance-name> | routing-table-name>
```

Description

Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.

Options

none	(Same as `brief`.) Display standard information about policy-based export for all instances and routing tables on all systems.
brief | detail	(Optional) Display the specified level of output.
instance <*instance-name*>	(Optional) Display a particular routing instance for which policy-based export is currently enabled.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
routing-table-name	(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the `show route export inet` command).

Required Privilege Level

view

Output Fields

Table 1085 on page 23613 lists the output fields for the `show route export` command. Output fields are listed in the approximate order in which they appear.

Table 1085: show route export Output Fields

Field Name	Field Description	Level of Output
Table or *table-name*	Name of the routing tables that either import or export routes.	All levels
Routes	Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.	`brief` none
Export	Whether the table is currently exporting routes to other tables: Y or N (Yes or No).	`brief` none

Table 1085: show route export Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Import	Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)	`detail`
Flags	(`instance` keyword only) Flags for this feature on this instance: • `config auto-policy`—The policy was deduced from the configured IGP export policies. • `cleanup`—Configuration information for this instance is no longer valid. • `config`—The instance was explicitly configured.	`detail`
Options	(`instance` keyword only) Configured option displays the type of routing tables the feature handles: • `unicast`—Indicates *instance*.inet.0. • `multicast`—Indicates *instance*.inet.2. • `unicast multicast`—Indicates *instance*.inet.0 and *instance*.inet.2.	`detail`
Import policy	(`instance` keyword only) Policy that `route export` uses to construct the import-export matrix. Not displayed if the instance type is `vrf`.	`detail`
Instance	(`instance` keyword only) Name of the routing instance.	`detail`
Type	(`instance` keyword only) Type of routing instance: `forwarding`, `non-forwarding`, or `vrf`.	`detail`

Sample Output

show route export

```
user@host> show route export
Table                           Export          Routes
```

```
inet.0                          N                    0
black.inet.0                    Y                    3
red.inet.0                      Y                    4
```

show route export detail

```
user@host> show route export detail
inet.0                                          Routes:        0
black.inet.0                                    Routes:        3
  Import: [ inet.0 ]
red.inet.0                                      Routes:        4
  Import: [ inet.0 ]
```

show route export instance detail

```
user@host> show route export instance detail
Instance: master                        Type: forwarding
  Flags: <config auto-policy> Options: <unicast multicast>
  Import policy: [ (ospf-master-from-red || isis-master-from-black) ]
Instance: black                         Type: non-forwarding
Instance: red                           Type: non-forwarding
```

Release Information

Command introduced before Junos OS Release 7.4.

show route export vrf-target

IN THIS SECTION

- Syntax | 23616
- Description | 23616
- Options | 23616

- Required Privilege Level | 23616
- Output Fields | 23617
- Sample Output | 23618
- Release Information | 23618

Syntax

```
show route export vrf-target
<brief | detail>
<community community--regular-expression>
<logical-system (all | logical-system-name)>
```

Description

Display the VPN routing and forwarding (VRF) target communities for which policy-based route export is currently distributing routes. This command is relevant when there are overlapping virtual private networks (VPNs).

Options

none	Display standard information about all target communities.
brief | detail	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.
community *community-regular-expression*	(Optional) Display information about the specified community.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1086 on page 23617 lists the output fields for the `show route export vrf-target` command. Output fields are listed in the approximate order in which they appear.

Table 1086: show route export vrf-target Output Fields

Field Name	Field Description	Level of Output
Route target	Target communities for which auto-export is currently distributing routes.	**brief** none
Family	Routing table entries for the specified family.	**brief** none
type-of-routing-table(s)	Type of routing tables the feature handles: • **unicast**—Indicates ***instance*.inet.0**. • **multicast**—Indicates ***instance*.inet.2**. • **unicast multicast**—Indicates ***instance*.inet.0** and ***instance*.inet.2**.	**brief** none
Import	Number of routing tables that are currently importing routes with this target community. Omitted for tables that are not importing routes.	**brief** none
Export	Number of routing tables that are currently exporting routes with this target community. Omitted for tables that are not exporting routes.	**brief** none
Target	Target communities, family, and options for which auto-export is currently distributing routes.	**detail**
Import table(s)	Name of the routing tables that are importing a particular route target.	**detail**
Export table(s)	Name of the routing tables that are exporting a particular route target.	**detail**

Sample Output

show route export vrf-target

```
user@host> show route export vrf-target
Route Target             Family                  Import     Export
69:1                     inet     unicast             2          2
69:2                     inet     unicast             2          2
```

show route export vrf-target community

```
user@host> show route export vrf-target community target:69:1
Route Target             Family                  Import     Export
69:1                     inet     unicast             2          2
```

show route export vrf-target detail

```
user@host> show route export vrf-target detail
Target: 1:12                              inet     unicast
  Import table(s): vrf-11.inet.0 vrf-12.inet.0
  Export table(s): vrf-12.inet.0
Target: 1:13                              inet     unicast
  Import table(s): vrf-12.inet.0 vrf-13.inet.0
  Export table(s): vrf-13.inet.0
```

Release Information

Command introduced before Junos OS Release 7.4.

show route extensive

IN THIS SECTION

- Syntax | 23619
- Syntax (EX Series Switches) | 23619
- Description | 23619
- Options | 23619
- Required Privilege Level | 23620
- Output Fields | 23620
- Sample Output | 23631
- Release Information | 23643

Syntax

```
show route extensive
<destination-prefix>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route extensive
<destination-prefix>
```

Description

Display extensive information about the active entries in the routing tables.

Options

none	Display all active entries in the routing table.

destination-prefix	(Optional) Display active entries for the specified address or range of addresses.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1087 on page 23620 describes the output fields for the `show route extensive` command. Output fields are listed in the approximate order in which they appear.

Table 1087: show route extensive Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number destinations	Number of destinations for which there are routes in the routing table.
number routes	Number of routes in the routing table and total number of routes in the following states: • `active` (routes that are active). • `holddown` (routes that are in the pending state before being declared inactive). • `hidden` (routes that are not used because of a routing policy).

Table 1087: show route extensive Output Fields *(Continued)*

Field Name	Field Description

route-destination (entry, announced)	Route destination (for example: 10.0.0.1/24). The `entry` value is the number of route for this destination, and the `announced` value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• *`MPLS-label`* (for example, 80001).
	• *`interface-name`* (for example, ge-1/0/2).
	• *`neighbor-address`*:*`control-word-status`*:*`encapsulation type`*:*`vc-id`*:*`source`* (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  • *`neighbor-address`*—Address of the neighbor.
	  • *`control-word-status`*—Whether the use of the control word has been negotiated for this virtual circuit: `NoCtrlWord` or `CtrlWord`.
	  • *`encapsulation type`*—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  • *`vc-id`*—Virtual circuit identifier.
	  • *`source`*—Source of the advertisement: `Local` or `Remote`.

TSI	Protocol header information.
label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • `S=0 route` indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no `S=` information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).

Table 1087: show route extensive Output Fields *(Continued)*

Field Name	Field Description
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a `to` line indicates the best subpath to the route. In every routing metric except for the BGP `LocalPref` attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the `LocalPref` value in the `Preference2` field. For example, if the `LocalPref` value for Route 1 is 100, the `Preference2` value is -101. If the `LocalPref` value for Route 2 is 155, the `Preference2` value is -156. Route 2 is preferred because it has a higher `LocalPref` value and a lower `Preference2` value.
Level	(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.

Table 1087: show route extensive Output Fields *(Continued)*

Field Name	Field Description
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word `Selected`. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be `pop` (where a label is removed from the top of the stack), `push` (where another label is added to the label stack), or `swap` (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.
label-operation	MPLS label and operation occurring at this routing device. The operation can be `pop` (where a label is removed from the top of the stack), `push` (where another label is added to the label stack), or `swap` (where a label is replaced by another label).

Table 1087: show route extensive Output Fields *(Continued)*

Field Name	Field Description
Indirect next hops	When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved. When BGP PIC Edge is enabled, the output lines that contain `Indirect next hop: weight` follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values: • 0x1 indicates active next hops. • 0x4000 indicates passive next hops.
State	State of the route (a route can be in more than one state).
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Weight	Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown.

Table 1087: show route extensive Output Fields *(Continued)*

Field Name	Field Description
Inactive reason	If the route is inactive, the reason for its current state is indicated. Typical reasons include: • `Active preferred`—Currently active route was selected over this route. • `Always compare MED`—Path with a lower multiple exit discriminator (MED) is available. • `AS path`—Shorter AS path is available. • `Cisco Non-deterministic MED selection`—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • `Cluster list length`—Path with a shorter cluster list length is available. • `Forwarding use only`—Path is only available for forwarding purposes. • `IGP metric`—Path through the next hop with a lower IGP metric is available. • `IGP metric type`—Path with a lower OSPF link-state advertisement type is available. • `Interior > Exterior > Exterior via Interior`—Direct, static, IGP, or EBGP path is available. • `Local preference`—Path with a higher local preference value is available. • `Next hop address`—Path with a lower metric next hop is available. • `No difference`—Paths from all neighbors are identical for all necessary route attributes in some special scenarios, such as BGP L2VPN/VPLS, and there is no difference. • `Not Best in its group`—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • `Number of gateways`—Path with a higher number of next hops is available. • `Origin`—Path with a lower origin code is available. • `OSPF version`—Path does not support the indicated OSPF version. • `RIB preference`—Route from a higher-numbered routing table is available. • `Route destinguisher`—64-bit prefix added to IP subnets to make them unique.

Table 1087: show route extensive Output Fields *(Continued)*

Field Name	Field Description
	• `Route metric or MED comparison`—Route with a lower metric or MED is available. • `Route preference`—Route with a lower preference value is available. • `Router ID`—Path through a neighbor with a lower ID is available. • `Unusable path`—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • `Update source`—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.
Task	Name of the protocol that has added the route.

Table 1087: show route extensive Output Fields *(Continued)*

Field Name	Field Description
Announcement bits	List of protocols that are consumers of the route. Using the following output as an example, `Announcement bits (3): 0-KRT 5-Resolve tree 2 8-BGP RT Background` there are (3) announcement bits to reflect the three clients (protocols) that have state for this route: Kernel (0-KRT), 5 (resolution tree process 2), and 8 (BGP). The notation *n*-`Resolve inet` indicates that the route is used for route resolution for next hops found in the routing table. *n* is an index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • `I`—IGP. • `E`—EGP. • `Recorded`—The AS path is recorded by the sample process (sampled). • `?`—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • `[ ]`—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • `{ }`—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • `( )`—Parentheses enclose a confederation. • `( [ ] )`—Parentheses and brackets enclose a confederation set. **NOTE**: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.

Table 1087: show route extensive Output Fields *(Continued)*

Field Name	Field Description
validation-state	(BGP-learned routes) Validation status of the route: • `Invalid`—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • `Unknown`—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • `Unverified`—Indicates that origin validation is not enabled for the BGP peers. • `Valid`—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.
route status	Indicates the status of a BGP route: • `Accepted`—The specified BGP route is imported by the default BGP policy. • `Import`—The route is imported into a Layer 3 VPN routing instance. • `Import-Protect`—A remote instance egress that is protected. • `Multipath`—A BGP multipath active route. • `MultipathContrib`—The route is not active but contributes to the BGP multipath. • `Protect`—An egress route that is protected. • `Stale`—A route that is marked stale due to graceful restart.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.

Table 1087: show route extensive Output Fields *(Continued)*

Field Name	Field Description
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route.
DeletePending	The DeletePending flag indicates that a BGP route needs to be processed due to a BGP peer down event.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).

Table 1087: show route extensive Output Fields *(Continued)*

Field Name	Field Description
control flags	Control flags: `none` or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.
Statistics ID Group	Indicates the Kernel ID number and Statistics ID number.
Statistics	Indicates the number of packets and data transferred.

Table 1087: show route extensive Output Fields *(Continued)*

Field Name	Field Description
Validation State	Displays the origin validation information for each route entry and a reason code if the route is marked invalid. The possible reasons for an invalid route entry include mismatch due to `origin-as` or due to `maximum-length`.

Sample Output

show route extensive

```
user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
203.0.113.10/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.10/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    64496
                Age: 1:34:06
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

203.0.113.30/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:40
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
```

```
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    64496
                Age: 1:32:40    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

203.0.113.103/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    644969
                Age: 1:32:43
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

203.0.113.203/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.203/30 -> {203.0.113.216}
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 203.0.113.216 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:19    Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.2/32 -> {}
```

```
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.22/32 -> {}
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

203.0.113.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

203.0.113.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
```

```
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0   (1 entry, 1 announced)
TSI:
KRT in-kernel 0      /36 -> {}
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:08    Metric: 1
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...
```

```
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.9 weight 0x1
                  Label-switched-path R2-to-R4-2p2mp
                  Label operation: Pop
                  Next hop type: Router, Next hop index: 1048574
                  Address: 0x92544f0
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.7 weight 0x1
                  Label-switched-path R2-to-R200-p2mp
                  Label operation: Pop
                  Next hop: 198.51.100.2 via lt-1/2/0.5 weight 0x8001
                  Label operation: Pop
                  State: <Active Int>
                  Age: 1:29       Metric: 1
                  Task: RSVP
                  Announcement bits (1): 0-KRT
                  AS path: I...

800010 (1 entry, 1 announced)

TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:31:53
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
```

```
TSI:
KRT in-kernel vt-3/2/0.32769.0      /16 -> {indirect(1048574)}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 203.0.113.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:31:53    Metric2: 2
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.103 Metric: 2
                        Push 800012
                        Indirect next hop: 87272e4 1048574
                        Indirect path forwarding next hops: 1
                                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1
                        203.0.113.103/32 Originating RIB: inet.3
                          Metric: 2                       Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 203.0.113.216 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

2001:db8::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
```

```
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
                Task: MLD
                Announcement bits (1): 0-KRT
```

```
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:34:07
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

203.0.113.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 203.0.113.103:1
                Next-hop reference count: 7
                Source: 203.0.113.103
                Protocol next hop: 203.0.113.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    64496 Peer AS:    64496
                Age: 1:28:12    Metric2: 1
                Task: BGP_69.203.0.113.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 203.0.113.103
                Primary Routing Table bgp.l2vpn.0

203.0.113.152:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
```

```
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

203.0.113.152:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

203.0.113.163:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 203.0.113.163 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    64499
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512
```

```
203.0.113.55/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
  198.51.100.0/24 -> {Push 300112}
        *BGP    Preference: 170/-101
                Next hop type: Router
                Address: 0x925c208
                Next-hop reference count: 2
                Source: 203.0.113.9
                Next hop: 203.0.113.9 via ge-1/2/0.15, selected
                Label operation: Push 300112
                Label TTL action: prop-ttl
                State: <Active Ext>
                Local AS:  64509 Peer AS: 65539
                Age: 1w0d 23:06:56
                AIGP: 25
                Task: BGP_65539.203.0.113.9+56732
                Announcement bits (1): 0-KRT
                AS path: 65539 64508 I
                Accepted
                Route Label: 300112
                Localpref: 100
                Router ID: 213.0.113.99
```

show route extensive (with Validation State)

```
user@host>show route extensive

inet.0: 48 destinations, 51 routes (48 active, 0 holddown, 0 hidden)
1.1.1.1/32 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 599
                Address: 0x7106094
                Next-hop reference count: 7
                Source: 10.1.1.1
                Next hop: 10.1.1.1 via ge-0/0/0.0, selected
                Session Id: 0x140
                State: <Active Ext>
                Local AS:     2 Peer AS:     1
```

```
                Age: 2w5d 4:29:09
                Validation State: [valid | invalid | unknown | unverified ]  +
                Task: RT
```

show route extensive (BGP-SRTE routes)

```
user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
9.9.9.9-1 <c>/64 (1 entry, 0 announced):
        **SPRING-TE Preference: 8
                Next hop type: Indirect, Next hop index: 0
                Address: 0xdc33080
                Next-hop reference count: 1
                Next hop type: Router, Next hop index: 0
                Next hop: 1.2.2.2 via ge-0/0/2.0, selected
                Label element ptr: 0xdf671d0
                Label parent element ptr: 0x0
                Label element references: 11
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x0
                Protocol next hop: 299920
                Label operation: Push 800040
                Label TTL action: prop-ttl
                Load balance label: Label 800040: None;
                Composite next hop: 0xcd4f950 - INH Session ID: 0x0
                Indirect next hop: 0xdc99a84 - INH Session ID: 0x0 Weight 0x1
                State: <Active Int>
                Local AS:   100
                Age: 5d 17:37:19     Metric: 1     Metric2: 16777215
                Validation State: unverified
                Task: SPRING-TE
                AS path:
                SRTE Policy State:
                SR Preference/Override: 200/100
                Tunnel Source: Static configuration
                Composite next hops: 1
                        Protocol next hop: 299920 Metric: 0
                        Label operation: Push 800040
                        Label TTL action: prop-ttl
                        Load balance label: Label 800040: None;
```

```
                        Composite next hop: 0xcd4f950 - INH Session ID: 0x0
                        Indirect next hop: 0xdc99a84 - INH Session ID: 0x0 Weight 0x1
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 1.2.2.2 via ge-0/0/2.0
                                Session Id: 0x0
                                299920 /52 Originating RIB: mpls.0
                                  Metric: 0 Node path count: 1
                                  Forwarding nexthops: 1
                                        Next hop type: Router
                                        Next hop: 1.2.2.2 via ge-0/0/2.0
                                        Session Id: 0x141
```

show route extensive (AIGP)

```
user@host> show route 3.3.3.3/32 extensive
inet.0: 18 destinations, 23 routes (18 active, 0 holddown, 1 hidden)
3.3.3.3/32 (2 entries, 1 announced)
TSI:
KRT in-kernel 3.3.3.3/32 -> {Push 299872}
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 570
                Address: 0x9760194
                Next-hop reference count: 7, key opaque handle: 0x0, non-key opaque handle: 0x0
                Source: 12.12.12.2
                Next hop: 12.12.12.2 via ge-0/0/0.0, selected
                Load balance label: Label 299872: None;
                Session Id: 333
                State: <Active Ext>
                Peer AS:   200
                Age: 3d 4:16:17
                Validation State: unverified
                AIGP: 67
                Effective AIGP: 67 (AIGP plus IGP metric)
                Task: BGP_200_100.12.12.12.2
                Announcement bits (3): 0-KRT 3-Resolve tree 2 4-Resolve_IGP_FRR task
                AS path: 200 300 I
                Accepted
                Localpref: 100
                Router ID: 128.51.141.40
                Session-IDs associated:
```

```
                  Session-id: 332 Version: 3
                  Thread: junos-main
```

Release Information

Command introduced before Junos OS Release 7.4.

`DeletePending` flag added to the command output in Junos OS Release 19.4R1.

show route extensive <route> expanded-nh

IN THIS SECTION

- Syntax | 23643
- Description | 23643
- Additional Information | 23644
- Required Privilege Level | 23644
- Output Fields | 23644
- Sample Output | 23644
- Release Information | 23648

Syntax

```
show route <IPv4/IPv6 route> extensive expanded-nh
```

Description

Displays the slice and/or the firewall filter attached to a specific route through routing-options forwarding-table export policy.

Additional Information

Issue `show route extensive expanded-nh` to display the slice and/or the firewall filter attached to all the routes through routing-options forwarding-table export policy.

Required Privilege Level

view

Output Fields

Table 1088 on page 23644 lists the output fields for the `show route <route> extensive expanded-nh` command.

Table 1088: show route <route> extensive expanded-nh Output Fields

Field Name	Field Description
slice-id	ID assigned to the slice
gencfg-id	ID of the slice-gencfg created within RPD
Filter	The firewall filter attached to the route through export routing policy action.

Sample Output

show route <IPv4/IPv6 route> extensive expanded-nh

```
user@host> show route 10.1.1.1 extensive expanded-nh
```

Sample output when slice action is applied

```
inet.0: 41 destinations, 41 routes (41 active, 0 holddown, 0 hidden)
          10.1.1.1/32 (1 entry, 1 announced)
          Installed-nexthop:
          Chain (0x8a1f238) Index:599 Slice: Gencfg-id 1, Slice-id 1
          Indr Composite (0x8a1e884) 1.1.1.1 Push 10002 Session-ID: 383
            Krt_cnh (0x6faa554) Index:598 Push 10002
              Krt_inh (0x716539c) Index:1048578 PNH: 1.1.1.1
```

```
      Router (0x8a1efb0) Index:580 11.1.1.2 Session-ID: 322 via ge-0/0/0.0
TSI:
KRT in-kernel 10.1.1.1/32 -> {Chain 0x8a1eed8 600}
      *FRR    Preference: 105
                Next hop type: Indirect, Next hop index: 0
                Address: 0x8a1e884
                Next-hop reference count: 4, key opaque handle: 0x0
                Next hop type: Router, Next hop index: 580
                Next hop: 11.1.1.2 via ge-0/0/0.0, selected
                Session Id: 142
                Protocol next hop: 1.1.1.1
                Label operation: Push 10002
                Label TTL action: prop-ttl
                Load balance label: Label 10002: None;
                Composite next hop: 0x6faa554 598 INH Session ID: 383
                Indirect next hop: 0x716539c 1048578 INH Session ID: 383
                State:
                Age: 6     Metric2: 0
                Validation State: unverified
                Task: MCNH_STUB_TASK
                Announcement bits (2): 0-KRT 2-Resolve tree 1
                AS path: I
                Route-nexthop:
                Indr (0x8a1e884) 1.1.1.1 Push 10002
                  Krt_cnh (0x6faa554) Index:598 Push 10002
                   Krt_inh (0x716539c) Index:1048578
                     Router (0x8a1efb0) Index:580 11.1.1.2
                 Thread: junos-main
                 Composite next hops: 1
                      Protocol next hop: 1.1.1.1
                      Label operation: Push 10002
                      Label TTL action: prop-ttl
                      Load balance label: Label 10002: None;
                      Composite next hop: 0x6faa554 598 INH Session ID: 383
                      Indirect next hop: 0x716539c 1048578 INH Session ID: 383
                      Indirect path forwarding next hops: 1
                       Next hop type: Router
                       Next hop: 11.1.1.2 via ge-0/0/0.0
                       Session Id: 142
                       1.1.1.1/32 Originating RIB: inet.0
                         Node path count: 1
                         Forwarding nexthops: 1
                               Next hop type: Router
```

```
                                        Next hop: 11.1.1.2 via ge-0/0/0.0
                                        Session Id: 142
```

Sample output when filter action is applied

```
inet.0: 41 destinations, 41 routes (41 active, 0 holddown, 0 hidden)
          10.1.1.1/32 (1 entry, 1 announced)
          Installed-nexthop:
          Chain (0x8a1eed8) Index:600 Filter: FWF1
            Indr Composite (0x8a1e884) 1.1.1.1 Push 10002 Session-ID: 383
              Krt_cnh (0x6faa554) Index:598 Push 10002
                Krt_inh (0x716539c) Index:1048578 PNH: 1.1.1.1
                  Router (0x8a1efb0) Index:580 11.1.1.2 Session-ID: 322 via ge-0/0/0.0
          TSI:
          KRT in-kernel 10.1.1.1/32 -> {Chain 0x8a1eed8 600}
                  *FRR    Preference: 105
                          Next hop type: Indirect, Next hop index: 0
                          Address: 0x8a1e884
                          Next-hop reference count: 4, key opaque handle: 0x0
                          Next hop type: Router, Next hop index: 580
                          Next hop: 11.1.1.2 via ge-0/0/0.0, selected
                          Session Id: 142
                          Protocol next hop: 1.1.1.1
                          Label operation: Push 10002
                          Label TTL action: prop-ttl
                          Load balance label: Label 10002: None;
                          Composite next hop: 0x6faa554 598 INH Session ID: 383
                          Indirect next hop: 0x716539c 1048578 INH Session ID: 383
                          State:
                          Age: 6     Metric2: 0
                          Validation State: unverified
                          Task: MCNH_STUB_TASK
                          Announcement bits (2): 0-KRT 2-Resolve tree 1
                          AS path: I
                          Route-nexthop:
                          Indr (0x8a1e884) 1.1.1.1 Push 10002
                            Krt_cnh (0x6faa554) Index:598 Push 10002
                              Krt_inh (0x716539c) Index:1048578
                                Router (0x8a1efb0) Index:580 11.1.1.2
                          Thread: junos-main
                          Composite next hops: 1
```

```
                                Protocol next hop: 1.1.1.1
                                Label operation: Push 10002
                                Label TTL action: prop-ttl
                                Load balance label: Label 10002: None;
                                Composite next hop: 0x6faa554 598 INH Session ID: 383
                                Indirect next hop: 0x716539c 1048578 INH Session ID: 383
                                Indirect path forwarding next hops: 1
                                    Next hop type: Router
                                    Next hop: 11.1.1.2 via ge-0/0/0.0
                                    Session Id: 142
                                    1.1.1.1/32 Originating RIB: inet.0
                                      Node path count: 1
                                      Forwarding nexthops: 1
                                        Next hop type: Router
                                        Next hop: 11.1.1.2 via ge-0/0/0.0
                                        Session Id: 142
```

Sample output when slice action and filter action is applied

```
inet.0: 41 destinations, 41 routes (41 active, 0 holddown, 0 hidden)
          10.1.1.1/32 (1 entry, 1 announced)
          Installed-nexthop:
          Chain (0x8a1f1cc) Index:601 Slice: Gencfg-id 1, Slice-id 1 Filter: FWF1
            Indr Composite (0x8a1e884) 1.1.1.1 Push 10002 Session-ID: 383
              Krt_cnh (0x6faa554) Index:598 Push 10002
                Krt_inh (0x716539c) Index:1048578 PNH: 1.1.1.1
                  Router (0x8a1efb0) Index:580 11.1.1.2 Session-ID: 322 via ge-0/0/0.0
          TSI:
          KRT in-kernel 10.1.1.1/32 -> {Chain 0x8a1f1cc 601}
                *FRR    Preference: 105
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0x8a1e884
                    Next-hop reference count: 4, key opaque handle: 0x0
                    Next hop type: Router, Next hop index: 580
                    Next hop: 11.1.1.2 via ge-0/0/0.0, selected
                    Session Id: 142
                    Protocol next hop: 1.1.1.1
                    Label operation: Push 10002
                    Label TTL action: prop-ttl
                    Load balance label: Label 10002: None;
                    Composite next hop: 0x6faa554 598 INH Session ID: 383
                    Indirect next hop: 0x716539c 1048578 INH Session ID: 383
```

```
                    State:
                    Age: 13     Metric2: 0
                    Validation State: unverified
                    Task: MCNH_STUB_TASK
                    Announcement bits (2): 0-KRT 2-Resolve tree 1
                    AS path: I
                    Route-nexthop:
                    Indr (0x8a1e884) 1.1.1.1 Push 10002
                      Krt_cnh (0x6faa554) Index:598 Push 10002
                        Krt_inh (0x716539c) Index:1048578
                          Router (0x8a1efb0) Index:580 11.1.1.2
                    Thread: junos-main
                    Composite next hops: 1
                        Protocol next hop: 1.1.1.1
                        Label operation: Push 10002
                        Label TTL action: prop-ttl
                        Load balance label: Label 10002: None;
                        Composite next hop: 0x6faa554 598 INH Session ID: 383
                        Indirect next hop: 0x716539c 1048578 INH Session ID: 383
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 11.1.1.2 via ge-0/0/0.0
                                Session Id: 142
                                1.1.1.1/32 Originating RIB: inet.0
                                  Node path count: 1
                                  Forwarding nexthops: 1
                                        Next hop type: Router
                                        Next hop: 11.1.1.2 via ge-0/0/0.0
                                        Session Id: 142
```

Release Information

Command introduced in JUNOS 22.2R1

show route flow validation

IN THIS SECTION

- Syntax | 23649
- Syntax (EX Series Switches) | 23649
- Description | 23649
- Options | 23650
- Required Privilege Level | 23650
- Output Fields | 23650
- Sample Output | 23651
- Release Information | 23652

Syntax

```
show route flow validation
<brief | detail>
<ip-prefix>
<table table-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route flow validation
<brief | detail>
<ip-prefix>
<table table-name>
```

Description

Display flow route information.

Options

none	Display flow route information.
brief | detail	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.
ip-prefix	(Optional) IP address for the flow route.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
table *table-name*	(Optional) Display flow route information for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the `show route flow validation inet` command).

Required Privilege Level

view

Output Fields

Table 1089 on page 23650 lists the output fields for the `show route flow validation` command. Output fields are listed in the approximate order in which they appear.

Table 1089: show route flow validation Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table (for example, inet.0).	All levels
prefix	Route address.	All levels
Active unicast route	Active route in the routing table.	All levels
Dependent flow destinations	Number of flows for which there are routes in the routing table.	All levels

Table 1089: show route flow validation Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Origin	Source of the route flow.	All levels
Neighbor AS	Autonomous system identifier of the neighbor.	All levels
Flow destination	Number of entries and number of destinations that match the route flow.	All levels
Unicast best match	Destination that is the best match for the route flow.	All levels
Flags	Information about the route flow.	All levels

Sample Output

show route flow validation

```
user@host> show route flow validation
inet.0:
10.0.5.0/24Active unicast route
Dependent flow destinations: 1
Origin: 192.168.224.218, Neighbor AS: 64501
Flow destination (3 entries, 1 match origin)
Unicast best match: 10.0.5.0/24
Flags: SubtreeApex Consistent
```

show route flow validation (IPv6)

```
user@host> show route flow validation
inet6.0:
2001:db8::11:11:11:0/120
                Active unicast route
                    Dependent flow destinations: 2
                    Origin: 2001:db8::13:14:2:2, Neighbor AS: 2000
```

```
2001:db8::11:11:11:10/128
                Flow destination (1 entries, 1 match origin, next-as)
                    Unicast best match: 2001:db8::11:11:11:0/120
                    Flags: Consistent
2001:db8::11:11:11:30/128
                Flow destination (1 entries, 1 match origin, next-as)
                    Unicast best match: 2001:db8::11:11:11:0/120
                    Flags: Consistent
```

Release Information

Command introduced before Junos OS Release 7.4.

show route forwarding-table

IN THIS SECTION

- Syntax | 23652
- Syntax (MX Series Routers) | 23653
- Syntax (TX Matrix and TX Matrix Plus Routers) | 23653
- Description | 23654
- Options | 23654
- Required Privilege Level | 23656
- Output Fields | 23656
- Sample Output | 23663
- Release Information | 23672

Syntax

```
show route forwarding-table
<detail | extensive | summary>
<all>
<ccc interface-name>
```

```
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<label name>
<matching matching>
<multicast>
<table (default | logical-system-name/routing-instance-name | routing-instance-name)>
<vlan (all | vlan-name)>
<vpn vpn>
```

Syntax (MX Series Routers)

```
show route forwarding-table
<detail | extensive | summary>
<all>
<bridge-domain (all | domain-name)>
<ccc interface-name>
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<label name>
<learning-vlan-id learning-vlan-id>
<matching matching>
<multicast>
<table (default | logical-system-name/routing-instance-name | routing-instance-name)>
<vlan (all |  vlan-name)>
<vpn vpn>
```

Syntax (TX Matrix and TX Matrix Plus Routers)

```
show route forwarding-table
<detail | extensive | summary>
<all>
<ccc interface-name>
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<matching matching>
<label name>
```

```
<lcc number>
<multicast>
<table routing-instance-name>
<vpn vpn>
```

Description

Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.

> **NOTE**: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the `show pfe route` command.

Options

none	Display the routes in the forwarding tables. By default, the `show route forwarding-table` command does not display information about private, or internal, forwarding tables.
detail | extensive | summary	(Optional) Display the specified level of output.
all	(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.
bridge-domain (all | *bridge-domain-name*)	(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.
ccc *interface-name*	(Optional) Display route entries for the specified circuit cross-connect interface.
destination *destination-prefix*	(Optional) Destination prefix.
family *family*	(Optional) Display routing table entries for the specified family: `bridge` (`ccc` | `destination` | `detail` | `extensive` | `interface-name` | `label` | `learning-vlan-id` | `matching` | `multicast` | `summary` | `table` | `vlan` | `vpn`), `ethernet-switching`, `evpn`, `fibre-channel`, `fmembers`, `inet`, `inet6`, `iso`, `mcsnoop-inet`, `mcsnoop-inet6`, `mpls`, `satellite-inet`, `satellite-inet6`, `satellite-vpls`, `tnp`, `unix`, `vpls`, or `vlan-classification`.

interface-name *interface-name*
(Optional) Display routing table entries for the specified interface.

label *name*
(Optional) Display route entries for the specified label.

lcc *number*
(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

learning-vlan-id *learning-vlan-id*
(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching *matching*
(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast
(Optional) Display routing table entries for multicast routes.

table
(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the `show route instance` command.

vlan (all | *vlan-name***)**
(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*
(Optional) Display routing table entries for a specified VPN.

Required Privilege Level

view

Output Fields

Table 1090 on page 23656 lists the output fields for the `show route forwarding-table` command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the `detail` keyword is used instead of the `extensive` keyword.

Table 1090: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the `table` *`logical-system-name/routing-instance-name`* option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels

Table 1090: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Enabled protocols	The features and protocols that have been enabled for a given routing table. This field can contain the following values: • BUM hashing—BUM hashing is enabled. • MAC Stats—Mac Statistics is enabled. • Bridging—Routing instance is a normal layer 2 bridge. • No VLAN—No VLANs are associated with the bridge domain. • All VLANs—The `vlan-id all` statement has been enabled for this bridge domain. • Single VLAN—Single VLAN ID is associated with the bridge domain. • MAC action drop—New MACs will be dropped when the MAC address limit is reached. • Dual VLAN—Dual VLAN tags are associated with the bridge domain • No local switching—No local switching is enabled for this routing instance.. • Learning disabled—Layer 2 learning is disabled for this routing instance. • MAC limit reached—The maximum number of MAC addresses that was configured for this routing instance has been reached. • VPLS—The VPLS protocol is enabled. • No IRB l2-copy—The no-irb-layer-2-copy feature is enabled for this routing instance. • ACKed by all peers—All peers have acknowledged this routing instance. • BUM Pruning—BUM pruning is enabled on the VPLS instance. • Def BD VXLAN—VXLAN is enabled for the default bridge domain. • EVPN—EVPN protocol is enabled for this routing instance.	All levels

Table 1090: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• Def BD OVSDB—Open vSwitch Database (OVSDB) is enabled on the default bridge domain. • Def BD Ingress replication—VXLAN ingress node replication is enabled on the default bridge domain. • L2 backhaul—Layer 2 backhaul is enabled. • FRR optimize—Fast reroute optimization • MAC pinning—MAC pinning is enabled for this bridge domain. • MAC Aging Timer—The MAC table aging time is set per routing instance. • EVPN VXLAN—This routing instance supports EVPN with VXLAN encapsulation. • PBBN—This routing instance is configured as a provider backbone bridged network. • PBN—This routing instance is configured as a provider bridge network. • ETREE—The ETREE protocol is enabled on this EVPN routing instance. • ARP/NDP suppression—EVPN ARP NDP suppression is enabled in this routing instance. • Def BD EVPN VXLAN—EVPN VXLAN is enabled for the default bridge domain. • MPLS control word—Control word is enabled for this MPLS routing instance.	
Address family	Address family (for example, `IP`, `IPv6`, `ISO`, `MPLS`, and `VPLS`).	All levels
Destination	Destination of the route.	`detail extensive`

Table 1090: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Route Type (Type)	How the route was placed into the forwarding table. When the `detail` keyword is used, the route type might be abbreviated (as shown in parentheses): • `cloned` (`clon`)—(TCP or multicast only) Cloned route. • `destination` (`dest`)—Remote addresses directly reachable through an interface. • `destination down` (`iddn`)—Destination route for which the interface is unreachable. • `interface cloned` (`ifcl`)—Cloned route for which the interface is unreachable. • `route down` (`ifdn`)—Interface route for which the interface is unreachable. • `ignore` (`ignr`)—Ignore this route. • `interface` (`intf`)—Installed as a result of configuring an interface. • `permanent` (`perm`)—Routes installed by the kernel when the routing table is initialized. • `user`—Routes installed by the routing protocol process or as a result of the configuration.	All levels
Route Reference (RtRef)	Number of routes to reference.	`detail extensive`

Table 1090: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Route type flags: • `none`—No flags are enabled. • `accounting`—Route has accounting enabled. • `cached`—Cache route. • `incoming-iface` *`interface-number`*—Check against incoming interface. • `prefix load balance`—Load balancing is enabled for this prefix. • `rt nh decoupled`—Route has been decoupled from the next hop to the destination. • `sent to PFE`—Route has been sent to the Packet Forwarding Engine. • `static`—Static route.	`extensive`
Next hop	IP address of the next hop to the destination. **NOTE**: For static routes that use point-to-point (P2P) outgoing interfaces, the next-hop address is not displayed in the output.	`detail extensive`

Table 1090: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Next hop Type (Type)	Next-hop type. When the `detail` keyword is used, the next-hop type might be abbreviated (as indicated in parentheses): • `broadcast (bcst)`—Broadcast. • `deny`—Deny. • `discard (dscd)` —Discard. • `hold`—Next hop is waiting to be resolved into a unicast or multicast type. • `indexed (idxd)`—Indexed next hop. • `indirect (indr)`—Indirect next hop. • `local (locl)`—Local address on an interface. • `routed multicast (mcrt)`—Regular multicast next hop. • `multicast (mcst)`—Wire multicast next hop (limited to the LAN). • `multicast discard (mdsc)`—Multicast discard. • `multicast group (mgrp)`—Multicast group member. • `receive (recv)`—Receive. • `reject (rjct)`—Discard. An ICMP unreachable message was sent. • `resolve (rslv)`—Resolving the next hop. • `unicast (ucst)`—Unicast. • `unilist (ulst)`—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. • `VxLAN Local`—EVPN Type 5 route in kernel.	`detail extensive`
Index	Software index of the next hop that is used to route the traffic for a given prefix.	`detail extensive` none

Table 1090: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	`extensive`
Reference (NhRef)	Number of routes that refer to this next hop.	`detail extensive` none
Next-hop interface (Netif)	Interface used to reach the next hop.	`detail extensive` none
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the `Balance` field description).	`extensive`
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	`extensive`
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when `rpf-check` is configured on the interface.	`extensive`
sftw	On Junos OS Evolved devices, `sftw` is another type of next-hop. A software next-hop can contain a link to another single next-hop, and hence it often works as an indirect next-hop. The usage of `sftw` helps to reduce churn in Junos OS Evolved objects by providing a common object where topology updates can be reflected. Further MPLS states can also be attached to these software next-hops. When a route is created, a `sftw` next-hop seems to be created whether it is the management port or the loopback interface.	`detail extensive` none

Table 1090: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Alternate forward nh index	Index number of the alternate next hop interface. Seen with **multicast** option only.	**extensive**
Next-hop L3 Interface	The next hop layer 3 interface. This option can be expressed as a VLAN name and is only seen with the **multicast** option.	**extensive**
Next-hop L2 Interfaces	The next hop layer 2 interfaces. Seen with **multicast** option only.	**extensive**

Sample Output

show route forwarding-table

```
user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     4
0.0.0.0/32         perm     0                    dscd    44     1
172.16.1.0/24         ifdn     0                    rslv   608     1 ge-2/0/1.0
172.16.1.0/32         iddn     0 172.16.1.0            recv   606     1 ge-2/0/1.0
172.16.1.1/32         user     0                    rjct    46     4
172.16.1.1/32         intf     0 172.16.1.1            locl   607     2
172.16.1.1/32         iddn     0 172.16.1.1            locl   607     2
172.16.1.255/32       iddn     0 ff:ff:ff:ff:ff:ff  bcst   605     1 ge-2/0/1.0
10.0.0.0/24        intf     0                    rslv   616     1 ge-2/0/0.0
10.0.0.0/32        dest     0 10.0.0.0           recv   614     1 ge-2/0/0.0
10.0.0.1/32        intf     0 10.0.0.1           locl   615     2
10.0.0.1/32        dest     0 10.0.0.1           locl   615     2
10.0.0.255/32      dest     0 10.0.0.255         bcst   613     1 ge-2/0/0.0
10.1.1.0/24        ifdn     0                    rslv   612     1 ge-2/0/1.0
10.1.1.0/32        iddn     0 10.1.1.0           recv   610     1 ge-2/0/1.0
10.1.1.1/32        user     0                    rjct    46     4
10.1.1.1/32        intf     0 10.1.1.1           locl   611     2
10.1.1.1/32        iddn     0 10.1.1.1           locl   611     2
```

```
10.1.1.255/32      iddn     0 ff:ff:ff:ff:ff:ff  bcst   609     1 ge-2/0/1.0
10.206.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0
10.209.0.0/16      user     1 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.0.0/18      intf     0                    rslv   418     1 fxp0.0
10.209.0.0/32      dest     0 10.209.0.0         recv   416     1 fxp0.0
10.209.2.131/32    intf     0 10.209.2.131       locl   417     2
10.209.2.131/32    dest     0 10.209.2.131       locl   417     2
10.209.17.55/32    dest     0 0:30:48:5b:78:d2   ucst   435     1 fxp0.0
10.209.63.42/32    dest     0 0:23:7d:58:92:ca   ucst   434     1 fxp0.0
10.209.63.254/32   dest     0 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.63.255/32   dest     0 10.209.63.255      bcst   415     1 fxp0.0
10.227.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0
...

Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct    27     1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf     0                   locl    28     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct     6     1
ff00::/8           perm     0                   mdsc     4     1
ff02::1/128        perm     0 ff02::1           mcst     3     1

Routing table: ccc
MPLS:
Interface.Label    Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 16     1
100004(top)fe-0/0/1.0
```

show route forwarding-table detail

```
user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     2 0:90:69:8e:b1:1b   ucst   132     4 fxp0.0
```

```
default            perm     0                    rjct    14     1
10.1.1.0/24        intf     0 ff.3.0.21          ucst   322     1 so-5/3/0.0
10.1.1.0/32        dest     0 10.1.1.0           recv   324     1 so-5/3/0.0
10.1.1.1/32        intf     0 10.1.1.1           locl   321     1
10.1.1.255/32      dest     0 10.1.1.255         bcst   323     1 so-5/3/0.0
10.21.21.0/24      intf     0 ff.3.0.21          ucst   326     1 so-5/3/0.0
10.21.21.0/32      dest     0 10.21.21.0         recv   328     1 so-5/3/0.0
10.21.21.1/32      intf     0 10.21.21.1         locl   325     1
10.21.21.255/32    dest     0 10.21.21.255       bcst   327     1 so-5/3/0.0
127.0.0.1/32       intf     0 127.0.0.1          locl   320     1
172.17.28.19/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0
172.17.28.44/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0
...

Routing table: private1__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     1
10.0.0.0/8         intf     0                    rslv   136     1 fxp1.0
10.0.0.0/32        dest     0 10.0.0.0           recv   134     1 fxp1.0
10.0.0.4/32        intf     0 10.0.0.4           locl   135     2
10.0.0.4/32        dest     0 10.0.0.4           locl   135     2

...

Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    38     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    22     1
ff00::/8           perm     0                    mdsc    21     1
ff02::1/128        perm     0 ff02::1            mcst    17     1

...

Routing table: mpls
MPLS:
```

```
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 28     1
```

show route forwarding-table summary

```
user@switch> show route forwarding-table summary

Routing table: default.inet
Internet:
         user:          6 routes
         perm:          5 routes
         intf:          8 routes
         dest:         12 routes
         ifdn:          1 routes
         iddn:          3 routes
```

show route forwarding-table extensive

```
user@switch> show route forwarding-table extensive

Routing table: default.inet [Index 0]
Internet:

Destination:  default
  Route type: user
  Route reference: 2                   Route interface-index: 0
  Flags: sent to PFE, rt nh decoupled
  Nexthop: 0:12:f2:21:cf:0
  Next-hop type: unicast               Index: 333      Reference: 5
  Next-hop interface: me0.0

Destination:  default
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: none
  Next-hop type: reject                Index: 36       Reference: 2

Destination:  0.0.0.0/32
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
```

```
  Flags: sent to PFE
  Next-hop type: discard               Index: 34       Reference: 1

Destination:  2.2.2.0/24
  Route type: interface
  Route reference: 0                   Route interface-index: 66
  Flags: sent to PFE
  Next-hop type: resolve               Index: 1309     Reference: 1
  Next-hop interface: ae0.0

Destination:  2.2.2.0/32
  Route type: destination
  Route reference: 0                   Route interface-index: 66
  Flags: sent to PFE
  Nexthop: 2.2.2.0
  Next-hop type: receive               Index: 1307     Reference: 1
  Next-hop interface: ae0.0

Destination:  2.2.2.1/32
  Route type: destination
  Route reference: 0                   Route interface-index: 66
  Flags: sent to PFE
  Nexthop: 0:21:59:cc:89:c0
  Next-hop type: unicast               Index: 1320     Reference: 1
  Next-hop interface: ae0.0

Destination:  2.2.2.2/32
  Route type: interface
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 2.2.2.2
  Next-hop type: local                 Index: 1308     Reference: 2

Destination:  2.2.2.2/32
  Route type: destination
  Route reference: 0                   Route interface-index: 66
  Flags: none
  Nexthop: 2.2.2.2
  Next-hop type: local                 Index: 1308     Reference: 2

Destination:  2.2.2.255/32
  Route type: destination
  Route reference: 0                   Route interface-index: 66
```

```
  Flags: sent to PFE
  Nexthop: 2.2.2.255
  Next-hop type: broadcast             Index: 1306     Reference: 1
  Next-hop interface: ae0.0
```

show route forwarding-table ccc

```
user@switch> show route forwarding-table ccc ge-0/0/0.10
Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
ge-0/0/0.10  (CCC) user     0 3.3.3.2           Push 300112  1343     2 ae1.0
```

show route forwarding-table family (MPLS)

```
user@switch> show route forwarding-table family mpls

Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd    50     1
0                  user     0                    recv    49     3
1                  user     0                    recv    49     3
2                  user     0                    recv    49     3
299776             user     0                    Pop    1334     2 ge-0/0/0.10
299792             user     0                    Pop    1339     2 ge-0/0/0.14
299808             user     0                    Pop    1341     2 ge-0/0/0.2
299824             user     0                    Pop    1344     2 ge-0/0/0.11
299840             user     0                    Pop    1345     2 ge-0/0/0.13
299856             user     0                    Pop    1346     2 ge-0/0/0.18
299872             user     0                    Pop    1347     2 ge-0/0/0.16
299888             user     0                    Pop    1348     2 ge-0/0/0.7
299904             user     0                    Pop    1349     2 ge-0/0/0.20
299920             user     0                    Pop    1350     2 ge-0/0/0.19
299936             user     0                    Pop    1351     2 ge-0/0/0.17
299952             user     0                    Pop    1352     2 ge-0/0/0.9
299968             user     0                    Pop    1353     2 ge-0/0/0.1
299984             user     0                    Pop    1354     2 ge-0/0/0.12
300000             user     0                    Pop    1355     2 ge-0/0/0.8
300016             user     0                    Pop    1356     2 ge-0/0/0.4
```

```
300032             user     0                    Pop    1357     2 ge-0/0/0.5
300048             user     0                    Pop    1358     2 ge-0/0/0.3
300064             user     0                    Pop    1359     2 ge-0/0/0.15
ge-0/0/0.1  (CCC)  user     0 3.3.3.2            Push 300064  1340     2 ae1.0
ge-0/0/0.2  (CCC)  user     0 3.3.3.2            Push 299872  1328     2 ae1.0
ge-0/0/0.3  (CCC)  user     0 3.3.3.2            Push 299792  1323     2 ae1.0
ge-0/0/0.4  (CCC)  user     0 3.3.3.2            Push 300016  1337     2 ae1.0
ge-0/0/0.5  (CCC)  user     0 3.3.3.2            Push 299824  1325     2 ae1.0
ge-0/0/0.7  (CCC)  user     0 3.3.3.2            Push 299920  1331     2 ae1.0
ge-0/0/0.8  (CCC)  user     0 3.3.3.2            Push 299840  1326     2 ae1.0
ge-0/0/0.9  (CCC)  user     0 3.3.3.2            Push 299888  1329     2 ae1.0
ge-0/0/0.10  (CCC) user     0 3.3.3.2            Push 300112  1343     2 ae1.0
ge-0/0/0.11  (CCC) user     0 3.3.3.2            Push 299776  1322     2 ae1.0
ge-0/0/0.12  (CCC) user     0 3.3.3.2            Push 299952  1333     2 ae1.0
ge-0/0/0.13  (CCC) user     0 3.3.3.2            Push 300096  1342     2 ae1.0
ge-0/0/0.14  (CCC) user     0 3.3.3.2            Push 299984  1335     2 ae1.0
ge-0/0/0.15  (CCC) user     0 3.3.3.2            Push 299936  1332     2 ae1.0
ge-0/0/0.16  (CCC) user     0 3.3.3.2            Push 299808  1324     2 ae1.0
ge-0/0/0.17  (CCC) user     0 3.3.3.2            Push 300000  1336     2 ae1.0
ge-0/0/0.18  (CCC) user     0 3.3.3.2            Push 300032  1338     2 ae1.0
ge-0/0/0.19  (CCC) user     0 3.3.3.2            Push 299904  1330     2 ae1.0
ge-0/0/0.20  (CCC) user     0 3.3.3.2            Push 299856  1327     2 ae1.0
```

show route forwarding-table family (IPv6)

```
user@switch> show route forwarding-table family inet6
Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    44     1
::/128             perm     0                    dscd    42     1
ff00::/8           perm     0                    mdsc    43     1
ff02::1/128        perm     0 ff02::1            mcst    39     1

Routing table: default-switch.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   530     1
::/128             perm     0                    dscd   528     1
2:1::3a00/312      user     0                    indr 131070     2
                                                 comp   572     1
```

```
2:1::3a82/320      user     0                    indr 131071     3
                                                 comp   573     1
2:1::3af0/320      user     0                    indr 131071     3
                                                 comp   573     1
2:1:0:ff00::/56    user     0                    mdsc   529     2
ff00::/8           perm     0                    mdsc   529     2
ff02::1/128        perm     0 ff02::1            mcst   526     1

Routing table: __master.anon__.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   554     1
::/128             perm     0                    dscd   552     1
ff00::/8           perm     0                    mdsc   553     1
ff02::1/128        perm     0 ff02::1            mcst   550     1
```

show route forwarding-table label

```
user@switch> show route forwarding-table label 29976

Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
299776             user     0                   Pop    1334     2 ge-0/0/0.10
```

show route forwarding-table matching

```
user@switch> show route forwarding-table matching 3

Routing table: default.inet
Internet:
```

show route forwarding-table multicast

```
user@switch> show route forwarding-table multicast

Routing table: default.inet
Internet:
```

```
Destination        Type RtRef Next hop           Type Index NhRef Netif
224.0.0.0/4        perm     1                    mdsc    35     1
224.0.0.1/32       perm     0 224.0.0.1          mcst    31     3
224.0.0.5/32       user     1 224.0.0.5          mcst    31     3

Routing table: __master.anon__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
224.0.0.0/4        perm     0                    mdsc  1289     1
224.0.0.1/32       perm     0 224.0.0.1          mcst  1285     1

Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
ff00::/8           perm     0                    mdsc    43     1
ff02::1/128        perm     0 ff02::1            mcst    39     1
```

show route forwarding-table destination extensive (EVPN Type 5 route with Type 2 and Type 5 route coexistence)

```
user@device> show route forwarding-table destination 10.1.1.20 table vrf1 extensive
Routing table: vrf1.inet [Index 9]
Internet:
Destination:  10.1.1.20/32
  Route type: user
  Route reference: 0                   Route interface-index: 0
  Multicast RPF nh index: 0
  P2mpidx: 0
  Flags: sent to PFE, VxLAN Local
  Nexthop:
  Next-hop type: composite             Index: 2694     Reference: 7
  Next-hop type: indirect              Index: 524326   Reference: 2
  Next-hop type: unilist               Index: 524288   Reference: 5
  Nexthop: 10.1.1.1
  Next-hop type: unicast               Index: 1724     Reference: 15
  Next-hop interface: xe-0/0/1.0    Weight: 0x0
  Nexthop: 10.1.1.4  Next-hop type: unicast               Index: 1725     Reference: 15
  Next-hop interface: xe-0/0/4.0    Weight: 0x0
```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```
so-1/1/0 {
    unit 0 {
        family inet {
            rpf-check;
            address 192.0.2.2/30;
        }
    }
}
```

show route forwarding-table (Junos OS Evolved devices with next-hop type 'sftw')

```
user@host> show route forwarding-table | grep sftw
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     0                    sftw     7021     1 re0:mgmt-0.0
10.21.1.0/32       user     0                    sftw     7022     1 et-0/0/32.0

user@host> show route forwarding-table | grep sftw
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     0                    sftw     5021     1 re0:mgmt-0.0
10.22.1.0/32       user     0                    sftw     5022     1 et-0/0/32.0
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `bridge-domain` introduced in Junos OS Release 7.5

Option `learning-vlan-id` introduced in Junos OS Release 8.4

Options `all` and `vlan` introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

show route instance | 23692

show route forwarding-table family fibre-channel

IN THIS SECTION

- Syntax | 23673
- Description | 23673
- Options | 23674
- Required Privilege Level | 23674
- Output Fields | 23674
- Sample Output | 23675
- Release Information | 23676

Syntax

```
show route forwarding-table family fibre-channel
<brief | detail | extensive>
<all>
<destination destination-prefix>
<interface-name interface-name>
<label label>
<matching ip-prefix>
<multicast>
<summary>
<table routing-table-name>
<vlan vlan-name>
<vpn vpn-instance-name>
```

Description

Display Fibre Channel family forwarding table route information.

Options

brief | detail | extensive	(Optional) Display the specified level of output.
all	Display all routing forwarding tables.
destination *destination-prefix*	Destination prefix.
interface-name *interface-name*	Name of the interface.
label *label*	Display route entries for the specified label name.
matching *ip-prefix*	Display route entries for the specified IP prefix or length.
multicast	Display multicast routes.
summary	Display route count instead of details.
table *routing-table-name*	Name of the routing table.
vlan *vlan-name*	Name of the VLAN.
vpn *vpn-instance-name*	Name of the VPN instance.

Required Privilege Level

view

Output Fields

Table 1091 on page 23674 lists the output fields for the `show route forwarding-table family fibre-channel` command. Output fields are listed in the approximate order in which they appear.

Table 1091: show route forwarding-table family fibre-channel Output Fields

Field Name	Field Description
`Routing table`	Name of the routing table.
`Destination`	Route destination.

Table 1091: show route forwarding-table family fibre-channel Output Fields *(Continued)*

Field Name	Field Description
Type	Type of route internal to Junos OS.
RtRef	Route reference count internal to Junos OS.
Next hop Type	Type of next hop internal to Junos OS.
Index	Next-hop index identifier.
NhRef	Number of routes that refer to the next hop.
Netif	Interface used to reach the next hop.

Sample Output

show route forwarding-table family fibre-channel

```
user@switch> show route forwarding-table family fibre-channel
```

```
Routing table: default.fibre-channel
Fibre Channel:
Destination      Type RtRef Next hop           Type Index NhRef Netif
default          perm     0                    dscd   126     1
0x30000/24       user     0                    ucst   674     2 fc-0/0/0.0
0x30001/24       user     0                    ucst   675     2 vlan.100
0x30002/24       user     0                    ucst   676     2 vlan.100
0x30003/24       user     0                    ucst   677     2 vlan.100
0x30004/24       user     0                    ucst   678     2 vlan.100
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel next-hops | 19985

show fibre-channel routes | 20006

show route forwarding-table interface-name

IN THIS SECTION

- Syntax | 23676
- Description | 23676
- Options | 23677
- Required Privilege Level | 23677
- Output Fields | 23677
- Sample Output | 23678
- Release Information | 23680

Syntax

```
show route forwarding-table interface-name interface-name
<detail | extensive>
<all>
```

Description

Display the interfaces in the Routing Engine's forwarding table.

Options

none	Display information for the specified interface.
detail | extensive	(Optional) Display the specified level of output.
all	(Optional) Display all interfaces in the routing table.

Required Privilege Level

view

Output Fields

Table 1092 on page 23677 lists the output fields for the `show route forwarding-table interface-name` command. Output fields are listed in the approximate order in which they appear.

Table 1092: show route forwarding-table interface-name Output Fields

Field Name	Field Description	Level of Output
Name	Name of the interface (for example **fe-0/1/1**, **lo0**, **ae0**, and so on).	All levels
MTU	Interface's maximum transmission unit (MTU).	All levels
Afam	Configured address family (for example **inet**, **tnp**, **inet6**, and so on).	**detail extensive**
Network	Network information: • **<Link>**—Physical interface, not a logical interface. • **<PtoP>**—Point-to-point network. • **ipaddress**—Network address.	All levels
Address	Address of the interface. The address can be a MAC address, IPv4 address, IPv6 address, and so on.	All levels

Table 1092: show route forwarding-table interface-name Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IPkts	Number of packets received on the interface.	All levels
Ierr	Number of packets received on the interface with errors.	All levels
Opkts	Number of packets transmitted or sent from the interface.	All levels
Oerr	Number of packets transmitted or sent from the interface with errors.	All levels
Coll	Number of packets that experienced collisions on the interface.	All levels

Sample Output

show route forwarding-table interface-name fe-0/1/1

```
user@host> show route forwarding-table interface-name fe-0/1/1
Name         Mtu Network      Address             Ipkts Ierr    Opkts Oerr  Coll
fe-0/1/1    1514 <Link>     00.05.85.88.cc.20         0    0        0    0     0
```

show route forwarding-table interface-name all

```
user@host> show route forwarding-table interface-name all
Name         Mtu Network      Address             Ipkts Ierr    Opkts Oerr  Coll
fxp0        1514 <Link>     00.a0.a5.56.03.83    180965    0    39907    0     0
 unit 0     1500 192.168.187.0/ 192.168.187.10
fxp1        1514 <Link>     02.00.00.00.00.04  33010497    0 30110800    0     0
 unit 0     1500 10.0.0.0/8     10.0.0.1
              10.0.0.0/8     10.0.0.4
              128.0.0.0/2    128.0.0.1
              128.0.0.0/2    128.0.0.4
            1500 fe80::/64      fe80::200:ff:fe0
              fec0::/64      fec0::a:0:0:4
            1500                4
```

```
lsi         1496 <Link>
dsc          max <Link>                              0    0        0    0     0
lo0          max <Link>                           8980    0     8980    0     0
 unit 0      max 127.0.0.1/8    127.0.0.1
              192.168.0.1/8  192.168.0.1
 unit 16384  max 127.0.0.1/8    127.0.0.1
 unit 16385  max
gre          max <Link>
ipip         max <Link>
tap          max <Link>
pime         max <Link>
pimd         max <Link>
mtun         max <Link>
so-0/0/0    4474 <Link>                        1679900    0  1068611    0     0
 unit 0     4470 <PtoP>         10.0.60.2               0    0        0    0     0
so-0/0/1    4474 <Link>                              0    0        0    0     0
 unit 0     4470 <PtoP>         10.0.80.2               0    0        0    0     0
so-0/0/2    4474 <Link>                              0    0        0    0     0
so-0/0/3    4474 <Link>                              0    0        0    0     0
fe-0/1/0    1514 <Link>     00.05.85.88.cc.1f    523120    0   623044    0     0
 unit 0     1500 10.0.90.12/30  10.0.90.14              0    0        0    0     0
fe-0/1/1    1514 <Link>     00.05.85.88.cc.20         0    0        0    0     0
fe-0/1/2    1514 <Link>     00.05.85.88.cc.21         0    0        0    0     0

...
```

show route forwarding-table interface-name all detail

```
user@host> show route forwarding-table interface-name all detail
Name         Mtu AFam    Network      Address             Ipkts Ierr    Opkts Oerr  Coll
fxp0        1514         <Link>     00.a0.a5.56.03.83    181005    0    39948    0     0
 unit 0     1500 inet    192.168.187.0/ 192.168.187.10
fxp1        1514         <Link>     02.00.00.00.00.04  33012676    0 30112468    0     0
 unit 0     1500 inet    10.0.0.0/8     10.0.0.1
                      10.0.0.0/8     10.0.0.4
                      128.0.0.0/2    128.0.0.1
                      128.0.0.0/2    128.0.0.4
            1500 inet6   fe80::/64      fe80::200:ff:fe0
                      fec0::/64      fec0::a:0:0:4
            1500 tnp                    4
lsi         1496         <Link>
```

```
dsc          max         <Link>                              0    0        0    0     0
lo0          max         <Link>                           8980    0     8980    0     0
 unit 0      max inet    127.0.0.1/8    127.0.0.1
                       192.168.0.1/8  192.168.0.1
 unit 16384  max inet    127.0.0.1/8    127.0.0.1
 unit 16385  max inet
gre          max         <Link>
ipip         max         <Link>
tap          max         <Link>
pime         max         <Link>
pimd         max         <Link>
mtun         max         <Link>
so-0/0/0    4474         <Link>                        1679980    0  1068661    0     0
 unit 0     4470 inet    <PtoP>         10.0.60.2               0    0        0    0     0

...
```

Release Information

Command introduced in Junos OS Release 9.6.

show route hidden

IN THIS SECTION

- Syntax | 23681
- Description | 23681
- Options | 23681
- Required Privilege Level | 23681
- Output Fields | 23681
- Sample Output | 23681
- Release Information | 23684

Syntax

```
show route hidden
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Description

Display only hidden route information. A hidden route is unusable, even if it is the best path.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to `brief`.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field table for the `show route` command, the `show route detail` command, the `show route extensive` command, or the `show route terse` command.

Sample Output

show route hidden

```
user@host> show route hidden
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
127.0.0.1/32        [Direct/0] 04:26:38
                    > via lo0.0
```

```
private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.5.5.5/32         [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: 100 I
                      Unusable
10.12.1.0/24        [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: 100 I
                      Unusable
10.12.80.4/30       [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: I
                      Unusable
...
```

show route hidden detail

```
user@host> show route hidden detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
127.0.0.1/32 (1 entry, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Hidden Martian Int>
                Local AS:     1
                Age: 4:27:37
                Task: IF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.5.5.5/32 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Route Distinguisher: 10.4.4.4:4
```

```
                Next hop type: Unusable
                Next-hop reference count: 6
                State: <Secondary Hidden Int Ext>
                Local AS:     1 Peer AS:     1
                Age: 3:45:09
                Task: BGP_1.10.4.4.4+2493
                AS path: 100 I
                Communities: target:1:999
                VPN Label: 100064
                Localpref: 100
                Router ID: 10.4.4.4
                Primary Routing Table bgp.l3vpn.0

...
```

show route hidden extensive

The output for the `show route hidden extensive` command is identical to that of the `show route hidden detail` command. For sample output, see "show route hidden detail" on page 23682.

show route hidden terse

```
user@host> show route hidden terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  127.0.0.1/32       D   0                       >lo0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.5.5.5/32        B 170        100             Unusable        100 I
  10.12.1.0/24       B 170        100             Unusable        100 I
  10.12.80.4/30      B 170        100             Unusable        I
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.4.4.4:4:10.5.5.5/32
                     B 170        100             Unusable        100 I
  10.4.4.4:4:10.12.1.0/24
                     B 170        100             Unusable        100 I
  10.4.4.4:4:10.12.80.4/30
                     B 170        100             Unusable        I

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show route | 23512
- show route detail | 23573
- show route extensive | 23619
- show route terse | 23824
- Understanding Hidden Routes

show route inactive-path

IN THIS SECTION

- Syntax | 23685
- Syntax (EX Series Switches) | 23685
- Description | 23685
- Options | 23685
- Required Privilege Level | 23686
- Output Fields | 23686
- Sample Output | 23686
- Release Information | 23689

Syntax

```
show route inactive-path
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route inactive-path
<brief | detail | extensive | terse>
```

Description

Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.

Options

none	Display all inactive routes.

brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to `brief`.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the `show route` command, the `show route detail` command, the `show route extensive` command, or the `show route terse` command.

Sample Output

show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30     [OSPF/10] 03:57:28, metric 1
                    > via so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/8          [Direct/0] 04:39:56
                    > via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
                      AS path: 100 I
```

```
                    > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-path detail

```
user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:     1
                Age: 3:58:24    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
```

```
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
                Inactive reason: No difference
                Age: 4:40:52
                Task: IF
                AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
         BGP    Preference: 170/-101
                Next-hop reference count: 6
                Source: 10.12.80.1
                Next hop: 10.12.80.1 via ge-6/3/2.0, selected
                State: <Ext>
                Inactive reason: Route Preference
                Peer AS:   100
                Age: 4:39:13
                Task: BGP_100.10.12.80.1+179
                AS path: 100 I
                Localpref: 100
                Router ID: 10.0.0.0
```

show route inactive-path extensive

The output for the `show route inactive-path extensive` command is identical to that of the `show route inactive-path detail` command.

show route inactive-path terse

```
user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.100.12/30    O  10          1            >so-0/3/0.0
```

```
private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.0.0.0/8         D   0                        >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.80.0/30      B 170        100            >10.12.80.1      100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show route | 23512
- show route active-path | 23526
- show route detail | 23573
- show route extensive | 23619
- show route terse | 23824

show route inactive-prefix

IN THIS SECTION

- Syntax | 23690
- Syntax (EX Series Switches) | 23690
- Description | 23690
- Options | 23690
- Required Privilege Level | 23691
- Output Fields | 23691
- Sample Output | 23691
- Release Information | 23692

Syntax

```
show route inactive-prefix
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route inactive-prefix
<brief | detail | extensive | terse>
```

Description

Display inactive route destinations in each routing table.

Options

none	Display all inactive route destination.

brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the "show route" on page 23512 command, the "show route detail" on page 23573 command, the "show route extensive" on page 23619 command, or the "show route terse" on page 23824 command.

Sample Output

show route inactive-prefix

```
user@host> show route inactive-prefix

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32        [Direct/0] 00:04:54
                    > via lo0.0
```

show route inactive-prefix detail

```
user@host> show route inactive-prefix detail

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
127.0.0.1/32 (1 entry, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Hidden Martian Int>
```

```
                Age: 4:51
                Task: IF
                AS path: I00:04:54
                    > via lo0.0
```

show route inactive-prefix extensive

The output for the `show route inactive-prefix extensive` command is identical to that of the `show route inactive-path detail` command. For sample output, see "show route inactive-prefix detail" on page 23691.

show route inactive-prefix terse

```
user@host> show route inactive-prefix terse

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  127.0.0.1/32       D   0                        >lo0.0
```

Release Information

Command introduced before Junos OS Release 7.4.

show route instance

IN THIS SECTION

- Syntax | 23693
- Syntax (EX Series Switches and QFX Series) | 23693
- Description | 23693
- Options | 23693
- Required Privilege Level | 23694
- Output Fields | 23694

Sample Output | 23696
Release Information | 23698

Syntax

```
show route instance
<brief | detail | summary>
<instance-name>
<logical-system (all | logical-system-name)>
<operational>
```

Syntax (EX Series Switches and QFX Series)

```
show route instance
<brief | detail | summary>
<instance-name>
<operational>
```

Description

Display routing instance information.

Options

none (Same as `brief`) Display standard information about all routing instances.

brief | detail | summary (Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to `brief`. (These options are not available with the `operational` keyword.)

instance-name (Optional) Display information for all routing instances whose name begins with this string (for example, `cust1`, `cust11`, and `cust111` are all displayed when you run the `show route instance cust1` command).

logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
operational	(Optional) Display operational routing instances.

Required Privilege Level

view

Output Fields

Table 1093 on page 23694 lists the output fields for the `show route instance` command. Output fields are listed in the approximate order in which they appear.

Table 1093: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or *instance-name*	Name of the routing instance.	All levels
Operational Routing Instances	(`operational` keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: `forwarding`, `l2vpn`, `no-forwarding`, `vpls`, `virtual-router`, or `vrf`.	All levels
State	State of the routing instance: `active` or `inactive`.	`brief` `detail` none
Interfaces	Name of interfaces belonging to this routing instance.	`brief` `detail` none
Restart State	Status of graceful restart for this instance: `Pending` or `Complete`.	`detail`
Path selection timeout	Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is `300`.	`detail`

Table 1093: show route instance Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Tables	Tables (and number of routes) associated with this routing instance.	`brief` `detail` none
Route-distinguisher	Unique route distinguisher associated with this routing instance.	`detail`
Vrf-import	VPN routing and forwarding instance import policy name.	`detail`
Vrf-export	VPN routing and forwarding instance export policy name.	`detail`
Vrf-import-target	VPN routing and forwarding instance import target community name.	`detail`
Vrf-export-target	VPN routing and forwarding instance export target community name.	`detail`
Vrf-edge-protection-id	Context identifier configured for edge-protection.	`detail`
Fast-reroute-priority	Fast reroute priority setting for a VPLS routing instance: `high`, `medium`, or `low`. The default is `low`.	`detail`
Restart State	Restart state: • `Pending:`*`protocol-name`*—List of protocols that have not yet completed graceful restart for this routing table. • `Complete`—All protocols have restarted for this routing table.	`detail`
Primary rib	Primary table for this routing instance.	`brief` none `summary`
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Sample Output

show route instance

```
user@host> show route instance
Instance             Type
         Primary RIB                                    Active/holddown/hidden
master               forwarding
         inet.0                                         16/0/1
         iso.0                                          1/0/0
         mpls.0                                         0/0/0
         inet6.0                                        2/0/0
         l2circuit.0                                    0/0/0
__juniper_private1__ forwarding
         __juniper_private1__.inet.0                    12/0/0
         __juniper_private1__.inet6.0                   1/0/0
```

show route instance detail (VPLS Routing Instance)

```
user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls              State: Active
  Interfaces:
    lsi.1048833
    lsi.1048832
    fe-0/1/0.513
  Route-distinguisher: 10.255.37.65:1
  Vrf-import: [ __vrf-import-test-vpls-internal__ ]
  Vrf-export: [ __vrf-export-test-vpls-internal__ ]
  Vrf-import-target: [ target:300:1 ]
  Vrf-export-target: [ target:300:1 ]
  Vrf-edge-protection-id:  166.1.3.1  Fast-reroute-priority: high
  Tables:
    test-vpls.l2vpn.0          : 3 routes (3 active, 0 holddown, 0 hidden)
```

show route instance operational

```
user@host> show route instance operational
Operational Routing Instances:

master
default
```

show route instance summary

```
user@host> show route instance summary
Instance            Type            Primary rib             Active/holddown/hidden
master              forwarding
                                    inet.0                  15/0/1
                                    iso.0                   1/0/0
                                    mpls.0                  35/0/0
                                    l3vpn.0                 0/0/0
                                    inet6.0                 2/0/0
                                    l2vpn.0                 0/0/0
                                    l2circuit.0             0/0/0
BGP-INET            vrf
                                    BGP-INET.inet.0         5/0/0
                                    BGP-INET.iso.0          0/0/0
                                    BGP-INET.inet6.0        0/0/0
BGP-L               vrf
                                    BGP-L.inet.0            5/0/0
                                    BGP-L.iso.0             0/0/0
                                    BGP-L.mpls.0            4/0/0
                                    BGP-L.inet6.0           0/0/0
L2VPN               l2vpn
                                    L2VPN.inet.0            0/0/0
                                    L2VPN.iso.0             0/0/0
                                    L2VPN.inet6.0           0/0/0
                                    L2VPN.l2vpn.0           2/0/0
LDP                 vrf
                                    LDP.inet.0              4/0/0
                                    LDP.iso.0               0/0/0
                                    LDP.mpls.0              0/0/0
                                    LDP.inet6.0             0/0/0
                                    LDP.l2circuit.0         0/0/0
```

```
OSPF                vrf
                                OSPF.inet.0            7/0/0
                                OSPF.iso.0             0/0/0
                                OSPF.inet6.0           0/0/0
RIP                 vrf
                                RIP.inet.0             6/0/0
                                RIP.iso.0              0/0/0
                                RIP.inet6.0            0/0/0
STATIC              vrf
                                STATIC.inet.0          4/0/0
                                STATIC.iso.0           0/0/0
                                STATIC.inet6.0         0/0/0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling
- *Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart*

show route label

IN THIS SECTION

- Syntax | 23699
- Syntax (EX Series Switches) | 23699
- Description | 23699
- Options | 23699
- Required Privilege Level | 23699
- Output Fields | 23699
- Sample Output | 23700
- Release Information | 23704

Syntax

```
show route label label
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route label label
<brief | detail | extensive | terse>
```

Description

Display the routes based on a specified Multiprotocol Label Switching (MPLS) label value.

Options

label	Value of the MPLS label.
brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field table for the "show route" on page 23512 command, the "show route detail" on page 23573 command, the "show route extensive" on page 23619 command, or the "show route terse" on page 23824 command.

Sample Output

show route label terse

```
user@host> show route label 100016 terse

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 100016             V 170                        >10.12.80.1
```

show route label

```
user@host> show route label 100016

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
100016             *[VPN/170] 03:25:41
                    > to 10.12.80.1 via ge-6/3/2.0, Pop
```

show route label detail

```
user@host> show route label 100016 detail

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
100016 (1 entry, 1 announced)
        *VPN    Preference: 170
                Next-hop reference count: 2
                Source: 10.12.80.1
                Next hop: 10.12.80.1 via ge-6/3/2.0, selected
                Label operation: Pop
                State: <Active Int Ext>
                Local AS:     65001
                Age: 3:23:31
```

```
                Task: BGP.0.0.0.0+179
                Announcement bits (1): 0-KRT
                AS path: 65100 I
                Ref Cnt: 2
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
                State: **Active Int>
                Local AS:  65001
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1, src
192.168.142.2
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
```

```
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
                State: **Active Int>
                Local AS:  65001
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1, src
192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 10.1.8.2 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 10.1.4.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    65010
                Age: 54:05      Metric: 1
                Validation State: unverified
```

```
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                FECs bound to route: P2MP root-addr 10.1.1.1, grp: 232.1.1.1, src: 192.168.219.11
                Primary Upstream : 10.1.1.3:0--10.1.1.2:0
                   RPF Nexthops  :
                        ge-1/2/15.0, 10.1.94.1, Label: 301568, weight: 0x1
                        ge-1/2/14.0, 10.1.3.1, Label: 301568, weight: 0x1
                Backup Upstream : 10.1.1.3:0--10. 1.1.6:0
                   RPF Nexthops :
                        ge-1/2/20.0, 10.1.96.1, Label: 301584, weight: 0xfffe
                        ge-1/2/19.0, 10.1.6.1, Label: 301584, weight: 0xfffe
```

show route label detail (Dynamic List Next Hop)

The output for `show route label detail` shows the two indirect next hop for an ESI.

```
user@host> show route label 299952 detail
mpls.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
299952 (1 entry, 1 announced)
TSI:
KRT in-kernel 299952 /52 -> {Dyn list:indirect(1048577), indirect(1048574)}
        *EVPN   Preference: 7
                Next hop type: Dynamic List, Next hop index: 1048575
                Address: 0x13f497fc
                Next-hop reference count: 5
                Next hop: ELNH Address 0xb7a3d90 uflags EVPN data
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xb7a3d90
                    Next-hop reference count: 3
                    Protocol next hop: 10.255.255.2
                    Label operation: Push 301344
                    Indirect next hop: 0x135b5c00 1048577 INH Session ID: 0x181
                        Next hop type: Router, Next hop index: 619
                        Address: 0xb7a3d30
                        Next-hop reference count: 4
                        Next hop: 10.1.0.4 via ge-0/0/1.0
                        Label operation: Push 301344, Push 299792(top)
                        Label TTL action: no-prop-ttl, no-prop-ttl(top)
                        Load balance label: Label 301344: None; Label 299792: None;
                        Label element ptr: 0xb7a3cc0
```

```
                Label parent element ptr: 0xb7a34e0
                Label element references: 1
                Label element child references: 0
                Label element lsp id: 0
        Next hop: ELNH Address 0xb7a37f0 uflags EVPN data
            Next hop type: Indirect, Next hop index: 0
            Address: 0xb7a37f0
            Next-hop reference count: 3
            Protocol next hop: 10.255.255.3
            Label operation: Push 301632
            Indirect next hop: 0x135b5480 1048574 INH Session ID: 0x180
                Next hop type: Router, Next hop index: 600
                Address: 0xb7a3790
                Next-hop reference count: 4
                Next hop: 10.1.0.4 via ge-0/0/1.0
                Label operation: Push 301632, Push 299776(top)
                Label TTL action: no-prop-ttl, no-prop-ttl(top)
                Load balance label: Label 301632: None; Label 299776: None;
                Label element ptr: 0xb7a3720
                Label parent element ptr: 0xb7a3420
                Label element references: 1
                Label element child references: 0
                Label element lsp id: 0
        State: <Active Int>
        Age: 1:18
        Validation State: unverified
        Task: evpn global task
        Announcement bits (2): 1-KRT 2-evpn global task
        AS path: I
        Routing Instance blue, Route Type Egress-MAC, ESI 00:11:22:33:44:55:66:77:88:99
```

show route label extensive

The output for the show route label extensive command is identical to that of the `show route label detail` command. For sample output, see "show route label detail" on page 23700.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs

show route label-switched-path

IN THIS SECTION

- Syntax | 23705
- Syntax (EX Series Switches) | 23705
- Description | 23705
- Options | 23706
- Required Privilege Level | 23706
- Output Fields | 23706
- Sample Output | 23706
- Release Information | 23707

Syntax

```
show route label-switched-path path-name
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route label-switched-path path-name
<brief | detail | extensive | terse>
```

Description

Display the routes used in an MPLS label-switched path (LSP).

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output.
path-name	LSP tunnel name.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the `show route` command, the `show route detail` command, the `show route extensive` command, or the `show route terse` command.

Sample Output

show route label-switched-path

```
user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.1/32          [MPLS/7] 00:00:06, metric 0
                    > to 10.111.1.9 via s0-0/0/0, label-switched-path sf-to-ny
10.3.3.3/32         *[MPLS/7] 00:00:06, metric 0
                    > to 10.111.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.2.2.2/32         *[MPLS/7] 00:00:06, metric 0
                    > to 10.111.1.9 via s0-0/0/0, label-switched-path sf-to-ny
10.4.4.4/32         *[MPLS/7] 00:00:06, metric 0
                      to 10.111.1.9 via s0-0/0/0, label-switched-path abc
                    > to 10.111.1.9 via s0-0/0/0, label-switched-path xyz
                      to 10.111.1.9 via s0-0/0/0, label-switched-path sf-to-ny
```

```
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0
                    > to 10.111.1.9 via s0-0/0/0, label-switched-path sf-to-ny

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

Release Information

Command introduced before Junos OS Release 7.4.

show route localization

IN THIS SECTION

- Syntax | 23707
- Description | 23707
- Options | 23708
- Required Privilege Level | 23708
- Output Fields | 23708
- Sample Output | 23708
- Release Information | 23709

Syntax

```
show route localization
```

Description

(T320, T640, and T1600 routers only) Display route localization details.

Options

detail Display detailed output.

Required Privilege Level

view

Output Fields

Table 1094 on page 23708 lists the output fields for the `show route localization` command. Output fields are listed in the approximate order in which they appear.

Table 1094: show route localization Output Fields

Field Name	Field Description
FIB-local	FPCs configured as FIB-local.
FIB-remote	FPCs configured as FIB-remote.
Normal	FPCs neither configured as **FIB-local** or **FIB-remote**.
Protocols	IPv4 (inet) or IPv6 (inet6) traffic configured for route localization.

Sample Output

command-name

```
user@R0> show route localization
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:  FPC2(4,5)
  FIB-remote: FPC0, FPC1
  Normal:     FPC3, FPC4, FPC5, FPC6, FPC7
```

command-name

```
user@R0> show route localization detail
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:  FPC2(4,5)
  FIB-remote: FPC0, FPC1
  Normal:     FPC3, FPC4, FPC5, FPC6, FPC7
FIB localization configuration
  Protocols:  inet, inet6
  FIB-local:  FPC2
  FIB-remote: FPC0, FPC1
Forwarding Engine addresses
  FPC0: 1
  FPC1: 2
  FPC2: 4, 5
  FPC3: 6
  FPC4: 8
  FPC5: 11
  FPC6: 13
  FPC7: 15
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring Packet Forwarding Engine FIB Localization

show route martians

IN THIS SECTION

- Syntax | 23710
- Syntax (EX Series Switches) | 23710

- Description | 23710
- Options | 23710
- Required Privilege Level | 23711
- Output Fields | 23711
- Sample Output | 23711
- Release Information | 23713

Syntax

```
show route martians
<logical-system (all | logical-system-name)>
<table routing-table-name>
```

Syntax (EX Series Switches)

```
show route martians
<table routing-table-name>
```

Description

Display the martian (invalid and ignored) entries associated with each routing table.

Options

none	Display standard information about route martians for all routing tables.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
table *routing-table-name*	(Optional) Display information about route martians for all routing tables whose name begins with this string (for example, **inet.0** and **inet6.0** are both displayed when you run the `show route martians table inet` command).

Required Privilege Level

view

Output Fields

Table 1095 on page 23711 lists the output fields for the `show route martians` command. Output fields are listed in the approximate order in which they appear

Table 1095: show route martians Output Fields

Field Name	Field Description
table-name	Name of the route table in which the route martians reside.
destination-prefix	Route destination.
match value	Route match parameter.
status	Status of the route: **allowed** or **disallowed**.

Sample Output

show route martians

```
user@host> show route martians

inet.0:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed

inet.1:
             0.0.0.0/0 exact -- allowed
```

```
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed

inet.2:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed

inet.3:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed
...

inet6.0:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed

inet6.1:
             ::1/128 exact -- disallowed

inet6.2:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed

inet6.3:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed
...
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Removing the Class E Prefix on Martian Addresses

Understanding Martian Addresses

show route match-prefix

IN THIS SECTION

- Syntax | 23713
- Description | 23713
- Options | 23714
- Additional Information | 23714
- Required Privilege Level | 23714
- Output Fields | 23714
- Sample Output | 23714
- Release Information | 23715

Syntax

```
show route match-prefix match-prefix;
```

Description

Allows you to search for routes using regular expressions based on the extended (modern) regular expressions as defined in POSIX 1003.2.

Options

match-prefix Regular expression to match formatted prefix.

Additional Information

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the `show route` command, the `show route detail` command, the `show route extensive` command, or the `show route terse` command.

Sample Output

show route match-prefix *:10.255.2.200:6:* (Show all routes matching route distributor 10.255.2.200:6)

```
user@host> show route match-prefix *:10.255.2.200:6:*
```

show route match-prefix 2* (Show all mvpn type-2 routes)

```
user@host> show route table vrf1.inet.0 match-prefix 2*
```

```
vrf1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.1.1.0/24        *[Direct/0] 00:48:39
                    >  via ge-3/0/8.0
20.1.1.1/32        *[Local/0] 00:48:39
                       Local via ge-3/0/8.0
```

show route match-prefix *:224.* (Show all routes matching group 224/4)

```
user@host> show route match-prefix *:224.*
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

> *Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands, Configuration Statements, and Hierarchies*

show route next-hop

IN THIS SECTION

- Syntax | 23715
- Syntax (EX Series Switches) | 23716
- Description | 23716
- Options | 23716
- Required Privilege Level | 23716
- Output Fields | 23716
- Sample Output | 23716
- Release Information | 23718

Syntax

```
show route next-hop next-hop
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route next-hop next-hop
<brief | detail | extensive | terse>
```

Description

Display the entries in the routing table that are being sent to the specified next-hop address.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of ouput.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
next-hop	Next-hop address.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the `show route` command, the `show route detail` command, the `show route extensive` command, or the `show route terse` command.

Sample Output

show route next-hop

```
user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 06:26:25
```

```
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
172.16.0.0/12      *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16     *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
192.168.102.0/23   *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.0/24    *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.192/32  *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop terse

```
user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                       >192.168.71.254
* 10.209.0.0/16      S   5                       >192.168.71.254
```

```
* 172.16.0.0/12      S   5                       >192.168.71.254
* 192.168.0.0/16     S   5                       >192.168.71.254
* 192.168.102.0/23   S   5                       >192.168.71.254
* 207.17.136.0/24    S   5                       >192.168.71.254
* 207.17.136.192/32  S   5                       >192.168.71.254

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show route | 23512
- show route detail | 23573
- show route extensive | 23619
- show route terse | 23824

show route no-community

IN THIS SECTION

- Syntax | 23719
- Syntax (EX Series Switches) | 23719
- Description | 23719
- Options | 23719
- Required Privilege Level | 23720
- Output Fields | 23720
- Sample Output | 23720
- Release Information | 23723

Syntax

```
show route no-community
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route no-community
<brief | detail | extensive | terse>
```

Description

Display the route entries in each routing table that are not associated with any community.

Options

none (Same as `brief`) Display the route entries in each routing table that are not associated with any community.

brief | detail | extensive | terse	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the "show route" on page 23512 command, the "show route detail" on page 23573 command, the "show route extensive" on page 23619 command, or the "show route terse" on page 23824 command.

Sample Output

show route no-community

```
user@host> show route no-community
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 00:36:27
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 00:36:27
                    > to 192.168.71.254 via fxp0.0
10.255.71.52/32    *[Direct/0] 00:36:27
                    > via lo0.0
10.255.71.63/32    *[OSPF/10] 00:04:39, metric 1
                    > to 35.1.1.2 via ge-3/1/0.0
10.255.71.64/32    *[OSPF/10] 00:00:08, metric 2
                    > to 35.1.1.2 via ge-3/1/0.0
10.255.71.240/32   *[OSPF/10] 00:05:04, metric 2
                       via so-0/1/2.0
                    > via so-0/3/2.0
10.255.71.241/32   *[OSPF/10] 00:05:14, metric 1
                    > via so-0/1/2.0
10.255.71.242/32   *[OSPF/10] 00:05:19, metric 1
                    > via so-0/3/2.0
```

```
172.16.12.0/24        *[OSPF/10] 00:05:14, metric 2
                    > via so-0/3/2.0
172.16.14.0/24        *[OSPF/10] 00:00:08, metric 3
                    > to 35.1.1.2 via ge-3/1/0.0
                      via so-0/1/2.0
                      via so-0/3/2.0
172.16.16.0/24        *[OSPF/10] 00:05:14, metric 2
                    > via so-0/1/2.0
.....
```

show route no-community detail

```
user@host> show route no-community detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 38:08
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 38:08
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
....
```

show route no-community extensive

```
user@host> show route no-community extensive
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:03:33
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:03:33
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```

show route no-community terse

```
user@host> show route no-community terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                       >192.168.71.254
* 10.209.0.0/16      S   5                       >192.168.71.254
* 10.255.71.52/32    D   0                       >lo0.0
* 10.255.71.63/32    O  10          1            >35.1.1.2
* 10.255.71.64/32    O  10          2            >35.1.1.2
```

```
* 10.255.71.240/32   O  10          2             so-0/1/2.0
                                                 >so-0/3/2.0
* 10.255.71.241/32   O  10          1            >so-0/1/2.0
* 10.255.71.242/32   O  10          1            >so-0/3/2.0
* 172.16.12.0/24        O  10          2            >so-0/3/2.0
* 172.16.14.0/24        O  10          3            >35.1.1.2
                                                  so-0/1/2.0
                                                  so-0/3/2.0
* 172.16.16.0/24        O  10          2            >so-0/1/2.0
...
```

Release Information

Command introduced before Junos OS Release 7.4.

show route output

IN THIS SECTION

- Syntax | 23723
- Syntax (EX Series Switches) | 23724
- Description | 23724
- Options | 23724
- Required Privilege Level | 23724
- Output Fields | 23724
- Sample Output | 23725
- Release Information | 23727

Syntax

```
show route output (address ip-address | interface interface-name)
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route output (address ip-address | interface interface-name)
<brief | detail | extensive | terse>
```

Description

Display the entries in the routing table learned through static routes and interior gateway protocols that are to be sent out the interface with either the specified IP address or specified name.

To view routes advertised to a neighbor or received from a neighbor for the BGP protocol, use the `show route advertising-protocol bgp` and `show route receive-protocol bgp` commands instead.

Options

address *ip-address*	Display entries in the routing table that are to be sent out the interface with the specified IP address.
brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to `brief`.
interface *interface-name*	Display entries in the routing table that are to be sent out the interface with the specified name.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route output address

```
user@host> show route output address 172.16.36.1/24

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.36.0/24        *[Direct/0] 00:19:56
                    > via so-0/1/2.0
                    [OSPF/10] 00:19:55, metric 1
                    > via so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output address detail

```
user@host> show route output address 172.16.36.1 detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
172.16.36.0/24 (2 entries, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via so-0/1/2.0, selected
                State: <Active Int>
                Age: 23:00
                Task: IF
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
```

```
                Next hop: via so-0/1/2.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Age: 22:59      Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output address extensive

The output for the `show route output address extensive` command is identical to that of the `show route output address detail` command. For sample output, see "show route output address detail" on page 23725.

show route output address terse

```
user@host> show route output address 172.16.36.1 terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 172.16.36.0/24        D   0                       >so-0/1/2.0
                     O  10          1            >so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- show route | 23512
- show route detail | 23573
- show route extensive | 23619
- show route terse | 23824

show route pfe-install-error

IN THIS SECTION

- Syntax | 23727
- Description | 23728
- Options | 23728
- Additional Information | 23728
- Required Privilege Level | 23728
- Output Fields | 23729
- Sample Output | 23729
- Release Information | 23730

Syntax

```
show route pfe-install-error
```

Description

(Junos OS Evolved only) When the forwarding plane is not forwarding traffic, use this command to see if the Packet Forwarding Engine has sent an error for the route that is dropping traffic. You must first configure the `pfe-install-error` statement at the `[edit routing-options forwarding-table]` hierarchy level and restart the rpd process to enable this error reporting.

Options

Additional Information

These errors are also displayed as part of the `show route extensive` command:

```
user@host> show route any-route-prefix extensive
inet.0: 11 destinations, 14 routes (11 active, 0 holddown, 0 hidden)
10.54.18.6/32 (1 entry, 1 announced)
TSI:
KRT in-kernel <>/32 -> {….}
        *BGP    Preference: 170/-101
                Route Error State: YES     Last error received: 20XX-XX-XX XX:XX:XX
                Next hop type: Router, Next hop index: 10001
```

Required Privilege Level

view

Output Fields

Table 1096: show route pfe-install-error Output Fields

Field Name	Field Description
Error category	There are two categories: • **PFE:** Error is received directly from the Packet Forwarding Engine. • **DERIVED:** Error is derived within the rpd process from input from the Packet Forwarding Engine. For example, if a route points to a next hop and the rpd process receives an error on that next hop from the Packet Forwarding Engine, then the rpd process also knows that there is an error on the route that points to the next hop.
Handle	Address of the object.
Route	IP address of the route.
Time	Time when error occurred.
Type	Error type is 0x1.

Sample Output

show route pfe-install-error

```
user@host> show route pfe-install-error
```

```
Route               Error category        Handle          Type     Time

10.31.1.4            PFE                  0x5594bd5523c0  0x1      20XX-XX-XX XX:XX:XX

10.31.1.1            PFE                  0x5594bd552328  0x1      20XX-XX-XX XX:XX:XX
```

Release Information

Command introduced in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

forwarding-table

show route protocol

IN THIS SECTION

- Syntax | 23730
- Syntax (EX Series Switches) | 23731
- Syntax (PTX and ACX Series Switches) | 23731
- Operational mode | 23731
- Description | 23731
- Options | 23731
- Required Privilege Level | 23733
- Output Fields | 23733
- Sample Output | 23733
- Release Information | 23739

Syntax

```
show route protocol protocol
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route protocol protocol
<brief | detail | extensive | terse>
```

Syntax (PTX and ACX Series Switches)

```
show route protocol protocol
(index index | prefix prefix)
<brief | detail | extensive | terse>
```

Operational mode

`cli` mode

`cli-pfe` mode.

Description

Display the route entries in the routing table that were learned from a particular protocol. Displays details about routes installed in `cli-pfe` mode after FIB compression of FIB compression statistics (PTX and ACX series devices only).

Options

brief | detail | extensive | terse (Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to `brief`.

logical-system (all | *logical-system-name*) (Optional) Perform this operation on all logical systems or on a particular logical system.

protocol Protocol from which the route was learned:

- `access`—Access route for use by DHCP application
- `access-internal`—Access-internal route for use by DHCP application
- `aggregate`—Locally generated aggregate route

- `arp`—Route learned through the Address Resolution Protocol
- `atmvpn`—Asynchronous Transfer Mode virtual private network
- `bgp`—Border Gateway Protocol
- `ccc`—Circuit cross-connect
- `direct`—Directly connected route
- `dvmrp`—Distance Vector Multicast Routing Protocol
- `esis`—End System-to-Intermediate System
- `flow`—Locally defined flow-specification route
- `frr`—Precomputed protection route or backup route used when a link goes down
- `isis`—Intermediate System-to-Intermediate System
- `ldp`—Label Distribution Protocol
- `l2circuit`—Layer 2 circuit
- `l2vpn`—Layer 2 virtual private network
- `local`—Local address
- `mpls`—Multiprotocol Label Switching
- `msdp`—Multicast Source Discovery Protocol
- `ospf`—Open Shortest Path First versions 2 and 3
- `ospf2`—Open Shortest Path First versions 2 only
- `ospf3`—Open Shortest Path First version 3 only
- `pim`—Protocol Independent Multicast
- `rip`—Routing Information Protocol
- `ripng`—Routing Information Protocol next generation
- `rsvp`—Resource Reservation Protocol
- `rtarget`—Local route target virtual private network

- `static`—Statically defined route
- `tunnel`—Dynamic tunnel
- `vpn`—Virtual private network

> **NOTE**: EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the `show route` command, the `show route detail` command, the `show route extensive` command, or the `show route terse` command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol arp

```
user@host>  show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)
```

```
inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.4/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.5/32        [ARP/4294967293] 00:04:32, from 20.20.1.1
                      Unusable
20.20.1.6/32        [ARP/4294967293] 00:04:34, from 20.20.1.1
                      Unusable
20.20.1.7/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.8/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.9/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.10/32       [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.11/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.12/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.13/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
...
```

show route protocol bgp

```
user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21       *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                      AS path: 10458 14203 2914 4788 4788 I
                    > to 192.168.167.254 via fxp0.0
```

show route protocol direct

```
user@host> show route protocol direct

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.8.0/24         *[Direct/0] 17w0d 10:31:49
                     > via fe-1/3/1.0
10.255.165.1/32    *[Direct/0] 25w4d 04:13:18
                     > via lo0.0
172.16.30.0/24      *[Direct/0] 17w0d 23:06:26
                     > via fe-1/3/2.0
192.168.164.0/22   *[Direct/0] 25w4d 04:13:20
                     > via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
                   *[Direct/0] 25w4d 04:13:21
                     > via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::10:255:165:1/128
                   *[Direct/0] 25w4d 04:13:21
                     > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
                   *[Direct/0] 25w4d 04:13:21
                     > via lo0.0
```

show route protocol frr

```
user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.3 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.4 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32       *[FRR/200] 00:05:35, from 20.20.1.1
                    > to 20.20.1.5 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32       *[FRR/200] 00:05:37, from 20.20.1.1
                    > to 20.20.1.6 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.7 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.8 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.9 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32      *[FRR/200] 00:05:38, from 20.20.1.1
...
```

show route protocol ldp

```
user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0

private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100080             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Swap 100000
```

show route protocol ospf (Layer 3 VPN)

```
user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30       *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.39.1.8/30        [OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
10.255.14.171/32   *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.255.14.179/32   *[OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
172.16.233.5/32       *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30       [OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
10.255.14.173/32   *[OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
172.16.233.5/32       *[OSPF/10] 20:26:20, metric 1
```

show route protocol rip

```
user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32   *[RIP/100] 20:24:34, metric 2
                    > to 10.39.1.22 via t3-0/2/2.0
172.16.233.9/32       *[RIP/100] 00:03:59, metric 1
```

show route protocol rsvp extensive

```
user@host> show route protocol rsvp extensive
inet.0: 37 destinations, 37 routes (36 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 4 routes (3 active, 0 holddown, 0 hidden)
192.168.4.4/32 (2 entries, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7/1
                Next hop type: Router, Next hop index: 0
                Address: 0xaad4fbc
                Next-hop reference count: 2, key opaque handle: 0x0, non-key opaque handle: 0x0
                Next hop: 192.168.0.1 via ge-0/0/0.0 weight 0x1, selected
                Label-switched-path PE1_to_PE2_1
                Label operation: Push 299920
                Label TTL action: prop-ttl
                Load balance label: Label 299920: None;
                Label element ptr: 0x9938be8
                Label parent element ptr: 0x0
                Label element references: 3
                Label element child references: 1
                Label element lsp id: 2
                Session Id: 0
                Next hop: 192.168.1.1 via ge-0/0/1.0 weight 0x8001 uflags Mark connection
protection tlv
                Label-switched-path Bypass->192.168.0.1
                Label operation: Push 299920
                Label TTL action: prop-ttl
                Load balance label: Label 299920: None;
                Label element ptr: 0x9938eb8
                Label parent element ptr: 0x9938be8
                Label element references: 2
                Label element child references: 0
```

```
                Label element lsp id: 2
                Session Id: 0
                State: <Active Int>
                Local AS:   100
                Age: 1:20       Metric: 100
                Validation State: unverified
                Actual IGP metric: 2
                Task: RSVP
                Announcement bits (2): 0-Resolve tree 3 3-Resolve tree 1
                AS path: I
                Thread: junos-main
```

show route protocol index prefix detail (ACX7100-32C, ACX7100-48L, and ACX7509 Devices)

Displays if a route is installed or not when FIB compression is enabledIf the route is not installed, the command displays the details of the parent route installed, that compressed the route.

```
user@root:pfe> show route proto 2 index 0 prefix 200.1.0.16 detail
Protocol: 2
Protocol     : IPv4
Table        : default
Prefix       : 200.1.0.16 (primary)
NH           : 25035 (software)
Flags        : 0x00008000
Details      :
guid          : 889058232562
type          :  user
nhid          :  25035
Forwarding state:
installed?    :  no
(Installed parent: 200.1.0/23)
```

Release Information

Command introduced before Junos OS Release 7.4.

`ospf2` and `ospf3` options introduced in Junos OS Release 9.2.

`ospf2` and `ospf3` options introduced in Junos OS Release 9.2 for EX Series switches.

`flow` option introduced in Junos OS Release 10.0.

`flow` option introduced in Junos OS Release 10.0 for EX Series switches.

RELATED DOCUMENTATION

- show route | 23512
- show route detail | 23573
- show route extensive | 23619
- show route terse | 23824

show route range

IN THIS SECTION

- Syntax | 23740
- Syntax (EX Series Switches) | 23741
- Description | 23741
- Options | 23741
- Required Privilege Level | 23741
- Output Fields | 23741
- Sample Output | 23741
- Release Information | 23746

Syntax

```
show route range
<brief | detail | extensive | terse>
<destination-prefix>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route range
<brief | detail | extensive | terse>
<destination-prefix>
```

Description

Display routing table entries using a prefix range.

Options

none	Display standard information about all routing table entries using a prefix range.
brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.
destination-prefix	Destination and prefix mask for the range.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the "show route" on page 23512 command, the "show route detail" on page 23573 command, the "show route extensive" on page 23619 command, or the "show route terse" on page 23824 command.

Sample Output

show route range

```
user@host> show route range
```

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
10.255.71.14/32    *[Direct/0] 00:30:01
                    > via lo0.0
172.16.0.0/12      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16     *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:30:01
                    > via fxp0.0
192.168.71.14/32   *[Local/0] 00:30:01
                      Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
...
```

show route range destination-prefix

```
user@host> show route range 192.168.0.0/16

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16     *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:31:14
                    > via fxp0.0
192.168.71.14/32   *[Local/0] 00:31:14
                      Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
```

show route range detail

```
user@host> show route range detail

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.255.71.14/32 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Age: 30:05
                Task: IF
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
```

```
                Announcement bits (1): 0-KRT
                AS path: I

...
```

show route range extensive

```
user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.255.71.14/32 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Age: 30:17
```

```
                Task: IF
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

...
```

show route range terse

```
user@host> show route range terse

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                       >192.168.71.254
* 10.209.0.0/16      S   5                       >192.168.71.254
* 10.255.71.14/32    D   0                       >lo0.0
* 172.16.0.0/12      S   5                       >192.168.71.254
* 192.168.0.0/16     S   5                       >192.168.71.254
* 192.168.64.0/21    D   0                       >fxp0.0
* 192.168.71.14/32   L   0                        Local
* 192.168.102.0/23   S   5                       >192.168.71.254
* 207.17.136.0/24    S   5                       >192.168.71.254
* 207.17.136.192/32  S   5                       >192.168.71.254

__juniper_private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.0.0.0/8         D   0                       >fxp2.0
```

```
                     D   0                       >fxp1.0
* 10.0.0.4/32        L   0                        Local

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152
*                    D   0                       >lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  abcd::10:255:71:14/128
*                    D   0                       >lo0.0
  fe80::280:42ff:fe11:226f/128
*                    D   0                       >lo0.0

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  fe80::280:42ff:fe11:226f/128
*                    D   0                       >lo0.16385
```

Release Information

Command introduced before Junos OS Release 7.4.

show route receive-protocol

IN THIS SECTION

- Syntax | 23747
- Syntax (EX Series Switches) | 23747
- Description | 23747
- Options | 23747

- Additional Information | 23748
- Required Privilege Level | 23748
- Output Fields | 23748
- Sample Output | 23751
- Release Information | 23753

Syntax

```
show route receive-protocol protocol neighbor-address
 <brief | detail | extensive | terse>
<logical-system (all | logical-system-name)
```

Syntax (EX Series Switches)

```
show route receive-protocol protocol neighbor-address
<brief | detail | extensive | terse>
```

Description

Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
protocol neighbor-address	Protocol transmitting the route (`bgp`, `dvmrp`, `msdp`, `pim`, `rip`, or `ripng`) and address of the neighboring router from which the route entry was received.

Additional Information

The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.

Required Privilege Level

view

Output Fields

Table 1097 on page 23748 describes the output fields for the `show route receive-protocol` command. Output fields are listed in the approximate order in which they appear.

Table 1097: show route receive-protocol Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, inet.0.	All levels
number destinations	Number of destinations for which there are routes in the routing table.	All levels
number routes	Number of routes in the routing table and total number of routes in the following states: • `active` • `holddown` (routes that are in pending state before being declared inactive) • `hidden` (routes that are not used because of a routing policy)	All levels
Prefix	Destination prefix.	none `brief`
MED	Multiple exit discriminator value included in the route.	none `brief`
destination-prefix (entry, announced)	Destination prefix. The `entry` value is the number of routes for this destination, and the `announced` value is the number of routes being announced for this destination.	`detail extensive`

Table 1097: show route receive-protocol Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.	`detail extensive`
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table	`detail extensive`
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.	`detail extensive`
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.	`detail extensive`
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.	`detail extensive`
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.	`detail extensive`
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	All levels
Localpref or Lclpref	Local preference value included in the route.	All levels

Table 1097: show route receive-protocol Output Fields *(Continued)*

Field Name	Field Description	Level of Output
AS path	Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • `I`—IGP. • `E`—EGP. • `?`—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • `[ ]`—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893. • `[ ]`—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • `{ }`—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • `( )`—Parentheses enclose a confederation. • `( [ ] )`—Parentheses and brackets enclose a confederation set. **NOTE**: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
Route Labels	Stack of labels carried in the BGP route update.	`detail extensive`
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	`detail extensive`

Table 1097: show route receive-protocol Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	`detail extensive`
Communities	Community path attribute for the route.	`detail extensive`
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	`detail extensive`
Attrset AS	Number, local preference, and path of the AS that originated the route. These values are stored in the `Attrset` attribute at the originating routing device.	`detail extensive`
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	`detail extensive`
control flags	Control flags: `none` or `Site Down`.	`detail extensive`
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	`detail extensive`

Sample Output

show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix                  Next hop              MED     Lclpref    AS path
10.22.1.0/24            10.255.245.215       0       100        I
10.22.2.0/24            10.255.245.215       0       100        I
```

show route receive-protocol bgp detail (BGP Multicast)

```
user@host> show route receive-protocol bgp 10.4.6.1 detail
bgpmcast.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 4:3:0:0:32:10.2.20.20:32:225.1.1.1:10.4.4.4:10.4.0.1/240 (1 entry, 1 announced)
     Accepted
     Nexthop: 10.4.0.1
     Localpref: 100
     AS path: 65002 I
     Communities: target:10.4.4.4:0
     Tunnel type: AnyEncap, RPF tunnel:, Remote end point: 10.1.4.4
     Tunnel type: AnyEncap, Remote end point: 10.2.4.4
     Tunnel type: AnyEncap, Remote end point: 10.3.4.4
     Tunnel type: AnyEncap, Remote end point: 10.4.6.4
     Tunnel type: AnyEncap, Remote end point: 10.4.20.4
```

Show route receive protocol (Segment Routing Traffic Engineering)

```
show route receive protocol bgp 10.1.1.4
bgp.inetcolor.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

* 50-10.4.4.4-1234<sr6>/96 (1 entry, 0 announced)
     Import Accepted
     Distinguisher: 50
     Color: 1234
     Nexthop: 10.1.1.4
     Localpref: 100
     AS path: 65003 I
     Communities: target:10.1.1.1:1


inetcolor.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
* 10.4.4.4-1234<c6>/64 (1 entry, 1 announced)
     Import Accepted
     Color: 1234
     Nexthop: 10.1.1.4
     Localpref: 100
     AS path: 65003 I
     Communities: target:10.1.1.1:1
```

```
user@host# run show route receive-protocol bgp 2001:db8:5001:1::4
bgp.inet6color.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

* 50-2001:db8:1::4-1234<sr6>/192 (1 entry, 0 announced)
     Import Accepted
     Distinguisher: 50
     Color: 1234
     Nexthop: ::ffff:10.1.1.4
     Localpref: 100
     AS path: 65003 I
     Communities: target:10.1.1.1:1


inet6color.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
* 2001:db8::5-1234<c6>/160 (1 entry, 1 announced)
     Import Accepted
     Color: 1234
     Nexthop: ::ffff:10.1.1.5
     Localpref: 100
     AS path: 65003 I
     Communities: target:10.1.1.1:1
```

Release Information

Command introduced before Junos OS Release 7.4.

show route resolution

IN THIS SECTION

- Syntax | 23754
- Syntax (EX Series Switches) | 23754
- Description | 23754
- Options | 23754
- Required Privilege Level | 23755

- Output Fields | 23755
- Sample Output | 23756
- Release Information | 23759

Syntax

```
show route resolution
<brief | detail | extensive | summary>
<index index>
<logical-system (all | logical-system-name)>
<prefix>
<table routing-table-name>
<unresolved>
```

Syntax (EX Series Switches)

```
show route resolution
<brief | detail | extensive | summary>
<index index>
<prefix>
<table routing-table-name>
<unresolved>
```

Description

Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.

Options

none	Display standard information about all entries in the next-hop resolution database.
brief | detail | extensive | summary	(Optional) Display the specified level of output.

index *index*	(Optional) Show the index of the resolution tree.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
prefix network/ destination-prefix	(Optional) Display database entries for the specified address.
table *routing-table-name*	(Optional) Display information about a particular routing table (for example, **inet.0**) where policy-based export is currently enabled.
unresolved	(Optional) Display routes that could not be resolved.

Required Privilege Level

view

Output Fields

Table 1098 on page 23755 describes the output fields for the `show route resolution` command. Output fields are listed in the approximate order in which they appear.

Table 1098: show route resolution Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.
Tree index	Tree index identifier.
Nodes	Number of nodes in the tree.
Reference count	Number of references made to the next hop.
Contributing routing tables	Routing tables used for next-hop resolution.

Table 1098: show route resolution Output Fields *(Continued)*

Field Name	Field Description
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of **inet.0** resolving through **inet.0** and **inet.3**, this field indicates which routing table, **inet.0** or **inet.3**, provided the best path for a particular prefix.
Metric	Metric associated with the forwarding next hop.
Node path count	Number of nodes in the path.
Forwarding next hops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it. Merged—Merged next hops when recursive resolution of multipath is configured.

Sample Output

show route resolution detail

```
user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
  Node path count: 1
```

```
  Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
  Metric: 2 Node path count: 1
  Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
```

show route resolution (Multipath Resolution)

```
user@host> show route resolution detail
user@host> show route resolution detail 10.1.1.2
Tree Index: 1, Nodes 36, Reference Count 3
Contributing routing tables: inet.0 inet.3
Policy: [ abc ]
10.1.1.2/32 Originating RIB: inet.0
  Node path count: 1
  Next hop subtype: INDIRECT
  Indirect next hops: 2
        Protocol next hop: 10.1.1.1
        Inode flags: 0x206 path flags: 0x08
        Path fnh link: 0xc9321c0 path inh link: 0x0
        Indirect next hop: 0xb2b20f0 1048574 INH Session ID: 0x143
        Indirect path forwarding next hops: 1
                Next hop type: Router
                Next hop: 12.1.1.2 via ge-2/0/1.0
                Session Id: 0x144
                Next hop: 13.1.1.2 via ge-2/0/2.0
                Session Id: 0x145

    10.1.1.1/32 Originating RIB: inet.0
      Node path count: 1
  Node flags: 1
      Forwarding nexthops: 1 (Merged)
        Nexthop: 12.1.1.2 via ge-2/0/1.0

        Nexthop: 13.1.1.2 via ge-2/0/2.0
```

```
user@host> show route resolution summary
Tree Index: 1, Nodes 7, Reference Count 2
Contributing routing tables: inet.3
Tree Index: 2, Nodes 7, Reference Count 8213
Contributing routing tables: inet.3
Policy: [ RRwM ]
Tree Index: 3, Nodes 7, Reference Count 2
Contributing routing tables: inet6.3
Tree Index: 4, Nodes 1, Reference Count 1
Contributing routing tables: iso.0
Tree Index: 5, Nodes 1000061, Reference Count 13
Contributing routing tables: inet.0 inet.3
Policy: [ Community-RRwM ]
Tree Index: 6, Nodes 2013, Reference Count 6
Contributing routing tables: inet6.0 inet6.3
Policy: [ RRwM ]
Tree Index: 7, Nodes 7, Reference Count 1501
Contributing routing tables: inet6.3
Policy: [ RRwM ]
Tree Index: 8, Nodes 1000061, Reference Count 2
Contributing routing tables: inet.0 inet.3
Policy: [ RRwM ]
```

show route resolution summary

```
user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
 Contributing routing tables: :voice.inet.0 :voice.inet.3
 Tree Index: 2, Nodes 2, Reference Count 1
 Contributing routing tables: inet.3
 Tree Index: 3, Nodes 43, Reference Count 1
 Contributing routing tables: inet.0 inet.3
```

show route resolution unresolved

```
user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
        Protocol Nexthop: 10.255.71.238 Push 800000
```

```
        Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
        Protocol Nexthop: 10.255.70.103 Push 800008
        Indirect nexthop: 0 -
Tree Index 2
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Example: Configuring Route Resolution on PE Routers*

show route rpm-tracking

IN THIS SECTION

- Syntax | 23759
- Description | 23760
- Options | 23760
- Required Privilege Level | 23760
- Output Fields | 23760
- Sample Output | 23761
- Release Information | 23764

Syntax

```
show route rpm-tracking
<destination IP-address>
<owner name>
<test test-name>
```

Description

Display a brief summary of state of RPM-tracked routes along with the current state for a given test.

Options

destination *IP-address*
(Optional) Show RPM-tracked routes for a particular IPv4 or IPv6 destination prefix.

owner *name*
The probe owner as configured at the `[edit services rpm probe owner]` hierarchy level.

test *test-name*
The test name as configured at the `[edit services rpm probe owner test test-name]` hierarchy level.

Required Privilege Level

view

Output Fields

Table 1099 on page 23760 lists the output fields for the `show route rpm-tracking` command. You can filter on routing table name, destination prefix, RPM probe owner, and RPM test name. If no filter is present all RPM-tracked routes are displayed. Output fields are listed in the approximate order in which they appear.

Table 1099: show route rpm-tracking Output Fields

Field Name	Field Description
`Destination`	Displays the IPv4 or IPv6 address and optional prefix length of the configured target address.
`Next-Hop` • ucst • ulst	Specifies the IPv4 or IPv6 next-hop address of the route to be injected during failover. When there are multiple next-hop entries, a type attribute is shown to indicate whether it is a single unicast next-hop, `ucst`, or a list of unicast next-hops, `ulst`.

Table 1099: show route rpm-tracking Output Fields *(Continued)*

Field Name	Field Description
`Metric`	Specifies a number that is associated with route preference; when multiple routes have the same preference, the route with the lowest metric is made active in the routing table.
`Preference`	Specifies a number that is used to select routes to destinations in external autonomous systems (ASs) or routing domains. The route with lowest preference value is selected by the routing protocol process as the active one. This column in the output appears only if you configured the `preference` statement at the [`edit routing-options rpm-tracking route` *`destination-prefix`*] hierarchy level. .
`Route-tag`	Specifies a number that is used to represent and advertise a group of routes throughout the routing domain. The route with the lowest tag value is selected by the routing protocol process as the active one. This column in the output appears only if you configured the `tag` statement at the [`edit routing-options rpm-tracking route` *`destination-prefix`*] hierarchy level.
`Owner`	Name of the probe owner.
`Test Name`	Name of the test.
`State`	Display the state of the route injection action. Routes added to the routing table appear as `active` in RPM. The initial state of an RPM-tracked route, that is, before the first completion of its associated RPM test, is `inactive`. Routes removed from the routing table appear as `inactive`.

Sample Output

show route rpm-tracking

```
user@host> show route rpm-tracking

Routing table: inet.0
```

```
Destination       Next-Hop      Metric    Owner    Test Name   State
10.10.10.0/24     10.10.10.11     1       probe1   test1       Active
10.10.20.0/24                     1
                  10.10.10.22             probe1   test2       Active
                  10.10.10.33             probe1   test3       Inactive
10.10.30.0/24                     1
                  10.1.010.11             probe1   test1       Active

Routing table: inet6.0
2001:db8:10::/64  2001:db8:10::11 1       probe1   test1       Active
2001:db8:20::/64                  1
                  2001:db8:10::22         probe1   test2       Active
                  2001:db8:10::33         probe1   test3       Inactive
10.10.20.0/24                     1
                  2001:db8:10::11         probe1   test1       Active
```

show route rpm-tracking when preference and tag are configured

```
user@host> show route rpm-tracking

Routing table: inet.0
Destination        Next-Hop      Metric   preference   route-tag    Owner      Test Name   State
10.10.10.0/0       -             5        50           12           -          -
-
                   10.10.10.11   -        -            -            RPM-OWNER  RPM-TEST    Active
Destination        Next-Hop      Metric   preference   route-tag    Owner      Test Name   State
10.10.10.0/16      -             2        18           15           -          -
-
                   10.10.10.22   -        -            -            RPM-OWNER  RPM-TEST1
Inactive
                   10.10.10.33   -        -            -            RPM-OWNER  RPM-TEST    Active

Routing table: inet6.0
Destination        Next-Hop         Metric   preference   route-tag    Owner      Test Name
State
2001:db8:10::/64   -                2        11           4            -          -
-
                   2001:db8:10::11  -        -            -            RPM-OWNER  RPM-TEST
Active
```

show route rpm-tracking destination *destination-prefix*

```
user@host> show route rpm-tracking destination 10.39.0.0/16

Routing table: inet.0
Destination              Next-Hop                Metric    Owner           Test Name       State
10.39.0.0/16             10.20.21.2              2         probe-delegate  test7984
Active
10.39.1.0/16             10.20.21.3              2         probe-delegate  test7985
Active
10.39.2.0/16             10.20.21.4              2         probe-delegate  test7986
Active
10.39.3.0/16             10.20.21.5              2         probe-delegate  test7987
Active
10.39.4.0/16             10.20.21.6              2         probe-delegate  test7988
Active
```

show route rpm-tracking destination *destination-prefix* when preference and tag are configured

```
user@router> show route rpm-tracking destination 10.10.10.0/32

Routing table: inet.0
Destination        Next-Hop       Metric   preference   route-tag    Owner     Test Name   State
10.10.10.0/32      -              1        14           20           -         -
-
                   10.10.10.2     -        -            -            probe-1   test-1-1    Active
```

show route rpm-tracking destination *destination-prefix* owner *name* test *test-name*

```
user@host> show route rpm-tracking destination 10.39.0.0/16 owner probe-delegate test test7998

Destination              Next-Hop                Metric    Owner           Test Name       State
```

```
10.39.14.0/24            10.20.21.2               2         probe-delegate  test7998
Active        inet.0
```

show route rpm-tracking destination *destination-prefix* owner *name* test *test-name* when preference and tag are configured

```
user@router> show route rpm-tracking destination 10.10.10.0/32 owner probe-1 test test-1-1

Destination        Next-Hop       Metric   preference   route-tag    Owner     Test Name   State
10.10.10.0/32      -              1        14           20           -         -
-           inet.0
                   10.10.10.2     -        -            -            probe-1   test-1-1    Active
```

Release Information

Command introduced in Junos OS Release 18.4 R1.

Output showing multiple next hops added in Junos OS Release 19.1R1.

Output showing new `preference` and `route-tag` fields added in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

rpm-tracking | 16038

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

show route sla-tracking

IN THIS SECTION

- Syntax | 23765
- Description | 23765
- Options | 23765
- Additional Information | 23765

- Required Privilege Level | 23766
- Output Fields | 23766
- Sample Output | 23767
- Release Information | 23767

Syntax

```
show route sla-tracking
<destination IP-address>
<next-hop IP-address>
<sla (rpm | twamp)>
<state (active | inactive)>
```

Description

(Junos OS Evolved only) Display a brief summary of the state of SLA-tracked routes along with the current state for a given test. You configure SLA-tracking with the `sla-tracking` configuration statement at the `[edit routing-options]` hierarchy level.

Options

destination *IP-address* (Optional) Choose to see SLA-tracked routes for a particular IPv4 or IPv6 destination prefix.

sla (rpm | twamp) (Optional) Choose to see only RPM-tracked routes or only TWAMP-tracked routes.

next-hop *IP-address* (Optional) Next-hop address for a particular SLA-tracked route for a particular IPv4 or IPv6 destination prefix.

state (active | inactive) (Optional) Choose to see only active or only inactive SLA-tracked static routes.

Additional Information

Required Privilege Level

view

Output Fields

Table 1100 on page 23766 lists the output fields for the `show route sla-tracking` command.

Table 1100: show route sla-tracking Output Fields

Field Name	Field Description
Routing table	Name of the routing table
Destination	Destination prefix for the static route
Next-hop	IP address of the next hop
Metric	Route metric value for the destination prefix
Pref.	Route preference value for the destination prefix
Tag	Route tag value for the destination prefix
SLA	Protocol used for the SLA tracking: `rpm` or `twamp`
SLA test name	For an RPM test, a combination of the owner name and test name. For a TWAMP test, a combination of the connection name and test name.
State	Status of the test: • `inactive`: Either the test has not yet been run, or the test has failed. • `active`: The test was successful.

Sample Output

show route sla-tracking

```
user@host> show route sla-tracking
```

```
Routing table: IN.inet.0
Destination        Next-hop    Metric    Pref.    Tag    SLA    SLA test name        State
10.50.50.0/24                  -         -        -
                10.1.1.1                            rpm    owner1/test_ipv4       active
                10.2.2.2                            twamp     connection1/test_ipv4     inactive

Routing table: IN.inet6.0
Destination        Next-hop    Metric    Pref.    Tag    SLA    SLA test name        State
2001:db8:50::/64                  -         -        -
                2001:db8::1                         rpm    owner1/test_ipv6        active
                2001:db8::2                         twamp     connection2/test_ipv6      inactive
```

Release Information

Command introduced in Junos OS Evolved Release 24.4R1.

show route snooping

IN THIS SECTION

- Syntax | 23768
- Description | 23768
- Options | 23768
- Required Privilege Level | 23769
- Output Fields | 23769
- Sample Output | 23769

Release Information | 23771

Syntax

```
show route snooping
<brief | detail | extensive | terse>
<all>
<best address/prefix>
<exact address>
<logical-system logical-system-name>
<range prefix-range>
<summary>
<table table-name>
```

Description

Display the entries in the routing table that were learned from snooping.

Options

none	Display the entries in the routing table that were learned from snooping.
brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.
all	(Optional) Display all entries, including hidden entries.
best ***address/prefix***	(Optional) Display the longest match for the provided address and optional prefix.
exact ***address/prefix***	(Optional) Display exact matches for the provided address and optional prefix.
logical-system ***logical-system-name***	(Optional) Display information about a particular logical system, or type 'all'.
range ***prefix-range***	(Optional) Display information for the provided address range.

summary	(Optional) Display route snooping summary statisitics.
table ***table-name***	(Optional) Display information for the named table.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the "show route" on page 23512 command, the "show route detail" on page 23573 command, the "show route extensive" on page 23619 command, or the "show route terse" on page 23824 command.

Sample Output

show route snooping detail

```
 user@host> show route snooping detail
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
224.0.0.2/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next hop type: MultiRecv
                Next-hop reference count: 4
                State: <Active NoReadvrt Int>
                Age: 2:24
                Task: IGMP
                Announcement bits (1): 0-KRT
                AS path: I

224.0.0.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next hop type: MultiRecv
                Next-hop reference count: 4
                State: <Active NoReadvrt Int>
                Age: 2:24
                Task: IGMP
                Announcement bits (1): 0-KRT
                AS path: I
```

```
__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4), Next hop index: 1048584
                Next-hop reference count: 4
                State: <Active Int>
                Age: 2:24
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

<snip>
```

show route snooping logical-system all

```
 user@host> show route snooping logical-system all

logical-system: default

inet.1: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
Restart Unsupported
+ = Active Route, - = Last Active, * = Both

0.0,0.1,0.0,232.1.1.65,100.1.1.2/112*[Multicast/180] 00:07:36
                      Multicast (IPv4) Composite
0.0,0.1,0.0,232.1.1.66,100.1.1.2/112*[Multicast/180] 00:07:36
                      Multicast (IPv4) Composite
0.0,0.1,0.0,232.1.1.67,100.1.1.2/112*[Multicast/180] 00:07:36

<snip>

default-switch.inet.1: 237 dest, 237 rts (237 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

0.15,0.1,0.0,0.0.0.0,0.0.0.0,2/120*[Multicast/180] 00:08:21
                      Multicast (IPv4) Composite
0.15,0.1,0.0,0.0.0.0,0.0.0.0,2,17/128*[Multicast/180] 00:08:21
                      Multicast (IPv4) Composite
```

```
<snip>
```

Release Information

Command introduced in Junos OS Release 8.5.

show route source-gateway

IN THIS SECTION

- Syntax | 23771
- Syntax (EX Series Switches) | 23771
- Description | 23772
- Options | 23772
- Required Privilege Level | 23772
- Output Fields | 23772
- Sample Output | 23772
- Release Information | 23779

Syntax

```
show route source-gateway address
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route source-gateway address
<brief | detail | extensive | terse>
```

Description

Display the entries in the routing table that were learned from a particular address. The **Source** field in the `show route detail` command output lists the source for each route, if known.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.
address	IP address of the system.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the "show route" on page 23512 command, the "show route detail" on page 23573 command, the "show route extensive" on page 23619 command, or the "show route terse" on page 23824 command.

Sample Output

show route source-gateway

```
 user@host> show route source-gateway 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
```

```
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
                   *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
                   *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
                   *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
                   *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3
```

show route source-gateway detail

```
user@host> show route source-gateway 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0
```

```
red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-red-l2vpn
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down, mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Secondary Tables: green.l2vpn.0
```

```
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Secondary Tables: red.l2vpn.0
```

show route source-gateway extensive

```
user@host> show route source-gateway 10.255.70.103 extensive
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

```
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-red-l2vpn
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down, mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0
```

```
bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Secondary Tables: green.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.255.70.103 Metric: 2
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 1
                Next hop:        via so-0/3/0.0 weight 0x1
                        10.255.70.103/32 Originating RIB: inet.3
                          Metric: 2                       Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: via so-0/3/0.0

10.255.70.103:2:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
```

```
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Secondary Tables: red.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.255.70.103 Metric: 2
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 1
                Next hop:        via so-0/3/0.0 weight 0x1
                        10.255.70.103/32 Originating RIB: inet.3
                          Metric: 2                       Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: via so-0/3/0.0
```

Release Information

Command introduced before Junos OS Release 7.4.

show route summary

IN THIS SECTION

- Syntax | 23780
- Syntax (EX Series Switches) | 23780
- Operational mode | 23780
- Description | 23780
- Options | 23780
- Required Privilege Level | 23781
- Output Fields | 23781
- Sample Output | 23784
- Release Information | 23789

Syntax

```
show route summary
<logical-system (all | logical-system-name)>
<table routing-table-name>
<rib-sharding (main | rib-shard-name)>
```

Syntax (EX Series Switches)

```
show route summary
```

Operational mode

`cli` mode

`cli-pfe` mode.

Description

Display summary statistics about the entries in the routing table. Additionally, displays summary of FIB compression statistics in `cli-pfe` mode. (PTX and ACX series devices).

CPU utilization might increase while the device learns routes. We recommend that you use the `show route summary` command after the device learns and enters the routes into the routing table. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the `show route summary` command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the command-line interface (CLI).

Options

none	Display summary statistics about the entries in the routing table.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
table *routing-table-name*	(Optional) Display summary statistics for all routing tables whose name begins with this string (for example, **inet.0** and **inet6.0** are both displayed when you run the `show route summary table inet` command). If you only want to display statistics

	for a specific routing table, make sure to enter the exact name of that routing table.
rib-sharding (main | *rib-shard-name*)	(Optional) Display name of the rib shard.

Required Privilege Level

view

Output Fields

Table 1101 on page 23781 lists the output fields for the `show route summary` command. Output fields are listed in the approximate order in which they appear.

Table 1101: show route summary Output Fields

Field Name	Field Description
Router ID	Address of the local routing device.
Highwater Mark	High-water mark data of routing and forwarding (RIB/FIB) table routes and VRFs in a system. Shows the following details: • RIB Unique destination routes • RIB routes • FIB routes • VRF type routing instances The high-water mark data can also be viewed in the syslog at the **LOG_NOTICE** level.

Table 1101: show route summary Output Fields *(Continued)*

Field Name	Field Description
routing-table-name	Name of the routing table example: • **inet.0**— List all routes for IPv4. • **junos-main::inet.0**—Lists all active routes for Ipv4. • **junos-bgpshard0::inet.0**—Lists IPv4 routes present in bgpshard0 thread. • **__raass__inet.inet.0**—Lists IPv4 routes stored in RaaS Server that are received from the clients. • **__raasc__inet.inet.0**—Lists IPv4 routes stored in RaaS Clients that are used for local resolution. • **inet.6**— List all routes for IPv6. • **junos-main::inet6.0**— Lists all active routes for IPv6. • **junos-bgpshard0::inet6.0**—Lists IPv6 routes present in bgpshard0 thread. • **__raass__inet.inet6.0**—Lists IPv6 routes stored in RaaS Server that are received from the clients. • **__raasc__inet.inet6.0**—Lists IPv6 routes stored in RaaS Clients that are used for local resolution.
destinations	Number of destinations for which there are routes in the routing table.
routes	Number of routes in the routing table: • **active**—Number of routes that are active. • **holddown**—Number of routes that are in the hold-down state before being declared inactive. • **hidden**—Number of routes that are not used because of routing policy.

Table 1101: show route summary Output Fields *(Continued)*

Field Name	Field Description

Restart complete	All protocols have restarted for this routing table.
	
	Restart state:
	
	• `Pending:`*`protocol-name`*—List of protocols that have not yet completed graceful restart for this routing table.
	
	• `Complete`—All protocols have restarted for this routing table.
	
	For example, if the output shows-
	
	• `LDP.inet.0: 5 routes (4 active, 1 holddown, 0 hidden)`
	`Restart Pending: OSPF LDP VPN`
	
	This indicates that `OSPF`, `LDP`, and `VPN` protocols did not restart for `LDP.inet.0` routing table.
	
	• `vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)`
	`Restart Complete`
	
	This indicates that all protocols have restarted for `vpls_1.l2vpn.0` routing table.

Limit/Threshold	Displays the configured route limits for the routing table set with the **maximum-prefixes** and the `maximum-paths` statements. If you do not configure route limits for the routing table, the show output does not display this information. • **destinations**—The first number represents the maximum number of route prefixes installed in the routing table. The second number represents the number of route prefixes that trigger a warning message. • **routes**—The first number represents the maximum number of routes. The second number represents the number of routes that trigger a warning message.
Direct	Routes on the directly connected network.
Local	Local routes.

Table 1101: show route summary Output Fields *(Continued)*

Field Name	Field Description
protocol-name	Name of the protocol from which the route was learned. For example, **OSPF**, **RSVP**, and **Static**.
Output fields for CLI-PFE mode (ACX7100-32C, ACX7100-48L, and ACX7509 Devices) that display route table statistics at the PFE (for example, IPv4, IPv6, MPLS Route Tables)	
Index	Routing table index number.
Routes	Number of active routes
Size(b)	Size of the packets #chk w SME.
Prefixes	Actual number of routes published by the peer devices.
Aggr	Number of aggregated routes generated by the PFE after compression that includes valid next hops
Installed	Number of routes installed on the ACX device.
Comp(%)	Route table compression percentage.

Sample Output

show route summary

```
user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52

Highwater Mark (All time / Time averaged watermark)
    RIB Unique destination routes: 51 at 2020-06-15 13:54:01 / 51
    RIB routes                   : 52 at 2020-06-15 13:54:01 / 52
    FIB routes                   : 33 at 2020-06-15 13:54:01 / 33
    VRF type routing instances   : 0 at 2020-06-15 13:54:00
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
```

```
Restart Complete
              Direct:      6 routes,      5 active
               Local:      4 routes,      4 active
                OSPF:      5 routes,      4 active
              Static:      7 routes,      7 active
                IGMP:      1 routes,      1 active
                 PIM:      2 routes,      2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
                RSVP:      2 routes,      2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
              Direct:      1 routes,      1 active

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
                MPLS:      3 routes,      3 active
                VPLS:      4 routes,      2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
              Direct:      2 routes,      2 active
                 PIM:      2 routes,      2 active
                 MLD:      1 routes,      1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      2 routes,      2 active
               L2VPN:      2 routes,      2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      2 routes,      2 active
               L2VPN:      1 routes,      1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      4 routes,      4 active
```

show route summary (cli-pfe mode in ACX7100-32C, ACX7100-48L, and ACX7509 Devices)

The following command shows the route summary information in `cli-pfe` mode only:

```
user@host:pfe> show route summary

IPv4 Route Tables:
Index         Routes     Size(b)  Prefixes     Aggr     Installed   Comp(%)
--------  ----------  ----------  ---------  ---------  ----------  ------
Default     10095472  1569927632   5000016    4999986       11525    100
1                  0           0         0          0           0      0
51                10        1040         5          0           5      0
36738             10        1040         5          0           5      0

IPv6 Route Tables:
Index         Routes     Size(b)  Prefixes     Aggr     Installed   Comp(%)
--------  ----------  ----------  ---------  ---------  ----------  ------
0       IPv6       1300025       1299963          62      100

1       IPv6             0             0           0        0

51      IPv6             6             0           6        0

36738   IPv6             6             0           6        0
MPLS Route Tables:
Index         Routes     Size(b)  Prefixes     Aggr     Installed   Comp(%)
--------  ----------  ----------  ---------  ---------  ----------  ------
Default           14        1456         7          0           7      -
52                 2         208         1          0           1      -

IPv6 Route Tables:
Index         Routes     Size(b)  Prefixes     Aggr     Installed   Comp(%)
--------  ----------  ----------  ---------  ---------  ----------  ------
Default           27        2808        13          0          16  4294967273
1                  0           0         0          0           0      0
51                12        1248         6          0           6      0
53                12        1248         6          0           6      0
36738             12        1248         6          0           6      0

CLNP Route Tables:
Index         Routes     Size(b)  Prefixes     Aggr     Installed   Comp(%)
--------  ----------  ----------  ---------  ---------  ----------  ------
```

```
Default            4         416         2          0          2      -
51                 2         208         1          0          1      -
```

show route summary table (with rib-sharding configured)

The following command shows the route summary information in junos-bgpshard0 thread only:

```
user@host> show route summary rib-sharding junos-bgpshard0
Autonomous system number: 1
Router ID: 2.2.2.1

inet.0: 774078 destinations, 1547602 routes (774078 active, 0 holddown, 0 hidden)
              Direct:    108 routes,    108 active
               Local:      7 routes,      7 active
                OSPF:    408 routes,    408 active
                 BGP: 1547048 routes, 773524 active
              Static:     31 routes,     31 active

junos-main::inet.0: 774078 destinations, 774078 routes (774078 active, 0 holddown, 0 hidden)
              Direct:    108 routes,    108 active
               Local:      7 routes,      7 active
                OSPF:    408 routes,    408 active
                 BGP: 773524 routes, 773524 active
              Static:     31 routes,     31 active

junos-bgpshard0::inet.0: 258448 destinations, 516635 routes (258448 active, 0 holddown, 0 hidden)
                 BGP: 516374 routes, 258187 active

junos-bgpshard1::inet.0: 258129 destinations, 515997 routes (258129 active, 0 holddown, 0 hidden)
                 BGP: 515736 routes, 257868 active

junos-bgpshard2::inet.0: 257731 destinations, 515200 routes (257731 active, 0 holddown, 0 hidden)
                 BGP: 514938 routes, 257469 active

__raass__inet.inet.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
                RaaS:      1 routes,      1 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Direct:      1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
              Direct:      2 routes,      2 active
               INET6:      1 routes,      1 active

junos-main::inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
              Direct:      2 routes,      2 active
               INET6:      1 routes,      1 active
```

show route summary table

```
user@host> show route summary table inet
Router ID: 192.168.0.1

inet.0: 32 destinations, 34 routes (31 active, 0 holddown, 1 hidden)
              Direct:      6 routes,      5 active
               Local:      9 routes,      9 active
                OSPF:      3 routes,      1 active
              Static:     13 routes,     13 active
                IGMP:      1 routes,      1 active
                 PIM:      2 routes,      2 active

inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
           Multicast:      1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
               Local:      1 routes,      1 active
                 PIM:      2 routes,      2 active

inet6.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
           Multicast:      1 routes,      1 active
```

show route summary table (with Route Limits Configured for the Routing Table)

```
user@host> show route summary table VPN-A.inet.0
Autonomous system number: 100
Router ID: 10.255.182.142

VPN-A.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
Limit/Threshold: 2000/200 destinations 20/12 routes
              Direct:      2 routes,      2 active
               Local:      1 routes,      1 active
```

```
                 OSPF:      4 routes,      3 active
                  BGP:      4 routes,      4 active
                 IGMP:      1 routes,      1 active
                  PIM:      2 routes,      2 active
```

show route summary rib-sharding

```
user@host> show route summary rib-sharding junos-bgpshard14
Autonomous system number: 100
Router ID: 20.255.255.10

inet.0: 54 destinations, 54 routes (54 active, 0 holddown, 0 hidden)
              Direct:     29 routes,     29 active
               Local:     25 routes,     25 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Direct:      1 routes,      1 active

inet6.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
               Local:     14 routes,     14 active
```

Release Information

Command introduced before Junos OS Release 7.4.

`rib-sharding` option introduced for cRPD Release 20.1R1.

show route table

IN THIS SECTION

- Syntax | 23790
- Syntax (EX Series Switches, QFX Series Switches) | 23790
- Description | 23790
- Options | 23790

- Required Privilege Level | 23790
- Output Fields | 23791
- Sample Output | 23807
- Release Information | 23821

Syntax

```
show route table routing-table-name
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches, QFX Series Switches)

```
show route table routing-table-name
<brief | detail | extensive | terse>
```

Description

Display the route entries in a particular routing table.

Options

brief | detail | extensive | terse	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS.
routing-table-name	Display route entries for all routing tables whose names begin with this string (for example, inet.0 and inet6.0 are both displayed when you run the `show route table inet` command).

Required Privilege Level

view

Output Fields

Table 1102 on page 23791 describes the output fields for the `show route table` command. Output fields are listed in the approximate order in which they appear.

Table 1102: show route table Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).

Restart complete	All protocols have restarted for this routing table.
	
	Restart state:
	
	• `Pending:`*`protocol-name`*—List of protocols that have not yet completed graceful restart for this routing table.
	
	• `Complete`—All protocols have restarted for this routing table.
	
	For example, if the output shows-
	
	• `LDP.inet.0 : 5 routes (4 active, 1 holddown, 0 hidden)`
	`Restart Pending: OSPF LDP VPN`
	
	This indicates that `OSPF`, `LDP`, and `VPN` protocols did not restart for the `LDP.inet.0` routing table.
	
	• `vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)`
	`Restart Complete`
	
	This indicates that all protocols have restarted for the `vpls_1.l2vpn.0` routing table.

number destinations	Number of destinations for which there are routes in the routing table.
number routes	Number of routes in the routing table and total number of routes in the following states: • `active` (routes that are active) • `holddown` (routes that are in the pending state before being declared inactive) • `hidden` (routes that are not used because of a routing policy)

Table 1102: show route table Output Fields *(Continued)*

Field Name	Field Description

route-destination (entry, announced)	Route destination (for example:10.0.0.1/24). The `entry` value is the number of routes for this destination, and the `announced` value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• *MPLS-label* (for example, 80001).
	• *interface-name* (for example, ge-1/0/2).
	• *neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source* (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  • *neighbor-address*—Address of the neighbor.
	  • *control-word-status*—Whether the use of the control word has been negotiated for this virtual circuit: `NoCtrlWord` or `CtrlWord`.
	  • *encapsulation type*—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  • `vc-id`—Virtual circuit identifier.
	  • *source*—Source of the advertisement: `Local` or `Remote`.
	• *`inclusive multicast Ethernet tag route`*—Type of route destination represented by (for example, 3:100.100.100.10:100::0::10::100.100.100.10/384):
	  • *route distinguisher*—(8 octets) Route distinguisher (RD) must be the RD of the EVPN instance (EVI) that is advertising the NLRI.
	  • *Ethernet tag ID*—(4 octets) Identifier of the Ethernet tag. Can set to 0 or to a valid Ethernet tag value.
	  • *IP address length*—(1 octet) Length of IP address in bits.
	  • *originating router's IP address*—(4 or 16 octets) Must set to the provider edge (PE) device's IP address. This address should be common for all EVIs on the PE device, and may be the PE device's loopback address.

Table 1102: show route table Output Fields *(Continued)*

Field Name	Field Description
label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • `S=0 route` indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no `S=` information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • `+`—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • `-`—A hyphen indicates the last active route. • `*`—An asterisk indicates that the route is both the active and the last active route. An asterisk before a `to` line indicates the best subpath to the route. In every routing metric except for the BGP `LocalPref` attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the `LocalPref` value in the `Preference2` field. For example, if the `LocalPref` value for Route 1 is 100, the `Preference2` value is -101. If the `LocalPref` value for Route 2 is 155, the `Preference2` value is -156. Route 2 is preferred because it has a higher `LocalPref` value and a lower `Preference2` value.
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).

Table 1102: show route table Output Fields *(Continued)*

Field Name	Field Description
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 1103 on page 23799.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word `Selected`. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be `pop` (where a label is removed from the top of the stack), `push` (where another label is added to the label stack), or `swap` (where a label is replaced by another label).

Table 1102: show route table Output Fields *(Continued)*

Field Name	Field Description
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 1104 on page 23801.
Local AS	AS number of the local routing devices.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.
Task	Name of the protocol that has added the route.

Table 1102: show route table Output Fields *(Continued)*

Field Name	Field Description
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the kernel routing table (KRT) for installing the route into the Packet Forwarding Engine, to a resolve tree, a Layer 2 VC, or even a VPN. For example, *n*-`Resolve inet` indicates that the specified route is used for route resolution for next hops found in the routing table. • *n*—An index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • `I`—IGP. • `E`—EGP. • `Recorded`—The AS path is recorded by the sample process (sampled). • `?`—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • `[ ]`—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • `[ ]`—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • `{ }`—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • `( )`—Parentheses enclose a confederation. • `( [ ] )`—Parentheses and brackets enclose a confederation set. **NOTE**: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.

Table 1102: show route table Output Fields *(Continued)*

Field Name	Field Description
validation-state	(BGP-learned routes) Validation status of the route: • `Invalid`—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • `Unknown`—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • `Unverified`—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • `Valid`—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Indicates point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, indicates the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, indicates the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.

Table 1102: show route table Output Fields *(Continued)*

Field Name	Field Description
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 1105 on page 23805 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: `none` or `Site Down`.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.

Table 1102: show route table Output Fields *(Continued)*

Field Name	Field Description
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 1103 on page 23799 describes all possible values for the Next-hop Types output field.

Table 1103: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.

Table 1103: Next-hop Types Output Field Values *(Continued)*

Next-Hop Type	Description
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.

Table 1103: Next-hop Types Output Field Values *(Continued)*

Next-Hop Type	Description
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	A specific node or set of nodes to which the routing device forwards packets that match the route prefix. To qualify as a next-hop type router, the route must meet the following criteria: • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 1104 on page 23801 describes all possible values for the State output field. A route can be in more than one state (for example, `<Active NoReadvrt Int Ext>`).

Table 1104: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.

Table 1104: State Output Field Values *(Continued)*

Value	Description
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.

Table 1104: State Output Field Values *(Continued)*

Value	Description
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.

Table 1104: State Output Field Values *(Continued)*

Value	Description
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available.
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.
VxlanLocalRT	Route is an EVPN Type 5 route (IP prefix route).

Table 1105 on page 23805 describes the possible values for the Communities output field.

Table 1105: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: *local AS number:link-bandwidth-number*	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
DF election	Designated forwarder (DF) election mode for multihoming configurations. Available options on a device might include mod-based or preference-based DF election.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
evpn-mcast-flags	Identifies the value in the multicast flags extended community, which includes bits that indicate EVPN multicast snooping and multicast optimized intersubnet multicast (OISM) information. For example: • 0x01—IGMP snooping bit (display includes keyword `igmp-snooping-enabled`) • 0x02—MLD snooping bit (display includes keyword `mld-snooping-enabled`) • 0x08—OISM bit (display includes keyword `oism`) • 0x10—PIM EVPN gateway (PEG) bit (display includes keyword `peg`) • 0x100—Supplemental bridge domain (SBD) bit (display includes keyword `sbd`)
evpn-l2-info	Identifies whether Multihomed Proxy MAC and IP Address Route Advertisement is enabled. A value of 0x20 indicates that the proxy bit is set. . Use the `show bridge mac-ip-table extensive` statement to determine whether the MAC and IP address route was learned locally or from a PE device.
GBP Tag	Assigned Group Based Policy (GBP) tag from 1 through 65535.

Table 1105: Communities Output Field Values *(Continued)*

Value	Description
link-bandwidth-number	Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).
local AS number	Local AS number: from 1 through 65,535.
options	1 byte. Currently this is only used if the route type is 5 or 7. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be0); 7 for NSSA routes; or 129 for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000. The format is *area-number:ospf-route-type:options*.
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306. The format is *area-number:ospf-route-type:options*.
target	Defines which VPN the route participates in; target has the format *32-bit IP address:16-bit number*. For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between 0x1 and 0x7fff. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above 0x8000. This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route table bgp.l2vpn.0

```
user@host> show route table bgp.l2vpn.0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
                   *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table evpn-vs2.evpn.0 extensive

```
user@host> show route table evpn-vs2.evpn.0 extensive
evpn-vs2.evpn.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
Restart Complete
2:10.3.3.3:2::601: :2c:6b:f5:f2:87:f0/304 MAC/IP (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.3.3.3:2
                Next hop type: Indirect, Next hop index: 0
                Address: 0x76b66dc
                Next-hop reference count: 32
                Source: 10.2.2.2
                Protocol next hop: 10.3.3.3
                Indirect next hop: 0x2 no-forward INH Session ID: 0
                State: <Secondary Active Ext>
                Local AS:   102 Peer AS:   201
                Age: 22:56:12     Metric2: 1
                Validation State: unverified
                Task: BGP_201.10.2.2.2
                Announcement bits (1): 0-evpn-vs2-evpn
                AS path: 201 203 I
                Communities: target:5:5 encapsulation:vxlan(0x8) evpn-default-gateway gbp-tag:
200
```

show route table bgp.evpn.0 match-prefix

```
user@hostshow route table bgp.evpn.0 match-prefix *172.23.226.*
bgp.evpn.0: 1343 destinations, 1343 routes (1343 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:172.23.226.134:0::052261342261350010::FFFF:FFFF/192 AD/ESI
                   *[BGP/170] 26w4d 14:18:02, localpref 100, from 172.23.226.134
                      AS path: I, validation-state: unverified
                    >  to 172.23.226.24 via ae1.0
1:172.23.226.134:0::052261342261350020::FFFF:FFFF/192 AD/ESI
                   *[BGP/170] 26w4d 14:18:02, localpref 100, from 172.23.226.134
                      AS path: I, validation-state: unverified
                    >  to 172.23.226.24 via ae1.0
1:172.23.226.134:0::052261342261350022::FFFF:FFFF/192 AD/ESI
                   *[BGP/170] 26w4d 14:18:02, localpref 100, from 172.23.226.134
                      AS path: I, validation-state: unverified
                    >  to 172.23.226.24 via ae1.0
<snip>
2:172.23.226.131:9999::800::b4:7a:f1:58:05:46/304 MAC/IP
                   *[BGP/170] 19w0d 07:48:24, localpref 100, from 172.23.226.131
                      AS path: I, validation-state: unverified
                    >  to 172.23.226.7 via et-0/1/2.0
                    [BGP/170] 17w1d 23:19:37, localpref 100, from 172.23.226.129
                      AS path: I, validation-state: unverified
                    >  to 172.23.226.7 via et-0/1/2.0
2:172.23.226.131:9999::800::b4:7a:f1:58:05:48/304 MAC/IP
                   *[BGP/170] 19w0d 07:48:25, localpref 100, from 172.23.226.131
                      AS path: I, validation-state: unverified
                    >  to 172.23.226.7 via et-0/1/2.0
                    [BGP/170] 17w1d 23:19:38, localpref 100, from 172.23.226.129
                      AS path: I, validation-state: unverified
                    >  to 172.23.226.7 via et-0/1/2.0
```

show route table bgpmcast.0 extensive

```
user@host> show route table bgpmcast.0 extensive
bgpmcast.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
4:3:0:0:32:10.2.20.20:32:225.1.1.1:10.4.4.4:10.20.20.20/240 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Tunnel type: AnyEncap, RPF tunnel:, Remote end point: 10.1.4.4
                Tunnel type: AnyEncap, Remote end point: 10.2.4.4
                Tunnel type: AnyEncap, Remote end point: 10.3.4.4
                Tunnel type: AnyEncap, Remote end point: 10.4.6.4
                Next hop type: Indirect, Next hop index: 0
                Address: 0xc54639c
                Next-hop reference count: 1
                Source: 10.1.1.1
                Protocol next hop: 10.20.20.20
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                State: <Active Int Ext>
                Local AS:   65100 Peer AS:   65100
                Age: 30         Metric2: 0
                Validation State: unverified
                Task: BGP_100.1.1.1.1
                Announcement bits (1): 0-bgpmcast global task
                AS path: I
                Communities: target:10.4.4.4:0
                Accepted
                Localpref: 100
                Router ID: 10.1.1.1
                Indirect next hops: 1
                        Protocol next hop: 10.20.20.20
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.4.20.20 via ge-0/0/5.0
                                Session Id: 0x140
                                10.20.20.20/32 Originating RIB: inet.0
                                  Node path count: 1
                                  Forwarding nexthops: 1
                                        Next hop type: Router
                                        Next hop: 10.4.20.20 via ge-0/0/5.0
                                        Session Id: 0x140
```

show route table inet.0

```
user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:51:57
                    > to 172.16.5.254 via fxp0.0
10.0.0.1/32         *[Direct/0] 00:51:58
                    > via at-5/3/0.0
10.0.0.2/32         *[Local/0] 00:51:58
                     Local
10.12.12.21/32     *[Local/0] 00:51:57
                     Reject
10.13.13.13/32     *[Direct/0] 00:51:58
                    > via t3-5/2/1.0
10.13.13.14/32     *[Local/0] 00:51:58
                     Local
10.13.13.21/32     *[Local/0] 00:51:58
                     Local
10.13.13.22/32     *[Direct/0] 00:33:59
                    > via t3-5/2/0.0
127.0.0.1/32        [Direct/0] 00:51:58
                    > via lo0.0
10.222.5.0/24     *[Direct/0] 00:51:58
                    > via fxp0.0
10.222.5.81/32    *[Local/0] 00:51:58
                     Local
```

show route table inet.3

```
user@host> show route table inet.3
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32        *[LDP/9] 00:25:43, metric 10, tag 200
                      to 10.2.94.2 via lt-1/2/0.49
                    > to 10.2.3.2 via lt-1/2/0.23
```

show route table inet.3 protocol ospf

```
user@host> show route table inet.3 protocol ospf
inet.3: 9 destinations, 18 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.20/32         [L-OSPF/10] 1d 00:00:56, metric 2
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800020
                      to 10.0.6.60 via lt-1/2/0.12, Push 800020, Push 800030(top)
1.1.1.30/32         [L-OSPF/10] 1d 00:01:01, metric 3
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800030
                      to 10.0.6.60 via lt-1/2/0.12, Push 800030
1.1.1.40/32         [L-OSPF/10] 1d 00:01:01, metric 4
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800040
                      to 10.0.6.60 via lt-1/2/0.12, Push 800040
1.1.1.50/32         [L-OSPF/10] 1d 00:01:01, metric 5
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800050
                      to 10.0.6.60 via lt-1/2/0.12, Push 800050
1.1.1.60/32         [L-OSPF/10] 1d 00:01:01, metric 6
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800060
                      to 10.0.6.60 via lt-1/2/0.12, Pop
```

show route table inet6.0

```
user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0
```

show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                   *[LDP/9] 00:00:22, metric 1
                    > via so-1/0/0.0
::10.255.245.196/128
                   *[LDP/9] 00:00:08, metric 1
                    > via so-1/0/0.0, Push 100008
```

show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
10.1.1.195:CtrlWord:1:2:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
```

show route table lsdist.0

```
user@host> show route table lsdist.0

lsdist.0: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both

NODE { AS:100 ISO:0003.0003.0303.00 ISIS-L2:0 }/1216
                   *[IS-IS/18] 10:19:16
                       Fictitious
NODE { AS:100 ISO:1282.0404.9202.00 ISIS-L2:0 }/1216
                   *[IS-IS/18] 10:19:07
                       Fictitious
NODE { AS:100 ISO:1282.0404.9202.02 ISIS-L2:0 }/1216
                   *[IS-IS/18] 10:19:07
                       Fictitious
LINK { Local { AS:100 ISO:0003.0003.0303.00 }.{ IPv4:23.0.0.2 IPv6:23::2 } Remote { AS:100
ISO:1282.0404.9202.02 }.{ } ISIS-L2:0 }/1216
                   *[IS-IS/18] 10:19:07
                       Fictitious
LINK { Local { AS:100 ISO:1282.0404.9202.00 }.{ IPv4:23.0.0.1 IPv6:23::1 } Remote { AS:100
ISO:1282.0404.9202.02 }.{ } ISIS-L2:0 }/1216
                   *[IS-IS/18] 10:19:07
                       Fictitious
LINK { Local { AS:100 ISO:1282.0404.9202.02 }.{ } Remote { AS:100 ISO:0003.0003.0303.00 }.{ }
ISIS-L2:0 }/1216
                   *[IS-IS/18] 10:19:07
                       Fictitious
LINK { Local { AS:100 ISO:1282.0404.9202.02 }.{ } Remote { AS:100 ISO:1282.0404.9202.00 }.{ }
ISIS-L2:0 }/1216
                   *[IS-IS/18] 10:19:07
                       Fictitious
PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv4:3.3.3.3/32 } ISIS-L2:0 }/1216
                   *[IS-IS/18] 10:19:16
                       Fictitious
PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv4:32.32.32.32/32 } ISIS-L2:0 }/
1216
                   *[IS-IS/18] 10:19:16
                       Fictitious
PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv4:33.33.33.33/32 } ISIS-L2:0 }/
1216
                   *[IS-IS/18] 10:19:16
                       Fictitious
PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv4:100.100.100.0/32 } ISIS-L2:0 }/
1216
                   *[IS-IS/18] 10:19:16
                       Fictitious
PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv4:128.204.49.193/32 } ISIS-L2:0 }/
```

```
1216
                   *[IS-IS/18] 10:19:16
                       Fictitious
PREFIX { Node { AS:100 ISO:1282.0404.9202.00 } { IPv4:2.2.2.2/32 } ISIS-L2:0 }/1216
                   *[IS-IS/18] 10:19:07
                       Fictitious
PREFIX { Node { AS:100 ISO:1282.0404.9202.00 } { IPv4:128.204.49.202/32 } ISIS-L2:0 }/
1216
                   *[IS-IS/18] 10:19:07
                       Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv6:3::3/128 } ISIS-L2:0 }/
1216
                   *[IS-IS/18] 10:19:16
                       Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv6:23::/64 } ISIS-L2:0 }/
1216
                   *[IS-IS/18] 10:19:16
                       Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv6:100:100:100::100/128 } ISIS-L2:0 }/
1216
                   *[IS-IS/18] 10:19:16
                       Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv6:abcd::128:204:49:193/128 } ISIS-
L2:0 }/1216
                   *[IS-IS/18] 10:19:16
                       Fictitious
IPV6 PREFIX { Node { AS:100 ISO:1282.0404.9202.00 } { IPv6:2::2/128 } ISIS-L2:0 }/
1216
                   *[IS-IS/18] 10:19:07
                       Fictitious
IPV6 PREFIX { Node { AS:100 ISO:1282.0404.9202.00 } { IPv6:23::/64 } ISIS-L2:0 }/
1216
                   *[IS-IS/18] 10:19:07
                       Fictitious
IPV6 PREFIX { Node { AS:100 ISO:1282.0404.9202.00 } { IPv6:abcd::128:204:49:202/128 } ISIS-
L2:0 }/1216
                   *[IS-IS/18] 10:19:07
                       Fictitious
```

show route table lsdist.0 (BGP Confederation Enabled)

```
user@host>show route table lsdist.0 detail
lsdist.0: 72 destinations, 72 routes (72 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

NODE { AS:17 Member-ASN:65002 ISO:0004.0404.0400.00 ISIS-L2:0 }/1216
                   *[BGP/170] 00:00:23, localpref 100
                     AS path: (65002) I, validation-state: unverified
                    >  to 11.14.1.2 via ge-0/0/0.0
NODE { AS:17 Member-ASN:65002 ISO:0005.0505.0500.00 ISIS-L2:0 }/1216
                   *[BGP/170] 00:00:23, localpref 100
                      AS path: (65002) I, validation-state: unverified
                    >  to 11.14.1.2 via ge-0/0/0.0

LINK { Local { AS:17 Member-ASN:65001 ISO:0003.0303.0300.00 }.{ IPv4:13.14.1.1
IPv6:abcd::13:14:1:1 } Remote { AS:17 Member-ASN:65001 ISO:0004.0404.0400.00 }.
{ IPv4:13.14.1.2 } ISIS-L2:0 }/1216
                   *[IS-IS/18] 00:00:48
                       Fictitious
LINK { Local { AS:17 Member-ASN:65001 ISO:0003.0303.0300.00 }.{ IPv4:13.14.2.1
IPv6:abcd::13:14:2:1 } Remote { AS:17 Member-ASN:65001 ISO:0004.0404.0400.00 }.
{ IPv4:13.14.2.2 } ISIS-L2:0 }/1216
                   *[IS-IS/18] 00:00:48
                       Fictitious
```

show route table lsdist.0 detail

```
user@host> show route table lsdist.0 detail
lsdist.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
NODE { AS:200 ISO:1282.2113.1154.00 ISIS-L1:0 }/1216 (1 entry, 1 announced)
*IS-IS  Preference: 15
Level: 1
Next hop type: Fictitious, Next hop index: 0
Address: 0xc5b3054
Next-hop reference count: 14
……….. .
……….. .
……….
```

```
Area membership:
47 00 05 80 ff f8 00 00 00 01 08 00 01
SPRING-Capabilities:
- SRGB block [Start: 800000, Range: 4096, Flags: 0xc0]
SPRING-Algorithms:
- Algo: 0
SPRING Flex-Algorithms Definition:
- Flex-Algo: 129
Metric: 0, Calc: 0, priority: 129
- Flags: 0x02, - Inc Any: 0x00040000, - Exclude: 0x00008000, - Inc All: 0x00004000
…….
…….
…….
PREFIX { Node { AS:200 ISO:1282.2113.3158.00 } { IPv4:128.220.13.196/32 } ISIS-L1:0 }/1216 (1
entry, 1 announced)
*IS-IS  Preference: 15
Level: 1
Next hop type: Fictitious, Next hop index: 0
Address: 0xc5b3054
Next-hop reference count: 14
Next hop:
State: <Active NotInstall>
Local AS:   200
Age: 16:16:25
…….
…….
Prefix SID: 10, Flags: 0xe0, Algo: 0
Prefix SID: 780, Flags: 0xe0, Algo: 129
Flex Algo: 129, Flex Algo Metric: 10
…….
…….
IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:4000::/64 } ISIS-L1:0 }/1216 (1
entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 588
                Address: 0x7660f64
                Next-hop reference count: 12, key opaque handle: 0x0, non-key opaque handle: 0x0
                Source: 21.0.2.1
                Next hop: 21.0.2.1 via ge-0/0/0.0, selected
                Session Id: 320
                State: <Active Ext>
                Local AS:   200 Peer AS:   100
                Age: 1d 10:20:38
```

```
                Validation State: unverified
                Task: BGP_100.21.0.2.1
                Announcement bits (1): 0-TED Export
                AS path: 100 I
                Accepted
                SRv6 Locator Flags: 0, Algo: 0, Metric: 0
                Localpref: 100
                Router ID: 100.100.100.0
                Thread: junos-main
```

show route table lsdist.0 detail

```
user@host> show route table lsdist.1 detail
SPRING Flex-Algorithms Definition:
                 - Flex-Algo: 128
                   Metric: 2, Calc: 1, Priority: 10
                    - Flags: 0X80000000, - Inc Any: 0x00000002, - Exclude: 0x00000004, - Inc
All: 0x00000002
```

show route table lsdist.0 extensive

The output of the `show route table lsdist.0 extensive` command to include IPv6 Prefix attributes.

```
user@host> show route table lsdist.0 extensive
lsdist.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
NODE { AS:100 ISO:0100.0100.0100.00 ISIS-L1:0 }/1216 (1 entry, 1 announced)
TSI:
LINK-STATE attribute handle 0x75ace70
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0x7661124
                Next-hop reference count: 5, key opaque handle: 0x0, non-key opaque handle: 0x0
                Source: 21.0.2.1
                Next hop: 21.0.2.1 via ge-0/0/0.0, selected
                Session Id: 0
                State: <Active Ext>
                Local AS:   200 Peer AS:   100
                Age: 17
                Validation State: unverified
```

```
                Task: BGP_100.21.0.2.1
                Announcement bits (1): 0-TED Export
                AS path: 100 I
                Accepted
                IPv4 Router-ids:
                  100.100.100.0
                Area border router: No
                External router: No
                Attached: No
                Overload: No
                Hostname: R0
                Area membership:
                  49 00 05
                SPRING-Algorithms:
                 - Algo: 0
                 - Algo: 1
                SRV6 Capable: - Flags: 0
                SRV6 Node MSD:
                 - Type: 41, Value: 6
                 - Type: 42, Value: 7
                 - Type: 43, Value: 5
                 - Type: 44, Value: 6
                 - Type: 45, Value: 6
                Localpref: 100
                Router ID: 100.100.100.0
                Thread: junos-main

PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv4:100.100.100.0/32 } ISIS-L1:0 }/1216 (1
entry, 1 announced)
TSI:
LINK-STATE attribute handle 0x0
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0x7661124
                Next-hop reference count: 5, key opaque handle: 0x0, non-key opaque handle: 0x0
                Source: 21.0.2.1
                Next hop: 21.0.2.1 via ge-0/0/0.0, selected
                Session Id: 0
                State: <Active Ext>
                Local AS:   200 Peer AS:   100
                Age: 28
                Validation State: unverified
                Task: BGP_100.21.0.2.1
```

```
                Announcement bits (1): 0-TED Export
                AS path: 100 I
                Accepted
                Localpref: 100
                Router ID: 100.100.100.0
                Thread: junos-main

IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:10::10/128 } ISIS-L1:0 }/1216 (1
entry, 1 announced)
TSI:
LINK-STATE attribute handle 0x0
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0x7661124
                Next-hop reference count: 5, key opaque handle: 0x0, non-key opaque handle: 0x0
                Source: 21.0.2.1
                Next hop: 21.0.2.1 via ge-0/0/0.0, selected
                Session Id: 0
                State: <Active Ext>
                Local AS:   200 Peer AS:   100
                Age: 28
                Validation State: unverified
                Task: BGP_100.21.0.2.1
                Announcement bits (1): 0-TED Export
                AS path: 100 I
                Accepted
                Localpref: 100
                Router ID: 100.100.100.0
                Thread: junos-main
```

show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1                  *[MPLS/0] 00:13:55, metric 1
                      Receive
2                  *[MPLS/0] 00:13:55, metric 1
```

```
                      Receive
1024               *[VPN/0] 00:04:18
                      to table red.inet.0, Pop
```

show route table mpls.0 protocol ospf

```
user@host> show route table mpls.0 protocol ospf
mpls.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299952             *[L-OSPF/10] 23:59:42, metric 0
                    > to 10.0.10.70 via lt-1/2/0.14, Pop
                      to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299952(S=0)        *[L-OSPF/10] 23:59:42, metric 0
                    > to 10.0.10.70 via lt-1/2/0.14, Pop
                      to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299968             *[L-OSPF/10] 23:59:48, metric 0
                    > to 10.0.6.60 via lt-1/2/0.12, Pop
```

show route table VPN-AB.inet.0

```
user@host> show route table VPN-AB.inet.0
      VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.39.1.0/30       *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
      10.39.1.4/30       *[Direct/0] 00:08:42
                          > via so-5/1/0.0
      10.39.1.6/32       *[Local/0] 00:08:46
                            Local
      10.255.71.16/32    *[Static/5] 00:07:24
                          > via so-2/0/0.0
      10.255.71.17/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                             AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.18/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                             AS path: I
```

```
                        > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32  *[BGP/170] 00:08:35, localpref 100
                      AS path: 2 I
                        > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32  *[OSPF/10] 00:07:24, metric 1
                        > via so-7/3/1.0
```

Release Information

Command introduced before Junos OS Release 7.4.

Show route table evpn statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.

RELATED DOCUMENTATION

- show route summary

show route tenant

IN THIS SECTION

- Syntax | 23822
- Description | 23822
- Options | 23822
- Required Privilege Level | 23822
- Output Fields | 23822
- Sample Output | 23823
- Release Information | 23824

Syntax

```
show route tenant
< (all | tenant-name)>
```

Description

Displays the routing table information for the tenant systems.

A tenant system is a logical partitioning of the device into multiple security domains similar to logical systems, to provide high scalability. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. A set of interfaces that belong to the routing instances and the routing protocol parameters control the information in the routing tables. A tenant system can configure the assigned routing instances and the interfaces within a tenant system.

Options

all Displays the summary statistics about the all entries in the routing table.

tenant-name Specify the tenant system name.

Required Privilege Level

view

Output Fields

Table 1106 on page 23822 lists the output fields for the `show route tenant` command. Output fields are listed in the approximate order in which they appear.

Table 1106: show route summary Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.o).
destinations	Number of destinations that correspond to the routes in the routing table.

Table 1106: show route summary Output Fields *(Continued)*

Field Name	Field Description
routes	Number of routes in the routing table: • **active**—Number of routes that are active. • **holddown**—Number of routes that are in the hold-down state before being declared inactive. • **hidden**—Number of routes that are not used because of routing policy.
Direct	Routes on the directly connected network.
Local	Specify the local routes.
OSPF	Name of the protocol from which the route was learnt.

Sample Output

show route tenant tenant1

```
user@host> show route tenant tenant1
tenant1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
12.1.1.0/24        *[Direct/0] 00:33:07
> via ge-0/0/0.0
12.1.1.1/32        *[Local/0] 00:33:07
Local via ge-0/0/0.0
224.0.0.5/32       *[OSPF/10] 00:30:31, metric 1
MultiRecv
tenant1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
ff02::2/128        *[INET6/0] 00:33:33
MultiRecv
```

Release Information

Statement introduced in Junos OS Release 18.3R1.

show route terse

IN THIS SECTION

- Syntax | 23824
- Syntax (EX Series Switches) | 23824
- Description | 23824
- Options | 23825
- Required Privilege Level | 23825
- Output Fields | 23825
- Sample Output | 23828
- Release Information | 23829

Syntax

```
show route terse
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route terse
```

Description

Display a high-level summary of the routes in the routing table.

> **NOTE**: For BGP routes, the `show route terse` command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.
>
> To display the metric1 and metric2 value of a BGP route, use the `show route extensive` command.

Options

none	Display a high-level summary of the routes in the routing table.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1107 on page 23825 describes the output fields for the `show route terse` command. Output fields are listed in the approximate order in which they appear.

Table 1107: show route terse Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number destinations	Number of destinations for which there are routes in the routing table.
number routes	Number of routes in the routing table and total number of routes in the following states: • `active` (routes that are active) • `holddown` (routes that are in the pending state before being declared inactive) • `hidden` (routes that are not used because of a routing policy)

Table 1107: show route terse Output Fields *(Continued)*

Field Name	Field Description
route key	Key for the state of the route: • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.
A	Active route. An asterisk (*) indicates this is the active route.
V	Validation status of the route: • ?—Not evaluated. Indicates that the route was not learned through BGP. • I—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • N—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database. • V—Valid. Indicates that the prefix and autonomous system pair are found in the database.
Destination	Destination of the route.

Table 1107: show route terse Output Fields *(Continued)*

Field Name	Field Description
P	Protocol through which the route was learned: • A—Aggregate • B—BGP • C—CCC • D—Direct • G—GMPLS • I—IS-IS • L—L2CKT, L2VPN, LDP, Local • K—Kernel • M—MPLS, MSDP • O—OSPF • P—PIM • R—RIP, RIPng • S—Static • T—Tunnel
Prf	Preference value of the route. In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.
Metric 1	First metric value in the route. For routes learned from BGP, this is the MED metric.
Metric 2	Second metric value in the route. For routes learned from BGP, this is the IGP metric.

Table 1107: show route terse Output Fields *(Continued)*

Field Name	Field Description
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated.

Sample Output

show route terse

```
user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.16.1.1/32         O  10          1            >10.0.0.2
  ?                     B 170        100                             I
  unverified                                        >10.0.0.2
* ? 172.16.1.1/32         D   0                        >lo0.2
* V 2.2.0.2/32          B 170        110                             200 I
  valid                                             >10.0.0.2
* ? 10.0.0.0/30         D   0                        >lt-1/2/0.1
  ?                     B 170        100                             I
  unverified                                        >10.0.0.2
* ? 10.0.0.1/32         L   0                         Local
* ? 10.0.0.4/30         B 170        100                             I
  unverified                                        >10.0.0.2
* ? 10.0.0.8/30         B 170        100                             I
  unverified                                        >10.0.0.2
* I 172.16.1.1/32       B 170         90                             200 I
```

```
  invalid                                          >10.0.0.2
* N 192.168.2.3/32     B 170        100                              200 I
  unknown                                          >10.0.0.2
* ? 172.16.233.5/32       O  10          1             MultiRecv
```

Release Information

Command introduced before Junos OS Release 7.4.

show route vpn-localization

IN THIS SECTION

- Syntax | 23829
- Description | 23829
- Options | 23830
- Required Privilege Level | 23830
- Output Fields | 23830
- Sample Output | 23832
- Release Information | 23835

Syntax

```
show route vpn-localization
<brief | detail | terse>
<family (inet| inet6)>
<vpn-name vpn-name>
```

Description

Display the localization information of all the VRFs in the system.

Options

none	Display the localization information of all the VRFs in the system.
brief | detail | terse	(Optional) Display the specified level of output.
family (inet | inet6)	(Optional) Display the localized information of all the VRFs in the system for inet or inet6 family.
vpn-name ***vpn-name***	(Optional) Display the localized information of the specific VPN.

Required Privilege Level

view

Output Fields

Table 1108 on page 23830 describes the output fields for the **show route vpn-localization** command. Output fields are listed in the approximate order in which they appear.

Table 1108: show route vpn-localization Output Fields

Field Name	Field Description	Level of Output
Routing table	Name of the routing table	none **detail**
Index	Index of the routing table.	none **detail**
Address Family	Address family of the routing table (inet/inet6).	none **detail**

Table 1108: show route vpn-localization Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Localization status	Localization status of the routing table and can be one of the following: • Complete - The routing table successfully installed on all the local FPCs. • Pending - The localization change on this routing table is yet to be processed in the pending state. The routing table might not be present on all the local FPCs or can be present on the non-local FPCs. • Replaying - The routing table is being installed on the local FPCs. A localization change involving the addition of new local FPCs requires the system to replay the routing table for newly added FPCs.	none **detail**
Local FPC's	List of FPCs on which the routing table should be installed, as per the configuration.	none
Local FPC (configured)	List of FPCs on which the routing table should be installed, as per the configuration.	**detail**
Local FPC (Current state)	List of FPCs on which the routing table is currently installed.	**detail**
Table Name	Name of the routing table.	**terse**
Local'd	The localization status of VRF. If VRF is localized, then the status is Yes; otherwise, it is No.	**terse**
Table Index	Index of the routing table.	**terse**
Local FPC's Configured	List of FPCs, in hexadecimal, on which the routing table should be installed as per the configuration.	**terse**
Local FPC's Actual	List of FPCs, in hexadecimal, on which the routing table is currently installed.	**terse**

Table 1108: show route vpn-localization Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Localization status of the routing table and can be one of the following: • Complete - The routing table successfully installed on the local FPCs. • Pending - The localization change on this routing table is yet to be processed in the pending state. The routing table might not be present on all the local FPCs or can be present on the non-local FPCs. • Replaying - The routing table is being installed on the local FPCs. A localization change involving addition of new local FPCs requires the system to replay the routing table for the newly added FPCs.	**terse**

Sample Output

show route vpn-localization

```
user@host> show route vpn-localization

Routing table: CE-1.inet, Localized
  Index: 7, Address Family: inet, Localization status: Complete
  Local FPC's: 4 5 13 18

Routing table: CE-1.inet6, Localized
  Index: 7, Address Family: inet6, Localization status: Complete
  Local FPC's: 4 5 13 18

Routing table: CE-10.inet, Localized
  Index: 8, Address Family: inet, Localization status: Complete
  Local FPC's: 4 5 13

Routing table: CE-10.inet6, Localized
  Index: 8, Address Family: inet6, Localization status: Complete
  Local FPC's: 4 5 13

Routing table: CE-2.inet, Localized
  Index: 9, Address Family: inet, Localization status: Complete
```

```
  Local FPC's: 4 5 10 13

Routing table: CE-2.inet6, Localized
  Index: 9, Address Family: inet6, Localization status: Complete
  Local FPC's: 4 5 10 13

Routing table: CE-3.inet, Localized
  Index: 10, Address Family: inet, Localization status: Complete
  Local FPC's: 4 5 9 13

Routing table: CE-3.inet6, Localized
  Index: 10, Address Family: inet6, Localization status: Complete
  Local FPC's: 4 5 9 13

Routing table: CE-4.inet, Localized
  Index: 11, Address Family: inet, Localization status: Complete
  Local FPC's: 3 4 5 13

Routing table: CE-4.inet6, Localized
  Index: 11, Address Family: inet6, Localization status: Complete
  Local FPC's: 3 4 5 13

Routing table: CE-5.inet, Localized
  Index: 12, Address Family: inet, Localization status: Complete
  Local FPC's: 4 5 13

Routing table: CE-5.inet6, Localized
  Index: 12, Address Family: inet6, Localization status: Complete
  Local FPC's: 4 5 13

Routing table: CE-6.inet, Localized
  Index: 13, Address Family: inet, Localization status: Complete
  Local FPC's: 4 5 13 18

Routing table: CE-6.inet6, Localized
  Index: 13, Address Family: inet6, Localization status: Complete
  Local FPC's: 4 5 13 18

Routing table: CE-7.inet, Localized
  Index: 14, Address Family: inet, Localization status: Complete
  Local FPC's: 4 5 10 13

Routing table: CE-7.inet6, Localized
```

```
  Index: 14, Address Family: inet6, Localization status: Complete
  Local FPC's: 4 5 10 13

Routing table: CE-8.inet, Localized
  Index: 15, Address Family: inet, Localization status: Complete
  Local FPC's: 4 5 9 13

Routing table: CE-8.inet6, Localized
  Index: 15, Address Family: inet6, Localization status: Complete
  Local FPC's: 4 5 9 13

Routing table: CE-9.inet, Localized
  Index: 16, Address Family: inet, Localization status: Complete
  Local FPC's: 3 4 5 13

Routing table: CE-9.inet6, Localized
  Index: 16, Address Family: inet6, Localization status: Complete
  Local FPC's: 3 4 5 13
```

show route vpn-localization vpn-name vpn-name family inet

```
user@host> show route vpn-localization vpn-name vpn-m200 family inet

Routing table: vpn-m200.inet, Localized
  Index: 7, Address Family: inet, Localization status: Complete
  Local FPC's: 0 1
```

show route vpn-localization vpn-name vpn-name family inet detail

```
user@host> show route vpn-localization vpn-name vpn-m200 family inet detail

Routing table: vpn-m200.inet, Non-localized
  Index: 7, Address Family: inet
  Localization status: Complete
  Local FPC (configured):
    List: All
  Local FPC (Current state):
    List: All
```

show route vpn-localization terse

```
user@host> show route vpn-localization terse

Table              Local'd  Table   Local FPC's
Name                         Index    Configured         Actual         State
vpn-m200.inet         Yes     7     0000000000000003 0000000000000003 Complete
vpn-m200.inet6        Yes     7     0000000000000003 0000000000000003 Complete
vpn-m300.inet         Yes     8     0000000000000005 0000000000000005 Complete
vpn-m300.inet6        Yes     8     0000000000000005 0000000000000005 Complete
vpn-m500.inet         Yes     9     0000000000000021 0000000000000021 Complete
vpn-m500.inet6        Yes     9     0000000000000021 0000000000000021 Complete
```

Release Information

Command introduced in Junos OS Release 14.2.

show routing nhdb pfe-install-error

IN THIS SECTION

- Syntax | 23835
- Description | 23836
- Required Privilege Level | 23836
- Output Fields | 23836
- Sample Output | 23837
- Release Information | 23837

Syntax

```
show routing nhdb pfe-install-error
```

Description

(Junos OS Evolved only) Use this command to check for next-hop errors. You must first configure the `pfe-install-error` statement at the `[edit routing-options forwarding-table]` hierarchy level and restart the rpd process to enable this error reporting.

Errors are supported for routes pointing to a limited set of next hops: RT_NH_INDIRECT, RT_NH_ROUTER, RT_NH_CHAIN, RT_NH_INDXD, RT_NH_LIST, RT_NH_TUNNEL_COMP, and RT_NH_FRR_INDIRECT.

Required Privilege Level

view

Output Fields

Table 1109: show routing nhdb pfe-install-error Output Fields

Field Name	Field Description
Error category	There are two categories: • **PFE:** Error is received directly from the Packet Forwarding Engine. • **DERIVED:** Error is derived within the rpd process from input from the Packet Forwarding Engine. For example, if a route points to a next hop and the rpd process receives an error on that next hop from the Packet Forwarding Engine, then the rpd process also knows that there is an error on the route that points to the next hop.
Error Type	Error type is 0x1.
Handle	Address of the object.
Index	Unique identifier of a next hop.
Time	Time when error occurred.

Sample Output

show routing nhdb pfe-install-error

```
user@host> show routing nhdb pfe-install-error
```

```
Handle              Error Type      Error category       Index      Time

0x5581d16b979c      0x1             PFE                  12002      20XX-XX-XX XX:XX:XX

0x5581d16b987c      0x1             PFE                  12001      20XX-XX-XX XX:XX:XX
```

Release Information

Command introduced in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

forwarding-table

show rsvp interface

IN THIS SECTION

- Syntax | 23838
- Syntax (EX Series Switches) | 23838
- Description | 23838
- Options | 23838
- Required Privilege Level | 23839
- Output Fields | 23839
- Sample Output | 23843

Release Information | 23847

Syntax

```
show rsvp interface
<brief | detail | extensive>
<instance instance-name>
<link-management>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show rsvp interface
<brief | detail | extensive>
<link-management>
```

Description

Display the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics. The RSVP input/input module collects statistics for certain events on a per-interface basis. Most of these events were tracked on a routing-instance basis in Junos OS releases earlier than Release 17.2. The `show rsvp interface detail` command displays these event counters under the `Events` section of the output only when the values of these fields are higher than zero. These statistics are also maintained at the global level (per routing-instance) and are also displayed in the output of the `show rsvp statistics` command.

Options

none	Display standard information about the status of RSVP-enabled interfaces and packet statistics.
brief | detail | extensive | link-management	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display RSVP status information for the specified instance. If `instance-name` is omitted, RSVP status information is displayed for the master instance.

link-management (Optional) Use the link-management option to display the control peers and corresponding TE-link information created by the Link Management Protocol (LMP).

logical-system (all | *logical-system-name*) (Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1110 on page 23839 lists the output fields for the `show rsvp interface` command. Output fields are listed in the approximate order in which they appear.

Table 1110: show rsvp interface Output Fields

Field Name	Field Description	Level of Output
`RSVP interface`	Number of interfaces on which RSVP is active. Each interface has one line of output.	All levels
`Interface`	Name of the interface.	All levels
`Index`	Index of the interface.	`detail`
`State`	State of the interface. • `Disabled`—No traffic engineering information is displayed. • `Down`—Interface is not operational. • `Enabled`—Displays traffic engineering information. • `Up`—Interface is operational.	All levels
`NoAuthentication`	Interface does not support RSVP authentication.	`detail`

Table 1110: show rsvp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`NoAggregate`	Interface does not support refresh reduction.	`detail`
`NoReliable`	Interface does not support refresh reduction message ID extension.	`detail`
`NoLinkProtection`	Interface does not support link protection.	`detail`
`HelloInterval`	Frequency at which RSVP hellos are sent on this interface (in seconds). Prior to Junos OS Release 18.2R2, when the `no-interface-hello` statement is configured under the `[edit protocols rsvp]` hierarchy, and there is no interface-specific configuration for the hello interval, the `HelloInteval` output field displayed the default hello interval time of 9 seconds. Starting in Junos OS Release 18.2R2, with a similar configuration, the `HelloInterval` output field displays 0 as the hello interval.	`detail`
`Address`	IP address of the local interface.	`detail`
`Active control channel`	Next-hop link address to transmit messages.	None specified
`TElink`	Traffic-engineered links that are managed by the peer they are associated with.	None specified
`Active resv`	Number of reservations that are actively reserving bandwidth on the interface.	All levels
`PreemptionCnt`	Number of times an RSVP session was preempted on this interface.	`detail`
`Update threshold`	Percentage change in reserved bandwidth to trigger an IGP update.	`detail`
`Subscription`	User-configured subscription factor.	All levels

Table 1110: show rsvp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Actual`	Available RSVP bandwidth that is recalculated after considering SPRING bandwidth utilization.	`extensive`
`bc` *number*	Bandwidth allocated for the specified bandwidth constraint.	`extensive`
`ct` *number*	Bandwidth allocated for the specified class type.	`extensive`
`Static BW`	Total interface bandwidth, in bps.	All levels
`Available BW`	Amount of bandwidth that RSVP is allowed to reserve, in bps. It is equal to (static bandwidth * subscription factor).	al levels
`Reserved BW`	Currently reserved bandwidth, in bps.	All levels
`SoftPreemptionCnt`	Number of times a soft preemption occurred on this interface. This number is not included in the `PreemptionCnt` value.	`detail`
`Overbooked BW`	Currently overbooked bandwidth, in bps, by class type (`ct0` through `ct3`).	`detail`
`Highwater mark`	Highest bandwidth that has ever been reserved on this interface, in bps.	`brief`
`PacketType`	Type of RSVP packet.	`detail`
`Total Sent`	Total number of packets sent.	`detail`
`Total Received`	Total number of packets received since RSVP was enabled.	`detail`
`Last 5 seconds Sent`	Number of packets sent in the last 5 seconds.	`detail`

Table 1110: show rsvp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Last 5 seconds Received`	Number of packets received in the last 5 seconds.	`detail`
`Path`	Statistics about Path messages, which are sent from the RSVP sender along the data paths and store path state information in each node along the path.	`detail`
`PathErr`	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.	`detail`
`PathTear`	Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.	`detail`
`Resv`	Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and store reservation state information in each node along the path.	`detail`
`ResvErr`	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.	`detail`
`ResvTear`	Statistics about ResvTear messages, which remove reservation states along a path.	`detail`
`Hello`	Number of RSVP hello packets that have been sent to and received from the neighbor.	`detail`
`Ack`	Acknowledge message for refresh reductions.	`detail`
`Srefresh`	Summary refresh messages.	`detail`
`EndtoEnd RSVP`	Statistics for the number of end-to-end RSVP messages sent.	`detail`

Table 1110: show rsvp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Queue	CoS transmit queue number and its associated forwarding class designation.	extensive
TxRate	Configured bandwidth in Mbps and configured bandwidth as a percentage of the specified queue.	extensive
Priority	Weight of the queue relative to other configured queues, in percentage.	extensive
queue-priority-value	Low, High, None, or Exact. None indicates no rate limiting. Exact indicates the queue transmits at the configured rate only.	extensive
subscription*priority percentage*	Indicates per-priority subscription percentage.	extensive

Sample Output

show rsvp interface brief

```
user@host> show rsvp interface brief
RSVP interface: 1 active
                  Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
de0.0     Up         1    23%    10Mbps      989.992kbps 1.31Mbps    1.31Mbps
```

show rsvp interface detail

Starting in Junos OS Release 15.2, this command also shows conditional PathTear statistics and Node Hellos.

```
user@host> show rsvp interface detail
so-0/1/1.0  Index 6, State: Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 3(second)
```

```
  Address 192.168.207.29, 10.255.245.194
  ActiveResv 0, PreemptionCnt 0, SoftPreemptionCnt 0, Update threshold 10%
  Subscription 100%, StaticBW 155.52Mbps, AvailableBW 155.52Mbps
ReservedBW [0] 155Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
SoftPreemptionCnt1
OverbookedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 155Mbps[5] 0bps[6] 0bps[7] 0bps
PacketType              Total                  Last 5 seconds
                   Sent      Received        Sent      Received
  Path               16             0           1             0
  PathErr             0             0           0             0
  PathTear            1             0           0             0
  Resv                0            11           0             1
  ResvErr             0             0           0             0
  ResvTear            0             0           0             0
  Hello              66            67           1             1
  Ack                 0             0           0             0
  Srefresh            0             0           0             0
  EndtoEnd RSVP       0             0           0             0
  Node Hello        100           100           0             0
  PathTear(Condl.)    0             3           0             0
```

show rsvp interface extensive

```
user@host> show rsvp interface extensive
so-1/0/0.0 Index 72, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 192.168.213.22, 10.255.240.175
  ActiveResv 1, PreemptionCnt 0, SoftPreemptionCnt 0, Update threshold 10%
  Subscription 100%, Actual 60%
  bc0 = (ct0+ct1+ct2+ct3), StaticBW 622.08Mbps
  bc1 = (ct1+ct2+ct3), StaticBW 466.56Mbps
  bc2 = (ct2+ct3), StaticBW 311.04Mbps
  bc3 = ct3, StaticBW 155.52Mbps
  ct0: StaticBW 155.52Mbps, AvailableBW 522.08Mbps
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  ct1: StaticBW 155.52Mbps, AvailableBW 366.56Mbps
  ReservedBW [0] 100Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  ct2: StaticBW 155.52Mbps, AvailableBW 311.04Mbps
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
```

```
  ct3: StaticBW 155.52Mbps, AvailableBW 155.52Mbps
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Queue          TxRate             Priority Exact
    0            155.52Mbps         25%      Low
    1            155.52Mbps         25%      Low
    2            155.52Mbps         25%      Low
    3            155.52Mbps         25%      Low
```

show rsvp interface link-management

```
user@host> show rsvp interface link-management
RSVP interface: 2 active
PEER-C  State: Up
Active Control Channel: so-0/1/0.0

  TElink: TElnk1, Link ID: 37811
  ActiveResv 0, PreemptionCnt 0
  StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

  TElink: TElnk2, Link ID: 37808
  ActiveResv 1, PreemptionCnt 0
  StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

PEER-B  State: Up
Active Control Channel: so-1/0/0.0

  TElink: TElnkAB1, Link ID: 1598
  ActiveResv 0, PreemptionCnt 0
  StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

  TElink: TElnkAB2, Link ID: 1597
  ActiveResv 0, PreemptionCnt 0
  StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps
```

show rsvp interface detail RSVP interface: 9 active

```
user@host> show rsvp interface detail RSVP interface: 9 active
fxp0.0 Index 4, State Dis/Up
NoAuthentication, Aggregate, Reliable, NoLinkProtection HelloInterval 9(second) Address
10.9.148.47
```

```
Event                              Count
bad packet length                    1
bad packet version                   1
authentication fail                  1
bad checksum                         1
bad packet format                    1
rcv pkt disabled intf                1
state timeout                        1
message out-of-order                 1
unknown ack                          1
unknown nack                         1
received nack                        1
send failure                         1

PacketType         Total      Last 5 seconds

Sent      Received        Sent      Received
Path                 0             0           0             0
PathErr              0             0           0             0
PathTear             0             0           0             0
Resv                 0             0           0             0
ResvErr              0             0           0             0
ResvTear             0             0           0             0
ResvConf             0             0           0             0
Bundle               0             0           0             0
Hello                0             0           0             0
Ack                  0             0           0             0
Srefresh             0             0           0             0
Notify               0             0           0             0
Unknown              0             0           0             0
EndtoEnd RSVP        0             0           0             0
Backup Path          0             0           0             0
Cnd PathTear         0             0           0             0
```

show rsvp interface extensive (With subscription priority configured)

```
user@host> show rsvp interface extensive ge-0/0/0.0
ge-0/0/0.0 Index 361, State Ena/Up
  NoAuthentication, Aggregate, Reliable, LinkProtection
  HelloInterval 9(second)
```

```
  Address 10.1.2.1
  ActiveResv 1, PreemptionCnt 0, MaxResvTh 0bps, 0%
  Update threshold 10.000%, UpdateThresholdValue 100Mbps
  Subscription 100%, Actual 100%
 Priority: [0] 100% [1] 100% [2] 100% [3] 100% [4] 100% [5] 70% [6] 80% [7] 90%
  bc0 = ct0, StaticBW 1000Mbps
  ct0: StaticBW 1000Mbps, AvailableBW 900Mbps
    MaxAvailableBW 1000Mbps = (bc0*subscription)
    ReservedBW [0] 100Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: On, Bypass: 1, LSP: 1, Protected LSP: 1, Unprotected LSP: 0
      1 Mar  2 22:01:16 New bypass Bypass->10.1.2.2
    Bypass: Bypass->10.1.2.2, State: Up, Type: LP, LSP: 1, Backup: 0
      4 Mar  2 22:01:17 Up
      3 Mar  2 22:01:17 Record Route:  10.1.2.18(Label=3)
      2 Mar  2 22:01:17 CSPF: computation result accepted
      1 Mar  2 22:01:16 Originate Call
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` option added in Junos OS Release 15.1 for the MX Series.

show rsvp neighbor

IN THIS SECTION

- Syntax | 23848
- Syntax (EX Series Switches) | 23848
- Description | 23848
- Options | 23848
- Required Privilege Level | 23848
- Output Fields | 23848
- Sample Output | 23853
- Release Information | 23854

Syntax

```
show rsvp neighbor
<brief | detail | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show rsvp neighbor
<brief | detail>
```

Description

Display Resource Reservation Protocol (RSVP) neighbors that were discovered dynamically during the exchange of RSVP packets.

Options

none	Display standard information about RSVP neighbors.
brief | detail	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display the RSVP neighbor information for the specified instance. If `instance-name` is omitted, RSVP neighbor information is displayed for the master instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1111 on page 23849 lists the output fields for the `show rsvp neighbor` command. Output fields are listed in the approximate order in which they appear.

Table 1111: show rsvp neighbor Output Fields

Field Name	Field Description	Level of Output
`RSVP neighbor`	Number of neighbors that the routing device has learned of. Each neighbor has one line of output.	All levels
`via`	Name of the interface where the neighbor has been detected. In the case of generalized MPLS (GMPLS) LSPs, the name of the peer where the neighbor has been detected.	`detail`
`Address`	Address of a learned neighbor.	All levels
`Idle`	Length of time the neighbor has been idle, in seconds. **NOTE**: Until Junos OS Release 15.1, in the output of the show rsvp neighbor command, the value under the Idle field immediately reflects the changed idle time when a link in the neighboring router is brought down. Starting with Junos OS Release 15.2, a router does not declare a neighbor as idle when a hello adjacency exists and has not timed out. When an interface is brought down, RSVP brings down the neighbor because of the notification it receives from IGP. The reason for considering the IGP-down notification is to support BFD-triggered fast reroute (FRR) and RSVP-TE is not directly a client for BFD notifications. When RSVP brings down the neighbor, the input/output process is not impacted. As a result, the idle time in the output of the show command is not immediately updated.	All levels
`Up/Dn`	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	All levels
`Up cnt and Down cnt`	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	`detail`

Table 1111: show rsvp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output

status	State of the RSVP neighbor:
	• Up—Routing device can detect RSVP Hello messages from the neighbor.
	• Down—Routing device has received one of the following indications:
	  • Communication failure from the neighbor.
	  • Communication from IGP that the neighbor is unavailable.
	  • Change in the sequence numbers in the RSVP Hello messages sent by the neighbor.
	• Restarting—RSVP neighbor is unavailable and might be restarting. The neighbor remains in this state until it has restarted or is declared dead. This state is possible only when graceful restart is enabled.
	• Restarted—RSVP neighbor has restarted and is undergoing state recovery (graceful restart) procedures.
	• Dead—Routing device has lost all communication with the RSVP neighbor. Any RSVP sessions with that neighbor are torn down.
		detail

LastChange	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is *hh*:*mm*:*ss*.	All levels
Last changed time	Time elapsed since the neighbor state changed either from up to down or from down to up.	detail
HelloInt	Frequency at which RSVP hellos are sent on this interface (in seconds).	All levels
HelloTx/Rx	Number of hello packets sent to and received from the neighbor.	All levels
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.	detail

Table 1111: show rsvp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Message received`	Number of Path and Resv messages that this routing device has received from the neighbor.	`detail`
`Remote Instance`	Identification provided by the remote routing device during Hello message exchange.	`detail`
`Local Instance`	Identification sent to the remote routing device during Hello message exchange.	`detail`
`Refresh reduction`	Measure of processing overhead requests of refresh messages. Refresh reduction extensions improve routing device performance by reducing the process overhead, thus increasing the number of LSPs a routing device can support. `Refresh reduction` can have the following values: • `operational`—All four RSVP refresh reduction extensions—message ack, bundling, summary refresh, and staged refresh timer—are functional between the two neighboring routing devices. For a detailed explanation of these extensions, see RFC 2961. • `incomplete`—Some RSVP refresh reduction extensions are functional between the two neighboring routing devices. • `not operational`—Either the refresh reduction feature has been turned off, or the remote routing device cannot support the refresh reduction extensions.	`detail`
`Remote end`	Neighboring routing device's status with regard to refresh reduction: • `enabled`—Remote routing device has requested refresh reduction during RSVP message exchanges. • `disabled`—Remote routing device does not require refresh reduction.	`detail`
`Pop label`	Pop labels of the RSVP-TE pop-and-forward LSP tunnels.	`detail`

Table 1111: show rsvp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Ack-extension`	An RSVP refresh reduction extension: • `enabled`—Both local and remote routing devices support the ack-extension (RFC 2961). • `disabled`—Remote routing device does not support the ack-extension.	`detail`
`Link protection`	Status of the MPLS fast reroute mechanism that protects traffic from link failure: • `enabled`—Link protection feature has been turned on, protecting the neighbor with a bypass LSP. • `disabled`—No link protection feature has been enabled for this neighbor.	`detail`
`LSP name`	Name of the bypass LSP.	`detail`
`Bypass LSP`	Status of the bypass LSP. It can have the following values: • `does not exist`—Bypass LSP is not available. • `connecting`—Routing device is in the process of establishing a bypass LSP, and the LSP is not available for link protection at the moment. • `operational`—Bypass LSP is up and running. • `down`—Bypass LSP has gone down, with the most probable cause a node or a link failure on the bypass path.	`detail`
`Backup routes`	Number of user LSPs (or routes) that are being protected by a bypass LSP (before link failure).	`detail`
`Backup LSPs`	Number of LSPs that have been temporarily established to maintain traffic by refreshing the downstream LSPs during link failure (not a one-to-one correspondence).	`detail`

Table 1111: show rsvp neighbor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Bypass explicit route`	Explicit route object's (ERO) path that is taken by the bypass LSP.	`detail`
`Restart time`	Length of time a neighbor waits to receive a Hello from the restarting node before declaring the node dead and deleting the states (in milliseconds).	`detail`
`Recovery time`	Length of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds). Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.	`detail`

Sample Output

show rsvp neighbor

```
user@host> show rsvp neighbor
RSVP neighbor: 2 learned
Address            Idle Up/Dn LastChange HelloInt HelloTx/Rx
192.168.207.203       0  3/2       13:01        3   366/349
192.168.207.207       0  1/0       22:49        3   448/448
```

show rsvp neighbor detail

Starting in Junos OS Release 16.1, this command also shows whether enhanced FRR procurers are enabled on the neighbor. Neighbors with Point of Local Repair (PLR) or Node Protecting Merge Point (NP-MP) also show the Hellos sent /received count.

```
user@host> show rsvp neighbor detail
RSVP neighbor: 2 learned
Address: 192.168.207.203   via: ecstasyl status: Up
   Last changed time: 28:47, Idle: 0 sec, Up cnt: 3, Down cnt: 2
   Message received: 632
```

```
    Hello: sent 673, received 656, interval 3 sec
    Remote instance: 0x6432838a, Local instance: 0x74b72e36
    Refresh reduction:  operational
      Remote end: enabled,  Ack-extension: enabled
    Enhanced FRR local protection: enabled
      LSPs (total 76): Phop 0, PPhop 0, Nhop 76, NNhop 0
    Pop Label: 299808(unprotected) 299840(link-protected)
 Link protection:  enabled
      LSP name: Bypass_to_192.168.207.203
      Bypass LSP: operational,  Backup routes: 1,  Backup LSPs: 0
      Bypass explicit route: 192.168.207.207 192.168.207.224
   Restart time: 60000 msec, Recovery time: 0 msec
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` option added in Junos OS Release 15.1 for the MX Series.

show rsvp pop-and-forward

IN THIS SECTION

- Syntax | 23854
- Description | 23855
- Options | 23855
- Required Privilege Level | 23855
- Sample Output | 23856
- Release Information | 23857

Syntax

```
show rsvp pop-and-forward
<brief | detail | extensive>
```

```
<instance routing-instance-name>
<label label>
<logical-system (all | logical-system-name)>
```

Description

Display RSVP-TE pop-and-forward LSP tunnel information. This information includes the set of in-labels (one-hop pop label or a delegation label), the number of session using each label and the next segment-label (if there is another delegation hop downstream), and whether the in-label is used for unprotected or protected LSPs.

Options

none	Display the standard level of information for the RSVP-TE pop-and-forward LSP tunnels.
brief | detail | extensive	(Optional) Display the desired level of output. The brief option is the default level of output. The detail option provides more information about the hops in a delegation segment (whether its one-hop or multi-hop). The extensive option lists the set of LSPs that are using a given pop or delegation label.
instance *routing-instance-name*	(Optional) Display the RSVP-TE pop-and-forward LSP tunnel information for the specified routing instance.
label *label*	(Optional) Display the RSVP-TE pop-and-forward LSP tunnel information for the specified label.
logical-system (all | *logical-system-name*)	(Optional) Display the RSVP-TE pop-and-forward LSP tunnel information for all or the specified logical system.

Required Privilege Level

view

Sample Output

show rsvp pop-and-forward

```
user@host> show rsvp pop-and-forward
RSVP pop-and-forward: 2 shared labels
Label-in     Hop-count  Next-segment- Protection     Session-
                        label                        count
299840       3          299808        unprotected    100
299872       3          299824        unprotected    50
```

show rsvp pop-and-forward extensive

```
user@host> show rsvp pop-and-forward extensive
RSVP pop-and-forward: 2 shared labels
299840 (shared-label)
  Next-segment-label: 299808, Hop-count: 3
  Protection: unprotected, Session-count: 2
  Segment-id:
    Hop 1: 70.1.1.2(label=299808)
    Hop 2: 92.1.1.1(label=299808)
    Hop 3: 93.1.1.2
  Segment route:
    Primary: 70.1.1.2, OutIf: ge-0/0/2.0
  Lsp-session list (name, dest-ip, sender-ip, lsp-id):
    pop1, 10.10.10.10, 2.2.2.2, 2
    pop2, 10.10.10.10, 2.2.2.2, 1

299872 (shared-label)
  Next-segment-label: 299824, Hop-count: 3
  Protection: unprotected, Session-count: 4
  Segment-id:
    Hop 1: 70.1.1.2(label=299808)
    Hop 2: 92.1.1.1(label=299808)
    Hop 3: 93.1.1.2
  Segment route:
    Primary: 70.1.1.2, OutIf: ge-0/0/2.0
  Lsp-session list (name, dest-ip, sender-ip, lsp-id):
    pop147, 9.9.9.9, 2.2.2.2, 1
    pop148, 9.9.9.9, 2.2.2.2, 1
```

```
    pop150, 9.9.9.9, 2.2.2.2, 1
    pop149, 9.9.9.9, 2.2.2.2, 1
```

show rsvp pop-and-forward label

```
user@host> show rsvp pop-and-forward label 299872
RSVP pop-and-forward: 2 shared labels
Label-in     Hop-count  Next-segment- Protection     Session-
                        label                        count
299872       3          299824        unprotected    4
```

Release Information

Command introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Pop-and-Forward LSP Configuration

pop-and-forward (Protocols RSVP) | **12803**

pop-and-forward (Protocols MPLS) | **12802**

show rsvp route-session-id

IN THIS SECTION

- Syntax | **23858**
- Description | **23858**
- Options | **23858**
- Required Privilege Level | **23858**
- Output Fields | **23858**
- Sample Output | **23859**
- Release Information | **23859**

Syntax

```
show rsvp route-session-id
```

Description

Display the session ID and the version information associated with the ingress route added by the Resource Reservation Protocol (RSVP) in the inet.3 table.

Session ID is a pre-populated identifier used for indirect next hops in BGP Prefix Independent Convergence (PIC) enabled router. Session ID is used to identify the session or path.

NOTE: `protect core` configuration is not required to display the `route-session-id`.

Options

none Validate and display RSVP route session details.

Required Privilege Level

view

Output Fields

Table 1112 on page 23858 describes the output fields for the `show rsvp route-session-id` command. Output fields are listed in the approximate order in which they appear.

Table 1112: show rsvp route-session-id Output Fields

Field Name	Field Description
Ingress Route Destination	Destination (egress routing device) of the session.
Ingress Route Preference	RSVP preference value of the ingress session.

Table 1112: show rsvp route-session-id Output Fields *(Continued)*

Field Name	Field Description
Ingress Route Metric 1	Metric 1 associated with the RSVP ingress route.
Ingress Route Metric 2	Metric 2 associated with the RSVP ingress route.
Ingress Route Session ID	Session ID associated with the RSVP ingress route.
Version	Version number associated with the RSVP ingress route.

Sample Output

show rsvp route-session-id

```
user@host> show rsvp route-session-id

RSVP Ingress Route Session ID Database:
======================================
   Ingress Route Destination: 1.1.1.5/32
   Ingress Route Preference:  7
   Ingress Route Metric 1:    20,        Metric 2:  0
   Ingress Route Session ID:  0x146,     Version:   0
```

Release Information

Command introduced in Junos OS Release 16.1.

show rsvp session

IN THIS SECTION

- Syntax | 23860
- Syntax (EX and QFX Series Switches) | 23861
- Description | 23861
- Options | 23861
- Required Privilege Level | 23863
- Output Fields | 23863
- Sample Output | 23869
- Release Information | 23877

Syntax

```
show rsvp session
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<down | up>
<externally-provisioned>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
<lsp-type>
<name session-name>
<p2mp>
<session-type>
<statistics>
<te-link te-link>
```

Syntax (EX and QFX Series Switches)

```
show rsvp session
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<down | up>
<externally-provisioned>
<interface interface-name>
<lsp-type>
<name session-name>
<p2mp>
<session-type>
<statistics>
<te-link te-link>
```

Description

Display information about RSVP sessions.

Options

none	Display standard information about all RSVP sessions.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
bidirectional | unidirectional	(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.
bypass	(Optional) Display RSVP sessions for bypass LSPs.
down | up	(Optional) Display only LSPs that are inactive or active, respectively.
externally-provisioned	(Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).
instance *instance-name*	(Optional) Display RSVP sessions for the specified instance. If `instance-name` is omitted, RSVP session information is displayed for the master instance.
interface *interface-name*	(Optional) Display RSVP sessions for the specified interface only.

RSVP reserves resources only for outgoing LSPs of an interface. Because resources are not reserved for incoming LSPs, the `show rsvp sessions interface` *`interface-name`* command output displays only those RSVP sessions whose next hops correspond to the specified interface.

To identify the number of RSVP sessions formed over the uplink interface on the egress label-switching router (LSR), you can use the following command:

```
user@host> show rsvp session egress extensive | match "PATH rcvfrom:"
| match interface-name | count
```

logical-system (all | *logical-system-name*)

(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type

(Optional) Display information about RSVP sessions with regard to LSPs:

- `bypass`—Sessions used for bypass LSPs.
- `lsp`—Sessions used to set up LSPs.
- `nolsp`—Sessions not used to set up LSPs.

name *session-name*

(Optional) Display information about the named session.

p2mp

(Optional) Display point-to-multipoint information.

session-type

(Optional) Display information about a particular session type:

- `egress`—Sessions that terminate on this routing device.

  To identify the number of RSVP sessions formed over the uplink interface on the egress label-switching router (LSR), you can use the following command:

  ```
  user@host> show rsvp session egress extensive | match "PATH
  rcvfrom:" | match interface-name | count
  ```

- `ingress`—Sessions that originate from this routing device.
- `transit`—Sessions that transit through this routing device.

statistics

(Optional) Display packet statistics.

te-link ***te-link*** (Optional) Display sessions with reservations on the specified TE link.

Required Privilege Level

view

Output Fields

Table 1113 on page 23863 describes the output fields for the `show rsvp session` command. Output fields are listed in the approximate order in which they appear.

Table 1113: show rsvp session Output Fields

Field Name	Field Description	Level of Output
`Ingress RSVP`	Information about ingress RSVP sessions.	`detail`
`Ingress RSVP`	Information about ingress RSVP sessions. Each session has one line of output.	All levels
`Egress RSVP`	Information about egress RSVP sessions.	All levels
`Transit RSVP`	Information about the transit RSVP sessions.	All levels
`P2MP name`	(Appears only when the `p2mp` option is specified). Name of the point-to-multipoint LSP path.	All levels
`P2MP branch count`	(Appears only when the `p2mp` option is specified). Number of LSPs receiving packets from the point-to-multipoint LSP.	All levels
`To`	Destination (egress routing device) of the session.	All levels
`From`	Source (ingress routing device) of the session.	All levels
`State`	State of the path: `Up`, `Down`, or `AdminDn`. `AdminDn` indicates that the LSP is being taken down gracefully.	All levels

Table 1113: show rsvp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Address`	Destination (egress routing device) of the LSP.	`detail`
`From`	Source (ingress routing device) of the session.	`detail`
`LSPstate`	State of the LSP that is being handled by this RSVP session. It can be either `Up`, `Dn` (down), or `AdminDn`. `AdminDn` indicates that the LSP is being taken down gracefully.	`brief detail`
`Rt`	Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (`inet.0`). For transit and egress RSVP sessions, the routing table is the primary MPLS table `mpls.0`).	`brief`
`Active Route`	Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table (`inet.0`). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (`mpls.0`).	`detail`
`LSPname`	Name of the LSP.	`brief detail`
`LSPpath`	Indicates whether the RSVP session is for the primary or secondary LSP path. `LSPpath` can be either `primary` or `secondary` and can be displayed on the ingress, egress, and transit routing devices. `LSPpath` can also indicate when a graceful LSP deletion has been triggered.	`detail`
`Bypass`	(Egress routing device) Destination address for the bypass LSP.	`detail`
`Bidir`	(When LSP is bidirectional) LSP will allow data to travel in both directions between GMPLS devices.	`detail`
`Bidirectional`	(When LSP is bidirectional) LSP will allow data to travel both ways between GMPLS devices.	`detail`

Table 1113: show rsvp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Upstream label in`	(When LSP is bidirectional) Incoming label for reverse direction traffic for this LSP.	`detail`
`Upstream label out`	(When LSP is bidirectional) Outgoing label for reverse direction traffic for this LSP.	`detail`
`Recovery label received`	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	`detail`
`Recovery label sent`	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.	`detail`
`Suggested label received`	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	`detail`
`Suggested label sent`	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.	`detail`
`Resv style` or `Style`	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be `FF` (fixed filter), `SE` (shared explicit), or `WF` (wildcard filter).	`brief detail`
`Label in`	Incoming label for this LSP.	`brief detail`
`Label out`	Outgoing label for this LSP.	`brief detail`
`Time left`	Number of seconds remaining in the lifetime of the reservation.	`brief detail`
`Since`	Date and time when the RSVP session was initiated.	`detail`
`Tspec`	Sender's traffic specification, which describes the sender's traffic parameters.	`detail`

Table 1113: show rsvp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Fspec`	Indicates signaling of LSP bandwidth change with the same lsp-id completion within a stipulated time duration	`detail`
`DiffServ info`	Indicates whether the LSP is a multiclass LSP (`multiclass diffServ-TE LSP`) or a Differentiated-Services-aware traffic engineering LSP (`diffServ-TE LSP`).	`detail`
`bandwidth`	Bandwidth for each class type (`ct0`, `ct1`, `ct2`, or `ct3`).	`detail`
`Port number`	Protocol ID and sender/receiver port used in this RSVP session.	`detail`
`Attrib flags`	`Non-PHP` indicates that ultimate hop popping has been requested by the LSP using this RSVP session	extensive
`FastReroute desired`	Fast reroute has been requested by the ingress routing device.	`detail`
`Soft preemption desired`	Soft preemption has been requested by the ingress routing device.	`detail`
`FastReroute desired`	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Fast reroute (one-to-one backup) has been requested by the ingress routing device.	`detail extensive`
`Link protection desired`	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Link protection (many-to-one backup) has been requested by the ingress routing device.	`detail extensive`
`Node/Link protection desired`	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Node and link protection (many-to-one backup) has been requested by the ingress routing device.	`detail extensive`

Table 1113: show rsvp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Type`	LSP type: • `Link protected LSP`—LSP has been protected by link protection at the outgoing interface. The name of the bypass used is also listed here (`extensive`). • `Node/Link protected LSP`—LSP has been protected by node and link protection at the outgoing interface. The name of the bypass used is also listed here (`extensive`). • `Protection down`—LSP is not currently protected. • `Bypass LSP`—LSP that is used to protected one or more user LSPs in case of link failure. • `Backup LSP at Point-of-Local-Repair (PLR)`—LSP that has been temporarily established to protected a user LSP at the ingress of a failed link. • `Backup LSP at Merge Point (MP)`—LSP that has been temporarily established to protected a user LSP at the egress of a failed link.	`detail extensive`
`New bypass`	New bypass (the bypass name is also displayed) has been activated to protect the LSP.	`extensive`
`Link protection up, using` *`bypass-name`*	Link protection (the bypass name is also displayed) has been activated for the LSP.	`extensive`
`Creating backup LSP, link down`	A `link down` event occurred, and traffic is being switched over to the bypass LSP.	`extensive`
`Deleting backup LSP, protected LSP restored`	Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.	`extensive`

Table 1113: show rsvp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Path mtu`	Displays the value of the path MTU received from the network (through signaling) and the value used for forwarding. This value is only displayed on ingress routing devices with the `allow-fragmentation` statement configured at the `[edit protocols mpls path-mtu]` hierarchy level. If there is a detour LSP, the path MTU for the detour is also displayed.	`detail`
Egress protection PLR as protector	RSVP state on the Protector or the point-of-local-repair (PLR) routing device: • `Active`— Egress protection is available at the Protector or the PLR routing device. • `In Use`— Local repair has been completed.	`detail`
`PATH rcvfrom`	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor.	`detail`
`Adspec`	MTU signaled from the ingress routing device to the egress routing device by means of the adspec object.	`detail`
`PATH sentto`	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor (or peer-name in the GMPLS LSP case), and number of packets sent to the downstream routing device.	`detail`
`Explct hop` *<ip-address>* `expanded`	Displays the specific hop in the EROs that has been expanded by the router.	`extensive`
`Explct route`	Explicit route for the session. Normally this value will be the same as that of record route. Differences indicate that path rerouting has occurred, typically during fast reroute.	`detail`
`Record route`	Recorded route for the session, taken from the record route object. Normally this value will be the same as that of explct route. Differences indicate that path rerouting has occurred, typically during fast reroute.	`detail`

Sample Output

show rsvp session

```
user@host> show rsvp session
Ingress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
10.255.245.214  10.255.245.212  AdminDn  0  1 FF       -    22293 LSP Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 2 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.194  10.255.245.195  Up     0  1 FF   39811        - Gpro3-ba Bidir
10.255.245.194  10.255.245.195  Up     0  1 FF       3        - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.198  10.255.245.197  Up     0  1 SE  100000        3 pro3-de
Total 1 displayed, Up 1, Down 0
```

show rsvp session statistics

```
user@host> show rsvp session statistics
Ingress RSVP: 2 sessions
To              From           State   Packets      Bytes    LSPname
10.255.245.24   10.255.245.22     Up         0        0      pro3-bd
10.255.245.24   10.255.245.22     Up     44868    2333136  pro3-bd-2
Total 2 displayed, Up 2, Down 0
Egress RSVP: 2 sessions
To              From           State   Packets      Bytes    LSPname
10.255.245.22   10.255.245.24     Up         0          0    pro3-db
10.255.245.22   10.255.245.24     Up         0          0  pro3-db-2
Total 2 displayed, Up 2, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show rsvp session detail

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.1.1.1
  From: 10.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left:    -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt
```

show rsvp session detail (When Egress Protection Is in Standby Mode)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.1.1.1
  From: 10.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left:    -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  Egress protection PLR as protector: Active
PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt
```

show rsvp session detail (When Egress Protection Is in Effect During a Local Repair)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.1.1.1
  From: 10.2.2.2, LSPstate: Down, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left:    -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
Egress protection PLR as protector: In Use
PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt
```

show rsvp session detail (Path MTU Output Field)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.255.245.3
  From: 10.255.245.5, LSPstate: Up, ActiveRoute: 3
  LSPname: to-c, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100432
  Resv style: 1 FF, Label in: -, Label out: 100432
  Time left:    -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
  Port number: sender 1 receiver 57843 protocol 0
  FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 4470
  Path mtu: received 4470, using 4458 for forwarding
  PATH sentto: 192.168.37.89 (so-0/2/3.0) 11 pkts
  RESV rcvfrom: 192.168.37.89 (so-0/2/3.0) 10 pkts
  Explct route: 192.168.37.89
  Record route: <self> 192.168.37.89 192.168.37.87
```

```
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
    Detour adspec: sent MTU 1512
    Path mtu: received 1512, using 1500 for forwarding
```

show rsvp session detail (GMPLS)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
192.168.4.1
  From: 192.168.1.1, LSPstate: Dn, ActiveRoute: 0
  LSPname: gmpls-r1-to-r3, LSPpath: Primary
  Bidirectional, Upstream label in: 21253, Upstream label out: –
  Suggested label received: -, Suggested label sent: 21253
  Recovery label received: -, Recovery label sent: –
  Resv style: 0 –, Label in: -, Label out: –
  Time left:    -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
  Port number: sender 2 receiver 46115 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH MTU: received 0
  PATH sentto: 10.35.1.5 (so-0/2/3.0) 11 pkts
  Explct route: 10.100.100.100 10.93.93.93
  Record route: <self> 10.100.100.100 10.93.93.93
  Total 1 displayed, Up 0, Down 1
  Egress RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0
  Transit RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0
```

show rsvp session detail (Fspec Output Field)

```
user@host> show rsvp session detail
IIngress RSVP: 2 sessions

192.168.255.4
  From: 192.168.255.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R4-1, LSPpath: Primary
  LSPtype: Static Configured
```

```
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299840
  Resv style: 1 SE, Label in: -, Label out: 299840
  Time left:    -, Since: Mon May  4 21:33:34 2020
  Tspec: rate 500Mbps size 500Mbps peak Infbps m 20 M 1500
     Fspec: rate 600Mbps size 600Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 58554 protocol 0
  SelfID: RSVP-2, Route Statistics ID: RSVP-3
  SessionID: 2
  Link protection desired
  Type: Link protected LSP, using Bypass->10.1.2.18
      2 May  4 21:33:37 Link protection up, using Bypass->10.1.2.18
      1 May  4 21:33:36 New bypass Bypass->10.1.2.18
  Enhanced FRR: Enabled (Downstream), LP-MP is 192.168.255.2
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.1.2.18 (ge-0/0/1.0) 4 pkts
       outgoing message state: refreshing, Message ID: 20, Epoch: 5199652
  RESV rcvfrom: 10.1.2.18 (ge-0/0/1.0) 7 pkts, Entropy label: Yes
       incoming message handle: R-1/7, Message ID: 77, Epoch: 5203461
  Explct route: 10.1.2.18 10.2.3.3 10.3.4.20
  Record route: 10.2.2.2 (node-id) 10.1.2.18 192.168.255.3 (node-id) 10.2.3.3 192.168.255.4
(node-id) 10.3.4.20
Total 1 displayed, Up 1, Down 0
```

show rsvp session extensive

Starting in Junos OS Release 16.1, this command includes additional details for both the incoming and outgoing Path and Resv messages. The information includes the internal message handle and revision number, as well as the message ID included by the neighbor in the signaling message.

```
user@host> show rsvp session extensive
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 1 sessions
```

```
10.0.0.5
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: 1to5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299856
  Resv style: 1 SE, Label in: 299776, Label out: 299856
  Time left:  123, Since: Sat Nov 29 10:39:15 2014
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 52631 protocol 0
  PATH rcvfrom: 10.1.2.1 (ge-0/0/0.0) 2 pkts
  incoming message handle: P-1/2, ID: 0xc82fd7/322
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.2.4.4 (ge-0/0/1.0) 1 pkts
  outgoing message state: refreshing, ID: 0xcacec0/22
  RESV rcvfrom: 10.2.4.4 (ge-0/0/1.0) 1 pkts, Entropy label: Yes
  incoming message handle: R-2/1, ID: 0xc82f3e/217
  RESV
  outgoing message state: refreshing, ID: 0xcacec0/17
  Explct route: 10.2.4.4 10.99.0.5
  Record route: 10.1.2.1 <self>  10.2.4.4 10.99.0.5
  Total 1 displayed, Up 1, Down 0
```

Starting in Junos OS Release 21.2, this command displays the specific hop in the incoming Explicit Route Objects (EROs) that has been expanded by the router.

```
user@host> show rsvp session extensive
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 1 sessions
10.50.50.50
  From: 10.40.40.40, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp1-40-to-50, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299824
  Resv style: 1 SE, Label in: 299824, Label out: 299824
  Time left: 5336, Since: Mon Jun 29 20:53:06 2020
  Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 44412 protocol 0
```

```
  Link protection desired
  Type: Link protected LSP, using Bypass->10.21.31.31
      1 Jun 29 20:53:06 Link protection up, using Bypass->10.21.31.31
  Enhanced FRR: Enabled (Upstream), LP-MP for 10.40.40.40
  Enhanced FRR: Enabled (Downstream), LP-MP is 10.31.31.31
  PATH rcvfrom: 10.40.21.21 (ge-0/0/0.0) 2 pkts
       incoming message handle: P-13/2, Message ID: 31, Epoch: 430841
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.21.31.31 (ge-1/0/0.0) 3 pkts
       outgoing message state: refreshing, Message ID: 64, Epoch: 435515
  RESV rcvfrom: 10.21.31.31 (ge-1/0/0.0) 4 pkts, Entropy label: Yes
       incoming message handle: R-14/4, Message ID: 50, Epoch: 435661
  RESV
       outgoing message state: refreshing, Message ID: 63, Epoch: 435515
  Explct route: 10.21.31.31 10.28.31.28 10.28.33.33 10.23.33.23 10.28.23.23 10.23.50.50
  Explct hop 10.28.23.23 expanded
  Record route: 10.40.21.21 10.40.40.40 <self> 10.31.31.31 (node-id) 10.21.31.31 10.28.28.28
(node-id) 10.28.31.28 10.33.33.33 (node-id) 10.28.33.33 10.23.23.23 (node-id) 10.23.33.23
10.50.50.50 (node-id) 10.23.50.50

Total 1 displayed, Up 1, Down 0
```

show rsvp session extensive transit

Starting in Junos OS Release 16.1, this command also shows node-related details, including whether enhanced local protection is enabled for the LSP and whether the node is a merge point. If the latter is true, both the IP address of the Point of Local Repair (PLR) and the status (LP-MP, NP-MP, or Non-MP) are shown.

```
user@host> show rsvp session extensive transit
  From: 10.1.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: A-D-1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299776
  Resv style: 1 SE, Label in: 299776, Label out: 299776
  Time left:  117, Since: Tue May  6 08:39:44 2014
  Tspec: rate 700Mbps size 700Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 24131 protocol 0
  Node/Link protection desired
  Type: Node/Link protected LSP, using Bypass->10.2.3.3->10.3.4.4
```

```
      2 May  6 08:39:47 Node protection up, using Bypass->10.2.3.3->10.3.4.4
      1 May  6 08:39:44 New bypass Bypass->10.2.3.3->10.3.4.4
  Enhanced Local Protection: Enabled, LP-MP for 10.2.2.2, NP-MP for 10.1.1.1
  PATH rcvfrom: 10.1.2.1 (lt-0/2/0.201) 5371 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.2.3.3 (lt-0/2/0.203) 5374 pkts
  RESV rcvfrom: 10.2.3.3 (lt-0/2/0.203) 5372 pkts, Entropy label: No
  Record route: 10.1.2.1 <self>  10.3.3.3 (node-id) 10.2.3.3 10.4.4.4 (node-id) 10.3.4.4
  Total 1 displayed, Up 1, Down 0
```

If enhanced FRR is not enabled (either because it is disabled on the router itself or one of the neighbors along the LSP path does not support it), either of the following lines might be displayed:

```
Enhanced Local Protection: Disabled, Reason: User Config
Enhanced Local Protection: Disabled, Reason: Backward Compatibility
```

If enhanced FRR is not enabled and the router is not an MP, the following line is displayed:

```
Enhanced Local Protection: Enabled, Non-MP
```

show rsvp session p2mp (Ingress Router)

```
user@host> show rsvp session p2mp
Ingress RSVP: 3 sessions
P2MP name: test, P2MP branch count: 1
To              From            State   Rt Style Labelin Labelout LSPname
10.255.10.95    10.255.10.2     Up       0  1 SE    -           3 to-pe1
P2MP name: test2, P2MP branch count: 2
To              From            State   Rt Style Labelin Labelout LSPname
10.255.10.23    10.255.10.2     Up       0  1 SE    -      299776 to-pe3
10.255.10.16    10.255.10.2     Up       0  1 SE    -      299776 to-pe4
Total 3 displayed, Up 3, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show rsvp session p2mp (Transit Router)

```
user@host> show rsvp session p2mp
Ingress RSVP: 1 sessions
P2MP name: test, P2MP branch count: 1
To              From            State   Rt Style Labelin  Labelout LSPname
10.255.10.23    10.255.10.95    Up       0  1 SE    -       299792 to-pe2
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
P2MP name: test, P2MP branch count: 1
To              From            State   Rt Style Labelin Labelout  LSPname
10.255.10.95    10.255.10.2     Up       0  1 SE    3        -     to-pe1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 2 sessions
P2MP name: test2, P2MP branch count: 2
To              From            State   Rt Style Labelin Labelout  LSPname
10.255.10.23    10.255.10.2     Up       0  1 SE  299776   299808  to-pe3
10.255.10.16    10.255.10.2     Up       0  1 SE  299776   299856  to-pe4
Total 2 displayed, Up 2, Down 0
```

Release Information

Command introduced before Junos OS Release 7.4.

`externally-provisioned` option added in Junos OS Release 13.3.

`instance` option added in Junos OS Release 15.1 for the MX Series.

`Explct hop <ip-address> expanded` fields are introduced in the output field of the `show rsvp session` command in Junos OS Release 21.2R1.

`Explct hop <ip-address> expanded` fields are introduced in the output field of the `show rsvp session` command in Junos OS Evolved Release 22.1R1.

RELATED DOCUMENTATION

clear rsvp session | 3846

show rsvp statistics

IN THIS SECTION

- Syntax | 23878
- Syntax (EX Series Switches) | 23878
- Description | 23878
- Options | 23879
- Required Privilege Level | 23879
- Output Fields | 23879
- Sample Output | 23883
- Release Information | 23886

Syntax

```
show rsvp statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show rsvp statistics
```

Description

Display Resource Reservation Protocol (RSVP) packet and error statistics. The RSVP input/input module collects statistics for certain events on a per-interface basis. Most of these events were tracked on a routing-instance basis in Junos OS releases earlier than Release 17.2. The "show rsvp interface detail"

command displays these event counters under the Events section of the output only when the values of these fields are higher than zero. These statistics are also maintained at the global level (per routing-instance) and are also displayed in the output of the "show rsvp statistics" command.

Options

none	Display RSVP packet and error statistics.
instance *instance-name*	(Optional) Display RSVP packet and error statistics for the specified instance. If `instance-name` is omitted, RSVP statistics are displayed for the master instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1114 on page 23879 describes the output fields for the `show rsvp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1114: show rsvp statistics Output Fields

Field Name	Field Description
Packet Type	Statistics about different RSVP messages.
Total Sent	Total number of packets sent since RSVP was enabled.
Total Received	Total number of packets received since RSVP was enabled.
Last 5 seconds Sent	Total number of packets sent in the last 5 seconds.
Last 5 seconds Received	Number of packets received in the last 5 seconds.

Table 1114: show rsvp statistics Output Fields *(Continued)*

Field Name	Field Description
Path	Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path.
PathErr	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.
PathTear	Statistics about PathTear messages, which remove path states and dependent reservation states in any routing devices along a path.
Resv FF	Statistics about fixed-filter reservation style messages, which consist of distinct reservations among explicit senders.
Resv WF	Statistics about wildcard-filter reservation style messages, which consist of shared reservations among wildcard senders.
Res SE	Statistics about shared-explicit reservation style messages, which consist of shared reservations among explicit senders.
ResvErr	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.
ResvTear	Statistics about ResvTear messages, which remove reservation states along a path.
ResvConf	Statistics about ResvConfirm messages, which are responses to confirm a reservation request.
Ack	Acknowledge message for refresh reductions.
SRefresh	Summary refresh messages.
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.
EndtoEnd RSVP	Statistics for the number of End-to-end RSVP messages.

Table 1114: show rsvp statistics Output Fields *(Continued)*

Field Name	Field Description
Errors	Statistics about errored RSVP packets.
Rcv pkt bad length	The packet was not processed because its length is inappropriate.
Rcv pkt unknown type	The packet is not one of the well-known RSVP types, as defined in RFC 2205, *Resource ReSerVation Protocol (RSVP)*.
Rcv pkt bad version	The packet is not an RSVP version 1 packet.
Rcv pkt auth fail	The packet failed authentication checks.
Rcv pkt bad checksum	The RSVP checksum check failed.
Rcv pkt bad format	General packet processing failed because the packet was badly formed.
Memory allocation fail	An internal resource failure occurred.
No path information	A reservation was received, but no sender is active.
Resv style conflict	The same session contains inconsistent reservation styles.
Port conflict	There were inconsistent port numbers for the same session.
Resv no interface	An interface for the receive reservation packets cannot be located.
PathErr to client	Number of PathErr packets delivered to the local client.
ResvErr to client	Number of ResvErr packets delivered to the local client.

Table 1114: show rsvp statistics Output Fields *(Continued)*

Field Name	Field Description
Path timeout	Number of times the sender timed out because the path was removed.
Resv timeout	Number of times the receiver timed out because the reservation was removed.
Message out-of-order	Records the number of RSVP incoming messages that are considered out of order. This is detected from the message ID object's sequence number.
Unknown ack msg	A neighboring routing device replies with an ACK object that contains an unknown message ID. This can indicate a message ID handshake problem. For example, a router receives an ACK for message IDs 1, 2, and 3. However, it only has state for message IDs 1 and 3. The router increments the unknown ack counter by 1.
Recv nack	If a neighboring router receives an unknown message ID in an RSVP refresh message, the router sends a Resv nack message back to the sender. This can happen if that neighbor has been rebooted. For this case, the router sends a regular RSVP refresh message to recover the state and start the message-ID handshake process again.
Recv duplicated msg-id	Number of times the same message ID is used by two different RSVP messages. This duplication is usually caused when a neighboring routing device restarts.
No TE-link to recv Hop	Counter of packets discarded because a TE link was not found.
Rcv pkt disabled interface	Number of RSVP packets received on an interface that is not enabled for RSVP.
Transmit buffer full	Number of times the buffer for assembling an outgoing RSVP message was not large enough.
Transmit failure	Number of times the RSVP task failed to send out a packet.
Receive failure	Number of times the RSVP task failed to read an incoming packet.

Table 1114: show rsvp statistics Output Fields *(Continued)*

Field Name	Field Description
P2MP RESV discarded by appl	Number of Resv messages discarded because the MPLS label is not valid for the P2MP LSP application.
Rate limit	Number of RSVP packets dropped due to rate limiting.
Err msg loop detected	Number of RSVP error messages that have looped back to their originator. This is detected by checking the error node address in the ERROR_SPEC object.

Sample Output

show rsvp statistics

Starting in Junos OS Release 16.1, this command also shows conditional PathTear statistics and the number of times an LSP state has been retained because of Link Protecting Merge Point (LP-MP) or Node Protecting Merge Point (NP-MP) determination.

```
user@host> show rsvp statistics
   PacketType              Total                  Last 5 seconds
                     Sent      Received        Sent      Received
   Path               355           408           0             0
   PathErr              2            13           0             0
   PathTear           101           139           0             0
   Resv FF              0             0           0             0
   Resv WF              0             0           0             0
   Resv SE            419           225           0             0
   ResvErr              0             0           0             0
   ResvTear             0            13           0             0
   ResvConf             0             0           0             0
   Bundle             455           378           0             0
   Ack                682          1414           0             0
   SRefresh        395198        236030           5             2
   Hello           578809        578221           4             4
   EndtoEnd RSVP        0             0           0             0
   Node Hello          50            50           0             0
   PathTear(Condl.)     0             3           0             0
```

```
Errors                          Total            Last 5 seconds
Rcv pkt bad length                0                         0
Rcv pkt unknown type              0                         0
Rcv pkt bad version               0                         0
Rcv pkt auth fail                 0                         0
Rcv pkt bad checksum              0                         0
Rcv pkt bad format                0                         0
Memory allocation fail            0                         0
No path information               10                         0
Resv style conflict               0                         0
Port conflict                     0                         0
Resv no interface                 0                         0
PathErr to client                38                         0
ResvErr to client                 0                         0
Path timeout                      8                         0
Resv timeout                     57                         0
Message out-of-order              0                         0
Unknown ack msg                2978                         0
Recv nack                        86                         0
Recv duplicated msg-id            5                         0
No TE-link to recv Hop            0                         0
Rcv pkt disabled interface        0                         0
Transmit buffer full              0                         0
Transmit failure                  0                         0
Receive failure                   0                         0
P2MP RESV discarded by appl       0                         0
Rate limit                      306                         0
Err msg loop detected             0                         0
MP Path LP-avail rcved             0                         0
MP Path NP-avail rcved             0                         0
PLR bk RSB life ext                0                         0
MP bk PSB life ext                 0                         0
LP-MP state retained on failure    0                         0
NP-MP state retained on failure    0                         0
Fast refresh skipped               0                         0
MP bk Srefresh Nack rcved          0                         0
RSB life extended for nh FRR       0                         0
```

show rsvp statistics

```
user@host> show rsvp statistics
PacketType              Total                  Last 5 seconds
                    Sent      Received        Sent      Received
  Path                21             0           0             0
  PathErr              0             4           0             0
  PathTear             9             0           0             0
  Resv                 0             9           0             0
  ResvErr              0             0           0             0
  ResvTear             0             2           0             0
  ResvConf             0             0           0             0
  Bundle              28             2           0             0
  Hello           172814        172802           5             5
  Ack                 11            12           0             0
  Srefresh           142           143           0             0
  Notify               0             0           0             0
  Unknown              0             0           0             0
  EndtoEnd RSVP        0             0           0             0
  Backup Path          0             0           0             0
  Backup Tear          0             0           0             0
  Cnd PathTear         0             0           0             0
  Rmt PathTear         0             0           0             0
  Rmt Backup           0             0           0             0

  Errors                          Total            Last 5 seconds
  Rcv pkt bad length                 0                         0
  Rcv pkt unknown type               0                         0
  Rcv pkt bad version                0                         0
  Rcv pkt auth fail                  0                         0
  Rcv pkt bad checksum               0                         0
  Rcv pkt bad format                 0                         0
  Message out-of-order               0                         0
  Unknown msg-id ack                 0                         0
  Unknown msg-id nack                0                         0
  Rcv msg-id nack                    0                         0
  Rcv pkt disabled interface         0                         0
  Transmit failure                   3                         0
  Memory allocation fail             0                         0
  ID allocation failed               0                         0
  No path information                0                         0
  Resv style conflict                0                         0
```

```
    Port conflict                     0                        0
    Resv no interface                 0                        0
    PathErr to client                 4                        0
    ResvErr to client                 0                        0
    Path timeout                      0                        0
    Resv timeout                      0                        0
    No TE-link to recv Hop            0                        0
    Transmit buffer full              0                        0
    P2MP RESV discarded by app        0                        0
    Rate limit                        0                        0
    Err msg loop detected             0                        0
    MP Path LP-avail rcved            0                        0
    MP Path NP-avail rcved            0                        0
    PLR bk RSB life ext               0                        0
    RSB life ext for nh FRR           0                        0
    MP pri PSB life ext LP            0                        0
    Rcvd state rejected               0                        0
    No matching senders               0                        0
    Del from client                   5                        0

    Enhanced FRR Stats              Total           Last 5 seconds
    LP-MP LSPs retained               0                        0
    NP-MP LSPs retained               0                        0
    Non-MP LSPs deleted               0                        0
    LSPs deleted on Phop down         0                        0
    LSPs deleted on PPhop down        0                        0
    LP avail signaled LSPs            0                        0
    NP avail signaled LSPs            0                        0
    NP flag reset for Phop            0                        0
    LSPs retained on Cnd tear         0                        0
    Upstr long refresh LSPs           0                        0
    Upstr short refresh LSPs          0                        0
    Dnstr long refresh LSPs           6                        0
    Dnstr short refresh LSPs          0                        0
    PathTear ignored on MP            0                        0
    RRO change Remote PathTear        0                        0
    Primary down Rmt PathTear         0                        0
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` option added in Junos OS Release 15.1 for the MX Series.

RELATED DOCUMENTATION

clear rsvp statistics | 3848

show rsvp version

IN THIS SECTION

- Syntax | 23887
- Syntax (EX Series Switches) | 23887
- Description | 23888
- Options | 23888
- Required Privilege Level | 23888
- Output Fields | 23888
- Sample Output | 23891
- Release Information | 23891

Syntax

```
show rsvp version
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show rsvp version
```

Description

Display information about the Resource Reservation Protocol (RSVP) protocol settings, such as the version of the RSVP software, the refresh timer and keep multiplier, and local RSVP graceful restart capabilities on a routing device.

Options

none	Display RSVP protocol settings.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1115 on page 23888 describes the output fields for the `show rsvp version` command. Output fields are listed in the approximate order in which they appear.

Table 1115: show rsvp version Output Fields

Field Name	Field Description
Resource ReSerVation Protocol, version	RSVP software version.
RSVP protocol	Status of RSVP: **Enabled** or **Disabled**.
R(refresh timer)	Configured time interval used to generate periodic RSVP messages.
K(keep multiplier)	Number of RSVP messages that can be lost before an RSVP state is declared stale.
Preemption	Currently configured preemption capability: **Aggressive**, **Disabled**, or **Normal**. The default is **Normal**.

Table 1115: show rsvp version Output Fields *(Continued)*

Field Name	Field Description
Soft-preemption cleanup	Time, in seconds, that an LSP is kept after it has been soft preempted. This is a global property of the RSVP protocol.
Graceful deleting timeout	Currently configured value for the `graceful-deletion-timeout` statement. The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down.
NSR Mode	Status of the nonstop active routing feature for RSVP on the restarting device: **Disabled**, **Enabled/Master**, or **Enabled/Standby**.
NSR State	State of the nonstop active routing feature for RSVP on the restarting device. Possible values are: • **Idle** • **TE-link sync complete** • **Neighbor sync complete** • **Path state sync complete** • **Resv state sync complete** • **Bypass sync complete** • **Init sync complete**
Setup protection	Status of point-to-point and point-to-multipoint LSP setup protection configuration on the device: **Enabled** or **Disabled**.
Route Session-Id count	Total count of session IDs associated with the combination of all the RSVP ingress routes. **NOTE**: Starting in Junos OS Release 16.1, the `show rsvp version` command output displays the `Route Session-Id count` output field by default, irrespective of the presence of associated session IDs. When there are no session IDs associated with any RSVP ingress route, the *Route Session-Id count* value is zero (0).

Table 1115: show rsvp version Output Fields *(Continued)*

Field Name	Field Description
Graceful restart	Status of the graceful restart feature for RSVP on the restarting routing device: **Enabled** or **Disabled**.
Restart helper mode	Status of the helper mode feature: **Enabled** or **Disabled**. When this feature is enabled, the restarting routing device can help the neighbor with its RSVP restart procedures.
Maximum helper restart time	Number of milliseconds (ms) configured for the maximum helper restart time. The maximum helper restart time is the length of time the routing device waits before declaring that an RSVP neighbor attempting to restart gracefully is down.
Maximum helper recovery time	Number of milliseconds configured for the maximum helper recovery time. The maximum helper recovery time is the amount of time the routing device maintains the state of an RSVP neighbor attempting to restart gracefully.
Restart time	Number of milliseconds that a neighbor waits to receive a Hello message from the restarting node before declaring the node dead and deleting the states.
Recovery time	Number of milliseconds during which the restarting node attempts to recover its lost states with help from its neighbors. Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.
P2p transit LSP nexthop mode	Point-to-point transit LSP next-hop mode on PTX Series devices. The possible values are **Chained** or **Unchained**.
P2mp transit LSP nexthop mode	Point-to-multipoint transit LSP next-hop mode on PTX Series devices. The possible values are **Chained** or **Unchained**.

Sample Output

show rsvp version

Starting with Junos OS Release 16.1, this command also shows whether enhanced FRR procurers are enabled on the router.

```
user@host> show rsvp version
Resource ReSerVation Protocol, version 1. rfc2205
   RSVP protocol:                Enabled
   R(refresh timer):             30 seconds
   K(keep multiplier):           3
   Preemption:                   Normal
   Soft-preemption cleanup:      30 seconds
   Graceful deletion timeout:    30 seconds
   NSR mode:                     Enabled/Master
   NSR state:                    Init sync complete
   Setup protection:             Disabled
   Route Session-Id count:       1
   Graceful restart:             Disabled
   Restart helper mode:          Enabled
   Maximum helper restart time:  20000 msec
   Maximum helper recovery time: 180000 msec
   Restart time:                 0 msec
   P2p transit LSP nexthop mode: Unchained
   P2mp transit LSP nexthop mode: Unchained
   Enhanced FRR local protection: Enabled
```

Release Information

Command introduced before Junos OS Release 7.4.

show sap listen

IN THIS SECTION

- Syntax | 23892
- Description | 23892
- Options | 23892
- Required Privilege Level | 23893
- Output Fields | 23893
- Sample Output | 23893
- Release Information | 23894

Syntax

```
show sap listen
<brief | detail>
<logical-system (all | logical-system-name)>
```

Description

Display the addresses that the router is listening to in order to receive multicast Session Announcement Protocol (SAP) session announcements.

Options

none	Display standard information about the addresses that the router is listening to in order to receive multicast SAP session announcements.
brief | detail	(Optional) Display the specified level of output.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1116 on page 23893 describes the output fields for the `show sap listen` command. Output fields are listed in the approximate order in which they appear.

Table 1116: show sap listen Output Fields

Field Name	Field Description
Group address	Address of the group that the local router is listening to for SAP messages.
Port	UDP port number used for SAP.

Sample Output

show sap listen

```
user@host> show sap listen
Group address   Port
224.2.127.254   9875
239.255.255.255 9875
```

show sap listen brief

The output for the `show sap listen brief` command is identical to that for the `show sap listen` command. For sample output, see "show sap listen" on page 23893.

show sap listen detail

The output for the `show sap listen detail` command is identical to that for the `show sap listen` command. For sample output, see "show sap listen" on page 23893.

Release Information

Command introduced before Junos OS Release 7.4.

show security advance-policy-based-routing detail

IN THIS SECTION

- Syntax | 23894
- Description | 23894
- Options | 23895
- Required Privilege Level | 23895
- Output Fields | 23895
- Sample Output | 23897
- Sample Output | 23897
- Release Information | 23898

Syntax

```
show security advance-policy-based-routing detail
```

Description

Display a summary of all APBR policies configured on the device.

You can use this command to understand the details of an APBR policy such as:

- Name, status, zone-context of the APBR policy.
- The number of times traffic matches the APBR policy and the APBR profile is applied for a session.

You can use the `show security advance-policy-based-routing detail` command only when you have configured advanced policy-based routing (APBR) profile with the APBR policy. You might not get the accurate results if you have an APBR profile attached the security zone.

Options

count Display the number of configured APBR policies.

- **Range:** 1 to 65535

detail Display a detailed view of all of the APBR policies configured on the device.

from-zone Display specific zone details applicable to the APBR policy.

logical-system Displaythe logical system name.

root-logical-system Display information about the default root-logical-system.

start Display the policy from the given position.

- **Range:** 1 to 65535

Required Privilege Level

view

Output Fields

Table 1117 on page 23895 lists the output fields for the `show security advance-policy-based-routing detail` command. Output fields are listed in the approximate order in which they appear.

Table 1117: show security advance-policy-based-routing statistics

Field Name	Field Description
Policy name	Name of the APBR policy
Enabled	Status of the policy (enabled or disabled)
Policy type	Type of the policy.
Index	An internal number associated with the policy.

Table 1117: show security advance-policy-based-routing statistics *(Continued)*

Field Name	Field Description
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zone A-to-zone B context might be ordered with sequence numbers 1, 2, and 3. Also, in a from-zone C-to-zone D context, four policies might have sequence numbers 1, 2, 3, and 4.
From zone	The zone on which APBR profile is applied to.
Source addresses	The names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	The names and corresponding IP addresses of the destination addresses (or address sets) for a policy as entered in the destination zone's address book. A packet's destination address must match one of these addresses for the policy to apply to it.
Application	Name of a preconfigured or custom application, or any if no application is specified.
ALG	If an ALG is associated with the session, the name of the ALG. Otherwise, 0.
protocol	Protocol name or numeric value of the traffic.
Inactivity timeout	Elapsed time without activity after which the application is terminated.
Source port range	Range of matching source ports defined in the policy.
Destination port range	Range of matching destination ports defined in the policy.
APBR-Profile	Name of the APBR profile
Source identities	User details specified in the source-identity field of the named policy.

Table 1117: show security advance-policy-based-routing statistics *(Continued)*

Field Name	Field Description
Scheduler name	Name of the scheduler associated with APBR policy.

Sample Output

show security advance-policy-based-routing statistics

```
user@host>  show security advance-policy-based-routing detail
Policy: SLA1, State: enabled, Index: 5
Policy Type: Configured
Sequence number: 1
From zone: trust
Source addresses:
any-ipv4(global): 0.0.0.0/0
any-ipv6(global): ::/0
Destination addresses:
any-ipv4(global): 0.0.0.0/0
any-ipv6(global): ::/0
Application: any
IP protocol: 0, ALG: 0, Inactivity timeout: 0
Source port range: [0-0]
Destination port range: [0-0]
APBR-Profile: profile1
Scheduler name: scheduler-1
```

Sample Output

show security advanced-policy-based-routing detail (Junos OS Release 19.1R1)

```
user@host> show security advanced-policy-based-routing detail
Policy: p1, State: enabled, Index: 4
  Sequence number: 1
  From zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
```

```
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  APBR Profile: apbr-pr1

Source identities:
      dev_user
```

Release Information

Command introduced in Junos OS Release 15.1X49-D60. The option scheduler is added in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

> *Example: Configuring Advanced Policy-Based Routing for Application-Aware Traffic Management Solution*

show security advance policy based routing multipath policy statistics

IN THIS SECTION

- Syntax | 23899
- Description | 23899
- Required Privilege Level | 23899
- Output Fields | 23899
- Sample Output | 23901
- Release Information | 23901

Syntax

```
show security advance-policy-based-routing multipath policy statistics
```

Description

Display the multipath policy statistics.

Required Privilege Level

view

Output Fields

Table 1118 on page 23899 lists the output fields for the `show security advance policy based routing multipath policy statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1118: show security advance policy based routing multipath policy statistics

Field Name	Field Description
Multipath Policies Information:	
Total policies	Total number of multipath policies
Policy ID	Policy ID number
State	State of the rule
Total-Sessions	Total number of sessions
Rule	Multipath rule
Application	Application name
Multipath rule	Multipath rule name
Statistics	

Table 1118: show security advance policy based routing multipath policy statistics *(Continued)*

Field Name	Field Description
Current Sessions	Number of sessions active currently
Ignored Sessions	Number of sessions ignored
Applications Matched	Number of sessions where application traffic matched
Applications Switched	Number of sessions where application traffic changed after the final match
Stopped due to Bandwidth Limit	Number of sessions stopped because of bandwidth limits.
Packets in path inactive state	Number of packets transferred when the rule is not active.
Packets in path active state	Number of packets transferred when the rule is active
Less than Configured Paths	Number of paths not configured.
Midstream Packets Ignored	Number of sessions ignored in midstream
Total Packets Processed	Total number of packets processed
Total Packets Copied	Total number of packets duplicated
SLA Violated	Number of times SLA was violated
SLA Met	Number of times SLA was met

Sample Output

show security advance-policy-based-routing multipath policy statistics

```
user@host> show security advance-policy-based-routing multipath policy
statistics
 Multipath Policies Information:
Total policies: 6
Policy-id: 1, State: config-no-sla, Total-Sessions: 0,
Rule: 131072, Application: junos:YAHOO, Multipath rule: multi1
Statistics:
Current Sessions 0
Ignored Sessions 0
Applications Matched 0
Applications Switched 0
Stopped due to Bandwidth Limit 0
Packets in path inactive state 0
Packets in path active state 0
Less than Configured Paths 0
Midstream Packets Ignored 0
Total Packets Processed 0
Total Packets Copied 0
SLA Violated 0
SLA Met 0
```

Release Information

Command introduced in Junos OS Release 21.2R1.

show security advance-policy-based-routing multipath rule

IN THIS SECTION

- Syntax | 23902
- Description | 23902
- Required Privilege Level | 23902

- Output Fields | 23902
- Sample Output | 23910
- Release Information | 23913

Syntax

```
show security advance-policy-based-routing multipath rule rule-name
```

Description

Display the multipath rule details, status , and statistics.

Required Privilege Level

view

Output Fields

Table 1119 on page 23902 lists the output fields for the `show security advance-policy-based-routing multipath rule` command. Output fields are listed in the approximate order in which they appear.

Table 1119: show security advance-policy-based-routing multipath rule Output Fields

Field Name	Field Description
Multipath Rule Information	
Multipath rule name	Name of the multipath rule
Multipath rule type	Type of multipath rule. Multipath at forward direction or multipath at reverse direction
Multipath rule state	State of the rule (active-config or active-dynamic)
Configured number of paths	Number of paths configured in the multipath rule

Table 1119: show security advance-policy-based-routing multipath rule Output Fields *(Continued)*

Field Name	Field Description
Configured applications	Application configured in the multipath rule
Path Group Information	
Total path groups	Destination groups. You can group all the overlay paths terminating at the same destination under a destination group
Path-Group-Id	ID of the path group
State	State of the paths (if active or not)
Avl-Num-Paths	Available number of paths
Statistics	
Current Sessions	Number of sessions active currently
Ignored Sessions	Number of sessions ignored
Applications Matched	Number of sessions where application traffic matched
Applications Switched	Number of sessions where application traffic changed after the final match
Stopped due to Bandwidth Limit	Number of sessions stopped because of bandwidth limits.
Packets in path inactive state	Number of packets transferred when the rule is not active.
Packets in path active state	Number of packets transferred when the rule is active
Less than Configured Paths	Number of paths not configured.
Midstream Packets Ignored	Number of sessions ignored in midstream

Table 1119: show security advance-policy-based-routing multipath rule Output Fields *(Continued)*

Field Name	Field Description
Total Packets Processed	Total number of packets processed
Total Packets Copied	Total number of packets duplicated
SLA Violated	Number of times SLA was violated
SLA Met	Number of times SLA was met
Status	
Policy reference count	Current sessions count on the policy.
Credit Limit (Mbps)	Credit limit. The policers begin to drop packets when the current credit exceeds the limit. Total credit available in Mbs per second.
Policer Rate (Kbits per ms)	Policer rate. Total credit available in Kbs per millisecond.
Bandwidth Limit	Configured bandwidth limit.
Maximum Wait Time (secs)	Configured maximum wait time
Time to Reinforce (secs)	Time after which the multipath routing resumes when bandwidth limit has reached.
Application Hit List	Matched application list
Sender Path Groups Information:	
Total sender path groups	Path groups from sending device
Path-Group-Id	Identification number of the paths
Cur-Num-Paths	Current number of paths

Table 1119: show security advance-policy-based-routing multipath rule Output Fields *(Continued)*

Field Name	Field Description
Path Information	
Src-IP	Source IP address
Pkts-Sent	Packets sent for destination IP address
Link-type	Link type (MPLS or IP) selected for destination IP address
Sender Application Groups Information	
Total sender applications	Number of application traffic sessions
Application Name:	Name of application
Application Id	Application ID
Pkts-Sent	Number of application traffic packet sent
Path Information	
Src-IP	Source IP address
Pkts-Sent	Packets sent for destination IP address
Link-type	Link type (MPLS or IP) selected for destination IP address
Sender Cos Q Statistics	
Total sender cos queues	Total number CoS queues.
Pkts-Sent	Packets sent in the CoS queue ID for the entire multipath rule.
Path Information	

Table 1119: show security advance-policy-based-routing multipath rule Output Fields *(Continued)*

Field Name	Field Description
Src-IP	Source IP address
Pkts-Sent	Packets sent for destination IP address
Link-type	Link type (MPLS or IP) selected for destination IP address

Table 1120: show security advance-policy-based-routing multipath rule (Device Configured to Receive Traffic)

Field Name	Field Description
Multipath Rule Information	
Multipath rule name	Name of the multipath rule
Multipath rule type	Type of multipath rule. Multipath at forward direction or multipath at reverse direction
Multipath rule state	State of the rule (active-config or active-dynamic)
Configured number of paths	Number of paths configured in the multipath rule
Configured applications	Application configured in the multipath rule
Path Group Information	
Total path groups	Destination groups. You can group all the overlay paths terminating at the same destination under a destination group
Path-Group-Id	ID of the path group
State	State of the paths (if active or not)
Avl-Num-Paths	Available number of paths

Table 1120: show security advance-policy-based-routing multipath rule (Device Configured to Receive Traffic) *(Continued)*

Field Name	Field Description
Path Groups Information	
Total path groups	Destination groups. You can group all the overlay paths terminating at the same destination under a destination group
Path-Group-Id	Identification number of the paths
State	State of the paths (if active or not)
Avl-Num-Paths	Available number of routes for the destination path groups
Receiver Information (Path Groups Information)	
Total receiver path groups	Total number of path group
id	Identification number of path group
Cur-Num-Paths:	Number of available paths
Pkts-Rcvd	Number of packets received
Out-of-seq pkt drop:	Number of out-of-squence packets dropped
Avg-Pkt-Loss	Average packet loss for the combined links
Avg-Ingress-Jitter(us)	Average ingress jitter for the combined links
Path Information	
Dst-IP	Destination IP address
Pkt-Loss(%)	Percentage of packet lost in this link

Table 1120: show security advance-policy-based-routing multipath rule (Device Configured to Receive Traffic) *(Continued)*

Field Name	Field Description
Pkts-Rcvd	Number of packets received
Ingress-Jitter(us)	Ingress jitter experienced in the link
Reduction-Pkt-Loss(%)	Reduction in packet loss over the link (because of multipath routing)
Reduction-Ingress-Jitter(us)	Reduction in ingress jitter over the link (because of multipath routing)
Receiver Application Groups Information:	
Total receiver applications	Number of application traffic received
id	Application ID
Name	Application Name
Pkts-Rcvd	Number of packets received
Avg-Pkt-Loss	Average packet loss for the combined links
Avg-Ingress-Jitter(us)	Average ingress jitter for the combined links
Out-of-seq pkt drop:	Number of out-of-squence packets dropped
Path Information	
Dst-IP	Destination IP address
Pkts-Rcvd	Number of packets received
Pkt-Loss(%)	Percentage of packet lost in this link
Ingress-Jitter(us)	Ingress jitter experienced in the link

Table 1120: show security advance-policy-based-routing multipath rule (Device Configured to Receive Traffic) *(Continued)*

Field Name	Field Description
Reduction-Pkt-Loss(%)	Reduction in packet loss over the link (because of multipath routing)
Reduction-Ingress-Jitter(us)	Reduction in ingress jitter over the link (because of multipath routing)
Cos Q Statistics	
Total receiver cos queues	Total number CoS queues.
Id	CoS queue identification number.
Name	Cos queue group name
Pkts-Rcvd	Packets sent in the CoS queue ID for the entire multipath rule.
Ingress-Jitter(us)	Ingress jitter experienced in the link
Reduction-Pkt-Loss(%)	Reduction in packet loss over the link (because of multipath routing)
Reduction-Ingress-Jitter(us)	Reduction in ingress jitter over the link (because of multipath routing)
Number of out-of-sync packets dropped in the queue	Packets sent in the CoS queue ID for the entire multipath rule.
Path Information	
Dst-IP	Destination IP address
Pkts-Rcvd	Number of packets received
Pkt-Loss(%)	Percentage of packet lost in this link

Table 1120: show security advance-policy-based-routing multipath rule (Device Configured to Receive Traffic) *(Continued)*

Field Name	Field Description
Ingress-Jitter(us)	Ingress jitter experienced in the link
Reduction-Pkt-Loss(%)	Reduction in packet loss over the link (because of multipath routing)
Reduction-Ingress-Jitter(us)	Reduction in ingress jitter over the link (because of multipath routing)

Sample Output

show security advance-policy-based-routing multipath rule (For Spoke Device) (Junos OS Release 21.2R1)

```
user@host> show security advance-policy-based-routing multipath rule multi1
       Multipath Rule Status:
  Multipath Rule Information:
    Logical system                    root-logical-system
    Multipath rule name               amr-rule1
    Multipath rule type               Packet-copy
    Multipath rule state              Active
    Configured number of paths        2
    Configured applications           junos:FTP, junos:HTTP, junos:SSH, junos:YAHOO
  Path Group Information:
    Total path groups : 1
    Path-Group-Id  State              Avl-Num-Paths
    2234384386     Active             3
    Statistics:
        Current Sessions                  0
        Ignored Sessions                  0
        Applications Matched              3
        Applications Switched             0
        Stopped due to Bandwidth Limit    0
        Packets in path inactive state    8
        Packets in path active state      15457
        Less than Configured Paths        0
        Midstream Packets Ignored         0
```

```
        Total Packets Processed            15465
        Total Packets Copied               15455
        SLA Violated                       0
        SLA Met                            0
    Status:
      Policy reference count             4
      Credit Limit (Mbps)                300
      Policer Rate (Kbits per ms)        0
      Bandwidth Limit                    Not-Reached
      Maximum Wait Time (secs)           60
      Time to Reinforce (secs)           0
      Application Hit List               junos:HTTP, junos:SSH
Sender Path Groups Information:
      Total sender path groups : 1
      Path-Group-Id : 2234384386, Cur-Num-Paths: 2
        Path Information:
          Src-IP          Pkts-Sent       Link-type
          40.1.1.2        15455          IP
          41.1.1.2        15011          IP
          42.1.1.2        444            IP
        Sender Application Groups Information:
          Total sender applications: 2
            Application Name: junos:HTTP, Application Id: 67, Pkts-Sent: 15427
              Path Information:
                Src-IP          Pkts-Sent       Link-type
                40.1.1.2        15427          IP
                41.1.1.2        14983          IP
                42.1.1.2        444            IP
            Application Name: junos:SSH, Application Id: 198, Pkts-Sent: 28
              Path Information:
                Src-IP          Pkts-Sent       Link-type
                40.1.1.2        28             IP
                41.1.1.2        28             IP
        Sender Cosq Groups Information:
          Total sender cosq's: 1
            Cos Q : 0, Pkts-Sent: 15455
              Path Information:
                Src-IP          Pkts-Sent       Link-type
                40.1.1.2        15455          IP
                41.1.1.2        15011          IP
                42.1.1.2        444            IP
```

show security advance-policy-based-routing multipath rule (Device Configured to Receive Traffic) (Junos OS Release 21.2R1)

```
user@host> show security advance-policy-based-routing multipath rule
multi1
 Multipath Rule Status:
  Multipath Rule Information:
    Logical system                    root-logical-system
    Multipath rule name               amr-rule1
    Multipath rule type               Packet-copy
    Multipath rule state              Active
    Configured number of paths        2
    Configured applications           junos:FTP, junos:HTTP, junos:SSH, junos:YAHOO
  Path Group Information:
    Total path groups : 1
    Path-Group-Id  State              Avl-Num-Paths
    2234384386     Active             3
    Receiver Path Groups Information:
      Total receiver path groups : 1
      Id: 2234384386, Cur-Num-Paths: 2, Pkts-Rcvd: 15411, Avg-Pkt-Loss(%): 0, Avg-Ingress-
Jitter(us): 73, Out-of-seq pkt drop: 15411
      Path Information:
        Dst-IP          Pkts-Rcvd      Pkt-Loss(%)    Ingress-Jitter(us)       Reduction-Pkt-
Loss(%)    Reduction-Ingress-Jitter(us)
        40.1.1.1        15412          0              69
0                        -4
        41.1.1.1        14967          0              255
0                        182
        42.1.1.1        443            0              111
0                        38
        Receiver Application Groups Information:
          Total receiver applications: 2
            Id: 67, Name: junos:HTTP, Pkts-Rcvd: 15383, Avg-Pkt-Loss(%): 0, Avg-Ingress-
Jitter(us): 73, Out-of-seq-drop: 15383
             Path Information:
                Dst-IP          Pkts-Rcvd      Pkt-Loss(%)    Ingress-Jitter(us)       Reduction-
Pkt-Loss(%)    Reduction-Ingress-Jitter(us)
                40.1.1.1        15384          0              69
```

```
0                       -4
                41.1.1.1        14939          0              62
0                       -11
                42.1.1.1        443            0              111
0                       38
            Id: 198, Name: junos:SSH, Pkts-Rcvd: 28, Avg-Pkt-Loss(%): 0, Avg-Ingress-Jitter(us):
178, Out-of-seq-drop: 28
             Path Information:
                Dst-IP          Pkts-Rcvd      Pkt-Loss(%)    Ingress-Jitter(us)       Reduction-
Pkt-Loss(%)    Reduction-Ingress-Jitter(us)
                40.1.1.1        28             0              265
0                       87
                41.1.1.1        28             0              253
0                       75
        Receiver Cosq Groups Information:
          Total receiver cosq's: 1
            Id: 0, Name: CosQ-0, Pkts-Rcvd: 15411, Avg-Pkt-Loss(%): 0, Avg-Ingress-Jitter(us):
73, Out-of-seq-drop: 15411
             Path Information:
                Dst-IP          Pkts-Rcvd      Pkt-Loss(%)    Ingress-Jitter(us)       Reduction-
Pkt-Loss(%)    Reduction-Ingress-Jitter(us)
                40.1.1.1        15412          0              69
0                       -4
                41.1.1.1        14967          0              255
0                       182
                42.1.1.1        443            0              111
0                       38
```

Release Information

Command introduced in Junos OS Release 15.1X49-D160.

show security advance-policy-based-routing profile

IN THIS SECTION

Syntax | 23914

- Description | 23914
- Required Privilege Level | 23914
- Output Fields | 23914
- Sample Output | 23915
- Release Information | 23915

Syntax

```
show security advance-policy-based-routing profile
```

Description

Display the advanced policy-based routing (APBR) profile-to-zone mapping.

Required Privilege Level

view

Output Fields

Table 1121 on page 23914 lists the output fields for the `show security advance-policy-based-routing profile` command. Output fields are listed in the approximate order in which they appear.

Table 1121: show security advance-policy-based-routing profile

Field Name	Field Description
`pic`	PIC number of the accumulated statistics. **NOTE**: The PIC number is always displayed as 0 for SRX300, SRX320, SRX340, SRX345, SRX550M, and SRX1500 devices.
`Profile`	The name of the advanced policy-based (APBR) routing profile.

Table 1121: show security advance-policy-based-routing profile *(Continued)*

Field Name	Field Description
Zone	The zone on which APBR profile is applied to.

Sample Output

show security advance-policy-based-routing profile

```
user@host> show security advance-policy-based-routing profile

pic: 0/0
Profile    Zone
Profile1  trust
```

Release Information

Command introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Example: Configuring Advanced Policy-Based Routing for Application-Aware Traffic Management Solution

show security advance-policy-based-routing sla active-probe-statistics

IN THIS SECTION

- Syntax | 23916
- Description | 23916
- Required Privilege Level | 23916

- Output Fields | 23916
- Sample Output | 23917
- Release Information | 23918

Syntax

```
show security advance-policy-based-routing sla active-probe-statistics  active-probe-params-name
probe-name
```

Description

Displays the details of active probe parameters. Active probe parameters are used by AppQoE to evaluate the SLA of the link. In the active probing, custom packets are sent between a spoke device and a hub device on multiple routes to measure RTT, jitter, and packet loss between two SRX Series Firewalls. The active probes are sent periodically on all the active and passive links. A configured number of samples is collected and a running average for each such application's probe path is measured. If there is a violation detected for any application traffic, the probe metrics are evaluated to determine the best link that satisfies the SLA.

Required Privilege Level

view

Output Fields

Table 1122 on page 23916 lists the output fields for the `show` command. Output fields are listed in the approximate order in which they appear.

Table 1122: show security advance-policy-based-routing sla active-probe-statistics

Field Name	Field Description
`Src-IP`	Probe IP addresses used as probes' start point.

Table 1122: show security advance-policy-based-routing sla active-probe-statistics *(Continued)*

Field Name	Field Description
Dst-IP	Probe IP addresses used as probes' end point.
PKT-LOSS	Percentage of number of packets lost.
RTT(us)	Round-trip time (in microseconds)
2way-Jit	Two-way jitter (in microseconds).
Ing-Jit	Ingress jitter (in microseconds).
Egr-Jit	Egress jitter (in microseconds).
Outgoing-IP	Outgoing IP address for the particular probe.

Sample Output

show security advance-policy-based-routing sla active-probe-statistics (Junos OS 21.3R1)

```
user@host> show security advance-policy-based-routing sla active-probe-statistics active-probe-
params-name probe-name
Active Probe Statistics:

Src-IP   Dst-IP   PKT-LOSS(%) RTT(us) 2way-Jit(us) Ing-Jit(us) Egr-Jit(us) Outgoing-IP
4000:5::1 4000:5::2       0     2808       538         120         315     4000:5::1
4001:5::1 4001:5::2       0     4367       360         203         140     4000:5::2
```

show security advance-policy-based-routing sla active-probe-statistics (Junos OS 20.4R1)

```
user@host> show security advance-policy-based-routing sla active-probe-statistics active-probe-
params-name probe-name
Active Probe Statistics:
```

```
Src-IP   Dst-IP   PKT-LOSS(%) RTT(us) 2way-Jit(us) Ing-Jit(us) Egr-Jit(us) Outgoing-IP
19.0.0.1 119.0.0.2       0        2808    538      4294967295  4294967295  170.0.0.2
19.0.0.1 119.0.0.2       0        4367    360      4294967295  4294967295  42.1.1.2
19.0.0.1 119.0.0.2.      0        2620    307      4294967295  4294967295  180.0.0.2
19.0.0.1 119.0.0.2       0        4209    429      4294967295  4294967295  41.1.1.2
19.0.0.1 119.0.0.2       0        4265    691      4294967295  4294967295  40.1.1.2
19.0.0.1 119.0.0.2       0        2718    409      4294967295  4294967295  160.0.0.2
```

show security advance-policy-based-routing sla active-probe-statistics (Prior to Junos OS 20.4R1 Releases)

show security advance-policy-based-routing sla active-probe-statistics

```
Active Probe Statistics:
  Src-IP          Dst-IP          PKT-LOSS(%)     RTT(us)       2way-Jit(us)    Ing-Jit(us)
Egr-Jit(us)
  42.1.1.2        42.1.1.1        100             0             0               0
0
  41.1.1.2        41.1.1.1        100             0             0               0
0
  40.1.1.2        40.1.1.1        100             0             0               0
0
```

Release Information

Command introduced in Junos OS Release 18.2R1.

Command modified to include outgoing IP addresses field in Junos OS Release 20.4R1.

Command modified to include IPv6 addresses in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

Application Quality of Experience

Advanced Policy-Based Routing

show security advance-policy-based-routing sla profile (Application Name)

IN THIS SECTION

- Syntax | 23919
- Description | 23919
- Required Privilege Level | 23919
- Output Fields | 23919
- Sample Output | 23920
- Release Information | 23921

Syntax

```
show security advance-policy-based-routing sla profile profile-name application application-name
```

Description

Displays the details of the best path selected for routing the application traffic. You can use this command to learn additional details of the selected path such as the destination group it belongs to, next-hop IP address, and the server IP in case of SaaS applications.

Required Privilege Level

view

Output Fields

Table 1123 on page 23920 lists the output fields for the `show` command. Output fields are listed in the approximate order in which they appear.

Table 1123: show security advance-policy-based-routing sla profile profile-name application application-name

Field Name	Field Description
`Best-Path Local IP Address`	The best link selected among all the links in the selected destination path group to send the application traffic.
`Destination-group name`	The destination path group name from which the link is selected.
`Next-Hop ID`	Next hop by ID number. It is the address of the next station to which the packet is sent on the way to its final destination. The range of values is 1 through 65,535.
`Server IP`	IP address of the server. Displayed as `N/A` for non-SaaS applications.

Sample Output

show security advance-policy-based-routing sla profile *profile-name* application *application-name* (Junos OS 21.3R1)

```
user@host> show security advance-policy-based-routing sla profile profile-1 application any
Best-Path Local IP Address  4000:5::1
Destination-group name      N/A
Next-Hop ID                 262145
Server IP                   10:1::2
```

show security advance-policy-based-routing sla profile *profile-name* application *application-name* (Junos OS 20.2R1)

```
user@host> show security advance-policy-based-routing sla profile profile-1 application JUNOS:ssh

Best-Path Local IP Address  172.16.4.2
Destination-group name      DPG-1
Next-Hop ID                 262142
Server IP                   N/A
```

```
Best-Path Local IP Address  182.17.5.3
Destination-group name      DPG-2
Next-Hop ID                 263132
Server IP                   N/A
```

Release Information

Command is modified in Junos OS Release 20.2R1.

Command is modified in Junos OS Release 21.3R1 to include IPv6 addresses.

RELATED DOCUMENTATION

Application Quality of Experience

show security advance-policy-based-routing sla profile (Next-Hop)

IN THIS SECTION

- Syntax | 23921
- Syntax | 23922
- Description | 23922
- Required Privilege Level | 23922
- Output Fields | 23922
- Sample Output | 23923
- Release Information | 23927

Syntax

```
show security advance-policy-based-routing sla profile profile-name application application-name
next-hop next-hop-number
```

Syntax

Syntax for Junos OS Releases prior 20.2R1

```
show security advance-policy-based-routing sla profile sla-profile-name application application-name destination-group-name destination-group-name
```

Description

Displays the details about the SLA applied for the application traffic. You can use this command to get all the details associated with an SLA such as SLA rule, APBR profile, APBR rules, routing-instance, probes, metrics and so on.

Required Privilege Level

view

Output Fields

Table 1124 on page 23922 lists the output fields for the `show` command. Output fields are listed in the approximate order in which they appear.

Table 1124: show security advance-policy-based-routing sla profile

Field Name	Field Description
`Application Name`	Name of the application.
`Application ID`	ID of the application
`DSCP`	DSCP value. This value corresponds to decimal values 0-63. This field is introduced in Junos OS Release 19.4R1.
`APBR Profile Name`	Name of the advanced policy-based (APBR) routing profile.
`APBR Rule Name`	Name of the APBR rule.

Table 1124: show security advance-policy-based-routing sla profile *(Continued)*

Field Name	Field Description
Application State	State of the application traffic.
Path Switch Idle State	Path switch idle state where no subsequent switching of application traffic path occurred.
Routing Instance Name	Name of the routing instance applied.
SLA Rule Name	Name of the SLA rule applied.
Active Probe Name	Name of the active probe parameter configured.
Selected Tunnel Destination	Selected tunnel destination where active probes are sent.
SLA Metrics	SLA metrics parameters, that are used by AppQoE to evaluate the SLA of the link. The SLA metric includes following parameters such as packet loss, RTT, jitter, and jitter type. Starting in Junos OS Release 19.2, With application-level summarization feature, each application's maximum, minimum, and average values of all the SLA metrics are displayed.

Sample Output

show security advance-policy-based-routing sla profile *profile-name* application *application-name* next-hop *number* (Junos OS Release 21.3R1)

```
user@host> show security advance-policy-based-routing sla profile apbr1 application junos:SSH
next-hop 262145

Application Details:

  Application Name            junos:SSH
  Application ID              198
  DSCP                        N/A
```

```
  APBR Profile Name           apbr1
  APBR Rule Name              rule1
  Application State           SLA MET
  Path Switch Idle State      0
  Routing Instance Name       TC1_VPN
  SLA Rule Name               sla1
  Active Probe Name           probe1
  Best-Path Local IP Address  4000:5::1
  Link Tag                    MPLS
  Link Priority               1
Average SLA Metrics:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  255             4294967295      4294967295      4294967295      4294967295
Min SLA Metrics:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  255             4294967295      4294967295      4294967295      4294967295
Max SLA Metrics:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  255             4294967295      4294967295      4294967295      4294967295
```

show security advance-policy-based-routing sla profile *profile-name* application *application-name* next-hop *number* server-ip *ip-address* (Junos OS 20.4R1)

```
user@host> show security advance-policy-based-routing sla profile profile-name application ANY
next-hop 263132 server-ip 10.1.1.1 (Junos OS 20.4R1)

Application Details:
  Application Name            ANY
  Application ID              0
  DSCP                        N/A
  APBR Profile Name           profile1
  APBR Rule Name              r1
  Application State           SLA MET
  Path Switch Idle State      0
  Routing Instance Name       saas_vrf
  SLA Rule Name               sla1
  Active Probe Name           probe1
  Selected Source             42.1.1.2
Average SLA Metrics:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  255             4294967295      4294967295      4294967295      4294967295
```

```
Min SLA Metrics:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  255             4294967295      4294967295      4294967295      4294967295
Max SLA Metrics:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  255             4294967295      4294967295      4294967295      4294967295
```

show security advance-policy-based-routing sla profile *profile-name* application *application-name* next-hop *number* (Junos OS Release 20.2R1)

```
user@host> show security advance-policy-based-routing sla profile apbr1 application junos:ssh
next-hop 262142

Application Details:
  Application Name            junos:SSH
  Application ID              198
  DSCP                        N/A
  APBR Profile Name           apbr1
  APBR Rule Name              rule1
  Application State           SLA MET
  Path Switch Idle State      0
  Routing Instance Name       TC1_VPN
  SLA Rule Name               sla1
  Active Probe Name           probe1
  Best-Path Local IP Address  40.1.1.2
Average SLA Metrics:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  0               1979            3038            3091            52
Min SLA Metrics:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  0               1167            23              17              16
Max SLA Metrics:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  0               36921           32053           26460           5593
```

show security advance-policy-based-routing sla profile *profile-name* destination-group-name *group-name* dscp *number* (Junos OS Release 19.4 and Later)

```
user@host> show security advance-policy-based-routing sla profile p2 destination-group-name
site3 dscp 15
Application Details:
  Application Name            N/A
  Application ID              N/A
  DSCP                        15
  APBR Profile Name           p2
  APBR Rule Name              def
  Application State           SLA MET
  Path Switch Idle State      0
  Routing Instance Name       TC1_VPN
  SLA Rule Name               sla1
  Active Probe Name           probe1
  Selected Tunnel Destination 111.111.114.1
Average SLA Metrics:
PKT-LOSS(%)     RTT(us)       2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  0               4645        1190            1318            722
Min SLA Metrics:
PKT-LOSS(%)     RTT(us)       2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  0               3589        104             4               195
Max SLA Metrics:
PKT-LOSS(%)     RTT(us)       2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  0               7329        3000            4452            1884
```

show security advanced-policy-based-routing sla profile *profile-name* application *application-name* destination-group-name *group-name* (Junos OS Release 19.2 and Later)

```
user@host> show security advanced-policy-based-routing sla profile apbrProf1 application
junos:CNN destination-group-name p1)
Application Details:
  Application Name            junos:CNN
  Application ID              988
  APBR Profile Name           apbrProf1
  APBR Rule Name              rule1
  Application State           SLA MET
  Path Switch Idle State      0
  Routing Instance Name       ri3
```

```
  SLA Rule Name               SLA1
  Active Probe Name           PP1
  Selected Tunnel Destination 5.1.1.1
SLA Metrics:
Average:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  0             1118            34              70              36
Minimum:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  0             1000            34              70              36
Maximum:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  0             1236            34              70              36
```

show security advanced-policy-based-routing sla profile *profile-name* application *application-name* destination-group-name *group-name* (Junos OS Release Prior 19.2R1)

```
user@host> show security advance-policy-based-routing sla profile apbr-1 application junos:HTTP
destination-group-name d1 (Junos OS Release Prior 19.2R1)
Application Details:
  Application Name            junos:HTTP
  Application ID              67
  APBR Profile Name           apbr1
  APBR Rule Name              rule1
  Application State           NO PATH SELECTED
  Path Switch Idle State      0
  Routing Instance Name       appqoe-vrf
  SLA Rule Name               sla1
  Active Probe Name           probe1
  Selected Tunnel Destination 0.0.0.0
SLA Metrics:
PKT-LOSS(%)     RTT(us)         2way-Jit(us)    Ing-Jit(us)     Egr-Jit(us)
  0               0               0               0               0
```

Release Information

Command introduced in Junos OS Release 18.2R1.

Command is modified in Junos OS Release 19.2R1 to display average, minimum, and maximum packet loss.

Command is modified in Junos OS Release 19.4R1 to include DSCP value in the output.

Command is modified in Junos OS Release 20.2R1 to include the next-hop *next-hop-number* in the syntax.

Command is modified in Junos OS Release 20.4R1 to include server IP addresses.

Command is modified in Junos OS Release 21.3R1 to display IPv6 addresses in the output.

RELATED DOCUMENTATION

Application Quality of Experience

Advanced Policy-Based Routing

show security advance-policy-based-routing sla profile (Status)

IN THIS SECTION

- Syntax | 23928
- Syntax | 23929
- Description | 23929
- Required Privilege Level | 23929
- Output Fields | 23929
- Sample Output | 23930
- Release Information | 23931

Syntax

```
show security advance-policy-based-routing sla profile <sla-profile-name> application
<application-name status
dscp dscp-value.
next-hop next-hop-id.
```

Syntax

Syntax prior to Junos OS Release 20.2R1

```
show security advance-policy-based-routing sla profile sla-profile-name application application-name destination-group-name destination-group-name
```

Description

Displays the number of times SLA violations occurred, application traffic switched route path, and monitored sessions.

Required Privilege Level

view

Output Fields

Table 1125 on page 23929 lists the output fields for the `show` command. Output fields are listed in the approximate order in which they appear.

Table 1125: show security advance-policy-based-routing sla profile

Field Name	Field Description
`Num of SLA Violations`	Number of times SLA violations occurred.
`Num of Path Switches`	Number of times application traffic switched route path. the value for the counter displayed as incremented on both the nodes although the path switch happened only once. When an SRX Series Firewall is operating in chassis cluster mode, the `Num of Path Switches` value depends on: If the path switch is due to AppQoE trigger: number of path switches counter is incremented on the specific node. If the path switch is due to other events such as a route event: number of path switches counter can be incremented on the specific node or incremented on both the nodes.

Table 1125: show security advance-policy-based-routing sla profile *(Continued)*

Field Name	Field Description
Num of monitored sessions	Number of monitored sessions by passive probes.
Num of sessions	Number of sessions.
Num of Violated Probes	Number of violations collected through probes.

Sample Output

command-name

```
user@host> show security advance-policy-based-routing sla profile apbr1 application application
any next-hop 263132 server-ip 10.1.1.1 status (Junos OS Release 20.4R1)
Application status:
  Num of SLA Violations       0
  Num of Path Switches        0
  Num of monitored sessions   0
  Num of sessions             0
  Num of Violated Probes      0
```

command-name

```
user@host> show security advance-policy-based-routing sla profile apbr1 application junos:ssh
next-hop 262142 status (Junos OS Release 20.2R1)
Application status:
Num of SLA Violations       0
Num of Path Switches        0
Num of monitored sessions   1
Num of sessions             1
Num of Violated Probes      0
```

command-name

```
user@host> show security advanced-policy-based-routing sla profile apbr1 application junos:ssh
destination-group-name p1 status (Prior to Junos OS Release 20.2R1)
Application status:
  Num of SLA Violations       2
  Num of Path Switches        0
  Num of monitored sessions   0
  Num of sessions             0
  Num of Violated Probes      6
```

Release Information

Command introduced in Junos OS Release 18.2R1. The syntax is changed to include the `next-hop` option in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Application Quality of Experience

Advanced Policy-Based Routing

show security advance-policy-based-routing sla statistics

IN THIS SECTION

- Syntax | 23932
- Description | 23932
- Required Privilege Level | 23932
- Output Fields | 23932
- Sample Output | 23933
- Release Information | 23933

Syntax

```
show security advance-policy-based-routing sla statistics
```

Description

Display the SLA statistics.

Required Privilege Level

view

Output Fields

Table 1126 on page 23932 lists the output fields for the `show security advance-policy-based-routing sla statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1126: show security advance-policy-based-routing sla statistics

Field Name	Field Description
`Passive Probe Session Processed`	Number of sessions on which passive probes are sent.
`Possible Passive Probe Sessions`	Number of sessions considered for passive probes.
`Passive Probe Sessions Sampled`	Number of sessions, from which, data is subjected to sampling.
`Passive Probe Ongoing Sessions`	Number of sessions on which passive probes are active.
`SLA violations`	Number of SLA violations detected.
`Active Probe Paths`	Number of links identified for active probe.
`Active Probe Session`	Number of sessions on which active probes are sent.

Table 1126: show security advance-policy-based-routing sla statistics *(Continued)*

Field Name	Field Description
`Active Probes Sent`	Number of active probes sent.
`Active Probe Paths down`	Number of links on which active probes are sent, are not active.

Sample Output

show security show security advance-policy-based-routing sla statistics

```
user@host> show security advance-policy-based-routing sla statistics

Advance Profile Based Routing SLA statistics:
Passive Probe Statistics
Passive Probe Session Processed 7040
Possible Passive Probe Sessions 0
Passive Probe Sessions Sampled 0
Passive Probe Ongoing Sessions 0
SLA violations 0
Active Probe Statistics
Active Probe Paths 0
Active Probe Session 3
Active Probes Sent 18360
Active Probe Paths down 3
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Application Quality of Experience

Advanced Policy-Based Routing

show security advance-policy-based-routing sla status

IN THIS SECTION

- Syntax | 23934
- Description | 23934
- Required Privilege Level | 23934
- Sample Output | 23934
- Release Information | 23935

Syntax

```
show security advance-policy-based-routing sla status
```

Description

Display the status of enabling switching of application path to an alternate route.

When local route switching is enabled, switching of application traffic to other route is enabled and also SLA monitoring and reporting functionality is available. By enabling local switch routing, the best possible link is selected for the application traffic to meet performance requirements as specified in SLA (service-level agreement).

Required Privilege Level

view

Sample Output

show security advance-policy-based-routing sla status

```
user@host> show security advance-policy-based-routing sla status
Local Switching is enabled.
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

> *Example: Configuring Advanced Policy-Based Routing for Application-Aware Traffic Management Solution*

show security advance-policy-based-routing sla version

IN THIS SECTION

- Syntax | 23935
- Description | 23935
- Required Privilege Level | 23935
- Output Fields | 23936
- show security advance-policy-based-routing sla version | 23936
- Release Information | 23936

Syntax

```
show security advance-policy-based-routing sla version
```

Description

Displays AppQoE version details. This information helps verify that the SLA version on both hub device and spoke device is same.

Required Privilege Level

view

Output Fields

show security advance-policy-based-routing sla version

command-name

```
user@host> show security advance-policy-based-routing sla version
SLA version: APPQOE.VERS.1.0.0.0
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Application Quality of Experience

Advanced Policy-Based Routing

show security advance-policy-based-routing statistics

IN THIS SECTION

- Syntax | 23937
- Description | 23937
- Required Privilege Level | 23937
- Output Fields | 23937
- Sample Output | 23940
- Release Information | 23943

Syntax

```
show security advance-policy-based-routing statistics
```

Description

Displays the statistics counter for the APBR.

You can use this command to understand the details on traffic handling with APBR such as:

- Sessions processed for the application-based routing.
- The number of times the application traffic matches the APBR profile and APBR is applied for the session.
- The number of times AppID is consulted to identify application traffic.

Required Privilege Level

view

Output Fields

Table 1127 on page 23937 lists the output fields for the `show security advance-policy-based-routing statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1127: show security advance-policy-based-routing statistics

Field Name	Field Description
`Session Processed`	The number of sessions processed for the application-based routing.
`ASC Success`	The number of times the presence of an entry in the application system cache (ASC) is found.
`Rule match success`	The number of times the application traffic matches the APBR profile.
`Route modified`	The number of times the APBR is applied for the session.

Table 1127: show security advance-policy-based-routing statistics *(Continued)*

Field Name	Field Description
AppID Requested	The number of times AppID is consulted to identify application traffic.

Table 1128 on page 23938 lists the output fields for the `show security advance-policy-based-routing statistics` command for midstream support. Output fields are listed in the approximate order in which they appear.

Table 1128: show security advance-policy-based-routing statistics (Advanced Policy-Based Routing Midstream Support)

Field Name	Field Description
Session Processed	The number of sessions processed for the application-based routing.
AppID cache hits	The number of times the presence of an entry in the application system cache (ASC) is found.
AppID Requested	The number of times AppID was consulted to identify application traffic.
Rule matches	The number of times the application traffic matches the APBR profile.
Route changed on cache hits	The number of times the APBR is applied for the session.
Route changed midstream	Number of times a route is changed for a session.
Zone mismatch	No of times a zone for an interface is changed in the middle of a session.
Drop on zone mismatch	Number of times a session is terminated because of change of zone in the middle of the session.

Table 1129 on page 23939 lists the output fields for the `show security advance-policy-based-routing statistics` command starting in Junos OS Release 19.3R1 and later releases. Output fields are listed in the approximate order in which they appear.

Table 1129: show security advance-policy-based-routing statistics

Field Name	Field Description
`Session Processed`	The number of sessions processed for the application-based routing.
`App rule hit on cache hit`	The number of times a rule with a matching entry in the application system cache (ASC) is found.
`App rule hit on HTTP Proxy/ALG`	The number of times a rule matching with the application obtained from ALG or secure Web (HTTP) proxy is found.
`Midstream disabled rule hit on cache hit`	The number of times a rule with a disabled midstream has matching entry in the ASC.
`URL cat rule hit on cache hit`	The number of times a rule with defined URL categories in ASC is found.
`DSCP rule hit on first packet`	The number of times the rule with defined DSCP value is matched for the first session.
`App and DSCP hit on first packet`	The number of times the rule with defined DSCP value and application is matched for the first session.
`App rule hit midstream`	The number of times a route is changed in the middle of a session because of the rule with defined application is matched.
`URL cat rule hit midstream`	The number of times a route is changed in the middle of a session because of the rule with defined URL categories is matched.
`App and DSCP rule hit midstream`	The number of times the rule with DSCP value and application is matched for the midstream session.
`Midstream disabled rule hit midstream`	The number of times a route remains unchanged in the middle of a session after rule with defined application is matched.

Table 1129: show security advance-policy-based-routing statistics *(Continued)*

Field Name	Field Description
DSCP rule hit midstream	The number of times the rule with DSCP value is matched for the midstream session.
Route changed on cache hits	Number of times a route is changed for a session because of the APBR applied for the session.
Route changed on HTTP Proxy/ALG	Number of times a route is changed because of the rule match for secure Web (HTTP) proxy or ALG applied for the session.
Route changed midstream	Number of times a route is changed in the middle of a session because of the APBR applied for the session.
Zone mismatch	No of times a zone for an interface is changed in the middle of a session.
Drop on zone mismatch	Number of times a session is terminated because of change of zone in the middle of the session.
Next hop not found	Number of times a session is terminated because next-hop IP address was not reachable.
Application Services Bypass	The number of times the application services are bypassed for the session.

Sample Output

show security advance-policy-based-routing statistics

```
user@host>  show security advance-policy-based-routing statistics
Advance Profile Based Routing statistics:
  Session Processed:                    5529
  ASC Success:                          3113
  Rule match success:                   107
```

```
  Route  modified:                        107
  AppID Requested:                        2416
```

show security advance-policy-based-routing statistics (Midstream Support)

```
user@host>  show security advance-policy-based-routing statistics
Advance Profile Based Routing statistics:
  Sessions Processed                      0
  AppID cache hits                        0
  AppID requested                         0
  Rule matches                            0
  Route changed on cache hits             0
  Route changed midstream                 0
  Zone mismatch                           0
  Drop on zone mismatch                   0
```

show security advance-policy-based-routing statistics (Changed Options from Junos OS Release 18.4R1)

```
user@host>  show security advance-policy-based-routing statistics
Advance Profile Based Routing statistics:
  Sessions Processed                      2
  App rule hit on cache hit               1
  URL cat rule hit on cache hit           0
  App rule hit midstream                  1
  URL cat rule hit midstream              0
  Route changed on cache hits             1
  Route changed midstream                 1
  Zone mismatch                           0
  Drop on zone mismatch                   0
  Next hop not found                      0
```

show security advance-policy-based-routing statistics (Changed Options from Junos OS Release 19.1R1)

```
user@host>  show security advance-policy-based-routing statistics
Advance Profile Based Routing statistics:
  Sessions Processed                      110
```

```
  AppID cache hits                      110
  AppID requested                       0
  Rule matches                          2
  Route changed on cache hits           1
  Route changed midstream               1
  Zone mismatch                         0
  Drop on zone mismatch                 0
  Next hop not found                    0
  Application Services Bypass           1
```

show security advance-policy-based-routing statistics (Changed Options from Junos OS Release 19.3R1)

```
user@host>  show security advance-policy-based-routing statistics

Advance Profile Based Routing statistics:
  Sessions Processed                        0
  App rule hit on cache hit                 0
  App rule hit on HTTP Proxy/ALG            0
  URL cat rule hit on cache hit             0
  DSCP rule hit on first packet             0
  App and DSCP hit on first packet          0
  App rule hit midstream                    0
  URL cat rule hit midstream                0
  App and DSCP rule hit midstream           0
  DSCP rule hit midstream                   0
  Route changed on cache hits               0
  Route changed on HTTP Proxy/ALG           0
  Route changed midstream                   0
  Zone mismatch                             0
  Drop on zone mismatch                     0
  Next hop not found                        0
  Application services bypass               0
```

show security advance-policy-based-routing statistics (Changed Options from Junos OS Release 19.4R1)

```
user@host>  show security advance-policy-based-routing statistics
Advance Profile Based Routing statistics:
```

```
Sessions Processed                        9
App rule hit on cache hit                 0
App rule hit on HTTP Proxy/ALG            0
Midstream disabled rule hit on cache hit  2
URL cat rule hit on cache hit             0
DSCP rule hit on first packet             2
App and DSCP hit on first packet          0
App rule hit midstream                    1
Default rule match                        0
Midstream disabled rule hit midstream     1
URL cat rule hit midstream                0
App and DSCP rule hit midstream           0
DSCP rule hit midstream                   0
Route changed on cache hits               2
Route changed on HTTP Proxy/ALG           0
Route changed midstream                   0
Default rule applied                      0
Zone mismatch                             0
Drop on zone mismatch                     0
Next hop not found                        0
Application services bypass               0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D60. Support.

RELATED DOCUMENTATION

Example: Configuring Advanced Policy-Based Routing for Application-Aware Traffic Management Solution

show security advance-policy-based-routing status

IN THIS SECTION

Syntax | 23944

- Description | 23944
- Required Privilege Level | 23944
- Sample Output | 23944
- Release Information | 23944

Syntax

```
show security advance-policy-based-routing status
```

Description

Check if the advanced policy-based routing (APBR) is enabled.

You can create an advanced policy-based routing (APBR) profile (application profile) to match applications and application groups and redirect those matching traffic to the specified routing instance for the route lookup. The application profile is attached to a security zone or it can be attached to a specific logical or physical interface associated with the security zone.

Required Privilege Level

view

Sample Output

show security advance-policy-based-routing status

```
user@host> show security advance-policy-based-routing status
Advance Policy Based Routing is enabled.
```

Release Information

Command introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

> *Example: Configuring Advanced Policy-Based Routing for Application-Aware Traffic Management Solution*

show security advanced-policy-based-routing policy-name

IN THIS SECTION

- Syntax | 23945
- Description | 23945
- Options | 23945
- Required Privilege Level | 23946
- Output Fields | 23946
- Sample Output | 23950
- Release Information | 23951

Syntax

```
show security advanced-policy-based-routing policy-name policy-name
<count | detail | from-zone |logical-system |root-logical-system | start>
```

Description

Display a summary of all APBR policies configured on the device.

You can use this command to understand the details of an APBR policy such as:

- Name, status, zone-context of the APBR policy.
- The number of times the traffic matches the APBR policy and APBR profile applied for the session.

Options

count Display the number of configured APBR policies.

- **Range:** 1 to 65535

detail Display a detailed view of all of the APBR policies configured on the device.

from-zone Display specific zone details applicable to the APBR policy.

logical-system Displaythe logical system name.

root-logical-system Display information about the default root-logical-system.

start Display the policy from the given position.

- **Range:** 1 to 65535

Required Privilege Level

view

Output Fields

Table 1130 on page 23946 lists the output fields for the `show security advanced-policy-based-routing` *`policy-name`* command. Output fields are listed in the approximate order in which they appear.

Table 1130: show security advanced-policy-based-routing policy-name

Field Name	Field Description
Policy	Name of the APBR policy.
State	Status of the policy. The policy is in one of the following state: • enabled: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • disabled: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.

Table 1130: show security advanced-policy-based-routing policy-name *(Continued)*

Field Name	Field Description
Sequence Number	Number of the policy within a given context. For example, three policies that are applicable in a from-zone A-to-zone B context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zone C-to-zone D context, four policies might have sequence numbers 1, 2, 3, 4.
From zone	Name of the source zone.
Source addresses	The names of the source addresses for a policy. Address sets are resolved to their individual names.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book
Applications	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time.
APBR Profile	Name of the applicable ABPR profile.

Table 1131 on page 23947 lists the output fields for the `show security advanced-policy-based-routing detail` command. Output fields are listed in the approximate order in which they appear.

Table 1131: show security advanced-policy-based-routing detail

Field Name	Field Description
APBR Policy	Name of the APBR policy.
State	Status of the policy. The policy is in one of the following state: • enabled: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • disabled: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.

Table 1131: show security advanced-policy-based-routing detail *(Continued)*

Field Name	Field Description
Index	Internal number associated with the policy.
Sequence Number	Number of the policy within a given context. For example, three policies that are applicable in a from-zone A-to-zone B context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zone C-to-zone D context, four policies might have sequence numbers 1, 2, 3, 4.
From zone	Name of the source zone.
Source addresses	The names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.

Table 1131: show security advanced-policy-based-routing detail *(Continued)*

Field Name	Field Description
Applications	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • IP protocol: The Internet protocol used by the application—for example, TCP, UDP, ICMP. • ALG: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If application-protocol ignore is configured, ignore is displayed. Otherwise, 0 is displayed. However, even if this command shows ALG: 0, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when application-protocol ignore is not configured for custom applications. • Inactivity timeout: Elapsed time without activity after which the application is terminated. • Source port range: The low-high source port range for the session application. • Destination port range: The low-high destination port range for the session application.
APBR Profile	Name of the applicable ABPR profile.

Table 1132 on page 23949 lists the output fields for the `show security advanced-policy-based-routing from-zone` command. Output fields are listed in the approximate order in which they appear.

Table 1132: show security advanced-policy-based-routing from-zone

Field Name	Field Description
From zone	Name of the source zone.
Policy count	Number of APBR policies configured for the zone.

Table 1133 on page 23950 lists the output fields for the `show security advanced-policy-based-routing hit-count` command. Output fields are listed in the approximate order in which they appear.

Table 1133: show security advanced-policy-based-routing hit-count

Field Name	Field Description
Logical system	Name of the associated logical system.
Index	Internal number associated with the policy.
From zone	Name of the source zone.
Name	Name of the APBR policy.
Policy count	Number of hits for each security policy.
Number of policy	Number of security policies for which hit counts are displayed.

Sample Output

show security advanced-policy-based-routing detail

```
user@host> show security advanced-policy-based-routing detail
Policy: p1, State: enabled, Index: 4
  Sequence number: 1
  From zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  APBR Profile: apbr-pr1
```

show security advanced-policy-based-routing from-zone

```
user@host> show security advanced-policy-based-routing from-zone trust
From zone: trust
  Policy: p1, State: enabled, Index: 4, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Applications: any
  APBR Profile: apbr-pr1
                                  1
```

show security advanced-policy-based-routing hit-count

```
user@host> show security advanced-policy-based-routing hit-count
Logical system: root-logical-system
 Index   From zone        Name           Hit count
 1       trust            p1             0

Number of policy: 1
```

show security advanced-policy-based-routing policy-name

```
user@host>  show security advanced-policy-based-routing policy-name sla_policy1
From zone: trust
  APBR Policy: sla_policy1, State: enabled, Index: 7, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Applications: any
    APBR profile: apbr-pr-default
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Example: Configuring Advanced Policy-Based Routing for Application-Aware Traffic Management Solution

show security alarms

IN THIS SECTION

- Syntax | 23952
- Description | 23953
- Options | 23953
- Required Privilege Level | 23954
- Output Fields | 23954
- Sample Output | 23955
- Release Information | 23958

Syntax

```
show security alarms
<detail>
<alarm-id id-number>
<alarm-type [ types ]>
<newer-than YYYY-MM-DD.HH:MM:SS>
<older-than YYYY-MM-DD.HH:MM:SS>
<process process>
<severity severity>
```

Description

Display the alarms that are active on the device. Run this command when the CLI prompt indicates that a security alarm has been raised, as shown here:

```
[1 SECURITY ALARM] user@host#
```

Options

none Display all active alarms.

detail (Optional) Display detailed output.

alarm-id *id-number* (Optional) Display the specified alarm.

alarm-type [*types* **]** (Optional) Display the specified alarm type or a set of types.

You can specify one or more of the following alarm types:

- `authentication`
- `cryptographic-self-test`
- `decryption-failures`
- `encryption-failures`
- `ike-phase1-failures`
- `ike-phase2-failures`
- `key-generation-self-test`
- `non-cryptographic-self-test`
- `policy`
- `replay-attacks`

newer-than *YYYY-MM-DD.HH:MM:SS* (Optional) Display active alarms that were raised after the specified date and time.

older-than *YYYY-MM-DD.HH:MM:SS* (Optional) Display active alarms that were raised before the specified date and time.

process *process*

(Optional) Display active alarms that were raised by the specified system process.

severity *severity*

(Optional) Display active alarms of the specified severity.

You can specify the following severity levels:

- `alert`
- `crit`
- `debug`
- `emerg`
- `err`
- `info`
- `notice`
- `warning`

Required Privilege Level

security—To view this statement in the configuration.

Output Fields

Table 1134 on page 23954 lists the output fields for the `show security alarms` command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified or when the `detail` keyword is used.

Table 1134: show security alarms

Field Name	Field Description	Level of Output
`ID`	Identification number of the alarm.	All levels
`Alarm time`	Date and time the alarm was raised..	All levels

Table 1134: show security alarms *(Continued)*

Field Name	Field Description	Level of Output
Message	Information about the alarm, including the alarm type, username, IP address, and port number.	All levels
Process	System process (For example, login or sshd) and process identification number associated with the alarm.	detail
Severity	Severity level of the alarm.	detail

Sample Output

show security alarms

```
[3 SECURITY ALARMS] user@router> show security alarms

ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2’
2     2010-01-19 13:41:52 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2’
3     2010-01-19 13:42:13 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2’
```

show security alarms detail

```
[3 SECURITY ALARMS] user@router> show security alarms detail

Alarm ID   : 1
Alarm Type : authentication
Time       : 2010-01-19 13:41:36 PST
Message    : SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user'
reached from '203.0.113.2’
Process    : sshd (pid 1414)
Severity   : notice
```

```
Alarm ID   : 2
Alarm Type : authentication
Time       : 2010-01-19 13:41:52 PST
Message    : SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user'
reached from '203.0.113.2'
Process    : sshd (pid 1414)
Severity   : notice

Alarm ID   : 3
Alarm Type : authentication
Time       : 2010-01-19 13:42:13 PST
Message    : SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user'
reached from '203.0.113.2'
Process    : sshd (pid 1414)
Severity   : notice
```

show security alarms alarm-id

```
[3 SECURITY ALARMS] user@router> show security alarms alarm-id 1

ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2'
```

show security alarms alarm-type authentication

```
[3 SECURITY ALARMS] user@router> show security alarms alarm-type authentication

ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2'
2     2010-01-19 13:41:52 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2'
3     2010-01-19 13:42:13 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2'
```

show security alarms newer-than <time>

```
[3 SECURITY ALARMS] user@router> show security alarms newer-than 2010-01-19.13:41:59

3     2010-01-19 13:42:13 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2’
```

show security alarms older-than <time>

```
[3 SECURITY ALARMS] user@router> show security alarms older-than 2010-01-19.13:41:59

ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2’
2     2010-01-19 13:41:52 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2’
```

show security alarms process <process>

```
[3 SECURITY ALARMS] user@router> show security alarms process sshd

ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2’
2     2010-01-19 13:41:52 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2’
3     2010-01-19 13:42:13 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2’
```

show security alarms severity <severity>

```
[3 SECURITY ALARMS] user@router> show security alarms severity notice

ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2’
2     2010-01-19 13:41:52 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
```

```
for user 'user' reached from '203.0.113.2'
3     2010-01-19 13:42:13 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1)
for user 'user' reached from '203.0.113.2'
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

- clear security alarms
- Example: Generating a Security Alarm in Response to Policy Violations

show security alg h323 counters

IN THIS SECTION

- Syntax | 23958
- Description | 23959
- Options | 23959
- Required Privilege Level | 23959
- Output Fields | 23959
- Sample Output | 23960
- Release Information | 23962

Syntax

```
show security alg h323 counters
<logical-system | root-logical-system | tenant>
```

Description

This command displays H.323 Application Layer Gateway (ALG) counters information.

Options

- `logical-system`—Name of the logical system.
- `root-logical-system`—Displays root logical system as default.
- `tenant`—Name of the tenant system.

Required Privilege Level

view

Output Fields

Table 1135 on page 23959 lists the output fields for the `show security alg h323 counters` command. Output fields are listed in the approximate order in which they appear.

Table 1135: show security alg h323 counters Output Fields

Field Name	Field Description
`Packets received`	Number of H.323 ALG packets received.
`Packets dropped`	Number of H.323 ALG packets dropped.
`RAS message received`	Number of incoming RAS (Endpoint Registration, Admission, and Status) messages per second per gatekeeper received and processed.
`Q.931 message received`	Counter for Q.931 message received.
`H.245 message received`	Counter for H.245 message received.

Table 1135: show security alg h323 counters Output Fields *(Continued)*

Field Name	Field Description
Number of calls	Total number of H.323 ALG calls. **NOTE**: This counter displays the number of call legs and may not display the exact number of voice calls that are active. For instance, for a single active voice call between two endpoints, this counter might display a value of 2.
Number of active calls	Number of active H.323 ALG calls.
Number of DSCP marked	Number of DSCP marked H.323 ALG.
Decoding errors	Number of decoding errors.
Message flood dropped	Error counter for message flood dropped.
NAT errors	H.323 ALG Network Address Translation (NAT) errors.
Resource manager errors	H.323 ALG resource manager errors.
DSCP marked errors	H.323 ALG DSCP marked errors.

Sample Output

show security alg h323 counters

```
user@host> show security alg h323 counters
H.323 counters summary:
Packets received           :4060
Packets dropped            :24
RAS message received       :3690
Q.931 message received       :202
H.245 message received       :145
Number of calls        :25
```

```
Number of active calls       :0
H.323 Error Counters:
Decoding errors            :24
Message flood dropped       :0
NAT errors           :0
Resource manager errors       :0
H.323 Message Counters:
RRQ      : 431   RCF      :  49 ARQ  : 60  ACF      : 33
URQ      : 34   UCF      : 25 DRQ  : 55  DCF      : 44
oth RAS  : 2942  Setup    : 28 Alert    : 9 Connect    : 25
CallPrcd  : 18  Info    : 0  RelCmpl  : 39  Facility   : 14
Progress  : 0  Empty    : 65  OLC      : 20 OLC-ACK  :   20
oth H245  : 16
```

show security alg h323 counters tenant TN1

```
user@host> show security alg h323 counters tenant TN1
H.323 counters summary:
Packets received            :0
Packets dropped             :0
RAS message received        :0
Q.931 message received      :0
H.245 message received      :0
Number of calls             :0
Number of active calls      :0
Number of dscp marked       :0
H.323 Error Counters:
Decoding errors             :0
Message flood dropped       :0
NAT errors                  :0
Resource manager errors     :0
DSCP marked errors          :0
H.323 Message Counters:
RRQ      : 0   RCF      :  0 ARQ  : 0  ACF      : 0
URQ      : 0   UCF      : 0 DRQ  : 0  DCF      : 0
oth RAS  : 0  Setup    : 0 Alert    : 0 Connect    : 0
CallPrcd  : 0  Info    : 0  RelCmpl  : 0  Facility   : 0
  Empty    : 0  OLC      : 0  OLC-ACK  :   0
oth H245  : 0
```

show security alg h323 counters tenant all

```
user@host> show security alg h323 counters tenant all
H.323 counters summary:
Packets received            :0
Packets dropped             :0
RAS message received        :0
Q.931 message received      :0
H.245 message received      :0
Number of calls             :0
Number of active calls      :0
Number of dscp marked       :0
H.323 Error Counters:
Decoding errors             :0
Message flood dropped       :0
NAT errors                  :0
Resource manager errors     :0
DSCP marked errors          :0
H.323 Message Counters:
RRQ      : 0   RCF      :  0 ARQ  : 0  ACF      : 0
URQ      : 0   UCF      : 0 DRQ  : 0  DCF      : 0
oth RAS  : 0  Setup    : 0 Alert    : 0 Connect    : 0
CallPrcd  : 0  Info    : 0  RelCmpl  : 0  Facility   : 0
  Empty    : 0  OLC      : 0  OLC-ACK  :   0
oth H245  : 0
```

Release Information

Statement introduced in Junos OS Release 8.5.

The `logical-system` option is introduced in Junos OS Release 18.2R1.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

h323 | **7961**

clear security alg h323 counters | **3855**

show security alg ike-esp-nat summary

IN THIS SECTION

- Syntax | 23963
- Description | 23963
- Required Privilege Level | 23963
- Output Fields | 23963
- Sample Output | 23963
- Release Information | 23964

Syntax

```
show security alg ike-esp-nat summary
```

Description

This command displays Application Layer Gateway (ALG) for IKE and ESP information summary.

Required Privilege Level

view

Output Fields

Sample Output

show security alg ike-esp-nat summary

```
user@host> security alg ike-esp-nat summary
Initiator cookie: d5732d9b4114de1a
Responder cookie: 4776fe31164ef
Session-ID:       13
```

```
ALG state :       1
Timeout:          6292
Used IKE cookies: 0
Maximum IKE cookies: 2400
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

clear security alg ike-esp-nat | 3857

show security alg msrpc

IN THIS SECTION

- Syntax | 23964
- Description | 23965
- Options | 23965
- Required Privilege Level | 23965
- Output Fields | 23965
- Sample Output | 23966
- Release Information | 23966

Syntax

```
show security alg msrpc
<object-id-map> <node ( node-id | all | local | primary)>
```

Description

This command displays Microsoft (MS) remote procedure call (RPC) Application Layer Gateway (ALG) information.

Options

- none—Display all MS-RPC ID (UUID) to object ID (OID) table information.
- `object-id-map`—(Optional) Display information from the MS-RPC ID (UUID) to object ID (OID) table.
- `node`—(Optional) For chassis cluster configurations, display MS-RPC UUID-to-object-ID mapping information on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1136 on page 23965 lists the output fields for the `show security alg msrpc` command. Output fields are listed in the approximate order in which they appear.

Table 1136: show security alg msrpc Output Fields

Field Name	Field Description
`UUID`	MS-RPC ID.
`OID`	MS-RPC object ID.

Sample Output

show security alg msrpc object-id-map

```
user@host> show security alg msrpc object-id-map
UUID                                    OID
1be617c0-31a5-11cf-a7d8-00805f48a135    0x80000020
e3514235-4b06-11d1-ab04-00c04fc2dcd2    0x80000002
67df7c70-0f04-11ce-b13f-00aa003bac6c    0x80000014
```

Release Information

Command modified in Junos OS Release 10.1.

RELATED DOCUMENTATION

msrpc | 10999

show security alg sccp calls

IN THIS SECTION

- Syntax | 23967
- Description | 23967
- Options | 23967
- Required Privilege Level | 23967
- Output Fields | 23967
- Sample Output | 23968
- Sample Output | 23968
- Release Information | 23969

Syntax

```
show security alg sccp calls
<                brief                /                detail | node
(        node-id        | all | local | primary)>
```

Description

This command displays information about Skinny Client Control Protocol (SCCP) Application Layer Gateway (ALG) calls.

Options

- none | `brief`—Display brief call information.
- `detail`—(Optional) Display detailed call information.
- `node`—(Optional) For chassis cluster configurations, display SCCP calls on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1137 on page 23968 lists the output fields for the `show security alg sccp calls` command. Output fields are listed in the approximate order in which they appear.

Table 1137: show security alg sccp calls Output Fields

Field Name	Field Description
Client IP address	IP address of the client.
Client zone	Client zone ID.
Call manager	IP address of the call manager.
Resource manager group	Resource manager group ID.

Sample Output

show security alg sccp calls

```
user@host> show security alg sccp calls
Client IP          Zone      Call Manager       Conference ID  RM group
11.0.102.91        7         13.0.99.226       16789504       2047
12.0.102.96        8         13.0.99.226       16789505       2048
```

Sample Output

show security alg sccp calls detail

```
user@host> show security alg sccp calls detail
Client IP address: 11.0.102.91
Client zone: 7
Call Manager IP: 13.0.99.226
Conference ID: 16789504
Resource manager group: 2048
SCCP channel information:
    Media transmit channel address (IP address/Port): 0.0.0.0:0
    Media transmit channel translated address (IP address/Port): 0.0.0.0:0
    Media transmit channel pass-through party ID (PPID): 0
    Media transmit channel resource ID: 0
```

```
    Media receive channel address (IP address/Port): 11.0.102.91:20060
    Media receive channel translated address (IP address/Port): 25.0.0.1:1032
    Media receive channel pass-through party ID (PPID): 16934451
    Media receive channel resource ID: 8185
    Multimedia transmit channel address (IP address/Port): 0.0.0.0:0
    Multimedia transmit channel translated address (IP address/Port): 0.0.0.0:0
    Multimedia transmit channel pass-through party ID (PPID): 0
    Multimedia transmit channel resource ID: 0
    Multimedia receive channel address (IP address/Port): 0.0.0.0:0
    Multimedia receive channel translated address (IP address/Port): 0.0.0.0:0
    Multimedia receive channel pass-through party ID (PPID): 0
    Multimedia receive channel resource ID: 0
Total number of calls = 1
```

Release Information

Command introduced in Junos OS Release 8.5.

The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

sccp | 16341

clear security alg sccp calls | 3858

show security alg sccp counters

IN THIS SECTION

- Syntax | 23970
- Description | 23970
- Options | 23970
- Required Privilege Level | 23970
- Output Fields | 23970
- Sample Output | 23972

Release Information | 23972

Syntax

```
show security alg sccp counters
<node (          node-id         | all | local | primary)>
```

Description

This command displays information about Skinny Client Control Protocol (SCCP) Application Layer Gateway (ALG) counters.

Options

- none—Display all SCCP ALG counters.
- `node`—(Optional) For chassis cluster configurations, display SCCP counters on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1138 on page 23971 lists the output fields for the `show security alg sccp counters` command. Output fields are listed in the approximate order in which they appear.

Table 1138: show security alg sccp counters Output Fields

Field Name	Field Description
`Active client sessions`	Number of active SCCP ALG client sessions.
`Active calls`	Number of active SCCP ALG calls.
`Total calls`	Total number of SCCP ALG calls.
`Packets received`	Number of SCCP ALG packets received.
`PDUs processed`	Number of SCCP ALG protocol data units (PDU) processed.
`Current call rate`	Number of calls per second.
`Packets dropped`	Number of packets dropped by the SCCP ALG.
`Decode errors`	Number of decoding errors.
`Protocol errors`	Number of protocol errors.
`Address translation errors`	Number of NAT errors.
`Policy lookup errors`	Number of errors occurring during policy lookups.
`Unknown PDUs`	Number of unknown protocol data units (PDUs).
`Maximum calls exceed`	Number of times the maximum number of calls was exceeded.
`Maximum call rate exceed`	Number of times the maximum call rate was exceeded.
`Initialization errors`	Number of call initialization errors.

Table 1138: show security alg sccp counters Output Fields *(Continued)*

Field Name	Field Description
`Internal errors`	Number of internal errors.
`Nonspecific error`	Number of nonspecific errors.

Sample Output

show security alg sccp counters

```
user@host> show security alg sccp counters
SCCP call statistics:
  Active client sessions         : 4
  Active calls                   : 2
  Total calls                    : 3
  Packets received               : 232
  PDUs processed                 : 232
  Current call rate              : 0
Error counters:
  Packets dropped                : 0
  Decode errors                  : 0
  Protocol errors                : 0
  Address translation errors     : 0
  Policy lookup errors           : 0
  Unknown PDUs                   : 0
  Maximum calls exceeded         : 0
  Maximum call rate exceeded     : 0
  Initialization errors          : 0
  Internal errors                : 0
  Nonspecific error              : 0
```

Release Information

Command introduced in Junos OS Release 8.5.

The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

sccp | 16341

clear security alg sccp counters | 3860

show security alg sip calls

IN THIS SECTION

- Syntax | 23973
- Description | 23973
- Options | 23973
- Required Privilege Level | 23974
- Output Fields | 23974
- Sample Output | 23974
- Sample Output | 23975
- Release Information | 23975

Syntax

```
show security alg sip calls
<brief | detail / node ( node-id | all | local
| primary) >
```

Description

This command displays information about Session Initiation Protocols (SIP) Application Layer Gateway (ALG) calls.

Options

- none | `brief`—Display brief call information.
- `detail`—(Optional) Display detailed information about SIP ALG calls.

- `node`—(Optional) For chassis cluster configurations, display SIP calls on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1139 on page 23974 lists the output fields for the `show security alg sip calls` command. Output fields are listed in the approximate order in which they appear.

Table 1139: show security alg sip calls Output Fields

Field Name	Field Description
`UAS callid`	Call Identifier for the SIP ALG user agent server.
`State`	State of the SIP ALG user agent server.

Sample Output

show security alg sip calls

```
user@host> show security alg sip calls
Total number of calls: 1
           Call ID  : 1-2468@20.0.0.174
            Method   : INVITE
```

Sample Output

show security alg sip calls detail

```
user@host> show security alg sip calls detail
Total number of calls: 1
       Call ID  : 1-2468@20.0.0.174
              Method   : INVITE
              State    : SETUP
              Group ID : 307
```

Release Information

Command modified in Junos OS Release 9.2.

The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

sip (Security) | 27977

clear security alg sip calls | 3862

show security alg sip counters

IN THIS SECTION

- Syntax | 23976
- Description | 23976
- Options | 23976
- Required Privilege Level | 23977
- Output Fields | 23977
- Sample Output | 23981
- Release Information | 23996

Syntax

```
show security alg sip counters
<all-logical-systems-tenants>
<logical-system (logical-system-name | all)>
<root-logical-system>
<tenant (tenant-name | all)
<node (node-id | all | local | primary)>
```

Description

This command displays information about Session Initiation Protocol (SIP) Application Layer Gateway (ALG) counters.

Options

- none—Displays the information about SIP ALG counters of root logical system.
- `all-logical-systems-tenants`—(Optional) Displays the sum of SIP ALG counters of all logical systems and tenant systems.
- `logical-system` *`logical-system-name`*—(Optional) Displays information about SIP ALG counters of a specified logical system.
- `logical-system all`—(Optional) Displays the sum of SIP counters of all logical systems.
- `root-logical-system`—(Optional) Displays information about SIP ALG counters of root logical system as default.
- `tenant` *`tenant-name`*—(Optional) Displays information about SIP ALG counters of a specified tenant system.
- `tenant all`—(Optional) Displays the sum of SIP ALG counters of all tenant systems.
- `node`—(Optional) For chassis cluster configurations, displays the information about SIP ALG counters on a specific node (device) in the cluster.
  - *`node-id`*—Identification number of the node. It can be 0 or 1.
  - `all`—Displays information about all nodes.
  - `local`—Displays information about the local node.
  - `primary`—Displays information about the primary node.

Required Privilege Level

view

Output Fields

Table 1140 on page 23977 lists the output fields for the `show security alg sip counters` command. Output fields are listed in the approximate order in which they appear.

Table 1140: show security alg sip counters Output Fields

Field Name	Field Description
`INVITE`	Number of INVITE requests sent. An INVITE request is sent to invite another user to participate in a session.
`CANCEL`	Number of CANCEL requests sent. A user can send a CANCEL request to cancel a pending INVITE request. A CANCEL request has no effect if the SIP server processing the INVITE had sent a final response for the INVITE before it received the CANCEL
`ACK`	Number of ACK requests sent. The user from whom the INVITE originated sends an ACK request to confirm reception of the final response to the INVITE request.
`BYE`	Number of BYE requests sent. A user sends a BYE request to abandon a session. A BYE request from either user automatically terminates the session.
`RR header exceeded max`	Number of times the SIP ALG RR (Record-Route) headers exceeded the maximum limit.
`REGISTER`	Number of REGISTER requests sent. A user sends a REGISTER request to a SIP registrar server to inform it of the current location of the user. A SIP registrar server records all the information it receives in REGISTER requests and makes this information available to any SIP server attempting to locate a user.

Table 1140: show security alg sip counters Output Fields *(Continued)*

Field Name	Field Description
OPTIONS	Number of OPTIONS requests sent. An OPTION request is used by the User Agent (UA) to obtain information about the capabilities of the SIP proxy. A server responds with information about what methods, session description protocols, and message encoding it supports.
INFO	Number of INFO requests sent. An INFO message is used to communicate mid-session signaling information along the signaling path for the call.
MESSAGE	Number of MESSAGE requests sent. SIP messages consist of requests from a client to a server and responses to the requests from a server to a client with the purpose of establishing a session (or a call).
NOTIFY	Number of NOTIFY requests sent. NOTIFY requests are sent to inform subscribers of changes in state to which the subscriber has a subscription.
PRACK	Number of PRACK requests sent. The PRACK request plays the same role as ACK, but for provisional responses.
PUBLISH	Number of PUBLISH requests sent. The PUBLISH request used for publishing event state. PUBLISH is similar to REGISTER in that it allows a user to create, modify, and remove state in another entity which manages this state on behalf of the user.
REFER	Number of REFER requests sent. A REFER request is used to refer the recipient (identified by the Request-URI) to a third party by the contact information provided in the request.
SUBSCRIBE	Number of SUBSCRIBE requests sent. A SUBCRIBE request is used to request current state and state updates from a remote node.
UPDATE	Number of UPDATE requests sent. AN UPDATE request is used to create a temporary opening in the firewall (pinhole) for new or updates Session Description Protocol (SDP) information. The following fields are modified: Via, From, To, Call-ID, Contact, Route, and Record-Route.

Table 1140: show security alg sip counters Output Fields *(Continued)*

Field Name	Field Description
BENOTIFY	Number of BENOTIFY requests sent. A BENOTIFY request is used to reduce the unnecessary SIP signaling traffic on application servers. Applications that do not need a response for a NOTIFY request can enhance performance by enabling BENOTIFY.
SERVICE	Number of SERVICE requests sent. The SERVICE method used by a SIP client to request a service of a SIP server. It is a standard SIP message and will be forwarded until it reaches the server or end user which is performing the service.
OTHER	Number of OTHER requests sent.
Total Pkt-in	Number of SIP ALG total packets received.
Total Pkt dropped on error	Number of SIP ALG total packets dropped while transmission and retransmission of messages.
Call error	Number of SIP ALG call errors.
IP resolve error	Number of SIP ALG IP address resolution errors.
NAT error	Number of SIP ALG NAT errors.
Resource manager error	Number of SIP ALG resource manager errors.
RR header exceeded max	Number of times the SIP ALG RR (Record-Route) headers exceeded the maximum limit.
Contact header exceeded max	Number of times the SIP ALG contact headers exceeded the maximum limit.
Call Dropped due to limit	SIP ALG calls dropped because of call limits.

Table 1140: show security alg sip counters Output Fields *(Continued)*

Field Name	Field Description
SIP stack error	Number of SIP ALG stack errors.
SIP decode error	Number of SIP ALG decode errors.
SIP unknown method error	Number of SIP ALG unknown method errors.
SIP dscp marked	SIP ALG DSCP marked.
Invite Dropped due to call limit	Number of SIP ALG invite dropped due to call limits.
SIP msg not processed by stack	Number of SIP ALG stack errors.
SIP msg not processed by alg	Number of SIP ALG messages not processed by ALGs.
SIP unknown method dropped	Number of SIP ALG unknown method errors.
Decoding error	Number of SIP ALG decoding errors.
Request for disconnected call	Number of SIP ALG calls disconnected.
Request out of state	Number of SIP ALG messages out of state errors.
RTO message sent	Number of SIP ALG marked RTO messages sent.
RTO message received	Number of SIP ALG RTO messages received.
RTO buffer allocation failure	Number of SIP ALG RTO buffer allocation failures.
RTO buffer transmit failure	Number of SIP ALG RTO buffer transmit failures.

Table 1140: show security alg sip counters Output Fields *(Continued)*

Field Name	Field Description
RTO send processing error	Number of SIP ALG RTO send processing errors.
RTO receive processing error	Number of SIP ALG RTO receiving processing errors.
RTO receive invalid length	Number of SIP ALG RTOs receiving invalid length.
RTO receive call process error	Number of SIP ALG RTO receiving call process errors.
RTO receive call allocation error	Number of SIP ALG RTO receiving call allocation error.
RTO receive call register error	Number of SIP ALG RTO receiving call register errors.
RTO receive invalid status error	Number of SIP ALG RTO receiving register errors.

Sample Output

show security alg sip counters

```
user@host> show security alg sip counters
SIP message counters(T:Transmit, RT:Retransmit):
   Method     T   1xx   2xx   3xx  4xx   5xx   6xx
             RT   RT    RT    RT   RT    RT    RT
   INVITE     4    4     3     0    0     0     0
              0    0     0     0    0     0     0
   CANCEL     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
      ACK     3    0     0     0    0     0     0
              0    0     0     0    0     0     0
      BYE     3    0     3     0    0     0     0
              0    0     0     0    0     0     0
 REGISTER     7    0     7     0    0     0     0
```

```
                0    0     0     0    0     0     0
  OPTIONS       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
     INFO       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
  MESSAGE       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
   NOTIFY       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
    PRACK       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
  PUBLISH       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
    REFER       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
SUBSCRIBE       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
   UPDATE       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
 BENOTIFY       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
  SERVICE       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
    OTHER       0    0     0     0    0     0     0
                0    0     0     0    0     0     0
SIP Error Counters:
Total Pkt-in                        :34
Total Pkt dropped on error          :0
Call error                          :0
IP resolve error                    :0
NAT error                           :0
Resource manager error              :0
RR header exceeded max              :0
Contact header exceeded max         :0
Call Dropped due to limit                  : 0
SIP stack error                                       : 0
SIP decode error                                 : 0
SIP unknown method error               : 0
RTO message sent                           : 0
RTO message received                    : 0
RTO buffer allocation failure             : 0
RTO buffer transmit failure               : 0
RTO send processing error             : 0
```

```
RTO receive processing error          : 0
RTO receive invalid length                : 0
RTO receive call process error        : 0
RTO receive call allocation error     : 0
RTO receive call register error         : 0
RTO receive invalid status error      : 0
```

show security alg sip counters all-logical-systems-tenants

```
user@host> show security alg sip counters all-logical-systems-tenants
Method              T            1xx            2xx            3xx            4xx
5xx            6xx
                      RT             RT             RT             RT             RT
RT             RT

   INVITE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
   CANCEL              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
      ACK              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
      BYE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
 REGISTER              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
  OPTIONS              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
     INFO              0              0              0              0
0              0              0
```

```
                        0             0             0             0
0             0             0
  MESSAGE               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
   NOTIFY               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
    PRACK               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
  PUBLISH               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
    REFER               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
SUBSCRIBE               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
   UPDATE               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
 BENOTIFY               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
  SERVICE               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
    OTHER               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
```

```
SIP Error Counters:
  Total Pkt-in                      : 0
  Total Pkt dropped on error        : 0
  Call error                        : 0
  IP resolve error                  : 0
  NAT error                         : 0
  Resource manager error            : 0
  RR header exceeded max            : 0
  Contact header exceeded max       : 0
  Call Dropped due to limit         : 0
  SIP stack error                   : 0
  SIP decode error                  : 0
  SIP unknown method error          : 0
  SIP dscp marked                   : 0
  SIP dscp marked error             : 0
  RTO message sent                  : 0
  RTO message received              : 0
  RTO buffer allocation failure     : 0
  RTO buffer transmit failure       : 0
  RTO send processing error         : 0
  RTO receive processing error      : 0
  RTO receive invalid length        : 0
  RTO receive call process error    : 0
  RTO receive call allocation error : 0
  RTO receive call register error   : 0
  RTO receive invalid status error  : 0
```

show security alg sip counters logical-system lsys1

```
user@host> show security alg sip counters lsys1
   Method              T            1xx            2xx            3xx            4xx
5xx            6xx
                      RT             RT             RT             RT             RT
RT             RT

   INVITE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
   CANCEL              0              0              0              0
0              0              0
```

```
                        0             0             0             0
0             0             0
      ACK               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
      BYE               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
 REGISTER               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
  OPTIONS               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
     INFO               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
  MESSAGE               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
   NOTIFY               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
    PRACK               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
  PUBLISH               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
    REFER               0             0             0             0
0             0             0
                        0             0             0             0
0             0             0
SUBSCRIBE               0             0             0             0
```

```
0              0              0
                       0              0              0              0
0              0              0
   UPDATE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
 BENOTIFY              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
  SERVICE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
    OTHER              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0

SIP Error Counters:
  Total Pkt-in                      : 0
  Total Pkt dropped on error        : 0
  Call error                        : 0
  IP resolve error                  : 0
  NAT error                         : 0
  Resource manager error            : 0
  RR header exceeded max            : 0
  Contact header exceeded max       : 0
  Call Dropped due to limit         : 0
  SIP stack error                   : 0
  SIP decode error                  : 0
  SIP unknown method error          : 0
  SIP dscp marked                   : 0
  SIP dscp marked error             : 0
  RTO message sent                  : 0
  RTO message received              : 0
  RTO buffer allocation failure     : 0
  RTO buffer transmit failure       : 0
  RTO send processing error         : 0
  RTO receive processing error      : 0
  RTO receive invalid length        : 0
  RTO receive call process error    : 0
```

```
  RTO receive call allocation error : 0
  RTO receive call register error   : 0
  RTO receive invalid status error  : 0
```

show security alg sip counters logical-system all

```
user@host> show security alg sip counters logical-system all
   Method              T            1xx            2xx            3xx            4xx
5xx            6xx
                      RT             RT             RT             RT             RT
RT             RT

   INVITE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
   CANCEL              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
      ACK              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
      BYE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
 REGISTER              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
  OPTIONS              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
     INFO              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
  MESSAGE              0              0              0              0
```

```
0              0              0
                       0              0              0              0
0              0              0
   NOTIFY              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
    PRACK              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
  PUBLISH              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
    REFER              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
SUBSCRIBE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
   UPDATE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
 BENOTIFY              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
  SERVICE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
    OTHER              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0

SIP Error Counters:
  Total Pkt-in                      : 0
  Total Pkt dropped on error        : 0
```

```
  Call error                        : 0
  IP resolve error                  : 0
  NAT error                         : 0
  Resource manager error            : 0
  RR header exceeded max            : 0
  Contact header exceeded max       : 0
  Call Dropped due to limit         : 0
  SIP stack error                   : 0
  SIP decode error                  : 0
  SIP unknown method error          : 0
  SIP dscp marked                   : 0
  SIP dscp marked error             : 0
  RTO message sent                  : 0
  RTO message received              : 0
  RTO buffer allocation failure     : 0
  RTO buffer transmit failure       : 0
  RTO send processing error         : 0
  RTO receive processing error      : 0
  RTO receive invalid length        : 0
  RTO receive call process error    : 0
  RTO receive call allocation error : 0
  RTO receive call register error   : 0
  RTO receive invalid status error  : 0
```

show security alg sip counters root-logical-system

```
user@host> show security alg sip counters root-logical-system
  Method              T            1xx            2xx            3xx            4xx
5xx            6xx
                      RT             RT             RT             RT             RT
RT             RT

   INVITE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
   CANCEL              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
      ACK              0              0              0              0
```

```
0             0             0
                      0             0             0             0
0             0             0
      BYE             0             0             0             0
0             0             0
                      0             0             0             0
0             0             0
 REGISTER             0             0             0             0
0             0             0
                      0             0             0             0
0             0             0
  OPTIONS             0             0             0             0
0             0             0
                      0             0             0             0
0             0             0
     INFO             0             0             0             0
0             0             0
                      0             0             0             0
0             0             0
  MESSAGE             0             0             0             0
0             0             0
                      0             0             0             0
0             0             0
   NOTIFY             0             0             0             0
0             0             0
                      0             0             0             0
0             0             0
    PRACK             0             0             0             0
0             0             0
                      0             0             0             0
0             0             0
  PUBLISH             0             0             0             0
0             0             0
                      0             0             0             0
0             0             0
    REFER             0             0             0             0
0             0             0
                      0             0             0             0
0             0             0
SUBSCRIBE             0             0             0             0
0             0             0
                      0             0             0             0
0             0             0
```

```
   UPDATE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
 BENOTIFY              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
  SERVICE              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0
    OTHER              0              0              0              0
0              0              0
                       0              0              0              0
0              0              0

SIP Error Counters:
  Total Pkt-in                      : 0
  Total Pkt dropped on error        : 0
  Call error                        : 0
  IP resolve error                  : 0
  NAT error                         : 0
  Resource manager error            : 0
  RR header exceeded max            : 0
  Contact header exceeded max       : 0
  Call Dropped due to limit         : 0
  SIP stack error                   : 0
  SIP decode error                  : 0
  SIP unknown method error          : 0
  SIP dscp marked                   : 0
  SIP dscp marked error             : 0
  RTO message sent                  : 0
  RTO message received              : 0
  RTO buffer allocation failure     : 0
  RTO buffer transmit failure       : 0
  RTO send processing error         : 0
  RTO receive processing error      : 0
  RTO receive invalid length        : 0
  RTO receive call process error    : 0
  RTO receive call allocation error : 0
```

```
  RTO receive call register error   : 0
  RTO receive invalid status error  : 0
```

show security alg sip counters tenant TN1

```
user@host> show security alg sip counters tenant TN1
SIP message counters(T:Transmit, RT:Retransmit):
   Method     T   1xx   2xx   3xx  4xx   5xx   6xx
             RT   RT    RT    RT   RT    RT    RT
   INVITE     4    4     3     0    0     0     0
              0    0     0     0    0     0     0
   CANCEL     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
      ACK     3    0     0     0    0     0     0
              0    0     0     0    0     0     0
      BYE     3    0     3     0    0     0     0
              0    0     0     0    0     0     0
 REGISTER     7    0     7     0    0     0     0
              0    0     0     0    0     0     0
  OPTIONS     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
     INFO     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
  MESSAGE     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
   NOTIFY     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
    PRACK     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
  PUBLISH     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
    REFER     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
SUBSCRIBE     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
   UPDATE     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
 BENOTIFY     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
  SERVICE     0    0     0     0    0     0     0
              0    0     0     0    0     0     0
```

```
    OTHER      0    0     0     0    0     0     0
               0    0     0     0    0     0     0
SIP Error Counters:
Total Pkt-in                        :0
Total Pkt dropped on error          :0
Call error                          :0
IP resolve error                    :0
NAT error                           :0
Resource manager error              :0
RR header exceeded max              :0
Contact header exceeded max         :0
Call Dropped due to limit           :0
SIP stack error                     :0
SIP decode error                    :0
SIP unknown method error            :0
SIP dscp marked                     :0
SIP dscp marked error               :0
RTO message sent                    :0
RTO message received                :0
RTO buffer allocation failure       :0
RTO buffer transmit failure         :0
RTO send processing error           :0
RTO receive processing error        :0
RTO receive invalid length          :0
RTO receive call process error      :0
RTO receive call allocation error   :0
RTO receive call register error     :0
RTO receive invalid status error    :0
```

show security alg sip counters tenant all

```
user@host> show security alg sip counters tenant all
SIP message counters(T:Transmit, RT:Retransmit):
   Method     T   1xx   2xx   3xx  4xx   5xx   6xx
              RT   RT    RT    RT   RT    RT    RT
   INVITE      4    4     3     0    0     0     0
               0    0     0     0    0     0     0
   CANCEL      0    0     0     0    0     0     0
               0    0     0     0    0     0     0
      ACK      3    0     0     0    0     0     0
               0    0     0     0    0     0     0
```

```
       BYE      3    0     3     0    0     0     0
                0    0     0     0    0     0     0
  REGISTER      7    0     7     0    0     0     0
                0    0     0     0    0     0     0
   OPTIONS      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
      INFO      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
   MESSAGE      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
    NOTIFY      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
     PRACK      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
   PUBLISH      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
     REFER      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
 SUBSCRIBE      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
    UPDATE      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
  BENOTIFY      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
   SERVICE      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
     OTHER      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
SIP Error Counters:
Total Pkt-in                        :0
Total Pkt dropped on error          :0
Call error                          :0
IP resolve error                    :0
NAT error                           :0
Resource manager error              :0
RR header exceeded max              :0
Contact header exceeded max         :0
Call Dropped due to limit           :0
SIP stack error                     :0
SIP decode error                    :0
SIP unknown method error            :0
SIP dscp marked                     :0
SIP dscp marked error               :0
```

```
RTO message sent                    :0
RTO message received                :0
RTO buffer allocation failure       :0
RTO buffer transmit failure         :0
RTO send processing error           :0
RTO receive processing error        :0
RTO receive invalid length          :0
RTO receive call process error      :0
RTO receive call allocation error   :0
RTO receive call register error     :0
RTO receive invalid status error    :0
```

Release Information

Command modified in Junos OS Release 9.2.

The `node` option introduced in Junos OS Release 9.0.

The `logical-system` option introduced in Junos OS Release 18.2R1.

The `tenant` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

sip (Security) | 27977

clear security alg sip counters | 3863

show security alg sip rate

IN THIS SECTION

- Syntax | 23997
- Description | 23997
- Options | 23997
- Required Privilege Level | 23997
- Output Fields | 23997

Sample Output | 23998
Release Information | 23998

Syntax

```
show security alg sip rate
<node (          node-id         | all | local | primary)>
```

Description

This command displays rate information for Session Initiation Protocol (SIP) Application Layer Protocol (ALG) messages.

Options

- none—Display all SIP ALG rate information.
- `node`—(Optional) For chassis cluster configurations, display SIP rate on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1141 on page 23998 lists the output fields for the `show security alg sip rate` command. Output fields are listed in the approximate order in which they appear.

Table 1141: show security alg sip rate Output Fields

Field Name	Field Description
`CPU ticks`	SIP ALG CPU ticks per microsecond.
`Time taken`	Time, in microseconds, that the last SIP ALG message needed to transit the network.
`Total time`	Total time, in microseconds, during an interval of less than 10 minutes for the specified number of SIP ALG messages to transit the network.
`Rate`	Number of SIP ALG messages per second transiting the network.

Sample Output

show security alg sip rate

```
user@host> show security alg sip rate
CPU ticks per us is 166
Time taken for the last message is 1103 us
Total time taken for 3124 messages is 6221482 us (in less than 10 minutes)
Rate: 502 messages/second
```

Release Information

Command modified in Junos OS Release 9.2.

The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

sip (Security) | 27977

show security alg status

IN THIS SECTION

- Syntax | 23999
- Description | 23999
- Options | 23999
- Required Privilege Level | 24000
- Output Fields | 24000
- Sample Output | 24001
- Release Information | 24007

Syntax

```
show security alg status
```

Description

This command displays the status (enabled/disabled) of the supported Application Layer Gateway (ALG) transactions.

The following list describes the default status on each of these devices:

- On all SRX Series Firewalls — The RSH, SQL, and IKE-ESP ALGs are disabled by default.
- SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices—FTP, TFTP, DNS, MS-RPC, PPTP, SUNRPC, and TALK ALGs are enabled by default. All other ALGs are disabled.
- SRX300, SRX320, SRX340, and SRX380 devices—All supported ALGs except the IKE-ESP, RSH, SQL, and TWAMP ALGs are enabled by default.

Options

- node—Display the ALG status on a specific node.
- logical-system—Display the ALG configuration status for a specific logical system.

- root-logical-system—Display the default ALG configuration status for a root logical system.

Required Privilege Level

view

Output Fields

The following list describes the output fields for the `show security alg status` command. Output fields are listed in the approximate order in which they appear.

- `DNS`–Domain Name Server
- `FTP`–File Transfer Protocol
- `H323`–H.323 protocol
- `MGCP`–Media Gateway Control Protocol
- `MSRPC`–Microsoft remote procedure call
- `PPTP`–Point-to-Point Tunneling Protocol
- `RSH`–UNIX remote shell services
- `RTSP`–Real-Time Streaming Protocol
- `SCCP`–Skinny Client Control Protocol
- `SIP`–Session Initiation Protocol
- `SQL`–Oracle SQL
- `SUNRPC`–Sun Microsystems remote procedure call
- `TALK`–TALK program
- `TFTP`–Trivial File Transfer Protocol
- `IKE-ESP-nat`–Internet Key Exchange and Encapsulating Security Payload. Configure IKE-ESP ALG with NAT
- `TWAMP`–Two-Way Active Measurement Protocol

Sample Output

ALG status

```
user@host> show security alg status
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled
```

ALG status for a specific logical system

```
user@host> show security alg status logical-system LSYS1
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
```

```
  IKE-ESP  : Disabled
  TWAMP    : Disabled
```

ALG status for all logical systems including root logical system

```
user@host> show security alg status logical-system all
Logical system: root-logical-system
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled

Logical system: LSYS3
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
```

```
  IKE-ESP  : Disabled
  TWAMP    : Disabled

Logical system: LSYS1
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled

Logical system: LSYS2
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled

Logical system: LSYS0
ALG Status:
```

```
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled
```

ALG status for a specific tenant system

```
user@host> show security alg status tenant TN1
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Enabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled
```

ALG status for all tenant systems

```
user@host> show security alg status tenant all
Tenant : TN1
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Enabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled

Tenant : TN2
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Enabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled
```

ALG status for all logical systems and tenant systems

```
user@host> show security alg status all-logical-systems-tenants
Logical system : root-logical-system
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled

Logical system : LSYS1
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Enabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled
Tenant : TN1
ALG Status:
```

```
DNS      : Enabled
FTP      : Enabled
H323     : Disabled
MGCP     : Disabled
MSRPC    : Enabled
PPTP     : Enabled
RSH      : Disabled
RTSP     : Enabled
SCCP     : Disabled
SIP      : Enabled
SQL      : Disabled
SUNRPC   : Enabled
TALK     : Enabled
TFTP     : Enabled
IKE-ESP  : Disabled
TWAMP    : Disabled
```

Release Information

Command modified in Junos OS Release 9.5.

Command output is modified to display the details for the TWMAP ALG in Junos OS Release 18.2.

RELATED DOCUMENTATION

alg | 2130

show security alg status logical-system

IN THIS SECTION

- Syntax | 24008
- Description | 24008
- Options | 24008
- Additional Information | 24008

- Required Privilege Level | 24008
- Output Fields | 24009
- Sample Output | 24009
- Release Information | 24011

Syntax

```
show security alg status logical-system
<logical-system-name>
```

Description

Display the ALG status for a specific logical system or for all logical systems on the device.

Options

logical-system-name	Display ALG status for specific logical system.
all	Display ALG status for all logical systems.

Additional Information

The `show security alg status` command is used to view the ALG status in root logical system. The `show security alg status logical-system lsys1` command is used to view the ALG status in logical system lsys1. The `show security alg status logical-system all` command is used to view the ALG status of all existing logical systems.

> **NOTE**: Only users under root logical system can view the ALG status for all logical systems. The keyword logical-system is not required in the command `show security alg status logical-system lsys1` when you log in to a particular logical system.

Required Privilege Level

view

Output Fields

Sample Output

show security alg status logical-system all

```
user@host> show security alg status logical-system all
Logical system: root-logical-system
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled

Logical system: LSYS2
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
```

```
  TFTP     : Enabled
  IKE-ESP  : Disabled

Logical system: LSYS0
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled

Logical system: LSYS1
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled

{secondary:node0}
```

show security alg status logical-system LSYS1

```
user@host> show security alg status logical-system LSYS1
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled

{secondary:node0}
```

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

ALG for Logical Systems

alg | 2130

ALG for Logical Systems

show security application-firewall rule-set

IN THIS SECTION

- Syntax | 24012
- Description | 24012
- Options | 24013
- Required Privilege Level | 24013
- Output Fields | 24013
- Sample Output | 24015
- Sample Output | 24015
- Sample Output | 24016
- Release Information | 24016

Syntax

```
show security application-firewall rule-set (<rule-set-name> | all)
show security application-firewall rule-set (rule-set-name | all) | (logical-system logical-
system-name  | all) | all-logical-systems-tenants | root-logical-system | tenant (tenant-name  |
all)
```

Description

Display information about the specified rule set defined in the application firewall.

The application firewall is defined by a collection of rule sets. A rule set defines the rules that specify match criteria, including dynamic applications, and the action to be taken for matching traffic.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Options

rule-set-name	Display the name of the rule set.
all	(default) Display all rule sets for all logical systems. The user logical system administrator can display all rule sets only for the logical system they can access.
logical-system-name	Display application firewall rule set information for a specific logical system.
root-logical-system	Display application firewall rule set information for the root logical system (primary administrator only).
all-logical-systems-tenants	Display application firewall rule set information for all the logical systems and tenants.
tenant	Display application firewall rule set information for the tenant systems.

Required Privilege Level

view

Output Fields

Table 1142 on page 24013 lists the output fields for the `show security application-firewall rule-set` command. Output fields are listed in the approximate order in which they appear.

Table 1142: show security application-firewall rule-set Output Fields

Field Name	Field Description
Rule-set	Name of the rule set.
Logical system	Name of the logical system of the rule set.
Tenant	Name of the tenant system of the rule set.
Profile	The redirect profile to be used for rules requiring redirection for reject or deny actions.

Table 1142: show security application-firewall rule-set Output Fields *(Continued)*

Field Name	Field Description

Rule	Name of the rule
	• **Dynamic applications**—Name of the applications.
	• **Dynamic application groups**—Name of the application groups.
	• **SSL-Encryption**—Setting for SSL traffic.
	• **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  • **permit**
	  • **deny**
	  • **reject**
	  • **redirect**
	• **Number of sessions matched**—Number of sessions matched with the application firewall rule.
	• **Number of sessions redirected**—Number of sessions redirected by the application firewall rule.

Default rule	The default rule applied when the identified application is not specified in any rules of the rule set. • **Number of sessions matched**—Number of sessions matched with the application firewall default rule. • **Number of sessions redirected**—Number of sessions redirected by the application firewall rule.
Number of sessions with appid pending	Number of sessions that are pending application identification processing

Sample Output

show security application-firewall rule-set my_ruleset1

```
user@host>show security application-firewall rule-set my_ruleset1
Rule-set: my_ruleset1
     Rule: rule1
         Dynamic Applications: junos:FACEBOOK-ACCESS, junos:YMSG
         Dynamic Application Groups: junos:web, junos:chat
         SSL-Encryption: any
         Action: deny or redirect
         Number of sessions matched: 10
         Number of sessions redirected: 10
     Default rule: permit
         Number of sessions matched: 200
         Number of sessions redirected: 0
     Number of sessions with appid pending: 2
```

Sample Output

show security application-firewall rule-set all

```
user@host> show security application-firewall rule-set all

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:TELNET
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r2
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
```

```
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: ls-product-design-rs2
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:FACEBOOK-ACCESS
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

Sample Output

show security application-firewall rule-set ruleset1 tenant all

```
user@host> show security application-firewall rule-set ruleset1 tenant all

Rule-set: ruleset1
    Logical system: root-logical-system
    Tenant: TSYS1
    Rule: rule1
        Dynamic Applications: junos:HTTP, junos:FTP
        SSL-Encryption: any
        Action:permit
        Number of sessions matched: 0
        Number of sessions redirected: 0
Default rule:permit
        Number of sessions matched: 0
        Number of sessions redirected: 0
Number of sessions with appid pending: 0
```

Release Information

Command introduced in Junos OS Release 11.1. Updated in Junos OS Release 12.1X44-D10 with output format changes. Updated in Junos OS Release 12.1X45-D10 with redirection counters.

The `tenant` and `all-logical-systems-tenants` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

clear security application-firewall rule-set statistics

show security application-firewall rule-set logical-system

IN THIS SECTION

- Syntax | 24017
- Description | 24017
- Options | 24018
- Required Privilege Level | 24018
- Output Fields | 24018
- Sample Output | 24019
- Release Information | 24021

Syntax

The primary, or root, administrator can issue the following statements:

```
show security application-firewall rule-set all
show security application-firewall rule-set rule-set-name | all | logical-system logical-system-name | all | root-logical-system [logical-system-name | all ]
```

The user logical system administrator can issue the following statement:

```
show security application-firewall rule-set all
```

Description

Display information about application firewall rule set(s) associated with a specific logical system, all logical systems, or the root logical system configured on a device.

> **NOTE**: The primary administrator can configure and view application firewall rule sets for the root logical system and all user logical systems configured on the device. User logical system administrators can configure and view application firewall rule set information only for the user logical systems for which they have access. For information about primary and user administrator roles in logical systems, see Understanding Logical Systems for SRX Series Services Gateways.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Options

`rule-set-name`—Name of a specific rule set.

`logical-system-name`—Name of a specific logical system.

`all`—(default) Display all rule sets for all logical systems. The user logical system administrator can display all rule sets only for the logical system they can access.

`root-logical-system`—Display application firewall rule set information for the root logical system (primary administrator only).

Required Privilege Level

view

Output Fields

Table 1143 on page 24018 lists the output fields for the `show security application-firewall rule-set logical-system` command. Output fields are listed in the approximate order in which they appear.

Table 1143: show security application-firewall rule-set logical-system Output Fields

Field Name	Field Description
`Rule-set`	Name of the rule set.

Table 1143: show security application-firewall rule-set logical-system Output Fields *(Continued)*

Field Name	Field Description
`Logical system`	Name of the logical system.

`Rule`	Name of the rule.
	• `Dynamic applications`—Name of the applications.
	• `Dynamic application groups`—Name of the application groups.
	• `Action`—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  • `permit`
	  • `deny`
	• `Number of sessions matched`—Number of sessions matched with the application firewall rule.

`Default rule`	The default rule applied when the identified application is not specified in any rules of the rule set. • `Number of sessions matched`—Number of sessions matched with the application firewall default rule.
`Number of sessions with appid pending`	Number of sessions that are pending with the application ID processing.

Sample Output

show security application-firewall rule-set logical-system all

```
root@host> show security application-firewall rule-set logical-system all

Rule-set: root_rs1
    Logical system: root-logical-system
    Rule: r1
        Dynamic Applications: junos:FTP
```

```
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 4

Rule-set: root-rs2
    Logical system: root-logical-system
    Rule: r1
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 10
```

show security application-firewall rule-set all

```
root@host> show security application-firewall rule-set all

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:TELNET
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r2
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4
```

```
Rule-set: ls-product-design-rs2
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:FACEBOOK-ACCESS
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

clear security application-firewall rule-set statistics logical-system | 3866

show security application-tracking counters

IN THIS SECTION

- Syntax | 24021
- Description | 24022
- Required Privilege Level | 24022
- Output Fields | 24022
- Sample Output | 24022
- Release Information | 24023

Syntax

```
show security application-tracking counters
```

Description

Display the status of AppTrack counters. These counters provide number of times an Apptrack message —that is—AppTrack session create, session close, route changes, and volume are generated.

Required Privilege Level

view

Output Fields

Table 1144 on page 24022 lists the output fields for the `show security application-tracking counters` command. Output fields are listed in the approximate order in which they appear.

Table 1144: show security application-tracking counters

Field Name	Field Description
`Session create messages`	The number of log messages generated when a session was created.
`Session close messages`	The number of log messages generated when a session was closed.
`Session volume updates`	The number of log messages generated when an update interval was exceeded.
`Session route updates`	The number of log messages generated when an egress interface was selected based on application carried in the session by APBR.
`Failed messages`	The number of messages that were not generated due to memory or session constraints.

Sample Output

show security application-tracking counters

```
user@host> show security application-tracking counters

Application tracking counters:
```

```
   AppTrack counter type                              Value
Session create messages                                 1
Session close messages                                 1
Session volume updates                                 0
Session route updates                                  1
Failed messages                                        0
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding AppTrack

Example: Configuring AppTrack

show security casb policies

IN THIS SECTION

- Syntax | 24023
- Description | 24024
- Options | 24024
- Required Privilege Level | 24024
- Output Fields | 24024
- Sample Output | 24024
- Release Information | 24025

Syntax

```
show security casb policies policy-name
```

Description

Display information about CASB policies configured on your device.

Options

- policy-name—Name of the specific policy.

Required Privilege Level

view

Output Fields

Following table lists the output fields for the command. Output fields are listed in the approximate order in which they appear.

Table 1145: show security casb policies

Field Name	Field Description
`Casb Policies`	Number of CASB policies on device
`PIC`	Indicate the PIC and FPC used.
`Policy Name`	Name of the CASB policy
`Policy ID`	ID of the CASB policy.

Sample Output

show security flow session session-state

```
user@host> show security casb casb-policies
Casb Policies: 1
```

```
  Policy Name                                   ID
  cp1                                            1
```

```
user@host> show security casb casb-policies policy-name
PIC : FPC 0 PIC 0
 Policy Name: cp1
 Policy ID:   1
```

Release Information

Command introduced in Junos OS Release 23.4R1.

show security datapath-debug capture

IN THIS SECTION

- Syntax | 24025
- Description | 24025
- Required Privilege Level | 24026
- Output Fields | 24026
- Sample Output | 24026
- Release Information | 24026

Syntax

```
show security datapath-debug capture
```

Description

Display details of the data path debugging capture file.

Required Privilege Level

view

Output Fields

Output fields are listed in the approximate order in which they appear.

Sample Output

show security datapath—debug capture

```
user@host> show security datapath-debug capture
Packet 1, len 120: (C0/F0/P0/SEQ:71:lbt)
91 00 00 47 11 00 10 00 9a 14 00 19 03 00 00 00
00 00 00 00 00 01 00 47 10 00 00 00 00 00 00 00
00 1f 12 f8 dd 29 00 21 59 84 f4 01 81 00 02 1e
08 00 45 60 01 f4 00 00 00 00 3f 06 73 9f 01 01
01 02 03 01 01 02 d4 31 d4 31 00 00 00 00 00 00
00 00 50 02 00 00 ff ad 00 00 00 00
Packet 2, len 120: (C0/F0/P0/SEQ:71:lbt)
90 00 00 47 04 00 00 00 00 00 00 00 02 02 00 47
10 00 00 00 00 00 00 00 50 00 a6 1c 00 00 00 00
00 00 00 0a 00 00 00 00 00 00 09 d9 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 1f 12 f8
dd 29 00 21 59 84 f4 01 81 00 02 1e
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show security datapath-debug counter | 24027

Understanding Data Path Debugging for Logical Systems

show security datapath-debug counter

IN THIS SECTION

- Syntax | 24027
- Description | 24027
- Required Privilege Level | 24027
- Output Fields | 24027
- Sample Output | 24027
- Release Information | 24028

Syntax

```
show security datapath-debug counter
```

Description

Display details of the data path debugging counter.

Required Privilege Level

view

Output Fields

Output fields are listed in the approximate order in which they appear.

Sample Output

show security datapath-debug counter

```
user@host> show security datapath-debug counter
Datapath debug counters
Packet Filter 1:
```

```
np-ingress
Chassis 0 FPC 4 : 1
np-ingress
Chassis 0 FPC 3 : 0
np-egress
Chassis 0 FPC 4 : 1
np-egress
Chassis 0 FPC 3 : 0
jexec
Chassis 0 FPC 0 PIC 1: 0
jexec
Chassis 0 FPC 0 PIC 0: 1
lbt
Chassis 0 FPC 0 PIC 1: 0
lbt
Chassis 0 FPC 0 PIC 0: 2
pot
Chassis 0 FPC 0 PIC 1: 0
pot
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show security datapath-debug capture | 24025

Understanding Data Path Debugging for Logical Systems

show security dns-cache

IN THIS SECTION

- Syntax | 24029
- Description | 24029
- Options | 24029

- Required Privilege Level | 24029
- Output Fields | 24029
- Sample Output | 24030
- Sample Output | 24030
- Release Information | 24031

Syntax

```
show security dns-cache  <dns-name dns-name>
```

Description

Display DNS cache information.

> **NOTE**: This command is only available to the primary administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

Options

- `dns-name`—Display DNS cache information for the specified name.

Required Privilege Level

view

Output Fields

Table 1146 on page 24030 lists the output fields for the `show security dns-cache` command. Output fields are listed in the approximate order in which they appear.

Table 1146: show security dns-cache Output Fields

Field Name	Field Description
DNS Name	DNS name.
Address Family	IPv4 or IPv6.
TTL	Time-to-live value.
IP Address	IP address for the DNS name.

Sample Output

show security dns-cache

```
user@host> show security dns-cache
DNS Name: dns1.test.com:
  Address Family: IPv4, TTL: 10
    IP Address: 10.1.1.1
  Address Family: IPv6: TTL = 15
    IP Address: 2001:10.1.1.1
DNS Name: dns2.test.com:
  Address Family: IPv4, TTL: 20
    IP Address: 10.2.2.2
    IP Address: 10.2.2.3
```

Sample Output

show security dns-cache dns-name dns2.test.com

```
user@host> show security dns-cache dns-name dns2.test.com
DNS Name: dns2.test.com:
  Address Family: IPv4, TTL: 20
```

```
    IP Address: 10.2.2.2
    IP Address: 10.2.2.3
```

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

clear security dns-cache | 3869

show security dynamic-address

IN THIS SECTION

- Syntax | 24031
- Description | 24032
- Options | 24032
- Required Privilege Level | 24034
- Output Fields | 24034
- Sample Output | 24035
- Release Information | 24045

Syntax

```
show security dynamic-address
<summary>
<category-name (Blacklist | CC | GeoIP | IPFilter | Infected-Hosts | SecProfiling | Whitelist) >
<family inet | inet6>
<feed-name (feed-name) >
<address-name (address-name) >
<ip-start (starting-IP-address) >
<ip-end (ending-IP-address) >
```

```
<instance (advanced-anti-malware | default | geoip) >
<logical-system (logical-system-name | all)>
<tenant (tenant-name | all)>
```

Description

Displays information about dynamic addresses. Each dynamic address belongs to only one instance. Within that instance is a set of categories to which the dynamic address further belongs.

A dynamic address entry provides dynamic IP address information to security policies. A dynamic address entry is a group of IP addresses, not just a single IP prefix, that can be imported in from external sources. These IP addresses are for specific domains or for entities that have a common attribute such as a particular undesired location that poses a threat. The administrator can then configure security policies to use the DAE within a security policy.

Options

none

Display source category (feed) and dynamic address name for all nodes (primary and backup nodes in case of HA.) The same as the `show security dynamic-address ip-start 1.0.0.0 ip-end 255.255.255.255` command.

summary

(Optional) Display basic information of dynamic-address including their name, feeds, properties and number of IPv4 and IPv6 entries.

If a Dynamic Address Entry (DAE) contains duplicate IP addresses, the output displays the IP address once, ignoring the duplicate entry.

category-name (*category-name*)

(Optional) Display the source category (feed) and dynamic address name for the specified threat type (category name). A category is basically a list of feeds of the same type. The type defines SRX Series enforcement point criteria for feed lookup and enforcement. Supported category names are:

- Blocklist—A list of locations (IP addresses, URLs, etc.) that you do not trust. A blocklist allows everyone access except for the members on the blocklist.
- CC—A list of known C&C servers that are able to send commands to members of a botnet.
- GeoIP—A list giving you the ability to filter traffic to and from specific geographies in the world.
- IPFilter—A list of addresses and ranges of malicious sites that can send junk data.

- Infected-Hosts—A lists of hosts within your network that may have been compromised and require attention. This list is generated from ATP Cloud based on hosts that have downloaded malware.
- SecProfiling—A security intelligence category. You configure the security policy or IDP rule to add the source IP or destination IP of the traffic to the secProfiling feed in the cloud. You can also configure the security policy or IDP to match the source IP or destination IP of the traffic against the secProfiling feed in the cloud.
- Allowlist—A list of locations (IP addresses, URLs, etc.) that you trust. An allowlist denies everyone access except for the members on the allowlist.

family

(Optional) Show the dynamic-address for specified protocol-family. Both IPv4 and IPv6 are displayed if no family is specified. 'inet' and 'inet6' can be combined with other options of the show command. For example, show security dynamic-address family inet6 ip-start 1111::1 ip-end 3333::3.

feed-name

(Optional) User-defined name of the source feed.

address-name

(Optional) The dynamic address name. If you do not specify an `address-name`, then information related to all dynamic addresses downloaded to this SRX Series Firewall is displayed.

ip-start

(Optional) The numerical minimum IP address where you want to investigate. Specifying `ip-start` and optionally `ip-end` (it is not required to specify `ip-end` if you use `ip-start`) is helpful to filter the output to a specific range instead of having to review the entire list which can be very long.

ip-end

(Optional) The numerical maximum IP address where you want to investigate. Specifying `ip-start` and optionally `ip-end` is helpful to filter the output to a specific range instead of having to review the entire list which can be very long. If you specify `ip-end`, you must specify `ip-start`.

instance (*instance-name*)

(Optional) The physically separated database. Supported instance names are:

- advanced-anti-malware—The IP-based allowlists and blocklists.
- default—The default instance holds the following data: blocklist, allowlist, C&C, infected host, and IPfilter.
- geoip—The geoip data.

logical-system (logical-system-name | all)

Perform this operation on all logical systems or on a particular logical system.

`tenant (tenant-name | all)`	Perform this operation on all tenant systems or on a particular tenant system.

Required Privilege Level

View

Output Fields

Table 1147 on page 24034 describes the output fields for the `show security dynamic-address` command. Output fields are listed in the approximate order in which they appear.

Table 1147: show security dynamic-address Output Fields

Field Name	Field Description
Address name	Dynamic address entry name.
Address ID	Internal ID used to uniquely identify the dynamic address entry.
IPv4 entries	The number of IPv4 entries in the specific dynamic address.
IPv6 entries	The number of IPv6 entries in the specific dynamic address.
Category/feed	The threat type associated with the dynamic address. See description of `category-name` and `feed-name` above.
Country code	Two-letter code (alpha-2) ISO 3166 country code in capital ASCII letters; Example" US or CN
Total number of IPv4 entries	The number of IPv4 entries in the database.
Total number of IPv4 entries from feed	The number of IPv4 entries in the feed. An entry in a feed can correspond to multiple entries in the database.
Total number of IPv6 entries	The number of IPv6 entries in the database.

Table 1147: show security dynamic-address Output Fields *(Continued)*

Field Name	Field Description
Total number of IPv6 entries from feed	The number of IPv6 entries in the feed. An entry in a feed can correspond to multiple entries in the database.
Instance default	Total number of default matching entries.
Instance geoip	Total number of geoip data matching entries.
Instance advanced-anti-malware	Total number of the IP-based allowlists and blocklists matching entries.
Session-scan	Status of the session scan.

Sample Output

show security dynamic-address summary

```
user@host> show security dynamic-address summary
Dynamic-address session scan status            : Disable
Hold-interval for dynamic-address session scan : 10 seconds




  Server Name                 : fd_svr1
    Hostname/IP               : https://192.208.132.20
    Update interval           : 1000
    Hold   interval           : 1200
    TLS Profile Name          : tls_profile_1
    User        Name          : ---




    Feed Name                             : fd1
        Mapped dynamic address name       : ad1
```

```
        URL                                : https://192.208.132.20/fd1.gz
        Feed update interval               : 1000     Feed hold interval :1200
        Total update                       : 3
        Total IPv4 entries                 : 1
        Total IPv6 entries                 : 1
        Total download   errors            : 0        Last occurence N/A
        Total db         errors            : 0        Last occurence N/A
        Total other      errors            : 0        Last occurence N/A
        Total ageout                       : 0        Last occurence N/A
        Next update time                   : Mon Aug 31 16:19:44 2020
        Next expire time                   : Mon Aug 31 16:53:04 2020
        Flags                              : 0x0
        Last update file size              : 47
        Last update IPv4 entries           : 1
        Last update IPv6 entries           : 1
        Last update   begin time           : Mon Aug 31 16:03:04 2020
        Last update   end   time           : Mon Aug 31 16:03:04 2020
        Last update   time cost(s)         : --
        Last download begin time           : Mon Aug 31 16:03:04 2020
        Last download end   time           : Mon Aug 31 16:03:04 2020
        Last update   status               : 255
        Last download time cost(s)         : --
    Feed Name                              : fd2
        Mapped dynamic address name        : ad2
        URL                                : https://192.208.132.20/fd2.gz
        Feed update interval               : 1000     Feed hold interval :1200
        Total update                       : 2
        Total IPv4 entries                 : 1
        Total IPv6 entries                 : 1
        Total download   errors            : 0        Last occurence N/A
        Total db         errors            : 0        Last occurence N/A
        Total other      errors            : 0        Last occurence N/A
        Total ageout                       : 0        Last occurence N/A
        Next update time                   : Mon Aug 31 16:19:44 2020
        Next expire time                   : Mon Aug 31 16:53:04 2020
        Flags                              : 0x0
        Last update file size              : 47
        Last update IPv4 entries           : 1
        Last update IPv6 entries           : 1
        Last update   begin time           : Mon Aug 31 16:03:04 2020
        Last update   end   time           : Mon Aug 31 16:03:04 2020
        Last update   time cost(s)         : --
        Last download begin time           : Mon Aug 31 16:03:04 2020
```

```
        Last download end   time          : Mon Aug 31 16:03:04 2020
        Last update   status              : 255
        Last download time cost(s)        : --




    Address name                : ad1
      Address id                : 11
      IPv4 entries              : 1
      IPv6 entries              : 1
      Session-scan              : Disable
      Mapping Feed        : fd1
      Category/feed       : ---     / ---
            property name : feed
                    value : fd1




    Address name                : ad2
      Address id                : 12
      IPv4 entries              : 1
      IPv6 entries              : 1
      Session-scan              : Disable
      Mapping Feed        : fd2
      Category/feed       : ---     / ---
            property name : feed
                    value : fd2




Instance Name                             : default
    Total number of IPv4 entries          : 3
    Total number of IPv4 entries from feed : 2
    Total number of IPv6 entries          : 3
    Total number of IPv6 entries from feed : 2



Instance Name                             : geoip
```

```
    Total number of IPv4 entries           : 0
    Total number of IPv6 entries           : 0



Instance Name                              : advanced-anti-malware
    Total number of IPv4 entries           : 0
    Total number of IPv6 entries           : 0
```

show security dynamic-address summary

```
user@host> show security dynamic-address summary
node1:
--------------------------------------------------------------------------



    Address name          : a1
    Address id            : 11
      IPv4 entries        : 13778
      IPv6 entries        : 0
      Category/feed       : GeoIP   / ---
            property name : country
                    value : AU
                    value : CN


    Address name          : a2
    Address id            : 12
      IPv4 entries        : 0
      IPv6 entries        : 0
      Category/feed       : IPFilter / ---
            property name : test
                    value : test


Total number of IPv4 entries           : 13778
Total number of IPv4 entries from feed : 0
Total number of IPv4 except entries    : 0
Total number of IPv6 entries           : 0
Total number of IPv6 entries from feed : 0
```

show security dynamic-address summary

```
user@host> show security dynamic-address summary
Dynamic-address session scan status            : Disable
Hold-interval for dynamic-address session scan : 10 seconds


    Address name                : ipfilter_zscaler_zpa
      Address id                : 11
      IPv4 entries              : 0
      IPv6 entries              : 0
      Session-scan              : Disable
      Category/feed       : IPFilter / ipfilter_zscaler_zpa
            property name : feed
                    value : ipfilter_zscaler_zpa


Instance Name                              : default
    Total number of IPv4 entries           : 0
    Total number of IPv4 entries from feed : 0
    Total number of IPv6 entries           : 0
    Total number of IPv6 entries from feed : 0

Instance Name                              : geoip
    Total number of IPv4 entries           : 0
    Total number of IPv6 entries           : 0

Instance Name                              : advanced-anti-malware
    Total number of IPv4 entries           : 3
    Total number of IPv6 entries           : 0
```

show security dynamic-address (Junos OS Release 23.2R1)

```
user@host> show security dynamic-address
node0:
--------------------------------------------------------------------------

Instance default Total number of matching entries: 0
No.      IP-start       IP-end        Feed             Address      CountryCode
1        1.0.0.0        1.0.0.255     geoip_country    my-geoip1    AU
```

```
2        1.0.0.0        1.0.0.255     geoip_country    my-geoip2    AU
3        1.0.1.0        1.0.1.255     geoip_country    my-geoip1    CN
4        1.0.1.0        1.0.1.255     geoip_country    my-geoip2    CN
5        1.0.2.0        1.0.3.255     geoip_country    my-geoip1    CN
```

show security dynamic-address summary (Junos OS Release 23.2R1)

```
user@host> show security dynamic-address summary
 node0:
--------------------------------------------------------------------------



Dynamic-address session scan status : Disable
Hold-interval for dynamic-address session scan : 10 seconds




Instance Name : default
Total number of IPv4 entries : 0
Total number of IPv4 entries from feed : 0
Total number of IPv6 entries : 0
Total number of IPv6 entries from feed : 0


Address name : my-geoip1
Address id : 11
IPv4 entries : 17820
IPv6 entries : 6398
Session-scan : Disable
Category/feed : GeoIP / ---
property name : country
value : AU
value : CN


Address name : my-geoip2
Address id : 12
IPv4 entries : 17820
IPv6 entries : 6398
Session-scan : Disable
Category/feed : GeoIP / ---
```

```
property name : country
value : AU
value : CN




Instance Name : geoip
Total number of IPv4 entries : 35640
Total number of IPv6 entries : 12796

Instance Name : advanced-anti-malware
Total number of IPv4 entries : 2
Total number of IPv6 entries : 0
```

show security dynamic-address instance geoip (Junos OS Release 23.2R1)

```
user@host> show security dynamic-address instance geoip
 No.     IP-start             IP-end               Feed
Address                       CountryCode
1       1.0.0.0              1.0.0.255            geoip_country                    my-
geoip1                        AU
2       1.0.0.0              1.0.0.255            geoip_country                    my-
geoip2                        AU
3       1.0.1.0              1.0.1.255            geoip_country                    my-
geoip1                        CN
4       1.0.1.0              1.0.1.255            geoip_country                    my-
geoip2                        CN
5       1.0.2.0              1.0.3.255            geoip_country                    my-
geoip1                        CN
6       1.0.2.0              1.0.3.255            geoip_country                    my-
geoip2                        CN
7       1.0.4.0              1.0.7.255            geoip_country                    my-
geoip1                        AU
8       1.0.4.0              1.0.7.255            geoip_country                    my-
geoip2                        AU
9       1.0.8.0              1.0.15.255           geoip_country                    my-
geoip1                        CN
10      1.0.8.0              1.0.15.255           geoip_country                    my-
geoip2                        CN
11      1.0.32.0             1.0.63.255           geoip_country                    my-
geoip1                        CN
```

```
12      1.0.32.0            1.0.63.255          geoip_country                   my-
geoip2                       CN
13      1.1.0.0             1.1.0.255           geoip_country                   my-
geoip1                       CN
14      1.1.0.0             1.1.0.255           geoip_country                   my-
geoip2                       CN
15      1.1.1.1             1.1.1.1             geoip_country                   my-
geoip1                       AU
16      1.1.1.1             1.1.1.1             geoip_country                   my-
geoip2                       AU
17      1.1.1.4             1.1.1.7             geoip_country                   my-
geoip1                       AU
18      1.1.1.4             1.1.1.7             geoip_country                   my-
geoip2                       AU
19      1.1.1.8             1.1.1.15            geoip_country                   my-
geoip1                       AU
20      1.1.1.8             1.1.1.15            geoip_country                   my-
geoip2                       AU
21      1.1.1.16            1.1.1.31            geoip_country                   my-
geoip1                       AU
22      1.1.1.16            1.1.1.31            geoip_country                   my-
geoip2                       AU
23      1.1.1.32            1.1.1.63            geoip_country                   my-
geoip1                       AU
24      1.1.1.32            1.1.1.63            geoip_country                   my-
geoip2                       AU
25      1.1.1.64            1.1.1.127           geoip_country                   my-
geoip1                       AU
26      1.1.1.64            1.1.1.127           geoip_country                   my-
geoip2                       AU
27      1.1.1.128           1.1.1.255           geoip_country                   my-
geoip1                       AU
28      1.1.1.128           1.1.1.255           geoip_country                   my-
geoip2                       AU
29      1.1.2.0             1.1.3.255           geoip_country                   my-
geoip1                       CN
30      1.1.2.0             1.1.3.255           geoip_country                   my-
geoip2                       CN
31      1.1.4.0             1.1.7.255           geoip_country                   my-
geoip1                       CN
32      1.1.4.0             1.1.7.255           geoip_country                   my-
geoip2                       CN
33      1.1.8.0             1.1.15.255          geoip_country                   my-
```

```
geoip1                          CN
34      1.1.8.0              1.1.15.255           geoip_country                    my-
geoip2                          CN
35      1.1.16.0             1.1.31.255           geoip_country                    my-
geoip1                          CN
36      1.1.16.0             1.1.31.255           geoip_country                    my-
geoip2                          CN
37      1.1.32.0             1.1.63.255           geoip_country                    my-
geoip1                          CN
```

show security dynamic-address instance advanced-anti-malware

```
user@host> show security dynamic-address instance advanced-anti-malware
node1:
--------------------------------------------------------------------------
No.      IP-start        IP-end          Feed             Address
1        5.5.0.0         5.5.0.10        global_whitelist ID-00000003
2        11.11.0.0       11.11.0.10      global_blacklist ID-00000004
```

show security dynamic-address instance geoip

```
user@host> show security dynamic-address instance geoip
node1:
--------------------------------------------------------------------------
No.      IP-start        IP-end          Feed             Address
1        1.0.0.0         1.0.0.255       geoip_country    a1
2        1.0.1.0         1.0.1.255       geoip_country    a1
3        1.0.2.0         1.0.3.255       geoip_country    a1
4        1.0.4.0         1.0.7.255       geoip_country    a1
5        1.0.8.0         1.0.15.255      geoip_country    a1
6        1.0.32.0        1.0.63.255      geoip_country    a1
7        1.1.0.0         1.1.0.255       geoip_country    a1
8        1.1.1.0         1.1.1.255       geoip_country    a1
9        1.1.2.0         1.1.3.255       geoip_country    a1
10       1.1.4.0         1.1.7.255       geoip_country    a1
11       1.1.8.0         1.1.15.255      geoip_country    a1
12       1.1.16.0        1.1.31.255      geoip_country    a1
13       1.1.32.0        1.1.63.255      geoip_country    a1
14       1.2.0.0         1.2.1.255       geoip_country    a1
```

show security dynamic-address category-name Infected-Hosts

```
user@host> show security dynamic-address category-name Infected-Hosts

node1:
--------------------------------------------------------------------------
No.          IP-start            IP-end          Feed                     Address
1               1.0.0.7                 1.0.0.7               Infected-Hosts/1
ID-21500011
2             1.0.0.10             1.0.0.10         Infected-Hosts/1         ID-21500011
3             1.0.0.21             1.0.0.21         Infected-Hosts/1         ID-21500011
4             1.0.0.11             1.0.0.11         Infected-Hosts/1         ID-21500012
5             1.0.0.12             1.0.0.12         Infected-Hosts/1         ID-21500012
6             1.0.0.22             1.0.0.22         Infected-Hosts/1         ID-21500012
7             1.0.0.6                 1.0.0.6             Infected-Hosts/1         ID-21500013
8             1.0.0.9                 1.0.0.9             Infected-Hosts/1         ID-21500013
9             1.0.0.13             1.0.0.13         Infected-Hosts/1         ID-21500013
10             1.0.0.23             1.0.0.23         Infected-Hosts/1         ID-21500013
```

show security dynamic-address logical-system LSYS1

```
user@host> show security dynamic-address logical-system LSYS1

Instance default Total number of matching entries: 0
Instance geoip   Total number of matching entries: 0
Instance advanced-anti-malware Total number of matching entries: 0
```

show security dynamic-address logical-system all

```
user@host> show security dynamic-address logical-system all

Instance default Total number of matching entries: 0
Instance geoip   Total number of matching entries: 0
Instance advanced-anti-malware Total number of matching entries: 0
```

show security dynamic-address tenant TSYS1

```
user@host> show security dynamic-address tenant TSYS1

Instance default Total number of matching entries: 0
Instance geoip   Total number of matching entries: 0
Instance advanced-anti-malware Total number of matching entries: 0
```

show security dynamic-address tenant all

```
user@host> show security dynamic-address tenant all

Instance default Total number of matching entries: 0
Instance geoip   Total number of matching entries: 0
Instance advanced-anti-malware Total number of matching entries: 0
```

Release Information

Command introduced in Junos OS Release 12.1X46-D25.

Command `<family inet | inet6>` introduced in Junos OS Release 18.1.

The `logical-system` and `tenant` options are introduced in Junos OS Release 18.4R1.

show security dynamic-policies

IN THIS SECTION

- Syntax | 24046
- Description | 24046
- Options | 24046
- Required Privilege Level | 24046
- Output Fields | 24047
- Sample Output | 24048

- Sample Output | 24049
- Sample Output | 24051
- Sample Output | 24051
- Sample Output | 24051
- Sample Output | 24052
- Release Information | 24053

Syntax

```
show security dynamic-policies [detail] [from-zone zone] [scope-id id] [to-zone zone]
```

Description

Display dynamic policies downloaded on the group member. This command is supported on SRX100, SRX110, SRX210, SRX220, SRX240, and SRX650 devices.

Options

- none—Display basic information about all policies installed on the group member.
- `detail`—(Optional) Display a detailed view of all of the policies installed on the group member.
- `from-zone`—(Optional) Display information about the policies installed on the group member for the specified source zone.
- `scope-id`—(Optional) Display information about the policies installed on the group member for the specified policy identifier.
- `to-zone`—(Optional) Display information about the policies installed on the group member for the specified destination zone.

Required Privilege Level

view

Output Fields

Table 1148 on page 24047 lists the output fields for the `show security dynamic-policies` command. Output fields are listed in the approximate order in which they appear.

Table 1148: show security dynamic-policies Output Fields

Field Name	Field Description
`Policy`	Name of the applicable Policy.
`State`	Status of the policy: • `enabled`: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • `disabled`: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
`Index`	An internal number associated with the policy.
`Scope Policy`	Policy identifier.
`Sequence number`	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, and 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, and 4.
`Source addresses`	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names. (In this case, only the names are given, not their IP addresses.) For detail display mode, the names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
`Destination addresses`	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.

Table 1148: show security dynamic-policies Output Fields *(Continued)*

Field Name	Field Description
`Application`	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • `IP protocol`: The IP protocol used by the application—for example, TCP, UDP, ICMP. • `ALG`: If an ALG is associated with the session, the name of the ALG. Otherwise, 0. • `Inactivity timeout`: Elapse time without activity after which the application is terminated. • `Source port range`: The low-high source port range for the session application. • `Destination port range`: The low-high destination port range for the session application.
`action-type`	Must be permit.
`Policy Type`	Must be dynamic.
`From zone`	Name of the source zone.
`To zone`	Name of the destination zone.
`Tunnel`	Tunnel name, type (IPsec), and index number.

Sample Output

show security dynamic-policies

```
user@host> show security dynamic-policies
Policy: policy_forward-0001, State: enabled, Index: 1048580, Scope Policy: 4
  Sequence number: 1
  Source addresses:192.168.10.0/24
```

```
  Destination addresses:192.168.20.0/24
    Applications: Unknown
action-type: permit, tunnel:
Policy: policy_forward-0002, State: enabled, Index: 2097156, Scope Policy: 4
  Sequence number: 2
  Source addresses:192.168.10.0/24
  Destination addresses:192.168.20.0/24
    Applications: Unknown
action-type: permit, tunnel:
```

Sample Output

show security dynamic-policies detail

```
user@host> show security dynamic-policies detail
Policy: policy_forward-0001, action-type: permit, State: enabled, Index: 1048580,AI: disabled,
Scope Policy: 4
  Policy Type: Dynamic
  Sequence number: 1
  From zone: Host, To zone: untrust
  Source addresses:192.168.10.0/24
  Destination addresses:192.168.20.0/24
  Application: Unknown
    IP protocol: 6, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [23-23]
  Tunnel: Test Tunnel, Type: IPSec, Index: 1001
Policy: policy_backward-0001, action-type: permit, State: enabled, Index: 1048582,AI: disabled,
Scope Policy: 6
  Policy Type: Dynamic
  Sequence number: 1
  From zone: untrust, To zone: Host
  Source addresses:192.168.10.0/24
  Destination addresses:192.168.20.0/24
  Application: Unknown
    IP protocol: 6, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [80-80]
  Tunnel: Test Tunnel, Type: IPSec, Index: 1003
Policy: policy_internal-0001, action-type: permit, State: enabled, Index: 1048583,AI: disabled,
Scope Policy: 7
```

```
  Policy Type: Dynamic
  Sequence number: 1
  From zone: Internal, To zone: Host
  Source addresses:192.168.1.0/24
  Destination addresses:192.168.20.0/24
  Application: Unknown
    IP protocol: 6, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [80-80]
  Tunnel: Test Tunnel, Type: IPSec, Index: 1005
Policy: policy_external-0001, action-type: permit, State: enabled, Index: 1048584,AI: disabled,
Scope Policy: 8
  Policy Type: Dynamic
  Sequence number: 1
  From zone: Internal, To zone: untrust
  Source addresses:192.168.1.0/24
  Destination addresses:192.168.20.0/24
  Application: Unknown
    IP protocol: 6, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [80-80]
  Tunnel: Test Tunnel, Type: IPSec, Index: 1006
Policy: policy_forward-0002, action-type: permit, State: enabled, Index: 2097156,AI: disabled,
Scope Policy: 4
  Policy Type: Dynamic
  Sequence number: 2
  From zone: Host, To zone: untrust
  Source addresses:192.168.10.0/24
  Destination addresses:192.168.20.0/24
  Application: Unknown
    IP protocol: 6, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [80-80]
  Tunnel: Test Tunnel, Type: IPSec, Index: 1002
Policy: policy_backward-0002, action-type: permit, State: enabled, Index: 2097158,AI: disabled,
Scope Policy: 6
  Policy Type: Dynamic
  Sequence number: 2
  From zone: untrust, To zone: Host
  Source addresses:192.168.10.0/24
  Destination addresses:192.168.20.0/24
  Application: Unknown
    IP protocol: 6, ALG: 0, Inactivity timeout: 0
```

```
      Source port range: [0-0]
      Destination port range: [23-23]
  Tunnel: Test Tunnel, Type: IPSec, Index: 1004
```

Sample Output

show security dynamic-policies from-zone Internal

```
user@host> show security dynamic-policies  from-zone Internal
Policy: policy_internal-0001, State: enabled, Index: 1048583, Scope Policy: 7
  Sequence number: 1
    Applications: Unknown
action-type: permit, tunnel:
Policy: policy_external-0001, State: enabled, Index: 1048584, Scope Policy: 8
  Sequence number: 1
    Applications: Unknown
action-type: permit, tunnel:
```

Sample Output

show security dynamic-policies scope-id 8 from-zone Internal

```
user@host> show security dynamic-policies scope-id 8 from-zone Internal
Policy: policy_external-0001, State: enabled, Index: 1048584, Scope Policy: 8
  Sequence number: 1
    Applications: Unknown
action-type: permit, tunnel:
```

Sample Output

show security dynamic-policies detail from-zone Internal

```
user@host> show security dynamic-policies detail from-zone Internal
Policy: policy_internal-0001, action-type: permit, State: enabled, Index: 1048583,AI: disabled,
Scope Policy: 7
  Policy Type: Dynamic
  Sequence number: 1
```

```
  From zone: Internal, To zone: Host
  Source addresses:192.168.1.0/24
  Destination addresses:192.168.20.0/24
  Application: Unknown
    IP protocol: 6, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [80-80]
  Tunnel: Test Tunnel, Type: IPSec, Index: 1005
Policy: policy_external-0001, action-type: permit, State: enabled, Index: 1048584,AI: disabled,
Scope Policy: 8
  Policy Type: Dynamic
  Sequence number: 1
  From zone: Internal, To zone: untrust
  Source addresses:192.168.1.0/24
  Destination addresses:192.168.20.0/24
  Application: Unknown
    IP protocol: 6, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [80-80]
  Tunnel: Test Tunnel, Type: IPSec, Index: 1006
```

Sample Output

show security dynamic-policies detail from-zone Internal to-zone Host

```
user@host> show security dynamic-policies detail from-zone Internal to-zone Host
Policy: policy_internal-0001, action-type: permit, State: enabled, Index: 1048583,AI: disabled,
Scope Policy: 7
  Policy Type: Dynamic
  Sequence number: 1
  From zone: Internal, To zone: Host
  Source addresses:192.168.1.0/24
  Destination addresses:192.168.20.0/24
  Application: Unknown
    IP protocol: 6, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [80-80]
  Tunnel: Test Tunnel, Type: IPSec, Index: 1005
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security policies | 25028

Group VPNv2 Overview

show security dynamic-vpn users

IN THIS SECTION

- Syntax | 24053
- Description | 24053
- Required Privilege Level | 24054
- Output Fields | 24054
- Sample Output | 24054
- Release Information | 24055

Syntax

```
show security dynamic-vpn users
```

Description

Display all relevant user information. This feature is supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

In Junos OS Release 21.4R1, we've deprecated the `show security dynamic-vpn users` operational mode command and we might remove it completely in a future release.

Required Privilege Level

view

Output Fields

Table 1149 on page 24054 lists the output fields for the `show security dynamic-vpn users` command. Output fields are listed in the approximate order in which they appear.

Table 1149: show security dynamic-vpn users Output Fields

Field Name	Field Description
`User`	Username.
`User-groups`	Remote IPSec VPN usergroups
`Number of connections`	Number of connections currently active.
`Remote IP`	IP address of the client.
`IPsec VPN`	Name of the IPsec VPN.
`IKE gateway`	Name of the IKE gateway.
`IKE ID`	IKE ID configured for the client.
`Status`	Status of the connection.

Sample Output

command-name

```
user@host> show security dynamic-vpn users
User: alice , User group: group-one , Number of connections: 1
Remote IP: 192.168.2.10
```

```
    IPSEC VPN: dyn_vpn2
    IKE gateway: gw2
    IKE ID   : alicegw2.example.net
    IKE Lifetime: 72000
    IPSEC Lifetime: 3600
    Status: CONNECTED
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

- show security dynamic-vpn users terse | 24055
- clear security dynamic-vpn user | 3872
- clear security dynamic-vpn all | 3871
- *Dynamic VPN Overview*

show security dynamic-vpn users terse

IN THIS SECTION

- Syntax | 24056
- Description | 24056
- Required Privilege Level | 24056
- Output Fields | 24056
- Sample Output | 24057
- Release Information | 24057

Syntax

```
show security dynamic-vpn users terse
```

Description

Display all relevant user information. This feature is supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

In Junos OS Release 21.4R1, we've deprecated the `show security dynamic-vpn users terse` operational mode command and we might remove it completely in a future release.

Required Privilege Level

view

Output Fields

Table 1150 on page 24056 lists the output fields for the `show security dynamic-vpn users terse` command. Output fields are listed in the approximate order in which they appear.

Table 1150: show security dynamic-vpn users terse Output Fields

Field Name	Field Description
`User`	Username.
`User-groups`	Remote IPSec VPN usergroups
`Remote IP`	IP address of the client.
`IKE ID`	IKE ID configured for the client.
`Status`	Status of the connection.
`Client Config Name`	Name of the client configuration.

Table 1150: show security dynamic-vpn users terse Output Fields *(Continued)*

Field Name	Field Description
Time Established	Time that the user connection was established.

Sample Output

command-name

```
user@host> show security dynamic-vpn users terse

User   User     Remote       IKE      Status    IKE      IPSEC    Client Time
      Groups    IP           ID                 Lifetime Lifetime Config Established
                                                               Name
alice group-one 192.168.2.10 alicegw2.CONNECTED 72000    3600     group   Wed
                             example.                                    Aug 8 10:
                             net                                         26:39 2012
```

Release Information

This command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show security dynamic-vpn users | 24053

clear security dynamic-vpn user | 3872

clear security dynamic-vpn all | 3871

Dynamic VPN Overview

show security firewall-authentication history

IN THIS SECTION

- Syntax | 24058
- Description | 24058
- Options | 24058
- Required Privilege Level | 24059
- Output Fields | 24059
- Sample Output | 24060
- Sample Output | 24060
- Release Information | 24061

Syntax

```
show security firewall-authentication history
<address (address)>
<from-zone (from-zone)>
<identifier (identifier)>
<logical-system (logical-system-name | all)>
<node (node-id | all | local | primary)>
<root-logical-system (address | from-zone | identifier | tenant | to-zone)>
<tenant (tenant-name | all)>
<to-zone (to-zone)>
```

Description

Displays security firewall authentication user history information and verify the number of firewall users who successfully authenticated and the number of firewall users who failed to log in.

Options

- none—Display history of firewall authentication information.
- address—Display authentication entries based on IP address.

- from-zone—Display authentication entries matching the given source zone, null for web-authentication and userfw-authentication.
- identifier—Display authentication entries by user identifier.
- logical-system—Display firewall authentication tables based on logical system name.
- `node`—(Optional) For chassis cluster configurations, display all firewall authentication history on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.
- root-logical-system—Display firewall authentication tables for root logical system.
- tenant—Display firewall authentication tables based on tenant name.
- to-zone—Display authentication entry matching the given destination zone, null for web-auth and userfw-auth.

Required Privilege Level

view

Output Fields

Table 1151 on page 24059 lists the output fields for the `show security firewall-authentication history` command. Output fields are listed in the approximate order in which they appear.

Table 1151: show security firewall-authentication history Output Fields

Field Name	Field Description
`Authentications`	Number of authentications.
`Id`	Identification number.
`Source IP`	IP address of the authentication source.

Table 1151: show security firewall-authentication history Output Fields *(Continued)*

Field Name	Field Description
Date	Authentication date.
Time	Authentication time.
Duration	Authentication duration.
Status	Authentication status success or failure.
User	Name of the user.

Sample Output

show security firewall-authentication history

```
user@host> show security firewall-authentication history
History of firewall authentication data:
  Authentications: 1
          Id Source Ip       Date       Time     Duration   Status   User
           1 203.0.113.1     2007-04-03 11:43:06 00:00:45   Success  hello
```

Sample Output

show security firewall-authentication history node all

```
user@host> show security firewall-authentication history node all
node0:
--------------------------------------------------------------------------
History of firewall authentication data:
Authentications: 2
Id Source Ip       Date       Time     Duration   Status   User
1 203.0.113.1      2008-01-04 12:00:10 0:05:49    Success  local1
2 203.0.113.1      2008-01-04 14:36:52 0:01:03    Success  local1
```

```
node1:
--------------------------------------------------------------------------
History of firewall authentication data:
  Authentications: 1
          Id Source Ip       Date       Time     Duration   Status   User
           203.0.113.1       2008-01-04 14:59:43 1193046:06: Success local1
```

show security firewall-authentication history tenant tn1

```
user@host> show security firewall-authentication history tenant tn1
History of firewall authentication data:
  Authentications: 0
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` option is added in Junos OS Release 9.0. The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

> *Understanding Logical System Firewall Authentication*
>
> *Firewall User Authentication Overview*

show security firewall-authentication history address

IN THIS SECTION

- Syntax | 24062
- Description | 24062
- Options | 24062
- Required Privilege Level | 24062
- Output Fields | 24062
- Sample Output | 24064

Sample Output | 24064
Release Information | 24065

Syntax

```
show security firewall-authentication history address ip-address
<node ( node-id | all | local | primary)>
```

Description

Display security firewall authentication history for this source IP address.

Options

- `address` *`ip-address`* —IP address of the authentication source.
- none—Display all firewall authentication history for this address.
- `node`—(Optional) For chassis cluster configurations, display firewall authentication history for this address on a specific node.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1152 on page 24063 lists the output fields for the `show security firewall-authentication history address` command. Output fields are listed in the approximate order in which they appear.

Table 1152: show security firewall-authentication history address Output Fields

Field Name	Field Description
`Username`	User ID.
`Source IP`	IP address of the authentication source.
`Authentication state`	Status of authentication (success or failure).
`Authentication method`	Path chosen for authentication.
`Access start date`	Date when user authenticated.
`Access start time`	Time when user authenticated.
`Duration of user access`	Time duration of the accessing firewall.
`Policy name`	Name of the policy.
`Source zone`	User traffic received from the zone.
`Destination zone`	User traffic destined to the zone.
`Access profile`	Name of profile used for authentication.
`Bytes sent by this user`	Number of bytes sent by the user.
`Bytes received by this user`	Number of bytes received by the user.

Sample Output

show security firewall-authentication history address 198.51.100.17

```
user@host> show security firewall-authentication history address 198.51.100.17
Username: u1
Source IP: 198.51.100.17
Authentication state: Success
Authentication method: Pass-through using HTTP
Access start date: 2007-09-12
Access start time: 15:33:29
Duration of user access: 0:00:48
Policy name: Z1-Z2
Source zone: Z1
Destination zone: Z2
Access profile: profile-local
Bytes sent by this user: 0
Bytes received by this user: 449
```

Sample Output

show security firewall-authentication history address 198.51.100.17 node local

```
user@host> show security firewall-authentication history address 198.51.100.17 node local
node0:
--------------------------------------------------------------------------
Username: local1
Source IP: 198.51.100.17
Authentication state: Success
Authentication method: Pass-through using Telnet
Access start date: 2008-01-04
Access start time: 12:00:10
Duration of user access: 0:05:49
Policy name: POL1
Source zone: z1
Destination zone: z2
Access profile: p1
Bytes sent by this user: 0
Bytes received by this user: 0
Username: local1
```

```
Source IP: 198.51.100.17
Authentication state: Success
Authentication method: Pass-through using Telnet
Access start date: 2008-01-04
Access start time: 14:36:52
Duration of user access: 0:01:03
Policy name: POL1
Source zone: z1
Destination zone: z2
Access profile: p1
Bytes sent by this user: 2178
Bytes received by this user: 4172
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show security firewall-authentication history identifier

IN THIS SECTION

- Syntax | 24066
- Description | 24066
- Options | 24066
- Required Privilege Level | 24066
- Output Fields | 24066
- Sample Output | 24068
- Sample Output | 24068
- Release Information | 24069

Syntax

```
show security firewall-authentication history identifier identifier
<node (node-id | all | local | primary)>
```

Description

Display security firewall authentication history information for the authentication with this identifier.

Options

- `identifier` *`identifier`*—Identifying number of the authentication process.
- none—Display all firewall authentication history information for the authentication with this identifier.
- `node`—(Optional) For chassis cluster configurations, display firewall authentication history on a specific node for the authentication with this identifier.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1153 on page 24067 lists the output fields for the `show security firewall-authentication history identifier` command. Output fields are listed in the approximate order in which they appear.

Table 1153: show security firewall-authentication history identifier Output Fields

Field Name	Field Description
`Username`	User ID.
`Source IP`	IP address of the authentication source.
`Authentication state`	Status of authentication (success or failure).
`Authentication method`	Path chosen for authentication.
`Access start date`	Date when user authenticated.
`Access start time`	Time when user authenticated.
`Duration of user access`	Time duration of the accessing firewall.
`Policy index`	Identification number of the policy.
`Policy name`	Name of the policy.
`Source zone`	User traffic received from the zone.
`Destination zone`	User traffic destined to the zone.
`Access profile`	Name of profile used for authentication.
`Bytes sent by this user`	Number of bytes sent by the user.
`Bytes received by this user`	Number of bytes received by the user.
`Client-groups`	Name of the client group.

Sample Output

show security firewall-authentication history identifier 1

```
user@host> show security firewall-authentication history identifier 1
Username: hello
Source IP: 192.0.2.5
Authentication state: Success
Authentication method: Pass-through using Telnet
Access start date: 2007-04-03
Access start time: 11:43:06
Duration of user access: 00:00:45
Policy index: 4
Source zone: z2
Destination zone: z1
Access profile: profile1
Bytes sent by this user: 0
Bytes received by this user: 1050
Client-groups: Sunnyvale Bengaluru
```

Sample Output

show security firewall-authentication identifier 1 node primary

```
user@host> show security firewall-authentication history identifier 1 node primary
node0:
--------------------------------------------------------------------------
Username: local1
Source IP: 192.0.2.5
Authentication state: Success
Authentication method: Pass-through using Telnet
Access start date: 2008-01-04
Access start time: 12:00:10
Duration of user access: 0:05:49
Policy name: POL1
Source zone: z1
Destination zone: z2
Access profile: p1
```

```
Bytes sent by this user: 0
Bytes received by this user: 0
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show security firewall-authentication jims

IN THIS SECTION

- Syntax | 24069
- Description | 24069
- Required Privilege Level | 24070
- Output Fields | 24070
- Sample Output | 24070
- Sample Output | 24070
- Release Information | 24071

Syntax

```
show security firewall-authentication jims (statistics | display)
```

Description

Display statistics of primary and secondary Juniper Identity Management Service (JIMS) server.

Required Privilege Level

view

Output Fields

show security firewall-authentication jims (statistics | display) on page 24070 Output Fields lists the output fields for the `show security firewall-authentication jims (statistics | display)` command. Output fields are listed in the approximate order in which they appear.

Table 1154: show security firewall-authentication jims (statistics | display) Output Fields

Field Name	Field Description
`Push success counter`	Number of authentication entries successfully pushed to JIMS server.
`Push failure counter`	Number of authentication entries failed to be pushed to JIMS server.

Sample Output

show security firewall-authentication jims statistics

```
user@host> show security firewall-authentication jims statistics
  Push success counter: 1
  Push failure counter: 0
```

Sample Output

Starting in Junos OS Release 18.3R2, the output for `show security firewall-authentication jims statistics` operational command is changed to display the statistics of both primary and secondary JIMS server.

show security firewall-authentication jims statistics

```
user@host> show security firewall-authentication jims statistics
node0:
--------------------------------------------------------------------------
```

```
Primary server:
  Push success counter: 0
  Push failure counter: 0

Secondary server:
  Push success counter: 0
  Push failure counter: 0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Configure Juniper Identity Management Service to Obtain User Identity Information

show security firewall-authentication users

IN THIS SECTION

- Syntax | 24072
- Description | 24072
- Options | 24072
- Required Privilege Level | 24073
- Output Fields | 24073
- Sample Output | 24074
- Sample Output | 24074
- Sample Output | 24074
- show security firewall-authentication users tenant all | 24075
- Release Information | 24075

Syntax

```
show security firewall-authentication users
<address (ip-address )>
<auth-type (pass-through | user-firewall | web-authentication)>
<from-zone (from-zone)>
<identifier (identifier )>
<logical-system (logical-system-name | all)>
<node (node-id | all | local | primary)>
<root-logical-system (address | auth-type | from-zone | identifier | tenant | to-zone>
<tenant (tenant-name |all)>
<to-zone (to-zone )>
```

Description

Display firewall authentication details about all users and verify the number of firewall users who successfully authenticated and firewall users who failed to log in.

Options

- none—Display details about all firewall authentication users.
- address—Display authentication entries based on ip address.
- auth-type—Display authentication entries matching the given auth-type.
- from-zone—Display authentication entries matching the given source zone, null for web-auth and userfw-auth.
- identifier—Display authentication entries by id.
- logical-system—Display firewall authentication tables based on logical system name.
- `node`—(Optional) For chassis cluster configurations, display firewall authentication details for all users on a specific node.
  - *node-id*—Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

- root-logical-system—Display firewall authentication tables for root logical system.
- tenant—Display firewall authentication tables based on tenant name.
- to-zone—Display authentication entry matching the given destination zone, null for web-auth and userfw-auth.

Required Privilege Level

view

Output Fields

Table 1155 on page 24073 lists the output fields for the `show security firewall-authentication users` command. Output fields are listed in the approximate order in which they appear.

Table 1155: show security firewall-authentication users Output Fields

Field Name	Field Description
`Total users in table`	Gives count of how many entries/users the command will display.
`Id`	Identification number.
`Source IP`	IP address of the authentication source.
`Src zone`	User traffic received from the zone.
`Dst zone`	User traffic destined to the zone.
`Profile`	Name of profile used for authentication.
`Age`	Idle timeout for the user.
`Status`	Authentication status success or failure.
`User`	Name of the user.

Sample Output

show security firewall-authentication users

```
user@host> show security firewall-authentication users
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           1 192.0.2.5/24         z1       z2       p1       0 Success  local1
```

Sample Output

show security firewall-authentication users node 0

```
user@host> show security firewall-authentication users node 0
node0:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           3 192.0.2.5/24        z1       z2       p1           1 Success  local1
```

Sample Output

show security firewall-authentication users node all

```
user@host> show security firewall-authentication users node all
node0:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           3 192.0.2.5        z1       z2       p1           1 Success  local1
node1:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
```

```
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           2 192.0.2.5       z1       z2       p1           1 Success  local1
```

show security firewall-authentication users tenant all

command-name

```
user@host> show security firewall-authentication users tenant all
Firewall authentication data:
  Total users in table: 1
          Id Source Ip                        Src zone Dst zone Profile    Age Status
User
           2 192.0.2.10                      N/A      N/A      test-rad     1 Success
b1
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0. The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show security firewall-authentication users address

IN THIS SECTION

- Syntax | 24076
- Description | 24076
- Options | 24076
- Required Privilege Level | 24076
- Output Fields | 24077

- Sample Output | 24078
- Sample Output | 24078
- Release Information | 24079

Syntax

```
show security firewall-authentication users address ip-address
<node (node-id | all | local | primary)>
```

Description

Display information about the users at the specified IP address that are currently authenticated.

Options

- `address` *`ip-address`*—IP address of the authentication source.
- none—Display all the firewall authentication information for users at this IP address.
- `node`—(Optional) For chassis cluster configurations, display user firewall authentication entries on a specific node.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1156 on page 24077 lists the output fields for the `show security firewall-authentication users address` command. Output fields are listed in the approximate order in which they appear.

Table 1156: show security firewall-authentication users address Output Fields

Field Name	Field Description
`Username`	User ID.
`Source IP`	IP address of the authentication source.
`Authentication state`	Status of authentication (success or failure).
`Authentication method`	Path chosen for authentication.
`Access time remaining`	Duration for which the connection exists.
`Lsys`	The logical system where the traffic was received.
`Source zone`	User traffic received from the zone.
`Destination zone`	User traffic destined to the zone.
`Policy index`	Identification number of the policy.
`Policy name`	Name of the policy.
`Access profile`	Name of profile used for authentication.
`Interface Name`	Name of the interface.
`Bytes sent by this user`	Number of bytes sent by the user.

Table 1156: show security firewall-authentication users address Output Fields *(Continued)*

Field Name	Field Description
`Bytes received by this user`	Number of bytes received by the user.
`Client-groups`	Name of the client group.

Sample Output

show security firewall-authentication users address 192.0.2.9

```
user@host>show security firewall-authentication users address 192.0.2.9
Username: hello
Source IP: 192.0.2.9
Authentication state: Success
Authentication method: Pass-through using Telnet
Access time remaining: 0
Source zone: z2
Destination zone: z1
Policy index: 5
Access profile: profile1
Interface Name: ge-0/0/2.0
Bytes sent by this user: 0
Bytes received by this user: 0
Client-groups: my-group1-example, my-group2-example
```

Sample Output

show security firewall-authentication users address 192.0.2.9 node local

```
user@host> show security firewall-authentication users address 192.0.2.9 node local
node0:
--------------------------------------------------------------------------
Username: local1
Source IP: 192.0.2.9
Authentication state: Success
```

```
Authentication method: Pass-through using Telnet
Age: 2
Access time remaining: 4
Source zone: z1
Destination zone: z2
Policy name: POL1
Access profile: p1
Interface Name: reth1.0
Bytes sent by this user: 614
Bytes received by this user: 1880
```

show security firewall-authentication users address 198.51.100.29

```
user@host> show security firewall-authentication users address 198.51.100.29
Username: hello
Source IP: 198.51.100.29/24
Authentication state: Success
Authentication method: User-firewall
Age: 0
Access time remaining: 10
Lsys: root-logical-system
Source zone: N/A
Destination zone: N/A
Access profile: test
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Understanding User Role Firewalls

show security firewall-authentication users auth-type

IN THIS SECTION

- Syntax | 24080
- Description | 24080
- Options | 24080
- Required Privilege Level | 24080
- Output Fields | 24081
- Sample Output | 24081
- Release Information | 24082

Syntax

```
show security firewall-authentication users auth-type [user-firewall | pass-through | web-
authentication]
```

Description

Display statistics about the users authenticated by the selected method.

Options

- `user-firewall`—Lists all users authenticated for user firewall.
- `pass-through`—Lists all users authenticated by the pass-through method.
- `web-authentication`—Lists all users authenticated by the user authentication method.

Required Privilege Level

view

Output Fields

Table 1157 on page 24081 lists the output fields for the `show security firewall-authentication users auth-type` command. Output fields are listed in the approximate order in which they appear.

Table 1157: show security firewall-authentication users auth-type Output Fields

Field Name	Field Description
`Total users in table`	Total number of users authenticated by this method.
`Id`	The ID assigned to the entry.
`Source IP`	The source IP address of the traffic.
`Src zone`	The source zone of the traffic.
`Dst zone`	The destination zone of the traffic.
`Profile`	The profile used to authenticate the used.
`Age`	The length of time since authentication.
`Status`	The status of the authentication.
`User`	The username associated with the traffic.

Sample Output

show security firewall-authentication users auth-type user-firewall

```
user@host> show security firewall-authentication users auth-type user-firewall
User-firewall authentication data:
  Total users in table: 2
          Id Source Ip   Src zone Dst zone Profile    Age Status   User
```

```
            1 10.208.16.1  N/A      N/A      test        0 Success  jasonliu
            2 10.208.16.6  N/A      N/A      test6       0 Failed   jason
```

show security firewall-authentication users auth-type pass-through

```
user@host> show security firewall-authentication users auth-type pass-through
Pass-through firewall authentication data:
  Total users in table: 1
          Id Source Ip   Src zone Dst zone Profile    Age Status   User
           1 10.208.16.2  zone1    zone2    test2       0 Success  jasonliu2
```

show security firewall-authentication users auth-type pass-through

```
user@host> show security firewall-authentication users auth-type web-authentication
Web firewall authentication data:
  Total users in table: 1
          Id Source Ip   Src zone Dst zone Profile    Age Status   User
           1 10.208.16.3  N/A      N/A      test3       0 Success  jasonliu3
```

Release Information

Command introduced in Junos OS Release 12.1X45-D10

RELATED DOCUMENTATION

Understanding User Role Firewalls

show security firewall-authentication users identifier

IN THIS SECTION

- Syntax | 24083
- Description | 24083

- Options | 24083
- Required Privilege Level | 24083
- Output Fields | 24084
- Sample Output | 24085
- Sample Output | 24085
- Release Information | 24086

Syntax

```
show security firewall-authentication users identifier identifier
<node (node-id | all | local | primary)>
```

Description

Display firewall authentication details about the user with this identification number.

Options

- `identifier` *`identifier`*—Identification number of the user for which to display authentication details.
- `node`—(Optional) For chassis cluster configurations, display the firewall authentication details security firewall authentication entry on a specific node (device) in the cluster for the user with this identification number.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1158 on page 24084 lists the output fields for the `show security firewall-authentication users identifier` command. Output fields are listed in the approximate order in which they appear.

Table 1158: show security firewall-authentication users identifier Output Fields

Field Name	Field Description
`Username`	User ID.
`Source IP`	IP address of the authentication source.
`Authentication state`	Status of authentication (success or failure).
`Authentication method`	Path chosen for authentication.
`Age`	Idle timeout for the user.
`Access time remaining`	Duration for which the connection exists.
`Source zone`	User traffic received from the zone.
`Destination Zone`	User traffic destined to the zone.
`Policy Name`	Name of the policy.
`Access profile`	Name of profile used for authentication.
`Interface Name`	Name of the interface
`Bytes sent by this user`	Number of bytes sent by the user.
`Bytes received by this user`	Number of bytes received by the user.

Sample Output

show security firewall-authentication users identifier 3

```
user@host> show security firewall-authentication users identifier 3
Username: u1
Source IP: 198.51.100.39
Authentication state: Success
Authentication method: Pass-through using HTTP
Age: 1
Access time remaining: 254
Source zone: Z1
Destination zone: Z2
Policy name: Z1-Z2
Access profile: profile-local
Interface Name: ge-0/0/1.0
Bytes sent by this user: 0
Bytes received by this user: 449
```

Sample Output

show security firewall-authentication users identifier 3 node primary

```
user@host> show security firewall-authentication users identifier 3 node primary
node0:
--------------------------------------------------------------------------
Username: local1
Source IP: 198.51.100.39
Authentication state: Success
Authentication method: Pass-through using Telnet
Age: 1
Access time remaining: 5
Source zone: z1
Destination zone: z2
Policy name: POL1
Access profile: p1
Interface Name: reth1.0
Bytes sent by this user: 614
Bytes received by this user: 1880
```

show security firewall-authentication users identifier 10

```
user@host> show security firewall-authentication users  identifier 10
Username: test
Source IP: 192.0.2.231
Authentication state: Success
Authentication method: Web-authentication using HTTP
Age: 1
Access time remaining: 9
Lsys: root-logical-system
Tenant: tenant-aa
Source zone: N/A
Destination zone: N/A
Access profile: test
Bytes sent by this user: 0
Bytes received by this user: 0
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show security flow cp-session

IN THIS SECTION

- Syntax | 24087
- Description | 24087
- Options | 24087
- Required Privilege Level | 24088
- Output Fields | 24088

- Sample Output | 24089
- Sample Output | 24089
- Release Information | 24091

Syntax

```
show security flow cp-session [<filter>] [summary | terse]
<node (          node-id         | all | local | primary)>
```

Description

Display central point session-related flow information. This command is supported on the SRX5800, SRX5600, and SRX5400 devices.

Options

- `conn`-tag—Session connection tag (0..4294967295)
- `destination-port`—Destination port (1..65535)
- `destination-prefix`—Destination prefix
- `family`—Display session by family.
- `logical-system`—Logical-system name
- `node`—(Optional) For chassis cluster configurations, display security flow cp-session information on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all` —Display information about all nodes.
  - `local` —Display information about the local node.
  - `primary`—Display information about the primary node.
- `protocol`—IP protocol number
- `root-logical-system`—Root logical-system (default)

- `source-port`—Source port (1..65535)
- `source-prefix`—Source IP prefix or address
- `summary | terse`–Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1159 on page 24088 lists the output fields for the `show security flow cp-session` command. Output fields are listed in the approximate order in which they appear.

Table 1159: show security flow cp-session Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Maximum sessions	Number of maximum central point sessions.
Maximum inet6 sessions	Number of maximum inet6 central point sessions.
Session ID	Number that identifies the session. Use this ID to get more information about the session.

Table 1159: show security flow cp-session Output Fields *(Continued)*

Field Name	Field Description
Conn Tag	A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U) and the Stream Control Transmission Protocol (STCP) sessions. The connection tag for GTP-U is the tunnel endpoint identifier (TEID) and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Sample Output

show security flow cp-session

```
root> show security flow cp-session
DCP Flow Sessions on FPC0 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC0 PIC1:

Session ID: 10320276, SPU: 1, Valid
  In: 203.0.113.1/1000 --> o    203.0.113.2/2000;udp, Conn Tag: 0x0,
  Out: 0.0.0.0/0 --> 0.0.0.0/0;0, Conn Tag: 0x0,
Total sessions: 1
```

Sample Output

show security flow cp-session summary

```
root> show security flow cp-session summary
DCP Flow Sessions on FPC10 PIC0:
```

```
Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2
Maximum sessions: 7549747
Maximum inet6 sessions: 7549747

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2
Maximum sessions: 7549747
Maximum inet6 sessions: 7549747

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
Maximum sessions: 7549747
Maximum inet6 sessions: 7549747
```

show security flow cp-session terse

```
root> show security flow cp-session terse
DCP Flow Sessions on FPC0 PIC1:
```

```
Session ID: 10000038, SPU: 1, Valid
  In: 203.0.113.6/1 --> 198.51.100.13/1;pim, Conn Tag: 0x0,
  Out: 198.51.100.13/1 --> 203.0.113.6/1;pim, Conn Tag: 0x0,

Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 10.2. Support.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security flow cp-session destination-port

IN THIS SECTION

- Syntax | 24091
- Description | 24092
- Options | 24092
- Required Privilege Level | 24092
- Output Fields | 24092
- Sample Output | 24093
- Release Information | 24094

Syntax

```
show security flow cp-session destination-port destination-port-number [summary | terse]
```

Description

Display central point session-related flow information for the specified destination port.

Options

- *destination-port-number*—Number of the destination port for which to display central point session information.

  Range: 1 through 65,535

- summary | terse—Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1160 on page 24092 lists the output fields for the `show security flow cp-session destination-port` command. Output fields are listed in the approximate order in which they appear.

Table 1160: show security flow cp-session destination-port Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Session ID	Number that identifies the session. Use this ID to get more information about the session.

Table 1160: show security flow cp-session destination-port Output Fields *(Continued)*

Field Name	Field Description
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Sample Output

show security flow cp-session destination-port summary

```
root> show security flow cp-session destination-port 21 summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

```
DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow cp-session destination-port terse

```
root> show security flow cp-session destination-port 21 terse
DCP Flow Sessions on FPC10 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Session ID: 410003298, SPU: 41, Valid
  In: 203.0.113.10/26182 --> 198.51.100.1/21;tcp,
  Out: 198.51.100.1/21 --> 203.0.113.10/26182;tcp,
Total sessions: 1

DCP Flow Sessions on FPC10 PIC2:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 10.2. This command is supported.

RELATED DOCUMENTATION

show security flow cp-session | 24086

show security flow cp-session destination-prefix | 24095

show security flow cp-session destination-prefix

IN THIS SECTION

- Syntax | 24095
- Description | 24095
- Options | 24095
- Required Privilege Level | 24095
- Output Fields | 24096
- Sample Output | 24096
- Release Information | 24098

Syntax

```
show security flow cp-session destination-prefix destination-IP-prefix  [summary | terse]
```

Description

Display central point session-related flow information for the specified destination prefix.

Options

- *destination-IP-prefix* —Destination IP prefix or address for which to display central point session information.

  Range: 1 through 65,535.

- summary | terse—Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1161 on page 24096 lists the output fields for the `show security flow cp-session destination-prefix` command. Output fields are listed in the approximate order in which they appear.

Table 1161: show security flow cp-session destination-prefix Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Session ID	Number that identifies the session. Use this ID to get more information about the session.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Sample Output

show security flow cp-session destination-prefix summary

```
root> show security flow cp-session destination-prefix 60/8 summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
```

```
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow cp-session destination-prefix terse

```
root> show security flow cp-session destination-prefix 60/8 terse
DCP Flow Sessions on FPC10 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Session ID: 420002660, SPU: 42, Valid
```

```
  In: 203.0.113.10/26183 --> 192.0.2.1/21;tcp,
  Out:192.0.2.1/21 --> 203.0.113.10/26183;tcp,
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 10.2. This command is supported.

RELATED DOCUMENTATION

- show security flow cp-session | 24086
- show security flow cp-session destination-port | 24091

show security flow cp-session family

IN THIS SECTION

- Syntax | 24098
- Description | 24099
- Options | 24099
- Required Privilege Level | 24099
- Output Fields | 24099
- Sample Output | 24100
- Release Information | 24101

Syntax

```
show security flow cp-session family family [summary | terse]
```

Description

Display central point session-related flow information for the specified family.

Options

- *family*–Display session by family.
- inet—Display IPv4 sessions.
- inet6–Display IPv6 and IPv6-NATPT sessions.
- summary | terse–Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1162 on page 24099 lists the output fields for the `show security flow cp-session` family command. Output fields are listed in the approximate order in which they appear.

Table 1162: show security flow cp-session family Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Session ID	Number that identifies the session. Use this ID to get more information about the session.

Table 1162: show security flow cp-session family Output Fields *(Continued)*

Field Name	Field Description
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Sample Output

show security flow cp-session family summary

```
root> show security flow cp-session family inet summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

```
DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow cp-session family terse

```
root> show security flow cp-session family inet terse
DCP Flow Sessions on FPC10 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Session ID: 420002660, SPU: 42, Valid
  In: 198.51.100.1/26183 --> 203.0.113.2/21;tcp,
  Out: 203.0.113.2/21 --> 198.51.100.1/26183;tcp,
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 10.2. This command is supported.

RELATED DOCUMENTATION

show security flow cp-session | 24086

show security flow cp-session protocol

IN THIS SECTION

- Syntax | 24102
- Description | 24102
- Options | 24102
- Required Privilege Level | 24103
- Output Fields | 24103
- Sample Output | 24104
- Release Information | 24105

Syntax

```
show security flow cp-session protocol protocol-name  [summary | terse]
```

Description

Display central point session-related flow information for the specified protocol.This command is supported on the SRX1500, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall.

Options

- *protocol-name* —Protocol to use as a central point session filter. Information about the central point session that uses this protocol is displayed.

Possible protocols are:

- ah—IP Security Authentication Header
- egp—Exterior gateway protocol
- esp—IPsec Encapsulating Security Payload
- gre—Generic routing encapsulation
- icmp—Internet Control Message Protocol

- icmp6—Internet Control Message Protocol
- igmp—Internet Group Management Protocol
- ipip—IP over IP
- ospf—Open Shortest Path First
- pim—Protocol Independent Multicast
- rsvp—Resource Reservation Protocol
- sctp—Stream Control Transmission Protocol
- tcp—Transmission Control Protocol
- udp—User Datagram Protocol
- summary | terse–Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1163 on page 24103 lists the output fields for the `show security flow cp-session` protocol command. Output fields are listed in the approximate order in which they appear.

Table 1163: show security flow cp-session protocol Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.

Table 1163: show security flow cp-session protocol Output Fields *(Continued)*

Field Name	Field Description
Session ID	Number that identifies the session. Use this ID to get more information about the session.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Sample Output

show security flow cp-session protocol summary

```
root> show security flow cp-session protocol tcp summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
```

```
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

show security flow cp-session protocol terse

```
root> show security flow cp-session protocol tcp terse
Session ID: 160000015, SPU: 17, Valid
  In: 203.0.113.9/32838 --> 198.51.100.26/21;tcp,
  Out: 198.51.100.26/21 --> 203.0.113.2/32838;tcp,
Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow cp-session | 24086

show security flow cp-session source-port

IN THIS SECTION

- Syntax | 24106
- Description | 24106
- Options | 24106

- Required Privilege Level | 24106
- Output Fields | 24106
- Sample Output | 24107
- Release Information | 24109

Syntax

```
show security flow cp-session source-port source-port-number [summary | terse]
```

Description

Display central point session-related flow information for the specified source-port. This command is supported on the SRX1500, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall.

Options

source-port-number	Number of the source port about which to display central point session information. **Range**: 1 through 65,535
summary | terse	Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1164 on page 24107 lists the output fields for the `show security flow cp-session source-port` command. Output fields are listed in the approximate order in which they appear.

Table 1164: show security flow cp-session source-port Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Session ID	Number that identifies the session. Use this ID to get more information about the session.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Sample Output

show security flow cp-session source-port summary

```
root> show security flow cp-session source-port 7000 summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

```
DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow cp-session source-port terse

```
root> show security flow cp-session source-port 7000 terse
DCP Flow Sessions on FPC10 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Session ID: 420002661, SPU: 42, Valid
  In: 203.0.113.64/7000 --> 192.0.2.9/8000;udp,
  Out: 192.0.2.9/8000 --> 203.0.113.64/7000;udp,
Total sessions: 1
```

```
DCP Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow cp-session | 24086

show security flow cp-session source-prefix | 24109

show security flow cp-session source-prefix

IN THIS SECTION

- Syntax | 24109
- Description | 24109
- Options | 24110
- Required Privilege Level | 24110
- Output Fields | 24110
- Sample Output | 24111
- Release Information | 24112

Syntax

```
show security flow cp-session source-prefix source-IP-prefix [summary | terse]
```

Description

Display central point session related flow information for the specified source-prefix. This is supported on the SRX1500, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall.

Options

- *source-IP-prefix*—Source IP prefix or address for which to display central point session information.
- summary | terse—Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1165 on page 24110 lists the output fields for the `show security flow cp-session source-prefix` command. Output fields are listed in the approximate order in which they appear.

Table 1165: show security flow cp-session source-prefix Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Session ID	Number that identifies the session. Use this ID to get more information about the session.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Sample Output

show security flow cp-session source-prefix summary

```
root> show security flow cp-session source-prefix 203/8 summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessionsons on FPC10 PIC1:

: 0

DCP Flow SessiValid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow cp-session source-prefix terse

```
root> show security flow cp-session source-prefix 203/8 terse
DCP Flow Sessions on FPC10 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Session ID: 420002663, SPU: 42, Valid
  In: 203.0.113.10/7000 --> 198.51.100.2/8000;udp,
  Out: 198.51.100.2/8000 --> 203.0.113.10/7000;udp,
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow cp-session | 24086

show security flow cp-session source-port | 24105

show security flow gate

IN THIS SECTION

- Syntax | 24113
- Description | 24113
- Options | 24113

- Required Privilege Level | 24113
- Output Fields | 24114
- Sample Output | 24115
- Release Information | 24118

Syntax

```
show security flow gate
[<filter>] [brief | summary]
```

Description

This command displays information about temporary openings known as pinholes or gates in the security firewall.

Pinholes are used by applications that commonly have both control and data sessions and must create openings in the firewall for the data sessions based on information from the parent sessions.

Options

- destination-port—Destination port
- destination-prefix—Destination IP prefix or address
- protocol—IP protocol number
- source-port—Source port
- source-prefix—Source IP prefix or address
- brief | summary—Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1166 on page 24114 lists the output fields for the `show security flow gate` command. Output fields are listed in the approximate order in which they appear.

Table 1166: show security flow gate Output Fields

Field Name	Field Description
`Hole`	Range of flows permitted by the pinhole.
`Translated`	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
`Protocol`	Application protocol, such as UDP or TCP.
`Application`	Name of the application.
`Age`	Idle timeout for the pinhole.
`Flags`	Internal debug flags for the pinhole.
`Zone`	Incoming zone.
`Reference count`	Number of resource manager references to the pinhole.
`Resource`	Resource manager information about the pinhole.
`Valid gates`	Number of valid gates.
`Pending gates`	Number of pending gates.
`Invalidated gates`	Number of invalid gates.

Table 1166: show security flow gate Output Fields *(Continued)*

Field Name	Field Description
`Gates in other states`	Number of gates in other states.
`Total gates`	Number of gates in total.
`Maximum gates`	Number of maximum gates

Sample Output

show security flow gate

```
user@host> show security flow gate
Hole: 0.0.0.0-0.0.0.0/0-0->40.1.1.198.51.100.252/64515-64515
Translated: 0.0.0.0/0->10.0.31.161/25415
Protocol: udp
Application: none/0
Age: 101 seconds
Flags: 0xe001
Zone: untrust
Reference count: 1
Resource: 5-1024-8185
Hole: 0.0.0.0-0.0.0.0/0-198.51.100.252/1046-1046
Translated: 198.51.100.252/36039-> 203.0.113.1/5060
Protocol: udp
Application: junos-sip/63
Age: 65535 seconds
Flags: 0xe200
Zone: untrust
Reference count: 1
Resource: 5-1024-8189
Hole: 0.0.0.0-0.0.0.0/0-0->198.51.100.252-198.51.100.252/24101-24101
Translated: 0.0.0.0/0-> 198.51.100.252/24101
Protocol: udp
Application: none/0
Age: 93 seconds
Flags: 0xe001
```

```
Zone: trust
Reference count: 1
Resource: 5-1024-8188
Hole: 0.0.0.0-0.0.0.0/0-0->40.1.1.5-198.51.100.252/24100-24100
Translated: 0.0.0.0/0->198.51.100.252/24100
Protocol: udp
Application: none/0
Age: 93 seconds
Flags: 0xe001
Zone: trust
Reference count: 1
Resource: 5-1024-8191
Hole: 0.0.0.0-0.0.0.0/0-0->198.51.100.252-198.51.100.252/5060-5060
Translated: 0.0.0.0/0->198.51.100.252/5060
Protocol: udp
Application: junos-sip/63
Age: 65535 seconds
Flags: 0xe200
Zone: trust
Reference count: 1
Resource: 5-1024-8190
```

show security flow gate brief

```
root> show security flow gate brief
Flow Gates on FPC4 PIC1:

Hole: 192.0.2.1-192.0.2.1/0->192.0.2.100-192.0.2.100/38143-38143
  Translated: 192.0.2.1->192.0.2.100/38143
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65532 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
```

```
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate summary

```
root> show security flow gate summary
Flow Gates on FPC4 PIC1:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC5 PIC1:
```

```
Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

Release Information

Command introduced in Junos OS Release 8.5.

Filter and display options added in Junos OS Release 10.2.

RELATED DOCUMENTATION

- show security flow gate brief node | 24118
- show security flow gate destination-port | 24127
- show security flow gate destination-prefix | 24131
- show security flow gate protocol | 24135
- show security flow gate summary node | 24140

show security flow gate brief node

IN THIS SECTION

- Syntax | 24119
- Description | 24119
- Options | 24119
- Required Privilege Level | 24119
- Output Fields | 24119
- Sample Output | 24121
- Release Information | 24126

Syntax

```
show security flow gate brief node (node-id | all | local | primary)
```

Description

Display information about temporary openings known as pinholes or gates in the security firewall for the specified node options in brief mode.

Options

`node`—(Optional) For chassis cluster configurations, display gate information on a specific node.

- *`node-id`* —Identification number of the node. It can be 0 or 1.
- `all`—Display information about all nodes.
- `local`—Display information about the local node.
- `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1167 on page 24119 lists the output fields for the `show security flow gate brief node` command. Output fields are listed in the approximate order in which they appear.

Table 1167: show security flow gate brief node Output Fields

Field Name	Field Description
`Valid gates`	Number of valid gates.
`Pending gates`	Number of pending gates.
`Invalidated gates`	Number of invalid gates.

Table 1167: show security flow gate brief node Output Fields *(Continued)*

Field Name	Field Description
`Gates in other states`	Number of gates in other states.
`Total gates`	Number of gates in total.
`Hole`	Range of flows permitted by the pinhole.
`Translated`	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
`Protocol`	Application protocol, such as UDP or TCP.
`Application`	Name of the application.
`Age`	Idle timeout for the pinhole.
`Flags`	Internal debug flags for the pinhole.
`Zone`	Incoming zone.
`Reference count`	Number of resource manager references to the pinhole.
`Resource`	Resource manager information about the pinhole.

Sample Output

show security flow gate brief node 0

```
root@antbert> show security flow gate brief node 0
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65518 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate brief node 1

```
root@antbert> show security flow gate brief node 1
node1:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65514 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate brief node all

```
root@antbert> show security flow gate brief node all
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65512 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

node1:
--------------------------------------------------------------------------
```

```
Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65510 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate brief node local

```
root@antbert> show security flow gate brief node local
node0:
--------------------------------------------------------------------------
```

```
Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65504 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate brief node primary

```
root@antbert> show security flow gate brief node primary
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:
```

```
Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65500 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

Release Information

Command introduced in Junos OS Release 8.5; node options added in Junos OS Release 9.0. Filter options added in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow gate | 24112

show security flow gate summary node | 24140

show security flow gate destination-port

IN THIS SECTION

- Syntax | 24127
- Description | 24127
- Options | 24127
- Required Privilege Level | 24128
- Output Fields | 24128
- Sample Output | 24129
- Release Information | 24131

Syntax

```
show security flow gate destination-port destination-port-number [brief | summary]
```

Description

Display information about temporary openings known as pinholes or gates in the security firewall that for the specified destination port.

> **NOTE**: Destination port filter matches the gate only if the given port falls within the range of ports specified in the gate.

Options

- *destination-port-number*—Number of the destination port for which to display gate information.

  Range: 1 through 65,535

- brief | summary—Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1168 on page 24128 lists the output fields for the `show security flow gate destination-port` command. Output fields are listed in the approximate order in which they appear.

Table 1168: show security flow gate destination-port Output Fields

Field Name	Field Description
Hole	Range of flows permitted by the pinhole.
Translated	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
Protocol	Application protocol, such as UDP or TCP.
Application	Name of the application.
Age	Idle timeout for the pinhole.
Flags	Internal debug flags for the pinhole.
Zone	Incoming zone.
Reference count	Number of resource manager references to the pinhole.
Resource	Resource manager information about the pinhole.
Valid gates	Number of valid gates.

Table 1168: show security flow gate destination-port Output Fields *(Continued)*

Field Name	Field Description
Pending gates	Number of pending gates.
Invalidated gates	Number of invalid gates.
Gates in other states	Number of gates in other states.
Total gates	Number of gates in total.
Maximum gates	Number of maximum gates.

Sample Output

show security flow gate destination-port brief

```
root> show security flow gate destination-port 33253 brief
Flow Gates on FPC4 PIC1:

Hole: 40.0.0.111-40.0.0.111/0-0->30.0.0.100-30.0.0.100/33253-33253
  Translated: 40.0.0.111/0->30.0.0.100/33253
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65526 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC5 PIC0:
```

```
Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate destination-port summary

```
root> show security flow gate destination-port 33253 summary
Flow Gates on FPC4 PIC1:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
```

```
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow gate | 24112

show security flow gate destination-prefix | 24131

show security flow gate destination-prefix

IN THIS SECTION

- Syntax | 24131
- Description | 24132
- Options | 24132
- Required Privilege Level | 24132
- Output Fields | 24132
- Sample Output | 24133
- Release Information | 24135

Syntax

```
show security flow gate destination-prefix destination-IP-prefix [brief | summary]
```

Description

Display information about temporary openings known as pinholes or gates in the security firewall for the specified destination prefix.

NOTE: Destination prefix must match both the starting and ending address in the gate.

Options

- *destination-IP-prefix*—Destination IP prefix or address for which to display gate information.
- brief | summary—Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1169 on page 24132 lists the output fields for the `show security flow gate destination-prefix` command. Output fields are listed in the approximate order in which they appear.

Table 1169: show security flow gate destination-prefix Output Fields

Field Name	Field Description
Hole	Range of flows permitted by the pinhole.
Translated	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
Protocol	Application protocol, such as UDP or TCP.
Application	Name of the application.

Table 1169: show security flow gate destination-prefix Output Fields *(Continued)*

Field Name	Field Description
Age	Idle timeout for the pinhole.
Flags	Internal debug flags for the pinhole.
Zone	Incoming zone.
Reference count	Number of resource manager references to the pinhole.
Resource	Resource manager information about the pinhole.
Valid gates	Number of valid gates.
Pending gates	Number of pending gates.
Invalidated gates	Number of invalid gates.
Gates in other states	Number of gates in other states.
Total gates	Number of gates in total.

Sample Output

show security flow gate destination-prefix brief

```
root> show security flow gate destination-prefix 192.0.2.1 brief
Hole: 203.0.113.1-203.0.113.1/0-0->192.0.2.1-192.0.2.1/37308-37308
  Translated: 203.0.113.1/0->192.0.2.1/37308
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65456 seconds
  Flags: 0x0080
  Zone: trust
```

```
  Reference count: 1
  Resource: 1-24575-86015

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate destination-prefix summary

```
root> show security flow gate destination-prefix 192.0.2.1 summary
Flow Gates on FPC4 PIC1:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
```

```
Gates in other states: 0
Total gates: 0

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow gate | 24112

show security flow gate destination-port | 24127

show security flow gate protocol

IN THIS SECTION

- Syntax | 24136
- Description | 24136
- Options | 24136
- Required Privilege Level | 24137
- Output Fields | 24137
- Sample Output | 24138
- Release Information | 24140

Syntax

```
show security flow gate protocol protocol-name [brief | summary]
```

Description

Display information about temporary openings known as pinholes or gates in the security firewall for the specified protocol.

Options

- *protocol-name* —Protocol to use as a gate filter. Information about gates that use this protocol is displayed.

Possible protocols are:

- ah—IP Security Authentication Header
- egp—Exterior gateway protocol
- esp—IPsec Encapsulating Security Payload
- gre—Generic routing encapsulation
- icmp—Internet Control Message Protocol
- icmp6—Internet Control Message Protocol
- igmp—Internet Group Management Protocol
- ipip—IP over IP
- ospf—Open Shortest Path First
- pim—Protocol Independent Multicast
- rsvp—Resource Reservation Protocol
- sctp—Stream Control Transmission Protocol
- tcp—Transmission Control Protocol
- udp—User Datagram Protocol
- brief | summary—Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1170 on page 24137 lists the output fields for the `show security flow gate protocol` command. Output fields are listed in the approximate order in which they appear.

Table 1170: show security flow gate protocol Output Fields

Field Name	Field Description
Hole	Range of flows permitted by the pinhole.
Translated	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
Protocol	Application protocol, such as UDP or TCP.
Application	Name of the application.
Age	Idle timeout for the pinhole.
Flags	Internal debug flags for the pinhole.
Zone	Incoming zone.
Reference count	Number of resource manager references to the pinhole.
Resource	Resource manager information about the pinhole.
Valid gates	Number of valid gates.

Table 1170: show security flow gate protocol Output Fields *(Continued)*

Field Name	Field Description
Pending gates	Number of pending gates.
Invalidated gates	Number of invalid gates.
Gates in other states	Number of gates in other states.
Total gates	Number of gates in total.

Sample Output

show security flow gate protocol brief

```
root> root> show security flow gate protocol tcp brief
Hole: 203.0.113.1-40.0.0.111/0-0->192.0.2.1-192.0.2.1/37308-37308
  Translated: 203.0.113.1/0->30.0.0.100/37308
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65414 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24575-86015

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
```

```
Gates in other states: 0
Total gates: 0

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate protocol summary

```
root> show security flow gate protocol tcp summary
Flow Gates on FPC4 PIC1:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow gate | 24112

show security flow gate summary node

IN THIS SECTION

- Syntax | 24140
- Description | 24140
- Options | 24140
- Required Privilege Level | 24141
- Output Fields | 24141
- Sample Output | 24142
- Release Information | 24147

Syntax

```
show security flow gate summary node (node-id | all | local | primary)
```

Description

Display information about temporary openings known as pinholes or gates in the security firewall for the specified node options in summary mode.

Options

`node`—(Optional) For chassis cluster configurations, display gate information on a specific node.

- *node-id* —Identification number of the node. It can be 0 or 1.
- `all`—Display information about all nodes.
- `local`—Display information about the local node.
- `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1171 on page 24141 lists the output fields for the `show security flow gate summary node` command. Output fields are listed in the approximate order in which they appear.

Table 1171: show security flow gate summary node Output Fields

Field Name	Field Description
`Valid gates`	Number of valid gates.
`Pending gates`	Number of pending gates.
`Invalidated gates`	Number of invalid gates.
`Gates in other states`	Number of gates in other states.
`Total gates`	Number of gates in total.
`Hole`	Range of flows permitted by the pinhole.
`Translated`	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port

Table 1171: show security flow gate summary node Output Fields *(Continued)*

Field Name	Field Description
Protocol	Application protocol, such as UDP or TCP.
Application	Name of the application.
Age	Idle timeout for the pinhole.
Flags	Internal debug flags for the pinhole.
Zone	Incoming zone.
Reference count	Number of resource manager references to the pinhole.
Resource	Resource manager information about the pinhole.

Sample Output

show security flow gate summary node 0

```
root@antbert> show security flow gate summary node 0
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:
```

```
Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

show security flow gate summary node 1

```
root@antbert> show security flow gate summary node 1
node1:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC4 PIC1:
```

```
Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

show security flow gate summary node all

```
root@antbert> show security flow gate summary node all
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

```
node1:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

show security flow gate summary node local

```
root@antbert> show security flow gate summary node local
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
```

```
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

show security flow gate summary node primary

```
root@antbert> show security flow gate summary node primary
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:

Valid gates: 1
Pending gates: 0
```

```
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

Release Information

Command introduced in Junos OS Release 8.5; node options added in Junos OS Release 9.0. Filter options added in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow gate brief node | 24118

show security flow ip-action

IN THIS SECTION

- Syntax | 24148
- Description | 24148
- Options | 24148
- Required Privilege Level | 24149
- Output Fields | 24149
- Sample Output | 24150
- Release Information | 24158

Syntax

```
show security flow ip-action  [ <filter> ] [ summary family (inet  | inet6) ]
```

Description

Display the current IP-action settings, based on filtered options, for IP sessions running on the device.

Options

- *filter*—Filter the display based on the specified criteria.

  The following filters display those sessions that match the criteria specified by the filter. Refer to the sample output for filtered output examples.

  `all` | [*filter*]—All active sessions on the device.

  `destination-port` *`destination-port`*—Destination port number of the traffic. Range is 1 through 65,535.

  `destination-prefix` *`destination-prefix`*—Destination IP prefix or address.

  `family (inet | inet6)` [*filter*]—IPv4 traffic or IPv6-NATPT traffic and filtered options.

  `logical-system` *`logical-system-name`* | `all` [*filter*]—Specified logical system or all logical systems.

  `protocol` *`protocol-name`* / *`protocol-number`* [*filter*]—Protocol name or number and filtered options.

  - `ah` or `51`
  - `egp` or `8`
  - `esp` or `50`
  - `gre` or `47`
  - `icmp` or `1`
  - `icmp6` or `58`
  - `ipip` or `4`
  - `ospf` or `89`
  - `pim` or `103`
  - `rsvp` or `46`

- `sctp` or `132`
- `tcp` or `6`
- `udp` or `17`

`root-logical-system` [*filter*]—Default logical system information and filtered options.

`source-port` *`source-port`*—Source port number of the traffic. Range is 1 through 65,535.

`source-prefix` *`source-prefix`*—Source IP prefix or address of the traffic.

- `summary` —Summary information about IP-action entries.

  `family`—Display summary of IP-action entries by family. This option is used to filter the output.

  - `inet`—Display summary of IPv4 entries.
  - `inet6`—Display summary of IPv6 entries.

Required Privilege Level

view

Output Fields

Table 1172 on page 24149 lists the output fields for the `show security flow ip-action` command. Output fields are listed in the approximate order in which they appear.

Table 1172: show security flow ip-action Output Fields

Field Name	Field Description
`Src-Addr`	Source address of outbound IP traffic.
`Src-Port`	Source port number of outbound IP traffic.
`Dst-Addr`	Destination address of inbound IP traffic.
`Dst-Port/Proto`	Destination port number and protocol type of inbound IP traffic.
`Timeout (sec)`	Configured timeouts and time remaining for an IP session.

Table 1172: show security flow ip-action Output Fields *(Continued)*

Field Name	Field Description
Zone	Security zone associated with an IP session.
Action	Configured action type, for example, block, close, and notify.
State	The active mode and passive mode describe the states of the ip-action entry.
IPv4 action count	The total number of IPv4 entries.
IPv6 action count	The total number of IPv6 entries.

Sample Output

show security flow ip-action

```
user@host> show security flow ip-action
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
    203.0.113.1       *       203.0.113.4         21/tcp         293/300      *
close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *      203.0.113.4          21/tcp         293/300      *     close
Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *      203.0.113.4          21/tcp         293/300      *     close
Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
```

```
State
203.0.113.1         *       203.0.113.4          21/tcp         293/300      *    close
Passive
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *      203.0.113.4          21/tcp         293/300       *      close
Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *    203.0.113.4          21/tcp         292/300      *        close
Passive
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
 203.0.113.1         *    203.0.113.4           21/tcp         292/300      *      close
Active
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
 IPv6 action count: 0 on FPC0.PIC1
 IPv6 action count: 0 on FPC0.PIC2
 IPv6 action count: 0 on FPC0.PIC3
 IPv6 action count: 0 on FPC1.PIC0
 IPv6 action count: 0 on FPC1.PIC1
 IPv6 action count: 0 on FPC1.PIC2
 IPv6 action count: 0 on FPC1.PIC3
 IPv6 action count: Active mode 0 on all PICs
```

show security flow ip-action destination-port

```
user@host> show security flow ip-action destination-port 21

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *    close
Passive
 IPv4 action count: 1 on FPC0.PIC1
```

```
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *   close
Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *    close
Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4        21/tcp         274/300      *     close
Passive
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *    close
Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1        *        203.0.113.4         21/tcp         274/300      *     close
Passive
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         21/tcp         273/300      *    close
Active
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
 IPv6 action count: 0 on FPC0.PIC1
 IPv6 action count: 0 on FPC0.PIC2
 IPv6 action count: 0 on FPC0.PIC3
 IPv6 action count: 0 on FPC1.PIC0
 IPv6 action count: 0 on FPC1.PIC1
 IPv6 action count: 0 on FPC1.PIC2
```

```
 IPv6 action count: 0 on FPC1.PIC3
 IPv6 action count: Active mode 0 on all PICs
```

show security flow ip-action destination-prefix

```
user@host>  show security flow ip-action destination-prefix 203.0.113.4/8

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *   close
Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1        *        203.0.113.4         21/tcp         245/300      *     close
Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *   close
Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
192.0.2.3         *        203.0.113.4         21/tcp         245/300      *     close
Passive
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
192.0.2.3         *        203.0.113.4         21/tcp         245/300      *     close
Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1        *        203.0.113.4         21/tcp         245/300      *     close
Passive
```

```
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *   close
Active
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action family inet protocol

```
user@host>  show security flow ip-action family inet protocoludp
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         69/udp         287/300     *     close
Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *      close
Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
close        Active
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
close        Passive
 IPv4 action count: 1 on FPC1.PIC1
```

```
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1        *        203.0.113.4         69/udp         287/300      *       close
Passive
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *      close
Passive
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action family inet logical-system all

```
user@host>  show security flow ip-action family inet logical-system all

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *
close        Passive           root-logical-system
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *   close
Passive           root-logical-system
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State             Logical-System
203.0.113.1         *       203.0.113.4         69/udp         267/300      *     close
Passive           root-logical-system
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *    close
Active            root-logical-system
 IPv4 action count: 1 on FPC1.PIC0
```

```
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *    close
Passive           root-logical-system
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State             Logical-System
203.0.113.1         *        203.0.113.4        69/udp         266/300      *     close
Passive           root-logical-system
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State             Logical-System
203.0.113.1         *        203.0.113.4        69/udp         266/300      *     close
Passive           root-logical-system
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action source-prefix

```
user@host>  show security flow ip-action source-prefix 192.0.2.3/8

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *      close
Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *       close
Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        192.0.2.4        69/udp         244/300      *         close
Passive
 IPv4 action count: 1 on FPC0.PIC3
```

```
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
close        Active
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *        close
Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
close        Passiveo
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone       Action
State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
close        Passive
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action summary

```
user@host>  show security flow ip-action summary

 IPv4 action count: 1 on FPC0.PIC1
 IPv4 action count: 1 on FPC0.PIC2
 IPv4 action count: 1 on FPC0.PIC3
 IPv4 action count: 1 on FPC1.PIC0
 IPv4 action count: 1 on FPC1.PIC1
 IPv4 action count: 1 on FPC1.PIC2
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
 IPv6 action count: 0 on FPC0.PIC1
 IPv6 action count: 0 on FPC0.PIC2
```

```
 IPv6 action count: 0 on FPC0.PIC3
 IPv6 action count: 0 on FPC1.PIC0
 IPv6 action count: 0 on FPC1.PIC1
 IPv6 action count: 0 on FPC1.PIC2
 IPv6 action count: 0 on FPC1.PIC3
 IPv6 action count: Active mode 0 on all PICs
```

show security flow ip-action summary family inet

```
user@host>  show security flow ip-action summary inet

 IPv4 action count: 1 on FPC0.PIC1
 IPv4 action count: 1 on FPC0.PIC2
 IPv4 action count: 1 on FPC0.PIC3
 IPv4 action count: 1 on FPC1.PIC0
 IPv4 action count: 1 on FPC1.PIC1
 IPv4 action count: 1 on FPC1.PIC2
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action summary family inet6

```
user@host> show security flow ip-action summary family inet6

 IPv6 action count: 1 on FPC0.PIC1
 IPv6 action count: 1 on FPC0.PIC2
 IPv6 action count: 1 on FPC0.PIC3
 IPv6 action count: 1 on FPC1.PIC0
 IPv6 action count: 1 on FPC1.PIC1
 IPv6 action count: 1 on FPC1.PIC2
 IPv6 action count: 1 on FPC1.PIC3
 IPv6 action count: Active mode 1 on all PICs
```

Release Information

Command introduced in Junos OS Release 10.1. Logical systems option added in Junos OS Release 11.2 . Summary option introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow ip-action | 3887

clear security flow session destination-port | 3901

show security flow pmi statistics

IN THIS SECTION

- Syntax | 24159
- Description | 24159
- Options | 24159
- Required Privilege Level | 24160
- Output Fields | 24160
- Sample Output | 24161
- Release Information | 24162

Syntax

```
show security flow pmi statistics
<node (node-id | all | local | primary)>
```

Description

The `show security flow statistics` command displays information for individual PIC.

Options

- none—Display the security flow statistics information.
- `node`—(Optional) For chassis cluster configurations, display all security flow statistics on a specific node (device) in the cluster.

- *node-id*—Identification number of the node. It can be 0 or 1.
- `all`—Display information about all nodes.
- `local`—Display information about the local node.
- `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1173 on page 24160 lists the output fields for the `show security flow pmi statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1173: show security flow statistics Output Fields

Field Name	Field Description
PMI received packets	Number of PMI inbound packets received.
PMI transmitted packets	Number of PMI outbound packets transmitted.
PMI regular flow path	Number of PMI packets returned using flow path.
PMI dropped packets	Number of PMI packets dropped
PMI encapsulation bytes	Number of PMI bytes encapsulated.
PMI decapsulation bytes	Number of PMI bytes decapsulated.
PMI encapsulation packets	Number of PMI packets encapsulated.
PMI decapsulation packets	Number of PMI packets decapsulated.

Sample Output

show security flow statistics

```
user@host> show security flow pmi statistics
show security flow pmi statistics
  Flow power-mode-IPSec statistics:
  PMI statistics of FPC2 PIC0:
    PMI received packets:                0
    PMI transmitted packets:             0
    PMI regular flow path:               0
    PMI dropped packets:                 0
    PMI encapsulation bytes:             0
    PMI decapsulation bytes:             0
    PMI encapsulation packets:           0
    PMI decapsulation packets:           0

  PMI statistics of FPC2 PIC1:
    PMI received packets:                0
    PMI transmitted packets:             0
    PMI regular flow path:               0
    PMI dropped packets:                 0
    PMI encapsulation bytes:             0
    PMI decapsulation bytes:             0
    PMI encapsulation packets:           0
    PMI decapsulation packets:           0

  PMI statistics summary:
    PMI received packets:                0
    PMI transmitted packets:             0
    PMI regular flow path:               0
    PMI dropped packets:                 0
    PMI encapsulation bytes:             0
    PMI decapsulation bytes:             0
    PMI encapsulation packets:           0
    PMI decapsulation packets:           0
```

show security flow pmi statistics (PowerMode)

```
user@host> show security flow pmi statistics
node0:
--------------------------------------------------------------------------
  Flow power-mode-IPSec statistics:
    PMI received packets:               40
    PMI transmitted packets:            33
    PMI regular flow path:               7
    PMI dropped packets:                 0
    PMI encapsulation bytes:             0
    PMI decapsulation bytes:             0
    PMI encapsulation packets:           0
    PMI decapsulation packets:           0
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Traffic Processing on SRX Series Devices Overview

show security flow session

IN THIS SECTION

- Syntax | **24163**
- Description | **24163**
- Options | **24163**
- Required Privilege Level | **24166**
- Output Fields | **24166**
- Sample Output | **24170**

Release Information | 24176

Syntax

```
show security flow session [filter] [brief | extensive | summary]
<node (node-id  | all | local | primary)>
<nat number>
<nat-port-overload-index number>
<source-nat-pool source-nat-pool-name>
```

Description

Display information about all currently active security sessions on the device. For the normal flow sessions, the `show security flow session` command displays byte counters based on IP header length. However, for sessions in Express Path mode, the statistics are collected from the IOC2 (SRX5K-MPC), IOC3 (SRX5K-MPC3-100G10G and SRX5K-MPC3-40G10G), and IOC4 (SRX5K-IOC4-MRAT and SRX5K-IOC4-10G) ASIC hardware engines and include full packet length with L2 headers. Because of this, the output displays slightly larger byte counters for sessions in Express Path mode than for the normal flow session.

Options

- *filter*—Filter the display by the specified criteria.

  The following filters reduce the display to those sessions that match the criteria specified by the filter. Refer to the specific `show` command for examples of the filtered output.

advanced-anti-malware	Show advanced-anti-malware sessions. For details on the `advanced-anti-malware` option, see the Sky Advanced Threat Prevention CLI Reference Guide.
all-logical-systems-tenants	All multitenancy systems.
application	Predefined application name.
application-firewall	Application firewall enabled.
application-firewall-rule-set	Application firewall enabled with the specified rule set.

application-traffic-control	Application traffic control session.
application-traffic-control-rule-set	Application traffic control rule set name and rule name.
bytes-less-than	Define session's bytes-count less than a value (1..4294967295).
bytes-more-than	Define session's bytes-count more a value (1..4294967295).
conn-tag	Session connection tag (0..4294967295).
content-filtering	Display the content filtering session details.
curr-less-than	Define session's current-timeout value less than a value (1..100000).
curr-more-than	Define session's current-timeout value more than a value (1..100000).
destination-port	Destination port.
destination-prefix	Destination IP prefix or address.
dynamic-application	Dynamic application.
dynamic-application-group	Dynamic application.
duration-less-than	Define session's duration time less than a value (1..100000).
duration-more-than	Define session's duration time more than a value (1..100000).
encrypted	Encrypted traffic.
family	Display session by family.
ha-link	Display HA link session information.
idp	IDP-enabled sessions.
interface	Name of incoming or outgoing interface.
logical-system (all | *logical-system-name*)	Name of a specific logical system or `all` to display all logical systems.
nat	Display sessions with network address translation.
nat-port-overload-index	Displays NAT port overload index; the range is 0 through 127.
node	(Optional) For chassis cluster configurations, display security flow session information on a specific node (device) in the cluster. • *node-id* —Identification number of the node. It can be 0 or 1.

- `all` —Display information about all nodes.
- `local` —Display information about the local node.
- `primary`—Display information about the primary node.

packets-less-than	Define session's packets-count less than a value (1..4294967295).
packets-more-than	Define session's packets-count more than a value (1..4294967295).
plugin-name	Plugin name.
plugin-status	Plugin status.
plugins	Display the flow session information of plugins.
policy-id	Display session information based on policy ID; the range is 1 through 4,294,967,295.
pretty	Display the flow session information in a list to make it easy for you to read and monitor.
protocol	IP protocol number.
resource-manager	Resource manager.
root-logical-system	Display root logical system as default.
security-intelligence	Display security intelligence sessions.
services-offload	Display services offload sessions.
session-identifier	Display session with specified session identifier.
session-state	Session state.
source-nat-pool	Displays the source NAT pool name.
source-port	Source port.
source-prefix	Source IP prefix.
ssl	Display the SSL proxy sessions information.
tenant	Displays the security flow session information for a tenant system.
timeout-less-than	Define session's timeout value less than a value (1..100000).
timeout-more-than	Define session's timeout value more than a value (1..100000).
tunnel	Tunnel sessions.

tunnel-inspection-type	Tunnel inspection type	
	gre	Displays gre tunnel inspection
	ipip	Displays ipip tunnel inspection
	vxlan	Displays vxlan tunnel inspection
vxlan-vni	It only lists the tunnel session which vni matches the one you specify in the command.	
url-category	Display flow session information by url-category.	
vrf-group	Display flow session information by L3VPN VRF Group.	
web-filtering	Display the web filtering sessions details.	

- brief | extensive | summary—Display the specified level of output.
- none—Display information about all active sessions.

Required Privilege Level

view

Output Fields

Table 1174 on page 24166 lists the output fields for the `show security flow session` command. Output fields are listed in the approximate order in which they appear.

Table 1174: show security flow session Output Fields

Field Name	Field Description	Level of Output
`Session ID`	Number that identifies the session. Use this ID to get more information about the session.	brief extensive none
If	Interface name.	brief none

Table 1174: show security flow session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Status of security flow session.	brief extensive none
`Conn Tag`	A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U) and the Stream Control Transmission Protocol (STCP) sessions. The connection tag for GTP-U is the tunnel endpoint identifier (TEID) and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.	brief extensive none
`CP Session ID`	Number that identifies the central point session. Use this ID to get more information about the central point session.	brief extensive none
`Policy name`	Name and ID of the policy that the first packet of the session matched.	brief extensive none
`Timeout`	Idle timeout after which the session expires.	brief extensive none
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).	brief extensive none

Table 1174: show security flow session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bytes	Number of received and transmitted bytes.	brief extensive none
Pkts	Number of received and transmitted packets.	brief extensive none
`Total sessions`	Total number of sessions.	brief extensive none
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).	brief extensive none
`Status`	Session status.	extensive
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.	extensive
`Source NAT pool`	The name of the source pool where NAT is used.	extensive
`Dynamic application`	Name of the application.	extensive
`Application traffic control rule-set`	AppQoS rule set for this session.	extensive
`Rule`	AppQoS rule for this session.	extensive

Table 1174: show security flow session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Maximum timeout`	Maximum session timeout.	extensive
`Current timeout`	Remaining time for the session unless traffic exists in the session.	extensive
`Session State`	Session state.	extensive
`Start time`	Time when the session was created, offset from the system start time.	extensive
`Unicast-sessions`	Number of unicast sessions.	Summary
`Multicast-sessions`	Number of multicast sessions.	Summary
`Services-offload-sessions`	Number of services-offload sessions.	Summary
`Failed-sessions`	Number of failed sessions.	Summary
`Sessions-in-use`	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states	Summary
`Maximum-sessions`	Maximum number of sessions permitted.	Summary

Sample Output

show security flow session

```
root> show security flow session
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 56, Valid
  In: 203.0.113.1/1000 --> 203.0.113.11/2000;udp, Conn Tag: 0x0, If: reth1.0, Pkts: 1, Bytes:
86, CP Session ID: 10320276
  Out: 203.0.113.11/2000 --> 203.0.113.1/1000;udp, Conn Tag: 0x0, If: reth0.0, Pkts: 0, Bytes:
0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session (with default policy)

```
root> show security flow session
Session ID: 36, Policy name: pre-id-default-policy/n, Timeout: 2, Valid
  In: 10.10.10.2/61606 --> 10.10.10.1/179;tcp, Conn Tag: 0x0, If: ge-0/0/2.0, Pkts: 1, Bytes: 64,
  Out: 10.10.10.1/179 --> 10.10.10.2/61606;tcp, Conn Tag: 0x0, If: .local..0, Pkts: 1, Bytes: 40,
```

show security flow session (drop flow)

Shows dropped flows for SRX5400.

```
root> show security flow session
Outgoing wing:   CP session ID: 12, CP sess SPU Id: 4617   1.0.0.1/55069 <- 1.0.0.254/23;6,
Conn, Drop Flow

Tag: 0x0, VRF GRP ID: 0(0), If: xe-1/0/0.0 (7), Flag: 0x40000020, Vector index: 0x00000002
WSF: 1, Diff: 0, Sequence: 0, Ack: 0, Port sequence: 0, FIN sequence: 0, FIN state: 0   Zone Id:
7, NH: 0x40010, NSP tunnel: 0x0, NP info: 0xffthread id:255
```

show security flow session (IPv6 tunnel)

```
root> show security flow session
Session ID: 44, Policy name: N/A, State: Stand-alone, Timeout: N/A, Valid
  In: 9001::4/1 --> 9001::3/1;ipip, Conn Tag: 0x0, If: ge-0/0/0.0, Pkts: 0, Bytes: 0,

Session ID: 45, Policy name: N/A, State: Stand-alone, Timeout: N/A, Valid
  In: 9001::4/1 --> 9001::3/1;ipv6, Conn Tag: 0x0, If: ge-0/0/0.0, Pkts: 0, Bytes: 0,

Session ID: 57, Policy name: default-policy-logical-system-00/2, State: Stand-alone, Timeout:
1796, Valid
  In: 20.0.0.2/37628 --> 30.0.0.2/22;tcp, Conn Tag: 0x0, If: ge-0/0/1.0, Pkts: 22, Bytes: 4409,
  Out: 30.0.0.2/22 --> 20.0.0.2/37628;tcp, Conn Tag: 0x0, If: ip-0/0/0.1, Pkts: 30, Bytes: 5209,

Session ID: 58, Policy name: default-policy-logical-system-00/2, State: Stand-alone, Timeout:
1784, Valid
  In: 2001::2/58602 --> 3001::2/22;tcp, Conn Tag: 0x0, If: ge-0/0/1.0, Pkts: 31, Bytes: 5569,
  Out: 3001::2/22 --> 2001::2/58602;tcp, Conn Tag: 0x0, If: ip-0/0/0.1, Pkts: 28, Bytes: 6249,
Total sessions: 4
```

show security flow session brief

```
root> show security flow session brief
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 62, Valid
  In: 203.0.113.11/1000 --> 203.0.113.1/2000;udp, Conn Tag: 0x0, If: reth1.0, Pkts: 1, Bytes:
86, CP Session ID: 10320276
  Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp, Conn Tag: 0x0, If: reth0.0, Pkts: 0, Bytes:
0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session content-filtering

```
root> show security flow session content-filtering
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 62, Valid
```

```
  In: 192.0.2.0/24/1000 --> 203.0.113.0/2000;udp, Conn Tag: 0x0, If: reth1.0, Pkts: 1, Bytes:
86, CP Session ID: 10320276
  Out: 203.0.113.0/2000 --> 192.0.2.0/24/1000;udp, Conn Tag: 0x0, If: reth0.0, Pkts: 0, Bytes:
0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session extensive

```
root> show security flow session extensive
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Status: Normal, State: Active
Flags: 0x8000040/0x18000000/0x12000003
Policy name: SG/4
Source NAT pool: Null, Application: junos-gprs-gtp-v0-udp/76
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 90, Current timeout: 54
Session State: Valid
Start time: 6704, Duration: 35
   In: 203.0.113.11/1000 --> 201.11.0.100/2000;udp,
    Conn Tag: 0x0, Interface: reth1.0,
    Session token: 0x6, Flag: 0x40000021
    Route: 0x86053c2, Gateway: 201.10.0.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 86
    CP Session ID: 10320276
   Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp,
    Conn Tag: 0x0, Interface: reth0.0,
    Session token: 0x7, Flag: 0x50000000
    Route: 0x86143c2, Gateway: 203.0.113.11, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10320276
Total sessions: 1
```

show security flow session extensive

```
root> show security flow session extensive
Flow Sessions on FPC0 PIC0:

Session ID: 10000059, Status: Normal
Flags: 0x10000/0x0/0x10/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 642, Duration: 369
   In: 3.0.0.2/64387 --> 2.0.0.1/8940;esp,
  Conn Tag: 0x0, Interface: xe-2/0/2.0,
    Session token: 0x7, Flag: 0x80100621
    Route: 0xc0010, Gateway: 2.0.0.2, Tunnel: 0
    ESP/AH frag Rx: 0, Generated: 0
    Inner IPv4 frag Rx: 0, Tx: 0, Generated: 0,
    Inner IPv6 frag Rx: 0, Tx: 0, Generated: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 25, Bytes: 3760
    CP Session ID: 0

Session ID: 10000060, Status: Normal
Flags: 0x10000/0x0/0x10/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 642, Duration: 369
   In: 3.0.0.2/0 --> 2.0.0.1/0;esp,
  Conn Tag: 0x0, Interface: xe-2/0/2.0,
    Session token: 0x7, Flag: 0x621
    Route: 0xc0010, Gateway: 2.0.0.2, Tunnel: 0
    ESP/AH frag Rx: 0, Generated: 0
```

```
    Inner IPv4 frag Rx: 0, Tx: 0, Generated: 0,
    Inner IPv6 frag Rx: 0, Tx: 0, Generated: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 0
Total sessions: 2
```

show security flow session summary

```
root> show security flow session summary
Flow Sessions on FPC10 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC10 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC10 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
```

```
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

show security flow session tunnel-inspection-type

```
root> show security flow session tunnel-inspection-type vxlan
Session ID: 335544369, Policy name: p1/7, Timeout: 2, Valid
In: 192.168.200.100/19183 --> 192.168.200.101/2;icmp, Conn Tag: 0xfcd, If: xe-7/0/0.0, Pkts: 2,
Bytes: 2048, CP Session ID: 30, Tunnel Session ID: 268435486, Type: VXLAN, VNI: 1000
Out: 192.168.200.101/2 --> 192.168.200.100/19183;icmp, Conn Tag: 0xfcd, If: xe-7/0/1.0, Pkts: 2,
Bytes: 2048, CP Session ID: 30, Tunnel Session ID: 268435488, Type: VXLAN, VNI: 1000
```

show security flow session tunnel-inspection-type

```
root> show security flow session vxlan-vni 400
Session ID: 1677861258, Policy name: pset1_p1/6, Timeout: 2, Valid
In: 192.150.0.12/55908 --> 192.160.0.66/80;tcp, Conn Tag: 0xfcd, If: xe-3/0/0.0, Pkts: 5, Bytes:
465, CP Session ID: 7021087,
Type: VXLAN, VNI: 400, Tunnel Session ID: 1680264845
Out: 192.160.0.66/80 --> 192.150.0.12/55908;tcp, Conn Tag: 0xfcd, If: xe-3/0/1.0, Pkts: 3,
Bytes: 328, CP Session ID: 7021087,
Type: VXLAN, VNI: 400, Tunnel Session ID: 1679640460



Session ID: 1678454648, Policy name: pset1_p1/6, Timeout: 2, Valid
In: 192.150.0.13/56659 --> 192.160.0.67/80;tcp, Conn Tag: 0xfcd, If: xe-3/0/0.0, Pkts: 5, Bytes:
465, CP Session ID: 5589311,
Type: VXLAN, VNI: 400, Tunnel Session ID: 1679698941
Out: 192.160.0.67/80 --> 192.150.0.13/56659;tcp, Conn Tag: 0xfcd, If: xe-3/0/1.0, Pkts: 3,
Bytes: 328, CP Session ID: 5589311,
Type: VXLAN, VNI: 400, Tunnel Session ID: 1679872223
```

show security flow session web-filtering

```
root> show security flow session web-filtering
Session ID: 256, Policy name: p/4, Timeout: 1794, Session State: Valid
  In: 198.51.100.0/33170 --> 203.0.113.0/443;tcp, Conn Tag: 0x0, If: ge-0/0/0.0, Pkts: 3, Bytes:
351,
  Out: 203.0.113.0/443 --> 192.0.2.0/13089;tcp, Conn Tag: 0x0, If: ge-0/0/2.0, Pkts: 1, Bytes:
44,
Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 8.5.

Support for filter and view options added in Junos OS Release 10.2.

Application firewall, dynamic application, and logical system filters added in Junos OS Release 11.2.

Policy ID filter added in Junos OS Release 12.3X48-D10.

Support for connection tag added in Junos OS Release 15.1X49-D40.

The `tenant` option introduced in Junos OS Release 18.3R1.

The `tunnel-inspection-type` option is introduced in Junos OS Release 20.4R1.

The `content filtering` and `Web filtering` filtering options are introduced in Junos OS Release 23.1R1.

The `nat`, `nat-port-overload-index`, and `source-nat-pool` are introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session all

show security flow session advanced-anti-malware

IN THIS SECTION

- Syntax | 24177
- Description | 24177
- Options | 24177
- Required Privilege Level | 24177
- Output Fields | 24177
- Sample Output | 24179
- Release Information | 24184

Syntax

```
show security flow session advanced-anti-malware [ brief | extensive | summary ]
```

Description

Display information about all currently active advanced anti-malware (AAMW) sessions on the device.

Options

brief | extensive | summary	Display the specified level of output.
none	Display information about all active AAMW sessions.

Required Privilege Level

View

Output Fields

Table 1175 on page 24178 lists the output fields for the `show security flow session advanced-anti-malware` command. Output fields are listed in alphabetical order.

Table 1175: show security flow session advanced-anti-malware Output Fields

Field Name	Field Description
Application traffic control rule-set	AppQoS rule set for this session.
Conn Tag	A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U) and the Stream Control Transmission Protocol (STCP) sessions. The connection tag for GTP-U is the tunnel endpoint identifier (TEID) and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.
Current timeout	Remaining time for the session unless traffic exists in the session.
Dynamic application	Name of the application.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Maximum timeout	Maximum session timeout.
Policy name	Policy that permitted the traffic.
Rule	AppQoS rule for this session.
Session ID	Number that identifies the session. Use this ID to get more information about the session.
Source NAT pool	The name of the source pool where NAT is used.

Table 1175: show security flow session advanced-anti-malware Output Fields *(Continued)*

Field Name	Field Description
Start time	Time when the session was created, offset from the system start time.
Status	Session status.
Timeout	Idle timeout after which the session expires.
Total sessions	Total number of sessions.
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).

Sample Output

show security flow session advanced-anti-malware

```
user@host> show security flow session advanced-anti-malware
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10002664, Policy name: p1/6, State: Backup, Timeout: 14048, Valid
  In: 6.0.0.100/1 --> 192.88.99.1/1;ipv6, Conn Tag: 0x0, If: reth0.0, Pkts: 0, Bytes: 0, CP
Session ID: 10002826
  Out: 192.88.99.1/1 --> 19.255.7.100/27341;ipv6, Conn Tag: 0x0, If: reth2.0, Pkts: 0, Bytes: 0,
CP Session ID: 20002505

Session ID: 10002677, Policy name: p1/6, State: Backup, Timeout: 1734, Valid
  In: 6.0.0.100/61471 --> 10.155.191.252/53;udp, Conn Tag: 0x0, If: reth0.0, Pkts: 0, Bytes: 0,
CP Session ID: 10002839
  Out: 10.155.191.252/53 --> 19.255.7.100/26957;udp, Conn Tag: 0x0, If: reth2.0, Pkts: 0, Bytes:
0, CP Session ID: 10002839
```

```
Session ID: 10002681, Policy name: p1/6, State: Backup, Timeout: 1758, Valid
  In: 6.0.0.100/53578 --> 10.155.191.252/53;udp, Conn Tag: 0x0, If: reth0.0, Pkts: 0, Bytes: 0,
CP Session ID: 10002843
  Out: 10.155.191.252/53 --> 19.255.7.100/2641;udp, Conn Tag: 0x0, If: reth2.0, Pkts: 0, Bytes:
0, CP Session ID: 20002521
Total sessions: 3

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10011868, Policy name: p1/6, State: Active, Timeout: 1758, Valid
  In: 6.0.0.100/1 --> 192.88.99.1/1;ipv6, Conn Tag: 0x0, If: reth0.0, Pkts: 3, Bytes: 204, CP
Session ID: 10011941
  Out: 192.88.99.1/1 --> 19.255.7.100/27341;ipv6, Conn Tag: 0x0, If: reth2.0, Pkts: 0, Bytes: 0,
CP Session ID: 20002763

Session ID: 10011931, Policy name: p1/6, State: Active, Timeout: 12, Valid
  In: 6.0.0.100/53578 --> 10.155.191.252/53;udp, Conn Tag: 0x0, If: reth0.0, Pkts: 5, Bytes:
310, CP Session ID: 10012001
  Out: 10.155.191.252/53 --> 19.255.7.100/2641;udp, Conn Tag: 0x0, If: reth2.0, Pkts: 0, Bytes:
0, CP Session ID: 20002786
Total sessions: 2

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0

{primary:node0}
```

show security flow session advanced-anti-malware summary

```
user@host> show security flow session advanced-anti-malware summary
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Valid sessions: 3
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 3

Flow Sessions on FPC0 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC0 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Valid sessions: 3
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 3

Flow Sessions on FPC0 PIC2:
```

```
Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC0 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

{primary:node0}
```

show security flow session advanced-anti-malware extensive

```
user@host> show security flow session advanced-anti-malware extensive
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10002664, Status: Normal, State: Backup
Flags: 0x10000000/0x0/0x103
Policy name: p1/6
Source NAT pool: interface
Dynamic application: INCONCLUSIVE,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 14184
Session State: Valid
Start time: 76795, Duration: 226
   In: 6.0.0.100/1 --> 192.88.99.1/1;ipv6,
    Conn Tag: 0x0, Interface: reth0.0,
    Session token: 0x7, Flag: 0x60000621
    Route: 0x86153c2, Gateway: 6.0.0.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
```

```
    Pkts: 0, Bytes: 0
    CP Session ID: 10002826
   Out: 192.88.99.1/1 --> 19.255.7.100/27341;ipv6,
    Conn Tag: 0x0, Interface: reth2.0,
    Session token: 0x6, Flag: 0x60000620
    Route: 0x860b3c2, Gateway: 19.255.254.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 20002505
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10011868, Status: Normal, State: Active
Flags: 0x8000002/0x8000000/0x103
Policy name: p1/6
Source NAT pool: interface
Dynamic application: INCONCLUSIVE,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1576
Session State: Valid
Start time: 76756, Duration: 225
   In: 6.0.0.100/1 --> 192.88.99.1/1;ipv6,
    Conn Tag: 0x0, Interface: reth0.0,
    Session token: 0x7, Flag: 0x40002621
    Route: 0x86173c2, Gateway: 6.0.0.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 68
    CP Session ID: 10011941
   Out: 192.88.99.1/1 --> 19.255.7.100/27341;ipv6,
    Conn Tag: 0x0, Interface: reth2.0,
```

```
    Session token: 0x6, Flag: 0x60002620
    Route: 0x860d3c2, Gateway: 19.255.254.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 20002763
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0

{primary:node0}
```

Release Information

Command introduced in Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

show security flow session

show security flow session application

IN THIS SECTION

- Syntax | 24185
- Description | 24185
- Options | 24185
- Required Privilege Level | 24186
- Output Fields | 24186
- Sample Output | 24188
- Release Information | 24190

Syntax

```
show security flow session application
application-name [brief | extensive | summary]
```

Description

This command displays information about each session of the specified application type.

Options

- *application-name*—Type of application about which to display sessions information. Possible values are:
  - dns–Domain Name System
  - ftp–File Transfer Protocol
  - ignore–Ignore application type
  - mgcp-ca–Media Gateway Control Protocol with Call Agent
  - mgcp-ua–MGCP with User Agent
  - pptp–Point-to-Point Tunneling Protocol
  - q931–ISDN connection control protocol
  - ras–Remote Access Server
  - realaudio–RealAudio
  - rsh–UNIX remote shell services
  - rtsp–Real-Time Streaming Protocol
  - sccp–Skinny Client Control Protocol
  - sip–Session Initiation Protocol
  - sqlnet-v2–Oracle SQLNET
  - talk–TALK program
  - tftp–Trivial File Transfer Protocol
  -

- brief | extensive | summary–Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1176 on page 24186 lists the output fields for the `show security flow session application` command. Output fields are listed in the approximate order in which they appear.

Table 1176: show security flow session application Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Total sessions`	Total number of sessions.
`Status`	Session status.
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.
`Policy name`	Name and ID of the policy that the first packet of the session matched.

Table 1176: show security flow session application Output Fields *(Continued)*

Field Name	Field Description
`Source NAT pool`	The name of the source pool where NAT is used.
`Application`	Name of the application.
`Maximum timeout`	Maximum session timeout.
`Current timeout`	Remaining time for the session unless traffic exists in the session.
`Session State`	Session state.
`Start time`	Time when the session was created, offset from the system start time.
`Unicast-sessions`	Number of unicast sessions.
`Multicast-sessions`	Number of multicast sessions.
`Failed-sessions`	Number of failed sessions.
`Sessions-in-use`	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
`Maximum-sessions`	Number of maximum sessions.

Sample Output

show security flow session application telnet

```
root> show security flow session application telnet
Flow Sessions on FPC4 PIC1:
Total sessions: 0

Flow Sessions on FPC5 PIC0:
Total sessions: 0

Flow Sessions on FPC5 PIC1:

Session ID: 210067547, Policy name: default-policy/2, Timeout: 1796, Valid
  In: 203.0.113.2/32781 --> 192.0.2.5/23;tcp, If: ge-0/0/2.0, Pkts: 10, Bytes: 610
  Out: 192.0.2.5/23 --> 203.0.113.2/32781;tcp, If: ge-0/0/1.0, Pkts: 9, Bytes: 602
Total sessions: 1
```

show security flow session application telnet brief

```
root> show security flow session application telnet brief
Flow Sessions on FPC4 PIC1:
Total sessions: 0

Flow Sessions on FPC5 PIC0:
Total sessions: 0

Flow Sessions on FPC5 PIC1:

Session ID: 210067547, Policy name: default-policy/2, Timeout: 1796, Valid
  In: 203.0.113.2/32781 --> 192.0.2.5/23;tcp, If: ge-0/0/2.0, Pkts: 10, Bytes: 610
  Out: 192.0.2.5/23 --> 203.0.113.2/32781;tcp, If: ge-0/0/1.0, Pkts: 9, Bytes: 602
Total sessions: 1
```

show security flow session application telnet extensive

```
root> show security flow session application telnet extensive
Flow Sessions on FPC4 PIC1:
```

```
Total sessions: 0

Flow Sessions on FPC5 PIC0:
Total sessions: 0

Flow Sessions on FPC5 PIC1:

Session ID: 210067547, Status: Normal
Flag: 0x40
Policy name: default-policy/2
Source NAT pool: Null, Application: junos-telnet/10
Maximum timeout: 1800, Current timeout: 1788
Session State: Valid
Start time: 670184, Duration: 33
   In: 203.0.113.2/32781 --> 192.0.2.5/23;tcp,
    Interface: ge-0/0/2.0,
    Session token: 0x180, Flag: 0x0x21
    Route: 0x60010, Gateway: 203.0.113.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 10, Bytes: 610
   Out: 192.0.2.5/23 --> 203.0.113.2/32781;tcp,
    Interface: ge-0/0/1.0,
    Session token: 0x1c0, Flag: 0x0x20
    Route: 0x70010, Gateway: 192.0.2.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 602
Total sessions: 1
```

show security flow session application telnet summary

```
root> show security flow session application telnet summary
Flow Sessions on FPC4 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

```
Flow Sessions on FPC5 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC5 PIC1:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 8.5.

Filter and view options added in Junos OS Release 10.2.

RELATED DOCUMENTATION

clear security flow session application | 3893

show security flow session application-firewall

IN THIS SECTION

- Syntax | 24191
- Description | 24191
- Options | 24191
- Required Privilege Level | 24192
- Output Fields | 24192

Sample Output | 24194
Release Information | 24199

Syntax

```
show security flow session application-firewall
< dynamic-application (dyn-app-name | junos:UNKNOWN) >
< dynamic-application-group (dyn-app-group | junos:UNASSIGNED) >
< application-firewall-rule-set rule-set-name >
< rule rule-name >
< brief | extensive | summary >
```

Description

Display all sessions where application firewall is enabled.

Include options to filter the output and display only those enabled sessions with the specified features.

Options

- `dynamic-application (dyn-app-name | junos:UNKNOWN)`–Display only those enabled sessions with the specified dynamic application. Enter `junos:UNKNOWN` to display all enabled sessions where no dynamic application can be determined.
- `dynamic-application-group (dyn-app-group | junos:UNASSIGNED)`– Display only those enabled session with the specified dynamic application group. Enter `junos:UNASSIGNED` to display all enabled sessions where no dynamic application group can be determined.
- `application-firewall-rule-set rule-set-name`–Display only those enabled sessions that match the specified rule set.
- `rule rule-name`–Display only those enabled sessions that match the specified rule.
- `brief | extensive | summary`–Specify the level of detail for the display.

  The output fields for the `brief` and `summary` options are the same as those of the `show security flow session` command. Only the `extensive` display is different and is shown in the following output table and examples.

Required Privilege Level

view

Output Fields

Table 1177 on page 24192 lists the output fields for the `show security flow session application-firewall extensive` command. Output fields are listed in the approximate order in which they appear in the extensive display.

Table 1177: show security flow session application-firewall extensive Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. Use this ID to display more information about a session.
`Status`	Session status.
`State`	Current state of the session: Active, Pending, Closed, Unknown.
`Flag`	Internal flag depicting the state of the session. It is used for debugging purposes.
`Policy name`	The name of the policy that permitted the traffic.
`Source NAT pool`	The name of the source pool where NAT is used.
`Dynamic application`	Name of the dynamic application of the session. If the dynamic application has yet to be determined, the output indicates Pending. If the dynamic application cannot be determined, the output indicates junos:UNKNOWN.
`Dynamic application group`	Name of the dynamic application group of the session. If the dynamic application cannot be determined, the output indicates junos:UNASSIGNED.

Table 1177: show security flow session application-firewall extensive Output Fields *(Continued)*

Field Name	Field Description
`Dynamic nested application`	Name of the dynamic nested application of the session if one exists. If the dynamic nested application is yet to be determined, the output indicates Pending. If the dynamic nested application cannot be determined, the output indicates junos:UNKNOWN.
`Application firewall rule-set`	Name of the rule set that the session matched.
`Rule`	Name of the rule that the session matched. If the match has not yet been made, the output indicates Pending. If the rule has been deleted since the match was made, the output indicates the rule is invalid.
`Maximum timeout`	Maximum amount of idle time allowed for the session.
`Current timeout`	Number of seconds that the current session has been idle.
`Session State`	Session state.
`Start time`	Time when the session was created. Start time is indicated as an offset from the system start time.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets, and bytes).
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Total sessions`	Total number of sessions per PIC that fit the display criteria.

Sample Output

show security flow session application-firewall extensive

The displayed information is similar to the `show security flow session` output but includes dynamic application and application firewall details for the session.

```
user@host> show security flow session application-firewall extensive
Flow Sessions on FPC9 PIC0:

     Session ID: 3729, Status: Normal, State: Active
      Policy name: self-traffic-policy/1
      Source NAT pool: Null
      Dynamic application: junos:HTTP, Dynamic nested application: junos:FACEBOOK-ACCESS
      Application firewall rule-set: rule-set1, Rule: rule2
      Maximum timeout: 300, Current timeout: 276
      Session State: Valid
      Start time: 18292, Duration: 603536
         In: 192.0.2.1/1 --> 203.0.113.1/1;pim,
          Interface: reth1.0,
          Session token: 0x1c0, Flag: 0x0x21
          Route: 0x0, Gateway: 192.0.2.4, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 21043, Bytes: 1136322
         Out: 203.0.113.1/1 --> 192.0.2.1/1;pim,
          Interface: .local..0,
          Session token: 0x80, Flag: 0x0x30
          Route: 0xfffd0000, Gateway: 203.0.113.13, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 0, Bytes: 0

     Total sessions: 1
```

show security flow session application-firewall dynamic-application junos:FTP extensive

Entering a specific dynamic application in the command line filters the output and displays only those sessions with the specified application.

```
user@host> show security flow session application-firewall dynamic-application junos:FTP
extensive
Flow Sessions on FPC3 PIC0:

     Session ID: 180013338, Policy name: policy1/4, Timeout: 1776, Valid
      Dynamic application: junos:FTP
      Application firewall rule-set: rule-set1, Rule: rule1
      Maximum timeout: 300, Current timeout: 276
      Session State: Valid
      Start time: 18292, Duration: 603536
         In: 192.0.2.4/1 --> 203.0.113.13/1;pim,
          Interface: reth1.0,
          Session token: 0x1c0, Flag: 0x0x21
          Route: 0x0, Gateway: 192.0.2.4, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 21043, Bytes: 1136322
         Out: 203.0.113.13/1 --> 192.0.2.4/1;pim,
          Interface: .local..0,
          Session token: 0x80, Flag: 0x0x30
          Route: 0xfffd0000, Gateway: 203.0.113.13, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 0, Bytes: 0

      Total sessions: 1
```

show security flow session application-firewall dynamic-application junos:UNKNOWN extensive

Using the keyword `junos:UNKNOWN` displays those enabled sessions where the dynamic application cannot be determined.

```
user@host> show security flow session application-firewall dynamic-application junos:UNKNOWN
extensive
```

```
Flow Sessions on FPC9 PIC0:

     Session ID: 180013338, Policy name: policy1/4, Timeout: 1776, Valid
      Dynamic application: junos:UNKNOWN
      Application firewall rule-set: rule-set1, Rule:rule1
      Maximum timeout: 300, Current timeout: 276
      Session State: Valid
      Start time: 18292, Duration: 603536
         In: 192.0.2.4/1 --> 203.0.113.13/1;pim,
          Interface: reth1.0,
          Session token: 0x1c0, Flag: 0x0x21
          Route: 0x0, Gateway: 192.0.2.4, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 21043, Bytes: 1136322
         Out: 203.0.113.13/1 --> 192.0.2.4/1;pim,
          Interface: .local..0,
          Session token: 0x80, Flag: 0x0x30
          Route: 0xfffd0000, Gateway: 203.0.113.13, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 0, Bytes: 0

      Session ID: 180013339, Policy name: policy1/4, Timeout: 1776, Valid
      Dynamic application: junos:HTTP, Dynamic nested application: junos:UNKNOWN
      Application firewall rule-set: rule-set1, Rule:rule1
      Maximum timeout: 300, Current timeout: 276
      Session State: Valid
      Start time: 18292, Duration: 603536
         In: 192.0.2.4/1 --> 203.0.113.13/1;pim,
          Interface: reth1.0,
          Session token: 0x1c0, Flag: 0x0x21
          Route: 0x0, Gateway: 192.0.2.4, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 21043, Bytes: 1136322
         Out: 203.0.113.13/1 --> 192.0.2.4/1;pim,
          Interface: .local..0,
          Session token: 0x80, Flag: 0x0x30
          Route: 0xfffd0000, Gateway: 203.0.113.13, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 0, Bytes: 0
```

```
      Total sessions: 2
```

show security flow session application-firewall dynamic-application-group junos:WEB extensive

Entering a specific dynamic application group in the command line filters the output and displays only those sessions with the specified application group.

```
user@host> show security flow session application-firewall dynamic-application-group junos:WEB
extensive
Flow Sessions on FPC9 PIC0:

     Session ID: 180013338, Policy name: policy1/4, Timeout: 1776, Valid
      Dynamic application: junos:HOTMAIL
      Application firewall rule-set: rule-set1, Rule: rule1
      Maximum timeout: 300, Current timeout: 276
      Session State: Valid
      Start time: 18292, Duration: 603536
         In: 192.0.2.4/1 --> 203.0.113.13/1;pim,
          Interface: reth1.0,
          Session token: 0x1c0, Flag: 0x0x21
          Route: 0x0, Gateway: 192.0.2.4, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 21043, Bytes: 1136322
         Out: 203.0.113.13/1 --> 192.0.2.4/1;pim,
          Interface: .local..0,
          Session token: 0x80, Flag: 0x0x30
          Route: 0xfffd0000, Gateway: 203.0.113.13, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 0, Bytes: 0

      Total sessions: 1
```

show security flow session application-firewall application-firewall-rule-set rule-set1 extensive

Specifying a rule set name reduces the display to only those sessions matching the specified rule set.

```
user@host> show security flow session application-firewall application-firewall-rule-set rule-
set1 extensive
Flow Sessions on FPC9 PIC0:

     Session ID: 180013338, Policy name: policy1/4, Timeout: 1776, Valid
      Dynamic application: junos:FTP
      Application firewall rule-set: rule-set1, Rule: rule1
      Maximum timeout: 300, Current timeout: 276
      Session State: Valid
      Start time: 18292, Duration: 603536
         In: 192.0.2.4/1 --> 203.0.113.13/1;pim,
          Interface: reth1.0,
          Session token: 0x1c0, Flag: 0x0x21
          Route: 0x0, Gateway: 192.0.2.4, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 21043, Bytes: 1136322
         Out: 203.0.113.13/1 --> 192.0.2.4/1;pim,
          Interface: .local..0,
          Session token: 0x80, Flag: 0x0x30
          Route: 0xfffd0000, Gateway: 203.0.113.13, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 0, Bytes: 0

     Session ID: 180013339, Policy name: policy1/4, Timeout: 1776, Valid
      Dynamic application: junos:HTTP, Dynamic nested application: junos:FACEBOOK-ACCESS
      Application firewall rule-set: rule-set1, Rule: rule2
      Maximum timeout: 300, Current timeout: 276
      Session State: Valid
      Start time: 18292, Duration: 603536
         In: 192.0.2.4/1 --> 203.0.113.13/1;pim,
          Interface: reth1.0,
          Session token: 0x1c0, Flag: 0x0x21
          Route: 0x0, Gateway: 192.0.2.4, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
```

```
            Pkts: 21043, Bytes: 1136322
           Out: 203.0.113.13/1 --> 192.0.2.4/1;pim,
            Interface: .local..0,
            Session token: 0x80, Flag: 0x0x30
            Route: 0xfffd0000, Gateway: 203.0.113.13, Tunnel: 0
            Port sequence: 0, FIN sequence: 0,
            FIN state: 0,
            Pkts: 0, Bytes: 0



        Total sessions: 2
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

> *Example: Configuring an Application Group for Application Firewall*
>
> *show security flow session*

show security flow session brief node

IN THIS SECTION

- Syntax | 24200
- Description | 24200
- Options | 24200
- Required Privilege Level | 24200
- Output Fields | 24200
- Sample Output | 24201
- Release Information | 24204

Syntax

```
show security flow session brief node (node-id | all | local | primary)
```

Description

Display information about all currently active security sessions on the device for the specified node options in brief mode.

Options

`node`—(Optional) For chassis cluster configurations, display session information on a specific node.

- *`node-id`* —Identification number of the node. It can be 0 or 1.
- `all`—Display information about all nodes.
- `local`—Display information about the local node.
- `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1178 on page 24200 lists the output fields for the `show security flow session brief node` command. Output fields are listed in the approximate order in which they appear.

Table 1178: show security flow session brief node Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. Use this ID to get more information about the session.
`Policy name`	Policy that permitted the traffic.

Table 1178: show security flow session brief node Output Fields *(Continued)*

Field Name	Field Description
State	Session state.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.
Total sessions	Total number of sessions.

Sample Output

show security flow session brief node 0

```
root@host> show security flow session brief node 0
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000001, Policy name: default-policy-00/2, State: Active, Timeout: 1696, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 14, Bytes: 626, CP Session
ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 13, Bytes: 744, CP Session
ID: 10000001
```

```
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0



Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session brief node 1

```
root@host> show security flow session brief node 1
node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000001, Policy name: default-policy-00/2, State: Active, Timeout: 1696, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 14, Bytes: 626, CP Session
ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 13, Bytes: 744, CP Session
ID: 10000001
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session brief node all

```
root@host> show security flow session brief node all
node0:
--------------------------------------------------------------------------

Session ID: 10000001, Policy name: default-policy-00/2, State: Active, Timeout: 1696, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 14, Bytes: 626, CP Session
```

```
ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 13, Bytes: 744, CP Session
ID: 10000001
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000001, Policy name: default-policy-00/2, State: Active, Timeout: 1696, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 14, Bytes: 626, CP Session
ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 13, Bytes: 744, CP Session
ID: 10000001
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session brief node local

```
root@host> show security flow session brief node local
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000001, Policy name: default-policy-00/2, State: Active, Timeout: 1696, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 14, Bytes: 626, CP Session
```

```
ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 13, Bytes: 744, CP Session
ID: 10000001

Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session brief node primary

```
root@host> show security flow session brief node primary
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000001, Policy name: default-policy-00/2, State: Active, Timeout: 1696, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 14, Bytes: 626, CP Session
ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 13, Bytes: 744, CP Session
ID: 10000001

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 8.5; node options added in Junos OS Release 9.0. Filter options added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security flow session | 24162

show security flow session destination-port

IN THIS SECTION

- Syntax | 24205
- Description | 24205
- Options | 24205
- Required Privilege Level | 24206
- Output Fields | 24206
- Sample Output | 24208
- Release Information | 24210

Syntax

```
show security flow session destination-port
            destination-port-number [brief | extensive | summary]
```

Description

Display information about each session that uses the specified destination port.

Options

- *destination-port-number* —Number of the destination port for which to display sessions information.
- **Range:** 1 through 65,535

brief | extensive | summary	Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1179 on page 24206 lists the output fields for the `show security flow session destination-port` command. Output fields are listed in the approximate order in which they appear.

Table 1179: show security flow session destination-port Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Total sessions`	Total number of sessions.
`Status`	Session status.
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.
`Policy name`	Name and ID of the policy that the first packet of the session matched.
`Source NAT pool`	The name of the source pool where NAT is used.

Table 1179: show security flow session destination-port Output Fields *(Continued)*

Field Name	Field Description
`Application`	Name of the application.
`Maximum timeout`	Maximum session timeout.
`Current timeout`	Remaining time for the session unless traffic exists in the session.
`Session State`	Session state.
`Start time`	Time when the session was created, offset from the system start time.
`Unicast-sessions`	Number of unicast sessions.
`Multicast-sessions`	Number of multicast sessions.
`Failed-sessions`	Number of failed sessions.
`Sessions-in-use`	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
`Maximum-sessions`	Number of maximum sessions.

Sample Output

show security flow session destination-port 23

```
root> show security flow session destination-port 23
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000098, Policy name: default-policy-00/2, Timeout: 1778, Valid
  In: 198.51.100.10/15190 --> 198.51.100.2/23;tcp, If: ge-7/1/0.0, Pkts: 109, Bytes: 5874, CP
Session ID: 430000093
  Out: 198.51.100.2/23 --> 198.51.100.10/15190;tcp, If: ge-7/1/1.0, Pkts: 64, Bytes: 4015, CP
Session ID: 430000093
Total sessions: 1
```

show security flow session destination-port 23 brief

```
root> show security flow session destination-port 23 brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000098, Policy name: default-policy-00/2, Timeout: 1778, Valid
  In: 198.51.100.10/15190 --> 198.51.100.2/23;tcp, If: ge-7/1/0.0, Pkts: 109, Bytes: 5874, CP
Session ID: 430000093
  Out: 198.51.100.2/23 --> 198.51.100.10/15190;tcp, If: ge-7/1/1.0, Pkts: 64, Bytes: 4015, CP
Session ID: 430000093
Total sessions: 1
```

show security flow session destination-port 23 extensive

```
root> show security flow session destination-port 23 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000098, Status: Normal
Flags: 0x40/0x0/0x2008003
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1630
Session State: Valid
Start time: 65490, Duration: 207
   In: 198.51.100.10/15190 --> 198.51.100.2/23;tcp
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0001021
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 109, Bytes: 5874
    CP Session ID: 430000093
   Out: 198.51.100.2/23 --> 2198.51.100.10/15190;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xc0001020
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 64, Bytes: 4015
    CP Session ID: 430000093
Total sessions: 1
```

show security flow session destination-port 23 summary

```
root> show security flow session destination-port 23 summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 8.5; Filter and view options added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session destination-port | **3901**

show security flow session destination-prefix

IN THIS SECTION

- Syntax | 24211
- Description | 24211
- Options | 24211
- Required Privilege Level | 24211
- Output Fields | 24211
- Sample Output | 24213
- Release Information | 24217

Syntax

```
show security flow session destination-prefix
          destination-IP-prefix  [brief | extensive | summary]
```

Description

Display information about each session that matches the specified IP destination prefix.

Options

- *destination-IP-prefix*—Destination IP prefix or address for which to display session information.
- brief | extensive | summary–Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1180 on page 24212 lists the output fields for the `show security flow session destination-prefix` command. Output fields are listed in the approximate order in which they appear.

Table 1180: show security flow session destination-prefix Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`HA State`	High Availability state of the session (Active or Warm).
`Timeout`	Idle timeout after which the session expires.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Total sessions`	Total number of sessions.
`Status`	Session status.
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.
`Policy name`	Name and ID of the policy that the first packet of the session matched.
`Source NAT pool`	The name of the source pool where NAT is used.
`Application`	Name of the application.
`Maximum timeout`	Maximum session timeout.

Table 1180: show security flow session destination-prefix Output Fields *(Continued)*

Field Name	Field Description
`Current timeout`	Remaining time for the session unless traffic exists in the session.
`Session State`	Session state.
`Start time`	Time when the session was created, offset from the system start time.
`Unicast-sessions`	Number of unicast sessions.
`Multicast-sessions`	Number of multicast sessions.
`Failed-sessions`	Number of failed sessions.
`Sessions-in-use`	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
`Maximum-sessions`	Number of maximum sessions.

Sample Output

show security flow session destination-prefix 198.51.100.1/24

```
root> show security flow session destination-prefix 198.51.100.1/24
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
```

```
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000098, Policy name: default-policy-00/2, Timeout: 1450, Valid
  In: 192.0.2.10/15190 --> 198.51.100.1/23;tcp, If: ge-7/1/0.0, Pkts: 109, Bytes: 5874, CP
Session ID: 430000093
  Out: 198.51.100.1/23 --> 192.0.2.10/15190;tcp, If: ge-7/1/1.0, Pkts: 64, Bytes: 4015, CP
Session ID: 430000093
Total sessions: 1
```

show security flow session destination-prefix 198.51.100.1/24 brief

```
root> show security flow session destination-prefix 198.51.100.1/24 brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000098, Policy name: default-policy-00/2, Timeout: 1450, Valid
  In: 192.0.2.10/15190 --> 198.51.100.1/23;tcp, If: ge-7/1/0.0, Pkts: 109, Bytes: 5874, CP
Session ID: 430000093
  Out: 198.51.100.1/23 --> 192.0.2.10/15190;tcp, If: ge-7/1/1.0, Pkts: 64, Bytes: 4015, CP
Session ID: 430000093
Total sessions: 1
```

show security flow session destination-prefix 198.51.100.1/24 extensive

```
root> show security flow session destination-prefix 198.51.100.1/24 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
```

```
Session ID: 430000098, Status: Normal
Flags: 0x40/0x0/0x2008003
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1172
Session State: Valid
Start time: 65490, Duration: 666
   In: 192.0.2.10/15190 --> 198.51.100.1/23;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0001021
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 109, Bytes: 5874
    CP Session ID: 430000093
   Out: 198.51.100.1/23 --> 200.0.0.10/15190;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xc0001020
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 64, Bytes: 4015
    CP Session ID: 430000093
Total sessions: 1
```

show security flow session destination-prefix 198.51.100.1/24 summary

```
root> show security flow session destination-prefix 198.51.100.1/24 summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:
```

```
Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

show security flow session destination-prefix 2001:db8::/32

```
user@host> show security flow session destination-prefix 2001:db8::/32
Session ID: 50000004, Policy name: self-traffic-policy/1, Timeout: 2
In: 10::11/42756 --> 2001:db8::1 ;icmp, If: .local..0
Out: 2001:db8::1 --> 10::11/42756;icmp, If: ge-0/3/0.0

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

For Multinode High Availability Scenario:

show security flow session destination-prefix

Backup Node

```
user@host> show security flow session destination-prefix 10.11.1.1
Session ID: 16445, Policy name: default-permit/5, HA State: Warm, Timeout: 13726, Session State:
Valid
  In: 22.1.1.2/36816 --> 11.1.1.1/22;tcp, Conn Tag: 0x0, If: ge-0/0/1.0, Pkts: 0, Bytes: 0,
  Out: 11.1.1.1/22 --> 22.1.1.2/36816;tcp, Conn Tag: 0x0, If: ge-0/0/2.0, Pkts: 0, Bytes: 0,
Total sessions: 1
```

Active Node

```
user@host> show security flow session destination-prefix 10.11.1.1
Session ID: 89282, Policy name: default-permit/5, HA State: Active, Timeout: 1488, Session
State: Valid
  In: 22.1.1.2/36816 --> 11.1.1.1/22;tcp, Conn Tag: 0x0, If: ge-0/0/1.0, Pkts: 39, Bytes: 4681,
  Out: 11.1.1.1/22 --> 22.1.1.2/36816;tcp, Conn Tag: 0x0, If: ge-0/0/2.0, Pkts: 24, Bytes: 3977,
Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 8.5. Support.

Support for IPv6 addresses added in Junos OS Release 10.2.

Filter and view options added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session destination-port | 3901

show security flow session extensive node

IN THIS SECTION

- Syntax | 24218
- Description | 24218
- Options | 24218
- Required Privilege Level | 24218
- Output Fields | 24218
- Sample Output | 24220
- Release Information | 24227

Syntax

```
show security flow session extensive node (node-id | all | local | primary)
```

Description

Display information about all currently active security sessions on the device for the specified node options in extensive mode.

Options

`node`—(Optional) For chassis cluster configurations, display session information on a specific node.

- *`node-id`* —Identification number of the node. It can be 0 or 1.
- `all`—Display information about all nodes.
- `local`—Display information about the local node.
- `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1181 on page 24218 lists the output fields for the `show security flow session extensive node` command. Output fields are listed in the approximate order in which they appear.

Table 1181: show security flow session extensive node Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Status`	Session status.

Table 1181: show security flow session extensive node Output Fields *(Continued)*

Field Name	Field Description
State	Session state.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Policy that permitted the traffic.
Source NAT pool	The name of the source pool where NAT is used.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Start time	Time when the session was created, offset from the system start time.
Duration	Length of time for which the session is active.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, flag, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, flag, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.

Sample Output

show security flow session extensive node 0

```
root@host> show security flow session extensive node 0
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Active
Flags: 0x8000042/0x8000000/0x110103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1778
Session State: Valid
Start time: 6466, Duration: 28
   In: 10.0.2.1/52080 --> 203.0.113.1/24;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x40002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 414
    CP Session ID: 10000004
   Out: 203.0.113.1/24 --> 10.0.2.1/52080;tcp,
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x40002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 420
    CP Session ID: 10000004
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0
```

```
Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session extensive node 1

```
root@host> show security flow session extensive node 1
node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Backup
Flags: 0x10000042/0x0/0x10103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 14324
Session State: Valid
Start time: 6248, Duration: 90
   In: 110.0.2.1/52080 --> 203.0.113.1/24;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x60002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10000003
   Out: 203.0.113.1/24 --> 10.0.2.1/52080;tcp,
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x60002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10000003
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0
```

```
Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session extensive node all

```
root@host> show security flow session extensive node all
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Active
Flags: 0x8000042/0x8000000/0x110103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1692
Session State: Valid
Start time: 6466, Duration: 113
   In: 10.0.2.1/52080 --> 203.0.113.1/21;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x40002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 414
    CP Session ID: 10000004
   Out: 203.0.113.1/21 --> 10.0.2.1/52080;tcp,
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x40002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 420
    CP Session ID: 10000004
Total sessions: 1

Flow Sessions on FPC0 PIC2:
```

```
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Backup
Flags: 0x10000042/0x0/0x10103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 14298
Session State: Valid
Start time: 6248, Duration: 115
   In: 10.0.2.1/52080 --> 203.0.113.1/21;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x60002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10000003
  Out: 203.0.113.1/21 --> 10.0.2.1/52080;tcp,
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x60002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10000003
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0
```

```
Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session extensive node local

```
root@host> show security flow session extensive node local
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Active
Flags: 0x8000042/0x8000000/0x110103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1584
Session State: Valid
Start time: 6466, Duration: 221
   In: 100.0.0.2/52080 --> 120.0.0.2/21;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x40002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 414
    CP Session ID: 10000004
   Out: 120.0.0.2/21 --> 100.0.0.2/52080;tcp,
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x40002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 420
    CP Session ID: 10000004
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0
```

```
Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session extensive node primary

```
root@host> show security flow session extensive node primary
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Active
Flags: 0x8000042/0x8000000/0x110103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1554
Session State: Valid
Start time: 6466, Duration: 252
   In: 100.0.0.2/52080 --> 120.0.0.2/21;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x40002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 414
    CP Session ID: 10000004
   Out: 120.0.0.2/21 --> 100.0.0.2/52080;tcp,
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x40002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 420
    CP Session ID: 10000004
Total sessions: 1

Flow Sessions on FPC0 PIC2:
```

```
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session extensive node 0 (PowerMode)

```
root@host> show security flow session extensive node 0
node0:
--------------------------------------------------------------------------

Session ID: 13607, Status: Normal, State: Active
Flags: 0x88000040/0x8000000/0x2/0x908003
Policy name: p1/4
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Url-category:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 150, Current timeout: 2
Session State: Valid
Start time: 443962, Duration: 3
   In: 20.1.1.2/10000 --> 30.1.1.2/2001;tcp,
  Conn Tag: 0x0, Interface: reth0.0,
    Session token: 0x8, Flag: 0x201021,
    Power-Mode Active: True
    Route: 0x33a3c2, Gateway: 20.1.1.2, Tunnel ID: 0, Tunnel type: None
    Port sequence: 0, FIN sequence: 0,
    FIN state: 2,
    Pkts: 24, Bytes: 1896
   Out: 30.1.1.2/2001 --> 20.1.1.2/10000;tcp,
  Conn Tag: 0x0, Interface: reth1.0,
    Session token: 0x7, Flag: 0x201020,
    Power-Mode Active: True
    Route: 0x326bc2, Gateway: 30.1.1.2, Tunnel ID: 0, Tunnel type: None
    Port sequence: 0, FIN sequence: 0,
    FIN state: 2,
    Pkts: 13, Bytes: 1324
Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 8.5; node options added in Junos OS Release 9.0. Filter options added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security flow session

show security flow session family

IN THIS SECTION

- Syntax | 24227
- Description | 24227
- Options | 24228
- Required Privilege Level | 24228
- Output Fields | 24228
- Sample Output | 24230
- Release Information | 24234

Syntax

```
show security flow session family (inet | inet6)
[brief | extensive | summary]
```

Description

Display filtered summary of information about existing sessions, including types of sessions, active and failed sessions, and the maximum allowed number of sessions.

Options

- `inet`—Display details summary of IPv4 sessions.
- `inet6`—Display details summary of IPv6 sessions.
- brief | extensive | summary–Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1182 on page 24228 lists the output fields for the `show security flow session family` command. Output fields are listed in the approximate order in which they appear.

Table 1182: show security flow session family Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. Use this ID to get more information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Total sessions`	Total number of sessions.
`Status`	Session status.

Table 1182: show security flow session family Output Fields *(Continued)*

Field Name	Field Description
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.
`Policy name`	Name and ID of the policy that the first packet of the session matched.
`Source NAT pool`	The name of the source pool where NAT is used.
`Application`	Name of the application.
`Maximum timeout`	Maximum session timeout.
`Current timeout`	Remaining time for the session unless traffic exists in the session.
`Session State`	Session state.
`Start time`	Time when the session was created, offset from the system start time.
`Unicast-sessions`	Number of unicast sessions.
`Multicast-sessions`	Number of multicast sessions.
`Failed-sessions`	Number of failed sessions.
`Sessions-in-use`	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states

Table 1182: show security flow session family Output Fields *(Continued)*

Field Name	Field Description
Maximum-sessions	Number of maximum sessions.

Sample Output

show security flow session family inet

```
root> show security flow session family inet
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000107, Policy name: default-policy-00/2, Timeout: 4, Valid
  In: 203.0.113.0/3 --> 203.0.113.5/24;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
420000202
  Out: 203.0.113.4/24 --> 203.0.113.6/24;icmp, If: .local..0, Pkts: 1, Bytes: 84, CP Session ID:
420000202
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Session ID: 430000115, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 203.0.113.0/4 --> 203.0.113.5/24;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
430000110
  Out: 203.0.113.5/24 --> 203.0.113.6/24;icmp, If: .local..0, Pkts: 1, Bytes: 84, CP Session ID:
430000110

Session ID: 430000117, Policy name: default-policy-00/2, Timeout: 4, Valid
  In: 203.0.113.0/4 --> 203.0.113.5/24;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
430000111
  Out: 203.0.113.5/24 --> 203.0.113.6/24;icmp, If: .local..0, Pkts: 1, Bytes: 84, CP Session ID:
430000111
Total sessions: 2
```

show security flow session family inet brief

```
root> show security flow session family inet brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000115, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 203.0.113.0/3 --> 203.0.113.5/24;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
420000206
  Out: 203.0.113.4/24 --> 203.0.113.6/24;icmp, If: .local..0, Pkts: 1, Bytes: 84, CP Session ID:
420000206

Session ID: 420000117, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 203.0.113.0/4 --> 203.0.113.5/24;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
420000207
  Out: 203.0.113.5/24 --> 203.0.113.6/24;icmp, If: .local..0, Pkts: 1, Bytes: 84, CP Session ID:
420000207
Total sessions: 2

Flow Sessions on FPC10 PIC3:

Session ID: 430000119, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 203.0.113.0/4 --> 203.0.113.5/24;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
430000112
  Out: 203.0.113.5/24 --> 203.0.113.6/24;icmp, If: .local..0, Pkts: 1, Bytes: 84, CP Session ID:
430000112
Total sessions: 1
```

show security flow session family inet extensive

```
root> show security flow session family inet extensive
Flow Sessions on FPC10 PIC1:

Session ID: 410000111, Status: Normal
Flags: 0x80400040/0x0/0x2800023
Policy name: default-policy-00/2
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
```

```
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 4, Current timeout: 4
Session State: Valid
Start time: 76455, Duration: 0
   In: 203.0.113.0/24 --> 203.0.113.1/24;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
    Route: 0xa0010, Gateway: 203.0.113.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 410000242
   Out: 203.0.113.1/24 --> 203.0.113.10/4;icmp,
    Interface: .local..0,
    Session token: 0x2, Flag: 0x40000030
    Route: 0xfffb0006, Gateway: 203.0.113.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 410000242
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Session ID: 420000123, Status: Normal
Flags: 0x80400040/0x0/0x2800023
Policy name: default-policy-00/2
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 4, Current timeout: 2
Session State: Valid
Start time: 76454, Duration: 2
   In: 203.0.113.10/24 --> 203.0.113.11/24;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
    Route: 0xa0010, Gateway: 20010, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 420000210
```

```
   Out: 203.0.113.11/24 --> 203.0.113.12/24;icmp,
    Interface: .local..0,
    Session token: 0x2, Flag: 0x40000030
    Route: 0xfffb0006, Gateway: 203.0.113.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 420000210
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Session ID: 430000131, Status: Normal
Flags: 0x80400040/0x0/0x2800023
Policy name: default-policy-00/2
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 4, Current timeout: 4
Session State: Valid
Start time: 76421, Duration: 1
   In: 203.0.113.10/24 --> 203.0.113.11/24;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
    Route: 0xa0010, Gateway: 203.0.113.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 430000118
   Out: 203.0.113.12/24 --> 203.0.113.13/24;icmp,
    Interface: .local..0,
    Session token: 0x2, Flag: 0x40000030
    Route: 0xfffb0006, Gateway: 203.0.113.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 430000118
Total sessions: 1
```

show security flow session family inet summary

```
root> show security flow session family inet summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 2
Sessions in other states: 0
Total sessions: 4

Flow Sessions on FPC10 PIC2:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 2
Sessions in other states: 0
Total sessions: 4

Flow Sessions on FPC10 PIC3:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 2
Sessions in other states: 0
Total sessions: 4
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session family | 3905

show security flow session idp family

IN THIS SECTION

- Syntax | 24235
- Description | 24235
- Options | 24235
- Required Privilege Level | 24235
- Output Fields | 24236
- Sample Output | 24236
- Release Information | 24237

Syntax

```
show security flow session idp family (inet | inet6)
```

Description

Display filtered summary of information about existing sessions, including types of sessions, active and failed sessions, and the maximum allowed number of sessions.

Options

inet Display details summary of IPv4 sessions.

inet6 Display details summary of IPv6 sessions.

Required Privilege Level

view

Output Fields

Table 1183 on page 24236 lists the output fields for the `show security flow session summary family` command. Output fields are listed in the approximate order in which they appear.

Table 1183: show security flow session summary Output Fields

Field Name	Field Description
`Valid sessions`	Count of valid sessions.
`Pending sessions`	Count of pending sessions.
`Invalidated sessions`	Count of sessions the security device has determined to be invalid.
`Sessions in other states`	Count of sessions not in valid, pending, or invalidated state.
`Total sessions`	Total of the above counts.

Sample Output

show security flow session summary family inet

```
user@host> show security flow session summary family inet
Flow Sessions on FPC4 PIC0:
Valid sessions: 3
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 3

Flow Sessions on FPC5 PIC0:
Valid sessions: 4
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 4
```

show security flow session summary family inet6

```
user@host> show security flow session summary family inet6

Flow Sessions on FPC1 PIC1:
Valid sessions: 20
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 20
```

Release Information

Command introduced in Junos OS Release 10.2.

Support for family inet6 added in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

Intrusion Detection and Prevention Overview

show security flow session idp summary

IN THIS SECTION

- Syntax | 24238
- Description | 24238
- Options | 24238
- Required Privilege Level | 24238
- Output Fields | 24238
- Sample Output | 24239
- Release Information | 24239

Syntax

```
show security flow session idp summary
```

Description

Display summary output.

Options

- application—Application name
- destination-port—Destination port
- destination-prefix—Destination IP prefix or address
- family—Display session by family.
- interface–Name of incoming or outgoing interface
- protocol—IP protocol number
- source-port—Source port
- source-prefix—Source IP prefix

Required Privilege Level

view

Output Fields

Table 1184 on page 24238 lists the output fields for the `show security flow session idp summary` command. Output fields are listed in the approximate order in which they appear.

Table 1184: show security flow session idp summary Output Fields

Field Name	Field Description
Valid session	Number of valid sessions.

Table 1184: show security flow session idp summary Output Fields *(Continued)*

Field Name	Field Description
Pending sessions	Number of pending sessions.
Invalidated sessions	Number of invalid sessions.
Sessions in other states	Number of sessions in other states.
Total sessions	Total number of sessions.

Sample Output

show security flow session idp summary

```
root@ show security flow session idp summary
Flow Sessions on FPC4 PIC0:

Valid sessions: 3
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 3

Flow Sessions on FPC5 PIC0:

Valid sessions: 4
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 4
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security flow session

show security flow session interface

IN THIS SECTION

- Syntax | 24240
- Description | 24240
- Options | 24240
- Required Privilege Level | 24241
- Output Fields | 24241
- Sample Output | 24243
- Release Information | 24245

Syntax

```
show security flow session interface
interface-name [brief | extensive | summary]
```

Description

Display information about each session that uses the specified interface. The interface name can be a session's incoming or outgoing interface.

Options

- *interface-name*—Name of the interface on the device for which to display sessions information.
- brief | extensive | summary–Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1185 on page 24241 lists the output fields for the `show security flow session interface` command. Output fields are listed in the approximate order in which they appear.

Table 1185: show security flow session interface Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Total sessions`	Total number of sessions.
`Status`	Session status.
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.
`Policy name`	Name and ID of the policy that the first packet of the session matched.
`Source NAT pool`	The name of the source pool where NAT is used.

Table 1185: show security flow session interface Output Fields *(Continued)*

Field Name	Field Description
`Application`	Name of the application.
`Maximum timeout`	Maximum session timeout.
`Current timeout`	Remaining time for the session unless traffic exists in the session.
`Session State`	Session state.
`Start time`	Time when the session was created, offset from the system start time.
`Unicast-sessions`	Number of unicast sessions.
`Multicast-sessions`	Number of multicast sessions.
`Failed-sessions`	Number of failed sessions.
`Sessions-in-use`	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
`Maximum-sessions`	Number of maximum sessions.

Sample Output

show security flow session interface ge-0/0/2.0

```
root> show security flow session interface ge-7/1/1.0
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000146, Policy name: default-policy-00/2, Timeout: 58, Valid
  In: 200.0.0.10/9 --> 60.0.0.2/21562;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
420000247
  Out: 60.0.0.2/21562 --> 200.0.0.10/9;icmp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID:
420000247
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Session ID: 430000146, Policy name: default-policy-00/2, Timeout: 56, Valid
  In: 200.0.0.10/8 --> 60.0.0.2/21562;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
430000131
  Out: 60.0.0.2/21562 --> 200.0.0.10/8;icmp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID:
430000131
Total sessions: 1
```

show security flow session interface ge-0/0/2.0 brief

```
root> show security flow session interface ge-7/1/1.0 brief
Flow Sessions on FPC10 PIC1:

Session ID: 410000137, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 200.0.0.10/5 --> 60.0.0.2/23354;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
410000269
  Out: 60.0.0.2/23354 --> 200.0.0.10/5;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84, CP Session ID:
410000269
Total sessions: 1

Flow Sessions on FPC10 PIC2:
```

```
Session ID: 420000151, Policy name: default-policy-00/2, Timeout: 54, Valid
  In: 200.0.0.10/1 --> 60.0.0.2/23354;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
420000252
  Out: 60.0.0.2/23354 --> 200.0.0.10/1;icmp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID:
420000252
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session interface ge-0/0/2.0 extensive

```
root> show security flow session interface ge-7/1/1.0 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000151, Status: Normal
Flags: 0x40/0x0/0x2000003
Policy name: default-policy-00/2
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 60, Current timeout: 48
Session State: Valid
Start time: 83328, Duration: 12
   In: 200.0.0.10/1 --> 60.0.0.2/23354;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 420000252
   Out: 60.0.0.2/23354 --> 200.0.0.10/1;icmp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xc0000020
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
```

```
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000252
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session interface ge-7/1/1.0 summary

```
root> show security flow session interface ge-7/1/1.0 summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2
```

Release Information

Command introduced in Junos OS Release 8.5; Filter and view options added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session interface | 3910

show security flow session nat

IN THIS SECTION

- Syntax | 24246
- Description | 24246
- Options | 24246
- Required Privilege Level | 24246
- Output Fields | 24247
- Sample Output | 24248
- Release Information | 24250

Syntax

```
show security flow session nat [brief | extensive | summary]
```

Description

Display sessions with network address translation.

Options

brief | extensive | summary	Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1186 on page 24247 lists the output fields for the `show security flow session nat` command. Output fields are listed in the approximate order in which they appear.

Table 1186: show security flow session nat Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.
`Resource information`	Information about the session particular to the resource manager, including the name of the ALG, the group ID. and the resource ID.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Total sessions`	Total number of sessions.
`Status`	Session status.
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.
`Policy name`	Name and ID of the policy that the first packet of the session matched.
`Source NAT pool`	The name of the source pool where NAT is used.

Table 1186: show security flow session nat Output Fields *(Continued)*

Field Name	Field Description
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Valid sessions	Number of valid sessions.
Pending sessions	Number of pending sessions.
Invalidated sessions	Number of invalidated sessions.

Sample Output

show security flow session nat brief

```
root> show security flow session nat brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000390, Policy name: default-policy-00/2, Timeout: 1778, Valid
  In: 200.0.0.10/41043 --> 60.0.0.2/21;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 414, CP Session ID:
420001090
  Out: 60.0.0.2/21 --> 60.0.0.1/19473;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 479, CP Session ID:
430000964
Total sessions: 1
```

```
Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session nat extensive

```
root> show security flow session nat extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000390, Status: Normal
Flags: 0x2/0x0/0x2010103
Policy name: default-policy-00/2
Source NAT pool: interface, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1770
Session State: Valid
Start time: 151971, Duration: 55
   In: 200.0.0.10/41043 --> 60.0.0.2/21;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0002621
    Route: 0x70010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 414
    CP Session ID: 420001090
   Out: 60.0.0.2/21 --> 60.0.0.1/19473;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xe0002620
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 479
    CP Session ID: 430000964
Total sessions: 1
```

```
Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session nat summary

```
root> show security flow session nat summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security flow session

show security flow session plugins

IN THIS SECTION

- Syntax | 24251
- Description | 24251
- Required Privilege Level | 24251
- Output Fields | 24251
- Sample Output | 24253
- Release Information | 24254

Syntax

```
show security flow session plugins
```

Description

Displays the detailed information about security flow session of plugins. The output displays the information in a list to make it easy for you to read and monitor the flow session.

Required Privilege Level

View

Output Fields

Table 1187 on page 24251 lists the output fields for the `show security flow session plugins` command.

Table 1187: show security flow session plugins Output Fields

Field Name	Field Description
Bandwidth	Client to server and server to client bytes added together and converted to MB or GB.

Table 1187: show security flow session plugins Output Fields *(Continued)*

Field Name	Field Description
Base Layer7 Application	Base layer7 application details. • True: SSL proxy decrypted. • False: SSL proxy not decrypted.
Decrypted	SSL proxy decryption status.
Decrypted-Mirrored	Decryption mirror status.
Duration	Duration of the session.
Final Layer7 Application	The last classified application.
Forwarding Type	Show the type of current flow session.
Forward Direction	Interface, IP address, protocol, tag, and gateway information of this direction.
From Zone	Source zone of the flow session.
Hardware Accelerated	Status of the Express Path flow session.
Layer4 Application	Layer4 applications details.
Plugin List	List of the plugins used in the flow session.
Policy	Policy details.
Reverse Direction	Interface, IP address, protocol, tag, and gateway information of the flowdirection.
Session ID	Number that identifies the session. Use this ID to get more information about the session.

Table 1187: show security flow session plugins Output Fields *(Continued)*

Field Name	Field Description
Session HA Status	Session HA status such as `Stand-alone`, `Forward`, `Warm`, `Active`, `Backup`, or`Unknown`.
Session State	Status of security flow session.
Start time	Time when the session was created, offset from the system start time.
Time Out	Idle time-out after which the session expires.
To Zone	Destination zone of the flow session.
Tunneled Session	Status of the tunnel session. • True: Session is a tunnel session. • False: Session is not a tunnel session.

Sample Output

show security flow session plugins

```
user@host> show security flow session plugins
```

```
Session ID              : 19
Forward Direction       : Interface: .local..0, 192.0.2.0/10002 --> 198.51.100.0/80;tcp, conn
tag: 0x0, gateway: 192.0.2.0
Reverse Direction       : Interface: ge-0/0/15.0, 198.51.100.0/80 --> 192.0.2.0/10002;tcp, conn
tag: 0x0, gateway: 203.0.113.0
Forwarding Type         : Route
Bandwidth               : 0.02 MB
From Zone               : junos-host
To Zone                 : internet
Layer4 Application      : junos-http
Base Layer7 Application : UNKNOWN
Final Layer7 Application: UNKNOWN
```

```
Hardware Accelerated    : False
Decrypted               : False
Decrypted-Mirrored      : False
Tunneled Session        : False
Session State           : Valid
Session HA Status       : Stand-alone
Policy                  : self-traffic-policy/1
Time Out                : 300 s
Time-To-Live            : 272 s
Start Time              : 83 s
Duration                : 3811 s
Plugin List:
junos-tcp-svr-emul: C2S Interested, S2C Interested
junos-rtcom-stream: C2S Interested, S2C Interested
junos-tcp-clt-emul: C2S Interested, S2C Interested
```

Release Information

Command introduced in Junos OS Release 21.1R1.

show security flow session policy-id

IN THIS SECTION

- Syntax | 24255
- Description | 24255
- Options | 24255
- Required Privilege Level | 24255
- Output Fields | 24255
- Sample Output | 24257
- Release Information | 24259

Syntax

```
show security flow session policy-id policy-id-number  [brief | extensive | summary ]
```

Description

Display information about each session by using policy id of the session.

Options

- *policy-id-number* —ID of the policy that the first packet of the session matches with.

  Range: 1through 4294967295

- brief | extensive | summary–Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1188 on page 24255 lists the output fields for the `show security flow session policy-id` command. Output fields are listed in the approximate order in which they appear.

Table 1188: show security flow session policy-id Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.

Table 1188: show security flow session policy-id Output Fields *(Continued)*

Field Name	Field Description
`In`	For the input flow: • Source and destination addresses and protocol tuple for the input flow. • `Interface`: Input flow interface. • `Session token`: Internal token derived from the virtual routing instance. • `Flag`: Internal debugging flags. • `Route`: Internal next hop of the route to be used by the flow. • `Gateway`: Next-hop gateway of the flow. • `Tunnel`: If the flow is going into a tunnel, the tunnel ID. Otherwise, 0 (zero). • `Port Sequence, FIN sequence, FIN state, Cookie`: Internal TCP state tracking information.
`Out`	For the reverse flow: • Source and destination addresses, and protocol tuple for the reverse flow. • `Interface`: Reverse flow interface. • `Session token`: Internal token derived from the virtual routing instance. • `Flag`: Internal debugging flags. • `Route`: Internal next hop of the route to be used by the flow. • `Gateway`: Next-hop gateway of the flow. • `Tunnel`: If the flow is going into a tunnel, the tunnel ID. Otherwise, 0 (zero). • `Port Sequence, FIN sequence, FIN state, Cookie`: Internal TCP state tracking information.
`Total sessions`	Total number of sessions.
`Status`	Session status.

Table 1188: show security flow session policy-id Output Fields *(Continued)*

Field Name	Field Description
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Dynamic application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.

Sample Output

show security flow session policy-id 4

```
root> show security flow session policy-id 4
Flow Sessions on FPC1 PIC0:

Session ID: 20093273, Policy name: p1/4, Timeout: 1784, Valid
  In: 101.0.0.2/1 --> 111.0.0.3/1;0, If: ge-0/0/0.0, Pkts: 1, Bytes: 84
  Out: 111.0.0.3/1 --> 201.0.0.1/22643;0, If: ge-0/0/1.0, Pkts: 0, Bytes: 0
Total sessions: 1
```

show security flow session policy-id 4 extensive

```
root> show security flow session policy-id 4 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000428, Status: Normal
Flags: 0x0/0x0/0x2008003
Policy name: p1/4
Source NAT pool: interface, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1740
Session State: Valid
Start time: 152305, Duration: 64
   In: 200.0.0.10/15192 --> 60.0.0.2/23;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0001021
    Route: 0x70010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 40, Bytes: 2251
    CP Session ID: 420001128
   Out: 60.0.0.2/23 --> 60.0.0.1/8078;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xe0001020
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 28, Bytes: 1714
    CP Session ID: 430000965
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 12.3X48-D10.

show security flow session pretty

IN THIS SECTION

- Syntax | 24259
- Description | 24259
- Required Privilege Level | 24259
- Output Fields | 24259
- Sample Output | 24262
- Release Information | 24263

Syntax

```
show security flow session pretty
```

Description

Displays the detailed information about security flow session. The output displays the information in a list to make it easy for you to read and monitor the flow session.

Required Privilege Level

View

Output Fields

Table 1189 on page 24260 lists the output fields for the `show security flow session pretty` command.

Table 1189: show security flow session pretty Output Fields

Field Name	Field Description
Advanced Security Feature Status	Status of the advanced security feature status such as Content Security, IDP, Advance Anti-Malware, Security-Intelligence, ICAP-Redirect , and App-Routing.
Bandwidth	Client to server and server to client bytes added together and converted to MB or GB.
Base Layer7 Application	Base layer7 application details.
Decrypted	SSL proxy decryption status. • True: SSL proxy decrypted. • False: SSL proxy not decrypted.
Decrypted-Mirrored	Decryption mirror status.
Destination-User	Destination user details.
Destination-Zone	Flow session destination zone details.
Duration	Duration of the session.
Final Layer7 Application	The last classified application.
Forwarding Type	Show the type of current flow session.
Forward Direction	Interface, IP address, protocol, tag, and gateway information of this direction.
Hardware Accelerated	Status of the Express Path flow session.
Inspection Mode	Flow session inspection mode.

Table 1189: show security flow session pretty Output Fields *(Continued)*

Field Name	Field Description
JDPI Classification	JDPI classification status.
Layer4 Application	Layer4 applications details.
LSYS	Logical system details.
Plugin List	List of the plugins used in the flow session.
Policy	Policy details.
Reverse Direction	Interface, IP address, protocol, tag, and gateway information of the flowdirection.
Session ID	Number that identifies the session. Use this ID to get more information about the session.
Session HA Status	Show the HA status.
Session State	Show the HA status.
Session State	Status of security flow session.
Start time	Time when the session was created, offset from the system start time.
Source-User	Shows the user source address.
Source-Zone	Flow session source zone details.
Time Out	Idle time-out after which the session expires.
TSYS	Tenant system details.
Tunnel Inspection	Tunnel inspection status.

Table 1189: show security flow session pretty Output Fields *(Continued)*

Field Name	Field Description
Tunneled Session	Status of the tunnel session. • True: Session is a tunnel session. • False: Session is not a tunnel session.

Sample Output

show security flow session pretty

```
user@host> show security flow session pretty
```

```
Session ID                   : 438249
Forward Direction            : Interface:ge-0/0/0.1, 192.0.2.0/45750 --> 203.0.113.0/443;tcp,
tag 0x0, Gateway 192.0.2.0
Reverse Direction            : Interface:ge-0/0/2.0, 203.0.113.0/443 -->
198.51.100.04/45750;tcp, tag 0x0, Gateway 198.51.100.0
Forwarding Type              : Route/Nat/Transparent/Secure-Wire/Tap
Session State                : Valid
Session HA Status            : Stand-alone
Start Time                   : 350 s
Duration                     : 19 s
Timeout                      : 1800 s
Time-to-Live                 : 1799 s
Bandwidth                    : 18MB
Source-User                  : example1.com\user123
Destination-User             : Unknown
LSYS                         : LSYS-X
TSYS                         : TSYS-Y
Source-Zone                  : Trust
Destination-Zone             : Untrust
Policy                       : permit_urls/5
Layer4 Application           : junos:https
Base Layer7 Application      : junos:SSL
Final Layer7 Application     : junos:CNN
```

```
Inspection Mode              : Flow
Hardware Accelerated         : False
Decrypted                    : False
Decryption-Mirrored          : False
JDPI Classification          : Complete
Tunneled Session             : False
Tunnel Inspection            : N/A
Advanced Security Feature Status:
UTM                          : Flow Mode
IDP                          : Flow Mode
Advance Anti-Malware         : Flow Mode
Security-Intelligence        : Flow Mode
ICAP-Redirect                : Flow Mode
App-Routing                  : Flow Mode
```

Release Information

Command introduced in Junos OS Release 21.1R1.

show security flow session protocol

IN THIS SECTION

- Syntax | 24264
- Description | 24264
- Options | 24264
- Required Privilege Level | 24265
- Output Fields | 24265
- Sample Output | 24267
- Release Information | 24270

Syntax

```
show security flow session protocol ( protocol-name | protocol-number )
[brief | extensive | summary]
```

Description

Display information about each session that uses the specified protocol.

Options

protocol-name —(Optional) Protocol to use as a sessions filter. Information about sessions that use this protocol is displayed. Possible protocols are:

- `ah`—IP Security Authentication Header
- `egp`—Exterior gateway protocol
- `esp`—IPsec Encapsulating Security Payload
- `gre`—Generic routing encapsulation
- `icmp`—Internet Control Message Protocol
- `igmp`—Internet Group Management Protocol
- `ipip`—IP over IP
- `ospf`—Open Shortest Path First
- `pim`—Protocol Independent Multicast
- `rsvp`—Resource Reservation Protocol
- `sctp`—Stream Control Transmission Protocol
- `tcp`—Transmission Control Protocol
- `udp`—User Datagram Protocol

protocol-number —(Optional) Numeric protocol value. For a complete list of possible numeric values, see RFC 1700, *Assigned Numbers* (*for the Internet Protocol Suite*).

Range: 0 through 255

brief | extensive | summary Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1190 on page 24265 lists the output fields for the `show security flow session protocol` command. Output fields are listed in the approximate order in which they appear.

Table 1190: show security flow session protocol Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Total sessions`	Total number of sessions.
`Status`	Session status.
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.

Table 1190: show security flow session protocol Output Fields *(Continued)*

Field Name	Field Description
`Policy name`	Name and ID of the policy that the first packet of the session matched.
`Source NAT pool`	The name of the source pool where NAT is used.
`Application`	Name of the application.
`Maximum timeout`	Maximum session timeout.
`Current timeout`	Remaining time for the session unless traffic exists in the session.
`Session State`	Session state.
`Start time`	Time when the session was created, offset from the system start time.
`Unicast-sessions`	Number of unicast sessions.
`Multicast-sessions`	Number of multicast sessions.
`Failed-sessions`	Number of failed sessions.
`Sessions-in-use`	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
`Maximum-sessions`	Number of maximum sessions.

Sample Output

show security flow session protocol icmp

```
root> show security flow session protocol icmp
Flow Sessions on FPC10 PIC1:

Session ID: 410000654, Policy name: p1/4, Timeout: 2, Valid
  In: 200.0.0.10/2 --> 60.0.0.2/15685;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
410001264
  Out: 60.0.0.2/15685 --> 200.0.0.10/2;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84, CP Session ID:
410001264
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000399, Policy name: p1/4, Timeout: 2, Valid
  In: 200.0.0.10/3 --> 60.0.0.2/15685;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
430001053
  Out: 60.0.0.2/15685 --> 200.0.0.10/3;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84, CP Session ID:
430001053
Total sessions: 1
```

show security flow session protocol icmp brief

```
root> show security flow session protocol icmp brief
Flow Sessions on FPC10 PIC1:

Session ID: 410000658, Policy name: p1/4, Timeout: 4, Valid
  In: 200.0.0.10/4 --> 60.0.0.2/16453;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
410001268
  Out: 60.0.0.2/16453 --> 200.0.0.10/4;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84, CP Session ID:
410001268
Total sessions: 1

Flow Sessions on FPC10 PIC2:
```

```
Session ID: 420000612, Policy name: p1/4, Timeout: 2, Valid
  In: 200.0.0.10/5 --> 60.0.0.2/16453;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
420001316
  Out: 60.0.0.2/16453 --> 200.0.0.10/5;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84, CP Session ID:
420001316
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Session ID: 430000405, Policy name: p1/4, Timeout: 2, Valid
  In: 200.0.0.10/6 --> 60.0.0.2/16453;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84, CP Session ID:
430001059
  Out: 60.0.0.2/16453 --> 200.0.0.10/6;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84, CP Session ID:
430001059
Total sessions: 1
```

show security flow session protocol icmp extensive

```
root> show security flow session protocol icmp extensive
Flow Sessions on FPC10 PIC1:

Session ID: 410000660, Status: Normal
Flags: 0x80000040/0x0/0x2800003
Policy name: p1/4
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 4, Current timeout: 2
Session State: Valid
Start time: 153201, Duration: 3
   In: 200.0.0.10/8 --> 60.0.0.2/16453;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
    Route: 0x70010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 410001270
   Out: 60.0.0.2/16453 --> 200.0.0.10/8;icmp,
    Interface: ge-7/1/1.0,
```

```
    Session token: 0x7, Flag: 0xc0000020
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 410001270
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session protocol icmp summary

```
root> show security flow session protocol icmp summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 1
Sessions in other states: 0
Total sessions: 3

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 1
Sessions in other states: 0
Total sessions: 3
```

Release Information

Command introduced in Junos OS Release 8.5; Filter and view options introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session protocol | 3915

show security flow session resource-manager

IN THIS SECTION

- Syntax | 24270
- Description | 24270
- Options | 24271
- Required Privilege Level | 24271
- Output Fields | 24271
- Sample Output | 24273
- Release Information | 24275

Syntax

```
show security flow session resource-manager
[brief | extensive | summary]
```

Description

This command displays information about sessions created by the resource manager.

Options

none—Display all resource manager sessions.

brief | extensive | summary Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1191 on page 24271 lists the output fields for the `show security flow session resource-manager` command. Output fields are listed in the approximate order in which they appear.

Table 1191: show security flow session resource-manager Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.
`Resource information`	Information about the session particular to the resource manager, including the name of the ALG, the group ID. and the resource ID.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Total sessions`	Total number of sessions.

Table 1191: show security flow session resource-manager Output Fields *(Continued)*

Field Name	Field Description
`Status`	Session status.
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.
`Policy name`	Name and ID of the policy that the first packet of the session matched.
`Source NAT pool`	The name of the source pool where NAT is used.
`Application`	Name of the application.
`Maximum timeout`	Maximum session timeout.
`Current timeout`	Remaining time for the session unless traffic exists in the session.
`Session State`	Session state.
`Start time`	Time when the session was created, offset from the system start time.
`Valid sessions`	Number of valid sessions.
`Pending sessions`	Number of pending sessions.
`Invalidated sessions`	Number of invalidated sessions.
`Sessions in other states`	Number of sessions in other states.
`CP Session ID`	Number that identifies the central point session. Use this ID to get more information about the central point session.

Sample Output

show security flow session resource-manager

```
root> show security flow session resource-manager
Flow Sessions on FPC10 PIC1:

Session ID: 410000664, Policy name: p1/4, Timeout: 1734, Valid
Resource information : FTP ALG, 1, 0
  In: 200.0.0.10/41047 --> 60.0.0.2/21;tcp, If: ge-7/1/0.0, Pkts: 13, Bytes: 586, CP Session ID:
410001274
  Out: 60.0.0.2/21 --> 200.0.0.10/41047;tcp, If: ge-7/1/1.0, Pkts: 13, Bytes: 803, CP Session
ID: 410001274
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session resource-manager brief

```
root> show security flow session resource-manager brief
Flow Sessions on FPC10 PIC1:

Session ID: 410000664, Policy name: p1/4, Timeout: 1704, Valid
Resource information : FTP ALG, 1, 0
  In: 200.0.0.10/41047 --> 60.0.0.2/21;tcp, If: ge-7/1/0.0, Pkts: 13, Bytes: 586, CP Session ID:
410001274
  Out: 60.0.0.2/21 --> 200.0.0.10/41047;tcp, If: ge-7/1/1.0, Pkts: 13, Bytes: 803, CP Session
ID: 410001274
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session resource-manager extensive

```
root> show security flow session resource-manager extensive
Flow Sessions on FPC10 PIC1:

Session ID: 410000664, Status: Normal
Flags: 0x42/0x0/0x2010103
Policy name: p1/4
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1682
Session State: Valid
Start time: 160496, Duration: 153
Client: FTP ALG, Group: 1, Resource: 0
   In: 200.0.0.10/41047 --> 60.0.0.2/21;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0002621
    Route: 0x70010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 13, Bytes: 586
    CP Session ID: 410001274
   Out: 60.0.0.2/21 --> 200.0.0.10/41047;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xc0002620
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 13, Bytes: 803
    CP Session ID: 410001274
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session resource-manager summary

```
root> show security flow session resource-manager summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 8.5.

Filter and view options introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session resource-manager | 3918

show security flow session services-offload

IN THIS SECTION

- Syntax | 24276
- Description | 24276
- Options | 24276
- Required Privilege Level | 24277
- Output Fields | 24277
- Sample Output | 24279
- Release Information | 24283

Syntax

```
show security flow session services-offload
[filter] [brief | extensive | summary]
```

Description

Display information about all currently active services-offload security sessions on the device.

Options

- *filter*—Filter the display by the specified criteria.

  The following filters reduce the display to those sessions that match the criteria specified by the filter:

`application`	Application name.
`application-firewall-rule-set`	Application firewall enabled with the specified rule set.
`application-traffic-control-rule-set`	Application traffic control enabled with the specified rule set.
`destination-port`	Destination port.

destination-prefix	Destination IP prefix or address.
dynamic-application	Dynamic application name.
dynamic-application-group	Dynamic application group name.
encrypted	Show encrypted traffic.
family	Protocol family.
interface	Name of incoming or outgoing interface.
logical-system	Logical system name.
protocol	IP protocol number.
root-logical-system	Root logical system name.
source-port	Source port.
source-prefix	Source IP prefix or address.
tenant	Tenant system name.

- `brief | extensive | summary`—Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1192 on page 24277 lists the output fields for the `show security flow session services-offload` command. Output fields are listed in the approximate order in which they appear.

Table 1192: show security flow session services-offload Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the services-offload session. Use this ID to get more information about the session.

Table 1192: show security flow session services-offload Output Fields *(Continued)*

Field Name	Field Description
`Policy name`	Policy that permits the services-offload traffic.
`Timeout`	Idle timeout period after which the services-offload session expires.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets, and bytes).
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets, and bytes).
`Total sessions`	Total number of services-offload sessions.
`Status`	Services-offload session status.
`Flag`	Internal flag depicting the state of the services-offload session, used for debugging purposes.
`Policy name`	Name and ID of the policy that the first packet of the services-offload session matched.
`Source NAT pool`	The name of the source pool where NAT is used.
`Application`	Name of the application.
`Dynamic application`	Name of the dynamic application.
`Maximum timeout`	Maximum amount of idle time allowed for the services-offload session.
`Current timeout`	Number of seconds that the current services-offload session has been idle.

Table 1192: show security flow session services-offload Output Fields *(Continued)*

Field Name	Field Description
Session State	Services-offload session state.
Start time	Time when the services-offload session was created, offset from the system start time.
Duration	Duration of the services-offload session.
Valid sessions	Number of valid services-offload sessions.
Pending sessions	Number of pending services-offload sessions.
Invalidated sessions	Number of invalidated services-offload sessions.
Sessions in other states	Number of services-offload sessions in other states.
Total sessions	Total number of services-offload sessions.

Sample Output

show security flow session services-offload

```
user@host>show security flow session services-offload
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 1788, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507, CP Session ID:
420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462, CP Session ID:
420000002
```

```
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session services-offload brief

```
user@host>show security flow session services-offload brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 1748, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507, CP Session ID:
420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462, CP Session ID:
420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session services-offload extensive

```
user@host>show security flow session services-offload extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Status: Normal
Flags: 0x40/0x0/0x2408003, services-offload
Policy name: p1/4
Source NAT pool: Null, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1718
Session State: Valid
```

```
Start time: 165, Duration: 89
   In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0x42001021
    Route: 0x80010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 507
    CP Session ID: 420000002
   Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x42001020
    Route: 0x70010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 462
    CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session services-offload summary

```
user@host>show security flow session services-offload summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

```
Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

show security flow session services-offload summary logical-system LSYS1

```
user@host>show security flow session services-offload summary logical-system LSYS1
Valid sessions: 500
           Pending sessions: 0
           Invalidated sessions: 0
           Sessions in other states: 0
           Total sessions: 500
```

show security flow session services-offload summary

```
user@host:LSYS1>show security flow session services-offload summary
Valid sessions: 500
           Pending sessions: 0
           Invalidated sessions: 0
           Sessions in other states: 0
           Total sessions: 500
```

show security flow session services-offload summary tenant TSYS1

```
user@host>show security flow session services-offload summary tenant TSYS1
Valid sessions: 10
           Pending sessions: 0
           Invalidated sessions: 0
           Sessions in other states: 0
           Total sessions: 10
```

show security flow session services-offload summary

```
user@host:TSYS1>show security flow session services-offload summary
Valid sessions: 10
           Pending sessions: 0
           Invalidated sessions: 0
           Sessions in other states: 0
           Total sessions: 10
```

Release Information

Command introduced in Junos OS Release 11.4.

Low-latency option introduced in Junos OS Release 12.1X44-D10.

Starting in Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) with Express Path (formerly known as *services offloading*) support are introduced for SRX5400, SRX5600, and SRX5800 Series devices.

Support added on SRX5800, SRX5600, SRX5400 Series devices and vSRX Virtual Firewall.

Support added on SRX4600 Series device in Junos OS Release 20.1R1.

Support added at the logical system and tenant system level in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session services-offload | **3920**

show security flow session session-identifier

IN THIS SECTION

- Syntax | **24284**
- Description | **24284**
- Options | **24284**

- Required Privilege Level | 24284
- Output Fields | 24284
- Sample Output | 24290
- Release Information | 24291

Syntax

```
show security flow session session-identifier session-identifier
```

Description

Display detailed information for the identified session.

Options

session-identifier —Identifier of the session about which to display information.

Required Privilege Level

view

Output Fields

Table 1193 on page 24284 lists the output fields for the `show security flow session session-identifier` command. Output fields are listed in the approximate order in which they appear.

Table 1193: show security flow session session-identifier Output Fields

Field Name	Field Description
`Session ID`	A unique number that a server assigns a specific user for the duration of that session.

Table 1193: show security flow session session-identifier Output Fields *(Continued)*

Field Name	Field Description
`Flags`	Internal flag depicting the state of the session, used for debugging purposes. It is internal. The three available flags are: Flags: 0x42/0x0/0x1/0x8103 • natflag: • natflag2: • natflag3: • flag:
`Policy name`	Name and ID of the policy that the first packet of the session matched and that permitted the traffic. Session log entries are tied to policy configuration. Each main session event—create, close, and deny—creates a log entry if the controlling policy has enabled logging.
`Source NAT pool`	The name of the source pool where NAT is used. A NAT pool is a user-defined set of IP addresses that are used for translation. Unlike static NAT, where there is a one-to-one mapping that includes destination IP address translation in one direction and source IP address translation in the reverse direction, with source NAT, the original source IP address is translated to an IP address in the address pool. Source NAT is used to allow hosts with private IP addresses to access a public network.
`Dynamic application`	Dynamic application: INCONCLUSIVE. If the dynamic application has yet to be determined, the output indicates Pending. If the dynamic application cannot be determined, the output indicates junos: UNKNOWN. Traffic with an application ID of junos: UNKNOWN matches a dynamic application of junos: UNKNOWN. If there is no such rule defined, the default rule is applied. The term junos: UNKNOWN is a reserved keyword.

Table 1193: show security flow session session-identifier Output Fields *(Continued)*

Field Name	Field Description
`Encryption`	Type of encryption, if the application traffic is encrypted. Encryption: Unknown.
`Application traffic control rule-set`	Name of the application traffic control rule set. Application traffic control rule-set: INVALID.
`Maximum timeout`	Maximum session timeout, in seconds.
`Session state`	Session state: Valid.
`Start time`	Time, in seconds, when the session was created, offset from the system start time.

Table 1193: show security flow session session-identifier Output Fields *(Continued)*

Field Name	Field Description
`In`	For the input flow: • Source and destination addresses, ports, and protocol tuple for the input flow. • `Interface`: Input flow interface. • `Session token`: Internal token derived from the virtual routing instance. • `Flag`: Internal debugging flags. • `Route`: Internal next hop of the route to be used by the flow. • `Gateway`: Next-hop gateway of the flow. • `Tunnel`: Used for internal debugging. If the flow is going into a tunnel, the decimal format of the tunnel ID, plus the tunnel type. Otherwise, 0 (zero). See Table 1194 on page 24289 for tunnel type identification. • `Port Sequence, FIN sequence, FIN state, Cookie`: Internal TCP state tracking information. • `Pkts, Bytes, CP Session ID`: Packets and bytes matched on the wing, and the associated CP session ID of the wing. • `Conn tag`: Session connection tag for GRPS tunneling protocol, user plane (GTP-U) flow sessions and Stream Control Transmission Protocol (SCTP) flow sessions.

Table 1193: show security flow session session-identifier Output Fields *(Continued)*

Field Name	Field Description
`Out`	For the reverse flow: • Source and destination IP addresses, and application protocol for the reverse flow. • `Interface`: Reverse flow interface. • `Session token`: Internal token derived from the virtual routing instance. • `Flag`: Internal debugging flags. • `Route`: Internal next hop of the route to be used by the flow. • `Gateway`: Next-hop gateway of the flow. • `Tunnel`: Tunnel type identifier. Used for internal debugging. If the flow is going into a tunnel, species the decimal format of the tunnel ID, plus the tunnel type. Otherwise, 0 (zero). See Table 1194 on page 24289 for tunnel type identification. • `Port Sequence, FIN sequence, FIN state, Cookie`: Internal TCP state tracking information. • `Pkts, Bytes, CP Session ID`: Packets and bytes matched on the wing, and the associated CP session ID of the wing. • `Conn tag`: Session connection tag for GRPS tunneling protocol, user plane (GTP-U) flow sessions and Stream Control Transmission Protocol (SCTP) flow sessions.
`Status`	Session status: • Auth (NAT flag with NAT_AUTH) • Transparent (NAT flag with NAT_TRANSPARENT) • Expired (NAT flag with NAT_INVALID) • Normal (no flag)
`Virtual system`	Virtual system to which the session belongs (it is optional).

Table 1193: show security flow session session-identifier Output Fields *(Continued)*

Field Name	Field Description
`Application`	Application match for applying the rule. Application: junos-http/6.
`Rule`	Name of the application traffic control rule. Rule: INVALID.
`Current timeout`	Remaining time for the session unless traffic exists in the session.
`Duration`	Length of time, in seconds, for which the session is active.
`Client`	Name of the ALG, if there is resource manager.
`Group`	Group identification number, if there is resource manager.
`Resource`	Resource identification number, if there is resource manager.

Table 1194: Tunnel Type Identification

Binary (first 3 bits)	Hexadecimal (the first 4 bits)	Tunnel Type
0x000	0x0	no tunnel
0x001	0x2	TUNNEL_TYPE_IPSEC
0x010	0x4	TUNNEL_TYPE_L2TP
0x011	0x6	TUNNEL_TYPE_NATT
0x100	0x8	TUNNEL_TYPE_DS_LITE

Table 1194: Tunnel Type Identification *(Continued)*

Binary (first 3 bits)	Hexadecimal (the first 4 bits)	Tunnel Type
0x101	0xa	TUNNEL_TYPE_MCNH

Sample Output

show security flow session session-identifier 20595

```
root> show security flow session session-identifier 20595
Flow Sessions on FPC10 PIC2:

Session ID: 20595, Status: Normal
Flags: 0x42/0x0/0x1/0x8103
Policy name: pre-id-default-policy-logical-system-00/3
Source NAT pool: Null, Application: junos-http/6
Dynamic application: INCONCLUSIVE,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1788
Session State: Valid
Start time: 247736, Duration: 14
   In: 4.0.0.1/47931 --> 5.0.0.2/80;tcp,
    Conn Tag: 0x0, Interface: ge-0/0/2.0,
    Session token: 0x7, Flag: 0x3621
    Route: 0xe0010, Gateway: 4.0.0.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 3, Bytes: 176
    Out: 5.0.0.2/80 --> 4.0.0.1/47931;tcp,
    Conn Tag: 0x0, Interface: ge-0/0/3.0,
     Session token: 0x8, Flag: 0x3620
    Route: 0xc0010, Gateway: 5.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 2, Bytes: 120
 Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 8.5. Output changed to support natflag2 and flag in Junos OS Release 12.3X48-D10.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session session-identifier | 3924

show security flow session session-state

IN THIS SECTION

- Syntax | 24291
- Description | 24291
- Options | 24292
- Required Privilege Level | 24292
- Output Fields | 24292
- Sample Output | 24294
- Release Information | 24294

Syntax

```
show security flow session session-state session-state
```

Description

Display detailed information for the Multinode High Availability session.

Options

- active-warm—Multinode High Availability session with one active wing and one warm wing.
- backup—Multinode High Availability Layer 2 deployment backup session
- warm—Multinode High Availability Layer 3 deployment warm session

Required Privilege Level

view

Output Fields

Table 1195 on page 24292 lists the output fields for the command. Output fields are listed in the approximate order in which they appear.

Table 1195: show security flow session session-state

Field Name	Field Description
`Session ID`	A unique number that a server assigns a specific user for the duration of that session.
`Policy name`	Name and ID of the policy that the first packet of the session matched and that permitted the traffic.
state	State of the session.
`timeout`	Idle timeout after which the session expires.
`Session state`	Session state: Valid.

Table 1195: show security flow session session-state *(Continued)*

Field Name	Field Description
`In`	For the input flow: • Source and destination addresses, ports, and protocol tuple for the input flow. • `Conn tag`: Session connection tag for GRPS tunneling protocol, user plane (GTP-U) flow sessions and Stream Control Transmission Protocol (SCTP) flow sessions. • `Interface`: Input flow interface. • `Pkts, Bytes, CP Session ID`: Packets and bytes matched on the wing, and the associated CP session ID of the wing. • `state`: Session state. The state can be `Active` or `Warm`. • `HA Wing state`: High Availability wing state. The state can be `Active` or `Warm`.
`Out`	For the reverse flow: • Source and destination addresses, ports, and protocol tuple for the input flow. • `Conn tag`: Session connection tag for GRPS tunneling protocol, user plane (GTP-U) flow sessions and Stream Control Transmission Protocol (SCTP) flow sessions. • `Interface`: reverse flow interface. • `Pkts, Bytes, CP Session ID`: Packets and bytes matched on the wing, and the associated CP session ID of the wing. • `state`: Session state. The state can be `Active` or `Warm`. • `HA Wing state`: High Availability wing state. The state can be `Active` or `Warm`.
`total sessions`	Total number of sessions.

Sample Output

show security flow session session-state

```
user@host> show security flow session session-state active-warm
Flow Sessions on FPC3 PIC0:
Session ID: 805306411, Policy name: All/4, State: Active, Timeout: 1794, Valid
  In: 20.11.1.1/32876 --> 20.11.2.1/22;tcp, Conn Tag: 0x0, If: xe-1/0/3.0, Pkts: 25, Bytes:
5285, CP Session ID: 46, State: Active
  Out: 20.11.2.1/22 --> 20.11.1.1/32876;tcp, Conn Tag: 0x0, If: st0.0, Pkts: 26, Bytes: 6581, CP
Session ID: 46, State: Warm
Total sessions: 1
Flow Sessions on FPC3 PIC1:
Total sessions: 0
```

```
user@host> show security flow session session-state active-warm
Flow Sessions on FPC3 PIC0:
Session ID: 1099513705861, Policy name: default-policy-logical-system-00/2, HA State: Active,
Timeout: 1800, Session State: Valid
  In: 1.0.1.2/3000 --> 6.0.1.2/3000;tcp, Conn Tag: 0x0, If: xe-4/0/2.0, Pkts: 1239, Bytes:
121240, HA Wing State: Warm, CP Session ID: 2046202
  Out: 6.0.1.2/3000 --> 1.0.1.2/3000;tcp, Conn Tag: 0x0, If: xe-4/0/5.0, Pkts: 622, Bytes:
80568, HA Wing State: Active, CP Session ID: 2046202

Total sessions: 1
Flow Sessions on FPC3 PIC1:
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 23.4R1.

show security flow session source-port

IN THIS SECTION

- Syntax | 24295
- Description | 24295
- Options | 24295
- Required Privilege Level | 24295
- Output Fields | 24296
- Sample Output | 24297
- Release Information | 24300

Syntax

```
show security flow session source-port
source-port-number
[brief | extensive | summary]
```

Description

Display information about each session that uses the specified source port.

Options

- *source-port-number* —Number of the source port about which to display sessions information.
- brief | extensive | summary–Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1196 on page 24296 lists the output fields for the `show security flow session source-port` command. Output fields are listed in the approximate order in which they appear.

Table 1196: show security flow session source-port Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.
`Resource information`	Information about the session particular to the resource manager, including the name of the ALG, the group ID, and the resource ID.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
`Total sessions`	Total number of sessions.
`Status`	Session status.
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.
`Policy name`	Name and ID of the policy that the first packet of the session matched.
`Source NAT pool`	The name of the source pool where NAT is used.

Table 1196: show security flow session source-port Output Fields *(Continued)*

Field Name	Field Description
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Valid sessions	Number of valid sessions.
Pending sessions	Number of pending sessions.
Invalidated sessions	Number of invalidated sessions.
Sessions in other states	Number of sessions in other states.

Sample Output

show security flow session source-port 15198

```
root> show security flow session source-port 15198
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 770, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507, CP Session ID:
420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462, CP Session ID:
```

```
420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session source-port 15198 brief

```
root> show security flow session source-port 15198 brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 740, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507, CP Session ID:
420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462, CP Session ID:
420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session source-port 15198 extensive

```
root> show security flow session source-port 15198 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Status: Normal
Flags: 0x40/0x0/0x2408003, services-offload
Policy name: p1/4
Source NAT pool: Null, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 750
```

```
Session State: Valid
Start time: 165, Duration: 1056
   In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0x42001021
    Route: 0x80010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 507
    CP Session ID: 420000002
   Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x42001020
    Route: 0x70010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 462
    CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session source-port 15198 summary

```
root> show security flow session source-port 15198 summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

```
Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 8.5; Filter and view options introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session source-port | 3926

show security flow session source-prefix

IN THIS SECTION

- Syntax | 24301
- Description | 24301
- Options | 24301
- Required Privilege Level | 24301
- Output Fields | 24301
- Sample Output | 24303
- Release Information | 24305

Syntax

```
show security flow session source-prefix
          source-prefix-number
[brief | extensive | summary]
```

Description

Display information about each session that uses the specified source prefix.

Options

source-prefix-number—Source IP prefix or address for which to display sessions information.

brief | extensive | summary Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1197 on page 24301 lists the output fields for the `show security flow session source-prefix` command. Output fields are listed in the approximate order in which they appear.

Table 1197: show security flow session source-prefix Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.

Table 1197: show security flow session source-prefix Output Fields *(Continued)*

Field Name	Field Description
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Valid sessions	Number of valid sessions.
Pending sessions	Number of pending sessions.

Table 1197: show security flow session source-prefix Output Fields *(Continued)*

Field Name	Field Description
Invalidated sessions	Number of invalidated sessions.
Sessions in other states	Number of sessions in other states.

Sample Output

show security flow session source-prefix 200.0.0.10

```
root> show security flow session source-prefix 200.0.0.10
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 488, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507, CP Session ID:
420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462, CP Session ID:
420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session source-prefix 200.0.0.10 brief

```
root> show security flow session source-prefix 200.0.0.10 brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 482, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507, CP Session ID:
```

```
420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462, CP Session ID:
420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session source-prefix 200.0.0.10 extensive

```
root> show security flow session source-prefix 200.0.0.10 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Status: Normal
Flags: 0x40/0x0/0x2408003, services-offload
Policy name: p1/4
Source NAT pool: Null, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 436
Session State: Valid
Start time: 165, Duration: 1370
   In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0x42001021
    Route: 0x80010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 507
    CP Session ID: 420000002
   Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x42001020
    Route: 0x70010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 462
```

```
    CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session source-prefix 200.0.0.10 summary

```
root> show security flow session source-prefix 200.0.0.10 summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 8.5.

Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4.

Filter and view options introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session source-prefix | 3928

show security flow session ssl

IN THIS SECTION

- Syntax | 24306
- Description | 24306
- Options | 24306
- Required Privilege Level | 24307
- Output Fields | 24307
- Sample Output | 24309
- Release Information | 24310

Syntax

```
show security flow session ssl [brief | extensive]
```

Description

Display information about the active SSL sessions on the device.

Options

brief | extensive	Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1198 on page 24307 lists the output fields for the `show security flow session ssl` command. Output fields are listed in the approximate order in which they appear.

Table 1198: show security flow session ssl Output Fields

Field Name	Field Description	Displayed with Option
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.	brief, extensive
`Status`	Status of the session.	brief, extensive
`Policy name`	Policy that permitted the traffic. Name and ID of the policy that the first packet of the session matched.	brief, extensive
`Timeout`	Idle timeout after which the session expires.	brief, extensive
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).	brief, extensive
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).	brief, extensive
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.	extensive
`Source NAT pool`	The name of the source pool where NAT is used.	extensive

Table 1198: show security flow session ssl Output Fields *(Continued)*

Field Name	Field Description	Displayed with Option
`dynamic-application`	Name of the application.	extensive
`encryption`	Encryption applied.	extensive
`Application traffic control rule-set: INVALID, Rule: INVALID`	Name of the application quality of service rule.	extensive
`Maximum timeout`	Maximum session timeout.	extensive
`Current timeout`	Remaining time for the session unless traffic exists in the session.	extensive
`Session State`	Status of security ssl session.	extensive
`Start time`	Time when the session was created, offset from the system start time.	extensive
`duration`	Duration of the session	extensive
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).	extensive
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).	extensive
`Total Sessions`	Total number of sessions.	extensive

Sample Output

show security flow session ssl brief

```
user@host> show security flow session ssl brief

Output:

Session ID: 1, Policy name: default-permit/5, Timeout: 1746, Valid
In: 4.0.0.1/37369 --> 5.0.0.1/4433;tcp, Conn Tag: 0x0, If: xe-0/0/0.0, Pkts: 6, Bytes: 671,
Out: 5.0.0.1/4433 --> 4.0.0.1/37369;tcp, Conn Tag: 0x0, If: xe-0/0/1.0, Pkts: 7, Bytes: 1635,
```

show security flow session ssl extensive

```
user@host> show security flow session ssl extensive
Output:
Session ID: 1, Status: Normal
Flags: 0x42/0x20000000/0x2/0x10103
Policy name: 1/5
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption: Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1636
Session State: Valid
Start time: 587131, Duration: 163
In: 4.0.0.1/37369 --> 5.0.0.1/4433;tcp,
Conn Tag: 0x0, Interface: xe-0/0/0.0,
Session token: 0x7, Flag: 0x2621
Route: 0xa0010, Gateway: 4.0.0.1, Tunnel: 0
Port sequence: 0, FIN sequence: 0,
FIN state: 0,
Pkts: 6, Bytes: 671
Out: 5.0.0.1/4433 --> 4.0.0.1/37369;tcp,
Conn Tag: 0x0, Interface: xe-0/0/1.0,
Session token: 0x8, Flag: 0x2620
Route: 0xb0010, Gateway: 5.0.0.1, Tunnel: 0
Port sequence: 0, FIN sequence: 0,
FIN state: 0,
```

```
Pkts: 7, Bytes: 1635
Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show security flow session summary family

IN THIS SECTION

- Syntax | 24310
- Description | 24310
- Options | 24311
- Required Privilege Level | 24311
- Output Fields | 24311
- Sample Output | 24311
- Release Information | 24313

Syntax

```
show security flow session summary family (inet | inet6)
```

Description

Display filtered summary of information about existing sessions, including types of sessions, active and failed sessions, and the maximum allowed number of sessions.

Options

- `inet`—Display details summary of IPv4 sessions.
- `inet6`—Display details summary of IPv6 sessions.

Required Privilege Level

view

Output Fields

Table 1199 on page 24311 lists the output fields for the `show security flow session summary family` command. Output fields are listed in the approximate order in which they appear.

Table 1199: show security flow session summary Output Fields

Field Name	Field Description
`Valid sessions`	Count of valid sessions.
`Pending sessions`	Count of pending sessions.
`Invalidated sessions`	Count of sessions the security device has determined to be invalid.
`Sessions in other states`	Count of sessions not in valid, pending, or invalidated state.
`Total sessions`	Total of the above counts.

Sample Output

show security flow session summary family inet

```
user@host> show security flow session summary family inet
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
```

```
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow session summary family inet6

```
user@host> show security flow session summary family inet6
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:
```

```
Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 10.2.

Support on SRX Series Firewalls for flow-based mode for family inet6 added in Junos OS Release 10.2.

Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session | 3890

show security flow session summary node

IN THIS SECTION

- Syntax | 24314
- Description | 24314
- Options | 24314
- Required Privilege Level | 24314
- Output Fields | 24314
- Sample Output | 24315
- Release Information | 24321

Syntax

```
show security flow session summary node (node-id | all | local | primary)
```

Description

Display information about all currently active security sessions on the device for the specified node options in summary mode.

Options

`node`—(Optional) For chassis cluster configurations, display session information on a specific node.

- *`node-id`* —Identification number of the node. It can be 0 or 1.
- `all`—Display information about all nodes.
- `local`—Display information about the local node.
- `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1200 on page 24314 lists the output fields for the `show security flow session summary node` command. Output fields are listed in the approximate order in which they appear.

Table 1200: show security flow session summary node Output Fields

Field Name	Field Description
`Unicast-sessions`	Number of unicast sessions.
`Multicast-sessions`	Number of multicast sessions.
`Failed-sessions`	Number of failed sessions.

Table 1200: show security flow session summary node Output Fields *(Continued)*

Field Name	Field Description
`Sessions-in-use`	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
`Maximum-sessions`	Maximum number of sessions permitted.

Sample Output

show security flow session summary node 0

```
root@host> show security flow session summary node 0
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
```

```
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

show security flow session summary node 1

```
root@host> show security flow session summary node 1
node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
```

```
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

show security flow session summary node all

```
root@host> show security flow session summary node all
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
```

```
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
```

```
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

show security flow session summary node local

```
root@host> show security flow session summary node local
node0:
--------------------------------------------------------------------------



Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
```

```
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

show security flow session summary node primary

```
root@host> show security flow session summary node primary
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
```

```
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

Release Information

Command introduced in Junos OS Release 8.5; node options added in Junos OS Release 9.0. Filter options added in Junos OS Release 10.2.

Support on SRX Series Firewalls for flow-based mode for family inet6 added in Junos OS Release 10.2.

Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security flow session summary services-offload

IN THIS SECTION

- Syntax | 24322
- Description | 24322
- Options | 24322
- Required Privilege Level | 24323
- Output Fields | 24323
- Sample Output | 24324
- Release Information | 24326

Syntax

```
show security flow session summary services-offload  [filter]
```

Description

Display information about all currently active services-offload security sessions on the device in summary mode.

Options

- *filter*—Filter the display by the specified criteria.

  The following filters reduce the display to those sessions that match the criteria specified by the filter:

`application`	Application name.
`application-firewall-rule-set`	Application firewall enabled with the specified rule set.
`application-traffic-control-rule-set`	Application traffic control enabled with the specified rule set.
`destination-port`	Destination port.

destination-prefix	Destination IP prefix or address.
dynamic-application	Dynamic application name.
dynamic-application-group	Dynamic application group name.
family	Protocol family.
interface	Name of incoming or outgoing interface.
logcal-system	Logical system name.
protocol	IP protocol number.
root-logical-system	Root logical system name.
source-port	Source port.
source-prefix	Source IP prefix or address.

Required Privilege Level

view

Output Fields

Table 1201 on page 24323 lists the output fields for the `show security flow session summary services-offload` command. Output fields are listed in the approximate order in which they appear.

Table 1201: show security flow session summary services-offload Output Fields

Field Name	Field Description
`Unicast-sessions`	Number of unicast sessions.
`Multicast-sessions`	Number of multicast sessions.
`Services-offload-sessions`	Number of services-offload sessions.
`Failed-sessions`	Number of failed sessions.

Table 1201: show security flow session summary services-offload Output Fields *(Continued)*

Field Name	Field Description
`Sessions-in-use`	Number of sessions in use: • Valid • Pending • Invalidated • Sessions in other states
`Maximum-sessions`	Maximum number of sessions.

Sample Output

show security flow session summary services-offload

```
user@host> show security flow session summary services-offload
Flow Sessions on FPC1 PIC0:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 409600

Flow Sessions on FPC2 PIC0:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 1
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
```

```
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 819200

Flow Sessions on FPC3 PIC0:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 819200

Flow Sessions on FPC5 PIC0:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 819200
```

show security flow session summary services-offload application

```
user@host> show security flow session summary services-offload application telnet
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:
```

```
Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

show security flow session summary services-offload destination-port

```
user@host> show security flow session summary services-offload destination-port 23
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000004, Policy name: p1/4, Timeout: 1500, Valid
  In: 200.0.0.10/15200 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 13, Bytes: 718, CP Session ID:
430000003
  Out: 60.0.0.2/23 --> 200.0.0.10/15200;tcp, If: ge-7/1/1.0, Pkts: 12, Bytes: 677, CP Session
ID: 430000003
Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 11.4.

Starting with Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) with Express Path (formerly known as *services offloading*) support are introduced for SRX5400, SRX5600, and SRX5800 devices.

This command is supported on the SRX1500, SRX 5800, SRX 5600, and SRX 5400 devices, and vSRX Virtual Firewall.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session services-offload | 3920

show security flow session tenant

IN THIS SECTION

- Syntax | 24327
- Description | 24327
- Options | 24328
- Required Privilege Level | 24328
- Output Fields | 24328
- Sample Output | 24329
- Release Information | 24330

Syntax

```
show security flow session tenant (tenant-name | all)
```

Description

Displays the information about the currently active security flow sessions of the tenant systems on the device. You can either view the currently active security flow sessions information for a specific tenant system or for all the tenant systems.

Options

tenant-name Name of the tenant system.

all Displays the security flow session information for all the tenant systems.

Required Privilege Level

view

Output Fields

Table 1202 on page 24328 lists the output fields for the `show security flow session tenant` command. Output fields are listed in the approximate order in which they appear.

Table 1202: show security flow session tenant

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.
`In`	Incoming security flow session details. The incoming security flow session details include the source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets, and bytes.
`Out`	Reverse security flow session details. The reverse security flow session details include the source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets, and bytes.
`Total sessions`	Total number of security flow sessions.

Table 1202: show security flow session tenant *(Continued)*

Field Name	Field Description
`Policy name`	Name and ID of the policy that the first packet of the security flow session matched.
`Tenant`	Name of the tenant system.

Sample Output

show security flow session tenant T1

```
root@host> show security flow session tenant T1
Flow Sessions on FPC7 PIC1:

Session ID: 290000224, Policy name: default-policy-logical-system-32/2, Timeout: 1790, Valid
Tenant: T1
  In: 203.0.113.0/39767 --> 203.0.113.1/23;tcp, Conn Tag: 0x0, If: xe-3/0/1.0, Pkts: 39, Bytes:
2136, CP Session ID: 1225556754
  Out: 203.0.113.1/23 --> 203.0.113.0/39767;tcp, Conn Tag: 0x0, If: lt-0/0/0.101, Pkts: 31,
Bytes: 1872, CP Session ID: 1225556754
```

show security flow session tenant all

```
root@host> show security flow session tenant all
Flow Sessions on FPC7 PIC1:

Session ID: 290000224, Policy name: default-policy-logical-system-32/2, Timeout: 1790, Valid
Tenant: T1
  In: 203.0.113.0/39767 --> 203.0.113.1/23;tcp, Conn Tag: 0x0, If: xe-3/0/1.0, Pkts: 39, Bytes:
2136, CP Session ID: 1225556754
  Out: 203.0.113.1/23 --> 203.0.113.0/39767;tcp, Conn Tag: 0x0, If: lt-0/0/0.101, Pkts: 31,
Bytes: 1872, CP Session ID: 1225556754

Session ID: 290000225, Policy name: default-policy-logical-system-33/2, Timeout: 1790, Valid
Tenant: T2
```

```
  In: 203.0.113.3/39767 --> 203.0.113.4/23;tcp, Conn Tag: 0x0, If: lt-0/0/0.103, Pkts: 39,
Bytes: 2136, CP Session ID: 1225556755
  Out: 203.0.113.4/23 --> 203.0.113.3/39767;tcp, Conn Tag: 0x0, If: xe-9/0/0.0, Pkts: 31, Bytes:
1872, CP Session ID: 1225556755
Total sessions: 2

Flow Sessions on FPC7 PIC2:
Total sessions: 0

Flow Sessions on FPC7 PIC3:
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

clear security flow session tenant | 3931

show security flow session tunnel

IN THIS SECTION

- Syntax | 24331
- Description | 24331
- Options | 24331
- Required Privilege Level | 24331
- Output Fields | 24331
- Sample Output | 24333
- Release Information | 24341

Syntax

```
show security flow session tunnel
[brief | extensive | summary]
```

Description

Display information about all tunnel sessions.

Options

none—Display the `brief` (default) level of output.

brief—Display the specified brief level of output.

extensive—Display the information about all current active sessions.

summary—Display the session information on each FPC.

Required Privilege Level

view

Output Fields

Table 1203 on page 24331 lists the output fields for the `show security flow session tunnel` command. Output fields are listed in the approximate order in which they appear.

Table 1203: show security flow session tunnel Output Fields

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic. NA (Not Applicable) for a tunnel session.
`Source NAT pool`	The name of the source pool where NAT is used.

Table 1203: show security flow session tunnel Output Fields *(Continued)*

Field Name	Field Description
`Timeout`	Idle timeout after which the session expires. NA (Not Applicable) for a tunnel session.
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, encapsulation and authentication header fragments generated, inner IPv4 fragments generated, inner IPv6 fragments generated, port sequence, FIN sequence, FIN state, packets and bytes).
`Total sessions`	Total number of sessions.
`Status`	Session status.
`Flags`	Internal flag depicting the state of the session, used for debugging purposes.
`Source NAT pool`	The name of the source pool where NAT is used.
`Application`	Name of the application.
`Maximum timeout`	Maximum session timeout.
`Current timeout`	Remaining time for the session unless traffic exists in the session.
`Encryption`	Encryption traffic name.
`Session State`	Session state.
`Start time`	Time when the session was created, offset from the system start time.
`Session token`	Internal token derived from the virtual routing instance.

Table 1203: show security flow session tunnel Output Fields *(Continued)*

Field Name	Field Description
`Route`	Internal next hop of the route to be used by the flow.
`Valid sessions`	Number of valid sessions.
`Pending sessions`	Number of pending sessions.
`Invalidated sessions`	Number of invalidated sessions.
`Sessions in other states`	Number of sessions in other states.
`ESP/AH frag Rx:` *number*, `Generated:` *number*	For IPsec tunnels, the number of Encapsulating Security Payload (ESP) or Authentication Header (AH) fragments that were received and the number that were generated.
`Inner IPv4 frag Rx:` *number*, `Tx:` *number*, `Generated:` *number*	For tunnels with IPv4 fragments, the number of fragments associated with the tunnel that were received, transmitted, and generated.
`Inner IPv6 frag Rx:` *number*, `Tx:` *number*, `Generated:` *number*	For tunnels with IPv6 fragments, the number of fragments associated with the tunnel that were received, transmitted, and generated.

Sample Output

show security flow session tunnel

```
root> show security flow session tunnel
Flow Sessions on FPC10 PIC1:

Session ID: 410000001, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/43405 --> 60.0.0.3/494;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID:
420000000

Session ID: 410000002, Policy name: N/A, Timeout: N/A, Valid
```

```
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID: 420000000
Total sessions: 2

Flow Sessions on FPC10 PIC2:

Session ID: 420000003, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID: 420000000

Session ID: 420000004, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;ah, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID: 420000000
Total sessions: 2

Flow Sessions on FPC10 PIC3:

Session ID: 430000005, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID: 420000000

Session ID: 430000006, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;ah, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID: 420000000
Total sessions: 2
```

show security flow session tunnel brief

```
root> show security flow session tunnel brief
Flow Sessions on FPC10 PIC1:

Session ID: 410000001, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/43405 --> 60.0.0.3/494;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID:
420000000

Session ID: 410000002, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID: 420000000
Total sessions: 2

Flow Sessions on FPC10 PIC2:

Session ID: 420000003, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID: 420000000

Session ID: 420000004, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;ah, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID: 420000000
```

```
Total sessions: 2

Flow Sessions on FPC10 PIC3:

Session ID: 430000005, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID: 420000000

Session ID: 430000006, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;ah, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session ID: 420000000
Total sessions: 2
```

show security flow session tunnel extensive

```
root> show security flow session tunnel extensive
Flow Sessions on FPC10 PIC1:

Session ID: 410000001, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3548, Duration: 797
   In: 60.0.0.2/43405 --> 60.0.0.3/494;esp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x80100621
    Route: 0x60010, Gateway: 60.0.0.2, Tunnel: 0
    ESP/AH frag Rx: 0, Generated: 0
      Inner IPv4 frag Rx: 4, Tx: 4, Generated: 4,
    Inner IPv6 frag Rx: 0, Tx: 0, Generated: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000000

Session ID: 410000002, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
```

```
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3548, Duration: 797
   In: 60.0.0.2/0 --> 60.0.0.3/0;esp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x621
    Route: 0x60010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000000
Total sessions: 2

Flow Sessions on FPC10 PIC2:

Session ID: 420000003, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3513, Duration: 798
   In: 60.0.0.2/0 --> 60.0.0.3/0;esp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x621
    Route: 0x0, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000000


Session ID: 420000004, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
Source NAT pool: Null
```

```
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3513, Duration: 798
   In: 60.0.0.2/0 --> 60.0.0.3/0;ah,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x621
    Route: 0x0, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000000
Total sessions: 2

Flow Sessions on FPC10 PIC3:

Session ID: 430000005, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3513, Duration: 799
   In: 60.0.0.2/0 --> 60.0.0.3/0;esp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x621
    Route: 0x0, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000000

Session ID: 430000006, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
```

```
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3513, Duration: 799
   In: 60.0.0.2/0 --> 60.0.0.3/0;ah,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x621
    Route: 0x0, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000000
Total sessions: 2
```

show security flow session tunnel summary extensive (with fragmentation counters output)

```
root> show security flow session tunnel extensive
node0:
Flow Sessions on FPC2 PIC1:
  Session ID: 90000004, Status: Normal, State: Active
  Flags: 0x10000/0x0/0x1
  Policy name: N/A
  Source NAT pool: Null
  Dynamic application: junos:UNKNOWN,
  Encryption:  Unknown
  Application traffic control rule-set: INVALID, Rule: INVALID
  Maximum timeout: N/A, Current timeout: N/A
  Session State: Valid
  Start time: 6251, Duration: 167168
    In: 2.2.2.2/0 --> 2.2.2.1/10203;esp,
    Conn Tag: 0x0, Interface: reth1.0,
    Session token: 0x7, Flag: 0x80100621
    Route: 0x867f3c1, Gateway: 2.2.2.2, Tunnel: 0
    ESP/AH frag Rx: 0, Generated: 0
    Inner IPv4 frag Rx: 27, Tx: 27, Generated: 18,
    Inner IPv6 frag Rx: 0, Tx: 0, Generated: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 90000000
```

```
    Session ID: 90000005, Status: Normal, State: Active
    Flags: 0x10000/0x0/0x1
    Policy name: N/A
    Source NAT pool: Null
    Dynamic application: junos:UNKNOWN,
    Encryption:  Unknown
    Application traffic control rule-set: INVALID, Rule: INVALID
    Maximum timeout: N/A,
    Current timeout: N/A
    Session State: Valid
    Start time: 6251, Duration: 167168
       In: 2.2.2.2/0 --> 2.2.2.1/0;esp,
       Conn Tag: 0x0, Interface: reth1.0,
       Session token: 0x7, Flag: 0x100621
       Route: 0x867f3c1, Gateway: 2.2.2.2, Tunnel: 0
       ESP/AH frag Rx: 0, Generated: 0
       Inner IPv4 frag Rx: 0, Tx: 0, Generated: 0,
       Inner IPv6 frag Rx: 0, Tx: 0, Generated: 0
       Port sequence: 0,
       FIN sequence: 0,
       FIN state: 0,
       Pkts: 0, Bytes: 0

       CP Session ID: 90000000
      Total sessions: 2
```

show security flow session tunnel summary (with fragmentation counters output)

```
root> show security flow session tunnel summary
node0:
Flow Sessions on FPC2 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2

Tunnel fragment summary:
```

```
Tunnels with ESP/AH frag Rx: 0 (0)
Tunnels with ESP/AH frag generated: 0 (0)
Tunnels with IPv4 frag Rx: 1 (27)
Tunnels with IPv4 frag Tx: 1 (27)
Tunnels with IPv4 frag generated: 1 (18)
Tunnels with IPv6 frag Rx: 0 (0)
Tunnels with IPv6 frag Tx: 0 (0)
Tunnels with IPv6 frag generated: 0 (0)

Flow Sessions on FPC2 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2



Tunnel fragment summary:
Tunnels with ESP/AH frag Rx: 0 (0)
Tunnels with ESP/AH frag generated: 0 (0)
Tunnels with IPv4 frag Rx: 0 (0)
Tunnels with IPv4 frag Tx: 0 (0)
Tunnels with IPv4 frag generated: 0 (0)
Tunnels with IPv6 frag Rx: 0 (0)
Tunnels with IPv6 frag Tx: 0 (0)
Tunnels with IPv6 frag generated: 0 (0)



Flow Sessions on FPC2 PIC3:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2

Tunnel fragment summary:
Tunnels with ESP/AH frag Rx: 0 (0)
Tunnels with ESP/AH frag generated: 0 (0)
Tunnels with IPv4 frag Rx: 0 (0)
Tunnels with IPv4 frag Tx: 0 (0)
Tunnels with IPv4 frag generated: 0 (0)
```

```
Tunnels with IPv6 frag Rx: 0 (0)
Tunnels with IPv6 frag Tx: 0 (0)
Tunnels with IPv6 frag generated: 0 (0)


Tunnel fragment summary:
Tunnels with ESP/AH frag Rx: 0 (0)
Tunnels with ESP/AH frag generated: 0 (0)
Tunnels with IPv4 frag Rx: 1 (27)
Tunnels with IPv4 frag Tx: 1 (27)
Tunnels with IPv4 frag generated: 1 (18)
Tunnels with IPv6 frag Rx: 0 (0)
Tunnels with IPv6 frag Tx: 0 (0)
Tunnels with IPv6 frag generated: 0 (0)
```

Release Information

Command introduced in Junos OS Release 8.5; Filter and view options introduced in Junos OS Release 10.2. Fragmentation counters options introduced in Junos OS Release 15.1X49-90.

NOTE: Only **show security flow session tunnel extensive** and **show security flow session tunnel summary** provide fragmentation counters output.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security flow statistics

IN THIS SECTION

- Syntax | **24342**
- Description | **24342**

- Options | 24343
- Required Privilege Level | 24344
- Output Fields | 24344
- Sample Output | 24345
- Release Information | 24347

Syntax

```
show security flow statistics
<node (node-id | all | local | primary)>
<logical-system (logical-system-name | all)>
<tenant (tenant-name | all)>
```

Description

Display security flow statistics on a specific SPU. A flow is a stream of related packets that meet the same matching criteria and share the same characteristics.

A packet undergoes flow-based processing after packet-based filters and some screens have been applied to it. A System Processing Unit (SPU) processes the packets of a flow according to the security features and other services configured for the session. Flow-based packet processing treats related packets, or a stream of packets, in the same way. Packet treatment depends on characteristics that were established for the first packet of the packet stream.

The `show security flow statistics` command displays information for individual SPUs. For each SPU, the active sessions on the SPU, packets received, packets transmitted, packets forwarded/queued, packets copied, packets dropped, packet fragments received in a flow on the SPU, pre-fragmented packets generated, and post-fragmented packets generated are displayed in terms of numbers.

There are many conditions that can cause a packet to be dropped. Here are some of them:

- A screen module detects IP spoofing
- The IPSec Encapsulating Security Payload (ESP) or the Authentication Header (AH) authentication failed. For example, incoming NAT errors could cause this to happen.
- A packet matches more than one security policy that specifies user authentication. (Sometimes packets are looped through the system more than once. Each time a packet passes through the system, that packet must be permitted by a policy.)

- A time constraint setting expires. For example, multicast streams with a packet interval of more than 60 seconds would experience premature aging-out of flow sessions. (In most cases, you can configure higher time-out value to prevent packet drop.)

Packet fragmentation can occur for a number of reasons, and, in some cases, it can be controlled through a configuration setting. Every link has a maximum transmission unit (MTU) size that specifies the size of the largest packet that the link can transmit. A larger MTU size means that fewer packets are required to transmit a certain amount of data. However, for a packet to successfully traverse the path from the source node to the destination node, the MTU size of the source node egress interface must be no larger than that of the smallest MTU size of all nodes on the path between the source and destination. This value is referred to as the path maximum transmission unit (path MTU).

When a packet is larger that the MTU size on any link in the data path, the link might fragment it or drop it.

- For IPv4, if a node within the path between a source node and a destination node receives a packet that is larger than its MTU size, it can fragment the packet and transmit the resulting smaller packets.
- For IPv6, an intermediate node cannot fragment a packet. If a packet is larger than a link's MTU size, it is likely that the link will drop it. However, the source node (the node that sent the packet) can fragment a packet, and this is done to accommodate a path MTU size-adjustment requirement. Nodes along the path of a packet cannot fragment the packet to transmit it.

The fragmentation counters feature for IPsec tunnels provides the show output information for the pre-fragments generated and post-fragments generated fields.

Starting in Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) are introduced for SRX5400, SRX5600, and SRX5800 devices that perform hash-based datapath packet forwarding to interconnect with all existing IOC and SPC cards using the XL chip (packet-processing chip). The IOC3 XL chip uses a hash-based method to distribute ingress traffic to a pool of SPUs by default.

Options

none Displays the security flow statistics information.

node (Optional) For chassis cluster configurations, displays all security flow statistics on a specific node (device) in the cluster.

- *node-id*—Identification number of the node. It can be 0 or 1.
- `all`—Displays information about all nodes.
- `local`—Displays information about the local node.

- `primary`—Displays information about the primary node.

`logical-system` ***`logical-system-name`***	(Optional) Displays information about the specified logical system.
`logical-system` ***`all`***	(Optional) Displays information about all the logical systems.
`tenant` ***`tenant-name`***	(Optional) Displays information about the specified tenant system.
`tenant` ***`all`***	(Optional) Displays information about all tenant systems.

Required Privilege Level

view

Output Fields

Table 1204 on page 24344 lists the output fields for the `show security flow statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1204: show security flow statistics Output Fields

Field Name	Field Description
Current sessions	Number of active sessions on the SPU.
Packets received	Number of packets received in a security flow of a specific SPU. The packets are processed and forwarded on that SPU.
Packets transmitted	Number of packets returned to Jexec for transmission.
Packets forwarded/queued	Number of packets fowarded or number of packets queued up by other modules. **NOTE**: Dropped packets are not captured by this field.
Packets copied	Number of packets copied by other modules including fragmentation and tcp proxy.

Table 1204: show security flow statistics Output Fields *(Continued)*

Field Name	Field Description
Packets dropped	Number of packets dropped in a flow on a specific SPU. The packets are received in the flow. However, during processing, the system discovers sanity check errors, security violations, or other conditions that caused the packet to be dropped. See the description for some of the conditions and events that can cause a packet to be dropped.
Fragment packets	Number of fragments received in a flow on the SPU. See the description for information about packet fragments.
Pre fragments generated	For IPsec tunnels, the number of fragments that are self-generated by the SRX Series Firewall before it encapsulates the packet with the IPsec encryption header.
Post fragments generated	For IPsec tunnels, the number of fragments that are received by the SRX Series Firewall and packets that are fragmented after encryption.

Sample Output

show security flow statistics

```
user@host> show security flow statistics
node0:
--------------------------------------------------------------------------

    Current sessions: 0
    Packets received: 2677
    Packets transmitted: 2278
    Packets forwarded/queued: 0
    Packets copied: 99
    Packets dropped: 300
    Fragment packets: 0
```

```
    Pre fragments generated: 0
    Post fragments generated: 0

node1:
--------------------------------------------------------------------------

    Current sessions: 0
    Packets received: 1267
    Packets transmitted: 904
    Packets forwarded/queued: 0
    Packets copied: 0
    Packets dropped: 363
    Fragment packets: 0
    Pre fragments generated: 0
    Post fragments generated: 0
```

show security flow statistics logical-system LSYS1

```
user@host> show security flow statistics logical-system LSYS1
Current sessions: 1000
           Packets received: 177888012
           Packets transmitted: 66705150
           Packets forwarded/queued: 14
           Packets copied: 77831798
           Packets dropped: 33351074
           Services-offload packets processed: 16574037502
           Fragment packets: 0
           Pre fragments generated: 0
           Post fragments generated: 0
```

show security flow statistics

```
user@host:LSYS1> show security flow statistics
Current sessions: 1000
           Packets received: 177888012
           Packets transmitted: 66705150
           Packets forwarded/queued: 14
           Packets copied: 77831798
           Packets dropped: 33351074
           Services-offload packets processed: 16574037502
```

```
            Fragment packets: 0
            Pre fragments generated: 0
            Post fragments generated: 0
```

show security flow statistics tenant TSYS1

```
user@host> show security flow statistics tenant TSYS1
Current sessions: 10
            Packets received: 28711
            Packets transmitted: 10
            Packets forwarded/queued: 0
            Packets copied: 9590
            Packets dropped: 19541
            Services-offload packets processed: 95909
            Fragment packets: 0
            Pre fragments generated: 0
            Post fragments generated: 0
```

show security flow statistics

```
user@host:TSYS1> show security flow statistics
Current sessions: 10
            Packets received: 28711
            Packets transmitted: 10
            Packets forwarded/queued: 0
            Packets copied: 9590
            Packets dropped: 19541
            Services-offload packets processed: 95909
            Fragment packets: 0
            Pre fragments generated: 0
            Post fragments generated: 0
```

Release Information

Command introduced in Junos OS Release 10.2.

Fragmentation counters options introduced in Junos OS Release 15.1X49-90.

Support added at the logical system and tenant system level in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security flow status

IN THIS SECTION

- Syntax | 24348
- Description | 24348
- Required Privilege Level | 24348
- Output Fields | 24348
- Sample Output | 24350
- Release Information | 24356

Syntax

```
show security flow status
```

Description

Display the flow processing modes and logging status.

Required Privilege Level

view

Output Fields

Table 1205 on page 24349 lists the output fields for the `show security flow status` command. Output fields are listed in the approximate order in which they appear.

Table 1205: show security flow status Output Fields

Field Name	Field Description
Flow forwarding mode	Flow processing mode. • Inet forwarding mode • Inet6 forwarding mode • MPLS forwarding mode • ISO forwarding mode • Session distribution mode • Enhanced route scaling mode • Tap mode: enabled, inspect ipip
Flow trace status	Flow logging status. • Flow tracing status • Flow tracing options
flow session distribution	SPU load distribution mode. • RR-based • Hash-based GTP-U distribution • Enabled
Flow packet ordering	packet-ordering mode. • Hardware • Software
Flow ipsec performance acceleration	IPsec VPN performance acceleration status.

Table 1205: show security flow status Output Fields *(Continued)*

Field Name	Field Description
Flow power mode IPsec	Flow PowerMode IPsec status.
Flow Enhanced Service Mode	Flow Enhanced Service Mode status • Disabled • Enabled • Disabled (reboot needed to enable Enhanced Service Mode) • Enabled (reboot needed to disable Enhanced Service Mode)
Flow gre performance acceleration	GRE performance acceleration status.

Sample Output

show security flow status

```
user@host> show security flow status
  Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: flow based
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
    Tap mode: disabled (default)
    Enhanced services mode: Disabled
Flow trace status
    Flow tracing status: on
    Flow tracing options: all
    Flow tracing level: detail
    Flow trace root-override status: on
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Disabled
  Flow ipsec performance acceleration: off
```

```
  Flow gre performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware
  Flow power mode IPsec: Disabled
  Flow power mode IPsec QAT: Disabled
  Fat core group status: off
  Flow inline fpga crypto: Disabled
```

```
node1:
--------------------------------------------------------------------------
  Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: flow based
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
    Tap mode: disabled (default)
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: RR-based
    GTP-U distribution: Disabled
    SCTP distribution: Enabled
  Flow ipsec performance acceleration: off
  Flow gre performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware
  Flow power mode: Enabled
  Flow power mode IPsec: Enabled
  Flow power mode IPsec QAT: Disabled
  Fat core group status: off
  Flow inline fpga crypto: Disabled
```

Logical system or tenant system user output.

```
Flow trace status
Flow tracing status: on
Flow tracing options: all
```

```
Flow tracing level: detail
Flow trace root-override status: on
```

```
user@host> show security flow status
Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: drop
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
    TAP mode: enabled
   Flow trace status
   Flow tracing status: off
   Flow session distribution
    Distribution mode: RR-based
    GTP-U distribution: Enabled Flow packet ordering
    Ordering mode: Software (reboot needed to change to software)
   Flow power mode IPsec: Enabled
```

show security flow status (IPsec Performance Acceleration)

```
user@host> show security flow status
node0:
--------------------------------------------------------------------------
   Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: flow based
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
    Tap mode: disabled (default)
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Disabled
  Flow ipsec performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware
  Flow power mode IPsec: Enabled
  Fat core group status: on
node1:
```

```
-------------------------------------------------------------------------
    Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: flow based
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
    Tap mode: disabled (default)
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Disabled
  Flow ipsec performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware
  Flow power mode IPsec: Enabled
  Fat core group status: on
```

show security flow status (for hash-based datapath forwarding using SRX5K-MPC3-40G10G (IOC3) and SRX5K-MPC3-100G10G (IOC3)

```
user@host> show security flow status
  Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: flow based
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
    Tap mode: disabled (default)
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Disabled
  Flow ipsec performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware
  Flow power mode IPsec: Disabled
  Fat core group status: on
```

show security flow status (SRX5400, SRX5600, SRX5600, and vSRX Virtual Firewall)

```
user@host> show security flow status
 Flow forwarding mode:
Inet forwarding mode: flow based
Inet6 forwarding mode: flow based
MPLS forwarding mode: drop
ISO forwarding mode: drop
Tap mode: enabled, inspect ipip
…
```

show security flow status (Tap mode enabled with IP-IP or GRE tunnel inspection on SRX300, SRX320, SRX340, SRX345, SRX1500, SRX4100, and SRX4200)

```
user@host> show security flow status
 Flow forwarding mode:
    Inet forwarding mode: flow based
….
  Flow packet ordering
    Ordering mode: Hardware
  Flow power mode IPsec: Disabled
 Flow Enhanced Service Mode: Enabled
…
```

show security flow status (Enhanced Service Mode enabled on vSRX Virtual Firewall 3.0)

```
user@host> show security flow status
Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: flow based
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
    Tap mode: disabled (default)
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Disabled
  Flow ipsec performance acceleration: off
```

```
  Flow packet ordering
    Ordering mode: Hardware
  Flow power mode IPsec: Disabled
  Flow power mode IPsec QAT: Enabled
  Fat core group status: off
  Flow inline fpga crypto: Disabled
```

show security flow status (GRE Performance Acceleration)

```
user@host> show security flow status
Flow forwarding mode:
Inet forwarding mode: flow based
Inet6 forwarding mode: drop
MPLS forwarding mode: drop
ISO forwarding mode: drop
TAP mode: enabled
Flow trace status
Flow tracing status: off
Flow session distribution
Distribution mode: RR-based
GTP-U distribution: Enabled Flow packet ordering
Ordering mode: Software (reboot needed to change to software)
Flow power mode IPsec: Disabled
Flow gre performance acceleration: on
```

show security flow status (PowerMode)

```
user@host> show security flow status
node0:
--------------------------------------------------------------------------
  Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: flow based
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
    Tap mode: disabled (default)
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
```

```
    GTP-U distribution: Disabled
    SCTP distribution: Enabled
  Flow ipsec performance acceleration: off
  Flow gre performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware
  Flow power mode: Enabled
  Flow power mode IPsec: Enabled
  Flow power mode IPsec QAT: Disabled
  Fat core group status: on
  Flow inline fpga crypto: Disabled
```

show security flow status (PowerMode IPsec QAT)

show security flow status (MX-SPC3 services card)

```
user@host> show security flow status
Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: flow based
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
    Tap mode: disabled (default)
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Disabled
  Flow ipsec performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware
  Flow power mode IPsec: Enabled
  Fat core group status: on
```

Release Information

Command introduced in Junos OS Release 10.2; session distribution mode option added in Junos OS Release 12.1X44-D10; enhanced route scaling mode option added in Junos OS Release 12.1X45-D10. GTP-U distribution option added in Junos OS Release 15.1X49-D40.

Starting in Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) are introduced for SRX5400, SRX5600, and SRX5800 devices that perform hash-based data path packet forwarding to interconnect with all existing IOC and SPC cards using the XL chip (packet-processing chip).

The IOC3 XL chip uses a hash-based method to distribute ingress traffic to a pool of SPUs by default. Selection of hash keys depends on application protocols.

Starting in Junos OS Release 18.3R1, flow PowerMode IPsec support is introduced on vSRX Virtual Firewall instances.

Starting in Junos OS Release 18.4R1, flow PowerMode IPsec support is introduced on SRX4100 and SRX 4200 devices.

Starting in Junos OS Release 18.2R2, flow PowerMode IPsec support is introduced on SRX5400, SRX5600, and SRX5800 devices.

Starting in Junos OS Release 19.1R1, flow PowerMode IPsec support is introduced on SRX4600 devices.

Starting in Junos OS Release 19.4R1, symmetric fat tunnel support is introduced on SRX5400, SRX5600, SRX5800, and vSRX Virtual Firewall instances.

Starting in Junos OS Release 20.3R1, Enhanced Service Mode is supported on vSRX Virtual Firewall 3.0. In this mode of service vSRX Virtual Firewall 3.0 can support upto 128K L7 service sessions with increased service memory. By default, ESM is disabled and the vSRX Virtual Firewall 3.0 is in basic firewall mode. You can enable ESM using the `set security forwarding-process enhanced-services-mode` command. After enabling this mode, you need to reboot the instance. When you enable this configuration, you will receive a warning message **warning: You have changed enhanced services mode. You must reboot the system for your change to take effect. If you have deployed a cluster, be sure to reboot all nodes**.

Starting in Junos OS Release 20.4R1, the `jexec`, `minimal`, and `root-override` support is introduced on SRX1500, SRX4100, SRX4200, SRX4600, SRX5000 line of devices, and vSRX Virtual Firewall. The options `route`, `session-scan`, and `tunnel` at `[edit security flow traceoptions flag]` hierarchy are configured only by root user.

Starting in Junos OS Release 21.1R1, we support the PMI and GRE acceleration solutions to improve the software-defined WAN (SD-WAN) performance.

For PMI solution, include the `power-mode-ipsec` and `gre-performance-acceleration` statements at the `[edit security flow]` hierarchy level. PMI supports both IPsec and GRE. In this case, traffic flows through the PMI data path.

For GRE acceleration solution, include the `gre-performance-acceleration` statement at the `[edit security flow]` hierarchy level. By default, `gre-performance-acceleration` is turned off. In this case, traffic flows through the GRE acceleration data path.

Starting in Junos OS Release 21.2R1, we support clear text traffic over generic routing encapsulation (GRE) tunnels and MPLS-over-GRE tunnels.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security flow tunnel inspection statistics

IN THIS SECTION

- Syntax | 24358
- Description | 24358
- Required Privilege Level | 24358
- Output Fields | 24359
- Sample Output | 24360
- Release Information | 24360

Syntax

```
show security flow tunnel-inspection statistics
```

Description

Required Privilege Level

Output Fields

Table 1206: show security flow tunnel-inspection statistics Output Fields

Field Name	Field Description
Tunnel inspection type	Tunnel inspection type. Displays VXLAN for VXLAN tunnel traffic
overlay session active	Number of active overlay sessions
overlay session create	Number of overlay sessions created
overlay session close	Number of overlay sessions closed
underlay session active	Number of active underlay sessions
underlay session create	Number of underlay sessions created
underlay session close	Number of underlay sessions closed
input packets	Total number of input packets
input bytes	Total number of input Bytes
output packets	Total number of output packets
output bytes	Total number of output bytes
bypass packets	Number of packets bypassing security policy
bypass bytes	Number of packets bypassing security policy

Sample Output

show security flow tunnel-inspection statistics

```
user@host> show security flow tunnel-inspection statistics
```

```
Flow Tunnel-inspection statistics:
  Tunnel-inspection type VXLAN:
    overlay session active:             0
    overlay session create:             275
    overlay session close:              275
    underlay session active:            0
    underlay session create:            615
    underlay session close:             615
    input packets:            0
    input bytes:              0
    output packets:           0
    output bytes:             0
    bypass packets:           0
    bypass bytes:             0
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Tunnel Inspection for EVPN-VXLAN on SRX Series Devices

show security forward-options resource-manager

IN THIS SECTION

Syntax | 24361

- Description | 24361
- Options | 24361
- Required Privilege Level | 24361
- Output Fields | 24362
- Sample Output | 24362
- Release Information | 24363

Syntax

```
show security forward-options resource-manager [settings | <node (node-id | all | local |
primary)> | status]
```

Description

Display forward option status and the CPU and memory allocated for the advance services and to verify the CPU allocation between routing engine and flow RT threads.

Options

- *setting* —Display resource manager settings
- node—(Optional) For chassis cluster configurations, display resource manager settings on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - all—Display information about all nodes.
  - local—Display information about the local node.
  - primary—Display information about the primary node.
- status—Verify the CPU allocation between routing engine and flow RT threads.

Required Privilege Level

view

Output Fields

Table 1207 on page 24362 lists the output fields for the `show security forward-options resource-manager settings` command. Output fields are listed in the approximate order in which they appear.

Table 1207: show security forward-options resource-manager settings Output Fields

Field Name	Field Description
`Owner`	Displays owner of the setting. Usually it is the service that requests the resource. The service utilizing resources such as Content Security, IDP, SWRSS, SWRSS-IO, NICRSS_Q, or RE.
`Type`	Displays resource type. This can be either CPU core number, CPU core percentage, regular memory, service memory, heap memory, SWRSS mode or value.
`Current settings` and `Next settings`	Displays the current effective and next boot effective value. Value "-" indicates that the setting is not effective or takes default value.

Sample Output

show security forward-options resource-manager status

```
user@host> show security forward-options resource-manager status
Resource manager details:
vCPU number for IO threads: 0
vCPU number for flow threads: 1
```

show security forward-options resource-manager settings

```
user@host> show security forward-options resource-manager settings
Owner         Type                    Current settings      Next settings
RE            CPU core number         1                     1
```

show security forward-options resource-manager status (SRX4600)

```
user@host> show security forward-options resource-manager status
Resource manager details:
vCPU number for IO threads: 0
vCPU number for flow threads: 23
Dedicated resources for logging: Enabled
```

Release Information

Command introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

Traffic Processing on SRX Series Devices Overview

show security forwarding-options mirror-filter

IN THIS SECTION

- Syntax | 24363
- Description | 24364
- Options | 24364
- Required Privilege Level | 24364
- Output Fields | 24365
- Sample Output | 24366
- Release Information | 24366

Syntax

```
show security forwarding-options mirror-filter (all | filter-name)
```

Description

Displays status information about all configured mirror filters or that of a specific mirror filter. Each mirror filter contains a set of parameters against which traffic is matched. For each mirror filter, the output identifies the number of packets that were matched by the filter for mirroring and the number of packets that were sent to the packet analyzer. It also shows the parameters that were configured for the mirror filter.

Network operators need a way to monitor X2 traffic to debug any handover issues across eNodeBs. The mirror filter feature allows you to do that. To use the mirror filter feature to monitor X2 traffic, you configure mirror filters. Traffic coming out of an IPsec tunnel is decrypted, mirrored, and analyzed by a packet analyzer, and then encrypted again to go into the outbound IPsec tunnel.

> **NOTE**: The SRX Series mirror filter feature is bidirectional, much like a session. X2 traffic flowing through an IPSec VPN from devices that match the configured filter conditions is mirrored and analyzed.

Starting in Junos OS Release 18.4R1, if the output X2 interface of a mirror filter is configured for an st0 interface to filter traffic that you want to analyze, the packet is duplicated and encrypted by the IPsec tunnel bound to the st0 interface. This enhancement supports the SRX Series Firewalls to send traffic mirrored from a port on an IPsec tunnel.

You can configure up to 15 different mirror filters to be used concurrently.

> **NOTE**: Although there is no minimum required number of parameters for a mirror filter, please be mindful that if you specify too few criteria or accidentally commit an incomplete filter, an over-proportional amount of traffic flow through the system could be mirrored.

Options

all	Display counters for all mirror filters.
filter-name	Name of the mirror filter for which the counters are displayed.

Required Privilege Level

view

Output Fields

Lists the output fields for the `show security forward-options mirror-filter` command. Output fields are listed in the approximate order in which they appear in the output.

Table 1208: show security forward-options mirror-filter

Field Name	Field Description
`mirror-filter-name`	Name of the mirror filter configured on the device.
`interface-in`	Name of the incoming logical interface to be matched for mirroring.
`interface-out`	Name of the outgoing logical interface to be matched for mirroring.
`protocol`	Networking protocol name or number to be matched for mirroring.
`source-port`	Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port number to be matched for mirroring.
`source-prefix`	Source IP prefix or address to be matched for mirroring.
`destination-port`	Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number to be matched for mirroring.
`destination-prefix`	Destination IP prefix or address to be matched for mirroring.
`filter-counters`	Number of packets matched for mirroring.
`output-counter`	Number of packets sent to the packet analyzer.

Sample Output

show security forward-options mirror-filter

```
user@host> show security forward-options mirror-filter traffic-https

  Security mirror status

    mirror-filter-name: traffic-https
    interface-in: st0.1
    interface-out: st0.2
    destination-port: 443
    protocol: 132
    source-prefix: 192.0.2.0
    destination-prefix: 203.0.113.0
    filter-counters: 2
    output-counters: 2
```

Release Information

Command introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

mirror-filter (Security Forwarding Options) | 10768

clear security forward-options mirror filter | 3935

show security group-vpn member ike security-associations

IN THIS SECTION

- Syntax | 24367
- Description | 24367
- Options | 24367

- Required Privilege Level | 24367
- Output Fields | 24368
- Sample Output | 24370
- Sample Output | 24371
- Release Information | 24371

Syntax

```
show security group-vpn member ike security-associations [brief | detail] [index sa-index] [peer-ipaddress]
```

Description

Display IKE security associations (SAs) for group members. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Options

- none—Display summary information about all IKE SAs for the group members.
- `brief`—(Optional) Display summary output.
- `detail`—(Optional) Display detailed output.
- `index`*sa-index*—(Optional) Display detailed information about the specified SA identified by index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.
- *peer-ipaddress*—(Optional) Display information about the SA with the specified peer.

Required Privilege Level

view

Output Fields

Table 1209 on page 24368 lists the output fields for the `show security group-vpn member ike security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 1209: show security group-vpn member ike security-associations Output Fields

Field Name	Field Description
`Index`	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
`State`	State of the IKE security associations: • `DOWN`—SA has not been negotiated with the peer. • `UP`—SA has been negotiated with the peer.
`Initiator cookie`	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.
`Responder cookie`	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received. A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.
`Mode`	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between themselves. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are • `main`—The exchange is done with six messages. This mode or exchange type encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • `aggressive`—The exchange is done with three messages. This mode or exchange type does not encrypt the payload, leaving the identity of the neighbor unprotected.
`Remote Address`	IP address of the destination peer with which the local peer communicates.

Table 1209: show security group-vpn member ike security-associations Output Fields *(Continued)*

Field Name	Field Description
`IKE Peer`	IP address of the destination peer with which the local peer communicates.
`Exchange type`	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between themselves. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are • `main`—The exchange is done with six messages. This mode or exchange type encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • `aggressive`—The exchange is done with three messages. This mode or exchange type does not encrypt the payload, leaving the identity of the neighbor unprotected.
`Authentication method`	Method the server uses to authenticate the source of IKE messages: • `pre-shared-keys`—Preshared key for encryption and decryption that both participants must have before beginning tunnel negotiations.
`Local`	Address of the local peer.
`Lifetime`	Number of seconds remaining until the IKE SA expires.

Table 1209: show security group-vpn member ike security-associations Output Fields *(Continued)*

Field Name	Field Description

`Algorithms`	Internet Key Exchange (IKE) algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:
	• `Authentication`—Type of authentication algorithm used.
	  • `sha-256`—Secure Hash Algorithm 256 authentication.
	  • `sha-384`—Secure Hash Algorithm 384 authentication.
	• `Encryption`—Type of encryption algorithm used.
	  • `aes-256-cbc`—Advanced Encryption Standard (AES) 256-bit encryption.
	  • `aes-192-cbc`— AES192-bit encryption
	  • `aes-128-cbc`—AES 128-bit encryption.

`Traffic statistics`	• `Input bytes`–Number of bytes received. • `Output bytes`–Number of bytes transmitted. • `Input packets`–Number of packets received. • `Output packets`–Number of packets transmitted.

Sample Output

show security group-vpn member ike security-associations

```
user@host> show security group-vpn member ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address
4736345 UP     70611c65603d53da  6e0888777ad10f8d  Main           192.0.2.3
```

Sample Output

show security group-vpn member ike security-associations detail

```
user@host> show security group-vpn member ike security-associations detail
IKE peer 192.0.2.5, Index 5824842, Gateway Name: group1_2
  Role: Initiator, State: UP
  Initiator cookie: fc866556b8afe4cd, Responder cookie: 1238de6b8a89de44
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 192.0.2.7:848, Remote: 192.0.2.5:848
  Lifetime: Expires in 2 seconds
  Peer ike-id: 192.0.2.5
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-2
  Traffic statistics:
   Input  bytes  :                 2044
   Output bytes  :                  900
   Input  packets:                    7
   Output packets:                    7
  Flags: IKE SA is created
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

clear security group-vpn member ike security-associations | 3938

Group VPNv2 Overview

show security group-vpn member ipsec inactive-tunnels

IN THIS SECTION

- Syntax | 24372
- Description | 24372
- Options | 24372
- Required Privilege Level | 24372
- Output Fields | 24373
- Sample Output | 24374
- Release Information | 24375

Syntax

```
show security group-vpn member ipsec inactive-tunnels <brief> <detail> <group-id group-id>
```

Description

Show inactive Group VPNs. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Options

none	Display information for all groups.
brief	(Optional) Display summary output.
detail	(Optional) Display detailed output.
group-id *group-id*	(Optional) Display information for the specified group identifier.

Required Privilege Level

view

Output Fields

Table 1210 on page 24373 lists the output fields for the `show security group-vpn member ipsec inactive-tunnels` command. Output fields are listed in the approximate order in which they appear.

Table 1210: show security group-vpn member ipsec inactive-tunnels Output Fields

Field Name	Field Description
Server	Server on which group member is registered.
Port	UDP port number.
GId	Group identifier.
lsys	Logical system.
Reason	Reason that the tunnel is inactive: • The tunnel was cleared through the CLI. • The hard lifetime has expired. • There are too many TEKs. • There was a configuration change. • There was an SA installation error. • The TEK is stale. • The tunnel was deleted from the server.
Virtual-system	Logical system name.
Group VPN Name	Name of the Group VPN.
Local Gateway	IP address of the local IKE gateway.
GDOI Server	IP address of the group server.

Table 1210: show security group-vpn member ipsec inactive-tunnels Output Fields *(Continued)*

Field Name	Field Description
Group Id	Group identifier.
Recovery Probe	Status of the recovery probe, either enabled or disabled (default).
DF-bit	Fragmentation of IPsec traffic on the group member—clear (default), copy, or set.
Stats	Statistics for GDOI groupkey-pull and groupkey-push exchanges, server failovers, deletes received, number of times the maximum number of keys and policies were exceeded, and the number of unsupported algorithms received.
Down Reason	Reason that the tunnel is inactive: • The tunnel was cleared through the CLI. • The hard lifetime has expired. • There are too many TEKs. • There was a configuration change. • There was an SA installation error. • The TEK is stale. • The tunnel was deleted from the server. • The tunnel is not initiated.

Sample Output

show security group-vpn member ipsec inactive-tunnels

```
user@host> show security group-vpn member ipsec inactive-tunnels
  Total inactive tunnels: 1
  Server           Port  GId lsys  Reason
  192.168.1.50       848   1000 root uninitiated
```

show security group-vpn member ipsec inactive-tunnels detail

```
user@host> show security group-vpn member ipsec inactive-tunnels detail
  Virtual-system: root Group VPN Name: group1000
  Local Gateway: 192.168.1.101, GDOI Server: 192.168.1.50
  Group Id: 1000
  Recovery Probe: Disabled
  DF-bit: clear
  Stats:
      Pull Succeeded             :   0
      Pull Failed                :   8841
      Pull Timeout               :   7996
      Pull Aborted               :   0
      Push Succeeded             :   0
      Push Failed                :   0
      Server Failover            :   0
      Delete Received            :   0
      Exceed Maximum Keys(4)     :   0
      Exceed Maximum Policies(10):   0
      Unsupported Algo           :   0
  Down Reason: uninitiated
```

Release Information

Command introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Group VPNv2 Overview

show security group-vpn member ipsec security-associations

IN THIS SECTION

Syntax | 24376

- Description | 24376
- Options | 24376
- Required Privilege Level | 24376
- Output Fields | 24377
- Sample Output | 24379
- Sample Output | 24380
- Release Information | 24381

Syntax

```
show security group-vpn member ipsec security-associations [brief | detail] [index sa-index]
```

Description

Display group VPN security associations (SAs) for a group member. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Options

- none—Display information about all group VPN SAs for the group member.
- `brief`—(Optional) Display summary output.
- `detail`—(Optional) Display detailed output.
- `index` *sa-index*—(Optional) Display detailed information about the specified SA identified by index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.

Required Privilege Level

view

Output Fields

Table 1211 on page 24377 lists the output fields for the `show security group-vpn member ipsec security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 1211: show security group-vpn member ipsec security-associations

Field Name	Field Description
`Total active tunnels`	Total number of active IPsec tunnels.
`ID`	Index number of the SA. You can use this number to get additional information about the SA.
`Server`	IP address of the group server (remote gateway).
`Port`	If Network Address Translation-Traversal (NAT-T) is used, this value is 4500. Otherwise it is the standard IKE port, 500.
`Algorithm`	Cryptography used to secure exchanges between peers during the IKE Phase 2 negotiations includes • An authentication algorithm used to authenticate exchanges between the peers. Options are `sha-256` or `sha-384` • An encryption algorithm used to encrypt data traffic. Options are `aes-128`, `aes-192`, and `aes-256`.
`SPI`	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI.
`Life: sec/kb`	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.
`GId`	Group identifier.
`vsys or Virtual-system`	The root system.

Table 1211: show security group-vpn member ipsec security-associations *(Continued)*

Field Name	Field Description
`Local Gateway`	Gateway address of the local system.
`GDOI Server`	IP address of the group server.
`Local Identity`	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IPv4 address, fully qualified domain name, e-mail address, or distinguished name.
`Remote Identity`	IPv4 address of the destination peer gateway.
`DF-bit`	State of the don't fragment bit: set or cleared.
Forward-policy-mismatch	Enable the support for forwarding policy-mismatched packets
`Policy name`	Name of the applicable policy.
`Direction`	Direction of the security association; it can be inbound or outbound.
`AUX-SPI`	Value of the auxiliary security parameter index. • When the value is AH or ESP, AUX-SPI is always 0. • When the value is AH+ESP, AUX-SPI is always a positive integer.
`Hard lifetime`	The hard lifetime specifies the lifetime of the SA. • `Expires in seconds`—Number of seconds left until the SA expires.
`Lifesize Remaining`	The lifesize remaining specifies the usage limits in kilobytes. If there is no lifesize specified, it shows unlimited. • `Expires in kilobytes`—Number of kilobytes left until the SA expires.

Table 1211: show security group-vpn member ipsec security-associations *(Continued)*

Field Name	Field Description
`Soft lifetime`	The soft lifetime informs the IPsec key management system that the SA is about to expire. Each lifetime of a security association has two display options, hard and soft, one of which must be present for a dynamic security association. This allows the key management system to negotiate a new SA before the hard lifetime expires. • `Expires in seconds`—Number of seconds left until the SA expires.
`Mode`	Mode of the security association: • transport—Protects host-to-host connections. • tunnel—Protects connections between security gateways.
`Protocol`	Protocol supported. Transport mode supports Encapsulation Security Protocol (ESP).
`Anti-replay service`	State of the service that prevents packets from being replayed. It can be `Enabled` or `Disabled`.

Sample Output

show security group-vpn member ipsec security-associations

```
user@host> show security group-vpn member ipsec security-associations
  Total active tunnels: 2
  ID    Server             Port  Algorithm       SPI      Life:sec/kb  GId lsys
  <>49157 192.168.1.53     848   ESP:3des/sha1   c0792f86 114/  unlim   2000 root
  <>49156 192.168.1.53     848   ESP:aes-256/md5 7def169d 18/   unlim   2000 root
  <>49156 192.168.1.53     848   ESP:aes-256/md5 86c48448 146/  unlim   2000 root
```

Sample Output

show security group-vpn member ipsec security-associations detail

```
user@host> show security group-vpn member ipsec security-associations detail
  Virtual-system: root Group VPN Name: group2000
  Local Gateway: 192.168.1.70, GDOI Server: 192.168.1.53
  Group Id: 2000
  Routing Instance: vr1
  Recovery Probe: Enabled
  DF-bit: clear
Forward-policy-mismatch:Enabled

  Stats:
      Pull Succeeded            :   3
      Pull Failed               :   0
      Pull Timeout              :   6
      Pull Aborted              :   0
      Push Succeeded            :   1773
      Push Failed               :   0
      Server Failover           :   0
      Delete Received           :   0
      Exceed Maximum Keys(4)    :   0
      Exceed Maximum Policies(10):   0
      Unsupported Algo          :   0
  Flags:
      Rekey Needed:   no

    List of policies received from server:
    Tunnel-id: 49157
      Source IP: ipv4_subnet(any:900,[0..7]=192.168.1.0/24)
      Destination IP: ipv4_subnet(any:901,[0..7]=192.168.1.0/24)

      Direction: bi-directional, SPI: c0792f86
      Protocol: ESP, Authentication: sha1, Encryption: 3des
      Hard lifetime: Expires in 81 seconds, Activated
      Lifesize Remaining:  Unlimited
      Soft lifetime: Expired
      Mode: Tunnel, Type: Group VPN, State: installed
      Anti-replay service: D3P enabled, Window size: 3000 milliseconds

      Direction: bi-directional, SPI: a645b381
```

```
        Protocol: ESP, Authentication: sha1, Encryption: 3des
        Hard lifetime: Expires in 207 seconds, Activated in 51 seconds
        Lifesize Remaining:  Unlimited
        Soft lifetime: Expires in 117 seconds
        Mode: Tunnel, Type: Group VPN, State: installed
        Anti-replay service: D3P enabled, Window size: 3000 milliseconds
```

Release Information

Command introduced in Junos OS Release 10.2.

Command introduced in Junos OS Release 18.2R1 for MX-series.

RELATED DOCUMENTATION

clear security group-vpn member ipsec security-associations | 3939

Group VPNv2 Overview

show security group-vpn member ipsec statistics

IN THIS SECTION

- Syntax | 24382
- Description | 24382
- Options | 24382
- Required Privilege Level | 24382
- Output Fields | 24382
- Sample Output | 24383
- Release Information | 24384

Syntax

```
show security group-vpn member ipsec statistics <index index>
```

Description

Show IPsec statistics. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Options

none Display information for all IPsec SAs.

index *index* (Optional) Display detailed information about the specified SA, identified by index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.

Required Privilege Level

view

Output Fields

Table 1212 on page 24382 lists the output fields for the `show security group-vpn member ipsec statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1212: show security group-vpn member ipsec statistics Output Fields

Field Name	Field Description
ESP Statistics	Numbers of encrypted and decrypted bytes and encrypted and decrypted packets.
AH Statistics	Numbers of input and output bytes and input and output packets.
Errors	Numbers of AH failures, replay errors, ESP authentication failures, ESP decryption failures, bad headers, and bad trailers.

Table 1212: show security group-vpn member ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
D3P Statistics	Numbers of old timestamp packets, new timestamp packets, no timestamp packets, unexpected D3P header packets, invalid type packets, invalid length packets, and invalid next header packets.
Exclude Statistics	Numbers of created and invalidated sessions.
Dynamic Policy Statistics	Numbers of created and invalidated sessions.
Fail-Open Statistics	Numbers of created and invalidated sessions.
Fail-Close Statistics	Number of dropped packets.
Forward-policy-mismatch Statistics	Number of bypassed packets.

Sample Output

show security group-vpn member ipsec statistics

```
user@host> show security group-vpn member ipsec statistics
ESP Statistics:
  Encrypted bytes:            54712
  Decrypted bytes:            16800
  Encrypted packets:            381
  Decrypted packets:            200
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
```

```
  Bad headers: 0, Bad trailers: 0
D3P Statistics:
  Old timestamp packets:                  0
  New timestamp packets:                  0
  No timestamp packets:                   0
  Unexpected D3P header packets:          0
  Invalid type packets:                   0
  Invalid length packets:                 0
  Invalid next header packets:            0
Exclude Statistics:
  Created sessions:                       0
  Invalidated sessions:                   0
Dynamic Policy Statistics:
  Created sessions:                     381
  Invalidated sessions:                   0
Fail-Open Statistics:
  Created sessions:                       0
  Invalidated sessions:                   0
Fail-Close Statistics:
  Dropped packets:                        0
Forward-policy-mismatch Statistics:
  Input Packets:                  0
  Output packets:                 0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D30.

Command introduced in Junos OS Release 18.2R1 for MX-series.

RELATED DOCUMENTATION

Group VPNv2 Overview

show security group-vpn member kek security-associations

IN THIS SECTION

- Syntax | 24385
- Description | 24385
- Options | 24385
- Required Privilege Level | 24386
- Output Fields | 24386
- Sample Output | 24389
- Sample Output | 24389
- Release Information | 24391

Syntax

```
show security group-vpn member kek security-associations [brief | detail | display xml] [index
sa-index] [peer-ipaddress]
```

Description

Display Group VPNv2 security associations (SAs) for a group member. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Group VPNv2 is the name of the Group VPN technology on MX5, MX10, MX40, MX80, MX240, MX480, and MX960 routers. Group VPNv2 is different from the Group VPN technology implemented on SRX Security Gateways.

For more information about Group VPN on SRX Security Gateway devices, see *Group VPNv2 Overview*.

Options

- none—Display information about all Group VPNv2 SAs for the group member.
- `brief`—(Optional) Display summary output.

- `detail`—(Optional) Display detailed output.
- `display xml`—(Optional) Display xml.
- `index` *sa-index*—(Optional) Display detailed information about the specified SA identified by index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.
- *peer-ipaddress*—(Optional) Display information about the SA with the specified peer.

Required Privilege Level

view

Output Fields

Table 1213 on page 24386 lists the output fields for the `show security group-vpn member kek security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 1213: show security group-vpn member kek security-associations

Field Name	Field Description
`Index`	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
`Remote Address`	IP address of the destination peer with which the local peer communicates.
`State`	State of the KEK security associations: • `DOWN`—SA is not active. • `UP`—SA is active.
`Initiator cookie`	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.
`Responder cookie`	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received.

Table 1213: show security group-vpn member kek security-associations *(Continued)*

Field Name	Field Description
SPI	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI.
GroupID	Group identifier.
KEK Peer	IP address of the destination peer with which the local peer communicates.
Role	For the member, it is always responder.
State	State of the KEK security associations, which is always up.
Authentication method	RSA is the supported authentication method.
Local	Address of the local peer.
Remote	Address of the remote peer.
Lifetime	Number of seconds remaining until the IKE SA expires.

Table 1213: show security group-vpn member kek security-associations *(Continued)*

Field Name	Field Description

`Algorithms`	Internet Key Exchange (IKE) algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:
	• `Sig-hash`—Type of authentication algorithm used.
	  • `sha-256`–Secure Hash Algorithm 256 (sha-256) authentication.
	  • `sha-384`–Secure Hash Algorithm 394 (sha-384) authentication.
	• `Sig key length (bits)`—Size of signature key in bits.
	• `Encryption`—Type of encryption algorithm used.
	  • `aes-256-cbc`—Advanced Encryption Standard (AES) 256-bit encryption.
	  • `aes-192-cbc`— AES192-bit encryption
	  • `aes-128-cbc`—AES 128-bit encryption.
	  • `3des-cbc`—3 Data Encryption Standard (DES) encryption.
	  • `des-cbc`—DES encryption.

`Traffic statistics`	• `Input bytes`–Number of bytes received. • `Output bytes`–Number of bytes transmitted. • `Input packets`–Number of packets received. • `Output packets`–Number of packets transmitted.
`Server Info Version`	Identify the latest set of information maintained in the server.
`Server Heartbeat Interval`	Interval in seconds at which the server sends heartbeats to group members.
`Member Heartbeat Threshold`	The heartbeat threshold configured on the group member for the IPsec VPN. If this number of heartbeats is missed on the member, the member reregisters with the server.

Table 1213: show security group-vpn member kek security-associations *(Continued)*

Field Name	Field Description
`Heartbeat Timeout Left`	Number of heartbeats until the heartbeat threshold is reached, at which time the member reregisters with the server. When this number reaches 0, reregistration happens within 60 seconds.
`Server Activation Delay`	Number of seconds before a group member can use a new key when the member reregisters with the server.
`Server Multicast Group`	Multicast IP address to which the server sends rekey messages.
`Server Replay Window`	Antireplay time window value in milliseconds. 0 means antireplay is disabled.
`Group Key Push sequence number`	Sequence number of the KEK SA groupkey-push message. This number is incremented with every groupkey-push message.

Sample Output

show security group-vpn member kek security-associations

```
user@host> show security group-vpn member kek security-associations
Index   Server Address  Life:sec  Initiator cookie  Responder cookie  GroupId
5824843 192.168.2.53      166       46871e26227f08f3  f0a463a4d5c3737b  1
```

Sample Output

show security group-vpn member kek security-associations detail

```
user@host> show security group-vpn member kek security-associations detail
  Index 5824843, Group Id: 1
  Group VPN Name: group1_2
  Local Gateway: 192.168.2.170, GDOI Server: 192.168.2.53
  Initiator cookie: 46871e26227f08f3, Responder cookie: f0a463a4d5c3737b
  Lifetime: Expires in 155 seconds
```

```
  Group Key Push Sequence number: 0

  Algorithms:
   Sig-hash              : hmac-md5-96
   Encryption            : 3des-cbc
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Stats:
      Push received            :   0
      Delete received          :   0
```

show security group-vpn member kek security-associations detail | display xml

```
user@host> show security group-vpn member kek security-associations detail | display xml

<rpc-reply xmlns:junos="http://xml.example.net/junos/15.1/junos">
    <gvpn-kek-security-associations-information junos:style="detail">
        <kek-security-associations-block>
            <security-association-index>2987691</security-association-index>
            <group-id>400</group-id>
            <group-vpn-name>gvpn400</group-vpn-name>
            <local-address>192.168.1.100</local-address>
            <server-address>192.168.1.1</server-address>
            <initiator-cookie>510f854307a03675</initiator-cookie>
            <responder-cookie>690e5f121fba6de7</responder-cookie>
            <lifetime-remaining>Expires in 23729 seconds</lifetime-remaining>
            <push-sequence-number>364</push-sequence-number>
            <ike-security-associations>
                <ike-sa-algorithms>
                    <ike-sa-authentication-algorithm>hmac-sha1-96</ike-sa-authentication-
algorithm>
                    <ike-sa-sig-key-length>2048</ike-sa-sig-key-length>
                    <ike-sa-encryption-algorithm>aes128-cbc</ike-sa-encryption-algorithm>
                </ike-sa-algorithms>
                <ike-sa-traffic-statistics>
                    <ike-sa-input-bytes>3012</ike-sa-input-bytes>
                    <ike-sa-output-bytes>252</ike-sa-output-bytes>
                    <ike-sa-input-packets>3</ike-sa-input-packets>
```

```
                    <ike-sa-output-packets>3</ike-sa-output-packets>
                </ike-sa-traffic-statistics>
            </ike-security-associations>
            <gvpn-kek-security-association-statistics>
                <kek-security-association-statistics>    Push received           :   3</kek-
security-association-statistics>
                <kek-security-association-statistics>    Delete received         :   0</kek-
security-association-statistics>
            </gvpn-kek-security-association-statistics>
        </kek-security-associations-block>
    </gvpn-kek-security-associations-information>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Group VPNv2 Overview

show security group-vpn member policy

IN THIS SECTION

- Syntax | 24392
- Description | 24392
- Options | 24392
- Required Privilege Level | 24392
- Output Fields | 24392
- Sample Output | 24393
- Release Information | 24394

Syntax

```
show security group-vpn member policy <vpn vpn-name> <group-id group-id>
```

Description

Show Group VPN policies. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Options

none	Display information for all groups.
vpn *vpn-name*	(Optional) Display policy information for the specified group name.
group-id *group-id*	(Optional) Display policy information for the specified group identifier.

Required Privilege Level

view

Output Fields

Table 1214 on page 24392 lists the output fields for the `show security group-vpn member policy` command. Output fields are listed in the approximate order in which they appear.

Table 1214: show security group-vpn member policy Output Fields

Field Name	Field Description
Group VPN Name	Group name.
Group Id	Group identifier.
From-zone	From zone configured for the policy.

Table 1214: show security group-vpn member policy Output Fields *(Continued)*

Field Name	Field Description
To-zone	To zone configured for the policy.
Tunnel-id	Tunnel identifier.
Policy type	Secure, fail-open, fail-close, or exclude.
Source	IP address, port, and protocol of the source traffic.
Destination	IP address, port, and protocol of the destination traffic.

Sample Output

show security group-vpn member policy

```
user@host> show security group-vpn member policy
Group VPN Name: group1000, Group Id: 1000
From-zone: trust_1, To-zone: untrust
  Tunnel-id: 63490, Policy type: Exclude
    Source      : IP <192.168.0.0 - 192.168.255.255>, Port <0 - 65535>, Protocol <17>
    Destination : IP <0.0.0.0 - 255.255.255.255>, Port <0 - 65535>, Protocol <17>

  Tunnel-id: 49153, Policy type: Secure
    Source      : IP 192.168.0.0 - 192.168.255.255>, Port <0 - 65535>, Protocol <0>
    Destination : IP <192.0.2.0 - 192.0.2.255>, Port <0 - 65535>, Protocol <0>

  Tunnel-id: 49152, Policy type: Secure
    Source      : IP <192.0.2.0 - 192.0.2.255>, Port <0 - 65535>, Protocol <1>
    Destination : IP <192.0.2.0 - 192.0.2.255>, Port <0 - 65535>, Protocol <1>

  Tunnel-id: 63491, Policy type: Fail-open (Inactivated)
    Source      : IP 192.168.0.0 - 192.168.255.255>, Port <0 - 65535>, Protocol <17>
    Destination : IP <192.168.0.0 - 192.168.255.255>, Port <0 - 65535>, Protocol <17>

  Tunnel-id: 63489, Policy type: Fail-close
```

```
    Source      : IP <0.0.0.0 - 255.255.255.255>, Port <0 - 65535>, Protocol <0>
Destination : IP <0.0.0.0 - 255.255.255.255>, Port <0 - 65535>, Protocol <0>
```

Release Information

Command introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Group VPNv2 Overview

show security group-vpn server ike security-associations

IN THIS SECTION

- Syntax | 24394
- Description | 24395
- Options | 24395
- Required Privilege Level | 24395
- Output Fields | 24395
- Sample Output | 24399
- Sample Output | 24400
- Release Information | 24400

Syntax

```
show security group-vpn server ike security-associations [brief | detail] [group group-name |
group-id group-id] [index sa-index]
```

Description

Display IKE security associations (SAs). Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Options

- none—Display all IKE SAs for all groups.
- brief—(Optional) Display summary output.
- `detail`—(Optional) Display detailed level of output.
- `group`—(Optional) Display IKE SAs for the specified group.
- `group-id`—(Optional) Display IKE SAs for the specified group.

  An IKE SA can be used by a group member to register to multiple groups. When you specify the `group` or `group-id` options to list the IKE SAs for a specified group, all existing IKE SAs that could be used to register to the group are displayed.
- `index`—(Optional) Display information for a particular SA based on the index number of the SA. To obtain the index number for a particular SA, display the list of existing SAs by using the command with no options.

Required Privilege Level

view

Output Fields

Table 1215 on page 24395 lists the output fields for the `show security group-vpn server ike security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 1215: show security group-vpn server ike security-associations Output Fields

Field Name	Field Description
`Index`	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.

Table 1215: show security group-vpn server ike security-associations Output Fields *(Continued)*

Field Name	Field Description
`Remote Address`	IP address of the destination peer with which the local peer communicates.
`State`	State of the IKE security associations: • `DOWN`—SA has not been negotiated with the peer. • `UP`—SA has been negotiated with the peer.
`Initiator cookie`	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.
`Responder cookie`	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received. A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.
`Mode`	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between themselves. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are • `main`—The exchange is done with six messages. This mode or exchange type encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • `aggressive`—The exchange is done with three messages. This mode or exchange type does not encrypt the payload, leaving the identity of the neighbor unprotected.
`IKE Peer`	IP address of the destination peer with which the local peer communicates.

Table 1215: show security group-vpn server ike security-associations Output Fields *(Continued)*

Field Name	Field Description
`Exchange type`	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between themselves. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are • `main`—The exchange is done with six messages. This mode or exchange type encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • `aggressive`—The exchange is done with three messages. This mode or exchange type does not encrypt the payload, leaving the identity of the neighbor unprotected.
`Authentication method`	Method the server uses to authenticate the source of IKE messages: • `pre-shared-keys`—Preshared key for encryption and decryption that both participants must have before beginning tunnel negotiations. • `rsa-signatures`—Digital signature, a certificate that confirms the identity of the certificate holder.
`Local`	Address of the local peer.
`Remote`	Address of the remote peer.
`Lifetime`	Number of seconds remaining until the IKE SA expires.

Table 1215: show security group-vpn server ike security-associations Output Fields *(Continued)*

Field Name	Field Description

`Algorithms`	Internet Key Exchange (IKE) algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:
	• `Authentication`—Type of authentication algorithm used.
	  • `sha-256`—Secure Hash Algorithm 256 authentication.
	  • `sha-384`—Secure Hash Algorithm 384 authentication..
	• `Encryption`—Type of encryption algorithm used.
	  • `aes-256-cbc`—Advanced Encryption Standard (AES) 256-bit encryption.
	  • `aes-192-cbc`— AES192-bit encryption
	  • `aes-128-cbc`—AES 128-bit encryption.

`Traffic statistics`	• `Input bytes`–Number of bytes received. • `Output bytes`–Number of bytes transmitted. • `Input packets`–Number of packets received. • `Output packets`–Number of packets transmitted.
`IPSec security associations`	• *`number`* `created`: The number of SAs created. • *`number`* `deleted`: The number of SAs deleted.

Table 1215: show security group-vpn server ike security-associations Output Fields *(Continued)*

Field Name	Field Description

`Phase 2 negotiations in progress`	Number of Phase 2 IKE negotiations in progress and status information:
	• `Negotiation type`—Type of Phase 2 negotiation. Junos OS currently supports quick mode.
	• `Message ID`—Unique identifier for a Phase 2 negotiation.
	• `Local identity`—Identity of the local Phase 2 negotiation. The format is id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)
	• `Remote identity`—Identity of the remote Phase 2 negotiation. The format is id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)
	• `Flags`—Notification to the key management process of the status of the IKE negotiation:
	  • `caller notification sent`—Caller program notified about the completion of the IKE negotiation.
	  • `waiting for done`—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.
	  • `waiting for remove`—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.
	  • `waiting for policy manager`—Negotiation is waiting for a response from the policy manager.

Sample Output

show security group-vpn server ike security-associations

```
user@host> show security group-vpn server ike security-associations
     Index   State  Initiator cookie  Responder cookie  Mode          Remote Address
     738879  UP     0fa7c5fdcb74669f  8c21f5d1b533010c  Aggressive    192.168.1.120
```

Sample Output

show security group-vpn server ike security-associations detail

```
user@host> show security group-vpn server ike security-associations detail
IKE peer 192.168.1.120, Index 738879, Gateway Name: gvpn
  Role: Responder, State: UP
  Initiator cookie: 0fa7c5fdcb74669f, Responder cookie: 8c21f5d1b533010c
  Exchange type: Aggressive, Authentication method: Pre-shared-keys
  Local: 192.168.1.50:848, Remote: 192.168.1.120:848
  Lifetime: Expires in 3541 seconds
  Peer ike-id: test
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha-256-128
   Encryption            : aes-256-cbc
   Pseudo random function: hmac-sha-256
   Diffie-Hellman group  : DH-group-14
  Traffic statistics:
   Input  bytes  :                  600
   Output bytes  :                  932
   Input  packets:                    4
   Output packets:                    3
  Flags: IKE SA is created
  IPSec security associations: 0 created, 0 deleted
  Phase 2 negotiations in progress: 0

Flags: IKE SA is created
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security group-vpn member ike security-associations | 24366

Group VPNv2 Overview

show security group-vpn server ipsec security-associations

IN THIS SECTION

- Syntax | 24401
- Description | 24401
- Options | 24401
- Required Privilege Level | 24402
- Output Fields | 24402
- Sample Output | 24403
- Sample Output | 24403
- Release Information | 24404

Syntax

```
show security group-vpn server ipsec security-associations [brief | detail] [group group-name |
group-id group-id]
```

Description

Display IPsec security associations (SAs). Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Options

- none—Display all IPsec SAs for all groups.
- `brief`—(Optional) Display summary output.
- `detail`—(Optional) Display detailed level of output.
- `group`—(Optional) Display IPsec SAs for the specified group.
- `group-id`—(Optional) Display IPsec SAs for the specified group.

Required Privilege Level

view

Output Fields

Table 1216 on page 24402 lists the output fields for the `show security group-vpn server ipsec security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 1216: show security group-vpn server ipsec security-associations

Field Name	Field Description
`Group`	Group name.
`Group ID`	Group identifier.
`Total IPsec SAs`	The total number of IPsec SAs for each group is shown.
`IPsec SA`	Name of the SA.
`Protocol`	Protocol supported. Transport mode supports Encapsulation Security Protocol (ESP).
`Algorithm`	Cryptography used to secure exchanges between peers during the IKE Phase 2 negotiations includes • An authentication algorithm used to authenticate exchanges between the peers. Options are `sha-256` and `sha-384`. • An encryption algorithm used to encrypt data traffic. Options are `aes-128-cbc`, `aes-192-cbc`, or `aes-256-cbc`.
`SPI`	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI.
`Lifetime`	The lifetime of the SA, after which it expires, expressed in seconds.

Table 1216: show security group-vpn server ipsec security-associations *(Continued)*

Field Name	Field Description
Policy Name	Group policy associated with the IPsec SA. The source address, destination address, source port, destination port, and protocol defined for the policy are displayed.

Sample Output

show security group-vpn server ipsec security-associations

```
user@host> show security group-vpn server ipsec security-associations
     Group: group200, Group Id: 200
       Total IPsec SAs: 1
       IPsec SA          Algorithm             SPI              Lifetime
       sa1               ESP:aes-256/sha-256    55837dfe         17
       sa1               ESP:aes-256/sha1-256   760088d          137
```

Sample Output

show security group-vpn server ipsec security-associations detail

```
user@host> show security group-vpn server ipsec security-associations detail
Group: group1, Group Id: 1
Total IPsec SAs: 10
  IPsec SA: sa1
    Protocol: ESP, Authentication: sha-256, Encryption: aes-256
    Anti-replay: D3P enabled, window size 10 milliseconds
    SPI: e68c9525
    Lifetime: Expires in 66 seconds, Activated
    Policy Name: pol1
      Source: 192.168.1.0/24
      Destination: 192.168.1.0/24
      Source Port: 0
      Destination Port: 0
      Protocol: 0
  IPsec SA: sa1
```

```
    Protocol: ESP, Authentication: sha-256, Encryption: aes-256
    Anti-replay: D3P enabled, window size 10 milliseconds
    SPI: 7ee14902
    Lifetime: Expires in 276 seconds, Activated in 36 seconds
    Rekey in 186 seconds
    Policy Name: pol1
      Source: 192.168.1.0/24
      Destination: 192.168.1.0/24
      Source Port: 0
      Destination Port: 0
      Protocol: 0
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security group-vpn member ipsec security-associations | 24375

Group VPNv2 Overview

show security group-vpn server kek security-associations

IN THIS SECTION

- Syntax | 24405
- Description | 24405
- Options | 24405
- Required Privilege Level | 24405
- Output Fields | 24405
- Sample Output | 24408
- Sample Output | 24408
- Release Information | 24409

Syntax

```
show security group-vpn server kek security-associations [brief | detail] [group group-name |
group-id group-id | index sa-index]
```

Description

Display configured server-member communications. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Options

- none—Display server-member communications configured for all groups.
- `brief`—(Optional) Display summary output.
- `detail`—(Optional) Display detailed output.
- `group`—(Optional) Display server-member communications configured for the specified group.
- `group-id`—(Optional) Display server-member communications configured for the specified group.
- `index`—(Optional) Display information for a particular SA based on the index number of the SA. To obtain the index number for a particular SA, display the list of existing SAs by using the command with no options.

Required Privilege Level

view

Output Fields

Table 1217 on page 24406 lists the output fields for the `show security group-vpn server kek security-assocations` command. Output fields are listed in the approximate order in which they appear.

Table 1217: show security group-vpn server kek security-associations Output Fields

Field Name	Field Description
`Index`	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
`Remote Address`	Identifier of the remote/peer. Because there could be multiple members, the remote address always contains the IP address 0.0.0.0.
`State`	State of the KEK security associations: • `DOWN`—SA is not active. • `UP`—SA is active.
`Initiator cookie`	Random number generated by the server. This is used when the server needs to push data to a member, or a member needs to reply to the server.
`Responder cookie`	Random number generated by the server. This is used when the server needs to push data to a member, or a member needs to reply to the server.
`GroupId`	Group identifier.
`KEK Peer`	IP address of the destination peer with which the local peer communicates. For KEK SAs, it always contains 0.0.0.0 which means any IP address.
`Role`	For the server, it is always initiator.
`Authentication method`	RSA is the supported authentication method.
`Local`	Address of the local peer.
`Remote`	Address of the remote peer.
`Lifetime`	Number of seconds remaining until the IKE SA expires.

Table 1217: show security group-vpn server kek security-associations Output Fields *(Continued)*

Field Name	Field Description

`Algorithms`	Internet Key Exchange (IKE) algorithms used to encrypt and secure exchanges between the peers during the Phase 2 process:
	• `Sig-hash`—Type of authentication algorithm used.
	  • `sha-256`—Secure Hash Algorithm 256 authentication.
	  • `sha-384`—Secure Hash Algorithm 384 authentication.
	• `Encryption`—Type of encryption algorithm used.
	  • `aes-256-cbc`—Advanced Encryption Standard (AES) 256-bit encryption.
	  • `aes-192-cbc`— AES192-bit encryption
	  • `aes-128-cbc`—AES 128-bit encryption.

`Traffic statistics`	• `Input bytes`–Number of bytes received. • `Output bytes`–Number of bytes transmitted. • `Input packets`–Number of packets received. • `Output packets`–Number of packets transmitted.
`Server Info Version`	Identify the latest set of information maintained in the server.
The following fields are the configured `server-member-communication` options:	
`Server Replay Window`	Antireplay time in milliseconds. This is 0 if antireplay is disabled.
`Retransmission Period`	Number of seconds between a rekey transmission and the first retransmission when there is no reply from the member.
`Number of Retransmissions`	For unicast communications, the number of times the server retransmits rekey messages to a member when there is no reply.

Table 1217: show security group-vpn server kek security-associations Output Fields *(Continued)*

Field Name	Field Description
`Lifetime Seconds`	Configured lifetime, in seconds, for the KEK.
`Group Key Push sequence number`	Sequence number of the KEK SA groupkey-push message. This number is incremented with every groupkey-push message.

Sample Output

show security group-vpn server kek security-associations

```
user@host> show security group-vpn server kek security-associations
Index   Life:sec  Initiator cookie  Responder cookie  GroupId
     739031  18995     7e17278bf0a65975  0616de443d1beb77  200
```

Sample Output

show security group-vpn server kek security-associations detail

```
user@host> show security group-vpn server kek security-associations detail
Index 738879, Group Name: GROUP_ID-0001, Group Id: 1
Initiator cookie: 114e4a214891e42f, Responder cookie: 4b2848d14372e5bd
Authentication method: RSA
Lifetime: Expires in 4186 seconds, Activated
Rekey in 3614 seconds
  Algorithms:
   Sig-hash              : sha256
   Encryption            : aes256-cbc
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Server Member Communication: Unicast
  Retransmission Period: 10, Number of Retransmissions: 2
```

```
  Group Key Push sequence number: 0

PUSH negotiations in progress: 0
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security group-vpn member kek security-associations | 24385

Group VPNv2 Overview

show security group-vpn server registered-members

IN THIS SECTION

- Syntax | 24409
- Description | 24410
- Options | 24410
- Required Privilege Level | 24410
- Output Fields | 24410
- Sample Output | 24411
- Sample Output | 24411
- Release Information | 24412

Syntax

```
show security group-vpn server registered-members <group group-name> <group-id group-id> <detail>
```

Description

Display currently registered group members. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Options

- none—Display all group members for all groups.
- brief—(Optional) Display summary output.
- detail—(Optional) Display detailed output.
- `group`—(Optional) Display group members for the specified group.
- `group-id`—(Optional) Display group members for the specified group.

Required Privilege Level

view

Output Fields

Table 1218 on page 24410 lists the output fields for the `show security group-vpn server registered-members`command. Output fields are listed in the approximate order in which they appear.

Table 1218: show security group—vpn server registered-members Output Fields

Field Name	Field Description
`Group`	Group name.
`Group Id`	Group identifier.
`Member Gateway`	IP address of the gateway for the group member.
`Member IP`	IP address of the group member.

Table 1218: show security group—vpn server registered-members Output Fields *(Continued)*

Field Name	Field Description
Last Update	The last time that members registered or sent acknowledgements to the server.
Vsys	The root system.

Sample Output

show security group-vpn server registered-members

```
user@host> show security group-vpn server registered-members
    Group: group200, Group Id: 200
       Total number of registered members: 1
       Member Gateway                   Member IP         Last Update              Vsys
       gvpn_simpleman                   192.168.1.100     Fri Dec 20 2013 07:27:33 root
```

Sample Output

show security group-vpn server registered-members detail

```
user@host> show security group-vpn server registered-members detail
Group: group1, Group Id: 1
         Total number of registered members: 1

  Member gateway: gateway_group1_1, Member IP: 192.168.1.2, Vsys: root
  Last Update: Fri May 16 2014 03:37:17
  Stats:
      Pull Succeeded                : 321
      Pull Failed                   : 0
      Push Sent                     : 0
      Push Acknowledged             : 0
      Push Unacknowledged           : 0
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

clear security group-vpn server | 3944

Group VPNv2 Overview

show security group-vpn server server-cluster

IN THIS SECTION

- Syntax | 24412
- Description | 24412
- Options | 24413
- Required Privilege Level | 24413
- Output Fields | 24413
- Sample Output | 24414
- Release Information | 24416

Syntax

```
show security group-vpn server server-cluster <brief> <detail> <group group-name> <group-id group-id> <peer-
gateway gateway-name>
```

Description

Show information about servers in the Group VPNv2 server cluster. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Options

none	Display Group VPNv2 server cluster information for all groups.
brief	(Optional) Display summary output.
detail	(Optional) Display detailed output, including information about exchanges with peer servers in the cluster.
group ***group-name***	(Optional) Display Group VPNv2 server cluster information for the specified group name.
group-id ***group-id***	(Optional) Display Group VPNv2 server cluster information for the specified group identifier.
peer-gateway ***gateway-name***	(Optional) Display Group VPNv2 server cluster information for the specified peer.

Required Privilege Level

view

Output Fields

Table 1219 on page 24413 lists the output fields for the `show security group-vpn server server-cluster` command. Output fields are listed in the approximate order in which they appear.

Table 1219: show security group-vpn server server-cluster Output Fields

Field Name	Field Description
Group	Group name.
Group Id	Group identifier.
Role	Role of this server in the Group VPNv2 server cluster.

Table 1219: show security group-vpn server server-cluster Output Fields *(Continued)*

Field Name	Field Description
Version Number	32-bit version number included in `cluster-update` exchanges and DPD probes to support anti-replay. The first `cluster-update` message sent from the root-server has version number 1. Subsequent `cluster-update` messages increment the version number by one. (Retransmit messages do not increment the version number.) Upon receipt of a `cluster-update` message, the sub-server validates the received version number. The received version number must be greater than the version number in the last received message, otherwise the message is discarded. The sub-server responds to a `cluster-update` message with an ACK message that contains the same version number as the received message. Upon receipt of the ACK message, the root-server checks that the version number is the same as in the message it sent. If the version number is valid, the exchange is considered successful. If the version number is not valid, the original message is retransmitted or the exchange is considered failed.
Peer Gateway	Name of the peer server in the Group VPNv2 server cluster.
Peer IP	IP address of the remote peer server in the Group VPNv2 server cluster.
Role	Role of the peer server in the Group VPNv2 server cluster.
Status	Status of the peer server in the Group VPNv2 server cluster.

Sample Output

show security group-vpn server server-cluster

```
user@host> show security group-vpn server server-cluster
Group: group200, Group Id: 200
Role: Root-server, Version Number: 1,
  Peer Gateway                    Peer IP          Role                    Status
  sub_server1                     192.168.1.112    Sub-server              Active
  sub_server2                     192.168.1.113    Sub-server              Active
```

show security group-vpn server server-cluster detail

```
user@host> show security group-vpn server server-cluster detail
GGroup: group200, Group Id: 200
Role: Root-server, Version Number: 1,

Peer gateway: sub_server1,
  Peer IP: 192.168.1.112, Local IP: 192.168.1.111, VR: vr1,
  Role: Sub-server, Status: Active,
  CLUSTER-INIT send:                 0
  CLUSTER-INIT recv:                 1
  CLUSTER-INIT success:              1
  CLUSTER-INIT fail:                 0
  CLUSTER-INIT dup:                  0
  CLUSTER-INIT abort:                0
  CLUSTER-INIT timeout:              0
  CLUSTER-UPDATE send:               1
  CLUSTER-UPDATE recv:               0
  CLUSTER-UPDATE success:            1
  CLUSTER-UPDATE fail:               0
  CLUSTER-UPDATE abort:              0
  CLUSTER-UPDATE timeout:            0
  CLUSTER-UPDATE pending:            0
  CLUSTER-UPDATE max retry reached:  0
  DPD send:                          5
  DPD send fail:                     0
  DPD ACK recv:                      5
  DPD ACK invalid seqno:             0
  IPsec SA policy mismatch:          0
  IPsec SA proposal mismatch:        0
  KEK SA proposal mismatch:          0

Peer gateway: sub_server2,
  Peer IP: 192.168.1.113, Local IP: 192.168.1.111, VR: default,
  Role: Sub-server, Status: Active,
  CLUSTER-INIT send:                 0
  CLUSTER-INIT recv:                 1
  CLUSTER-INIT success:              1
  CLUSTER-INIT fail:                 0
  CLUSTER-INIT dup:                  0
  CLUSTER-INIT abort:                0
  CLUSTER-INIT timeout:              0
```

```
CLUSTER-UPDATE send:                1
CLUSTER-UPDATE recv:                0
CLUSTER-UPDATE success:             1
CLUSTER-UPDATE fail:                0
CLUSTER-UPDATE abort:               0
CLUSTER-UPDATE timeout:             0
CLUSTER-UPDATE pending:             0
CLUSTER-UPDATE max retry reached:   0
DPD send:                           6
DPD send fail:                      0
DPD ACK recv:                       6
DPD ACK invalid seqno:              0
IPsec SA policy mismatch:           0
IPsec SA proposal mismatch:         0
KEK SA proposal mismatch:           0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Group VPNv2 Overview

Understanding Group VPNv2 Server Clusters

show security group-vpn server statistics

IN THIS SECTION

- Syntax | 24417
- Description | 24417
- Options | 24417
- Required Privilege Level | 24417
- Output Fields | 24417
- Sample Output | 24418

Release Information | 24418

Syntax

```
show security group-vpn server statistics <group group-name> <group-id group-id>
```

Description

Show Group VPNv2 server statistics. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 Series Firewalls and vSRX Virtual Firewall instances.

Options

none	Display Group VPNv2 server statistics for all groups.
group *group-name*	(Optional) Display Group VPNv2 server statistics for the specified group name.
group-id *group-id*	(Optional) Display Group VPNv2 server statistics for the specified group identifier.

Required Privilege Level

view

Output Fields

Table 1220 on page 24417 lists the output fields for the `show security group-vpn server statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1220: show security group-vpn server statistics Output Fields

Field Name	Field Description
Group	Group name.

Table 1220: show security group-vpn server statistics Output Fields *(Continued)*

Field Name	Field Description
Group Id	Group identifier.
Stats	Server events and number of occurrences.

Sample Output

show security group-vpn server statistics

```
user@host> show security group-vpn server statistics
Group: group1, Group Id: 1
  Stats:
      Pull Succeeded                : 321
      Pull Failed                   : 0
      Pull Exceed Member Threshold  : 0
      Push Sent                     : 0
      Push Acknowledged             : 0
      Push Unacknowledged           : 0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Group VPNv2 Overview

Understanding Group VPNv2 Server Clusters

show security gtp

IN THIS SECTION

- Syntax | 24419
- Description | 24419
- Options | 24419
- Required Privilege Level | 24419
- Sample Output | 24419

Syntax

```
show security gtp apn-control {<apn-ctrl-name> | all}
```

Description

Use this command to display the APN profile configuration.

Options

- aon-control-name —Displays specific APN profile.
- all —Displays all the APN profile list.

Required Privilege Level

view

Sample Output

Refer the following sample outputs for Junos OS Release 21.2R1.

show security gtp apn-control all (GTP Profile)

```
user@host> show security gtp apn-control all
GTP Profile List (id, name):
            1 apn1
            2 apn2

root> show security gtp apn-control apn1
APN: IoT-Greece uid: 1
Rate-Limit:
Alarm Threshold     : 200
Drop Threshold       : 400
GTPv0 Message    : 16, 18, 20, 22, 24
GTPv1 Message    : 16, 20
GTPv2 Message    : 32, 34, 36, 95, 99
```

show security gtp counters

IN THIS SECTION

- Syntax | 24420
- Description | 24421
- Options | 24421
- Required Privilege Level | 24421
- Output Fields | 24421
- Sample Output | 24424
- Release Information | 24432

Syntax

```
show security gtp counters <all | data-packet | error | ha | message <message-name> | packet |
request | tunnel | path-rate-limit | rate-limit>
```

Description

Use this command to display counters that can be used to indicate the number of GPRS tunneling protocol (GTP) tunnel counters (allocated and freed), GTP packet counters (received, passed, and dropped), brief message counters (receive, forward, and drop), error counters, request counters, HA counters, and path-rate-limit counters (drop and alarm).

From Junos OS Release 20.4R1 onwards, `show security gprs gtp counters` is replaced by `show security gtp counters`.

Options

- `all`—Displays all GTP counters.
- `data-packet`— Displays GTP-U data packet counters.
- `error`—Displays GTP error counters.
- `ha`—Displays GTP HA counters.
- `message` *`message-name`*—Displays GTP message counters.
- `packet`—Displays GTP packet counters.
- `request`—Displays GTP request counters.
- `tunnel`—Displays GTP tunnel counters.
- `path-rate-limit`—Displays path-rate-limit counters.
- `rate-limit`—Displays rate-limit counters.

Required Privilege Level

view

Output Fields

Table 1221 on page 24422 lists the output fields for the `show security gtp counters` command.

Table 1221: show security gtp counters all Output Fields

Field Name	Field Description
`Tunnel counters`	Tunnel counters are used to track the number of tunnels that are created on the device. There are two entries: • Allocated • Freed Active tunnel number = number of allocated counters - number of freed counters
`Packet counters`	Packet counters indicate the number of GTP packets that are received and processed on the device. There are three entries: • Received–Number of GTP packet messages received. • Passed–Number of GTP packet messages passed. • Dropped–Number of GTP packet messages dropped because of an error. Number of received counters = number of dropped counters + number of passed counters
`Data-packet counters`	GTP-U data packet counters are used to track the number of data packets that are received, passed, dropped, no-tunnel dropped, sequence error dropped, end-user dropped, or dropped otherwise.

Table 1221: show security gtp counters all Output Fields *(Continued)*

Field Name	Field Description
`Brief message counters`	GTP messages counters indicate the number of GTP messages that are received and processed on the device. There are three entries: • Received–Number of GTP messages received. • Forwarded–Number of GTP messages forwarded. • Dropped–Number of GTP messages dropped because of an error. Number of received counters = number of dropped counters + number of forward counters
`Error counters`	Drop reason and drop counters indicate the number of GTP packets that are dropped as a result of an error. Total error = Sum of all the following errors
`Request counters`	Request counters indicate the number of GTP request messages that are received and processed on the device. This information can be used for debugging purpose.
`HA counters`	HA counters indicate the number of messages that are received or sent by the device.

Table 1221: show security gtp counters all Output Fields *(Continued)*

Field Name	Field Description
`Path-rate-limit counters`	Path-rate-limit counters indicate the number of packet data protocol (PDP) create, delete, echo, and other messages that are received and processed on the device after drop-threshold and alarm-threshold are reached. • Create Request–Number of create PDP messages. • Delete Request–Number of delete PDP messages. • Echo Request–Number of PDP echo messages. • Others–Control messages other than the above messages. Drop–Indicate the number of packets dropped. Alarm–Indicate the number of packets transferred after the alarm threshold is reached.

Table 1221: show security gtp counters all Output Fields *(Continued)*

rate-limit counters	Rate-limit counters indicate the number of drop and alarm packets. • Drop indicates the number of messages dropped after reaching the drop threshold. • Alarm indicates the number of messages went through after reaching the alarm threshold.

Sample Output

Refer the following sample outputs for Junos OS 20.3R1 or previous releases.

show security gtp counters all

```
user@host> show security gtp counters all
Tunnel counters:
                 Total      GTPv0    GTPv1-c    GTPv1-u    GTPv2-c    GTPv2-u
Allocated      0          0         0         0          0           0
Freed          0          0         0         0          0           0
```

```
Packet counters:
                 Total      GTPv0      GTPv1      GTPv2       GTP'
Received       0          0          0          0           0
Passed         0          0          0          0           0
Dropped        0          0          0          0           0

Data-packet counters:
        gtpv1 data pkt received               : 0
        gtpv1 data pkt passed                 : 0
        gtpv1 data pkt dropped                : 0
        gtpv1 data pkt no-tunnel dropped      : 0
        gtpv1 data pkt sequence error dropped : 0
        gtpv1 data pkt end-user dropped       : 0
        gtpv1 data pkt other dropped          : 0
        gtpv0 data pkt received               : 0
        gtpv0 data pkt passed                 : 0
        gtpv0 data pkt dropped                : 0
        gtpv0 data pkt no-tunnel dropped      : 0
        gtpv0 data pkt sequence error dropped : 0
        gtpv0 data pkt end-user dropped       : 0
        gtpv0 data pkt other dropped          : 0


Brief message counters:
                                    Received    Forwarded       Dropped
GTPv0
  Create PDP Request                0           0             0
  Create PDP Response               0           0             0
  Update PDP Request                0           0             0
  Update PDP Response               0           0             0
  Delete PDP Request                0           0             0
  Delete PDP Response               0           0             0
  Create AA PDP Request             0           0             0
  Create AA PDP Response            0           0             0
  Delete AA PDP Request             0           0             0
  Delete AA PDP Response            0           0             0
  SGSN Context Request              0           0             0
  SGSN Context Response             0           0             0
  SGSN Context Acknowledge          0           0             0
  Others                            0           0             0
GTPv1
  Create PDP Request                0           0             0
```

```
  Create PDP Response               0          0             0
  Update PDP Request                0          0             0
  Update PDP Response               0          0             0
  Delete PDP Request                0          0             0
  Delete PDP Response               0          0             0
  SGSN Context Request              0          0             0
  SGSN Context Response             0          0             0
  SGSN Context Acknowledge          0          0             0
  Forward Relocation Request        0          0             0
  Forward Relocation Response       0          0             0
  Others                            0          0             0
GTPv2
  Create Session Request            0          0             0
  Create Session Response           0          0             0
  Delete Session Request            0          0             0
  Delete Session Response           0          0             0
  Create Bearer Request             0          0             0
  Create Bearer Response            0          0             0
  Modify Bearer Request             0          0             0
  Modify Bearer Response            0          0             0
  Delete Bearer Request             0          0             0
  Delete Bearer Response            0          0             0
  Context Request                   0          0             0
  Context Response                  0          0             0
  Context Acknowledge               0          0             0
  Forward Relocation Request        0          0             0
  Forward Relocation Response       0          0             0
  Create Indirect Tunnel Request    0          0             0
  Create Indirect Tunnel Response   0          0             0
  Others                            0          0             0

Error counters:
        Total error                  : 0
        Exception                    : 0
        Invalid header               : 0
        Message length               : 0
        Zero IMSI                    : 0
        Zero charge ID               : 0
        Sequence                     : 0
        APN filter                   : 0
        Port not match               : 0
        GTP-in-GTP                   : 0
        Message too short            : 0
```

```
        Message too long              : 0
        GSN not exist                 : 0
        Over GSN rate limit           : 0
        Request not found             : 0
        Retransmit response           : 0
        Missing IE                    : 0
        Unexpected IE                 : 0
        Unknown IE type               : 0
        IE order                      : 0
        IE length                     : 0
        Duplicate IE                  : 0
        Non-digit TID/TEID            : 0
        Non-zero TID/TEID             : 0
        Zero TID/TEID                 : 0
        Control TID/TEID              : 0
        Data TID/TEID                 : 0
        Control GSN IE                : 0
        Data GSN IE                   : 0
        End user IE                   : 0
        GGSN IP for handover          : 0
        Disallowed v0 message         : 0
        Disallowed v1 message         : 0
        Disallowed v2 message         : 0
        Invalid message type          : 0
        No tunnel0                    : 0
        No control tunnel             : 0
        No user tunnel                : 0
        Invalid tunnel0               : 0
        Invalid control tunnel        : 0
        Invalid user tunnel           : 0
        Create tunnel0                : 0
        Create control tunnel         : 0
        Create user tunnel            : 0
        No request                    : 0
        Out of request                : 0
        No action                     : 0
        Out of action                 : 0
        GTPv2 TEID not exist          : 0
        GTPv2 Missing TEID            : 0
        GTPv2 Non-zero EBI            : 0
        GTPv2 EBI not found           : 0
        GTPv2 IE context              : 0
        Duplicate new jbuf            : 0
```

```
        Out of jmpi cookie            : 0
        Send jmpi message             : 0
        JMPI target                   : 0
        Out of callback cookie        : 0
        Reinject packet               : 0
        Distribute wrong spu          : 0
        System under reset            : 0
        Source IP                     : 0
        Destination IP                : 0
        Invalid EBI                   : 0
        SPU not all up                : 0
        Interface not support         : 0
        Out of path                   : 0
        Over path rate limit          : 0
        New utnl key                  : 0
        No utnl key                   : 0
        End user                      : 0
        New sync action               : 0
        NSAPI                         : 0
        New conflict sync action      : 0
        Primary u-tunnel id not exist : 0
        Too many same type ie         : 0
        Invalid v2 lbi                : 0
        Remove conflict utnl          : 0
        Conflict primary u-tunnel     : 0
        Ctnl recovery error           : 0
        Link ctnl conflict            : 0
        Link utnl conflict            : 0
        Link tnl0 conflict            : 0
        Error ggsn ip for v0 packet   : 0
        Invalid payload(NAT)          : 0
        Wrong SPU since NULL sinfor   : 0
        Wrong SPU since no match wings: 0
        Wrong SPU since no anchor GSN : 0
        Bad reqeust retransmit        : 0
        Cross group handover deny     : 0
        Handover default deny         : 0

HA counters:
        Total message received              : 0
        Message received success            : 0
        Bad message received                : 0
        Unknown message type received       : 0
```

```
        Unknown message version received    : 0
        Total message send                  : 0
        Message send success                : 0
        Message send failed                 : 0
        Memory allocate failed              : 0
        Message received during tunnel reset: 0

Request counters:
        Request allocated               : 0
        Request freed                   : 0
        Request allocated fail          : 0
        Request hit by wrong SPU        : 0
        Request pending for JMPI ACK    : 0


Path-rate-limit counters:
                            Drop               Alarm
Create Request            0                  0
Delete Request            0                  0
Echo Request              0                  0
Others                    0                  0

Rate-limit counters:
                                     Drop                        Alarm
Packets                               100                          200
```

Refer the following sample output for Junos OS 21.2R1 or later releases.

show security gtp counters error

```
user@host> show security gtp counters error
Error counters:
        Total error                          : 0
        Exception                            : 0
        Invalid header                       : 0
        Message length                       : 0
……
              NE group check failure                : 0
        UE group check failure               : 0
        Over APN-controlled rate limit       : 0
```

show security gtp counters ha

```
user@host> show security gtp counters ha
HA counters:
        Total message received              : 0
        Message received success            : 0
        Bad message received                : 0
        Unknown message type received       : 0
        Unknown message version received    : 0
        Total message send                  : 0
        Message send success                : 0
        Message send failed                 : 0
        Memory allocate failed              : 0
```

show security gtp counters message v0-create-aa-pdp-req

```
user@host> show security gtp counters message v0-create-aa-pdp-req
Message counters:
  Received               0
  Forwarded              0
  Dropped                0
```

show security gtp counters packet

```
user@host> show security gtp counters packet
Packet counters:
                 Total      GTPv0      GTPv1      GTPv2       GTP'
  Received           0          0          0          0          0
  Passed             0          0          0          0          0
  Dropped            0          0          0          0          0
```

show security gtp counters request

```
user@host> show security gtp counters request

Request counters:
        Request allocated           : 0
```

```
        Request freed               : 0
        Request activated           : 0
        Request died                : 0
        Request action allocated    : 0
        Request action freed        : 0
```

show security gtp counters tunnel

```
user@host> show security gtp counters tunnel
Tunnel counters:
                 Total      GTPv0    GTPv1-c    GTPv1-u    GTPv2-c    GTPv2-u
  Allocated          0          0          0          0          0          0
  Freed              0          0          0          0          0          0
```

show security gtp counters path-rate-limit

```
user@host> show security gtp counters path-rate-limit
Path-rate-limit counters:
                            Drop           Alarm
  Create Request             0               0
  Delete Request             0               0
  Echo Request               0               0
  Others                     0               0
```

show security gtp counters rate-limit

```
user@host> show security gtp counters rate-limit
Rate-limit counters:

                                   Drop                        Alarm
Packets                             100                          200
```

Refer the following sample output for Junos OS 20.4R1 or later releases.

Release Information

Command introduced in Junos OS Release 12.1X44-D10. Support for rate-limit option added in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

clear security gtp counters | 3949

show security gtp counters path-rate-limit

IN THIS SECTION

- Syntax | 24432
- Description | 24432
- Required Privilege Level | 24432
- Output Fields | 24433
- Sample Output | 24433
- Release Information | 24434

Syntax

```
show security gtp counters path-rate-limit
```

Description

Use this command to display information about path-rate-limit counters.

Required Privilege Level

view

Output Fields

Table 1222 on page 24433 lists the output fields for the `show security gtp counters path-rate-limit` command.

Table 1222: show security gtp counters path-rate-limit Output Fields

Field Name	Field Description
`Create Request`	Specify the number of create request messages received in a second after the alarm-threshold or drop-threshold is reached.
`Delete Request`	Specify the number of delete request messages received in a second after the alarm-threshold or drop-threshold is reached.
`Echo Request`	Specify the number of echo request messages received in a minute after the alarm-threshold or drop-threshold is reached.
`Other messages`	Specify the number of other GTP control messages received in a second after the alarm-threshold or drop-threshold is reached.
`Drop`	Displays the number of packets dropped after the drop-threshold is reached.
`Alarm`	Displays the number of packets received after the alarm-threshold is reached.

Sample Output

show security gtp counters path-rate-limit

```
user@host> show security gtp counters path-rate-limit

Path-rate-limit counters:
                                    Drop                        Alarm
Create Request                      200                          100
Delete Request                      300                          200
Echo Request                        600                          400
Others                              900                          800
```

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

path-rate-limit | 12263

show security gtp grouped-ie-profile

IN THIS SECTION

- Syntax | 24434
- Description | 24434
- Options | 24434
- Required Privilege Level | 24435
- Sample Output | 24435
- Release Information | 24435

Syntax

```
show security gtp grouped-ie-profile (all | <grpie-prf-name>)
```

Description

Use this command to display the list of all grouped-ie profile and the list of IE number under specific grouped-ie-profile.

Options

all Displays the list of all GTP grouped-ie-profile.

grpie-prf-name Displays a specifc GTP grouped-ie-profile.

Required Privilege Level

View

Sample Output

show security gtp grouped-ie-profile all

```
user@host> show security gtp grouped-ie-profile all
GTP Profile List (id, name):
            1 Bearer-ctxt-crt
            2 Bearer-ctxt-rmv
```

show security gtp grouped-ie-profile Bearer-ctxt-crt

```
user@host> show security gtp grouped-ie-profile Bearer-ctxt-crt
Profile Bearer-ctxt-crt, uid 1
Grouped-IE Number   IE number/Grouped-IE   Instance numbers
93                   73                     0
                     80                     0
                     87                     0
```

show security gtp grouped-ie-profile Bearer-ctxt-rmv

```
user@host> show security gtp grouped-ie-profile Bearer-ctxt-rmv
Profile Bearer-ctxt-rmv, uid 2
Grouped-IE Number   IE number/Grouped-IE   Instance numbers
93                   73                     0
```

Release Information

Command introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Example: Configure Must-IE check for GTPv1 and GTPv2

grouped-ie-profile | 7915

show security gtp gsn statistics

IN THIS SECTION

- Syntax | 24436
- Description | 24436
- Required Privilege Level | 24436
- Sample Output | 24437
- Release Information | 24437

Syntax

```
show security gtp gsn statistics
```

Description

Use this command to display a brief summary of GPRS support node (GSN) statistics, including active GSNs, obsolete GSNs, and the usage rate of each SPU.

Required Privilege Level

view

Sample Output

show security gtp gsn statistics

```
user@host> show security gtp gsn statistics
FPC 1 PIC 0:

Active GSNs: 0 Obsolete GSNs: 0 Use rate: 0%

FPC 2 PIC 0:

Active GSNs: 0 Obsolete GSNs: 0 Use rate: 0%
```

Release Information

Command introduced in Junos OS Release 12.1X46-D25.

RELATED DOCUMENTATION

show security gtp counters path-rate-limit | 24432

show security gtp handover-group

IN THIS SECTION

- Syntax | 24438
- Description | 24438
- Required Privilege Level | 24438
- Sample Output | 24438
- Release Information | 24438

Syntax

```
show security gtp handover-group
```

Description

Use this command to display the handover-group of the GTP profile. A GPRS tunneling protocol (GTP) handover group is a set of SGSNs or serving gateway (SGW) with a common address-book library. An administrator can configure a GTP profile and associate an GTP handover group to the GTP profile. When a GTP handover group name is referenced by a GTP profile, the device checks to see if the current SGSN/SGW address and the proposed SGSN/SGW address are both contained within the same GTP handover group. If both SGSN/SGW addresses are contained within the same GTP handover group, then the handover is allowed. If both the current and proposed SGSN/SGW addresses are not within the same GTP handover group, then the profile for the default handover group is used.

Required Privilege Level

view

Sample Output

show security gtp handover-group

```
user@host> show security gtp handover-group
All configured handover group:
    Group name                  Address book name          Address set name
handover_group_v1v2                addr_book_h              roam_group_v1v2
    hg_v1_h                        addr_book_h                 roam_v1_h
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

show security gtp counters | 24420

handover-group | 7966

show security gtp ie-set

IN THIS SECTION

- Syntax | 24439
- Description | 24439
- Options | 24439
- Required Privilege Level | 24439
- Output Fields | 24439
- Sample Output | 24440
- Release Information | 24440

Syntax

```
show security gtp ie-set (all | <ieset-name>)
```

Description

Use this command to display the list of all the GTP ieset-names and IE numbers.

Options

all	Displays the list of all the GTP ieset-names.
ieset-name	Displays the IE numbers in an ieset.

Required Privilege Level

View

Output Fields

Sample Output

show security gtp ie-set all

```
user@host> show security gtp ie-set all
GTP Profile List (id, name):
            1 ieset-v1-r7
            2 ieset-v2
```

show security gtp ie-set ieset-v1-r7

```
user@host> show security gtp ie-set ieset-v1-r7
Profile ieset-v1-r7, uid 1
IE numbers:
172, 180, 181, 182, 183, 184, 199
```

show security gtp ie-set ieset-v2

```
user@host> show security gtp ie-set ieset-v2
Profile ieset-v2, uid 2
IE numbers:
255
```

Release Information

Command introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

> *Example: Configure IE removal for GTPv1 and GTPv2*
>
> ie-set | 8343

show security gtp ip-group

IN THIS SECTION

- Syntax | 24441
- Description | 24441
- Required Privilege Level | 24441
- Output Fields | 24441
- Sample Output | 24442
- Release Information | 24442

Syntax

```
show security gtp ip-group
```

Description

Use this command to get the details of the configured IP group. IP group is a list of IP addresses that belongs to all kinds of network equipment. IP group name(s) are referenced in GTP profiles. The GTP firewall applies configured policies in incoming and outgoing IP addresses in GPRS tunneling protocol (GTP) messages.

Required Privilege Level

view

Output Fields

Table 1223 on page 24442 lists the output fields for the `show security gtp ip-group` command. Output fields are listed in the approximate order in which they appear.

Table 1223: show security gtp ip-group

Field Name	Field Description
Group name	The name of the handover group on the GTP profile.
Address book name	The name of the common address-book name.
Address set name	The name of the address set for the handover group.

Sample Output

show security gtp ip-group

```
user@host> show security gtp ip-group
All configured IP group:
    Group name              Address book name        Address set name
    ng1                     global                   ne-group-as
    ug1                     global                   ue-group-as
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

gtp | 7931

show security gtp message-ie-profile-v1

IN THIS SECTION

- Syntax | 24443
- Description | 24443
- Options | 24443
- Required Privilege Level | 24443
- Sample Output | 24444
- Release Information | 24444

Syntax

```
show security gtp message-ie-profile-v1 (all | <msgie-prf-v1-name> )
```

Description

Use this command to display the list of all GTPv1 message-ie profile and the list of IE numbers under a specific GTPv1 message-ie profile.

Options

all	Displays the list of all GTPv1 Message-IE profile.
msgie-prf-v1-name	Displays the list of IE numbers defined under a specific GTPv1 Message-IE profile.

Required Privilege Level

View

Sample Output

show security gtp message-ie-profile-v1 all

```
user@host> show security gtp message-ie-profile-v1 all
GTPv1 Message-IE Profile List (id, name):
    1 msgie-v1
```

show security gtp message-ie-profile-v1 msgie-v1

```
user@host> show security gtp message-ie-profile-v1 msgie-v1
Profile msgie-v1, uid 1

Message Number 2
IE numbers:
14

Message Number 16
IE numbers:
2, 3, 16, 17, 20, 133
```

Release Information

Command introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

message-ie-profile-v1 | **10647**

Example: Configure Must-IE check for GTPv1 and GTPv2

show security gtp message-ie-profile-v2

IN THIS SECTION

- Syntax | 24445
- Description | 24445
- Options | 24445
- Required Privilege Level | 24445
- Output Fields | 24445
- Sample Output | 24446
- Release Information | 24446

Syntax

```
show security gtp message-ie-profile-v2 ( all | <msgie-prf-v2-name>)
```

Description

Use this command to display the list of all GTPv2 message-ie profile and the list of IE numbers under a specific GTPv2 message-ie profile.

Options

all	Displays the list of all GTPv2 Message-IE profile.
msgie-prf-v2-name	Displays information of a specific GTPv2 Message-IE profile.

Required Privilege Level

View

Output Fields

Sample Output

show security gtp message-ie-profile-v2 all

```
user@host> show security gtp message-ie-profile-v2 all
GTPv2 Message-IE Profile List (id, name):
    1 msgie-v2
```

show security gtp message-ie-profile-v2 msgie-v2

```
user@host> show security gtp message-ie-profile-v2 msgie-v2
Profile msgie-v2, uid 1

Message Number   IE number/Grouped-IE   Instance numbers
2
                 3                      0
32
                 1                      0
                 71                     0
                 82                     0
                 87                     0
                 87                     1
                 Bearer-ctxt-crt        0
                 Bearer-ctxt-rmv        1
```

Release Information

Command introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

> *Example: Configure Must-IE check for GTPv1 and GTPv2*
>
> message-ie-profile-v2 | **10648**

show security gtp message-list

IN THIS SECTION

- Syntax | 24447
- Description | 24447
- Options | 24447
- Required Privilege Level | 24447
- Output Fields | 24447
- Sample Output | 24448
- Release Information | 24448

Syntax

```
show security gtp message-list (all |<msg-list-name>)
```

Description

Use this command to display the list of all the GTP message-list profile list.

Options

all	Displays the list of all GTP message-list profile.
msg-list-name	Displays a specific message-list profile name.

Required Privilege Level

View

Output Fields

Sample Output

show security gtp message-list all

```
user@host> show security gtp message-list all
```

```
GTP Message-List Profile List (id, name):
1 msg-v1
```

show security gtp message-list msg-v1

```
user@host> show security gtp message-list msg-v1
```

```
Profile msg-v1, uid 1
Message Number:
16, 20
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

message-list | 10650

show security gtp profile

IN THIS SECTION

Syntax | 24449
Description | 24449

- Options | 24449
- Required Privilege Level | 24449
- Output Fields | 24449
- Sample Output | 24455
- Release Information | 24460

Syntax

```
show security gtp profile (profile-name | all)
```

Description

Use this command to display the GTP configuration. IP addresses in GPRS tunneling protocol (GTP) message on Gp or the S8 interface are validated with the configured IP group list to prevent attacks. The GTP firewall determines the IP addresses in GTP messages and matches with the configured IP group list. Based on the match criteria, valid GTP messages are forwarded to Packet and Forwarding Engine, and invalid GTP messages are dropped.

From Junos OS Release 20.4R1 onwards, `show security gprs gtp configuration` is replaced by `show security gtp profile <profile_name>`. The `identifier` option is replaced by `profile-name` in Junos OS Release 20.4R1.

Options

- profile-name —Displays specific GTP profile.
- all —Displays all the GTP profile list.

Required Privilege Level

view

Output Fields

Table 1224 on page 24450 lists the output fields for the `show security gtp profile GTP` command. Output fields are listed in the approximate order in which they appear.

Table 1224: show security gtp profile GTP

Field Name	Field Description
Index	An internal number associated with the GTP message.
Min Message Length	Displays minimum message payload length (in bytes).
Max Message Length	Displays maximum message payload length (in bytes).
Timeout	Elapsed time without activity after which the profile is terminated.
Rate Limit Alarm Threshold Rate Limit Drop Threshold	Displays limit rate of control traffic to any GSN defined in a GTP profile. A drop threshold option and an alarm threshold option reduces the duplicate drop logs for destination GSN.
Remove R6 Remove R7 Remove R8 Remove R9	Displays count of IEs that are removed from GTP messages.
Deny Nested GTP	Represents the deny of nested GTP profiles.
Validated	Represents validated address of the end user.
Restart Path	Represents the restart status of the GTP path.
Log Forwarded	Represents packets that the security device transmitted because it was valid.
Log event	A packet that the security device dropped because it failed stateful inspection, log alarms, and tunnel management events.
Log Prohibited	packet that the security device dropped because it was invalid.
Log Ratelimited	A packet that the security device dropped because it exceeded the maximum rate limit of the destination GSN.

Table 1224: show security gtp profile GTP *(Continued)*

Field Name	Field Description
Drop AA Create PDU	Represents Create AA PDU Context Request and Create AA PDU Context Response messages.
Drop AA Delete PDU	Represents Delete AA PDP Context Request and Delete AA PDP Context Response messages.
Drop Bearer Resource	Represents Bearer Resource Command and Bearer Resource Failure messages.
Drop Change Notification	Represents Change Notification Request and Change Notification Response messages.
Drop Config Transfer	Represents Configuration Transfer Tunnel messages.
Drop Context	Represents Context Request and Context Response messages.
Drop Create Bear	Represents Create Bearer Request and Create Bearer Response messages.
Drop Create Data Forwarding	Represents Create Indirect Data Forwarding Request and Create Indirect Data Forwarding Response messages.
Drop Create PDU	Represents Create PDU Context Request and Create PDU Context Response messages.
Drop Create Session	Represents Create Session Request and Create Session Response messages.
Drop Create Forwarding Tnl	Represents Create Forwarding Tunnel Request and Create Forwarding Tunnel Response messages.
Drop CS Paging	Represents CS Paging Indication messages.
Drop Data Record	Represents Data Record Request and Data Record Response messages.

Table 1224: show security gtp profile GTP *(Continued)*

Field Name	Field Description
Drop Delete Bearer	Represents Delete Bearer Request and Delete Bearer Response messages.
Drop Delete Command	Represents Delete Bearer Command and Delete Bearer Failure messages.
Drop Delete Data Forwarding	Represents Delete Indirect Data Forwarding Request and Delete Indirect Data Forwarding Response messages.
Drop Delete PDN	Represents Delete PDN Connection Set Request and Delete PDN Connection Set Response messages
Drop Delete PDP	Represents Delete PDP Context Request and Delete PDP Context Response messages.
Drop Delete Session	Represents Delete Session Request and Delete Session Response messages.
Drop Detach	Represents Detach Notification and Detach Acknowledgement messages.
Drop Downlink Notification	Represents Downlink Data Notification, Downlink Data Acknowledgement, and Downlink Data Notification Failure Indication messages.
Drop Echo	Represents Echo Request and Echo Response messages.
Drop Error Indication	Represents Error Indication messages.
Drop Failure Report	Represents Failure Report Request and Failure Report Response messages.
Drop FWD Access	Represents Forward Access Context Notification and Forward Access Context Acknowledgment messages.
Drop FWD Relocation	Represents Forward Relocation Request, Forward Relocation Response, Forward Relocation Complete, and Forward Relocation Complete Acknowledge messages.

Table 1224: show security gtp profile GTP *(Continued)*

Field Name	Field Description
Drop FWD SRNS Context	Represents Forward SRNS Context Request and Forward SRNS Context Response messages.
Drop G-PDU	Represents G-PDU and T-PDU messages.
Drop Identification	Represents Identification Request and Identification Response messages.
Drop MBMS Sess Start	Represents MBMS Session Start Request and MBMS Session Start Response messages.
Drop MBMS Sess Stop	Represents MBMS Session Stop Request and MBMS Session Stop Response messages.
Drop MBMS Sess Update	Represents MBMS Session Update Request and MBMS Session Update Response messages.
Drop Modify Bearer	Represents Modify Bearer Request and Modify Bearer Response messages.
Drop Modify Command	Represents Modify Bearer Command and Modify Bearer Failure messages.
Drop Node Alive	Represents Node Alive Request and Node Alive Response messages.
Drop Note MS Present	Represents Note MS GPRS Present Request and Note MS GPRS Present Response messages.
Drop PDU Notification	Represents PDU Notification request and PDU Notification response messages.
Drop Ran Info	Represents Ran Info Relay messages.
Drop Redirection	Represents Redirection Request and Redirection Response messages.
Drop Release Access	

Table 1224: show security gtp profile GTP *(Continued)*

Field Name	Field Description
Drop Relocation Cancel	Represents Relocations Cancel Request and Relocation Cancel Response messages.
Drop Resume	Represents Resume Notification and Resume Acknowledgement messages.
Drop Send Route	Represents Send Route Info Request and Send Route Info Response messages.
Drop SGSN Context	Represents SGSN Context Request and SGSN Context Response messages.
Drop Stop Paging	Represents Stop Paging Indication messages.
Drop Supported Extension	Represents Supported Extension Headers Notification messages.
Drop Suspend	Represents Suspend Notification and Suspend Acknowledgement messages.
Drop Trace Session	Represents Trace Session Activation in GTP.
Drop Update Bearer	Represents Update Bearer Request and Update Bearer Response messages.
Drop Update PDN	Represents Update PDN Set Connection Request and PDN Set Connection Response messages.
Drop Update PDP	Represents Update PDP Request and Update PDP Response messages.
Drop Ver Not Supported	Represents Version Not Supported messages.
Handover group name	Name of the handover IP address group.
NE group name	Name of the network equipment group.
UE group name	Name of the user equipment group.

Table 1224: show security gtp profile GTP *(Continued)*

Field Name	Field Description
Must-IE profile V1	Represents GTPv1 Must-IE check.
Must-IE profile V2	Represents GTPv2 Must-IE check.
Remove-ie-set V1	Represents GTPv1 IE Removal.
Remove-ie-set V2	Represents GTPv2 IE Removal.
Listening-mode	Represents listening-mode for GTP inspection.

Sample Output

Refer the following sample outputs for Junos OS 20.3R1 or previous releases

show security gtp configuration all (GTP Profile)

```
user@host> show security gtp configuration all
GTP Profile List (id, name):
            1 GTP
```

show security gtp configuration 1 (GTP Profile)

```
user@host>  show security gtp configuration 1
Profile Details:

        Index                       : 1
        Min Message Length          : 0
        Max Message Length          : 65535
        Timeout                     : 1000
        Rate Limit                  : 0
        Request Timeout             : 0
        Remove IE V1                : 172, 180, 181, 182, 183, 184, 199
```

```
        Remove IE V2                : 255
        Deny Nested GTP             : 0
        Validated                   : 0
        Passive learning enable     : 0
        Restart Path                : 0
        Log Forwarded               : 0
        Log State Invalid           : 0
        Log Prohibited              : 0
        Log Ratelimited             : 0
        Frequency Number            : 0
        Drop AA Create PDU          : 0
        Drop AA Delete PDU          : 0
        Drop Bearer Resource        : 0
        Drop Change Notification    : 0
        Drop Config Transfer        : 0
        Drop Context                : 0
        Drop Create Bear            : 0
        Drop Create Data Forwarding : 0
        Drop Create PDU             : 0
        Drop Create Session         : 0
        Drop Create Forwarding Tnl  : 0
        Drop CS Paging              : 0
        Drop Data Record            : 0
        Drop Delete Bearer          : 0
        Drop Delete Command         : 0
        Drop Delete Data Forwarding : 0
        Drop Delete PDN             : 0
        Drop Delete PDP             : 0
        Drop Delete Session         : 0
        Drop Detach                 : 0
        Drop Downlink Notification  : 0
        Drop Echo                   : 0
        Drop Error Indication       : 0
        Drop Failure Report         : 0
        Drop FWD Access             : 0
        Drop FWD Relocation         : 0
        Drop FWD SRNS Context       : 0
        Drop G-PDU                  : 0
        Drop Identification         : 0
        Drop MBMS Sess Start        : 0
        Drop MBMS Sess Stop         : 0
        Drop MBMS Sess Update       : 0
        Drop Modify Bearer          : 0
```

```
        Drop Modify Command         : 0
        Drop Node Alive             : 0
        Drop Note MS Present        : 0
        Drop PDU Notification       : 0
        Drop Ran Info               : 0
        Drop Redirection            : 0
        Drop Release Access         : 0
        Drop Relocation Cancel      : 0
        Drop Resume                 : 0
        Drop Send Route             : 0
        Drop SGSN Context           : 0
        Drop Stop Paging            : 0
        Drop Supported Extension    : 0
        Drop Suspend                : 0
        Drop Trace Session          : 0
        Drop Update Bearer          : 0
        Drop Update PDN             : 0
        Drop Update PDP             : 0
        Drop Ver Not Supported      : 0
        Handover group name         : N/A
        NE group name               : N/A
        UE group name               : N/A
        Must-ie profile V1          : msgie-v1
        Must-ie profile V2          : msgie-v2
        Remove-ie-set V1            : ieset-v1-r7
        Remove-ie-set V2            : ieset-v2
```

Refer the following sample output for Junos OS 20.4R1 or later releases.

show security gtp profile name (listening-mode, rate-limit, log event)

```
user@host> show security gtp profile gtp1
  Profile Details:

        Index                       : 1
        Min Message Length          : 0
        Max Message Length          : 65535
        Timeout                     : 1000
        Rate Limit Alarm Threshold  : 400
        Rate Limit Drop Threshold   : 800
  Rate Limit GTPv0 Message    : 16, 18, 20, 22, 24
  Rate Limit GTPv1 Message    : 16, 20
```

```
  Rate Limit GTPv2 Message    : 32, 34, 36, 95, 99
        Request Timeout             : 6
        Deny Nested GTP             : 0
        Validated                   : 1
……
      APN Control            : apn1
```

Refer the following sample output for Junos OS 21.2R1 or later releases.

show security gtp profile name (listening-mode, rate-limit, log event)

```
user@host> show security gtp profile gtp1
Profile Details:

        Index                       : 1
        Min Message Length          : 0
        Max Message Length          : 65535
        Timeout                     : 1000
        Rate Limit Alarm Threshold  : 2
        Rate Limit Drop Threshold   : 0
        Request Timeout             : 0
        GTP-in-GTP Denied           : 1
        User Tunnel Validation      : 1
        Sequnecy Number Validation  : 1
        End User Address Validation : 1
        Listening-mode              : enabled
        Restart Path                : 0
        Log Forwarded               : 0
        Log Event                   : 1
        Log Prohibited              : 0
        Log Ratelimited             : 0
        Log GTP-U                   : none
        Drop AA Create PDU          : 0
        Drop AA Delete PDU          : 0
        Drop Bearer Resource        : 0
        Drop Change Notification    : 0
        Drop Config Transfer        : 0
        Drop Context                : 0
        Drop Create Bear            : 0
        Drop Create Data Forwarding : 0
        Drop Create PDU             : 0
        Drop Create Session         : 0
```

```
        Drop Create Forwarding Tnl  : 0
        Drop CS Paging              : 0
        Drop Data Record            : 0
        Drop Delete Bearer          : 0
        Drop Delete Command         : 0
        Drop Delete Data Forwarding : 0
        Drop Delete PDN             : 0
        Drop Delete PDP             : 0
        Drop Delete Session         : 0
        Drop Detach                 : 0
        Drop Downlink Notification  : 0
        Drop Echo                   : 0
        Drop Error Indication       : 0
        Drop Failure Report         : 0
        Drop FWD Access             : 0
        Drop FWD Relocation         : 0
        Drop FWD SRNS Context       : 0
        Drop G-PDU                  : 0
        Drop Identification         : 0
        Drop MBMS Sess Start        : 0
        Drop MBMS Sess Stop         : 0
        Drop MBMS Sess Update       : 0
        Drop Modify Bearer          : 0
        Drop Modify Command         : 0
        Drop Node Alive             : 0
        Drop Note MS Present        : 0
        Drop PDU Notification       : 0
        Drop Ran Info               : 0
        Drop Redirection            : 0
        Drop Release Access         : 0
        Drop Relocation Cancel      : 0
        Drop Resume                 : 0
        Drop Send Route             : 0
        Drop SGSN Context           : 0
        Drop Stop Paging            : 0
        Drop Supported Extension    : 0
        Drop Suspend                : 0
        Drop Trace Session          : 0
        Drop Update Bearer          : 0
        Drop Update PDN             : 0
        Drop Update PDP             : 0
        Drop Ver Not Supported      : 0
        Handover group name         : N/A
```

```
        NE group name                : N/A
        UE group name                : N/A
        Must-ie profile V1           : N/A
        Must-ie profile V2           : N/A
        Remove-ie-set V1             : N/A
        Remove-ie-set V2             : N/A
```

Release Information

Command introduced in Junos OS Release 19.3R1. The `identifier` option is replaced by `profile-name` in Junos OS Release 20.4R1.

show security gtp rate-limit default

IN THIS SECTION

- Syntax | 24460
- Description | 24460
- Options | 24461
- Required Privilege Level | 24461
- Output Fields | 24461
- Sample Output | 24461
- Release Information | 24461

Syntax

```
show security gtp rate-limit default
```

Description

Use this command to display the default message number of GTPv0, GTPv1, and GTPv2.

Options

default Displays the default message number of GTPv0, GTPv1, and GTPv2.

Required Privilege Level

View

Output Fields

Sample Output

show security gtp rate-limit default

```
user@host> show security gtp rate-limit default
```

```
GTPv0 Message Number:
1, 16, 18, 20, 22, 24
GTPv1 Message Number:
1, 16, 20
GTPv2 Message Number:
1, 32, 34, 36, 95, 99
```

Release Information

Command introduced in Junos OS Release 21.1R1.

show security gtp sctp association

IN THIS SECTION

Syntax | 24462

- Description | 24462
- Options | 24462
- Required Privilege Level | 24463
- Output Fields | 24463
- Sample Output | 24463
- Release Information | 24464

Syntax

```
show security gtp sctp association
<all>
<destination-ip>
<destination-port>
<guid>
<init-state>
<source-ip>
<source-port>
<summary>
```

Description

Use this command to display the Stream Control Transmission Protocol (SCTP) association information.

Options

none	Displays the live security SCTP association.
all	Displays information about all the SCTP associations, both initiated and live.
destination-ip	Displays information about the destination IP address associations.
destination-port	Displays information about the destination port associations.
guid	Displays information about the globally unique identifier associations.
init	Displays information about initiated associations.

source-ip	Displays information about the source IP address associations.
source-port	Displays information about the source port associations.
summary	Displays the output summary.

Required Privilege Level

view

Output Fields

Table 1225 on page 24463 lists the output fields for the **show security gtp sctp association** command. Output fields are listed in the approximate order in which they appear.

Table 1225: show security gprs sctp association

Field Name	Field Description
Association Information	Association Information of FPC and PIC.
SCTP association numbers • **Total association**	Number of established SCTP associations. The SCTP association numbers field contains the total number of associations.
Association GUID	Globally unique association identifier information.

Sample Output

show security gprs sctp association

```
user@host>show security gtp sctp association


SCTP association numbers:
Total association 0

Association Information for FPC: 0      PIC: 1
```

```
SCTP association numbers:

Association GUID: 502a161a-a063bc44-0108000000001402
  source:
    10.3.202.118 (10.57.68.118)
    10.3.202.218 (10.57.68.218)
    port: 4215, state: SCTP_ESTABLISHED, tag: 0xe5d562d2;
  destination:
    172.28.34.206 (172.28.34.206)
    192.168.24.2 (192.168.24.2)
    port: 4215, state: SCTP_ESTABLISHED, tag: 0x631b82e4;
time left: 1786 s, access time: 45370 s;
policy id: sctp_policy/1, cfg live timeout: 30 min, handshake timeout: 20 s;

SCTP association numbers:
Total association 1

Association Information for FPC: 1      PIC: 0
SCTP association numbers:
Total association 0

Association Information for FPC: 1      PIC: 1
SCTP association numbers:
Total association 0
```

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

The **all**, **destination-ip**, **destination-port**, **guid**, **init**, **source-ip**, **source-port**, and **summary** options introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

sctp | **16396**

clear security gtp sctp association | **3952**

show security gtp sctp counters

IN THIS SECTION

- Syntax | 24465
- Description | 24465
- Options | 24465
- Required Privilege Level | 24465
- Output Fields | 24466
- Sample Output | 24467
- Sample Output | 24470
- Release Information | 24473

Syntax

```
show security gtp sctp counters <detail>
```

Description

Use this command to display the statistics of the received and dropped Stream Control Transmission Protocol (SCTP) chunks.

Options

none Displays the statistics of all received and dropped SCTP chunks.

detail Displays detailed debugging counters for SCTP chunks.

Required Privilege Level

view

Output Fields

Table 1226 on page 24466 lists the output fields for the **show security gtp sctp counters** command. Output fields are listed in the approximate order in which they appear.

Table 1226: show security gtp sctp counters

Field Name	Field Description
`Name`	Name of the SCTP payload protocol identifier.
`Received Counter`	Number of SCTP chunk counters received.
`Drop Counter`	Number of SCTP chunk counters dropped due to error.
`Counter Information`	Association information of FPC and PIC.
`Association detail counters`	(detail output only) Number of total and dying associations.
`Dbg records`	(detail output only) Number and type of debugging records.
`Packet error`	(detail output only) Number and type of packet errors.
`Association matching error`	(detail output only) Number of association matching errors.
`Association state error`	(detail output only) Number of state errors.
`Over rate drop`	(detail output only) Number of messages over the rate limit.
`Memory counters`	(detail output only) Number and type of memory counters.
`Other error`	(detail output only) Number and type of other errors.

Sample Output

show security sctp counters

```
user@host> show security gtp sctp counters
Counter Information for FPC: 1 PIC: 0
Association detail counters:
Total association: 0
Dying association: 0
Ready wrap: 0

Dbg records:
pak-without-profile : 0
pak-nat-only : 0
pak-inspection : 0
drop-at-clearing-all : 0
src-pnat : 0
dst-pnat : 0
hostname : 0
dup-init : 0
dup-initack : 0
tag-null-abort : 0
error-chunk : 0
bad-interest : 0
wing-attach : 0
wing-detach : 0
wrap-with-assoc : 0
unwrap-from-assoc : 0
conflict-assoc : 0
conflict-redr : 0
wrong-distribution : 0

Packet error:
chunk-unsupport : 0
cookie-invalid : 0
pkt-len : 0
chunk-len : 0
tag-error : 0
bad-len : 0
bad-chk-hdr : 0

Association matching error:
```

```
ha-assoc : 0
data-assoc : 0
initack-assoc : 0
sack-assoc : 0
hb-assoc : 0
hb-ack-assoc : 0
abort-assoc : 0
shutdown-assoc : 0
shutdown-ack-assoc : 0
err-assoc : 0
cookie-echo-assoc : 0
cookie-ack-assoc : 0
shutdown-complete-assoc : 0
lookup-no-assoc : 0
dup-init-diff-ip-list : 0
dup-init-diff-dst-ip : 0
dup-initack-src-ip-invalid : 0
dup-initack-diff-ip-lis : 0

Associaiton state error:
data-state : 0
init-state : 0
initack-state : 0
sack-state : 0
shutdown-state : 0
shutdown-ack-state : 0
cookie-echo-state : 0
cookie-ack-state : 0
shutdown-complete-state : 0
cookie-echo-retrans-timeout : 0
cookie-ack-retrans-timeout : 0

Association LoadBalance counter:
redirect-assoc-request-send : 0
redirect-assoc-request-ack-recv : 0
redirect-assoc-request-nack-recv : 0
redirect-assoc-request-ack-timeoute : 0
redirect-assoc-request-recv : 0
redirect-assoc-request-ack-send : 0
redirect-assoc-request-nack-send : 0

Over rate drop:
sccp : 0
```

```
ssp : 0
sst : 0

Memory counters:
alloc-assoc : 0
free-assoc : 0
alloc-redr : 0
free-redr : 0
alloc-assoc-wrap : 0
free-assoc-wrap : 0
alloc-cookie : 0
free-cookie : 0
alloc-addr : 0
free-addr : 0

HA counters:
invalid-type : 0
bad-msg : 0
no-assoc-info : 0
send-fail : 0
dup-create : 0
no-policy : 0
no-profile : 0
alloc-fail : 0
non-established-issu: 0

Other error:
over-max : 0
over-min : 0
del-error : 0
sess-cookie-set-fail : 0
sess-cookie-get-fail : 0
no-assoc-install-redr-cb : 0
wrap-alloced-failure : 0
wrap-null-assocp : 0
assoc-alloced-failure : 0
redr-assoc-alloced-failure : 0
invalid-pkt-pointer : 0
nat-jbuf-alloc-fail : 0
evt-cookie-alloc-fail : 0
```

Sample Output

show security gtp sctp counters detail

```
user@host> show security gtp sctp counters detail
Counter Information for FPC: 1 PIC: 0
Association detail counters:
Total association: 0
Dying association: 0
Ready wrap: 0

Dbg records:
pak-without-profile : 0
pak-nat-only : 0
pak-inspection : 0
drop-at-clearing-all : 0
src-pnat : 0
dst-pnat : 0
hostname : 0
dup-init : 0
dup-initack : 0
tag-null-abort : 0
error-chunk : 0
bad-interest : 0
wing-attach : 0
wing-detach : 0
wrap-with-assoc : 0
unwrap-from-assoc : 0
conflict-assoc : 0
conflict-redr : 0
wrong-distribution : 0

Packet error:
chunk-unsupport : 0
cookie-invalid : 0
pkt-len : 0
chunk-len : 0
tag-error : 0
bad-len : 0
bad-chk-hdr : 0

Association matching error:
```

```
ha-assoc : 0
data-assoc : 0
initack-assoc : 0
sack-assoc : 0
hb-assoc : 0
hb-ack-assoc : 0
abort-assoc : 0
shutdown-assoc : 0
shutdown-ack-assoc : 0
err-assoc : 0
cookie-echo-assoc : 0
cookie-ack-assoc : 0
shutdown-complete-assoc : 0
lookup-no-assoc : 0
dup-init-diff-ip-list : 0
dup-init-diff-dst-ip : 0
dup-initack-src-ip-invalid : 0
dup-initack-diff-ip-lis : 0

Associaiton state error:
data-state : 0
init-state : 0
initack-state : 0
sack-state : 0
shutdown-state : 0
shutdown-ack-state : 0
cookie-echo-state : 0
cookie-ack-state : 0
shutdown-complete-state : 0
cookie-echo-retrans-timeout : 0
cookie-ack-retrans-timeout : 0

Association LoadBalance counter:
redirect-assoc-request-send : 0
redirect-assoc-request-ack-recv : 0
redirect-assoc-request-nack-recv : 0
redirect-assoc-request-ack-timeoute : 0
redirect-assoc-request-recv : 0
redirect-assoc-request-ack-send : 0
redirect-assoc-request-nack-send : 0

Over rate drop:
sccp : 0
```

```
ssp : 0
sst : 0

Memory counters:
alloc-assoc : 0
free-assoc : 0
alloc-redr : 0
free-redr : 0
alloc-assoc-wrap : 0
free-assoc-wrap : 0
alloc-cookie : 0
free-cookie : 0
alloc-addr : 0
free-addr : 0

HA counters:
invalid-type : 0
bad-msg : 0
no-assoc-info : 0
send-fail : 0
dup-create : 0
no-policy : 0
no-profile : 0
alloc-fail : 0
non-established-issu: 0

Other error:
over-max : 0
over-min : 0
del-error : 0
sess-cookie-set-fail : 0
sess-cookie-get-fail : 0
no-assoc-install-redr-cb : 0
wrap-alloced-failure : 0
wrap-null-assocp : 0
assoc-alloced-failure : 0
redr-assoc-alloced-failure : 0
invalid-pkt-pointer : 0
nat-jbuf-alloc-fail : 0
evt-cookie-alloc-fail : 0
```

Release Information

Command introduced in Junos OS Release 10.2.

Support for the **detail** option added in Junos OS Release 12.1X45-D10.

Support for SCTP payload protocols chunk counters added in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

clear security gtp sctp counters | 3954

sctp | 16396

show security host-vpn security-associations

IN THIS SECTION

- Syntax | 24473
- Description | 24474
- Options | 24474
- Additional Information | 24474
- Required Privilege Level | 24474
- Output Fields | 24474
- Sample Output | 24476
- Release Information | 24477

Syntax

```
show security host-vpn security-associations
<connection-name>
```

Description

Display the protection details about a specified security association or all security associations.

Options

connection-name Specify for which connection the connection information is to be displayed. If no connection-name is specified, information for all security associations is displayed.

Additional Information

The Security Parameters Index (SPI) is an arbitrary value which is used (together with the destination IP address) to identify the security association of the receiving party. Each IPsec datagram has a special field for the SPI. All datagrams in the SA will use the same SPI value in this field.

Required Privilege Level

view

Output Fields

Table 1227 on page 24474 describes the output fields for the `how security host-vpn security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 1227: show security host-vpn security-associations Output Fields

Field Name	Description
`IKE SA`	Name of the security association connection.
`ID`	Identifier of the security association.

Table 1227: show security host-vpn security-associations Output Fields *(Continued)*

Field Name	Description
State	State of the parent SA connection. Values include the following: • CREATED—IKE SA just got created, but is not yet initiating or responding. • CONNECTING—IKE SA gets initiated actively or passively. • DESTROYING—IKE SA object gets destroyed. • ESTABLISHED— IKE SA is fully established. • PASSIVE—IKE SA is managed externally and does not process messages. • REKEYING—IKE SA rekeying is in progress.
(I:R)	Initiator and responder cookie.
local	Local endpoint information and identities.
remote	Remote endpoint information and identities.
crypto	Negotiated encryption details in effect (one for each IKE SA and child SA).
established	How long ago the SA was established, and when it rekeys.
Child SA	Name of the child SA.
State	State of the child SA connection. Values include the following: • CREATED—Child SA is just created, but is not yet installed. • DESTROYING—Child SA object gets destroyed. • INSTALLED— Child SA is installed and in use. • REKEYING—Child SA rekeying is in progress.

Table 1227: show security host-vpn security-associations Output Fields *(Continued)*

Field Name	Description
mode	IPsec mode: (transport | tunnel).
in spi	Inbound SPI values. Also, shows the number of bytes and packets encrypted.
out spi	Outbound SPI values. Also, shows the number of bytes and packets encrypted.
local ts	The local traffic selector (that is, what local traffic is protected).
remote ts	The remote traffic selector (that is, what remote traffic is protected).

Sample Output

show security host-vpn security-associations

```
user@host> show security host-vpn security-associations
IKE SA  : leftT1, ID:1, State:ESTABLISHED, IKEv2, (I:R):96e7757f275c3aa1:ff01ca9e7c4590b2
  local : 10.102.227.201, id:vm1@juniper.net
  remote: 10.102.228.200, id:vm1@juniper.net
  crypto: AES_CBC-256/HMAC_SHA2_384_192-0/PRF_HMAC_SHA2_384/ECP_384
  established 57s ago, rekey in 3295s
  Child SA   : childLeft1, ID:1, State:INSTALLED, mode:TUNNEL
    crypto   : ESP: AES_GCM_16-256-0
    in spi   : c5dfd0be, 5541188 bytes, 105772 packets
    out spi  : c39dbd67, 322089572 bytes, 224729 packets
    installed: 58 s ago, rekey in 3264 s, expires in 3903 s
    local ts : [10.102.227.201/32[tcp]]
    remote ts: [10.102.228.200/32[tcp/afs3-callback]]
IKE SA  : leftT2, ID:2, State:ESTABLISHED, IKEv2, (I:R):2bd786adf65eb875:0546171950dbb490
  local : 10.102.227.201, id:vm2@juniper.net
  remote: 10.102.228.200, id:vm2@juniper.net
  crypto: AES_CBC-256/HMAC_SHA2_384_192-0/PRF_HMAC_SHA2_384/ECP_384
  established 57s ago, rekey in 3475s
  Child SA   : childLeft2, ID:2, State:INSTALLED, mode:TUNNEL
    crypto   : ESP: AES_GCM_16-256-0
```

```
    in spi   : c0a912ee, 40 bytes, 1 packets
    out spi  : c52e4bf0, 60 bytes, 1 packets
    installed: 57 s ago, rekey in 3262 s, expires in 3903 s
    local ts : [10.102.227.201/32[tcp]]
    remote ts: [10.102.228.200/32[tcp/afs3-prserver]]
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- clear security host-vpn security-associations | 3956
- show security host-vpn version
- host-vpn | 8174

show security host-vpn version

IN THIS SECTION

- Syntax | 24477
- Description | 24478
- Required Privilege Level | 24478
- Sample Output | 24478
- Release Information | 24478

Syntax

```
show security host-vpn version
```

Description

Display the version of IPsec being used in the system.

Required Privilege Level

view

Sample Output

command-name

```
user@host> show security host-vpn version
Version: host IPsec 5.3.5 charon-systemd (Linux, 4.1.15-juniper-00909-g6846316, x86_64)
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

host-vpn | 8174

show security idp active-policy

IN THIS SECTION

- Syntax | 24479
- Description | 24479
- Required Privilege Level | 24479
- Output Fields | 24479
- Sample Output | 24479
- Release Information | 24480

Syntax

```
show security idp active-policy
```

Description

Display information about the policy name and running detector version with which the policy is compiled from the IDP data plane module.

Required Privilege Level

view

Output Fields

Table 1228 on page 24479 lists the output fields for the `show security idp active-policy` command. Output fields are listed in the approximate order in which they appear.

Table 1228: show security idp active-policy Output Fields

Field Name	Field Description
`Policy Name`	Name of the running policy.
`Running Detector Version`	Current version of the running detector.

Sample Output

show security idp active-policy

```
user@host> show security idp active-policy
  Policy Name : viking-policy
  Running Detector Version : 9.1.140080300
```

Release Information

Command introduced in Junos OS Release 9.2.

Starting with Junos OS Release 18.2R1, IDP policy is directly assigned in the security policy rule. This is to simplify IDP policy usage and to provide flexibility to have multiple policies active at the same time. As a part of session interest check IDP will enabled if IDP policy is present in any of the matched rules. IDP policy is activated in security policies, by permitting the IDP policy within the application services using the `set security policies from-zone zone-name to-zone zone-name policy policy-name then permit application-services idp-policy idp-policy-name` command. Since IDP policy name is directly use in the security policy rule, the `[edit security idp active-policy policy-name]` statement is deprecated.

RELATED DOCUMENTATION

request security idp security-package download | 14754

request security idp security-package install | 14757

show security idp attack attack-list

IN THIS SECTION

- Syntax | 24480
- Description | 24481
- Options | 24481
- Required Privilege Level | 24481
- Sample Output | 24481
- show security idp attack attack-list predefined-group | 24482
- Release Information | 24484

Syntax

```
show security idp attack attack-list attack-group (custom-group | dynamic-group | predefined-group)attack-group-name (recursive)
```

Description

Display list of all attacks present in the attack group specified.

You can view the attacks that are available in an attack group (predefined, dynamic, and custom attack groups). The attack option has a sub option named attack list that allows you to view attacks in an attack group. The attack list option accommodates three new options (custom, dynamic, and predefined). You can select any of these groups and provide a valid group name to see the list of attacks that belong to that group.

Starting in Junos OS Release 18.3R1, to which an attack belongs.

Options

- `custom-group` *`custom-group`*—Custom group name.
- `dynamic-group` *`dynamic-group`*—Dynamic group name.
- `predefined-group` *`predefined-group`*—Predefined group name.
- `recursive`—Recursively search through nested attack-groups.

Required Privilege Level

view

Sample Output

show security idp attack attack-list predefined-group FTP

```
user@host> show security idp attack attack-list predefined-group FTP
Processing your request, results will show up shortly
FTP:AUDIT:REP-BINARY-DATA
FTP:AUDIT:REP-INVALID-REPLY
FTP:AUDIT:REP-NESTED-REPLY
FTP:MS-FTP:STAT-GLOB
FTP:WS-FTP:CPWD
FTP:OVERFLOW:PATH-LINUX-X86-3
FTP:OVERFLOW:K4FTP-OF1
```

show security idp attack attack-list predefined-group

Check attack list against each group one at a time.

```
user@host> show security idp attack attack-list predefined-group Recommended-CTS-1-Year-Old
Processing your request, results will show up shortly!
Please use show security idp attack attack-list predefined-group command on each group listed
below to further display attacks
Recommended-CTS-Critical-1-Year-Old
Recommended-CTS-Major-1-Year-Old
Recommended-CTS-Minor-1-Year-Old
```

show security idp attack attack-list predefined-group idpengine

```
user@host> show security idp attack attack-list predefined-group idpengine
Processing your request, results will show up shortly!
Please use show security idp attack attack-list predefined-group/dynamic-group command if there
are any nested attack-groups listed below to further display attacks
RULEBASE IPS ATTACKS
  HTTP:AUDIT:REQ-LONG-UTF8CODE
  HTTP:CISCO:VOIP:STREAM-ID-REQ
  HTTP:BROWSER:ICQ
  HTTP:INFO-LEAK:SNOOP-DISCLOSURE
  HTTP:CGI:NULL-ENCODING
  HTTP:INFO:MWS-SEARCH-OF1
  HTTP:INFO:TMICRO-PROXY-REQ
  HTTP:AUDIT:URL
  HTTP:TOMCAT:REAL-PATH-REQ
  HTTP:TOMCAT:JSP-BUFFER
  HTTP:TOMCAT:JSP-COMMENTS
  HTTP:TOMCAT:JSP-PAGE
  HTTP:TOMCAT:JSP-DEC-INT-OF
  HTTP:TOMCAT:SOURCE-MAL-REQ
  HTTP:REQERR:BIN-DATA-ACC-ENC
  HTTP:TUNNEL:TELNET
  HTTP:TUNNEL:CHAT-YIM
  HTTP:TUNNEL:CHAT-AOL-IM
  HTTP:UNIX-CMD:UNIX-CMD-A-L
  HTTP:UNIX-CMD:UNIX-CMD-M-Z
  HTTP:TUNNEL:ALTNET-OVER-HTTP
  HTTP:TUNNEL:PROXY
```

```
  HTTP:MISC:MOODLOGIC-CLIENT
  HTTP:STREAM:QUICKTIME-CLIENT
  HTTP:TUNNEL:CHAT-MSN-IM
  HTTP:AUDIT:FW1-SCHEME-OF
  HTTP:HOTMAIL:FILE-DOWNLOAD
  HTTP:HOTMAIL:ZIP-DOWNLOAD
  HTTP:INFO:HTTPPOST-GETSTYLE
  HTTP:EXT:DOT-CHM
  HTTP:INFO-LEAK:HTTP-SHARE-ENUM
  HTTP:3COM:ADMIN-LOGOUT
  HTTP:PROXY:HTTP-PROXY-GET
  HTTP:HOTMAIL:FILE-UPLOAD
  HTTP:EXT:DOT-RAT
  HTTP:GMAIL:FILE-UPLOAD
  HTTP:PHP:BZOPEN-OF
  HTTP:COLDFUSION:CF-CLASS-DWLD
  HTTP:AUDIT:ROBOTS.TXT
  HTTP:STREAM:GOOGLE-VIDEO
  HTTP:STREAM:ITUNES-USERAGENT
  HTTP:INFO-LEAK:CC-CLEAR-VAR
  HTTP:IIS:ENCODING:UNICODE
  HTTP:DOMINO:INFO-LEAK
  HTTP:STREAM:YOUTUBE-REQ
HTTP:PASSWD:COMMON
  HTTP:PROXY:LIST:PUBWEBPROXIES
  HTTP:PROXY:ANON:PROXY-2
  HTTP:PROXY:LIST:PROXYFIND
  HTTP:PROXY:ANON:CGIPROXY
  HTTP:EXT:DOT-VML
  HTTP:EXT:DOT-RPT
  HTTP:PROXY:ANON:CONCEAL-WS
  HTTP:PROXY:WPAD-CONNECTION
  HTTP:PROXY:CAW-URI-RES
  HTTP:XDOMAINXML
  HTTP:INFO-LEAK:SSN-CLEARTEXT
  HTTP:AUDIT:LENGTH-OVER-256
  HTTP:AUDIT:LENGTH-OVER-512
  HTTP:AUDIT:LENGTH-OVER-1024
  HTTP:AUDIT:LENGTH-OVER-2048
  HTTP:INFO:FACEBOOK
  HTTP:INFO:MS-UPDATE
  HTTP:YAHOO:ATTACHMENT-UPLOAD
  HTTP:YAHOO:ATTACHMENT-DOWNLOAD
```

```
HTTP:INFO:YOUTUBE
HTTP:INFO:FARK
HTTP:HOTMAIL:LIVE-ACTIVITY
HTTP:YAHOO:ACTIVITY
HTTP:EXT:DOT-PPT
HTTP:INFO:SPIDER-ROBOT
HTTP:PROXY:ANON:PHPROXY
HTTP:UA:WGET
HTTP:UA:CURL
HTTP:TUNNEL:ANCHORFREE-CLIENT
HTTP:PHP:PHPINFO-QUERY
HTTP:UA:SKIPFISH
HTTP:STREAM:AAJTAK-STREAM
HTTP:STREAM:FLV
HTTP:STREAM:STARTV-STREAM
HTTP:MISC:APPLE-MAPS-APP
HTTP:AUDIT:HTTP-VER-1.0
HTTP:INFO:YOUTUBE-APP
HTTP:UA:MOBILE
HTTP:UA:CRAZY-BROWSER
HTTP:UA:GOOGLEBOT
HTTP:UA:MSN-BINGBOT
HTTP:UA:NUTCH
HTTP:UA:MOREOVER
HTTP:EK-RED-SIMPLETDS-GO
HTTP:TUNNEL:PSIPHON-TUNNEL
FTP:AUDIT:REQ-BINARY-DATA
FTP:AUDIT:REQ-INVALID-CMD-SEQ
FTP:AUDIT:REQ-NESTED-REQUEST
FTP:AUDIT:REQ-UNKNOWN-CMD
FTP:AUDIT:LOGIN-FAILED
FTP:USER:ANONYMOUS
FTP:PASSWORD:COMMON-PASSWD
FTP:PASSWORD:DEFAULT-USERNM-PW
FTP:EXT:DOT-PDF
FTP:FILE:RETR
FTP:FILE:STOR
```

Release Information

Command introduced in Junos OS Release 18.3R1.

`recursive` option introduced in Junos OS and EVO Release 21.3.

RELATED DOCUMENTATION

show security idp attack detail | 24494

show security idp attack attack-list policy

IN THIS SECTION

- Syntax | 24485
- Description | 24485
- Options | 24485
- Required Privilege Level | 24486
- With just Rule base IDP Attacks Configured | 24486
- Sample Output | 24486
- With both Rule Base and Rule Base Exempt Configured | 24486
- Release Information | 24489

Syntax

```
show security idp attack attack-list (policy Default | predefined-policy policy-name) (recursive)
```

Description

Display a list of all attacks that belong to a specified IDP policy.

Specify any configured IDP policy name to determine the attacks available in that particular IDP policy.

Options

`policy Default`—Specify the default IDP policy name.

`predefined-policy` *`policy-name`*—Specify the predefined IDP policy name.

`predefined-group` *`group-name`*—Specify the predefined IDP policy group name.

`recursive`—Recursively search through nested policies and attack-groups.

Required Privilege Level

view

With just Rule base IDP Attacks Configured

show security idp attack attack-list policy Default

Check attack-list against the default policy.

```
user@host> show security idp attack attack-list policy Default
Processing your request, results will show up shortly!
Please use show security idp attack attack-list predefined-group/dynamic-group command if there
are any nested attack-groups listed below to further display attacks
RULEBASE IPS ATTACKS
 Recommended-CTS-1-Year-Old
 Recommended-STC-1-Year-Old
 Recommended-CTS-2-Years-Old
 Recommended-STC-2-Years-Old
```

Sample Output

With both Rule Base and Rule Base Exempt Configured

run show security idp attack attack-list predefined-group FTP

```
user@host# run show security idp attack attack-list policy idpengine
Processing your request, results will show up shortly!
Please use show security idp attack attack-list predefined-group/dynamic-group command if there
are any nested attack-groups listed below to further display attacks
RULEBASE IPS ATTACKS
  HTTP:AUDIT:REQ-LONG-UTF8CODE
  HTTP:CISCO:VOIP:STREAM-ID-REQ
```

```
HTTP:BROWSER:ICQ
HTTP:INFO-LEAK:SNOOP-DISLOSURE
HTTP:CGI:NULL-ENCODING
HTTP:INFO:MWS-SEARCH-OF1
HTTP:INFO:TMICRO-PROXY-REQ
HTTP:AUDIT:URL
HTTP:TOMCAT:REAL-PATH-REQ
HTTP:TOMCAT:JSP-BUFFER
HTTP:TOMCAT:JSP-COMMENTS
HTTP:TOMCAT:JSP-PAGE
HTTP:TOMCAT:JSP-DEC-INT-OF
HTTP:TOMCAT:SOURCE-MAL-REQ
HTTP:REQERR:BIN-DATA-ACC-ENC
HTTP:TUNNEL:TELNET
HTTP:TUNNEL:CHAT-YIM
HTTP:TUNNEL:CHAT-AOL-IM
HTTP:UNIX-CMD:UNIX-CMD-A-L
HTTP:UNIX-CMD:UNIX-CMD-M-Z
HTTP:TUNNEL:ALTNET-OVER-HTTP
HTTP:TUNNEL:PROXY
HTTP:MISC:MOODLOGIC-CLIENT
HTTP:STREAM:QUICKTIME-CLIENT
HTTP:TUNNEL:CHAT-MSN-IM
HTTP:AUDIT:FW1-SCHEME-OF
HTTP:HOTMAIL:FILE-DOWNLOAD
HTTP:HOTMAIL:ZIP-DOWNLOAD
HTTP:INFO:HTTPPOST-GETSTYLE
HTTP:EXT:DOT-CHM
HTTP:INFO-LEAK:HTTP-SHARE-ENUM
HTTP:3COM:ADMIN-LOGOUT
HTTP:PROXY:HTTP-PROXY-GET
HTTP:HOTMAIL:FILE-UPLOAD
HTTP:EXT:DOT-RAT
HTTP:GMAIL:FILE-UPLOAD
HTTP:PHP:BZOPEN-OF
HTTP:COLDFUSION:CF-CLASS-DWLD
HTTP:AUDIT:ROBOTS.TXT
HTTP:STREAM:GOOGLE-VIDEO
HTTP:STREAM:ITUNES-USERAGENT
HTTP:INFO-LEAK:CC-CLEAR-VAR
HTTP:IIS:ENCODING:UNICODE
HTTP:DOMINO:INFO-LEAK
HTTP:STREAM:YOUTUBE-REQ
```

```
HTTP:PASSWD:COMMON
HTTP:PROXY:LIST:PUBWEBPROXIES
HTTP:PROXY:ANON:PROXY-2
HTTP:PROXY:LIST:PROXYFIND
HTTP:PROXY:ANON:CGIPROXY
HTTP:EXT:DOT-VML
HTTP:EXT:DOT-RPT
HTTP:PROXY:ANON:CONCEAL-WS
HTTP:PROXY:WPAD-CONNECTION
HTTP:PROXY:CAW-URI-RES
HTTP:XDOMAINXML
HTTP:INFO-LEAK:SSN-CLEARTEXT
HTTP:AUDIT:LENGTH-OVER-256
HTTP:AUDIT:LENGTH-OVER-512
HTTP:AUDIT:LENGTH-OVER-1024
HTTP:AUDIT:LENGTH-OVER-2048
HTTP:INFO:FACEBOOK
HTTP:INFO:MS-UPDATE
HTTP:YAHOO:ATTACHMENT-UPLOAD
HTTP:YAHOO:ATTACHMENT-DOWNLOAD
HTTP:INFO:YOUTUBE
HTTP:INFO:FARK
HTTP:HOTMAIL:LIVE-ACTIVITY
HTTP:YAHOO:ACTIVITY
HTTP:EXT:DOT-PPT
HTTP:INFO:SPIDER-ROBOT
HTTP:PROXY:ANON:PHPROXY
HTTP:UA:WGET
HTTP:UA:CURL
HTTP:TUNNEL:ANCHORFREE-CLIENT
HTTP:PHP:PHPINFO-QUERY
HTTP:UA:SKIPFISH
HTTP:STREAM:AAJTAK-STREAM
HTTP:STREAM:FLV
HTTP:STREAM:STARTV-STREAM
HTTP:MISC:APPLE-MAPS-APP
HTTP:AUDIT:HTTP-VER-1.0
HTTP:INFO:YOUTUBE-APP
HTTP:UA:MOBILE
HTTP:UA:CRAZY-BROWSER
HTTP:UA:GOOGLEBOT
HTTP:UA:MSN-BINGBOT
HTTP:UA:NUTCH
```

```
    HTTP:UA:MOREOVER
    HTTP:EK-RED-SIMPLETDS-GO
    HTTP:TUNNEL:PSIPHON-TUNNEL
    FTP:AUDIT:REQ-BINARY-DATA
    FTP:AUDIT:REQ-INVALID-CMD-SEQ
    FTP:AUDIT:REQ-NESTED-REQUEST
    FTP:AUDIT:REQ-UNKNOWN-CMD
    FTP:AUDIT:LOGIN-FAILED
    FTP:USER:ANONYMOUS
    FTP:PASSWORD:COMMON-PASSWD
    FTP:PASSWORD:DEFAULT-USERNM-PW
    FTP:EXT:DOT-PDF
    FTP:FILE:RETR
    FTP:FILE:STOR
  RULEBASE EXEMPT ATTACKS
    FTP:AUDIT:REQ-BINARY-DATA
    FTP:AUDIT:REQ-INVALID-CMD-SEQ
    FTP:AUDIT:REQ-NESTED-REQUEST
    FTP:AUDIT:REQ-UNKNOWN-CMD
    FTP:AUDIT:LOGIN-FAILED
    FTP:USER:ANONYMOUS
    FTP:PASSWORD:COMMON-PASSWD
    FTP:PASSWORD:DEFAULT-USERNM-PW
    FTP:EXT:DOT-PDF
    FTP:FILE:RETR
    FTP:FILE:STOR
```

Release Information

Command introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

show security idp attack detail | 24494

show security idp attack context

IN THIS SECTION

- Syntax | 24490
- Description | 24490
- Options | 24490
- Required Privilege Level | 24490
- Release Information | 24490

Syntax

```
show security idp attack context
```

Description

Display the contexts associated with the traffic passed through the interfaces only through the CLI `pcap-analysis`. It shows information about where in the packet the attack occured.

Specify any configured IDP policy name to determine the attacks available in that particular IDP policy.

Options

`context`—Specify the context-name, hits, and data.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show security idp attack detail | 24494

show security idp attack deprecated-list

IN THIS SECTION

- Syntax | 24491
- Description | 24491
- Required Privilege Level | 24491
- Sample Output | 24492
- Release Information | 24492

Syntax

```
show security idp attack deprecated-list
```

Description

Displays the list of signatures which are deprecated from the signature update file. The list of deprecated attacks is updated when a new signature database is used. Use this command to find deprecated attacks for the configured policy and deprecated lists. If any of the attacks are used in an IDP policy and that policy is associated to "Firewall Rules", then the command output displays the deprecated signature.

Required Privilege Level

view

Sample Output

show security idp attack deprecated-list

```
user@host> show security idp attack deprecated-list
FTP:USER:ANONYMOUS
FTP:USER:FORMAT-STRING
FTP:USER:ROOT
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

IDP Signature Database Overview

show security idp attack description

IN THIS SECTION

- Syntax | 24493
- Description | 24493
- Options | 24493
- Required Privilege Level | 24493
- Output Fields | 24493
- Sample Output | 24493
- Release Information | 24494

Syntax

```
show security idp attack description attack-name
```

Description

Display description of a specified IDP attack.

Options

- *attack-name* —IDP attack name.

Required Privilege Level

view

Output Fields

Table 1229 on page 24493 lists the output fields for the `show security idp attack description` command. Output fields are listed in the approximate order in which they appear.

Table 1229: show security idp attack description Output Fields

Field Name	Field Description
`Description`	IDP attack description.

Sample Output

show security idp attack description

```
user@host> show security idp attack description FTP:USER:ROOT
Description: This signature detects attempts to login to an FTP server using the "root"
account.  This can indicate an attacker trying to gain root-level access, or it can indicate
poor security practices. FTP typically uses plain-text passwords, and using the root account to
FTP could expose sensitive data over the network.
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

clear security idp attack table | 3959

show security idp attack detail

IN THIS SECTION

- Syntax | 24494
- Description | 24494
- Options | 24494
- Required Privilege Level | 24495
- Output Fields | 24495
- Sample Output | 24496
- Release Information | 24498

Syntax

```
show security idp attack detail attack-name
```

Description

Display details of a specified IDP attack.

Options

- *attack-name* —IDP attack name.

Required Privilege Level

view

Output Fields

Table 1230 on page 24495 lists the output fields for the `show security idp attack detail` command. Output fields are listed in the approximate order in which they appear.

Table 1230: show security idp attack detail Output Fields

Field Name	Field Description
`Display Name`	Display name of the IDP attack.
`Severity`	Severity level of the IDP attack.
`Category`	IDP attack category.
`Recommended`	Specifies whether a default action for the IDP attack is recommended by Juniper Networks (true or false).
`Recommended Action`	Recommended action for the IDP attack.
`Type`	Type of IDP attack.
`Direction`	Direction of the IDP attack.
`False Positives`	Specifies whether the IDP attack produces false positive on the network.
`Service`	IDP service configured for the IDP attack. If a service is configured for the IDP attack, the IDP service name is displayed. Otherwise, `Not available` is displayed.
`Member Name`	Name of attack member in IDP attack

Table 1230: show security idp attack detail Output Fields *(Continued)*

Field Name	Field Description
Expression	Specifies the Boolean expression of attack members used to identify the way(for example, OR, AND, or oAND) attack members should be matched.
PCRE Expression	Specifies the Boolean expression of PCRE format based attack members used to identify the way(for example, OR, AND, or oAND) attack members should be matched. If this field is not present “Expression” is used as a Boolean expression for attack matching.
Shellcode	Signifies if the IDP attack is a shellcode attack.
Flow	Signifies the channel(control, data) of IDP attack.
Context	Name of the context under which IDP attack has to be matched.
Negate	Signifies if the signature in the IDP attack is a negate signature.
TimeBinding	Specifies count and scope under which the attack is valid.
Pattern	Specifies the regular expression in the IDP attack.
PCRE Pattern	Specifies the regular expression in PCRE format in the IDP attack.
Hidden Pattern	Specifies if the attack pattern is hidden.

Sample Output

show security idp attack detail FTP:USER:ROOT

```
user@hostt> run show security idp attack detail FTP:USER:ROOT
Display Name: FTP: "root" Account Login
Severity: Minor
Category: FTP
```

```
Recommended: false
Recommended Action: None
Type: signature
Direction: CTS
False Positives: unknown
Shellcode: no
Flow: control
Context: ftp-username
Negate: false
TimeBinding:
            Scope: none
        Count: 1
Hidden Pattern: False
Pattern: \[root\]
```

show security idp attack detail TROJAN:MISC:ROOTBEER-CLIENT

```
user@host> show security idp attack detail TROJAN:MISC:ROOTBEER-CLIENT
Display Name: TROJAN: Digital Rootbeer Client Connect
Severity: Minor
Category: TROJAN
Recommended: false
Recommended Action: None
Type: chain
False Positives: unknown
Service: TCP/2600
Expression: m01 oAND m02
Order: no
Reset: no
Scope: session
TimeBinding:
Members:
        Member Name: m01
        Type: Signature
        Direction: CTS
        Flow: control
        Shellcode: no
        Context: stream256
        Negate: false
        Hidden Pattern: False
        Pattern: .*/QUE,who are you\.\.\.\?.*
```

```
        PCRE Pattern: ^(.)*\/QUE,who are you\.\.\.\?

        Member Name: m02
        Type: Signature
        Direction: STC
        Flow: control
        Shellcode: no
        Context: stream256
        Negate: false
        Hidden Pattern: False
        Pattern: .*/QUE,billy the kid.*
        PCRE Pattern: ^(.)*\/QUE,billy the kid
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

clear security idp attack table | 3959

show security idp attack group-list

IN THIS SECTION

- Syntax | 24499
- Description | 24499
- Options | 24499
- Required Privilege Level | 24499
- Sample Output | 24499
- Release Information | 24499

Syntax

```
show security idp attack group-list attack-name
```

Description

Display list of predefined attack groups to which the predefined attack belongs.

All the available predefined attacks are listed. You can select any attack and find the group to which that attack belongs.

Options

- *attack-name*—IDP attack name.

Required Privilege Level

view

Sample Output

show security idp attack group-list FTP:USER:ANONYMOUS

```
user@host#> show security idp attack group-list FTP:USER:ANONYMOUS
Processing your request , results will show up shortly
"Additional Web Services - Info"
"Category"
"FTP"
"FTP - All"
"FTP - Info"
"Info"
"Info - FTP"
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

IDP Policy Rules and IDP Rule Bases

show security idp attack table

IN THIS SECTION

- Syntax | 24500
- Description | 24500
- Options | 24500
- Required Privilege Level | 24501
- Output Fields | 24501
- Sample Output | 24501
- Release Information | 24502

Syntax

```
show security idp attack table
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the detailed information of IDP attack table and displays the IPv6 covert channels which are identified and mitigated.

Options

none	Displays the details of the IDP attack table.
logical-system *logical-system-name*	(Optional) Displays the details of the IDP attack table for a specific logical system.

logical-system all (Optional) Displays the details of the IDP attack table for all logical systems.

tenant *tenant-name* (Optional) Displays the details of the IDP attack table for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1231 on page 24501 lists the output fields for the `show security idp attack table` command. Output fields are listed in the approximate order in which they appear.

Table 1231: show security idp attack table Output Fields

Field Name	Field Description
`Attack name`	Name of the attack that you want to match in the monitored network traffic.
`Hits`	Total number of attack matches. On SRX Series Firewalls, for brute force and time-binding-related attacks, the logging is to be done only when the match `count` is equal to the `threshold`. That is, only one log is generated within the 60-second period in which the threshold is measured. This process prevents repetitive logs from being generated and ensures consistency with other IDP platforms, such as IDP-standalone. When no attack is seen within the 60-second period and the BFQ entry is flushed out, the match count starts over the new attack match shows up in the attack table, and the log is generated.

Sample Output

show security idp attack table

```
user@host> show security idp attack table
IDP attack statistics:
```

```
  Attack name                                   #Hits
  HTTP:OVERFLOW:PI3WEB-SLASH-OF                 1
```

show security idp attack table tenant TSYS1

```
user@host> show security idp attack table tenant TSYS1
IDP attack statistics:

  Attack name                                   #Hits
  FTP:USER:ROOT                                 1
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

IPv6 covert channels are detected in Junos OS Release 19.1R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp attack table | 3959

show security idp attacks deprecated-attacks policy policy_name

IN THIS SECTION

- Syntax | 24503
- Description | 24503
- Required Privilege Level | 24503
- Sample Output | 24503
- Release Information | 24503

Syntax

```
show security idp attacks deprecated-attacks policy policy_name (recursive)
```

Description

Displays the list of deprecated signatures configured in the policy. The list of deprecated attacks is updated when a new signature database is used.

Required Privilege Level

view

Sample Output

show security idp attack deprecated-attack-list policy

```
user@host> show security idp attacks deprecated-attacks policy policy_name
FTP:USER:ANONYMOUS
FTP:USER:FORMAT-STRING
FTP:USER:ROOT
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

IDP Signature Database Overview

show security idp counters action

IN THIS SECTION

- Syntax | 24504
- Description | 24504
- Options | 24504
- Required Privilege Level | 24505
- Output Fields | 24505
- Sample Output | 24505
- Release Information | 24507

Syntax

```
show security idp counters action
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the detailed information of IDP counter type and value.

Options

none	Displays the detailed information of IDP counter type and value.
logical system ***logical-system-name***	(Optional) Displays the detailed information of IDP counter type and value for a specific logical system.
logical system all	(Optional) Displays the detailed information of IDP counter type and value for all logical systems.
tenant ***tenant-name***	(Optional) Displays the detailed information of IDP counter type and value for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1232 on page 24505 lists the output fields for the `show security idp counters action` command. Output fields are listed in the approximate order in which they appear.

Table 1232: show security idp counters action Output Fields

Field Name	Field Description
IDP counter type	Name of the action
Value	Number of packets dropped, recommended, and ignored based on the action in the IDP counters.

Sample Output

show security idp counters action

```
user@host> show security idp counters action
IDP counters:

  IDP counter type                                                    Value
 None                                                                  0
 Recommended                                                           0
 Ignore                                                                0
 Diffserv                                                              0
 Drop packet                                                           0
 Drop                                                                  0
 Close                                                                 0
 Close server                                                          0
 Close client                                                          0
 IP action rate limit                                                  0
 IP action drop                                                        0
 IP action close                                                       0
```

```
 IP action nofity                                                    0
 IP action failed                                                    0
```

show security idp counters action logical-system LSYS0

```
user@host> show security idp counters action logical-system LSYS0
IDP counters:

  IDP counter type                                                 Value
 None                                                                 0
 Recommended                                                          0
 Ignore                                                               0
 Diffserv                                                             0
 Drop packet                                                          0
 Drop                                                                 0
 Close                                                                0
 Close server                                                         0
 Close client                                                         0
 IP action rate limit                                                 0
 IP action drop                                                       0
 IP action close                                                      0
 IP action nofity                                                     0
 IP action failed                                                     0
```

show security idp counters action tenant TSYS1

```
user@host> show security idp counters action tenant TSYS1
IDP counters:

  IDP counter type                                                 Value
 None                                                                 1
 Recommended                                                          0
 Ignore                                                               0
 Diffserv                                                             0
 Drop packet                                                          0
 Drop                                                                 0
 Close                                                                0
 Close server                                                         0
 Close client                                                         0
 IP action rate limit                                                 0
```

```
IP action drop                                                    0
IP action close                                                   0
IP action nofity                                                  0
IP action failed                                                  0
```

Release Information

Command introduced for user logical system in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters action | 3961

show security idp counters application-identification

IN THIS SECTION

- Syntax | 24507
- Description | 24508
- Options | 24508
- Required Privilege Level | 24508
- Output Fields | 24508
- Sample Output | 24511
- Release Information | 24513

Syntax

```
show security idp counters application-identification
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all IDP application identification (AI) counter values.

Options

none	Displays the status of all IDP application identification (AI) counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all IDP application identification (AI) counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all IDP application identification (AI) counter values for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all IDP application identification (AI) counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1233 on page 24508 lists the output fields for the `show security idp counters application-identification` command. Output fields are listed in the approximate order in which they appear.

Table 1233: show security idp counters application-identification Output Fields

Field Name	Field Description
`AI matches`	Number of sessions with an AI signature match.
`AI no-matches`	Number of sessions with no AI signature match.
`AI-enabled sessions`	Number of sessions with AI enabled.
`AI-disabled sessions`	Number of sessions with AI disabled.

Table 1233: show security idp counters application-identification Output Fields *(Continued)*

Field Name	Field Description
`AI-disabled sessions due to ssl encapsulated flows`	Number of sessions with AI disabled due to SSL encapsulated flows.
`AI-disabled sessions due to cache hit`	Number of sessions with AI disabled due to a cache match.
`AI-disabled sessions due to configuration`	Number of sessions with AI disabled because the configured session limit was reached.
`AI-disabled sessions due to protocol remapping`	Number of sessions with AI disabled due to protocol remapping.
`AI-disabled sessions due to RPC match`	Number of sessions with AI disabled due to an RPC match.
`AI-disabled sessions due to gate match`	Number of sessions with AI disabled due to a gate match.
`AI-disabled sessions due to non-TCP/UDP flows`	Number of sessions with AI disabled due to non-TCP or non-UDP flows.
`AI-disabled sessions due to session limit`	Number of sessions with AI disabled because the maximum session limit was reached.
`AI-disabled sessions due to session packet memory limit`	Number of sessions with AI disabled because the memory usage limit per session was reached.
`AI-disabled sessions due to global packet memory limit`	Number of sessions with AI disabled because the global memory usage limit was reached.
`AI sessions current global reass packet memory usage`	Number of AI sessions with current global reassembler packet memory usage limit

Table 1233: show security idp counters application-identification Output Fields *(Continued)*

Field Name	Field Description
`AI sessions peak global reass packet memory usage`	Number of AI sessions with peak global reassembler packet memory usage limit
`AI sessions current global packet memory usage`	Number of AI sessions with current global packet memory usage limit
`AI sessions peak global packet memory usage`	Number of AI sessions with peak global packet memory usage limit
`AI-sessions dropped due to malloc failure before session create`	Number of AI sessions dropped because the malloc failure occurred before session create.
`AI-sessions dropped due to malloc failure after create`	Number of AI sessions dropped because the malloc failure occurred after session create.
`AI-Packets received on sessions marked for drop due to malloc failure`	Number of AI packets received on sessions that are marked to be dropped because the malloc failure.
`Packets cloned for AI`	Number of packets cloned for application identification.
`Policy update`	Number of times the IDP policy has been updated.
`Total PME prematch job ignored`	Number of jobs ignored because of pattern matching engine (PME) not matching.
`Total packets for which prematch job were ignored`	Number of packets for which signature matching was ignored as prematch found.
`Prematch busy packet count`	Number of packets saved as they are handed off for signature matching during prematch reprocess.

Table 1233: show security idp counters application-identification Output Fields *(Continued)*

Field Name	Field Description
`Final match busy packet count`	Number of packets saved as they are handed off for signature matching during final match reprocess.
`Total AI busy packet count`	Number of times AI saved packet handed off for signature matching.
`Final match processed busy packet count`	Number of times a packet processed for final matching before signature matching.
`Prematch processed busy packet count`	Number of times a packet processed for prematch before signature match.
`Prematch ignored busy packet count`	Number of packets ignored for signature matching as prematch found.
`AI done busy packet count`	Number of packets signature matching not completed before AI done.
`JPME flow for Ignored jobs destroyed`	Number of jobs destroyed because of flow mismatch due to policy relookup.
`Set AI done for prematch`	Number of sessions set for AI applied.
`AI done for prematch`	Number of sessions with AI applied.

Sample Output

show security idp counters application-identification

```
user@host> show security idp counters application-identification

IDP counter type                                                    Value
 AI matches                                                            0
 AI no-matches                                                         0
```

```
  AI-enabled sessions                                                  0
  AI-disabled sessions                                                 0
  AI-disabled sessions due to ssl encapsulated flows                   0
  AI-disabled sessions due to cache hit                                0
  AI-disabled sessions due to configuration                            0
  AI-disabled sessions due to protocol remapping                       0
  AI-disabled sessions due to RPC match                                0
  AI-disabled sessions due to gate match                               0
  AI-disabled sessions due to non-TCP/UDP flows                        0
  AI-disabled sessions due to session limit                            0
  AI-disabled sessions due to session packet memory limit              0
  AI-disabled sessions due to global packet memory limit               0
  AI sessions current global reass packet memory usage                 0
  AI sessions peak global reass packet memory usage                    0
  AI sessions current global packet memory usage                       0
  AI sessions peak global packet memory usage                          0
  AI-sessions dropped due to malloc failure before session create      0
  AI-sessions dropped due to malloc failure after create               0
  AI-Packets received on sessions marked for drop due to malloc failure  0
  Packets cloned for AI                                                0
  Policy update                                                        0
  Total PME prematch job ignored                                       0
  Total packets for which prematch job were ignored                    0
  Prematch busy packet count                                           0
  Final match busy packet count                                        0
  Total AI busy packet count                                           0
  Final match processed busy packet count                              0
  Prematch processed busy packet count                                 0
  Prematch ignored busy packet count                                   0
  AI done busy packet count                                            0
  JPME flow for Ignored jobs destroyed                                 0
  Set AI done for prematch                                             0
  AI done for prematch                                                 0
                                                                       0
```

show security idp counters application-identification tenant TSYS1

```
user@host> show security idp counters application-identification tenant TSYS1
IDP counters:

  IDP counter type                                                  Value
```

```
AI matches                                                          0
AI no-matches                                                       0
AI-enabled sessions                                                 0
AI-disabled sessions                                                1
AI-disabled sessions due to ssl encapsulated flows                  0
AI-disabled sessions due to cache hit                               1
AI-disabled sessions due to configuration                           0
AI-disabled sessions due to protocol remapping                      0
AI-disabled sessions due to RPC match                               0
AI-disabled sessions due to gate match                              0
AI-disabled sessions due to non-TCP/UDP flows                       0
AI-disabled sessions due to global packet memory limit              0
AI sessions current global packet memory usage                      0
AI sessions peak global packet memory usage                         0
AI-sessions dropped due to malloc failure before session create     0
AI-sessions dropped due to malloc failure after create              0
AI-Packets received on sessions marked for drop due to malloc failure   0
Packets cloned for AI                                               0
Policy update                                                       0
Total PME prematch job ignored                                      0
Total packets for which prematch job were ignored                   0
Prematch busy packet count                                          0
Final match busy packet count                                       0
Total AI busy packet count                                          0
Final match processed busy packet count                             0
Prematch processed busy packet count                                0
Prematch ignored busy packet count                                  0
AI done busy packet count                                           0
JPME flow for Ignored jobs destroyed                                0
Set AI done for prematch                                            0
AI done for prematch                                                0
```

Release Information

Command introduced in Junos OS Release 9.2. Modified in Junos OS Release 12.1.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters application-identification | 3962

show security idp counters dfa

IN THIS SECTION

- Syntax | 24514
- Description | 24514
- Options | 24514
- Required Privilege Level | 24515
- Output Fields | 24515
- Sample Output | 24515
- Release Information | 24516

Syntax

```
show security idp counters dfa
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all DFA counter values.

Options

none	Displays the status of all DFA counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all DFA counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all DFA counter values for all logical systems.

tenant *tenant-name* (Optional) Displays the status of all DFA counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1234 on page 24515 lists the output fields for the `show security idp counters dfa` command. Output fields are listed in the approximate order in which they appear.

Table 1234: show security idp counters dfa Output Fields

Field Name	Field Description
`DFA Group Merged Usage`	Number of DFA groups merged.
`DFA Matches`	Number of DFA matches found.

Sample Output

show security idp counters dfa

```
user@host> show security idp counters dfa
IDP counters:

  IDP counter type                                                    Value
 DFA Group Merged Usage                                                0
 DFA Matches                                                           0
 DFA compressed                                                        0
 DFA group compressed                                                  0
 DFA uncompressed                                                      0
 DFA group uncompressed                                                0
```

show security idp counters dfa logical-system LSYS1

```
user@host> show security idp counters dfa logical-system LSYS1
IDP counters:

  IDP counter type                                                    Value
 DFA Group Merged Usage                                                 0
 DFA Matches                                                            0
 DFA compressed                                                         0
 DFA group compressed                                                   0
 DFA uncompressed                                                       0
 DFA group uncompressed                                                 0
```

show security idp counters dfa tenant TSYS1

```
user@host> show security idp counters dfa tenant TSYS1
IDP counters:

  IDP counter type                                                    Value
 DFA Group Merged Usage                                                 0
 DFA Matches                                                            1
 DFA compressed                                                         0
 DFA group compressed                                                   0
 DFA uncompressed                                                       0
 DFA group uncompressed                                                 0
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters dfa | 3964

show security idp counters flow

IN THIS SECTION

- Syntax | 24517
- Description | 24517
- Options | 24517
- Required Privilege Level | 24518
- Output Fields | 24518
- Sample Output | 24526
- Release Information | 24533

Syntax

```
show security idp counters flow
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all IDP flow counter values.

NOTE: On SRX Series Firewalls with IDP enabled, if IDP attacks are configured for a single direction (server or client), a flow in the opposite direction does not need IDP processing. For TCP traffic, the TCP optimization feature ensures minimal processing for these flows without running into reassembly errors.

Options

none	Displays the status of all IDP flow counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all IDP flow counter values for a specific logical system.

logical-system all (Optional) Displays the status of all IDP flow counter values for all logical systems.

tenant *tenant-name* (Optional) Displays the status of all IDP flow counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1235 on page 24518 lists the output fields for the `show security idp counters flow` command. Output fields are listed in the approximate order in which they appear.

Table 1235: show security idp counters flow Output Fields

Field Name	Description
`Fast-path packets`	Number of packets that are set through fast path after completing IDP policy lookup.
`Slow-path packets`	Number of packets that are sent through slow path during IDP policy lookup.
`Session construction failed` (Unsupported)	Number of times the packet failed to establish the session.
`Session limit reached`	Number of sessions that reached IDP sessions limit.
`Session inspection depth reached`	Number of sessions that reached inspection depth.
`Memory limit reached`	Number of sessions that reached memory limit.
`Not a new session` (Unsupported)	Number of sessions that extended beyond time limit.

Table 1235: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
`Invalid index at age-out` (Unsupported)	Invalid session index in session age-out message.
`Packet logging`	Number of packets saved for packet logging.
`Policy cache hits`	Number of sessions that matched policy cache.
`Policy cache misses`	Number of sessions that did not match policy cache.
`Policy cache entries`	Number of policy cache entries.
`Maximum flow hash collisions`	Maximum number of packets, of one flow, that share the same hash value.
`Flow hash collisions`	Number of packets that share the same hash value.
`Gates added`	Number of gate entries added for dynamic port identification.
`Gate matches` (Unsupported)	Number of times a gate is matched.
`Sessions deleted`	Number of sessions deleted.
`Sessions aged-out` (Unsupported)	Number of sessions that are aged out if no traffic is received within session timeout value.
`Sessions in-use while aged-out` (Unsupported)	Number of sessions in use during session age-out.

Table 1235: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
`TCP flows marked dead on RST/FIN`	Number of sessions marked dead on TCP RST/FIN.
`policy init failed`	Policy initiation failed.
`Number of times Sessions exceed high mark`	Number of times sessions exceeded the high mark.
`Number of sessions exceeds high mark`	Number of sessions that exceed high mark.
`Number of sessions drops below low mark`	Number of sessions that fall below low mark.
`Memory of sessions exceeds high mark`	Session memory exceeds high mark.
`Memory of sessions drops below low mark`	Session memory drops below low mark.
`SM Sessions encountered memory failures`	Number of SM sessions that encountered memory failures.
`SM Packets on sessions with memory failures`	Number of SM packets that encountered memory failures.
`Sessions constructed`	Number of sessions established.
`SM Sessions dropped`	Number of SM sessions dropped.
`SM sessions ignored`	Number of sessions ignored in Security Module (SM).

Table 1235: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
SM sessions interested	Number of SM sessions interested.
SM sessions not interested	Number of SM sessions not interested.
SM sessions interest error	Number of errors created for SM sessions interested.
Sessions destructed	Number of sessions destructed.
SM Session Create	Number of SM sessions created.
SM Packet Process	Number of packets processed from SM.
SM FTP data session ignored by IDP	Number of SM FTP data sessions that are ignored by IDP.
SM Session close	Number of SM sessions closed.
SM client-to-server packets	Number of SM client-to-server packets.
SM server-to-client packets	Number of SM server-to-client packets.
SM client-to-server L7 bytes	Number of SM client-to-server Layer 7 bytes.
SM server-to-client L7 bytes	Number of SM server-to-client Layer 7 bytes.
Client-to-server flows ignored	Number of client-to-server flow sessions that are ignored.
Server-to-client flows ignored	Number of server-to-client flow sessions that are ignored.

Table 1235: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
`Server-to-client flows tcp optimized`	Number of server-to-client flow TCP sessions that are optimized.
`Client-to-server flows tcp optimized`	Number of client-to-server flow TCP sessions that are optimized.
`Both directions flows ignored`	Number of server-to-client and client-to-server flow sessions that are ignored.
`Fail-over sessions dropped`	Number of failover sessions dropped.
`Sessions dropped due to no policy`	Number of sessions dropped because there was no active IDP policy.
`IDP Stream Sessions dropped due to memory failure`	Number of IDP stream sessions that are dropped because of memory failure.
`IDP Stream Sessions ignored due to memory failure`	Number of IDP stream sessions that are ignored because of memory failure.
`IDP Stream Sessions closed due to memory failure`	Number of IDP stream sessions that are closed because of memory failure.
`IDP Stream Sessions accepted`	Number of IDP stream sessions that are accepted.
`IDP Stream Sessions constructed`	Number of IDP stream sessions that are constructed.
`IDP Stream Sessions destructed`	Number of IDP stream sessions that are destructed.
`IDP Stream Move Data`	Number of stream data events handled by IDP.

Table 1235: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
`IDP Stream Sessions ignored on JSF SSL Event`	Number of IDP stream sessions that are ignored because of a JSF SSL proxy event.
`IDP Stream Sessions not processed for no matching rules`	Number of IDP stream sessions that are not processed for no matching rules.
`IDP Stream stbuf dropped`	Number of IDP stream plug-in buffers dropped.
`IDP Stream stbuf reinjected`	Number of IDP stream plug-in buffers injected.
`Busy packets from stream plugin`	Number of packets saved as one or more packets of this session from stream plug-in.
`Busy packets from packets plugin`	Number of saved packets for IDP stream plug-in sessions.
`Bad kpp`	Number of internal marked packets logged for IDP processing.
`Lsys policy id lookup failed sessions`	Number of sessions that failed logical systems policy lookup.
`Busy packets`	Number of packets saved as one or more packets of this session are handed off for asynchronous processing.
`Busy packet errors`	Number of packets found with IP checksum error after asynchronous processing is completed.
`Dropped queued packets` (async mode)	Number of queued packets dropped based on policy action, reinjection failures, or if the session is marked to destruct.

Table 1235: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
`Dropped queued packets failed` (async mode)	Not used currently.
`Reinjected packets (async mode)`	Number of packets reinjected into the queue.
`Reinjected packets failed(async mode)`	Number of failed reinjected packets.
`AI saved processed packet`	Number of AI packets saved for which the asynchronous processing is completed.
`Busy packet count incremented`	Number of times the busy packet count incremented in asynchronous processing.
`busy packet count decremented`	Number of times the busy packet count decremented in asynchronous processing.
`session destructed in pme`	Number of sessions destructed as a part of asynchronous result processing.
`session destruct set in pme`	Number of sessions set to be destructed as a result of asynchronous processing.

Table 1235: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
`KQ op`	Number of sessions with one of the following status: • KQ op hold–number of times packets held by IDP. • KQ op drop–number of times packets dropped by IDP. • KQ op route–number of times IDP decided to be route the packet directly. • KQ op Continue–number of times IDP decided to continue to process the packet. • KQ op error–number of times error occurred while IPD processing packet. • KQ op stop–number of times IDP decided to stop processing the packet.
`PME wait not set`	Number of AI saved packets given for signature matching.
`PME wait set`	Number of packets given for signature matching without AI save.
`PME KQ run not called`	Number of times signature matching results processed out of packet receiving order.
`IDP sessions ignored for content decompression in intel inspect mode`	Number of IDP session ignored for content decompression in the IDP intelligent inspection mode.
`IDP sessions ignored for bytes depth limit in intel inspect mode`	Number of IDP session ignored for bytes depth in the IDP intelligent inspection mode.
`IDP sessions ignored for protocol decoding in intel inspect mode`	Number of IDP session ignored for protocol decoding in the IDP intelligent inspection mode.

Table 1235: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
IDP sessions detected CPU usage crossed intel inspect CPU threshold	Number of IDP session detected when the CPU usage crosses the CPU threshold of the IDP intelligent inspection.
IDP sessions detected mem drop below intel inspect low mem threshold	Number of IDP session detected when memory drops below the IDP intelligent inspect low memory threshold.

Sample Output

show security idp counters flow

```
user@host> show security idp counters flow
IDP counters:

  IDP counter type                                                    Value
 Fast-path packets                                                     40252
 Slow-path packets                                                     127
 Session construction failed                                           0
 Session limit reached                                                 0
 Session inspection depth reached                                      0
 Memory limit reached                                                  0
 Not a new session                                                     0
 Invalid index at ageout                                               0
 Packet logging                                                        0
 Policy cache hits                                                     92
 Policy cache misses                                                   67
 Policy cache entries                                                  67
 Maximum flow hash collisions                                          0
 Flow hash collisions                                                  0
 Gates added                                                           0
 Gate matches                                                          0
 Sessions deleted                                                      127
 Sessions aged-out                                                     0
 Sessions in-use while aged-out                                        0
```

```
TCP flows marked dead on RST/FIN                                        13
Policy init failed                                                      0
Number of times Sessions exceed high mark                               0
Number of times Sessions drop below low mark                            0
Memory of Sessions exceeds high mark                                    0
Memory of Sessions drops below low mark                                 0
SM Sessions encountered memory failures                                 0
SM Packets on sessions with  memory failures                            0
IDP session gate creation requests                                      0
IDP session gate creation acknowledgements                              0
IDP session gate hits                                                   0
IDP session gate timeouts                                               0
Number of times Sessions crossed the CPU threshold value that is set    0
Number of times Sessions crossed the CPU upper threshold                0
Sessions constructed                                                    127
SM Sessions ignored                                                     0
SM Sessions dropped                                                     0
SM Sessions interested                                                  168
SM Sessions not interested                                              4
SM Sessions interest error                                              0
Sessions destructed                                                     127
SM Session Create                                                       127
SM Packet Process                                                       52257
SM ftp data session ignored by idp                                      0
SM Session close                                                        127
SM Client-to-server packets                                             20066
SM Server-to-client packets                                             32191
SM Client-to-server L7 bytes                                            167292
SM Server-to-client L7 bytes                                            28523514
Client-to-server flows ignored                                          1
Server-to-client flows ignored                                          1
Server-to-client flows tcp optimized                                    3
Client-to-server flows tcp optimized                                    0
Both directions flows ignored                                           32
Fail-over sessions dropped                                              0
Sessions dropped due to no policy                                       0
IDP Stream Sessions dropped due to memory failure                       0
IDP Stream Sessions ignored due to memory failure                       0
IDP Stream Sessions closed due to memory failure                        0
IDP Stream Sessions accepted                                            0
IDP Stream Sessions constructed                                         0
IDP Stream Sessions destructed                                          0
IDP Stream Move Data                                                    0
```

```
 IDP Stream Sessions ignored on JSF SSL Event                          0
 IDP Stream Sessions not processed for no matching rules               0
 IDP Stream stbuf dropped                                              0
 IDP Stream stbuf reinjected                                           0
 Busy pkts from stream plugin                                          0
 Busy pkts from pkt plugin                                             0
 bad kpp                                                               0
 Lsys policy id lookup failed sessions                                 0
 Busy packets                                                          0
 Busy packet Errors                                                    0
 Dropped queued packets (async mode)                                   0
 Dropped queued packets failed(async mode)                             0
 Reinjected packets (async mode)                                       0
 Reinjected packets failed(async mode)                                 0
 AI saved processed packet                                             0
 busy packet count incremented                                         0
 busy packet count decremented                                         0
 session destructed in pme                                             0
 session destruct set in pme                                           0
 kq op hold                                                            0
 kq op drop                                                            0
 kq op route                                                           0
 kq op continue                                                        35155
 kq op error                                                           0
 kq op stop                                                            0
 PME wait not set                                                      0
 PME wait set                                                          0
 PME KQ run not called                                                 0
 IDP sessions ignored for content decompression in intel inspect mode  0
 IDP sessions ignored for bytes depth limit in intel inspect mode      0
 IDP sessions ignored for protocol decoding in intel inspect mode      0
 IDP sessions detected CPU usage crossed intel inspect CPU threshold   0
 IDP sessions detected mem drop below intel inspect low mem threshold  0
```

show security idp counters flow tenant TSYS1

```
user@host> show security idp counters flow tenant TSYS1
IDP counters:

  IDP counter type                                                   Value
 Fast-path packets                                                    38
```

```
Slow-path packets                                                    1
Session construction failed                                          0
Session limit reached                                                0
Session inspection depth reached                                     0
Memory limit reached                                                 0
Not a new session                                                    0
Invalid index at ageout                                              0
Packet logging                                                       0
Policy cache hits                                                    0
Policy cache misses                                                  1
Policy cache entries                                                 0
Maximum flow hash collisions                                         0
Flow hash collisions                                                 0
Gates added                                                          0
Gate matches                                                         0
Sessions deleted                                                     1
Sessions aged-out                                                    0
Sessions in-use while aged-out                                       0
TCP flows marked dead on RST/FIN                                     1
Policy init failed                                                   0
Policy reinit failed                                                 0
Number of times Sessions exceed high mark                            0
Number of times Sessions drop below low mark                         0
Memory of Sessions exceeds high mark                                 0
Memory of Sessions drops below low mark                              0
SM Sessions encountered memory failures                              0
SM Packets on sessions with  memory failures                         0
IDP session gate creation requests                                   0
IDP session gate creation acknowledgements                           0
IDP session gate hits                                                0
IDP session gate timeouts                                            0
Number of times Sessions crossed the CPU threshold value that is set    0
Number of times Sessions crossed the CPU upper threshold             0
Sessions constructed                                                 1
SM Sessions ignored                                                  0
SM Sessions dropped                                                  0
SM Sessions interested                                               2
SM Sessions not interested                                           0
SM Sessions interest error                                           0
Sessions destructed                                                  1
SM Session Create                                                    1
SM Packet Process                                                    38
SM ftp data session ignored by idp                                   1
```

```
SM Session close                                                    1
SM Client-to-server packets                                         15
SM Server-to-client packets                                         23
SM Client-to-server L7 bytes                                        99
SM Server-to-client L7 bytes                                        367
Client-to-server flows ignored                                      0
Server-to-client flows ignored                                      0
Server-to-client flows tcp optimized                                0
Client-to-server flows tcp optimized                                0
Both directions flows ignored                                       1
Fail-over sessions dropped                                          0
Sessions dropped due to no policy                                   0
IDP Stream Sessions dropped due to memory failure                   0
IDP Stream Sessions ignored due to memory failure                   0
IDP Stream Sessions closed due to memory failure                    0
IDP Stream Sessions accepted                                        0
IDP Stream Sessions constructed                                     0
IDP Stream Sessions destructed                                      0
IDP Stream Move Data                                                0
IDP Stream Sessions ignored on JSF SSL Event                        0
IDP Stream Sessions not processed for no matching rules             0
IDP Stream stbuf dropped                                            0
IDP Stream stbuf reinjected                                         0
Busy pkts from stream plugin                                        0
Busy pkts from pkt plugin                                           0
bad kpp                                                             0
Lsys policy id lookup failed sessions                               0
NGAppID Events with no L7 App                                       0
NGAppID Events with no active-policy                                0
NGAppID Detector failed from event handler                          0
NGAppID Detector failed from API                                    0
Busy packets                                                        0
Busy packet Errors                                                  0
Dropped queued packets (async mode)                                 0
Dropped queued packets failed(async mode)                           0
Reinjected packets (async mode)                                     0
Reinjected packets failed(async mode)                               0
AI saved processed packet                                           0
busy packet count incremented                                       0
busy packet count decremented                                       0
session destructed in pme                                           0
session destruct set in pme                                         0
kq op hold                                                          0
```

```
 kq op drop                                                          0
 kq op route                                                         0
 kq op continue                                                      37
 kq op error                                                         0
 kq op stop                                                          0
 PME wait not set                                                    0
 PME wait set                                                        0
 PME KQ run not called                                               0
 IDP sessions ignored for content decompression in intel inspect mode    0
 IDP sessions ignored for bytes depth limit in intel inspect mode        0
 IDP sessions ignored for protocol decoding in intel inspect mode        0
 IDP sessions detected CPU usage crossed intel inspect CPU threshold     0
 IDP sessions detected mem drop below intel inspect low mem threshold    0
```

show security idp counters flow (unified parser)

```
user@host> show security idp counters flow
IDP counters:

  IDP counter type                                                  Value
 Fast-path packets                                                   38
 Slow-path packets                                                   1
 Session construction failed                                         0
 Session limit reached                                               0
 Session inspection depth reached                                    0
 Memory limit reached                                                0
 Not a new session                                                   0
 Invalid index at ageout                                             0
 Packet logging                                                      0
 Policy cache hits                                                   0
 Policy cache misses                                                 1
 Policy cache entries                                                0
 Maximum flow hash collisions                                        0
 Flow hash collisions                                                0
 Gates added                                                         0
 Gate matches                                                        0
 Sessions deleted                                                    1
 Sessions aged-out                                                   0
 Sessions in-use while aged-out                                      0
 TCP flows marked dead on RST/FIN                                    1
 Policy init failed                                                  0
```

```
  Policy reinit failed                                                  0
  Number of times Sessions exceed high mark                             0
  Number of times Sessions drop below low mark                          0
  Memory of Sessions exceeds high mark                                  0
  Memory of Sessions drops below low mark                               0
  SM Sessions encountered memory failures                               0
  SM Packets on sessions with  memory failures                          0
  IDP session gate creation requests                                    0
  IDP session gate creation acknowledgements                            0
  IDP session gate hits                                                 0
  IDP session gate timeouts                                             0
  Number of times Sessions crossed the CPU threshold value that is set  0
  Number of times Sessions crossed the CPU upper threshold              0
  Sessions constructed                                                  1
  SM Sessions ignored                                                   0
  SM Sessions dropped                                                   0
  SM Sessions interested                                                2
  SM Sessions not interested                                            0
  SM Sessions interest error                                            0
  Sessions destructed                                                   13
 Sessions processed with Unified Parser                                11
  Sessions processed without Unified Parser                             4
  SM Session Create                                                     15
  SM Packet Process                                                     38
  SM ftp data session ignored by idp                                    1
  SM Session close                                                      1
  SM Client-to-server packets                                           15
  SM Server-to-client packets                                           23
  SM Client-to-server L7 bytes                                          99
  SM Server-to-client L7 bytes                                          367
  Client-to-server flows ignored                                        0
  Server-to-client flows ignored                                        0
  Server-to-client flows tcp optimized                                  0
  Client-to-server flows tcp optimized                                  0
  Both directions flows ignored                                         1
  Fail-over sessions dropped                                            0
  Sessions dropped due to no policy                                     0
  IDP Stream Sessions dropped due to memory failure                     0
  IDP Stream Sessions ignored due to memory failure                     0
  IDP Stream Sessions closed due to memory failure                      0
  IDP Stream Sessions accepted                                          0
  IDP Stream Sessions constructed                                       0
  IDP Stream Sessions destructed                                        0
```

```
IDP Stream Move Data                                                0
IDP Stream Sessions ignored on JSF SSL Event                        0
IDP Stream Sessions not processed for no matching rules             0
IDP Stream stbuf dropped                                            0
IDP Stream stbuf reinjected                                         0
Busy pkts from stream plugin                                        0
Busy pkts from pkt plugin                                           0
bad kpp                                                             0
Lsys policy id lookup failed sessions                               0
NGAppID Events with no L7 App                                       0
NGAppID Events with no active-policy                                0
NGAppID Detector failed from event handler                          0
NGAppID Detector failed from API                                    0
Busy packets                                                        0
Busy packet Errors                                                  0
Dropped queued packets (async mode)                                 0
Dropped queued packets failed(async mode)                           0
Reinjected packets (async mode)                                     0
Reinjected packets failed(async mode)                               0
AI saved processed packet                                           0
busy packet count incremented                                       0
busy packet count decremented                                       0
session destructed in pme                                           0
session destruct set in pme                                         0
kq op hold                                                          0
kq op drop                                                          0
kq op route                                                         0
kq op continue                                                      37
kq op error                                                         0
kq op stop                                                          0
PME wait not set                                                    0
PME wait set                                                        0
PME KQ run not called                                               0
IDP sessions ignored for content decompression in intel inspect mode    0
IDP sessions ignored for bytes depth limit in intel inspect mode        0
IDP sessions ignored for protocol decoding in intel inspect mode        0
IDP sessions detected CPU usage crossed intel inspect CPU threshold     0
IDP sessions detected mem drop below intel inspect low mem threshold    0
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

flow (Security IDP) | 7339

clear security idp counters flow | 3965

show security idp counters http-decoder

IN THIS SECTION

- Syntax | 24534
- Description | 24534
- Options | 24535
- Required Privilege Level | 24535
- Output Fields | 24535
- Sample Output | 24536
- Release Information | 24537

Syntax

```
show security idp counters http-decoder
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all HTTP decoders.

Options

none	Displays the status of all HTTP decoders.
logical-system ***logical-system-name***	(Optional) Displays the status of all HTTP decoders for a specific logical system.
logical-system all	(Optional) Displays the status of all HTTP decoders for all logical systems.
tenant ***tenant-name***	(Optional) Displays the status of all HTTP decoders for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1236 on page 24535 lists the output fields for the `show security idp counters http-decoder` command. Output fields are listed in the approximate order in which they appear.

Table 1236: show security idp counters http-decoder Output Fields

Field Name	Field Description
`No of file-decoder requests from MIME over HTTP`	Number of active file decoder requests sent over HTTP from MIME.
`No of pending file-decoder requests from MIME over HTTP`	Number of pending file decoder requests sent over HTTP from MIME.
`No of completed file-decoder requests from MIME over HTTP`	Number of completed file decoder requests sent over HTTP from MIME.
`No of unrecognized file type from MIME over HTTP`	Number of unrecognized file types sent over HTTP from MIME.

Table 1236: show security idp counters http-decoder Output Fields *(Continued)*

Field Name	Field Description
No of compressed payload transferred over HTTP	Number of compressed files transferred over HTTP from MIME.

Sample Output

show security idp counters http-decoder

```
user@host> show security idp counters http-decoder
IDP counters:
IDP counter type                                                    Value
No of file-decoder requests from MIME over HTTP                        0
No of pending file-decoder requests from MIME over HTTP                0
No of completed file-decoder requests from MIME over HTTP              0
No of unrecognized file type from MIME over HTTP                       0
No of compressed payload transferred over HTTP                         0
No of bypassed files over HTTP                                        0
```

show security idp counters http-decoder logical-system LSYS1

```
user@host> show security idp counters http-decoder logical-system LSYS1
IDP counters:

  IDP counter type                                                    Value
 No of file-decoder requests from MIME over HTTP                        0
 No of pending file-decoder requests from MIME over HTTP                0
 No of completd file-decoder requests from MIME over HTTP               0
 No of unrecognized file type from MIME over HTTP                       0
 No of compressed payload transferred over HTTP                         0
 No of bypassed files over HTTP                                         0
```

show security idp counters http-decoder tenant TSYS1

```
user@host> show security idp counters http-decoder tenant TSYS1
IDP counters:

  IDP counter type                                                    Value
 No of file-decoder requests from MIME over HTTP                       0
 No of pending file-decoder requests from MIME over HTTP               0
 No of completd file-decoder requests from MIME over HTTP              0
 No of unrecognized file type from MIME over HTTP                      0
 No of compressed payload transferred over HTTP                        0
 No of bypassed files over HTTP                                        0
```

Release Information

Command introduced in Junos OS Release 11.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters http-decoder | 3966

show security idp counters ips

IN THIS SECTION

- Syntax | 24538
- Description | 24538
- Options | 24538
- Required Privilege Level | 24538
- Output Fields | 24538
- Sample Output | 24541

Release Information | 24545

Syntax

```
show security idp counters ips
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all IPS counter values.

Options

none	Displays the status of all IPS counter values for root-system.
logical-system *logical-system-name*	(Optional) Displays the status of all IPS counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all IPS counter values for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all IPS counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1237 on page 24539 lists the output fields for the `show security idp counters ips` command. Output fields are listed in the approximate order in which they appear.

Table 1237: show security idp counters ips Output Fields

Field Name	Field Description
`TCP fast path`	Number of TCP packets skipped for IDS processing.
`Layer-4 anomalies`	Number of Layer-4 protocol error or anomaly.
`Anomaly hash misses`	Number of times look failed on anomaly hash.
`Line context matches`	Number of attempts to match line based attacks in traffic stream.
`Stream256 context matches`	Number of attempts to match stream based attacks in first 256 bytes of traffic stream.
`Stream context matches`	Number of attempts to match stream based attacks in traffic stream.
`Packet context matches`	Number of attempts to match packet based attacks in traffic packet.
`Packet header matches`	Number of attempts to match packet header based attacks in traffic packet.
`Context matches`	Number of attempts to match protocol context based attacks in traffic stream.
`Regular expression matches`	Number of attempts to match PCRE expressions in traffic stream.
`Tail DFAs`	Number of attempts to match an attack on tail DFA group matches.
`Exempted attacks`	Number of attacks exempted from match as per exempt rulebase.
`Out of order chains`	Number of times attack is excluded from match due to member attacks in an attack group did not complete chain.
`Partial chain matches`	Number of attacks in partial chain match with attack scope as transaction.

Table 1237: show security idp counters ips Output Fields *(Continued)*

Field Name	Field Description
`IDS device FIFO size`	Number of IDS contexts in virtual IDS device.
`IDS device FIFO overflows`	Number of times an IDS context can not be written as the IDS device is full.
`Brute force queue size`	Number of entries in the brute force queue.
`IDS cache hits` (Unsupported)	Number of sessions those found attack instance in IDS cache.
`IDS cache misses` (Unsupported)	Number of sessions those did not find attack instance in IDS cache.
`Shellcode detection invocations`	Number of times shell code match is attempted.
`Wrong offsets`	Number of times attack's offset is not within the service offset range.
`No peer MAC` (Unsupported)	Number of times flow peer MAC address is not available.
`Number of times custom feed updated`	Number of times the custom feeds are updated.
`Number of times custom feed update failed due to error`	Number of times the custom feed updates failed due to an error.
`Number of times custom feed update failed due to out of memory`	Number of times custom feed updates failed due to memory capacity.

Table 1237: show security idp counters ips Output Fields *(Continued)*

Field Name	Field Description
`Number of times custom feed update failed due to feed not found`	Number of times custom feed updates failed due to the feed not found.
`Number of times custom feed update returned unexpected value`	Number of times custom feed updates returned an unexpected value.

Sample Output

show security idp counters ips

```
user@host> show security idp counters ips
IDP counters:
 IDP counter type                                                    Value
 TCP fast path                                                         15
 Layer-4 anomalies                                                      0
 Anomaly hash misses                                                    3
 Line context matches                                                   5
 Stream256 context matches                                              5
 Stream context matches                                                 5
 Packet context matches                                                 0
 Packet header matches                                                  0
 Context matches                                                       12
 Regular expression matches                                             0
 Tail DFAs                                                              0
 Exempted attacks                                                       0
 Out of order chains                                                    0
 Partial chain matches                                                  0
 IDS device FIFO size                                                   0
 IDS device FIFO overflows                                              0
 Brute force queue size                                                 0
 IDS cache hits                                                         0
 IDS cache misses                                                       0
 Shellcode detection invocations                                        0
```

```
 Wrong offsets                                                          0
 No peer MAC                                                            0
 Content-decompression memory usage in KB                               0
 Content-decompression memory over limit                                0
 Content-decompression gunzip called                                    0
 Content-decompression gunzip failed                                    0
 Content-decompression others called                                    0
 Content-decompression others failed                                    0
 Content-decompression input bytes                                      0
 Content-decompression output bytes                                     0
 Content-decompression ratio over limit                                 0
 Content-decompression type mismatch                                    0
 URL track session bypassed                                             0
 Exceeded max Tail DFA transition limit                                 0
 Number of times HS stream close failed                                 0
 Number of times HS stream open failed                                  0
 Number of times HS scan stream failed                                  0
 Number of times HS scan failed                                         0
 Number of times custom feed updated                                    0
 Number of times custom feed update failed due to error                 0
 Number of times custom feed update failed due to out of memory         0
 Number of times custom feed update failed due to feed not found        0
 Number of times custom feed update returned unexpected value           0
```

show security idp counters ips logical-system LSYS1

```
user@host> show security idp counters ips logical-system LSYS1
IDP counters:

  IDP counter type                                                    Value
 TCP fast path                                                         40
 Layer-4 anomalies                                                      0
 Anomaly hash misses                                                    4
 Line context matches                                                   0
 Stream256 context matches                                              0
 Stream context matches                                                 0
 Packet context matches                                                 0
 Packet header matches                                                  0
 Context matches                                                        4
 Context reset                                                          0
 Regular expression matches                                             0
```

```
  Tail DFAs                                                          0
  Exempted attacks                                                   0
  Out of order chains                                                0
  Partial chain matches                                              0
  IDS device FIFO size                                               0
  IDS device FIFO overflows                                          0
  Brute force queue size                                             2
  IDS cache hits                                                     0
  IDS cache misses                                                   0
  Shellcode detection invocations                                    0
  Wrong offsets                                                      0
  No peer MAC                                                        0
  Content-decompression memory usage in KB                           0
  Content-decompression memory over limit                            0
  Content-decompression gunzip called                                0
  Content-decompression gunzip failed                                0
  Content-decompression others called                                0
  Content-decompression others failed                                0
  Content-decompression input bytes                                  0
  Content-decompression output bytes                                 0
  Content-decompression ratio over limit                             0
  Content-decompression type mismatch                                0
  URL track session bypassed                                         0
  Exceeded max Tail DFA transition limit                             0
  Number of times HS stream close failed                             0
  Number of times HS stream open failed                              0
  Number of times HS scan stream failed                              0
  Number of times HS scan failed                                     0
  Number of times custom feed updated                                0
  Number of times custom feed update failed due to error             0
  Number of times custom feed update failed due to out of memory     0
  Number of times custom feed update failed due to feed not found    0
  Number of times custom feed update returned unexpected value       0
```

show security idp counters ips tenant TSYS1

```
user@host> show security idp counters ips tenant TSYS1
IDP counters:

   IDP counter type                                                Value
  TCP fast path                                                     16
```

```
Layer-4 anomalies                                                    0
Anomaly hash misses                                                  1
Line context matches                                                 0
Stream256 context matches                                            0
Stream context matches                                               0
Packet context matches                                               0
Packet header matches                                                0
Context matches                                                      1
Context reset                                                        0
Regular expression matches                                           0
Tail DFAs                                                            0
Exempted attacks                                                     0
Out of order chains                                                  0
Partial chain matches                                                0
IDS device FIFO size                                                 0
IDS device FIFO overflows                                            0
Brute force queue size                                               0
IDS cache hits                                                       0
IDS cache misses                                                     0
Shellcode detection invocations                                      0
Wrong offsets                                                        0
No peer MAC                                                          0
Content-decompression memory usage in KB                             0
Content-decompression memory over limit                              0
Content-decompression gunzip called                                  0
Content-decompression gunzip failed                                  0
Content-decompression others called                                  0
Content-decompression others failed                                  0
Content-decompression input bytes                                    0
Content-decompression output bytes                                   0
Content-decompression ratio over limit                               0
Content-decompression type mismatch                                  0
URL track session bypassed                                           0
Exceeded max Tail DFA transition limit                               0
Number of times HS stream close failed                               0
Number of times HS stream open failed                                0
Number of times HS scan stream failed                                0
Number of times HS scan failed                                       0
Number of times custom feed updated                                  0
Number of times custom feed update failed due to error               0
Number of times custom feed update failed due to out of memory       0
```

```
Number of times custom feed update failed due to feed not found         0
Number of times custom feed update returned unexpected value            0
```

Release Information

Command modified in Junos OS Release 11.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

ips | 9295

clear security idp counters ips | 3968

show security idp counters log

IN THIS SECTION

- Syntax | 24545
- Description | 24546
- Options | 24546
- Required Privilege Level | 24546
- Output Fields | 24546
- Sample Output | 24549
- Release Information | 24552

Syntax

```
show security idp counters log
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all IDP log counter values.

Options

none	Displays the status of all IDP log counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all IDP log counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all IDP log counter values for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all IDP log counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1238 on page 24546 lists the output fields for the `show security idp counters log` command. Output fields are listed in the approximate order in which they appear.

Table 1238: show security idp counters log Output Fields

Field Name	Field Description
`Logs dropped`	Number of logs that are dropped.
`Suppressed log count`	Number of logs that are suppressed.
`Logs waiting for post-window packets` (Unsupported)	Number of logs waiting for post-window packets.

Table 1238: show security idp counters log Output Fields *(Continued)*

Field Name	Field Description
`Logs ready to be sent` (Unsupported)	Number of logs ready to be sent.
`Logs in suppression list` (Unsupported)	Number of logs considered for suppression list.
`Log timers created`	Number of times the log timer is created.
`Logs timers expired`	Number of times the log timer is expired.
`Log timers cancelled`	Number of times the log timer is canceled.
`Logs ready to be sent high watermark` (Unsupported)	Number of packets that are ready to be sent with high degree watermark.
`Log receive buffer full` (Unsupported)	Number of times the buffer is full.
`Packet log too big` (Unsupported)	Number of packet logs that exceeded allowed packet log size.
`Reads per second` (Unsupported)	Number of packets that are read per second.
`Logs in read buffer high watermark` (Unsupported)	Number of high watermark packets that are in read buffer.

Table 1238: show security idp counters log Output Fields *(Continued)*

Field Name	Field Description
`Packets logged`	Number of packets that are logged,
`Packets lost` (Unsupported)	Number of packets that are failed to log.
`Packets copied` (Unsupported)	Number of packets copied during packet log.
`Packets held` (Unsupported)	Number of packets held for packet log.
`Packets released`	Number of packets that are released from hold.
`IP Action Messages` (Unsupported)	Number of IP action messages.
`IP Action Drops` (Unsupported)	Number of IP action messages dropped.
`IP Action Exists` (Unsupported)	Number of exits during IP action creation.
`NWaits` (Unsupported)	Number of logs waiting for post window packets.
`Match vectors`	Number of attacks in IDS match vector.
`Supercedes`	Number of attacks in supercede vector.

Sample Output

show security idp counters log

```
user@host> show security idp counters log
IDP counters:
IDP counter type                                                    Value
Logs dropped                                                          0
Suppressed log count                                                  0
Logs waiting for post-window packets                                  0
Logs ready to be sent                                                 0
Logs in suppression list                                              0
Log timers created                                                    0
Logs timers expired                                                   0
Log timers cancelled                                                  0
Logs ready to be sent high watermark                                  0
Log receive buffer full                                               0
Packet log too big                                                    0
Reads per second                                                      1
Logs in read buffer high watermark                                    0
Log Bytes in read buffer high watermark                               0
Packets logged                                                        0
Packets lost                                                          0
Packets copied                                                        0
Packets held                                                          0
Packets released                                                      0
IP Action Messages                                                    0
IP Action Drops                                                       0
IP Action Exists                                                      0
NWaits                                                                0
Match vectors                                                         0
Supercedes                                                            0
Kpacket too big                                                       0
```

show security idp counters log logical-system LSYS1

```
user@host> show security idp counters log logical-system LSYS1
IDP counters:
IDP counter type                                                    Value
Logs dropped                                                          0
```

```
Suppressed log count                                                    0
Logs waiting for post-window packets                                    0
Logs ready to be sent                                                   0
Logs in suppression list                                                0
Log timers created                                                      0
Logs timers expired                                                     0
Log timers cancelled                                                    0
Logs ready to be sent high watermark                                    0
Log receive buffer full                                                 0
Packet log too big                                                      0
Reads per second                                                        0
Logs in read buffer high watermark                                      0
Log Bytes in read buffer high watermark                                 0
Packets logged                                                          0
Packets lost                                                            0
Packets copied                                                          0
Packets held                                                            0
Packets released                                                        0
IP Action Messages                                                      0
IP Action Drops                                                         0
IP Action Exists                                                        0
NWaits                                                                  0
Match vectors                                                           0
Supercedes                                                              0
send succeed                                                            0
send fail                                                               0
retries on send failures                                                0
uac send succeed                                                        0
uac send fail                                                           0
idpd to flowd alloc msg fail                                            0
idpd to flowd enqueue log  msg fail                                     0
idpd to flowd enqueue log msg succeed                                   0
idpd to flowdlog msg dequeued                                           0
idpd to flowdlog unknown msg type                                       0
flowd send succeed                                                      0
flowd send fail                                                         0
objcache alloc failure for sc_pcap_mbuf_info_t                         0
pcap mbuf alloc fail counter                                            0
pcap mbuf reinj failed                                                  0
pcap fragmented packets count                                           0
idpd to flowd pcap messages count in dedicated mode                     0
idpd pcap type1 messages count                                          0
idpd pcap type2 messages count                                          0
```

```
idpd pcap type3 messages count                                          0
Kpacket too big                                                         0
```

show security idp counters log tenant TSYS1

```
user@host> show security idp counters log tenant TSYS1
IDP counters:

  IDP counter type                                                     Value
 Logs dropped                                                           0
 Suppressed log count                                                   0
 Logs waiting for post-window packets                                   0
 Logs ready to be sent                                                  0
 Logs in suppression list                                               0
 Log timers created                                                     0
 Logs timers expired                                                    0
 Log timers cancelled                                                   0
 Logs ready to be sent high watermark                                   0
 Log receive buffer full                                                0
 Packet log too big                                                     0
 Reads per second                                                       0
 Logs in read buffer high watermark                                     0
 Log Bytes in read buffer high watermark                                0
 Packets logged                                                         0
 Packets lost                                                           0
 Packets copied                                                         0
 Packets held                                                           0
 Packets released                                                       0
 IP Action Messages                                                     0
 IP Action Drops                                                        0
 IP Action Exists                                                       0
 NWaits                                                                 0
 Match vectors                                                          0
 Supercedes                                                             0
 send succeed                                                           1
 send fail                                                              0
 retries on send failures                                               0
 uac send succeed                                                       0
 uac send fail                                                          0
 idpd to flowd alloc msg fail                                           0
 idpd to flowd enqueue log  msg fail                                    0
```

```
idpd to flowd enqueue log msg succeed                                  0
idpd to flowdlog msg dequeued                                          0
idpd to flowdlog unknown msg type                                      0
flowd send succeed                                                     0
flowd send fail                                                        0
 objcache alloc failure for sc_pcap_mbuf_info_t                        0
pcap mbuf alloc fail counter                                           0
pcap mbuf reinj failed                                                 0
pcap fragmented packets count                                          0
idpd to flowd pcap messages count in dedicated mode                    0
idpd pcap type1 messages count                                         0
idpd pcap type2 messages count                                         0
idpd pcap type3 messages count                                         0
Kpacket too big                                                        0
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

event-rate | 6531

clear security idp counters log | 3969

show security idp counters memory

IN THIS SECTION

- Syntax | 24553
- Description | 24553
- Options | 24553
- Required Privilege Level | 24553

- Output Fields | 24553
- Sample Output | 24554
- Release Information | 24556

Syntax

```
show security idp counters memory
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays detailed information of allocated, reallocated IDP counters memory values.

Options

none	Displays detailed information of allocated, reallocated IDP counters memory values.
logical-system ***logical-system-name***	(Optional) Displays detailed information of allocated, reallocated IDP counters memory values for a specific logical system.
logical-system all	(Optional) Displays detailed information of allocated, reallocated IDP counters memory values for all logical systems.
tenant ***tenant-name***	(Optional) Displays detailed information of allocated, reallocated IDP counters memory values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1239 on page 24554 lists the output fields for the `show security idp counters memory` command. Output fields are listed in the approximate order in which they appear.

Table 1239: show security idp counters memory Output Fields

Field Name	Field Description
IDP counter type	Name of the counter type
Value	Number of requests made for memory allocation and reallocation,

Sample Output

show security idp counters memory

```
user@host> show security idp counters memory
IDP counters:

  IDP counter type                                                    Value
 Memory allocation requested                                           928058
 Memory reallocation requested                                         0
 Memory allocation failed                                              0
 Memory reallocation failed                                            0
 Memory free requested                                                 889749
 Memory free failed                                                    0
 IDP Arena requested                                                   0
 IDP Arena failed                                                      0
 IDP Arena freed                                                       0
 Objcache requested                                                    132
 Objcache failed                                                       0
 Objcache freed                                                        2
 Objcache over limit                                                   0
 Objcache invalid record                                               0
 Detector Objcache requested                                           21032
 Detector Objcache failed                                              0
 Detector Objcache freed                                               20324
 Detector Objcache invalid record                                      0
 Detector Arena requested                                              0
 Detector Arena failed                                                 0
 Detector Arena freed                                                  0
 Kzalloc requested                                                     910823
 Kzalloc failed                                                        0
```

```
 Kzalloc freed                                                    885264
 Pool alloc called                                                0
 Pool alloc subscribed size                                       0
 Pool alloc done                                                  0
 Pool free called                                                 0
 Pool free done                                                   0
 Malloc over limit                                                0
 Remote free deadbeef                                             0
```

show security idp counters memory tenant TSYS1

```
user@host> show security idp counters memory tenant TSYS1
IDP counters:

  IDP counter type                                               Value
 Memory allocation requested                                      78
 Memory reallocation requested                                    0
 Memory allocation failed                                         0
 Memory reallocation failed                                       0
 Memory free requested                                            75
 Memory free failed                                               0
 IDP Arena requested                                              0
 IDP Arena failed                                                 0
 IDP Arena freed                                                  0
 Objcache requested                                               70
 Objcache failed                                                  0
 Objcache freed                                                   68
 Objcache over limit                                              0
 Objcache invalid record                                          0
 Detector Objcache requested                                      13
 Detector Objcache failed                                         0
 Detector Objcache freed                                          13
 Detector Objcache invalid record                                 0
 Detector Arena requested                                         0
 Detector Arena failed                                            0
 Detector Arena freed                                             0
 Kzalloc requested                                                8
 Kzalloc failed                                                   0
 Kzalloc freed                                                    7
 Pool alloc called                                                0
 Pool alloc subscribed size                                       0
```

```
Pool alloc done                                                    0
Pool free called                                                   0
Pool free done                                                     0
Malloc over limit                                                  0
Remote free deadbeef                                               0
```

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters memory | 3971

show security idp counters packet

IN THIS SECTION

- Syntax | 24557
- Description | 24557
- Options | 24557
- Required Privilege Level | 24557
- Output Fields | 24557
- Sample Output | 24560
- Release Information | 24566

Syntax

```
show security idp counters packet
<logical-system (logical-system-name | all)>
<tenant (tenant-name | all)>
```

Description

Displays the status of all IDP packet counter values.

Options

none	Displays the status of all IDP packet counter values.
logical-system ***logical-system-name***	(Optional) Displays the status of all IDP packet counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all IDP packet counter values for all logical systems.
tenant ***tenant-name***	(Optional) Displays the status of all IDP packet counter values for a specific tenant system.
tenant all	(Optional) Displays the status of all IDP packet counter values for a all tenant systems.

Required Privilege Level

view

Output Fields

Table 1240 on page 24558 lists the output fields for the `show security idp counters packet` command. Output fields are listed in the approximate order in which they appear.

Table 1240: show security idp counters packet Output Fields

Field Name	Field Description
`Processed packets`	Number of packets processed by the IDP service.
`Dropped packets`	Number of packets dropped by the IDP service. The counter for all dropped packets.
`Dropped by IDP policy`	Number of packets dropped by the IDP policy. The counter for dropped packets due to the action specified in the IDP policy (starting with the attack detection).
`Dropped by Error`	Number of packets dropped by error. The difference between `Dropped packets` and `Dropped by IDP policy`. IDS drops are primarily due to policy actions. Reassembly errors lead to packet drops. So all drops shown in `show security idp counters ips`, `show security idp counters flow` and `show security idp counters tcp-reassembler` add to `Dropped by Error`. All drops includes reassembly errors, anomalies similar to bad ip header and TTL errors.
`Dropped sessions` (Unsupported)	Number of sessions dropped.
`Bad IP headers`	Number of packets that fail IP header length validity check.
`Packets with IP options`	Number of packets that contain the optional header fields.
`Decapsulated packets`	Number of packets that are decapsulated.
`GRE decapsulations` (Unsupported)	Number of packets that are generic routing encapsulation (GRE) decapsulated.

Table 1240: show security idp counters packet Output Fields *(Continued)*

Field Name	Field Description
`PPP decapsulations` (Unsupported)	Number of packets that are Point-to-Point Protocol (PPP) decapsulated.
`TCP decompression uncompressed IP` (Unsupported)	Number of uncompressed IP headers that are to be TCP decompressed.
`TCP decompression compressed IP` (Unsupported)	Number of compressed IP headers that are to be TCP decompressed.
`Deferred-send packets` (Unsupported)	Number of deferred IP packets that are sent out.
`IP-in-IP packets` (Unsupported)	Number of packets that are IP-in-IP encapsulated.
`TTL errors` (Unsupported)	Number of packets with TTL error in the header.
`Routing loops` (Unsupported)	Number of packets that continue to be routed in an endless circle due to an inconsistent routing state.
`No-route packets` (Unsupported)	Number of packets that could not be routed further.
`Flood IP` (Unsupported)	Number of packets that are identified as IP flood packets.

Table 1240: show security idp counters packet Output Fields *(Continued)*

Field Name	Field Description
`Invalid ethernet headers` (Unsupported)	Number of packets that are identified with an invalid Ethernet header.
`Packets attached`	Number of packets attached.
`Packets cloned`	Number of packets that are cloned.
`Packets allocated`	Number of packets allocated.
`Packets destructed`	Number of packets destructed.

Sample Output

show security idp counters packet

```
user@host> show security idp counters packet
IDP counters:
 IDP counter type                                                    Value
 Processed packets                                                     27
 Dropped packets                                                        0
 Dropped by IDP policy                                                  0
 Dropped by error                                                       0
 Dropped sessions                                                       0
 Bad IP headers                                                         0
 Packets with IP options                                                0
 Decapsulated packets                                                   0
 GRE decapsulations                                                     0
 PPP decapsulations                                                     0
 TCP decompression uncompressed IP                                      0
 TCP decompression compressed IP                                        0
 Deferred-send packets                                                  0
 IP-in-IP packets                                                       0
 TTL errors                                                             0
```

```
 Routing loops                                                    0
 STP drops                                                        0
 No-route packets                                                 0
 Flood IP                                                         0
 Invalid ethernet headers                                         0
 Packets attached                                                 28
 Packets cloned                                                   28
 Packets allocated                                                0
 Packets destructed                                               55
```

show security idp counters packet logical-system LSYS1

```
user@host> show security idp counters packet logical-system LSYS1
IDP counters:
Logical System: LSYS1
  IDP counter type                                               Value
 Processed packets                                                64
 Dropped packets                                                  0
 Dropped ICMP packets                                             0
 Dropped TCP packets                                              0
 Dropped UDP packets                                              0
 Dropped Other packets                                            0
 Dropped by IDP Policy                                            0
 Dropped by Error                                                 0
 Dropped sessions                                                 0
 Bad IP headers                                                   0
 Packets with IP options                                          0
 Decapsulated packets                                             0
 GRE decapsulations                                               0
 PPP decapsulations                                               0
 GTP decapsulations                                               0
 GTP flows                                                        0
 TCP decompression uncompressed IP                                0
 TCP decompression compressed IP                                  0
 Deferred-send packets                                            0
 IP-in-IP packets                                                 0
 TTL errors                                                       0
 Routing loops                                                    0
 STP drops                                                        0
 No-route packets                                                 0
 Flood IP                                                         0
```

```
 Invalid ethernet headers                                          0
 Packets attached                                                  64
 IP Packet attach failed                                           0
 Packets cloned                                                    25
 Packets allocated                                                 0
 Packets destructed                                                89
 Packet data buffer allocated                                      24
 Packet data buffer released                                       24
 Buffer allocation on clone avoided                                0
 Late buffer allocation on clone                                   0
 Distinct clone request                                            0
 KPP clone buf cache allocated                                     0
 KPP clone buf cache released                                      0
 KPP clone buf cache used                                          0
 KQMSG constructed                                                 69
 KQMSG destructed                                                  69
 jbuf copy failed                                                  0
 jbuf pullup failed                                                0
 jbuf copy done                                                    0
 jbuf copy freed                                                   0
 jbuf copy reinjected                                              0
```

show security idp counters packet tenant TSYS1

```
user@host> show security idp counters packet tenant TSYS1
IDP counters:
Tenant: TSYS1
  IDP counter type                                                Value
 Processed packets                                                 38
 Dropped packets                                                   0
 Dropped ICMP packets                                              0
 Dropped TCP packets                                               0
 Dropped UDP packets                                               0
 Dropped Other packets                                             0
 Dropped by IDP Policy                                             0
 Dropped by Error                                                  0
 Dropped sessions                                                  0
 Bad IP headers                                                    0
 Packets with IP options                                           0
 Decapsulated packets                                              0
 GRE decapsulations                                                0
```

```
  PPP decapsulations                                                    0
  GTP decapsulations                                                    0
  GTP flows                                                             0
  TCP decompression uncompressed IP                                     0
  TCP decompression compressed IP                                       0
  Deferred-send packets                                                 0
  IP-in-IP packets                                                      0
  TTL errors                                                            0
  Routing loops                                                         0
  STP drops                                                             0
  No-route packets                                                      0
  Flood IP                                                              0
  Invalid ethernet headers                                              0
  Packets attached                                                      38
  IP Packet attach failed                                               0
  Packets cloned                                                        21
  Packets allocated                                                     0
  Packets destructed                                                    59
  Packets destructed in pipeline                                        0
  Packet data buffer allocated                                          21
  Packet data buffer released                                           21
  Buffer allocation on clone avoided                                    0
  Late buffer allocation on clone                                       0
  Distinct clone request                                                0
  KPP clone buf cache allocated                                         0
  KPP clone buf cache released                                          0
  KPP clone buf cache used                                              0
  KQMSG constructed                                                     38
  KQMSG destructed                                                      38
  KQMSG destructed in pipeline                                          0
  jbuf copy failed                                                      0
  jbuf pullup failed                                                    0
  jbuf copy done                                                        0
  jbuf copy freed                                                       0
  jbuf copy reinjected                                                  0
```

show security idp counters packet tenant all

```
Tenant: Tn3
IDP counters:

  IDP counter type                                                    Value
```

```
Processed packets                                                    0
Dropped packets                                                      0
Dropped ICMP packets                                                 0
Dropped TCP packets                                                  0
Dropped UDP packets                                                  0
Dropped Other packets                                                0
Dropped by IDP Policy                                                0
Dropped by Error                                                     0
Dropped sessions                                                     0
Bad IP headers                                                       0
Packets with IP options                                              0
Decapsulated packets                                                 0
GRE decapsulations                                                   0
PPP decapsulations                                                   0
GTP decapsulations                                                   0
GTP flows                                                            0
TCP decompression uncompressed IP                                    0
TCP decompression compressed IP                                      0
Deferred-send packets                                                0
Ktimer entry optimized                                               0
TTL errors                                                           0
Routing loops                                                        0
STP drops                                                            0
No-route packets                                                     0
Flood IP                                                             0
Invalid ethernet headers                                             0
Packets attached                                                     0
IP Packet attach failed                                              0
Packets cloned                                                       0
Packets allocated                                                    0
Packets destructed                                                   0
Packets destructed in pipeline                                       0
Packet data buffer allocated                                         0
Packet data buffer released                                          0
Buffer allocation on clone avoided                                   0
Late buffer allocation on clone                                      0
Distinct clone request                                               0
KPP clone buf cache allocated                                        0
KPP clone buf cache released                                         0
KPP clone buf cache used                                             0
KQMSG constructed                                                    0
KQMSG destructed                                                     0
KQMSG destructed in pipeline                                         0
```

```
jbuf copy failed                                                    0
jbuf pullup failed                                                  0
jbuf copy done                                                      0
jbuf copy freed                                                     0
jbuf copy reinjected                                                0

Tenant: Tn2
IDP counters:

  IDP counter type                                                  Value
Processed packets                                                   0
Dropped packets                                                     0
Dropped ICMP packets                                                0
Dropped TCP packets                                                 0
Dropped UDP packets                                                 0
Dropped Other packets                                               0
Dropped by IDP Policy                                               0
Dropped by Error                                                    0
Dropped sessions                                                    0
Bad IP headers                                                      0
Packets with IP options                                             0
Decapsulated packets                                                0
GRE decapsulations                                                  0
PPP decapsulations                                                  0
GTP decapsulations                                                  0
GTP flows                                                           0
TCP decompression uncompressed IP                                   0
TCP decompression compressed IP                                     0
Deferred-send packets                                               0
Ktimer entry optimized                                              0
TTL errors                                                          0
Routing loops                                                       0
STP drops                                                           0
No-route packets                                                    0
Flood IP                                                            0
Invalid ethernet headers                                            0
Packets attached                                                    0
IP Packet attach failed                                             0
Packets cloned                                                      0
Packets allocated                                                   0
Packets destructed                                                  0
Packets destructed in pipeline                                      0
Packet data buffer allocated                                        0
```

```
Packet data buffer released                                    0
Buffer allocation on clone avoided                             0
Late buffer allocation on clone                                0
Distinct clone request                                         0
KPP clone buf cache allocated                                  0
KPP clone buf cache released                                   0
KPP clone buf cache used                                       0
KQMSG constructed                                              0
KQMSG destructed                                               0
KQMSG destructed in pipeline                                   0
jbuf copy failed                                               0
jbuf pullup failed                                             0
jbuf copy done                                                 0
jbuf copy freed                                                0
jbuf copy reinjected                                           0
```

Usage

The command "show security idp counters" with options such as action, application-identification, dfa, flow, http-decoder, ips, log, memory, packet, packet-log, pdf-decoder, policy-manager, tcp-reassembler, now displays system names for all IDP counters in addition to root-logical-system counters.

Release Information

Command introduced in Junos OS Release 9.2.

The fields `Dropped by IDP policy` and `Dropped by Error` added in Junos OS Release 10.1.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters packet | 3972

show security idp counters packet-log

IN THIS SECTION

- Syntax | 24567
- Description | 24567
- Options | 24567
- Required Privilege Level | 24568
- Output Fields | 24568
- Sample Output | 24571
- Release Information | 24581

Syntax

```
show security idp counters packet-log
<logical-system (logical-system-name | all)>
<tenant (tenant-name | all)>
```

Description

Displays the values of all IDP packet-log counters.

Starting in Junos OS Release 22.1R1, you can enable a secure SSL or TLS connection to send encrypted IDP packet capture log to the packet capture receiver. To establish the SSL or TLS connection, you must specify the SSL initiation profile that you want to use in the IDP packet log configuration. With this encryption support the output of this command is also enhanced to provide information on maximum number of RTCOM ports.

Options

none	Displays the values of all IDP packet-log counters.
logical-system *logical-system-name*	(Optional) Displays the values of all IDP packet-log counters for a specific logical system.

logical-system all (Optional) Displays the values of all IDP packet-log counters for all logical systems. Displays values for all IDP counters from Junos release version 20.3R3.

tenant *tenant-name* (Optional) Displays the values of all IDP packet-log counters for a specific tenant system.

Required Privilege Level

view

Output Fields

The following table lists the output fields for the `show security idp counters packet-log` command. Output fields are listed in the approximate order in which they appear.

Field Name	Field Description
`Total packets captured since packet capture was activated`	Number of packets captured by the device by the IDP service.
`Total sessions enabled since packet capture was activated`	Number of sessions that have performed packet capture since the capture facility was activated.
`Sessions currently enabled for packet capture`	Number of sessions that are actively capturing packets at this time.
`Packets currently captured for enabled sessions`	Number of packets that have been captured by active sessions.
`Packet clone failures`	Number of packet capture failures due to cloning error.
`Session log object failures`	Number of objects containing log messages generated during packet capture that were not successfully transmitted to the host.
`Session packet log object failures`	Number of objects containing captured packets that were not successfully transmitted to the host.

(Continued)

Field Name	Field Description
`Sessions skipped because session limit exceeded`	Number of sessions that could not initiate packet capture because the maximum number of sessions specified for the device were conducting captures at that time.
`Packets skipped because packet limit exceeded`	Number of packets not captured because the packet limit specified for this device was reached.
`Packets skipped because total memory limit exceeded`	Number of packets not captured because the memory allocated for packet capture on this device was exceeded.
`Packet log host route lookup failures`	Number of times captured packet failures due to missing route details.
`Number of Packet-Log RTCOM create req successful`	This counter is incremented each time RTCOM com create request is successful. This does not mean that the TCP + SSL handshake is complete, and connection is UP.
`Number of Packet-Log RTCOM create req failed due to invalid param`	This counter is incremented each time RTCOM com create request is failed due to an invalid parameter.
`Number of Packet-Log RTCOM create req failed due to quota exceed`	This counter is incremented each time RTCOM com create request is failed due to lack of ports availability.
`Number of Packet-Log RTCOM create req failed due to status not ready`	This counter is incremented each time RTCOM com create request is failed due to IDP plugin is not registered with RTCOM or registration was failed.
`Number of Packet-Log RTCOM create done event received`	This counter is incremented each time IDP plugin receives an event CREATE DONE from RTCOM plugin. This event is basically received after the handshake was complete and connection is established to the host.

(Continued)

Field Name	Field Description
`Number of Packet-Log RTCOM peer close event received`	This counter is incremented each time IDP plugin receives an event PEER CLOSE from RTCOM plugin. This event is basically received if the SSL/TLS connection is closed by host due to some reason.
`Number of Packet-Log RTCOM conn abort event received`	This counter is incremented each time IDP plugin receives an event CONN ABORT from RTCOM plugin. This event is basically received if the SSL/TLS connection is aborted due to some reason.
`Number of Packet-Log RTCOM data rcvd event received`	This counter is incremented each time IDP plugin receives an event DATA RCVD from RTCOM plugin. This event is basically received if some data is received on the established connection from the host.
`Number of Packet-Log RTCOM connections established successfully`	This counter is incremented when the SSL/TLS connection is successful to the host. Basically, handshake is completed, and data can be send over the established connection.
`Number of Packet-Log RTCOM connections failed after retries`	This counter is incremented when SSL/TLS connection can not be established to the specified host even after 3 retries. Only 3 connection attempts are made to the configured host initially and on receiving CONN ABORT or PEER CLOSE event. If connection is no successful even after 3 attempts, no more further attempts are performed and this counter is incremented.
`Number of Packet-Log RTCOM connections closed successfully`	Number of SSL/TLS connections closed successfully. This counter is incremented only for such connections on which connection is up and gets closed later.
`Number of Packet-Log RTCOM data send successful`	This counter is incremented each time RTCOM send is successful. For a single packet log, it will be incremented multiple time as packet log is send as a chunk of 1024 bytes.
`Number of Packet-Log RTCOM data send failed`	This counter is incremented each time RTCOM send is failed.

(Continued)

Field Name	Field Description
Number of Packet-Log RTCOM contexts created	This counter is incremented each time a new RTCOM context structure is created for a specific host ip, host port, source ip, LSYS identifier, SSL identifier combination.
Number of Packet-Log RTCOM contexts reused	This counter is incremented each time a RTCOM context structure is reused for a specific host IP, host port, source IP, LSYS identifier, SSL identifier combination. This happens when attacks are getting detected simultaneously on a specific LSYS across different sessions. All the packet logs corresponding to those different sessions are sent using same SSL/TLS connection.
Number of Packet-Log RTCOM contexts destroyed	This counter is incremented each time a RTCOM context structure is destroyed. This happens only when no session or no packet log is referring that context to send packet logs.
Number of Packet-Log RTCOM packet log send successful	This counter is incremented each time a packet log is send successfully.
Number of Packet-Log RTCOM packet log send failed	This counter is incremented each time a packet log is send failed.

Sample Output

show security idp counters packet-log

```
user@host> show security idp counters packet-log
IDP counters:                                                                    Value
 Total packets captured since packet capture was activated              0
 Total sessions enabled since packet capture was activated              0
 Sessions currently enabled for packet capture                          0
 Packets currently captured for enabled sessions                        0
 Packet clone failures                                                  0
 Session log object failures                                            0
 Session packet log object failures                                     0
 Sessions skipped because session limit exceeded                        0
 Packets skipped because packet limit exceeded                          0
```

```
 Packets skipped because total memory limit exceeded                    0
 Packet log host route lookup failures                                  0
Number of Packet-Log RTCOM create req successful                       1
Number of Packet-Log RTCOM create req failed due to invalid param      0
Number of Packet-Log RTCOM create req failed due to quota exceed       0
Number of Packet-Log RTCOM create req failed due to status not ready   0
Number of Packet-Log RTCOM create done event received                  1
Number of Packet-Log RTCOM peer close event received                   0
Number of Packet-Log RTCOM conn abort event received                   0
Number of Packet-Log RTCOM data rcvd event received                    0
Number of Packet-Log RTCOM connections established successfully        1
Number of Packet-Log RTCOM connections failed after retries            0
Number of Packet-Log RTCOM connections closed successfully             1
Number of Packet-Log RTCOM data send successful                        20
Number of Packet-Log RTCOM data send failed                            0
Number of Packet-Log RTCOM contexts created                            1
Number of Packet-Log RTCOM contexts reused                             9
Number of Packet-Log RTCOM contexts destroyed                          1
Number of Packet-Log RTCOM packet log send successful                  10
Number of Packet-Log RTCOM packet log send failed                      0
```

show security idp counters packet-log logical-system LSYS1

```
user@host> show security idp counters packet-log logical-system LSYS1
IDP counters:

  IDP counter type                                                    Value
 No of file-decoder requests from MIME over HTTP                        0
 No of pending file-decoder requests from MIME over HTTP                0
 No of completd file-decoder requests from MIME over HTTP               0
 No of unrecognized file type from MIME over HTTP                       0
 No of compressed payload transferred over HTTP                         0
 No of bypassed files over HTTP                                         0
 Packet log host route lookup failures                                  0
 Number of Packet-Log RTCOM create req successful                       1
 Number of Packet-Log RTCOM create req failed due to invalid param      0
 Number of Packet-Log RTCOM create req failed due to quota exceed       0
 Number of Packet-Log RTCOM create req failed due to status not ready   0
 Number of Packet-Log RTCOM create done event received                  1
 Number of Packet-Log RTCOM peer close event received                   0
 Number of Packet-Log RTCOM conn abort event received                   0
```

```
 Number of Packet-Log RTCOM data rcvd event received                    0
 Number of Packet-Log RTCOM connections established successfully        1
 Number of Packet-Log RTCOM connections failed after retries            0
 Number of Packet-Log RTCOM connections closed successfully             1
 Number of Packet-Log RTCOM data send successful                        20
 Number of Packet-Log RTCOM data send failed                            0
 Number of Packet-Log RTCOM contexts created                            1
 Number of Packet-Log RTCOM contexts reused                             9
 Number of Packet-Log RTCOM contexts destroyed                          1
 Number of Packet-Log RTCOM packet log send successful                  10
 Number of Packet-Log RTCOM packet log send failed                      0

user@host> show security idp counters packet-log logical-system Tn1
IDP counters:

  IDP counter type                                                    Value
 Total packets captured since packet capture was activated              0
 Total sessions enabled since packet capture was activated              0
 Sessions currently enabled for packet capture                          0
 Packets currently captured for enabled sessions                        0
 Packet clone failures                                                  0
 Session log object failures                                            0
 Session packet log object failures                                     0
 Sessions skipped because session limit exceeded                        0
 Packets skipped because packet limit exceeded                          0
 Packets skipped because total memory limit exceeded                    0
Number of Packet-Log RTCOM create req successful                       1
 Number of Packet-Log RTCOM create req failed due to invalid param      0
 Number of Packet-Log RTCOM create req failed due to quota exceed       0
 Number of Packet-Log RTCOM create req failed due to status not ready   0
 Number of Packet-Log RTCOM create done event received                  1
 Number of Packet-Log RTCOM peer close event received                   0
 Number of Packet-Log RTCOM conn abort event received                   0
 Number of Packet-Log RTCOM data rcvd event received                    0
 Number of Packet-Log RTCOM connections established successfully        1
 Number of Packet-Log RTCOM connections failed after retries            0
 Number of Packet-Log RTCOM connections closed successfully             1
 Number of Packet-Log RTCOM data send successful                        20
 Number of Packet-Log RTCOM data send failed                            0
 Number of Packet-Log RTCOM contexts created                            1
 Number of Packet-Log RTCOM contexts reused                             9
 Number of Packet-Log RTCOM contexts destroyed                          1
```

```
 Number of Packet-Log RTCOM packet log send successful                  10
 Number of Packet-Log RTCOM packet log send failed                      0
```

show security idp counters packet-log tenant TSYS1

```
user@host> show security idp counters packet-log tenant TSYS1
IDP counters:

  IDP counter type                                                    Value
 Total packets captured since packet capture was activated              0
 Total sessions enabled since packet capture was activated              0
 Sessions currently enabled for packet capture                          0
 Packets currently captured for enabled sessions                        0
 Packet clone failures                                                  0
 Session log object failures                                            0
 Session packet log object failures                                     0
 Sessions skipped because session limit exceeded                        0
 Packets skipped because packet limit exceeded                          0
 Packets skipped because total memory limit exceeded                    0
Number of Packet-Log RTCOM create req successful                       1
 Number of Packet-Log RTCOM create req failed due to invalid param      0
 Number of Packet-Log RTCOM create req failed due to quota exceed       0
 Number of Packet-Log RTCOM create req failed due to status not ready   0
 Number of Packet-Log RTCOM create done event received                  1
 Number of Packet-Log RTCOM peer close event received                   0
 Number of Packet-Log RTCOM conn abort event received                   0
 Number of Packet-Log RTCOM data rcvd event received                    0
 Number of Packet-Log RTCOM connections established successfully        1
 Number of Packet-Log RTCOM connections failed after retries            0
 Number of Packet-Log RTCOM connections closed successfully             1
 Number of Packet-Log RTCOM data send successful                        20
 Number of Packet-Log RTCOM data send failed                            0
 Number of Packet-Log RTCOM contexts created                            1
 Number of Packet-Log RTCOM contexts reused                             9
 Number of Packet-Log RTCOM contexts destroyed                          1
 Number of Packet-Log RTCOM packet log send successful                  10
 Number of Packet-Log RTCOM packet log send failed                      0
```

show security idp counters packet-log tenant all

```
user@host> show security idp counters packet-log logical-system all
root@idpdevesx16-vsrx2-12> show security idp counters packet logical-system all



Logical System: root-logical-system
IDP counters:



IDP counter type Value
Processed packets 0
Dropped packets 0
Dropped ICMP packets 0
Dropped TCP packets 0
Dropped UDP packets 0
Dropped Other packets 0
Dropped by IDP Policy 0
Dropped by Error 0
Dropped sessions 0
Bad IP headers 0
Packets with IP options 0
Decapsulated packets 0
GRE decapsulations 0
PPP decapsulations 0
GTP decapsulations 0
GTP flows 0
TCP decompression uncompressed IP 0
TCP decompression compressed IP 0
Deferred-send packets 0
Ktimer entry optimized 0
TTL errors 0
Routing loops 0
STP drops 0
No-route packets 0
Flood IP 0
Invalid ethernet headers 0
Packets attached 0
IP Packet attach failed 0
```

```
Packets cloned 0
Packets allocated 0
Packets destructed 0
Packets destructed in pipeline 0
Packet data buffer allocated 0
Packet data buffer released 0
Buffer allocation on clone avoided 0
Late buffer allocation on clone 0
Distinct clone request 0
KPP clone buf cache allocated 0
KPP clone buf cache released 0
KPP clone buf cache used 0
KQMSG constructed 0
KQMSG destructed 0
KQMSG destructed in pipeline 0
jbuf copy failed 0
jbuf pullup failed 0
jbuf copy done 0
jbuf copy freed 0
jbuf copy reinjected 0



Logical System: LSYS0
IDP counters:



IDP counter type Value
Processed packets 0
Dropped packets 0
Dropped ICMP packets 0
Dropped TCP packets 0
Dropped UDP packets 0
Dropped Other packets 0
Dropped by IDP Policy 0
Dropped by Error 0
Dropped sessions 0
Bad IP headers 0
Packets with IP options 0
Decapsulated packets 0
GRE decapsulations 0
PPP decapsulations 0
```

```
GTP decapsulations 0
GTP flows 0
TCP decompression uncompressed IP 0
TCP decompression compressed IP 0
Deferred-send packets 0
Ktimer entry optimized 0
TTL errors 0
Routing loops 0
STP drops 0
No-route packets 0
Flood IP 0
Invalid ethernet headers 0
Packets attached 0
IP Packet attach failed 0
Packets cloned 0
Packets allocated 0
Packets destructed 0
Packets destructed in pipeline 0
Packet data buffer allocated 0
Packet data buffer released 0
Buffer allocation on clone avoided 0
Late buffer allocation on clone 0
Distinct clone request 0
KPP clone buf cache allocated 0
KPP clone buf cache released 0
KPP clone buf cache used 0
KQMSG constructed 0
KQMSG destructed 0
KQMSG destructed in pipeline 0
jbuf copy failed 0
jbuf pullup failed 0
jbuf copy done 0
jbuf copy freed 0
jbuf copy reinjected 0



Logical System: LSYS2
IDP counters:



IDP counter type Value
```

```
Processed packets 0
Dropped packets 0
Dropped ICMP packets 0
Dropped TCP packets 0
Dropped UDP packets 0
Dropped Other packets 0
Dropped by IDP Policy 0
Dropped by Error 0
Dropped sessions 0
Bad IP headers 0
Packets with IP options 0
Decapsulated packets 0
GRE decapsulations 0
PPP decapsulations 0
GTP decapsulations 0
GTP flows 0
TCP decompression uncompressed IP 0
TCP decompression compressed IP 0
Deferred-send packets 0
Ktimer entry optimized 0
TTL errors 0
Routing loops 0
STP drops 0
No-route packets 0
Flood IP 0
Invalid ethernet headers 0
Packets attached 0
IP Packet attach failed 0
Packets cloned 0
Packets allocated 0
Packets destructed 0
Packets destructed in pipeline 0
Packet data buffer allocated 0
Packet data buffer released 0
Buffer allocation on clone avoided 0
Late buffer allocation on clone 0
Distinct clone request 0
KPP clone buf cache allocated 0
KPP clone buf cache released 0
KPP clone buf cache used 0
KQMSG constructed 0
KQMSG destructed 0
KQMSG destructed in pipeline 0
```

```
jbuf copy failed 0
jbuf pullup failed 0
jbuf copy done 0
jbuf copy freed 0
jbuf copy reinjected 0



Logical System: LSYS1
IDP counters:

IDP counter type     Value
Processed packets     0
Dropped packets       0
Dropped ICMP packets  0
Dropped TCP packets   0
Dropped UDP packets   0
Dropped Other packets 0
Dropped by IDP Policy 0
Dropped by Error 0
Dropped sessions 0
Bad IP headers 0
Packets with IP options 0
Decapsulated packets 0
GRE decapsulations 0
PPP decapsulations 0
GTP decapsulations 0
GTP flows 0
TCP decompression uncompressed IP 0
TCP decompression compressed IP 0
Deferred-send packets 0
Ktimer entry optimized 0
TTL errors 0
Routing loops 0
STP drops 0
No-route packets 0
Flood IP 0
Invalid ethernet headers 0
Packets attached 0
IP Packet attach failed 0
Packets cloned 0
Packets allocated 0
Packets destructed 0
```

```
Packets destructed in pipeline 0
Packet data buffer allocated 0
Packet data buffer released 0
Buffer allocation on clone avoided 0
Late buffer allocation on clone 0
Distinct clone request 0
KPP clone buf cache allocated 0
KPP clone buf cache released 0
KPP clone buf cache used 0
KQMSG constructed 0
KQMSG destructed 0
KQMSG destructed in pipeline 0
jbuf copy failed 0
jbuf pullup failed 0
jbuf copy done 0
jbuf copy freed 0
jbuf copy reinjected 0

IDP counters:
LSYS1:
  IDP counter type                                                    Value
 No of file-decoder requests from MIME over HTTP                       0
 No of pending file-decoder requests from MIME over HTTP               0
 No of completd file-decoder requests from MIME over HTTP              0
 No of unrecognized file type from MIME over HTTP                      0
 No of compressed payload transferred over HTTP                        0
 No of bypassed files over HTTP                                        0
 Packet log host route lookup failures                                 0
 Number of Packet-Log RTCOM create req successful                      1
 Number of Packet-Log RTCOM create req failed due to invalid param     0
 Number of Packet-Log RTCOM create req failed due to quota exceed      0
 Number of Packet-Log RTCOM create req failed due to status not ready  0
 Number of Packet-Log RTCOM create done event received                 1
 Number of Packet-Log RTCOM peer close event received                  0
 Number of Packet-Log RTCOM conn abort event received                  0
 Number of Packet-Log RTCOM data rcvd event received                   0
 Number of Packet-Log RTCOM connections established successfully       1
 Number of Packet-Log RTCOM connections failed after retries           0
 Number of Packet-Log RTCOM connections closed successfully            1
 Number of Packet-Log RTCOM data send successful                       20
 Number of Packet-Log RTCOM data send failed                           0
 Number of Packet-Log RTCOM contexts created                           1
 Number of Packet-Log RTCOM contexts reused                            9
```

```
Number of Packet-Log RTCOM contexts destroyed                          1
Number of Packet-Log RTCOM packet log send successful                  10
Number of Packet-Log RTCOM packet log send failed                      0
```

Release Information

Command introduced in Junos OS Release 10.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

IDP Security Packet Capture

clear security idp counters packet-log | 3974

show security idp counters pdf-decoder

IN THIS SECTION

- Syntax | 24582
- Description | 24582
- Options | 24582
- Required Privilege Level | 24582
- Output Fields | 24582
- Sample Output | 24583
- Release Information | 24585

Syntax

```
show security idp counters pdf-decoder
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the IDP counters value for PDF-Decode enabled sessions, requests, and memory limit IDP counter types.

Options

none	Displays the IDP counters value for PDF-Decode enabled sessions, requests, and memory limit IDP counter types.
logical-system *logical-system-name*	(Optional) Displays the IDP counters value for PDF-Decode enabled sessions, requests, and memory limit IDP counter types for a specific logical system.
logical-system all	(Optional) Displays the IDP counters value for PDF-Decode enabled sessions, requests, and memory limit IDP counter types for all logical systems.
tenant *tenant-name*	(Optional) Displays the IDP counters value for PDF-Decode enabled sessions, requests, and memory limit IDP counter types for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1241 on page 24583 lists the output fields for the `show security idp counters pdf-decoder` command. Output fields are listed in the approximate order in which they appear.

Table 1241: show security idp counters pdf-decoder Output Fields

Field Name	Field Description
IDP counter type	Name of the action
Value	Number of PDF-decode enabled sessions and requests for the IDP counters.

Sample Output

show security idp counters pdf-decoder

```
user@host>  show security idp counters pdf-decoder
IDP counters:

  IDP counter type                                                    Value
 PDF-Decode enabled sessions                                           0
 PDF-Decode requests                                                   0
 PDF-decode in pending state                                           0
 PDF-decode finished sucessfully                                       0
 PDF-decode per file memory limit reached                              0
 PDF-decode memory limit reached                                       0
 PDF-decode session limit reached                                      0
 PDF-decode malloc count                                               0
 PDF-decode free count                                                 0
 PDF-decode bypassed  -  Global memory limit                           0
 PDF-decode bypassed  -  Per file memory limit                         0
 PDF-decode bypassed  -  Document encrypted                            0
 PDF-decode bypassed  -  Unsupported filter decode                     0
 PDF-decode bypassed  -  corrupted PDF file                            0
 PDF-decode bypassed  -  unsupported obj-stream length                 0
 PDF-decode bypassed  -  filter decode failed                          0
```

show security idp counters pdf-decoder logical-system LSYS1

```
user@host>  show security idp counters pdf-decoder logical-system LSYS1
IDP counters:
```

```
  IDP counter type                                                    Value
 PDF-Decode enabled sessions                                           0
 PDF-Decode requests                                                   0
 PDF-decode in pending state                                           0
 PDF-decode finished sucessfully                                       0
 PDF-decode per file memory limit reached                              0
 PDF-decode memory limit reached                                       0
 PDF-decode session limit reached                                      0
 PDF-decode malloc count                                               0
 PDF-decode free count                                                 0
 PDF-decode bypassed  -  Global memory limit                           0
 PDF-decode bypassed  -  Per file memory limit                         0
 PDF-decode bypassed  -  Document encrypted                            0
 PDF-decode bypassed  -  Unsupported filter decode                     0
 PDF-decode bypassed  -  corrupted PDF file                            0
 PDF-decode bypassed  -  unsupported obj-stream length                 0
 PDF-decode bypassed  -  filter decode failed                          0
```

show security idp counters pdf-decoder tenant TSYS1

```
user@host>  show security idp counters pdf-decoder tenant TSYS1
IDP counters:

  IDP counter type                                                    Value
 PDF-Decode enabled sessions                                           0
 PDF-Decode requests                                                   0
 PDF-decode in pending state                                           0
 PDF-decode finished sucessfully                                       0
 PDF-decode per file memory limit reached                              0
 PDF-decode memory limit reached                                       0
 PDF-decode session limit reached                                      0
 PDF-decode malloc count                                               0
 PDF-decode free count                                                 0
 PDF-decode bypassed  -  Global memory limit                           0
 PDF-decode bypassed  -  Per file memory limit                         0
 PDF-decode bypassed  -  Document encrypted                            0
 PDF-decode bypassed  -  Unsupported filter decode                     0
 PDF-decode bypassed  -  corrupted PDF file                            0
 PDF-decode bypassed  -  unsupported obj-stream length                 0
 PDF-decode bypassed  -  filter decode failed                          0
```

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters pdf-decoder | 3975

show security idp counters policy-manager

IN THIS SECTION

- Syntax | 24585
- Description | 24585
- Options | 24586
- Required Privilege Level | 24586
- Output Fields | 24586
- Sample Output | 24586
- Release Information | 24587

Syntax

```
show security idp counters policy-manager
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all IDP policies counter values.

Options

none	Displays the status of all IDP policies counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all IDP policies counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all IDP policies counter values for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all IDP policies counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1242 on page 24586 lists the output fields for the `show security idp counters policy-manager` command. Output fields are listed in the approximate order in which they appear.

Table 1242: show security idp counters policy-manager Output Fields

Field Name	Field Description
`Number of policies`	Number of policies installed.
`Number of aged out policies`	Number of IDP policies that are expired.

Sample Output

show security idp counters policy-manager

```
user@host> show security idp counters policy-manager
IDP counters:
 IDP counter type                                                    Value
```

```
 Number of policies                                                    0
 Number of aged out policies                                           0
```

show security idp counters policy-manager logical-system LSYS1

```
user@host> show security idp counters policy-manager logical-system LSYS1
IDP counters:

  IDP counter type                                                   Value
 Number of policies                                                   1
 Number of aged out policies                                          0
 Policy compile failure due to memory                                 0
```

show security idp counters policy-manager tenant TSYS1

```
user@host> show security idp counters policy-manager tenant TSYS1
IDP counters:

  IDP counter type                                                   Value
 Number of policies                                                   0
 Number of aged out policies                                          0
 Policy compile failure due to memory                                 0
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters policy-manager | 3977

show security idp counters ssl-inspection

IN THIS SECTION

- Syntax | 24588
- Description | 24588
- Options | 24588
- Required Privilege Level | 24589
- Output Fields | 24589
- Sample Output | 24589
- Release Information | 24591

Syntax

```
show security idp counters ssl-inspection
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the IDP counters value for decrypted and encrypted sessions.

Options

none	Displays the IDP counters value for decrypted and encrypted sessions.
logical-system ***logical-system-name***	(Optional) Displays the IDP counters value for decrypted and encrypted sessions for a specific logical system.
logical-system all	(Optional) Displays the IDP counters value for decrypted and encrypted sessions for all logical systems.
tenant ***tenant-name***	(Optional) Displays the IDP counters value for decrypted and encrypted sessions for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1243 on page 24589 lists the output fields for the `show security idp counters ssl-inspection` command. Output fields are listed in the approximate order in which they appear.

Table 1243: show security idp counters ssl-inspection Output Fields

Field Name	Field Description
IDP counter type	Name of the action
Value	Number of packets and sessions decrypted, sessions not decrypted.

Sample Output

show security idp counters ssl-inspection

```
user@host>  show security idp counters ssl-inspection
IDP counters:

  IDP counter type                                                    Value
 Packets Decrypted                                                     0
 Sessions Decrypted                                                    0
 Sessoins Not Decrypted                                                0
 Sessions Not Decrypted - Configuration                                0
 Sessions Not Decrypted - No Key                                       0
 Sessions Not Decrypted - Unsupported Ciphers                          0
 Sessions Not Decrypted - Unsupported Compression                      0
 Sessions Not Decrypted - Unsupported Key Exchange                     0
 Sessions Not Decrypted - Bulk Decryption Failure                      0
 Sessions Not Decrypted - Key Generation Failure                       0
 Sessions Not Decrypted - Temporary Certificate                        0
 Sessions Not Decrypted - Handshake Verification Failure               0
 Sessions Not Decrypted - ID Cache Miss                                0
 Sessions Not Decrypted - Session Limit                                0
```

```
 Sessions Not Decrypted - Message Size                                  0
 Sessions Not Decrypted - No Memory                                     0
 Sessions New Key                                                       0
 Sessions Used Key                                                      0
 Session ID Cache Hits                                                  0
 Session ID Cache Misses                                                0
 Sessions Used XLR RSA SAE for Key Decryption                           0
 Sessions - Error when XLR RSA SAE used                                 0
```

show security idp counters ssl-inspection logical-system LSYS1

```
user@host>  show security idp counters ssl-inspection logical-system LSYS1
IDP counters:

  IDP counter type                                                    Value
 Packets Decrypted                                                      0
 Sessions Decrypted                                                     0
 Sessoins Not Decrypted                                                 0
 Sessions Not Decrypted - Configuration                                 0
 Sessions Not Decrypted - No Key                                        0
 Sessions Not Decrypted - Unsupported Ciphers                           0
 Sessions Not Decrypted - Unsupported Compression                       0
 Sessions Not Decrypted - Unsupported Key Exchange                      0
 Sessions Not Decrypted - Bulk Decryption Failure                       0
 Sessions Not Decrypted - Key Generation Failure                        0
 Sessions Not Decrypted - Temporary Certificate                         0
 Sessions Not Decrypted - Handshake Verification Failure                0
 Sessions Not Decrypted - ID Cache Miss                                 0
 Sessions Not Decrypted - Session Limit                                 0
 Sessions Not Decrypted - Message Size                                  0
 Sessions Not Decrypted - No Memory                                     0
 Sessions New Key                                                       0
 Sessions Used Key                                                      0
 Session ID Cache Hits                                                  0
 Session ID Cache Misses                                                0
 Sessions Used XLR RSA SAE for Key Decryption                           0
 Sessions - Error when XLR RSA SAE used                                 0
```

show security idp counters ssl-inspection tenant TSYS1

```
user@host> show security idp counters ssl-inspection tenant TSYS1
IDP counters:

  IDP counter type                                                    Value
 Packets Decrypted                                                     0
 Sessions Decrypted                                                    0
 Sessoins Not Decrypted                                                0
 Sessions Not Decrypted - Configuration                                0
 Sessions Not Decrypted - No Key                                       0
 Sessions Not Decrypted - Unsupported Ciphers                          0
 Sessions Not Decrypted - Unsupported Compression                      0
 Sessions Not Decrypted - Unsupported Key Exchange                     0
 Sessions Not Decrypted - Bulk Decryption Failure                      0
 Sessions Not Decrypted - Key Generation Failure                       0
 Sessions Not Decrypted - Temporary Certificate                        0
 Sessions Not Decrypted - Handshake Verification Failure               0
 Sessions Not Decrypted - ID Cache Miss                                0
 Sessions Not Decrypted - Session Limit                                0
 Sessions Not Decrypted - Message Size                                 0
 Sessions Not Decrypted - No Memory                                    0
 Sessions New Key                                                      0
 Sessions Used Key                                                     0
 Session ID Cache Hits                                                 0
 Session ID Cache Misses                                               0
 Sessions Used XLR RSA SAE for Key Decryption                          0
 Sessions - Error when XLR RSA SAE used                                0
```

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters ssl-inspection | 3978

show security idp counters tcp-reassembler

IN THIS SECTION

- Syntax | 24592
- Description | 24592
- Options | 24592
- Required Privilege Level | 24593
- Output Fields | 24593
- Sample Output | 24596
- Release Information | 24598

Syntax

```
show security idp counters tcp-reassembler
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all TCP reassembler counter values.

NOTE: On SRX Series Firewalls with IDP enabled, if IDP attacks are configured for a single direction (server or client), a flow in the opposite direction does not need IDP processing. For TCP traffic, the TCP optimization feature ensures minimal processing for these flows without running into reassembly errors.

Options

none	Displays the status of all TCP reassembler counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all TCP reassembler counter values for a specific logical system.

logical-system all (Optional) Displays the status of all TCP reassembler counter values for all logical systems.

tenant *tenant-name* (Optional) Displays the status of all TCP reassembler counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1244 on page 24593 lists the output fields for the `show security idp counters tcp-reassembler` command. Output fields are listed in the approximate order in which they appear.

Table 1244: show security idp counters tcp-reassembler Output Fields

Field Name	Field Description
`Bad TCP checksums` (Unsupported)	Number of packets that have incorrect TCP checksums.
`Bad TCP headers`	Number of bad TCP headers detected.
`Slow path segments`	Number of segments that are sent through the slow path if the TCP segment does not pass fast-path segment validation.
`Fast path segments`	Number of segments that are sent through the fast path after passing a predefined TCP validation sequence.
`Tcp Optimized s2c segments`	Number of TCP segments that are sent through optimized re-assembly process from server to client.
`Tcp Optimized c2s segments`	Number of TCP segments that are sent through optimized re-assembly process from server to client.

Table 1244: show security idp counters tcp-reassembler Output Fields *(Continued)*

Field Name	Field Description
`Sequence number wrap around errors`	Number of packets that wrap around of the sequence number.
`Session reuses`	Number of sessions that reused an already established TCP session.
`SYN retransmissions`	Number of SYN packets that are retransmitted.
`Bad three way handshake acknowledgements`	Number of packets that have incorrect three-way handshake acknowledgements (ACK packet).
`Sequence number out of sync flows`	Number of packets that have out-of-sync sequence numbers.
`Fast path pattern matches in queued up streams`	Number of queued packets that have fast path pattern match.
`New segments with no overlaps with old segment`	Number of new segments that do not overlap with old segment.
`New segment overlaps with beginning of old segment`	Number of new segments that overlap with beginning of old segment.
`New segment overlaps completely with old segment`	Number of new segments that overlap completely with old segment.
`New segment is contained in old segment`	Number of new segments contained in old segment.
`New segment overlaps with end of old segment`	Number of new segments that overlap with the end of old segment.

Table 1244: show security idp counters tcp-reassembler Output Fields *(Continued)*

Field Name	Field Description
`New segment begins after end of old segment`	Number of new segments that overlap after the end of old segment.
`Memory consumed by new segment`	Memory that is consumed by the new segment.
`Peak memory consumed by new segments`	Peak memory that is consumed by the new segment.
`Segments in memory`	Number of segments that are stored in memory for processing.
`Per-flow memory overflows`	Number of segments dropped after reaching per flow memory limit.
`Global memory overflows`	Number of segments dropped after reaching reassembler global memory limit.
`Overflow drops`	Number of packets that are dropped due to memory overflow.
`Copied packets` (Unsupported)	Number of packets copied in reassembler.
`Closed Acks`	Number of Ack packets seen without having seen SYN on the same session.
`Ack Validation failures`	Number of Invalid ACKs received from server during 3-way handshake.
`Simultanious syn`	Number of simultaneous syn packets seen.
`C2S synack`	Number of C2S Syn/Ack packets seen.
`Segment to left of receiver window`	Number of segments falling left of receive window.

Table 1244: show security idp counters tcp-reassembler Output Fields *(Continued)*

Field Name	Field Description
Segment to right of receiver window	Number of segments falling right of receive window.
SYN seen in the window	Number of Syn packets seen after connection establishment.
ACK bit is off	Number of packets seen without ACK after connection establishment.
Unexpected FIN	Number of unexpected FIN packets seen.
Duplicate Syn/Ack with different SEQ	Number of Syn/Ack packets with different SEQ numbers.

Sample Output

show security idp counters tcp-reassembler

```
user@host> show security idp counters tcp-reassembler
 IDP counters:

  IDP counter type                                                    Value
 Bad TCP checksums                                                     0
 Bad TCP headers                                                       0
 Slow path segments                                                    90
 Fast path segments                                                    7099
 Tcp Optimized s2c segments                                            0
 Tcp Optimized c2s segments                                            0
 Sequence number wrap around errors                                    0
 Session reuses                                                        0
 SYN retransmissions                                                   0
 Bad three way handshake acknowledgements                              0
 Sequence number out of sync flows                                     0
 Fast path pattern matches in queued up streams                        0
 New segments with no overlaps with old segment                        0
 New segment overlaps with beginning of old segment                    0
```

```
 New segment overlaps completely with old segment                        0
 New segment is contained in old segment                                 0
 New segment overlaps with end of old segment                            0
 New segment begins after end of old segment                             3
 Memory consumed by new segment                                          0
 Peak memory consumed by new segments                                    3821
 Segments in memory                                                      0
 Per-flow memory overflows                                               0
 Global memory overflows                                                 0
 Overflow drops                                                          0
 Copied packets                                                          0
 Closed Acks                                                             3
 Ack Validation failure                                                  0
 Simultanious syn                                                        0
 C2S synack                                                              0
 segment to left of receiver window                                      0
 segment to right of receiver window                                     0
 SYN seen in the window                                                  0
 ACK bit is off                                                          0
 Unexpected FIN                                                          0
 Duplicate Syn/Ack with different SEQ                                    0
```

show security idp counters tcp-reassembler logical-system LSYS1

```
user@host> show security idp counters tcp-reassembler logical-system LSYS1
IDP counters:

  IDP counter type                                                     Value
 Bad TCP checksums                                                      0
 Bad TCP headers                                                        0
 Slow path segments                                                     37
 Fast path segments                                                     27
 Tcp Optimized s2c segments                                             0
 Tcp Optimized c2s segments                                             0
 Sequence number wrap around errors                                     0
 Session reuses                                                         0
 SYN retransmissions                                                    0
 Bad three way handshake acknowledgements                               0
 Sequence number out of sync flows                                      0
 Fast path pattern matches in queued up streams                         0
 New segments with no overlaps with old segment                         0
```

```
New segment overlaps with beginning of old segment                    0
New segment overlaps completely with old segment                      0
New segment is contained in old segment                               0
New segment overlaps with end of old segment                          0
New segment begins after end of old segment                           0
Memory consumed by new segment                                        0
Peak memory consumed by new segments                                  2021
Segments in memory                                                    0
Per-flow memory overflows                                             0
Global memory overflows                                               0
Overflow drops                                                        0
Overflow drops - missing packets                                      0
Copied packets                                                        0
Closed Acks                                                           0
Ack Validation failure                                                0
Simultanious syn                                                      0
C2S synack                                                            0
segment to left of receiver window                                    0
segment to right of receiver window                                   0
SYN seen in the window                                                0
ACK bit is off                                                        0
Unexpected FIN                                                        0
Duplicate Syn/Ack with different SEQ                                  0
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

re-assembler | 14079

clear security idp counters tcp-reassembler | 3980

show security idp jist-package-version

IN THIS SECTION

- Syntax | 24599
- Description | 24599
- Options | 24599
- Required Privilege Level | 24599
- Output Fields | 24599
- Sample Output | 24600
- Release Information | 24600

Syntax

```
show security idp jist-package-version
```

Description

Displays information of the currently installed JIST package version.

Options

none Displays information of the currently installed JIST package version.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

show security idp jist-package-version

```
user@host> show security idp jist-package-version
```

```
JIST package version: 3.1
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

request security idp jist-conversion | 14751

show security idp logical system

IN THIS SECTION

- Syntax | 24600
- Description | 24601
- Options | 24601
- Required Privilege Level | 24601
- Sample Output | 24601
- Release Information | 24601

Syntax

```
show security idp  logical-system
```

Description

Display information about the logical systems and the IDP policies associated to the logical systems.

Options

- `logical-system`— Show the IDP policy assigned to a logical system. The IDP policy is assigned to a logical system through the security profile.
  - `policy-association`

Required Privilege Level

view

Sample Output

show security idp logical-system policy-association

```
user@host> show security idp logical-system policy-association
Logical system                  IDP policy

root-logical-system             policy1
LSYS1                           idpengine
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

clear security idp | 3957

show security idp logical-system policy-association

IN THIS SECTION

- Syntax | 24602
- Description | 24602
- Required Privilege Level | 24602
- Output Fields | 24602
- Sample Output | 24603
- Release Information | 24603

Syntax

```
show security idp logical-system policy-association
```

Description

Display the IDP policy assigned to a logical system. The IDP policy is assigned to a logical system through the security profile.

Required Privilege Level

view

Output Fields

Table 1245 on page 24603 lists the output fields for the `show security idp logical-system policy-association` command.

Table 1245: show security idp logical-system policy-association Output Fields

Field Name	Field Description
`Logical system`	Name of the logical system to which an IDP policy is assigned.
`IDP policy`	Name of the IDP policy that is specified in the security profile that is bound to the logical system.

Sample Output

show security idp logical-system policy-association

```
user@host> show security idp logical-system policy-association
Logical system         IDP policy
root-logical-system    idp-policy1
lsys1                  idp-policy2
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

security-profile | 16507

show security idp memory

IN THIS SECTION

- Syntax | 24604
- Description | 24604

- Required Privilege Level | 24604
- Output Fields | 24604
- Sample Output | 24605
- Release Information | 24605

Syntax

```
show security idp memory
```

Description

Display the status of all IDP data plane memory.

Required Privilege Level

view

Output Fields

Table 1246 on page 24604 lists the output fields for the `show security idp memory` command. Output fields are listed in the approximate order in which they appear.

Table 1246: show security idp memory Output Fields

Field Name	Field Description
`PIC`	Name of the PIC.
`Total IDP data plane memory`	Total memory space that is allocated for the IDP data plane. **NOTE**: IDP requires a minimum of 5 MB of memory for session inspection.
`Used`	Used memory space in the data plane.

Table 1246: show security idp memory Output Fields *(Continued)*

Field Name	Field Description
`Available`	Available memory space in the data plane.

Sample Output

show security idp memory

```
user@host> show security idp memory
    IDP data plane memory statistics:
                          PIC : FPC 0 PIC 0:
  Total IDP data plane memory : 196 MB
                         Used : 8 MB ( 8192 KB ) ( 4.08% )
                    Available : 188 MB ( 192512 KB ) (95.91%)
```

Release Information

Command introduced in Junos OS Release 9.2. Percentage outputs added in Junos OS Release 10.1.

show security idp policies

IN THIS SECTION

- Syntax | 24606
- Description | 24606
- Options | 24606
- Required Privilege Level | 24606
- Output Fields | 24606
- Sample Output | 24606
- Release Information | 24607

Syntax

```
show security idp policies
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the list of currently installed policies.

Options

none	Displays the list of currently installed policies.
logical-system ***logical-system-name***	(Optional) Displays the list of currently installed policies for a specific logical system.
logical-system all	(Optional) Displays the list of currently installed policies for all logical systems.
tenant ***tenant-name***	(Optional) Displays the list of currently installed policies for a specific tenant system.

Required Privilege Level

view

Output Fields

Sample Output

show security idp policies

```
user@host>show security idp policies
PIC : FPC 0 PIC 0:
  ID    Name                    Sessions    Memory      Detector
      0        idp-policy-unified         0        10179      12.6.130180509
```

show security idp policies logical-system LSYS0

```
user@host> show security idp policies logical-system LSYS0
 PIC : FPC 0 PIC 0:
ID    Name                    Sessions    Memory      Detector
53      idp_one policy             0           189712      12.6.130180509
```

Release Information

Command introduced in Junos OS Release 10.1.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp active-policy | 24478

show security idp policy-commit-status

IN THIS SECTION

- Syntax | 24608
- Description | 24608
- Options | 24608
- Required Privilege Level | 24608
- Output Fields | 24608
- Sample Output | 24608
- Sample Output | 24609
- Sample Output | 24609
- Release Information | 24609

Syntax

```
show security idp policy-commit-status
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the IDP policy commit status. For example, status of policy compilation or load.

Options

none	Displays the IDP policy commit status.
logical-system ***logical-system-name***	(Optional) Displays the IDP policy commit status for a specific logical system.
logical-system all	(Optional) Displays the IDP policy commit status for all logical systems.
tenant ***tenant-name***	(Optional) Displays the IDP policy commit status for a specific tenant system.

Required Privilege Level

view

Output Fields

Sample Output

show security idp policy-commit-status

```
user@host> show security idp policy-commit-status
IDP policy[/var/db/idpd/bins/test.bin.gz.v] and detector[/var/db/idpd/sec-repository/installed-
detector/libidp-detector.so.tgz.v] loaded successfully.

The loaded policy size is:45583070 Bytes
```

Sample Output

show security idp policy-commit-status (on vSRX Virtual Firewall when you configure dynamic attack groups filters)

```
user@host> show security idp policy-commit-status
 Last good policy file does not exist. Aborted
```

Sample Output

show security idp policy-commit-status logical-system LSYS1

```
user@host> show security idp policy-commit-status logical-system LSYS1
 IDP policy[/var/db/idpd/bins//idp-policy-combined.bin.gz.v] and detector[/var/db/idpd/sec-
repository/installed-detector/libidp-detector.so.tgz.v] loaded successfully.
 The loaded policy size is:7416 Bytes
```

Release Information

Command introduced in JUNOS OS Release 10.4.

Starting with Junos OS Release 12.3X48-D15 and Junos OS Release 17.3R1, a new pattern matching engine is introduced for the SRX Series IDP feature. This scanning mechanism helps improve performance and policy loading. The new engine is 9.223 times faster than the existing DFA engine.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp status | 24622

show security idp policy-commit-status clear | 24610

show security idp policy-commit-status clear

IN THIS SECTION

- Syntax | 24610
- Description | 24610
- Options | 24610
- Required Privilege Level | 24610
- Output Fields | 24611
- Release Information | 24611

Syntax

```
show security idp policy-commit-status clear
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Clears the IDP policy commit status.

Options

none	Clears the IDP policy commit status.
logical-system ***logical-system-name***	(Optional) Clears the IDP policy commit status for a specific logical system.
logical-system all	(Optional) Clears the IDP policy commit status for all logical systems.
tenant ***tenant-name***	(Optional) Clears the IDP policy commit status for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 10.4.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp policy-commit-status | 24607

show security idp policy-templates-list

IN THIS SECTION

- Syntax | 24611
- Description | 24612
- Required Privilege Level | 24612
- Output Fields | 24612
- Sample Output | 24612
- Release Information | 24612

Syntax

```
show security idp policy-templates-list
```

Description

This command displays the list of available policy templates for logical systems.

Required Privilege Level

view

Output Fields

Sample Output

show security idp policy-templates-list

```
user@host>show security idp policy-templates-list
 Web_Server
DMZ_Services
DNS_Service
File_Server
Getting_Started
IDP_Default
Server-Protection
Server-Protection-1G
Client-Protection
Client-Protection-1G
Client-And-Server-Protection
Client-And-Server-Protection-1G
Recommended
```

Release Information

Command introduced in Junos OS Release 10.1.

Command introduced for user logical system in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

show security idp active-policy | 24478

show security idp predefined-attacks

IN THIS SECTION

- Syntax | 24613
- Description | 24613
- Options | 24613
- Required Privilege Level | 24614
- Output Fields | 24614
- Sample Output | 24614
- Release Information | 24615

Syntax

```
show security idp predefined-attacks
filters ( category | severity | direction)
```

Description

Display information about predefined attacks using optional filters.

Options

filters (Optional)

- `category`—Show predefined attacks in different categories.
- `severity`—Show predefined attacks based on different severities.
  - `critical`
  - `info`
  - `major`
  - `minor`

  - warning

- direction — Show predefined attacks for different directions.
  - any
  - client-to-server
  - exclude-any
  - exclude-client-to-server
  - exclude-server-to-client
  - server-to-client

Required Privilege Level

view

Output Fields

```
user@host> show security idp predefined-attacks filters category APP
```

Sample Output

command-name

```
APP:AMANDA:AMANDA-ROOT-OF1
APP:AMANDA:AMANDA-ROOT-OF2
APP:ARKEIA:TYPE-77-OF
APP:CA:ALERT-SRV-OF
APP:CA:ARCSRV:TCP-BOF
APP:CA:ARCSRV:UA-OF
APP:CA:IGATEWAY-BOF
APP:CA:LIC-COMMAND-OF
APP:CA:LIC-GCR-OF
APP:CA:LIC-GETCONFIG-OF
APP:CA:LIC-GETCONFIG-OF2
APP:CA:LIC-PUTOLF-OF
APP:CDE-DTSPCD-OF
APP:DOUBLETAKE
APP:ETHEREAL:DISTCC-OF
```

```
APP:HPOVNNM:HPOVTRACE-OF
APP:KERBEROS:GSS-ZERO-TOKEN
APP:KERBEROS:KBR-DOS-TCP-2
APP:MDAEMON:FORM2RAW-OF
APP:MERCURY-BOF
APP:MISC:MCAFFEE-SRV-HDR
APP:NTOP-WEB-FS1
APP:PPTP:MICROSOFT-PPTP
APP:REMOTE:TIMBUKTU-AUTH-OF
```

```
user@host> show security idp security-package predefined-attacks filters category FTP severity critical direction
client-to-server
```

```
FTP:COMMAND:WZ-SITE-EXEC
FTP:DIRECTORY:TILDE-ROOT
FTP:EXPLOIT:OPENFTPD-MSG-FS
FTP:OVERFLOW:OPENBSD-FTPD-GLOB
FTP:OVERFLOW:PATH-LINUX-X86-3
FTP:OVERFLOW:WFTPD-MKD-OVERFLOW
FTP:OVERFLOW:WUBSD-SE-RACE
FTP:PROFTP:OVERFLOW1
FTP:PROFTP:PPC-FS2
FTP:SERVU:CHMOD-OVERFLOW
FTP:SERVU:LIST-OVERFLOW
FTP:SERVU:MDTM-OVERFLOW
FTP:WU-FTP:IREPLY-FS
```

Release Information

Command introduced in Junos OS Release 10.1.

show security idp security-package-version

IN THIS SECTION

Syntax | 24616

- Description | 24616
- Options | 24616
- Required Privilege Level | 24616
- Output Fields | 24617
- Sample Output | 24617
- Release Information | 24618

Syntax

```
show security idp security-package-version
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays information of the currently installed security package version and detector version.

Options

none	Displays information of the currently installed security package version and detector version.
logical-system ***logical-system-name***	(Optional) Displays information of the currently installed security package version and detector version for a specific logical system.
logical-system all	(Optional) Displays information of the currently installed security package version and detector version for all logical systems.
tenant ***tenant-name***	(Optional) Displays information of the currently installed security package version and detector version for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 1247 on page 24617 lists the output fields for the `show security idp security-package-version` command. Output fields are listed in the approximate order in which they appear.

Table 1247: show security idp security-package-version Output Fields

Field Name	Field Description
`Attack database version`	Attack database version number that is currently installed on the system.
`Detector version`	Detector version number that is currently installed on the system.
`Policy template version`	Policy template version number that is currently installed on the system.
`Rollback version`	(SRX) Rollback version number on the system.
`Rollback Detector version`	Rollback Detector version number on the system.

Sample Output

show security idp security-package-version (Output After Initial Bootup)

```
user@host> show security idp security-package-version
Attack database version:N/A(N/A)
Detector version :N/A
Policy template version :N/A
Rollback Attack database version :N/A(N/A)
Rollback Detector version : N/A
```

show security idp security-package-version (Output After Security Package is Installed)

```
user@host> show security idp security-package-version
 Attack database version :3525(Tue Sep 13 17:40:07 2022 UTC)
  Detector version :12.6.180220128
  Policy template version :N/A
```

```
  Rollback Attack database version :N/A(N/A)
  Rollback Detector version :N/A
```

show security idp security-package-version

```
user@host> show security idp security-package-version
  Attack database version :3533(Mon Oct 10 14:57:03 2022 UTC)
  Detector version :12.6.180220128
  Policy template version :3521
  Rollback Attack database version :3527(Tue Sep 20 14:21:36 2022 UTC)
  Rollback Detector version :12.6.180220128
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

- security-package | 16505
- request security idp security-package download | 14754
- request security idp security-package install | 14757

show security idp ssl-inspection key

IN THIS SECTION

- Syntax | 24619
- Description | 24619
- Options | 24619

- Required Privilege Level | 24619
- Output Fields | 24619
- Sample Output | 24620
- Sample Output | 24620
- Release Information | 24620

Syntax

```
show security idp ssl-inspection key [<key-name> [server <server-ip>]]
```

Description

Display SSL keys added to the system along with their associated server IP addresses.

Options

- *key-name* —(Optional) Name of SSL private key.
- server *server-ip* —(Optional) Server IP address associated for specified key.

Required Privilege Level

view

Output Fields

Table 1248 on page 24619 lists the output fields for the `show security idp ssl-inspection key` command. Output fields are listed in the approximate order in which they appear.

Table 1248: show security idp ssl-inspection key Output Fields

Field Name	Field Description
Total SSL keys	Total number of SSL keys.

Table 1248: show security idp ssl-inspection key Output Fields *(Continued)*

Field Name	Field Description
key	Name of the SSL private key.
server	Server IP address associated with the SSL keys.

Sample Output

show security idp ssl-inspection key

```
user@host> show security idp ssl-inspection key
Total SSL keys : 4

SSL Server key and ip address:

  Key : key1, server : 1.1.0.1
  Key : key1, server : 1.1.0.2
  Key : key2, server : 2.2.0.1
  key : key3
```

Sample Output

show security idp ssl-inspection key key2

```
user@host> show security idp ssl-inspection key key2
SSL Server key and ip address:

  Key : key2, server : 2.2.0.1
```

Release Information

Command introduced in Junos OS Release 9.3.

show security idp ssl-inspection session-id-cache

IN THIS SECTION

- Syntax | 24621
- Description | 24621
- Required Privilege Level | 24621
- Output Fields | 24621
- Sample Output | 24622
- Release Information | 24622

Syntax

```
show security idp ssl-inspection session-id-cache
```

Description

Display all the SSL session IDs in the session ID cache. Each cache entry is 32 bytes long.

Required Privilege Level

view

Output Fields

Table 1249 on page 24621 lists the output fields for the `show security idp ssl-inspection session-id-cache` command. Output fields are listed in the approximate order in which they appear.

Table 1249: show security idp ssl-inspection session-id-cache Output Fields

Field Name	Field Description
`Total SSL session identifiers`	Total number of SSL session identifiers stored in the session ID cache.

Sample Output

show security idp ssl-inspection session-id-cache

```
user@host> show security idp ssl-inspection session-id-cache
SSL session identifiers :

c98396c768f983b515d93bb7c421fb6b8ce5c2c5c230b8739b7fcf8ce9c0de4e
a211321a3242233243c3dc0d421fb6b8ce5e4e983b515d932c5c230b87392c

Total SSL session identifiers : 2
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

clear security idp ssl-inspection session-id-cache | 3981

show security idp status

IN THIS SECTION

- Syntax | 24623
- Description | 24623
- Required Privilege Level | 24623
- Output Fields | 24623
- Sample Output | 24624
- Release Information | 24625

Syntax

```
show security idp status
```

Description

Display the status of the current IDP policy.

Required Privilege Level

view

Output Fields

Table 1250 on page 24623 lists the output fields for the `show security idp status` command. Output fields are listed in the approximate order in which they appear. The following information displayed in this command are captured at the point when IDP processes the most recent traffic:

- `Intelligent Inspection State`
- `Packets/second`
- `KBits/second`
- `Latency`

The system may reflect past values of these fields if the system is idle when the command is run.

Table 1250: show security idp status Output Fields

Field Name	Field Description
`Intelligent Inspection State Details`	Status of the IDP intelligent inspection.
`State of IDP`	Status of current IDP policy.
`Packets/second`	The aggregated throughput (packets per second) for the system.

Table 1250: show security idp status Output Fields *(Continued)*

Field Name	Field Description
KBits/second	The aggregated throughput (kilobits per second) for the system.
Latency	• min—Minimum delay for a packet to receive and return by a node in microseconds. • max—Maximum delay for a packet to receive and return by a node in microseconds. • ave—Average delay for a packet to receive and return by a node in microseconds.
Packet Statistics	Statistics for ICMP, TCP, and UDP packets.
Flow Statistics	Flow-related system statistics for ICMP, TCP, and UDP packets.
Session Statistics	Session-related system statistics for ICMP, TCP, and UDP packets.
Number of SSL Sessions	Number of current SSL sessions.
Policy Name	Name of the running policy. If IDP is configured for logical systems, idp-policy-combined is displayed.
Running Detector Version	Current version of the running detector.
Forwarding process mode	IDP dedicated mode: default, equal, idp, or firewall.

Sample Output

show security idp status

```
user@host> show security idp status
Intelligent Inspection State Details:
```

```
 State: Inactive

State of IDP: 2-default,  Up since: 2010-02-04 13:37:16 UTC (17:15:02 ago)

Packets/second: 5               Peak: 11 @ 2010-02-05 06:51:58 UTC
KBits/second  : 2               Peak: 5 @ 2010-02-05 06:52:06 UTC
Latency (microseconds): [min: 0] [max: 0] [avg: 0]

Packet Statistics:
[ICMP: 0] [TCP: 82] [UDP: 0] [Other: 0]

Flow Statistics:
  ICMP: [Current: 0] [Max: 0 @ 2010-02-05 06:49:51 UTC]
  TCP: [Current: 2] [Max: 6 @ 2010-02-05 06:52:08 UTC]
  UDP: [Current: 0] [Max: 0 @ 2010-02-05 06:49:51 UTC]
  Other: [Current: 0] [Max: 0 @ 2010-02-05 06:49:51 UTC]

Session Statistics:
[ICMP: 0] [TCP: 1] [UDP: 0] [Other: 0]

Policy Name : sample
Running Detector Version : 10.4.160091104
```

Release Information

Command introduced in Junos OS Release 9.2. Multiple detector information introduced in Junos OS Release 10.1. Output changed to support IDP dedicated mode in Junos OS Release 11.2. Output changed to support IDP intelligent inspection mode in Junos OS Release 19.2R1.

show security idp status detail

IN THIS SECTION

- Syntax | 24626
- Description | 24626
- Options | 24626
- Required Privilege Level | 24626

- Output Fields | 24626
- Sample Output | 24627
- Release Information | 24630

Syntax

```
show security idp status detail
```

Description

Display statistics for each Services Processing Unit (SPU), including multiple detector information for each SPU.

Options

Required Privilege Level

view

Output Fields

Table 1251 on page 24626 lists the output fields for the `show security idp attack detail` command. Output fields are listed in the approximate order in which they appear.

Table 1251: show security idp status detail Output Fields

Field Name	Field Description
`PIC and FPC`	Indicate the PIC and FPC used.
`State of IDP`	Status of current IDP policy.
`Packets/second`	The aggregated throughput (packets per second) for the system.

Table 1251: show security idp status detail Output Fields *(Continued)*

Field Name	Field Description
KBits/second	The aggregated throughput (kilobits per second) for the system.
Latency	• min—Minimum delay for a packet to receive and return by a node in microseconds. • max—Maximum delay for a packet to receive and return by a node in microseconds. • ave—Average delay for a packet to receive and return by a node in microseconds.
Packet Statistics	Statistics for ICMP, TCP, and UDP packets.
Flow Statistics	Flow-related system statistics for ICMP, TCP, and UDP packets.
Session Statistics	Session-related system statistics for ICMP, TCP, and UDP packets.
Number of SSL Sessions	Number of current SSL sessions.
Policy Name	Name of the running policy. If IDP is configured for logical systems, idp-policy-combined is displayed.
Running Detector Version	Current version of the running detector.
Forwarding process mode	IDP dedicated mode: default, equal, idp, or firewall.

Sample Output

show security idp status detail

```
user@host> show security idp status detail
  PIC : FPC 1 PIC 1:
```

```
State of IDP: Default,  Up since: 2011-03-29 17:25:07 UTC (00:02:48 ago)

Packets/second: 0               Peak: 0 @ 2011-03-29 17:25:07 UTC
KBits/second  : 0               Peak: 0 @ 2011-03-29 17:25:07 UTC
Latency (microseconds): [min: 0] [max: 0] [avg: 0]



Packet Statistics:
[ICMP: 0] [TCP: 0] [UDP: 0] [Other: 0]

Flow Statistics:
  ICMP: [Current: 0] [Max: 0 @ 2011-03-29 17:25:07 UTC]
  TCP: [Current: 0] [Max: 0 @ 2011-03-29 17:25:07 UTC]
  UDP: [Current: 0] [Max: 0 @ 2011-03-29 17:25:07 UTC]
  Other: [Current: 0] [Max: 0 @ 2011-03-29 17:25:07 UTC]

Session Statistics:
 [ICMP: 0] [TCP: 0] [UDP: 0] [Other: 0]

Number of SSL Sessions : 0

 PIC : FPC 1 PIC 0:

State of IDP: Default,  Up since: 2011-03-29 17:25:08 UTC (00:02:47 ago)

Packets/second: 0               Peak: 0 @ 2011-03-29 17:25:08 UTC
KBits/second  : 0               Peak: 0 @ 2011-03-29 17:25:08 UTC
Latency (microseconds): [min: 0] [max: 0] [avg: 0]

Packet Statistics:
 [ICMP: 0] [TCP: 0] [UDP: 0] [Other: 0]

Flow Statistics:
  ICMP: [Current: 0] [Max: 0 @ 2011-03-29 17:25:08 UTC]
  TCP: [Current: 0] [Max: 0 @ 2011-03-29 17:25:08 UTC]
  UDP: [Current: 0] [Max: 0 @ 2011-03-29 17:25:08 UTC]
  Other: [Current: 0] [Max: 0 @ 2011-03-29 17:25:08 UTC]

Session Statistics:
 [ICMP: 0] [TCP: 0] [UDP: 0] [Other: 0]

Number of SSL Sessions : 0
```

```
 PIC : FPC 0 PIC 1:

State of IDP: Default,  Up since: 2011-03-29 17:25:04 UTC (00:02:51 ago)


Packets/second: 0               Peak: 0 @ 2011-03-29 17:25:04 UTC
KBits/second  : 0               Peak: 0 @ 2011-03-29 17:25:04 UTC
Latency (microseconds): [min: 0] [max: 0] [avg: 0]

Packet Statistics:
 [ICMP: 0] [TCP: 0] [UDP: 0] [Other: 0]

Flow Statistics:
  ICMP: [Current: 0] [Max: 0 @ 2011-03-29 17:25:04 UTC]
  TCP: [Current: 0] [Max: 0 @ 2011-03-29 17:25:04 UTC]
  UDP: [Current: 0] [Max: 0 @ 2011-03-29 17:25:04 UTC]
  Other: [Current: 0] [Max: 0 @ 2011-03-29 17:25:04 UTC]

Session Statistics:
 [ICMP: 0] [TCP: 0] [UDP: 0] [Other: 0]


Number of SSL Sessions : 0

  PIC : FPC 1 PIC 1:

  Policy Name : none

  PIC : FPC 1 PIC 0:

  Policy Name : none

  PIC : FPC 0 PIC 1:

  Policy Name : none

Forwarding process mode : maximizing sessions  firewall
```

Release Information

Command introduced in Junos OS Release 10.1. Output changed to support IDP dedicated mode in Junos OS Release 11.2.

show security ike active-peer

IN THIS SECTION

- Syntax | 24630
- Description | 24631
- Options | 24631
- Required Privilege Level | 24632
- Output Fields | 24632
- Sample Output | 24634
- Release Information | 24638

Syntax

```
show security ike active-peer
<peer-address>
<aaa-username username>
<brief | detail>
<debug>
local-address IP address
local-ike-id IKE ID
local-port port number (1..65535)
<fpc slot-number pic slot-number>
<ike-id IKE-ID>
<kmd-instance (all | kmd-instance-name)>
<node-local>
<pic slot-number fpc slot-number>
<port port-number peer-address>
<srg-id id-number>
routing-instance name of the local gateway routing instance
```

```
stats
<ha-link-encryption>
```

Description

Display the list of connected active users with details about the peer addresses and ports they are using.

Options

peer-address	(Optional) Display details about the user with the specified peer address.
aaa-username *username*	(Optional) Display information about the user with the specified authentication, authorization, and accounting (AAA) username.
brief	(Optional) Display standard information about all users. (Default)
detail	(Optional) Display detailed information about all users.
debug	(Optional) Display debug information about all users.
local-address	Display information about the user with the specified local gateway IP address.
local-ike-id	Display information about the user with the specified local gateway IKE ID.
local-port *port-number*	Display information about users on the specified local gateway port number for specified local gateway IP address.
fpc *slot-number* **pic** *slot-number*	(Optional) Display information about users on the specified Flexible PIC Concentrator (FPC) slot and PIC slot.
ike-id *IKE-ID*	(Optional) Display information about the user with the specified IKE ID.
kmd-instance (all | *kmd-instance-name***)**	(Optional) Display information about users in the key management process (KMD) identified by FPC *slot-number* and PIC *slot-number*. • `all`—All KMD instances running on the Services Processing Unit (SPU). • *`kmd-instance-name`*—Name of the KMD instance running on the SPU.
node-local	—(Optional) Display information about users for node-local tunnels in a Multinode High Availability setup.
pic *slot-number* **fpc** *slot-number*	(Optional) Display information about users on the specified PIC slot and FPC slot.

port ***port-number peer-address***	(Optional) Display information about users on the specified port for the specified peer address.
routing-instance	Display information about users on the specified local gateway routing instance.
stats	Display detailed output along with IKE SA stats information accumulated at the peer.
ha-link-encryption	(Optional) Display information related to interchassis link (ICL) tunnel only. See "ipsec (High Availability)" on page 9301 and "show security ike active-peer ha-link-encryption (SRX5400, SRX5600, SRX5800)" on page 24636.
srg-id ***number***	(Optional) Display information related to a specific services redundancy group (SRG) in a Multinode High Availability setup.

Required Privilege Level

view

Output Fields

Table 1252 on page 24632 lists the output fields for the `show security ike active-peer` command. Output fields are listed in the approximate order in which they appear.

Table 1252: show security ike active-peer Output Fields

Field Name	Field Description	Level of Output
`Remote Address`	IP address of the peer.	`brief`
`Port`	Port used by the peer.	All levels
`Peer IKE-ID`	IKE ID used by the peer.	All levels
`AAA username`	Username of the peer.	All levels
`Assigned IP`	IP address assigned to the peer.	`brief`

Table 1252: show security ike active-peer Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Assigned network attributes`	Network attributes assigned to the peer can include the IP address and netmask, and DNS and WINS server addresses.	`detail`
`Previous Peer address`	IP address previously assigned to the peer.	`detail`
`Active IKE SA indexes`	Index number of the SA associated with the peer. This number is an internally generated number.	`detail`
`IKE SA negotiated`	Number of IKE SAs negotiated.	`detail`
`IPSec tunnels active`	Number of IPsec tunnels active.	`detail`
`IPSec Tunnel IDs`	IDs of the active IPsec tunnels.	`detail`
`DPD Config Info`	DPD configuration values.	`detail`
`DPD Statistics`	Information about DPD operations.	`detail`
`Local gateway interface`	Interface name of the local gateway.	`detail`
`Routing instance`	Name of the local gateway routing instance.	`detail`
`Local address`	IP address of the local gateway.	`detail`
`Local IKE-ID`	IKE ID used by local gateway.	`detail`

Sample Output

show security ike active-peer

```
user@host> show security ike active-peer

Remote Address    Port     Peer IKE-ID        AAA username     Assigned IP
192.168.6.136     8034     user1tac@650a      user1
192.168.80.225
```

show security ike active-peer stats

```
user@host> show security ike active-peer stats
Local gateway interface: xe-1/1/2
Routing instance: default
Local address: 192.0.2.1, Port: 500,
Local IKE-ID : device.example.net
Peer address: 198.51.100.2, Port: 500,
Peer IKE-ID : device1.example.net
AAA username: not available
Assigned network attributes:
IP Address      : 192.0.2.10 ,   netmask         : 255.255.255.0
DNS Address     : 19851.100.25 ,   DNS2 Address    : 198.51.100.26
WINS Address    : 203.0.113.25 ,   WINS2 Address   : 203.0.113.26
Assigned network attributes (IPv6):
IP Address      : :: ,   prefix          : 0
DNS Address     : 2001:db8:::ffff ,   DNS2 Address    : 2001:db8::1001
Previous Peer address   : 0.0.0.0, Port               : 0
Active IKE SA indexes   : 1
IKE SA negotiated       : 1
IPSec tunnels active    : 1, IPSec Tunnel IDs   : 500001
IKE_SA_INIT exchange stats:
Initiator stats:                                  Responder stats:
  Request Out             : 0                       Request In               : 1
  Response In             : 0                       Response Out             : 1
  Invalid KE Payload In   : 0                       Invalid KE Payload Out   : 0
  No Proposal Chosen In   : 0                       No Proposal Chosen Out   : 0
  Cookie Request In       : 0                       Cookie Request Out       : 0
  Cookie Response Out     : 0                       Cookie Response In       : 0
  Res Invalid IKE SPI     : 0                       Res DH Gen Key Fail      : 0
```

```
  Res Verify SA Fail      : 0                       Res Invalid DH Group Conf: 0
  Res IKE SA Fill Fail    : 0                       Res Get CAs Fail         : 0
  Res Verify DH Group Fail: 0                       Res Get VID Fail         : 0
  Res DH Compute Key Fail : 0                       Res DH Compute Key Fail  : 0
IKE_AUTH exchange stats:
Initiator stats:                                  Responder stats:
  Request Out             : 0                       Request In               : 1
  Response In             : 0                       Response Out             : 1
  No Proposal Chosen In   : 0                       No Proposal Chosen Out   : 0
  TS Unacceptable In      : 0                       TS Unacceptable Out      : 0
  Authentication Failed In: 0                       Authentication Failed Out: 0
IKE SA Rekey CREATE_CHILD_SA exchange stats:
Initiator stats:                                  Responder stats:
  Request Out             : 0                       Request In             : 0
  Response In             : 0                       Response Out           : 0
  No Proposal Chosen In   : 0                       No Proposal Chosen Out : 0
  Invalid KE In           : 0                       Invalid KE Out         : 0
  Res DH Compute Key Fail : 0                       Res DH Compute Key Fail: 0
  Res Verify SA Fail      : 0
  Res Fill IKE SA Fail    : 0
  Res Verify DH Group Fail: 0
IPsec SA Rekey CREATE_CHILD_SA exchange stats:
Initiator stats:                                  Responder stats:
  Request Out             : 0                       Request In             : 0
  Response In             : 0                       Response Out           : 0
  No Proposal Chosen In   : 0                       No Proposal Chosen Out : 0
  Invalid KE In           : 0                       Invalid KE Out         : 0
  TS Unacceptable In      : 0                       TS Unacceptable Out    : 0
  Res DH Compute Key Fail : 0                       Res DH Compute Key Fail: 0
  Res Verify SA Fail      : 0
  Res Verify DH Group Fail: 0
  Res Verify TS Fail      : 0
```

show security ike active-peer detail

```
user@host> show security ike active-peer detail
Local gateway interface: xe-1/1/2
Routing instance: default
Local address: 192.0.2.1, Port: 500,
Local IKE-ID : device.example.net
Peer address: 198.51.100.2, Port: 500,
```

```
Peer IKE-ID : device1.example.net
AAA username: not available
Assigned network attributes:
IP Address      : 192.0.2.10 ,   netmask         : 255.255.255.0
DNS Address     : 198.51.100.25 ,   DNS2 Address    : 198.51.100.26
WINS Address    : 203.0.113.25 ,   WINS2 Address   : 203.0.113.26
Assigned network attributes (IPv6):
IP Address      : 5000::1 ,   prefix          : 112
DNS Address     : 1000::ffff:ffff ,   DNS2 Address    : 1100::ffff:ffff

Previous Peer address   : 0.0.0.0, Port               : 0
Active IKE SA indexes   : 1
IKE SA negotiated       : 1
IPSec tunnels active    : 1, IPSec Tunnel IDs   : 500001
```

show security ike active-peer ha-link-encryption (SRX5400, SRX5600, SRX5800)

Starting in Junos OS Release 20.4R1, when you configure the high availability (HA) feature, you can use this show command to view only interchassis link tunnel details. The following command displays only interchassis link active peers and not regular active peers.

```
user@host> show security ike active-peer ha-link-encryption

Remote Address  Port     Peer IKE-ID   AAA username   Assigned IP
23.0.0.2        500      23.0.0.2      not available     0.0.0.0
```

show security ike active-peer srg-id 1

```
user@host> show security ike active-peer srg-id 1
Remote Address                      Port     Peer IKE-ID                         AAA
username                        Assigned IP
10.112.0.1                          500      10.112.0.1                                  not
available                0.0.0.0
```

show security ike active-peer node-local

```
user@host> show security ike active-peer node-local
Remote Address    Port     Peer IKE-ID                        AAA
username                        Assigned IP
6.0.0.2           500      DC=juniper, CN=r0, OU=marketing, O=juniper, L=sunnyvale,
ST=california, C=usnot available 0.0.0.0
```

```
user@host> show security ike active-peer node-local detail
Local gateway interface: xe-0/0/2.0
Routing instance: default
Local address: 4.0.0.1, Port: 500,
Local IKE-ID : DC=juniper, CN=r0, OU=marketing, O=juniper, L=sunnyvale, ST=california, C=us
Peer address: 6.0.0.2, Port: 500,
Peer IKE-ID : DC=juniper, CN=r0, OU=marketing, O=juniper, L=sunnyvale, ST=california, C=us
AAA username: not available
Assigned network attributes:
IP Address      : 0.0.0.0 ,   netmask         : 0.0.0.0
DNS Address     : 0.0.0.0 ,   DNS2 Address    : 0.0.0.0
WINS Address    : 0.0.0.0 ,   WINS2 Address   : 0.0.0.0
Assigned network attributes (IPv6):
IP Address      : :: ,   prefix          : 0
DNS Address     : :: ,   DNS2 Address    : ::

Previous Peer address   : 0.0.0.0, Port               : 0
Active IKE SA indexes   : 25
IKE SA negotiated       : 1
IPSec tunnels active    : 1, IPSec Tunnel IDs   : 500003

DPD Config Mode    : always-send
DPD Config Interval: 10
DPD Config Treshold: 3
DPD Config P1SA IDX: 25
DPD Stats Req sent: 5, DPD Stats Resp rcvd: 5
DPD Statistics          : DPD TTL                  :3    DPD seq-no               :0
DPD Statistics          : DPD triggerd p1SA        :0    DPD Reserved             :0
```

Release Information

Command introduced in Junos OS Release 10.4. Support to display dead peer detection (DPD) statistics added in Junos OS Release 12.3X48-D10.

Support for the `ha-link-encryption` option added in Junos OS Release 20.4R1.

Support for the `srg-id` option added in Junos OS Release 22.4R1.

Support for the `node-local` option added in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

show security ike security-associations | 24648

show security ipsec security-associations | 24701

show security ike debug-status

IN THIS SECTION

- Syntax | 24638
- Description | 24638
- Required Privilege Level | 24639
- Output Fields | 24639
- Sample Output | 24639
- Release Information | 24639

Syntax

```
show security ike debug-status
```

Description

Display debug status for currently enabled Internet Key Exchange (IKE) tracing.

Required Privilege Level

view

Output Fields

Table 1253 on page 24639 lists the output fields for the `show security ike debug-status` command. Output fields are listed in the approximate order in which they appear.

Table 1253: show security ike debug-status Output Fields

Field Name	Field Description
Enabled/Disabled	Status of the IKE per-tunnel tracing.
flag	Trace operation; the default is all.
level	Level of logging; the default is 7.
Local IP	Local IP address of the VPN tunnel endpoint.
Remote IP	Remote IP address of the VPN tunnel endpoint.

Sample Output

show security ike debug-status

```
user@host> show security ike debug-status
Enabled
flag: all
level: 7
Local IP: 192.0.2.1, Remote IP: 203.0.113.2
```

Release Information

Command introduced in Junos OS Release 11.4R3.

RELATED DOCUMENTATION

request security ike debug-disable | 14770

request security ike debug-enable | 14771

show security ike peers

IN THIS SECTION

- Syntax | 24640
- Description | 24640
- Options | 24641
- Required Privilege Level | 24641
- Output Fields | 24641
- Sample Output | 24643
- Release Information | 24644

Syntax

```
show security ike peers <backoff | blocked | failed | in-progress>  <brief | detail> role
<initiator | responder> index <value> remote-ip-address <value> local-ip-address <value> gateway-
name <value>
```

```
show security ike peers blocked blocklist-name <blocklist-name> <brief | detail> role <initiator
| responder> index <value> remote-ip-address <value> local-ip-address <value> gateway-name
<value>
```

Description

Display information about IKE peer such as the backoff peers, blocked peers, failed peers, and the in progress peers.

Options

backoff	Display the IKE backoff peers. Display details of the remote peers (initiators) that are temporarily denied permission to start a new IKE negotiation due to the previous SA_INIT or IKE_AUTH phase failures.
blocked	Display the IKE blocked peers. Display the latest 100 IKE blocked negotiations. The blocked negotiations are due to the blocklist configuration at the [`edit security ike peer blocklist`] hierarchy level.
failed	Display the IKE failed peers. Display the latest 200 IKE negotiation failed peers.
in-progress	Display the IKE in progress peers.
blocklist-name *blocklist-name*	(Optional) Provide the IKE blocklist name. Applicable with `blocked` option.
brief	(Optional) Display brief output
detail	(Optional) Display detailed output.
index *value*	(Optional) Enter the index number.
role *value*	(Optional) Enter `initiator` or `responder`.
gateway-name *value*	(Optional) Enter the IKE gateway name.
local-ip-address *value*	(Optional) Enter the local IP address.
remote-ip-address *value*	(Optional) Enter the remote IP address.

Required Privilege Level

View

Output Fields

The Table 1254 on page 24642 lists the output fields of the `show security ike peers in-progress brief` command. You'll see similar output with `backoff`, `blocked`, and `failed` options. The Table 1255 on page

24642 lists the output fields for the `show security ike peers in-progress detail` command. You'll see similar output with `backoff`, `blocked`, and `failed` options.

Table 1254: show security ike peers in-progress brief Output Fields

Field Name	Field Description
Index	Index number
Started at	IKE negotiation process start time.
IKE peer	Remote IKE peer IP address.
Gateway name	IKE gateway name of the remote peer.
Role	The initiator or the responder role.

Table 1255: show security ike peers in-progress detail Output Fields

Field Name	Field Description
Index	Index number
Role	The initiator or the responder role.
Exchange Type	Negotiation mode—either IKEv2 or Main or Aggressive
Authentication Method	Certificates or Pre-shared-keys based authentication
Remote port	Remote port number
Local Address	Local IP address
Local Port	Local port number
Negotiation Phase	INIT
Routing Instance	default
Gateway name	IKE gateway name of the remote peer.

Table 1255: show security ike peers in-progress detail Output Fields *(Continued)*

Field Name	Field Description
NATT Detection	Not Detected or Detected at remote end
Started at	IKE negotiation process start time.

Sample Output

show security ike peers in-progress brief

```
user@host> show security ike peers in-progress brief
```

```
 Index   Started at                       IKE peer       Gateway name    Role

10        2023-01-31 00:45:30     10.0.1.1         IKE_GW1            Initiator
11        2023-01-31 00:45:30     200.1.1.2       IKE_GW1            Responder
```

show security ike peers in-progress detail

```
user@host> show security ike peers in-progress detail
```

```
 IKE peer 10.0.1.1
   Index: 10, Role: Initiator, Exchange Type: IKEv2, Authentication Method:
Certificates
   Remote Port: 500, Local Address: 2.0.0.1, Local Port: 500, Negotiation Phase: INIT
   Routing Instance: default, Gateway Name: IKE_GW1, NATT Detection: Not Detected
   Started At: 2023-01-31 00:45:30


IKE peer 200.1.1.2
   Index: 11, Role: Responder, Exchange Type: Main, Authentication Method: Pre-shared-
keys
   Remote Port: 500, Local Address: 2.0.0.1, Local Port: 500, Negotiation Phase:
INIT
```

```
    Routing Instance: default, Gateway Name: IKE_GW1, NATT Detection: Detected(Remote-end)
    Started At: 2023-01-31 00:45:30
```

Release Information

Command introduced in Junos OS Release 23.4R1

RELATED DOCUMENTATION

IKE Protection from DDoS Attacks

show security ike peers statistics

IN THIS SECTION

- Syntax | 24644
- Description | 24644
- Options | 24645
- Required Privilege Level | 24645
- Output Fields | 24645
- Sample Output | 24646
- Release Information | 24646

Syntax

```
show security ike peers statistics <gateway-name value | summary>
```

Description

Display the global or per IKE gateway level statistics.

Options

statistics Display global IKE gateway statistics for the remote peers to show details for all the gateways.

gateway-name *value* (Optional) Enter the IKE gateway name to display the statistics for a specific gateway.

summary (Optional) Display the IKE statistics summary with active, in progress, failed, blocked, and backoff peers.

Required Privilege Level

View

Output Fields

The Table 1256 on page 24645 lists the output fields for the `show security ike peers statistics` command.

Table 1256: show security ike peers statistics Output Fields

Field Name	Field Description
Gateway name	IKE gateway name of the remote peer. This field is not seen with the CLI option `summary`.
Active	Number of successful IKE negotiations with the remote peers.
In-progress	Number of in progress IKE negotiations with the remote peers.
Failed	Number of failed IKE negotiations with the remote peers.
Blocked	Number of blocked IKE negotiations with the remote peers.
Backoff	Number of backed off IKE negotiations with the remote peers.

Sample Output

show security ike peers statistics

```
user@host> show security ike peers statistics
```

```
 Gateway name     Active   In-progress   Failed    Blocked
Backoff

 IKE_GW1             2          10                 2           3
4
 IKE_GW2             3          11                 1           3              2
```

show security ike peers statistics summary

```
user@host> show security ike peers statistics summary
```

```
 Active   In-progress   Failed    Blocked   Backoff

 2          10                 2           3             4
```

Release Information

Command introduced in Junos OS Release 23.4R1

RELATED DOCUMENTATION

- IKE Protection from DDoS Attacks

show security ike pre-shared-key

IN THIS SECTION

- Syntax | 24647
- Description | 24647
- Options | 24647
- Required Privilege Level | 24647
- Sample Output | 24648
- Release Information | 24648

Syntax

```
show security ike pre-shared key
master-key      <master-key>
user-id           <user-id>
gateway         <gateway_name>
```

Description

Display the Internet Key Exchange (IKE) preshared key used by the Virtual Private network (VPN) gateway to authenticate the remote access user. Use either master-key or gateway option to get the master presharedkey.

Options

- `master-key` *`master-key`* —(Optional) Primary preshared key.
- `user-id` *`user-id`* —(Optional) IKE user ID value.
- `gateway` *`gateway_name`*—(Optional) Label of the VPN gateway set with master preshared key.

Required Privilege Level

view

Sample Output

show security ike pre-shared-key user-id

```
user@host> show security ike pre-shared-key user-id a@example.net master-key example
Preshared Key:3b33ec3631a561ec5a710f5d02f208033b108bb4
```

show security ike pre-shared-key gateway gateway_name user-id user-id

```
user@host> show security ike pre-shared-key gateway HUB_GW user-id user1@juniper.net
Pre-shared key: 79e4ea39f5c06834a3c4c031e37c6de24d46798a
```

Release Information

Command introduced in Junos OS Release 8.5.

`gateway` option is introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

policy (Security IKE) | 12713

show security ike security-associations

IN THIS SECTION

- Syntax | 24649
- Description | 24649
- Options | 24649
- Required Privilege Level | 24650
- Output Fields | 24651
- Sample Output | 24658

Release Information | 24670

Syntax

```
show security ike security-associations
<peer-address>
<brief | detail>
<family (inet  | inet6)>
<fpc slot-number>
<index SA-index-number>
<kmd-instance (all | kmd-instance-name)>
<node-local>
<pic slot-number>
<sa-type shortcut >
<srg-id id-number>
<ha-link-encryption>
```

Description

Display information about Internet Key Exchange security associations (IKE SAs).

Options

- none—Display standard information about existing IKE SAs, including index numbers.
- *peer-address*—(Optional) Display details about a particular SA based on the IPv4 or IPv6 address of the destination peer. This option and `index` provide the same level of output.
- `brief`—(Optional) Display standard information about all existing IKE SAs. (Default)
- `detail`—(Optional) Display detailed information about all existing IKE SAs.
- `family`—(Optional) Display IKE SAs by family. This option is used to filter the output.
  - `inet`—IPv4 address family.
  - `inet6`—IPv6 address family.
- `fpc` *slot-number*—(Optional) Display information about existing IKE SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.

In a chassis cluster, when you execute the CLI command `show security ike security-associations pic <slot-number> fpc <slot-number>` in operational mode, only the primary node information about the existing IPsec SAs in the specified Flexible PIC Concentrator (FPC) slot and PIC slot is displayed.

Starting Junos OS Release 23.4R1, this option is not available when `junos-ike` package is installed for running IPsec VPN using `IKED` process.

- `index` *`SA-index-number`*—(Optional) Display information for a particular SA based on the index number of the SA. For a particular SA, display the list of existing SAs by using the command with no options. This option and *`peer-address`* provide the same level of output.

- `kmd-instance` —(Optional) Display information about existing IKE SAs in the key management process (in this case, it is KMD) identified by FPC *slot-number* and PIC *slot-number*. This option is used to filter the output.

  - `all`—All KMD instances running on the Services Processing Unit (SPU).

  - *`kmd-instance-name`*—Name of the KMD instance running on the SPU.

  Starting Junos OS Release 23.4R1, this option is not available when `junos-ike` package is installed for running IPsec VPN using `IKED` process.

node-local —(Optional) Display information about IKE SAs for node-local tunnels in a Multinode High Availability setup.

- `pic` *`slot-number`* —(Optional) Display information about existing IKE SAs in this PIC slot. This option is used to filter the output.

  Starting Junos OS Release 23.4R1, this option is not available when `junos-ike` package is installed for running IPsec VPN using `IKED` process.

- `sa-type shortcut`—(Optional) It's applicable for ADVPN. Display information about IKE SAs by type `shortcut`.

- `ha-link-encryption`—(Optional) Display information related to interchassis link tunnel only. See "ipsec (High Availability)" on page 9301 and "show security ike security-associations ha-link-encryption (SRX5400, SRX5600, SRX5800)" on page 24667.

- `srg-id`—(Optional) Display information related to a specific services redundancy group (SRG).

Required Privilege Level

view

Output Fields

Table 1257 on page 24651 lists the output fields for the `show security ike security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 1257: show security ike security-associations Output Fields

Field Name	Field Description
`IKE Peer or Remote Address`	IP address of the destination peer with which the local peer communicates.
`Index`	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
`Gateway Name`	Name of the IKE gateway.
`Location`	• `FPC`—Flexible PIC Concentrator (FPC) slot number. • `PIC`—PIC slot number. • `KMD-Instance`—The name of the KMD instance running on the SPU, identified by FPC *slot-number* and PIC *slot-number*. Currently, 4 KMD instances are running on each SPU, and any particular IKE negotiation is carried out by a single KMD instance.
`Role`	Part played in the IKE session. The device triggering the IKE negotiation is the initiator, and the device accepting the first IKE exchange packets is the responder.
`State`	State of the IKE SAs: • `DOWN`—SA has not been negotiated with the peer. • `UP`—SA has been negotiated with the peer.
`Initiator cookie`	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.

Table 1257: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description
`Responder cookie`	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received. A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.
`Exchange type`	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between one another. Each exchange type or mode determines the number of messages and the payload types that are contained in each message. The modes are: • `main`—The exchange is done with six messages. This mode encrypts the payload, protecting the identity of the neighbor. • `aggressive`—The exchange is done with three messages. This mode does not encrypt the payload, leaving the identity of the neighbor unprotected. IKEv2 protocol does not use the mode configuration for negotiation. Therefore, the mode displays the version number of the security association.
`Authentication method`	Method used to authenticate the source of IKE messages, which can be either `Pre-shared-keys` or digital certificates, such as `DSA-signatures`, `ECDSA-signatures-256`, `ECDSA-signatures-384`, `RSA-signatures`, or `digital-signature`.
`Local`	Address of the local peer.
`Remote`	Address of the remote peer.
`Lifetime`	Number of seconds remaining until the IKE SA expires.
`Reauth Lifetime`	When enabled, number of seconds remaining until reauthentication triggers a new IKEv2 SA negotiation.

Table 1257: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description
`IKE Fragmentation`	`Enabled` means that both the IKEv2 initiator and responder support message fragmentation and have negotiated the support during the IKE_SA_INIT message exchange. `Size` shows the maximum size of an IKEv2 message before it is fragmented.

Table 1257: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description

`Algorithms`	IKE algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:
	• `Authentication`—Type of authentication algorithm used:
	  • `sha1`—Secure Hash Algorithm 1 authentication.
	  • `md5`—MD5 authentication.
	• `Encryption`—Type of encryption algorithm used:
	  • `aes-256-cbc`—Advanced Encryption Standard (AES) 256-bit encryption.
	  • `aes-192-cbc`— AES192-bit encryption.
	  • `aes-128-cbc`—AES 128-bit encryption.
	  • `3des-cbc`—3 Data Encryption Standard (DES) encryption.
	  • `aes-128-gcm`—Advanced Encryption Standard (AES) 256-bit encryption.
	  • `chacha20-poly1305`—ChaCha20-Poly1305 authenticated encryption.
	  • `des-cbc`—DES encryption.
	  Starting in Junos OS Release 19.4R2, when you configure `aes-128-gcm` or `aes-256-gcm` as an encryption algorithm at the `[edit security ipsec proposalproposal-name]` hierarchy level, the authentication algorithm field of the `show security ikesecurity-associations detail` command displays the same configured encryption algorithm.
	• `Pseudo random function`—Function that generates highly unpredictable random numbers: `hmac-md5` or `hmac-sha1`.
	• `Diffie-Hellman group`—Specifies the type of Diffie-Hellman group when performing the new Diffie-Hellman exchange. It can be one of the following:
	  • `group1`—768-bit Modular Exponential (MODP) algorithm.
	  • `group2`—1024-bit MODP algorithm.

Table 1257: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description
	• `group14`—2048-bit MODP group. • `group15`—3072-bit MODP algorithm. • `group16`—4096-bit MODP algorithm. • `group19`—256-bit random Elliptic Curve Groups modulo a prime (ECP group) algorithm. • `group20`—384-bit random ECP group algorithm. • `group21`—521-bit random ECP group algorithm. • `group24`—2048-bit MODP group with 256-bit prime order subgroup. • `Signature hash algo`—Specifies the signature hash algorithm used by the local and remote endpoints.
`Traffic statistics`	• `Input bytes`–Number of bytes received. • `Output bytes`–Number of bytes transmitted. • `Input packets`–Number of packets received. • `Output packets`–Number of packets transmitted. • `Input fragmented packets`—Number of IKEv2 fragmented packets received. • `Output fragmented packets`—Number of IKEv2 fragmented packets transmitted.

Table 1257: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description
`Flags`	Notification to the key management process of the status of the IKE negotiation: • `caller notification sent`—Caller program notified about the completion of the IKE negotiation. • `waiting for done`—Negotiation is done. The library is waiting for the remote end retransmission timers to expire. • `waiting for remove`—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation. • `waiting for policy manager`—Negotiation is waiting for a response from the policy manager.
`IPSec security associations`	• *`number`* `created`: The number of SAs created. • *`number`* `deleted`: The number of SAs deleted.

Table 1257: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description

`Phase 2 negotiations in progress`	Number of Phase 2 IKE negotiations in progress and status information:
	• `Negotiation type`—Type of Phase 2 negotiation. Junos OS currently supports quick mode.
	• `Message ID`—Unique identifier for a Phase 2 negotiation.
	• `Local identity`—Identity of the local Phase 2 negotiation. The format is *id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)*.
	• `Remote identity`—Identity of the remote Phase 2 negotiation. The format is *id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)*.
	• `Flags`—Notification to the key management process of the status of the IKE negotiation:
	  • `caller notification sent`—Caller program notified about the completion of the IKE negotiation.
	  • `waiting for done`—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.
	  • `waiting for remove`—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.
	  • `waiting for policy manager`—Negotiation is waiting for a response from the policy manager.

`Local gateway interface`	Interface name of the local gateway.
`Routing instance`	Name of the local gateway routing instance.
`IPsec Tunnel IDs`	Indicates the list of child IPsec tunnel IDs

Sample Output

show security ike security-associations (IPv4)

```
user@host> show security ike security-associations
Index        Remote Address        State        Initiator cookie        Responder cookie    Mode
8        192.168.1.2        UP        3a895f8a9f620198    9040753e66d700bb    Main
Index        Remote Address        State        fInitiator cookie    Responder cookie    Mode
9        192.168.1.3         UP        5ba96hfa9f65067        70890755b65b80b        Main
```

show security ike security-associations (IPv6)

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address
5       UP     e48efd6a444853cf  0d09c59aafb720be  Aggressive     2001:db8::1112
```

show security ike security-associations detail (SRX300, SRX320, SRX340, SRX345, and SRX550HM Devices)

```
user@host> show security ike security-associations detail
IKE peer 192.168.134.245, Index 2577565, Gateway Name: tropic
  Role: Initiator, State: UP
  Initiator cookie: b869b3424513340a, Responder cookie: 4cb3488cb19397c3
  Exchange type: Main, Authentication method: Pre-shared-keys Trusted CA group: xyz_ca_grp
  Local: 192.168.134.241:500, Remote: 192.168.134.245:500
  Local gateway interface: ge-0/0/0
  Routing instance: default
  Lifetime: Expires in 169 seconds
  Peer ike-id: 192.168.134.245
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes-128-gcm
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 1012
   Output bytes  :                 1196
```

```
   Input  packets:                    4
   Output packets:                    5
  Flags: IKE SA is created
  IPSec security associations: 1 created, 0 deleted
  Phase 2 negotiations in progress: 0

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 192.168.134.241:500, Remote: 192.168.134.245:500
    Local identity: 192.168.134.241
    Remote identity: 192.168.134.245
    Flags: IKE SA is created
IPsec SA Rekey CREATE_CHILD_SA exchange stats:
   Initiator stats:                                  Responder stats:
    Request Out             : 1                       Request In             : 0
    Response In             : 1                       Response Out           : 0
    No Proposal Chosen In   : 0                       No Proposal Chosen Out : 0
    Invalid KE In           : 0                       Invalid KE Out         : 0
    TS Unacceptable In      : 0                       TS Unacceptable Out    : 0
    Res DH Compute Key Fail : 0                       Res DH Compute Key Fail: 0
    Res Verify SA Fail      : 0
    Res Verify DH Group Fail: 0
    Res Verify TS Fail      : 0
```

show security ike security-associations detail (SRX5400, SRX5600, and SRX5800 Devices)

```
user@host> show security ike security-associations detail
IKE peer 2.0.0.2, Index 2068, Gateway Name: IKE_GW
  Role: Responder, State: DOWN
  Initiator cookie: aa08091f3d4f1fb6, Responder cookie: 08c89a7add5f9332
  Exchange type: IKEv2, Authentication method: Pre-shared-keys
  Local gateway interface: ge-0/0/3
  Routing instance: default
  Local: 2.0.0.1:500, Remote: 2.0.0.2:500
  Lifetime: Expires in 186 seconds
  Reauth Lifetime: Disabled
  IKE Fragmentation: Enabled, Size: 576
  Remote Access Client Info: Unknown Client
  Peer ike-id: 2.0.0.2
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha256-128
```

```
    Encryption            : aes128-cbc
    Pseudo random function: hmac-sha256
    Diffie-Hellman group  : DH-group-5
   Traffic statistics:
    Input  bytes  :                  704
    Output bytes  :                 1408
    Input  packets:                    4
    Output packets:                    4
    Input  fragmented packets:       0
    Output fragmented packets:       0
   IPSec security associations: 4 created, 2 deleted
   Phase 2 negotiations in progress: 1
   IPSec Tunnel IDs:  500766, 500767


 Negotiation type: Quick mode, Role: Responder, Message ID: 0
 Local: 2.0.0.1:500, Remote: 2.0.0.2:500
 Local identity: 2.0.0.1
 Remote identity: 2.0.0.2
 Flags: IKE SA is created

 IPsec SA Rekey CREATE_CHILD_SA exchange stats:
    Initiator stats:                                  Responder stats:
     Request Out             : 0                       Request In             : 0
     Response In             : 0                       Response Out           : 0
     No Proposal Chosen In   : 0                       No Proposal Chosen Out : 0
     Invalid KE In           : 0                       Invalid KE Out         : 0
     TS Unacceptable In      : 0                       TS Unacceptable Out    : 0
     Res DH Compute Key Fail : 0                       Res DH Compute Key Fail: 0
     Res Verify SA Fail      : 0
     Res Verify DH Group Fail: 0
     Res Verify TS Fail      : 0
```

command-name

The "show security ike stats" on page 24671 topic lists the output fields for the `show security ike security-associations detail` command.

show security ike security-associations family inet6

```
user@host> show security ike security-associations family inet6
  IKE peer 2001:db8:1212::1112, Index 5, Gateway Name: tropic
  Role: Initiator, State: UP
  Initiator cookie: e48efd6a444853cf, Responder cookie: 0d09c59aafb720be
  Exchange type: Aggressive, Authentication method: Pre-shared-keys
  Local: 2001:db8:1212::1111:500, Remote: 2001:db8:1212::1112:500
  Lifetime: Expires in 19518 seconds
  Peer ike-id: not valid
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : sha1
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 1568
   Output bytes  :                 2748
   Input  packets:                    6
   Output packets:                   23
  Flags: Caller notification sent
  IPSec security associations: 5 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 2900338624
    Local: 2001:db8:1212::1111:500, Remote: 2001:db8:1212::1112:500
    Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    Flags: Caller notification sent, Waiting for done
```

show security ike security-associations index 222075191 detail

```
user@host> show security ike security-associations index 222075191 detail
node0:
-
IKE peer 192.168.1.2, Index 222075191, Gateway Name: ZTH_HUB_GW
  Location: FPC 0, PIC 3, KMD-Instance 2
  Auto Discovery VPN:
   Type: Static, Local Capability: Suggester, Peer Capability: Partner
```

```
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent    :    2
     Suggestions accepted:    4
     Suggestions declined:    1
  Role: Responder, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.1.1:500, Remote: 192.168.1.2:500
  Lifetime: Expires in 828 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=cssvk36-d
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                20474
   Output bytes  :                21091
   Input  packets:                  237
   Output packets:                  237
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 192.168.1.1:500, Remote: 192.168.1.2:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=host1
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=host2
    Flags: IKE SA is created
```

show security ike security-associations index 788674 detail

```
user@host> show security ike security-associations index 788674 detail
IKE peer 192.168.1.1, Index 788674, Gateway Name: ZTH_SPOKE_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:    2
     Suggestions accepted:    2
     Suggestions declined:    0
```

```
  Role: Initiator, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.1.2:500, Remote: 192.168.1.1:500
  Lifetime: Expires in 734 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=test
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                22535
   Output bytes  :                21918
   Input  packets:                  256
   Output packets:                  256
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 192.168.1.2:500, Remote: 192.168.1.1:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=host1
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=host2
    Flags: IKE SA is created
```

show security ike security-associations 192.168.1.2

```
user@host> show security ike security-associations 192.168.1.2
Index     State  Initiator cookie  Responder cookie  Mode Remote Address
   8        UP     3a895f8a9f620198  9040753e66d700bb  Main 192.168.1.2
```

show security ike security-associations fpc 6 pic 1 kmd-instance all (SRX Series Firewalls)

```
user@host> show security ike security-associations fpc 6 pic 1 kmd-instance all
Index      Remote Address  State  Initiator cookie  Responder cookie  Mode

1728053250 192.168.1.2     UP     fc959afd1070d10b  bdeb7e8c1ea99483  Main
```

show security ike security-associations detail (ADVPN Suggester, Static Tunnel)

```
user@host> show security ike security-associations detail
IKE peer 192.168.0.105, Index 13563297, Gateway Name: zth_hub_gw
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
 Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent            :  12
     Suggestion response accepted:  12
     Suggestion response declined:   0
  Role: Responder, State: UP
  Initiator cookie: 4d3f4e4b2e75d727, Responder cookie: 81ab914e13cecd21
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.0.154:500, Remote: 192.168.0.105:500
  Lifetime: Expires in 26429 seconds
  Peer ike-id: DC=example, CN=host02, L=Sunnyvale, ST=CA, C=US
```

show security ike security-associations detail (ADVPN Partner, Static Tunnel)

```
user@host> show security ike security-associations detail
IKE peer 192.168.0.154, Index 4980720, Gateway Name: zth_spoke_gw
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
 Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:  12
     Suggestions accepted:  12
     Suggestions declined:   0
  Role: Initiator, State: UP
  Initiator cookie: 4d3f4e4b2e75d727, Responder cookie: 81ab914e13cecd21
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.0.105:500, Remote: 192.168.0.154:500
  Lifetime: Expires in 26252 seconds
  Peer ike-id: DC=example, CN=host01, OU=SBU, O=example, L=Sunnyvale, ST=CA, C=US
```

show security ike security-associations detail (ADVPN Partner, Shortcut)

```
user@host> show security ike security-associations detail
IKE peer 192.168.0.106, Index 4980737, Gateway Name: GW-ADVPN-GT-ADVPN-zth_spoke_vpn-268173323
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
   Type: Shortcut, Local Capability: Partner, Peer Capability: Partner
  Role: Responder, State: UP
  Initiator cookie: e1ed0c655929debc, Responder cookie: 437de6ed784ba63e
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.0.105:500, Remote: 192.168.0.106:500
  Lifetime: Expires in 28796 seconds
  Peer ike-id: DC=example, CN=paulyd, L=Sunnyvale, ST=CA, C=US
```

show security ike security-associations sa-type shortcut (ADVPN)

```
user@host> show security ike security-associations sa-type shortcut
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address
4980742 UP     vb56fbe694eaee5b6 064dbccbfa3b2aab  IKEv2          192.168.0.106
```

show security ike security-associations sa-type shortcut detail (ADVPN)

```
user@host> show security ike security-associations sa-type shortcut detail
IKE peer 192.168.0.106, Index 4980742, Gateway Name: GW-ADVPN-GT-ADVPN-zth_spoke_vpn-268173327
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
   Type: Shortcut, Local Role: Partner, Peer Role: Partner
  Role: Responder, State: UP
```

show security ike security-associations detail (IKEv2 Reauthentication)

```
user@host> show security ike security-associations detail
IKE peer 10.1.2.11, Index 6009224, Gateway Name: GW
  Role: Responder, State: UP
  Initiator cookie: 2c74d14c798a9d70, Responder cookie: 83cbb49bfbcb80cb
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 10.1.1.11:500, Remote: 10.1.2.11:500
```

```
  Lifetime: Expires in 173 seconds
  Reauth Lifetime: Expires in 600 seconds
  Peer ike-id: vsrx@example.net
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes128-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-2
  Traffic statistics:
   Input  bytes  :                 1782
   Output bytes  :                 1743
   Input  packets:                    2
```

show security ike security-associations detail (IKEv2 Fragmentation)

```
user@host> show security ike security-associations detail
IKE peer 172.24.23.157, Index 11883008, Gateway Name: routebased_s2s_gw-552_1
  Role: Responder, State: UP
  Initiator cookie: f3255e720f162e3a, Responder cookie: 17555e3ff7451841
  Exchange type: Main, Authentication method: Pre-shared-keys Trusted CA group: xyz_ca_grp
  Local: 192.168.254.1:500, Remote: 172.24.23.157:500
  Lifetime: Expires in 530 seconds
  Reauth Lifetime: Disabled
  IKE Fragmentation: Enabled, Size: 576
  Peer ike-id: 172.24.23.157
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 1004
   Output bytes  :                  756
   Input  packets:                    6
   Output packets:                    4
   Input  fragmented packets:  3
   Output fragmented packets: 3
  IPSec security associations: 1 created, 1 deleted
  Phase 2 negotiations in progress: 1
```

```
    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 192.168.254.1:500, Remote: 172.24.23.157:500
    Local identity: 192.168.254.1
    Remote identity: 172.24.23.157
    Flags: IKE SA is created
```

show security ike security-associations ha-link-encryption (SRX5400, SRX5600, SRX5800)

Starting in Junos OS Release 20.4R1, when you configure the high availability (HA) feature, you can use this show command to view only interchassis link tunnel details. The following command displays only the link encryption SAs on both nodes.

```
user@host> show security ike security-associations ha-link-encryption

Index      State  Initiator cookie    Responder cookie  Mode   Remote Address
4294966287  UP    7b77b4e2fd5a87e5    ab4a398e6a28687a  IKEv2   23.0.0.2
```

show security ike security-associations srg-id

```
user@host> show security ike security-associations srg-id 1
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address
16778113 UP    16d1f4efae91608c  53f234767bdd0b9b  IKEv2          10.112.0.1
```

show security ike security-associations node-local

```
user@host> show security ike security-associations node-local
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address
24      UP     c982a43f5dd03bf0  c37ae96722a0e1bc  IKEv2          6.0.0.2
```

show security ike security-associations node-local detail

```
user@host> show security ike security-associations node-local
IKE peer 6.0.0.2, Index 25, Gateway Name: IKEv1_GW
  Role: Responder, State: UP
  Initiator cookie: 34b2b16c3dd35442, Responder cookie: 91fc9975f83e932d
  Exchange type: IKEv2, Authentication method: RSA-signatures
```

```
  Local gateway interface: xe-0/0/2.0
  Routing instance: default
  Local: 4.0.0.1:500, Remote: 6.0.0.2:500
  Lifetime: Expires in 1159 seconds
  Reauth Lifetime: Disabled
  IKE Fragmentation: Enabled, Size: 576
  Remote Access Client Info: Unknown Client
  Peer ike-id: DC=juniper, CN=r0, OU=marketing, O=juniper, L=sunnyvale, ST=california, C=us
  AAA assigned IP: 0.0.0.0
  PPK-profile: None
  Algorithms:
   Authentication        : hmac-sha384-192
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha384
   Diffie-Hellman group  : DH-group-19
  Traffic statistics:
   Input  bytes  :                 3434
   Output bytes  :                 3427
   Input  packets:                   15
   Output packets:                   15
   Input  fragmented packets:       4
   Output fragmented packets:       4
  IPSec security associations: 4 created, 1 deleted
  Phase 2 negotiations in progress: 1
  IPSec Tunnel IDs: 500003

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 4.0.0.1:500, Remote: 6.0.0.2:500
    Local identity: DC=juniper, CN=r0, OU=marketing, O=juniper, L=sunnyvale, ST=california, C=us
    Remote identity: DC=juniper, CN=r0, OU=marketing, O=juniper, L=sunnyvale, ST=california, C=us
    Flags: IKE SA is created

  IPsec SA Rekey CREATE_CHILD_SA exchange stats:
   Initiator stats:                                  Responder stats:
    Request Out             : 0                       Request In             :
0
    Response In             : 0                       Response Out           :
0
    No Proposal Chosen In   : 0                       No Proposal Chosen Out :
0
    Invalid KE In           : 0                       Invalid KE Out         :
0
    TS Unacceptable In      : 0                       TS Unacceptable Out    :
```

```
0
    Res DH Compute Key Fail : 0                       Res DH Compute Key Fail:
0
    Res Verify SA Fail      : 0
    Res Verify DH Group Fail: 0
    Res Verify TS Fail      : 0
```

show security ike security-associations detail (ChaCha20-Poly1305)

```
user@host> show security ike security-associations detail
IKE peer 10.1.1.10, Index 1, Gateway Name: ike_gw
  Role: Responder, State: UP
  Initiator cookie: 26b772950a8bea9c, Responder cookie: be644ff0068a0e05
  Exchange type: IKEv2, Authentication method: Pre-shared-keys
  Local gateway interface: ge-0/0/1.0
  Routing instance: default
  Local: 192.168.1.20:500, Remote: 10.1.1.10:500
  Lifetime: Expires in 28790 seconds
  Reauth Lifetime: Disabled
  IKE Fragmentation: Enabled, Size: 576
  Remote Access Client Info: Unknown Client
  Peer ike-id: 10.1.1.10
  AAA assigned IP: 0.0.0.0
  PPK-profile: None
  Algorithms:
   Authentication        : chacha20-poly1305
   Encryption            : chacha20-poly1305
   Pseudo random function: hmac-sha384
   Diffie-Hellman group  : DH-group-2
```

show security ike security-associations detail (IKEv2 digital signature authentication method and signature hash algorithm)

```
user@host> show security ike security-associations detail
  IKE peer 10.0.0.1, Index 83, Gateway Name: SPK_GW1
  Role: Initiator, State: UP
  Initiator cookie: 4c5c36dc34e24093, Responder cookie: 69abe1c4b45a2a9f
  Exchange type: Main, Authentication method: digital-signature(ECDSA)
  Local gateway interface: ge-0/0/2.0
```

```
Routing instance: default
Local: 192.168.1.2:500, Remote: 10.0.0.1:500
Lifetime: Expires in 28515 seconds
Reauth Lifetime: Disabled
IKE Fragmentation: Disabled, Size: 0
Remote Access Client Info: Unknown Client
Peer ike-id: 10.0.0.1
AAA assigned IP: 0.0.0.0
PPK-profile: None
Algorithms:
 Authentication        : hmac-sha256-128
 Encryption            : aes256-cbc
 Pseudo random function: hmac-sha256
 Diffie-Hellman group  : DH-group-14
 Signature hash algo   : sha256(local), sha1(remote)
```

Release Information

Command introduced in Junos OS Release 8.5. Support for the `fpc`, `pic`, and `kmd-instance` options added in Junos OS Release 9.3. Support for the `family` option added in Junos OS Release 11.1. Support for Auto Discovery VPN added in Junos OS Release 12.3X48-D10. Support for IKEv2 reauthentication added in Junos OS Release 15.1X49-D60. Support for IKEv2 fragmentation added in Junos OS Release 15.1X49-D80.

Support for the `ha-link-encryption` option added in Junos OS Release 20.4R1.

Support for the `srg-id` option added in Junos OS Release 22.4R1.

Support for the `node-local` option added in Junos OS Release 23.2R1.

Support for the `chacha20-poly1305` option added in Junos OS Release 24.2R1.

Support for the `digital-signature` and `Signature hash algo` options added in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

Example: Configuring a Route-Based VPN Tunnel in a User Logical Systems

show security ike stats

IN THIS SECTION

- Syntax | 24671
- Description | 24671
- Options | 24671
- Required Privilege Level | 24672
- Output Fields | 24672
- Sample Output | 24677
- Sample Output | 24678
- Sample Output | 24680
- Release Information | 24681

Syntax

```
show security ike stats <brief | detail>
```

Description

Display information about global IKE (Internet Key Exchange) statistics for the tunnels such as in-progress, established, and expired negotiations using IKEv2 on your SRX5000 line with SPC3 card.

Options

- **Default:** `brief`

  Displays tunnel count statistics and non-zero counters of the global IKE statistics.

`detail`

Displays all the global IKE and tunnel count statistics.

Required Privilege Level

view

Output Fields

Table 1258 on page 24672 lists the output fields of total IKE SA and tunnel count statistics. Table 1259 on page 24673 lists the output fields of `IKE_SA_INIT`, `IKE_AUTH`, `IKE SA Rekey CREATE_CHILD_SA`, `IPsec SA Rekey CREATE_CHILD_SA` exchanges statistics. Table 1260 on page 24677 lists total IKE message failure statistics for the `show security ike stats` command. Output fields are listed in the approximate order in which they appear.

Table 1258: total-IKE-SA-and-tunnel-count-statistics Output Fields

Field Name	Field Description
`Number of IKE SAs`	Number of IKE SAs currently active.
`Number of IPsec Tunnels`	Number of IPsec tunnels currently active.
`Number of DPD failovers`	Shows number of times there is a DPD failover.

Table 1259: IKEV2_negotiaton_exchange_statistics

Field Name	Field Description for Output Fields of Initiator Statistics	Field Description for Output Fields of Responder Statistics
IKE_SA_INIT exchange stats	• Request Out —Number of IKE_SA_INIT request message sent by initiator. • Response In—Number of IKE_SA_INIT response message received by initiator. • Invalid KE Payload In—Number of IKE_SA_INIT INVALID_KE_PAYLOAD notification message received by initiator. • No Proposal Chosen In—Number of IKE_SA_INIT NO_PROPSAL_CHOSEN notification message received by initiator. • Cookie Request In—Number of IKE_SA_INIT cookie request notification message received by initiator. • Cookie Response Out—Number of IKE_SA_INIT cookie response notification message sent by responder. • Res Invalid IKE SPI—Number of IKE_SA_INIT response message containing invalid SPI received by initiator. • Res Verify SA Fail—Number of IKE_SA_INIT response message processing failed during verification of peer SA at initiator. • Res IKE SA Fill Fail—Number of IKE_SA_INIT response message processing failed during verification of IKE SA fill operation at initiator.	• Request In—Number of IKE_SA_INIT request message received by responder. • Response Out—Number of IKE_SA_INIT response message sent by responder. • Invalid KE Payload Out—Number of IKE_SA_INIT INVALID_KE_PAYLOAD notification message sent by responder. • No Proposal Chosen Out—Number of IKE_SA_INIT NO_PROPSAL_CHOSEN notification message sent by responder. • Cookie Request Out—Number of IKE_SA_INIT cookie request notification message sent by responder. • Cookie Response In—Number of IKE_SA_INIT cookie response notification message received by responder. • Res DH Gen Key Fail—Number of IKE_SA_INIT response message processing failed during Diffie-Hellman generate key at responder. • Res Invalid DH Group Conf—Number of IKE_SA_INIT response message processing failed due to invalid Diffie-Hellman group configured at responder. • Res Get CAs Fail—Number of IKE_SA_INIT response message processing failed during get CAs operation at responder. • Res Get VID Fail—Number of IKE_SA_INIT response message processing failed during get vendor ID request operation at responder.

Table 1259: IKEV2_negotiaton_exchange_statistics *(Continued)*

Field Name	Field Description for Output Fields of Initiator Statistics	Field Description for Output Fields of Responder Statistics
	• `Res Verify DH Group Fail`—Number of `IKE_SA_INIT` response message processing failed during verification of Diffie-Hellman group at initiator. • `Res DH Compute Key Fail`—Number of `IKE_SA_INIT` response message processing failed during verification of Diffie-Hellman compute key at initiator.	• `Res DH Compute Key Fail`—Number of `IKE_SA_INIT` response message processing failed during Diffie-Hellman compute key at responder.
`IKE_AUTH exchange stats`	• `Request Out`—Number of `IKE_AUTH` request message sent by initiator. • `Response In`—Number of `IKE_AUTH` response message received by initiator. • `No Proposal Chosen In`—Number of `IKE_AUTH NO_PROPSAL_CHOSEN` notification message received by initiator. • `TS Unacceptable In`—Number of `IKE_AUTH TS_UNACCEPTABLE` notification message received by initiator. • `Authentication Failed In`—Number of `IKE_AUTH AUTHENTICATION_FAILED` notification message received by initiator.	• `Request In`—Number of `IKE_AUTH` request message received by responder. • `Response Out`—Number of `IKE_AUTH` response message sent by responder. • `No Proposal Chosen Out`—Number of `IKE_AUTH NO_PROPSAL_CHOSEN` notification message sent by responder. • `TS Unacceptable out`—Number of `IKE_AUTH TS_UNACCEPTABLE` notification message sent by responder. • `Authentication Failed Out`—Number of `IKE_AUTH AUTHENTICATION_FAILED` notification message sent by responder.

Table 1259: IKEV2_negotiaton_exchange_statistics *(Continued)*

Field Name	Field Description for Output Fields of Initiator Statistics	Field Description for Output Fields of Responder Statistics
`IKE SA Rekey CREATE_CHILD_SA exchange stats`	• `Request Out`—Number of IKE SA rekey `CREATE_CHILD_SA` request message sent by initiator. • `Response In`—Number of IKE SA rekey `CREATE_CHILD_SA` response message received by initiator. • `No Proposal Chosen In`—Number of IKE SA rekey `CREATE_CHILD_SA NO_PROPSAL_CHOSEN` notification message received by initiator. • `Invalid KE In`—Number of IKE SA rekey `CREATE_CHILD_SA INVALID_KE_PAYLOAD` notification message received by initiator. • `Res DH Compute Key Fail`—Number of IKE SA rekey `CREATE_CHILD_SA` response message processing failed during verification of Diffie-Hellman compute key at initiator. • `Res Verify SA Fail`—Number of IKE SA rekey `CREATE_CHILD_SA` response message processing failed during verification of peer SA failed at initiator. • `Res Fill IKE SA Fail`—Number of IKE SA rekey `CREATE_CHILD_SA` response message processing failed during IKE SA fill operation at initiator. • `Res Verify DH Group Fail`—Number of IKE SA rekey `CREATE_CHILD_SA` response message processing failed during verification of Diffie-Hellman group at initiator.	• `Request In`—Number of IKE SA rekey `CREATE_CHILD_SA` request message received by responder. • `Response Out`—Number of IKE SA rekey `CREATE_CHILD_SA` response message sent by responder. • `No Proposal Chosen Out`—Number of IKE SA rekey `CREATE_CHILD_SA NO_PROPSAL_CHOSEN` notification message sent by responder. • `Invalid KE Out`—Number of IKE SA rekey `CREATE_CHILD_SA INVALID_KE_PAYLOAD` notification message sent by responder. • `Res DH Compute Key Fail`—Number of IKE SA rekey `CREATE_CHILD_SA` response message processing failed during Diffie-Hellman compute key at responder.

Table 1259: IKEV2_negotiaton_exchange_statistics *(Continued)*

Field Name	Field Description for Output Fields of Initiator Statistics	Field Description for Output Fields of Responder Statistics
`IPsec SA Rekey CREATE_CHILD_SA exchange stats`	• `Request Out`—Number of IPsec SA rekey `CREATE_CHILD_SA` request message sent by initiator. • `Response In`—Number of IPsec SA rekey `CREATE_CHILD_SA` response message received by initiator. • `No Proposal Chosen In`—Number of IPsec SA rekey `CREATE_CHILD_SA NO_PROPSAL_CHOSEN` notification message received by initiator. • `Invalid KE In`—Number of IPsec SA rekey `CREATE_CHILD_SA INVALID_KE_PAYLOAD` notification message received by initiator. • `TS Unacceptable In`—Number of IPsec SA rekey `CREATE_CHILD_SA TS_UNACCEPTABLE` notification message received by initiator. • `Res DH Compute Key Fail`—Number of IPsec SA rekey `CREATE_CHILD_SA` response message processing failed during verification of Diffie-Hellman compute key at initiator. • `Res Verify SA Fail`—Number of IPsec SA rekey `CREATE_CHILD_SA` response message processing failed during verification of peer SA at initiator. • `Res Verify DH Group Fail`—Number of IPsec SA rekey `CREATE_CHILD_SA` response message processing failed during verification of Diffie-Hellman group at initiator.	• `Request In`—Number of IPsec SA rekey `CREATE_CHILD_SA` request message received by responder. • `Response Out`—Number of IPsec SA rekey `CREATE_CHILD_SA` response message sent by responder. • `No Proposal Chosen Out`—Number of IPsec SA rekey `CREATE_CHILD_SA NO_PROPSAL_CHOSEN` notification message sent by responder. • `Invalid KE Out`—Number of IPsec SA rekey `CREATE_CHILD_SA INVALID_KE_PAYLOAD` notification message sent by responder. • `TS Unacceptable Out`—Number of IPsec SA rekey `CREATE_CHILD_SA TS_UNACCEPTABLE` notification message sent by responder. • `Res DH Compute Key Fail`—Number of IPsec SA rekey `CREATE_CHILD_SA` response message processing failed during Diffie-Hellman compute key at responder.

Table 1259: IKEV2_negotiaton_exchange_statistics *(Continued)*

Field Name	Field Description for Output Fields of Initiator Statistics	Field Description for Output Fields of Responder Statistics
	• `Res Verify TS Fail`—Number of IPsec SA rekey `CREATE_CHILD_SA` response message processing failed during verification of TS at initiator.	

Table 1260: IKEv2_negotiation_message_failure_statistics

Field Name	Field Description
`Discarded`	The total number of discarded messages.
`Integrity fail`	The total number of messages with integrity check failure.
`Invalid exchange type`	The total number of messages with invalid exchange type failure.
`Disorder`	The total number of messages failure due to disorder.
`ID error`	The total number of messages with ID error.
`Invalid SPI`	The total number of messages with invalid SPI failure.
`Invalid length`	The total number of messages with invalid length failure.

Sample Output

show security ike stats brief

```
user@host> show security ike stats brief
Total IKE SA and Tunnel Count Statistics:
  Number of IKE SAs: 2             Number of IPsec Tunnels: 2
```

```
IKE_SA_INIT exchange stats:
 Initiator stats:                                    Responder stats:
                                                      Request In               : 4
                                                      Response Out             : 4

IKE_AUTH exchange stats:
 Initiator stats:                                    Responder stats:
                                                      Request In               : 4
                                                      Response Out             : 4

IKE SA Rekey CREATE_CHILD_SA exchange stats:
 Initiator stats:                                    Responder stats:
  Request Out             : 1                         Request In             : 1
  Response In             : 1                         Response Out           : 1

IPsec SA Rekey CREATE_CHILD_SA exchange stats:
 Initiator stats:                                    Responder stats:
  Request Out             : 1537
  Response In             : 1537
```

Sample Output

show security ike stats detail

```
user@host> show security ike stats detail
Total IKE SA and Tunnel Count Statistics:
  Number of IKE SAs: 2             Number of IPsec Tunnels: 2

IKE_SA_INIT exchange stats:
 Initiator stats:                                    Responder stats:
  Request Out             : 0                         Request In               : 4
  Response In             : 0                         Response Out             : 4
  Invalid KE Payload In   : 0                         Invalid KE Payload Out   : 0
  No Proposal Chosen In   : 0                         No Proposal Chosen Out   : 0
  Cookie Request In       : 0                         Cookie Request Out       : 0
  Cookie Response Out     : 0                         Cookie Response In       : 0
  Res Invalid IKE SPI     : 0                         Res DH Gen Key Fail      : 0
  Res Verify SA Fail      : 0                         Res Invalid DH Group Conf: 0
  Res IKE SA Fill Fail    : 0                         Res Get CAs Fail         : 0
  Res Verify DH Group Fail: 0                         Res Get VID Fail         : 0
  Res DH Compute Key Fail : 0                         Res DH Compute Key Fail  : 0
```

```
IKE_AUTH exchange stats:
 Initiator stats:                                  Responder stats:
  Request Out             : 0                       Request In                : 4
  Response In             : 0                       Response Out              : 4
  No Proposal Chosen In   : 0                       No Proposal Chosen Out    : 0
  TS Unacceptable In      : 0                       TS Unacceptable Out       : 0
  Authentication Failed In: 0                       Authentication Failed Out: 0

IKE SA Rekey CREATE_CHILD_SA exchange stats:
 Initiator stats:                                  Responder stats:
  Request Out             : 1                       Request In              : 1
  Response In             : 1                       Response Out            : 1
  No Proposal Chosen In   : 0                       No Proposal Chosen Out : 0
  Invalid KE In           : 0                       Invalid KE Out          : 0
  Res DH Compute Key Fail : 0                       Res DH Compute Key Fail: 0
  Res Verify SA Fail      : 0
  Res Fill IKE SA Fail    : 0
  Res Verify DH Group Fail: 0

IPsec SA Rekey CREATE_CHILD_SA exchange stats:
 Initiator stats:                                  Responder stats:
  Request Out             : 1537                    Request In              : 0
  Response In             : 1537                    Response Out            : 0
  No Proposal Chosen In   : 0                       No Proposal Chosen Out : 0
  Invalid KE In           : 0                       Invalid KE Out          : 0
  TS Unacceptable In      : 0                       TS Unacceptable Out     : 0
  Res DH Compute Key Fail : 0                       Res DH Compute Key Fail: 0
  Res Verify SA Fail      : 0
  Res Verify DH Group Fail: 0
  Res Verify TS Fail      : 0

Total IKE message failure stats:
  Discarded            : 0                          ID error      : 0
  Integrity fail       : 0                          Invalid SPI   : 0
  Invalid exchange type: 0                          Invalid length: 0
  Disorder             : 0
```

Sample Output

show security ike stats brief

The command displays IKE statistics for the VPN tunnel. Notice that the output displays the number of times there is a DPD failover. The output field `Number of DPD failovers` is displayed only if there is a failover.

```
user@host> show security ike stats brief
Total IKE SA and Tunnel Count Statistics:
  Number of IKE SAs: 1             Number of IPsec Tunnels: 1
  Number of DPD failovers: 16

IKE_SA_INIT exchange stats:
 Initiator stats:                                  Responder stats:
  Request Out             : 9                       Request In               :
169
  Response In             : 9
                                                    No Proposal Chosen Out   :
169

IKE_AUTH exchange stats:
 Initiator stats:                                  Responder stats:
  Request Out             : 9
  Response In             : 9

IKE SA Rekey CREATE_CHILD_SA exchange stats:
 Initiator stats:                                  Responder stats:
  Request Out             : 1
  Response In             : 1

IPsec SA Rekey CREATE_CHILD_SA exchange stats:
 Initiator stats:                                  Responder stats:
  Request Out             : 5                       Request In             : 4
  Response In             : 5                       Response Out           : 4

Total IKE message failure stats:
  Discarded            : 7
                                                    Invalid length: 2
```

show security ike stats

The command displays IKE statistics for the VPN tunnel. Starting Junos OS 23.4R1, you can see the count of half open IKE SAs listed in the output field `DOWN`.

```
user@host> show security ike stats
Total IKE SA and Tunnel Count Statistics:
  Total IKE SAs: 1500   UP: 1136   DOWN: 364
    Initiator stats:
      Total IKE SAs: 1000   UP: 718   DOWN: 282
    Responder stats:
      Total IKE SAs: 500   UP: 418   DOWN: 82

  Total Tunnels: 10
    Active Tunnels (SAs): 5 (5)   Inactive Tunnels: 5
```

Release Information

Command introduced in Junos OS Release 19.4R1.

CLI options `brief` and `detail` are introduced in Junos OS Release 20.1R1.

The output field `Number of DPD failovers` is introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

Phase 1 of IKE Tunnel Negotiation

Phase 2 of IKE Tunnel Negotiation

show security ike tunnel-map

IN THIS SECTION

- Syntax | 24682
- Description | 24682
- Options | 24682

- Required Privilege Level | 24683
- Output Fields | 24683
- Sample Output | 24683
- Release Information | 24685

Syntax

```
show security ike tunnel-map (<brief | summary>) <fpc slot-number> <kmd-instance (all | kmd-
instance-name)> <pic slot-number>
```

Description

Display the tunnel mapping on different Services Processing Units (SPUs) for site-to-site and manual VPNs. You can insert an SPC on a device in a chassis cluster without disrupting traffic on the existing VPN tunnels. After inserting the SPC, you can view the tunnel mapping using this command. This feature is supported only on SRX5400, SRX5600, and SRX5800 Series Firewalls and vSRX Virtual Firewall instances.

Options

brief Display standard information about all existing IKE SAs. This is the default.

fpc *slot-number* Display information about existing IKE SAs in the specified Flexible PIC Concentrator (FPC) slot.

kmd-instance (all | *kmd-instance-name*) (Optional) Display information about existing IKE SAs in the key management process (KMD) identified by FPC *slot-number* and PIC *slot-number*. This option is used to filter the output. You can specify one of the following options:

- all—All KMD instances running on the Services Processing Unit (SPU).
- *kmd-instance-name*—Name of the KMD instance running on the SPU.

pic *slot-number* Display information about existing IKE SAs in the specified PIC slot.

summary Display the tunnel-mapping load on each SPU. The load is the number of times an SPU has been chosen as an anchor SPU. For site-to-site VPNs, the load should be equal to the number of gateways mapped to an SPU.

Required Privilege Level

view

Output Fields

Table 1261 on page 24683 lists the output fields for the `show security ike tunnel-map` command. Output fields are listed in the approximate order in which they appear.

Table 1261: show security ike tunnel-map Output Fields

Field Name	Field Descripton
Gateway ID	Gateway identifier. This is a nondeterministic number that is constant as long as the configuration is present. This number does not appear in any other outputs.
Gateway Name	Name of the IKE gateway.
FPC	FPC slot number.
PIC	PIC slot number.
IKED Instance	IKE process instance identifier.
SPU Load	Number of times an SPU has been chosen as an anchor SPU.

Sample Output

show security ike tunnel-map

```
user@host> show security ike tunnel-map
Gateway ID   Gateway Name FPC  PIC  IKED Instance
    2        ike_gw1       4    0       1
```

```
    3        ike_gw2       7    0       1
    4        ike_gw3       7    0       2
    5        ike_gw4       4    0       2
```

show security ike tunnel-map brief

```
user@host> show security ike tunnel-map brief
Gateway ID   Gateway Name FPC  PIC  IKED Instance
    2        gw-01         1    0       1
    3        LAN_1         1    0       2
    4        LAN_2         1    0       1
    5        LAN_3         1    0       2
    6        LAN_4         1    0       1
```

show security ike tunnel-map fpc 1 pic 0

```
user@host> run show security ike tunnel-map fpc 1 pic 0
Gateway ID   Gateway Name FPC  PIC  IKED Instance
    2        gw-01         1    0       1
    3        LAN_1         1    0       2
    4        LAN_2         1    0       1
    5        LAN_3         1    0       2
    6        LAN_4         1    0       1
```

show security ike tunnel-map kmd-instance kmd1

```
user@host> show security ike tunnel-map kmd-instance kmd1
Gateway ID   Gateway Name FPC  PIC  IKED Instance
    2        gw-01         1    0       1
    4        LAN_2         1    0       1
    6        LAN_4         1    0       1
```

show security ike tunnel-map kmd-instance all

```
user@host> show security ike tunnel-map kmd-instance all
Gateway ID   Gateway Name FPC  PIC  IKED Instance
```

```
    2        gw-01        1    0       1
    3        LAN_1        1    0       2
    4        LAN_2        1    0       1
    5        LAN_3        1    0       2
    6        LAN_4        1    0       1
```

show security ike tunnel-map summary

```
user@host> show security ike tunnel-map summary
FPC  PIC  SPU Load
1    0      5
```

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

VPN Support for Inserting Services Processing Cards

show security integrity

IN THIS SECTION

- Syntax | 24686
- Description | 24686
- Options | 24686
- Required Privilege Level | 24686
- Release Information | 24686

Syntax

```
show security integrity (extended-keyring | system-keystore)
```

Description

Display the location of public/private keys on the disk or the list of keys loaded into the IMA extended keyring.

Options

extended-keyring	Display the contents of the IMA extended keyring.
system-keystore	Display the location of public/private keys on disk.

Required Privilege Level

maintenance

Release Information

Statement introduced in Junos OS Evolved Release 20.1R1.

show security internal-security-association

IN THIS SECTION

- Syntax | 24687
- Description | 24687
- Required Privilege Level | 24687
- Output Fields | 24687
- Sample Output | 24687
- Release Information | 24688

Syntax

```
show security internal-security-association
```

Description

Provide secure login by enabling the internal security association in a chassis cluster configuration.

Required Privilege Level

view

Output Fields

Table 1262 on page 24687 lists the output fields for the `show security internal-security-association` command. Output fields are listed in the approximate order in which they appear.

Table 1262: show security internal-security-association Output Fields

Field Name	Field Description
`Internal SA Status`	State of the internal SA option on the chassis cluster control link: `enabled` or `disabled`.
`Iked Encryption Status`	State of the iked encryption.

Sample Output

show security internal-security-association

```
user@host>show security internal-security-association

node0:
--------------------------------------------------------------------------
Internal SA Status : Enabled
Iked Encryption Status : Enabled
```

```
node1:
--------------------------------------------------------------------------
Internal SA Status : Enabled
Iked Encryption Status : Enabled
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Chassis Cluster User Guide for SRX Series Devices

show security ipsec control-plane-security-associations

IN THIS SECTION

- Syntax | 24688
- Description | 24689
- Options | 24689
- Required Privilege Level | 24689
- Output Fields | 24689
- Sample Output | 24690
- Release Information | 24691

Syntax

```
show security ipsec control-plane-security-associations
<brief | detail>
<sa-name sa-name>
```

Description

Display information about manual IPsec security associations (SAs) applied to OSPF or OSPFv3 interfaces or virtual links.

Options

- `brief` | `detail`—(Optional) Display the specified level of output.
- `sa-name` *`sa-name`*—Name of the manual SA.

Required Privilege Level

view

Output Fields

Table 1263 on page 24689 lists the output fields for the `show security ipsec control-plane-security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 1263: show security ipsec control-plane-security-associations Output Fields

Field Name	Field Description
`Name`	Name of the SA.
`Algorithm`	IPsec protocol followed by encryption algorithm and authentication algorithm.
`SPI`	SPI value.
`Total active security-associations`	Total number of active manual SAs for application to OSPF or OSPFv3 interfaces or virtual links.

Sample Output

show security ipsec control-plane-security-associations

```
user@host> show security ipsec control-plane-security-associations
Name      Algorithm       SPI
test_sa   ESP:3des/md5   3e8
test_sa   ESP:3des/md5   3e8
test_sa2  ESP:3des/sha1  7d1
test_sa2  ESP:3des/sha1  7d1
Total active security-associations: 2
```

show security ipsec control-plane-security-associations sa-name

```
user@host> show security ipsec control-plane-security-associations sa-name test_sa
Name      Algorithm       SPI
test_sa   ESP:3des/md5    3e8
test_sa   ESP:3des/md5    3e8
Total active security-associations: 1
```

show security ipsec control-plane-security-associations detail

```
user@host> show security ipsec control-plane-security-associations detail
Direction: inbound, SA Name: test_sa,
Protocol: ESP:, Authentication: md5
SPI: 3e8, AUX-SPI: 0,
Mode: transport, Type: manual,
ID: 1,

Direction: outbound, SA Name: test_sa,
Protocol: ESP:, Authentication: md5
SPI: 3e8, AUX-SPI: 0,
Mode: transport, Type: manual,
ID: 2,

Direction: inbound, SA Name: test_sa2,
Protocol: ESP:, Authentication: sha1
SPI: 7d1, AUX-SPI: 0,
```

```
Mode: transport, Type: manual,
ID: 3,

Direction: outbound, SA Name: test_sa2,
Protocol: ESP:, Authentication: sha1
SPI: 7d1, AUX-SPI: 0,
Mode: transport, Type: manual,
ID: 4,
```

Release Information

Command introduced in Junos OS Release 12.1X46-D20.

RELATED DOCUMENTATION

Understanding OSPF and OSPFv3 Authentication on SRX Series Devices

show security ipsec inactive-tunnels

IN THIS SECTION

- Syntax | 24691
- Description | 24692
- Options | 24692
- Required Privilege Level | 24693
- Output Fields | 24693
- Sample Output | 24695
- Release Information | 24698

Syntax

```
show security ipsec inactive-tunnels
brief | detail
```

```
family (inet  | inet6)
fpc slot-number
index index-number
kmd-instance (all | kmd-instance-name)
node-local
pic slot-number
srg-id id-number
sa-type shortcut
vpn-name vpn-name
```

Description

Display security information about the inactive tunnel.

Options

- none—Display information about all inactive tunnels.
- `brief` | `detail`—(Optional) Display the specified level of output.
- `family`—(Optional) Display the inactive tunnel by family. This option is used to filter the output.
  - `inet`—IPv4 address family.
  - `inet6`—IPv6 address family.
- `fpc` *`slot-number`*—(Optional) Display information about inactive tunnels in the Flexible PIC Concentrator (FPC) slot.
- `index` *`index-number`*—(Optional) Display detailed information about the specified inactive tunnel identified by this index number. For a list of all inactive tunnels with their index numbers, use the command with no options.
- `kmd-instance` —(Optional) Display information about inactive tunnels in the key management process (in this case, it is KMD) identified by FPC *slot-number* and PIC *slot-number*.
  - `all`—All KMD instances running on the Services Processing Unit (SPU).
  - *`kmd-instance-name`*—Name of the KMD instance running on the SPU.
- `node-local`—(Optional) Display information about inactive tunnels for node-local tunnels in a Multinode High Availability setup.
- `pic` *`slot-number`*—Display information about inactive tunnels in the PIC slot.

- `sa-type shortcut`—(Optional) It's applicable for ADVPN. Display information about inactive tunnels by type `shortcut`.
- `vpn-name` *`vpn-name`*—(Optional) Name of the VPN.
- `srg-id`*`id-number`*—(Optional) Display information related to a specific services redundancy group (SRG) in a Multinode High Availability setup.

The `fpc` *`slot-number`*, `kmd-instance (all |` *`kmd-instance-name`*`)`, and `pic` *`slot-number`* parameters apply to SRX5600 and SRX5800 devices only.

Required Privilege Level

view

Output Fields

Table 1264 on page 24693 lists the output fields for the `show security ipsec inactive-tunnels` command. Output fields are listed in the approximate order in which they appear.

Table 1264: show security ipsec inactive-tunnels Output Fields

Field Name	Field Description
`Total inactive tunnels`	Total number of inactive IPsec tunnels.
`Total inactive tunnels which establish immediately`	Total number of inactive IPsec tunnels that can establish a session immediately.
`ID`	Identification number of the inactive tunnel. You can use this number to get more information about the inactive tunnel.
`Gateway`	IP address of the remote gateway.
`Port`	If Network Address Translation (NAT) is used, this value is 4500. Otherwise, it is the standard IKE port, 500.
`Def-Del#`	Number of deferred deletions of a dial-up IPsec VPN.

Table 1264: show security ipsec inactive-tunnels Output Fields *(Continued)*

Field Name	Field Description
`Virtual system`	Virtual system to which the VPN belongs.
`VPN name`	Name of the IPsec VPN.
`Local gateway`	Gateway address of the local system.
`Remote gateway`	Gateway address of the remote system.
`Traffic Selector Name`	For IPsec running KMD process - • Displays the name only when traffic selector is configured. • Doesn't display anything if traffic selector is not configured. For IPsec running IKED-NG process, by default - • Displays the name when traffic selector is configured. • Displays the name as *default_proxyid* when `proxy-identity` is configured. • Displays the name as *default_any_any* when traffic selector is not configured. See "show security ipsec inactive-tunnels detail" on page 24696, for more details.
`Local identity`	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IP address, fully qualified domain name, e-mail address, or distinguished name (DN). Displays `proxy-identity` when configured for IPsec running either KMD or IKED-NG process. Displays *0.0.0.0* when `proxy-identity` is not configured. See "show security ipsec inactive-tunnels detail" on page 24696, for more details.

Table 1264: show security ipsec inactive-tunnels Output Fields *(Continued)*

Field Name	Field Description
`Remote identity`	Identity of the destination peer gateway. The value is specified as an IP address, fully qualified domain name, e-mail address, or distinguished name (DN). Displays `proxy-identity` when configured for IPsec running either KMD or IKED-NG process. Displays *0.0.0.0* when `proxy-identity` is not configured. See "show security ipsec inactive-tunnels detail" on page 24696, for more details.
`Version`	Version of IKE.
`Passive Mode Tunneling`	IPsec tunneling of malformed packets; enabled if set or disabled if not set.
`DF-bit`	State of the don't fragment bit: `set` or `clear`.
`Bind-interface`	The tunnel interface to which the route-based VPN is bound.
`Policy-name`	Name of the applicable policy.
`Tunnel Down Reason`	Reason for which the tunnel is inactive.
`Tunnel events`	Tunnel event and the number of times the event has occurred. See Tunnel Events for descriptions of tunnel events and the action you can take.

Sample Output

show security ipsec inactive-tunnels

```
user@host> show security ipsec inactive-tunnels
Total inactive tunnels: 1
  Total inactive tunnels with establish immediately: 0
```

```
  ID     Gateway    Port  Tunnel down reason
  131073 192.168.1.2  500   Phase1 proposal mismatch detected
```

show security ipsec inactive-tunnels detail

For IPsec running KMD process, when both `proxy-identity` and `traffic-selector` are not configured.

```
user@host> show security ipsec inactive-tunnels detail
ID: 131073 Virtual-system: root, VPN Name: vpn1
  Local Gateway: 192.12.0.20, Remote Gateway: 192.12.0.10
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
```

For IPsec running KMD process, when `proxy-identity` is configured.

```
user@host> show security ipsec inactive-tunnels detail
ID: 131074 Virtual-system: root, VPN Name: vpn1
  Local Gateway: 192.12.0.20, Remote Gateway: 192.12.0.10
  Local Identity: ipv4_subnet(any:0,[0..7]=1.0.0.0/8)
  Remote Identity: ipv4_subnet(any:0,[0..7]=4.0.0.0/8)
```

For IPsec running KMD process, when `traffic-selector` is configured.

```
user@host> show security ipsec inactive-tunnels detail
ID: 67108865 Virtual-system: root, VPN Name: vpn1
  Local Gateway: 192.12.0.20, Remote Gateway: 192.12.0.10
  Traffic Selector Name: ts1
  Local Identity: ipv4(1.0.0.0-1.255.255.255)
  Remote Identity: ipv4(4.0.0.0-4.255.255.255)
```

For IPsec running IKED-NG process, when `traffic-selector` is configured.

```
user@host> show security ipsec inactive-tunnels detail
ID: 105 Virtual-system: root, VPN Name: IPSEC_VPN
  Local Gateway: 2.0.0.1, Remote Gateway: 2.0.0.5
  Traffic Selector Name: ts1
  Local Identity: ipv4(1.0.0.0-1.255.255.255)
  Remote Identity: ipv4(4.0.0.0-4.255.255.255)
```

For IPsec running IKED-NG process, when `traffic-selector` is not configured.

```
user@host> show security ipsec inactive-tunnels detail
ID: 107 Virtual-system: root, VPN Name: IPSEC_VPN
  Local Gateway: 2.0.0.1, Remote Gateway: 2.0.0.5
  Traffic Selector Name: default_any_any_v4
  Local Identity: ipv4(0.0.0.0-255.255.255.255)
  Remote Identity: ipv4(0.0.0.0-255.255.255.255)
```

For IPsec running IKED-NG process, when `proxy-identity` is configured.

```
user@host> show security ipsec inactive-tunnels detail
ID: 110 Virtual-system: root, VPN Name: IPSEC_VPN
  Local Gateway: 2.0.0.1, Remote Gateway: 2.0.0.5
  Traffic Selector Name: default_proxyid
  Local Identity: ipv4(1.0.0.0-1.255.255.255)
  Remote Identity: ipv4(4.0.0.0-4.255.255.255)
```

show security ipsec inactive-tunnels index 131073

```
user@host> show security ipsec inactive-tunnels index 131073
ID: 131073 Virtual-system: root, VPN Name: vpn1
  Local Gateway: 192.168.1.100, Remote Gateway: 192.168.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.0
  Port: 500, Nego#: 2, Fail#: 0, Def-Del#: 0 Flag: 600a29
  Tunnel events:
    Wed Jul 16 2014 06:18:02 +0800: User cleared IPSec SA from CLI (1 times)
    Wed Jul 16 2014 06:17:58 +0800: IPSec SA negotiation successfully completed (1 times)
    Wed Jul 16 2014 06:17:54 +0800: User cleared IPSec SA from CLI (1 times)
    Wed Jul 16 2014 06:16:58 +0800: IPSec SA negotiation successfully completed (1 times)
    Wed Jul 16 2014 06:16:58 +0800: Bind interface's address received. Information updated (1
times)
    Wed Jul 16 2014 06:16:58 +0800: Tunnel is ready. Waiting for trigger event or peer to
trigger negotiation (1 times)
    Wed Jul 16 2014 06:16:58 +0800: External interface's address received. Information updated
(1 times)
```

```
    Wed Jul 16 2014 06:16:58 +0800: Bind interface's zone received. Information updated (1 times)
    Wed Jul 16 2014 06:16:58 +0800: IKE SA negotiation successfully completed (1 times)
```

show security ipsec inactive-tunnels sa-type shortcut

```
user@host> show security ipsec inactive-tunnels sa-type shortcut
  Total inactive tunnels: 1
  Total inactive tunnels with establish immediately: 0
  ID     Port  Nego#  Fail#  Flag      Gateway          Tunnel Down Reason
  268173322 500 0     0      40608aa9  192.168.0.105       Cleared via CLI
```

show security ipsec inactive-tunnels (MX-SPC3 and SRX4600) with passive mode tunneling

```
user@host>show security ipsec inactive-tunnels
  ID: 6 Virtual-system: root, VPN Name: vpn2
  Local Gateway: 10.0.0.2, Remote Gateway: 30.0.0.2
  Traffic Selector Name: ts2
  Local Identity: ipv4(50.0.1.0-50.0.1.255)
  Remote Identity: ipv4(140.0.1.0-140.0.1.255)
  Version: IKEv2
  Passive mode tunneling: Disabled
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1, Policy-name: ipsec_policy
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
```

show security ipsec inactive-tunnels node-local

```
user@host>show security ipsec inactive-tunnels node-local
  Total inactive tunnels: 0
  Total inactive tunnels with establish immediately: 0
```

Release Information

Command introduced in Junos OS Release 11.4R3.

Support for the `passive-mode-tunneling` option on MX-SPC3 is added in Junos OS Release 23.1R1.

Support for the `node-local` option is added in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

show security ipsec security-associations

show security ipsec next-hop-tunnels

IN THIS SECTION

- Syntax | 24699
- Description | 24699
- Options | 24699
- Required Privilege Level | 24700
- Output Fields | 24700
- Sample Output | 24701
- Release Information | 24701

Syntax

```
show security ipsec next-hop-tunnels {
    family (inet | inet6);
    index;
    interface-name;
}
```

Description

Display security information about the secure tunnel interface.

Options

family Display IPSec next-hop-tunnel entries by family.

index Index of security association.

- Range:
  - 1 through 4294967295

inet Displays IPv4 protocol parameters.

inet6 Displays IPv6 protocol parameters.

interface-name Name of the secure tunnel logical interface.

Required Privilege Level

view

Output Fields

Table 1265 on page 24700 lists the output fields for the `show security ipsec next-hop-tunnels` command. Output fields are listed in the approximate order in which they appear.

Table 1265: show security ipsec next-hop-tunnels Output Fields

Field Name	Field Description
`Next-hop gateway`	IP address of the next gateway.
`Interface`	Name of the secure tunnel logical interface.
`IPsec VPN name`	Name of the IPsec VPN tunnel.
`Flag`	• `Static`—IP address manually configured. • `Auto`—IP address obtained from the remote peer automatically.

Sample Output

show security ipsec next-hop-tunnels family inet

```
user@host> show security ipsec next-hop-tunnels inet
Next-hop gateway     interface   IPsec VPN name                    Flag
192.168.1.2          st0.0       autokey                           Static
192.168.1.3          st0.0       pbd-4-6                           Auto
```

show security ipsec next-hop-tunnels family inet6

```
user@host> show security ipsec next-hop-tunnels family inet6
Next-hop gateway                interface   IPSec VPN name            Flag
2001:db8::2                     st0.1       IPSEC_VPNA_1              Auto
2001:db8::3                     st0.1       IPSEC_VPNA_1              Auto
2001:fe80::5668:ad10:fcd8:59db  st0.1       IPSEC_VPNA_1              Auto
2001:fe80::5668:ad10:fcd8:5aa5  st0.1       IPSEC_VPNA_1              Auto
```

Release Information

Command introduced in Junos OS Release 8.5.

The `family inet6` option is introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

show security ipsec security-associations | 24701

show security ipsec security-associations

IN THIS SECTION

- Syntax | 24702
- Description | 24702

- Options | 24702
- Required Privilege Level | 24704
- Output Fields | 24704
- Sample Output | 24715
- Release Information | 24743

Syntax

```
show security ipsec security-associations
<brief | detail | extensive>
<family (inet  | inet6)>
<fpc slot-number pic slot-number>
<index SA-index-number>
<kmd-instance (all | kmd-instance-name)>
<node-local>
<pic slot-number fpc slot-number>
<sa-type shortcut>
<traffic-selector traffic-selector-name>
<srg-id id-number>
<vpn-name vpn-name>
<ha-link-encryption>
```

Description

Display information about the IPsec security associations (SAs).

In Junos OS Releases 20.1R2, 20.2R2, 20.3R2, 20.3R1, and later, when you execute the `show security ipsec security-associations detail` command, a new output field `IKE SA Index` corresponding to every IPsec SA within a tunnel is displayed under each IPsec SA information. See "show security ipsec security-associations detail (SRX5400, SRX5600, SRX5800)" on page 24730.

Options

none	Display information about all SAs.
brief | detail | extensive	(Optional) Display the specified level of output. The default is `brief`.

family

(Optional) Display SAs by family. This option is used to filter the output.

- `inet`—IPv4 address family.
- `inet6`—IPv6 address family.

fpc *slot-number* **pic** *slot-number*

(Optional) Display information about existing IPsec SAs in the specified Flexible PIC Concentrator (FPC) slot and PIC slot.

In a chassis cluster, when you execute the CLI command `show security ipsec security-associations pic <slot-number> fpc <slot-number>` in operational mode, only the primary node information about the existing IPsec SAs in the specified Flexible PIC Concentrator (FPC) slot and PIC slot is displayed.

index *SA-index-number*

(Optional) Display detailed information about the specified SA identified by this index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.

kmd-instance

(Optional) Display information about existing IPsec SAs in the key management process (in this case, it is KMD) identified by the FPC *slot-number* and PIC *slot-number*.

This option is applicable when you have `kmd` process for IPsec VPN features. This option is not available when you enable `iked` process using `junos-ike` package for running IPsec VPN features.

- `all`—All KMD instances running on the Services Processing Unit (SPU).
- *kmd-instance-name*—Name of the KMD instance running on the SPU.

node-local

—(Optional) Display information about IPsec SAs for node-local tunnels in a Multinode High Availability setup.

pic *slot-number* **fpc** *slot-number*

(Optional) Display information about existing IPsec SAs in the specified PIC slot and FPC slot.

sa-type shortcut

(Optional) It's applicable for ADVPN. Display information about IPsec SAs by type `shortcut`.

traffic-selector *traffic-selector-name*

(Optional) Display information about the specified traffic selector.

vpn-name *vpn-name*

(Optional) Display information about the specified VPN.

ha-link-encryption (Optional) Display information related to interchassis link tunnel only. See "ipsec (High Availability)" on page 9301, "show security ipsec security-associations ha-link-encryption (SRX5400, SRX5600, SRX5800)" on page 24731, and "show security ipsec sa detail ha-link-encryption (SRX5400, SRX5600, SRX5800)" on page 24732.

srg-id (Optional) Display information related to a specific services redundancy group (SRG) in a Multinode High Availability setup.

Required Privilege Level

view

Output Fields

Table 1266 on page 24704 lists the output fields for the `show security ipsec security-associations` command, Table 1267 on page 24710 lists the output fields for the `show security ipsec sa` command and Table 1268 on page 24712. lists the output fields for the `show security ipsec sa detail`. Output fields are listed in the approximate order in which they appear.

Table 1266: show security ipsec security-associations

Field Name	Field Description	Level of Output
`Total active tunnels`	Total number of active IPsec tunnels.	`brief`
`ID`	Index number of the SA. You can use this number to get additional information about the SA.	All levels
`Algorithm`	Cryptography used to secure exchanges between peers during the IKE negotiations includes: • An authentication algorithm used to authenticate exchanges between the peers. • An encryption algorithm used to encrypt data traffic.	`brief`

Table 1266: show security ipsec security-associations *(Continued)*

Field Name	Field Description	Level of Output
`SPI`	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, and keys. The peer gateways each have two SAs, one resulting from each of the two phases of negotiation: IKE and IPsec.	`brief`
`Life: sec/kb`	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.	`brief`
`Mon`	The Mon field refers to VPN monitoring status. If VPN monitoring is enabled, then this field displays `U` (up) or `D` (down). A hyphen (-) means VPN monitoring is not enabled for this SA. A `V` means that IPsec datapath verification is in progress.	`brief`
`lsys`	The root system.	`brief`
`Port`	If Network Address Translation (NAT) is used, this value is 4500. Otherwise, it is the standard IKE port, 500.	All levels
`Gateway`	IP address of the remote gateway.	`brief`
`Virtual-system`	Name of the logical system.	`detail`, `extensive`
`VPN name`	IPsec name for VPN.	`detail`, `extensive`

Table 1266: show security ipsec security-associations *(Continued)*

Field Name	Field Description	Level of Output
State	State has two options, Installed and Not Installed. • Installed—The SA is installed in the SA database. • Not Installed—The SA is not installed in the SA database. For transport mode, the value of State is always Installed.	detail, extensive
Local gateway	Gateway address of the local system.	detail, extensive
Remote gateway	Gateway address of the remote system.	detail, extensive
Traffic selector	Name of the traffic selector.	detail, extensive
Local identity	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IP address, fully qualified domain name, e-mail address, or distinguished name (DN).	detail, extensive
Remote identity	IP address of the destination peer gateway.	detail, extensive
Term	Defines local IP range, remote IP range, source port range, destination port range, and protocol.	detail, extensive
Source-port	Source port range configured for a term.	detail, extensive
Destination-Port	Destination port range configured for a term.	detail, extensive

Table 1266: show security ipsec security-associations *(Continued)*

Field Name	Field Description	Level of Output
`Version`	IKE version, either `IKEv1` or `IKEv2`.	`detail`, `extensive`
`DF-bit`	State of the don't fragment bit: `set` or `cleared`.	`detail`, `extensive`
`Location`	`FPC`—Flexible PIC Concentrator (FPC) slot number. `PIC`—PIC slot number. `KMD-Instance`—The name of the KMD instance running on the SPU, identified by FPC *slot-number* and PIC *slot-number*. Currently, 4 KMD instances running on each SPU, and any particular IPsec negotiation is carried out by a single KMD instance. This option is applicable when you have `kmd` process for IPsec VPN features. This option is not available when you enable `iked` process using `junos-ike` package for running IPsec VPN features.	`detail`, `extensive`
`Tunnel events`	Tunnel event and the number of times the event has occurred. See Tunnel Events for descriptions of tunnel events and the action you can take. • The `detail` option displays upto ten tunnel events in reverse chronological order. • The `extensive` option displays all the tunnel events.	`detail`, `extensive`
`Anchorship`	Anchor thread ID for the SA (for SRX4600 Series devices with the `detail` option).	

Table 1266: show security ipsec security-associations *(Continued)*

Field Name	Field Description	Level of Output
`Direction`	Direction of the SA; it can be inbound or outbound.	`detail`, `extensive`
`AUX-SPI`	Value of the auxiliary security parameter index(SPI). • When the value is `AH` or `ESP`, `AUX-SPI` is always 0. • When the value is `AH+ESP`, `AUX-SPI` is always a positive integer.	`detail`, `extensive`
`Mode`	Mode of the SA: • `transport`—Protects host-to-host connections. • `tunnel`–Protects connections between security gateways.	`detail`, `extensive`
`Type`	Type of the SA: • `manual`—Security parameters require no negotiation. They are static and are configured by the user. • `dynamic`—Security parameters are negotiated by the IKE protocol. Dynamic SAs are not supported in transport mode.	`detail`, `extensive`
`State`	State of the SA: • `Installed`—The SA is installed in the SA database. • `Not Installed`—The SA is not installed in the SA database. For transport mode, the value of State is always `Installed`.	`detail`, `extensive`

Table 1266: show security ipsec security-associations *(Continued)*

Field Name	Field Description	Level of Output
`Protocol`	Protocol supported. • Transport mode supports Encapsulation Security Protocol (ESP) and Authentication Header (AH). • Tunnel mode supports ESP and AH.	`detail`, `extensive`
`Authentication`	Type of authentication used.	`detail`, `extensive`
`Encryption`	Type of encryption used. Starting in Junos OS Release 19.4R2, when you configure `aes-128-gcm` or `aes-256-gcm` as an encryption algorithm at the `[edit security ipsec proposal proposal-name]` hierarchy level, the authentication algorithm field of the `show security ipsec security-associations detail` command displays the same configured encryption algorithm.	`detail`, `extensive`
`Soft lifetime`	The soft lifetime informs the IPsec key management system that the SA is about to expire. Each lifetime of an SA has two display options, hard and soft, one of which must be present for a dynamic SA. This allows the key management system to negotiate a new SA before the hard lifetime expires. • `Expires in seconds`—Number of seconds left until the SA expires.	`detail`, `extensive`
`Hard lifetime`	The hard lifetime specifies the lifetime of the SA. • `Expires in seconds`—Number of seconds left until the SA expires.	`detail`, `extensive`

Table 1266: show security ipsec security-associations *(Continued)*

Field Name	Field Description	Level of Output
`Lifesize Remaining`	The lifesize remaining specifies the usage limits in kilobytes. If there is no lifesize specified, it shows unlimited. • `Expires in kilobytes`—Number of kilobytes left until the SA expires.	`detail`, `extensive`
`Anti-replay service`	State of the service that prevents packets from being replayed. It can be `Enabled` or `Disabled`.	`detail`, `extensive`
`Replay window size`	Size of the antireplay service window, which is 64 bits.	`detail`, `extensive`
`Bind-interface`	The tunnel interface to which the route-based VPN is bound.	`detail`, `extensive`
`Copy-Outer-DSCP`	Indicates if the system copies the outer DSCP value from the IP header to the inner IP header.	`detail`, `extensive`
`tunnel-establishment`	Indicates how the IKE is activated.	`detail`, `extensive`
`IKE SA index`	Indicates the list of parent IKE security associations.	`detail`, `extensive`

Table 1267: show security ipsec sa Output Fields

Field Name	Field Description
`Total active tunnels`	Total number of active IPsec tunnels.

Table 1267: show security ipsec sa Output Fields *(Continued)*

Field Name	Field Description
`ID`	Index number of the SA. You can use this number to get additional information about the SA.
`Algorithm`	Cryptography used to secure exchanges between peers during the IKE Phase 2 negotiations includes: • An authentication algorithm used to authenticate exchanges between the peers. Options are `hmac-md5-96`, `hmac-sha-256-128`, or `hmac-sha1-96`. • An encryption algorithm used to encrypt data traffic. Options are `3des-cbc`, `aes-128-cbc`, `aes-192-cbc`, `aes-256-cbc`, `chacha20-poly1305`, or `des-cbc`.
`SPI`	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, and keys. The peer gateways each have two SAs, one resulting from each of the two phases of negotiation: Phase 1 and Phase 2.
`Life:sec/kb`	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.
`Mon`	The Mon field refers to VPN monitoring status. If VPN monitoring is enabled, then this field displays U (up) or D (down). A hyphen (-) means VPN monitoring is not enabled for this SA. A V means that IPSec datapath verification is in progress.
`lsys`	The root system.
`Port`	If Network Address Translation (NAT) is used, this value is 4500. Otherwise, it is the standard IKE port, 500.
`Gateway`	Gateway address of the system.

Table 1268: show security ipsec sa detail Output Fields

Field Name	Field Description
`FC-name`	Forwarding class (FC) name for the child security association.
`ID`	Index number of the SA. You can use this number to get additional information about the SA.
`Virtual-system`	The virtual system name.
`VPN Name`	IPSec name for VPN.
`Local Gateway`	Gateway address of the local system.
`Remote Gateway`	Gateway address of the remote system.
`Local Identity`	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IP address, fully qualified domain name, e-mail address, or distinguished name (DN).
`Remote Identity`	IP address of the destination peer gateway.
`Version`	IKE version. For example, IKEv1, IKEv2.
`Passive Mode Tunneling`	IPsec tunneling of malformed packets. You can either enable or disable the option.
`DF-bit`	State of the don't fragment bit: `set` or `cleared`.
`Bind-interface`	The tunnel interface to which the route-based VPN is bound.
`identity-management`	Count for the VPN tunnel create, delete, and rekey events when push to identity management is enabled.

Table 1268: show security ipsec sa detail Output Fields *(Continued)*

Field Name	Field Description
Tunnel Events	
`Direction`	Direction of the SA; it can be inbound or outbound.
`AUX-SPI`	Value of the auxiliary security parameter index(SPI). • When the value is `AH` or `ESP`, `AUX-SPI` is always 0. • When the value is `AH+ESP`, `AUX-SPI` is always a positive integer.
`VPN Monitoring`	If VPN monitoring is enabled, then the `Mon` field displays `U (up)` or `D (down)`. A hyphen (-) means VPN monitoring is not enabled for this SA. A V means that IPsec datapath verification is in progress. Starting Junos OS release 23.4R1 , the output displays `Interval` and `Threshold` details when the firewall runs IPsec VPN services with the iked process.
`Hard lifetime`	The hard lifetime specifies the lifetime of the SA. • `Expires in seconds` - Number of seconds left until the SA expires.
`Lifesize Remaining`	The lifesize remaining specifies the usage limits in kilobytes. If there is no lifesize specified, it shows unlimited.
`Soft lifetime`	The soft lifetime informs the IPsec key management system that the SA is about to expire. Each lifetime of an SA has two display options, hard and soft, one of which must be present for a dynamic SA. This allows the key management system to negotiate a new SA before the hard lifetime expires. • `Expires in seconds` - Number of seconds left until the SA expires.
`Mode`	Mode of the SA: • `transport` - Protects host-to-host connections. • `tunnel` - Protects connections between security gateways.

Table 1268: show security ipsec sa detail Output Fields *(Continued)*

Field Name	Field Description
`Type`	Type of the SA: • `manual` - Security parameters require no negotiation. They are static and are configured by the user. • `dynamic` - Security parameters are negotiated by the IKE protocol. Dynamic SAs are not supported in transport mode.
`State`	State of the SA: • `Installed` - The SA is installed in the SA database. • `Not Installed` - The SA is not installed in the SA database. For transport mode, the value of State is always Installed.

`Protocol`	Protocol supported.
	• Transport mode supports Encapsulation Security Protocol (ESP) and Authentication Header (AH).
	• Tunnel mode supports ESP and AH.
	  • `Authentication` - Type of authentication used.
	  • `Encryption` - Type of encryption used.

`Anti-replay service`	State of the service that prevents packets from being replayed. It can be `Enabled` or `Disabled`.
`Replay window size`	Configured size of the antireplay service window. It can be 32 or 64 packets. If the replay window size is 0, the antireplay service is disabled. The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets.
Interchassis Link Tunnel	

Table 1268: show security ipsec sa detail Output Fields *(Continued)*

Field Name	Field Description
HA Link Encryption Mode	High availability mode supported. Displays `Multi-Node` when multi-node high availability feature is enabled.

Sample Output

For brevity, the show command outputs does not display all the values of the configuration. Only a subset of the configuration is displayed. Rest of the configuration on the system has been replaced with ellipses (...).

show security ipsec security-associations (IPv4)

```
user@host> show security ipsec security-associations
  Total active tunnels: 14743 Total Ipsec sas: 14743
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <511672 ESP:aes-cbc-128/sha1 0x071b8cd2      -   root 500   10.21.45.152
  >503327 ESP:aes-cbc-128/sha1 0x69d364dd 1584/ unlim - root 500 10.21.12.255
  <503327 ESP:aes-cbc-128/sha1 0x0a577f2d 1584/ unlim - root 500 10.21.12.255
  >512896 ESP:aes-cbc-128/sha1 0xd2f51c81 1669/ unlim - root 500 10.21.50.96
  <512896 ESP:aes-cbc-128/sha1 0x071b8d9e 1669/ unlim - root 500 10.21.50.96
  >513881 ESP:aes-cbc-128/sha1 0x95955834 1696/ unlim - root 500 10.21.54.57
```

show security ipsec security-associations (IPv6)

```
user@host> show security ipsec security-associations
Total active tunnels: 1
ID     Algorithm         SPI      Life:sec/kb  Mon  vsys Port  Gateway
131074 ESP:aes256/sha256 14caf1d9 3597/ unlim   -   root 500   2001:db8::1112
131074 ESP:aes256/sha256 9a4db486 3597/ unlim   -   root 500   2001:db8::1112
```

show security ipsec security-associations index 511672

```
user@host> show security ipsec security-associations index 511672
ID: 511672 Virtual-system: root, VPN Name: ipsec_vpn
```

```
  Local Gateway: 10.20.0.1, Remote Gateway: 10.21.45.152
  Traffic Selector Name: ts
  Local Identity: ipv4(10.191.151.0-10.191.151.255)
  Remote Identity: ipv4(10.40.151.0-10.40.151.255)
  Version: IKEv2
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.0, Policy-name: IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Location: FPC 0, PIC 1, KMD-Instance 0
  Anchorship: Thread 10
  Direction: inbound, SPI: 0x835b8b42, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1639 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1257 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 0x071b8cd2, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1639 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1257 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations index 131073 detail

```
user@host> show security ipsec security-associations index 131073 detail
ID: 131073 Virtual-system: root, VPN Name: IPSEC_VPN1
  Local Gateway: 10.4.0.1, Remote Gateway: 10.5.0.1
  Local Identity: ipv4_subnet(any:0,[0..7]=10.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=10.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1
  Port: 500, Nego#: 18, Fail#: 0, Def-Del#: 0 Flag: 0x600a39
  Multi-sa, Configured SAs# 9, Negotiated SAs#: 9
  Tunnel events:
    Mon Apr 23 2018 22:20:54 -0700: IPSec SA negotiation successfully completed (1 times)
    Mon Apr 23 2018 22:20:54 -0700: IKE SA negotiation successfully completed (2 times)
```

```
    Mon Apr 23 2018 22:20:18 -0700: User cleared IKE SA from CLI, corresponding IPSec SAs
cleared (1 times)
    Mon Apr 23 2018 22:19:55 -0700: IPSec SA negotiation successfully completed (2 times)
    Mon Apr 23 2018 22:19:23 -0700: Tunnel is ready. Waiting for trigger event or peer to
trigger negotiation (1 times)
    Mon Apr 23 2018 22:19:23 -0700: Bind-interface's zone received. Information updated (1 times)
    Mon Apr 23 2018 22:19:23 -0700: External interface's zone received. Information updated (1
times)
  Direction: inbound, SPI: 2d8e710b, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1930 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1563 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    Multi-sa FC Name: default
  Direction: outbound, SPI: 5f3a3239, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1930 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1563 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes-256-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    Multi-sa FC Name: default
  Direction: inbound, SPI: 5d227e19, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1930 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1551 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes-256-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    Multi-sa FC Name: best-effort
  Direction: outbound, SPI: 5490da, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1930 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1551 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes-256-cbc
```

```
    Anti-replay service: counter-based enabled, Replay window size: 64
...
```

Starting with Junos OS Release 18.2R1, the CLI `show security ipsec security-associations index` *`index-number`* `detail` output displays all the child SA details including forwarding class name.

show security ipsec sa

```
user@host> show security ipsec sa
Total active tunnels: 2
ID Algorithm SPI Life:sec/kb Mon lsys Port Gateway
>67108885 ESP:aes-gcm-256/None fdef4dab 2918/ unlim - root 500 2001:db8:3000::2
>67108885 ESP:aes-gcm-256/None e785dadc 2918/ unlim - root 500 2001:db8:3000::2
>67108887 ESP:aes-gcm-256/None 34a787af 2971/ unlim - root 500 2001:db8:5000::2
>67108887 ESP:aes-gcm-256/None cf57007f 2971/ unlim - root 500 2001:db8:5000::2
```

show security ipsec sa detail

```
user@host> show security ipsec sa detail
ID: 500201 Virtual-system: root, VPN Name: IPSEC_VPN
  Local Gateway: 10.2.0.1, Remote Gateway: 10.2.0.2
  Local Identity: ipv4(10.0.0.0-255.255.255.255)
  Remote Identity: ipv4(10.0.0.0-255.255.255.255)
  Version: IKEv1
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1, Policy-name: IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Location: FPC 0, PIC 1, KMD-Instance 0
  Anchorship: Thread 1
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0x0a25c960, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 91 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 44 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    tunnel-establishment: establish-tunnels-responder-only-no-rekey
  Direction: outbound, SPI: 0x43e34ad3, AUX-SPI: 0
```

```
                              , VPN Monitoring: -
    Hard lifetime: Expires in 91 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 44 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    tunnel-establishment: establish-tunnels-responder-only-no-rekey
...
```

Starting with Junos OS Release 19.1R1, a new field **tunnel-establishment** in the output of the CLI `show security ipsec sa detail` displays the option configured under `ipsec vpn establish-tunnels` hierarchy.

Starting with Junos OS Release 21.3R1, a new field **Tunnel MTU** in the output of the CLI `show security ipsec sa detail` displays the option configured under `ipsec vpn hub-to-spoke-vpn tunnel-mtu` hierarchy.

Starting in Junos OS Release 22.1R3, on SRX5000 line of devices, the Tunnel MTU is not displayed in the CLI output if the tunnel MTU is not configured.

show security ipsec sa detail (MX-SPC3)

```
user@host> show security ipsec sa detail
ID: 500055 Virtual-system: root, VPN Name: IPSEC_VPN
  Local Gateway: 10.2.0.1, Remote Gateway: 10.2.0.2
  Local Identity: ipv4(10.0.0.0-255.255.255.255)
  Remote Identity: ipv4(10.0.0.0-255.255.255.255)
  Version: IKEv2
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1, Tunnel MTU: 1420  Policy-name:
IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Location: FPC 0, PIC 0, KMD-Instance 0
  Anchorship: Thread 15
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0x229b998e, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 23904 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 23288 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-md5-96, Encryption: aes-cbc (128 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

```
    Extended-Sequence-Number: Enabled
    tunnel-establishment: establish-tunnels-immediately
  Direction: outbound, SPI: 0xb2e843a3, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 23904 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 23288 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-md5-96, Encryption: aes-cbc (128 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Enabled
    tunnel-establishment: establish-tunnels-immediately
```

show security ipsec sa detail (MX-SPC3 and SRX4600) with passive mode tunneling

```
user@host> show security ipsec sa detail
ID: 500054 Virtual-system: root, VPN Name: TUN_3
  Local Gateway: 100.0.0.3, Remote Gateway: 200.0.0.3
  Traffic Selector Name: ts1
  Local Identity: ipv4(11.0.0.3-11.0.0.3)
  Remote Identity: ipv4(75.0.0.3-75.0.0.3)
  TS Type: traffic-selector
  Version: IKEv2
  Quantum Secured: No
  PFS group: N/A
  SRG ID: 0
  Passive mode tunneling: Enabled
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.3, Policy-name: IPSEC_POLICY
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Tunnel events:
    Mon Sep 19 2022 19:27:44: IPsec SA negotiation succeeds (1 times)
  Location: FPC 3, PIC 1, KMD-Instance 0
  Anchorship: Thread 15
  Distribution-Profile: vms-3/1/0
  Direction: inbound, SPI: 0x25c03740, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expired
    Lifesize Remaining: Expired
    Soft lifetime: Expires in 2920 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
```

```
    Protocol: ESP, Authentication: aes256-gcm, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 512
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 122
  Direction: outbound, SPI: 0x8e8f2009, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expired
    Lifesize Remaining: Expired
    Soft lifetime: Expires in 2920 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes256-gcm, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 512
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 122
```

show security ipsec security-association

```
user@host>show security ipsec security-association
Total active tunnels: 1     Total IPsec sas: 1
  ID      Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <500006 ESP:aes-gcm-128/aes128-gcm 0x782b233c 1432/ unlim - root 500 10.2.0.2
```

show security ipsec security-associations brief

```
user@host> show security ipsec security-associations brief
Total active tunnels: 2     Total Ipsec sas: 18
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <131073 ESP:aes256/sha256 89e5098  1569/ unlim   -   root 500   10.5.0.1
  >131073 ESP:aes256/sha256 fcee9d54 1569/ unlim   -   root 500   10.5.0.1
  <131073 ESP:aes256/sha256 f3117676 1609/ unlim   -   root 500   10.5.0.1
  >131073 ESP:aes256/sha256 6050109f 1609/ unlim   -   root 500   10.5.0.1
  <131073 ESP:aes256/sha256 e01f54b1 1613/ unlim   -   root 500   10.5.0.1
  >131073 ESP:aes256/sha256 29a05dd6 1613/ unlim   -   root 500   10.5.0.1
  <131073 ESP:aes256/sha256 606c90f6 1616/ unlim   -   root 500   10.5.0.1
  >131073 ESP:aes256/sha256 9b5b059d 1616/ unlim   -   root 500   10.5.0.1
  <131073 ESP:aes256/sha256 b8116d6d 1619/ unlim   -   root 500   10.5.0.1
  >131073 ESP:aes256/sha256 b7ed6bfd 1619/ unlim   -   root 500   10.5.0.1
  <131073 ESP:aes256/sha256 4f5ce754 1619/ unlim   -   root 500   10.5.0.1
```

```
  >131073 ESP:aes256/sha256 af8984b6 1619/ unlim   -   root 500   10.5.0.1
...
```

show security ipsec security-associations detail

```
user@host> show security ipsec security-associations detail

ID: 500009 Virtual-system: root, VPN Name: IPSEC_VPN
  Local Gateway: 10.2.0.2, Remote Gateway: 10.2.0.1
  Local Identity: ipv4(10.0.0.0-255.255.255.255)
  Remote Identity: ipv4(10.0.0.0-255.255.255.255)
  Version: IKEv1
  PFS group: DH-group-14
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1, Policy-name: IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Location: FPC 0, PIC 0, KMD-Instance 0
  Anchorship: Thread 0
  Distribution-Profile: default-profile
  IKE SA Index: 2068
  Direction: inbound, SPI: 0xba7bb1f2, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 146 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 101 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-on-traffic
  Direction: outbound, SPI: 0x41650a1b, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 146 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 101 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-on-traffic
```

show security ipsec security-associations with VPN monitoring enabled

```
user@host> show security ipsec security-associations
  Total active tunnels: 1     Total Ipsec sas: 1
  ID    Algorithm       SPI      Life:sec/kb   Mon lsys Port  Gateway
  <131073 ESP:aes-gcm-256/None 15e8de2c 2086/ unlim U root 500 3.0.0.1
  >131073 ESP:aes-gcm-256/None 2a1f46fb 2086/ unlim U root 500 3.0.0.1
  VPN Monitoring: UP
```

show security ipsec security-associations detail with VPN monitoring enabled

```
user@host> show security ipsec security-associations detail
ID: 131073 Virtual-system: root, VPN Name: SPK_VPN1
  Local Gateway: 3.0.0.2, Remote Gateway: 3.0.0.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1
  Port: 500, Nego#: 2, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Multi-sa, Configured SAs# 1, Negotiated SAs#: 1

  Direction: inbound, SPI: 15e8de2c, AUX-SPI: 0
                              , VPN Monitoring: UP Mode: optimized  Interval:10sec Threshold: 3
    Hard lifetime: Expires in 1314 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 708 seconds
    Mode: Tunnel(10 5), Type: dynamic, State: installed
    Protocol: ESP, Authentication: None, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 2a1f46fb, AUX-SPI: 0
                              , VPN Monitoring: UP Mode: optimized  Interval:10sec Threshold: 3
    Hard lifetime: Expires in 1313 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 707 seconds
    Mode: Tunnel(10 5), Type: dynamic, State: installed
    Protocol: ESP, Authentication: None, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations family inet6

```
user@host> show security ipsec security-associations family inet6
  Virtual-system: root
  Local Gateway: 2001:db8:1212::1111, Remote Gateway: 2001:db8:1212::1112
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    DF-bit: clear
    Direction: inbound, SPI: 14caf1d9, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3440 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2813 seconds
    Mode: tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64

    Direction: outbound, SPI: 9a4db486, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3440 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2813 seconds
    Mode: tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations fpc 6 pic 1 kmd-instance all (SRX Series Firewalls)

```
user@host> show security ipsec security-associations fpc 6 pic 1 kmd-instance all
  Total active tunnels: 1

ID    Gateway          Port  Algorithm           SPI      Life:sec/kb  Mon vsys

<2    192.168.1.2      500   ESP:aes256/sha256   67a7d25d 28280/unlim   -   0

>2    192.168.1.2      500   ESP:aes256/sha256   a23cbcdc 28280/unlim   -   0
```

show security ipsec security-associations detail (ADVPN Suggester, Static Tunnel)

```
user@host> show security ipsec security-associations detail
ID: 70516737 Virtual-system: root, VPN Name: ZTH_HUB_VPN
  Local Gateway: 192.168.1.1, Remote Gateway: 192.168.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear
  Bind-interface: st0.1

  Port: 500, Nego#: 5, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
  Tunnel events:
  Tue Nov 03 2015 01:24:27 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:24:27 -0800: IKE SA negotiation successfully completed (4 times)
  Tue Nov 03 2015 01:23:38 -0800: User cleared IPSec SA from CLI (1 times)
  Tue Nov 03 2015 01:21:32 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:21:31 -0800: IPSec SA delete payload received from peer, corresponding
IPSec SAs cleared (1 times)
  Tue Nov 03 2015 01:21:27 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:21:13 -0800: Tunnel configuration changed. Corresponding IKE/IPSec SAs are
deleted (1 times)
  Tue Nov 03 2015 01:19:27 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:19:27 -0800: Tunnel is ready. Waiting for trigger event or peer to trigger
negotiation (1 times)
  Location: FPC 0, PIC 3, KMD-Instance 2
  Direction: inbound, SPI: 43de5d65, AUX-SPI: 0
  Hard lifetime: Expires in 1335 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 996 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 3, KMD-Instance 2
  Direction: outbound, SPI: 5b6e157c, AUX-SPI: 0
  Hard lifetime: Expires in 1335 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 996 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
```

```
  Anti-replay service: counter-based enabled

  , Replay window size: 64
```

show security ipsec security-associations detail (ADVPN Partner, Static Tunnel)

```
user@host> show security ipsec security-associations detail
ID: 67108872 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 192.168.1.2, Remote Gateway: 192.168.1.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x8608a29
  Tunnel events:
  Tue Nov 03 2015 01:24:26 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:24:26 -0800: IKE SA negotiation successfully completed (4 times)
  Tue Nov 03 2015 01:23:37 -0800: IPSec SA delete payload received from peer, corresponding
IPSec SAs cleared (1 times)
  Tue Nov 03 2015 01:21:31 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:21:31 -0800: Tunnel is ready. Waiting for trigger event or peer to trigger
negotiation (1 times)
  Tue Nov 03 2015 01:18:26 -0800: Key pair not found for configured local certificate.
Negotiation failed (1 times)
  Tue Nov 03 2015 01:18:13 -0800: CA certificate for configured local certificate not found.
Negotiation not initiated/successful (1 times)
  Direction: inbound, SPI: 5b6e157c, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 43de5d65, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations sa-type shortcut (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut
Total active tunnels: 1
ID         Algorithm         SPI      Life:sec/kb  Mon lsys Port  Gateway
<268173318 ESP:aes256/sha256 6f164ee0 3580/ unlim - root 500 192.168.0.111
>268173318 ESP:aes256/sha256 e6f29cb0 3580/ unlim - root 500 192.168.0.111
```

show security ipsec security-associations sa-type shortcut detail (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut detail
node0:
--------------------------------------------------------------------------

ID: 67108874 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 192.168.1.2, Remote Gateway: 192.168.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Auto Discovery VPN:
    Type: Shortcut, Shortcut Role: Initiator
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 4500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x40608a29
  Tunnel events:
    Tue Nov 03 2015 01:47:26 -0800: IPSec SA negotiation successfully completed (1 times)
    Tue Nov 03 2015 01:47:26 -0800: Tunnel is ready. Waiting for trigger event or peer to
trigger negotiation (1 times)
    Tue Nov 03 2015 01:47:26 -0800: IKE SA negotiation successfully completed (1 times)
  Direction: inbound, SPI: b7a5518, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: b7e0268, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
```

```
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations family inet detail

```
user@host> show security ipsec security-associations family inet detail
ID: 131073 Virtual-system: root, VPN Name: ike-vpn
  Local Gateway: 192.168.1.1, Remote Gateway: 192.168.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
  DF-bit: clear
  , Copy-Outer-DSCP Enabled
  Bind-interface: st0.99

  Port: 500, Nego#: 116, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
  Fri Oct 30 2015 15:47:21 -0700: IPSec SA rekey successfully completed (115 times)
  Fri Oct 30 2015 11:38:35 -0700: IKE SA negotiation successfully completed (12 times)
  Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1 times)
  Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or peer to trigger
negotiation (1 times)
  Mon Oct 26 2015 16:40:56 -0700: External interface's address received. Information updated (1
times)
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: inbound, SPI: 81b9fc17, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1090 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: outbound, SPI: 727f629d, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1090 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
```

```
  Anti-replay service: counter-based enabled


  , Replay window size: 64
```

show security ipsec security-associations detail (SRX4600)

```
user@host> show security ipsec security-associations detail
ID: 131073 Virtual-system: root, VPN Name: ike-vpn
  Local Gateway: 10.62.1.3, Remote Gateway: 10.62.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.0
  Port: 500, Nego#: 25, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
    Fri Jan 12 2007 07:50:10 -0800: IPSec SA rekey successfully completed (23 times)
  Location: FPC 0, PIC 0, KMD-Instance 0
  Anchorship: Thread 6
  Direction: inbound, SPI: 812c9c01, AUX-SPI: 0
    Hard lifetime: Expires in 2224 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1598 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Location: FPC 0, PIC 0, KMD-Instance 0
  Anchorship: Thread 7
  Direction: outbound, SPI: c4de0972, AUX-SPI: 0
    Hard lifetime: Expires in 2224 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1598 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations detail (SRX5400, SRX5600, SRX5800)

A new output field `IKE SA Index` corresponding to every IPsec SA within a tunnel is displayed under each IPsec SA information.

```
user@host> show security ipsec security-associations detail
ID: 500005 Virtual-system: root, VPN Name: 85BX5-OAM
  Local Gateway: 10.217.0.4, Remote Gateway: 10.200.254.118
  Traffic Selector Name: TS_DEFAULT
  Local Identity: ipv4(0.0.0.0-255.255.255.255)
  Remote Identity: ipv4(10.181.235.224-10.181.235.224)
  Version: IKEv2
  PFS group: N/A
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.0, Policy-name: MACRO-IPSEC-POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Location: FPC 7, PIC 1, KMD-Instance 0
  Anchorship: Thread 15
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0xe2eb3838, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 644 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 159 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes128-gcm, Encryption: aes-gcm (128 bits)
    Anti-replay service: disabled
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-responder-only
    IKE SA Index: 22
  Direction: outbound, SPI: 0x4f7c3101, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 644 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 159 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes128-gcm, Encryption: aes-gcm (128 bits)
    Anti-replay service: disabled
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-responder-only
    IKE SA Index: 22
  Direction: inbound, SPI: 0x30b6d66f, AUX-SPI: 0
```

```
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1771 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1391 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes128-gcm, Encryption: aes-gcm (128 bits)
    Anti-replay service: disabled
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-responder-only
          IKE SA Index: 40
Direction: outbound, SPI: 0xd2db4108, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1771 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1391 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes128-gcm, Encryption: aes-gcm (128 bits)
    Anti-replay service: disabled
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-responder-only
          IKE SA Index: 40
```

show security ipsec security-associations ha-link-encryption (SRX5400, SRX5600, SRX5800)

Starting in Junos OS Release 20.4R1, when you configure the high availability (HA) feature, you can use this show command to view only interchassis link tunnel details.

```
user@host> show security ipsec security-associations ha-link-encryption
  Total active tunnels: 1     Total IPsec sas: 91
  ID      Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <495001 ESP:aes-gcm-256/aes256-gcm 0x0047658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x0046c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x0447658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x0446c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x0847658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x0846c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x0c47658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x0c46c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x1047658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x1046c5cd 298/ unlim - root 500 10.23.0.2
```

```
  <495001 ESP:aes-gcm-256/aes256-gcm 0x1447658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x1446c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x1847658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x1846c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x1c47658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x1c46c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x2047658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x2046c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x2447658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x2446c5cd 298/ unlim - root 500 10.23.0.2
...
```

show security ipsec sa detail ha-link-encryption (SRX5400, SRX5600, SRX5800)

Starting in Junos OS Release 20.4R1, when you configure the high availability (HA) feature, you can use this show command to view only interchassis link tunnel details. It displays the multi SAs created for interchassis link encryption tunnel.

```
user@host> show security ipsec sa detail ha-link-encryption
ID: 495001 Virtual-system: root, VPN Name: L3HA_IPSEC_VPN
  Local Gateway: 10.23.0.1, Remote Gateway: 10.23.0.2
  Traffic Selector Name: __L3HA_IPSEC_VPN__multi_node__
  Local Identity: ipv4(180.100.1.1-180.100.1.1)
  Remote Identity: ipv4(180.100.1.2-180.100.1.2)
  Version: IKEv2
  PFS group: DH-Group-24
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.16000, Policy-name: L3HA_IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  HA Link Encryption Mode: Multi-Node
  Location: FPC -, PIC -, KMD-Instance -
  Anchorship: Thread -
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0x00439cf8, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 294 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 219 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes256-gcm, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

```
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    Location: FPC 1, PIC 0, KMD-Instance 0
    Anchorship: Thread 15
    IKE SA Index: 4294966297
  Direction: outbound, SPI: 0x004cfceb, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 294 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 219 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes256-gcm, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    Location: FPC 1, PIC 0, KMD-Instance 0
    Anchorship: Thread 15
    IKE SA Index: 4294966297
  Direction: inbound, SPI: 0x04439cf8, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 294 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 219 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes256-gcm, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    Location: FPC 1, PIC 0, KMD-Instance 0
    Anchorship: Thread 16
    IKE SA Index: 4294966297
  Direction: outbound, SPI: 0x044cfceb, AUX-SPI: 0
                              , VPN Monitoring: -

...
```

In Junos OS Release 22.3R1 and later, when you configure the Chassis Cluster HA control link encryption feature, you can execute the `show security ike sa ha-link-encryption detail`, `show security ipsec sa ha-link-encryption detail`, and `show security ipsec sa ha-link-encryption` commands to view the Chassis cluster control link encryption tunnel details.

show security ike sa ha-link-encryption detail

```
user@host> show security ike sa ha-link-encryption detail
IKE peer 10.2.0.1, Index 4294966274, Gateway Name: IKE_GW_HA_0
  Role: Initiator, State: UP
  Initiator cookie: ae5bcb5540d388a1, Responder cookie: 28bbae629ceb727f
  Exchange type: IKEv2, Authentication method: Pre-shared-keys
  Local gateway interface: em0
  Routing instance: __juniper_private1__
  Local: 10.7.0.2:500, Remote: 10.2.0.1:500
  Lifetime: Expires in 24856 seconds
  Reauth Lifetime: Disabled
  IKE Fragmentation: Enabled, Size: 576
  Remote Access Client Info: Unknown Client
  Peer ike-id: 10.2.0.1
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-2
  Traffic statistics:
   Input  bytes  :               200644
   Output bytes  :               200644
   Input  packets:                 2635
   Output packets:                 2635
   Input  fragmented packets:       0
   Output fragmented packets:       0
  IPSec security associations: 6 created, 3 deleted
  Phase 2 negotiations in progress: 1
  IPSec Tunnel IDs: 495002
    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 10.7.0.2:500, Remote: 10.2.0.1:500
    Local identity: 10.7.0.2
    Remote identity: 10.2.0.1
    Flags: IKE SA is created
 IPsec SA Rekey CREATE_CHILD_SA exchange stats:
   Initiator stats:                                  Responder stats:
    Request Out             : 1                       Request In             : 1
    Response In             : 1                       Response Out           : 1
    No Proposal Chosen In   : 0                       No Proposal Chosen Out : 0
    Invalid KE In           : 0                       Invalid KE Out         : 0
```

```
    TS Unacceptable In      : 0                       TS Unacceptable Out    : 0
    Res DH Compute Key Fail : 0                       Res DH Compute Key Fail: 0
    Res Verify SA Fail      : 0
    Res Verify DH Group Fail: 0
    Res Verify TS Fail      : 0
```

show security ipsec sa ha-link-encryption detail

```
user@host> show security ipsec sa ha-link-encryption detail
ID: 495002 Virtual-system: root, VPN Name: IPSEC_VPN_HA_0
  Local Gateway: 10.7.0.2, Remote Gateway: 10.2.0.1
  Traffic Selector Name: __IPSEC_VPN_HA_0__l2_chassis_clu
  Local Identity: ipv4(10.7.0.2-10.7.0.2)
  Remote Identity: ipv4(10.2.0.1-10.2.0.1)
  TS Type: traffic-selector
  Version: IKEv2
  PFS group: DH-group-24
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.16000, Tunnel MTU: 0, Policy-
name: IPSEC_POL_HA_0
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  HA Link Encryption Mode: L2 Chassis Cluster
  Location: FPC -, PIC -, KMD-Instance -
  Anchorship: Thread -
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0x35fae26b, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3435 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2818 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 4294966274
  Direction: outbound, SPI: 0x0a2b9927, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3435 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2818 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
```

```
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 4294966274
```

show security ipsec sa ha-link-encryption

```
user@host> show security ipsec sa ha-link-encryption
Total active tunnels: 1     Total IPsec sas: 1
  ID      Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <495002 ESP:aes-cbc-256/sha1 0x35fae26b 3484/ unlim - root 500 10.2.0.1
  >495002 ESP:aes-cbc-256/sha1 0x0a2b9927 3484/ unlim - root 500 10.2.0.1
```

show security ipsec security-associations (ChaCha20-Poly1305)

```
user@host> show security ipsec security-associations
  Total active tunnels: 1     Total IPsec sas: 1
  ID      Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <500001 ESP:chacha20-poly1305/chacha20-poly1305 0x8db47f67 86372/ unlim - root 500 30.1.1.10
  >500001 ESP:chacha20-poly1305/chacha20-poly1305 0x6bfe4971 86372/ unlim - root 500 30.1.1.10
```

show security ipsec security-associations detail (ChaCha20-Poly1305)

```
user@host> show security ipsec security-associations detail
ID: 500001 Virtual-system: root, VPN Name: ipsec_vpn
  Local Gateway: 30.1.1.20, Remote Gateway: 30.1.1.10
  Traffic Selector Name: ts1
  Local Identity: ipv4(40.1.1.0-40.1.1.255)
  Remote Identity: ipv4(20.1.1.0-20.1.1.255)
  TS Type: traffic-selector
  Version: IKEv2
  Quantum Secured: No
  PFS group: N/A
  Passive mode tunneling: Disabled
  DF-bit: clear, Copy-Outer-DSCP: Disabled, Bind-interface: st0.0, Policy-name: ipsec_pol
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Tunnel events:
    Sun Sep 10 2023 23:41:25: IPsec SA negotiation succeeds (1 times)
  Location: FPC 0, PIC 0
```

```
  Anchorship: Thread 1
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0x8db47f67, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 86367 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 85755 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: chacha20-poly1305, Encryption: chacha20-poly1305
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Enabled
    tunnel-establishment: establish-tunnels-responder-only-no-rekey
```

show security ipsec security-associations detail (identity-management)

```
user@host> show security ipsec sa detail
ID: 500008 Virtual-system: root, VPN Name: juniper-vpn
  Local Gateway: 5.0.0.254, Remote Gateway: 5.0.0.1
  Traffic Selector Name: TS1
  Local Identity: ipv4(0.0.0.0-255.255.255.255)
  Remote Identity: ipv4(80.1.1.8-80.1.1.8)
  TS Type: traffic-selector
  Version: IKEv1
  Quantum Secured: No
  PFS group: DH-group-2
  Passive mode tunneling: Disabled
  DF-bit: clear, Copy-Outer-DSCP: Disabled, Bind-interface: st0.1, Policy-name: juniper-ipsec-
policy
  Port: 10952, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Tunnel events:
    Fri Jan 12 2024 04:37:27: IPSec SA is deleted because received DEL notification from peer (1
times)
    Mon Apr 05 2021 12:35:32: IPsec SA identity-management tunnel create event sent (1 times)
    Mon Apr 05 2021 12:35:32: IPsec SA identity-management tunnel rekey event sent (3 times)
```

show security ipsec security-associations detail (multi-sa with forwarding class details)

```
user@host> show security ipsec sa detail
ID: 67108933 Virtual-system: root, VPN Name: v2
  Local Gateway: 2.2.2.1, Remote Gateway: 3.3.3.3
  Traffic Selector Name: ts1
  Local Identity: ipv4(1.1.1.0-1.1.1.255)
  Remote Identity: ipv4(4.4.4.0-4.4.4.255)
  Version: IKEv2
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.0
  Port: 500, Nego#: 5353, Fail#: 0, Def-Del#: 0 Flag: 0x2c608b29
  Multi-sa, pending SAs:1

  Direction: inbound, SPI: 1b37340d, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 107 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 15 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    FC-name:voip-data
  Direction: outbound, SPI: 40c545e3, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 107 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 15 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    FC-name:voip-data
```

show security ipsec security-associations detail (SRX Series Firewalls and MX Series Routers)

In Junos OS Release 20.4R2, 21.1R1, and later, you can execute the `show security ipsec security-associations detail` command to view the traffic selector type for a VPN.

```
user@host> show security ipsec security-associations detail
ID: 500024 Virtual-system: root, VPN Name: S2S_VPN2
  Local Gateway: 10.7.0.2, Remote Gateway: 10.2.0.1
```

```
  Traffic Selector Name: ts1
  Local Identity: ipv4(10.20.20.0-10.20.20.255)
  Remote Identity: ipv4(10.10.10.0-10.10.10.255)
  TS Type: traffic-selector
  Version: IKEv2
  PFS group: DH-group-14
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.2, Policy-name: IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Tunnel events:
    Tue Jan 19 2021 04:43:49: IPsec SA negotiation succeeds (1 times)
  Location: FPC 0, PIC 0, KMD-Instance 0
  Anchorship: Thread 1
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0xf8642fae, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1798 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1397 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 17
  Direction: outbound, SPI: 0xb2a26969, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1798 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1397 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 17
ID: 500025 Virtual-system: root, VPN Name: S2S_VPN1
  Local Gateway: 10.7.0.1, Remote Gateway: 10.2.0.1
  Local Identity: ipv4(0.0.0.0-255.255.255.255)
  Remote Identity: ipv4(0.0.0.0-255.255.255.255)
  TS Type: proxy-id
  Version: IKEv2
  PFS group: DH-group-14
```

```
    DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1, Policy-name: IPSEC_POL
    Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
    Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
    Tunnel events:
      Tue Jan 19 2021 04:44:41: IPsec SA negotiation succeeds (1 times)
    Location: FPC 0, PIC 0, KMD-Instance 0
    Anchorship: Thread 1
    Distribution-Profile: default-profile
    Direction: inbound, SPI: 0xe293762a, AUX-SPI: 0
                                , VPN Monitoring: -
      Hard lifetime: Expires in 1755 seconds
      Lifesize Remaining:  Unlimited
      Soft lifetime: Expires in 1339 seconds
      Mode: Tunnel(0 0), Type: dynamic, State: installed
      Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
      Anti-replay service: counter-based enabled, Replay window size: 64
      Extended-Sequence-Number: Disabled
      tunnel-establishment: establish-tunnels-immediately
      IKE SA Index: 18
    Direction: outbound, SPI: 0x7aef9d7f, AUX-SPI: 0
                                , VPN Monitoring: -
      Hard lifetime: Expires in 1755 seconds
      Lifesize Remaining:  Unlimited
      Soft lifetime: Expires in 1339 seconds
      Mode: Tunnel(0 0), Type: dynamic, State: installed
      Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
      Anti-replay service: counter-based enabled, Replay window size: 64
      Extended-Sequence-Number: Disabled
      tunnel-establishment: establish-tunnels-immediately
      IKE SA Index: 18
```

show security ipsec security-associations detail (SRX5400, SRX5600, SRX5800)

Starting in Junos OS Release 21.1R1, you can view the traffic selector details, that includes, local identity, remote identity, protocol, source-port range, destination port range for multiple terms defined for an IPsec SA.

In the earlier Junos Releases, traffic selection for a particular SA is performed using existing IP range defined using IP address or netmask. From Junos OS Release 21.1R1 onwards, additionally traffic is

selected through protocol specified using *protocol_name*. And also, low and high port range specified for source and destination port numbers.

```
user@host> show security ipsec security-associations detail

ID: 500075 Virtual-system: root, VPN Name: pkn-r0-r1-ipsec-vpn-1
Local Gateway: 10.1.1.1, Remote Gateway: 10.1.1.2

Traffic Selector Name: ts1

   Local Identity:
   Protocol         Port                 IP
   17/UDP         100-200             198.51.100.0-198.51.100.255
   6/TCP         250-300             198.51.100.0-198.51.100.255
   Remote Identity:
   Protocol         Port                 IP
   17/UDP         150-200             10.80.0.1-10.80.0.1
   6/TCP         250-300             10.80.1.1-10.80.1.1
Version: IKEv2
DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.0, Policy-name: pkn-r0-r1-ipsec-
policy
Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
Location: FPC 0, PIC 0, KMD-Instance 0
Anchorship: Thread 1
Distribution-Profile: default-profile
Direction: inbound, SPI: ……
Direction: outbound, SPI: ……
```

show security ipsec security-associations srg-id

```
user@host> show security ipsec security-associations srg-id 1

Total active tunnels: 1     Total IPsec sas: 2
  ID      Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <17277217 ESP:aes-cbc-256/sha256 0xc7faee3e 1440/ unlim - root 500 10.112.0.1
  >17277217 ESP:aes-cbc-256/sha256 0x7921d472 1440/ unlim - root 500 10.112.0.1
  <17277217 ESP:aes-cbc-256/sha256 0xf1a01dd4 1498/ unlim - root 500 10.112.0.1
  >17277217 ESP:aes-cbc-256/sha256 0xa0b77273 1498/ unlim - root 500 10.112.0.1
```

show security ipsec security-associations node-local

```
user@host> show security ipsec security-associations node-local
Total active tunnels: 1     Total IPsec sas: 1
  ID      Algorithm              SPI      Life:sec/kb  Mon lsys Port  Gateway
  <500001 ESP:aes-cbc-256/sha256 0x5f2fdf60 3093/ unlim -  root 500   6.0.0.2
  >500001 ESP:aes-cbc-256/sha256 0x293e67e0 3093/ unlim -  root 500   6.0.0.2
```

show security ipsec security-associations node-local detail

```
user@host> show security ipsec security-associations node-local detail
ID: 500003 Virtual-system: root, VPN Name: IPSEC_VPN
  Local Gateway: 4.0.0.1, Remote Gateway: 6.0.0.2
  Local Identity: ipv4(0.0.0.0-255.255.255.255)
  Remote Identity: ipv4(0.0.0.0-255.255.255.255)
  TS Type: proxy-id
  Version: IKEv2
  Quantum Secured: No
  PFS group: DH-group-19
  Passive mode tunneling: Disabled
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1, Policy-name: IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Tunnel events:
    Sun Apr 09 2023 22:22:27: IPsec SA negotiation succeeds (1 times)
  Location: FPC 1, PIC 1, KMD-Instance 0
  Anchorship: Thread 1
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0x8c8c3761, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3564 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2884 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-responder-only
    IKE SA Index: 25
```

```
Direction: outbound, SPI: 0x0e798f8c, AUX-SPI: 0
                            , VPN Monitoring: -
  Hard lifetime: Expires in 3564 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 2884 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
  Extended-Sequence-Number: Disabled
  tunnel-establishment: establish-tunnels-responder-only
  IKE SA Index: 25
```

Release Information

Command introduced in Junos OS Release 8.5. Support for the `family` option added in Junos OS Release 11.1.

Support for the `vpn-name` option added in Junos OS Release 11.4R3. Support for the `traffic-selector` option and traffic selector field added in Junos OS Release 12.1X46-D10.

Support for Auto Discovery VPN (ADVPN) added in Junos OS Release 12.3X48-D10.

Support for IPsec datapath verification added in Junos OS Release 15.1X49-D70.

Support for thread anchorship added in Junos OS Release 17.4R1.

Starting in Junos OS Release 18.2R2 the `show security ipsec security-assocations detail` command output will include thread anchorship information for the security associations (SAs).

Starting in Junos OS Release 19.4R1, we have deprecated the CLI option `fc-name` (COS Forward Class name) in the new **iked** process that displays the security associations (SAs) under show command `show security ipsec sa`.

Support for the `ha-link-encryption` option added in Junos OS Release 20.4R1.

Support for the `srg-id` option added in Junos OS Release 22.4R1.

Support for the `passive-mode-tunneling` option added in Junos OS Release 23.1R1.

Support for the `node-local` option is added in Junos OS Release 23.2R1.

Starting in Junos OS Release 23.4R1, `kmd-instance` option is available only when you have `kmd` process for IPsec VPN. When you enable `iked` process using`junos-iked` package, this option is not available.

Support for lifesize in kilobytes, lifesize remaining, and VPN monitoring information in the command output with IPsec VPN running the iked process is added in Junos OS Release 23.4R1.

Support for the `chacha20-poly1305` option is added in Junos OS Release 24.2R1.

Support for `fc-name` in the displayed output added in Junos OS Release 24.4R1.

Support for `identity-management` in the displayed output added in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

Example: Configuring a Route-Based VPN Tunnel in a User Logical Systems

show security ipsec statistics

IN THIS SECTION

- Syntax | 24744
- Description | 24745
- Options | 24745
- Required Privilege Level | 24745
- Output Fields | 24745
- Sample Output | 24747
- Release Information | 24752

Syntax

```
show security ipsec statistics
<fpc slot-number>
<index SA-index-number>
<pic slot-number>
<srg-id id-number>
<ha-link-encryption>
```

Description

Display standard IPsec statistics.

Options

- none—Display statistics about all IPsec security associations (SAs).
- `fpc` *`slot-number`* —Specific to SRX Series Firewalls. Display statistics about existing IPsec SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.
- `index` *`SA-index-number`* —(Optional) Display statistics for the SA with this index number.
- `srg-id` *`id-number`* —(Optional) Display information related to a specific services redundancy group (SRG) in a Multinode High Availability setup.
- `pic` *`slot-number`* —Specific to SRX Series Firewalls. Display statistics about existing IPsec SAs in this PIC slot. This option is used to filter the output.
- `ha-link-encryption`—(Optional) Display information related to interchassis link tunnel only. See "ipsec (High Availability)" on page 9301 and "show security ipsec statistics ha-link-encryption (SRX5400, SRX5600, SRX5800)" on page 24749.

Required Privilege Level

view

Output Fields

Table 1269 on page 24745 lists the output fields for the `show security ipsec statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1269: show security ipsec statistics Output Fields

Field Name	Field Description
`Virtual-system`	The root system.

Table 1269: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
`ESP Statistics`	• `Encrypted bytes`—Total number of bytes encrypted by the local system across the IPsec tunnel. • `Decrypted bytes`—Total number of bytes decrypted by the local system across the IPsec tunnel. • `Encrypted packets`—Total number of packets encrypted by the local system across the IPsec tunnel. • `Decrypted packets`—Total number of packets decrypted by the local system across the IPsec tunnel.
`AH Statistics`	• `Input bytes`—Total number of bytes received by the local system across the IPsec tunnel. • `Output bytes`—Total number of bytes transmitted by the local system across the IPsec tunnel. • `Input packets`—Total number of packets received by the local system across the IPsec tunnel. • `Output packets`—Total number of packets transmitted by the local system across the IPsec tunnel.

Table 1269: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
`Errors`	• `AH authentication failures`—Total number of authentication header (AH) failures. An AH failure occurs when there is a mismatch of the authentication header in a packet transmitted across an IPsec tunnel. • `Replay errors`—Total number of replay errors. A replay error is generated when a duplicate packet is received within the replay window. • `ESP authentication failures`—Total number of Encapsulation Security Payload (ESP) failures. An ESP failure occurs when there is an authentication mismatch in ESP packets. • `ESP decryption failures`—total number of ESP decryption errors. • `Bad headers`—Total number of invalid headers detected. • `Bad trailers`—Total number of invalid trailers detected. • `Invalid SPI`— Total number of invalid SPIs packets detected. • `TS check fail`— Total number of TS check fail detected. • `Discarded`— Total number of discarded packets detected.
Multi-sa tunnel statistics	• `FC-name`—Forwarding class name for the child security association. • `Encrypted pkts`—Total number of encrypted packets for a forwarding class. • `Encrypted bytes`—Total encrypted bytes for a forwarding class. • `Decrypted pkts`—Total number of encrypted packets for a forwarding class. • `Decrypted bytes`—Total encrypted bytes for a forwarding class.

Sample Output

show security ipsec statistics

```
user@host> show security ipsec statistics
Virtual-system: Root
```

```
ESP Statistics:
  Encrypted bytes:                0
  Decrypted bytes:                0
  Encrypted packets:              0
  Decrypted packets:              0
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
  Invalid SPI: 0, TS check fail: 0
  Discarded: 0
```

show security ipsec statistics index 131073

```
user@host> show security ipsec statistics index 131073
ESP Statistics:
  Encrypted bytes:              952
  Decrypted bytes:              588
  Encrypted packets:              7
  Decrypted packets:              7
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
  Invalid SPI: 0, TS check fail: 0
  Discarded: 0

  FC Name     Encrypted Pkts  Decrypted Pkts  Encrypted bytes  Decrypted bytes
  best-effort 7                 7                 952                588
  custom_q1   0                 0                 0                  0
  custom_q2   0                 0                 0                  0
```

```
  network-control 0             0                 0                 0
  custom_q4   0                 0                 0                 0
  custom_q5   0                 0                 0                 0
  custom_q6   0                 0                 0                 0
  custom_q7   0                 0                 0                 0
  default     0                 0                 0                 0
```

Starting with Junos OS Release 18.2R1, the CLI `show security ipsec statistics index 131073` *`index-number`* output displays statistics for each forwarding class name.

show security ipsec statistics fpc 6 pic 1 (SRX Series Firewalls)

```
user@host> show security ipsec statistics fpc 6 pic 1
ESP Statistics:
Encrypted bytes:           536408
Decrypted bytes:           696696
Encrypted packets:           1246
Decrypted packets:            888
AH Statistics:
Input bytes:                    0
Output bytes:                   0
Input packets:                  0
Output packets:                 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0
Invalid SPI: 0, TS check fail: 0
Discarded: 0
```

show security ipsec statistics ha-link-encryption (SRX5400, SRX5600, SRX5800)

Starting in Junos OS Release 20.4R1, when you configure the high availability (HA) feature, you can use this show command to view only interchassis link tunnel details. The following command displays only link encryption tunnel statistics on both nodes.

```
user@host> show security ipsec statistics ha-link-encryption
ESP Statistics:
  Encrypted bytes:            10376
  Decrypted bytes:             4996
```

```
  Encrypted packets:             96
  Decrypted packets:             96
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
  Invalid SPI: 0, TS check fail: 0
  Discarded: 0
```

show security ipsec statistics (MX-SPC3)

Starting with Junos OS Release 21.3R1, a new field **Tunnel MTU** in the output of the CLI `show security ipsec statistics` displays the option configured under `ipsec vpn hub-to-spoke-vpn tunnel-mtu` hierarchy.

```
user@host> show security ipsec statistics
  Encrypted bytes:                0
  Decrypted bytes:                0
  Encrypted packets:              0
  Decrypted packets:              0
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
  Invalid SPI: 0, TS check fail: 0
  Exceeds tunnel MTU: 0       --------  New counter
  Discarded: 0
```

show security ipsec statistics srg-id <srg-id>

```
user@host> show security ipsec statistics srg-id 1
ESP Statistics:
  Encrypted bytes:            10646
  Decrypted bytes:             4296
  Encrypted packets:             96
  Decrypted packets:             96
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
  Invalid SPI: 0, TS check fail: 0
  Exceeds tunnel MTU: 0
  Discarded: 0
```

show security ipsec statistics (MX304)

```
user@host> show security ipsec statistics
 ESP Statistics:
Encrypted bytes: 724710
Decrypted bytes: 950040
Encrypted packets: 3045
Decrypted packets: 3045
AH Statistics:
Input bytes: 0
Output bytes: 0
Input packets: 0
Output packets: 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0
```

show security ipsec statistics index <index-number> (forwarding class details)

```
user@host> show security ipsec statistics index 67108865
ESP Statistics:
  Encrypted bytes:                0
  Decrypted bytes:                0
  Encrypted packets:              0
  Decrypted packets:              0
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0

FC-name    Encrypted pkts       Encrypted bytes    Decrypted pkts    Decrypted bytes
default        10           840                      10        840
voip-data    100           8400                          100        8400
web-data    200           16800                                200        16800
```

Release Information

Command introduced in Junos OS Release 8.5. `fpc` and `pic` options added in Junos OS Release 9.3.

Support for the `ha-link-encryption` option added in Junos OS Release 20.4R1.

Support for the `srg-id` option added in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

clear security ipsec statistics | 3995

show security ipsec traffic-selector

IN THIS SECTION

- Syntax | 24753
- Description | 24753
- Options | 24753
- Required Privilege Level | 24754
- Output Fields | 24754
- Sample Output | 24755
- Release Information | 24755

Syntax

```
show security ipsec traffic-selector interface-name interface-name
<brief | detail>
<destination-address address>
<fpc slot-number pic slot-number>
<kmd-instance (all | kmd-instance-name)>
<pic slot-number fpc slot-number>
<source-address address>
```

Description

Display information about the traffic selectors that have been negotiated between the initiator and responder.

Options

interface-name *interface-name*	Name of the secure tunnel logical interface.
brief | detail	(Optional) Display the specified level of output. The default is `brief`.
destination-address *address*	(Optional) Destination IP address.

fpc *slot-number* **pic** *slot-number*	(Optional) Display information about existing traffic selectors on the specified Flexible PIC Concentrator (FPC) slot and PIC slot.
kmd-instance	(Optional) Display information about existing traffic selectors in the key management process (in this case, it is KMD) identified by FPC slot-number and PIC slot-number. This option is used to filter the output. • `all`—All KMD instances running on the Services Processing Unit (SPU). • *`kmd-instance-name`*—Name of the KMD instance running on the SPU.
pic *slot-number* **fpc** *slot-number*	(Optional) Display information about existing traffic selectors on the specified PIC slot and FPC slot.
source-address *address*	(Optional) Source IP address.

Required Privilege Level

view

Output Fields

Table 1270 on page 24754 lists the output fields for the `show security ipsec traffic-selector` command. Output fields are listed in the approximate order in which they appear.

Table 1270: show security ipsec traffic-selector Output Fields

Field Name	Field Description	Level of Output
Tunnel-id	Tunnel ID.	All levels
Interface	Secure tunnel (st0) interface for the traffic selector.	All levels
IKE-ID	Peer IKE ID for the negotiated traffic selector.	All levels
Source IP	Source IP address for the negotiated traffic selector.	All levels
Destination IP	Destination IP address for the negotiated traffic selector.	All levels

Sample Output

show security ipsec traffic-selector interface-name st0.1 detail

```
user@host> show security ipsec traffic-selector interface-name st0.1 detail
Tunnel ID: 6920601, Interface: st0.1
IKE-ID: DC=Common_component, CN=enodeA, OU=Dept, O=Company, L=City, ST=CA, C=US
Source IP: ipv4 (192.0.2.0-192.0.2.255)
Destination IP: ipv4 (198.51.100.0-198.51.100.255)

Tunnel ID: 77594626, Interface: st0.1
IKE-ID: DC=Common_component, CN=enodeB, OU=Det, O=Company, L=City, ST=CA, C=US
Source IP: ipv4 (192.0.2.0-192.0.2.255)
Destination IP: ipv4 (203.0.113.0-203.0.113.255)
```

Release Information

Command introduced in Junos OS Release 12.3X48-D10.

RELATED DOCUMENTATION

IPsec Overview

show security ipsec tunnel-distribution

IN THIS SECTION

- Syntax | 24756
- Description | 24756
- Options | 24756
- Required Privilege Level | 24756
- Output Fields | 24757
- Sample Output | 24757
- Release Information | 24761

Syntax

```
show security ipsec tunnel-distribution
<brief | summary | summary-cpuload>
<srg-id number>
```

Description

Display the number of IPsec VPN tunnels that are anchored in each thread. An IPsec tunnel session is assigned an anchor thread, based on the load during the tunnel session installation. When a new tunnel session is created, the least loaded thread is chosen to anchor the new tunnel. When the tunnel is deleted, the anchor mapping is removed from the control plane.

Tunnel distribution across different Services Processing Unit (SPU) or equivalent is based on the number of tunnels and not on throughput in each tunnel. Tunnels anchored in a SPU are not transferred to a different SPU or equivalent during SPU failure.

The distribution profile shows any assigned IPSec distribution profile without any distribution profiles assigned to a vpn object. This tab shows `default_profiile`, else the associated profile is displayed.

Options

none	Display thread information about all active tunnels.
brief	(Optional) Display thread information about all active tunnels. (Default)
fpc	FPC slot number (0..5).
pic	PIC slot number (0..3).
summary	(Optional) Display the number of tunnels anchored to each thread.
summary-cpuload	(Optional) Displays the load on each FPC and PIC. You can use this option to check the load on each FPC and PIC before or after redistributing the tunnel. See "show security ipsec tunnel-distribution summary-cpuload" on page 24760.
srg-id	(Optional) Display information related to a specific services redundancy group (SRG) in a Multinode High Availability setup.

Required Privilege Level

view

Output Fields

Table 1271 on page 24757 lists the output fields for the `show security ipsec tunnel-distribution` command. Output fields are listed in the approximate order in which they appear.

Table 1271: show security ipsec tunnel-distribution Output Fields

Field Name	Field Description	Level of Output
Tunnel-ID	VPN tunnel identifier.	`brief`
Thread-ID	Thread identifier.	All levels
Number of Tunnels	The number of tunnels anchored to the thread.	`summary`
CPU:1m	CPU load average for last 1 minute for FPC or PIC.	`summary-cpuload`
CPU:1h	CPU load average for last 1 hour for FPC or PIC.	`summary-cpuload`
CPU:1d	CPU load average for last 1 day for FPC or PIC.	`summary-cpuload`

Sample Output

show security ipsec tunnel-distribution

```
user@host> show security ipsec tunnel-distribution
Tunnel-ID          FPC        PIC       Thread-ID
-------------------------------------------------------------------
  500006              0        1        4
  500012              0        1        8
  500009              0        1        6
  500002              0        1        1
  500005              0        1        3
  500001              0        0        15
  500008              0        1        5
  500010              0        0        18
  500004              0        0        16
```

```
  500003              0        1        2
  500011              0        1        7
  500007              0        0        17


Tunnel-ID                      FPC         PIC     Thread-ID  Distribution-profile
-----------------------------------------------------------------
  500755                  0         1         1       spc-3
  500756                  2         0         0       spc-2
  500758                  0         1         1       default_profile
```

show security ipsec tunnel-distribution summary

```
user@host> show security ipsec tunnel-distribution summary
Number of Tunnels       FPC       PIC       Thread-ID
-----------------------------------------------------------------
  1                       0         0         15
  1                       0         0         16
  1                       0         0         17
  1                       0         0         18
  1                       0         1         1
  1                       0         1         2
  1                       0         1         3
  1                       0         1         4
  1                       0         1         5
  1                       0         1         6
  1                       0         1         7
  1                       0         1         8
```

show security ipsec tunnel-distribution fpc 0 pic 0

```
user@host> show security ipsec tunnel-distribution fpc 0 pic 0
Tunnel-ID          FPC        PIC       Thread-ID
-----------------------------------------------------------------
  500001              0        0        15
  500010              0        0        18
  500004              0        0        16
  500007              0        0        17
```

show security ipsec tunnel-distribution fpc 0 pic 1

```
user@host> show security ipsec tunnel-distribution fpc 0 pic 1
Tunnel-ID          FPC        PIC       Thread-ID
-----------------------------------------------------------------
  500006              0        1        4
  500012              0        1        8
  500009              0        1        6
  500002              0        1        1
  500005              0        1        3
  500008              0        1        5
  500003              0        1        2
  500011              0        1        7
```

show security ipsec tunnel-distribution summary fpc 0 pic 0

```
user@host> show security ipsec tunnel-distribution summary fpc 0 pic 0
Number of Tunnels       FPC       PIC       Thread-ID
-----------------------------------------------------------------
  1                       0         0         15
  1                       0         0         16
  1                       0         0         17
  1                       0         0         18
  0                       0         0         19
  0                       0         0         20
  0                       0         0         21
  0                       0         0         22
  0                       0         0         23
  0                       0         0         24
  0                       0         0         25
  0                       0         0         26
  0                       0         0         27
```

show security ipsec tunnel-distribution summary fpc 0 pic 1

```
user@host> show security ipsec tunnel-distribution summary fpc 0 pic 1
Number of Tunnels       FPC       PIC       Thread-ID
-----------------------------------------------------------------
  1                       0         1         1
```

```
  1                       0         1         2
  1                       0         1         3
  1                       0         1         4
  1                       0         1         5
  1                       0         1         6
  1                       0         1         7
  1                       0         1         8
  0                       0         1         9
  0                       0         1         10
  0                       0         1         11
  0                       0         1         12
  0                       0         1         13
  0                       0         1         15
  0                       0         1         16
  0                       0         1         17
  0                       0         1         18
  0                       0         1         19
  0                       0         1         20
  0                       0         1         21
  0                       0         1         22
  0                       0         1         23
  0                       0         1         24
  0                       0         1         25
  0                       0         1         26
  0                       0         1         27
```

show security ipsec tunnel-distribution summary-cpuload

This command displays the same output as `show security ipsec tunnel-distribution summary`, but includes load averages (last 1 minute, 1 hour, and 1 day) of all threads for each FPC and PIC.

```
user@host> show security ipsec tunnel-distribution summary-cpuload

node0:

--------------------------------------------------------------------------------------------------
----------
   Number of Tunnels  FPC    PIC  Thread-ID  CPU:1m   CPU:1h  CPU:1d

--------------------------------------------------------------------------------------------------
----------
```

```
    1                 0       0     15         0        0       0
    1                 0       0     16         0        0       0
```

show security ipsec tunnel-distribution srg-id

```
user@host> show security ipsec tunnel-distribution srg-id 1
Tunnel-ID          FPC        PIC       Thread-ID Distribution Profile
-----------------------------------------------------------------
  17277221            0        0        1        default-profile
```

Release Information

Command introduced in Junos OS Release 17.4R1.

summary-cpuload option introduced in Junos OS Release 20.4R1.

srg-id option introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

show security ipsec security-associations | **24701**

request security re-distribution ipsec-vpn | **14863**

show security re-distribution ipsec-vpn | **25093**

show security ipsec tunnel-events-statistics

IN THIS SECTION

- Syntax | **24762**
- Description | **24762**
- Required Privilege Level | **24762**
- Sample Output | **24762**
- Release Information | **24763**

Syntax

```
show security ipsec tunnel-events-statistics
```

Description

Show tunnel event statistics.

Required Privilege Level

view

Sample Output

show security ipsec tunnel-events statistics

```
user@host> show security ipsec tunnel-events statistics
IPSec SA delete payload received from peer                                : 153
Configuration change triggered clearing of IPSec SA                       : 1
Peer's remote IKE-ID validation failed during negotiation                 : 2
Phase1 proposal mismatch detected                                         : 2
Phase2 proposal mismatch detected                                         : 2
Peer proposed traffic-selectors are not in configured range               : 8576
Negotiation failed as peer did not respond                                : 4
IKE SA negotiation successfully completed                                 : 19
IPSec SA negotiation successfully completed                               : 154
PKI validation failed: Peer's CA not configured in trusted-CA-group in IKE policy : 1
Tunnel is ready. Waiting for trigger event or peer to trigger negotiation  : 1
```

show security ipsec tunnel-events statistics (identity-management)

```
user@host> show security ipsec tunnel-events statistics
IPsec SA negotiation
succeeds                                                                  : 1
IPsec SA rekey
succeeds                                                                       : 11
IPsec SA identity-management tunnel create event
sent                                             : 1
```

```
IPsec SA identity-management tunnel rekey event
sent                                            : 11
IPsec SA identity-management tunnel delete event
sent                                           : 1
```

Release Information

Command introduced in Junos OS Release 12.3X48-D10.

Starting with Junos OS Release 15.1X49-D120, you can configure the CLI option `reject-duplicate-connection` at the [`edit security ike gateway` *`gateway-name`* `dynamic`] hierarchy level to retain an existing tunnel session and reject negotiation requests for a new tunnel with the same IKE ID. By default, an existing tunnel is tear down when a new tunnel with the same IKE ID is established. The `reject-duplicate-connection` option is only supported when `ike-user-type group-ike-id` or `ike-user-type shared-ike-id` is configured for the IKE gateway; the `aaa access-profile` *`profile-name`* configuration is not supported with this option.

Use the CLI option `reject-duplicate-connection` only when you are certain that reestablishment of a new tunnel with the same IKE ID should be rejected.

Support for `identity-management` statistics added in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

clear security ipsec tunnel-events-statistics | 3997

show security ipsec vpn vpnname ike idle time

IN THIS SECTION

- Syntax | 24764
- Description | 24764
- Options | 24764
- Additional Information | 24764
- Required Privilege Level | 24764
- Sample Output | 24764

Release Information | 24764

Syntax

```
show security ipsec vpn vpn-nameike idle-time
```

Description

Displays the idle time to delete SAs.

Options

seconds The range of values, in seconds, is 0 through 999,999.

Additional Information

Required Privilege Level

view

Sample Output

show security ipsec vpn *vpn-name* ike idle-time

```
user@host> show security ipsec vpn vpn-name ike idle-time
```

Release Information

Command introduced in Junos OS Release 22.3R1.

show security keychain

IN THIS SECTION

- Syntax | 24765
- Description | 24765
- Options | 24765
- Required Privilege Level | 24765
- Output Fields | 24766
- Sample Output | 24768
- Release Information | 24769

Syntax

```
show security keychain
<brief | detail>
```

Description

Display information about authentication keychains configured for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol.

Options

none	Display information about authentication keychains.
brief | detail	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1272 on page 24766 describes the output fields for the `show security keychain` command. Output fields are listed in the approximate order in which they appear.

Table 1272: show security keychain Output Fields

Field Name	Field Description	Level of Output
keychain	The name of the keychain in operation.	All levels
Active-ID Send	Number of routing protocols packets sent with the active key.	All levels
Active-ID Receive	Number of routing protocols packets received with the active key.	All levels
Next-ID Send	Number of routing protocols packets sent with the next key.	All levels
Next-ID Receive	Number of routing protocols packets received with the next key.	All levels
Transition	Amount of time until the current key will be replaced with the next key in the keychain.	All levels
Tolerance	Configured clock-skew tolerance, in seconds, for accepting keys for a key chain.	All levels
Id	Identification number configured for the current key.	`detail`
Algorithm	Authentication algorithm configured for the current key.	`detail`

Table 1272: show security keychain Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the current key. The value can be: • `receive` • `send` • `send-receive` For the active key, the `State` can be `send-receive`, `send`, or `receive`. For keys that have a future start time, the `State` is `inactive`. Compare the `State` field to the `Mode` field.	`detail`
Option	For IS-IS only, the option determines how Junos OS encodes the message authentication code in routing protocol packets. The values can be: • `basic`—Based on RFC 5304. • `isis-enhanced`—Based on RFC 5310. The default value is `basic`. When you configure the `isis-enhanced` option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices. When you configure `basic` (or do not include the `options` statement in the key configuration) Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices. Because this setting is for IS-IS only, the TCP and the BFD protocol ignore the encoding option configured in the key.	`detail`
Start-time	Time that the current key became active.	`detail`

Table 1272: show security keychain Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Mode	Mode of each key (Informational only.) The value can be • `receive` • `send` • `send-receive` The mode of the key is based on the configuration. Suppose you configure two keys, one with a start-time of today and the other with a start-time of next week. For both keys, the `Mode` can be `send-receive`, `send`, or `receive`, regardless of the configured start-time. Compare the `Mode` field to the `State` field.	`detail`

Sample Output

show security keychain brief

```
user@host> show security keychain brief
keychain                 Active-ID       Next-ID       Transition  Tolerance
                       Send  Receive   Send  Receive
 hakr                    3     3        1     1        1d 23:58    3600
```

show security keychain detail

```
user@host> show security keychain detail
keychain                 Active-ID       Next-ID       Transition  Tolerance
                       Send  Receive   Send  Receive
 hakr                    3     3        1     1        1d 23:58    3600
  Id 3, Algorithm hmac-md5, State send-receive, Option basic
  Start-time Wed Aug 11 16:28:00 2010, Mode send-receive
  Id 1, Algorithm hmac-md5, State inactive, Option basic
  Start-time Fri Aug 20 11:30:57 2010, Mode send-receive
```

Release Information

Command introduced in Junos OS Release 11.2.

show security log

IN THIS SECTION

- Syntax | 24769
- Description | 24769
- Options | 24769
- Required Privilege Level | 24770
- Output Fields | 24771
- Sample Output | 24771
- Release Information | 24773

Syntax

```
show security log {all| destination-address| destination-port| event-id| failure|interface-name|
newer-than| older-than| process| protocol|report| severity| sort-by| source-address| source-
port| success| user}
```

Description

Display security event logs. This command continuously displays security events on the screen. To stop the display, press Ctrl+c.

Options

all	Display all audit event logs stored in the device memory.
destination-address	Display audit event logs with the specified destination address.

destination-port	Display audit event logs with the specified destination port.
event-id	Display audit event logs with the specified event identification number.
failure	Display failed audit event logs.
file	Display security logs in binary format.
interface-name	Display audit event logs with the specified interface.
newer-than	Display audit event logs newer than the specified date and time.
older-than	Display audit event logs older than the specified date and time.
process	Display audit event logs with the specified process that generated the event.
profile	Display security log profile information.
protocol	Display audit event logs generated through the specified protocol.
query	Query logs from database.
report	Display on-box reports for system traffic logs.
severity	Display audit event logs generated with the specified severity.
sort-by	Display audit event logs generated sorted with the specified options.
source-address	Display audit event logs with the specified source address.
source-port	Display audit event logs with the specified source port.
stream	Display security logs stream information.
success	Display successful audit event logs.
username	Display audit event logs generated for the specified user.

Required Privilege Level

view

Output Fields

Table 1273 on page 24771 lists the output fields for the `show security log` command. Output fields are listed in the approximate order in which they appear.

Table 1273: show security log Output Fields

Field Name	Field Description
`Event time`	The timestamp of the events received. Security logs were always timestamped using the UTC time zone by running `set system time-zone utc` and `set security log utc-timestamp` CLI commands. Now, time zone can be defined using the local time zone by running the `set system time-zone` **`time-zone`** command to specify the local time zone that the system should use when timestamping the security logs.
`Message`	Security events are listed.

Sample Output

show security log

```
user@host> show security log
Event time               Message
2010-10-22 13:28:37 CST  session created 1.1.1.2/1-->2.2.2.2/1308
 icmp 1.1.1.2/1-->2.2.2.2/1308
 None None 1 policy1 trustZone untrustZone 52 N/A(N/A) ge-0/0/1.0
2010-10-22 13:28:38 CST  session created 1.1.1.2/1-->2.2.2.2/1308 icmp 1.1.1.2/1-->2.2.2.2/1308
None None 1 policy1 trustZone untrustZone 54 N/A(N/A) ge-0/0/1.0

...
2010-10-22 13:36:12 CST  session denied m icmp 1(8) policy1 trustZone untrustZone N/A(N/A)
ge-0/0/1.0
2010-10-22 13:36:14 CST  session denied 1.1.1.2/2-->2.2.2.2/54812  icmp 1(8) policy1 trustZone
untrustZone N/A(N/A) ge-0/0/1.0

...
2010-10-27 15:50:11 CST  IP spoofing! source: 2.2.2.20, destination: 2.2.2.2, protocol-id: 17,
zone name: trustZone, interface name: ge-0/0/1.0, action: drop
```

```
2010-10-27 15:50:11 CST  IP spoofing! source: source: 2.2.2.20, destination: 2.2.2.2, protocol-
id: 17, zone name: trustZone, interface name: ge-0/0/1.0, action: drop

...
2011-02-18 15:53:34 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory from /var/db/
certs/common/certification-authority/ca-profile1-ca1.cert
2011-02-18 15:53:35 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory from /var/db/
certs/common/crl/ca-profile1.crl
2011-02-18 15:53:35 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory from /var/db/
certs/system-key-pair/system-generated.priv
2011-02-18 15:53:35 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory from /var/db/
certs/system-cert/system-generated.cert
2011-02-18 15:53:35 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory from /var/db/
certs/common/key-pair/cert1.priv
2011-02-18 15:53:42 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory from /var/db/
certs/common/key-pair/test2.priv

...
2011-03-14 23:00:40 PDT  IDP_COMMIT_COMPLETED: IDP policy commit is complete.
                         IDP_POLICY_LOAD_FAILED: IDP policy loading failed ;poli
cy[/var/db/idpd/bins/.bin.gz.v], detector[/usr/libdata/libidp-detector.so.tgz.v]
 ,failure detail[Policy loading failed :: Policy file not found
2011-03-14 23:00:58 PDT  ]
                         IDP_POLICY_LOAD_FAILED: IDP policy loading failed ;poli
cy[/var/db/idpd/bins/.bin.gz.v], detector[/usr/libdata/libidp-detector.so.tgz.v]
 ,failure detail[Policy loading failed :: Policy file not found
2011-03-14 23:00:58 PDT  ]
                         IDP_POLICY_LOAD_FAILED: IDP policy loading failed ;poli
cy[/var/db/idpd/bins/.bin.gz.v], detector[/usr/libdata/libidp-detector.so.tgz.v]
 ,failure detail[Policy loading failed :: Policy file not found
2011-03-14 23:00:58 PDT  ]

...
Event time               Message
2011-03-21 14:21:49 CST  UI_CMDLINE_READ_LINE: User 'root', command 'set date ntp 9.9.9.1 source-
address 6.6.6.1 '
2011-03-21 14:23:01 CST  UI_CMDLINE_READ_LINE: User 'root', command 'set date ntp 9.9.9.1 source-
address 6.6.6.1 .5 '
2011-03-21 14:23:05 CST  KMD_PM_SA_ESTABLISHED: Local gateway: 7.7.7.1, Remote gateway: 8.8.8.1,
Local ID: ipv4(any:0,[0..3]=6.6.6.1), Remote ID: ipv4(any:0,[0..3]=9.9.9.1), Direction: inbound,
SPI: 37a2a179, AUX-SPI: 0, Mode: tunnel, Type: dynamic
2011-03-21 14:23:05 CST  KMD_PM_SA_ESTABLISHED: Local gateway: 7.7.7.1, Remote gateway: 8.8.8.1,
Local ID: ipv4(any:0,[0..3]=6.6.6.1), Remote ID: ipv4(any:0,[0..3]=9.9.9.1), Direction:
```

```
outbound, SPI: b2231c1f, AUX-SPI: 0, Mode: tunnel, Type: dynamic
2011-03-21 14:23:08 CST  UI_CMDLINE_READ_LINE: User 'root', command 'set date ntp 9.9.9.1 source-
address 6.6.6.1 '
2011-03-21 14:23:13 CST  UI_CMDLINE_READ_LINE: User 'root', command 'show security log '
```

Release Information

Command introduced in Junos OS Release 11.2 .

RELATED DOCUMENTATION

exclude (Security Log) | 6627

clear security log | 3998

show security log file

IN THIS SECTION

- Syntax | 24773
- Description | 24773
- Required Privilege Level | 24774
- Output Fields | 24774
- Sample Output | 24774
- Release Information | 24781

Syntax

```
show security log file
```

Description

Enables you to view event-mode log files stored on the device in binary and protobuf format.

Required Privilege Level

view

Output Fields

Table 1274 on page 24774 lists the output fields for the `show security log file` command. Output fields are listed in the approximate order in which they appear.

Table 1274: show security log file Output Fields

Field Name	Field Description
`Event time`	The timestamp when the security event was received.
`Message`	The message describing the security event.

Sample Output

show security log file

```
user@host> show security log file

<14>1 2011-08-28T21:14:43 topstar RT_FLOW - RT_FLOW_SESSION_CREATE [junos@2636.1.1.1.2.34 source-
address="7.7.7.2" source-port="1" destination-address="8.8.8.2" destination-port="5636" service-
name="icmp" nat-source-address="7.7.7.2" nat-source-port="1" nat-destination-address="8.8.8.2"
nat-destination-port="5636" src-nat-rule-name="None" dst-nat-rule-name="None" protocol-id="1"
policy-name="client-to-server" source-zone-name="client" destination-zone-name="server" session-
id-32="60000442" username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/0.0"]

<14>1 2011-08-28T21:14:45 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE [junos@2636.1.1.1.2.34
reason="response received" source-address="7.7.7.2" source-port="0" destination-
address="8.8.8.2" destination-port="5636" service-name="icmp" nat-source-address="7.7.7.2" nat-
source-port="0" nat-destination-address="8.8.8.2" nat-destination-port="5636" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="1" policy-name="client-to-server" source-zone-
name="client" destination-zone-name="server" session-id-32="60000441" packets-from-client="1"
bytes-from-client="84" packets-from-server="1" bytes-from-server="84" elapsed-time="3"
application="UNKNOWN" nested-application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-
interface="ge-0/0/0.0"]
```

```
...

user@host> show security log file

<14>1 2011-11-17T23:41:46 topstar RT_FLOW - RT_FLOW_SESSION_CREATE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE [junos@2636.1.1.1.2.34
reason="response received" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
packets-from-client="1" bytes-from-client="104" packets-from-server="1" bytes-from-server="104"
elapsed-time="3" application="UNKNOWN" nested-application="UNKNOWN" username="N/A" roles="N/A"
packet-incoming-interface="ge-0/0/2.0" encrypted="No "]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" reason="response received" source-address="3001::2" source-port="0"
destination-address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-
address="3001::2" nat-source-port="0" nat-destination-address="5001::2" nat-destination-
port="17420" src-nat-rule-name="None" dst-nat-rule-name="None" protocol-id="58" policy-
name="lsys1trust-to-lsys1trust" source-zone-name="lsys1-trust" destination-zone-name="lsys1-
trust" session-id-32="60000218" packets-from-client="1" bytes-from-client="104" packets-from-
server="1" bytes-from-server="104" elapsed-time="3" application="UNKNOWN" nested-
application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"
encrypted="No "]

 ...

user@host>show security log file bin_msg

<14>1 2011-11-17T23:41:46 topstar RT_FLOW - RT_FLOW_SESSION_CREATE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
```

```
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE [junos@2636.1.1.1.2.34
reason="response received" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
packets-from-client="1" bytes-from-client="104" packets-from-server="1" bytes-from-server="104"
elapsed-time="3" application="UNKNOWN" nested-application="UNKNOWN" username="N/A" roles="N/A"
packet-incoming-interface="ge-0/0/2.0" encrypted="No "]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" reason="response received" source-address="3001::2" source-port="0"
destination-address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-
address="3001::2" nat-source-port="0" nat-destination-address="5001::2" nat-destination-
port="17420" src-nat-rule-name="None" dst-nat-rule-name="None" protocol-id="58" policy-
name="lsys1trust-to-lsys1trust" source-zone-name="lsys1-trust" destination-zone-name="lsys1-
trust" session-id-32="60000218" packets-from-client="1" bytes-from-client="104" packets-from-
server="1" bytes-from-server="104" elapsed-time="3" application="UNKNOWN" nested-
application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"
encrypted="No "]

 ...

user@host>show security log file bin_msg logical-system LSYS1

<14>1 2011-11-17T23:41:46 topstar RT_FLOW - RT_FLOW_SESSION_CREATE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE [junos@2636.1.1.1.2.34
reason="response received" source-address="3001::2" source-port="0" destination-
```

```
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
packets-from-client="1" bytes-from-client="104" packets-from-server="1" bytes-from-server="104"
elapsed-time="3" application="UNKNOWN" nested-application="UNKNOWN" username="N/A" roles="N/A"
packet-incoming-interface="ge-0/0/2.0" encrypted="No "]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" reason="response received" source-address="3001::2" source-port="0"
destination-address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-
address="3001::2" nat-source-port="0" nat-destination-address="5001::2" nat-destination-
port="17420" src-nat-rule-name="None" dst-nat-rule-name="None" protocol-id="58" policy-
name="lsys1trust-to-lsys1trust" source-zone-name="lsys1-trust" destination-zone-name="lsys1-
trust" session-id-32="60000218" packets-from-client="1" bytes-from-client="104" packets-from-
server="1" bytes-from-server="104" elapsed-time="3" application="UNKNOWN" nested-
application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"
encrypted="No "]

 ...

user@host>show security log file bin_msg tenant TSYS1

<14>1 2011-11-17T23:41:46 topstar RT_FLOW - RT_FLOW_SESSION_CREATE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE [junos@2636.1.1.1.2.34
reason="response received" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
packets-from-client="1" bytes-from-client="104" packets-from-server="1" bytes-from-server="104"
elapsed-time="3" application="UNKNOWN" nested-application="UNKNOWN" username="N/A" roles="N/A"
packet-incoming-interface="ge-0/0/2.0" encrypted="No "]
```

```
<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" reason="response received" source-address="3001::2" source-port="0"
destination-address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-
address="3001::2" nat-source-port="0" nat-destination-address="5001::2" nat-destination-
port="17420" src-nat-rule-name="None" dst-nat-rule-name="None" protocol-id="58" policy-
name="lsys1trust-to-lsys1trust" source-zone-name="lsys1-trust" destination-zone-name="lsys1-
trust" session-id-32="60000218" packets-from-client="1" bytes-from-client="104" packets-from-
server="1" bytes-from-server="104" elapsed-time="3" application="UNKNOWN" nested-
application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"
encrypted="No "]

 ...

user@host>show security log stream file s1_f1

<14>1 2011-11-17T23:41:46 topstar RT_FLOW - RT_FLOW_SESSION_CREATE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE [junos@2636.1.1.1.2.34
reason="response received" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
packets-from-client="1" bytes-from-client="104" packets-from-server="1" bytes-from-server="104"
elapsed-time="3" application="UNKNOWN" nested-application="UNKNOWN" username="N/A" roles="N/A"
packet-incoming-interface="ge-0/0/2.0" encrypted="No "]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" reason="response received" source-address="3001::2" source-port="0"
destination-address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-
address="3001::2" nat-source-port="0" nat-destination-address="5001::2" nat-destination-
port="17420" src-nat-rule-name="None" dst-nat-rule-name="None" protocol-id="58" policy-
name="lsys1trust-to-lsys1trust" source-zone-name="lsys1-trust" destination-zone-name="lsys1-
trust" session-id-32="60000218" packets-from-client="1" bytes-from-client="104" packets-from-
```

```
server="1" bytes-from-server="104" elapsed-time="3" application="UNKNOWN" nested-
application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"
encrypted="No "]

 ...

user@host>show security log stream file s1_f1 logical-system LSYS1

<14>1 2011-11-17T23:41:46 topstar RT_FLOW - RT_FLOW_SESSION_CREATE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE [junos@2636.1.1.1.2.34
reason="response received" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
packets-from-client="1" bytes-from-client="104" packets-from-server="1" bytes-from-server="104"
elapsed-time="3" application="UNKNOWN" nested-application="UNKNOWN" username="N/A" roles="N/A"
packet-incoming-interface="ge-0/0/2.0" encrypted="No "]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" reason="response received" source-address="3001::2" source-port="0"
destination-address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-
address="3001::2" nat-source-port="0" nat-destination-address="5001::2" nat-destination-
port="17420" src-nat-rule-name="None" dst-nat-rule-name="None" protocol-id="58" policy-
name="lsys1trust-to-lsys1trust" source-zone-name="lsys1-trust" destination-zone-name="lsys1-
trust" session-id-32="60000218" packets-from-client="1" bytes-from-client="104" packets-from-
server="1" bytes-from-server="104" elapsed-time="3" application="UNKNOWN" nested-
application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"
encrypted="No "]

 ...

user@host>show security log stream file s1_f1 tenant TSYS1
```

```
<14>1 2011-11-17T23:41:46 topstar RT_FLOW - RT_FLOW_SESSION_CREATE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"]
<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE [junos@2636.1.1.1.2.34
reason="response received" source-address="3001::2" source-port="0" destination-
address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-address="3001::2"
nat-source-port="0" nat-destination-address="5001::2" nat-destination-port="17420" src-nat-rule-
name="None" dst-nat-rule-name="None" protocol-id="58" policy-name="lsys1trust-to-lsys1trust"
source-zone-name="lsys1-trust" destination-zone-name="lsys1-trust" session-id-32="60000218"
packets-from-client="1" bytes-from-client="104" packets-from-server="1" bytes-from-server="104"
elapsed-time="3" application="UNKNOWN" nested-application="UNKNOWN" username="N/A" roles="N/A"
packet-incoming-interface="ge-0/0/2.0" encrypted="No "]


<14>1 2011-11-17T23:41:48 topstar RT_FLOW - RT_FLOW_SESSION_CLOSE_LS [junos@2636.1.1.1.2.34
logical-system-name="LSYS1" reason="response received" source-address="3001::2" source-port="0"
destination-address="5001::2" destination-port="17420" service-name="icmpv6" nat-source-
address="3001::2" nat-source-port="0" nat-destination-address="5001::2" nat-destination-
port="17420" src-nat-rule-name="None" dst-nat-rule-name="None" protocol-id="58" policy-
name="lsys1trust-to-lsys1trust" source-zone-name="lsys1-trust" destination-zone-name="lsys1-
trust" session-id-32="60000218" packets-from-client="1" bytes-from-client="104" packets-from-
server="1" bytes-from-server="104" elapsed-time="3" application="UNKNOWN" nested-
application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/2.0"
encrypted="No "]
 ...
```

show security log file test1.pb

```
user@host> show security log file test1.pb
<14>1 2023-03-17T00:06:55 10.53.78.91 RT_LOG_SELF_TEST - SECINTEL_ACTION_LOG
[junos@2636.1.1.1.2.129 category="secintel" sub-category="CC" action="block" action-
detail="test" http-host="test" threat-severity="5" source-address="1.16.16.16" source-
port="16384" destination-address="2.16.16.16" destination-port="32768" protocol-id="17"
application="test" nested-application="test" feed-name="test" policy-name="test" profile-
name="test" username="Fake username" roles="test" session-id="1" source-zone-name="Fake src
zone" destination-zone-name="Fake dst zone" occur-count="3"]
<14>1 2023-03-17T00:06:55 10.53.78.91 RT_LOG_SELF_TEST - AAMW_ACTION_LOG [junos@2636.1.1.1.2.129
```

```
hostname="test" file-category="virus" verdict-number="5" malware-info="Test-File" action="block"
list-hit="test" file-hash-lookup="test" source-address="1.16.16.16" source-port="16384"
destination-address="2.16.16.16" destination-port="32768" protocol-id="17" application="test"
nested-application="test" policy-name="test" username="Fake username" roles="test" session-
id="1" source-zone-name="Fake src zone" destination-zone-name="Fake dst zone" sample-
sha256="da26ba1e13ce4702bd5154789ce1a699ba206c12021d9823380febd795f5b002" file-name="test_name"
url="www.test.com"]
<14>1 2023-03-17T00:06:55 10.53.78.91 RT_LOG_SELF_TEST - AAMW_MALWARE_EVENT_LOG
[junos@2636.1.1.1.2.129 timestamp="fake-time" tenant-id="test" sample-
sha256="da26ba1e13ce4702bd5154789ce1a699ba206c12021d9823380febd795f5b002" source-
address="1.16.16.16" verdict-number="5" malware-info="Test-File" username="Fake username"
hostname="test" session-id="1" url="www.test.com"]
...
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

show security log | 24769

show security log profile

IN THIS SECTION

- Syntax | 24782
- Description | 24782
- Required Privilege Level | 24782
- Sample Output | 24782
- Release Information | 24791

Syntax

```
show security log profile <config|template>
```

Description

Use this command to see the log template details.

Required Privilege Level

view

Sample Output

show security log profile traditional-firewall

```
user@host> show security log profile traditional-firewall
The template: traditional-firewall
==========================
category:     session
fields selected:
reason
source-address
source-port
destination-address
destination-port
service-name
nat-source-address
nat-source-port
nat-destination-address
nat-destination-port
protocol-id
policy-name
source-zone-name
destination-zone-name
session-id
packets-from-client
bytes-from-client
packets-from-server
```

```
bytes-from-server
elapsed-time
username
packet-incoming-interface
session-flag
==========================
category:     webfilter
fields selected:
field-all
==========================
category:     antivirus
fields selected:
field-all
==========================
category:     content-filter
fields selected:
field-all
==========================
category:     antispam
fields selected:
field-all
==========================
category:     idp
fields selected:
field-all
==========================
category:     secintel
fields selected:
field-all
==========================
category:     aamw
fields selected:
field-all
==========================
category:     ssl-proxy
fields selected:
field-all
==========================
category:     apptrack
fields selected:
field-all
```

show security log profile sd-wan

```
user@host> show security log profile sd-wan
The template: sd-wan
=========================
category:     session
fields selected:
session-id
reason
source-address
source-port
destination-address
destination-port
nat-source-address
nat-source-port
nat-destination-address
nat-destination-port
protocol-id
policy-name
source-zone-name
destination-zone-name
packets-from-client
bytes-from-client
packets-from-server
bytes-from-server
elapsed-time
application
nested-application
username
packet-incoming-interface
application-category
application-sub-category
application-risk
session-flag
=========================
category:     webfilter
fields selected:
session-id
category
reason
profile
url
```

```
obj
username
urlcategory-risk
==========================
category:     antivirus
fields selected:
session-id
profile-name
filename
temporary-filename
name
url
username
action
==========================
category:     content-filter
fields selected:
session-id
argument
profile-name
action
reason
username
filename
==========================
category:     antispam
fields selected:
session-id
profile-name
action
reason
username
==========================
category:     idp
fields selected:
session-id
message-type
rule-name
rulebase-name
policy-name
export-id
repeat-count
action
```

```
threat-severity
attack-name
packet-log-id
alert
username
xff-header
message
==========================
category:     secintel
fields selected:
session-id
category
sub-category
action
action-detail
http-host
threat-severity
feed-name
policy-name
profile-name
username
==========================
category:     aamw
fields selected:
session-id
hostname
file-category
verdict-number
malware-info
action
list-hit
file-hash-lookup
policy-name
username
url
==========================
category:     ssl-proxy
fields selected:
session-id
profile-name
message
username
==========================
```

```
category:     apptrack
fields selected:
source-address
source-port
destination-address
destination-port
service-name
application
nested-application
policy-name
source-zone-name
destination-zone-name
packets-from-client
bytes-from-client
packets-from-server
bytes-from-server
elapsed-time
profile-name
rule-name
routing-instance
destination-interface-name
uplink-incoming-interface-name
uplink-tx-bytes
uplink-rx-bytes
multipath-rule-name
src-vrf-grp
dst-vrf-grp
username
session-id
reason
```

show security log profile unified-ngfw

```
user@host> show security log profile unified-ngfw
The template: unified-ngfw
==========================
category:     session
fields selected:
session-id
reason
source-address
```

```
source-port
destination-address
destination-port
nat-source-address
nat-source-port
nat-destination-address
nat-destination-port
protocol-id
policy-name
source-zone-name
destination-zone-name
packets-from-client
bytes-from-client
packets-from-server
bytes-from-server
elapsed-time
application
nested-application
username
packet-incoming-interface
application-category
application-sub-category
application-risk
session-flag
==========================
category:     webfilter
fields selected:
session-id
category
reason
profile
url
obj
username
urlcategory-risk
==========================
category:     antivirus
fields selected:
session-id
profile-name
filename
temporary-filename
name
```

```
url
username
action
==========================
category:     content-filter
fields selected:
session-id
argument
profile-name
action
reason
username
filename
==========================
category:     antispam
fields selected:
session-id
profile-name
action
reason
username
==========================
category:     idp
fields selected:
session-id
message-type
rule-name
rulebase-name
policy-name
export-id
repeat-count
action
threat-severity
attack-name
packet-log-id
alert
username
xff-header
message
==========================
category:     secintel
fields selected:
session-id
```

```
category
sub-category
action
action-detail
http-host
threat-severity
feed-name
policy-name
profile-name
username
==========================
category:     aamw
fields selected:
session-id
hostname
file-category
verdict-number
malware-info
action
list-hit
file-hash-lookup
policy-name
username
url
==========================
category:     ssl-proxy
fields selected:
session-id
profile-name
message
username
==========================
category:     apptrack
fields selected:
session-id
application
nested-application
username
profile-name
rule-name
routing-instance
destination-interface-name
uplink-incoming-interface-name
```

```
uplink-tx-bytes
uplink-rx-bytes
apbr-policy-name
amr-rule-name
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

profile (security) | 13292

show security log query

IN THIS SECTION

- Show Security Log Query | 24791
- Show Security Log Stream | 24792
- Description | 24792
- Options | 24792
- Required Privilege Level | 24793
- Sample Output | 24793
- Release Information | 24794

Show Security Log Query

```
show security log query {category all | utm | idp | alg | appqos | flow | fw-auth | gtp | ipsec
| nat | pst-ds-lite | rtlog | screen | sctp | secintel} count < count>
[src-ip <src-ip>]
[dst-ip <dst-ip>]
[src-port <src-port>]
[dst-port <dst-port>]
```

```
[application <application>]
[user <user>]
[event-type <event-type>]
[service <service>]
[start-time <start-time>]
[stop-time <stop-time>]
```

Show Security Log Stream

```
show security log stream
file <filename>
```

Description

- `show security log query`—View the security log from the database with query conditions.
- `show security log stream file`—View all the security log messages in the specified log file. Use the `/var/log/ hostlogs` directory to search the specified log file, and use the `show security log stream file` command to view logs in log files in the `/var/log/hostlogs` directory.

Options

- count—The log number to output.
- scr-ip—The source IP address of log messages.
- dst-ip—The destination IP address of log messages.
- src-port—The source port of log messages.
- dst-port—The destination port of log messages.
- application—The application of log messages.
- user—The user of log messages.
- event-type—The event type of log messages.
- service—The service of log messages.
- start-time—The earliest timestamp of log messages; the format for time is YYYY-MM-DDTHH:MM:SS.

- stop-time—The latest timestamp of log messages.

Required Privilege Level

view

Sample Output

show security log query

```
rootr@dut> show security log query category flow count 20 src-ip 211.0.0.2 start-time
2013-11-29T00:00:00 end-time 2013-11-29T23:59:00
<14>1 2013-11-29T16:01:26.820+08:00 plat02 RT_FLOW - RT_FLOW_SESSION_CLOSE reason="CLI" source-
address="211.0.0.2" source-port="20263" destination-address="211.0.1.3" destination-port="4903"
service-name="None" nat-source-address="30.0.11.11" nat-source-port="27140" nat-destination-
address="211.0.1.3" nat-destination-port="4903" src-nat-rule-name="src_rs2_rule1" dst-nat-rule-
name="None" protocol-id="17" policy-name="p1" source-zone-name="green" destination-zone-
name="red" session-id-32="30" packets-from-client="1" bytes-from-client="60" packets-from-
server="0" bytes-from-server="0" elapsed-time="92683" application="UNKNOWN" nested-
application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/1.0"
encrypted="UNKNOWN"
```

show security log stream file <file-name>

```
root@dut> show security log stream file traffic.log
<14>1 2013-11-29T16:01:26.820+08:00 plat02 RT_FLOW - RT_FLOW_SESSION_CLOSE reason="CLI" source-
address="211.0.0.2" source-port="20263" destination-address="211.0.1.3" destination-port="4903"
service-name="None" nat-source-address="30.0.11.11" nat-source-port="27140" nat-destination-
address="211.0.1.3" nat-destination-port="4903" src-nat-rule-name="src_rs2_rule1" dst-nat-rule-
name="None" protocol-id="17" policy-name="p1" source-zone-name="green" destination-zone-
name="red" session-id-32="30" packets-from-client="1" bytes-from-client="60" packets-from-
server="0" bytes-from-server="0" elapsed-time="92683" application="UNKNOWN" nested-
application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-interface="ge-0/0/1.0"
encrypted="UNKNOWN"] session closed CLI: 211.0.0.2/20263->211.0.1.3/4903 None 30.0.11.11/27140-
>211.0.1.3/4903 src_rs2_rule1 None 17 p1 green red 30 1(60) 0(0) 92683 UNKNOWN UNKNOWN N/A(N/A)
ge-0/0/1.0 UNKNOWN
```

Release Information

Command introduced in Junos OS Release 15.1X49-D70.

show security log report

IN THIS SECTION

- Syntax | 24794
- Description | 24794
- Options | 24795
- Required Privilege Level | 24795
- Output Fields | 24795
- Sample Output | 24795
- Release Information | 24795

Syntax

```
in-detail
in-interval
summary
top
```

Description

Displays the security log report settings.

On-box reporting offers a comprehensive reporting facility where your security management team can spot a security event when it occurs. Immediately access and review pertinent details about the event, and quickly decide appropriate remedial action.

Options

in-detail	Displays the detail log content
in-interval	Displays the count in intervals
summary	Displays the summary information
top	Displays the top number to be calculated

Required Privilege Level

view

Output Fields

Sample Output

show security log report summary session-all tenant TSYS1

```
user@host> show security log report summary session-all tenant TSYS1
total-count: 4
```

Show security log report summary session-all root-logical-system

```
user@host> Show security log report summary session-all root-logical-system
total-count: 4
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

exclude (Security Log) | 6627

clear security log | 3998

show security log report in-detail

IN THIS SECTION

- Syntax | 24796
- Description | 24796
- Options | 24796
- Required Privilege Level | 24797
- Output Fields | 24797
- Sample Output | 24797
- Release Information | 24802

Syntax

```
show security log report in-detail {all | atp | dns | idp | ipsec-vpn | screen | session |
threat | utm}
```

Description

This command allows you to generate detail log content.

Options

all Display all audit event logs stored in the device memory.

atp Display all ATP event logs.

dns Display all DNS event logs.

idp Display all IDP event logs.

ipsec-vpn Display all IPsec-VPN event logs.

screen Display all sceen event logs.

session Display all session logs.

threat Display all threat related logs.

Content Security Display all Content Security event logs.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show security log report in-detail` command.

Table 1275: name-of-command Output Fields

Field Name	Field Description
`Event time`	The timestamp when the security event was received.
`Message`	The message describing the security event.

Sample Output

show security log report in-detail

```
user@host> show security log report in-detail all
<14>1 2023-01-13T07:40:06 host RT_LOG_SELF_TEST - IDP_ATTACK_LOG_EVENT [message-type="0" policy-
name="policy_5_7_5" repeat-count="0" elapsed-time="0" source-interface-name="ge-4/1/7.0"
destination-interface-name="ge-4/3/7.0" packet-log-id="0" xff-header="-" source-
address="192.0.2.0/24" destination-address="198.51.100.0/24" application-name="TELNET"
username="NONE" rule-name="r15" action="action15" threat-severity="Critical" attack-
name="attack15" cve-id="CVE-2014-0160" session-id="1" source-zone-name="Fake src zone"
destination-zone-name="Fake dst zone" source-port="16384" destination-port="32768" protocol-
name="TCP"]
<14>1 2023-01-13T07:40:06 host RT_LOG_SELF_TEST - IDP_ATTACK_LOG_EVENT [message-type="0" policy-
name="policy_5_7_5" repeat-count="0" elapsed-time="0" source-interface-name="ge-4/1/7.0"
destination-interface-name="ge-4/3/7.0" packet-log-id="0" xff-header="-" source-
```

```
address="192.0.2.0/24" destination-address="198.51.100.0/24" application-name="TELNET"
username="NONE" rule-name="r15" action="action15" threat-severity="Critical" attack-
name="attack15" cve-id="CVE-2014-0160" session-id="1" source-zone-name="Fake src zone"
destination-zone-name="Fake dst zone" source-port="16384" destination-port="32768" protocol-
name="TCP"]
<14>1 2023-01-23T00:41:36 host RT_LOG_SELF_TEST - RT_FLOW_SESSION_CLOSE [nat-source-
address="192.168.7.0/24" nat-source-port="31000" nat-destination-address="235.13.2.1/24" nat-
destination-port="32768" packets-from-client="4294967295" application="Fake application" session-
flag="0" source-address="192.0.2.2/24" destination-address="198.51.100.2/24" nested-
application="ANSWERS" username="Lex" policy-name="Fake policy" bytes-from-server="4294967292"
packets-from-server="4294967294" bytes-from-client="4294967293" packet-incoming-interface="Fake
packet incoming if" source-zone-name="Fake src zone" destination-zone-name="Fake dst zone"
session-id="1" source-port="16384" destination-port="32768" protocol-id="UDP" elapsed-
time="88888888" reason="Test reason"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - ANTISPAM_SPAM_DETECTED_MT [action="BLOCKED"
source-address="192.0.2.1/24" source-name=spamtest@spamtest.com username="192.0.2.1/24"
reason="TESTSPAM" profile-name="PROFILE" session-id="1" source-zone-name="src-zone" destination-
zone-name="dst-zone"]
<12>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - AV_VIRUS_DETECTED_MT [action="BLOCKED" source-
address="192.0.2.1/24" destination-address="198.51.100.2/24" name="EICAR-Test-File" url=http://
www.viruslist.com/en/search?VN=EICAR-Test-File username="192.0.2.1/24" profile-name="PROFILE"
session-id="1" filename="Fake file" source-zone-name="src-zone" destination-zone-name="dst-zone"
source-port="16384" destination-port="32768"]
<12>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - WEBFILTER_URL_BLOCKED [application="TEST-
APPLICATION" urlcategory-risk="0" source-address="192.0.2.1/24" destination-
address="198.51.100.2/24" url=http://www.viruslist.com/en/search?VN=EICAR-Test-File
username="192.0.2.1/24" reason="TESTSPAM" profile="PROFILE" category="N/A" session-id="1" nested-
application="SQUIRRELMAIL" source-zone-name="src-zone" destination-zone-name="dst-zone" source-
port="16384" destination-port="32768"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - WEBFILTER_URL_PERMITTED [application="TEST-
APPLICATION" urlcategory-risk="0" source-address="192.0.2.1/24" destination-
address="198.51.100.2/24" url=http://www.viruslist.com/en/search?VN=EICAR-Test-File
username="192.0.2.1/24" reason="TESTSPAM" profile="PROFILE" category="N/A" session-id="1" nested-
application="SQUIRRELMAIL" source-zone-name="src-zone" destination-zone-name="dst-zone" source-
port="16384" destination-port="32768"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - CONTENT_FILTERING_BLOCKED_MT [action="BLOCKED"
source-address="192.0.2.1/24" destination-address="198.51.100.2/24" username="192.0.2.1/24"
reason="TESTSPAM" profile-name="PROFILE" argument="TCP" session-id="1" filename="Fake file"
source-zone-name="src-zone" destination-zone-name="dst-zone" source-port="16384" destination-
port="32768"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_SCREEN_ICMP [interface-name="test"
action="test" source-address="192.0.2.1/24" destination-address="198.51.100.2/24" attack-
name="test" source-zone-name="Fake src zone"]
```

```
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_SCREEN_IP [interface-name="test_interface"
action="test_action" source-address="192.0.2.1/24" destination-address="198.51.100.2/24" attack-
name="test" source-zone-name="Fake src zone" protocol-id="17"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_SCREEN_SESSION_LIMIT [interface-
name="test_interface" action="test_action" source-address="192.0.2.1/24" destination-
address="198.51.100.2/24" attack-name="test" source-zone-name="Fake src zone" source-
port="16384" destination-port="32768" protocol-id="17"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_SCREEN_TCP [interface-name="test_interface"
action="test_action" source-address="192.0.2.1/24" destination-address="198.51.100.2/24" attack-
name="test" source-zone-name="Fake src zone" source-port="16384" destination-port="32768"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_SCREEN_TCP_DST_IP [interface-
name="test_interface" action="test_action" destination-address="198.51.100.2/24" attack-
name="test" source-zone-name="Fake src zone"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_SCREEN_TCP_SRC_IP [interface-
name="test_interface" action="test_action" source-address="192.0.2.1/24" attack-name="test"
source-zone-name="Fake src zone"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_SCREEN_UDP [interface-name="test_interface"
action="test_action" source-address="192.0.2.1/24" destination-address="198.51.100.2/24" attack-
name="test" source-zone-name="Fake src zone" source-port="16384" destination-port="32768"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_FLOW_SESSION_CREATE [nat-source-
address="192.168.7.0/24" nat-source-port="30000" nat-destination-address="235.13.2.1/24" nat-
destination-port="32768" application="Fake application" source-address="192.168.7.0/24"
destination-address="235.13.2.1/24" nested-application="ANSWERS" username="Fake username" policy-
name="Fake policy" packet-incoming-interface="Fake packet incoming if" source-zone-name="Fake
src zone" destination-zone-name="Fake dst zone" session-id="1" source-port="16384" destination-
port="32768" protocol-id="UDP"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_FLOW_SESSION_CLOSE [nat-source-
address="192.168.7.0/24" nat-source-port="30000" nat-destination-address="235.13.2.1/24" nat-
destination-port="32768" packets-from-client="4294967295" application="Fake application" session-
flag="0" source-address="192.168.7.0/24" destination-address="235.13.2.1/24" nested-
application="ANSWERS" username="Fake username" policy-name="Fake policy" bytes-from-
server="4294967292" packets-from-server="4294967294" bytes-from-client="4294967293" packet-
incoming-interface="Fake packet incoming if" source-zone-name="Fake src zone" destination-zone-
name="Fake dst zone" session-id="1" source-port="16384" destination-port="32768" protocol-
id="UDP" elapsed-time="88888888" reason="Test reason"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_FLOW_SESSION_DENY [application="Fake
application" source-address="192.168.7.0/24" destination-address="235.13.2.1/24" nested-
application="ANSWERS" username="Fake username" policy-name="Fake policy" packet-incoming-
interface="Fake packet incoming if" source-zone-name="Fake src zone" destination-zone-name="Fake
dst zone" session-id="1" source-port="16384" destination-port="32768" protocol-id="UDP"
reason="Test reason"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_IPSEC_BAD_SPI [interface-name="Fake
interface" tunnel-id="1" length="1024" type="Fake_Type" index="1010" sequence-number="4000"
```

```
source-address="192.168.7.0/24" destination-address="235.13.2.1/24"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_IPSEC_REPLAY [interface-name="Fake
interface" tunnel-id="1" length="1024" type="Fake_Type" index="1010" sequence-number="4000"
source-address="192.168.7.0/24" destination-address="235.13.2.1/24"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_IPSEC_PV_REPLAY [interface-name="Fake
interface" tunnel-id="1" length="1024" type="Fake_Type" index="1010" sequence-number="4000"
source-address="192.168.7.0/24" destination-address="235.13.2.1/24"]
<8>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_IPSEC_D3P_ERR [tunnel-id="1" length="1024"
type="Fake_Type" index="1010" D3P error="Fake error" source-address="192.168.7.0/24" destination-
address="235.13.2.1/24"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_IPSEC_DECRYPT_BAD_PAD [tunnel-id="1"
index="1010" source-address="192.168.7.0/24" destination-address="235.13.2.1/24"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_IPSEC_PV_SYM_KEYGEN [error-message="Fake
error"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_IPSEC_AUTH_FAIL [error-message="Fake error"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_IPSEC_PV_ENCRYPTION [error-message="Fake
error"]
<11>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - RT_IPSEC_PV_DECRYPTION [error-message="Fake
error"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - IDP_ATTACK_LOG_EVENT [message-type="0" policy-
name="policy_5_7_5" repeat-count="0" elapsed-time="0" source-interface-name="ge-4/1/7.0"
destination-interface-name="ge-4/3/7.0" packet-log-id="0" xff-header="-" source-
address="192.0.2.1/24" destination-address="198.51.100.2/24" application-name="TELNET"
username="NONE" rule-name="1" action="TRAFFIC_IPACTION_NOTIFY" threat-severity="INFO" attack-
name="-" cve-id="CVE-2014-0160" session-id="1" source-zone-name="Fake src zone" destination-zone-
name="Fake dst zone" source-port="16384" destination-port="32768" protocol-name="TCP"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - SECINTEL_ACTION_LOG [http-host="test"
application="test" feed-name="test" policy-name="test" profile-name="test" occur-count="3"
source-address="192.0.2.1/24" destination-address="198.51.100.2/24" threat-severity="5" sub-
category="CC" action="block" session-id="1" username="Fake username" nested-
application="ANSWERS" source-zone-name="Fake src zone" destination-zone-name="Fake dst zone"
source-port="16384" destination-port="32768" protocol-id="17"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - AAMW_ACTION_LOG [hostname="test" list-
hit="test" file-hash-lookup="test" application="test" policy-name="test" source-
address="192.0.2.1/24" destination-address="198.51.100.2/24" verdict-number="5" file-
category="virus" malware-info="Test-File" action="block" file-name="test_name" url=www.test.com
session-id="1" username="Fake username" nested-application="ANSWERS" source-zone-name="Fake src
zone" destination-zone-name="Fake dst zone" source-port="16384" destination-port="32768"
protocol-id="17" sample-
sha256="da26ba1e13ce4702bd5154789ce1a699ba206c12021d9823380febd795f5b002"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - AAMW_MALWARE_EVENT_LOG [timestamp="fake-time"
tenant-id="test" hostname="test" source-address="192.0.2.1/24" verdict-number="5" malware-
info="Test-File" url=www.test.com session-id="1" username="Fake username" sample-
```

```
sha256="da26ba1e13ce4702bd5154789ce1a699ba206c12021d9823380febd795f5b002"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - AAMW_SMTP_ACTION_LOG [list-hit="test"
application="test" policy-name="test" source-address="192.0.2.1/24" destination-
address="198.51.100.2/24" verdict-number="5" file-category="virus" malware-info="Test-File"
action="block" file-name="test_name" session-id="1" username="Fake username" nested-
application="ANSWERS" source-zone-name="Fake src zone" destination-zone-name="Fake dst zone"
source-port="16384" destination-port="32768" protocol-id="17" send-
from=send_from_selftest@email.com send-to=send_to_selftest@email.com sample-
sha256="da26ba1e13ce4702bd5154789ce1a699ba206c12021d9823380febd795f5b002"]
<14>1 2023-02-26T23:15:52 host RT_LOG_SELF_TEST - AAMW_IMAP_ACTION_LOG [list-hit="test"
application="test" policy-name="test" source-address="192.0.2.1/24" destination-
address="198.51.100.2/24" verdict-number="5" file-category="virus" malware-info="Test-File"
action="block" file-name="test_name" session-id="1" username="Fake username" nested-
application="ANSWERS" source-zone-name="Fake src zone" destination-zone-name="Fake dst zone"
source-port="16384" destination-port="32768" protocol-id="17" send-
from=send_from_selftest@email.com send-to=send_to_selftest@email.com sample-
sha256="da26ba1e13ce4702bd5154789ce1a699ba206c12021d9823380febd795f5b002"]
```

show security log report in-detail dns

```
user@host> show security log report in-detail dns
<14>1 2022-08-22T00:37:34 host RT_LOG_SELF_TEST - SMS_FALLBACK_EVENT [application="test" query-
type="test" source-address="192.0.2.1/24" destination-address="198.51.100.2/24" query="test"
action="test" source-port="36733" destination-port="80" protocol-id="unkn" policy-name="test"
username="test" session-id="1" source-zone-name="test" destination-zone-name="test" nested-
application="DNS" dns-model="test" reason="test"]
```

show security log report in-detail threat

```
user@host> show security log report in-detail threat
<14>1 2022-08-22T00:37:34 host RT_LOG_SELF_TEST - SMS_FALLBACK_EVENT [application="test" query-
type="test" source-address="192.0.2.1/24" destination-address="198.51.100.2/24" query="test"
action="test" source-port="36733" destination-port="80" protocol-id="unkn" policy-name="test"
username="test" session-id="1" source-zone-name="test" destination-zone-name="test" nested-
application="DNS" dns-model="test" reason="test"]
```

show security log report in-detail all expression

```
user@host> show security log report in-detail all expression ( source-address=192.0.3.2/24 &&
( session-id=785 || session-id=786 ))
<14>1 2022-06-15T05:04:35 10.49.18.230 RT_FLOW - RT_FLOW_SESSION_CREATE [nat-source-
address="192.0.3.3/24" nat-source-port="47452" nat-destination-address="198.51.100.3/24 nat-
destination-port="80" protocol-id="6" application="HTTP" source-address="192.0.3.2/24"
destination-address="198.51.100.3/24" nested-application="BING" username="192.0.3.2/24" policy-
name="policy1" packet-incoming-interface="ge-0/0/0.0" source-zone-name="trust" destination-zone-
name="untrust" session-id="785" source-port="0" destination-port="0"]
<14>1 2022-06-15T05:04:59 10.49.18.230 RT_FLOW - RT_FLOW_SESSION_CLOSE [reason="TCP FIN" nat-
source-address="192.0.3.3/24" nat-source-port="47452" nat-destination-address="198.51.100.3/24"
nat-destination-port="80" protocol-id="6" packets-from-client="15840" elapsed-time="24"
application="HTTP" session-flag="0" source-address="192.0.3.2/24" destination-
address="198.51.100.3/24" nested-application="BING" username="192.0.3.2/24" policy-
name="policy1" bytes-from-server="218102537" packets-from-server="157892" bytes-from-
client="828078" packet-incoming-interface="ge-0/0/0.0" source-zone-name="trust" destination-zone-
name="untrust" session-id="785" source-port="0" destination-port="0"]
```

Release Information

Command introduced in Junos OS Release 23.1R1 .

RELATED DOCUMENTATION

On-Box Logging and Reporting

show security log report in-interval

IN THIS SECTION

- Syntax | 24803
- Description | 24803
- Options | 24803
- Required Privilege Level | 24803

Sample Output | 24804
Release Information | 24804

Syntax

```
show security log report in-interval {atp | dns | idp | ipsec-vpn | screen | session | threat |
utm}
```

Description

This command allows you to generate the events logged between certain time intervals.

Options

atp Display all ATP event logs.

dns Display all DNS event logs.

idp Display all IDP event logs.

ipsec-vpn Display all IPsec-VPN event logs.

screen Display all sceen event logs.

session Display all session logs.

threat Display all threat related logs.

Content Security Display all Content Security event logs.

Required Privilege Level

view

Sample Output

show security log report in-interval

```
user@host> show security log report in-interval session all start-time 2023-02-26T23:15:48 stop-time 2023-02-26T23:15:55
Time-intervals        Blocked/Denied events   All events
-------------------------------------------------------
2023-02-26T23:15:48   0                       0
2023-02-26T23:15:49   0                       0
2023-02-26T23:15:50   0                       0
2023-02-26T23:15:51   0                       0
2023-02-26T23:15:52   1                       3
2023-02-26T23:15:53   0                       0
2023-02-26T23:15:54   0                       0
2023-02-26T23:15:55   0                       0
```

Release Information

Command introduced in Junos OS Release 23.1R1 .

RELATED DOCUMENTATION

On-Box Logging and Reporting

show security log report summary

IN THIS SECTION

- Syntax | 24805
- Description | 24805
- Options | 24805
- Required Privilege Level | 24805
- Sample Output | 24806

Release Information | 24807

Syntax

```
show security log report summary {all | atp | dns | idp | ipsec-vpn | screen | session | threat
| utm}
```

Description

This command allows you to generate the summary of the events. In this way, you can fine-tune the report to your needs, and displays only the data that you want to use.

Options

all Display summary of all audit event logs stored in the device memory.

atp Display summary of all ATP event logs.

dns Display summary of all DNS event logs.

idp Display summary of all IDP event logs.

ipsec-vpn Display summary of all IPsec-VPN event logs.

screen Display summary of all sceen event logs.

session Display summary of all session logs.

threat Display summary of all threat related logs.

Content Security Display summary of all Content Security event logs.

Required Privilege Level

view

Sample Output

show security log report summary

```
user@host> show security log report summary all
total-count: 33
```

```
user@host> show security log report summary session all
total-count: 4
```

show security log report summary dns

```
user@host> show security log report summary dns
total-count: 1
```

show security log report summary threat

```
user@host> show security log report summary threat
total-count: 1
KNOWN        UNKNOWN
--------------------
1          0
Application    IDP        SCREEN      antispam  antivirus webfilter cntfilter sky
infecth   C&C       dns         Source-IP   Destination-IP
-----------------------------------------------------------------------------------------------
-----------------------------------------------
0              0          0           0         0         0         0         0
0         0         1           0           0
```

show security log report summary all expression

```
user@host> show security log report summary all expression (source-address= <source-address/
mask> && (session-id=<session-id> || session-id=<session-id>))
total-count: 32
```

Release Information

Command introduced in Junos OS Release 23.1R1 .

RELATED DOCUMENTATION

On-Box Logging and Reporting

show security log severity

IN THIS SECTION

- Syntax | 24807
- Description | 24807
- Required Privilege Level | 24807
- Output Fields | 24808
- Release Information | 24808

Syntax

```
show security log severity
```

Description

Display severity information for the event.

Required Privilege Level

view

Output Fields

Table 1276 on page 24808 lists the output fields for the `show security log severity` command. Output fields are listed in the approximate order in which they appear.

Table 1276: show security log severity Output Fields

Field Name	Field Description
alert	Alert severity
crit	Critical severity
debug	Debug severity
emerg	Emergency severity
err	Error severity
info	Information severity
notice	Notice severity
warning	Warning severity

Release Information

Command introduced in Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

show security log | 24769

show security macsec connections (SRX)

IN THIS SECTION

- Syntax | 24809
- Description | 24809
- Options | 24809
- Required Privilege Level | 24809
- Output Fields | 24810
- Sample Output | 24811
- Release Information | 24812

Syntax

```
show security macsec connections
<interface interface-name>
```

Description

Display the status of the active MACsec connections on the device.

Options

none	Display MACsec connection information for all interfaces on the device.
interface *interface-name*	(Optional) Display MACsec connection information for the specified interface only.

Required Privilege Level

view

Output Fields

Table 1277 on page 24810 lists the output fields for the `show security macsec connections` command. Output fields are listed in the approximate order in which they appear.

Table 1277: show security macsec connections Output Fields

Field Name	Field Description
Fields for Interface	
`Interface name`	Name of the interface.
`CA name`	Name of the connectivity association. A connectivity association is named using the `connectivity-association` statement when you are enabling MACsec.
`Cipher suite`	Name of the cipher suite used for encryption.
`Key server offset`	Offset setting. The offset is set using the `offset` statement when configuring the connectivity association when using static connectivity association key (CAK) or dynamic security mode.
`Replay protect`	Replay protection setting. Replay protection is enabled when this output is `on` and disabled when this output is `off`. You can enable replay protection using the `replay-protect` statement in the connectivity association.
Outbound secure channels	Displays outgoing packet number.
Inbound secure channels	Displays source identifier and secure associations detail.

Sample Output

show security macsec connections

```
user@host> show security macsec connections
Interface name: fxp1
        CA name: ca1
        Cipher suite: GCM-AES-128   Encryption: on
        Key server offset: 0        Include SCI: no
        Replay protect: off         Replay window: 0
```

show security macsec connections (SRX1600 and SRX2300)

```
user@host> show security macsec connections
    Interface name: em0
        CA name: ca_mka_01
        Cipher suite: GCM-AES-128   Encryption: on
        Key server offset: 0        Include SCI: no
        Replay protect: off         Replay window: 0
          Outbound secure channels
            SC Id: 02:00:00:01:01:04/1
            Outgoing packet number: 1914287
            Secure associations
            AN: 0 Status: inuse Create time: 07:33:26
          Inbound secure channels
            SC Id: 02:00:00:02:01:04/1
            Secure associations
            AN: 0 Status: inuse Create time: 07:33:26

    Interface name: em1
        CA name: ca_mka_01
        Cipher suite: GCM-AES-128   Encryption: on
        Key server offset: 0        Include SCI: no
        Replay protect: off         Replay window: 0
          Outbound secure channels
            SC Id: 02:00:01:01:01:04/1
            Outgoing packet number: 108885
            Secure associations
            AN: 0 Status: inuse Create time: 07:33:26
```

```
            Inbound secure channels
              SC Id: 02:00:01:02:01:04/1
              Secure associations
              AN: 0 Status: inuse Create time: 07:33:26
```

Release Information

Command introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

show security mka statistics | 24857

show security macsec connections

IN THIS SECTION

- Syntax | 24812
- Description | 24813
- Options | 24813
- Required Privilege Level | 24813
- Output Fields | 24813
- Sample Output | 24815
- Release Information | 24815

Syntax

```
show security macsec connections
<interface interface-name>
```

Description

Display the status of the active MACsec connections on the switch.

This command does not display output when MACsec is enabled using static secure association key (SAK) security mode.

Options

none	Display MACsec connection information for all interfaces on the switch.
interface *interface-name*	(Optional) Display MACsec connection information for the specified interface only.

Required Privilege Level

view

Output Fields

Table 1278 on page 24813 lists the output fields for the `show security macsec connections` command. Output fields are listed in the approximate order in which they appear.

Table 1278: show security macsec connections Output Fields

Field Name	Field Description
Fields for Interface	
`Interface name`	Name of the interface.
`CA name`	Name of the connectivity association. A connectivity association is named using the `connectivity-association` statement when you are enabling MACsec.
`Cipher suite`	Name of the cipher suite used for encryption.

Table 1278: show security macsec connections Output Fields *(Continued)*

Field Name	Field Description
`Encryption`	Encyption setting. Encryption is enabled when this output is `on` and disabled when this output is `off`. The encryption setting is set using the `no-encryption` statement in the connectivity association when using static connectivity association key (CAK) security mode and is set using the `encryption` statement in the secure channel when using static secure association key (SAK) or dynamic security mode.
`Key server offset`	Offset setting. The offset is set using the `offset` statement when configuring the connectivity association when using static connectivity association key (CAK) or dynamic security mode or the secure channel when using static secure association key (SAK) security mode.
`Include SCI`	SCI tagging. The SCI tag is included on packets in a secure channel when this output is `yes`, and not included on packets in a secure channel when this output is `no`. You can enable SCI tagging using the `include-sci` statement in the connectivity association. **NOTE**: SCI tags are automatically appended to packets leaving a MACsec-enabled interface on an EX4300 switch. The `include-sci` option is, therefore, not available on EX4300 switches. The output for the `Include SCI` field is `yes`.
`Replay protect`	Replay protection setting. Replay protection is enabled when this output is `on` and disabled when this output is `off`. You can enable replay protection using the `replay-protect` statement in the connectivity association.
`Replay window`	Replay protection window setting. This output is set to `0` when replay protection is disabled, and is the size of the replay window, in number of packets, when replay protection is enabled. The size of the replay window is configured using the `replay-window-size` statement in the connectivity association.

Sample Output

show security macsec connections

```
user@host> show security macsec connections
Interface name: xe-0/1/0
        CA name: CA1
        Cipher suite: GCM-AES-128   Encryption: on
        Key server offset: 0        Include SCI: no
        Replay protect: off         Replay window: 0
```

Release Information

Command introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

show security macsec statistics | 24819

show security macsec connections (MX Series)

IN THIS SECTION

- Syntax | 24816
- Description | 24816
- Options | 24816
- Required Privilege Level | 24816
- Output Fields | 24816
- Sample Output | 24818
- Release Information | 24819

Syntax

```
show security macsec connections
<interface interface-name>
```

Description

Display the status of the active MACsec connections on the router.

Options

none	Display MACsec connection information for all interfaces on the switch.
interface *interface-name*	(Optional) Display MACsec connection information for the specified interface only.

Required Privilege Level

view

Output Fields

Table 1279 on page 24816 lists the output fields for the `show security macsec connections` command. Output fields are listed in the approximate order in which they appear.

Table 1279: show security macsec connections Output Fields

Field Name	Field Description
Fields for Interface	
`Interface name`	Name of the interface.
`CA name`	Name of the connectivity association. A connectivity association is named using the `connectivity-association` statement when you are enabling MACsec.

Table 1279: show security macsec connections Output Fields *(Continued)*

Field Name	Field Description
`Cipher suite`	Name of the cipher suite used for encryption.
`Encryption`	Encryption setting. Encryption is enabled when this output is `on` and disabled when this output is `off`. The encryption setting is set using the `no-encryption` statement in the connectivity association when using static connectivity association key (CAK) security mode and is set using the `encryption` statement in the secure channel when using static secure association key (SAK) or dynamic security mode.
`Key server offset`	The offset value in a packet from which encryption can be performed. The offset is set using the `offset` statement when configuring the connectivity association when using static connectivity association key (CAK) or dynamic security mode or the secure channel when using static secure association key (SAK) security mode.
`Include SCI`	SCI tagging. The SCI tag is included on packets in a secure channel when this output is `yes`, and not included on packets in a secure channel when this output is `no`. SCI tagging is automatically enabled on MX Series routers. By default, include SCI tag is disabled. You can enable SCI tagging using the `include-sci` statement in the connectivity association configuration.
`Replay protect`	By default, replay protection is disabled. Replay protection ensures that a snooped packet is not replayed or a packet number is reused. Replay protection is enabled when this output is `on` and disabled when this output is `off`. You can enable replay protection using the `replay-protect` statement in the connectivity association configuration.

Table 1279: show security macsec connections Output Fields *(Continued)*

Field Name	Field Description
`Replay window`	Number of packets that can be replayed. Must be configured with replay protection. This output is set to `0` when replay protection is disabled, and is the size of the replay window, in number of packets, when replay protection is enabled. The size of the replay window is configured using the `replay-window-size` statement in the connectivity association configuration.

Sample Output

show security macsec connections

```
user@host> show security macsec connections
Interface name: xe-0/1/0
        CA name: CA1
        Cipher suite: GCM-AES-128   Encryption: on
        Key server offset: 0        Include SCI: no
        Replay protect: off         Replay window: 0
```

show security macsec connections (MX480 routers with MPC7E-10G)

```
user@host> show security macsec connections
 Interface name: xe-4/0/18
        CA name: ca1
        Cipher suite: GCM-AES-128   Encryption: on
        Key server offset: 30       Include SCI: no
        Replay protect: off         Replay window: 0
          Outbound secure channels
            SC Id: 54:1E:56:B4:0D:3A/1
            Outgoing packet number: 11
            Secure associations
            AN: 1 Status: inuse Create time: 1d 17:31:10
          Inbound secure channels
            SC Id: 54:1E:56:B3:CA:A7/1
```

```
            Secure associations
            AN: 1 Status: inuse Create time: 1d 17:31:10
```

show security macsec connections (MX480 routers with MPC7E-10G)

```
user@host> show security macsec connections interface xe-1/0/7
CA name: caae1
        Cipher suite: AES_GCM_128   Encryption: off
        Key server offset: 0        Include SCI: no
        Replay protect: off         Replay window: 0
          Outbound secure channels
            SC Id: 54:1E:56:B3:CA:9C/1
            Outgoing packet number: 1
            Secure associations
            AN: 0 Status: inuse Create time: 4d 05:56:06
          Inbound secure channels
            SC Id: 54:1E:56:B4:0D:2F/1
            Secure associations
            AN: 0 Status: inuse Create time: 4d 05:56:06
```

Release Information

Command introduced in Junos OS Release 15.1.

Support for MPC7E-10G introduced in Junos OS Release 16.1R1 for MX240, MX480, and MX960 routers.

show security macsec statistics

IN THIS SECTION

- Syntax | 24820
- Description | 24820
- Options | 24820
- Required Privilege Level | 24820

- Output Fields | 24821
- Sample Output | 24824
- Release Information | 24825

Syntax

```
show security macsec statistics
<brief | detail>
<interface interface-name>
```

Description

Display Media Access Control Security (MACsec) statistics.

This command does not display output when MACsec is enabled using static secure association key (SAK) security mode.

Options

none Display MACsec statistics in brief form for all interfaces on the switch.

brief | detail (Optional) Display the specified level of output. Using the `brief` option is equivalent to entering the command with no options (the default). The `detail` option displays additional fields that are not visible in the `brief` output.

> **NOTE**: The field names that only appear in this command output when you enter the `detail` option are mostly useful for debugging purposes by Juniper Networks support personnel.

interface interface-name (Optional) Display MACsec statistics for the specified interface only.

Required Privilege Level

view

Output Fields

Table 1280 on page 24821 lists the output fields for the `show security macsec statistics` command. Output fields are listed in the approximate order in which they appear.

The field names that appear in this command output only when you enter the `detail` option are mostly useful for debugging purposes by Juniper Networks support personnel. Those field names are, therefore, not included in this table.

Table 1280: show security macsec statistics Output Fields

Field Name	Field Description	Level of Output
`Interface name`	Name of the interface.	All levels
Fields for Secure Channel transmitted		
`Encrypted packets`	Total number of packets transmitted out of the interface in the secure channel that were secured and encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).	All levels
`Encrypted bytes`	Total number of bytes transmitted out of the interface in the secure channel that were secured and encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).	All levels
`Protected packets`	Total number of packets transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).	All levels

Table 1280: show security macsec statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Protected bytes`	Total number of bytes transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).	All levels
Fields for Secure Association transmitted		
`Encrypted packets`	Total number of packets transmitted out of the interface in the connectivity association that were secured and encrypted using MACsec. The total includes the data packets transmitted in the secure channel and secured using a SAK and the control packets secured using a connectivity association key (CAK).	All levels
`Protected packets`	Total number of packets transmitted out of the interface in the connectivity association that were secured but not encrypted using MACsec. The total includes the data packets transmitted in the secure channel and secured using a SAK and the control packets secured using a connectivity association key (CAK).	All levels
Fields for Secure Channel received		
`Accepted packets`	The number of received packets that have been accepted by the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link. A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check. This counter increments for traffic that is and is not encrypted using MACsec.	All levels

Table 1280: show security macsec statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Validated bytes`	The number of bytes that have been validated by the MACsec integrity check and received on the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link. This counter increments for traffic when MACsec encryption is disabled.	All levels
`Decrypted bytes`	The number of bytes received in the secure channel on the interface that have been decrypted. The secure channel is used to send all data plane traffic on a MACsec-enabled link. An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.	All levels
Fields for Secure Association received		
`Accepted packets`	The number of received packets that have been accepted in the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface. A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check.	All levels
`Validated bytes`	The number of bytes that have been validated by the MACsec integrity check and received on the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface. This counter increments for traffic when MACsec encryption is disabled.	All levels

Table 1280: show security macsec statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Decrypted bytes`	The number of bytes received in the connectivity association on the interface that have been decrypted. The counter includes all control and data plane traffic accepted on the interface. An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.	All levels

Sample Output

show security macsec statistics interface xe-0/1/0 detail

```
user@host> show security macsec statistics interface xe-0/1/0 detail

  Interface name: xe-0/1/0
    Secure Channel transmitted
        Encrypted packets: 123858
        Encrypted bytes:   32190903
        Protected packets: 0
        Protected bytes:   0
    Secure Association transmitted
        Encrypted packets: 123858
        Protected packets: 0
    Secure Channel received
        Accepted packets:  123877
        Validated bytes:   0
        Decrypted bytes:   32196238
    Secure Association received
        Accepted packets:  123877
        Validated bytes:   0
        Decrypted bytes:   32196238
    Error and debug
    Secure Channel transmitted packets
        Untagged: 0, Too long: 0
    Secure Channel received packets
```

```
        Control: 0, Tagged miss: 3202804
        Untagged hit: 0, Untagged: 0
        No tag: 0, Bad tag: 0
        Unknown SCI: 0, No SCI: 0
        Control pass: 0, Control drop: 0
        Uncontrol pass: 123877, Uncontrol drop: 0
        Hit dropped: 0, Invalid accept: 0
        Late drop: 0, Delayed accept: 0
        Unchecked: 0, Not valid drop: 0
        Not using SA drop: 0, Unused SA accept: 0
```

Release Information

Command introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

show security macsec connections | 24812

show security macsec statistics (SRX Series Firewalls)

IN THIS SECTION

- Syntax | 24826
- Description | 24826
- Options | 24826
- Required Privilege Level | 24826
- Output Fields | 24826
- Sample Output | 24830
- Release Information | 24830

Syntax

```
show security macsec statistics
<brief | detail>
<interface interface-name>
```

Description

Display Media Access Control Security (MACsec) statistics.

Options

none Display MACsec statistics in brief form for all interfaces on the switch.

brief | detail (Optional) Display the specified level of output. Using the `brief` option is equivalent to entering the command with no options (the default). The `detail` option displays additional fields that are not visible in the `brief` output.

> **NOTE**: The field names that only appear in this command output when you enter the `detail` option are mostly useful for debugging purposes by Juniper Networks support personnel.

interface interface-name (Optional) Display MACsec statistics for the specified interface only.

Required Privilege Level

view

Output Fields

Table 1281 on page 24827 lists the output fields for the `show security macsec statistics` command. Output fields are listed in the approximate order in which they appear.

The field names that appear in this command output only when you enter the `detail` option are mostly useful for debugging purposes by Juniper Networks support personnel. Those field names are, therefore, not included in this table.

Table 1281: show security macsec statistics Output Fields

Field Name	Field Description	Level of Output
`Interface name`	Name of the interface.	All levels
Fields for Secure Channel transmitted		
`Encrypted packets`	Total number of packets transmitted out of the interface in the secure channel that were secured and encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a connectivity association key (CAK).	All levels
`Encrypted bytes`	Total number of bytes transmitted out of the interface in the secure channel that were secured and encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a connectivity association key (CAK).	All levels
`Protected packets`	Total number of packets transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a connectivity association key (CAK).	All levels
`Protected bytes`	Total number of bytes transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a connectivity association key (CAK).	All levels
Fields for Secure Association transmitted		

Table 1281: show security macsec statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Encrypted packets`	Total number of packets transmitted out of the interface in the connectivity association that were secured and encrypted using MACsec. The total includes the data packets transmitted in the secure channel and the control packets secured using a connectivity association key (CAK).	All levels
`Protected packets`	Total number of packets transmitted out of the interface in the connectivity association that were secured but not encrypted using MACsec. The total includes the data packets transmitted in the secure channel and the control packets secured using a connectivity association key (CAK).	All levels
Fields for Secure Channel received		
`Accepted packets`	The number of received packets that have been accepted by the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link. A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check. This counter increments for traffic that is and is not encrypted using MACsec.	All levels
`Validated bytes`	The number of bytes that have been validated by the MACsec integrity check and received on the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link. This counter does not increment when MACsec encryption is disabled.	All levels

Table 1281: show security macsec statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Decrypted bytes`	The number of bytes received in the secure channel on the interface that have been decrypted. The secure channel is used to send all data plane traffic on a MACsec-enabled link. An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.	All levels
Fields for Secure Association received		
`Accepted packets`	The number of received packets that have been accepted in the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface. A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check.	All levels
`Validated bytes`	The number of bytes that have been validated by the MACsec integrity check and received on the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface. This counter does not increment when MACsec encryption is disabled.	All levels
`Decrypted bytes`	The number of bytes received in the connectivity association on the interface that have been decrypted. The counter includes all control and data plane traffic accepted on the interface. An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.	All levels

Sample Output

show security macsec statistics interface

```
user@host> show security macsec statistics interface fxp1 detail

  Interface name: fxp1
    Secure Channel transmitted
        Encrypted packets: 2397305
        Encrypted bytes:   129922480
        Protected packets: 0
        Protected bytes:   0
    Secure Association transmitted
        Encrypted packets: 2397305
        Protected packets: 0
    Secure Channel received
        Accepted packets:  2395850
        Validated bytes:   0
        Decrypted bytes:   131715088
    Secure Association received
        Accepted packets:  2395850
        Validated bytes:   0
        Decrypted bytes:   0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

- *Understanding Media Access Control Security (MACsec)*
- *Configuring Media Access Control Security (MACsec)*
- macsec | **10278**
- show interfaces (Gigabit Ethernet)
- show security mka sessions (SRX Series Device | **24849**

show security match-policies

IN THIS SECTION

- Syntax | 24831
- Description | 24832
- Options | 24832
- Required Privilege Level | 24833
- Output Fields | 24834
- Sample Output | 24836
- Release Information | 24839

Syntax

```
show security match-policies
destination-ip <ip-address>
destination-port < port-number>
destination-vrf <destination-vrf>
from-zone <zone-name>
global
logical-system <logical-system-name>
protocol <protocol-name | protocol-number>
result-count <number>
root-logical-system
source-end-user-profile <device-identity-profile-name>
source-identity <role-name>
source-ip <ip-address>
source-port <port-number>
tenant <tenant-name>
to-zone <zone-name>
url category<url-category>
```

Description

The `show security match-policies` command allows you to troubleshoot traffic problems using the match criteria: source port, destination port, source IP address, destination IP address, and protocol. For example, if your traffic is not passing because either an appropriate policy is not configured or the match criteria is incorrect, then the `show security match-policies` command allows you to work offline and identify where the problem actually exists. It uses the search engine to identify the problem and thus enables you to use the appropriate match policy for the traffic.

The `result-count` option specifies how many policies to display. The first enabled policy in the list is the policy that is applied to all matching traffic. Other policies below it are "shadowed" by the first and are never encountered by matching traffic.

> **NOTE**: The `show security match-policies` command is applicable only to security policies; IDP policies are not supported.

Options

- `destination-ip` *`destination-ip`*—Displays the destination IP address of the traffic.
- `destination-port` *`destination-port`*–Displays the destination port number of the traffic. Range is 1 through 65,535.
- `destination-vrf` *`destination-vrf`*—(Optional) Displays the destination VRF information.
- `from-zone` *`zone-name`*—Displays the name or ID of the source zone of the traffic. This is optional for global policies.
- `global`—Displays information about global policies.
- `logical-system`—Displays the logical system name.
- `protocol` *`protocol-name`* | *`protocol-number`*–Displays the protocol name or numeric value of the traffic.
  - `ah` or `51`
  - `egp` or `8`
  - `esp` or `50`
  - `gre` or `47`
  - `icmp` or `1`
  - `igmp` or `2`

  - `igp` or `9`
  - `ipip` or `94`
  - `ipv6` or `41`
  - `ospf` or `89`
  - `pgm` or `113`
  - `pim` or `103`
  - `rdp` or `27`
  - `rsvp` or `46`
  - `sctp` or `132`
  - `tcp` or `6`
  - `udp` or `17`
  - `vrrp` or `112`
- `result-count` *`number`*—(Optional) Displays the number of policy matches. Valid range is from 1 through 16. The default value is 1.
- `root-logical-system`—Displays root logical system as default.
- `source-end-user-profile` *`device-identity-profile-name`*—(Optional) Displays the device identity profile that specifies characteristics that can apply to one or more devices.
- `source-identity` *`role-name`*—(Optional) Displays the source identity of the traffic determined by the user role.
- `source-ip` *`source-ip`*—Displays the source IP address of the traffic.
- `source-port` *`source-port`*—Displays the source port number of the traffic. Range is 1 through 65,535.
- `tenant`—Displays the name of the tenant system.
- `to-zone` *`zone-name`*—Displays the name or ID of the destination zone of the traffic. This is optional for global policies.

Required Privilege Level

view

Output Fields

Table 1282 on page 24834 lists the output fields for the `show security match-policies` command. Output fields are listed in the approximate order in which they appear.

Table 1282: show security match-policies Output Fields

Field Name	Field Description
`Policy`	Name of the applicable policy.
`Action` or `Action-type`	The action to be taken for traffic that matches the policy's match criteria. Actions include the following: • `permit` • `firewall-authentication` • `tunnel ipsec-vpn` *`vpn-name`* • `pair-policy` *`pair-policy-name`* • `deny` • `reject`
`State`	Status of the policy: • `enabled`: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • `disabled`: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
`Index`	An internal number associated with the policy.
`Sequence number`	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, and 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, and 4.
`From zone`	Name of the source zone.

Table 1282: show security match-policies Output Fields *(Continued)*

Field Name	Field Description
`To zone`	Name of the destination zone.
`Source addresses`	The names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
`Destination addresses`	The names and corresponding IP addresses of the destination addresses (or address sets) for a policy as entered in the destination zone's address book. A packet's destination address must match one of these addresses for the policy to apply to it.
`Application`	Name of a preconfigured or custom application, or `any` if no application is specified.
`IP protocol`	Numeric value for the IP protocol used by the application, such as 6 for TCP or 1 for ICMP.
`ALG`	If an ALG is associated with the session, the name of the ALG. Otherwise, 0.
`Inactivity timeout`	Elapsed time without activity after which the application is terminated.
`Source-port range`	Range of matching source ports defined in the policy.
`Destination-port range`	Range of matching destination ports defined in the policy.
`Source identities`	One or more user roles defined in the matching policy.
`global`	Display information about global policies.
`device-identity-profile-name`	Device identity profile that specifies characteristics that can apply to one or more devices.

Sample Output

Example 1: show security match-policies

```
user@host> show security match-policies from-zone z1 to-zone z2 source-ip 10.10.10.1 destination-
ip 192.0.2.1 source-port 1 destination-port 21 protocol tcp
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: z1, To zone: z2
  Source addresses:
    a2: 198.51.100.0/24
    a3: 10.10.10.1/32
  Destination addresses:
    d2: 203.0.113.0/24
    d3: 192.0.2.1/32
  Application: junos-ftp
    IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
            Source port range: [0-0]
            Destination port range: [21-21]
```

Example 2: show security match policies ... result-count

```
user@host> show security match-policies from-zone zone-A to-zone zone-B source-ip 10.10.10.1
destination-ip 192.0.2.5 source_port 1004 destination_port 80 protocol tcp result_count 5
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: zone-A, To zone: zone-B
  Source addresses:
    sa1: 10.10.0.0/16
  Destination addresses:
    da5: 192.0.2.0/24
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
            Source port range: [1000-1030]
          Destination port range: [80-80]

Policy: p15, action-type: deny, State: enabled, Index: 18
  Sequence number: 15
  From zone: zone-A, To zone: zone-B
  Source addresses:
```

```
    sa11: 10.10.10.1/32
  Destination addresses:
    da15: 192.0.2.5/32
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
            Source port range: [1000-1030]
          Destination port range: [80-80]
```

Example 3: show security match policies ... source-identity

```
user@host> show security match-policies from-zone untrust to-zone trust  source-ip 10.10.10.1
destination-ip 192.0.2.1 destination_port 21 protocol 6 source-port 1234 source-identity role1
Policy: p1, action-type: permit, State: enabled, Index: 40
      Policy Type: Configured
      Sequence number: 1
      From zone: untrust, To zone: trust
      Source addresses:
        a1: 10.0.0.0/8
      Destination addresses:
        d1: 192.0.2.0/24
      Application: junos-ftp
        IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
          Source port range: [0-0]
          Destination port range: [21-21]
      Source identities: role1
      Per policy TCP Options: SYN check: No, SEQ check: No
```

Example 4: show security match policies ... global

```
user@host> show security match-policies global source-ip 10.10.10.1 destination-ip 192.0.2.5
source_port 1004 destination_port 80 protocol tcp result_count 5
Policy: gp1, action-type: permit, State: enabled, Index: 6, Scope Policy: 0
  Policy Type: Configured, global
  Sequence number: 1
  From zones:
    Any
  To zones:
    Any
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
```

```
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security match-policies tenant TN1 from-zone trust to-zone untrust source-ip 10.10.10.1 destination-ip 192.0.2.1 source-port 1 destination-port 21 protocol tcp

```
user@host> show security match-policies tenant TN1 from-zone trust to-zone untrust source-ip
10.10.10.1 destination-ip 192.0.2.1 source-port 1 destination-port 21 protocol tcp

Policy: p1, action-type: permit, State: enabled, Index: 4
Sequence number: 1
From zone: trust, To zone: untrust
Source addresses:
a2: 198.51.100.0/24
a3: 10.10.10.1/32
Destination addresses:
d2: 203.0.113.0/24
d3: 192.0.2.1/32
Application: junos-ftp
IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
Source port range: [0-0]
Destination port range: [21-21]
```

show security match-policies from-zone client to-zone svr source-ip 10.1.1.1 source-port 88 destination-ip 10.2.2.2 destination-port 80 protocol tcp url-category Enhanced_Games

```
user@host> show security match-policies from-zone client to-zone svr source-ip 10.1.1.1 source-
port 88 destination-ip 10.2.2.2 destination-port 80 protocol tcp url-category Enhanced_Games

Policy: p1, action-type: permit, State: enabled, Index: 7
0
  Policy Type: Configured
```

```
Sequence number: 1
From zone: client, To zone: server
Source vrf group:
  any
Destination vrf group:
  any
Source addresses:
  any-ipv4(global): 0.0.0.0/0
  any-ipv6(global): ::/0
Destination addresses:
  any-ipv4(global): 0.0.0.0/0
  any-ipv6(global): ::/0
Application: any
  IP protocol: 0, ALG: 0, Inactivity timeout: 0
    Source port range: [0-0]
    Destination ports: [0-0]
Url-category:
  Enhanced_Sex: 234881056
  Enhanced_Games: 234881037
Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
Intrusion Detection and Prevention: disabled
Unified Access Control: disabled
Unified Threat Management: enabled
```

Release Information

Command introduced in Junos OS Release 10.3.

Command updated in Junos OS Release 10.4.

Command updated in Junos OS Release 12.1.

Command updated to include optional from-zone and to-zone global match options in Junos OS Release 12.1X47-D10.

The `tenant` option is introduced in Junos OS Release 18.3R1.

The `url category` option is introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

clear security policies statistics | 4037

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

show security mka sessions

IN THIS SECTION

- Syntax | 24840
- Description | 24840
- Options | 24840
- Required Privilege Level | 24841
- Output Fields | 24841
- Sample Output | 24843
- Release Information | 24843

Syntax

```
show security mka sessions
<interface interface-name>
```

Description

Display MACsec Key Agreement (MKA) session information for all interfaces. The MKA protocol is responsible for maintaining MACsec on the link, and decides which router on the point-to-point link becomes the key server.

Options

- `interface` *`interface-name`*—Display the MKA session information for the specified interface only.
- `summary | brief | detail`—Display the specified level of output.
- none (same as `brief`)—Display the MKA session information for all interfaces.

Required Privilege Level

view

Output Fields

Table 1283 on page 24841 lists the output fields for the `show security mka sessions` command. Output fields are listed in the approximate order in which they appear.

Table 1283: show security mka sessions Output Fields

Field Name	Field Description
`Interface name`	Name of the interface.
`Interface state`	Shows whether the interface is secured or not. If it is secured, the CAK type is also displayed.
`Member identifier`	Name of the member identifier.
`CAK name`	Name of the connectivity association key (CAK). The CAK is configured using the `cak` keyword when configuring the pre-shared key.
`CAK type`	The CAK type: primary, fallback, or preceding.
`Transmit interval`	The transmit interval. Both ends of the point-to-point link should be configured to the same value. Default value is 2000 seconds. Possible values: 2000 through 6000 milliseconds.
`Outbound SCI`	Name of the outbound secure channel identifier.
`Message number`	Number of the last data message.
`Key number`	Key number.

Table 1283: show security mka sessions Output Fields *(Continued)*

Field Name	Field Description
`Key server`	Key server status. The router is the key server when this output is `yes`. The router is not the key server when this output is `no`.
`Key server priority`	Displays the priority of the key server. Lower value indicates higher priority. Use the `key-server-priority` statement to set the priority. Possible values: 0 through 255.
`Latest SAK AN`	Name of the latest secure association key (SAK) association number.
`Latest SAK KI`	Name of the latest secure association key (SAK) key identifier.
Fields for Peer list	
`Member identifier`	Name of the member identifier.
`Hold time`	Hold time, in seconds.
`Message number`	Number of the last data message
`SCI`	Name of the secure channel identifier.
`Lowest acceptable PN`	Number of the lowest acceptable packet number (PN).
Fields for CAK list (detail only)	
`CAK name`	Name of the connectivity association key (CAK).
`CAK type`	The CAK type: primary, fallback, or preceding.
`Status`	The CAK status: live, active, or in-progress.

Table 1283: show security mka sessions Output Fields *(Continued)*

Field Name	Field Description
`Member identifier`	Name of the member identifier.
`Message number`	Number of the last data message

Sample Output

show security mka sessions

```
user@host> show security mka sessions
Member identifier: ABC09234C234245345
CAK Name: EF00132234324ABCDE2342352345DC
Send period : 2000 (ms)
Key server priority: 16

Message number: 132     Outbound SCI: 01:01:02:02:03:04/1968
Key Server: Yes        Key Server priority: 16
Latest SAK AN : 2        Latest SAK KI: ABC09090EFAA1212
Previous SAK AN: 1    Pervious SAK KI: CEE090A07FAA3223

Peer list
1. MI: ABC09234C234245345 (Live/Potential)     MN: 2345
   SCI: 01:02:02:02:04:04/1990       Hold time: 6 sec
   Lowest Acceptable PN: 243235
2. MI: ACC0926C334245341 (Potential)         MN: 2784
   SCI: 04:02:02:02:05:04/1340       Hold time: 6 sec
   Lowest Acceptable PN: 645236
```

Release Information

Command introduced in Junos OS Release 13.2X50-D15.

show security mka sessions (MX Series)

IN THIS SECTION

- Syntax | 24844
- Description | 24844
- Options | 24844
- Required Privilege Level | 24844
- Output Fields | 24845
- Sample Output | 24847
- Release Information | 24849

Syntax

```
show security mka sessions
<interface interface-name>
<summary | brief | detail>
```

Description

Display MACsec Key Agreement (MKA) session information for all interfaces. The MKA protocol is responsible for maintaining MACsec on the link, and decides which router on the point-to-point link becomes the key server.

Options

- `interface` *`interface-name`*—Display the MKA session information for the specified interface only.
- `summary | brief | detail`—Display the specified level of output.
- none (same as `brief`)—Display the MKA session information for all interfaces.

Required Privilege Level

view

Output Fields

Table 1284 on page 24845 lists the output fields for the `show security mka sessions` command. Output fields are listed in the approximate order in which they appear.

Table 1284: show security mka sessions Output Fields

Field Name	Field Description
`Interface name`	Name of the interface.
`Interface state`	State of the interface: • Secured • Secured-Suspended (during GRES) • Unsecured If the interface is in secured or secured-suspended state, the CAK type is also displayed.
`Member identifier`	Name of the member identifier.
`CAK name`	Name of the connectivity association key (CAK). The CAK is configured using the `cak` keyword when configuring the pre-shared key.
`CAK type`	The CAK type: primary, fallback, or preceding.
`MKA suspended`	The number of seconds the MACsec session can be suspended during GRES. This count decrements until the remote node comes out of suspension.
`Transmit interval`	The transmit interval. Both ends of the point-to-point link should be configured to the same value. Default value is 2000 seconds. Possible values: 2000 through 6000 milliseconds.
`SAK rekey interval`	The timer-based refresh interval for the secure association key (SAK). Default value is 0 seconds. Possible values: 60 through 86,400 seconds.

Table 1284: show security mka sessions Output Fields *(Continued)*

Field Name	Field Description
`Preceding key`	Shows whether preceding key is enabled or not.
`Bounded delay`	Shows whether bounded delay is enabled or not.
`Outbound SCI`	Name of the outbound secure channel identifier.
`Message number`	Number of the last data message.
`Key number`	Key number.
`Key server`	Key server status. The router is the key server when this output is `yes`. The router is not the key server when this output is `no`.
`Key server priority`	Displays the priority of the key server. Lower value indicates higher priority. Use the `key-server-priority` statement to set the priority. Possible values: 0 through 255.
`Latest SAK AN`	Name of the latest secure association key (SAK) association number.
`Latest SAK KI`	Name of the latest secure association key (SAK) key identifier.
`MKA Suspend For`	Shows whether MKA session suspensions are enabled or disabled. Configure the `suspend-for` statement to enable suspensions during GRES.
`MKA Suspend On Request`	Shows whether the key server is enabled to accept MKA session suspension requests from the peer server. Configure the `suspend-on-request` statement to enable the key server to accept suspension requests for GRES.
Fields for Peer list	

Table 1284: show security mka sessions Output Fields *(Continued)*

Field Name	Field Description
`Member identifier`	Name of the member identifier.
`Hold time`	Hold time, in seconds.
`Message number`	Number of the last data message
`SCI`	Name of the secure channel identifier.
`Lowest acceptable PN`	Number of the lowest acceptable packet number (PN).
Fields for CAK list (detail only)	
`CAK name`	Name of the connectivity association key (CAK).
`CAK type`	The CAK type: primary, fallback, or preceding.
`Status`	The CAK status: live, active, or in-progress.
`Member identifier`	Name of the member identifier.
`Message number`	Number of the last data message

Sample Output

show security mka sessions

```
user@host> show security mka sessions

  Interface name: xe-0/2/0
     Interface State: Secured-Suspended
     Member identifier: 9D8976C83B8DCB101430AC8B
```

```
     CAK name: 3333
     CAK type: primary
     MKA suspended: 99(s)
     Transmit interval: 2000(ms)
     SAK rekey interval: 0(s)
     Preceding Key: enabled
     Outbound SCI: 88:E0:F3:1F:40:64/1
     Message number: 236        Key number: 3
     Key server: yes            Key server priority: 16
     Latest SAK AN: 2           Latest SAK KI: 9D8976C83B8DCB101430AC8B/3
     Previous SAK AN: 1         Previous SAK KI: 9D8976C83B8DCB101430AC8B/2
     MKA Suspend For: enabled   MKA Suspend On Request: enabled
```

show security mka sessions interface ge-0/0/2 detail

```
user@host> show security mka sessions interface ge-0/0/2 detail
     Interface state: Secured - Primary
     Member identifier: 20E8DB2EA6A09291E497BA41
     CAK name: 1111
     CAK type: primary
     MKA suspended: 0(s)
     Transmit interval: 6000(ms)
     SAK rekey interval: 0(s)
     Preceding Key: enabled
     Bounded Delay: disabled
     Outbound SCI: 30:B6:4F:6A:C8:02/1
     Message number: 42265      Key number: 1
     Key server: yes            Key server priority: 16
     Latest SAK AN: 1           Latest SAK KI: 20E8DB2EA6A09291E497BA41/1
     Previous SAK AN: 0         Previous SAK KI: 00000000000000000000000/0
     MKA Suspend For: enabled   MKA Suspend On Request: enabled
     CAK list: (2)
        1. CAK name: 1111
           CAK type: primary                            Status: live
           Member identifier: 20E8DB2EA6A09291E497BA41   Message number: 42265
           Peer list: (1)
              1. Member identifier: 053303E64A2B6207EF54CF12 (live)
                 Message number: 51887       Hold time: 13000 (ms)
                 SCI: B0:A8:6E:A1:B8:5A/1
                 Lowest acceptable PN: 0
        2. CAK name: FFF1
```

```
            CAK type: fallback                            Status: active
            Member identifier: 15CCDF84E92E90FFA1541A87   Message number: 42261
            Peer list: (1)
               1. Member identifier: 24007B2D31E69974AD3E8416 (live)
                  Message number: 51870       Hold time: 16000 (ms)
                  SCI: B0:A8:6E:A1:B8:5A/1
                  Lowest acceptable PN: 0
```

Release Information

Command introduced in Junos OS Release 15.1.

show security mka sessions (SRX Series Device

IN THIS SECTION

- Syntax | 24849
- Description | 24849
- Options | 24850
- Required Privilege Level | 24850
- Output Fields | 24850
- Sample Output | 24851
- Release Information | 24851

Syntax

```
show security mka sessions
<interface interface-name>
```

Description

Display MACsec Key Agreement (MKA) session information.

Options

- `interface` *`interface-name`*—(Optional) Display the MKA information for the specified interface only.

Required Privilege Level

view

Output Fields

Table 1285 on page 24850 lists the output fields for the `show security mka sessions` command. Output fields are listed in the approximate order in which they appear.

Table 1285: show security mka sessions Output Fields

Field Name	Field Description
`Interface name`	Name of the interface.
`Member identifier`	Name of the member identifier.
`CAK name`	Name of the Connectivity Association Key (CAK.. The CAK is configured using the `cak` keyword when configuring the pre-shared key.
`Transmit interval`	The transmit interval.
`Outbound SCI`	Name of the outbound secure channel identifier.
`Message number`	Number of the last data message.
`Key server`	Key server status. The switch is the key server when this output is `yes`. The switch is not the key server when this output is `no`.
`Key number`	Key number.

Table 1285: show security mka sessions Output Fields *(Continued)*

Field Name	Field Description
`Key server priority`	The key server priority. The key server priority can be set using the `key-server-priority` statement.
`Latest SAK AN`	Name of the latest secure association key (SAK) association number.
`Latest SAK KI`	Name of the latest secure association key (SAK) key identifier.
`Previous SAK AN`	Name of the previous secure association key (SAK) association number.
`Previous SAK KI`	Name of the previous secure association key (SAK) key identifier.

Sample Output

show security mka sessions

```
user@host> show security mka sessions

    Interface name: fxp1
      Member identifier: 71235CA1B78D0AF7B3F29CFB
      CAK name: AAAA
      Transmit interval: 10000(ms)
      Outbound SCI: 30:7C:5E:44:98:42/1
      Message number: 2326       Key number: 2
      Key server: yes            Key server priority: 16
      Latest SAK AN: 1           Latest SAK KI: 71235CA1B78D0AF7B3F29CFB/2
      Previous SAK AN: 0         Previous SAK KI: 71235CA1B78D0AF7B3F29CFB/1
```

Release Information

Command introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | 10278

show interfaces (Gigabit Ethernet)

show security macsec statistics (SRX Series Firewalls) | 24825

show security mka sessions summary

IN THIS SECTION

- Syntax | 24852
- Description | 24852
- Required Privilege Level | 24852
- Output Fields | 24853
- Sample Output | 24853
- Release Information | 24854

Syntax

```
show security mka sessions summary
```

Description

Display MACsec Key Agreement (MKA) session information to see the number of MKAs that are in progress, connectivity association key (CAK) type, CAK status, and MKA packet count activity.

Required Privilege Level

view

Output Fields

Table 1286 on page 24853 lists the output fields for the `show security mka sessions summary` command. Output fields are listed in the approximate order in which they appear.

Table 1286: show security mka sessions summary Output Fields

Field Name	Field Description
`Interface name`	Name of the interface.
`Member ID`	Member identifier.
`Type`	The CAK type: primary, fallback, or preceding.
`Status`	The CAK status: live, active, or in-progress.
`Tx`	MKA packets transmitted on the interface.
`Rx`	MKA packets received on the interface.
`CAK name`	Name of the connectivity association key (CAK). The CAK is configured using the `cak` keyword when configuring the pre-shared key.

Sample Output

show security mka sessions summary

```
user@host> show security mka sessions summary

Interface  Member-ID                      Type            Status            Tx        Rx
CAK Name
ge-0/0/1   E752CAEAE8DDFB82D4EA4BF7  preceding    live                8887    8951    8888
ge-0/0/1   0F2D5171F38EAB16C2E0CB62  fallback    active            8959    8952    FFFF
ge-0/0/1   6B49BD5CF7188F3CD9A29D30  primary        in-progress        2439    0        AAAA
```

Release Information

Command introduced in Junos OS Release 19.2.

show security mka statistics

IN THIS SECTION

- Syntax | 24854
- Description | 24854
- Options | 24854
- Required Privilege Level | 24855
- Output Fields | 24855
- Sample Output | 24856
- Release Information | 24857

Syntax

```
show security mka statistics
<interface interface-name>
```

Description

Display MACsec Key Agreement (MKA) protocol statistics.

The output for this command does not include statistics for MACsec data traffic. For MACsec data traffic statistics, see "show security macsec statistics" on page 24819.

Options

- `interface` *`interface-name`*—(Optional) Display the MKA information for the specified interface only.

Required Privilege Level

view

Output Fields

Table 1287 on page 24855 lists the output fields for the `show security mka statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1287: show security mka statistics Output Fields

Field Name	Field Description
`Received packets`	Number of received MKA control packets. This counter increments for received MKA control packets only. This counter does not increment when data packets are received.
`Transmitted packets`	Number of transmitted MKA packets This counter increments for transmitted MKA control packets only. This counter does not increment when data packets are transmitted.
`Version mismatch packets`	Number of version mismatch packets.
`CAK mismatch packets`	Number of Connectivity Association Key (CAK) mismatch packets. This counter increments when the connectivity association key (CAK) and connectivity association key name (CKN), which are user-configured values that have to match to enable MACsec, do not match for an MKA control packet.
`ICV mismatch packets`	Number of ICV mismatched packets. This counter increments when the connectivity association key (CAK) value does not match on both ends of a MACsec-secured Ethernet link.
`Duplicate message identifier packets`	Number of duplicate message identifier packets.

Table 1287: show security mka statistics Output Fields *(Continued)*

Field Name	Field Description
`Duplicate message number packets`	Number of duplicate message number packets.
`Duplicate address packets`	Number of duplicate source MAC address packets.
`Invalid destination address packets`	Number of invalid destination MAC address packets.
`Formatting error packets`	Number of formatting error packets.
`Old Replayed message number packets`	Number of old replayed message number packets.

Sample Output

show security mka statistics

```
user@host> show security mka statistics

        Received packets:                     1525844
        Transmitted packets:                  1525841
        Version mismatch packets:             0
        CAK mismatch packets:                 0
        ICV mismatch packets:                 0
        Duplicate message identifier packets: 0
        Duplicate message number packets:     0
        Duplicate address packets:            0
        Invalid destination address packets:  0
        Formatting error packets:             0
        Old Replayed message number packets:  0
```

Release Information

Command introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

show security mka sessions | 24840
show security macsec statistics | 24819
show security macsec connections | 24812

show security mka statistics

IN THIS SECTION

- Syntax | 24857
- Description | 24857
- Options | 24858
- Required Privilege Level | 24858
- Output Fields | 24858
- Sample Output | 24859
- Release Information | 24860

Syntax

```
show security mka statistics
<interface interface-name>
```

Description

Display MACsec Key Agreement (MKA) protocol statistics.

The output for this command does not include statistics for MACsec data traffic. For MACsec data traffic statistics, see "show security macsec statistics (SRX Series Devices)" on page 24825.

Options

- `interface` *`interface-name`*—(Optional) Display the MKA information for the specified interface only.

Required Privilege Level

view

Output Fields

Table 1288 on page 24858 lists the output fields for the `show security mka statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1288: show security mka statistics Output Fields

Field Name	Field Description
`Received packets`	Number of received MKA control packets. This counter increments for received MKA control packets only. This counter does not increment when data packets are received.
`Transmitted packets`	Number of transmitted MKA packets This counter increments for transmitted MKA control packets only. This counter does not increment when data packets are transmitted.
`Version mismatch packets`	Number of version mismatch packets.
`CAK mismatch packets`	Number of Connectivity Association Key (CAK) mismatch packets. This counter increments when the connectivity association key (CAK) and connectivity association key name (CKN), which are user-configured values that have to match to enable MACsec, do not match for an MKA control packet.
`ICV mismatch packets`	Number of ICV mismatched packets. This counter increments when the connectivity association key (CAK) value does not match on both ends of a MACsec-secured Ethernet link.

Table 1288: show security mka statistics Output Fields *(Continued)*

Field Name	Field Description
Duplicate message identifier packets	Number of duplicate message identifier packets.
Duplicate message number packets	Number of duplicate message number packets.
Duplicate address packets	Number of duplicate source MAC address packets.
Invalid destination address packets	Number of invalid destination MAC address packets.
Formatting error packets	Number of formatting error packets.
Old Replayed message number packets	Number of old replayed message number packets.

Sample Output

show security mka statistics

```
user@host> show security mka statistics

        Interface name: fxp1
        Received packets:                    3
        Transmitted packets:                 14
        Version mismatch packets:            0
        CAK mismatch packets:                0
        ICV mismatch packets:                0
        Duplicate message identifier packets: 0
        Duplicate message number packets:    0
        Duplicate address packets:           0
        Invalid destination address packets:  0
```

```
        Formatting error packets:             0
        Old Replayed message number packets:  0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

- *Understanding Media Access Control Security (MACsec)*
- *Configuring Media Access Control Security (MACsec)*
- macsec | **10278**
- show interfaces (Gigabit Ethernet)
- show security macsec statistics (SRX Series Firewalls) | **24825**

show security mka statistics (MX Series)

IN THIS SECTION

- Syntax | **24860**
- Description | **24861**
- Options | **24861**
- Required Privilege Level | **24861**
- Output Fields | **24861**
- Sample Output | **24862**
- Release Information | **24864**

Syntax

```
show security mka statistics
<interface interface-name>
```

Description

Display MACsec Key Agreement (MKA) protocol statistics.

The output for this command does not include statistics for MACsec data traffic. For MACsec data traffic statistics, see "show security macsec statistics" on page 24819.

Options

- `interface` *`interface-name`*—(Optional) Display the MKA information for the specified interface only.
- none—Display the MKA information for all interfaces.

Required Privilege Level

view

Output Fields

Table 1289 on page 24861 lists the output fields for the `show security mka statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1289: show security mka statistics Output Fields

Field Name	Field Description
`Received packets`	Number of received MKA control packets. This counter increments for received MKA control packets only. This counter does not increment when data packets are received.
`Transmitted packets`	Number of transmitted MKA packets This counter increments for transmitted MKA control packets only. This counter does not increment when data packets are transmitted.
`Version mismatch packets`	Number of version mismatch packets.

Table 1289: show security mka statistics Output Fields *(Continued)*

Field Name	Field Description
`CAK mismatch packets`	Number of Connectivity Association Key (CAK) mismatch packets. This counter increments when the connectivity association key (CAK) and connectivity association key name (CKN), which are user-configured values that have to match to enable MACsec, do not match for an MKA control packet.
`ICV mismatch packets`	Number of ICV mismatched packets. This counter increments when the connectivity association key (CAK) value does not match on both ends of a MACsec-secured Ethernet link.
`Duplicate message identifier packets`	Number of duplicate message identifier packets.
`Duplicate message number packets`	Number of duplicate message number packets.
`Duplicate address packets`	Number of duplicate source MAC address packets.
`Invalid destination address packets`	Number of invalid destination MAC address packets.
`Formatting error packets`	Number of formatting error packets.
`Old Replayed message number packets`	Number of old replayed message number packets.

Sample Output

show security mka statistics

```
user@host> show security mka statistics
```

```
        Received packets:                      1525844
        Transmitted packets:                   1525841
        Version mismatch packets:              0
        CAK mismatch packets:                  0
        ICV mismatch packets:                  0
        Duplicate message identifier packets: 0
        Duplicate message number packets:      0
        Duplicate address packets:             0
        Invalid destination address packets:   0
        Formatting error packets:              0
        Old Replayed message number packets:   0
```

show security mka statistics (MX480 routers with MPC7E-10G)

```
user@host> show security mka statistics
 Interface name: xe-4/0/18
        Received packets:                      73009
        Transmitted packets:                   73011
        Version mismatch packets:              0
        CAK mismatch packets:                  1
        ICV mismatch packets:                  0
        Duplicate message identifier packets: 0
        Duplicate message number packets:      0
        Duplicate address packets:             0
        Invalid destination address packets:   0
        Formatting error packets:              0
        Old Replayed message number packets:   0
```

show security mka statistics (MX480 routers with MPC7E-10G)

```
user@host> show security mka statistics interface xe-1/0/7
Received packets:                      179211
        Transmitted packets:                   179186
        Version mismatch packets:              0
        CAK mismatch packets:                  0
        ICV mismatch packets:                  0
        Duplicate message identifier packets: 0
        Duplicate message number packets:      0
        Duplicate address packets:             0
        Invalid destination address packets:   0
```

```
        Formatting error packets:             0
        Old Replayed message number packets:  0
```

Release Information

Command introduced in Junos OS Release 15.1.

Support for MPC7E-10G introduced in Junos OS Release 16.1R1 for MX240, MX480, and MX960 routers.

show security monitoring

IN THIS SECTION

- Syntax | 24864
- Description | 24864
- Required Privilege Level | 24865
- show security monitoring | 24865
- show security monitoring (SRX1400) | 24865
- show security monitoring (vSRX Virtual Firewall) | 24866
- show security monitoring (vSRX Virtual Firewall in a Chassis Cluster) | 24866
- Release Information | 24866

Syntax

```
show security monitoring
```

Description

Displays a count of security flow and central point (CP) sessions, CPU utilization (as a percentage of maximum), and memory in use (also as a percentage of maximum) at the moment the command is run. This command is supported on SRX1400, SRX1500, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall.

Required Privilege Level

View

show security monitoring

command-name

user@host>**show security monitoring**

```
user@host> show security monitoring
-----------------------------------------------------------------------------
                  Flow session   Flow session     CP session     CP session
FPC PIC CPU Mem        current        maximum        current        maximum
-----------------------------------------------------------------------------
  1   0   0  11              0              0              0              0
  1   1   0   5              3        6291456              1        7549747
  1   2   0   5              2        6291456              0        7549747
  1   3   0   5              3        6291456              1        7549747
  8   0   0  65              4           6963              2           8355
  8   1   0  65              2           6963              0           8355
Total Sessions:             14       18888294              4       22665951
```

show security monitoring (SRX1400)

command-name

user@host>**show security monitoring**

```
user@host> show security monitoring
-----------------------------------------------------------------------------
                  Flow session   Flow session     CP session     CP session
FPC PIC CPU Mem        current        maximum        current        maximum
-----------------------------------------------------------------------------
  1   0   0  75              0        1048576            0           1048576
```

show security monitoring (vSRX Virtual Firewall)

command-name

user@host>**show security monitoring**

```
user@host> show security monitoring
-------------------------------------------------------------------------
                  Flow session   Flow session     CP session     CP session
FPC PIC CPU Mem        current        maximum        current        maximum
-------------------------------------------------------------------------
  0   0   0  68              2         524288            N/A            N/A
```

show security monitoring (vSRX Virtual Firewall in a Chassis Cluster)

command-name

user@host>**show security monitoring**

```
node0:
-------------------------------------------------------------------------
                  Flow session   Flow session     CP session     CP session
FPC PIC CPU Mem        current        maximum        current        maximum
-------------------------------------------------------------------------
  0   0   0  67              0         524288            N/A            N/A

node1:
-------------------------------------------------------------------------
                  Flow session   Flow session     CP session     CP session
FPC PIC CPU Mem        current        maximum        current        maximum
-------------------------------------------------------------------------
  0   0   0  67              0         524288            N/A            N/A
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show security monitoring fpc fpc-number | 24867
show security monitoring performance session | 24871
show security monitoring performance spu | 24873

show security monitoring fpc fpc-number

IN THIS SECTION

- Syntax | 24867
- Description | 24867
- Options | 24868
- Additional Information | 24868
- Required Privilege Level | 24868
- Output Fields | 24868
- Sample Output | 24869
- Sample Output | 24870
- Sample Output | 24870
- Release Information | 24871

Syntax

```
show security monitoring fpc fpc-number
<node (node-id | all | local | primary)>
```

Description

Display security monitoring information about the FPC slot.

Options

- *fpc-number* —Display security monitoring information for the specified FPC slot. It can be in the range from 0 to 11.
- node—(Optional) For chassis cluster configurations, display security monitoring information for the specified FPC on a specific node (device) in the cluster.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - all—Display information about all nodes.
  - local—Display information about the local node.
  - primary—Display information about the primary node.

Additional Information

For complete list of slot numbering, physical port, and logical interface numbering for SRX Series Firewalls in chassis cluster, see Chassis Cluster User Guide for SRX Series Devices.

Required Privilege Level

view

Output Fields

Table 1290 on page 24868 lists the output fields for the show security monitoring fpc *fpc-number* command. Output fields are listed in the approximate order in which they appear.

Table 1290: show security monitoring fpc fpc-number Output Fields

Field Name	Field Description
FPC	Slot number in which the FPC is installed.
PIC	Slot number in which the PIC is installed.
CPU Utilization (%)	Total percentage of CPU being used by the PIC's processors.

Table 1290: show security monitoring fpc fpc-number Output Fields *(Continued)*

Field Name	Field Description
`Memory Utilization (%)`	Percentage of heap space (dynamic memory) being used by the PIC's processor. If this number exceeds 80 percent, there might be a software problem (memory leak).
`Current flow session`	The current number of flow sessions. When SRX Series Firewalls operate in packet mode, flow sessions will not be created and this field will remain zero.
`Max flow session`	The maximum number of flow sessions allowed. This number will differ from one device to another.
`SPU current cp session`	The current number of cp sessions for the SPU (on SRX5600, and SRX5800 devices only).
`SPU max cp session`	The maximum number of cp sessions allowed for the SPU (on SRX5600, and SRX5800 devices only).

Sample Output

show security monitoring fpc 0

```
user@host> show security monitoring fpc 0
FPC 0
  PIC 0
    CPU utilization      :    0 %
    Memory utilization   :   82 %
    Current flow session :    0
    Max flow session     :    0
    Current CP session   :    0
    Max CP session       : 12000000
Session Creation Per Second (for last 96 seconds on average):    0
  PIC 1
    CPU utilization      :    0 %
    Memory utilization   :   54 %
    Current flow session :    0
    Max flow session     : 819200
```

```
    Current CP session   :    0
    Max CP session       :    0
Session Creation Per Second (for last 96 seconds on average):    0
```

Sample Output

show security monitoring fpc 1

```
user@host> show security monitoring fpc 1
FPC 1
  PIC 0
    CPU utilization      :    0 %
    Memory utilization   :   21 %
    Current flow session :    0
    Max flow session     : 524288
    Current CP session   :    0
    Max CP session       : 1048576
Session Creation Per Second (for last 96 seconds on average):    0
```

Sample Output

show security monitoring fpc 8

```
user@host> show security monitoring fpc 5
FPC 5
  PIC 0
    CPU utilization      :    0 %
    Memory utilization   :   64 %
    Current flow session :    0
    Max flow session     : 524288
    Current CP session   :    0
    Max CP session       : 2359296
Session Creation Per Second (for last 96 seconds on average):    0
  PIC 1
    CPU utilization      :    0 %
    Memory utilization   :   65 %
    Current flow session :    0
    Max flow session     : 1048576
    Current CP session   :    0
```

```
    Max CP session       :    0
Session Creation Per Second (for last 96 seconds on average):    0
```

Release Information

Command introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

show services ip-monitoring status | 25604

show security monitoring performance session

IN THIS SECTION

- Syntax | 24871
- Description | 24871
- Options | 24872
- Required Privilege Level | 24872
- show security monitoring performance session | 24872
- Release Information | 24872

Syntax

```
show security monitoring performance session
<fpc slot-number>
<pic slot-number>
```

Description

Display the current session (total number of sessions at that time) for the last 60 seconds.

Options

- `fpc slot-number` — Display information about the FPC slot. Use this option to filter the output based on the slot number.
- `pic slot-number` — Display information about existing PIMs or Mini-PIMs in a particular PIC slot. Use this option to filter the output based on PIC slot.

> **NOTE**: The fpc slot-number and pic slot-number options are not available on SRX300, SRX320, and SRX340 devices.

Required Privilege Level

View

show security monitoring performance session

command-name

```
user@host> show security monitoring performance session
```

```
fpc  0  pic  0
Last 60 seconds:
 0:       8   1:       8   2:       8   3:       8   4:       8   5:       7
 6:       7   7:       7   8:       7   9:       7  10:       7  11:       8
12:       8  13:       8  14:       7  15:       7  16:       7  17:       7
18:       7  19:       7  20:       7  21:       5  22:       5  23:       5
24:       5  25:       5  26:       5  27:       5  28:       5  29:       4
30:       4  31:       4  32:       3  33:       3  34:       3  35:       3
36:       5  37:       5  38:       6  39:       6  40:       5  41:       5
42:       5  43:       5  44:       5  45:       5  46:       5  47:       5
48:       7  49:       7  50:       6  51:       8  52:       8  53:       6
54:       5  55:       7  56:       7  57:       5  58:       5  59:       8
```

Release Information

Command introduced in Junos OS Release of 10.2.

RELATED DOCUMENTATION

show services ip-monitoring status | 25604

show security monitoring performance spu

IN THIS SECTION

- Syntax | 24873
- Description | 24873
- Options | 24874
- Required Privilege Level | 24874
- Sample Output | 24874
- Release Information | 24877

Syntax

```
show security monitoring performance spu (extensive | summary)
<fpc fpc>
<pic pic>
<thread thread>
```

Description

Display the services processing unit (SPU) percent utilization for all FPC slots over the last 60 seconds. Use this command to track the percent utilization statistics per second for the past 60 seconds for each FPC slot and PIC.

Starting in Junos OS Release 20.2R1, you can use the following operational commands to monitor the average CPU usage information for the last minute, hour, or day of an SPC3 card:

- `show security monitoring performance spu summary fpc fpc pic pic`
- `show security monitoring performance spu summary fpc fpc pic pic thread thread`

You can monitor the CPU usage information only when the PIC is online.

Options

- `fpc` *`fpc`* — Display information about the FPC slot. Use this option to filter the output based on the FPC slot number.
- `pic` *`pic`* — Display information about existing PIMs or Mini-PIMs in a particular PIC slot. Use this option to filter the output based on PIC slot number.

> **NOTE**: The *fpc* and *pic* options are not available on SRX300, SRX320, SRX340, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, and SRX4600 devices, nor on vSRX Virtual Firewall instances.

- `thread` *`thread`* — Display information about the SPU based on the thread number.
- `extensive | summary` — Display the specified level of output.
- None — Display the SPU percent utilization for all FPC slots over the last 60 seconds.

Required Privilege Level

View

Sample Output

show security monitoring performance spu

This sample shows 46% utilization of the SPU for second 42 in the past 60 seconds for FPC 0 and PIC 0.

`user@host>` **show security monitoring performance spu**

```
fpc  0  pic  0
Last 60 seconds:
 0:  48   1:  48   2:  48   3:  48   4:  48   5:  48
 6:  48   7:  48   8:  49   9:  48  10:  48  11:  48
12:  48  13:  48  14:  48  15:  48  16:  48  17:  48
18:  48  19:  48  20:  48  21:  48  22:  49  23:  48
24:  49  25:  49  26:  48  27:  48  28:  48  29:  48
30:  48  31:  48  32:  48  33:  48  34:  48  35:  48
36:  46  37:  47  38:  46  39:  46  40:  46  41:  46
42:  46  43:  46  44:  46  45:  46  46:  46  47:  46
```

```
48:  46  49:  46  50:  46  51:  46  52:  46  53:  46
54:  46  55:  46  56:  46  57:  46  58:  46  59:  46
```

show security monitoring performance spu (vSRX Virtual Firewall and vSRX Virtual Firewall 3.0)

`user@host>` **show security monitoring performance spu**

```
fpc  0  pic  0
Last 60 seconds:
  0:    0    1:   0    2:    0    3:    0    4:    0    5:    0
  6:    0    7:   0    8:    0    9:    0   10:    0   11:    0
 12:    0   13:   0   14:    0   15:    0   16:    0   17:    0
 18:    0   19:   0   20:    0   21:    0   22:    0   23:    0
 24:    0   25:   0   26:    0   27:    0   28:    0   29:    0
 30:    0   31:   0   32:    0   33:    0   34:    0   35:    0
 36:    0   37:   0   38:    0   39:    0   40:    0   41:    0
 42:    0   43:   0   44:    0   45:    0   46:    0   47:    0
 48:    0   49:   0   50:    0   51:    0   52:    0   53:    0
 54:    0   55:   0   56:    0   57:    0   58:    0   59:    0
```

show security monitoring performance spu summary fpc 1 pic 0 (SRX5400, SRX5600, and SRX5800)

`user@host>` **show security monitoring performance spu summary fpc 1 pic 0**

```
FPC 1 PIC 0THREAD 0
Last Minute Average %: 0
Last Hour Average %: 0
Last Day Average %: 0

THREAD 1
Last Minute Average %: 0
Last Hour Average %: 0
Last Day Average %: 0

THREAD 2
Last Minute Average %: 0
```

```
Last Hour Average %: 0
Last Day Average %: 0

... ...

THREAD 23
Last Minute Average %: 0
Last Hour Average %: 0
Last Day Average %: 0

THREAD 24
Last Minute Average %: 0
Last Hour Average %: 0
Last Day Average %: 0

THREAD 25
Last Minute Average %: 0
Last Hour Average %: 0
Last Day Average %: 0
```

show security monitoring performance spu summary fpc 1 pic 1 thread 8 (SRX5400, SRX5600, and SRX5800)

user@host> **show security monitoring performance spu summary fpc 1 pic 1 thread 8**

```
FPC 1 PIC 1
THREAD 8
Last Minute Average %: 0
Last Hour Average %: 5
Last Day Average %: 4
```

show security monitoring performance spu extensive (SRX4200)

user@host> **show security monitoring performance spu extensive**

```
         CPU   Util   Wutil   Status   SchedCounter
FPC0
    PIC0
            0      0       0    alive         574289
```

```
          1      0       0    alive         574289
          2      0       0    alive         574289
          3      0       0    alive         574289
          4      0       0    alive         574289
          5      0       0    alive         574289
          6      0       0    alive         574289
          7      0       0    alive         574289
          8      0       0    alive         574289
          9      0       0    alive         574289
         10      0       0    alive         574289
         11      0       0    alive         574289
         12      0       0    alive         574289
         13      0       0    alive         574289
         14      0       0    alive         574289
         15      0       0    alive         574289
         16      0       0    alive         574289
         17      0       0    alive         574289
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

show services ip-monitoring status | 25604

show security nat destination pool

IN THIS SECTION

- Syntax | 24878
- Description | 24878
- Options | 24878
- Required Privilege Level | 24878
- Output Fields | 24879

- Sample Output | 24879
- Sample Output | 24880
- Release Information | 24881

Syntax

```
show security nat destination pool
pool-name
all
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display information about the specified Network Address Translation (NAT) destination address pool.

Options

pool-name	Name of the destination address pool.
all	Display information about all the destination NAT address pools.
logical-system (*logical-system-name*)	Display information about the destination NAT pools for a specified logical system. Specify `all` to display information for all logical systems.
root-logical-system	Display information about the destination NAT pools for the primary (root) logical system.
tenant (*tenant-name*)	Display information about the destination NAT pools for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 1291 on page 24879 lists the output fields for the **show security nat destination pool** command. Output fields are listed in the approximate order in which they appear.

Table 1291: show security nat destination pool Output Fields

Field Name	Field Description
Pool name	Name of the destination pool.
Description	Description of the destination pool.
Pool id	Pool identification number.
Routing instance	Name of the routing instance.
Total address	Number of IP addresses that are in use.
Translation hits	Number of translation hits.
Address range	IP address or IP address range for the pool.

Sample Output

show security nat destination pool dst-nat-pool1

```
user@host> show security nat destination pool dst-p1

Pool name       : dst-p1
Description     : The destination pool dst-p1 is for the sales team
Pool id         : 1
Routing instance: default
Total address   : 1
Translation hits: 0
Address range                                Port
     203.0.113.1 -203.0.113.1               0
```

Sample Output

show security nat destination pool all

```
user@host> show security nat destination pool all

Total destination-nat pools: 2

Pool name       : dst_v4_pool-namelength-is-now-changed-upto-63-characters-length
Description     : The destination pool dst-p1 is for the sales team
Pool id         : 1
Routing instance: default
Total address   : 1
Translation hits: 0
Address range                                Port
     203.0.113.1 -203.0.113.1               0

Pool name       : dst_v4_pool-namelength-is-now-changed-upto-63-characters-length
Description     : The destination pool dst-p2 is for the sales team
Pool id         : 2
Routing instance: default
Total address   : 1
Translation hits: 0
Address range                                Port
     2001:db8::1 - 2001:db8::1               0
```

show security nat destination pool all tenant

```
user@host> show security nat destination pool all tenant tn1

Total destination-nat pools: 1
Pool name       : dst_v4_pool-namelength-is-now-changed-upto-63-characters-length
Pool id         : 1
Total address   : 1
Translation hits: 0
Address range                        Port
  192.168.1.200 - 192.168.1.200         0
```

Release Information

Command introduced in Junos OS Release 9.2.

The `Description` output field added in Junos OS Release 12.1.

Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

pool (Security Destination NAT) | 12779

show security nat destination rule

IN THIS SECTION

- Syntax | 24881
- Description | 24882
- Options | 24882
- Required Privilege Level | 24882
- Output Fields | 24882
- Sample Output | 24884
- Sample Output | 24885
- Release Information | 24886

Syntax

```
show security nat destination rule
rule-name
all
logical-system (logical-system-name)
```

```
root-logical-system
tenant (tenant-name)
```

Description

Display information about the specified destination Network Address Translation (NAT) rule. Destination NAT rules are processed after static NAT rules but before source NAT rules.

Options

rule-name	Display information about the specified destination NAT rule.
all	Display information about all the destination NAT rules.
logical-system	Display information about the destination NAT rules for a specified logical system. Specify `all` to display information for all logical systems.
root-logical-system	Display information about the destination NAT rules for the primary (root) logical system.
tenant	Display information about the destination NAT rules for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 1292 on page 24882 lists the output fields for the **show security nat destination rule** command. Output fields are listed in the approximate order in which they appear.

Table 1292: show security nat destination rule Output Fields

Field Name	Field Description
Total destination-nat rules	Number of destination NAT rules.

Table 1292: show security nat destination rule Output Fields *(Continued)*

Field Name	Field Description
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly as address names and as address set names in the rule.
Destination NAT rule	Name of the destination NAT rule.
Description	Description of the destination NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the destination NAT rule.
From routing instance	Name of the routing instance from which the packets flow.
From interface	Name of the interface from which the packets flow.
From zone	Name of the zone from which the packets flow.
Source addresses	Name of the source addresses that match the rule. The default value is any.
Destination addresses	Name of the destination addresses that match the rule. The default value is any.
Action	The action taken when a packet matches the rule's tuples. Actions include the following: • **destination NAT pool**—Use user-defined destination NAT pool to perform destination NAT. • **off**—Do not perform destination NAT.
Destination ports	Destination ports number that match the rule. The default value is any.

Table 1292: show security nat destination rule Output Fields *(Continued)*

Field Name	Field Description
Application	Indicates whether the application option is configured.
Translation hits	Number of translation hits.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat destination rule namelength-is-now-changed-upto-63-characterslength

```
user@host>show security nat destination rule namelength-is-now-changed-upto-63-characterslength

 Destination NAT rule: namelength-is-now-changed-upto-63-characterslength           Rule-set:
namelength-is-now-changed-upto-63-characterslength
  Description                : The destination rule dst2-rule is for the sales team
  Rule-Id                    : 1
  Rule position              : 1
  From routing instance      : ri1
                             : ri2
  Match
    Source addresses         : add1
                               add2
    Destination addresses    : add9
  Action                     : off

  Destination port           : 0
 Translation hits           : 68
   Successful sessions      : 25
   Failed sessions          : 43
```

```
 Number of sessions         : 2
```

Sample Output

show security nat destination rule all

```
user@host> show security nat destination rule all

Total destination-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0

Destination NAT rule: r4                   Rule-set: rs4
  Rule-Id                    : 2
  Rule position              : 2
  From zone                  : untrust
  Match
    Source addresses         : 192.0.2.0 - 192.0.2.255
    Destination addresses    : 198.51.100.0 - 198.51.100.255
    Application              : configured
  Action                     : off
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

show security nat destination rule all tenant

```
user@host> show security nat destination rule all tenant tn1

Total destination-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
Destination NAT rule: r1                     Rule-set: from_zone
  Rule-Id                    : 1
  Rule position              : 1
  From zone                  : untrust
  Match
    Source addresses         : 192.0.2.0       - 192.0.2.255
```

```
    Destination addresses    : 203.0.113.202   - 203.0.113.202
  Action                     : h1
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

Release Information

Command introduced in Junos OS Release 9.2. The `Description` output field added in Junos OS Release 12.1.

Support for IPv6 logical systems and the `Successful sessions`, `Failed sessions` and `Number of sessions` output fields added in Junos OS Release 12.1X45-D10.

Output for multiple destination ports and the `application` option field added in Junos OS Release 12.1X47-D10.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

rule (Security Destination NAT) | **16105**

show security nat destination rule-application

IN THIS SECTION

- Syntax | **24887**
- Description | **24887**
- Options | **24887**
- Required Privilege Level | **24887**
- Output Fields | **24887**
- Sample Output | **24888**
- Sample Output | **24889**
- Release Information | **24889**

Syntax

```
show security nat destination rule-application
rule-name
all
logical-system logical-system-name
root-logical-system
```

Description

Display information about the specified destination Network Address Translation (NAT) rule application.

Options

rule-name	Display information about the specified destination NAT rule application.
all	Display information about all the destination NAT rule applications.
logical-system *logical-system-name*	Display information about the destination NAT rule applications for the specified logical system.
root-logical-system	Display information about the destination NAT rule applications for the primary (root) logical system.

Required Privilege Level

view

Output Fields

Table 1293 on page 24887 lists the output fields for the **show security nat destination rule-application** command. Output fields are listed in the approximate order in which they appear.

Table 1293: show security nat destination rule-application Output Fields

Field Name	Field Description
Destination NAT rule	Name of the destination NAT rule.

Table 1293: show security nat destination rule-application Output Fields *(Continued)*

Field Name	Field Description
Rule-set	Rule set identification number.
Rule-Id	Rule identification number.
Application	Name of the application or application set.
IP protocol	IP protocol identifier.
Source port range	Source port range identifier.
Destination port range	Destination port identifier.
ICMP information	• type—ICMP message type. • code—Code corresponding to the ICMP message type.

Sample Output

show security nat destination rule-application for port application

```
user@host>show security nat destination rule-application all

 Destination NAT rule: r4                   Rule-set: rs4
  Rule-Id                    : 2
    Application: app-set1
      IP protocol: 17
      Source port range: [40-50]
      Destination port range: [50-60]
      IP protocol: 17
      Source port range: [100-200]
      Destination port range: [300-500]
```

Sample Output

show security nat destination rule-application for ICMP application

```
user@host>show security nat destination rule-application all

Destination NAT rule: r1                    Rule-set: rs1
  Rule-Id                    : 1
    Application: junos-icmp-all
      IP protocol: icmp
      ICMP Information: type=255, code=0
    Application: icmp1
      IP protocol: icmp
      ICMP Information: type=1, code=1
    Application: junos-icmp6-all
      IP protocol: 58
      ICMP Information: type=255, code=0
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

- Logical Systems and Tenant Systems User Guide for Security Devices

show security nat destination summary

IN THIS SECTION

- Syntax | 24890
- Description | 24890
- Options | 24890
- Required Privilege Level | 24890

- Output Fields | 24890
- Sample Output | 24891
- Release Information | 24892

Syntax

```
show security nat destination summary
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display a summary of Network Address Translation (NAT) destination pool information.

Options

none	Display summary information about the destination NAT pool.
logical-system	Display summary information about the destination NAT for a specified logical system. Specify `all` to display information for all logical systems.
root-logical-system	Display summary information about the destination NAT for the primary (root) logical system.
tenant	Display information about the destination NAT for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 1294 on page 24891 lists the output fields for the **show security nat destination summary** command. Output fields are listed in the approximate order in which they appear.

Table 1294: show security nat destination summary Output Fields

Field Name	Field Description
Total destination nat pool number	Number of destination NAT pools.
Pool name	Name of the destination address pool.
Address range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.
Port	Port number.
Total	Number of IP addresses that are in use.
Available	Number of IP addresses that are free for use.
Total destination nat rule number	Number of destination NAT rules.
Total hit times	Number of times a translation in the translation table is used for all the destination NAT rules.
Total fail times	Number of times a translation in the translation table failed to translate for all the destination NAT rules.

Sample Output

show security nat destination summary

```
user@host> show security nat destination summary
Total pools: 1
Pool name            Address                          Routing        Port  Total
                     Range                            Instance             Address
```

```
dst_v4_pool-namelength-is-now-changed-upto-63-characters-length 192.0.2.2-192.0.2.2 0 1

Total rules: 1
Rule name : dstrule_namelen_is_now_changed_upto_63_characters_length
    Rule set : dst_ruleset_namelength_changed_upto_63_characters_length
    Action   : dst_v4_pool-namelength-is-now-changed-upto-63-characters-length
    From     : trust
```

show security nat destination summary tenant

```
user@host> show security nat destination summary tenant tn1

Total pools: 1
Pool name            Address                          Routing        Port  Total
                     Range                            Instance             Address
h1                   192.168.1.200  - 192.168.1.200                  0     1

Total rules: 1
Rule name            Rule set       From                              Action
r1                   from_zone      untrust                           h1
```

Release Information

Command introduced in Junos OS Release 9.2.

Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

pool (Security Destination NAT) | **12779**

rule (Security Destination NAT) | **16105**

show security nat incoming-table

IN THIS SECTION

- Syntax | 24893
- Description | 24893
- Options | 24893
- Required Privilege Level | 24894
- Output Fields | 24894
- Sample Output | 24895
- Release Information | 24895

Syntax

```
show security nat incoming-table
<node (          node-id         | all | local | primary)>
```

Description

Display Network Address Translation (NAT) table information.

NOTE: The incoming dip NAT table is replaced with ALG cone NAT binding table and the `show security nat incoming-table` command is obsolete from Junos OS Release 11.2 onward. The `show security nat incoming-table` command works as is in the previous releases.

Options

- none—Display all information NAT incoming table.
- `node`—(Optional) For chassis cluster configurations, display incoming table information on a specific node.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.

- `local`—Display information about the local node.
- `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1295 on page 24894 lists the output fields for the `show security nat incoming-table` command. Output fields are listed in the approximate order in which they appear.

Table 1295: show security nat incoming-table Output Fields

Field Name	Field Description
`In use`	Number of entries in the NAT table.
`Maximum`	Maximum number of entries possible in the NAT table.
`Entry allocation failed`	Number of entries failed for allocation.
`Destination`	Destination IP address and port number.
`Host`	Host IP address and port number that the destination IP address is mapped.
`References`	Number of sessions referencing the entry.
`Timeout`	Timeout, in seconds, of the entry in the NAT table.
`Source-pool`	Name of source pool where translation is allocated.

Sample Output

show security nat incoming-table

```
user@host> show security nat incoming-table
In use: 1, Maximum: 1024, Entry allocation failed: 0
Destination           Host                                References Timeout Source-pool
10.1.1.26:1028        203.0.113.10:5060                    1     3600     p1
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

clear security nat incoming-table | 4007

show security nat interface-nat-ports

IN THIS SECTION

- Syntax | 24896
- Description | 24896
- Options | 24896
- Required Privilege Level | 24896
- Output Fields | 24896
- Sample Output | 24897
- Sample Output | 24898
- Release Information | 24899

Syntax

```
show security nat interface-nat-ports
<node (node-id | all | local | primary)>
<logical-system (logical-system-name | all)>
```

Description

Display port usage for an interface source pool for Network Address Translation (NAT).

Options

none Display all port usage information for an interface source pool.

node For chassis cluster configurations, display interface NAT ports information on a specific node.

- ***node-id*** Identification number of the node. It can be 0 or 1.
- **all** Display information about all nodes.
- **local** Display information about the local node.
- **primary** Display information about the primary node.

logical-system (*logical-system-name* | all) Display port usage information for the specified logical system or for all logical systems.

Required Privilege Level

view

Output Fields

Table 1296 on page 24897 lists the output fields for the **show security nat interface-nat-ports** command. Output fields are listed in the approximate order in which they appear.

Table 1296: show security nat interface-nat-ports Output Fields

Field Name	Field Description
Pool Index	Port pool index.
Total Ports	Total number of ports in a port pool. In SRX Series Firewalls, 10 interface NAT ports are supported.
Single Ports Allocated	Number of ports allocated one at a time that are in use.
Single Ports Available	Number of ports allocated one at a time that are free for use.
Twin Ports Allocated	Number of ports allocated two at a time that are in use.
Twin Ports Available	Number of ports allocated two at a time that are free for use.

Sample Output

show security nat interface-nat-ports

```
user@host> show security nat interface-nat-ports
Pool   Total  Single ports  Single ports  Twin ports  Twin ports
    index  ports     allocated     available   allocated   available
     0  64510             0         63486           0        1024
     1  64510             0         63486           0        1024
     2  64510             0         63486           0        1024
     3  64510             0         63486           0        1024
     4  64510             0         63486           0        1024
     5  64510             0         63486           0        1024
     6  64510             0         63486           0        1024
     7  64510             0         63486           0        1024
     8  64510             0         63486           0        1024
     9  64510             0         63486           0        1024
```

Sample Output

show security nat interface-nat-ports logical-system all

```
user@host> show security nat interface-nat-ports logical-system all
Logical system: root-logical-system
Pool   Total  Single ports  Single ports  Twin ports  Twin ports
index  ports     allocated     available   allocated   available
    0  64510             0         63486           0        1024
Logical system: LSYS1
Pool   Total  Single ports  Single ports  Twin ports  Twin ports
index  ports     allocated     available   allocated   available
    0  64510             0         63486           0        1024
    1  64510             0         63486           0        1024
    2  64510             0         63486           0        1024
    3  64510             0         63486           0        1024
    4  64510             0         63486           0        1024
    5  64510             0         63486           0        1024
    6  64510             0         63486           0        1024
    7  64510             0         63486           0        1024
    8  64510             0         63486           0        1024
    9  64510             0         63486           0        1024
   10  64510             0         63486           0        1024
   11  64510             0         63486           0        1024
   12  64510             0         63486           0        1024
   13  64510             0         63486           0        1024
   14  64510             0         63486           0        1024
   15  64510             0         63486           0        1024
   16  64510             0         63486           0        1024
   17  64510             0         63486           0        1024
   18  64510             0         63486           0        1024
   19  64510             0         63486           0        1024
   20  64510             0         63486           0        1024
   21  64510             0         63486           0        1024
   22  64510             0         63486           0        1024
   23  64510             0         63486           0        1024
   24  64510             0         63486           0        1024
   25  64510             0         63486           0        1024
   26  64510             0         63486           0        1024
   27  64510             0         63486           0        1024
   28  64510             0         63486           0        1024
   29  64510             0         63486           0        1024
```

```
    30  64510          0        63486          0        1024
    31  64510          0        63486          0        1024
    32  64510          0        63486          0        1024
    33  64510          0        63486          0        1024
    34  64510          0        63486          0        1024
    35  64510          0        63486          0        1024
    36  64510          0        63486          0        1024
    37  64510          0        63486          0        1024
    38  64510          0        63486          0        1024
    39  64510          0        63486          0        1024
    40  64510          0        63486          0        1024
    41  64510          0        63486          0        1024
    42  64510          0        63486          0        1024
    43  64510          0        63486          0        1024
    44  64510          0        63486          0        1024
    45  64510          0        63486          0        1024
    45  64510          0        63486          0        1024
    46  64510          0        63486          0        1024
    47  64510          0        63486          0        1024
    48  64510          0        63486          0        1024
    49  64510          0        63486          0        1024
    50  64510          0        63486          0        1024
    51  64510          0        63486          0        1024
    52  64510          0        63486          0        1024
    53  64510          0        63486          0        1024
    54  64510          0        63486          0        1024
    55  64510          0        63486          0        1024
    56  64510          0        63486          0        1024
    57  64510          0        63486          0        1024
    58  64510          0        63486          0        1024
    59  64510          0        63486          0        1024
```

Release Information

Command modified in Junos OS Release 9.2. The **node** options added in Junos OS Release 9.0. Logical system support added in Junos OS Release 12.1X45-D10.

show security nat resource-usage source-pool

IN THIS SECTION

- Syntax | 24900
- Description | 24900
- Options | 24900
- Required Privilege Level | 24901
- Output Fields | 24901
- Sample Output | 24902
- Release Information | 24903

Syntax

```
show security nat resource-usage source-pool
all | source-pool-name
logical-system logical-system-name | root logical system
```

Description

Display source NAT pool usage information. In pools without Port Address Translation (PAT), information about IP addresses is displayed. In pools with PAT, information about ports is displayed.

Options

all	Display resource use information for all source NAT pools.
source-pool-name	Display resource use information for the specified source NAT pool.
logical-system *logical-system-name*	Display resource use information for the source NAT pools in the specified logical system.
root-logical-system	Display resource use information for the source NAT pools in the root logical system.

Required Privilege Level

view

Output Fields

Table 1297 on page 24901 lists the output fields for the **show security nat resource-usage source-pool** command. Output fields are listed in the approximate order in which they appear. You can use the **clear security nat statistics** command to reset the peak usage statistics.

Table 1297: show security nat resource-usage source-pool Output Fields

Field Name	Field Description
Pool	Name of the pool.
Address	Address of the pool.
Factor-index	Port pool index.
Total address	Number of addresses in the pool. Starting in Junos OS Release 15.1X49-D90 and Junos OS Release 17.3R1, the total number of addresses for pools with IPv6 prefixes is shown as zero (0).
Port-range	Number of ports allocated at a time.
Used	Number of used resources in the pool.
Avail	Number of available resources in the pool. Starting in Junos OS Release 15.1X49-D90 and Junos OS Release 17.3R1, the total number of available resources for pools with IPv6 prefixes is shown as 0.
Usage	Percent of resources used. In a PAT pool, use includes single and twin ports.
Current usage	Percent of current resources used.

Table 1297: show security nat resource-usage source-pool Output Fields *(Continued)*

Field Name	Field Description
Peak usage	Percent of resources used during the peak date and time.
Total	Number of used and available resources.
Total ports	Number of used and available ports.
Port-overloading-factor	Port overloading capacity for the pool.

Sample Output

show security nat resource-usage resource-pool all

```
user@host> show security nat resource-usage source-pool all

PAT pools(including address-shared pool) port utilization:
Pool      Address          Used       Avail      Total      Usage
SpoolA        512       2387968    29593600   31981568         7%
SpoolB        128        393216      655360    1048576        38%

Non-PAT pools address utilization:
Pool         Used    Avail    Total   Usage
Spool1        300     3796     4096      7%
Spool2        512      512     1024     50%
```

show security nat resource-usage resource-pool pool-name (Without PAT)

```
user@host> show security nat resource-usage source-pool Spool1
Logical system: root
Peak usage: 60% @ 2012-08-26 20:16:20 UTC

Pool         Used     Avail     Total    Usage
```

```
Spool1        300      3796      4096       7%
```

show security nat resource-usage resource-pool pool-name (with PAT)

```
user@host> show security nat resource-usage source-pool sp3
Logical system: root
Pool name: sp3
Total address: 2
Port-overloading-factor: 2
Total ports: 258048 Used: 60563  Avail: 197485
Current usage: 23%   Peak usage: 35% at 2012-11-12 20:15:26 CST
  Address    Factor-index Port-range          Used       Avail       Total Usage
  192.0.2.113
             0            Single Ports       30001       32463       62464   48%
             -            Alg Ports            462        1586        2048   22%
             1            Single Ports           0       62464       62464    0%
             -            Alg Ports              0        2048        2048    0%
             Sum          Single Ports       30001       94927      124928   24%
             -            Alg Ports            462        3634        4096   11%
  192.0.2.114
             0            Single Ports       29600       32864       62464   47%
             -            Alg Ports            500        1548        2048   24%
             1            Single Ports           0       62464       62464    0%
             -            Alg Ports              0        2048        2048    0%
             Sum          Single Ports       29600       95328      124928   23%
             -            Alg Ports            500        3596        4096   12%
```

Release Information

Command introduced in Junos OS Release 12.1X45-D10.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D90	Starting in Junos OS Release 15.1X49-D90 and Junos OS Release 17.3R1, the total number of addresses for pools with IPv6 prefixes is shown as zero (0).

15.1X49-D90	Starting in Junos OS Release 15.1X49-D90 and Junos OS Release 17.3R1, the total number of available resources for pools with IPv6 prefixes is shown as 0.

RELATED DOCUMENTATION

clear security nat statistics source pool | 4013

show security nat source deterministic

IN THIS SECTION

- Syntax | 24904
- Description | 24905
- Options | 24905
- Required Privilege Level | 24905
- Output Fields | 24905
- Sample Output | 24906
- Release Information | 24907

Syntax

```
show security nat source deterministic
pool-name
host-ip host ip address
host-address-range
xlated-ip xlated-ip-address
xlated-port xlated-port
node
root-logical-system | logical-system {lsys-name | all }
```

Description

Verify the mapping relation when Deterministic-Nat is on.

Options

host-address-range	Display deterministic host address range without overlap.
pool-name	Display Deterministic NAT port block table for the specified source pool name.
node	Display source NAT deterministic port block table on specific node.
host ip address	Display deterministic NAT port block table based on internal host ip address.
xlated ip address	Display deterministic NAT port block table based on translated IP address.
xlated-port	Display deterministic NAT port block table based on translated IP and port; *xlated-port* can be used only with *xlated-ip* together for display.
root-logical-system	Display information about the source NAT pools for the primary (root) logical system.
logical-system (*lsys-name* | all)	Display information about the specified logical system source NAT pools or all logical system source NAT pools.

Required Privilege Level

view

Output Fields

Table 1298 on page 24905 lists the output fields for the `show security nat source deterministic` command. Output fields are listed in the approximate order in which they appear.

Table 1298: show security nat source deterministic Output Fields

Field Name	Field Description
Pool name	Name of pool.

Table 1298: show security nat source deterministic Output Fields *(Continued)*

Field Name	Field Description
Port-overloading-factor	Factor of port overloading for the source pool.
Port block size	Number of ports that a port block contains.
Used/total port blocks	Port block used number and port block total number for this source NAT pool.
Host IP	IP address of host.
External IP	IP address of external router.
Port_Range	The range of ports in a block, ranging from lowest to highest.
Ports_Used/Total	Number of ports used and total ports.

Sample Output

show security nat source deterministic

```
user@host> show security nat source deterministic
Pool name: source_pool1_name_length_can_be_configured_upto_63_chars_length
Port-overloading-factor:     1     Port block size: 10000
Used/total port blocks: 0/12
Host_IP                      External_IP            Port_Range     Ports_Used/Total
192.0.2.25                     203.0.113.1            1024-11023            0/10000*1
192.0.2.26                     203.0.113.1            11024-21023           0/10000*1
```

show services nat source deterministic (MX-SPC3)

```
user@host> show services nat source deterministic
Pool name: NAT_POOL1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-POOL
Port-overloading-factor:     1     Port block size:    1
Used/total port blocks: 0/128
```

```
Host_IP                        External_IP                        Port_Block
Ports_Used/
                                                                    Range
Ports_Total
192.0.2.27                      203.0.113.1                         3035-3035
0/1*1
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

show security nat source pool | 24914

show security nat source port-block | 24927

show security nat source paired-address

IN THIS SECTION

- Syntax | 24908
- Description | 24908
- Options | 24908
- Additional Information | 24908
- Required Privilege Level | 24908
- Output Fields | 24908
- Sample Output | 24909
- Release Information | 24910

Syntax

```
show security nat source paired-address
<internal-ip internal-ip-address>
<logical-system logical-system-name>
<pool-name pool-name>
<root-logical-system>
<xlated-ip x-lated-ip-address>
```

Description

Display information about the Network Address Translation (NAT) source paired addresses.

Options

none	Display all paired IP address information.
internal ip *internal-ip-address*	Display information about the specified internal IP address.
logical-system *logical-system-name*	Display information about the source NAT pools for the specified logical system.
pool-name *pool-name*	Display paired address information for the specified pool.
root-logical-system	Display information about the source NAT pools for the primary (root) logical system.
x-lated-ip *x-lated-ip-address*	Display information about the specified translated external IP address.

Additional Information

Required Privilege Level

view

Output Fields

Table 1299 on page 24909 lists the output fields for the `show security nat source paired-address` command. Output fields are listed in the approximate order in which they appear.

Table 1299: show security nat source paired-address Output Fields

Field Name	Field Description
Pool name	Name of the source pool.
Internal address	Internal IP address.
External address	External IP address.

Sample Output

show security nat source paired-address

```
user@host> show security nat source paired-address
Pool name: pool-namelength-is-now-changed-upto-63-characters-lengtht
Internal address          External address
192.0.2.0                 203.0.113.105

Pool name: pool-namelength-is-now-changed-upto-63-characters-length
Internal address          External address
192.0.2.1                 203.0.113.105
192.0.2.2                 203.0.113.105
192.0.2.3                 203.0.113.105
192.0.2.4                 203.0.113.105
192.0.2.5                 203.0.113.105
192.0.2.6                 203.0.113.105

Pool name: pool-namelength-is-now-changed-upto-63-characters-length
Internal address          External address
192.0.2.9                 203.0.113.105
```

show security nat source paired-address pool-name

```
user@host> show security nat source paired-address pool-name sp1
Pool name: sp1
Internal address        External address
```

```
192.168.1.1             192.0.2.1
192.168.1.2             192.0.2.2
192.168.1.3             192.0.2.3
```

show security nat source paired-address pool-name internal-ip

```
user@host> show security nat source paired-address pool-name sp1 internal-ip 192.168.1.1
Pool name: sp1
Internal address        External address
192.168.1.1             192.0.2.1
```

show security nat source paired-address pool-name xlated-ip

```
user@host> show security nat source paired-address pool-name sp1 xlated-ip 192.0.2.2
Pool name: sp1
Internal address        External address
192.168.1.2             192.0.2.2
```

Release Information

Command introduced in Junos OS Release 12.1X45-D10.

show security nat source persistent-nat-table

IN THIS SECTION

- Syntax | 24911
- Description | 24911
- Options | 24911
- Required Privilege Level | 24911
- Output Fields | 24911

- Sample Output | 24912
- Sample Output | 24913
- Sample Output | 24913
- Release Information | 24913

Syntax

```
show security nat source persistent-nat-table ( all | interface | internal-ip ip-address
<internal-port port> | pool poolname )
```

Description

Display a summary of persistent Network Address Translation (NAT) information.

Options

- **all**—Display all persistent NAT bindings.
- **interface**—Display persistent NAT bindings for the interface.
- **internal-ip** ***ip-address***—Display persistent NAT bindings for the specified internal IP address.
- **internal-ip** ***ip-address*** **internal-port** ***port***—Display persistent NAT bindings for the specified internal IP address and port.
- **pool**—Display persistent NAT bindings for the specified source NAT pool.
- **summary**—Display persistent NAT bindings summary.

Required Privilege Level

view

Output Fields

Table 1300 on page 24912 lists the output fields for the **show security nat source persistent-nat-table** command. Output fields are listed in the approximate order in which they appear.

Table 1300: show security nat source persistent–nat–table Output Fields

Field Name	Field Description
Enode	To support restricted cone NAT scenarios, the enode entries contain remote host and port information.
Internal IP/Port	Internal transport IP address and port number of the outgoing session from internal to external.
Reflexive IP/Port	Translated IP address and port number of the source IP address and port.
Source NAT Pool	The name of the source pool where persistent NAT is used.
Type	Persistent NAT type.
Left_time/Conf_time	The inactivity timeout period that remains and the configured timeout value.
Current_Sess_Num/ Max_Sess_Num	The number of current sessions associated with the persistent NAT binding. The persistent NAT timer starts only when all sessions using that binding are closed. Even if one session is using that binding, the entry displays with "-" as Left_Time value in the output.
Source NAT Rule	Name of the source NAT rule to which this persistent NAT binding applies.

Sample Output

show security nat source persistent–nat–table internal-ip internal-port

```
user@host> show security nat source persistent-nat-table internal-ip 192.0.2.1 internal-port
60784

Internal                        Reflective        Source     Type         Left_time/
Curr_Sess_Num/ Source
 In_IP  In_Port I_Proto Ref_IP    Ref_Port R_Proto NAT Pool                Conf_time
Max_Sess_Num  NAT Rule
```

```
192.0.2.1  60784   udp  198.51.100.68  60784     udp   dynamic-customer-source any-remote-host
254/300  0/30 105
```

Sample Output

show security nat source persistent–nat–table all

```
user@host> show security nat source persistent-nat-table all
 Internal             Reflective                  Source     Type          Left_time/
Curr_Sess_Num/  Source
 In_IP     In_Port I_Proto Ref_IP       Ref_Port R_Proto NAT Pool                    Conf_time
Max_Sess_Num    NAT Rule
192.0.2.3    14000   udp     203.51.0.68   64933     udp    dynamic-customer-source any-remote-
host  -/300     1/32
```

Sample Output

show security nat source persistent-nat-table summary

```
user@host> show security nat source persistent-nat-table summary
Persistent NAT Table Statistics on FPC5 PIC0:
binding total : 65536
binding in use : 0
enode total : 524288
enode in use : 0
```

Release Information

Command introduced in Junos OS Release 9.6. Support.

show security nat source pool

IN THIS SECTION

- Syntax | 24914
- Description | 24914
- Options | 24914
- Required Privilege Level | 24915
- Output Fields | 24915
- Sample Output | 24917
- Sample Output | 24919
- Release Information | 24926

Syntax

```
show security nat source pool
pool-name
all
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display information about the specified Network Address Translation (NAT) source address pool and the configured twin port range per pool.

Options

pool-name	Display source NAT information for the specified address pool.
all	Display information about all source NAT address pools.

logical-system Display information about the source NAT pools for a specified logical system. Specify `all` to display information for all logical system.

root-logical-system Display information about the source NAT pools for the primary (root) logical system.

tenant Display information about the source NAT pools for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 1301 on page 24915 lists the output fields for the `show security nat source pool` command. Output fields are listed in the approximate order in which they appear.

Table 1301: show security nat source pool Output Fields

Field Name	Field Description
Pool name	Name of the source pool.
Description	Description of the source pool.
Pool id	Pool identification number.
Routing Instance	Name of the routing instance.
Host address base	Base address of the original source IP address range.
Port	Port numbers used for the source pool.
Twin port	Upper and lower limits of the twin port.
port overloading	Number of port overloading for the source pool.

Table 1301: show security nat source pool Output Fields *(Continued)*

Field Name	Field Description
Address assignment	Type of address assignment.
Total addresses	Number of IP addresses that are in use. Starting in Junos OS Release 15.1X49-D90 and Junos OS Release 17.3R1, the total number of addresses for pools with IPv6 prefixes is shown as zero (0).
Translation hits	Number of translation hits.
Port block size	Block size for the deterministic pool.
Last block recycle timeout	Amount of time before the last active block is released.
Interim logging interval	Time interval for which additional system log messages are sent for active blocks and for inactive blocks with existing sessions.
Determ host range num	Host range for the deterministic pool.
Address range	IP address or IP address range for the source pool.
Address-Persistent	Address-persistent information for IPv4 source pools: • IPv6 prefix length–Configured IPv6 prefix length. • IPv6 subscriber out of port–Number of port allocation failures.
Single Ports	Number of allocated single ports.
Twin Ports	Number of allocated twin ports.

Sample Output

show security nat source pool src_v4_pool-namelength-is-now-changed-upto-63-characters-length

```
user@host> show security nat source pool src_v4_pool-namelength-is-now-changed-upto-63-
characters-length
Pool name          : src_v4_pool-namelength-is-now-changed-upto-63-characters-length
Pool id            : 8
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Twin port          : [63488, 65535]
Port overloading   : 1
Address assignment : paired
Total addresses    : 64
Translation hits   : 0
Address range                        Single Ports   Twin Ports
         203.0.203.1 - 203.0.203.1         0              0
         203.0.203.2 - 203.0.203.2         0              0
         203.0.203.3 - 203.0.203.3         0              0
         203.0.203.4 - 203.0.203.4         0              0
         203.0.203.5 - 203.0.203.5         0              0
         203.0.203.6 - 203.0.203.6         0              0
         203.0.203.7 - 203.0.203.7         0              0
         203.0.203.8 - 203.0.203.8         0              0
         203.0.203.9 - 203.0.203.9         0              0
        203.0.203.10 - 203.0.203.10        0              0
        203.0.203.11 - 203.0.203.11        0              0
        203.0.203.12 - 203.0.203.12        0              0
        203.0.203.13 - 203.0.203.13        0              0
        203.0.203.14 - 203.0.203.14        0              0
        203.0.203.15 - 203.0.203.15        0              0
        203.0.203.16 - 203.0.203.16        0              0
        203.0.203.17 - 203.0.203.17        0              0
        203.0.203.18 - 203.0.203.18        0              0
        203.0.203.19 - 203.0.203.19        0              0
        203.0.203.20 - 203.0.203.20        0              0
        203.0.203.21 - 203.0.203.21        0              0
        203.0.203.22 - 203.0.203.22        0              0
        203.0.203.23 - 203.0.203.23        0              0
        203.0.203.24 - 203.0.203.24        0              0
```

```
        203.0.203.25 - 203.0.203.25        0              0
        203.0.203.26 - 203.0.203.26        0              0
        203.0.203.27 - 203.0.203.27        0              0
        203.0.203.28 - 203.0.203.28        0              0
        203.0.203.29 - 203.0.203.29        0              0
        203.0.203.30 - 203.0.203.30        0              0
        203.0.203.31 - 203.0.203.31        0              0
        203.0.203.32 - 203.0.203.32        0              0
        203.0.203.33 - 203.0.203.33        0              0
        203.0.203.34 - 203.0.203.34        0              0
        203.0.203.35 - 203.0.203.35        0              0
        203.0.203.36 - 203.0.203.36        0              0
        203.0.203.37 - 203.0.203.37        0              0
        203.0.203.38 - 203.0.203.38        0              0
        203.0.203.39 - 203.0.203.39        0              0
        203.0.203.40 - 203.0.203.40        0              0
        203.0.203.41 - 203.0.203.41        0              0
        203.0.203.42 - 203.0.203.42        0              0
        203.0.203.43 - 203.0.203.43        0              0
        203.0.203.44 - 203.0.203.44        0              0
        203.0.203.45 - 203.0.203.45        0              0
        203.0.203.46 - 203.0.203.46        0              0
        203.0.203.47 - 203.0.203.47        0              0
        203.0.203.48 - 203.0.203.48        0              0
        203.0.203.49 - 203.0.203.49        0              0
        203.0.203.50 - 203.0.203.50        0              0
        203.0.203.51 - 203.0.203.51        0              0
        203.0.203.52 - 203.0.203.52        0              0
        203.0.203.53 - 203.0.203.53        0              0
        203.0.203.54 - 203.0.203.54        0              0
        203.0.203.55 - 203.0.203.55        0              0
        203.0.203.56 - 203.0.203.56        0              0
        203.0.203.57 - 203.0.203.57        0              0
        203.0.203.58 - 203.0.203.58        0              0
        203.0.203.59 - 203.0.203.59        0              0
        203.0.203.60 - 203.0.203.60        0              0
        203.0.203.61 - 203.0.203.61        0              0
        203.0.203.62 - 203.0.203.62        0              0
        203.0.203.63 - 203.0.203.63        0              0
        203.0.203.64 - 203.0.203.64        0              0
Total used ports   :                     0              0
```

Sample Output

show security nat source pool all (SRX Series Firewalls)

```
user@host> show security nat source pool all

Total pools: 4

Pool name          : src_v4_pool-namelength-is-now-changed-upto-63-characters-length
Pool id            : 8
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Twin port          : [63488, 65535]
Port overloading   : 1
Address assignment : paired
Total addresses    : 64
Translation hits   : 0Address range                        Single Ports   Twin Ports
          203.0.203.1 - 203.0.203.1         0              0
          203.0.203.2 - 203.0.203.2         0              0
          203.0.203.3 - 203.0.203.3         0              0
          203.0.203.4 - 203.0.203.4         0              0
          203.0.203.5 - 203.0.203.5         0              0
          203.0.203.6 - 203.0.203.6         0              0
          203.0.203.7 - 203.0.203.7         0              0
          203.0.203.8 - 203.0.203.8         0              0
          203.0.203.9 - 203.0.203.9         0              0
         203.0.203.10 - 203.0.203.10        0              0
         203.0.203.11 - 203.0.203.11        0              0
         203.0.203.12 - 203.0.203.12        0              0
         203.0.203.13 - 203.0.203.13        0              0
         203.0.203.14 - 203.0.203.14        0              0
         203.0.203.15 - 203.0.203.15        0              0
         203.0.203.16 - 203.0.203.16        0              0
         203.0.203.17 - 203.0.203.17        0              0
         203.0.203.18 - 203.0.203.18        0              0
         203.0.203.19 - 203.0.203.19        0              0
         203.0.203.20 - 203.0.203.20        0              0
         203.0.203.21 - 203.0.203.21        0              0
         203.0.203.22 - 203.0.203.22        0              0
         203.0.203.23 - 203.0.203.23        0              0
         203.0.203.24 - 203.0.203.24        0              0
```

```
        203.0.203.25 - 203.0.203.25        0              0
        203.0.203.26 - 203.0.203.26        0              0
        203.0.203.27 - 203.0.203.27        0              0
        203.0.203.28 - 203.0.203.28        0              0
        203.0.203.29 - 203.0.203.29        0              0
        203.0.203.30 - 203.0.203.30        0              0
        203.0.203.31 - 203.0.203.31        0              0
        203.0.203.32 - 203.0.203.32        0              0
        203.0.203.33 - 203.0.203.33        0              0
        203.0.203.34 - 203.0.203.34        0              0
        203.0.203.35 - 203.0.203.35        0              0
        203.0.203.36 - 203.0.203.36        0              0
        203.0.203.37 - 203.0.203.37        0              0
        203.0.203.38 - 203.0.203.38        0              0
        203.0.203.39 - 203.0.203.39        0              0
        203.0.203.40 - 203.0.203.40        0              0
        203.0.203.41 - 203.0.203.41        0              0
        203.0.203.42 - 203.0.203.42        0              0
        203.0.203.43 - 203.0.203.43        0              0
        203.0.203.44 - 203.0.203.44        0              0
        203.0.203.45 - 203.0.203.45        0              0
        203.0.203.46 - 203.0.203.46        0              0
        203.0.203.47 - 203.0.203.47        0              0
        203.0.203.48 - 203.0.203.48        0              0
        203.0.203.49 - 203.0.203.49        0              0
        203.0.203.50 - 203.0.203.50        0              0
        203.0.203.51 - 203.0.203.51        0              0
        203.0.203.52 - 203.0.203.52        0              0
        203.0.203.53 - 203.0.203.53        0              0
        203.0.203.54 - 203.0.203.54        0              0
        203.0.203.55 - 203.0.203.55        0              0
        203.0.203.56 - 203.0.203.56        0              0
        203.0.203.57 - 203.0.203.57        0              0
        203.0.203.58 - 203.0.203.58        0              0
        203.0.203.59 - 203.0.203.59        0              0
        203.0.203.60 - 203.0.203.60        0              0
        203.0.203.61 - 203.0.203.61        0              0
        203.0.203.62 - 203.0.203.62        0              0
        203.0.203.63 - 203.0.203.63        0              0
        203.0.203.64 - 203.0.203.64        0              0
Total used ports   :                     0              0


Pool name          : src-p2
```

```
Description        : The source pool src-p2 is for the sales team
Pool id            : 5
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Address assignment : no-paired
port overloading   : 1
Total addresses    : 1
Translation hits   : 0
Address range                        Single Ports   Twin Ports
            192.0.2.0 - 192.0.2.3            0              0

Pool name          : src-p3
Description        : The source pool src-p3 is for the sales team
Pool id            : 6
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Address assignment : no-paired
port overloading   : 1
Total addresses    : 1
Translation hits   : 0
Address range                                Single Ports   Twin Ports
            2001:db8::1 - 2001:db8::1            0              0

Pool name          : src-p4
Description        : The source pool src-p4 is for the sales team
Pool id            : 7
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Address assignment : no-paired
port overloading   : 1
Total addresses    : 1
Translation hits   : 0
Address range                                Single Ports   Twin Ports
            2001:db8::2 - 2001:db8::2            0              0
```

show services nat source pool all (MX-SPC3)

```
user@host> show services nat source pool all
Total pools: 1
Interface: vms-4/0/0 , Service set: ss1
Pool name          : NAT_POOL1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-POOL
Pool id            : 4
Host address base  : 0.0.0.0
Port               : [3035, 3036]
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 64
Translation hits   : 0
Ei-mapping-timeout : 300
Mapping-timeout    : 300
EIF Inbound session count: 0
EIF Inbound session limit exceeded drops: 0
Address range                        Ports
          203.1.1.0 - 203.1.1.0          0
          203.2.1.0 - 203.2.1.0          0
          203.3.1.0 - 203.3.1.0          0
          203.4.1.0 - 203.4.1.0          0
          203.5.1.0 - 203.5.1.0          0
          203.6.1.0 - 203.6.1.0          0
          203.7.1.0 - 203.7.1.0          0
          203.8.1.0 - 203.8.1.0          0
          203.9.1.0 - 203.9.1.0          0
         203.10.1.0 - 203.10.1.0         0
         203.11.1.0 - 203.11.1.0         0
         203.12.1.0 - 203.12.1.0         0
         203.13.1.0 - 203.13.1.0         0
         203.14.1.0 - 203.14.1.0         0
         203.15.1.0 - 203.15.1.0         0
         203.16.1.0 - 203.16.1.0         0
         203.17.1.0 - 203.17.1.0         0
         203.18.1.0 - 203.18.1.0         0
         203.19.1.0 - 203.19.1.0         0
         203.20.1.0 - 203.20.1.0         0
         203.21.1.0 - 203.21.1.0         0
         203.22.1.0 - 203.22.1.0         0
         203.23.1.0 - 203.23.1.0         0
         203.24.1.0 - 203.24.1.0         0
```

```
          203.25.1.0 - 203.25.1.0         0
          203.26.1.0 - 203.26.1.0         0
          203.27.1.0 - 203.27.1.0         0
          203.28.1.0 - 203.28.1.0         0
          203.29.1.0 - 203.29.1.0         0
          203.30.1.0 - 203.30.1.0         0
          203.31.1.0 - 203.31.1.0         0
          203.32.1.0 - 203.32.1.0         0
          203.33.1.0 - 203.33.1.0         0
          203.34.1.0 - 203.34.1.0         0
          203.35.1.0 - 203.35.1.0         0
          203.36.1.0 - 203.36.1.0         0
          203.37.1.0 - 203.37.1.0         0
          203.38.1.0 - 203.38.1.0         0
          203.39.1.0 - 203.39.1.0         0
          203.40.1.0 - 203.40.1.0         0
          203.41.1.0 - 203.41.1.0         0
          203.42.1.0 - 203.42.1.0         0
          203.43.1.0 - 203.43.1.0         0
          203.44.1.0 - 203.44.1.0         0
          203.45.1.0 - 203.45.1.0         0
          203.46.1.0 - 203.46.1.0         0
          203.47.1.0 - 203.47.1.0         0
          203.48.1.0 - 203.48.1.0         0
          203.49.1.0 - 203.49.1.0         0
          203.50.1.0 - 203.50.1.0         0
          203.51.1.0 - 203.51.1.0         0
          203.52.1.0 - 203.52.1.0         0
          203.53.1.0 - 203.53.1.0         0
          203.54.1.0 - 203.54.1.0         0
          203.55.1.0 - 203.55.1.0         0
          203.56.1.0 - 203.56.1.0         0
          203.57.1.0 - 203.57.1.0         0
          203.58.1.0 - 203.58.1.0         0
          203.59.1.0 - 203.59.1.0         0
          203.60.1.0 - 203.60.1.0         0
          203.61.1.0 - 203.61.1.0         0
          203.62.1.0 - 203.62.1.0         0
          203.63.1.0 - 203.63.1.0         0
          203.64.1.0 - 203.64.1.0         0
Total used ports   :                      0
Error Counters:
    Out of port errors                     : 0
```

```
    Out of address errors                  : 0
    Parity port errors                     : 0
    Preserve Range errors                  : 0
    APP port allocation errors             : 0
    APP port limit allocation errors       : 0
```

show security nat source pool all tenant

```
user@host> show security nat source pool all tenant tn1

Total pools: 1
Pool name          : pat
Pool id            : 4
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Twin port          : [63488, 65535]
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 24
Translation hits   : 0
Address range                        Single Ports   Twin Ports
          192.0.2.1 - 192.0.2.24         0              0
Total used ports   :                      0              0
```

show security nat source pool P_1

```
user@host>show security nat source pool P_1
Pool name          : P_1
Pool id            : 4
Routing instance   : default
Port               : [12345, 17890]
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 256
Translation hits   : 0
Port block size      : 1000
Determ host range num: 3
```

```
Address range                                Single Ports   Twin Ports
         203.0.203.0 - 203.0.203.255             0             0
```

show security nat source pool src-nat-v4-with-pat

```
user@host>how security nat source pool src-nat-v4-with-pat

Pool name          : src-nat-v4-with-pat
Pool id            : 5
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 1
Translation hits   : 0
Address-persistent
        IPv6 prefix length: 64
        IPv6 subscriber out of port: 0
Address range                                Single Ports   Twin Ports
            203.0.203.1 - 203.0.203.10           0              0
```

show security nat source pool (with map-e confidentiality enabled)

```
user@host> show security nat source pool p1

Pool name          : p1
Pool id            : 4
Routing instance   : default
Host address base  : 0.0.0.0
Map-e domain name  : domain_1
Map-e rule name    : r1
Address assignment : no-paired
Total addresses    : 1
Translation hits   : 0
Address range       Single Ports   Twin Ports
confidential          0              0
Total used ports   :    0              0
```

show security nat source pool (without map-e confidentiality enabled)

```
user@host> show security nat source pool p1

Pool name          : p1
Pool id            : 4
Routing instance   : default
Host address base  : 0.0.0.0
Map-e domain name  : domain_1
Map-e rule name    : r1
PSID offset        : 4
PSID length        : 8
PSID               : 0x34
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 1
Translation hits   : 0
Address range       Single Ports   Twin Ports
10.0.0.1 - 10.0.0.1     0              0
Total used ports   :    0              0
```

Release Information

Command introduced in Junos OS Release 9.2.

The `Description` output field added in Junos OS Release 12.1.

The `Address assignment` output field and IPv6 logical system support added in Junos OS Release 12.1X45-D10.

The `twin-port` output field added in Junos OS Release 12.1X47-D10.

The `Address-persistent` output field added in Junos OS Release 12.3X48-D10.

The `Last block recycle timeout` and `Interim logging interval` output fields added in Junos OS Release 15.1X49-D60.

The tenant option is introduced in Junos OS Release 18.3R1.

show security nat source port-block

IN THIS SECTION

- Syntax | 24927
- Description | 24927
- Options | 24927
- Required Privilege Level | 24928
- Output Fields | 24928
- Sample Output (Junos OS) | 24929
- Release Information | 24930

Syntax

```
show security nat source port-block
pool-name
host-ip host ip address
xlated-ip xlated-ip-address
xlated-port xlated-port
root-logical-system | logical-system {lsys-name | all}
```

Description

Display the port blocks allocated by the host.

Options

pool-name	Display the PBA port block table for the specified source pool.
host ip address	Display the PBA port block table based on the host IP address.
xlated ip address	Display the PBA port block table based on the translated IP address.
xlated-port	Display the PBA port block table based on the translated IP address and the translated port information.

root-logical-system Display the PBA port block table for the primary (root) logical system.

logical-system (*lsys-name* | all) Display information about the specified logical system source NAT pools or all logical system source NAT pools.

Required Privilege Level

view

Output Fields

The following table lists all possible fields for the `show security nat source port-block` and `show services nat source port-block` commands. The fields that appear may vary depending on the platform and software release.

Table 1302: show security nat source port-block Output Fields

Field Name	Field Description
Pool name	Name of pool.
Port-overloading-factor	Factor of port overloading for the source pool.
Port block size	Number of ports that a port block contains.
Max port blocks per host	Maximum number of blocks that one host can use for translation.
Port block active timeout	Longest duration that a block remains active for port allocation.
Used/total port blocks	Current number of used ports and the total number of ports in this source pool.
Host IP	Address of the host IP.
External IP	Address of an external IP.

Table 1302: show security nat source port-block Output Fields *(Continued)*

Field Name	Field Description
Port_Block Range	Port range of one PBA port block entry from the lowest to the highest port number that can be allowed to allocate ports for this block.
Ports_Used/Ports_Total	Current number of used ports and total number of ports in this source pool.
Block_State/Left_Time(s)	PBA port block entry state for NAT port allocation, including Active, Inactive, Query, and the time left for a port block that is in the Active state or Query state. • Active—When an internal subscriber initiates a NAT request, a port block is allocated from the pool, and the status is set to Active. When there is a subsequent request from the same subscriber, a port is allocated from the existing Active block. • Inactive—When there is a request from an internal subscriber who has previously had a port allocated from this port block, but the time on the Active port block has expired or the ports are used up, the port block status changes from Active to Inactive. • InactiveB—When a chassis cluster is in active/passive mode, and a port block is created on the active node, the status for the synced port block on the backup node is InactiveB. • Query—When no ports are used in an Active port block, the status changes from Active to Query.
Last active block recycle timeout	Amount of time before the last active block is released when `active-port-block-timeout` is set to zero.

Sample Output (Junos OS)

show security nat source port-block (SRX Series Firewall)

```
user@host> show security nat source port-block
Pool name: pool1
Port-overloading-factor:     1     Port block size:   10
Max port blocks per host:    2     Port block active timeout:    0
```

```
Last active block recycle timeout:    0
Used/total port blocks: 1/6451
Host_IP                       External_IP                    Port_Block
Ports_Used/        Block_State/
                                                                Range
Ports_Total        Left_Time(s)
10.10.10.3                     192.0.2.19                    24994-25003
6/10*1           Active/-
```

show services nat source port-block (MX-SPC3)

```
user@host> show services nat source port-block
Pool name: NAT_POOL1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-POOL
Port-overloading-factor:     1     Port block size:    2
Max port blocks per host:    1     Port block active timeout:    0
Used/total port blocks: 0/64
```

show services nat source port-block (MX240, MX480, and MX960)

```
user@host> show services nat source port-block
Interface: vms-3/0/0
Pool name: Pool2
Port-overloading-factor:     1    Port block size:    5
Max port blocks per host:    8    Port block active timeout:    0
Used/total port blocks: 2/8
Host_IP                       External_IP                    Port_Block
Ports_Used/        Block_State/
                                                             Range
Ports_Total        Left_Time(s)
10.10.10.2                     192.0.2.18                    2005-2009
5/5*1          Inactive/-
10.10.10.2                     192.0.2.18                    2020-2024
5/5*1            Active/-
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10. The `Last active block recycle timeout` output field added in Junos OS Release 15.1X49-D60.

show security nat source rule

IN THIS SECTION

- Syntax | 24931
- Description | 24931
- Options | 24931
- Required Privilege Level | 24932
- Output Fields | 24932
- Sample Output | 24934
- Sample Output | 24935
- Sample Output | 24936
- Release Information | 24937

Syntax

```
show security nat static rule
rule-name
all
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display information about the specified source Network Address Translation (NAT) rule.

Options

rule-name	Name of the rule.
all	Display information about all the source NAT rules.

logical-system Display information about the source NAT rules for a specified logical system. Specify `all` to display information for all logical systems.

root-logical-system Display information about the source NAT rules for the primary (root) logical system.

tenant Display information about the source NAT rules for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 1303 on page 24932 lists the output fields for the **show security nat source rule** command. Output fields are listed in the approximate order in which they appear

Table 1303: show security nat source rule Output Fields

Field Name	Field Description
Source NAT rule	Name of the source NAT rule.
Total rules	Number of source NAT rules.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly, as address names, and as address set names in the rule.
Description	Description of the source NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the source NAT rule.
From zone	Name of the zone from which the packets flow.

Table 1303: show security nat source rule Output Fields *(Continued)*

Field Name	Field Description
To zone	Name of the zone to which the packets flow.
From routing instance	Name of the routing instance from which the packets flow.
To routing instance	Name of the routing instance to which the packets flow.
From interface	Name of the interface from which the packets flow.
To interface	Name of the interface to which the packets flow.
Source addresses	Name of the source addresses that match the rule.
Source port	Source port numbers that match the rule.
Destination address	Name of the destination addresses that match the rule.
Destination ports	Destination port numbers that match the rule.
Application	Indicates whether the application option is configured.
Action	The action taken in regard to a packet that matches the rule's tuples. Actions include the following: • **off**—Do not perform source NAT. • **source NAT pool**—Use user-defined source NAT pool to perform source NAT • **interface**—Use egress interface's IP address to perform source NAT.
Persistent NAT type	Persistent NAT type.

Table 1303: show security nat source rule Output Fields *(Continued)*

Field Name	Field Description
Persistent NAT mapping type	Persistent NAT mapping type.
Inactivity timeout	Inactivity timeout for persistent NAT binding.
Max session number	Maximum number of sessions.
Translation hits	Number of translation hits.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat source rule rule_namelength_is_now_changed_upto_63_characters_length

```
user@host> show security nat source rule rule_namelength_is_now_changed_upto_63_characters_length

source NAT rule: rule_namelength_is_now_changed_upto_63_characters_length
  Rule set                   : ruleset_namelength_now_changed_upto_63_characters_length
  Rule Id                    : 1
  Rule position              : 1
  From zone                  : trust
  To zone                    : untrust
  Match
    Source addresses         : 192.0.2.0 - 192.0.2.255
  Action                        : src_v4_pool-namelength-is-now-changed-upto-63-characters-length
    Persistent NAT type         : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout          : 0
    Max session number          : 0
```

```
  Translation hits           : 0
    Successful sessions      : 0
  Number of sessions         : 0
```

Sample Output

show security nat source rule all (SRX Series Firewalls)

```
user@host> show security nat source rule all
Logical system: root
Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 3/0

source NAT rule: r2                   Rule-set: rs2
  Rule-Id                    : 2
  Rule position              : 1
  From zone                  : trust
  To zone                    : untrust
  Match
    Source addresses         : 192.0.2.0 - 192.0.2.255
    Destination addresses    : 203.0.113.0 - 203.0.113.255
                               198.51.100.0 - 198.51.100.255
    Application              : configured
  Action                        : off
    Persistent NAT type         : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout          : 0
    Max session number          : 0
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

show services nat source rule all (MX-SPC3)

```
user@host> show services nat source rule all
Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 1/0
Interface: vms-4/0/0 , Service set: ss1
```

```
source NAT rule: NAT_RULE1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-RULE
  Rule set                  : NAT_RULE-SET1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_RSET
  Rule Id                   : 1
  Rule position             : 1
  Match
    Source addresses        : 0.0.0.0         - 255.255.255.255
    Application             : configured
  Action                       : NAT_POOL1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-POOL
    Syslog                     : enabled
    Mapping-type               : N/A
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0
```

Sample Output

show security nat source rule all tenant

```
user@host> show security nat source rule all tenant tn1
Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
source NAT rule: r1                     Rule-set: from_intf
  Rule-Id                   : 1
  Rule position             : 1
  From interface            : ge-0/0/0.0
  To interface              : ge-0/0/1.0
  Match
    Source addresses        : 192.168.1.0     - 192.168.1.255
    Destination addresses   : 203.0.113.200   - 203.0.113.200
  Action                       : pat
    Persistent NAT type        : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout         : 0
    Max session number         : 0
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0
```

Release Information

Command introduced in Junos OS Release 9.2. Support.

The **Description** output field added in Junos OS Release 12.1.

Support for IPv6 logical systems and the **Source port**, **Successful sessions**, **Failed sessions**, and **Number of sessions** output fields added in Junos OS Release 12.1X45-D10.

Output for multiple destination ports and the `application` output field added in Junos OS Release 12.1X47-D10.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

rule (Security Source NAT) | **16111**

show security nat source rule-application

IN THIS SECTION

- Syntax | **24937**
- Description | **24938**
- Options | **24938**
- Required Privilege Level | **24938**
- Output Fields | **24938**
- Sample Output | **24939**
- Sample Output | **24940**
- Release Information | **24940**

Syntax

```
show security nat source rule-application
rule-name
```

```
all
logical-system logical-system-name
root-logical-system
```

Description

Display information about the specified source Network Address Translation (NAT) rule application.

Options

rule-name	Display information about the specified source NAT rule application.
all	Display information about all the source NAT rule applications.
logical-system *logical-system-name*	Display information about the source NAT rule applications for the specified logical system.
root-logical-system	Display information about the source NAT rule applications for the primary (root) logical system.

Required Privilege Level

view

Output Fields

Table 1304 on page 24938 lists the output fields for the **show security nat source rule-application** command. Output fields are listed in the approximate order in which they appear.

Table 1304: show security nat source rule-application Output Fields

Field Name	Field Description
Destination NAT rule	Name of the source NAT rule.
Rule-set	Rule set identification number.
Rule-Id	Rule identification number.

Table 1304: show security nat source rule-application Output Fields *(Continued)*

Field Name	Field Description
Application	Name of the application or application set.
IP protocol	IP protocol identifier.
Source port range	Source port range identifier.
Destination port range	Destination port identifier.
ICMP information	• type—ICMP message type. • code—Code corresponding to the ICMP message type.

Sample Output

show security nat source rule-application for port application

```
user@host>show security nat source rule-application all

source NAT rule: r2                   Rule-set: rs2
Rule-Id                    : 2
  Application: app1
    IP protocol: 3
    Source port range: [90-90]
    Destination port range: [0-0]
    IP protocol: 4
    Source port range: [100-100]
    Destination port range: [200-200]
  Application: app2
    IP protocol: 7
    Source port range: [400-500]
    Destination port range: [80-80]
```

Sample Output

show security nat source rule-application for ICMP application

```
user@host>show security nat source rule-application all

source NAT rule: r1                    Rule-set: rs1
  Rule-Id                    : 1
    Application: junos-icmp-all
      IP protocol: icmp
      ICMP Information: type=255, code=0
    Application: icmp1
      IP protocol: icmp
      ICMP Information: type=1, code=1
    Application: junos-icmp6-all
      IP protocol: 58
      ICMP Information: type=255, code=0
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Logical Systems and Tenant Systems User Guide for Security Devices

show security nat source summary

IN THIS SECTION

- Syntax | 24941
- Description | 24941
- Options | 24941
- Required Privilege Level | 24941

- Output Fields | 24942
- Sample Output | 24943
- show security nat source summary (SRX Series Firewalls) | 24944
- Release Information | 24948

Syntax

```
show security nat source summary
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display a summary of Network Address Translation (NAT) source information.

Options

none	Display summary source NAT information.
logical-system	Display summary information about the source NAT for a specified logical system. Specify `all` to display information for all logical systems.
root-logical-system	Display summary information about the source NAT for the primary (root) logical system.
tenant	Display summary information about the source NAT for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 1305 on page 24942 lists the output fields for the **show security nat source summary** command. Output fields are listed in the approximate order in which they appear.

Table 1305: show security nat source summary Output Fields

Field Name	Field Description
Total source nat pool number	Number of source NAT pools.
Pool name	Name of the source address pool.
Address range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.
PAT	Whether Port Address Translation (PAT) is enabled (yes or no).
Total Address	Number of IP addresses that are in use. Starting in Junos OS Release 12.3X48-D55, and Junos OS Release 15.1X49-D90, and Junos OS Release 17.3R1, the total number of addresses that are in use for pools with IPv6 prefixes is shown as zero (0).
Total source nat rule number	Number of source NAT rules.
Total port number usage for port translation pool	Number of ports assigned to the pool.
Maximum port number for port translation pool	Maximum number of NAT or PAT transactions done at any given time.

Sample Output

show security nat source summary

```
user@host> show security nat source summary logical-system all



Logical system: root-logical-system
Total port number usage for port translation pool: 67108864
Maximum port number for port translation pool: 134217728

Logical system: lsys1
Total port number usage for port translation pool: 193536
Maximum port number for port translation pool: 134217728
Total pools: 2

Logical system: root-logical-system
Pool               Address                   Routing   PAT   Total
Name               Range                     Instance        Address
pool1              10.1.1.0-10.1.4.255-      default   yes   2048
                   10.1.5.0-10.1.8.255



Logical system: lsys1
Pool               Address                    Routing   PAT   Total
Name               Range                      Instance        Address
pool2              203.0.203.1-203.0.203.3     default   yes   3



Total rules: 1

Logical system: root-logical-system
Rule name         Rule set      From         To             Action
rule 1            ruleset1      ge-2/2/2.0   ge-2/2/3.0     pool1
rule 1                          ge-2/2/4.0   ge-2/2/5.0
```

show security nat source summary tenant

```
user@host> show security nat source summary tenant tn1
```

```
Total port number usage for port translation pool: 1548288
Maximum port number for port translation pool: 268435456
Total pools: 1
Pool                 Address                  Routing              PAT  Total
Name                 Range                    Instance                  Address
pat                  192.0.2.1-192.0.2.24     default              yes  24

Total rules: 1
Rule name          Rule set       From              To                   Action
r1                 from_intf      ge-0/0/0.0        ge-0/0/1.0           pat
```

show security nat source summary (SRX Series Firewalls)

```
user@host> show security nat source summary
Total port number usage for port translation pool: 4128768
Maximum port number for port translation pool: 50331648
Total pools: 1
Pool                 Address                  Routing              PAT  Total
Name                 Range                    Instance                  Address
src_v4_pool-namelength-is-now-changed-upto-63-characters-length 192.0.2.1-192.0.2.1
 -                   192.0.2.2-192.0.2.2
 -                   192.0.2.3-192.0.2.3
 -                   192.0.2.4-192.0.2.4
 -                   192.0.2.5-192.0.2.5
 -                   192.0.2.6-192.0.2.6
 -                   192.0.2.7-192.0.2.7
 -                   192.0.2.8-192.0.2.8
 -                   192.0.2.9-192.0.2.9
 -                   192.0.2.10-192.0.2.10
 -                   192.0.2.11-192.0.2.11
 -                   192.0.2.12-192.0.2.12
 -                   192.0.2.13-192.0.2.13
 -                   192.0.2.14-192.0.2.14
 -                   192.0.2.15-192.0.2.15
 -                   192.0.2.16-192.0.2.16
 -                   192.0.2.17-192.0.2.17
 -                   192.0.2.18-192.0.2.18
 -                   192.0.2.19-192.0.2.19
 -                   192.0.2.20-192.0.2.20
```

```
-                    192.0.2.21-192.0.2.21
-                    192.0.2.22-192.0.2.22
-                    192.0.2.23-192.0.2.23
-                    192.0.2.24-192.0.2.24
-                    192.0.2.25-192.0.2.25
-                    192.0.2.26-192.0.2.26
-                    192.0.2.27-192.0.2.27
-                    192.0.2.28-192.0.2.28
-                    192.0.2.29-192.0.2.29
-                    192.0.2.30-192.0.2.30
-                    192.0.2.31-192.0.2.31
-                    192.0.2.32-192.0.2.32
-                    192.0.2.33-192.0.2.33
-                    192.0.2.34-192.0.2.34
-                    192.0.2.35-192.0.2.35
-                    192.0.2.36-192.0.2.36
-                    192.0.2.37-192.0.2.37
-                    192.0.2.38-192.0.2.38
-                    192.0.2.39-192.0.2.39
-                    192.0.2.40-192.0.2.40
-                    192.0.2.41-192.0.2.41
-                    192.0.2.42-192.0.2.42
-                    192.0.2.43-192.0.2.43
-                    192.0.2.44-192.0.2.44
-                    192.0.2.45-192.0.2.45
-                    192.0.2.46-192.0.2.46
-                    192.0.2.47-192.0.2.47
-                    192.0.2.48-192.0.2.48
-                    192.0.2.49-192.0.2.49
-                    192.0.2.50-192.0.2.50
-                    192.0.2.51-192.0.2.51
-                    192.0.2.52-192.0.2.52
-                    192.0.2.53-192.0.2.53
-                    192.0.2.54-192.0.2.54
-                    192.0.2.55-192.0.2.55
-                    192.0.2.56-192.0.2.56
-                    192.0.2.57-192.0.2.57
-                    192.0.2.58-192.0.2.58
-                    192.0.2.59-192.0.2.59
-                    192.0.2.60-192.0.2.60
-                    192.0.2.61-192.0.2.61
-                    192.0.2.62-192.0.2.62
-                    192.0.2.63-192.0.2.63
```

```
 -                   192.0.2.64-192.0.2.64  default              yes  64

Total rules: 1
Rule name : rule_namelength_is_now_changed_upto_63_characters_length
    Rule set  : ruleset_namelength_now_changed_upto_63_characters_length
    Action    : src_v4_pool-namelength-is-now-changed-upto-63-characters-length
    From      : trust                 To : untrust
```

show services nat source summary

```
user@host> show services nat source summary
Total port number usage for port translation pool: 128
Maximum port number for port translation pool: 2576980378
Total pools: 1
Interface: vms-4/0/0 , Service set: ss1
Pool Name            Address Range            PAT  Total Address
NAT_POOL1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-POOL 203.1.1.0-203.1.1.0
 -                   203.2.1.0-203.2.1.0
 -                   203.3.1.0-203.3.1.0
 -                   203.4.1.0-203.4.1.0
 -                   203.5.1.0-203.5.1.0
 -                   203.6.1.0-203.6.1.0
 -                   203.7.1.0-203.7.1.0
 -                   203.8.1.0-203.8.1.0
 -                   203.9.1.0-203.9.1.0
 -                   203.10.1.0-203.10.1.0
 -                   203.11.1.0-203.11.1.0
 -                   203.12.1.0-203.12.1.0
 -                   203.13.1.0-203.13.1.0
 -                   203.14.1.0-203.14.1.0
 -                   203.15.1.0-203.15.1.0
 -                   203.16.1.0-203.16.1.0
 -                   203.17.1.0-203.17.1.0
 -                   203.18.1.0-203.18.1.0
 -                   203.19.1.0-203.19.1.0
 -                   203.20.1.0-203.20.1.0
 -                   203.21.1.0-203.21.1.0
 -                   203.22.1.0-203.22.1.0
 -                   203.23.1.0-203.23.1.0
 -                   203.24.1.0-203.24.1.0
 -                   203.25.1.0-203.25.1.0
```

```
 -                   203.26.1.0-203.26.1.0
 -                   203.27.1.0-203.27.1.0
 -                   203.28.1.0-203.28.1.0
 -                   203.29.1.0-203.29.1.0
 -                   203.30.1.0-203.30.1.0
 -                   203.31.1.0-203.31.1.0
 -                   203.32.1.0-203.32.1.0
 -                   203.33.1.0-203.33.1.0
 -                   203.34.1.0-203.34.1.0
 -                   203.35.1.0-203.35.1.0
 -                   203.36.1.0-203.36.1.0
 -                   203.37.1.0-203.37.1.0
 -                   203.38.1.0-203.38.1.0
 -                   203.39.1.0-203.39.1.0
 -                   203.40.1.0-203.40.1.0
 -                   203.41.1.0-203.41.1.0
 -                   203.42.1.0-203.42.1.0
 -                   203.43.1.0-203.43.1.0
 -                   203.44.1.0-203.44.1.0
 -                   203.45.1.0-203.45.1.0
 -                   203.46.1.0-203.46.1.0
 -                   203.47.1.0-203.47.1.0
 -                   203.48.1.0-203.48.1.0
 -                   203.49.1.0-203.49.1.0
 -                   203.50.1.0-203.50.1.0
 -                   203.51.1.0-203.51.1.0
 -                   203.52.1.0-203.52.1.0
 -                   203.53.1.0-203.53.1.0
 -                   203.54.1.0-203.54.1.0
 -                   203.55.1.0-203.55.1.0
 -                   203.56.1.0-203.56.1.0
 -                   203.57.1.0-203.57.1.0
 -                   203.58.1.0-203.58.1.0
 -                   203.59.1.0-203.59.1.0
 -                   203.60.1.0-203.60.1.0
 -                   203.61.1.0-203.61.1.0
 -                   203.62.1.0-203.62.1.0
 -                   203.63.1.0-203.63.1.0
 -                   203.64.1.0-203.64.1.0    yes  64

Total rules: 1
Interface: vms-4/0/0 , Service set: ss1
Rule name : NAT_RULE1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-RULE
```

```
    Rule set  : NAT_RULE-SET1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_RSET
    Action    : NAT_POOL1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-POOL
```

Release Information

Command introduced in Junos OS Release 9.2.

Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.

The tenant option is introduced in Junos OS Release 18.3R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
12.3X48-D55	Starting in Junos OS Release 12.3X48-D55, and Junos OS Release 15.1X49-D90, and Junos OS Release 17.3R1, the total number of addresses that are in use for pools with IPv6 prefixes is shown as zero (0).

RELATED DOCUMENTATION

pool (Security Source NAT) | 12781

rule (Security Source NAT) | 16111

show security nat static rule

IN THIS SECTION

- Syntax | 24949
- Description | 24949
- Options | 24949
- Required Privilege Level | 24950
- Output Fields | 24950
- Sample Output | 24951

- Sample Output | 24952
- Sample Output | 24953
- Sample Output | 24953
- Release Information | 24954

Syntax

```
show security nat static rule
rule-name
all
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display information about the specified static Network Address Translation (NAT) rule. Traffic directions allows you to specify from interface, from zone, or from routing-instance and packet information can be source addresses and ports, and destination addresses and ports.

Options

rule-name	Name of the rule.
all	Display information about all the static NAT rules.
logical-system	Display information about the static NAT rules for a specified logical system. Specify `all` to display information for all logical systems.
root-logical-system	Display information about the static NAT rules for the primary (root) logical system.
tenant	Display information about the static NAT rules for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 1306 on page 24950 lists the output fields for the **show security nat static rule** command. Output fields are listed in the approximate order in which they appear.

Table 1306: show security nat static rule Output Fields

Field Name	Field Description
Static NAT rule	Name of the static NAT rule.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly, as address names, and as address set names in the rule.
Rule-set	Name of the rule set. Currently, you can configure 8 rules within the same rule set.
Description	Description of the static NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the rule that indicates the order in which it applies to traffic.
From interface	Name of the interface from which the packets flow.
From routing instance	Name of the routing instance from which the packets flow.
From zone	Name of the zone from which the packets flow.
Destination addresses	Name of the destination addresses that match the rule.
Destination NPTv6 addr	Destination address that matches the rule.

Table 1306: show security nat static rule Output Fields *(Continued)*

Field Name	Field Description
Source addresses	Name of the source addresses that match the rule.
Host addresses	Name of the host addresses that match the rule.
Netmask	Subnet IP address.
Destination NPTv6 Netmask	Subnet IPv6 address.
Host routing-instance	Name of the host routing instance.
Destination port	Destination port numbers that match the rule. The default value is any.
Source port	Source port numbers that match the rule.
Total static-nat rules	Number of static NAT rules.
Translation hits	Number of times a translation in the translation table is used for a static NAT rule.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat static rule

```
user@host> show security nat static rule namelength-is-now-changed-upto-63-
characterslength
```

```
 Static NAT rule: namelength-is-now-changed-upto-63-characterslength              Rule-set:
namelength-is-now-changed-upto-63-characterslength  Description               : The static rule
sta-r2 is for the sales team
  Rule-Id                    : 1
  Rule position              : 1
  From zone                  : zone9
  Destination addresses      : add3
  Host addresses             : add4
  Netmask                    : 24
  Host routing-instance      : N/A
  Translation hits           : 2
   Successful sessions      : 2
   Failed sessions          : 0
 Number of sessions         : 2
```

Sample Output

show security nat static rule all tenant namelength-is-now-changed-upto-63-characterslength

```
user@host> show security nat static rule all tenant namelength-is-now-changed-upto-63-
characterslength

 Total static-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
Static NAT rule: namelength-is-now-changed-upto-63-characterslength                    Rule-
set: from_zone
  Rule-Id                    : 1
  Rule position              : 1
  From zone                  : untrust
  Source addresses           : 192.0.2.0       - 192.0.2.255
  Destination addresses      : 203.0.113.203
  Host addresses             : 192.168.1.203
  Netmask                    : 32
  Host routing-instance      : N/A
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

Sample Output

show security nat static rule (IPv6)

```
user@host> show security nat static rule namelength-is-now-changed-upto-63-
characterslength

Static NAT rule: namelength-is-now-changed-upto-63-characterslength                  Rule-set:
namelength-is-now-changed-upto-63-characterslength
   Rule-Id                    : 1
   Rule position              : 1
   From zone                  : trust
   Destination NPTv6 addr     : 2001:db8::
   Destination NPTv6 Netmask  : 48
   Host addresses             : 2001:db8::3000
   Netmask                    : 48
   Host routing-instance      : N/A
   Translation hits           : 0
      Successful sessions      : 0
      Failed sessions          : 0
   Number of sessions         : 0
```

Sample Output

show security nat static rule all (SRX Series Firewalls)

```
user@host> show security nat static rule all
Total static-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
Static NAT rule: static_rule_length_can_be_configured_upto_63_characters_lengthh
  Rule set                   : static_rule-set_length_can_be_configured_upto_63_characters_len
  Rule Id                    : 1
  Rule position              : 1
  From zone                  : trust
  Destination addresses      : 198.51.100.0
  Host addresses             : 203.0.113.0
  Netmask                    : 32
  Host routing-instance      : N/A
  Translation hits           : 0
```

```
    Successful sessions      : 0
  Number of sessions         : 0
```

show services nat static rule (MX-SPC3)

```
user@host> show services nat static rule
Total static-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 1/1
Interface: vms-4/0/0 , Service set: ss5
Static NAT rule: NAT_RULE12_WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-RULE
  Rule set                   : NAT_RULE-SET12_WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_RSET
  Rule Id                    : 1
  Rule position              : 1
  From zone                  : ss5-ZoneIn
  Destination addresses      : 198.51.100.7
  Host addresses             : 400::
  Netmask                    : 24
  Host routing-instance      : N/A
  Translation hits           : 1
    Successful sessions      : 1
    Failed sessions          : 0
  Number of sessions         : 0
```

Release Information

Command introduced in Junos OS Release 9.3.

The `Description` output field added in Junos OS Release 12.1.

Support for IPv6 logical systems and the `Successful sessions`, `Failed sessions`, `Number of sessions`, `Source addresses` and `Source ports` output fields added in Junos OS Release 12.1X45-D10.

The `Destination NPTv6 addr` and `Destination NPTv6 Netmask` output fields added in Junos OS Release 12.3X48-D25.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

rule (Security Static NAT) | **16113**

show security np-cache summary

IN THIS SECTION

- Syntax | 24955
- Description | 24955
- Required Privilege Level | 24955
- Output Fields | 24955
- Sample Output | 24956
- Sample Output | 24956
- Release Information | 24957

Syntax

```
show security np-cache summary
```

Description

This command is supported on the SRX4600, SRX5400, SRX5600, and SRX5800 devices. This command shows the IOC NP session cache information.

Required Privilege Level

view

Output Fields

Table 1307 on page 24956 lists the output fields for the `show security flow cp-session` command. Output fields are listed in the approximate order in which they appear.

Table 1307: show security flow cp-session Output Fields

Field Name	Field Description
NP-cache/services-offload Wings	Number of wings used (a wing is one half or direction of a session),
NP-cache Capacity	Capacity of the NP in number of wings.
NP-cache Usage	Percentage of wings used.

Sample Output

show security np-cache summary (SRX4600)

```
user@host> show security np-cache summary
NP-cache/services-offload Wings: 10
    NP-cache Capacity: 5242880
    NP-cache Usage: 0%
2
```

Sample Output

show security np-cache summary (SRX5400, SRX5600, and SRX5800)

```
user@host> show security np-cache summary
    IO Card NP cache usage of FPC9:
    PIC0
    NP-cache/services-offload Wings: 0
    NP-cache Capacity: 5242880
    NP-cache Usage: 0%
    PIC1
    NP-cache/services-offload Wings: 0
    NP-cache Capacity: 5242880
    NP-cache Usage: 0%
```

Release Information

Command introduced in Junos OS Release 21.2 R1.

show security pki ca-certificate

IN THIS SECTION

- Syntax | 24957
- Description | 24957
- Options | 24957
- Required Privilege Level | 24958
- Output Fields | 24958
- Sample Output | 24960
- Release Information | 24962

Syntax

```
show security pki ca-certificate
<brief | detail>
<ca-profile ca-profile-name>
```

Description

Display information about certificate authority (CA) digital certificates installed in the router.

Options

none	(Same as brief) Display information about all CA digital certificates.
brief | detail	(Optional) Display the specified level of output.
ca-profile *ca-profile-name*	(Optional) Display information about only the specified CA profile.

Required Privilege Level

view

Output Fields

Table 1308 on page 24958 lists the output fields for the **show security pki ca-certificate** command. Output fields are listed in the approximate order in which they appear.

Table 1308: show security pki ca-certificate Output Fields

Field Name	Field Description	Level of Output
Certificate identifier	Name of the digital certificate.	All levels
Certificate version	Revision number of the digital certificate.	**detail**
Serial number	Unique serial number of the digital certificate.	**detail**
Issued by	Authority that issued the digital certificate.	**none brief**
Issued to	Device that was issued the digital certificate.	**none brief**
Issuer	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • **Common name**—Name of the authority. • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **State**—State of origin. • **Country**—Country of origin.	**detail**

Table 1308: show security pki ca-certificate Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Subject	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • **Common name**—Name of the requestor. • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **State**—State of origin. • **Country**—Country of origin.	**detail**
Validity	Time period when the digital certificate is valid. Values are: • **Not before**—Start time when the digital certificate becomes valid. • **Not after**—End time when the digital certificate becomes invalid.	All levels
Public key algorithm	Encryption algorithm used with the private key, such as **rsaEncryption(1024 bits)**.	All levels
Signature algorithm	Encryption algorithm that the CA used to sign the digital certificate, such as **sha1WithRSAEncryption**.	**detail**
Fingerprint	Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate.	**detail**
Distribution CRL	Distinguished name information and the URL for the certificate revocation list (CRL) server.	**detail**
Use for key	Use of the public key, such as **Certificate signing**, **CRL signing**, **Digital signature**, or **Key encipherment**.	**detail**

Sample Output

show security pki ca-certificate

```
user@host> show security pki ca-certificate
Certificate identifier: abc
  Issued to: example, Issued by: exmple
  Validity:
    Not before: 2005 Oct 18th, 23:54:22 GMT
    Not after: 2025 Oct 19th, 00:24:22 GMT
  Public key algorithm: rsaEncryption(1024 bits)

Certificate identifier: entrust
  Issued to: First Officer, Issued by: example
  Validity:
    Not before: 2005 Oct 18th, 23:55:59 GMT
    Not after: 2008 Oct 19th, 00:25:59 GMT
  Public key algorithm: rsaEncryption(1024 bits)

Certificate identifier:abe
  Issued to: First Officer, Issued by: example
  Validity:
    Not before: 2005 Oct 18th, 23:55:59 GMT
    Not after: 2008 Oct 19th, 00:25:59 GMT
  Public key algorithm: rsaEncryption(1024 bits)
```

show security pki ca-certificate detail

```
user@host> show security pki ca-certificate detail
Certificate identifier: entrust
  Certificate version: 3
  Serial number: 4355 9235
  Issuer:
    Organization: example, Country: us
  Subject:
    Organization: example, Country: us
  Validity:
    Not before: 2005 Oct 18th, 23:54:22 GMT
    Not after: 2025 Oct 19th, 00:24:22 GMT
  Public key algorithm: rsaEncryption(1024 bits)
```

```
    cb:9e:2d:c0:70:f8:ea:3c:f2:b5:f0:02:48:87:dc:68:99:a3:57:4f
    0e:b9:98:0b:95:47:0d:1f:97:7c:53:17:dd:1a:f8:da:e5:08:d1:1c
    78:68:1f:2f:72:9f:a2:cf:81:e3:ce:c5:56:89:ce:f0:97:93:fa:36
    19:3e:18:7d:8c:9d:21:fe:1f:c3:87:8d:b3:5d:f3:03:66:9d:16:a7
    bf:18:3f:f0:7a:80:f0:62:50:43:83:4f:0e:d7:c6:42:48:c0:8a:b2
    c7:46:30:38:df:9b:dc:bc:b5:08:7a:f3:cd:64:db:2b:71:67:fe:d8
    04:47:08:07:de:17:23:13
  Signature algorithm: sha1WithRSAEncryption
  Fingerprint:
    00:8e:6f:58:dd:68:bf:25:0a:e3:f9:17:70:d6:61:f3:53:a7:79:10 (sha1)
    71:6f:6a:76:17:9b:d6:2a:e7:5a:72:97:82:6d:26:86 (md5)
  Distribution CRL:
    C=us, O=example, CN=CRL1
    http://CA-1/CRL/example_us_crlfile.crl
  Use for key: CRL signing, Certificate signing
Certificate identifier: entrust
  Certificate version: 3
  Serial number: 4355 925c
  Issuer:
    Organization: example, Country: us
  Subject:
    Organization: example, Country: us, Common name: First Officer
  Validity:
    Not before: 2005 Oct 18th, 23:55:59 GMT
    Not after: 2008 Oct 19th, 00:25:59 GMT
  Public key algorithm: rsaEncryption(1024 bits)
    c0:a4:21:32:95:0a:cd:ec:12:03:d1:a2:89:71:8e:ce:4e:a6:f9:2f
    1a:9a:13:8c:f6:a0:3d:c9:bd:9d:c2:a0:41:77:99:1b:1e:ed:5b:80
    34:46:f8:5b:28:34:38:2e:91:7d:4e:ad:14:86:78:67:e7:02:1d:2e
    19:11:b7:fa:0d:ba:64:20:e1:28:4e:3e:bb:6e:64:dc:cd:b1:b4:7a
    ca:8f:47:dd:40:69:c2:35:95:ce:b8:85:56:d7:0f:2d:04:4d:5d:d8
    42:e1:4f:6b:bf:38:c0:45:1e:9e:f0:b4:7f:74:6f:e9:70:fd:4a:78
    da:eb:10:27:bd:46:34:33
  Signature algorithm: sha1WithRSAEncryption
  Fingerprint:
    bc:78:87:9b:a7:91:13:20:71:db:ac:b5:56:71:42:ad:1a:b6:46:17 (sha1)
    23:79:40:c9:6d:a6:f0:ca:e0:13:30:d4:29:6f:86:79 (md5)
  Distribution CRL:
    C=us, O=example, CN=CRL1
    http://CA-1/CRL/example_us_crlfile.crl
  Use for key: Key encipherment
Certificate identifier: entrust
  Certificate version: 3
```

```
Serial number: 4355 925b
Issuer:
  Organization: example, Country: us
Subject:
  Organization: example, Country: us, Common name: First Officer
Validity:
  Not before: 2005 Oct 18th, 23:55:59 GMT
  Not after: 2008 Oct 19th, 00:25:59 GMT
Public key algorithm: rsaEncryption(1024 bits)
  ea:75:c4:f3:58:08:ea:65:5c:7e:b3:de:63:0a:cf:cf:ec:9a:82:e2
  d7:e8:b9:2f:bd:4b:cd:86:2f:f1:dd:d8:a2:95:af:ab:51:a5:49:4e
  00:10:c6:25:ff:b5:49:6a:99:64:74:69:e5:8c:23:5b:b4:70:62:8e
  e4:f9:a2:28:d4:54:e2:0b:1f:50:a2:92:cf:6c:8f:ae:10:d4:69:3c
  90:e2:1f:04:ea:ac:05:9b:3a:93:74:d0:59:24:e9:d2:9d:c2:ef:22
  b9:32:c7:2c:29:4f:91:cb:5a:26:fe:1d:c0:36:dc:f4:9c:8b:f5:26
  af:44:bf:53:aa:d4:5f:67
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
  46:71:15:34:f0:a6:41:76:65:81:33:4f:68:47:c4:df:78:b8:e3:3f (sha1)
  ee:cc:c7:f4:5d:ac:65:33:0a:55:db:59:72:2c:dd:16 (md5)
Distribution CRL:
  C=us, O=example, CN=CRL1
  http://CA-1/CRL/example_us_crlfile.crl
Use for key: Digital signature
```

Release Information

Command introduced in Junos OS Release 7.5.

show security pki ca-certificate (View)

IN THIS SECTION

- Syntax | 24963
- Description | 24963
- Options | 24963
- Required Privilege Level | 24963

- Output Fields | **24963**
- Sample Output | **24966**
- Release Information | **24967**

Syntax

```
show security pki ca-certificate
<brief | detail>
<ca-profile ca-profile-name>
```

Description

Display information about the certificate authority (CA) public key infrastructure (PKI) digital certificates configured on the device.

The FIPS image does not permit the use of MD5 fingerprints. Therefore, MD5 fingerprints are not included when a certificate is displayed using this command. The SHA-1 fingerprint that is currently displayed is retained in the FIPS image. The Simple Certificate Enrollment Protocol (SCEP) is disabled in the FIPS image.

Options

- none—Display basic information about all configured CA certificates.
- `brief` | `detail`—(Optional) Display the specified level of output.
- `ca-profile` *`ca-profile-name-`* (Optional) Display information about only the specified CA certificate.

Required Privilege Level

view

Output Fields

Table 1309 on page 24964 lists the output fields for the `show security pki ca-certificate` command. Output fields are listed in the approximate order in which they appear.

Table 1309: show security pki ca-certificate Output Fields

Field Name	Field Description
`CA profile`	Name of the CA profile in the CA certificate. Starting in Junos OS Release 21.4R1, you can view this information by executing the `show security pki ca-certificate <brief | detail>` command.
`Certificate identifier`	Name of the digital certificate.
`Certificate version`	Revision number of the digital certificate.
`Serial number`	Unique serial number of the digital certificate.
`Issuer`	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • `Organization`—Organization of origin. • `Organizational unit`—Department within an organization. • `Country`—Country of origin. • `Locality`—Locality of origin. • `Common name`—Name of the authority.
`Subject`	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • `Organization`—Organization of origin. • `Organizational unit`—Department within an organization. • `Country`—Country of origin. • `Locality`—Locality of origin. • `Common name`—Name of the authority. If the certificate contains multiple subfield entries, all entries are displayed.

Table 1309: show security pki ca-certificate Output Fields *(Continued)*

Field Name	Field Description
`Subject string`	Subject field as it appears in the certificate.
`Validity`	Time period when the digital certificate is valid. Values are: • `Not before`—Start time when the digital certificate becomes valid. • `Not after`—End time when the digital certificate becomes invalid.
`Public key algorithm`	Encryption algorithm used with the private key, such as `rsaEncryption(1024 bits)`.
`Signature algorithm`	Encryption algorithm that the CA used to sign the digital certificate, such as `sha1WithRSAEncryption`.
`Certificate Policy`	`Policy Identifier`—One or more policy object identifiers (OIDs).
`Use for key`	Use of the public key, such as `Certificate signing`, `CRL signing`, `Digital signature`, or `Data encipherment`.
`Fingerprint`	Secure Hash Algorithm (`SHA1`) and Message Digest 5 (`MD5`) hashes used to identify the digital certificate. Starting in Junos OS Release 21.4R1, you can also view the `SHA256` fingerprint for a local certificate along with `SHA1` and `MD5` fingerprints.
`Distribution CRL`	Distinguished name information and the URL for the certificate revocation list (CRL) server.

Sample Output

show security pki ca-certificate (MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall)

Starting in Junos OS Release 21.4R1, execute the `show security pki ca-certificate <ca-profile ca-profile-name>` command to view the CA profile name printed in the CA. The `CA profile` field in the output represents the CA profile name printed in the CA. In this sample, the CA profile name printed in the CA certificate is a `Root-CA`.

```
user@host> show security pki ca-certificate ca-profile Root-CA
LSYS: root-logical-system
  CA profile: Root-CA
Certificate identifier: Root-CA
  Issued to: Root-CA, Issued by: C = us, O = juniper, CN = Root-CA
  Validity:
    Not before: 05-19-2021 08:05 UTC
    Not after: 05-17-2031 08:05 UTC
  Public key algorithm: rsaEncryption(2048 bits)
  Keypair Location: Keypair generated locally
```

show security pki ca-certificate ca-profile detail (MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall)

Starting in Junos OS Release 21.4R1, execute the `show security pki ca-certificate <ca-profile ca-profile-name> detail` command to view:

- the CA profile name printed in the CA. The `CA profile` field in the output represents the CA profile name printed in the CA. In this sample, the CA profile name printed in the CA certificate is `Root-CA`.
- the SHA256 fingerprint for a CA certificate.

```
user@host> show security pki ca-certificate ca-profile Root-CA detail
LSYS: root-logical-system
  CA profile: Root-CA
Certificate identifier: Root-CA
  Certificate version: 3

  Serial number:
    hexadecimal: 0x00000d87
    decimal: 3463
```

```
  Issuer:
    Organization: juniper, Country: us, Common name: Root-CA
  Subject:
    Organization: juniper, Country: us, Common name: Root-CA
  Subject string:
    C=us, O=juniper, CN=Root-CA

  Validity:
    Not before: 05-19-2021 08:05 UTC
    Not after: 05-17-2031 08:05 UTC
  Public key algorithm: rsaEncryption(2048 bits)
    30:82:01:0a:02:82:01:01:00:cf:28:0c:04:ae:f0:89:f1:0a:cc:b3
    5a:0a:d9:c7:0a:f3:90:2e:7d:06:73:a4:65:94:3d:53:d4:25:2e:40
    11:98:4e:2f:52:53:1e:b3:69:2b:80:89:2e:b0:17:3a:3d:96:b3:70
    26:f7:da:ae:4e:ba:15:50:db:42:bd:bc:8c:0c:fd:5b:8e:f5:fb:74
    3c:48:8f:ec:c0:6a:5f:46:b3:1f:19:10:10:c4:e2:7e:e7:c5:ed:e1
    ff:64:01:01:f5:69:82:47:7a:2f:4c:6f:52:df:a4:06:fb:f8:ac:04
    3c:46:51:08:b4:5d:71:f3:69:a1:22:cb:53:18:74:bc:bf:4d:6b:4a
    b0:cd:4c:60:38:5f:ec:a8:6d:6c:77:dd:ed:14:a1:5f:c7:84:a7:74
    7a:6c:45:fa:4e:8a:db:8d:6c:ec:6a:25:fa:38:54:97:ac:0e:d0:12
    48:e5:0f:10:b2:3d:b0:de:95:53:d3:c8:a5:dc:6f:ed:f5:7d:49:e3
    b5:68:98:24:a7:8b:5d:a7:e5:98:de:51:b5:20:68:15:22:64:f1:c3
    cc:c4:1a:1a:be:bf:cb:fb:a7:79:92:a8:45:a3:ef:0d:2e:0f:21:f4
    5e:9d:77:1f:32:68:45:e1:93:ab:27:88:a6:c6:b2:81:55:a1:6d:c6
    81:85:1b:7f:61:02:03:01:00:01
  Signature algorithm: sha256WithRSAEncryption
  Distribution CRL:
    http://10.48.148.132:8080/crl-as-der/currentcrl-11.crl?id=11
  Authority Information Access OCSP:
    http://10.48.148.132:8090/Root-CA/
  Use for key: CRL signing, Certificate signing, Key encipherment, Digital signature
  Fingerprint:
    b4:65:6b:a2:28:01:b1:76:26:8b:8f:4f:53:b9:50:a6:eb:df:39:3a (sha1)
    14:c9:4f:da:96:15:94:6f:fa:5e:fd:60:ce:47:90:97 (md5)

49:ee:63:56:72:0b:f4:87:08:75:c9:1a:fa:6c:4d:c7:7c:2f:a2:21:31:68:30:67:87:37:cd:c0:86:34:1c:76
(sha256)
```

Release Information

Command modified in Junos OS Release 8.5.

Subject string output field added in Junos OS Release 12.1X44-D10. Policy identifier output field added in Junos OS Release 12.3X48-D10.

`CA profile` and `(sha256)` for `Fingerprint` output field added in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

- ca-profile (Security PKI)
- request security pki ca-certificate verify (Security)

show security pki certificate-request

IN THIS SECTION

- Syntax | 24968
- Description | 24968
- Options | 24969
- Required Privilege Level | 24969
- Output Fields | 24969
- Sample Output | 24971
- Release Information | 24971

Syntax

```
show security pki certificate-request
<brief | detail>
<certificate-id certificate-id-name>
```

Description

Display information about manually generated local digital certificate requests that are stored in the router.

Options

none	(same as brief) Display information about all local digital certificate requests.
brief | detail	(Optional) Display the specified level of output.
certificate-id ***certificate-id-name***	(Optional) Display information about only the specified local digital certificate request

Required Privilege Level

view

Output Fields

Table 1310 on page 24969 lists the output fields for the `show security pki certificate-request` command. Output fields are listed in the approximate order in which they appear.

Table 1310: show security pki certificate-request Output Fields

Field Name	Field Description	Level of Output
Certificate identifier	Name of the digital certificate.	All levels
Certificate version	Revision number of the digital certificate.	**detail**
Issued to	Device that was issued the digital certificate.	**none brief**

Table 1310: show security pki certificate-request Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Subject	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • **Common name**—Name of the authority. • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **State**—State of origin. • **Country**—Country of origin.	**detail**
Alternate subject	Domain name or IP address of the device related to the digital certificate.	**detail**
Validity	Time period when the digital certificate is valid. Values are: • **Not before**—Time when the digital certificate becomes valid. • **Not after**—End time when the digital certificate becomes invalid.	All levels
Public key algorithm	Encryption algorithm used with the private key, such as **rsaEncryption(1024 bits)**.	All levels
Public key verification status	Public key verification status: **Failed** or **Passed**. The **detail** output also provides the verification hash.	All levels
Fingerprint	Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate.	**detail**
Use for key	Use of the public key, such as **Certificate signing**, **CRL signing**, **Digital signature**, or **Key encipherment**.	**detail**

Sample Output

show security pki certificate-request

```
user@host> show security pki certificate-request
Certificate identifier: local-microsoft-2
  Issued to: router2.example.com
  Public key algorithm: rsaEncryption(1024 bits)
  Public key verification status: Passed
```

show security pki certificate-request detail

```
user@host> show security pki certificate-request detail
Certificate identifier: local-entrust3
  Certificate version: 3
  Subject:
    Common name: router3.example.com
  Alternate subject: router3.example.com
  Public key algorithm: rsaEncryption(1024 bits)
  Public key verification status: Passed
    fb:79:df:d4:a9:03:0f:d3:69:7e:c1:e4:27:35:9c:d9:b1:a2:47:78
    d2:6d:f3:e5:f4:68:4f:b3:04:45:88:57:99:82:39:a6:51:9e:5f:42
    23:3f:d7:6e:3d:a5:54:a9:b1:2d:6e:90:dd:12:8a:bf:ef:2b:20:50
    ba:f0:da:d9:0c:ad:5e:d6:c6:98:3a:ae:3f:90:dd:94:78:c1:ea:2e
    7c:f0:2d:d4:79:d4:cd:f0:52:df:5e:72:f2:e7:ae:66:f7:61:f4:bc
    72:57:3e:6c:6d:d3:24:58:8b:f4:ef:da:2a:6a:fa:eb:98:f8:34:84
    79:54:da:4f:d3:6f:52:1f
  Fingerprint:
    7c:e8:f9:45:93:8d:a3:92:7f:18:29:02:f1:c8:e2:85:3d:ad:df:1f (sha1)
    00:4e:df:a0:6b:ad:8c:50:da:7c:a1:cf:5d:37:b0:ea (md5)
  Use for key: Digital signature
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

clear security pki certificate-request | 4018

show security pki certificate-request (View)

IN THIS SECTION

- Syntax | 24972
- Description | 24972
- Options | 24972
- Required Privilege Level | 24973
- Output Fields | 24973
- Sample Output | 24974
- Sample Output | 24974
- Release Information | 24975

Syntax

```
show security pki certificate-request
<brief|detail>
<certificate-id certificate-id-name>
```

Description

Display information about manually generated local digital certificate requests that are stored on the device.

Options

- none—Display basic information about all local digital certificate requests.
- `brief` / `detail`—(Optional) Display the specified level of output.

- `certificate-id` *`certificate-id-name`* —(Optional) Display information about only the specified local digital certificate requests.

Required Privilege Level

view

Output Fields

Table 1311 on page 24973 lists the output fields for the `show security pki certificate-request` command. Output fields are listed in the approximate order in which they appear.

Table 1311: show security pki certificate-request Output Fields

Field Name	Field Description
`Certificate identifier`	Name of the digital certificate.
`Certificate version`	Revision number of the digital certificate.
`Issued to`	Device that was issued the digital certificate.
`Subject`	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • `Organization`—Organization of origin. • `Organizational unit`—Department within an organization. • `Country`—Country of origin. • `Locality`—Locality of origin. • `Common name`—Name of the authority.
`Alternate subject`	Domain name or IP address of the device related to the digital certificate.
`Public key algorithm`	Encryption algorithm used with the private key, such as `rsaEncryption(1024 bits)`.

Table 1311: show security pki certificate-request Output Fields *(Continued)*

Field Name	Field Description
`Public key verification status`	Public key verification status: `Failed` or `Passed`. The `detail` output also provides the verification hash.
`Fingerprint`	Secure Hash Algorithm (`SHA1`) and Message Digest 5 (`MD5`) hashes used to identify the digital certificate.
`Use for key`	Use of the public key, such as `Certificate signing`, `CRL signing`, `Digital signature`, or `Data encipherment`.

Sample Output

show security pki certificate-request certificate-id user brief

```
user@host> show security pki certificate-request certificate-id hassan brief
Certificate identifier: user
      Issued to: user@example.net
      Public key algorithm: rsaEncryption(1024 bits)
```

Sample Output

show security pki certificate-request certificate-id user detail

```
user@host> show security pki certificate-request certificate-id hassan detail
    Certificate identifier: user
      Certificate version: 3
      Subject:
        Organization: example, Organizational unit: example, Country: IN,
                        Common name: user1
      Alternate subject: 192.168.72.124
      Public key algorithm: rsaEncryption(1024 bits)
      Public key verification status: Passed
        c7:a4:fb:e7:8c:4f:31:e7:eb:01:d8:32:65:21:f2:eb:6f:7d:49:1a:c3:9b
        63:47:e2:4f:f6:db:f6:c8:75:dd:e6:ec:0b:35:0a:62:32:45:6b:35:1f:65
```

```
        c9:66:b7:40:b2:f9:2a:ab:5b:60:f7:c7:73:36:da:68:25:fc:40:4b:12:3c
        d5:c8:c6:66:f6:10:1e:86:67:a8:95:9b:7f:1c:ae:a7:55:b0:28:95:a7:9a
        a2:24:28:e4:5a:b2:a9:06:7a:69:37:20:15:e1:b6:66:eb:22:b5:b6:77:f6
        65:88:b0:94:2b:91:4b:99:78:4a:e3:56:cc:14:45:d7:97:fd
      Fingerprint:
        8f:22:1a:f2:9f:27:b0:21:6c:da:46:64:31:34:1f:68:42:5a:39:e0 (sha1)
        09:15:11:aa:ea:f9:5a:b5:70:d7:0b:8e:be:a6:d3:cb (md5)
      Use for key: Digital signature
```

Release Information

Command modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

clear security pki key-pair (Local Certificate) | 4023

show security pki crl

IN THIS SECTION

- Syntax | 24976
- Description | 24976
- Options | 24976
- Required Privilege Level | 24976
- Output Fields | 24976
- Sample Output | 24977
- Release Information | 24978

Syntax

```
show security pki crl
<brief | detail>
<ca-profile ca-profile-name>
```

Description

Display information about the certificate revocation lists (CRLs) that are stored in the router.

Options

none	(same as brief) Display information about all CRLs.
brief | detail	(Optional) Display the specified level of output.
ca-profile *ca-profile-name*	(Optional) Display CRL information about only the specified CA profile.

Required Privilege Level

view

Output Fields

Table 1312 on page 24976 shows the output fields for the `show security pki crl` command. Output fields are listed in the approximate order in which they appear.

Table 1312: show security pki crl Output Fields

Field Name	Field Description	Level of Output
CA profile	Name of the configured CA profile.	All levels
CRL version	Revision number of the certificate revocation list.	All levels
CRL number	Number of the certificate revocation list	All levels

Table 1312: show security pki crl Output Fields *(Continued)*

Field Name	Field Description	Level of Output
CRL Issuer	Device that was issued the certificate revocation list.	All levels
Issuer	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • **Common name**—Name of the authority. • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **State**—State of origin. • **Country**—Country of origin.	**detail**
Effective date	Date and time the certificate revocation list becomes valid.	All levels
Next update	Date and time the router will download the latest version of the certificate revocation list.	All levels
Revocation List	List of digital certificates that have been revoked before their expiration date. Values are: • **Serial number**—Unique serial number of the digital certificate • **Revocation date**—Date and time that the digital certificate was revoked.	**detail**

Sample Output

show security pki crl

```
user@host> show security pki crl
    CA profile entrust
    CRL version: V2
    CRL number: 24
    CRL issuer: C=CA, O=juniper
```

```
    Effective date: 2006 May 31st, 05:35:25 GMT
    Next update: 2006 Jun 1st, 06:35:25 GMT
```

show security pki crl detail

```
user@host> show security pki crl detail
     CA profile: entrust
    CRL version: V2
    CRL number: 24
    Issuer:
        Organization: juniper, Country: ca
    Validity:
        Effective date: 2006 May 31st, 05:35:25 GMT
        Next update: 2006 Jun 1st, 06:35:25 GMT
    Revocation List:
        Serial number      Revocation date
        4451aca3 2006      May 25th, 09:13:38 GMT
        4451aca4 2006      May 25th, 10:11:33 GMT
        4451acb4 2006      May 29th, 11:28:54 GMT
        4451aceb 2006      May 29th, 11:29:01 GMT
        4451acfe 2006      May 29th, 11:29:17 GMT
        4451acff 2006      May 31st, 05:29:55 GMT
```

Release Information

Command introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

clear security pki crl | 4020

show security pki crl (View)

IN THIS SECTION

- Syntax | 24979
- Description | 24979
- Options | 24979
- Required Privilege Level | 24979
- Output Fields | 24980
- Sample Output | 24981
- Sample Output | 24981
- Sample Output | 24981
- Release Information | 24982

Syntax

```
show security pki crl
                    <                  brief                    /                  detail>
<ca-profile           ca-profile-name        >
```

Description

Display information about the certificate revocation lists (CRLs) configured on the device.

Options

- none—Display basic information about all CRLs.
- `brief` | `detail`—(Optional) Display the specified level of output.
- `ca-profile` *`ca-profile-name`*- (Optional) Display information about only the specified CA profile.

Required Privilege Level

view

Output Fields

Table 1313 on page 24980 lists the output fields for the `show security pki crl` command. Output fields are listed in the approximate order in which they appear.

Table 1313: show security pki crl Output Fields

Field Name	Field Description
`CA profile`	Name of the configured CA profile.
`CRL version`	Revision number of the certificate revocation list.
`CRL issuer`	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • `emailAddress`—Mail address of the issuing authority. • `C`—Country of origin. • `ST`—State of origin. • `L`—Locality of origin. • `O`—Organization of origin. • `OU`—Department within an organization. • `CN`—Name of the authority.
`Effective date`	Date and time the certificate revocation list becomes valid.
`Next update`	Date and time the routing platform will download the latest version of the certificate revocation list.
`Revocation List`	List of digital certificates that have been revoked before their expiration date. Values are: • `Serial number`—Unique serial number of the digital certificate. • `Revocation date`—Date and time that the digital certificate was revoked.

Sample Output

show security pki crl ca-profile ca2

```
user@host> show security pki crl ca-profile ca2
CA profile: ca2
  CRL version: V00000001
  CRL issuer: emailAddress = user@example.net, C = US, ST = ca, L = sunnyvale, O = , OU = SPG
QA, CN = 2000-spg-example-net
  Effective date: 04-26-2007 18:47
  Next update: 05- 4-2007 07:07
```

Sample Output

show security pki crl ca-profile ca2 brief

```
user@host> show security pki crl ca-profile ca2 brief
CA profile: ca2
  CRL version: V00000001
  CRL issuer: emailAddress = user@example.net, C = US, ST = ca, L = sunnyvale, O = example
networks, OU = SPG QA, CN = 2000-spg-example-net
  Effective date: 04-26-2007 18:47
  Next update: 05- 4-2007 07:07
```

Sample Output

show security pki crl ca-profile ca2 detail

```
user@host> show security pki crl ca-profile ca2 detail
CA profile: ca2
  CRL version: V00000001
  CRL issuer: emailAddress = user@example.net, C = US, ST = ca, L = sunnyvale, O = example, OU =
SPG QA, CN = 2000-spg-example-net
  Effective date: 04-26-2007 18:47
  Next update: 05- 4-2007 07:07
  Revocation List:
    Serial number              Revocation date
    174e6399000000000506       03-16-2007 23:09
```

```
    174ef3f3000000000507       03-16-2007 23:09
    17529cd6000000000508       03-16-2007 23:09
    1763ac26000000000509       03-16-2007 23:09
    21904e5700000000050a       03-16-2007 23:09
    2191cf7900000000050b       03-16-2007 23:09
    21f10eb600000000050c       03-16-2007 23:09
    2253ca2a00000000050f       03-16-2007 23:09
    2478939b000000000515       03-16-2007 23:09
    24f35004000000000516       03-16-2007 23:09
    277ddfa8000000000517       03-16-2007 23:09
    277e97bd000000000518       03-16-2007 23:09
    27846a76000000000519       03-16-2007 23:09
    2785176f00000000051a       03-16-2007 23:09
```

Release Information

Command modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

crl (Security) | 4721

show security pki local-certificate

IN THIS SECTION

- Syntax | 24983
- Description | 24983
- Options | 24983
- Required Privilege Level | 24983
- Output Fields | 24983
- Sample Output | 24985
- Release Information | 24987

Syntax

```
show security pki local-certificate
<brief | detail>
<certificate-id certificate-id-name>
<system-generated>
```

Description

Display information about the local digital certificates and the corresponding public keys installed in the router.

Options

none	(same as brief) Display information about all local digital certificates and corresponding public keys.
brief | detail	(Optional) Display the specified level of output.
certificate-id *certificate-id-name*	(Optional) Display information about only the specified the local digital certificate and corresponding public keys.
system-generated	(Optional) Auto-generated self-signed certificate.

Required Privilege Level

view

Output Fields

Table 1314 on page 24984 lists the output fields for the `show security pki local-certificate` command. Output fields are listed in the approximate order in which they appear.

Table 1314: show security pki local-certificate Output Fields

Field Name	Field Description	Level of Output
Certificate identifier	Name of the digital certificate.	All levels
Certificate version	Revision number of the digital certificate.	**detail**
Serial number	Unique serial number of the digital certificate.	**detail**
Issued by	Authority that issued the digital certificate.	none **brief**
Issued to	Device that was issued the digital certificate.	none **brief**
Issuer	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • **Common name**—Name of the authority. • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **State**—State of origin. • **Country**—Country of origin.	**detail**
Subject	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • **Common name**—Name of the authority. • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **State**—State of origin. • **Country**—Country of origin.	**detail**

Table 1314: show security pki local-certificate Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Alternate subject	Domain name or IP address of the device related to the digital certificate.	**detail**
Validity	Time period when the digital certificate is valid. Values are: • **Not before**—Start time when the digital certificate becomes valid. • **Not after**—End time when the digital certificate becomes invalid.	All levels
Public key algorithm	Encryption algorithm used with the private key, such as **rsaEncryption (1024 bits)**.	All levels
Public key verification status	Public key verification status: **Failed** or **Passed**. The **detail** output also provides the verification hash.	All levels
Signature algorithm	Encryption algorithm that the CA used to sign the digital certificate, such as **sha1WithRSAEncryption**.	**detail**
Fingerprint	Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate.	**detail**
Distribution CRL	Distinguished name information and URL for the certificate revocation list (CRL) server.	**detail**
Use for key	Use of the public key, such as **Certificate signing**, **CRL signing**, **Digital signature**, or **Key encipherment**.	**detail**

Sample Output

show security pki local-certificate

```
user@host> show security pki local-certificate
Certificate identifier: local-entrust2
```

```
  Issued to: router2.example.com, Issued by: juniper
  Validity:
    Not before: 2005 Nov 21st, 23:28:22 GMT
    Not after: 2008 Nov 21st, 23:58:22 GMT
  Public key algorithm: rsaEncryption(1024 bits)
  Public key verification status: Passed
```

show security pki local-certificate detail

```
user@host> show security pki local-certificate detail
Certificate identifier: local-entrust3
  Certificate version: 3
  Serial number: 4355 94f9
  Issuer:
    Organization: juniper, Country: us
  Subject:
    Organization: juniper, Country: us, Common name: router3.example.com
  Alternate subject: router3.example.com
  Validity:
    Not before: 2005 Nov 21st, 23:33:58 GMT
    Not after: 2008 Nov 22nd, 00:03:58 GMT
  Public key algorithm: rsaEncryption(1024 bits)
  Public key verification status: Passed
    fb:79:df:d4:a9:03:0f:d3:69:7e:c1:e4:27:35:9c:d9:b1:a2:47:78
    d2:6d:f3:e5:f4:68:4f:b3:04:45:88:57:99:82:39:a6:51:9e:5f:42
    23:3f:d7:6e:3d:a5:54:a9:b1:2d:6e:90:dd:12:8a:bf:ef:2b:20:50
    ba:f0:da:d9:0c:ad:5e:d6:c6:98:3a:ae:3f:90:dd:94:78:c1:ea:2e
    7c:f0:2d:d4:79:d4:cd:f0:52:df:5e:72:f2:e7:ae:66:f7:61:f4:bc
    72:57:3e:6c:6d:d3:24:58:8b:f4:ef:da:2a:6a:fa:eb:98:f8:34:84
    79:54:da:4f:d3:6f:52:1f
  Signature algorithm: sha1WithRSAEncryption
  Fingerprint:
    61:3a:d0:b4:7a:16:9b:39:ba:81:3f:9d:ab:34:e5:c8:be:3b:a1:6d (sha1)
    60:a0:ff:58:05:4a:65:73:9d:74:3a:e1:83:6f:1b:c8 (md5)
  Distribution CRL:
    C=us, O=juniper, CN=CRL1
    http://CA-1/CRL/juniper_us_crlfile.crl
  Use for key: Digital signature
```

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

clear security pki local-certificate | 4024

show security pki local-certificate

IN THIS SECTION

- Syntax | 24987
- Description | 24987
- Options | 24988
- Required Privilege Level | 24988
- Output Fields | 24988
- Sample Output | 24990
- Release Information | 24991

Syntax

```
show security pki local-certificate
<brief | detail>
<certificate-id certificate-id-name>
<system-generated>
```

Description

Display information about the local digital certificates and the corresponding public keys installed in the switch.

Options

none	(Same as brief) Display information about all local digital certificates and corresponding public keys.
brief | detail	(Optional) Display information about local digital certificates and corresponding public keys for the specified level of output.
certificate-id *certificate-id-name*	(Optional) Display information about only the specified the local digital certificate and corresponding public keys.
system-generated	(Optional) Display information about the automatically generated self-signed certificate.

Required Privilege Level

view

Output Fields

Table 1315 on page 24988 lists the output fields for the `show security pki local-certificate` command. Output fields are listed in the approximate order in which they appear.

Table 1315: show security pki local-certificate Output Fields

Field Name	Field Description	Level of Output
Certificate identifier	Name of the digital certificate.	All levels
Certificate version	Revision number of the digital certificate.	**detail**
Serial number	Unique serial number of the digital certificate.	**detail**
Issued by	Authority that issued the digital certificate.	none **brief**
Issued to	Device that was issued the digital certificate.	none **brief**

Table 1315: show security pki local-certificate Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Issuer	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • **Common name**—Name of the authority. • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **State**—State of origin. • **Country**—Country of origin.	**detail**
Subject	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • **Common name**—Name of the authority. • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **State**—State of origin. • **Country**—Country of origin.	**detail**
Alternate subject	Domain name or IP address of the device related to the digital certificate.	**detail**
Validity	Time period when the digital certificate is valid. Values are: • **Not before**—Start time when the digital certificate becomes valid. • **Not after**—End time when the digital certificate becomes invalid.	All levels
Public key algorithm	Encryption algorithm used with the private key, such as **rsaEncryption (1024 bits)**.	All levels

Table 1315: show security pki local-certificate Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Public key verification status	Public key verification status: **Failed** or **Passed**. The **detail** output also provides the verification hash.	All levels
Signature algorithm	Encryption algorithm that the CA used to sign the digital certificate, such as **sha1WithRSAEncryption**.	**detail**
Fingerprint	Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate.	**detail**
Distribution CRL	Distinguished name information and URL for the certificate revocation list (CRL) server.	**detail**
Use for key	Use of the public key, such as **Certificate signing**, **CRL signing**, **Digital signature**, or **Key encipherment**.	**detail**

Sample Output

show security pki local-certificate

```
user@switch> show security pki local-certificate
Certificate identifier: local-entrust2
  Issued to: router2.juniper.net, Issued by: juniper
  Validity:
    Not before: 2005 Nov 21st, 23:28:22 GMT
    Not after: 2008 Nov 21st, 23:58:22 GMT
  Public key algorithm: rsaEncryption(1024 bits)
  Public key verification status: Passed
```

show security pki local-certificate detail

```
user@switch> show security pki local-certificate detail
Certificate identifier: local-entrust3
```

```
Certificate version: 3
Serial number: 4355 94f9
Issuer:
  Organization: juniper, Country: us
Subject:
  Organization: juniper, Country: us, Common name: switch1.juniper.net
Alternate subject: switch1.juniper.net
Validity:
  Not before: 2005 Nov 21st, 23:33:58 GMT
  Not after: 2008 Nov 22nd, 00:03:58 GMT
Public key algorithm: rsaEncryption(1024 bits)
Public key verification status: Passed
  fb:79:df:d4:a9:03:0f:d3:69:7e:c1:e4:27:35:9c:d9:b1:a2:47:78
  d2:6d:f3:e5:f4:68:4f:b3:04:45:88:57:99:82:39:a6:51:9e:5f:42
  23:3f:d7:6e:3d:a5:54:a9:b1:2d:6e:90:dd:12:8a:bf:ef:2b:20:50
  ba:f0:da:d9:0c:ad:5e:d6:c6:98:3a:ae:3f:90:dd:94:78:c1:ea:2e
  7c:f0:2d:d4:79:d4:cd:f0:52:df:5e:72:f2:e7:ae:66:f7:61:f4:bc
  72:57:3e:6c:6d:d3:24:58:8b:f4:ef:da:2a:6a:fa:eb:98:f8:34:84
  79:54:da:4f:d3:6f:52:1f
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
  61:3a:d0:b4:7a:16:9b:39:ba:81:3f:9d:ab:34:e5:c8:be:3b:a1:6d (sha1)
  60:a0:ff:58:05:4a:65:73:9d:74:3a:e1:83:6f:1b:c8 (md5)
Distribution CRL:
  C=us, O=juniper, CN=CRL1
  http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: Digital signature
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Manually Generating Self-Signed Certificates on Switches (CLI Procedure)

show security pki local-certificate (View)

IN THIS SECTION

- Syntax | 24992
- Description | 24992
- Options | 24992
- Required Privilege Level | 24993
- Output Fields | 24993
- Sample Output | 24995
- Release Information | 24998

Syntax

```
show security pki local-certificate
<brief /detail>
<certificate-id certificate-id-name>
<system-generated>
```

Description

Display information about the local digital certificates, corresponding public keys, and the automatically generated self-signed certificate configured on the device.

Options

- none—Display basic information about all configured local digital certificates, corresponding public keys, and the automatically generated self-signed certificate.
- `brief` | `detail`—(Optional) Display the specified level of output.
- certificate-id *certificate-id-name* —(Optional) Display information about only the specified local digital certificates and corresponding public keys.
- `system-generated`—Display information about the automatically generated self-signed certificate.

Required Privilege Level

view

Output Fields

Table 1316 on page 24993 lists the output fields for the `show security pki local-certificate` command. Output fields are listed in the approximate order in which they appear.

Table 1316: show security pki local-certificate Output Fields

Field Name	Field Description
`Certificate identifier`	Name of the digital certificate.
`Certificate version`	Revision number of the digital certificate.
`Serial number`	Unique serial number of the digital certificate. Starting in Junos OS Release 20.1R1, PKI local certificate serial number is displayed with **0x** as prefix to indicate that the PKI local certificate is in the hexadecimal format. Starting in Junos OS Release 21.4R1, you can view the serial number of the digital certificate in both hexadecimal and decimal formats.
`Issued to`	Device that was issued the digital certificate.
`Issued by`	Authority that issued the digital certificate.
`Issuer`	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • `Organization`—Organization of origin. • `Organizational unit`—Department within an organization. • `Country`—Country of origin. • `Locality`—Locality of origin. • `Common name`—Name of the authority.

Table 1316: show security pki local-certificate Output Fields *(Continued)*

Field Name	Field Description
`LSYS`	Name of the logical systems.
`Subject`	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • `Organization`—Organization of origin. • `Organizational unit`—Department within an organization. • `Country`—Country of origin. • `Locality`—Locality of origin. • `Common name`—Name of the authority. • `Serial number`—Serial number of the device. If the certificate contains multiple subfield entries, all entries are displayed.
`Subject string`	Subject field as it appears in the certificate.
`Alternate subject`	Domain name or IP address of the device related to the digital certificate. For multiple FQDNs, displays only the last FQDN details. Starting Junos OS Release 22.4R2, with multiple FQDNs, this option shows all domain names, IPv4 or IPv6 addresses and email addresses related to the digital certificate configured on the device.
`Cert-Chain`	Starting in Junos OS Release 21.4R1, you can view the certificate chain for a given local certificate.
`Validity`	Time period when the digital certificate is valid. Values are: • `Not before`—Start time when the digital certificate becomes valid. • `Not after`—End time when the digital certificate becomes invalid.

Table 1316: show security pki local-certificate Output Fields *(Continued)*

Field Name	Field Description
`Public key algorithm`	Encryption algorithm used with the private key, such as `rsaEncryption(1024 bits)`.
`Public key verification status`	Public key verification status: `Failed` or `Passed`. The `detail` output also provides the verification hash.
`Signature algorithm`	Encryption algorithm that the CA used to sign the digital certificate, such as `sha1WithRSAEncryption`.
`Fingerprint`	Secure Hash Algorithm (`SHA1`) and Message Digest 5 (`MD5`) hashes used to identify the digital certificate. Starting in Junos OS Release 21.4R1, you can also view the SHA-256 fingerprint for a local certificate along with SHA-1 and MD-5 fingerprints.
`Distribution CRL`	Distinguished name information and URL for the certificate revocation list (`CRL`) server.
`Use for key`	Use of the public key, such as `Certificate signing`, `CRL signing`, `Digital signature`, or `Data encipherment`.

Sample Output

show security pki local-certificate certificate-id hello

```
user@host> show security pki local-certificate certificate-id hello
LSYS: root-logical-system
Certificate identifier: hello
  Issued to: tc5-5-1, Issued by: DC = Juniper, CN = root-551-AAA
  Validity:
    Not before: 10-14-2021 21:41 UTC
    Not after: 02-13-2026 14:27 UTC
```

```
  Public key algorithm: rsaEncryption(1024 bits)
  Keypair Location: Keypair generated locally
```

show security pki local-certificate system-generated

```
user@host> show security pki local-certificate system-generated

LSYS: root-logical-system
Certificate identifier: system-generated
Issued to: 4a505bb373d7, Issued by: CN = 4a505bb373d7, CN = system generated, CN = self-signed
Validity:
Not before: 07-12-2019 22:23 UTC
Not after: 07-10-2024 22:23 UTC
Public key algorithm: rsaEncryption(2048 bits)
Keypair Location: Keypair generated locally
```

show security pki local-certificate system-generated detail

```
user@host> show security pki local-certificate system-generated detail
LSYS: root-logical-system
Certificate identifier: system-generated
  Certificate version: 3

  Serial number:
    hexadecimal: 0x23171f4f104463e2847bc792c39eb614
    decimal: 46643037698975347221422984685160412692
  Issuer:
    Common name: 4a505bb373d7, Common name: system generated, Common name: self-signed
  Subject:
    Common name: 4a505bb373d7, Common name: system generated, Common name: self-signed
  Subject string:
    CN=4a505bb373d7, CN=system generated, CN=self-signed

  Validity:
    Not before: 07-12-2019 22:23 UTC
    Not after: 07-10-2024 22:23 UTC
  Public key algorithm: rsaEncryption(2048 bits)
    30:82:01:0a:02:82:01:01:00:d5:7e:5e:7a:15:90:e3:23:07:8e:e3
    4b:40:0e:95:33:31:8c:17:0b:d1:78:48:2e:b5:e8:cb:44:03:f1:fd
    00:57:af:e9:d9:2c:78:96:04:37:3c:4a:65:d9:f1:fb:72:14:7f:b2
```

```
    d3:42:d3:84:be:e8:c5:6c:e2:f5:91:8a:41:02:30:a7:8b:2f:10:5e
    ab:5e:4e:d7:d6:f1:e7:ad:e3:6c:16:8d:6b:3c:0e:11:e9:26:8a:38
    99:78:0a:57:67:cc:0a:ea:fa:35:2b:f3:51:4e:cc:30:ee:e9:a7:0a
    26:14:42:fc:1b:22:ec:2d:0c:3b:10:d5:fb:e3:e6:ae:c6:cc:e7:de
    0f:cf:4d:a7:87:11:e1:4e:7f:33:69:c0:16:4e:80:c8:57:b4:9a:f8
    90:15:d8:e6:3e:06:7a:1c:a3:34:91:92:a6:88:9f:14:f5:89:39:da
    0f:88:1c:b0:bd:7d:46:23:b2:42:e8:6f:d2:34:9e:f2:bd:00:34:23
    99:4e:bb:39:0e:e4:bb:b2:9b:53:02:36:30:10:b7:28:e3:c4:8c:0e
    4c:fd:cf:4f:58:81:72:91:b4:82:18:cf:ba:f6:76:59:f2:d5:36:e1
    3a:29:20:72:02:5b:26:45:6f:92:0c:8e:dc:6c:d4:1c:78:55:db:66
    3a:e9:9a:9c:81:02:03:01:00:01
  Signature algorithm: sha256WithRSAEncryption
  Fingerprint:
    0b:08:f8:bc:c6:a3:c1:41:75:2b:48:da:5d:a7:0f:d8:99:45:cd:8a (sha1)
    8a:1b:b9:79:19:c6:c3:88:05:a8:05:28:3c:f2:b0:e9 (md5)

a3:9b:c1:c4:55:a8:f8:79:6f:a9:27:fc:f8:5a:af:45:37:dd:42:5f:2f:2b:bb:85:e3:f0:d7:99:9d:93:65:b1
(sha256)
```

show security pki local-certificate detail (MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall)

Starting in Junos OS Release 21.4R1, execute the `show security pki local-certificate detail` command to view:

- the CA certificate chain for a local certificate. The output field `cert-chain` displays the CA certificate chain.

  if there is no certificate chain available for a given local certificate, then the `cert-chain` field displays the *Issuer/Root CA name*. If certificate chain exists, then `cert-chain` displays the Root-CA, followed by intermediate CA's.

- the local certificate serial number in both hexadecimal and decimal format.

- the SHA-256 fingerprint for a local certificate.

```
user@host> show security pki local-certificate certificate-id localcert-Sub11 detail
LSYS: root-logical-system
Certificate identifier: localcert-Sub11
  Certificate version: 3

  Serial number:
    hexadecimal: 0x0000202f
```

```
    decimal: 8239
  Issuer:
    Organization: juniper, Country: us, Common name: Sub11-CA
  Subject:
    Organizational unit: net_name, Common name: localcert-Sub11, Domain component: Juniper
  Subject string:
    DC=Juniper, CN=localcert-Sub11, OU=net_name
  Alternate subject: "localcert-Sub11@juniper.net", localcert-Sub11.juniper.net, 3.3.3.1, ipv6
empty
  Cert-Chain: Root-CA , Sub1-CA , Sub11-CA
  Validity:
    Not before: 05-19-2021 16:30 UTC
    Not after: 05-17-2031 08:05 UTC
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:ae:16:b6:d7:72:34:9e:ef:4b:9b:e2:c8:d1
    8b:2a:e4:04:16:7a:06:ac:d6:be:96:e3:2f:2b:ac:b9:28:42:1b:c4
    ef:10:1e:7d:76:a5:8f:c4:fa:b5:b6:c1:7d:53:15:b7:85:f0:aa:4c
    af:9d:35:1e:06:dc:38:ce:40:70:b3:63:b9:4c:55:eb:ba:61:85:40
    71:32:ec:5a:3a:83:1f:e3:bf:0f:8d:cd:f7:29:44:e2:c6:a3:10:62
    bb:aa:f1:ae:cc:6e:ef:8a:4e:cc:03:cf:e9:35:c5:8f:7a:21:a9:ee
    9b:c1:2d:a3:7b:94:6f:db:2a:d7:01:0a:1c:1b:c3:02:03:01:00:01
  Signature algorithm: sha256WithRSAEncryption
  Distribution CRL:
    http://10.48.148.132:8080/crl-as-der/currentcrl-23.crl?id=23
  Authority Information Access OCSP:
    http://10.48.148.132:8090/Sub11-CA/
  Fingerprint:
    4b:04:da:b1:03:a6:a2:fc:24:d4:e3:ec:61:7a:d0:10:97:10:25:9e (sha1)
    e4:6a:3d:90:a1:a2:ec:5b:3b:de:c6:3f:16:1d:02:d5 (md5)

40:d3:95:c6:3c:5e:0e:cd:32:ca:63:76:e9:83:8e:ca:ec:8a:c7:0e:84:bb:e5:a5:bc:e4:25:0c:54:0c:23:51
(sha256)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

Release Information

Command modified in Junos OS Release 9.1.

Subject string output field added in Junos OS Release 12.1X44-D10.

`Cert-Chain`, `hexadecimal` and `decimal` for `Serial Number`, `(sha256)` for `Fingerprint` output fields are added in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

- clear security pki local-certificate (Device) | **4027**
- request security pki local-certificate generate-self-signed (Security) | **14827**
- show security ipsec statistics (MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall) | **25023**
- clear security pki statistics | **4033**

show security pki node-local certificate-request

IN THIS SECTION

- Syntax | **24999**
- Description | **25000**
- Options | **25000**
- Required Privilege Level | **25000**
- Output Fields | **25000**
- Sample Output | **25001**
- Sample Output | **25002**
- Release Information | **25002**

Syntax

```
show security pki node-local certificate-request
<brief|detail>
<certificate-id certificate-id-name>
```

Description

Display information about manually generated local digital certificate requests that are stored on the local device in your Multinode High Availability setup.

Options

- none—Display basic information about all local digital certificate requests.
- `brief` / `detail`—(Optional) Display the specified level of output.
- `certificate-id` *`certificate-id-name`* —(Optional) Display information about only the specified local digital certificate requests.

Required Privilege Level

view

Output Fields

Table 1317 on page 25000 lists the output fields for the `show security pki node-local certificate-request` command. Output fields are listed in the approximate order in which they appear.

Table 1317: show security pki node-local certificate-request Output Fields

Field Name	Field Description
`Certificate identifier`	Name of the digital certificate.
`Certificate version`	Revision number of the digital certificate.
`Issued to`	Device that was issued the digital certificate.

Table 1317: show security pki node-local certificate-request Output Fields *(Continued)*

Field Name	Field Description
`Subject`	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • `Organization`—Organization of origin. • `Organizational unit`—Department within an organization. • `Country`—Country of origin. • `Locality`—Locality of origin. • `Common name`—Name of the authority.
`Alternate subject`	Domain name or IP address of the device related to the digital certificate.
`Public key algorithm`	Encryption algorithm used with the private key, such as `rsaEncryption(1024 bits)`.
`Public key verification status`	Public key verification status: `Failed` or `Passed`. The `detail` output also provides the verification hash.
`Fingerprint`	Secure Hash Algorithm (`SHA1`) and Message Digest 5 (`MD5`) hashes used to identify the digital certificate.
`Use for key`	Use of the public key, such as `Certificate signing`, `CRL signing`, `Digital signature`, or `Data encipherment`.

Sample Output

show security pki node-local certificate-request certificate-id user brief

```
user@host> show security pki node-local certificate-request certificate-id user-1
brief
Certificate identifier: user-1
```

```
      Issued to: user@example.net
      Public key algorithm: rsaEncryption(1024 bits)
```

Sample Output

show security pki node-local certificate-request certificate-id user detail

```
user@host> show security pki node-local certificate-request certificate-id user-1
detail
    Certificate identifier: user
      Certificate version: 3
      Subject:
        Organization: example, Organizational unit: example, Country: IN,
                        Common name: user1
      Alternate subject: 192.168.72.124
      Public key algorithm: rsaEncryption(1024 bits)
      Public key verification status: Passed
        c7:a4:fb:e7:8c:4f:31:e7:eb:01:d8:32:65:21:f2:eb:6f:7d:49:1a:c3:9b
        63:47:e2:4f:f6:db:f6:c8:75:dd:e6:ec:0b:35:0a:62:32:45:6b:35:1f:65
        c9:66:b7:40:b2:f9:2a:ab:5b:60:f7:c7:73:36:da:68:25:fc:40:4b:12:3c
        d5:c8:c6:66:f6:10:1e:86:67:a8:95:9b:7f:1c:ae:a7:55:b0:28:95:a7:9a
        a2:24:28:e4:5a:b2:a9:06:7a:69:37:20:15:e1:b6:66:eb:22:b5:b6:77:f6
        65:88:b0:94:2b:91:4b:99:78:4a:e3:56:cc:14:45:d7:97:fd
      Fingerprint:
        8f:22:1a:f2:9f:27:b0:21:6c:da:46:64:31:34:1f:68:42:5a:39:e0 (sha1)
        09:15:11:aa:ea:f9:5a:b5:70:d7:0b:8e:be:a6:d3:cb (md5)
      Use for key: Digital signature
```

Release Information

Command modified in Junos OS Release 22.3R1

RELATED DOCUMENTATION

request security pki node-local generate-certificate-request

Multinode High Availability

show security pki node-local local-certificate

IN THIS SECTION

- Syntax | 25003
- Description | 25003
- Options | 25003
- Required Privilege Level | 25004
- Output Fields | 25004
- Sample Output | 25006
- Release Information | 25008

Syntax

```
show security pki node-local local-certificate
<brief|detail>
<certificate-id certificate-id-name | idev-id>
<system-generated>
```

Description

Display information about the local digital certificates, corresponding public keys, and the automatically generated self-signed certificate configured on the local device in a Multinode High Availability setup.

Options

- none—Display basic information about all configured local digital certificates, corresponding public keys, and the automatically generated self-signed certificate.
- `brief` | `detail`—(Optional) Display the specified level of output.
- certificate-id *certificate-id-name* —(Optional) Display information about only the specified local digital certificates and corresponding public keys.

  Use `idev-id` as the *certificate-id-name* to display TPM-based certificate with the SRX1600, SRX2300, and SRX4300 Series Firewalls to securely identify your device.

- `system-generated`—Display information about the automatically generated self-signed certificate.

Required Privilege Level

view

Output Fields

Table 1318 on page 25004 lists the output fields for the `show security pki node-local local-certificate` command. Output fields are listed in the approximate order in which they appear.

Table 1318: show security pki node-local local-certificate Output Fields

Field Name	Field Description
`Certificate identifier`	Name of the digital certificate.
`Certificate version`	Revision number of the digital certificate.
`Serial number`	Unique serial number of the digital certificate. Starting in Junos OS Release 20.1R1, PKI local certificate serial number is displayed with **0x** as prefix to indicate that the PKI local certificate is in the hexadecimal format. Starting in Junos OS Release 21.4R1, you can view the serial number of the digital certificate in both hexadecimal and decimal formats.
`Issued to`	Device that was issued the digital certificate.
`Issued by`	Authority that issued the digital certificate.

Table 1318: show security pki node-local local-certificate Output Fields *(Continued)*

Field Name	Field Description
`Issuer`	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • `Organization`—Organization of origin. • `Organizational unit`—Department within an organization. • `Country`—Country of origin. • `Locality`—Locality of origin. • `Common name`—Name of the authority.
`LSYS`	Name of the logical systems.
`Subject`	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • `Organization`—Organization of origin. • `Organizational unit`—Department within an organization. • `Country`—Country of origin. • `Locality`—Locality of origin. • `Common name`—Name of the authority. • `Serial number`—Serial number of the device. If the certificate contains multiple subfield entries, all entries are displayed.
`Subject string`	Subject field as it appears in the certificate.
`Alternate subject`	Domain name or IP address of the device related to the digital certificate.
`Cert-Chain`	Starting in Junos OS Release 21.4R1, you can view the certificate chain for a given local certificate.

Table 1318: show security pki node-local local-certificate Output Fields *(Continued)*

Field Name	Field Description
`Validity`	Time period when the digital certificate is valid. Values are: • `Not before`—Start time when the digital certificate becomes valid. • `Not after`—End time when the digital certificate becomes invalid.
`Public key algorithm`	Encryption algorithm used with the private key, such as `rsa Encryption(1024 bits)`.
`Public key verification status`	Public key verification status: `Failed` or `Passed`. The `detail` output also provides the verification hash.
`Signature algorithm`	Encryption algorithm that the CA used to sign the digital certificate, such as `sha1WithRSAEncryption`.
`Fingerprint`	Secure Hash Algorithm (`SHA1`) and Message Digest 5 (`MD5`) hashes used to identify the digital certificate. Starting in Junos OS Release 21.4R1, you can also view the SHA-256 fingerprint for a local certificate along with SHA-1 and MD-5 fingerprints.
`Distribution CRL`	Distinguished name information and URL for the certificate revocation list (`CRL`) server.
`Use for key`	Use of the public key, such as `Certificate signing`, `CRL signing`, `Digital signature`, or `Data encipherment`.

Sample Output

show security pki node-local local-certificate certificate-id hello

```
user@host> show security pki node-local local-certificate certificate-id cert-1234
LSYS: root-logical-system
Certificate identifier: cert-1234
```

```
  Issued to: tc5-5-1, Issued by: DC = Juniper, CN = root-551-AAA
  Validity:
    Not before: 10-14-2021 21:41 UTC
    Not after: 02-13-2026 14:27 UTC
  Public key algorithm: rsaEncryption(1024 bits)
  Keypair Location: Keypair generated locally
```

show security pki node-local local-certificate system-generated

```
user@host> show security pki node-local local-certificate system-generated
LSYS: root-logical-system
Certificate identifier: system-generated
Issued to: 4a505bb373d7, Issued by: CN = 4a505bb373d7, CN = system generated, CN = self-signed
Validity:
Not before: 07-12-2019 22:23 UTC
Not after: 07-10-2024 22:23 UTC
Public key algorithm: rsaEncryption(2048 bits)
Keypair Location: Keypair generated locally
```

show security pki node-local local-certificate system-generated detail

```
user@host> show security pki node-local local-certificate system-generated detail
LSYS: root-logical-system
Certificate identifier: system-generated
  Certificate version: 3

  Serial number:
    hexadecimal: 0x23171f4f104463e2847bc792c39eb614
    decimal: 46643037698975347221422984685160412692
  Issuer:
    Common name: 4a505bb373d7, Common name: system generated, Common name: self-signed
  Subject:
    Common name: 4a505bb373d7, Common name: system generated, Common name: self-signed
  Subject string:
    CN=4a505bb373d7, CN=system generated, CN=self-signed

  Validity:
    Not before: 07-12-2019 22:23 UTC
    Not after: 07-10-2024 22:23 UTC
  Public key algorithm: rsaEncryption(2048 bits)
```

```
    30:82:01:0a:02:82:01:01:00:d5:7e:5e:7a:15:90:e3:23:07:8e:e3
    4b:40:0e:95:33:31:8c:17:0b:d1:78:48:2e:b5:e8:cb:44:03:f1:fd
    00:57:af:e9:d9:2c:78:96:04:37:3c:4a:65:d9:f1:fb:72:14:7f:b2
    d3:42:d3:84:be:e8:c5:6c:e2:f5:91:8a:41:02:30:a7:8b:2f:10:5e
    ab:5e:4e:d7:d6:f1:e7:ad:e3:6c:16:8d:6b:3c:0e:11:e9:26:8a:38
    99:78:0a:57:67:cc:0a:ea:fa:35:2b:f3:51:4e:cc:30:ee:e9:a7:0a
    26:14:42:fc:1b:22:ec:2d:0c:3b:10:d5:fb:e3:e6:ae:c6:cc:e7:de
    0f:cf:4d:a7:87:11:e1:4e:7f:33:69:c0:16:4e:80:c8:57:b4:9a:f8
    90:15:d8:e6:3e:06:7a:1c:a3:34:91:92:a6:88:9f:14:f5:89:39:da
    0f:88:1c:b0:bd:7d:46:23:b2:42:e8:6f:d2:34:9e:f2:bd:00:34:23
    99:4e:bb:39:0e:e4:bb:b2:9b:53:02:36:30:10:b7:28:e3:c4:8c:0e
    4c:fd:cf:4f:58:81:72:91:b4:82:18:cf:ba:f6:76:59:f2:d5:36:e1
    3a:29:20:72:02:5b:26:45:6f:92:0c:8e:dc:6c:d4:1c:78:55:db:66
    3a:e9:9a:9c:81:02:03:01:00:01
  Signature algorithm: sha256WithRSAEncryption
  Fingerprint:
    0b:08:f8:bc:c6:a3:c1:41:75:2b:48:da:5d:a7:0f:d8:99:45:cd:8a (sha1)
    8a:1b:b9:79:19:c6:c3:88:05:a8:05:28:3c:f2:b0:e9 (md5)

a3:9b:c1:c4:55:a8:f8:79:6f:a9:27:fc:f8:5a:af:45:37:dd:42:5f:2f:2b:bb:85:e3:f0:d7:99:9d:93:65:b1
(sha256)
```

show security pki node-local local-certificate certificate-id idev-id

```
user@host> show security pki node-local local-certificate certificate-id idev-id
LSYS: root-logical-system
Certificate identifier: idev-id
  Issued to: SRX4300-CHAS, Issued by: C = US, ST = CA, L = Sunnyvale, O = Juniper Networks Inc.,
OU = Juniper CA, CN = InitialDeviceIDRSA2048CA
  Validity:
    Not before: 09- 7-2023 12:28 UTC
    Not after: 09- 7-2123 12:38 UTC
  Public key algorithm: rsaEncryption(2048 bits)
  Keypair Location: Keyppair generated locally
```

Release Information

Command modified in Junos OS Release 22.3R1.

Support for `idev-id` option added in Junos OS Release 24.2R1

RELATED DOCUMENTATION

Multinode High Availability

request security pki node-local local-certificate verify

request security pki node-local local-certificate re-enroll

request security pki node-local local-certificate load

request security pki node-local local-certificate export

request security pki node-local local-certificate enroll

PKI Components In Junos OS

show security pki statistics

IN THIS SECTION

- Syntax | 25009
- Description | 25009
- Options | 25010
- Required Privilege Level | 25010
- Output Fields | 25010
- Sample Output | 25023
- Release Information | 25027

Syntax

```
show security pki statistics
```

Description

Display standard PKI statistics.

Options

None

Required Privilege Level

view

Output Fields

Table 1319 on page 25010 lists the output fields for the `show security ipsec statistics` command.

Table 1319: show security ipsec statistics Output Fields

Field Name	Field Description
`iked_msgs_inv`	Invalid messages from `iked` process.
`iked_msgs_rxd`	Messages received from `iked` process.
`iked_msgs_txd`	Messages sent to `iked` process.
`cc_kp_req`	Certificate chain keypair requests.
`cc_kp_success`	Certificate chain keypair success.
`cc_kp_fail`	Certificate chain keypair fails (counter of no of certificate key-pair get failure).
`cc_id_ip`	Peer ID type is IP.
`cc_id_dn`	Peer ID type is DN (Domain Name).
`cc_id_fqdn`	Peer ID type is FQDN(Fully Qualified Domain Name).
`cc_id_user_fqdn`	User ID type is FQDN.
`cc_verify_req`	Number of certificate chain verification requests.
`cc_verify_success`	Number of successful certificate verifications.
`cc_verify_fail`	Number of failed certificate verifications

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
cc_inv_ids	IKE IDs did not match EE sub-alt-name .
cc_inv_cert_count	Invalid number of CA's in the certificate request.
ocsp_requests_duplicate	OCSP duplicate requests.
ocsp_requests_sent	OCSP requests sent.
ocsp_resp_success	Successful OCSP response.
ocsp_resp_timeout	OCSP response timed out.
ocsp_action_fail	OCSP next action failed on connection failure.
ocsp_get_req_fail	Failed to get OCSP request for a certificate.
ocsp_resp_malformed_req	Malformed OCSP response.
ocsp_resp_internal_error	OCSP response has an internal error.
ocsp_this_update_failed	OCSP response is not valid yet.
ocsp_next_update_failed	Invalid next update time in OCSP response.
ocsp_resp_try_later	Busy OCSP responder or server. Try again later.
ocsp_resp_sign_required	OCSP responder requires signed request.
ocsp_sign_verify_failed	OCSP responder signature verification failed.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
ocsp_http_parse_error	HTTP parsing error for OCSP response.
ocsp_missing_cert_id	OCSP response does not have responses for given certificate.
ocsp_resp_unauthorized	The OCSP responder does not accept requests from unauthorized clients.
ocsp_rev_status_success	OCSP certificate revocation check success.
ocsp_rev_status_revoked	OCSP certificate is revoked.
ocsp_rev_status_unknown	OCSP certificate revocation status is unknown.
ocsp_nonce_check_failed	Nonce check failed for OCSP responder.
crl_entries_created	Number of CRL entry created.
crl_entries_deleted	Number of CRL entry deleted.
mem_alloc_fails	Memory allocation failure.
crl_requests_sent	Number of CRL requests sent.
crl_responses_rcd	Number of CRL responses received.
crl_download_stop	Number of CRL downloads stopped.
crl_timer_start	Number of times CRL timer started.
crl_timer_stop	Number of times CRL timer stopped.
crl_revoked_certs	Number of times certificates revoked due to CRl check.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
crl_revoke_skip	Number of times CRL revocation check is skipped.
crl_larger_size	Received large CRL file greater than maximum file size limit.
crl_download_failed	Number of CRL download failures.
crl_mem_alloc_fails	Number of CRL entry memory allocation failures.
crl_timer_mem_alloc_fails	Number of CRL timer memory allocation failures.
cmpv2_resp_invalid	Number of Invalid CMPv2 responses.
cmpv2_resp_invalid_status	Failed to get valid CMPv2 response.
cmpv2_resp_http_failed	HTTP parsing failed for CMPv2 response.
cmpv2_resp_validation_failed	Number of CMPv2 response validation failures.
cmpv2_resp_null	Number of NULL CMPv2 response received.
cmpv2_resp_ca_cert_validation_failed	Number of CMPv2 CA certificate validation success.
cmpv2_resp_kup_ca_cert_missing	CA certificate not found to validate CMPv2 response.
cmpv2_resp_kup_ee_cert_missing	EE or local certificate not found to validate CMPv2 response.
cmpv2_resp_null_poll_resp	CMPv2 poll-response is null.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
cmpv2_resp_no_trusted_ca	Trusted CA is not available to validate received CA in CMPv2 response.
cmpv2_resp_success	Received valid CMPv2 response.
cmpv2_ctx_set_caPubs_failed	Failed to set ca-certificates received flag in CMPv2 context.
cmpv2_ctx_set_extraCerts_failed	Failed to set extraCerts field in CMPv2 context.
cmpv2_load_local_failed	CMPv2 local certificate load has failed.
cmpv2_load_ca_failed	CMPv2 CA certificate load has failed.
cmpv2_poll_reached_max_retries	No response from CMPv2 server after maximum configured retries.
cmpv2_send_req_failed	Failed to send CMPv2 requests.
cmpv2_resp_nonce_check_failed	CMPv2 responder nonce check failed.
cmpv2_resp_stack_missing_issuer	Failed to get Issuer certificate for CMPv2 local certificate.
cmpv2_enroll_keypair_missing	CMPv2 Keypair does not exist for certificate.
cmpv2_auto_reenroll_new_keypair_missing	New key missing during CMPv2 auto-reenrollment.
cmpv2_auto_reenroll_keypair_missing	Key pair missing during CMPv2 auto-reenrollment.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
cmpv2_auto_reenroll_cert_missing	Local certificate is missing during CMPv2 auto-reenrollment.
cmpv2_auto_reenroll_ca_profile_missing	CA profile configuration missing during CMPv2 auto-reenrollment.
cmpv2_send_http_req_failed	Failed to send CMPv2 HTTP request.
cmpv2_context_init_failed	CMPv2 context initialization failed.
cmpv2_context_search_failed	CMpv2 context search failed.
cmpv2_context_search_invalid_input	CMpv2 context search failed: due to invalid inputs.
cmpv2_context_create_invalid_input	CMPv2 context creation failed due to invalid inputs.
cmpv2_context_create_context_exists	CMPv2 context creation failed as CMPv2 context already exists.
cmpv2_context_freed	CMPv2 context freed.
cmpv2_gen_http_req_i2d_failed:	CMPv2 message into DER format failed.
cmpv2_gen_http_req_invalid_pkt_len	CMPv2 HTTP request length is invalid.
cmpv2_gen_http_req_failed	Failed to generate CMPv2 HTTP request
cmpv2_gen_http_req_invalid_msg_len	Failed to generate CMPv2 HTTP request: invalid message length.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
cmpv2_search_timer_invalid_input	Failed to get CMPv2 timer entry: invalid input.
cmpv2_search_timer_failed	Failed to get CMPv2 timer entry.
cmpv2_stop_timer_failed	Failed to stop CMPv2 timer.
cmpv2_start_timer_failed	Failed to start CMPv2 timer.
cmpv2_send_message_failed	Failed to send CMPv2 request to server.
cmpv2_connection_failed	Failed to connect to CMPv2 server.
mem_alloc_failed	pkid_malloc - failed to allocate memory.
mem_alloc_type_invalid	pkid_malloc - invalid type parameter.
mem_free_type_invalid	pkid_free - invalid type parameter.
mem_free_alloc_external	pkid_free - not allocated by pkid_malloc.
ldap_state_pending_release	Pending LDAP state.
ldap_state_released	LDAP state is released or freed.
scep_state_pending_release	LDAP state needs to be released.
scep_state_released	SCEP state structure released or freed.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
scep_state_pkey3_initialised	SCEP state keypair initialized.
scep_state_pkey3_added	Added SCEP state keypair.
scep_state_pkey3_deleted	Deleted SCEP state keypair.
scep_ca_query_send_fail	Failed to send SCEP request to server.
scep_x509_lu_ca_obj_case:	Received SCEP CA certificate case.
scep_x509_lu_pkey_rs_ds_obj_case	Received SCEP keypair case.
scep_err_p_subject_is_null	Missing subject in SCEP cert request.
scep_p_err_keypair_is_null	Keypair missing for certificate during SCEP process.
scep_free_cert_req	Freed SCEP certificate request.
scep_reenroll_free_cert_req_info	Freed SCEP certificate request information during SCEP re-enrollment.
crl_state_pending_release	SCEP CRL check pending.
crl_state_released	SCEP CRL state freed.
ca_cert_issuer_verification_fail	Failed to CA certificate for given CA.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
ae_cn_for_ca_cert_fail	Failed to get CA name for given CA certificate.
ae_cn_for_local_cert_fail	Failed to get CA name for given local certificate.
ae_get_cert_dn_fail	Failed to get subject DN field for given certificate id.
ae_x509_issuer_fail	Failed to get issuer certificate for given local certificate.
tpm_ae_key_null	TPM key is missing.
tpm_ae_key_gen_fail	TPM key generation failed.
tpm_key_gen_failure_uncaught	TPM key generation failure not captured.
pkid_db_open	PKI configuration DB is opened.
pkid_db_close	PKI configuration DB is closed
pkid_db_close_fail	Failed to close PKI configured DB.
tpm_ae_success_failure	TPM: failed to store keypair to file.
tpm_pkid_opendir_fail	Failed to open keypair directory in case of TPM.
hsm_session_create_success	HSM session creation success.
hsm_session_create_failure	HSM session creation failure.
hsm_key_create_success	HSM key creation success.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
hsm_key_create_failure	HSM key creation failed.
hsm_key_sign_success	HSM signature sign success.
hsm_key_sign_failure	HSM signature sign failed.
hsm_cert_sign_verify_success	HSM signature verification success.
hsm_cert_sign_verify_failure	HSM signature verification failed.
hsm_pki_to_ike_success	HSM keypair sent to iked process.
hsm_pki_to_ike_failure	HSM keypair sent to IKED failed.
hsm_key_sign_verify_failure	HSM: private key signing failed at HSM.
hsm_function_initialize_failure:	HSM initialization function failed.
hsm_pub_key_retrieval_failure	HSM failed to retrieve public key.
hsm_cleanup_failure	HSM failed to cleanup data structures.
hsm_session_sign_re_create_success	Re-create HSM signature for given session.
hsm_session_sign_re_create_failure	Re-create HSM signature for the given session failed.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
hsm_ss_key_sign_success	HSM self-signed key signature success.
hsm_ss_key_sign_failure	HSM self-signed key signature failure.
hsm_ae_local_cert_delete_failure	HSM local certificate deletion failure.
hsm_ae_local_cert_verif_failure	HSM local certificate verification failure.
hsm_ss_cert_load_failure	HSM failed to load the self-signed certificate.
hsm_dummy_key_delete_fail	HSM failed to create dummy keypair.
pkid_ha_file_replicate_fail	HSM failed to copy file to other node.
pkid_mnha_ae_cert_load_fail	MNHA certificate load failed.
pkid_mnha_ae_cert_verification_fail	MNHA certificate verification failed.
mnha_file_sync_fail	MNHA failed to synchronize file to other node.
kqueue_init_error	kqueue initialization failure.
kqueue_cacert_hash_alloc_fail	kqueue failed to generate memory for CA certificate hash.
kqueue_cacert_file_open_fail	kqueue: failed to open CA certificate file.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
kqueue_cacert_start_fail	kqueue failed.
kqueue_cacert_kevent_fail	kqueue: failed to add kevent.
kqueue_cacert_handler_register_fail	kqueue: CA certificate handler function failed.
kqueue_cacrl_hash_alloc_fail	kqueue: failed to allocate memory for CRL hash.
kqueue_cacrl_file_open_fail	kqueue: failed to open CRL file.
kqueue_cacrl_start_fail	kqueue: failed to get CRL.
kqueue_cacrl_kevent_fail	kqueue: failed to add kevent for CRL.
kqueue_cacrl_handler_register_fail	kqueue: CRL handler function failed.
kqueue_untrusted_ca_hash_alloc_fail	kqueue: failed to allocate memory for untrusted CA certificate hash.
kqueue_untrusted_ca_file_open_fail	kqueue: failed to open untrusted CA certificate file.
kqueue_untrusted_ca_start_fail	kqueue failed for untrusted CA certificate.
kqueue_untrusted_ca_kevent_fail	kqueue failed to add untrusted CA certificate event .
kqueue_untrusted_ca_handler_register_fail	kqueue: untrusted CA handler function failed.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
kqueue_eecert_hash_alloc_fail	kqueue: failed to allocate memory for local certificate hash.
kqueue_eecert_file_open_fail	kqueue: failed to open local certificate file.
kqueue_eecert_start_fail	kqueue: failed to get local certificate.
kqueue_eecert_kevent_fail	kqueue failed to add local certificate event.
kqueue_eecert_handler_register_fail	kqueue: local certificate handler function failed.
kqueue_key_hash_alloc_fail	kqueue: failed to allocate memory for keypair hash.
kqueue_key_file_open_fail	kqueue: failed to open keypair file.
kqueue_key_start_fail	kqueue: failed to get keypair.
kqueue_key_kevent_fail	kqueue failed to add keypair kevent.
kqueue_key_handler_register_fail	kqueue: keypair handler function failed.
pkid_certchain_cacert_fail	Cannot find the signing certificate in the certificate store.
pkid_certs_less_than_min	The chain has less than two certificates. A chain must contain a minimum of two certificates.
pkid_untrust_certs_less_than_min	The untrusted certificate chain has less than two certificates.

Table 1319: show security ipsec statistics Output Fields *(Continued)*

Field Name	Field Description
pkid_ocsp_cert_issuer_null	OCSP failed to get the certificate issuer name.

Sample Output

show security ipsec statistics (MX240, MX480, MX960, SRX Series Firewalls and vSRX Virtual Firewall)

```
user@host> show security ipsec statistics
Statistic Name                 Value
--------------                 -----
iked_msgs_inv                    0
iked_msgs_rxd                    1862
iked_msgs_txd                    1869
cc_kp_req                        1862
cc_kp_success                    0
cc_kp_fail                       1862
cc_id_ip                         0
cc_id_dn                         0
cc_id_fqdn                       0
cc_id_user_fqdn                  0
cc_verify_req                    0
cc_verify_success                0
cc_verify_fail                   0
cc_inv_ids                       0
cc_inv_cert_count                0
ocsp_requests_duplicate          0
ocsp_requests_sent               0
ocsp_resp_success                0
ocsp_resp_timeout                0
ocsp_action_fail                 0
ocsp_get_req_fail                0
ocsp_resp_malformed_req          0
ocsp_resp_internal_error         0
ocsp_this_update_failed          0
ocsp_next_update_failed          0
```

```
ocsp_resp_try_later              0
ocsp_resp_sign_required          0
ocsp_sign_verify_failed          0
ocsp_http_parse_error            0
ocsp_missing_cert_id             0
ocsp_resp_unauthorized           0
ocsp_rev_status_success          0
ocsp_rev_status_revoked          0
ocsp_rev_status_unknown          0
ocsp_nonce_check_failed          0
crl_entries_created              0
crl_entries_deleted              0
mem_alloc_fails                  0
crl_requests_sent                0
crl_responses_rcd                0
crl_download_stop                0
crl_timer_start                  0
crl_timer_stop                   0
crl_revoked_certs                1
crl_revoke_skip                  0
crl_larger_size                  0
crl_download_failed              0
crl_mem_alloc_fails              0
crl_timer_mem_alloc_fails        0
cmpv2_resp_invalid               0
cmpv2_resp_invalid_status        0
cmpv2_resp_http_failed           0
cmpv2_resp_validation_failed     0
cmpv2_resp_null                  0
cmpv2_resp_ca_cert_validation_failed 0
cmpv2_resp_kup_ca_cert_missing   0
cmpv2_resp_kup_ee_cert_missing   0
cmpv2_resp_null_poll_resp        0
cmpv2_resp_no_trusted_ca         0
cmpv2_resp_success               0
cmpv2_ctx_set_caPubs_failed      0
cmpv2_ctx_set_extraCerts_failed  0
cmpv2_load_local_failed          0
cmpv2_load_ca_failed             0
cmpv2_poll_reached_max_retries   0
cmpv2_send_req_failed            0
cmpv2_resp_nonce_check_failed    0
cmpv2_resp_stack_missing_issuer  0
```

```
cmpv2_enroll_keypair_missing     0
cmpv2_auto_reenroll_new_keypair_missing 0
cmpv2_auto_reenroll_keypair_missing 0
cmpv2_auto_reenroll_cert_missing 0
cmpv2_auto_reenroll_ca_profile_missing 0
cmpv2_send_http_req_failed       0
cmpv2_context_init_failed        0
cmpv2_context_search_failed      0
cmpv2_context_search_invalid_input 0
cmpv2_context_create_invalid_input 0
cmpv2_context_create_context_exists 0
cmpv2_context_freed              0
cmpv2_gen_http_req_i2d_failed    0
cmpv2_gen_http_req_invalid_pkt_len 0
cmpv2_gen_http_req_failed        0
cmpv2_gen_http_req_invalid_msg_len 0
cmpv2_search_timer_invalid_input 0
cmpv2_search_timer_failed        0
cmpv2_stop_timer_failed          0
cmpv2_start_timer_failed         0
cmpv2_send_message_failed        0
cmpv2_connection_failed          0
cmpv2_ee_cert_get_keypair_failed 0
mem_alloc_failed                 0
mem_alloc_type_invalid           0
mem_free_type_invalid            0
mem_free_alloc_external          0
ldap_state_pending_release       0
ldap_state_released              0
scep_state_pending_release       0
scep_state_released              0
scep_state_pkey3_initialised     0
scep_state_pkey3_added           0
scep_state_pkey3_deleted         0
scep_ca_query_send_fail          0
scep_x509_lu_ca_obj_case         0
scep_x509_lu_pkey_rs_ds_obj_case 0
scep_err_p_subject_is_null       0
scep_p_err_keypair_is_null       0
scep_free_cert_req               0
scep_reenroll_free_cert_req_info 0
crl_state_pending_release        0
crl_state_released               0
```

```
ca_cert_issuer_verification_fail 0
ae_cn_for_ca_cert_fail           0
ae_cn_for_local_cert_fail        0
ae_get_cert_dn_fail              0
ae_x509_issuer_fail              0
tpm_ae_key_null                  0
tpm_ae_key_gen_fail              0
tpm_key_gen_failure_uncaught     0
pkid_db_open                     7
pkid_db_close                    7
pkid_db_close_fail               0
tpm_ae_success_failure           0
tpm_pkid_opendir_fail            0
hsm_session_create_success       0
hsm_session_create_failure       0
hsm_key_create_success           0
hsm_key_create_failure           0
hsm_key_sign_success             0
hsm_key_sign_failure             0
hsm_cert_sign_verify_success     0
hsm_cert_sign_verify_failure     0
hsm_pki_to_ike_success           0
hsm_pki_to_ike_failure           0
hsm_key_sign_verify_failure      0
hsm_function_initialize_failure  0
hsm_pub_key_retrieval_failure    0
hsm_cleanup_failure              0
hsm_session_sign_re_create_success 0
hsm_session_sign_re_create_failure 0
hsm_ss_key_sign_success          0
hsm_ss_key_sign_failure          0
hsm_ae_local_cert_delete_failure 0
hsm_ae_local_cert_verif_failure  0
hsm_ss_cert_load_failure         0
hsm_dummy_key_delete_fail        0
pkid_ha_file_replicate_fail      0
pkid_mnha_ae_cert_load_fail      0
pkid_mnha_ae_cert_verification_fail 0
mnha_file_sync_fail              0
kqueue_init_error                0
kqueue_cacert_hash_alloc_fail    0
kqueue_cacert_file_open_fail     0
kqueue_cacert_start_fail         0
```

```
kqueue_cacert_kevent_fail        0
kqueue_cacert_handler_register_fail 0
kqueue_cacrl_hash_alloc_fail     0
kqueue_cacrl_file_open_fail      0
kqueue_cacrl_start_fail          0
kqueue_cacrl_kevent_fail         0
kqueue_cacrl_handler_register_fail 0
kqueue_untrusted_ca_hash_alloc_fail 0
kqueue_untrusted_ca_file_open_fail 0
kqueue_untrusted_ca_start_fail   0
kqueue_untrusted_ca_kevent_fail  0
kqueue_untrusted_ca_handler_register_fail 0
kqueue_eecert_hash_alloc_fail    0
kqueue_eecert_file_open_fail     0
kqueue_eecert_start_fail         0
kqueue_eecert_kevent_fail        0
kqueue_eecert_handler_register_fail 0
kqueue_key_hash_alloc_fail       0
kqueue_key_file_open_fail        0
kqueue_key_start_fail            0
kqueue_key_kevent_fail           0
kqueue_key_handler_register_fail 0
pkid_certchain_cacert_fail       0
pkid_certs_less_than_min         0
pkid_untrust_certs_less_than_min 0
pkid_ocsp_cert_issuer_null       0
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

clear security pki statistics | 4033

show security policies

IN THIS SECTION

- Syntax | 25028
- Description | 25029
- Options | 25029
- Required Privilege Level | 25030
- Output Fields | 25030
- Sample Output | 25036
- Release Information | 25055

Syntax

```
show security policies
<all-logical-systems-tenants>
<checksum>
<count>
<detail>
<from-zone zone-name>
<global>
<hit-count>
<information>
<logical-system logical-system-name>
<policy-name policy-name>
<root-logical-system>
<service-set>
<start>
<tenant tenant-name>
<to-zone zone-name>
<unknown-source-identity>
<zone-context>
```

Description

Displays a summary of all security policies configured on the device. If a particular policy is specified, display information specific to that policy. The existing show commands for displaying the policies configured with multiple tenant support are enhanced. A security policy controls the traffic flow from one zone to another zone. The security policies allow you to deny, permit, reject (deny and send a TCP RST or ICMP port unreachable message to the source host), encrypt and decrypt, authenticate, prioritize, schedule, filter, and monitor the traffic attempting to cross from one security zone to another.

Options

- `all-logical-systems-tenants`—Displays all multitenancy systems.
- `checksum`—Displays the policy information checksum.
- `count`—Displays the number of policies to show. Range is 1 through 65,535.
- `detail`—(Optional) Displays a detailed view of all of the policies configured on the device.
- `from-zone`—Displays the policy information matching the given source zone.
- `global`—(Optional) Displays the policy information about global policies.
- `hit-count`—Displays the policies hit count.
- `information`—Displays the policy information.
- `logical-system`—Displays the logical system name.
- `policy-name`—(Optional) Displays the policy information matching the given policy name.
- `root-logical-system`—Displays root logical system as default.
- `service-set`—Displays the name of the service set.
- `start`—Displays the policies from a given position. Range is 1 through 65,535.
- `tenant`—Displays the name of the tenant system.
- `to-zone`—Displays the policy information matching the given destination zone.
- `unknown-source-identity`—Displays the unknown-source-identity of a policy.
- `zone-context`—Displays the count of policies in each context (from-zone and to-zone).

Required Privilege Level

view

Output Fields

Table 1320 on page 25030 lists the output fields for the `show security policies` command. Output fields are listed in the approximate order in which they appear.

Table 1320: show security policies Output Fields

Field Name	Field Description
`From zone`	Name of the source zone.
`To zone`	Name of the destination zone.
`Policy-name`	Name of the applicable policy.
`Description`	Description of the applicable policy.
`State`	Status of the policy: • `enabled`: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • `disabled`: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
`Index`	Internal number associated with the policy.
`Sequence number`	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, 4.

Table 1320: show security policies Output Fields *(Continued)*

Field Name	Field Description
`Source addresses`	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names. For detail display mode, the names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
`Destination addresses`	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
`source-end-user-profile`	Name of the device identity profile (referred to as `end-user-profile` in the CLI) that contains attributes, or characteristics of a device. Specification of the device identity profile in the `source-end-user-profile` field is part of the device identity feature. If a device matches the attributes specified in the profile and other security policy parameters, then the security policy's action is applied to traffic issuing from the device.
`Source addresses (excluded)`	Name of the source address excluded from the policy.
`Destination addresses (excluded)`	Name of the destination address excluded from the policy.
`Source identities`	One or more user roles specified for a policy.

Table 1320: show security policies Output Fields *(Continued)*

Field Name	Field Description
`Applications`	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • `IP protocol`: The Internet protocol used by the application—for example, TCP, UDP, ICMP. • `ALG`: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If `application-protocol ignore` is configured, `ignore` is displayed. Otherwise, `0` is displayed. However, even if this command shows `ALG: 0`, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when `application-protocol ignore` is not configured for custom applications. • `Inactivity timeout`: Elapsed time without activity after which the application is terminated. • `Source port range`: The low-high source port range for the session application.
`Source identity feeds`	Name of a source identity (user name) added as match criteria
`Destination identity feeds`	Name of a destination identity (user name) added as match criteria
`Dynamic Applications`	Application identification-based Layer 7 dynamic applications.
`Destination Address Translation`	Status of the destination address translation traffic: • `drop translated`—Drop the packets with translated destination addresses. • `drop untranslated`—Drop the packets without translated destination addresses.

Table 1320: show security policies Output Fields *(Continued)*

Field Name	Field Description

`Application Firewall`	An application firewall includes the following:
	• `Rule-set`—Name of the rule set.
	• `Rule`—Name of the rule.
	  • `Dynamic applications`—Name of the applications.
	  • `Dynamic application groups`—Name of the application groups.
	  • `Action`—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  • `permit`
	  • `deny`
	• `Default rule`—The default rule applied when the identified application is not specified in any rules of the rule set.

Table 1320: show security policies Output Fields *(Continued)*

Field Name	Field Description
`Action or Action-type`	• The action taken for a packet that matches the policy's tuples. Actions include the following: • `permit` • `feed` • `firewall-authentication` • `tunnel ipsec-vpn` *`vpn-name`* • `pair-policy` *`pair-policy-name`* • `source-nat pool` *`pool-name`* • `pool-set` *`pool-set-name`* • `interface` • `destination-nat` *`name`* • `deny` • `reject` • `services-offload`
`Session log`	Session log entry that indicates whether the `at-create` and `at-close` flags were set at configuration time to log session information.
`Scheduler name`	Name of a preconfigured scheduler whose schedule determines when the policy is active and can be used as a possible match for traffic.

Table 1320: show security policies Output Fields *(Continued)*

Field Name	Field Description

`Policy statistics`	• `Input bytes`—The total number of bytes presented for processing by the device.
	  • `Initial direction`—The number of bytes presented for processing by the device from the initial direction.
	  • `Reply direction`—The number of bytes presented for processing by the device from the reply direction.
	• `Output bytes`—The total number of bytes actually processed by the device.
	  • `Initial direction`—The number of bytes from the initial direction actually processed by the device.
	  • `Reply direction`—The number of bytes from the reply direction actually processed by the device.
	• `Input packets`—The total number of packets presented for processing by the device.
	  • `Initial direction`—The number of packets presented for processing by the device from the initial direction.
	  • `Reply direction`—The number of packets presented for processing by the device from the reply direction.
	• `Output packets`—The total number of packets actually processed by the device.
	  • `Initial direction`—The number of packets actually processed by the device from the initial direction.
	  • `Reply direction`—The number of packets actually processed by the device from the reply direction.
	• `Session rate`—The total number of active and deleted sessions.
	• `Active sessions`—The number of sessions currently present because of access control lookups that used this policy.
	• `Session deletions`—The number of sessions deleted since system startup.
	• `Policy lookups`—The number of times the policy was accessed to check for a match.

Table 1320: show security policies Output Fields *(Continued)*

Field Name	Field Description
dynapp-redir-profile	Displays unified policy redirect profile. See "profile(dynamic-application)" on page 13275.
Per policy TCP Options	Configured syn and sequence checks, and the configured TCP MSS value for the initial direction, the reverse direction or, both.
Feed	Feeds details added in the security policy. The supported feeds are: • add-source-ip-to-feed • add-destination-ip-to-feed • add-source-identity-to-feed • add-destination-identity-to-feed

Sample Output

show security policies

```
user@host> show security policies

  From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Sequence number: 1
    Source addresses:
    sa-1-ipv4: 198.51.100.11/24
    sa-2-ipv6: 2001:db8:a0b:12f0::1/32
    sa-3-ipv6: 2001:db8:a0b:12f0::22/32
    sa-4-wc:   203.0.113.1/255.255.0.255
    Destination addresses:
    da-1-ipv4: 10.2.2.2/24
    da-2-ipv6: 2001:db8:a0b:12f0::8/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
    Source identities: role1, role2, role4
    Applications: any
```

```
    Action: permit, application services, log, scheduled
    Application firewall : my_ruleset1
  Policy: p2, State: enabled, Index: 5, Sequence number: 2
    Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::22/32
    Destination addresses:
    da-1-ipv4:  10.2.2.2/24
    da-2-ipv6:  2001:db8:a0b:12f0::1/32
    da-3-ipv6:  2001:db8:a0b:12f0::9/32
    Source identities: role1, role4
    Applications: any
    Action: deny, scheduled
```

show security policies (Dynamic Applications)

```
user@host>show security policies

  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: junos:YAHOO
    Action: deny, log
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 2
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: junos:web, junos:web:social-networking:facebook,
    junos:TFTP, junos:QQ
    Action: permit, log
  Policy: p3, State: enabled, Index: 6, Scope Policy: 0, Sequence number: 3
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: junos:HTTP, junos:SSL
    Action: permit, application services, log
```

The following example displays the output with unified policies configured.

```
user@host> show security policies

Default policy: deny-all
Pre ID default policy: permit-all
From zone: trust, To zone: untrust
  Policy: p2, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Applications: junos-defaults
    Dynamic Applications: junos:GMAIL, junos:FACEBOOK-CHAT
    dynapp-redir-profile: profile1
```

show security policies policy-name p2

```
user@host> show security policies policy-name p2
  Policy: p2, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    From zones: any
    To zones: any
    Source vrf group: any
    Destination vrf group: any
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: any
    Action: permit, application services, feed
```

show security policies policy-name detail

```
user@host> show security policies policy-name p2 detail

Policy: p2, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured, global
  Sequence number: 1
  From zones:
    any
  To zones:
    any
```

```
  Source vrf group:
    any
  Destination vrf group:
    any
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination ports: [0-0]
  Dynamic Application:
    any: 0
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
  Intrusion Detection and Prevention: disabled
  Unified Access Control: disabled
  Feed: add-source-ip-to-feed

user@host> show security policies policy-name p1 detail

  Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::9/32
    sa-4-wc:    203.0.113.1/255.255.0.255
  Destination addresses:
    da-1-ipv4: 192.0.2.0/24
    da-2-ipv6: 2001:db8:a0b:12f0::1/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
    Source identities:
        role1
        role2
        role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
```

```
      Source port range: [0-0]
      Destination port range: [0-0]
  Destination Address Translation: drop translated
  Application firewall :
     Rule-set: my_ruleset1
       Rule: rule1
         Dynamic Applications: junos:FACEBOOK-ACCESS, junos:YMSG
         Dynamic Application groups: junos:web, junos:chat
         Action: deny
       Default rule: permit
  Session log: at-create, at-close
  Scheduler name: sch20
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :                18144                  545 bps
      Initial direction:                 9072                    272 bps
      Reply direction  :                 9072                  272 bps
    Output bytes       :                18144                  545 bps
      Initial direction:                 9072                  272 bps
      Reply direction  :                 9072                  272 bps
    Input  packets     :                  216                    6 pps
      Initial direction:                  108                    3 bps
      Reply direction  :                  108                    3 bps
    Output packets     :                  216                    6 pps
      Initial direction:                  108                    3 bps
      Reply direction  :                  108                    3 bps
    Session rate       :                  108                    3 sps
    Active sessions    :                   93
    Session deletions  :                   15
    Policy lookups     :                  108
```

show security policies (Services-Offload)

```
user@host> show security policies

  Policy: p1, action-type: reject, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
```

```
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  dynapp-redir-profile: profile1(1)
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security policies (Device Identity)

```
user@host> show security policies
From zone: trust, To zone: untrust
       Policy: dev-id-marketing, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
         Source addresses: any
         Destination addresses: any
         source-end-user-profile: marketing-profile
         Applications: any
         Action: permit
```

show security policies detail

```
user@host> show security policies detail

Default policy: deny-all
Policy: p1, action-type: permit, services-offload:enabled , State: enabled, Index: 4, Scope
Policy: 0
  Policy Type: Configured
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global):  0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global):  0.0.0.0/0
    any-ipv6(global): ::/0
    Source identities:
```

```
        role1
        role2
        role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :                18144                  545 bps
      Initial direction:                 9072                    272 bps
      Reply direction  :                 9072                  272 bps
    Output bytes       :                18144                  545 bps
      Initial direction:                 9072                  272 bps
      Reply direction  :                 9072                  272 bps
    Input  packets     :                  216                    6 pps
      Initial direction:                  108                    3 bps
      Reply direction  :                  108                    3 bps
    Output packets     :                  216                    6 pps
      Initial direction:                  108                    3 bps
      Reply direction  :                  108                    3 bps
    Session rate       :                  108                    3 sps
    Active sessions    :                   93
    Session deletions  :                   15
    Policy lookups     :                  108
Policy: p2, action-type: permit, services-offload:enabled , State: enabled, Index: 5, Scope
Policy: 0
  Policy Type: Configured
  Description: The policy p2 is for the sales team
  Sequence number: 1
  From zone: untrust, To zone: trust
  Source addresses:
    any-ipv4(global):  0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global):  0.0.0.0/0
    any-ipv6(global): ::/0
    Source identities:
        role1
        role2
        role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
```

```
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

The following example displays the output with unified policies configured.

```
user@host> show security policies detail

Default policy: deny-all
Pre ID default policy: permit-all
Policy: p2, action-type: reject, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: junos-defaults
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [443-443]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [5432-5432]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [80-80]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [3128-3128]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [8000-8000]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [8080-8080]
    IP protocol: 17, ALG: 0, Inactivity timeout: 60
      Source port range: [0-0]
      Destination port range: [1-65535]
```

```
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [443-443]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [5432-5432]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [80-80]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [3128-3128]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [8000-8000]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [8080-8080]
    IP protocol: 17, ALG: 0, Inactivity timeout: 60
      Source port range: [0-0]
      Destination port range: [1-65535]
  Dynamic Application:
    junos:FACEBOOK-CHAT: 10704
    junos:GMAIL: 51
  dynapp-redir-profile: profile1(1)
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security policies detail (TCP Options)

```
user@host> show security policies policy-name p2 detail
node0:
--------------------------------------------------------------------------
Policy:p2, action-type:permit, State: enabled,Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
```

```
    any-ipv6(global): ::/0
  Application: junos-defaults
    IP protocol: tcp, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
    Destination port range: [80-80]
Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
  Dynamic-application: junos:HTTP
```

show security policies policy-name (Negated Address)

```
user@host> show security policies policy-name p1
node0:
--------------------------------------------------------------------------
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses(excluded): as1
    Destination addresses(excluded): as2
    Applications: any
    Action: permit
```

show security policies policy-name detail (Negated Address)

```
user@host> show security policies policy-name p1 detail
node0:
--------------------------------------------------------------------------
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses(excluded):
    ad1(ad): 255.255.255.255/32
    ad2(ad):  198.51.100.1/24
    ad3(ad):  198.51.100.6  ~ 198.51.100.56
    ad4(ad):  192.0.2.8/24
    ad5(ad):  198.51.100.99 ~ 198.51.100.199
    ad6(ad):  203.0.113.9/24
    ad7(ad):  203.0.113.23/24
  Destination addresses(excluded):
    ad13(ad2): 198.51.100.76/24
    ad12(ad2): 198.51.100.88/24
```

```
    ad11(ad2): 192.0.2.23 ~ 192.0.2.66
    ad10(ad2): 192.0.2.93
    ad9(ad2):  203.0.113.76 ~ 203.0.113.106
    ad8(ad2):  203.0.113.199
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies global

```
user@host> show security policies global policy-name Pa
node0:
--------------------------------------------------------------------------
  Global policies:
  Policy: Pa, State: enabled, Index: 6, Scope Policy: 0, Sequence number: 1
    From zones: any
    To zones: any
    Source addresses: H0
    Destination addresses: H1
    Applications: junos-http
    Action: permit
```

show security policies detail tenant

```
user@host> show security policies detail tenant TN1

Default policy: deny-all
Pre ID default policy: permit-all
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
Policy Type: Configured
Sequence number: 1
From zone: trust, To zone: untrust
Source addresses: any
Destination addresses: any
Application: junos-ping
IP protocol: 1, ALG: 0, Inactivity timeout: 60
ICMP Information: type=255, code=0
```

```
Application: junos-telnet
IP protocol: tcp, ALG: 0, Inactivity timeout: 1800
Source port range: [0-0]
Destination port range: [23-23]
Application: app_udp
IP protocol: udp, ALG: 0, Inactivity timeout: 1800
Source port range: [0-0]
Destination port range: [5000-5000]
Application: junos-icmp6-all
IP protocol: 58, ALG: 0, Inactivity timeout: 60
ICMP Information: type=255, code=0
Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
Session log: at-create, at-close
Policy statistics:
Input  bytes     :                    0                    0 bps
Initial direction:                    0                    0 bps
Reply direction  :                    0                    0 bps
Output bytes     :                    0                    0 bps
Initial direction:                    0                    0 bps
Reply direction  :                    0                    0 bps
Input  packets   :                    0                    0 pps
Initial direction:                    0                    0 bps
Reply direction  :                    0                    0 bps
Output packets   :                    0                    0 pps
Initial direction:                    0                    0 bps
Reply direction  :                    0                    0 bps
Session rate     :                    0                    0 sps
Active sessions  :                    0
Session deletions:                    0
Policy lookups   :                    0
```

show security policies (threat profile feeds)

```
user@host> show security policies policy-name p2
From zone: trust, To zone: untrust
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 2
    Source vrf group: any
    Destination vrf group: any
    Source addresses: any
    Destination addresses: any
```

```
    Applications: any
Source identity feeds: user_feed_1, user_feed_2
Destination identity feeds: user_feed_3, user_feed_4
    Action: permit, application services, feed
```

show security policies detail (threat profile feeds)

```
user@host> show security policies policy-name p2 detail
Policy: p2, action-type: permit, State: enabled, Index: 5, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
  From zone: trust, To zone: untrust
  Source vrf group:
    any
  Destination vrf group:
    any
  Source addresses:
    any-ipv4(bob_addrbook_1): 0.0.0.0/0
    any-ipv6(bob_addrbook_1): ::/0
  Destination addresses:
    any-ipv4(bob_addrbook_1): 0.0.0.0/0
    any-ipv6(bob_addrbook_1): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination ports: [0-0]
  Source identity feeds:
user_feed_1
user_feed_2
  Destination identity feeds:
user_feed_3
user_feed_4
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
  Intrusion Detection and Prevention: disabled
  Unified Access Control: disabled
  Feed: add-source-ip-to-feed
  Feed: add-destination-ip-to-feed
  Feed: add-source-identity-to-feed
  Feed: add-destination-identity-to-feed
```

show security policies detail (services-offload enabled)

```
user@host> show security policies detail
Default policy: deny-all
Default policy log Profile ID: 0
Pre ID default policy: permit-all
Policy: SOF-enable, action-type: permit, services-offload:enabled , State: enabled, Index: 5,
Scope Policy: 0
Policy Type: Configured
Sequence number: 1
From zone: trust, To zone: untrust1
Source vrf group:
any
Destination vrf group:
any
Source addresses:
any-ipv4(global): 0.0.0.0/0
any-ipv6(global): ::/0
 Destination addresses:
 any-ipv4(global): 0.0.0.0/0
 any-ipv6(global): ::/0
 Application: any
 IP protocol: 0, ALG: 0, Inactivity timeout: 0
 Source port range: [0-0]
 Destination ports: [0-0]
 Dynamic Application:
 any: 0
 Source identity feeds:
 any
Destination identity feeds:
 any
Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security policies policy-name SOF-enable

```
user@host> show security policies policy-name SOF-enable
From zone: trust, To zone: untrust1
  Policy: SOF-enable, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1, Log Profile
ID: 0
    Source vrf group: any
```

```
    Destination vrf group: any
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: any
    Source identity feeds: any
    Destination identity feeds: any
    Action: permit, services-offload
```

show security policies detail (services-offload disabled)

```
user@host> show security policies detail
Default policy: deny-all
Default policy log Profile ID: 0
Pre ID default policy: permit-all
Policy: SOF-disable, action-type: permit, services-offload:disabled , State: enabled, Index: 5,
Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust1
  Source vrf group:
    any
  Destination vrf group:
    any
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination ports: [0-0]
  Dynamic Application:
    any: 0
  Source identity feeds:
    any
  Destination identity feeds:
```

```
    any
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security policies policy-name SOF-disable

```
user@host> show security policies policy-name SOF-disable
 From zone: trust, To zone: untrust1
  Policy: SOF-disable, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1, Log
Profile ID: 0
    Source vrf group: any
    Destination vrf group: any
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: any
    Source identity feeds: any
    Destination identity feeds: any
    Action: permit, no-services-offload
```

show security policies (destination-identity)

```
user@host> show security policies
Default policy: deny-all
Default policy log Profile ID: 0
Pre ID default policy: permit-all
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1, Log Profile ID: 0
    Source vrf group: any
    Destination vrf group: any
    Source addresses: any
    Destination addresses: any
    Applications: junos-http, junos-https
    Dynamic Applications: junos:HTTP, junos:HTTP-VIDEO, junos:HTTP-AUDIO-CONTENT, junos:BMFF,
junos:SSL
    Source identities: role1, role3
    Source identity feeds: any
    Destination identity context: role2, role4
    Destination identity context profile: hr
```

```
    Destination identity feeds: any
    Action: permit
```

show security policies from-zone trust to-zone untrust detail (destination-identity)

```
user@host> show security policies from-zone trust to-zone untrust detail
Policy: p2, action-type: permit, services-offload:not-configured , State: enabled, Index: 6,
Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source vrf group:
    any
Destination vrf group:
    any
  Source addresses:
    any-ipv4: 0.0.0.0/0
    any-ipv6: ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: junos-http
    IP protocol: tcp, ALG: 0, Inactivity timeout: 300
      Source port range: [0-0]
      Destination ports: 80
  Application: junos-https
    IP protocol: tcp, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination ports: 443
  Application: junos-ssh
    IP protocol: tcp, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination ports: 22
  Source identity feeds:
    any
  Destination identity context:
    role2
    role4
  Destination identity context profile:
    hr
  Destination identity feeds:
```

```
    any
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security policies detail (destination-identity)

```
user@host> show security policies detail
Default policy: deny-all
Default policy log Profile ID: 0
Pre ID default policy: permit-all
Policy: p1, action-type: permit, services-offload:not-configured , State: enabled, Index: 4,
Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source vrf group:
    any
  Destination vrf group:
    any
  Source addresses:
    any-ipv4: 0.0.0.0/0
    any-ipv6: ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: junos-http
    IP protocol: tcp, ALG: 0, Inactivity timeout: 300
      Source port range: [0-0]
      Destination ports: 80
  Application: junos-https
    IP protocol: tcp, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination ports: 443
  Dynamic Application:
    junos:SSL: 199
    junos:BMFF: 1293
    junos:HTTP-AUDIO-CONTENT: 10806
    junos:HTTP-VIDEO: 11032
    junos:HTTP: 67
  Source identities:
    role1
    role3
```

```
  Source identity feeds:
    any
  Destination identity context:
    role2
    role4
  Destination identity context profile:
    hr
  Destination identity feeds:
    any
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security policies global detail (destination-identity)

```
user@host> show security policies global detail
Policy: g1, action-type: reject, services-offload:not-configured , State: enabled, Index: 8,
Scope Policy: 0
  Policy Type: Configured, global
  Sequence number: 1
  From zones:
    any
  To zones:
    any
  Source vrf group:
any
  Destination vrf group:
    any
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: junos-http
    IP protocol: tcp, ALG: 0, Inactivity timeout: 300
      Source port range: [0-0]
      Destination ports: 80
  Application: junos-https
    IP protocol: tcp, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination ports: 443
  Dynamic Application:
```

```
    junos:HTTP: 67
  Source identities:
    unauthenticated-user
    role1
  Source identity feeds:
    any
  Destination identity context:
    role2
    role4
  Destination identity context profile:
    hr
  Destination identity feeds:
    any
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

Release Information

Command modified in Junos OS Release 9.2.

Support for IPv6 addresses is added in Junos OS Release 10.2.

Support for wildcard addresses is added in Junos OS Release 11.1.

Support for global policy and services offloading is added in Junos OS Release 11.4.

Support for source-identities and the `Description` output field is added in Junos OS Release 12.1.

Support for negated address added in Junos OS Release 12.1X45-D10.

The output fields for Policy Statistics expanded, and the output fields for the `global` and `policy-name` options are expanded to include from-zone and to-zone global match criteria in Junos OS Release 12.1X47-D10.

Support for the `initial-tcp-mss` and `reverse-tcp-mss` options is added in Junos OS Release 12.3X48-D20.

Output field and description for `source-end-user-profile` option is added in Junos OS Release 15.1x49-D70.

Output field and description for `dynamic-applications` option is added in Junos OS Release 15.1x49-D100.

Output field and description for `dynapp-redir-profile` option is added in Junos OS Release 18.2R1.

The `tenant` option is introduced in Junos OS Release 18.3R1.

The `<all-logical-systems-tenants>` option is introduced in Junos OS Release 18.4R1.

The `information` option is introduced in Junos OS Release 18.4R1.

The `checksum` option is introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

Unified Policies Configuration Overview

show security policies checksum

IN THIS SECTION

- Syntax | 25056
- Description | 25056
- Options | 25057
- Additional Information | 25057
- Required Privilege Level | 25058
- Sample Output | 25058
- Release Information | 25058

Syntax

```
show security policies checksum
<logical-system (logical-system name | all)>
<root-logical-system>
<tenant tenant-name>
```

Description

Displays the security policy checksum value.

Verifying the checksum helps in validating the security policy sync status between the Routing Engine and the Packet Forwarding Engine. The checksum value should be the same for the Routing Engine and the Packet Forwarding Engine. If the checksum value is not same, then the different values indicates that the security policies on the Routing Engine and the Packet Forwarding Engine are out-of-sync.

> **NOTE**: The `show security policies checksum` command can only be used to ensure if the security policies are out of sync but cannot confirm if they are in-sync. Use the `request security policies check` command to get a list of all polices in-sync and/or out-of-sync. Use the `request security policies resync` command to synchronize the configuration of security policies in the Routing Engine and Packet Forwarding Engine.

Options

logical-system (*logical-system-name* | all)	Displays the security policy checksum value for the security policies configured on a logical system or on all logical systems.
root-logical-system	Displays the security policy checksum value for the security policies configured on the root logical system. This is the default outcome.
tenant *tenant-name*	Displays the security policy checksum value for the security policies configured on a tenant.

Additional Information

The checksum value is a 32-character hexadecimal number that is computed for the security policy on the device.

Security policies are stored in the routing engine and the packet forwarding engine. Security policies are pushed from the Routing Engine to the Packet Forwarding Engine when you commit configurations. If the security policies on the Routing Engine are out of sync with the Packet Forwarding Engine, the commit of a configuration fails. Core dump files may be generated if the commit is tried repeatedly. The out of sync can be due to:

- A policy message from Routing Engine to the Packet Forwarding Engine is lost in transit.
- An error with the routing engine, such as a reused policy UID.

When the policy configuration are modified and the policies are out of sync, the following error message displays - `error: Warning: policy might be out of sync between RE and PFE <SPU-name(s)>. Please request security policies check/resync.`

Required Privilege Level

view

Sample Output

show security policies checksum (RE)

```
user@host> show security policies checksum
Logical system: root-logical-system
    From zone          To zone           Checksum
    trust              untrust           0xe0fc5791-d7ec5b89-cbc66724-35d706c
```

show security policies checksum (PFE)

```
FLOWD_OCTEON(vty)#  show usp policy checksum
Logical system: root-logical-system
    From zone          To zone           Checksum
    trust              untrust           0xe0fc5791-d7ec5b89-cbc66724-35d706c
```

show security policies checksum logical-system all

```
user@host> show security policies checksum logical-system all
Logical system: LSYS1
    From zone          To zone           Checksum
    lsys1-trust        lsys1-untrust     0x7ff147ff-f9df0820-081fb02e-226455e3

Logical system: LSYS2
    From zone          To zone           Checksum
    lsys2-trust        lsys2-untrust     0x4bfd47fa-64fbb3e7-3a9444b5-94dd9db5
```

Release Information

Command introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

request security policies check | 14857

request security policies resync | 14861

show security policies file-serialization counters

IN THIS SECTION

- Syntax | 25059
- Description | 25059
- Required Privilege Level | 25059
- Output Fields | 25059
- Sample Output | 25061
- Release Information | 25061

Syntax

```
show security policies file-serialization counters
```

Description

Displays the file-based configuration and operational status details.

Required Privilege Level

view

Output Fields

Table lists the output fields for the `show security policies file-serialization counters` command. Output fields are listed in the approximate order in which they appear.

Table 1321: show security policies file-serialization counters

Field Name	Field Description
`Status`	Status of file serialization configuration.
`Start Transaction`	Number of file-write transactions started.
`End Transaction`	Number of file-write transactions completed.
`Files sent to dataplane`	Number of file-read messages sent to PFE.
`Ack from dataplane`	Number of file-read done messages received from PFE.
`Total entries written`	Total policy data entries written.
`Total bytes written`	Total bytes written.
`Files deleted`	Number of files deleted.
`Files ageout`	Files aged out as the file TTL exceeded – case of missing ACK from PFE
`File limit reached`	File size limit of 400MB reached and the new rest of the data written to subsequent file.
`Serialization errors`	Errors on Routing Engine while writing – space issues, write errors, data conversion errors.
`Deserialization errors`	Errors on PFE while reading – file open failures, invalid tags, malloc failures, data conversion.
`Total bytes excluded`	Total bytes excluded from the file-serialization.

Sample Output

show security policies file-serialization counters

```
user@host> show security policies file-serialization counters
Policy file-serialization status:
    Status                 :      Enabled

  Policy file-serialization counters:
    Start Transaction                   :            1
    End Transaction                     :            1
    Files sent to dataplane             :            1
    Ack from dataplane                  :            1
    Total entries written               :          147
    Total bytes written                 :        33147
    Files deleted                       :            1
    Files ageout                        :            0
    File limit reached                  :            0
    Serialization errors                :            0
    Deserialization errors              :            0
    Total bytes excluded                :          144
```

Release Information

Command introduced in Junos OS Release 24.4R1

show security policies hit-count

IN THIS SECTION

- Syntax | 25062
- Description | 25062
- Options | 25063
- Required Privilege Level | 25064

- Output Fields | 25064
- Sample Output | 25065
- Sample Output | 25065
- Sample Output | 25065
- Sample Output | 25066
- Sample Output | 25066
- Sample Output | 25066
- Sample Output | 25066
- Sample Output | 25067
- Release Information | 25067

Syntax

```
show security policies hit-count
<ascending>
<descending>
<dynamic-applications>
<from-zone zone-name>
<greater-than count>
<less-than count>
<logical-system (logical-system-name | all)>
<root-logical-system>
<tenant (tenant-name | all)>
<to-zone zone-name>
<url-categories>
```

Description

Display the utility rate of security policies by listing the number of times a security policy rule matches the traffic (number of hits). You can specify the options to list the output in ascending or descending order. You can specify the range to display security policies with certain number of hits. You can filter the output by zones, logical or tenant systems, dynamic applications, and URL categories.

When the device is operating in chassis cluster mode, the count displayed is a sum of all the Services Processing Cards (SPC) hit counts in the cluster setup. The security device retains the count if a Packet Forwarding Engine (PFE) in a node is in failover mode, but does not reboot. . The device clears the count

if a node reboots and the PFE in the node also reboots. During an in-service software upgrade (ISSU), all PFEs reboot, therefore all counters are cleared.

Use this command without options to display the number of hits in random order for all security policies and for all zones.

The `show security policies hit-count` command with `to-zone` and `from-zone` options work for zone-based policies only and not for global policies.

Options

- `ascending`—(Optional) Displays the number of hits for security policies in ascending order.
- `descending`—(Optional) Displays the number of hits for security policies in descending order.
- `dynamic-applications`—(Optional) Displays the number of hits for security policies configured with dynamic applications.

  When you display the policy count for the dynamic applications, the device considers the count for the final matched application identification. For example, if the traffic's classification path is: HTTP:FACEBOOK-ACCESS:FACEBOOK-CHAT, then the count increases only for FACEBOOK-CHAT.
- `from-zone` *`zone-name`*—(Optional) Displays the number of hits for security policies associated with the named source zone.
- `greater-than` *`count`*—(Optional) Displays security policies for which the number of hits is greater than the specified number.

  Range: 0 through 4,294,967,295
- `less-than` *`count`*—(Optional) Displays security policies for which the number of hits is less than the specified number.

  Range: 0 through 4,294,967,295
- `logical-system`—Displays the logical system name.
- `root-logical-system`—Displays the number of hits for security policies configured for a root logical system.
- `tenant`—Displays the number of hits for security policies configured for the tenant system.
- `to-zone` *`zone-name`*—(Optional) Displays the number of hits for security policies associated with the named destination zone.
- `url-categories`—(Optional) Displays the number of hits for security policies based on the matching URL categories.

Required Privilege Level

view

Output Fields

"No Link Title" on page 25064 lists the output fields for the `show security policies hit-count` command. Output fields are listed in the approximate order in which they appear.

show security policies hit-count Output Fields

Field Name	Field Description
`index`	Displays the line number
`from-zone`	Name of the source zone
`to-zone`	Name of the destination zone
`name`	Name of the security policy
`policy count`	Number of hits for each security policy
`tenant`	Displays the name of the tenant system.
Action	Policy action - Permit or Deny.
Redirect	Number of permitted sessions that are redirected by the policy
Dynamic-applications	Details of the dynamic applications. • Name - Dynamic-application name • Count- Number of hits for each dynamic application
url-categories	Details of the URL categories. • Name - URL category name • Count- Number of hits for each URL category

Sample Output

show security policies hit-count

```
user@host> show security policies hit-count
index   from zone    to zone       name       policy count
 1       untrust      vrtrust      policy1         40
 2       untrust      trust        policy2         20
 3       untrust      trust        policy3         80

Number of policy: 3
```

Sample Output

show security policies hit-count ascending

```
user@host> show security policies hit-count ascending
index   from zone    to zone       name       policy count
 2       untrust      trust        policy2         20
 1       untrust      vrtrust      policy1         40
 3       untrust      trust        policy3         80

Number of policy: 3
```

Sample Output

show security policies hit-count descending greater-than 70 less-than 100

```
user@host> show security policies hit-count descending greater-than 70 less-than 100
index   from zone    to zone       name       policy count
 2       untrust      trust        policy2         100
 1       untrust      vrtrust      policy1         90
 3       untrust      trust        policy3         80

Number of policy: 3
```

Sample Output

show security policies hit-count from-zone untrust to-zone trust

```
user@host> show security policies hit-count from-zone untrust to-zone trust
index   from zone    to zone       name       policy count
 2       untrust      trust        policy2         20
 3       untrust      trust        policy3         80

Number of policy: 2
```

Sample Output

show security policies hit-count

```
user@host> show security policies hit-count
Logical system: root-logical-system
Index  From zone        To zone          Name                  Policy count  Action
1      junos-global    junos-global      policy1                5202         Permit
```

Sample Output

show security policies hit-count detail

```
user@host> show security policies hit-count detail
Logical system: root-logical-system
Index  From zone        To zone          Name                  Policy count  Action  Redirect
1      junos-global    junos-global       policy1                5202        Permit    0
2      junos-global    junos-global       policy2                5202        Reject    3
```

Sample Output

show security policies hit-count policy-name `policy-name` detail url-categories

```
user@host> show security policies hit-count policy-name policy1 detail url-categories
Logical system: root-logical-system
```

```
Index  From zone        To zone          Name                  Policy count  Action  Redirects
1      junos-global    junos-global       policy1              5202            Permit        0
URL-Categories
Name                                  Count
Enhanced_Information_Technology       1000
Enhanced_Collaboration_Office         2000
```

Sample Output

show security policies hit-count policy-name `policy-name` detail dynamic-applications

```
user@host> show security policies hit-count policy-name policy1 detail dynamic-applications
Logical system: root-logical-system
Index  From zone        To zone          Name                  Policy count  Action  Redirect
1      junos-global    junos-global       policy1                5202        Permit    0
Dynamic-applications
Name    Application Count
Junos:SSH         2234
Junos:HTTP        2502
Junos:TWITTER      529
```

Release Information

Command introduced in Junos OS Release 12.1.

The `tenant` option is introduced in Junos OS Release 18.3R1.

The `dynamic-applications` and `url-categories` options are introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

clear security policies hit-count | 4035

Security Policies Overview

Unified Policies Configuration Overview

show security policies information

IN THIS SECTION

- Syntax | 25068
- Description | 25068
- Options | 25068
- Required Privilege Level | 25069
- Output Fields | 25069
- Sample Output | 25072
- Release Information | 25077

Syntax

```
show security policies information
<logical-system (logical-system-name | all)>
<root-logical-system>
<tenant tenant-name>
```

Description

Displays detailed information about the security policies configured on the device.

NOTE: Dynamic policy counters are only supported in the root logical system.

Options

logical-system Displays detailed information about the security policies configured on a logical system or on all logical systems.

root-logical-system Displays detailed information about the security policies configured on the root logical system. This is the default option.

tenant Displays detailed information about the security policies configured on a tenant.

Required Privilege Level

view

Output Fields

Table 1322 on page 25069 lists the output fields for the `show security policies information` command. Output fields are listed in the approximate order in which they appear.

Table 1322: show security policies Output Fields

Field Name	Field Description
`Number of global policies`	Number of global policies configured on the device.
`Number of policies with scheduler`	Number of policies with schedulers configured on the device.
`Number of policies with statistics enabled`	Number of policies configured with statistics enabled on the device and the maximum number of policies which can be configured with statistics enabled on the device.
`Number of unified policies`	Number of unified policies configured on the device.
`Number of policy contexts`	Number of policy contexts configured on the device.
`Number of Policies per context`	Number of policies per context configured on the device and the maximum number of policies per context that can be configured on the device.
`Number of Source addresses per policy`	Number of source addresses configured per policy and the maximum number of source addresses configured per policy. The source address in the match criteria is composed of one or more address names or address set names in the `from-zone`.

Table 1322: show security policies Output Fields *(Continued)*

Field Name	Field Description
`Number of Destination addresses per policy`	Number of destination addresses configured per policy and maximum of destination addresses that can be configured per policy. The destination address of the match criteria is composed of one or more address names or address set names in the `to-zone`.
`Number of Applications per policy`	Number of applications per policy and the maximum number of applications per policy.
`Number of Dynamic applications per policy`	Number of dynamic applications per policy and the maximum number of dynamic applications per policy.
`Number of Source identities per policy`	Number of source identities per policy and the maximum number of source identities per policy.
`Number of Match source/destination identity feeds per policy`	Number of source and destination identities feeds per policy matching traffic.
`Messages received`	Number of messages received.
`Messages rejected`	Number of messages rejected.
`Add messages received`	Number of add messages received.
`Delete messages received`	Number of delete messages received.
`Clear messages received`	Number of clear messages received.
`Invalid messages received`	Number of invalid messages received.

Table 1322: show security policies Output Fields *(Continued)*

Field Name	Field Description
`Add messages sent to PFE`	Number of add messages sent from Routing Engine to Packet Forwarding Engine.
`Delete messages sent to PFE`	Number of delete messages sent from Routing Engine to Packet Forwarding Engine.
`Clear messages sent to PFE`	Number of clear messages sent from Routing Engine to Packet Forwarding Engine.
`Policy added successfully`	Number of policies successfully added.
`Policy deleted successfully`	Number of policies successfully deleted.
`Policy cleared successfully`	Number of policies successfully cleared.
`Policy add failures`	Number of policies unsuccessfully added.
`Policy delete failures`	Number of policies unsuccessfully deleted.
`Policy clear failures`	Number of policies unsuccessfully cleared.
`SSAM send attempted`	Number of SSAM message attempts sent to Internet Key Exchange Protocol Daemon (IKED).
`SSAM send succeeded`	Number of SSAM messages sent to IKED.
`SSAM send failed`	Number of SSAM messages unsucessfully sent to IKED.
`Policy failures - bad configuration`	Number of messages with invalid dynamic policy configurations provided.

Table 1322: show security policies Output Fields *(Continued)*

Field Name	Field Description
`Policy failures - bad scope policy`	Number of messages with invalid scope policy provided.
`Dependent-dynamic-application-lookup`	Value of the unified policy dependent match flag's value (`Dependent-dynamic-application-lookup`).
`Unified-policy-implicit-match`	Value of the unified policy implicit match flag's value (`Unified-policy-implicit-match`).

Sample Output

show security policies information

```
user@host> show security policies information

    Number of policies                            :            0 / 16000
        Number of global policies                 :            0
        Number of policies with scheduler         :            0
        Number of policies with statistics enabled:            0 / 256
        Number of unified policies                :            0
    Number of policy contexts                     :            0 / 4096
    Number of Policies per context                :            0 / 10240
    Number of Source addresses per policy         :            0 / 4096
    Number of Destination addresses per policy    :            0 / 4096
    Number of Applications per policy             :            0 / 3072
    Number of Dynamic applications per policy     :            0 / 4096
    Number of Source identities per policy        :            0 / 128

  Dynamic policy counters:
    Messages received                   :            0
    Messages rejected                   :            0
    Add messages received               :            0
    Delete messages received            :            0
    Clear messages received             :            0
    Invalid messages received           :            0
```

```
    Add messages sent to PFE           :            0
    Delete messages sent to PFE        :            0
    Clear messages sent to PFE         :            0
    Policy added successfully          :            0
    Policy deleted successfully        :            0
    Policy cleared successfully        :            0
    Policy add failures                :            0
    Policy delete failures             :            0
    Policy clear failures              :            0
    SSAM send attempted                :            0
    SSAM send succeeded                :            0
    SSAM send failed                   :            0
    Policy failures - bad configuration :            0
    Policy failures - bad scope policy  :            0
Unified policy status:
  Dependent-dynamic-application-lookup: disable
  Unified-policy-implicit-match: enable
```

show security policies information logical-system all

```
user@host> show security policies information logical-system all

   Policy utilization counters:
    Number of policies                            :           10 / 60000
        Number of global policies                 :            0
        Number of policies with scheduler         :            0
        Number of policies with statistics enabled:            0 / 1024
        Number of unified policies                :            1
    Number of policy contexts                     :            2 / 4096
    Number of Policies per context                :            9 / 60000
    Number of Source addresses per policy         :            4 / 4096
    Number of Destination addresses per policy    :            4 / 4096
    Number of Applications per policy             :            8 / 3072
    Number of Dynamic applications per policy     :           10 / 4096
    Number of URL Category per policy             :            0 / 1024
    Number of Source identities per policy        :            0 / 128

  Dynamic policy counters:
    Messages received                  :            0
    Messages rejected                  :            0
    Add messages received              :            0
```

```
    Delete messages received           :            0
    Clear messages received            :            0
    Invalid messages received          :            0
    Add messages sent to PFE           :            0
    Delete messages sent to PFE        :            0
    Clear messages sent to PFE         :            0
    Policy added successfully          :            0
    Policy deleted successfully        :            0
    Policy cleared successfully        :            0
    Policy add failures                :            0
    Policy delete failures             :            0
    Policy clear failures              :            0
    SSAM send attempted                :            0
    SSAM send succeeded                :            0
    SSAM send failed                   :            0
    Policy failures - bad configuration :            0
    Policy failures - bad scope policy  :            0
Unified policy status:
  Dependent-dynamic-application-lookup: disable
  Unified-policy-implicit-match: enable

Logical system: lsys1
  Policy utilization counters:
    Number of policies                           :            1 / 60000
        Number of global policies                :            0
        Number of policies with scheduler        :            0
        Number of policies with statistics enabled:           0 / 1024
        Number of unified policies               :            0
    Number of policy contexts                    :            1 / 4096
    Number of Policies per context               :            1 / 60000
    Number of Source addresses per policy        :            1 / 4096
    Number of Destination addresses per policy   :            1 / 4096
    Number of Applications per policy            :            1 / 3072
    Number of Dynamic applications per policy    :            0 / 4096
    Number of URL Category per policy            :            0 / 1024
    Number of Source identities per policy       :            0 / 128
Unified policy status:
  Dependent-dynamic-application-lookup: disable
  Unified-policy-implicit-match: enable
```

show security policies information (Threat Profiling)

```
user@host> show security policies information

   Policy utilization counters:
    Number of policies                                          :            3 / 10240
        Number of global policies                               :            0
        Number of policies with scheduler                       :            0
        Number of policies with statistics enabled              :            0 / 1024
        Number of unified policies                              :            0
    Number of policy contexts                                   :            1 / 2048
    Number of Policies per context                              :            3 / 10240
    Number of Source addresses per policy                       :            1 / 1024
    Number of Destination addresses per policy                  :            1 / 1024
    Number of Applications per policy                           :            1 / 128
    Number of Dynamic applications per policy                   :            0 / 4096
    Number of URL Category per policy                           :            0 / 1024
    Number of Source identities per policy                      :            0 / 128
    Number of Match source/destination identity feeds per policy      :            4 / 1024

  Dynamic policy counters:
    Messages received                      :            0
    Messages rejected                      :            0
    Add messages received                  :            0
    Delete messages received               :            0
    Clear messages received                :            0
    Invalid messages received              :            0
    Add messages sent to PFE               :            0
    Delete messages sent to PFE            :            0
    Clear messages sent to PFE             :            0
    Policy added successfully              :            0
    Policy deleted successfully            :            0
    Policy cleared successfully            :            0
    Policy add failures                    :            0
    Policy delete failures                 :            0
    Policy clear failures                  :            0
    SSAM send attempted                    :            0
    SSAM send succeeded                    :            0
    SSAM send failed                       :            0
    Policy failures - bad configuration    :            0
    Policy failures - bad scope policy     :            0
Unified policy status:
```

```
  Dependent-dynamic-application-lookup: disable
  Unified-policy-implicit-match: enable
```

show security policies information (destination-identity)

```
user@host> show security policies information
Policy utilization counters:
    Number of policies                                          :            4 / 20480
    Number of global policies                                   :            1
    Number of policies with scheduler                           :            0
    Number of policies with statistics enabled                  :            0 / 1024
    Number of unified policies                                  :            2
    Number of policy contexts                                   :            2 / 4096
    Number of Policies per context                              :            3 / 20480
    Number of Source addresses per policy                       :            1 / 2048
    Number of Destination addresses per policy                  :            1 / 2048
    Number of Applications per policy                           :            3 / 512
    Number of Dynamic applications per policy                   :            5 / 4096
    Number of URL Category per policy                           :            0 / 1024
    Number of Source/destination identities per policy          :            2 / 128
    Number of Match source/destination identity feeds per policy :            0 / 1024

Dynamic policy counters:
    Messages received                   :            0
    Messages rejected                   :            0
    Add messages received               :            0
    Delete messages received            :            0
    Clear messages received             :            0
    Invalid messages received           :            0
    Add messages sent to PFE            :            0
    Delete messages sent to PFE         :            0
    Clear messages sent to PFE          :            0
    Policy added successfully           :            0
    Policy deleted successfully         :            0
    Policy cleared successfully         :            0
    Policy add failures                 :            0
    Policy delete failures              :            0
    Policy clear failures               :            0
    SSAM send attempted                 :            0
    SSAM send succeeded                 :            0
    SSAM send failed                    :            0
```

```
    Policy failures - bad configuration :            0
    Policy failures - bad scope policy  :            0
Unified policy status:
  Dependent-dynamic-application-lookup  :         disable
  Unified-policy-implicit-match         :         enable
```

Release Information

Command introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Configuring Security Policies

Unified Security Policies

show security policies lookup-intact-on-commit

IN THIS SECTION

- Syntax | 25077
- Description | 25078
- Options | 25078
- Required Privilege Level | 25078
- Output Fields | 25078
- Sample Output | 25079
- Release Information | 25080

Syntax

```
show security policies lookup-intact-on-commit [ eligibility | status ]
```

Description

Displays the information about the `lookup-intact-on-commit` configuration status on all forwarding planes on the device. The output also displays eligibility of the device based on the option.

Options

eligibility Display the eligibility of the device to enable the (lookup-intact-on-commit) configuration. The command checks the memory and platform requirements of the device and displays the eligibility information along with the system memory availability. The device is eligible only if it has enough free space (20% of total memory) after accommodating the required memory utilization of this feature.

status Display the activation status of the configuration (lookup-intact-on- commit) on all the forwarding planes of the device. The output shows the details per forwarding plane. To activate the (**lookup-intact-on-commit**) option on the forwarding plane(s), you must restart all the forwarding planes and the forwarding plane(s) on the chassis cluster pair.

Required Privilege Level

view

Output Fields

Table lists the output fields for the `show security policies lookup-intact-on-commit` command. Output fields are listed in the approximate order in which they appear.

Table 1323: show security policies lookup-intact-on-commit

Field Name	Field Description
`FPC name`	Displays the forwarding plane name along with Layer2 chassis cluster node information. This ou is shown for the (status and eligibility) options

Table 1323: show security policies lookup-intact-on-commit *(Continued)*

Field Name	Field Description
Status	Displays the forwarding plane configuration activation status. This output is shown for the (stat option. The output displays one of the following status related to the : • **Not configured**—The option is not enabled on the system. • **Configured, restart forwarding to Activate**—The option is configured and committed, but the forwarding plane is not restarted. • **Activated**—The configuration is activated on the forwarding plane. • **Activation failure on dataplane**—The configuration is enabled, but the forwarding plane faile activate it as the memory requirements are not met.
Eligibility	Displays the memory eligibility of the system per FPC. This output indicates if the particular FP eligible for configuring the set security policies lookup-intact-on-commit configuration. The fiel displays the verdict based on forwarding plane inputs as Eligible (or) Ineligible with Kernel heap (%) after validating the required memory for the feature.

Sample Output

show security policies lookup-intact-on-commit status

```
user@host> show security policies lookup-intact-on-commit status
FPC NAME     Status
master       Not Configured
```

show security policies lookup-intact-on-commit eligibility

```
user@host> show security policies lookup-intact-on-commit eligibility
FPC NAME         Eligibility
master           Eligible, kernel heap free:34%, is more than threshold
                 Eligible, user heap free:15%, is more than threshold
```

show security policies lookup-intact-on-commit eligibility (Chassis Cluster)

```
user@host> show security policies lookup-intact-on-commit eligibility
node0:
--------------------------------------------------------------------------
            FPC NAME         Eligibility
            node0.fpc0       Eligible, kernel heap free:58%, is more than threshold
                             Eligible, user heap free:3%, is more than threshold
            node1.fpc0       Eligible, kernel heap free:58%, is more than threshold
                             Eligible, user heap free:3%, is more than threshold
```

Release Information

Command introduced in Junos OS Release 24.4R1

show security policies policy set

IN THIS SECTION

- Syntax | 25080
- Description | 25081
- Required Privilege Level | 25081
- Output Fields | 25081
- Sample Output | 25082
- Release Information | 25083

Syntax

```
show security policies policy-set
```

Description

Displays a summary of all policy set you have created for inner session tunnel inspection for EVPN-VXLAN tunnel traffic.

Required Privilege Level

view

Output Fields

show security policies policy-set Output Fields

Field Name	Field Description
`From zone`	Name of the source zone.
`To zone`	Name of the destination zone.
`Policy-name`	Name of the policy-set
`State`	Status of the policy: • `enabled`: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • `disabled`: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
`Index`	Internal number associated with the policy.
`Sequence number`	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, 4.
Scope Policy	Policy identifier.
Log Profile ID	Internal log profile number.
from-zone	source zone of the traffic.
to-zone	Destination zone of the traffic.

Field Name	Field Description
`Source vrf group`	Source virtual routing and forwarding (VRF). One or many source VRF instances, for example, the VRF routing instance associated with an incoming packet
`Destination vrf group`	Destination (VRF) . One or many destination VRF instances.
Source address	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names.
Destination Address	Name of the destination address (or address set) as it was entered in the destination zone's address book.
Application	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time.
`Source identity feeds`	One or more user roles specified for a policy.
`Destination identity feeds`	One or more user roles specified for a policy.
Action	The action taken for a packet that matches the policy's tuples. Actions include : deny reject permit

Sample Output

show security policies policy-set

```
user@host> show security policies policy-set
```

```
From zone: PSET-1, To zone: PSET-1
  Policy: PSET-1-P1, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1, Log Profile
ID: 0
    From zones: any
    To zones: any
    Source vrf group: any
    Destination vrf group: any
    Source addresses: a2-untrust
```

```
    Destination addresses: a2-trust
    Applications: any
    Source identity feeds: any
    Destination identity feeds: any
    Action: permit, application services
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Tunnel Inspection for EVPN-VXLAN on SRX Series Devices

show security policies unknown-source-identity

IN THIS SECTION

- Syntax | 25083
- Description | 25084
- Required Privilege Level | 25084
- Output Fields | 25084
- Sample Output | 25085
- Release Information | 25085

Syntax

```
show security policies unknown-source-identity
```

Description

Display a list of any user or role that is referenced in a policy as a source-identity, but is not yet included in the role provisioning table.

The role provisioning table is created from the local authentication table, UAC authentication tables, and user-identification authentication tables. The UAC and user-identification authentication tables are dynamic and contain only those users currently authenticated. Because of this, a role can be listed as unknown because no user associated with the role has authenticated yet. There is no consequence if a role remains unknown.

Required Privilege Level

view

Output Fields

Table 1324 on page 25084 lists the output fields for the `show security policies unknown-source-identity` command. Output fields are listed in the approximate order in which they appear.

Table 1324: show security policies unknown-source-identity Output Fields

Field Name	Field Description
`From zone`	Part of the zone pair that identifies the source of the traffic to which a policy applies. Affected policies are grouped by their zone pair.
`To zone`	Part of the zone pair that identifies the destination of the traffic to which a policy applies. Affected policies are grouped by their zone pair.
`Policy`	The name of the policy that contains the unknown source identity.
Unknown source identities	A list of user names and roles specified in the source-identity field of the named policy that are unknown.

Sample Output

show security policies unknown-source-identity

In the following sample output, policy p1 which controls traffic from the untrust zone to the trust zone specifies two roles, r1 and r3, that are not yet provisioned. Similarly, policy p2 affecting traffic from the trust zone to the trust zone also contains two roles that are not provisioned, role1 and abc.

```
user@host> show security policies unknown-source-identity
From zone: untrust, To zone: trust
  Policy: p1
  Unknown source identities: r1, r3
From zone: trust, To zone: trust
  Policy: p2
  Unknown source identities: role1, abc
```

Release Information

Command introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Security Policies Overview

show security policies zone-context

IN THIS SECTION

- Syntax | 25086
- Description | 25086
- Options | 25086
- Required Privilege Level | 25087
- Output Fields | 25087

Sample Output | 25088
Release Information | 25088

Syntax

```
show security policies zone-context
application-firewall
count
detail
from-zone <zone-name>
interface
logical-system <logical-system-name>
policy-name
root-logical-system
service-set
start
tenant <tenant-name>
to-zone <zone-name>
```

Description

Displays the security policy that applies the security rules to the transit traffic within a context (from-zone to to-zone). From the perspective of security policies, traffic enters into one security zone and goes out on another security zone. This combination of a `from-zone` and a `to-zone` is defined as a context. Each context contains an ordered list of policies. The existing show command for security policies zone-context is enhanced with tenant support.

Options

`application-firewall` Displays the information of application-firewall.

`count` Displays the number of policies to show.

- **Range:** 1 through 65535

`detail` Displays the detailed information.

from-zone	Displays the policy information matching the given source zone.
interface	Displays the name of adaptive services interface.
logical-system	Displays the name of the logical system.
policy-name	Displays the policy information matching the given policy name.
root-logical-system	Displays root logical system as default.
service-set	Displays the name of service set.
start	Displays the policies from a given position. • **Range:** 1 through 65535
tenant	Displays the name of the tenant system.
to-zone	Displays the policy information matching the given destination zone.

Required Privilege Level

view

Output Fields

Table 1325 on page 25087 lists the output fields for the `show security policies zone-context` command. Output fields are listed in the approximate order in which they appear.

Table 1325: show security zone-context Output Fields

Field Name	Field Description
`From zone`	Name or ID of the source zone of the traffic.
`To zone`	Name or ID of the destination zone of the traffic.
`Policy count`	Number of hits for each security policy.
`Tenant`	Name of the tenant system.

Sample Output

show security policies zone-context tenant TN1

```
user@host> show security policies zone-context tenant TN1
From zone                To zone                  Policy count
   z1                           z2                       1
   z2                           z3                       1
```

show security policies zone-context tenant all

```
user@host> show security policies zone-context tenant all
Tenant: TN1
    From zone                To zone                  Policy count
    z1                       z2                       1

Tenant: TN2
    From zone                To zone                  Policy count
    z1                       z2                       1
```

Release Information

Command introduced in Junos OS Release 8.5.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Security Policies Overview

Configuring Security Policies

show security policy-report

IN THIS SECTION

- Syntax | 25089
- Description | 25089
- Options | 25089
- Required Privilege Level | 25091
- Sample Output | 25091
- Release Information | 25092

Syntax

```
show security policy-report
<from-zone zone-name>
<global>
<report-type report-type>
<to-zone zone-name>
```

Description

Displays detailed security policy reports.

Optimizing security policies ensure that the policies are efficient. Over time, policies become disorganised and hence ineffective. You can use the `show security policy-report` command to notify end users when you create new policies or change existing policies which adversely affect other security policies.

Options

from-zone Displays the policy report matching the given source zone.

- **Default:** any

global Displays the security policy report of the global policies.

report-type Displays the type of the policy report.

You can configure the following options for `report-type`:

1-year-not-hit Displays the policy report for policies that have not been hit in the last 1 year.

30d-not-hit Displays the policy report for policies that have not been hit in the last 30 days.

60d-not-hit Displays the policy report for policies that have not been hit in the last 60 days.

90d-not-hit Displays the policy report for policies that have not been hit in the last 90 days.

consolidation Displays the policy report for policies which can be consolidated.

When two policies share most fields in common and only one of the fields contains difference, they are eligible for consolidation.

expired Displays the policy report for expired policies.

generalization Displays the policy report for generalized policies.

A policy is generalized when a following policy matches all the packets of the current policy and the action for each policy is different.

least-hit Displays the policy report for policies with the least hit count.

most-hit Displays the policy report for policies with the most hit count.

no-comments Displays the policy report for policies which have no comments.

no-logging Displays the policy report for policies which do not have any logging

redundant Displays the policy report for redundant policies.

A policy is redundant if there exists a preceding policy within the policy which performs the same action on the same packets as performed by the current policy.

scheduler Displays the policy report for policies in which the scheduled has expired/ policies that have an active or inactive schedule.

shadowing Displays the policy report for shadowed policies.

A policy is shadowed when a preceding policy matches all the packets of the current policy and the action for each policy is different.

unused Displays the policy report for policies which have a hit count value of zero.

- **Default:** If you don't configure a specific `report-type`, then all the reports are displayed.

to-zone Displays the policy report matching the given destination zone.

- **Default:** any

> **NOTE**: SRX series devices only analyze the following fields of a policy for the shadowing, redundant, generalization, and consolidation reports:
>
> - Source address (Ipv4 only)
> - Destination address (Ipv4 only)
> - Applications

Required Privilege Level

view

Sample Output

show security policy-report report-type consolidation

```
user@host> show security policy-report report-type consolidation
Logical system: root-logical-system
  Problem: Policy p2 is consolidated with policy p1
  Recommendation: Please check if policy p2 and p1
    Index: 129
    Policy: p2
      Source zone: trust
      Destination zone: untrust
      Source addresses: s_ad2
      Destination addresses: d_ad2
      Applications: app2
      Action: permit
```

```
    Policy: p1
      Source zone: trust
      Destination zone: untrust
      Index: 180
      Source addresses: s_ad1
      Destination addresses: d_ad1
      Applications
```

show security policy-report report-type scheduler

```
user@host> show security policy-report report-type scheduler
Logical system: root-logical-system
  Problem: Policy p2 has scheduler
     Policy: p2, Index: 129
     Source zone: trust
     Destination zone: untrust
     Source addresses: s_ad2
     Destination addresses: d_ad2
     Applications: app2
     Action: permit
     Scheduler: sch
```

show security policy-report report-type 1-year-not-hit

```
user@host> show security policy-report report-type 1-year-not-hit
Logical system: root-logical-system
  Problem: Not been hit for one year
  Recommendation: Please check if you need to remove policy p2
     Policy: p2, Index: 129
     Source zone: trust
     Destination zone: untrust
     Source addresses: s_ad2
     Destination addresses: d_ad2
     Applications: app2
     Action: permit
```

Release Information

Command introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

report-skip | 14403

show security re-distribution ipsec-vpn

IN THIS SECTION

- Syntax | 25093
- Description | 25093
- Options | 25093
- Required Privilege Level | 25094
- Output Fields | 25094
- Sample Output | 25094
- Sample Output | 25095
- Release Information | 25095

Syntax

```
show security re-distribution ipsec-vpn <gateway-name gateway-name>
```

Description

After executing `request security re-distribution ipsec-vpn gateway-name` *`gateway-name`* command, it may take some time to establish a new tunnel. This command displays the commands for which the tunnels are in the pending or awaiting state to get established.

Options

none	Display information for all the gateways.
gateway-name	(Optional) Display information for the specified gateway.

Required Privilege Level

view

Output Fields

Table 1326 on page 25094 lists the output fields for the `show security re-distribution ipsec-vpn` command. Output fields are listed in the approximate order in which they appear.

Table 1326: show security re-distribution ipsec-vpn Output Fields

Field Name	Field Description	Level of Output
`Gateway-name`	Name of the IKE gateway.	All levels
`FPC`	FPC slot number.	All levels
`PIC`	PIC slot number.	All levels
`Thread-id`	Thread identifier.	All levels
`Remote-id`	Remote identifier.	All levels

Sample Output

show security re-distribution ipsec-vpn

```
user@host> show security re-distribution ipsec-vpn

Gateway-name        FPC        PIC        Thread-id        Remote-id
Gateway-name-1       3          1             3              n/a
Gateway-name         2          2             5              n/a
Gateway-name-2       1          3             7            ike-id-3
```

Sample Output

show security re-distribution ipsec-vpn gateway-name

```
user@host> show security re-distribution ipsec-vpn gateway-name gateway-name-1

Gateway-name        FPC        PIC      Thread-id         Remote-id
Gateway-name-1       3          1            3             ike-id-1
Gateway-name-1       3          0            4             ike-id-3
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

- request security re-distribution ipsec-vpn | 14863
- show security ipsec tunnel-distribution | 24755

show security resource-manager group active

IN THIS SECTION

- Syntax | 25096
- Description | 25096
- Options | 25096
- Required Privilege Level | 25096
- Output Fields | 25096
- Sample Output | 25097
- Sample Output | 25097
- Sample Output | 25098
- Sample Output | 25098
- Sample Output | 25098

Release Information | 25099

Syntax

```
show security resource-manager group active
                  <group-number        >
<node (          node-id         | all | local | primary)>
```

Description

This command displays security information about active groups created through the resource manager.

Options

- none—Display resource manager group service information for all active groups.
- *group-number* —(Optional) Display resource manager group service information for a specific group identification number.
- node—(Optional) For chassis cluster configurations, display active resource manager group service information on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - all—Display information about all nodes.
  - local—Display information about the local node.
  - primary—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1327 on page 25097 lists the output fields for the show security resource-manager group command. Output fields are listed in the approximate order in which they appear.

Table 1327: show security resource-manager group Output Fields

Field Name	Field Description
`Total groups`	Total number of groups in the system.
`active groups`	Number of active groups.
`Group ID`	Identification number whose group information is displayed.

Sample Output

show security resource-manager group active

```
user@host> show security resource-manager group active
Total groups 32, active groups 0
```

Sample Output

show security resource-manager group active 2048

```
user@host> show security resource-manager group active 2048
Total groups 2048, active groups 1
Group ID 2048: state - Active
          : Virtual System      - root
          : Application         - SIP ALG
          : Group Timeout       - 65535
          : Number of resources - 3
               Resource ID - 8190
               Resource ID - 8188
               Resource ID - 8187
```

Sample Output

show security resource-manager group active node primary

```
user@host> show security resource-manager group active node primary
node0:
--------------------------------------------------------------------------
Group ID 1024: Application - SIP ALG
Total groups 1024, active groups 1
```

Sample Output

show security resource-manager group active node all

```
user@host> show security resource-manager group active node all
node0:
--------------------------------------------------------------------------
Group ID 1024: Application - SIP ALG
Total groups 1024, active groups 1
node1:
--------------------------------------------------------------------------
Group ID 1024: Application - SIP ALG
Total groups 1024, active groups 1
```

Sample Output

show security resource-manager group active 1024 node all

```
user@host> show security resource-manager group active 1024 node all
node0:
--------------------------------------------------------------------------
Group ID 1024: state - Active
          : Application         - SIP ALG
          : Group Timeout       - 65535
          : Number of resources - 3
               Resource ID - 8192
               Resource ID - 8188
               Resource ID - 8187
```

```
node1:
--------------------------------------------------------------------------
Group ID 1024: state - Active
          : Application         - SIP ALG
          : Group Timeout       - 65535
          : Number of resources - 3
               Resource ID - 8187
               Resource ID - 8186
               Resource ID - 8190
```

Release Information

Command introduced in Junos OS Release 8.5.

The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security resource-manager resource active

IN THIS SECTION

- Syntax | 25100
- Description | 25100
- Options | 25100
- Required Privilege Level | 25100
- Output Fields | 25100
- Sample Output | 25101
- Sample Output | 25101
- Sample Output | 25102
- Sample Output | 25102
- Release Information | 25103

Syntax

```
show security resource-manager resource active
                   <resource-id        >
<node (          node-id         | all | local | primary)>
```

Description

This command displays security information about active resources created through the resource manager.

Options

- none—Display information for all active resources.
- *resource-id* —(Optional) Display information for a resource with a specific identification number.
- `node`—(Optional) For chassis cluster configurations, display active resource manager information on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1328 on page 25101 lists the output fields for the `show security resource-manager resource` command. Output fields are listed in the approximate order in which they appear.

Table 1328: show security resource-manager resource Output Fields

Field Name	Field Description
`Total resources`	Total number of resources in the system.
`active resources`	Number of active resources.
`Resource ID`	Identification number whose resource information is displayed.

Sample Output

show security resource-manager resource active

```
user@host> show security resource-manager resource active
    Resource ID 7: Group ID - 2, Application - JSF_sip

    Resource ID 6: Group ID - 2, Application - JSF_sip

    Resource ID 5: Group ID - 2, Application - JSF_sip

    Resource ID 4: Group ID - 2, Application - JSF_sip

    Resource ID 3: Group ID - 2, Application - JSF_sip

    Resource ID 1: Group ID - 2, Application - JSF_sip

    Resource ID 2: Group ID - 2, Application - JSF_sip
Total Resources 4326, active resources 7
```

Sample Output

show security resource-manager resource active 5

```
user@host> show security resource-manager resource active 5
Resource ID 5: state - Active
    Application         - asl_client
```

```
    Parent group        - 2
    Policy              - 5
    From zone           - untrust
    To zone             - trust
    Resource timeout    - 0
    Number of sessions  - 0
    Number of Holes     - 1
        Source IP range        - {0.0.0.0, 0.0.0.0}
        Source port range      - {0, 0}
        Destination IP range   - {33.1.0.200, 33.1.0.200}
        Destination port range - {5060, 5060}
        Translated             - {0.0.0.0/0 -> 33.1.0.200/5060}
        Protocol               - 17
        Reference count        - 1
```

Sample Output

show security resource-manager resource active node local

```
user@host> show security resource-manager resource active node local
node0:
--------------------------------------------------------------------------
    Resource ID 8192: Group ID - 1024, Application - SIP ALG
    Resource ID 8188: Group ID - 1024, Application - SIP ALG
    Resource ID 8187: Group ID - 1024, Application - SIP ALG
Total Resources 8192, active resources 3
```

Sample Output

show security resource-manager resource active node primary

```
user@host> show security resource-manager resource active node primary
node0:
--------------------------------------------------------------------------
    Resource ID 8192: Group ID - 1024, Application - SIP ALG
    Resource ID 8188: Group ID - 1024, Application - SIP ALG
    Resource ID 8187: Group ID - 1024, Application - SIP ALG
Total Resources 8192, active resources 3
```

Release Information

Command introduced in Junos OS Release 8.5.

The `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security resource-manager settings

IN THIS SECTION

- Syntax | 25103
- Description | 25103
- Options | 25104
- Required Privilege Level | 25104
- Output Fields | 25104
- Sample Output | 25105
- Sample Output | 25105
- Sample Output | 25105
- Release Information | 25105

Syntax

```
show security resource-manager settings
<node (          node-id         | all | local | primary)>
```

Description

Display resource manager settings.

Options

`node`—(Optional) For chassis cluster configurations, display resource manager settings on a specific node.

- *`node-id`* —Identification number of the node. It can be 0 or 1.
- `all`—Display information about all nodes.
- `local`—Display information about the local node.
- `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 1329 on page 25104 lists the output fields for the `show security resource-manager settings` command. Output fields are listed in the approximate order in which they appear.

Table 1329: show security resource-manager settings Output Fields

Field Name	Field Description
`Client Heartbeat`	Time after which idle an resource manager client is timed out.
`Count`	Number of active clients.
`Pinhole age`	Duration for which the temporary opening in the security firewall (pinhole) is open for specified traffic. If the specified traffic does not exist during this time period, the pinhole is timed out.

Sample Output

show security resource-manager settings

```
user@host> show security resource-manager settings
Client Heartbeat: timeout 600 seconds, count 5
Pinhole age: 32 seconds
```

Sample Output

show security resource-manager settings node primary

```
user@host> show security resource-manager settings node primary
node0:
--------------------------------------------------------------------------
Client heartbeat: timeout 600 seconds, count 5
Pinhole age: 120 seconds
```

Sample Output

show security resource-manager settings node all

```
user@host> show security resource-manager settings node all
node0:
--------------------------------------------------------------------------
Client heartbeat: timeout 600 seconds, count 5
Pinhole age: 120 seconds
node1:
--------------------------------------------------------------------------
Client heartbeat: timeout 600 seconds, count 5
Pinhole age: 120 seconds
```

Release Information

Command introduced in Junos OS Release 8.5; `node` options added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security resource-manager summary

IN THIS SECTION

- Syntax | 25106
- Description | 25106
- Required Privilege Level | 25106
- Output Fields | 25106
- Sample Output | 25107
- Release Information | 25107

Syntax

```
show security resource-manager summary
```

Description

This command displays summary information about active resources, clients, groups, and sessions created through the resource manager.

Required Privilege Level

view

Output Fields

Table 1330 on page 25107 lists the output fields for the `show security resource-manager summary` command. Output fields are listed in the approximate order in which they appear.

Table 1330: show security resource-manager summary Output Fields

Field Name	Field Description
`Active resource-manager clients`	Number of active resource manager clients.
`Active resource-manager groups`	Number of active resource manager groups.
`Active resource-manager resources`	Number of active resource manager resources.
`Active resource-manager sessions`	Number of active resource manager sessions.

Sample Output

show security resource-manager summary

```
user@host> show security resource-manager summary

Active resource-manager clients   : 15
Active resource-manager groups    : 1
Active resource-manager resources : 1
Active resource-manager sessions  : 0
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security screen ids-option

IN THIS SECTION

- Syntax | 25108
- Description | 25108
- Options | 25108
- Required Privilege Level | 25109
- Output Fields | 25109
- Sample Output | 25112
- Sample Output | 25113
- Sample Output | 25114
- Release Information | 25116

Syntax

```
show security screen ids-option
screen-name
logical-system
root-logical-system
tenant
```

Description

Display the configuration information about the specified security screen. You can configure a `ids-option` to enable screen protection on the SRX Series Firewalls.

Options

- *`screen-name`* —Name of the screen.
- `logical-system`—Name of the logical system.
- `root-logical-system`—Displays root logical system as default.

- `tenant`—Name of the tenant system.

Required Privilege Level

view

Output Fields

Table 1331 on page 25109 lists the output fields for the `show security screen ids-option` command. Output fields are listed in the approximate order in which they appear.

Table 1331: show security screen ids-option Output Fields

Field Name	Field Description
`TCP address sweep threshold`	Number of microseconds for which the device accepts 10 TCP packets from the same remote source to different destination addresses.
`TCP port scan threshold`	Number of microseconds during which the device accepts packets from the same remote source with up to 10 different port numbers.
`ICMP address sweep threshold`	Number of microseconds during which up to 10 ICMP echo requests from the same host are allowed into the device.
`UDP flood threshold`	Number of UDP packets per second allowed to ping the same destination address before the device rejects further UDP packets.
`UDP port scan threshold`	Number of microseconds during which the device accepts packets from the same remote source IP with up to 10 different destination port numbers.
`TCP winnuke`	Enable or disable the detection of TCP WinNuke attacks.
`TCP SYN flood attack threshold`	Number of SYN packets per second required to trigger the SYN proxy response.
`TCP SYN flood alarm threshold`	Number of half-complete proxy connections per second at which the device makes entries in the event alarm log.

Table 1331: show security screen ids-option Output Fields *(Continued)*

Field Name	Field Description
`TCP SYN flood source threshold`	Number of SYN segments to be received per second before the device begins dropping connection requests.
`TCP SYN flood destination threshold`	Number of SYN segments received per second before the device begins dropping connection requests.
`TCP SYN flood timeout`	Maximum length of time before a half-completed connection is dropped from the queue.
`TCP SYN flood queue size`	Number of proxy connection requests that can be held in the proxy connection queue before the device begins rejecting new connection requests.
`ICMP large packet`	Enable or disable the detection of any ICMP frame with an IP length greater than 1024 bytes.
`UDP address sweep threshold`	Number of microseconds for which the device accepts 10 UDP packets from the same remote source to different destination addresses.
`IPv6 extension routing`	Enable or disable the IPv6 extension routing screen option.
`IPv6 extension shim6`	Enable or disable the IPv6 extension shim6 screen option.
`IPv6 extension fragment/IP block fragment`	Enable or disable the IPv6 extension fragment screen option.
`IPv6 extension AH`	Enable or disable the IPv6 extension Authentication Header Protocol screen option.
`IPv6 extension ESP`	Enable or disable the IPv6 extension Encapsulating Security Payload screen option.
`IPv6 extension mobility`	Enable or disable the IPv6 extension mobility screen option.

Table 1331: show security screen ids-option Output Fields *(Continued)*

Field Name	Field Description
`IPv6 extension HIP`	Enable or disable the IPv6 extension Host Identify Protocol screen option.
`IPv6 extension no next`	Enable or disable the IPv6 extension no-next screen option.
`IPv6 extension user-defined`	Enable or disable the IPv6 extension user-defined screen option.
`IPv6 extension HbyH jumbo`	Enable or disable the IPv6 extension HbyH jumbo screen option.
`IPv6 extension HbyH RPL`	Enable or disable the IPv6 extension HbyH RPL screen option.
`IPv6 extension HbyH router alert`	Enable or disable the IPv6 extension HbyH router screen option.
`IPv6 extension HbyH quick start`	Enable or disable the IPv6 extension HbyH quick-start screen option.
`IPv6 extension HbyH CALIPSO`	Enable or disable the IPv6 extension HbyH Common Architecture Label IPv6 Security Screen option.
`IPv6 extension HbyH SMF DPD`	Enable or disable the IPv6 extension HbyH Simplified Multicast Forwarding IPv6 Duplicate Packet Detection screen option.
`IPv6 extension HbyH user-defined`	Enable or disable the IPv6 extension HbyH user-defined screen option.
`IPv6 extension Dst tunnel encap limit`	Enable or disable the IPv6 extension distributed (network) storage tunnel encapsulation limit screen option.
`IPv6 extension Dst home address`	Enable or disable the IPv6 extension DST home address screen option.
`IPv6 extension Dst ILNP nonce`	Enable or disable the IPv6 extension DST Identifier-Locator Network Protocol nonce screen option.

Table 1331: show security screen ids-option Output Fields *(Continued)*

Field Name	Field Description
IPv6 extension Dst line-id	Enable or disable the IPv6 extension DST line-ID screen option.
IPv6 extension Dst user-defined	Enable or disable the IPv6 extension DST user-defined screen option.
IPv6 extension header limit	Threshold for the number of IPv6 extension headers that can pass through the screen.
IPv6 malformed header	Enable or disable the IPv6 malformed header screen option.
ICMPv6 malformed header	Enable or disable the ICMPv6 malformed packet screen option.
UDP flood white-list	Allowlist of IP addresses to bypass UDP flood detection.
IP block fragment white-list	Allowlist of IP addresses to bypass IP block fragmentation check.
Session source limit threshold	Limit the number of concurrent sessions the device can initiate from a single source IP address or the number of sessions it can direct to a single destination IP address.
Logical system/Tenant	Name of the logical system or tenant system.

Sample Output

show security screen ids-option jscreen

```
user@host> show security screen ids-option jscreen
Screen object status:
Name                                        Value
  TCP port scan threshold                   5000
  UDP port scan threshold                   10000
  ICMP address sweep threshold              5000
```

Sample Output

show security screen ids-option jscreen (IPv6)

```
user@host> show security screen ids-option jscreen

Screen object status:

   Name                                        Value
  ICMP ping of death                          enabled
  ......

  IPv6 extension routing                      enabled
  IPv6 extension shim6                        enabled
  IPv6 extension fragment                     enabled
  IPv6 extension AH                           enabled
  IPv6 extension ESP                          enabled
  IPv6 extension mobility                     enabled
  IPv6 extension HIP                          enabled
  IPv6 extension no next                      enabled
  IPv6 extension user-defined                 enabled
  IPv6 extension HbyH jumbo                   enabled
  IPv6 extension HbyH RPL                     enabled
  IPv6 extension HbyH router alert            enabled
  IPv6 extension HbyH quick start             enabled
  IPv6 extension HbyH CALIPSO                 enabled
  IPv6 extension HbyH SMF DPD                 enabled
  IPv6 extension HbyH user-defined            enabled
  IPv6 extension Dst tunnel encap limit       enabled
  IPv6 extension Dst home address             enabled
  IPv6 extension Dst ILNP nonce               enabled
  IPv6 extension Dst line-id                  enabled
  IPv6 extension Dst user-defined             enabled
  IPv6 extension header limit                 20
  IPv6 Malformed header                       enabled
  ICMPv6 malformed packet                     enabled
```

Sample Output

show security screen ids-option jscreen1 node all

```
user@host> show security screen ids-option jscreen1 node all
node0:
--------------------------------------------------------------------------
Screen object status:
Name                                          Value
  UDP flood threshold                         1000
  TCP winnuke                                 enabled
  TCP SYN flood attack threshold              200
  TCP SYN flood alarm threshold               512
  TCP SYN flood source threshold              4000
  TCP SYN flood destination threshold         4000
  TCP SYN flood timeout                       20
  TCP SYN flood queue size                    1024
  ICMP large packet                           enabled
node1:
--------------------------------------------------------------------------
Screen object status:
Name                                          Value
  UDP flood threshold                         1000
  TCP winnuke                                 enabled
  TCP SYN flood attack threshold              200
  TCP SYN flood alarm threshold               512
  TCP SYN flood source threshold              4000
  TCP SYN flood destination threshold         4000
  TCP SYN flood timeout                       20
  TCP SYN flood queue size                    1024
  ICMP large packet                           enabled
```

show security screen ids-option jscreen tenant TN1

```
user@host> show security screen ids-option jscreen tenant TN1

Screen object status:

Name                                         value
UDP flood threshold                          1000
```

```
UDP flood white-list                        a1
UDP flood white-list                        a2
```

show security screen ids-option jscreen tenant all

```
user@host> show security screen ids-option jscreen tenant all

Logical system: root-logical-system
Screen object status:

Name                                        value
UDP flood threshold                         1
UDP flood white-list                        a1
UDP flood white-list                        a2
IP block fragment                           enabled
Session source limit threshold              5

Tenant: TN1
Screen object status:

Name                                        value
UDP flood threshold                         1000
UDP flood white-list                        a1
UDP flood white-list                        a2
```

show security screen ids-option jscreen (IP block fragment screen)

```
user@host> show security screen ids-option jscreen

Screen object status:

Name                                        value
IP block fragment                           enabled
IP block fragment white-list                a1
IP block fragment white-list                a2
```

Release Information

Command introduced in Junos OS Release 8.5. Support for UDP `port scan` added in Junos OS Release 12.1X47-D10.

Support for `node` option added in Junos OS Release 9.0.

Support for IPv6 extension header screens added in Junos OS Release 12.1X46-D10.

The `tenant` option is introduced in Junos OS Release 18.3R1.

The IP block fragment allowlist option added in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

ids-option | 8335

Example: Configuring Multiple Screening Options

show security screen statistics

IN THIS SECTION

- Syntax | 25117
- Description | 25117
- Options | 25117
- Required Privilege Level | 25117
- Output Fields | 25117
- Sample Output | 25121
- Sample Output | 25122
- Sample Output | 25122
- Sample Output | 25123
- Sample Output | 25124
- Sample Output | 25125
- Release Information | 25131

Syntax

```
show security screen statistics <zone zone-name / interface interface-name>
logical-system <logical-system-name | all>
root-logical-system
tenant <tenant-name >
```

Description

Display intrusion detection service (IDS) security screen statistics.

Options

- `zone` *`zone-name`*—Display screen statistics for this security zone.
- `interface` *`interface-name`* —Display screen statistics for this interface.
- *`logical-system-name`*—Display screen statistics for the named logical system.
- `root-logical-system`—(Optional) Display screen statistics for the primary logical system only.
- `tenant`—Display the name of the tenant system.

Required Privilege Level

view

Output Fields

Table 1332 on page 25117 lists the output fields for the `show security screen statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1332: show security screen statistics Output Fields

Field Name	Field Description
`ICMP flood`	Internet Control Message Protocol (ICMP) flood counter. An ICMP flood typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.

Table 1332: show security screen statistics Output Fields *(Continued)*

Field Name	Field Description
`UDP flood`	User Datagram Protocol (UDP) flood counter. UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled.
`TCP winnuke`	Number of Transport Control Protocol (TCP) WinNuke attacks. WinNuke is a denial-of-service (DoS) attack targeting any computer on the Internet running Windows.
`TCP port scan`	Number of TCP port scans. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
`ICMP address sweep`	Number of ICMP address sweeps. An IP address sweep can occur with the intent of triggering responses from active hosts.
`IP tear drop`	Number of teardrop attacks. Teardrop attacks exploit the reassembly of fragmented IP packets.
`TCP SYN flood`	Number of TCP SYN attacks.
`IP spoofing`	Number of IP spoofs. IP spoofing occurs when an invalid source address is inserted in the packet header to make the packet appear to come from a trusted source.
`ICMP ping of death`	ICMP ping of death counter. Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).
`IP source route option`	Number of IP source route attacks.
`TCP address sweep`	Number of TCP address sweeps.

Table 1332: show security screen statistics Output Fields *(Continued)*

Field Name	Field Description
`TCP land attack`	Number of land attacks. Land attacks occur when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.
`TCP SYN fragment`	Number of TCP SYN fragments.
`TCP no flag`	Number of TCP headers without flags set. A normal TCP segment header has at least one control flag set.
`IP unknown protocol`	Number of IPs.
`IP bad options`	Number of invalid options.
`IP record route option`	Number of packets with the IP record route option enabled. This option records the IP addresses of the network devices along the path that the IP packet travels.
`IP timestamp option`	Number of IP timestamp option attacks. This option records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination.
`IP security option`	Number of IP security option attacks.
`IP loose source route option`	Number of IP loose source route option attacks. This option specifies a partial route list for a packet to take on its journey from source to destination.
`IP strict source route option`	Number of IP strict source route option attacks. This option specifies the complete route list for a packet to take on its journey from source to destination.
`IP stream option`	Number of stream option attacks. This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support streams.

Table 1332: show security screen statistics Output Fields *(Continued)*

Field Name	Field Description
`ICMP fragment`	Number of ICMP fragments. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.
`ICMP large packet`	Number of large ICMP packets.
`TCP SYN FIN`	Number of TCP SYN FIN packets.
`TCP FIN no ACK`	Number of TCP FIN flags without the acknowledge (ACK) flag.
`Source session limit`	Number of concurrent sessions that can be initiated from a source IP address.
`TCP SYN-ACK-ACK proxy`	Number of TCP flags enabled with SYN-ACK-ACK. To prevent flooding with SYN-ACK-ACK sessions, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold and SRX Series Firewalls running Junos OS reject further connection requests from that IP address.
`IP block fragment`	Number of IP block fragments.
`Destination session limit`	Number of concurrent sessions that can be directed to a single destination IP address.
`UDP address sweep`	Number of UDP address sweeps.

Table 1332: show security screen statistics Output Fields *(Continued)*

`IPv6 extension header`	Number of packets filtered for the defined IPv6 extension headers.
`IPv6 extension hop by hop option`	Number of packets filtered for the defined IPv6 hop-by-hop option types.

IPv6 extension destination option	Number of packets filtered for the defined IPv6 destination option types.
IPv6 extension header limit	Number of packets filtered for crossing the defined IPv6 extension header limit.
IPv6 malformed header	Number of IPv6 malformed headers defined for the intrusion detection service (IDS).
ICMPv6 malformed packet	Number of ICMPv6 malformed packets defined for the IDS options.

Sample Output

show security screen statistics zone scrzone

```
user@host> show security screen statistics zone scrzone
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              91
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
  TCP land attack                            0
  TCP SYN fragment                           0
  TCP no flag                                0
  IP unknown protocol                        0
  IP bad options                             0
  IP record route option                     0
  IP timestamp option                        0
  IP security option                         0
```

```
  IP loose source route option              0
  IP strict source route option             0
  IP stream option                          0
  ICMP fragment                             0
  ICMP large packet                         0
  TCP SYN FIN                               0
  TCP FIN no ACK                            0
  Source session limit                      0
  TCP SYN-ACK-ACK proxy                     0
  IP block fragment                         0
  Destination session limit                 0
```

Sample Output

show security screen statistics zone untrust (IPv6)

```
user@host>show security screen statistics zone untrust
Screen statistics:

IDS attack type                           Statistics
  ICMP flood                                    0
  UDP  flood                                    0
  TCP  winnuke                                  0
  ......
  IPv6   extension header                        0
  IPv6   extension hop by hop option             0
  IPv6   extension destination option            0
  IPv6   extension header limit                  0
  IPv6   malformed header                        0
  ICMPv6 malformed packet                        0
```

Sample Output

show security screen statistics interface ge-0/0/3

```
user@host> show security screen statistics interface ge-0/0/3
Screen statistics:
IDS attack type                              Statistics
```

```
  ICMP flood                                0
  UDP flood                                 0
  TCP winnuke                               0
  TCP port scan                             91
  ICMP address sweep                        0
  TCP sweep                                 0
  UDP sweep                                 0
  IP tear drop                              0
  TCP SYN flood                             0
  IP spoofing                               0
  ICMP ping of death                        0
  IP source route option                    0
  TCP land attack                           0
  TCP SYN fragment                          0
  TCP no flag                               0
  IP unknown protocol                       0
  IP bad options                            0
  IP record route option                    0
  IP timestamp option                       0
  IP security option                        0
  IP loose source route option              0
  IP strict source route option             0
  IP stream option                          0
  ICMP fragment                             0
  ICMP large packet                         0
  TCP SYN FIN                               0
  TCP FIN no ACK                            0
  Source session limit                      0
  TCP SYN-ACK-ACK proxy                     0
  IP block fragment                         0
  Destination session limit                 0
```

Sample Output

show security screen statistics interface ge-0/0/1 (IPv6)

```
user@host> show security screen statistics interface ge-0/0/1

Screen statistics:

IDS attack type                                 Statistics
```

```
  ICMP flood                                          0
  UDP flood                                           0
  ......
  IPv6 extension header                               0
  IPv6 extension hop by hop option                    0
  IPv6 extension destination option                   0
  IPv6 extension header limit                         0
  IPv6 malformed header                               0
  ICMPv6 malformed packet                             0
```

Sample Output

show security screen statistics interface ge-0/0/1 node primary

```
user@host> show security screen statistics interface ge-0/0/1 node primary
node0:
--------------------------------------------------------------------------
Screen statistics:
IDS attack type                               Statistics
  ICMP flood                                  1
  UDP flood                                   1
  TCP winnuke                                 1
  TCP port scan                               1
  ICMP address sweep                          1
  TCP sweep                                   1
  UDP sweep                                   1
  IP tear drop                                1
  TCP SYN flood                               1
  IP spoofing                                 1
  ICMP ping of death                          1
  IP source route option                      1
  TCP land attack                             1
  TCP SYN fragment                            1
  TCP no flag                                 1
  IP unknown protocol                         1
  IP bad options                              1
  IP record route option                      1
  IP timestamp option                         1
  IP security option                          1
  IP loose source route option                1
  IP strict source route option               1
```

```
  IP stream option                          1
  ICMP fragment                             1
  ICMP large packet                         1
  TCP SYN FIN                               1
  TCP FIN no ACK                            1
  Source session limit                      1
  TCP SYN-ACK-ACK proxy                     1
  IP block fragment                         1
  Destination session limit                 1
```

Sample Output

show security screen statistics zone trust logical-system all

```
user@host> show security screen statistics zone trust logical-system all
Logical system: root-logical-system
Screen statistics:

IDS attack type                             Statistics
  ICMP flood                                0
  UDP flood                                 0
  TCP winnuke                               0
  TCP port scan                             0
  ICMP address sweep                        0
  TCP sweep                                 0
  UDP sweep                                 0
  IP tear drop                              0
  TCP SYN flood                             0
  IP spoofing                               0
  ICMP ping of death                        0
  IP source route option                    0
  TCP land attack                           0
  TCP SYN fragment                          0
  TCP no flag                               0
  IP unknown protocol                       0
  IP bad options                            0
  IP record route option                    0
  IP timestamp option                       0
  IP security option                        0
  IP loose source route option              0
  IP strict source route option             0
```

```
  IP stream option                          0
  ICMP fragment                             0
  ICMP large packet                         0
  TCP SYN FIN                               0
  TCP FIN no ACK                            0
  Source session limit                      0
  TCP SYN-ACK-ACK proxy                     0
  IP block fragment                         0
  Destination session limit                 0

Logical system: ls1
Screen statistics:

IDS attack type                             Statistics
  ICMP flood                                0
  UDP flood                                 0
  TCP winnuke                               0
  TCP port scan                             0
  ICMP address sweep                        0
  TCP sweep                                 0
  UDP sweep                                 0
  IP tear drop                              0
  TCP SYN flood                             0
  IP spoofing                               0
  ICMP ping of death                        0
  IP source route option                    0
  TCP land attack                           0
  TCP SYN fragment                          0
  TCP no flag                               0
  IP unknown protocol                       0
  IP bad options                            0
  IP record route option                    0
  IP timestamp option                       0
  IP security option                        0
  IP loose source route option              0
  IP strict source route option             0
  IP stream option                          0
  ICMP fragment                             0
  ICMP large packet                         0
  TCP SYN FIN                               0
  TCP FIN no ACK                            0
  Source session limit                      0
  TCP SYN-ACK-ACK proxy                     0
```

```
  IP block fragment                          0
  Destination session limit                  0

Logical system: ls2
Screen statistics:

IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              0
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
  TCP land attack                            0
  TCP SYN fragment                           0
  TCP no flag                                0
  IP unknown protocol                        0
  IP bad options                             0
  IP record route option                     0
  IP timestamp option                        0
  IP security option                         0
  IP loose source route option               0
  IP strict source route option              0
  IP stream option                           0
  ICMP fragment                              0
  ICMP large packet                          0
  TCP SYN FIN                                0
  TCP FIN no ACK                             0
  Source session limit                       0
  TCP SYN-ACK-ACK proxy                      0
  IP block fragment                          0
  Destination session limit                  0
```

show security screen statistics zone trust tenant TN1

```
user@host> show security screen statistics zone trust tenant TN1

  Screen statistics:

  IDS attack type                             Statistics
  ICMP flood                                  0
  UDP flood                                   0
  TCP winnuke                                 0
  TCP port scan                               0
  UDP port scan                               0
  ICMP address sweep                          0
  TCP sweep                                   0
  UDP sweep                                   0
  IP tear drop                                0
  TCP SYN flood                               0
  SYN flood source                            0
  SYN flood destination                       0
  IP spoofing                                 0
  ICMP ping of death                          0
  IP source route option                      0
  TCP land attack                             0
  TCP SYN fragment                            0
  TCP no flag                                 0
  IP unknown protocol                         0
  IP bad options                              0
  IP record route option                      0
  IP timestamp option                         0
  IP security option                          0
  IP loose source route option                0
  IP strict source route option               0
  IP stream option                            0
  ICMP fragment                               0
  ICMP large packet                           0
  TCP SYN FIN                                 0
  TCP FIN no ACK                              0
  Source session limit                        0
  TCP SYN-ACK-ACK proxy                       0
  IP block fragment                           0
  Destination session limit                   0
  IPv6 extension header                       0
```

```
  IPv6 extension hop by hop option           0
  IPv6 extension destination option          0
  IPv6 extension header limit                0
  IPv6 malformed header                      0
  ICMPv6 malformed packet                    0
  IP tunnel summary                          0
```

show security screen statistics zone trust tenant all

```
user@host> show security screen statistics zone trust tenant all

  Logical system: root-logical-system
  creen statistics:

  IDS attack type                            Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              0
  UDP port scan                              0
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  SYN flood source                           0
  SYN flood destination                      0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
  TCP land attack                            0
  TCP SYN fragment                           0
  TCP no flag                                0
  IP unknown protocol                        0
  IP bad options                             0
  IP record route option                     0
  IP timestamp option                        0
  IP security option                         0
  IP loose source route option               0
  IP strict source route option              0
  IP stream option                           0
```

```
ICMP fragment                              0
ICMP large packet                          0
TCP SYN FIN                                0
TCP FIN no ACK                             0
Source session limit                       0
TCP SYN-ACK-ACK proxy                      0
IP block fragment                          0
Destination session limit                  0
IPv6 extension header                      0
IPv6 extension hop by hop option           0
IPv6 extension destination option          0
IPv6 extension header limit                0
IPv6 malformed header                      0
ICMPv6 malformed packet                    0
IP tunnel summary                          0

Tenant: TN1
Screen statistics:

IDS attack type                            Statistics
ICMP flood                                 0
UDP flood                                  0
TCP winnuke                                0
TCP port scan                              0
UDP port scan                              0
ICMP address sweep                         0
TCP sweep                                  0
UDP sweep                                  0
IP tear drop                               0
TCP SYN flood                              0
SYN flood source                           0
SYN flood destination                      0
IP spoofing                                0
ICMP ping of death                         0
IP source route option                     0
TCP land attack                            0
TCP SYN fragment                           0
TCP no flag                                0
IP unknown protocol                        0
IP bad options                             0
IP record route option                     0
IP timestamp option                        0
IP security option                         0
```

```
IP loose source route option              0
IP strict source route option             0
IP stream option                          0
ICMP fragment                             0
ICMP large packet                         0
TCP SYN FIN                               0
TCP FIN no ACK                            0
Source session limit                      0
TCP SYN-ACK-ACK proxy                     0
IP block fragment                         0
Destination session limit                 0
IPv6 extension header                     0
IPv6 extension hop by hop option          0
IPv6 extension destination option         0
IPv6 extension header limit               0
IPv6 malformed header                     0
ICMPv6 malformed packet                   0
IP tunnel summary                         0
```

Release Information

Command introduced in Junos OS Release 8.5.

The `node` option added in Junos OS Release 9.0.

The `logical-system all` option added in Junos OS Release 11.2R6.

Support for IPv6 extension header screens added in Junos OS Release 12.1X46-D10.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

clear security screen statistics | 4038

clear security screen statistics interface | 4039

clear security screen statistics zone | 4042

Example: Configuring Multiple Screening Options

show security screen status

IN THIS SECTION

- Syntax | 25132
- Description | 25132
- Required Privilege Level | 25132
- Sample Output | 25132
- Release Information | 25132

Syntax

```
show security screen status
```

Description

Show screen status data.

Required Privilege Level

view

Sample Output

show security screen status

```
user@host> show security screen status
Screen status:
      Screen trap interval : 2 second(s)
```

Release Information

Command introduced in Junos OS Release 12.3X48-D20.

show security screen white-list

IN THIS SECTION

- Syntax | 25133
- Description | 25133
- Options | 25133
- Required Privilege Level | 25134
- Output Fields | 25134
- Sample Output | 25134
- Sample Output | 25134
- Release Information | 25135

Syntax

```
show security screen white-list
<wlist-name>
logical-system
root-logical-system
tenant
```

Description

Display a set of IP addresses for allowlist. Allowlists allows users to download files that are known to be safe. Allowlists can be added to in order to decrease false positives.

Options

wlist-name	The name of the allowlist.
logical-system	Name of the logical system.
root-logical-system	(Default) Displays root logical system as default.
tenant	The name of the tenant system.

Required Privilege Level

view

Output Fields

Table 1333 on page 25134 lists the output fields for the `show security screen ids-option` command. Output fields are listed in the approximate order in which they appear.

Table 1333: show security screen ids-option Output Fields

Field Name	Field Description
`Logical system`	Name of the logical system.
`Tenant`	Name of the tenant system.
`Screen white list`	Display the allowlist of IP addresses.

Sample Output

show security screen allowlist a1 tenant TN1

```
user@host> show security screen white-list a1 tenant TN1

Screen white list:

21.23.24.25/32
21.23.24.251/32
```

Sample Output

show security screen allowlist a1 tenant all

```
user@host> show security screen white-list a1 tenant all

Logical system: root-logical-system
```

```
Screen white list:

2001::2/128
2001::23/128

Tenant: TN1
Screen white list:

21.23.24.25/32
21.23.24.251/32
```

Release Information

Statement introduced in Junos OS Release 12.1.

Statement for `UDP flood` screen allowlist introduced in Junos OS Release 17.4.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

white-list (Security Screen) | **31107**

Example: Configuring Multiple Screening Options

show security shadow-policies

IN THIS SECTION

- Syntax | **25136**
- Description | **25136**
- Options | **25136**
- Required Privilege Level | **25136**
- Output Fields | **25137**
- Sample Output | **25137**
- Release Information | **25138**

Syntax

```
show security shadow-policies
<from-zone zone-name>
<global>
<logical-system logical-system-name>
<policy policy-name>
<reverse>
<root-logical-system>
<tenant tenant-name>
<to-zone zone-name>
```

Description

Displays the shadowing and shadowed policies in a policy list. The output displays the list of all policies that shadows other policies. The concept of policy shadowing refers to the situation where a policy higher in the policy list always takes effect before a subsequent policy. Because the policy lookup always uses the first policy it finds that matches the five-part tuple of the source and destination zone, source and destination address, and application type, if another policy applies to the same tuple (or a subset of the tuple), the policy lookup uses the first policy in the list and never reaches the second one. The existing show command for security shadow-policy is enhanced with tenant support.

Options

- `from-zone` *`zone-name`*—Displays the name or ID of the source zone of the traffic.
- `global`—Displays the information about global policies.
- `logical-system`—Displays the name of the logical system.
- `policy`—Displays the shadow policy information for the given policy.
- `reverse`—Displays the policies which shadows the given policy.
- `root-logical-system`—Displays root logical system as default.
- `tenant`—Displays the name of the tenant system.
- `to-zone` *`zone-name`*—Displays the shadow policy information for the given destination zone.

Required Privilege Level

view

Output Fields

Table 1334 on page 25137 lists the output fields for the `show security shadow-policies logical-system` command. Output fields are listed in the approximate order in which they appear.

Table 1334: show security shadow-policies logical-system Output Fields

Field Name	Field Description
`Policies`	The policies shadowing one or more policies in the policy list.
`Shadowed policies`	The policies shadowed by one or more policies in the policy list.

Sample Output

show security shadow-policies from-zone zone-a to-zone zone-b

```
root@host> show security shadow-policies from-zone zone-a to-zone zone-b
    Policies              Shadowed policies
    P1                    P3
    P1                    P4
    P2                    P5
```

show security shadow-policies from-zone zone-a to-zone zone-b policy P1

```
root@host> show security shadow-policies from-zone zone-a to-zone zone-b policy P1
   Policies              Shadowed policies
    P1                    P3
    P1                    P4
```

show security shadow-policies from-zone zone-a to-zone zone-b policy P4 reverse

```
root@host> show security shadow-policies from-zone zone-a to-zone zone-b policy P4 reverse
   Policies              Shadowed policies
    P1                    P4
```

show security shadow-policies tenant TN1 from-zone trust to-zone untrust

```
user@host> show security shadow-policies tenant TN1 from-zone trust to-zone untrust
Policies                  Shadowed policies
    p12                       p11
```

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

The `tenant` option is added in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

show security policies | 25028

show security softwires

IN THIS SECTION

- Syntax | 25138
- Description | 25139
- Options | 25139
- Required Privilege Level | 25139
- Sample Output | 25139
- Release Information | 25140

Syntax

```
show security softwires <softwire-name softwire-name>
<logical-system (all | logical-system-name)>
```

Description

Display a summary of information of all the softwire concentrators and details on concentrators with specified name.

Options

softwire-name ***softwire-name***	Display the details of the specified softwire concentrator.
logical-system (all | *logical-system-name*)	Display softwire information for all logical systems or for a specified logical system. This option is only available to the primary administrator.

Required Privilege Level

view

Sample Output

command-name

```
user@host> show security softwires
Softwire Name            SC Address         Status   Number of SI connected
  SC-CSSI-1                 3001::1          Connected      2
  SC-CSSI-str00             3100::1          Active         0
  SC-CSSI-str01             3101::1          Inactive       0
  SC-CSSI-str02             3001::1          Connected      2520
```

command-name

```
user@host> show security softwires softwire-name SC-CSSI-1
Name of softwire: SC-CSSI-1
   SC status: Connected
   SC address: 3001::1
   Zone: trust
   VR ID: 0
    SI Address               SI Status                    SPU
    3001::2                  Active                       spu-1
```

```
    3001::2                  Active                      spu-21
    SI number: 2
```

command-name

```
user@host> show security softwires logical-system ls-product-design
Softwire Name            SC Address         Status   Number of SI connected
  sc_1                      3000::1         Connected      1
```

Release Information

Command introduced in Junos OS Release 10.4. The `logical-system` option introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security softwires map-e confidentiality status

IN THIS SECTION

- Syntax | 25141
- Description | 25141
- Required Privilege Level | 25141
- Output Fields | 25141
- Sample Output | 25141
- Release Information | 25142

Syntax

```
show security softwires map-e confidentiality status
```

Description

Displays the MAP-E confidentiality status is enabled or disabled.

Required Privilege Level

admin

Output Fields

Table 1335 on page 25141 lists the output fields for the `show security softwires map-e confidentiality status` command. Output fields are listed in the approximate order in which they appear.

Table 1335: show security softwires map-e confidentiality status

Field Name	Field Description
`Domain`	Displays the specific MAP-E domain for which the confidentiality status is being shown.
`Confidentiality`	Displays the MAP-E confidentiality status is enabled or disabled. This option allows you to hide the MAP-E parameters in the show command output for non-super users. When enabled, the MAP-E parameters are hidden. When disabled, the parameters are visible in the command output.

Sample Output

show security softwires map-e confidentiality status

```
user@host> show security softwires map-e confidentiality status

Domain           : domain1
Confidentiality  : Enabled
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

map-e | 10329

show security softwires map-e domain

IN THIS SECTION

- Syntax | 25142
- Description | 25142
- Required Privilege Level | 25142
- Output Fields | 25143
- Sample Output | 25143
- Release Information | 25144

Syntax

```
show security softwires map-e domain
```

Description

Displays the details of the MAP-E BMR and FMR rules that provides mapping between the IPv4 network and IPv6 network.

Required Privilege Level

admin

Output Fields

Sample Output

show security softwires map-e domain (with confidentiality enabled)

```
user@host> show security softwires map-e domain domain_1

Role            : CE
Domain Name     : domain_1
BR Address      : confidential
End User Ipv6 prefix: 2001:db6:12:3400::/56
BMR Mapping Rule :
    Rule Name            : bmr
FMR Mapping Rule :
    Rule Name            : fmr-1
    Rule Name            : fmr-2
    Rule Name            : fmr-3
    …
    Rule Name            : fmr-n
```

show security softwires map-e domain (without confidentiality enabled)

```
user@host> show security softwires map-e domain domain_1

Role                  : CE
Domain Name           : domain_1
BR Address            : 2001:db6:ffff::1/128
End User Ipv6 prefix  : 2001:db6:12:3400::/56
BMR Mapping Rule :
    Rule Name            : bmr
    Rule Ipv4 Prefix     : 6.0.0.1/32
    Rule Ipv6 Prefix     : 2001::3400::/56
    PSID offset          : 4
    PSID length          : 0
    EA bit length        : 0
    Port SetID           : 0
FMR Mapping Rule :
    Rule Name            : fmr
```

```
    Rule Ipv4 Prefix     : 9.0.0.1/32
    Rule Ipv6 Prefix     : 2002::3400::/56
    PSID offset          : 4
    PSID length          : 0
    EA bit length        : 0
    Port SetID           : 0
```

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

map-e | 10329

show security ssh key-pair-identity

IN THIS SECTION

- Syntax | 25144
- Description | 25145
- Options | 25145
- Required Privilege Level | 25145
- Output Fields | 25145
- Sample Output | 25145
- Release Information | 25146

Syntax

```
show security ssh key-pair-identity
    ( brief  <identity-name>  | public  identity-name  )
```

Description

Display the SSH key pair identity information.

Options

- `brief` *`identity-name`*—Display the brief information for a specified identity. The *identity-name* variable is optional, if an identity is not specified, the command will list brief information of all identities.
- `public` *`identity-name`* —Display the public key for a specified identity.

> **NOTE**: The `public` and `brief` options are mutually exclusive

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

show security ssh key-pair-identity brief

```
user@host> show security ssh key-pair-identity brief
SSH Key Pair Identity Information:
  Name             Create Time      Encrypted
  sample           Dec 28, 17:40    yes
  identity-name    Dec 28, 17:26    yes
```

show security ssh key-pair-identity brief sample

```
user@host> show security ssh key-pair-identity brief sample
SSH Key Pair Identity Information:
  Name             Create Time      Encrypted
  sample           Dec 28, 17:34    yes
```

Release Information

Command introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

request security ssh key-pair-identity generate | 14865
clear security ssh key-pair-identity | 4044

show security tcp-encap connection

IN THIS SECTION

- Syntax | 25146
- Description | 25146
- Options | 25147
- Required Privilege Level | 25147
- Output Fields | 25147
- Sample Output | 25148
- Release Information | 25148

Syntax

```
show security tcp-encap connection
<brief | detail>
<session-id session-id>
```

Description

Display information about TCP encapsulation sessions.

Options

none	Display information about TCP encapsulation sessions.
brief | detail	(Optional) Display the specified level of output.
session-id ***session-id***	(Optional) Display information for the specified session identifier.

Required Privilege Level

view

Output Fields

Table 1336 on page 25147 lists the output fields for the `show security tcp-encap connection` command. Output fields are listed in the approximate order in which they appear.

Table 1336: show security tcp-encap connection Output Fields

Field Name	Field Description
Session-Id	Session identifier.
Client	Name of the remote access client.
Gateway	IP address of the remote gateway.
Local Gateway	IP address of the local gateway.
Remote Gateway	IP address of the remote gateway.
Started	Date and time the connection started.
Anchor spu	Services Processing Unit (SPU) on which the connection is anchored.

Sample Output

show security tcp-encap connection

```
user@host> show security tcp-encap connection
Session-Id    Client            Gateway
   34         NCP-1              10.4.0.1
   644        NCP-1              10.5.0.1
```

show security tcp-encap connection detail

```
user@host> show security tcp-encap connection detail
Session id: 34
    Local Gateway: 10.4.0.2:500 , Remote Gateway: 10.4.0.1:9500
    Client: NCP-1
    Started: Sun Jan 08 2017 21:32:58
    Anchor spu: 1

Session id: 644
    Local Gateway: 10.4.0.2:443 , Remote Gateway: 10.5.0.1:9500
    Client: NCP-1
    Started: Sun Jan 08 2017 21:32:58
    Anchor spu: 1
```

show security tcp-encap connection session-id 644

```
user@host> show security tcp-encap connection session-id 644
Session id: 644
    Local Gateway: 10.4.0.2:443 , Remote Gateway: 10.5.0.1:9500
    Client: NCP-1
    Started: Sun Jan 08 2017 21:32:58
    Anchor spu: 1
```

Release Information

Command introduced in Junos OS Release 15.1X49-D80.

RELATED DOCUMENTATION

tcp-encap | 28980

show security tcp-encap statistics

IN THIS SECTION

- Syntax | 25149
- Description | 25149
- Required Privilege Level | 25149
- Output Fields | 25149
- Sample Output | 25150
- Release Information | 25150

Syntax

```
show security tcp-encap statistics
```

Description

Display TCP encapsulation statistics.

Required Privilege Level

view

Output Fields

Table 1337 on page 25150 lists the output fields for the `show security tcp-encap statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1337: show security tcp-encap statistics Output Fields

Field Name	Field Description
Policy Matched	Number of policies matched.
TCP sessions	Number of TCP sessions.

Sample Output

show security tcp-encap statistics

```
user@host> show security tcp-encap statistics
TCP encapsulation statistics:
  Policy Matched:                    16
  TCP sessions:                      16
```

Release Information

Command introduced in Junos OS Release 15.1X49-D80.

RELATED DOCUMENTATION

clear security tcp-encap statistics | 4046

show security tpm status

IN THIS SECTION

- Syntax | 25151
- Description | 25151
- Options | 25151

- Required Privilege Level | 25151
- Output Fields | 25151
- Sample Output | 25152
- Release Information | 25153

Syntax

```
show security tpm status
```

Description

Display the current status of the Trusted Platform Module (TPM). You can use this `show security tpm status` command to check the status of TPM ownership, master binding key, master encryption password, family version, and firmware version.

Options

This command has no options.

Required Privilege Level

security

Output Fields

Table 1338 on page 25151 lists the output fields for the `show security tpm status` command.

Table 1338: show security tpm status Output Fields

Field Name	Field Description
`Enabled`	Specifies whether TPM is enabled or disabled.

Table 1338: show security tpm status Output Fields *(Continued)*

Field Name	Field Description
Owned	Specifies the TPM ownership. TPM can be owned even if the Master Encryption Key and Master Encryption Key are not created/configured.
Master Binding Key	Displays the TPM's Master Binding Key status whether it is created or not created. TPM generates cryptographic keys and encrypts them so that those can only be decrypted by the TPM. This process is know as binding. Each TPM has a master binding key, which is also know as storage root key.
Master Encryption Key	Displays Master Encryption Password status whether it is set or not set. The encrypted data and the hash of the configuration is protected by the TPM module using the master encryption password.
TPM Family	Displays Trusted Computing Group's (TCG) TPM family version.
TPM Firmware version	Displays the firmware version loaded in TPM.

Sample Output

show security tpm status

```
user@host> show security tpm status
TPM Status:
  Enabled: yes
  Owned: yes
  Master Binding Key: not-created
  Master Encryption Key: not-configured
  TPM Family: 1.2
  TPM Firmware version: 4.43
```

Release Information

Command introduced in Junos OS Release 15.1X49-D80.

Command introduced in Junos OS Release 20.1R1 for SRX5400, SRX5600, and SRX5800 devices with SRX5K-RE3-128G Routing Engine (RE3).

TPM family and TPM firmware version details are introduced in Junos OS Release 15.1X49-D120.

RELATED DOCUMENTATION

Using Trusted Platform Module to Bind Secrets on SRX Series Devices

request security tpm master-encryption-password set | **14871**

show security tunnel inspection

IN THIS SECTION

- Syntax | **25154**
- Description | **25154**
- Required Privilege Level | **25154**
- Output Fields | **25154**
- Sample Output | **25155**
- Release Information | **25156**

Syntax

```
show security tunnel-inspection profiles
```

```
<profiles>
```

```
<vnis>
```

Description

Display details of the tunnel inspection profile created for EVPN-VXLAN tunnel traffic.

Required Privilege Level

View

Output Fields

Table 1339: show security tunnel-inspection profiles Output Fields

Field Name	Field Description
Profile count	Number of tunnel inspection profiles.
Profile	Tunnel inspection profile name
Type	Type of tunnel inspection profile. Displays VXLAN for the profile created to inspect EVPN-VXLAN tunnel traffic.
Vxlan count	Number of associated VXLAN identifiers.
Vxlan name	VXLAN tunnel identifier used to uniquely identify for security inspection.
VNI count	Number of virtual network identifier (VNI) associated.

Table 1339: show security tunnel-inspection profiles Output Fields *(Continued)*

Field Name	Field Description
VNI	Name of the VNI.
VNI id count	VNI range.
Policy set	Associated policy set for inner session inspection.
Inspection level	Level of inspection.

Sample Output

show security tunnel-inspection profiles

```
user@host> show security tunnel-inspection profiles
```

```
Logical system: root-logical-system
  Profile count: 1
  Profile: TP-1
    Type: VXLAN
    Vxlan count: 1
    Vxlan name: VXT-1
    VNI count: 1
      VNI:VNI-1
      Policy set: PSET-1
      Inspection level: 1
```

show security tunnel-inspection vnis

```
user@host> show security tunnel-inspection vnis
```

```
Logical system: root-logical-system
```

```
VNI count: 1
VNI name: VNI-1
    VNI id count: 1
    [50 - 100]
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Tunnel Inspection for EVPN-VXLAN on SRX Series Devices

show security user-identification device-provision authentication-source active-directory start 1 count 9 (match-string|prefix)

IN THIS SECTION

- Syntax | 25156
- Description | 25157
- Options | 25157
- Required Privilege Level | 25157
- Output Fields | 25157
- Sample Output | 25157
- Release Information | 25157

Syntax

```
show security user-identification device-provision authentication-source active-directory start
1 count 9 match-string <match-string-name>$
```

Description

All the devices from Active Directory (AD) are listed with matching string. Displays the list and matched string pattern devices.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the " show security user-identification user-provision authentication-source jims start 1 count 9 (match-string|prefix)" on page 25174 command.

Sample Output

command-name

```
user@host> show security user-identification device-provision authentication-source active-
directory start 1 count 9 match-string <match-string-name>$
node0:
--------------------------------------------------------------------------
Total num: 1
domain-name\device-name$
Total list count: 1
Total matched count: 1
```

Release Information

Command introduced in Junos OS Release 20.2R1.

show security user-identification device-provision authentication-source jims start 1 count 9 (match-string|prefix)

IN THIS SECTION

- Syntax | 25158
- Description | 25158
- Options | 25158
- Required Privilege Level | 25158
- Output Fields | 25159
- Sample Output | 25159
- Release Information | 25159

Syntax

```
show security user-identification device-provision authentication-source jims start 1 count 9
match-string <match-string-name>
```

Description

All the devices from JIMS in SRX Series Firewall are listed with the list of pattern matching involves strings of characters.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the " show security user-identification user-provision authentication-source jims start 1 count 9 (match-string|prefix)" on page 25174 command.

Sample Output

command-name

```
user@host> show security user-identification device-provision authentication-source jims start 1
count 9 match-string <match-string-name>
node0:
--------------------------------------------------------------------------
Total num: 2
domain-name\device-name_1$
domain-name\device-name_5$
Total list count: 2
Total matched count: 2
```

Release Information

Command introduced in Junos OS Release 20.2R1.

show security user-identification local-authentication-table

IN THIS SECTION

- Syntax | 25160
- Description | 25160
- Required Privilege Level | 25160
- Output Fields | 25160
- Sample Output | 25161
- Release Information | 25162

Syntax

```
show security user-identification local-authentication-table [ ( all [brief | extensive]) |ip-address ip-address | role role-name | start value | count value | user user-name]
```

Description

This command displays the content of the local authentication table by IP address.

all (Optional) All entries displayed from the beginning of the table or from the specified starting entry.

- **brief** (Default) Uses a tabular format and truncates longer entries: username—displays up to 13 characters, roles—displays up to 32 characters.
- **extensive** (Optional) Displays the full names and all items.

count *value* (Optional) The total number of entries to display.

ip-address *ip-address* (Optional) The IP address of the entry to display.

role *role-name* (Optional) The role name of the entries to display.

start *value* (Optional) The first entry to display.

user *user-name* (Optional) The username of the entry to display.

Required Privilege Level

view

Output Fields

Table 1340 on page 25161 lists the output fields for the `show security user-identification local-authentication-table` command. Output fields are listed in the approximate order in which they appear.

Table 1340: show security user-identification local-authentication-table Output Fields

Field Name	Field Description
`Total entries`	The number of entries in the table.
`IP address`	IP address of the associated user. **NOTE**: Only one user can be associated with an IP address.
`Username`	User associated with the specified IP address. Username is case insensitive in local authenitcation table.
`Roles`	A comma-separated list of all roles associated with this IP address and user. Role names are case insensitive in local authentication table.

Sample Output

show security user-identification local-authentication-table all

```
user@host> show security user-identification local-authentication-table all
Total entries: 3
Source IP       Username     Roles
192.0.2.1       user1        role1
203.0.113.2     user1        role2
198.51.100.3    user3        role1, role2
```

show security user-identification local-authentication-table ip-address

```
user@host> show security user-identification local-authentication-table ip-address 203.0.113.2
Ip-address: 203.0.113.2
    Username: user2
    Roles: role2, role3, role1
```

show security user-identification local-authentication-table start

```
user@host> show security user-identification local-authentication-table start 2 count 2
Total entries: 2
    Ip-address: 203.0.113.2
    Username: user2
    Roles: role2, role3, role1

    Ip-address: 198.51.100.3    Username: user3
    Roles: role2, role3
```

show security user-identification local-authentication-table role

```
user@host> show security user-identification local-authentication-table role qa3456
Total entries: 3
    Ip-address: 203.0.113.2
    Username: dev-grp-3
    Roles: qa432, qa3456, qa84, qa794

    Ip-address: 198.51.100.3
    Username: dev-qa
    Roles: qa3456, qa3985, qa23

    Ip-address: 203.0.113.2
    Username: brandall
    Roles: qa3456
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

request security user-identification local-authorization-table add | **14877**

Understanding the User Identification Table

show security user-identification role-provision all

IN THIS SECTION

- Syntax | 25163
- Description | 25163
- Required Privilege Level | 25163
- Output Fields | 25163
- Sample Output | 25164
- Release Information | 25164

Syntax

```
show security user-identification role-provision all
```

Description

Display all the available user roles for policy provisioning. The output combines user roles from all available UITs.

Required Privilege Level

view

Output Fields

Table 1341 on page 25164 lists the output fields for the `show security user-identification role-provision all` command. Output fields are listed in the approximate order in which they appear.

Table 1341: show security user-identification role-provision all Output Fields

Field Name	Field Description
Roles	A comma-separated list of all user roles available for provisioning in user role policies. This list combines user roles from both the local authentication table and any UAC authentication tables that have been configured.

Sample Output

show security user-identification role-provision all

```
user@host> show security user-identification role-provision all
Roles: role1, role2, role3, role4, role_0_1, role_1_1
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

show security user-identification user-provision all | 25171

show security user-identification source-identity-provision all | 25169

show security user-identification role-provision authentication-source active-directory start 1 count 9 (match-string|prefix)

IN THIS SECTION

- Syntax | 25165
- Description | 25165
- Options | 25165

- Required Privilege Level | 25165
- Output Fields | 25165
- Sample Output | 25165
- Release Information | 25166

Syntax

```
show security user-identification role-provision authentication-source active-directory start 1
count 9 match-string <match-string-name>
```

Description

All the group (user-group or device-group) from Active Directory (AD) are listed using matching string.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the " show security user-identification user-provision authentication-source jims start 1 count 9 (match-string|prefix)" on page 25174 command.

Sample Output

command-name

```
user@host> show security user-identification role-provision authentication-source active-
directory start 1 count 9 match-string <match-string-name>
node0:
```

```
---------------------------------------------------------------------------
Total num: 1
domain-name\group-name_1
Total list count: 1
Total matched count: 1
```

Release Information

Command introduced in Junos OS Release 20.2R1.

show security user-identification role-provision authentication-source jims start 1 count 9 (match-string|prefix)

IN THIS SECTION

- Syntax | 25166
- Description | 25167
- Options | 25167
- Required Privilege Level | 25167
- Output Fields | 25167
- Sample Output | 25167
- Release Information | 25169

Syntax

```
show security user-identification role-provision authentication-source jims start 1 count 9
match-string <match-string-name>
```

Description

All the groups (user-group or device-group) in SRX Series Firewall from JIMS are listed with the pattern matching.

Match the Active Directory domain name to the primary DNS suffix of the computer name.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the " show security user-identification user-provision authentication-source jims start 1 count 9 (match-string|prefix)" on page 25174 command.

Sample Output

command-name

```
user@host> show security user-identification role-provision authentication-source jims start 1
count 9 match-string <match-string-name>
node0:
--------------------------------------------------------------------------
Total num: 58
domain-name\Access Control Assistance Operators
domain-name\Account Operators
domain-name\Administrators
domain-name\Alloyoud RODC Password Replication Group
domain-name\Backup Operators
domain-name\Cert Publishers
domain-name\Certificate Service DCOM Access
domain-name\Cloneable Domain Controllers
domain-name\Cryptographic Operators
domain-name\Denied RODC Password Replication Group
domain-name\Distributed COM Users
```

```
domain-name\DnsAdmins
domain-name\DnsUpdateProxy
domain-name\Domain Admins
domain-name\Domain Computers
domain-name\Domain Controllers
domain-name\Domain Guests
domain-name\Domain Users
domain-name\Enterprise Admins
domain-name\Enterprise Read-only Domain Controllers
domain-name\Event Log Readers
domain-name\Group Policy Creator Owners
domain-name\Guests
domain-name\Hyper-V Administrators
domain-name\IIS_IUSRS
domain-name\Incoming Forest Trust Builders
domain-name\Network Configuration Operators
domain-name\Performance Log Users
domain-name\Performance Monitor Users
domain-name\Pre-Windows 2000 Compatible Access
domain-name\Print Operators
domain-name\Protected Users
domain-name\RAS and IAS Servers
domain-name\RDS Endpoint Servers
domain-name\RDS Management Servers
domain-name\RDS Remote Access Servers
domain-name\Read-only Domain Controllers
domain-name\Remote Desktop Users
domain-name\Remote Management Users
domain-name\Replicator
domain-name\Schema Admins
domain-name\Server Operators
domain-name\TelnetClients
domain-name\Terminal Server License Servers
domain-name\Users
domain-name\WinRMRemoteWMIUsers__
domain-name\Windows Authorization Access Group
domain-name\administrators
domain-name\denied rodc password replication group
domain-name\domain admins
domain-name\domain users
domain-name\enterprise admins
domain-name\group policy creator owners
domain-name\group-name
```

```
domain-name\posture-healthy
domain-name\remote desktop users
domain-name\schema admins
domain-name\users
Total list count: 58
Total matched count: 58
```

Release Information

Command introduced in Junos OS Release 20.2R1.

show security user-identification source-identity-provision all

IN THIS SECTION

- Syntax | 25169
- Description | 25169
- Required Privilege Level | 25170
- Output Fields | 25170
- Sample Output | 25170
- Release Information | 25170

Syntax

```
show security user-identification source-identity-provision all
```

Description

Display all source identities used in user identification tables (UITs) untill now. The output includes all historical users and groups which are used till now in UITs.

Display all source identities used in user identification tables (UITs), except local authentication table, until now. For local authentication table, the command displays only currently present source identities

in the table. The output includes all historical users and groups which are used till now in UITs except local authentication table.

Required Privilege Level

view

Output Fields

Table 1342 on page 25170 lists the output fields for the `show security user-identification source-identity-provision all` command. Output fields are listed in the approximate order in which they appear.

Table 1342: show security user-identification source-identity-provision all Output Fields

Field Name	Field Description
`Source identities`	This is a list of all users and user roles used in UITs, except local authentication table, until now. The list includes users and user roles from current local authentication table along with all users and roles in other UITs till now. The list is comma-separated. To use users in policy provisioning, you can utilize the user details currently present in UIT tables. We can view the list of users from the command line interface (CLI) by running the following commands: show services user-identification authentication-table authentication-source all show security user-identification local-authentication-table all .

Sample Output

show security user-identification source-identity-provision all

```
user@host> show security user-identification source-identity-provision all
Source identities: ariana, ben, guest5, role1, role2, role3, role4, role_0_1, role_1_1,u1, user2
```

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

show security user-identification role-provision all | 25163

show security user-identification user-provision all | 25171

show security user-identification user-provision all

IN THIS SECTION

- Syntax | 25171
- Description | 25171
- Required Privilege Level | 25171
- Output Fields | 25171
- Sample Output | 25172
- Release Information | 25172

Syntax

```
show security user-identification user-provision all
```

Description

Display the available user names for policy provisioning. The output combines users from all available UITs.

Required Privilege Level

view

Output Fields

Table 1343 on page 25172 lists the output fields for the `show security user-identification user-provision all` command. Output fields are listed in the approximate order in which they appear.

Table 1343: show security user-identification user-provision all Output Fields

Field Name	Field Description
Users	A comma-separated list of all users available for policy provisioning. This list combines users from both the local authentication table and any UAC authentication tables that have been configured.

Sample Output

show security user-identification user-provision all

```
user@host> show security user-identification user-provision all
Users: ariana, ben, guest5, u1 user2, ...
```

Release Information

Command introduced in Release Junos OS 12.1X44-D10.

RELATED DOCUMENTATION

show security user-identification role-provision all | 25163

show security user-identification source-identity-provision all | 25169

show security user-identification user-provision authentication-source active-directory start 1 count 9 (match-string|prefix)

IN THIS SECTION

- Syntax | 25173
- Description | 25173
- Options | 25173

- Required Privilege Level | 25173
- Output Fields | 25173
- Sample Output | 25173
- Release Information | 25174

Syntax

```
show security user-identification user-provision authentication-source active-directory start 1
count 9 match-string <match-string-value>
```

Description

All user from Active Directory (AD) in SRX Series Firewall are listed with matching string.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the " show security user-identification user-provision authentication-source jims start 1 count 9 (match-string|prefix)" on page 25174 command.

Sample Output

command-name

```
user@host> show security user-identification user-provision authentication-source active-
directory start 1 count 9 match-string <match-string-value>
node0:
```

```
---------------------------------------------------------------------------
Total num: 1
domain-name\user-name
Total list count: 1
Total matched count: 1
```

Release Information

Command introduced in Junos OS Release 20.2R1.

show security user-identification user-provision authentication-source jims start 1 count 9 (match-string|prefix)

IN THIS SECTION

- Syntax | 25174
- Description | 25175
- Options | 25175
- Required Privilege Level | 25175
- Output Fields | 25175
- Sample Output | 25176
- Release Information | 25176

Syntax

```
show security user-identification user-provision authentication-source jims start 1 count 9
match-string <domain-name>
```

Description

All users from JIMS are listed with search pattern matching (string match). Use *match-string* to extract the first text string matching a pattern and match string to extract all text strings that match. *start* indicates first count index and *count* indicates total number to display the string match pattern.

Options

match-string Match string in regular expression. The match-string does not allow regression string matching. The match-string supported is [1-9] or [a-f] format. The match-string is case sensitive.

prefix Prefix

Required Privilege Level

view

Output Fields

Table 1344 on page 25175 lists the output fields for the `show security user-identification user-provision authentication-source jims start 1 count 9 match-string <domain-name>` command. Output fields are listed in the approximate order in which they appear.

Table 1344: show security user-identification user-provision authentication-source jims start 1 count 9 (match-string|prefix)

Field Name	Field Description
`Total list count`	The actual listed count for current command.
`Total matched count`	The total matched number for current match condition.
`Total num`	The total number of items for current xml file.

Sample Output

command-name

```
user@host> show security user-identification user-provision authentication-source jims start 1
count 9 match-string <domain-name>
node0:
--------------------------------------------------------------------------
Total num: 3
domain-name\user-name_1
domain-name\user-name_2
domain-name\user-name_3
Total list count: 3
Total matched count: 3
```

Release Information

Command introduced in Junos OS Release 20.2R1.

show security utm anti-spam statistics

IN THIS SECTION

- Syntax | 25177
- Description | 25177
- Options | 25177
- Required Privilege Level | 25178
- Sample Output | 25178
- Release Information | 25181

Syntax

```
show security utm anti-spam statistics
<root-logical-system>
<logical-system (logical-system-name | all)>
<all-logical-systems-tenants>
<tenant (tenant-name | all)>
```

Description

Displays antispam statistics for connections including total e-mail scanned, tagged, and dropped connections.

Statistics from both the nodes (with full chassis cluster support for Content Security) are displayed.

Starting in Junos OS Release 18.2R1, on SRX5000 line devices, the options `pic` and `fpc` to view the physical interface cards (PICs) and Flexible PIC Concentrator (FPC) statistics are not supported.

Starting in Junos OS Release 18.3R1, you can view the antispam statistics for the primary logical system or for a specific user logical system or for all the user logical systems.

Starting in Junos OS Release 19.2R1, you can view the antispam statistics for the tenant system.

Options

none Displays antispam statistics for the primary logical system.

root-logical-system (Optional) Displays antispam statistics for the primary logical system.

logical-system *logical-system-name* (Optional) Displays antispam statistics for a specific user logical system.

all (Optional) Displays antispam statistics for all the user logical systems.

<all-logical-systems-tenants> (Optional) Displays antispam statistics for all the logical systems and tenant systems.

tenant *tenant-name* (Optional) Displays antispam statistics for a specific tenant system.

all (Optional) Displays antispam statistics for all the tenant systems.

Required Privilege Level

view

Sample Output

show security utm anti-spam statistics

```
user@host> show security utm anti-spam statistics
Total connections:     0
Denied connections:    0
Total greetings:       0
Denied greetings:      0
Total e-mail scanned:  0
White list hit:        0
Black list hit:        0
Spam total:            0
Spam tagged:           0
Spam dropped:          0
DNS errors:            0
Timeout errors:        0
Return errors:         0
Invalid parameter errors: 0
```

show security utm anti-spam statistics root-logical-system

```
user@host> show security utm anti-spam statistics root-logical-system
 UTM Anti Spam statistics:

Total connections:     0
Denied connections:    0
Total greetings:       0
Denied greetings:      0
Total e-mail scanned:  0
White list hit:        0
Black list hit:        0
Spam total:            0
Spam tagged:           0
Spam dropped:          0
DNS errors:            0
```

```
Timeout errors:        0
Return errors:         0
Invalid parameter errors: 0
```

show security utm anti-spam statistics logical-system LSYS1

```
user@host> show security utm anti-spam statistics logical-system LSYS1
 UTM Anti Spam statistics:

Total connections:     0
Denied connections:    0
Total greetings:       0
Denied greetings:      0
Total e-mail scanned:  0
White list hit:        0
Black list hit:        0
Spam total:            0
Spam tagged:           0
Spam dropped:          0
DNS errors:            0
Timeout errors:        0
Return errors:         0
Invalid parameter errors: 0
```

show security utm anti-spam statistics logical-system all

```
user@host> show security utm anti-spam statistics logical-system all
 UTM Anti Spam statistics:

Total connections:     0
Denied connections:    0
Total greetings:       0
Denied greetings:      0
Total e-mail scanned:  0
White list hit:        0
Black list hit:        0
Spam total:            0
Spam tagged:           0
Spam dropped:          0
DNS errors:            0
```

```
Timeout errors:        0
Return errors:         0
Invalid parameter errors: 0
```

show security utm anti-spam statistics tenant TSYS1

```
user@host> show security utm anti-spam statistics tenant TSYS1
 UTM Anti Spam statistics:

Total connections:     0
Denied connections:    0
Total greetings:       0
Denied greetings:      0
Total e-mail scanned:  0
White list hit:        0
Black list hit:        0
Spam total:            0
Spam tagged:           0
Spam dropped:          0
DNS errors:            0
Timeout errors:        0
Return errors:         0
Invalid parameter errors: 0
```

show security utm anti-spam statistics tenant all

```
user@host> show security utm anti-spam statistics tenant all
 UTM Anti Spam statistics:

Total connections:     0
Denied connections:    0
Total greetings:       0
Denied greetings:      0
Total e-mail scanned:  0
White list hit:        0
Black list hit:        0
Spam total:            0
Spam tagged:           0
Spam dropped:          0
DNS errors:            0
```

```
Timeout errors:        0
Return errors:         0
Invalid parameter errors: 0
```

show security utm anti-spam statistics all-logical-systems-tenants

```
user@host> show security utm anti-spam statistics all-logical-systems-tenants
 UTM Anti Spam statistics:

Total connections:     0
Denied connections:    0
Total greetings:       0
Denied greetings:      0
Total e-mail scanned:  0
White list hit:        0
Black list hit:        0
Spam total:            0
Spam tagged:           0
Spam dropped:          0
DNS errors:            0
Timeout errors:        0
Return errors:         0
Invalid parameter errors: 0
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for Content Security in chassis cluster added in Junos OS Release 11.4.

Support for Content Security in logical system added in Junos OS Release 18.3R1.

Support for Content Security in tenant system added in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security utm anti-spam statistics | 4048

show security utm anti-spam status | 25182

show security utm anti-spam status

IN THIS SECTION

- Syntax | 25182
- Description | 25182
- Required Privilege Level | 25182
- Output Fields | 25182
- show security utm anti-spam status | 25183
- Release Information | 25183

Syntax

```
show security utm anti-spam status
```

Description

Display antispam status for connections including allowlist and blocklist server information. Status of both the nodes (with full chassis cluster support for Content Security) is displayed.

Required Privilege Level

view

Output Fields

show security Content Security anti-spam status

Output fields are listed in the approximate order in which they appear.

show security utm anti-spam status

command-name

```
user@host> show security utm anti-spam status
SBL Whitelist Server:
SBL Blacklist Server:
    msgsecurity.example.net

DNS Server:
    Primary  :     1.2.3.4, Src Interface: ge-0/0/0
    Secondary:     0.0.0.0, Src Interface: ge-0/0/1
    Ternary  :     0.0.0.0, Src Interface: fe-0/0/2
```

Release Information

Command introduced in Junos OS Release 9.5 .

Support for Content Security in chassis cluster added in Junos OS Release 11.4 .

RELATED DOCUMENTATION

clear security utm anti-spam statistics | 4048

show security utm anti-spam statistics | 25176

show security utm anti-virus statistics

IN THIS SECTION

- Syntax | 25184
- Description | 25184
- Options | 25184
- Required Privilege Level | 25185

Sample Output | 25185
Release Information | 25190

Syntax

```
show security utm anti-virus statistics
<root-logical-system>
<logical-system (logical-system-name | all)>
<tenant (tenant-name | all)>
<all-logical-systems-tenants>
<fpc <fpc-slot fpc-slot pic-slot pic-slot>>
```

Description

Displays antivirus statistics for connections including clean and infected files, scan engine status, and aggregated statistics from all FPCs and PICs. Statistics from both the nodes (with full chassis cluster support for Content Security) are displayed.

Starting in Junos OS Release 18.3R1, you can view the antivirus statistics for the primary logical system or for a specific user logical system or for all the user logical systems.

Starting in Junos OS Release 19.2R1, you can view the antivirus statistics for a specific tenant system or for all the tenant systems.

Options

none Displays antivirus statistics for the primary logical system.

root-logical-system (Optional) Displays antivirus statistics for the primary logical system.

logical-system *logical-system-name* (Optional) Displays antivirus statistics for a specific user logical system.

all (Optional) Displays antivirus statistics for all the user logical systems.

all-logical-systems-tenants (Optional) Displays antivirus statistics for all the logical systems and tenant systems.

tenant ***tenant-name*** (Optional) Displays antivirus statistics for a specific tenant system.

all (Optional) Displays antispam statistics for all the tenant systems.

Required Privilege Level

view

Sample Output

show security utm anti-virus statistics

```
user@host> show security utm anti-virus statistics
UTM Anti Virus statistics:

 Intelligent-prescreening passed:    0
 MIME-whitelist passed:              0
 URL-whitelist passed:               0
 Session abort:                      0
 Scan Request:

  Total           Clean         Threat-found    Fallback
    14              1                1             1

 File Reputation:        11
                                  Log-and-Permit        Block             Permit
Potentially Unwanted Application:        1                 1                 2
Unknown:               2                 2                 3

 Fallback:
                        Log-and-Permit       Block             Permit
 Engine not ready:            0                    0                        0
 Out of resources:            0                    0                        0
 Timeout:                     0                    0                        0
 Maximum content size:        0                    0                        0
 Too many requests:           0                    0                        0
 Decompress error:            0                    0                        0
 Server connection error:     0                    0                        0
 Others:                      0                    0                        0
```

show security utm anti-virus statistics fpc

```
user@host> show security utm anti-virus statistics fpc

 fpc-slot 5 pic-slot 0
 UTM Anti Virus statistics:
 MIME-whitelist passed:                0
 URL-whitelist passed:                 0
 Scan Request:

  Total           Clean         Threat-found    Fallback
       0               0              0               0

 Fallback:
                              Log-and-Permit    Block             Permit
 Engine not ready:                0                 0                 0
 Out of resources:                0                 0                 0
 Timeout:                         0                 0                 0
 Maximum content size:            0                 0                 0
 Too many requests:               0                 0                 0
 Others:                          0                 0                 0
```

show security utm anti-virus statistics fpc fpc-slot 5 pic-slot 0

```
user@host> show security utm anti-virus statistics fpc fpc-slot 5 pic-slot 0
 UTM Anti Virus statistics:
 MIME-whitelist passed:                0
 URL-whitelist passed:                 0
 Scan Request:

  Total           Clean         Threat-found    Fallback
       0               0              0               0

 Fallback:
                              Log-and-Permit    Block             Permit
 Engine not ready:                0                 0                 0
 Out of resources:                0                 0                 0
 Timeout:                         0                 0                 0
 Maximum content size:            0                 0                 0
```

```
 Too many requests:              0                0                0
 Others:                         0                0                0
```

show security utm anti-virus statistics root-logical-system

```
user@host> show security utm anti-virus statistics root-logical-system
 UTM Anti Virus statistics:
 MIME-whitelist passed:                0
 URL-whitelist passed:                 0
 Session abort:                        0
 Scan Request:

  Total           Clean         Threat-found    Fallback
       0               0              0               0
 Fallback:

                           Log-and-permit         Block
  Engine not ready:               0                     0
  Out of resources:               0                     0
  Timeout:                        0                     0
  Maximum content size:           0                     0
  Too many requests:              0                     0
  Others:                         0                     0
```

show security utm anti-virus statistics logical-system LSYS1

```
user@host> show security utm anti-virus statistics logical-system LSYS1
 UTM Anti Virus statistics:
 MIME-whitelist passed:                0
 URL-whitelist passed:                 0
 Session abort:                        0
 Scan Request:

  Total           Clean         Threat-found    Fallback
       0               0              0               0
 Fallback:

                           Log-and-permit         Block
  Engine not ready:               0                     0
  Out of resources:               0                     0
```

```
  Timeout:                          0                       0
  Maximum content size:             0                       0
  Too many requests:                0                       0
  Others:                           0                       0
```

show security utm anti-virus statistics logical-system all

```
user@host> show security utm anti-virus statistics logical-system all
 UTM Anti Virus statistics:
 MIME-whitelist passed:                0
 URL-whitelist passed:                 0
 Session abort:                        0
 Scan Request:

  Total           Clean         Threat-found    Fallback
       0               0              0               0
 Fallback:

                           Log-and-permit         Block
  Engine not ready:                0                       0
  Out of resources:                0                       0
  Timeout:                         0                       0
  Maximum content size:            0                       0
  Too many requests:               0                       0
  Others:                          0                       0
```

show security utm anti-virus statistics tenant TSYS1

```
user@host> show security utm anti-virus statistics tenant TSYS1
 UTM Anti Virus statistics:
 MIME-whitelist passed:                0
 URL-whitelist passed:                 0
 Session abort:                        0
 Scan Request:

  Total           Clean         Threat-found    Fallback
       0               0              0               0
 Fallback:

                           Log-and-permit         Block
```

```
  Engine not ready:                0                      0
  Out of resources:                0                      0
  Timeout:                         0                      0
  Maximum content size:            0                      0
  Too many requests:               0                      0
  Decompress error:                0                      0
  Others:                          0                      0
```

show security utm anti-virus statistics tenant all

```
user@host> show security utm anti-virus statistics tenant all
 UTM Anti Virus statistics:
 MIME-whitelist passed:                0
 URL-whitelist passed:                 0
 Session abort:                        0
 Scan Request:

  Total           Clean         Threat-found    Fallback
       0               0              0               0
 Fallback:

                           Log-and-permit         Block
  Engine not ready:                0                      0
  Out of resources:                0                      0
  Timeout:                         0                      0
  Maximum content size:            0                      0
  Too many requests:               0                      0
  Decompress error:                0                      0
  Others:                          0                      0
```

show security utm anti-virus statistics all-logical-systems-tenants

```
user@host> show security utm anti-virus statistics all-logical-systems-tenants
 UTM Anti Virus statistics:
 MIME-whitelist passed:                0
 URL-whitelist passed:                 0
 Session abort:                        0
 Scan Request:

  Total           Clean         Threat-found    Fallback
```

```
      0               0              0               0
Fallback:

                          Log-and-permit         Block
 Engine not ready:                0                     0
 Out of resources:                0                     0
 Timeout:                         0                     0
 Maximum content size:            0                     0
 Too many requests:               0                     0
 Decompress error:                0                     0
 Others:                          0                     0
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for Sophos Antivirus added in Junos OS Release 11.1.

Support for Content Security in chassis cluster added in Junos OS Release 11.4.

Support for Flexible PIC Concentrator (FPC) and PIC status added in Junos OS Release 12.1X46-D10.

Starting in Junos OS Release 18.2R1, on SRX5000 line, the options `pic` and `fpc` are deprecated—rather than immediately removed—to provide backward compatibility.

Support for Content Security in logical system added in Junos OS Release 18.3R1.

Support for Content Security in tenant system added in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

- clear security utm antivirus statistics | **4051**
- show security utm anti-virus status | **25191**
- request security utm anti-virus juniper-express-engine | **14881**
- request security utm anti-virus kaspersky-lab-engine | **14883**

show security utm anti-virus status

IN THIS SECTION

- Syntax | 25191
- Description | 25191
- Required Privilege Level | 25191
- Sample Output | 25191
- Release Information | 25194

Syntax

```
show security utm anti-virus status <fpc <fpc-slot fpc-slot pic-slot pic-slot>>
```

Description

Display antivirus status for connections including clean and infected files, scan engine status, and aggregated status from all FPCs and PICs. Status of both the nodes (with full chassis cluster support for Content Security) is displayed.

Required Privilege Level

view

Sample Output

show security utm anti-virus status for Avira engine

```
user@host> show security utm anti-virus status
 UTM anti-virus status:

    Anti-virus key expire date: 2021-05-10 18:23:43
    Update server: https://update.juniper-updates.net/AVIRA/VSRX
           Interval: 1440 minutes
```

```
            Pattern update status: next update in 1228 minutes
            Last result: Downloading file failed
    Forwarding-mode: hold
    Onbox AV load flavor: running Light, configure Light
    Scan engine type: avira-engine
    Scan engine information: 8.3.60.28
    Anti-virus signature version: 8.16.46.24
```

show security utm anti-virus status for Sophos engine (Prior to Junos OS Release 23.1R1)

```
user@host> show security utm anti-virus status
UTM anti-virus status:

    Anti-virus key expire date: 2021-06-07 00:00:00
    Update server: https://update.juniper-updates.net/SAV/
            Interval: 1440 minutes
            Pattern update status: next update in 467 minutes
            Last result: already have latest database
    Forwarding-mode: continuous delivery, inline-tap
    Scan engine type: sophos-engine
    Scan engine information: last action result: No error
    Anti-virus signature version: 1.13 (1.02)
```

show security utm anti-virus status for Sophos engine (Starting in Junos OS Release 23.1R1)

With Web proxy support:

```
user@host> show security utm anti-virus status
    UTM anti-virus status:

        Anti-virus key expire date: 2024-03-01 05:30:00
         Connection Mode: Via Proxy
         Connection Details: host1.example.com:3128 (192.168.2.2:3128)
         Forwarding-mode: continuous delivery
         Scan engine type: sophos-engine
         Scan engine information: running
```

Without Web proxy support:

```
user@host> show security utm anti-virus status
    UTM anti-virus status:

        Anti-virus key expire date: 2024-03-01 05:30:00
        Connection Mode: Direct
        Connection Details: 4.sophosxl.net:443 (203.0.113.0:443)
        Forwarding-mode: continuous delivery
        Scan engine type: sophos-engine
        Scan engine information: running
```

show security utm anti-virus status fpc

```
user@host> show security utm anti-virus status fpc
 fpc-slot 5 pic-slot 0
 UTM anti-virus status:

    Anti-virus key expire date: 2021-06-07 00:00:00
    Update server: https://update.juniper-updates.net/SAV/
           Interval: 1440 minutes
           Pattern update status: next update in 467 minutes
           Last result: already have latest database
    Forwarding-mode: continuous delivery, inline-tap
    Scan engine type: sophos-engine
    Scan engine information: last action result: No error
    Anti-virus signature version: 1.13 (1.02)
```

show security utm anti-virus status fpc fpc-slot 5 pic-slot 0

```
user@host> show security utm anti-virus status fpc fpc-slot 5 pic-slot 0
 UTM anti-virus status:

    Anti-virus key expire date: license not installed
    Update server: http://update.juniper-updates.net/SAV/
           Interval: 1440 minutes
           Pattern update status: update disabled due to no license
           Last result: already have latest database
```

```
    Anti-virus signature version: 000000_00
    Forwarding-mode: continuous delivery, inline-tap
    Scan engine type: sophos-engine
    Scan engine information: last action result: No error
```

show security utm anti-virus status

Refer the sample output for Avira scan engine. Support for Avira is added in 18.4R1 release.

```
UTM anti-virus status:
    Anti-virus key expire date: 2021-03-12 08:00:00
    Update server: https://update.juniper-updates.net/AVIRA/VSRX
            Interval: 1440 minutes
            Pattern update status: av updater is running
            Last result: downloading signature files
    Forwarding-mode: hold
    Scan engine type: sophos-engine
    Scan engine information: last action result: No error
    Anti-virus signature version: 1.13 (1.02)
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for Content Security in chassis cluster added in Junos OS Release 11.4. Support for Flexible PIC Concentrator (FPC) and PIC status added in Junos OS Release 12.1X46-D10.

Starting in Junos OS Release 18.2R1, on SRX5000 line, the options `pic` and `fpc` to display PIC and FPC statistics are not supported.

RELATED DOCUMENTATION

clear security utm antivirus statistics | 4051

show security utm anti-virus statistics | 25183

show security utm content-filtering statistics

IN THIS SECTION

- Syntax | 25195
- Description | 25195
- Options | 25196
- Required Privilege Level | 25196
- Sample Output | 25197
- Release Information | 25206

Syntax

```
show security utm content-filtering statistics
<root-logical-system> utm-policy utm-policy-name;
<logical-system (logical-system-name utm-policy utm-policy-name | all)>;
<all-logical-systems-tenants>;
<tenant (tenant-name | all)>;
utm-policy utm-policy-name;
```

Description

Displays the content filtering statistics for connections including lists of blocked files and the reasons for blocking. Statistics from both the nodes (with full chassis cluster support for Content Security) are displayed.

Starting in Junos OS Release 18.3R1, you can view the content filtering statistics for the primary logical system or for a specific user logical system or for all the user logical systems.

Starting in Junos OS Release 19.2R1, you can view the content filtering statistics for a specific tenant system or for all the tenant systems.

Starting in Junos OS Release 21.4R1, we have improved the Content Security content filtering efficiency by adding pre-detect technique to detect file types accurately. This techinque allows you to create rulesets and rules to define content filter criteria. This enhnacemnet is a replacement of the existing file type detection based on file name extensions and profile based filtering on application profiles. This command is enhnaced to display the enhnaced content filtering statistics in a Content Security policy.

Options

none Displays content filtering statistics for the primary logical system.

root-logical-system (Optional) Displays content filtering statistics for the primary logical system.

`utm-policy` *`utm-policy-name`*

Starting in Junos OS Release 21.4R1, this option is added under the root-logical-system option, to display content filtering statistics in a specified Content Security policy within the root logical system.

logical-system ***logical-system-name*** (Optional) Displays content filtering statistics for a specific user logical system.

`utm-policy` *`utm-policy-name`*

Starting in Junos OS Release this option is added under the root-logical-system option, to display content filtering statistics in a specified Content Security policy within the specified logical system.

all (Optional) Displays content filtering statistics for all the user logical systems.

all-logical-systems-tenants (Optional) Displays content filtering statistics for all logical systems and tenant systems.

tenant ***tenant-name*** (Optional) Displays content filtering statistics for a specific tenant system.

all (Optional) Displays content filtering statistics for all the tenant systems.

`utm-policy` ***`utm-policy-name`*** Starting in Junos OS Release 21.4R1, this option is added to display content filtering statistics in a specified Content Security policy.

Required Privilege Level

view

Sample Output

show security utm content-filtering statistics

```
user@host> show security utm content-filtering statistics
Content-filtering-statistic:         Blocked
     Base on command list:                    0
     Base on mime list:                       0
     Base on extension list:                  0
     ActiveX plugin:                          0
     Java applet:                             0
     EXE files:                               0
     ZIP files:                               0
     HTTP cookie:                             0
     Number of Rule-sets :                    1
     Rule-set add failures :                  0
     Rule-set delete failures :               0
     Number of Rules :                        1
     Rule add failures :                      0
     Rule delete failures :                   0
     Fatal errors :                           0
     Malloc failures :                        0
     Sanity errors :                          0
     Malloc blocks in use :                   10
```

show security utm content-filtering statistics utm-policy <utm-policy-name>

```
user@host> show security utm content-filtering statistics utm-policy pol1
 Number of rules configured w.r.t application

                         Any :         0
                        http :         1
                         ftp :         1
                        smtp :         0
                        imap :         0
                        pop3 :         0

Number of times content blocked w.r.t direction

                    Download :         0
```

```
                     Upload :         0

Number of times content blocked w.r.t application

                       http :         0
                        ftp :         0
                       smtp :         0
                       imap :         0
                       pop3 :         0

Number of rules matched w.r.t action

                  no-action :         0
                      block :         0
               close-client :         0
               close-server :         0
        close-client-server :         0

Number of times content blocked w.r.t file-type
                    unknown :         0
                         7z :         0
                        ace :         0
                applesingle :         0
                        arj :         0
                       bzip :         0
                   diskdupe :         0
                        dos :         0
                      eicar :         0
                        elf :         0
                        emf :         0
                        eml :         0
                      flash :         0
                        gea :         0
                       gzip :         0
                         ha :         0
                        hlp :         0
                     hybris :         0
                       itsf :         0
                       java :         0
                        jmp :         0
                       jpeg :         0
                        lha :         0
                        lnk :         0
```

```
                        map :         0
                       mime :         0
                      msasf :         0
                       mscf :         0
                        msi :         0
                        mso :         0
                   ntdetect :         0
                        ole :         0
                     paquet :         0
                        pdf :         0
                         pe :         0
                       perl :         0
                        pif :         0
                        png :         0
                        rar :         0
                        reg :         0
                        rtf :         0
                     script :         0
                         sh :         0
                        sis :         0
                       spis :         0
                        sys :         0
                        tar :         0
                        vba :         0
                        wmf :         0
                        wsf :         0
                        xml :         0
                         xz :         0
                        zip :         0
```

show security utm content-filtering statistics root-logical-system

```
user@host> show security utm content-filtering statistics root-logical-system
 Content-filtering-statistic:         Blocked
     Base on command list:                    0
     Base on mime list:                       0
     Base on extension list:                  0
     ActiveX plugin:                          0
     Java applet:                             0
     EXE files:                               0
```

```
     ZIP files:                              0
     HTTP cookie:                            0
```

show security utm content-filtering statistics root-logical-system utm-policy <utm-policy-name> (Commands to display CF statistics in a Content Security policy within a root-logical-system)

```
user@host> show security utm content-filtering statistics root-logical-system utm-policy pol1
 Number of rules configured w.r.t application

                        Any :         0
                       http :         1
                        ftp :         1
                       smtp :         0
                       imap :         0
                       pop3 :         0

Number of times content blocked w.r.t direction

                   Download :         0
                     Upload :         0

Number of times content blocked w.r.t application

                       http :         0
                        ftp :         0
                       smtp :         0
                       imap :         0
                       pop3 :         0

Number of rules matched w.r.t action

                  no-action :         0
                      block :         0
               close-client :         0
               close-server :         0
        close-client-server :         0

Number of times content blocked w.r.t file-type

                    unknown :         0
```

```
                      dos :         0
                      pe1 :         0
                      pe2 :         0
                shockwave :         0
                     jpeg :         0
                      elf :         0
                      hlp :         0
                     java :         0
                     ole2 :         0
            ole2-inverted :         0
                       sh :         0
                    perl1 :         0
                    perl2 :         0
                     itsf :         0
                      reg :         0
                      pdf :         0
                      wsf :         0
                     mime :         0
                      zip :         0
                      lnk :         0
                      vba :         0
                      xml :         0
                   ms-asf :         0
             script-file :         0
            script-file2 :         0
                    eicar :         0
                 ntdetect :         0
     decrypt-hybris-plugin :         0
    decrypt-hybris-plugin2 :         0
                      png :         0
                      wmf :         0
                     wmf2 :         0
                      emf :         0
                      map :         0
                      jmp :         0
                      ace :         0
                      lha :         0
                   paquet :         0
             applesingle :         0
                     bzip :         0
                 diskdupe :         0
                      rtf :         0
                      gea :         0
```

```
                      gzip :         0
                       mso :         0
                       sis :         0
                      spis :         0
                       tar :         0
                       arj :         0
                        ha :         0
                       rar :         0
                    zip-pk :         0
                      mscf :         0
                       msi :         0
                       pif :         0
             hybris-plugin :         0
             dos-full-scan :         0
     script-file-full-scan :         0
                       swf :         0
                   swf-cws :         0
                       sys :         0
                       eml :         0
                        7z :         0
                        xz :         0
```

show security utm content-filtering statistics logical-system LSYS1

```
user@host> show security utm content-filtering statistics logical-system LSYS1
 Content-filtering-statistic:         Blocked
     Base on command list:                    0
     Base on mime list:                       0
     Base on extension list:                  0
     ActiveX plugin:                          0
     Java applet:                             0
     EXE files:                               0
     ZIP files:                               0
     HTTP cookie:                             0
```

show security utm content-filtering statistics logical-system <logical-system-name> utm-policy <utm-policy-name>

```
user@host> show security utm content-filtering statistics root-logical-system utm-policy pol1
Number of rules configured w.r.t application

                        Any :         0
                       http :         1
                        ftp :         1
                       smtp :         0
                       imap :         0
                       pop3 :         0

Number of times content blocked w.r.t direction

                   Download :         0
                     Upload :         0

Number of times content blocked w.r.t application

                       http :         0
                        ftp :         0
                       smtp :         0
                       imap :         0
                       pop3 :         0

Number of rules matched w.r.t action

                  no-action :         0
                      block :         0
               close-client :         0
               close-server :         0
        close-client-server :         0

Number of times content blocked w.r.t file-type

                    unknown :         0
                        dos :         0
                        pe1 :         0
                        pe2 :         0
                  shockwave :         0
                       jpeg :         0
```

```
                    elf :         0
                    hlp :         0
                   java :         0
                   ole2 :         0
           ole2-inverted :         0
                     sh :         0
                  perl1 :         0
                  perl2 :         0
                   itsf :         0
                    reg :         0
                    pdf :         0
                    wsf :         0
                   mime :         0
                    zip :         0
                    lnk :         0
                    vba :         0
                    xml :         0
                 ms-asf :         0
            script-file :         0
           script-file2 :         0
                  eicar :         0
               ntdetect :         0
 decrypt-hybris-plugin :         0
decrypt-hybris-plugin2 :         0
                    png :         0
                    wmf :         0
                   wmf2 :         0
                    emf :         0
                    map :         0
                    jmp :         0
                    ace :         0
                    lha :         0
                 paquet :         0
            applesingle :         0
                   bzip :         0
               diskdupe :         0
                    rtf :         0
                    gea :         0
                   gzip :         0
                    mso :         0
                    sis :         0
                   spis :         0
                    tar :         0
```

```
                        arj :         0
                         ha :         0
                        rar :         0
                     zip-pk :         0
                       mscf :         0
                        msi :         0
                        pif :         0
              hybris-plugin :         0
              dos-full-scan :         0
      script-file-full-scan :         0
                        swf :         0
                    swf-cws :         0
                        sys :         0
                        eml :         0
                         7z :         0
                         xz :         0
```

show security utm content-filtering statistics logical-system all

```
user@host> show security utm content-filtering statistics logical-system all
 Content-filtering-statistic:         Blocked
     Base on command list:                    0
     Base on mime list:                       0
     Base on extension list:                  0
     ActiveX plugin:                          0
     Java applet:                             0
     EXE files:                               0
     ZIP files:                               0
     HTTP cookie:                             0
```

show security utm content-filtering statistics tenant TSYS1

```
user@host> show security utm content-filtering statistics tenant TSYS1
 Content-filtering-statistic:         Blocked
     Base on command list:                    0
     Base on mime list:                       0
     Base on extension list:                  0
     ActiveX plugin:                          0
     Java applet:                             0
     EXE files:                               0
```

```
     ZIP files:                              0
     HTTP cookie:                            0
```

show security utm content-filtering statistics tenant all

```
user@host> show security utm content-filtering statistics tenant all
 Content-filtering-statistic:         Blocked
     Base on command list:                   0
     Base on mime list:                      0
     Base on extension list:                 0
     ActiveX plugin:                         0
     Java applet:                            0
     EXE files:                              0
     ZIP files:                              0
     HTTP cookie:                            0
```

show security utm content-filtering statistics all-logical-systems-tenants

```
user@host> show security utm content-filtering statistics all-logical-systems-tenants
 Content-filtering-statistic:         Blocked
     Base on command list:                   0
     Base on mime list:                      0
     Base on extension list:                 0
     ActiveX plugin:                         0
     Java applet:                            0
     EXE files:                              0
     ZIP files:                              0
     HTTP cookie:                            0
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for Content Security in chassis cluster added in Junos OS Release 11.4.

Starting in Junos OS Release 18.2R1, on SRX5000 line, the options `pic` and `fpc` to display physical interface cards (PICs) and Flexible PIC Concentrator (FPC) statistics are not supported.

Support for Content Security in logical system added in Junos OS Release 18.3R1.

Support for Content Security in tenant system added in Junos OS Release 19.2R1.

Starting in Junos OS Release 21.4R1, This command is enhnaced to display the enhnaced content filtering statistics in a Content Security policy.

RELATED DOCUMENTATION

clear security utm content-filtering statistics | 4054

show security utm custom-objects url-feed status

IN THIS SECTION

- Syntax | 25207
- Description | 25207
- Options | 25208
- Required Privilege Level | 25208
- Output Fields | 25208
- Sample Output | 25210
- Release Information | 25211

Syntax

```
show security utm custom-objects url-feed status feed-name
 <options [invalid-urls | ignored-urls]>
```

Description

Display the URL feed status. You can check the invalid and ignored URLs status individually.

Options

feed-name Define the name of the URL feed object

invalid-urls Provides the invalid URLs that is available in the URL feed file.

ignored-urls Provides the ignored URLs that available in the URL feed file.

If the URL pattern limit exceeds then, the system considers the further URLs in feed file as ignored URLs. The system does not consider whether the additional URL is valid or invalid.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Output Fields

Table 1345 on page 25208 lists the output fields for the `show security utm custom-objects url-feed status` *`feed-name`* command.

Table 1345: show security utm custom-objects url-feed status *feed-name* Output Fields

Field Name	Field Description
Auto-update	Status of the URL feed automatic update
Current Status	Current status of the URL feed
Error	Status of the URL feed creation error. See Table 1346 on page 25209 for different values for current status, corresponding errors and root cause.
Feed	Name of the URL feed
Feed URL	URL feed file location
Ignored URLS Count	Number of ignored URLs
Interval	Auto-update interval

Table 1345: show security utm custom-objects url-feed status *feed-name* Output Fields *(Continued)*

Field Name	Field Description
Invalid URLS Count	Number of the invalid URLs
Last Update Time	Last URL pattern updated time
Proxy Settings	Status of the proxy settings
Routing Instance	Status of the routing instance

Table 1346: URL Feed Status and Corresponding Errors and Root Cause

Current Status	Error	Root cause
Downloading	No error	Not available
Download Failed	Internal Server Error	Server side internal error
	Domain resolution failed	The feed URL contains domain. The name-servers did not resolve the domain name.
	File is not available on server.	File is not located on given feed URL.
	Operation timed out	Server did not respond to the request.
	No route to feed server	Route is not present to DUT to reach the feed server.
	Forbidden	File might not have appropriate permissions.

Table 1346: URL Feed Status and Corresponding Errors and Root Cause *(Continued)*

Current Status	Error	Root cause
	File size limit exceeded	Size of file to be downloaded is more than 512 KB.
	Feed file empty	Feed file size is 0.
	Connection refused	Server process is not running or server machine does not have corresponding port opened.
Installing	No Error	Not available
Update failed	Commit failed.	Not available
	No valid entry found in feed file	No valid entry available in the feed file to commit.
	Unable to access configuration database	There could be uncommitted configuration in CLI.
Updated	No Error.	Not available

Sample Output

show security utm custom-objects url-feed status *feed-name*

```
user@host> show security utm custom-objects url-feed status temp
URL-Feed Status:
    Feed                          :temp
    Feed URL                      :https//192.168.12.1/temp.txt
    Routing Instance              :Not Configured
    Proxy Settings                :Not Configured
```

```
    Auto-update                   :enabled
    Interval                      :12
    URL Pattern                   :temp_feed
    Last Update Time              :2023-05-12 03:39
    Current Status                :Updated
    Error                         :No error
    Invalid URLS Count            :1
    Ignored URLS Count            :7
```

show security utm custom-objects url-feed status *feed-name* options invalid-urls

```
user@host> show security utm custom-objects url-feed status feed-name options invalid-urls
URL-Feed InvalidURLs:
    *.*.*.^.*.^.com
```

show security utm custom-objects url-feed status *feed-name* options ignored-urls

```
user@host> show security utm custom-objects url-feed status feed-name options ignored-urls
URL-Feed Ignored URLs:
    www.xyz.com
```

Release Information

Statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

url-feed | 30489

show security utm session

IN THIS SECTION

- Syntax | 25212
- Description | 25212
- Required Privilege Level | 25212
- Output Fields | 25212
- show security utm session | 25213
- Release Information | 25213

Syntax

```
show security utm session
```

Description

Display general Content Security session information including all allocated sessions and active sessions. Also, display information from both nodes in a chassis cluster.

Required Privilege Level

view

Output Fields

show security utm session

When you enter this command, you are provided feedback on the status of your request.

show security utm session

command-name

```
user@host> show security utm session
Maximum sessions:           4000
 Total allocated sessions:         0
 Total freed sessions:             0
 Active sessions:                  0
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for Content Security in chassis cluster added in Junos OS Release 11.4.

Starting in Junos OS Release 18.2R1, on SRX5000 line, the options `pic` and `fpc` to display physical interface cards (PICs) and Flexible PIC Concentrator (FPC) statistics are not supported.

RELATED DOCUMENTATION

clear security utm session | 4057

show security utm status | 25213

show security utm status

IN THIS SECTION

- Syntax | 25214
- Description | 25214
- Required Privilege Level | 25214
- Output Fields | 25214
- show security utm status | 25214
- Release Information | 25214

Syntax

```
show security utm status
```

Description

Display whether the Content Security service is running or not and status of both the nodes (with full chassis cluster support for Content Security).

Required Privilege Level

view

Output Fields

show security utm status

When you enter this command, you are provided feedback on the status of your request.

show security utm status

command-name

```
user@host> show security utm status
UTM service status: Running
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for Content Security in chassis cluster added in Junos OS Release 11.4.

RELATED DOCUMENTATION

clear security utm session | 4057

show security utm session | 25212

show security utm web-filtering category base-filter

IN THIS SECTION

- Syntax | 25215
- Description | 25215
- Required Privilege Level | 25215
- Sample Output | 25215
- Release Information | 25217

Syntax

```
show security utm web-filtering category base-filter
```

Description

Show the list of predefined base filters. A base filter is an object that contains a category-action pair for all categories defined in the category file. A base filter is a structured object, and is defined with the help of a filter name and an array of category-action pairs. Each Enhanced Web Filtering (EWF) category has a default action in a base filter, which is attached to the user profile to act as a backup filter. If the categories are not configured in the user profile, the base filter takes the action. Junos OS Release 17.4R1 also supports online upgradation of base filters.

Required Privilege Level

view

Sample Output

show security utm web-filtering category base-filter

```
user@host> show security utm web-filtering category base-filter
Base-filter: ewf-default-filter
    Enhanced_Adult_Material                                                 block
```

```
    Enhanced_Business_and_Economy                                   permit
    Enhanced_Education                                              permit
    Enhanced_Government                                             permit
    Enhanced_News_and_Media                                         permit
    Enhanced_Religion                                               permit
    Enhanced_Society_and_Lifestyles                                 permit
    Enhanced_Special_Events                                         permit
    Enhanced_Information_Technology                                 permit
    Enhanced_Abortion                                               block
    Enhanced_Advocacy_Groups                                        permit
    Enhanced_Entertainment                                          permit
    Enhanced_Gambling                                               block
    Enhanced_Games                                                  block
    Enhanced_Illegal_or_Questionable                                block
    Enhanced_Job_Search                                             permit
    Enhanced_Shopping                                               permit
    Enhanced_Sports                                                 permit
    Enhanced_Tasteless                                              permit
    Enhanced_Travel                                                 permit
    Enhanced_Vehicles                                               permit
    Enhanced_Violence                                               block
    Enhanced_Weapons                                                block
    Enhanced_Drugs                                                  block
    Enhanced_Militancy_and_Extremist                                block
    Enhanced_Intolerance                                            permit
    Enhanced_Health                                                 permit
    Enhanced_Website_Translation                                    permit
    Enhanced_Advertisements                                         permit
    Enhanced_User_Defined                                           permit
    Enhanced_Nudity                                                 block
    Enhanced_Adult_Content                                          block
    Enhanced_Sex                                                    block
    Enhanced_Financial_Data_and_Services                            permit
    Enhanced_Cultural_Institutions                                  permit
    Enhanced_Media_File_Download                                    permit
    Enhanced_Military                                               permit
    Enhanced_Political_Organizations                                permit
    Enhanced_General_Email                                          permit
    Enhanced_Proxy_Avoidance                                        block
    Enhanced_Search_Engines_and_Portals                             permit
    Enhanced_Web_Hosting                                            permit
    Enhanced_Web_Chat                                               permit
    Enhanced_Hacking                                                block
```

```
    Enhanced_Alternative_Journals                                   permit
    Enhanced_Non_Traditional_Religions                              block
    Enhanced_Traditional_Religions                                  permit
    Enhanced_Restaurants_and_Dining                                 permit
    Enhanced_Gay_or_Lesbian_or_Bisexual_Interest                    permit
    Enhanced_Personals_and_Dating                                   permit
    Enhanced_Alcohol_and_Tobacco                                    permit
    Enhanced_Prescribed_Medications                                 permit
```

Release Information

Command introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

- category (Security Web Filtering) | 3226
- request security utm web-filtering category install | 14893
- show security utm web-filtering category status | 25219

show security utm web-filtering category category

IN THIS SECTION

- Syntax | 25217
- Description | 25218
- Required Privilege Level | 25218
- Sample Output | 25218
- Release Information | 25219

Syntax

```
show security utm web-filtering category category
```

Description

Show the list of categories predefined by Websense. A category list is available on the device. This list consists of categories, each containing a category code, a name, and a parent ID. Categories can also be user-defined. Each category consists of a list of URLs or IP addresses. Categories are not updated dynamically and are tied to the Junos OS release because they have to be compiled into the Junos OS image. Any update in categories needs to be synchronized with the Junos OS release cycle.

> **NOTE**: Starting with Junos OS Release 17.4R1, you can download and dynamically load new Enhanced Web Filtering (EWF) categories. The downloading and dynamic loading of the new EWF categories do not require a software upgrade.

Required Privilege Level

view

Sample Output

show security utm web-filtering category category

```
user@host> show security utm web-filtering category category
Enhanced_Adult_Material
    Enhanced_Business_and_Economy
    Enhanced_Education
    Enhanced_Government
    Enhanced_News_and_Media
    Enhanced_Religion
    Enhanced_Society_and_Lifestyles
    Enhanced_Special_Events
    Enhanced_Information_Technology
    Enhanced_Abortion
    Enhanced_Advocacy_Groups
    Enhanced_Entertainment
    Enhanced_Gambling
    Enhanced_Games
    Enhanced_Illegal_or_Questionable
    Enhanced_Job_Search
    Enhanced_Shopping
    Enhanced_Sports
    Enhanced_Tasteless
```

```
    Enhanced_Travel
    Enhanced_Vehicles
    Enhanced_Violence
```

Release Information

Command introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

- category (Security Web Filtering) | 3226
- request security utm web-filtering category install | 14893
- show security utm web-filtering category base-filter | 25215
- *Predefined Category Upgrading and Base Filter Configuration Overview*

show security utm web-filtering category status

IN THIS SECTION

- Syntax | 25219
- Description | 25220
- Required Privilege Level | 25220
- Sample Output | 25220
- Release Information | 25220

Syntax

```
show security utm web-filtering category status
```

Description

Show the current running version of the downloaded category file or the status of the installed predefined file.

Required Privilege Level

view

Sample Output

show security utm web-filtering category status

```
user@host> show security utm web-filtering category status
 Installed version:  1
             Download version:   0
             Update status:      Done
```

Release Information

Command introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

- category (Security Web Filtering) | 3226
- request security utm web-filtering category install | 14893
- show security utm web-filtering category base-filter | 25215

show security utm web-filtering statistics

IN THIS SECTION

- Syntax | 25221
- Description | 25221

- Options | 25221
- Required Privilege Level | 25222
- Sample Output | 25222
- Release Information | 25227

Syntax

```
show security utm web-filtering statistics
<root-logical-system>
<logical-system (logical-system-name | all)>
<all-logical-systems-tenants>
<tenant (tenant-name | all)>
<fpc <fpc-slot fpc-slot pic-slot pic-slot>>
```

Description

Displays Web filtering statistics for connections including allowlist and blocklist hits and custom category hits. The aggregated statistics from all FPCs and PICs and statistics from both the nodes (with full chassis cluster support for Content Security) are also displayed.

Starting in Junos OS Release 18.3R1, you can view the Web filtering statistics for the primary logical system or for a specific user logical system or for all the user logical systems.

Starting in Junos OS Release 19.2R1, you can view the Web filtering statistics for a specific tenant system or for all the tenant systems.

Options

none Displays Web filtering statistics for the primary logical system.

root-logical-system (Optional) Displays Web filtering statistics for the primary logical system.

logical-system *logical-system-name* (Optional) Displays Web filtering statistics for a specific user logical system.

all (Optional) Displays Web filtering statistics for all the user logical systems.

all-logical-systems-tenants (Optional) Displays Web filtering statistics for all the logical systems and tenant systems.

tenant *tenant-name* (Optional) Displays Web filtering statistics for a specific tenant system.

all (Optional) Displays Web filtering statistics for all the tenant systems.

Required Privilege Level

view

Sample Output

show security utm web-filtering statistics

```
user@host> show security utm web-filtering statistics
UTM web-filtering statistics:
    Total requests:                     0
    white list hit:                     0
    Black list hit:                     0
    Default action hit:                 0
    No license permit:                  0
    Queries to server:                  0
    Server reply permit:                0
    Server reply block:                 0
    Server reply quarantine:            0
    Server reply quarantine block:      0
    Server reply quarantine permit:     0
    Custom category permit:             0
    Custom category block:              0
    Custom category quarantine:         0
    Custom category qurantine block:    0
    Custom category quarantine permit:  0
    Site reputation permit:             0
    Site reputation block:              0
    Site reputation quarantine:         0
    Site reputation quarantine block:   0
    Site reputation quarantine permit:  0
    Site reputation by Category         0
    Site reputation by Global           0
```

```
    Cache hit permit:                   0
    Cache hit block:                    0
    Cache hit quarantine:               0
    Cache hit quarantine block:         0
    Cache hit quarantine permit:        0
Safe-search redirect:               0
+Safe-search rewrite:               0
    SNI pre-check queries to server:    0
    SNI pre-check server responses:     0
    Web-filtering sessions in total:    64000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit           block
          Default                                 0               0
          Timeout                                 0               0
     Connectivity                                 0               0
Too-many-requests                                 0               0
```

show security utm web-filtering statistics fpc

```
user@host> show security utm web-filtering statistics fpc

 fpc-slot 5 pic-slot 0
 UTM web-filtering statistics:
    Total requests:                     0
    white list hit:                     0
    Black list hit:                     0
    Queries to server:                  0
    Server reply permit:                0
    Server reply block:                 0
    Server reply quarantine:            0
    Server reply quarantine block:      0
    Server reply quarantine permit:     0
    Custom category permit:             0
    Custom category block:              0
    Custom category quarantine:         0
    Custom category qurantine block:    0
    Custom category quarantine permit:  0
    Site reputation permit:             0
    Site reputation block:              0
    Site reputation quarantine:         0
    Site reputation quarantine block:   0
```

```
    Site reputation quarantine permit:  0
    Site reputation by Category         0
    Site reputation by Global           0
    Cache hit permit:                   0
    Cache hit block:                    0
    Cache hit quarantine:               0
    Cache hit quarantine block:         0
    Cache hit quarantine permit:        0
    Safe-search redirect:               0
        SNI pre-check queries to server:   1
    SNI pre-check server responses:     1
    Web-filtering sessions in total:    128000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit           block
          Default                                 0               0
          Timeout                                 0               0
     Connectivity                                 0               0
Too-many-requests                                 0               0
```

show security utm web-filtering statistics fpc fpc-slot 5 pic-slot 0

```
user@host> show security utm web-filtering statistics fpc fpc-slot 5 pic-slot 0
 UTM web-filtering statistics:
    Total requests:                     0
    white list hit:                     0
    Black list hit:                     0
    Queries to server:                  0
    Server reply permit:                0
    Server reply block:                 0
    Server reply quarantine:            0
    Server reply quarantine block:      0
    Server reply quarantine permit:     0
    Custom category permit:             0
    Custom category block:              0
    Custom category quarantine:         0
    Custom category qurantine block:    0
    Custom category quarantine permit:  0
    Site reputation permit:             0
    Site reputation block:              0
    Site reputation quarantine:         0
    Site reputation quarantine block:   0
```

```
    Site reputation quarantine permit:  0
    Site reputation by Category         0
    Site reputation by Global           0
    Cache hit permit:                   0
    Cache hit block:                    0
    Cache hit quarantine:               0
    Cache hit quarantine block:         0
    Cache hit quarantine permit:        0
    Safe-search redirect:               0
        SNI pre-check queries to server:   1
    SNI pre-check server responses:     1
    Web-filtering sessions in total:    128000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit           block
          Default                                 0               0
          Timeout                                 0               0
     Connectivity                                 0               0
Too-many-requests                                 0               0
```

show security utm web-filtering statistics root-logical-system

```
user@host> show security utm web-filtering statistics root-logical-system
 UTM web-filtering statistics:
    Web-filtering sessions in total:    2048000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit           block
          Default                                 0               0
          Timeout                                 0               0
     Connectivity                                 0               0
Too-many-requests                                 0               0
```

show security utm web-filtering statistics logical-system LSYS1

```
user@host> show security utm web-filtering statistics logical-system LSYS1
 UTM web-filtering statistics:
    Web-filtering sessions in total:    2048000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit           block
          Default                                 0               0
          Timeout                                 0               0
```

```
     Connectivity                                0               0
Too-many-requests                                0               0
```

show security utm web-filtering statistics logical-system all

```
user@host> show security utm web-filtering statistics logical-system all
 UTM web-filtering statistics:
    Web-filtering sessions in total:    2048000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit           block
          Default                                0               0
          Timeout                                0               0
     Connectivity                                0               0
Too-many-requests                                0               0
```

show security utm web-filtering statistics tenant TSYS1

```
user@host> show security utm web-filtering statistics tenant TSYS1
 UTM web-filtering statistics:
    Web-filtering sessions in total:    1536000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit           block
          Default                                0               0
          Timeout                                0               0
     Connectivity                                0               0
Too-many-requests                                0               0
```

show security utm web-filtering statistics tenant all

```
user@host> show security utm web-filtering statistics tenant all
 UTM web-filtering statistics:
    Web-filtering sessions in total:    1536000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit           block
          Default                                0               0
          Timeout                                0               0
```

```
     Connectivity                                0               0
Too-many-requests                                0               0
```

show security utm web-filtering statistics all-logical-system-tenants

```
user@host> show security utm web-filtering statistics all-logical-system-tenants
 UTM web-filtering statistics:
    Web-filtering sessions in total:    1536000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit           block
          Default                                0               0
          Timeout                                0               0
     Connectivity                                0               0
Too-many-requests                                0               0
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for Content Security in chassis cluster added in Junos OS Release 11.4.

Support for Flexible PIC Concentrator (FPC) and PIC statistics added in Junos OS Release 12.1X46-D10.

Starting in Junos OS Release 18.2R1, on SRX5000 line, the options `pic` and `fpc` are deprecated—rather than immediately removed—to provide backward compatibility.

Support for Content Security in logical system added in Junos OS Release 18.3R1.

Support for Content Security in tenant system added in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security utm web-filtering statistics | 4058

show security utm web-filtering status | 25228

show security utm web-filtering status

IN THIS SECTION

- Syntax | 25228
- Description | 25228
- Required Privilege Level | 25228
- Output Fields | 25228
- Output Fields | 25228
- Sample Output | 25229
- Release Information | 25231

Syntax

```
show security utm web-filtering status <fpc <fpc-slot fpc-slot pic-slot pic-slot>>
```

Description

Display whether the Web filtering server connection is up or not. The aggregated status from all FPCs and PICs and status of both the nodes (with full chassis cluster support for Content Security) are also displayed.

Required Privilege Level

view

Output Fields

show security utm web-filtering status

Output fields are listed in the approximate order in which they appear.

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Table 1347 on page 25229 lists the output fields for the `show security utm web-filtering status` command.

Table 1347: show security utm web-filtering status

Field Name	Field Description
`Connectivity status`	Status of the connection: • `UP` • `DOWN`
`JDPI Parser`	Status of the JDPI parser: • `Enabled`: When the `performance-mode` is not configured at the [`edit security utm default-configuration web-filtering`] hierarchy level. • `Disabled`: When the `performance-mode` is configured at the [`edit security utm default-configuration web-filtering`] hierarchy level.
`Server status`	Status of the Web filtering server.

Sample Output

show security utm web-filtering status (EWF)

With a valid EWF license:

```
user@host> show security utm web-filtering status
 UTM web-filtering status:
    Server status: Juniper Enhanced using Websense server UP
    JDPI Parser: Enabled
```

Without a valid EWF license:

```
UTM web-filtering status:
    Web-filtering type: Juniper Enhanced
    Server status     : Juniper Enhanced server DOWN
    Status reason: EWF license is invalid
```

```
    JDPI Parser       : Enabled
    Dynapp-profile-selection : Enabled
```

show security utm web-filtering status (NGWF)

```
user@host> show security utm web-filtering status
 UTM web-filtering status:
    Web-filtering type: NG-Juniper
    Server status: Juniper Cloud Server UP
    JDPI Parser: Enabled
    Dynapp-profile-selection : Enabled
```

show security utm web-filtering status fpc

```
user@host> show security utm web-filtering status fpc
 UTM web-filtering status fpc:
 fpc-slot 5 pic-slot 0
 Connectivity status: UP
 fpc-slot 0 pic-slot 1
 Connectivity status: UP
```

show security utm web-filtering status fpc fpc-slot 5 pic-slot 0

```
user@host> show security utm web-filtering status fpc fpc-slot 5 pic-slot 0
 UTM web-filtering status:
 Connectivity status: UP
```

show security utm web-filtering chassis cluster status

```
{primary:node0}
user@host>  show security utm web-filtering status
node0:
--------------------------------------------------------------------------
 UTM web-filtering status:
    Server status: Juniper Enhanced using Websense server UP
```

```
node1:
--------------------------------------------------------------------------
 UTM web-filtering status:
    Server status: Juniper Enhanced using Websense server DOWN
```

Starting with 12.3X48-D10 and Junos OS Release 17.3R1, on SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, SRX345, and SRX550M devices, the Content Security process has been moved to the Packet Forwarding Engine (PFE). Starting with 12.1X46-D10 and Junos OS Release 17.3R1, on SRX1400, SRX1500, SRX3400, SRX3600, SRX4100, SRX4200, SRX4600, SRX5400, and SRX5600 devices, the Content Security process has been moved to the PFE. Hence, the status shows down on the secondary node of the cluster.

show security utm web-filtering chassis cluster status (NGWF)

```
{primary:node0}
user@host>  show security utm web-filtering status
node0:
--------------------------------------------------------------------------
 UTM web-filtering status:
    Web-filtering type: NG-Juniper
    Server status: Juniper Cloud server UP
    JDPI Parser: Enabled
    Dynapp-profile-selection : Enabled

node1:
--------------------------------------------------------------------------
 UTM web-filtering status:
    Web-filtering type: NG-Juniper
    Server status: Juniper Cloud server UP
    JDPI Parser: Enabled
    Dynapp-profile-selection : Enabled
```

Release Information

Command introduced in Junos OS Release 9.5.

Support for Content Security in chassis cluster added in Junos OS Release 11.4. Support for Flexible PIC Concentrator (FPC) and PIC status added in Junos OS Release 12.1X46-D10.

Starting in Junos OS Release 18.2R1, on SRX5000 line, the options `pic` and `fpc` to display PIC and FPC statistics are not supported.

RELATED DOCUMENTATION

clear security utm web-filtering statistics | 4058

show security utm web-filtering statistics | 25220

show security zones

IN THIS SECTION

- Syntax | 25232
- Description | 25232
- Options | 25233
- Required Privilege Level | 25233
- Output Fields | 25233
- Sample Output | 25235
- Release Information | 25241

Syntax

```
show security zones
<zone-name>
<all-logical-systems-tenants>
<detail>
<logical-system (logical-system-name | all)>
<root-logical-system>
<tenant (tenant-name | all)>
<terse>
<type (functional | security)>
```

Description

This command displays the information about the security zones. You can define a security zone, which allows you to divide the network into different segments and apply different security options to each segment.

Options

- *none*—Displays information about all the security zones configured.
- *zone-name*—(Optional) Displays information about the specified security zone.
- `all-logical-systems-tenants`—(Optional) Displays the information about the security zone of all logical systems and tenant systems.
- `detail`—(Optional) Displays the detail level information about the security zone.
- `logical-system` *logical-system-name*—(Optional) Displays the information about the security zones of a specified logical system.
- `logical-system all`—(Optional) Displays the information about the security zones of all logical systems.
- `root-logical-system`—(Optional) Displays the information about the security zones of the root logical system.
- `tenant` *tenant-name*—(Optional) Displays the information about the security zones of a specified tenant system.
- `tenant all`—(Optional) Displays the information about the security zones of all tenant systems.
- `terse`—(Optional) Displays the specified level information about the security zone.
- `type functional`—(Optional) Displays the information for functional zones.
- `type security`—(Optional) Displays the information for security zones.

Required Privilege Level

view

Output Fields

Table 1348 on page 25233 lists the output fields for the `show security zones` command. Output fields are listed in the approximate order in which they appear.

Table 1348: show security zones Output Fields

Field Name	Field Description	Level of Output
`Functional zone`	Name of the functional zone.	none

Table 1348: show security zones Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Security zone`	Name of the security zone.	detail none
`Description`	Description of the security zone.	detail none
`Policy configurable`	Whether the policy can be configured or not.	detail none
`Interfaces bound`	Number of interfaces in the zone.	detail none
`Interfaces`	List of the interfaces in the zone.	detail none
`Zone`	Name of the zone.	terse
`Type`	Type of the zone.	terse
`Logical system`	Name of the logical system.	detail
`Tenant`	Name of the tenant system.	detail
`Screen whitelist`	Allowlist for all IP screen options to bypass the IP screening check at zone level.	detail

Sample Output

show security zones

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

show security zones abc

```
user@host> show security zones abc
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
```

```
    Interfaces bound: 1
    Interfaces:
      ge-0/0/1.0
```

show security zones all-logical-systems-tenants

```
user@host> show security zones all-logical-systems-tenants

Logical system: root-logical-system

Security zone: HOST
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: all
  Interfaces:

Security zone: junos-host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Logical system: LSYS1

Security zone: z1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Logical system: TSYS1

Security zone: z3a
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Security zone: z3b
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
```

```
  Interfaces bound: 0
  Interfaces:

Security zone: z3c
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zones abc detail

```
user@host> show security zones abc detail
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```

show security zones logical-system LSYS1

```
user@host> show security zones logical-system LSYS1

Security zone: z1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zones logical-system all

```
user@host> show security zones logical-system all

Logical system: root-logical-system

Security zone: HOST
  Send reset for non-SYN session TCP packets: Off
```

```
  Policy configurable: Yes
  Interfaces bound: all
  Interfaces:

Security zone: junos-host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Logical system: LSYS1

Security zone: z1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zones root-logical-system

```
user@host> show security zones root-logical-system

Security zone: HOST
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: all
  Interfaces:

Security zone: junos-host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zones tenant TSYS1

```
user@host> show security zones tenant TSYS1

Security zone: z3a
  Send reset for non-SYN session TCP packets: Off
```

```
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Security zone: z3b
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Security zone: z3c
 Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zone tenant all

```
user@host> show security zone tenant all

  Tenant: TSYS1



  Security zone: Host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

  Security zone: abc
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:xe-0/0/1.0

  Security zone: def
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:xe-0/0/3.0
```

show security zones terse

```
user@host> show security zones terse
Zone                    Type
my-internal             Security
my-external             Security
dmz                     Security
```

show security zones type security

```
user@host> show security zones type security

Security zone: HOST
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: all
  Interfaces:

Security zone: junos-host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zones TRUST

```
user@host> show security zones TRUST
Security zone: TRUST
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
xe-7/0/3.0
  unidirectional-session-refreshing: Yes
```

show security zones (Screen allowlist)

```
user@host> show security zones
Logical system: root-logical-system

Security zone: trust
  Zone ID: 7
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: jscreen
  Screen whitelist: 2
    a1
    a2
  Interfaces bound: 3
  Interfaces:
ge-0/0/2.0
  Advanced-connection-tracking timeout: 1800
  Unidirectional-session-refreshing: No
```

Release Information

Command introduced in Junos OS Release 8.5.

`tenant` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Security Zones Overview

Supported System Services for Host Inbound Traffic

security-zone | **16516**

show security zones type

IN THIS SECTION

- Syntax | 25242
- Description | 25242
- Options | 25242
- Required Privilege Level | 25243
- Output Fields | 25243
- Sample Output | 25244
- Sample Output | 25244
- Sample Output | 25245
- Sample Output | 25245
- Release Information | 25245

Syntax

```
          show security zones type
          (functional | security)
<detail | terse>
```

Description

This command displays information about security zones of the specified type.

Options

- `functional`—Display functional zones.
- `security`—Display security zones.
- `detail` | `terse`—(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1349 on page 25243 lists the output fields for the `show security zones type` command. Output fields are listed in the approximate order in which they appear.

Table 1349: show security zones type Output Fields

Field Name	Field Description	Level of Output
`Security zone`	Zone name.	All levels
`Description`	Description of the security zone.	none detail
`Policy configurable`	Whether the policy can be configured or not.	none detail
`Interfaces bound`	Number of interfaces in the zone.	none detail
`Interfaces`	List of the interfaces in the zone.	none detail
`Zone`	Name of the zone.	All levels
`Type`	Type of the zone.	All levels

Sample Output

show security zones type functional

```
user@host> show security zones type functional
Functional zone: management
  Description: management zone
  Policy configurable: No
  Interfaces bound: 0
  Interfaces:
```

Sample Output

show security zones type security

```
user@host> show security zones type security
Security zone: trust
  Description: trust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/0.0
Security zone: untrust
  Description: untrust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/1.0
Security zone: junos-host
  Description: junos-host zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

Sample Output

show security zones type security terse

```
user@host> show security zones type security terse
Zone                    Type
trust                   Security
untrust                 Security
junos-host              Security
```

Sample Output

show security zones type security detail

```
user@host> show security zones type security detail
Security zone: trust
  Description: trust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/0.0
Security zone: untrust
  Description: untrust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/1.0
Security zone: junos-host
  Description: junos-host zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

Release Information

Command introduced in Junos OS Release 8.5.

The `Description` output field added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Security Zones Overview

Supported System Services for Host Inbound Traffic

security-zone | **16516**

show service l2tp session routing-instance

IN THIS SECTION

- Syntax | **25246**
- Description | **25246**
- Options | **25247**
- Required Privilege Level | **25247**
- Output Fields | **25247**
- Sample Output | **25247**
- Release Information | **25248**

Syntax

```
show service l2tp session routing-instance name
```

Description

Enables you to view all the LNS L2TP subscriber sessions details associated with a routing instance at once. This feature supports IPv4, IPv6, and dual stack subscribers.

Options

routing-instance *name*
Name of the routing instance on which you need to view the L2TP subscriber sessions details.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show service l2tp session routing-instance *name*` command.

Table 1350: show service l2tp session routing-instance Output Fields

Field Name	Field Description
Interface Name	Name of the interface
Interface unit	Logical interface for this session.
Local ID	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.
Remote ID	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).
State	State of the L2TP session.
Tunnel local ID	Identifier of the local endpoint of the tunnel, as assigned by the L2TP network server (LNS).

Sample Output

show service l2tp session routing-instance *name*

```
user@host> show service l2tp session routing-instance cust-vrf
Tunnel local ID: 28014
```

```
    Local ID  Remote ID   State        Interface unit   Interface Name
    50924     50924       Established  3221225523       si-0/0/0
```

Release Information

Command introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

clear service l2tp session routing-instance | 4062

show service user-identification authentication-source aruba-clearpass user-query counters

IN THIS SECTION

- Syntax | 25248
- Description | 25249
- Options | 25249
- Required Privilege Level | 25249
- Output Fields | 25249
- Sample Output | 25250
- Release Information | 25250

Syntax

```
show service user-identification authentication-source aruba-clearpass user-query counters
```

Description

Display statistics on the counters maintained by the user query function. The output identifies the ClearPass webserver as the destination of the user query requests. It displays the number of requests sent from the SRX Series Firewall to the ClearPass webserver and the number of responses that the SRX Series Firewall received from it. You can use this command to identify that a problem exists–the number of responses received is less than the number of requests sent.—and then analyze and correct it.

If there are no problems with the communication between the ClearPass Policy Manager (CPPM) and the SRX Series Firewall, the number of requests sent is equal to the number of responses received and the number of error responses.

number-of-requests = number-of-responses + error-message-responses

The user query function is part of the SRX Series integrated ClearPass authentication and enforcement feature. The SRX Series Firewall can automatically send requests for individual user authentication and identity information to ClearPass in the event that ClearPass does not post that information to it. For this to occur, you must have configured the user query function.

The SRX Series Firewall exposes to ClearPass a Web API (webapi) that ClearPass uses to send POST request messages to it automatically. These messages contain user authentication and identity information.

The user query function supplements use of the SRX Series Web API function.

Options

authentication-source Specify aruba-clearpass to identifies Aruba ClearPass as the authentication source.

Required Privilege Level

view

Output Fields

- Webserver Address—The IP address of the ClearPass webserver.
- Access token—The token string that the SRX Series Firewall obtains from ClearPass which allows the SRX Series Firewall to query the ClearPass webserver for an individual user's authentication and identity information.
- Requests sent number—A counter that shows the number of individual user authentication information queries that the SRX Series Firewall sent to the ClearPass webserver.

- Total response received number—A counter that shows the number of returns from the ClearPass webserver in response to the individual user authentication information queries that the SRX Series Firewall sent to it. The number of responses should match the number of requests unless an error occurred.
- Error response received number—The number errors that occurred in relation to requests.
- Time of last response—A timestamp showing when the last response from the ClearPass webserver was received.

Sample Output

show service user-identification authentication-source aruba-clearpass user-query counters

```
user@host> show service user-identification authentication-source aruba-clearpass user-query
counters

    Web server Address: 4.0.0.20
    Access token: 433feffae5c3eb3ff8ffdc49f968b03437ca1ce5
    Request sent number: 7
    Total response received number: 7
    Error response received number: 0
    Time of last response: 2000-01-01 11:57:17
```

Release Information

Command introduced in Junos OS Release 12.3X48-D30.

show service user-identification authentication-source aruba-clearpass user-query status

IN THIS SECTION

- Syntax | 25251
- Description | 25251
- Options | 25251

Required Privilege Level | 25251
Release Information | 25251

Syntax

```
show service user-identification authentication-source authentication-source user-query status
```

Description

Checks to determine if the ClearPass webserver is online. The SRX Series Firewall sends user query requests to the ClearPass webserver. The user query function is part of the SRX Series ClearPass Authentication and Enforcement feature.

Options

authentication-source	Identifies the authentication source. For the integrated ClearPass feature, you must specify the predefined term aruba-clearpass to determine if the ClearPass webserver is online.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 12.3X48-D30.

show service user-identification identity-management

IN THIS SECTION

- Syntax | 25252
- Description | 25252
- Options | 25252
- Required Privilege Level | 25253
- Output Fields | 25253
- Sample Output | 25254
- Release Information | 25256

Syntax

```
show service user-identification identity-management (counter | status)
```

Description

Display statistical data about the advanced user query function batch queries and IP queries, or show status on the Juniper Identity Management Service servers.

Options

The following information is displayed for the primary server and the secondary server separately.

counter Display counters for batch and IP queries send to the Juniper Identity Management Service device and responses received from the Juniper Identity Management Service server. This is displayed separately for the primary server and the secondary server, if more than one is configured.

status Verify that the Juniper Identity Management Service server is online and which server is responding to queries from the SRX Series Firewall.

Required Privilege Level

view

Output Fields

Address	IP address of the Juniper Identity Management Service server.
Port	Listening port of the Juniper Identity Management Service server.
Source	Source address used by SRX Series Firewall when connecting to Juniper Identity Management Service server.
Interface	Outgoing interface used by SRX Series Firewall when connecting to Juniper Identity Management Service server.
Routing-instance	Routing-instance used by SRX Series Firewall when connecting to Juniper Identity Management Service server.
Connection method	HTTP or HTTPS connection method when connecting to Juniper Identity Management Service server.
Connection status	Network reachability status from SRX Series Firewall when connecting to Juniper Identity Management Service server.
Last received status message	Status of the last HTTP transaction of SRX Series Firewall when connecting to Juniper Identity Management Service server.
Access token	When SRX Series Firewall gets authenticated with Juniper Identity Management Service server, an access-token is received with specified validity. SRX Series Firewall uses this token for subsequent transactions with the Juniper Identity Management Service server.
Batch queries sent	A number indicating how many batch queries the SRX Series Firewall sent to the Juniper Identity Management Service server.
Batch queries returned	A number indicating how many responses the SRX Series Firewall received from the Juniper Identity Management Service server in response to its batch queries.

Batch query errors received	A number indicating how many errors the SRX Series Firewall received from the Juniper Identity Management Service server in response to its batch queries.
Primary server address	For the status option, the IP address of the primary server.
Secondary server address	For the status option, the IP address of the secondary server.

Sample Output

show service user-identification identity-management counter

```
user@host> show service user-identification identity-management counter
Primary server :
            Address                                     : 12.0.0.15
            Batch queries sent                          : 14
            Batch queries returned                      : 14
            Batch query error received                  : 1
            109028
            Auth entry lookup queries sent              : 0
            Auth entry lookup queries returned          : 0
            Auth entry lookup query errors encountered  : 0
            Auth entry lookup time, average(ms)         : 0
            Auth entry lookup time, max(ms)             : 0
            Certificate revocation requests sent        : 0
            Certificate revocation responses received   : 0
            Certificates revoked                        : 0
Secondary server :
            Address                                     : 12.0.0.2
            Batch queries sent                          : 10899
            Batch queries returned                      : 10899
            Batch query error received                  : 1
            Auth entry lookup queries sent              : 0
            Auth entry lookup queries returned          : 0
            Auth entry lookup query errors encountered  : 0
            Auth entry lookup time, average(ms)         : 0
            Auth entry lookup time, max(ms)             : 0
            Certificate revocation requests sent        : 0
            Certificate revocation responses received   : 0
```

```
            Certificates revoked                        : 0
10
Secondary server :
            Address                                          : Not configured
```

show service user-identification identity-management status

```
user@host> show service user-identification identity-management status
Primary server :
            Address                      : 12.0.0.15 172.16.1.1
            Address                      : 172.16.2.1
            Address                      : 172.16.3.1*

            Address                      : 172.16.4.1
            Port                         : 443
            Source                       : Automatic
            Interface                    : Automatic
            Routing-instance             : Automatic
            Connection method            : HTTPS
            Connection status            : Online
            Last received status message : OK (200)
            Access token                 : fjAYg0nZ8PfDR2IAodygb6b8G0eFFBa3DRzTJ1u
            Token expire time            : 2023-10-02 08:06:15
Secondary server :
            Address                      : 12.0.0.2*
            Port                          172.16.10.1
            Address                      : 443
            Source                        172.16.20.1*
            Address                      : Automatic
            Interface                     172.16.30.1
            Address                      : Automatic
            Routing-instance             : Automatic
            Connection method            : HTTPS
            Connection status            : Online
            Last received status message : OK (200)
            Access token                 : aabe843b-f3f0-40b9-bcfd-734933e8b0d6
            Token expire time            : 2023-10-02 08:06:15 172.16.40.1
```

NOTE: *Asterisk defines active IP.

Release Information

Command introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

Understanding Advanced Query Feature for Obtaining User Identity Information from JIMS

primary connection (Identity Management Advanced Query) | **13131**

show services accounting

IN THIS SECTION

- Syntax | **25256**
- Description | **25257**
- Options | **25257**
- Required Privilege Level | **25257**
- Output Fields | **25257**
- Sample Output | **25258**
- Release Information | **25259**

Syntax

```
show services accounting
aggregation
errors
<inline-jflow | inline-jflow fpc-slot slot number>
flow
<inline-jflow | inline-jflow fpc-slot slot number>
flow-detail
memory
packet-size-distribution
status
```

```
<inline-jflow | inline-jflow fpc-slot slot number>
usage
```

Description

Display sampled accounting service.

Options

- aggregation—Display aggregation information.
- errors —Display error statistics.
  - inline-jflow — Display service accounting inline flow monitoring parameters.
  - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
- flow—Display flow information.
  - inline-jflow — Display service accounting inline flow monitoring parameters.
  - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
- flow-detail—Display flow detail.
- memory—Display memory information.
- packet-size-distribution—Display packet size distribution.
- status–Display service accounting parameters.
  - inline-jflow — Display service accounting inline flow monitoring parameters.
  - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
- usage–Display CPU usage.

Required Privilege Level

view

Output Fields

Lists the output fields for the **show services accounting** command.

Sample Output

show services accounting status inline-jflow

```
user@host> show services accounting status inline-jflow
Status information
    FPC Slot: 5
    Export format: IP-FIX(V9)
    IPv4 Route Record Count: 16, IPv6 Route Record Count: 5
    Route Record Count: 21, AS Record Count: 1
    Route-Records Set: Yes, Config Set: Yes
```

show services accounting errors inline-jflow

```
user@host> show services accounting errors inline-jflow
Error Information
  FPC Slot:  5
  PIC Slot:  0
  Flow Creation Failures: 0
  Route Record Lookup Failures: 0
  AS Lookup Failures: 0
  Export Packet Failures: 0
  Memory Overload: No

  IPv4 Errors:
  IPv4 Flow Creation Failures: 0
  IPv4 Route Record Lookup Failures: 0
  IPv4 AS Lookup Failures: 0
  IPv4 Export Packet Failures: 0

  IPv6 Errors:
  IPv6 Flow Creation Failures: 0
  IPv6 Route Record Lookup Failures: 0
  IPv6 AS Lookup Failures: 0
  IPv6 Export Packet Failures: 0
```

show service accounting flow inline-jflow

```
user@host> show service accounting flow inline-jflow
Flow Information
FPC Slot:  5
PIC Slot:  0
Flow Packets: 2  Flow Bytes: 0
Active Flows: 1  Total Flows: 2
Flows Exported: 0 Flow Packets Exported: 231
Flows Inactive Timed Out: 1 Flows Active Timed Out: 2

IPv4 Flows:
IPv4 Flow Packets: 1  IPv4 Flow Bytes: 0
IPv4 Active Flows: 1  IPv4 Total Flows: 1
IPv4 Flows Exported: 0 IPv4 Flow Packets Exported: 132
IPv4 Flows Inactive Timed Out: 0 IPv4 Flows Active Timed Out: 1

IPv6 Flows:
IPv6 Flow Packets: 1  IPv6 Flow Bytes: 0
IPv6 Active Flows: 0  IPv6 Total Flows: 1
IPv6 Flows Exported: 0 IPv6 Flow Packets Exported: 99
IPv6 Flows Inactive Timed Out: 1 IPv6 Flows Active Timed Out: 1
```

Release Information

Command introduced in Junos OS Release 10.4. The **inline-jflow** and **fpc-slot** options are added in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

- Flow Aggregation to Use Version 9 Flow Templates

show services accounting aggregation

IN THIS SECTION

- Syntax | 25260
- Description | 25260
- Options | 25260
- Additional Information | 25261
- Required Privilege Level | 25261
- Output Fields | 25261
- Sample Output | 25263
- Release Information | 25265

Syntax

```
show services accounting aggregation aggregation-type <aggregation-value>
<detail | extensive | terse>
<limit limit-value>
 <  name service-name>
<order (bytes | packets)>
```

Description

Display information about the aggregated active flows being processed by the accounting service.

Options

aggregation-type <aggregation-value>

Display information for the specified aggregation type and optional value:

- **as** ***<source-as-value | destination-as-value | input-snmp-interface-index-value | output-snmp-interface-index-value>***—Aggregate by autonomous system (AS).

- **destination-prefix <*destination-prefix-value* | *destination-as-value* | *output-snmp-interface-index-value*>**—Aggregate by destination prefix.
- **protocol-port <*protocol-value* | *source-port-value* | *destination-port-value*>**—Aggregate by protocol and port.
- **source-destination-prefix <*source-prefix-value* | *destination-prefix-value* | *destination-as-value* | *source-as-value* | *input-snmp-interface-index-value* | *output-snmp-interface-index-value*>**—Aggregate by source and destination prefix.
- **source-prefix <*source-prefix-value* | *source-as-value* | *input-snmp-interface-index-value*>**—Aggregate by source prefix.

detail | extensive | terse	(Optional) Display the specified level of output.
limit *limit-value*	(Optional) Limit the display output to the specified number of flows. The default is no limit.
name *service-name*	(Optional) Display information about the aggregated flows for a specified service name.
order (bytes | packets)	(Optional) Display the flow with the ordering of the highest number, either by byte count or by packet count.

Additional Information

For information about aggregation configuration options, see the Junos OS Services Interfaces Library for Routing Devices.

Required Privilege Level

view

Output Fields

Table 1351 on page 25262 lists the output fields for the `show services accounting aggregation` command. Output fields are listed in the approximate order in which they appear.

Table 1351: show services accounting aggregation Output Fields

Field Name	Field Description
Service Accounting interface	Name of the service accounting interface.
Local interface index	Index corresponding to the service accounting interface.
Service name	Name of a service that was configured at the [`edit forwarding-options accounting`] hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the [**edit forwarding-options sampling-level**] hierarchy level.
Protocol	Protocol identifier and number.
Source Port	Source port identifier and number.
Destination Port	Destination port identifier and number.
Source-AS	Source autonomous system (AS) number.
Destination-AS	Destination AS number.
Source Prefix	Source prefix.
Destination Prefix	Destination prefix.
Source address	Source address.
Source prefix length	Source prefix length.
Destination address	Destination address.
Destination prefix length	Destination prefix length.

Table 1351: show services accounting aggregation Output Fields *(Continued)*

Field Name	Field Description
Input SNMP interface index	SNMP index of the interface the packet came in on.
Output SNMP interface index	SNMP index of the interface the packet went out on.
Start time	Actual time when the packet in this aggregation was first seen.
End time	Actual time when the packet in this aggregation was last seen.
Flow count	Number of flows in the aggregation.
Packet count	Number of packets in the aggregation.
Byte count	Number of bytes in the aggregation.

Sample Output

show services accounting aggregation protocol-port detail

```
user@host> show service accounting aggregation protocol-port detail
Service Accounting interface: mo-2/0/0, Local interface index: 468
Service name: (default sampling)
  Protocol: 6, Source port: 20, Destination port: 20
  Start time: 442349, End time: 6425714
  Flow count: 194, Packet count: 4294964388, Byte count: 4294781184

  Protocol: 0, Source port: 0, Destination port: 0
  Start time: 442349, End time: 6425749
  Flow count: 204, Packet count: 4294964324, Byte count: 4294777088

  Protocol: 17, Source port: 123, Destination port: 123
```

```
  Start time: 442364, End time: 6425784
  Flow count: 186, Packet count: 4294964152, Byte count: 4294766080
```

show services accounting aggregation source-destination-prefix

```
user@host> show service accounting aggregation source-destination-prefix
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Source         Destination       Input        Output      Flow   Packet  Byte
prefix         prefix            interface    interface   count  count   count
192.0.2.0/20   198.51.100.0/24   ge-5/0/1.0   ge-5/0/0.0  256    491761  31472704
192.0.2.0/20   203.0.113.36/32   ge-5/0/1.0   ge-5/0/0.0  1      1926    123264
192.0.2.0/20   203.0.113.59/32   ge-5/0/1.0   ge-5/0/0.0  1      1926    123264
192.0.2.0/20   192.168.0.63/32   ge-5/0/1.0   ge-5/0/0.0  1      1925    123200
192.0.2.0/20   192.168.0.32/32   ge-5/0/1.0   ge-5/0/0.0  1      1925
```

show services accounting aggregation source-destination- prefix order packet detail

```
user@host> show service accounting aggregation source-destination-prefix order packet detail
name t2 input-snmp-interface-index 538
Service Accounting interface: mo-2/0/0, Local interface index: 468
Service name: t2
Source         Destination  Input SNMP  Output  SNMP   Flow   Packet Byte
Prefix            Prefix       Index      Index Count  Count    Count
10.1.1.2/20   192.168.167.1/0  538         432    1     60    46483
10.1.1.2/20   192.168.168.1/0  538         432    1     60     5191
10.1.1.2/20   192.168.154.1/0  538         432    2     60    45504
10.1.1.2/20   192.168.76.1/0   538         432    1     60    42177
10.1.1.2/20   192.168.149.1/0  538         432    1     60    49184
10.1.1.2/20   192.168.113.1/0  538         432    2     60    48757
```

show services accounting aggregation source-destination- prefix extensive limit

```
user@host> show service accounting aggregation source-destination-prefix name t2 extensive limit
3
Service Accounting interface: mo-2/0/0, Local interface index: 542
Service name: t2
```

```
  Source address: 10.1.1.2, Source prefix length: 20
  Destination address: 192.168.200.176.1, Destination prefix length: 0
  Input SNMP interface index: 24, Output SNMP interface index: 26
  Source-AS: 69, Destination-AS: 69
  Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003
  Flow count: 0, Packet count: 6, Byte count: 5340

  Source address: 10.1.1.2, Source prefix length: 20
  Destination address: 192.168.160.1, Destination prefix length: 0
  Input SNMP interface index: 24, Output SNMP interface index: 26
  Source-AS: 69, Destination-AS: 69
  Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003
  Flow count: 0, Packet count: 6, Byte count: 5490

  Source address: 10.1.1.2, Source prefix length: 20
  Destination address: 192.168.160.1, Destination prefix length: 0
  Input SNMP interface index: 24, Output SNMP interface index: 26
  Source-AS: 69, Destination-AS: 69
  Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003
  Flow count: 0, Packet count: 6, Byte count: 4079
```

show services accounting aggregation source-destination-prefix name terse

```
user@host> show service accounting aggregation source-destination-prefix name T3 terse
Service Accounting interface: rsp0, Local interface index: 171
Service name: T3
Interface state: Accounting
Source        Destination       Input        Output       Flow   Packet  Byte
prefix        prefix            interface    interface    count  count   count
10.1.0.0/20   192.168.3.0/24    ge-5/0/1.0   ge-5/0/0.0   256    639822  40948608
10.1.0.0/20   192.168.2.67/32   ge-5/0/1.0   ge-5/0/0.0   1      2485    159040
10.1.0.0/20   192.168.2.92/32   ge-5/0/1.0   ge-5/0/0.0   1      2485
```

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting aggregation template

IN THIS SECTION

- Syntax | 25266
- Description | 25266
- Options | 25266
- Required Privilege Level | 25266
- Output Fields | 25267
- Sample Output | 25267
- Release Information | 25267

Syntax

```
show services accounting aggregation template
<template-name template-name>
```

Description

Display information for flow aggregation version 9 templates.

Options

none	Display information for all flow aggregation version version 9 templates.
template-name *template-name*	(Optional) Display information for the specified template only.

Required Privilege Level

view

Output Fields

Table 1352 on page 25267 lists the output fields for the `show services accounting aggregation template` command. Output fields are listed in the approximate order in which they appear.

Table 1352: show services accounting aggregation template Output Fields

Field Name	Field Description
MPLS Label 1	Position of first MPLS label.
MPLS Label 2	Position of second MPLS label.
MPLS Label 3	Position of third MPLS label.
MPLS Top Level Address	Outer top label FEC IP address.
Packet Count	Number of packets sent.

Sample Output

show services accounting aggregation template template-name

```
user@host> show services accounting aggregation template template-name mpls
MPLS label 1: 299808, MPLS label 2: 0, MPLS label 3: 0
Source address: 192.0.2.2, Destination address: 10.255.15.22, Top Label Address: 198.51.100.10
Source port: 0, Destination port: 0
Protocol: 61, TOS: 0, TCP flags: 0
Source mask: 24, Destination mask: 32
Input SNMP interface index: 503, Output SNMP interface index: 505
Start time: 40780, End time: 157330
Packet count: 3949198, Byte count: 181663062
```

Release Information

Command introduced in Junos OS Release 8.3.

show services accounting aggregation template (View)

IN THIS SECTION

- Syntax | 25268
- Description | 25268
- Options | 25268
- Required Privilege Level | 25268
- Release Information | 25268

Syntax

```
show services accounting aggregation template
```

Description

Display aggregation type template.

Options

- detail—Display detailed output.
- extensive—Display extensive output.
- template-name—Display name of the template.
- terse—Display terse output (default).

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Flow Aggregation to Use Version 9 Flow Templates

show services accounting aggregation (View)

IN THIS SECTION

- Syntax | 25269
- Description | 25269
- Options | 25269
- Required Privilege Level | 25270
- Release Information | 25270

Syntax

```
show services accounting aggregation
```

Description

Display aggregation information for the accounting service.

Options

- as—Display aggregation type AS.
- destination-prefix—Display aggregation type destination-prefix.
- protocol-port—Display aggregation type protocol-port.
- source-destination-prefix—Display aggregation type source-destination-prefix.
- source-prefix—Display aggregation type source-prefix.
- template—Display aggregation type template.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Flow Aggregation to Use Version 9 Flow Templates

show services accounting errors

IN THIS SECTION

- Syntax | 25270
- Description | 25271
- Options | 25271
- Required Privilege Level | 25271
- Output Fields | 25271
- Sample Output | 25273
- Sample Output | 25273
- Release Information | 25277

Syntax

```
show services accounting  errors
<inline-jflow | name (* | all | service-name)>
```

Description

Display active flow error statistics.

Options

none	Display error statistics for all services accounting instances.
inline-jflow fpc-slot *slot-number*	(Optional) Display error statistics for inline jflow.
name (* | all | *service-name*)	(Optional) Display active flow error statistics. Use a wildcard character, specify all services, or provide a specific service name.

Required Privilege Level

view

Output Fields

Table 1353 on page 25271 lists the output fields for the `show services accounting errors` command. Output fields are listed in the approximate order in which they appear.

Table 1353: show services accounting errors Output Fields

Field	Field Description
Service Accounting interface	Name of the service accounting interface.
Local interface index	Index counter of the local interface.
FPC slot	Slot number of the FPC for which the flow information is displayed. (Available only when the `inline-jflow fpc-slot` *`slot-number`* option is used.)
Service name	Name of a service that was configured at the [**edit forwarding-options accounting**] hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the [**edit forwarding-options sampling-level**] hierarchy level.
Error Information	

Table 1353: show services accounting errors Output Fields *(Continued)*

Field	Field Description
Packets dropped (no memory)	Number of packets dropped because of memory shortage.
Packets dropped (not IP)	Number of non-IP packets dropped.
Packets dropped (not IPv4)	Number of packets dropped because they failed the IPv4 version check.
Packets dropped (header too small)	Number of packets dropped because the packet length or IP header length was too small.
Memory allocation failures	Number of flow record memory allocation failures. A small number reflects failures to replenish the free list. A large number indicates the monitoring station is almost out of memory space.
Memory free failures	Number of flow record memory free failures.
Memory free list failures	Number of flow records received from the free list that failed. Memory is nearly exhausted, or too many new flows greater than 128 KB are being created per second.
Memory overload	Whether the memory has been overloaded. The response can be **Yes** or **No**.
PPS overload	Whether the PIC is receiving more packets per second than the configured threshold. The response can be **Yes** or **No**.
BPS overload	Whether the PIC is receiving more bits per second than the configured threshold. The response can be **Yes** or **No**.
Flow Creation Failures	Number of times flow creation failed.

Table 1353: show services accounting errors Output Fields *(Continued)*

Field	Field Description
Route Record Lookup Failures	Number of times the route record lookup failed.
AS Lookup Failures	Number of times autonomous system lookup failed.
Export Packet Failures	Number of times packet export failed.

Sample Output

show services accounting errors (Monitoring PIC interface)

```
user@host> show services accounting errors
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: (default sampling)
  Error information
    Packets dropped (no memory): 0, Packets dropped (not IP): 0
    Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
    Memory allocation failures: 0, Memory free failures: 0
    Memory free list failures: 0
    Memory overload: No, PPS overload: No, BPS overload: No
```

Sample Output

show services accounting errors (Service PIC interface)

```
user@host> show services accounting errors
Service Accounting interface: sp-0/1/0
Service name: (default sampling)
  Error information
    Service sets dropped: 0, Active timeout failures: 0
```

```
    Export packet failures: 0, Flow creation failures: 0
    Memory overload: No
```

```
Service Accounting interface: sp-1/0/0
Service name: (default sampling)
  Error information
    Service sets dropped: 0, Active timeout failures: 0
    Export packet failures: 0, Flow creation failures: 0
    Memory overload: No
```

show services accounting errors inline-jflow fpc-slot (When Only IPv6 Is Configured)

```
user@host> show services accounting errors inline-jflow fpc-slot 5
  Error information
    FPC Slot: 5
    Flow Creation Failures: 0
    Route Record Lookup Failures: 0, AS Lookup Failures: 0
    Export Packet Failures: 0
    Memory Overload: No, Memory Alloc Fail Count: 0
```

show services accounting errors inline-jflow fpc-slot (When IPv4, IPv6, VPLS, and Bridge Are Configured)

```
user@host> show services accounting errors inline-jflow fpc-slot 5
  Error information
    FPC Slot: 5
    Flow Creation Failures: 0
    Route Record Lookup Failures: 0, AS Lookup Failures: 0
    Export Packet Failures: 0
    Memory Overload: No, Memory Alloc Fail Count: 0

    IPv4:
    IPv4 Flow Creation Failures: 0
    IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
    IPv4 Export Packet Failures: 0

    IPv6:
```

```
    IPv6 Flow Creation Failures: 0
    IPv6 Route Record Lookup Failures: 0, IPv6 AS Lookup Failures: 0
    IPv6 Export Packet Failures: 0



    VPLS:
    VPLS Flow Creation Failures: 0
    VPLS Export Packet Failures: 0

    BRIDGE:
    BRIDGE Flow Creation Failures: 0
    BRIDGE Route Record Lookup Failures: 0, BRIDGE AS Lookup Failures: 0
    BRIDGE Export Packet Failures: 0
```

show services accounting errors inline-jflow (MX80 Router When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting errors inline-jflow
  Error information
    TFEB Slot: 0
    Flow Creation Failures: 0
    Route Record Lookup Failures: 0, AS Lookup Failures: 0
    Export Packet Failures: 0
    Memory Overload: No

         IPv4:
    IPv4 Flow Creation Failures: 0
    IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
    IPv4 Export Packet Failures: 0

    IPv6:
    IPv6 Flow Creation Failures: 0
    IPv6 Route Record Lookup Failures: 0, IPv6 AS Lookup Failures: 0
    IPv6 Export Packet Failures: 0
```

show services accounting errors inline-jflow fpc-slot(PTX1000 Router When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting errors inline-jflow fpc-slot 0
Error information
FPC Slot: 0
Flow Creation Failures: 0
Route Record Lookup Failures: 0, AS Lookup Failures: 0
Export Packet Failures: 0
Memory Overload: No, Memory Alloc Fail Count: 0

IPv4:
IPv4 Flow Creation Failures: 0
IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
IPv4 Export Packet Failures: 0

IPv6:
IPv6 Flow Creation Failures: 0
IPv6 Route Record Lookup Failures: 0, IPv6 AS Lookup Failures: 0
IPv6 Export Packet Failures: 0
```

show services accounting errors inline-jflow (SRX Series Firewalls When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting errors inline-jflow
Error information
    FPC Slot: 0
    Flow Creation Failures: 0
    Route Record Lookup Failures: 0, AS Lookup Failures: 0
    Export Packet Failures: 0
    Memory Overload: No, Memory Alloc Fail Count: 0

    IPv4:
    IPv4 Flow Creation Failures: 0
    IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
    IPv4 Export Packet Failures: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show services accounting flow | 25277

show services accounting flow

IN THIS SECTION

- Syntax | 25277
- Description | 25277
- Options | 25278
- Required Privilege Level | 25278
- Output Fields | 25278
- Sample Output | 25280
- Release Information | 25286

Syntax

```
show services accounting flow
<inline-jflow fpc-slot slot-number | logical-system (all | logical-system) | name (* | all |
service-name)>
```

Description

Display active flow statistics.

Options

none	Display active flow statistics for all service instances.
logical-system (all | *logical-system*)	(Optional) Display active flow statistics for the specified logical system or all logical systems on the device.
inline-jflow (fpc-slot *slot-number*)	(Optional) Display inline flow statistics for the specified FPC. For PTX Series, starting in Junos OS Evolved Release 21.2R1 and Junos OS Release 21.3R1, every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. Therefore, when you specify this option, the command now displays 0 for all of the various `Active Flows` and `Timed Out` fields. The values of the various `Total Flows` fields are now equal to their respective `Flow Packets` field values. The values of the various `Flows Inactive Timed Out` fields are now equal to their respective `Flow Packets` field values.
name (* | all | *service-name*)	(Optional) Display services accounting active flow statistics. Use a wildcard character, specify all services, or provide a specific service name.

Required Privilege Level

view

Output Fields

Table 1354 on page 25278 lists the output fields for the `show services accounting flow` command. Output fields are listed in the approximate order in which they appear.

Table 1354: show services accounting flow Output Fields

Output Field	Output Field Description
Service Accounting interface	Name of the service accounting interface.
Local interface index	Index counter of the local interface.

Table 1354: show services accounting flow Output Fields *(Continued)*

Output Field	Output Field Description
Service name	Name of a service that was configured at the `[edit forwarding-options accounting]` hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the [**edit forwarding-options sampling-level**] hierarchy level.
Flow Information	
FPC Slot	Slot number of the FPC for which the flow information is displayed. (Available only when the `inline-jflow fpc-slot` *`slot-number`* option is used.)
Flow packets	Number of packets received by an operational PIC.
Flow bytes	Number of bytes received by an operational PIC.
Flow packets 10-second rate	Number of packets per second handled by the PIC and displayed as a 10-second average.
Flow bytes 10-second rate	Number of bytes per second handled by the PIC and displayed as a 10-second average.
Active flows	Number of currently active flows tracked by the PIC.
Total flows	Total number of flows received by an operational PIC.
Flows exported	Total number of flows exported by an operational PIC.
Flows packets exported	Total number of cflowd packets exported by an operational PIC.
Flows inactive timed out	Total number of flows that are exported because of inactivity.
Flows active timed out	Total number of long-lived flows that are exported because of an active timeout.

Sample Output

show services accounting flow (Flow Aggregation v5/v8 Configuration)

```
user@host> show services accounting flow
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
  Flow information
    Flow packets: 87168293, Flow bytes: 5578770752
    Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928962
    Active flows: 1000, Total flows: 2000
    Flows exported: 19960, Flows packets exported: 582
    Flows inactive timed out: 1000, Flows active timed out: 29000
```

show services accounting flow (Flow Aggregation v9 Configuration)

```
user@host> show services accounting flow
  Flow information
    Service Accounting interface: sp-7/1/0, Local interface index: 149
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0
    Flows exported: 0, Flows packets exported: 1
    Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow name

```
user@host> show services accounting flow name count2
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: count2
  Flow information
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0
    Flows exported: 0, Flows packets exported: 0
    Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow name all

```
user@host> show services accounting flow name all
Service Accounting interface: rsp0, Local interface index: 171
Service name: T2
Interface state: Accounting
  Flow information
    Flow packets: 37609891, Flow bytes: 2407033024
    Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928953
    Active flows: 1000, Total flows: 1000
    Flows exported: 6705, Flows packets exported: 198
    Flows inactive timed out: 0, Flows active timed out: 13000

Service Accounting interface: rsp0, Local interface index: 171
Service name: T3
Interface state: Accounting
  Flow information
    Flow packets: 37750807, Flow bytes: 2416051712
    Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928940
    Active flows: 1000, Total flows: 1000
    Flows exported: 13437, Flows packets exported: 378
    Flows inactive timed out: 0, Flows active timed out: 13000

Service Accounting interface: rsp0, Local interface index: 171
Service name: T4
 Interface state: Accounting
  Flow information
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0
    Flows exported: 0, Flows packets exported: 0
    Flows inactive timed out: 0, Flows active timed out: 0

Service Accounting interface: rsp0, Local interface index: 171
Service name: count1
Interface state: Accounting
  Flow information
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0
```

```
    Flows exported: 0, Flows packets exported: 0
    Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow (Multiple Sampling Instances)

```
user@host> show services accounting flow
  Flow information
    Service Accounting interface: sp-2/0/0, Local interface index: 215
    Flow packets: 9867, Flow bytes: 631488
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 628
    Active flows: 2, Total flows: 10
    Flows exported: 4028, Flows packets exported: 6150
    Flows inactive timed out: 8, Flows active timed out: 4026

    Service Accounting interface: sp-2/1/0, Local interface index: 223
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0
    Flows exported: 0, Flows packets exported: 1
    Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow inline-jflow fpc-slot (for IPv4 Flow)

```
user@host> show services accounting flow inline-jflow fpc-slot 5
  Flow information
    FPC Slot: 5
    Flow Packets: 0, Flow Bytes: 0
    Active Flows: 0, Total Flows: 0
    Flows Exported: 0, Flow Packets Exported: 0
    Flows Inactive Timed Out: 0, Flows Active Timed Out: 0
```

show services accounting flow inline-jflow fpc-slot (with IPv4, IPv6, VPLS, and Bridge Configuration)

```
user@host> show services accounting flow inline-jflow fpc-slot 5
Flow information
    FPC Slot: 5
    Flow Packets: 0, Flow Bytes: 0
```

```
    Active Flows: 0, Total Flows: 0
    Flows Exported: 0, Flow Packets Exported: 0
    Flows Inactive Timed Out: 0, Flows Active Timed Out: 0

    IPv4 Flows:
    IPv4 Flow Packets: 0, IPv4 Flow Bytes: 0
    IPv4 Active Flows: 0, IPv4 Total Flows: 0
    IPv4 Flows Exported: 0, IPv4 Flow Packets exported: 0
    IPv4 Flows Inactive Timed Out: 0, IPv4 Flows Active Timed Out: 0

    IPv6 Flows:
    IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
    IPv6 Active Flows: 0, IPv6 Total Flows: 0
    IPv6 Flows Exported: 0, IPv6 Flow Packets Exported: 0
    IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0

    VPLS Flows:
    VPLS Flow Packets: 0, VPLS Flow Bytes: 0
    VPLS Active Flows: 0, VPLS Total Flows: 0
    VPLS Flows Exported: 0, VPLS Flow Packets Exported: 0
    VPLS Flows Inactive Timed Out: 0, VPLS Flows Active Timed Out: 0

    BRIDGE Flows:
    BRIDGE Flow Packets: 0, BRIDGE Flow Bytes: 0
    BRIDGE Active Flows: 0, BRIDGE Total Flows: 0
    BRIDGE Flows Exported: 0, BRIDGE Flow Packets Exported: 0
    BRIDGE Flows Inactive Timed Out: 0, BRIDGE Flows Active Timed Out: 0
    BRIDGE Flow Insert Count: 0
```

show services accounting flow inline-jflow (MX80 Router with IPv4 and IPv6 Configuration)

```
user@host> show services accounting flow inline-jflow
Flow information
    TFEB Slot: 0
    Flow Packets: 0, Flow Bytes: 0
    Active Flows: 0, Total Flows: 0
    Flows Exported: 0, Flow Packets Exported: 0
    Flows Inactive Timed Out: 0, Flows Active Timed Out: 0

    IPv4 Flows:
    IPv4 Flow Packets: 0, IPv4 Flow Bytes: 0
```

```
    IPv4 Active Flows: 0, IPv4 Total Flows: 0
    IPv4 Flows Exported: 0, IPv4 Flow Packets exported: 0
    IPv4 Flows Inactive Timed Out: 0, IPv4 Flows Active Timed Out: 0

    IPv6 Flows:
    IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
    IPv6 Active Flows: 0, IPv6 Total Flows: 0
    IPv6 Flows Exported: 0, IPv6 Flow Packets Exported: 0
    IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0
```

show services accounting flow inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting flow inline-jflow fpc-slot 0
Flow information
FPC Slot: 0
Flow Packets: 47427946, Flow Bytes: 5217074060
Active Flows: 0, Total Flows: 2
Flows Exported: 194, Flow Packets Exported: 7045
Flows Inactive Timed Out: 2, Flows Active Timed Out: 192

IPv4 Flows:
IPv4 Flow Packets: 47427946, IPv4 Flow Bytes: 5217074060
IPv4 Active Flows: 0, IPv4 Total Flows: 2
IPv4 Flows Exported: 194, IPv4 Flow Packets exported: 7045
IPv4 Flows Inactive Timed Out: 2, IPv4 Flows Active Timed Out: 192

IPv6 Flows:
IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
IPv6 Active Flows: 0, IPv6 Total Flows: 0
IPv6 Flows Exported: 0, IPv6 Flow Packets Exported: 0
IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0
```

show services accounting flow inline-jflow (Junos OS Evolved 21.2R1 and later, for a PTX10003 router)

```
user@host> show services accounting flow inline-jflow fpc-slot 0
Flow information
  FPC Slot: 0
```

```
  Flow Packets: 0, Flow Bytes: 0
  Active Flows: 0, Total Flows: 0
  Flows Exported: 0, Flow Packets Exported: 0
  Flows Inactive Timed Out: 0, Flows Active Timed Out: 0

  IPv4 Flows:
  IPv4 Flow Packets: 0, IPv4 Flow Bytes: 0
  IPv4 Active Flows: 0, IPv4 Total Flows: 0
  IPv4 Flows Exported: 0
  IPv4 Flows Inactive Timed Out: 0, IPv4 Flows Active Timed Out: 0

  IPv6 Flows:
  IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
  IPv6 Active Flows: 0, IPv6 Total Flows: 0
  IPv6 Flows Exported: 0
  IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0

  MPLS Flows:
  MPLS Flow Packets: 0, MPLS Flow Bytes: 0
  MPLS Active Flows: 0, MPLS Total Flows: 0
  MPLS Flows Exported: 0
  MPLS Flows Inactive Timed Out: 0, MPLS Flows Active Timed Out: 0
```

show services accounting flow inline-jflow (SRX Series When IPv4 is configured)

```
user@host> show services accounting flow inline-jflow
Flow information
    FPC Slot: 0
    Flow Packets: 462680, Flow Bytes: 45433206
    Active Flows: 34, Total Flows: 61093
    Flows Exported: 138936, Flow Packets Exported: 96649
    Flows Inactive Timed Out: 61083, Flows Active Timed Out: 138936
    Total Flow Insert Count: 0

    IPv4 Flows:
    IPv4 Flow Packets: 462680, IPv4 Flow Bytes: 45433206
    IPv4 Active Flows: 34, IPv4 Total Flows: 61093
    IPv4 Flows Exported: 138936, IPv4 Flow Packets exported: 96649
    IPv4 Flows Inactive Timed Out: 61083, IPv4 Flows Active Timed Out: 138936
    IPv4 Flow Insert Count: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

Junos OS Release 10.0 added the capability to display output from multiple sampling instances.

RELATED DOCUMENTATION

show services accounting status | 25300

show services accounting flow-detail

IN THIS SECTION

- Syntax | 25286
- Description | 25287
- Options | 25287
- Additional Information | 25288
- Required Privilege Level | 25288
- Output Fields | 25288
- Sample Output | 25290
- Release Information | 25293

Syntax

```
show services accounting flow-detail
<detail | extensive | terse>
<filters>
<limit limit-value>
<name (* | all | service-name)>
 <order (bytes | packets)>
```

Description

Display information about the flows being processed by the accounting service.

Options

none Display information about all flows.

detail | extensive | terse (Optional) Display the specified level of output.

filters (Optional) Filter the display output of the currently active flow records. The following filters query actively changing data structures and result in different results for multiple invocations:

- **destination-as**—Display flow records filtered by destination autonomous system information.
- **destination-port**—Display flow records filtered by destination port information.
- **destination-prefix**—Display flow records filtered by destination prefix information.
- **input-snmp-interface-index**—Display flow records filtered by SNMP input interface index information.
- **output-snmp-interface-index**—Display flow records filtered by SNMP output interface index information.
- **proto**—Display flow records filtered by protocol type.
- **source-as**—Display flow records filtered by source autonomous system information.
- **source-port**—Display flow records filtered by source port information.
- **source-prefix**—Display flow records filtered by source prefix information.
- **tos**—Display flow records filtered by type of service classification.

limit *limit-value* (Optional) Limit the display output to the specified number of flows. The default is no limit.

name (* | all | *service-name*) (Optional) Display information about the flows being processed. Use a wildcard character, specify all services, or provide a specific services name.

order (bytes | packets)	(Optional) Display the flow with the ordering of the highest number, either by byte count or by packet count.

Additional Information

When no PIC is active, or when no route record has been downloaded from the PIC, this command reports no flows, even though packets are being sampled. This command displays information about two concurrent sessions only. If a third session is attempted, the command pauses with no output until one of the previous sessions is completed.

Required Privilege Level

view

Output Fields

Table 1355 on page 25288 lists the output fields for the `show services accounting flow-detail` command. Output fields are listed in the approximate order in which they appear.

Table 1355: show services accounting flow-detail Output Fields

Field Name	Field Description	Output Level
Service Accounting interface	Name of the service accounting interface.	All levels
Service name	Name of a service that was configured at the `[edit forwarding-options accounting]` hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the `[edit forwarding-options sampling]` hierarchy level.	All levels
Local interface index	Index counter of the local interface.	All levels
TOS	Type-of-service value from the IP header.	**extensive**

Table 1355: show services accounting flow-detail Output Fields *(Continued)*

Field Name	Field Description	Output Level
Input SNMP interface index	SNMP index of the interface on which the packet came in.	**extensive**
Output SNMP interface index	SNMP index of the interface on which the packet went out.	**extensive**
Source-AS	Source AS number.	**extensive**
Destination-AS	Destination AS number.	**extensive**
Protocol	Name of the protocol used for the packet flow from the corresponding source address.	All levels
Input interface	Interface on which the packets were received.	All levels
Output interface	Interface on which the packets were transmitted.	All levels
TCP flags	Number of TCP header flags detected in the flow.	**extensive**
Source address	Address where the flow originated.	All levels
Source port	Name of the source port.	All levels
Source prefix length	Source prefix length.	**extensive**
Destination address	Address where the flow is sent.	All levels
Destination prefix length	Destination prefix length.	**extensive**

Table 1355: show services accounting flow-detail Output Fields *(Continued)*

Field Name	Field Description	Output Level
Destination port	Name of the destination port.	All levels
Start time	Actual time when the packet in this aggregation was first seen.	**detail extensive**
End time	Actual time when the packet in this aggregation was last seen.	**detail extensive**
Packet count	Number of packets in the aggregation.	All levels
Byte count	Number of bytes in the aggregation.	All levels
Time since last active timeout	Amount of time elapsed since the last active timeout, in the format *hh:mm:ss*.	None specified
Packet count for last active timeout	Number of packets in the aggregation since the last active timeout.	None specified
Byte count for last active timeout	Number of bytes in the aggregation since the last active timeout.	None specified

Sample Output

show services accounting flow-detail

In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```
user@host> show services accounting flow-detail
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Protocol   Input          Source                   Source   Output
```

```
           interface      address                    port   interface...
tcp(6)     ge-5/0/1.0     192.0.2.2                     0   ge-5/0/0.0
tcp(6)     ge-5/0/1.0     192.0.2.2                     0   ge-5/0/0.0

Destination          Destination      Packet        Byte  Time since last
address                     port       count       count   active timeout...
198.51.100.149                     0        2660      170240         00:00:58
198.51.100.138                     0        2660      170240         00:00:58

Packet count for       Byte count for
last active timeout    last active timeout
2805                      179520
2805                      179520
```

show services accounting flow-detail limit

In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```
user@host> show services accounting flow-detail limit 1
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Protocol   Input          Source                   Source   Output
           interface      address                  port     interface...
tcp(6)     ge-5/0/1.0     192.0.2.2                      0   ge-5/0/0.0

Destination          Destination      Packet        Byte    Time since last
address              port             count         count   active timeout...
 198.51.100.149                    0        2158        138112         00:00:47

Packet count for            Byte count for
last active timeout      last active timeout
                     2827                 180928
```

show services accounting flow-detail name extensive

```
user@host> show services accounting flow-detail name cf-2 extensive
Service Accounting interface: mo-0/2/0, Local interface index: 145
Service name: cf-2
```

```
  TOS: 0, Protocol: udp(17), TCP flags: 0
  Source address: 10.10.10.1, Source prefix length: 0, Destination address: 203.0.113.20,
Destination prefix length: 0, Source port: 1173, Destination port: 69
  Input SNMP interface index: 65, Output SNMP interface index: 0, Source-AS: 0,
Destination-AS: 0
  Start time: 62425, End time: 635265, Packet count: 165845, Byte count: 9453165
```

show services accounting flow-detail limit order bytes

The output of the following command is displayed over 141 columns, not the standard 80 columns. In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```
user@host> show services accounting flow-detail limit 5 order bytes
Service Accounting interface: mo-2/0/0, Local interface index: 356
Service name: (default sampling)
           Input          Source                   Source   Output
Protocol   interface      address                  port     interface...
icmp(1)    ge-2/3/0.0     192.0.2.2                      0 .local.
icmp(1)    ge-2/3/0.0     192.0.2.2                      0 .local.
icmp(1)    ge-2/3/0.0     192.0.2.2                      0 .local.
icmp(1)    ge-2/3/0.0     192.0.2.2                      0 .local.
icmp(1)    ge-2/3/0.0     192.0.2.2                      0 .local.

Destination          Destination       Packet       Byte    Time since last
address              port              count        count   active timeout...
192.168.128.2                    0          16        12148   Not applicable
192.168.144.2                    0          16        15229   Not applicable
192.168.192.2                   0          16        13296   Not applicable
192.168.16.2                    0          16        13924   Not applicable
192.168.48.2                    0          16        13428   Not applicable

Packet count for          Byte count for
last active timeout       last active timeout
Not applicable            Not applicable
Not applicable            Not applicable
Not applicable            Not applicable
Not applicable            Not applicable
Not applicable            Not applicable
```

show services accounting flow-detail name detail source-port

```
user@host> show services accounting flow-detail name cf-2 detail source-port 1173
Service Accounting interface: mo-0/2/0, Local interface index: 145
Service name: cf-2
  Protocol: udp(17), Source address: 10.10.10.1, Source port: 1173, Destination address:
203.0.113.20, Destination port: 69
  Start time: 62425, End time: 811115, Packet count: 142438, Byte count: 8118966
```

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting flow-detail (View)

IN THIS SECTION

- Syntax | 25293
- Description | 25293
- Options | 25294
- Required Privilege Level | 25294
- Release Information | 25294

Syntax

```
show services accounting flow-detail
```

Description

Display flow detail

Options

- destination-as—Filter term destination AS.
- destination-port—Filter term destination port.
- destination-prefix—Filter term destination prefix.
- detail—Display detailed output.
- extensive–Display extensive output.
- input-snmp-interface-index–Filter term input SNMP interface index.
- limit–Display maximum number of flows to display.
- name-Display name of the service, wildcard, or “all”.
- order–Display order for displaying flows.
- output-snmp-interface-index–Filter term output SNMP interface index.
- proto–Filter term protocol.
- source-as–Filter term source AS.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Flow Aggregation to Use Version 9 Flow Templates

show services accounting memory

IN THIS SECTION

- Syntax | 25295
- Description | 25295
- Options | 25295
- Required Privilege Level | 25295
- Output Fields | 25295
- Sample Output | 25296
- Release Information | 25297

Syntax

```
show services accounting  memory
```

Description

Display memory and flow record statistics.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1356 on page 25296 lists the output fields for the `show services accounting memory` command. Output fields are listed in the approximate order in which they appear.

Table 1356: show services accounting memory Output Fields

Output Field	Output Field Description
Service Accounting interface	Name of the service accounting interface.
Memory Utilization	
Local interface index	Index counter of the local interface.
Allocation count	Number of flow records allocated.
Free count	Number of flow records freed.
Maximum allocated	Maximum number of flow records allocated since the monitoring station booted. This number represents the peak number of flow records allocated at a time.
Allocations per second	Flow records allocated per second during the last statistics interval on the PIC.
Frees per second	Flow records freed per second during the last statistics interval on the PIC.
Total memory used	Total amount of memory currently used (in bytes).
Total memory free	Total amount of memory currently free (in bytes).

Sample Output

show services accounting memory (Monitoring PIC Interface)

```
user@host> show                  services accounting                  memory
Service Accounting interface: mo-2/0/0, Local interface index: 468
  Memory utilization
    Allocation count: 437340, Free count: 433699, Maximum allocated: 6782
    Allocations per second: 3366, Frees per second: 6412
```

```
    Total memory used (in bytes): 133460320,
    Total memory free (in bytes): 133918352
```

show services accounting memory (Service PIC Interface)

```
user@host> show                  services accounting                  memory
Service Accounting interface: sp-0/1/0
  Memory utilization
    Allocation count: 1000, Free count: 0
    Allocations per second: 0, Frees per second: 0
    Total memory used (in bytes): 218158272
    Total memory free (in bytes): 587147696
```

```
Service Accounting interface: sp-1/0/0
  Memory utilization
    Allocation count: 1000, Free count: 0
    Allocations per second: 0, Frees per second: 0
    Total memory used (in bytes): 218157592
    Total memory free (in bytes): 587148376
```

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting packet-size-distribution

IN THIS SECTION

- Syntax | 25298
- Description | 25298
- Options | 25298
- Required Privilege Level | 25298

- Output Fields | 25298
- Sample Output | 25299
- Release Information | 25299

Syntax

```
show services accounting packet-size-distribution
<name (* | all | service-name)>
```

Description

Display a packet size distribution histogram.

Options

none	Display a packet size distribution histogram of all accounting services.
name (* | all | *service-name*)	(Optional) Display a packet size distribution histogram. Use a wildcard character, specify all services, or provide a specific services name.

Required Privilege Level

view

Output Fields

Table 1357 on page 25299 lists the output fields for the `show services accounting packet-size-distribution` command. Output fields are listed in the approximate order in which they appear.

Table 1357: show services accounting packet-size-distribution Output Fields

Field Name	Field Description
Service Accounting interface	Name of the service accounting interface.
Service name	Name of a service that was configured at the [**edit-forwarding-options accounting**] hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the [**edit-forwarding-options sampling-level**] hierarchy level.
Local interface index	Index counter of the local interface.
Range start	Smallest packet length (in bytes) to count.
Range end	Largest packet length (in bytes) to count.
Number of packets	Count of packets detected in the size between **Range start** and **Range end**.
Percentage packets	Percentage of the total number of packets that are in this size range.

Sample Output

show services accounting packet-size-distribution name

```
user@host> show services accounting packet-size-distribution name test3
Service Accounting interface: mo-0/2/0, Local interface index: 163
Service name: test3
Range start     Range end     Number of packets     Percentage packets
           32            64                  2924                    100
```

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting status

IN THIS SECTION

- Syntax | 25300
- Description | 25300
- Options | 25300
- Required Privilege Level | 25301
- Output Fields | 25301
- Sample Output | 25302
- Sample Output | 25303
- Release Information | 25305

Syntax

```
show services accounting  status
<inline-jflow fpc-slot slot-number | name (* | all | service-name)>
```

Description

Display available Physical Interface Cards (PICs) for accounting services.

Options

none	Display available PICs for all accounting services.
inline-jflow fpc-slot *slot-number*	(Optional) Display inline flow accounting status for the specified FPC. For a two-member MX Series Virtual Chassis or EX9200 Virtual Chassis, the primary router or switch uses FPC slot numbers 0 through 11 with no offset; the backup router or switch uses FPC slot numbers 12 through 23, with an offset of 12.
name (* | all | *service-name*)	(Optional) Display available PICs. Use a wildcard character, specify all services, or provide a specific services name.

Required Privilege Level

view

Output Fields

Table 1358 on page 25301 lists the output fields for the `show services accounting status` command. Output fields are listed in the approximate order in which they appear.

Table 1358: show services accounting status Output Fields

Field	Field Description
`Service Accounting interface`	Name of the service accounting interface.
`Service name`	Name of a service that was configured at the [`edit-forwarding-options accounting`] hierarchy level. The default display,(`default sampling`), indicates the service was configured at the [`edit-forwarding-options sampling-level`] hierarchy level.
`FPC Slot`	Slot number of the FPC for which the flow information is displayed.
`Local interface index`	Index counter of the local interface.
`Interface state`	Accounting state of the passive monitoring interface. • `Accounting`—PIC is actively accounting. • `Disabled`—PIC has been disabled from the CLI. • `Not accounting`—PIC is up but not accounting. This can happen while the PIC is coming online, or when the PIC is up but has no logical unit configured under the physical interface. • Unknown
`Group index`	Integer that represents the monitoring group of which the PIC is a member. `Group index` is a mapping from the group name to an index. It is not related to the number of monitoring groups.

Table 1358: show services accounting status Output Fields *(Continued)*

Field	Field Description
`Export interval (in seconds)`	Configured export interval for cflowd records, in seconds.
`Export format`	Configured export format.
`Protocol`	Protocol the PIC is configured to monitor.
`Engine type`	Configured engine type that is inserted in output cflowd packets.
`Engine ID`	Configured engine ID that is inserted in output cflowd packets.
`Route Record Count`	Number of routes recorded.
`AS Record Count`	Number of autonomous systems recorded.
`Route Records Set`	Status of route recording; whether routes are recorded or not.
Configuration Set	Status of monitoring configuration; whether monitoring configuration is set or not.

Sample Output

show services accounting status name (Monitoring PIC Interface)

```
user@host> show services accounting status name count1
Service Accounting interface: mo-2/0/0, Local interface index: 468
Service name: count1
Interface state: Accounting
  Group index: 0
  Export interval (in seconds): 60, Export format: cflowd v8
  Protocol: IPv4, Engine type: 55, Engine ID: 5
```

Sample Output

show services accounting status name (Service PIC Interface)

```
user@host> show services accounting status name
Service Accounting interface: sp-0/1/0
Interface state: Accounting
  Export format: 9, Route record count: 0
  IFL to SNMP index count: 7, AS count: 0
  Configuration set: Yes, Route record set: No, IFL SNMP map set: Yes
```

```
Service Accounting interface: sp-1/0/0
Interface state: Accounting
  Export format: 9, Route record count: 33
  IFL to SNMP index count: 7, AS count: 1
  Configuration set: Yes, Route record set: Yes, IFL SNMP map set: Yes
```

show services accounting status inline-jflow fpc-slot (When IPv4, IPv6 and Bridge Family Are Configured)

```
user@host> show services accounting status inline-jflow fpc-slot 0
FPC Slot: 0
     IPV4 export format: Version-IPFIX, IPV6 export format: Not set
     BRIDGE export format: Version-IPFIX, MPLS export format: Version-IPFIX
     IPv4 Route Record Count: 31, IPv6 Route Record Count: 0, MPLS Route Record Count: 13
     Route Record Count: 44, AS Record Count: 1
     Route-Records Set: Yes, Config Set: Yes
     Service Status: PFE-0: Steady PFE-1: Steady
     Using Extended Flow Memory?: PFE-0: No PFE-1: No
     Flex Flow Sizing ENABLED?: PFE-0: No PFE-1: No
     IPv4 MAX FLOW Count: 1024, IPv6 MAX FLOW Count: 512
     BRIDGE MAX FLOW Count: 1024, MPLS MAX FLOW Count: 1024
     MAX Flow Table size: 15
```

show services accounting status inline-jflow (MX80 Router When Both IPv4 and IPv6)

```
user@host> show services accounting status inline-jflow

Status information
          TFEB Slot: 0
    Export format: IP-FIX
    IPv4 Route Record Count: 6, IPv6 Route Record Count: 8
    Route Record Count: 14, AS Record Count: 1
    Route-Records Set: Yes, Config Set: Yes
```

show services accounting status inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting status inline-jflow fpc-slot 0
Status information
FPC Slot: 0
IPV4 export format: Version-IPFIX, IPV6 export format: Version-IPFIX
MPLS export format: Not set
IPv4 Route Record Count: 23, IPv6 Route Record Count: 3, MPLS Route Record Count: 0
Route Record Count: 26, AS Record Count: 1
Route-Records Set: Yes, Config Set: Yes
```

show services accounting status inline-jflow (SRX Series Firewalls When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting status inline-jflow
Status information
    FPC Slot: 0
    IPV4 export format: Version9, IPV6 export format: Version9
    BRIDGE export format: Not set, MPLS export format: Not set
    IPv4 Route Record Count: 24, IPv6 Route Record Count: 0, MPLS Route Record Count: 0
    Route Record Count: 24, AS Record Count: 1
    Route-Records Set: Yes, Config Set: Yes
    Service Status: PFE-0: Steady
    Using Extended Flow Memory?: PFE-0: No
    Flex Flow Sizing ENABLED?: PFE-0: No
```

```
    IPv4 MAX FLOW Count: 0, IPv6 MAX FLOW Count: 0
    BRIDGE MAX FLOW Count: 0, MPLS MAX FLOW Count: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show services accounting flow | 25277

Inline Flow Monitoring for Virtual Chassis Overview

show services accounting usage

IN THIS SECTION

- Syntax | 25305
- Description | 25306
- Options | 25306
- Additional Information | 25306
- Required Privilege Level | 25306
- Output Fields | 25306
- Sample Output | 25307
- Release Information | 25308

Syntax

```
show services accounting usage
<name service-name>
```

Description

Display the CPU usage of PIC used for active flow monitoring.

Options

none Display CPU usage for all service names.

name *service-name* (Optional) Display CPU usage for the specified service name.

Additional Information

When no route record has been downloaded from the PIC, this command reports no flows, even though packets are being sampled.

Required Privilege Level

view

Output Fields

Table 1359 on page 25306 lists the output fields for the `show services accounting usage` command. Output fields are listed in the approximate order in which they appear.

Table 1359: show services accounting usage Output Fields

Output Field	Output Field Description
Service Accounting interface	Name of the service accounting interface.
Service name	Name of a service that was configured at the [**edit-forwarding-options accounting**] hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the [**edit-forwarding-options sampling-level**] hierarchy level.
Local interface index	Index counter of the local interface.
Uptime	Time that the PIC has been operational (in milliseconds).

Table 1359: show services accounting usage Output Fields *(Continued)*

Output Field	Output Field Description
Interrupt time	Total time that the PIC has spent processing packets since the last PIC reset (in microseconds).
Load (5 second)	CPU load on the PIC, averaged more than 5 seconds. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.
Load (1 minute)	CPU load on the PIC, averaged more than 1 minute. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.

Sample Output

show services accounting usage (Monitoring PIC Interface)

```
user@host> show services accounting usage
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: (default sampling)
  CPU utilization
    Uptime: 600413856 milliseconds, Interrupt time: 2403 microseconds
    Load (5 second): 43%, Load (1 minute): 24%
```

show services accounting usage (Service PIC Interface)

```
user@host> show services accounting usage
Service Accounting interface: sp-0/1/0
Service name: (default sampling)
  CPU utilization
    Uptime: 7853940 milliseconds, Interrupt time: 0 microseconds
    Load (5 second): 2%, Load (1 minute): 0%



Service Accounting interface: sp-0/1/0
Service name: (default sampling)
  CPU utilization
```

```
    Uptime: 331160 milliseconds, Interrupt time: 0 microseconds
    Load (5 second): 2%, Load (1 minute): 0%
```

Release Information

Command introduced before Junos OS Release 7.4.

show services advanced-anti-malware dynamic-filter status

IN THIS SECTION

- Syntax | 25308
- Description | 25308
- Required Privilege Level | 25308
- Output Fields | 25309
- Sample Output | 25310
- Release Information | 25311

Syntax

```
show services advanced-anti-malware dynamic-filter status
```

Description

Displays the connection status between the Juniper Advanced Threat Prevention Cloud service and the SRX Series Firewall.

Required Privilege Level

View

Output Fields

Table 1360 on page 25309 lists the output fields for the `show services advanced-anti-malware dynamic-filter status` command. Output fields are listed in the approximate order in which they appear.

Table 1360: show services advanced-anti-malware dynamic-filter status Output Fields

Field Name	Field Description
Server hostname	Juniper Advanced Threat Prevention Cloud server hostname.
Server port	Juniper Advanced Threat Prevention Cloud server port.
Proxy hostname	Web proxy hostname
Proxy port	Web proxy port number

Table 1360: show services advanced-anti-malware dynamic-filter status Output Fields *(Continued)*

Field Name	Field Description

Control Plane	• Last Successful Connect—Displays when the control plane last connected to the cloud.
	• Connection Status—Displays the current status, such as:
	  • Not connected
	  • Initializing
	  • Connecting
	  • Connected
	  • Disconnected
	  • Connect failed
	  • Client certificate not configured
	  • Request client certificate failed
	  • Request server certificate validation failed
	  • Server certificate validation succeeded
	  • Server certificate validation failed
	  • Server hostname lookup failed
	• Pkts Sent—Number of packets sent to the server.
	• Pkts Received—Number of packets received from the server.

Sample Output

show services advanced-anti-malware dynamic-filter status

```
user@host> show services advanced-anti-malware dynamic-filter status
Dynamic Filter Server Connection Status:
Server Hostname: juniper.abc.cloud.com
Server Port: 443
```

```
Proxy Hostname: None
Proxy Port: None
Control Plane
Connection Status: Connected
Last Successful Connect: 2022-02-12 09:51:50 PST
Pkts Sent: 3
Pkts Received: 42
```

Release Information

Command introduced in Junos OS Release 22.1R1.

RELATED DOCUMENTATION

show services advanced-anti-malware statistics | 25321

show services advanced-anti-malware machine-learning-scan

IN THIS SECTION

- Syntax | 25311
- Description | 25312
- Options | 25312
- Additional Information | 25312
- Required Privilege Level | 25312
- Output Fields | 25312
- Sample Output | 25313
- Release Information | 25313

Syntax

```
show services advanced-anti-malware machine-learning-scan
```

Description

Options

Additional Information

Required Privilege Level

View

Output Fields

Table 1 lists the output fields for the `show services advanced-anti-malware machine-learning-scan` command.

Table 1361: show services advanced-anti-malware machine-learning-scan Output Fields

Field Name	Field Description
Version	
Update	
File scanned	
Virus found	
Virus blocked	
Virus permitted	

Sample Output

show services advanced-anti-malware machine-learning-scan

```
user@host> show services advanced-anti-malware machine-learning-scan
```

```
Machine Learning Scan Engine:
Version:1644355400
Update: Fri 11 Nov 2022 08:46:49 PM PST
                   Total      HTTP       HTTPS      SMTP       SMTPS      IMAP       IMAPS
SMB
  File scanned:    0          0          0          0          0          0          0          0
  Virus found:     0          0          0          0          0          0          0          0
  Virus blocked:   0          0          0          0          0          0          0          0
  Virus permitted: 0          0          0          0          0          0          0          0
```

Release Information

Command introduced in Junos OS Release 24.1R1.

show services advanced-anti-malware policy

IN THIS SECTION

- Syntax | 25314
- Description | 25314
- Options | 25314
- Required Privilege Level | 25314
- Output Fields | 25314
- Sample Output | 25316
- Release Information | 25316

Syntax

```
show services advanced-anti-malware policy policy-name
```

Description

The connection to the Juniper Advanced Threat Prevention Cloud is launched on-demand. It is established only when a condition is met and a file or URL must be sent to the cloud. The cloud inspects the file and returns a verdict number (1 through 10). A verdict number is a score or threat level. The higher the number, the higher the malware threat. The SRX Series Firewall compares this verdict number to the Juniper Advanced Threat Prevention Cloud policy settings and either permits or denies the session. If the session is denied, a reset packet is sent to the client and the packets are dropped from the server.

Juniper Advanced Threat Prevention Cloud policies append to the Junos OS security policies by defining the actions to take when a file is considered malware or when an attempt is made to download a file from a location that's on a custom blocklist or allowlist.

Use this command for debugging purposes to verify the policy on the SRX Series Firewall. For example, if files are being downloaded that shouldn't be, then use this command to verify the Juniper Advanced Threat Prevention Cloud policy settings are correct.

Options

policy *policy-name*	(Optional) Display information about the specified policy. If you do not specify a policy, then basic information about all configured Juniper Advanced Threat Prevention Cloud policies are displayed.

Required Privilege Level

view

Output Fields

Table 1362 on page 25315 lists the output fields for the `show services advanced-anti-malware policy` command. Output fields are listed in the approximate order in which they appear.

Table 1362: show services advanced-anti-malware policy Output Fields

Field Name	Field Description
Policy Name	Name of the Juniper Advanced Threat Prevention Cloud policy.
Inspection-profile	Name of the Juniper Advanced Threat Prevention Cloud profile. The profile determines which file types should be sent to the Juniper Advanced Threat Prevention Cloud service.
Protocols	Juniper Networks supports HTTP, HTTPS, SMTP, SMTPS, IMAP, IMAPS, and SMB protocols.
Verdict-threshold	The threshold determines when a file is considered malware. If the cloud service returns a file verdict higher than this threshold, then that file is considered malware.
Action	The action to take, permit or block, when the cloud services reports a file verdict that is higher than the threshold.
Notification	The notification action, log or no log, when a file verdict is higher than the threshold.
Default-notification	The notification action, log or no log, when a file verdict is lower than the threshold.
Whitelist-notification	The notification action, log or no log, when a client attempts to access a hostname, IP address or URL that matches an entry in the allowlist.
Blocklist-notification	The notification action, log or no log, when a client attempts to access a hostname, IP address or URL that matches an entry in the blocklist.
Fallback Options	The actions to take, permit or block and log or no log, when resources are out of limits or when error conditions occur. For example, when the connection to the cloud is broken.

Sample Output

show services advanced-anti-malware policy

```
user@host> show services advanced-anti-malware policy
Advanced-anti-malware configuration:
Policy Name: SkyATP-default-log-only
  Default-notification  : Log
  Whitelist-notification: No Log
  Blacklist-notification: No Log
  Fallback options:
    Action: permit
    Notification: Log
  Protocol: HTTP
  Verdict-threshold: recommended
    Action: permit
    Notification: Log
    Inspection-profile: default
  Protocol: SMTP
  Verdict-threshold: recommended
    Action: User-Defined-in-Cloud (Unknown)
    Notification: Log
    Inspection-profile: default
  Protocol: IMAP
  Verdict-threshold: recommended
    Action: User-Defined-in-Cloud (Unknown)
    Notification: Log
    Inspection-profile: test
  Protocol: SMB
    Notification: Log
    Inspection-profile: smb-test
```

Release Information

Command introduced in Junos OS Release 15.1X49-D33.

RELATED DOCUMENTATION

- Juniper Advanced Threat Prevention Cloud Policy Overview

advanced-anti-malware policy | 1997

show services advanced-anti-malware profile

IN THIS SECTION

- Syntax | 25317
- Description | 25317
- Options | 25318
- Required Privilege Level | 25318
- Output Fields | 25318
- Sample Output | 25320
- Release Information | 25321

Syntax

```
show services advanced-anti-malware profile
    <profile-name profile-name>
```

Description

Displays information about configured Juniper Advanced Threat Prevention Cloud profiles on the SRX Series Firewall. Juniper Advanced Threat Prevention Cloud profiles let you define which files to send to the cloud for inspection. You can create Juniper Advanced Threat Prevention Cloud profiles only with the cloud graphical interface; you cannot create the profile using CLI commands. You can only view the profile on the SRX Series Firewall.

Use this command during troubleshooting to verify you are sending the correct files to the cloud.

Options

profile-name *profile-name* (optional) Display information about the specified profile. If you do not specify a profile, then display basic information about all configured Juniper Advanced Threat Prevention Cloud profiles on the SRX Series Firewall.

Required Privilege Level

View

Output Fields

Table 1363 on page 25318 lists the output fields for the `show services advanced-anti-malware profile` command. Output fields are listed in the approximate order in which they appear.

Table 1363: show services advanced-anti-malware profile Output Fields

Field Name	Field Description
profile_name	The Juniper Advanced Threat Prevention Cloud profile name.
version	The Juniper Advanced Threat Prevention Cloud profile version number. The number is automatically generated. You can compare the version number on the SRX Series Firewall with the version number in the cloud to verify the correct version is loaded.
disabled_file_types	This section lists file types that are not sent to the cloud for inspection. If this section is blank, then all supported file types are sent.

Table 1363: show services advanced-anti-malware profile Output Fields *(Continued)*

Field Name	Field Description
disabled_categories	This section lists file categories that are not sent to the cloud for inspection. If this section is blank, then all supported file categories are sent. This section lets you list general categories instead of having to specify all file types. For example, you can specify the `archive` category instead of listing `ZIP`, `TAR`, `RAR` and so forth as the file types. The following are supported categories: • archive—Archive files such as TAR, RAR, ZIP and so forth. • media—Audio video formats such as WMV (Windows Media Video) and ASF (Advanced Systems Format). • mobile—Mobile applications for iOS and Android. • active media—Flash and Silverlight applications. • script—Scripts such as JS, Python and so forth. • OS package—OS specific update applications such as RPM, DEB, DMG and so forth. • portable documents—PDF, email and MBOX files. • document—All document types except PDF. • Java—Java applications, archives and libraries. • config—Configuration files such as INI, INF, REG, LNK and so forth. • code—Source code files such as C, Java and so forth. • library—Dynamic and static libraries and kernel modules. • executables—Executable binaries. • emerging threat—A special category that includes known threat source file types.
category_thresholds	The minimum and maximum file sizes to be sent to the cloud. If the file size exceeds the maximum limit, the SRX Series Firewall stops sending the file to the cloud, ignores the rest of the file content and sends a file termination notice.

Sample Output

show services advanced-anti-malware profile

```
user@host> show services advanced-anti-malware profile
Advanced Anti-malware inspection profile:
Profile Name:default_profile
version:   1422221827
  disabled_file_types:
  {
  }
  disabled categories:
  [
  ]
  category_thresholds:[
  {
    Category       archive
      min size:    102603
      max size:    1048576
  }]
Profile Name:profile1
version:   1422315427
  disabled_file_types:
  {
    application/mbox
  }
  disabled categories:
  [
        "document"
        "code"
  ]
  category_thresholds:[
  {
    Category       portable_document
      min size:    0
      max size:    1048576
    Category       document
      min size:    0
      max size:    1048576
    Category       executable
      min size:    0
      max size:    5062000
```

```
  }]
Profile Name:profile2
version:   1422308227
  disabled_file_types:
  {
  }
  disabled categories:
  [
        "portable_document"
  ]
  category_thresholds:[
  {
    Category       library
      min size:    4096
      max size:    1048576
  }]
```

Release Information

Command introduced in Junos OS Release 15.1X49-D33.

RELATED DOCUMENTATION

show services advanced-anti-malware policy | 25313

show services advanced-anti-malware statistics

IN THIS SECTION

- Syntax | 25322
- Description | 25322
- Required Privilege Level | 25322
- Output Fields | 25322
- Sample Output | 25326
- Release Information | 25328

Syntax

```
show services advanced-anti-malware statistics
show services advanced-anti-malware malware-db-statistics
```

Description

Displays Juniper Advanced Threat Prevention Cloud statistics, such as total number of sessions processed, number of sessions blocked because they contained malware or were considered C&C sites. Use this command to get an overview of how much malware is being blocked on your site.

You can reset these statistics, for example when you change the Juniper Advanced Threat Prevention Cloud profile or policy, using the `clear services advanced-anti-malware statistics` command.

Required Privilege Level

View

Output Fields

Table 1364 on page 25322 lists the output fields for the `show services advanced-anti-malware statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1364: show services advanced-anti-malware statistics Output Fields

Field Name	Field Description
Session interested	Number of sessions that match both the firewall policy and the Juniper Advanced Threat Prevention Cloud policy.
Session ignored	Total number of sessions where the traffic is not HTTP or HTTPS.
Session hit blocklist	Total number of sessions where the contacted server is on the Juniper ATP Cloud blocklist.
Session hit allowlist	Total number of sessions where the contacted server is on the Juniper ATP Cloud allowlist.

Table 1364: show services advanced-anti-malware statistics Output Fields *(Continued)*

Field Name	Field Description
Session active	Number of current active sessions.
Session blocked	Number of sessions blocked for any reason.
Session permitted	Number of sessions permitted, such as when allowed through the allowlist. Basically, this is any session that is not blocked.
File submission success	Number of files successfully submitted to the Juniper ATP Cloud for inspection.
File submission failure	Number of files that were not submitted successfully to the Juniper ATP Cloud for inspection. This may occur when an error occurs while files are in the submission queue. Number of files that were blocked due to signature match.
File submission not needed	Files not submitted to the Juniper ATP Cloud for inspection because they were below minimum or above the maximum size, above the sample rate, or the submission was paused.
File verdict meets threshold	Number of files where the returned verdict is greater than or equal to the specified threshold.
File verdict under threshold	Number of files where the returned verdict is less than the specified threshold.
File fallback blocked	Number of files that were blocked due to a fallback condition.
File fallback permitted	Number of files permitted during a fallback condition.
File hit submission limit	Number of files that could not be submitted to the Juniper ATP Cloud because a rate limit fallback occurred.
Email processed	Total number of e-mail attachments sent to the Juniper ATP Cloud.

Table 1364: show services advanced-anti-malware statistics Output Fields *(Continued)*

Field Name	Field Description
Email permitted	Total number of e-mails sent to their recipients.
Email blocked	Total number of e-mails blocked from being sent to their recipients.
Email tag-and-delivered	Number of e-mails delivered with warning headers attached. If you configured this option, headers are added to e-mails that most mail servers recognize and filter into Spam or Junk folders.
Email quarantined	Total number of e-mails quarantined due to their attachments containing malware.
Email fallback blocked	Total number of e-mails blocked from being sent to their recipient due to configured fallback settings.
Email fallback permitted	Total number of e-mails sent to their recipient due to configured fallback settings.
Email hit allowlist	E-mails are checked against administrator-configured blocklists and allowlists using information such as Envelope From (MAIL FROM), Envelope To (RCPT TO), Body Sender, Body Receiver. If an e-mail matches the allowlist, that e-mail is allowed through without any scanning. If an email matches the blocklist, it is considered to be malicious and is handled the same way as an e-mail with a malicious attachment.
Email hit blocklist	E-mails are checked against administrator-configured blocklists and allowlists using information such as Envelope From (MAIL FROM), Envelope To (RCPT TO), Body Sender, Body Receiver. If an e-mail matches the allowlist, that email is allowed through without any scanning. If an e-mail matches the blocklist, it is considered to be malicious and is handled the same way as an e-mail with a malicious attachment.

show services advanced-anti-malware malware-database-statistics Output Fields on page 25325 lists the output fields for the `show services advanced-anti-malware malware-database-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1365: show services advanced-anti-malware malware-database-statistics Output Fields

Field Name	Field Description
Malware DB type	The malware signature database type. Currently only hot-db is supported. Hot DB consists exclusively of signatures that are generated from files submitted to Juniper ATP cloud from SRX Series Firewalls and various sources.
Total signatures	Total number of malware signatures at any given time.
Malware DB version	The signature database version.
Malware DB update time	The date and time when the malware signature database was last updated.
Malware scan statistics	
File scanned	Total number of files scanned by the device.
Malware found	Total number of malwares signature hits.
Malware blocked	Number of malwares blocked based on the action set for the malware signature. The detected malware file will be immediately blocked after the signature match. Full file submission to Juniper ATP cloud is blocked. SRX Series Firewall shares a notification of the malware hit event to Juniper ATP Cloud.
Malware permitted	Number of malwares permitted based on the action set for the malware signature.

Sample Output

show services advanced-anti-malware statistics

```
user@host> show services advanced-anti-malware statistics
Advanced-anti-malware session statistics:
  Session interested:    4
  Session ignored:       0
  Session hit blacklist: 0
  Session hit whitelist: 0
                         Total      HTTP       HTTPS      SMTP       SMTPS      IMAP
IMAPS     SMB
  Session active:        0          0          0          0          0          0
0         0
  Session blocked:       3          1          0          0          0          1
0         1
  Session permitted:     1          0          0          1          0          0
0         0

Advanced-anti-malware file statistics:
                                 Total      HTTP       HTTPS      SMTP       SMTPS
IMAP       IMAPS     SMB
  File submission success:      1          0          0          1          0
0          0         0
  File submission failure:      3          1          0          0          0
1          0         1
  File submission not needed:   0          0          0          0          0
0          0         0
  File verdict meets threshold: 4          1          0          1          0
1          0         1
  File verdict under threshold: 0          0          0          0          0
0          0         0
  File fallback blocked:        0          0          0          0          0
0          0         0
  File fallback permitted:      0          0          0          0          0
0          0         0
  File hit submission limit:    0          0          0          0          0
0          0         0

Advanced-anti-malware file hash statistics:
                                 Total      HTTP       HTTPS
  File hash eligible samples:   0          0          0
```

```
  File hash selected samples:   0          0          0
  File hash submitted samples:  0          0          0
  File hash failed samples:     0          0          0
  File hash known samples:      0          0          0
  File hash unknown samples:    0          0          0
  File hash verdict timeout:    0          0          0

Advanced-anti-malware email statistics:
                            Total      SMTP       SMTPS      IMAP       IMAPS
  Email processed:          2          1          0          1          0
  Email permitted:          1          1          0          0          0
  Email blocked:            1          0          0          1          0
  Email tag-and-delivered:  0          0          0          0          0
  Email quarantined:        0          0          0          0          0
  Email fallback blocked:   0          0          0          0          0
  Email fallback permitted: 0          0          0          0          0
  Email hit whitelist:      0          0          0          0          0
  Email hit blacklist:      0          0          0          0          0
```

show services advanced-anti-malware malware-database-statistics

```
user@host>show services advanced-anti-malware malware-database-statistics
Malware scan statistics:
  Malware DB type: hot_db
  Total signatures: 1168
  Malware DB version: 1632950572
  Malware DB update time: 2021-09-29 14:24:56 PDT
                     Total      HTTP       HTTPS      SMTP       SMTPS      IMAP       IMAPS
SMB
  File scanned:      4          1          0          1          0          1          0
1
  Malware found:     4          1          0          1          0          1          0
1
  Malware blocked:   3          1          0          0          0          1          0
1
  Malware permitted: 1          0          0          1          0          0          0
0
```

Shown below is an example that is applicable only for Firestorm customers.

```
user@host>show services advanced-anti-malware malware-database-statistics
Malware scan statistics:
  Malware DB type: full_db
  Total signatures: 47971
  Malware DB version: 1674475590
  Malware DB update time: 2023-01-23 04:12:15 PST
                    Total      HTTP       HTTPS      SMTP       SMTPS      IMAP       IMAPS
SMB
  File scanned:     2          2          0          0          0          0          0
1
  Malware found:    1          1          0          0          0          0          0
1
  Malware blocked:  0          0          0          0          0          0          0
1
  Malware permitted: 1         1          0          0          0          0          0
0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D33.

RELATED DOCUMENTATION

clear services advanced-anti-malware statistics | 4066

show services advanced-anti-malware status

IN THIS SECTION

- Syntax | 25329
- Description | 25329
- Required Privilege Level | 25329
- Output Fields | 25329

Sample Output | 25331
Release Information | 25331

Syntax

```
show services advanced-anti-malware status
```

Description

Displays the connection status between the Juniper Advanced Threat Prevention Cloud service and the SRX Series Firewall.

Use this command when you are seeing a lot of failover files with the `show services advanced-anti-malware statistics` command. Also check the system log messages for server connection errors. See Viewing Juniper Advanced Threat Prevention Cloud System Log Messages.

Required Privilege Level

View

Output Fields

Table 1366 on page 25329 lists the output fields for the `show services advanced-anti-malware status` command. Output fields are listed in the approximate order in which they appear.

Table 1366: show services advanced-anti-malware status Output Fields

Field Name	Field Description
Server hostname	Juniper Advanced Threat Prevention Cloud server hostname.
Server port	Juniper Advanced Threat Prevention Cloud server port.
Proxy hostname	Web proxy hostname

Table 1366: show services advanced-anti-malware status Output Fields *(Continued)*

Field Name	Field Description
Proxy port	Web proxy port number

Control Plane	• Connection Time—Displays when the control plane last connected to the cloud.
	• Connection Status—Displays the current status, such as:
	  • Not connected
	  • Initializing
	  • Connecting
	  • Connected
	  • Disconnected
	  • Connect failed
	  • Client certificate not configured
	  • Request client certificate failed
	  • Request server certificate validation failed
	  • Server certificate validation succeeded
	  • Server certificate validation failed
	  • Server hostname lookup failed

Service Plane	• Connection Active Number—Number of active connection to the cloud.
	• Connection Failures—Displays number of times the service plane failed to connect to the cloud. Failures can be due to:
	  • Physical connection problems
	  • SSL authentication issues
	  • SRX overload (for example, no memory)

Sample Output

show services advanced-anti-malware status

```
user@host> show services advanced-anti-malware status
Server connection status:
Server connection status:
  Server hostname: juniper.cloud.com
  Server port:     443
    Control Plane:
      Connection Time: 2015-09-17 20:37:20 UTC
      Connection Status: Connected
    Service Plane:
      fpc0
        Connection Active Number: 16
        Connection Failures : 41
```

Release Information

Command introduced in Junos OS Release 15.1X49-D33.

RELATED DOCUMENTATION

show services advanced-anti-malware statistics | 25321

show services agf fiveqi remarking

IN THIS SECTION

- Syntax | 25332
- Description | 25332
- Options | 25332
- Required Privilege Level | 25332
- Output Fields | 25332

Sample Output | 25333
Release Information | 25333

Syntax

```
show services agf fiveqi-remarking
<session-id session-id>
```

Description

Displays the 5QI marking table for the subscribers

Options

none Displays the shadow 5QI Table marking table.

session-id *session-id* ID of the subscriber session.

Required Privilege Level

view

Output Fields

Table 1367 on page 25332 lists the output fields for the `name-of-command` command.

Table 1367: show services agf fiveqi remarking

Field Name	Field Description
5QI	5QI value
QFI	QoS Flow Identifier. The 5GC supports up to 64 QFIs for a single PDU session.

Table 1367: show services agf fiveqi remarking *(Continued)*

Field Name	Field Description
Forwarding Class	Name of the forwarding class.
Upstream DSCP	IPv4 and IPv6 DSCP mapping on the IP packet containing user data that is sent on the N3 interface. The output displays the DSCP value using 6 bits.
Upstream N3-DSCP	IPv4 and IPv6 DSCP mapping on the encapsulated GTP-U packet that is the N3 interface. The output displays the N3-DSCP value using 6 bits.
PLP	Packet loss priority (PLP) for dropping packets. • high—High probability of a packet being dropped at times of congestion. • low—Low probability of a packet being dropped at times of congestion.

Sample Output

show services agf fiveqi-remarking session-id

```
user@host> show services agf fiveqi-remarking session-id 8995
```

```
5QI    QFI    Forwarding       Upstream    Upstream    PLP
              Class            DSCP        N3-DSCP
6      16                                  011110      high
7      17     EXPEDITED        100000
8      18                      100010      100011
9      19     PREMIUM                                  low
```

Release Information

Command introduced in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

fiveqi-map | 7199

upstream-rewrite | 30481

show services agf nas

IN THIS SECTION

- Syntax | 25334
- Description | 25334
- Options | 25334
- Required Privilege Level | 25335
- Output Fields | 25335
- Sample Output | 25336
- Release Information | 25338

Syntax

```
show services agf nas (statistics | user-equipment <ue-id>
<session-id session-id>
<brief | detail | extensive | summary > )
```

Description

Display information about the Non-Access Stratum (NAS) operation between the Access Gateway Function (AGF) and the user equipment (UE).

Options

statistics (Optional) Display the statistics for different NAS messages that are sent and received on the AGF.

user-equipment *ue-id*	(Optional) Display NAS connection information on user equipment.
session-id *session-id*	Display information for a specified session ID.
brief | detail | extensive | summary	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1368 on page 25335 lists the output fields for the `show services agf nas` command.

Table 1368: show-services-agf-nas

Field Name	Field Description
Message counts	Number of messages sent and received by the Access Gateway Function (AGF).
Messages received	Type and number of NAS signaling messages that has received sent by the AGF.
Messages sent	Type and number of NAS signaling messages that has been sent by the AGF.
PDU ID	PDU identifier
Session ID	Session identifier

Table 1368: show-services-agf-nas *(Continued)*

Field Name	Field Description
Session State	Status of the session. The session can have one of the following states: • Inactive • Activating • Active • Deactivating • Deactivated
UE ID	User equipment (UE) identifier
UE State	Status of the UE. The UE can have one of the following states: • Initial • Registering • Registered • Deregistering • Deregistered • Releasing

Sample Output

show services agf nas statistics

```
user@host> show services agf nas user-equipment
```

```
UE id      UE State         PDU id   Session State  Session id
10076162   Registered       1        Active         4
```

show services agf nas statistics

```
user@host> show services agf nas statistics
```

```


Message counts:
    Sent                                13
    Sent discards                       0
    Received                            6
    Receive discards                    0

Messages sent:
    Registration Request                2
    Registration Complete               2
    Deregistration Request              1
    Deregistration Accept               0
    Identity Response                   0
    Authentication Response             0
    Authentication Failure              0
    Security Mode Complete              0
    Security Mode Reject                0
    Service Request                     0
    UL Nas Transport                    4
    5GMM Status                         0
    PDU Session Establishment Request   2
    PDU Session Modification Complete   0
    PDU Session Modification Reject     0
    PDU Session Release Request         1
    PDU Session Release Complete        1
    5GSM Status                         0

Messages received:
    Registration Accept                 2
    Registration Reject                 0
    Deregistration Accept               1
    Deregistration Request              0
    Identity Request                    0
    Authentication Request              0
    Authentication Result               0
    Authentication Reject               0
    Security Mode Command               0
```

```
    Service Accept                        0
    Service Reject                        0
    DL Nas Transport                      3
    5GMM Status                           0
    PDU Session Establishment Accept      2
    PDU Session Establishment Reject      0
    PDU Session Modification Command      0
    PDU Session Release Command           1
    PDU Session Release Reject            0
    5GSM Status                           0
```

Release Information

Command introduced in Junos OS Release 22.3R1.

show services agf ngap

IN THIS SECTION

- Syntax | 25338
- Description | 25339
- Options | 25339
- Required Privilege Level | 25339
- Output Fields | 25339
- Sample Output | 25346
- Release Information | 25348

Syntax

```
show services agf ngap (amf <amf-name>| user-equipment <ue-id>
<brief | detail | summary>)
```

Description

Display information about the Next Generation Application Protocol (NGAP) operations.

Options

`amf` *`amf-name`*	(Optional) Display the Next Generation Application Protocol (NGAP) information about the Access and Mobility Management Function (AMF).
`user-equipment` *`ue-id`*	(Optional) Display NGAP information about user equipment (UE).
`brief | detail| summary`	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1369 on page 25339 lists the Access and Mobility Management Function (AMF) output fields for the `show services agf ngap` command.

Table 1369: show services agf ngap (AMF outputs)

Field Name	Field Description
AMF	AMF name in the configuration.
Backup AMF name	Fully qualified domain name (FQDN) of the backup AMF for the Globally Unique AMF ID (GUAMI). The AMF provides the GUAMI during the NG Setup Procedure. If the AMF does not provide a backup AMF name, then Junos OS displays `None` is displayed.
Default	Identifies whether the AMF is configured as a default AMF. • Yes—This is configured as the default AMF. • No—This is not a default AMF.
IP address	IP address configured for the AMF.

Table 1369: show services agf ngap (AMF outputs) *(Continued)*

Field Name	Field Description
MCC	Mobile Country Code
MNC	Mobile Network Code
Name	Globally unique FQDN provided by the AMF during the NG Setup Procedure. This name must be published in the Domain Name System (DNS).
Node id	Uniquely configured AMF node ID used internally by the AGF to identify the AMF.
Overload State	Current AMF overload state. When the AMF is in the overload state, the AGF does not select the AMF for any new user equipment (UE) setups
Port	If you do not configure the number, Junos OS uses the default port 38412.
Routing Instance	Configured AMF routing instance. The IP address for the AMF must be reachable in the configured routing instance.
SD	Slice Differentiator
Served GUAMI	List of Global Unique AMF Identifier (GUAMIs). The GUAMI consists of: • PLMN ID—Public land mobile network identifier which consists of an MCC and MNC. • Region ID—AMF region identifier. • AMF Set ID—AMF set identifier. The AMF set ID identifies the AMF set within the AMF region. • Pointer—AMF pointer. The AMF pointer identifies one or more AMFs within the AMF set.

Table 1369: show services agf ngap (AMF outputs) *(Continued)*

Field Name	Field Description
SST	Service Slice Type • Embb—Enhanced mobile broadband • Miot— Massive IoT • Urllc—Ultra-reliable, low latency communications • V2x—Vehicle-to-everything
State	The AMF can have one of the following states: • Offline—The AGF is not connected to the AMF. • Offline (Admin) —The AGF is administratively disabled. The AMF configuration was not found. • Connecting—The AGF is in the process of connecting to the AMF. • Initializing—The AGF is in the initial state of NG Setup. • Online—The AGF has completed the NG setup and is ready to service UE contexts.
Tracking areas	Configured AMF tracking areas.
UE count	Current number of UEbeing serviced by the AMF.

Table 1371 on page 25343 lists the TNLA output fields for the `show services agf ngap` command.

Table 1370: show services agf ngap (TNLA outputs)

Field Name	Field Description
TNLA	Transport Layer Network Association Identifier

Table 1370: show services agf ngap (TNLA outputs) *(Continued)*

Field Name	Field Description
Type	Type of TNLA association . • Static—Initial TNLA established. Static TNLA is based on the AGF configuration. • Dynamic—Subsequent TNLA added by the AMF as a result of N2 AMF signal update messages.
Usage Type	Identifies the type of NGAP packets that are permitted on the TNLA. • UE associated signalling • Non-UE associated signalling • Non-UE associated signalling
Weight Factor	The AMF sets theTNLA weight factor. The AGF uses the weight factor to load balance NGAP UE-associated messages across all TNLAs. The Range of values is 0 through 255. When the AGF receives a weight factor of zero, AGF prevents the allocation of TNLA from future UE registrations.
State	Current TNLA status: • Establishing—Indicates the association is in the process of being established. • Established—TNLA has been established and can exchange NGAP packets with the AMF.
Bound UE Count	Current number of UEs bound to the TNLA
Packets Sent	Number of NGAP packet sent on the TNLA.
Packets Received	Number of NGAP packet received on the TNLA.

Table 1370: show services agf ngap (TNLA outputs) *(Continued)*

Field Name	Field Description
Local IP Address	The IP address of the local endpoint.
Port	The port number of the local endpoint.
Peer IP Address	IP address of the peer SCTP endpoint.
Port	Port number of the peer SCTP endpoint.
State	Current status of the peer IP address. • Inactive—IP address cannot be reached. • Active—IP address is active and reachable. • Unconfirmed—The heartbeat connection has not been established between AGF and AMF. Data has not been received from this IP address.

Table 1371 on page 25343 lists the UE output fields for the `show services agf ngap` command.

Table 1371: show services agf ngap Output Fields (UE outputs)

Field Name	Field Description
AMF TNLA	AMF TNLA that the UE is currently bound to. A value of -1 indicates that the UE is currently not bound to any TNLA..
AMF GUAMI	Globally Unique AMF Identifier assigned to service the UE context.
AMF UE ID	UE identifier allocated by the AMF.
AMF GUAMI	Globally Unique AMF Identifier assigned to service the UE context.

Table 1371: show services agf ngap Output Fields (UE outputs) *(Continued)*

Field Name	Field Description
RM State	Current registration management state. Registration management state is defined by the Third-Generation Partnership Project (3GPP). • RM_REGISTERED—The UE is registered with the 5GC. • RM_DEREGISTERED—The UE is not registered with the 5GC.
State	Current internal Next Generation Application Protocol (NGAP) UE state. • Establishing—The UE context is being set up in the AMF. • Established—The UE context setup is complete. • Releasing —The UE has requested the AMF to release the UE logical NG connection. • Released —The UE context has been released.
Type	UE Type. • Fixed Network Residential Gateway (FN-RG) • 5G Residential Gateway (5G-RG)
UE identifier	User equipment identifier. The identifier uniquely identifies the UE within the AGF.

Table 1372 on page 25344 lists the PDU session output fields for the `show services agf ngap` command.

Table 1372: show services agf ngap Output Fields (PDU Session)

Field Name	Field Description
AGF IP address	Local AGF IP address.
Local GTP tunnel	The local GPRS tunneling protocol (GTP) attributes for the Protocol Data Unit (PDU) session allocated by the AGF.
PDU Session id	PDU session identifier allocated by the 5GC.

Table 1372: show services agf ngap Output Fields (PDU Session) ***(Continued)***

Field Name	Field Description
Remote GTP tunnel	The remote GTP tunnel attributes for the PDU session allocated by the user plane function (UPF).
SD	Slice Differentiator
S-NSSAI	Single Network Service Slice Assistance Information. It consists of SST and SD values.
SST	Service Slice Type • Embb—Enhanced mobile broadband • Miot— Massive IoT • Urllc—Ultra-reliable, low latency communications • V2x—Vehicle-to-everything
State	Current state of the PDU session. • Activating—The PDU session resource setup is in progress. • Active—The PDU session has been set up and the session is active on the AGF. • Failed—The PDU session setup has failed on the AGF.
TEID	GTP tunnel endpoint identifier.
UPF IP address	Remote GTP tunnel endpoint IP address.

Sample Output

show services agf ngap amf

```
user@host> show services agf ngap amf
```

```
AMF                 Routing instance      Default   State          Overload   UEs
amf1                default               Yes       Online         Off        1
```

show services agf ngap amf detail

```
user@host> show services agf ngap amf detail
```

```
AMF: amf.spirent.com
  Name:             amf.spirent.com
  Node id:          1
  Default:          Yes
  Ip address:       10.1.1.7
  Port:             38412
  Routing instance: default
  Tracking areas:   0
  State:            Online
  Overload state:   Off
  UE count:         1

  TNLA: 1
    Type:              Static
    Usage type:        UE and Non-UE associated signalling
    Weight factor:     128
    State:             Established
    Bound UE count:    1
    Packets sent:      16
    Packets received:  10
    Local ip address   Port
    10.1.1.1           55465
```

```
    Peer ip address    Port      State
    10.1.1.7*          38412     Active

  Served GUAMI     PLMN id   Region id   Set id   Pointer   Backup AMF name
  0x123456090207   123456    9           8        7
None

  PLMN id: 123456 (MCC: 123, MNC: 456)
    SST: Embb, SD: 0
    SST: Miot, SD: 0
    SST: Urllc, SD: 0
```

show services agf ngap user-equipment detail

```
user@host> show services agf ngap user-equipment detail
```

```
UE identifier: 33
  State:          Established
  RM state:       RM-REGISTERED
  Type:           FN-RG
  AMF:            amf1
  AMF TNLA:       1
  AMF UE id:      3690987520
  AMF GUAMI:      0x123456090207
  PLMN id:        MCC: 123, MNC: 456
  AMF region id:  9
  AMF set id:     8
  AMF pointer:    7

  PDU session id: 1
    State: Active
    S-NSSAI: SST: Embb, SD: 0
    Remote GTP tunnel
      TEID: 0x001e8480
      UPF ip address: 10.250.1.62
      QoS flow identifier: 1
    Local GTP tunnel
      TEID: 0x000000e9
      AGF ip address: 10.1.1.51
```

Release Information

Command introduced in Junos OS Release 22.3R1.

show services alg conversations

IN THIS SECTION

- Syntax | 25348
- Description | 25348
- Options | 25349
- Required Privilege Level | 25350
- Output Fields | 25350
- Sample Output | 25351
- Release Information | 25356

Syntax

```
show services alg conversations
<brief >
<application-protocol protocol>
<extensive>
<interface interface-name>
```

Description

Display ALG information for Junos OS extension-provider packages.

NOTE: In Junos OS releases earlier than 12.3, the extension-provider packages were variously referred to as Junos Services Framework (JSF), MP-SDK, and eJunos.

Options

none Display standard information about all Junos OS extension-provider packages ALG sessions.

brief (Optional) Display the specified level of output.

application-protocol (Optional) Display information about one of the following application protocols:

- **dce-rpc** Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap** Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns** Domain Name System protocol
- **ftp** File Transfer Protocol
- **h323** H323 protocol
- **ike-esp-nat** IKE ALG
- **pptp** Point-to-Point Tunneling Protocol
- **rpc** Remote Procedure Call protocol
- **rpc-portmap** Remote Procedure Call protocol portmap service
- **rtsp** Real-Time Streaming Protocol
- **rsh** Remote Shell
- **sip** Session Initiation Protocol
- **sql** SQLNet
- **talk** Talk Program

extensive Display extensive information

interface *interface-name* (Optional) Display information about a particular interface.

Required Privilege Level

view

Output Fields

Table 1373 on page 25350 lists the output fields for the `show services alg conversations` command. Output fields are listed in the approximate order in which they appear.

Table 1373: show services alg conversations Output Fields

Field Name	Field Description
`Interface`	Name of the interface.
`ALG`	Name of the ALG in use.
`Number of conversations`	Number of ALG conversations open. A conversation is a group of parent and child sessions.
`Group ID`	Numeric identifier for the session.
`Parent session status`	Status of the parent session: • `Active` • `Closed`
`Parent session ID`	Numeric identifier for the parent session.
`Protocol`	Protocol used for the parent session.
`Forward Flow`	The source and destination prefixes for forward flow.
`Reverse Flow`	The source and destination prefixes for reverse flow.

Table 1373: show services alg conversations Output Fields *(Continued)*

Field Name	Field Description
`Child session status`	Status of the child session: • `Active` • `Closed`
`Child session ID`	Numeric identifier for the child session.
`Number of Resources`	Total number of active child sessions associated with the parent session.
`Resource ID`	Numeric identifier for the resources associated with the parent session.
`Protocol`	Protocol used for the child session.

Sample Output

show services alg conversations

```
user@host> show services alg conversations
Interface name: ms-2/1/0
ALG : SQLV2 ALG, State : active
Number of conversations: 1
Parent session status: closed
Child session : 1, protocol: TCP
Forward Flow : {10.50.50.2:37244 -> 10.40.40.10:4334}
Reverse Flow : {10.40.40.10:4334 -> 10.11.11.10:37244}
```

show services alg conversations brief

The output for the `show services alg conversations brief` command is identical to that for the `show services alg conversations` command. For sample output, see "show services alg conversations" on page 25351.

show services alg conversations extensive

```
user@host> show services alg conversations extensive
Interface name: ms-1/0/0
ALG : H323 ALG, State : active
Number of conversations: 1
Group ID : 3499913712, State : active
Parent session state: active
Parent session ID: 33554433, protocol : TCP
Forward Flow : {198.51.100.2:30000 -> 192.0.2.2:1720}
Reverse Flow : {192.0.2.2:1720 -> 203.0.113.1:57730}
Number of resources: 4
Resource ID: 3499927656, State: active
Number of sessions: 1
Child session ID: 33554436, protocol : UDP
Forward Flow : {198.51.100.2:5086 -> 192.0.2.2:5090}
Reverse Flow : {192.0.2.2:5090 -> 203.0.113.3:55916}
Resource ID: 3499927376, State: active
Number of sessions: 1
Child session ID: 67108867, protocol : UDP
Forward Flow : {192.0.2:5091 -> 203.0.113.3:55917}
Reverse Flow : {198.51.100.2:5087 -> 192.0.2.2:5091}
Resource ID: 3499926816, State: active
Number of sessions: 1
Child session ID: 33554438, protocol : UDP
Forward Flow : {198.51.100.2:5089 -> 192.0.2.2:5093}
Reverse Flow : {192.0.2.2:5093 -> 203.0.113.2:63435}
Resource ID: 3499926536, State: active
Number of sessions: 1
Child session ID: 33554437, protocol : UDP
Forward Flow : {198.51.100.2:5088 -> 192.0.2.2:5092}
Reverse Flow : {192.0.2.2:5092 -> 203.0.113.2:63434}
ALG : RAS ALG, State : active
Number of conversations: 1
Group ID : 799037592, State : active
Parent session state: closed
Number of resources: 0
```

show services alg conversations application-protocol

This command has the same output for the rpc, dce-rpc, rpc-portmap and dce-rpc-portmap ALGs.

```
user@router> show services alg conversations application-protocol rpc
Interface name: ms-1/1/0
ALG : SUNRPC ALG, State : active
  Number of conversations: 2
    Parent session status: closed
      Child session : 1, protocol: UDP
        Forward Flow : {192.168.203.198:1019 -> 192.168.203.194:2049}
        Reverse Flow : {192.168.203.194:2049 -> 192.168.203.198:1019}
      Child session : 2, protocol: UDP
        Forward Flow : {192.168.203.198:36595 -> 192.168.203.194:2049}
        Reverse Flow : {192.168.203.194:2049 -> 192.168.203.198:36595}
    Parent session status: closed
      Child session : 1, protocol: UDP
        Forward Flow : {192.168.203.198:954 -> 192.168.203.194:613}
        Reverse Flow : {192.168.203.194:613 -> 192.168.203.198:954}
      Child session : 2, protocol: UDP
        Forward Flow : {192.168.203.198:53836 -> 192.168.203.194:613}
        Reverse Flow : {192.168.203.194:613 -> 192.168.203.198:53836}

user@router> show services alg conversations application-protocol dns
Interface name: ms-1/1/0
ALG : DNS ALG, State : active
  Number of conversations: 1
    Parent session status: closed
      Child session : 1, protocol: UDP
        Forward Flow : {192.168.203.198:1019 -> 192.168.203.194:2049}
        Reverse Flow : {192.168.203.194:2049 -> 192.168.203.198:1019}

user@router> show services alg conversations application-protocol ftp
Interface name: ms-1/1/0
ALG : DNS ALG, State : active
  Number of conversations: 1
    Parent session status: closed
      Child session : 1, protocol: UDP
        Forward Flow : {192.168.203.198:53836 -> 192.168.203.194:613}
        Reverse Flow : {192.168.203.194:613 -> 192.168.203.198:53836}
user@router> show services alg conversations application-protocol ike-esp-nat
Interface name: ms-2/2/0
```

```
ALG : IKE ALG, State : active
  Number of conversations: 1
    Parent session status: closed
    Child session : 1, protocol: ESP
      Forward Flow : {198.51.100.101:2623 -> 203.0.113.1:46838}
      Reverse Flow : {192.0.2.101:46838 -> 198.51.10.101:2623}
    Child session : 2, protocol: ESP
      Forward Flow : {192.0.2.101:2666 -> 198.51.10.101:57882}
      Reverse Flow : {198.51.10.101:57882 -> 203.0.113.1:2666}
user@router> show services alg conversations application-protocol pptp
Interface name: ms-2/0/0
ALG : PPTP ALG, State : active
  Number of conversations: 1
    Parent session status: active
    Parent session : 1, protocol : TCP
      Forward Flow : {192.0.2.10:1511 -> 198.51.100.10:1723}
      Reverse Flow : {198.51.100.10:1723 -> 192.0.2.10:1511}
      Child session : 1, protocol: GRE
        Forward Flow : {192.0.2.10:0 -> 198.51.100.10:49913}
        Reverse Flow : {198.51.100.10:49913 -> 192.0.2.10:65001}
      Child session : 2, protocol: GRE
        Forward Flow : {198.51.100.10:0 -> 192.0.2.10:0}
        Reverse Flow : {192.0.2.10:0 -> 198.51.100.10:65000}

user@router> show services alg conversations application-protocol rtsp
Interface name: ms-0/1/0
ALG : RTSP ALG, State : active
  Number of conversations: 1
    Parent session : 1, protocol : TCP
      Forward Flow : {198.51.100.2:3985 -> 192.0.2.1:554}
      Reverse Flow : {203.0.113.2:554 -> 198.51.100.2:3985}
    Child session : 1, protocol: UDP
      Forward Flow : {203.0.113.2:35859 -> 198.51.100.2:38159}
      Reverse Flow : {198.51.100.2:38159 -> 192.0.2.1:35859}
    Child session : 2, protocol: UDP
      Forward Flow : {203.0.113.2:35859 -> 198.51.100.2:37391}
      Reverse Flow : {198.51.100.2:37391 -> 192.0.2.1:35859}
user@router> show services alg conversations application-protocol rsh
Interface name: ms-0/1/0
ALG : RSH ALG, State : active
  Number of conversations: 1
    Parent session : 1, protocol : TCP
      Forward Flow : {198.51.100.2:3985 -> 192.0.2.1:554}
```

```
      Reverse Flow : {203.0.113.2:554 -> 198.51.100.2:3985}
    Child session : 1, protocol: UDP
      Forward Flow : {203.0.113.2:35859 -> 198.51.100.2:38159}
      Reverse Flow : {198.51.100.2:38159 -> 192.0.2.1:35859}

user@router> show services alg conversations application-protocol sip
Interface name: ms-1/1/0
ALG : SIP ALG, State : active
  Number of conversations: 1
    Parent session status: active
    Parent session : 1, protocol : UDP
      Forward Flow : {192.0.2.2:5060 -> 198.51.100.2:5060}
      Reverse Flow : {198.51.100.2:5060 -> 203.0.113.2:5060}
    Child session : 1, protocol: UDP
        Forward Flow : {192.0.2.2:6000 -> 198.51.100.2:12442}
        Reverse Flow : {198.51.100.2:12442 -> 203.0.113.2:6000}

user@router> show services alg conversations application-protocol sql
Interface name: ms-2/0/0
ALG : SQLV2 ALG, State : active
  Number of conversations: 1
    Parent session : 1, protocol : 0
      Forward Flow : {0.0.0.0:0 -> 0.0.0.0:0}
      Reverse Flow : {0.0.0.0:0 -> 0.0.0.0:0}
    Child session : 1, protocol: TCP
      Forward Flow : {203.0.113.2:19099 -> 198.51.100.10:32773}
      Reverse Flow : {198.51.100.10:32773 -> 192.0.2.1:19099}
user@router> show services alg conversations application-protocol talk
Interface name: ms-0/1/0
ALG : TALK ALG, State : active
  Number of conversations: 1
    Parent session : 1, protocol : TCP
      Forward Flow : {198.51.2:3985 -> 192.0.2.1:554}
      Reverse Flow : {203.0.113.2:554 -> 198.51.2:3985}
    Child session : 1, protocol: UDP
      Forward Flow : {203.0.113.2:35859 -> 198.51.2:38159}
      Reverse Flow : {198.51.2:38159 -> 192.0.2.1:35859}
```

show services alg conversations interface

```
user@router> show services alg conversations interface ms-1/1/0

ALG : FTP ALG, State : active
Number of conversations: 1
Parent session status: active
Parent session : 1, protocol : TCP
Forward Flow : {10.20.20.10:47164 -> 10.30.30.30:21}
```

Release Information

Command introduced in Junos OS Release 10.4.

`h323` option introduced in Junos OS Release 17.1.

`ike-esp-nat` option introduced in Junos OS Release 17.1.

show services alg statistics

IN THIS SECTION

- Syntax | 25357
- Description | 25357
- Options | 25357
- Required Privilege Level | 25358
- Output Fields | 25358
- Sample Output | 25369
- Release Information | 25374

Syntax

```
show services alg statistics
<application-protocol protocol>
<interface interface-name>
```

Description

Display ALG statistics for Junos OS extension-provider packages.

> **NOTE**: In Junos OS releases earlier than 12.3, the extension-provider packages were variously referred to as Junos Services Framework (JSF), MP-SDK, and eJunos.

Options

application-protocol (Optional) Display statistics for one of the following application protocols:

dce-rpc	Distributed Computing Environment-Remote Procedure Call protocols
dce-rpc-portmap	Distributed Computing Environment-Remote Procedure Call protocols portmap service
dns	Domain Name System protocol
ftp	File Transfer Protocol
h323	H323 protocol
ike-esp-nat	IKE ALG
pptp	Point-to-Point Tunneling Protocol
rpc	Remote Procedure Call protocol
rpc-portmap	Remote Procedure Call protocol portmap service
rtsp	Real-Time Streaming Protocol
rsh	Remote Shell
sip	Session Initiation Protocol

sql	SQLNet
talk	Talk Program
tftp	Trivial File Transfer Protocol

interface ***interface-name*** (Optional) Display information about a particular interface.

Required Privilege Level

view

Output Fields

Table 1374 on page 25358 lists the output fields for the `show services alg statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1374: show services alg statistics Output Fields

Field Name	Field Description
`Interface`	Name of the interface.
`ALG statistics`	Name of the ALG for which the statistics are displayed.
`Packets with wrong header`	Number of packets with wrong header.
`Non epm 3.0 packets`	Number of non epm 3.0 packets.
`Packets with type mismatch`	Number of packets with type mismatch.
`Packets with id mismatch`	Number of packets with id mismatch.

Table 1374: show services alg statistics Output Fields *(Continued)*

Field Name	Field Description
`Packets with call mismatch`	Number of packets with call mismatch.
`Packets fragmented`	Number of packets fragmented.
`Packets queued`	Number of packets queued.
`Packets dropped`	Number of packets dropped.
`Packets released`	Number of packets released.
`Invalid packets received`	Number of invalid packets received.
`Reply packets received`	Number of reply packets received.
`Oversized packets received`	Number of oversized packets received.
`ALG parser errors`	Number of parsing failed errors.
`Packets translated`	Number of packets translated.
`H323 total calls`	Total number of audio/video calls that have been established.
`H323 active calls`	Current number of active H.323 calls.

Table 1374: show services alg statistics Output Fields *(Continued)*

Field Name	Field Description
`H323 gate install failed`	Number of gate installation failures for child sessions.
`H323 pinhole opened too late`	Number of H323 parent sessions that released the resources before pinhole creation.
`H323 pinhole hit dropped`	Number of H323 gate hits that have been dropped.
`H323 gate timeout failed`	Number of gate timeout failures due to an error.
`H323 packets dropped`	Number of packets dropped.
`H323 get virtual ctx failed`	Number of failures to get the session virtualization ctx information.
`H323 obj alloc failed`	Number of memory allocation failures for H323 session cookie.
`H323 group alloc failed`	Number of H323 session resource/group memory allocation failures.
`H323 ce alloc failed`	Number of H323 session call entity object memory allocation failures.
`H323 Q931 decode error`	Number of errors in decoding Q931 packets.
`H323 H245 decode error`	Number of errors in decoding H245 packets.

Table 1374: show services alg statistics Output Fields *(Continued)*

Field Name	Field Description
`H323 Q931 process error`	Number of errors in processing Q931 packets.
`H323 H245 process error`	Number of errors in processing H245 packets.
`H323 do nat failed`	Number of NAT translation failures after packet decode.
`H323 do rm failed`	Number of H323 vsip table creation failures.
`H323 dscp marked`	Number of Differentiated Services code point (DSCP) packets marked.
`H323 dscp marked error`	Number of Differentiated Services code point (DSCP) packets marked as errors.
`RAS obj alloc failed`	Number of RAS session object memory allocation failures.
`RAS group alloc failed`	Number of RAS session group memory allocation failures.
`RAS packets dropped`	Number of RAS packets dropped.
`RAS packet exists in cookie error`	Number of times that some packets exist in existing RAS sessions cookie.
`RAS decode error`	Number of errors in decoding RAS packets.
`RAS flood error`	Number of gatekeeper requests that were dropped because of too many RAS request messages.

Table 1374: show services alg statistics Output Fields *(Continued)*

Field Name	Field Description
RAS do nat failed	Number of RAS session payload IP translation errors.
PPTP Objects Active	Number of PPTP objects active.
PPTP Objects Total	Number of PPTP objects in total.
PPTP Objects Error	Number of PPTP objects having errors.
PPTP ASL Group Active	Number of PPTP groups active.
PPTP ASL Group Total	Number of PPTP groups in total.
PPTP ASL Group Error	Number of PPTP groups having errors.
PPTP Packets received	Number of PPTP packets received.
PPTP Packets Discarded	Number of PPTP packets discarded.
PPTP Packets Free	Number of PPTP packets freed.
PPTP OCRQ Received	Number of Outgoing Call Requests received.

Table 1374: show services alg statistics Output Fields *(Continued)*

Field Name	Field Description
PPTP OCRQ Discarded	Number of Outgoing Call Requests discarded.
PPTP OCRP Received	Number of Outgoing Call Packets received.
PPTP OCRP Discarded	Number of Outgoing Call Packets discarded.
PPTP WEN(SLI) Received	Number of WEN (SLI) packets received.
PPTP WEN(SLI) Discarded	Number of WEN (SLI) packets discarded.
PPTP CCRQ-CDSN Received	Number of Call Clear Requests received.
PPTP CDSN Received	Number of Call Disconnection Notifications received.
PPTP CCRQ-CDSN Discarded	Number of Call Clear Requests discarded.
PPTP Session Create	Number of PPTP sessions created.
PPTP Session Destroy	Number of PPTP sessions destroyed.
PPTP Gate Create	Number of PPTP gates created.

Table 1374: show services alg statistics Output Fields *(Continued)*

Field Name	Field Description
PPTP Gate Hit	Number of PPTP gates hit.
PPTP Gate Timeout	Number of PPTP gates timed out.
PPTP NAT Events	Number of NAT events.
PPTP DO-NAT Total	Number of DO NATs in total.
PPTP DO-NAT Ok	Number of DO NATs okay.
PPTP DO-NAT Pending	Number of DO NATs pending.
PPTP DO-NAT Fail	Number of DO NATs failed.
PPTP DO-RM Total	Number of DO RMs in total.
PPTP DO-RM Ok	Number of DO RMs okay.
PPTP DO-RM Pending	Number of DO RMs pending.
PPTP DO-RM Fail	Number of DO RMs failed.
PPTP NAT-ASYNC Total	Number of NAT-ASYNCs in total.
PPTP NAT-ASYNC Invalid	Number of NAT-ASYNCs invalid.

Table 1374: show services alg statistics Output Fields *(Continued)*

Field Name	Field Description
PPTP NAT-ASYNC Error1	Number of NAT-ASYNCs error1.
PPTP NAT-ASYNC Error2	Number of NAT-ASYNCs error2.
PPTP ASL Hole Ok	Number of ASYNC holes okay.
PPTP ASL Hole Error	Number of ASYNC hole errors.
PPTP ASL First Hit	Number of ASYNC holes first hit.
PPTP ASL Hole Timeout	Number of ASYNC holes timed out.
PPTP ASL Invalid	Number of ASYNC holes invalid.
PPTP NAT Ctx Free	Number of NAT Ctxs free.
PPTP Create Resource Error	Number of create resource errors.
PPTP set S2C hole error	Number of server-to-client hole errors.
PPTP set C2S hole error	Number of client-to-server hole errors.
PPTP lnbrk error	Number of PPTP lnbrk errors.

Table 1374: show services alg statistics Output Fields *(Continued)*

Field Name	Field Description
`PPTP Mpool Create Error`	Number of Mpool create errors.
`PPTP RM register client Error`	Number of client registration errors.
`Call packet with rpcbind2`	Number of call packets with rpcbind2.
`Call packet with rpcbind3`	Number of call packets with rpcbind3.
`Call packet with rpcbind4`	Number of call packets with rpcbind4.
`Invalid rpcbind call`	Number of invalid rpcbind calls.
`Reply packet with rpcbind2`	Number of reply packets with rpcbind2.
`Reply packet with rpcbind3`	Number of reply packets with rpcbind3.
`Reply packet with rpcbind4`	Number of reply packets with rpcbind4.
`Invalid rpcbind reply`	Number of invalid rpcbind replies.
`Packets exceeded maximum length`	Number of packets exceeding maximum length.

Table 1374: show services alg statistics Output Fields *(Continued)*

Field Name	Field Description
`Packets dropped by ALG`	Number of packets dropped by the ALG.
`Number of describe messages received`	Number of describe messages received.
`Number of setup messages received`	Number of setup messages received.
`Number of teardown messages received`	Number of teardown messages received.
`Total packets dropped`	Total number of SIP packets dropped.
`Unexpected requests dropped`	Number of unexpected requests dropped.
`Unexpected responses dropped`	Number of unexpected responses dropped.
`Packets DSCP marked`	Number of Differentiated Services code point (DSCP) packets marked.
`Packets DSCP marked error`	Number of Differentiated Services code point (DSCP) packets marked as error.
`NAT errors`	Number of Network Address Translation errors.

Table 1374: show services alg statistics Output Fields *(Continued)*

Field Name	Field Description
RR headers exceeded maximum limits	Number of RR headers exceeded maximum limits.
Contact headers exceeded maximum limits	Number of contact headers exceeded maximum limits.
Invite dropped due to call limit	Number of invites dropped due to call limit.
Messages not processed by sip stack	Number of messages not processed by sip stack.
Unknown packets dropped	Number of unknown packets dropped.
Decoding Errors	Number of decoding errors.
Packets received in out of state	Number of packets received in out of state.
Packets received	Number of packets received.
Packets freed by ALG	Number of packets freed by ALG.
Gate fail errors	Number of gate fail errors.
Lookup packets	Number of lookup packets.

Table 1374: show services alg statistics Output Fields *(Continued)*

Field Name	Field Description
Announce packets	Number of announce packets.
Delete packets	Number of delete packets.
Number of packets received	Number of packets received.
Number of Invalid packets	Number of invalid packets.
Total number of sessions	Total number of sessions.
Number of actives sessions	Number of active sessions.

Sample Output

show services alg statistics application-protocol

While the statistics are the same for dce-rpc and dce-rpc-portmap, both rpc and rpc-portmap have the same output too.

```
user@router> show services alg statistics application-protocol dce-rpc
Interface name: ms-1/1/0
  DCE-RPC ALG statistics:
    Packets with wrong header : 0
    Non epm 3.0 packets       : 0
    Packets with type mismatch: 0
    Packets with id mismatch  : 0
    Packets with call mismatch: 0
    Packets fragmented        : 0
    Packets queued            : 0
```

```
    Packets dropped           : 0
    Packets released          : 0

user@router> show services alg statistics application-protocol dns
Interface name: ms-2/0/0
  DNS ALG statistics:
    Invalid packets received   : 0
    Reply packets received     : 3509
    Oversized packets received : 0

user@router> show services alg statistics application-protocol ftp
Interface name: ms-1/1/0
FTP ALG statistics:
Packets dropped           : 0
ALG parser errors         : 0
Packets translated        : 0

user@router> show services alg statistics  application-protocol h323
Interface name: ms-1/0/0
H323 ALG statistics:
  H323 total calls: 1
  H323 active calls: 1
  H323 gate install failed: 0
  H323 pinhole opened too late: 0
  H323 pinhole hit dropped: 0
  H323 gate timeout failed: 0
  H323 packets dropped: 0
  H323 get virtual ctx failed: 0
  H323 obj alloc failed: 0
  H323 group alloc failed: 0
  H323 ce alloc failed: 0
  H323 Q931 decode error: 0
  H323 H245 decode error: 0
  H323 Q931 process error: 0
  H323 H245 process error: 0
  H323 do nat failed: 0
  H323 do rm failed: 0
  H323 dscp marked: 0
  H323 dscp marked error: 0
  RAS obj alloc failed: 0
  RAS group alloc failed: 0
  RAS packets dropped: 0
  RAS packet exists in cookie error: 0
```

```
  RAS decode error: 0
  RAS flood error: 0
  RAS do nat failed: 0
user@router> show services alg statistics application-protocol ike-esp-nat
Interface name: ms-4/1/0
  IKE ESP ALG statistics:
    Session interests processed: 2
    Sessions created: 2
    Sessions destroyed: 1
    Control sessions created: 2
    Control sessions destroyed: 1
    Data sessions created: 0
    Data sessions destroyed: 0
    Gates created: 4
    Gate hits: 0
    Gates timedout: 4
user@router> show services alg statistics application-protocol pptp
Interface name: ms-2/0/0
  PPTP ALG statistics:
    PPTP Objects Active  : 1
    PPTP Objects Total  : 1
    PPTP Objects Error  : 0
    PPTP ASL Group Active  : 1
    PPTP ASL Group Total  : 1
    PPTP ASL Group Error  : 0
    PPTP Packets received  : 11
    PPTP Packets Discarded  : 0
    PPTP Packets Free  : 0
    PPTP OCRQ Received  : 1
    PPTP OCRQ Discarded  : 0
    PPTP OCRP Received  : 1
    PPTP OCRP Discarded  : 0
    PPTP WEN(SLI) Received  : 3
    PPTP WEN(SLI) Discarded  : 0
    PPTP CCRQ-CDSN Received  : 0
    PPTP CDSN Received  : 0
    PPTP CCRQ-CDSN Discarded  : 0
    PPTP Session Create  : 3
    PPTP Session Destroy  : 0
    PPTP Gate Create  : 0
    PPTP Gate Hit  : 2
    PPTP Gate Timeout  : 0
    PPTP NAT Events  : 0
```

```
    PPTP DO-NAT Total  : 1
    PPTP DO-NAT Ok  : 1
    PPTP DO-NAT Pending  : 0
    PPTP DO-NAT Fail  : 0
    PPTP DO-RM Total  : 1
    PPTP DO-RM Ok  : 2
    PPTP DO-RM Pending  : 0
    PPTP DO-RM Fail  : 0
    PPTP NAT-ASYNC Total  : 0
    PPTP NAT-ASYNC Invalid  : 0
    PPTP NAT-ASYNC Error1  : 0
    PPTP NAT-ASYNC Error2  : 0
    PPTP ASL Hole Ok  : 2
    PPTP ASL Hole Error  : 0
    PPTP ASL First Hit  : 2
    PPTP ASL Hole Timeout  : 0
    PPTP ASL Invalid  : 0
    PPTP NAT Ctx Free  : 0
    PPTP Create Resource Error  : 0
    PPTP set S2C hole error  : 0
    PPTP set C2S hole error  : 0
    PPTP lnbrk error  : 0
    PPTP Mpool Create Error  : 0
    PPTP RM register client Error  : 0

user@router> show services alg statistics application-protocol rpc
Interface name: ms-1/1/0
  RPC ALG statistics:
    Call packet with rpcbind2 : 2
    Call packet with rpcbind3 : 0
    Call packet with rpcbind4 : 0
    Invalid rpcbind call      : 0
    Reply packet with rpcbind2: 2
    Reply packet with rpcbind3: 0
    Reply packet with rpcbind4: 0
    Invalid rpcbind reply     : 0
    Packets fragmented        : 0
    Packets dropped           : 0
    Packets released          : 0

user@router> show services alg statistics application-protocol rtsp
Interface name: ms-0/1/0
  RTSP ALG statistics:
```

```
    Packets exceeded maximum length : 0
    Packets dropped by ALG : 0
    Number of describe messages received : 8
    Number of setup messages received : 30
    Number of teardown messages received : 7
user@router> show services alg statistics application-protocol rsh
Interface name: ms-2/0/0
  RSH ALG statistics:
    Invalid packets received   : 0
    Packets dropped by ALG   : 0
    ALG parser errors  : 0
    Packets freed by ALG     : 0

user@router> show services alg statistics application-protocol sip
Interface name: ms-2/0/0
  SIP ALG statistics:
    Total packets dropped     : 0
    Unexpected requests dropped : 0
    Unexpected responses dropped : 0
    Packets DSCP marked     : 0
    Packets DSCP marked error     : 0
    NAT errors : 0
    RR headers exceeded maximum limits : 0
    Contact headers exceeded maximum limits     : 0
    Invite dropped due to call limit : 0
    Messages not processed by sip stack : 0
    Unknown packets dropped     : 0
    Decoding Errors : 0
    Packets received in out of state     : 0

user@router> show services alg statistics application-protocol sql
Interface name: ms-2/0/0
  SQLNET ALG statistics:
    Packets received     : 5
    ALG parser errors     : 0
    Packets freed by ALG : 0
    Gate fail errors     : 0

user@router> show services alg statistics application-protocol talk
Interface name: ms-2/0/0
  TALK ALG statistics:
    Lookup packets     : 5
    Announce packets     : 0
```

```
    Delete packets : 0

user@router> show services alg statistics application-protocol tftp
Interface name: ms-0/0/0
  TFTP ALG statistics:
    Number of packets received : 0
    Number of Invalid packets     : 0
    Total number of sessions :   0
    Number of active sessions:   0
```

show services alg statistics interface

```
user@router> show services alg statistics interface ms-1/1/0
Interface name: ms-1/1/0
FTP ALG statistics:
Packets dropped           : 0
ALG parser errors         : 0
Packets translated        : 0
```

Release Information

Command introduced in Junos OS Release 10.4.

`h323` option introduced in Junos OS Release 17.1.

`ike-esp-nat` option introduced in Junos OS Release 17.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

show services analytics agent

IN THIS SECTION

- Syntax | 25375
- Description | 25375

- Options | 25375
- Required Privilege Level | 25375
- Output Fields | 25375
- Sample Output | 25377
- Release Information | 25377

Syntax

```
show services analytics agent
<brief | detail>
```

Description

Display information about running instances of Network Telemetry Framework (NTF) agent.

Options

none (Same as brief) Display summary information about analytics agents.

brief | detail (Optional) Display information about analytics agents for the specified level of output.

Required Privilege Level

view

Output Fields

Table 1375 on page 25376 lists the output fields for the `show services analytics agent` command. Output fields are listed in the approximate order in which they appear.

Table 1375: show services analytics agent Output Fields

Field Name	Field Description	Level of Output
`Agent ID`	Name of the agent.	`brief` none
`Output Plugins`	Number of output plug-ins configured for the agent.	`brief` none
`Input Plugins`	Number of input plug-ins configured for the agent.	`brief` none
`Process ID`	Number that uniquely identifies the active process for the service agent at the `brief` and none levels. At the `detail` level, the process ID is displayed for the analytics agent (the parent NTF agent) and for the active service agents.	All levels
`Analytics agent`	Information about the parent NTF agent.	`detail`
`Configuration File`	Path where the NTF agent configuration file is located.	`detail`
`Log File`	Path where logs are stored for the NTF agent.	`detail`
`Service Agent Count`	Number of active service agents.	`detail`
`Analytics Service agent(s)`	Information about the active service agents.	`detail`
`Agent Name`	Name of the service agent.	`detail`
`Input Plugin/s`	Name of all input plug-ins configured for the service agent.	`detail`
`Output Plugin/s`	Name of all output plug-ins configured for the service agent.	`detail`

Sample Output

show services analytics agent

```
user@host> show services analytics agent

Agent ID        Output Plugins    Input Plugins    Process ID
ipfix           1                 2                8368
```

show services analytics agent (Brief)

```
user@host> show services analytics agent brief
Agent ID        Output Plugins    Input Plugins    Process ID
ipfix           1                 2                8368
```

show services analytics agent (Detail)

```
user@host> show services analytics agent detail
Analytics agent:
Process ID          : 6246
Configuration File  : /var/etc/ntf-agent.conf
Log File            : /var/log/ntf-agent.log
Service Agent Count : 1
Analytics service agent(s):
  Agent Name        : ipfix
  Input Plugin/s    : input-ipfix
  Output Plugin/s   : output-ipfix
  Process ID        : 8368
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

IPFIX Mediation on the BNG

Configuring NTF Agent

show services anti-virus statistics

IN THIS SECTION

- Syntax | 25378
- Description | 25378
- Required Privilege Level | 25378
- Output Fields | 25378
- Sample Output | 25380
- Release Information | 25381

Syntax

```
show services anti-virus statistics
```

```
show services anti-virus machine-learning-scan-statistics
```

Description

Displays the statistics of antivirus database. Use this command to get an overview of how many files were scanned, how many viruses were found and blocked.

Required Privilege Level

View

Output Fields

Table 1376 on page 25379 lists the output fields for the `show services anti-virus statistics` command.

Table 1376: show services anti-virus statistics Output Fields

Field Name	Field Description
Virus DB type	The antivirus signature database type. Currently full_db is supported.
Total signatures	Total number of antivirus signatures at any given time.
Anti-virus DB version	The signature database version.
Anti-virus DB update time	The date and time when the antivirus signature database was last updated.
File scanned	Total number of files scanned by the device.
Virus found	Total number of virus signature hits.
Virus blocked	Number of viruses blocked based on the action set for the antivirus signature. The detected virus file will be immediately blocked after the signature match.
Virus permitted	Number of viruses permitted based on the action set for the antivirus signature.

Table 1377: show services anti-virus machine-learning-scan-statistics Output Fields

Field Name	Field Description
Machine learning scan engine version	The machine learning scan engine version.
Machine learning scan engine update time	The date and time when the machine learning scan engine was last updated.
File scanned	Total number of files scanned by the device.
Virus found	Total number of machine learning scan hits.

Table 1377: show services anti-virus machine-learning-scan-statistics Output Fields *(Continued)*

Field Name	Field Description
Virus blocked	Number of viruses blocked based on the action set for the machine learning scan. The detected virus file will be immediately blocked after the match.
Virus permitted	Number of viruses permitted based on the action set for the machine learning scan.

Sample Output

show services anti-virus statistics

```
show services anti-virus statistics
```

```
Anti-virus scan statistics:
  Virus DB type: anti-virus
  Total signatures: 11
  Anti-virus DB version: 1654594666
  Anti-virus DB update time: 2022-08-25 13:03:58 PDT
                     Total      HTTP       HTTPS      SMTP       SMTPS      IMAP       IMAPS
SMB
  File scanned:      419382     81947      177549     16067      31591      15994      31925
64309
  Virus found:       290713     1613       161485     15940      31591      15994      31925
32165
  Virus blocked:     290713     1613       161485     15940      31591      15994      31925
32165
  Virus permitted:   0          0          0          0          0          0          0
0
```

show services anti-virus machine-learning-scan-statistics

```
show services anti-virus machine-learning-scan-statistics
```

```
Anti-virus machine learning scan statistics:
  Machine learning scan engine version: 1696526121
  Machine learning scan engine update time: 2023-10-05 22:48:50 UTC
                     Total      HTTP       HTTPS      SMTP       SMTPS      IMAP       IMAPS
SMB
  File scanned:      359382     68947      154549     14367      24591      12494      20025
52309
  Virus found:       187713     1417       146795     13840      24591      12494      20025
25165
  Virus blocked:     187713     1417       146795     13840      24591      12494      20025
25165
  Virus permitted:   0          0          0          0          0          0          0
0
```

Release Information

Command introduced in Junos OS Release 23.4R1.

`show services anti-virus machine-learning-scan-statistics` command introduced in Junos OS Release 24.2R1.

show services application-aware-access-list flows

IN THIS SECTION

- Syntax | 25382
- Description | 25382
- Options | 25382
- Required Privilege Level | 25382
- Output Fields | 25382
- Sample Output | 25384

Release Information | 25386

Syntax

```
show services application-aware-access-list flows
<interface interface-name>
<subscriber subscriber-name>
```

Description

Display application-aware-access-list (AACL) flows. Offloading using JFM is supported only on MX Series routers with Modular Port Concentrators (MPCs).

Options

interface *interface-name*	Displays AACL flows for the specified interfaces only. The keyword, `interface`, must be appended to the command.
subscriber *subscriber-name*	Displays AACL flows for the specified subscribers only. The keyword, `subscriber`, must be appended to the command.

Required Privilege Level

view

Output Fields

Table 1378 on page 25383 lists the output fields for the `show services application-aware-access-list flows` command. Output fields are listed in the approximate order in which they appear.

Table 1378: show services application-aware-access-list flows Output Fields

Field Name	Field Description	Level of Output
`5-tuple`	This field comprises five components of the given flow. The components are: • Src IP • Dest IP • Src Port • Dest Port • Protocol	All levels
`Application-ID`	The identification number associated with the application.	All levels
`Dir`	The direction in terms of input or output. • Input (I) • Output (O)	All levels
`Off`	The status of offload to Packet Forwarding Engine. The various options are: • Not Offloaded (-) • Policer Offloaded, Flow Not Offloaded (P) • Policer Not Offloaded, Flow Offloaded (F) • Policer and Offloaded (P+F)	All levels
Off	The status of offload to Packet Forwarding Engine using JFM. The various options are: • Not Offloaded (-) • Offload requested but not completed (R) • Offload requested and completed (O)	All levels

Table 1378: show services application-aware-access-list flows Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Actions	The types of actions displayed are: • discard: (D) • accept : A • accept, count [T]: C-A or C-G or C-T • accept, fwd-class [C]: FC • accept, policer [P]: P • accept, count [T], fwd-class [C]: C-T+FC • accept, count [T], policer [P]: C-T+P • accept, fwd-class [C], policer [P]: FC+P • accept, count[T],fwd-class[C],policer[P]: C-T+FC+P	All levels

Sample Output

show services application-aware-access-list flows interface

```
user@host>show services application-aware-access-list flows interface ge-1/0/5.0
Interface: ge-1/0/5.0
service-set: aacl-countApps
service-set interface: ms-0/0/0
Currently active flows: 2
High watermark flows: 2

5-tuple                                          Application-ID                Dir Off Action
------------------------------------------------ ------------------------------ --- --- --------
  198.51.100.2:47072->  10.10.254.116:80   ,6     junos:http [64]               I   -   C-T
  10.10.254.116:80   ->    198.51.100.2:47072,6   junos:http [64]               O   -   C-T
```

show services application-aware-access-list flows subscriber

```
user@host>show services application-aware-access-list flows subscriber user@example.com
Subscriber: user@example.com

Service-set: ss1
Service-set interface: ms-2/0/0
Currently active flows: 4
High watermark flows: 40


5-tuple                                                          Application-ID
Dir    Off    Action

192.0.2.100:20109->160.200.200.200:80,17        junos:http [64]        I        -        C-T+FC+P
203.0.113.200:80->192.0.2..100:20109,17        junos:http [64]        O        -        C-T+FC+P
192.0.2.100:20108->203.0.113.100:80,17        junos:http [64]        I        P+F    C-T+FC+P
203.0.113.100:80->192.0.2.100:20108,17        junos:http [64]        O        P+F    C-T+FC
+P
```

show services application-aware-access-list flows subscriber (Offloading Using JFM)

```
user@host>show services application-aware-access-list flows subscriber user@example.com
Subscriber: user@example.com

Service-set: ss1
Service-set interface: ms-2/0/0
Currently active flows: 4
High watermark flows: 40


5-tuple                                                          Application-ID
Dir    Off    Action

192.0.2.100:20109->160.200.200.200:80,17     junos:http [64]              I   -   C-T+FC+P
203.0.113.200:80   ->192.0.2.100:20109,17  junos:http [64]              O   -   C-T+FC+P
192.0.2..100:20108->203.0.113.100:80,17     junos:http [64]              I   R   C-T+FC+P
203.0.113.100:80   ->192.0.2.100:20108,17  junos:http [64]              O   O   C-T+FC+P
```

Release Information

Command introduced in Junos OS Release 10.1.

Offload status for flows using Juniper Forwarding Mechanism (JFM) added in Junos OS Release 12.1.

RELATED DOCUMENTATION

AACL Overview

show services application-aware-access-list statistics

IN THIS SECTION

- Syntax | 25386
- Description | 25386
- Options | 25387
- Required Privilege Level | 25387
- Output Fields | 25387
- Sample Output | 25388
- Release Information | 25388

Syntax

```
show services application-aware-access-list statistics
<interface interface-name>
<subscriber subscriber-name>
```

Description

Display application-aware access list (AACL) statistics.

Options

`interface` *`interface-name`*	(Optional) Display AACL statistics for the specified interface only.
`subscriber` *`subscriber-name`*	(Optional) Display AACL statistics for the specified subscriber only.

Required Privilege Level

view

Output Fields

Table 1379 on page 25387 lists the output fields for the `show services application-aware-access-list statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1379: show services application-aware-access-list statistics Output Fields

Field Name	Field Description	Level of Output
`Interface`	Interface name.	Subscriber option
`Subscriber`	Subscriber identifier.	Interface option
`Service-set-interface`	Service set interface name.	All levels
`Service set`	Service set name.	All levels
`Application group`	Application group identifier.	All levels
`Packets in`	Number of ingress packets.	All levels
`Bytes in`	Number of ingress bytes.	All levels
`Packets out`	Number of egress packets.	All levels
`Bytes out`	Number of egress bytes.	All levels

Sample Output

show services application-aware-access-list statistics interface

```
user@host> show services application-aware-access-list statistics interface ge-0/0/0.100
Subscriber: user@example.com

service-set: IDP
service-set interface: ms-2/0/0

Application group     Application           Packets in            Bytes in              Packets
out           Bytes out

                      junos:ftp [63]        5                     334
6                     346
```

show services application-aware-access-list statistics subscriber

```
user@host> show services application-aware-access-list statistics subscriber user@example.com
Interface: ge-1/1/0.0

Service-set-interface: ms-1/3/0
Service set: aacl-svc-set

Application-aware-access-list statistics

Application group       Packets in       Bytes in         Packets out       Bytes out
P2P                                  400                    32025
200                     16284
FTP                                  20000               5231000       100
8700
```

Release Information

Command introduced in Junos OS Release 9.5.

show services application-identification application

IN THIS SECTION

- Syntax | 25389
- Description | 25389
- Options | 25389
- Required Privilege Level | 25389
- Output Fields | 25390
- Sample Output | 25393
- Sample Output | 25395
- Sample Output | 25398
- Release Information | 25401

Syntax

```
show services application-identification application (detail | summary)
```

Description

Display detailed information about a specified application signature, detailed information about all application signatures, or a summary of the existing application signatures.

Options

detail Display detailed information for all application signatures.

summary Display summary information for all application signatures.

Required Privilege Level

view

Output Fields

Table 1380 on page 25390 shows the output details for the `show services application-identification application detail` command.

Table 1380: show services application-identification application summary Output Fields

Field Name	Field Description
`Application(s)`	The number of applications present.
`Application`	Name of the custom application.
`Disabled`	The status of the application and whether the mapping method is currently used to identify this application.
`ID`	The unique ID number of an application. ID numbers 1 through 32,767 are automatically generated for applications; these IDs do not change. ID numbers for custom applications use 16,777,216 to 33,554,431.
`Order`	Number used to specify priority when multiple applications match the traffic. The lowest order number takes the highest priority.

show services application-identification application Output Fields on page 25390 lists the output fields for the `show services application-identification application` command. Output fields are listed in the approximate order in which they appear.

Table 1381: show services application-identification application Output Fields

Field Name	Field Description
`Application Name`	Name of the application.
`Application Type`	The basic application type, such as HTTP.
`Description`	A description of the application.

Table 1381: show services application-identification application Output Fields *(Continued)*

Field Name	Field Description
`Application ID`	The unique ID number of an application signature. ID numbers 1 through 32,767 are automatically generated for application; these IDs do not change. ID numbers for custom applications use 16,777,216 to 33,554,431.
`Priority`	Priority over other signature applications.
`Order`	
`Disabled`	The status of the application and whether the mapping method is currently used to identify this application.
`Cacheable`	The status whether the application identification results caching is enabled or not for the application. When this option is enabled, you can cache the application detection result in an ASC table.
`Configurable`	The status whether application is configurable or not.
`Activation Date`	Date when the application was activated for the first time.
`Last Modified`	Date when the application was last updated.
`Number of Parent Group(s)`	Total number of parent groups in this application signature group or cluster.
`Underlying consolidated Protocols/ports application is dependent on`	List of default protocols and ports for dependent applications of the specified application. • `Protocols`—List of default protocols. • `TCP ports`—List of default TCP ports.
`Layer-7 Immediate Protocol(s)`	List of applications over which that dynamic application can be identified.

Table 1381: show services application-identification application Output Fields *(Continued)*

Field Name	Field Description
`Application Specific Ports:`	The default port for this application type.
`Signature:`	Signature mapping criteria for application identification
`Protocol`	Application protocol
`Port range`	Port range. This option is applicable for TCP or UDP-based applications only.
`Member(s)`	Member name for a custom application signature. • Depth–Maximum number of bytes to check for context match. Byte limit for AppID to identify custom application pattern for applications running over TCP or UDP or Layer 7 applications. • Context–Service-specific context, such as http-header-content-type. • Pattern–Deterministic finite automaton (DFA) pattern matched on the context. The DFA pattern specifies the pattern to be matched for the signature. • Direction–Connection direction of the packets to match pattern (example : CTS [client-to-server])
Number of Parent Group	Number of parents application groups.
Application Groups	Application groups names.
Application tags	Application tag groups created to group related applications based on the attributes.
group-tags	Name of the tag group
characteristic	characteristic of application
risk	Associated risk
subcategory	Subcategory of the application

Table 1381: show services application-identification application Output Fields *(Continued)*

Field Name	Field Description
category	Category of the application.
Attribute	Attribute of the application. Shows `Obsolete` for the deprecated application

Sample Output

show services application-identification application summary

```
user@host> show services application-identification application summary
Application(s): 3616
  Applications                                  Disabled         ID      Order
  junos:SLACKER                                  No               1179    1
  junos:GOOGLE-TRUSTED-STORE                     No               2819    5
  junos:AMJILT                                   No               2272    4
  junos:DSI                                      No               2644    3
  junos:HLN                                      No               2096    2
  junos:ETSI-LI                                  No               537     1
  junos:CRAZYSALOON                              No               1720    5
  junos:EKSISOZLUK                               No               2436    4
  junos:SABAH                                    No               2574    3
  junos:AFREECA                                  No               2373    2
  junos:SENEWEB                                  No               2068    1
  junos:DIINO                                    No               776     5
  junos:CARE2                                    No               376     4
  junos:MOBAGE                                   No               1456    3
  junos:CARTOONNETWORK                           No               982     2
  junos:AVATARS-UNITED                           No               363     1
  junos:CONVIVA                                  No               2015    5
  junos:DREAMORA                                 No               1725    4
  junos:ELWATANNEWS                              No               2381    3
  junos:REUTERS                                  No               1044    2
  junos:BABYCENTER                               No               364     1
  junos:SOUTHWEST                                No               289     5
  junos:ONEDIO                                   No               2517    4
```

```
........
........
```

show services application-identification application detail

```
user@host> show services application-identification application detail junos:FTP

Application Name: junos:FTP
Application type: FTP
Description: This signature detects the File Transfer Protocol (FTP), which provides facilities
for transferring files to and from remote computer systems. It usually runs on TCP port 21.
Application ID: 45
Priority: high
Order: 0
Disabled: Yes
Cacheable: Yes
Activation Date: 2003-05-05
Last Modified: 2016-04-11
Number of Parent Group(s): 1
Application Groups:
    junos:infrastructure:file-servers
Application Tags:
    characteristic        : Supports File Transfer
    characteristic        : Known Vulnerabilities
    characteristic        : Capable of Tunneling
    risk                  : 3
    subcategory           : File-Servers
    category              : Infrastructure
Layer-7 Protocol(s):
    Protocol: TCP         / 205
    Protocol: SPDY        / 1469
    Protocol: SOCKS5      / 193
    Protocol: SOCKS4      / 192
    Protocol: HTTPS       / 68
    Protocol: HTTP2       / 2553
    Protocol: HTTP        / 67
Port Mapping:
    Default ports: TCP/21
```

show services application-identification application detail (Custom Applications)

```
user@host> show services application-identification application detail my-custom-app

Application Name: my-custom-app
Application type: MY-CUSTOM-APP
Description: custom App
Application ID: 16777216
Priority: high
Order: 65500
Disabled: No
Cacheable: No
Activation Date: N/A
Last Modified: N/A
Layer-7 Protocol(s):
    Protocol: http        / http
    Port range: N/A
    Member(s): 1
        Member m01
            Context: http-header-host
            Pattern: MY-SERVER.COM
            Direction: CTS
```

Sample Output

show services application-identification application detail (Unified Policies)

```
user@host> show services application-identification application detail


Application Name: junos:GOOGLE
Application type: GOOGLE
Description: This signature detects SSL connections to Google.com. Google is a
             company best known for their search engine but offers many cloud
             based services.
Application ID: 54
Priority: high
Order: 0
Disabled: No
```

```
Cacheable: No
Activation Date: 2003-05-05
Last Modified: 2017-06-28
Number of Parent Group(s): 2
Application Groups:
    junos:web:applications
    junos:web:portal
Application Tags:
    characteristic        : Can Leak Information
    characteristic        : Loss of Productivity
    characteristic        : Supports File Transfer
    risk                  : 3
    subcategory           : Applications
    category              : Web
Underlying consolidated Protocols/ports application is dependent on:
    Protocols:
        Protocol: junos:GOOGLE-GEN / 943
        Protocol: junos:STUN  / 201
        Protocol: junos:UDP   / 216
        Protocol: junos:TCP   / 205
        Protocol: junos:HTTP-PROXY / 2956
        Protocol: junos:SSL   / 199
        Protocol: junos:SPDY  / 1469
        Protocol: junos:POSTGRESQL / 150
        Protocol: junos:HTTPS / 68
        Protocol: junos:HTTP  / 67
        Protocol: junos:NET-PROXY / 2629
        Protocol: junos:HTTP2 / 2553
        Protocol: junos:HTTP-TUNNEL / 750
        Protocol: junos:COTP  / 22
        Protocol: junos:RTSP  / 176
        Protocol: junos:RTP   / 175
        Protocol: junos:DTLS  / 1291
        Protocol: junos:RTMP  / 337
        Protocol: junos:QUIC  / 2521
        Protocol: junos:JABBER / 94
    TCP Ports:
        Port: 443
        Port: 554
        Port: 80
    UDP Ports:
        Port: 554
Layer-7 Immediate Protocol(s):
```

```
    Protocol: GOOGLE-GEN  / 943
Alias List:
    junos:GOOGLE-SSL
Application Specific Ports:
    Default ports: N/A
Signature:
    Port range: N/A
    Client-to-server
    Order: 1
```

show services application-identification application detail (Junos OS Release 20.2R1)

```
user@host> show services application-identification application detail
Application Name: test
Application type: TEST
Description: N/A
Application ID: 16777221
Priority: high
Order: 65500
Disabled: No
Cacheable: No
Activation Date: N/A
Last Modified: N/A
Underlying consolidated Protocols/ports application is dependent on:
    Protocols:
        Protocol: junos:HTTP  / 67
        Protocol: junos:UDP   / 216
        Protocol: junos:TCP   / 205
        Protocol: junos:NET-PROXY / 2629
        Protocol: junos:SPDY  / 1469
        Protocol: junos:SSL   / 199
        Protocol: junos:LIBJINGLE-PSEUDOTCP / 3237
        Protocol: junos:STUN  / 201
        Protocol: junos:HTTPS / 68
        Protocol: junos:HTTP  / 67
        Protocol: junos:HTTP2 / 2553
        Protocol: junos:HTTP-TUNNEL / 750
        Protocol: junos:HTTP-PROXY / 2956
        Protocol: junos:HAPROXY / 3331
```

```
        Protocol: junos:COTP  / 22
    TCP Ports:
        Port: 80
        Port: 3128
        Port: 8000
        Port: 8080
Layer-7 Immediate Protocol(s):
    Protocol: HTTP        / 67
    Signature: fgnm
    Port range: N/A
    Member(s): 1
        Member m01
            Depth: 4
            Context: http-get-url-parsed-param-parsed
            Pattern: ads
            Direction: CTS
```

Sample Output

show services application-identification application detail (Junos OS Release 20.3 R1)

```
user@host> show services application-identification application detail junos:PDF
Application Name: junos:PDF
Application type: PDF
Description: This signature detects the download of PDF documents.
Application ID: 11046
Priority: high
Order: 0
Disabled: No
Cacheable: Yes
Configurable: No
Activation Date: 1999-12-31
Last Modified: 2019-12-31
Application Tags:
    characteristic        : Known Vulnerabilities
    characteristic        : Carrier of Malware
    characteristic        : Bandwidth Consumer
    risk                  : 4
    subcategory           : Multimedia
    category              : Web
Layer-7 Immediate Protocol(s):
```

```
    Protocol: SPDY        / 1469
    Protocol: HTTPS       / 68
    Protocol: HTTP2       / 2553
    Protocol: HTTP        / 67
Application Specific Ports:
    Default ports: N/A
Signature:
    Port range: N/A
    Client-to-server
    Order: 3
```

show services application-identification application detail junos:DNS-ENCRYPTED

```
user@host> show services application-identification application detail junos:DNS-ENCRYPTED


Application Name: junos:DNS-ENCRYPTED
Application type: DNS-ENCRYPTED
Description: This application is used to represent DNS Queries over HTTPS (DoH) and DNS over
Transport Layer Security (TLS)
Application ID: 33554507
Priority: high
Order: 0
Disabled: No
Cacheable: Yes
Configurable: Yes
Activation Date: N/A
Last Modified: N/A
Number of Parent Group(s): 1
Application Groups:
    junos:unassigned
Application Tags:
    risk                  : 4
    subcategory           : Networking
    category              : Infrastructure
Underlying consolidated Protocols/ports application is dependent on:
    Protocols:
        Protocol: junos:SSL   / 199
        Protocol: junos:TCP   / 205
        Protocol: junos:SPDY  / 1469
        Protocol: junos:LIBJINGLE-PSEUDOTCP / 3237
```

```
        Protocol: junos:UDP   / 216
        Protocol: junos:STUN  / 201
        Protocol: junos:HTTP-PROXY / 2956
        Protocol: junos:HTTPS / 68
        Protocol: junos:HTTP  / 67
        Protocol: junos:NET-PROXY / 2629
        Protocol: junos:HTTP2 / 2553
        Protocol: junos:HTTP-TUNNEL / 750
        Protocol: junos:HAPROXY / 3331
        Protocol: junos:COTP  / 22
    TCP Ports:
        Port: 853
        Port: 443
Layer-7 Immediate Protocol(s):
    Protocol: SSL         / 199
```

show services application-identification application detail junos:RLOGIN (Junos OS Release 21.1R1)

```
Application Name: junos:RLOGIN
Application type: RLOGIN
Description: This signature detects RLOGIN, a remote access protocol.
Application ID: 165
Priority: high
Order: 0
Disabled: No
Cacheable: Yes
Configurable: Yes
Activation Date: 2003-05-05
Last Modified: 2017-01-25
Number of Parent Group(s): 2
Application Groups:
    junos:all-new-apps
    junos:remote-access:command
Application Tags:
    group-tags            : new-app
    group-tags            : standardized
    group-tags            : remote_access
    group-tags            : enterprise
    characteristic        : Known Vulnerabilities
    risk                  : 1
```

```
    subcategory           : Command
    category              : Remote-Access
```

show services application-identification application detail junos:MXIT (Junos OS Release 21.1R1)

```
Application Name: junos:MXIT
Application type: MXIT
Description: This protocol plug-in is deprecated.
Application ID: 535
Priority: high
Order: 0
Disabled: No
Cacheable: Yes
Configurable: Yes
Activation Date: 2010-08-06
Last Modified: 2010-10-01
Number of Parent Group(s): 1
Application Groups:
junos:web:messaging:instant-messaging
Application Tags:
characteristic : Loss of Productivity
characteristic : Bandwidth Consumer
risk : 2
subcategory : Messaging
category : Web
attribute : Obsolete
Alias List:
junos:MXIT-TCP
Application Specific Ports:
Default ports: N/A
Signature:
Port range: N/A
Client-to-server
Order: 3
```

Release Information

Command introduced in Junos OS Release 11.4.

Starting in Junos OS Release 15.1X49-D100, the options `Cacheable`, `Activation Date`, and `Last modified` are introduced for `show services application-identification application detail` command.

The `Underlying consolidated Protocols/ports application is dependent on` and `Layer-7 Immediate Protocol(s)` options are introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

request services application-identification application | 14916

show services application-identification application micro-applications

IN THIS SECTION

- Syntax | 25402
- Description | 25402
- Options | 25403
- Required Privilege Level | 25403
- Output Fields | 25403
- Sample Output | 25403
- Release Information | 25404

Syntax

```
show services application-identification application micro-applications
```

Description

Displays the list of all micro-applications.

Micro-applications are a part of application signature package. You must enable micro-application detection in application identification and then use them as matching criteria in a security policy.

Options

Required Privilege Level

view

Output Fields

This command output displays the list of micro-applications.

Sample Output

show services application-identification application micro-applications

```
user@host> show services application-identification application micro-applications




Micro Applications
  junos:BACNET-GET-EVENT-INFORMATION
  junos:BACNET-SUBSCRIBE-COV-PROPERTY
  junos:BACNET-LIFE-SAFETY-OPERATION
  junos:BACNET-READ-RANGE
  junos:BACNET-REQUEST-KEY
  junos:BACNET-AUTHENTICATE
  junos:BACNET-VT-DATA
  junos:BACNET-VT-CLOSE
  junos:BACNET-VT-OPEN
  junos:BACNET-REINITIALIZE-DEVICE
  junos:BACNET-CONFIRMED-TEXT-MESSAGE
  junos:BACNET-CONFIRMED-PRIVATE-XFER
  junos:BACNET-DEVICE-COMM-CONTROL
  junos:BACNET-WRITE-PROP-MULTIPLE
  junos:BACNET-WRITE-PROPERTY
  junos:BACNET-READ-PROP-MULTIPLE
  junos:BACNET-READ-PROP-CONDITIONAL
  junos:BACNET-READ-PROPERTY
  junos:BACNET-DELETE-OBJECT
  junos:BACNET-CREATE-OBJECT
  junos:BACNET-REMOVE-LIST-ELEMENT
```

```
  junos:BACNET-ADD-LIST-ELEMENT
  junos:BACNET-ATOMIC-WRITE-FILE
  junos:BACNET-ATOMIC-READ-FILE
  junos:BACNET-SUBSCRIBE-COV
  junos:SIEMENS-S7-SETUP-COMM
  junos:SIEMENS-S7-UPLOAD-START
......
```

Release Information

Command introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Application Identification Support for Micro-Applications

show services application-identification application | 25389

show services application-identification application non-configurable

IN THIS SECTION

- Syntax | 25404
- Description | 25405
- Options | 25405
- Required Privilege Level | 25405
- Output Fields | 25405
- Sample Output | 25405
- Release Information | 25406

Syntax

```
show services application-identification application non-configurable
```

Description

Displays the list of non-configurable applications.

Non-configurable applications are the applications that cannot be configured as dynamic-application in a unified policy. These applications include the Layer 3 / Layer 4 applications and few Layer 7 applications which are set as non-configurable or deprecated based on the requirements. These non-configurable applications are available in signature package for backward compatibility with legacy application firewall configuration.

Options

Required Privilege Level

view

Output Fields

This command output displays the list of non-configurable applications.

Sample Output

show services application-identification application non-configurable

```
user@host> show services application-identification application non-configurable



Non-configurable Applications
  junos:INVALID
  junos:PPP
  junos:RIP2
  junos:8021Q
  junos:BASE
  junos:CDP
  junos:ETH
  junos:TCP
  junos:RIP1
  junos:ESTABLISHED
  junos:POP3S
```

```
junos:802-11
junos:MALFORMED
junos:SMTPS
junos:UDP
junos:IMAPS
junos:INCOMPLETE
junos:HTTPS
```

Release Information

Command introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Application Identification Support for Micro-Applications

show services application-identification application | 25389

show services application identification application obsolete applications

IN THIS SECTION

- Syntax | 25406
- Description | 25407
- Required Privilege Level | 25407
- Sample Output | 25407
- Release Information | 25407

Syntax

```
show services application-identification application obsolete-applications
```

Description

Displays the application signatures that are deprecated in the installed application signature package version on your security device.

Required Privilege Level

View

Sample Output

show services application-identification application obsolete-applications

```
user@host> show services application-identification application obsolete-applications
```

```
Obsolete Applications
  junos:CRAZYSALOON
  junos:MEGASHARES-COM
  junos:JOQ
  junos:SNAP-VPN
  junos:DIRECTDOWNLOADLINKS
  junos:CINEMAGEDDON
  junos:PRODAVALNIK
  junos:BLOCKBUSTER
  junos:DEALFISH
  junos:PRESENT
  junos:ONEWORLDTV
  junos:SOCIALCAM
  junos:SCIENCESTAGE
  junos:VPN-MASTERPRO
  junos:SLANDO
  junos:GOGOYOKO
  junos:BT-CHAT

..............
```

Release Information

Command introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Predefined Application Signatures for Application Identification

show services application-identification application-system-cache

IN THIS SECTION

- Syntax | 25408
- Description | 25408
- Options | 25408
- Required Privilege Level | 25409
- Output Fields | 25409
- Sample Output | 25412
- Release Information | 25413

Syntax

```
show services application-identification application-system-cache
<interface interface-name>
```

Description

Display the database of cached values stored by the application identification system.

NOTE: The `show services application-identification application-system-cache` command gives the information only when the application identifier (AI) is matched with the signature.

Options

none Display the database of cached values for the all services interfaces.

interface ***interface-name***	(Optional) Display the database of cached values for the specified services interface.

Required Privilege Level

view

Output Fields

Table 1382 on page 25409 lists the output fields for the `show services application-identification application-system-cache` command. Output fields are listed in the approximate order in which they appear.

Table 1382: show services application-identification application-system-cache Output Fields

Field Name	Field Description
`application-cache`	Status (**on** or **off**) of the application cache.
`cache-entry-timeout`	Number of seconds the mapping information is saved.
`pic`	PIC number of the accumulated statistics.
`IP address`	IP address of the traffic flow for which application-identification is enabled.
`Port`	Port number of the traffic flow for which application-identification is enabled.
`Protocol`	Protocol name of the flow for which application-identification is enabled.
`Application`	Application number, which is a unique identifier that denotes the application or service for which identification of traffic flows is enabled.
`Classification Path`	Protocols or nested applications that denote the paths traversed for classified packets.

Table 1382: show services application-identification application-system-cache Output Fields *(Continued)*

Field Name	Field Description
PIC	PIC number of the accumulated statistics. For the interface on which deep packet inspection (DPI) application is not running, that detail is also displayed for the corresponding interface.
Unknown applications	Number of unknown applications.
Cache hits	Number of sessions that matched the application in the application identification cache.
Cache misses	Number of sessions that did not find the application in the application identification cache.
Client-to-server packets processed	Number of client-to-server packets processed.
Server-to-client packets processed	Number of server-to-client packets processed.
Client-to-server bytes processed	Number of client-to-server payload bytes processed.
Server-to-client layer bytes processed	Number of server-to-client payload bytes processed.
Client-to-server packets processed	Number of client-to-server packets processed.
Server-to-client packets processed	Number of server-to-client packets processed.
Client-to-server bytes processed	Number of client-to-server payload bytes processed.

Table 1382: show services application-identification application-system-cache Output Fields *(Continued)*

Field Name	Field Description
`Server-to-client layer bytes processed`	Number of server-to-client payload bytes processed.
`Sessions bypassed due to resource allocation failure`	Number of sessions bypassed due to resource allocation failure.
`Segment case 1 - New segment to left`	Number of TCP segments contained before the previous segment.
`Segment case 2 - New segment overlap right`	Number of TCP segments that start before the previous segment and are contained in it.
`Segment case 3 - Old segment overlapped`	Number of TCP segments that start before the previous segment and extend beyond it.
`Segment case 4 - New segment overlapped`	Number of TCP segments that start and end within the previous segment.
`Segment case 5 - New segment overlap left`	Number of TCP segments that start within the previous segments and extend beyond it.
`Segment case 6 - New segment to right`	Number of TCP segments that start after the previous segment. This is the normal case.

Sample Output

show services application-identification application-system-cache

```
user@host> show services application-identification application-system-cache
Application System Cache Configurations:
  application-cache: on
  cache-entry-timeout: 3600 seconds
pic: ams0
pic: ms-0/3/0
ms-0/3/0 is not running DPI engine
pic: ams1
pic: ms-0/0/0
IP address: 192.0.2.2                                Port: 80     Protocol: TCP
Application: HTTP:YOUTUBE
Classification Path: IP:TCP:HTTP:YOUTUBE
```

show services application-identification application-system-cache interface

```
user@host> show services application-identification application-system-cache interface ms-1/0/0
Application System Cache Configurations:
  application-cache: on
  cache-entry-timeout: 3600 seconds
pic: ms-0/0/0
IP address: 192.0.2.2                                Port: 80     Protocol: TCP
Application: HTTP:YOUTUBE
Classification Path: IP:TCP:HTTP:YOUTUBE
user@host> show services application-identification counter

pic: ms-0/3/0
ms-0/3/0 is not running DPI engine

pic: ms1
  Counter type                                                     Value
Unknown applications                                                32682
Cache hits                                                          323504
Cache misses                                                        400
Client-to-server packets processed                                  2034
Server-to-client packets processed                                  1982
Client-to-server bytes processed                                    258786
```

```
Server-to-client bytes processed                                    1314722
Sessions bypassed due to resource allocation failure                0
Segment case 1 - New segment to left                                0
Segment case 2 - New segment overlap right                          0
Segment case 3 - Old segment overlapped                             0
Segment case 4 - New segment overlapped                             0
Segment case 5 - New segment overlap left                           0
Segment case 6 - New segment to right                               0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

request services application-identification application | 14916

show services application-identification application-system-cache (View)

IN THIS SECTION

- Syntax | 25414
- Description | 25414
- Options | 25414
- Required Privilege Level | 25414
- Output Fields | 25414
- Sample Output | 25416
- Release Information | 25418

Syntax

```
show services application-identification application-system-cache
<logical-system (logical-system-name | all | root-logical-system)>
<tenant (tenant-name | all)>
```

Description

Displays application details saved in application system cache. Use this command to get details about application detected in the traffic, the IP address from where the traffic was initiated, the protocol details, and the application signature that matched at that time stored in application cache.

Options

none	Displays application system cache for the root logical system, all logical systems, and all tenant systems.
logical-system *logical-system-name*	(Optional) Displays application system cache for the specified logical system.
logical-system *all*	(Optional) Displays application system cache for all the logical systems.
root-logical-system	(Optional) Displays application system cache for the root logical system.
tenant *tenant-name*	(Optional) Displays application system cache for the specified tenant system.
tenant *all*	(Optional) Displays application system cache for all the tenant systems.

Required Privilege Level

view

Output Fields

Table 1383 on page 25415 and Table 1384 on page 25415 list the output fields for the `show services application-identification application-system-cache` command. Output fields are listed in the approximate order in which they appear.

Table 1383: show services application-identification application-system-cache Output Fields

Field Name	Field Description
`application-cache`	On or Off status of the application cache.
`nested-application-cache`	On or Off status of the nested application cache.
`cache-unknown-result`	On or Off status for caching unknown results.
`cache-entry-timeout`	The number of seconds the mapping information is saved.
pic	PIC number of the accumulated statistics.
`Logical system name`	Name of a specific logical system.
`IP address`	IP address.
`Port`	Port number.
`Protocol`	Type of protocol.
`Application`	Name of the application.
`Encrypted`	Yes or No to identify the traffic as encrypted or not.

Table 1384: show services application-identification application-system-cache Output Fields (For Unified Policies)

Field Name	Field Description
`application-cache`	On or Off status of the application cache.

Table 1384: show services application-identification application-system-cache Output Fields (For Unified Policies) *(Continued)*

Field Name	Field Description
`Cache lookup for security-services`	On or Off status of the application cache for security services such as security policies, application firewall (AppFW), Juniper ATP Cloud, IDP, and Content Security. By default, the ASC is disabled for the security services.
`Cache lookup for miscellaneous-services`	On or Off status of the application cache for miscellaneous services such as APBR and AppTrack. By default, the ASC is enabled for the miscellaneous services.
`cache-entry-timeout`	The number of seconds the mapping information is saved.

Sample Output

show services application-identification application-system-cache

```
user@host> show services application-identification application-system-cache
Application System Cache Configurations:
  application-cache: on
  nested-application-cache: on
  cache-unknown-result: on
  cache-entry-timeout: 3600 seconds
  pic: 1/0
  Logical system name: root-logical-system
  IP address: 192.0.2.1                              Port: 443    Protocol: TCP
  Application: SSL                               Encrypted: Yes

  pic: 1/1
  Logical system name: root-logical-system
  IP address: 192.0.2.2                              Port: 80     Protocol: TCP
  Application: HTTP                              Encrypted: No
```

show services application-identification application-system-cache (Application System Cache Changes with Unified Policy Support)

```
user@host> show services application-identification application-system-cache

Application System Cache Configurations:
  application-cache: on
    Cache lookup for security-services: off
    Cache lookup for miscellaneous-services: on
  cache-entry-timeout: 3600 seconds
```

show services application-identification application-system-cache

```
user@host:TSYS1> show services application-identification application-system-cache
Application System Cache Configurations:
  application-cache: on
    Cache lookup for security-services: off
    Cache lookup for miscellaneous-services: on
  cache-entry-timeout: 3600 seconds
pic: 0/0
Logical system name: TSYS1
IP address: 4.0.0.1                                  Port: 22     Protocol: TCP
Application: SSH                                     Encrypted: No
Classification Path: IP:TCP:SSH
```

show services application-identification application-system-cache (When application identification is disabled)

```
user@host> show services application-identification application-system-cache

warning: application-identification subsystem has been disabled by the user
```

show services application-identification application-system-cache tenant TSYS1

```
user@host> show services application-identification application-system-cache tenant TSYS1
Application System Cache Configurations:
  application-cache: on
```

```
    Cache lookup for security-services: off
    Cache lookup for miscellaneous-services: on
  cache-entry-timeout: 3600 seconds
pic: 0/0
Logical system name: TSYS1
IP address: 192.0.2.0                          Port: 22     Protocol: TCP
Application: SSH                           Encrypted: No
Classification Path: IP:TCP:SSH
```

show services application-identification application-system-cache tenant all

```
user@host> show services application-identification application-system-cache tenant all
Application System Cache Configurations:
  application-cache: on
    Cache lookup for security-services: off
    Cache lookup for miscellaneous-services: on
  cache-entry-timeout: 3600 seconds
pic: 0/0
Logical system name: TSYS1
IP address: 192.0.2.0                       Port: 22     Protocol: TCP
Application: SSH                        Encrypted: No
Classification Path: IP:TCP:SSH

pic: 0/0
Logical system name: TSYS2
IP address: 203.0.113.0                     Port: 22     Protocol: TCP
Application: SSH                        Encrypted: No
Classification Path: IP:TCP:SSH
```

Release Information

Command introduced in Junos OS Release 10.2.

Command updated in Junos OS Release 12.1X47-D10.

Output updated in Junos OS Release 12.1X47-D15.

The `Cache lookup for security-services` and the `Cache lookup for miscellaneous-services` are introduced in Junos OS Release 18.2R1.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

clear services application-identification application-system-cache (Junos OS) | 4076

show services application-identification commit-status

IN THIS SECTION

- Syntax | 25419
- Description | 25419
- Required Privilege Level | 25419
- Sample Output | 25419
- Release Information | 25420

Syntax

```
show services application-identification commit-status]
```

Description

Display information about the commit status. Because the custom signatures commit is performed asynchronously, the command output shows the current status of your configuration commit.

Required Privilege Level

view

Sample Output

show services application-identification commit-status

```
user@host> show services application-identification commit-status
Custom signatures commit is in progress
```

show services application-identification commit-status

```
user@host> show services application-identification commit-status
Custom signatures committed successfully
```

show services application-identification commit-status

```
user@host> show services application-identification commit-status
Custom signatures serialization failed
```

Release Information

Command introduced in Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

request services application-identification application | 14916

show services application-identification commit-status (Next Gen Services)

IN THIS SECTION

- Syntax | 25421
- Description | 25421

- Required Privilege Level | 25421
- Sample Output | 25421
- Release Information | 25422

Syntax

```
show services application-identification commit-status]
```

Description

Display information about the commit status. Because the custom signatures commit is performed asynchronously, the command output shows the current status of your configuration commit.

Required Privilege Level

view

Sample Output

show services application-identification commit-status

```
user@host> show services application-identification commit-status
Custom signatures commit is in progress
```

show services application-identification commit-status

```
user@host> show services application-identification commit-status
Custom signatures committed successfully
```

show services application-identification commit-status

```
user@host> show services application-identification commit-status
Custom signatures serialization failed
```

Release Information

Statement introduced in Junos OS Releases 19.3R2 and 19.4R1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

show services application-identification counter

IN THIS SECTION

- Syntax | 25422
- Description | 25423
- Options | 25423
- Required Privilege Level | 25423
- Output Fields | 25423
- Sample Output | 25425
- Release Information | 25426

Syntax

```
show services application-identification counter
<interface interface-name>
```

Description

Display application identification counter statistics.

Options

none	Display counter statistics for all services interfaces.
`interface` *`interface-name`*	(Optional) Display counter statistics for the specified services interface.

Required Privilege Level

view

Output Fields

Table 1385 on page 25423 lists the output fields for the `show services application-identification counter` command. Output fields are listed in an approximate order in which they appear.

Table 1385: show services application-identification counter Output Fields

Field Name	Field Description
`PIC`	PIC number of the accumulated statistics.
`Unknown applications`	Number of unknown applications.
`Cache hits`	Number of sessions that matched the application in the application identification cache.
`Cache misses`	Number of sessions that did not find the application in the application identification cache.
`Client-to-server packets processed`	Number of client-to-server packets processed.

Table 1385: show services application-identification counter Output Fields *(Continued)*

Field Name	Field Description
`Server-to-client packets processed`	Number of server-to-client packets processed.
`Client-to-server bytes processed`	Number of client-to-server payload bytes processed.
`Server-to-client layer bytes processed`	Number of server-to-client payload bytes processed.
`Sessions bypassed due to resource allocation failure`	Number of sessions bypassed due to resource allocation failure.
`Segment case 1 - New segment to left`	Number of TCP segments contained before the previous segment.
`Segment case 2 - New segment overlap right`	Number of TCP segments that start before the previous segment and are contained in it.
`Segment case 3 - Old segment overlapped`	Number of TCP segments that start before the previous segment and extend beyond it.
`Segment case 4 - New segment overlapped`	Number of TCP segments that start and end within the previous segment.
`Segment case 5 - New segment overlap left`	Number of TCP segments that start within the previous segments and extend beyond it.
`Segment case 6 - New segment to right`	Number of TCP segments that start after the previous segment. This is the normal case.

Sample Output

show services application-identification counter

```
user@host> show services application-identification counter


pic: 5/0
  Counter type                                                    Value
 Unknown applications                                                 0
 Cache hits                                                           0
 Cache misses                                                         36
 Client-to-server packets processed                                   16
 Server-to-client packets processed                                   101
 Client-to-server bytes processed                                     3494
 Server-to-client bytes processed                                     112493
 Sessions bypassed due to resource allocation failure                 0
 Segment case 1 - New segment to left                                 11
 Segment case 2 - New segment overlap right                           8
 Segment case 3 - Old segment overlapped                              0
 Segment case 4 - New segment overlapped                              0
 Segment case 5 - New segment overlap left                            0
 Segment case 6 - New segment to right                                7

pic: 5/1
  Counter type                                                    Value
 Unknown applications                                                 0
 Cache hits                                                           0
 Cache misses                                                         0
 Client-to-server packets processed                                   0
 Server-to-client packets processed                                   0
 Client-to-server bytes processed                                     0
 Server-to-client bytes processed                                     0
 Sessions bypassed due to resource allocation failure                 0
 Segment case 1 - New segment to left                                 0
 Segment case 2 - New segment overlap right                           0
 Segment case 3 - Old segment overlapped                              0
 Segment case 4 - New segment overlapped                              0
 Segment case 5 - New segment overlap left                            0
 Segment case 6 - New segment to right                                0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

show services application-identification counter (AppSecure)

IN THIS SECTION

- Syntax | 25426
- Description | 25427
- Options | 25427
- Required Privilege Level | 25427
- Output Fields | 25427
- Sample Output | 25430
- Release Information | 25435

Syntax

```
show services application-identification counter
<ssl-encrypted-sessions>
<logical-system (logical-system-name | all | root-logical-system)>
<tenant (tenant-name | all)>
```

Description

Display the status of all Junos OS application identification counter values per SPU.

Options

none	Displays the application identification counter for the root logical system, all logical systems, and all tenant systems.
ssl-encrypted-sessions	(Optional) Displays counters for SSL encrypted sessions.
`logical-system` *`logical-system-name`*	(Optional) Displays the application identification counter for the specified logical system.
`logical-system` *`all`*	(Optional) Displays the application identification counter for all the logical systems.
`root-logical-system`	(Optional) Displays the application identification counter for the root logical system.
`tenant` *`tenant-name`*	(Optional) Displays the application identification counter for the specified tenant system.
`tenant` *`all`*	(Optional) Displays the application identification counter for all the tenant systems.

Required Privilege Level

view

Output Fields

Table 1386 on page 25428 lists the output fields for the `show services application-identification counter` command. Output fields are listed in an approximate order in which they appear.

Table 1386: show services application-identification counter Output Fields

Field Name	Field Description
`PIC`	PIC number of the accumulated statistics. **NOTE**: The PIC number is always displayed as 0 for SRX300, SRX320, SRX340, and SRX345 devices.
`Unknown applications`	Number of unknown applications.
`Encrypted unknown applications`	Number of encrypted unknown applications.
`Cache hits`	Number of sessions that matched the application in the AI cache.
`Cache hits pkt-plugin`	Number of packet plugin hits in a session.
`Cache hits stream-plugin`	Number of stream plugin hits in a session.
`Cache misses`	Number of sessions that did not find the application in the AI cache.
`Cache misses pkt-plugin`	Number of packet plugin miss in a session.
`Cache misses stream-plugin`	Number of stream plugin miss in a session
`Client-to-server packets processed`	Number of client-to-server packets processed.
`Server-to-client packets processed`	Number of server-to-client packets processed.
`Client-to-server bytes processed`	Number of client-to-server payload bytes processed.

Table 1386: show services application-identification counter Output Fields *(Continued)*

Field Name	Field Description
`Server-to-client layer bytes processed`	Number of server-to-client payload bytes processed.
`Client-to-server encrypted packets processed`	Number of client-to-server encrpted packets processed.
`Server-to-client encrypted packets processed`	Number of server-to-client encrpted packets processed.
`Client-to-server encrypted bytes processed`	Number of client-to-server encrpted payload bytes processed.
`Server-to-client encrypted bytes processed`	Number of server-to-client encrpted payload bytes processed.
`Sessions bypassed due to resource allocation failure`	Number of sessions bypassed due to resource allocation failure.
`Segment case 1 - New segment to left`	Number of TCP segments contained before the previous segment.
`Segment case 2 - New segment overlap right`	Number of TCP segments that start before the previous segment and are contained in it
`Segment case 3 - Old segment overlapped`	Number of TCP segments that start before the previous segment and extend beyond it.
`Segment case 4 - New segment overlapped`	Number of TCP segments that start and end within the previous segment.
`Segment case 5 - New segment overlap left`	Number of TCP segments that start within the previous segments and extend beyond it.

Table 1386: show services application-identification counter Output Fields *(Continued)*

Field Name	Field Description
`Segment case 6 - New segment to right`	Number of TCP segments that start after the previous segment. This is the normal case.

Sample Output

show services application-identification counter ssl-encrypted-sessions

```
user@host> show services application-identification counter ssl-encrypted-sessions

 pic: 1/0
Counter type                                                        Value
AI cache hits                                                         0
AI cache hits by nested application                                   0
AI cache misses                                                       0
AI matches                                                            0
AI uni-matches                                                        0
AI no-matches                                                         0
AI partial matches                                                    0
AI no-partial matches                                                 0
Sessions that triggered Appid create session API                      0
Sessions that do not incur signature match or decoding                0
Sessions that incur signature match or decoding                       0
Client-to-server packets processed                                    0
Server-to-client packets processed                                    0
Client-to-server layer-7 bytes processed                              0
Server-to-client layer-7 bytes processed                              0
Terminal first data packets on both direction                         0
pic: 1/1
Counter type                                                      Value
AI cache hits                                                         0
AI cache hits by nested application                                   0
AI cache misses                                                       0
AI matches                                                            0
AI uni-matches                                                        0
AI no-matches                                                         0
AI partial matches                                                    0
```

```
AI no-partial matches                                                  0
Sessions that triggered Appid create session API                       0
Sessions that do not incur signature match or decoding                 0
Sessions that incur signature match or decoding                        0
Client-to-server packets processed                                     0
Server-to-client packets processed                                     0
Client-to-server layer-7 bytes processed                               0
Server-to-client layer-7 bytes processed                               0
Terminal first data packets on both direction                          0
```

show services application-identification counter

```
user@host> show services application-identification counter
Logical System: root-logical-system

pic: 0/0
  Counter type                                                    Value
 Unknown applications                                                 0
 Encrpted unknown applications                                        0
 Cache hits pkt-plugin                                                0
 Cache hits stream-plugin                                             0
 Cache misses pkt-plugin                                              0
 Cache misses stream-plugin                                           0
 Client-to-server packets processed                                   0
 Server-to-client packets processed                                   0
 Client-to-server bytes processed                                     0
 Server-to-client bytes processed                                     0
 Client-to-server encrypted packets processed                         0
 Server-to-client encrypted packets processed                         0
 Client-to-server encrypted bytes processed                           0
 Server-to-client encrypted bytes processed                           0
 Sessions bypassed due to resource allocation failure                 0
 Segment case 1 - New segment to left                                 0
 Segment case 2 - New segment overlap right                           0
 Segment case 3 - Old segment overlapped                              0
 Segment case 4 - New segment overlapped                              0
 Segment case 5 - New segment overlap left                            0
 Segment case 6 - New segment to right                                0

Tenant: TSYS1
```

```
pic: 0/0
  Counter type                                                    Value
 Unknown applications                                                 0
 Encrpted unknown applications                                        0
 Cache hits pkt-plugin                                                0
 Cache hits stream-plugin                                             0
 Cache misses pkt-plugin                                              0
 Cache misses stream-plugin                                           0
 Client-to-server packets processed                                   983
 Server-to-client packets processed                                   0
 Client-to-server bytes processed                                     82572
 Server-to-client bytes processed                                     0
 Client-to-server encrypted packets processed                         0
 Server-to-client encrypted packets processed                         0
 Client-to-server encrypted bytes processed                           0
 Server-to-client encrypted bytes processed                           0
 Sessions bypassed due to resource allocation failure                 0
 Segment case 1 - New segment to left                                 0
 Segment case 2 - New segment overlap right                           0
 Segment case 3 - Old segment overlapped                              0
 Segment case 4 - New segment overlapped                              0
 Segment case 5 - New segment overlap left                            0
 Segment case 6 - New segment to right                                0
```

show services application-identification counter logical-system all

```
user@host> show services application-identification counter logical-system all
Logical System: root-logical-systempic: 0/0
  Counter type                                                    Value
 Unknown applications                                                 0
 Encrpted unknown applications                                        0
 Cache hits pkt-plugin                                                0
 Cache hits stream-plugin                                             0
 Cache misses pkt-plugin                                              0
 Cache misses stream-plugin                                           0
 Client-to-server packets processed                                   0
 Server-to-client packets processed                                   0
 Client-to-server bytes processed                                     0
 Server-to-client bytes processed                                     0
 Client-to-server encrypted packets processed                         0
 Server-to-client encrypted packets processed                         0
```

```
 Client-to-server encrypted bytes processed                              0
 Server-to-client encrypted bytes processed                              0
 Sessions bypassed due to resource allocation failure                    0
 Segment case 1 - New segment to left                                    0
 Segment case 2 - New segment overlap right                              0
 Segment case 3 - Old segment overlapped                                 0
 Segment case 4 - New segment overlapped                                 0
 Segment case 5 - New segment overlap left                               0
 Segment case 6 - New segment to right                                   0
```

show services application-identification counter

```
user@host:TSYS1> show services application-identification counter
Tenant: TSYS1

pic: 0/0
  Counter type                                                       Value
 Unknown applications                                                     0
 Encrpted unknown applications                                            0
 Cache hits pkt-plugin                                                    0
 Cache hits stream-plugin                                                 0
 Cache misses pkt-plugin                                                  0
 Cache misses stream-plugin                                               0
 Client-to-server packets processed                                       5
 Server-to-client packets processed                                       1
 Client-to-server bytes processed                                         1169
 Server-to-client bytes processed                                         73
 Client-to-server encrypted packets processed                             0
 Server-to-client encrypted packets processed                             0
 Client-to-server encrypted bytes processed                               0
 Server-to-client encrypted bytes processed                               0
 Sessions bypassed due to resource allocation failure                     0
 Segment case 1 - New segment to left                                     0
 Segment case 2 - New segment overlap right                               0
 Segment case 3 - Old segment overlapped                                  0
 Segment case 4 - New segment overlapped                                  0
 Segment case 5 - New segment overlap left                                0
 Segment case 6 - New segment to right                                    0
```

show services application-identification counter tenant all

```
user@host> show services application-identification counter tenant all
Tenant: TSYS1

pic: 0/0
  Counter type                                                    Value
 Unknown applications                                                 0
 Encrpted unknown applications                                        0
 Cache hits pkt-plugin                                                0
 Cache hits stream-plugin                                             0
 Cache misses pkt-plugin                                              0
 Cache misses stream-plugin                                           0
 Client-to-server packets processed                                   1006
 Server-to-client packets processed                                   0
 Client-to-server bytes processed                                     84504
 Server-to-client bytes processed                                     0
 Client-to-server encrypted packets processed                         0
 Server-to-client encrypted packets processed                         0
 Client-to-server encrypted bytes processed                           0
 Server-to-client encrypted bytes processed                           0
 Sessions bypassed due to resource allocation failure                 0
 Segment case 1 - New segment to left                                 0
 Segment case 2 - New segment overlap right                           0
 Segment case 3 - Old segment overlapped                              0
 Segment case 4 - New segment overlapped                              0
 Segment case 5 - New segment overlap left                            0
 Segment case 6 - New segment to right                                0

Tenant: TSYS2

pic: 0/0
  Counter type                                                    Value
 Unknown applications                                                 0
 Encrpted unknown applications                                        0
 Cache hits pkt-plugin                                                0
 Cache hits stream-plugin                                             0
 Cache misses pkt-plugin                                              0
 Cache misses stream-plugin                                           0
 Client-to-server packets processed                                   1006
 Server-to-client packets processed                                   0
 Client-to-server bytes processed                                     84504
```

```
Server-to-client bytes processed                                    0
Client-to-server encrypted packets processed                        0
Server-to-client encrypted packets processed                        0
Client-to-server encrypted bytes processed                          0
Server-to-client encrypted bytes processed                          0
Sessions bypassed due to resource allocation failure                0
Segment case 1 - New segment to left                                0
Segment case 2 - New segment overlap right                          0
Segment case 3 - Old segment overlapped                             0
Segment case 4 - New segment overlapped                             0
Segment case 5 - New segment overlap left                           0
Segment case 6 - New segment to right                               0
```

Release Information

Command introduced in Junos OS Release 10.2.

Output updated in Junos OS Release 12.1X47-D10.

Command and output updated in Junos OS Release 12.1X47-D15.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

clear services application-identification counter (Values) | 4078

show services application-identification entries

IN THIS SECTION

- Syntax | 25436
- Description | 25436
- Options | 25436
- Required Privilege Level | 25436

- Output Fields | 25437
- Sample Output | 25437
- Release Information | 25438

Syntax

```
show services application-identification entries (detail | filter)
```

Description

Displays detailed information or filtered information about a specified application signature or group signature, detailed information about all application signatures or application group signatures. Used to Support and improve the J-Web search mechanism and to search the applications easily.

Options

detail Displays detailed information for all application signatures or application group signatures.

filter Displays filtered information about a specified application signatures or application group signatures. Apply filter details like the `show services application-identification entries filter type application limit 1 offset 5` command. Similarly, any combination from allowed filters can be used.

category-list Displays the list of categories of available application signatures or application group signatures.

subcategory-list Displays the list of subcategories of available application signatures or application group signatures.

Required Privilege Level

view

Output Fields

The below table lists the output fields for the `show services application-identification entries` command.

Table 1387: show services application-identification entries Output Fields

Field Name	Field Description
`Entry-name`	Entry by name of the application or the group.
`Entry Type`	Type the application name or the group name to filter by application or group entry.
`Entry Category`	Entry by entry category name.
`Entry Subcategory`	Entry by subcategory name of the application.
`Entry Risk`	Entry risk. ID numbers 1 to 5. The default value is -1 when the risk is not configured.
`Entry Characteristic`	Entry by entry characteristic name.
`Entry Status`	Entry status enabled or disabled.
`Entry Predefined`	Entry by predefined or custom entry.
`Total Entries`	Number of entries in the application or group.

Sample Output

show services application-identification entries detail

```
user@host> show services application-identification entries detail
Entry Name: c1
Entry Type: application
Entry Category: (null)
```

```
Entry Subcategory: (null)
Entry Risk: -1
Entry Characteristic: (null)
Entry Status: enabled
Entry Predefined: custom
Total Entries: 1
```

show services application-identification entries filter

```
user@host> show services application-identification entries filter type application limit 1
offset 5
Entry Name: junos:104COM
Entry Type: application
Entry Category: Web
Entry Subcategory: miscellaneous
Entry Risk: 2
Entry Status: enabled
Entry Predefined: predefined
Total Entries: 1
```

Release Information

Command introduced in Junos OS Release 18.4R1.

The `category-list` and `subcategory-list` options are introduced in the Junos OS Release 19.1R1.

RELATED DOCUMENTATION

request services application-identification application | 14916

show services application-identification group

IN THIS SECTION

Syntax | 25439

- Description | 25439
- Options | 25439
- Required Privilege Level | 25439
- Output Fields | 25439
- Sample Output | 25440
- Release Information | 25443

Syntax

```
show services application-identification group [detail application-group name | summary]
```

Description

Display detailed or summary information about a specified application signature group or all application signature groups. Both custom and predefined application signature groups can be displayed.

Options

detail ***application-group name***	(Optional) Display detailed information for the specified application signature group.
summary	(Optional) Display summary information for all application signature groups.

Required Privilege Level

view

Output Fields

Table 1388 on page 25440 lists the output fields for the `show services application-identification group` command. Output fields are listed in the approximate order in which they appear.

Table 1388: show services application-identification group Output Fields

Field Name	Field Description
Description	Description of the specified application in the detailed display.
Group ID or ID	The unique ID number of an application signature or application signature group. ID numbers 1 through 32,767 are automatically generated for predefined application signatures and application signature groups; these IDs do not change. ID numbers for custom application signatures and application signature groups use ID numbers 32,768 to 65,534.
Disabled	The status of the application signature group and whether the signature method is currently used to identify this application. The default is No.
Application Group(s)	The application signature groups present.
Applications	The application signatures associated with this application signature group.
Number of Applications	Number of applications in the group
Number of Sub-Groups	Number of subgroups belonging to the application group.
Number of Parent-Groups	Number of parent group in the application group.
Tag Group	Tag group created to group applications based on the application attributes.
Tag group applications:	The application signatures associated with the tag group.

Sample Output

show services application-identification group summary

```
user@host> show services application-identification group summary

Application Group(s): 24
Application Groups                                  Disabled  ID
  my:enterprise                                      No        32770
```

```
  junos:enterprise:voip                          No        25
  junos:peer-to-peer:voip                        No        24
  junos:peer-to-peer:chat                        No        23
  junos:peer-to-peer:file-sharing                No        22
...
```

show services application-identification group detail

```
user@host> show services application-identification group detail junos:social-networking
Group Name: junos:social-networking
Group ID: 36
Description: N/A
Disabled: No
Number of Applications: 0
Number of Sub-Groups: 2
Number of Parent-Groups: 1
Sub Groups:
    junos:social-networking:applications
    junos:social-networking:business
```

show services application-identification group detail (Junos OS 21.1R1)

```
user@host> show services application-identification group detail junos:all-new-apps
  junos:all-new-apps
Group Name: junos:all-new-apps
Group ID: 32766
Description: N/A
Disabled: No
Number of Applications: 77
Number of Sub-Groups: 0
Number of Parent-Groups: 1
Applications:
    junos:RLOGIN
    junos:LINKEDIN
```

show services application-identification group detail (Junos OS 21.1)

```
user@host> show services application-identification group detail remote_acc_web_tags
 Group Name: remote_acc_web_tags
Group ID: 32770
Description: N/A
Disabled: No
Number of Applications: 75
Number of Sub-Groups: 0
Number of Parent-Groups: 0
Tag Group: tg2
    Applications Tags:
        social_network
Tag Group: tg1
    Applications Tags:
        web
        remote_access
Tag group applications:
    junos:FLIPBOARD
    junos:GATHER
    junos:FLICKR
    junos:YAMMER
    junos:TWITCH-VIDEO-STREAM
    junos:IMVU
    junos:FACEBOOK-APP
    junos:BEBO
    junos:ORKUT
    junos:SLIDESHARE
    junos:FACEBOOK-ACCESS
    junos:FUBAR
    junos:ATHLINKS
    junos:GOOGLE-PLUS
    junos:REDDIT
    junos:CLOOB
    junos:VIADEO-COOKIE
    junos:MYSPACE
```

show services application-identification group summary (Junos OS Release 21.1)

```
user@host> show services application-identification group summary
  Application Group(s): 92
Application Groups                                  Disabled  ID
  MY-GROUP                                           No        32780
  junos:all-new-apps                                 No        32766
  junos:behavioral                                   No        94
  junos:unassigned                                   No        89
  junos:web:proxy                                    No        48
  junos:remote-access:interactive-desktop            No        34
(...............)
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Predefined Application Signatures for Application Identification

show services application-identification packet-capture counters

IN THIS SECTION

- Syntax | 25444
- Description | 25444
- Required Privilege Level | 25444
- Output Fields | 25444
- Sample Output | 25445
- Release Information | 25446

Syntax

```
show services application-identification packet-capture counters
```

Description

Display the packet capture counter details for unknown application traffic.

Required Privilege Level

view

Output Fields

Table 1389 on page 25444 lists the output fields for the `show services application-identification packet-capture counters` command. Output fields are listed in the approximate order in which they appear.

Table 1389: show services application-identification packet-capture counters

Field Name	Field Description
Total sessions captured	Total number of sessions captured in the packet capture file
Total packets captured	Total number of packets captured in the packet capture file
Active sessions being captured	Number of active sessions currently being captured in the packet capture file
Sessions ignored because of memory allocation failures	Number of sessions not captured in the packet capture file because of memory allocation failure
Packets ignored because of memory allocation failures	Number of packets not captured in the packet capture file because of memory allocation failure
Ipc messages ignored because of storage limit	Number of interprocess communication (IPC) messages ignored because the storage limit is reached

Table 1389: show services application-identification packet-capture counters *(Continued)*

Field Name	Field Description
Sessions ignored because of buffer-packets limit	Number of sessions not captured in the packet capture file because the buffer packet limit is reached
Packets ignored because of buffer-packets limit	Number of packets not captured in packet capture file because buffer packet limit has reached
Inconclusive sessions captured	Number of inconclusive sessions captured in the packet capture file
Inconclusive sessions ignored	Number of inconclusive sessions not captured in the packet capture file
Cache entries timed out	Number of times the cache entries timeout value is reached

Sample Output

show services application-identification packet-capture counters

```
user@host> show services application-identification packet-capture counters

pic: 0/0
  Counter type                                                    Value
Total sessions captured                                               1
Total packets captured                                                6
Active sessions being captured                                        0
Sessions ignored because of memory allocation failures                0
Packets ignored because of memory allocation failures                 0
Ipc messages ignored because of storage limit                         0
Sessions ignored because of buffer-packets limit                      0
Packets ignored because of buffer-packets limit                       0
Inconclusive sessions captured                                        0
Inconclusive sessions ignored                                         0
Cache entries timed out                                               0
```

Release Information

Command introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

- packet-capture | 12137
- request services application-identification clear packet-capture all | 14918
- clear services application-identification packet-capture counters | 4081

show services application-identification statistics application-groups

IN THIS SECTION

- Syntax | 25446
- Description | 25446
- Options | 25447
- Required Privilege Level | 25447
- Output Fields | 25447
- Sample Output | 25448
- Release Information | 25450

Syntax

```
show services application-identification statistics applications
<interval>
<logical-system (logical-system-name | all | root-logical-system)>
<tenant (tenant-name | all)>
```

Description

Displays application group usage statistics.

Options

none	Displays application group usage statistics for the root logical system, all logical systems, and all tenant systems.
`interval`	(Optional) Displays interval statistics per application group. Interval statistics are displayed in Top-N format, such that the first application group displayed has the largest byte count. The default interval is 1, if no parameter is specified. The current interval is 1.
`logical-system` *`logical-system-name`*	(Optional) Displays application group usage statistics for the specified logical system.
`logical-system` *`all`*	(Optional) Displays application group usage statistics for all the logical systems.
`root-logical-system`	(Optional) Displays application group usage statistics for the root logical system.
`tenant` *`tenant-name`*	(Optional) Displays application group usage statistics for the specified tenant system.
`tenant` *`all`*	(Optional) Displays application group usage statistics for all the tenant systems.

Required Privilege Level

view

Output Fields

Table 1390 on page 25447 lists the output fields for the `show services application-identification statistics application-groups` command. Output fields are listed in the approximate order in which they appear.

Table 1390: show services application-identification statistics application-groups Output Fields

Field Name	Field Description
`Last Reset`	Displays date, time, and how long ago the statistics for the sessions were cleared. The format None specified is in *year-month-day hour:minute:second timezone*. If you did not clear the statistics previously at any point, `Never` is displayed.
`Application Group`	Displays the name of the application group.

Table 1390: show services application-identification statistics application-groups Output Fields *(Continued)*

Field Name	Field Description
`Sessions`	Displays the number of sessions for the application group.
`Kilo Bytes`	Displays the size of the application group in kilobytes. **NOTE**: When an SRX Series Firewall is operating in Chassis Cluster mode (Active/Active mode - Z mode), the `show services application-identification statistics application-groups` command output does not provide complete statistics for bytes count for the session in application/application group statistics. This is because, ingress and egress traffic byte counts are updated separately on the primary and secondary nodes in the chassis cluster setup for a given application.

Sample Output

show services application-identification statistics application-groups

```
user@host> show services application-identification statistics application-groups

Last Reset: 2014-02-19 00:38:01 PST
                Application Group        Sessions      Kilo Bytes
             junos:infrastructure               2              18
             junos:encryption                   1               2
  junos:infrastructure:monitoring               2              18
```

show services application-identification statistics application-groups interval 1

```
user@host> show services application-identification statistics application-groups interval 1
Logical System: root-logical-system
Interval Start: 2018-07-16 16:11:27 PDT
Elapsed time: 04:56:01
```

show services application-identification statistics application-groups logical-system all

```
user@host> show services application-identification statistics application-groups logical-system
all
Logical System: root-logical-system
Last Reset: 2018-06-21 16:11:21 PDT
```

show services application-identification statistics application-groups

```
user@host:TSYS1> show services application-identification statistics application-groups
Tenant: TSYS1
Last Reset: 2019-05-12 23:47:58 PDT
               Application Group           Sessions         Kilo Bytes
            junos:infrastructure                  3                  0
 junos:infrastructure:networking                  3                  0
 junos:infrastructure:networking:icmp             3                  0
             junos:remote-access                  1                 10
     junos:remote-access:command                  1                 10
   junos:remote-access:tunneling                  1                 10
```

show services application-identification statistics application-groups tenant TSYS1

```
user@host> show services application-identification statistics application-groups tenant TSYS1
Tenant: TSYS1
Last Reset: 2019-05-12 23:47:58 PDT
                        Application Group   Sessions   Kilo Bytes
                     junos:infrastructure          3            0
          junos:infrastructure:networking          3            0
     junos:infrastructure:networking:icmp          3            0
                      junos:remote-access          1           10
              junos:remote-access:command          1           10
            junos:remote-access:tunneling          1           10
```

show services application-identification statistics application-groups tenant all

```
user@host> show services application-identification statistics application-groups tenant all
Tenant: TSYS1
```

```
Last Reset: 2019-05-12 23:47:58 PDT
                 Application Group           Sessions         Kilo Bytes
              junos:infrastructure                  3                  0
   junos:infrastructure:networking                  3                  0
junos:infrastructure:networking:icmp                3                  0
               junos:remote-access                  1                 10
       junos:remote-access:command                  1                 10
     junos:remote-access:tunneling                  1                 10

Tenant: TSYS2
Last Reset: 2019-05-12 23:47:58 PDT
                 Application Group           Sessions         Kilo Bytes
              junos:infrastructure                  3                  0
   junos:infrastructure:networking                  3                  0
junos:infrastructure:networking:icmp                3                  0
               junos:remote-access                  1                 10
       junos:remote-access:command                  1                 10
     junos:remote-access:tunneling                  1                 10
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Command introduced in Junos OS Release 11.4.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

statistics (Services) | **28599**

clear services application-identification application-statistics | **4072**

show services application-identification statistics applications

IN THIS SECTION

- Syntax | 25451
- Description | 25451
- Options | 25451
- Required Privilege Level | 25452
- Output Fields | 25452
- Sample Output | 25453
- Sample Output | 25455
- Release Information | 25456

Syntax

```
show services application-identification statistics applications
interval
detail
<logical-system (logical-system-name | all | root-logical-system)>
<tenant (tenant-name | all)>
```

Description

Displays application usage statistics. Application statistics allow an administrator to access cumulative statistics as well as statistics accumulated over user-defined intervals. The security devices support a history of eight intervals that an administrator can use to display application session and byte counts. Starting in Junos OS 18.3R1, the security devices support a history of one interval.

Options

none Displays the application identification statistics for the root logical system, all logical systems, and all tenant systems.

interval (Optional) Displays interval statistics per application. Interval statistics are displayed in Top-N format, such that the first application displayed has the largest byte count.

- **Default:** 1440 minutes

  Starting in Junos OS Release 15.1X49-D120 and Junos OS Release 18.1R1, the default interval value is changed from 1 minute to 1440 minutes.

`logical-system` *`logical-system-name`*	(Optional) Displays the application identification statistics for the specified logical system.
`logical-system` *`all`*	(Optional) Displays the application identification statistics for all the logical systems.
`root-logical-system`	(Optional) Displays the application identification statistics for the root logical system.
`tenant` *`tenant-name`*	(Optional) Displays the application identification statistics for the specified tenant system.
`tenant` *`all`*	(Optional) Displays the application identification statistics for all the tenant systems.
`details`	(optional) Displays the details such as sessions completed final application classification, sessions closed prematurely, and number sessions forced classified as final-match on reaching any of the TCP/UDP or global limits.

Required Privilege Level

view

Output Fields

Table 1391 on page 25452 lists the output fields for the `show services application-identification statistics applications` command. Output fields are listed in the approximate order in which they appear.

Table 1391: show services application-identification statistics applications Output Fields

Field Name	Field Description
`Application`	Name of the application or micro-application.
`Sessions`	Number of sessions for the application.

Table 1391: show services application-identification statistics applications Output Fields *(Continued)*

Field Name	Field Description
`Bytes`	Size of the application in bytes. **NOTE**: When an SRX Series Firewall is operating in chassis cluster mode (Active/Active mode - Z mode), the `show services application-identification statistics applications` command output does not provide complete statistics for bytes count for the session in application/application group statistics. This is because, ingress and egress traffic byte counts are updated separately on the primary and secondary nodes in the chassis cluster setup for a given application.
`Encrypted`	Yes or No identifying the traffic as encrypted or not.
`Last Reset`	Displays date, time, and how long ago the statistics for the sessions were cleared. The format None specified is in *year-month-day hour:minute:second timezone*. If you did not clear the statistics previously at any point, `Never` is displayed.
`DPI Final-match`	Number of sessions completed final application classification.
`Pre-matched`	Number sessions closed pre-maturely before reaching the final classification.
`Limits Final-matched`	Number sessions forced classified as final-match on reaching any of the limits (TCP or UDP or Global).

Sample Output

show services application-identification statistics applications

```
user@host> show services application-identification statistics applications

Last Reset: 2014-02-19 00:38:01 PST
                      Application           Sessions              Bytes    Encrypted
                           SYSLOG                  2              18610           No
```

show services application-identification statistics applications interval 1

```
user@host> show services application-identification statistics applications interval 1
Logical System: root-logical-system
Interval Start: 2018-07-16 16:11:27 PDT
Elapsed time: 04:47:50
```

show services application-identification statistics applications logical-system all

```
user@host> show services application-identification statistics applications logical-system all
Logical System: root-logical-system
Last Reset: 2018-06-21 16:11:21 PDT
```

show services application-identification statistics applications

```
user@host> show services application-identification statistics applications
Last Reset: 2019-05-21 22:48:53 PDT
Application           Sessions        Bytes    Encrypted
HTTP                   1              6022           No
ICMP-ECHO              12             1764           No
```

show services application-identification statistics applications

```
user@host:TSYS1> show services application-identification statistics applications
Tenant: TSYS1
Last Reset: 2019-05-13 03:02:56 PDT
        Application   Sessions  Bytes  Encrypted
          ICMP-ECHO         10   1680         No
```

show services application-identification statistics applications tenant TSYS1

```
user@host> show services application-identification statistics applications tenant TSYS1
Tenant: TSYS1
Last Reset: 2019-05-12 23:47:58 PDT
                      Application    Sessions   Bytes  Encrypted
```

```
                        ICMP-ECHO           3     504         No
                              SSH           1   10890         No
```

show services application-identification statistics applications tenant all

```
user@host> show services application-identification statistics applications tenant all
Tenant: TSYS1
Last Reset: 2019-05-12 23:47:58 PDT
                      Application    Sessions   Bytes  Encrypted
                        ICMP-ECHO           3     504         No
                              SSH           1   10890         No
Tenant: TSYS2
Last Reset: 2019-05-12 23:47:58 PDT
                      Application    Sessions   Bytes  Encrypted
                        ICMP-ECHO           3     504         No
                              SSH           1   10890         No
```

show services application-identification statistics applications interval 1

```
user@host:TSYS1> show services application-identification statistics applications interval 1
Tenant: TSYS1
Interval Start: 2019-05-13 03:04:16 PDT
Elapsed time: 00:00:04
    Application           Sessions           Bytes    Encrypted
      ICMP-ECHO                  4             672           No
```

Sample Output

show services application-identification statistics applications details (Junos OS Release 20.3R1)

```
user@host> show services application-identification statistics applications details
```

command-name

```
Logical System: root-logical-system
Last Reset: 2020-05-08 08:55:31 PDT
            Application  Enc    DPI final-match    Pre-match  Limits
                                                                final-match
                    NTP   No                  1            0            0
                 SYSLOG   No                  5            0            0
```

Release Information

Command introduced in Junos OS Release 11.4.

Command updated in Junos OS Release 12.1.

Statement introduced in Junos OS Release 17.1 on MX Series.

`logical-system` option introduced in Junos OS Release 18.3R1.

Command is updated to include micro-applications in Junos OS Release 19.2R1.

`tenant` option introduced in Junos OS Release 19.4R1.

`detail` option introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

statistics (Services) | 28599

clear services application-identification application-statistics | 4072

show services application-identification status

IN THIS SECTION

Syntax | 25457

- Description | 25457
- Required Privilege Level | 25457
- Output Fields | 25457
- Sample Output | 25460
- Release Information | 25469

Syntax

```
show services application-identification status
```

Description

Displays detailed information about application identification status.

Displays the application identification details on your security device. The output includes information such as status of application identification and application system cache on your device, application signature package details, download server, and so on.

Required Privilege Level

view

Output Fields

"No Link Title" on page 25457 lists the output fields for the `show services application-identification status` command. Output fields are listed in the approximate order in which they appear.

show services application-identification status Output Fields

Field Name	Field Description
`Status`	Status of application identification: `Enabled` or `Disabled`.
`Sessions under app detection`	Sessions undergoing application identification detection.

Field Name	Field Description
`Engine Version`	Application identification detector engine version. This field displays `0` when there is no JDPI-Decoder engine installed or uninstalled, and displays the JDPI-Decoder engine version when it is installed.
`Max TCP session packet memory`	Maximum number of TCP sessions that application identification maintains.
`Force packet plugin`	Force packet plugin status: `Enabled` or `Disabled`.
`Force stream plugin`	Force stream plugin status: `Enabled` or `Disabled`.
`DPI Performance mode`	DPI performance mode status. This field is displayed only if the DPI performance mode is enabled.
`Statistics collection interval`	Frequency (in minutes) for collecting statistics.
`Status`	Status of application system cache: `Enabled` or `Disabled`.
`Negative cache status`	Status on the number of sessions that reach the Unknown cache entry: `Enabled` or `Disabled`.
`Max Number of entries in cache`	Maximum number of cache entries.
`Cache timeout`	Idle timeout after which the cache entries expires.
`Download Server CGI`	Name of the server from where protocol bundle was downloaded.
`Auto Update`	Status of auto update to receive protocol bundle updates from the server: `Enabled` or `Disabled`.
Release Date	Display release date of the application signature package.
`Status`	Status of protocol bundle: `Active` or `Free`.
Application Package Version	Version of the application signature package.

`Version`
Or
`PB Version`
	Version of protocol bundle.
	NOTE: Starting from Junos OS Release 17.4R1, the field `PB Version` is used for displaying version of the protocol bundle.

Field Name	Field Description
`Proxy Profile`	Display the proxy profile name. If you have disabled proxy server for downloading application signature package, the `Proxy Profile` displays `Not Configured`.
`Proxy Address`	Display the IP address and the port number of the proxy server.
`Session`	The number of active sessions.
`Micro apps version`	The version of micro-applications.
`Session`	The number of active sessions.
`Rollback version details`	Rolled back version of the application signature package details such as application package version, protocol bundle version, application identification detector engine version, micro-applications version.
`Custom-App Infra Version`	Custom application infra version.

Starting from Junos OS Release 17.4R1, Juniper Networks Deep Packet Inspection-Decoder (JDPI-Decoder) engine, is packaged along with the application signature package version 534 that includes protobundle version 1.270.0.48.005. When you upgrade to Junos OS Release 17.4R1 or later from the earlier versions of Junos OS, the application identification security package installed is of version 534.

Starting in Junos OS Release 12.3X48-D80, on SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, SRX650, SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800 Series devices, the Juniper Networks Deep Packet Inspection-Decoder (JDPI-Decoder) engine is separated from Junos OS and allows you to download the JDPI-Decoder engine along with the protobundle. This implementation allows you to upgrade the JDPI-Decoder engine separately without upgrading Junos OS.

However, if you require latest versions of the protocol bundle, you must download and install the application signature package separately.

Starting in Junos OS Release 19.2R1, Junos OS Release 15.1X49-D200, and Junos OS Release 12.3X48-D95, the Juniper Networks Deep Packet Inspection-Decoder (JDPI-Decoder) engine comes with a default application signature package version 999 that includes the protobundle version 1.380.0-64.005 and the JDPI-Decoder engine version 5.3.0-56. You can upgrade the application signature package when a new signature package version is available.

Sample Output

show services application-identification status

```
user@host> show services application-identification status
pic: 5/0

Application Identification
 Status                            Enabled
 Sessions under app detection      0
 Engine Version                    4.18.1-20 (build date Feb 15 2014)
 Max TCP session packet memory     30000
 Force packet plugin               Disabled
 Force stream plugin               Disabled
 Statistics collection interval    1 (in minutes)

Application System Cache
 Status                            Enabled
 Negative cache status             Disabled
 Max Number of entries in cache    131072
 Cache timeout                     3600 (in seconds)

Protocol Bundle
 Download Server                   https://signatures.juniper.net/cgi-bin/index.cgi
 AutoUpdate                        Disabled
Slot 1:
 Status                            Active
 Version                           1.30.4-22.005 (build date Jan 17 2014)
 Sessions                          0
Slot 2
 Status                            Free
```

show services application-identification status (Junos OS Release 19.2R1 and Later)

```
user@host> show services application-identification status
Application Identification
 Status                            Enabled
 Sessions under app detection      0
 Max TCP session packet memory     0
 Force packet plugin               Disabled
```

```
 Force stream plugin               Disabled
 DPI Performance mode:             Enabled
 Statistics collection interval    1440 (in minutes)

Application System Cache
 Status                            Enabled
  Cache lookup security-services   Enabled
  Cache lookup miscellaneous-services Enabled
 Max Number of entries in cache    0
 Cache timeout                     3600 (in seconds)

Protocol Bundle
 Download Server                   https://devdb.secteam.juniper.net/cgi-bin/index.cgi
 AutoUpdate                        Disabled

Proxy Details
 Proxy Profile                     Not Configured
Slot 1:
 Application package version       50041
 Status                            Active
 PB Version                        1.380.0-64.005 (build date May  6 2019)
 Engine version                    5.3.0-56 (build date Mar  6 2019)
 Micro-App Version                 0
 Sessions                          0
```

show services application-identification status (DPI Performance Mode Enabled)

```
user@host> show services application-identification status
pic: 2/1

Application Identification
Status                            Enabled
Sessions under app detection      0
Engine Version                    4.18.2-24.006 (build date Jul 30 2014)
Max TCP session packet memory     30000
Force packet plugin               Disabled
Force stream plugin               Disabled
DPI Performance mode:             Enabled
Statistics collection interval    1 (in minutes)

Application System Cache
```

```
Status                              Enabled
Negative cache status               Disabled
Max Number of entries in cache      262144
Cache timeout                       3600 (in seconds)

Protocol Bundle
Download Server                     https://signatures.juniper.net/cgi-bin/index.cgi
AutoUpdate                          Disabled
Slot 1:
Application package version         2399
Status                              Active
Version                             1.40.0-26.006 (build date May 1 2014)
Sessions                            0
Slot 2
Application package version         0
Status                              Free
Version
Sessions                            0
```

show services application-identification status (Application Identification Detector Engine Version)

```
user@host> show services application-identification status
Application Identification
 Status                              Enabled
 Sessions under app detection        0
 Max TCP session packet memory       0
 Force packet plugin                 Disabled
 Force stream plugin                 Disabled
 Statistics collection interval      1 (in minutes)

Application System Cache
 Status                              Enabled
 Max Number of entries in cache      131072
 Cache timeout                       3600 (in seconds)

Protocol Bundle
 Download Server                     https://indiavm-sigdb2.englab.juniper.net/cgi-bin/index.cgi
 AutoUpdate                          Disabled
Slot 1:
 Application package version         534
```

```
 Status                             Active
 PB Version                         1.270.0-48.005 (build date May 22 2017)
 Engine version                     4.20.0-49.005 (build date May 22 2017)
 Sessions                           0
```

show services application-identification status (Download Through Proxy Server)

```
user@host> show services application-identification status
Application Identification
 Status                             Enabled
 Sessions under app detection       0
 Max TCP session packet memory      0
 Force packet plugin                Disabled
 Force stream plugin                Disabled
 DPI Performance mode:              Enabled
 Statistics collection interval     1440 (in minutes)

Application System Cache
 Status                             Enabled
  Cache lookup security-services    Enabled
  Cache lookup miscellaneous-services Enabled
 Max Number of entries in cache     131072
 Cache timeout                      3600 (in seconds)

Protocol Bundle
 Download Server                    https://signatures.juniper.net/cgi-bin/index.cgi
 AutoUpdate                         Disabled

Proxy Details
 Proxy Profile                      p1
 Proxy Address                      http://5.0.0.1:3128
Slot 1:
 Application package version        3058
 Status                             Active
 PB Version                         1.340.0-57.005 (build date Apr 19 2018)
 Engine version                     4.20.0-91 (build date Feb 27 2018)
 Sessions                           0
```

show services application-identification status (Logical Systems)

```
user@host> show services application-identification status
Application Identification
 Status                            Enabled
 Sessions under app detection      0
 Max TCP session packet memory     0
 Force packet plugin               Disabled
 Force stream plugin               Disabled
DPI Performance mode:              Enabled
 Statistics collection interval    1440 (in minutes)

Application System Cache
 Status                            Enabled
Cache lookup security-services     Enabled
Cache lookup miscellaneous-services Enabled
 Max Number of entries in cache    131072
 Cache timeout                     3600 (in seconds)

Protocol Bundle
 Download Server                   https://signatures.juniper.net/cgi-bin/index.cgi
 AutoUpdate                        Disabled
Proxy Details
Proxy Profile                      Not Configured
Slot 1:
 Application package version       534
 Status                            Active
 PB Version                        1.270.0-48.005 (build date May 22 2017)
 Engine version                    4.20.0-49.005 (build date May 22 2017)
 Sessions                          0
```

show services application-identification status (Micro-Applications)

```
user@host> show services application-identification status
Application Identification
 Status                            Enabled
 Sessions under app detection      0
 Max TCP session packet memory     0
 Force packet plugin               Disabled
 Force stream plugin               Disabled
```

```
 Statistics collection interval    1440 (in minutes)

Application System Cache
 Status                            Enabled
  Cache lookup security-services   Disabled
  Cache lookup miscellaneous-services Disabled
 Max Number of entries in cache    0
 Cache timeout                     3600 (in seconds)

Protocol Bundle
 Download Server                   https://signatures.juniper.net/cgi-bin/index.cgi
 AutoUpdate                        Disabled

Proxy Details
 Proxy Profile                     Not Configured
Slot 1:
 Application package version       3172
 Status                            Active
 PB Version                        1.380.0-64.005 (build date May 13 2019)
 Engine version                    5.3.0-56 (build date May 13 2019)
 Micro-App Version                 1.0.0-0
 Sessions                          0
```

show services application-identification status (Rollback Status)

```
user@host> show services application-identification status
pic: 0/1

Application Identification
Status Enabled
Sessions under app detection 0
Max TCP session packet memory 0
Force packet plugin Disabled
Force stream plugin Disabled
Statistics collection interval 1440 (in minutes)

Application System Cache
Status Disabled
Cache lookup security-services Disabled
Cache lookup miscellaneous-services Enabled
Max Number of entries in cache 262144
```

```
Cache timeout 3600 (in seconds)

Protocol Bundle Download Server https://signatures.juniper.net/cgi-bin/index.cgi
AutoUpdate Disabled
Slot 1:
Application package version 3250
Status Active
PB Version 1.380.0-68.005 (build date Dec 9 2019)
Engine version 4.20.0-107 (build date Dec 9 2019)
Sessions 0
Slot 2:
Application package version 0
Status Free
PB Version N/A
Engine version 0
Sessions 0
Rollback version details
Application package version 3240
PB Version 1.380.0-65.005 (build date Jan 13 2020)
Engine version 5.3.0-56 (build date Jan 13 2020)
Micro-App Version 1.0.0-0
```

show services application-identification status (Junos OS Release 20.1 or Later)

```
Application Identification
Status                            Enabled
Sessions under app detection      0
Max TCP session packet memory     2097152
Force packet plugin               Disabled
Force stream plugin               Disabled
Statistics collection interval    1440 (in minutes)

Application System Cache
Status                            Enabled
  Cache lookup security-services   Disabled
  Cache lookup miscellaneous-services Enabled
Max Number of entries in cache    131072
Cache timeout                     3600 (in seconds)

Protocol Bundle
Download Server                   https://signatures.juniper.net/cgi-bin/index.cgi
```

```
AutoUpdate                        Enabled
Interval                          24 (in hours)

Proxy Details
Proxy Profile                     Not Configured
Slot 1:
Application package version       3437
Release date                      Thu Nov 18 14:52:06 2021 UTC
Status                            Active
PB Version                        1.530.2-40 (build date Sep 30 2021)
Engine version                    5.6.0-62 (build date Sep 30 2021)
Micro-App Version                 1.1.0-0
Sessions                          0
Custom-App Infra Version          1.0.0-0
```

show services application-identification status (Junos OS Release 20.3R1 or Later)

```
Application Identification
Status                            Enabled
Sessions under app detection      0
Max TCP session packet memory     2097152
Force packet plugin               Disabled
Force stream plugin               Disabled
Statistics collection interval    1440 (in minutes)

Application System Cache
Status                            Enabled
  Cache lookup security-services   Disabled
  Cache lookup miscellaneous-services Enabled
Max Number of entries in cache    131072
Cache timeout                     3600 (in seconds)

Protocol Bundle
Download Server                   https://signatures.juniper.net/cgi-bin/index.cgi
AutoUpdate                        Enabled
Interval                          24 (in hours)

Proxy Details
Proxy Profile                     Not Configured
Slot 1:
Application package version       3437
```

```
Release date                      Thu Nov 18 14:52:06 2021 UTC
Status                            Active
PB Version                        1.530.2-40 (build date Sep 30 2021)
Engine version                    5.6.0-62 (build date Sep 30 2021)
Micro-App Version                 1.1.0-0
Sessions                          0
Custom-App Infra Version          1.0.0-0

Rollback version details:
Application package version       3437
PB Version                        1.530.2-40
Engine version                    5.6.0-62
Micro-App Version                 1.1.0-0
Custom-App Infra Version          1.0.0-0
```

show services application-identification status (Junos OS Release 24.4R1 and Later)

```
user@host> show services application-identification status
Application Identification
Status                            Enabled
Sessions under app detection      0
Max TCP session packet memory     2097152
Force packet plugin               Disabled
Force stream plugin               Disabled
Statistics collection interval    1440 (in minutes)

Application System Cache
Status                            Enabled
  Cache lookup security-services   Disabled
  Cache lookup miscellaneous-services Enabled
Max Number of entries in cache    131072
Cache timeout                     3600 (in seconds)

Protocol Bundle
Download Server                   https://signatures.juniper.net/cgi-bin/index.cgi
AutoUpdate                        Disabled

Proxy Details
Proxy Profile                     Not Configured
Slot 1:
```

```
Application package version       3533 (Major)
Release date                      Mon Oct 10 14:55:29 2022 UTC
Status                            Active
PB Version                        1.550.2-31 (build date Oct 10 2022)
Engine version                    5.7.0-47 (build date Mar 25 2022)
Micro-App Version                 1.1.0-0
Sessions                          0
Custom-App Infra Version          1.0.0-0
Rollback version details:
Application package version       3662 (Major)
PB Version                        1.550.2-43
Engine version                    5.7.1-47
Micro-App Version                 1.2.0-1
Custom-App Infra Version          1.0.0-1
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

Command introduced in Junos OS Release 18.3R1 for logical systems.

Statement introduced in Junos OS Release 17.1 on MX Series.

RELATED DOCUMENTATION

- request services application-identification application | 14916
- No Link Title
- No Link Title

show services application-identification version

IN THIS SECTION

- Syntax | 25470
- Description | 25470
- Required Privilege Level | 25470

- Sample Output | **25470**
- show services application-identification version (Junos OS Release 24.2R1) | **25471**
- Release Information | **25471**

Syntax

```
show services application-identification version
```

Description

Displays the application signature package version installed on your security device.

Required Privilege Level

view

Sample Output

show services application-identification version

The following output shows that the application package version is 1608.

```
user@host> show services application-identification version
Application package version: 1608
```

show services application-identification version (Logical Systems)

The following output shows that the application package version is 534.

```
user@host> show services application-identification version
Application package version: 534
```

show services application-identification version (Junos OS Release 21.1R1)

The following output shows that the application package version is 3345 and release date as 12th January, 2021.

```
user@host> show services application-identification version
Application package version: 3345
Release date: Tue Jan 12 14:56:26 2021 UTC
```

show services application-identification version (Junos OS Release 24.2R1)

The following output displays the failed application package version due to validation errors.

```
user@host> show services application-identification version detail
  Application package version: 3654
  Release date: Thu Nov 23 14:07:48 2023 UTC

Dataplane validation failure version details:
  Application package version       3620
  PB Version                        1.550.2-43 (build date Apr  5 2023)
  Engine version                    5.7.1-47 (build date Mar 30 2023)
```

show services application-identification version (Junos OS Release 24.4R1)

The following output displays the major version of application package.

```
user@host> show services application-identification version
Application package version: 3666 (Major)
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

No Link Title

No Link Title

show services application-identification group obsolete-groups

IN THIS SECTION

- Syntax | 25472
- Description | 25472
- Required Privilege Level | 25472
- Sample Output | 25472
- Release Information | 25473

Syntax

```
show services application-identification group obsolete-groups
```

Description

Displays the application groups that are deprecated in the installed application signature package version on your security device.

Required Privilege Level

View

Sample Output

show services application-identification application obsolete-applications

```
user@host> show services application-identification group obsolete-groups
```

```
  Obsolete Application-groups
  junos:infrastructure:file-sharing
  junos:social-networking:applications
```

```
  junos:web:multimedia:adult
Obsolete application-group(s): 3
```

Release Information

Command introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

> *Predefined Application Signatures for Application Identification*

show services application-identification recent-appid-sigpack-versions

IN THIS SECTION

- Syntax | 25473
- Description | 25473
- Required Privilege Level | 25474
- Sample Output | 25474
- Release Information | 25474

Syntax

```
show services application-identification recent-appid-sigpack-versions
```

Description

Display the list of application signature packages available for downloading. The output displays both major and minor versions. To view the current application signature package version running on the system, use the show services application-identification version command.

Required Privilege Level

maintenance

Sample Output

show services application-identification recent-appid-sigpack-versions

```
user@host> show services application-identification recent-appid-sigpack-versions
appid sigpack version: 3642
appid sigpack version: 3615
appid sigpack version: 3604
appid sigpack version: 3533
appid sigpack version: 3470
appid sigpack version: 3429
appid sigpack version: 3405
appid sigpack version: 3390
appid sigpack version: 3372
appid sigpack version: 3351
```

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

No Link Title

show services application-identification recent-appid-sigpack-versions

IN THIS SECTION

Syntax | 25475

- Description | 25475
- Required Privilege Level | 25475
- Sample Output | 25475
- Release Information | 25476

Syntax

```
show services application-identification recent-appid-sigpack-versions
```

Description

Display the list of application signature packages available for downloading. The output displays both major and minor versions. To view the current application signature package version running on the system, use the show services application-identification version command.

Required Privilege Level

maintenance

Sample Output

show services application-identification recent-appid-sigpack-versions

```
user@host> show services application-identification recent-appid-sigpack-versions
appid sigpack version: 3642
appid sigpack version: 3615
appid sigpack version: 3604
appid sigpack version: 3533
appid sigpack version: 3470
appid sigpack version: 3429
appid sigpack version: 3405
appid sigpack version: 3390
appid sigpack version: 3372
appid sigpack version: 3351
```

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

No Link Title

No Link Title

show services captive-portal-content-delivery

IN THIS SECTION

- Syntax | 25476
- Description | 25477
- Options | 25477
- Required Privilege Level | 25477
- Output Fields | 25477
- Sample Output | 25479
- Release Information | 25482

Syntax

```
show services captive-portal-content-delivery
<pic pic-name>
<profile profile-name>
<rule rule-name> <term term-name>
<ruleset ruleset-name>
<sset sset-name> <brief> <detail> <summary>
<statistics <interface pic-name>>
```

Description

Display the current operational state of all captive portal interfaces.

Options

brief (Optional) Display brief service set database information.

detail (Optional) Display detailed service set database information.

pic Display the PIC database.

profile Display the profile database.

rule Display the rule database.

ruleset Display the rule set database.

sset Display the service set database.

statistics Display captive portal content delivery statistics about a PIC.

summary (Optional) Display a summary of service set database information.

term (Optional) Display term information for the rule database.

Required Privilege Level

view

Output Fields

Table 1392 on page 25477 lists the output fields for the `show services captive-portal-content-delivery` command. Output fields are listed in the approximate order in which they appear.

Table 1392: show services captive-portal-content-delivery Output Fields

Field Name	Field Description	Level of Output
`Name`	Name of the interface.	none

Table 1392: show services captive-portal-content-delivery Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Index`		none
`Profile`	Name of the service profile for the HTTP redirect services that contains the rules or rule sets specifying the service.	none
`Rules or Rule Sets`	List of rules or rule sets contained in the HTTP redirect service profile.	none
`Rule Name`	Name of an HTTP redirect service rule.	none
`Term Name`	Name of a rule term.	none
`Rule match direction`	Traffic direction on the interface where the rule match is applied, `input`, `output`, or `input-output`.	none
`Term action`	Action performed on packets when the rule term is matched: • Accept the packets. • Redirect the packets to a new destination URL. • Rewrite the packets with a new destination address and optionally a new destination port. • Log information about the packet to a system log file.	none
`Term action option`	Additional information related to the term action. • A new destination URL for a redirect action. • A new destination address for a rewrite action. • A new destination port for a rewrite action.	none
`Service Sets`	Name of service sets contained in a profile.	none

Table 1392: show services captive-portal-content-delivery Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Id`	Identifier number for a service set.	none
`Compiled Rules`		none
`service-set interface`	Interface on which the service set rules are applied.	none
`Packets received`	Number of packets received on the service-set interface.	none
`Packets altered`	Number of packets redirected or rewritten on the service-set interface.	none
`Packets dropped`	Number of packets dropped on the interface.	`detail`
`Received`	Number of packets received for the listed action: `Redirect`, `Rewrite`, or `Insert`.	`detail`
`Altered`	Number of packets altered by the listed action: `Redirect` or `Rewrite`.	`detail`
`Redirected`	Number of packets redirected by the `Insert` action.	`detail`
`Tag-inserted`	Number of packets for which one or more tags has been inserted by the `Insert` action.	`detail`

Sample Output

show services captive-portal-content-delivery

```
user@host> show services captive-portal-content-delivery pic si-5/0/0
Name          Index
si-5/0/0         20
```

show services captive-portal-content-delivery (Profile)

```
user@host> show services captive-portal-content-delivery profile
Profile         Rules or Rule Sets
http-redirect   1
ipda-rewrite    1
```

show services captive-portal-content-delivery (Profile HTTP Redirect)

```
user@host> show services captive-portal-content-delivery profile http-redirect
Profile         Rules or Rule Sets
http-redirect   1
```

show services captive-portal-content-delivery (Profile IPDA Rewrite)

```
user@host> show services captive-portal-content-delivery profile ipda-rewrite
Profile         Rules or Rule Sets
ipda-rewrite    1
```

show services captive-portal-content-delivery (Rules)

```
user@host> show services captive-portal-content-delivery rule
Rule Name          Term Name
redirect           t2
rewrite            t1
```

show services captive-portal-content-delivery (Rewrite Term)

```
user@host> show services captive-portal-content-delivery rule rewrite term t1
Rule name: rewrite
Rule match direction: input
Term name: t1
```

```
Term action: rewrite
Term action option: null
```

show services captive-portal-content-delivery (Redirect Term)

```
user@host> show services captive-portal-content-delivery rule redirect term t2
Rule name: redirect
Rule match direction: input
Term name: t2
Term action: redirect
Term action option: http://www.example.net
```

show services captive-portal-content-delivery (Service Set Detail)

```
user@host> show services captive-portal-content-delivery sset sset1 detail
Service Set        Id       Profile         Compiled Rules
sset1              1        ipda-rewrite    1
```

show services captive-portal-content-delivery (Interface)

```
user@host> show services captive-portal-content-delivery statistics interface sis-5/0/0
service-set interface: si-5/0/0

Packets received   Packets altered
5                     3
```

show services captive-portal-content-delivery (Subscriber with Insert Action Detail)

```
user@host> show services captive-portal-content-delivery statistics interface si-5/0/0 detail
service-set interface: si-5/0/0

Packets dropped:  10
```

```
Redirect:
Received           Altered
0                   0

Rewrite:
Received           Altered
0                    0

Insert:
Received           Redirected    Tag-inserted
2                    1               1
```

show services captive-portal-content-delivery (Subscriber with Redirect Action Detail)

```
user@host> show services captive-portal-content-delivery statistics interface si-4/0/0 detail
service-set interface: si-4/0/0

Packets dropped:  10

Redirect:
Received           Altered
4                   4

Rewrite:
Received           Altered
0                    0

Insert:
Received           Redirected    Tag-inserted
0                    0               0
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

clear services captive-portal-content-delivery statistics | 4085

show services cos statistics

IN THIS SECTION

- Syntax | 25483
- Description | 25483
- Options | 25484
- Required Privilege Level | 25484
- Output Fields | 25484
- Sample Output | 25485
- Release Information | 25487

Syntax

```
show services cos statistics
 <brief | detail | extensive>
<diffserv | forwarding-class>
<interface interface-name>
<service-set service-set-name>
<summary>
```

Description

Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns and the mapping of forwarding class names to queue numbers as configured in CoS services for the Multiservices PIC, MS-MIC, or MS-MPC.

Options

none	Display all services CoS statistics.
brief | detail | extensive	(Optional) Display the specified level of output.
diffserv | forwarding-class	(Optional) Display only the selected information, either DiffServ codepoints or forwarding classes.
interface *interface-name*	(Optional) Display statistics for the specified interface only.
service-set *service-set-name*	(Optional) Display statistics for the specified service set only.
summary	(Optional) Display summary of statistics on a per-interface basis.

Required Privilege Level

view

Output Fields

Table 1393 on page 25484 describes the output fields for the `show services cos statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1393: show services cos statistics Output Fields

Field Name	Field Description	Level of Output
Interface	Name of interface.	All levels
Service set	Name of service set.	All levels
DSCP	DiffServ code point bit pattern.	All levels
Packets in	Number of packets received.	All levels
Packets out	Number of packets transmitted.	All levels

Table 1393: show services cos statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Forwarding class	Forwarding class queue number.	All levels

Sample Output

show services cos statistics

```
user@host> show services cos statistics
Interface: sp-1/0/0, Service set: scos
DSCP               Packets in           Packets out
000000                      0                     0
000001                      0                     0
000010                      0                     0
000011                      0                     0
000100                      0                     0
000101                      0                     0
000110                      0                     0
000111                      0                     0
001000                      0                     0
001001                      0                     0
001010                      0                     0
001011                      0                     0
001100                      0                     0
001101                      0                     0
001110                      0                     0
001111                      0                     0
010000                      0                     0
010001                      0                     0
010010                      0                     0
010011                      0                     0
010100                      0                     0
010101                      0                     0
010110                      0                     0
010111                      0                     0
011000                      0                     0
011001                      0                     0
```

```
011010                    0                    0
011011                    0                    0
011100                    0                    0
011101                    0                    0
011110                    0                    0
011111                    0                    0
100000                    0                    0
100001                    0                    0
100010                    0                    0
100011                    0                    0
100100                    0                    0
100101                    0                    0
100110                    0                    0
100111                    0                    0
101000                    0                    0
101001                    0                    0
101010                    0                    0
101011                    0                    0
101100                    0                    0
101101                    0                    0
101110                    0                    0
101111                    0                    0
110000                    0                    0
110001                    0                    0
110010                    0                    0
110011                    0                    0
110100                    0                    0
110101                    0                    0
110110                    0                    0
110111                    0                    0
111000                    0                    0
111001                    0                    0
111010                    0                    0
111011                    0                    0
111100                    0                    0
111101                    0                    0
111110                    0                    0
111111                    0                    0
Forwarding class          Packets in          Packets out
0                                  0                    0
1                                  0                    0
2                                  0                    0
3                                  0                    0
```

```
4                                   0                   0
5                                   0                   0
6                                   0                   0
7                                   0                   0
8                                   0                   0
9                                   0                   0
10                                  0                   0
11                                  0                   0
12                                  0                   0
13                                  0                   0
14                                  0                   0
15                                  0                   0
```

show services cos statistics brief

The output for the `show services cos statistics brief` command is identical to that for the `show services cos statistics` command.

show services cos statistics detail

The output for the `show services cos statistics detail` command is identical to that for the `show services cos statistics` command.

show services cos statistics extensive

The output for the `show services cos statistics extensive` command is identical to that for the `show services cos statistics` command.

Release Information

Command introduced in Junos OS Release 8.1.

show services cos statistics (Next Gen Services)

IN THIS SECTION

Syntax | 25488

- Description | 25488
- Options | 25488
- Required Privilege Level | 25489
- Output Fields | 25489
- Sample Output | 25489
- Release Information | 25492

Syntax

```
show services cos statistics
 <brief | detail | extensive>
<diffserv | forwarding-class>
<interface interface-name>
<service-set service-set-name>
<summary>
```

Description

Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns and the mapping of forwarding class names to queue numbers as configured in CoS services for Next Gen Services services PICs.

Options

none	Display all services CoS statistics.
brief | detail | extensive	(Optional) Display the specified level of output.
diffserv | forwarding-class	(Optional) Display only the selected information, either DiffServ codepoints or forwarding classes.
interface *interface-name*	(Optional) Display statistics for the specified interface only.
service-set *service-set-name*	(Optional) Display statistics for the specified service set only.
summary	(Optional) Display summary of statistics on a per-interface basis.

Required Privilege Level

view

Output Fields

Table 1394 on page 25489 describes the output fields for the `show services cos statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1394: show services cos statistics Output Fields

Field Name	Field Description	Level of Output
Interface	Name of interface.	All levels
Service set	Name of service set.	All levels
DSCP	DiffServ code point bit pattern.	All levels
Packets in	Number of packets received.	All levels
Packets out	Number of packets transmitted.	All levels
Forwarding class	Forwarding class queue number.	All levels

Sample Output

show services cos statistics

```
user@host> show services cos statistics details
Interface: vms-0/2/0, Service set: ss1
DSCP                Packets in           Packets out
DSCP               Packets in           Packets out
000000                      0                     0
000001                      0                     0
000010                      0                     0
```

000011	0	0
000100	0	0
000101	0	0
000110	0	0
000111	0	0
001000	0	0
001001	0	0
001010	0	0
001011	0	0
001100	0	0
001101	0	0
001110	0	0
001111	0	0
010000	0	0
010001	0	0
010010	0	0
010011	0	0
010100	0	0
010101	0	0
010110	0	0
010111	0	0
011000	0	0
011001	0	0
011010	0	0
011011	0	0
011100	0	0
011101	0	0
011110	0	0
011111	0	0
100000	0	0
100001	0	0
100010	0	0
100011	0	0
100100	0	0
100101	0	0
100110	0	0
100111	0	0
101000	0	0
101001	0	0
101010	0	0
101011	0	0
101100	0	0
101101	0	0

```
101110                    0                    0
101111                    0                    0
110000                    0                    0
110001                    0                    0
110010                    0                    0
110011                    0                    0
110100                    0                    0
110101                    0                    0
110110                    0                    0
110111                    0                    0
111000                    0                    0
111001                    0                    0
111010                    0                    0
111011                    0                    0
111100                    0                    0
111101                    0                    0
111110                    0                    0
111111                    0                    0
Forwarding class          Packets in          Packets out
0                                  0                    0
1                                  0                    0
2                                  0                    0
3                                  0                    0
4                                  0                    0
5                                  0                    0
6                                  0                    0
7                                  0                    0
8                                  0                    0
9                                  0                    0
10                                 0                    0
11                                 0                    0
12                                 0                    0
13                                 0                    0
14                                 0                    0
15                                 0                    0
```

show services cos statistics brief

The output for the `show services cos statistics brief` command is identical to that for the `show services cos statistics` command.

show services cos statistics detail

The output for the `show services cos statistics detail` command is identical to that for the `show services cos statistics` command.

show services cos statistics extensive

The output for the `show services cos statistics extensive` command is identical to that for the `show services cos statistics` command.

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

show services crtp

IN THIS SECTION

- Syntax | 25492
- Description | 25493
- Options | 25493
- Required Privilege Level | 25493
- Output Fields | 25493
- Sample Output | 25495
- Release Information | 25495

Syntax

```
show services crtp
<extensive>
<interface interface-name>
```

Description

Display Compressed Real-Time Transport Protocol (CRTP) extensive output.

Options

none Display CRTP extensive output for all interfaces.

extensive (Optional) Display extensive CRTP information.

interface *interface-name* (Optional) Display CRTP flow statistics for the specified interface. On M Series and T Series routers, a link services IQ (**lsq-***fpc/pic/port*) or redundant link services IQ (**rlsq-***fpc/pic/port*) interface.

Required Privilege Level

view

Output Fields

Table 1395 on page 25493 lists the output fields for the `show services crtp` command. Output fields are listed in the approximate order in which they appear.

Table 1395: show services crtp Output Fields

Field Name	Field Description
Interface	Name of the physical interface.
Port minimum **Port maximum**	Compression is applied to UDP packets with even ports in the specified range.
Maximum UDP compressed sessions	Maximum value of a context identifier in the space of context identifiers allocated for UDP.
CRTP maximum period	Maximum interval between full headers. Suggested value is 256.

Table 1395: show services crtp Output Fields *(Continued)*

Field Name	Field Description
CRTP maximum time	Maximum time interval between full headers. Suggested value is 5 seconds.
Compression ratio	Ratio of received packet size to compressed packet size, in percentage. For example, if the packet size is 100 bytes when it is received, and is 40 bytes after compression, the compression ratio is 100 ? 40/100 * 100 = 60%.
Decompression ratio	Ratio of received packet size to decompressed packet size, in percentage. For example, if the packet size is 40 bytes when it is received, and is 100 bytes after compression, the decompression ratio is 100 ? 40/100 * 100 = 60%.
Discards	Number of frames that the incoming packet match code discarded because they were not recognized.
Sessions	Total number of active CRTP sessions.
IP bytes	Number of IP bytes sent and received.
Compressed bytes	Number of compressed IP header bytes sent and received.
CRTP packets	Number of CRTP packets sent and received.
CUDP/CNTCP packets	Number of compressed UDP packets and compressed non-TCP packets sent and received.
Full header packets	Number of full header packets sent and received. Full header packets communicate the uncompressed IP header plus any following headers and data to establish the uncompressed header state in the decompressor for a particular context.
Context state packet	Number of context state packets sent and received. Context state packets are sent from the decompressor to the compressor to communicate a list of context IDs for which synchronization is lost or might be lost.

Table 1395: show services crtp Output Fields *(Continued)*

Field Name	Field Description
IP packets	Number of IP packets sent and received.
Compressed packets	Number of compressed packets sent and received.

Sample Output

show services crtp extensive

```
user@host>                 show services crtp extensive
Interface: lsq-1/1/0.1
  Port minimum: 2000, Port maximum: 64009
  Maximum UDP compressed sessions: 256
  CRTP maximum period: 256, CRTP maximum time: 5
  Compression ratio: 0, Decompression ratio: 0, Discards: 0
  CRTP stats                       Receive         Transmit
  Sessions                               1                1
  IP bytes                              60               60
  Compressed bytes                      61               60
  CRTP packets                           0                0
  CUDP/CNTCP packets                     0                0
  Full header packets                    1                1
  Context state packets                  0                0
  IP packets                             1                1
  Compressed packets                     1                1
```

Release Information

Command introduced before Junos OS Release 7.4.

show services crtp flows

IN THIS SECTION

- Syntax | 25496
- Description | 25496
- Options | 25496
- Required Privilege Level | 25496
- Output Fields | 25497
- Sample Output | 25497
- Release Information | 25498

Syntax

```
show services crtp flows
<interface interface-name>
```

Description

Display Compressed Real-Time Transport Protocol (CRTP) flows.

Options

none	Display CRTP flows for all interfaces.
interface *interface-name*	(Optional) Display CRTP flows for the specified interface. On M Series and T Series routers, a link services IQ (**lsq-***fpc/pic/port*) or redundant link services IQ (**rlsq-***fpc/pic/port*) interface.

Required Privilege Level

view

Output Fields

Table 1396 on page 25497 lists the output fields for the `show services crtp flows` command. Output fields are listed in the approximate order in which they appear.

Table 1396: show services crtp flows Output Fields

Field Name	Field Description
Interface	Name of the physical interface.
Flow	Received or transmitted flow.
Source	IP source address.
Destination	IP destination address.
SSRC ID	Synchronization source (SSRC) identifier. One of the fields in the RTP header used to select the context. The SSRC identifier is a randomly chosen value unique within a particular CRTP session.
Ctx ID	Session context ID. Indicates the session context in which to interpret the packet. The decompressor can use the context ID to index its table of stored session contexts directly.

Sample Output

show services crtp flows

```
user@host>                 show services crtp flows
Interface: lsq-1/1/0.1
  Flow           Source                Destination          SSRC ID  Ctx ID
  Receive        192.0.2.3:28004       198.51.100.3:26000       123       0
  Transmit       198.51.100.3:26000    192.0.2.3:28004          123       2
```

Release Information

Command introduced before Junos OS Release 7.4.

show services dns-filtering cache

IN THIS SECTION

- Syntax | 25498
- Description | 25499
- Required Privilege Level | 25499
- Output Fields | 25499
- Sample Output | 25499
- Release Information | 25502

Syntax

```
show services dns-filtering cache
show services dns-filtering cache all-logical-systems-tenants
show services dns-filtering cache logical-system all
show services dns-filtering cache logical-system <logical-system-name>
show services dns-filtering cache tenant all
show services dns-filtering cache tenant <tenant-name>
show services dns-filtering cache summary
show services dns-filtering cache summary all-logical-systems-tenants
show services dns-filtering cache summary logical-system <logical-system-name>
show services dns-filtering cache summary tenant <tenant-name>
show services dns-filtering cache benign
show services dns-filtering cache benign all-logical-systems-tenants
show services dns-filtering cache benign logical-system <logical-system-name>
show services dns-filtering cache benign tenant <tenant-name>
show services dns-filtering cache c2
show services dns-filtering cache c2 all-logical-systems-tenants
```

```
show services dns-filtering cache c2 logical-system <logical-system-name>
show services dns-filtering cache c2 tenant <tenant-name>
```

Description

Show all entries within the DNS cache.

Required Privilege Level

view

Output Fields

Table 1397 on page 25499 lists the output fields for the `show services dns-filtering cache` command.

Table 1397: show services dns-filtering cache Output Fields

Field Name	Field Description
Domain Name	List of domain names (with or without wildcard).
TTL	The time-to-live (TTL) value in seconds.
Prevalence	The prevalence value 0 or 1. 0 being unique to your realm, 1 being seen elsewhere.
Verdict	Provides the verdict for ATP Cloud provided DGA verdict and tunnel detected domains. The available options are benign and c2.
Hitcount	Displays how frequently each domain entry is hit.

Sample Output

show services dns-filtering cache

```
user@host> show services dns-filtering cache
Logical System:root-logical-system
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC0
```

```
Domain-Name, TTL, Prevalence , Verdict, Hitcount
*.google.com,, 172800, 0, Benign, 1
*.xyz.com,, 172800, 0, C2, 1
*.fb.com, 172800, 0, Benign, 1
*.abc.com, 172800, 0, C2, 1
```

show services dns-filtering cache logical-system lsys1

```
user@host> show services dns-filtering cache logical-system lsys1
Logical System:lsys1
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC0
Domain-Name, TTL, Prevalence , Verdict, Hitcount
a1.com, 480, 1, Benign, 1
a3.com, 480, 1, C2, 1

Logical System:lsys1
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC1
Domain-Name, TTL, Prevalence , Verdict, Hitcount
a11.com, 480, 1, Benign, 1
a13.com, 480, 1, C2, 1
```

show services dns-filtering cache all-logical-systems-tenants

```
user@host> show services dns-filtering cache all-logical-systems-tenants
Logical System:root-logical-system
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC0
Domain-Name, TTL, Prevalence , Verdict, Hitcount
a0.com, 480, 1, Benign, 1
a4.com, 480, 1, C2, 1

Logical System:lsys1
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC0
Domain-Name, TTL, Prevalence , Verdict, Hitcount
a1.com, 480, 1, Benign, 1
a3.com, 480, 1, C2, 1

Logical System:t1
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC0
Domain-Name, TTL, Prevalence , Verdict, Hitcount
a2.com, 480, 1, Benign, 1
```

```
a5.com, 480, 1, C2, 1

Logical System:root-logical-system
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC1
Domain-Name, TTL, Prevalence , Verdict, Hitcount
a10.com, 480, 1, Benign, 1
a14.com, 480, 1, C2, 1


Logical System:lsys1
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC1
Domain-Name, TTL, Prevalence , Verdict, Hitcount
a11.com, 480, 1, Benign, 1
a13.com, 480, 1, C2, 1

Logical System:t1
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC1
Domain-Name, TTL, Prevalence , Verdict, Hitcount
a12.com, 480, 1, Benign, 1
a15.com, 480, 1, C2, 1
```

show services dns-filtering cache summary

```
show services dns-filtering cache summary
Logical System:root-logical-system
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC0
Total Entries in Cache:4
Total Benign Entries in Cache:2
Total C2 Entries in Cache:2
```

show services dns-filtering cache benign

```
show services dns-filtering cache benign
Logical System:root-logical-system
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC0
Domain-Name, TTL, Prevalence , Verdict, Hitcount
*.abc.com, 172800, 0, BENIGN, 1
*.gavin.com, 172800, 0, BENIGN, 1
```

show services dns-filtering cache c2

```
show services dns-filtering cache c2
Logical System:root-logical-system
DNS Cache Refresh Rate:5 Minutes on FPC0 PIC0
Domain-Name, TTL, Prevalence , Verdict, Hitcount
*.fb.com, 172800, 0, C2, 1
*.youtube.com, 172800, 0, C2, 1
```

Release Information

dns-filtering cache command introduced in Junos OS Release 21.2R1

show services dns-filtering cache summary, show services dns-filtering cache c2 and show services dns-filtering cache benign commands were introduced in Junos OS Release 22.2R1

show services dynamic-flow-capture content-destination

IN THIS SECTION

- Syntax | 25502
- Description | 25503
- Options | 25503
- Required Privilege Level | 25503
- Output Fields | 25503
- Sample Output | 25504
- Release Information | 25504

Syntax

```
show services dynamic-flow-capture content-destination capture-group group-name  destination-
identifier identifier
<terse>
```

Description

(M320 Series routers and T Series routers only) Display information about the content destination that receives packets from the dynamic flow capture (DFC) interface.

Options

capture-group ***group-name***	Display information for the specified capture-group identifier.
destination-identifier ***identifier***	Display information for the specified content destination identifier.
terse	(Optional) Display summary information.

Required Privilege Level

view

Output Fields

Table 1398 on page 25503 lists the output fields for the `show services dynamic-flow-capture content-destination` command. Output fields are listed in the approximate order in which they appear.

Table 1398: show services dynamic-flow-capture content-destination Output Fields

Output Field	Output Field Description
Capture group	Name of the capture group.
Content destination	Name of the content destination.
Criteria	Number of criteria specified.
Bandwidth	Bandwidth used by the matched traffic.
Matched packets	Number of matched packets sent to the content destination.
Matched bytes	Number of matched bytes sent to the content destination.

Table 1398: show services dynamic-flow-capture content-destination Output Fields *(Continued)*

Output Field	Output Field Description
Congestion notifications	Number of notification messages sent.

Sample Output

show services dynamic-flow-capture content-destination capture-group

```
user@host> show services dynamic-flow-capture content-destination capture-group g1 destination-identifier cd1 terse
  Capture group: g1, Content destination: cd1, Criteria: 0, Bandwidth: 0, Matched packets: 0, Matched bytes: 0, Congestion notifications: 0
```

Release Information

Command introduced in Junos OS Release 7.4.

show services dynamic-flow-capture control-source

IN THIS SECTION

- Syntax | 25505
- Description | 25505
- Options | 25505
- Required Privilege Level | 25505
- Output Fields | 25505
- Sample Output | 25506
- Release Information | 25507

Syntax

```
show services dynamic-flow-capture control-source capture-group group-name   control-source
source-identifier identifier
<detail | terse>
```

Description

(M320 Series routers and T Series routers only) Display information about the control source that makes dynamic flow capture requests to the dynamic flow capture interface.

Options

capture-group *group-name*	Capture group identifier.
source-identifier *identifier*	Control source identifier.
detail | terse	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1399 on page 25505 lists the output fields for the `show services dynamic-flow-capture control-source` scommand. Output fields are listed in the approximate order in which they appear.

Table 1399: show services dynamic-flow-capture control-source Output Fields

Output Field	Output Field Description
Capture group	Name of the capture group.
Control source	Name of the control source.
Criteria added, Criteria add failed	Number of criteria added or added and failed.

Table 1399: show services dynamic-flow-capture control-source Output Fields *(Continued)*

Output Field	Output Field Description
Active criteria	Number of active criteria.
Static criteria, Dynamic criteria	Number of static or dynamic criteria.
Control protocol requests	Total number of control protocol requests.
Requests	Number of **Add**, **Delete**, **List**, **Refresh**, and **No-op** control protocol requests.
Failed	Number of **Add**, **Delete**, **List**, **Refresh**, and **No-op** failed control protocol requests.
Add request rate	Rate of add requests.
Add request peak rate	Peak rate of add requests.
Bandwidth across all criteria	Bandwidth used by all the requests.
Total notifications	Total number of notifications sent and the number of notifications by category: **Restart**, **Rollover**, **Timeout**, **Congestion**, **Congestion delete**, and **Dups** (duplicates) dropped.
Criteria deleted	Total number of criteria deleted and the number of deleted criteria by category: **Timeout idle**, **Timeout total**, **Packets**, and **Bytes**.
Sequence number	Sequence number.

Sample Output

show services dynamic-flow-capture control-source source-identifier capture-group

```
user@host> show services dynamic-flow-capture control-source source-identifier cs0_cg0 capture-
group cg_0
```

```
Capture group: cg_0, Control source: cs0_cg0
    Criteria added: 28, Criteria add failed: 0, Active criteria: 0, Control protocol requests:
28,     Add request rate: 0,
    Add request peak rate: 1, Bandwidth across all criteria: 0, Total notifications: 1, Criteria
deleted: 28,    Sequence number: 0
```

show services dynamic-flow-capture control-source ource-identifier capture-group detail

```
user@host> show services dynamic-flow-capture control-source source-identifier cs0_cg0 capture-
group cg_0 detail
Capture group: cg_0, Control source: cs0_cg0
  Criteria added: 28, Criteria add failed: 0
  Active criteria: 0
    Static criteria: 0, Dynamic criteria: 0
  Control protocol requests: 28
                     Add        Delete          List       Refresh         No-op
    Requests          28             0             0             0             0
    Failed             0             0             0             0             0
  Add request rate: 0
  Add request peak rate: 1
  Bandwidth across all criteria: 0
  Total notifications: 1
    Restart: 1, Rollover: 0, No-op: 0, Timeout: 0, Congestion: 0, Congestion delete: 0, Dups
dropped: 0
  Criteria deleted: 28
    Timeout idle: 0, Timeout total: 0, Packets: 0, Bytes: 0
  Sequence number: 0
```

Release Information

Command introduced in Junos OS Release 7.4.

show services dynamic-flow-capture statistics

IN THIS SECTION

- Syntax | 25508
- Description | 25508
- Options | 25508
- Required Privilege Level | 25508
- Output Fields | 25508
- Sample Output | 25511
- Release Information | 25512

Syntax

```
show services dynamic-flow-capture statistics capture-group group-name
```

Description

(M320 Series routers and T Series routers only) Display statistics information about the capture group specified for dynamic flow capture.

Options

capture-group *group-name*	Display information for the specified capture group identifier.

Required Privilege Level

view

Output Fields

Table 1400 on page 25509 lists the output fields for the `show services dynamic-flow-capture statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1400: show services dynamic-flow-capture statistics Output Fields

Output Field	Output Field Description
Input	Incoming dynamic flow capture packet statistics: • **Control protocol packets**—Number of control protocol packets received. • **Captured data packets**—Number of data packets captured. • **Control IRI packets**—Number of control IRI packets received.
Control protocol drops	Control protocol packets dropped for the following reasons: • **Not IP packets**—Dropped packets were not IP packets. • **Not UDP packets**—Dropped packets were not User Datagram Protocol (UDP) packets. • **Invalid destination address**—Dropped packets had invalid destination addresses. • **No memory**—Packets dropped because of insufficient memory. • **Unauthorized control source**—Packets dropped because the control source was not authenticated. • **Bad request**—Packets dropped because the request was invalid. • **Unknown control source**—Packets dropped because the control source was not known. • **Not DTCP**—Dropped packets did not adhere to the control protocol format. • **Bad command line**—Packets dropped because of a version mismatch. • **Bandwidth exceeded**—Packets dropped because the bandwidth was exceeded. • **Drop rate due to exceeded bandwidth**—Rate of traffic dropped because the bandwidth was exceeded. • **Other**—Packets dropped for other reasons or undetermined causes.

Table 1400: show services dynamic-flow-capture statistics Output Fields *(Continued)*

Output Field	Output Field Description
Input drops	Incoming dynamic flow capture packets dropped for the following reasons: • **Unknown packets**—Packets dropped because the packet type was not recognized. • **Captured data not IPv4**—Packets dropped because they were not IPv4 packets. • **Captured data too small**—Packets dropped because they were smaller than the size reported in their headers. • **Captured data drops**—Data packets dropped because of undetermined causes. • **Captured data not matched**—Packets dropped because they did not match filter criteria. • **Bandwidth exceeded**—Packets dropped because the bandwidth was exceeded. • **Drop rate due to exceeded bandwidth**—Rate of traffic dropped because the bandwidth was exceeded.
Output	Outgoing dynamic flow capture packet statistics: • **Control protocol packets**—Number of control protocol packets sent. • **Captured data packets**—Number of captured data packets sent.
Output drops	Outgoing packets dropped: • **Control protocol drops**—Number of control protocol packets dropped. • **Captured data drops**—Number of captured data packets dropped.

Table 1400: show services dynamic-flow-capture statistics Output Fields *(Continued)*

Output Field	Output Field Description
Flow Statistics	DFC flow statistics: • **Active flow cache entries** • **Active flow cache usage percentage** • **Flow cache entries allocated** • **Number of control sources** • **Number of content destinations** • **Number of criteria** • **Maximum criteria matching one flow** • **Cached flows purged for memory** • **Maximum filters matching one packet**

Sample Output

show services dynamic-flow-capture statistics capture-group

```
user@host> show services dynamic-flow-capture statistics capture-group g1
Input:

  Control protocol packets: 643, Captured data packets: 69977, Control IRI packets: 337

Control protocol drops:

  Not IP packets: 0, Not UDP packets: 3, Invalid destination address: 0, No memory: 0,
Unauthorized control source: 0,

  Bad request: 0, Unknown control source: 0, Not DTCP: 0, Bad command line: 0, Bandwidth
exceeded: 0,

  Drop rate due to exceeded bandwidth: 0, Other: 0
```

```
Input drops:

  Unknown packets: 0, Captured data not IPv4: 0, Captured data too small: 0, Captured data
drops: 0, Captured data not matched: 0,

  Bandwidth exceeded: 0, Drop rate due to exceeded bandwidth: 0

Output:

  Control protocol packets: 644, Captured data packets: 1119624

Output drops:

  Control protocol drops: 0, Captured data drops: 0

Flow Statistics:

  Active flow cache entries: 40, Active flow cache usage percentage: 0, Flow cache entries
allocated: 40,

  Number of control sources: 4, Number of content destinations: 64, Number of criteria: 640,

  Maximum criteria matching one flow: 16, Cached flows purged for memory: 0, Maximum filters
matching one packet: 16
```

Release Information

Command introduced in Junos OS Release 7.4.

show services fixed-wireless-access statistics

IN THIS SECTION

- Syntax | 25513
- Description | 25513
- Required Privilege Level | 25513

- Output Fields | 25513
- Sample Output | 25514
- Release Information | 25515

Syntax

```
show services fixed-wireless-access statistics
```

Description

Display statistics for fixed wireless access messages. The statistics are collected on the primary Routing Engine because that is where the UDP port is maintained.

Required Privilege Level

view

Output Fields

Table 1401 on page 25514 lists the output fields for the `show system subscriber-management resiliency` command. Output fields are listed in the approximate order in which they appear.

Table 1401: show system subscriber-management resiliency Output

Field Name	Field Description
`GTPv2 Message`	Message type supported by the General Packet Radio Service (GPRS) Tunnelling Protocol, version 2. • `Create Session`—When the SAEGW receives a Create Session request message from the MME, it creates the session. The SAEGW then sends a Create Session response to the MME with the L-TEID-C and L-TEID-U values for the session. • `Modify Bearer`—When the SAEGW receives a Modify Bearer request message from the MME, it creates the dynamic pseudowire interface that receives the GTP-U encapsulated data packets from the UE by means of the eNodeB. When the dynamic pseudowire interface is created, the SAEGW sends a Modify Bearer response message to the MME. • `Delete Session`—When the SAEGW receives a Delete Session request message from the MME to delete the default bearer for a subscriber, the SAEGW initiates a subscriber logout. When the logout completes, the SAEGW sends a Delete Session response to the MME. • `Delete Bearer`—When the SAEGW receives a Delete Bearer request message from the MME or the eNodeB to delete a specific bearer context, the SAEGW deletes the specified context and sends a Delete Bearer response to the appropriate entity. • `Echo`—The Echo request/response messages are path management messages. Some networks use them as a keepalive mechanism to determine whether the path to the SAEGW is still alive and available. The SAEGW does not originate any Echo requests, but it does send an Echo response to any Echo request that it receives.
`Request Rcv`	Number of requests received on the BNG.
`Response Xmit`	Number of response transmitted by the BNG.

Sample Output

show services fixed-wireless-access statistics

```
user@host> show services fixed-wireless-access statistics

GTPv2 Message            Request Rcv       Response Xmit
```

```
Create Session            650            0
Modify Bearer             650            0
Delete Session            327            0
Delete Bearer               0            0
Error Indication            0            0
Echo                        0            0
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Fixed Wireless Access Networks

show services flow-collector file interface

IN THIS SECTION

- Syntax | 25515
- Description | 25516
- Options | 25516
- Additional Information | 25516
- Required Privilege Level | 25516
- Output Fields | 25516
- Sample Output | 25518
- Release Information | 25519

Syntax

```
show services flow-collector file interface (all | cp-fpc/pic/port)
<detail | extensive | terse>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display information about flow collector files.

Options

none	Display file information for all configured flow collector interfaces.
all | cp-*fpc/pic/port*	Display file information for all configured flow collector interfaces or for the specified interface.
detail | extensive | terse	(Optional) Display the specified level of output.

Additional Information

No entries are displayed for files that have been successfully transferred.

Required Privilege Level

view

Output Fields

Table 1402 on page 25516 lists the output fields for the `show services flow-collector file interface` command. Output fields are listed in the approximate order in which they appear.

Table 1402: show services flow-collector file interface Output Fields

Output Field	Output Field Description	Level of Output
Filename	Name of the file created on the flow collector interface.	All levels
Flows	Total number of collector flows for which records are present in the file.	none specified

Table 1402: show services flow-collector file interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output

Throughput	Throughput statistics:
	• **Flow records**—Number of flow records in the file.
	  • **per second**—Average number of flow records per second.
	  • **peak per second**—Peak number of flow records per second.
	• **Uncompressed bytes**—Total file size before compression.
	  • **per second**—Average number of uncompressed bytes per second.
	  • **peak per second**—Peak number of uncompressed bytes per second.
	• **Compressed bytes**—Total file size after compression.
	  • **per second**—Average number of compressed bytes per second.
	  • **peak per second**—Peak number of compressed bytes per second.
		extensive

Table 1402: show services flow-collector file interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output

Status	File statistics:
	• **Compressed blocks**—(**extensive** output only) Data blocks in the file that have been compressed. The file is exported only when the compressed block count and block count become the same.
	• **Block count**—(**extensive** output only) Total number of data blocks in the file.
	• **State**—Processing state of the file.
	  • **Active**—The flow collector interface is writing to the file.
	  • **Export 1**—File export is in progress to the primary server.
	  • **Export 2**—File export is in progress to the secondary server.
	  • **Wait**—File is pending export.
	  • **Transfer attempts 0**.—Number of attempts made to transfer the file. If the file is successfully transferred in the first attempt, this field is 0.
		All levels

Sample Output

show services flow-collector file interface extensive

```
user@host> show services flow-collector file interface cp-3/2/0 extensive
Filename: cFlowd-py69Ni69-0-20031112_014301-so_3_0_0_0.bcp.bi.gz
  Throughput:
    Flow records: 188365, per second: 238, peak per second: 287
    Uncompressed bytes: 21267756, per second: 27007, peak per second: 32526
    Compressed bytes: 2965643, per second: 0, peak per second: 22999
  Status:
    Compressed blocks: 156, Block count: 156
    State: Active, Transfer attempts: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

show services flow-collector input interface

IN THIS SECTION

- Syntax | 25519
- Description | 25519
- Options | 25519
- Required Privilege Level | 25520
- Output Fields | 25520
- Sample Output | 25520
- Release Information | 25521

Syntax

```
show services flow-collector input interface (all | cp-fpc/pic/port)
<detail | extensive | terse>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display the number of packets received by collector interfaces from monitoring interfaces.

Options

none	Display packets received by all configured flow collector interfaces.
all | cp-*fpc/pic/port*	Display packets received by all configured flow collector interfaces or by the specified interface.

detail | extensive | terse (Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1403 on page 25520 lists the output fields for the `show services flow-collector input interface` command. Output fields are listed in the approximate order in which they appear.

Table 1403: show services flow-collector input interface Output Fields

Output Field	Output Field Description
Interface	Name of the monitoring interface.
Packets	Number of packets traveling from the monitoring interface to the flow collector interface.
Bytes	Number of bytes traveling from the monitoring interface to the flow collector interface.

Sample Output

show services flow-collector input interface

```
user@host> show services flow-collector input interface cp-3/2/0
Interface                      Packets        Bytes
mo-3/0/0.0                       21706     32328568
mo-3/1/0.0                       21706     32329096
```

show services flow-collector input interface all

```
user@host> show services flow-collector input interface all
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
Interface                      Packets        Bytes
```

```
mo-3/0/0.0                          274       416232
mo-3/3/0.0                          274       416184
mo-1/0/0.0                          274       416232
mo-1/1/0.0                          274       416232
mo-1/2/0.0                          274       416232
mo-1/3/0.0                          274       416232
mo-3/1/0.0                          274       416232
mo-4/0/0.0                          274       416232
mo-4/1/0.0                          274       416232
mo-4/2/0.0                          274       416184
mo-4/3/0.0                          274       416232
mo-5/0/0.0                          274       416232
mo-5/1/0.0                          274       416232
mo-5/2/0.0                          274       416232
mo-5/3/0.0                          274       416232
mo-6/0/0.0                          274       416232

Flow collector interface: cp-6/3/0
Interface state: Collecting flows
```

Release Information

Command introduced before Junos OS Release 7.4.

show services flow-collector interface

IN THIS SECTION

- Syntax | 25522
- Description | 25522
- Options | 25522
- Required Privilege Level | 25522
- Output Fields | 25522
- Sample Output | 25527
- Release Information | 25532

Syntax

```
show services flow-collector interface (all | cp-fpc/pic/port)
<detail | extensive | terse>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display overall statistics for the flow collector application.

Options

none	Display statistics for flow collector applications on all interfaces.
all | cp-*fpc/pic/port*	Display statistics for flow collector applications on all interfaces or for the specified interface.
detail | extensive | terse	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1404 on page 25522 lists the output fields for the `show services flow-collector interface` command. Output fields are listed in the approximate order in which they appear.

Table 1404: show services flow-collector interface Output Fields

Output Field	Output Field Description	Level of Output
Flow collector interface	Name of the flow collector interface.	All levels
Interface state	Collecting flow state for the interface.	All levels

Table 1404: show services flow-collector interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output
Packets	Total number of packets received.	none specified
Flows Uncompressed Bytes	Total uncompressed data size for all files created on this PIC.	none specified
Compressed Bytes	Total compressed data size for all files created on this PIC.	none specified
FTP bytes	Total number of bytes transferred to the FTP server, including those dropped during transfer.	none specified
FTP files	Total number of FTP transfers attempted by the server.	none specified
Memory	Bytes used on the PIC and bytes free.	**detail extensive**

Input	Incoming flow collector packet statistics:
	• **Packets**—Number of packets received on the unit.
	  • **per second**—Average number of packets per second.
	  • **peak per second**—Peak number of packets per second.
	• **Bytes**—Number of bytes received on the unit.
	  • **per second**—Average number of bytes per second.
	  • **peak per second**—Peak number of bytes per second.
	• **Flow records processed**—Number of records in the flow collector packets that were processed by the flow-collector interface.
	  • **per second**—Average number of flow records processed per second.
	  • **peak per second**—Peak number of flow records per second.
		detail extensive

Table 1404: show services flow-collector interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output

Allocation	Data block statistics:
	• **Blocks allocated**—Total number of data blocks (containing flow records) allocated to the files created on this PIC.
	  • **per second**—Average number of blocks allocated per second.
	  • **peak per second**—Peak number of blocks allocated per second.
	• **Blocks freed**—Total number of data blocks freed.
	  • per second—Average number of blocks freed per second.
	  • peak per second—Peak number of blocks freed per second.
	• **Blocks unavailable**—Total number of data block requests denied, typically because of a memory shortage.
	  • per second—Average number of blocks unavailable per second.
	  • peak per second—Peak number of blocks unavailable per second.
		extensive

Files	File statistics, incremented since the PIC last booted: • **Files created**—Total number of files created on this PIC. • **Files exported**— Number of files successfully created and exported. • **Files destroyed**—(**extensive** output only) Number of files successfully exported and files dropped by the flow collection interface.	**detail extensive**

Table 1404: show services flow-collector interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output

Throughput	Throughput statistics:
	• **Uncompressed bytes**—Total uncompressed data size for all files created on this PIC.
	  • **per second**—Average number of uncompressed bytes per second.
	  • **peak per second**—Peak number of uncompressed bytes per second.
	• **Compressed bytes**—Total compressed data size for all files created on this PIC.
	  • **per second**—Average number of compressed bytes per second.
	  • **peak per second**—Peak number of compressed bytes per second.
		detail extensive

Packet drops	Number of packets dropped for the following causes: • **No memory**—Packets dropped because of insufficient memory. • **Not IP**—Packets dropped because they are not IP packets. • **Not IPv4**—Packets dropped because they are not IP version 4 packets. • **Too small**—Packets dropped because each packet was smaller than the size reported in its header. • **Fragments**—Packets dropped because of fragmentation. Fragments are not reassembled. • **ICMP**—Packets dropped because they are not ICMP packets. • **TCP**—Packets dropped because they are not TCP packets. • **Unknown**—Packets dropped because of undetermined causes. • **Not Junos flow**—Packets dropped because they are not interpreted by Junos OS. Junos OS interprets only IPv4, UDP cflowd version 5 packets.	**extensive**

Table 1404: show services flow-collector interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output
File transfer	File transfer statistics: • **FTP bytes**—Total number of bytes transferred to the FTP server, including those dropped during transfer. • **FTP files**—Total number of FTP transfers attempted by the server. • **FTP failure**—Total number of FTP failures encountered by the server.	**detail extensive**
Flow collector interface	Physical interface acting as a flow collector.	**detail**

Table 1404: show services flow-collector interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output

Export channel	Export channel **0** is unit 0. Export channel **1** is unit 1. Flow receive channel is unit 2. Server status statistics are the following:
	• **Current server Primary** or **Secondary**—Current FTP server being used. Value is
	• **Primary server state**—State of the server:
	  • **OK**—Server is operating without problems.
	  • **FTP error**—Server encountered an FTP protocol error while sending files.
	  • **Network error**—Flow-collector interface has errors when contacting the primary FTP server.
	  • **Unknown**—First file transfer has not been sent to the primary server.
	• **Secondary server state**—State of the server:
	  • **OK**—Server is operating without errors.
	  • **FTP error**—Server encountered an FTP protocol error while sending files.
	  • **Network error**—Flow-collector interface has errors when contacting the secondary FTP server.
	  • **Unknown**—First file transfer has not been sent to the secondary server.
	• **Not configured**—Secondary server is not configured.
		detail extensive

Sample Output

show services flow-collector interface all detail

```
user@host> show services flow-collector interface all detail
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
```

```
Memory:
    Used: 51452732, Free: 440329088
Input:
    Packets: 4384, per second: 0, peak per second: 156
    Bytes: 6659616, per second: 0, peak per second: 249695
    Flow records processed: 131070, per second: 0, peak per second: 4914
Files:
    Files created: 1, per second: 0, peak per second: 0
    Files exported: 1, per second: 0, peak per second: 0
Throughput:
    Uncompressed bytes: 13742307, per second: 0, peak per second: 593564
    Compressed bytes: 3786177, per second: 0, peak per second: 162826
File Transfer:
    FTP bytes: 3786247, per second: 0, peak per second: 378620
    FTP files: 1, per second: 0, peak per second: 0
    FTP failure: 0
Export channel: 0
    Current server: Primary
    Primary server state: OK, Secondary server state: OK
Export channel: 1
    Current server: Primary
    Primary server state: Unknown, Secondary server state: OK

Flow collector interface: cp-6/3/0
Interface state: Collecting flows
Memory:
    Used: 51452732, Free: 440329088
Input:
    Packets: 0, per second: 0, peak per second: 0
    Bytes: 0, per second: 0, peak per second: 0
    Flow records processed: 0, per second: 0, peak per second: 0
Files:
    Files created: 0, per second: 0, peak per second: 0
    Files exported: 0, per second: 0, peak per second: 0
Throughput:
    Uncompressed bytes: 0, per second: 0, peak per second: 0
    Compressed bytes: 0, per second: 0, peak per second: 0
File Transfer:
    FTP bytes: 70, per second: 0, peak per second: 6
    FTP files: 0, per second: 0, peak per second: 0
    FTP failure: 0
Export channel: 0
    Current server: Primary
```

```
    Primary server state: Unknown, Secondary server state: OK
Export channel: 1
    Current server: Primary
    Primary server state: Unknown, Secondary server state: OK
```

show services flow-collector interface all extensive

```
user@host> show services flow-collector interface all extensive
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
Memory:
    Used: 51452732, Free: 440329088
Input:
    Packets: 4384, per second: 0, peak per second: 156
    Bytes: 6659616, per second: 0, peak per second: 249695
    Flow records processed: 131070, per second: 0, peak per second: 4914
Allocation:
    Blocks allocated: 108, per second: 0, peak per second: 0
    Blocks freed: 108, per second: 0, peak per second: 10
    Blocks unavailable: 0, per second: 0, peak per second: 0
Files:
    Files created: 1, per second: 0, peak per second: 0
    Files exported: 1, per second: 0, peak per second: 0
    Files destroyed: 1, per second: 0, peak per second: 0
Throughput:
    Uncompressed bytes: 13742307, per second: 0, peak per second: 593564
    Compressed bytes: 3786177, per second: 0, peak per second: 162826
Packet drops:
    No memory: 0, Not IP: 0
    Not IPv4: 0, Too small: 0
    Fragments: 0, ICMP: 0
    TCP: 0, Unknown: 0
    Not JUNOS flow: 0
File Transfer:
    FTP bytes: 3786247, per second: 0, peak per second: 378620
    FTP files: 1, per second: 0, peak per second: 0
    FTP failure: 0
Export channel: 0
    Current server: Primary
    Primary server state: OK, Secondary server state: OK
Export channel: 1
```

```
    Current server: Primary
    Primary server state: Unknown, Secondary server state: OK

Flow collector interface: cp-6/3/0
Interface state: Collecting flows
Memory:
    Used: 51452732, Free: 440329088
Input:
    Packets: 0, per second: 0, peak per second: 0
    Bytes: 0, per second: 0, peak per second: 0
    Flow records processed: 0, per second: 0, peak per second: 0
Allocation:
    Blocks allocated: 0, per second: 0, peak per second: 0
    Blocks freed: 0, per second: 0, peak per second: 0
    Blocks unavailable: 0, per second: 0, peak per second: 0
Files:
    Files created: 0, per second: 0, peak per second: 0
    Files exported: 0, per second: 0, peak per second: 0
    Files destroyed: 0, per second: 0, peak per second: 0
Throughput:
    Uncompressed bytes: 0, per second: 0, peak per second: 0
    Compressed bytes: 0, per second: 0, peak per second: 0
Packet drops:
    No memory: 0, Not IP: 0
    Not IPv4: 0, Too small: 0
    Fragments: 0, ICMP: 0
    TCP: 0, Unknown: 0
    Not JUNOS flow: 0
File Transfer:
    FTP bytes: 70, per second: 0, peak per second: 6
    FTP files: 0, per second: 0, peak per second: 0
    FTP failure: 0
Export channel: 0
    Current server: Primary
    Primary server state: Unknown, Secondary server state: OK
Export channel: 1
    Current server: Primary
    Primary server state: Unknown, Secondary server state: OK
```

show services flow-collector interface all terse

```
user@host> show services flow-collector interface all terse
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
  Packets     Bytes     Flows Uncompressed  Compressed     FTP bytes FTP files
                                   Bytes       Bytes
     4384   6659616    131070     13742307     3786177       3786247         1

Flow collector interface: cp-6/3/0
Interface state: Collecting flows
  Packets     Bytes     Flows Uncompressed  Compressed     FTP bytes FTP files
                                   Bytes       Bytes
        0         0         0            0           0            70         0
```

show services flow-collector interface extensive

```
user@host> show services flow-collector interface cp-5/2/0 extensive
Flow collector interface: cp-5/2/0
Interface state: Collecting flows
Memory:
     Used: 458311860, Free: 40810008
Input:
     Packets: 922629, per second: 2069, peak per second: 3266
     Bytes: 1376559252, per second: 3096940, peak per second: 4880051
     Flow records processed: 25764957, per second: 42564, peak per second: 98124
Allocation:
     Blocks allocated: 20862, per second: 31, peak per second: 72
     Blocks freed: 17161, per second: 40, peak per second: 202
     Blocks unavailable: 58786, per second: 652, peak per second: 1120
Files:
     Files created: 52, per second: 0, peak per second: 0
     Files exported: 42, per second: 0, peak per second: 0
     Files destroyed: 42, per second: 0, peak per second: 0
Throughput:
     Uncompressed bytes: 2592070401, per second: 7297307,
     peak per second: 8630023
     Compressed bytes: 659600068, per second: 1858458, peak per second: 2198471
Packet drops:
     No memory: 58786, Not IP: 0
```

```
     Not IPv4: 0, Too small: 0
     Fragments: 0, ICMP: 0
     TCP: 0, Unknown: 0
     Not JUNOS flow: 0
File Transfer:
     FTP bytes: 585981447, per second: 1313320, peak per second: 4857798
     FTP files: 48, per second: 0, peak per second: 0
     FTP failure: 8
Export channel: 0
     Current server: Primary
     Primary server state: FTP error, Secondary server state: Not configured Export channel: 1
     Current server: Primary
     Primary server state: OK, Secondary server state: Not configured
```

Release Information

Command introduced before Junos OS Release 7.4.

show services ha detail

IN THIS SECTION

- Syntax | 25533
- Description | 25533
- Options | 25533
- Required Privilege Level | 25533
- Output Fields | 25533
- Sample Output | 25534
- Release Information | 25535

Syntax

```
show services ha detail
<interface interface-name>
```

Description

Display detailed information for stateful sync processing for a specified interface or for all interfaces.

Options

none Display detailed information for stateful sync processing for all interfaces.

interface-name (Optional) Name of a specific interface.

Required Privilege Level

view

Output Fields

Table 1405 on page 25533 lists the output fields for the `show services ha detail` command. Output fields are listed in the approximate order in which they appear.

Table 1405: show services ha detail Output Fields

Field Name	Field Description
Interface	Name of the interface for which information is reported.
Inter-chassis	
Role	Role of the interface. • `active`—Active interface. • `backup`–Backup interface.

Table 1405: show services ha detail Output Fields *(Continued)*

Field Name	Field Description
Connection	Status of the peer connection. • Up • Down
Synchronization	Synchronization state of peers. • Off–Peers are not currently engaged in synchronization.. • Cold–Peers are in a pre-synchronization state. • Hot—Peers are ready for synchronization.
Peers	
Local	Local peer IP address.
Port	Local peer port number.
Remote	Remote peer IP address.
Port	Remote peer port number.

Sample Output

show services ha detail

```
user@host> show services ha detail
Interface:        ms-7/0/0
Inter-chassis:    Role: active, Connection: Up, Synchronization: Hot
Peers:            Local: 192.0.2.1 Port: 4001, Remote: 192.0.2.2 Port: 4001

Interface:        ms-7/1/0
Inter-chassis:    Role: active, Connection: Down, Synchronization: Off
```

```
Peers:            Local: 198.51.100.1 Port: 4001, Remote: 198.51.100.2 Port: 4001

Interface:        ms-8/0/0
Inter-chassis:    Role: active, Connection: Up, Synchronization: Cold
Peers:            Local: 203.0.113.1 Port: 4001, Remote: 203.0.113.2 Port: 4001

Interface:        ms-8/1/0
Inter-chassis:    Role: active, Connection: Up, Synchronization: Hot
Peers:            Local: 10.10.10.1 Port: 4001, Remote: 10.10.10.2 Port: 4001
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> *Inter-Chassis Stateful Synchronization for Long Lived NAT and Stateful Firewall Flows (MS-MPC, MS-MIC) Overview (Release 16.1 and later)*

show services ha statistics

IN THIS SECTION

- Syntax | 25536
- Description | 25536
- Options | 25536
- Required Privilege Level | 25536
- Output Fields | 25536
- Sample Output | 25541
- Release Information | 25542

Syntax

```
show services ha statistics
<interface interface-name>
```

Description

Display detailed statistics for stateful sync processing for a specified interface or for all interfaces.

Options

none Display detailed statistics for stateful sync processing for all interfaces.

interface-name (Optional) Name of a specific interface.

Required Privilege Level

view

Output Fields

Table 1406 on page 25536 lists the output fields for the `show services ha statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1406: show services ha statistics Output Fields

Field Name	Field Description
Interface	Interface name.
Inter-chassis	
Role	Role of the interface. • `active`—Active interface. • `backup`–Backup interface.

Table 1406: show services ha statistics Output Fields *(Continued)*

Field Name	Field Description
Connection	Status of the peer connection. • Up • Down
Synchronization	Synchronization state of peers. • Off–Peers are not currently engaged in synchronization. • Cold–Peers are in a pre-synchronization state. • Hot—Peers are ready for synchronization.
Peers	
Local	Local peer IP address.
Port	Local peer port number.
Remote	Remote peer IP address.
Port	Remote peer port number.
Connection Status	
TCP connection establish	Number of times a TCP connection is established.
TCP connection teardown	Number of times a TCP connection is torn down.
UDP address exchange sent	Number of times a UDP address is sent.
Stateful sync start sent	Number of stateful sync start messages sent by the backup PIC, indicating the start of the cold sync phase.

Table 1406: show services ha statistics Output Fields *(Continued)*

Field Name	Field Description
Stateful sync start received	Number of stateful sync start messages received by active PIC, indicating the start of the cold sync phase.
Cold sync completed count	Number of times the PIC has successfully completed the cold sync phase.
Session Add Statistics	
Sent	Number of session add statistics sent by the active PIC.
Received	Number of session add statistics received by the backup PIC.
Completed	Number of session adds completed on the active and backup PICs.
rate	Number of sessions currently added per second.
Nack sent	Number of times that a session add failed on the backup PIC, resulting in the sending of a Nack message to the active PIC.
Nack received	Number of Nack messages received from backup PIC due to session add failure.
Add pending	Number of sessions eligible for synchronization, but not yet synchronized.
Session Delete Statistics	
Sent	Number of session deletes sent by the active PIC.
Received	Number of session deletes received by the backup PIC.
Completed	Number of session deletes completed on the active and backup PICs.
rate	Number of sessions currently deleted per second.

Table 1406: show services ha statistics Output Fields *(Continued)*

Field Name	Field Description
Nack sent	Number of times that a session add failed on the backup PIC, resulting in the sending of a Nack message to the active PIC.
Nack received	Number of Nack messages received from backup PIC due to session add failure.
Session not found	Number of sessions not found when session delete was attempted.
Session Error Statistics	
Session attach failures	Number of high-availability extension creation failures on the active PIC.
Session detach failures	Number of high-availability extension deletion failures on the active PIC.
Session extension get failures	Number of times that the high-availability extension is not available when requested.
Session nullify	Number of times the high-availability session creation failed on the active PIC.
Lookup fail	Number of times session lookup failed because the session has already been released by the infrastructure.
Initiate fail	Number of times session creation failed on the backup PIC.
Activate fail	Number of times session activation failed on the backup PIC.
Illegal flow type	Number of times an illegal flow type occurred on the active and backup PICs.
Illegal service set	Number of times service set extraction failed on backup and active PICs.

Table 1406: show services ha statistics Output Fields *(Continued)*

Field Name	Field Description
Unsupported protocol	Number of times that a session was not backed up because the protocol was neither TCP or UDP.
Send overflow	Number of times buffer overflowed when the high-availability session was created on the active PIC.
Send discard	Number of sessions that not synchronized to the backup, even though they were eligible for synchronization. This occurs whe at least one plugin in the service set indicates that a session should not be synchronized.
Spurious	Number of packets received on the backup PIC for which there are no existing sessions
Process incoming failed	Number of times JMUX header processing failed.
Session ignored	Number of sessions that were eligible for synchronization, but are ignored because stateful sync is not supported for them, such as ALG sessions
JMUX Error Statistics	Synchronization statistics related to the JMUX library.
JMUX begin fail	Number of times that JMUX key verification or header creation failed.
JMUX commit fail	Number of times addition of JMUX data failed.
JMUX flush fail	Number of times a send of JMUX data failed.
Invalid plugin header	Number of times stateful sync messages were rejected due to an invalid plugin header (internal error).
Invalid plugin name	Number of times stateful sync messages were rejected due to an invalid plugin name (internal error).

Table 1406: show services ha statistics Output Fields *(Continued)*

Field Name	Field Description
Invalid plugin length	Number of times stateful sync messages were rejected due to invalid plugin length (internal error).
Plugin receive error	Number of times installation of plugin information failed on the backup.
Plugin send error	Number of times the plugin failed to pack the extension.
IDL Error Statistics	Statistics concerning encode or decode errors at the backup.
IDL encode fail	Number of times IDL encoding failed on the active and backup PICs.
IDL decode fail	Number of times IDL decoding failed on the active and backup PICs.

Sample Output

show services ha statistics

```
user@host> show services ha statistics
Interface:        ms-5/0/0
Inter-chassis:    Role: active, Connection: Up, Synchronization: Hot
Peers:            Local: 192.0.2.2 Port: 4001, Remote: 192.0.2.1 Port: 4001
Connection Status:
    TCP connection establish: 8, Teardown: 8
    UDP address exchange sent: 8, Received: 8
    Stateful sync start sent: 0, Received: 8
    Cold sync completed count: 0
Session Add Statistics:
    Sent: 255, Received: 0
    Completed: 255, Rate: 0
    Nack sent: 0, Nack received: 0
    Add pending: 0
Session Delete Statistics:
    Sent: 255, Received: 0
    Completed: 255, Rate: 0
```

```
    Nack sent: 0, Nack received: 0
    Session not found: 0
Session Error Statistics:
   Session attach failures: 0, Session detach failures: 0
    Session extension get failures: 0, Session nullify: 0
    Lookup fail: 0, Initiate fail: 0, Activate fail: 0
    Illegal flow type: 0, Illegal service set: 0
    Unsupported protocol: 0, Send overflow: 0, Send discard: 0
    Spurious: 0, Process incoming failed: 0, Session ignored: 0
JMUX Error Statistics:
    JMUX begin fail: 0, JMUX commit fail: 0, JMUX flush fail: 0
    Invalid plugin header: 0, Invalid plugin name: 0
    Invalid plugin length: 0, Plugin receive error: 0, Plugin send error: 0
IDL Error Statistics:
    IDL encode fail: 0, IDL decode fail: 0
```

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Inter-Chassis Stateful Synchronization for Long Lived NAT and Stateful Firewall Flows (MS-MPC, MS-MIC) Overview (Release 16.1 and later)

show services hcm pic-statistics

IN THIS SECTION

- Syntax | 25543
- Description | 25543
- Options | 25543
- Required Privilege Level | 25543
- Output Fields | 25543
- Sample Output | 25547

Sample Output | 25549
Release Information | 25550

Syntax

```
show services hcm pic-statistics
<interface interface-name>
```

Description

Display the statistics collected (from the services PICs) for HTTP header enrichment.

Options

none	Display the statistics for all the services PICs.
interface *interface-name*	(Optional) Display the statistics for the specified services PIC.

Required Privilege Level

view

Output Fields

Table 1407 on page 25543 lists the output fields for the `show services hcm pic-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1407: show services hcm pic-statistics Output Fields

Field Name	Field Description
`Interface`	Name of the services PIC interface for which statistics are displayed.

Table 1407: show services hcm pic-statistics Output Fields *(Continued)*

Field Name	Field Description
`Session statistics`—For each services PIC.	
`Number of Session Interest events`	Number of Session Interest events.
`Number of Session Create events`	Number of Session Create events.
`Number of Session Close events`	Number of Session Close events.
`Number of Session Destroy events`	Number of Session Destroy events.
`Number of Session Data events`	Number of Session Data events.
`Number of Session Handle failures`	Number of Session Handle failures.
`Number of Session Extension allocations`	Number of Session Extension allocations that were successful.
`Number of Session Extension alloc failures`	Number of Session Extension allocations that failed.
`Number of Session Extension frees`	Number of Session Extension frees (memory releases).
`TCP Proxy statistics`	
`Number of missing stbuf`	Number of missing stream buffers.
`Number of stbuf initializations`	Number of stream buffer initializations that were successful.
`Number of stbuf initialization failures`	Number of stream buffer initializations that failed.
`Number of stbuf store failures`	Number of stream buffer store failures.

Table 1407: show services hcm pic-statistics Output Fields *(Continued)*

Field Name	Field Description
`Number of stbuf frees`	Number of stream buffer frees (memory releases) that were successful.
`Number of stbuf free failures`	Number of stream buffer frees that failed.
`Number of stbuf sends`	Number of stream buffer sends that were successful.
`Number of stbuf send failures`	Number of stream buffer sends that failed.
`Number of stbuf receives`	Number of stream buffer receives that were successful.
`Number of stbuf throttles`	Number of stream buffer throttles. Throttles are done when the stream buffer queue is full.
`Number of invalid stbuf`	Number of invalid stream buffers.
`THR statistics`	
`Number of THR creates`	Number of successful TCP Header Rewriter (THR) Create Requests.
`Number of missing THR handles`	Number of missing THR handles.
`Number of THR create failures`	Number of THR Create Requests that failed.
`Number of THR store failures`	Number of THR store failures.
`Number of THR short circuit failures`	Number of THR short circuit (packet bypass) failures.
`Number of THR update failures`	Number of THR updates that failed.

Table 1407: show services hcm pic-statistics Output Fields *(Continued)*

Field Name	Field Description
Number of THR state updates	Number of THR state updates.
Number of THR destroy failures	Number of THR destroys that failed.
Number of THR destroys	Number of THR Cleanup Requests that were successful.
JCPP statistics	
Number of JCPP handle allocations	Number of Juniper Content and Protocol Parsers (JCPP) handle allocations that were successful.
Number of JCPP handle allocation failures	Number of JCPP handle allocations that failed.
Header Insertion statistics	
Number of HCM Header Insertions	Number of times that tags were successfully inserted into HTTP headers.
Number of HCM Header Insertion failures	Number of times that the insertion of tags into HTTP headers failed.
Number of HCM Header Renamed	Number of times that HTTP headers were successfully renamed.
Number of HCM Header Rename failures	Number of times that HTTP header rename attempts failed.
Number of HCM IPV4 Mask modifications	Number of times IPv4 address mask was inserted.
Number of HCM IPV6 Mask modifications	Number of times IPv6 address mask was inserted.
Number of HCM Tags too large	Number of tags that were not inserted into HTTP headers because the tag size was larger than the maximum allowed size.

Table 1407: show services hcm pic-statistics Output Fields *(Continued)*

Field Name	Field Description
Number of HCM Tag encryption failures	Number of times that the encryption of HTTP tags used for header insertion failed.
Number of HCM requests	Number of HTTP header enrichment requests.
Number of missing Subscribers in HCM	Number of times that tags were not inserted because subscriber was missing.
Number of HCM missing subscriber attributes	Number of times that tags were not inserted because subscriber attributes were missing.
Number of HCM missing IPV4 attributes	Number of times that tags were not inserted because subscriber IPv4 user address attributes were missing.
Number of HCM missing IPV6 attributes	Number of times that tags were not inserted because subscriber IPv6 user address attributes were missing.
Number of HCM IPV4 / IPV6 tag insertions	Number of times that an IPv4 or an IPv6 user address tag was successfully inserted into HTTP headers when the tag rule included both IPv4 and IPv6 user address tags.

Sample Output

show services hcm pic-statistics (mams interface)

```
user@host> show services hcm pic-statistics
Interface: mams-3/0/0
Session statistics
    Number of Session Interest events              :224590
    Number of Session Create events                :224590
    Number of Session Close events                 :224590
    Number of Session Destroy events               :224590
    Number of Session Data events                  :224589
    Number of Session Handle failures              :0
```

```
    Number of Session Extension allocations         :224590
    Number of Session Extension alloc failures      :0
    Number of Session Extension frees               :224590
TCP Proxy statistics
    Number of missing stbuf                         :0
    Number of stbuf initializations                 :0
    Number of stbuf initialization failures         :0
    Number of stbuf store failures                  :0
    Number of stbuf frees                           :0
    Number of stbuf free failures                   :0
    Number of stbuf sends                           :0
    Number of stbuf send failures                   :0
    Number of stbuf receives                        :0
    Number of stbuf throttles                       :0
    Number of invalid stbuf                         :0
THR statistics
    Number of THR creates                           :224590
    Number of missing THR handles                   :0
    Number of THR create failures                   :0
    Number of THR store failures                    :0
    Number of THR short circuit failures            :0
    Number of THR update failures                   :0
    Number of THR state updates                     :449180
    Number of THR destroy failures                  :0
    Number of THR destroys                          :0
JCPP statistics
    Number of JCPP handle allocations               :0
    Number of JCPP handle allocation failures       :0
Header Insertion statistics
    Number of HCM Header Insertions                 :224589
    Number of HCM Header Insertion failures         :0
    Number of HCM Header Renamed                    :0
    Number of HCM Header Rename failures            :0
    Number of HCM IPV4 Mask modifications           :0
    Number of HCM IPV6 Mask modifications           :0
    Number of HCM Tags too large                    :0
    Number of HCM Tag encryption failures           :0
    Number of HCM requests                          :224589
    Number of missing Subscribers in HCM            :0
    Number of HCM missing subscriber attributes     :0
    Number of HCM missing IPV4 attributes           :0
    Number of HCM missing IPV6 attributes           :0
```

```
    Number of HCM IPV4 / IPV6 tag insertions        :0
```

Sample Output

show services hcm pic-statistics (vms- interface)

```
user@host> show services hcm pic-statistics
Interface: vms-5/2/0
Session statistics
    Number of Session Interest events               :90064
    Number of Session Create events                 :90064
    Number of Session Close events                  :90064
    Number of Session Destroy events                :90064
    Number of Session Data events                   :90064
    Number of Session Handle failures               :0
    Number of Session Extension allocations         :90064
    Number of Session Extension alloc failures      :0
    Number of Session Extension frees               :90064
TCP Proxy statistics
    Number of missing stbuf                         :0
    Number of stbuf initializations                 :0
    Number of stbuf initialization failures         :0
    Number of stbuf store failures                  :0
    Number of stbuf frees                           :0
    Number of stbuf free failures                   :0
    Number of stbuf sends                           :0
    Number of stbuf send failures                   :0
    Number of stbuf receives                        :0
    Number of stbuf throttles                       :0
    Number of invalid stbuf                         :0
THR statistics
    Number of THR creates                           :90064
    Number of missing THR handles                   :0
    Number of THR create failures                   :0
    Number of THR store failures                    :0
    Number of THR short circuit failures            :0
    Number of THR update failures                   :0
    Number of THR state updates                     :180128
    Number of THR destroy failures                  :0
    Number of THR destroys                          :0
```

```
JCPP statistics
    Number of JCPP handle allocations              :0
    Number of JCPP handle allocation failures      :0
Header Insertion statistics
    Number of HCM Header Insertions                :90061
    Number of HCM IP Mask modifications            :90061
    Number of HCM Header Insertion failures        :0
    Number of HCM Tags too large                   :0
    Number of HCM Tag encryption failures          :0
    Number of HCM requests                         :90061
    Number of missing Subscribers in HCM           :90061
```

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2 and 19.4R1 on MX Series routers MX240, MX480 and MX960.

RELATED DOCUMENTATION

show services hcm statistics | 25550

show services hcm statistics

IN THIS SECTION

- Syntax | 25551
- Description | 25551
- Options | 25551
- Required Privilege Level | 25551
- Output Fields | 25551
- Sample Output | 25552
- Release Information | 25552

Syntax

```
show services hcm statistics rule rule-name
```

Description

Display the statistics collected for HTTP header enrichment for a specified tag rule.

> **NOTE**: This command displays output only if the `count` statement is configured for the term in a tag rule at the `[edit services hcm tag-rule rule-name term term-name then]` hierarchy level.
>
> If you change the configuration of tag rules during an existing subscriber data session and commit the change, the tag rule statistics are reset to 0 and stop incrementing for the existing TCP sessions.

Options

none	Display detailed statistics about stateful sync processing for all interfaces.
rule *rule-name*	Display statistics for the specified tag rule.

Required Privilege Level

view

Output Fields

Table 1408 on page 25551 lists the output fields for the `show services hcm statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1408: show services hcm statistics Output Fields

Field Name	Field Description
`Interface`	Name of the interface for which the statistics are displayed.

Table 1408: show services hcm statistics Output Fields *(Continued)*

Field Name	Field Description
`Term id`	Identifier for the term (in the tag rule) for which the statistics are displayed.
`Hits`	Number of times that the term was matched. This field displays the aggregate number of occurrences in service sets that include the term.

Sample Output

show services hcm statistics rule

```
user@host> show services hcm statistics rule rule1
Interface: mams-3/1/0
Term id              Hits
1                  58
Interface: mams-4/1/0
Term id              Hits
1                  144
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

count (HTTP Header Enrichment) | **4686**

Configuring HTTP Header Enrichment Overview

show services hcm pic-statistics | **25542**

show services icap-redirect server status

IN THIS SECTION

- Syntax | 25553
- Description | 25553
- Required Privilege Level | 25553
- Sample Output | 25553
- Release Information | 25554

Syntax

```
show services icap-redirect server status
```

Description

Display the status of On-Premises in DLP.

Required Privilege Level

view

Sample Output

show services icap-redirect server status

```
user@host> show services icap-redirect server status
 ICAP Status :
         Spu-1 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
         Spu-1 Profile: icap-pf1 Server: icap-svr2 : UP
ICAP Status :
         Spu-2 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
```

```
         Spu-2 Profile: icap-pf1 Server: icap-svr2 : UP
ICAP Status :
         Spu-3 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
         Spu-3 Profile: icap-pf1 Server: icap-svr2 : UP
```

show services icap-redirect server status logical-system

```
user@host> show services icap-redirect server status logical-system LSYS1
 ICAP Status :
         spu-1 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
         spu-2 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
         spu-3 Profile: icap-pf1 Server: icap-svr1 : UP
```

Release Information

Command introduced in Junos OS Release 18.1R1.

The logical system option is introduced in Junos OS Release 18.3R1.

show services icap-redirect statistic

IN THIS SECTION

- Syntax | 25555
- Description | 25555
- Options | 25555
- Required Privilege Level | 25555
- Output Fields | 25556
- Sample Output | 25556
- Release Information | 25558

Syntax

```
show services icap-redirect statistic
<all-logical-systems-tenants>
<root-logical-system>
<logical-system (logical-system-name | all)>
<tenant (tenant-name | all)>
```

Description

Displays the ICAP services redirects statistic. ICAP services redirect redirects the HTTP or HTTPS traffic to any third-party server. The security device acts as an SSL proxy server and decrypts the pass-through traffic with the proper SSL profile under a security policy.

Options

all-logical-systems-tenants	(Optional) Displays the ICAP services redirects statistic for the root logical system, all logical systems, and all tenant systems.
logical-system *logical-system-name*	(Optional) Displays the ICAP services redirects statistic for the specified logical system.
logical-system *all*	(Optional) Displays the ICAP services redirects statistic for all the logical systems.
root-logical-system	(Optional) Displays the ICAP services redirects statistic for the root logical system.
tenant *tenant-name*	(Optional) Displays the ICAP services redirects statistic for the specified tenant system.
tenant *all*	(Optional) Displays the ICAP services redirects statistic for all the tenant systems.

Required Privilege Level

view

Output Fields

Table 1409 on page 25556 lists the output fields for the `show services icap-redirect statistic` command. Output fields are listed in an approximate order in which they appear.

Table 1409: show services icap-redirect statistic

Field Name	Field Description
`Message Redirected`	Number of messages redirected.
`Message Received`	Number of messages received.

Sample Output

show services icap-redirect statistic root-logical-system

```
user@host> show services icap-redirect statistic root-logical-system
ICAP Redirect statistic:
  Message Redirected                  : 0
      Message REQMOD Redirected       : 0
      Message RESPMOD Redirected      : 0
  Message Received                    : 38
      Message REQMOD Received         : 0
      Message RESPMOD Received        : 0
Fallback:      permit           log-permit       reject
  Timeout         0                 0                 0
  Connectivity    0                 0                 0
  Default         0                 0                 0
```

show services icap-redirect statistic all-logical-systems-tenants

```
user@host> show services icap-redirect statistic all-logical-systems-tenants
Logical system: root-logical-system
ICAP Redirect statistic:
  Message Redirected                  : 0
      Message REQMOD Redirected       : 0
      Message RESPMOD Redirected      : 0
```

```
  Message Received                    : 38
      Message REQMOD Received         : 0
      Message RESPMOD Received        : 0
Fallback:      permit           log-permit       reject
  Timeout         0                 0                 0
  Connectivity    0                 0                 0
  Default         0                 0                 0
Logical system: LSYS1
ICAP Redirect statistic:
  Message Redirected                  : 0
      Message REQMOD Redirected       : 0
      Message RESPMOD Redirected      : 0
  Message Received                    : 38
      Message REQMOD Received         : 0
      Message RESPMOD Received        : 0
Fallback:      permit           log-permit       reject
  Timeout         0                 0                 0
  Connectivity    0                 0                 0
  Default         0                 0                 0
Tenants: TSYS1
ICAP Redirect statistic:
  Message Redirected                  : 0
      Message REQMOD Redirected       : 0
      Message RESPMOD Redirected      : 0
  Message Received                    : 38
      Message REQMOD Received         : 0
      Message RESPMOD Received        : 0
Fallback:      permit           log-permit       reject
  Timeout         0                 0                 0
  Connectivity    0                 0                 0
  Default         0                 0                 0
```

show services icap-redirect statistic logical-system LSYS1

```
user@host> show services icap-redirect statistic logical-system LSYS1
ICAP Redirect statistic:
              Message Redirected                  : 0
                      Message REQMOD Redirected       : 0
                      Message RESPMOD Redirected      : 0
              Message Received                    : 0
                      Message REQMOD Received         : 0
```

```
                        Message RESPMOD Received        : 0
        Fallback:      permit           log-permit       reject
          Timeout         0                 0                 0
          Connectivity    0                 0                 0
          Default         0                 0                 0
```

show services icap-redirect statistic tenant TSYS1

```
user@host> show services icap-redirect statistic tenant TSYS1
ICAP Redirect statistic:
              Message Redirected                  : 0
                      Message REQMOD Redirected       : 0
                      Message RESPMOD Redirected      : 0
              Message Received                    : 0
                      Message REQMOD Received         : 0
                      Message RESPMOD Received        : 0
        Fallback:      permit           log-permit       reject
          Timeout         0                 0                 0
          Connectivity    0                 0                 0
          Default         0                 0                 0
```

show services icap-redirect statistic

```
user@host:TSYS1> show services icap-redirect statistic
ICAP Redirect statistic:
              Message Redirected                  : 0
                      Message REQMOD Redirected       : 0
                      Message RESPMOD Redirected      : 0
              Message Received                    : 0
                      Message REQMOD Received         : 0
                      Message RESPMOD Received        : 0
        Fallback:      permit           log-permit       reject
          Timeout         0                 0                 0
          Connectivity    0                 0                 0
          Default         0                 0                 0
```

Release Information

Command introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

ICAP Service Redirect

show services icap-redirect status

IN THIS SECTION

- Syntax | 25559
- Description | 25559
- Options | 25560
- Required Privilege Level | 25560
- Output Fields | 25560
- Sample Output | 25561
- Release Information | 25562

Syntax

```
show services icap-redirect status
<all-logical-systems-tenants>
<root-logical-system>
<logical-system (logical-system-name | all)>
<tenant (tenant-name | all)>
```

Description

Displays the status of ICAP services redirects. ICAP services redirect redirects the HTTP or HTTPS traffic to any third-party server. The security device acts as an SSL proxy server and decrypts the pass-through traffic with the proper SSL profile under a security policy.

Options

`all-logical-systems-tenants`	(Optional) Displays the status of ICAP services redirects for the root logical system, all logical systems, and all tenant systems.
`logical-system` *`logical-system-name`*	(Optional) Displays the status of ICAP services redirects for the specified logical system.
`logical-system` *`all`*	(Optional) Displays the status of ICAP services redirects for all the logical systems.
`root-logical-system`	(Optional) Displays the status of ICAP services redirects for the root logical system.
`tenant` *`tenant-name`*	(Optional) Displays the status of ICAP services redirects for the specified tenant system.
`tenant` *`all`*	(Optional) Displays the status of ICAP services redirects for all the tenant systems.

Required Privilege Level

view

Output Fields

Table 1410 on page 25560 lists the output fields for the `show services icap-redirect status` command. Output fields are listed in an approximate order in which they appear.

Table 1410: show services icap-redirect status

Field Name	Field Description
`ICAP Status`	Status of the ICAP services redirect.
`Profile`	Name of the security profile assigned to the tenant systems or logical systems.
`Server`	Name of the server associated with the tenant systems or logical systems.

Sample Output

show services icap-redirect status root-logical-system

```
user@host> show services icap-redirect status root-logical-system
ICAP Status :
         Profile: p1 Server: s1 : UP
ICAP Status :
         Profile: p1 Server: s1 : UP
ICAP Status :
         Profile: p2 Server: s2 : UP
ICAP Status :
         Profile: p2 Server: s2 : UP
```

show services icap-redirect status all-logical-systems-tenants

```
user@host> show services icap-redirect status all-logical-systems-tenants
LSYS1
ICAP Status :
    Spu-1 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
    Spu-1 Profile: icap-pf1 Server: icap-svr2 : UP
TSYS1
ICAP Status :
    Spu-1 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
    Spu-1 Profile: icap-pf1 Server: icap-svr2 : UP
```

show services icap-redirect status logical-system LSYS1

```
user@host> show services icap-redirect status logical-system LSYS1
ICAP Status :
    Spu-1 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
    Spu-1 Profile: icap-pf1 Server: icap-svr2 : UP
```

show services icap-redirect status tenant TSYS1

```
user@host> show services icap-redirect status tenant TSYS1
ICAP Status :
    Spu-1 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
    Spu-1 Profile: icap-pf1 Server: icap-svr2 : UP
```

show services icap-redirect status

```
user@host:TSYS1> show services icap-redirect status
ICAP Status :
    Spu-1 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
    Spu-1 Profile: icap-pf1 Server: icap-svr2 : UP
```

Release Information

Command introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

ICAP Service Redirect

show services ids

IN THIS SECTION

- Syntax | 25563
- Description | 25563
- Options | 25563
- Required Privilege Level | 25564

- Output Fields | 25564
- Sample Output | 25569
- Release Information | 25573

Syntax

```
show services ids (destination-table | pair-table | source-table)
<brief | extensive | terse>
<destination-prefix destination-prefix-name>
<interface interface-name>
<limit number>
<order  (anomalies | bytes | flows | packets)>
<service-set service-set-name>
<source-prefix source-prefix-name>
<threshold number>
```

Description

Display information about intrusion detection service (IDS) events. All events gathered by IDS are reported as anomalies. For example, events such as **create forward or watch flow**, **FTP passive**, and **FTP active** are genuinely allowed by the stateful firewall but are logged as anomalies to track the rates and number for these events.

Options

destination-table	Display information for an address under possible attack.
pair-table	Display information for a particular suspected attack source and destination address pair.
source-table	Display information for an address that is a suspected attacker.
brief | extensive | terse	(Optional) Display the specified level of output.
destination-prefix ***destination-prefix-name***	(Optional) Display information for a particular destination prefix.

interface *interface-name*	(Optional) On M Series and T Series routers, the *interface-name* can be **sp-***fpc/pic/port* or **rsp***number*.
limit *number*	(Optional) Maximum number of entries to display. By default, all tables display the top 32 entries sorted by the number of events for the criteria chosen. To display additional entries, configure the limit option to set up to 256 entries.
order	(Optional) Display events according to one of the following table-ordering criteria. The default is anomalies. • **anomalies**—Display information for particular anomalies. • **bytes**—Order output by number of bytes received. • **flows**—Order output by number of flows. • **packets**—Order output by number of packets received.
service-set *service-set-name*	(Optional) Display information about a particular service set.
source-prefix *source-prefix-name*	(Optional) Display information about a particular source prefix.
threshold *number*	(Optional) Limit the display to events with this number of anomalies, bytes, flows, or packets, whichever criterion you specify for order. For example, to display all events with more than 100 flows, specify order flows and threshold 100.

Required Privilege Level

view

Output Fields

Table 1411 on page 25564 lists the output fields for the `show services ids` command. Output fields are listed in the approximate order in which they appear.

Table 1411: show services ids Output Fields

Field Name	Field Description	Output Level
Interface	Name of an adaptive services interface.	All levels

Table 1411: show services ids Output Fields *(Continued)*

Field Name	Field Description	Output Level
Service set	Name of a service set. Individual empty service sets are not displayed, but if no service set has any flows, a flow table header is printed for each service set.	All levels
Sorting order	Primary mode to display information: **Anomalies**, **Bytes**, **Flows**, or **Packets**.	All levels
Source address	Name of the source address.	All levels
Dest address	Name of the destination address.	All levels
Time	Total time the information has been in the table.	All levels
Flags	**Flags** can be **Forced**, **F** (terse output only), **SYNcookie**, **S** (terse output only), **Forced+SYNcookie**, and **F+S** (terse output only). The **SYNcookie** flag is visible only in the destination table.	All levels
Application	Configured application, such as **FTP** or **Telnet**.	All levels
Bytes	Total number of bytes sent from the source to the destination address, in thousands (**k**) or millions (**m**).	All levels
Packets	Total number of packets sent from the source to the destination address, in thousands (**k**) or millions (**m**).	All levels
Flows	Total number of flows of packets sent from the source to the destination address, in thousands (**k**) or millions (**m**).	All levels
Anomalies	Total number of packets in the anomaly table, in thousands (**k**) or millions (**m**).	All levels

Table 1411: show services ids Output Fields *(Continued)*

Field Name	Field Description	Output Level
Anomaly description	One or more of the following types of anomalies. For more information, see the detailed descriptions in the stateful firewall section of the System Log Explorer. • **First packet of TCP session not SYN** • **ICMP echo request dropped, because sequence number duplicated** • **ICMP echo reply dropped. No matching sequence number** • **ICMP echo request dropped. Too many echo requests without echo reply** • **ICMP header length check failed** • **ICMP packet length greater than 64K** • **IP fragment assembly timeout** • **IP fragment length error** • **IP fragment overlap** • **IP packet length greater than 64K** • **IP packet too short** • **IP packet with broadcast destination address** • **IP packet with checksum error** • **IP packet with incorrect length** • **IP packet with TTL equal to 0**	**extensive**

Table 1411: show services ids Output Fields *(Continued)*

Field Name	Field Description	Output Level
Anomaly description (continued)	• **IP packet with version other than 4** • **Land attack (IP src address = dest address)** • **No matching SFW rule; attempting to create discard flow** • **Number of open sessions exceeds IDS limit; packet dropped** • **Packet rate exceeds IDS limit; packet dropped** • **Session creation rate exceeds IDS limit; packet dropped** • **SFW application message too long** • **SFW discard packet contains non-configured IP option types** • **SFW drop packet because of discard flow** • **SFW dropped TCP watch packet** • **SFW rules request FTP active mode data packets to be accepted; attempting to create forward flow** • **SFW rules request FTP passive mode data packets to be accepted; attempting to create forward flow** • **SFW rules request packet to be accepted; attempting to create forward or watch flow** • **SFW rules request packet to be discarded; attempting to create discard flow** • **SFW rules request packet to be rejected; attempting to create reject flow** • **SFW discard flow requires packet to be dropped** • **SFW SYN defense** • **Smurf attack (ping to IP broadcast address)** • **TCP FIN/RST or SYN/(URG|FIN|RST) flags set**	**extensive**

Table 1411: show services ids Output Fields *(Continued)*

Field Name	Field Description	Output Level
	• **TCP header length check failed** • **TCP port scan (port not in LISTEN state)** • **TCP seq number zero and FIN/PSH/RST flags set** • **TCP seq number zero and no flags set** • **TCP source or destination port zero** • **TCP SYN flood attack** • **UDP header length check failed** • **UDP port scan (port not in LISTEN state)** • **UDP source or destination port zero**	
Count	Number of times that a particular anomaly occurred, in thousands (**k**) or millions (**M**).	**extensive**
Rate (eps)	Anomaly events per second. The IDS subsystem attempts to maintain a weighted average of rates, which might not reflect the exact incoming rate of attack at low rates. However, at high rates exceeding 160 events per second, the rates generally match.	**extensive**
Elapsed	Time since the same type of event last occurred.	**extensive**
Total IDS table entries	Number of entries in the IDS table. This number is not necessarily the sum of all entries displayed.	All levels
Total failed IDS table entry insertions	Number of IDS entries not allowed into the table because the table was full	All levels

Table 1411: show services ids Output Fields *(Continued)*

Field Name	Field Description	Output Level
Total number of events (closed flows and anomalies detected)	Total number of events since the system was started or since the `show ids services` command was executed.	All levels

Sample Output

show services ids destination-table

```
user@host> show services ids destination-table
Interface: sp-1/3/0, Service set: null-sfw
Sorting order: Packets
Source address       Dest address   Time    Flags             Application

any             ->   10.58.255.146   36m12s SYN cookie
  Bytes:   35.0 m, Packets:  822.0 k, Flows:  274.0 k, Anomalies: 2251.0 k


Total IDS table entries: 87
Total failed IDS table entry insertions 0
Total number of events (closed flows and anomalies detected): 2606018
```

show services ids destination-table extensive

```
user@host> show services ids destination-table extensive
Interface: sp-1/3/0, Service set: null-sfw
Sorting order: Packets
Source address       Dest address   Time    Flags             Application

any             ->   10.58.255.146   35m52s SYN cookie
  Bytes:   34.0 m, Packets:  798.0 k, Flows:  266.0 k, Anomalies: 2251.0 k
    Anomalies                                        Count   Rate(eps) Elapsed
    First packet of TCP session not SYN              160.0 k       0       14s
    TCP source or destination port zero              634.0 k   154.6     3m37s
```

```
    UDP source or destination port zero             633.0 k   170.0     3m37s
    ICMP header length check failed                    2875     0.9     3m37s
    IP fragment assembly timeout                    820.0 k    12.8     3m18s
    UDP header length check failed                     385     0.5     3m53s
    TCP header length check failed                     383     0.5     3m53s

Total IDS table entries:
87
Total failed IDS table entry insertions
0
Total number of events (closed flows and anomalies detected):
2598063
```

show services ids destination-table extensive order anomalies

```
user@host> show services ids destination-table extensive order anomalies
Interface: sp-0/2/0, Service set: ss1
IDS sorting order: Anomalies
Source address       Dest address      Time Flags             Application
192.0.2.1        ->  198.51.100.1         1m28s               junos-ftp
  Bytes: 1065, Packets: 18, Flows: 1, Anomalies: 10
   Anomaly description                                Count   Rate(eps)  Elapsed
   creating forward or watch flow                       1     15.6       1m28s
   Number of open sessions exceeds IDS limit            9      0.8         18s

Total IDS table entries:                                      3
Total failed IDS table entry insertions                       0
Total number of events (closed flows and anomalies):         11
```

show services ids pair-table extensive

```
user@host> show services ids pair-table extensive
Interface: sp-3/2/0, Service set: ss_all_limits
IDS sorting order: Packets
Source address       Dest address      Time Flags             Application
198.51.100.4         198.51.100.4      2m20s                  junos-ftp
  Bytes: 5.7k, Packets: 102.0, Flows: 41.0, Anomalies: 462.0
    Anomaly description                                Count     Rate   Elapsed
    creating forward or watch flow                      41.0      8.8     2m17s
    Packet rate exceeds IDS src limit                   21.0      7.1     2m17s
```

```
    Session creation rate exceeds IDS src limit       359.0     99.7     2m16s
    TCP SYN flood attack                               41.0      1.9     1m30s


Total IDS table entries:                                    3
Total failed IDS table entry insertions                     0
Total number of events (closed flows and anomalies):        462
```

show services ids pair-table extensive limit

```
user@host> show services ids pair-table extensive limit 3
Interface: sp-1/3/0, Service set: null-sfw
Sorting order: Packets
Source address       Dest address   Time    Flags             Application
10.58.255.18    ->   10.58.255.146   38m41s SYN cookie
  Bytes:  286.0 m, Packets: 2823.0 k, Flows:  324.0 k, Anomalies:  387.0 k
    Anomalies                                         Count   Rate(eps) Elapsed
    First packet of TCP session not SYN             160.0 k     0.1       25s
    TCP source or destination port zero              69.0 k    14.1     6m26s
    UDP source or destination port zero              68.0 k    12.7     6m26s
    ICMP header length check failed                     318     0.1      7m6s
    IP fragment assembly timeout                     88.0 k     1.3      6m7s
    UDP header length check failed                       39     0.0     6m58s
    TCP header length check failed                       46     0.0     6m45s

10.58.255.23    ->   10.58.255.146   18m48s SYN cookie
  Bytes:  104.0 m, Packets:  421.0 k, Flows:      230, Anomalies:  124.0 k
    Anomalies                                         Count   Rate(eps) Elapsed
    TCP source or destination port zero              37.0 k     9.8     6m26s
    UDP source or destination port zero              37.0 k     8.4     6m26s
    IP fragment assembly timeout                     48.0 k     1.0      6m7s
    ICMP header length check failed                     190     0.2     6m47s
    UDP header length check failed                       29     0.0     6m51s
    TCP header length check failed                       23     0.0     6m59s

10.58.255.25    ->   10.58.255.146   18m48s SYN cookie
  Bytes:  104.0 m, Packets:  420.0 k, Flows:      232, Anomalies:  123.0 k
    Anomalies                                         Count   Rate(eps) Elapsed
    TCP source or destination port zero              37.0 k     9.8      6m26s
    UDP source or destination port zero              37.0 k     8.6      6m26s
    IP fragment assembly timeout                     48.0 k     1.5       6m7s
```

```
    ICMP header length check failed                     173     0.1      6m43s
    UDP header length check failed                       24     0.0      6m43s
    TCP header length check failed                       19     0.0      6m56s

Total IDS table entries:
87
Total failed IDS table entry insertions
0
Total number of events (closed flows and anomalies detected):
2659291
```

show services ids source-table extensive

```
user@host> show services ids source-table extensive
Interface: sp-3/2/0, Service set: ss_all_limits
IDS sorting order: Packets
Source address       Dest address      Time Flags             Application
198.51.100.4          any              2m43s                  junos-ftp
  Bytes: 5.7k, Packets: 102.0, Flows: 41.0, Anomalies: 462.0
    Anomaly description                                Count     Rate   Elapsed
    creating forward or watch flow                      41.0      8.8     2m40s
    Packet rate exceeds IDS src limit                   21.0      7.1     2m40s
    Session creation rate exceeds IDS src limit        359.0     99.7     2m39s
    TCP SYN flood attack                                41.0      1.9     1m53s


Total IDS table entries:                                     3
Total failed IDS table entry insertions                      0
Total number of events (closed flows and anomalies):        462
```

show services ids source-table extensive limit

```
user@host> show services ids source-table extensive limit 3
Interface: sp-1/3/0, Service set: null-sfw
Sorting order: Packets
Source address       Dest address   Time    Flags             Application

10.58.255.18    ->             any   40m 0s SYN cookie
  Bytes:  250.0 m, Packets: 1978.0 k, Flows:  356.0 k, Anomalies:  387.0 k
    Anomalies                                        Count   Rate(eps) Elapsed
```

```
    TCP source or destination port zero              37.0 k     9.8     6m26s
    First packet of TCP session not SYN             160.0 k     0.0       40s
    TCP source or destination port zero              69.0 k    62.5     7m45s
    UDP source or destination port zero              68.0 k    56.2     7m45s
    ICMP header length check failed                     319     0.1     7m49s
    IP fragment assembly timeout                     89.0 k     4.4     7m26s
    UDP header length check failed                       39     0.0     8m17s
    TCP header length check failed                       46     0.0      8m4s

10.58.255.30    ->             any   20m 7s SYN cookie
  Bytes:  107.0 m, Packets:  427.0 k, Flows:      264, Anomalies:  125.0 k
    Anomalies                                        Count   Rate(eps) Elapsed
    UDP source or destination port zero              38.0 k    65.5     7m45s
    TCP source or destination port zero              37.0 k    38.1     7m45s
    IP fragment assembly timeout                     49.0 k     4.1     7m26s
    TCP header length check failed                       24     0.0     9m23s
    ICMP header length check failed                     165     0.1      8m6s
    UDP header length check failed                       26     0.0     8m13s

10.58.255.17    ->             any   20m10s SYN cookie
  Bytes:  107.0 m, Packets:  426.0 k, Flows:      262, Anomalies:  125.0 k
    Anomalies                                        Count   Rate(eps) Elapsed
    TCP source or destination port zero              38.0 k    55.      7m45s
    UDP source or destination port zero              38.0 k    55.1     7m45s
    ICMP header length check failed                     147    0.1      7m50s
    IP fragment assembly timeout                     49.0 k    2.8      7m26s
    TCP header length check failed                       22    0.0      9m33s
    UDP header length check failed                       22    0.0       8m1s
Total IDS table entries:
87
Total failed IDS table entry insertions
0
Total number of events (closed flows and anomalies detected):
2691423
Interface: sp-1/3/0, Service set: blue
NAT pool                  Address                   Port         Ports in use
d2-pool             10.59.16.100-10.59.16.100     4000-4002          1
```

Release Information

Command introduced before Junos OS Release 7.4.

show services inband-flow-telemetry

IN THIS SECTION

- Syntax | 25574
- Description | 25574
- Options | 25574
- Required Privilege Level | 25574
- Output Fields | 25575
- Sample Output | 25576
- Release Information | 25576

Syntax

```
show services inband-flow-telemetry (global | profile | stats)
```

Description

Display Inband Flow Analyzer (IFA) 2.0 statistics.

Options

global Display information about IFA global parameters configured in the IFA node. The IFA global parameters include the global device ID, metadata stack length, and hop limit.

profile Display information about the IFA profile configured in the IFA node. The IFA profile configuration parameters include the profile name, sample rate, and so on.

stats Display information about the number of IFA probe copies that traversed the IFA initiator, transit, and terminating nodes.

Required Privilege Level

view

Output Fields

Table 1412 on page 25575 lists the output fields for the `show services inband-flow-telemetry` command.

Table 1412: show services inband-flow-telemetry Output Fields

Field Name	Field Description	Level of Output
`Global Device ID`	Global device identifier, which is a unique identifier configured in the IFA node.	global
`Meta-data Stack Length`	Metadata stack length configured in the IFA node.	global
`Hop Limit`	Maximum number of hops in an IFA zone.	global
`Profile Name`	Name configured for the IFA profile. A profile name uniquely identifies each profile.	profile
`Sample rate`	Sample rate configured on the IFA node.	profile
`Source Address`	IFA terminating node's IPv4 source address.	profile
`Destination Address`	Collector's IPv4 destination address.	profile
`Destination Port`	Collector's destination port number.	profile
`IFA Init Packets`	Number of IFA probes initiated from the device.	stats
`IFA Transit Packets`	Number of IFA probes transited.	stats
`IFA Terminate Rx Packets`	Number of IFA probes received at the terminating node.	stats
`IFA Terminate Tx Packets`	Number of IFA copies transmitted to the collector application.	stats

Sample Output

show services inband-flow-telemetry global (QFX5120-48Y and QFX5120-32C)

```
user@host> show services inband-flow-telemetry global
Global Device ID            : 20
Meta-data Stack Length      : 160
Hop Limit                   : 255
```

show services inband-flow-telemetry profile (QFX5120-48Y and QFX5120-32C)

```
user@host> show services inband-flow-telemetry profile
Profile Name            : p1
Sample rate             : 1
Source Address          : 10.80.1.1
Destination Address     : 10.80.1.2
Destination Port        : 2055
```

show services inband-flow-telemetry stats (QFX5120-48Y and QFX5120-32C)

```
user@host> show services inband-flow-telemetry stats
IFA Init Packets            : 10
IFA Transit Packets         : 20
IFA Terminate Rx Packets    : 30
IFA Terminate Tx Packets    : 30
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

clear inband-flow-telemetry stats | 3635

inband-flow-telemetry | 8512

Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Performance Monitoring

show services inline ip-reassembly statistics

IN THIS SECTION

- Syntax | 25577
- Description | 25577
- Options | 25577
- Required Privilege Level | 25578
- Output Fields | 25578
- Sample Output | 25583
- Sample Output | 25585
- Release Information | 25586

Syntax

```
show services inline ip-reassembly statistics
<fpc fpc-slot>
<pfe pfe-slot>
```

Description

Display the inline IP reassembly statistics for the Packet Forwarding Engines on one or more MPCs or Next Gen Services MX-SPC3 services card. Inline IP reassembly statistics are collected at the Packet Forwarding Engine level.

NOTE: For more information on MPCs that support inline IP reassembly, refer to Protocols and Applications Supported on the MPC1E for MX Series Routers.

Options

none Displays standard inline IP reassembly statistics for all MPCs or MX-SPC3 services card.

fpc *fpc* (Optional) Displays inline IP reassembly statistics for the specified MPC or MX-SPC3 services card.

> **NOTE**: Starting with Junos OS Release 14.2, the FPC option is not displayed for MX Series routers that do not contain switch fabrics, such as MX80 and MX104 routers.

pfe *pfe* (Optional) Displays inline IP reassembly for the specified Packet Forwarding Engine slot. You must specify an FPC slot number before specifying a Packet Forwarding Engine slot.

Required Privilege Level

view

Output Fields

Table 1413 on page 25578 lists the output fields for the `show services inline ip-reassembly statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1413: show services inline ip-reassembly statistics Output Fields

Field Name	Field Description
FPC	MPC or MX-SPC3 services card slot number for which the statistics are displayed.
PFE	Packet Forwarding Engine on the MPC or MX-SPC3 services card for which the statistics are displayed.

> **NOTE**: The output fields displayed (per Packet Forwarding Engine) are arranged in a logical sequence from top to bottom to enable users to understand how the inline IP reassembly statistics are gathered.
>
> The information about total number of fragments received is displayed first, and then the information about the reassembled packets and those pending reassembly are displayed. Then, the reasons why the fragments were dropped or not reassembled are displayed. Finally, the information about the fragments reassembled, fragments dropped, and fragments sent to the backup user plane PIC (services PIC) are displayed.

Table 1413: show services inline ip-reassembly statistics Output Fields *(Continued)*

Field Name	Field Description
`Total Fragments Received`	Total number of fragments received and the current rate of fragments received for inline IP reassembly. The following information is also displayed: • `First Fragments`—Number of first fragments received and current rate of first fragments processed. • `Intermediate Fragments`—Number of intermediate fragments received and current rate of intermediate fragments processed. • `Last Fragments`—Number and rate of last fragments received. **NOTE**: Current rate refers to the current number of fragments processed per second in the instant preceding the command's execution.
`Total Packets Reassembled`	Total number of packets reassembled and current rate, in the instant preceding the command's execution, at which the packets are reassembled.
Packets Fully Reassembled	Total number of packets fully reassembled.
Packets Partially Reassembled	Total number of packets partially reassembled.
`Approximate Packets Pending Reassembly`	Approximate number of packets pending reassembly.

Table 1413: show services inline ip-reassembly statistics Output Fields *(Continued)*

Field Name	Field Description
`Fragments Dropped Reasons`	Total number of fragments dropped reasons and the current rate of total fragment dropped reasons. The number of dropped reasons and rate corresponding to each of the following reasons are also displayed: • Buffers not available • Fragments per packet exceeded • Packet length exceeded • Record insert error (Account errors caused while trying to add duplicate entry or when hash bucket is full.) • Record in use error (Pre-processing errors and count of any new fragment hash lookup results leading to existing fragments with, "Marked for Delete" or "Reassem in Progress".) • Duplicate first fragments • Duplicate last fragments • Missing first fragment **NOTE**: • These fields indicate *why* a fragment was dropped. When a fragment is dropped, the corresponding reason field is incremented by 1. For example, when a fragment is dropped because the memory runs out, the `Buffers not available` field increases by 1. • The maximum number of fragments allowed for reassembly is 16. If the interface encounters a 17th fragment, it drops the entire packet and increments the `Fragment per packet exceeded` field by 17. • Current rate refers to the current number of fragment dropped reasons per second in the instant preceding the command's execution.

Table 1413: show services inline ip-reassembly statistics Output Fields *(Continued)*

Field Name	Field Description
`Reassembly Errors Reasons`	Number of errors during reassembly and the current rate of reassembly errors. The number of errors and the rate for each of the following types of errors are also displayed: • Fragment not found • Fragment not in sequence • ASIC errors **NOTE**: Current rate refers to the current number of reassembly errors processed per second in the instant preceding the command's execution.
`Aged out packets`	Number of aged out packets and the current number of packets aged out per second in the instant preceding the command's execution. **NOTE**: In some cases, aged out packets can refer to aged out fragments. If previous fragments of the packet have already been discarded then linking of the dropped fragments to the aged out fragments cannot occur.
`Total Fragments Successfully Reassembled`	Number of fragments successfully reassembled and the current number of fragments reassembled per second in the instant preceding the command's execution.

Table 1413: show services inline ip-reassembly statistics Output Fields *(Continued)*

Field Name	Field Description
`Total Fragments Dropped`	Total number of fragments dropped and the current rate of total number of fragments dropped. The number of fragments dropped and rate corresponding to each of the following reasons are also displayed: • Buffers not available • Fragments per packet exceeded • Packet length exceeded • Record insert error • Record in use error • Duplicate first fragments • Duplicate last fragments • Missing first fragment • Fragment not found • Fragment not in sequence • ASIC errors • Aged out fragments
`Total fragments punted to UPIC`	Number of fragments sent to the backup user plane PIC (services PIC) and current rate of fragments sent per second in the instant preceding the command's execution

The following information applies to the `Total Fragments Dropped` field.

- These fields indicate *how many* of the packet fragments received were then dropped due to a particular reason.

  For example, consider a packet that has 10 fragments, 9 of which have been received and stored in memory. When the tenth fragment arrives, if the memory runs out (Buffers not available), then this fragment is dropped. Because the tenth fragment has been dropped, the other 9 fragments must also be dropped. In this case, the `Buffers not available` field (under the `Fragments Dropped Reasons` field) is

incremented by 1 and the `Buffers not available` field (under the `Total Fragments Dropped` field) is incremented by 10.

For the next packet arriving, which also has 10 fragments, the first four fragments are stored but the memory runs out for the fifth fragment. Then the first 5 fragments (fifth and the first four) are dropped. In this case, the `Buffers not available` field (under the `Fragments Dropped Reasons` field) is incremented by 1 and the `Buffers not available` field (under the `Total Fragments Dropped` field) is incremented by 5.

For fragments of the packet, if memory becomes available, the next 5 fragments (6 through 10) that arrive are stored in memory. The fragments are stored until the timeout period elapses, and are eventually dropped. In this case, the `Aged out packets` field is incremented by 1 and the `Aged out fragments` field (under the `Total Fragments Dropped` field) is incremented by 5.

The fragment counters (after both packets have been processed) are as follows:

- `Fragments Dropped Reasons`
  - `Buffers not available` 2
  - `Aged out packets` 1
- `Total Fragment Dropped`
  - `Buffers not available` 15
  - `Aged out packets` 5

- Current rate refers to the current total number fragments dropped per second in the instant preceding the command's execution.

Sample Output

show services inline ip-reassembly statistics fpc

```
user@host> show services inline ip-reassembly statistics fpc 0
FPC: 0 PFE: 0
=============
                                                  Total   Current Rate
Total Fragments Received                      728177644          83529
   First Fragments                            260759430          29924
   Intermediate Fragments                     206658784          23681
   Last Fragments                             260759430          29924

Total Packets Successfully Reassembled        260746982          29924
```

```
Approximate Packets Pending Reassembly                4

Fragments Dropped Reasons                          34558              3
   Buffers not available                               0              0
   Fragments per packet exceeded                       0              0
   Packet length exceeded                              0              0
   Record insert error                                 0              0
   Record in use error                             34558              3
   Duplicate first fragments                           0              0
   Duplicate last fragments                            0              0
   Missing first fragment                              0              0

Reassembly Errors Reasons                              0              0
   Fragment not found                                  0              0
   Fragment not in sequence                            0              0
   ASIC errors                                         0              0

Aged out packets                                      63              0

Total Fragments Successfully Reassembled       728142977          83528

Total Fragments Dropped                            34673              3
   Buffers not available                               0              0
   Fragments per packet exceeded                       0              0
   Packet length exceeded                              0              0
   Record insert error                                 0              0
   Record in use error                             34558              3
   Duplicate first fragments                           0              0
   Duplicate last fragments                            0              0
   Missing first fragment                              0              0
   Fragment not found                                  0              0
   Fragment not in sequence                            0              0
   ASIC errors                                         0              0
   Aged out fragments                                115              0

Total fragments punted to UPIC                         0              0
```

When partial reassembly of IPv4 packets for MAP-E is enabled the output is enhanced to display `Total Packets Successfully Reassembled` which includes `Packets Fully Reassembled` and `Packets Partially Reassembled`.

Sample Output

show services inline ip-reassembly statistics fpc

```
user@host> show services inline ip-reassembly statistics fpc 2 pfe-slot 0 FPC: 2 PFE: 0
=============
                                                    Total   Current Rate
Total Fragments Received                                0              0
   First Fragments                                      0              0
   Intermediate Fragments                               0              0
   Last Fragments                                       0              0

Total Packets Successfully Reassembled                  0              0
   Packets Partially Reassembled                        0              0

Total Fragments Successfully Reassembled                0              0

Approximate Packets Pending Reassembly                  0

Fragments Dropped Reasons                               0              0
   Buffers not available                                0              0
   Fragments per packet exceeded                        0              0
   Packet length exceeded                               0              0
   Record insert error                                  0              0
   Record in use error                                  0              0
   Duplicate first fragments                            0              0
   Duplicate last fragments                             0              0
   Missing first fragment                               0              0

Reassembly Errors Reasons                               0              0
   Fragment not found                                   0              0
   ASIC errors                                          0              0

Aged out packets                                        0              0

Total Fragments Dropped                                 0              0
   Buffers not available                                0              0
   Fragments per packet exceeded                        0              0
   Packet length exceeded                               0              0
   Record insert error                                  0              0
   Record in use error                                  0              0
   Duplicate first fragments                            0              0
```

```
    Duplicate last fragments                          0              0
    Missing first fragment                            0              0
    Fragment not found                                0              0
    ASIC errors                                       0              0
    Aged out fragments                                0              0
```

Release Information

Statement introduced in Junos OS Release 12.2X49.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

ip-reassembly | 9262

show services inline-monitoring feature-profile-mapping fpc-slot

IN THIS SECTION

- Syntax | 25586
- Description | 25587
- Options | 25587
- Required Privilege Level | 25587
- Output Fields | 25587
- Sample Output | 25588
- Release Information | 25588

Syntax

```
show services inline-monitoring feature-profile-mapping fpc-slot slot-number
```

Description

(EX4100 and EX4100-F switches only) Show what features you configured for flow-based telemetry and how they are ordered in the flow.

Options

fpc-slot *slot-number*

Specify the slot number to view the feature mapping for that Flexible PIC Concentrator (FPC). You must have already configured a feature profile for inline monitoring for this command to contain any information.

Required Privilege Level

view

Output Fields

Table 1414 on page 25587 lists the output fields for the `show services inline-monitoring feature-profile-mapping` command.

Table 1414: show services inline-monitoring feature-profile-mapping Output Fields

Field Name	Field Description
`FPC Slot`	Slot configured for flow-based telemetry.
`Flow counter Instance Name`	Name of the inline-monitoring instance you configured for flow-based telemetry.
`EnterpriseElemId`	Information Element number for a feature you configured for the inline-monitoring feature profile.
`Feature/Counter Name`	Name of a feature you configured in an inline-monitoring feature profile.

Sample Output

show services inline-monitoring feature-profile-mapping fpc-slot 0

```
user@host> show services inline-monitoring feature-profile-mapping fpc-slot 0

FPC Slot: 0, Flow counter Instance Name: i1
-------------------------------------------------------------------+
EnterpriseElemId   |            Feature/Counter Name
-------------------------------------------------------------------+

       34                       Inter-arrival-time(4bytes)
       35                       Inter-departure-delay  (4bytes)
       36                       Chip-delay  (4 bytes)
       37                       shared-pool-congestion-level(1 byte)
       38                       Queue-congestion-level (1 byte)
       16                       Dos-attack( 4 bytes)
       40                       Ingress-drop-reason(2 bytes)
       41                       Egress-drop-reason(2 bytes)
       42                       aggregate-intf-member-id(1 byte)
```

Release Information

Command introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Flow-Based Telemetry

show services inline-monitoring statistics fpc-slot

IN THIS SECTION

- Syntax | 25589
- Description | 25589

- Options | 25589
- Required Privilege Level | 25589
- Output Fields | 25589
- Sample Output | 25590
- Release Information | 25591

Syntax

```
show services inline-monitoring statistics fpc-slot fpc-slot
```

Description

Display inline-monitoring services statistical information, including for features that use inline-monitoring services, such as flow-based telemetry.

Options

fpc-slot Display inline-monitoring statistical information for the specified Flexible PIC Concentrator (FPC) number.

- **Range:** 0 through 11

Required Privilege Level

view

Output Fields

Table 1415 on page 25590 lists the output fields for the `show services inline-monitoring statistics fpc-slot` command. Output fields are listed in the approximate order in which they appear.

Table 1415: show services inline-monitoring statistics fpc-slot Output Fields

Field Name	Field Description
FPC Slot	Flexible PIC Concentrator (FPC) slot level counters: • Packets—Number of packets serviced under FPC slot. • Bytes—Number of bytes serviced under FPC slot.
Instance Name	Instance level counters: • Packets—Number of packets serviced under instance number under the FPC slot. • Bytes—Number of bytes serviced under instance number under the FPC slot.
Collector Name	Collector level counters: • Packets—Number of packets serviced under collector number of an instance under the FPC slot. • Bytes—Number of bytes serviced under collector number of an instance under the FPC slot.

Sample Output

show services inline-monitoring services fpc-slot

```
user@host> show services inline-monitoring services fpc-slot fpc-slot
IMON Statistics
   FPC Slot     :  <FPC_SLOT>
   Packets      :  <F_P>       Bytes  :  <F_B>

     Instance Name  :  <I1>
     Packets  :  <I1_P>      Bytes  :  <I1_B>

            Collector Name  :  <I1_C1>
            Packets  :  <I1_C1_P>    Bytes  :  <I1_C1_B>

            Collector Name  :  <I1_C2>
            Packets  :  <I1_C2_P>    Bytes  :  <I1_C2_B>
```

```
    Instance Name  :  <I2>
    Packets  :  <I2_P>      Bytes  :  <I2_B>

           Collector Name  :  <I2_C1>
           Packets  :  <I2_C1_P>    Bytes  :  <I2_C1_B>
```

show services inline-monitoring services fpc-slot (flow-based telemetry)

```
user@host> show services inline-monitoring statistics fpc-slot 0
  IMON Statistics
    FPC Slot: 0
        Instance Name: i1
        Flow exceeded count: 0
        Flow missed count: 0
        Flow aged out count: 1
        Flow learnt count: 3
        Flow meter exceeded count: 0
```

Release Information

Command introduced in Junos OS Release 19.4R1.

Command introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

Understanding Inline Monitoring Services

Flow-Based Telemetry

show services inline nat pool

IN THIS SECTION

Syntax | 25592

- Description | 25592
- Options | 25592
- Required Privilege Level | 25592
- Output Fields | 25592
- Sample Output | 25593
- Release Information | 25594

Syntax

```
show services inline nat pool
<pool pool--name>
```

Description

Display information about inline Network Address Translation (NAT) pool.

Options

pool-name Display information about the specified services-inline interface NAT pool.

Required Privilege Level

view

Output Fields

Table 1416 on page 25593 lists the output fields for the `show services inline nat pool` command. Output fields are listed in the order in which they appear.

Table 1416: show services inline nat pool Output Fields

Field Name	Field Description
Interface	Name of an **si** interace hosted on a Trio-based line card.
NAT pool	Name of the pool used for address translations.
Translation type	Translation type specified in the applicable NAT rule for the service set.
Address range	Starting and ending public NAT addresses available for translation.
NATed packets	Number of packets translated for the specified pool.
un-NATed packets	Number of received packets that were not translated.
`deNATed packets`	Number of packets that were not translated for the specified service PIC.
Errors	Number of packets with translation errors.

Sample Output

show services inline nat pool

```
user@host> show services inline nat pool p1
Interface: si-5/0/0, Service set: ss-inat
  NAT pool: p1, Translation type: BASIC NAT44
    Address range: 192.0.2.0-192.0.2.255
    NATed packets: 0, Un-NATed packets: 0, Errors: 0
```

show services inline nat pool (Network Prefix Translation for IPv6)

```
user@host> show services inline nat pool ss_nptv6_pool1

     Interface: si-4/0/0, Service set: ss_nptv6
```

```
            NAT pool: ss_nptv6_pool1, Translation type: NPTV6
                 Address range: 2001:db8:3456::/48
                 NATed packets: 0, deNATed packets: 0, Errors: 0
```

Release Information

Command introduced in Junos OS Release 11.4.

show services inline nat statistics

IN THIS SECTION

- Syntax | 25594
- Description | 25594
- Options | 25595
- Required Privilege Level | 25595
- Output Fields | 25595
- Sample Output | 25596
- Release Information | 25597

Syntax

```
show services inline nat statistics
<interface interface-name>
```

Description

Display information about inline Network Address Translation (NAT) address translations.

Options

interface-name (Optional) Display information about the specified NAT services-inline interface only. When a specific interface is not specified, statistics for all services-inline interfaces are shown.

Required Privilege Level

view

Output Fields

Table 1417 on page 25595 lists the output fields for the `show services inline nat statistics` command. Output fields are listed in the order in which they appear.

Table 1417: show services inline nat statistics Output Fields

Field Name	Field Description	Level of Output
Service PIC	Name of an **si** interface hosted on a Trio-based line card.	All levels
Slow path packets received	Number of ICMP exception packets received for NAT translation.	All levels
Slow path packets dropped	Number of received ICMP exception packets that were dropped.	All levels
`Service PIC Name`	FPC and PIC slots for the service PIC on which NAT processing is performed	All levels
`Data Plane Statistics`	Information about packets processed by the data plane for NAT operations	All levels
`Control Plane Statistics`	Information about packets processed by the control plane for NAT operations	All levels

Table 1417: show services inline nat statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`ICMPv4 errors packets pass through`	Number of ICMPv4 error packets that were passed through without being subjected to rules	All levels
`ICMPv4 errors packets locally generated`	Number of ICMPv4 error packets that were locally generated	All levels
`ICMPv6 errors packets pass through`	Number of ICMPv6 error packets that were passed through without being subjected to rules	All levels
`ICMPv6 errors packets locally generated`	Number of ICMPv6 error packets that were locally generated	All levels
`Dropped packets`	Number of packets dropped during inline NAT processing	All levels
`NATed packets`	Number of packets translated for the specified service PIC.	All levels
`deNATed packets`	Number of packets that were not translated for the specified service PIC.	All levels
`Errors`	Number of packets with translation errors.	All levels

Sample Output

show services inline nat statistics

```
user@host> show services inline nat statistics
Service PIC Name                                    :si-5/0/0
```

```
     Slow path packets received                          :0
     Slow path packets dropped                           :0
```

show services inline nat statistics (Network Prefix Translation for IPv6)

```
user@host> show services inline nat statistics
Service PIC Name                                                    :si-4/0/0

    Control Plane Statistics
        ICMPv4 errors packets pass through                                  :0
        ICMPv4 errors packets locally generated                             :0
        ICMPv6 errors packets pass through                                  :0
        ICMPv6 errors packets locally generated                             :0
        Dropped packets                                                     :0

     Data Plane Statistics
        NATed packets                                                       :0
        deNATed packets                                                     :0
        Errors                                                              :0

Service PIC Name                                                         :si-4/1/0

    Control Plane Statistics
        ICMPv4 errors packets pass through                                  :0
        ICMPv4 errors packets locally generated                             :0
        ICMPv6 errors packets pass through                                  :0
        ICMPv6 errors packets locally generated                             :0
        Dropped packets                                                     :0

     Data Plane Statistics
        NATed packets                                                       :0
        deNATed packets                                                     :0
        Errors                                                              :0
```

Release Information

Command introduced in Junos OS Release 11.4.

show services inline softwire statistics

IN THIS SECTION

- Syntax | 25598
- Description | 25598
- Options | 25598
- Required Privilege Level | 25599
- Output Fields | 25599
- Sample Output | 25601
- Release Information | 25604

Syntax

```
show services inline softwire statistics
<interface interface-name>
<mape name>
<mapt name>
<v6rd>
```

Description

Display information about inline softwire activity.

Options

interface *interface-name*	(Optional) Display information about the specified services-inline interface only. When a specific interface is not specified, statistics for all services-inline interfaces are shown.
mape *name*	(Optional) Display map-e information on per physical service interface basis.
mapt *name*	(Optional) Display map-t information on per physical service interface basis.
v6rd	(Optional) Display information for 6rd.

Required Privilege Level

view

Output Fields

Table 1418 on page 25599 lists the output fields for the `show services inline softwire statistics` command. Output fields are listed in the order in which they appear.

Table 1418: show services inline softwire statistics Output Fields

Field Name	Field Description
`Service PIC Name`	Name of the service PIC for which statistics are displayed.
`Control Plane Statistics`	Statistics on the control plane.
`ICMPv4 echo requests to softwire concentrator`	Number of ICMPv4 echo received by the softwire concentrator. IPv6 ICMP type = 128, code =0. destined to BR IPv6 address
`ICMPv4 echo responses from softwire concentrator`	Number of ICMPv4 echo responses sent from the softwire concentrator or BR. IPv6 ICMP type = 129
`Dropped ICMPv4 packets to softwire concentrator`	Number of ICMP packets (except ICMP request) received by the softwire concentrator or BR. All these packets are dropped in by the packet forwarding engine Ukernel.
`Trace route UDP packets to softwire concentrator`	Number of UDP trace route packets (port numbers 33434 through 33534) received by the softwire concentrator.
`ICMPv4 Port unreachable errors sent from softwire concentrator`	Number of ICMP port unreachable errors sent by the softwire concentrator after receiving the UDP trace route packets.

Table 1418: show services inline softwire statistics Output Fields *(Continued)*

Field Name	Field Description
`Other dropped IPv4 packets to softwire concentrator`	Number of non-ICMPpackets that were received and dropped because of fragmentation during encapsulation or decapsulation.
`Data Plane Statistics`	Statistics of the data plane.
`6rd decaps`	Number of 6rd decapsulated packets and bytes in the data plane. Decapsulation includes removing the outer IPv4 header and routing the inner IPv6 packet.
`6rd encaps`	Number of 6rd encapsulted (IPv4) packets and bytes in the data plane.
`6rd decap errors`	Number of all the packets and bytes that are not IPv4-IPv6, IPv4-UDP, or IPV4-ICMP packets.
`6rd decap fragment errors`	Number of IPv4 fragmented packets and bytes.
`6rd decap spoof attacks`	Number of spoof attack packets and bytes, which includes packets for which the 6rd derived IPv4 address does not match with the source IPv4 address and packets for which the source IPv6 prefix does not match the 6rd IPv6 prefix.
`6rd encap v4 mtu errors`	Count of packets and bytes with IPv4 encapsulation MTU errors. For downlink packets after encapsulating with an IPv4 header, if the packet length is more than Tunnel MTU then it is dropped as v4 MTU errors. For these packet drops, an ICMPv6 packet too big error is sent back to the sender.
Data Plane Statistics (MAP-E upstream)	
MAPE decaps	IPv6 packets successfully decapsulated by BR (includes reassembled IPv6)
MAPE ICMP decap errors	IPv6 packets dropped due to unsupported type/code of inner ICMPv4

Table 1418: show services inline softwire statistics Output Fields *(Continued)*

Field Name	Field Description
MAPE decap spoof errors	IPv6 Packets that failed MAPE spoof check

Sample Output

show services inline softwire statistics

```
user@host> show services inline softwire statistics
 Border Router v6rd statistics:

     Service PIC Name                                      si-0/0/0

     Control Plane Statistics
         ICMPv4 echo requests to softwire concentrator                    0
         ICMPv4 echo responses from softwire concentrator                 0
         Dropped ICMPv4 packets to softwire concentrator                  0
         Trace route UDP packets to softwire concentrator                 0
         ICMPv4 Port unreachable errors sent from softwire concentrator   0
         Other dropped IPv4 packets to softwire concentrator              0

      Data Plane Statistics           Packets                  Bytes
     6rd decaps                    32222173891             3061106519645
     6rd encaps                    415480622               28252710148
     6rd decap errors              0                       0
     6rd decap fragment errors     0                       0
     6rd decap spoof attacks       0                       0


     Service PIC Name                                      si-0/2/0

     Control Plane Statistics
         ICMPv4 echo requests to softwire concentrator                    0
         ICMPv4 echo responses from softwire concentrator                 0
         Dropped ICMPv4 packets to softwire concentrator                  0
         Trace route UDP packets to softwire concentrator                 0
         ICMPv4 Port unreachable errors sent from softwire concentrator   0
```

```
        Other dropped IPv4 packets to softwire concentrator             0

    Data Plane Statistics           Packets                  Bytes
        6rd decaps                    0                        0
        6rd encaps                    0                        0
        6rd decap errors              0                        0
        6rd decap fragment errors     0                        0
        6rd decap spoof attacks       0                        0
        6rd encap v4 mtu errors       0                        0
```

show services inline softwire statistics mape (Adaptive Services si- interfaces)

```
user@host> show services inline softwire statistics mape
Service PIC Name                                        si-0/0/0

Statistics                        Packets                  Bytes

     MAP-E     decaps                        0                        0
     MAP-E     encaps                        0                        0
     MAP-E     decap errors                  0                        0
     MAP-E     encap errors                  0                        0
     MAP-E decap spoof attacks             0                        0
     MAP-E decap v4 fragmented            0                        0
     MAP-E decap v4 reassembled           0                        0
     MAP-E encap v4 mtu errors             0                        0
```

show services inline softwire statistics mape (Next Gen Services si- interfaces)

```
user@host> show services inline softwire statistics mape
Service PIC Name                                        si-2/0/0

 Control Plane Statistics
     MAPE ICMPv6 echo requests to softwire concentrator          0
     MAPE ICMPv6 echo responses from softwire concentrator       0
     MAPE Dropped ICMPv6 packets to softwire concentrator        0

 Data Plane Statistics (v6-to-v4)      Packets                  Bytes
     MAPE decaps                        0                        0
     MAPE ICMP decap errors             0                        0
```

```
    MAPE decap spoof errors           0                       0
    MAPE v6 reassembled               0                       0
    MAPE dropped v6 fragments         0                       0
    MAPE v6 unsupp protocol drops     0                       0


 Data Plane Statistics (v4-to-v6)     Packets                 Bytes
    MAPE encaps                       0                       0
    MAPE ICMP encap errors            0                       0
    MAPE v6 mtu errors                0                       0
    MAPE v4 reassembled               0                       0
    MAPE dropped v4 fragments         0                       0
```

show services inline softwire statistics mapt (Adaptive Services si- interfaces)

```
user@host> show services inline softwire statistics mapt
              Service PIC Name                                    si-2/0/0

 Control Plane Statistics
    MAPT ICMPv4 translated to ICMPv6                           0
    MAPT ICMPv6 translated to ICMPv4                           0
    MAPT Dropped ICMPv4 packets                                0
    MAPT Dropped ICMPv6 packets                                0

 Data Plane Statistics (v6-to-v4)     Packets                 Bytes
    MAPT v6 translated to v4          0                       0
    MAPT spoof errors                 0                       0
    MAPT v6 unsupported discards      0                       0

 Data Plane Statistics (v4-to-v6)     Packets                 Bytes
    MAPT v4 translated to v6          0                       0
    MAPT BR MTU errors                0                       0
    MAPT v4 unsupported discards      0                       0
```

show services inline softwire statistics mapt (Full reassembly support)

```
user@host> show services inline softwire statistics mapt
              Service PIC Name                                    si-0/0/0
```

```
Control Plane Statistics
    MAPT ICMPv6 translated to ICMPv4                          0
    MAPT ICMPv4 translated to ICMPv6                          0
    MAPT ICMPv4 discards                                      0
    MAPT ICMPv6 discards                                      0

Data Plane Statistics (v6-to-v4)      Packets                 Bytes
    MAPT v6 translated to v4           0                      0
    MAPT v6 spoof errors               0                      0
    MAPT v6 fragment drops             0                      0
    MAPT v6 unsupported drops          0                      0
    MAPT v6 reassembled                0                      0

Data Plane Statistics (v4-to-v6)      Packets                 Bytes
    MAPT v4 translated to v6           0                      0
    MAPT v6 MTU exceed drops           0                      0
    MAPT v4 fragment drops             0                      0
    MAPT v4 unsupported drops          0                      0
    MAPT v4 reassembled                0                      0
```

hemapriya, removed sample output for MX-SPC3 as per mdenny's inputs.

Release Information

Command introduced in Junos OS Release 13.3R3.

`map-e` option introduced in Junos OS Release 18.2R1 for MX Series Routers with MPC and MIC interfaces.

`map-e` option introduced in Junos OS Release 20.2R1 for Next Gen Services on MX240, MX480 and MX960 routers.

show services ip-monitoring status

IN THIS SECTION

Syntax | 25605

- Description | 25605
- Required Privilege Level | 25605
- Output Fields | 25605
- Sample Output | 25607
- Sample Output | 25607
- Sample Output | 25608
- Sample Output | 25609
- Sample Output | 25609
- Sample Output | 25610
- Release Information | 25610

Syntax

```
show services ip-monitoring status
```

Description

Display a brief summary of IP monitoring status along with the current state for a given policy.

Required Privilege Level

view

Output Fields

Table 1419 on page 25605 lists the output fields for the `show services ip-monitoring status` command. Output fields are listed in the approximate order in which they appear.

Table 1419: show services ip-monitoring status Output Fields

Field Name	Field Description
`Policy`	Name of the policy configured.

Table 1419: show services ip-monitoring status Output Fields *(Continued)*

Field Name	Field Description
`Probe Name`	Name of the probe configured.
`Address`	Displays the configured target address.
`Status`	Displays the status of the probe on the target address. If the status is PASS, then the target address is reached.
`Route-Action`	Displays route injection information configured for the policy and its failover status.
`Route-Instance`	Displays the routing instance of the route to be injected during failover.
`Route`	Routing address of the route to be injected during failover.
`Next-Hop`	Specifies the next-hop address of the route to be injected during failover. P2P interfaces only.
`State`	Display the state of the route injection action. If the state is APPLIED, then the ip-monitoring policy is in failover state.
`Interface Action`	Displays the interface action type as enable or disable.
`Policy Action`	Displays the policy action type as enable or disable.
`Admin State`	Displays the current admin state of the interface.
`Action Status`	Displays the current action status of the interface.

Sample Output

show services ip-monitoring status

```
user@host> show services ip-monitoring status

Policy - policy1 (Non-preemptive. Status: FAIL)
  RPM Probes:
    Probe name             Test Name       Address          Status
    ---------------------- --------------- ---------------- ---------
    probe_a                a1              15.1.1.10        FAIL
    probe_a                a2              200.1.1.1        FAIL
  Route-Action:
    route-instance    route             next-hop         State
    ----------------- ----------------- ---------------- -------------
    inet.0            200.1.1.0         150.1.1.1        APPLIED
  Interface-Action:
    interface         policy action   admin state action status
    ----------------- --------------- ----------- -----------------
    fe-0/0/5.2        Enable          UP          FAILOVER
    fe-0/0/5.4        Disable         DOWN        FAILOVER
    t1-1/0/0          Enable          UP          FAILOVER
    dl0               Enable          UP          FAILOVER
    ge-0/0/1          Enable          UP          FAILOVER
```

Sample Output

show services ip-monitoring status

In this example, the policy is in the failback state, and the no-preempt option is not configured.

```
user@host> show services ip-monitoring status

Policy - policy1 (Status: PASS)
  RPM Probes:
    Probe name             Test Name       Address          Status
    ---------------------- --------------- ---------------- ---------
    probe1                 a1              99.1.1.2         PASS
```

```
  Route-Action:
    route-instance    route             next-hop         state
    ----------------- ----------------- ---------------- -------------
    inet.0            99.1.1.0          12.12.12.2       NOT-APPLIED
  Interface-Action:
    interface         policy action   admin state action status
    ----------------- --------------- ----------- -----------------
    at-2/0/0          Enable          DOWN        MARKED-DOWN
    ge-0/0/2.2        Enable          DOWN        MARKED-DOWN
    ge-0/0/2.3        Enable          DOWN        MARKED-DOWN
```

Sample Output

show services ip-monitoring status

In this example, the policy is in the failover state, and the primary is restored. The no-preempt option is configured.

```
user@host> show services ip-monitoring status

Policy - policy1 (Non-preemptive. Status: FAILOVER-NO-PREEMPT)
  RPM Probes:
     Probe name             Test Name       Address          Status
    ---------------------- --------------- ---------------- ---------
     probe1                 a1              99.1.1.2         PASS
  Route-Action:
    route-instance    route             next-hop         state
    ----------------- ----------------- ---------------- -------------
    inet.0            99.1.1.0          12.12.12.2       APPLIED
  Interface-Action:
    interface         policy action   admin state action status
    ----------------- --------------- ----------- -----------------
    at-2/0/0          Enable          UP          FAILOVER
    ge-0/0/2.2        Enable          UP          FAILOVER
    ge-0/0/2.3        Enable          UP          FAILOVER
```

Sample Output

show services ip-monitoring status

When the probe succeeds and the policy is not applied, the output is as follows:

```
user@host> show services ip-monitoring status
Policy payment (Status: PASS)
 RPM Probes:
     Probe name             Test Name       Address          Status
    ---------------------- --------------- ---------------- ---------
     Probe-Payment-Server   paysvr          9.9.9.2          PASS
 Route-Action:
    route-instance    route             next-hop         state
    ----------------- ----------------- ---------------- -------------
    inet.0            9.9.9.0/24        e1-6/0/0.0       NOT-APPLIED
```

Sample Output

show services ip-monitoring status

When the probe fails and the policy is applied, the output is as follows:

```
user@host> show services ip-monitoring status
Policy payment (Status: FAIL)
 RPM Probes:
     Probe name             Test Name       Address          Status
    ---------------------- --------------- ---------------- ---------
     Probe-Payment-Server   paysvr          9.9.9.2          FAIL
 Route-Action:
    route-instance    route             next-hop         state
    ----------------- ----------------- ---------------- -------------
    inet.0            9.9.9.0/24        e1-6/0/0.0       APPLIED
```

Sample Output

show services ip-monitoring status

When the probe fails and the policy is applied, the output is as follows. The status of the route `WITHDRAWED` indicates that this route is removed.

```
user@host> show services ip-monitoring status
Policy - p1 (Status: FAIL)
  RPM Probes:
    Probe name             Test Name       Address          Status
    ---------------------- --------------- ---------------- ---------
    p1                     t1              200.20.2.2       FAIL
  Route-Action:
    route-instance    route             next-hop         state
    ----------------- ----------------- ---------------- -------------
     inet.0         111.111.111.1/32   10.208.128.1     WITHDRAWED
```

Release Information

Command modified in Junos OS Release 11.4 R2. Next-hop functionality added in Junos OS Release 12.1X46-D15.

RELATED DOCUMENTATION

show security monitoring | 24864

show services ipsec-vpn certificates

IN THIS SECTION

- Syntax | 25611
- Description | 25611
- Options | 25611
- Required Privilege Level | 25611

- Output Fields | 25611
- Sample Output | 25613
- Release Information | 25615

Syntax

```
show services ipsec-vpn certificates
<brief | detail>
<service-set service-set>
```

Description

(Adaptive services interfaces only) Display local and remote certificates installed in the IPsec configuration memory cache that are used for the IKE negotiation.

Options

none	(same as brief) Display information about local and remote certificates associated with all service sets.
brief | detail	(Optional) Display the specified level of output.
service-set *service-set*	(Optional) Display information about local and remote certificates associated with only the specified service set.

Required Privilege Level

view

Output Fields

Table 1420 on page 25612 lists the output fields for the `show services ipsec-vpn certificates` command. Output fields are listed in the approximate order in which they appear.

Table 1420: show services ipsec-vpn certificates Output Fields

Field Name	Field Description	Level of Output
Service set	Name of the IPsec service set.	All levels
Total entries	Number of certificate cache entries.	All levels
Certificate cache entry	Identification number of the certificate cache entry.	All levels
Flags	Information about the digital certificate, including whether the certificate is a root certificate and trusted.	none **brief**
Issued to	Device that was issued the digital certificate.	none **brief**
Issued by	Authority that issued the digital certificate.	none **brief**
Certificate version	Revision number of the digital certificate.	**detail**
Serial number	Unique serial number of the digital certificate.	**detail**
Alternate subject	Domain name or IP address of the device related to the digital certificate.	All levels
Validity	Time period when the digital certificate is valid. Values are: • **Not before**—Start time when the digital certificate becomes valid. • **Not after**—End time when the digital certificate becomes invalid.	none **brief**
Public key algorithm	Specifies the encryption algorithm used with the private key, such as **rsaEncryption (1024 bits)**.	**detail**
Signature algorithm	Encryption algorithm that the CA used to sign the digital certificate, such as **sha1WithRSAEncryption**.	**detail**

Table 1420: show services ipsec-vpn certificates Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Fingerprint	Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate.	**detail**
Distribution CRL	Distinguished name information and the URL for the certificate revocation list (CRL) server.	**detail**
Use for key	Use of the public key, such as **Certificate signing**, **CRL signing**, **Digital signature**, or **Key encipherment**.	**detail**

Sample Output

show services ipsec-vpn certificates

```
user@host> show services ipsec-vpn certificates
Service set: serviceset-dynamic-BiEspsha3des, Total entries: 3
  Certificate cache entry: 3
    Flags: Non-root Trusted
    Issued to: router3.example.com, Issued by: juniper
    Alternate subject: router3.example.com
    Validity:
      Not before: 2005 Nov 21st, 23:33:58 GMT
      Not after: 2008 Nov 22nd, 00:03:58 GMT

 Certificate cache entry: 2
    Flags: Non-root Trusted
    Issued to: router2.example.com, Issued by: juniper
    Alternate subject: router2.example.com
    Validity:
      Not before: 2005 Nov 21st, 23:28:22 GMT
      Not after: 2008 Nov 21st, 23:58:22 GMT

  Certificate cache entry: 1
    Flags: Root Trusted
    Issued to: juniper, Issued by: juniper
    Validity:
```

```
      Not before: 2005 Oct 18th, 23:54:22 GMT
      Not after: 2025 Oct 19th, 00:24:22 GMT
```

show security ipsec-vpn certificates detail

```
user@host> show services ipsec-vpn certificates detail
Service set: serviceset-dynamic-BiEspsha3des, Total entries: 3
  Certificate cache entry: 3
    Certificate version: 3
    Serial number: 4355 94f9
    Alternate subject: router3.example.com
    Public key algorithm: rsaEncryption
    Signature algorithm: sha1WithRSAEncryption
    Fingerprint:
      61:3a:d0:b4:7a:16:9b:39:ba:81:3f:9d:ab:34:e5:c8:be:3b:a1:6d (sha1)
      60:a0:ff:58:05:4a:65:73:9d:74:3a:e1:83:6f:1b:c8 (md5)
    Distribution CRL:
      C=us, O=juniper, CN=CRL1
      http://CA-1/CRL/juniper_us_crlfile.crl
    Use for key: Digital signature

  Certificate cache entry: 2
    Certificate version: 3
    Serial number: 4355 94f8
    Alternate subject: router2.example.com
    Public key algorithm: rsaEncryption
    Signature algorithm: sha1WithRSAEncryption
    Fingerprint:
      30:c3:a4:04:da:33:9d:60:23:5a:48:75:48:2c:f0:c6:96:6c:31:fa (sha1)
      9a:a2:ce:ef:7e:10:80:a0:c8:4d:2f:e7:e1:d3:69:9d (md5)
    Distribution CRL:
      C=us, O=juniper, CN=CRL1
      http://CA-1/CRL/juniper_us_crlfile.crl
    Use for key: Digital signature

  Certificate cache entry: 1
    Certificate version: 3
    Flags: Root
    Serial number: 4355 9235
    Public key algorithm: rsaEncryption
    Signature algorithm: sha1WithRSAEncryption
```

```
    Fingerprint:
      00:8e:6f:58:dd:68:bf:25:0a:e3:f9:17:70:d6:61:f3:53:a7:79:10 (sha1)
      71:6f:6a:76:17:9b:d6:2a:e7:5a:72:97:82:6d:26:86 (md5)
    Distribution CRL:
      C=us, O=juniper, CN=CRL1
      http://CA-1/CRL/juniper_us_crlfile.crl
    Use for key: CRL signing, Certificate signing
```

Release Information

Command introduced in Junos OS Release 7.5.

show services ipsec-vpn ike security-associations

IN THIS SECTION

- Syntax | 25615
- Description | 25616
- Options | 25616
- Required Privilege Level | 25616
- Output Fields | 25616
- Sample Output | 25620
- Release Information | 25622

Syntax

```
show services ipsec-vpn ike security-associations
<brief | detail>
<peer-address>
```

Description

(Adaptive services interface only) Display information for Internet Key Exchange (IKE) security associations. If no security association is specified, the information for all security associations is displayed.

Options

none (same as brief) Display standard information for all IPsec security associations.

brief | detail (Optional) Display the specified level of output.

peer-address (Optional) Display information about a particular security association address.

Required Privilege Level

view

Output Fields

Table 1421 on page 25616 lists the output fields for the `show services ipsec-vpn ike security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 1421: show services ipsec-vpn ike security-associations Output Fields

Field Name	Field Description	Level of Output
IKE peer	Remote end of the IKE negotiation.	**detail**
Role	Part played in the IKE session. The router triggering the IKE negotiation is the initiator, and the router accepting the first IKE exchange packets is the responder.	**detail**
Remote Address	Responder's address.	none specified

Table 1421: show services ipsec-vpn ike security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the IKE security association: • **Matured**—IKE security association is established. • **Not matured**—The IKE security association is in the process of negotiation.	none specified
Initiator cookie	When the IKE negotiation is triggered, a random number is sent to the remote node.	All levels
Responder cookie	The remote node generates its own random number and sends it back to the initiator as a verification that the packets were received. Of the numerous security services available, protection against denial of service (DoS) is one of the most difficult to address. A “cookie” or anticlogging token (ACT) is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity. An exchange prior to CPU-intensive public key operations can thwart some DoS attempts (such as simple flooding with invalid IP source addresses).	All levels
Exchange type	Specifies the number of messages in an IKE exchange, and the payload types that are contained in each message. Each exchange type provides a particular set of security services, such as anonymity of the participants, perfect forward secrecy of the keying material, and authentication of the participants. Junos OS supports two types of exchanges: • **Main**—The exchange is done with six messages. **Main** encrypts the payload, protecting the identity of the neighbor. • **Aggressive**—The exchange is done with three messages. **Aggressive** does not encrypt the payload, leaving the identity of the neighbor unprotected. • **IKEv2**—The exchange is negotiated using IKE version 2.	All levels
PIC	The services PIC for which the IKE security associations are displayed.	All levels

Table 1421: show services ipsec-vpn ike security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Authentication method	Authentication method that determines which payloads are exchanged and when they are exchanged. Value can be ECDSA-signatures (256 bit key), ECDSA-signatures (384 bit key), Pre-shared-keys, or RSA-signatures. **NOTE**: In Junos FIPS mode, ECDSA is not supported in Junos OS Release 17.3R1. Starting in Junos OS Release 17.4R1, ECDSA is supported in Junos FIPS mode.	**detail**
Local	Prefix and port number of the local end.	**detail**
Remote	Prefix and port number of the remote end.	**detail**
Lifetime	Number of seconds remaining until the IKE security association expires.	**detail**
Algorithms	Header for the IKE algorithms output. • **Authentication**—(**detail** output only) Type of authentication algorithm used: **md5** or **sha1** • **Encryption**—(**detail** output only) Type of encryption algorithm used: **des-cbc**, **3des-cbc**, or **None**. • **Pseudo random function**—Function that generates highly unpredictable random numbers: **hmac-md5** or **hmac-sha1**.	**detail**
Traffic statistics	Number of bytes and packets received and transmitted on the IKE security association. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the IKE security association. • **Input packets, Output packets**—Number of packets received and transmitted on the IKE security association.	**detail**

Table 1421: show services ipsec-vpn ike security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Notification to the key management process of the status of the IKE negotiation: • **caller notification sent**—Caller program notified about the completion of the IKE negotiation. • **waiting for done**—Negotiation is done. The library is waiting for the remote end retransmission timers to expire. • **waiting for remove**—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation. • **waiting for policy manager**—Negotiation is waiting for a response from the policy manager.	**detail**
IPsec security associates	Number of IPsec security associations created and deleted with this IKE security association.	**detail**

Table 1421: show services ipsec-vpn ike security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Phase 2 negotiations in progress	Number of phase 2 negotiations in progress and status information:
	• Negotiation type—Type of phase 2 negotiation. The Junos OS currently supports **quick mode**.
	• Message ID—Unique identifier for a phase 2 negotiation.
	• Local identity—Identity of the local phase 2 negotiation. The format is *id-type-name* (*proto-name*:*port-number*,[0..*id-data-len*] = *iddata-presentation*).
	• Remote identity—Identity of the remote phase 2 negotiation. The format is *id-type-name* (*proto-name*:*port-number*,[0..*id-data-len*] = *iddata-presentation*)
	• Flags—Notification to the key management process of the status of the IKE negotiation:
	  • **caller notification sent**—Caller program notified about the completion of the IKE negotiation.
	  • **waiting for done**—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.
	  • **waiting for remove**—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.
	  • **waiting for policy manager**—Negotiation is waiting for a response from the policy manager.
		detail

Sample Output

show services ipsec-vpn ike security-associations

```
user@host> show services ipsec-vpn ike security-associations
Remote Address  State         Initiator cookie  Responder cookie  Exchange type

192.0.2.1         Matured       062d291d21275fc7  82ef00e3d1f1c981  Main
```

```
192.0.2.2         Matured       cd6d581d7bb1664d  88a707779f3ad8d1  Main

192.0.2.3         Matured       86621051e3e78360  6bc5cc83fd67baa4  IKEv2

PIC: sp-0/3/0

192.0.2.7         Matured       565e2813075e6fdb  67886757a74edcd6  IKEv2
```

show services ipsec-vpn ike security-associations detail

```
user@host> show services ipsec-vpn ike security-associations detail
IKE peer 198.51.100.2
      Role: Responder, State: Matured
      Initiator cookie: d91c9f20f78e1d4e, Responder cookie: 727a04ed8d5021a1
      Exchange type: IKEv2, Authentication method: Pre-shared-keys
      Local: 2013.0.113.2:500, Remote: 198.51.100:500
      Lifetime: Expires in 1357 seconds
      Algorithms:
       Authentication        : sha1
       Encryption            : 3des-cbc
       Pseudo random function: hmac-sha1
      Traffic statistics:
       Input  bytes  :                22244
       Output bytes  :                22236
       Input  packets:                  263
       Output packets:                  263
      Flags: Caller notification sent
      IPSec security associations: 0 created, 0 deleted
      Phase 2 negotiations in progress: 0

IKE peer 192.0.2.4
  Role: Initiator, State: Matured
  Initiator cookie: cf22bd81a7000001, Responder cookie: fe83795c2800002e
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 192.0.2.5:500, Remote: 192.0.2.4:500
  Lifetime: Expires in 187 seconds
  Algorithms:
   Authentication        : md5
   Encryption            : 3des-cbc
   Pseudo random function: hmac-md5
  Traffic statistics:
```

```
    Input  bytes  :                 1000
    Output bytes  :                 1280
    Input  packets:                    5
    Output packets:                    9
   Flags: Caller notification sent
   IPsec security associations: 2 created, 0 deleted
   Phase 2 negotiations in progress: 1

 Negotiation type: Quick mode, Role: Initiator, Message ID: 3582889153
     Local: 192.0.2.5:500, Remote: 192.0.2.4:500
     Local identity: ipv4_subnet(tcp:80,[0..7]=10.1.1.0/24)
     Remote identity: ipv4_subnet(tcp:100,[0..7]=10.1.2.0/24)
     Flags: Caller notification sent, Waiting for done
```

show services ipsec-vpn ike security-associations (on ACX500 Routers)

```
user@host> show services ipsec-vpn ike security-associations

Remote Address  State         Initiator cookie  Responder cookie  Exchange type
192.168.10.130  Matured       90864887dfecb178  9a2ee2ab786f960d  Main
192.168.20.130  Matured       1dd17732a8c9b13a  b06e5072ac7362bf  Main
192.0.2.7       Matured       565e2813075e6fdb  67886757a74edcd6  IKEv2
```

Release Information

Command introduced before Junos OS Release 7.4.

Statistics for Internet Key Exchange (IKE) security associations for each services PIC introduced in Junos OS Release 12.1.

show services ipsec-vpn ipsec security-associations

IN THIS SECTION

- Syntax | 25623
- Description | 25623

- Options | 25623
- Required Privilege Level | 25623
- Output Fields | 25623
- Sample Output | 25629
- Release Information | 25631

Syntax

```
show services ipsec-vpn ipsec security-associations
<brief | detail | extensive>
<service-set service-set-name>
```

Description

(Adaptive services interface only) Display IPsec security associations for the specified service set. If no service set is specified, the security associations for all service sets are displayed.

Options

none	Display standard information about IPsec security associations for all service sets.
brief | detail | extensive	(Optional) Display the specified level of output.
service-set *service-set-name*	(Optional) Display information about a particular service set.

Required Privilege Level

view

Output Fields

Table 1422 on page 25624 lists the output fields for the `show services ipsec-vpn ipsec security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 1422: show services ipsec-vpn ipsec security-associations Output Fields

Field Name	Field Description	Level of Output
`Service set`	Name of the service set for which the IPsec security associations are defined. If appropriate, includes the outside service interface VRF name.	All levels
`Rule`	Name of the rule set applied to the security association.	`detail extensive`
`Term`	Name of the IPsec term applied to the security association.	`detail extensive`
`Tunnel index`	Numeric identifier of the specific IPsec tunnel for the security association.	`detail extensive`
`Local gateway`	Gateway address of the local system.	All levels
`Remote gateway`	Gateway address of the remote system.	All levels
`IPsec inside interface`	Name of the logical interface hosting the IPsec tunnels.	All levels
`Tunnel MTU`	MTU of the IPsec tunnel.	All levels
`Total uptime`	Total amount of time that an IPsec tunnel has been up across security association rekeys.	`detail`

Table 1422: show services ipsec-vpn ipsec security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Local identity`	Protocol, address or prefix, and port number of the local entity of the IPsec association. The format is `id-type-name (proto-name:port-number, [0..`*`id-data-len`*`] = `*`iddata-presentation`*`)`. The protocol is always displayed as any because it is not user-configurable in the IPsec rule. Similarly, the port number field in the output is always displayed as 0 because it is not user-configurable in the IPsec rule. The value of the *`id-data-len`* parameter can be one of the following, depending on the address configured in the IPsec rule: • For an IPv4 address, the length is 4 and the value displayed is 3. • For a subnet mask of an IPv4 address, the length is 8 and the value displayed is 7. • For a range of IPv4 addresses, the length is 8 and the value displayed is 7. • For an IPv6 address prefix, the length is 16 and the value displayed is 15. • For a subnet mask of an IPv6 address prefix, the length is 32 and the value displayed is 31. • For a range of IPv6 address prefixes, the length is 32 and the value displayed is 31. The value of the *`id-data-presentation`* field denotes the IPv4 address or IPv6 prefix details. If the fully qualified domain name (FQDN) is specified instead of the address for the local peer of the IPsec association, it is displayed instead of the address details.	All levels

Table 1422: show services ipsec-vpn ipsec security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Remote identity`	Protocol, address or prefix, and port number of the remote entity of the IPsec association. The format is `id-type-name (proto-name:port-number, [0..`*`id-data-len`*`] = `*`iddata-presentation`*`)`. The protocol is always displayed as any because it is not user-configurable in the IPsec rule. Similarly, the port number field in the output is always displayed as 0 because it is not user-configurable in the IPsec rule. The value of the *`id-data-len`* parameter can be one of the following, depending on the address configured in the IPsec rule: • For an IPv4 address, the length is 4 and the value displayed is 3. • For a subnet mask of an IPv4 address, the length is 8 and the value displayed is 7. • For a range of IPv4 addresses, the length is 8 and the value displayed is 7. • For an IPv6 address prefix, the length is 16 and the value displayed is 15. • For a subnet mask of an IPv6 address prefix, the length is 32 and the value displayed is 31. • For a range of IPv6 address prefixes, the length is 32 and the value displayed is 31. The value of the *`id-data-presentation`* field denotes the IPv4 address or IPv6 prefix details. If the fully qualified domain name (FQDN) is specified instead of the address for the remote peer of the IPsec association, it is displayed instead of the address details.	All levels
`Primary remote gateway`	IP address of the configured primary remote peer.	All levels
`Backup remote gateway`	IP address of the configured backup remote peer.	All levels

Table 1422: show services ipsec-vpn ipsec security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the primary or backup interface: Active, Offline, or Standby. Both ES PICs are initialized to Offline. For primary and backup peers, State can be Active or Standby. If both peers are in a state of Standby, no connection exists yet between the two peers.	All levels
Failover counter	Number of times a PIC switched between primary and backup interfaces, or the number of times the tunnel switched between the primary and remote peers since the software has been activated.	All levels
Direction	Direction of the security association: inbound or outbound.	All levels
SPI	Value of the security parameter index.	All levels
AUX-SPI	Value of the auxiliary security parameter index. • When the value of Protocol is AH or ESP, AUX-SPI is always 0. • When the value of Protocol is AH+ESP, AUX-SPI is always a positive integer.	All levels
Mode	Mode of the security association: • transport—Protects single host-to-host protections. • tunnel—Protects connections between security gateways.	detail extensive
Type	Type of security association: • manual—Security parameters require no negotiation. They are static, and are configured by the user. • dynamic—Security parameters are negotiated by the IKE protocol. Dynamic security associations are not supported in transport mode.	detail extensive

Table 1422: show services ipsec-vpn ipsec security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Status of the security association: • Installed—The security association is installed in the security association database. (For transport mode security associations, the value of State must always be Installed.) • Not installed—The security association is not installed in the security association database.	detail extensive
Protocol	Protocol supported: • transport mode supports Encapsulation Security Protocol (ESP) or Authentication Header (AH). • tunnel mode supports ESP or AH+ESP.	All levels
Authentication	Type of authentication used: hmac-md5-96, hmac-sha1-96, or none.	detail extensive
Encryption	Type of encryption algorithm used: can be 3des-cbc, aes-cbc (128 bits), aes-cbc (192 bits), aes-cbc (256 bits), aes-gcm (128 bits), aes-gcm(192 bits), aes-gcm (256 bits), des-cbc, or None. **NOTE**: In Junos FIPS mode, AES-GCM is not supported in Junos OS Release 17.3R1. Starting in Junos OS Release 17.4R1, AES-GCM is supported in Junos FIPS mode.	detail
Soft lifetime Hard lifetime	Each lifetime of a security association has two display options, hard and soft, one of which must be present for a dynamic security association. The hard lifetime specifies the lifetime of the SA. The soft lifetime, which is derived from the hard lifetime, informs the IPsec key management system that the SA is about to expire. This information allows the key management system to negotiate a new SA before the hard lifetime expires. • Expires in *seconds* seconds—Number of seconds left until the security association expires. • Expires in *kilobytes* kilobytes—Number of kilobytes left until the security association expires.	detail extensive

Table 1422: show services ipsec-vpn ipsec security-associations Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Anti-replay service`	State of the service that prevents packets from being replayed: `Enabled` or `Disabled`.	`detail extensive`
`SA lifetime`	Configured hard lifetime (total lifetime), in seconds, for the security association.	`detail`
`Replay window size`	Configured size, in packets, of the antireplay service window: `32` or `64`. The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets. If the replay window size is `0`, antireplay service is disabled.	`detail`
`disable-natt`	Configure to disable NAT-T functionality. By default the NAT-T is enabled.	All levels.
`nat-keepalive`	Specify the interval at which NAT keepalive packets can be sent so that NAT translation continues.	All levels.

Sample Output

show services ipsec-vpn ipsec security associations extensive

```
user@host> show services ipsec-vpn ipsec security-associations extensive
Service set: service-set-1
  Rule: _junos_, Term: term-1, Tunnel index: 1
  Local gateway: 192.0.2.2, Remote gateway: 198.51.100.4
  IPSec inside interface: sp-2/0/0.1  Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Primary remote gateway: 192.0.2.1, State: Standby
  Backup remote gateway: 198.51.100.4, State: Active
  Failover counter: 1

   Direction: inbound, SPI: 3743521590, AUX-SPI: 0
   Mode: tunnel, Type: dynamic, State: Installed
   Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
```

```
    Soft lifetime: Expires in 23043 seconds
    Hard lifetime: Expires in 23178 seconds
    Anti-replay service: Enabled, Replay window size: 64

    Direction: outbound, SPI: 2551045240, AUX-SPI: 0
    Mode: tunnel, Type: dynamic, State: Installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Soft lifetime: Expires in 23043 seconds
    Hard lifetime: Expires in 23178 seconds
    Anti-replay service: Enabled, Replay window size: 64

disable-natt: No, nat-keepalive: 10
```

show services ipsec-vpn ipsec security associations detail

```
user@host> show services ipsec-vpn ipsec security-associations detail
Service set: ipsec-sset-0, IKE Routing-instance: default

  Rule: ipsec-rule-0, Term: term0, Tunnel index: 1
  Local gateway: 192.0.2.1, Remote gateway: 192.0.2.2
  IPSec inside interface: ms-3/0/0.1, Tunnel MTU: 1500
  UDP encapsulate: Disabled, UDP Destination port: 0
  Local identity: ipv4_subnet(any:0,[0..7]=198.51.100.0/16)
  Remote identity: ipv4_subnet(any:0,[0..7]=203.0.113.0/16)
  NATT Detection: Not Detected, NATT keepalive interval: 0
  Total uptime:  0 days 0 hrs 1 mins 4 secs

  Direction: inbound, SPI: 4004530393, AUX-SPI: 0
  Mode: tunnel, Type: dynamic, State: Installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Soft lifetime: Expires in 27885 seconds
  Hard lifetime: Expires in 28736 seconds
  Anti-replay service: Enabled, Replay window size: 4096
  Copy ToS: Enabled
  Copy TTL: Disabled, TTL value: 64
  SA lifetime: 28800 seconds

  Direction: outbound, SPI: 1323638473, AUX-SPI: 0
  Mode: tunnel, Type: dynamic, State: Installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Soft lifetime: Expires in 27885 seconds
```

```
  Hard lifetime: Expires in 28736 seconds
  Anti-replay service: Enabled, Replay window size: 4096
  Copy ToS: Enabled
  Copy TTL: Disabled, TTL value: 64
  SA lifetime: 28800 seconds
```

show services ipsec-vpn ipsec security associations (on ACX500 Routers)

```
user@host> show services ipsec-vpn ipsec security-associations

Service set: SS_1, IKE Routing-instance: Customer-1

  Rule: rule_1, Term: 1, Tunnel index: 2
  Local gateway: 192.168.1.11, Remote gateway: 192.168.10.130
  IPSec inside interface: ms-0/2/0.8, Tunnel MTU: 1300
  UDP encapsulate: Disabled, UDP Destination port: 0
    Direction SPI         AUX-SPI     Mode       Type     Protocol
    inbound   2204677182  0           tunnel     dynamic  ESP
    outbound  3015420439  0           tunnel     dynamic  ESP

Service set: SS_2, IKE Routing-instance: Customer-1

  Rule: Customer-1_rule_1, Term: 1, Tunnel index: 1
  Local gateway: 192.168.1.12, Remote gateway: 192.168.20.130
  IPSec inside interface: ms-0/2/0.7, Tunnel MTU: 1300
  UDP encapsulate: Disabled, UDP Destination port: 0
    Direction SPI         AUX-SPI     Mode       Type     Protocol
    inbound   2093089828  0           tunnel     dynamic  ESP
    outbound  2160146627  0           tunnel     dynamic  ESP
```

Release Information

Command introduced before Junos OS Release 7.4.

show services ipsec-vpn ipsec statistics

IN THIS SECTION

- Syntax | 25632
- Description | 25632
- Options | 25632
- Required Privilege Level | 25633
- Output Fields | 25633
- Sample Output | 25636
- Release Information | 25638

Syntax

```
show services ipsec-vpn ipsec statistics
<brief | detail>
<remote-gw remote-peer-address>
<service-set service-set-name>
```

Description

(Adaptive services interface only) Display IPsec statistics for the specified service set. If no service set is specified, the statistics for all service sets are displayed.

Options

none	Display standard IPsec statistics for all service sets.
brief | detail	(Optional) Display the specified level of output.
remote-gw *remote-peer-address*	(Optional) Display IPsec statistics for an individual IPsec tunnel and an individual remote host.
service-set *service-set-name*	(Optional) Display information about a particular service set.

Required Privilege Level

view

Output Fields

Table 1423 on page 25633 lists the output fields for the `show services ipsec-vpn ipsec statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1423: show services ipsec-vpn ipsec statistics Output Fields

Field Name	Field Description	Level of Output
PIC	The physical interface on which the IPsec tunnel is configured.	All levels
Service set	Name of the service set for which the IPsec tunnel is defined.	All levels
Local gateway	Gateway address of the local system.	All levels
Remote gateway	Gateway address of the remote system.	All levels
Tunnel index	Numeric identifier of the specific IPsec tunnel for the security association.	All levels
ESP statistics	Encapsulation Security Payload (ESP) statistics: • **Encrypted bytes**—Total number of bytes encrypted by the local system across the IPsec tunnel. • **Decrypted bytes**—Total number of bytes decrypted by the local system across the IPsec tunnel. • **Encrypted packets**—Total number of packets encrypted by the local system across the IPsec tunnel. • **Decrypted packets**—Total number of packets decrypted by the local system across the IPsec tunnel	All levels

Table 1423: show services ipsec-vpn ipsec statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
AH Statistics	Authentication Header statistics: • **Input bytes**—Total number of bytes received by the local system across the IPsec tunnel. • **Output bytes**—Total number of bytes transmitted by the local system across the IPsec tunnel. • **Input packets**—Total number of packets received by the local system across the IPsec tunnel. • **Output packets**—Total number of packets transmitted by the local system across the IPsec tunnel.	All levels

Table 1423: show services ipsec-vpn ipsec statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Errors	• **AH authentication failures**—Number of authentication header (AH) failures. An AH failure occurs when there is a mismatch of the authentication header in a packet transmitted across an IPsec tunnel. • **ESP authentication failures**—Number of Encapsulation Security Payload (ESP) failures. An ESP failure occurs when there is an authentication mismatch in ESP packets. • **ESP Decryption failures**—Number of ESP decryption failures. • **Bad headers**—Number of invalid headers detected. • **Bad trailers**—Number of invalid trailers detected. • **Replay before window drops**—Number of replay errors. A replay error is generated when a duplicate packet is received within the replay window. • **Replayed pkts**—Number of packets replayed. • **IP integrity errors**—Number of IP integrity errors. • **Exceeds tunnel MTU**—Number of times the tunnel maximum transmission unit (MTU) value was exceeded. • **Rule lookup failures**—Number of rule lookup failures. • **No SA errors**—Number of errors resulting from a missing security association (SA). • **Flow errors**—Number of flow errors. • **Misc errors**—Number of miscellaneous errors.	All levels

Sample Output

show services ipsec-vpn ipsec statistics detail

```
user@host> show services ipsec-vpn ipsec statistics
PIC: sp-0/2/0, Service set: ss0

ESP Statistics:
  Encrypted bytes:                0
  Decrypted bytes:                0
  Encrypted packets:              0
  Decrypted packets:              0
AH Statistics:
  Input bytes:                  168
  Output bytes:                 168
  Input packets:                  2
  Output packets:                 2
Errors:
  AH authentication failures:     0
  ESP authentication failures:    0
  ESP decryption failures:        0
  Bad headers: 0, Bad trailers: 0
  Replay before window drops: 0, Replayed pkts: 0
  IP integrity errors: 0, Exceeds tunnel MTU: 0
  Rule lookup failures: 0, No SA errors: 0
  Flow errors: 0, Misc errors: 0
```

show services ipsec-vpn ipsec statistics remote-gw

```
user@host> show services ipsec-vpn ipsec statistics remote-gw 192.0.2.1
PIC: sp-3/1/0, Service set: service-set-2
Local gateway: 198.51.100.1, Remote gateway: 192.0.2.1, Tunnel index: 2
ESP Statistics:
 Encrypted bytes:                0
 Decrypted bytes:                0
 Encrypted packets:              0
 Decrypted packets:              0
AH Statistics:
 Input bytes:                    0
 Output bytes:                   0
```

```
 Input packets:                  0
 Output packets:                 0
Errors:
  AH authentication failures:     0
  ESP authentication failures:    0
  ESP decryption failures:        0
  Bad headers: 0, Bad trailers: 0
  Replay before window drops: 0, Replayed pkts: 0
  IP integrity errors: 0, Exceeds tunnel MTU: 0
  Rule lookup failures: 0, No SA errors: 0
  Flow errors: 0, Misc errors: 0
```

show services ipsec-vpn ipsec statistics (on ACX500)

```
user@host> show services ipsec-vpn ipsec statistics

PIC: ms-0/2/0, Service set: SS_1

ESP Statistics:
  Encrypted bytes:          4121664
  Decrypted bytes:           151584
  Encrypted packets:          64162
  Decrypted packets:           1579
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures:     0
  ESP authentication failures:    0
  ESP decryption failures:        0
  Bad headers: 0, Bad trailers: 0
  Replay before window drops: 0, Replayed pkts: 0
  IP integrity errors: 0, Exceeds tunnel MTU: 0
  Rule lookup failures: 3, No SA errors: 0
  Flow errors: 0, Misc errors: 0


PIC: ms-0/2/0, Service set: SS_2
```

```
ESP Statistics:
  Encrypted bytes:              576
  Decrypted bytes:              576
  Encrypted packets:              6
  Decrypted packets:              6
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures:     0
  ESP authentication failures:    0
  ESP decryption failures:        0
  Bad headers: 0, Bad trailers: 0
  Replay before window drops: 0, Replayed pkts: 0
  IP integrity errors: 0, Exceeds tunnel MTU: 0
  Rule lookup failures: 0, No SA errors: 0
  Flow errors: 0, Misc errors: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

New fields added in Junos OS Release 10.0.

show services l2tp client

IN THIS SECTION

- Syntax | 25639
- Description | 25639
- Options | 25639
- Required Privilege Level | 25639
- Output Fields | 25639
- Sample Output | 25640

Release Information | 25640

Syntax

```
show services l2tp client
<client-name>
```

Description

Display information about all L2TP clients or a specific L2TP client.

Options

client-name	(Optional) Name of a client.

Required Privilege Level

view

Output Fields

Table 1424 on page 25639 lists the output fields for the `show services l2tp client` command. Output fields are listed in the approximate order in which they appear.

Table 1424: show services l2tp client Output Fields

Field Name	Field Description
`Client` `Client Name`	Name of the client.
`Tunnels`	Number of tunnels in the tunnel group.

Table 1424: show services l2tp client Output Fields *(Continued)*

Field Name	Field Description
Sessions	Number of L2TP sessions established for tunnels in the tunnel group.
Tunnel-group	Name of a tunnel group to which the client belongs.
Session-limit-group	Name of a session-limit group to which the client belongs.

Sample Output

show services l2tp client

```
user@host> show services l2tp client
Client           Tunnels    Sessions      Tunnel-group         Session-limit-group
entA-serviceA    2          20            l2tp-tunnel-group1   enterpriseA
entA-serviceB    3          120           l2tp-tunnel-group2   enterpriseB
```

show services l2tp client (Client Name)

```
user@host> show services l2tp client entA-serviceA
Client Name      Tunnels    Sessions    Tunnel-group          Session-limit-group
entA-serviceA    2          20          l2tp-tunnel-group1    enterpriseA
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

show services l2tp session-limit-group | 25674

show services l2tp tunnel-group | 25694

L2TP LAC Tunneling for Subscribers

show services l2tp destination

IN THIS SECTION

- Syntax | 25641
- Description | 25641
- Options | 25641
- Required Privilege Level | 25642
- Output Fields | 25642
- Sample Output | 25644
- Release Information | 25646

Syntax

```
show services l2tp destination
<brief | detail | extensive>
<local-gateway gateway-address>
<peer-gateway gateway-address>
<statistics>
```

Description

Display information about L2TP tunnel destinations.

Options

brief | detail | extensive	(Optional) Display the specified level of information.
local-gateway *gateway-address*	(Optional) Display L2TP session information for only the specified local gateway address.
peer-gateway *gateway-address*	(Optional) Display L2TP session information for only the specified peer gateway address.

statistics (Optional) Display the number of control packets and bytes transmitted and received for the destination. You cannot include this option with any of the level options, `brief`, `detail`, or `extensive`.

Required Privilege Level

view

Output Fields

Table 1425 on page 25642 lists the output fields for the `show services l2tp destination` command. Output fields are listed in the approximate order in which they appear.

Table 1425: show services l2tp destination Output Fields

Field Name	Field Description	Level of Output
`Local Name`	Name of this destination.	All levels
`Remote IP`	IP address of the remote peer (LNS).	All levels
`Tunnels`	Number of tunnel connections for the destination in the following categories: • total • active • failed	All levels for total `extensive` for active and failed
`Sessions`	Number of session connections for the destination in the following categories: • total • active • failed	All levels for total `extensive` for active and failed

Table 1425: show services l2tp destination Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Administrative state of the L2TP destination: • Enabled—No restrictions exist on creation or operation of sessions and tunnels for this destination. • Disabled—Existing sessions and tunnels for this destination have been disabled and no new sessions or tunnels are created while in the Disabled state. • Drain—Creation of new sessions and tunnels is disabled for this destination.	All levels
Local IP	IP address of the local gateway (LAC).	detailextensive
Transport	Medium used for tunneling. Only ipUdp is supported.	detailextensive
Logical System	Logical system in which the tunnel is configured.	detailextensive
Router Instance	Routing instance in which the tunnel is configured.	detailextensive
Lockout State	Reachability state of the destination: • not locked—Destination is considered reachable. • waiting for lockout timeout—Destination is locked out by L2TP because it is unreachable, so no attempts are made to reach the destination until the lockout timeout (300 seconds) expires, unless this is the only destination available for tunneling the subscriber.	detailextensive
Access Line Information	State of the LAC per-destination configuration for forwarding subscriber line information to the LNS, Enabled or Disabled. Starting in Junos OS Release 17.4R1, this information is displayed on the LNS for information it receives from the LAC, Enabled or Disabled.	detailextensive

Table 1425: show services l2tp destination Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Speed Updates	State of the LAC per-destination configuration for including connection speed updates when it forwards subscriber line information to the LNS, Enabled or Disabled. Starting in Junos OS Release 17.4R1, this information is displayed on the LNS for updates it receives from the LAC, Enabled or Disabled.	detailextensive
Connections	Number of total, active, and failed tunnel and session connections for the destination.	extensive
Control Tx	Amount of control information transmitted, in packets and bytes.	statistics
Control Rx	Amount of control information received, in packets and bytes.	statistics
Data Tx	Amount of data transmitted, in packets and bytes.	statistics
Data Rx	Amount of data received, in packets and bytes.	statistics
Error Tx	Number of errors transmitted, in packets.	statistics
Error Rx	Number of errors received, in packets.	statistics

Sample Output

show services l2tp destination

```
user@host> show services l2tp destination
  Local Name    Remote IP        Tunnels       Sessions  State
  1             203.0.113.101    1             1         Enabled
```

show services l2tp destination detail

```
user@host> show services l2tp destination detail
 Local name: 1
    Remote IP: 203.0.113.101
    Tunnels: 1, Sessions: 1
    State: Enabled
    Local IP: 203.0.113.102
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
    Access Line Information: Enabled, Speed Updates: Enabled
 Local name: 1
    Remote IP: 203.0.113.108
    Tunnels: 1, Sessions: 1
    State: Enabled
    Local IP: 203.0.113.2
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: waiting for lockout timeout
    Access Line Information: Enabled, Speed Updates: Enabled
```

show services l2tp destination extensive (LAC)

```
user@host> show services l2tp destination extensive
  Local name: 1
    Remote IP: 203.0.113.101
    State: Enabled
    Local IP: 203.0.113.102
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
    Access Line Information: Enabled, Speed Updates: Enabled
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
```

show services l2tp destination extensive (LNS)

```
user@host> show services l2tp destination extensive
  Local name: 3
    Remote IP: 203.0.113.103
```

```
    State: Enabled
    Local IP: 203.0.113.102
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
    Access Line Information: Enabled, Speed Updates: Disabled
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
```

show services l2tp destination statistics (LAC only on MX Series Routers)

```
user@host> show services l2tp destination statistics
 Local name: 2, Tunnels: 1, Sessions: 210
                       Packets        Bytes
    Control Tx             680        63.3k
    Control Rx             283        10.6k
    Data Tx               1129        14.3k
    Data Rx                877        10.9k
    Errors Tx                0
    Errors Rx                0
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- clear services l2tp destination | 4117
- show services l2tp destination lockout | 25647
- show services l2tp session | 25661
- show services l2tp summary | 25676
- show services l2tp tunnel | 25685

show services l2tp destination lockout

IN THIS SECTION

- Syntax | 25647
- Description | 25647
- Options | 25647
- Required Privilege Level | 25647
- Output Fields | 25647
- Sample Output | 25648
- Release Information | 25648

Syntax

```
show services l2tp destination lockout
```

Description

Display a list of destinations that are currently locked out and the time remaining for each to remain in the lockout state.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1426 on page 25648 lists the output fields for the `show services l2tp destination lockout` command. Output fields are listed in the approximate order in which they appear.

Table 1426: show services l2tp destination lockout Output Fields

Field Name	Field Description
Destination	Name of the destination.
Time Remaining	Time remaining for the destination to be locked out.
L2TP lockout destinations found	Total count of lockout destinations.

Sample Output

show services l2tp destination lockout

```
user@host> show services l2tp destination lockout
  Destination   Time Remaining
  4             45
  5             43
  6             8
3 L2TP lockout destinations found
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

clear services l2tp destination | 4117

request services l2tp destination unlock | 14951

show services l2tp destination | 25641

show services l2tp session | 25661

show services l2tp summary | 25676

show services l2tp tunnel | 25685

show services l2tp multilink

IN THIS SECTION

- Syntax | 25649
- Description | 25649
- Options | 25649
- Required Privilege Level | 25650
- Output Fields | 25650
- Sample Output | 25654
- Release Information | 25655

Syntax

```
show services l2tp multilink
<brief | detail | extensive | statistics>
<bundle-id number>
```

Description

(M10i and M7i routers only) Display L2TP output organized by multilink bundle.

Options

none	Same as brief.
brief | detail | extensive | statistics	(Optional) Display the specified level of output. Use the **statistics** option to display packets and bytes that have been encapsulated in the Multilink Protocol. Nonmultilink packets received on member sessions are not counted here.
bundle-id *number*	(Optional) Display L2TP multilink bundle information for only the specified bundle.

Required Privilege Level

view

Output Fields

Table 1427 on page 25650 lists the output fields for the `show services l2tp multilink` command. Output fields are listed in the approximate order in which they appear.

Table 1427: show services l2tp multilink Output Fields

Field Name	Field Description	Level of Output
Bundle ID	Bundle identifier.	All levels
Links	Number of links in the multilink bundle.	All levels
Bundle endpoint	Endpoint discriminator that represents the device transmitting the packet.	All levels
Input MRRU	Maximum packet size that the input interface can process.	detail
Output MRRU	Maximum packet size that the output interface can process.	detail
Session local ID	Identifier of the local endpoint of the L2TP session, as assigned by the L2TP network server (LNS).	detail
Session remote ID	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	detail

Table 1427: show services l2tp multilink Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Status of the L2TP session: • **Established**—The session is operating. • **closed**—The session is being closed. • **destroyed**—The session is being destroyed. • **clean-up**—The session is being cleaned up. • **lns-ic-accept-new**—A new session is being accepted. • **lns-ic-idle**—The session has been created and is idle. • **lns-ic-reject-new**—The new session is being rejected. • **lns-ic-wait-connect**—The session is waiting for the peer's incoming call connected (ICCN) message.	detail
Username	Name of the user logged in to the session.	detail
Mode	Mode of the interface representing the multilink bundle: **dedicated** or **shared**.	extensive
Local IP	IP address of the local endpoint of the Point-to-Point Protocol (PPP) session.	extensive
Remote IP	IP address of the remote endpoint of the PPP session.	extensive

Table 1427: show services l2tp multilink Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Local name	Name of the LNS instance in which the session was created.	extensive
Remote name	Name of the LAC from which the session was created.	extensive
Local MRU	Maximum receive unit (MRU) setting of the local device, in bytes.	extensive
Remote MRU	MRU setting of the remote device, in bytes.	extensive

Table 1427: show services l2tp multilink Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Statistics since	Date and time when collection of the following statistics began: • **Control Tx**—Amount of control information transmitted, in packets and bytes. • **Control Rx**—Amount of control information received, in packets and bytes. • **Data Tx**—Amount of data transmitted, in packets and bytes. • **Data Rx**—Amount of data received, in packets and bytes. • **Errors Tx**—Number of errors transmitted, in packets. • **Errors Rx**—Number of errors received, in packets. • **Lcp Echo Req Tx**—Number of LCP echo requests transmitted, in packets. • **Lcp Echo Req Rx**—Number of LCP echo requests received, in packets. • **Lcp Echo Rep Tx**—Number of LCP echo responses transmitted, in packets. • **Lcp Echo Rep Rx**—Number of LCP echo responses received, in packets.	extensive

Table 1427: show services l2tp multilink Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **Lcp Echo Req Timeout**—Number of LCP echo requests that timed out. • **Lcp Echo Req Error**—Number of errors received for LCP echo packets. • **Lcp Echo Rep Error** —Number of errors transmitted for LCP echo packets. • **MRRU**—Maximum packet size processed. • **TX**—Number of packets transmitted. • **RX**—Number of packets received. • **link**—Link of the multilink bundle associated with the L2TP session.	

Sample Output

show services l2tp multilink extensive

```
user@host> show services l2tp multilink extensive
Bundle ID: 1
  Links: 2, Bundle endpoint: user@example.com
  Input MRRU: 1524, Output MRRU: 1524
  Session local ID: 46122, Session remote ID: 39307
    State: Established, Username: user1@example.com, Mode: dedicated
    Local IP: 10.58.255.129:1701, Remote IP: 10.58.255.131:1701
    Local name: router3, Remote name: router4
  Session local ID:  4254, Session remote ID: 39308
    State: Established, Username: user2@example.com, Mode: dedicated
```

```
    Local IP: 10.1.255.1:1701, Remote IP: 10.1.255.2:1701
    Local name: router1, Remote name: router2
  Statistics since: Mon May 17 11:47:35 2004
                        Packets        Bytes
      Control Tx              7          196
      Control Rx              3           90
      Data Tx                 0            0
      Data Rx                 0            0
      Errors Tx               0
      Errors Rx               0
         Lcp Echo Req Tx         0
         Lcp Echo Req Rx         0
         Lcp Echo Rep Tx         0
         Lcp Echo Rep Rx         0
         Lcp Echo Req Timeout         0
         Lcp Echo Req Error           0
         Lcp Echo Rep Error           0
         MRRU 1486 droptime 0 maxfrag 0   minfrag 32  minmru 1482 maxqlen 3000
        TX: Packets 0    Frags 0    Txseq 0x0
        RX: Packets 24   Frags 24   Rxseq 0x18   mseq 23  maxdiff 1   reass 24
        fragments copied 0
        link 0 : seq 0x17 mru 1482 encaplen 8  qlen 0   context 0xea01eb0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

L2TP Services Configuration Overview

L2TP Minimum Configuration

clear services l2tp multilink | **4123**

show services l2tp radius

IN THIS SECTION

- Syntax | 25656
- Description | 25656
- Options | 25656
- Required Privilege Level | 25657
- Output Fields | 25657
- Sample Output | 25659
- Release Information | 25661

Syntax

```
show services l2tp radius
<accounting (servers | statistics)>
<authentication (servers | statistics)>
<servers>
<statistics>
```

Description

(M7i, M10i, and M120 routers only) Display RADIUS servers and statistics information for the RADIUS servers configured on the router.

Options

You must include one of the following keywords to provide a valid completion for the command:

accounting (servers | statistics)	(Optional) Display RADIUS servers or statistical accounting information only.
authentication (servers | statistics)	(Optional) Display RADIUS servers or statistical authentication information only.

servers (Optional) Display RADIUS authentication and accounting server information only.

statistics (Optional) Display RADIUS authentication and accounting statistics information only.

Required Privilege Level

view

Output Fields

Table 1428 on page 25657 lists the output fields for the `show services l2tp radius` command. Output fields are listed in the approximate order in which they appear.

Table 1428: show services l2tp radius Output Fields

Field Name	Field Description
`IP Address`	IP address of the server.
`State`	(`servers` keyword only) Present state of the server.
`UDP Port`	Number of the UDP port used to send authentication or accounting messages to the server.
`Retry Count`	(`servers` keyword only) Number of times the RADIUS client resends a packet if no ACK is received.
`Timeout`	(`servers` keyword only) Length of time the client waits for an ACK before retransmission.
`Pending Requests`	(`servers` keyword only) Number of client pending authentication or accounting requests.
`Maximum Sessions`	(`servers` keyword only) Maximum number of pending requests on each RADIUS client before the server moves to the next RADIUS client, which is 200 times the maximum number of clients that can be created on a server (which is 12).

Table 1428: show services l2tp radius Output Fields *(Continued)*

Field Name	Field Description
Dead Time	(servers keyword only) Interval to wait before retrying a server after it fails to send a response to an authentication or accounting request.
Secret Type	(servers keyword only) Secret type configured on the RADIUS server.
Profile	(servers keyword only) Name of profile configured for the RADIUS server.
Access requests	(statistics keyword only) Number of access requests sent to the server.
Rollover requests	(statistics keyword only) Number of requests coming into the server as a result of the previous server timing out.
Retransmissions	(statistics keyword only) Number of retransmissions.
Access accepts	(statistics keyword only) Number of access accept messages received from the server.
Access rejects	(statistics keyword only) Number of access reject messages received from the server.
Access challenges	(statistics keyword only) Number of access challenges received from the server.
Malformed responses	(statistics keyword only) Number of responses with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).
Bad authenticators	(statistics keyword only) Number of responses in which the authenticator is incorrect for the matching request. This can occur if the RADIUS secrets for the client and server do not match.
Requests pending	(statistics keyword only) Number of requests waiting for a response.
Request timeouts	(statistics keyword only) Number of requests that timed out.

Table 1428: show services l2tp radius Output Fields *(Continued)*

Field Name	Field Description
`Unknown responses`	(`statistics` keyword only) Number of unknown responses. The RADIUS response type in the header is invalid or unsupported.
`Packets dropped`	(`statistics` keyword only) Number of packets dropped because they are too short or because the router receives a response for which there is no corresponding request. For example, if the router sends a request that times out, the router removes the request from the list and sends a new request. If the server is slow and sends a response to the first request after the router removes the request, the packet is dropped.

Sample Output

show services l2tp radius servers

```
user@host> show services l2tp radius servers
                        RADIUS Authentication Servers

                            UDP  Retry         Pending  Maximum  Dead   Secret
     IP Address     State   Port Count Timeout Requests Sessions Time   Type
   192.0.2.1        Active   1812 2     25      0        2400     300  radius-key
   198.51.100.1     Active   1812 5     35      0        2400     300  radius-key
   203.0.113.1     Active   1812 2     25      0        2400     300  radius-key
   172.28.30.174   Active   1812 7     75      0        2400     300  radius-key
   172.28.30.175   Active   1812 7     75      0        2400     300  radius-key
   172.28.30.176   Active   1812 4     55      0        2400     300  radius-key
   172.31.30.176  Active   1812 3     3       0        2400     300  none-set
   172.31.130.174 Active   1812 7     75      0        2400     300  radius-key

                        RADIUS Accounting Servers

                            UDP  Retry         Pending  Maximum  Dead   Secret
     IP Address     State   Port Count Timeout Requests Sessions Time   Type
   192.0.2.1        Active   1813 2     25      0        2400     300  radius-key
   198.51.100.1     Active   1813 5     35      0        2400     300  radius-key
   203.0.113.1     Active   1813 2     25      0        2400     300  radius-key
   172.28.30.174   Active   1813 7     75      0        2400     300  radius-key
   172.28.30.175   Active   1813 7     75      0        2400     300  radius-key
```

```
   172.28.30.176   Active   1813 4     55      0        2400     300  radius-key
   172.31.30.176  Active   1813 3     3       0        2400     300  none-set
   172.31.130.174 Active   1813 7     75      0        2400     300  radius-key

                          RADIUS Accounting Servers

    Profile: user1
```

show services l2tp radius statistics

```
user@host> show services l2tp radius statistics
     RADIUS Authentication Statistics

Authentication statistics:
  Server 192.0.2.1, UDP port: 1812
    Access requests     : 40
    Rollover requests   : 5
    Retransmissions     : 2
    Access accepts      : 39
    Access rejects      : 1
    Access challenges   : 3
    Malformed responses : 0
    Bad authenticators  : 0
    Requests pending    : 1
    Request timeouts    : 0
    Unknown responses   : 0
    Packets dropped     : 0

     RADIUS Accounting Statistics

 Accounting statistics:
  Server 172.31.130.174, UDP port: 1813
    Total requests      : 9
    Start requests      : 6
    Interim requests    : 1
    Stop requests       : 2
    Rollover requests   : 0
    Retransmissions     : 1
    Total response      : 9
    Start responses     : 6
    Interim responses   : 1
```

```
    Stop responses      : 2
    Malformed responses : 0
    Bad authenticators  : 0
    Requests pending    : 1
    Request timeouts    : 0
    Unknown responses   : 0
    Packets dropped     : 0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- L2TP Services Configuration Overview
- L2TP Minimum Configuration

show services l2tp session

IN THIS SECTION

- Syntax | 25661
- Description | 25662
- Options | 25662
- Required Privilege Level | 25663
- Output Fields | 25663
- Sample Output | 25670
- Release Information | 25674

Syntax

```
show services l2tp session
<brief | detail | extensive>
```

```
<interface interface-name>
<local-gateway gateway-address>
<local-gateway-name gateway-name>
<local-session-id session-id>
<local-tunnel-id tunnel-id>
<peer-gateway gateway-address>
<peer-gateway-name gateway-name>
<statistics>
<tunnel-group group-name>
<user username>
```

Description

(M10i and M7i routers only) Display information about active L2TP sessions for LNS.

(MX Series routers only) Display information about active L2TP sessions for LAC and LNS.

Options

none	Display standard information about all active L2TP sessions.
brief | detail | extensive	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display L2TP session information for only the specified adaptive services or inline services interface. The interface type depends on the line card as follows: • `si-`*fpc/pic/port*— MPCs on MX Series routers only. This option is not available for L2TP on M Series routers. • `sp-`*fpc/pic/port*—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers.
local-gateway *gateway-address*	(Optional) Display L2TP session information for only the specified local gateway address.
local-gateway-name *gateway-name*	(Optional) Display L2TP session information for only the specified local gateway name.
local-session-id *session-id*	(Optional) Display L2TP session information for only the specified local session identifier.

local-tunnel-id *tunnel-id* — (Optional) Display L2TP session information for only the specified local tunnel identifier.

peer-gateway *gateway-address* — (Optional) Display L2TP session information for only the specified peer gateway address.

peer-gateway-name *gateway-name* — (Optional) Display L2TP session information for only the specified peer gateway name.

statistics — (Optional) Display the number of control packets and bytes transmitted and received for the session. You cannot include this option with any of the level options, `brief`, `detail`, or `extensive`.

tunnel-group *group-name* — (Optional) Display L2TP session information for only the specified tunnel group. To display information about L2TP CPU and memory usage, you can include the tunnel group name in the `show services service-sets memory-usage` *`group-name`* and `show services service-sets cpu-usage` *`group-name`* commands. This option is not available for L2TP LAC on MX Series routers.

user *username* — (M Series routers only) (Optional) Display L2TP session information for only the specified username.

Required Privilege Level

view

Output Fields

Table 1429 on page 25663 lists the output fields for the `show services l2tp session` command. Output fields are listed in the approximate order in which they appear.

Table 1429: show services l2tp session Output Fields

Field Name	Field Description	Level of Output
`Interface`	(LNS only) Name of an adaptive services interface.	All levels
`Tunnel group`	(LNS only) Name of a tunnel group.	All levels

Table 1429: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Tunnel local ID`	Identifier of the local endpoint of the tunnel, as assigned by the L2TP network server (LNS).	All levels
`Session local ID`	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	All levels
`Session remote ID`	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	All levels
`State`	State of the L2TP session: • `Established`—Session is operating. This is the only state supported for the LAC. • `closed`—Session is being closed. • `destroyed`—Session is being destroyed. • `clean-up`—Session is being cleaned up. • `lns-ic-accept-new`—New session is being accepted. • `lns-ic-idle`—Session has been created and is idle. • `lns-ic-reject-new`—New session is being rejected. • `lns-ic-wait-connect`—Session is waiting for the peer's incoming call connected (ICCN) message.	All levels
`Bundle ID`	(LNS only) Bundle identifier. Indicates the session is part of a multilink bundle. Sessions that have a blank `Bundle` field are not participating in the Multilink Protocol. Sessions in a multilink bundle might belong to different L2TP tunnels. For L2TP output organized by bundle ID, issue the `show services l2tp multilink extensive` command.	All levels

Table 1429: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Mode`	(LNS) Mode of the interface representing the session: `shared` or `exclusive`. (LAC) Mode of the interface representing the session: `shared` or `dedicated`. Only `dedicated` is currently supported for the LAC.	`extensive`
`Local IP`	IP address of local endpoint of the Point-to-Point Protocol (PPP) session.	`extensive`
`Remote IP`	IP address of remote endpoint of the PPP session.	`extensive`
`Username`	(LNS only) Name of the user logged in to the session.	All levels
`Assigned IP address`	(LNS only) IP address assigned to remote client.	`extensive`
`Local name`	For LNS, name of the LNS instance in which the session was created. For LAC, name of the LAC.	`extensive`
`Remote name`	For LNS, name of the LAC from which the session was created. For LAC, name of the LAC instance.	`extensive`
`Local MRU`	(LNS only) Maximum receive unit (MRU) setting of the local device, in bytes.	`extensive`
`Remote MRU`	(LNS only) MRU setting of the remote device, in bytes.	`extensive`

Table 1429: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Tx speed`	Transmit speed of the session conveyed from the LAC to the LNS, in bits per second (bps) and the source method from which the speed is derived. Starting in Junos OS Release 14.1, either the initial (`initial`) line speed or both the initial and current (`update`) line speeds can be displayed on MX Series routers: • When connection speed updates are not enabled, then only the initial line speed is displayed. • When connection speed updates are enabled, then both the initial and the current speeds are displayed. For Junos OS Release 17.2 and Release 17.3, only the current (`update`) line speed can be displayed on MX Series routers. Starting in Junos OS Release 17.4R1, once again either the initial (`initial`) line speed or both the initial and current (`update`) line speeds can be displayed on MX Series routers. Starting in Junos OS Release 15.1, when the Tx connect speed method is set to none, the value of zero (0) is displayed.	`extensive`

Table 1429: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Rx speed`	Receive speed of the session conveyed from the LAC to the LNS, in bits per second (bps) and the source method from which the speed is derived. Starting in Junos OS Release 14.1, either the initial (`initial`) line speed or both the initial and current (`update`) line speeds can be displayed on MX Series routers: • When connection speed updates are not enabled, then only the initial line speed is displayed. • When connection speed updates are enabled, then both the initial and the current speeds are displayed. For Junos OS Release 17.2 and Release 17.3, only the current (`update`) line speed can be displayed on MX Series routers. Starting in Junos OS Release 17.4R1, once again either the initial (`initial`) line speed or both the initial and current (`update`) line speeds can be displayed on MX Series routers. Starting in Junos OS Release 15.1, when the Tx connect speed method is set to none, the value of zero (0) is displayed.	`extensive`
`Bearer type`	Type of bearer enabled: • `0`—Might indicate that the call was not received over a physical link (for example, when the LAC and PPP are located in the same subsystem). • `1`—Digital access requested. • `2`—Analog access requested. • `4`—Asynchronous Transfer Mode (ATM) bearer support.	`extensive`
`Framing type`	Type of framing enabled: • `1`—Synchronous framing • `2`—Asynchronous framing	`extensive`

Table 1429: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`LCP renegotiation`	(LNS only) Whether Link Control Protocol (LCP) renegotiation is configured: `On` or `Off`.	`extensive`
`Authentication`	Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).	`extensive`
`Interface ID`	(LNS only) Identifier used to look up the logical interface for this session.	`extensive`
`Interface unit`	Logical interface for this session.	All levels
`Call serial number`	Unique serial number assigned to the call.	`extensive`
`Policer bandwidth`	Maximum policer bandwidth configured for this session.	`extensive`
`Policer burst size`	Maximum policer burst size configured for this session.	`extensive`
`Firewall filter`	Configured firewall filter name.	`extensive`
`Session encapsulation overhead`	Overhead allowance configured for this session, in bytes.	`extensive`
`Session cell overhead`	Cell overhead activation (`On` or `Off`).	`extensive`
`Create time`	Date and time when the call was created.	`extensive`

Table 1429: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Up time`	Length of time elapsed since the call became active, in hours, minutes, and seconds.	`extensive`
`Idle time`	Length of time elapsed since the call became idle, in hours, minutes, and seconds.	`extensive`
`Statistics since`	Date and time when collection of the following statistics began: • `Control Tx`—Amount of control information transmitted, in packets and bytes. • `Control Rx`—Amount of control information received, in packets and bytes. • `Data Tx`—Amount of data transmitted, in packets and bytes. • `Data Rx`—Amount of data received, in packets and bytes. • `Errors Tx`—Number of errors transmitted, in packets. • `Errors Rx`—Number of errors received, in packets. • `LCP echo req Tx`—Number of LCP echo requests transmitted, in packets. • `LCP echo req Rx`—Number of LCP echo requests received, in packets. • `LCP echo rep Tx`—Number of LCP echo responses transmitted, in packets. • `LCP echo rep Rx`—Number of LCP echo responses received, in packets. • `LCP echo Req timout`—Number of LCP echo requests that timed out. • `LCP echo Req error`—Number of errors received for LCP echo packets. • `LCP echo Rep error` —Number of errors transmitted for LCP echo packets.	`extensive`

Sample Output

show services l2tp session (LNS on M Series Routers)

```
user@host> show services l2tp session
Interface: sp-1/2/0, Tunnel group: group1, Tunnel local ID:  8802
  Local Remote Interface State               Bundle Username
  ID    ID     unit
  37966      5         2 Established
```

show services l2tp session (LNS on MX Series Routers)

```
user@host> show services l2tp session
Tunnel local ID: 40553
  Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
  17967  1        Established          1073749824        si-5/2/0
```

show services l2tp session (LAC)

```
user@host> show services l2tp session
Tunnel local ID: 31889
  Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
   31694      1   Established                311         pp0
```

show services l2tp session detail (LAC)

```
user@host> show services l2tp session detail
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:     1, Interface unit: 311
    State: Established, Interface: pp0, Mode: Dedicated
    Local IP: 203.0.113.2:1701, Remote IP: 203.0.113.1:1701
    Local name: ce-lac, Remote name: ce-lns
```

show services l2tp session extensive (LAC)

```
user@host> show services l2tp session extensive
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:     1
    Interface unit: 311
    State: Established, Mode: Dedicated
    Local IP: 203.0.113.2:1701, Remote IP: 203.0.113.1:1701
    Local name: ce-lac, Remote name: ce-lns
    Tx speed: 0, Rx speed: 0
    Bearer type: 1, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Interface unit: 311, Call serial number: 0
    Policer bandwidth: 0, Policer burst size: 0
    Policer exclude bandwidth: 0, Firewall filter: 0
    Session encapsulation overhead: 0, Session cell overhead: 0
    Create time: Tue Aug 24 14:38:23 2010, Up time: 01:06:25
    Idle time: N/A
```

show services l2tp session extensive (LAC on MX Series Routers)

```
user@host> show services l2tp session extensive
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:     1
    Interface unit: 311
    State: Established, Mode: Dedicated
    Local IP: 203.0.113.102:1701, Remote IP: 203.0.113.101:1701
    Local name: ce-lac, Remote name: ce-lns
    Tx speed: 256000, source service-profile
    Rx speed: 128000, source ancp
    Bearer type: 1, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Interface unit: 311, Call serial number: 0
    Policer bandwidth: 0, Policer burst size: 0
    Policer exclude bandwidth: 0, Firewall filter: 0
    Session encapsulation overhead: 0, Session cell overhead: 0
    Create time: Tue Aug 24 14:38:23 2010, Up time: 01:06:25
    Idle time: N/A
```

show services l2tp session extensive (LNS on M Series Routers)

```
user@host> show services l2tp session extensive
Interface: sp-1/2/0, Tunnel group: group1, Tunnel local ID: 62746
  Session local ID: 56793, Session remote ID: 53304
    State: Established, Bundle ID: 5, Mode: shared
    Local IP: 203.0.113.121:1701, Remote IP: 203.0.113.202:1701
    Username: user@example.com, Assigned IP address: 203.0.113.51/32
    Local MRU: 4000, Remote MRU: 1500, Tx speed: 64000, Rx speed: 64000
    Bearer type: 2, Framing type: 1
    LCP renegotiation: Off, Authentication: CHAP, Interface ID: unit_20
    Interface unit: 20, Call serial number: 4137941434
    Policer bandwidth: 64000, Policer burst size: 51200
    Firewall filter: f1
    Session encapsulation overhead: 16, Session cell overhead: On
    Create time: Tue Mar 23 14:13:15 2004, Up time: 01:16:41
    Idle time: 00:00:00
    Statistics since: Tue Mar 23 14:13:13 2004
                        Packets        Bytes
      Control Tx              4           88
      Control Rx              2           28
      Data Tx                 0            0
      Data Rx               461        29.0k
      Errors Tx               0
      Errors Rx               0

Interface: sp-1/2/0, Tunnel group: group_company_dns, Tunnel local ID: 37266
  Session local ID: 39962, Session remote ID: 53303
    State: Established, Bundle ID: 5, Mode: shared
    Local IP: 203.0.113.121:1701, Remote IP: 203.0.113.222:1701
    Username: usr1@company.example.com, Assigned IP address: 203.0.113.3/24
    Local name: router-1, Remote name: router-2
    Local MRU: 4470, Remote MRU: 4470, Tx speed: 155000000, Rx speed: 155000000
    Bearer type: 2, Framing type: 1
    LCP renegotiation: Off, Authentication: CHAP, Interface ID: unit_31
    Interface unit: 31, Call serial number: 4137941433
    Policer bandwidth: 64000, Policer burst size: 51200
    Firewall filter: f1
    Create time: Tue Mar 23 14:13:17 2004, Up time: 01:16:39
    Idle time: 01:16:36
    Statistics since: Tue Mar 23 14:13:15 2004
                        Packets        Bytes
```

```
      Control Tx              6          196
      Control Rx              4          150
      Data Tx                 0            0
      Data Rx                 1           80
      Errors Tx               0
      Errors Rx               0
```

show services l2tp session extensive (LNS on MX Series Routers)

```
user@host> show services l2tp session extensive
Tunnel local ID: 40553
  Session local ID: 17967, Session remote ID: 1
    Interface unit: 1073749824
    State: Established
    Interface: si-5/2/0
    Mode: Dedicated
    Local IP: 192.0.2.2:1701, Remote IP: 192.0.2.3:1701
    Local name: lns-mx960, Remote name: testlac
    Tx speed: initial 64000, Update 256000
    Rx speed: initial 64000, Update 256000
    Bearer type: 2, Framing type: 1
    LCP renegotiation: Off, Authentication: None
    Call serial number: 1
    Create time: Mon Apr 25 20:27:50 2011, Up time: 00:01:48
    Idle time: N/A
    Statistics since: Mon Apr 25 20:27:50 2011
                        Packets        Bytes
      Control Tx              4          219
      Control Rx              4          221
      Data Tx                 0            0
      Data Rx                10          228
      Errors Tx               0
      Errors Rx               0
```

show services l2tp session statistics (MX Series Routers)

```
user@host>show services l2tp session statistics local session-id 1
Tunnel local ID: 17185
Session local ID: 1, Session remote ID: 14444, Interface unit: 1073788352
State: Established
```

```
Statistics since: Mon Aug 1 13:27:47 2011
                Packets         Bytes
Data Tx         4             51
Data Rx         3             36
```

Release Information

Command introduced before Junos OS Release 7.4.

Support for LAC on MX Series routers introduced in Junos OS Release 10.4.

Support for LNS on MX Series routers introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- L2TP Services Configuration Overview
- L2TP Minimum Configuration
- clear services l2tp session | **4125**

show services l2tp session-limit-group

IN THIS SECTION

- Syntax | **25675**
- Description | **25675**
- Options | **25675**
- Required Privilege Level | **25675**
- Output Fields | **25675**
- Sample Output | **25676**
- Release Information | **25676**

Syntax

```
show services l2tp session-limit-group
<limit-group-name>
```

Description

Display information about all session-limit groups or a specific session limit group.

Options

limit-group-name	(Optional) Name of a session-limit group.

Required Privilege Level

view

Output Fields

Table 1430 on page 25675 lists the output fields for the `show services l2tp session-limit-group` command. Output fields are listed in the approximate order in which they appear.

Table 1430: show services l2tp session-limit-group Output Fields

Field Name	Field Description
`Session-limit-group`	Name of a session-limit group.
`Tunnels`	Number of tunnels associated with the session-limit group in the tunnel group.
`Sessions`	Number of L2TP sessions established for session-limit group.
`Maximum limit`	Maximum number of sessions allowed for the session-limit group.

Sample Output

show services l2tp session-limit-group

```
user@host> show services l2tp session-limit-group

Session-limit-group     Tunnels       Sessions       Maximum limit
enterpriseA             2             10             1000
enterpriseB             10            120            2000
```

show services l2tp session-limit-group (Limit Group Name)

```
user@host> show services l2tp session-limit-group enterpriseA

Session-limit-group     Tunnels       Sessions       Maximum limit
enterpriseA             2             10             1000
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

show services l2tp client | 25638

show services l2tp tunnel-group | 25694

L2TP LAC Tunneling for Subscribers

show services l2tp summary

IN THIS SECTION

- Syntax | 25677
- Description | 25677

- Options | 25677
- Required Privilege Level | 25677
- Output Fields | 25678
- Sample Output | 25682
- Release Information | 25684

Syntax

```
show services l2tp summary
<interface sp-fpc/pic/port>
<statistics>
```

Description

(M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Display Layer 2 Tunneling Protocol (L2TP) summary information.

Options

none	Display complete L2TP summary information. For LNS on M Series routers, display L2TP summary information for all adaptive services interfaces. For LNS on MX Series routers, display L2TP summary information for all inline services interfaces.
interface sp-*fpc/pic/port*	(Optional) Display L2TP summary information for only the specified adaptive services interface. This option is not available for L2TP on MX Series routers.
statistics	(Optional) Display a summary of control packets and bytes transmitted and received.

Required Privilege Level

view

Output Fields

Table 1431 on page 25678 lists the output fields for the `show services l2tp summary` command. Output fields are listed in the approximate order in which they appear.

Table 1431: show services l2tp summary Output Fields

Field Name	Field Description
`Administrative state`	Administrative state of the tunnel is drain. In this state you cannot configure new sessions, destinations, or tunnels at the LAC or LNS.
`Failover within a preference level`	State of this tunnel selection method on the LAC. When enabled, tunnel selection fails over within a preference level. When disabled, tunnel selection drops to the next lower preference level. Not displayed for LNS on M Series routers.
`Weighted load balancing`	State of this tunnel selection method on the LAC. When enabled, the maximum session limit of a tunnel determines its weight within a preference level. Tunnel selection proceeds from greatest to least weight. When disabled, selection defaults to a round robin method. Not displayed for LNS on M Series routers.
`Destination equal load balancing`	State of this tunnel selection method on the LAC. When enabled, the LAC selects tunnels based on the session count for destinations and the tunnel session count. Not displayed for LNS on M Series routers.
`Tunnel authentication challenge`	State of tunnel authentication, indicating whether the LAC and LNS exchange an authentication challenge and response during the establishment of the tunnel. The state is `Enabled` when a secret is configured in the tunnel profile or on the RADIUS server in the Tunnel-Password attribute [69]. The state is `Disabled` when the secret is not present. Not displayed for LNS on M Series routers.
`Calling number avp`	When the state is `Enabled`, the LAC includes the value of the Calling Number AVP 22 in ICRQ packets sent to the LNS. When the state is `Disabled`, the attribute is not sent to the LNS. Not displayed for LNS on M Series routers.

Table 1431: show services l2tp summary Output Fields *(Continued)*

Field Name	Field Description
`Failover Protocol`	When the state is enabled, the LAC operates in the default *failover-protocol-fall-back-to-silent-failover* manner. When the state is disabled, the `disable-failover-protocol` statement has been issued and the LAC operates only in silent failover mode. Not displayed for LNS on M Series routers.
`Tx connect speed method`	The connection speed method configured to send the speed values in the L2TP Tx Connect Speed (AVP 24) and L2TP Rx Connect Speed (AVP 38). Possible values are: • `actual` This is the default value in Junos OS Releases 15.1, 16.1, 16.2, and 17.1. It is deprecated in Junos Releases 17.2 and higher. • `ancp` • `none` • `pppoe-ia-tag` • `service-profile` • `static` This is the default value in Junos Releases 13.3, 14.1, 14.2, 17.2 and higher. It is deprecated in Junos OS Releases 15.1, 16.1, 16.2, and 17.1.
`Rx speed avp when equal`	Indicates if the Rx connect speed when equal configuration is `enabled` or `disabled`.
`Tunnel assignment id`	Format of the tunnel name. Format of the tunnel name, based on RADIUS attributes returned from the AAA server: • `authentication-id`—Name consists of only Tunnel Assignment-Id [82]. This is the default value. • `client-server-id`—Name is a combination of Tunnel-Client-Auth-Id [90], Tunnel-Server-Endpoint [67], and Tunnel-Assignment-Id [82]. This format is available only on MX Series routers.

Table 1431: show services l2tp summary Output Fields *(Continued)*

Field Name	Field Description
`Tunnel Tx Address Change`	Action taken by LAC when it receives a request from a peer to change the destination IP address, UDP port, or both: • `accept`—Accepts change requests for the IP address or UDP port. This is the default action. • `ignore`—Ignores all change requests. • `ignore-ip-address`—Ignores change requests for the IP address but accepts them for the UDP port. • `ignore-udp-port`—Ignores change requests for the UDP port but accepts them for the IP address.
`Min Retransmission Timeout for control packets`	Minimum number of seconds that the local peer waits for the initial response after transmitting an L2TP control packet. If no response has been received by the time the period expires, the local peer retransmits the packet.
`Min Retransmission Timeout for control packets`	Minimum number of seconds that the local peer waits for the initial response after transmitting an L2TP control packet. If no response has been received by the time the period expires, the local peer retransmits the packet.
`Max Retransmissions for Established Tunnel`	Maximum number of times control messages are retransmitted for established tunnels.
`Max Retransmissions for Not Established Tunnel`	Maximum number of times control messages are retransmitted for tunnels that are not established.
`Tunnel Idle Timeout`	Period that a tunnel can be inactive–that is, carrying no traffic–before it times out and is torn down.
`Destruct Timeout`	Period that the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed.
`Reassembly Service Set`	Indicates active IP reassembly configured for the interface.

Table 1431: show services l2tp summary Output Fields *(Continued)*

Field Name	Field Description
`Destination Lockout Timeout`	Timeout period for which all future destinations are locked out, meaning that they are not considered for selection when a new tunnel is created.
`Access Line Information`	State of LAC global configuration for forwarding subscriber line information to the LNS, `Enabled` or `Disabled`. Indicates active IP reassembly configured for the interface. Starting in Junos OS Release 17.4R1, this information can also be displayed on the LNS for information it receives from the LAC.
`IPv6 Services for LAC Sessions`	State of LAC IPv6 service configuration for creating the IPv6 (inet6) address family for LAC subscribers, allowing the application of IPv6 firewall filters, `Enabled` or `Disabled`.
`Speed Updates`	State of LAC global configuration for including connection speed updates when it forwards subscriber line information to the LNS, `Enabled` or `Disabled`. Starting in Junos OS Release 17.4R1, this information can also be displayed on the LNS for updates it receives from the LAC.
`Destinations`	Number of L2TP destinations for the LAC. Not displayed for LNS on M Series routers.
`Tunnels`	Number of L2TP tunnels established on the router.
`Sessions`	Number of L2TP sessions established on the router.
`Switched sessions`	Number of L2TP tunnel-switched sessions established on the router.
`Control`	Count of L2TP control packets and bytes sent and received.
`Data`	Count of L2TP data packets and bytes sent and received.

Table 1431: show services l2tp summary Output Fields *(Continued)*

Field Name	Field Description
Errors	Count of L2TP error packets and bytes sent and received.

Sample Output

show services l2tp summary (LAC on M Series routers)

```
user@host> show services l2tp summary
Administrative state is Drain
Failover within a preference level is Disabled
Weighted load balancing is Enabled
Destination equal load balancing is Disabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Disabled
Tunnel assignment id format is authentication-id
Destinations: 1 Tunnels: 1, Sessions: 1
      Tx packets     Rx packets  Memory (bytes)
Control    260             144              11513856
Data       7.5k           16.9k             8.3k
Errors        0              0
```

show services l2tp summary (LAC on MX Series routers)

```
user@host> show services l2tp summary
Administrative state is Drain
      Failover within a preference level is Disabled
  Weighted load balancing is Disabled
  Destination equal load balancing is Enabled
  Tunnel authentication challenge is Enabled
  Calling number avp is Enabled
  Failover Protocol is Disabled
  Tx Connect speed method is static
  Rx speed avp when equal is enabled
  Tunnel Tx Address Change is Accept
```

```
  Min Retransmissions Timeout for control packets is 2 seconds
  Max Retransmissions for Established Tunnel is 7
  Max Retransmissions for Not Established Tunnel is 5
  Tunnel Idle Timeout is 60 seconds
  Destruct Timeout is 300 seconds
  Destination Lockout Timeout is 300 seconds
  Reassembly Service Set is ssnr3
  Access Line Information is Enabled, Speed Updates is Enabled
  IPv6 Services For LAC Sessions is Enabled
  Destinations: 0, Tunnels: 0, Sessions: 0, Switched sessions: 0
```

show services l2tp summary (LNS on MX Series routers)

```
user@host show services l2tp summary
 Administrative state is Drain
  Failover within a preference level is Disabled
  Weighted load balancing is Disabled
  Destination equal load balancing is Disabled
  Tunnel authentication challenge is Enabled
  Calling number avp is Enabled
  Failover Protocol is Enabled
  Tx Connect speed method is static
  reassembly Service Set is ssnr3
  Destinations: 4, Tunnels: 19, Sessions: 65, Switched sessions: 2
  Access Line Information is Enabled, Speed Updates is Enabled
```

show services l2tp summary (LNS on M Series routers)

```
user@host> show services l2tp summary
Tunnels: 2, Sessions: 2, Errors: 0
    Tx packets   Rx packets    Memory (bytes)
Control     6k           9k                 688k
Data       70k          70k                 3054
```

show services l2tp summary statistics (MX Series routers)

```
user@host>show services l2tp summary statistics
Administrative state is Drain
```

```
Failover within a preference level is Disabled
Weighted load balancing is Disabled
Destination equal load balancing is Disabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Enabled
Tx Connect speed method is advisory
Tunnel assignment id format is assignment-id
Tunnel Tx Address Change is Accept
Min Retransmissions Timeout for control packets is 4 seconds
Max Retransmissions for Established Tunnel is 7
Max Retransmissions for Not Established Tunnel is 5
Tunnel Idle Timeout is 60 seconds
Destruct Timeout is 300 seconds
Destination Lockout Timeout is 300 seconds
Destinations: 1, Tunnels: 1, Sessions: 31815, Switched sessions: 0
                    Tx packets    Rx packets  Memory (bytes)
  Control          90.4k           32.0k       245678080
  Data             127.3k          100.8kk                  0
  Errors                 0                     0
```

Release Information

Command introduced before Junos OS Release 7.4.

Support for LAC on MX Series routers introduced in Junos OS Release 10.4.

Support for LNS on MX Series routers introduced in Junos OS Release 11.4.

Support for **statistics** option introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

- L2TP Services Configuration Overview
- L2TP Minimum Configuration

show services l2tp tunnel

IN THIS SECTION

- Syntax | 25685
- Description | 25685
- Options | 25686
- Required Privilege Level | 25686
- Output Fields | 25686
- Sample Output | 25690
- Release Information | 25694

Syntax

```
show services l2tp tunnel
<brief | detail | extensive>
<interface sp-fpc/pic/port>
<local-gateway gateway-address>
<local-gateway-name gateway-name>
<local-tunnel-id tunnel-id>
<peer-gateway gateway-address>
<peer-gateway-name gateway-name>
<statistics>
<tunnel-group group-name>
```

Description

(M10i and M7i routers only) Display information about active Layer 2 Tunneling Protocol (L2TP) tunnels for LNS.

(MX Series routers only) Display information about L2TP tunnels for LAC and LNS; the tunnels may or may not have active sessions.

Options

none	Display standard information about all active L2TP tunnels.
brief | detail | extensive	(Default) Display the specified level of output.
interface sp-*fpc/pic/port*	(Optional) Display L2TP tunnel information for only the specified adaptive services interface. This option is not available for L2TP on MX Series routers.
local-gateway *gateway-address*	(Optional) Display L2TP tunnel information for only the specified local gateway address.
local-gateway-name *gateway-name*	(Optional) Display L2TP tunnel information for only the specified local gateway name.
local-tunnel-id *tunnel-id*	(Optional) Display L2TP tunnel information for only the specified local tunnel identifier.
peer-gateway *gateway-address*	(Optional) Display L2TP tunnel information for only the specified peer gateway address.
peer-gateway-name *gateway-name*	(Optional) Display L2TP tunnel information for only the specified peer gateway name.
statistics	(Optional) Display the number of control packets and bytes transmitted and received for the tunnel. The statistics for a tunnel are retained until the tunnel is disconnected, rather than until the last session in the tunnel is cleared. Retaining the statistics enables them to increment in the event a new session subsequently uses the tunnel. You cannot include this option with any of the level options, `brief`, `detail`, or `extensive`.
tunnel-group *group-name*	(Optional) Display L2TP tunnel information for only the specified tunnel group.

Required Privilege Level

view

Output Fields

Table 1432 on page 25687 lists the output fields for the `show services l2tp tunnel` command. Output fields are listed in the approximate order in which they appear.

Table 1432: show services l2tp tunnel Output Fields

Field Name	Field Description
`Interface`	(LNS only) Name of an adaptive services interface.
`Tunnel group`	(LNS only) Name of a tunnel group.
`Local ID`	On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS. On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.
`Remote ID`	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.
`Remote IP`	IP address of the peer endpoint of the tunnel.
`Sessions`	Number of L2TP sessions established through the tunnel.

Table 1432: show services l2tp tunnel Output Fields *(Continued)*

Field Name	Field Description
`State`	State of the L2TP tunnel: • `cc_responder_accept_new`—The tunnel has received and accepted the start control connection request (SCCRQ). • `cc_responder_reject_new`—The tunnel has received and rejected the SCCRQ. • `cc_responder_idle`—The tunnel has just been created. • `cc_responder_wait_ctl_conn`—The tunnel has sent the start control connection response (SCCRP) and is waiting for the start control connection connected (SCCCN) message. • `clean-up`—The tunnel is being cleaned up. • `closed`—The tunnel is being closed. • `destroyed`—The tunnel is being destroyed. • Drain—Creation of new sessions and destinations is disabled for this tunnel. • `Established`—The tunnel is operating. This is the only state supported for the LAC. • `Terminate`—The tunnel is terminating. • `Unknown`—The tunnel is not connected to the router.
`Tunnel Name`	(LAC only) Name of the created tunnel. This value includes the destination name followed by the value of the RADIUS Tunnel-Assignment-ID VSA [82].
`Local IP`	IP address of the local endpoint of the tunnel.
`Local name`	Name used for local tunnel endpoint during tunnel negotiation.
`Remote name`	Name used for remote tunnel endpoint during tunnel negotiation.

Table 1432: show services l2tp tunnel Output Fields *(Continued)*

Field Name	Field Description
`Effective Peer Resync Mechanism`	(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel: • Failover protocol • Silent failover—Recovery takes place in the failed endpoint only using the proprietary silent failover protocol.
`Nas Port Method`	NAS port method (type), which indicates whether the LAC sends Cisco NAS Port Info AVP (100) in ICRQs to the LNS: • `cisco-avp`—sends the AVP. • `none`—does not send the AVP.
`Tunnel Logical System`	Logical system in which the L2TP tunnel is brought up.
`Tunnel Routing Instance`	Routing instance in which the L2TP tunnel is brought up.
`Max sessions`	Maximum number of sessions that can be established on this tunnel. The displayed limit for configured sessions is set to the lowest of the following configured session values for either LAC or LNS: • Global (chassis)—`set services l2tp tunnel maximum-sessions` *number* • Tunnel profile (individual tunnel)—`set access tunnel-profile` *profile-name* `tunnel` *tunnel-id* `max-sessions` *number*] • RADIUS—Value of VSA 26–33, Tunnel-Max-Sessions For LNS only, the following configuration is also considered: • Host profile—`access profile l2tp-profile client default l2tp maximum-sessions-per-tunnel`
`Window size`	Number of control messages that can be sent without receipt of an acknowledgment.

Table 1432: show services l2tp tunnel Output Fields *(Continued)*

Field Name	Field Description
`Hello interval`	Interval between the transmission of hello messages, in seconds.
`Create time`	Date and time when the tunnel was created. While the LNS and LAC are connected, this value should correspond to the when the call was created. If connection to the LAC is severed, the State changes to `Unknown` and the `Create time` value resets.
`Up time`	Amount of time elapsed since the tunnel became active, in hours, minutes, and seconds.
`Idle time`	Amount of time elapsed since the tunnel became idle, in hours, minutes, and seconds.
`Statistics since`	Date and time when collection of the following statistics began: • `Control Tx`—Amount of control information transmitted, in packets and bytes. • `Control Rx`—Amount of control information received, in packets and bytes. • `Data Tx`—Amount of data transmitted, in packets and bytes. • `Data Rx`—Amount of data received, in packets and bytes. • `Errors Tx`—Number of errors transmitted, in packets. • `Errors Rx`—Number of errors received, in packets.

Sample Output

show services l2tp tunnel (LAC)

```
user@host> show services l2tp tunnel
  Local ID  Remote ID  Remote IP               Sessions  State
     17185          1  203.0.113.101:1701               1  Established
```

show services l2tp tunnel detail (LAC)

```
user@host> show services l2tp tunnel detail
  Tunnel local ID: 31889, Tunnel remote ID:     1
    Remote IP: 203.0.113.101:1701
    Sessions: 1, State: Established
    Tunnel Name: 1/tunnel-to-LNS-1
    Local IP: 192.0.2.2:1701
    Local name: ce-lac, Remote name: ce-lns
    Effective Peer Resync Mechanism: silent failover
```

show services l2tp tunnel detail (LAC on MX Series Routers)

```
user@host> show services l2tp tunnel detail
Tunnel local ID: 17301, Tunnel remote ID: 1
    Remote IP: 203.0.113.101:1701
    Sessions: 1, State: Established
    Tunnel Name: 2/tunnel-to-LNS-2
    Local IP: 192.0.2.2:1701
    Local name: ce-lac, Remote name: ce-lns
    Effective Peer Resync Mechanism: silent failover
    Tunnel Logical System: default, Tunnel Routing Instance: default
```

show services l2tp tunnel detail (LNS on MX Series Routers)

```
user@host> show services l2tp tunnel detail
Tunnel local ID: 17301, Tunnel remote ID: 1
    Remote IP: 198.51.100.15:1701
    Sessions: 1, State: Established
    Tunnel Name: 2/2
    Local IP: 198.51.100.5:1701
    Local name: ce-bras-mx240-e, Remote name: testlac2
    Effective Peer Resync Mechanism: silent failover
    Tunnel Logical System: default, Tunnel Routing Instance: vrf1
```

show services l2tp tunnel extensive (LAC)

```
user@host> show services l2tp tunnel extensive
  Tunnel local ID: 17185, Tunnel remote ID:     1
    Remote IP: 203.0.113.101:1701
    Sessions: 1, State: Established
    Tunnel Name: 2/tunnel-to-LNS-2
    Local IP: 192.0.2.22:1701
    Local name: ce-lac, Remote name: ce-lns
    Effective Peer Resync Mechanism: failover protocol
    Max sessions: 32000, Window size: 4, Hello interval: 60
    Create time: Tue Nov  9 15:23:29 2010, Up time: 00:00:26
    Idle time: 00:00:00
```

show services l2tp tunnel extensive (LNS on M Series Routers)

```
user@host> show services l2tp tunnel extensive
Interface: sp-1/2/0, Tunnel group: group1
  Tunnel local ID: 62746, Tunnel remote ID: 16930
    Remote IP: 203.0.113.202:1701
    Sessions: 1, State: Established
    Local IP: 203.0.113.121:1701
    Local name: router-1, Remote name: router-2
    Max sessions: 50, Window size: 32, Hello interval: 60
    Create time: Tue Mar 23 14:13:15 2004, Up time: 01:14:58
    Idle time: 00:00:07
    Statistics since: Tue Mar 23 14:13:13 2004
                        Packets        Bytes
      Control Tx             80         1152
      Control Rx              3          272
      Data Tx                 0            0
      Data Rx               450        28.0k
      Errors Tx               0
      Errors Rx               0

Interface: sp-1/2/0, Tunnel group: group_company_dns
  Tunnel local ID: 37266, Tunnel remote ID: 36217
    Remote IP: 203.0.113.222:1701
    Sessions: 1, State: Established
    Local IP: 203.0.113.111:1701
```

```
    Local name: router-1, Remote name: router-2
    Max sessions: unlimited, Window size: 32, Hello interval: 60
    Create time: Tue Mar 23 14:13:15 2004, Up time: 01:14:59
    Idle time: 01:14:55
    Statistics since: Tue Mar 23 14:13:13 2004
                        Packets        Bytes
      Control Tx             81         1164
      Control Rx              3          273
      Data Tx                 0            0
      Data Rx                 1           80
      Errors Tx               0
      Errors Rx               0
```

show services l2tp tunnel extensive (LNS on MX Series Routers)

```
user@host> show services l2tp tunnel extensive
Tunnel local ID: 40553, Tunnel remote ID: 1
    Remote IP: 192.0.2.3:1701
    Sessions: 1, State: Established
    Tunnel Name: 3/1838
    Local IP: 203.0.113.2:1701
    Local name: lns-mx960, Remote name: testlac
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: vrf1
    Max sessions: 60000, Window size: 4, Hello interval: 60
    Create time: Mon Apr 25 20:27:50 2011, Up time: 00:01:11
    Idle time: 00:00:00, ToS Reflect: Enabled
    Tunnel Group Name: tg1
    Statistics since: Mon Apr 25 20:27:50 2011
                        Packets        Bytes
      Control Tx              4          219
      Control Rx              4          221
      Data Tx                 0            0
      Data Rx                 6           64
      Errors Tx               0
      Errors Rx               0
```

show services l2tp tunnel statistics (MX Series Routers)

```
user@host>show services l2tp tunnel statistics
Tunnel local ID: 17185, Tunnel remote ID: 1
Sessions: 31.8k, State: Established
Statistics since: Mon Aug 1 13:21:38 2011
                    Packets         Bytes
Control Tx           90.3k            9.0M
Control Rx           32.0k            1296.9k
Data Tx          127.3k         1591.6k
Data Rx             100.8k         1273.4k
Errors Tx            0
Errors Rx               0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- L2TP Services Configuration Overview
- L2TP Minimum Configuration

show services l2tp tunnel-group

IN THIS SECTION

- Syntax | 25695
- Description | 25695
- Options | 25695
- Required Privilege Level | 25695
- Output Fields | 25695
- Sample Output | 25696
- Release Information | 25696

Syntax

```
show services l2tp tunnel-group
<group-name>
```

Description

Display information about all L2TP tunnel groups or a specific L2TP tunnel group.

Options

group-name	(Optional) Name of a tunnel group.

Required Privilege Level

view

Output Fields

Table 1433 on page 25695 lists the output fields for the `show services l2tp tunnel-group` command. Output fields are listed in the approximate order in which they appear.

Table 1433: show services l2tp tunnel-group Output Fields

Field Name	Field Description
`Tunnel-group`	Name of a tunnel group.
`Tunnels`	Number of tunnels in the tunnel group.
`Sessions`	Number of L2TP sessions established for tunnels in the tunnel group.

Sample Output

show services l2tp tunnel-group

```
user@host> show services l2tp tunnel-group
Tunnel-group            Tunnels        Sessions
l2tp-tunnel-group1      2              20
l2tp-tunnel-group2      3              120
```

show services l2tp tunnel-group (Group Name)

```
user@host> show services l2tp tunnel-group l2tp-tunnel-group1
Tunnel-group            Tunnels        Sessions
l2tp-tunnel-group1      2              20
```

Release Information

Command introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

show services l2tp client | 25638

show services l2tp session-limit-group | 25674

L2TP LAC Tunneling for Subscribers

show services l2tp tunnel-switch destination

IN THIS SECTION

- Syntax | 25697
- Description | 25697
- Options | 25697

- Required Privilege Level | 25697
- Output Fields | 25697
- Sample Output | 25700
- Release Information | 25702

Syntax

```
show services l2tp tunnel-switch destination
< detail | extensive>
<statistics>
```

Description

Display information about L2TP switched tunnel destinations.

Options

none	Display standard information for all L2TP switched tunnel destinations.
detail | extensive	(Optional) Display the specified level of information.
statistics	(Optional) Display the number of control packets and bytes transmitted and received for the destination. You cannot include this option with either of the level options, `detail` or `extensive`.

Required Privilege Level

view

Output Fields

Table 1434 on page 25698 lists the output fields for the `show services l2tp tunnel-switch destination` command. Output fields are listed in the approximate order in which they appear.

Table 1434: show services l2tp tunnel-switch destination Output Fields

Field Name	Field Description	Level of Output
`Local Name`	Name of this destination.	All levels
`Remote IP`	IP address of the remote peer (LNS).	All levels
`Tunnels`	Number of tunnel connections for the destination in the following categories: • total • active • failed	All levels for total `extensive` for active and failed
`Sessions`	Number of session connections for the destination in the following categories: • total • active • failed	All levels for total `extensive` for active and failed
`Switched-sessions`	Number of L2TP sessions established by tunnel switching.	All levels
`State`	Administrative state of the L2TP destination: • `Enabled`—No restrictions exist on creation or operation of sessions and tunnels for this destination. • `Disabled`—Existing sessions and tunnels for this destination have been disabled and no new sessions or tunnels are created while in the `Disabled` state.	All levels
`Local IP`	IP address of the local gateway (LAC).	`detailextensive`
`Transport`	Medium used for tunneling. Only `ipUdp` is supported.	`detailextensive`

Table 1434: show services l2tp tunnel-switch destination Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Logical System`	Logical system in which the tunnel is configured.	`detailextensive`
`Router Instance`	Routing instance in which the tunnel is configured.	`detailextensive`
`Lockout State`	Reachability state of the destination: • `not locked`—Destination is considered reachable. • `waiting for lockout timeout`—Destination is locked out by L2TP because it is unreachable, so no attempts are made to reach the destination until the lockout timeout (300 seconds) expires, unless this is the only destination available for tunneling the subscriber.	`detailextensive`
`Connections`	Number of total, active, and failed tunnel and session connections for the destination.	`extensive`
`Control Tx`	Amount of control information transmitted, in packets and bytes.	`extensivestatistics`
`Control Rx`	Amount of control information received, in packets and bytes.	`extensivestatistics`
`Data Tx`	Amount of data transmitted, in packets and bytes.	`extensivestatistics`
`Data Rx`	Amount of data received, in packets and bytes.	`extensivestatistics`
`Error Tx`	Number of errors transmitted, in packets.	`extensivestatistics`
`Error Rx`	Number of errors received, in packets.	`extensivestatistics`

Sample Output

show services l2tp tunnel-switch destination

```
user@host> show services l2tp tunnel-switch destination
  Local Name   Remote IP       Tunnels  Sessions Switched-sessions State
  1            192.0.2.3       1        1        1                 Enabled
  2            203.0.113.10    1        1        1                 Enabled
```

show services l2tp tunnel-switch destination detail

```
user@host> show services l2tp tunnel-switch destination detail
  Local name: 1
    Remote IP: 192.0.2.3
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP: 203.0.113.51
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
  Local name: 2
    Remote IP: 198.51.100.10
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP: 203.0.113.31
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
```

show services l2tp tunnel-switch destination extensive

```
user@host> show services l2tp tunnel-switch destination extensive
Waiting for statistics...
  Local name: 1
    Remote IP: 192.0.2.3
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP: 203.0.113.51
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
```

```
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
                        Packets        Bytes
      Control Tx              6          239
      Control Rx              6          267
      Data Tx                67          815
      Data Rx                 0            0
      Errors Tx               0
      Errors Rx               0
  Local name: 2
    Remote IP: 198.51.100.10
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP:203.0.113.31
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
                        Packets        Bytes
      Control Tx              7          462
      Control Rx              6          171
      Data Tx                 0            0
      Data Rx                66          798
      Errors Tx               0
      Errors Rx               0
```

show services l2tp tunnel-switch destination statistics

```
user@host> show services l2tp tunnel-switch destination statistics
Waiting for statistics...
  Local name: 2, Tunnels: 1, Sessions: 1
                        Packets        Bytes
      Control Tx              5          452
      Control Rx              4          147
      Data Tx                 0            0
      Data Rx                 4           54
      Errors Tx               0
      Errors Rx               0
  Local name: 1, Tunnels: 1, Sessions: 1
```

```
                        Packets        Bytes
    Control Tx                4          184
    Control Rx                4          243
    Data Tx                   5           71
    Data Rx                   0            0
    Errors Tx                 0
    Errors Rx                 0
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

- show services l2tp tunnel-switch session | 25702
- show services l2tp tunnel-switch summary | 25709
- show services l2tp tunnel-switch tunnel | 25711

show services l2tp tunnel-switch session

IN THIS SECTION

- Syntax | 25703
- Description | 25703
- Options | 25703
- Required Privilege Level | 25703
- Output Fields | 25703
- Sample Output | 25707
- Release Information | 25708

Syntax

```
show services l2tp tunnel-switch session
<detail | extensive>
<statistics>
```

Description

Display information about L2TP switched tunnel sessions.

Options

none Display standard information about all active L2TP switched tunnel sessions.

detail | extensive (Optional) Display the specified level of output.

statistics (Optional) Display the number of control packets and bytes transmitted and received for the session. You cannot include this option with either of the level options, `detail` or `extensive`.

Required Privilege Level

view

Output Fields

Table 1435 on page 25703 lists the output fields for the `show services l2tp tunnel-switch session` command. Output fields are listed in the approximate order in which they appear.

Table 1435: show services l2tp tunnel-switch session Output Fields

Field Name	Field Description	Level of Output
`Tunnel local ID`	Identifier of the local endpoint of the tunnel, as assigned by the L2TP network server (LNS).	All levels
`Local ID`	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	none

Table 1435: show services l2tp tunnel-switch session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Remote ID`	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	none
`State`	State of the L2TP session: • `Established`—Session is operating. This is the only state supported for the LAC. • `closed`—Session is being closed. • `destroyed`—Session is being destroyed. • `clean-up`—Session is being cleaned up. • `lns-ic-accept-new`—New session is being accepted. • `lns-ic-idle`—Session has been created and is idle. • `lns-ic-reject-new`—New session is being rejected. • `lns-ic-wait-connect`—Session is waiting for the peer's incoming call connected (ICCN) message.	All levels
`Interface unit`	Logical interface for this session.	All levels
`Interface Name`	(LNS only) Name of an adaptive services interface.	none
`Session local ID`	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	`detailextensive`
`Session remote ID`	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	`detailextensive`
`Tunnel switch profile name`	Name of a tunnel switch profile.	`detailextensive`

Table 1435: show services l2tp tunnel-switch session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Mode`	(LNS) Mode of the interface representing the session: `shared` or `exclusive`. (LAC) Mode of the interface representing the session: `shared` or `dedicated`. Only `dedicated` is currently supported for the LAC.	`detailextensive`
`Local IP`	IP address of local endpoint of the Point-to-Point Protocol (PPP) session.	`detailextensive`
`Remote IP`	IP address of remote endpoint of the PPP session.	`detailextensive`
`Local name`	For LNS, name of the LNS instance in which the session was created. For LAC, name of the LAC.	`detailextensive`
`Remote name`	For LNS, name of the LAC from which the session was created. For LAC, name of the LAC instance.	`detailextensive`
`Bearer type`	Type of bearer enabled: • `0`—Might indicate that the call was not received over a physical link (for example, when the LAC and PPP are located in the same subsystem). • `1`—Digital access requested. • `2`—Analog access requested. • `4`—Asynchronous Transfer Mode (ATM) bearer support.	`extensive`
`Framing type`	Type of framing enabled: • `1`—Synchronous framing • `2`—Asynchronous framing	`extensive`
`LCP renegotiation`	(LNS only) Whether Link Control Protocol (LCP) renegotiation is configured: `On` or `Off`.	`extensive`

Table 1435: show services l2tp tunnel-switch session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Authentication`	Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).	`extensive`
`Interface ID`	(LNS only) Identifier used to look up the logical interface for this session.	`extensive`
`Call serial number`	Unique serial number assigned to the call.	`extensive`
`Tx speed`	Transmit speed of the session conveyed from the LAC to the LNS, in bits per second (bps).	`extensive`
`Rx speed`	Receive speed of the session conveyed from the LAC to the LNS, in bits per second (bps).	`extensive`
`Create time`	Day, date, and time when the call was created.	`extensive`
`Up time`	Length of time elapsed since the call became active, in hours, minutes, and seconds.	`extensive`
`Idle time`	Length of time elapsed since the call became idle, in hours, minutes, and seconds.	`extensive`
`ToS Reflect`	Status of IP ToS value reflection, `Disabled` or `Enabled`.	`extensive`
`Statistics since`	Date and time when collection of the following statistics began: • `Data Tx`—Amount of data transmitted, in packets and bytes. • `Data Rx`—Amount of data received, in packets and bytes.	`extensive`

Sample Output

show services l2tp tunnel-switch session

```
user@host> show services l2tp tunnel-switch session
Tunnel local ID: 37602
  Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
  13545  1        Established          1073741842        si-2/1/0

Tunnel local ID: 37060
  Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
  58296  1        Established          1073741843        si-2/1/0
```

show services l2tp tunnel-switch session detail

```
user@host> show services l2tp tunnel-switch session detail
Tunnel local ID: 37602
  Session local ID: 13545, Session remote ID: 1, Interface unit: 1073741842
    State: Established, Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.51:1701, Remote IP: 192.0.2.3:1701
    Local name: ce-bras-mx240-f, Remote name: testlac

Tunnel local ID: 37060
  Session local ID: 58296, Session remote ID: 1, Interface unit: 1073741843
    State: Established, Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.31:1701, Remote IP: 198.51.100.10:1701
    Local name: lns, Remote name: lns
```

show services l2tp tunnel-switch session extensive

```
user@host> show services l2tp tunnel-switch session extensive
Tunnel local ID: 37602
```

```
  Session local ID: 13545, Session remote ID: 1
    Interface unit: 1073741842
    State: Established
    Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.51:1701, Remote IP: 192.0.2.3:1701
    Local name: ce-bras-mx240-f, Remote name: testlac
    Bearer type: 2, Framing type: 1
    LCP renegotiation: On, Authentication: None, Interface ID: si-2/1/0
    Call serial number: 0
    Tx speed: 56000, Rx speed: 0
    Create time: Fri Jan 18 03:01:11 2013, Up time: 00:06:50
    Idle time: N/A, ToS Reflect: Disabled
    Statistics since: Fri Jan 18 03:01:11 2013
                        Packets        Bytes
      Data Tx                85         1031
      Data Rx                 0            0

Tunnel local ID: 37060
  Session local ID: 58296, Session remote ID: 1
    Interface unit: 1073741843
    State: Established
    Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.31:1701, Remote IP: 198.51.100.10:1701
    Local name: lns, Remote name: lns
    Bearer type: 2, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Call serial number: 0
    Tx speed: 56000, Rx speed: 0
    Create time: Fri Jan 18 03:01:14 2013, Up time: 00:06:48
    Idle time: N/A
    Statistics since: Fri Jan 18 03:01:14 2013
                        Packets        Bytes
      Data Tx                 0            0
      Data Rx                84         1014
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

show services l2tp tunnel-switch destination | 25696
show services l2tp tunnel-switch summary | 25709
show services l2tp tunnel-switch tunnel | 25711

show services l2tp tunnel-switch summary

IN THIS SECTION

- Syntax | 25709
- Description | 25709
- Options | 25709
- Required Privilege Level | 25710
- Output Fields | 25710
- Sample Output | 25710
- Release Information | 25711

Syntax

```
show services l2tp tunnel-switch summary
<statistics>
```

Description

Display L2TP tunnel switch summary information.

Options

none Display complete L2TP switched tunnel summary information.

statistics (Optional) Display the number of control packets and bytes transmitted and received for all switched tunnels and sessions.

Required Privilege Level

view

Output Fields

Table 1436 on page 25710 lists the output fields for the `show services l2tp tunnel-switch summary` command. Output fields are listed in the approximate order in which they appear.

Table 1436: show services l2tp tunnel-switch summary Output Fields

Field Name	Field Description
`Tunnel switch profile name`	Name of a tunnel switch profile.
`LNS local session id`	Identifier assigned by the LNS function on the LTS to the local endpoint of the L2TP session originating on a remote LAC (the first session)
`LAC local session id`	Identifier assigned by the LAC function on the LTS to the local endpoint of the L2TP session originating on the LTS (the second session).
`LNS state`	State of the L2TP session (the first session) between a remote LAC and the LNS function on the LTS.
`LAC state`	State of the L2TP session (the second session) between the LAC function on the LTS and a remote LNS.

Sample Output

show services l2tp tunnel-switch summary

```
user@host> show services l2tp tunnel-switch summary
Tunnel switch profile name: ce-lts-profile
  LNS local   LAC local   LNS state      LAC state      Interface
  session ID  session ID                                name
  13545       58296       established    established    si-2/1/0
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

- show services l2tp tunnel-switch destination | 25696
- show services l2tp tunnel-switch session | 25702
- show services l2tp tunnel-switch tunnel | 25711

show services l2tp tunnel-switch tunnel

IN THIS SECTION

- Syntax | 25711
- Description | 25711
- Options | 25712
- Required Privilege Level | 25712
- Output Fields | 25712
- Sample Output | 25716
- Release Information | 25717

Syntax

```
show services l2tp tunnel-switch tunnel
<detail | extensive>
<statistics>
```

Description

Display information about L2TP switched tunnels.

Options

none Display standard information about all active L2TP tunnels.

detail | extensive (Default) Display the specified level of output.

statistics (Optional) Display the number of control packets and bytes transmitted and received for the tunnel. You cannot include this option with either of the level options, `detail` or `extensive`.

Required Privilege Level

view

Output Fields

Table 1437 on page 25712 lists the output fields for the `show services l2tp tunnel-switch tunnel` command. Output fields are listed in the approximate order in which they appear.

Table 1437: show services l2tp tunnel-switch tunnel Output Fields

Field Name	Field Description	Level of Output
`Local ID`	On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS. On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.	none
`Remote ID`	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.	none
`Remote IP`	IP address of the peer endpoint of the tunnel.	All levels
`Sessions`	Number of L2TP sessions established through the tunnel.	All levels
`Switched-sessions`	Number of L2TP sessions established by tunnel switching.	All levels

Table 1437: show services l2tp tunnel-switch tunnel Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the L2TP tunnel: • cc_responder_accept_new—The tunnel has received and accepted the start control connection request (SCCRQ). • cc_responder_reject_new—The tunnel has received and rejected the SCCRQ. • cc_responder_idle—The tunnel has just been created. • cc_responder_wait_ctl_conn—The tunnel has sent the start control connection response (SCCRP) and is waiting for the start control connection connected (SCCCN) message. • clean-up—The tunnel is being cleaned up. • closed—The tunnel is being closed. • destroyed—The tunnel is being destroyed. • Established—The tunnel is operating. This is the only state supported for the LAC. • Terminate—The tunnel is terminating. • Unknown—The tunnel is not connected to the router.	All levels
Tunnel local ID	On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS. On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.	detailextensive
Tunnel remote ID	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.	detailextensive

Table 1437: show services l2tp tunnel-switch tunnel Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Tunnel Name`	(LAC only) Name of the created tunnel. This value includes the destination name followed by the value of the RADIUS Tunnel-Assignment-ID VSA [82].	`detailextensive`
`Local IP`	IP address of the local endpoint of the tunnel.	`detailextensive`
`Local name`	Name used for local tunnel endpoint during tunnel negotiation.	`detailextensive`
`Remote name`	Name used for remote tunnel endpoint during tunnel negotiation.	`detailextensive`
`Effective Peer Resync Mechanism`	(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel: • Failover protocol • Silent failover—Recovery takes place in the failed endpoint only using the proprietary silent failover protocol.	`detailextensive`
`NAS Port Method`	(LAC only) Status of interoperation with Cisco LNS devices: • none—NAS port method is not enabled for interoperation. • cisco-avp—NAS port method is enabled for interoperation.	`detailextensive`
`Tunnel Logical System`	Logical system in which the L2TP tunnel is brought up.	`detailextensive`
`Tunnel Routing Instance`	Routing instance in which the L2TP tunnel is brought up.	`detailextensive`
`Max sessions`	Maximum number of sessions that can be established on this tunnel.	`extensive`
`Window size`	Number of control messages that can be sent without receipt of an acknowledgment.	`extensive`

Table 1437: show services l2tp tunnel-switch tunnel Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Hello interval`	Interval between the transmission of hello messages, in seconds.	`extensive`
`Create time`	Date and time when the tunnel was created. While the LNS and LAC are connected, this value should correspond to the router's uptime. If connection to the LAC is severed, the State changes to `Unknown` and the `Create time` value resets.	`extensive`
`Up time`	Amount of time elapsed since the tunnel became active, in hours, minutes, and seconds.	`extensive`
`Idle time`	Amount of time elapsed since the tunnel became idle, in hours, minutes, and seconds.	`extensive`
`ToS Reflect`	Status of IP ToS value reflection, `Disabled` or `Enabled`.	`extensive`
`Interface Name`	(LNS only) Name of an adaptive services interface.	`extensive`
`Tunnel Group Name`	(LNS only) Name of a tunnel group.	`extensive`
`Statistics since`	Date and time when collection of the following statistics began: • `Control Tx`—Amount of control information transmitted, in packets and bytes. • `Control Rx`—Amount of control information received, in packets and bytes. • `Data Tx`—Amount of data transmitted, in packets and bytes. • `Data Rx`—Amount of data received, in packets and bytes. • `Errors Tx`—Number of errors transmitted, in packets. • `Errors Rx`—Number of errors received, in packets.	`extensive`

Sample Output

show services l2tp tunnel-switch tunnel

```
user@host> show services l2tp tunnel-switch tunnel
  Local ID Remote ID Remote IP              Sessions Switched-sessions State
  37602    1         192.0.2.3:1701                 1                 1 Established
  37060    1         198.51.100.10:1701             1                 1 Established
```

show services l2tp tunnel-switch tunnel detail

```
user@host> show services l2tp tunnel-switch tunnel detail
  Tunnel local ID: 37602, Tunnel remote ID: 1
    Remote IP: 192.0.2.3:1701
    Sessions: 1, Switched sessions: 1, State: Established
    Tunnel Name: 1/1
    Local IP: 203.0.113.51:1701
    Local name: ce-bras-mx240-f, Remote name: testlac
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
  Tunnel local ID: 37060, Tunnel remote ID: 1
    Remote IP: 198.51.100.10:1701
    Sessions: 1, Switched sessions: 1, State: Established
    Tunnel Name: 2/1
    Local IP: 203.0.113.31:1701
    Local name: lns, Remote name: lns
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
```

show services l2tp tunnel-switch tunnel extensive

```
user@host> show services l2tp tunnel-switch tunnel extensive
Waiting for statistics...
  Tunnel local ID: 37602, Tunnel remote ID: 1
    Remote IP: 192.0.2.3:1701
    Sessions: 1, Switched sessions: 1, State: Established
```

```
    Tunnel Name: 1/1
    Local IP: 203.0.113.51:1701
    Local name: ce-bras-mx240-f, Remote name: testlac
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
    Max sessions: 128100, Window size: 4, Hello interval: 60
    Create time: Fri Jan 18 03:01:11 2013, Up time: 00:07:49
    Idle time: 00:00:00, ToS Reflect: Disabled
    Interface Name: si-2/1/0, Tunnel Group Name: ce-l2tp-tunnel-group
    Statistics since: Fri Jan 18 03:01:11 2013
                        Packets        Bytes
      Control Tx              7          259
      Control Rx              7          279
      Data Tx                97         1175
      Data Rx                 0            0
      Errors Tx               0
      Errors Rx               0
  Tunnel local ID: 37060, Tunnel remote ID: 1
    Remote IP: 198.51.100.10:1701
    Sessions: 1, Switched sessions: 1, State: Established
    Tunnel Name: 2/1
    Local IP: 203.0.113.31:1701
    Local name: lns, Remote name: lns
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
    Max sessions: 128100, Window size: 4, Hello interval: 60
    Create time: Fri Jan 18 03:01:14 2013, Up time: 00:07:46
    Idle time: 00:00:00
    Statistics since: Fri Jan 18 03:01:14 2013
                        Packets        Bytes
      Control Tx              8          482
      Control Rx              7          183
      Data Tx                 0            0
      Data Rx                96         1158
      Errors Tx               0
      Errors Rx               0
```

Release Information

Command introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

show services l2tp tunnel-switch destination | 25696

show services l2tp tunnel-switch session | 25702

show services l2tp tunnel-switch summary | 25709

show services l2tp user

IN THIS SECTION

- Syntax | 25718
- Description | 25718
- Options | 25718
- Required Privilege Level | 25719
- Output Fields | 25719
- Sample Output | 25722
- Release Information | 25723

Syntax

```
show services l2tp user
<brief | detail | extensive | statistics>
<user username>
```

Description

(M10i and M7i routers only) Display a list of active Layer 2 Tunneling Protocol (L2TP) users.

Options

none	Display all active L2TP users.

brief | detail | extensive | statistics	(Optional) Display the specified level of output. Use the **statistics** option to display L2TP user statistics.
user *username*	(Optional) Display L2TP user information for only the specified username.

Required Privilege Level

view

Output Fields

Table 1438 on page 25719 lists the output fields for the `show services l2tp user` command. Output fields are listed in the approximate order in which they appear.

Table 1438: show services l2tp user Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Tunnel group	Name of a tunnel group.
Tunnel local ID	Local identifier of the tunnel, as assigned by the L2TP network server (LNS).
Session local ID	Local identifier of the session, as assigned by the L2TP network server (LNS).
Session remote ID	Remote identifier of the session, as assigned by the L2TP access concentrator (LAC).

Table 1438: show services l2tp user Output Fields *(Continued)*

Field Name	Field Description
State	State of the L2TP session: • **Established**—The session is operating. • **closed**—The session is being closed. • **destroyed**—The session is being destroyed. • **clean-up**—The session is being cleaned up. • **lns-ic-accept-new**—A new session is being accepted. • **lns-ic-idle**—The session has been created and is idle. • **lns-ic-reject-new**—The new session is being rejected. • **lns-ic-wait-connect**—The session is waiting for the peer's incoming call connected (ICCN) message.
Mode	Mode of the interface representing the session: **shared** or **exclusive**.
Local IP	IP address of the local endpoint of the tunnel.
Remote IP	IP address of the peer endpoint of the tunnel.
Username	Name of the user logged in to the session.
Assigned IP address	IP address assigned to remote client.
Local name	Name of the local device.
Remote name	Name of the remote device.
Local MRU	Maximum receive unit (MRU) setting of the local device, in bytes.

Table 1438: show services l2tp user Output Fields *(Continued)*

Field Name	Field Description
Remote MRU	MRU setting of the remote device, in bytes.
Tx speed	Transmit speed of the tunnel session, in bps.
Rx speed	Receive speed of the tunnel session, in bps.
Bearer type	Type of bearer enabled: • **0**—Might indicate that the call was not received over a physical link (for example, when the LAC and PPP are located in the same subsystem) • **1**—Digital access requested • **2**—Analog access requested • **4**—Asynchronous Transfer Mode (ATM) bearer support
Framing type	Type of framing enabled: • **1**—Synchronous framing • **2**—Asynchronous framing
LCP renegotiation	Whether Link Control Protocol (LCP) renegotiation is configured: **On** or **Off**.
Authentication	Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).
Interface ID	Name of the logical unit.
Interface unit	Logical unit number.
Call serial number	Unique serial number assigned to the call.

Table 1438: show services l2tp user Output Fields *(Continued)*

Field Name	Field Description
Create time	Date and time when the call was created.
Up time	Amount of time elapsed since the call became active, in hours, minutes, and seconds.
Idle time	Amount of time elapsed since the call became idle, in hours, minutes, and seconds.
Statistics sine	Date and time when collection of the following statistics began: • **Control Tx**—Amount of control information transmitted, in packets and bytes. • **Control Rx**—Amount of control information received, in packets and bytes. • **Data Tx**—Amount of data transmitted, in packets and bytes. • **Data Rx**—Amount of data received, in packets and bytes. • **Errors Tx**—Number of errors transmitted, in packets. • **Errors Rx**—Number of errors received, in packets.

Sample Output

show services l2tp user extensive

```
user@host> show services l2tp user extensive
Interface: sp-1/2/0, Tunnel group: group1, Tunnel local ID: 62746
  Session local ID: 56793, Session remote ID: 53304
    State: Established, Mode: shared
    Local IP: 10.128.1.1:1701, Remote IP: 10.128.1.2:1701
    Username: usr1@example.com, Assigned IP address: 10.50.2.1/32
    Local name: router-1, Remote name: router-2
    Local MRU: 4000, Remote MRU: 1500, Tx speed: 64000, Rx speed: 64000
    Bearer type: 2, Framing type: 1
    LCP renegotiation: Off, Authentication: CHAP, Interface ID: unit_20
    Interface unit: 20, Call serial number: 4137941434
    Create time: Tue Mar 23 14:13:15 2004, Up time: 01:16:41
```

```
    Idle time: 00:00:00
    Statistics since: Tue Mar 23 14:13:13 2004
                        Packets        Bytes
      Control Tx              4           88
      Control Rx              2           28
      Data Tx                 0            0
      Data Rx               461        29.0k
      Errors Tx               0
      Errors Rx               0
Interface: sp-1/2/0, Tunnel group: group_company_dns, Tunnel local ID: 37266
  Session local ID: 39962, Session remote ID: 53303
    State: Established, Username: usr1@company_dns.com, Mode: shared
    Local IP: 10.128.11.1:1701, Remote IP: 10.128.11.2:1701
    Username: usr1@company_dns.com, Assigned IP address: 10.48.1.1/32
    Local name: router-1, Remote name: router-2
    Local MRU: 4470, Remote MRU: 4470, Tx speed: 155000000,
                                        Rx speed: 155000000
    Bearer type: 2, Framing type: 1
    LCP renegotiation: Off, Authentication: CHAP, Interface ID: unit_31
    Interface unit: 31, Call serial number: 4137941433
    Create time: Tue Mar 23 14:13:17 2004, Up time: 01:16:39
    Idle time: 01:16:36
    Statistics since: Tue Mar 23 14:13:15 2004
                        Packets        Bytes
      Control Tx              6          196
      Control Rx              4          150
      Data Tx                 0            0
      Data Rx                 1           80
      Errors Tx               0
      Errors Rx               0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

L2TP Services Configuration Overview

L2TP Minimum Configuration

show services link-services cpu-usage

IN THIS SECTION

- Syntax | 25724
- Description | 25724
- Options | 25724
- Required Privilege Level | 25724
- Output Fields | 25725
- Sample Output | 25727
- Release Information | 25730

Syntax

```
show services link-services cpu-usage
<brief | detail>
<interface interface-name>
```

Description

(M Series and T Series routers only) Display information about Link Services IQ (LSQ) CPU usage.

Options

none	Display standard information about CPU usage for all LSQ interfaces.
brief | detail	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display information about the specified LSQ interface.

Required Privilege Level

view

Output Fields

Table 1439 on page 25725 lists the output fields for the `show services link-services cpu-usage` command. Output fields are listed in the approximate order in which they appear.

Table 1439: show services link-services cpu-usage Output Fields

Field Name	Field Description	Level of Output
Role	CPU functional category.	**brief**
1 Second Average	Percentage of usage during 1-second duration.	All levels
5 Second Average	Percentage of usage during 5-second duration.	All levels
QoS	Quality of service (QoS) CPU, which takes care of queuing and scheduling of incoming IP packets on a per-bundle basis. It schedules packets with higher QoS values first.	All levels
Sequencer	Assigns sequence numbers to outgoing MLPPP fragments and interleaves link fragmentation and interleaving (LFI) traffic.	All levels
Load Balancer	Distributes load across different fragmenter CPUs.	All levels
Fragmenter	Main LSQ CPU; fragments IP packets into MLPPP fragments and also reassembles MLPPP fragments into IP packets.	All levels
Total	Sum of all CPU functions.	**brief**
Idle	Counts idle cycles when the CPU does not have any work.	**detail**
Timer	Takes care of periodic events driven by a timer, such as timeouts.	**detail**
System	System housekeeping thread.	**detail**

Table 1439: show services link-services cpu-usage Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Input (QoS)	Acquires and queues incoming IP frames from hardware interfaces.	**detail**
Output (QoS)	Sends scheduled frames to the next processing CPU.	**detail**
Output Frags (QoS)	Sends outstanding frames to the fragmenter CPU.	**detail**
Bypass (QoS)	Sends outstanding frames for LFI.	**detail**
Free frame (QoS)	Frees dropped frames.	**detail**
CPU*number*	Identifier number of specific CPU.	**detail**
Drop (Fragmenter)	Drops frames that have been marked by the QoS CPU.	**detail**
Frag (Fragmenter)	Fragments IP frames into MLPPP fragments.	**detail**
Reass (Fragmenter)	Reassembles MLPPP fragments into IP frames.	**detail**
Freeback (Fragmenter)	Handles freeback of credits from other CPUs (MultiServices PICs only).	**detail**
Input LFI (Sequencer)	Receives LFI traffic from QoS CPU and transmits it with strict priority over MLPPP.	**detail**
Input Frag (Sequencer)	Receives MLPPP fragments from fragmenter CPUs, assigns sequence numbers, and appends MLPPP headers.	**detail**
Output Frag (Sequencer)	Load-balances and transmits fragments across links.	**detail**

Table 1439: show services link-services cpu-usage Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Retry (Sequencer)	Retries transmission if hardware was busy in the previous attempt.	**detail**
Input Alloc (Load Balancer)	Acquires frames from hardware interfaces and validates them.	**detail**
Input (Load Balancer)	Performs error and sanity checks and check frames for PortMapping.	**detail**
Output (Load Balancer)	Sends frame to next processing CPU.	**detail**
Freeback (Load Balancer)	Handles freeback of credits from other CPUs.	**detail**

Sample Output

show services link-services cpu-usage brief (AS PIC)

```
user@host> show services link-services cpu-usage interface lsq-0/0/0 brief
Role             1 Second Average       5 Second Average
QOS                         1.0%                   1.0%
Sequencer                  0.1%                   0.1%
Fragmenter                  0.1%                   0.1%
Total                       0.1%                   0.1%
```

show services link-services cpu-usage brief (MultiServices PIC)

```
user@host> show services link-services cpu-usage interface lsq-0/0/0 brief
  Role           1 Second Average       5 Second Average

  QoS                       0.1%                   0.1%
  Fragmenter                0.1%                   0.1%
```

```
  Load Balancer              0.0%                    0.0%
  Total                      0.1%                    0.1%
```

show services link-services cpu-usage detail (AS PIC)

```
user@host> show services link-services cpu-usage interface lsq-0/0/0 detail

  QoS              Idle   Timer  System  Input Output Output Bypass  Free
                                                     Frags          frame

  CPU0            99.1%    0.9%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
  CPU1            99.8%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
  1 sec ave       99.5%    0.5%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
  5 sec ave       99.5%    0.5%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%

  Fragmenter       Idle   Timer  System   Drop   Frag  Reass   Free
                                                             back

  CPU0            96.6%    0.1%    0.0%   0.0%   0.0%   3.3%   0.0%
  CPU1            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU2            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU3            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU4            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU5            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU6            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU7            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU8            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  1 sec ave       99.5%    0.1%    0.0%   0.0%   0.0%   0.4%   0.0%
  5 sec ave       99.5%    0.1%    0.0%   0.0%   0.0%   0.4%   0.0%

  Sequencer        Idle   System  Input  Input Output  Retry
                                    LFI   Frag   Frag

  CPU0            99.9%     0.1%   0.0%   0.0%   0.0%   0.0%
  CPU1           100.0%     0.0%   0.0%   0.0%   0.0%   0.0%
  1 sec ave       99.9%     0.1%   0.0%   0.0%   0.0%   0.0%
  5 sec ave       99.9%     0.1%   0.0%   0.0%   0.0%   0.0%
```

show services link-services cpu-usage detail (MultiServices PIC)

```
user@host> show services link-services cpu-usage interface lsq-0/0/0 detail
  QoS              Idle   Timer  System  Input Output Output Bypass  Free
                                                      Frags         frame

  CPU0            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
  CPU1            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
  CPU2            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
  CPU3            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
  CPU4            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
  1 sec ave       99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
  5 sec ave       99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%

  Fragmenter       Idle   Timer  System   Drop   Frag  Reass   Free
                                                             back

  CPU0            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU1            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU2            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU3            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU4            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU5            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU6            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU7            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU8            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU9            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU10           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU11           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU12           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU13           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU14           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU15           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU16           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  CPU17           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  1 sec ave       99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
  5 sec ave       99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%

  Load-Balancer    Idle   System  Input  Input Output   Free
                                  Alloc                  back

  CPU0           100.0%     0.0%   0.0%   0.0%   0.0%   0.0%
```

```
CPU1           100.0%     0.0%   0.0%   0.0%   0.0%   0.0%
1 sec ave      100.0%     0.0%   0.0%   0.0%   0.0%   0.0%
5 sec ave      100.0%     0.0%   0.0%   0.0%   0.0%   0.0%
```

Release Information

Command introduced in Junos OS Release 8.4.

show services lrf collector statistics

IN THIS SECTION

- Syntax | 25730
- Description | 25730
- Options | 25731
- Required Privilege Level | 25731
- Output Fields | 25731
- Sample Output | 25732
- Release Information | 25732

Syntax

```
show services lrf collector statistics
<collector-name>
```

Description

Display LRF statistics for one or more collectors. If a collector is not specified, statistics are displayed for all collectors.

Options

none	Display LRF statistics for all collectors.
collector-name	(Optional) Display LRF statistics for the specified collector.

Required Privilege Level

view

Output Fields

Table 1440 on page 25731 lists the output fields for the `show services lrf collector statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1440: show services lrf collector statistics Output Fields

Field Name	Field Description
Interface	Name of the interface from which data records are sent to the collector.
Templates registered	Number of templates registered with the collector.
Template registration failures	Number of template registration failures.
Templates active	Number of active templates.
Sessions received	Number of data sessions received for logging of data.
Sessions ignored	Number of data sessions received for logging of data that were ignored.
Records logged	Number of logs sent to the collector.
Records exported	Number of data records exported to the collector.
Record export failures	Number of data record export attempts that failed.

Sample Output

show services lrf collector statistics

```
user@host> show services lrf collector statistics
LRF Collector Statistics
    Interface: ms-2/1/0
    Templates registered: 0, Template registration failures: 0, Templates active: 1
    Sessions received: 0, Sessions ignored: 0, Records logged: 0
    Records exported: 0, Record export failures: 0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Logging and Reporting Function for Subscribers

show services lrf rule statistics

IN THIS SECTION

- Syntax | 25733
- Description | 25733
- Options | 25733
- Required Privilege Level | 25733
- Output Fields | 25733
- Sample Output | 25734
- Release Information | 25734

Syntax

```
show services lrf rule statistics
<rule-name>
```

Description

Display LRF statistics for one or more LRF rules. If an LRF rule is not specified, statistics are displayed for all LRF rules.

Options

none Display LRF statistics for all LRF rules.

rule-name (Optional) Display LRF statistics for the specified LRF rule.

Required Privilege Level

view

Output Fields

Table 1441 on page 25733 lists the output fields for the `show services lrf rule statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1441: show services lrf rule statistics Output Fields

Field Name	Field Description
Interface	Name of the interface from which data records are sent to the collector.
Rule	Name of the LRF rule that caused data records to be exported to the collector.
Template	Name of the template that was used to export data records to the collector.
Templates registered	Number of templates registered with the collector.

Table 1441: show services lrf rule statistics Output Fields *(Continued)*

Field Name	Field Description
Template registration failures	Number of template registration failures.
Collector	Name of the collector to which data records were sent.
Sessions received	Number of data sessions received for logging of data.
Sessions ignored	Number of data sessions received for logging of data that were ignored.
Sessions logged	Number of data sessions that had data records exported to the collector.
Records exported	Number of data records exported to the collector.
Record export failures	Number of data record export attempts that failed.

Sample Output

show services lrf rule statistics

```
user@host> show services lrf rule statistics
  LRF Rule Statistics
    Interface: ms-3/1/0
    Rule: r1
    Template: temp1
    Templates registered: 2, Template registration failures: 0
    Collector: coll1
    Sessions received: 115, Sessions ignored: 0, Sessions logged: 134
    Records exported: 134, Record export failures: 0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Logging and Reporting Function for Subscribers

show services lrf statistics

IN THIS SECTION

- Syntax | 25735
- Description | 25735
- Required Privilege Level | 25735
- Output Fields | 25735
- Sample Output | 25736
- Release Information | 25736

Syntax

```
show services lrf statistics
```

Description

Display number of bytes, packets, and flows for carrying data records to the collector.

Required Privilege Level

view

Output Fields

Table 1442 on page 25736 lists the output fields for the `show services lrf statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1442: show services lrf statistics Output Fields

Field Name	Field Description
Interface	Name of the interface from which data records are sent to the collector.
Flow packets	Number of packets carrying data records to the collector.
Flow bytes	Number of bytes carrying data records to the collector.
Active flows	Number of active flows carrying data records to the collector.
Total flows	Total number of flows for carrying data records to the collector.

Sample Output

show services lrf statistics

```
user@host> show services lrf statistics
LRF Statistics
    Interface: ms-3/1/0
    Flow packets: 31125, Flow bytes: 15335751
    Active flows: 0, Total flows: 1887

    Interface: ms-3/2/0
    Flow packets: 0, Flow bytes: 0
    Active flows: 0, Total flows: 0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Logging and Reporting Function for Subscribers

show services lrf template

IN THIS SECTION

- Syntax | 25737
- Description | 25737
- Options | 25737
- Required Privilege Level | 25738
- Sample Output | 25738
- Release Information | 25739

Syntax

```
show services lrf template option
```

Description

Display the fields for a template type. You must specify a template type.

Options

option Specify one of the following template types:

- device-data—Display the fields for the Device Data template type.
- flow-id—Display the fields for the Flow ID template type.
- http—Display the fields for the HTTP template type.
- ifl-subscriber—Display the fields for the IFL Subscriber template type.
- ipflow—Display the fields for the IPFlow template type.
- ipflow-extended—Display the fields for the IPFlow Extended template type.
- ipflow-tcp—Displays the fields for the IPFlow TCP template type.

- ipflow-tcp-ts—Displays the fields for the IPFlow TCP Timestamp template type.
- ipflow-ts—Display the fields for the IPFlow Timestamp template type.
- ipv4—Display the fields for the IPv4 template type.
- ipv4-extended—Display the fields for the IPv4 Extended template type.
- ipv6—Display the fields for the IPv6 template type.
- ipv6-extended—Display the fields for the IPv6 Extended template type.
- l7-app—Display the fields for the L7 Application template type.
- mobile-subscriber—Display the fields for the Mobile Subscriber template type.
- pcc—Display the fields for the PCC template type.
- subscriber-data—Display the fields for the Subscriber Data template type.
- wireline-subscriber—Display the fields for the Wireline Subscriber template type.

Required Privilege Level

view

Sample Output

show services lrf template ipv4

```
user@host> show services lrf template ipv4
LRF Template fields
    Ipv4 source address
    Ipv4 destination address
    TCP/UDP source port
    TCP/UDP destination port
```

show services lrf template ipflow-extended

```
user@host> show services lrf template ipflow-extended
Field                   Element Id       Length(bytes)      Vendor
```

```
Service set name        520             16                 Juniper
Routing-instance        521             16                 Juniper
```

show services lrf template ipflow-tcp-ts

```
user@host> show services lrf template ipflow-tcp-ts
Field                            Element Id    Length(bytes)  Vendor
Smooth RTT uplink                10000         4              Juniper
Smooth RTT downlink              10001         4              Juniper
Client setup Time                10002         4              Juniper
Server Setup time                10003         4              Juniper
Client first payload timestamp   10004         8              Juniper
Upload time                      10005         4              Juniper
Server first payload timestamp   10006         8              Juniper
Download time                    10007         4              Juniper
Acknowledged volumes uplink      10008         8              Juniper
Acknowledged volumes downlink    10009         8              Juniper
```

show services lrf template ipflow-tcp

```
user@host> show services lrf template ipflow-tcp
Field                              Element Id  Length(bytes)  Vendor
Retransmitted TCP packets uplink   115         4              Juniper
Retransmitted TCP packets downlink 116         4              Juniper
TCP flow creation timestamp        121         8              Juniper
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Logging and Reporting Function for Subscribers

show services mobile-edge anchor class-of-service

IN THIS SECTION

- Syntax | 25740
- Description | 25740
- Required Privilege Level | 25740
- Output Fields | 25740
- Sample Output | 25741
- Release Information | 25742

Syntax

```
show services mobile-edge anchor class-of-service
```

Description

This command displays information about the configuration of the downlink forwarding queues.

Required Privilege Level

view

Output Fields

Table 1443 on page 25740 lists the output fields for the `show services mobile-edge anchor class-of-service` command.

Table 1443: show services mobile-edge anchor class-of-service Output Fields

Field Name	Field Description
Excess priority	Best-effort priority level assigned to a given queue.

Table 1443: show services mobile-edge anchor class-of-service Output Fields *(Continued)*

Field Name	Field Description
Guaranteed priority	Guaranteed bit-rate priority level assigned to a given queue.
Queue	Queue number and name.
Reserved bandwidth	Percentage of total bandwidth reserved for a given queue.
Weight	Importance assigned to each best-effort queue. This field is used to determine how often the system serves a queue.

Sample Output

show services mobile-edge anchor class-of-service

```
user@host> show services mobile-edge anchor class-of-service
Downlink forwarding queues/classes of service:
    Queue 0 (BE-LOW-0)
      Guaranteed priority:  low
      Excess priority:      low
      Reserved Bandwidth:   0 %
      Weight:               10

    Queue 1 (BE-LOW-1)
      Guaranteed priority:  low
      Excess priority:      low
      Reserved Bandwidth:   0 %
      Weight:               20

    Queue 2 (BE-HIGH-0)
      Guaranteed priority:  low
      Excess priority:      high
      Reserved Bandwidth:   0 %
      Weight:               20

    Queue 3 (BE-HIGH-1)
      Guaranteed priority:  low
      Excess priority:      high
```

```
      Reserved Bandwidth:   0 %
      Weight:               50

    Queue 4 (GBR-LOW)
      Guaranteed priority:  low
      Excess priority:      low
      Reserved Bandwidth:   30 %
      Weight:               0

    Queue 5 (GBR-MEDIUM)
      Guaranteed priority:  medium-high
      Excess priority:      low
      Reserved Bandwidth:   30 %
      Weight:               0

    Queue 6 (GBR-HIGH)
      Guaranteed priority:  high
      Excess priority:      low
      Reserved Bandwidth:   30 %
      Weight:               0

    Queue 7 (EXPRESS)
      Guaranteed priority:  strict-high
      Excess priority:      low
      Reserved Bandwidth:   10 %
      Weight:               0
```

Release Information

Juniper Networks introduced this command in Junos OS Release 22.2R1.

show services mobile-edge peers

IN THIS SECTION

- Syntax | 25743
- Description | 25743

- Options | 25743
- Required Privilege Level | 25743
- Output Fields | 25743
- Sample Output | 25745
- Release Information | 25747

Syntax

```
show services mobile-edge peers
<statistics>
```

Description

This command gives information on SAEGW-C and access network peers known to the SAEGW-U.

Options

none This command gives information on SAEGW-C and access network peers known to the SAEGW-U.

statistics This command gives protocol statistics on SAEGW-C and access network peers known to the SAEGW-U.

Required Privilege Level

view

Output Fields

Table 1444 on page 25744 describes the output fields for the `show services mobile-edge peers` command. Output fields are listed in the approximate order in which they appear.

Table 1444: show services mobile-edge peers Output Fields

Field Name	Field Description
Peers Summary	The `Peers Summary` field provides the following information: • Total number of control (SAEGW-C) peers. • Total number of access network peers. • Total rejected association setup requests.
Control Peer Information	The `Control Peer Information` field provides the following information about each SAEGW-C: • IP address • Hostname • Routing-Instance • Peer Group name
Access Peer Information	The `Access Peer Information` field provides the following information about each access network peer: • IP address • Hostname • Routing-Instance • Peer Group name

Table 1444: show services mobile-edge peers Output Fields *(Continued)*

Field Name	Field Description
Control Peer Statistics	The `Control Peer Statistics` field provides the following statistics about each SAEGW-C: • Heartbeat requests received and sent • Association setup requests received and sent • Association release requests received and sent • Session establishment requests received and sent (accepted and rejected) • Session modification requests received and sent (accepted and rejected) • Session deletion requests received and sent (accepted and rejected)
Access Peer Statistics	The `Access Peer Statistics` field provides the following statistics about each access network peer: • Echo requests received and sent

Sample Output

show services mobile-edge peers

```
user@host> show services mobile-edge peers
Peers Summary:
    Total control peers: 1
    Total access peers:  1
    Total association setup request rejects:  0

Control Peer information:
    IP address:          30.71.1.3
    Routing-Instance:    default

Access Peer information:
    IP address:          40.71.1.2
    Routing-Instance:    default
```

show services mobile-edge peers statistics

```
user@host> show services mobile-edge peers statistics
Peers Summary:
    Total control peers: 1
    Total access peers:  1
    Total association setup request rejects:  0

Control Peer Statistics:
    IP address:         13.1.0.4
    Hostname:           saegw-c1
    Routing-Instance:   default

    Heartbeat Requests Received:             11
    Heartbeat Responses Sent:                11

    Heartbeat Requests Sent:                 2
    Heartbeat Responses Received:            2

    Association Setup Requests Received:     1
    Association Setup Responses Sent:        1

    Association Release Requests Received:   0
    Association Release Responses Sent:      0

    Session Establishment Requests Received: 30000
    Session Establishment Responses Sent (Accepted):    30000
    Session Establishment Responses Sent (Rejected):    0

    Session Modification Requests Received:             30000
    Session Modification Responses Sent (Accepted):     30000
    Session Modification Responses Sent (Rejected):     0

    Session Deletion Requests Received:                 23169
    Session Deletion Responses Sent (Accepted):         22968
    Session Deletion Responses Sent (Rejected):         0
    Load Control Enabled                                Yes
    Load Control Information Sent                       10
    Overload Control Enabled                            No
    Overload Control Information Sent                   5

Access Peer Statistics:
```

```
    IP address:         12.1.0.4
    Routing-Instance:   default

    Echo Requests Received:                 0
    Echo Responses Sent:                    0
    Echo Requests Sent:                     0
    Echo Responses Received:                0
```

Release Information

Command introduced in Junos OS 19.4R1.

Fields for Load and Overload Control introduced in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Example: Configuring an MX Router as an SAEGW-U

show services mobile-edge resources call-admission-control statistics

IN THIS SECTION

- Syntax | 25747
- Description | 25748
- Options | 25748
- Required Privilege Level | 25748
- Output Fields | 25748
- Sample Output | 25749
- Release Information | 25749

Syntax

```
show services mobile-edge resources call-admission-control statistics
```

Description

This command displays information about the bandwidth currently allocated to the guaranteed bit rate queues on each anchor Packet Forwarding Engine. It also displays the Call Admission Control (CAC) statistics.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1445 on page 25748 lists the output fields for the `show services mobile-edge resources call-admission-control statistics` command.

Table 1445: show services mobile-edge resources call-admission-control statistics Output Fields

Field Name	Field Description
Allocated bandwidth	Amount of bandwidth currently allocated to GBR bearers and flows on the Packet Forwarding Engine.
Anchor	Name of the anchor Packet Forwarding Engine interface and its direction (downlink or uplink).
Calls admitted	Number of calls admitted to the anchor Packet Forwarding Engine during its uptime.
Calls rejected	Number of calls rejected by the anchor Packet Forwarding Engine due to no available bandwidth during its uptime.
High watermark	Maximum aggregate guaranteed bit rate bandwidth allocated on the anchor Packet Forwarding Engine through its uptime.
Oversubscription events	Number of times a new guaranteed bit rate bearer or flow was added when the aggregate bandwidth across all guranteed bit rate bearers and flows exceeded the reserved bandwidth.

Table 1445: show services mobile-edge resources call-admission-control statistics Output Fields *(Continued)*

Field Name	Field Description
Reserved bandwidth	Amount of bandwidth reserved for guaranteed bit rate bearers and flows on the Packet Forwarding Engine.

Sample Output

show services mobile-edge resources call-admission-control statistics

```
user@host> show services mobile-edge resources call-admission-control statistics
Global CAC Statistics
    Calls admitted:       153943454
    Calls rejected (no bandwidth):  6
    Oversubscription events:    80

Anchor: pfe-2/1/0, Downlink
       Reserved bandwidth:    5000000 Kbps
       Allocated bandwidth:    24235 Kbps
       High watermark:       5804400 Kbps
       Oversubscription events:    10
```

Release Information

Juniper Networks introduced this command in Junos OS Release 22.2R1.

show services mobile-edge sessions

IN THIS SECTION

- Syntax | 25750
- Description | 25750

- Options | 25750
- Required Privilege Level | 25751
- Output Fields | 25751
- Sample Output | 25754
- Release Information | 25771

Syntax

```
show services mobile-edge sessions
<control-plane-peers | detail | extensive | histogram | summary>
show services mobile-edge sessions control-plane-peers
<address>
show services mobile-edge sessions detail
<control-plane-peers | transient-sessions>
show services mobile-edge sessions extensive
<local-seid | transient-sessions>
show services mobile-edge sessions summary
<access-network-peers | anchor-group | apns | control-plane-peers | core-network-peers | pic |
slot | slot | transient-sessions>
```

Description

This command gives information on mobile sessions active on the SAEGW-U/SGW-U/PGW-U/UPF.

> **NOTE**: The backup routing engine (RE) supports only the top level of the `show services mobile-edge sessions summary` command and not any of its sub-options.

Options

none Display information on mobile sessions active on the SAEGW-U/SGW-U/PGW-U/UPF.

control-plane-peers Display information on mobile sessions for a specific or all SAEGW-Cs/SGW-C/PGW-C/SMF.

detail	Display detailed information on mobile sessions or on a specific SAEGW-C/SGW-C/PGW-C/SMF.
extensive	Display extensive information on mobile sessions or by local SEID.
histogram	Display total established sessions in five minute and hourly intervals.
summary	Display summary information on mobile sessions active on the SAEGW-U/SGW-U/PGW-U/UPF.
access-network-peers	Display session summary output by access network peer.
core-network-peers	Display session summary output by core network peer.
anchor-group	Display session summary output by anchor group.
apns	Display session summary output by APN.
local-seid	Display session extensive output for a local SEID.
pic	Display session summary output by PIC.
slot	Display session summary output by FPC slot.
statistics	Display detailed session statistics.
transient-sessions	Display detailed data for any transient sessions.

Required Privilege Level

view

Output Fields

Table 1446 on page 25751 describes the output fields for the `show services mobile-edge sessions` command. Output fields are listed in the approximate order in which they appear.

Table 1446: show services mobile-edge sessions Output Fields

Field Name	Field Description
Session-address	IP address of the session.

Table 1446: show services mobile-edge sessions Output Fields *(Continued)*

Field Name	Field Description
State	State of the session. ESTABLISHED or DELETING.
Num-bearers	Number of bearers for the session.
Num-QoS-Flows	Number of 5G QoS flows for the session.
VRF-ID	The ID number (in hexadecimal format) of the routing-instance that the mobile-edge session's APN is attached to.
APN	Access Point Name
CPF-peer	IP address of the control plane peer for the session.
Access-peer	IP address of the access peer (eNodeB) for the session.
Core-peer	IP address of the core network peer for the session.
Anchor-PFE	Anchor PFE for the session.
Secondary-anchor-PFE	Backup anchor PFE for the session.
Local-SEID	Local session ID (SEID) for the session.
Remote-SEID	Remote SEID for the session.
User ID	IE that can be present in a PFCP Session Establishment Request. This IE is useful for troubleshooting problems in the UPF affecting a subscriber.
Bearer EBI/NSAPI	EPS Bearer ID / Network Service Access Point Identifier.
Loc TEID	The local TEID allocated by the SAEGW-U for a given access peer.

Table 1446: show services mobile-edge sessions Output Fields *(Continued)*

Field Name	Field Description
Rem TEID	The remote TEID allocated for a given access peer. This is signaled over PFCP.
FAR-ID	Forwarding Action Rule ID.
Destination Interface	Destination Interface for a given FAR, either `Access` or `Core`.
PDR ID	Packet Detection Rule ID.
QER ID	QoS Enforcement Rule ID.
Uplink gate	Uplink QER gate status, either `Open` or `Closed`.
Downlink gate	Downlink QER gate status, either `Open` or `Closed`.
Uplink mbr	Uplink QER maximum bit rate before policing starts, in kbps.
Downlink mbr	Downlink QER maximum bit rate before policing starts, in kbps.
SDF Filter	Service data flow filter(s) for a given PDR. Flows matching these filters will be forwarded. All else are discarded.
Downlink Forwarding Class	Name of the downlink forwarding queue assigned to a given flow.
Sessions by State	The current number of sessions in state: `ESTABLISHED`, `DELETING`, and `Total`.
Bearers by State	The current number of bearers in state: `BEARER_WAIT`, `ESTABLISHED`, `DELETING`, and `Total`.

Table 1446: show services mobile-edge sessions Output Fields *(Continued)*

Field Name	Field Description
Bearers by Downlink FAR state	The current number of bearers by downlink FAR state. The allowable states are: • FORWARD– Downlink FAR apply action is set to Forward. • DROP – Downlink FAR apply action is set to Drop. • IDLE – Downlink FAR apply action is set to Notify Control Plane and/or Buffer, but the first packet has not been detected by the PFE. • BUFFER - Downlink FAR apply action is set to Notify Control Plane and/or Buffer, and the first packet has been detected by the PFE.
Bearers by Resource Type	The current number of bearers by the type of resource: guaranteed bit rate (`GBR`) or best effort (`Non-GBR`).
5QoS Flows by Resource Type	The current number of 5G QoS flows by the type of resource: guaranteed bit rate (`GBR`) or best effort (`Non-GBR`).
APFE PIC	The anchor PFE PIC.
Slot	The FPC slot.

Sample Output

show services mobile-edge sessions

```
user@host> show services mobile-edge sessions
    Session-address: 23.0.21.163 State: ESTABLISHED Num-bearers: 1
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.3.1.3 Access-peer: 10.4.1.6 Core-peer: 30.4.1.2
        Anchor-PFE: apfe0:pfe-3/0/0 Secondary-anchor-PFE: apfe0:pfe-2/0/0
        Local-SEID: 0x20015a2 Remote-SEID: 0x3cb2

    Session-address: 23.0.47.237 State: ESTABLISHED Num-bearers: 1
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.3.1.3 Access-peer: 10.4.1.6 Core-peer: 30.4.1.2
```

```
        Anchor-PFE: apfe0:pfe-3/0/0 Secondary-anchor-PFE: apfe0:pfe-2/0/0
        Local-SEID: 0x2fec Remote-SEID: 0x56fc

    Session-address: 23.0.21.49 State: ESTABLISHED Num-bearers: 1
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.3.1.3 Access-peer: 10.4.1.6 Core-peer: 30.4.1.2
        Anchor-PFE: apfe0:pfe-3/0/0 Secondary-anchor-PFE: apfe0:pfe-2/0/0
        Local-SEID: 0x1531 Remote-SEID: 0x3c40

    Session-address: 23.0.29.83 State: ESTABLISHED Num-bearers: 1
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.3.1.3 Access-peer: 10.4.1.6 Core-peer: 30.4.1.2
        Anchor-PFE: apfe0:pfe-3/0/0 Secondary-anchor-PFE: apfe0:pfe-2/0/0
        Local-SEID: 0x2001d53 Remote-SEID: 0x4462

....
```

show services mobile-edge sessions (with framed IPv6 Route)

```
user@host> show services mobile-edge sessions
    Session-address: 60.0.0.2 State: ESTABLISHED Num-bearers: 1
        IPv6-prefix: db43:2019::/64
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.3.1.3 Access-peer: 20.3.1.11   CPF-U-peer Count: 1
        Anchor-PFE: apfe0:pfe-3/0/0 Secondary-anchor-PFE: apfe0:pfe-2/1/0
        Local-SEID: 0x1 Remote-SEID: 0x2710
        Framed-IPv6-Route: 2000:0:0:106::/64 2000::106:a00:20ff:fe99:a998 1

....
```

show services mobile-edge sessions control-plane-peers

```
user@host> show services mobile-edge sessions control-plane-peers
Peer Address: 20.3.1.3

    Session-address: 23.0.21.163 State: ESTABLISHED Num-bearers: 1
        Num-QoS-Flows: 2
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.3.1.3 Access-peer: 10.4.1.6
        Anchor-PFE: apfe0:pfe-3/0/0 Secondary-anchor-PFE: apfe0:pfe-2/0/0
```

```
        Local-SEID: 0x20015a2 Remote-SEID: 0x3cb2

Peer Address: 20.3.1.3

    Session-address: 23.0.47.237 State: ESTABLISHED Num-bearers: 1
        Num-QoS-Flows: 2
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.3.1.3 Access-peer: 10.4.1.6
        Anchor-PFE: apfe0:pfe-3/0/0 Secondary-anchor-PFE: apfe0:pfe-2/0/0
        Local-SEID: 0x2fec Remote-SEID: 0x56fc

Peer Address: 20.3.1.3

    Session-address: 23.0.21.49 State: ESTABLISHED Num-bearers: 1
        Num-QoS-Flows: 2
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.3.1.3 Access-peer: 10.4.1.6
        Anchor-PFE: apfe0:pfe-3/0/0 Secondary-anchor-PFE: apfe0:pfe-2/0/0
        Local-SEID: 0x1531 Remote-SEID: 0x3c40

Peer Address: 20.3.1.3

    Session-address: 23.0.29.83 State: ESTABLISHED Num-bearers: 1
        Num-QoS-Flows: 2
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.3.1.3 Access-peer: 10.4.1.6
        Anchor-PFE: apfe0:pfe-3/0/0 Secondary-anchor-PFE: apfe0:pfe-2/0/0
        Local-SEID: 0x2001d53 Remote-SEID: 0x4462

...
```

show services mobile-edge sessions detail

```
user@host> show services mobile-edge sessions detail
    Session-address: 99.0.6.2 State: ESTABLISHED Time: Tue Mar 21 22:47:55 2023 UTC Num-bearers:
1
        Num-QoS-Flows: 2
        VRF-ID: 0x0 APN: default
        CPF-peer: 35.1.0.4 Routing-Instance: default
        Access-peer: 35.1.0.5 Routing-INstance: ri-data1
        Anchor-PFE: apfe1:pfe-0/1/0 Secondary-anchor-PFE: apfe1:pfe-3/1/0
```

```
        Local-SEID: 0x2000608 Remote-SEID: 0x2d12

        Bearer EBI/NSAPI: 5 State: ESTABLISHED
            Access: Address: 35.1.0.5 Local TEID: 0x504000 Remote TEID: 0x602

        QFI: 1
            Downlink Forwarding Class: GBR-LOW

        QFI: 2
            Downlink Forwarding Class: BE-HIGH-0

    Session-address: 99.0.4.219 State: ESTABLISHED Num-bearers: 1
        Num-QoS-Flows: 2
        VRF-ID: 0x0 APN: default
        CPF-peer: 35.1.0.4 Routing-Instance: default
        Access-peer: 35.1.0.5 Routing-Instance: data1
        Anchor-PFE: apfe1:pfe-0/1/0 Secondary-anchor-PFE: apfe1:pfe-3/1/0
        Local-SEID: 0x30004da Remote-SEID: 0x2beb

        Bearer EBI/NSAPI: 5 State: ESTABLISHED
            Access: Address: 35.1.0.5 Local TEID: 0x701860 Remote TEID: 0x4db

        QFI: 1
            Downlink Forwarding Class: GBR-LOW

        QFI: 2
            Downlink Forwarding Class: BE-HIGH-0

   ....
```

show services mobile-edge sessions detail (with framed IPv6 route)

```
user@host> show services mobile-edge sessions detail
    Session-address: 60.0.0.2 State: ESTABLISHED Time: Tue Mar 21 22:47:55 2023 UTC Num-bearers:
1
        IPv6-prefix: db43:2019::/64
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.3.1.3 Routing-Instance: default
        Access-peer: 20.3.1.11 Routing-Instance: data1
        CPF-U-peer Count: 1
        Anchor-PFE: apfe0:pfe-3/0/0 Secondary-anchor-PFE: apfe0:pfe-2/1/0
```

```
        Local-SEID: 0x1 Remote-SEID: 0x2710
        Framed-IPv6-Route: 2000:0:0:106::/64 2000::106:a00:20ff:fe99:a998 1
        Framed-IPv4-Route: 66.0.1.0/24 60.0.0.2 1

        Bearer EBI/NSAPI: 5 State: ESTABLISHED
            Access: Address: 20.3.1.11 Local TEID: 0x100000 Remote TEID: 0x1f4240
            CPF-U: Address: 20.3.1.4 Local TEID: 0x100700 Remote TEID: 0x7530

....
```

show services mobile-edge sessions detail control-plane-peers

```
user@host> show services mobile-edge sessions detail control-plane-peers address
Peer Address: 35.1.0.4
    Session-address: 99.0.6.2 State: ESTABLISHED Time: Tue Mar 21 22:47:55 2023 UTC Num-bearers:
1
        Num-QoS-Flows: 2
        VRF-ID: 0x0 APN: default
        CPF-peer: 35.1.0.4 Routing-Instance: default
        Access-peer: 35.1.0.5 Routing-INstance: ri-data1
        Anchor-PFE: apfe1:pfe-0/1/0 Secondary-anchor-PFE: apfe1:pfe-3/1/0
        Local-SEID: 0x2000608 Remote-SEID: 0x2d12

        Bearer EBI/NSAPI: 5 State: ESTABLISHED
            Access: Address: 35.1.0.5 Local TEID: 0x504000 Remote TEID: 0x602

        QFI: 1
            Downlink Forwarding Class: GBR-LOW

        QFI: 2
            Downlink Forwarding Class: BE-HIGH-0

    Session-address: 99.0.4.219 State: ESTABLISHED Num-bearers: 1
        Num-QoS-Flows: 2
        VRF-ID: 0x0 APN: default
        CPF-peer: 35.1.0.4 Access-peer: 35.1.0.5
        Anchor-PFE: apfe1:pfe-0/1/0 Secondary-anchor-PFE: apfe1:pfe-3/1/0
        Local-SEID: 0x30004da Remote-SEID: 0x2beb

        Bearer EBI/NSAPI: 5 State: ESTABLISHED
            Access: Address: 35.1.0.5 Local TEID: 0x701860 Remote TEID: 0x4db
```

```
        QFI: 1
            Downlink Forwarding Class: GBR-LOW

        QFI: 2
            Downlink Forwarding Class: BE-HIGH-0
....
```

show services mobile-edge sessions detail transient-sessions

```
user@host> show services mobile-edge sessions detail transient-sessions
    Session-address: 23.0.0.1 State: ESTABLISHED Time: Tue Mar 21 22:47:55 2023 UTC Num-bearers:
1
        Num-QoS-Flows: 2
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.7.1.3 Routing-Instance: default
        Access-peer: 10.8.1.6 Routing-Instance: ri-data1
        Anchor-PFE: apfe0:pfe-2/0/0 Secondary-anchor-PFE: apfe0:pfe-3/0/0
        Local-SEID: 0x1 Remote-SEID: 0x2710

        Bearer EBI/NSAPI: 5 State: ESTABLISHED
            Access: Address: 10.8.1.6 Local TEID: 0x100000 Remote TEID: 0x5c4b40

        QFI: 1
            Downlink Forwarding Class: GBR-LOW



   ....
```

show services mobile-edge sessions extensive

```
user@host> show services mobile-edge sessions extensive
     Session-address: 99.0.6.2 State: ESTABLISHED Time: Tue Mar 21 22:47:66 2023 UTC Num-
bearers: 1
        Num-QoS-Flows: 2
        VRF-ID: 0x0 APN: default
        CPF-peer: 198.18.26.1 Routing-Instance: default
        Access-peer: 198.18.25.2 Routing-Instance: ri-data1
        Anchor-PFE: apfe1:pfe-2/1/0 Secondary-anchor-PFE: apfe1:pfe-3/1/0
        Local-SEID: 0x2000608 Remote-SEID: 0x2d12
```

```
Bearer EBI/NSAPI: 5 State: ESTABLISHED
    Access: Address: 198.18.25.2 Local TEID: 0x504000 Remote TEID: 0x602

    FAR-ID: 1
    Destination Interface: Core
    Transport Level Marking: No marking
    Apply Action: FORWARD

    FAR-ID: 2
    Destination Interface: Access
    Transport Level Marking: No marking
    Apply Action: FORWARD
    Outer Header Creation: GTP-U/UDP/IPv4
    Remote Address: 35.1.0.5 Remote TEID: 0x602
    Downlink forwarding class: BE-LOW-1
    Downlink egress forwarding class: best-effort

    FAR-ID: 3
    Destination Interface: Access
    Transport Level Marking: 0x20
    Apply Action: Notify CP, Buffer

    PDR ID: 4
      Source Interface: Core
      Precedence: 250
      FAR ID: 3
      TFE ID: 2
      QER ID: 2
        Uplink gate : Open      Downlink gate   : Open
        Uplink mbr  : 20000 kbps    Downlink mbr    : 20000 kbps
        QFI : 2
      SDF Filter:
        permit out 17 from any 4002 to assigned 4003
      Downlink Forwarding Class: BE-HIGH-0

    PDR ID: 2
      Source Interface: Access
        Local TEID: 0x504000
      Precedence: 252
      FAR ID: 1
      URR ID: 2
      QER ID: 2
```

```
    Uplink gate : Open      Downlink gate   : Open
    Uplink mbr  : 20000 kbps    Downlink mbr    : 20000 kbps
    QFI : 2
  SDF Filter:
    permit out 17 from any 4002 to assigned 4003
  QFI: 2

PDR ID: 3
  Source Interface: Core
  Precedence: 255
  FAR ID: 2
  URR ID: 1
  QER ID: 1
    Uplink gate : Open      Downlink gate   : Open
    Uplink mbr  : 10000 kbps    Downlink mbr    : 10000 kbps
    Uplink gbr  : 500 kbps    Downlink gbr    : 500 kbps
    QFI : 1
  SDF Filter:
    permit out 17 from any 3002 to assigned 3003
  Downlink Forwarding Class: GBR-LOW

PDR ID: 1
  Source Interface: Access
    Local TEID: 0x504000
  Precedence: 255
  FAR ID: 1
  URR ID: 1
  QER ID: 1
    Uplink gate : Open      Downlink gate   : Open
    Uplink mbr  : 10000 kbps    Downlink mbr    : 10000 kbps
    Uplink gbr  : 500 kbps    Downlink gbr    : 500 kbps
    QFI : 1
  SDF Filter:
    permit out 17 from any 3002 to assigned 3003
  QFI: 1

URR ID: 0x2
    Measurement Method:            Volume
    Reporting Trigger:             Volume Threshold
    Total volume threshold:        1000000 octets

    Usage report trigger:          Volume Threshold
    Usage report start timestamp:  Tue May 24 21:37:52 2022 UTC
```

```
                Usage report stop timestamp:   Tue May 24 21:39:41 2022 UTC
                Usage report uplink volume:    999948 octets
                Usage report downlink volume:  0 octets
                Usage report total volume:     999948 octets


            URR ID: 0x1
                Measurement Method:            Volume
                Reporting Trigger:             Volume Threshold
                Total volume threshold:        500000 octets

                Usage report trigger:          Volume Threshold
                Usage report start timestamp:  Tue May 24 21:40:44 2022 UTC
                Usage report stop timestamp:   Tue May 24 21:40:58 2022 UTC
                Usage report uplink volume:    250056 octets
                Usage report downlink volume:  249964 octets
                Usage report total volume:     500020 octets


        URR statistics:
            URR messages:
            URR ID          Usage Report Count   Last Report Timestamp
            0x2             160                  Tue May 24 21:39:41 2022

            0x1             1284                 Tue May 24 21:40:58 2022

        Session Reports:
            Type              Report Count         Last Report Timestamp
            Usage             1444                 Tue May 24 21:40:58 2022
            Total             1444
```

show services mobile-edge sessions extensive (with framed IPv6 route)

```
user@host> show services mobile-edge sessions extensive
     Session-address: 60.0.0.2 State: ESTABLISHED Num-bearers: 1
        IPv6 Prefix: db43:2019::/64
        VRF-ID: 0x0 APN: default
        CPF-peer: 198.18.26.1 Routing-Instance: default
        Access-peer: 198.18.25.2 Routing-Instance: ri-data1
        CPF-U-peer Count: 1
        Anchor-PFE: apfe0:pfe-3/0/0 Secondary-anchor-PFE: apfe1:pfe-2/1/0
```

```
        Local-SEID: 0x1 Remote-SEID: 0x2710
        Framed-IPv6-Route: 2000:0:0:106::/64 2000::106:a00:20ff:fe99:a998 1
        Framed-IPv4-Route: 66.0.1.0/24 60.0.0.2 1

        Bearer EBI/NSAPI: 5 State: ESTABLISHED
            Access: Address: 20.3.1.11 Local TEID: 0x100000 Remote TEID: 0x1f4240
            CPF-U: Address: 20.3.1.4 Local TEID: 0x100700 Remote TEID: 0x7530

            FAR-ID: 1
            Destination Interface: Core
            Transport Level Marking: No marking
            Apply Action: FORWARD

            FAR-ID: 2
            Destination Interface: Access
            Transport Level Marking: No marking
            Apply Action: FORWARD
            Outer Header Creation: GTP-U/UDP/IPv4

            FAR-ID: 3
            Destination Interface: CP-Func
            Transport Level Marking: No Marking
            Apply Action: FORWARD
            Outer Header Creation: GTP-U/UDP/IPv4

            PDR ID: 4
              Source Interface: Access
                Local TEID: 0x100000
              FAR ID: 3
              QER ID: 1
                Uplink gate : Open      Downlink gate   : Open
                Uplink mbr  : 5000 kbps    Downlink mbr    : 5000 kbps
                Precedence: 1
              SDF Filter:
                permit out 58 from ff02::2 to assigned

            PDR ID: 3
              Source Interface: CP-FUNC
                Local TEID: 0x100700
              FAR ID: 2
              QER ID: 1
                Uplink gate : Open    Downlink gate    : Open
                Uplink mbr : 5000 kbps    Downlink mbr    : 5000 kbps
```

```
            Precedence: 1000
          SDF Filter:
            permit out 58 from any to assigned

        PDR ID: 2
          Source Interface: Core
          FAR ID: 2
          TFE ID: 2
          URR ID: 1
          QER ID: 1
            Uplink gate : Open      Downlink gate   : Open
            Uplink mbr  : 5000 kbps    Downlink mbr    : 5000 kbps
            Precedence: 255
          SDF Filter:
            permit out ip from any to assigned

        PDR ID: 1
          Source Interface: Access
            Local TEID: 0x100000
          FAR ID: 1
          TFE ID: 1
          URR ID: 1
          QER ID: 1
            Uplink gate : Open      Downlink gate   : Open
            Uplink mbr  : 5000 kbps    Downlink mbr    : 5000 kbps
            Precedence: 255
          SDF Filter:
            permit out ip from any to assigned

        URR ID: 0x1
          Measurement Method:            Volume
            Reporting Trigger:        Volume Threshold
            Total volume threshold:        10000 octets

    Session Reports:
        Type              Report Count         Last Report Timestamp
        Total             0
```

show services mobile-edge sessions extensive transient-sessions

```
user@host> show services mobile-edge sessions extensive transient-sessions
     Session-address: 23.0.0.1 State: ESTABLISHED Time: Tue Mar 21 22:47:66 2023 UTC Num-
bearers: 1
        Num-QoS-Flows: 2
        VRF-ID: 0x0 APN: default
        CPF-peer: 20.7.1.3 Routing-Instance: default
        Access-peer: 10.8.1.6 Routing-Instance: ri-data1
        Anchor-PFE: apfe1:pfe-2/1/0 Secondary-anchor-PFE: apfe1:pfe-3/1/0
        Local-SEID: 0x1 Remote-SEID: 0x2710

        Bearer EBI/NSAPI: 5 State: ESTABLISHED
            Access: Address: 10.8.1.6 Local TEID: 0x100000 Remote TEID: 0x5c4b40

            FAR-ID: 1
            Destination Interface: Core
            Transport Level Marking: No marking
            Apply Action: FORWARD

            FAR-ID: 2
            Destination Interface: Access
            Transport Level Marking: No marking
            Apply Action: FORWARD
            Outer Header Creation: GTP-U/UDP/IPv4
            Remote Address: 10.8.1.6 Remote TEID: 0x5c4b40
            Downlink forwarding class: BE-LOW-0
            Downlink egress forwarding class: best-effort

            PDR ID: 1
              Source Interface: Access
                Local TEID: 0x504000
              Precedence: 4294967295
              FAR ID: 2
              URR ID: 1
              QER ID: 1
                Uplink gate : Open      Downlink gate   : Open
                Uplink mbr  : 10000000 kbps    Downlink mbr    : 10000000 kbps
                QFI : 2
              SDF Filter:
                permit out ip from any to assigned
              QFI: 2
```

```
    PDR ID: 2
      Source Interface: Core
        Local TEID: 0x504000
      Precedence: 4294967295
      FAR ID: 2
      URR ID: 1
      QER ID: 1
        Uplink gate : Open      Downlink gate   : Open
        Uplink mbr  : 10000000 kbps    Downlink mbr    : 10000000 kbps
        QFI : 2
      SDF Filter:
        permit out ip from any to assigned
      QFI: 2

    PDR ID: 1
      URR ID: 1
        Measurement Method:            Volume
        Reporting Trigger:             Volume Threshold
        Total volume threshold:        8000000 octets

        Usage report trigger:          Volume Threshold
        Usage report start timestamp:  Tue May 24 21:37:52 2022 UTC
        Usage report stop timestamp:   Tue May 24 21:39:41 2022 UTC
        Usage report uplink volume:    999948 octets
        Usage report downlink volume:  0 octets
        Usage report total volume:     999948 octets

URR statistics:
    URR messages:
    URR ID          Usage Report Count   Last Report Timestamp
    0x2             160                  Tue May 24 21:39:41 2022

    0x1             1284                 Tue May 24 21:40:58 2022

Session Reports:
    Type              Report Count         Last Report Timestamp
    Usage             1444                 Tue May 24 21:40:58 2022
    Total             1444
```

show services mobile-edge sessions histogram

```
user@host> show services mobile-edge sessions histogram
Established sessions 5-Minutely(UTC):
   18:25    200500   200550   20058    200600   200510   200525
   17:55    320050   320055   42005    420050   520050   520055

Established sessions Hourly(UTC):
   16:31    200500   320055   350055   340055   330055   320056
   10:31    331055   331057   331058   331059   331050   331051
   04:31    350052   350053   350054   350055   350056   350057
   18:31    350058   350059   350050   320051   310022   300023
```

show services mobile-edge sessions summary

```
user@host> show services mobile-edge sessions summary
Sessions by State:
    ESTABLISHED: 2000
    Total: 2000

Bearers by State:
    ESTABLISHED: 2000
    Total: 2000

Bearers by Downlink FAR State:
    FORWARD: 2000
    Total: 2000

Bearers by Resource Type:
    Non-GBR: 2000
    Total:   2000

5G QoS Flows by State:
    ESTABLISHED: 4000
    Total: 4000

5G QoS Flows by Resource Type:
    GBR:     2000
    Non-GBR: 2000
    Total:   4000
```

show services mobile-edge sessions summary access-network-peers

```
user@host> show services mobile-edge sessions summary access-network-peers
Summary by Access Peer:
Peer Address: 40.71.1.2
    Sessions by State:
    ESTABLISHED: 6426
    Total: 6426

    Bearers by State:
    ESTABLISHED: 6426
    DELETING: 4832
    Total: 11258
```

show services mobile-edge sessions summary anchor-group

```
user@host> show services mobile-edge sessions summary anchor-group
Summary by pic:
APFE PIC: apfe0:pfe-4/0
    Sessions by State:
    ESTABLISHED: 6421
    DELETING: 4832
    Total: 11253

    Bearers by State:
    BEARER_WAIT: 58696
    ESTABLISHED: 6421
    DELETING: 4927
    Total: 70044

Summary by pic:
APFE PIC: apfe0:pfe-1/0
    Sessions by State:
    ESTABLISHED: 6421
    DELETING: 4832
    Total: 11253

    Bearers by State:
    BEARER_WAIT: 58696
    ESTABLISHED: 6421
```

```
    DELETING: 4927
    Total: 70044
```

show services mobile-edge sessions summary apns

```
user@host> show services mobile-edge sessions summary apns
Summary by Access Point Names:
APN: default
    Sessions by State:
    ESTABLISHED: 6421
    DELETING: 4832
    Total: 11253

    Bearers by State:
    ESTABLISHED: 6421
    DELETING: 4832
    Total: 11253

Summary by Access Point Names:
APN: test.internet.488
    Sessions by State:
    Total: 0

    Bearers by State:
    Total: 0
```

show services mobile-edge sessions summary control-plane-peers

```
user@host> show services mobile-edge sessions summary control-plane-peers
Summary by Control Peer:
Peer Address: 30.71.1.3
    Sessions by State:
    ESTABLISHED: 6421
    DELETING: 4832
    Total: 11253

    Bearers by State:
    BEARER_WAIT: 58696
    ESTABLISHED: 6421
```

```
    DELETING: 4927
    Total: 70044
```

show services mobile-edge sessions summary pic

```
user@host> show services mobile-edge sessions summary pic
Summary by pic:
APFE PIC: pfe-1/0
    Sessions by State:
    ESTABLISHED: 6421
    DELETING: 4832
    Total: 11253

    Bearers by State:
    BEARER_WAIT: 58696
    ESTABLISHED: 6421
    DELETING: 4927
    Total: 70044

Summary by pic:
APFE PIC: pfe-4/0
    Sessions by State:
    ESTABLISHED: 6421
    DELETING: 4832
    Total: 11253

    Bearers by State:
    BEARER_WAIT: 58696
    ESTABLISHED: 6421
    DELETING: 4927
    Total: 70044
```

show services mobile-edge sessions summary slot

```
user@host> show services mobile-edge sessions summary slot
Summary by slot:
Slot: pfe-1
    Sessions by State:
    ESTABLISHED: 6421
    DELETING: 4832
```

```
    Total: 11253

    Bearers by State:
    BEARER_WAIT: 58696
    ESTABLISHED: 6421
    DELETING: 4927
    Total: 70044

Summary by slot:
Slot: pfe-4
    Sessions by State:
    ESTABLISHED: 6421
    DELETING: 4832
    Total: 11253

    Bearers by State:
    BEARER_WAIT: 58696
    ESTABLISHED: 6421
    DELETING: 4927
    Total: 70044
```

show services mobile-edge session summary transient-sessions

```
user@host> show services mobile-edge sessions summary transient-sessions
Total sessions: 3
```

Release Information

Command introduced in Junos OS 19.4R1.

core-network-peers option introduced in Junos OS 20.4R1.

Framed IPv6 Route options introduced in Junos OS 22.4R1.

transient-sessions and *histogram*options introduced in Junos OS 23.4R1.

RELATED DOCUMENTATION

No Link Title

show services mobile-edge summary

IN THIS SECTION

- Syntax | 25772
- Description | 25772
- Required Privilege Level | 25772
- Output Fields | 25772
- Sample Output | 25774
- Release Information | 25774

Syntax

```
show services mobile-edge summary
```

Description

This command gives summary information on the SAEGW-U.

Required Privilege Level

view

Output Fields

Table 1447 on page 25772 describes the output fields for the `show services mobile-edge summary` command. Output fields are listed in the approximate order in which they appear.

Table 1447: show services mobile-edge summary Output Fields

Field Name	Field Description
Graceful-Restart	Whether graceful-restart is enabled or disabled.

Table 1447: show services mobile-edge summary Output Fields *(Continued)*

Field Name	Field Description
Primary Role	Primary Role state for the given routing engine (RE) this command is executed on. For the primary RE, `Master` is displayed. For the backup RE, `Standby` is displayed. If GRES not enabled, `Standalone` is displayed.
State	State can be either `Running` or `Bulk Sync`. `Bulk Sync` implies we are waiting for synchronization.
Bulk Sync	Status of `Bulk Sync`. Must be `Synchronized` with GRES enabled to preserve state through daemon restart. Other states are `Synchronizing`or `N/A` (when GRES is not enabled).
CPU Load	Provides the percentage of the CPU in use.
Memory Load	The amount of memory in use.
Session Load	The percentage of the maximum session capacity in use. The date and time of the highest ever session load is also provided.
GBR CAC Load	The percentage of the maximum system bandwidth in use on the system. The date and time of the highest ever GBR CAC Load is also provided.
Inflight Load	The percentage of the maximum active registration procedures on the system. The system will be in overload status if this number exceeds 80%. The date and time of the highest ever Inflight Load is also provided.
PFE queue Load	The percentage of traffic on the PFE that is on the Waiting-to-Send (WTS) queue. This can be used to monitor system congestion.The date and time of the highest ever PFE Queue Load is also provided.
Aggregate Load	The total load of the system in use, as compared to the maximum load use case. The date and time of the highest ever Aggregate Load is also provided.
Aggregate Overload	The total overload percentage of the system. The date and time of the highest ever Aggregate Overload is also given.

Table 1447: show services mobile-edge summary Output Fields *(Continued)*

Field Name	Field Description
Processed Session Messages/Sec	The number of session messages per second that the system has processed. The date and time of the highest ever Processed Session Messages/Sec is also provided.
Received Session Messages/Sec	The number of system messages per second that the system has received. The date and time of the highest ever Received Session Messages/Sec is also given.

Sample Output

show services mobile-edge summary

```
user@host> show services mobile-edge summary
Graceful-Restart        Enabled
Mastership              Master
State                   Running
Bulk Sync               Synchronized
CPU Load                        40%
Memory Load                     10%
Session Load                    10% (top 26% at Jan.24 12:34)
GBR CAC Load                    15% (top 20% at Jan.24 12:34)
Inflight Load                   5% (top 6% at Jan.24 12:31)
PFE queue Load                  0% (top 2% at Jan.23 21:08)
Aggregate Load                  7% (top 7% at Jan.24 12:53)
Aggregate Overload              7% (top 7% at Jan.25 12:53)
Processed Session Messages/Sec  0 (top 1117 at Jan.24 12:44)
Received Session Messages/Sec   0 (top 1631 at Jan.24 12:44)
```

Release Information

Command introduced in Junos OS 19.4R1.

Additional fields for Load and Overload Information added in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Example: Configuring an MX Router as an SAEGW-U

show services monitoring rfc2544

IN THIS SECTION

- Syntax | 25775
- Description | 25775
- Options | 25775
- Additional Information | 25776
- Required Privilege Level | 25776
- Output Fields | 25776
- Sample Output | 25777
- Release Information | 25779

Syntax

```
show services monitoring rfc2544
(active-tests <extensive>|  terminated-tests <extensive> | summary)
```

Description

Display information about the results of each category or state of the RFC 2544-based benchmarking tests, such as terminated tests and active tests. You can also display summary statistics about the total number of tests of each state for a high-level, quick analysis. The values in the output displayed vary, depending on the state in which the test is passing through, when you issue the command.

Options

active-tests <extensive> Display the results of the set of tests that are currently running.

terminated-tests <extensive> Display the list of tests that were terminated or stopped. This list includes tests that failed due to various error conditions and tests that you terminated by entering the `test services monitoring rfc2544 test` *`test-name`* `stop` command. The `Status` field in the output specifies the reason for the termination of the test.

summary Display summary output.

Additional Information

Required Privilege Level

view

Output Fields

Table 1448 on page 25776 lists the output fields for the `show services monitoring rfc2544` command.

Table 1448: show services monitoring rfc2544 Output Fields

Field Name	Field Description
`Counters last cleared`	Date, time, and how long ago the statistics for the test were cleared. The format is *year-month-day hour:minute:second timezone* . For example, 2021-02-11 07:51:28 PDT. If you did not clear the statistics previously at any point, a dash (-) is displayed.
`Status`	Indicates whether the test is currently in progress or has been terminated. This field is displayed for tests that are in progress or were terminated by entering the `test services monitoring rfc2544 test <`*`test-name`* `/` *`test-id`*`> stop` command.
`Test family`	The family type configured for the test; for example, INET.
`Test id`	Unique identifier configured for the test.
`Test mode`	Mode configured for the test on the router. Test mode is: • Reflect: Test frames that originate from one end are reflected at the other end on the selected service, such as IPv4.

Table 1448: show services monitoring rfc2544 Output Fields *(Continued)*

Field Name	Field Description
`Test name`	Name configured for the test.
`Test start time` and `Test finish time`	Time at which the test started and finished in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). A dash (-) is displayed if the test is still running.
`Test state`	State of the test that is in progress or active when the output is displayed.

Sample Output

show services monitoring rfc2544 summary

```
user@host> show services monitoring rfc2544 summary
```

```
  Tests summary :
      Number of active tests: 1, Number of terminated tests: 2
```

This output indicates that 1 test iteration is currently in progress (at the time of issue of the command) and 2 tests were halted.

show services monitoring rfc2544 active-tests

```
user@host> show services monitoring rfc2544 active-tests
```

```
Active tests:

Test id: 3, Test name: mytest
    Test mode: Reflect
    Test family: INET
    Test state: Active
    Status: Running
    Test start time: 2021-02-11 04:39:10 PST
```

```
    Test finish time: -
    Counters last cleared: 2021-02-11 04:39:12 PST
```

show services monitoring rfc2544 active-tests extensive

```
user@host> show services monitoring rfc2544 active-tests extensive
```

```
Active tests:

Test id: 3, Test name: mytest
    Test mode: Reflect
    Test family: INET
    Test state: Active
    Status: Running
    Test start time: 2021-02-11 04:39:10 PST
    Test finish time: -
    Counters last cleared: 2021-02-11 04:39:12 PST

    INET family Configuration:
        Destination IPv4 address: 192.0.2.2
        Destination UDP port: 7890
        Source IPv4 address: 192.0.2.1

Elapsed        Reflected            Reflected
time           Packets               Bytes
   271                0                     0
```

show services monitoring rfc2544 terminated-tests

```
user@host> show services monitoring rfc2544 terminated-tests
```

```
Terminated tests:

Test id: 1, Test name: mytest
    Test mode: Reflect
    Test family: INET
    Test state: Terminated
```

```
    Status: Stopped from CLI
    Test start time: 2021-02-11 03:14:19 PST
    Test finish time: 2021-02-11 03:17:26 PST
    Counters last cleared: -
```

show services monitoring rfc2544 terminated-tests extensive

```
user@host> show services monitoring rfc2544 terminated-tests extensive
```

```
Terminated tests:

Test id: 1, Test name: mytest
    Test mode: Reflect
    Test family: INET
    Test state: Terminated
    Status: Stopped from CLI
    Test start time: 2021-02-11 03:14:19 PST
    Test finish time: 2021-02-11 03:17:26 PST
    Counters last cleared: -

    INET family Configuration:
        Destination IPv4 address: 192.0.2.2
        Destination UDP port: 7890
        Source IPv4 address: 192.0.2.1

Elapsed        Reflected            Reflected
time           Packets               Bytes
   187                0                     0
```

Release Information

Command introduced in Junos OS Evolved Release 21.1R1.

show services monitoring rfc2544

IN THIS SECTION

- Syntax | 25780
- Description | 25780
- Options | 25780
- Additional Information | 25781
- Required Privilege Level | 25781
- Output Fields | 25781
- Sample Output | 25782
- Release Information | 25786

Syntax

```
show services monitoring rfc2544
<active-tests>
<completed-tests>
<detail>
<extensive>
<summary>
<terminated-tests>
<test-id test-id>
```

Description

Display information about the results of the RFC 2544-based benchmarking test for a specific test ID or for a category of tests. The values in the output displayed vary, depending on the state in which the test is passing through, when you issue the command.

Options

active-tests (Optional) Display results for all active RFC 2544 benchmarking tests.

completed-tests (Optional) Display results for all completed RFC 2544 benchmarking tests.

detail (Optional) When specified with the `test-id` option, display detailed results. When specified on its own, display detailed results of all RFC 2544 benchmarking tests (active, completed, and terminated).

extensive (Optional) When specified with the `test-id` option, display extensively detailed results. When specified on its own, displays extensive details of all RFC 2544 benchmarking tests (active, completed, and terminated).

summary (Optional) When specified with the `test-id` option, display summarized results. When specified on its own, display a summary of all RFC 2544 benchmarking tests (active, completed, and terminated).

terminated-tests (Optional) Display results for all terminated RFC 2544 benchmarking tests.

`test-id` *`test-id`* Display test results for the specified unique identifier.

Additional Information

Required Privilege Level

view

Output Fields

Table 1448 on page 25776 lists the output fields for the `show services monitoring rfc2544 test-id` command.

Table 1449: show services monitoring rfc2544 test-id Output Fields

Field Name	Field Description
`Counters last cleared`	Date, time, and how long ago the statistics for the test were cleared. The format is *year-month-day hour:minute:second timezone* . For example, 2021-02-11 07:51:28 PDT. If you did not clear the statistics previously at any point, a dash (-) is displayed.
`Status`	Indicates whether the test is currently in progress or has been terminated. This field is displayed for tests that are in progress or were terminated by entering the `test services monitoring rfc2544 test <test-name / test-id> stop` command.

Table 1449: show services monitoring rfc2544 test-id Output Fields *(Continued)*

Field Name	Field Description
`Test family`	The family type configured for the test; for example, INET.
`Test id`	Unique identifier configured for the test.
`Test mode`	Mode configured for the test on the router. Test mode is: • Initiate-and-Terminate: Test frames are initiated from one end and terminated at the same end. This mode requires a reflector to be configured at the peer end to enable the test frames to be returned to the source. This mode is supported only on ACX Series routers. • Reflect: Test frames that originate from one end are reflected at the other end on the selected service, such as IPv4.
`Test name`	Name configured for the test.
`Test start time` and `Test finish time`	Time at which the test started and finished in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). A dash (-) is displayed if the test is still running.
`Test state`	State of the test that is in progress or active when the output is displayed.

Sample Output

show services monitoring rfc2544 test-id *test-id*

```
user@host> show services monitoring rfc2544 test-id 2
```

```
Test id: 2, Test name: mytest
    Test mode: Reflect
    Test family: INET
    Test state: Terminated
    Status: Stopped from CLI
    Test start time: 2021-02-11 03:38:25 PST
```

```
    Test finish time: 2021-02-11 03:38:43 PST
    Counters last cleared: -
```

show services monitoring rfc2544 test-id *test-id* detail

```
user@host> show services monitoring rfc2544 test-id 1 detail
Test id: 1, Test name: t1
    Test mode: reflect
    Test family: ethernet-switching
    Test state: Active
    Status: Running
    Test start time: 2022-05-30 13:23:21 PDT
    Test finish time: -
    Counters last cleared: -

    Bridge family Configuration:
        Interface: et-0/0/3.0
        Direction: egress
        Service type: elan
        Destination MAC address: 00:00:00:08:00:08
        Source MAC address: 00:00:00:07:00:07

Elapsed        Reflected            Reflected
 time           Packets               Bytes
  1016       5196751848         2868607020096
```

show services monitoring rfc2544 test-id *test-id* extensive

```
user@host> show services monitoring rfc2544 test-id 2 extensive
```

```
Test id: 2, Test name: mytest
    Test mode: Reflect
    Test family: INET
    Test state: Terminated
    Status: Stopped from CLI
    Test start time: 2021-02-11 03:38:25 PST
    Test finish time: 2021-02-11 03:38:43 PST
    Counters last cleared: -
```

```
    INET family Configuration:
        Destination IPv4 address: 192.0.2.4
        Destination UDP port: 7890
        Source IPv4 address: 192.0.2.3

Elapsed        Reflected            Reflected
time           Packets               Bytes
    18                0                     0
```

show services monitoring rfc2544 summary

```
user@host> show services monitoring rfc2544 summary
[Number of active tests: active-test-count],
[Number of completed tests: completed-test-count],
[Number of terminated tests: terminated-test-count]
```

show services monitoring rfc2544 active-tests

```
user@host> show services monitoring rfc2544 active-tests
```

```
Active tests:

Test id: 3, Test name: mytest
    Test mode: Reflect
    Test family: INET
    Test state: Active
    Status: Running
    Test start time: 2021-02-11 04:39:10 PST
    Test finish time: -
    Counters last cleared: 2021-02-11 04:39:12 PST
```

show services monitoring rfc2544 active-tests extensive

```
user@host> show services monitoring rfc2544 active-tests extensive
```

```
Active tests:

Test id: 3, Test name: mytest
    Test mode: Reflect
    Test family: INET
    Test state: Active
    Status: Running
    Test start time: 2021-02-11 04:39:10 PST
    Test finish time: -
    Counters last cleared: 2021-02-11 04:39:12 PST

    INET family Configuration:
        Destination IPv4 address: 192.0.2.2
        Destination UDP port: 7890
        Source IPv4 address: 192.0.2.1

Elapsed        Reflected            Reflected
time           Packets               Bytes
   271                0                     0
```

show services monitoring rfc2544 terminated-tests

```
user@host> show services monitoring rfc2544 terminated-tests
```

```
Terminated tests:

Test id: 1, Test name: mytest
    Test mode: Reflect
    Test family: INET
    Test state: Terminated
    Status: Stopped from CLI
    Test start time: 2021-02-11 03:14:19 PST
```

```
    Test finish time: 2021-02-11 03:17:26 PST
    Counters last cleared: -
```

show services monitoring rfc2544 terminated-tests extensive

```
user@host> show services monitoring rfc2544 terminated-tests extensive
```

```
Terminated tests:

Test id: 1, Test name: mytest
    Test mode: Reflect
    Test family: INET
    Test state: Terminated
    Status: Stopped from CLI
    Test start time: 2021-02-11 03:14:19 PST
    Test finish time: 2021-02-11 03:17:26 PST
    Counters last cleared: -

    INET family Configuration:
        Destination IPv4 address: 192.0.2.2
        Destination UDP port: 7890
        Source IPv4 address: 192.0.2.1

Elapsed        Reflected            Reflected
time           Packets               Bytes
   187               0                     0
```

Release Information

Command introduced in Junos OS Evolved Release 21.1R1.

`active-tests`, `completed-tests`, `detail`, `summary`, and `terminated-tests` options introduced in Junos OS Evolved Release 22.4R1.

show services monitoring rpm history-results

IN THIS SECTION

- Syntax | 25787
- Description | 25787
- Options | 25787
- Required Privilege Level | 25788
- Output Fields | 25788
- Sample Output | 25790
- Release Information | 25794

Syntax

```
show services monitoring rpm history-results
owner name
<brief | detail>
<since time>
<source-address address>
<target address>
<test name>
```

Description

Display the results stored for the specified real-time performance monitoring (RPM) probes.

Options

owner *name*

(Required) Display information for probes with the specified probe owner.

brief | detail

(Optional) Display the specified level of output.

- **Default:** brief

since *time* (Optional) Display information from the specified time. Specify time as *yyyy-mm-dd.hh:mm:ss*.

source-address *address* (Optional) Display information only for probes with the specified source address.

target *address* (Optional) Display information only for probes with the specified target address.

test *name* (Optional) Display information only for the specified test.

Do not configure `test` if you configure any of the following options: `source-address`or `target`. These options do not work when you configure `test`.

Required Privilege Level

view

Output Fields

Table 1450 on page 25788 lists the output fields for the `show services monitoring rpm history-results` command. Output fields are listed in the approximate order in which they appear.

Table 1450: show services monitoring rpm history-results Output Fields

Field Name	Field Description	Level of Output
`Owner`	Owner name. When you configure the probe owner statement at the `[edit services monitoring rpm]` hierarchy level, this field displays the configured owner name.	All levels
`Test`	Name of a test representing a collection of probes. When you configure the `test` *name* statement at the `[edit services monitoring rpm owner]` hierarchy level, the field displays the configured test name.	All levels
`Probe sent`	Timestamp when the probe was sent.	`owner, brief`
`Probe received`	Timestamp when the probe result was determined.	`owner, brief`
`Round trip time`	Average ping round-trip time (RTT), in microseconds.	`owner, brief`

Table 1450: show services monitoring rpm history-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Probe-type`	Configured probe type for the test.	`detail`
`Probe results`	Result of a particular probe performed by a remote host, including whether a probe response was received. The following information is contained in the results: • `Probe sent time`—Timestamp for when the probe was sent. • `Probe rcvd time`—Timestamp for when the probe was sent, and whether offload timestamping is configured. • `Rtt`—Average ping round-trip time (RTT), in microseconds. • `Rtt jitter`—Difference, in microseconds, between the maximum and minimum RTT. • `Egress jitter`—For the probe type icmp-ping-timestamp, the egress delay measured for this probe, in microseconds. • `Ingress jitter`—For the probe type icmp-ping-timestamp, the ingress delay measured for this probe, in microseconds.	`detail`
`Results over current test`	Displays the results for the current test by probe at the time each probe was completed.	`detail`
`Probes sent`	Number of probes sent with the current test.	`detail`
`Probes received`	Number of probe responses received within the current test.	`detail`
`Loss percentage`	Percentage of lost probes for the current test.	`detail`

Table 1450: show services monitoring rpm history-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Measurement	Measurement for round trip time: • Samples—Number of samples. • Minimum—Minimum RTT measured over the course of the current test. • Maximum—Maximum RTT measured over the course of the current test. • Average—Average RTT measured over the course of the current test. • Stddev—Standard deviation of the round-trip time, in microseconds, measured over the course of the current test.	detail

Sample Output

show services monitoring rpm history-results owner

```
user@host> show services monitoring rpm history-results owner icmp-evo

Owner, Test                  Probe sent                 Probe received             Round trip
time
icmp-evo, icmp-evo-2         11/18/20 12:08:28.621987   11/18/20 12:08:28.648443   26444 usec
icmp-evo, icmp-evo-4         11/18/20 12:08:28.642527   11/18/20 12:08:28.660700   18171 usec
icmp-evo, icmp-evo-1         11/18/20 12:08:28.642533   11/18/20 12:08:28.663284   20749 usec
icmp-evo, icmp-evo-3         11/18/20 12:08:28.653151   11/18/20 12:08:28.669098   15945 usec
icmp-evo, icmp-evo-2         11/18/20 12:08:33.616890   11/18/20 12:08:33.628069   11176 usec
icmp-evo, icmp-evo-1         11/18/20 12:08:33.637811   11/18/20 12:08:33.656368   18556 usec
icmp-evo, icmp-evo-4         11/18/20 12:08:33.637819   11/18/20 12:08:33.652862   15039 usec
icmp-evo, icmp-evo-3         11/18/20 12:08:33.648109   11/18/20 12:08:33.662094   13983 usec
icmp-evo, icmp-evo-2         11/18/20 12:09:33.624341   11/18/20 12:09:33.636556   12205 usec
icmp-evo, icmp-evo-3         11/18/20 12:09:33.655442   11/18/20 12:09:33.681148   25704 usec
icmp-evo, icmp-evo-1         11/18/20 12:09:33.655454   11/18/20 12:09:33.690099   34644 usec
icmp-evo, icmp-evo-4         11/18/20 12:09:33.655456   11/18/20 12:09:33.679532   24073 usec
icmp-evo, icmp-evo-2         11/18/20 12:09:38.624449   11/18/20 12:09:38.636945   12493 usec
```

```
icmp-evo, icmp-evo-4         11/18/20 12:09:38.655661   11/18/20 12:09:38.675016   19351 usec
icmp-evo, icmp-evo-1         11/18/20 12:09:38.655690   11/18/20 12:09:38.676942   21250 usec
icmp-evo, icmp-evo-3         11/18/20 12:09:38.656034   11/18/20 12:09:38.682802   26766 usec
icmp-evo, icmp-evo-2         11/18/20 12:10:38.634733   11/18/20 12:10:38.646682   11929 usec
icmp-evo, icmp-evo-3         11/18/20 12:10:38.675453   11/18/20 12:10:38.695945   20489 usec
icmp-evo, icmp-evo-1         11/18/20 12:10:38.675859   11/18/20 12:10:38.700085   24225 usec
icmp-evo, icmp-evo-4         11/18/20 12:10:38.675945   11/18/20 12:10:38.701973   26026 usec
icmp-evo, icmp-evo-2         11/18/20 12:10:43.641129   11/18/20 12:10:43.657255   16108 usec
icmp-evo, icmp-evo-4         11/18/20 12:10:43.682202   11/18/20 12:10:43.701403   19190 usec
icmp-evo, icmp-evo-1         11/18/20 12:10:43.682684   11/18/20 12:10:43.704907   22220 usec
icmp-evo, icmp-evo-3         11/18/20 12:10:43.682981   11/18/20 12:10:43.710825   27843 usec
icmp-evo, icmp-evo-2         11/18/20 12:11:43.658117   11/18/20 12:11:43.680440   22312 usec
icmp-evo, icmp-evo-4         11/18/20 12:11:43.699092   11/18/20 12:11:43.716991   17896 usec
icmp-evo, icmp-evo-3         11/18/20 12:11:43.709647   11/18/20 12:11:43.733250   23601 usec
icmp-evo, icmp-evo-1         11/18/20 12:11:43.709653   11/18/20 12:11:43.737341   27687 usec
icmp-evo, icmp-evo-2         11/18/20 12:11:48.656468   11/18/20 12:11:48.670180   13708 usec
icmp-evo, icmp-evo-4         11/18/20 12:11:48.697172   11/18/20 12:11:48.708817   11642 usec
icmp-evo, icmp-evo-1         11/18/20 12:11:48.707382   11/18/20 12:11:48.729500   22116 usec
icmp-evo, icmp-evo-3         11/18/20 12:11:48.707621   11/18/20 12:11:48.735767   28144 usec
icmp-evo, icmp-evo-2         11/18/20 12:12:48.669000   11/18/20 12:12:48.688289   19279 usec
icmp-evo, icmp-evo-4         11/18/20 12:12:48.710047   11/18/20 12:12:48.721077   11027 usec
icmp-evo, icmp-evo-1         11/18/20 12:12:48.731012   11/18/20 12:12:48.749973   18959 usec
... (output truncated)
```

show services monitoring rpm history-results owner detail

```
user@host> show services monitoring rpm history-results owner icmp-junos detail

Owner: icmp-junos, Test: icmp-junos-2, Probe type: icmp-ping
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:12:40.911234
    Probe rcvd time: 11/18/20 12:12:40.918908, Client offload timestamping
    Rtt: 7674 usec, Rtt jitter: 0 usec
  Results over current test:
    Probes sent: 1, Probes received: 1, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 1, Minimum: 7674, Maximum: 7674, Average: 7674, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-3, Probe type: icmp-ping-timestamp
  Probe results:
```

```
    Probe response received
    Probe sent time: 11/18/20 12:12:41.821260
    Probe rcvd time: 11/18/20 12:12:41.832135, Client offload timestamping
    Rtt: 10875 usec, Rtt jitter: 0 usec
    Egress jitter: 0 usec, Ingress jitter: 0 usec
  Results over current test:
    Probes sent: 1, Probes received: 1, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 1, Minimum: 10875, Maximum: 10875, Average: 10875, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-4, Probe type: icmp-ping-timestamp
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:12:41.821327
    Probe rcvd time: 11/18/20 12:12:41.831576, Client offload timestamping
    Rtt: 10249 usec, Rtt jitter: 0 usec
    Egress jitter: 0 usec, Ingress jitter: 0 usec
  Results over current test:
    Probes sent: 1, Probes received: 1, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 1, Minimum: 10249, Maximum: 10249, Average: 10249, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-1, Probe type: icmp-ping
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:12:45.167546
    Probe rcvd time: 11/18/20 12:12:45.174508, Client offload timestamping
    Rtt: 6962 usec, Rtt jitter: 0 usec
  Results over current test:
    Probes sent: 1, Probes received: 1, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 1, Minimum: 6962, Maximum: 6962, Average: 6962, Stddev: 0
Owner: icmp-junos, Test: icmp-junos-2, Probe type: icmp-ping
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:12:45.906709
    Probe rcvd time: 11/18/20 12:12:45.912594, Client offload timestamping
    Rtt: 5885 usec, Rtt jitter: 1789 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 5885, Maximum: 7674, Average: 6779, Stddev: 898
    Measurement: Round trip jitter (usec)
```

```
      Samples: 2, Minimum: 1789, Maximum: 1789, Average: 1789, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-4, Probe type: icmp-ping-timestamp
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:12:46.815590
    Probe rcvd time: 11/18/20 12:12:46.823159, Client offload timestamping
    Rtt: 7569 usec, Rtt jitter: 2680 usec
    Egress jitter: 2263 usec, Ingress jitter: 417 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 7569, Maximum: 10249, Average: 8909, Stddev: 1340
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 2680, Maximum: 2680, Average: 2680, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-3, Probe type: icmp-ping-timestamp
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:12:46.815597
    Probe rcvd time: 11/18/20 12:12:46.823684, Client offload timestamping
    Rtt: 8087 usec, Rtt jitter: 2788 usec
    Egress jitter: 2337 usec, Ingress jitter: 451 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 8087, Maximum: 10875, Average: 9481, Stddev: 1394
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 2788, Maximum: 2788, Average: 2788, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-1, Probe type: icmp-ping
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:12:50.170153
    Probe rcvd time: 11/18/20 12:12:50.177224, Client offload timestamping
    Rtt: 7071 usec, Rtt jitter: 109 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 6962, Maximum: 7071, Average: 7016, Stddev: 99
    Measurement: Round trip jitter (usec)
```

```
      Samples: 2, Minimum: 109, Maximum: 109, Average: 109, Stddev: 0
... (output truncated)
```

Release Information

Command introduced in Junos OS Evolved Release 20.1R1.

show services monitoring rpm probe-results

IN THIS SECTION

- Syntax | 25794
- Description | 25794
- Options | 25795
- Required Privilege Level | 25795
- Output Fields | 25795
- Sample Output | 25801
- Release Information | 25804

Syntax

```
show services monitoring rpm probe-results
<owner name>
<source-address address>
<status (FAIL | PASS) >
<target address>
<test name>
```

Description

Display the results of the most recent real-time performance monitoring (RPM) probes.

Options

owner *name* (Optional) Display information only for probes with the specified probe owner name. `owner` is required if you want to specify the `test` option.

source-address *address* (Optional) Display information only for probes with the specified source address.

status (Optional) Display information only for probes with the specified type of test result. Specify one of the following:

- **FAIL** Failed tests
- **PASS** Passed tests

target *address* (Optional) Display information only for probes with the specified target address.

test *name* (Optional) Display information only for the specified test. You must also specify the `owner` option.

Required Privilege Level

view

Output Fields

Table 1451 on page 25795 lists the output fields for the `show services monitoring rpm probe-results` command. Output fields are listed in the approximate order in which they appear.

Table 1451: show services rpm probe-results Output Fields

Field Name	Field Description	Level of Output
`Owner`	Owner name. When you configure the probe owner statement at the `[edit services monitoring rpm]` hierarchy level, this field displays the configured owner name.	`owner`, `source-address`, `target`, `test`
`Test`	Name of a test representing a collection of probes. When you configure the `test` *`name`* statement at the `[edit services monitoring rpm owner]` hierarchy level, the field displays the configured test name.	All levels

Table 1451: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Target address`	Destination IPv4 address used for the probes. This field is displayed when the probes are sent to the configured IPv4 or IPv6 targets or RPM servers.	`owner`, `source-address`, `target`, `test`
`Source address`	Source address used for the probes.	`owner`, `source-address`, `target`, `test`
`Probe type`	Protocol configured on the receiving probe server: `icmp-ping`, `icmp-ping-timestamp`, `udp-ping`, `udp-ping-timestamp`.	`owner`, `source-address`, `target`, `test`
`Test size`	Number of probes within a test.	`owner`, `source-address`, `target`, `test`,

Table 1451: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Probe results`	Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement:
	• `Probe response received`
	  • `Probe sent time`—Timestamp when the probe results were sent.
	  • `Probe rcvd time`—Timestamp when the probe results were received.
	• `Client and server hardware timestamps` or `Client offload timestamping`—If timestamps are configured, an entry appears at this point.
	• `Rtt`—Average ping round-trip time (RTT), in microseconds.
	• `Round trip jitter`—Round-trip jitter, in microseconds.
	• `Egress jitter`—Egress delay, in microseconds.
	• `Ingress jitter`—Ingress delay, in microseconds.
		`owner`, `source-address`, `target`, `test`

Table 1451: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Results over current test`	Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes.
	• `Probes sent`—Number of probes sent within the current test.
	• `Probes received`—Number of probe responses received within the current test.
	•
	• `Loss percentage`—Percentage of lost probes for the current test.
	• `Measurement`—Measurement type: round-trip time or round-trip jitter
	For each measurement type, the following individual calculated results are provided:
	  • `Samples`—Number of probes.
	  • `Minimum`—Minimum RTT measured over the course of the current test.
	  • `Maximum`—Maximum RTT measured over the course of the current test.
	  • `Average`—Average RTT measured over the course of the current test.
	  • `Stddev`—Standard deviation, in microseconds.
		`owner`, `source-address`, `target`, `test`

Table 1451: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Results over last test`	Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed
	• `Probes sent`—Number of probes sent for the most recently completed test.
	• `Probes received`—Number of probe responses received for the most recently completed test.
	• `Loss percentage`—Percentage of lost probes for the most recently completed test.
	• `Measurement`—Measurement type: round-trip time or round-trip jitter
	For each measurement type, the following individual calculated results are provided:
	  • `Samples`—Number of probes.
	  • `Minimum`—Minimum RTT measured for the most recently completed test.
	  • `Maximum`—Maximum RTT measured for the most recently completed test.
	  • `Average`—Average RTT for the most recently completed test.
	  • `Stddev`—Standard deviation, in microseconds.
		`owner`, `source-address`, `target`, `test`

Table 1451: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Results over all tests`	Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test.
	• `Probes sent`—Number of probes sent in all tests.
	• `Probes received`—Number of probe responses received in all tests.
	• `Loss percentage`—Percentage of lost probes in all tests.
	• `Measurement`—Measurement type: round-trip time or round-trip jitter
	For each measurement type, the following individual calculated results are provided:
	  • `Samples`—Number of probes.
	  • `Minimum`—Minimum RTT measured over the course of the current test.
	  • `Maximum`—Maximum RTT measured over the course of the current test.
	  • `Average`—Average RTT measured over the course of the current test.
	  • `Stddev`—Standard deviation, in microseconds.
		`owner`, `source-address`, `target`, `test`

`Error Stats`	Displays error statistics for each probe. • `Invalid client recv timestamp`—Number of client receive timestamp less than client send timestamp. • `Invalid server send timestamp`—Number of server send timestamp less than server receive timestamp. • `Invalid server processing time`—Number of server side spent time greater than RTT. **NOTE**: `Error Stats` is displayed in the output only if non-zero statistics exist.	`owner`, `source-address`, `target`, `test`

Sample Output

show services monitoring rpm probe-results

show services monitoring rpm probe-results owner

```
user@host> show services monitoring rpm probe-results owner icmp-junos

Owner: icmp-junos, Test: icmp-junos-1
Target address: 10.3.1.2, Probe type: icmp-ping, Test size: 2
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:42:05.426631
    Probe rcvd time: 11/18/20 12:42:05.438092, Client offload timestamping
    Rtt: 11461 usec, Rtt jitter: 4982 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 6479, Maximum: 11461, Average: 8970, Stddev: 2491
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 4982, Maximum: 4982, Average: 4982, Stddev: 0
  Results over last test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 5845, Maximum: 7279, Average: 6562, Stddev: 717
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 1434, Maximum: 1434, Average: 1434, Stddev: 0
  Results over all tests:
    Probes sent: 442, Probes received: 438, Loss percentage: 0.9
    Measurement: Round trip time (usec)
      Samples: 438, Minimum: 4783, Maximum: 41643, Average: 8081, Stddev: 3003
    Measurement: Round trip jitter (usec)
      Samples: 438, Minimum: 3, Maximum: 35766, Average: 2564, Stddev: 3372

Owner: icmp-junos, Test: icmp-junos-2
Target address: 192.168.22.22, Source address: 192.168.11.11, Probe type: icmp-ping, Test size: 2
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:42:01.146976
    Probe rcvd time: 11/18/20 12:42:01.160419, Client offload timestamping
```

```
    Rtt: 13443 usec, Rtt jitter: 6738 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 6705, Maximum: 13443, Average: 10074, Stddev: 3369
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 6738, Maximum: 6738, Average: 6738, Stddev: 0
  Results over last test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 5884, Maximum: 11405, Average: 8644, Stddev: 2762
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 5521, Maximum: 5521, Average: 5521, Stddev: 0
  Results over all tests:
    Probes sent: 442, Probes received: 439, Loss percentage: 0.67
    Measurement: Round trip time (usec)
      Samples: 439, Minimum: 5479, Maximum: 42406, Average: 8884, Stddev: 3884
    Measurement: Round trip jitter (usec)
      Samples: 439, Minimum: 7, Maximum: 36177, Average: 3384, Stddev: 4437

Owner: icmp-junos, Test: icmp-junos-3
Target address: 10.4.1.2, Probe type: icmp-ping-timestamp, Test size: 2
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:42:02.128854
    Probe rcvd time: 11/18/20 12:42:02.142671, Client offload timestamping
    Rtt: 13817 usec, Rtt jitter: 6024 usec
    Egress jitter: 2366 usec, Ingress jitter: 3658 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 7793, Maximum: 13817, Average: 10805, Stddev: 3012
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 6024, Maximum: 6024, Average: 6024, Stddev: 0
  Results over last test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 8043, Maximum: 10271, Average: 9157, Stddev: 1114
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 2228, Maximum: 2228, Average: 2228, Stddev: 0
  Results over all tests:
    Probes sent: 442, Probes received: 440, Loss percentage: 0.45
    Measurement: Round trip time (usec)
```

```
      Samples: 440, Minimum: 5486, Maximum: 102674, Average: 12544, Stddev: 11088
    Measurement: Round trip jitter (usec)
      Samples: 440, Minimum: 1, Maximum: 92729, Average: 6896, Stddev: 13940

Owner: icmp-junos, Test: icmp-junos-4
Target address: 10.0.22.22, Source address: 10.0.11.11, Probe type: icmp-ping-timestamp, Test
size: 2
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:42:02.098433
    Probe rcvd time: 11/18/20 12:42:02.140279, Client offload timestamping
    Rtt: 41846 usec, Rtt jitter: 33904 usec
    Egress jitter: 19853 usec, Ingress jitter: 14051 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 7942, Maximum: 41846, Average: 24894, Stddev: 16952
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 33904, Maximum: 33904, Average: 33904, Stddev: 0
  Results over last test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 8066, Maximum: 13446, Average: 10756, Stddev: 2690
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 5380, Maximum: 5380, Average: 5380, Stddev: 0
  Results over all tests:
    Probes sent: 442, Probes received: 439, Loss percentage: 0.67
    Measurement: Round trip time (usec)
      Samples: 439, Minimum: 5771, Maximum: 100477, Average: 12726, Stddev: 11125
    Measurement: Round trip jitter (usec)
      Samples: 439, Minimum: 6, Maximum: 92178, Average: 7331, Stddev: 14240
```

show services monitoring rpm probe-results owner test

```
user@host> show services monitoring rpm probe-results owner icmp-junos test icmp-junos-1

Owner: icmp-junos, Test: icmp-junos-1
Target address: 10.3.1.2, Probe type: icmp-ping, Test size: 2
  Probe results:
    Probe response received
    Probe sent time: 11/18/20 12:48:35.484952
```

```
    Probe rcvd time: 11/18/20 12:48:35.493256, Client offload timestamping
    Rtt: 8304 usec, Rtt jitter: 2524 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 5780, Maximum: 8304, Average: 7042, Stddev: 1262
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 2524, Maximum: 2524, Average: 2524, Stddev: 0
  Results over last test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 7124, Maximum: 7768, Average: 7446, Stddev: 322
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 644, Maximum: 644, Average: 644, Stddev: 0
  Results over all tests:
    Probes sent: 454, Probes received: 450, Loss percentage: 0.88
    Measurement: Round trip time (usec)
      Samples: 450, Minimum: 4783, Maximum: 41643, Average: 8072, Stddev: 2976
    Measurement: Round trip jitter (usec)
      Samples: 450, Minimum: 2, Maximum: 35766, Average: 2535, Stddev: 3340
```

Release Information

Command introduced in Junos OS Evolved Release 20.1R1.

show services monitoring twamp client control-info

IN THIS SECTION

- Syntax | 25805
- Description | 25805
- Options | 25805
- Required Privilege Level | 25805
- Output Fields | 25805
- Sample Output | 25806
- Release Information | 25806

Syntax

```
show services monitoring twamp client control-info
<control-connection control-connection-name>
```

Description

Display information about the control connections established between the Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, all established sessions are displayed, unless you specify the control-connection option when you issue the command.

Options

control-connection *control-connection-name*	(Optional) Display information about the specified control connection.

Required Privilege Level

view

Output Fields

Table 1452 on page 25805 lists the output fields for the `show services monitoring twamp client control-info` command. Output fields are listed in the approximate order in which they appear.

Table 1452: show services monitoring twamp client control-info Output Fields

Field Name	Field Description	Levels of Output
`Control name`	Control-connection name that uniquely identifies the connection between the TWAMP server and a particular client	All
`Client address:port`	Client IP address and port number	All
`Server address:port`	Server IP address and port number	All

Table 1452: show services monitoring twamp client control-info Output Fields *(Continued)*

Field Name	Field Description	Levels of Output
`Active tests`	Number of active tests	All
`Control status`	The status of the connection; can be sleeping, testing, or disconnected	All
`Auth mode`	Authentication mode	All

Sample Output

show services monitoring twamp client control-info

```
user@host> show services monitoring twamp client control-info

Control                 Client                  Server                  Active   Control       Auth
name                    address:port            address:port            tests    status        mode
mngd-deflt-1            10.6.1.1:57629          10.6.1.2:60000          2        testing       none
mngd-deflt-2            10.5.1.1:52777          10.5.1.2:862            0        sleeping      none
mngd-inst-1             10.0.1.1:43033          10.0.3.3:60001          2        testing       none
mngd-inst-2             10.0.1.1:34167          10.0.2.2:862            2        testing       none
```

show services monitoring twamp client control-info control-connection

```
user@host> show services monitoring twamp client control-info control-connection mngd-inst-1

Control                 Client                  Server                  Active   Control       Auth
name                    address:port            address:port            tests    status        mode
mngd-inst-1             10.0.1.1:43033          10.0.3.3:60001          2        testing       none
```

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

control-connection (Junos OS Evolved) | 4627

Understanding Two-Way Active Measurement Protocol on Routers

show services monitoring twamp client history-results

IN THIS SECTION

- Syntax | 25807
- Description | 25807
- Options | 25808
- Required Privilege Level | 25808
- Output Fields | 25808
- Sample Output | 25810
- Release Information | 25814

Syntax

```
show services rpm twamp client history-results
control-connection control-connection-name
<detail>
<since time>
<source-address address>
<target address>
<test-session test-session-name>
```

Description

Display standard information about the results of the last 50 probes for a Two-Way Active Measurement Protocol (TWAMP) control connection. You can also view the history results of the probes or test packets sent from a TWAMP client to a TWAMP server by source address, by target address, or for a specific test session associated with the control connection. To change the number of probe results

displayed, configure the history-size statement at the [edit monitoring twamp client control-connection *control-connection-name* test-session *test-session-name* hierarchy level.

Options

<control-connection *control-connection-name*>	(Required) Display information for the specified control connection between a TWAMP client and a TWAMP server.
detail	(Optional) Display detailed information about the control connection.
since *time*	(Optional) Display information from the specified time. Specify time as *yyyy-mm-dd.hh:mm:ss*.
source-address *address*	(Optional) Display information only for those probes with the specified IPv4 source address.
target *address*	(Optional) Display information only for those probes with the specified IPv4 target address.
test-session *test-session-name*	(Optional) Display information for the specified test session associated with the control connection between a TWAMP client and a TWAMP server.

Required Privilege Level

view

Output Fields

Table 1453 on page 25808 lists the output fields for the `show services monitoring twamp client history-results` command. Output fields are listed in the approximate order in which they appear.

Table 1453: show services monitoring twamp client history-results Output Fields

Field Name	Field Description	Level of Output
`Owner`	Probe owner or the TWAMP client.	All levels
`Test`	Name of a test for a TWAMP probe instance.	All levels

Table 1453: show services monitoring twamp client history-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Probe received`	Timestamp when the probe result was determined.	All levels
`Round trip time`	Average ping round-trip time (RTT), in microseconds.	All levels
`Probe results`	Result of a particular probe performed by a remote host. The following information is contained in the results: • `Probe response received`—Timestamp when the probe result was determined. • `Probe sent time`—Timestamp when the probe was sent. • `Probe rcvd time`—Timestamp when the probe was received. • `Rtt`—Average ping round-trip time (RTT), in microseconds. • `Rtt jitter`—Average ping round-trip time (RTT) jitter, in microseconds. • `Egress jitter`—Average ping egress jitter, in microseconds. • `Ingress jitter`—Average ping ingress jitter, in microseconds.	`detail`
`Results over current test`	Displays the results for the current test by probe at the time each probe was completed, as well as the status of the current test at the time the probe was completed.	`detail`
`Probes sent`	Number of probes sent with the current test.	`detail`
`Probes received`	Number of probe responses received within the current test.	`detail`
Egress loss inferred	Computed unidirectional loss towards server (available only when `stateful-sequence` is enabled on the server)	detail
`Loss percentage`	Percentage of lost probes for the current test.	`detail`

Table 1453: show services monitoring twamp client history-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Measurement`	Measurements of minimum, maximum, average, and standard deviation for each type of measurement: . • `Minimum`—Minimum measured over the course of the current test. • `Maximum`—Maximum measured over the course of the current test. • `Average`—Average measured over the course of the current test. • `Stddev`—Standard deviation measured over the course of the current test. Types of measurements: • `Round trip time (usec)`—Ping round trip time (RTT) • `Egress delay (usec)`—Unidirectional one way delay from client to server • `Egress jitter (usec)`—Jitter between successive measures on egress • `Ingress delay (usec)`—Unidirectional one way delay from server to client • `Ingress jitter (usec)`—Jitter between successive measures on ingress	`detail`

Sample Output

show services monitoring twamp client history-results control-connection

```
user@host> show services monitoring twamp client history-results control-connection mngd-deflt
Owner, Test                  Probe sent                 Probe received             Round trip
time
mngd-deflt, test-2        08/27/20 18:32:40.532435   08/27/20 18:32:40.553237   20216 usec
mngd-deflt, test-1        08/27/20 18:32:40.532536   08/27/20 18:32:40.555498   22921 usec
mngd-deflt, test-1        08/27/20 18:33:40.623973   08/27/20 18:33:40.647633   23597 usec
mngd-deflt, test-2        08/27/20 18:33:40.624073   08/27/20 18:33:40.648176   24070 usec
mngd-deflt, test-2        08/27/20 18:34:40.630696   08/27/20 18:34:40.641727   10992 usec
mngd-deflt, test-1        08/27/20 18:34:40.630741   08/27/20 18:34:40.644139   13350 usec
```

```
mngd-deflt, test-1        08/27/20 18:35:40.696545   08/27/20 18:35:40.711118   14501 usec
mngd-deflt, test-2        08/27/20 18:35:40.696878   08/27/20 18:35:40.711692   14778 usec
mngd-deflt, test-2        08/27/20 18:36:40.694764   08/27/20 18:36:40.711443   16643 usec
mngd-deflt, test-1        08/27/20 18:36:40.694885   08/27/20 18:36:40.713718   18792 usec
mngd-deflt, test-1        08/27/20 18:37:40.765843   08/27/20 18:37:40.778095   12209 usec
mngd-deflt, test-2        08/27/20 18:37:40.766109   08/27/20 18:37:40.778821   12678 usec
mngd-deflt, test-2        08/27/20 18:38:40.768726   08/27/20 18:38:40.781763   13001 usec
mngd-deflt, test-1        08/27/20 18:38:40.768772   08/27/20 18:38:40.782283   13478 usec
mngd-deflt, test-1        08/27/20 18:39:40.838053   08/27/20 18:39:40.862990   24891 usec
mngd-deflt, test-2        08/27/20 18:39:40.838387   08/27/20 18:39:40.863068   24618 usec
mngd-deflt, test-2        08/27/20 18:40:40.841528   08/27/20 18:40:40.859848   18216 usec
mngd-deflt, test-1        08/27/20 18:40:40.841584   08/27/20 18:40:40.861503   19865 usec
mngd-deflt, test-1        08/27/20 18:41:40.929306   08/27/20 18:41:40.947766   18424 usec
mngd-deflt, test-2        08/27/20 18:41:40.929475   08/27/20 18:41:40.948378   18871 usec
mngd-deflt, test-2        08/27/20 18:42:40.923137   08/27/20 18:42:40.941889   18673 usec
mngd-deflt, test-1        08/27/20 18:42:40.923192   08/27/20 18:42:40.943715   20475 usec
mngd-deflt, test-1        08/27/20 18:43:41.003976   08/27/20 18:43:41.021443   17453 usec
mngd-deflt, test-2        08/27/20 18:43:41.004567   08/27/20 18:43:41.022093   17492 usec
mngd-deflt, test-2        08/27/20 18:44:41.010953   08/27/20 18:44:41.032575   21564 usec
mngd-deflt, test-1        08/27/20 18:44:41.011250   08/27/20 18:44:41.034401   23110 usec
mngd-deflt, test-1        08/27/20 18:45:41.106478   08/27/20 18:45:41.124683   18173 usec
mngd-deflt, test-2        08/27/20 18:45:41.106667   08/27/20 18:45:41.125132   18433 usec
mngd-deflt, test-2        08/27/20 18:46:41.109541   08/27/20 18:46:41.132243   21855 usec
mngd-deflt, test-1        08/27/20 18:46:41.110129   08/27/20 18:46:41.137104   26935 usec
mngd-deflt, test-1        08/27/20 18:47:41.200747   08/27/20 18:47:41.217530   16731 usec
mngd-deflt, test-2        08/27/20 18:47:41.201048   08/27/20 18:47:41.218069   16982 usec
mngd-deflt, test-2        08/27/20 18:48:41.200511   08/27/20 18:48:41.226252   25662 usec
mngd-deflt, test-1        08/27/20 18:48:41.200678   08/27/20 18:48:41.227703   27010 usec
mngd-deflt, test-1        08/27/20 18:49:41.273230   08/27/20 18:49:41.289343   16094 usec
mngd-deflt, test-2        08/27/20 18:49:41.273428   08/27/20 18:49:41.289450   16003 usec
mngd-deflt, test-2        08/27/20 18:50:41.283700   08/27/20 18:50:41.296895   13179 usec
mngd-deflt, test-1        08/27/20 18:50:41.283950   08/27/20 18:50:41.297335   13373 usec
mngd-deflt, test-1        08/27/20 18:51:41.342941   08/27/20 18:51:41.368757   25778 usec
mngd-deflt, test-2        08/27/20 18:51:41.343179   08/27/20 18:51:41.368799   25585 usec
mngd-deflt, test-2        08/27/20 18:52:41.360424   08/27/20 18:52:41.376456   15989 usec
mngd-deflt, test-1        08/27/20 18:52:41.360560   08/27/20 18:52:41.377131   16539 usec
mngd-deflt, test-1        08/27/20 18:53:41.426307   08/27/20 18:53:41.440569   14224 usec
mngd-deflt, test-2        08/27/20 18:53:41.426425   08/27/20 18:53:41.440605   14137 usec
mngd-deflt, test-2        08/27/20 18:54:41.431346   08/27/20 18:54:41.443324   11958 usec
mngd-deflt, test-1        08/27/20 18:54:41.432825   08/27/20 18:54:41.443912   11074 usec
mngd-deflt, test-1        08/27/20 18:55:41.498060   08/27/20 18:55:41.511760   13664 usec
mngd-deflt, test-2        08/27/20 18:55:41.498084   08/27/20 18:55:41.512597   14469 usec
mngd-deflt, test-2        08/27/20 18:56:41.497135   08/27/20 18:56:41.513473   15792 usec
```

```
mngd-deflt, test-1        08/27/20 18:56:41.497152   08/27/20 18:56:41.513544   16380 usec
mngd-deflt, test-1        08/27/20 18:57:41.579541   08/27/20 18:57:41.592352   12750 usec
mngd-deflt, test-2        08/27/20 18:57:41.579718   08/27/20 18:57:41.592905   13143 usec
mngd-deflt, test-2        08/27/20 18:58:41.575705   08/27/20 18:58:41.587480   11738 usec
mngd-deflt, test-1        08/27/20 18:58:41.575725   08/27/20 18:58:41.591966   16195 usec
mngd-deflt, test-1        08/27/20 18:59:41.644790   08/27/20 18:59:41.661570   16763 usec
mngd-deflt, test-2        08/27/20 18:59:41.644812   08/27/20 18:59:41.662216   17390 usec
mngd-deflt, test-2        08/27/20 19:00:41.646420   08/27/20 19:00:41.658933   12495 usec
mngd-deflt, test-1        08/27/20 19:00:41.646725   08/27/20 19:00:41.659386   12646 usec
mngd-deflt, test-1        08/27/20 19:01:41.710108   08/27/20 19:01:41.725249   15097 usec
mngd-deflt, test-2        08/27/20 19:01:41.710156   08/27/20 19:01:41.726253   16059 usec
mngd-deflt, test-2        08/27/20 19:02:41.710661   08/27/20 19:02:41.728264   17552 usec
mngd-deflt, test-1        08/27/20 19:02:41.711725   08/27/20 19:02:41.728481   16725 usec
mngd-deflt, test-1        08/27/20 19:03:41.775533   08/27/20 19:03:41.792378   16816 usec
mngd-deflt, test-2        08/27/20 19:03:41.775601   08/27/20 19:03:41.792795   17133 usec
mngd-deflt, test-2        08/27/20 19:04:41.781272   08/27/20 19:04:41.799808   18327 usec
mngd-deflt, test-1        08/27/20 19:04:41.781391   08/27/20 19:04:41.801613   20187 usec
mngd-deflt, test-1        08/27/20 19:05:41.862516   08/27/20 19:05:41.883037   20497 usec
mngd-deflt, test-2        08/27/20 19:05:41.862527   08/27/20 19:05:41.883585   21018 usec
mngd-deflt, test-2        08/27/20 19:06:41.869384   08/27/20 19:06:41.883258   13837 usec
mngd-deflt, test-1        08/27/20 19:06:41.869466   08/27/20 19:06:41.884977   15475 usec
mngd-deflt, test-1        08/27/20 19:07:41.945802   08/27/20 19:07:41.964598   18758 usec
mngd-deflt, test-2        08/27/20 19:07:41.945817   08/27/20 19:07:41.965232   19379 usec
mngd-deflt, test-2        08/27/20 19:08:41.953929   08/27/20 19:08:41.972951   18950 usec
mngd-deflt, test-1        08/27/20 19:08:41.954886   08/27/20 19:08:41.974187   19258 usec
mngd-deflt, test-1        08/27/20 19:09:42.037577   08/27/20 19:09:42.056201   18570 usec
mngd-deflt, test-2        08/27/20 19:09:42.037629   08/27/20 19:09:42.056842   19171 usec
mngd-deflt, test-2        08/27/20 19:10:42.039524   08/27/20 19:10:42.056397   16793 usec
mngd-deflt, test-1        08/27/20 19:10:42.039601   08/27/20 19:10:42.059688   20038 usec
mngd-deflt, test-1        08/27/20 19:11:42.100698   08/27/20 19:11:42.118606   17881 usec
mngd-deflt, test-2        08/27/20 19:11:42.100942   08/27/20 19:11:42.119158   18188 usec
mngd-deflt, test-2        08/27/20 19:12:42.102936   08/27/20 19:12:42.122253   18728 usec
mngd-deflt, test-1        08/27/20 19:12:42.103062   08/27/20 19:12:42.125590   22511 usec
mngd-deflt, test-1        08/27/20 19:13:42.170855   08/27/20 19:13:42.192140   21239 usec
mngd-deflt, test-2        08/27/20 19:13:42.171132   08/27/20 19:13:42.192333   21158 usec
mngd-deflt, test-2        08/27/20 19:14:42.162464   08/27/20 19:14:42.182361   19794 usec
mngd-deflt, test-1        08/27/20 19:14:42.162601   08/27/20 19:14:42.184849   22206 usec
mngd-deflt, test-1        08/27/20 19:15:42.239051   08/27/20 19:15:42.257361   18283 usec
mngd-deflt, test-2        08/27/20 19:15:42.239516   08/27/20 19:15:42.257933   18388 usec
mngd-deflt, test-2        08/27/20 19:16:42.242182   08/27/20 19:16:42.263899   21610 usec
mngd-deflt, test-1        08/27/20 19:16:42.242245   08/27/20 19:16:42.265736   23434 usec
mngd-deflt, test-1        08/27/20 19:17:42.346073   08/27/20 19:17:42.363597   17477 usec
mngd-deflt, test-2        08/27/20 19:17:42.346432   08/27/20 19:17:42.364867   18345 usec
```

```
mngd-deflt, test-2        08/27/20 19:18:42.349453   08/27/20 19:18:42.378702   29204 usec
mngd-deflt, test-1        08/27/20 19:18:42.349502   08/27/20 19:18:42.378729   29171 usec
mngd-deflt, test-1        08/27/20 19:19:42.430651   08/27/20 19:19:42.444923   13792 usec
mngd-deflt, test-2        08/27/20 19:19:42.430842   08/27/20 19:19:42.446115   15236 usec
mngd-deflt, test-2        08/27/20 19:20:42.428240   08/27/20 19:20:42.443861   15571 usec
mngd-deflt, test-1        08/27/20 19:20:42.428294   08/27/20 19:20:42.444358   16020 usec
mngd-deflt, test-1        08/27/20 19:21:42.494229   08/27/20 19:21:42.509517   15217 usec
mngd-deflt, test-2        08/27/20 19:21:42.494436   08/27/20 19:21:42.510934   16478 usec
```

show services monitoring twamp client history-results control-connection detail

```
user@host> show services monitoring twamp client history-results control-connection mngd-deflt
detail
Owner: mngd-deflt, Test: test-1, Probe type: twamp
  Probe results:
    Probe response received
    Probe sent time: 08/27/20 18:35:40.696545
    Probe rcvd time: 08/27/20 18:35:40.711118, Client and server offload timestamping
    Rtt: 14501 usec, Rtt jitter: 0 usec
    Egress jitter: 0 usec, Ingress jitter: 0 usec
  Results over current test:
    Probes sent: 1, Probes received: 1, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 1, Minimum: 14501, Maximum: 14501, Average: 14501, Stddev: 0

Owner: mngd-deflt, Test: test-2, Probe type: twamp
  Probe results:
    Probe response received
    Probe sent time: 08/27/20 18:35:40.696878
    Probe rcvd time: 08/27/20 18:35:40.711692, Client and server offload timestamping
    Rtt: 14778 usec, Rtt jitter: 0 usec
    Egress jitter: 0 usec, Ingress jitter: 0 usec
  Results over current test:
    Probes sent: 1, Probes received: 1, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 1, Minimum: 14778, Maximum: 14778, Average: 14778, Stddev: 0
...
```

Release Information

Command introduced in Junos OS Evolved 20.3R1.

Command output fields for inferred egress loss, egress delay and jitter, and ingress delay and jitter introduced in Junos OS Evolved 23.4R1.

RELATED DOCUMENTATION

control-connection (Junos OS Evolved) | 4627

Understanding Two-Way Active Measurement Protocol on Routers

show services monitoring twamp client probe-results

IN THIS SECTION

- Syntax | 25814
- Description | 25815
- Options | 25815
- Required Privilege Level | 25815
- Output Fields | 25815
- Sample Output | 25820
- Release Information | 25825

Syntax

```
show services monitoring twamp client probe-results
<control-connection control-connection-name>
<source-address>
<status (FAIL | PASS)>
<target>
<test-session test-session-name>
```

Description

Display the results of the most recent Two-Way Active Measurement Protocol (TWAMP) probes. By default, all probe results are displayed, unless you specify the control-connection option when you issue the command. The control-connection option is required when you want to issue the command with one of the other options. For example, to see the probe results for a particular test session, you must also include the control-connection option to specify which connection contains the test session: `show services monitoring twamp client probe-results control-connection` *`control-connection-name`* `test-session` *`test-session-name`*.

Options

control-connection ***control-connection-name***
(Optional, but required if specifying one of other options as well) Display probe results for the specified control connection between a TWAMP client and a TWAMP server.

source-address ***address***
(Optional) Display results only for those probes with the specified IPv4 source address associated with a particular control connection.

status (FAIL | PASS)
(Optional) Display results only for those probes with the specified status associated with a particular control connection.

- **Values:** Specify one of the following:
  - FAIL—Display results for failed probes.
  - PASS—Display results for passed probes.

target ***address***
(Optional) Display results only for those probes with the specified IPv4 target address associated with a particular control connection.

test-session ***test-session-name***
(Optional) Display probe results for the specified test session associated with the control connection between a TWAMP client and a TWAMP server.

Required Privilege Level

view

Output Fields

Table 1454 on page 25816 lists the output fields for the `show services monitoring twamp client probe-results` command. Output fields are listed in the approximate order in which they appear.

Table 1454: show services monitoring twamp client probe-results Output Fields

Field Name	Field Description	Level of Output
`Owner`	Owner name. Entries are listed in alphabetical order by owner name and then by test name. When you configure the control-connection statement at the `[edit services monitoring twamp client]` hierarchy level, this field displays the configured control-connection name.	All levels
`Test`	Name of a test representing a collection of probes. When you configure the test-session test-name statement at the `[edit services monitoring twamp client control-connection]` hierarchy level, the field displays the configured test-session name.	All levels
`Target address`	Destination address used for the probes. This field is displayed when the probes are sent to the configured targets.	All levels
`Source address`	Source address used for the probes.	All levels
`Probe type`	Protocol configured on the receiving probe server: `twamp`	All levels
`Test size`	Number of probes within a test.	All levels

Table 1454: show services monitoring twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Probe results`	Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement:
	• `Probe response received`
	  • `Probe sent time`—Timestamp when the probe's results were sent.
	  • `Probe rcvd time`—Timestamp when the probe's results were received.
	• `Client and server offload timestamping`—If PFE timestamping is configured, an entry appears at this point.
	• `Rtt`—Average ping round-trip time (RTT), in microseconds.
	• `Rtt jitter`—Round-trip jitter, in microseconds.
	• `Egress jitter`—Egress jitter, in microseconds.
	• `Ingress jitter`—Ingress jitter, in microseconds.
		All levels

Table 1454: show services monitoring twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Results over current test`	Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes.
	• `Probes sent`—Number of probes sent within the current test.
	• `Probes received`—Number of probe responses received within the current test.
	• `Egress loss inferred`—Computed unidirectional loss towards server (available only when `stateful-sequence` is enabled on the server)
	• `Loss percentage`—Percentage of lost probes for the current test.
	• `Measurement`—Measurement type: round-trip time, round-trip jitter, egress delay, egress jitter, ingress delay, or ingress jitter.
	For each measurement type, the following individual calculated results are provided:
	  • `Samples`—Number of probes.
	  • `Minimum`—Minimum RTT measured over the course of the current test.
	  • `Maximum`—Maximum RTT measured over the course of the current test.
	  • `Average`—Average RTT measured over the course of the current test.
	  • `Stddev`—Standard deviation, in microseconds.
		All levels

Table 1454: show services monitoring twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Results over last test`	Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed.
	• `Probes sent`—Number of probes sent for the most recently completed test.
	• `Probes received`—Number of probe responses received for the most recently completed test.
	• `Egress loss inferred`—Computed unidirectional loss towards server (available only when `stateful-sequence` is enabled on the server)
	• `Loss percentage`—Percentage of lost probes for the most recently completed test.
	• `Measurement`—Measurement type: round-trip time, round-trip jitter, egress delay, egress jitter, ingress delay, or ingress jitter.
	For each measurement type, the following individual calculated results are provided:
	  • `Samples`—Number of probes.
	  • `Minimum`—Minimum RTT measured for the most recently completed test.
	  • `Maximum`—Maximum RTT measured for the most recently completed test.
	  • `Average`—Average RTT measured for the most recently completed test.
	  • `Stddev`—Standard deviation, in microseconds.
		All levels

Table 1454: show services monitoring twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Results over all tests`	Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test.
	• `Probes sent`—Number of probes sent in all tests.
	• `Probes received`—Number of probe responses received in all tests.
	• `Egress loss inferred`—Computed unidirectional loss towards server (available only when `stateful-sequence` is enabled on the server)
	• `Loss percentage`—Percentage of lost probes in all tests.
	• `Measurement`—Measurement type: round-trip time, round-trip jitter, egress delay, egress jitter, ingress delay, or ingress jitter.
	For each measurement type, the following individual calculated results are provided:
	  • `Samples`—Number of probes.
	  • `Minimum`—Minimum RTT measured over the course of the current test.
	  • `Maximum`—Maximum RTT measured over the course of the current test.
	  • `Average`—Average RTT measured over the course of the current test.
	  • `Stddev`—Standard deviation, in microseconds.
		All levels

Sample Output

show services monitoring twamp probe-results (Junos OS Evolved Release 23.4R1 and later)

```
user@host> show services monitoring twamp probe-results
Owner: light-remote-dflt-v64, Test: dflt-3-v4 Target address: 10.61.1.1, Source address:
10.61.1.2, Probe type: twamp, Test size: 100
  Probe results:
    Probe response received
    Probe sent time: 01/17/24 02:19:55.774425
    Probe rcvd time: 01/17/24 02:19:55.800658, Client and server offload timestamping
    Rtt: 14806 usec, Rtt jitter: 413 usec
```

```
    Egress jitter: 56 usec, Ingress jitter: 357 usec
  Results over current test:
    Probes sent: 41, Probes received: 19, Egress loss inferred: 22, Loss percentage: 53.66
    Measurement: Round trip time (usec)
      Samples: 19, Minimum: 12356, Maximum: 17668, Average: 14818, Stddev: 1136
    Measurement: Round trip jitter (usec)
      Samples: 18, Minimum: 251, Maximum: 5312, Average: 1612, Stddev: 1261
    Measurement: Egress delay (usec)
      Samples: 19, Minimum: 6333, Maximum: 10255, Average: 8279, Stddev: 854
    Measurement: Egress jitter (usec)
      Samples: 18, Minimum: 56, Maximum: 3922, Average: 1087, Stddev: 889
    Measurement: Ingress delay (usec)
      Samples: 19, Minimum: 5934, Maximum: 7413, Average: 6538, Stddev: 434
    Measurement: Ingress jitter (usec)
      Samples: 18, Minimum: 70, Maximum: 1390, Average: 603, Stddev: 411
  Results over last test:
    Probes sent: 100, Probes received: 100, Loss percentage: 0.00
    Measurement: Round trip time (usec)
      Samples: 100, Minimum: 12993, Maximum: 22654, Average: 15905, Stddev: 1860
    Measurement: Round trip jitter (usec)
      Samples: 99, Minimum: 6, Maximum: 8249, Average: 1840, Stddev: 1807
    Measurement: Egress delay (usec)
      Samples: 100, Minimum: 6584, Maximum: 15577, Average: 9123, Stddev: 1486
    Measurement: Egress jitter (usec)
      Samples: 99, Minimum: 4, Maximum: 7587, Average: 1475, Stddev: 1586
    Measurement: Ingress delay (usec)
      Samples: 100, Minimum: 5722, Maximum: 12514, Average: 6782, Stddev: 964
    Measurement: Ingress jitter (usec)
      Samples: 99, Minimum: 0, Maximum: 6321, Average: 773, Stddev: 1022
  Results over all tests:
    Probes sent: 141, Probes received: 119, Egress loss inferred: 22, Loss percentage: 15.60
    Measurement: Round trip time (usec)
      Samples: 119, Minimum: 12356, Maximum: 22654, Average: 15732, Stddev: 1804
    Measurement: Round trip jitter (usec)
      Samples: 118, Minimum: 6, Maximum: 8249, Average: 1796, Stddev: 1732
    Measurement: Egress delay (usec)
      Samples: 119, Minimum: 6333, Maximum: 15577, Average: 8988, Stddev: 1439
    Measurement: Egress jitter (usec)
      Samples: 118, Minimum: 4, Maximum: 7587, Average: 1406, Stddev: 1504
    Measurement: Ingress delay (usec)
      Samples: 119, Minimum: 5722, Maximum: 12514, Average: 6743, Stddev: 906
```

```
    Measurement: Ingress jitter (usec)
      Samples: 118, Minimum: 0, Maximum: 6321, Average: 744, Stddev: 952
```

show services monitoring twamp client probe-results

```
user@host> show services monitoring twamp client probe-results

Owner: light-deflt, Test: test-5
Target address: 10.0.1.2, Source address: 10.0.1.1, Probe type: twamp, Test size: 2
  Probe results:
    Probe response received
    Probe sent time: 08/27/20 19:50:25.549385
    Probe rcvd time: 08/27/20 19:50:25.564693, Client and server offload timestamping
    Rtt: 15224 usec, Rtt jitter: 673 usec
    Egress jitter: 850 usec, Ingress jitter: 177 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 15224, Maximum: 15897, Average: 15560, Stddev: 358
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 673, Maximum: 673, Average: 673, Stddev: 0
  Results over last test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 14762, Maximum: 16395, Average: 15578, Stddev: 825
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 1633, Maximum: 1633, Average: 1633, Stddev: 0
  Results over all tests:
    Probes sent: 714, Probes received: 714, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 714, Minimum: 8740, Maximum: 307066, Average: 20625, Stddev: 12145
    Measurement: Round trip jitter (usec)
      Samples: 714, Minimum: 6, Maximum: 292383, Average: 5420, Stddev: 15797

Owner: light-deflt, Test: test-6
Target address: 192.168.33.33, Source address: 192.168.11.11, Probe type: twamp, Test size: 2
  Probe results:
    Probe response received
    Probe sent time: 08/27/20 19:50:25.549497
    Probe rcvd time: 08/27/20 19:50:25.564740, Client and server offload timestamping
    Rtt: 15180 usec, Rtt jitter: 125 usec
```

```
    Egress jitter: 114 usec, Ingress jitter: 11 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 15180, Maximum: 15305, Average: 15242, Stddev: 138
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 125, Maximum: 125, Average: 125, Stddev: 0
  Results over last test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 14823, Maximum: 16362, Average: 15592, Stddev: 779
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 1539, Maximum: 1539, Average: 1539, Stddev: 0
  Results over all tests:
    Probes sent: 712, Probes received: 712, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 712, Minimum: 9961, Maximum: 57700, Average: 20115, Stddev: 5615
    Measurement: Round trip jitter (usec)
      Samples: 712, Minimum: 7, Maximum: 44088, Average: 4691, Stddev: 4920

Owner: mngd-deflt, Test: test-1
Target address: 10.0.1.2, Source address: 10.0.1.1, Probe type: twamp, Test size: 2
  Probe results:
    Probe response received
    Probe sent time: 08/27/20 19:50:43.572219
    Probe rcvd time: 08/27/20 19:50:43.597556, Client and server offload timestamping
    Rtt: 25299 usec, Rtt jitter: 12621 usec
    Egress jitter: 1312 usec, Ingress jitter: 11309 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 12678, Maximum: 25299, Average: 18988, Stddev: 6312
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 12621, Maximum: 12621, Average: 12621, Stddev: 0
  Results over last test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 13690, Maximum: 13953, Average: 13821, Stddev: 176
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 263, Maximum: 263, Average: 263, Stddev: 0
  Results over all tests:
    Probes sent: 714, Probes received: 714, Loss percentage: 0
    Measurement: Round trip time (usec)
```

```
      Samples: 714, Minimum: 8496, Maximum: 69366, Average: 18243, Stddev: 5540
    Measurement: Round trip jitter (usec)
      Samples: 714, Minimum: 10, Maximum: 57276, Average: 5810, Stddev: 5782

Owner: mngd-deflt, Test: test-2
Target address: 192.168.33.33, Source address: 192.168.11.11, Probe type: twamp, Test size: 2
  Probe results:
    Probe response received
    Probe sent time: 08/27/20 19:50:43.572124
    Probe rcvd time: 08/27/20 19:50:43.597511, Client and server offload timestamping
    Rtt: 25338 usec, Rtt jitter: 11895 usec
    Egress jitter: 75 usec, Ingress jitter: 11820 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 13443, Maximum: 25338, Average: 19390, Stddev: 5949
  Results over last test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 12612, Maximum: 15279, Average: 13945, Stddev: 1338
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 2667, Maximum: 2667, Average: 2667, Stddev: 0
  Results over all tests:
    Probes sent: 714, Probes received: 714, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 714, Minimum: 8403, Maximum: 69232, Average: 18341, Stddev: 5290
    Measurement: Round trip jitter (usec)
      Samples: 714, Minimum: 0, Maximum: 56810, Average: 4766, Stddev: 5358
...
```

show services monitoring twamp client probe-results control-connection test-session

```
user@host> show services monitoring twamp client probe-results control-connection mngd-deflt
test-session test-2

Owner: mngd-deflt, Test: test-2
Target address: 192.168.33.33, Source address: 192.168.11.11, Probe type: twamp, Test size: 2
  Probe results:
    Probe response received
    Probe sent time: 08/27/20 19:50:43.572124
    Probe rcvd time: 08/27/20 19:50:43.597511, Client and server offload timestamping
```

```
    Rtt: 25338 usec, Rtt jitter: 11895 usec
    Egress jitter: 75 usec, Ingress jitter: 11820 usec
  Results over current test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 13443, Maximum: 25338, Average: 19390, Stddev: 5949
  Results over last test:
    Probes sent: 2, Probes received: 2, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 2, Minimum: 12612, Maximum: 15279, Average: 13945, Stddev: 1338
    Measurement: Round trip jitter (usec)
      Samples: 2, Minimum: 2667, Maximum: 2667, Average: 2667, Stddev: 0
  Results over all tests:
    Probes sent: 714, Probes received: 714, Loss percentage: 0
    Measurement: Round trip time (usec)
      Samples: 714, Minimum: 8403, Maximum: 69232, Average: 18341, Stddev: 5290
    Measurement: Round trip jitter (usec)
      Samples: 714, Minimum: 0, Maximum: 56810, Average: 4766, Stddev: 5358
```

Release Information

Command introduced in Junos OS Evolved 20.3R1.

Command output fields for inferred egress loss, egress delay and jitter, and ingress delay and jitter introduced in Junos OS Evolved 23.4R1.

RELATED DOCUMENTATION

control-connection (Junos OS Evolved) | **4627**

Understanding Two-Way Active Measurement Protocol on Routers

show services monitoring twamp client test-info

IN THIS SECTION

- Syntax | **25826**

- Description | 25826
- Options | 25826
- Required Privilege Level | 25827
- Output Fields | 25827
- Sample Output | 25827
- Release Information | 25828

Syntax

```
show services monitoring twamp client test-info
<control-connection control-connection-name>
<test-session test-session-name>
```

Description

Display information about the test sessions established between the Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, information for all established control connections and test sessions are displayed, unless you specify the control-connection option when you issue the command. The control-connection option is required when you want to issue the command with one of the other options. For example, to display the information for a particular test session, you must also include the control-connection option to specify which connection contains the test session: **show services monitoring twamp client test-info control-connection *control-connection-name* test-session *test-session-name***.

Options

control-connection *control-connection-name*	(Optional, but required if specifying the test-session option) Display information about the test sessions on the specified control connection, which is established for control packets exchanged between a TWAMP client and a TWAMP server.
test-session *test-session-name*	(Optional) Display information about the specified test session, which is established for data packets transmitted between a TWAMP client and a TWAMP server, and the control connection associated with the test session.

Required Privilege Level

view

Output Fields

Table 1455 on page 25827 lists the output fields for the `show services monitoring twamp client test-info` command. Output fields are listed in the approximate order in which they appear.

Table 1455: show services monitoring twamp client test-info Output Fields

Field Name	Field Description	Level of Output
`Control name`	Name that uniquely identifies the control connection between the TWAMP server and a particular client	All
`Test name`	Name that uniquely identifies the control connection between the TWAMP server and a particular client	
`Client address:port`	Client IP address and port number	All
`Server address:port`	Server IP address and port number	All
`Test status`	Status of the test session; can be active or idle.	All

Sample Output

show services monitoring twamp client test-info

```
user@host> show services monitoring twamp client test-info

Control                Test                   Client                 Server                 Test
name                   name                   address:port           address:port           status
light-deflt            test-5                 10.0.1.1:51274         10.0.1.2:61000         active
light-deflt            test-6                 192.168.11.11:35197    192.168.33.33:61000    active
light-inst             test-7                 10.2.1.1:45441         10.0.1.2:61001         active
light-inst             test-8                 10.0.11.11:43598       10.0.33.33:61001       active
mngd-deflt             test-1                 10.0.1.1:50508         10.0.1.2:10015         idle
```

```
mngd-deflt          test-2              192.168.11.11:44246   192.168.33.33:10017   idle
mngd-inst           test-3              10.2.1.1:42503        10.2.1.2:10016        idle
mngd-inst           test-4              10.0.11.11:39008      10.0.33.33:10018      idle
```

show services monitoring twamp client test-info control-connection test-session

```
user@host> show services monitoring twamp client test-info control-connection light-deflt  test-
session test-6

Control             Test                Client              Server              Test
name                name                address:port        address:port        status
light-deflt         test-6              192.168.11.11:35197 192.168.33.33:61000 active
```

Release Information

Command introduced in Junos Evolved 20.3R1 .

RELATED DOCUMENTATION

control-connection (Junos OS Evolved) | 4627

Understanding Two-Way Active Measurement Protocol on Routers

show services monitoring twamp server control-info

IN THIS SECTION

- Syntax | 25829
- Description | 25829
- Options | 25829
- Required Privilege Level | 25829
- Output Fields | 25829
- Sample Output | 25830
- Release Information | 25830

Syntax

```
show services monitoring twamp server control-info
<control-connection control-identifier>
```

Description

Display information about the control connections established between the Two-Way Active Measurement Protocol (TWAMP) server and control-clients for managed servers. By default, all established control connections are displayed, unless you specify the control-connection option when you issue the command. Because TWAMP light servers are stateless, information about them is not included in the output of this command; only information about managed servers is included.

Options

control-connection ***control-identifier***	(Optional) Specify the numeric identifier for a control connection to display information about that control connection.

Required Privilege Level

view

Output Fields

Table 1456 on page 25829 lists the output fields for the `show services monitoring twamp server control-info` command. Output fields are listed in the approximate order in which they appear.

Table 1456: show services monitoring twamp server control-info Output Fields

Field Name	Field Description	Level of Output
`Control identifier`	Numeric identifier that uniquely identifies the session between the TWAMP server and a particular client	All
`Client address:port`	Client IP address and port number	All
`Server address:port`	Server IP address and port number	All

Table 1456: show services monitoring twamp server control-info Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Active tests`	Number of active tests	All
`Control status`	Status of the control connection; can be sleeping or testing	All
`Auth mode`	Authentication mode	All

Sample Output

show services monitoring twamp server control-info

```
user@host> show services monitoring twamp server control-info

Control                 Client                  Server                  Active   Control       Auth
identifier              address:port            address:port            tests    status        mode
143                     10.0.3.3:36289          10.0.1.1:862            2        testing       none
146                     10.1.1.2:40365          10.1.1.1:862            2        testing       none
149                     10.5.1.2:49613          10.5.1.1:862            0        sleeping      none
152                     10.2.1.2:56958          10.0.1.1:862            2        testing       none
```

show services monitoring twamp server control-info control-connection

```
user@host> show services monitoring twamp server control-info control-connection 143

Control                 Client                  Server                  Active   Control       Auth
identifier              address:port            address:port            tests    status        mode
143                     10.0.3.3:36289          10.0.1.1:862            2        testing       none
```

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

show services monitoring twamp server test-info

IN THIS SECTION

- Syntax | 25831
- Description | 25831
- Options | 25832
- Required Privilege Level | 25832
- Output Fields | 25832
- Sample Output | 25833
- Release Information | 25833

Syntax

```
show services monitoring twamp server test-info
<control-connection control-identifier>
<test-session test-identifier>
```

Description

Display information about the test sessions established between the Two-Way Active Measurement Protocol (TWAMP) server and control-clients. By default, all established test sessions are displayed, unless you specify the control-connection option when you issue the command. The control-connection option is required when you want to issue the command with the test-session option. For example, to display the information for a particular test session, you must also include the control-connection option to specify which connection contains the test session: **show services monitoring twamp server test-info control-connection *control-identifier* test-session *test-identifier***.

Options

control-connection *control-identifier* — (Optional, but required if specifying the test-session option) Specify a numeric identifier to display information about the specified control connection, which is established for control packets exchanged between a TWAMP client and a TWAMP server.

test-session *test-identifier* — (Optional) Specify a numeric identifier to display information about a specific test session.

Required Privilege Level

view

Output Fields

Table 1457 on page 25832 lists the output fields for the `show services monitoring twamp server test-info` command. Output fields are listed in the approximate order in which they appear.

Table 1457: show services monitoring twamp server test-info Output Fields

Field Name	Field Description	Level of Output
`Control identifier`	Identification number that uniquely identifies the control session between the TWAMP server and a particular client	All
`Test identifier`	Identification number that uniquely identifies the test session between the TWAMP server and a particular client	All
`Client address:port`	Client IP address and port number	All
`Server address:port`	Server IP address and port number	All
`Test status`	Status of the test session; can be active or idle	All
`Auth mode`	Authentication mode	All

Sample Output

show services monitoring twamp server test-info

```
user@host> show services monitoring twamp server test-info

Control                 Test                    Client                  Server                  Test
identifier              identifier              address:port            address:port            status
149                     7045                    10.5.1.2:13474          192.168.1.1:13474       active
149                     7046                    10.5.1.2:13475          10.5.1.1:13475          active
152                     7043                    10.2.1.2:13472          10.0.1.1:13472          active
152                     7044                    10.2.1.2:13473          10.2.1.1:13473          active
```

show services monitoring twamp server test-info control-connection

```
user@host> show services monitoring twamp server test-info control-connection 149

Control                 Test                    Client                  Server                  Test
identifier              identifier              address:port            address:port            status
149                     7045                    10.5.1.2:13474          192.168.1.1:13474       active
149                     7046                    10.5.1.2:13475          10.5.1.1:13475          active
```

show services monitoring twamp server test-info control-connection test-session

```
user@host> show services monitoring twamp server test-info control-connection 149 test-session
7045

Control                 Test                    Client                  Server                  Test
identifier              identifier              address:port            address:port            status
149                     7045                    10.5.1.2:13474          192.168.1.1:13474       active
```

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

show services nat destination pool

IN THIS SECTION

- Syntax | 25834
- Description | 25834
- Options | 25834
- Required Privilege Level | 25835
- Output Fields | 25835
- Sample Output | 25835
- Release Information | 25836

Syntax

```
show services nat destination pool
<interface interface-name>
<service-set service set>
<all>
```

Description

Display destination NAT address pool information.

Options

interface *interface-name*>	Optional. Display destination NAT information specific to the interface.
service-set *service-set*>	Optional. Display destination NAT information specific to the service set.
all	Optional. Display all destination NAT address pool information.

Required Privilege Level

view

Output Fields

Table 1458 on page 25835lists the output fields for the `show services nat destination pool` command. Output fields are listed in the approximate order in which they appear.

Table 1458: show services nat destination pool Output Fields

Field Name	Description
`Interface`	Interface name.
`Service set`	Service set name.
`Pool name`	Pool name.
`Pool id`	Pool identification.
`Total address`	Number of IP addresses that are in use.
`Translation hits`	Number of times a translation in the translation table is used for a source NAT rule.
`Address range`	IP address range in the source pool.
`Port`	Port number used to access the pool.

Sample Output

show services nat destination pool

```
user@host> show services nat destination pool service-set ss1_interface_style1 interface
vms-0/2/0 all
```

```
ss1_interface_style1 interface vms-0/2/0 all | no-more
Interface: vms-0/2/0 , Service set: ss1_interface_style1
Pool name       : dest_pool
Pool id         : 1
Total address   : 253
Translation hits: 11
Address range                        Port
       30.1.1.2 - 30.1.1.254            0
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services nat destination rule

IN THIS SECTION

- Syntax | 25836
- Description | 25837
- Options | 25837
- Required Privilege Level | 25837
- Output Fields | 25837
- Sample Output | 25838
- Release Information | 25839

Syntax

```
show services nat destination rule
rule-name
<service-set service-set>
<interface interface-name>
<all>
```

Description

Display destination NAT rule-set information.

Options

rule-name	Display information about the specified destination NAT rule.
service-set *service-set*	Display information specific to the service-set.
interface *interface-name*	Display information specific to the interface.
all	Display all NAT rule-set information.

Required Privilege Level

view

Output Fields

Table 1459 on page 25837 lists the output fields for the `show services nat destination rule` command. Output fields are listed in the approximate order in which they appear.

Table 1459: show services nat destination rule Output Fields

Field Name	Description
`Interface`	Interface name.
`Service set`	Service set name.
`Destination NAT rule`	Name of the destination NAT rule.
`Rule-Id`	Rule identification number.
`Rule-position`	Position of the destination NAT rule.

Table 1459: show services nat destination rule Output Fields *(Continued)*

Field Name	Description
`Match-direction`	Three options: • input—Apply the rule match on the input side of the interface. • input-output—Apply the rule match bidirectionally. • output—Apply the rule match on the output side of the interface.
`Destination addresses`	Name of the destination addresses that match the rule. The default value is any.
`Action`	The action taken when a packet matches the rule's tuples. Actions include the following: • destination NAT pool—Use user-defined destination NAT pool to perform destination NAT. • off—Do not perform destination NAT.
`Translation hits`	Number of times a translation in the translation table is used for a source NAT rule.
`Successful sessions`	Number of successful session installations after the NAT rule is matched.
`Failed sessions`	Number of unsuccessful session installations after the NAT rule is matched.
`Number of sessions`	Number of sessions that reference the specified rule.

Sample Output

show services nat destination rule service-set ss1_interface_style1 interface vms-0/2/0 all | no-more

```
user@host> show services nat destination rule service-set ss1_interface_style1 interface
vms-0/2/0 all | no-more
ss1_interface_style1 interface vms-0/2/0 all | no-more
```

```
Interface: vms-0/2/0 , Service set: ss1_interface_style1
Destination NAT rule: r1                      Rule-set: rs2
  Rule-Id                    : 2
  Rule position              : 1
  Match-direction            : input
    Destination addresses    : 50.1.1.2        - 50.1.1.2
  Action                     : dest_pool
  Translation hits           : 34
    Successful sessions      : 34
    Failed sessions          : 0
  Number of sessions         : 0
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services nat destination summary

IN THIS SECTION

- Syntax | 25840
- Description | 25840
- Options | 25840
- Required Privilege Level | 25840
- Output Fields | 25840
- Sample Output | 25842
- Release Information | 25842

Syntax

```
show services nat destination summary
<interface interface-name>
<service-set service-set>
```

Description

Display summary destination NAT information.

Options

interface *interface-name* Display summary destination NAT information for the specified service interface.

service-set *service-set* Display summary destination NAT information for the specified service set.

Required Privilege Level

view

Output Fields

Table 1460 on page 25840 lists the output fields for the `show services nat destination summary` command. Output fields are listed in the approximate order in which they appear.

Table 1460: show services nat destination summary Output Fields

Field Name	Description
`Interface`	Interface name.
`Service set`	Service set name.
`Pool name`	Name of the destination address pool.

Table 1460: show services nat destination summary Output Fields *(Continued)*

Field Name	Description
`Address Range`	IP address or IP address range for the pool.
`Routing Instance`	Name of the routing instance.
`Port`	Port number.
`Total Address`	Number of IP addresses that are in use.
`Rule name`	Rule name.
`Rule set`	The set of rules for destination NAT.
`Match-direction`	Three options: • input—Apply the rule match on the input side of the interface. • input-output—Apply the rule match bidirectionally. • output—Apply the rule match on the output side of the interface.
`Action`	The action taken when a packet matches the rule's tuples. Actions include the following: • destination NAT pool—Use user-defined destination NAT pool to perform destination NAT. • off—Do not perform destination NAT.

Sample Output

show services nat destination summary service-set ss1_interface_style1 interface vms-0/2/0

```
user@host> show services nat destination summary service-set ss1_interface_style1 interface
vms-0/2/0
Interface: vms-0/2/0 , Service set: ss1_interface_style1
Pool name            Address                           Routing        Port  Total
                     Range                             Instance             Address
dest_pool            30.1.1.2       - 30.1.1.254                      0     253
Interface: vms-0/2/0 , Service set: ss1_interface_style1
Rule name    Rule set                        Match-direction   Action
r1            rs2                              input         dest_pool
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services nat deterministic-nat internal-host

IN THIS SECTION

- Syntax | 25843
- Description | 25843
- Options | 25843
- Required Privilege Level | 25843
- Output Fields | 25843
- Sample Output | 25844
- Release Information | 25844

Syntax

```
show services nat deterministic-nat internal-host
nat-address
nat-port
```

Description

This commands prints the internal host address and algorithmically determined port ranges for the specified NAT IP address and port number. The results are calculated on the PIC and the results are sent to RE.

Options

nat-address	NAT address of the internal host.
nat-port	NAT port of the internal host.

Required Privilege Level

view

Output Fields

Table 1461 on page 25843 lists the output fields for the `nat deterministic-nat internal-host` command. Output fields are listed in the approximate order in which they appear.

Table 1461: show services nat deterministic-nat internal-host Output Fields

Field Name	Field Description
`Interface`	Name of a service interface.
`Service set`	Name of a service set. Individual empty service sets are not displayed, but if none of the service sets has any flows, a flow table header is printed for each service set.
`Internal Host`	Private IP address of a subscriber on the access network.

Table 1461: show services nat deterministic-nat internal-host Output Fields *(Continued)*

Field Name	Field Description
NAT IP address	NAT public IP address
NAT Port Start	Lowest port number in range of assigned ports.
NAT Port End	Highest port number in range of assigned ports.

Sample Output

show services nat deterministic-nat internal-host

```
user@host> show services nat deterministic-nat internal-host 203.0.113.1 2000
Service set: ss1
Interface: sp-2/0/0
NAT pool: pool1
Internal Host: 192.0.2.4, NAT IP Address: 203.0.113.1, NAT Port Start: 1792, NAT Port End: 2047
```

Release Information

Command introduced in Junos OS Release 12.1.

show services nat deterministic-nat nat-port-block

IN THIS SECTION

- Syntax | 25845
- Description | 25845
- Options | 25845

- Required Privilege Level | 25845
- Output Fields | 25845
- Sample Output | 25846
- Release Information | 25846

Syntax

```
show services nat deterministic-nat nat-port-block
internal-host
```

Description

Display the translated NAT address and port ranges for the given internal host.

Options

internal-host	IP address of the internal host.

Required Privilege Level

view

Output Fields

Table 1462 on page 25845 lists the output fields for the `show services nat deterministic-nat nat-port-block` command. Output fields are listed in the approximate order in which they appear.

Table 1462: show services nat deterministic-nat nat-port-block Output Fields

Field Name	Field Description
`Interface`	Name of a service interface.

Table 1462: show services nat deterministic-nat nat-port-block Output Fields *(Continued)*

Field Name	Field Description
`Service set`	Name of a service set. Individual empty service sets are not displayed, but if none of the service sets has any flows, a flow table header is printed for each service set.
`Internal Host`	Private IP address of a subscriber on the access network.
`NAT IP address`	NAT public IP address
`NAT Port Start`	Lowest port number in range of assigned ports.
`NAT Port End`	Highest port number in range of assigned ports.

Sample Output

run show services nat deterministic-nat nat-port-block

```
user@host> show services nat deterministic-nat nat-port-block 192.0.2.1
Service set: ss1
Interface: sp-2/0/0
NAT pool: pool1
Internal Host: 192.0.2.1, NAT IP Address: 203.0.113.1, NAT Port Start: 1024, NAT Port End: 1279
```

Release Information

Command introduced in Junos OS Release 12.1.

show services nat ipv6-multicast-interfaces

IN THIS SECTION

- Syntax | 25847
- Description | 25847
- Required Privilege Level | 25847
- Output Fields | 25847
- Sample Output | 25848
- Release Information | 25850

Syntax

```
show services nat ipv6-multicast-interfaces
```

Description

Displays a list of interfaces enabled for IPv6 multicast.

Required Privilege Level

view

Output Fields

Table 1463 on page 25847 lists the output fields for the `show services nat ipv6-multicast-interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 1463: show services nat ipv6-multicast-interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a service interface.	All levels

Table 1463: show services nat ipv6-multicast-interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Admin State	Configured IPv6 multicast capability of an interface ,	All levels
Operational State	Operation IPv6 multicast status of an interface.	All levels

Sample Output

show services nat ipv6-multicast-interfaces

```
user@host> show services nat ipv6-multicast-interfaces
Interface         Admin         Operational
                  State            State
ge-5/1/9          Enabled       Enabled
ge-5/1/8          Enabled       Enabled
ge-5/1/7          Enabled       Enabled
ge-5/1/6          Enabled       Enabled
ge-5/1/5          Enabled       Enabled
ge-5/1/4          Enabled       Enabled
ge-5/1/3          Enabled       Enabled
ge-5/1/2          Enabled       Enabled
ge-5/1/1          Enabled       Enabled
ge-5/1/0          Enabled       Enabled
ge-5/0/9          Enabled       Enabled
ge-5/0/8          Enabled       Enabled
ge-5/0/7          Enabled       Enabled
ge-5/0/6          Enabled       Enabled
ge-5/0/5          Enabled       Enabled
ge-5/0/4          Enabled       Enabled
ge-5/0/3          Enabled       Enabled
ge-5/0/2          Enabled       Enabled
ge-5/0/1          Enabled       Enabled
ge-5/0/0          Enabled       Enabled
ge-1/3/9          Enabled       Enabled
ge-1/3/8          Enabled       Enabled
ge-1/3/7          Enabled       Enabled
```

```
ge-1/3/6          Enabled       Enabled
ge-1/3/5          Enabled       Enabled
ge-1/3/4          Enabled       Enabled
ge-1/3/3          Enabled       Enabled
ge-1/3/2          Enabled       Enabled
ge-1/3/1          Enabled       Enabled
ge-1/3/0          Enabled       Enabled
ge-1/2/9          Enabled       Enabled
ge-1/2/8          Enabled       Enabled
ge-1/2/7          Enabled       Enabled
ge-1/2/6          Enabled       Enabled
ge-1/2/5          Enabled       Enabled
ge-1/2/4          Enabled       Enabled
ge-1/2/3          Enabled       Enabled
ge-1/2/2          Enabled       Enabled
ge-1/2/1          Enabled       Enabled
ge-1/2/0          Enabled       Enabled
ge-1/1/9          Enabled       Enabled
ge-1/1/8          Enabled       Enabled
ge-1/1/7          Enabled       Enabled
ge-1/1/6          Enabled       Enabled
ge-1/1/5          Enabled       Enabled
ge-1/1/4          Enabled       Enabled
ge-1/1/3          Enabled       Enabled
ge-1/1/2          Enabled       Enabled
ge-1/1/1          Enabled       Enabled
ge-1/1/0          Enabled       Enabled
ge-1/0/9          Enabled       Enabled
ge-1/0/8          Enabled       Enabled
ge-1/0/7          Enabled       Enabled
ge-1/0/6          Enabled       Enabled
ge-1/0/5          Enabled       Enabled
ge-1/0/4          Enabled       Enabled
ge-1/0/3          Enabled       Enabled
ge-1/0/2          Enabled       Enabled
ge-1/0/1          Enabled       Enabled
ge-1/0/0          Enabled       Enabled
xe-0/3/0          Enabled       Enabled
xe-0/2/0          Enabled       Enabled
xe-0/1/0          Enabled       Enabled
xe-0/0/0          Enabled       Enabled
```

Release Information

Command introduced in Junos OS Release 8.5.

show services nat pool

IN THIS SECTION

- Syntax | 25850
- Description | 25850
- Options | 25851
- Required Privilege Level | 25851
- Output Fields | 25851
- Sample Output | 25855
- Release Information | 25858

Syntax

```
show services nat pool
<brief | detail>
<pool-name>
pgcp <ports-per-session | remotely-controlled>
```

Description

Display information about Network Address Translation (NAT) pools.

NOTE: On MS-MPCs and MS-MICs, if the line cards receive a packet immediately after the active port block timeout interval has expired, a new port block is allocated and the old port block is released thereafter (if no more ports are being used from that block). In such a scenario, you might notice that the `Max number of port blocks used` field displays a

higher value than the value shown for the `Unique pool users` field in the output of the `show services nat pool detail` command. This behavior is expected with port block allocation.

With MS-MPCs and MS-MICs, in the output of the `show services nat pool detail` command, the `Max ports used` and the `Ports in use` fields display values that indicate a higher number than the number of active subscribers on the member interfaces of an `ams` interface. This behavior of an increased value displayed for the number of ports allocated and maximum number of ports used is expected after you perform a Graceful Routing Engine switchover (GRES) and a restart of the MPC.

With MS-MPCs and MS-MICs on MX Series routers with AMS interfaces, it is observed that the subscriber and port count details are displayed only after a long time in the output of the `show services nat pool detail` command. This behavior is expected with NAT pool counters and occurs, regardless of port block allocation being configured.

Options

none	Display standard information about all NAT pools.
brief | detail	(Optional) Display the specified level of output.
pool-name	(Optional) Display information about the specified NAT pool.
pgcp	(Optional) Display information about a NAT pool that is exclusive to the BGF.
ports-per-session	(Optional) Display the number of ports allocated per session from the NAT pool.
remotely-controlled	(Optional) Display if the NAT pool is explicitly specified by the gateway controller.

Required Privilege Level

view

Output Fields

Table 1464 on page 25852 lists the output fields for the `show services nat pool` command. Output fields are listed in the approximate order in which they appear.

Table 1464: show services nat pool Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of an adaptive services interface.	All levels
`Service set`	Name of a service set. Individual empty service sets are not displayed, but if none of the service sets has any flows, a flow table header is printed for each service set.	All levels
`NAT pool`	Name of the Network Address Translation pool.	All levels
`Type` or `Translation type`	Address translation type: `basic-nat-pt`, Y, Y, Y, Y, Y, Y, Y, Y, Y, Y, Y, Y.	All levels
`Address` or `Address range`	IPv4 address range of the pool.	All levels
`Port` or `Port range`	Port range of the pool. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	All levels
`Ports used` or `Ports in use`	Number of ports allocated in this pool with this name. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	All levels
`Port block type`	Type of port block allocation: secured or deterministic	All levels
`Available addresses`	Number of free addresses in the NAT pool.	`detail`
`Configured port range`	The range of ports configured to be used for NAT pool.	`detail`
`Out of port errors`	Number of port allocation errors. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	`detail`

Table 1464: show services nat pool Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Parity port errors`	Number of port allocations that failed because a port number of the desired parity was not available.	`detail`
`Preserve Range errors`	Number of port allocations that failed because a port in the desired range was not available.	`detail`
`Max ports used`	Maximum number of ports used. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	`detail`
`Addresses in use`	Number of addresses in use for dynamic source address NAT pools.	`detail`
`AP-P port allocation errors`	When address pooling paired (AP-P) is configured, a private IP is paired to a public IP. This is a counter of translation errors where there are free ports available in the NAT pool, but none for the NAT IP to which the private IP is paired.	`detail`
`AP-P port limit allocation errors`	When AP-P is configured, this is a counter of out-of-port errors that are due to a configured limit for the number of allocated ports in the `limit-ports-per-address` statement at the [`edit services nat pool` *`nat-pool-name`*] hierarchy level.	`detail`
`Memory allocation errors`	Number of memory allocation failures.	`detail`
`EIF Inbound session count`	Current number of EIF inbound sessions.	`detail`
`EIF Inbound session Limit exceeded drops`	Number of inbound connections that were dropped because the EIF limit was exceeded.	`detail`
`Port block size`	Number of ports in a port block.	none `brief`

Table 1464: show services nat pool Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Max port blocks per address`	Maximum number of port blocks per private address.	none `brief`
`Active block timeout`	Activity timeout of port block.	none `brief`
`Effective port range`	Effective range of port numbers.	none `brief`
`Effective number of port blocks`	Effective number of port blocks.	none `brief`
`Effective number of ports`	Effective number of ports.	none `brief`
`Port block efficiency`	Port block efficiency.	none `brief`
`Port blocks limit exceeded errors`	The total number of times when a request for more than the allowed port blocks allocated for a user arrives from a user.	All levels
`Preserve range enabled`	Whether the capability to preserve the privileged port range after translation is enabled. One of the following is displayed: • `Is active`—Preservation of port range is enabled. • `Not active`—Preservation of port range is not enabled.	`detail`
`AP-P out of port errors`	When AP-P is configured, a private IP is paired to a public IP. This is a counter of translation errors where there are free ports available in the NAT pool, but none for the NAT IP to which the private IP is paired.	`detail`
`MAX number of port blocks used`	The maximum number of port blocks used.	All levels

Table 1464: show services nat pool Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Current number of port blocks in use	Current count of the port blocks that are being used.	detail
Port block allocation errors	The consolidated number of port block allocation errors.	All levels
Port block memory allocation errors	The number of memory allocation errors for port blocks.	All levels
DetNat subscriber exceeded port limits	The number of times a subscriber exceeded its port limits for a NAT pool that uses deterministic port block allocation.	All levels
Unique pool users	The number of different users of the NAT pools.	All levels
Current EIF Inbound flows count	Current count of EIF inbound flows, including all EIF flows per pool.	detail
EIF flow limit exceeded drops	Current number of flow drops due to exceeded flow limit. This number is per pool, not per EIF mapping.	detail

Sample Output

show services nat pool brief

```
user@host> show services nat pool brief

Interface: ms-1/0/0, Service set: s1
NAT pool          Type    Address                         Port        Ports used
dest-pool         DNAT-44 10.10.10.2-10.10.10.2
napt-pool         NAPT-44 50.50.50.1-50.50.50.254         1024-63487  0
```

```
source-dynamic-pool DYNAMIC NAT44 40.40.40.1-40.40.40.254
source-static-pool BASIC NAT44 30.30.30.1-30.30.30.254
```

show services nat pool detail

```
user@host> show services nat pool detail
Interface: ms-4/0/0, Service set: ss1
     NAT pool: srcpool, Translation type: NAPT-44
      Address range: 100.0.0.1-100.0.0.254
      Available addresses: 254
      Configured port range: 1024-65535
      Port range: 1024-65535, Ports in use: 0, Out of port errors: 0
      Parity port errors: 0, Preserve Range errors: 0
      Max ports used: 0
      AP-P port allocation errors: 0, AP-P port limit allocation errors: 0
      Memory allocation errors: 0
      EIF Inbound session count: 0
      EIF Inbound session Limit exceeded drops: 0
```

show services nat pool (Secured Port Block Allocation)

```
user@host> show services nat pool
```

```
Interface: sp-2/0/0, Service set: in
NAT pool          Type    Address                         Port        Ports used
mypool            dynamic 3.3.3.3-3.3.3.10                512-65535   0
                          3.3.3.15-3.3.3.20
                          3.3.3.25-3.3.3.30
                          3.3.3.95-3.3.3.200
Port block size: 64, Max port blocks per address: 1, Active block timeout: 86400, Effective port
range: 1024-65471,
Effective number of port blocks: 126882, Effective number of ports: 8120448, Port block
efficiency: nan

Interface: sp-2/1/0, Service set: in1
NAT pool          Type    Address                         Port        Ports used
mypool1           dynamic 9.9.9.1-9.9.9.254               512-65535   0
Port block size: 64, Max port blocks per address: 1, Active block timeout: 86400, Effective port
```

```
range: 1024-65471,
Effective number of port blocks: 255778, Effective number of ports: 16369792, Port block
efficiency: nan
```

show services nat pool detail (Deterministic Port Block Allocation)

```
user@host> show services nat pool detail
Interface: sp-2/0/0, Service set: ss1
  NAT pool: napt_pool, Translation type: dynamic
    Address range: 5.5.5.1-5.5.5.254
    Configured port range: 1-60000, Preserve range enabled: Is active
    Port range: 2000-2002, Ports in use: 2, Out of port errors: 0, Max ports used: 2
    AP-P out of port errors: 188
    Max number of port blocks used: 1, Current number of port blocks in use: 1, Port block
allocation errors: 0,
    Port block memory allocation errors: 0
    DetNAT subscriber exceeded port limits: 1
    Unique pool users: 1
```

show services nat pool (Deterministic Port Block Allocation)

```
user@host> show services nat pool

Interface: sp-2/0/0, Service set: ss2
NAT pool          Type    Address                         Port        Ports Used
pba               dynamic 33.33.33.1-33.33.33.128         512-65535   6604
Port block type: Deterministic port block, Port block size: 200
```

show services nat pool detail (Port Block Allocation)

```
user@host> show services nat pool detail

Interface: sp-2/0/0, Service set: s
  NAT pool: napt_pool, Translation type: dynamic
    Address range: 44.1.1.1-44.1.1.1
    Configured port range: 1-60000
    Port range: 1024-65535, Ports in use: 0, Out of port errors: 0,
    Max ports used: 0
```

```
    AP-P out-of-port errors: 0
    Current EIF Inbound flows count: 0
    EIF flow limit exceeded drops: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

pgcp option added in Junos OS Release 8.5.

show services nat resource-usage source-pool

IN THIS SECTION

- Syntax | 25858
- Description | 25858
- Options | 25859
- Required Privilege Level | 25859
- Output Fields | 25859
- Sample Output | 25860
- Release Information | 25860

Syntax

```
show services nat resource-usage source-pool
<all>
pool-name
```

Description

Display NAT resource usage.

Options

<all>	Display all NAT resource usage statistics.
pool-name	Display NAT resource usage statistics for the specified pool.

Required Privilege Level

view

Output Fields

Table 1465 on page 25859 lists the output fields for the `show services nat resource-usage` command. Output fields are listed in the approximate order in which they appear.

Table 1465: show services nat resource-usage Output Fields

Field Name	Description
`Pool`	Name of the pool.
`Address`	Address of the pool.
`Used`	Number of used resources in the pool.
`Available`	Number of available resources in the pool.
`Total`	Total number of addresses in the pool.
`Usage`	Percent of resources used.

Sample Output

show services nat resource-usage source-pool all

```
user@host> show services nat resource-usage source-pool all
PAT pools(including address-shared pool) port utilization:
  Pool                                Address      Used     Avail     Total Usage
  src-nat-pool-1                            1        64         0        64  100%
  src-nat-pool-2                            4         0    258048    258048    0%
```

show services nat resource-usage source-pool src-nat-pool-2

```
show services nat resource-usage source-pool src-nat-pool-2
Pool name: src-nat-pool-2
Total address: 4
Port-overloading-factor: 1
Total ports: 258048 Used: 0  Avail: 258048
Current usage: 0%   Peak usage: 0% at 1970-01-01 00:00:00 UTC
  Address    Factor-index Port-range        Used     Avail     Total Usage
  1.1.1.20
             0            Single Ports         0     64512     64512    0%
  1.1.1.21
             0            Single Ports         0     64512     64512    0%
  1.1.1.22
             0            Single Ports         0     64512     64512    0%
  1.1.1.23
             0            Single Ports         0     64512     64512    0%
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services nat source deterministic

IN THIS SECTION

- Syntax | 25861
- Description | 25861
- Options | 25861
- Required Privilege Level | 25862
- Output Fields | 25862
- Sample Output | 25863
- Release Information | 25863

Syntax

```
show services source nat deterministic
host-address-range
host-ip ip-address
pool pool-name
xlated-ip translated-ip-address
xlated-port translated-port-number
```

Description

Display deterministic port block allocation information.

Options

host-address-range	Display the deterministic host address range without overlap.
host-ip *ip-address*	Display the internal host IP address.
pool *pool-name*	Display the source NAT pool.
xlated-ip *translated-ip-address*	Display translated IP address.

xlated-port ***translated-port-number*** Display the translated port number.

Required Privilege Level

view

Output Fields

Table 1466 on page 25862 lists the output fields for the command. Output fields are listed in the approximate order in which they appear.

Table 1466: show services nat source deterministic Output Fields

Field Name	Field Description
`Pool name`	Name of the NAT source pool.
`Port overloading factor`	Factor of port overloading for the source pool.
`Used/total port blocks`	Port block used number and port block total number for this source NAT pool.
`Host IP`	Host IP address.
`External IP`	IP address of external router.
`Port Block Range`	The range of ports in a block, ranging from lowest to highest.
`Ports Used/Ports Total`	Number of ports used and total ports.
`Total host ranges number`	Host ranges in total.
`Min Host Address`	Minimum host address.

Table 1466: show services nat source deterministic Output Fields *(Continued)*

Field Name	Field Description
Max Host Address	Maximum host address.

Sample Output

show services nat source deterministic

```
user@host> show services nat source deterministic

Pool name: src-nat-pool-1
Port-overloading-factor: 1 Port block size: 256
Used/total port blocks: 0/12
Host_IP External_IP Port_Block Ports_Used/
                           Range           Ports_Total
10.1.1.1 202.0.0.1     1280-1535             0/256*1
10.1.1.2 202.0.0.1     1536-1791             0/256*1
```

show services nat source deterministic host-address-range

```
user@host> show services nat source deterministic host-address-range
Pool name: src-nat-pool-1
Total host ranges number: 1
Min Host Address Max Host Address
10.1.1.1 10.1.1.2
```

Release Information

This command was introduced in Junos OS 19.3R2.

show services nat source mappings address-pooling-paired

IN THIS SECTION

- Syntax | 25864
- Description | 25864
- Options | 25864
- Required Privilege Level | 25865
- Sample Output | 25865
- Release Information | 25868

Syntax

```
show services nat source mappings address-pooling-paired
```

Description

Displays NAT source mappings address pooling information

Options

address-pooling-paired	(Optional) Display only information about address-pooling paired mappings.
endpoint-independent	(Optional) Display only information about endpoint-independent mappings.
pcp	(Optional) Display only information about port control protocol mappings.

NOTE: PCP requests with the prefer-failure option request a particular external IP address and port. When the request cannot be fulfilled, the mapping is not created. In this case, the subscriber does not have a mapped IP address. Such a subscriber is counted in the summary of the number or address mappings, but is not displayed in the list of address mappings, as shown in the following examples:

```
user@host# show services nat mappings summary
Service Interface:                                  sp-2/0/0
Total number of address mappings:                          1
Total number of endpoint independent port mappings:        0
Total number of endpoint independent filters:              0


user@host# show services nat mappings address-pooling-paired
[edit]
```

This is expected behavior because unfulfilled address mappings (IP of 0.0.0.0) are not displayed in the output of the second CLI command. These address mappings will time out based on configured or default values.

Required Privilege Level

view

Sample Output

show services nat source mappings address-pooling-paired

```
user@host> show services nat source mappings address-pooling-paired
Interface: ms-2/0/0, Service set: ss1
Pool name: sp1
Internal address          External address     Session Count       Mapping State
1.1.1.100                     30.30.30.1           1                     Active
1.1.1.101                     30.30.30.2           1                     Active
```

show services nat source mappings address-pooling-paired private 1.1.1.100

```
user@host> show services nat source mappings address-pooling-paired private
1.1.1.100
Interface: ms-2/0/0, Service set: ss1
Pool name: sp1
Internal address          External address     Session Count       Mapping State
1.1.1.100                     30.30.30.1            1                  Active
```

show services nat source mappings address-pooling-paired public 30.30.30.2

```
user@host> show services nat source mappings address-pooling-paired public
30.30.30.2
Interface: ms-2/0/0, Service set: ss1
Pool name: sp1
Internal address          External address     Session Count       Mapping State
1.1.1.101                     30.30.30.2            1                  Active
```

show services nat source mappings address-pooling-paired pool-name sp1

```
user@host> show services nat source mappings address-pooling-paired pool-name sp1
  Interface: ms-2/0/0, Service set: ss1
Pool name: sp1
Internal address          External address     Session Count       Mapping State
1.1.1.100                     30.30.30.1             1                  Active
1.1.1.101                     30.30.30.2             1                   Active
```

show services nat mappings address-pooling-paired

```
user@host> show services nat mappings address-pooling-paired
Interface: sp-3/0/0, Service set: NAPT44-SS1
NAT pool: napt44-SS1-p1
Mapping          : 29.32.38.255    --> 192.168.75.23
Ports In Use     :     9
Session Count    :     1
Mapping State    : Active
```

show services nat mappings address-pooling-paired (mapping of active B4 for a subscriber)

```
user@host> show services nat mappings address-pooling-paired
Interface: sp-0/0/0, Service set: sset_1

NAT pool: nat_pool1

Mapping          : 2001::          --> 33.33.33.2
Ports In Use     :     1
```

```
Session Count    :     9
Mapping State    : Timeout
```

show services nat mappings endpoint-independent

```
user@host> show services nat mappings endpoint-independent
Interface: sp-3/0/0, Service set: NAPT44-SS1
NAT pool: napt44-SS1-p1
Mapping          : 29.32.38.255:10000    --> 192.168.75.23:1024
Session Count    : 1
Mapping State    : Active
```

show services nat mappings pcp

```
user@host> show services nat mappings pcp
PCP Client       : 172.16.0.1        PCP Lifetime : 45
Mapping          : 29.32.38.255:10000    --> 192.168.75.23:1024
Session Count    : 1
Mapping State    : Active
```

show services nat mappings nptv6 internal

```
user@host> show services nat mappings nptv6 internal 1111:2222:3333:aaaa:bbbb::1

    Interface       Service-set     NAT-Pool           Address Mapping
    vms-0/1/0       ss_nptv6       ss_nptv6_pool   1111:2222:3333:aaaa:bbbb::1 ->
aaaa:bbbb:cccc:dddd:bbbb::1
```

show services nat mappings nptv6 external

```
user@host> show services nat mappings nptv6 external aaaa:bbbb:cccc:dddd:bbbb::1

    Interface       Service-set     NAT-Pool              Address Mapping
    vms-0/1/0        ss_nptv6            ss_nptv6_pool   1111:2222:3333:aaaa:bbbb::1 ->
aaaa:bbbb:cccc:dddd:bbbb::1
```

Release Information

show services nat source mappings endpoint-independent

IN THIS SECTION

- Syntax | 25868
- Description | 25868
- Options | 25868
- Required Privilege Level | 25869
- Output Fields | 25869
- Sample Output | 25870
- Sample Output | 25871
- Release Information | 25871

Syntax

```
show services nat source mappings endpoint-independent
<pool-name>
<private | public>
```

Description

Displays NAT endpoint independent mapping.

Options

<pool-name>	Name of address pool.
<private>	Private IPv4/IPv6 prefix to use as a filter.
<public>	Public IP prefix to use as a filter.

Required Privilege Level

view

Output Fields

Table 1467 on page 25869 lists the output fields from the `show services nat source mappings endpoint-independent` command. Output fields are listed in the approximate order in which they appear.

Table 1467: show services nat source mappings endpoint-independent Output Fields

Field Name	Description
`Interface`	Name of the interface.
`Service set`	Name of the service set.
`NAT pool`	Name of the NAT pool.
`Mapping`	Shows the mapping of IP addresses.
`Session Count`	Number of sessions currently using the mapping.
`Mapping State`	NAT mapping state. The following states are possible: • ACTIVE—Indicates that the entry is active and in use. • TIMEOUT—Indicates that the mapping is not is in use. After the mapping-timeout, configured at the `[edit services nat pool pool-name]` hierarchy level, lapses, the mapping is deleted. This fields also displays the number of seconds after which the timeout occurs.

Sample Output

show services nat source mappings endpoint-independent (ms- interfaces)

```
user@host> show services nat source mappings endpoint independent
Interface: ms-2/0/0, Service set: ss1
NAT pool: test-pool
Mapping          : 2.1.1.1         : 1026  --> 123.0.0.5       :10926
Session Count    :     1
Mapping State    : Active
```

show services nat source mappings endpoint-independent private 15.4.4.2 public 20.20.20.1 (ms-interfaces)

```
user@host> show services nat source mappings endpoint-independent private 15.4.4.2 public
20.20.20.1
Interface: ms-2/0/0, Service set: ss1
NAT pool: p1
Mapping          : 15.4.4.2        :12841  --> 20.20.20.1      :11205
Session Count    :     1
Mapping State    : Active
```

show services nat source mappings endpoint-independent pool-name p1 (ms-interfaces)

```
user@host> show services nat source mappings endpoint-independent pool-name p1
Interface: ms-2/0/0, Service set: ss1
NAT pool: p1
Mapping          : 15.4.4.2        :12841  --> 20.20.20.1      :11205
Session Count    :     1
Mapping State    : Active
```

show services nat source mappings address-pooling-paired pool-name sp1 (sp- interfaces)

```
user@host> show services nat source mappings address-pooling-paired pool-name sp1
  Interface: ms-2/0/0, Service set: ss1
Pool name: sp1
Internal address          External address     Session Count       Mapping State
```

```
1.1.1.100                    30.30.30.1            1                 Active
1.1.1.101                    30.30.30.2            1                  Active
```

Sample Output

show services nat source mappings endpoint-independent (vms- interfaces)

```
user@host> show services nat source mappings endpoint-independent
Interface: vms-2/0/0, Service set: vms-sset10
Pool name: napt44-pool12

Mapping          : 20.1.0.101     : 1024  --> 50.0.12.1       : 1024
Session Count    :     1
Mapping State    : Active
B4 Address       : 2002:2010::1401:4 >>>>>>>> B4 Address in mapping
```

Release Information

Command introduced in Junos OS 19.3R2.

Support for Next Gen Services with the MX-SPC3 security services card added in Junos OS Release 20.2.

show services nat source mappings pcp

IN THIS SECTION

- Syntax | 25872
- Description | 25872
- Options | 25872
- Required Privilege Level | 25872
- Sample Output | 25872
- Release Information | 25873

Syntax

```
show services nat source mappings pcp
<interface interface-name>
<service-set service-set.>
```

Description

Display NAT source mapping for PCP.

Options

interface *interface-name*	Display PCP source NAT mapping for the specified interface.
service-set *service-set*	Display PCP source NAT mapping for the specified service set.

Required Privilege Level

view

Sample Output

show services nat source mappings pcp

```
user@host> show services nat source mappings pcp Interface: vms-0/0/0, Service set: in
NAT pool: p
PCP Client       : 10.1.1.2                PCP lifetime : 995
Mapping          : 10.1.1.2        : 9000  --> 8.8.8.8         : 1025
Session Count    :     1
Mapping State    : Active

DS-LITE output:
===============
PCP Client       : 2222::1                PCP lifetime : 106
Mapping          : 88.1.0.47       :   47  --> 70.70.70.1      :41972
Session Count    :     1
Mapping State    : Active
B4 Address       : 2222::1
```

Release Information

Command introduced in Junos OS 20.1R1.

show services nat source mappings summary

IN THIS SECTION

- Syntax | 25873
- Description | 25873
- Options | 25873
- Required Privilege Level | 25874
- Output Fields | 25874
- Sample Output | 25874
- Release Information | 25874

Syntax

```
show services nat source mappings summary
<interface interface-name>
<service-set service-set.>
```

Description

Display NAT mapping summary information.

Options

interface *interface-name*	Display source NAT mapping information for the specified interface.
service-set *service-set*	Display source NAT mapping information for the specified service set.

Required Privilege Level

view

Output Fields

Table 1468 on page 25874 lists the output fields for the `show services nat source mappings summary` command. Output fields are listed in the approximate order in which they appear.

Table 1468: show services nat source mappings summary Output Fields

Field Name	Field Description
`Service Interface`	Name of the service interface.
`Total number of address mappings`	Displays total number of address mappings.
`Total number of endpoint independent port mappings`	Displays total number of endpoint independent port mappings.
`Total number of endpoint independent filters`	Displays total number of endpoint independent filters.

Sample Output

show services nat source mappings summary

```
user@host> show services nat source mappings summary
Service Interface:                              ms-2/0/0
Total number of address mappings:                          2
Total number of endpoint independent port mappings:        1
Total number of endpoint independent filters:              1
```

Release Information

Command introduced in Junos OS 19.3R2.

show services nat source pool

IN THIS SECTION

- Syntax | 25875
- Description | 25875
- Options | 25875
- Required Privilege Level | 25876
- Output Fields | 25876
- Sample Output | 25878
- Release Information | 25881

Syntax

```
show services nat source pool pool-name
<all>
<interface interface-name>
<service-set service-set>
```

Description

Display source NAT information for a pool.

Options

pool-name	Display information about the specified pool.
all	Display all source NAT pool information.
interface *interface-name*	Display information specific to the adaptive services interface.
service-set *service-set*	Display information specific to the service set.

Required Privilege Level

view

Output Fields

Table 1469 on page 25876 lists the output fields for the `show services nat source pool` command. Output fields are listed in the approximate order in which they appear.

Table 1469: show services nat source pool Output Fields

Field Name	Description
`Pool name`	Name of the source pool.
`Pool id`	Pool identification number.
`Routing instance`	Name of the routing instance.
`Host address base`	Base address of the original source IP address range.
`Port`	Port numbers used for the source pool.
`Port overloading`	Number of port overloading for the source pool.
`Address assignment`	Type of address assignment.
`Total addresses`	Number of IP addresses that are in use.
`Translation hits`	Number of times there is traffic that matches the source rule.
`Limit ports per host`	
`Include-boundary-addresses`	Include the lowest and highest addresses in the source address range of the NAT rule to be translated when the NAT pool is used.

Table 1469: show services nat source pool Output Fields *(Continued)*

Field Name	Description
Ei-mapping-timeout	Duration for endpoint independent translations that use the specified NAT pool.
Mapping-timeout	Duration for mappings that use the specified NAT pool.
EIF Inbound session count	Number of EIF inbound sessions.
EIF Inbound session limit exceeded drops	Number of EIF inbound sessions that exceed the drop limit.
Address range	IP address range for the source pool.
Ports	
Total used ports	
Error Counters • Out of port errors • Out of address errors • Parity port errors • Preserve Range errors • APP port allocation errors • App port limit allocation errors • Port block allocation errors • Port blocks limit exceeded errors	The following bullets describe the fields: • No ports available. • No room in the pool for another address. • • • • • •

Sample Output

show services nat source pool JNPR-CGNAT-PUB-POOL (NAT Pool)

```
user@host> show services nat source pool JNPR-CGNAT-PUB-POOL
Interface: vms-0/2/0 , Service set: JNPR-IF-SSET
Pool name          : JNPR-CGNAT-PUB-POOL
Pool id            : 4
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 65535]
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 254
Translation hits   : 0
+Limit ports per host : 10
Include-boundary-addresses: Disable
Ei-mapping-timeout : 300
Mapping-timeout    : 300
EIF Inbound session count: 0
EIF Inbound session limit exceeded drops: 0
Address range                          Ports
         20.20.20.1 - 20.20.20.254       0
Total used ports   :                     0
+Error Counters:
+    Out of port errors               : 0
+    Out of address errors            : 0
+    Parity port errors               : 0
+    Preserve Range errors            : 0
+    APP port allocation errors       : 0
+    APP port limit allocation errors : 0
+    Port block allocation errors     : 0
+    Port blocks limit exceeded errors: 0
```

show services nat source pool JNPR-CGNAT-PUB-POOL (PBA Pool)

```
user@host> show services nat source pool JNPR-CGNAT-PUB-POOL
Interface: vms-0/2/0 , Service set: JNPR-IF-SSET
Pool name          : JNPR-CGNAT-PUB-POOL
Pool id            : 4
```

```
Routing instance   : default
Port               : [1024, 65535]
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 510
Translation hits   : 0
Port block size      : 256
Max blocks per host  : 8
Active block timeout : 0
Interim logging interval   : 0
PBA block log        : Enable
Used/total port blocks: 0/128520
+Max number of port blocks used: 0
Include-boundary-addresses: Disable
Ei-mapping-timeout : 300
Mapping-timeout    : 300
EIF Inbound session count: 0
EIF Inbound session limit exceeded drops: 0
Address range                          Ports
          100.0.0.1 - 100.0.1.254        0
Total used ports   :                     0
Error Counters:
    Out of port errors                : 0
    Out of address errors             : 0
    Parity port errors                : 0
    Preserve Range errors             : 0
    APP port allocation errors        : 0
    APP port limit allocation errors  : 0
    Port block allocation errors      : 0
Port blocks limit exceeded errors : 0
```

show services nat source pool JNPR-CGNAT-PUB-POOL (Deterministic)

```
user@host> show services nat source pool JNPR-CGNAT-PUB-POOL
Interface: vms-0/2/0 , Service set: JNPR-IF-SSET
Pool name          : JNPR-CGNAT-PUB-POOL
Pool id            : 4
Routing instance   : default
Port               : [1024, 65535]
Port overloading   : 1
Address assignment : no-paired
```

```
Total addresses    : 510
Translation hits   : 0
Port block size      : 256
Determ host range num: 1
+Unique pool users: 0
Include-boundary-addresses: Disable
Ei-mapping-timeout : 300
Mapping-timeout    : 300
EIF Inbound session count: 0
EIF Inbound session limit exceeded drops: 0
Address range                        Single Ports   Twin Ports
          100.0.0.1 - 100.0.1.254        0              0
Total used ports   :                     0              0
Error Counters:
    Out of port errors                : 0
    Out of address errors             : 0
    Parity port errors                : 0
    Preserve Range errors             : 0
    APP port allocation errors        : 0
    APP port limit allocation errors  : 0
    Port block allocation errors      : 0
    Port blocks limit exceeded errors : 0
```

show services nat source pool service-set ss1_interface_style1 interface vms-0/2/0 all

```
user@router>show services nat source pool service-set ss1_interface_style1 interface vms-0/2/0
all
Interface: vms-0/2/0 , Service set: ss1_interface_style1
Pool name          : src_pool1
Pool id            : 4
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Twin port          : [63488, 65535]
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 254
Translation hits   : 3
Address range                        Single Ports   Twin Ports
```

```
            44.0.0.1 - 44.0.0.254         1              0
Total used ports   :                      1              0
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services nat source port-block

IN THIS SECTION

- Syntax | 25881
- Description | 25881
- Options | 25882
- Required Privilege Level | 25882
- Output Fields | 25882
- Sample Output | 25884
- Release Information | 25884

Syntax

```
show services nat source port-block
<host-ip ip-address>
<pool pool-name>
<xlated-ip translated-ip-address>
<xlated-port translated-port-number>
```

Description

Display port block allocation information.

Options

host-ip *ip-address*	Display port block allocation information for the specified host.
pool *pool-name*	Display port block allocation information for the specified pool.
xlated-ip *translated-ip-address*	Display port block allocation information for the specified translated IP address.
xlated-port *translated-port-number*	Display port block allocation information for the specified translated port number.

Required Privilege Level

view

Output Fields

Table 1470 on page 25882 lists the output fields for the `show services nat source port block` command. Output fields are listed in the approximate order in which they appear.

Table 1470: show services nat source port block Output Fields

Field Name	Field Description
Pool name	Name of the pool.
Port-overloading-factor	Factor of port overloading for the source pool.
Port block size	Number of ports that a port block contains.
Max port blocks per host	Maximum number of blocks that one host can use for translation.
Port block active timeout	Longest duration that a block remains active for port allocation.
Used/total port blocks	Current number of used ports and total number of ports in this source pool.
Host IP	Host IP address.

Table 1470: show services nat source port block Output Fields *(Continued)*

Field Name	Field Description
External IP	External IP address.
Port Block Range	Port range of one PBA port block entry from the lowest to the highest port number that can be allowed to allocate ports for this block.
Ports Used/Ports Total	Current number of used ports and total number of ports in this source pool.
Block State/Left Time (s)	PBA port block entry state for NAT port allocation, including Active, Inactive, Query, and the time left for a port block that is in the Active or Query state. • Active–When an internal subscriber initiates a NAT request, a port block is allocated from the pool, and the status is set to Active. When there is a subsequent request from the same subscriber, a port is allocated from the existing Active block. • Inactive–When there is a request from an internal subscriber who had previously had a port allocated from this port block, but the time on the Active port block has expired or the ports are used up, the port block status changes from Active to Inactive. • InactiveB–When a chassis cluster is in active/passive mode, and a port block is created on the active node, the status for the synced port block on the backup node is InactiveB. • Query–When no ports are used in an Active port block, the status changes from Active to Query.
Failed sessions	Number of failed sessions.
Number of sessions	Total number of sessions.

Sample Output

show services nat source port-block

```
user@host> show services nat source port-block
Pool name: sp1
Port-overloading-factor:     1     Port block size:  512
Max port blocks per host:    8     Port block active timeout:  100
Used/total port blocks: 1/64260
Host_IP           External_IP                    Port_Block          Ports_Used/
Block_State/
                                                   Range             Ports_Total
Left_Time(s)
1.1.1.100          30.30.30.1                     13824-14335          1/512*1
Active/71

    Failed sessions          : 0
  Number of sessions         : 0
```

Release Information

Command introduced in Junos OS 19.3R2.

show services nat source rule

IN THIS SECTION

- Syntax | 25885
- Description | 25885
- Options | 25885
- Required Privilege Level | 25885
- Output Fields | 25885
- Sample Output | 25887
- Release Information | 25888

Syntax

```
show services nat source rule
rule-name
<all>
<interface interface-name>
<service-set service-set>
```

Description

Display source NAT rule-set information.

Options

rule-name	Display source NAT rule-set information for the specified rule.
all	Display all source NAT rule-set information.
interface *interface-name*	Display rule-set information about the adaptive services interface.
service-set *service-set*	Display rule-set information about the service set.

Required Privilege Level

view

Output Fields

Table 1471 on page 25885 lists the output fields for the `show services nat source rule` command. Output fields are described in the approximate order in which they appear.

Table 1471: show services nat source rule Output Fields

Field Name	Description
`Interface`	Interface name.
`Service set`	Service set name.

Table 1471: show services nat source rule Output Fields *(Continued)*

Field Name	Description
`Rule Id`	Rule identification number.
`Rule position`	Position of the source NAT rule.
`Match-direction`	Specifies the direction in which to match traffic that meets the rule conditions.
`Match` • `Source address` • `Destination address` • `Application`	Match the following: • Name of the source address that matches the rule. • Name of the destination address that matches the rule. • Indicates whether the application option is configured.
`Action` • `Persistent NAT type` • `Persistent NAT mapping type` • `Inactivity timeout` • `Max session number`	
`Translation hits` • `Successful sessions` • `Failed sessions`	Use this field to check for traffic that matches the rule. Note the successful or failed sessions.
`Number of sessions`	Number of active sessions.

Sample Output

show services nat source rule

```
user@host> show services nat source rule all
ss1_interface_style1 interface vms-0/2/0 all | no-more
Interface: vms-0/2/0 , Service set: ss1_interface_style1
source NAT rule: r1                     Rule-set: rs1
  Rule-Id                    : 1
  Rule position              : 1
  Match-direction            : input
  Match
    Source addresses         : 0.0.0.0         - 255.255.255.255
    Destination addresses    : 0.0.0.0         - 255.255.255.255
    Application              : configured
  Action                        : src_pool1
    Persistent NAT type         : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout          : 0
    Max session number          : 0
  Translation hits           : 3
    Successful sessions      : 3
    Failed sessions          : 0
  Number of sessions         : 1
```

show services nat source rule (Mapping and EIF Configuration)

```
show services nat source rule all
Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 1/0
source NAT rule: r1                     Rule-set: rs1
  Rule-Id                    : 1
  Rule position              : 1
  From zone                  : nh-JNPR-NH-SSET-ZoneIn
  To zone                    : nh-JNPR-NH-SSET-ZoneOut
  Match
    Source addresses         : 30.30.30.0      - 30.30.30.255
  Action                        : p2
  +Mapping-type : endpoint-independent;
```

```
  +Mapping-refresh : inbound
  +Filtering-type: endpoint-independent
  +Prefix-list :
1.2.2.0    --- 2.2.2.3
3.3.3.0   --- 3.3.3.3  except
Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

Release Information

Command introduced in Junos OS 19.3R2.

show services nat source rule-application

IN THIS SECTION

- Syntax | 25888
- Description | 25889
- Options | 25889
- Required Privilege Level | 25889
- Output Fields | 25889
- Sample Output | 25890
- Release Information | 25890

Syntax

```
show services nat source rule-application
<all>
<interface interface-name>
<service-set service-set>
```

Description

Display source NAT rule application information.

Options

all	Display all source NAT rule application information.
interface ***interface-name***	Display source NAT rule application information for the specified interface.
service-set ***service-set***	Display source NAT rule application information for the specified service set.

Required Privilege Level

view

Output Fields

Table 1472 on page 25889 lists the output fields for the `show services nat source rule-application` command. Output fields are described in the approximate order in which they appear.

Table 1472: show services nat source rule-application Output Fields

Field Name	Description
Interface	Displays rule application for the specified interface.
Service set	Displays rule application for the specified service set.

Table 1472: show services nat source rule-application Output Fields *(Continued)*

Field Name	Description
Source NAT rule	The name of the source NAT rule.
• Rule-set	• Set of rules for matching traffic.
• Rule-Id	• Rule identification number.
• Match-direction	• Specifies the direction in which to match traffic that meets the rule conditions.
• Application	• Name of the application or application set.
• IP Protocol	• IP protocol identifier.
• Source port range	• Source port range identifier.
• Destination port range	• Destination port range identifier.

Sample Output

show services nat source rule-application

```
user@host> show services nat source rule-application service-set ss1_interface_style1 interface
vms-0/2/0 all
Interface: vms-0/2/0 , Service set: ss1_interface_style1
source NAT rule: r1                     Rule-set: rs1
  Rule-Id                    : 1
  Match-direction            : input
    Application: any
      IP protocol: 0
      Source port range: [0-0]
      Destination port range: [0-0]
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services nat source summary

IN THIS SECTION

- Syntax | 25891
- Description | 25891
- Options | 25891
- Required Privilege Level | 25891
- Output Fields | 25892
- Sample Output | 25893
- Release Information | 25893

Syntax

```
show services nat source summary
<interface interface-name>
<service-set service-set>
```

Description

Displays source NAT summary information.

Options

interface *interface-name*	Display source NAT summary information for the specified interface.
service-set *service-set*	Display source NAT summary information for the specified service set.

Required Privilege Level

view

Output Fields

Table 1473 on page 25892 lists the output fields for the `show services nat source summary` command. Output fields are listed in the approximate order in which they appear.

Table 1473: show services nat source summary Output Fields

Field Name	Description
`Interface`	Interface name.
`Service set`	Service set name.
`Pool Name`	Name of the source address pool.
`Address Range`	IP address or IP address range for the pool.
`Routing Instance`	Name of the routing instance.
`PAT`	Whether Port Address Translation (PAT) is enabled (yes or no).
`Total Address`	Number of IP addresses that are in use.
`Rule name`	Name of the rule.
`Rule set`	Set of rules.
`Match-direction`	Specifies the direction in which to match traffic that meets the rule conditions.
`Action`	Action taken for a packet that matches a rule.

Sample Output

show services nat source summary

```
user@host> show services nat source summary service-set ss1_interface_style11 interface vms-0/2/0
Interface: vms-0/2/0 , Service set: ss1_interface_style1
Pool                Address                  Routing              PAT  Total
Name                Range                    Instance                  Address
src_pool1           44.0.0.1-44.0.0.254      default              yes  254
Interface: vms-0/2/0 , Service set: ss1_interface_style1
Rule name  Rule set                Match-direction      Action
r1         rs1                          input          src_pool1
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services paa status

IN THIS SECTION

- Syntax | 25893
- Description | 25894
- Additional Information | 25894
- Required Privilege Level | 25894
- Output Fields | 25894
- Sample Output | 25896
- Release Information | 25898

Syntax

```
show services paa status
```

Description

Get the status of the Paragon Active Assurance (PAA) test agent. Because the test agent software is in a Docker container, the states reported with this command are the same as the states reported by a `docker inspect` command.

Additional Information

You can also use NETCONF to get the status of the PAA test agent:

```
<rpc>
  <get-paa-status>
  </get-paa-status>
</rpc>
]]>]]>
```

Required Privilege Level

view

Output Fields

Table 1474 on page 25895 lists the output fields for the `show services paa status` command.

Table 1474: show services paa status Output Fields

Field Name	Field Description
`Container Status`	Status of the PAA test agent and its Docker container. The possible statuses are: • Created—The Docker container was never started; the container is not using any CPU or memory. • Running—The PAA test agent is running inside the container. • Restarting—The container is restarting. • Exited—The PAA test agent has exited or encountered an exception and is no longer operational. This status also occurs if someone has stopped the container with the `docker stop` command. The container is not using any CPU or memory. • Paused—The container has been paused with the `docker pause` command. The container is still occupying memory, but the CPU has been released. • Dead—The container is not working. Often, a container stops working when we try to remove the container and some resources are still in use by an external process. The container is not using any memory or CPU. You cannot restart a dead container. You must uninstall the test agent and re-install.
`Control Center`	IP address or DNS hostname of the PAA Control Center to which this test agent reports.
`Dead`	If dead is `true`, then the Docker container is dead and the test agent needs to be uninstalled and installed again. This field is `false` if the Docker container is operational.
`Finished At`	Time of the last exit. If the Docker container is running normally, nothing has exited, and so this field shows zeroes. If this field is non-zero, you need to check the log messages to understand why.
`Image`	Name of the Docker container image and the PAA version number.
`OOMKilled`	If `true`, the Docker container was not created because the system did not have enough memory to add a new container. See Running Third-Party Applications in Containers for information about the default limits for memory and other resources that Docker containers can use from Junos OS Evolved.

Table 1474: show services paa status Output Fields *(Continued)*

Field Name	Field Description
Paused	If true, the container was paused with the docker pause command.
Pid	Process ID for this PAA test agent. If Pid is 0, the test agent did not install, possibly due to an incorrect password or user name for the PAA Control Center.
Restarting	If true, the container is restarting. The most common cause is incorrectly specifying your PAA Control Center credentials when installing the PAA test agent. Check the log messages to understand why. When the test agent fails to connect to the PAA Control Center, the container goes into Restarting state, because the test agent exits if the connection fails and the container is set to always restart.
Running	If false, the Docker container is not running. Check the Status field or log messages to understand why.
Started At	The time at which the Docker container was started.
Status	Shows the installation status of the PAA test agent.The possible statuses are: Uninstalling, Uninstalled, Installing, Installed.

Sample Output

show services paa status

```
user@host> show services paa status
```

```
Status: Installed
Error message: None
Control center: 10.49.23.49
Image: paa/test-agent-application:4.3.0.20
Container status: running
Running: true
Paused: false
```

```
Restarting: false
OOMKilled: false
Dead: false
Pid: 25078
ExitCode: 0
Error: None
Started At: 2023-12-12T15:02:11.993510852Z Finished At: 0001-01-01T00:00:00Z
VRF: vrf0
```

show services paa status with debug log messages (test agent was installed with the ta-debug option)

```
user@host> show services paa status
```

```
Status: Installed
Error message: None
Control center: 10.49.23.49
Image: paa/test-agent-application:4.3.0.20
Container status: running
Running: true
Paused: false
Restarting: false
OOMKilled: false
Dead: false
Pid: 25078
ExitCode: 0
Error: None
Started At: 2023-12-12T15:02:11.993510852Z Finished At: 0001-01-01T00:00:00Z
VRF: vrf0
Logs: 2023-12-12 15:02:13.082784Z DEBUG: populate_interface_state: ifname=lo0_0 name=lo0.0
Logs: 2023-12-12 15:02:13.084369Z DEBUG: Sending interface state to CC, type=InterfaceState
Logs: 2023-12-12 15:02:32.333126Z DEBUG: Received pong.
```

show services paa status with log messages showing a failed install because the address for the Control Center is incorrect

```
user@host> show services paa status
```

```
Status: Error
Error message: Connection to control center failed
Logs: curl: (7) Failed to connect to example.com port 443 after 3081 ms: Error
Logs: 2023-12-15 10:39:57Z ERROR:
Logs: 2023-12-15 10:39:57Z ERROR: Can't reach control center at address example.com
```

show services paa status with log messages showing a failed install because the user name or password was incorrect

```
user@host> show services paa status
```

```
Status: Error
Error message: Incorrect login credentials
Logs: 2023-12-12 15:06:27Z INFO: Setting environment
Logs: 2023-12-12 15:06:30Z ERROR: Incorrect user name or password
```

Release Information

Command introduced in Junos OS Evolved Release 22.3R1.

RELATED DOCUMENTATION

> *Install the Paragon Active Assurance (PAA) Test Agent*
>
> request services paa install | **14954**

show services pcef pic

IN THIS SECTION

- Syntax | 25899
- Description | 25899
- Options | 25899
- Required Privilege Level | 25899
- Output Fields | 25900
- Sample Output | 25900
- Release Information | 25900

Syntax

```
show services pcef pic
<fpc-slot slot-number>
<pic-slot pic-number>
```

Description

Displays the number of Junos OS Subscriber Management subscribers present on each service PIC who are a using policy control and enforcement function (PCEF) profile to define the treatment to apply to packets associated with specific applications (for example, Facebook) or to specific service data flows.

Options

fpc-slot *slot-number* (Optional) Display the number of PCEF subscribers present on the specified FPC.

pic-slot *pic-number* (Optional) Display the number of PCEF subscribers present on the specified PIC.

Required Privilege Level

view

Output Fields

Table 1475 on page 25900 lists the output fields for the `show services pcef pic` command. Output fields are listed in the approximate order in which they appear.

Table 1475: show services pcef pic Output Fields

Field Name	Field Description
`FPC Slot`	FPC slot on which the PCEF subscribers are present.
`PIC Slot`	PIC slot on which the PCEF subscribers are present.
`Active Subscribers`	Number of active PCEF subscribers that are present on the PIC slot.

Sample Output

show services pcef pic

```
user@host> show services pcef pic
FPC Slot    PIC Slot     Active Subscribers
    2           0                    1
    2           1                    1
    2           2                    0
    2           3                    0
```

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

show services pcef subscribers | 25901

show services pcef subscribers

IN THIS SECTION

- Syntax | 25901
- Description | 25902
- Options | 25902
- Required Privilege Level | 25904
- Output Fields | 25904
- Sample Output | 25909
- Sample Output | 25910
- Release Information | 25910

Syntax

```
show services pcef subscribers
<detail | extensive | summary | terse>
<aci-interface-set-name aci-interface-set-name>
<address address>
<agent-circuit-identifier agent-circuit-identifier-substring>
<client-type client-type>
<count>
<id>
<interface interface>
<logical-system logical-system>
<mac-address mac-address>
<physical-interface physical-interface-name>
<profile-name profile-name>
<routing-instance routing-instance>
<stacked-vlan-id stacked-vlan-id>
<subscriber-state subscriber-state>
<user-name user-name>
<vci vci-identifier>
<vpi vpi-identifier>
<vlan-id vlan-id>
```

Description

Displays information for Junos OS Subscriber Management subscribers who are a using policy control and enforcement function (PCEF) profile to define the treatment to apply to packets associated with specific applications (for example, Facebook) or to specific service data flows.

Options

detail | extensive | summary | terse (Optional) Display the specified level of output.

aci-interface-set-name (Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the DHCP or PPPoE control packets.

address (Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.0.2.0). If you specify the IP address as a prefix with a netmask (for example, 192.0.2.0/32), the router displays a message that the IP address is invalid, and rejects the command.

agent-circuit-identifier-substring (Optional) Display all dynamic subscriber sessions whose ACI value matches the specified substring.

client-type (Optional) Display subscribers whose client type matches one of the following client types:

- `dhcp`—DHCP clients only.
- `dotlx`—Dotlx clients only.
- `essm`—ESSM clients only.
- `fwauth`—FwAuth (authenticated across a firewall) clients only.
- `l2tp`—L2TP clients only.
- `mlppp`—MLPPP clients only.
- `ppp`—PPP clients only.
- `pppoe`—PPPoE clients only.

- `static`—Static clients only.
- `vlan`—VLAN clients only.
- `vpls-pw`—VPLS pseudowire clients only.
- `xauth`—Xauth clients only.

count (Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the `count` option alone or with the `address`, `client-type`, `interface`, `logical-system`, `mac-address`, `profile-name`, `routing-instance`, `stacked-vlan-id`, `subscriber-state`, or `vlan-id` options.

id (Optional) Display a specific subscriber session whose session id matches the specified subscriber ID. You can display subscriber IDs by using the `show subscribers extensive` or the `show subscribers interface extensive` commands.

interface (Optional) Display subscribers whose interface matches the specified interface.

logical-system (Optional) Display subscribers whose logical system matches the specified logical system.

mac-address (Optional) Display subscribers whose MAC address matches the specified MAC address.

physical-interface-name (M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

profile-name (Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance (Optional) Display subscribers whose routing instance matches the specified routing instance.

stacked-vlan-id (Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

subscriber-state (Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

user-name (M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

vci-identifier (MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is 0 through 255.

vpi-identifier (MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is 0 through 65,535.

vlan-id (Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the `stacked-vlan-id` *`stacked-vlan-id`* option to match the outer VLAN tag.

NOTE: Because of display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

Output Fields

Table 1476 on page 25904 lists the output fields for the `show services pcef subscribers` command. Output fields are listed in the approximate order in which they appear.

Table 1476: show services pcef subscribers Output Fields

Field Name	Field Description	Level of Output

`Interface`	Interface associated with the subscriber.	`detail`
		`extensive`
		`terse`

Table 1476: show services pcef subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`IP Address`	Subscriber IPv4 or IPv6 address.	`detail`
		`extensive`
		`terse`

`User Name`	Name of subscriber.	`detail`
		`extensive`
		`terse`

`Packets`	Number of the subscriber's packets that are processed by a PCEF profile. • `Input`—Input packets. • `Output`—Output packets.	`terse`
`Packet Drops`	Number of the subscriber's packets that were dropped as a result of being processed by a PCEF profile. • `Input`—Input packets. • `Output`—Output packets.	`terse`
`Session ID`	ID number for a subscriber service session.	`extensive` `detail`
`PFE Flow ID`	Variable-based forwarding flow ID.	`extensive` `detail`
`PCEF profile`	PCEF profile that is assigned to the subscriber.	`extensive` `detail`

Table 1476: show services pcef subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Routing Instance`	Routing instance associated with the subscriber.	`extensive` `detail`
`Service VRF`		`extensive` `detail`
`Service set`	Service set that is performing policy control.	`extensive` `detail`
`Input packets`	Number of the subscriber's input packets that are processed by a PCEF profile.	`extensive` `detail`
`Input octet`	Number of the subscriber's input octets that are processed by a PCEF profile.	`extensive` `detail`
`Output packets`	Number of the subscriber's output packets that are processed by a PCEF profile.	`extensive` `detail`
`Output octet`	Number of the subscriber's output octets that are processed by a PCEF profile.	`extensive` `detail`
`Input drops`	Number of the subscriber's input packets that were dropped as a result of being processed by a PCEF profile.	`extensive` `detail`

Table 1476: show services pcef subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input drop bytes`	Number of the subscriber's input bytes that were dropped as a result of being processed by a PCEF profile.	`extensive` `detail`
`Output drops`	Number of the subscriber's output packets that were dropped as a result of being processed by a PCEF profile.	`extensive` `detail`
`Output drop bytes`	Number of the subscriber's output bytes that were dropped as a result of being processed by a PCEF profile.	`extensive` `detail`
`Rule count`	Number of PCC rules that were applied to subscriber's traffic.	`extensive` `detail`

Table 1476: show services pcef subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Rule information`	For each PCC rule that is assigned to the subscriber, the following information appears: • `Rule name`—Name of PCC rule. • `In sessions`—Number of incoming sessions for the subscriber that are processed by the rule. • `Out sessions`—Number of outgoing sessions for the subscriber that are processed by the rule. • `Input packets`—Number of input packets and bytes for the subscriber that are processed by the rule. • `Output packets`—Number of output packets and bytes for the subscriber that are processed by the rule. • `Input drop packets`—Number of input packets and bytes for the subscriber that are dropped by the rule. • `Output drop packets`—Number of output packets and bytes for the subscriber that are dropped by the rule.	`extensive` `detail`
`Total subscribers`	Total number of subscribers.	`summary`

Sample Output

show services pcef subscribers terse

```
user@host> show services pcef subscribers terse
Interface Name     IP Address   User Name      Packets       Packet Drops
                                             Input   Output   Input  Output
demux0.3221225518  192.0.2.26   pcefuser   4215521 5155789      0  589900
```

show services pcef subscribers extensive

```
user@host> show services pcef subscribers extensive
Session ID: 60
PFE flow ID: 73
Interface: demux0.3221225518
IP address: 192.0.2.26
Username: pcefuser
PCEF profile: pcef-prof-1
Routing Instance: default
Service VRF: 1
Service set: 0
Input packets: 4229161
Input octet: 1632292118
Output packets: 5171863
Output octet: 2199274163
Input drops: 0
Input drop bytes: 0
Output drops: 591751
Output drop bytes: 882496606
Rule count: 2
  Rule information           :
    Rule name                :          limit-fb
    In sessions              :                13680
    Out sessions             :                    0
    Input packets            :              2629243            244119982 bytes
    Output packets           :              3702363           1117717893 bytes
    Input drop packets       :                    0                    0 bytes
    Output drop packets      :               591751            882496606 bytes
    Rule name                :          default
    In sessions              :                84545
```

```
    Out sessions            :                   0
    Input packets           :             1599918          1388172136 bytes
    Output packets          :             1469500          1081556270 bytes
    Input drop packets      :                   0                   0 bytes
    Output drop packets     :                   0                   0 bytes
```

Sample Output

show services pcef subscribers summary

```
user@host> show services pcef subscribers summary
Total subscribers: 1
```

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

show services pcef pic | 25899

show services pcp statistics

IN THIS SECTION

- Syntax | 25911
- Description | 25911
- Options | 25911
- Required Privilege Level | 25911
- Output Fields | 25911
- Sample Output | 25913
- Release Information | 25914

Syntax

```
show services pcp statistics
```

Description

Display information PCP mappings.

Options

Required Privilege Level

view

Output Fields

Table 1477 on page 25911 lists the output fields for the `show services pcp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1477: show services pcp statistics Output Fields

Field Name	Field Description
`Services PIC Name`	Name of a service interface.
`Protocol Statistics`	Overall PCP statistics, consisting of: operational, option, and results statistics.
`Operational Statistics`	Operational statistics group.
`Map request received`	Total PCP MAP requests received from PCP clients.
`Peer request received`	Number of peer requests received.
`Option Statistics`	Number of requests using available options.

Table 1477: show services pcp statistics Output Fields *(Continued)*

Field Name	Field Description
`Unprocessed requests received`	Number of requests received with no option specified.
`Third party requests received`	Number of third-party requests received.
`Prefer fail option received`	Number of prefer fail requests received.
`Filter option received`	Number of filter option requests received.
`Other options counters`	Number of packets received with options other than `prefer-fail` and `third-party`.
`Other optional received`	
`Results Statistics`	Information about the results of PCP requests.
`PCP success`	Number of PCP MAP requests successfully processed by the server.
`PCP unsupported version`	Number of PCP packets received with version other than 1.
`Not authorized`	Number of unauthorized MAP delete requests.
`Bad requests`	Number of requests with invalid PCP packets.
`Unsupported opcode`	Number of packets that have an unsupported opcode.
`Unsupported option`	Number of packets that have an unsupported option.
`Bad option`	Number of packet that have a malformed option.

Table 1477: show services pcp statistics Output Fields *(Continued)*

Field Name	Field Description
`Network failure`	Number of times a mapping could not be provided due to a network failure.
`Out of resources`	Number of times a mapping could not be provided because the PCP server ran out of pool resources.
`Unsupported protocol`	Number of requests for which the protocol was neither TCP nor UDP.
`User exceeded quota`	Number of requests for which the PCP client requested more than the configured number of ports.
`Cannot provide external`	Number of requests for which the PCP server cannot provide the external address or port requested by the client.
`Address mismatch`	Number of requests for which the PCP client IP address and the layer-3 source IP do not match.
`Excessive number of remote peers`	This counter is not currently used.
`Processing error`	Number of requests with malformed PCP packets information, such as an invalid IP address in a `third-party` request .
`Other result counters`	Not currently used.

Sample Output

show services pcp statistics pcp

```
user@host> show services pcp statistics pcp
  Services PIC Name:    sp-2/1/0

Protocol Statistics:
```

```
Operational Statistics

  Map request received                                    : 0
  Peer request received                                   : 0
  Other operational counters                              : 0

 Option Statistics

  Unprocessed requests received                           : 0
  Third party requests received                           : 0
  Prefer fail option received                             : 0
  Filter option received                                  : 0
  Other options counters                                  : 0
  Option optional received                                : 0

Result Statistics

  PCP success                                             : 0
  PCP unsupported version                                 : 0
  Not authorized                                          : 0
  Bad requests                                            : 0
  Unsupported opcode                                      : 0
  Unsupported option                                      : 0
  Bad option                                              : 0
  Network failure                                         : 0
  Out of resources                                        : 0
  Unsupported protocol                                    : 0
  User exceeded quota                                     : 0
  Cannot provide external                                 : 0
  Address mismatch                                        : 0
  Excessive number of remote peers                        : 0
  Processing error                                        : 0
  Other result counters                                   : 0
```

Release Information

Command introduced in Junos OS Release 13.2

show services policies

IN THIS SECTION

- Syntax | 25915
- Description | 25915
- Required Privilege Level | 25915
- Output Fields | 25915
- Sample Output | 25917
- Release Information | 25917

Syntax

```
show services policies
```

Description

Display services policy information.

Required Privilege Level

view

Output Fields

Table 1478 on page 25915 lists the output fields for the `show services policies` command. Fields are listed in the approximate order in which they appear.

Table 1478: show services policies Output Fields

Field Name	Description
`Default policy`	

Table 1478: show services policies Output Fields *(Continued)*

Field Name	Description
`Policy`	Name of the applicable policy.
`State`	Status of the policy: • `enabled`: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • `disabled`: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
`Index`	Internal number associated with the policy.
`Scope policy`	
`Sequence number`	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1,2,3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1,2,3,4.
`Stateful firewall rule`	
`Service set`	Name of the service set.
`Interface`	Name of the interface.
`Match direction`	
`Source addresses`	Names of the source addresses for a policy. Address sets are resolved to their individual Names of the source addresses for a policy. Address sets are resolved to their individual
`Destination addresses`	Name of the destination address (or address set as it was entered om the destination zone's address book.

Table 1478: show services policies Output Fields *(Continued)*

Field Name	Description
Application	

Sample Output

show services policies

```
user@host> show services policies
Default policy: deny-all
  Policy: p1, State: enabled, Index: 1007, Scope Policy: 0, Sequence number: 1
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0, Match
Direction: input
    Source addresses: any-ipv4
    Destination addresses: any
    Applications: junos-ftp
  Policy: p2, State: enabled, Index: 1008, Scope Policy: 0, Sequence number: 2
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0, Match
Direction: input
    Source addresses: any
    Destination addresses: any
    Applications: any
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services policies detail

IN THIS SECTION

Syntax | 25918

- Description | 25918
- Required Privilege Level | 25918
- Output Fields | 25918
- Sample Output | 25920
- Release Information | 25921

Syntax

```
show services policies detail
```

Description

Display detailed information about configured services policies.

Required Privilege Level

view

Output Fields

Table 1479 on page 25918 lists the output fields for the `show services policies detail` command. Output fields are listed in the approximate order in which they appear.

Table 1479: show services policies detail

Field Name	Description
`Default policy`	
`Policy`	
`Action type`	

Table 1479: show services policies detail *(Continued)*

Field Name	Description
State	Status of the policy: • enabled: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • disabled: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Scope policy	
Policy type	
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1,2,3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1,2,3,4.
Stateful firewall rule	
Service set	Service set name.
Interface	Interface name.
Source addresses	The names and corresponding IP addresses for the policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
Application	

Table 1479: show services policies detail *(Continued)*

Field Name	Description
IP protocol	
Inactivity timeout	
Source port range	
Destination port range	
Per policy TCP Options	

Sample Output

show services policies detail

```
user@host> show services policies detail
Default policy: deny-all
Policy: p1, action-type: permit, State: enabled, Index: 1007, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0, Match
Direction: input
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: junos-ftp
    IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [21-21]
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
Policy: p2, action-type: permit, State: enabled, Index: 1008, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
```

```
  Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0, Match
Direction: input
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services policies hit-count

IN THIS SECTION

- Syntax | 25921
- Description | 25922
- Required Privilege Level | 25922
- Output Fields | 25922
- Sample Output | 25922
- Release Information | 25922

Syntax

```
show services policies hit-count
```

Description

Display the hit count of policies.

Required Privilege Level

view

Output Fields

Sample Output

show services policies hit-count

```
user@host> show services policies hit-count
Index   Service Set      Interface         Name           Sfw rule        Direction       Policy
count
         1       JNPR-NH-SSET     vms-0/2/0         p1             sfw1
input           0
         2       JNPR-NH-SSET     vms-0/2/0         p2             sfw1
input           0
Number of policy: 2
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services policies interface

IN THIS SECTION

- Syntax | 25923
- Description | 25923
- Required Privilege Level | 25923

- Output Fields | 25923
- Sample Output | 25923
- Release Information | 25924

Syntax

```
show services policies interface interface-name
```

Description

Display services policies for the specified interface.

Required Privilege Level

view

Output Fields

Sample Output

show services policies interface vms-0/2/0

```
user@host> show services policies interface vms-0/2/0
Default policy: deny-all
  Policy: p1, State: enabled, Index: 1007, Scope Policy: 0, Sequence number: 1
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0, Match
Direction: input
    Source addresses: any-ipv4
    Destination addresses: any
    Applications: junos-ftp
  Policy: p2, State: enabled, Index: 1008, Scope Policy: 0, Sequence number: 2
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0, Match
Direction: input
    Source addresses: any
```

```
    Destination addresses: any
    Applications: any
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services policies service-set

IN THIS SECTION

- Syntax | 25924
- Description | 25924
- Required Privilege Level | 25924
- Output Fields | 25924
- Sample Output | 25925
- Release Information | 25925

Syntax

```
show services policies service-set service-set
```

Description

Display policy information for the specified service set.

Required Privilege Level

view

Output Fields

Sample Output

show services policies service-set

```
user@host> show services policies service-set JNPR-NH-SSET
Default policy: deny-all
  Policy: p1, State: enabled, Index: 1007, Scope Policy: 0, Sequence number: 1
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0, Match
Direction: input
    Source addresses: any-ipv4
    Destination addresses: any
    Applications: junos-ftp
  Policy: p2, State: enabled, Index: 1008, Scope Policy: 0, Sequence number: 2
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0, Match
Direction: input
    Source addresses: any
    Destination addresses: any
    Applications: any
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services redundancy-group

IN THIS SECTION

- Syntax | 25926
- Description | 25926
- Options | 25926
- Required Privilege Level | 25926
- Output Fields | 25926
- Sample Output | 25933
- Release Information | 25937

Syntax

```
show services redundancy-group
<rg-id>
<brief | extensive | terse>
```

Description

Display redundancy group status information for all redundancy groups or a specified redundancy group.

Options

rg-id (Optional) Name of a specific redundancy group.

brief | extensive | terse (Optional) Display the specified level of output. When no level is specified, display terse level output.

- **Default:** terse

Required Privilege Level

view

Output Fields

Table 1480 on page 25926 lists the output fields for the `show services redundancy-group` command. Output fields are listed in the approximate order in which they appear.

Table 1480: show services redundancy-group Output Fields

Field Name	Field Description	Level of Output
ICCP process connection	Status of the connection between the srd and iccpd. • Connected • Not connected	all levels
Redundancy Group ID	Identifier of the redundancy group.	all levels

Table 1480: show services redundancy-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Number of peer RG connections	Total number of peers in the redundancy group.	brief, extensive
Local RG IP	IP address of the local redundancy group.	all levels
RS ID		terse
Local RS state	State of the local redundancy set. • MASTER • STANDBY • INITIALIZING • STANDBY (WARNED)	terse
Peer RS state	State of the peer redundancy set. • MASTER • STANDBY • INITIALIZING • STANDBY (WARNED)	terse
Peer RG IP	Peer redundancy group IP address.	all
Status	Status of redundancy group connection with this peer. • Connected • Not Connected	terse
Number of peer RG connections	Total number of peers in the redundancy group.	brief

Table 1480: show services redundancy-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Redundancy Set ID	Identifier of the redundancy set.	brief, extensive
Connection status	Status of the connection between the srd and iccpd. • Connected • Not Connected	brief, extensive
Redundancy Set state	State of the local redundancy set state. • INITIALIZING • MASTER • STANDBY • STANDBY (WARNED)	brief, extensive
Redundancy Set peer state	State of the peer redundancy set state. • INITIALIZING • MASTER • STANDBY • STANDBY (WARNED)	brief, extensive
Redundancy Set health status	• Passed • Failed	brief, extensive
Number of Monitored interface down	Number of monitored interfaces that are d	brief, extensive
Failed Interfaces	List of all monitored interfaces that are down.	brief, extensive

Table 1480: show services redundancy-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Service Set	Service set used for stateful sync.	brief, extensive
Service Interface	Service set used for	brief, extensive
Type	Type of redundancy and stateful sync for the listed service interface. • Inter-chassis • Intra-chassis	brief, extensive
Role	Role of the listed service interface. • active • backup	brief, extensive
Connection	Status of connection with peer service PIC. • Up • Down	brief, extensive
Synchronization	Type of synchronization. When all eligible sessions are still synchronizing, it is cold synchronization. When all current existing sessions are synchronized, it is a HOT synchronization, When long lived sessions are eligible, they are synchronized. • Hot—All current existing sessions are synced. When long-lived sessions are eligible, they are synchronized. • Cold–Eligible sessions are in the processing of synchronizing.	brief, extensive

Table 1480: show services redundancy-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
ICCP process connection open complete count	Number of completed opens of ICCP process connections.	extensive
ICCP process connection close complete count	Number of completed closes of ICCP process connections.	
ICCP packet sent count	Number of ICCP packets sent.	extensive
ICCP packet receive count	Number of ICCP packets received.	extensive
ICCP process keepalive receive count	Number of ICCP process keepalive messages received.	extensive
ICCP process keepalive sent count	Number of ICCP process keepalive messages sent.	extensive
ICCP redundancy group add count	Number of redundancy group add messages received by srd from ICCP.	extensive
ICCP redundancy group delete count	Number of redundancy group delete messages received by srd from ICCP.	extensive
RG connection up count	Number of redundancy group connection up messages received by srd from ICCP.	extensive
RG connection down count	Number of redundancy group connection down messages received by srd from ICCP.	extensive
RG join count	Number of redundancy group join messages sent from srd to ICCP.	extensive

Table 1480: show services redundancy-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
RG data receive count	Number of packets of messages received by srd from a peer.	extensive
RG data sent count	Number of packets of messages sent from srd to a peer.	extensive
RG connect message sent count	Number of connect messages sent from srd to ICCP.	extensive
RG connect message receive count	Number of connect messages received by srd from ICCP.	extensive
RG disconnect message sent count	Number of disconnect messages sent from srd to ICCP.	extensive
RG disconnect message receive count	Number of disconnect messages received by srd from ICCP.	extensive
RG ack sent count	Number of RG ack messages sent.	extensive
RG nack sent count	Number of RG nack messages sent.	extensive
RG nack receive count	Number of RG nack messages received.	extensive
Transition Events Received	Number of transition events received in each of the following categories: • Acquire primary role auto • Acquire primary role manual • Release primary role auto • Release primary role manual	extensive

Table 1480: show services redundancy-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Transition Events Ignored	Number of transition events ignored in each of the following categories: • Acquire primary role auto • Acquire primary role manual • Release primary role auto • Release primary role manual In a high-availability or redundancy pair of SDGs, in which one SDG is the primary and the other is the standby, when perform a double failover of the SDGs, the second failover event is not ignored, which is the expected behavior. The event is not disregarded because it arrives as a critical redundancy-event based on the redundancy-policy. However, because the SDG is already be in Standby state, the finite state machine transitions to the Standby-Warned state until it recovers. Therefore, the event is honored and not ignored. Although there was no primary role transition, it is because of a valid reason that the SDG is already in Standby state. The redundancy-event is associated with to a primary role release policy based on the configuration and the Release primary role field under the Transition Events Ignored column displays a number that corresponds to the redundancy event. The services redundancy daemon (SRD) finite state machine quickly recovers (transitions from Standby-Warned to Standby) during restart-routing because the rpd restart-handling and recovery are fast and the following critical event is not ignored. However, disabling or deactivating the interface results in the FSM remaining in Standby-Warned until the interface is up. Any critical events during the time when the interface is down are ignored because the state is already Standby-Warned and does not transition to a different state. In summary, the following is the manner in which critical events are analyzed during state transitions:	extensive

Table 1480: show services redundancy-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• Standby -> Standby Warned = Critical Event Not ignored [valid state transition] • Standby Warned -> Standby Warned = Critical Event Ignored [no state transition]	
Monitored Events Received	Number of monitored events received in each of the following categories: • Link-down • Routing restart/terminate • Route update error • Peer primary-role-acquire • Peer primary-role-release	extensive
Monitored Events Ignored	Number of monitored events ignored in each of the following categories: • Link-down • Routing restart/terminate • Route update error • Peer primary-role-acquire • Peer primary-role-release	extensive

Sample Output

show services redundancy-group terse

```
user@host> show services redundancy-group terse
ICCP process connection                : Connected
```

```
Redundancy Group ID              : 1
  Number of peer RG connections  : 1
  Local RG IP                    : 172.19.39.70
  RS ID     Local RS state     Peer RS state     Peer RG IP        Status
  1         MASTER             STANDBY           172.19.39.69      Connected
```

show services redundancy-group brief (Health Status Passed)

```
user@host> show services redundancy-group brief
ICCP process connection                  : Connected
Redundancy Group ID                      : 1
  Number of peer RG connections          : 1
  Local RG IP                            : 172.19.39.70
  Redundancy Set ID                      : 1
    Connection status                    : Connected
    Redundancy Set state                 : MASTER
    Redundancy Set peer state            : STANDBY
    Peer RG IP                           : 172.19.39.69
    Redundancy Set health status         : Passed
      Service Set : IPv6-SFW
        Service interface     Type              Role       Connection     Synchronization
        ms-1/3/0              Inter-chassis     active     Up             Hot

        ms-1/2/0              Inter-chassis     active     Up             Hot

        ms-1/1/0              Inter-chassis     active     Up             Hot

        ms-1/0/0              Inter-chassis     active     Up             Hot

      Service Set : NAPT44-SS1-SS4
        Service interface     Type              Role       Connection     Synchronization
        ms-1/3/0              Inter-chassis     active     Up             Hot

        ms-1/2/0              Inter-chassis     active     Up             Hot

        ms-1/1/0              Inter-chassis     active     Up             Hot

        ms-1/0/0              Inter-chassis     active     Up             Hot
```

show services redundancy-group brief (Health Status Failed)

```
user@host> show services redundancy-group brief
ICCP Process Connection                : Connected
Redundancy Group ID   : 1
  Number of Members   : 2
  Redundancy Set ID       : 1
    Remote IP address                              : 203.0.113.2
    Connection Status                               : Connected
    Redundancy Set State                         : STANDBY (WAIT)
    Redundancy Set Peer State               : MASTER
    Redundancy Set Health Status          : Failed
      Number of  Monitored interface down : 1             <<<<<<< Failure Reasons
      Failed Interfaces                                                                <<<<<<<  Name
of the monitored interfaces which have gone  down
        ms-2/3/0
      Service Set : ss2
        Service Interface     Type                  Role           Connection     Synchronization
        ms-2/2/0                  Inter-chassis     backup     Up                     Hot
        ms-2/1/0                  Inter-chassis     backup     Down                Off
        ms-2/0/0                  Inter-chassis     backup     Down                Off
      Service Set : ss_new
        Service Interface     Type                  Role          Connection     Synchronization
        ms-2/3/0
```

show services redundancy-group extensive

```
user@host> show services redundancy-group extensive
ICCP process connection                        : Connected
ICCP process connection close count            : 0
ICCP process connection open complete count    : 1
ICCP packet sent count                         : 7303
ICCP packet receive count                      : 7321
ICCP process keepalive receive count           : 7253
ICCP process keepalive sent count              : 7253
ICCP redundancy group add count                : 0
ICCP redundancy group delete count             : 0
Redundancy Group ID                      : 1
  Number of peer RG connections          : 1
  Local RG IP                            : 172.19.39.70
```

```
RG connection up count                 : 4
RG connection down count               : 2
RG join count                          : 4
RG data receive count                  : 37
RG data sent count                     : 0
RG connect message sent count          : 4
RG connect message receive count       : 4
RG disconnect message sent count       : 0
RG disconnect message receive count    : 4
RG ack sent count                      : 4
RG nack sent count                     : 0
RG nack receive count                  : 4
Redundancy Set ID                      : 1
  Connection status                    : Connected
  Redundancy Set state                 : MASTER
  Redundancy Set peer state            : STANDBY
  Peer RG IP                           : 172.19.39.69
  Redundancy Set health status         : Passed
    Service Set : IPv6-SFW
      Service interface     Type              Role       Connection     Synchronization
      ms-1/3/0              Inter-chassis     active     Up             Hot

      ms-1/2/0              Inter-chassis     active     Up             Hot

      ms-1/1/0              Inter-chassis     active     Up             Hot

      ms-1/0/0              Inter-chassis     active     Up             Hot

    Service Set : NAPT44-SS1-SS4
      Service interface     Type              Role       Connection     Synchronization
      ms-1/3/0              Inter-chassis     active     Up             Hot

      ms-1/2/0              Inter-chassis     active     Up             Hot

      ms-1/1/0              Inter-chassis     active     Up             Hot

      ms-1/0/0              Inter-chassis     active     Up             Hot

  Transition events              Received         Ignored
    Acquire mastership auto      3                0
    Acquire mastership manual    0                0
    Release mastership auto      3                0
    Release mastership manual    0                0
```

```
Monitored events                Received         Ignored
  Link-down                     145              31
  Routing restart/abort         1                0
  Route update error            0                0
  Peer mastership-acquire       3                0
  Peer mastership-release       3                0
```

Release Information

Statement introduced in Junos OS Release 16.1.

show services rpm active-servers

IN THIS SECTION

- Syntax | 25937
- Description | 25937
- Options | 25938
- Required Privilege Level | 25938
- Output Fields | 25938
- Sample Output | 25938
- Release Information | 25938

Syntax

```
show services rpm active-servers
```

Description

Display the protocols and corresponding ports for which a router or switch is configured as a real-time performance monitoring (RPM) server.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1481 on page 25938 lists the output fields for the `show services rpm active-servers` command. Output fields are listed in the approximate order in which they appear.

Table 1481: show services rpm active-servers Output Fields

Field Name	Field Description
Protocol	Protocol configured on the receiving probe server. The protocol can be the User Datagram Protocol (UDP) or the Transmission Control Protocol (TCP).
Port	Port configured on the receiving probe server.
Destination interface name	Output interface name for the probes.

Sample Output

show services rpm active-servers

```
user@host> show services rpm active-servers
    Protocol: TCP, Port: 50000, Destination interface name: lt-0/0/0.0
    Protocol: UDP, Port: 50001, Destination interface name: lt-0/0/0.0
```

Release Information

Command introduced before Junos OS Release 7.4.

show services rpm history-results

IN THIS SECTION

- Syntax | 25939
- Description | 25939
- Options | 25939
- Required Privilege Level | 25940
- Output Fields | 25940
- Sample Output | 25942
- Release Information | 25944

Syntax

```
show services rpm history-results
<brief | detail>
<dst-interface interface-name>
<owner owner>
<limit number>
<since time>
<source-address address>
<target-address address>
<test name>
```

Description

Display the results stored for the specified real-time performance monitoring (RPM) probes.

Options

none (Optional) Display the results of the last 50 probes for all RPM instances.

brief | detail (Optional) Display the specified level of output.

dst-interface *interface-name*
(Optional) Display information only for RPM probes that are generated on this MS-MPC or MS-MIC services interface. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. You must also configure the `owner` option.

limit *number*
(Optional) Limit the number of results that are displayed. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. You must also configure the `owner` option.

- **Range:** 1 through 4,294,967,295
- **Default:** 100

owner *owner*
(Optional) Display information only for probes with the specified probe owner. You must configure `owner` if you configure any of the following options: `dst-interface`, `limit`, `source-address`, or `target-address`.

since *time*
(Optional) Display information from the specified time. Specify time as *yyyy-mm-dd.hh:mm:ss*.

source-address *address*
(Optional) Display information only for probes with the specified source address. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. You must also configure the `owner` option.

target-address *address*
(Optional) Display information only for probes with the specified target address. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. You must also configure the `owner` option.

test *name*
(Optional starting in Junos OS Release 18.1R1) Display information only for the specified test.

Do not configure `test` if you configure any of the following options: `dst-interface`, `limit`, `source-address`, or `target-address`. These options do not work when you configure `test`.

Required Privilege Level

view

Output Fields

Table 1482 on page 25941 lists the output fields for the `show services rpm history-results` command. Output fields are listed in the approximate order in which they appear.

Table 1482: show services rpm history-results Output Fields

Field Name	Field Description	Level of Output
`Owner`	Probe owner.	All levels
`Test`	Name of a test for a probe instance.	All levels
`Probe received`	Timestamp when the probe result was determined.	All levels
`Round trip time`	Average ping round-trip time (RTT), in microseconds.	All levels
`Probe results`	Result of a particular probe performed by a remote host. The following information is contained in the results: • `Response received`—Timestamp when the probe result was determined. • `Rtt`—Average ping round-trip time (RTT), in microseconds.	`detail`
`Results over current test`	Displays the results for the current test by probe at the time each probe was completed, as well as the status of the current test at the time the probe was completed.	`detail`
`Probes sent`	Number of probes sent with the current test.	`detail`
`Probes received`	Number of probe responses received within the current test.	`detail`
`Loss percentage`	Percentage of lost probes for the current test.	`detail`

Table 1482: show services rpm history-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Measurement`	Increment of measurement. Possible values are round-trip time delay and, for the probe type icmp-pin-timestamp, the egress and ingress delay: • `Minimum`—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • `Maximum`—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • `Average`—Average RTT, ingress delay, or egress delay measured over the course of the current test. • `Jitter`—Difference, in microseconds, between the maximum and minimum RTT measured over the course of the current test. • `Stddev`—Standard deviation of the round-trip time, in microseconds, measured over the course of the current test.	`detail`

Sample Output

show services rpm history-results owner test

```
user@host> show services rpm history-results owner p1 test t1
     Owner, Test                 Probe received              Round trip time
    p1, t1                      Wed Aug 12 01:02:35 2009              315 usec
    p1, t1                      Wed Aug 12 01:02:36 2009              266 usec
    p1, t1                      Wed Aug 12 01:02:37 2009              314 usec
    p1, t1                      Wed Aug 12 01:02:38 2009              388 usec
    p1, t1                      Wed Aug 12 01:02:39 2009              316 usec
    p1, t1                      Wed Aug 12 01:02:40 2009              271 usec
    p1, t1                      Wed Aug 12 01:02:41 2009              314 usec
    p1, t1                      Wed Aug 12 01:02:42 2009             1180 usec
```

show services rpm history-results owner test detail

```
user@host> show services rpm history-results owner p1 test t1 detail
    Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
    Probe results:
      Response received, Wed Aug 12 01:02:35 2009,
      Client and server hardware timestamps
      Rtt: 315 usec
    Results over current test:
      Probes sent: 1, Probes received: 1, Loss percentage: 0
      Measurement: Round trip time
        Samples: 1, Minimum: 315 usec, Maximum: 315 usec, Average: 315 usec,
        Peak to peak: 0 usec, Stddev: 0 usec, Sum: 315 usec

    Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
    Probe results:
      Response received, Wed Aug 12 01:02:36 2009,
      Client and server hardware timestamps
      Rtt: 266 usec, Round trip jitter: -50 usec,
      Round trip interarrival jitter: 3 usec
    Results over current test:
      Probes sent: 2, Probes received: 2, Loss percentage: 0
      Measurement: Round trip time
        Samples: 2, Minimum: 266 usec, Maximum: 315 usec, Average: 291 usec,
        Peak to peak: 49 usec, Stddev: 24 usec, Sum: 581 usec
      Measurement: Negative round trip jitter
        Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
        Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec

    Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
    Probe results:
      Response received, Wed Aug 12 01:02:37 2009,
      Client and server hardware timestamps
      Rtt: 314 usec, Round trip jitter: 49 usec,
      Round trip interarrival jitter: 6 usec
    Results over current test:
      Probes sent: 3, Probes received: 3, Loss percentage: 0
      Measurement: Round trip time
        Samples: 3, Minimum: 266 usec, Maximum: 315 usec, Average: 298 usec,
        Peak to peak: 49 usec, Stddev: 23 usec, Sum: 895 usec
      Measurement: Positive round trip jitter
        Samples: 1, Minimum: 49 usec, Maximum: 49 usec, Average: 49 usec,
```

```
        Peak to peak: 0 usec, Stddev: 0 usec, Sum: 49 usec
      Measurement: Negative round trip jitter
        Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
        Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec

    Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
    Probe results:
      Response received, Wed Aug 12 01:02:38 2009,
      Client and server hardware timestamps
      Rtt: 388 usec, Round trip jitter: 74 usec,
      Round trip interarrival jitter: 10 usec
    Results over current test:
      Probes sent: 4, Probes received: 4, Loss percentage: 0
      Measurement: Round trip time
        Samples: 4, Minimum: 266 usec, Maximum: 388 usec, Average: 321 usec,
        Peak to peak: 122 usec, Stddev: 44 usec, Sum: 1283 usec
      Measurement: Positive round trip jitter
        Samples: 2, Minimum: 49 usec, Maximum: 74 usec, Average: 62 usec,
        Peak to peak: 25 usec, Stddev: 12 usec, Sum: 123 usec
      Measurement: Negative round trip jitter
        Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
        Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec
```

Release Information

Command introduced before Junos OS Release 7.4.

`dst-interface`, `limit`, `source-address`, and `target-address` options introduced in Junos OS Release 18.1R1 on MX Series.

`owner` and `test` options became optional in Junos OS Release 18.1R1 on MX Series.

show services rpm probe-results

IN THIS SECTION

- Syntax | 25945
- Description | 25945

- Options | 25945
- Required Privilege Level | 25946
- Output Fields | 25946
- Sample Output | 25955
- Release Information | 25960

Syntax

```
show services rpm probe-results
<dst-interface interface-name>
<limit number>
<owner owner>
<source-address address>
<status (fail | pass) >
<target-address address>
<terse>
<test name>
```

Description

Display the results of the most recent real-time performance monitoring (RPM) probes.

Options

All the following options require that you also configure the `owner` option.

dst-interface *interface-name*
(Optional) Display information only for RPM probes that are configured on this MS-MPC or MS-MIC services interface. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface.

limit *number*
(Optional) Limit the number of results that are displayed. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface.

- **Range:** 1 through 4,294,967,295
- **Default:** 100

none	Display information for all of the most recent RPM probes.
owner *owner*	(Optional) Display information only for probes with the specified probe owner. You must configure `owner` if you configure any other options.
source-address *address*	(Optional) Display information only for probes with the specified source address. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface.
status	(Optional) Display information only for probes with the specified type of test result. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. Specify one of the following: **fail** — Failed tests **pass** — Passed tests
target-address *address*	(Optional) Display information only for probes with the specified target address. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface.
terse	(Optional) Display summary information. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface.
test *name*	(Optional) Display information only for the specified test. Do not configure `test` if you configure any of the following options: `dst-interface`, `source-address`, or `target-address`. These options do not work when you configure `test`.

Required Privilege Level

view

Output Fields

Table 1483 on page 25947 lists the output fields for the `show services rpm probe-results` command. Output fields are listed in the approximate order in which they appear.

Table 1483: show services rpm probe-results Output Fields

Field Name	Field Description	Level of Output
`Owner`	Owner name. When you configure the probe owner statement at the `[edit services rpm]` hierarchy level, this field displays the configured owner name. When you configure BGP neighbor discovery through RPM, the output for this field is `Rpm-Bgp-Owner`.	none `dst-interface` `limit` `owner` `source-address` `target-address` `test`
`Test`	Name of a test representing a collection of probes. When you configure the test test-name statement at the `[edit services rpm probe owner]` hierarchy level, the field displays the configured test name. When you configure BGP neighbor discovery through RPM, the output for this field is `Rpm-BGP-Test-` *n*, where *n* is a cumulative number.	All levels
`Target address`	Destination IPv4 address used for the probes. This field is displayed when the probes are sent to the configured IPv4 or IPv6 targets or RPM servers.	none `dst-interface` `limit` `owner` `source-address` `target-address` `terse` `test`
`Target inet6-address`	Destination IPv6 address used for the probes. This field is displayed when the probes are sent to the configured IPv6 targets or RPM servers.	none `dst-interface` `limit` `owner` `source-address` `target-address` `terse` `test`
`Source address`	Source address used for the probes.	none `dst-interface` `limit` `owner` `source-address` `target-address` `test`
`Probe type`	Protocol configured on the receiving probe server: `http-get`, `http-metadata-get`, `icmp-ping`, `icmp-ping-timestamp`, `tcp-ping`, `udp-ping`, or `udp-ping-timestamp`.	none `dst-interface` `limit` `owner` `source-address` `target-address` `test`
`Test size`	Number of probes within a test.	none `dst-interface` `limit` `owner` `source-address` `target-address` `test`

Table 1483: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Routing Instance Name`	(BGP neighbor discovery) Name of the configured (if any) routing instance, logical system name, or both, in which the probe is configured: • When a routing instance is defined within a logical system, the logical system name is followed by the routing instance name. A slash (/) is used to separate the two entities. For example, if the routing instance called `R1` is configured within the logical system called `LS`, the name in the output field is `LS/R1`. • When a routing instance is configured but the default logical system is used, the name in the output field is the name of the routing instance. • When a logical system is configured but the default routing instance is used, the name in the output field is the name of the logical system followed by `default`. A slash (/) is used to separate the two entities. For example, `LS/default`.	none `dst-interface` `limit` `owner` `source-address` `target-address` `test`

Table 1483: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Probe results`	Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement:
	• `Response received`
	  • `Probe sent time`—Timestamp when the probe's results was sent.
	  • `Probe rcvd/timeout time`—Timestamp when the probe's results was received.
	• `Client and server hardware timestamps`—If timestamps are configured, an entry appears at this point.
	• `Rtt`—Average ping round-trip time (RTT), in microseconds.
	• `Egress jitter`—Egress jitter, in microseconds.
	• `Ingress jitter`—Ingress jitter, in microseconds.
	• `Round trip jitter`—Round-trip jitter, in microseconds.
	• `Egress interarrival jitter`—Egress interarrival jitter, in microseconds.
	• `Ingress interarrival jitter`—Ingress interarrival jitter, in microseconds.
	• `Round trip interarrival jitter`—Round-trip interarrival jitter, in microseconds.
		none `dst-interface` `limit` `owner` `source-address` `target-address` `test`

Table 1483: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Results over current test`	Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes. • `Probes sent`—Number of probes sent within the current test. • `Probes received`—Number of probe responses received within the current test. • `Loss percentage`—Percentage of lost probes for the current test. • `Measurement`—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type `icmp-ping-timestamp`, the egress delay and ingress delay. For each measurement type, the following individual calculated results are provided: • `Samples`—Number of probes. • `Minimum`—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • `Maximum`—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • `Average`—Average RTT, ingress delay, or egress delay measured over the course of the current test.	none `dst-interface` `limit` `owner` `source-address` `target-address` `test`

Table 1483: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Peak to peak`—Peak-to-peak difference, in microseconds. • `Stddev`—Standard deviation, in microseconds. • `Sum`—Statistical sum.	

Table 1483: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Results over last test`	Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed • `Probes sent`—Number of probes sent for the most recently completed test. • `Probes received`—Number of probe responses received for the most recently completed test. • `Loss percentage`—Percentage of lost probes for the most recently completed test. • `Test completed`—Time the most recent test was completed. • `Measurement`—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type `icmp-ping-timestamp`, the egress delay and ingress delay. For each measurement type, the following individual calculated results are provided: • `Samples`—Number of probes. • `Minimum`—Minimum RTT, ingress delay, or egress delay measured for the most recently completed test. • `Maximum`—Maximum RTT, ingress delay, or egress delay measured for the most recently completed test. • `Average`—Average RTT, ingress delay, or egress delay measured for the most recently completed test.	none `dst-interface` `limit` `owner` `source-address` `target-address` `test`

Table 1483: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Peak to peak`—Peak-to-peak difference, in microseconds. • `Stddev`—Standard deviation, in microseconds. • `Sum`—Statistical sum.	

Table 1483: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Results over all tests`	Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test. • `Probes sent`—Number of probes sent in all tests. • `Probes received`—Number of probe responses received in all tests. • `Loss percentage`—Percentage of lost probes in all tests. • `Measurement`—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe types `icmp-ping-timestamp` and `udp-ping-timestamp`, the egress delay and ingress delay. For each measurement type, the following individual calculated results are provided: • `Samples`—Number of probes. • `Minimum`—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • `Maximum`—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • `Average`—Average RTT, ingress delay, or egress delay measured over the course of the current test. • `Peak to peak`—Peak-to-peak difference, in microseconds. • `Stddev`—Standard deviation, in microseconds. • `Sum`—Statistical sum.	none `dst-interface` `limit` `owner` `source-address` `target-address` `test`

Table 1483: show services rpm probe-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Error Stats	Displays error statistics for each probe. • Invalid client recv timestamp—Number of client receive timestamp less than client send timestamp. • Invalid server send timestamp—Number of server send timestamp less than server receive timestamp. • Invalid server processing time—Number of server side spent time greater than RTT. **NOTE**: Error Stats is displayed in the output only if non-zero statistics exists.	none dst-interface limit owner source-address target-address test
Last Probe Status	Status of the last probe that was sent for the current test (fail or pass).	status
Status	Status of the last completed test (up or down).	status terse
Source-IF	The MS-MPC or MS-MIC services interface that generates the RPM probes.	terse

Sample Output

show services rpm probe-results (IPv4 Targets)

```
user@host> show services rpm probe-results
    Owner: ADSN-J4300.ADSN-J2300.D2, Test: 75300002
    Target address: 172.16.54.172, Source address: 10.206.0.1,
    Probe type: udp-ping-timestamp, Test size: 10 probes
    Probe results:
      Response received
      Probe sent time: Tue Feb  6 14:53:15 2007,
      Probe rcvd/timeout time: Tue Feb 6 14:53:15 2007
      Client and server hardware timestamps
      Rtt: 575 usec, Egress jitter: 5 usec, Ingress jitter: 8 usec,
```

```
      Round trip jitter: 12 usec, Egress interarrival jitter: 8 usec,
      Ingress interarrival jitter: 7 usec, Round trip interarrival jitter: 7 usec,
      Round trip interarrival jitter: 669 usec
    Results over current test:
      Probes sent: 10, Probes received: 10, Loss percentage: 0
      Measurement: Round trip time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec,  Sum: xxxx usec
      Measurement: Positive round trip jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative round trip jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
      Measurement: Egress time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
      Measurement: Positive Egress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative Egress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
      Measurement: Ingress time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
      Measurement: Positive Ingress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative Ingress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
    Results over last test:
      Probes sent: 10, Probes received: 10, Loss percentage: 0
      Test completed on Tue Feb  6 14:53:16 2007
      Measurement: Round trip time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
      Measurement: Positive round trip jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative round trip jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
```

```
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
      Measurement: Egress time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
      Measurement: Positive Egress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative Egress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
      Measurement: Ingress time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
      Measurement: Positive Ingress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative Ingress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
    Results over all tests:
      Probes sent: 560, Probes received: 560, Loss percentage: 0
      Measurement: Round trip time
        Samples: 560, Minimum: 805 usec, Maximum: 3114 usec, Average: 1756 usec,
        Peak to peak: 2309 usec, Stddev: 519 usec, Sum: xxxx usec
      Measurement: Positive round trip jitter
        Samples: 257, Minimum: 0 usec, Maximum: 2054 usec, Average: 597 usec,
        Peak to peak: 2054 usec, Stddev: 427 usec, Sum: xxxx usec
      Measurement: Negative round trip jitter
        Samples: 302, Minimum: 1 usec, Maximum: 1812 usec, Average: 511 usec,
        Peak to peak: 1811 usec, Stddev: 408 usec, Sum: xxxx usec
      Measurement: Egress time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
      Measurement: Positive Egress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative Egress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
      Measurement: Ingress time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
      Measurement: Positive Ingress  jitter
```

```
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative Ingress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
    Error Stats:
        Invalid client recv timestamp: 3, Invalid server send timestamp: 0
        Invalid server processing time: 0
```

show services rpm probe-results (IPv6 Targets)

```
user@host> show services rpm probe-results
    Owner: p, Test: t1
    Target inet6-address: 2001:db8:0:1:2a0:a502:0:1da,
    Target Port : 34567 Test size: 1000000 probes
    Probe results:
      Response received
      Probe sent time: Mon Dec 16 10:48:07 2013
      Probe rcvd/timeout time: Mon Dec 16 10:48:07 2013
      Client and server hardware timestamps
      Rtt: 236 usec, Round trip jitter: -10 usec, Round trip interarrival jitter: 484 usec
    Results over current test:
      Probes sent: 10, Probes received: 10, Loss percentage: 0
      Measurement: Round trip time
        Samples: 10, Minimum: 231 usec, Maximum: 298 usec, Average: 268 usec, Peak to peak: 67
usec, Stddev: 24 usec, Sum: 2682 usec
      Measurement: Positive round trip jitter
        Samples: 3, Minimum: 15 usec, Maximum: 1841 usec, Average: 750 usec, Peak to peak: 1826
usec, Stddev: 787 usec, Sum: 2251 usec
      Measurement: Negative round trip jitter
        Samples: 7, Minimum: 10 usec, Maximum: 1244 usec, Average: 709 usec, Peak to peak: 1234
usec, Stddev: 466 usec, Sum: 4961 usec
    Results over last test:
      Probes sent: 10, Probes received: 10, Loss percentage: 0
      Test completed on Mon Dec 16 10:48:07 2013
      Measurement: Round trip time
        Samples: 10, Minimum: 231 usec, Maximum: 298 usec, Average: 268 usec, Peak to peak: 67
usec, Stddev: 24 usec, Sum: 2682 usec
      Measurement: Positive round trip jitter
        Samples: 3, Minimum: 15 usec, Maximum: 1841 usec, Average: 750 usec, Peak to peak: 1826
usec, Stddev: 787 usec, Sum: 2251 usec
```

```
      Measurement: Negative round trip jitter
        Samples: 7, Minimum: 10 usec, Maximum: 1244 usec, Average: 709 usec, Peak to peak: 1234
usec, Stddev: 466 usec, Sum: 4961 usec
    Results over all tests(From start of current control session):
      Probes sent: 490, Probes received: 488, Loss percentage: 0
      Measurement: Round trip time
        Samples: 488, Minimum: 231 usec, Maximum: 306 usec, Average: 270 usec, Peak to peak: 75
usec, Stddev: 16 usec, Sum: 131586 usec
      Measurement: Positive round trip jitter
        Samples: 254, Minimum: 0 usec, Maximum: 10151 usec, Average: 157 usec, Peak to peak:
10151 usec, Stddev: 873 usec, Sum: 39817 usec
      Measurement: Negative round trip jitter
        Samples: 233, Minimum: 1 usec, Maximum: 10170 usec, Average: 171 usec, Peak to peak:
10169 usec, Stddev: 888 usec, Sum: 39889 usec
```

show services rpm probe-results owner terse

```
user@host> show services rpm probe-results owner owner1 terse
Test Name      Source-IF            Target Address      Status       Last Change
t_1            ms-2/2/0.10          192.0.2.1           UP           0D0H1M29S
t_2            ms-2/2/0.10          192.0.2.1           UP           0D0H1M29S
t_3            ms-2/2/0.10          192.0.2.1           UP           0D0H1M29S
```

show services rpm probe-results owner status fail

```
user@host> show services rpm probe-results owner owner1 status fail
Test Name               Last Probe Status         Status
t_1                     FAIL                      DOWN
t_2                     FAIL                      DOWN
t_3                     FAIL                      DOWN
```

show services rpm probe-results (BGP Neighbor Discovery)

```
user@host> show services rpm probe-results
Owner: Rpm-Bgp-Owner, Test: Rpm-Bgp-Test-1
    Target address: 10.209.152.37, Probe type: icmp-ping, Test size: 5 probes
    Routing Instance Name: LS1/RI1
    Probe results:
```

```
    Response received
    Probe sent time: Fri Oct 28 05:20:23 2005
    Probe rcvd/timeout time: Fri Oct 28 05:20:23 2005
    Rtt: 662 usec
  Results over current test:
    Probes sent: 5, Probes received: 5, Loss percentage: 0
    Measurement: Round trip time
      Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
      Jitter: 133 usec, Stddev: 53 usec
  Results over all tests:
    Probes sent: 5, Probes received: 5, Loss percentage: 0
    Measurement: Round trip time
      Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
      Jitter: 133 usec, Stddev: 53 usec
```

Release Information

Command introduced before Junos OS Release 7.4.

`dst-interface`, `limit`, `source-address`, `status`, `target-address`, and `terse` options introduced in Junos OS Release 18.1R1 on MX Series.

show services rpm rfc2544-benchmarking

IN THIS SECTION

- Syntax | 25961
- Description | 25961
- Options | 25961
- Required Privilege Level | 25962
- Output Fields | 25962
- Sample Output | 25963
- Release Information | 25969

Syntax

```
show services rpm rfc2544-benchmarking
<aborted-tests (test-id [test-id] | brief | detail)>
<active-tests (test-id [test-id] | brief | detail)>
<completed-tests (test-id [test-id] | brief | detail)>
<summary>
```

Description

Display information about the results of each category or state of the RFC 2544-based benchmarking test, such as terminated tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance. You can view the results of each test state for all of the configured test IDs or for a specific test ID. Also, you can display statistics about the total number of tests of each state for a high-level, quick analysis. The values in the output displayed vary, depending on the state in which the test is passing through, when you issue the command.

You can view the test results of multiple test IDs at the same time by entering the IDs in a single command. If you enter multiple test ID values, you must separate each number with a space.

Options

none Display test results for all categories.

aborted-tests (Optional) Display the list of tests that were terminated or stopped. This list includes tests that failed due to various error conditions and tests that you terminated by entering the `test service rpm rfc2544-benchmarking test` *`test-name`* `stop` command. The `Status` field in the output specifies the reason for the termination of the test.

active-tests (Optional) Display the results of the set of tests that are currently running.

brief | detail | extensive Display the specified level of output.

completed-tests (Optional) Display the results of the set of tests that were successfully completed. A completed test is one that passes through all the test steps or states specified in RFC 2544. A test that is marked as completed after it went through all the states from the beginning to the end can still be reported as a failed test. For example, a failed test can be a test that sends the desired number of packets, but does not receive the frames back from the other end.

summary (Optional) Display summary output.

test-id *test-id* Display test results for the specified unique identifier of the test.

Required Privilege Level

view

Output Fields

Table 1484 on page 25962 lists the output fields for the `show services rpm rfc2544-benchmarking` command. Output fields are listed in the approximate order in which they appear.

Table 1484: show services rpm rfc2544-benchmarking Output Fields

Field Name	Field Description
`Test information`	Details of the performed RFC 2544 benchmarking test.
`Test id`	Unique identifier configured for the test.
`Test name`	Name configured for the test.
`Test type`	The type of statistical detail that is collected for the test, based on the configured test type. Throughput-related, latency, frame-loss, or back-to-back frames-related information is displayed for ACX Series routers. Reflected packets-related information is displayed for MX104 Series routers.
`Test mode`	Mode configured for the test on the router. Test modes are: • Initiate-and-Terminate: Test frames are initiated from one end and terminated at the same end. This mode requires a reflector to be configured at the peer end to enable the test frames to be returned to the source. This mode is supported only on ACX Series routers. • Reflect: Test frames that originate from one end are reflected at the other end on the selected service, such as IPv4 or Ethernet.
`Test packet size`	Size of the test packets in bytes. This field is valid only when the test mode is Initiate-and-Terminate.

Table 1484: show services rpm rfc2544-benchmarking Output Fields *(Continued)*

Field Name	Field Description
`Test state`	State of the test that is in progress or active when the output is displayed.
`Status`	Indicates whether the test is currently in progress or has been terminated. This field is displayed for tests that are in progress or were terminated by entering the `test services rpm rfc2544-benchmarking test <test-name | test-id> stop` command.
`Test start time`	Time at which the test started in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS).
`Test finish time`	Time at which the test completed.
`Counters last cleared`	Date, time, and how long ago the statistics for the test were cleared. The format is *year-month-day hour:minute:second:timezone* (*hour:minute:second* ago). For example, 2010-05-17 07:51:28 PDT (00:04:33 ago). If you did not clear the statistics previously at any point, `Never` is displayed.
`Number of active tests`	Total number of tests that are currently running.
`Number of completed tests`	Total number of tests that were successfully completed
`Number of aborted tests`	Total number of tests that were terminated or halted.

Sample Output

show services rpm rfc2544-benchmarking summary

```
user@host> show services rpm rfc2544-benchmarking summary
  Tests summary :
```

```
      Number of active tests: 0, Number of completed tests: 4,
      Number of aborted tests: 52
```

This output indicates that no test iteration is currently in progress (at the time of issue of the command), 4 tests were completed successfully, and 52 tests were halted.

show services rpm rfc2544-benchmarking aborted-tests (ACX Series Router)

```
user@host> show services rpm rfc2544-benchmarking aborted-tests
Test information :
      Test id: 1, Test name: test1, Test type: Throughput
      Test mode: Initiate-and-Terminate
      Test packet size: 64 1280
      Test state: RFC2544_TEST_STATE_STOPPED
      Status: User-aborted-via-cli
      Test start time: 2005-08-05 03:19:58 UTC
      Test finish time: 2005-08-05 03:20:00 UTC
      Counters last cleared: Never

      Test id: 2, Test name: test1, Test type: Throughput
      Test mode: Initiate-and-Terminate
      Test packet size: 64 1280
      Test state: RFC2544_TEST_STATE_STOPPED
      Status: User-aborted-via-cli
      Test start time: 2005-08-05 03:20:00 UTC
      Test finish time: 2005-08-05 03:20:02 UTC
      Counters last cleared: Never
```

show services rpm rfc2544-benchmarking completed-tests (ACX Series Router)

```
user@host> show services rpm rfc2544-benchmarking completed-tests
 Test information :
      Test id: 18, Test name: test1, Test type: Throughput
      Test mode: Initiate-and-Terminate
      Test packet size: 64 1280
      Test state: RFC2544_TEST_STATE_COMPLETED
      Test start time: 2005-08-05 03:20:34 UTC
      Test finish time: 2005-08-05 03:21:23 UTC
      Counters last cleared: Never
```

show services rpm rfc2544-benchmarking active-tests (ACX Series Router)

```
user@host> show services rpm rfc2544-benchmarking active-tests
Test information :
      Test id: 57, Test name: test1, Test type: Back-Back-Frames
      Test mode: Initiate-and-Terminate
      Test packet size: 64 1280
      Test state: RFC2544_TEST_STATE_RUNNING
      Status: Running
      Test start time: 2005-08-05 20:15:41 UTC
      Test finish time: TEST_RUNNING
      Counters last cleared: Never
```

show services rpm rfc2544-benchmarking aborted-tests (MX104 Router)

```
user@host> show services rpm rfc2544-benchmarking aborted-tests
  Test information :
      Test id: 1, Test name: prof_tput1, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: TEST_STATE_STOPPED
      Status: Test-intf-ifl-change
      Test start time: 2013-12-16 22:54:27 PST
      Test finish time: 2013-12-16 23:30:28 PST
      Counters last cleared: Never

      Test id: 2, Test name: prof_tput1, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: TEST_STATE_STOPPED
      Status: User-aborted-via-cli
      Test start time: 2013-12-16 23:31:06 PST
      Test finish time: 2013-12-16 23:36:22 PST
      Counters last cleared: Never

      Test id: 3, Test name: prof_tput1, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: TEST_STATE_STOPPED
```

```
      Status: User-aborted-via-cli
      Test start time: 2013-12-16 23:36:24 PST
      Test finish time: 2013-12-17 01:49:24 PST
      Counters last cleared: Never
```

show services rpm rfc2544-benchmarking completed-tests (MX104 Router)

```
user@host> show services rpm rfc2544-benchmarking completed-tests
 Test information :
      Test id: 18, Test name: test1, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: TEST_STATE_COMPLETED
      Test start time: 2005-08-05 03:20:34 UTC
      Test finish time: 2005-08-05 03:21:23 UTC
      Counters last cleared: Never
```

show services rpm rfc2544-benchmarking active-tests (MX104 Router)

```
user@host> show services rpm rfc2544-benchmarking active-tests
  Test information :
      Test id: 4, Test name: prof_tput1, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: TEST_STATE_RUNNING
      Status: Running
      Test start time: 2013-12-17 01:49:26 PST
      Test finish time: TEST_RUNNING
      Counters last cleared: Never
```

show services rpm rfc2544-benchmarking aborted-tests (SRX300 and SRX550HM)

```
user@host> show services rpm rfc2544-benchmarking aborted-tests
  Test information :
      Test id: 5, Test name: ts14, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: TEST_STATE_STOPPED
```

```
        Status: User-aborted-via-cli
        Test start time: 2020-07-13 16:19:31 CST
        Test finish time: 2020-07-13 16:19:37 CST
        Counters last cleared: Never
```

show services rpm rfc2544-benchmarking active-tests (SRX300 and SRX550HM)

```
user@host> show services rpm rfc2544-benchmarking active-tests
  Test information :
      Test id: 1, Test name: ts1, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: TEST_STATE_RUNNING
      Status: Running
      Test start time: 2020-06-15 12:34:52 CST
      Test finish time: TEST_RUNNING
      Counters last cleared: Never

      Test id: 2, Test name: ts14, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: TEST_STATE_RUNNING
      Status: Running
      Test start time: 2020-07-08 20:46:35 CST
      Test finish time: TEST_RUNNING
      Counters last cleared: Never
```

show services rpm rfc2544-benchmarking completed-tests (SRX300 and SRX550HM)

```
user@host> show services rpm rfc2544-benchmarking completed-tests
  Test information :
      Test id: 6, Test name: ts1, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: TEST_STATE_COMPLETED
      Status: Test-Completed
      Test start time: 2020-07-13 16:28:13 CST
      Test finish time: 2020-07-13 16:30:17 CST
      Counters last cleared: Never
```

show services rpm rfc2544-benchmarking aborted-tests detail (SRX300 and SRX550HM)

```
user@host> show services rpm rfc2544-benchmarking aborted-tests detail
  Test information :
      Test id: 1, Test name: ts13, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: TEST_STATE_STOPPED
      Status: User-aborted-via-cli
      Test start time: 2020-07-10 12:38:26 CST
      Test finish time: 2020-07-10 12:49:59 CST
      Counters last cleared: Never

      Test Configuration:
          Test mode: Reflect
          Duration in seconds: 864000
          Test finish wait duration in seconds: 1
          Test family: INET
          Test iterator pass threshold: 0.50 %
          Test receive failure threshold: 0.00 %
          Test transmit failure threshold: 0.50 %
          Routing Instance Name: default

      Inet family Configuration:
          Egress Interface : ge-0/0/13.0
          Destination ipv4 address: 10.0.2.1
          Destination udp port: 400

Elapsed        Reflected            Reflected
 time           Packets               Bytes
 692               0                     0

  Test information :
      Test id: 5, Test name: ts14, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: TEST_STATE_STOPPED
      Status: User-aborted-via-cli
      Test start time: 2020-07-13 16:19:31 CST
      Test finish time: 2020-07-13 16:19:37 CST
      Counters last cleared: Never
```

```
        Test Configuration:
            Test mode: Reflect
            Duration in seconds: 864000
            Test finish wait duration in seconds: 1
            Test family: INET
            Test iterator pass threshold: 0.50 %
            Test receive failure threshold: 0.00 %
            Test transmit failure threshold: 0.50 %
            Routing Instance Name: default

        Inet family Configuration:
            Egress Interface : ge-0/0/14.0
            Destination ipv4 address: 10.0.3.1
            Destination udp port: 400

 Elapsed        Reflected            Reflected
  time           Packets               Bytes
    5               0                    0
```

Release Information

Command introduced in Junos OS Release 12.3X52.

RELATED DOCUMENTATION

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

rfc2544-benchmarking | **15759**

show services rpm rfc2544-benchmarking test-id

IN THIS SECTION

- Syntax | **25970**
- Description | **25970**

- Options | 25970
- Required Privilege Level | 25970
- Output Fields | 25970
- Sample Output | 25983
- Release Information | 25994

Syntax

```
show services rpm rfc2544-benchmarking test-id test-id
<brief | detail>
```

Description

Display information about the results of the RFC 2544-based benchmarking test for a specific test ID for each real-time performance monitoring (RPM) instance. The values in the output displayed vary, depending on the state in which the test is passing through, when you issue the command.

Options

test-id	Display test results for the specified unique identifier.
brief | detail	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1485 on page 25971 lists the output fields for the `show services rpm rfc2544-benchmarking test-id` command. Output fields are listed in the approximate order in which they appear.

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields

Field Name	Field Description	Level of Output
`Test information`	Details of the performed RFC 2544 benchmarking test.	None specified
`Test id`	Unique identifier configured for the test.	None specified
`Test name`	Name configured for the test.	None specified
`Test type`	The type of actual test run that is collected for the test, based on the configured test type. Throughput-related, latency, frame-loss, or back-to-back frames-related information is displayed for ACX Series routers. Reflected packets-related information is displayed for MX104 Series routers.	None specified
`Test mode`	Mode configured for the test on the router. Test modes are: • Initiate-and-Terminate: Test frames are initiated from one end and terminated at the same end. This mode requires a reflector to be configured at the peer end to enable the test frames to be returned to the source. This mode is supported only on ACX Series routers. • Reflect: Test frames that originate from one end are reflected back to the originator, such as IPv4 or Ethernet.	None specified
`Test packet size`	Size of the test packets in bytes. This field is valid only when the test mode is Initiate-and-Terminate.	None specified
`Test state`	State of the test that is in progress or active when the output is displayed. For details about the states, see *RFC 2544-Based Benchmarking Test States*.	None specified
Status	Indicates whether the test is currently in progress or has been terminated.	None specified
`Test start time`	Time at which the test started in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS).	None specified

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Test finish time`	Time at which the test completed.	None specified
`Counters last cleared`	Date, time, and how long ago the statistics for the test were cleared. The format is *year-month-day hour:minute:second:timezone* (*hour:minute:second* ago). For example, 2010-05-17 07:51:28 PDT (00:04:33 ago). If you did not clear the statistics previously at any point, `Never` is displayed.	None specified
`Test-profile Configuration`	(ACX Series routers only) Details of the specified test profile	`detail`
`Test-profile name`	(ACX Series routers only) Name of the configured test profile that contains the parameters for the test	`detail`
`Test packet size`	(ACX Series routers only) Size of the test packets in bytes	`detail`
`Theoretical max bandwidth`	(ACX Series routers only) Theoretical maximum bandwidth configured for the test. This value is typically set to the bandwidth of the server being tested. Valid values are 1 Kbps through 1,000,000 Kbps (1 Gbps). The value defined is the highest bandwidth value tested for this test.	`detail`
`Test Configuration`	Details of the configured test ID.	`detail`
`Test mode`	Mode configured for the test. Test modes are Initiate-and-Terminate and Reflect.	`detail`
`Duration in seconds`	Period in seconds for which the test has been performed.	`detail`

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Test family`	The underlying service on which the test is run. Test families are: • INET: Indicates that the test is run on a IPV4 service. • CCC: Indicates that the test is run on a circuit cross-connect (CCC) or pseudowire service.	`detail`
`Routing Instance Name`	(ACX Series routers only) Name of the routing instance for the test	`detail`
`Inet family Configuration`	Details of the configured `inet` family for an IPv4 service	`detail`
`Egress Interface`	Name of the egress interface from which the test frames are sent	`detail`
`Source ipv4 address`	Source IPv4 address used in the IP header of the generated test frame.	`detail`
`Destination ipv4 address`	Destination IPv4 address used in the IP header of the generated test frame.	`detail`
`Source udp port`	Source UDP port number used in the UDP header of the generated test frame.	`detail`
`Destination udp port`	Destination UDP port number used in the UDP header of the generated test frame.	`detail`
`Ccc family Configuration`	Details of the configured CCC family for an Ethernet service	`detail`
`Source MAC address`	(ACX Series routers only) Source MAC address used in generated test frames for a CCC or Ethernet pseudowire service.	`detail`

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Destination MAC address	(ACX Series routers only) Destination MAC address used in generated test frames for a CCC or Ethernet pseudowire service.	detail
Ivlan-id	(ACX Series routers only) Inner VLAN ID for test-frames.	detail
Ovlan-id	(ACX Series routers only) Outer VLAN ID for test-frames.	detail
Direction egress	Test is run in the egress direction of the interface (NNI)	detail
Direction ingress	Test is run in the ingress direction of the interface (UNI)	detail
Rfc2544 throughput test information	**(ACX Series routers only) Details of the throughput test**	**detail**
Initial test load percentage	Percentage of the steady state load for the test.	detail
Test iteration mode	Mode of the test iteration: Binary or step-down.	detail
Test iteration step percent	The test step percentage for tests. If not specified, the default step-percent is 10 percent. This parameter is ignored for all type of tests other than frame-loss tests.	detail
Theoretical max bandwidth	The theoretical limit of the media for the frame size configured for the test. This value is typically set to the bandwidth of the server being tested.	detail
Test packet size:	Packet size of the test frames in bytes.	detail
Iteration	Number of the test iteration.	detail

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Duration (sec)`	Period in seconds for which the test iteration is run	`detail`
`Elapsed time`	Amount of time that has passed, in seconds, since the start of the test.	`detail`
`pps`	Total count of packets-per-second (pps) transmitted during the test.	`detail`
`Tx Packets`	Number of transmitted test packets.	`detail`
`Rx Packets`	Number of received test packets.	`detail`
`Tx Bytes`	Number of transmitted bytes.	`detail`
`Rx Bytes`	Number of received bytes.	`detail`
`Percentage throughput`	Percentage of throughput for the test iteration.	`detail`
`Result of the iteration runs (Throughput) :`	Results of the completed throughput test for a particular packet size.	`detail`
`Best iteration`	Number of the iteration with the highest throughout, among the listed iterations.	`detail`
`Best iteration (pps)`	Packets-per-second (pps) count of the iteration with the highest throughout, among the listed iterations.	`detail`
`Best iteration throughput`	Percentage of throughput of the iteration with the highest throughout, among the listed iterations.	`detail`

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
RFC2544 Throughput test results summary	Consolidated information of the throughput test.	detail summary
Packet Size	Size of the test packet in bytes.	detail summary
Theoretical rate (pps)	Theoretical frame rate in packets-per-second.	detail summary
Tx Packets	Number of transmitted packets.	detail summary
Rx Packets	Number of received packets.	detail summary
Offered throughput (percentage)	The offered throughput in percentage of the chosen service (such as Layer 3 or Ethernet pseudowire).	detail summary
Measured bandwidth (kbps)	Available bandwidth of the service based on the calculated throughput.	detail summary
Rfc2544 latency test information :	**(ACX Series routers only) Details of the latency test**	**detail**
Theoretical max bandwidth	Theoretical maximum bandwidth configured for the test. This value is typically set to the bandwidth of the server being tested. Valid values are 1 Kbps through 1,000,000 Kbps (1 Gbps). The value defined is the highest bandwidth value used for this test.	detail
Initial test load percentage	Percentage of the steady state load for the test.	detail
Duration in seconds	Period in seconds for which the test has been performed.	detail

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Test packet size`	Size of the test packet in bytes.	`detail`
`Iteration`	Number of the test iteration.	`detail`
`Duration (sec)`	Period in seconds for which the test iteration is run.	`detail`
`Elapsed time`	Amount of time that has passed, in seconds, since the start of the test.	`detail`
`pps`	Total count of packets-per-second (pps) transmitted during the test.	`detail`
`Tx Packets`	Number of transmitted test packets.	`detail`
`Rx Packets`	Number of received test packets.	`detail`
`Latency`	Displays the latency parameters.	`detail`
`Min(ns)`	Aggregated minimum latency in nanoseconds.	`detail`
`Avg(ns)`	Aggregated average latency in nanoseconds.	`detail`
`Max(ns)`	Aggregated maximum latency in nanoseconds.	`detail`
`Probe(ns)`	Aggregated probe latency in nanoseconds.	`detail`
`Result of the iteration runs (Latency)`	Results of the latency test completed for a particular packet size.	`detail`
`Avg (min) Latency`	Average of the minimum latency in nanoseconds.	`detail`

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Avg (avg) latency`	Average of the average latency in nanoseconds.	`detail`
`Avg (Max) latency`	Average of the maximum latency in nanoseconds.	`detail`
`Avg (probe) latency`	Average of the probe latency in nanoseconds.	`detail`
`RFC2544 Latency test results summary:`	Consolidated statistics of the latency test.	`detail summary`
`Packet Size`	Size of the test packet in bytes.	`detail summary`
`Theoretical rate (pps)`	Theoretical frame rate in packets-per-second.	`detail summary`
`Tx Packets`	Number of transmitted packets.	`detail summary`
`Rx Packets`	Number of received packets.	`detail summary`
`Latency`	Displays the latency parameters.	`detail summary`
`Min(ns)`	Aggregated minimum latency in nanoseconds.	`detail summary`
`Avg(ns)`	Aggregated average latency in nanoseconds.	`detail summary`
`Max(ns)`	Aggregated maximum latency in nanoseconds.	`detail summary`
`Probe(ns)`	Aggregated probe latency in nanoseconds.	`detail summary`

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Rfc2544 Back-Back test information :	**(ACX Series routers only) Details of the back-to-back frames or bursty frames test.**	**detail**
`Initial burst length:`	Length of the first burst when test frames are sent, as a measure of number of seconds at the rate of Kbps.	`detail`
`Test iteration mode :`	Mode of the test iteration: Binary or step-down.	`detail`
`Test iteration step percent`	The test step percentage for tests. If not specified, the default step-percent is 10 percent. This parameter is ignored for all type of tests other than frame-loss tests.	`detail`
`Theoretical max bandwidth`	The theoretical limit of the media for the frame size configured for the test. This value is typically set to the bandwidth of the server being tested.	`detail`
`Test packet size:`	Packet size of the test frames in bytes.	`detail`
`Iteration`	Number of the test iteration.	`detail`
`Burst Length (Packets)`	Number of packets in the burst.	`detail`
`Elapsed time`	Amount of time that has passed, in seconds, since the start of the test.	`detail`
`Tx Packets`	Number of transmitted test packets.	`detail`
`Rx Packets`	Number of received test packets.	`detail`
`Tx Bytes`	Number of transmitted bytes.	`detail`

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Rx Bytes`	Number of received bytes.	`detail`
`Result of the iteration runs :`	Results of the back-to-back frames test completed for a certain packet size.	`detail`
`Best iteration :`	Number of the iteration with the longest burst.	`detail`
`Measured burst (num sec)`	Time in seconds of the burst of the iteration with the longest burst.	`detail`
`Measured burst (num pkts)`	Number of packets during the burst of the iteration with the longest burst.	`detail`
`RFC2544 Back-Back test results summary:`	Consolidated statistics of the back-to-back frames test.	`detail summary`
`Packet Size`	Size of the test packets in bytes.	`detail summary`
`Measure Burst length (Packets)`	Computed burst length in terms of number of packets.	`detail summary`
Rfc2544 frame-loss test information :	**(ACX Series routers only) Details of the frame-loss test.**	**detail**
`Initial burst length:`	Length of the first burst when test frames are sent, as a measure of number of seconds at the rate of Kbps.	`detail`
`Test iteration mode :`	Mode of the test iteration: Binary or step-down.	`detail`

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Test iteration step percent`	The test step percentage for tests. If not specified, the default step-percent is 10 percent. This parameter is ignored for all type of tests other than frame-loss tests.	`detail`
`Theoretical max bandwidth`	The theoretical limit of the media for the frame size configured for the test. This value is typically set to the bandwidth of the server being tested.	`detail`
`Test packet size`	Size of the test packets in bytes.	`detail`
`Iteration`	Number of the test iteration.	`detail`
`Duration (sec)`	Period, in seconds, for which the test iteration is run.	`detail`
`Offered throughput (percentage)`	The offered throughput in percentage of the chosen service (such as Layer 3 or Ethernet pseudowire)	`detail`
`Elapsed time`	Amount of time that has passed, in seconds, since the start of the test.	`detail`
`pps`	Theoretical frame rate in packets-per-second.	`detail`
`Tx Packets`	Number of transmitted test packets.	`detail`
`Rx Packets`	Number of received test packets.	`detail`
`Tx Bytes`	Number of transmitted bytes.	`detail`
`Rx Bytes`	Number of received bytes.	`detail`

Table 1485: show services rpm rfc2544-benchmarking test-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Frame-loss rate %`	Percentage of frames that must been forwarded by the router under steady state (constant) load, but were not forwarded due to lack of resources.	`detail`
`Result of the iteration runs :`	Results of the frame-loss test completed for a certain packet size.	`detail`
`Frame-loss rate (percent) :`	Percentage of dropped frames for the specified packet size	`detail`
`RFC2544 Frame-loss test results summary`	Consolidated statistics of the frame-loss test	`detail`
`Packet Size`	Size of the test packet in bytes.	`detail summary`
`Theoretical rate (pps)`	Theoretical frame rate in packets-per-second.	`detail summary`
`Percentage throughput`	Percentage of throughput for the test iteration.	`detail summary`
`Tx Packets`	Number of transmitted packets.	`detail summary`
`Rx Packets`	Number of received packets.	`detail summary`
`Frame Loss rate percent`	Percentage of dropped frames for the specified packet size	`detail summary`

Sample Output

show services rpm rfc2544-benchmarking test-id detail (Throughput Test on ACX Series Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 19 detail
Test information :
      Test id: 19, Test name: test1, Test type: Throughput
      Test mode: Initiate-and-Terminate
      Test packet size: 64 1280
      Test state: RFC2544_TEST_STATE_COMPLETED
      Test start time: 2005-07-29 10:25:00 UTC
      Test finish time: 2005-07-29 10:26:02 UTC
      Counters last cleared: Never

      Test-profile Configuration:
          Test-profile name: prof_tput
          Test packet size: 64 1280
          Therotical max bandwidth : 993000 kbps

      Test Configuration:
          Test mode: Initiate-and-Terminate
          Duration in seconds: 20
          Test family: INET
          Routing Instance Name: default

      Inet family Configuration:
          Egress Interface : ge-0/1/1.0
          Source ipv4 address: 192.0.2.1
          Destination ipv4 address: 192.0.2.2
          Source udp port: 2020
          Destination udp port: 3030

  Rfc2544 throughput test information :
      Initial test load percentage : 100.00 %
      Test iteration mode : Binary
      Test iteration step percent : 50.00 %
      Therotical max bandwidth : 993000 kbps

Test packet size: 64
Iteration Duration Elapsed pps      Tx       Rx       Tx         Rx         Percentage
           (sec)    time            Packets  Packets  Bytes      Bytes      throughput
```

```
   1       3        3      134918   404754   404754   27523272   27523272   10.00 %
   2       20       20     1349184  26983501 26983501 1834878068 1834878068 100.00 %

Result of the iteration runs : Throughput Test complete for packet size 64
Best iteration : 2, Best iteration (pps) : 1349184
Best iteration throughput : 100.00 %

Test packet size: 1280
Iteration Duration Elapsed pps    Tx       Rx       Tx         Rx         Percentage
           (sec)    time          Packets  Packets  Bytes      Bytes      throughput
   1       3        3      9489   28467    28467    36551628   36551628   10.00 %
   2       20       20     94896  1897920  1897920  2436929280 2436929280 100.00 %

Result of the iteration runs : Throughput Test complete for packet size 1280
Best iteration : 2, Best iteration (pps) : 94896
Best iteration throughput : 100.00 %

RFC2544 Throughput test results summary:
----------------------------------------

  Packet  Theoretical   Tx         Rx        Offered throughput   Measured
   Size   rate (pps)    Packets    Packets       (percentage)     bandwidth (kbps)
   64     1349184       26983501   26983501        100.00 %       993000
   1280   94896         1897920    1897920         100.00 %       993000
```

show services rpm rfc2544-benchmarking test-id detail (Latency Test on ACX Series Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 37 detail
  Test information :
      Test id: 37, Test name: test1, Test type: Latency
      Test mode: Initiate-and-Terminate
      Test packet size: 64 1280
      Test state: RFC2544_TEST_STATE_COMPLETED
      Test start time: 2005-07-29 10:26:41 UTC
      Test finish time: 2005-07-29 10:36:15 UTC
      Counters last cleared: Never

      Test-profile Configuration:
          Test-profile name: prof_latency
          Test packet size: 64 1280
```

```
            Therotical max bandwidth : 993000 kbps

        Test Configuration:
            Test mode: Initiate-and-Terminate
            Duration in seconds: 10
            Test family: INET
            Routing Instance Name: default

        Inet family Configuration:
            Egress Interface : ge-0/1/1.0
            Source ipv4 address: 192.0.2.1
            Destination ipv4 address: 192.0.2.2
            Source udp port: 2020
            Destination udp port: 3030

   Rfc2544 latency test information :
        Therotical max bandwidth : 993000 kbps
        Initial test load percentage : 100.00 %
        Duration in seconds: 10

Test packet size: 64

Iteration  Duration   Elapsed     pps            Tx           Rx
            (sec)      time                    Packets      Packets
     1           3         3       134918       404754       404754
     2          10        10      1349184     13491751     13491751
     3          10        10      1349184     13491751     13491751
     4          10        10      1349184     13491751     13491751
     5          10        10      1349184     13491751     13491751
     6          10        10      1349184     13491751     13491751
     7          10        10      1349184     13491751     13491751
     8          10        10      1349184     13491751     13491751
     9          10        10      1349184     13491751     13491751
    10          10        10      1349184     13491751     13491751
    11          10        10      1349184     13491751     13491751
    12          10        10      1349184     13491751     13491751
    13          10        10      1349184     13491751     13491751
    14          10        10      1349184     13491751     13491751
    15          10        10      1349184     13491751     13491751
    16          10        10      1349184     13491751     13491751
    17          10        10      1349184     13491751     13491751
    18          10        10      1349184     13491751     13491751
    19          10        10      1349184     13491751     13491751
```

```
    20          10        10      1349184     13491751      13491751
    21          10        10      1349184     13491751      13491751

-------------- Latency ---------------
Min(ns)    Avg(ns)    Max(ns)    Probe(ns)
17464       18770      18880        18784
17472       18799      20488        18848
17472       18799      20416        18816
17472       18799      20440        18704
17464       18799      20376        18880
17464       18799      20232        18832
17464       18799      20400        18848
17472       18799      20240        18864
17472       18799      20264        18848
17464       18799      20264        18880
17472       18800      20320        18864
17464       18799      20176        18864
17464       18800      20248        18864
17464       18800      20272        18864
17464       18799      20472        18832
17464       18799      20256        18880
17464       18799      20336        18848
17464       18800      20688        18848
17472       18800      20504        18864
17464       18799      20448        18768
17472       18799      20240        18864

Result of the iteration runs : Latency Test complete for packet size 64
Avg (min) Latency                  : 17466
Avg (avg) latency                  : 18799
Avg (Max) latency                  : 20360
Avg (probe) latency                : 18844

Test packet size: 1280

Iteration  Duration   Elapsed     pps            Tx           Rx
            (sec)      time                    Packets       Packets
     1           3         3         9489        28467         28467
     2          10        10        94896       948960        948960
     3          10        10        94896       948960        948960
     4          10        10        94896       948960        948960
     5          10        10        94896       948960        948960
     6          10        10        94896       948960        948960
```

```
     7          10        10        94896       948960        948960
     8          10        10        94896       948960        948960
     9          10        10        94896       948960        948960
    10          10        10        94896       948960        948960
    11          10        10        94896       948960        948960
    12          10        10        94896       948960        948960
    13          10        10        94896       948960        948960
    14          10        10        94896       948960        948960
    15          10        10        94896       948960        948960
    16          10        10        94896       948960        948960
    17          10        10        94896       948960        948960
    18          10        10        94896       948960        948960
    19          10        10        94896       948960        948960
    20          10        10        94896       948960        948960
    21          10        10        94896       948960        948960

-------------- Latency ---------------
Min(ns)    Avg(ns)    Max(ns)    Probe(ns)
68712       70031      70576        69456
68728       70344      71808        70512
68720       70344      71744        70352
68720       70344      71680        70112
68720       70345      71856        70352
68720       70344      71808        70384
68720       70344      71752        70480
68720       70344      71880        70112
68720       70344      71792        70320
68728       70345      73344        70336
68720       70344      71688        70560
68728       70345      71896        70496
68720       70344      71760        70096
68720       70344      71776        70320
68720       70344      71760        70400
68712       70345      71920        70352
68720       70344      71792        70576
68720       70345      71840        70320
68720       70344      71792        70368
68720       70345      71824        70464
68712       70345      71904        70512


Result of the iteration runs : Latency Test complete for packet size 1280
Avg (min) Latency                  : 68720
```

```
Avg (avg) latency                  : 70344
Avg (Max) latency                  : 71880
Avg (probe) latency                : 70371

RFC2544 Latency test results summary:
---------------------------------------

  Packet  Theoretical Tx        Rx        ------------- Latency -------------
   Size   rate (pps)  Packets   Packets   Min(ns)  Avg(ns)  Max(ns) Probe(ns)
   64     1349184     269835020 269835020 17466    18799    20360    18844
   1280   94896       18979200  18979200  68720    70344    71880    70371
```

show services rpm rfc2544-benchmarking test-id detail (Frame Loss Test on ACX Series Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 73 detail
  Test information :
       Test id: 73, Test name: test1, Test type: Frame-Loss
       Test mode: Initiate-and-Terminate
       Test packet size: 64 1280
       Test state: RFC2544_TEST_STATE_COMPLETED
       Test start time: 2005-07-29 10:38:41 UTC
       Test finish time: 2005-07-29 10:41:19 UTC
       Counters last cleared: Never

       Test-profile Configuration:
           Test-profile name: prof_fl
           Test packet size: 64 1280
           Therotical max bandwidth : 993000 kbps

       Test Configuration:
           Test mode: Initiate-and-Terminate
           Duration in seconds: 20
           Test family: INET
           Routing Instance Name: default

       Inet family Configuration:
           Egress Interface : ge-0/1/1.0
           Source ipv4 address: 192.0.2.1
           Destination ipv4 address: 192.0.2.2
```

```
            Source udp port: 2020
            Destination udp port: 3030

   Rfc2544 frame-loss test information :
       Initial test load percentage : 100.00 %
       Test iteration mode : step-down
       Test iteration step percent : 10 %
       Therotical max bandwidth : 993000 kbps

 Test packet size: 64
 Iteration Duration  Elapsed  Offered     pps    Tx       Rx       Tx         Rx        Frame-
loss
             (sec)  time    throughput%         Packets  Packets  Bytes      Bytes      rate %
  1       3         3       10.00 %     134918  404754   404754   27523272   27523272   0.00 %
  2       20        20     100.00 %     1349184 26983501 26983501 1834878068 1834878068 0.00 %
  3       20        20     100.00 %     1349184 26983501 26983501 1834878068 1834878068 0.00 %
  4       20        20     100.00 %     1349184 26983501 26983501 1834878068 1834878068 0.00 %


 Result of the iteration runs : Frame-loss test complete for packet size 64
 Frame-loss rate (percent) : 0.00 %

 Test packet size: 1280
 Iteration Duration  Elapsed  Offered     pps    Tx       Rx       Tx         Rx        Frame-
loss
             (sec)  time    throughput%         Packets  Packets  Bytes      Bytes      rate %
  1       3         3       10.00 %     9489    404754   28467    36551628   36551628   0.00 %
  2       20        20     100.00 %     94896   1897920  1897920  2436929280 2436929280 0.00 %
  3       20        20     100.00 %     94896   1897920  1897920  2436929280 2436929280 0.00 %
  4       20        20     100.00 %     94896   1897920  1897920  2436929280 2436929280 0.00 %


 Result of the iteration runs : Frame-loss test complete for packet size 1280
 Frame-loss rate (percent) : 0.00 %

 RFC2544 Frame-loss test results summary:
 ----------------------------------------

   Packet      Theoretical     Percentage      Tx             Rx           Frame Loss
    Size       rate (pps)      throughput    Packets        Packets       rate percent
       64         1349184       100.00 %      26983501       26983501        0.00 %
     1280           94896       100.00 %       1897920        1897920        0.00 %
```

show services rpm rfc2544-benchmarking test-id detail (Back-to-Back Frames Test on ACX Series Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 55 detail
Test information :
       Test id: 55, Test name: test1, Test type: Back-Back-Frames
       Test mode: Initiate-and-Terminate
       Test packet size: 64 1280
       Test state: RFC2544_TEST_STATE_COMPLETED
       Test start time: 2005-07-29 10:36:54 UTC
       Test finish time: 2005-07-29 10:37:57 UTC
       Counters last cleared: Never

       Test-profile Configuration:
           Test-profile name: prof_b2b
           Test packet size: 64 1280
           Therotical max bandwidth : 993000 kbps

       Test Configuration:
           Test mode: Initiate-and-Terminate
           Duration in seconds: 20
           Test family: INET
           Routing Instance Name: default

       Inet family Configuration:
           Egress Interface : ge-0/1/1.0
           Source ipv4 address: 192.0.2.1
           Destination ipv4 address: 192.0.2.2
           Source udp port: 2020
           Destination udp port: 3030

   Rfc2544 Back-Back test information :
       Initial burst length: 20 seconds at 993000 kbps
       Test iteration mode : Binary
       Test iteration step percent : 50.00 %

 Test packet size: 64
 Iteration   Burst Length    Elapsed         Tx           Rx              Tx             Rx
```

```
              (Packets)       time        Packets       Packets         Bytes          Bytes
    1           404754            3        404754        404754       27523272       27523272
    2         26983680           20      26983680      26983680     1834890240     1834890240

Result of the iteration runs : Back-Back-Frames Test complete for packet size 64
Best iteration : 2
Measured burst (num sec)  : 20 sec,
Measured burst (num pkts)  : 26983680 packets
Result of the iteration runs : Back-Back-Frames Test complete for packet size 64
Best iteration : 2
Measured burst (num sec)  : 20 sec,
Measured burst (num pkts)  : 26983680 packets

Test packet size: 1280
Iteration   Burst Length    Elapsed         Tx           Rx              Tx             Rx
              (Packets)       time        Packets       Packets         Bytes          Bytes
    1            28467            3         28467         28467       36551628       36551628
    2          1897920           20       1897920       1897920     2436929280     2436929280

Result of the iteration runs : Back-Back-Frames Test complete for packet size 12
Best iteration : 2
Measured burst (num sec)  : 20 sec,
Measured burst (num pkts)  : 1897920 packets

RFC2544 Back-Back test results summary:
----------------------------------------

  Packet      Measure Burst
   Size      length (Packets)
     64       26983680 packets
   1280       1897920 packets
```

show services rpm rfc2544-benchmarking test-id detail (Reflection Test on MX104 Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 1 detail
  Test information :
      Test id: 1, Test name: fort_uni_inet_ref, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: RFC2544_TEST_STATE_RUNNING
```

```
      Status: Running
      Test start time: 2013-12-09 16:24:52 IST
      Test finish time: TEST_RUNNING
      Counters last cleared: Never

      Test Configuration:
          Test mode: Reflect
          Duration in seconds: 864000
          Test family: INET
          Routing Instance Name: default

      Inet family Configuration:
          Egress Interface : ge-0/3/1.0
          Destination ipv4 address: 198.51.100.2
          Destination udp port: 200

Elapsed        Reflected            Reflected
 time           Packets               Bytes
    176         8977917            9031784502
```

show services rpm rfc2544-benchmarking test-id brief (Reflection Test on MX104 Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 1 brief
  Test information :
      Test id: 1, Test name: fort_uni_inet_ref, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: RFC2544_TEST_STATE_RUNNING
      Status: Running
      Test start time: 2013-12-09 16:24:52 IST
      Test finish time: TEST_RUNNING
      Counters last cleared: Never
```

show services rpm rfc2544-benchmarking test-id detail (Reflection Test on MX104 Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 2 detail
  Test information :
      Test id: 2, Test name: fort_uni_inet_ref, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
```

```
        Test state: RFC2544_TEST_STATE_RUNNING
        Status: Running
        Test start time: 2013-12-09 16:39:18 IST
        Test finish time: TEST_RUNNING
        Counters last cleared: Never

        Test Configuration:
            Test mode: Reflect
            Duration in seconds: 864000
            Test family: CCC
            Routing Instance Name: default

        CCC family Configuration:
            Interface : ge-0/3/2.0
            Test direction: Egress

Elapsed        Reflected            Reflected
 time           Packets               Bytes
     23          809137             825319740
```

show services rpm rfc2544-benchmarking test-id brief (Reflection Test on MX104 Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 2 brief
  Test information :
      Test id: 2, Test name: fort_uni_inet_ref, Test type: Reflect
      Test mode: Reflect
      Test packet size: 0
      Test state: RFC2544_TEST_STATE_RUNNING
      Status: Running
      Test start time: 2013-12-09 16:39:18 IST
      Test finish time: TEST_RUNNING
      Counters last cleared: Never
```

show services rpm rfc2544-benchmarking test-id detail (Reflection Test on SRX300 and SRX550HM)

```
user@host> show services rpm rfc2544-benchmarking test-id 1 detail
  Test information :
      Test id: 1, Test name: ts13, Test type: Reflect
```

```
        Test mode: Reflect
        Test packet size: 0
        Test state: TEST_STATE_RUNNING
        Status: Running
        Test start time: 2020-07-08 19:53:23 CST
        Test finish time: TEST_RUNNING
        Counters last cleared: Never

        Test Configuration:
            Test mode: Reflect
            Duration in seconds: 864000
            Test finish wait duration in seconds: 1
            Test family: INET
            Test iterator pass threshold: 0.50 %
            Test receive failure threshold: 0.00 %
            Test transmit failure threshold: 0.50 %
            Routing Instance Name: default

        Inet family Configuration:
            Egress Interface : ge-0/0/13.0
            Destination ipv4 address: 10.0.2.1
            Destination udp port: 400

 Elapsed        Reflected            Reflected
  time           Packets               Bytes
   62053       220259916          134445283434
```

Release Information

Command introduced in Junos OS Release 12.3X52.

RELATED DOCUMENTATION

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

rfc2544-benchmarking | **15759**

show services rpm twamp client connection

IN THIS SECTION

- Syntax | 25995
- Description | 25995
- Options | 25995
- Required Privilege Level | 25996
- Output Fields | 25996
- Sample Output | 25996
- Release Information | 25997

Syntax

```
show services rpm twamp client connection
<connection-name>
```

Description

Display information about the connections established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control-clients. By default, all established sessions are displayed, unless you specify a control-connection name when you issue the command. Because TWAMP light servers are stateless, information about them is not included in the output of this command; only information about managed servers is included.

Options

none	Display information about all TWAMP client connection sessions.
connection-name	(Optional) Display information about the specified control-connection or TWAMP control-client.

Required Privilege Level

view

Output Fields

Table 1486 on page 25996 lists the output fields for the `show services rpm twamp client connection` command. Output fields are listed in the approximate order in which they appear.

Table 1486: show services rpm twamp client connection Output Fields

Field Name	Field Description
Connection Name	Connection name that uniquely identifies the connection between the TWAMP server and a particular client.
Client address	Client IP address.
Client port	Client port number.
Server address	Server IP address.
Server port	Server port number.
Session count	Session count.
Auth mode	Authentication mode.

Sample Output

show services rpm twamp client connection (with a managed TWAMP server configured)

```
user@host> show services rpm twamp client connection
 Connection  Client       Client  Server          Server   Session  Auth
    ID       address       port   address          port     count   mode
     4     192.0.2.1       12345  192.168.219.203  890        16    none
    78   198.51.100.55     345    203.0.113.2      89022       5    none
```

```
   234   192.168.219.203   2345   192.168.22.2     3333       16    none
     5   192.168.3.1       82345  192.168.2.2      45909      16    authenticated
     1   192.168.1.1       645    192.168.4.23     2394       16    encrypted
```

show services rpm twamp client connection (with TWAMP light server configured)

```
user@host> show services rpm twamp client connection c1
error: Control-connection c1 does not exist.
```

Release Information

Command introduced in Junos OS Release 15.1.

show services rpm twamp client history-results

IN THIS SECTION

- Syntax | 25997
- Description | 25998
- Options | 25998
- Required Privilege Level | 25998
- Output Fields | 25998
- Sample Output | 26000
- Release Information | 26002

Syntax

```
show services rpm twamp client history-results
<brief | detail>
<control-connection control-connection-name>
<since time>
<test-session test-session-name>
```

Description

Display standard information about the results of the last 50 probes for each real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) instance. You can also view the historical results of the probes or test packets sent from a TWAMP client to a TWAMP server for a particular control-connection, or a test-session associated with a control-connection.

Options

none	Display the results of the last 50 probes for all RPM TWAMP instances.
brief | detail	(Optional) Display the specified level of output. Default: brief
control-connection *control-connection-name*	(Optional) Display information for the specified control-connection between a TWAMP client and a TWAMP server.
since *time*	(Optional) Display information from the specified time. Specify time as *yyyy-mm-dd.hh:mm:ss*.
test-session *test-session-name*	(Optional) Display information for the specified test session associated with a control-connection between a TWAMP client and a TWAMP server.

Required Privilege Level

view

Output Fields

Table 1487 on page 25998 lists the output fields for the `show services rpm twamp client history-results` command. Output fields are listed in the approximate order in which they appear.

Table 1487: show services rpm twamp client history-results Output Fields

Field Name	Field Description	Level of Output
Owner	Probe owner or the TWAMP client.	All levels
Test	Name of a test for a TWAMP probe instance.	All levels

Table 1487: show services rpm twamp client history-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Probe received	Timestamp when the probe result was determined.	All levels
Round trip time	Average ping round-trip time (RTT), in microseconds.	All levels
Probe results	Result of a particular probe performed by a remote host. The following information is contained in the results: • **Response received**—Timestamp when the probe result was determined. • **Rtt**—Average ping round-trip time (RTT), in microseconds.	**detail**
Results over current test	Displays the results for the current test by probe at the time each probe was completed, as well as the status of the current test at the time the probe was completed.	**detail**
Probes sent	Number of probes sent with the current test.	**detail**
Probes received	Number of probe responses received within the current test.	**detail**
Loss percentage	Percentage of lost probes for the current test.	**detail**

Table 1487: show services rpm twamp client history-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Measurement	Increment of measurement. Possible values are round-trip time delay and, for the probe type icmp-ping-timestamp, the egress and ingress delay: • **Minimum**—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • **Maximum**—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • **Average**—Average RTT, ingress delay, or egress delay measured over the course of the current test. • **Peak to peak**—Difference between two peak values of RTT, ingress delay, or egress delay measured over the course of the current test. • **Jitter**—Difference, in microseconds, between the maximum and minimum RTT measured over the course of the current test. • **Sum**—Total round-trip time, in microseconds, measured over the course of the current test.	**detail**

Sample Output

show services rpm twamp client history-results

```
user@host> show services rpm twamp client history-results
 Aug 02 19:11:38
 Owner, Test         Probe Sent              Probe received              Round trip time
 c1, t2              Sun Aug  2 18:26:58 2020 Sun Aug  2 18:26:58 2020            6455 usec
 c1, t2              Sun Aug  2 18:26:59 2020 Sun Aug  2 18:26:59 2020            6450 usec
 c1, t2              Sun Aug  2 18:27:00 2020 Sun Aug  2 18:27:00 2020            6456 usec
 c1, t2              Sun Aug  2 18:27:01 2020 Sun Aug  2 18:27:01 2020            6574 usec
 c2, t1              Sun Aug  2 18:59:05 2020 Sun Aug  2 18:59:05 2020             879 usec
 c2, t1              Sun Aug  2 18:59:06 2020 Sun Aug  2 18:59:06 2020            6582 usec
 c2, t1              Sun Aug  2 18:59:07 2020 Sun Aug  2 18:59:07 2020            7211 usec
 c2, t1              Sun Aug  2 18:59:10 2020 Sun Aug  2 18:59:10 2020            6551 usec
 c2, t1              Sun Aug  2 18:59:11 2020 Sun Aug  2 18:59:11 2020            6547 usec
 c2, t2              Sun Aug  2 18:58:59 2020 Sun Aug  2 18:58:59 2020            6585 usec
```

```
 c2, t2              Sun Aug  2 18:59:00 2020 Sun Aug  2 18:59:00 2020             6581 usec
 c2, t2              Sun Aug  2 18:59:01 2020 Sun Aug  2 18:59:01 2020             6586 usec
 c2, t2              Sun Aug  2 18:59:02 2020 Sun Aug  2 18:59:02 2020             6592 usec
 c2, t2              Sun Aug  2 18:59:03 2020 Sun Aug  2 18:59:03 2020             6591 usec
```

show services rpm twamp client history-results detail

```
user@host> show services rpm twamp-client history-results detail
    Owner: p, Test: t
    Probe results:
      Response received, Tue Jan  7 05:11:49 2014,
      Rtt: 184 usec, Round trip jitter: -96 usec, Round trip interarrival jitter: 57 usec
    Results over current test:
      Probes sent: 4, Probes received: 4, Loss percentage: 0
      Measurement: Round trip time
        Samples: 4, Minimum: 174 usec, Maximum: 196 usec, Average: 183 usec, Peak to peak: 22
usec, Stddev: 8 usec, Sum: 732 usec
      Measurement: Positive round trip jitter
        Samples: 1, Minimum: 110 usec, Maximum: 110 usec, Average: 110 usec, Peak to peak: 0
usec, Stddev: 0 usec, Sum: 110 usec
      Measurement: Negative round trip jitter
        Samples: 2, Minimum: 96 usec, Maximum: 811 usec, Average: 454 usec, Peak to peak: 715
usec, Stddev: 358 usec, Sum: 907 usec

    Owner: p, Test: t
    Probe results:
      Response received, Tue Jan  7 05:11:50 2014,      Rtt: 174 usec, Round trip jitter: -8
usec, Round trip interarrival jitter: 54 usec
    Results over current test:
      Probes sent: 5, Probes received: 5, Loss percentage: 0
      Measurement: Round trip time
        Samples: 5, Minimum: 174 usec, Maximum: 196 usec, Average: 181 usec, Peak to peak: 22
usec, Stddev: 8 usec, Sum: 906 usec
      Measurement: Positive round trip jitter
        Samples: 1, Minimum: 110 usec, Maximum: 110 usec, Average: 110 usec, Peak to peak: 0
usec, Stddev: 0 usec, Sum: 110 usec
      Measurement: Negative round trip jitter
        Samples: 3, Minimum: 8 usec, Maximum: 811 usec, Average: 305 usec, Peak to peak: 803
usec, Stddev: 360 usec, Sum: 915 usec
```

Release Information

Command introduced in Junos OS Release 15.1.

show services rpm twamp client probe-results

IN THIS SECTION

- Syntax | 26002
- Description | 26002
- Options | 26002
- Required Privilege Level | 26003
- Output Fields | 26003
- Sample Output | 26009
- Release Information | 26013

Syntax

```
show services rpm twamp client probe-results
<control-connection control-connection-name>
<test-session test-session-name>
```

Description

Display the results of the most recent real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) probes sent from the TWAMP client to the TWAMP server. You can also view the results of the probes or test packets sent from a TWAMP client to a TWAMP server for a particular control-connection, or a test-session associated with a control-connection.

Options

none	Display all results of the most recent TWAMP probes.

control-connection *control-connection-name* (Optional) Display information for the specified control-connection between a TWAMP client and a TWAMP server.

test-session *test-session-name* (Optional) Display information for the specified test session associated with a control-connection between a TWAMP client and a TWAMP server.

Required Privilege Level

view

Output Fields

Table 1488 on page 26003 lists the output fields for the `show services twamp client probe-results` command. Output fields are listed in the approximate order in which they appear.

Table 1488: show services twamp client probe-results Output Fields

Field Name	Field Description
`Owner`	Name of the session-sender or the control-client, which is the TWAMP client. When you configure the *control-client-name* option at the [`edit services rpm twamp client control-connection` *control-connection-name*] hierarchy level, this field displays the configured owner name or the client name.
`Test`	Name of a test representing a collection of probes. When you configure the `test-session` *test-name* statement at the [`edit services rpm twamp client control-connection` *control-connection-name*] hierarchy level, the field displays the configured test name.
`server-address`	Destination address used for the probes.
`server-port`	Destination port used for the probes.
`Client address`	Source or TWAMP client address used for the probes.
`Client port`	Source or TWAMP client port used for the probes.

Table 1488: show services twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description
`TWAMP-Server-Status`	Possible values: • `Light`: the `control-type light` statement is configured for the control connection. • `Connected` or `Not Connected`: for control connections configured with the `control-type managed` statement, displays whether or not the connection is up.
`Number-of-Retries-with-TWAMP-Server`	The number of times the system has tried to connect to the TWAMP server.
`Reflector address`	Session-reflector or TWAMP server address used for the probes.
`Reflector port`	Session-reflector or TWAMP server port used for the probes.
`Sender address`	Session-sender or TWAMP client address used for the probes.
`Sender-port`	Session-sender or TWAMP client port used for the probes.
`Local-link`	Egress logical interface for the link-local address
`Routing Instance Name`	Name configured on the `routing-instance` statement for the control connection.
`Probe type`	Protocol configured on the receiving probe server: `http-get`, `http-metadata-get`, `icmp-ping`, `icmp-ping-timestamp`, `tcp-ping`, `udp-ping`, or `udp-ping-timestamp`.
`Test size`	Number of probes within a test.
`Destination Interface Name`	Name of the interface configured on the TWAMP server or the session-reflector on which the TWAMP probe packets sent from the TWAMP client are received.

Table 1488: show services twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description
`Probe results`	Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement: • `Response received`—Timestamp when the probe result was determined. • `Client and server hardware timestamps`—If timestamps are configured, an entry appears at this point. • `Rtt`—Average ping round-trip time (RTT), in microseconds. • `Egress jitter`—Egress jitter, in microseconds. • `Ingress jitter`—Ingress jitter, in microseconds. • `Round trip jitter`—Round-trip jitter, in microseconds. • `Egress interarrival jitter`—Egress interarrival jitter, in microseconds. • `Ingress interarrival jitter`—Ingress interarrival jitter, in microseconds. • `Round trip interarrival jitter`—Round-trip interarrival jitter, in microseconds.

Table 1488: show services twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description

`Results over current test`	Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes.
	• `Probes sent`—Number of probes sent within the current test.
	• `Probes received`—Number of probe responses received within the current test.
	• `Loss percentage`—Percentage of lost probes for the current test.
	• `Measurement`—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type `icmp-ping-timestamp`, the egress delay and ingress delay.
	For each measurement type, the following individual calculated results are provided:
	  • `Samples`—Number of probes.
	  • `Minimum`—Minimum RTT, ingress delay, or egress delay measured over the course of the current test.
	  • `Maximum`—Maximum RTT, ingress delay, or egress delay measured over the course of the current test.
	  • `Average`—Average RTT, ingress delay, or egress delay measured over the course of the current test.
	  • `Peak to peak`—Peak-to-peak difference, in microseconds.
	  • `Stddev`—Standard deviation, in microseconds.
	  • `Sum`—Statistical sum.

Table 1488: show services twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description

`Results over last test`	Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed
	• `Probes sent`—Number of probes sent for the most recently completed test.
	• `Probes received`—Number of probe responses received for the most recently completed test.
	• `Loss percentage`—Percentage of lost probes for the most recently completed test.
	• `Test completed`—Time the most recent test was completed.
	• `Measurement`—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type `icmp-ping-timestamp`, the egress delay and ingress delay.
	For each measurement type, the following individual calculated results are provided:
	  • `Samples`—Number of probes.
	  • `Minimum`—Minimum RTT, ingress delay, or egress delay measured for the most recently completed test.
	  • `Maximum`—Maximum RTT, ingress delay, or egress delay measured for the most recently completed test.
	  • `Average`—Average RTT, ingress delay, or egress delay measured for the most recently completed test.
	  • `Peak to peak`—Peak-to-peak difference, in microseconds.
	  • `Stddev`—Standard deviation, in microseconds.
	  • `Sum`—Statistical sum.

Table 1488: show services twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description

`Results over all tests`	Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test.
	• `Probes sent`—Number of probes sent in all tests.
	• `Probes received`—Number of probe responses received in all tests.
	• `Loss percentage`—Percentage of lost probes in all tests.
	• `Measurement`—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe types `icmp-ping-timestamp` and `udp-ping-timestamp`, the egress delay and ingress delay.
	For each measurement type, the following individual calculated results are provided:
	  • `Samples`—Number of probes.
	  • `Minimum`—Minimum RTT, ingress delay, or egress delay measured over the course of the current test.
	  • `Maximum`—Maximum RTT, ingress delay, or egress delay measured over the course of the current test.
	  • `Average`—Average RTT, ingress delay, or egress delay measured over the course of the current test.
	  • `Peak to peak`—Peak-to-peak difference, in microseconds.
	  • `Stddev`—Standard deviation, in microseconds.
	  • `Sum`—Statistical sum.

`Error stats`	Error statistics

Sample Output

show services rpm twamp client probe-results

```
user@host> show services rpm twamp client probe-results
Owner: c1, Test: t1
    TWAMP-Server-Status: Light,    Number-Of-Retries-With-TWAMP-Server: 0
    Reflector address: 10.44.6.22, Reflector port: 862,  Sender address: 10.2.2.1, sender-port:
64109
    Routing Instance Name: IN
    Test size: 100 probes
    Probe results:
      Response received
      Probe sent time: Fri Mar  5 09:59:34 2021
      Probe rcvd/timeout time: Fri Mar  5 09:59:34 2021
      Rtt: 1728 usec, Egress jitter: -3082 usec, Ingress jitter: 742 usec, Round trip jitter:
-339 usec
      Egress interarrival jitter: 1047 usec, Ingress interarrival jitter: 826 usec, Round trip
interarrival jitter: 1139 usec
    Results over current test:
      Probes sent: 11, Probes received: 11, Loss percentage: 0.000000
      Measurement: Round trip time
        Samples: 11, Minimum: 1135 usec, Maximum: 4445 usec, Average: 2994 usec, Peak to peak:
3310 usec, Stddev: 1039 usec, Sum: 32938 usec
      Measurement: Positive egress jitter
        Samples: 7, Minimum: 17 usec, Maximum: 3101 usec, Average: 1426 usec, Peak to peak: 3084
usec, Stddev: 958 usec, Sum: 9981 usec
      Measurement: Negative egress jitter
        Samples: 4, Minimum: 1317 usec, Maximum: 3082 usec, Average: 2205 usec, Peak to peak:
1765 usec, Stddev: 868 usec, Sum: 8821 usec
      Measurement: Positive ingress jitter
        Samples: 3, Minimum: 742 usec, Maximum: 3306 usec, Average: 2222 usec, Peak to peak:
2564 usec, Stddev: 1083 usec, Sum: 6665 usec
      Measurement: Negative ingress jitter
        Samples: 8, Minimum: 3 usec, Maximum: 3198 usec, Average: 791 usec, Peak to peak: 3195
usec, Stddev: 1122 usec, Sum: 6325 usec
      Measurement: Positive round trip jitter
        Samples: 4, Minimum: 650 usec, Maximum: 4964 usec, Average: 2580 usec, Peak to peak:
4314 usec, Stddev: 1551 usec, Sum: 10319 usec
      Measurement: Negative round trip jitter
        Samples: 7, Minimum: 230 usec, Maximum: 2952 usec, Average: 919 usec, Peak to peak: 2722
usec, Stddev: 923 usec, Sum: 6430 usec
```

```
    Results over last test:
      Probes sent: 100, Probes received: 100, Loss percentage: 0.000000
      Test completed on Fri Mar  5 09:58:50 2021
      Measurement: Round trip time
        Samples: 100, Minimum: 192 usec, Maximum: 5425 usec, Average: 635 usec, Peak to peak:
5233 usec, Stddev: 948 usec, Sum: 63507 usec
      Measurement: Positive egress jitter
        Samples: 75, Minimum: 9 usec, Maximum: 1810 usec, Average: 223 usec, Peak to peak: 1801
usec, Stddev: 346 usec, Sum: 16735 usec
      Measurement: Negative egress jitter
        Samples: 25, Minimum: 23 usec, Maximum: 3871 usec, Average: 603 usec, Peak to peak: 3848
usec, Stddev: 805 usec, Sum: 15064 usec
      Measurement: Positive ingress jitter
        Samples: 19, Minimum: 1 usec, Maximum: 3710 usec, Average: 837 usec, Peak to peak: 3709
usec, Stddev: 922 usec, Sum: 15909 usec
      Measurement: Negative ingress jitter
        Samples: 81, Minimum: 4 usec, Maximum: 3834 usec, Average: 274 usec, Peak to peak: 3830
usec, Stddev: 577 usec, Sum: 22206 usec
      Measurement: Positive round trip jitter
        Samples: 48, Minimum: 5 usec, Maximum: 4382 usec, Average: 778 usec, Peak to peak: 4377
usec, Stddev: 1077 usec, Sum: 37365 usec
      Measurement: Negative round trip jitter
        Samples: 52, Minimum: 3 usec, Maximum: 4915 usec, Average: 787 usec, Peak to peak: 4912
usec, Stddev: 1239 usec, Sum: 40922 usec
    Results over all tests:
      Probes sent: 38611, Probes received: 38581, Loss percentage: 0.077698
      Measurement: Round trip time
        Samples: 38556, Minimum: 2 usec, Maximum: 22356 usec, Average: 3275 usec, Peak to peak:
22354 usec, Stddev: 3272 usec, Sum: 126256294 usec
      Measurement: Positive egress jitter
        Samples: 23770, Minimum: 0 usec, Maximum: 11266 usec, Average: 1054 usec, Peak to peak:
11266 usec, Stddev: 1496 usec, Sum: 25053834 usec
      Measurement: Negative egress jitter
        Samples: 14781, Minimum: 1 usec, Maximum: 9994 usec, Average: 1533 usec, Peak to peak:
9993 usec, Stddev: 1609 usec, Sum: 22659276 usec
      Measurement: Positive ingress jitter
        Samples: 12989, Minimum: 0 usec, Maximum: 12438 usec, Average: 1639 usec, Peak to peak:
12438 usec, Stddev: 1673 usec, Sum: 21293431 usec
      Measurement: Negative ingress jitter
        Samples: 25562, Minimum: 1 usec, Maximum: 13883 usec, Average: 927 usec, Peak to peak:
13882 usec, Stddev: 1458 usec, Sum: 23689224 usec
      Measurement: Positive round trip jitter
        Samples: 19298, Minimum: 0 usec, Maximum: 21037 usec, Average: 1793 usec, Peak to peak:
```

```
21037 usec, Stddev: 2146 usec, Sum: 34594283 usec
      Measurement: Negative round trip jitter
        Samples: 19257, Minimum: 1 usec, Maximum: 21288 usec, Average: 1796 usec, Peak to peak:
21287 usec, Stddev: 2148 usec, Sum: 34589725 usec
    Error Stats:
      Invalid client recv timestamp (T4 < T1): 1, Invalid client send(T1), recv(T4) timestamps: 0
      Invalid server send timestamp (T3 < T2): 1, Invalid server processing time (T4 -T1) < (T3 -
T2): 23
```

show services rpm twamp client probe-results (with local-link address information)

```
user@host> show services rpm twamp client probe-results
Owner: tc7, Test: ts1
    TWAMP-Server-Status: Light, Number-Of-Retries-With-TWAMP-Server: 0
    Reflector address: 2001:db8::9f16, Reflector port: 862, Sender address: 2001:db8::167d,
sender-port: 54399
    Local-link: ge-2/2/9.1
    Test size: 1 probes
    Probe results:
      Response received
      Probe sent time: Mon Aug 30 10:30:51 2021
      Probe rcvd/timeout time: Mon Aug 30 10:30:51 2021
      Rtt: 264 usec, Egress jitter: -66 usec, Ingress jitter: -14 usec, Round trip jitter: -59
usec
      Egress interarrival jitter: 73 usec, Ingress interarrival jitter: 45 usec, Round trip
interarrival jitter: 111 usec
    Results over current test:
      Probes sent: 1, Probes received: 1, Loss percentage: 0.000000
      Measurement: Round trip time
        Samples: 1, Minimum: 264 usec, Maximum: 264 usec, Average: 264 usec, Peak to peak: 0
usec, Stddev: 0 usec, Sum: 264 usec
      Measurement: Negative egress jitter
        Samples: 1, Minimum: 66 usec, Maximum: 66 usec, Average: 66 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 66 usec
      Measurement: Negative ingress jitter
        Samples: 1, Minimum: 14 usec, Maximum: 14 usec, Average: 14 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 14 usec
      Measurement: Negative round trip jitter
        Samples: 1, Minimum: 59 usec, Maximum: 59 usec, Average: 59 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 59 usec
```

```
    Results over last test:
      Probes sent: 1, Probes received: 1, Loss percentage: 0.000000
      Test completed on Mon Aug 30 10:30:51 2021
      Measurement: Round trip time
        Samples: 1, Minimum: 264 usec, Maximum: 264 usec, Average: 264 usec, Peak to peak: 0
usec, Stddev: 0 usec, Sum: 264 usec
      Measurement: Negative egress jitter
        Samples: 1, Minimum: 66 usec, Maximum: 66 usec, Average: 66 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 66 usec
      Measurement: Negative ingress jitter
        Samples: 1, Minimum: 14 usec, Maximum: 14 usec, Average: 14 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 14 usec
      Measurement: Negative round trip jitter
        Samples: 1, Minimum: 59 usec, Maximum: 59 usec, Average: 59 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 59 usec
    Results over all tests:
      Probes sent: 46059, Probes received: 46059, Loss percentage: 0.000000
      Measurement: Round trip time
        Samples: 46025, Minimum: 209 usec, Maximum: 4730 usec, Average: 361 usec, Peak to peak:
4521 usec, Stddev: 351 usec, Sum: 16593381 usec
      Measurement: Positive egress jitter
        Samples: 14538, Minimum: 0 usec, Maximum: 4319 usec, Average: 73 usec, Peak to peak:
4319 usec, Stddev: 275 usec, Sum: 1056470 usec
      Measurement: Negative egress jitter
        Samples: 31486, Minimum: 1 usec, Maximum: 4559 usec, Average: 40 usec, Peak to peak:
4558 usec, Stddev: 191 usec, Sum: 1251936 usec
      Measurement: Positive ingress jitter
        Samples: 34378, Minimum: 0 usec, Maximum: 4391 usec, Average: 59 usec, Peak to peak:
4391 usec, Stddev: 386 usec, Sum: 2029749 usec
      Measurement: Negative ingress jitter
        Samples: 11646, Minimum: 1 usec, Maximum: 4396 usec, Average: 158 usec, Peak to peak:
4395 usec, Stddev: 652 usec, Sum: 1834338 usec
      Measurement: Positive round trip jitter
        Samples: 23383, Minimum: 0 usec, Maximum: 10709 usec, Average: 276 usec, Peak to peak:
10709 usec, Stddev: 772 usec, Sum: 6454130 usec
      Measurement: Negative round trip jitter
        Samples: 22641, Minimum: 1 usec, Maximum: 11233 usec, Average: 285 usec, Peak to peak:
11232 usec, Stddev: 788 usec, Sum: 6454183 usec
    Error Stats:
      Invalid client recv timestamp (T4 < T1): 11, Invalid client send(T1), recv(T4) timestamps:
0
```

```
      Invalid server send timestamp (T3 < T2): 3, Invalid server processing time (T4 -T1) < (T3 -
T2): 20
```

show services rpm twamp client probe-results control-connection

```
user@host> show services rpm twamp client probe-results control-connection c2
    Owner: c2, Test: t1
    server-address: 192.0.2.14, server-port: 862, Client address: 192.0.2.13,
    Client port: 57170
    Reflector address: 192.0.2.14, Reflector port: 10010,
    Sender address: 192.0.2.13, sender-port: 10010
    Destination interface name: si-1/1/0.10
    Test size: 500 probes
    Probe results:
      Request timed out, Fri Feb 13 00:07:14 2015
    Results over current test:
      Probes sent: 188, Probes received: 0, Loss percentage: 100.00000
    Results over last test:
      Probes sent: 500, Probes received: 0, Loss percentage: 100.00000
    Results over all tests:
      Probes sent: 3688, Probes received: 0, Loss percentage: 100.00000
```

Release Information

Command introduced in Junos OS Release 15.1.

`Local-link` field added in Junos OS Release 21.4R1.

show services rpm twamp client session

IN THIS SECTION

- Syntax | 26014
- Description | 26014
- Options | 26014

- Required Privilege Level | 26014
- Output Fields | 26015
- Sample Output | 26015
- Release Information | 26016

Syntax

```
show services rpm twamp client session
<control-connection control-connection-name>
<test-session test-session-name>
```

Description

Display information about the sessions established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients for control packets and data packets. By default, all established control-connection and data-connection or test sessions are displayed, unless you specify a control-connection name or a test-session name when you issue the command. Because TWAMP light servers are stateless, information about them is not included in the output of this command; only information about managed servers is included.

Options

none	Display information about all established connections and sessions.
control-connection ***control-connection-name***	(Optional) Display information about the specified control-connection, which is established for control-packets exchanged between a TWAMP client and a TWAMP server.
test-session ***test-session-name***	(Optional) Display information about the specified test session, which is established for data packets transmitted between a TWAMP client and a TWAMP server, associated with a control-connection..

Required Privilege Level

view

Output Fields

Table 1489 on page 26015 lists the output fields for the `show services rpm twamp client session` command. Output fields are listed in the approximate order in which they appear.

Table 1489: show services rpm twamp client session Output Fields

Field Name	Field Description
`Connection Name`	Name of the control connection that uniquely identifies the connection between the TWAMP server and the TWAMP client.
`Session Name`	Name of the test session that uniquely identifies the data-session between the TWAMP server and the TWAMP client.
`Sender address`	Sender IP address.
`Sender port`	Sender port number.
`Reflector address`	Reflector IP address.
`Reflector port`	Reflector port number.

Sample Output

show services rpm twamp client session

```
user@host> show services rpm twamp client session

Connection  Session   Sender       Sender    Reflector    Reflector
      Name     Name   address       port      address       port
      cs1       ts1   198.51.100.1  41998   198.51.100.2    5008
      cs2       ts1   198.51.100.1  53710   198.51.100.2    5009
```

show services rpm twamp client session control-connection

```
user@host> show services rpm twamp client session control-connection c2

Connection  Session   Sender        Sender    Reflector  Reflector
      Name     Name   address         port      address     port
        c2       t1   192.0.2.13     10008   192.0.2.14    10008
```

Release Information

Command introduced in Junos OS Release 15.1.

show services rpm twamp server connection

IN THIS SECTION

- Syntax | 26016
- Description | 26017
- Options | 26017
- Required Privilege Level | 26017
- Output Fields | 26017
- Sample Output | 26018
- Release Information | 26018

Syntax

```
show services rpm twamp server connection
<connection-id>
```

Description

Display information about the connections established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control-clients. By default, all established sessions are displayed, unless you specify a session ID when you issue the command. Because TWAMP light servers are stateless, information about them is not included in the output of this command; only information about managed servers is included.

Options

none Display connection information about all established sessions.

connection-id (Optional) Identifier of the connection that you want to display information about.

Required Privilege Level

view

Output Fields

Table 1490 on page 26017 lists the output fields for the `show services rpm twamp server connection` command. Output fields are listed in the approximate order in which they appear.

Table 1490: show services rpm twamp server connection Output Fields

Field Name	Field Description
Connection ID	Connection ID that uniquely identifies the connection between the TWAMP server and a particular client.
Client address	Client IP address.
Client port	Client port number.
Server address	Server IP address.
Server port	Server port number.

Table 1490: show services rpm twamp server connection Output Fields *(Continued)*

Field Name	Field Description
Session count	Session count.
Auth mode	Authentication mode.

Sample Output

show services rpm twamp server connection

```
user@host> show services rpm twamp server connection
    Connection  Client       Client  Server           Server  Session Auth
    ID          address      port    address          port    count   mode
     4         192.0.2.1     12345  192.168.219.203  890      16      none
    78        198.51.100.55  345    203.0.113.2      89022    5       none
    234      192.168.219.203 2345   192.168.22.2     3333     16      none
     5        192.168.3.1    82345  192.168.2.2      45909    16   authenticated
     1        192.168.1.1    645    192.168.4.23     2394     16   encrypted
```

Release Information

Command introduced in Junos OS Release 9.3.

show services rpm twamp server session

IN THIS SECTION

- Syntax | 26019
- Description | 26019
- Options | 26019
- Required Privilege Level | 26019

- Output Fields | 26019
- Sample Output | 26020
- Release Information | 26020

Syntax

```
show services rpm twamp server session
<session-id>
```

Description

Display information about the sessions established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, all established sessions are displayed, unless you specify a session ID when you issue the command.

Options

none Display information about all established sessions.

session-id (Optional) Identifier of the session that you want to display information about.

Required Privilege Level

view

Output Fields

Table 1491 on page 26020 lists the output fields for the `show services rpm twamp server session` command. Output fields are listed in the approximate order in which they appear.

Table 1491: show services rpm twamp server session Output Fields

Field Name	Field Description
Session ID	Session ID that uniquely identifies the session between the TWAMP server and a particular client.
Connection ID	Connection ID that uniquely identifies the connection between the TWAMP server and a particular client.
Sender address	Sender IP address.
Sender port	Sender port number.
Reflector address	Reflector IP address.
Reflector port	Reflector port number.

Sample Output

show services rpm twamp server session

```
user@host> show services rpm twamp server session
    Session  Connection  Sender          Sender  Reflector      Reflector
    ID       ID          address         port    address        port
     4       44         192.0.2.1        12345  192.168.219.203    890
    78       44         198.51.100.55    345    203.0.113.2       89022
   234       423        192.168.219.203  2345   192.168.22.2       3333
     5       423        192.168.3.1      82345  192.168.2.2       45909
     1       423        192.168.1.1      645    192.168.4.23       2394
```

Release Information

Command introduced in Junos OS Release 9.3.

show services screen ids-option

IN THIS SECTION

- Syntax | 26021
- Description | 26021
- Required Privilege Level | 26021
- Output Fields | 26021
- Sample Output | 26023
- Release Information | 26024

Syntax

```
show services screen ids-option <ids-rules>
screen-name
```

Description

Display the configuration information about the specified services screen. You can configure a `ids-option` to enable screen protection on the MX Series devices.

Required Privilege Level

view

Output Fields

Table 1492 on page 26022 lists the output fields for the `show services screen ids-option` command. Output fields are listed in the approximate order in which they appear.

Table 1492: show services screen ids-option Output Fields

Field Name	Field Description
`IP block fragment`	Number of IP block fragments.
`Limit session source max sessions`	The maximum number of concurrent sessions that can be initiated from a single source IP address.
`Limit session destination max sessions`	The maximum number of concurrent sessions that can be directed to a single destination IP address.
`Limit session source packet rate`	Number of packets per second that can be initiated from a single source IP address.
`Limit session destination packet rate`	Number of packets per second that can be directed to a single destination IP address.
`Limit session source session rate`	Number of concurrent sessions per second that can be initiated from a single source IP address.
`Limit session destination session rate`	Number of concurrent sessions per second that can be directed to a single destination IP address.
`whitelist`	Allowlist of IP addresses. Specifies the list of IP addresses for a given category that are to be bypassed for scanning.
`Limit session source number of whitelists`	Number of allowed concurrent sessions that a device can initiate from a single source IP address.
`Limit session destination number of whitelists`	Number of allowed concurrent sessions that can be directed to the same destination IP address.

Sample Output

show services screen ids-option

```
user@host> show services screen ids-option <ids-rules>
Screen object status:

Name                                                        value
  IP block fragment                                         enabled
  Limit session source max sessions                         1
  Limit session source packet rate                          1
  Limit session source session rate                         1
  Limit session source number of whitelists                 2
  whitelist3
  whitelist1
  Limit session source tcp max sessions                     455
  Limit session source tcp packet rate                      5
  Limit session source tcp session rate                     99
  Limit session source tcp syn cookies mss                  65
  Limit session source tcp syn cookies threshold rate       1
  Limit session source tcp syn cookies threshold sessions   1
  Limit session source number of tcp whitelists             2
  whitelist2
  whitelist1
  Limit session source udp max sessions                     444
  Limit session source udp packet rate                      33
  Limit session source udp session rate                     77
  Limit session source number of udp whitelists             2
  whitelist1
  whitelist2
  Limit session source icmp max sessions                    2
  Limit session source icmp packet rate                     2
  Limit session source icmp session rate                    3
  Limit session source number of icmp whitelists            2
  whitelist3
  whitelist1
  Limit session destination max sessions                    1
  Limit session destination packet rate                     1
  Limit session destination session rate                    1250
  Limit session destination number of whitelists            2
  whitelist2
  whitelist1
```

```
Limit session destination tcp max sessions                  455
Limit session destination tcp packet rate                   5
Limit session destination tcp session rate                  99
Limit session destination tcp syn cookies mss               65
Limit session destination tcp syn cookies threshold rate    1
Limit session destination tcp syn cookies threshold sessions 1
Limit session destination number of tcp whitelists          2
whitelist2
whitelist1
Limit session destination udp max sessions                  444
Limit session destination udp packet rate                   33
Limit session destination udp session rate                  77
Limit session destination number of udp whitelists          2
whitelist1
whitelist2
Limit session destination icmp max sessions                 2
Limit session destination icmp packet rate                  2
Limit session destination icmp session rate                 3
Limit session destination number of icmp whitelists         2
whitelist3
whitelist1
```

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

ids-option | 8335

show services screen-statistics service-set

IN THIS SECTION

Syntax | 26025

- Description | 26025
- Required Privilege Level | 26025
- Output Fields | 26025
- Sample Output | 26029
- Release Information | 26031

Syntax

```
show services screen statistics service-set service-set
```

Description

Display intrusion detection service (IDS) screen statistics.

Required Privilege Level

view

Output Fields

Table 1493 on page 26025 lists the output fields for the `show services screen statistics service-set` command. Output fields are listed in the approximate order in which they appear.

Table 1493: show services screen statistics service-set Output Fields

Field Name	Field Description
`ICMP flood`	Internet Control Message Protocol (ICMP) flood counter. An ICMP flood typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.
`UDP flood`	User Datagram Protocol (UDP) flood counter. UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled.

Table 1493: show services screen statistics service-set Output Fields *(Continued)*

Field Name	Field Description
`TCP winnuke`	Number of Transport Control Protocol (TCP) WinNuke attacks. WinNuke is a denial-of-service (DoS) attack targeting any computer on the Internet running Windows.
`TCP port scan`	Number of TCP port scans. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
`UDP port scan`	Number of UDP port scans. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
`ICMP address sweep`	Number of ICMP address sweeps. An IP address sweep can occur with the intent of triggering responses from active hosts.
`IP tear drop`	Number of teardrop attacks. Teardrop attacks exploit the reassembly of fragmented IP packets.
`TCP SYN flood`	Number of TCP SYN attacks.
`IP spoofing`	Number of IP spoofs. IP spoofing occurs when an invalid source address is inserted in the packet header to make the packet appear to come from a trusted source.
`ICMP ping of death`	ICMP ping of death counter. Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).
`IP source route option`	Number of IP source route attacks.
`TCP address sweep`	Number of TCP address sweeps.
`TCP land attack`	Number of land attacks. Land attacks occur when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.

Table 1493: show services screen statistics service-set Output Fields *(Continued)*

Field Name	Field Description
`TCP SYN fragment`	Number of TCP SYN fragments.
`TCP no flag`	Number of TCP headers without flags set. A normal TCP segment header has at least one control flag set.
`IP unknown protocol`	Number of IPs.
`IP bad options`	Number of invalid options.
`IP record route option`	Number of packets with the IP record route option enabled. This option records the IP addresses of the network devices along the path that the IP packet travels.
`IP timestamp option`	Number of IP timestamp option attacks. This option records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination.
`IP security option`	Number of IP security option attacks.
`IP loose source route option`	Number of IP loose source route option attacks. This option specifies a partial route list for a packet to take on its journey from source to destination.
`IP strict source route option`	Number of IP strict source route option attacks. This option specifies the complete route list for a packet to take on its journey from source to destination.
`IP stream option`	Number of stream option attacks. This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support streams.
`ICMP fragment`	Number of ICMP fragments. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.

Table 1493: show services screen statistics service-set Output Fields *(Continued)*

Field Name	Field Description
`ICMP large packet`	Number of large ICMP packets.
`TCP SYN FIN`	Number of TCP SYN FIN packets.
`TCP FIN no ACK`	Number of TCP FIN flags without the acknowledge (ACK) flag.
`Source session limit`	Number of concurrent sessions that can be initiated from a source IP address.
`TCP SYN-ACK-ACK proxy`	Number of TCP flags enabled with SYN-ACK-ACK. To prevent flooding with SYN-ACK-ACK sessions, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold and SRX Series Firewalls running Junos OS reject further connection requests from that IP address.
`IP block fragment`	Number of IP block fragments.
`Destination session limit`	Number of concurrent sessions that can be directed to a single destination IP address.
`IPv6 extension header`	Number of IPv6 extension headers for the intrusion detection service (IDS). These headers are placed between the IPv6 header and the upper layer header in a packet. .
`IPv6 extension hop by hop option`	Number of IPv6 extension header that specifies delivery parameters at each hop on the path to the destination host.
`IPv6 extension destination option`	Number of IPv6 destination option for the extension header.
`IPv6 extension header limit`	Number of IPv6 extension headers that can pass through the screen.

Table 1493: show services screen statistics service-set Output Fields *(Continued)*

Field Name	Field Description
`IPv6 malformed header`	Number of IPv6 malformed header that is used to enable the IPv6 malformed header intrusion detection service (IDS) option. The screen doesn't pass any communication using the subnet 2001:db8 :: / 32.
`ICMPv6 malformed packet`	Number of malformed ICMPv6 packets, which might cause damage to the device and network.

Sample Output

show services screen statistics service-set

```
user@host> show services screen statistics service-set <usf-service-set>
Service-set: ipsec_ss1

Match-Direction: input
Screen statistics:

IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              0
  UDP port scan                              0
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
      SYN flood source                       0
      SYN flood destination                  0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
  TCP land attack                            0
  TCP SYN fragment                           0
  TCP no flag                                0
```

```
  IP unknown protocol                        0
  IP bad options                             0
  IP record route option                     0
  IP timestamp option                        0
  IP security option                         0
  IP loose source route option               0
  IP strict source route option              0
  IP stream option                           0
  ICMP fragment                              0
  ICMP large packet                          0
  TCP SYN FIN                                0
  TCP FIN no ACK                             0
  Source session limit                       0
  TCP SYN-ACK-ACK proxy                      0
  IP block fragment                          0
  Destination session limit                  0
  IPv6 extension header                      0
  IPv6 extension hop by hop option           0
  IPv6 extension destination option          0
  IPv6 extension header limit                0
  IPv6 malformed header                      0
  ICMPv6 malformed packet                    0
  IP tunnel summary                          0

Service-set: ipsec_ss1

Match-Direction: output
Screen statistics:

IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              0
  UDP port scan                              0
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
      SYN flood source                       0
      SYN flood destination                  0
  IP spoofing                                0
```

```
    ICMP ping of death                          0
    IP source route option                      0
    TCP land attack                             0
    TCP SYN fragment                            0
    TCP no flag                                 0
    IP unknown protocol                         0
    IP bad options                              0
    IP record route option                      0
    IP timestamp option                         0
    IP security option                          0
    IP loose source route option                0
    IP strict source route option               0
    IP stream option                            0
    ICMP fragment                               0
    ICMP large packet                           0
    TCP SYN FIN                                 0
    TCP FIN no ACK                              0
    Source session limit                        0
    TCP SYN-ACK-ACK proxy                       0
    IP block fragment                           0
    Destination session limit                   0
    IPv6 extension header                       0
    IPv6 extension hop by hop option            0
    IPv6 extension destination option           0
    IPv6 extension header limit                 0
    IPv6 malformed header                       0
    ICMPv6 malformed packet                     0
    IP tunnel summary                           0
```

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

ids-option | 8335

Example: Configuring Multiple Screening Options

show services security-intelligence

IN THIS SECTION

- Syntax | 26032
- Description | 26032
- Required Privilege Level | 26032
- Sample Output | 26032
- Release Information | 26033

Syntax

```
show services security-intelligence sec-profiling-feed status
show services security-intelligence category detail feed-name custom_userid_data start
start_number count total_number category-name SecProfiling
```

Description

Displays the status of security profiling feeds.

Required Privilege Level

View

Sample Output

show services security-intelligence sec-profiling-feed status

```
user@host> show services security-intelligence sec-profiling-feed status
Category SecProfiling
feed-name1
Feed type         : IP
     Last post time    : 2019-08-13 22:47:09 PDT
```

```
     Last post ret code: 200
     Last post status  : succeeded
feed-name2
Feed type         : UserID
     Last post time    : 2019-08-13 22:47:09 PDT
     Last post ret code: 200
     Last post status  : succeeded
```

show services security-intelligence category detail feed-name *custom_userid_data* start *start_number* count *total_number* category-name SecProfiling

```
user@host> show services security-intelligence category detail feed-name custom_userid_data
start 1 count 5 category-name SecProfiling
Category name   :SecProfiling
  Feed name     :custom_userid_data
  Version       :20200623.2
  Objects number:477
  Create time   :2020-06-22 21:51:56 PDT
  Update time   :2020-06-22 22:00:26 PDT
  Update status :Store succeeded
  Expired       :No
  Options       :N/A
  { UserID_A threat_level:10}
  { UserID_B threat_level:10}
  { UserID_C threat_level:10}
  { UserID_D threat_level:10}
  { UserID_E threat_level:10}
```

Release Information

Statement introduced before Junos OS Release 21.1.

`show services security-intelligence category detail` command added in Junos OS Release 21.1.

RELATED DOCUMENTATION

security-intelligence | **16464**

show services security-intelligence category summary

IN THIS SECTION

- Syntax | 26034
- Description | 26034
- Options | 26034
- Required Privilege Level | 26034
- Output Fields | 26034
- Sample Output | 26035
- Release Information | 26036

Syntax

```
show services security-intelligence category summary category-name
```

Description

Display summary for the specified Security Intelligence category.

Options

category-name	Name of the category.

Required Privilege Level

View

Output Fields

Table 1494 on page 26035 lists the output fields for the `show services security-intelligence category summary` command. Output fields are listed in the approximate order in which they appear.

Table 1494: show services security-intelligence category summary Output Fields

Field Name	Field Description
Category name	Name of the Security Intelligence category.
Status	Status of the Security Intelligence category.
Description	Description of the Security Intelligence category
Update interval	Amount of time after which Policy Enforcer sends an update for the feed.
TTL	Length of time (in minutes) the file remains open, receiving statistics before it is closed, transferred, and rotated. When either the time or the file size is exceeded, the file is closed and a new one is opened, whether or not a transfer site is specified.
Feed name	Information about the feed, including: • **Version** • **Object umber** • **Create time** • **Update time** • **Update status** • **Expired** • **Options** • **Status**

Sample Output

show services security-intelligence category summary

```
user@host> show services security-intelligence category summary
node1:
```

```
---------------------------------------------------------------------------

Category name     :CC
  Status          :Enable
  Description     :Command and Control data schema
  Update interval :1800s
  TTL             :3456000s
  Feed name       :cc_ip_data
    Version       :N/A
    Objects number:0
    Create time   :2018-03-16 05:57:39 PDT
    Update time   :2018-03-19 12:30:32 PDT
    Update status :N/A
    Expired       :No
    Options       :N/A
    Status        :Enabled
  Feed name       :cc_ipv6_data
    Version       :20180228.1
    Objects number:1
    Create time   :2018-03-16 05:57:39 PDT
    Update time   :2018-03-16 06:19:47 PDT
    Update status :Store succeeded
    Expired       :No
    Options       :N/A
    Status        :Disabled
```

Release Information

Statement introduced before Junos OS Release 18.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480, and MX960 with the MX-SPC3 services card.

Support for threat feed status (enabled, disabled, or user disabled) is added in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

security-intelligence (Services) | 16462

show services security-intelligence dns-statistics

IN THIS SECTION

- Syntax | 26037
- Description | 26037
- Required Privilege Level | 26037
- Output Fields | 26037
- Sample Output | 26037
- Release Information | 26041

Syntax

```
show services security-intelligence dns-statistics
show services security-intelligence dns-statistics logical-system all
show services security-intelligence dns-statistics tenant all
```

Description

Displays the DNS profile statistics.

Required Privilege Level

View

Output Fields

Sample Output

show services security-intelligence dns-statistics

```
user@host> show services security-intelligence dns-statistics
Logical system: root-logical-system
```

```
Category Whitelist:
UDP Counter:
Requests        Permit
  0              0

TCP Counter:
Request        Permit
  0              0

Category DNS:
  Profile p1:
    DNS Query Type       Requests      Sinkhole      Drop          Permit        Log

  UDP Counter:
            A              0             0             0             0             0
         AAAA              0             0             0             0             0
           MX              6             0             0             0             0
        CNAME             12             0             0             0             0
          SRV              0             0             0             0             0
    SRV NoErr              0             0             0             0             0
          TXT              9             0             0             0             0
    TXT NoErr              0             0             0             0             0
          ANY              0             0             0             0             0
         MISC              0             0             0             0             0

  TCP Counter:
            A              0             0             0             0             0
         AAAA              0             0             0             0             0
           MX              0             0             0             0             0
        CNAME              0             0             0             0             0
          SRV              0             0             0             0             0
    SRV NoErr              0             0             0             0             0
          TXT              0             0             0             0             0
    TXT NoErr              0             0             0             0             0
          ANY              0             0             0             0             0
         MISC              0             0             0             0             0

        Total             27             0             0             0             0
```

show services security-intelligence dns-statistics logical-system all

```
user@host> show services security-intelligence dns-statistics logical-system all
Logical system: root-logical-system
Category Whitelist:
UDP Counter:
Requests        Permit
  0              0

TCP Counter:
Request        Permit
  0              0

Category DNS:
  Profile p1:
    DNS Query Type       Requests      Sinkhole      Drop          Permit         Log

  UDP Counter:
            A              0             0             0             0             0
         AAAA              0             0             0             0             0
           MX              6             0             0             0             0
        CNAME             12             0             0             0             0
          SRV              0             0             0             0             0
    SRV NoErr              0             0             0             0             0
          TXT              9             0             0             0             0
    TXT NoErr              0             0             0             0             0
          ANY              0             0             0             0             0
         MISC              0             0             0             0             0
  TCP Counter:
            A              0             0             0             0            0
         AAAA              0             0             0             0            0
           MX              0             0             0             0            0
        CNAME              0             0             0             0            0
          SRV              0             0             0             0            0
    SRV NoErr              0             0             0             0            0
          TXT              0             0             0             0            0
    TXT NoErr              0             0             0             0            0
          ANY              0             0             0             0            0
         MISC              0             0             0             0            0

        Total             27             0             0             0            0
```

run show services security-intelligence dns-statistics tenant all

```
user@host> show services security-intelligence dns-statistics tenant all
Logical system: root-logical-system
Category Whitelist:
UDP Counter:
Requests        Permit
  0              0

TCP Counter:
Request        Permit
  0              0

Category DNS:
  Profile p1:
    DNS Query Type       Requests      Sinkhole      Drop          Permit        Log

  UDP Counter:
            A              0             0             0             0             0
         AAAA              0             0             0             0             0
           MX              6             0             0             0             0
        CNAME             12             0             0             0             0
          SRV              0             0             0             0             0
    SRV NoErr              0             0             0             0             0
          TXT              9             0             0             0             0
    TXT NoErr              0             0             0             0             0
          ANY              0             0             0             0             0
         MISC              0             0             0             0             0

  TCP Counter:
            A              0             0             0             0             0
         AAAA              0             0             0             0             0
           MX              0             0             0             0             0
        CNAME              0             0             0             0             0
          SRV              0             0             0             0             0
    SRV NoErr              0             0             0             0             0
          TXT              0             0             0             0             0
    TXT NoErr              0             0             0             0             0
          ANY              0             0             0             0             0
         MISC              0             0             0             0             0
```

```
        Total             27             0             0             0             0
```

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

security-intelligence(services) | 16469

clear services security-intelligence dns-statistics | 4160

show services security-intelligence update status

IN THIS SECTION

- Syntax | 26041
- Description | 26041
- Required Privilege Level | 26042
- Sample Output | 26042
- Release Information | 26042

Syntax

```
show services security-intelligence update status
```

Description

Display the status of the connection with Policy Enforcer.

Required Privilege Level

View

Sample Output

show services security-intelligence update status

```
user@host> show services security-intelligence update status
node1:
--------------------------------------------------------------------------
Current action        :Start downloading the latest manifest.
Last update status    :Download manifest failed.
Last connection status:succeeded
Last update time      :2018-03-21 16:59:59 PDT
```

Release Information

Statement introduced before Junos OS Release 18.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480, and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

security-intelligence (Services) | 16462

show services security-metadata-streaming

IN THIS SECTION

- Syntax | 26043
- Description | 26043
- Required Privilege Level | 26043
- Output Fields | 26043

Sample Output | 26046
Release Information | 26047

Syntax

```
show services security-metadata-streaming dns statistics
show services security-metadata-streaming dns statistics logical-system <logical-system-name>
show services security-metadata-streaming dns statistics logical-system all
show services security-metadata-streaming dns statistics tenant <tenant-name>
show services security-metadata-streaming dns statistics tenant all
show services security-metadata-streaming http statistics
show services security-metadata-streaming http whitelist
```

Description

The `show services security-metadata-streaming http statistics` command displays the statistics of security metadata streaming sessions for HTTP protocol.

The `show services security-metadata-streaming http whitelist` command displays the allowlist servers that are configured by the user for HTTP protocol.

The `show services security-metadata-streaming dns statistics` displays the statistics of security metadata streaming sessions for DNS protocol.

Required Privilege Level

view

Output Fields

Table 1495 on page 26044 lists the output fields for the `show services security-metadata-streaming http statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1495: show services security-metadata-streaming http statistics Output Fields

Field Name	Field Description
HTTP session statistics	
Session inspected	Number of sessions inspected.
Session allowlisted	Number of sessions allowlisted for encrypted traffic analysis.
Session detected	Number of sessions detected as potentially malicious.
HTTP submission statistics	
Records submission success	Number of records that were successfully submitted to ATP Cloud.
Records submission failure	Number of records submissions that failed while submitting to ATP Cloud.

Table 1496 on page 26044 lists the output fields for the `show services security-metadata-streaming dns statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1496: show services security-metadata-streaming dns statistics Output Fields

Field Name	Field Description
DNS session statistics	
Cache Hits	Number of cache hits (domain is present in the cache).
Cache Misses	Number of cache miss(domain is not present in the cache).
C2 Sessions Permitted	Number of C2 sessions that are permitted.
C2 Sessions Dropped	Number of C2 sessions that are dropped.

Table 1496: show services security-metadata-streaming dns statistics Output Fields *(Continued)*

Field Name	Field Description
C2 Sessions Sinkholed	Number of C2 sessions that are sinkholed.
DNS submission statistics	
Domain Submission Success	Number of successful domain submissions.
Domain Submission Failures	Number of failed domain submissions.
Safe Verdicts Received	Number of safe verdicts received by Juniper ATP cloud.
C2 Verdicts Received	Number of C2 verdicts received by Juniper ATP cloud.
DNS Tunnels Detected	Number of DNS tunnels detected.
Latency Fallback Triggered	Number of times fallback options were triggered due to not receiving a verdict from Juniper ATP cloud within the configured verdict timeout.
ATP latency statistics	
Average Latency	The average response time (in milliseconds) taken by Juniper ATP Cloud to provide a verdict to the SRX Series Firewall.
Maximum Latency	The maximum response time (in milliseconds) taken by Juniper ATP Cloud to provide a verdict to the SRX Series Firewall.
Minimum Latency	The minimum response time (in millisecond) taken by Juniper ATP Cloud to provide a verdict to the SRX Series Firewall.

Table 1496: show services security-metadata-streaming dns statistics Output Fields *(Continued)*

Field Name	Field Description
sub-50ms response	Number of responses with latency is less than 50ms.
sub-100ms response	Number of responses with latency is less than 100ms.
sub-250ms response	Number of responses with latency is less than 250ms.
sub-500ms response	Number of responses with latency is less than 500ms.

Sample Output

show services security-metadata-streaming http statistics

```
user@host> show services security-metadata-streaming http statistics
Security Metadata Streaming session statistics:
  Session inspected:    10
  Session whitelisted:   0
  Session detected:      6
Security Metadata Streaming submission statistics:
  Records submission success:         8
  Records submission failure:         2
```

show services security-metadata-streaming http whitelist

```
user@host> show services security-metadata-streaming http whitelist

No. IP-start IP-end Feed Address
1 192 0.5.0 192.0.5.1 eta_custom_whitelist ID-80001400
```

show services security-metadata-streaming dns statistics

```
user@host> show services security-metadata-streaming dns statistics
Logical system: root-logical-system
DNS session statistics:
  Cache Hits:                0
  Cache Misses:              116
  C2 Sessions Permitted:     0
  C2 Sessions Dropped:       0
  C2 Sessions Sinkholed:     12

DNS submission statistics:
  Domain Submission Success:   43
  Domain Submission Failures:  0
  Safe Verdicts Received:      8
  C2 Verdicts Received:        12
  DNS Tunnels Detected:        0
  Latency Fallback Triggered:  0

ATP latency statistics:
  Average Latency:             63ms
  Maximum Latency:             119ms
  Minimum Latency:             39ms
  sub-50ms response:           52 (6%)
  sub-100ms response:          4000 (88%)
  sub-250ms response:          30 (4%)
  sub-500ms response:          2 (2%)
```

Release Information

Command introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Juniper Advanced Threat Prevention Cloud Policy Overview

security-metadata-streaming | **16487**

show services service-redirect statistic

IN THIS SECTION

- Syntax | 26048
- Description | 26048
- Required Privilege Level | 26048
- Sample Output | 26048
- Sample Output | 26049
- Release Information | 26049

Syntax

```
show services service-redirect statistic
```

Description

Display the Service Redirect statistic.

Required Privilege Level

view

Sample Output

show services service-redirect statistic

```
user@host> show services service-redirect statistic
ICAP Redirect statistic:
  Message Redirected                  : 4
      Message REQMOD Redirected       : 2
      Message RESPMOD Redirected      : 2
  Message Received                    : 4
      Message REQMOD Received         : 2
```

```
      Message RESPMOD Received        : 2
Fallback:      permit           log-permit       reject
  Timeout         0                 0                 0
  Connectivity    0                 0                 0
  Default         0                 0                 0
```

Sample Output

show services icap-redirect statistic logical-system

```
user@host> show services icap-redirect statistic logical-system LSYS1
ICAP Redirect statistic:
  Message Redirected                  : 12
      Message REQMOD Redirected       : 6
      Message RESPMOD Redirected      : 6
  Message Received                    : 12
      Message REQMOD Received         : 6
      Message RESPMOD Received        : 6
Fallback:      permit           log-permit       reject
  Timeout         0                 0                 0
  Connectivity    0                 0                 0
  Default         0                 0                 0
```

Release Information

Command introduced in Junos OS Release 18.1R1.

`logical-system` option is introduced in Junos OS Release 18.3R1.

show services service-set statistics ids drops

IN THIS SECTION

- Syntax | 26050
- Description | 26050

- Options | 26050
- Required Privilege Level | 26050
- Output Fields | 26050
- Sample Output | 26056
- Release Information | 26058

Syntax

```
show services service-set statistics ids drops
<interface interface-name>
<service-set service-set-name>
```

Description

Display statistics for packet drops resulting from header-integrity, suspicious packet pattern, and session-limit checks performed by an MS-MPC or MS-MIC.

Options

none	Display statistics for all configured service interfaces and service sets.
interface *interface-name*	(Optional) Display statistics for the specified services interface.
service-set *service-set-name*	(Optional) Display statistics for the specified service set.

Required Privilege Level

view

Output Fields

Table 1497 on page 26051 lists the output fields for the `show services service-set integrity-drops` command. Output fields are listed in the approximate order in which they appear.

Table 1497: show services service-set statistics ids drops Output Fields

Field Name	Field Description
`Interface`	Name of an adaptive services interface.
`Service set`	Name of a service set.
`Errors`	Total errors, categorized by protocol: • `IP`—Total IP version 4 errors. • `TCP`—Total Transmission Control Protocol (TCP) errors. • `UDP`—Total User Datagram Protocol (UDP) errors. • `ICMP`—Total Internet Control Message Protocol (ICMP) errors.

Table 1497: show services service-set statistics ids drops Output Fields *(Continued)*

Field Name	Field Description
`IP Errors`	Number of IPv4 errors for the following categories: • `IP packet length inconsistencies`—IP packet length did not match the Layer 2 reported length. • `Minimum IP header length check failures`—Minimum IP header length is 20 bytes. The received packet contained less than 20 bytes. • `Reassembled packet exceeds maximum IP length`—After fragment reassembly, the reassembled IP packet length exceeded 65,535. • `Illegal source address 0`—Source address is not a valid address. Invalid addresses are loopback, broadcast, multicast, and reserved addresses. Source address `0`, however, is allowed to support BOOTP and the destination address `0xffffffff`. • `Illegal destination address` —Destination address was not a valid address. The address is reserved. • `TTL zero errors`—Received packet had a time-to-live (TTL) value of 0. • `Illegal IP protocol number 0 or 255`—IP protocol is 0 or 255. • `Land attack`—IP source address is the same as the destination address. • `Non-IP packets`—Packet did not conform to the IP standard. • `IP option`—Packet had a non-allowed IP option. • `Non-IPv4 packets`—Packet was not of the IPv4 type. • `Non-IPv6 packets`—Packet was not of the IPv6 type. • `Bad checksum`—Packet had an invalid IP checksum. • `Illegal IP fragment length`—Illegal fragment length. All fragments (other than the last fragment) must have a length that is a multiple of 8 bytes. • `IP fragment overlap`—Fragments had overlapping fragment offsets. • `IP fragment limit exceeded` —Configured number of allowed fragments for a packet was exceeded.

Table 1497: show services service-set statistics ids drops Output Fields *(Continued)*

Field Name	Field Description
	• `IP fragment reassembly timeout`—Some of the fragments for an IP packet were not received in time, and the reassembly handler dropped partial fragments. Whenever a fragment is received, it is maintained in a chain until all other fragments are received. If other fragments do not arrive within the configured value of reassembly-timeout, this packet is dropped and the value of the counter shown in this field is incremented. If other fragments arrive in time but the total number of fragments is more than the configured value of fragment-limit, all the fragments (of this packet) are dropped and the value of the counter shown in this field is incremented. • `IPv4 bad options`—Packet IP header contained IPv4 option that is not allowed. • `IPv6 bad extension headers`—Packet contained IPv6 extension header type that is not allowed. • `session-limit exceeded for source`—Number of concurrent sessions from an individual source address or subnet exceeded limit. • `session-limit exceeded for destination`—Number of concurrent sessions to an individual destination address or subnet exceeded limit. • `connections/second limit exceeded for source`—Number of connections per second for an individual source address or subnet exceeded limit. • `connections/second limit exceeded for destination`—Number of connections per second for an individual destination address or subnet exceeded limit. • `packets/second limit exceeded for source`—Number of packets per second for an individual source address or subnet exceeded limit. • `packet/second limit exceeded for destination`—Number of packets per second for an individual destination address or subnet exceeded limit. • `Unknown` —Unknown fragments.

Table 1497: show services service-set statistics ids drops Output Fields *(Continued)*

Field Name	Field Description
`TCP Errors`	Number of TCP protocol errors for the following categories: • `TCP header length inconsistencies`—Minimum TCP header length is 20 bytes, and the IP packet received did not contain at least 20 bytes. • `Source or destination port number is zero`—TCP source or destination port was zero. • `Illegal sequence number, flags combination`—Packet had any type of TCP header anomaly. • `TCP winnuke`—TCP segments destined for port 139 with the urgent (URG) flag set. • `TCP SYN Fragment`—TCP SYN packet was a fragment. • `TCP connection closed due to SYN defense`—Unestablished TCP connection closed because `open-timeout` value expired. • `TCP session-limit exceeded for source`—Number of concurrent TCP sessions from an individual source address or subnet exceeded limit. • `TCP session-limit exceeded for destination`—Number of concurrent TCP sessions to an individual destination address or subnet exceeded limit. • `TCP connections/second limit exceeded for source`—Number of TCP connections per second for an individual source address or subnet exceeded limit. • `TCP connections/second limit exceeded for destination`—Number of TCP connections per second for an individual destination address or subnet exceeded limit. • `TCP packets/second limit exceeded for source`—Number of TCP packets per second for an individual source address or subnet exceeded limit. • `TCP packet/second limit exceeded for destination`—Number of TCP packets per second for an individual destination address or subnet exceeded limit.

Table 1497: show services service-set statistics ids drops Output Fields *(Continued)*

Field Name	Field Description
`UDP Errors`	Number of UDP protocol errors for the following categories: • `IP data length less than minimum UDP header length (8 bytes)`—Minimum UDP header length is 8 bytes. The received IP packets contained less than 8 bytes. • `Source or destination port is zero`—UDP source or destination port was `0`. • `UDP session-limit exceeded for source`—Number of concurrent UDP sessions from an individual source address or subnet exceeded limit. • `UDP session-limit exceeded for destination`—Number of concurrent UDP sessions to an individual destination address or subnet exceeded limit. • `UDP connections/second limit exceeded for source`—Number of UDP connections per second for an individual source address or subnet exceeded limit. • `UDP connections/second limit exceeded for destination`—Number of UDP connections per second for an individual destination address or subnet exceeded limit. • `UDP packets/second limit exceeded for source`—Number of UDP packets per second for an individual source address or subnet exceeded limit. • `UDP packet/second limit exceeded for destination`—Number of UDP packets per second for an individual destination address or subnet exceeded limit.

Table 1497: show services service-set statistics ids drops Output Fields *(Continued)*

Field Name	Field Description
`ICMP Errors`	Number of ICMP protocol errors for the following categories: • `IP data length less than minimum ICMP header length (8 bytes)`—ICMP header length contained less than 8 bytes. • `ICMP error length inconsistencies`—ICMP error packet length was outside range of 48 bytes through 576 bytes. • `ICMP fragments`— ICMP packet was an IP fragment. • `ICMP session-limit exceeded for source`—Number of concurrent ICMP sessions from an individual source address or subnet exceeded limit. • `ICMP session-limit exceeded for destination`—Number of concurrent ICMP sessions to an individual destination address or subnet exceeded limit. • `ICMP connections/second limit exceeded for source`—Number of ICMP connections per second for an individual source address or subnet exceeded limit. • `ICMP connections/second limit exceeded for destination`—Number of ICMP connections per second for an individual destination address or subnet exceeded limit. • `ICMP packets/second limit exceeded for source`—Number of ICMP packets per second for an individual source address or subnet exceeded limit. • `ICMP packet/second limit exceeded for destination`—Number of ICMP packets per second for an individual destination address or subnet exceeded limit.

Sample Output

show services service-set statistics ids drops

```
user@host> show services service-set statistics ids drops
 Interface: ms-1/0/0
  Service set: sset1
    Errors:
      IP: 0, TCP: 0
```

```
      UDP: 0, ICMP: 0
    IP errors:
      IP packet length inconsistencies: 0
      Illegal source address: 0
      Illegal destination address: 0
      TTL zero errors: 0, Illegal IP protocol number (0 or 255): 0
      Land attack: 0
      Non-IPv4 packets: 0
      Non-IPv6 packets: 0
      Bad checksum: 0
      Illegal IP fragment length: 0
      IP fragment overlap: 0
      IP fragment reassembly timeout: 0
      IP fragment limit exceeded: 0
      IPv4 bad options: 0
      IPv6 bad extension headers: 0
      session-limit exceeded for source: 0
      session-limit exceeded for destination: 0
      connections/second limit exceeded for source: 0
      connections/second limit exceeded for destination: 0
      packets/second limit exceeded for source: 0
      packet/second limit exceeded for destination: 0
      Unknown: 0
    TCP errors:
      TCP header length inconsistencies: 0
      Source or destination port number is zero: 0
      Illegal sequence number and flags combinations: 0
      TCP winnuke: 0
      TCP SYN Fragment: 0
      TCP connection closed due to SYN defense: 0
      TCP session-limit exceeded for source: 0
      TCP session-limit exceeded for destination: 0
      TCP connections/second limit exceeded for source: 0
      TCP connections/second limit exceeded for destination: 0
      TCP packets/second limit exceeded for source: 0
      TCP packet/second limit exceeded for destination: 0
    UDP errors:
      IP data length less than minimum UDP header length (8 bytes): 0
      Source or destination port number is zero: 0
      UDP session-limit exceeded for source: 0
      UDP session-limit exceeded for destination: 0
      UDP connections/second limit exceeded for source: 0
      UDP connections/second limit exceeded for destination: 0
```

```
    UDP packets/second limit exceeded for source: 0
    UDP packet/second limit exceeded for destination: 0
  ICMP errors:
    IP data length less than minimum ICMP header length (8 bytes): 0
    ICMP error length inconsistencies: 0
    ICMP fragments: 0
    ICMP session-limit exceeded for source: 0
    ICMP session-limit exceeded for destination: 0
    ICMP connections/second limit exceeded for source: 0
    ICMP connections/second limit exceeded for destination: 0
    ICMP packets/second limit exceeded for source: 0
    ICMP packet/second limit exceeded for destination: 0
```

Release Information

Command introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

show services service-sets cpu-usage

IN THIS SECTION

- Syntax | 26059
- Description | 26059
- Options | 26059
- Required Privilege Level | 26059
- Output Fields | 26059
- Sample Output | 26060
- Release Information | 26060

Syntax

```
show services service-sets cpu-usage
<interface interface-name>
<service-set service-set-name>
```

Description

Display service set CPU usage as a percentage. The command is supported only on Adaptive Services PICs (SP PICs).

Options

none	Display CPU usage for all adaptive services interfaces and service sets.
interface *interface-name*	(Optional) Display CPU usage for a particular interface. On M Series and T Series routers, the *interface-name* parameter can have the value **sp-***fpc/pic/port* or **rsp***number*.
service-set *service-set-name*	(Optional) Display CPU usage for a particular service set. For the Layer 2 Tunneling Protocol (L2TP), you can use a tunnel group to represent a service set.

Required Privilege Level

view

Output Fields

Table 1498 on page 26059 lists the output fields for the `show services service-sets cpu-usage` command. Output fields are listed in the approximate order in which they appear.

Table 1498: show services service-sets cpu-usage Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface

Table 1498: show services service-sets cpu-usage Output Fields *(Continued)*

Field Name	Field Description
Service set (system category)	Name of the CPU usage category: • idp_recommended—Name of the service sets (displays all the service sets attached to the service PICs) • Idle • System • Receive • Transmit
CPU utilization %	Percentage of the CPU resources being used

Sample Output

show services service-sets cpu-usage

```
user@host> show services service-sets cpu-usage
Interface   Service set (system category)                  CPU utilization %
sp-4/1/0    idp_recommended                                           18.20 %
sp-4/1/0    Idle                                                      44.69 %
sp-4/1/0    System                                                     7.01 %
sp-4/1/0    Receive                                                   15.10 %
sp-4/1/0    Transmit                                                  15.00 %
```

Release Information

Command introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

show services service-sets memory-usage

IN THIS SECTION

- Syntax | 26061
- Description | 26061
- Options | 26061
- Required Privilege Level | 26062
- Output Fields | 26062
- Sample Output | 26063
- Release Information | 26063

Syntax

```
show services service-sets memory-usage
<interface interface-name>
<service-set service-set-name>
<zone>
```

Description

Display service set memory usage.

Options

none Display service set memory usage.

interface ***interface-name*** (Optional) Display memory usage for a particular interface. On M Series and T Series routers, the ***interface-name*** can be **sp-*fpc/pic/port***, or **rsp*number***.

> **NOTE**: This command is not supported on Multilink Protocol–based services PICs.
>
> The interface option is not supported on Multiservice PICs.

service-set *service-set-name*	(Optional) Display memory usage for a particular service set. For Layer 2 Tunneling Protocol (L2TP), you can use a tunnel group to represent a service set.
zone	(Optional) Display the memory usage zone of the adaptive services interface or an individual service set.

Required Privilege Level

view

Output Fields

Table 1499 on page 26062 lists the output fields for the `show services service-sets memory-usage` command. Output fields are listed in the approximate order in which they appear.

Table 1499: show services service-sets memory-usage Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface
Service set	Name of a service set
Bytes Used	Number of bytes of memory being used
Memory zone	Memory zone in which the adaptive services interface is currently operating: • **Green**—All new flows are allowed. • **Yellow**—Unused memory is reclaimed. All new flows are allowed. • **Orange**—New flows are allowed only for service sets that are using less than their equal share of memory. • **Red**—No new flows are allowed.

Sample Output

show services service-sets memory-usage

```
user@host> show services service-sets memory-usage
Interface   Service set                                Bytes Used
ms-4/0/0    N/A                                                14817036
ms-4/1/0    N/A                                                14691700
```

show services service-sets memory-usage zone

```
user@host> show services service-sets memory-usage zone
Interface   Memory zone
```

show services service-sets memory-usage interface

```
user@host> show services service-sets memory-usage interface ms-4/1/0
Interface   Service Set                                        Bytes Used
 ms-4/1/0    N/A                                                       14691700
```

Release Information

Command introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

show services service-sets plug-ins

IN THIS SECTION

- Syntax | 26064
- Description | 26064

- Options | 26064
- Required Privilege Level | 26064
- Output Fields | 26064
- Sample Output | 26064
- Release Information | 26065

Syntax

```
show services service-sets plug-ins <interface interface-name>
```

Description

Display service set plug-ins summary.

Options

interface *interface-name*	Display service set plug-ins information for the specified interface.

Required Privilege Level

view

Output Fields

Sample Output

show services service-sets plug-ins

```
user@host> show services service-sets plug-ins
Interface: vms-0/2/0
  Service-set: ss1_interface_style1, State: Ready
```

```
  Plugins configured: 1
Plugin: junos-alg, ID: 25
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services service-sets statistic screen-drops (Next Gen Services)

IN THIS SECTION

- Syntax | 26065
- Description | 26065
- Options | 26065
- Required Privilege Level | 26066
- Output Fields | 26066
- Sample Output | 26071
- Release Information | 26073

Syntax

```
show services service-sets statistic screen-drops [service-set| interface interface-name]
```

Description

Display statistics for packet drops resulting from header-integrity, suspicious packet pattern, and session-limit checks performed by an MS-MPC or MS-MIC.

Options

none	Display statistics for all configured service interfaces and service sets.

<interface *interface-name*>	(Optional) Display statistics for the specified services interface.
<service-set *service-set-name* >	(Optional) Display statistics for the specified service set.

Required Privilege Level

view

Output Fields

Table 1500 on page 26066 lists the output fields for the `show services service-set integrity-drops` command. Output fields are listed in the approximate order in which they appear.

Table 1500: show services service-set statistics screen drops Output Fields

Field Name	Field Description
`Interface`	Name of an adaptive services interface.
`Service set`	Name of a service set.
`Errors`	Total errors, categorized by protocol: • `IP`—Total IP version 4 errors. • `TCP`—Total Transmission Control Protocol (TCP) errors. • `UDP`—Total User Datagram Protocol (UDP) errors. • `ICMP`—Total Internet Control Message Protocol (ICMP) errors.

Table 1500: show services service-set statistics screen drops Output Fields *(Continued)*

Field Name	Field Description
`IP Errors`	Number of IPv4 errors for the following categories: • `IP packet length inconsistencies`—IP packet length did not match the Layer 2 reported length. • `Minimum IP header length check failures`—Minimum IP header length is 20 bytes. The received packet contained less than 20 bytes. • `Reassembled packet exceeds maximum IP length`—After fragment reassembly, the reassembled IP packet length exceeded 65,535. • `Illegal source address 0`—Source address is not a valid address. Invalid addresses are loopback, broadcast, multicast, and reserved addresses. Source address `0`, however, is allowed to support BOOTP and the destination address `0xffffffff`. • `Illegal destination address` —Destination address was not a valid address. The address is reserved. • `TTL zero errors`—Received packet had a time-to-live (TTL) value of 0. • `Illegal IP protocol number 0 or 255`—IP protocol is 0 or 255. • `Land attack`—IP source address is the same as the destination address. • `Non-IP packets`—Packet did not conform to the IP standard. • `IP option`—Packet had a non-allowed IP option. • `Non-IPv4 packets`—Packet was not of the IPv4 type. • `Non-IPv6 packets`—Packet was not of the IPv6 type. • `Bad checksum`—Packet had an invalid IP checksum. • `Illegal IP fragment length`—Illegal fragment length. All fragments (other than the last fragment) must have a length that is a multiple of 8 bytes. • `IP fragment overlap`—Fragments had overlapping fragment offsets. • `IP fragment limit exceeded` —Configured number of allowed fragments for a packet was exceeded.

Table 1500: show services service-set statistics screen drops Output Fields *(Continued)*

Field Name	Field Description
	• `IP fragment reassembly timeout`—Some of the fragments for an IP packet were not received in time, and the reassembly handler dropped partial fragments. Whenever a fragment is received, it is maintained in a chain until all other fragments are received. If other fragments do not arrive within the configured value of reassembly-timeout, this packet is dropped and the value of the counter shown in this field is incremented. If other fragments arrive in time but the total number of fragments is more than the configured value of fragment-limit, all the fragments (of this packet) are dropped and the value of the counter shown in this field is incremented. • `IPv4 bad options`—Packet IP header contained IPv4 option that is not allowed. • `IPv6 bad extension headers`—Packet contained IPv6 extension header type that is not allowed. • `session-limit exceeded for source`—Number of concurrent sessions from an individual source address or subnet exceeded limit. • `session-limit exceeded for destination`—Number of concurrent sessions to an individual destination address or subnet exceeded limit. • `connections/second limit exceeded for source`—Number of connections per second for an individual source address or subnet exceeded limit. • `connections/second limit exceeded for destination`—Number of connections per second for an individual destination address or subnet exceeded limit. • `packets/second limit exceeded for source`—Number of packets per second for an individual source address or subnet exceeded limit. • `packet/second limit exceeded for destination`—Number of packets per second for an individual destination address or subnet exceeded limit. • `Unknown` —Unknown fragments.

Table 1500: show services service-set statistics screen drops Output Fields *(Continued)*

Field Name	Field Description
`TCP Errors`	Number of TCP protocol errors for the following categories: • `TCP header length inconsistencies`—Minimum TCP header length is 20 bytes, and the IP packet received did not contain at least 20 bytes. • `Source or destination port number is zero`—TCP source or destination port was zero. • `Illegal sequence number, flags combination`—Packet had any type of TCP header anomaly. • `TCP winnuke`—TCP segments destined for port 139 with the urgent (URG) flag set. • `TCP SYN Fragment`—TCP SYN packet was a fragment. • `TCP connection closed due to SYN defense`—Unestablished TCP connection closed because `open-timeout` value expired. • `TCP session-limit exceeded for source`—Number of concurrent TCP sessions from an individual source address or subnet exceeded limit. • `TCP session-limit exceeded for destination`—Number of concurrent TCP sessions to an individual destination address or subnet exceeded limit. • `TCP connections/second limit exceeded for source`—Number of TCP connections per second for an individual source address or subnet exceeded limit. • `TCP connections/second limit exceeded for destination`—Number of TCP connections per second for an individual destination address or subnet exceeded limit. • `TCP packets/second limit exceeded for source`—Number of TCP packets per second for an individual source address or subnet exceeded limit. • `TCP packet/second limit exceeded for destination`—Number of TCP packets per second for an individual destination address or subnet exceeded limit.

Table 1500: show services service-set statistics screen drops Output Fields *(Continued)*

Field Name	Field Description
`UDP Errors`	Number of UDP protocol errors for the following categories: • `IP data length less than minimum UDP header length (8 bytes)`—Minimum UDP header length is 8 bytes. The received IP packets contained less than 8 bytes. • `Source or destination port is zero`—UDP source or destination port was `0`. • `UDP session-limit exceeded for source`—Number of concurrent UDP sessions from an individual source address or subnet exceeded limit. • `UDP session-limit exceeded for destination`—Number of concurrent UDP sessions to an individual destination address or subnet exceeded limit. • `UDP connections/second limit exceeded for source`—Number of UDP connections per second for an individual source address or subnet exceeded limit. • `UDP connections/second limit exceeded for destination`—Number of UDP connections per second for an individual destination address or subnet exceeded limit. • `UDP packets/second limit exceeded for source`—Number of UDP packets per second for an individual source address or subnet exceeded limit. • `UDP packet/second limit exceeded for destination`—Number of UDP packets per second for an individual destination address or subnet exceeded limit.

Table 1500: show services service-set statistics screen drops Output Fields *(Continued)*

Field Name	Field Description
`ICMP Errors`	Number of ICMP protocol errors for the following categories: • `IP data length less than minimum ICMP header length (8 bytes)`—ICMP header length contained less than 8 bytes. • `ICMP error length inconsistencies`—ICMP error packet length was outside range of 48 bytes through 576 bytes. • `ICMP fragments`— ICMP packet was an IP fragment. • `ICMP session-limit exceeded for source`—Number of concurrent ICMP sessions from an individual source address or subnet exceeded limit. • `ICMP session-limit exceeded for destination`—Number of concurrent ICMP sessions to an individual destination address or subnet exceeded limit. • `ICMP connections/second limit exceeded for source`—Number of ICMP connections per second for an individual source address or subnet exceeded limit. • `ICMP connections/second limit exceeded for destination`—Number of ICMP connections per second for an individual destination address or subnet exceeded limit. • `ICMP packets/second limit exceeded for source`—Number of ICMP packets per second for an individual source address or subnet exceeded limit. • `ICMP packet/second limit exceeded for destination`—Number of ICMP packets per second for an individual destination address or subnet exceeded limit.

Sample Output

show services service-sets statistic screen-drops

```
user@host> show services service-sets statistic screen-drops USF-Service-Set-X interface
vms-0/2/0
Interface: vms-0/2/0
  Service set: sset1
    Errors:
```

```
      IP: 0, TCP: 0
      UDP: 0, ICMP: 0
    IP errors:
      IP packet length inconsistencies: 0
      Illegal source address: 0
      Illegal destination address: 0
      TTL zero errors: 0, Illegal IP protocol number (0 or 255): 0
      Land attack: 0
      Non-IPv4 packets: 0
      Non-IPv6 packets: 0
      Bad checksum: 0
      Illegal IP fragment length: 0
      IP fragment overlap: 0
      IP fragment reassembly timeout: 0
      IP fragment limit exceeded: 0
      IPv4 bad options: 0
      IPv6 bad extension headers: 0
      session-limit exceeded for source: 0
      session-limit exceeded for destination: 0
      connections/second limit exceeded for source: 0
      connections/second limit exceeded for destination: 0
      packets/second limit exceeded for source: 0
      packet/second limit exceeded for destination: 0
      Unknown: 0
    TCP errors:
      TCP header length inconsistencies: 0
      Source or destination port number is zero: 0
      Illegal sequence number and flags combinations: 0
      TCP winnuke: 0
      TCP SYN Fragment: 0
      TCP connection closed due to SYN defense: 0
      TCP session-limit exceeded for source: 0
      TCP session-limit exceeded for destination: 0
      TCP connections/second limit exceeded for source: 0
      TCP connections/second limit exceeded for destination: 0
      TCP packets/second limit exceeded for source: 0
      TCP packet/second limit exceeded for destination: 0
    UDP errors:
      IP data length less than minimum UDP header length (8 bytes): 0
      Source or destination port number is zero: 0
      UDP session-limit exceeded for source: 0
      UDP session-limit exceeded for destination: 0
      UDP connections/second limit exceeded for source: 0
```

```
      UDP connections/second limit exceeded for destination: 0
      UDP packets/second limit exceeded for source: 0
      UDP packet/second limit exceeded for destination: 0
    ICMP errors:
      IP data length less than minimum ICMP header length (8 bytes): 0
      ICMP error length inconsistencies: 0
      ICMP fragments: 0
      ICMP session-limit exceeded for source: 0
      ICMP session-limit exceeded for destination: 0
      ICMP connections/second limit exceeded for source: 0
      ICMP connections/second limit exceeded for destination: 0
      ICMP packets/second limit exceeded for source: 0
      ICMP packet/second limit exceeded for destination: 0
```

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

show services service-sets statistic screen-session-limit-counters (Next Gen Services)

IN THIS SECTION

- Syntax | 26074
- Description | 26074
- Options | 26074
- Required Privilege Level | 26074
- Output Fields | 26074
- Sample Output | 26082

Release Information | 26084

Syntax

```
show services service-set statistic screen-session-limit-counters
<interface interface>
<service-set service-set>
```

Description

Display counters for session drops and packet drops resulting from session-limit checks performed by an IDS rule on an MS-MPC or MS-MIC.

Options

none	Display statistics for all configured services interfaces.
interface *interface-name*	(Optional) Display statistics for the specified services interface.
service *service-set*	Display statistics for the specified service set.

Required Privilege Level

view

Output Fields

Table 1501 on page 26075 lists the output fields for the `show services service-set statistics ids session-limits counters` command. Output fields are listed in the approximate order in which they appear.

Table 1501: show services service-sets statistics ids session-limits counters Output Fields

Field Name	Field Description
`TCP Counters`	Session-limit TCP counters in the ingress direction for the following: • `Sessions allowed`—Number of TCP sessions allowed by the IDS rule. • `Sessions ignored`—Number of TCP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of TCP sessions dropped because the number of TCP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of TCP sessions dropped because the number of TCP connections per second exceeded the limit. • • `Packets allowed`—Number of TCP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of TCP packets dropped because the number of TCP packets per second exceeded the limit.

Table 1501: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`UDP Counters`	Session-limit UDP counters in the ingress direction for the following: • `Sessions allowed`—Number of UDP sessions allowed by the IDS rule. • `Sessions ignored`—Number of UDP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of UDP sessions dropped because the number of UDP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of UDP sessions dropped because the number of UDP connections per second exceeded the limit. • `Packets allowed`—Number of UDP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of UDP packets dropped because the number of TCP packets per second exceeded the limit.

Table 1501: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`ICMP Counters`	Session-limit ICMP counters in the ingress direction for the following: • `Sessions allowed`—Number of ICMP sessions allowed by the IDS rule. • `Sessions ignored`—Number of ICMP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of ICMP sessions dropped because the number of ICMP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of ICMP sessions dropped because the number of ICMP connections per second exceeded the limit. • `Packets allowed`—Number of ICMP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of ICMP packets dropped because the number of ICMP packets per second exceeded the limit.

Table 1501: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Other-Protocols Counters`	Session-limit counters in the ingress direction for protocols other than TCP, UDP, and ICMP for the following: • `Sessions allowed`—Number of sessions allowed by the IDS rule. • `Sessions ignored`—Number of sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of sessions dropped because the number of sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of sessions dropped because the number of connections per second exceeded the limit. • `Packets allowed`—Number of packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of packets dropped because the number of packets per second exceeded the limit.
`Egress General Info`	Information for IDS rules for the service set in the egress direction. • `Match-direction`—Displays `output`. • `Rule name`—Name of the IDS rule. • `Term name`—Name of the term in the IDS rule.

Table 1501: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Egress TCP Counters`	Session-limit TCP counters in the egress direction for the following: • `Sessions allowed`—Number of TCP sessions allowed by the IDS rule. • `Sessions ignored`—Number of TCP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of TCP sessions dropped because the number of TCP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of TCP sessions dropped because the number of TCP connections per second exceeded the limit. • `Packets allowed`—Number of TCP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of TCP packets dropped because the number of TCP packets per second exceeded the limit.

Table 1501: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Egress UDP Counters`	Session-limit UDP counters in the egress direction for the following: • `Sessions allowed`—Number of UDP sessions allowed by the IDS rule. • `Sessions ignored`—Number of UDP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of UDP sessions dropped because the number of UDP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of UDP sessions dropped because the number of UDP connections per second exceeded the limit. • `Packets allowed`—Number of UDP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of UDP packets dropped because the number of TCP packets per second exceeded the limit.

Table 1501: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Egress ICMP Counters`	Session-limit ICMP counters in the egress direction for the following: • `Sessions allowed`—Number of ICMP sessions allowed by the IDS rule. • `Sessions ignored`—Number of ICMP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of ICMP sessions dropped because the number of ICMP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of ICMP sessions dropped because the number of ICMP connections per second exceeded the limit. • `Packets allowed`—Number of ICMP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of ICMP packets dropped because the number of ICMP packets per second exceeded the limit.

Table 1501: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Egress Other-Protocols Counters`	Session-limit counters in the egress direction for protocols other than TCP, UDP, and ICMP for the following: • `Sessions allowed`—Number of sessions allowed by the IDS rule. • `Sessions ignored`—Number of sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of sessions dropped because the number of sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of sessions dropped because the number of connections per second exceeded the limit. • `Packets allowed`—Number of packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of packets dropped because the number of packets per second exceeded the limit.

Sample Output

show services service-sets statistic screen-session-limit-counters

```
user@host> show services service-sets statistic screen-session-limit-counters
IDS Option Name: option-1
---------------------------------------------------------------
TCP Counters:
    Sessions allowed: 0
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
    Packets allowed: 0
    Packets dropped due to high pps: 0
UDP Counters:
    Sessions allowed: 0
    Sessions ignored: 0
```

```
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
    Packets allowed: 0
    Packets dropped due to high pps: 0
ICMP Counters:
    Sessions allowed: 0
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
    Packets allowed: 0
    Packets dropped due to high pps: 0
Other-Protocols Counters:
    Sessions allowed: 0
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
    Packets allowed: 0
    Packets dropped due to high pps: 0

IDS Option Name: option-2
----------------------------------------------------------------
TCP Counters:
    Sessions allowed: 0
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
    Packets allowed: 0
    Packets dropped due to high pps: 0
UDP Counters:
    Sessions allowed: 0
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
    Packets dropped due to high pps: 0
ICMP Counters:
    Sessions allowed: 0
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
    Packets allowed: 0
    Packets dropped due to high pps: 0
Other-Protocols Counters:
    Sessions allowed: 0
```

```
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
    Packets allowed: 0
    Packets dropped due to high pps: 0 Destination session limit       0
```

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

show services service-sets statistics ids session-limits counters

IN THIS SECTION

- Syntax | 26084
- Description | 26084
- Options | 26085
- Required Privilege Level | 26085
- Output Fields | 26085
- Sample Output | 26093
- Release Information | 26095

Syntax

```
show services service-sets statistics ids session-limits counters
<interface interface-name>
```

Description

Display counters for session drops and packet drops resulting from session-limit checks performed by an IDS rule on an MS-MPC or MS-MIC.

Options

none	Display statistics for all configured services interfaces.
interface *interface-name*	(Optional) Display statistics for the specified services interface.

Required Privilege Level

view

Output Fields

Table 1502 on page 26085 lists the output fields for the `show services service-set statistics ids session-limits counters` command. Output fields are listed in the approximate order in which they appear.

Table 1502: show services service-sets statistics ids session-limits counters Output Fields

Field Name	Field Description
`Interface`	Name of the service interface assigned to the service set.
`Service set`	Name of the service set to which the IDS rule is applied.
`Ingress General Info`	Information for IDS rules for the service set in the ingress direction. • `Match-direction`—Displays `input`. • `Rule name`—Name of the IDS rule. • `Term name`—Name of the term in the IDS rule.

Table 1502: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Ingress TCP Counters`	Session-limit TCP counters in the ingress direction for the following: • `Sessions allowed`—Number of TCP sessions allowed by the IDS rule. • `Sessions ignored`—Number of TCP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of TCP sessions dropped because the number of TCP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of TCP sessions dropped because the number of TCP connections per second exceeded the limit. • `Sessions dropped due to suspicious packets`—Number of TCP sessions dropped because suspicious TCP packets were found. • `Packets allowed`—Number of TCP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of TCP packets dropped because the number of TCP packets per second exceeded the limit.

Table 1502: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Ingress UDP Counters`	Session-limit UDP counters in the ingress direction for the following: • `Sessions allowed`—Number of UDP sessions allowed by the IDS rule. • `Sessions ignored`—Number of UDP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of UDP sessions dropped because the number of UDP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of UDP sessions dropped because the number of UDP connections per second exceeded the limit. • `Sessions dropped due to suspicious packets`—Number of UDP sessions dropped because suspicious UDP packets were found. • `Packets allowed`—Number of UDP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of UDP packets dropped because the number of TCP packets per second exceeded the limit.

Table 1502: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Ingress ICMP Counters`	Session-limit ICMP counters in the ingress direction for the following: • `Sessions allowed`—Number of ICMP sessions allowed by the IDS rule. • `Sessions ignored`—Number of ICMP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of ICMP sessions dropped because the number of ICMP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of ICMP sessions dropped because the number of ICMP connections per second exceeded the limit. • `Sessions dropped due to suspicious packets`—Number of ICMP sessions dropped because suspicious ICMP packets were found. • `Packets allowed`—Number of ICMP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of ICMP packets dropped because the number of ICMP packets per second exceeded the limit.

Table 1502: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Ingress Other-Protocols Counters`	Session-limit counters in the ingress direction for protocols other than TCP, UDP, and ICMP for the following: • `Sessions allowed`—Number of sessions allowed by the IDS rule. • `Sessions ignored`—Number of sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of sessions dropped because the number of sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of sessions dropped because the number of connections per second exceeded the limit. • `Sessions dropped due to suspicious packets`—Number of sessions dropped because suspicious packets were found. • `Packets allowed`—Number of packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of packets dropped because the number of packets per second exceeded the limit.
`Egress General Info`	Information for IDS rules for the service set in the egress direction. • `Match-direction`—Displays `output`. • `Rule name`—Name of the IDS rule. • `Term name`—Name of the term in the IDS rule.

Table 1502: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Egress TCP Counters`	Session-limit TCP counters in the egress direction for the following: • `Sessions allowed`—Number of TCP sessions allowed by the IDS rule. • `Sessions ignored`—Number of TCP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of TCP sessions dropped because the number of TCP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of TCP sessions dropped because the number of TCP connections per second exceeded the limit. • `Sessions dropped due to suspicious packets`—Number of TCP sessions dropped because suspicious TCP packets were found. • `Packets allowed`—Number of TCP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of TCP packets dropped because the number of TCP packets per second exceeded the limit.

Table 1502: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Egress UDP Counters`	Session-limit UDP counters in the egress direction for the following: • `Sessions allowed`—Number of UDP sessions allowed by the IDS rule. • `Sessions ignored`—Number of UDP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of UDP sessions dropped because the number of UDP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of UDP sessions dropped because the number of UDP connections per second exceeded the limit. • `Sessions dropped due to suspicious packets`—Number of UDP sessions dropped because suspicious UDP packets were found. • `Packets allowed`—Number of UDP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of UDP packets dropped because the number of TCP packets per second exceeded the limit.

Table 1502: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Egress ICMP Counters`	Session-limit ICMP counters in the egress direction for the following: • `Sessions allowed`—Number of ICMP sessions allowed by the IDS rule. • `Sessions ignored`—Number of ICMP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of ICMP sessions dropped because the number of ICMP sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of ICMP sessions dropped because the number of ICMP connections per second exceeded the limit. • `Sessions dropped due to suspicious packets`—Number of ICMP sessions dropped because suspicious ICMP packets were found. • `Packets allowed`—Number of ICMP packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of ICMP packets dropped because the number of ICMP packets per second exceeded the limit.

Table 1502: show services service-sets statistics ids session-limits counters Output Fields *(Continued)*

Field Name	Field Description
`Egress Other-Protocols Counters`	Session-limit counters in the egress direction for protocols other than TCP, UDP, and ICMP for the following: • `Sessions allowed`—Number of sessions allowed by the IDS rule. • `Sessions ignored`—Number of sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included `accept skip-ids`. • `Sessions dropped due to maximum reached`—Number of sessions dropped because the number of sessions exceeded the limit. • `Sessions dropped due to high rate`—Number of sessions dropped because the number of connections per second exceeded the limit. • `Sessions dropped due to suspicious packets`—Number of sessions dropped because suspicious packets were found. • `Packets allowed`—Number of packets that the IDS rule allowed. • `Packets dropped due to high pps`—Number of packets dropped because the number of packets per second exceeded the limit.

Sample Output

show services service-sets statistics ids session-limits counters interface

```
user@host> show services service-sets statistics ids session-limits counters interface mams-4/0/0
Interface: mams-4/0/0
  Service set: ams_ss1
    Ingress General Info:
      Match-direction: input
      Rule name: ids_rule_1
      Term name: 0
    Ingress TCP Counters:
      Sessions allowed: 1000
      Sessions ignored: 0
      Sessions dropped due to maximum reached: 0
```

```
    Sessions dropped due to high rate: 0
    Sessions dropped due to suspicious packets: 0
    Packets allowed: 1000
    Packets dropped due to high pps: 0
  Ingress UDP Counters:
    Sessions allowed: 1000
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
    Sessions dropped due to suspicious packets: 0
    Packets allowed: 1000
    Packets dropped due to high pps: 0
  Ingress ICMP Counters:
    Sessions allowed: 100
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 50
    Sessions dropped due to high rate: 0
    Sessions dropped due to suspicious packets: 0
    Packets allowed: 100
    Packets dropped due to high pps: 0
  Ingress Other-Protocols Counters:
    Sessions allowed: 0
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
    Sessions dropped due to suspicious packets: 0
    Packets allowed: 0
    Packets dropped due to high pps: 0
  Egress General Info:
    Match-direction: output
  Egress TCP Counters:
    Sessions allowed: 0
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
    Sessions dropped due to suspicious packets: 0
    Packets allowed: 0
    Packets dropped due to high pps: 0
  Egress UDP Counters:
    Sessions allowed: 0
    Sessions ignored: 0
    Sessions dropped due to maximum reached: 0
    Sessions dropped due to high rate: 0
```

```
      Sessions dropped due to suspicious packets: 0
      Packets allowed: 0
      Packets dropped due to high pps: 0
    Egress ICMP Counters:
      Sessions allowed: 0
      Sessions ignored: 0
      Sessions dropped due to maximum reached: 0
      Sessions dropped due to high rate: 0
      Sessions dropped due to suspicious packets: 0
      Packets allowed: 1
      Packets dropped due to high pps: 0
    Egress Other-Protocols Counters:
      Sessions allowed: 0
      Sessions ignored: 0
      Sessions dropped due to maximum reached: 0
      Sessions dropped due to high rate: 0
      Sessions dropped due to suspicious packets: 0
      Packets allowed: 0
      Packets dropped due to high pps: 0
```

Release Information

Command introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

show services service-sets statistics integrity-drops

IN THIS SECTION

- Syntax | 26096
- Description | 26096
- Options | 26096
- Required Privilege Level | 26096

- Output Fields | 26096
- Sample Output | 26100
- Release Information | 26101

Syntax

```
show services service-sets statistics integrity-drops
<interface interface-name>
<service-set service-set-name>
```

Description

Display integrity-drops statistics for one adaptive services interface, for all adaptive services interfaces, or for one service-set. You can configure use the output of this command to verify the packet header for anomalies in IP, TCP, UDP, and IGMP information and to examine any anomalies and errors.

Options

none	Display integrity-drops statistics for all configured adaptive service interfaces/ service-set.
service-set *service-set-name*	(Optional) Display integrity-drops statistics for the specified service-set
interface *interface-name*	(Optional) Display integrity-drops statistics for the specified adaptive services interface.

Required Privilege Level

view

Output Fields

Table 1503 on page 26097 lists the output fields for the `show services service-sets integrity-drops` command. Output fields are listed in the approximate order in which they appear.

Table 1503: show services service-sets integrity-drops Output Fields

Field Name	Field Description
`Interface`	Name of an adaptive services interface.
`Service set`	Name of a service set.
`Errors`	Total errors, categorized by protocol: • `IP`—Total IP version 4 errors. • `TCP`—Total Transmission Control Protocol (TCP) errors. • `UDP`—Total User Datagram Protocol (UDP) errors. • `ICMP`—Total Internet Control Message Protocol (ICMP) errors.

Table 1503: show services service-sets integrity-drops Output Fields *(Continued)*

Field Name	Field Description
IP Errors	IPv4 errors: • `IP packet length inconsistencies`—IP packet length does not match the Layer 2 reported length. • `Minimum IP header length check failures`—Minimum IP header length is 20 bytes. The received packet contains less than 20 bytes. • `Reassembled packet exceeds maximum IP length`—After fragment reassembly, the reassembled IP packet length exceeds 65,535. • `Illegal source address 0`—Source address is not a valid address. Invalid addresses are, loopback, broadcast, multicast, and reserved addresses. Source address `0`, however, is allowed to support BOOTP and the destination address `0xffffffff`. • `Illegal destination address` —Destination address is not a valid address. The address is reserved. • `TTL zero errors`—Received packet had a time-to-live (TTL) value of 0. • `Illegal IP protocol number 0 or 255`—IP protocol is 0 or 255. • `Land attack`—IP source address is the same as the destination address. • `Non-IP packets`—Packet did not conform to the IP standard. • `IP option`—Packet dropped because of a nonallowed IP option. • `Non-IPv4 packets`—Packet was not of the IPv4 type. • `Non-IPv6 packets`—Packet was not of the IPv6 type. • `Bad checksum`—Packet had an invalid IP checksum. • `Illegal IP fragment length`—Illegal fragment length. All fragments (other than the last fragment) must have a length that is a multiple of 8 bytes. • `IP fragment overlap`—Fragments have overlapping fragment offsets. • `IP fragment limit exceeded:` —Fragments dropped because the configured number of allowed fragments for a packet was exceeded.

Table 1503: show services service-sets integrity-drops Output Fields *(Continued)*

Field Name	Field Description
	• `IP fragment reassembly timeout`—Some of the fragments for an IP packet were not received in time, and the reassembly handler dropped partial fragments. Whenever a fragment is received, it is maintained in a chain until all other fragments are received. If other fragments do not arrive within the configured value of reassembly-timeout, this packet is dropped and the value of the counter shown in this field is incremented. If other fragments arrive in time but the total number of fragments is more than the configured value of fragment-limit, all the fragments (of this packet) are dropped and the value of the counter shown in this field is incremented. • `Unknown:` —Unknown fragments.
`TCP Errors`	TCP protocol errors: • `TCP header length inconsistencies`—Minimum TCP header length is 20 bytes, and the IP packet received does not contain at least 20 bytes. • `Source or destination port number is zero`—TCP source or destination port is zero. • `Illegal sequence number, flags combination`—Dropped because of TCP errors, such as an illegal sequence number, which causes an illogical combination of flags to be set.
`UDP Errors`	UDP protocol errors: • `IP data length less than minimum UDP header length (8 bytes)`—Minimum UDP header length is 8 bytes. The received IP packets contain less than 8 bytes. • `Source or destination port is zero`—UDP source or destination port is `0`.

Table 1503: show services service-sets integrity-drops Output Fields *(Continued)*

Field Name	Field Description
`ICMP Errors`	ICMP protocol errors: • `IP data length less than minimum ICMP header length (8 bytes)`—ICMP header length is 8 bytes. This counter is incremented when received IP packets contain less than 8 bytes. • `ICMP error length inconsistencies`—Minimum length of an ICMP error packet is 48 bytes, and the maximum length is 576 bytes. This counter is incremented when the received ICMP error falls outside this range.

Sample Output

show services service-sets statistics integrity-drops

```
user@host> show services service-sets statistics integrity-drops
 Interface: ms-1/0/0
  Service set: sset1
    Errors:
      IP: 0, TCP: 0
      UDP: 0, ICMP: 0
    IP errors:
      IP packet length inconsistencies: 0
      Illegal source address: 0
      Illegal destination address: 0
      TTL zero errors: 0, Illegal IP protocol number (0 or 255): 0
      Land attack: 0
      Non-IPv4 packets: 0
      Non-IPv6 packets: 0
      Bad checksum: 0
      Illegal IP fragment length: 0
      IP fragment overlap: 0
      IP fragment limit exceeded: 0
      IP fragment reassembly timeout: 0
      Unknown: 0
    TCP errors:
      TCP header length inconsistencies: 0
      Source or destination port number is zero: 0
```

```
    Illegal sequence number and flags combinations: 0
  UDP errors:
    IP data length less than minimum UDP header length (8 bytes): 0
    Source or destination port number is zero: 0
  ICMP errors:
    IP data length less than minimum ICMP header length (8 bytes): 0
    ICMP error length inconsistencies: 0
```

Release Information

Command introduced in Junos OS Release 13.1

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

clear services service-sets statistics integrity-drops | 4166

show services service-sets statistics jflow-log

IN THIS SECTION

- Syntax | 26102
- Description | 26102
- Options | 26102
- Required Privilege Level | 26102
- Output Fields | 26102
- Sample Output | 26105
- Release Information | 26111

Syntax

```
show services service-sets statistics jflow-log
<interface interface-name>
<service-set service-set-name>
<brief | detail>
```

Description

Display statistical information on the logs or records generated in flow monitoring format with optional filtering by interface and service set name..

Options

none	Display the statistical details on flow monitoring logs for NAT events for all services interfaces and all service sets.
brief	(Default) (Optional) Display abbreviated flow monitoring log statistics.
detail	(Optional) Display detailed flow monitoring log statistics.
interface *interface-name*	(Optional) Display the flow monitoring log statistics for the specified adaptive service interface. On M Series and T Series routers, *interface-name* can be ms-*fpc*/*pic*/*port*. It is supported only on MS-MICs and MS-MPCs.
service-set *service-set name*	(Optional) Display the flow monitoring log statistics for the specified named service-set.

Required Privilege Level

view

Output Fields

Table 1504 on page 26103 lists the output fields for the `show services service-sets statistics jflow-log` command. Output fields are listed in the approximate order in which they appear.

Table 1504: show services service-sets statistics jflow-log Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of a services interface.	all
`Rate limit`	Maximum number of NAT error events for which records in flow monitoring format must be recorded.	all
`Template records`	Details of the template records in flow monitoring log messages.	all
`Sent`	Number of template records sent to a collector	all
`Messages dropped`	Number of template records dropped while transmission to a collector.	all
`Data records`	Details of the data records in flow monitoring log messages.	all
`Sent`	Number of data records sent to a collector.	all
`Dropped`	Number of data records dropped while transmission to a collector	all
`Service set`	Name of a service set.	all
`Unresolvable collectors`	Number of collectors that cannot be traced and be reached to export records for NAT events.	all

Table 1504: show services service-sets statistics jflow-log Output Fields *(Continued)*

Field Name	Field Description	Level of Output
class name	Logs created for events for each of the following classes: • NAT44 Session logs—Details of logs created for NAT44 sessions • NAT64 Session logs—Details of logs created for NAT64 sessions • NAT44 BIB logs—Details of logs created for NAT44 binding information bases, which is a table of bindings kept by a NAT44. Each NAT44 has a BIB for each translated protocol. • NAT64 BIB logs—Details of logs created for NAT44 binding information bases, which is a table of bindings kept by a NAT64. Each NAT64 has a BIB for each translated protocol. • NAT Address Exhausted logs—Details of logs created for exhaustion of NAT addresses • NAT Port Exhausted logs—Details of logs created for exhaustion of NAT pool • NAT44 Quota Exceeded logs—Details of logs created when allocated quota is exceeded for NAT44 events • NAT64 Quota Exceeded logs—Details of logs created when allocated quota is exceeded for NAT64 events • NAT44 Address Bind logs—Details of logs generated for address bindings for NAT44 events • NAT64 Address Bind logs—Details of logs generated for address bindings for NAT64 events • NAT44 PBA logs—Details of logs generated for NAT44 port block allocation events • NAT64 PBA logs—Details of logs generated for NAT64 port block allocation events The following information is displayed for flow monitoring log messages for each class of event that is logged:	detail

Table 1504: show services service-sets statistics jflow-log Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Template records`—Details of the template records in flow monitoring log messages • `Sent`—Number of template records sent to a collector • `Dropped`—Number of template records dropped while transmission to a collector • `Data records`—Details of the data records in flow monitoring log messages • `Sent`—Number of data records sent to a collector • `Dropped`—Number of data records dropped while transmission to a collector. Counts are provided for the drop reasons • `socket send error`—Number of times a socket was not opened for data transmission • `no memory`—Number of messages dropped because of insufficient memory • `invalid data`—Number of messages dropped because of invalid data in the records • `above rate limit`—The maximum number of flow monitoring log messages per second was exceeded.	

Sample Output

show services service-sets statistics jflow-log brief

```
user@host> show services service-sets statistics jflow-log brief
Interface: ms-5/0/0
  Rate limit: 1000
  Template records:
    Sent: 36
    Dropped: 0
  Data records:
```

```
    Sent: 2
    Dropped: 0

    Service-set: sset_44
      Unresolvable collectors: 0
      Template records:
        Sent: 36
        Dropped: 0
      Data records:
        Sent: 2
        Dropped: 0
```

show services service-sets statistics jflow-log detail

```
user@host> show services service-sets statistics jflow-log detail
Interface: ms-5/0/0
  Rate limit: 1000
  Template records:
    Sent: 48
    Dropped: 0
  Data records:
    Sent: 4
    Dropped: 0

    Service-set: sset_44
      Unresolvable collectors: 0
      Template records:
        Sent: 48
        Dropped: 0
      Data records:
        Sent: 4
        Dropped: 0
      NAT44 Session logs:
        Template records:
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 4
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT64 Session logs:
```

```
        Template records:
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT44 BIB logs:
        Template records:
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT64 BIB logs:
        Template records:
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT Address Exhausted logs:
        Template records:
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT Port Exhausted logs:
        Template records:
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT44 Quota Exceeded logs:
        Template records:
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT64 Quota Exceeded logs:
        Template records:
```

```
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT44 Address Bind logs:
        Template records:
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT64 Address Bind logs:
        Template records:
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT44 PBA logs:
        Template records:
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT64 PBA logs:
        Template records:
          Sent: 4
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
```

show services service-sets statistics jflow-log service-set

```
user@host> show services service-sets statistics jflow-log service-set sset_44
Interface: ms-5/0/0

    Service-set: sset_44
```

```
      Unresolvable collectors: 0
      Template records:
        Sent: 72
        Dropped: 0
      Data records:
        Sent: 4
        Dropped: 0
```

show services service-sets statistics jflow-log service-set detail

```
user@host> show services service-sets statistics jflow-log service-set sset_44 detail
Interface: ms-5/0/0

    Service-set: sset_44
      Unresolvable collectors: 0
      Template records:
        Sent: 84
        Dropped: 0
      Data records:
        Sent: 4
        Dropped: 0
      NAT44 Session logs:
        Template records:
          Sent: 7
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 4
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT64 Session logs:
        Template records:
          Sent: 7
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
          Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
      NAT44 BIB logs:
        Template records:
          Sent: 7
          Dropped: 0 (socket send error: 0, no memory: 0)
        Data records:
          Sent: 0
```

```
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 BIB logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT Address Exhausted logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT Port Exhausted logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 Quota Exceeded logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 Quota Exceeded logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 Address Bind logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
```

```
    NAT64 Address Bind logs:
      Template records:
        Sent: 7
        Dropped: 0 (socket send error: 0, no memory: 0)
      Data records:
        Sent: 0
        Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
    NAT44 PBA logs:
      Template records:
        Sent: 7
        Dropped: 0 (socket send error: 0, no memory: 0)
      Data records:
        Sent: 0
        Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
    NAT64 PBA logs:
      Template records:
        Sent: 7
        Dropped: 0 (socket send error: 0, no memory: 0)
      Data records:
        Sent: 0
        Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

clear services service-sets statistics syslog | 4171

show services service-sets statistics packet-drops

IN THIS SECTION

- Syntax | 26112
- Description | 26112

- Options | 26112
- Required Privilege Level | 26112
- Output Fields | 26112
- Sample Output | 26113
- Release Information | 26113

Syntax

```
show services service-sets statistics packet-drops
<interface interface-name>
```

Description

Display the number of dropped packets for service sets exceeding CPU limits or memory limits.

Options

none	Display the number of dropped service sets packets for all adaptive services interfaces.
interface ***interface-name***	(Optional) Display the number of dropped service sets packets for a particular interface. On M Series and T Series routers, *interface-name* can be *ms-fpc/pic/port*, *sp-fpc/pic/port*, or *rspnumber*.

Required Privilege Level

view

Output Fields

Table 1505 on page 26113 lists the output fields for the `show services service-sets packet-drops` command. Output fields are listed in the approximate order in which they appear.

Table 1505: show services service-sets packet-drops Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of a service set.
CPU limit Drops	Number of packets dropped because the service set exceeded the average CPU limit.
Memory limit Drops	Number of packets dropped because the service set exceeded the memory limit.
Flow limit Drops	Number of packets dropped because the service set exceeded the flow limit.

Sample Output

show services service-sets statistics packet-drops

```
user@host> show services service-sets statistics packet-drops
Interface: vms-1/0/0
  Service set: ss1
    CPU limit drops: 0
    Memory limit drops: 0
    Flow limit drops: 0
```

Release Information

Command introduced in Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

clear services flow-collector statistics | 4094

show services service-sets statistics syslog

IN THIS SECTION

- Syntax | 26114
- Description | 26114
- Options | 26114
- Required Privilege Level | 26115
- Output Fields | 26115
- Sample Output | 26120
- Sample Output | 26120
- For Next Gen Services MX-SPC3 Services Card | 26121
- Release Information | 26122

Syntax

```
show services service-sets statistics syslog
<interface interface-name>
<service-set service-set-name>
<brief | detail>
```

Description

Display the system log statistics with optional filtering by interface and service set name.

Options

none	Display the system log statistics for all services interfaces and all service sets.
brief	(Default) (Optional) Display abbreviated system log statistics.
detail	(Optional) Display detailed system log statistics.

interface ***interface-name***	(Optional) Display the system log statistics for a specific adaptive service interface. On M Series and T Series routers, *interface-name* can be `ms-fpc/pic/port`, `sp-fpc/pic/port`, or `rspnumber`.
service-set ***service-set-name***	(Optional) Display the system log statistics for a specific named service-set.

Required Privilege Level

view

Output Fields

Table 1506 on page 26115 lists the output fields for the `show services service-sets statistics syslog` command. Output fields are listed in the approximate order in which they appear.

Table 1506: show services service-sets statistics syslog Output Fields

Field Name	Field Description	Level
`Interface`	Name of a services interface.	all
`Rate limit`	Maximum number of messages per second written to the interface's system log.	all
`Sent`	Number of messages sent that are not associated with a service set.	all
`Dropped`	Number of messages dropped that are not associated with a service set.	all
Service-set		
`Service-set`	Name of a service set.	all
`Sent`	Number of sent messages that are associated with the service set.	all

Table 1506: show services service-sets statistics syslog Output Fields *(Continued)*

Field Name	Field Description	Level
`Dropped`	Number of dropped messages that are associated with the service set.	all

`Session open logs`	The following information is displayed for system log messages for session open events that are logged and are associated with the service set:
	• `Sent`—Number of messages sent.
	• `Dropped`—Number of messages dropped. Counts are given for these drop reasons:
	  • `low priority`—Priority of the message was too low for the message to be sent.
	  • `no class set`—Specific classes of event messages were configured and this class was not selected.
	  • `above rate limit`—Maximum number of system log messages per second was exceeded.
		`detail`

`Session close logs`	The following information is displayed for system log messages for session close events that are logged and are associated with the service set:
	• `Sent`—Number of messages sent.
	• `Dropped`—Number of messages dropped. Counts are given for these drop reasons:
	  • `low priority`—Priority of the message was too low for the message to be sent.
	  • `no class set`—Specific classes of event messages were configured and this class was not selected.
	  • `above rate limit`—Maximum number of system log messages per second was exceeded.
		`detail`

Table 1506: show services service-sets statistics syslog Output Fields *(Continued)*

Field Name	Field Description	Level

`Packet logs`	The following information is displayed for system log messages for packet events that are logged and are associated with the service set:
	• `Sent`—Number of messages sent.
	• `Dropped`—Number of messages dropped. Counts are given for these drop reasons:
	  • `low priority`—Priority of the message was too low for the message to be sent.
	  • `no class set`—Specific classes of event messages were configured and this class was not selected.
	  • `above rate limit`—Maximum number of system log messages per second was exceeded.
		`detail`

`Stateful firewall logs`	The following information is displayed for system log messages for stateful firewall events that are logged and are associated with the service set:
	• `Sent`—Number of messages sent.
	• `Dropped`—Number of messages dropped. Counts are given for these drop reasons:
	  • `low priority`—Priority of the message was too low for the message to be sent.
	  • `no class set`—Specific classes of event messages were configured and this class was not selected.
	  • `above rate limit`—Maximum number of system log messages per second was exceeded.
		`detail`

Table 1506: show services service-sets statistics syslog Output Fields *(Continued)*

Field Name	Field Description	Level

`ALG logs`	The following information is displayed for system log messages for ALG events that are logged and are associated with the service set:
	• `Sent`—Number of messages sent.
	• `Dropped`—Number of messages dropped. Counts are given for these drop reasons:
	  • `low priority`—Priority of the message was too low for the message to be sent.
	  • `no class set`—Specific classes of event messages were configured and this class was not selected.
	  • `above rate limit`—Maximum number of system log messages per second was exceeded.
		`detail`

`NAT logs`	The following information is displayed for system log messages for NAT events that are logged and are associated with the service set:
	• `Sent`—Number of messages sent.
	• `Dropped`—Number of messages dropped. Counts are given for these drop reasons:
	  • `low priority`—Priority of the message was too low for the message to be sent.
	  • `no class set`—Specific classes of event messages were configured and this class was not selected.
	  • `above rate limit`—Maximum number of system log messages per second was exceeded.
		`detail`

Table 1506: show services service-sets statistics syslog Output Fields *(Continued)*

Field Name	Field Description	Level

`IDS logs`	The following information is displayed for system log messages for IDS events that are logged and are associated with the service set:
	• `Sent`—Number of messages sent.
	• `Dropped`—Number of messages dropped. Counts are given for these drop reasons:
	  • `low priority`—Priority of the message was too low for the message to be sent.
	  • `no class set`—Specific classes of event messages were configured and this class was not selected.
	  • `above rate limit`—Maximum number of system log messages per second was exceeded.
		`detail`

`Other logs`	The following information is displayed for system log messages for other types of events that are logged and are associated with the service set:
	• `Sent`—Number of messages sent.
	• `Dropped`—Number of messages dropped. Counts are given for these drop reasons:
	  • `low priority`—Priority of the message was too low for the message to be sent.
	  • `no class set`—Specific classes of event messages were configured and this class was not selected.
	  • `above rate limit`—Maximum number of system log messages per second was exceeded.
		`detail`

Sample Output

show services service-sets statistics syslog brief

```
user@host> show services service-sets statistics syslog brief
Interface: sp-1/1/0
  Rate limit: 200000
  Sent: 0
  Dropped: 0
  Service-set: sset-sfw-sp1
    Sent: 20
    Dropped: 3488
  Service-set: sset-nat-sp1
    Sent: 18
    Dropped: 91
Interface: sp-1/2/0
  Rate limit: 15000
  Sent: 0
  Dropped: 0
  Service-set: sset-sfw-sp2
    Sent: 210
    Dropped: 579
```

Sample Output

show services service-sets statistics syslog detail

```
user@host> show services service-sets statistics syslog detail

Interface: ms-2/1/0
  Rate limit: 0
  Sent: 0
  Dropped: 0
  Service-set: sset1
    Sent: 0
    Dropped: 0
    Session open logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
    Session close logs:
```

```
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
    Packet logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
    Stateful firewall logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
    ALG logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
    NAT logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
    IDS logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
    PCP MAP logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
    PCP protocol logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
    PCP protocol error logs:
Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
    PCP debug logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
    Other logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate limit: 0)
```

For Next Gen Services MX-SPC3 Services Card

Following shows the output for the `show services service-sets statistics syslog` on the MX-SPC3 services cards `vms-x/y/z` interfaces.

command-name

```
user@host> show services service-sets statistics syslog
show services service-sets statistics syslog
Log Module Statistics
Interface-Name- vms-2/0/0
Service-set Name- Sset1
Name                Generated       Discarded
--------------------------------------------
UTM                 0               0
FW_AUTH             0               0
SCREEN              0               0
ALG                 0               0
NAT                 0               0
FLOW                0               0
SCTP                0               0
GTP                 0               0
IPSEC               0               0
IDP                 0               0
RTLOG               0               0
PST_DS_LITE         0               0
APPQOS              0               0
SECINTEL            0               0
AAMW                0               0
OTHERS              0               0

Log stream Statistics
Interface-Name- vms-2/0/0
Service-set Name- Sset1
Name                send            Fail
---------------------------------------------
database            0               0
```

Release Information

Command introduced in Junos OS Release 11.1.

Support for this command introduced in Junos OS Release 19.3R2 for Next Gen Services with the MX-SPC3 services card on MX240, MX480 and MX960 routers.

RELATED DOCUMENTATION

clear services service-sets statistics syslog | 4171

show services service-sets statistics tcp

IN THIS SECTION

- Syntax | 26123
- Description | 26123
- Options | 26123
- Required Privilege Level | 26124
- Output Fields | 26124
- Sample Output | 26124
- Release Information | 26124

Syntax

```
show services service-sets statistics tcp
<interface interface-name>
<service-set service-set-name>
```

Description

Display TCP-related statistics.

Options

`interface` *`interface-name`*	Name of adaptive services interface.
`service-set` *`service-set-name`*	Name of service set.

Required Privilege Level

view

Output Fields

Sample Output

show services service-sets statistics tcp

```
user@host> show services service-sets statistics tcp
Interface:vms-0/2/0
  Service set: ss1_interface_style1
    TCP open/close statistics:
      TCP first packet non-syn: 1
      TCP first packet reset: 0
      TCP first packet FIN: 0
      TCP non syn discard: 0
      TCP extension alloc fail: 0
      TFO SYN with cookie request: 0
      TFO SYN with cookie: 0
      TFO SYN ACK with cookie: 0
      TFO packets forwarded: 0
      TFO packets dropped: 0
      TFO packets stripped: 0
      TCP invalid syn ack: 0
      TCP invalid ack window check: 0
      TCP invalid syn transmit: 0
      TCP invalid reset in listen: 0
      TCP invalid reset in syn received: 0
      TCP invalid reset in syn sent: 0
      TCP invalid flags handshake: 0
    TCP MSS statistics:
      TCP SYN MSS Received: 0
      TCP SYN MSS Modified: 0
```

Release Information

Command introduced in Junos OS Release 17.2.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Configuring TFO

show services service-sets statistics tcp-mss

IN THIS SECTION

- Syntax | 26125
- Description | 26125
- Options | 26125
- Required Privilege Level | 26126
- Output Fields | 26126
- Sample Output | 26126
- Release Information | 26126

Syntax

```
show services service-sets statistics tcp-mss
<interface interface-name>
```

Description

(M Series and T Series routers only) Display TCP maximum segment size (MSS) statistics for service sets.

Options

none Display service set TCP MSS information for all adaptive services interfaces.

interface *interface-name*

(Optional) Display TCP MSS statistics for a particular interface. The *interface-name* can be ms-*fpc*/*pic*/*port*, sp-*fpc*/*pic*/*port*, or rsp *number*.

Required Privilege Level

view

Output Fields

Table 1507 on page 26126 lists the output fields for the show services service-sets statistics tcp-mss command. Output fields are listed in the approximate order in which they appear.

Table 1507: show services service-sets statistics tcp-mss Output Fields

Field Name	Field Description
Interface	Name of the adaptive services interface.
Service Set	Name of the configured service set.
SYN Received	Number of TCP SYN packets received.
SYN Modified	Number of TCP SYN packets with the MSS value modified to match the MSS value specified in the TCP MSS configuration.

Sample Output

show services service-sets statistics tcp-mss

```
user@host> show services service-sets statistics tcp-mss
Interface   Service Set                      SYN Received  SYN Modified
sp-1/2/0    asq_ipsec_svc_0                           500           220
```

Release Information

Command introduced in Junos OS Release 9.5.

show services service-sets summary

IN THIS SECTION

- Syntax | 26127
- Description | 26127
- Options | 26127
- Required Privilege Level | 26128
- Output Fields | 26128
- Sample Output | 26128
- Release Information | 26129

Syntax

```
show services service-sets summary
<interface interface-name>
```

Description

Display service set summary information.

Options

none

Display service set summary information for all adaptive services interfaces.

interface *interface-name*

(Optional) Display service set summary information for a particular interface. On M Series and T Series routers, *interface-name* can be `ms-`*fpc*/*pic*/*port*, `sp-`*fpc*/*pic*/*port*, or `rsp`*number*.

On MX Series MX240, MX480, and MX960 routers, *interface-name* can be `vms-`*fpc*/*pic*/*port* for the MX-SPC3 services card for Next Gen Services.

Required Privilege Level

view

Output Fields

Table 1508 on page 26128 lists the output fields for the `show services service-sets summary` command. Output fields are listed in the approximate order in which they appear.

Table 1508: show services service-sets summary Output Fields

Field Name	Field Description
`Interface`	Name of an adaptive services interface
`Service type`	Type of adaptive service, such as stateful firewall (SFW), Network Address Translation (NAT), intrusion detection service (IDS), Layer 2 Tunneling Protocol (L2TP), Compressed Real-Time Transport Protocol (CRTP), or IP Security (IPsec)
`Service sets configured`	Total number of service sets configured on the PIC that use internal service set IDs and do not consume external service sets, including CRTP and L2TP
`Bytes used`	Bytes used by a particular service or all services
`Policy bytes used`	Policy bytes used by a particular service or all services
`CPU utilization`	Percentage of the CPU resources being used

Sample Output

show services service-sets summary

```
user@host> show services service-sets summary
                 Service sets
CPU
Interface    configured            Bytes used        Session bytes used     Policy bytes
```

```
used      utilization
vms-3/0/0          1           3453621040 (24.93%)          0 ( 0.00%)       8161168
( 0.90%)        0.14 %
```

show services service-sets summary interface

```
user@host> show services service-sets summary interface sp-1/3/0
Interface: sp-1/3/0
                  Service sets                          CPU
  Service type    configured              Bytes used   utilization
  SFW/NAT/IDS                1           54 ( 0.00 %)        N/A
  L2TP                       1           58 ( 0.00 %)        N/A
  CRTP                       1           58 ( 0.00 %)        N/A
  System                     0       920831 ( 0.44 %)        N/A
  Idle                       0            0 ( 0.00 %)        N/A
  Total                      3       921001 ( 0.44 %)        N/A
```

Release Information

Command introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

show services sessions

IN THIS SECTION

- Syntax | 26130
- Description | 26130
- Options | 26130
- Required Privilege Level | 26133
- Output Fields | 26133
- Sample Output | 26134
- Release Information | 26140

Syntax

```
show services sessions
<brief | extensive | terse>
<application-protocol protocol>
<count>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<nat number>
<nat-port-overload-index number>
<protocol protocol>
<service-set service-set>
<source-nat-pool source-nat-pool-name>
<source-port source-port>
<source-prefix source-prefix>
<utilization>
```

Description

Display session information.

> **NOTE**: On MX Series routers (with interchassis redundancy configured), the idle timeout for every flow is displayed in the `show services session extensive` and `show services flows extensive` commands.

Options

none	Display standard information about all sessions.
brief | extensive | terse	(Optional) Display the specified level of output.
application-protocol *protocol*	(Optional) Display information about one of the following application protocols: • `bootp`—Bootstrap protocols • `dce-rpc`—Distributed Computing Environment-Remote Procedure Call protocols

- `dce-rpc-portmap`—Distributed Computing Environment-Remote Procedure Call protocols portmap service
- `dns`—Domain Name System protocol
- `exec`—Remote Execution Protocol
- `ftp`—File Transfer Protocol
- `h323`—H.323
- `icmp`—ICMP
- `icmpv6`—ICMPv6
- `iiop`—Internet Inter-ORB Protocol
- `ike-esp-nat`—IKE ALG
- `ip`—IP
- `login`—LOGIN
- `netbios`—NETBIOS
- `netshow`—NETSHOW
- `pptp`—Point-to-Point Tunneling Protocol
- `realaudio`—RealAudio
- `rpc`—Remote Procedure Call protocol
- `rpc-portmap`—Remote Procedure Call protocol portmap service
- `rtsp`—Real-Time Streaming Protocol
- `rsh`—Remote Shell
- `sip`—Session Initiation Protocol
- `shell`—Shell
- `snmp`—SNMP
- `sql`—SQLNet
- `talk`—Talk Program

- `tftp`—Trivial File Transfer Protocol
- `traceroute`—Traceroute
- `winframe`—WinFrame

> **NOTE**: You can use the `none` option with the `show services sessions count application-protocol` command to display information about sessions other than ALG sessions.

count
(Optional) Display a count of the matching entries.

destination-port *destination-port*
(Optional) Display information for the specified destination port. The range of values is from 0 to 65,535.

destination-prefix *destination-prefix*
(Optional) Display information for the specified destination prefix.

interface *interface-name*
(Optional) Display information about the specified interface. On M Series and T Series routers, *`interface-name`* can be `ms-`*`fpc`*`/`*`pic`*`/`*`port`* or `rsp`*`number`*. On J Series routers, *`interface-name`* is `ms-`*`pim`*`/0/`*`port`*.

limit *number*
(Optional) Maximum number of entries to display.

nat
Display sessions with network address translation.

nat-port-overload-index
Displays NAT port overload index; the range is 0 through 127.

protocol *protocol*
(Optional) Display information about one of the following IP types:

- *`number`*—Numeric protocol value from 0 to 255
- `ah`—IPsec Authentication Header protocol
- `egp`—An exterior gateway protocol
- `esp`—IPsec Encapsulating Security Payload protocol
- `gre`—A generic routing encapsulation protocol
- `icmp`—Internet Control Message Protocol
- `icmp6`—Internet Control Message Protocol version 6
- `igmp`—Internet Group Management Protocol

- `ipip`—IP-within-IP Encapsulation Protocol
- `ospf`—Open Shortest Path First protocol
- `pim`—Protocol Independent Multicast protocol
- `rsvp`—Resource Reservation Protocol
- `sctp`—Stream Control Transmission Protocol
- `tcp`—Transmission Control Protocol
- `udp`—User Datagram Protocol

service-set *service-set* (Optional) Display information for the specified service set.

source-nat-pool Displays the source NAT pool name.

source-port *source-port* (Optional) Display information for the specified source port. The range of values is from 0 to 65,535.

source-prefix *source-prefix* (Optional) Display information for the specified source prefix.

utilization (Optional) Display statistical details about session utilization.

Required Privilege Level

view

Output Fields

Table 1509 on page 26133 lists the output fields for the `show services sessions` command. Output fields are listed in the approximate order in which they appear.

Table 1509: show services sessions Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the interface.	`application-protocol`
`Session`	Session ID that uniquely identifies the session.	All levels

Table 1509: show services sessions Output Fields *(Continued)*

Field Name	Field Description	Level of Output
ALG	Name of the application.	terse
Flags	Session flag for the ALG: • 0x1—Found an existing session. • 0x2—Reached session or flow limit. • 0x3—No memory available for new sessions. • 0x4—No free session ID available. • 0x0000—No session ID found.	All levels
IP Action	Flag indicating whether IP action has been set for the session.	All levels
Offload	Flag indicating whether the session has been offloaded to the Packet Forwarding Engine.	All levels
Asymmetric	Flag indicating whether the session is uni-directional.	terse application-protocol
Service set	Name of a service set. Individual empty service sets are not displayed.	count
Sessions Count	Number of sessions.	count

Sample Output

show services sessions

```
user@host> show services sessions
ms-2/0/0
Session: 293, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
```

```
UDP        10.10.10.2:43677 ->    10.20.20.1:53    Forward  I           1
UDP       10.20.20.1:53    ->        192.0.2.1:43677 Forward  O           1
Session: 53, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        10.10.10.2:37494 ->    10.20.20.1:53    Forward  I           1
UDP       10.20.20.1:53    ->        10.11.11.11:37494 Forward  O           1
Session: 66, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        10.10.10.2:48161 ->    10.20.20.1:53    Forward  I           1
UDP       10.20.20.1:53    ->        10.11.11.11:48161 Forward  O           1
Session: 17, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        10.10.10.2:38908 ->    10.20.20.1:53    Forward  I           1
UDP       10.20.20.1:53    ->        10.11.11.11:38908 Forward  O           1
Session: 42, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        10.10.10.2:58189 ->    10.20.20.1:53    Forward  I           1
UDP       10.20.20.1:53    ->        10.11.11.11:58189 Forward  O           1
```

show services sessions brief

The output for the `show services flows brief` command is identical to that for the `show services sessions` command. For sample output, see "show services sessions" on page 26134.

show services sessions extensive

```
user@host> show services sessions extensive
ms-0/1/0
Session: 2, ALG: 0, Flags: 0x0080, IP Action: no, Offload: no
NAT PLugin Data:
  NAT Action:   Translation Type - DYNAMIC NAT44
    NAT source         192.0.21.2         ->    10.10.10.127
TCP            192.0.2.2:52145 ->        198.51.100.2:23    Forward  I              22
  Byte count: 1483
  Flow role: Unknown, Timeout: 0
TCP            198.51.100.2:23    ->   10.10.10.127:52145 Forward  O              18
  Byte count: 2712
  Flow role: Unknown, Timeout: 0
```

show services sessions terse

```
user@router> show services sessions terse
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              33
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              31
```

show services sessions application-protocol

This command has the same output for the rpc, dce-rpc, rpc-portmap and dce-rpc-portmap ALGs.

```
user@router> show services sessions application-protocol dce-rpc
Interface name: ms-1/1/0
Session: 8, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:1019  ->192.168.203.194:2049  Forward  I               4
UDP    192.168.203.194:2049  ->192.168.203.198:1019  Forward  O               4
Session: 7, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:954   ->192.168.203.194:613   Forward  I               1
UDP    192.168.203.194:613   ->192.168.203.198:954   Forward  O               1
Session: 6, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:53836 ->192.168.203.194:613   Forward  I               1
UDP    192.168.203.194:613   ->192.168.203.198:53836 Forward  O               1
Session: 5, ALG: portmapper, Flags: 0x1000, IP Action: no, Offload: no
UDP    192.168.203.198:59813 ->192.168.203.194:111   Forward  I               1
UDP    192.168.203.194:111   ->192.168.203.198:59813 Forward  O               1
Session: 4, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:36595 ->192.168.203.194:2049  Forward  I               1
UDP    192.168.203.194:2049  ->192.168.203.198:36595 Forward  O               1
Session: 3, ALG: portmapper, Flags: 0x1000, IP Action: no, Offload: no
UDP    192.168.203.198:56050 ->192.168.203.194:111   Forward  I               1
UDP    192.168.203.194:111   ->192.168.203.198:56050 Forward  O               1

user@router> show services sessions application-protocol dns
Interface name: ms-2/0/0
Session: 293, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        198.51.100.2:43677 ->    203.0.113.10:53    Forward  I           1
UDP       203.0.113.10:53    ->        192.0.2.1:43677 Forward  O           1
Session: 53, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        198.51.100.2:37494 ->    203.0.113.10:53    Forward  I           1
```

```
UDP       203.0.113.10:53    ->        192.0.2.1:37494 Forward  O           1
Session: 66, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        198.51.100.2:48161 ->    203.0.113.10:53    Forward  I           1
UDP       203.0.113.10:53    ->        192.0.2.1:48161 Forward  O           1
Session: 17, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        198.51.100.2:38908 ->    203.0.113.10:53    Forward  I           1
UDP       203.0.113.10:53    ->        192.0.2.1:38908 Forward  O           1
Session: 42, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        198.51.100.2:58189 ->    203.0.113.10:53    Forward  I           1
UDP       203.0.113.10:53    ->        192.0.2.1:58189 Forward  O           1
user@router> show services sessions application-protocol ftp
Interface name: ms-4/1/0
Session: 1, ALG: 1, Flags: 0x0040, IP Action: no, Offload: no
TCP          192.0.2.129:32843 ->       198.51.100.129:21    Forward  I          26
TCP          198.51.100.129:21    ->        192.0.2.0:32843 Forward  O          30
user@router> show services sessions application-protocol ike-esp-nat
Service Set: ss_ipv4, Session: 33554435, ALG: ike-esp-nat, Flags: 0x0800, IP Action: no,
Offload: no, Asymmetric: no
ESP 198.51.100.2:4689  ->        203.0.113.1:62108 Forward O 2199
ESP 192.0.2.2:62108  ->       198.51.100.2:4689 Forward I 0
Service Set: ss_ipv4, Session: 33554434, ALG: ike-esp-nat, Flags: 0x0800, IP Action: no,
Offload: no, Asymmetric: no
ESP 192.0.2.2:44179  ->       198.51.100.2:43809 Forward I 2199
ESP 198.51.100.2:43809  ->       203.0.113.1:44179 Forward O 0
Service Set: ss_ipv4, Session: 33554433, ALG: ike-esp-nat, Flags: 0x0000, IP Action: no,
Offload: no, Asymmetric: no
UDP 192.0.2.2:500  ->         198.51.100.2:500 Forward I 8
UDP 198.51.100.2:500  ->         203.0.113.1:57730 Forward O
user@router> show services sessions application-protocol pptp
Interface name: ms-2/0/0
Session: 3, ALG: pptp, Flags: 0x2800, IP Action: no, Offload: no, Asymmetric: no
GRE        203.0.113.138:0     ->    203.0.113.138:0     Forward  O          21
GRE        192.0.2.794:0     ->    203.0.113.138:0:65000 Forward  I           0
Session: 2, ALG: pptp, Flags: 0x2800, IP Action: no, Offload: no, Asymmetric: no
GRE        192.0.2.794:0     ->    203.0.113.138:0:49913 Forward  I          88
GRE        203.0.113.138:0:49913 ->    192.0.2.794:65001 Forward  O           0
Session: 1, ALG: pptp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP        192.0.2.794:1511  ->    203.0.113.138:0:1723  Forward  I          13
TCP        203.0.113.138:0:1723  ->    192.0.2.794:1511  Forward  O          12

user@router> show services sessions application-protocol rtsp
Interface name: ms-0/1/0
Session: 13, ALG: rtsp, Flags: 0x0800, IP Action: no, Offload: no
```

```
UDP           203.0.113.66:5004  ->        198.51.100.66:3989  Forward  O         152
UDP           198.51.100.66:3989  ->        192.0.2.161:5004  Forward  I           0
Session: 9, ALG: rtsp, Flags: 0x0800, IP Action: no, Offload: no
UDP           203.0.113.66:5004  ->        198.51.100.66:3986  Forward  O           3
UDP           198.51.100.66:3986  ->        192.0.2.161:5004  Forward  I           0

user@router> show services sessions application-protocol rsh
Interface name: ms-2/0/0
Session: 3, ALG: 2, Flags: 0x0840, IP Action: no, Offload: no
TCP       203.0.113.10:1023  ->     198.51.100.2:1020  Forward  O           4
TCP        198.51.100.2:1020  ->    203.0.113.10:1023  Forward  I           3
Session: 1, ALG: 2, Flags: 0x0040, IP Action: no, Offload: no
TCP        198.51.100.2:1021  ->    203.0.113.10:514   Forward  I        1331
TCP       203.0.113.10:514   ->     198.51.100.2:1021  Forward  O        2485
user@router> show services sessions application-protocol sip
Interface name: ms-2/0/0
Session: 4, ALG: sip, Flags: 0x0800, IP Action: no, Offload: no
UDP           198.51.100.130:6000  ->       192.0.2.129:12682 Forward  I             246
UDP           192.0.2.129:12682 ->       198.51.100.162:6000  Forward  O               0
Session: 1, ALG: sip, Flags: 0x0000, IP Action: no, Offload: no
UDP           198.51.100.130:5060  ->       192.0.2.130:5060  Forward  I              10
UDP           192.0.2.130:5060  ->       198.51.100.162:5060  Forward  O               9

user@router> show services sessions application-protocol sql
Interface name: ms-2/0/0
Session: 3934, ALG: sqlnet, Flags: 0x0800, IP Action: no, Offload: no
TCP        198.51.100.2:39754 ->    203.0.113.138:0:1408  Forward  I          26
TCP       203.0.113.138:0:1408  ->        192.0.2.1:39754 Forward  O          23

user@router> show services sessions application-protocol talk
Interface name: ms-0/2/0
Session: 4, ALG: 65, Flags: 0x0800, IP Action: no, Offload: no
TCP            203.0.113.162:36888 ->        192.0.2.2:33294 Forward  O               4
TCP            192.0.2.1:33294 ->        203.0.113.162:36888 Forward  I               3
Session: 7, ALG: 65, Flags: 0x0800, IP Action: no, Offload: no
UDP            203.0.113.162:1165  ->        192.0.2.2:518   Forward  O               1
UDP            192.0.2.2:518   ->        203.0.113.162:1165  Forward  I               1
Session: 8, ALG: 65, Flags: 0x0000, IP Action: no, Offload: no
UDP            192.0.2.2:1509  ->        203.0.113.162:518   Forward  I               3
UDP            203.0.113.162:518   ->        192.0.2.2:1509  Forward  O               3
Session: 6, ALG: 0, Flags: 0x0000, IP Action: no, Offload: no
UDP            192.0.2.1:123   ->        192.0.2.2:123   Forward  O               4
```

show services sessions count

```
user@host> show services sessions count
Interface   Service set                                        Sessions count
ms-1/1/0    ss                                                            2
```

show services sessions destination-port

```
user@router> show services sessions destination-port 21
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              25
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              24
```

show services sessions destination-prefix

```
user@router> show services sessions destination-prefix 10.1.1.2
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              25
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              24
```

show services sessions interface

```
user@router> show services sessions interface ms-1/1/0
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              30
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              29
```

show services sessions protocol

```
user@router> show services sessions protocol tcp
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
```

```
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              30
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              29
```

show services sessions service-set

```
user@router> show services sessions service-set sample
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              33
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              31
```

show services sessions source-port

```
user@router> show services sessions source-port 21
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              33
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              31
```

show services sessions source-prefix

```
user@router> show services sessions source-prefix 10.2.2.2
ms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              33
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              31
```

Release Information

Command introduced in Junos OS Release 10.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

The `nat`, `nat-port-overload-index`, and `source-nat-pool` are introduced in Junos OS Release 23.4R1.

show services sessions (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 26141
- Description | 26141
- Options | 26141
- Required Privilege Level | 26143
- Output Fields | 26143
- Sample Output | 26145
- Release Information | 26150

Syntax

```
show services sessions
<brief | extensive | terse>
<application-protocol protocol>
<count>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Description

Display the session information for each service set in each member interface of the AMS interface.

Options

none	Display standard information about all sessions.

brief | extensive | terse (Optional) Display the specified level of output.

application-protocol (Optional) Display information about one of the following application protocols:

- `ftp`—File Transfer Protocol
- `icmp`—Internet Control Message Protocol
- `pptp`—Point-to-Point Tunneling Protocol
- `rtsp`—Real-Time Streaming Protocol
- `sqlnet`—SQL *Net
- `tcp`—Transmission Control Protocol
- `traceroute`—Traceroute
- `tftp`—Trivial File Transfer Protocol
- `udp`—User Datagram Protocol

count (Optional) Display a count of the matching entries.

destination-port *destination-port* (Optional) Display information for a particular destination port. The range of values is from 0 through 65,535.

destination-prefix *destination-prefix* (Optional) Display information for a particular destination prefix.

interface *interface-name* (Optional) Display information about a particular interface. On M Series and T Series routers, *interface-name* can be `ms-`*fpc*`/`*pic*`/`*port* or `rsp`*number*. On J Series routers, *interface-name* is `ms-`*pim*`/0/`*port*.

limit *number* (Optional) Maximum number of entries to display.

protocol *protocol* (Optional) Display information about one of the following IP types:

- *number*—Numeric protocol value from 0 through 255
- `ah`—IPsec Authentication Header protocol
- `egp`—An exterior gateway protocol
- `esp`—IPsec Encapsulating Security Payload protocol
- `gre`—A generic routing encapsulation protocol

- `icmp`—Internet Control Message Protocol
- `icmp6`—Internet Control Message Protocol version 6
- `igmp`—Internet Group Management Protocol
- `ipip`—IP-over-IP encapsulation protocol
- `ospf`—Open Shortest Path First protocol
- `pim`—Protocol Independent Multicast protocol
- `rsvp`—Resource Reservation Protocol
- `sctp`—Stream Control Transmission Protocol
- `tcp`—Transmission Control Protocol
- `udp`—User Datagram Protocol

service-set *service-set* — (Optional) Display information for a particular service set.

source-port *source-port* — (Optional) Display information for a particular source port. The range of values is from 0 through 65,535.

source-prefix *source-prefix* — (Optional) Display information for a particular source prefix.

Required Privilege Level

view

Output Fields

Table 1510 on page 26143 lists the output fields for the `show services sessions` command. Output fields are listed in the approximate order in which they appear.

Table 1510: show services sessions Output Fields

Field Name	Field Description
`Interface`	Name of the member interface (`mams-`) and the aggregated multiservices interface (`ams`) to which it belongs.

Table 1510: show services sessions Output Fields *(Continued)*

Field Name	Field Description
`Session ID`	Session ID that uniquely identifies the session.
`ALG`	Name of the application.
`Flags`	Session flag for the ALG: • `0x1`—Found an existing session. • `0x2`—Reached session or flow limit. • `0x3`—No memory available for new sessions. • `0x4`—No free session ID available.
`IP Action`	Flag indicating whether IP action has been set for the session.
`Offload`	Flag indicating whether the session has been offloaded to the Packet Forwarding Engine.
`Asymmetric`	Flag indicating whether the session is unidirectional.
`Service set`	Name of a service set. Individual empty service sets are not displayed.
`Sessions Count`	Number of sessions.
`Flow or Flow Prot`	Protocol used for this session.
`Source`	Source prefix of the flow in the format *`source-prefix:port`*. For ICMP flows, port information is not displayed.
`Dest`	Destination prefix of the flow. For ICMP flows, port information is not displayed.

Table 1510: show services sessions Output Fields *(Continued)*

Field Name	Field Description
State	Status of the flow: • Drop—Drop all packets in the flow without response. • Forward—Forward the packet in the flow without looking at it. • Reject—Drop all packets in the flow with response. • Watch—Inspect packets in the flow. • Bypass—Bypass packets in the flow. • Unknown—Unknown flow status.
Packet Direction	Direction of the flow: ingress (I), egress (O), or unknown.
Frm count	Number of frames in the flow.

Sample Output

show services sessions brief

```
user@host> show services sessions brief
mams-1/0/0 (ams0)
Service Set: napt_set, Session: 16777217, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP         30.30.30.2:63    ->     40.40.40.2:63    Forward  I           85689
UDP         40.40.40.2:63    ->   30.30.30.160:6000  Forward  O               0
```

show services sessions interface mams-5/0/0 extensive

```
user@host> show services sessions interface mams-5/0/0 extensive
mams-1/0/0 (ams0)
Service Set: napt_set, Session: 16777235, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
```

```
NAT PLugin Data:
  NAT Action:   Translation Type - NAPT-44
    NAT source     30.30.30.62:63      ->    30.30.30.176:6003
UDP        30.30.30.62:63    ->    40.40.40.62:63    Forward  I            1805
  Byte count: 83030
  Flow role: Initiator, Timeout: 0
UDP        40.40.40.62:63    ->   30.30.30.176:6003  Forward  O               0
  Byte count: 0
  Flow role: Responder, Timeout: 0
Service Set: napt_set, Session: 16777234, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no

NAT PLugin Data:
  NAT Action:   Translation Type - NAPT-44
    NAT source     30.30.30.57:63      ->    30.30.30.163:6003
UDP        30.30.30.57:63    ->    40.40.40.57:63    Forward  I            1805
  Byte count: 83030
  Flow role: Initiator, Timeout: 0
UDP        40.40.40.57:63    ->   30.30.30.163:6003  Forward  O               0
  Byte count: 0
  Flow role: Responder, Timeout: 0

[...output truncated...]
mams-1/1/0 (ams0)
Service Set: napt_set, Session: 16777234, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no

NAT PLugin Data:
  NAT Action:   Translation Type - NAPT-44
    NAT source     30.30.30.63:63      ->    30.30.30.165:6004
UDP        30.30.30.63:63    ->    40.40.40.63:63    Forward  I            1805
  Byte count: 83030
  Flow role: Initiator, Timeout: 0
UDP        40.40.40.63:63    ->   30.30.30.165:6004  Forward  O               0
  Byte count: 0
  Flow role: Responder, Timeout: 0
Service Set: napt_set, Session: 16777233, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no

NAT PLugin Data:
  NAT Action:   Translation Type - NAPT-44
    NAT source     30.30.30.60:63      ->    30.30.30.164:6004
UDP        30.30.30.60:63    ->    40.40.40.60:63    Forward  I            1805
```

```
  Byte count: 83030
  Flow role: Initiator, Timeout: 0
UDP        40.40.40.60:63    ->   30.30.30.164:6004  Forward  O               0
  Byte count: 0
  Flow role: Responder, Timeout: 0
Service Set: napt_set, Session: 16777232, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no

[...output truncated...]
mams-5/0/0 (ams0)
Service Set: napt_set, Session: 16777225, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no

NAT PLugin Data:
  NAT Action:   Translation Type - NAPT-44
    NAT source     30.30.30.64:63      ->    30.30.30.168:6002
UDP        30.30.30.64:63    ->    40.40.40.64:63    Forward  I            1805
  Byte count: 83030
  Flow role: Initiator, Timeout: 0
UDP        40.40.40.64:63    ->   30.30.30.168:6002  Forward  O               0
  Byte count: 0
  Flow role: Responder, Timeout: 0
Service Set: napt_set, Session: 16777224, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no

NAT PLugin Data:
  NAT Action:   Translation Type - NAPT-44
    NAT source     30.30.30.56:63      ->    30.30.30.171:6001
UDP        30.30.30.56:63    ->    40.40.40.56:63    Forward  I            1805
  Byte count: 83030
  Flow role: Initiator, Timeout: 0
UDP        40.40.40.56:63    ->   30.30.30.171:6001  Forward  O               0
  Byte count: 0
  Flow role: Responder, Timeout: 0
Service Set: napt_set, Session: 16777223, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no

[...output truncated...]
mams-5/1/0 (ams0)
Service Set: napt_set, Session: 16777233, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no

NAT PLugin Data:
```

```
  NAT Action:   Translation Type - NAPT-44
    NAT source     30.30.30.61:63      ->    30.30.30.172:6004
UDP        30.30.30.61:63    ->    40.40.40.61:63    Forward  I            1805
  Byte count: 83030
  Flow role: Initiator, Timeout: 0
UDP        40.40.40.61:63    ->   30.30.30.172:6004  Forward  O               0
  Byte count: 0
  Flow role: Responder, Timeout: 0
Service Set: napt_set, Session: 16777232, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no

NAT PLugin Data:
  NAT Action:   Translation Type - NAPT-44
    NAT source     30.30.30.52:63      ->    30.30.30.175:6003
UDP        30.30.30.52:63    ->    40.40.40.52:63    Forward  I            1805
  Byte count: 83030
  Flow role: Initiator, Timeout: 0
UDP        40.40.40.52:63    ->   30.30.30.175:6003  Forward  O               0
  Byte count: 0
  Flow role: Responder, Timeout: 0
Service Set: napt_set, Session: 16777231, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no

[...output truncated...]
```

show services sessions terse

```
user@router> show services sessions terse
mams-1/0/0 (ams0)
Service Set: napt_set, Session: 16777235, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.62:63    ->    40.40.40.62:63    Forward  I            2541
UDP        40.40.40.62:63    ->   30.30.30.176:6003  Forward  O               0
Service Set: napt_set, Session: 16777234, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.57:63    ->    40.40.40.57:63    Forward  I            2541
UDP        40.40.40.57:63    ->   30.30.30.163:6003  Forward  O               0
Service Set: napt_set, Session: 16777233, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.50:63    ->    40.40.40.50:63    Forward  I            2541
UDP        40.40.40.50:63    ->   30.30.30.162:6003  Forward  O               0
```

```
Service Set: napt_set, Session: 16777232, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.48:63    ->    40.40.40.48:63    Forward  I            2541
UDP        40.40.40.48:63    ->   30.30.30.161:6003  Forward  O               0
[...output truncated...]
mams-1/1/0 (ams0)
Service Set: napt_set, Session: 16777234, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.63:63    ->    40.40.40.63:63    Forward  I            2543
UDP        40.40.40.63:63    ->   30.30.30.165:6004  Forward  O               0
Service Set: napt_set, Session: 16777233, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.60:63    ->    40.40.40.60:63    Forward  I            2543
UDP        40.40.40.60:63    ->   30.30.30.164:6004  Forward  O               0
Service Set: napt_set, Session: 16777232, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.59:63    ->    40.40.40.59:63    Forward  I            2543
UDP        40.40.40.59:63    ->   30.30.30.167:6003  Forward  O               0
Service Set: napt_set, Session: 16777231, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.58:63    ->    40.40.40.58:63    Forward  I            2543
UDP        40.40.40.58:63    ->   30.30.30.166:6003  Forward  O               0
[...output truncated...]
mams-5/0/0 (ams0)
Service Set: napt_set, Session: 16777225, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.64:63    ->    40.40.40.64:63    Forward  I            2543
UDP        40.40.40.64:63    ->   30.30.30.168:6002  Forward  O               0
Service Set: napt_set, Session: 16777224, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.56:63    ->    40.40.40.56:63    Forward  I            2543
UDP        40.40.40.56:63    ->   30.30.30.171:6001  Forward  O               0
Service Set: napt_set, Session: 16777223, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.55:63    ->    40.40.40.55:63    Forward  I            2543
UDP        40.40.40.55:63    ->   30.30.30.170:6001  Forward  O               0
Service Set: napt_set, Session: 16777222, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.51:63    ->    40.40.40.51:63    Forward  I            2543
UDP        40.40.40.51:63    ->   30.30.30.169:6001  Forward  O               0
[...output truncated...]
mams-5/1/0 (ams0)
Service Set: napt_set, Session: 16777233, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
```

```
Asymmetric: no
UDP        30.30.30.61:63    ->    40.40.40.61:63    Forward  I            2544
UDP        40.40.40.61:63    ->   30.30.30.172:6004  Forward  O               0
Service Set: napt_set, Session: 16777232, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.52:63    ->    40.40.40.52:63    Forward  I            2545
UDP        40.40.40.52:63    ->   30.30.30.175:6003  Forward  O               0
Service Set: napt_set, Session: 16777231, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.47:63    ->    40.40.40.47:63    Forward  I            2545
UDP        40.40.40.47:63    ->   30.30.30.174:6003  Forward  O               0
Service Set: napt_set, Session: 16777230, ALG: none, Flags: 0x2000, IP Action: no, Offload: no,
Asymmetric: no
UDP        30.30.30.46:63    ->    40.40.40.46:63    Forward  I            2545
UDP        40.40.40.46:63    ->   30.30.30.173:6003  Forward  O               0
[...output truncated...]
```

show services sessions count

```
user@host> show services sessions count
Interface   Service set                                          Sessions count
mams-1/0/0  napt_set                                                        19
mams-1/0/0  ss1                                                              0
mams-1/1/0  napt_set                                                        18
mams-1/1/0  ss1                                                              0
mams-5/0/0  napt_set                                                         9
mams-5/0/0  ss1                                                              0
mams-5/1/0  napt_set                                                        17
mams-5/1/0  ss1                                                              0
```

Release Information

Statement introduced in Junos OS Release 16.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

show services sessions analysis

IN THIS SECTION

- Syntax | 26151
- Description | 26151
- Options | 26151
- Required Privilege Level | 26151
- Output Fields | 26152
- Sample Output | 26154
- Release Information | 26156

Syntax

```
show services sessions analysis
<interface interface-name>
```

Description

Display session statistics.

Options

none	Display standard information about all session statistics.
interface *interface-name*	(Optional) Display information about the specified interface.

Required Privilege Level

view

Output Fields

Table 1511 on page 26152 lists the output fields for the `show services sessions analysis` command. Output fields are listed in the approximate order in which they appear.

Table 1511: show services sessions analysis Output Fields

Field Name	Field Description
Services PIC Name	FPC and PIC slots for the services PIC on which the sessions are running.
Session Analysis Statistics:	
Total Sessions Active	Total active sessions in the MS-PIC including TCP, UDP, ICMP and Softwires.
Total TCP Sessions Active	Total active TCP sessions in the MS-PIC.
Total UDP Sessions Active	Total active UDP session in the MS-PIC.
Total Other Sessions Active	Total other active sessions in the MS-PIC including ICMP and softwires.
Total Predicted Sessions Active	Predicted sessions are created only by the ALG traffic using the L3/L4 information available.
Created Sessions per Second	Session setup rate at the time of running the command.
Deleted Sessions per Second	Session deletion rate at the time of running the command.
Peak Total Sessions Active	Highest number of active sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Total TCP Sessions Active	Highest number of active TCP sessions since the last PIC restart or since the last time session stats are flushed.

Table 1511: show services sessions analysis Output Fields *(Continued)*

Field Name	Field Description
Peak Total UDP Sessions Active	Highest number of active UDP sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Total Other Sessions Active	Highest number of other active sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Created Sessions per Second	Maximum session setup rate observed since the last PIC restart or since the last time session statistics are flushed.
Peak Deleted Sessions per Second	Maximum session deletion rate observed since the last PIC restart or from the last time session statistics are flushed.
Packets received	Total number of packets received by the MS-PIC.
Packets transmitted	Total number of packets transmitted by the MS-PIC.
Slow path forward	Number of packets forwarded in the slow path (that is, after the successful rule match and session creation).
Slow path discard	Number of packets discarded before the session creation.
Session Rate Data: Number of Samples	Number of samples used to calculate the session rate since the last PIC restart or since the last time session statistics are flushed.
Session Rate Distribution(sec)	
Session Operation :Creation	Number of sampling intervals during which a number of sessions in the indicated range were created during the current sampling period.
Session Operation :Deletion	Number of sampling intervals during which a number of sessions in the indicated range were deleted during the current sampling period.

Table 1511: show services sessions analysis Output Fields *(Continued)*

Field Name	Field Description
Session Lifetime Distribution(sec):	Number of TCP, UDP, and HTTP sessions whose length was in the indicated range in seconds.

Sample Output

show services sessions analysis interface

```
user@host> show services sessions analysis interface ms-5/1/0
   Services PIC Name:    ms-5/1/0

Session Analysis Statistics:

  Total sessions Active                   :0
  Total TCP Sessions Active               :0
      Tcp sessions from gate              :0
      Tunneled TCP sessions               :0
      Regular TCP sessions                :0
      IPv4 active Session                 :0
      IPv6 active Session                 :0
  Total UDP sessions Active               :0
      UDP sessions from gate              :0
      Tunneled UDP sessions               :0
      Regular UDP sessions                :0
      IPv4 active Session                 :0
      IPv6 active Session                 :0
  Total Other sessions Active             :0
      IPv4 active Session                 :0
      IPv6 active Session                 :0
  Created sessions per Second             :0
  Deleted sessions per Second             :0
  Peak Total sessions Active              :0
  Peak Total TCP sessions Active          :0
  Peak Total UDP sessions Active          :0
  Peak Total Other sessions Active        :0
  Peak Created Sessions per Second        :0
  Peak Deleted Sessions per Second        :0
```

```
  Packets received                        :0
  Packets transmitted                     :0
  Slow path forward                       :0
  Slow path discard                       :0
Session Rate Data:
  Number of Samples: 3518

 Session Rate Distribution(sec)

 Session Operation :Creation

    400000+            :0
    350001  - 400000   :0
    300001  - 350000   :0
    250001  - 300000   :0
    200001  - 250000   :0
    150001  - 200000   :0
     50001  - 150000   :0
     40001  -  50000   :0
     30001  -  40000   :0
     20001  -  30000   :0
     10001  -  20000   :0
      1001  -  10000   :0
         1  -   1000   :0
                    0  :3518


 Session Operation :Deletion

    400000+            :0
    350001  - 400000   :0
    300001  - 350000   :0
    250001  - 300000   :0
    200001  - 250000   :0
    150001  - 200000   :0
     50001  - 150000   :0
     40001  -  50000   :0
     30001  -  40000   :0
     20001  -  30000   :0
     10001  -  20000   :0
      1001  -  10000   :0
         1  -   1000   :0
                    0  :3518
```

```
Session Lifetime Distribution(sec):

                 TCP                 UDP                 HTTP
  240+          :0                  0                   0
  120 - 240     :0                  0                   0
   60 - 120     :0                  0                   0
   30 - 60      :0                  0                   0
   15 - 30      :0                  0                   0
    5 - 15      :0                  0                   0
    1 - 5       :0                  0                   0
    0 - 1       :0                  0                   0
```

Release Information

Statement introduced in Junos OS Release 17.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

show services sessions analysis (USF)

IN THIS SECTION

- Syntax | 26157
- Description | 26157
- Options | 26157
- Required Privilege Level | 26157
- Output Fields | 26157
- Sample Output | 26159
- Release Information | 26161

Syntax

```
show services sessions analysis
<interface interface-name>
```

Description

Display session statistics.

Options

none	Display standard information about all session statistics.
interface *interface-name*	(Optional) Display information about the specified services interface.

Required Privilege Level

view

Output Fields

Table 1512 on page 26157 lists the output fields for the `show services sessions analysis` command. Output fields are listed in the approximate order in which they appear.

Table 1512: show services sessions analysis Output Fields

Field Name	Field Description
Services PIC Name	FPC and PIC slots for the services PIC on which the sessions are running.
Session Analysis Statistics:	
Total Sessions Active	Total active sessions in the services PIC, including TCP, UDP, ICMP and Softwires.
Total TCP Sessions Active	Total active TCP sessions in the services PIC.

Table 1512: show services sessions analysis Output Fields *(Continued)*

Field Name	Field Description
Total UDP Sessions Active	Total active UDP session in the services PIC.
Total Other Sessions Active	Total other active sessions in the services PIC, including ICMP and softwires.
Total Predicted Sessions Active	Predicted sessions are created only by the ALG traffic using the L3/L4 information available.
Created Sessions per Second	Session setup rate at the time of running the command.
Deleted Sessions per Second	Session deletion rate at the time of running the command.
Peak Total Sessions Active	Highest number of active sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Total TCP Sessions Active	Highest number of active TCP sessions since the last PIC restart or since the last time session stats are flushed.
Peak Total UDP Sessions Active	Highest number of active UDP sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Total Other Sessions Active	Highest number of other active sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Created Sessions per Second	Maximum session setup rate observed since the last PIC restart or since the last time session statistics are flushed.
Peak Deleted Sessions per Second	Maximum session deletion rate observed since the last PIC restart or from the last time session statistics are flushed.
Packets received	Total number of packets received by the services PIC.

Table 1512: show services sessions analysis Output Fields *(Continued)*

Field Name	Field Description
Packets transmitted	Total number of packets transmitted by the services PIC.
Slow path forward	Number of packets forwarded in the slow path (that is, after the successful rule match and session creation).
Slow path discard	Number of packets discarded before the session creation.
Session Rate Data: Number of Samples	Number of samples used to calculate the session rate since the last PIC restart or since the last time session statistics are flushed.
Session Rate Distribution(sec)	
Session Operation :Creation	Number of sampling intervals during which a number of sessions in the indicated range were created during the current sampling period.
Session Operation:Deletion	Number of sampling intervals during which a number of sessions in the indicated range were deleted during the current sampling period.
Session Lifetime Distribution(sec):	Number of TCP, UDP, and HTTP sessions whose length was in the indicated range in seconds.

Sample Output

show services sessions analysis interface

```
user@host> show services sessions analysis interface vms-5/1/0
   Services PIC Name:    vms-5/1/0

Session Analysis Statistics:

  Total sessions Active                   :0
  Total TCP Sessions Active               :0
      Tcp sessions from gate              :0
```

```
      Tunneled TCP sessions               :0
      Regular TCP sessions                :0
      IPv4 active Session                 :0
      IPv6 active Session                 :0
  Total UDP sessions Active               :0
      UDP sessions from gate              :0
      Tunneled UDP sessions               :0
      Regular UDP sessions                :0
      IPv4 active Session                 :0
      IPv6 active Session                 :0
  Total Other sessions Active             :0
      IPv4 active Session                 :0
      IPv6 active Session                 :0
  Created sessions per Second             :0
  Deleted sessions per Second             :0
  Peak Total sessions Active              :0
  Peak Total TCP sessions Active          :0
  Peak Total UDP sessions Active          :0
  Peak Total Other sessions Active        :0
  Peak Created Sessions per Second        :0
  Peak Deleted Sessions per Second        :0
  Packets received                        :0
  Packets transmitted                     :0
  Slow path forward                       :0
  Slow path discard                       :0
Session Rate Data:
  Number of Samples: 3518

 Session Rate Distribution(sec)

 Session Operation :Creation

    400000+            :0
    350001  - 400000   :0
    300001  - 350000   :0
    250001  - 300000   :0
    200001  - 250000   :0
    150001  - 200000   :0
     50001  - 150000   :0
     40001  -  50000   :0
     30001  -  40000   :0
     20001  -  30000   :0
     10001  -  20000   :0
```

```
      1001  -  10000   :0
         1  -   1000   :0
                   0   :3518


 Session Operation :Deletion

    400000+            :0
    350001  - 400000   :0
    300001  - 350000   :0
    250001  - 300000   :0
    200001  - 250000   :0
    150001  - 200000   :0
     50001  - 150000   :0
     40001  -  50000   :0
     30001  -  40000   :0
     20001  -  30000   :0
     10001  -  20000   :0
      1001  -  10000   :0
         1  -   1000   :0
                   0   :3518


Session Lifetime Distribution(sec):

                TCP                UDP               HTTP
  240+         :0                  0                 0
  120 - 240    :0                  0                 0
   60 - 120    :0                  0                 0
   30 - 60     :0                  0                 0
   15 - 30     :0                  0                 0
    5 - 15     :0                  0                 0
    1 - 5      :0                  0                 0
    0 - 1      :0                  0                 0
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services sessions count

IN THIS SECTION

- Syntax | 26162
- Description | 26162
- Required Privilege Level | 26162
- Output Fields | 26162
- Sample Output | 26162
- Release Information | 26163

Syntax

```
show services sessions count
```

Description

Display the count of matching entries.

Required Privilege Level

view

Output Fields

Sample Output

show services sessions count

```
user@host> show services sessions count
Interface   Service set     Valid    Invalid      Pending  Other state
vms-0/2/0   ss1_interface_style1     1         0            0            0
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services sessions (Next Gen Services)

IN THIS SECTION

- Syntax | 26163
- Description | 26164
- Options | 26164
- Required Privilege Level | 26166
- Output Fields | 26167
- Sample Output | 26168
- Release Information | 26176

Syntax

```
show services sessions
<brief | extensive | terse>
<application-protocol protocol>
<count>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
<utilization>
```

Description

Display session information.

> **NOTE**: On MX Series routers (with interchassis redundancy configured), the idle timeout for every flow is displayed in the `show services session extensive` and `show services flows extensive` commands.

Options

none

Display standard information about all sessions.

brief | extensive | terse

(Optional) Display the specified level of output.

application-protocol *protocol*

(Optional) Display information about one of the following application protocols:

- `bootp`—Bootstrap protocols
- `dce-rpc`—Distributed Computing Environment-Remote Procedure Call protocols
- `dce-rpc-portmap`—Distributed Computing Environment-Remote Procedure Call protocols portmap service
- `dns`—Domain Name System protocol
- `exec`—Remote Execution Protocol
- `ftp`—File Transfer Protocol
- `h323`—H.323
- `icmp`—ICMP
- `icmpv6`—ICMPv6
- `iiop`—Internet Inter-ORB Protocol
- `ike-esp-nat`—IKE ALG
- `ip`—IP
- `login`—LOGIN
- `netbios`—NETBIOS

- `netshow`—NETSHOW
- `pptp`—Point-to-Point Tunneling Protocol
- `realaudio`—RealAudio
- `rpc`—Remote Procedure Call protocol
- `rpc-portmap`—Remote Procedure Call protocol portmap service
- `rtsp`—Real-Time Streaming Protocol
- `rsh`—Remote Shell
- `sip`—Session Initiation Protocol
- `shell`—Shell
- `snmp`—SNMP
- `sql`—SQLNet
- `talk`—Talk Program
- `tftp`—Trivial File Transfer Protocol
- `traceroute`—Traceroute
- `winframe`—WinFrame

> **NOTE**: You can use the `none` option with the `show services sessions count application-protocol` command to display information about sessions other than ALG sessions.

count — (Optional) Display a count of the matching entries.

destination-port *destination-port* — (Optional) Display information for the specified destination port. The range of values is from 0 to 65,535.

destination-prefix *destination-prefix* — (Optional) Display information for the specified destination prefix.

interface *interface-name* — (Optional) Display information about the specified services interface.

limit *number* — (Optional) Maximum number of entries to display.

protocol *protocol* — (Optional) Display information about one of the following IP types:

- *number*—Numeric protocol value from 0 to 255
- ah—IPsec Authentication Header protocol
- egp—An exterior gateway protocol
- esp—IPsec Encapsulating Security Payload protocol
- gre—A generic routing encapsulation protocol
- icmp—Internet Control Message Protocol
- icmp6—Internet Control Message Protocol version 6
- igmp—Internet Group Management Protocol
- ipip—IP-within-IP Encapsulation Protocol
- ospf—Open Shortest Path First protocol
- pim—Protocol Independent Multicast protocol
- rsvp—Resource Reservation Protocol
- sctp—Stream Control Transmission Protocol
- tcp—Transmission Control Protocol
- udp—User Datagram Protocol

service-set ***service-set***	(Optional) Display information for the specified service set.
source-port ***source-port***	(Optional) Display information for the specified source port. The range of values is from 0 to 65,535.
source-prefix ***source-prefix***	(Optional) Display information for the specified source prefix.
utilization	(Optional) Display statistical details about session utilization.

Required Privilege Level

view

Output Fields

Table 1513 on page 26167 lists the output fields for the `show services sessions` command. Output fields are listed in the approximate order in which they appear.

Table 1513: show services sessions Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the services interface.	`application-protocol`
`Session`	Session ID that uniquely identifies the session.	All levels
`ALG`	Name of the application.	`terse`
Flags	Session flag for the ALG: • `0x1`—Found an existing session. • `0x2`—Reached session or flow limit. • `0x3`—No memory available for new sessions. • `0x4`—No free session ID available. • `0x0000`—No session ID found.	All levels
IP Action	Flag indicating whether IP action has been set for the session.	All levels
`Offload`	Flag indicating whether the session has been offloaded to the Packet Forwarding Engine.	All levels
`Asymmetric`	Flag indicating whether the session is uni-directional.	`terse` `application-protocol`
`Service set`	Name of a service set. Individual empty service sets are not displayed.	`count`

Table 1513: show services sessions Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Sessions Count	Number of sessions.	count

Sample Output

show services sessions

```
user@host> show services sessions
Session ID: 536870913, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 26, Valid
Logical system: root-logical-system
  In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.16391, Pkts: 1, Bytes: 110,
  Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip, Conn Tag: 0x0, If: vms-2/0/0.0,
Pkts: 0, Bytes: 0,

Session ID: 536870914, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 26, Valid
Logical system: root-logical-system
    Softwire        2002:2010::1401:4       -> 2002:2010::1401:1
  In: 30.1.0.101/1024 --> 30.2.0.101/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.16391, Pkts: 1,
Bytes: 70,
  Out: 30.2.0.101/1024 --> 50.0.12.1/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.0, Pkts: 0, Bytes: 0,
Total sessions: 2
```

show services sessions brief

The output for the `show services flows brief` command is identical to that for the `show services sessions` command. For sample output, see "show services sessions" on page 26168.

show services sessions extensive

```
user@host> show services sessions extensive
Session ID: 536870917, Service-set: vms-sset10, Status: Normal
```

```
Flags: 0x40/0x0/0x4000/0x2000103
Policy name: default-service-set-policy/32779
Source NAT pool: Null, Destination NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 30, Current timeout: 28
Session State: Valid
Logical system: root-logical-system
Start time: 1878, Duration: 2
   In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip,
  Conn Tag: 0x0, Interface: vms-2/0/0.16391,
    Session token: 0xfcc, Flag: 0x400023
    Route: 0x0, Gateway: 2002:2010::1401:4, Tunnel ID: 0, Tunnel type: None
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 110
   Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip,
  Conn Tag: 0x0, Interface: vms-2/0/0.0,
    Session token: 0x4fcc, Flag: 0x400022
    Route: 0x0, Gateway: 2002:2010::1401:1, Tunnel ID: 0, Tunnel type: None
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
Total sessions: 1
```

show services sessions terse

```
user@router> show services sessions terse
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              33
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              31
```

show services sessions analysis

```
user@router>show services sessions analysis
vms-1/0/0
  Interface:   vms-1/0/0
```

```
Session Analysis Statistics:

  Total sessions Active                   :0
  Total TCP Sessions Active               :0
      Tcp sessions from gate              :0
      Tunneled TCP sessions               :0
      Regular TCP sessions                :0
      IPv4 active Session                 :0
      IPv6 active Session                 :0
  Total UDP sessions Active               :0
      UDP sessions from gate              :0
      Tunneled UDP sessions               :0
      Regular UDP sessions                :0
      IPv4 active Session                 :0
      IPv6 active Session                 :0
  Total Other sessions Active             :0
      IPv4 active Session                 :0
      IPv6 active Session                 :0
  Created sessions per Second             :0
  Deleted sessions per Second             :0
  Peak Total sessions Active              :0
  Peak Total TCP sessions Active          :0
  Peak Total UDP sessions Active          :0
  Peak Total Other sessions Active        :0
  Peak Created Sessions per Second        :0
  Peak Deleted Sessions per Second        :0
  Packets received                        :0
  Packets transmitted                     :0
  Slow path forward                       :0
  Slow path discard                       :0

Session Rate Data:
  Number of Samples: 638051

 Session Rate Distribution(sec)

 Session Operation :Creation

    400000+            :0
    350001  - 400000   :0
    300001  - 350000   :0
    250001  - 300000   :0
    200001  - 250000   :0
```

```
    150001  - 200000   :0
     50001  - 150000   :0
     40001  -  50000   :0
     30001  -  40000   :0
     20001  -  30000   :0
     10001  -  20000   :0
      1001  -  10000   :0
         1  -   1000   :0
                    0   :638051


 Session Operation :Deletion

    400000+            :0
    350001  - 400000   :0
    300001  - 350000   :0
    250001  - 300000   :0
    200001  - 250000   :0
    150001  - 200000   :0
     50001  - 150000   :0
     40001  -  50000   :0
     30001  -  40000   :0
     20001  -  30000   :0
     10001  -  20000   :0
      1001  -  10000   :0
         1  -   1000   :0
                    0   :638051


Session Lifetime Distribution(sec):

                TCP                UDP               HTTP
  240+         :0                 0                 0
  120 - 240    :0                 0                 0
   60 - 120    :0                 0                 0
   30 - 60     :0                 0                 0
   15 - 30     :0                 0                 0
    5 - 15     :0                 0                 0
    1 - 5      :0                 0                 0
    0 - 1      :0                 0                 0
```

show services sessions application-protocol

This command has the same output for the rpc, dce-rpc, rpc-portmap and dce-rpc-portmap ALGs.

```
user@router> show services sessions application-protocol dce-rpc
Interface name: vms-1/1/0
Session: 8, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:1019  ->192.168.203.194:2049  Forward  I               4
UDP    192.168.203.194:2049  ->192.168.203.198:1019  Forward  O               4
Session: 7, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:954   ->192.168.203.194:613   Forward  I               1
UDP    192.168.203.194:613   ->192.168.203.198:954   Forward  O               1
Session: 6, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:53836 ->192.168.203.194:613   Forward  I               1
UDP    192.168.203.194:613   ->192.168.203.198:53836 Forward  O               1
Session: 5, ALG: portmapper, Flags: 0x1000, IP Action: no, Offload: no
UDP    192.168.203.198:59813 ->192.168.203.194:111   Forward  I               1
UDP    192.168.203.194:111   ->192.168.203.198:59813 Forward  O               1
Session: 4, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:36595 ->192.168.203.194:2049  Forward  I               1
UDP    192.168.203.194:2049  ->192.168.203.198:36595 Forward  O               1
Session: 3, ALG: portmapper, Flags: 0x1000, IP Action: no, Offload: no
UDP    192.168.203.198:56050 ->192.168.203.194:111   Forward  I               1
UDP    192.168.203.194:111   ->192.168.203.198:56050 Forward  O               1

user@router> show services sessions application-protocol dns
Interface name: vms-2/0/0
Session: 293, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        198.51.100.2:43677 ->    203.0.113.10:53    Forward  I           1
UDP       203.0.113.10:53    ->        192.0.2.1:43677 Forward  O           1
Session: 53, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        198.51.100.2:37494 ->    203.0.113.10:53    Forward  I           1
UDP       203.0.113.10:53    ->        192.0.2.1:37494 Forward  O           1
Session: 66, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        198.51.100.2:48161 ->    203.0.113.10:53    Forward  I           1
UDP       203.0.113.10:53    ->        192.0.2.1:48161 Forward  O           1
Session: 17, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        198.51.100.2:38908 ->    203.0.113.10:53    Forward  I           1
UDP       203.0.113.10:53    ->        192.0.2.1:38908 Forward  O           1
Session: 42, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP        198.51.100.2:58189 ->    203.0.113.10:53    Forward  I           1
UDP       203.0.113.10:53    ->        192.0.2.1:58189 Forward  O           1
```

```
user@router> show services sessions application-protocol ftp
Interface name: vms-4/1/0
Session: 1, ALG: 1, Flags: 0x0040, IP Action: no, Offload: no
TCP          192.0.2.129:32843 ->       198.51.100.129:21    Forward  I          26
TCP          198.51.100.129:21    ->        192.0.2.0:32843 Forward  O          30
user@router> show services sessions application-protocol ike-esp-nat
Service Set: ss_ipv4, Session: 33554435, ALG: ike-esp-nat, Flags: 0x0800, IP Action: no,
Offload: no, Asymmetric: no
ESP 198.51.100.2:4689  ->        203.0.113.1:62108 Forward O 2199
ESP 192.0.2.2:62108  ->       198.51.100.2:4689 Forward I 0
Service Set: ss_ipv4, Session: 33554434, ALG: ike-esp-nat, Flags: 0x0800, IP Action: no,
Offload: no, Asymmetric: no
ESP 192.0.2.2:44179  ->       198.51.100.2:43809 Forward I 2199
ESP 198.51.100.2:43809  ->       203.0.113.1:44179 Forward O 0
Service Set: ss_ipv4, Session: 33554433, ALG: ike-esp-nat, Flags: 0x0000, IP Action: no,
Offload: no, Asymmetric: no
UDP 192.0.2.2:500  ->         198.51.100.2:500 Forward I 8
UDP 198.51.100.2:500  ->         203.0.113.1:57730 Forward O
user@router> show services sessions application-protocol pptp
Interface name: vms-2/0/0
Session: 3, ALG: pptp, Flags: 0x2800, IP Action: no, Offload: no, Asymmetric: no
GRE        203.0.113.138:0     ->    203.0.113.138:0     Forward  O              21
GRE        192.0.2.794:0     ->    203.0.113.138:0:65000 Forward  I               0
Session: 2, ALG: pptp, Flags: 0x2800, IP Action: no, Offload: no, Asymmetric: no
GRE        192.0.2.794:0     ->    203.0.113.138:0:49913 Forward  I              88
GRE        203.0.113.138:0:49913 ->    192.0.2.794:65001 Forward  O               0
Session: 1, ALG: pptp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP        192.0.2.794:1511  ->    203.0.113.138:0:1723  Forward  I              13
TCP        203.0.113.138:0:1723  ->    192.0.2.794:1511  Forward  O              12

user@router> show services sessions application-protocol rtsp
Interface name: vms-0/1/0
Session: 13, ALG: rtsp, Flags: 0x0800, IP Action: no, Offload: no
UDP           203.0.113.66:5004  ->        198.51.100.66:3989  Forward  O         152
UDP           198.51.100.66:3989  ->        192.0.2.161:5004  Forward  I           0
Session: 9, ALG: rtsp, Flags: 0x0800, IP Action: no, Offload: no
UDP           203.0.113.66:5004  ->        198.51.100.66:3986  Forward  O           3
UDP           198.51.100.66:3986  ->        192.0.2.161:5004  Forward  I           0

user@router> show services sessions application-protocol rsh
Interface name: vms-2/0/0
Session: 3, ALG: 2, Flags: 0x0840, IP Action: no, Offload: no
TCP       203.0.113.10:1023  ->     198.51.100.2:1020  Forward  O           4
```

```
TCP        198.51.100.2:1020  ->    203.0.113.10:1023  Forward  I           3
Session: 1, ALG: 2, Flags: 0x0040, IP Action: no, Offload: no
TCP        198.51.100.2:1021  ->    203.0.113.10:514   Forward  I        1331
TCP       203.0.113.10:514   ->     198.51.100.2:1021  Forward  O        2485
user@router> show services sessions application-protocol sip
Interface name: vms-2/0/0
Session: 4, ALG: sip, Flags: 0x0800, IP Action: no, Offload: no
UDP           198.51.100.130:6000  ->       192.0.2.129:12682 Forward  I             246
UDP           192.0.2.129:12682 ->       198.51.100.162:6000  Forward  O               0
Session: 1, ALG: sip, Flags: 0x0000, IP Action: no, Offload: no
UDP           198.51.100.130:5060  ->       192.0.2.130:5060  Forward  I              10
UDP           192.0.2.130:5060  ->       198.51.100.162:5060  Forward  O               9

user@router> show services sessions application-protocol sql
Interface name: vms-2/0/0
Session: 3934, ALG: sqlnet, Flags: 0x0800, IP Action: no, Offload: no
TCP        198.51.100.2:39754 ->    203.0.113.138:0:1408  Forward  I          26
TCP       203.0.113.138:0:1408  ->        192.0.2.1:39754 Forward  O          23

user@router> show services sessions application-protocol talk
Interface name: vms-0/2/0
Session: 4, ALG: 65, Flags: 0x0800, IP Action: no, Offload: no
TCP            203.0.113.162:36888 ->        192.0.2.2:33294 Forward  O               4
TCP            192.0.2.1:33294 ->        203.0.113.162:36888 Forward  I               3
Session: 7, ALG: 65, Flags: 0x0800, IP Action: no, Offload: no
UDP            203.0.113.162:1165  ->        192.0.2.2:518   Forward  O               1
UDP            192.0.2.2:518   ->        203.0.113.162:1165  Forward  I               1
Session: 8, ALG: 65, Flags: 0x0000, IP Action: no, Offload: no
UDP            192.0.2.2:1509  ->        203.0.113.162:518   Forward  I               3
UDP            203.0.113.162:518   ->        192.0.2.2:1509  Forward  O               3
Session: 6, ALG: 0, Flags: 0x0000, IP Action: no, Offload: no
UDP            192.0.2.1:123   ->        192.0.2.2:123   Forward  O               4
```

show services sessions count

```
user@host> show services sessions count
Interface   Service set                   Valid      Invalid      Pending  Other state
vms-0/2/0   ss1_interface_style1                   1            0            0            0
```

show services sessions destination-port

```
user@router> show services sessions destination-port 21
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              25
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              24
```

show services sessions destination-prefix

```
user@router> show services sessions destination-prefix 10.1.1.2
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              25
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              24
```

show services sessions interface

```
user@router> show services sessions interface vms-1/1/0
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              30
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              29
```

show services sessions protocol

```
user@router> show services sessions protocol tcp
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              30
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              29
```

show services sessions service-set

```
user@router> show services sessions service-set ss1_interface_style1
Session ID: 3, Service-set: ss1_interface_style1, Policy name: R11/7, Timeout: 30, Valid
```

```
  In: 20.1.1.2/48102 --> 30.1.1.2/22;tcp, Conn Tag: 0x0, If: vms-0/2/0.16387, Pkts: 70, Bytes:
6257,
  Out: 30.1.1.2/22 --> 44.0.0.3/29071;tcp, Conn Tag: 0x0, If: vms-0/2/0.0, Pkts: 59, Bytes:
8193,
Total sessions: 1
```

show services sessions source-port

```
user@router> show services sessions source-port 21
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              33
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              31
```

show services sessions source-prefix

```
user@router> show services sessions source-prefix 10.2.2.2
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP           10.2.2.2:52138 ->       10.1.1.2:21    Forward  I              33
TCP           10.1.1.2:21    ->       10.2.2.2:52138 Forward  O              31
```

Release Information

Command introduced in Junos OS Release 19.3R2 on MX Series for Next Gen Services for CGNAT 6rd softwires running inline on the MPC card and specifying the si-1/0/0 interface naming convention.Support added in Junos OS Release 20.2R1 for Next Gen Services CGNAT DS-Lite softwires on the MX-SPC3 security services card .

show services sessions service-set

IN THIS SECTION

Syntax | 26177

- Description | 26177
- Required Privilege Level | 26177
- show services sessions service-set | 26177
- Release Information | 26178

Syntax

```
show services sessions service-set
```

Description

Display the open and close sessions for a service-set.

Required Privilege Level

show services sessions service-set

command-name

```
user@host> show services sessions service-set service-set-name

Session ID: 268436944, Policy name: self-traffic-policy/1, Timeout: 554, Valid
Logical system: root-logical-system
  In: 5.5.5.1/12253 --> 70.0.0.2/514;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2, Bytes: 84,
  Out: 70.0.0.2/514 --> 5.5.5.1/12253;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2, Bytes: 84,

Session ID: 268436945, Policy name: self-traffic-policy/1, Timeout: 554, Valid
Logical system: root-logical-system
  In: 5.5.5.1/12254 --> 70.0.0.2/514;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2, Bytes: 84,
  Out: 70.0.0.2/514 --> 5.5.5.1/12254;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2, Bytes: 84,

Session ID: 268436946, Policy name: self-traffic-policy/1, Timeout: 596, Valid
Logical system: root-logical-system
  In: 5.5.5.1/12255 --> 70.0.0.2/514;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2, Bytes: 84,
```

```
  Out: 70.0.0.2/514 --> 5.5.5.1/12255;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 1, Bytes: 44,

Session ID: 268436947, Policy name: self-traffic-policy/1, Timeout: 554, Valid
Logical system: root-logical-system
  In: 5.5.5.1/12256 --> 70.0.0.2/514;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2, Bytes: 84,
  Out: 70.0.0.2/514 --> 5.5.5.1/12256;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2, Bytes: 84,

Session ID: 268436948, Policy name: self-traffic-policy/1, Timeout: 596, Valid
Logical system: root-logical-system
  In: 5.5.5.1/12257 --> 70.0.0.2/514;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2, Bytes: 84,
  Out: 70.0.0.2/514 --> 5.5.5.1/12257;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 1, Bytes: 44,
Total sessions: 5
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services sessions softwire

IN THIS SECTION

- Syntax | 26179
- Description | 26179
- Options | 26179
- Required Privilege Level | 26179
- Output Fields | 26179
- Sample Output | 26179
- show services sessions softwire count | 26180
- show services sessions softwire ds-lite | 26180
- show services sessions softwire ds-lite count | 26181
- show services sessions softwire ds-lite aftr | 26181
- show services sessions softwire ds-lite b4 | 26182
- show services sessions softwire ds-lite b4 <ip-address> aftr <ip-address> | 26182

Show services sessions softwire flow-details | 26183
Release Information | 26183

Syntax

```
show services sessions softwire
interfaces interface-name
```

Description

Display session information for softwires.

Options

count Display statistics and information on the number of softwires.

ds-lite Display information about DS-Lite softwires.

Required Privilege Level

view

Output Fields

Sample Output

show services sessions softwire

```
user@host> show services sessions softwire
Session ID: 536870913, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 26, Valid
Logical system: root-logical-system
  In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.16391, Pkts: 1, Bytes: 110,
```

```
  Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip, Conn Tag: 0x0, If: vms-2/0/0.0,
Pkts: 0, Bytes: 0,

Session ID: 536870914, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 26, Valid
Logical system: root-logical-system
    Softwire        2002:2010::1401:4       -> 2002:2010::1401:1
  In: 30.1.0.101/1024 --> 30.2.0.101/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.16391, Pkts: 1,
Bytes: 70,
  Out: 30.2.0.101/1024 --> 50.0.12.1/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.0, Pkts: 0, Bytes: 0,
Total sessions: 2
```

show services sessions softwire count

show services sessions softwire count

```
user@host> show services sessions softwire count
Interface   Service set                        Valid      Invalid      Pending  Other state
vms-2/0/0   vms-sset10                             1            0            0            0
vms-2/0/0   vms-sset11
```

show services sessions softwire ds-lite

show services sessions softwire ds-lite

```
user@host> show services sessions softwire ds-lite
Session ID: 536870913, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 26, Valid
Logical system: root-logical-system
  In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.16391, Pkts: 1, Bytes: 110,
  Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip, Conn Tag: 0x0, If: vms-2/0/0.0,
Pkts: 0, Bytes: 0,

Session ID: 536870914, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 26, Valid
Logical system: root-logical-system
    Softwire        2002:2010::1401:4       -> 2002:2010::1401:1
```

```
  In: 30.1.0.101/1024 --> 30.2.0.101/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.16391, Pkts: 1,
Bytes: 70,
  Out: 30.2.0.101/1024 --> 50.0.12.1/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.0, Pkts: 0, Bytes: 0,
Total sessions: 2
```

show services sessions softwire ds-lite count

show services sessions softwire ds-lite count

```
user@host> show services sessions softwire ds-lite count
Interface   Service set                        Valid      Invalid      Pending  Other state
vms-2/0/0   vms-sset10                             1            0            0            0
vms-2/0/0   vms-sset11
```

show services sessions softwire ds-lite aftr

show services sessions softwire ds-lite aftr

```
user@host> show services sessions softwire ds-lite aftr
Session ID: 536870913, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 6, Valid
Logical system: root-logical-system
  In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.16391, Pkts: 1, Bytes: 110,
  Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip, Conn Tag: 0x0, If: vms-2/0/0.0,
Pkts: 0, Bytes: 0,

Session ID: 536870914, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 6, Valid
Logical system: root-logical-system
    Softwire        2002:2010::1401:4       -> 2002:2010::1401:1
  In: 30.1.0.101/1024 --> 30.2.0.101/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.16391, Pkts: 1,
Bytes: 70,
  Out: 30.2.0.101/1024 --> 50.0.12.1/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.0, Pkts: 0, Bytes: 0,
Total sessions: 2
```

show services sessions softwire ds-lite b4

show services sessions softwire ds-lite b4

```
user@host> show services sessions softwire ds-lite b4
Session ID: 536870913, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 6, Valid
Logical system: root-logical-system
  In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.16391, Pkts: 1, Bytes: 110,
  Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip, Conn Tag: 0x0, If: vms-2/0/0.0,
Pkts: 0, Bytes: 0,

Session ID: 536870914, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 6, Valid
Logical system: root-logical-system
    Softwire        2002:2010::1401:4       -> 2002:2010::1401:1
  In: 30.1.0.101/1024 --> 30.2.0.101/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.16391, Pkts: 1,
Bytes: 70,
  Out: 30.2.0.101/1024 --> 50.0.12.1/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.0, Pkts: 0, Bytes: 0,
Total sessions: 2
```

show services sessions softwire ds-lite b4 <ip-address> aftr <ip-address>

show services sessions softwire ds-lite b4 <ip address> aftr <ip-address>

```
user@host> show services sessions softwire ds-lite b4 ip address aftr ip-address
Session ID: 536870913, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 6, Valid
Logical system: root-logical-system
  In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.16391, Pkts: 1, Bytes: 110,
  Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip, Conn Tag: 0x0, If: vms-2/0/0.0,
Pkts: 0, Bytes: 0,

Session ID: 536870914, Service-set: vms-sset10, Policy name: default-service-set-policy/32779,
Timeout: 6, Valid
Logical system: root-logical-system
    Softwire        2002:2010::1401:4       -> 2002:2010::1401:1
  In: 30.1.0.101/1024 --> 30.2.0.101/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.16391, Pkts: 1,
```

```
Bytes: 70,
  Out: 30.2.0.101/1024 --> 50.0.12.1/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.0, Pkts: 0, Bytes: 0,
Total sessions: 2
```

Show services sessions softwire flow-details

Show services sessions softwire flow-details

```
user@host> show services sessions softwire flow-details
Interface: vms-2/0/0, Service set: vms-sset10
Softwire                                                    Direction     Flow count
2002:2010::1401:4->2002:2010::1401:1            In                  1
```

Release Information

Command introduced in Junos OS Release 20.2R1.

show services sessions tcp-log

IN THIS SECTION

- Syntax | 26183
- Description | 26184
- Required Privilege Level | 26184
- Output Fields | 26184
- Sample Output | 26184
- Release Information | 26184

Syntax

```
show services sessions   tcp-log interface interface-name
```

Description

Display information about the tcp log session details.

Required Privilege Level

view

Output Fields

Table 1514 on page 26184 lists the output fields for the `show services sessions tcp-log interface` *`interface-name`* command. Output fields are listed in the approximate order in which they appear.

Table 1514: show services sessions tcp-log interface Output Fields

Field Name	Field Description
Service-set	Name of a service set.
TCP	TCP log session details.

Sample Output

command-name

```
user@router> show services sessions tcp-log interface interface-name

Service Set: junos-tcplog, Session: 33554434, ALG: None, Flags: 0x200000, IP Action: no,
Offload: no, Asymmetric: no
TCP          18.1.1.1        ->       22.1.1.2        Forward  I               6
TCP          22.1.1.2        ->       18.1.1.1        Forward  O              6
```

Release Information

Command introduced in Junos OS Release 19.1.

show services sessions utilization

IN THIS SECTION

- Syntax | 26185
- Description | 26185
- Options | 26185
- Required Privilege Level | 26185
- Output Fields | 26185
- Sample Output | 26186
- Release Information | 26186

Syntax

```
show services sessions utilization
<interface interface-name>
```

Description

Display session utilization statistics.

Options

interface *interface-name* Display session utilization statistics specific to the interface.

Required Privilege Level

view

Output Fields

Sample Output

show services sessions utilization

```
user@host> show services sessions utilization
           Session  %Memory  %Session-Memory    Setup   %Rate   Drop     Teardown   %CPU
Interface  Count                               Rate            Rate     Rate
vms-3/0/0  0        24.96    0.00               0                        0          0.13    Green
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services soft-gre tunnel

IN THIS SECTION

- Syntax | 26186
- Description | 26187
- Options | 26187
- Required Privilege Level | 26187
- Output Fields | 26187
- Sample Output | 26189
- Release Information | 26189

Syntax

```
show services soft-gre tunnel
<brief | detail | extensive>
<interface interface-name>
<local-ip local-ip-address>
<remote-ip remote-ip-address>
```

```
<statistics>
<tunnel-group group-name>
```

Description

Display information about dynamic generic routing encapsulation (GRE) tunnels and pseudowire subscriber (ps*n*) interface devices.

Options

none	Display standard information about all active dynamic GRE tunnels.
brief | detail | extensive	(Optional) Display the specified level of detail.
interface *interface-name*	(Optional) Display information for a specific pseudowire subscriber (ps*n*) interface.
local-ip *local-ip-address*	(Optional) Display information for a specific local or source IP address of the GRE tunnel endpoint (Wi-Fi access gateway side).
remote-ip *remote-ip-address*	(Optional) Display information for a specific remote IP address of the GRE tunnel endpoint.
statistics	(Optional) Display dynamic GRE tunnel statistics.
tunnel-group *group-name*	(Optional) Display information for a specific dynamic GRE tunnel group.

Required Privilege Level

view

Output Fields

Table 1515 on page 26188 describes the output fields for the `show services soft-gre tunnel` command. Output fields are listed in the approximate order in which they appear.

Table 1515: show services soft-gre tunnel Output Fields

Field Name	Field Description	Level of Output
`Interface` or `Interface Name`	Name of the pseudowire subscriber (ps0) interface configured for the GRE tunnel	All levels
`Remote IP`	Remote IP address of the GRE tunnel endpoint	All levels
`Local IP`	Local or source IP address of the GRE tunnel endpoint (Wi-Fi access gateway side)	All levels
`Subscribers`	Number of subscribers accessing the GRE tunnel	All levels
`Group Name`	Name of the dynamic GRE tunnel group	`detail` `extensive`
`Routing Instance`	Type of routing instance used for the GRE tunnel	`detail` `extensive`
`Create time`	Date and time when the GRE tunnel was created: *yyyy-mm-dd hh:mm:ss timezone*	`detail` `extensive`
`Statistics since`	Day of the week, date, time, and year when statistics for packets received and transmitted were recorded: *Dayofweek Month date hh:mm:ss yyyy*	`extensive`
`Statistic`	Type of data through the GRE tunnel: Data Rx (data received) and Data Tx (transmitted)	`extensive`
`Packets`	Number of data packets received (Data Rx) and transmitted (Data Tx) through the GRE tunnel	`extensive`
`Bytes`	Number of data bytes received (Data Rx) and transmitted (Data Tx) through the GRE tunnel	`extensive`

Sample Output

show services soft-gre tunnel brief

```
user@host> show services soft-gre tunnel brief
Interface          Remote IP               Local IP                Subscribers
ps0.3221225475     192.0.2.10                198.51.100.1                    1
```

show services soft-gre tunnel detail

```
user@host> show services soft-gre tunnel detail
Interface Name: ps0.3221225475, Group Name: landslide_v4_1
    Local IP: 198.51.100.1
    Remote IP: 192.0.2.10
    Subscribers: 1
    Routing Instance: default
    Create time: 2015-03-02 08:06:28 PST
```

show services soft-gre tunnel extensive

```
user@host> show services soft-gre tunnel extensive
Interface Name: ps0.3221225475, Group Name: landslide_v4_1
    Local IP: 198.51.100.1
    Remote IP: 192.0.2.10
    Subscribers: 1
    Routing Instance: default
    Create time: 2015-03-02 08:06:28 PST
    Statistics since: Mon Mar  2 08:06:28 2015
      Statistic         Packets        Bytes
       Data Rx               31         3324
       Data Tx               32         3613
```

Release Information

Command introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Wi-Fi Access Gateway Deployment Model Overview

Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway

Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation

soft-gre | 28083

show services softwire

IN THIS SECTION

- Syntax | 26190
- Description | 26190
- Options | 26190
- Required Privilege Level | 26191
- Output Fields | 26191
- Sample Output | 26191
- Release Information | 26192

Syntax

```
show services softwire
```

Description

Display information about softwire services. Information is displayed on both 6rd and DS-Lite services.

Options

count *interface-name*	(Optional) Display the current softwire counts for a service set for both DS-Lite and 6rd.
count	(Optional) Display the number of created softwires.

Required Privilege Level

view

Output Fields

Table 1516 on page 26191 lists the output fields for the `command-name` command. Output fields are listed in the approximate order in which they appear.

Table 1516: show-services-softwire Output Fields

Field Name	Field Description	Level of Output
`Interface`	Interface for which information is displayed.	All levels
`Service Set`	Service set containing the softwire rules for the interface.	All levels
`Softwire`	Name of the softwire concentrator.	All levels
`Direction`	Direction of the flow.	All levels
`Flow count`	Number of flows.	All levels

Sample Output

show services softwire

```
user@host> show services softwire
Interface: sp-3/0/0, Service set: v6rd-dom1-dom3-service-set
Softwire                                    Direction     Flow count
10.10.10.2      ->     192.0.2.1              I                 13
```

show services softwire count (sp- interfaces)

```
user@host> show services softwire count
Interface   Service set      DS-Lite        6RD
sp-0/0/0    dslite-svc-set1     2            0
```

show services softwires count (vms- interfaces

```
user@host> show services softwire count
Interface   Service set    DS-Lite       6RD       MAPE
vms-2/0/0    vms-sset10       1          0
```

Release Information

Command introduced in Junos OS Release 10.4.

`count` option added in Junos OS Release 11.2.

Support added for Next Gen Services in Junos OS Release 20.2 on the MX-SPC3 security services card.

show services softwire flows

IN THIS SECTION

- Syntax | 26193
- Description | 26193
- Options | 26193
- Required Privilege Level | 26194
- Output Fields | 26194
- Sample Output | 26195
- Release Information | 26197

Syntax

```
show services softwire flows
(<interface interface-name> <service-set service-set-name>|
count <interface interface-name> <service-set service-set-name>|
ds-lite <B4 b4-address> <AFTR aftr-address>|
v6rd <initiator initiator-ip-address><concentrator concentrator-ip-address>)
```

Description

Display statistics information about the softwire flows.

> **NOTE**: Starting with Junos OS Release 14.1R4, the IPv6 prefix length associated with a subscriber's basic broadband bridging device that is subject to a limited number of sessions (`dslite-ipv6-prefix-length` attribute) is taken into account while the session count is calculated and displayed in the output of the show services softwire flows command. Until Junos OS Release 14.1R3, only IPv4 flows were counted and IPv6 flows were not considered for the statistics about softwire flows

Options

interface *interface-name*	(Optional) Display statistics information about the specified interface only.
service-set *service-set-name*	(Optional) Display statistics information about the specified service set only.
count <interface *interface-name*> <service-set *service-set-name*>|	(Optional) Display flow count information only, with optional filtering by interface and service set.
ds-lite <B4 *b4-address*> <AFTR *aftr-address*>|	(Optional) Display DS-Lite flow information, with optional filtering by B4 (softwire initiator) and AFTR (softwire concentrator).
v6rd <initiator *initiator-ip-address*><concentrator *concentrator-ip-address*>)	(Optional) Display v6rd flow information, with optional filtering by the softwire initiator and softwire concentrator.

Required Privilege Level

view

Output Fields

Table 1517 on page 26194 lists the output fields for the `show services softwire flows` command. Output fields are listed in the approximate order in which they appear.

Table 1517: show services softwire flows Output Fields

Field Name	Field Description
`Interface`	Name of the interface.
`Service set`	Name of the service set.
`Flow`	Description of flow, including protocol input and output interface addresses.
`State`	Flow state. Value is: • `Forward`
`Dir`	Flow direction. Values are: • `I`—inbound • `O`—outbound
`Frm count`	Number of frames transferred.
`NAT dest`	NAT translation of the decapsulated address.
`Softwire`	For outbound flows, the address of the local softwire initiator (B4 for DS-Lite) is shown first, followed by the address of the softwire concentrator (AFTR for DS-Lite). For inbound flows, the address of the software concentrator is shown first, followed by the address of the softwire initiator.

Sample Output

show services softwire flows

```
user@host> show services softwire flows
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow                                                State       Dir       Frm count
TCP      200.200.200.2:80    ->     33.33.33.1:1066  Forward  O         2005418
    NAT dest        33.33.33.1:1066    ->       20.20.1.2:1025
    Softwire           1001::1         ->         2001::2
TCP          20.20.1.2:1025  ->  200.200.200.2:80    Forward  I         2007168
    NAT source       20.20.1.2:1025    ->      33.33.33.1:1066
    Softwire           2001::2         ->         1001::1
TCP          20.20.1.2:1025  ->  200.200.200.2:80    Forward  I         2635998
    NAT source       20.20.1.2:1025    ->      33.33.33.1:1065
    Softwire           2001::3         ->         1001::1
DS-LITE         2001::2      ->        1001::1       Forward  I         2008157
TCP      200.200.200.2:80    ->     33.33.33.1:1065  Forward  O         2637909
    NAT dest        33.33.33.1:1065    ->       20.20.1.2:1025
    Softwire           1001::1         ->         2001::3
DS-LITE         2001::3      ->        1001::1       Forward  I         2640499
```

show services softwire flows count

```
user@host> show services softwire flows count
Interface   Service set                                                Flow count
sp-0/0/0    dslite-svc-set1                                           6
```

show services softwire flows ds-lite B4

```
user@host> show services softwire flows ds-lite B4 2001::2
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow                                                State       Dir       Frm count
TCP      200.200.200.2:80    ->     33.33.33.1:1066  Forward  O         2884037
    NAT dest        33.33.33.1:1066    ->       20.20.1.2:1025
    Softwire           1001::1         ->         2001::2
TCP          20.20.1.2:1025  ->  200.200.200.2:80    Forward  I         2885884
    NAT source       20.20.1.2:1025    ->      33.33.33.1:1066
```

```
    Softwire           2001::2         ->         1001::1
DS-LITE         2001::2      ->        1001::1       Forward  I         2886821
```

show services softwire flows ds-lite AFTR

```
user@host> show services softwire flows ds-lite AFTR 1001::1
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow                                                State       Dir       Frm count
TCP      200.200.200.2:80    ->     33.33.33.1:1066  Forward  O         3359356
    NAT dest        33.33.33.1:1066    ->       20.20.1.2:1025
    Softwire           1001::1         ->         2001::2
TCP          20.20.1.2:1025  ->  200.200.200.2:80    Forward  I         3361235
    NAT source       20.20.1.2:1025    ->      33.33.33.1:1066
    Softwire           2001::2         ->         1001::1
TCP          20.20.1.2:1025  ->  200.200.200.2:80    Forward  I         4479810
    NAT source       20.20.1.2:1025    ->      33.33.33.1:1065
    Softwire           2001::3         ->         1001::1
DS-LITE         2001::2      ->        1001::1       Forward  I         3362168
TCP      200.200.200.2:80    ->     33.33.33.1:1065  Forward  O         4481520
    NAT dest        33.33.33.1:1065    ->       20.20.1.2:1025
    Softwire           1001::1         ->         2001::3
DS-LITE         2001::3      ->        1001::1       Forward  I         4484094
```

services softwire flows ds-lite AFTR and B4

```
user@host> show services softwire flows ds-lite AFTR 1001::1 B4 2001::2
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow                                                State       Dir       Frm count
TCP      200.200.200.2:80    ->     33.33.33.1:1066  Forward  O         3931026
    NAT dest        33.33.33.1:1066    ->       20.20.1.2:1025
    Softwire           1001::1         ->         2001::2
TCP          20.20.1.2:1025  ->  200.200.200.2:80    Forward  I         3932792
    NAT source       20.20.1.2:1025    ->      33.33.33.1:1066
    Softwire           2001::2         ->         1001::1
DS-LITE         2001::2      ->        1001::1       Forward  I         3933782
```

show services softwires softwire-types map-e

```
user@host> show services softwires softwire-types map-e mape-tun1
br-address 2001:db8:ffff::1/128; //Mandatory
rule r1 {
    ipv4-prefix 192.0.2.0/24; //Mandatory
    ipv6-prefix 2001:db8:0000::/40; //Mandatory
    ea-bits-length 16; //Mandatory
    psid-offset 4; //Mandatory
    psid-len 8;
}
version 3;
```

Release Information

Command introduced in Junos OS Release 10.2.

Support added for Next Gen Services in Junos OS Release 20.2

show services softwire statistics

IN THIS SECTION

- Syntax | 26198
- Description | 26198
- Options | 26198
- Required Privilege Level | 26198
- Output Fields | 26198
- Sample Output | 26203
- Sample Output | 26207
- Release Information | 26209

Syntax

```
show services softwire statistics
<ds-lite>
<ds-lite>
<inferface interface-name>
<v6rd>
```

Description

Display information about softwire services.

Options

ds-lite	(Optional) Display only DS-Lite.
interface *interface-name*	(Optional) Name of the interface servicing the softwire. When you omit this option, data for all interfaces are shown.
v6rd	(Optional) Display only 6rd statistics.

Required Privilege Level

view

Output Fields

Table 1518 on page 26198 lists the output fields for the `command-name` command. Output fields are listed in the approximate order in which they appear.

Table 1518: command-name Output Fields

Field Name	Field Description	Level of Output
`Service PIC Name`	Name of service PIC for which statistics are shown.	`statistics`
`Softwires Created`	Number of softwires created.	`statistics`

Table 1518: command-name Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Softwires Created for EIF/HP`	Number of softwires created for endpoint-independent filtering (EIF) or hairpinning (HP).	`statistics` for `ds-lite` only
`Softwires Deleted`	Number of softwires deleted.	`statistics`
`Softwires Flows Created`	Number of flows created.	`statistics`
`Softwires Flows Deleted`	Number of flows deleted.	`statistics`
`Slow Path Packets Processed`	Number of packets processed as initial packets in a softwire session. These packets require a rule lookup and setting up of flows; this processing of an initial packet in a flow is called *the slow path*.	`statistics`
`Slow Path Packets Processed for EIF/HP`	Number of slow path EIF/HP packets processed.	`statistics` for `ds-lite` only
`Fast Path Packets Processed`	Number of packets processed that are not *slow path*.	`statistics`
`Fast Path Encapsulated`	Number of packets encapsulated in the fast path.	`statistics`
`Softwire EIF Accept`	Number of packets that matched an EIF entry that initiated the creation of a DS-Lite tunnel. The EIF entry was previously triggered by a DS-Lite packet.	`statistics` for `ds-lite` only
`Rule Match Succeeded`	Number of packets that matched a softwire rule.	`statistics`
`Rule Match Failed`	Number of packets that did not match any softwire rule.	`statistics`

Table 1518: command-name Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IPv6 Packets Fragmented	Number of packets fragmented by the services PIC.	statistics for ds-lite only
IPv4 Client Fragments	Number of IPv4 fragments received from the client end over the softwire tunnel destined to the server.	statistics for ds-lite only
IPv4 Server First Fragments	Number of IPv4 first fragments received from the server destined to go over the softwire tunnel to the client.	statistics for ds-lite only
IPv4 Server More Fragments	Number of IPv4 other fragments (excluding first and last fragment) received from the server destined to go over the softwire tunnel to the client.	statistics for ds-lite only
IPv4 Server Last Fragments	Number of IPv4 last fragments received from the server destined to go over the softwire tunnel to the client.	statistics for ds-lite only
ICMPv4 Packets sent	Number of ICMPv4 packets sent to the softwire concentrator.	statistics
ICMPv4 Error Packets sent	Number of ICMPv4 error packets sent to the softwire concentrator.	statistics
ICMPv6 Packets sent	Number of ICMPv6 packets sent to the softwire concentrator.	statistics
Dropped ICMPv6 packets destined to AFTR	Number of ICMPv6 packets dropped instead of sending to the softwire concentrator.	statistics
Softwire Creation Failed	Number of softwire creation failures.	statistics for ds-lite and 6rd

Table 1518: command-name Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Softwire Creation Failed for EIF/HP`	Number of softwire creation failures for EIF/HP.	`statistics` for `ds-lite` only
`Flow Creation Failed`	Number of flow creation failures.	`statistics`
`Flow Creation Failed for EIF/HP`	Number of flow creation failures for EIF/HP.	`statistics` for `ds-lite` only
`Flow Creation Failed - Retry`	Number of flow creations retried after failure.	`statistics`
`Slow Path Failed`	Number of failures detected in the slow path.	`statistics`
`Slow Path Failed - Retry`	Number of times processing of a packet was reprocessed in the slow path.	`statistics`
`Packet not IPv4-in-IPv6`	Number of IPv4 packets not encapsulated in IPv6.	`statistics` for `ds-lite` only
`IPv6 Fragmentation Error`	Number of IPv6 packets with fragmentation errors.	`statistics`
`Slow Path Failed-IPv6 Next Header Offset`	Number of IPv6 header errors detected in slow path processing.	`statistics` for `ds-lite` only
`Decapsulated Packet not IPv4`	Number of packets without IPv4 inner header.	`statistics` for `ds-lite` only

Table 1518: command-name Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Decap Failed - IPv6 Next Header Offset`	Decapsulation failure due to an unexpected inner header.	`statistics` for `ds-lite` only
`Decap Failed - IPv4 L3 Integrity`	Decapsulation failure due to incorrect Layer 3 data, such as not an IP packet, bad source or destination address, checksum error, or protocol error.	`statistics` for `ds-lite` only
`Decap Failed - IPv4 L4 Integrity`	Decapsulation failure due to incorrect Layer 4 data, such as errors in TCP, UDP, or TCP headers.	`statistics` for `ds-lite` only
`No Softwire ID`	Number of times a softwire ID was not found.	`statistics`
`No Flow Extension`	Number of times flow extensions were not found.	`statistics`
`ICMPv4 Dropped Packets`	Number of ICMPv4 packets dropped.	`statistics`
`Packet not IPv6-in-IPv4`	Number of IPv6 packets not encapsulated in IPv4.	`statistics` for `v6rd` only
`Decapsulated Packet not IPv6`	Number of packets without an IPv6 inner header.	`statistics` for `v6rd` only
`Encapsulation Failed - No packet memory`	Failed to encapsulate IPv6 packets in IPv4 due to low memory.	`statistics` for `v6rd` only
`Flow limit exceeded`	Flow not created because configured maximum flows per softwire is exceeded.	`statistics`

Table 1518: command-name Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Session limit exceeded`	Flow not created because configured maximum DS-Lite softwire sessions per IPv6 prefix is exceeded.	`statistics` for ds-lite only

Sample Output

show services softwire statistics (sp- interfaces)

```
user@host> show services softwire statistics
DS-Lite Statistics:


  Service PIC Name:                             :sp-0/0/0

  Statistics
  ----------

    Softwires Created                           :0
    Softwires Created for EIF/HP                :0
    Softwires Deleted                           :0
    Softwires Flows Created                     :0
    Softwires Flows Deleted                     :0
    Slow Path Packets Processed                 :0
    SLow Path Packets Processed for EIF/HP      :0
    Fast Path Packets Processed                 :0
    Fast Path Packets Encapsulated              :0
    Softwire EIF Accept                         :0
    Rule Match Succeeded                        :0
    Rule Match Failed                           :0
    IPv6 Packets Fragmented                     :0
    IPv4 Client Fragments                       :0
    IPv4 Server First Fragments                 :0
    IPv4 Server More Fragments                  :0
    IPv4 Server Last Fragments                  :0
    ICMPv4 Packets sent                         :0
    ICMPv4 Error Packets sent                   :0
    ICMPv6 Packets sent                         :0
```

```
    Dropped ICMPv6 packets destined to AFTR      :0

  Transient Errors
  ----------------

    Flow Creation Failed - Retry                 :0
    FLow Creation Failed - Retry for EIF/HP      :0
    Slow Path Failed - Retry                     :0

  Errors
  ------

    Softwire Creation Failed                            :0
    Softwire Creation Failed for EIF/HP                 :0
    Flow Creation Failed                                :0
    FLow Creation Failed For EIF/HP                     :0
    Slow Path Failed                                    :0
    Packet not IPv4-in-IPv6                             :0
    IPv6 Fragmentation Error                            :0
    Softwire Creation Failed - IPv6 Next Header Offset  :0
    Decapsulated Packet not IPv4                        :0
    Decap Failed - IPv6 Next Header Offset              :0
    Decap Failed - IPv4 L3 Integrity                    :0
    Decap Failed - IPv4 L4 Integrity                    :0
    No Softwire ID                                      :0
    No Flow Extension                                   :0
    Flow Limit Exceeded                                 :0

6rd Statistics:



 Service PIC Name                                :sp-0/0/0

  Statistics
  ----------

    Softwires Created                            :0
    Softwires Deleted                            :0
    Softwires Flows Created                      :0
    Softwires Flows Deleted                      :0
    Slow Path Packets Processed                  :0
    Fast Path Packets Processed                  :0
```

```
    Fast Path Packets Encapsulated              :0
    Rule Match Failed                           :0
    Rule Match Succeeded                        :0

  Transient Errors
  ----------------

    Flow Creation Failed - Retry                :0
    Slow Path Failed - Retry                    :0

  Errors
  ------

    Softwire Creation Failed                    :0
    Flow Creation Failed                        :0
    Slow Path Failed                            :0
    Packet not IPv6-in-IPv4                     :0
    Slow Path Failed - IPv6 Next Header Offset  :0
    Decapsulated Packet not IPv6                :0
    Encapsulation Failed - No packet memory     :0
    No Softwire ID                              :0
    No Flow Extension                           :0
    ICMPv4 Dropped Packets                      :0
```

show services softwire statistics ds-lite (sp- interfaces)

```
user@host> show services softwire statistics ds-lite
DS-Lite Statistics:


  Service PIC Name:                             :sp-0/0/0

  Statistics
  ----------

    Softwires Created                           :0
    Softwires Created for EIF/HP                :0
    Softwires Deleted                           :0
    Softwires Flows Created                     :0
    Softwires Flows Deleted                     :0
    Slow Path Packets Processed                 :0
```

```
  SLow Path Packets Processed for EIF/HP       :0
  Fast Path Packets Processed                  :0
  Fast Path Packets Encapsulated               :0
  Softwire EIF Accept                          :0
  Rule Match Succeeded                         :0
  Rule Match Failed                            :0
  IPv6 Packets Fragmented                      :0
  IPv4 Client Fragments                        :0
  IPv4 Server First Fragments                  :0
  IPv4 Server More Fragments                   :0
  IPv4 Server Last Fragments                   :0
  ICMPv4 Packets sent                          :0
  ICMPv4 Error Packets sent                    :0
  ICMPv6 Packets sent                          :0
  Dropped ICMPv6 packets destined to AFTR      :0

Transient Errors
----------------

  Flow Creation Failed - Retry                 :0
  FLow Creation Failed - Retry for EIF/HP      :0
  Slow Path Failed - Retry                     :0

Errors
------

  Softwire Creation Failed                            :0
  Softwire Creation Failed for EIF/HP                 :0
  Flow Creation Failed                                :0
  FLow Creation Failed For EIF/HP                     :0
  Slow Path Failed                                    :0
  Packet not IPv4-in-IPv6                             :0
  IPv6 Fragmentation Error                            :0
  Softwire Creation Failed - IPv6 Next Header Offset  :0
  Decapsulated Packet not IPv4                        :0
  Decap Failed - IPv6 Next Header Offset              :0
  Decap Failed - IPv4 L3 Integrity                    :0
  Decap Failed - IPv4 L4 Integrity                    :0
  No Softwire ID                                      :0
  No Flow Extension                                   :0
  Flow Limit Exceeded                                 :0
  Session Limit Exceeded                              :0
```

Sample Output

show services softwire statistics (vms- interfaces)

```
user@host> show services softwire statistics
vms-2/0/0
    Total Session Interest events                :3
    Total Session Destroy events                            :2
    Total Session Public Request events                     :0
    Total Session Accepts                                   :1
    Total Session Discards                                  :0
    Total Session Ignores                                   :0
    Total Session extension alloc failures                  :0
    Total Session extension set failures                    :0
Softwire statistics
    Total Softwire sessions created                         :1
    Total Softwire sessions deleted                         :2
    Total Softwire sessions created for reverse packets     :1
    Total Softwire session create failed for reverse pkts   :0
    Total Softwire rule match success                       :1
    Total Softwire rule match failed                        :0
    Softwire session limit exceeded                         :0
Softwire packet statistics
    Total Packets processed                                 :1
    Total packets encapsulated                              :1
    Total packets decapsulated                              :1
    Encapsulation errors                                    :0
    Decapsulation errors                                    :0
    Encapsulated pkts re-inject failures                    :0
    Decapsulated pkts re-inject failures                    :0
    DS-Lite ICMPv4 Echo replies sent                        :0
    DS-Lite ICMPv4 TTL exceeded messages sent               :0
    ICMPv6 ECHO request messages received destined to AFTR  :0
    ICMPv6 ECHO reply messages sent from AFTR               :0
    ICMPv6 ECHO requests to AFTR process failures           :0
    V6 untunnelled packets destined to AFTR dropped         :1
    Softwire policy add errors                              :0
    Softwire policy delete errors                           :0
    Softwire policy memory alloc failures                   :0
    Softwire Untunnelled packets ignored                    :0
```

```
Softwire Misc errors
    DS-Lite ICMPv4 TTL exceed message process errors        :0
```

show services softwire statistics ds-lite (vms- interfaces)

```
user@host> show services softwire statistics ds-lite interface vms-2/0/0
vms-2/0/0
    Total Session Interest events                :3
    Total Session Destroy events                            :2
    Total Session Public Request events                     :0
    Total Session Accepts                                   :1
    Total Session Discards                                  :0
    Total Session Ignores                                   :0
    Total Session extension alloc failures                  :0
    Total Session extension set failures                    :0
Softwire statistics
    Total Softwire sessions created                         :1
    Total Softwire sessions deleted                         :2
    Total Softwire sessions created for reverse packets     :1
    Total Softwire session create failed for reverse pkts   :0
    Total Softwire rule match success                       :1
    Total Softwire rule match failed                        :0
    Softwire session limit exceeded                         :0
Softwire packet statistics
    Total Packets processed                                 :1
    Total packets encapsulated                              :1
    Total packets decapsulated                              :1
    Encapsulation errors                                    :0
    Decapsulation errors                                    :0
    Encapsulated pkts re-inject failures                    :0
    Decapsulated pkts re-inject failures                    :0
    DS-Lite ICMPv4 Echo replies sent                        :0
    DS-Lite ICMPv4 TTL exceeded messages sent               :0
    ICMPv6 ECHO request messages received destined to AFTR  :0
    ICMPv6 ECHO reply messages sent from AFTR               :0
    ICMPv6 ECHO requests to AFTR process failures           :0
    V6 untunnelled packets destined to AFTR dropped         :1
    Softwire policy add errors                              :0
    Softwire policy delete errors                           :0
    Softwire policy memory alloc failures                   :0
    Softwire Untunnelled packets ignored                    :0
```

```
Softwire Misc errors
    DS-Lite ICMPv4 TTL exceed message process errors        :0
```

Release Information

Command introduced in Junos OS Release 10.4.

Support for Next Gen Services with the MX-SPC3 security services card added in Junos OS Release 20.2.

show services ssl certificate

IN THIS SECTION

- Syntax | 26209
- Description | 26209
- Options | 26210
- Required Privilege Level | 26210
- Output Fields | 26210
- Sample Outputs | 26212
- Release Information | 26215

Syntax

```
show services ssl certificate [all | brief | detail]
<pic-info fpc-slot slot number pic-slot slot-number>
```

Description

Display information about the SSL certificates available on the device.

NOTE: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

Options

all	Display information about all SSL certificates.
brief	Display brief information about SSL certificate.
detail	Display detail information about SSL certificates.
pic-info *fpc-slot slot number pic-slot slot-number*	Display the information for the FPC in the specified slot.

Required Privilege Level

view

Output Fields

Table 1519 on page 26210 lists the output fields for the `show services ssl certificate` command. Output fields are listed in the approximate order in which they appear.

Table 1519: show services ssl certificate Output Fields

Field Name	Field Description	Level of Output
`CertID`	Name of the local digital certificate.	all, brief, detail
`Certificate Type`	Type of certificate. That is—Signing certificate (LOCAL-CERT) which is used to sign other certificates or it is CA-CERT used to verify other certificates in context of SSL-proxy.	brief, detail
`cert modify time`	Indicates the time when the certificate data was last modified.	detail
`key modify time`	Indicates the time when the certificate key was last modified (displayed in local certificate only).	detail

Table 1519: show services ssl certificate Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`certificate version`	Version of the digital certificate.	detail
`Serial number`	Unique serial number of the digital certificate.	detail
`Issuer`	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • `C`—Country of origin. • `ST`—State or province name. • `L`—Locality. • `O`—Organization of origin. • `OU`—Organizational unit. • `CN`—Common name of the authority. • `emailAddress`—Common name of the authority.	brief, detail
`Subject`	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • `C`—Country of origin. • `ST`—State or province name. • `L`—Locality. • `O`—Organization of origin. • `OU`—Organizational unit. • `CN`—Common name of the authority. • `emailAddress`—Common name of the authority.	brief, detail

Table 1519: show services ssl certificate Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`validity`	Validity of the certificate (displayed in local certificate only). It includes: • `not before`—Start time when the digital certificate becomes valid. • `not after`—End time when the digital certificate becomes invalid.	detail
`Public Key algorithm`	Encryption algorithm used with the private key, such as `rsaEncryption (1024 bits)`.	brief, detail
`Signature algorithm`	Encryption algorithm that the CA used to sign the digital certificate, such as `sha1WithRSAEncryption`.	detail
`CRL`	Certificate revocation list related information (displayed for CA certificates only). It includes: • present—list of digital certificates that have been revoked before their expiration date are present or not. • check—CRL check status: enabled or disabled. • download-failed—Indicates the download status of the certificate revocation list. • check-on-download-fail—Indicates checking of the certificates against the local CRL file is enabled when the CRL download fails.	detail

Sample Outputs

show services ssl certificate all

```
user@host > show services ssl certificate all

Lsys Name : root-logical-system
PIC:fwdd0 fpc[0] pic[0] ------
CertId
----------------------------
```

```
ssl-inspect-ca
ssl-cert-4k
```

command-name

```
user@host > show services ssl certificate brief certificate-id

Lsys Name : root-logical-system

PIC:fpc0 fpc[0] pic[0] ------

CertID : trusted-ca
Certificate Type : CA-CERT
Issuer : /C=IN/ST=KA/L=BNG/O=XYZ/OU=ABC/CN=5.0.0.1/emailAddress=newca@test.com
Subject : /C=IN/ST=KA/L=BNG/O=XYZ/OU=ABC/CN=5.0.0.1/emailAddress=newca@test.com
Public Key algorithm : rsaEncryption
```

command-name

```
user@host> show services ssl certificate brief certificate-id ssl-inspect-ca
Lsys Name : root-logical-system

PIC:fpc0 fpc[0] pic[0] ------

CertID                : ssl-inspect-ca
Certificate Type      : LOCAL-CERT
Issuer                : /DC=dc/CN=xyz.com/OU=IT/O=abc/L=bng/ST=KA/C=IN
Subject               : /DC=dc/CN=xyz.com/OU=IT/O=abc/L=bng/ST=KAC=IN
Validity :
    Not before        : Mon 02/18/2019 07:30:37 AM
    Not after         : Sat 02/17/2024 07:30:37 AM
Public Key algorithm  : rsaEncryption
```

show services ssl certificate detail (Local Certificate)

```
user@host > show services ssl certificate detail
Lsys Name : root-logical-system
```

```
PIC:fpc0 fpc[0] pic[0] ------

CertID                : ssl-inspect-ca
Certificate Type      : LOCAL-CERT
cert modify time      : Mon 02/18/2019 07:30:37 AM
key modify time       : Mon 02/18/2019 07:30:23 AM
certificate version   : 3
serial number         : 72 a4 a8 12 0e a0 da 5f ee 27 47 d8 19 7c 76 b5
Issuer                : /DC=dc/CN=XYZ.com/OU=IT/O=jnpr/L=bng/ST=KA/C=IN
Subject               : /DC=dc/CN=XYZ.com/OU=IT/O=jnpr/L=bng/ST=KA/C=IN
Validity :
    Not before        : Mon 02/18/2019 07:30:37 AM
    Not after         : Sat 02/17/2024 07:30:37 AM
Public Key algorithm  : rsaEncryption
Signature Algorithm   : sha256WithRSAEncryption
```

show services ssl certificate detail (CA Certificate)

```
user@host > show services ssl certificate detail
Lsys Name : root-logical-system

PIC:fpc0 fpc[0] pic[0] ------

CertID                : test
Certificate Type      : CA-CERT
cert modify time      : Mon 09/02/2019 09:47:48 PM
certificate version   : 1
serial number         : 21 a8 d6 00 eb 24 1f 78 9a e5 0e ec 6a 39 ce 65 66 42 8c 0a
Issuer                : /C=IN/ST=KA/L=BLR/O=XYZ.com/OU=IT/CN=5.0.0.1/emailAddress=newca@test.com
Subject               : /C=IN/ST=KA/L=BLR/O=XYZ.con/OU=IT/CN=5.0.0.1/emailAddress=newca@test.com
Public Key algorithm  : rsaEncryption
Signature Algorithm   : sha256WithRSAEncryption
CRL :
    present           : no
    check             : enabled
    download-failed   : true
    check-on-download-fail : enabled
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl droplogs

IN THIS SECTION

- Syntax | 26215
- Description | 26215
- Options | 26216
- Required Privilege Level | 26216
- Sample Output | 26216
- Release Information | 26217

Syntax

```
show services ssl droplogs
pic-info fpc-slot slot number pic-slot slot-number
```

Description

Display the denied or dropped session details. The messages log file records the details about the dropped packets.

A new field is added in the existing droplog command. The field indicates the version that is negotiated between the client and server when any issue occurs.

> **NOTE**: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

Options

pic-info fpc0.pic0 ***fpc-slot slot number pic-slot slot-number***	Display the information for the FPC in the specified slot.

Required Privilege Level

view

Sample Output

command-name

```
user@host > show services ssl droplogs

Lsys Name : root-logical-system
PIC:fpc0 fpc[0] pic[0]-------

 ===========log mesg for cpu 0

 ===========log mesg for cpu 1

 log mesg is  File: ../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../../src/junos/jsf/plugin/ssl/jssl_common.c Function: jssl_X509_verify_cert Line: 3767 Message: unable to get local issuer certificate C2S plugin chain : [Plugin junos-jdpi: action: ignore]-> [Plugin junos-tcp-svr-emul: action: none]-> [Plugin junos-ssl-proxy: action: ignore]-> [Plugin junos-ssl-term: action: none]-> [Plugin junos-dpi-stream: action: none]-> [Plugin junos-idp-stream: action: ignore]-> [Plugin junos-ssl-init: action: none]-> [Plugin junos-tcp-clt-emul: action: none] S2C plugin chain: [Plugin junos-jdpi: action: ignore]-> [Plugin junos-tcp-clt-emul: action: none]-> [Plugin junos-ssl-init: action: none]-> [Plugin junos-dpi-stream: action: none]-> [Plugin junos-idp-stream: action: ignore]-> [Plugin junos-ssl-term: action: none]-> [Plugin junos-ssl-proxy: action: ignore]-> [Plugin junos-tcp-svr-emul: action: none] SourceIP:5.0.0.1 DestIP:4.0.0.1 Source Port:40281 Dest Port:443
```

```
source interface:ge-0/0/1.0 Destination interface:ge-0/0/0.0 source zone:untrust destination
Zone:trust
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl session | 26258

show services ssl initiation counters

IN THIS SECTION

- Syntax | 26217
- Description | 26217
- Options | 26218
- Required Privilege Level | 26218
- Output Fields | 26218
- Sample Output | 26220
- Release Information | 26222

Syntax

```
show services ssl initiation counters [all | error | handshake)
pic-info fpc-slot slot number pic-slot slot-number
```

Description

Display statistical counters for the SSL initiation sessions.

> **NOTE**: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

Options

pic-info fpc0.pic0 ***fpc-slot slot number pic-slot slot-number***	Display the information for the FPC in the specified slot.
all	Display all the counters generated during SSL initiation.
error	Display all the counters related to errors occurred during SSL initiation.
handshake	Display all the counters related to handshake during SSL initiation.

Required Privilege Level

view

Output Fields

Table 1520 on page 26218 lists the output fields for the `show services ssl initiation counters` command. Output fields are listed in the approximate order in which they appear.

Table 1520: show services ssl initiation counters Fields

Field Name	Field Description	Display Level
`Memory errors`	Errors related to memory allocation.	all, errors
`Handshake errors`	Number of errors occurred during handshake.	all, errors
`Cert Cache errors`	Number of certificate cache errors.	all, errors
`Server Protection errors`	Errors occurred during SSL reverse proxy.	all, errors
`Proxy errors`	Errors occurred in SSL proxy sessions.	all, errors

Table 1520: show services ssl initiation counters Fields *(Continued)*

Field Name	Field Description	Display Level
`Crypto errors`	Number of crypto errors.	all, errors
`Certificate errors`	Errors related to digital certificate	all, errors
`One-Crypto errors`	Number of one-crypto errors	all, errors
`Async-Crypto errors`	Number of async-crypto errors	all, errors
`Mirror errors`	SSL decryption mirrors	all, errors
`handshakes started`	Number of SSL handshakes started.	all, errors
`handshakes completed`	Number of SSL handshakes completed successfully.	all, errors
`active sessions`	Number of active SSL sessions	all, errors
`Interdicted cert generated`	Number of interdicted certificates generated	all, errors
`proxy: sessions created`	Number of proxy sessions created	all, errors
`proxy: sessions active`	Number of active proxy sessions	all, errors
`proxy: sessions ignored`	Number of proxy sessions ignored.	all, errors
`proxy: renegotiation ignored`	Number of renegotiation requests ignored.	all, errors
`proxy: session resumption`	Number of session resumption requests	all, errors
`proxy: secure renegotiation`	Number of SSL sessions with secure renegotiation	all, errors

Table 1520: show services ssl initiation counters Fields *(Continued)*

Field Name	Field Description	Display Level
`proxy: insecure renegotiation`	Number of SSL sessions with insecure renegotiation	all, errors
`proxy: multiple renegotiation`	Number os SSL sessions with multiple renegotiation	all, errors
`proxy: reneg after resumption`	Number os SSL sessions undergo renegotiation after resumption	all, errors
`init: passthrough requests`	Passthrough requests during initiation	all, errors
`init: start requests`	Start requests during initiation	all, errors
`proxy: ECDSA based srvr auth`	Sessions completed ECDSA-based server authentication	all, errors
`proxy: RSA based srvr auth`	Sessions completed RSA-based server authentication	all, errors

Sample Output

show services ssl initiation counters all

```
user@host >  show services ssl initiation counters all
Lsys Name : root-logical-system

PIC:fpc0 fpc[0] pic[0] ------

Memory errors                                   0
Handshake errors                                0
Cert Cache errors                               0
Server Protection errors                        0
Proxy errors                                    0
Crypto errors                                   0
Certificate errors                              0
One-Crypto errors                               0
Async-Crypto errors                             0
Mirror errors                                   0
```

```
handshakes started                              0
handshakes completed                            0
active sessions                                 0
Interdicted cert generated                      0
proxy: sessions created                         0
proxy: sessions active                          0
proxy: sessions ignored                         0
proxy: renegotiation ignored                    0
proxy: session resumption                       0
proxy: secure renegotiation                     0
proxy: insecure renegotiation                   0
proxy: multiple renegotiation                   0
proxy: reneg after resumption                   0
init: passthrough requests                      0
init: start requests                            0
proxy: ECDSA based srvr auth                    0
proxy: RSA based srvr auth                      0
```

show services ssl initiation counters error

```
user@host > show services ssl initiation counters error
Lsys Name : root-logical-system

PIC:fpc0 fpc[0] pic[0] ------

Memory errors 0
Handshake errors 0
Cert Cache errors 0
Server Protection errors 0
Proxy errors 0
Crypto errors 0
Certificate errors 0
One-Crypto errors 0
Async-Crypto errors 0
Mirror errors 0
```

show services ssl initiation counters handshake

```
user@host > show services ssl initiation counters handshake
Lsys Name : root-logical-system
```

```
PIC:fpc0 fpc[0] pic[0] ------

handshakes started 0
handshakes completed 0
active sessions 0
Interdicted cert generated 0
proxy: sessions created 0
proxy: sessions active 0
proxy: sessions ignored 0
proxy: renegotiation ignored 0
proxy: session resumption 0
proxy: secure renegotiation 0
proxy: insecure renegotiation 0
proxy: multiple renegotiation 0
proxy: reneg after resumption 0
init: passthrough requests 0
init: start requests 0
proxy: ECDSA based srvr auth 0
proxy: RSA based srvr auth 0
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl initiation profile | 26222

show services ssl initiation profile

IN THIS SECTION

- Syntax | 26223
- Description | 26223

- Options | 26223
- Required Privilege Level | 26223
- Output Fields | 26224
- Sample Output | 26226
- Release Information | 26227

Syntax

```
show services ssl initiation profile   [all | brief | detail]
pic-info fpc-slot slot number pic-slot slot-number
```

Description

Display the SSL initiation profiles details.

> **NOTE**: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

Options

pic-info *fpc-slot slot number pic-slot slot-number*	Display the information for the FPC in the specified slot.
all	Display all SSL initiation profiles configured on the device.
brief	Display brief information about SSL initiation profiles.
detail	Display detail information about SSL initiation profiles.

Required Privilege Level

view

Output Fields

Table 1521 on page 26224 lists the output fields for the `show services ssl initiation profile` command. Output fields are listed in the approximate order in which they appear.

Table 1521: show services ssl initiation profile Output Fields

Field Name	Field Description	Output Levels
`Profile`	SSL initiation profile name	brief, detail
`allow non-ssl session`	Allow or not allow (bypass) non-SSL sessions.	brief, detail
`preferred-ciphers`	SSL cipher that can be used with acceptable key strength. Possible values are strong, medium, weak, and custom.	brief, detail
`Num of url categories configured`	URL categories exempted from SSL proxy.	brief, detail
`Protocol-version`	SSL protocol version. Possible values are all, TLS version 1.0, TLS version 1.1, and TLS version 1.2.	detail
`Client authentication`	Status of client certificate verification process.	detail
`Server Authentication`	Status of server certificate verification process.	detail
`Crypto-mode`	Crypto mode used. Options are synchronous-hardware or software or asynchronous-hardware.	detail
`Session Resumption`	SSL session resumption status.	detail
`CRL check`	Status of the CRL checking of certificate validity.	detail
`Certificate`	Digital certificate used.	detail
`Renegotiation`	Renegotiation option. Possible values are allow, allow secure, and drop.	detail

Table 1521: show services ssl initiation profile Output Fields *(Continued)*

Field Name	Field Description	Output Levels
`Custom ciphers`	Custom ciphers configured.	detail
`Server Cert`	Server certificate configured.	detail
`Decrypt Mirror`	Status of decrypt mirroring functionality.	detail
`Trusted CA:`	Trusted CA configured for a profile	detail
`handshakes started`	Number of SSL handshakes started.	detail
`handshakes completed`	Number of SSL handshakes completed successfully.	detail
`active sessions`	Number of active SSL sessions	detail
`total handshake errors`	Number of errors occurred during handshake process.	detail
`data errors`	Cumulative errors in a single counter	
`session resumption`	Number of SSL session resumption count.	detail
`secure renegotiation`	Secure sessions allowed after renegotiation.	detail
`insecure renegotiation`	All sessions allowed after renegotiation.	detail
`multiple renegotiation`	Sessions with multiple renegotiation.	detail
`reneg after resumption`	Sessions undergoing renegotiation after resumption.	detail
`no_reneg alert by peer`	Number of times no renegotiation alerts received from peer.	detail

Table 1521: show services ssl initiation profile Output Fields *(Continued)*

Field Name	Field Description	Output Levels
drop on reneg	Sessions dropped after renegotiation.	detail

Sample Output

command-name

```
user@host > show services ssl initiation profile all
Lsys Name : root-logical-system

PIC:fpc0 fpc[0] pic[0] ------
ID          Name

65536  SSL_PROFILE_65536_proxy_i
```

command-name

```
user@host > show services ssl initiation profile brief profile-name
Lsys Name : root-logical-system

PIC: fpc0 fpc[0] pic[0] ----------

Profile                          : SSL_PROFILE_65536_proxy_i
allow non-ssl session            : true
preferred-ciphers                : medium
Num of url categories configured : 0
```

command-name

```
user@host > show services ssl initiation profile detail profile-name

Lsys Name : root-logical-system
```

```
PIC: fpc0 fpc[0] pic[0] ----------

Profile                          : SSL_PROFILE_65536_proxy_i
allow non-ssl session            : true
preferred-ciphers                : medium
Num of url categories configured : 0
Protocol version                 : all
Client Authentication            : notset
Server Authentication            : Ignore Failure
Crypto Mode                      : sw
Session Resumption               : Enabled
CRL check                        : Enabled
Certficate RSA : ssl-inspect-ca
Renegotiation        : only secure allowed
Custom ciphers       : 0
Server cert          : 0
Decrypt Mirror       : Disabled
Trusted CA           : 0
       handshakes started            0
       handshakes completed          0
       active sessions               0
       total handshake errors        0
       Data Errors                   0
       session resumption            0
       secure renegotiation          0
       insecure renegotiation        0
       multiple renegotiation        0
       reneg after resumption        0
       no_reneg alert by peer        0
       drop on reneg                 0
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl initiation counters | 26217

show services ssl proxy certificate-cache entries

IN THIS SECTION

- Syntax | 26228
- Description | 26228
- Options | 26228
- Required Privilege Level | 26229
- Output Fields | 26229
- Sample Output | 26230
- Release Information | 26230

Syntax

```
show services ssl proxy certificate-cache entries [detail |  summary]
<pic-info fpc-slot slot number pic-slot slot-number>
```

Description

Display information about the entries stored in the SSL proxy certificate cache.

> **NOTE**: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

Options

pic-info ***fpc-slot slot number pic-slot slot-number***	Display the information for the FPC in the specified slot.
detail	Display the detail information about the SSL proxy certificate cache entries.
summary	Display the summary of the SSL proxy certificate cache entries.

Required Privilege Level

view

Output Fields

Table 1522 on page 26229 lists the output fields for the `show services ssl proxy certificate-cache` command. Output fields are listed in the approximate order in which they appear.

Table 1522: show services ssl proxy certificate-cache Output Fields

Field Name	Field Description	Level of Output
`Cache Entries`	Index number of the entry.	summary, detail
`Serial number`	Serial number of the server certificate.	summary, detail
`SSL-I Profile Id`	SSL initiation profile identification number.	summary, detail
`Num of CRL updates`	Number of times the CRL updates done till the interdicted certificate is added to the certificate-cache.	summary, detail
`Status`	Status of the cache entry. That is—whether the cache entry has expired or not, because the cache entries are valid only for short interval.	detail
`Interdicted cert type`	Interdicted certificate details such as type and authentication status.	detail
`Server cert verification result`	Server certificate validation results.	detail
`Cert reference count`	Certificate reference count.	detail
`Issuer`	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format.	detail
`Subject`	Details of the digital certificate holder organized using the distinguished name format.	detail

Sample Output

command-name

```
user@host > show services ssl proxy certificate-cache entries summary
Lsys Name : root-logical-system
PIC:fwdd0 fpc[0] pic[0] ------
Cache Entries : 1
Serial number : 0x12345678
SSL-I Profile Id: 1
Num of CRL updates: 0
```

command-name

```
user@host > show services ssl proxy certificate-cache entries detail
Lsys Name : root-logical-system
PIC:fwdd0 fpc[0] pic[0] ------
Cache entrie : 1
Serial number : 0x12345678
SSL-I Profile Id: 1
Num of CRL updates: 0
Status: Active: Time to expire 570 seconds

Cert Info:
-------------
Interdicted cert type [0x0]: CA issued, Authentication failed
Server cert verification result: unable to get local issuer certificate [0x14]
Cert reference count: 2
Subject: /C=IN/ST=KA/O=XYZ Inc/CN=XYZ Root CA/emailAddress=host@xyz.com
Issuer: /CN=SSL-PROXY:DUMMY_CERT:GENERATED DUE TO SRVR AUTH FAILURE
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl proxy certificate–cache statistics

IN THIS SECTION

- Syntax | 26231
- Description | 26231
- Options | 26231
- Required Privilege Level | 26231
- Output Fields | 26232
- Sample Output | 26232
- Release Information | 26232

Syntax

```
show services ssl proxy certificate-cache statistics
<pic-info fpc-slot slot number pic-slot slot-number>
```

Description

Display SSL proxy certificate cache statistics.

Options

pic-info *fpc-slot slot number pic-slot slot-number*	Display the information for the FPC in the specified slot.

Required Privilege Level

view

Output Fields

Table 1523 on page 26232 lists the output fields for the `show services ssl proxy certificate-cache statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1523: show services ssl proxy certificate–cache statistics Output Fields

Field Name	Field Description
`cert cache hit`	Number times the certificate matched the entry in the certificate cache
`cert cache miss`	Number times the certificate did not find the match in the certificate cache
`cert cache full`	Number of times the certificate cache limit is reached.

Sample Output

command-name

```
user@host > show services ssl proxy certificate-cache statistics
Lsys Name : root-logical-system

PIC: fpc0 fpc[0] pic[0]------------

cert cache hit 0
cert cache miss 0
cert cache full
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl proxy certificate-cache entries | 26228

show services ssl proxy counters

IN THIS SECTION

- Syntax | 26233
- Description | 26233
- Options | 26233
- Required Privilege Level | 26234
- Output Fields | 26234
- Sample Output | 26237
- Release Information | 26239

Syntax

```
 show services ssl proxy counters  [all | errors | info]
<pic-info fpc-slot slot number pic-slot slot-number>
```

Description

Display statistical counters for the SSL proxy sessions.

NOTE: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

Options

all	Display information about counter values for all SSL proxy sessions
errors	Display information about counter values for all SSL proxy sessions errors.
info	Display some informational counters which are subset of all the counters

pic-info ***fpc-slot slot number pic-slot slot-number***	Display the information for the FPC in the specified slot.

Required Privilege Level

view

Output Fields

Table 1524 on page 26234 lists the output fields for the `show services ssl certificate` command. Output fields are listed in the approximate order in which they appear.

Table 1524: show services ssl proxy counters Output Fields

Field Name	Field Description	Level of Output
`Session create failed`	The number of failed proxy sessions	errors, all
`non SSL sessions received`	The number of non-SSL sessions received	errors, all
`memory failures`	The number of errors related to memory. Example, memory errors such as the device is on “low memory” is indicated by this counter.	errors, all
`session dropped`	The number of dropped proxy sessions.	errors, all
`sessions matched`	The number of matched proxy sessions.	info, all
`sessions created`	The number of newly created proxy sessions.	info, all
`sessions destroyed`	The number of dropped or destroyed proxy sessions.	info, all
`sessions ignored`	The number of proxy sessions that are ignored.	info, all

Table 1524: show services ssl proxy counters Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`sessions ignored : backup only`	The number of sessions ignored on the backup node in a chassis cluster setup. In chassis cluster or high-availability mode, the SSL session is processed only on the active node and on the backup node session is ignored. This counter indicates the session ignored on the backup node.	info, all
`sessions whitelisted : IP based`	The number of all sessions that are allowlisted based on IP addresses.	info, all
`sessions whitelisted : url based`	The number of all sessions that are allowlisted based on the URL categories.	info, all
`crl : data added`	The number of times CRL data is added.	info, all
`crl : certificate revoked`	The number of sessions dropped because of checking for revoked certificates from servers.	info, all
`crl : no crl info present`	The number of sessions dropped because no CRL information was present.	info, all
`crl : no CA certificate`	The number of sessions dropped because no CA certificate was present.	info, all
`SSL sessions`	Number of SSL sessions	info, all
`SMTP over STARTTLS`	Number of SMTP over STARTTLS sessions	info, all
`IMAP over STARTTLS`	Number of IMAP over STARTTLS sessions	info, all
`POP3 over STARTTLS`	Number of POP3 over STARTTLS sessions	info, all
`SMTP sessions`	Number of SMTP sessions	info, all
`IMAP sessions`	Number of IMAP sessions	info, all

Table 1524: show services ssl proxy counters Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`POP3 sessions`	Number of POP3 sessions	info, all
`Server not supporting STARTTLS`	Number of times server not supported STARTTLS sessions	info, all
`Client not supporting STARTTLS`	Number of times client not supported STARTTLS sessions	info, all
`Unified policy : default profile hit`	The number of times sessions matched default SSL proxy profile.	info, all
`Unified policy : no default profile`	The number of times sessions are dropped because no default SSL proxy profile available.	info, all
`proxy sess matched with early dynapp`	The number of times sessions matched after receiving the dynamic application details from SNI.	all
`proxy sess ignored with early dynapp`	The number of times proxy sessions are disengaged after receiving details from SNI because no SSL proxy profile was configured or the matched pre-identification default policy action was to ignore the session.	all
`proxy sess matched with ssl as dynapp`	The number of times sessions are matched because the sessions received unknown application details from SNI or the sessions have not received details from SNI.	all
`proxy sess ignored with ssl as dynapp`	The number of times proxy sessions are disengaged either because SSL proxy profile was not configured for the matched policy or the matched pre-identification default policy action was to ignore the session	all
`proxy sess matched with default fw policy`	The number of times sessions matched after receiving the dynamic application details from SNI and identified application matched with default security policy.	all

Table 1524: show services ssl proxy counters Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`proxy sess ignored with default fw policy`	The number of times sessions disengaged because the identified dynamic application details from SNI has not matched with the default security policy.	all
`proxy sess matched with pre-id fw policy`	The number of times sessions matched after receiving the dynamic application details from SNI and the application matched pre-identification default policy.	all
`proxy sess ignored with pre-id fw policy`	The number of times sessions disengaged because the identified dynamic application details from SNI has not matched with the pre-identification default security policy.	all

Sample Output

show services ssl proxy counters info

```
user@host > show services ssl proxy counters info

Lsys Name : root-logical-system
PIC:fpc0 ------

sessions matched 0
sessions created 0
sessions destroyed 0
sessions ignored 0
sessions ignored : backup only 0
sessions whitelisted : IP based 0
sessions whitelisted : url based 0
crl : data added 1
crl : certificate revoked 0
crl : no crl info present 0
crl : no CA certificate 0
SSL sessions 0
SMTP over STARTTLS 0
IMAP over STARTTLS 0
```

```
POP3 over STARTTLS 0
SMTP sessions 0
IMAP sessions 0
POP3 sessions 0
Server not supporting STARTTLS 0
Client not supporting STARTTLS 0
Unified policy : default profile hit 0
Unified policy : no default profile 0
```

show services ssl proxy counters errors

```
user@host > show services ssl proxy  counters errors
Lsys Name : root-logical-system
PIC:fpc0 ------

Session create failed 0
non SSL sessions received 0
memory failures 0
session dropped 7
```

show services ssl proxy counters all

```
user@host > show services ssl proxy  counters all
Lsys Name : root-logical-system

PIC:fpc0 fpc[0] pic[0] ------

session create failed                          0
non SSL sessions recieved                      0
Memory failures                                0
session dropped                                0
sessions matched                               0
sessions created                               0
sessions destroyed                             0
sessions ignored                               0
sessions ignored : backup only                 0
sessions whitelisted : IP based                0
sessions whitelisted : url based               0
crl : data added                               0
crl : certificate revoked                      0
```

```
crl : no crl info present                    0
crl : no CA certificate                      0
SSL sessions                                 0
SMTP over STARTTLS                           0
IMAP over STARTTLS                           0
POP3 over STARTTLS                           0
SMTP  sessions                               0
IMAP  sessions                               0
POP3  sessions                               0
Server not supporting STARTTLS               0
Client not supporting STARTTLS               0
Unified policy : default profile hit         0
Unified policy : no default profile          0

proxy sess matched with early dynapp        : 0
proxy sess ignored with early dynapp          : 1
proxy sess matched with ssl as dynapp        : 0
proxy sess ignored with ssl as dynapp     : 0
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl proxy status | 26255

show services ssl proxy session-cache entries | 26243

show services ssl proxy session-cache statistics | 26249

show services ssl proxy profile

IN THIS SECTION

Syntax | 26240

- Description | 26240
- Options | 26240
- Required Privilege Level | 26240
- Output Fields | 26241
- Sample Output | 26241
- Sample Output | 26242
- Release Information | 26242

Syntax

```
show services ssl proxy profile   [all | profile-name]
<pic-info fpc-slot slot number pic-slot slot-number>
```

Description

Display information about the SSL proxy profile details.

> **NOTE**: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

Options

all	Display all SSL proxy profiles configured on the device.
profile-name	Display information about SSL proxy profile.
pic-info ***fpc-slot slot number pic-slot slot-number***	Display the information for the FPC in the specified slot.

Required Privilege Level

view

Output Fields

Table 1525 on page 26241 lists the output fields for the `show services ssl proxy profile` command. Output fields are listed in the approximate order in which they appear.

Table 1525: show services ssl proxy profile Output Fields

Field Name	Field Description
`Profile`	SSL proxy profile name.
`enable-tracing`	Enable flow tracing option is set or not set for the profile.
`root-ca expired`	Enable or disable ignoring server authentication when root CA is expired.
`allow non-ssl session`	Allow or not allow (bypass) non-SSL sessions.
`ssl-termination-id`	SSL termination profile ID.
`ssl-initiation-id`	SSL initiation profile ID.
`TLS version supported`	Details of the TLS version supported for this proxy profile.
`Number of whitelist entries`	The number of allowlisted domains (both IP- based and DNS-based) that are configured for this particular SSL proxy profile.

Sample Output

show services ssl proxy profile

```
user@host > show services ssl proxy profile profile-name ssl-proxy

Lsys Name : root-logical-system
PIC:fpc0 fpc[0] pic[0] ------

Profile                : ssl-proxy
```

```
enable-tracing        : false
root-ca expired       : false
allow non-ssl session : true
ssl-termination id    :      65536
ssl-initiation id     :      65536
Number of Whitelist entries:          0
TLS version supported : ALL
```

Sample Output

show services ssl proxy all

```
user@host > show services ssl proxy all

Lsys Name : root-logical-system
PIC:fwdd0 fpc[0] pic[0] ------
ID          Name
10         p1
11         p2
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl proxy counters | 26233

show services ssl proxy status | 26255

show services ssl proxy session-cache entries

IN THIS SECTION

- Syntax | 26243
- Description | 26243
- Options | 26243
- Required Privilege Level | 26244
- Output Fields | 26244
- Sample Output | 26246
- Release Information | 26249

Syntax

```
show services ssl proxy session-cache entries [detail |  summary]
<pic-info fpc-slot slot number pic-slot slot-number>
```

Description

Display information about the entries stored in the SSL proxy session cache.

NOTE: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

Options

pic-info ***fpc-slot slot number pic-slot slot-number***	Display the information for the FPC in the specified slot.
detail	Display the detail information about the SSL proxy session cache entries.
summary	Display the summary of the SSL proxy session cache entries.

Required Privilege Level

view

Output Fields

Table 1526 on page 26244 lists the output fields for the `show services ssl proxy session-cache entries` command. Output fields are listed in the approximate order in which they appear.

Table 1526: show services ssl proxy session-cache entries Output Fields

Field Name	Field Description	Display Level
`Hash Entry`	Index number of the entry.	summary, detail
`Status`	Status of the cache entry--active or expired. The cache entries are valid only for short interval.	summary, detail
`Session Id Length`	Length of the session ID. 32-bit field that identifies an SSL session.	summary, detail
`Session Id`	SSL session identifier.	summary, detail
`Dst IP`	Destination IP address.	summary, detail
`Dst Port`	Destination port number.	summary, detail
`SSL-T Profile Id`	SSL termination profile identification number.	summary, detail
`SSL-I Profile Id`	SSL initiation profile identification number.	summary, detail
`Interdicted cert type [0x0]:`	Interdicted server certificate	detail
`Server cert verification result:`	Server certificate validation results.	detail

Table 1526: show services ssl proxy session-cache entries Output Fields *(Continued)*

Field Name	Field Description	Display Level
`Server name extn len`	Extension length in the TLS server name extension.	detail
`name`	Server name in the TLS server name extension	detail
`Server cert chain hash`	The hash value of the server certificate chain.	detail
`SSL-TERM session:`	SSL termination session details. It includes the following fields. • `SSL ver`—SSL/TLS protocol version • `Compression Method`—Agreed-upon compression method used to compress data and • `Cipher Id`—Identification number for the cipher • `Master Key Length`—Length of the primary secret key.	detail
`SSL-INIT session:`	SSL initiation session details. It includes the following fields. • `SSL ver`—SSL/TLS protocol version • `Compression Method`—Agreed-upon compression method used to compress data and • `Cipher Id`—Identification number for the cipher • `Master Key Length`—Length of the primary secret key.	detail
`SSL-T resumption type`	Session resumption type used in SSL termination	summary
`SSL-I resumption type`	Session resumption type used in an SSL session—PSK-based (TLS1.3) or session-ID based (TLS1.2)	summary
`Resumption type`	Session resumption type used in an SSL session—PSK-based (TLS1.3) or session-ID based (TLS1.2 or earlier versions)	detail
`Ticket valid time`	Validity time of the session ticket that includes pre-shared key (PSK) identity	detail

Table 1526: show services ssl proxy session-cache entries Output Fields *(Continued)*

Field Name	Field Description	Display Level
Tick_lifetime_hint	Life time of the session ticket	detail
Tick_age_add	Age of the session ticket	detail
Ticklen	Length of the session ticket	detail

Sample Output

show services ssl proxy session-cache entries summary

```
user@host > show services ssl proxy session-cache entries summary

Lsys Name : root-logical-system
PIC: fpc0 fpc[0] pic[0
Hash Entry 1
Status: ACTIVE, Time to expire 294 seconds
Session Id Length: 32
Session Id: 1b 2a 9f 5f d8 6e d2 cd 6b b8 89 e8 88 07 75 80 32 c2 54 5a c7 9b 12 a2 e6 5c f0 6d
85 c5 40 4b
Dst IP: 5.0.0.1, Dst Port: 20753
SSL-T Profile Id: 2, SSL-I Profile Id: 2
```

show services ssl proxy session-cache entries summary (Junos OS Release 22.1R1)

```
user@host > show services ssl proxy session-cache entries summaryLsys Name : root-logical-system
PIC:fpc0 fpc[0] pic[0] -------
Hash Entry        :          1
Status            : Active: Time to expire 86240 seconds
Session ID length :         32
Session ID        : 37 7f af 71 36 19 eb 9b 07 16 c0 1e db a3 7f 58 45 1c 61 5b 93 1c 34 28 58
d5 49 05 7e 77 ca 33
Dest IP           : 5.0.0.1
Dest Port         :       9090
SSL_T Profile ID  :          1
```

```
SSL_I Profile ID  :          1
SSL-T resumption type: Session-id
SSL-I resumption type: PSK
```

show services ssl proxy session-cache entries detail

```
user@host > show services ssl proxy session-cache entries detail

Lsys Name : root-logical-system
PIC: fpc0 fpc[0] pic[0
Hash Entry: 1
Status: ACTIVE, Time to expire 294 seconds
Session Id Length: 32
Session Id: c1 6e 88 65 43 9f 57 2f 0f 06 f7 4b 03 c5 38 58 74 b4 4f 43 66 9a 6f c7 a6 2a ae 22
ab f8 b4 ce
Dst IP: 5.0.0.1, Dst Port: 4433
SSL-T Profile Id: 2, SSL-I Profile Id: 2
Session Info:
Interdicted cert type [0x0]: CA issued, Authentication failed
Server cert verification result: unable to get local issuer certificate [0x14]
Server name extn len: 0, name: None
Server cert chain hash: b5 3d cd cb ca 35 81 5a db 6f 83 ab 5e a0 19 73

SSL-TERM session:
SSL ver: 0x303
Compression Method: 0
Cipher Id: 0x3000004
Master Key Length: 48

SSL-INIT session:
SSL ver: 0x303
Compression Method: 0
Cipher Id: 0x3000004
Master Key Length: 48

Hash Entry:2
Status: EXPIRED
Session Id Length: 32
Session Id: 1b 2a 9f 5f d8 6e d2 cd 6b b8 89 e8 88 07 75 80 32 c2 54 5a c7 9b 12 a2 e6 5c f0 6d
85 c5 40 4b
```

```
Dst IP: 5.0.0.1, Dst Port: 4433,
SSL-T Profile Id: 2, SSL-I Profile Id: 2
Session Info:
-------------
Interdicted cert type [0x0]: CA issued, Authentication failed
Server cert verification result: unable to get local issuer certificate [0x14]
Server name extn len: 0, name: None
Server cert chain hash: b5 3d cd cb ca 35 81 5a db 6f 83 ab 5e a0 19 73

SSL-TERM session:
----------------
SSL ver: 0x303
Compression Method: 0
Cipher Id: 0x3000004
Master Key Length: 48

SSL-INIT session:
----------------
SSL ver: 0x303
Compression Method: 0
Cipher Id: 0x3000004
Master Key Length: 48


Stale entry in cache: 1
```

show services ssl proxy session-cache entries detail (Junos OS Release 22.1R1)

```
user@host > show services ssl proxy session-cache entries detail

Lsys Name : root-logical-system
PIC:fpc0 fpc[0] pic[0] -------
Hash Entry        :          1
Status            : Active: Time to expire 86367 seconds
Session ID length :     0
Dest IP           : 9.0.0.1
Dest Port         :       8080
SSL_T Profile ID  :          1
SSL_I Profile ID  :          1
Session Info :
         Interdicted cert type           : [0x1]: CA issued, Authentication Successful
```

```
 Server cert verification result : ok [0x0]
 Server name extn len            :          4 name  : www,example.com
 Server cert chain hash          : 39 da 35 c6 b8 ec fe 9d 56 bf f1 ae 4a bf 93 3f
SSL-TERM Session :
         SSL ver            : 0x303
         Compression method : 0
         Cipher ID          : 0x300c030
         Master key length  : 48
         Resumption type: PSK
SSL-INIT Session :
         SSL ver            : 0x0304
         Compression method : 0
         Cipher ID          : 0x41084d0
         Master key length  : 48
         Resumption type : PSK
         Ticket valid time: 2 hr 23 min 20 sec
         Tick_lifetime_hint: 7200
         Tick_age_add: 2354718221,
         Ticklen: 208
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

show services ssl proxy session-cache statistics | 26249

Operational Commands to Troubleshoot SSL Sessions

show services ssl proxy session-cache statistics

IN THIS SECTION

- Syntax | 26250
- Description | 26250

- Options | 26250
- Required Privilege Level | 26250
- Output Fields | 26250
- Sample Output | 26251
- Release Information | 26251

Syntax

```
show services ssl proxy session-cache statistics
<pic-info fpc-slot slot number pic-slot slot-number>
```

Description

Display the data for the SSL proxy session cache.

Options

pic-info ***fpc-slot slot number pic-slot slot-number***	Display the information for the FPC in the specified slot.

Required Privilege Level

view

Output Fields

Table 1527 on page 26250 lists the output fields for the `show services ssl proxy session-cache statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1527: show services ssl proxy session-cache statistics Output Fields

Field Name	Field Description
`Session cache hit`	Number times the session matched the entry in the SSL proxy session cache.

Table 1527: show services ssl proxy session-cache statistics Output Fields *(Continued)*

Field Name	Field Description
`Session cache miss`	Number times the session did not find the match in the SSL proxy session cache.
`Session cache full`	Number of times the session cache limit is reached.

Sample Output

show services ssl proxy session-cache statistics

```
user@host > show services ssl proxy session-cache statistics

Lsys Name : root-logical-system
PIC: fpc0 fpc[0] pic[0]------------

Session cache hit                :          0
Session cache miss               :          0
Session cache full               :          0
```

Release Information

Command introduced in Junos OS Release 19.3R1.

> **NOTE**: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl proxy session-cache entries | 26243

show services ssl proxy statistics

IN THIS SECTION

- Syntax | 26252
- Description | 26252
- Options | 26252
- Required Privilege Level | 26252
- Output Fields | 26253
- Sample Output | 26254
- Release Information | 26255

Syntax

```
show services ssl proxy statistics
```

Description

Display information about the SSL proxy statistics. An SSL proxy profile defines SSL behavior for the SRX Series Firewall.

NOTE: When devices are operating in chassis cluster mode, the SSL proxy statistics increment only on the active node of the chassis cluster setup.

Options

logical-system Displays summary information about SSL proxy.

Required Privilege Level

view

Output Fields

Table 1528 on page 26253 describes the output fields for the `show services ssl proxy statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1528: show services ssl proxy statistics Output Fields

Field Name	Field Description
`Sessions matched`	The number of proxy sessions that are matched.
`Sessions bypassed: non SSL`	The number of proxy sessions that are bypassed because the non SSL sessions limit was exceeded
`Sessions bypassed: memory overflow`	The number of proxy sessions that are bypassed because the memory usage limit per session was reached.
`sessions bypassed: low memory`	The number of proxy sessions that are bypassed because of low memory on Packet Forwarding Engine.
`Sessions created`	The number of proxy sessions that are newly created.
`Sessions ignored`	The number of proxy sessions that are ignored.
`Sessions active`	The number of proxy sessions that are active.
`Sessions dropped`	The number of proxy sessions that are dropped.
`Sessions whitelisted`	The number of sessions that are allowlisted. Allowlist comprise addresses or domain names that you want to exempt from the SSL proxy processing.
`whitelisted url category match`	Allowlist comprise url hostnames that you want to exempt from the SSL proxy processing.
`default profile hit`	The number of default profiles that are matched when the sessions are allowlisted.

Table 1528: show services ssl proxy statistics Output Fields *(Continued)*

Field Name	Field Description
session dropped no default profile	The number of sessions dropped when no default profiles are matched.
policy hit no profile configured	The number of policies matched when no profile is configured.

Sample Output

show services ssl proxy statistics

```
user@host> show services ssl proxy statistics
PIC:fwdd0 fpc[0] pic[0] ------
        sessions matched                      30647
        sessions bypassed:non-ssl                 0
        sessions bypassed:mem overflow            0
         sessions bypassed:low memory               0
        sessions created                      25665
        sessions ignored                          6
        sessions active                           0
        sessions dropped                          0
        sessions whitelisted                          0
        whitelisted url category match             0
```

show services ssl proxy statistics logical-system

```
user@host> show services ssl proxy statistics logical-system LSYS1
PIC:spu-3 fpc[0] pic[3] ------
        sessions matched                                  1
        sessions bypassed:non-ssl                         0
        sessions bypassed:mem overflow                    0
        sessions bypassed:low memory                      0
        sessions created                                  1
        sessions ignored                                  0
        sessions active                                   1
```

```
        sessions dropped                                0
        sessions whitelisted                            0
        whitelisted url category match                  0
        default profile hit                             0
        session dropped no default profile              0
        policy hit no profile configured                0
```

Release Information

Command introduced in Junos OS Release 12.1.

The `logical system` option is added in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

clear services ssl proxy statistics | 4179

show services ssl proxy status

IN THIS SECTION

- Syntax | 26256
- Description | 26256
- Options | 26256
- Required Privilege Level | 26256
- Output Fields | 26256
- Sample Output | 26257
- Release Information | 26258

Syntax

```
show services ssl proxy status
```

Description

Display information about the SSL proxy status.

Options

pic-info ***fpc-slot slot number pic-slot slot-number***	Display the information for the FPC in the specified slot.

Required Privilege Level

view

Output Fields

Table 1529 on page 26256 list the output fields for the `show services ssl proxy status` command. Output fields are listed in the approximate order in which they appear.

Table 1529: show services ssl proxy status Output Fields

Field Name	Field Description
`One-Crypto`	One-crypto status: enabled or disabled.
`Async Crypto`	Async Crypto status: enabled or disabled.
`Proxy activation`	Status of proxy activation.
`Local logging`	Status of local logging.
`SSLFP <-> PKID Link Status`	SSL forward proxy to PKID link status

Table 1529: show services ssl proxy status Output Fields *(Continued)*

Field Name	Field Description
Certificate cache activated	Status of the certificate cache
Max cert cache nodes	Maximum number of certificates in cache nodes
Invalidate certificate cache on CRL update : Disabled	Status of invalidation of the existing certificate cache
Cert cache node in use	Number of cached certificates in in use
Session cache activated	Status of the session cache
Max session cache node	Maximum number of sessions in cache nodes
Session cache node in use	Number of cached sessions in use.

Sample Output

command-name

```
user@host > show services ssl proxy status
PIC:fwdd0 fpc[0] pic[0] ------
        One-Crypto       :  Enable
        Async Crypto     :  disable
        Proxy-activation :  Only if interested svcs configured
        Local Logging    :  disable
        SSLFP-PKID Link  :  Down
        Certificate cache : -
                Certificate Cache activated                : no
                Invalidate certificate cache on CRL update : Disabled
                Max cert cache nodes  :       4000
                Cert cache node in use :          0
        Session cache : -
                Session cache activated : Deactivated
```

```
                Max session cache node  :          0
                Session cache node in use     :          0
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl proxy counters | 26233

show services ssl proxy profile | 26239

show services ssl session

IN THIS SECTION

- Syntax | 26258
- Description | 26259
- Options | 26259
- Required Privilege Level | 26259
- Output Fields | 26259
- Sample Output | 26261
- Release Information | 26261

Syntax

```
show services ssl session flow-session-id
 pic-info fpc-slot slot number pic-slot slot-number
```

Description

Display information about the Secure Sockets Layer (SSL) session.

> **NOTE**: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

Options

pic-info ***fpc-slot slot number pic-slot slot-number***	Display the information for the FPC in the specified slot.

Required Privilege Level

view

Output Fields

Table 1530 on page 26259 lists the output fields for the `show services ssl session` command. Output fields are listed in the approximate order in which they appear.

Table 1530: show services ssl session Output Fields

Field Name	Field Description
`Session ID`	Session identifier.
`Connection Type`	SSL connection type.
`SSL Profile:`	SSL profile applied - proxy or termination or initiation.
`Resumed Session`	Session resumption applied.
`One-Crypto`	One-crypto status for this particular session: Enabled or Disabled.

Table 1530: show services ssl session Output Fields *(Continued)*

Field Name	Field Description
`Async-crypto`	Async-crypto status for this particular session: Enabled or Disabled.
`Renegotiation Count`	Number of times the session renegotiation was done.
`Server Certificate Subject Name`	Full subject name of the certificate.
`Server Cert verification status`	Status of the server certificate validation.
`Serial Number`	Serial number of the certificate.
`CRL check`	Status of cRL validation is enabled or disabled.
`Action`	Actions related to certification revocations checks.
`SSL_T Details:`	SSL termination details: • Key Size—Server certificates key size • Cipher—SSL Cipher. • TLS version—Protocol version used.
`SSL_I Details`	SSL initiation details: • Key Size—Server certificates key size • Cipher—SSL Cipher. • TLS version—Protocol version used.

Sample Output

show services ssl session

```
user@host > show services ssl session flow-session-id
Lsys Name : root-logical-system

PIC:fpc0 fpc[0] pic[0] ------

Session ID          : 56
Connection Type     : PROXY
SSL Profile         : SSL_PROFILE
Resumed Session     : No
One-crypto          : Disabled
Async-crypto        : Enabled
Renegotiation count : 0
Server Certificate Subject Name    : /C=IN/ST=KAR/L=BNG/O=JN/OU=JNPR/CN=server/emailAddress=ser
Server Cert verification status    : OK
CRL check           : Enabled
Action              : Allow
SSL_T Details :

      Key size      : 2048
      cipher        : ECDHE-RSA-AES256-GCM-SHA384
      TLS version   : 1.2
SSL_I Details :

      Key size      : 2048
      Cipher        : ECDHE-RSA-AES256-GCM-SHA384
      TLS version   : 1.2
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show security flow session ssl | **24306**

show services ssl termination counters

IN THIS SECTION

- Syntax | 26262
- Description | 26262
- Options | 26262
- Required Privilege Level | 26263
- Output Fields | 26263
- Sample Output | 26265
- Release Information | 26267

Syntax

```
show services ssl termination counters [all | errors | handshake]
pic-info fpc-slot slot number pic-slot slot-number
```

Description

Display statistical counters for the SSL termination session.

> **NOTE**: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

Options

pic-info fpc0.pic0 ***fpc-slot slot number pic-slot slot-number***	Display the information for the FPC in the specified slot.
all	Display all the counters generated during SSL termination.
error	Display all the counters related to errors occurred during SSL termination.

handshake Display all the counters related to handshake during SSL termination.

Required Privilege Level

view

Output Fields

Table 1531 on page 26263 lists the output fields for the `show services ssl termination counters` command. Output fields are listed in the approximate order in which they appear.

Table 1531: show services ssl termination counters Fields

Field Name	Field Description	Display Level
`Memory errors`	Errors related to memory allocation.	all, errors
`Handshake errors`	Number of errors occurred during handshake.	all, errors
`Cert Cache errors`	Number of certificate cache errors.	all, errors
`Server Protection errors`	Errors occurred during SSL reverse proxy.	all, errors
`Proxy errors`	Errors occurred in SSL proxy sessions.	all, errors
`Crypto errors`	Errors related to Cryptographic modules.	all, errors
`Certificate errors`	Errors related to digital certificates	all, errors
`One-Crypto errors`	Number of one-crypto errors	all, errors
`Async-Crypto errors`	Number of Async-crypto errors	all, errors
`Mirror errors`	Errors in SSL decryption mirroring	all, errors

Table 1531: show services ssl termination counters Fields *(Continued)*

Field Name	Field Description	Display Level
handshakes started	Number of SSL handshakes started.	all, handshake
handshakes completed	Number of SSL handshakes completed successfully.	all, handshake
active sessions	Number of active SSL sessions	all, handshake
Interdicted cert generated	Number of interdicted certificates generated	all, handshake
proxy: sessions created	Number of proxy sessions created	all, handshake
proxy: sessions active	Number of active proxy sessions	all, handshake
proxy: sessions ignored	Number of proxy sessions ignored.	all, handshake
proxy: renegotiation ignored	Number of renegotiation requests ignored.	all, handshake
proxy: session resumption	Number of session resumption requests	all, handshake
proxy: secure renegotiation	Number of SSL sessions with secure renegotiation	all, handshake
proxy: insecure renegotiation	Number of SSL sessions with insecure renegotiation	all, handshake
proxy: multiple renegotiation	Number os SSL sessions with multiple renegotiation	all, handshake

Table 1531: show services ssl termination counters Fields *(Continued)*

Field Name	Field Description	Display Level
`proxy: reneg after resumption`	Number os SSL sessions undergo renegotiation after resumption	all, handshake
`init: passthrough requests`	Passthrough requests during initiation	all, handshake
`init: start requests`	Start requests during initiation	all, handshake
`proxy: ECDSA based srvr auth`	Sessions completed ECDSA-based server authentication	all, handshake
`proxy: RSA based srvr auth`	Sessions completed RSA-based server authentication	all, handshake

Sample Output

show services ssl termination counters all

```
user@host > show services ssl termination counters all
Lsys Name : root-logical-system
PIC:fpc0 fpc[0] pic[0] ------

Memory errors                                  0
Handshake errors                               0
Cert Cache errors                              0
Server Protection errors                       0
Proxy errors                                   0
Crypto errors                                  0
Certificate errors                             0
One-Crypto errors                              0
Async-Crypto errors                            0
Mirror errors                                  0
handshakes started                             0
handshakes completed                           0
```

```
active sessions                                0
Interdicted cert generated                     0
proxy: sessions created                        0
proxy: sessions active                         0
proxy: sessions ignored                        0
proxy: renegotiation ignored                   0
proxy: session resumption                      0
proxy: secure renegotiation                    0
proxy: insecure renegotiation                  0
proxy: multiple renegotiation                  0
proxy: reneg after resumption                  0
init: passthrough requests                     0
init: start requests                           0
proxy: ECDSA based srvr auth                   0
proxy: RSA based srvr auth                     0
```

show services ssl termination counters error

```
user@host > show services ssl termination counters errors

Lsys Name : root-logical-system
PIC:fpc0 ------

Memory errors 0
Handshake errors 0
Cert Cache errors 0
Server Protection errors 0
Proxy errors 0
Crypto errors 0
Certificate errors 0
One-Crypto errors 0
Async-Crypto errors 0
Mirror errors 0
```

command-name

```
user@host > show services ssl termination counters handshake

Lsys Name : root-logical-system
PIC:fpc0 fpc[0] pic[0] ------
```

```
handshakes started 0
handshakes completed 0
active sessions 0
Interdicted cert generated 0
proxy: sessions created 0
proxy: sessions active 0
proxy: sessions ignored 0
proxy: renegotiation ignored 0
proxy: session resumption 0
proxy: secure renegotiation 0
proxy: insecure renegotiation 0
proxy: multiple renegotiation 0
proxy: reneg after resumption 0
init: passthrough requests 0
init: start requests 0
proxy: ECDSA based srvr auth 0
proxy: RSA based srvr auth 0
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl termination profile | 26267

show services ssl termination profile

IN THIS SECTION

- Syntax | 26268
- Description | 26268
- Options | 26268

- Required Privilege Level | 26268
- Output Fields | 26269
- Sample Output | 26271
- Release Information | 26273

Syntax

```
show services ssl termination profile   [all | brief | detail]
<pic-info fpc-slot slot number pic-slot slot-number>
```

Description

Display the SSL termination profile details.

> **NOTE**: When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

Options

pic-info ***fpc-slot slot number pic-slot slot-number***	Display the information for the FPC in the specified slot.
all	Display all SSL termination profiles configured on the device.
brief	Display brief information about SSL termination profile.
detail	Display detail information about SSL termination profiles.

Required Privilege Level

view

Output Fields

Table 1532 on page 26269 lists the output fields for the `show services ssl termination profile` command. Output fields are listed in the approximate order in which they appear.

Table 1532: show show services ssl termination profile Output Fields

Field Name	Field Description	Output Levels
`Profile`	SSL termination profile name.	detail
`allow non-ssl session`	Allow or not allow (bypass) non-SSL sessions.	brief, detail
`preferred-ciphers`	SSL cipher that can be used with acceptable key strength. Possible values are strong, medium, weak, and custom.	brief, detail
`Num of url categories configured`	URL categories exempted from SSL proxy.	brief, detail
`Number of whitelist entries`	Allowlisted sessions bypassing SSL proxy processing.	brief
`Protocol-version`	SSL protocol version. Possible values are all, TLS version 1.0, TLS version 1.1, and TLS version 1.2.	detail
`Client authentication`	Status of client certificate verification process.	detail
`Server Authentication`	Status of server certificate verification process.	detail
`Crypto-mode`	Crypto mode used. Options are synchronous-hardware or software or asynchronous-hardware.	detail
`Session Resumption`	SSL session resumption status.	detail
`CRL check`	Status of the CRL checking of certificate validity.	detail
`Certificate`	Types of certificates used.	detail

Table 1532: show show services ssl termination profile Output Fields *(Continued)*

Field Name	Field Description	Output Levels
`Renegotiation`	Renegotiation option. Possible values are allow, allow secure, and drop.	detail
`Custom ciphers`	Custom ciphers configured.	detail
`Server Cert`	Server certificate configured.	detail
`Custom ciphers`	Custom ciphers configured.	detail
`Server Cert`	Server certificate configured.	detail
`Decrypt Mirror`	Status of decrypt mirroring functionality.	detail
`Trusted CA:`	Trusted CA configured for a profile	detail
`Counters`	Details of the counters generated in the session	detail
`handshakes started`	Number of SSL handshakes started.	detail
`handshakes completed`	Number of SSL handshakes completed successfully.	detail
`active sessions`	Number of active SSL sessions	detail
`total handshake errors`	Number of errors occurred during handshake process.	detail
`Data Errors`	Cumulative errors in a single counter. Any errors related to data such as read or write errors.	detail
`session resumption`	Number of SSL session resumption count.	detail

Table 1532: show show services ssl termination profile Output Fields *(Continued)*

Field Name	Field Description	Output Levels
secure renegotiation	Secure sessions allowed after renegotiation.	detail
insecure renegotiation	All sessions allowed after renegotiation.	detail
multiple renegotiation	Sessions with multiple renegotiation.	detail
reneg after resumption	Sessions undergoing renegotiation after resumption.	detail
no_reneg alert by peer	Number of times no renegotiation alerts received from peer.	detail
drop on reneg	Sessions dropped after renegotiation.	detail

Sample Output

show services ssl termination profile all

```
user@host > show services ssl termination profile

Lsys Name : root-logical-system

PIC: fwdd0 fpc[0] pic[0] ----------

ID Name
----------------------------
10 ssl_t
65537 ssl-proxy_65537_proxy_t
```

show services ssl termination profile brief profile-name

```
user@host > show services ssl termination profile brief profile-name

Lsys Name : root-logical-system
```

```
PIC: fwdd0 fpc[0] pic[0] ----------

Profile: ssl-termination
allow non-ssl session: true
preferred-ciphers: medium
Num of url categories configured: NIL
Number of whitelist entries: 0
```

show services ssl termination profile detail profile-name

```
user@host > show services ssl termination profile detail profile-name
Lsys Name : root-logical-system

PIC: fwdd0 fpc[0] pic[0] ----------

Profile                          : root_profile_65536_proxy_t
allow non-ssl session            : true
preferred-ciphers                : medium
Num of url categories configured : 0
Protocol version                 : all
Client Authentication            : notset
Server Authentication            : Required
Crypto Mode                      : hw-sync
Session Resumption               : Enabled
CRL check                        : Enabled
Certficate RSA : p_5
Renegotiation        : disabled
Custom ciphers       : 0
Server cert          : 0
Decrypt Mirror       : Disabled
Trusted CA           : 0
       handshakes started            0
       handshakes completed          0
       active sessions               0
       total handshake errors        0
       Data Errors                   0
       session resumption            0
       secure renegotiation          0
       insecure renegotiation        0
       multiple renegotiation        0
```

```
        reneg after resumption        0
        no_reneg alert by peer        0
        drop on reneg                 0
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Operational Commands to Troubleshoot SSL Sessions

show services ssl termination counters | 26262

show services stateful-firewall conversations

IN THIS SECTION

- Syntax | 26273
- Description | 26274
- Options | 26274
- Required Privilege Level | 26276
- Output Fields | 26276
- Sample Output | 26278
- Release Information | 26279

Syntax

```
show services stateful-firewall conversations
<brief | extensive | terse>
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
```

```
<limit number>
<pgcp>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Description

Display information about stateful firewall conversations.

Options

none

Display standard information about all stateful firewall conversations.

brief | extensive | terse

(Optional) Display the specified level of output.

application-protocol *protocol*

(Optional) Display information about one of the following application protocols:

- `bootp`—Bootstrap protocol
- `dce-rpc`—Distributed Computing Environment-Remote Procedure Call protocols
- `dce-rpc-portmap`—Distributed Computing Environment-Remote Procedure Call protocols portmap service
- `dns`—Domain Name System protocol
- `exec`—Exec
- `ftp`—File Transfer Protocol
- `h323`—H.323 standards
- `icmp`—Internet Control Message Protocol
- `iiop`—Internet Inter-ORB Protocol
- `login`—Login
- `netbios`—NetBIOS

- `netshow`—NetShow
- `realaudio`—RealAudio
- `rpc`—Remote Procedure Call protocol
- `rpc-portmap`—Remote Procedure Call protocol portmap service
- `rtsp`—Real-Time Streaming Protocol
- `shell`—Shell
- `sip`—Session Initiation Protocol
- `snmp`—Simple Network Management Protocol
- `sqlnet`—SQLNet
- `tftp`—Trivial File Transfer Protocol
- `traceroute`—Traceroute
- `winframe`—WinFrame

destination-port *destination-port*

(Optional) Display information for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*

(Optional) Display information for a particular destination prefix.

interface *interface-name*

(Optional) Display information about a particular interface. On M Series and T Series routers, the *interface-name* can be `sp-`*fpc*/*pic*/*port* or `rsp`*number*.

limit *number*

(Optional) Maximum number of entries to display.

pgcp

(Optional) Display information about stateful firewall conversations for Packet Gateway Control Protocol (PGCP) flows.

protocol *protocol*

(Optional) Display information about one of the following IP types:

- *number*—Numeric protocol value from 0 to 255
- `ah`—IPsec Authentication Header protocol
- `egp`—An exterior gateway protocol
- `esp`—IPsec Encapsulating Security Payload protocol

- `gre`—A generic routing encapsulation protocol
- `icmp`—Internet Control Message Protocol
- `igmp`—Internet Group Management Protocol
- `ipip`—IP-within-IP Encapsulation Protocol
- `ospf`—Open Shortest Path First protocol
- `pim`—Protocol Independent Multicast protocol
- `rsvp`—Resource Reservation Protocol
- `sctp`—Stream Control Protocol
- `tcp`—Transmission Control Protocol
- `udp`—User Datagram Protocol

service-set *service-set* (Optional) Display information for the specific service set.

source-port *source-port* (Optional) Display information for a particular source port. The range of values is 0 to 65535.

source-prefix *source-prefix* (Optional) Display information for a particular source prefix.

Required Privilege Level

view

Output Fields

Table 1533 on page 26276 lists the output fields for the `show services stateful-firewall conversations` command. Output fields are listed in the approximate order in which they appear.

Table 1533: show services stateful-firewall conversations Output Fields

Field Name	Field Description
`Interface`	Name of an adaptive services interface.

Table 1533: show services stateful-firewall conversations Output Fields *(Continued)*

Field Name	Field Description
`Service set`	Name of a service set. Individual empty service sets are not displayed, but if no service set has any flows, a flow table header is printed for each service set.
`Conversation`	Information about a group of related flows. • `ALG Protocol`—Application-level gateway protocol. • `Number of initiators`—Number of flows that initiated a session. • `Number of responders`—Number of flows that responded in a session.
`Flow or Flow Prot`	Protocol used for this flow.
`Source`	Source prefix of the flow, in the format *`source-prefix-port`*.
`Destination`	Destination prefix of the flow.
`State`	Status of the flow: • `Drop`—Drop all packets in the flow without response. • `Forward`—Forward the packet in the flow without looking at it. • `Reject`—Drop all packets in the flow with response. • `Watch`—Inspect packets in the flow.
`Dir`	Direction of the flow: input (`I`) or output (`O`).
`Source NAT`	Original and translated source IPv4 or IPv6 addresses are displayed if Network Address Translation (NAT) is configured on this particular flow or conversation.
`Frm Count`	Number of frames in the flow.

Table 1533: show services stateful-firewall conversations Output Fields *(Continued)*

Field Name	Field Description
`Destin NAT`	Original and translated destination IPv4 or IPv6 addresses are displayed if NAT is configured on this particular flow or conversation.
`Byte count`	Number of bytes forwarded in the flow.
`TCP established`	Whether a TCP connection was established: `Yes` or `No`.
`TCP window size`	Negotiated TCP connection window size, in bytes.
`TCP acknowledge`	TCP acknowledgment sequence number.
`TCP tickle`	Whether TCP inquiry mode is on (`enabled` or `disabled`) and the time remaining to send the next inquiry, in seconds.
`Master flow`	Flow that initiated the conversation.
`TImeout`	Lifetime of the flow, in seconds.

Sample Output

show services stateful-firewall conversations

```
user@host> show services stateful-firewall conversations
Interface: sp-1/3/0, Service set: green
Conversation: ALG Protocol: any, Number of initiators: 1,
Number of responders: 1

Flow
Prot     Source                 Dest               State      Dir   Frm count
TCP      10.58.255.50:33005->   10.58.255.178:23   Forward    I      13
  Source NAT     10.58.255.50:33005->    10.59.16.100:4000
  Destin NAT     10.58.255.178:23  ->         0.0.0.0:4000
```

```
Byte count:          918
TCP established, TCP window size: 65535, TCP acknowledge: 2502627025
TCP tickle enabled, 0 seconds,
Master flow, Timeout: 30 seconds
TCP     10.58.255.178:23   ->    10.59.16.100:4000 Forward    O        8
```

show services stateful-firewall conversations destination-port

```
user@host> show services stateful-firewall conversations destination-port 21
Interface: sp-0/3/0, Service set: svc_set_trust

Interface: sp-0/3/0, Service set: svc_set_untrust
Conversation: ALG protocol: ftp
  Number of initiators: 1, Number of responders: 1
Flow                                                State    Dir       Frm count
TCP         10.50.10.2:2143  ->     10.50.20.2:21    Watch    O               0
TCP         10.50.20.2:21    ->     10.50.10.2:2143  Watch    I               0
TCP         10.50.20.2:21    ->     10.50.10.2:2143  Watch    I               0
```

Release Information

Command introduced before Junos OS Release 7.4.

`pgcp` option introduced in Junos OS Release 8.4.

show services stateful-firewall flow-analysis

IN THIS SECTION

- Syntax | 26280
- Description | 26280
- Options | 26280
- Required Privilege Level | 26280
- Output Fields | 26280
- Sample Output | 26282

Sample Output | 26284
Release Information | 26285

Syntax

```
show services stateful-firewall flow-analysis
<interface interface-name>
```

Description

Display stateful firewall flow statistics.

Options

none	Display standard information about all stateful firewall flow statistics.
interface *interface-name*	(Optional) Display information about a particular interface.

Required Privilege Level

view

Output Fields

Table 1534 on page 26280 lists the output fields for the `show services stateful-firewall flow-analysis` command. Output fields are listed in the approximate order in which they appear.

Table 1534: show services stateful-firewall flow-analysis Output Fields

Field Name	Field Description
Total Flows Active	Total active flows in the MS-PIC including TCP, UDP, ICMP and Softwires.
Total TCP Flows Active	Total active TCP flows in the MS-PIC.

Table 1534: show services stateful-firewall flow-analysis Output Fields *(Continued)*

Field Name	Field Description
Total UDP Flows Active	Total active UDP flows in the MS-PIC.
Total Other Flows Active	Total other active flows in the MS-PIC including ICMP and softwires.
Total Predicted Flows Active	Predicted flows are created only by the ALG traffic using the L3/L4 information available.
Created Flows per Second	Flow setup rate at the time of running the command.
Deleted Flows per Second	Flow deletion rate at the time of running the command.
Peak Total Flows Active	The highest number of active flows since the last PIC restart or since the last time flow statistics are flushed.
Peak Total TCP Flows Active	The highest number of active TCP flows since the last PIC restart or since the last time flow stats are flushed.
Peak Total UDP Flows Active	The highest number of active UDP flows since the last PIC restart or since the last time flow statistics are flushed.
Peak Total Other Flows Active	The highest number of other active flows since the last PIC restart or since the last time flow statistics are flushed.
Peak Created Flows per Second	The maximum flow setup rate observed since the last PIC restart or since the last time flow statistics are flushed.
Peak Deleted Flows per Second	The maximum flow deletion rate observed since the last PIC restart or from the last time flow statistics are flushed.
Average HTTP Flow Lifetime(ms)	Average HTTP Flow Lifetime in millisecond.
Packets received	The total number of packets received by the MS-PIC.

Table 1534: show services stateful-firewall flow-analysis Output Fields *(Continued)*

Field Name	Field Description
Packets transmitted	The total number of packets transmitted by the MS-PIC.
Slow path forward	The number of packets forwarded in the slow path (i.e. after the successful rule match and flow creation).
Slow path discard	The number of packets discarded before the flow creation.
Flow Rate Data: Number of Samples	The number of samples used to calculate the flow rate, since the last PIC restart or since the last time flow statistics are flushed.
Flow Rate Distribution(sec) Flow Operation :Creation Flow Operation :Deletion	Histogram of the samples used for flow rate calculation.
Flow Lifetime Distribution(sec):	Histogram of the samples used to calculate the flow life time in sec.

Sample Output

show services stateful-firewall flow-analysis

```
user@host> show services stateful-firewall flow-analysis
Services PIC Name: sp-3/0/0
 Flow Analysis Statistics:
      Total Flows Active               :40
      Total TCP Flows Active           :0
      Total UDP Flows Active           :40
      Total Other Flows Active         :0
      Total Predicted Flows Active     :0
      Created Flows per Second         :0
      Deleted Flows per Second         :0
      Peak Total Flows Active          :40
      Peak Total TCP Flows Active      :0
      Peak Total UDP Flows Active      :40
```

```
     Peak Total Other Flows Active    :0
     Peak Created Flows per Second    :20
     Peak Deleted Flows per Second    :20
     Average HTTP Flow Lifetime(ms)   :0
     Packets received                 :48682539117
     Packets transmitted              :48682502703
     Slow path forward                :6550
     Slow path discard                :0
Flow Rate Data:
     Number of Samples: 19720
Flow Rate Distribution(sec)
Flow Operation :Creation
     300000+            :0
     250000  - 300000   :0
     200000  - 250000   :0
     160000  - 200000   :0
     150000  - 160000   :0
     50000  - 150000   :0
     40000  -  50000   :0
     30000  -  40000   :0
     20000  -  30000   :0
     10000  -  20000   :0
     1000  -  10000   :0
     0  -   1000   :19720
Flow Operation :Deletion
     300000+            :0
     250000  - 300000   :0
     200000  - 250000   :0
     160000  - 200000   :0
     150000  - 160000   :0
     50000  - 150000   :0
     40000  -  50000   :0
     30000  -  40000   :0
     20000  -  30000   :0
     10000  -  20000   :0
     1000  -  10000   :0
     0  -   1000   :19720
Flow Lifetime Distribution(sec):
                TCP                UDP              HTTP
 240+         :0                 0                0
 120 - 240    :0                 0
  60 - 120    :0                 0
  30 - 60     :0                 0
```

```
   15 - 30     :0                  6530
    5 - 15     :0                  0
    1 - 5      :0                  0
    0 - 1      :0                  6530
```

Sample Output

show services stateful-firewall flow-analysis interface sp-3/0/0

```
user@host> show services stateful-firewall flow-analysis interface sp-3/0/0
Services PIC Name: sp-3/0/0
 Flow Analysis Statistics:
      Total Flows Active               :40
      Total TCP Flows Active           :0
      Total UDP Flows Active           :40
      Total Other Flows Active         :0
      Total Predicted Flows Active     :0
      Created Flows per Second         :0
      Deleted Flows per Second         :0
      Peak Total Flows Active          :40
      Peak Total TCP Flows Active      :0
      Peak Total UDP Flows Active      :40
      Peak Total Other Flows Active    :0
      Peak Created Flows per Second    :20
      Peak Deleted Flows per Second    :20
      Average HTTP Flow Lifetime(ms)   :0
      Packets received                 :54696856768
      Packets transmitted              :54696815873
      Slow path forward                :7350
      Slow path discard                :0
 Flow Rate Data:
      Number of Samples: 22139
 Flow Rate Distribution(sec)
 Flow Operation :Creation
      300000+            :0
      250000  - 300000   :0
      200000  - 250000   :0
      160000  - 200000   :0
      150000  - 160000   :0
      50000  - 150000   :0
      40000  -  50000   :0
```

```
      30000  -  40000   :0
      20000  -  30000   :0
      10000  -  20000   :0
      1000  -  10000   :0
      0  -   1000   :22139
  Flow Operation :Deletion
      300000+            :0
      250000  - 300000   :0
      200000  - 250000   :0
      160000  - 200000   :0
      150000  - 160000   :0
      50000  - 150000   :0
      40000  -  50000   :0
      30000  -  40000   :0
      20000  -  30000   :0
      10000  -  20000   :0
      1000  -  10000   :0
      0  -   1000   :22139
  Flow Lifetime Distribution(sec):
                 TCP                UDP               HTTP
   240+         :0                 0                 0
   120 - 240    :0                 0
    60 - 120    :0                 0
    30 - 60     :0                 0
    15 - 30     :0                 7330
     5 - 15     :0                 0
     1 - 5      :0                 0
     0 - 1      :0                 7330
```

Release Information

Command introduced in Junos OS Release 10.4R1.

show services stateful-firewall flows

IN THIS SECTION

- Syntax | 26286
- Description | 26286
- Options | 26287
- Required Privilege Level | 26289
- Output Fields | 26289
- Sample Output | 26290
- Release Information | 26293

Syntax

```
show services stateful-firewall flows
<brief | extensive | summary | terse>
<application-protocol protocol>
<count>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Description

Display stateful firewall flow table entries. When the interface is used for softwire processing, the type of softwire concentrator (`DS-LITE` or `6rd`) is shown, and frame counts are provided.

Options

none

Display standard information about all stateful firewall flows.

brief | extensive | summary | terse

(Optional) Display the specified level of output.

application-protocol *application-protocol*

(Optional) Display information about one of the following application-level gateway (ALG) protocol types:

- `bootp`—Bootstrap protocol
- `dce-rpc`—Distributed Computing Environment (DCE) remote procedure call (RPC) protocol

  > **NOTE**: Use this option to select Microsoft Remote Procedure Call (MSRPC).

- `dce-rpc-portmap`—Distributed Computing Environment (DCE) remote procedure call (RPC) portmap protocol
- `dns`—Domain Name Service protocol
- `exec`—Remote execution protocol
- `ftp`—File Transfer Protocol
- `h323`—H.323 protocol
- `icmp`—Internet Control Message Protocol
- `iiop`—Internet Inter-ORB Protocol
- `ip`—Internet protocol
- `netbios`—NetBIOS protocol
- `netshow`—Netshow protocol
- `pptp` —Point-to-Point Tunneling Protocol
- `realaudio`—RealAudio protocol
- `rpc`—Remote Procedure Call protocol

> **NOTE**: Use this option to select Sun Microsystems Remote Procedure Call protocol (SunRPC).

- `rpc-portmap`—Remote Procedure Call portmap protocol
- `rtsp`—Real-Time Streaming Protocol
- `sip`—Session Initiation Protocol
- `snmp`—Simple Network Management Protocol
- `talk`—Talk protocol
- `tftp`—Trivial File Transfer Protocol
- `traceroute`—Traceroute
- `winframe`—WinFrame

count

(Optional) Display a count of the matching entries.

destination-port *destination-port*

(Optional) Display information for a particular destination port. The range of values is from 0 to 65535.

destination-prefix *destination-prefix*

(Optional) Display information for a particular destination prefix.

interface *interface-name*

(Optional) Display information about a particular interface. On M Series and T Series routers, *`interface-name`* can be `ms-`*`fpc`*`/`*`pic`*`/`*`port`* or `rsp`*`number`*.

limit *number*

(Optional) Maximum number of entries to display.

protocol *protocol*

(Optional) Display information about one of the following IP types:

- *`number`*—Numeric protocol value from 0 to 255
- `ah`—IPsec Authentication Header protocol
- `egp`—An exterior gateway protocol
- `esp`—IPsec Encapsulating Security Payload protocol
- `gre`—A generic routing encapsulation protocol
- `icmp`—Internet Control Message Protocol

- `igmp`—Internet Group Management Protocol
- `ipip`—IP-within-IP Encapsulation Protocol
- `ospf`—Open Shortest Path First protocol
- `pim`—Protocol Independent Multicast protocol
- `rsvp`—Resource Reservation Protocol
- `sctp`—Stream Control Protocol
- `tcp`—Transmission Control Protocol
- `udp`—User Datagram Protocol

service-set *service-set* — (Optional) Display information for a particular service set.

source-port *source-port* — (Optional) Display information for a particular source port. The range of values is from 0 to 65535.

source-prefix *source-prefix* — (Optional) Display information for a particular source prefix.

Required Privilege Level

view

Output Fields

Table 1535 on page 26289 lists the output fields for the `show services stateful-firewall flows` command. Output fields are listed in the approximate order in which they appear.

Table 1535: show services stateful-firewall flows Output Fields

Field Name	Field Description
`Interface`	Name of the interface.
`Service set`	Name of a service set. Individual empty service sets are not displayed. If no service set has any flows, a flow table header is displayed for each service set.

Table 1535: show services stateful-firewall flows Output Fields *(Continued)*

Field Name	Field Description
`Flow Count`	Number of flows in a session.
`Flow or Flow Prot`	Protocol used for this flow.
`Source`	Source prefix of the flow in the format *`source-prefix:port`*. For ICMP flows, port information is not displayed.
`Dest`	Destination prefix of the flow. For ICMP flows, port information is not displayed.
`State`	Status of the flow: • `Drop`—Drop all packets in the flow without response. • `Forward`—Forward the packet in the flow without looking at it. • `Reject`—Drop all packets in the flow with response. • `Watch`—Inspect packets in the flow.
`Dir`	Direction of the flow: input (`I`) or output (`O`). For any configured stateful firewall rule, the reverse flow is dynamically created, so you will see an input and an output flow.
`Frm count`	Number of frames in the flow. If this value is zero, then that flow does not yet exist.

Sample Output

show services stateful-firewall flows

On the MX Series router, both input (I) and output (O) flow entries appear, even if traffic only flows in one direction. This applies to both NAT and non-NAT cases.

```
user@host> show services stateful-firewall flows
Interface: ms-1/3/0, Service set: green
```

```
Flow
Prot     Source                 Dest                State      Dir     Frm count
TCP     10.58.255.178:23   ->    10.59.16.100:4000 Forward    O
TCP      10.58.255.50:33005->   10.58.255.178:23   Forward    I              1
  Source NAT    10.58.255.50:33005->    10.59.16.100:4000
  Destin NAT    10.58.255.178:23   ->         0.0.0.0:4000
```

show services stateful-firewall flows (For Softwire Flows)

When a service set includes softwire processing, the following output format is used for the softwire flows:

```
user@host> show services stateful-firewall flows
Interface: sp-0/1/0, Service set: dslite-svc-set2
Flow                                                State    Dir       Frm count
TCP      200.200.200.2:80    ->     44.44.44.1:1025  Forward  O          219942
    NAT dest        44.44.44.1:1025    ->       20.20.1.4:1025
    Softwire           2001::2         ->         1001::1
TCP          20.20.1.2:1025  ->  200.200.200.2:80    Forward  I          110244
    NAT source       20.20.1.2:1025    ->      44.44.44.1:1024
    Softwire           2001::2         ->         1001::1
TCP      200.200.200.2:80    ->     44.44.44.1:1024  Forward  O          219140
    NAT dest        44.44.44.1:1024    ->       20.20.1.2:1025
    Softwire           2001::2         ->         1001::1
DS-LITE         2001::2      ->        1001::1       Forward  I          988729
TCP      200.200.200.2:80    ->     44.44.44.1:1026  Forward  O          218906
    NAT dest        44.44.44.1:1026    ->       20.20.1.3:1025
    Softwire           2001::2         ->         1001::1
TCP          20.20.1.3:1025  ->  200.200.200.2:80    Forward  I          110303
    NAT source       20.20.1.3:1025    ->      44.44.44.1:1026
    Softwire           2001::2         ->         1001::1
TCP          20.20.1.4:1025  ->  200.200.200.2:80    Forward  I          110944
    NAT source       20.20.1.4:1025    ->      44.44.44.1:1025
    Softwire           2001::2         ->         1001::1
```

show services stateful-firewall flows brief

The output for the `show services stateful-firewall flows brief` command is identical to that for the `show services stateful-firewall flows` command. For sample output, see "show services stateful-firewall flows" on page 26286.

show services stateful-firewall flows extensive

```
user@host> show services stateful-firewall flows extensive
Interface: ms-0/3/0, Service set: ss_nat
Flow                                                                State        Dir       Frm
count
TCP           16.1.0.1:2330  ->      16.49.0.1:21                    Forward
I              8
    NAT source        16.1.0.1:2330    ->       16.41.0.1:2330
    NAT dest         16.49.0.1:21      ->       16.99.0.1:21
  Byte count: 455, TCP established, TCP window size: 57344
  TCP acknowledge: 3251737524, TCP tickle enabled, tcp_tickle: 0
  Flow role: Master, Timeout: 720
TCP          16.99.0.1:21    ->      16.41.0.1:2330                  Forward
O              5
    NAT source       16.99.0.1:21      ->       16.49.0.1:21
    NAT dest         16.41.0.1:2330    ->        16.1.0.1:2330
  Byte count: 480, TCP established, TCP window size: 57344
  TCP acknowledge: 463128048, TCP tickle enabled, tcp_tickle: 0
  Flow role: Responder, Timeout: 720
```

show services stateful-firewall flows count

```
user@host> show services stateful-firewall flows count
Interface             Service set                                   Flow Count

ms-1/3/0              green                                                2
```

show services stateful-firewall flows destination port

```
user@host> show services stateful-firewall flows destination-port 21
Interface: ms-0/3/0, Service set: svc_set_trust
Flow
                                                  State    Dir       Frm count
Interface: ms-0/3/0, Service set: svc_set_untrust
Flow                                                 State    Dir       Frm count
TCP         10.50.10.2:2143  ->     10.50.20.2:21    Watch    O               0
```

show services stateful-firewall flows source port

```
user@host> show services stateful-firewall flows source-port 2143
Interface: ms-0/3/0, Service set: svc_set_trust
Flow
                                                State    Dir       Frm count
Interface: ms-0/3/0, Service set: svc_set_untrust
Flow                                                  State    Dir       Frm count
TCP         10.50.10.2:2143  ->     10.50.20.2:21    Watch    O               0
```

show services stateful-firewall flows (Twice NAT)

```
user@host> show services stateful-firewall flows
Flow                                                 State    Dir       Frm count
UDP          40.0.0.8:23439 ->     80.0.0.1:16485   Watch    I             20
    NAT source        40.0.0.8:23439   ->     172.16.1.10:1028
    NAT dest          80.0.0,1:16485   ->     192.16.1.10:22415
UDP      192.16.1.10:22415  ->  172.16.1.10:1028    Watch    O             20
    NAT source     192.16.1.10:22415   ->        80.0.0.1:16485
    NAT dest       172.16.1.10:1028    ->        40.0.0.8:23439
```

Release Information

Command introduced before Junos OS Release 7.4.

`pgcp` option introduced in Junos OS Release 8.4.

`application-protocol` option introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

clear services stateful-firewall flows | 4180

show services stateful-firewall sip-call

IN THIS SECTION

- Syntax | 26294
- Description | 26294
- Options | 26294
- Required Privilege Level | 26297
- Output Fields | 26297
- Sample Output | 26299
- Release Information | 26300

Syntax

```
show services stateful-firewall sip-call
<brief | extensive | terse>
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Description

Display stateful firewall Session Initiation Protocol (SIP) call information.

Options

count (Optional) Display a count of the matching entries.

brief (Optional) Display brief SIP call information.

extensive (Optional) Display detailed SIP call information.

terse (Optional) Display terse SIP call information.

application-protocol (Optional) Display information about one of the following application protocols:

- `bootp`—(SIP only) Bootstrap protocol
- `dce-rpc`—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- `dce-rpc-portmap`—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- `dns`—(SIP only) Domain Name System protocol
- `exec`—(SIP only) Exec
- `ftp`—(SIP only) File Transfer Protocol
- `h323`—H.323 standards
- `icmp`—Internet Control Message Protocol
- `iiop`—Internet Inter-ORB Protocol
- `login`—Login
- `netbios`—NetBIOS
- `netshow`—NetShow
- `realaudio`—RealAudio
- `rpc`—Remote Procedure Call protocol
- `rpc-portmap`—Remote Procedure Call protocol portmap service
- `rtsp`—Real-Time Streaming Protocol
- `shell`—Shell
- `sip`—Session Initiation Protocol
- `snmp`—Simple Network Management Protocol

- `sqlnet`—SQLNet
- `tftp`—Trivial File Transfer Protocol
- `traceroute`—Traceroute
- `winframe`—WinFrame

destination-port *destination-port*
(Optional) Display information for a particular destination port. The range of values is from 0 to 65535.

destination-prefix *destination-prefix*
(Optional) Display information for a particular destination prefix.

interface *interface-name*
(Optional) Display information about a particular adaptive services interface. On M Series and T Series routers, *`interface-name`* can be `sp-`*`fpc/pic/port`* or `rsp`*`number`*.

limit *number*
(Optional) Maximum number of entries to display.

protocol
(Optional) Display information about one of the following IP types:

- `ah`—IPsec Authentication Header protocol
- `egp`—An exterior gateway protocol
- `esp`—IPsec Encapsulating Security Payload protocol
- `gre`—A generic routing encapsulation protocol
- `icmp`—Internet Control Message Protocol
- `igmp`—Internet Group Management Protocol
- `ipip`—IP-within-IP Encapsulation Protocol
- `ipv6`—IPv6 within IP
- `ospf`—Open Shortest Path First protocol
- `pim`—Protocol Independent Multicast protocol
- `rsvp`—Resource Reservation Protocol
- `sctp`—Stream Control Protocol
- `tcp`—Transmission Control Protocol
- `udp`—User Datagram Protocol

service-set *service-set*	(Optional) Display information for a particular service set.
source-port *source-port*	(Optional) Display information for a particular source port. The range of values is from 0 to 65535.
source-prefix *source-prefix*	(Optional) Display information for a particular source prefix.

Required Privilege Level

view

Output Fields

Table 1536 on page 26297 lists the output fields for the `show services stateful-firewall sip-call` command. Output fields are listed in the approximate order in which they appear.

Table 1536: show services stateful-firewall sip-call Output Fields

Field Name	Field Description
`Interface`	Name of an adaptive services interface.
`Service set`	Name of a service set.
`From`	Initiator address.
`To`	Responder address.
`Call ID`	SIP call identification string.
`Number of initiator flows`	Number of `control`, `contact`, or `media` initiator flows.
`Number of responder flows`	Number of `control`, `contact`, or `media` responder flows.

Table 1536: show services stateful-firewall sip-call Output Fields *(Continued)*

Field Name	Field Description
protocol	Protocol used for this flow.
source-prefix	Source prefix of the flow in the format *source-prefix* : *port*.
destination-prefix	Destination prefix of the flow.
state	Status of the flow: • Drop—Drop all packets in the flow without a response. • Forward—Forward the packet in the flow without examining it. • Reject—Drop all packets in the flow with a response. • Unknown—Unknown status. • Watch—Inspect packets in the flow.
direction	Direction of the flow: input (I), output (O), or unknown (U).
frame-count	Number of frames in the flow.
Byte count	Number of bytes forwarded in the flow.
Flow role	Role of the flow that is under evaluation: Initiator, Master, Responder, or Unknown.
Timeout	Lifetime of the flow, in seconds.

Sample Output

show services stateful-firewall sip-call extensive

```
user@host> show services stateful-firewall sip-call extensive
Interface: sp-0/3/0, Service set: test_sip_777

From: : 6507771234@10.200.100.1:0;000ff73ac89900021bb231dc-3ef68435
To: : 4085551234@10.200.100.1:0;0011bb65c2a30007777bd0fc-5748b749
Call ID: : 000ff73a-c8990004-0741adac-3e027c7e@10.20.70.2
Number of control initiator flows: : 1, Number of control responder flows:
: 1
UDP         10.20.70.2:50354 ->   10.200.100.1:5060  Watch    I
2
  Byte count: 1112
  Flow role: Master, Timeout: 30
UDP       10.200.100.1:5060  ->  10.20.170.111:50354 Watch    O
0
  Byte count: 0
  Flow role: Responder, Timeout: 30
UDP            0.0.0.0:0     ->  10.20.170.111:5060  Watch    O
7
  Byte count: 2749
  Flow role: Responder, Timeout: 30
Number of contact initiator flows: 1, Number of contact responder flows: 1
UDP            0.0.0.0:0     ->   10.20.140.11:5060  Watch    I
1
  Byte count: 409
  Flow role: Master, Timeout: 30
UDP       10.20.140.11:31864 ->  10.20.170.111:18808 Forward  O
622
  Byte count: 124400
  Flow role: Master, Timeout: 30
UDP            0.0.0.0:0     ->  10.20.170.111:18809 Forward  O
0
  Byte count: 0
  Flow role: Initiator, Timeout: 30
Number of media initiator flows: 4, Number of media responder flows: 0
UDP         10.20.70.2:18808 ->   10.20.140.11:31864 Forward  I
628
  Byte count: 125600
  Flow role: Initiator, Timeout: 30
```

```
UDP            0.0.0.0:0     ->   10.20.140.11:31865 Forward  I
0
  Byte count: 0
  Flow role: Initiator, Timeout: 30
0              0.0.0.0:0     ->        0.0.0.0:0     Unknown  U
0
  Byte count: 0
  Flow role: Unknown, Timeout: 0
0              0.0.0.0:0     ->        0.0.0.0:0     Unknown  U
Interface: sp-0/3/0, Service set: test_sip_888
```

Release Information

Command introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

clear services stateful-firewall sip-call | 4183

show services stateful-firewall sip-register

IN THIS SECTION

- Syntax | 26301
- Description | 26301
- Options | 26301
- Required Privilege Level | 26303
- Output Fields | 26303
- Sample Output | 26304
- Release Information | 26305

Syntax

```
show services stateful-firewall sip-register
<brief | extensive | terse>
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Description

Display stateful firewall Session Initiation Protocol (SIP) register information.

Options

count (Optional) Display a count of the matching entries.

brief (Optional) Display brief SIP register information.

extensive (Optional) Display detailed SIP register information.

terse (Optional) Display terse SIP register information.

application-protocol (Optional) Display information about one of the following application protocols:

- `bootp`—(SIP only) Bootstrap protocol
- `dce-rpc`—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- `dce-rpc-portmap`—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- `dns`—(SIP only) Domain Name System protocol

- `exec`—(SIP only) Exec
- `ftp`—(SIP only) File Transfer Protocol
- `h323`—H.323 standards
- `icmp`—Internet Control Message Protocol
- `iiop`—Internet Inter-ORB Protocol
- `login`—Login
- `netbios`—NetBIOS
- `netshow`—NetShow
- `realaudio`—RealAudio
- `rpc`—Remote Procedure Call protocol
- `rpc-portmap`—Remote Procedure Call protocol portmap service
- `rtsp`—Real-Time Streaming Protocol
- `shell`—Shell
- `sip`—Session Initiation Protocol
- `snmp`—Simple Network Management Protocol
- `sqlnet`—SQLNet
- `tftp`—Trivial File Transfer Protocol
- `traceroute`—Traceroute
- `winframe`—WinFrame

destination-port *destination-port*
(Optional) Display information for a particular destination port.

destination-prefix *destination-prefix*
(Optional) Display information for a particular destination prefix. The range of values is from 0 to 65535.

interface *interface-name*
(Optional) Display information about a particular interface. On M Series and T Series routers, the *interface-name* can be `sp-`*fpc*/*pic*/*port* or `rsp`*number*.

limit *number*
(Optional) Maximum number of entries to display.

protocol
(Optional) Display information about one of the following IP types:

- `ah`—IPsec Authentication Header protocol
- `egp`—An exterior gateway protocol
- `esp`—IPsec Encapsulating Security Payload protocol
- `gre`—A generic routing encapsulation protocol
- `icmp`—Internet Control Message Protocol
- `igmp`—Internet Group Management Protocol
- `ipip`—IP-within-IP Encapsulation Protocol
- `ipv6`—IPv6 within IP
- `ospf`—Open Shortest Path First protocol
- `pim`—Protocol Independent Multicast protocol
- `rsvp`—Resource Reservation Protocol
- `sctp`—Stream Control Protocol
- `tcp`—Transmission Control Protocol
- `udp`—User Datagram Protocol

service-set *service-set* (Optional) Display information for a particular service set.

source-port *source-port* (Optional) Display information for a particular source port. The range of values is from 0 to 65535.

source-prefix *source-prefix* (Optional) Display information for a particular source prefix.

Required Privilege Level

view

Output Fields

Table 1537 on page 26304 lists the output fields for the `show services stateful-firewall sip-register` command. Output fields are listed in the approximate order in which they appear.

Table 1537: show services stateful-firewall sip-register Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of a service set.
SIP Register	Register information header.
Protocol	Protocol used for this flow.
Registered IP	Register IP address.
Port	Register port number.
Expiration timeout	Configured lifetime, in seconds.
Timeout remaining	Lifetime remaining, in seconds.
From	Initiator address.
To	Responder address.
Call ID	SIP call identification string.

Sample Output

show services stateful-firewall sip-register extensive

```
user@host> show services stateful-firewall sip-register extensive
Interface: sp-0/3/0, Service set: test_sip_777

SIP Register: Protocol: UDP, Registered IP: 10.20.170.111, Port: 5060, Acked
```

```
    Expiration timeout: 36000, Timeout remaining: 35544
    From: : 6507771234@10.200.100.1:0;
    To: : 6507771234@10.200.100.1:0;
    Call ID: : 000ff73a-c8990002-23b1d942-2ba1f91f@10.20.70.2

Interface: sp-0/3/0, Service set: test_sip_888

SIP Register: Protocol: UDP, Registered IP: 10.20.170.112, Port: 5060, Acked
    Expiration timeout: 36000, Timeout remaining: 35549
    From: : 8881234@10.200.100.1:0;
    To: : 8881234@10.200.100.1:0;
    Call ID: : 00112096-81fc0002-23b38905-7cb41f62@10.20.71.2
```

Release Information

Command introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

clear services stateful-firewall sip-register | 4187

show services stateful-firewall statistics

IN THIS SECTION

- Syntax | 26306
- Description | 26306
- Options | 26306
- Required Privilege Level | 26306
- Output Fields | 26306
- Sample Output | 26315
- Release Information | 26317

Syntax

```
show services stateful-firewall statistics
<application-protocol protocol>
<brief | detail | extensive | summary>
<interface interface-name>
<service-set service-set>
```

Description

Display stateful firewall statistics.

Options

none	Display standard information about all stateful firewall statistics.
brief | detail | extensive | summary	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display information about a particular interface. On M Series and T Series routers, the *interface-name* can be ms-*fpc*/*pic*/*port* or rsp*number*.
service-set *service-set*	(Optional) Display information about a particular service set.

Required Privilege Level

view

Output Fields

Table 1538 on page 26306 lists the output fields for the `show services stateful-firewall statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1538: show services stateful-firewall statistics Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.

Table 1538: show services stateful-firewall statistics Output Fields *(Continued)*

Field Name	Field Description
`Service set`	Name of a service set.
`New flows`	Rule match counters for new flows: • `Rule Accepts`—New flows accepted. • `Rule Discards`—New flows discarded. • `Rule Rejects`—New flows rejected.
`Existing flow types packet counters`	Rule match counters for existing flows: • `Accepts`—Match existing forward or watch flow. • `Drop`—Match existing discard flow. • `Rejects`—Match existing reject flow.
`Hairpinning Counters`	Hairpinning counters: • `Slow Path Hairpinned Packets`—Slow path packets that were hairpinned back to the internal network. • `Fast Path Hairpinned Packets`—Fast path packets that were hairpinned back to the internal network.
`Drops`	Drop counters: • `IP option`—Packets dropped in IP options processing. • `TCP SYN defense`—Packets dropped by SYN defender. • `NAT ports exhausted`—Hide mode. The router has no available Network Address Translation (NAT) ports for a given address or pool. • `Sessions dropped due to subscriber flow limit`—Sessions dropped because the subscriber's flow limit was exceeded.

Table 1538: show services stateful-firewall statistics Output Fields *(Continued)*

Field Name	Field Description
`Errors`	Total errors, categorized by protocol: • `IP`—Total IP version 4 errors. • `TCP`—Total Transmission Control Protocol (TCP) errors. • `UDP`—Total User Datagram Protocol (UDP) errors. • `ICMP`—Total Internet Control Message Protocol (ICMP) errors. • `Non-IP packets`—Total non-IPv4 errors. • `ALG`—Total application-level gateway (ALG) errors

Table 1538: show services stateful-firewall statistics Output Fields *(Continued)*

Field Name	Field Description
`IP Errors`	IPv4 errors: • `IP packet length inconsistencies`—IP packet length does not match the Layer 2 reported length. • `Minimum IP header length check failures`—Minimum IP header length is 20 bytes. The received packet contains less than 20 bytes. • `Reassembled packet exceeds maximum IP length`—After fragment reassembly, the reassembled IP packet length exceeds 65,535. • `Illegal source address 0`—Source address is not a valid address. Invalid addresses are, loopback, broadcast, multicast, and reserved addresses. Source address `0`, however, is allowed to support BOOTP and the destination address `0xffffffff`. • `Illegal destination address 0`—Destination address is not a valid address. The address is reserved. • `TTL zero errors`—Received packet had a time-to-live (TTL) value of 0. • `Illegal IP protocol number (0 or 255)`—IP protocol is 0 or 255. • `Land attack`—IP source address is the same as the destination address. • `Non-IPv4 packets`—Packet was not IPv4. (Only IPv4 is supported.) • `Bad checksum`—Packet had an invalid IP checksum. • `Illegal IP fragment length`—Illegal fragment length. All fragments (other than the last fragment) must have a length that is a multiple of 8 bytes. • `IP fragment overlap`—Fragments have overlapping fragment offsets. • `IP fragment reassembly timeout`—Some of the fragments for an IP packet were not received in time, and the reassembly handler dropped partial fragments. • `IP fragment limit exceeded: 0`—Fragments that exceeded the limit. • `Unknown: 0`—Unknown fragments.

Table 1538: show services stateful-firewall statistics Output Fields *(Continued)*

Field Name	Field Description
`TCP Errors`	TCP protocol errors: • `TCP header length inconsistencies`—Minimum TCP header length is 20 bytes, and the IP packet received does not contain at least 20 bytes. • `Source or destination port number is zero`—TCP source or destination port is zero. • `Illegal sequence number and flags combinations` — Dropped because of TCP errors, such as an illegal sequence number, which causes an illogical combination of flags to be set. • `SYN attack (multiple SYN messages seen for the same flow)`—Multiple SYN packets received for the same flow are treated as a SYN attack. The packets might be retransmitted SYN packets and therefore valid, but a large number is cause for concern. • `First packet not a SYN message`—First packets for a connection are not SYN packets. These packets might originate from previous connections or from someone performing an ACK/FIN scan. • `TCP port scan (TCP handshake, RST seen from server for SYN)`—In the case of a SYN defender, if an RST (reset) packet is received instead of a SYN/ACK message, someone is probably trying to scan the server. This behavior can result in false alarms if the RST packet is not combined with an intrusion detection service (IDS). • `Bad SYN cookie response`—SYN cookie generates a SYN/ACK message for all incoming SYN packets. If the ACK received for the SYN/ACK message does not match, this counter is incremented. • `TCP reconstructor sequence number error`—This counter is incremented in the following cases: The TCP seqno is 0 and all the TCP flags are also 0. The TCP seqno is 0 and FIN/PSH/URG TCP flags are set. • `TCP reconstructor retransmissions`—This counter is incremented for the retransmitted packets during connection 3-way handshake. • `TCP partially opened connection timeout (SYN)`—This counter is incremented when the SYN Defender is enabled and the 3-way handshake is not completed within the SYN DEFENDER TIMEOUT. The connection will be closed and resources will be released by sending RST to the responder.

Table 1538: show services stateful-firewall statistics Output Fields *(Continued)*

Field Name	Field Description
	• `TCP partially opened connection timeout (SYN-ACK)`—This counter is incremented when the SYN Defender is enabled and the 3-way handshake is not completed within the SYN DEFENDER TIMEOUT. The connection will be closed and resources will be released by sending RST to the responder. • `TCP partially closed connection reuse`—Not supported. • `TCP 3-way error - client sent SYN+ACK`—A SYN/ACK should be sent by the server on receiving a SYN. This counter is incremented when the first message received from the initiator is SYN+ACK. • `TCP 3-way error - server sent ACK`—ACK should be sent by the client on receiving a SYN/ACK from the server. This counter is incremented when the ACK is received from the Server instead of from the Client. • `TCP 3-way error - SYN seq number retransmission mismatch`—This counter is incremented when the SYN is received again with a different sequence number from the first SYN sequence number. • `TCP 3-way error - RST seq number mismatch`—A reset could be received from either side. The server could send a RST on receiving a SYN or the client could send a RST on receiving SYN/ACK. This counter is incremented when the RST is received either from the client or server with a non-matching sequence number. • `TCP 3-way error - FIN received`—This counter is incremented when the FIN is received during the 3-way handshake. • `TCP 3-way error - invalid flags (PSH, URG, ECE, CWR)`—This counter is incremented when any of the PSH, URG, ECE, or CWR flags were received during the 3-way handshake. • `TCP 3-way error - SYN recvd but no client flows`—This counter is incremented when SYN is received but not from the connection initiator. The counter is not incremented in the case of simultaneous open, when the SYN is received in both the directions. • `TCP 3-way error - first packet SYN+ACK`—The first packet received was SYN+ACK instead of SYN. • `TCP 3-way error - first packet FIN+ACK`—The first packet received was FIN+ACK instead of SYN.

Table 1538: show services stateful-firewall statistics Output Fields *(Continued)*

Field Name	Field Description
	• `TCP 3-way error - first packet FIN`—The first packet received was FIN instead of SYN. • `TCP 3-way error - first packet RST`—The first packet received was RST instead of SYN. • `TCP 3-way error - first packet ACK`—The first packet received was ACK instead of SYN. • `TCP 3-way error - first packet invalid flags (PSH, URG, ECE, CWR)`—The first packet received had invalid flags. • `TCP Close error - no final ACK`—This counter is incremented when ACK is not received after the FINs are received from both directions. • `TCP Resumed Flow`—Plain ACKs create flows if rule match permits, and these are classified as TCP Resumed Flows. This counter is incremented in the case of a TCP Resumed Flow.
`UDP Errors`	UDP protocol errors: • `IP data length less than minimum UDP header length (8 bytes)`—Minimum UDP header length is 8 bytes. The received IP packets contain less than 8 bytes. • `Source or destination port is zero`—UDP source or destination port is `0`. • `UDP port scan (ICMP error seen for UDP flow)`—ICMP error is received for a UDP flow. This could be a genuine UDP flow, but it is counted as an error.

Table 1538: show services stateful-firewall statistics Output Fields *(Continued)*

Field Name	Field Description
`ICMP Errors`	ICMP protocol errors: • `IP data length less than minimum ICMP header length (8 bytes)`—ICMP header length is 8 bytes. This counter is incremented when received IP packets contain less than 8 bytes. • `ICMP error length inconsistencies`—Minimum length of an ICMP error packet is 48 bytes, and the maximum length is 576 bytes. This counter is incremented when the received ICMP error falls outside this range. • `Duplicate ping sequence number`—Received ping packet has a duplicate sequence number. • `Mismatched ping sequence number`—Received ping packet has a mismatched sequence number. • `No matching flow`—No matching existing flow was found for the ICMP error.

Table 1538: show services stateful-firewall statistics Output Fields *(Continued)*

Field Name	Field Description
`ALG errors`	Accumulation of all the application-level gateway protocol (ALG) drops counted separately in the ALG context: • `BOOTP`—Bootstrap protocol errors • `DCE-RPC`—Distributed Computing Environment-Remote Procedure Call protocols errors • `DCE-RPC portmap`—Distributed Computing Environment-Remote Procedure Call protocols portmap service errors • `DNS`—Domain Name System protocol errors • `Exec`—Exec errors • `FTP`—File Transfer Protocol errors • `H323`—H.323 standards errors • `ICMP`—Internet Control Message Protocol errors • `IIOP`—Internet Inter-ORB Protocol errors • `Login`—Login errors • `NetBIOS`—NetBIOS errors • `Netshow`—NetShow errors • `Real Audio`—RealAudio errors • `RPC`—Remote Procedure Call protocol errors • `RPC portmap`—Remote Procedure Call protocol portmap service errors • `RTSP`—Real-Time Streaming Protocol errors • `Shell`—Shell errors • `SIP`—Session Initiation Protocol errors • `SNMP`—Simple Network Management Protocol errors

Table 1538: show services stateful-firewall statistics Output Fields *(Continued)*

Field Name	Field Description
	• `SQLNet`—SQLNet errors • `TFTP`—Trivial File Transfer Protocol errors • `Traceroute`—Traceroute errors
`Drop Flows`	• `Maximum Ingress Drop flows allowed`--Maximum number of ingress flow drops allowed. • `Maximum Egress Drop flows allowed`--Maximum number of egress flow drops allowed. • `Current Ingress Drop flows`--Current number of ingress flow drops. • `Current Egress Drop flows`--Current number of egress flow drops. • `Ingress Drop Flow limit drops count`--Number of ingress flow drops due to maximum number of ingress flow drops being exceeded. • `Egress Drop Flow limit drops count`--Number of egress flow drops due to maximum number of egress flow drops being exceeded.

Sample Output

show services stateful-firewall statistics extensive

```
user@host> show services stateful-firewall statistics extensive
Interface: ms-1/3/0
  Service set: interface-svc-set
    New flows:
      Rule Accepts: 907, Rule Discards: 0, Rule Rejects: 0
    Existing flow types packet counters:
      Accepts: 3535, Drop: 0, Rejects: 0
    Haripinning counters:
      Slow Path Hairpinned Packets: 0, Fast Path Hairpinned Packets: 0
    Drops:
      IP option: 0, TCP SYN defense: 0
      NAT ports exhausted: 0, Sessions dropped due to subscriber flow limit: 0
    Errors:
```

```
      IP: 0, TCP: 0
      UDP: 0, ICMP: 0
      Non-IP packets: 0, ALG: 0
    IP errors:
      IP packet length inconsistencies: 0
      Minimum IP header length check failures: 0
      Reassembled packet exceeds maximum IP length: 0
      Illegal source address: 0
      Illegal destination address: 0
      TTL zero errors: 0, Illegal IP protocol number (0 or 255): 0
      Land attack: 0
      Non-IPv4 packets: 0, Bad checksum: 0
      Illegal IP fragment length: 0
      IP fragment overlap: 0
      IP fragment reassembly timeout: 0
      IP fragment limit exceeded:0
      Unknown: 0
    TCP errors:
      TCP header length inconsistencies: 0
      Source or destination port number is zero: 0
      Illegal sequence number and flags combination: 0
      SYN attack (multiple SYN messages seen for the same flow): 0
      First packet not a SYN message: 0
      TCP port scan (TCP handshake, RST seen from server for SYN): 0
      Bad SYN cookie response: 0
      TCP reconstructor sequence number error: 0
      TCP reconstructor retransmissions: 0
      TCP partially opened connection timeout (SYN): 0
      TCP partially opened connection timeout (SYN-ACK): 0
      TCP partially closed connection reuse: 0
      TCP 3-way error - client sent SYN+ACK: 0
      TCP 3-way error - server sent ACK: 0
      TCP 3-way error - SYN seq number retransmission mismatch: 0
      TCP 3-way error - RST seq number mismatch: 0
      TCP 3-way error - FIN received: 0
      TCP 3-way error - invalid flags (PSH, URG, ECE, CWR): 0
      TCP 3-way error - SYN recvd but no client flows: 0
      TCP 3-way error - first packet SYN+ACK: 0
      TCP 3-way error - first packet FIN+ACK: 0
      TCP 3-way error - first packet FIN: 0
      TCP 3-way error - first packet RST: 0
      TCP 3-way error - first packet ACK: 0
      TCP 3-way error - first packet invalid flags (PSH, URG, ECE, CWR): 0
```

```
      TCP Close error - no final ACK: 0
      TCP Resumed Flow: 0
    UDP errors:
      IP data length less than minimum UDP header length (8 bytes): 0
      Source or destination port is zero: 0
      UDP port scan (ICMP error seen for UDP flow): 0
    ICMP errors:
      IP data length less than minimum ICMP header length (8 bytes): 0
      ICMP error length inconsistencies: 0
      Duplicate ping sequence number: 0
      Mismatched ping  sequence number: 0
      No matching flow: 0
    ALG errors:
      BOOTP: 0, DCE-RPC: 0, DCE-RPC portmap: 0
      DNS: 0, Exec: 0, FTP: 0
      H323: 0, ICMP: 0, IIOP: 0
      Login: 0, NetBIOS: 0, Netshow: 0
      Real Audio: 0, RPC: 0, RPC portmap: 0
      RTSP: 0, Shell: 0, SIP: 0
      SNMP: 0, SQLNet: 0, TFTP: 0
      Traceroute: 0
    Drop Flows:
      Maximum Ingress Drop flows allowed: 20
      Maximum Egress Drop flows allowed: 20
      Current Ingress Drop flows: 0
      Current Egress Drop flows: 0
      Ingress Drop Flow limit drops count: 0
      Egress Drop Flow limit drops count: 0

**If max-drop-flows is not configured, the following is shown**
    Drop Flows:
      Maximum Ingress Drop flows allowed: Default
      Maximum Egress Drop flows allowed: Default
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear services stateful-firewall statistics | 4191

show services stateful-firewall statistics application-protocol sip

IN THIS SECTION

- Syntax | 26318
- Description | 26318
- Options | 26318
- Required Privilege Level | 26318
- Output Fields | 26318
- Sample Output | 26321
- Release Information | 26322

Syntax

```
show services stateful-firewall application-protocol sip
```

Description

Display stateful firewall Session Initiation Protocol (SIP) statistics.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1539 on page 26319 lists the output fields for the `show services stateful-firewall statistics application-protocol-sip` command. Output fields are listed in the approximate order in which they appear.

Table 1539: show services stateful-firewall statistics application-protocol-sip Output Fields

Field Name	Field Description
`Interface`	Name of an adaptive services interface.
`Service set`	Name of the service set flow.
`ALG`	Name of the application-layer gateway.
`Active SIP call count`	Number of active SIP calls.
`Active SIP registration count`	Number of active SIP registrations.
`REGISTER`	Number of new, invalid, and retransmitted register requests sent to the SIP registrar.
`INVITE`	Number of new, invalid, and retransmitted invite messages sent by user agent clients.
`ReINVITE`	Number of new, invalid, and retransmitted reinvite messages sent by user agent clients.
`ACK`	Number of new, invalid, and retransmitted ACK messages received (in response to a SIP Call Invite message).
`BYE`	Number of new, invalid, and retransmitted requests to terminate SIP dialogues.
`CANCEL`	Number of new, invalid, and retransmitted SIP request cancellations.
`SUBSCRIBE`	Number of new, invalid, and retransmitted SIP requests to subscribe for event notifications.
`NOTIFY`	Number of new, invalid, and retransmitted event notifications in SIP dialogues.
`OPTIONS`	Number of new, invalid, and retransmitted requests to query SIP capabilities.

Table 1539: show services stateful-firewall statistics application-protocol-sip Output Fields *(Continued)*

Field Name	Field Description
INFO	Number of new, invalid, and retransmitted requests carrying application-level information.
UPDATE	Number of new, invalid, and retransmitted SIP dialogue updates.
REFER	Number of new, invalid, and retransmitted requests to the recipient to contact a third party.
Provisional responses	Number of new, invalid, and retransmitted responses from the user agent server to indicate the progress of a SIP transaction.
OK responses to INVITEs	OK responses sent from the user agent clients to user agent servers in response to Invite messages. The server can then return an ACK message.
OK responses to non-INVITES	OK responses to SIP messages other than an Invite message.
Redirection responses	Responses from the user agent server to a user agent client requesting the client to contact a different SIP uniform resource identifier (URI).
Request failure responses	Responses that indicate a definite failure from a particular server. The client must not retry the same request without modification after receiving this response.
Server failure responses	Responses that indicate a server failure.
Global failure responses	Responses that indicate a server has definitive information about a particular user, not just the particular instance indicated in the Request URI.
Invalid responses	Responses that are invalid.
Response (all) retransmits	Retransmissions of all responses.

Table 1539: show services stateful-firewall statistics application-protocol-sip Output Fields *(Continued)*

Field Name	Field Description
Parser	Syntax errors, content errors, and unknown methods counted by the message parser.

Sample Output

show services stateful-firewall statistics application-protocol-sip

```
user@host> show services stateful-firewall statistics application-protocol sip
Interface: sp-0/3/0
  Service set: test_sip_777, ALG: SIP
  Active SIP call count: 0, Active SIP registration count: 1
                    New      Invalid    Retransmit
    REGISTER          2
    INVITE            1                          0
    ReINVITE          1
    ACK               1            0             0
    BYE               0            0
    CANCEL            0            0
    SUBSCRIBE         0            0
    NOTIFY            0            0
    OPTIONS           0            0
    INFO              0            0
    UPDATE            0            0
    REFER             0            0
    Provisional responses (18x): 1, OK responses to INVITEs: 2
    OK responses to non-INVITEs: 2, Redirection (3xx) responses: 0
    Request failure (4xx) responses: 0, Server failure (5xx) responses: 0
    Global failure (6xx) responses: 0, Invalid responses: 0
    Response (all) retransmits: 0
    Parser:
      Syntax errors: 0, Content errors: 0, Unknown methods: 0
  Service set: test_sip_888, ALG: SIP
  Active SIP call count: 0, Active SIP registration count: 1
                    New      Invalid    Retransmit
    REGISTER          2
    INVITE            0                          0
    ReINVITE          0
```

```
    ACK                0            0             0
    BYE                0            0
    CANCEL             0            0
    SUBSCRIBE          0            0
    NOTIFY             0            0
    OPTIONS            0            0
    INFO               0            0
    UPDATE             0            0
    REFER              0            0
    Provisional responses (18x): 0, OK responses to INVITEs: 0
    OK responses to non-INVITEs: 2, Redirection (3xx) responses: 0
    Request failure (4xx) responses: 0, Server failure (5xx) responses: 0
    Global failure (6xx) responses: 0, Invalid responses: 0
    Response (all) retransmits: 0
    Parser:
      Syntax errors: 0, Content errors: 0, Unknown methods: 0
```

Release Information

Command introduced in Junos OS Release 7.4.

show services stateful-firewall subscriber-analysis

IN THIS SECTION

- Syntax | 26323
- Description | 26323
- Options | 26323
- Required Privilege Level | 26323
- Output Fields | 26323
- Sample Output | 26324
- Release Information | 26326

Syntax

```
show services stateful-firewall subscriber analysis
<interface interface-name>
```

Description

Display information about the number of active subscribers on the service physical interface card (PIC).

Options

none	Display standard information about all active subscribers on the PIC.
interface *interface-name*	(Optional) Display information about a particular interface.

Required Privilege Level

view

Output Fields

Table 1540 on page 26323 lists the output fields for the `show services stateful-firewall subscriber analysis` command. Output fields are listed in the approximate order in which they appear.

Table 1540: show services stateful-firewall subscriber-analysis Output Fields

Field Name	Field Description
Services PIC Name	Name of an adaptive services interface.
Total Subscribers Active	Total number of subscribers currently active on the service PIC.
Created Subscribers per Second	Rate at which subscribers are currently being created on the service PIC.
Deleted Subscribers per Second	Rate at which subscribers are currently being deleted on the service PIC.

Table 1540: show services stateful-firewall subscriber-analysis Output Fields *(Continued)*

Field Name	Field Description
Peak Total Subscribers Active	Highest number of subscribers that were active during the lifetime of the service PIC.
Peak Created Subscribers per Second	Highest rate at which subscribers were being created during the lifetime of the service PIC.
Peak Deleted Subscribers per Second	Highest rate at which subscribers were being deleted during the lifetime of the service PIC.
Number of Samples	The current sampling period lifetime.
Subscriber Operation: Creation	Number of sampling intervals during which a number of subscribers in the indicated range were created during the current sampling period.
Subscriber Operation: Deletion	Number of sampling intervals during which a number of subscribers in the indicated range were deleted during the current sampling period.

Sample Output

show services stateful-firewall subscriber analysis

```
user@host> show services stateful-firewall subscriber analysis
  Services PIC Name:    sp-2/0/0
Subscriber Analysis Statistics:
  Total Subscribers Active               :100000
  Created Subscribers per Second         :0
  Deleted Subscribers per Second         :0
  Peak Total Subscribers Active          :100000
  Peak Created Subscribers per Second    :2389
  Peak Deleted Subscribers per Second    :0

Subscriber Rate Data:
  Number of Samples: 55
```

```
Subscriber Rate Distribution(sec)
Subscriber Operation :Creation

    300000+            :0
    250000  - 300000   :0
    200000  - 250000   :0
    160000  - 200000   :0
    150000  - 160000   :0
     50000  - 150000   :0
     40000  -  50000   :0
     30000  -  40000   :0
     20000  -  30000   :0
     10000  -  20000   :0
      1000  -  10000   :42
         0  -   1000   :1
Subscriber Operation :Deletion
    300000+            :0
    250000  - 300000   :0
    200000  - 250000   :0
    160000  - 200000   :0
    150000  - 160000   :0
     50000  - 150000   :0
     40000  -  50000   :0
     30000  -  40000   :0
     20000  -  30000   :0
```

show services stateful-firewall subscriber-analysis

```
user@host> show services stateful-firewall subscriber analysis
  Services PIC Name:    sp-2/0/0

Subscriber Analysis Statistics:

  Total Subscribers Active               :23547
  Created Subscribers per Second         :2389
  Deleted Subscribers per Second         :0
  Peak Total Subscribers Active          :23547
  Peak Created Subscribers per Second    :2389
  Peak Deleted Subscribers per Second    :0

Subscriber Rate Data:
```

```
 Number of Samples: 16

Subscriber Rate Distribution(sec)

Subscriber Operation :Creation

   300000+            :0
   250000  - 300000   :0
   200000  - 250000   :0
   160000  - 200000   :0
   150000  - 160000   :0
    50000  - 150000   :0
    40000  -  50000   :0
    30000  -  40000   :0
    20000  -  30000   :0
    10000  -  20000   :0
     1000  -  10000   :9
        0  -   1000   :1


Subscriber Operation :Deletion

   300000+            :0
   250000  - 300000   :0
   200000  - 250000   :0
   160000  - 200000   :0
   150000  - 160000   :0
    50000  - 150000   :0
    40000  -  50000   :0
    30000  -  40000   :0
    20000  -  30000   :0
    10000  -  20000   :0
     1000  -  10000   :0
        0  -   1000   :0
```

Release Information

Command introduced in Junos OS Release 11.4.

show services subscriber analysis

IN THIS SECTION

- Syntax | 26327
- Description | 26327
- Options | 26327
- Required Privilege Level | 26327
- Output Fields | 26328
- Sample Output | 26329
- Release Information | 26330

Syntax

```
show services subscriber analysis
<interface interface-name>
```

Description

Display information about the number of active subscribers on the services PIC.

Options

none	Display standard information about all active subscribers on the PIC.
interface *interface-name*	(Optional) Display information about the specified interface.

Required Privilege Level

view

Output Fields

Table 1541 on page 26328 lists the output fields for the `show services subscriber analysis` command. Output fields are listed in the approximate order in which they appear.

Table 1541: show services subscriber analysis Output Fields

Field Name	Field Description
Services PIC Name	Name of an adaptive services interface.
Subscriber Analysis Statistics:	
Total Subscribers Active	Total number of subscribers currently active on the service PIC.
Created Subscribers per Second	Rate at which subscribers are currently being created on the service PIC.
Deleted Subscribers per Second	Rate at which subscribers are currently being deleted on the service PIC.
Peak Total Subscribers Active	Highest number of subscribers that were active during the lifetime of the service PIC.
Peak Created Subscribers per Second	Highest rate at which subscribers were being created during the lifetime of the service PIC.
Peak Deleted Subscribers per Second	Highest rate at which subscribers were being deleted during the lifetime of the service PIC.
Number of Samples	Number of samples during the current sampling period lifetime.
Subscriber Rate Distribution(sec)	
Subscriber Operation: Creation	Number of sampling intervals during which a number of subscribers in the indicated range were created during the current sampling period.
Subscriber Operation: Deletion	Number of sampling intervals during which a number of subscribers in the indicated range were deleted during the current sampling period.

Sample Output

show services subscriber analysis interface

```
user@host> show services subscriber analysis interface ms-5/1/0
Services PIC Name:    ms-5/1/0

Subscriber Analysis Statistics:

  Total Subscribers Active               :0
  Created Subscribers per Second         :0
  Deleted Subscribers per Second         :0
  Peak Total Subscribers Active          :0
  Peak Created Subscribers per Second    :0
  Peak Deleted Subscribers per Second    :0

Subscriber Rate Data:
  Number of Samples: 3916

 Subscriber Rate Distribution(sec)

 Subscriber Operation :Creation

    400000+            :0
    350001  - 400000   :0
    300001  - 350000   :0
    250001  - 300000   :0
    200001  - 250000   :0
    160001  - 200000   :0
    150001  - 160000   :0
     50001  - 150000   :0
     40001  -  50000   :0
     30001  -  40000   :0
     20001  -  30000   :0
     10001  -  20000   :0
      1001  -  10000   :0
         1  -   1000   :0
                   0   :3916


 Subscriber Operation :Deletion
```

```
400000+            :0
350001  - 400000   :0
300001  - 350000   :0
250001  - 300000   :0
200001  - 250000   :0
160001  - 200000   :0
150001  - 160000   :0
 50001  - 150000   :0
 40001  -  50000   :0
 30001  -  40000   :0
 20001  -  30000   :0
 10001  -  20000   :0
  1001  -  10000   :0
     1  -   1000   :0
                0   :3916
```

Release Information

Statement introduced in Junos OS Release 17.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

show services tcp-log

IN THIS SECTION

- Syntax | 26331
- Description | 26331
- Required Privilege Level | 26331
- Sample Output | 26331
- Release Information | 26331

Syntax

```
show services tcp-log
```

Description

Display the specified TCP log.

Required Privilege Level

Sample Output

show services tcp-log

```
user@host> show services tcp-log
user@hst> show services tcp-log log1
Interface: vms-1/0/0

State: Reconnect-In-Progress
    5.5.5.1 -> 70.0.0.2 : 514
```

Release Information

Command introduced in Junos OS Release 19.3R2.

show services tcp-log connections

IN THIS SECTION

- Syntax | 26332
- Description | 26332
- Required Privilege Level | 26332

- Output Fields | 26332
- Sample Output | 26332
- Release Information | 26333

Syntax

```
show services tcp-log connections interface interface-name
```

Description

Display the TCP connection status for the specified interface.

Required Privilege Level

view

Output Fields

Table 1542 on page 26332 lists the output fields for the `show services sessions tcp-log connections interface` *`interface-name`* command. Output fields are listed in the approximate order in which they appear.

Table 1542: show services sessions tcp-log connections interface Output Fields

Field Name	Field Description
Session Id	TCP connection status including the state, source IP address, destination IP address, and destination port.

Sample Output

```
user@router>show services tcp-log connections  interface interface-name
```

command-name

```
user@router>show services tcp-log connections interfaceinterface-name

Session Id: 1744830467 State: Established
    10.1.1.1 -> 10.40.0.2 : 10214
```

Release Information

Command introduced in Junos OS Release 19.1.

show services traffic-detection-function hcm statistics

IN THIS SECTION

- Syntax | 26333
- Description | 26334
- Options | 26334
- Required Privilege Level | 26334
- Output Fields | 26334
- Sample Output | 26336
- Sample Output | 26337
- Release Information | 26337

Syntax

```
show services traffic-detection-function hcm statistics
<ipv4-address v4-addr>
<ipv6-address v6-addr>
<routing-instance routing-instance>
```

Description

Display the statistics related to HTTP header enrichment for all the active HTTP sessions for the TDF subscriber.

Options

none	Display HTTP header enrichment statistics for all active HTTP sessions.
ipv4-address *v4-addr*	(Optional) Display HCM statistics for the specified IPv4 address of the subscriber's user equipment (UE).
ipv6-address *v6-addr*	(Optional) Display HCM statistics for the specified IPv6 address of the subscriber's user equipment.
routing-instance *routing-instance*	(Optional) Display HCM statistics for the specified routing instance of the subscriber's user equipment.

Required Privilege Level

view

Output Fields

Table 1543 on page 26334 lists the output fields for the `show services traffic-detection-function hcm statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1543: show services traffic-detection-function hcm statistics Output Fields

Field Name	Field Description
`Interface Name`	Name of the services PIC on which data sessions are being serviced. The HTTP header enrichment statistics sessions are displayed per services PIC.
`Session id`	Identifier for the session.

Table 1543: show services traffic-detection-function hcm statistics Output Fields *(Continued)*

Field Name	Field Description
`Subscriber-type`	Type of subscriber: • `ip`—IP-based subscriber. • `ifl`—Interface-based subscriber.
`IMSI`	International Mobile Subscriber Identity (IMSI) of the subscriber's user detail equipment (UE).
`MSISDN`	Mobile station ISDN of the subscriber's user equipment.
`Header inserted`	Number of times that tags were successfully inserted into HTTP headers for the data session.
`Header insert failed`	Number of times that the insertion of tags into HTTP headers failed for the data session.
`Header renamed`	Number of times an HTTP header was renamed.
`Header rename fail`	Number of times an attempt to rename an HTTP header failed.
`IPV4 mask modification`	Number of times IPv4 address mask was inserted.
`IPV6 mask modification`	Number of times IPv6 address mask was inserted.
`Tag too large`	Number of tags that cannot be inserted into HTTP headers because the tag size was larger than the maximum configured size for the data session.
`Tag encryption failed`	Number of times that the encryption of HTTP tags used for header insertion failed for the data session.

Table 1543: show services traffic-detection-function hcm statistics Output Fields *(Continued)*

Field Name	Field Description
Total Get request	Total number of HTTP Get Requests received for the data session.
Subscriber info unavailable	Number of times that subscriber attributes were missing during attempted header insertions for the data session.
Subscriber attribute missing	Number of times that tags were not inserted because subscriber attributes were missing.
IPV4 attribute missing	Number of times that tags were not inserted because subscriber IPv4 user address attributes were missing.
IPV6 attribute missing	Number of times that tags were not inserted because subscriber IPv6 user address attributes were missing.
IPV4 / IPV6 attribute	Number of times that IPv4 and IPv6 user address tags were successfully inserted into HTTP headers.

Sample Output

show services traffic-detection-function hcm statistics routing-instance

```
user@host> show services traffic-detection-function hcm statistics routing-instance r1
Interface Name: mams-2/3/0 (ams1)
Session id: 134217730, Subscriber-type: ip
    Header inserted              : 6
    Header insert failed         : 0
    Header renamed               : 36
    Header rename fail           : 0
    IPV4 mask modification       : 3
    IPV6 mask modification       : 0
    Tag too large                : 0
    Tag encryption failed        : 0
```

```
    Total Get request            : 3
    Subscriber info unavailable  : 9
    Subscriber attribute missing : 9
    IPV4 attribute missing       : 0
    IPV6 attribute missing       : 3
    IPV4 / IPV6 attributes       : 0
```

Sample Output

show services traffic-detection-function hcm statistics ipv4-address routing-instance

```
user@host> show services traffic-detection-function hcm statistics ipv4-address 192.0.2.1
routing-instance default
Interface Name: mams-2/0/0 (ams1)
Session id: 67108865, Subscriber Type: IP, IMSI: 324234324, MSISDN: 0
    Header inserted              : 0
    Header insert failed         : 0
    Header renamed               : 0
    Header rename fail           : 0
    IPV4 mask modification       : 0
    IPV6 mask modification       : 0
    Tag too large                : 0
    Tag encryption failed        : 0
    Total Get request            : 0
    Subscriber info unavailable  : 0
    Subscriber attribute missing : 0
    IPV4 attribute missing       : 0
    IPV6 attribute missing       : 0
    IPV4 / IPV6 attributes       : 0
```

Release Information

Statement introduced in Junos OS Release 17.1.

This statement is not supported for Next Gen Services.

RELATED DOCUMENTATION

show services traffic-detection-function sessions | 26338

show services traffic-detection-function sessions

IN THIS SECTION

- Syntax | 26338
- Description | 26338
- Options | 26338
- Required Privilege Level | 26339
- Output Fields | 26339
- Sample Output | 26340
- Sample Output | 26341
- Release Information | 26341

Syntax

```
show services traffic-detection-function sessions
<ipv4-address v4-addr>
<ipv6-address v6-addr>
<routing-instance routing-instance>
```

Description

Display the active data sessions (TCP or UDP flows) that are being serviced (passing through a services PIC) for a specified TDF subscriber.

Options

none	No output is displayed.

ipv4-address *v4-addr* (Optional) Display subscriber sessions for the specified IPv4 address of the subscriber's user equipment (UE).

ipv6-address *v6-addr* (Optional) Display subscriber sessions for the specified IPv6 address of the subscriber's user equipment.

routing-instance *routing-instance* (Optional) Display subscriber sessions for the specified routing instance.

Required Privilege Level

view

Output Fields

Table 1544 on page 26339 lists the output fields for the `show services traffic-detection-function sessions` command. Output fields are listed in the approximate order in which they appear.

Table 1544: show services traffic-detection-function sessions Output Fields

Field Name	Field Description
`Interface Name`	Name of the service PIC on which data sessions are being serviced. The data sessions are displayed per services PIC.
`Service Set`	Name of the service set on which the data session is being serviced.
`Session`	Identifier for the data session.
`ALG`	Identifier for the application-level gateway (ALG).
`Subscriber-type`	Type of subscriber: • `ip`—IP-based subscriber. • `ifl`—Interface-based subscriber.

Table 1544: show services traffic-detection-function sessions Output Fields *(Continued)*

Field Name	Field Description
IMSI	International Mobile Subscriber Identity (IMSI) of the subscriber's user detail equipment (UE).
MSISDN	Mobile station ISDN of the subscriber's user equipment.

For each session, the following information, pertaining to the flow, is displayed:

- Flow protocol: TCP, UDP, or ICMP
- Flow source IP address and source port address
- Flow destination IP address and destination port address
- Flow state: Forward or Drop
- Flow direction: input (I) or output (O)
- Number of packets transmitted

Sample Output

show services traffic-detection-function sessions routing-instance

```
user@host> show services traffic-detection-function sessions routing-instance r1
Interface Name: mams-5/1/0 (ams1)
Service Set: set-hcm, Session: 67258263, ALG: none, Subscriber-type: ip
TCP         192.0.2.8:17751 ->     198.51.100.5:80    Forward  I              31
TCP         198.51.100.5:80    ->     192.0.2.8:17751 Forward  O              53
Service Set: set-hcm, Session: 67269654, ALG: none, Subscriber-type: ifl
TCP         192.0.2.8:18572 ->     198.51.100.5:80    Forward  I              31
TCP         198.51.100.5:80    ->     192.0.2.8:18572 Forward  O              54
Service Set: set-hcm, Session: 83939629, ALG: none, Subscriber-type: ifl
TCP         192.0.2.8:20826 ->     198.51.100.5:80    Forward  I              31
TCP         198.51.100.5:80    ->     192.0.2.8:20826 Forward  O              53
```

Sample Output

show services traffic-detection-function sessions ipv4-address routing-instance

```
user@host> show services traffic-detection-function sessions ipv4-address 203.0.113.1 routing-
instance default
Interface Name: mams-2/0/0 (ams1)
Service Set: tdf-service-set, Session: 33554433, ALG: none, Subscriber Type: IP, IMSI:
324234324, MSISDN: 0
ICMP           203.0.113.1     ->    10.11.0.1     Forward  I              81
ICMP          10.11.0.1     ->    203.0.113.1      Forward  O               0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf subscribers | 27513

show services traffic-load-balance statistics

IN THIS SECTION

- Syntax | 26342
- Description | 26342
- Options | 26342
- Required Privilege Level | 26343
- Output Fields | 26343
- Sample Output | 26351
- Release Information | 26358

Syntax

```
show services traffic-load-balance statistics
<extensive>
<group group-name>
<instance instance-name>
<num-instances number>
<real-service real-service-name>
<summary>
<virtual-service virtual-service-name>
```

Description

The basic form of the command displays the list of real servers associated with this group and traffic statistics, including packet count and byte count

Options

none	Display information about the load-balancing statistics in brief.
extensive	(Optional) Display extensive information about the traffic load-balancing statistics.
group *group-name*	(Optional) Display load-balancing statistics for a specified group of load-balancer servers.
instance *instance-name*	(Optional) Display load-balancing statistics for a specific traffic load balancer (TLB) instance.
num-instances *number*	(Optional) Display load-balancing statistics for a specified number of TLB instances.
real-service *real-service-name*	(Optional) Display load-balancing statistics for a specified load balancer serve.
summary	(Optional) Display summary information about the traffic load-balancing statistics.
virtual-service *virtual-service-name*	(Optional) Display load-balancing statistics for a specified TLB virtual service.

Required Privilege Level

view

Output Fields

Table 1545 on page 26343 lists the output fields for the `show services traffic-load-balance statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1545: show services traffic-load-balance statistics Output Fields

Field Name	Field Description	Level of Output
`Traffic load balance instance name`	Name of the traffic load balancer (TLB) instance that contains the load-distribution-related configuration settings.	All levels
`Multi services interface name`	Name of the services interface used for the TLB instance to provide one-to-one redundancy for server health monitoring. For MS-MPC services card, this is the name of the aggregated multiservices (AMS) interface or “ms-slot/pic/port”. For Next Gen Services and the MX-SPC3 services card, this is the name of the VMS interface or “vms-slot/pic/port”.	All levels
`Interface state`	Inter-process communications (IPC) status between the TLB daemon (traffic-dird) and the health checking daemon (net-monitord). • `DOWN` • `UP`	All levels
`Interface type`	Logical interface type.	All levels
`Route hold timer`	Time that the programmed VIP routes are kept intact after connectivity between traffic-dird and net-monitord daemons is lost. If connectivity is not reestablished within this time, all the VIP routes are withdrawn.	All levels

Table 1545: show services traffic-load-balance statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Traffic load balance virtual svc name`	Name of the virtual service for the TLB instance. The virtual service provides an address that is associated with the group of servers to which traffic is directed.	none `extensive`
`Virtual service`	Name of the virtual service for the TLB instance. The virtual service provides an address that is associated with the group of servers to which traffic is directed.	`summary`
`Routing instance name`	Name of the routing instance used for the virtual service.	none `extensive`
`IP address`	IP address of the virtual service.	none `extensive`
`Address`	IP address of the virtual service.	`summary`
`Sts`	Operational state of the virtual service.	`summary`
`Packet Sent`	Number of packets originating from the clients that the TLB instance virtual service processes for load balancing to next-hop servers.	`summary`
`Byte Sent`	Number of bytes originating from the clients that the TLB instance virtual service processes for load balancing to next-hop servers.	`summary`
`Packet Recv`	Number of packets returning from the next-hop servers that the TLB instance virtual service processes and forwards to the clients.	`summary`
`Byte Recv`	Number of bytes returning from the next-hop servers that the TLB instance virtual service processes and forwards to the clients.	`summary`

Table 1545: show services traffic-load-balance statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Virtual service mode`	Virtual service processing mode. • `layer-2-direct-server-return`—Virtual service is in transparent mode with Layer 2 direct server return (DSR) • `direct-server-return`—Virtual service is in transparent mode with Layer 3 direct server return (DSR) • `translated`—Virtual service is in translated mode.	none `extensive`
`Traffic load balance group name`	Server group name used for the virtual service.	none `extensive`
`Health check interface subunit`	Number of the subunit of the multiservice interface used for health checking.	none `extensive`
`Traffic load balance group down count`	Number of times the status of the TLB server group was down.	`extensive`
`Protocol`	Virtual service protocol, either tcp or udp. In translated mode, packets destined to the virtual service IP address+port number+protocol are load balanced and then replaced by the real service IP address and server listening port number.	none `extensive`
`Port Number`	Virtual service port number. In translated mode, packets destined to the virtual service IP address+port number+protocol are load balanced and then replaced by the real service IP address and server listening port number.	none `extensive`
`Server Listening Port Number`	Real service port number that replaces the virtual service port number. In translated mode, packets destined to the virtual service IP address +port number+protocol are load balanced and then replaced by the real service IP address and server listening port number.	none `extensive`

Table 1545: show services traffic-load-balance statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Demux Nexthop index`	Index number of the demultiplexing next hop for the virtual service. Index number is unique for a VIP, routing-instance, and protocol combination. The demultiplexing next hop is responsible for port-based demultiplexing of traffic to the load-balancing next hop for session distribution.	none `extensive`
`DFW client-id`	Client connection identifier assigned to the TLB daemon (traffic-dird) by the firewall daemon (dfwd) when the daemons are successfully connected.	`extensive`
`Traffic load balance group warmup time`	Time, in seconds, that passes after the traffic-dird daemon comes up until the traffic-dird programs the distribution table on the Packet Forwarding Engine.	`extensive`
`Traffic load balance group auto-rejoin`	Indicates whether the option that allows a server to rejoin the group automatically when it comes up is enabled or not.	`extensive`
`Route metric`	Routing metric assigned to the virtual service. A lower metric makes a route more preferred.	`extensive`
`Virtual service down count`	Number of times the status of the virtual service was down.	`extensive`
`Traffic load balance hash method`	Hash key parameter used for load balancing. Hash keys supported in the ingress direction are protocol, source IP address, and destination IP address.	`extensive`
`Nexthop index`	Index number of the next-hop for the virtual service. A group of servers function as a pool for next-hop session distribution.	none `extensive`
`Up time`	Period of time for which the virtual service is up, in the format *number-of-days hh:mm:ss*.	none `extensive`

Table 1545: show services traffic-load-balance statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Real Server Up count`	Starting in Junos OS Release 16.1R6 and 18.2R1, number of real servers that are up for the specified virtual service or server group.	none
`Real Server Down count`	Starting in Junos OS Release 16.1R6 and 18.2R1, number of real servers that are down for the specified virtual service or server group.	none
`Total packet sent count`	Number of packets originating from the clients that the TLB instance virtual service processes for load balancing to next-hop servers.	none `extensive`
`Total byte sent count`	Number of bytes originating from the clients that the TLB instance virtual service processes for load balancing to next-hop servers.	none `extensive`
`Total packet received count`	Number of packets returning from the next-hop servers that the TLB instance virtual service processes and forwards to the clients.	none `extensive`
`Total byte received count`	Number of bytes returning from the next-hop servers that the TLB instance virtual service processes and forwards to the clients.	none `extensive`
`Network monitoring profile count`	Number of network monitoring profiles that are used to monitor the health of servers used in TLB session distribution.	`extensive`
`Active real service count`	Number of real services that are functional and active.	`extensive`
`Total real service count`	Total number of real services in different states.	`extensive`
`Network monitoring profile index`	Unique index number associated with the network monitoring profile. Network monitoring profiles are used to monitor the health of servers used in TLB session distribution.	`extensive`

Table 1545: show services traffic-load-balance statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Network monitoring profile name`	Name configured for the network monitoring profile.	`extensive`
`Probe type`	Probe type used to examine the health of servers. TLB supports ICMP, TCP, and HTTP health check probes to monitor the health of servers in a group.	`extensive`
`Probe interval`	Frequency, in number of seconds, at which health check probes are sent.	`extensive`
`Probe failure retry count`	Number of failure retries, after which the real service is tagged as down.	`extensive`
`Probe recovery retry count`	Number of successful retries after which the real service is tagged as up.	`extensive`
`Real service`	Name of the TLB server (also referred to as real service). The name is the identifier for a server to which sessions can be distributed using the server distribution table in conjunction with the session distribution API.	none
`Address`	IP address of the configured real service.	none
`Sts`	Operational state of the TLB server.	none
`Packet Sent`	Number of packets originating from the clients that the TLB instance virtual service sends to the real service.	none
`Byte Sent`	Number of bytes originating from the clients that the TLB instance virtual service sends to the real service next-hop server.	none

Table 1545: show services traffic-load-balance statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Packet Recv	Number of packets returning from the real service next-hop server that the TLB instance virtual service processes and forwards to the clients.	none
Byte Recv	Number of bytes returning from the real service next-hop server that the TLB instance virtual service processes and forwards to the clients.	none
Traffic load balance real svc name	Name of the real service used for traffic load-balancing.	extensive
Routing instance name	Name of the routing instance on which the real service is configured.	extensive
IP address	IP address of the configured real service.	extensive
Traffic load balance group name	Name of the server group for real service.	extensive
Admin state	Administrative state of the real service, such as Up or Down.	extensive
Oper state	Operational state of the real service, such as Up or Down.	extensive
Network monitoring probe up count	Number of probes for which the status of the server whose health is checked is observed to be up. If a server group is configured for dual health check, a real service is declared to be UP only if both health-check probes are simultaneously UP; otherwise a real service declared to be DOWN.	extensive
Network monitoring probe down count	Number of probes for which the status of the server whose health is checked is observed to be down.	extensive

Table 1545: show services traffic-load-balance statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Total rejoin event count`	Number of events that caused a server that was previously down and later operational to rejoin a group of real services for load-balancing.	`extensive`
`Total up event count`	Number of TLB events that identified a virtual service or real service to be up.	`extensive`
`Total down event count`	Number of TLB events that identified a virtual service or real service to be down.	`extensive`
`Real Service packet sent count`	Number of packets originating from the clients that the TLB instance virtual service sends to the real service.	`extensive`
`Real Service byte sent count`	Number of bytes originating from the clients that the TLB instance virtual service sends to the real service next-hop server.	`extensive`
`Real Service packet received count`	Number of packets returning from the real service next-hop server that the TLB instance virtual service processes and forwards to the clients.	`extensive`
`Real Service byte received count`	Number of bytes returning from the real service next-hop server that the TLB instance virtual service processes and forwards to the clients.	`extensive`
`Total probe sent`	Number of health-monitoring probes sent from the TLB health check daemon.	`extensive`
`Total probe success`	Number of health-monitoring probes sent from the TLB health check daemon that were successful.	`extensive`
`Total probe fail`	Number of health-monitoring probes attempted to be sent from the TLB health check daemon that failed.	`extensive`

Table 1545: show services traffic-load-balance statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Total probe sent fail`	Number of health-monitoring probes attempted to be sent from the TLB health check daemon that were unsuccessfully initiated.	`extensive`
`Probe state`	Status of the health-check probe, such as `Up` or `Down`.	`extensive`
`Probe sent`	Number of health-check probe requests transmitted from the TLB health check daemon.	`extensive`
`Probe success`	Number of successful health-check probe requests transmitted from the TLB health check daemon.	`extensive`
`Probe fail`	Number of failed health-check probe requests transmitted from the TLB health check daemon.	`extensive`
`Probe sent failed`	Number of times the TLB health check daemon was unable to initiate transmission of a extensive health-check probe.	`extensive`
`Probe consecutive success`	Number of health-check probe requests transmitted from the TLB health check daemon that were consecutively successful.	`extensive`
`Probe consecutive fail`	Number of health-check probe requests transmitted from the TLB health check daemon that failed for two successive times.	`extensive`

Sample Output

show services traffic-load-balance statistics

```
user@host> show services traffic-load-balance statistics
Traffic load balance instance name    :  lb1
Multi services interface name         :  ms-3/0/0
```

```
Interface state                       :  UP
Interface type                        :  Multi services
Route hold timer                      :  180
Active real service count             :  0
Total real service count              : 100
Traffic load balance virtual svc name :  v1
IP address                            :  0.0.0.0
Virtual service mode                  :  Layer-2 based Direct Server Return mode
Routing instance name                 :  internal-client-vrf
Traffic load balance group name       :  g1
Health check interface subunit        :  40
Demux Nexthop index                   :  N/A
Nexthop index                         :  840
Up time                               :  2d 19:09
Real Server Up count                  :  1
Real Server Down count                :  1
Total packet sent count               :  0
Total byte sent count                 :  0
Real service   Address       Sts  Packet Sent Byte Sent Packet Recv Byte Recv
r11            203.0.113.11    UP   0           0         0           0
r10            203.0.113.10    UP   0           0         0           0

Traffic load balance virtual svc name :  v2
IP address                            :  192.0.2.11
Virtual service mode                  :  Translate mode
Routing instance name                 :  msp-tproxy-forwarding1
Traffic load balance group name       :  g2
Health check interface subunit        :  50
Protocol                              :  tcp
Port number                           :  8080
Server Listening Port Number          :  8084
Demux Nexthop index                   :  536
Nexthop index                         :  539
Up time                               :  2d 19:06
Total packet sent count               :  0
Total byte sent count                 :  0
Total packet received count           :  0
Total byte received count             :  0
Real service     Address      Sts  Packet Sent  Byte Sent  Packet Recv  Byte Recv
r12              203.0.113.12   UP   0            0          0            0
r13              203.0.113.13   UP   0            0          0            0
```

show services traffic-load-balance statistics extensive

```
user@host> show services traffic-load-balance statistics extensive
Traffic Load Balance General Information
    DFW client-id                     :  39

Traffic load balance instance name    :  lb1
Multi services interface name         :  ms-3/0/0
Interface state                       :  UP
Interface type                        :  Multi services
Route hold timer                      :  180
Active real service count             :  0
Total real service count              : 100
Traffic load balance virtual svc name :  v1
IP address                            :  0.0.0.0
Virtual service mode                  :  Layer-2 based Direct Server Return mode
Routing instance name                 :  internal-client-vrf
Traffic load balance group name       :  g1
Traffic load balance group warmup time:  15
Traffic load balance group auto-rejoin:  TRUE
Health check interface subunit        :  40
Traffic load balance group down count :  1
Route metric                          :  1
Virtual service down count            :  1
Traffic load balance hash method      :  source
Network monitoring profile count      :  1
Active real service count             :  2
Total real service count              :  2
Demux Nexthop index                   :  N/A
Nexthop index                         :  840
Up time                               :  2d 19:09
Total packet sent count               :  0
Total byte sent count                 :  0
Total packet received count           :  0
Total byte received count             :  0
Network monitoring profile index      :  1
Network monitoring profile name       :  prof1
Probe type                            :  ICMP
Probe interval                        :  5
Probe failure retry count             :  5
Probe recovery retry count            :  3
```

```
Traffic load balance real svc name    :  r11
Routing instance name                 :  server-vrf10
IP address                            :  203.0.113.11
Traffic load balance group name       :  g1
Admin state                           :  UP
Oper state                            :  UP
Network monitoring probe up count     :  1
Network monitoring probe down count   :  0
Total rejoin event count              :  0
Total up event count                  :  1
Total down event count                :  0
Real Service packet sent count               :  0
Real Service byte sent count                 :  0
Total probe sent                      :  47939
Total probe success                   :  47918
Total probe fail                      :  21
Total probe sent failed               :  0
Network monitoring profile index      :  1
Network monitoring profile name       :  prof1
Probe type                            :  ICMP
Probe state                           :  UP
Probe sent                            :  47939
Probe success                         :  47918
Probe fail                            :  21
Probe sent failed                     :  0
Probe consecutive success             :  10090
Probe consecutive fail                :  0

Traffic load balance real svc name    :  r10
Routing instance name                 :  server-vrf10
IP address                            :  203.0.113.10
Traffic load balance group name       :  g1
Admin state                           :  UP
Oper state                            :  UP
Network monitoring probe up count     :  1
Network monitoring probe down count   :  0
Total rejoin event count              :  0
Total up event count                  :  1
Total down event count                :  0
Real Service packet sent count               :  0
Real Service byte sent count                 :  0
Total probe sent                      :  47939
Total probe success                   :  47917
```

```
Total probe fail                      :  22
Total probe sent failed               :  0
Network monitoring profile index      :  1
Network monitoring profile name       :  prof1
Probe type                            :  ICMP
Probe state                           :  UP
Probe sent                            :  47939
Probe success                         :  47917
Probe fail                            :  22
Probe sent failed                     :  0
Probe consecutive success             :  10090
Probe consecutive fail                :  0

Traffic load balance virtual svc name :  v2
IP address                            :  192.0.2.11
Virtual service mode                  :  Translate mode
Routing instance name                 :  msp-tproxy-forwarding1
Traffic load balance group name       :  g2
Traffic load balance group warmup time:  15
Traffic load balance group auto-rejoin:  TRUE
Health check interface subunit        :  50
Traffic load balance group down count :  1
Protocol                              :  tcp
Port number                           :  8080
Server Listening Port Number          :  8084
Route metric                          :  1
Virtual service down count            :  1
Traffic load balance hash method      :  source-destination
Network monitoring profile count      :  1
Active real service count             :  2
Total real service count              :  2
Demux Nexthop index                   :  536
Nexthop index                         :  539
Up time                               :  2d 19:07
Total packet sent count               :  0
Total byte sent count                 :  0
Total packet received count           :  0
Total byte received count             :  0
Network monitoring profile index      :  1
Network monitoring profile name       :  prof1
Probe type                            :  ICMP
Probe interval                        :  5
Probe failure retry count             :  5
```

```
Probe recovery retry count            :  3

Traffic load balance real svc name    :  r12
Routing instance name                 :  server-vrf10
IP address                            :  203.0.113.12
Traffic load balance group name       :  g2
Admin state                           :  UP
Oper state                            :  UP
Network monitoring probe up count     :  1
Network monitoring probe down count   :  0
Total rejoin event count              :  0
Total up event count                  :  1
Total down event count                :  0
Real Service packet sent count        :  0
Real Service byte sent count          :  0
Real Service packet received count    :  0
Real Service byte received count      :  0
Total probe sent                      :  47939
Total probe success                   :  47916
Total probe fail                      :  23
Total probe sent failed               :  0
Network monitoring profile index      :  1
Network monitoring profile name       :  prof1
Probe type                            :  ICMP
Probe state                           :  UP
Probe sent                            :  47939
Probe success                         :  47916
Probe fail                            :  23
Probe sent failed                     :  0
Probe consecutive success             :  10089
Probe consecutive fail                :  0

Traffic load balance real svc name    :  r13
Routing instance name                 :  server-vrf10
IP address                            :  203.0.113.13
Traffic load balance group name       :  g2
Admin state                           :  UP
Oper state                            :  UP
Network monitoring probe up count     :  1
Network monitoring probe down count   :  0
Total rejoin event count              :  0
Total up event count                  :  1
Total down event count                :  0
```

```
Real Service packet sent count              :  0
Real Service byte sent count                :  0
Real Service packet received count          :  0
Real Service byte received count            :  0
Total probe sent                      :  47939
Total probe success                   :  47910
Total probe fail                      :  29
Total probe sent failed               :  0
Network monitoring profile index      :  1
Network monitoring profile name       :  prof1
Probe type                            :  ICMP
Probe state                           :  UP
Probe sent                            :  47939
Probe success                         :  47910
Probe fail                            :  29
Probe sent failed                     :  0
Probe consecutive success             :  6283
Probe consecutive fail                :  0
```

show services traffic-load-balance statistics summary

```
user@host> show services traffic-load-balance statistics summary
Traffic load balance instance name    :  tlb_sdg
Multi services interface name         :  ms-8/3/0
Interface state                       :  UP
Interface type                        :  Multi services
Route hold timer                      :  180
Active real service count             :  0
Total real service count              : 100
Virtual service     Address       Sts Packet Sent Byte Sent   Packet Recv Byte Recv
DNS-VIP1-TCP        198.51.100.1   Up  13182260    709736171   11951566    732469940
DNS-VIP1-UDP        198.51.100.1   Up  2683203     163675383   2683101     262943898
HTTP-80-ADDRESS-VIP 203.0.113.156  Up  363080548   25152313876 282072340   280409712450
HTTP-8080-ADDR-VIP  203.0.113.157  Up  363198700   25318638843 282030640   280388777065
Secure-Ent-443-VIP  203.0.113.158  Up  30561467    3012763619  28007583    3992807922
Simple-Ent-80-VIP   203.0.113.159  Up  155857682   11558785554 89649255    79217609518

Traffic load balance instance name    :  tlb_sdg_v6
Multi services interface name         :  ms-8/3/0
Interface state                       :  UP
Interface type                        :  Multi services
```

```
Route hold timer                      :  180
Virtual service  Address           Sts Packet Sent Byte Sent   Packet Recv Byte Recv
DNS-VIP1-TCP-V6  2001:db8:a::300     Up  25118146    1829085032  24172053    2088425092
DNS-VIP1-UDP-V6  2001:db8:a::300     Up  1318497     108116747   1319249     386274267
HTTP-80-ADDR-VIP-V6 2001:db8:a::100  Up  368696950   33051271152 282178604   287789935055
HTTP-8080-ADD-VIP-V6 2001:db8:a::100 Up  368797597   33217998028 281989122   287768684085
Sec-Ent-443-VIP-V6 2001:db8:a::200   Up  0662649     3622545250  28080924    4531356641
```

Release Information

Statement introduced in Junos OS Release 16.1.

`num-instances` option added in Junos OS Release 16.1R6 and 18.2R1 on MX Series.

Support added in Junos OS 19.3R2 for Next Gen Services with the MX-SPC3 services card.

show services unified-access-control authentication-table

IN THIS SECTION

- Syntax | **26358**
- Description | **26359**
- Options | **26359**
- Required Privilege Level | **26359**
- Sample Output | **26359**
- Release Information | **26361**

Syntax

```
show services unified-access-control authentication-table
```

Description

Display a summary of the authentication table entries configured from the IC Series UAC Appliance. Authentication tables store mappings between traffic sessions and Unified Access Control (UAC) roles. The IC Series appliance uses the roles specified in the mappings to help determine which UAC policies to apply to a session.

Use this command when you have configured the SRX Series Firewall to act as a Junos OS Enforcer in a UAC deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series appliance.

You can also use this command to display the content of the authentication table in a user role firewall implementation. The table, pushed from a supporting UAC device, provides the user roles associated with incoming traffic.

Options

- `detail`—Display a detailed view of all authentication table entries.
- `extended`—Display a view of all authentication table entries with the user roles listed.
- `identifier` *`id`*—Display all authentication table entries with the specified identifier number.
- `ip` *`source-ip-address`*—Display any authentication table entry for the specified IP address.
- `role` *`role-name`*—Display all authentication table entries for the specified role name.
- `user` *`username`*—Display all authentication table entries for the specified user.

Required Privilege Level

view

Sample Output

show services unified-access-control authentication-table

```
user@host>show services unified-access-control authentication-table
Id    Source IP       Username     Age      Role identifier
1     198.51.100.22          user1        0        0000000001.000005.0
Total: 1
```

show services unified-access-control authentication-table detail

```
user@host>show services unified-access-control authentication-table detail
Identifier: 1
  Source IP: 198.51.100.22
  Username: john
  Age: 0
  Role identifier       Role name
    0000000001.000005.0 Users
    1113249951.100616.0 PersonalFirewall
    1183670148.427197.0 UAC
Total: 1
```

show services unified-access-control authentication-table extended

```
user@host>show services unified-access-control authentication-table extended
Id    Source IP       Username     Age   Role name
3     10.214.161.195          johna      60    Users, PersonalFirewall
6     10.214.161.183          mayb      60    role-1
Total: 2
```

show services unified-access-control authentication-table identifier id

```
user@host>show services unified-access-control authentication-table identifier 1
Identifier: 1
  Source IP: 10.214.161.195
  Username: johna
  Age: 0
  Role identifier       Role name
    0000000001.000005.0 Users
    1113249951.100616.0 PersonalFirewall
    1183670148.427197.0 UAC
Total: 1
```

show services unified-access-control authentication-table ip

```
user@host>show services unified-access-control authentication-table ip 10.214.161.183
Id    Source IP       Username     Age      Role identifier
8     10.214.161.183         mayb          0        1420298444.225667.0
Total: 1
```

show services unified-access-control authentication-table role

```
user@host>show services unified-access-control authentication-table role role-1
Id    Source IP       Username     Age      Role identifier
6     10.214.161.183          maybe         60       1420298444.225667.0
Total: 1
```

show services unified-access-control authentication-table user username

```
user@host>show services unified-access-control authentication-table user prasanta
Id    Source IP       Username     Age      Role identifier
7     10.214.161.195         paul1     0        0000000001.000005.0
Total: 1
```

Release Information

Command introduced in Junos OS Release 9.4. Options updated in Junos OS Release 12.1.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show services unified-access-control counters

IN THIS SECTION

- Syntax | 26362
- Description | 26362
- Required Privilege Level | 26362
- Output Fields | 26362
- Sample Output | 26364
- Release Information | 26365

Syntax

```
show services unified-access-control counters
```

Description

Display the number of sessions allowed, denied, and terminated by the Unified Access Control (UAC) service when invoked by a firewall policy with the uac-policy action. Counts are reported for each action taken by UAC. Sessions that were allowed, denied, or terminated by other firewall policy actions are not included in these statistics.

On SRX1500, SRX5400, SRX5600, and SRX5800 devices, UAC counts are grouped and displayed for each PIC on the device. On SRX 300, SRX 320, SRX 340, SRX 345 SRX Series Firewalls, UAC counts are accumulated by device only. There is no PIC specification on these devices.

Required Privilege Level

view

Output Fields

Table 1546 on page 26363 lists the output fields for the `show services unified-access-control counters` command. Output fields are listed in the approximate order in which they appear.

Table 1546: show services unified-access-control counters Output Fields

Field Name	Field Description
PIC	If applicable, the number of each PIC implementing UAC. UAC statistics are grouped by PIC.
Sessions allowed	The sessions permitted by UAC when invoked by a user role firewall policy.
Policy action	Number of sessions permitted by UAC based on the UAC policy action.
Timeout action	Number of sessions permitted by the timeout action while the SRX was disconnected from the UAC device.
Sessions denied	The sessions denied by UAC when invoked by a user role firewall policy.
Unauthenticated	Number of sessions denied by UAC because the user was not authenticated.
Policy action	Number of sessions denied by UAC based on the UAC policy action.
Policy not matched	Number of sessions denied because no UAC policy match was found.
Timeout action	Number of sessions denied by the timeout action while the SRX was disconnected from the access control device.
Sessions terminated	The sessions originally permitted that were later terminated.
Reevaluation	Number of sessions terminated due to a change in the UAC user roles associated with the session.
Signout	Number of sessions terminated due to the user signing out.

Sample Output

show services unified-access-control counters

```
user@host> show services unified-access-control counters
PIC: fpc2.pic0
  Sessions allowed
    Policy action: 0
    Timeout action: 0
  Sessions denied
    Unauthenticated: 0
    Policy action: 0
    Policy not matched: 0
    Timeout action: 0
  Sessions terminated
    Reevaluation: 0
    Signout: 0
```

command-name

Statistics on SRX 300, SRX 320, SRX 340, and SRX 345 devices are accumulated by device only. There is no PIC specification on these devices.

```
user@host> show services unified-access-control counters
Sessions allowed
    Policy action: 0
    Timeout action: 0
  Sessions denied
    Unauthenticated: 0
    Policy action: 0
    Policy not matched: 0
    Timeout action: 0
  Sessions terminated
    Reevaluation: 0
    Signout: 0
```

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show services unified-access-control policies

IN THIS SECTION

- Syntax | 26365
- Description | 26365
- Options | 26366
- Required Privilege Level | 26366
- Sample Output | 26366
- Sample Output | 26366
- Sample Output | 26367
- Release Information | 26367

Syntax

```
show services unified-access-control policies
```

Description

Display a summary of resource access policies configured from the IC Series UAC Appliance.

Use this command when you have configured the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series appliance.

Options

- `detail`—Display a detailed view of all policies.
- `identifier` *`id`*—Display information about a specific policy by identification number.

Required Privilege Level

view

Sample Output

show services unified-access-control policies

```
user@host> show services unified-access-control policies
Id    Resource                Action Apply        Role identifier
1     10.100.15.0/24:*        allow  selected     1113249951.100616.0
2     10.100.17.0/24:*        deny   all
```

Sample Output

show services unified-access-control policies detail

```
user@host> show services unified-access-control policies detail
Identifier: 1
  Resource: 10.100.15.0/24:*
  Resource: 10.100.16.23-10.100.16.60:*
  Action: allow
  Apply: selected
  Role identifier       Role name
    1113249951.100616.0 Personal Firewall
    1112927873.881659.0 Antivirus
    1183670148.427197.0 UAC
Identifier: 2
  Resource: 10.100.17.0/24:*
  Resource: 10.100.16.23-10.100.16.60:*
  Resource: 10.100.18.0/24:*
  Action: deny
  Apply: all
```

Sample Output

show services unified-access-control policies identifier 1

```
user@host> show services unified-access-control policies identifier 1
Identifier: 1
  Resource: 10.100.15.0/24:*
  Resource: 10.100.16.23-10.100.16.60:*
  Action: allow
  Apply: selected
  Role identifier       Role name
    1113249951.100616.0 Personal Firewall
    1112927873.881659.0 Antivirus
    1183670148.427197.0 UAC
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show services unified-access-control roles

IN THIS SECTION

- Syntax | 26368
- Description | 26368
- Required Privilege Level | 26368
- Output Fields | 26368
- Sample Output | 26368
- Release Information | 26369

Syntax

```
show services unified-access-control roles
```

Description

When implementing user role firewall, display a summary of the roles that have been pushed to the SRX Series Firewall from the access control service.

Required Privilege Level

view

Output Fields

Table 1547 on page 26368 lists the output fields for the `show services unified-access-control roles` command. Output fields are listed in the approximate order in which they appear.

Table 1547: show services unified-access-control roles Output Fields

Field Name	Field Description
`Name`	Name of the user role. The maximum length of role name is 39 characters.
`Identifier`	Unique identifier associated with the specified user role.
`Total`	Total number of user roles specified in the table.

Sample Output

show services unified-access-control roles

```
user@host> show services unified-access-control roles
Name                                    Identifier
Users                                   0000000001.000005.0
```

```
admin-1                                  1420298444.225667.0
Total: 2
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

> Security Policies User Guide for Security Devices
>
> *Firewall User Authentication Overview*

show services unified-access-control status

IN THIS SECTION

- Syntax | 26369
- Description | 26369
- Required Privilege Level | 26370
- Sample Output | 26370
- Release Information | 26370

Syntax

```
show services unified-access-control status
```

Description

Display the status of the connection between the SRX Series Firewall and the IC Series UAC Appliance as well as statistics to help debug connections to the IC Series appliance.

Use this command when you have configured the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series appliance.

Required Privilege Level

view

Sample Output

show services unified-access-control status

```
user@host> show services unified-access-control status
Host           Address         Port   Interface     State
dev106vm26     10.64.11.106    11123  ge-0/0/0.0    connected
dev107vm26     10.64.11.106    11123  ge-0/0/0.0    closed
```

Release Information

Command introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show services url-filter dns-resolution profile

IN THIS SECTION

- Syntax | 26371
- Description | 26371
- Options | 26371
- Required Privilege Level | 26371
- Output Fields | 26372

Sample Output | 26372
Release Information | 26374

Syntax

```
show services url-filter dns-resolution profile profile-name <template template-name>
<fpc-slot fpc-slot pic-slot pic-slot>
```

Description

Display URL filter domain name system (DNS) resolution information.

URL filtering resolves the blocklisted domains. The total number of domains are divided into chunks of 50 domains per chunk. The **filter term** in the command output is the name of a chunk.

> **NOTE**: Starting in Junos OS Release 18.3R1, the `show services url-filter dns-resolution profile` command is deprecated and has been replaced by the `show services web-filter dns-resolution profile` command.

Options

fpc-slot *fpc-slot* **pic-slot** *pic-slot*	(Optional) Specify the FPC and PIC for which you want URL filter information displayed.
profile *profile-name*	Specify the profile for which you want URL filter information displayed.
template *template-name*	(Optional) Specify the template for which you want URL filter information displayed.

Required Privilege Level

view

Output Fields

Table 1548 on page 26372 lists the output fields for the `show services url-filter dns-resolution profile` command. Output fields are listed in the approximate order in which they appear.

Table 1548: show services url-filter dns-resolution profile Output Fields

Field Name	Field Description
`Profile`	Name of profile.
`Template`	Name of template.
`Filter Term`	Name of the domains chunk. All domains are divided into chunks of 50 domains per chunk.
`IPv4 Address Count`	The number of IPv4 addresses resolved for all domains under the filter term.
`IPv6 Address Count`	The number of IPv6 addresses resolved for all domains under the filter term.
`Domain Name`	Name of domain.
`IPv4 Records`	Listing of IPv4 addresses.
`IPv6 Records`	Listing of IPv6 addresses.

Sample Output

show services url-filter dns-resolution profile

```
user@host> show services url-filter dns-resolution profile p1
URL filtering DNS resolution:
Profile: p1
Template: t1

 1). Filter Term: URLF_t1_0004
```

```
  IPv4 Address Count: 20
  IPv6 Address Count: 20

1 ). Domain Name: www.facebook.com

       IPv4 Records:
          31.13.77.36
          31.13.76.68

       IPv6 Records:
          2a03:2880:f122:83:face:b00c:0:25de
          2a03:2880:f111:83:face:b00c:0:25de

2 ). Domain Name: www.youtube.com

       IPv4 Records:
          216.58.193.78
          216.58.194.206

       IPv6 Records:
          2607:f8b0:400a:800::200e
          2607:f8b0:4005:809::200e

3 ). Domain Name: www.netflix.com

       IPv4 Records:
          50.112.200.248
          52.10.96.2
          52.25.242.211
          52.39.87.182
          52.38.44.92
          52.36.125.176
          52.40.2.42
          52.42.184.64
          52.5.80.199
          52.206.203.18
          52.5.231.14
          52.21.94.89
          52.71.118.87
          52.201.133.109
          52.71.122.233
          52.203.136.33
```

```
            IPv6 Records:
               2620:108:700f::342a:b840
               2620:108:700f::3644:fc64
               2620:108:700f::3459:2ce1
               2620:108:700f::3459:c025
               2620:108:700f::3459:f556
               2620:108:700f::3459:c5c5
               2620:108:700f::3644:c2a0
               2620:108:700f::342a:df11
               2406:da00:ff00::3404:d29c
               2406:da00:ff00::3415:a86e
               2406:da00:ff00::3415:fda4
               2406:da00:ff00::3414:91d2
               2406:da00:ff00::3403:73dd
               2406:da00:ff00::22c7:d016
               2406:da00:ff00::3400:290b
               2406:da00:ff00::3213:c65f
```

Release Information

Command introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

- show services url-filter dns-resolution-statistics profile template | 26374
- show services url-filter statistics profile template | 26381
- *Configuring URL Filtering*

show services url-filter dns-resolution-statistics profile template

IN THIS SECTION

- Syntax | 26375
- Description | 26375

- Options | 26375
- Required Privilege Level | 26375
- Output Fields | 26376
- Sample Output | 26378
- Release Information | 26381

Syntax

```
show services url-filter dns-resolution-statistics profile profile-name template template-name
(extensive | summary)
```

Description

Display URL filter domain name system (DNS) resolution statistics.

> **NOTE**: Starting in Junos OS Release 18.3R1, the `show services url-filter dns-resolution-statistics profile template` command is deprecated and has been replaced by the `show services web-filter dns-resolution-statistics profile template` command.

Options

(extensive | summary)	Specify the level of detail of information you want displayed.
profile *profile-name*	Specify the profile for which you want URL filter information displayed.
template *template-name*	Specify the template for which you want URL filter information displayed.

Required Privilege Level

view

Output Fields

Table 1549 on page 26376 lists the output fields for the `show services url-filter dns-resolution-statistics profile template` command. Output fields are listed in the approximate order in which they appear.

Table 1549: show services url-filter dns-resolution-statistics profile template Output Fields

Field Name	Field Description	Level of Detail
`Profile`	Name of profile.	all
`Template`	Name of template.	all
`DNS start time`	Start time of the DNS resolution.	`summary`
`Next DNS start time`	Start time of the next DNS resolution.	`summary`
`Number of resolved A addresses`	Number of resolved IPv4 addresses.	`summary`
`Number of resolved AAAA addresses`	Number of resolved IPv6 addresses.	`summary`
`Number of unresolved A addresses`	Number of unresolved IPv4 addresses.	`summary`
`Number of unresolved AAAA addresses`	Number of unresolved IPv6 addresses.	`summary`
`Number of resolved A domains`	Number of resolved IPv4 domains.	`summary`
`Number of resolved AAAA domains`	Number of resolved IPv6 domains.	`summary`
`Number of unresolved A domains`	Number of unresolved IPv4 domains.	`summary`

Table 1549: show services url-filter dns-resolution-statistics profile template Output Fields *(Continued)*

Field Name	Field Description	Level of Detail
`Number of unresolved AAAA domains`	Number of unresolved IPv6 domains.	`summary`
`Number of requests sent`	Number of DNS requests sent.	`summary`
`Number of responses received`	Number of DNS responses received.	`summary`
`Domain Name`	Name of domain.	`extensive`
`IPv4 Address information`	IPv4 address information includes the following fields: • `DNS server IP`—IPv4 address of DNS server. • `Req Sent`—Number of DNS requests sent. • `Resp Received`—Number of DNS responses received. • `DNS retries`—Number of times no DNS response was received and so retried.	`extensive`

Table 1549: show services url-filter dns-resolution-statistics profile template Output Fields *(Continued)*

Field Name	Field Description	Level of Detail
`IPv6 Address information`	IPv6 address information includes the following fields: • `DNS server IP`—IPv6 address of DNS server. • `Req Sent`—Number of DNS requests sent. • `Resp Received`—Number of DNS responses received. • `DNS retries`—Number of times no DNS response was received and so retried.	`extensive`

Sample Output

show services url-filter dns-resolution-statistics profile template summary

```
user@host> show services url-filter dns-resolution-statistics profile1 template t1 summary
URL filtering DNS resolution statistics:
Profile: p1
Template: t1

     DNS start time                      : May 01 16:40:24 PDT
     Next DNS start time                 : May 01 17:40:24 PDT
     Number of resolved A domains        : 114
     Number of resolved AAAA domains     : 114
     Number of unresolved A domains      : 0
     Number of unresolved AAAA domains   : 0
     Number of requests sent             : 246
     Number of responses received        : 228
```

show services url-filter dns-resolution-statistics profile template extensive

```
user@host> show services url-filter dns-resolution-statistics profile p1 template t1 extensive
URL filtering DNS resolution statistics:
Profile: p1
Template: t1

     1) Domain Name:    www.facebook.com

         IPv4 Address information:
              DNS server IP     8.8.8.8
              Req Sent          20
              Resp Received     20
              DNS retries       0

         IPv4 Address information:
              DNS server IP     172.29.131.60
              Req Sent          21
              Resp Received     20
              DNS retries       0

         IPv6 Address information:
              DNS server IP     8.8.8.8
              Req Sent          25
              Resp Received     20
              DNS retries       0

         IPv6 Address information:
              DNS server IP     172.29.131.60
              Req Sent          24
              Resp Received     20
              DNS retries       0
     2) Domain Name:    www.youtube.com

         IPv4 Address information:
              DNS server IP     8.8.8.8
              Req Sent          21
              Resp Received     20
              DNS retries       0

         IPv4 Address information:
              DNS server IP     172.29.131.60
```

```
            Req Sent          21
            Resp Received     20
            DNS retries       0

        IPv6 Address information:
            DNS server IP     8.8.8.8
            Req Sent          21
            Resp Received     20
            DNS retries       0

        IPv6 Address information:
            DNS server IP     172.29.131.60
            Req Sent          21
            Resp Received     20
            DNS retries       0
    3) Domain Name:    www.netflix.com

        IPv4 Address information:
            DNS server IP     8.8.8.8
            Req Sent          21
            Resp Received     20
            DNS retries       0

        IPv4 Address information:
            DNS server IP     172.29.131.60
            Req Sent          21
            Resp Received     20
            DNS retries       0

        IPv6 Address information:
            DNS server IP     8.8.8.8
            Req Sent          21
            Resp Received     20
            DNS retries       0

        IPv6 Address information:
            DNS server IP     172.29.131.60
            Req Sent          21
            Resp Received     20
            DNS retries       0
```

Release Information

Command introduced in Junos Os Release 17.2.

RELATED DOCUMENTATION

show services url-filter dns-resolution profile | 26370

show services url-filter statistics profile template | 26381

Configuring URL Filtering

show services url-filter statistics profile template

IN THIS SECTION

- Syntax | 26381
- Description | 26381
- Options | 26382
- Required Privilege Level | 26382
- Output Fields | 26382
- Sample Output | 26383
- Release Information | 26384

Syntax

```
show services url-filter statistics profile profile-name template template-name
<fpc-slot fpc-slot pic-slot pic-slot>
```

Description

Display URL filter statistics.

> **NOTE**: Starting in Junos OS Release 18.3R1, the `show services url-filter statistics profile template` command is deprecated and has been replaced by the `show services web-filter statistics profile` command.

Options

Required Privilege Level

view

Output Fields

Table 1550 on page 26382 lists the output fields for the `show services url-filter statistics profile template` command. Output fields are listed in the approximate order in which they appear.

Table 1550: show services url-filter statistics profile template Output Fields

Field Name	Field Description
`Accept`	Action counters for accepted packets.
`Custom page`	Action counters for custom page sent to recipient.
`Http scode`	Action counters for HTTP status code response.
`Redirect url`	Action counters for redirect URL response.
`TCP reset`	Action counters for TCP reset. Connection is closed.
`Bypass session count`	Number of sessions not blocked by URL filtering because the match criteria was not met.
`IPV4 Disable IP Blocking`	Action counters for IPv4 packets that were accepted because filtering is disabled for HTTP traffic that contains an embedded IP address belonging to a blocklisted domain name in the URL filter database.

Table 1550: show services url-filter statistics profile template Output Fields *(Continued)*

Field Name	Field Description
IPV6 Disable IP Blocking	Action counters for IPv6 packets that were accepted because filtering is disabled for HTTP traffic that contains an embedded IP address belonging to a blocklisted domain name in the URL filter database.
session count	The session of activity that a user with a unique IP address spends on a website during a specified period of time. A session, in this case, would be the packets going to the service PIC from the Packet Forwarding Engine and then back to the service PIC.
uplink packet count	Number of packets going from the Packet Forwarding Engine to the service PIC.
uplink bytes	Number of bytes passing uplink.
downlink packet count	Number of packets going from the service PIC to the service Packet Forwarding Engine.
downlink bytes	Number of bytes passing downlink.

Sample Output

show services url-filter statistics profile template

```
user@host> show services url-filter statistics profile p1 template t1
URL filtering action counters:


Accept session count                : 0
Accept uplink packet count          : 0
Accept uplink bytes                 : 0
Accept downlink packet count        : 0
Accept downlink bytes               : 0

Custom page session count           : 0
Custom page uplink packet count     : 0
Custom page uplink bytes            : 0
```

```
Custom page downlink packet count   : 0
Custom page downlink bytes          : 0

Http scode session count            : 0
Http scode uplink packet count      : 0
Http scode uplink bytes             : 0
Http scode dowlink packet count     : 0
Http scode downlink bytes           : 0

Redirect url session count          : 0
Redirect url uplink packet count    : 0
Redirect url uplink bytes           : 0
Redirect url downlink packet count  : 0
Redirect url downlink bytes         : 0

Tcp reset session count             : 0
Tcp reset uplink packet count       : 0
Tcp reset uplink bytes              : 0
Tcp reset downlink packet count     : 0
Tcp reset downlink bytes            : 0

Bypass session count                : 0

IPV4 Disable IP Blocking Sessions          : 0
IPV4 Disable IP Blocking uplink packets    : 0
IPV4 Disable IP Blocking uplink bytes      : 0
IPV4 Disable IP Blocking downlink packets  : 0
IPV4 Disable IP Blocking downlink bytes    : 0

IPV6 Disable IP Blocking Sessions          : 0
IPV6 Disable IP Blocking uplink packets    : 0
IPV6 Disable IP Blocking uplink bytes      : 0
IPV6 Disable IP Blocking downlink packets  : 0
IPV6 Disable IP Blocking downlink bytes    : 0
```

Release Information

Command introduced in Junos Os Release 17.2.

RELATED DOCUMENTATION

show services url-filter dns-resolution profile | 26370

show services url-filter dns-resolution-statistics profile template | 26374

Configuring URL Filtering

show services user-identification active-directory-access domain-controller status

IN THIS SECTION

- Syntax | 26385
- Description | 26385
- Options | 26386
- Required Privilege Level | 26386
- Output Fields | 26386
- Sample Output | 26387
- Sample Output | 26387
- Sample Output | 26388
- Release Information | 26388

Syntax

```
show services user-identification active-directory-access domain-controller status
<domain name> <node (node-id | all | local | primary)> <brief | extensive>
```

Description

Display status information for the Active Directory domain controllers configured for the integrated user firewall feature.

Options

- `domain` *`name`*—(Optional) Display the status of the domain controllers for a specific domain.
- `node`—(Optional) For chassis cluster configurations, display the status of the domain controllers for a specific node.
  - *`node-id`*—Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.
- `brief | extensive`—Display the specified level of output (the default is brief).

Required Privilege Level

view

Output Fields

Table 1551 on page 26386 lists the output fields for the `show services user-identification active-directory-access domain-controller status` command.

Table 1551: show services user-identification active-directory-access domain-controller Output Fields

Field Name	Field Description
`Domain controller`	Domain controller name.
`Address`	IP address of the domain controller.
`Status`	Connection status of the domain controller: connected or disconnected.
`Reason`	Reason for a disconnected status: network issue, authentication failed, or host unreachable.

Sample Output

show services user-identification active-directory-access domain-controller status

Displays brief information for domain controllers in all configured domains.

```
user@host> show services user-identification active-directory-access domain-controller status
Domain: example-domain-controller.com
    Domain controller       Address         Status
    DC1                    203.0.113.51     Connected
    DC2                    203.0.113.12     Connected
    DC3                    203.0.113.6      Connected
    DC4                    203.0.113.11     Disconnected
    DC5                    203.0.113.7      Disconnected

Domain: example-domain
    Domain controller       Address          Status
    example-domain10          10.1.1.1       Disconnected
    example-domain20          10.2.2.2       Disconnected
    example-domain30          10.3.3.3       Disconnected
```

Sample Output

show services user-identification active-directory-access domain-controller status brief domain

```
user@host> show services user-identification active-directory-access domain-controller status
brief domain example-domain-controller.com
Domain: example-domain-controller.com
    Domain controller       Address           Status
    DC1                    203.0.113.51       Connected
    DC2                    203.0.113.12       Connected
    DC3                    203.0.113.6        Connected
    DC4                    203.0.113.11       Disconnected
    DC5                    203.0.113.7        Disconnected
```

Sample Output

show services user-identification active-directory-access domain-controller status extensive domain

```
user@host> show services user-identification active-directory-access domain-controller status
extensive domain example-domain
Domain: example-domain
  Domain controller: example-domain10
    Address: 10.1.1.1
    Status: Disconnected
    Reason: Network issue
  Domain controller: example-domain20
    Address: 10.2.2.2
    Status: Disconnected
    Reason: Authentication failed
  Domain controller: example-domain30
    Address: 10.3.3.3
    Status: Disconnected
    Reason: Host unreachable
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

active-directory-access | 1802

show services user-identification active-directory-access active-directory-authentication-table

show services user-identification active-directory-access statistics

IN THIS SECTION

- Syntax | 26389
- Description | 26389
- Options | 26389
- Required Privilege Level | 26390
- Output Fields | 26390
- Sample Output | 26391
- Sample Output | 26392
- Sample Output | 26393
- Release Information | 26393

Syntax

```
show services user-identification active-directory-access statistics
(ip-user-mapping | ip-user-probe | user-group-mapping) <domain name>
```

Description

Display statistics about IP address-to-user mapping, user-to-group mapping, and IP user probes used for the integrated user firewall feature. If two domains are configured, output is provided per domain.

Options

- `ip-user-mapping`—Number of total queries and failed queries to the event log on the domain controller for address-to-user mappings. Includes additional information, such as the log scan interval and the timestamp of the last event read.
- `ip-user-probe`—Number of total PC probes and failed probes.
- `user-group-mapping`—Number of total queries and failed queries to the LDAP server for user-to-group mappings

- `domain` *`name`*—(Optional) Display the statistics for the specified domain.

Required Privilege Level

view

Output Fields

Table 1552 on page 26390 lists the output fields for the `show services user-identification active-directory-access statistics ip-user-mapping` command.

Table 1552: show services user-identification active-directory-access statistics ip-user-mapping Output Fields

Field Name	Field Description
`Host`	IP address of the domain controller.
`Initial event log timespan`	When the feature is first deployed, the number of previous hours for which the event log on the domain controller is read. A one means the last hour of the event log is read.
`Eventlog scan interval`	Number of seconds between event log scans.
`Total log query number`	Count of the queries on the event log.
`Failed log query number`	Count of the failed queries on the event log.
`Log read number`	Count of the times the event log was read.
`Latest timestamp`	Year:month:date:hours:minutes:seconds is the timestamp taken from the event log. Timestamp records the latest statistics updated time of the SRX Series Firewalls.

Table 1553 on page 26391 lists the output fields for the `show services user-identification active-directory-access statistics ip-user-probe` command.

Table 1553: show services user-identification active-directory-access statistics ip-user-probe Output Fields

Field Name	Field Description
`Total user probe number`	Count of the probes of IP addresses to get IP address-to-user mappings.
`Failed user probe number`	Count of failed probe attempts.

Table 1554 on page 26391 lists the output fields for the `show services user-identification active-directory-access statistics user-group-mapping` command.

Table 1554: show services user-identification active-directory-access statistics user-group-mapping Output Fields

Field Name	Field Description
`Host`	IP address and port being queried.
`Total query number`	Count of queries.
`Failed query number`	Count of failed query attempts.

Sample Output

show services user-identification active-directory-access statistics ip-user-mapping

```
user@host> show services user-identification active-directory-access statistics ip-user-mapping
Domain: example-domain1.com
     Host: 192.0.2.192
     Initial event log timespan : 1
     Eventlog scan interval : 60
     Total log query number : 240
     Failed log query number : 0
     Log read number : 838
     Latest timestamp :2013-10-11:15:11:54
     Host: 192.0.2.50
     Initial event log timespan : 1
```

```
     Eventlog scan interval : 60
     Total log query number : 273
     Failed log query number : 0
     Log read number : 2012
     Latest timestamp :2013-10-11:15:11:23

Domain: example-domain2.com
     Host: 192.0.2.39
     Initial event log timespan : 1
     Eventlog scan interval : 10
     Total log query number : 1596
     Failed log query number : 0
     Log read number : 6691
     Latest timestamp :2013-10-11:15:25:03
     Host: 192.0.2.1
     Initial event log timespan : 1
     Eventlog scan interval : 10
     Total log query number : 2628
     Failed log query number : 0
     Log read number : 114953
     Latest timestamp :2013-10-11:15:24:01
```

Sample Output

show services user-identification active-directory-access statistics ip-user-probe

```
user@host> show services user-identification active-directory-access statistics ip-user-probe
Domain: example-domain3.com
Total user probe number : 176116
Failed user probe number : 916

Domain: example-domain3.com
Total user probe number : 17632
Failed user probe number : 342
```

Sample Output

show services user-identification active-directory-access statistics user-group-mapping

```
user@host> show services user-identification active-directory-access statistics user-group-
mapping
Domain: example-domain3.com
Host: 192.0.2.1 Port  389
Total query number : 176116
Failed query number : 916

Domain: example-domain3.com
Host: 192.0.2.5 Port  389
Total query number : 8965
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

- clear services user-identification active-directory-access | 4194
- ip-user-mapping | 9283
- request services user-identification active-directory-access ip-user-probe | 14985
- user-group-mapping | 30555

show services user-identification active-directory-access user-group-mapping

IN THIS SECTION

- Syntax | 26394
- Description | 26394
- Options | 26394

- Required Privilege Level | 26394
- Output Fields | 26394
- Sample Output | 26396
- Sample Output | 26396
- Sample Output | 26397
- Release Information | 26397

Syntax

```
show services user-identification active-directory-access user-group-mapping
(group name | status | user name) domain name
```

Description

Display user-to-group mapping information used in the integrated user firewall feature. Note that the LDAP server is often part of the domain controller.

Options

- `group` *`group-name`*—Display the users mapped to the specified group.
- `status`—Display the status of the last query to the LDAP server for user-group mapping.
- `user` *`name`*—Display the groups for the specified username.
- `domain` *`name`*—(Optional) Display the group, status, or user information for the specified domain.

Required Privilege Level

view

Output Fields

Table 1555 on page 26395 lists the output fields for the `show services user-identification active-directory-access user-group-mapping group` command.

Table 1555: show services user-identification active-directory-access user-group-mapping group Output Fields

Field Name	Field Description
`Domain`	Domain of the specified group.
`Users`	Usernames mapped to the specified group.

Table 1556 on page 26395 lists the output fields for the `show services user-identification active-directory-access user-group-mapping status` command.

Table 1556: show services user-identification active-directory-access user-group-mapping status Output Fields

Field Name	Field Description
`Domain`	Domain for which the status is displayed.
`LDAP server`	IP address of the LDAP server.
`Port`	Port number on the LDAP server.
`Last-query-status`	Status of the last query from the SRX Series Firewall.
`Last-query-time`	Year-month-date:hour:minutes:seconds when the SRX Series Firewall last queried the LDAP server.

Table 1557 on page 26395 lists the output fields for the `show services user-identification active-directory-access user-group-mapping user` command.

Table 1557: show services user-identification active-directory-access user-group-mapping user Output Fields

Field Name	Field Description
`Domain controller`	Domain controller about which the user information is displayed.

Table 1557: show services user-identification active-directory-access user-group-mapping user Output Fields *(Continued)*

Field Name	Field Description
Groups	Groups to which the user belongs.
Referenced by policy	Groups to which the user belongs and that are referenced by a firewall policy.

Sample Output

show services user-identification active-directory-access user-group-mapping group domain

```
user@host> show services user-identification active-directory-access user-group-mapping group
finance domain www.apac-acme.net
show services user-identification active-directory-access user-group-mapping group finance-group
Domain: example-domain.net
Users: user1, user2
Domain: example2.domain.net
Users: user3
```

Sample Output

show services user-identification active-directory-access user-group-mapping status

```
user@host> show services user-identification active-directory-access user-group-mapping status
Domain: example-domain.net
LDAP server   Port      Last-query-status    Last-query-time
192.0.2.87     389       Query success        2014-02-07:15:50:52

Domain: example2.domain.net
LDAP server   Port      Last-query-status    Last-query-time
192.0.2.144    389       Idle                0
```

Sample Output

show services user-identification active-directory-access user-group-mapping user

```
user@host> show services user-identification active-directory-access user-group-mapping user
user1
Domain example-domain.net
Groups: Dev, NAT, SBU
Referenced by policy: SBU

Domain: example2.domain.net
Groups: HR, USA
```

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

LDAP Functionality in Integrated User Firewall

user-group-mapping | 30555

show services user-identification authentication-table

IN THIS SECTION

- Syntax | 26398
- Description | 26399
- Options | 26400
- Required Privilege Level | 26401
- Output Fields | 26401
- Active Directory | 26403
- Identity Management | 26415

- Identity Management | 26417
- Firewall Authentication Forced Age Timeout | 26420
- Release Information | 26420

Syntax

```
show services user-identification authentication-table
<authentication-source | counter | ip-address>
show services user-identification authentication-table authentication-source
<active-directory | all | aruba-clearpass | identity-management>
show services user-identification authentication-table authentication-source active-directory
<brief | domain | extensive | group | logical-system | root-logical-system |summary |user>
show services user-identification authentication-table authentication-source all
<brief | domain | extensive |group | logical-system | root-logical-system |summary | user>
<domain domain>
<group (group-name | brief | domain | extensive | logical-system | root-logical-system |
summary)>
<logical-system (logical-system-name| all)>
<node (node-id | all | local | primary)>
<root-logical-system (enter |brief | domain | extensive | node)>
<user (user-name | brief | domain | extensive | logical-system | node | root-logical-system |
summary)>
show services user-identification authentication-table authentication-source active-directory
<brief | domain | extensive | group | logical-system | root-logical-system |summary |user>
show services user-identification authentication-table authentication-source identity-management
source-name
show services user-identification authentication-table authentication-source identity-management
tenant <tenant-name> extensive
show services user-identification authentication-table counter
show services user-identification authentication-table ip-address
<summary>
<logical-system logical-system-name>
<root-logical-system>
<tenant tenant-name>
<node node-id>
<IP address ip-address>
```

Description

Display the user identity information authentication table entries for the specified authentication source. You can display the entire contents of the specified authentication source's authentication table, or you can constrain the displayed information to a specific domain, group, or user based on the user name. You can also display identity information for a user based on the IP address of the user's device. You can show brief or extensive information for all of these instances.

authentication-source

User authentication source whose authentication table or identity management server entries are to be displayed.

Authentication sources include:

active-directory

Display the SRX Series active-directory table contents. You can display all of the table's contents or you can delimit the display of user identity information by domain, group, or user name. You can display brief or extensive information for each of these categories.

- domain—Display the entries in the authentication table for the specified domain. You can display summary, group, or user entries for the specified domain.
- group—Display the entries from the authentication table for the specified group.
- user—Display the entries from the authentication table for the specified user based on the user name.

aruba-clearpass

Display the SRX Series Aruba ClearPass authentication table contents. You can display all of the table's contents or you can delimit the display of user information by domain, group, or user name. You can display brief or extensive information for each of these categories.

- domain—Display the entries in the authentication table for the specified domain. You can display summary, group, or user entries for the specified domain.
- group—Display the entries from the authentication table for the specified group.
- user—Display the entries from the authentication table for the specified user based on the user name.

identity-management Display user identity entries contained in the identity-management authentication system.

- source-name—Name of the identity -management source. This could be the Juniper Identity Management Service (JIMS) or any third-party authentication source.
  - If you specify a source, such as “JIMS – Active Directory” for Juniper Identity Management Service, the SRX Series Firewall will show entries only for that authentication source.

    Possible values include:
    - For JIMS: “JIMS – Active Directory”, “JIMS – Exchange”
    - For ClearPass: “Aruba ClearPass”
- domain—Display the entries in the identity management system for the specified domain. You can display summary, group, or user entries for the specified domain.
- group—Display the entries in the identity management system for the specified group.
- user—Display the entries in the identity management system for the specified user based on the user name.
- tenant—Display the entries in the identity management system for the specified tenant system.

Options

- `all`—Summary of the authentication entry information for all entries.
- `group` *`group-name`*—Entries from the authentication table or identity management system for the specified group.
- `ip-address` *`ip-address`*—Entries from the authentication table or identity management system for the specified IP address.
- `user` *`name`*—Entries from the authentication table for the specified username.
- `domain` *`name`*—Summary, group, or user entries for the specified domain.
- `node`—(Optional) For chassis cluster configurations, the summary, IP address, or user entries for a specific node.

- *node-id*—Identification number of the node. It can be 0 or 1.
- all—Display information about all nodes.
- local—Display information about the local node.
- primary—Display information about the primary node.

- brief | extensive—Display the specified level of output (the default is brief).
- logical-system—Display the authentication entries based on the logical system name.
- root-logical-system—Display the authentication entries based on the root logical system.
- tenant *tanant-name*—Display the authentication entries based on the specified tenant system name.

Required Privilege Level

view

Output Fields

Field Name	Field Description
Domain	Name of the domain that the users belong to. User identity and authentication information is display for all users who belong to the domain and for whom there are entries in the specified authentication source table or repository.
Total entries	Number of user entries in the authentication table, by domain.
For each entry:	
Source IP	The IP address of the user's device. If a user is logged in to the network with more than one device, a separate entry is created for the user for each device. It showing the devices IP address.
Username	The name by which the user is logged in to the network.

(Continued)

Field Name	Field Description
Groups	A list of the groups that the user belongs to. The list can include a group that identifies the device posture.
State	The state of the entry. There are four states for an authentication entry: initial, valid, invalid, and pending. • An initial state is a temporary state, and it can be created from either a valid or an invalid entry. The entry had not been pushed to the Packet Forwarding Engine. • A valid state indicates that the authentication entry has a valid IP address, domain, and username. The authentication entry is pushed to the Packet Forwarding Engine. • An invalid state indicates that the entry does not have a valid IP address, domain, and username. If the entry is invalid, it is put in the null domain. • A pending state indicates that the entry was created after the user query was sent and before the response was received. The IP address is being probed.
Source	Authentication source.
Access start date	The date when the authentication entry was created by the SRX Series Firewall.
Access start time	The time when the authentication entry was created by the SRX Series Firewall. The time value is in device local time zone.
Last updated timestamp	The time when the user information was created. This value is taken from the timestamp field in the user information. The time value is in device local time zone.

(Continued)

Field Name	Field Description
Age time	The time, in minutes, after which the entry expires, as configured by the authentication-entry-timeout statement. If a value of 0 was specified, the entry never expires.
Forced Age time	The rest value and the forced value. This information is made available if you configure the firewall-authentication-forced-timeout statement for active directory.

Active Directory

show services user-identification active-directory-access active-directory-authentication-table ip-address

Output of this command displays authentication and identity information for a specific user based on the IP address of the user's device.

```
user@host> show services user-identification active-directory-access active-directory-
authentication-table ip-address 198.51.100.3.
Domain: ad.example.net
 Source-ip: 198.51.100.3
  Username: user1
  Groups:group1
  State: Valid
  Source: wmic
  Access start date: 2014-03-10
  Access start time: 13:59:56
  Age time: 1437
```

show services user-identification authentication-table ip-address

Output of this command displays authentication and identity information for a specific user based on the IP address of the user's device.

```
user@host> show services user-identification authentication-table ip-address 2001:db8::1:1
Domain: ac.example.net
 Source-ip: 2001:db8::1:1
  Username: user1
  Groups:group1
  State: Valid
  Source: wmic
  Access start date: 2017-05-10
  Access start time: 13:59:56
  Age time: 1437
```

show services user-identification active-directory-access active-directory-authentication-table all

Output of this command displays user authentication and identity information for all users for whom there are entries in the active directory authentication table.

```
user@host> show services user-identification active-directory-access active-directory-
authentication-table all
Domain: www.engineering-example.net
Total count: 2
Source IP       Username      Groups          State
198.51.100.22     u2           r1, r3, r4      initial
198.51.100.23     u3           r5, r6, r4      pending

Domain: www.hr-example.net
Total count: 2
Source IP       Username      Groups          State
 198.51.100.26    u4           r1, r3, r4      initial
 198.51.100.27    u5           r5, r6, r4      pending
```

show services user-identification active-directory-access active-directory-authentication-table all extensive

Output of this command, which specifies the `extensive` option, shows state and access information for all entries in the active directory authentication table, in addition to basic information displayed when the `brief` option is used and by default.

```
user@host> show services user-identification active-directory-access active-directory-
authentication-table all extensive

Domain: www.mycompany-example.com
Total entries: 2

Source IP: 198.51.100.29
Username: u2
Groups: r1, r3, r4
State: inital
Access start date: 2013-05-22
Access start time: 10:56:58
Age time: 20 min

Source IP: 198.51.100.30
Username: u3
Groups: r5, r6, r4
State: pending
Access start date: 2013-05-22
Access start time: 10:56:58
Age time: 20 min

Domain: www.hr-example.net
Total entries: 2

Source IP: 198.51.100.31
Username: u2
Groups: r1, r3, r4
State: inital
Access start date: 2013-05-22
Access start time: 10:56:58
Age time: 20 min

Source IP: 198.51.100.32
Username: u3
```

```
Groups: r5, r6, r4
State: pending
Access start date: 2013-05-22
Access start time: 10:56:58
Age time: 20
```

show services user-identification active-directory-access active-directory-authentication-table all domain

Output of this command shows by default brief user identity and authentication information for all users for whom there are entries in the active directory authentication table and whose devices belong to the specified domain.

```
user@host> show services user-identification active-directory-access active-directory-
authentication-table all domain www.mydomain-example.com
Domain: www.mydomain-example.com
Total count: 2
Source IP       Username      Groups          State
 198.51.100.36    u2          r1, r3, r4      initial
 198.51.100.37    u3          r5, r6, r4      pending
```

Capacity of User-Identification Authentication Table

There is a certain limit for maximum number of authentication entries in user-identification authentication table. For example, the maximum capacity in user-identification authentication table is set at 5000 auth-entries for vSRX3 firewall and 2,56,000 auth-entries for SRX5000 line firewalls.

When the user-identification authentication table has maximum number of authentication entries, `ip-probe` will be disabled and `ip-query` cannot be sent for any new incoming sessions. The `ip-probe` gets enabled only after authentication table entry count reduces to 90 % of the the total capacity. When SRX user-identification authentication table has reached its capacity, the new authentication entries gets added only through batch response from JIMS or event log from Active Directory. When new authentication entry gets added, the LRU (last recently used) authentication entry with similar state gets deleted only if there are no active sessions associated with the LRU. The new authentication entry is then added in the authentication table. When no such LRU is found, the new authentication entry gets discarded.

All Authentication Sources

Output of this command shows extensive user identity and authentication information for all users with entries in authentication tables of any authentication source. This example shows only one entry to illustrate the content that is displayed with the extensive option.

```
user@host> show services user-identification authentication-table authentication-source all
extensive
Domain: ad-userfw-example.net
       Total entries: 1
       Source-ip: 198.51.100.1/24
       Username: administrator
       State: Valid
       Source: firewall-authentication
       Access start date: 2016-10-27
       Access start time: 09:30:27
       Age time: 30
```

command-name

```
user@host> show services user-identification authentication-table authentication-source all
logical-system
lsys1
node0:
--------------------------------------------------------------------------
Logical System: root-logical-system

Domain: ad2012.jims.com
Total entries: 18003
Source IP       Username       groups(Ref by policy)          state
bbbb:bbbb:bbbb: jimsuser18000                                 Valid
bbbb:bbbb:bbbb: jimsuser17999                                 Valid
bbbb:bbbb:bbbb: jimsuser17998                                 Valid
bbbb:bbbb:bbbb: jimsuser17997                                 Valid
bbbb:bbbb:bbbb: jimsuser17996                                 Valid
bbbb:bbbb:bbbb: jimsuser17995                                 Valid
bbbb:bbbb:bbbb: jimsuser17994                                 Valid
bbbb:bbbb:bbbb: jimsuser17993                                 Valid
```

command-name

```
user@host> show services user-identification authentication-table authentication-source all root-
logical-system
node0:
--------------------------------------------------------------------------
Logical System: root-logical-system

Domain: ad2012.jims.com
Total entries: 18003
Source IP       Username       groups(Ref by policy)          state
bbbb:bbbb:bbbb: jimsuser10745
bbbb:bbbb:bbbb: jimsuser18000                                  Valid
bbbb:bbbb:bbbb: jimsuser17999                                  Valid
bbbb:bbbb:bbbb: jimsuser17998                                  Valid
bbbb:bbbb:bbbb: jimsuser17997                                  Valid
bbbb:bbbb:bbbb: jimsuser17996                                  Valid
bbbb:bbbb:bbbb: jimsuser17995                                  Valid
bbbb:bbbb:bbbb: jimsuser17994                                  Valid
bbbb:bbbb:bbbb: jimsuser17993                                  Valid
bbbb:bbbb:bbbb: jimsuser17992                                  Valid
user@host> show services user-identification authentication-table authentication-source all node
0
node0:
--------------------------------------------------------------------------
Logical System: root-logical-system

Domain: ad2012.jims.com
Total entries: 18003
Source IP       Username       groups(Ref by policy)          state
bbbb:bbbb:bbbb: jimsuser14716
bbbb:bbbb:bbbb: jimsuser18000                                  Valid
bbbb:bbbb:bbbb: jimsuser17999                                  Valid
bbbb:bbbb:bbbb: jimsuser17998                                  Valid
bbbb:bbbb:bbbb: jimsuser17997                                  Valid
bbbb:bbbb:bbbb: jimsuser17996                                  Valid
bbbb:bbbb:bbbb: jimsuser17995                                  Valid
bbbb:bbbb:bbbb: jimsuser17994                                  Valid
bbbb:bbbb:bbbb: jimsuser17993                                  Valid
```

command-name

```
user@host> show services user-identification authentication-table authentication-source all node
0 logical-system lsys1
node0:
--------------------------------------------------------------------------
Logical System: root-logical-system

Domain: ad2012.jims.com
Total entries: 18003
Source IP       Username       groups(Ref by policy)          state
bbbb:bbbb:bbbb: jimsuser18000                                  Valid
bbbb:bbbb:bbbb: jimsuser17999                                  Valid
bbbb:bbbb:bbbb: jimsuser17998                                  Valid
bbbb:bbbb:bbbb: jimsuser17997                                  Valid
bbbb:bbbb:bbbb: jimsuser17996                                  Valid
bbbb:bbbb:bbbb: jimsuser17995                                  Valid
bbbb:bbbb:bbbb: jimsuser17994                                  Valid
bbbb:bbbb:bbbb: jimsuser17993                                  Valid
bbbb:bbbb:bbbb: jimsuser17992                                  Valid
```

command-name

```
user@host> show services user-identification authentication-table authentication-source all node
0
node0:
--------------------------------------------------------------------------
Logical System: root-logical-system

Domain: ad2012.jims.com
Total entries: 18003
Source IP       Username       groups(Ref by policy)          state
bbbb:bbbb:bbbb: jimsuser1213
bbbb:bbbb:bbbb: jimsuser18000                                  Valid
bbbb:bbbb:bbbb: jimsuser17999                                  Valid
bbbb:bbbb:bbbb: jimsuser17998                                  Valid
bbbb:bbbb:bbbb: jimsuser17997                                  Valid
bbbb:bbbb:bbbb: jimsuser17996                                  Valid
bbbb:bbbb:bbbb: jimsuser17995                                  Valid
```

```
bbbb:bbbb:bbbb: jimsuser17994                                Valid
bbbb:bbbb:bbbb: jimsuser17993                                Valid
```

show services user-identification authentication-table authentication-source all all-logical-systems-tenants

Output of this command displays brief user authentication and identity information for all users for whom there are entries in the identity-management authentication source.

```
user@host> show services user-identification authentication-table authentication-source all all-logical-systems-tenants
node0:
--------------------------------------------------------------------------
Logical System: ld1
Domain: ad03.net
Total entries: 4
Source IP       Username       groups(Ref by policy)          state
12.0.0.2        administrator  posture-healthy                Valid
12.0.0.15       administrator  posture-healthy                Valid
3000::5         N/A            posture-healthy                Valid
2001:db8:::302b N/A        posture-healthy                Valid

Logical System: tn1
Domain: ad03.net
Total entries: 4
Source IP       Username       groups(Ref by policy)          state
12.0.0.2        administrator  posture-healthy                Valid
12.0.0.15       administrator  posture-healthy                Valid
3000::5         N/A            posture-healthy                Valid
2001:db8:::302b N/A            posture-healthy                Valid
```

Aruba ClearPass

show services user-identification authentication-table authentication-source aruba-clearpass domain extensive

Output of this command shows extensive user identity and authentication information, when Aruba ClearPass is used as the authentication source, for all users whose devices belong to the GLOBAL domain.

```
user@host> show services user-identification authentication-table authentication-source aruba-
clearpass domain GLOBAL extensive
Domain: GLOBAL
Total entries: 7
  Source-ip: 203.0.113.21
    Username: vikiyr
    Groups:posture-healthy, accounting-grp, accounting-grp-and-company-device,
    corporate-limited, [user authenticated]
    Groups referenced by policy:accounting-grp-and-company-device,
    corporate-limited
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:20:30
    Last updated timestamp: 2015-12-22 04:02:48
    Age time: 0
  Source-ip: 203.0.113.89
    Username: abewhfy
    Groups:posture-unknown, marketing-access-limited-grp, [user authenticated]
    Groups referenced by policy:marketing-access-limited-grp
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:31:40
    Last updated timestamp: 2015-12-22 04:18:48
    Age time: 0
  Source-ip: 203.0.113.52
    Username: jjxchan
    Groups:posture-healthy, marketing-access-for-pcs-limited-group,
    marketing-general, sales-limited, corporate-limited, [user authenticated]
    Groups referenced by policy:marketing-access-for-pcs-limited-group,
    corporate-limited
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:22:48
    Last updated timestamp: 2015-12-22 05:46:21
    Age time: 0
```

```
Source-ip: 203.0.113.53
  Username: ltchen1
  Groups:posture-healthy, human-resources-grp, accounting-limited,
  corporate-limited, [user authenticated]
  Groups referenced by policy:corporate-limited
  State: Valid
  Source: Aruba ClearPass
  Access start date: 2016-03-08
  Access start time: 17:21:37
  Last updated timestamp: 2015-12-22 05:41:18
  Age time: 0
Source-ip: 203.0.113.54
  Username: guest1
  Groups:posture-healthy, guest, [user authenticated]
  State: Valid
  Source: Aruba ClearPass
  Access start date: 2016-03-08
  Access start time: 17:23:10
  Last updated timestamp: 2015-12-22 05:50:47
  Age time: 0
Source-ip: 203.0.113.55
  Username: guest2
  Groups:posture-healthy, guest-device-byod, [user authenticated]
  State: Valid
  Source: Aruba ClearPass
  Access start date: 2016-03-08
  Access start time: 17:23:21
  Last updated timestamp: 2015-12-22 05:52:44
  Age time: 0
Source-ip: 2001:db8:4136:e378:8000:63bf:3fff:fdd2
  Username: guest3
  Groups:posture-healthy, guest-device-grp, [user authenticated]
  State: Valid
  Source: Aruba ClearPass
  Access start date: 2016-03-08
  Access start time: 17:23:21
  Last updated timestamp: 2015-12-22 05:52:44
  Age time: 0
```

show services user-identification authentication-table authentication-source aruba-clearpass domain brief

Output of this command shows brief user identity and authentication information for users whose devices belong to the GLOBAL domain.

If you do not specify brief, the same information would be displayed. The default behavior is to show brief output.

```
user@host> show services user-identification authentication-table authentication-source aruba-
clearpass domain GLOBAL brief

Domain: GLOBAL
Total entries: 6
Source IP                                Username       groups(Ref by policy)          state
203.0.113.71                             taviki2          accounting-grp-and-company-dev  Valid
203.0.113.89                             gabewb1          marketing-access-limited-grp    Valid
203.0.113.92                             tljxchan         marketing-access-for-pcs-limit  Valid
203.0.113.93                             tjlchen1         corporate-limited               Valid
203.0.113.94                             guest1                                           Valid
203.0.113.95                             guest2                                           Valid
2001:db8:4136:e378:8000:63bf:3fff:fdd2   guest2                                           Valid
```

show services user-identification authentication-table authentication-source aruba-clearpass extensive

Output of the following command shows extensive user identity and authentication information for all users authenticated by Aruba ClearPass for whom entries exist in the aruba-clearpass authentication table.

```
user@host> show services user-identification authentication-table authentication-source aruba-
clearpass extensive

Domain: GLOBAL
Total entries: 7
  Source-ip: 203.0.113.31
    Username: vjki2
    Groups:posture-healthy, accounting-grp, accounting-grp-and-company-device,
    corporate-limited, [user authenticated]
    Groups referenced by policy:accounting-grp-and-company-device,
    corporate-limited
```

```
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:20:30
    Last updated timestamp: 2015-12-22 04:02:48
    Age time: 0
  Source-ip: 203.0.113.89
    Username: labew11
    Groups:posture-unknown, marketing-access-limited-grp, [user authenticated]
    Groups referenced by policy:marketing-access-limited-grp
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:31:40
    Last updated timestamp: 2015-12-22 04:18:48
    Age time: 0
  Source-ip: 203.0.113.62
    Username: dxchan45
    Groups:posture-healthy, marketing-access-for-pcs-limited-group,
    marketing-general, sales-limited, corporate-limited, [user authenticated]
    Groups referenced by policy:marketing-access-for-pcs-limited-group,
    corporate-limited
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:22:48
    Last updated timestamp: 2015-12-22 05:46:21
    Age time: 0
  Source-ip: 2001:db8:4136:e378:8000:63bf:3fff:fdd2
    Username: efchan47
    Groups:posture-healthy, marketing-access-for-pcs-limited-group,
    marketing-general, sales-limited, corporate-limited, [user authenticated]
    Groups referenced by policy:marketing-access-for-pcs-limited-group,
    corporate-limited
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:22:48
    Last updated timestamp: 2015-12-22 05:46:21
    Age time: 0
  Source-ip: 203.0.113.83
    Username: ljhen1
    Groups:posture-healthy, human-resources-grp, accounting-limited,
```

```
    corporate-limited, [user authenticated]
    Groups referenced by policy:corporate-limited
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:21:37
    Last updated timestamp: 2015-12-22 05:41:18
    Age time: 0
  Source-ip: 203.0.113.34
    Username: guest1
    Groups:posture-healthy, guest, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:23:10
    Last updated timestamp: 2015-12-22 05:50:47
    Age time: 0
  Source-ip: 203.0.113.95
    Username: guest2
    Groups:posture-healthy, guest-device-byod, [user authenticated]
    State: Valid
    Source: Aruba ClearPass
    Access start date: 2016-03-08
    Access start time: 17:23:21
    Last updated timestamp: 2015-12-22 05:52:44
    Age time: 0
```

Identity Management

show services user-identification authentication-table authentication-source identity-management brief

Output of this command displays brief user authentication and identity information for all users for whom there are entries in the identity-management authentication source.

```
user@host> show services user-identification authentication-table authentication-source identity-management brief
Domain: ad-domaine-example.net
Total entries: 5
Source IP       Username       groups(Ref by policy)           state
198.51.100.63    N/A                                            Valid
```

```
203.0.113.30     administrator                                Valid
203.0.113.18     N/A                                          Valid
198.51.100.69    N/A                                          Valid
198.51.100.66    administrator                                Valid

Domain: NULL
Total entries: 1
Source IP       Username       groups(Ref by policy
```

show services user-identification authentication-table authentication-source identity-management extensive

Output of this command displays extensive user authentication and identity information for all users for whom there are entries in the identity-management authentication source.

```
user@host> show services user-identification authentication-table authentication-source identity-
management extensive
Domain: ad-domain2-example.net
Total entries: 5
  Source-ip: 198.51.100.63
    Username: N/A
    Groups:posture-healthy
    State: Valid
    Source: JIMS - Active Directory
    Access start date: 2017-06-05
    Access start time: 09:28:45
    Last updated timestamp: 2017-06-06 08:41:56
    Age time: 0
  Source-ip: 198.51.100.66
    Username: administrator
    Groups:posture-healthy, group policy creator owners, enterprise admins, schema admins,
domain admins,
    administrators, denied rodc password replication group
    State: Valid
    Source: JIMS - Active Directory
    Access start date: 2017-06-05
    Access start time: 09:23:44
    Last updated timestamp: 2017-06-06 08:11:45
    Age time: 0
```

show services user-identification authentication-table authentication-source all extensive

Output of this command, which specifies the extensive option, shows state and access information for all entries.

```
user@host> show services user-identification authentication-table authentication-source identity-management extensive
Domain: jims-dom1.local
Total entries: 1
Source-ip: 2001:db8:4136:e378:8000:63bf:3fff:fdd2
Username: user1
Groups:posture-healthy
Groups referenced by policy:posture-healthy
State: Valid
Source: JIMS - Active Directory
Access start date: 2017-08-23
Access start time: 15:06:32
Last updated timestamp: 2017-06-07 02:50:10
Age time: 30
```

Identity Management

show services user-identification authentication-table authentication-source identity-management brief

Output of this command displays brief user authentication and identity information for all users for whom there are entries in the identity-management authentication source.

```
user@host> show services user-identification authentication-table authentication-source identity-management brief
Domain: ad-domaine-example.net
Total entries: 5
Source IP       Username       groups(Ref by policy)           state
198.51.100.63    N/A                                            Valid
203.0.113.30     administrator                                  Valid
203.0.113.18     N/A                                            Valid
198.51.100.69    N/A                                            Valid
198.51.100.66    administrator                                  Valid

Domain: NULL
```

```
Total entries: 1
Source IP       Username       groups(Ref by policy
```

show services user-identification authentication-table authentication-source identity-management extensive

Output of this command displays extensive user authentication and identity information for all users for whom there are entries in the identity-management authentication source.

```
user@host> show services user-identification authentication-table authentication-source identity-
management extensive
Domain: ad-domain2-example.net
Total entries: 5
  Source-ip: 198.51.100.63
    Username: N/A
    Groups:posture-healthy
    State: Valid
    Source: JIMS - Active Directory
    Access start date: 2017-06-05
    Access start time: 09:28:45
    Last updated timestamp: 2017-06-06 08:41:56
    Age time: 0
  Source-ip: 198.51.100.66
    Username: administrator
    Groups:posture-healthy, group policy creator owners, enterprise admins, schema admins,
domain admins,
    administrators, denied rodc password replication group
    State: Valid
    Source: JIMS - Active Directory
    Access start date: 2017-06-05
    Access start time: 09:23:44
    Last updated timestamp: 2017-06-06 08:11:45
    Age time: 0
```

show services user-identification authentication-table authentication-source identity-management tenant tn1 extensive

Output of this command, which specifies the extensive option, shows state and access information for all entries.

```
user@host> show services user-identification authentication-table authentication-source identity-
management tenant tn1 extensive
node0:
--------------------------------------------------------------------------
Logical System: root-logical-system

Domain: ad03.net
Total entries: 4
  Source-ip: 12.0.0.15
    Username: administrator
    Groups:posture-healthy, admin, group policy creator owners, domain admins, enterprise
admins, schema admins, administrators, denied rodc password replication group
    State: Valid
    Source: JIMS - Active Directory
    Access start date: 2017-12-05
    Access start time: 09:36:30
    Last updated timestamp: 2017-12-04 15:45:51
    Age time: 0
  Source-ip: 3000::12
    Username: jasonlee
    Groups:posture-healthy, domain users, users, group1
    State: Valid
    Source: JIMS - Active Directory
    Access start date: 2017-12-05
    Access start time: 09:36:30
    Last updated timestamp: 2017-12-04 15:46:46
    Age time: 0
  Source-ip: 3000::5
    Username: N/A
    Groups:posture-healthy
    State: Valid
    Source: JIMS - Active Directory
    Access start date: 2017-12-05
    Access start time: 09:36:30
    Last updated timestamp: 2017-12-04 16:01:18
    Age time: 0
```

```
  Source-ip: fe80::342c:302b:6cb4:e109
    Username: N/A
    Groups:posture-healthy
    State: Valid
    Source: JIMS - Active Directory
    Access start date: 2017-12-05
    Access start time: 09:36:30
Last updated timestamp: 2017-12-04 16:01:14
Age time: 0
```

Firewall Authentication Forced Age Timeout

Output shows the “Forced Age timeout” value is displayed when the firewall authentication forced timeout function is configured, but only for when the extensive option is used. The value shows the remaining time left based on the forced timeout setting.

show services user-identification authentication-table authentication-source all extensive

```
user@host> show services user-identification authentication-table authentication-source all
extensive
Domain: ad-userfw.net
Total entries: 1
  Source-ip: 198.51.100.98
    Username: administrator
    State: Valid
    Source: firewall-authentication
    Access start date: 2016-10-27
    Access start time: 09:30:27
    Age time: 30
    Forced Age time:  30/180
```

Release Information

Command introduced in Junos OS Release 12.

Support for Aruba ClearPass added in Junos OS release 12.3X48-D30.

Support added for identity-management as an authentication source in Junos OS Release 15.1X49-D100.

Support added for logical-system for `authentication-source all` in Junos OS Release 18.2R1.

Support added for tenant system for `authentication-source identity management` in Junos OS Release 19.1R1.

show services user-identification device-information table

IN THIS SECTION

- Syntax | 26421
- Description | 26421
- Options | 26422
- Required Privilege Level | 26422
- Output Fields | 26422
- Sample Output | 26424
- Release Information | 26425

Syntax

```
show services user-identification device-information table (all | brief | domain name|
extensive) | device-id device-id ( brief |domain name |extensive) | ip-address ip-address
(logical-system lsys-name | tenant tenant-name | all-logical-systems-tenants | root-logical-
system)
```

Description

Display the contents of the device identity authentication table. The device identity authentication table includes entries for authenticated devices whose information is obtained from external authentication sources. A device identity entry contains the device's IP address, the device ID, and a list of groups that the device belongs to. It also contains attributes that are configured in the device identity profile—for example, the type of device, the vendor, and the operating system that is running on the device and its version.

The device identity authentication table is separate from the Active Directory authentication table or any other local authentication table that is used for Junos OS features, or for specific third-party authentication sources. Also, unlike local user authentication tables, which are specific to an

authentication source, the device identity authentication table holds device identity information for devices authenticated by different sources.

Only one authentication source, such as Active Directory, can be active at a time. A result of this requirement is that there is less demand on the system to process information.

Options

all Display information for all authenticated devices with entries in the table.

device-id Display information for the authenticated device whose device ID is specified.

ip-address Display information for the authenticated device whose IP address is specified.

brief Display terse information for the entries in the device identity authentication table entries. You can specify **brief** as a keyword to the parameters **all** and **device-id**.

domain Display the name of domain and information for all authenticated devices that belong to the domain. You can specify **domain** as a keyword to the parameters **all** and **device-id**.

extensive Display**extensive** information for all of the authenticated devices for which there are table entries. It displays the domain name, the IP address of the device, the device's ID, the device category and vendor, the device type, and the operating system running on the device and its version.

Required Privilege Level

view

Output Fields

Table 1558: show services user-identification device-information table Output Fields

Field Name	Field Description
`Domain name`	The name of the domain to which the devices belong.

NOTE: For each authenticated device, the following information is displayed when the parameter **all** is specified after **table** and it is modified by the keyword **extensive**.

Table 1558: show services user-identification device-information table Output Fields *(Continued)*

Field Name	Field Description
`Source IP address`	The IP address of the device.
Device ID	The ID assigned to the device.
Device-Groups	The groups to which the device belongs.
device-category	The kind of device. For example, the device might be a laptop. You configured this value as part of the device identity profile.
device-vendor	The maker of the device. For example, the device vendor might be Lenovo.
device-type	The device type. If this device is a laptop made by Lenovo, it might be of type thinkpad-t430.
device-os	The operating system that is running on the device. The operating system might be Windows.
device-os-version	The version of the operating system running on the device. For example, for Windows, this might be 7.1.
Location1	The location where the device is being used. The location might be specified as United States.
Referred by	The security policy that refers to the device in its `source-end-user-profile` field. The `source-end-user-profile` that you configure might pertain to a group of devices or a single device.

Sample Output

show services user-identification device-information table

```
user@host> show services user-identification device-information table all extensive
Domain: example.net
Total entries: 3
  Source IP:192.0.2.11
    Device ID: dev01
    Device-Groups: device_group01, device_group02, device_group03, device_group04, device_group05
    device-category: laptop
    device-vendor: lenovo
    device-type: thinkpad-t430
    device-os: windows
    device-os-version: 7.1
    Location1: us1
    Referred by: My-pf_0
  Source IP: 192.0.2.12
    Device ID: dev02
    Device-Groups: device_group06, device_group07, device_group08, device_group09, device_group10
    device-category: laptop
    device-vendor: lenovo
    device-type: thinkpad-t430
    device-os: windows
    device-os-version: 7.1
    Location1: us1
    Referred by: My-pf_0
  Source IP: 192.0.2.14
    Device ID: dev03
    Device-Groups: device_group01, device_group02, device_group03, device_group04, device_group05
    device-category: laptop
    device-vendor: lenovo
    device-type: thinkpad-t430
    device-os: windows
    device-os-version: 7.1
    Location1: us1
    Referred by: My-pf_0

user@host> show services user-identification device-information table all
Domain: example.net
Total entries: 1
```

```
Source IP       Device ID      Device-Groups
2001:db8::1:1   dev04          device-group08
```

show services user-identification device-information table all extensive

```
user@host> show services user-identification device-information table all extensive
Domain: jims-dom1.local
Total entries: 1
Source IP: 2001:db8:4136:e378:8000:63bf:3fff:fdd2
Device ID: win-test$
Device-Groups: dev, pre-windows 2000 compatible access, cert publishers,
denied rodc password replication group
device-os: windows server 2012 r2 standard evaluation
device-os-version: 6.3 (9600)
Referred by: p1
```

show services user-identification device-information table all

```
user@host> show services user-identification device-information table all
example.net
Total entries: 1
Source IP                                Device ID Device-Groups
2001:db8:4136:e378:8000:63bf:3fff:fdd2  dev04     device-group08
```

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

Understanding the Device Identity Attributes and Profiles for the Integrated User Firewall Device Identity Authentication Feature

Understanding the Device Identity Authentication Table and Its Entries

Understanding Access Control to Network Resources Based on Device Identity Information

authentication-source (Services User Identification Device Identity) | **2662**

source-end-user-profile | **28204**

show services user-identification logical-domain-identity-management

IN THIS SECTION

- Syntax | 26426
- Description | 26426
- Options | 26426
- Required Privilege Level | 26426
- Output Fields | 26427
- Sample Output | 26428
- Release Information | 26429

Syntax

```
show services user-identification logical-domain-identity-management
<counters>
<status>
```

Description

Displays the information about the logical domain identity-management.

Options

counters Displays the logical domain identity management query counters.

status Displays logical domain identity management query status.

Required Privilege Level

view

Output Fields

Table 1559 on page 26427 lists the output fields for the `show services user-identification logical-domain-identity-management` command. Output fields are listed in the approximate order in which they appear.

Table 1559: show services user-identification logical-domain-identity-management Output Fields

Field Name	Field Description
Node	Node (device) in the Juniper Identity Management Service server (node0 or node1).
Query server name	A Juniper Identity Management Service server stores the records for a domain name and responds to queries from clients based on these records.
For primary server:	
Address	For the status option, the IP address of the primary server.
Port	The name of the port.
Connection method	The SRX Series Firewall that connects to the Juniper Identity Management Service server to obtain user identity information.
Connection status	The status of SRX Series Firewall connection to the Juniper Identity Management Service server.
Last received status message	The last message received.
Access token	Token string.
Token expiry time	The access token expiry time.
For secondary server:	
Address	For the status option, the IP address of the secondary server.

Table 1559: show services user-identification logical-domain-identity-management Output Fields *(Continued)*

Field Name	Field Description
Batch query sent number	A number indicating how many batch queries the SRX Series Firewall sent to the Juniper Identity Management Service server.
Batch query total response number	A number indicating how many responses the SRX Series Firewall received from the Juniper Identity Management Service server in response to its batch queries.
Batch query error response number	A number indicating how many error responses the SRX Series Firewall received from the Juniper Identity Management Service server in response to its batch queries.
Batch query last response time	Timestamp indicating when the last response was received.
IP query sent number	A number indicating how many IP queries the SRX Series Firewall sent to the Juniper Identity Management Service server.
IP query total response number	A number indicating how many responses the SRX Series Firewall received from the Juniper Identity Management Service server in response to its IP queries.
IP query error response number	A number indicating how many error responses the SRX Series Firewall received from the Juniper Identity Management Service server in response to its IP queries.
IP query last response time	Timestamp indicating when the last response was received.

Sample Output

show services user-identification logical-domain-identity-management status

```
user@host> show services user-identification logical-domain-identity-management status
node0:
--------------------------------------------------------------------------
```

```
Query server name                 :jims1
   Primary server :
     Address                      : 192.0.2.0/24
     Port                         : 443
     Connection method            : HTTPS
     Connection status            : Online
     Last received status message : OK (200)
     Access token                 : isdHIbl8BXwxFftMRubGVsELRukYXtW3rtKmHiL
     Token expire time            : 2017-11-27 23:45:22
  Secondary server :
     Address                      : Not configured
```

show services user-identification logical-domain-identity-management counters

```
user@host> show services user-identification logical-domain-identity-management counters
node0:
--------------------------------------------------------------------------
Query server name                      :jims1
  Primary server :
     Address                           : 203.0.113.0/24
     Batch query sent number           : 65381
     Batch query total response number : 64930
     Batch query error response number : 38
     Batch query last response time    : 2018-08-14 15:10:52
     IP query sent number              : 10
     IP query total response number    : 10
     IP query error response number    : 0
     IP query last response time       : 2018-08-13 12:41:56
  Secondary server :
 Address                            : Not configured
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

clear services user-identification logical-domain-identity-management counters | 4197

show services user-identification validate-statistics

IN THIS SECTION

- Syntax | 26430
- Description | 26430
- Required Privilege Level | 26430
- Output Fields | 26430
- Sample Output | 26431
- Release Information | 26431

Syntax

```
show services user-identification validate-statistics
```

Description

Displays the processed validating count for user or device or group.

Required Privilege Level

view

Output Fields

Table 1560 on page 26430 lists the output fields for the `show services user-identification validate-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1560: show services user-identification validate-statistics

Field Name	Field Description
`Count`	The total number of processed validating count.

Table 1560: show services user-identification validate-statistics *(Continued)*

Field Name	Field Description
Device	The total number of processed validating devices.
Group	The total number of processed validating groups.
User	The total number of processed validating users.

Sample Output

command-name

```
user@host> show services user-identification validate-statistics
node0:
--------------------------------------------------------------------------
            Processed validating count  : 157
                    User                : 77
                    Device              : 16
                    Group               : 64
```

Release Information

Command introduced in Junos OS Release 20.2R1.

show services video-monitoring mdi errors fpc-slot

IN THIS SECTION

- Syntax | 26432
- Description | 26432
- Options | 26432

- Required Privilege Level | 26432
- Output Fields | 26432
- Sample Output | 26433
- Release Information | 26433

Syntax

```
show services video-monitoring mdi errors fpc-slot fpc-slot
```

Description

Display video monitoring error statistics.

Options

fpc-slot Number of the fpc slot for which statistics are displayed.

Required Privilege Level

view

Output Fields

Table 1561 on page 26432 lists the output fields for the `show services video-monitoring mdi errors fpc-slot` command. Output fields are listed in the approximate order in which they appear.

Table 1561: show services video-monitoring mdi errors fpc-slot Output Fields

Field Name	Field Description
`FPC slot`	Slot number of the monitored FPC.
`Flow Insert Error`	Number of errors during new flow insert operations.

Table 1561: show services video-monitoring mdi errors fpc-slot Output Fields *(Continued)*

Field Name	Field Description
`Flow Policer Drops`	Number of packets dropped by flow policer process. **NOTE**: New flows usually arrive within a very short time interval (1.5 microseconds). These errors do not represent the loss of entire flows, because subsequent packets in the flow can establish the flow. All packets are monitored after a flow has been established. Packet forwarding occurs independently of the video monitoring, and packets are not dropped due to video monitoring errors.
`Unsupported Media Packets Count`	Number of packets dropped because they are not media packets or they are unsupported media packets.
`PID Limit Exceeded`	Number of packets unmonitored because the process identifier (PID) limit exceeded has been exceeded. **NOTE**: The current PID limit is 6.

Sample Output

show services video-monitoring mdi errors fpc-slot

```
user@host> show services video-monitoring mdi errors fpc-slot 2
 MDI Errors Information
    FPC Slot: 2
    Flow Insert Error: 0, Flow Policer Drops: 0
    Unsupported Media Packets Count: 0, PID Limit Exceeded: 202995
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Inline Video Monitoring on MX Series Routers

show services video-monitoring mdi flows fpc-slot

IN THIS SECTION

- Syntax | 26434
- Description | 26435
- Options | 26435
- Required Privilege Level | 26436
- Output Fields | 26436
- Sample Output | 26438
- Sample Output | 26438
- Sample Output | 26439
- Release Information | 26441

Syntax

```
show services video-monitoring mdi flows fpc-slot fpc-slot
<brief>
<count>
<destination-address>
<destination-port>
<detail>
<df-mlr-split-view>
<flow-over-ipv4 | flow-over-ipv4-over-mpls | flow-over ipv6 | flow-over-ipv6-mpls>
<input>
<interface-name>
<output>
<rtp>
<source-address>
<source-port>
<template-name>
<udp>
```

Description

Display inline video monitoring flow statistics.

Options

fpc-slot	Number of the slot for which flows are reported.
brief	(Optional) Display brief output (default).
count	(Optional) Display the number of flows.
destination-address	(Optional) Filter output by destination address.
destination-port	(Optional) Filter output by destination port.
detail	(Optional) Display output in detailed format including media delivery index records.
df-mlr-split-view	(Optional) Display detailed/brief flow output with DF and MLR split.
flow-over-ipv4 | flow-over-ipv4-over-mpls | flow-over ipv6 | flow-over-ipv6-mpls	(Optional) Display only IPv4 flows, only IPv4-over-MPLS flows, only IPv6 flows, or only IPv6-over-MPLS flows.
input	(Optional) Filter output by flow direction input.
interface-name	(Optional) Filter output by logical interface name.
output	(Optional) Filter output by flow direction output.
rtp	(Optional) Filter output by flow type rtp.
source-address	(Optional) Filter output by source IP address.
source-port	(Optional) Filter output by source port.
template-name	(Optional) Filter output by media delivery index template name.
udp	(Optional) Filter output by flow type MPEG-TS.
df-mlr-split-view	(Optional) Display detailed/brief flow output with DF and MLR split.

Required Privilege Level

view

Output Fields

Table 1562 on page 26436 lists the output fields for the `show services video-monitoring mdi flows fpc-slot` command. Output fields are listed in the approximate order in which they appear.

Table 1562: show services video-monitoring mdi flows fpc-slot Output Fields

Field Name	Field Description
`SIP`	Source IP address.
`DIP`	Destination IP address.
`SP`	Source port.
`DP`	Destination port.
`Di`	Direction of flow (I=Input, O=Output).
`Ty`	Type of flow.
`Last DF:MLR`	Delay factor and media loss rate value of last media delivery index record. **NOTE**: If you choose `df-mlr-split-view` option, then the DF and MLR values will be displayed in split format as: `Last DF:`*<xxx>*, `Last MLR:`*<xxx>*, `Avg DF:`*<xxx>*, `Avg MLR:`*<xxx>*
`Avg DF:MLR`	Average value of delay factor and media loss rate.
`Last MRV`	Media rate variation value of last media delivery index record.
`Avg MRV`	Average value of media rate variation.

Table 1562: show services video-monitoring mdi flows fpc-slot Output Fields *(Continued)*

Field Name	Field Description
IFL	Interface name on which flow is received.
Template Name	Name of template associated with flow.
Flow Identifier	Identifier for the flow.
Source Address	Source IP address.
Destination Address	Destination IP address.
Interface Name	Interface name on which flow is received.
Flow Direction	Direction of flow (I=Input, O=Output).
Flow Type	Flow type that is monitored. The value is one of the following: • RTP • RTP over IPv6 • RTP over IPv4-over-MPLS • RTP over IPv6-over-MPLS • UDP • UDP over IPv6 • UDP over IPv4-over-MPLS • UDP over IPv6-over-MPLS
Rec No	Record number.
PID	Process identifier.

Table 1562: show services video-monitoring mdi flows fpc-slot Output Fields *(Continued)*

Field Name	Field Description
MLR	Media loss rate.

Sample Output

show services video-monitoring mdi flows fpc-slot brief

```
user@host> show services video-monitoring mdi flows fpc-slot 2 brief
-----------------------------------------------------------------------------------------------
-----------------------------------------
Sno |SIP      |SP  |DIP         |DP  |Di|Ty |Last DF:MLR |Avg DF:MLR |Last MRV |Avg MRV |
IFL         |Template Name |Flow Identifier

-----------------------------------------------------------------------------------------------
-----------------------------------------
1   192.0.2.2 1024 198.51.100.2 2048 I  UDP  70.90:1      92.15:8205  -7.09      -9.36
xe-2/2/1.0        t1            16777216
```

Sample Output

show services video-monitoring mdi flows fpc-slot 0 detail

```
user@host> show services video-monitoring mdi flows fpc-slot 0 detail
Source Address: 192.0.2.22, Source Port: 1024
    Destination Address: 203.0.113.1, Destination Port: 2060
    Last DF:MLR: 3.56:0, Avg DF:MLR: 3.60:0
    Last MRV: 0.00, Avg MRV: 0.00
    Interface Name: ge-0/3/4.0, Template Name: t1
    Flow Direction: Input, Flow Type: RTP
    MDI Records Count: 10
    Flow Identifier: 16777216
---------------------------------------+
Rec No|        DF|       MLR|       MRV|
---------------------------------------+
```

```
     1       3.61          0       0.00
     2       3.64          0       0.00
     3       3.58          0       0.00
     4       3.62          0       0.00
     5       3.57          0       0.00
     6       3.60          0       0.00
     7       3.63          0       0.00
     8       3.58          0       0.00
     9       3.61          0       0.00
    10       3.56          0       0.00

Source Address: 192.0.2.22, Source Port: 1024
    Destination Address: 203.0.113.1, Destination Port: 2060
    Last DF:MLR: 3.57:0, Avg DF:MLR: 3.60:0
    Last MRV: 0.00, Avg MRV: -0.04
    Interface Name: ge-0/2/2.0, Template Name: t1
    Flow Direction: Output, Flow Type: RTP over IPv4-over-MPLS
    MPLS Labels: (299776,16,0)
    MDI Records Count: 10
    Flow Identifier: 16777217
---------------------------------------+
Rec No|        DF|       MLR|       MRV|
---------------------------------------+
     1       3.59          0       0.00
     2       3.62          0       0.00
     3       3.57          0       0.00
     4       3.60          0       0.00
     5       3.64          0       0.00
     6       3.58          0       0.00
     7       3.62          0       0.00
     8       3.57          0      -0.35
     9       3.62          0       0.00
    10       3.57          0       0.00
```

Sample Output

show services video-monitoring mdi flows fpc-slot 2 detail

```
user@host> show services video-monitoring flows fpc-slot 2 detail count 19
```

Format for RTP flows:

```
 Source Address: 192.0.2.2, Source Port: 1024
    Destination Address: 198.51.100.2, Destination Port: 2048
    Last DF:MLR: 3.58:0, Avg DF:MLR: 3.60:0
    Last MRV: 0.00, Avg MRV: 0.00
    Interface Name: xe-2/2/1.0, Template Name: t1
    Flow Direction: Input, Flow Type: RTP, MDI Records Count: 10
---------------------------------------+
Rec No|        DF|       MLR|       MRV|
---------------------------------------+
     1       3.58          0       0.00
     2       3.62          0       0.00
     3       3.59          0       0.00
     4       3.63          0       0.00
     5       3.60          0       0.00
     6       3.64          0       0.00
     7       3.61          0       0.00
     8       3.57          0       0.00
     9       3.62          0       0.00
    10       3.58          0       0.00
```

Format for MPEG2-TS over UDP flows:

```
    Source Address: 192.0.2.2, Source Port: 1024
    Destination Address: 198.51.100.2, Destination Port: 2048
    Last DF:MLR: 3.63:0, Avg DF:MLR: 3.61:4097
    Last MRV: 0.00, Avg MRV: 0.00
    Interface Name: xe-2/2/1.0, Template Name: t1
    Flow Direction: Input, Flow Type: UDP, MDI Records Count: 10
---------------------------------------
+-------------------------------------------------------------------------------------------+
Rec No|        DF|       MLR|       MRV|    PID-0      |    PID-1      |    PID-2      |
PID-3      |    PID-4      |    PID-5      |
---------------------------------------
+-------------------------------------------------------------------------------------------+
                                      |    Val|    MLR|    Val|    MLR|    Val|    MLR|
Val|    MLR|    Val|    MLR|    Val|    MLR|
-------------------------------------------------------------------------------------------
---------------------------------------+
```

```
     1       3.63          0       0.00  0x1f40       0  0x1f41       0    0x12       0
0x1f54       0    0x11       0  0x1020       0
     2       3.59          0       0.00  0x1f40       0  0x1f41       0    0x12       0
0x1f54       0    0x11       0  0x1020       0
     3       3.64          0       0.00  0x1f40       0  0x1f41       0    0x12       0
0x1f54       0    0x11       0  0x1020       0
     4       3.60          0       0.00  0x1f40       0  0x1f41       0    0x12       0
0x1f54       0    0x11       0  0x1020       0
     5       3.64          0       0.00  0x1f40       0  0x1f41       0    0x12       0
0x1f54       0    0x11       0  0x1020       0
     6       3.61          0       0.00  0x1f40       0  0x1f41       0    0x12       0
0x1f54       0    0x11       0  0x1020       0
     7       3.57          0       0.00  0x1f40       0  0x1f41       0    0x12       0
0x1f54       0    0x11       0  0x1020       0
     8       3.62          0       0.00  0x1f40       0  0x1f41       0    0x12       0
0x1f54       0    0x11       0  0x1020       0
     9       3.58      40977       0.00  0x1f40   40977  0x1f41       0    0x12       0
0x1f54       0    0x11       0  0x1020       0
    10       3.63          0       0.00  0x1f40       0  0x1f41       0    0x12       0
0x1f54       0    0x11       0  0x1020       0
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Inline Video Monitoring on MX Series Routers

show services video-monitoring mdi stats fpc-slot

IN THIS SECTION

- Syntax | 26442
- Description | 26442
- Options | 26442

- Required Privilege Level | 26442
- Output Fields | 26442
- Sample Output | 26444
- Release Information | 26444

Syntax

```
show services video-monitoring mdi stats fpc-slot fpc-slot
```

Description

Display inline video monitoring statistics.

Options

fpc-slot Number of the fpc slot for which statistics are displayed.

Required Privilege Level

view

Output Fields

Table 1563 on page 26442 lists the output fields for the `show services video-monitoring mdi stats fpc-slot` command. Output fields are listed in the approximate order in which they appear.

Table 1563: show services video-monitoring mdi stats fpc-slot Output Fields

Field Name	Field Description
`FPC Slot`	Slot number of the monitored FPC

Table 1563: show services video-monitoring mdi stats fpc-slot Output Fields *(Continued)*

Field Name	Field Description
`Active Flows`	Number of active flows currently monitored. active flows = inserted flows - deleted flows.
`Total Inserted Flows`	Number of flows initiated under video monitoring.
`Total Deleted Flows`	Number of flows deleted due to inactivity timeout.
`Total Packets Count`	Number of total packets monitored.
`Total Bytes Count`	Number of total bytes monitored.
`DF Alarm Count`	Number of delay factor alarms at each of the following levels: • `Info level` • `Warning level` • `Critical level`
`MLR Alarm Count`	Number of media loss rate (MLR) alarms at each of the following levels: • `Info level` • `Warning level` • `Critical level`
`MRV alarm count`	Number of media rate variation (MRV) alarms at each of the following levels: • `Info level` • `Warning level` • `Critical level`

Sample Output

show services video-monitoring mdi stats fpc-slot

```
user@host> show services video-monitoring mdi stats fpc-slot 2

May 13 22:51:38  blurr fpc2 user.crit aftd-trio: [MDI] MRV: -28.90, exceeded critical threshold
for flow (flow_id:16779635 src:4.4.4.2 dst:7.7.7.2 sport:1000 dport:9000) egressing from
interface index(345) with template mpls_ipv4_layer.
May 13 22:51:42  blurr fpc2 user.crit aftd-trio: [MDI] DF: 292.58 ms, exceeded critical
threshold for flow(flow_id:16777352 src:4.4.4.3 dst:7.7.7.2 sport:1000 dport:9000) egressing
from interface index(345) with template mpls_ipv4_layer.

May 13 23:30:18  blurr fpc2 user.warning aftd-trio: [MDI] MRV: -28.60, exceeded warning
threshold for flow (flow_id:16779635 src:4.4.4.2 dst:7.7.7.2 sport:1000 dport:9000) egressing
from interface index(345) with template mpls_ipv4_layer.
May 13 23:30:23  blurr fpc2 user.warning aftd-trio: [MDI] DF: 289.57 ms, exceeded warning
threshold for flow(flow_id:16777352 src:4.4.4.3 dst:7.7.7.2 sport:1000 dport:9000) egressing
from interface index(345) with template mpls_ipv4_layer.

May 13 23:32:58  blurr fpc2 user.info aftd-trio: [MDI] MRV: -29.30, exceeded info threshold for
flow (flow_id:16779635 src:4.4.4.2 dst:7.7.7.2 sport:1000 dport:9000) egressing from interface
index(345) with template mpls_ipv4_layer.
May 13 23:33:03  blurr fpc2 user.info aftd-trio: [MDI] DF: 291.26 ms, exceeded info threshold
for flow(flow_id:16777352 src:4.4.4.3 dst:7.7.7.2 sport:1000 dport:9000) egressing from
interface index(345) with template mpls_ipv4_layer.
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Inline Video Monitoring on MX Series Routers

show services web-filter dns-resolution profile

IN THIS SECTION

- Syntax | 26445
- Description | 26445
- Options | 26445
- Required Privilege Level | 26446
- Output Fields | 26446
- Sample Output | 26446
- Release Information | 26448

Syntax

```
show services web-filter dns-resolution profile profile-name <template template-name>
<fpc-slot fpc-slot pic-slot pic-slot>
```

Description

Display URL filter domain name system (DNS) resolution information.

URL filtering resolves the disallowed domains. The total number of domains are divided into chunks of 50 domains per chunk. The `filter term` in the command output is the name of a chunk.

Options

`fpc-slot` *`fpc-slot`* `pic-slot` *`pic-slot`*	(Optional) Specify the FPC and PIC for which you want URL filter information displayed.
`profile` *`profile-name`*	Specify the profile for which you want URL filter information displayed.
`template` *`template-name`*	(Optional) Specify the template for which you want URL filter information displayed.

Required Privilege Level

view

Output Fields

Table 1564 on page 26446 lists the output fields for the `show services web-filter dns-resolution profile` command. Output fields are listed in the approximate order in which they appear.

Table 1564: show services web-filter dns-resolution profile Output Fields

Field Name	Field Description
`Profile`	Name of profile.
`Template`	Name of template.
`Filter Term`	Name of the domains chunk. All domains are divided into chunks of 50 domains per chunk.
`IPv4 Address Count`	The number of IPv4 addresses resolved for all domains under the filter term.
`IPv6 Address Count`	The number of IPv6 addresses resolved for all domains under the filter term.
`Domain Name`	Name of domain.
`IPv4 Records`	Listing of IPv4 addresses.
`IPv6 Records`	Listing of IPv6 addresses.

Sample Output

show services web-filter dns-resolution profile

```
user@host> show services web-filter dns-resolution profile p1
URL filtering DNS resolution:
Profile: p1
```

```
Template: t1

 1). Filter Term: URLF_t1_0004

      IPv4 Address Count: 20
      IPv6 Address Count: 20

    1 ). Domain Name: www.example.com

          IPv4 Records:
             31.13.77.36
             31.13.76.68

          IPv6 Records:
             2a03:2880:f122:83:face:b00c:0:25de
             2a03:2880:f111:83:face:b00c:0:25de

    2 ). Domain Name: www.youtube.com

          IPv4 Records:
             216.58.193.78
             216.58.194.206

          IPv6 Records:
             2607:f8b0:400a:800::200e
             2607:f8b0:4005:809::200e

    3 ). Domain Name: www.netflix.com

          IPv4 Records:
             50.112.200.248
             52.10.96.2
             52.25.242.211
             52.39.87.182
             52.38.44.92
             52.36.125.176
             52.40.2.42
             52.42.184.64
             52.5.80.199
             52.206.203.18
             52.5.231.14
             52.21.94.89
             52.71.118.87
```

```
            52.201.133.109
            52.71.122.233
            52.203.136.33

         IPv6 Records:
            2620:108:700f::342a:b840
            2620:108:700f::3644:fc64
            2620:108:700f::3459:2ce1
            2620:108:700f::3459:c025
            2620:108:700f::3459:f556
            2620:108:700f::3459:c5c5
            2620:108:700f::3644:c2a0
            2620:108:700f::342a:df11
            2406:da00:ff00::3404:d29c
            2406:da00:ff00::3415:a86e
            2406:da00:ff00::3415:fda4
            2406:da00:ff00::3414:91d2
            2406:da00:ff00::3403:73dd
            2406:da00:ff00::22c7:d016
            2406:da00:ff00::3400:290b
            2406:da00:ff00::3213:c65f
```

Release Information

Command introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

show services web-filter dns-resolution-statistics profile template | **26449**

show services web-filter statistics profile | **26464**

Configuring URL Filtering

show services web-filter dns-resolution-statistics profile template

IN THIS SECTION

- Syntax | 26449
- Description | 26449
- Options | 26449
- Required Privilege Level | 26449
- Output Fields | 26450
- Sample Output | 26452
- Release Information | 26455

Syntax

```
show services web-filter dns-resolution-statistics profile profile-name template template-name
(extensive | summary)
```

Description

Display URL filter domain name system (DNS) resolution statistics.

Options

(extensive | summary)	Specify the level of detail of information you want displayed.
profile *profile-name*	Specify the profile for which you want URL filter information displayed.
template *template-name*	Specify the template for which you want URL filter information displayed.

Required Privilege Level

view

Output Fields

Table 1565 on page 26450 lists the output fields for the `show services web-filter dns-resolution-statistics profile template` command. Output fields are listed in the approximate order in which they appear.

Table 1565: show services web-filter dns-resolution-statistics profile template Output Fields

Field Name	Field Description	Level of Detail
`Profile`	Name of profile.	all
`Template`	Name of template.	all
`DNS start time`	Start time of the DNS resolution.	`summary`
`Next DNS start time`	Start time of the next DNS resolution.	`summary`
`Number of resolved A addresses`	Number of resolved IPv4 addresses.	`summary`
`Number of resolved AAAA addresses`	Number of resolved IPv6 addresses.	`summary`
`Number of unresolved A addresses`	Number of unresolved IPv4 addresses.	`summary`
`Number of unresolved AAAA addresses`	Number of unresolved IPv6 addresses.	`summary`
`Number of resolved A domains`	Number of resolved IPv4 domains.	`summary`
`Number of resolved AAAA domains`	Number of resolved IPv6 domains.	`summary`
`Number of unresolved A domains`	Number of unresolved IPv4 domains.	`summary`

Table 1565: show services web-filter dns-resolution-statistics profile template Output Fields *(Continued)*

Field Name	Field Description	Level of Detail
`Number of unresolved AAAA domains`	Number of unresolved IPv6 domains.	`summary`
`Number of requests sent`	Number of DNS requests sent.	`summary`
`Number of responses received`	Number of DNS responses received.	`summary`
`Domain Name`	Name of domain.	`extensive`
`IPv4 Address information`	IPv4 address information includes the following fields: • `DNS server IP`—IPv4 address of DNS server. • `Req Sent`—Number of DNS requests sent. • `Resp Received`—Number of DNS responses received. • `DNS retries`—Number of times no DNS response was received and so retried.	`extensive`

Table 1565: show services web-filter dns-resolution-statistics profile template Output Fields *(Continued)*

Field Name	Field Description	Level of Detail
`IPv6 Address information`	IPv6 address information includes the following fields: • `DNS server IP`—IPv6 address of DNS server. • `Req Sent`—Number of DNS requests sent. • `Resp Received`—Number of DNS responses received. • `DNS retries`—Number of times no DNS response was received and so retried.	`extensive`

Sample Output

show services web-filter dns-resolution-statistics profile template summary

```
user@host> show services web-filter dns-resolution-statistics profile1 template t1 summary
URL filtering DNS resolution statistics:
Profile: p1
Template: t1

     DNS start time                      : May 01 16:40:24 PDT
     Next DNS start time                 : May 01 17:40:24 PDT
     Number of resolved A domains        : 114
     Number of resolved AAAA domains     : 114
     Number of unresolved A domains      : 0
     Number of unresolved AAAA domains   : 0
     Number of requests sent             : 246
     Number of responses received        : 228
```

show services web-filter dns-resolution-statistics profile template extensive

```
user@host> show services web-filter dns-resolution-statistics profile p1 template t1 extensive
URL filtering DNS resolution statistics:
Profile: p1
Template: t1

     1) Domain Name:    www.facebook.com

         IPv4 Address information:
              DNS server IP     8.8.8.8
              Req Sent          20
              Resp Received     20
              DNS retries       0

         IPv4 Address information:
              DNS server IP     172.29.131.60
              Req Sent          21
              Resp Received     20
              DNS retries       0

         IPv6 Address information:
              DNS server IP     8.8.8.8
              Req Sent          25
              Resp Received     20
              DNS retries       0

         IPv6 Address information:
              DNS server IP     172.29.131.60
              Req Sent          24
              Resp Received     20
              DNS retries       0
     2) Domain Name:    www.youtube.com

         IPv4 Address information:
              DNS server IP     8.8.8.8
              Req Sent          21
              Resp Received     20
              DNS retries       0

         IPv4 Address information:
              DNS server IP     172.29.131.60
```

```
            Req Sent          21
            Resp Received     20
            DNS retries       0

        IPv6 Address information:
            DNS server IP     8.8.8.8
            Req Sent          21
            Resp Received     20
            DNS retries       0

        IPv6 Address information:
            DNS server IP     172.29.131.60
            Req Sent          21
            Resp Received     20
            DNS retries       0
    3) Domain Name:    www.netflix.com

        IPv4 Address information:
            DNS server IP     8.8.8.8
            Req Sent          21
            Resp Received     20
            DNS retries       0

        IPv4 Address information:
            DNS server IP     172.29.131.60
            Req Sent          21
            Resp Received     20
            DNS retries       0

        IPv6 Address information:
            DNS server IP     8.8.8.8
            Req Sent          21
            Resp Received     20
            DNS retries       0

        IPv6 Address information:
            DNS server IP     172.29.131.60
            Req Sent          21
            Resp Received     20
            DNS retries       0
```

Release Information

Command introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

show services web-filter dns-resolution profile | 26445

show services web-filter statistics profile | 26464

Configuring URL Filtering

show services web-filter secintel-policy status

IN THIS SECTION

- Syntax | 26455
- Description | 26456
- Options | 26456
- Required Privilege Level | 26456
- Sample Output | 26456
- Release Information | 26460

Syntax

```
show services web-filter secintel-policy status
 profile  profile-name
 template  template-name
```

Description

Display the IPv4 and IPv6 count per threat level received from the C&C feed from Policy Enforcer. It also displays the count of the number of terms used in the implicit filter per threat level.

Options

profile-name	Name of the profile
template-name	Name of the template

Required Privilege Level

view

Sample Output

show services web-filter secintel-policy status

```
user@host> show services web-filter secintel-policy status profile
URL Filtering SecIntel Policy Status:
Profile     : Profile1
C&C DB File : /var/db/url-filterd/urlf_si_cc_db.txt
Policy State: Ready
DB File Change Time : Tue Nov 27 11:01:10 2018
DB File Load Time   : Tue Nov 27 11:01:38 2018
C&C Prefix Count    : IPv4: 11093      IPv6: 5
Filters:
Threat level   Action       v4 Term Count    IPv4        v6 Term Count   IPv6

1              ACCEPT       23            1129         1               2
2              ACCEPT       11            1444         0               0
3              ACCEPT       6             996          0               0
4              ACCEPT       7             564          0               0
5              ACCEPT       7             451          0               0
6              ACCEPT       4             126          0               0
7              LOG          5             175          0               0
8              DROP AND LOG 4             396          1               1
9              ACCEPT       2             164          0               0
```

```
10             ACCEPT       33            5601         1               2
```

```
user@host> show services web-filter secintel-policy-status profile Profile1 url-filter-template
template200
Template    : template200
       C&C DB File : /var/db/url-filterd/urlf_si_ip_white_list_db.txt
       Policy State: NA
       DB File Change Time : NA
       DB File Load Time   : NA
       C&C Prefix Count    : IPv4: 0          IPv6: 0

       C&C DB File : /var/db/url-filterd/urlf_si_ip_black_list_db.txt
       Policy State: NA
       DB File Change Time : NA
       DB File Load Time   : NA
       C&C Prefix Count    : IPv4: 0          IPv6: 0

       C&C DB File : /var/db/url-filterd/urlf_si_ip_custom_db.txt
       Policy State: Ready
       DB File Change Time : Tue Feb 04 15:22:20 2020
       DB File Load Time   : Tue Feb 04 15:24:29 2020
       C&C Prefix Count    : IPv4: 16         IPv6: 0
       Filters:
       Threat level   Action              v4 Term Count    IPv4        v6 Term Count   IPv6

       0             ACCEPT AND SAMPLE       0             0              0            0
       255           DROP AND SAMPLE         0             0              0            0
       1             DROP AND SAMPLE         1             11             0            0
       2             ACCEPT                  0             0              0            0
       3             DROP AND SAMPLE         1             1              0            0
       4             DROP AND SAMPLE         1             1              0            0
       5             ACCEPT                  0             0              0            0
       6             ACCEPT                  1             1              0            0
       7             ACCEPT                  1             1              0            0
       8             DROP AND SAMPLE         0             0              0            0
       9             ACCEPT                  1             1              0            0
       10            DROP AND SAMPLE         0             0              0            0
```

show services web-filter secintel-policy status profile

To display GeoIP feed, allowlist and blocklist.

```
user@host> show services web-filter secintel-policy status profile Profile1
URL Filtering SecIntel Policy Status:
Profile     : Profile1
C&C DB File : /var/db/url-filterd/urlf_si_ip_global_db.txt
Policy State: Ready
DB File Change Time : Mon Nov 29 15:24:53 2021
DB File Load Time   : Mon Nov 29 15:25:09 2021
C&C Prefix Count    : IPv4: 151768     IPv6: 1
Filters:
Threat level   Action              v4 Term Count    IPv4        v6 Term Count   IPv6

 1             ACCEPT                  2             518            1            1
 2             ACCEPT                  35            8645           0            0
 3             ACCEPT                  30            7038           0            0
 4             ACCEPT                  41            10985          0            0
 5             ACCEPT                  2             361            0            0
 6             ACCEPT                  390           116291         0            0
 7             ACCEPT                  7             1663           0            0
 8             LOG AND SAMPLE          11            1852           0            0
 9             ACCEPT                  9             520            0            0
 10            ACCEPT                  15            3895           0            0


Global WL DB File : /var/db/url-filterd/urlf_si_ip_white_list_db.txt
DB File Change Time : Wed Nov 24 16:52:28 2021
DB File Load Time   : Mon Nov 29 15:25:09 2021
Global WL Prefix Count  : IPv4: 24         IPv6: 0

Global BL DB File : urlf_si_ip_global_bl_list_db.txt
DB File Change Time : Wed Nov 24 16:52:28 2021
Global BL Prefix Count  : IPv4: 1          IPv6: 0



Template    : template1
       C&C DB File : /var/db/url-filterd/urlf_si_ip_white_list_db.txt
       Policy State: NA
       DB File Change Time : NA
```

```
        DB File Load Time   : NA
        C&C Prefix Count    : IPv4: 0          IPv6: 0

        C&C DB File : /var/db/url-filterd/urlf_si_ip_black_list_db.txt
        Policy State: NA
        DB File Change Time : NA
        DB File Load Time   : NA
        C&C Prefix Count    : IPv4: 0          IPv6: 0

        C&C DB File : /var/db/url-filterd/urlf_si_ip_custom_db.txt
        Policy State: NA
        DB File Change Time : NA
        DB File Load Time   : NA
        C&C Prefix Count    : IPv4: 0          IPv6: 0
        Filters:
        Threat level   Action              v4 Term Count    IPv4        v6 Term Count   IPv6

        0              ACCEPT AND SAMPLE       0             0              0            0
        255            DROP AND SAMPLE         0             0              0            0
        1              ACCEPT                  0             0              0            0
        2              ACCEPT                  0             0              0            0
        3              ACCEPT                  0             0              0            0
        4              ACCEPT                  0             0              0            0
        5              ACCEPT                  0             0              0            0
        6              ACCEPT                  0             0              0            0
        7              ACCEPT                  0             0              0            0
        8              ACCEPT                  0             0              0            0
        9              ACCEPT                  0             0              0            0
        10             ACCEPT                  0             0              0            0


GeoIP    :
        GeoIP DB File : /var/db/url-filterd/urlf_si_ip_geo_db.txt
        Policy State: Ready
        DB File Change Time : Sat Nov 27 18:07:00 2021
        DB File Load Time   : Sat Nov 27 18:09:01 2021
        GeoIP Prefix Count    : IPv4: 382607     IPv6: 247093
        Filters:
        Country        Action              v4 Term Count    IPv4        v6 Term Count   IPv6

        AU             DROP                    1             300            1            300
```

show services web-filter secintel-policy status

To verify if an IP address is part of the GeoIP feed.

```
user@host> show services web-filter secintel-policy-db ip-prefix-information  192.168.1.1/24
profile Profile1
URL Filtering SecIntel Policy DB IP Prefix Info:
Profile     : Profile1
Downloaded Feed Category : GeoIP
Applied Feed Category    : GeoIP
Threat Level             : 255
Threat Level Action      : DROP
Add Time                 : Sat Nov 27 18:06:29 2021
Filter Name              : v4-si-prof-Profile1-gbl-geo-filter
Filter Index             : 201326592
Filter Term Name         : FILTER_TL_255_COUNTRY_AU_ID_71
Pending Delete           : FALSE
```

Release Information

Statement introduced before Junos OS Release 18.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480, and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

security-intelligence (Services) | **16462**

show services web-filter statistics dns-filter-template

IN THIS SECTION

- Syntax | 26461
- Description | 26461
- Options | 26461
- Required Privilege Level | 26461
- Output Fields | 26461
- Sample Output | 26462
- Release Information | 26463

Syntax

```
show services web-filter statistics dns-filter-template template-name
```

Description

Display statistics for DNS request filtering and URL filtering for the specified filter profile.

Options

dns-filter-template *template-name* (Optional) Display statistics for the specified DNS filter template.

Required Privilege Level

view

Output Fields

Table 1566 on page 26462 lists the output fields for the `show services web-filter statistics profile` command. Output fields are listed in the approximate order in which they appear.

Table 1566: show services web-filter statistics profile Output Fields

Field Name	Field Description
UDP DNS	Number of UDP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.
TCP DNS	Number of TCP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.

Sample Output

show services web-filter statistics dns-filter-template

```
user@host> show services web-filter statistics dns-filter-template DNS_CUSTOMER-A
       DNS filtering counters:


UDP DNS A req count                        : 0
UDP DNS A resp count                       : 0
UDP DNS A log only count                   : 0
UDP DNS AAAA req count                     : 0
UDP DNS AAAA resp count                    : 0
UDP DNS AAAA log only count                : 0
UDP DNS MX req count                       : 0
UDP DNS MX resp count                      : 0
UDP DNS MX log only count                  : 0
UDP DNS CNAME req count                    : 0
UDP DNS CNAME resp count                   : 0
UDP DNS CNAME log only count               : 0
UDP DNS SRV req count                      : 0
UDP DNS SRV resp count                     : 0
UDP DNS SRV resp count                     : 0

UDP DNS SRV resp count                     : 0
+ UDP DNS SRV Resp No Err count   : 0
+ UDP DNS SRV Resp Resp Refused Err count   : 0
UDP DNS SRV log only count                 : 0
UDP DNS TXT req count                      : 0
```

```
UDP DNS TXT resp count                      : 0
UDP DNS TXT log only count                  : 0
+ UDP DNS TXT Resp No Err count   : 0
+ UDP DNS TXT Resp Resp Refused Err count   : 0
UDP DNS ANY req count                       : 0
UDP DNS ANY resp count                      : 0
UDP DNS ANY log only count                  : 0
UDP DNS MISC req count                      : 0
UDP DNS MISC log only count                 : 0
TCP DNS A req count                         : 0
TCP DNS A resp count                        : 0
TCP DNS A log only count                    : 0
TCP DNS AAAA req count                      : 0
TCP DNS AAAA resp count                     : 0
TCP DNS AAAA log only count                 : 0
TCP DNS MX req count                        : 0
TCP DNS MX resp count                       : 0
TCP DNS MX log only count                   : 0
TCP DNS CNAME req count                     : 0
TCP DNS CNAME resp count                    : 0
TCP DNS CNAME log only count                : 0
TCP DNS SRV req count                       : 0
TCP DNS SRV resp count                      : 0
TCP DNS SRV log only count                  : 0
+ TCP DNS SRV Resp No Err count   : 0
+ TCP DNS SRV Resp Resp Refused Err count   : 0

TCP DNS TXT req count                       : 0
TCP DNS TXT resp count                      : 0
TCP DNS TXT log only count                  : 0
+ TCP DNS SRV Resp No Err count   : 0
+ TCP DNS SRV Resp Resp Refused Err count   : 0

TCP DNS ANY req count                       : 0
TCP DNS ANY resp count                      : 0
TCP DNS ANY log only count                  : 0
TCP DNS MISC req count                      : 0
TCP DNS MISC log only count                 : 0
```

Release Information

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

Configuring URL Filtering

show services web-filter statistics profile

IN THIS SECTION

- Syntax | 26464
- Description | 26464
- Options | 26465
- Required Privilege Level | 26465
- Output Fields | 26465
- Sample Output | 26467
- Sample Output | 26468
- Release Information | 26470

Syntax

```
show services web-filter statistics profile profile-name
<dns-filter-template template-name>
<dns-filter-term term-name>
<fpc-slot fpc-slot pic-slot pic-slot>
<url-filter-template template-name>
```

Description

Display statistics for DNS request filtering and URL filtering for the specified filter profile.

Options

dns-filter-template ***template-name***	(Optional) Display statistics for the specified DNS filter template.
dns-filter-term ***term-name***	(Optional) Display statistics for the specified term in the DNS filter template.
fpc-slot ***fpc-slot*** **pic-slot** ***pic-slot***	(Optional) Display statistics for the specified services PIC.
profile ***profile-name***	Display statistics for the specified filter profile.
url-filter-template ***template-name***	(Optional) Display statistics for the specified URL filter template.

Required Privilege Level

view

Output Fields

Table 1567 on page 26465 lists the output fields for the `show services web-filter statistics profile` command. Output fields are listed in the approximate order in which they appear.

Table 1567: show services web-filter statistics profile Output Fields

Field Name	Field Description
`UDP Counters`	Number of UDP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.
`TCP Counters`	Number of TCP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.
`Accept`	Action counters for accepted packets for URL filtering.
`Custom page`	Action counters for custom page sent to recipient for URL filtering.
`Http scode`	Action counters for HTTP status code response for URL filtering.

Table 1567: show services web-filter statistics profile Output Fields *(Continued)*

Field Name	Field Description
`Redirect url`	Action counters for redirect URL response for URL filtering.
`TCP reset`	Action counters for TCP reset for URL filtering. Connection is closed.
`Bypass session count`	Number of sessions not blocked by URL filtering because the match criteria was not met for URL filtering.
`IPV4 Disable IP Blocking`	Action counters for IPv4 packets that were accepted because filtering is disabled for HTTP traffic that contains an embedded IP address belonging to a disallowed domain name in the URL filter database.
`IPV6 Disable IP Blocking`	Action counters for IPv6 packets that were accepted because filtering is disabled for HTTP traffic that contains an embedded IP address belonging to a disallowed domain name in the URL filter database.
`session count`	The session of activity that a user with a unique IP address spends on a website during a specified period of time for URL filtering. A session, in this case, would be the packets going to the service PIC from the Packet Forwarding Engine and then back to the service PIC.
`uplink packet count`	Number of packets going from the Packet Forwarding Engine to the service PIC for URL filtering.
`uplink bytes`	Number of bytes passing uplink for URL filtering.
`downlink packet count`	Number of packets going from the service PIC to the service Packet Forwarding Engine for URL filtering.
`downlink bytes`	Number of bytes passing downlink for URL filtering.
`UDP DNS`	Number of UDP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.

Table 1567: show services web-filter statistics profile Output Fields *(Continued)*

Field Name	Field Description
TCP DNS	Number of TCP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.

Sample Output

show services web-filter statistics profile dns-filter-template

```
user@host> show services web-filter statistics profile pdns dns-filter-template tdns
    Query   Requests         Responses        Log
    Type                                      only

    UDP Counters:

     A      0                0                0
     AAAA   0                0                0
     MX     0                0                0
     CNAME  0                0                0
     SRV    0                0                0
     TXT    0                0                0
     MISC   0                0                0

    TCP Counters:

     A      0                0                0
     AAAA   0                0                0
     MX     0                0                0
     CNAME  0                0                0
     SRV    0                0                0
     TXT    0                0                0
     MISC   0                0                0
```

Sample Output

show services web-filter statistics profile

```
user@host> show services web-filter statistics profile Profile1
URL filtering action counters:


    Accept session count                : 0
    Accept uplink packet count          : 0
    Accept uplink bytes                 : 0
    Accept downlink packet count        : 0
    Accept downlink bytes               : 0

    Custom page session count           : 0
    Custom page uplink packet count     : 0
    Custom page uplink bytes            : 0
    Custom page downlink packet count   : 0
    Custom page downlink bytes          : 0

    Http scode session count            : 0
    Http scode uplink packet count      : 0
    Http scode uplink bytes             : 0
    Http scode dowlink packet count     : 0
    Http scode downlink bytes           : 0

    Redirect url session count          : 0
    Redirect url uplink packet count    : 0
    Redirect url uplink bytes           : 0
    Redirect url downlink packet count  : 0
    Redirect url downlink bytes         : 0

    Tcp reset session count             : 0
    Tcp reset uplink packet count       : 0
    Tcp reset uplink bytes              : 0
    Tcp reset downlink packet count     : 0
    Tcp reset downlink bytes            : 0

    Bypass session count                : 0

    IPV4 Disable IP Blocking Sessions          : 0
    IPV4 Disable IP Blocking uplink packets    : 0
```

```
IPV4 Disable IP Blocking uplink bytes      : 0
IPV4 Disable IP Blocking downlink packets  : 0
IPV4 Disable IP Blocking downlink bytes    : 0
IPV6 Disable IP Blocking Sessions          : 0
IPV6 Disable IP Blocking uplink packets    : 0
IPV6 Disable IP Blocking uplink bytes      : 0
IPV6 Disable IP Blocking downlink packets  : 0
IPV6 Disable IP Blocking downlink bytes    : 0

DNS filtering counters:

UDP DNS A req count                        : 0
UDP DNS A resp count                       : 0
UDP DNS A log only count                   : 0
UDP DNS AAAA req count                     : 0
UDP DNS AAAA resp count                    : 0
UDP DNS AAAA log only count                : 0
UDP DNS MX req count                       : 0
UDP DNS MX resp count                      : 0
UDP DNS MX log only count                  : 0
UDP DNS CNAME req count                    : 0
UDP DNS CNAME resp count                   : 0
UDP DNS CNAME log only count               : 0
UDP DNS SRV req count                      : 0
UDP DNS SRV resp count                     : 0
UDP DNS SRV log only count                 : 0
UDP DNS TXT req count                      : 0
UDP DNS TXT resp count                     : 0
UDP DNS TXT log only count                 : 0
UDP DNS ANY req count                      : 0
UDP DNS ANY resp count                     : 0
UDP DNS ANY log only count                 : 0
UDP DNS MISC req count                     : 0
UDP DNS MISC log only count                : 0
TCP DNS A req count                        : 0
TCP DNS A resp count                       : 0
TCP DNS A log only count                   : 0
TCP DNS AAAA req count                     : 0
TCP DNS AAAA resp count                    : 0
TCP DNS AAAA log only count                : 0
TCP DNS MX req count                       : 0
TCP DNS MX resp count                      : 0
TCP DNS MX log only count                  : 0
```

```
    TCP DNS CNAME req count                    : 0
    TCP DNS CNAME resp count                   : 0
    TCP DNS CNAME log only count               : 0
    TCP DNS SRV req count                      : 0
    TCP DNS SRV resp count                     : 0
    TCP DNS SRV log only count                 : 0
    TCP DNS TXT req count                      : 0
    TCP DNS TXT resp count                     : 0
    TCP DNS TXT log only count                 : 0
    TCP DNS ANY req count                      : 0
    TCP DNS ANY resp count                     : 0
    TCP DNS ANY log only count                 : 0
    TCP DNS MISC req count                     : 0
    TCP DNS MISC log only count                : 0
```

Release Information

Command introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

Configuring URL Filtering

show services web-proxy dns forwarding-cache

IN THIS SECTION

- Syntax | 26471
- Description | 26471
- Required Privilege Level | 26471
- Output Fields | 26471

Sample Output | 26472
Release Information | 26473

Syntax

```
show services web-proxy dns forwarding-cache
```

Description

Display DNS cache information available at the packet forwarding engine for a secure Web proxy session.

Required Privilege Level

view

Output Fields

Table 1568 on page 26471 and Table 1569 on page 26472 describe the output fields for the `show services web-proxy dns forwarding-cache` command. Output fields are listed in the approximate order in which they appear.

Table 1568: show services web-proxy dns forwarding-cache statistics Output Fields

Field Name	Field Description
`Active DNS Cache Entries`	Number of active DNS cache entries.
`Total DNS Cache Entries`	Total number of DNS cache entries.
`DNS Cache hits`	DNS requests finding the match in the cache.
`DNS Cache miss`	DNS requests missing in the cache.

Table 1568: show services web-proxy dns forwarding-cache statistics Output Fields *(Continued)*

Field Name	Field Description
DNS Cache add failed	DNS requests failed to add in the DNS cache memory.

Table 1569: show services web-proxy dns forwarding-cache Output Fields

Field Name	Field Description
DNS Name	Name of the Domain Name System (DNS).
Address Family	IPv4 or IPv6 address family.
IP Address	IP address for the DNS name.

Sample Output

show services web-proxy dns forwarding-cache statistics

```
user@host> show services web-proxy dns forwarding-cache statistics

        DNS status                                      Active
        Active DNS Cache entries                        4294967270
        Total DNS Cache Entries                         1
        DNS Cache hits                                  61
        DNS Cache miss                                  191
        DNS Cache add failed                            0
```

show services web-proxy dns forwarding-cache

```
user@host> show services web-proxy dns forwarding-cache

DNS Name: settings-win.data.microsoft.com
  Address Family: IPv4
    IP Address: 40.90.221.9
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show services web-proxy dns global-cache statistics | 26473
show services web-proxy session | 26476

show services web-proxy dns global-cache statistics

IN THIS SECTION

- Syntax | 26473
- Description | 26473
- Required Privilege Level | 26473
- Output Fields | 26474
- Sample Output | 26475
- Release Information | 26475

Syntax

```
show services web-proxy dns global-cache statistics
```

Description

Display DNS cache information available at the routing engine for a secure Web proxy session.

Required Privilege Level

view

Output Fields

Table 1570 on page 26474 and Table 1571 on page 26474describe the output fields for the `show services web-proxy dns global-cache` commands. Output fields are listed in the approximate order in which they appear.

Table 1570: show services web-proxy statistics Output Fields

Field Name	Field Description
`Active DNS Cache Entries`	Number of active DNS cache entries.
`Total DNS Cache Entries`	Total number of DNS cache entries.
`DNS Cache hits`	Number of DNS requests finding the match in the cache.
`DNS Cache miss`	Number of DNS requests missing the match in the cache.
`DNS resolve request send failed`	Number of failed DNS resolve requests.
`DNS resolve request Timeout`	Number of DNS resolve requests timed out.

Table 1571: show services web-proxy dns global-cache Output Fields

Field Name	Field Description
DNS Name	Name of the Domain Name System (DNS).
Address Family	IPv4 or IPv6 address family.
TTL	Time-to-live value.
IP Address	IP address for the DNS name.

Sample Output

show services web-proxy dns global-cache statistics

```
user@host> show services web-proxy dns global-cache statistics

        DNS Status                                      Active
        Active DNS Cache entries                        1
        Total DNS Cache Entries                         1
        DNS Cache hits                                  0
        DNS Cache miss                                  191
        DNS resolve request send failed                 0
        DNS resolve request Timeout                     6
```

show services web-proxy dns global-cache

```
user@host> show services web-proxy dns global-cache

DNS Name: settings-win.data.microsoft.com
Address Family: IPv4, TTL: 0
IP Address: 40.90.221.9
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show services web-proxy dns forwarding-cache | 26470

show services web-proxy session | 26476

show services web-proxy session

IN THIS SECTION

- Syntax | 26476
- Description | 26476
- Options | 26476
- Required Privilege Level | 26476
- Output Fields | 26477
- Sample Output | 26477
- Sample Output | 26478
- Release Information | 26478

Syntax

```
show services web-proxy session
```

Description

Display information about the secure Web proxy session.

Options

detail	Shows the web proxy session detail.
summary	Shows the web proxy session summary.

Required Privilege Level

view

Output Fields

Table 1572 on page 26477 describes the output fields for the `show services web-proxy session` command. Output fields are listed in the approximate order in which they appear.

Table 1572: show services web-proxy session Output Fields

Field Name	Field Description
`Client Session`	Session originating from the client to proxy server.
`Proxy Session`	Session originating from the proxy server to the client.
`Client Session ID`	Number that identifies the client session. Use this ID to get more information about the session.
`Proxy Session ID`	Number that identifies the proxy session. Use this ID to get more information about the session.
`Proxy Request`	Connect request details.
`Dynamic Web App`	Dynamic Web application details.

Sample Output

show services web-proxy session summary

```
user@host> show services web-proxy session summary

Web Proxy sessions:
Client Session                                Proxy Session
[34] 6.0.0.1/62746 ---> 5.0.0.1/8080         [35] 6.0.0.1/62746 ---> 208.80.154.224/443
[37] 6.0.0.1/62747 ---> 5.0.0.1/8080         [38] 6.0.0.1/62747 ---> 208.80.154.224/443
```

Sample Output

show services web-proxy session detail

```
user@host> show services web-proxy session detail

Web Proxy sessions:
Client Session ID: 36994, Proxy Session ID: 36995
Client: 6.0.0.1/61324 ---> 5.0.0.1/8080
Proxy : 6.0.0.1/61324 ---> 74.125.195.188/443
Proxy Request: CONNECT:mtalk.google.com:443
Dynamic Web App: junos:GOOGLE-MAPS

Client Session ID: 38037, Proxy Session ID: 38038
Client: 6.0.0.1/57342 ---> 5.0.0.1/8080
Proxy : 6.0.0.1/57342 ---> 216.58.194.202/443
Proxy Request: CONNECT:safebrowsing.googleapis.com:443
Dynamic Web App: junos:GOOGLE-GEN
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Secure Web Proxy

show sflow

IN THIS SECTION

- Syntax | 26479
- Description | 26479
- Options | 26479

- Required Privilege Level | 26479
- Output Fields | 26479
- Sample Output | 26481
- Release Information | 26481

Syntax

```
show sflow
<collector>
<interface>
```

Description

Display sFlow configuration information.

Options

none Display all sFlow configuration information.

collector (Optional) Display a list of configured sFlow collectors and their properties. See *show sflow collector*.

interface (Optional) Display the interfaces on which sFlow technology is enabled and the sampling parameters. See *show sflow interface*.

Required Privilege Level

view

Output Fields

Table 1573 on page 26480 lists the output fields for the `show sflow` command. Output fields are listed in the approximate order in which they appear.

Table 1573: show sflow Output Fields

Field Name	Field Description
`sFlow`	Status of the feature: `Enabled` or `Disabled`.
`Adaptive fallback`	Reverts back to the previously configured adaptive sample rate. Adaptive fallback is triggered when the number of samples per second for each FPC is less than the sample-limit threshold value. This option is disabled by default.
`Sample limit`	Number of packets sampled per second. This sample limit cannot be configured and is set to 300 packets per second.
`Sample limit threshold`	Threshold value that you explicitly configure for adaptive sampling fallback.
`Polling interval`	Interval at which the sFlow agent polls the interface.
`Sample rate egress`	Rate at which egress packets are sampled.
`Sample rate ingress`	Rate at which ingress packets are sampled.
`Agent ID`	IP address assigned to the sFlow agent.
`Source IP address`	Source IP address for the sFlow packets.
`Agent ID IPv6`	IPv6 address assigned to the sFlow agent.
`Source IPv6 address`	Source IPv6 address for the sFlow packets.
`Sample size`	Sample size of the raw packet header in bytes. The sample size ranges from 128 bytes to 512 bytes.

Sample Output

show sflow

```
user@host> show sflow
sFlow                    : Enabled
Adaptive fallback                          : Disabled
Sample limit             : 300 packets/second
Sample limit threshold   : 150 packets/second
Polling interval         : 20 second
Sample rate egress       : 1:2048: Disabled
Sample rate ingress      : 1:1000: Enabled
Agent ID                 : 10.93.54.7
Source IP address        : 10.93.54.7
```

```
user@router> show sflow
sFlow                    : Enabled
Adaptive fallback        : False
Sample limit             : 2000 packets/second
Sample limit Threshold   : 0 packets/second
Polling interval         : 0 second
Sample rate egress       : 1:2048:Disabled
Sample rate ingress      : 1:2048:Disabled
Agent ID                 : No valid agent IPv4
Agent ID IPv6            : No valid agent IPv6
Source IP address        : 45.1.1.1
Source IPv6 address      : No valid source IPv6
Sample Size              : 128 Bytes
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

show sflow interface

show sflow collector

clear sflow collector statistics (QFX Series) | 4204

show sflow collector

IN THIS SECTION

- Syntax | 26482
- Description | 26482
- Required Privilege Level | 26482
- Output Fields | 26482
- Sample Output | 26483
- Release Information | 26484

Syntax

```
show sflow collector
<brief | detail>
<address>
```

Description

Display a list of configured sFlow collectors and their properties. Use the `address` option to display the information for a single collector address.

Required Privilege Level

view

Output Fields

Table 1574 on page 26483 lists the output fields for the `show sflow collector` command. Output fields are listed in the approximate order in which they appear.

Table 1574: show sflow collector Output Fields

Field Name	Field Description	Level of Output
`Collector address`	IP address of the collector.	All levels
`Udp-Port`	UDP port number of the collector.	All levels
`Dscp`	Differentiated Services Code Point (DSCP) values applied to all packets for the collector.	All levels
`Forwarding-Class`	Forwarding-class values applied to all packets destined for the collector.	All levels
`No. of samples`	Number of samples collected.	All levels
`Number of Counter Samples`	Number of counter samples collected.	`detail`
`Number of Flow Sample`	Number of flow samples collected.	`detail`
`Number of Datagrams`	Number of datagrams collected.	`detail`

Sample Output

show sflow collector brief

```
user@host> show sflow collector brief
Collector            Udp-port    Dscp     Forwarding-Class     No. of samples
 address
10.213.0.191         6343        0        best-effort          23259
```

You get the same output for `show sflow collector address 10.213.0.191 brief`.

show sflow collector detail

```
user@host> show sflow collector detail
  Sflow Collector Information
    Collector Address: 10.213.0.191, Collector UDP Port: 6343
    DSCP Value: 0, Forwarding Class: best-effort
    Number of Counter Samples: 2724, Number of Flow Samples: 11
    Number of Samples: 23306, Number of Datagrams: 1509
```

You get the same output for `show sflow collector address 10.213.0.191 detail`.

show sflow collector detail (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)

```
user@host> show sflow collector detail
  Sflow Collector Information
    Collector Address: 10.213.1.13, Collector UDP Port: 6343
    DSCP Value: 0, Forwarding Class: best-effort
    Number of Counter Samples: 0, Number of Flow Samples: 0
    Number of Samples: 0, Number of Datagrams: 0
    Routing Instance Name: mgmt_default, Routing Instance Id:0
```

You get the same output for `show sflow collector address 10.213.1.13 detail`.

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

clear sflow collector statistics (QFX Series) | 4204

show sflow

show sflow interface

Example: Configure sFlow Technology to Monitor Network Traffic

show sflow interface

IN THIS SECTION

- Syntax | 26485
- Description | 26485
- Required Privilege Level | 26485
- Output Fields | 26485
- Sample Output | 26486
- Sample Output | 26488
- Sample Output | 26488
- Sample Output | 26488
- Release Information | 26488

Syntax

```
show sflow interface
<interface-name interface-name>
```

Description

Display the interfaces on which sFlow is enabled and the sampling parameters for the interface. Use the `interface-name` *`interface-name`* option to get the sFlow interfaces information for a specific interface.

Required Privilege Level

view

Output Fields

Table 1575 on page 26486 lists the output fields for the `show sflow interface` command. Output fields are listed in the approximate order in which they appear.

Table 1575: show sflow interface Output Fields

Field Name	Field Description	Level of Output
`Interface`	Interface on which sFlow technology is enabled.	All levels
`Status Egress`	Indicates whether an egress sample rate is enabled.	All levels
`Status Ingress`	Indicates whether an ingress sample rate is enabled.	All levels
`Sample rate Egress`	Rate at which egress packets are sampled.	All levels
`Sample rate Ingress`	Rate at which ingress packets are sampled.	All levels
`Adapted sample rate Egress`	Adapted rate at which egress packets are sampled.	All levels
`Adapted sample rate Ingress`	Adapted rate at which ingress packets are sampled.	All levels
`Polling-interval`	Interval at which the sFlow agent polls the interface.	All levels

Sample Output

show sflow interface interface-name

```
user@host> show sflow interface interface-name et-0/0/1
Interface       Status            Sample            Adapted          Polling
                                  rate              sample rate      interval

               Egress  Ingress    Egress Ingress  Egress Ingress
et-0/0/1.0     Enabled Enabled    100     100       800     800       20
```

show sflow interface (EX Series)

```
Interface        Status      Sample rate    Adapted sample rate   Polling-interval
            Egress Ingress   Egress Ingress   Egress Ingress
ge-0/0/0.0  Enabled  Disabled  1000    2048    1000     2048            20
```

show sflow interface (QFX Series Switch)

```
user@host> show sflow interface
Interface  Status         Sample rate    Adapted sample rate   Polling-interval
              Egress Ingress   Egress Ingress   Egress Ingress
ge-0/0/13.0   Enabled  Enabled  100    100   800    800           20
ge-1/0/15.0   Enabled  Enabled  100    100   400    400           20
ge-1/0/20.0   Enabled  Enabled  10     10    160    160           20
ge-1/0/21.0   Enabled  Enabled  10     10    10     10            20
```

show sflow interface (QFX3500 Switch in Standalone Mode)

```
user@host> show sflow interface
Interface  Status         Sample rate    Adapted sample rate   Polling-interval
              Egress Ingress   Egress Ingress   Egress Ingress
xe-0/0/0.0    Enabled  Disabled  1000    2048    1000     2048            20
xe-1/0/1.0    Enabled  Disabled  1000    2048    1000     2048            20
```

show sflow interface (QFX3500 Switch in Standalone Mode)

```
user@host> show sflow interface
Interface  Status         Sample rate    Adapted sample rate   Polling-interval
              Egress Ingress   Egress Ingress   Egress Ingress
xe-0/0/0.0    Enabled  Disabled  1000    2048    1000     2048            20
xe-1/0/1.0    Enabled  Disabled  1000    2048    1000     2048            20
```

Sample Output

show sflow interface (QFX3500 Switch in Standalone Mode)

```
user@host> show sflow interface
Interface  Status         Sample rate    Adapted sample rate   Polling-interval
              Egress Ingress   Egress Ingress   Egress Ingress
xe-0/0/0.0    Enabled  Disabled  1000    2048    1000     2048            20
xe-1/0/1.0    Enabled  Disabled  1000    2048    1000     2048            20
```

Sample Output

show sflow interface (QFX3500 Switch in Standalone Mode)

```
user@host> show sflow interface
Interface  Status         Sample rate    Adapted sample rate   Polling-interval
              Egress Ingress   Egress Ingress   Egress Ingress
xe-0/0/0.0    Enabled  Disabled  1000    2048    1000     2048            20
xe-1/0/1.0    Enabled  Disabled  1000    2048    1000     2048            20
```

Sample Output

show sflow interface (QFabric System)

```
user@host> show sflow interface
Interface  Status         Sample rate    Adapted sample rate   Polling-interval
              Egress Ingress   Egress Ingress   Egress Ingress
node1:xe-0/0/0.0    Enabled  Disabled  1000    2048    1000     2048            20
node2:xe-1/0/1.0    Enabled  Disabled  1000    2048    1000     2048            20
node4:xe-1/0/0.0    Enabled  Disabled  1000    2048    1000     2048            20
```

Release Information

Command introduced in Junos OS Release 9.3.

`interface-name` *`interface-name`* option introduced.

RELATED DOCUMENTATION

show sflow

show sflow collector

show snmp health-monitor

IN THIS SECTION

- Syntax | 26489
- Description | 26489
- Options | 26489
- Required Privilege Level | 26490
- Output Fields | 26490
- Sample Output | 26493
- Release Information | 26499

Syntax

```
show snmp health-monitor
<alarms <detail>> | <logs>
```

Description

Display information about Simple Network Management Protocol (SNMP) health monitor alarms and logs.

Options

none	Display information about all health monitor alarms and logs.
alarms <detail>	(Optional) Display detailed information about health monitor alarms.

logs (Optional) Display information about health monitor logs.

Required Privilege Level

view

Output Fields

Table 1576 on page 26490 describes the output fields for the `show snmp health-monitor` command. Output fields are listed in the approximate order in which they appear.

Table 1576: show snmp health-monitor Output Fields

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels
Variable description	Description of the health monitor object instance being monitored.	All levels
Variable name	Name of the health monitor object instance being monitored.	All levels
Value	Current value of the monitored variable in the most recent sample interval.	All levels

Table 1576: show snmp health-monitor Output Fields *(Continued)*

Field Name	Field Description	Level of Output

State	State of the alarm or event entry:
	• Alarms:
	  • **active**—Entry is fully configured and activated.
	  • **falling threshold crossed**—Value of the variable has crossed the lower threshold limit.
	  • **rising threshold crossed**—Value of the variable has crossed the upper threshold limit.
	  • **under creation**—Entry is being configured and is not yet activated.
	  • **startup**—Alarm is waiting for the first sample of the monitored variable.
	  • **object not available**—Monitored variable of that type is not available to the health monitor agent.
	  • **instance not available**—Monitored variable's instance is not available to the health monitor agent.
	  • **object type invalid**—Monitored variable is not a numeric value.
	  • **object processing errored**—An error occurred when the monitored variable was processed.
	  • **unknown**—State is not one of the above.
		All levels

Variable OID	Object ID to which the variable name is resolved. The format is **x.x.x.x**.	**detail**
Sample type	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of **absolute value** or **delta value**.	**detail**

Table 1576: show snmp health-monitor Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Startup alarm	Alarm that might be sent when this entry is first activated, depending on the following criteria:
	• Alarm is sent when one of the following situations exists:
	  • Value of the alarm is above or equal to the rising threshold and the startup type is either **rising alarm** or **rising or falling alarm.**
	  • Value of the alarm is below or equal to the falling threshold and the startup type is either **falling alarm** or **rising or falling alarm**.
	• Alarm is *not* sent when one of the following situations exists:
	  • Value of the alarm is above or equal to the rising threshold and the startup type is **falling alarm**.
	  • Value of the alarm is below or equal to the falling threshold and the startup type is **rising alarm**.
	  • Value of the alarm is between the thresholds.
		detail

Owner	Name of the entry configured by the user. If the entry was created through the CLI, the owner has **monitor** prepended to it.	**detail**
Creator	Mechanism by which the entry was configured (**Health Monitor**).	**detail**
Sample interval	Time period between samples (in seconds).	**detail**
Rising threshold	Upper limit threshold value as a percentage of the maximum possible value.	**detail**
Falling threshold	Lower limit threshold value as a percentage of the maximum possible value.	**detail**
Rising event index	Event triggered when the rising threshold is crossed.	**detail**

Table 1576: show snmp health-monitor Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Falling event index	Event triggered when the falling threshold is crossed.	**detail**

Sample Output

show snmp health-monitor

```
user@host> show snmp health-monitor

Alarm
Index  Variable description                          Value State

 32768 Health Monitor: root file system utilization
       jnxHrStoragePercentUsed.1                        58 active

 32769 Health Monitor: /config file system utilization
       jnxHrStoragePercentUsed.2                         0 active

 32770 Health Monitor: RE 0 CPU utilization
       jnxOperatingCPU.9.1.0.0                           0 active

 32773 Health Monitor: RE 0 Memory utilization
       jnxOperatingBuffer.9.1.0.0                       35 active

 32775 Health Monitor: jkernel daemon CPU utilization
       Init daemon                                       0 active
       Chassis daemon                                   50 active
       Firewall daemon                                   0 active
       Interface daemon                                  5 active
       SNMP daemon                                      11 active
       MIB2 daemon                                      42 active
       Sonet APS daemon                                  0 active
       VRRP daemon                                       0 active
       Alarm daemon                                      3 active
       PFE daemon                                        0 active
       CRAFT daemon                                      0 active
```

```
      Traffic sampling control daemon                    0 active
      Ilmi daemon                                        0 active
      Remote operations daemon                           0 active
      CoS daemon                                         0 active
      Pic Services Logging daemon                        0 active
      Internal Routing Service Daemon                    3 active
      Network Access Service daemon                      0 active
      Forwarding UDP daemon                              0 active
      Routing socket proxy daemon                        0 active
      Disk Monitoring daemon                             1 active
      Inet daemon                                        0 active
      Syslog daemon                                      0 active
      Adaptive Services PIC daemon                       0 active
      ECC parity errors logging Daemon                   0 active
      Layer 2 Tunneling Protocol daemon                  0 active
      PPPoE daemon                                       3 active
      Redundancy device daemon                           0 active
      PPP daemon                                         0 active
      Dynamic Flow Capture Daemon                        0 active

32776 Health Monitor: jroute daemon CPU utilization
      Routing protocol daemon                            1 active
      Management daemon                                  0 active
      Management daemon                                  0 active
      Command line interface                             4 active
      Periodic Packet Management daemon                  0 active
      Link Management daemon                             0 active
      Pragmatic General Multicast daemon                 0 active
      Bidirectional Forwarding Detection daemon          0 active
      SRC daemon                                         0 active
      audit daemon                                       0 active
      Event daemon                                       0 active

32777 Health Monitor: jcrypto daemon CPU utilization
      IPSec Key Management daemon                        0 active

32779 Health Monitor: jkernel daemon Memory utilization
      Init daemon                                    47384 active
      Chassis daemon                                 20204 active
      Firewall daemon                                 1956 active
      Interface daemon                                3340 active
      SNMP daemon                                     4540 active
      MIB2 daemon                                     3880 active
```

```
      Sonet APS daemon                                2632 active
      VRRP daemon                                     2672 active
      Alarm daemon                                    1856 active
      PFE daemon                                      2600 active
      CRAFT daemon                                    2000 active
      Traffic sampling control daemon                 3164 active
      Ilmi daemon                                     2132 active
      Remote operations daemon                        2964 active
      CoS daemon                                      3044 active
      Pic Services Logging daemon                     1944 active
      Internal Routing Service Daemon                 1392 active
      Network Access Service daemon                   1992 active
      Forwarding UDP daemon                           1876 active
      Routing socket proxy daemon                     1296 active
      Disk Monitoring daemon                          1180 active
      Inet daemon                                     1296 active
      Syslog daemon                                   1180 active
      Adaptive Services PIC daemon                    3220 active
      ECC parity errors logging Daemon                1100 active
      Layer 2 Tunneling Protocol daemon               3372 active
      PPPoE daemon                                    1424 active
      Redundancy device daemon                        1820 active
      PPP daemon                                      2060 active
      Dynamic Flow Capture Daemon                    10740 active
32780 Health Monitor: jroute daemon Memory utilization
      Routing protocol daemon                         8104 active
      Management daemon                              13360 active
      Management daemon                              19252 active
      Command line interface                          9912 active
      Periodic Packet Management daemon               1484 active
      Link Management daemon                          2016 active
      Pragmatic General Multicast daemon              1968 active
      Bidirectional Forwarding Detection daemon       1956 active
      SRC daemon                                      1772 active
      audit daemon                                    1772 active
      Event daemon                                    1808 active

32781 Health Monitor: jcrypto daemon Memory utilization
      IPSec Key Management daemon                     5600 active
```

show snmp health-monitor alarms detail

```
user@host> show snmp health-monitor alarms detail

Alarm Index 32768:
  Variable name                        jnxHrStoragePercentUsed.1
  Variable OID                         1.3.6.1.4.1.2636.3.31.1.1.1.1.1
  Sample type                          absolute value
  Startup alarm                        rising alarm
  Owner                                Health Monitor: root file system
                                       utilization
  Creator                              Health Monitor
  State                                active
  Sample interval                  300 seconds
  Rising threshold                  80
  Falling threshold                 70
  Rising event index             32768
  Falling event index            32768
    Instance Value: 58
    Instance State: active

Alarm Index 32769:
  Variable name                        jnxHrStoragePercentUsed.2
  Variable OID                         1.3.6.1.4.1.2636.3.31.1.1.1.1.2
  Sample type                          absolute value
  Startup alarm                        rising alarm
  Owner                                Health Monitor: /config file system
                                       utilization
  Creator                              Health Monitor
  State                                active
  Sample interval                  300 seconds
  Rising threshold                  80
  Falling threshold                 70
  Rising event index             32768
  Falling event index            32768
    Instance Value: 0
    Instance State: active

Alarm Index 32770:
  Variable name                        jnxOperatingCPU.9.1.0.0
  Variable OID                         1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0
  Sample type                          absolute value
```

```
  Startup alarm                        rising alarm
  Owner                                Health Monitor: RE 0 CPU utilization

  Creator                              Health Monitor
  State                                active
  Sample interval                  300 seconds
  Rising threshold                  80
  Falling threshold                 70
  Rising event index             32768
  Falling event index            32768
    Instance Value: 0
    Instance State: active

Alarm Index 32773:
  Variable name                        jnxOperatingBuffer.9.1.0.0
  Variable OID                         1.3.6.1.4.1.2636.3.1.13.1.11.9.1.0.0
  Sample type                          absolute value
  Startup alarm                        rising alarm
  Owner                                Health Monitor: RE 0 Memory utilization

  Creator                              Health Monitor
  State                                active
  Sample interval                  300 seconds
  Rising threshold                  80
  Falling threshold                 70
  Rising event index             32768
  Falling event index            32768
    Instance Value: 35
    Instance State: active

Alarm Index 32775:
  Variable name                        sysApplElmtRunCPU.3
  Variable OID                         1.3.6.1.2.1.54.1.2.3.1.9.3
  Sample type                          delta value
  Startup alarm                        rising alarm
  Owner                                Health Monitor: jkernel daemon CPU
                                       utilization
  Creator                              Health Monitor
  State                                active
  Sample interval                  300 seconds
  Rising threshold               24000
  Falling threshold              21000
  Rising event index             32768
```

```
Falling event index            32768
  Instance Name: sysApplElmtRunCPU.3.1.1
  Instance Description: Init daemon
  Instance Value: 0
  Instance State: active

  Instance Name: sysApplElmtRunCPU.3.2.2786
  Instance Description: Chassis daemon
  Instance Value: 50
  Instance State: active

  Instance Name: sysApplElmtRunCPU.3.3.2938
  Instance Description: Firewall daemon
  Instance Value: 0
  Instance State: active

  Instance Name: sysApplElmtRunCPU.3.4.2942
  Instance Description: Interface daemon
  Instance Value: 5
  Instance State: active

  Instance Name: sysApplElmtRunCPU.3.7.7332
  Instance Description: SNMP daemon
  Instance Value: 11
  Instance State: active

  Instance Name: sysApplElmtRunCPU.3.9.2914
  Instance Description: MIB2 daemon
  Instance Value: 42
  Instance State: active

  Instance Name: sysApplElmtRunCPU.3.12.2916
  Instance Description: Sonet APS daemon
  Instance Value: 0
  Instance State: active

  Instance Name: sysApplElmtRunCPU.3.13.2917
  Instance Description: VRRP daemon
  Instance Value: 0
  Instance State: active

  Instance Name: sysApplElmtRunCPU.3.14.2787
  Instance Description: Alarm daemon
```

```
    Instance Value: 3
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.15.2940
    Instance Description: PFE daemon
    Instance Value: 0
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.16.2788
    Instance Description: CRAFT daemon
    Instance Value: 0
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.17.2918
    Instance Description: Traffic sampling control daemon
---(more 23%)---
```

Release Information

Command introduced in Junos OS Release 8.0.

show snmp inform-statistics

IN THIS SECTION

- Syntax | 26500
- Description | 26500
- Options | 26500
- Required Privilege Level | 26500
- Output Fields | 26500
- Sample Output | 26501
- Release Information | 26501

Syntax

```
show snmp inform-statistics
```

Description

Display information about Simple Network Management Protocol (SNMP) inform requests.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1577 on page 26500 describes the output fields for the `show snmp inform-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1577: show snmp inform-statistics Output Fields

Field Name	Field Description
`Target Name`	Name of the device configured to receive and respond to SNMP informs.
`Address`	IP address of the target device.
`Sent`	Number of informs sent to the target device and acknowledged by the target device.
`Pending`	Number of informs held in memory pending a response from the target device.
`Discarded`	Number of informs discarded after the specified number of retransmissions to the target device were attempted.

Table 1577: show snmp inform-statistics Output Fields *(Continued)*

Field Name	Field Description
Timeouts	Number of informs that did not receive an acknowledgement from the target device within the timeout specified.
Probe Failures	Connection failures that occurred (for example, when the target server returned invalid content or you incorrectly configured the target address).

Sample Output

show snmp inform-statistics

```
user@host> show snmp inform-statistics
Inform Request Statistics:
    Target Name: TA1_v3_md5_none Address: 172.17.20.184
        Sent: 176, Pending: 0
        Discarded: 0, Timeouts: 0, Probe Failures: 0
    Target Name: TA2_v3_sha_none Address: 192.168.110.59
        Sent: 0, Pending: 4
        Discarded: 84, Timeouts: 0, Probe Failures: 258
    Target Name: TA5_v2_none Address: 172.17.20.184
        Sent: 0, Pending: 0
        Discarded: 2, Timeouts: 10, Probe Failures: 0
```

Release Information

Command introduced in Junos OS Release 7.4.

show snmp mib

IN THIS SECTION

- Syntax (Junos OS and Junos OS Evolved) | 26502
- Description | 26502
- Options | 26503
- Required Privilege Level | 26503
- Output Fields | 26503
- Sample Output | 26504
- Release Information | 26516

Syntax (Junos OS and Junos OS Evolved)

```
show snmp mib (get | get-next | walk) (ascii | decimal | hex) object-id
```

Description

Display local Simple Network Management Protocol (SNMP) Management Information Base (MIB) object values.

> **NOTE**: Starting in Junos OS Release 18.3R1, SNMP queue statistics reporting for static interface-sets configured over Aggregate Ethernet (AE) interfaces is supported.
>
> Starting in Junos OS Release 19.1R1, the sysName.0 MIB object displays the fully qualified domain name. That is, if the hostname and domain name are configured on the system, both will show up for the sysName.0 MIB object.

Starting in Junos OS Release 20.2R1, We’ve introduced the new SNMP MIBs `jnxJsSPUMonitoringSPUThreadsNumber`, `jnxJsSPUMonitoringSPUThreadIndex`, `jnxJsSPUMonitoringSPUThreadLastMinUsage`, `jnxJsSPUMonitoringSPUThreadLastHourUsage`, and `jnxJsSPUMonitoringSPUThreadLastDayUsage` to monitor the CPU usage information of an SPC3 card.

Options

get Retrieve and display one or more SNMP object values.

get-next Retrieve and display the next SNMP object values.

On Junos OS Evolved, when the system (router) receives `get-next` request from external server with 1000 varbinds, the system (router) does not support to process this request and it responds with a general error `genErr` every time.

walk Retrieve and display the SNMP object values that are associated with the requested object identifier (OID). When you use this option, the Junos OS displays the objects below the subtree that you specify.

> **NOTE**: As of Junos OS Release 18.4R1, the CLI configuration command `set snmp customization ether-stats-ifd-only` is introduced. When `ether-stats-ifd-only` is configured, the `show snmp mib walk etherstatsTable` command displays data only for physical interfaces (IFDs). See "customization (SNMP)" on page 4784.

ascii Display the SNMP object's string indices as an ASCII-key representation.

decimal Display the SNMP object values in the decimal (default) format. The `decimal` option is the default option for this command. Therefore, issuing the `show snmp mib (get | get-next | walk) decimal object-id` and the `show snmp mib (get | get-next | walk) object-id` commands display the same output.

hex Display the SNMP object values in the hexadecimal format.

object-id The object can be represented by a sequence of dotted integers (such as `1.3.6.1.2.1.2`) or by its subtree name (such as `interfaces`). When entering multiple objects, enclose the objects in quotation marks.

Required Privilege Level

snmp To view this statement in the configuration.

Output Fields

Table 1578 on page 26504 describes the output fields for the `show snmp mib` command. Output fields are listed in the approximate order in which they appear.

Table 1578: show snmp mib Output Fields

Field Name	Field Description
name	Object name and numeric instance value.
object value	Object value. The Junos OS translates OIDs into the corresponding object names.

Sample Output

show snmp mib get

```
user@host> show snmp mib get sysObjectID.0
sysObjectID.0 = jnxProductNameM20
```

show snmp mib get (Routing Engine)

```
user@router> show snmp mib get jnxFruName.9.1.0.0

jnxFruName.9.1.0.0 = Routing Engine 0
```

show snmp mib get (Routing Engine, PTX10003)

```
user@router> show snmp mib get jnxFruName.9.1.0.0

jnxFruName.9.1.0.0 = Routing Engine slot 0
```

show snmp mib get (Multiple Objects)

```
user@host> show snmp mib get “sysObjectID.0 sysUpTime.0”
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
```

show snmp mib get (Layer 2 Policer)

```
user@host> show snmp mib get ifInOctets.25970
ifInOctets.25970 = 7545720
```

show snmp mib get-next

```
user@host> show snmp mib get-next jnxMibs
jnxBoxClass.0 = jnxProductLineM20.0
```

show snmp mib get-next (Specify an OID)

```
user@host> show snmp mib get-next 1.3.6.1
sysDescr.0    = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release: 2004-1 Build date: build date UTC Copyright (c) 1996-2004 Juniper Networks,
Inc.
```

show snmp mib walk

```
user@host> show snmp mib walk system
sysDescr.0    = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release #0: 2004-1 Build date: build date UTC Copyright (c) 1996-2004 Juniper Networks,
Inc.
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
sysContact.0 = Your contact
sysName.0 = my router
sysLocation.0 = building 1
sysServices.0 = 4
```

show snmp mib walk (QFX Series)

```
user@switch> show snmp mib walk system
sysDescr.0    = Juniper Networks, Inc. qfx3500s internet router, kernel JUNOS 11.1-20100926.0
#0: 2010-09-26 06:17:38 UTC Build date: 2010-09-26 06:00:10
sysObjectID.0 = jnxProductQFX3500
```

```
sysUpTime.0   = 138980301
sysContact.0  = System Contact
sysName.0     = LabQFX3500
sysLocation.0 = Lab
sysServices.0 = 4
```

show snmp mib walk (Junos OS Evolved)

```
user@host> show snmp mib walk system
sysDescr.0    = Juniper Networks, Inc. ptx10003-80c , kernel JUNOS 22.1-202109210525.0-EVO,
Copyright (c) 1996-2021 Juniper Networks, Inc.
sysObjectID.0 = jnxProductNamePTX1000380c
sysUpTime.0   = 72814201
sysContact.0
sysName.0     = evoptx10k-a.englab.juniper.net
sysLocation.0
sysServices.0 = 4
```

show snmp mib walk (ASCII)

```
user@host> show snmp mib walk ascii jnxUtilData
jnxUtilCounter32Value."fred" = 100
```

show snmp mib walk (Multiple Indices)

```
user@host> show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.......
```

show snmp mib walk decimal

```
user@host>show snmp mib walk decimal jnxUtilData
jnxUtilCounter32Value.102.114.101.100 = 100
```

show snmp mib walk decimal (Multiple Indices)

```
user@host> show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.......
```

show snmp mib walk hex

```
user@host> show snmp mib walk hex jnxBgpM2PeerRemoteAddr
jnxBgpM2PeerRemoteAddr.0.1.10.0.3.3.1.10.0.3.1 = 0a 00 03 01
jnxBgpM2PeerRemoteAddr.0.1.32.32.32.9.1.32.32.32.32 = 00 00 00 00
jnxBgpM2PeerRemoteAddr.0.1.62.40.96.38.1.62.40.96.19 = 3e 28 60 13
jnxBgpM2PeerRemoteAddr.0.1.62.40.96.38.1.62.40.96.39 = 3e 28 60 27
```

show snmp mib walk (Queue Statistics)

To get interface-set queue statistics from jnxCosQstatQedPkts MIB with using interface-set SNMP index, use the following command:

```
show snmp mib walk jnxCosQstatQedPkts.interface-set snmp index
```

For example, if the interface-set SNMP index is 67108866

```
user@host> show snmp mib walk jnxCosQstatQedPkts.67108866
jnxCosQstatQedPkts.67108866.0 = 10
jnxCosQstatQedPkts.67108866.1 = 0
jnxCosQstatQedPkts.67108866.2 = 0
jnxCosQstatQedPkts.67108866.3 = 0
jnxCosQstatQedPkts.67108866.4 = 0
jnxCosQstatQedPkts.67108866.5 = 0
jnxCosQstatQedPkts.67108866.6 = 0
jnxCosQstatQedPkts.67108866.7 = 0
```

To get interface-set queue statistics from jnxCosIfsetQstatQedPkts MIB with using interface-set SNMP index, use the following command:

```
show snmp mib walk jnxCosIfsetQstatQedPkts.interface-set snmp index
```

For example, if the interface-set snmp index is 67108866

```
user@host> show snmp mib walk jnxCosIfsetQstatQedPkts.67108866
jnxCosIfsetQstatQedPkts.67108866.0 = 10
jnxCosIfsetQstatQedPkts.67108866.1 = 0
jnxCosIfsetQstatQedPkts.67108866.2 = 0
jnxCosIfsetQstatQedPkts.67108866.3 = 0
jnxCosIfsetQstatQedPkts.67108866.4 = 0
jnxCosIfsetQstatQedPkts.67108866.5 = 0
jnxCosIfsetQstatQedPkts.67108866.6 = 0
jnxCosIfsetQstatQedPkts.67108866.7 = 0
```

To get Interface-set queue statistics from jnxCosIfsetQstatQedPkts MIB using interface-set member IFL SNMP index, use the following command:

```
show snmp mib walk jnxCosIfsetQstatQedPkts.interface-set member IFL SNMP index
```

For example, if the interface-set member IFL SNMP is 519

```
user@host> show snmp mib walk jnxCosIfsetQstatQedPkts.519
jnxCosIfsetQstatQedPkts.519.0 = 10
jnxCosIfsetQstatQedPkts.519.1 = 0
jnxCosIfsetQstatQedPkts.519.2 = 0
jnxCosIfsetQstatQedPkts.519.3 = 0
jnxCosIfsetQstatQedPkts.519.4 = 0
jnxCosIfsetQstatQedPkts.519.5 = 0
jnxCosIfsetQstatQedPkts.519.6 = 0
jnxCosIfsetQstatQedPkts.519.7 = 0
```

show snmp mib walk bgp (MX Series)

Use the `show snmp mib walk bgp` to see the BGP versions.

```
user@host> show snmp mib walk bgp

bgpVersion.0  = 10
bgpLocalAs.0  = 1
bgpPeerIdentifier.2.2.2.2 = 2.2.2.2
bgpPeerIdentifier.3.3.3.3 = 3.3.3.3

bgpPeerInUpdates.2.2.2.2 = 3
bgpPeerInUpdates.3.3.3.3 = 2
bgpPeerOutUpdates.2.2.2.2 = 1
bgpPeerOutUpdates.3.3.3.3 = 1
bgpPeerInTotalMessages.2.2.2.2 = 11
bgpPeerInTotalMessages.3.3.3.3 = 10
bgpPeerOutTotalMessages.2.2.2.2 = 8
bgpPeerOutTotalMessages.3.3.3.3 = 8
```

show snmp mib walk mplsL3VpnMIB (MX Series)

Use the `show snmp mib walk mplsL3VpnMIB` to see attributes of L3VPN based MPLS.

```
user@host> show snmp mib walk mplsL3VpnMIB

mplsL3VpnConfiguredVrfs.0 = 1
mplsL3VpnActiveVrfs.0 = 1
mplsL3VpnConnectedInterfaces.0 = 1
mplsL3VpnNotificationEnable.0 = 2
mplsL3VpnVrfConfMaxPossRts.0 = 0
mplsL3VpnVrfConfRteMxThrshTime.0 = 0
mplsL3VpnIllLblRcvThrsh.0 = 0
mplsL3VpnIfVpnClassification.3.114.101.100.548 = 2
mplsL3VpnIfVpnRouteDistProtocol.3.114.101.100.548 = 40 7f
mplsL3VpnIfConfStorageType.3.114.101.100.548 = 5
mplsL3VpnIfConfRowStatus.3.114.101.100.548 = 1
mplsL3VpnVrfVpnId.3.114.101.100
mplsL3VpnVrfDescription.3.114.101.100 = red
mplsL3VpnVrfRD.3.114.101.100 = 2:1
```

```
mplsL3VpnVrfCreationTime.3.114.101.100 = 9111800
mplsL3VpnVrfOperStatus.3.114.101.100 = 1
mplsL3VpnVrfActiveInterfaces.3.114.101.100 = 1
mplsL3VpnVrfAssociatedInterfaces.3.114.101.100 = 1
mplsL3VpnVrfConfMidRteThresh.3.114.101.100 = 0
mplsL3VpnVrfConfHighRteThresh.3.114.101.100 = 0
mplsL3VpnVrfConfMaxRoutes.3.114.101.100 = 0
mplsL3VpnVrfConfLastChanged.3.114.101.100 = 9111800
```

show snmp mib walk (PTX10003)

On PTX10003-80C and PTX10003-160C devices, the `show snmp mib walk jnxFilledDescr` output shows only the fan tray number. This output does not show the number of fan slots present in each tray.

```
user@router> show snmp mib walk jnxFilledDescr

jnxFilledDescr.1.0.0.0 = Chassis
jnxFilledDescr.4.2.0.0 = Fan Tray 1
jnxFilledDescr.4.3.0.0 = Fan Tray 2
jnxFilledDescr.4.4.0.0 = Fan Tray 3
[...Output truncated...]
```

show snmp mib walk usmUserAuthProtocol (MX-Series and SRX-Series)

Use `show snmp mib walk usmUserAuthProtocol`to see the SHA authentication type.

```
user@router> show snmp mib walk usmUserAuthProtocol
usmUserAuthProtocol.< usmUserEngineID>.<usmUserName> =
usmHMAC128SHA224AuthProtocol

        New possible set of values for ‘usmUserAuthProtocol’ are:
            usmHMAC128SHA224AuthProtocol
            usmHMAC192SHA256AuthProtocol
            usmHMAC256SHA384AuthProtocol
            usmHMAC384SHA512AuthProtocol
```

show snmp mib walk ascii jnxWlanWAPStatusTable (SRX320, SRX340, SRX345, and SRX550M)

Use the `show snmp mib walk ascii jnxWlanWAPStatusTable` command to monitor the Wi-Fi Mini-Physical Interface Module (Mini-PIM) status.

```
user@host> show snmp mib walk ascii jnxWlanWAPStatusTable

jnxWAPStatusIfdIndex.161 = 161
jnxWAPStatusIfdIndex.162 = 162
jnxWAPStatusAccessPoint.161 = bj345b_wl3_wap
jnxWAPStatusAccessPoint.162 = bj345b_wap
jnxWAPStatusType.161 = Internal
jnxWAPStatusType.162 = Internal
jnxWAPStatusLocation.161 = Default Location
jnxWAPStatusLocation.162 = Default Location
jnxWAPStatusSerialNumber.161
jnxWAPStatusSerialNumber.162
jnxWAPStatusFirmwareVersion.161 = v1.1.0
jnxWAPStatusFirmwareVersion.162 = v1.1.0
jnxWAPStatusAlternateVersion.161 = v1.1.0
jnxWAPStatusAlternateVersion.162 = v1.1.0
jnxWAPStatusCountry.161 = US
jnxWAPStatusCountry.162 = US
jnxWAPStatusAccessInterface.161 = wl-3/0/0
jnxWAPStatusAccessInterface.162 = wl-4/0/0
jnxWAPStatusSystemTime.161 = Fri Jun 21 05:05:42 UTC 2019
jnxWAPStatusSystemTime.162 = Fri Jun 21 05:38:46 UTC 2019
jnxWAPStatusPacketCapture.161 = Off
jnxWAPStatusPacketCapture.162 = Off
jnxWAPStatusEthernetPortMAC.161 = 56:48:0d:5e:8f:c5
jnxWAPStatusEthernetPortMAC.162 = 8a:b7:0a:7e:ad:8a
jnxWAPStatusEthernetIPv4.161
jnxWAPStatusEthernetIPv4.162
jnxWAPStatusRadio1Status.161 = On
jnxWAPStatusRadio1Status.162 = On
[...Output truncated...]
```

show snmp mib walk jnxWlanWAPClientTable (SRX320, SRX340, SRX345, and SRX550M)

Use the `show snmp mib walk jnxWlanWAPClientTable` command to monitor the Wi-Fi Mini-PIM client information.

```
user@host> show snmp mib walk jnxWlanWAPClientTable

jnxWAPClientIfdIndex.161.1 = 161
jnxWAPClientIfdIndex.162.1 = 162
jnxWAPClientIfdIndex.162.2 = 162
jnxWAPClientIndex.161.1 = 1
jnxWAPClientIndex.162.1 = 1
jnxWAPClientIndex.162.2 = 2
jnxWAPClientRadioID.161.1 = 1
jnxWAPClientRadioID.162.1 = 1
jnxWAPClientRadioID.162.2 = 1
jnxWAPClientSSID.161.1 = bj345b_wl3_5g
jnxWAPClientSSID.162.1 = bj345b_5g
jnxWAPClientSSID.162.2 = bj345b_5g
jnxWAPClientMAC.161.1 = e8:4e:06:64:38:89
jnxWAPClientMAC.162.1 = e8:4e:06:64:38:a3
jnxWAPClientMAC.162.2 = e8:4e:06:63:9d:f6
jnxWAPClientAuth.161.1 = NO
[...Output truncated...]
```

show snmp mib walk jnxJsSPUMonitoringSPUThreadsNumber (SRX5400, SRX5600, and SRX5800)

Use the `show snmp mib walk jnxJsSPUMonitoringSPUThreadsNumber` command to list the indices of flow threads in an SPU.

```
user@host> show snmp mib walk jnxJsSPUMonitoringSPUThreadsNumber

jnxJsSPUMonitoringSPUThreadsNumber.4 = 26
jnxJsSPUMonitoringSPUThreadsNumber.5 = 26
jnxJsSPUMonitoringSPUThreadsNumber.16 = 0
jnxJsSPUMonitoringSPUThreadsNumber.17 = 0
jnxJsSPUMonitoringSPUThreadsNumber.18 = 0
jnxJsSPUMonitoringSPUThreadsNumber.19 = 0
```

show snmp mib walk jnxJsSPUMonitoringSPUThreadIndex (SRX5400, SRX5600, and SRX5800)

Use the `show snmp mib walk jnxJsSPUMonitoringSPUThreadIndex` command to list the indices of flow threads in an SPU.

```
user@host> show snmp mib walk jnxJsSPUMonitoringSPUThreadIndex

jnxJsSPUMonitoringSPUThreadIndex.4.0 = 0
jnxJsSPUMonitoringSPUThreadIndex.4.1 = 1
jnxJsSPUMonitoringSPUThreadIndex.4.2 = 2
jnxJsSPUMonitoringSPUThreadIndex.4.3 = 3

... ...

jnxJsSPUMonitoringSPUThreadIndex.5.22 = 22
jnxJsSPUMonitoringSPUThreadIndex.5.23 = 23
jnxJsSPUMonitoringSPUThreadIndex.5.24 = 24
jnxJsSPUMonitoringSPUThreadIndex.5.25 = 25
```

show snmp mib get jnxJsSPUMonitoringSPUThreadLastMinUsage.5.8 (SRX5400, SRX5600, and SRX5800)

Use the `show snmp mib get jnxJsSPUMonitoringSPUThreadLastMinUsage.5.8` command to monitor the average CPU usage for the last minute on the specified thread.

```
user@host> show snmp mib get jnxJsSPUMonitoringSPUThreadLastMinUsage.5.8

jnxJsSPUMonitoringSPUThreadLastMinUsage.5.8 = 0
```

show snmp mib get jnxJsSPUMonitoringSPUThreadLastHourUsage.5.8 (SRX5400, SRX5600, SRX5800)

Use the `show snmp mib get jnxJsSPUMonitoringSPUThreadLastHourUsage.5.8` command to monitor the average CPU usage for the last hour on the specified thread.

```
user@host> show snmp mib get jnxJsSPUMonitoringSPUThreadLastHourUsage.5.8

jnxJsSPUMonitoringSPUThreadLastHourUsage.5.8 = 0
```

show snmp mib get jnxJsSPUMonitoringSPUThreadLastDayUsage.5.8 (SRX5400, SRX5600, and SRX5800)

Use the `show snmp mib get jnxJsSPUMonitoringSPUThreadLastDayUsage.5.8` command to monitor the average CPU usage for the last 24 hours on the specified thread.

```
user@host> show snmp mib get jnxJsSPUMonitoringSPUThreadLastDayUsage.5.8

jnxJsSPUMonitoringSPUThreadLastDayUsage.5.8 = 0
```

show snmp mib walk ifJnxTable (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)

```
user@switch> show snmp mib walk ifJnxTable | grep Ecn
ifJnxOutEcnMarkedPackets.4 = 0
ifJnxOutEcnMarkedPackets.6 = 0
ifJnxOutEcnMarkedPackets.16 = 0
ifJnxOutEcnMarkedPackets.193 = 0
ifJnxOutEcnMarkedPackets.194 = 0
ifJnxOutEcnMarkedPackets.501 = 0
ifJnxOutEcnMarkedPackets.502 = 0
ifJnxOutEcnMarkedPackets.503 = 0
ifJnxOutEcnMarkedPackets.504 = 0

show interfaces et-0/0/0 extensive
...
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
```

```
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0,
  ECN Marked packets: 0
```

```
user@switch> show snmp mib walk ifJnxTable | grep InQ
ifJnxInQDrops.4 = 0
ifJnxInQDrops.6 = 0
ifJnxInQDrops.16 = 0
ifJnxInQDrops.193 = 0
ifJnxInQDrops.194 = 0
ifJnxInQDrops.501 = 0
ifJnxInQDrops.502 = 0
ifJnxInQDrops.503 = 0
ifJnxInQDrops.504 = 0
ifJnxInQDrops.505 = 0

show interfaces et-0/0/0 extensive
...
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
  Resource errors: 0
```

show snmp mib walk jnxCosPfcPriorityTable (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)

```
user@switch> show snmp mib walk jnxCosPfcPriorityTable
jnxCosPfcPriorityRequestsTx.503.0 = 0
jnxCosPfcPriorityRequestsTx.503.1 = 0
jnxCosPfcPriorityRequestsTx.503.2 = 0
jnxCosPfcPriorityRequestsTx.503.3 = 0
jnxCosPfcPriorityRequestsTx.503.4 = 0
jnxCosPfcPriorityRequestsTx.503.5 = 0
jnxCosPfcPriorityRequestsTx.503.6 = 0
jnxCosPfcPriorityRequestsTx.503.7 = 0

jnxCosPfcPriorityRequestsRx.503.0 = 0
jnxCosPfcPriorityRequestsRx.503.1 = 0
jnxCosPfcPriorityRequestsRx.503.2 = 0
jnxCosPfcPriorityRequestsRx.503.3 = 0
jnxCosPfcPriorityRequestsRx.503.4 = 0
```

```
jnxCosPfcPriorityRequestsRx.503.5 = 0
jnxCosPfcPriorityRequestsRx.503.6 = 0
jnxCosPfcPriorityRequestsRx.503.7 = 0

show interfaces et-0/0/0 extensive
...
MAC Priority Flow Control Statistics:
    Priority :  0                           0               0
    Priority :  1                           0               0
    Priority :  2                           0               0
    Priority :  3                           0               0
    Priority :  4                           0               0
    Priority :  5                           0               0
    Priority :  6                           0               0
    Priority :  7                           0               0
```

Release Information

Command introduced before Junos OS Release 7.4.

`ascii` and `decimal` options introduced in Junos OS Release 9.6.

`ascii` and `decimal` options introduced in Junos OS Release 9.6 for EX Series switches.

`hex` option introduced in Junos OS Release 19.2R2.

Behavior in regard to sysName.0 MIB object changed in Junos OS Release 19.1R1.

show snmp rmon

IN THIS SECTION

- Syntax | 26517
- Description | 26517
- Options | 26517
- Required Privilege Level | 26517
- Output Fields | 26517
- Sample Output | 26521

- Release Information | 26523

Syntax

```
show snmp rmon
<alarms <brief | detail> | events <brief | detail> | logs>
```

Description

Display information about Simple Network Management Protocol (SNMP) Remote Monitoring (RMON) alarms, events, and logs.

Options

none	Display information about all RMON alarms and events.
alarms	(Optional) Display information about RMON alarms.
brief | detail	(Optional) Display brief or detailed information about RMON alarms or events.
events	(Optional) Display information about RMON events.
logs	(Optional) Display information about RMON monitoring logs.

Required Privilege Level

view

Output Fields

Table 1579 on page 26518 describes the output fields for the `show snmp rmon` command. Output fields are listed in the approximate order in which they appear.

Table 1579: show snmp rmon Output Fields

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels
State	State of the alarm or event entry: Alarms: • **active**—Entry is fully configured and activated. • **falling threshold crossed**—Value of the variable has crossed the lower threshold limit. • **rising threshold crossed**—Value of the variable has crossed the upper threshold limit. • **under creation**—Entry is being configured and is not yet activated. • **startup**—Alarm is waiting for the first sample of the monitored variable. • **object not available**—Monitored variable of that type is not available to the SNMP agent. • **instance not available**—Monitored variable's instance is not available to the SNMP agent. • **object type invalid**—Monitored variable is not a numeric value. • **object processing errored**—An error occurred when the monitored variable was processed. • **unknown**—State is not one of the above. Events: • **active**—Entry has been fully configured and activated. • **under creation**—Entry is being configured and is not yet activated. • **unknown**—State is not one of the above.	All levels
Variable name	Name of the SNMP object instance being monitored.	All levels

Table 1579: show snmp rmon Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Event Index	Event identifier.	All levels
Type	Type of notification made when an event is triggered. It can be one of the following: • **log**—A system log message is generated and an entry is made to the log table. • **snmptrap**—An SNMP trap is sent to the configured destination. • **log and trap**—A system log message is generated, an entry is made to the log table, and an SNMP trap is sent to the configured destination. • **none**—Neither log nor trap will be sent.	**detail**
Last Event	Date and time of the last event. It has the format *yyyy-mm-dd hh:mm:ss timezone*.	**brief**
Community	Trap group used for sending the SNMP trap.	**detail**
Variable OID	Object ID to which the variable name is resolved. The format is **x.x.x.x**.	**detail**
Sample type	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of **absolute value** or **delta value**.	**detail**

Table 1579: show snmp rmon Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Startup alarm	Alarm that might be sent when this entry is first activated, depending on the following criteria:
	• Alarm is sent when one of the following situations exists:
	  • Value of the alarm is above or equal to the rising threshold and the startup type is either **rising alarm** or **rising or falling alarm.**
	  • Value of the alarm is below or equal to the falling threshold and the startup type is either **falling alarm** or **rising or falling alarm**.
	• Alarm is *not* sent when one of the following situations exists:
	  • Value of the alarm is above or equal to the rising threshold and the startup type is **falling alarm**.
	  • Value of the alarm is below or equal to the falling threshold and the startup type is **rising alarm**.
	  • Value of the alarm is between the thresholds.
		detail

Owner	Name of the entry configured by the user. If the entry was created through the CLI, the owner has **monitor** prepended to it.	**detail**
Creator	Mechanism by which the entry was configured (**CLI** or **SNMP**).	**detail**
Sample interval	Time period between samples (in seconds).	**detail**
Rising threshold	Upper limit threshold value configured by the user.	**detail**
Falling threshold	Lower limit threshold value configured by the user.	**detail**
Rising event index	Event triggered when the rising threshold is crossed.	**detail**
Falling event index	Event triggered when the falling threshold is crossed.	**detail**

Table 1579: show snmp rmon Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Current value	Current value of the monitored variable in the most recent sample interval.	**detail**

Sample Output

show snmp rmon

```
user@host> show snmp rmon
Alarm
Index  State                        Variable name
    1  falling threshold crossed   ifInOctets.1

Event
Index  Type                         Last Event
    1  log and trap                2002-01-30 01:13:01 PST
```

show snmp rmon (QFX Series)

```
user@host> show snmp rmon
Alarm
Index  Variable description                          Value State

     5 monitor
       jnxOperatingCPU.9.1.0.0                           5 falling threshold

Event
Index  Type                         Last Event
    1  log and trap                2009-07-10 11:34:17 PDT
Event Index: 1
    Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed, (variable:
jnxOperatingCPU.9.1.0.0, value: 100)
    Time: 2009-07-10 11:34:07 PDT
    Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed, (variable:
```

```
jnxOperatingCPU.9.1.0.0, value: 5)
    Time: 2009-07-10 11:34:17 PDT
```

show snmp rmon alarms detail

```
user@host> show snmp rmon alarms detail

Alarm Index 1:
  Variable name                        ifInOctets.1
  Variable OID                         1.3.6.1.2.1.2.2.1.10.1
  Sample type                          delta value
  Startup alarm                        rising or falling alarm
  Owner                                monitor
  Creator                              CLI
  State                                falling threshold crossed
  Sample interval                   60 seconds
  Rising threshold              100000
  Falling threshold              80000
  Rising event index                 1
  Falling event index                1
  Current value                      0
```

show snmp rmon events detail

```
user@host> show snmp rmon events detail
Event Index 1:
  Description                          rmon event
  Type                                 log and trap
  Community                            rmon-trap-group
  Last event                           2009-07-10 11:34:17 PDT
  Creator                              CLI
  State                                active
```

show snmp rmon logs (QFX Series)

```
user@host> show snmp rmon logs
Event Index: 1
    Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed, (variable:
```

```
jnxOperatingCPU.9.1.0.0, value: 100)
    Time: 2009-07-10 11:34:07 PDT
    Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed, (variable:
jnxOperatingCPU.9.1.0.0, value: 5)
    Time: 2009-07-10 11:34:17 PDT
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

RMON Overview

clear snmp statistics | 4208

clear snmp history | 4207

show snmp rmon history | 26523

show snmp rmon history

IN THIS SECTION

- Syntax | 26524
- Description | 26524
- Options | 26524
- Required Privilege Level | 26524
- Output Fields | 26524
- Sample Output | 26527
- Release Information | 26529

Syntax

```
show snmp rmon history
<history-index>
sample-index <sample-index>
```

Description

Display the contents of the RMON history group.

Options

none	Display all the entries in the RMON history group.
history-index	(Optional) Display the contents of the specified entry in the RMON history group.
sample-index ***sample-index***	(Optional) Display the statistics collected for the specified sample within the specified entry in the RMON history group.

Required Privilege Level

view

Output Fields

Table 1580 on page 26524 lists the output fields for the `show smp rmon history` command. Output fields are listed in the approximate order in which they appear.

Table 1580: show smp rmon history Output Fields

Field Name	Field Description
History Index	Identifies this RMON history entry within the RMON history group.
Owner	The entity that configured this entry. Range is 0 to 32 alphanumeric characters.
Status	The status of the RMON history entry.

Table 1580: show smp rmon history Output Fields *(Continued)*

Field Name	Field Description
Interface or **Data Source**	The ifndex object that identifies the interface that is being monitored.
Interval	The interval (in seconds) configured for this RMON history entry.
Buckets Requested	The requested number of buckets (**intervals**) configured for this RMON history entry.
Buckets Granted	The number of buckets granted for this RMON history entry.

Table 1580: show smp rmon history Output Fields *(Continued)*

Field Name	Field Description
Sample Index	The sample statistics taken at the specified interval. • **Drop Events**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Octets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • **Packets**—Total number of packets. • **Broadcast Packets**—Number of broadcast packets. • **Multicast Packets**—Number of multicast packets. • **CRC errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS error) or a bad FCS with a nonintegral number of octets (alignment error). • **Undersize Pkts**—Number of packets received during this sampling interval that were less than 64 octets long (excluding framing bits but including FCS octets) and were otherwise well formed. • **Oversize Pkts**—Number of packets received during the sampling interval that were longer than 1518 octets (excluding framing bits, but including FCS octets) but were otherwise well formed. • **Fragments**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **Jabbers**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.

Table 1580: show smp rmon history Output Fields *(Continued)*

Field Name	Field Description
	• **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Utilization(%)**—The best estimate of the mean physical layer network utilization on this interface during this sampling interval, in hundredths of a percent.

Sample Output

show snmp rmon history 1

```
user@host> show snmp rmon history 1

 History Index 1:
 Interface                              171
 Requested Buckets                      50
 Interval                               10

Sample Index 1: Interval Start: Tue Feb 12 04:12:32 2008
    Drop Events         0
    Octets            486
    Packets             2
    Broadcast Packet    0
    Multicast Packets   2
    CRC errors          0
    Undersize Pkts      0
    Oversize Pkts       0
    Fragments           0
    Jabbers             0
    Collisions          0
    Utilization(%)      0

Sample Index 2: Interval Start: Tue Feb 12 04:12:42 2008
    Drop Events         0
    Octets            486
    Packets             2
```

```
    Broadcast Packet    0
    Multicast Packets   2
    CRC errors          0
    Undersize Pkts      0
    Oversize Pkts       0
    Fragments           0
    Jabbers             0
    Collisions          0
    Utilization(%)      0

Sample Index 3: Interval Start: Tue Feb 12 04:12:52 2008
    Drop Events         0
    Octets             486
    Packets             2
    Broadcast Packet    0
    Multicast Packets   2
    CRC errors          0
    Undersize Pkts      0
    Oversize Pkts       0
    Fragments           0
    Jabbers             0
    Collisions          0
    Utilization(%)      0
```

show snmp rmon history 1 sample-index 15

```
user@host> show snmp rmon history 1 sample-index 15

 Index 1
Owner             = monitor
Status             = valid
Data Source        = ifIndex.17
Interval            = 1800
Buckets Requested    = 50
Buckets Granted    = 50


Sample Index 44: Interval Start: Thu Jan  1 00:08:35 1970
Drop Events         = 0
Octetes            = 0
Packets            = 0
```

```
Broadcast Pkts        = 0
Multicast Pkts        = 0
CRC Errors        = 0
Undersize Pkts        = 0
Oversize Pkts        = 0
Fragments        = 0
Jabbers            = 0
Collisions        = 0
Utilization (%)        = 0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

clear snmp history | 4207

RMON MIB Event, Alarm, Log, and History Control Tables

show snmp statistics

IN THIS SECTION

- Syntax | 26530
- Description | 26530
- Options | 26530
- Required Privilege Level | 26530
- Output Fields | 26530
- Sample Output | 26536
- Release Information | 26539

Syntax

```
show snmp statistics
<subagents>
```

Description

Display statistics about Simple Network Management Protocol (SNMP) packets sent and received by the router or switch.

Options

subagents (Optional) Display the statistics of the protocol data unit (PDU), the number of SNMP requests and responses per subagent, and the SNMP statistics received from each subagent per logical system.

Required Privilege Level

view

Output Fields

Table 1581 on page 26531 describes the output fields for the `show snmp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1581: show snmp statistics Output Fields

Field Name	Field Description
Input	Information about received packets: • **Packets(snmpInPkts)**—Total number of messages delivered to the SNMP entity from the transport service. • **Bad versions**—**(snmpInBadVersions**) Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version. • **Bad community names**—**(snmpInBadCommunityNames)** Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity. • **Bad community uses**—**(snmpInBadCommunityUses)** Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message. • **ASN parse errors**—**(snmpInASNParseErrs)** Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages. • **Too bigs**—**(snmpInTooBigs)** Total number of SNMP PDUs delivered to the SNMP entity with an error status field of **tooBig**. • **No such names**—**(snmpInNoSuchNames)** Total number of SNMP PDUs delivered to the SNMP entity with an error status field of **noSuchName**. • **Bad values**—**(snmpInBadValues)** Total number of SNMP PDUs delivered to the SNMP entity with an error status field of **badValue**. • **Read onlys**—**(snmpInReadOnlys)** Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of **readOnly**. Only incorrect implementations of SNMP generate this error.

Table 1581: show snmp statistics Output Fields *(Continued)*

Field Name	Field Description
Input (continued)	• **General errors—(snmpInGenErrs)** Total number of SNMP PDUs delivered to the SNMP entity with an error status field of **genErr**. • **Total requests varbinds—(snmpInTotalReqVars)** Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP **GetRequest** and **GetNext** PDUs. • **Total set varbinds—(snmpInSetVars)** Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP **SetRequest** PDUs. • **Get requests—(snmpInGetRequests)** Total number of SNMP **GetRequest** PDUs that have been accepted and processed by the SNMP entity. • **Get nexts—(snmpInGetNexts)** Total number of SNMP **GetNext** PDUs that have been accepted and processed by the SNMP entity. • **Set requests—(snmpInSetRequests)** Total number of SNMP **SetRequest** PDUs that have been accepted and processed by the SNMP entity. • **Get responses—(snmpInGetResponses)** Total number of SNMP **GetResponse** PDUs that have been accepted and processed by the SNMP entity. • **Traps—(snmpInTraps)** Total number of SNMP traps generated by the SNMP entity. • **Silent drops—(snmpSilentDrops)** Total number of **GetRequest**, **GetNextRequest**, **GetBulkRequest**, **SetRequests**, and **InformRequest** PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests. • **Proxy drops—(snmpProxyDrops)** Total number of **GetRequest**, **GetNextRequest**, **GetBulkRequest**, **SetRequests**, and **InformRequest** PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned.

Table 1581: show snmp statistics Output Fields *(Continued)*

Field Name	Field Description
	• **Commit pending drops**—Number of SNMP packets for **Set** requests dropped because of a previous pending SNMP **Set** request on the committed configuration. • **Throttle drops**—Number of SNMP packets for any requests dropped reaching the throttle limit.

Table 1581: show snmp statistics Output Fields *(Continued)*

Field Name	Field Description
V3 Input	Information about SNMP version 3 packets: • **Unknown security models—(snmpUnknownSecurityModels)** Total number of packets received by the SNMP engine that were dropped because they referenced a security model that was not known to or supported by the SNMP engine. • **Invalid messages—(snmpInvalidMsgs)** Number of packets received by the SNMP engine that were dropped because there were invalid or inconsistent components in the SNMP message. • **Unknown pdu handlers—(snmpUnknownPDUHandlers)** Number of packets received by the SNMP engine that were dropped because the PDU contained in the packet could not be passed to an application responsible for handling the PDU type. • **Unavailable contexts—(snmpUnavailableContexts)** Number of requests received for a context that is known to the SNMP engine, but is currently unavailable. • **Unknown contexts—(snmpUnknownContexts)** Total number of requests received for a context that is unknown to the SNMP engine. • **Unsupported security levels—(usmStatsUnsupportedSecLevels)** Total number of packets received by the SNMP engine that were dropped because they requested a security level unknown to the SNMP engine (or otherwise unavailable). • **Not in time windows—(usmStatsNotInTimeWindows)** Total number of packets received by the SNMP engine that were dropped because they appeared outside the authoritative SNMP engine's window. • **Unknown user names—(usmStatsUnknownUserNames)** Total number of packets received by the SNMP engine that were dropped because they referenced a user that was not known to the SNMP engine. • **Unknown engine ids—(usmStatsUnknownEngineIDs)** Total number of packets received by the SNMP engine that were dropped because they referenced an SNMP engine ID that was not known to the SNMP engine.

Table 1581: show snmp statistics Output Fields *(Continued)*

Field Name	Field Description
	• **Wrong digests—(usmStatsWrongDigests)** Total number of packets received by the SNMP engine that were dropped because they did not contain the expected digest value. • **Decryption errors—(usmStatsDecryptionErrors)** Total number of packets received by the SNMP engine that were dropped because they could not be decrypted.
Output	Information about transmitted packets: • **Packets—(snmpOutPkts)** Total number of messages passed from the SNMP entity to the transport service. • **Too bigs—(snmpOutTooBigs)** Total number of SNMP PDUs generated by the SNMP entity with an error status field of **tooBig**. • **No such names—(snmpOutNoSuchNames)** Total number of SNMP PDUs delivered to the SNMP entity with an error status field of **noSuchName**. • **Bad values—(snmpOutBadValues)** Total number of SNMP PDUs generated by the SNMP entity with an error status field of **badValue**. • **General errors—(snmpOutGenErrs)** Total number of SNMP PDUs generated by the SNMP entity with an error status field of **genErr**. • **Get requests—(snmpOutGetRequests)** Total number of SNMP **GetRequest** PDUs generated by the SNMP entity. • **Get nexts—(snmpOutGetNexts)** Total number of SNMP **GetNext** PDUs generated by the SNMP entity. • **Set requests—(snmpOutSetRequests)** Total number of SNMP **SetRequest** PDUs generated by the SNMP entity. • **Get responses—(snmpOutGetResponses)** Total number of SNMP **GetResponse** PDUs generated by the SNMP entity. • **Traps—(snmpOutTraps)** Total number of SNMP traps generated by the SNMP entity.

Table 1582 on page 26536 describes the output fields for the `show snmp statistics subagents` command. Output fields are listed in the approximate order in which they appear.

Table 1582: show snmp statistics subagents Output Fields

Field Name	Field Description
Subagent	Location of the SNMP subagent.
Request PDUs	Number of PDUs requested by the SNMP manager.
Response PDUs	Number of response PDUs sent by the SNMP subagent.
Request Variables	Number of variable bindings on the PDUs requested by the SNMP manager.
Response Variables	Number of variable bindings on the PDUs sent by the SNMP subagent.
Average Response Time	Average time taken by the SNMP subagent to send statistics response.
Maximum Response Time	Maximum time taken by the SNMP subagent to send the statistics response.

Sample Output

show snmp statistics

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 246213, Bad versions: 12, Bad community names: 12,
    Bad community uses: 0, ASN parse errors: 96,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 227084, Total set varbinds: 67,
    Get requests: 44942, Get nexts: 190371, Set requests: 10712,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0,
  V3 Input:
```

```
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 1
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
  Output:
    Packets: 246093, Too bigs: 0, No such names: 31561,
    Bad values: 0, General errors: 2,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 246025, Traps: 0
```

show snmp statistics subagents

```
user@host> show snmp statistics subagents

  Subagent: /var/run/cosd-20
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/pfed-30
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/rmopd-15
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/chassisd-30
    Request PDUs: 33116, Response PDUs: 33116,
    Request Variables: 33116, Response Variables: 33116,
    Average Response Time(ms): 1.83,
    Maximum Response Time(ms): 203.48

  Subagent: /var/run/pkid-13
    Request PDUs: 0, Response PDUs: 0,
```

```
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/apsd-13
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/dfcd-32
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/mib2d-33
    Request PDUs: 74211, Response PDUs: 74211,
    Request Variables: 74211, Response Variables: 74211,
    Average Response Time(ms): 2.30,
    Maximum Response Time(ms): 51.04

  Subagent: /var/run/license-check-16
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/craftd-14
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/bfdd-19
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/smihelperd-24
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
```

```
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/cfmd-18
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/rpd_snmp
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00

  Subagent: /var/run/l2tpd-18
    Request PDUs: 0, Response PDUs: 0,
    Request Variables: 0, Response Variables: 0,
    Average Response Time(ms): 0.00,
    Maximum Response Time(ms): 0.00
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `subagents` introduced in Junos OS Release 14.2.

show snmp stats-response-statistics

IN THIS SECTION

- Syntax | 26540
- Syntax (PTX Series for Junos OS Evolved) | 26540
- Description | 26540
- Options | 26540
- Required Privilege Level | 26540

- Output Fields | 26540
- Sample Output | 26541
- Sample Output (PTX Series for Junos OS Evolved) | 26543
- Release Information | 26543

Syntax

```
show snmp stats-response-statistics
```

Syntax (PTX Series for Junos OS Evolved)

```
show snmp stats-response-statistics
```

Description

Display statistics of SNMP statistics responses sent from the Packet Forwarding Engine during the MIB II process (mib2d).

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1583 on page 26541 describes the output fields for the `show snmp stats-response-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1583: show snmp stats-response-statistics Output Fields

Field Name	Field Description
`Average response time statistics`	Display the average response time in milliseconds per protocol data unit (PDU) by snmpd. It includes the following information: • `Stats Type`—Type of SNMP statistics. • `Stats Responses`—Number of SNMP statistics responses received from the Packet Forwarding Engine. • `Average Response Time`—Average time taken to receive the statistics response from the Packet Forwarding Engine in milliseconds.
`Bucket statistics`	Information about SNMP statistics responses: • `Bucket Type`—Category of time intervals in which SNMP statistics responses are received from the Packet Forwarding Engine. • `Stats Responses`—Number of SNMP statistics responses received from the Packet Forwarding Engine.
`Bad responses`	Information about top 20 bad responses from a subagent: • `Response`—Time taken to receive the SNMP statistics response from the Packet Forwarding Engine in milliseconds. • `Request Time`—Date and time of SNMP request. • `Key`—Display the attribute of SNMP `Stats Type`. For example, in the case of SNMP statistics responses for interfaces, the Key value is SNMP ifIndex, and for firewalls, the Key value is the filter name.

Sample Output

show snmp stats-response-statistics

```
user@host> show snmp stats-response-statistics
Average response time statistics:
Stats                Stats                    Average
Type                 Responses                Response
```

```
                                                Time (ms)
ifd(non ae)          34182                      175.48
ifd(ae)              0                          0.00
ifl(non ae)          5472                       5.40
ifl(ae)              0                          0.00
firewall             15                         1141.73

Bucket statistics:
Bucket               Stats
Type(ms)             Responses
0 - 10               39078
11 - 50              588
51 - 100             0
101 - 200            0
201 - 500            1
501 - 1000           2
1001 - 2000          0
2001 - 5000          0
More than 5001       0

Bad responses:
Response        Request                 Stats          Key
                Time                    Type
(ms)            (UTC)
927.80          2014-03-26 05:44:16     firewall       __default_arp_policer__
908.68          2014-03-26 05:44:16     firewall       __default_bpdu_filter__
421.00          2014-03-26 05:46:25     ifd(non ae)    504
49.76           2014-04-13 04:15:18     ifd(non ae)    503
49.62           2014-04-13 04:30:18     ifd(non ae)    504
48.52           2014-04-05 10:06:55     ifd(non ae)    504
47.61           2014-04-11 04:06:27     ifd(non ae)    505
47.38           2014-04-13 03:30:18     ifd(non ae)    501
47.22           2014-03-27 20:08:07     ifd(non ae)    502
46.26           2014-03-31 13:08:58     ifd(non ae)    506
46.00           2014-04-13 04:00:18     ifd(non ae)    503
45.95           2014-04-05 17:15:17     ifd(non ae)    503
45.75           2014-04-15 13:06:10     ifd(non ae)    507
45.60           2014-04-01 03:07:28     ifd(non ae)    517
45.56           2014-04-08 13:09:15     ifd(non ae)    502
45.23           2014-04-13 03:15:18     ifd(non ae)    501
45.15           2014-04-05 16:45:17     ifd(non ae)    501
44.74           2014-04-08 22:08:47     ifd(non ae)    505
```

```
44.10           2014-04-05 16:30:17    ifd(non ae)    501
44.00           2014-04-08 09:09:23    ifd(non ae)    524
```

Sample Output (PTX Series for Junos OS Evolved)

show snmp stats-response-statistics

```
user@host> show snmp stats-response-statistics
Average response time statistics:
Stats                Stats                    Average
Type                 Responses                Response
                                              Time (ms)
ifd(non ae)          105                      1.31
ifd(ae)              0                        0.00
ifl(non ae)          70                       0.73
ifl(ae)              0                        0.00

Bucket statistics:
Bucket               Stats
Type(ms)             Responses
0 - 10               209
11 - 50              20
51 - 100             0
101 - 200            0
201 - 500            0
501 - 1000           0
1001 - 2000          0
2001 - 5000          0
More than 5001       0

Bad responses:
Response        Request                Stats          Key
                Time                   Type
(ms)            (UTC)
34.49           2021-08-06 03:01:17    ifd(non ae)    529
10.62           2021-08-06 03:01:17    ifd(non ae)    530
```

Release Information

Command introduced in Junos OS Release 14.2.

show snmp v3

IN THIS SECTION

- Syntax | 26544
- Description | 26544
- Options | 26544
- Additional Information | 26545
- Required Privilege Level | 26545
- Output Fields | 26545
- Sample Output | 26547
- Release Information | 26548

Syntax

```
show snmp v3
<access <brief | detail> | community | general | groups | notify <filter> | target <address |
parameters> | users>
```

Description

Display the Simple Network Management Protocol version 3 (SNMPv3) operating configuration.

Options

none	Display all of the SNMPv3 operating configuration.
access	(Optional) Display SNMPv3 access information.
brief | detail	(Optional) Display brief or detailed information about SNMPv3 access information.
community	(Optional) Display SNMPv3 community information.
general	(Optional) Display SNMPv3 general information.

groups	(Optional) Display SNMPv3 security-to-group information.
notify <filter>	(Optional) Display SNMPv3 notify and, optionally, notify filter information.
target <address | parameters>	(Optional) Display SNMPv3 target and, optionally, either target address or target parameter information.
users	(Optional) Display SNMPv3 user information.

Additional Information

To edit the default display of the `show snmp v3` command, specify options in the `show` statement at the `[edit snmp v3]` hierarchy level.

Required Privilege Level

view

Output Fields

Table 1584 on page 26545 describes the output fields for the `show snmp v3` command. Output fields are listed in the approximate order in which they appear.

Table 1584: show snmp v3 Output Fields

Field Name	Field Description
Access control	Information about access control: • **Group**—Group name for which the configured access privileges apply. The group, together with the context prefix and the security model and security level, forms the index for this table. • **Context prefix**—SNMPv3 context for which the configured access privileges apply. • **Security model/level**—Security model and security level for which the configuration access privileges apply. • **Read view**—Identifies the MIB view applied to SNMPv3 read operations. • **Write view**—Identifies the MIB view applied to SNMPv3 write operations. • **Notify view**—Identifies the MIB view applied to outbound SNMP notifications.

Table 1584: show snmp v3 Output Fields *(Continued)*

Field Name	Field Description
Engine	Information about local engine configuration: • **Local engine ID**—Identifier that uniquely and unambiguously identifies the local SNMPv3 engine. • **Engine boots**—Number of times the local SNMPv3 engine has rebooted or reinitialized since the engine ID was last changed. • **Engine time**—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized. • **Max msg size**—Maximum message size the sender can accommodate.
Engine ID	Information about engine ID: • **Local engine ID**—Identifier that uniquely and unambiguously identifies the local SNMPv3 engine. • **Engine boots**—Number of times the local SNMPv3 engine has rebooted or reinitialized since the engine ID was last changed. • **Engine time**—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized. • **Max msg size**—Maximum message size the sender can accommodate. • **Engine ID**—SNMPv3 engine ID associated with each user. • **User**—SNMPv3 user. • **Auth/Priv**—Authentication and encryption algorithm available for use by each user. • **Storage**—Indicates whether a user is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status. • **Status**—Status of the conceptual row. Only rows with an active status are used by the SNMPv3 engine.
Group name	Name of the group to which this entry belongs.
Security model	Identifies the security model context for the security name.

Table 1584: show snmp v3 Output Fields *(Continued)*

Field Name	Field Description
Security name	Used with the security model; identifies a specific security name instance. Each security model/security name combination can be assigned to a specific group.
Storage type	Indicates whether a user is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status.
Status	Status of the conceptual row. Only rows with active status are used by the SNMPv3 engine.

Sample Output

show snmp v3

```
user@host> show snmp v3
Local engine ID: 80 00 0a 4c e04 31 32 33 34
Engine boots:          38
Engine time:        64583 seconds
Max msg size:        2048 bytes

Engine ID: local
    User                            Auth/Priv   Storage      Status
    user1                            md5/des    nonvolatile  active
    user2                            sha/none   nonvolatile  active
    user3                           none/none   nonvolatile  active

Engine ID: 81 00 0a 4c 04 64 64 64 64
    User                            Auth/Priv   Storage      Status
    UNEW                             md5/none   nonvolatile  active
Group name           Security  Security              Storage      Status
                     model     name                  type
g1                   usm       user1                 nonvolatile  active
g2                   usm       user2                 nonvolatile  active
g3                   usm       user3                 nonvolatile  active

Access control:
Group                Context Security      Read       Write     Notify
```

```
                    prefix  model/level   view       view      view
g1                           usm/privacy  v1         v1
g2                           usm/authent  v1         v1
g3                           usm/none     v1         v1
```

Release Information

Command introduced before Junos OS Release 7.4.

show spanning-tree bridge

IN THIS SECTION

- Syntax | 26548
- Syntax (QFX Series and EX Series) | 26549
- Description | 26549
- Options | 26549
- Required Privilege Level | 26549
- Output Fields | 26549
- Sample Output | 26551
- Release Information | 26555

Syntax

```
show spanning-tree bridge
<brief | detail>
<msti msti-id>
<routing-instance routing-instance-name>
<vlan-id vlan-id>
```

Syntax (QFX Series and EX Series)

```
show spanning-tree bridge
<brief | detail>
<msti msti-id>
<vlan-id vlan-id>
```

Description

Displays the configured or calculated spanning-tree protocol (can be either STP, RSTP, MSTP, or VSTP) parameters.

Options

none	(Optional) Display brief STP bridge information for all multiple spanning-tree instances (MSTIs).
brief | detail	(Optional) Display the specified level of output.
msti ***msti-id***	(Optional) Display spanning-tree protocol bridge information for the specified MSTI or Common and Internal Spanning Tree (CIST). Specify **0** for CIST. Specify a value from **1** through **4094** for an MSTI.
routing-instance ***routing-instance-name***	(Optional) Display STP bridge information for the specified routing instance.
vlan-id ***vlan-id***	(Optional) Display spanning-tree protocol bridge information for the specified VLAN. Specify a VLAN tag identifier from **1** through **4094**.

Required Privilege Level

view

Output Fields

Table 1585 on page 26550 lists the output fields for the `show spanning-tree bridge` command. Output fields are listed in the approximate order in which they appear.

Table 1585: show spanning-tree bridge Output Fields

Field Name	Field Description
`Routing instance name`	Name of the routing instance under which the bridge is configured.
`Enabled protocol`	Spanning Tree Protocol type enabled.
`Root ID`	Bridge ID of the elected spanning-tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
`Root cost`	Calculated cost to reach the root bridge from the bridge where the command is entered.
`Root port`	Interface that is the current elected root port for this bridge.
`CIST regional root`	Bridge ID of the elected MSTP regional root bridge.
`CIST internal root cost`	Calculated cost to reach the regional root bridge from the bridge where the command is entered.
`Hello time`	Configured number of seconds between transmissions of configuration BPDUs.
`Maximum age`	Configured maximum expected arrival time of hello bridge protocol data units (BPDUs).
`Forward delay`	How long an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.
`Hop count`	Configured maximum number of hops a BPDU can be forwarded in the MSTP region.
`Message age`	Number of elapsed seconds since the most recent BPDU was received.

Table 1585: show spanning-tree bridge Output Fields *(Continued)*

Field Name	Field Description
Number of topology changes	Total number of STP topology changes detected since the routing device last booted.
Time since last topology change	Number of elapsed seconds since the most recent topology change.
Bridge ID (Local)	Locally configured bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Extended system ID	System identifier.
MSTI regional root	Bridge ID of the elected MSTP regional root bridge.

Sample Output

show spanning-tree bridge routing-instance

```
user@host> show spanning-tree bridge routing-instance vs1 detail
STP bridge parameters
Routing instance name               : vs1
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 32768.00:13:c3:9e:c8:80
  Root cost                         : 0
  Root port                         : ge-10/2/0
  CIST regional root                : 32768.00:13:c3:9e:c8:80
  CIST internal root cost           : 22000
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 18
  Message age                       : 0
  Number of topology changes        : 1
  Time since last topology change   : 1191 seconds
```

```
  Local parameters
    Bridge ID                       : 32768.00:90:69:0b:7f:d1
    Extended system ID              : 1

STP bridge parameters for MSTI 1
  MSTI regional root                : 32769.00:13:c3:9e:c8:80
  Root cost                         : 22000
  Root port                         : ge-10/2/0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 18
  Number of topology changes        : 1
  Time since last topology change   : 1191 seconds
  Local parameters
    Bridge ID                       : 32769.00:90:69:0b:7f:d1
    Extended system ID              : 1

STP bridge parameters for MSTI 2
  MSTI regional root                : 32770.00:13:c3:9e:c8:80
  Root cost                         : 22000
  Root port                         : ge-10/2/0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 18
  Number of topology changes        : 1
  Time since last topology change   : 1191 seconds
  Local parameters
    Bridge ID                       : 32770.00:90:69:0b:7f:d1
    Extended system ID              : 1
```

show spanning-tree bridge msti

```
user@host> show spanning-tree bridge msti 1 routing-instance vs1 detail
STP bridge parameters
Routing instance name               : vs1
Enabled protocol                    : MSTP

STP bridge parameters for MSTI 1
  MSTI regional root                : 32769.00:13:c3:9e:c8:80
```

```
  Root cost                          : 22000
  Root port                          : xe-10/2/0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 18
  Number of topology changes         : 1
  Time since last topology change    : 1191 seconds
  Local parameters
    Bridge ID                        : 32769.00:90:69:0b:7f:d1
    Extended system ID               : 1
```

show spanning-tree bridge vlan-id (MSTP)

```
user@host> show spanning-tree bridge vlan-id 1 101 routing-instance vs1 detail
STP bridge parameters
Routing instance name                : vs1
Enabled protocol                     : MSTP

STP bridge parameters for CIST
  Root ID                            : 32768.00:13:c3:9e:c8:80
  Root cost                          : 0
  Root port                          : xe-10/2/0
  CIST regional root                 : 32768.00:13:c3:9e:c8:80
  CIST internal root cost            : 22000
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 18
  Message age                        : 0
  Number of topology changes         : 0
  Local parameters
    Bridge ID                        : 32768.00:90:69:0b:7f:d1
    Extended system ID               : 1
    Hello time                       : 2 seconds
    Maximum age                      : 20 seconds
    Forward delay                    : 15 seconds
    Path cost method                 : 32 bit
    Maximum hop count                : 20
```

show spanning-tree bridge (RSTP)

```
user@host> show spanning-tree bridge
STP bridge parameters
Routing instance name               : GLOBAL
Enabled protocol                    : RSTP
  Root ID                           : 28672.00:90:69:0b:3f:d0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Message age                       : 0
  Number of topology changes        : 58
  Time since last topology change   : 14127 seconds
  Local parameters
    Bridge ID                       : 28672.00:90:69:0b:3f:d0
    Extended system ID              : 0

 STP bridge parameters for bridge VLAN 10
   Root ID                          : 28672.00:90:69:0b:3f:d0
   Hello time                       : 2 seconds
   Maximum age                      : 20 seconds
   Forward delay                    : 15 seconds
   Message age                      : 0
   Number of topology changes       : 58
   Time since last topology change : 14127 seconds
   Local parameters
     Bridge ID                      : 28672.00:90:69:0b:3f:d0
     Extended system ID             : 0

 STP bridge parameters for bridge VLAN 20
   Root ID                          : 28672.00:90:69:0b:3f:d0
   Hello time                       : 2 seconds
   Maximum age                      : 20 seconds
   Forward delay                    : 15 seconds
   Message age                      : 0
   Number of topology changes       : 58
   Time since last topology change : 14127 seconds
   Local parameters
     Bridge ID                      : 28672.00:90:69:0b:3f:d0
     Extended system ID             : 0
```

show spanning-tree bridge vlan-id (RSTP)

```
user@host> show spanning-tree bridge vlan-id 10
STP bridge parameters
Routing instance name               : GLOBAL
Enabled protocol                    : RSTP

 STP bridge parameters for VLAN 10
   Root ID                          : 28672.00:90:69:0b:3f:d0
   Hello time                       : 2 seconds
   Maximum age                      : 20 seconds
   Forward delay                    : 15 seconds
   Message age                      : 0
   Number of topology changes       : 58
   Time since last topology change : 14127 seconds
   Local parameters
     Bridge ID                      : 28672.00:90:69:0b:3f:d0
     Extended system ID             : 0
```

Release Information

Command introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Understanding STP

Understanding RSTP

Example: Configuring Network Regions for VLANs with MSTP on Switches

show spanning-tree interface | 26555

show spanning-tree interface

IN THIS SECTION

Syntax | 26556

- Syntax (EX Series Switches and QFX Series Switches) | 26556
- Syntax (EX Series Switches) | 26556
- Description | 26557
- Options | 26557
- Required Privilege Level | 26557
- Output Fields | 26557
- Sample Output | 26558
- Release Information | 26568

Syntax

```
show spanning-tree interface
<brief | detail>
<msti msti-id>
<routing-instance routing-instance-name>
<vlan-id vlan-id>
```

Syntax (EX Series Switches and QFX Series Switches)

```
show spanning-tree interface
<brief | detail>
<msti msti-id>
<vlan-id vlan-id>
```

Syntax (EX Series Switches)

```
show spanning-tree interface
<brief | detail>
<interface-name interface-name>
<msti msti-id>
<vlan-id vlan-id>
```

Description

Display the configured or calculated interface-level spanning-tree protocol (can be either STP, RSTP, or MSTP) parameters. In **brief** mode, will not display interfaces that are administratively disabled or do not have a physical link.

Options

none	Display brief STP interface information.
brief | detail	(Optional) Display the specified level of output.
interface-name *interface-name*	(Optional) Name of an interface.
msti *msti-id*	(Optional) Display STP interface information for the specified MST instance.
routing-instance *routing-instance-name*	(Optional) Display STP interface information for the specified routing instance.
vlan-id *vlan-id*	(Optional) Display STP interface information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 1586 on page 26557 lists the output fields for the `show spanning-tree interface` command. Output fields are listed in the approximate order in which they appear.

Table 1586: show spanning-tree Interface Output Fields

Field Name	Field Description
Interface name	Interface configured to participate in the STP, RSTP, VSTP, or MSTP instance.
Port ID	Logical interface identifier configured to participate in the MSTP or VSTP instance.

Table 1586: show spanning-tree Interface Output Fields *(Continued)*

Field Name	Field Description
Designated port ID	Port ID of the designated port for the LAN segment to which this interface is attached.
Designated bridge ID	Bridge ID of the designated bridge for the LAN segment to which this interface is attached.
Port Cost	Configured cost for the interface.
Port State	STP port state: forwarding (**FWD**), blocking (**BLK**), listening, learning, or disabled.
Port Role	MSTP, VSTP, or RSTP port role: designated (**DESG**), backup (**BKUP**), alternate (**ALT**), (**ROOT**), or Root Prevented (**Root-Prev**).
Link type	MSTP, VSTP, or RSTP link type. Shared or point-to-point (pt-pt) and edge or nonedge.
Alternate	Identifies the interface as an MSTP, VSTP, or RSTP alternate root port (**Yes**) or nonalternate root port (**No**).
Boundary Port	Identifies the interface as an MSTP regional boundary port (**Yes**) or nonboundary port (**No**).
Edge delay while expiry count	Number of times the edge delay timer expired on that interface.
Rcvd info while expiry count	Number of times the rcvd info timer expired on that interface.

Sample Output

show spanning-tree interface

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

```
Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae1              128:1        128:1  32768.0090690b47d1      1000  FWD    DESG
ge-2/1/2         128:2        128:2  32768.0090690b47d1     20000  FWD    DESG
ge-2/1/5         128:3        128:3  32768.0090690b47d1     29999  FWD    DESG
ge-2/2/1         128:4       128:26  32768.0013c39ec880     20000  FWD    ROOT
xe-9/2/0         128:5        128:5  32768.0090690b47d1      2000  FWD    DESG
xe-9/3/0         128:6        128:6  32768.0090690b47d1      2000  FWD    DESG

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae1              128:1        128:1  32769.0090690b47d1      1000  FWD    DESG
ge-2/1/2         128:2        128:2  32769.0090690b47d1     20000  FWD    DESG
ge-2/1/5         128:3        128:3  32769.0090690b47d1     29999  FWD    DESG
ge-2/2/1         128:4       128:26  32769.0013c39ec880     20000  FWD    ROOT
xe-9/2/0         128:5        128:5  32769.0090690b47d1      2000  FWD    DESG
xe-9/3/0         128:6        128:6  32769.0090690b47d1      2000  FWD    DESG

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae1              128:1        128:1  32770.0090690b47d1      1000  FWD    DESG
ge-2/1/2         128:2        128:2  32770.0090690b47d1     20000  FWD    DESG
ge-2/1/5         128:3        128:3  32770.0090690b47d1     29999  FWD    DESG
ge-2/2/1         128:4       128:26  32770.0013c39ec880     20000  FWD    ROOT
xe-9/2/0         128:5        128:5  32770.0090690b47d1      2000  FWD    DESG
xe-9/3/0         128:6        128:6  32770.0090690b47d1      2000  FWD    DESG
```

show spanning-tree interface (QFX Series)

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae1              128:1        128:1  32768.0090690b47d1      1000  FWD    DESG
ge-2/1/2         128:2        128:2  32768.0090690b47d1     20000  FWD    DESG
```

```
ge-2/1/5         128:3        128:3  32768.0090690b47d1     29999  FWD    DESG
ge-2/2/1         128:4       128:26  32768.0013c39ec880     20000  FWD    ROOT
xe-9/2/0         128:5        128:5  32768.0090690b47d1      2000  FWD    DESG
xe-9/3/0         128:6        128:6  32768.0090690b47d1      2000  FWD    DESG

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae1              128:1        128:1  32769.0090690b47d1      1000  FWD    DESG
ge-2/1/2         128:2        128:2  32769.0090690b47d1     20000  FWD    DESG
ge-2/1/5         128:3        128:3  32769.0090690b47d1     29999  FWD    DESG
ge-2/2/1         128:4       128:26  32769.0013c39ec880     20000  FWD    ROOT
xe-9/2/0    128:5        128:5  32769.0090690b47d1      2000  FWD    DESG
xe-9/3/0    128:6        128:6  32769.0090690b47d1      2000  FWD    DESG

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae1              128:1        128:1  32770.0090690b47d1      1000  FWD    DESG
ge-2/1/2         128:2        128:2  32770.0090690b47d1     20000  FWD    DESG
ge-2/1/5         128:3        128:3  32770.0090690b47d1     29999  FWD    DESG
ge-2/2/1         128:4       128:26  32770.0013c39ec880     20000  FWD    ROOT
xe-9/2/0    128:5        128:5  32770.0090690b47d1      2000  FWD    DESG
xe-9/3/0    128:6        128:6  32770.0090690b47d1      2000  FWD    DESG
```

show spanning-tree interface (EX Series)

```
user@switch> show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated  Designated     Port    State  Role
                         port ID    bridge ID      Cost
ge-0/0/0.0  128:513  128:513   8192.0019e2500340    1000   FWD    DESG
ge-0/0/2.0  128:515  128:515   8192.0019e2500340    1000   BLK    DIS
ge-0/0/4.0  128:517  128:517   8192.0019e2500340    1000   FWD    DESG
ge-0/0/23.0 128:536  128:536   8192.0019e2500340    1000   FWD    DESG
```

```
Spanning tree interface parameters for instance 1

Interface    Port ID    Designated  Designated     Port    State  Role
                         port ID    bridge ID      Cost
ge-0/0/0.0  128:513  128:513   8193.0019e2500340    1000   FWD    DESG
ge-0/0/2.0  128:515  128:515   8193.0019e2500340    1000   BLK    DIS
ge-0/0/4.0  128:517  128:517   8193.0019e2500340    1000   FWD    DESG
ge-0/0/23.0 128:536  128:536   8193.0019e2500340    1000   FWD    DESG

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated  Designated     Port    State  Role
                         port ID    bridge ID      Cost
ge-0/0/0.0  128:513  128:1   8194.001b549fd000  1000  FWD    ROOT
ge-0/0/2.0  128:515  128:515  32770.0019e2500340  4000  BLK    DIS
ge-0/0/4.0  128:517  128:1    16386.001b54013080  1000  BLK    ALT
ge-0/0/23.0 128:536  128:536  32770.0019e2500340  1000  FWD    DESG
```

show spanning-tree interface detail

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0

Interface name                 : ae1
Port identifier                : 128.1
Designated port ID             : 128.1
Port cost                      : 1000
Port state                     : Forwarding
Designated bridge ID           : 32768.00:90:69:0b:47:d1
Port role                      : Designated
Link type                      : Pt-Pt/NONEDGE
Boundary port                   : No

Interface name                 : ge-2/1/2
Port identifier                : 128.2
Designated port ID             : 128.2
Port cost                      : 20000
Port state                     : Forwarding
Designated bridge ID           : 32768.00:90:69:0b:47:d1
Port role                      : Designated
```

```
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No

Interface name                  : ge-2/1/5
Port identifier                 : 128.3
Designated port ID              : 128.3
Port cost                       : 29999
Port state                      : Forwarding
Designated bridge ID            : 32768.00:90:69:0b:47:d1
Port role                       : Designated
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No

Interface name                  : ge-2/2/1
Port identifier                 : 128.4
Designated port ID              : 128.26
Port cost                       : 20000
Port state                      : Forwarding
Designated bridge ID            : 32768.00:13:c3:9e:c8:80
Port role                       : Root
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No

Interface name                  : xe-9/2/0
Port identifier                 : 128.5
Designated port ID              : 128.5
Port cost                       : 2000
Port state                      : Forwarding
Designated bridge ID            : 32768.00:90:69:0b:47:d1
Port role                       : Designated
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No

Interface name                  : xe-9/3/0
Port identifier                 : 128.6
Designated port ID              : 128.6
Port cost                       : 2000
Port state                      : Forwarding
Designated bridge ID            : 32768.00:90:69:0b:47:d1
Port role                       : Designated
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No
```

```
Spanning tree interface parameters for instance 1

Interface name                  : ae1
Port identifier                 : 128.1
Designated port ID              : 128.1
Port cost                       : 1000
Port state                      : Forwarding
Designated bridge ID            : 32768.00:90:69:0b:47:d1
Port role                       : Designated
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No

Interface name                  : ge-2/1/2
Port identifier                 : 128.2
Designated port ID              : 128.2
Port cost                       : 20000
Port state                      : Forwarding
Designated bridge ID            : 32768.00:90:69:0b:47:d1
Port role                       : Designated
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No

Interface name                  : ge-2/1/5
Port identifier                 : 128.3
Designated port ID              : 128.3
Port cost                       : 29999
Port state                      : Forwarding
Designated bridge ID            : 32768.00:90:69:0b:47:d1
Port role                       : Designated
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No

Interface name                  : ge-2/2/1
Port identifier                 : 128.4
Designated port ID              : 128.26
Port cost                       : 20000
Port state                      : Forwarding
Designated bridge ID            : 32768.00:13:c3:9e:c8:80
Port role                       : Root
Link type                       : Pt-Pt/NONEDGE
Boundary port                    : No
```

```
...
```

show spanning-tree interface detail (EX Series)

```
user@switch> show spanning-tree interface detail
Spanning tree interface parameters for instance 0

Interface name     : ge-1/0/0.0
Port identifier    : 128.625
Designated port ID     : 128.625
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/EDGE
Boundary port       : NA
Edge delay while expiry count        : 0
Rcvd info while expiry count            : 0

Interface name     : ge-1/0/1.0
Port identifier    : 128.626
Designated port ID     : 128.626
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA
Edge delay while expiry count        : 0
Rvcd info while expiry count            : 0

Interface name     : ge-1/0/2.0
Port identifier    : 128.627
Designated port ID     : 128.627
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA
```

```
Edge delay while expiry count        : 0
Rvcd info while expiry count            : 0

Interface name     : ge-1/0/10.0
Port identifier    : 128.635
Designated port ID     : 128.635
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA
Edge delay while expiry count        : 0
Rvcd info while expiry count            : 0

Interface name     : ge-1/0/20.0
Port identifier    : 128.645
Designated port ID     : 128.645
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA
Edge delay while expiry count        : 0
Rvcd info while expiry count            : 0
[output truncated]
```

show spanning-tree interface msti

```
user@host> show spanning-tree interface msti 1 routing-instance vs1 detail
 Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
xe-7/0/0         128:1        128:1  32769.0090690b4fd1      2000  FWD    DESG
ge-5/1/0         128:2        128:2  32769.0090690b4fd1     20000  FWD    DESG
ge-5/1/1         128:3        128:3  32769.0090690b4fd1     20000  FWD    DESG
ae1              128:4        128:1  32769.0090690b47d1     10000  BLK    ALT
```

```
ge-5/1/4         128:5        128:3  32769.0090690b47d1     20000  BLK    ALT
xe-7/2/0         128:6        128:6  32769.0090690b47d1      2000  FWD    ROOT
```

show spanning-tree interface vlan-id

```
user@host> show spanning-tree interface vlan-id 101 routing-instance vs1 detail
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-11/0/5        128:1        128:1  32768.0090690b7fd1     20000  FWD    DESG
ge-11/0/6        128:2        128:1  32768.0090690b7fd1     20000  BLK    BKUP
ge-11/1/0        128:3        128:2  32768.0090690b4fd1     20000  BLK    ALT
ge-11/1/1        128:4        128:3  32768.0090690b4fd1     20000  BLK    ALT
ge-11/1/4        128:5        128:1  32768.0090690b47d1     20000  BLK    ALT
xe-10/0/0        128:6        128:5  32768.0090690b4fd1      2000  BLK    ALT
xe-10/2/0        128:7        128:4  32768.0090690b47d1      2000  FWD    ROOT
```

show spanning-tree interface (VSTP)

```
user@host> show spanning-tree interface
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Cost    State  Role
                         port ID        bridge ID
ge-1/0/1         128:1        128:1  28672.0090690b3fe0     20000  FWD    DESG
ge-1/0/2         128:2        128:2  28672.0090690b3fe0     20000  FWD    DESG

Spanning tree interface parameters for VLAN 10

Interface    Port ID    Designated      Designated         Cost    State  Role
                         port ID        bridge ID
ge-1/0/1         128:1        128:1  28672.0090690b3fe0     20000  FWD    DESG
ge-1/0/2         128:2        128:2  28672.0090690b3fe0     20000  FWD    DESG

Spanning tree interface parameters for VLAN 20

Interface    Port ID    Designated      Designated         Cost    State  Role
                         port ID        bridge ID
```

```
ge-1/0/1         128:1        128:1  28672.0090690b3fe0     20000  FWD    DESG
ge-1/0/2         128:2        128:2  28672.0090690b3fe0     20000  FWD    DESG
```

show spanning-tree interface vlan-id (VSTP)

```
user@host> show spanning-tree  interface vlan-id 10
Spanning tree interface parameters for VLAN 10

Interface    Port ID    Designated      Designated         Cost    State  Role
                         port ID        bridge ID
ge-1/0/1         128:1        128:1  28672.0090690b3fe0     20000  FWD    DESG
ge-1/0/2         128:2        128:2  28672.0090690b3fe0     20000  FWD    DESG
```

show spanning-tree interface brief (EX Series)

```
user@switch> show spanning-tree interface brief
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated  Designated     Port    State  Role
     port ID    bridge ID  Cost
ge-1/0/0.0  128:625  128:625  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/1.0  128:626  128:626  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/2.0  128:627  128:627  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/10.0 128:635  128:635  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/20.0 128:645  128:645  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/30.0 128:655  128:655  32768.0019e25095a0 20000  BLK    DIS
```

show spanning-tree interface ge-1/0/0 (EX Series)

```
user@switch> show spanning-tree interface ge-1/0/0

Interface    Port ID    Designated  Designated     Port    State  Role
     port ID    bridge ID  Cost
ge-1/0/0.0 128:625  128:625  32768.0019e25095a0 20000  BLK    DIS
```

Release Information

Command introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

- show spanning-tree bridge | 26548
- *Example: Configuring Network Regions for VLANs with MSTP on Switches*
- *Understanding STP*
- *Understanding RSTP*
- *Understanding MSTP*

show spanning-tree mstp configuration

IN THIS SECTION

- Syntax | 26568
- Syntax (EX Series Switch and the QFX Series) | 26569
- Description | 26569
- Options | 26569
- Required Privilege Level | 26569
- Output Fields | 26569
- Sample Output | 26570
- Release Information | 26571

Syntax

```
show spanning-tree mstp configuration
<brief | detail>
<routing-instance routing-instance-name>
```

Syntax (EX Series Switch and the QFX Series)

```
show spanning-tree mstp configuration
<brief | detail>
```

Description

Display the MSTP configuration.

Options

none	Display MSTP configuration information.
brief | detail	(Optional) Display the specified level of output.
routing-instance *routing-instance-name*	(Optional) Display MSTP configuration information for the specified routing instance.

Required Privilege Level

view

Output Fields

Table 1587 on page 26569 lists the output fields for the `show spanning-tree mstp configuration` command. Output fields are listed in the approximate order in which they appear.

Table 1587: show spanning-tree mstp configuration Output Fields

Field Name	Field Description
Context id	Internally generated identifier.
Region name	MSTP region name carried in the MSTP BPDUs.
Revision	Revision number of the MSTP configuration.

Table 1587: show spanning-tree mstp configuration Output Fields *(Continued)*

Field Name	Field Description
Configuration digest	Numerical value derived from the VLAN-to-instance mapping table.
MSTI	MST instance identifier.
Member VLANs	VLAN identifiers associated with the MSTI.

Sample Output

show spanning-tree mstp configuration detail

```
user@host> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier     : 1
Region name            : henry
Revision               : 3
Configuration digest   : 0x6da4b5c4fd587757eef35675365e1


MSTI     Member VLANs
   0 0-99,101-199,201-4094
   1 100
   2 200
```

show spanning-tree mstp configuration detail (QFX Series)

```
user@lf0> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier     : 1
Region name            : henry
Revision               : 3
Configuration digest   : 0x6da4b5c4fd587757eef35675365e1
```

```
MSTI     Member VLANs
   0 0-99,101-199,201-4094
   1 100
   2 200
```

show spanning-tree mstp configuration (EX Series)

```
user@host> show spanning-tree mstp configuration
MSTP configuration information
Context identifier     : 0
Region name            : region1
Revision               : 0
Configuration digest   : 0xc92e7af9febb44d8df928b87f16b


MSTI     Member VLANs
   0 0-100,105-4094
   1 101-102
   2 103-104
```

Release Information

Command introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Understanding MSTP

Example: Configuring Network Regions for VLANs with MSTP on Switches

show spanning-tree bridge | **26548**

show spanning-tree statistics | **26572**

mstp | **11006**

show spanning-tree statistics

IN THIS SECTION

- Syntax | 26572
- Syntax (EX Series and QFX Series) | 26572
- Syntax (EX Series) | 26572
- Description | 26573
- Options | 26573
- Required Privilege Level | 26573
- Output Fields | 26573
- Sample Output | 26574
- Release Information | 26574

Syntax

```
show spanning-tree statistics
<brief | detail>
<interface interface-name>
<routing-instance routing-instance-name>
```

Syntax (EX Series and QFX Series)

```
show spanning-tree statistics
<brief | detail>
<interface interface-name | vlan vlan-id>
```

Syntax (EX Series)

```
show spanning-tree statistics
interface interface-name
```

```
vlan vlan-id
<brief | detail>
```

Description

Display STP statistics.

Options

none	Display brief STP statistics.
brief | detail	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display STP statistics for the specified interface.
routing-instance *routing-instance-name*	(Optional) Display STP statistics for the specified routing instance.

Required Privilege Level

view

Output Fields

Table 1588 on page 26573 lists the output fields for the `show spanning-tree statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1588: show spanning-tree statistics Output Fields

Field Name	Field Description
Message type	Type of message being counted.
BPDUs sent	Total number of BPDUs sent.
BPDUs received	Total number of BPDUs received.
BPDUs sent in last interval	Number of BPDUs sent within a specified interval.

Table 1588: show spanning-tree statistics Output Fields *(Continued)*

Field Name	Field Description
BPDUs received in last interval	Number of BPDUs received within a specified interval.
Interface	Interface for which the statistics are being displayed.
Next BPDU transmission	Number of seconds until the next BPDU is scheduled to be sent.

Sample Output

show spanning-tree statistics routing-instance

```
user@host> show spanning-tree statistics routing-instance vs1 detail
Routing instance level STP statistics
Message type                   : bpdus
BPDUs sent                     : 1396
BPDUs received                 : 1027
BPDUs sent in last interval    : 5        (duration: 4 sec)
BPDUs received in last interval: 4        (duration: 4 sec)
```

show spanning-tree statistics interface routing-instance detail

```
user@host> show spanning-tree statistics interface ge-11/1/4 routing-instance vs1 detail
Interface   BPDUs sent   BPDUs received      Next BPDU
                                            transmission
ge-11/1/4            7           190                0
```

Release Information

Command introduced in Junos OS Release 8.4.

show spanning-tree statistics bridge

IN THIS SECTION

- Syntax | 26575
- Description | 26575
- Required Privilege Level | 26575
- Output Fields | 26575
- Sample Output | 26576
- Release Information | 26576

Syntax

```
show spanning-tree statistics bridge
```

Description

Display the STP statistics of the bridge.

Required Privilege Level

view

Output Fields

Table 1589 on page 26575 describes the output fields for the **show spanning-tree statistics bridge** command. Output fields are listed in the approximate order in which they appear.

Table 1589: show spanning-tree statistics bridge Output Fields

Field Name	Field Description
`STP Context`	Context of STP instances saved for each routing instance. All STP instances in the same routing instances have same context.

Table 1589: show spanning-tree statistics bridge Output Fields *(Continued)*

Field Name	Field Description
STP Instance	Instance number that uniquely identifies each STP session per routing instance.
Number of Root Bridge Changes	Counts the number of Root Bridge change events.
Number of Root Port Changes	Counts the number of Root Port change events.
Recent TC Received	Details about the last topology change received.

Sample Output

show spanning-tree statistics bridge (MX Series)

```
user@host> show spanning-tree statistics bridge
STP Context  : default
STP Instance : 0
Number of Root Bridge Changes: 1            Last Changed: Wed Oct 23 07:10:05 2013
Number of Root Port Changes:   2            Last Changed: Wed Oct 23 07:10:05 2013
Recent TC  Received:  ge-3/1/4.32767        Received    : Wed Oct 23 07:10:07 2013
```

Release Information

Command introduced in Junos OS Release 14.1.

show spanning-tree statistics interface

IN THIS SECTION

- Syntax | 26577
- Description | 26577

- Required Privilege Level | 26577
- Output Fields | 26577
- Sample Output | 26578
- Release Information | 26578

Syntax

```
show spanning-tree statistics interface
```

Description

Display the STP statistics related to the interface.

Required Privilege Level

view

Output Fields

Table 1590 on page 26577 describes the output fields for the **show spanning-tree statistics interface** command. Output fields are listed in the approximate order in which they appear.

Table 1590: show spanning-tree statistics interface Output Fields

Field Name	Field Description
`Interface`	Interface name.
`BPDUs sent`	Total number of BPDUs sent from the bridge on the interface.
`BPDUs received`	Total number of BPDUs received by the bridge on the interface.
`Next BPDU Transmission`	Time after which the next BPDU is sent by the bridge on the interface.

Table 1590: show spanning-tree statistics interface Output Fields *(Continued)*

Field Name	Field Description
TC Tx/Rx	Total number of Topology Change BPDUs sent or received on the interface.
Proposal Tx/Rx	Total number of Proposal BPDUs sent or received on the interface.
Agreement Tx/Rx	Total number of Agreement BPDUs sent or received on the interface.

Sample Output

show spanning-tree statistics interface (MX Series)

```
user@host> show spanning-tree statistics interface

Interface BPDUs  BPDUs     Next BPDU     TCs    Proposal  Agreement
          Sent   Received  Transmission  Tx/Rx   Tx/Rx      Tx/Rx
xe-0/0/0   49       3            1        5/2      0/2      1/0

ge-1/0/0   48       1            1        5/1      0/1      1/1
```

Release Information

Command introduced in Junos OS Release 14.1.

show spanning-tree statistics message-queues

IN THIS SECTION

- Syntax | 26579
- Description | 26579

- Required Privilege Level | 26579
- Output Fields | 26579
- Sample Output | 26580
- Release Information | 26580

Syntax

```
show spanning-tree statistics message-queues
```

Description

Display the STP message queues-related statistics.

Required Privilege Level

view

Output Fields

Table 1591 on page 26579 describes the output fields for the **show spanning-tree statistics message-queues** command. Output fields are listed in the approximate order in which they appear.

Table 1591: show spanning-tree statistics message-queues Output Fields

Field Name	Field Description
`Queue`	PPMD name.
`Current size`	Number of packets currently present in the queue.
`High-watermark`	Maximum number of packets present in the queue at any time.
`max/avg wait time`	Maximum or average time packet waiting to be consumed.

Sample Output

show spanning-tree statistics message-queues (MX Series)

```
user@host> show spanning-tree statistics message-queues

Queue          Current size         High-watermark      max/avg wait time
PPMD-TX          15                       142             17636884/17636884
PPMD-RX          18                       83              18866272/18866272
```

Release Information

Command introduced in Junos OS Release 14.1.

show spanning-tree stp-buffer see-all

IN THIS SECTION

- Syntax | 26580
- Description | 26581
- Options | 26581
- Required Privilege Level | 26581
- Output Fields | 26581
- Sample Output | 26583
- Release Information | 26587

Syntax

```
show spanning-tree stp-buffer see-all
<stp-instance stp-instance-id routing-instance instance-name>
<vlan vlan-id routing-instance instance-name>
```

Description

Display the configured STP (RSTP, MSTP, VSTP) interface parameters.

Options

none	Display STP (RSTP, MSTP, VSTP) interface role/state changes that are logged into internal memory called the buffer. Entries in the buffer depend on the user configuration.
`stp-instance` *stp-instance-id* `routing-instance` *instance-name*	(Optional) Display the STP buffer for the specified stp-instance and routing instance.
`vlan` *vlan-id* `routing-instance` *instance-name*	(Optional) Display the STP buffer for the specified vlan id and routing instance.

Required Privilege Level

view

Output Fields

Table 1592 on page 26581 describes the output fields for the **show spanning-tree stp-buffer see-all** command. Output fields are listed in the approximate order in which they appear.

Table 1592: show spanning-tree stp-buffer see-all Output Fields

Field Name	Field Description	Level of Output
`Global Events`	Displays events when PPMD RX/TX queues reach 70 % of their maximum queue size. The following indicates received queue status: • GT — Greater than 70 % of their maximum queue size. • LT — Less than 70 % of their maximum queue size.	none, `stp-instance`, `vlan`

Table 1592: show spanning-tree stp-buffer see-all Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Per STP instance Information`	Information about every STP instance.	none, `stp-instance`, `vlan`
`Routing Inst`	Routing instance to which the STP instance belongs.	none, `stp-instance`, `vlan`
`STP Instance`	Instance number that uniquely identifies each STP session per routing-instance.	none, `stp-instance`
`Root Bridge`	Bridge priority and bridge ID of ROOT bridge in the topology.	none, `vlan`
`Root Port`	Information about ROOT port, if any, on the local bridge at the displayed timestamp.	none, `vlan`
`TC Received`	Time at which topology change was received and on which port.	none, `stp-instance`, `vlan`
`TC Generated`	Time at which topology change occurred and on which port.	none, `stp-instance`, `vlan`
`Port`	The interface where the event is occurring.	none, `stp-instance`, `vlan`

Table 1592: show spanning-tree stp-buffer see-all Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`State`	STP state of the port. The following are the types of state: • Forwarding — Port forwards the traffic and is included in active topology and learns MAC addresses. • Blocking — Port does not forward traffic and is not included in active topology. Does not learn MAC addresses.	none, `vlan`
`Role`	Role of the port. The following are the types of ports: • Root — Port closest to the Root bridge. • Designated — Port sends the best BPDU on the connected segment. • Alternate — Port blocked for receiving more useful BPDUs from another bridge. • Disabled — Port is disabled. Cannot send or receive BPDUs.	none, `stp-instance`, `vlan`

Sample Output

show spanning-tree stp-buffer see-all (MX Series)

```
user@host> show spanning-tree stp-buffer see-all
1. Global Events:
Time at which different Queue’s reached 70% of the Maximum Q-size
Mar 18 13:18:04 RCV_Q GT
Mar 18 13:18:57 RCV_Q LT
```

```
Mar 18 13:19:33 XMIT_Q GT

2. Per STP Instance Information :
Routing Inst    :  default
STP Instance  : 02
Root Bridge     : 12288.00:23:9c:f0:17:d0 Mar 18 13:18:04
STP Instance  : 03
Root Port        : ge-0/1/2 Mar 18 13:18:38
STP Instance  : 03
TC  Received  : ge-0/0/2 Mar 18 15:12:12
STP Instance  : 03
TC  Generated: ge-0/2/2 Mar 18 15:13:27

3. This section will print the time stamp for per Instance - port event changes.
STP Instance :  02
Port         :  ge-1/0/0 Mar 22 13:35:02
State       :  FWD  S
Role        :  DESG

STP Instance :  00
Port         :  ge-0/0/3 Mar 22 14:03:46
State       :  BLK
Role        :  ALT  R
STP Instance :  00
Port         :  ge-1/0/0 Mar 28 02:03:49
State       :  BLK  S
Role        :  ALT
```

show spanning-tree stp-buffer see-all stp-instance stp-instance-id routing-instance instance-name (MX Series)

```
user@host> show spanning-tree stp-buffer see-all stp-instance 0 routing-instance mstp_inst

1. Global Events:
Time at which different Queue's reached 70% of the Maximum Q-size
No Entry So far

2. Per STP Instance Information :

Routing Inst :
```

```
mstp_inst



STP Instance : 0
TC Generated : ge-3/0/5.32767 Tue Dec 17 06:00:50 2013
STP Instance : 0
TC Generated : ge-3/1/1.32767 Tue Dec 17 06:00:50 2013
STP Instance : 0
TC  Received : ge-3/0/5.32767 Tue Dec 17 06:00:50 2013
STP Instance : 0
TC  Received : ge-3/0/5.32767 Tue Dec 17 06:00:50 2013
STP Instance : 0
TC  Received : ge-3/0/5.32767 Tue Dec 17 06:00:51 2013
STP Instance : 0
TC  Received : ge-3/0/5.32767 Tue Dec 17 06:00:53 2013

3.  This section will print the time stamp for per Instance, Port Event changes.



STP Instance : 0
Port         : ge-3/0/5.32767 Tue Dec 17 06:00:49 2013
State        : BLK S
Role         : DIS

STP Instance : 0
Port         : ge-3/0/5.32767 Tue Dec 17 06:00:49 2013
State        : BLK
Role         : DESG R

STP Instance : 0
Port         : ge-3/1/1.32767 Tue Dec 17 06:00:49 2013
State        : BLK S
Role         : DIS

STP Instance : 0
Port         : ge-3/1/1.32767 Tue Dec 17 06:00:49 2013
State        : BLK
Role         : DESG R

STP Instance : 0
Port         : ge-3/0/5.32767 Tue Dec 17 06:00:50 2013
State        : FWD S
Role         : DESG
```

```
STP Instance : 0
Port         : ge-3/1/1.32767 Tue Dec 17 06:00:50 2013
State        : FWD S
Role         : DESG
```

show spanning-tree stp buffer see-all vlan vlan-id routing-instance instance-name (MX Series)

```
user@host> show spanning-tree stp-buffer see-all vlan 10 routing-instance vstp_inst
1. Global Events:
Time at which different Queue’s reached 70% of the Maximum Q-size
Mar 18 13:18:04 RCV_Q GT
Mar 18 13:18:57 RCV_Q LT
Mar 18 13:19:33 XMIT_Q GT

2. Per STP Instance Information :
Routing Inst    :  default
VLAN ID  : 02
Root Bridge     : 12288.00:23:9c:f0:17:d0 Mar 18 13:18:04
VLAN ID  : 03
Root Port        : ge-0/1/2 Mar 18 13:18:38
VLAN ID  : 03
TC  Received  : ge-0/0/2 Mar 18 15:12:12
VLAN ID  : 03
TC  Generated: ge-0/2/2 Mar 18 15:13:27

3. This section will print the time stamp for per Instance - port event changes.
VLAN ID :  02
Port         :  ge-1/0/0 Mar 22 13:35:02
State       :  FWD  S
Role        :  DESG

VLAN ID :  00
Port         :  ge-0/0/3 Mar 22 14:03:46
State       :  BLK
Role        :  ALT  R
VLAN ID :  00
Port         :  ge-1/0/0 Mar 28 02:03:49
State       :  BLK  S
Role        :  ALT
```

Release Information

Command introduced in Junos OS Release 14.1.

show spring-traffic-engineering

IN THIS SECTION

- Syntax | 26587
- Description | 26587
- Options | 26587
- Required Privilege Level | 26588
- Output Fields | 26588
- Sample Output | 26591
- Release Information | 26599

Syntax

```
show spring-traffic-engineering (lsp | overview | route | sbfd)
<brief | detail>
<color>
<destination>
<logical-system (all | logical-system-name)>
<name lsp-name>
```

Description

Display ingress details of SPRING traffic engineering.

Options

brief | detail (Optional) Display the specific level of output.

lsp Display details of SPRING traffic-engineered LSPs on the ingress router or the Path Computation Client (PCC).

route Display the spring traffic engineering route information.

color Display only SPRING traffic-engineered LSPs for a specific color. This filtering option is available only for LSP, not available for overview and sbfd.

destination Display only SPRING traffic-engineered LSPs to a specific destination. Destination can be IPv4/IPv6 to-address. This filtering option is available only for LSP, not available for overview and sbfd.

overview Display overview of SPRING traffic-engineered LSPs on the ingress router, or the PCC.

sbfd Display the SPRING traffic-engineered BFD session.

name *lsp-name* (Optional) Regular expression for LSP names to match for displaying SPRING traffic-engineering details.

Required Privilege Level

view

Output Fields

Table 1593 on page 26588 describes the output fields for the `show spring-traffic-engineering` command. Output fields are listed in the approximate order in which they appear.

Table 1593: show spring-traffic-engineering Output Fields

Field Name	Field Description
To	IP address of the SR-TE LSP destination.
State	State of the SR-TE LSP: • Up • Down
LSP Name	Name of the SR-TE LSP.

Table 1593: show spring-traffic-engineering Output Fields *(Continued)*

Field Name	Field Description
S-ERO	Source Explicit Route Object (ERO), or LSP path.
Bandwidth	Bandwidth allocated for the SR-TE LSP.

Delegation info	LSP control and routing status:
	• Control-status:
	  • Externally controlled—PCE has control of the source-routing path.
	  This can happen when:
	  • The `lsp-external-controller pccd` statement is configured either under the source-routing path or under the primary segment list.
	  • The `request path-computation-client retry-delegation` *`lsp-name`* command is issued for a delegated LSP that was not previously controlled by the PCE.
	  • Locally controlled—PCC has control of the source-routing path.
	  This can happen when:
	  • The PCE has returned the control of the source-routing path.
	  • Delegation timer with the PCE has expired.
	• Routing-status: Applicable to delegated source-routing paths only.
	  • Externally routed—PCE provided the ERO for the source-routing path for a delegated LSP through PCUpdate.
	  • Locally routed—PCE does not provide ERO for the source-routing path.

Route preference	Route preference of the SR-TE LSP.
Number of LSPs	Statistics of the total number of SR-TE LSPs and the LSP state.

Table 1593: show spring-traffic-engineering Output Fields *(Continued)*

Field Name	Field Description
External controllers	Name of the LSP external controller. By default the only supported external controller is `pccd`.
BFD name	Name of the BFD session. The name is auto-generated in the `V4-srte_bfd_session-`*`id`* for IPv6. The name is based on the Explicit Route Object (ERO) stack of the LSP path, that is, if multiple LSPs have same path they share the same BFD session name.
BFD status	Status of the BFD session: UP, DOWN.
BFD remote-discriminator	Remote discriminator of the reflector (auto-derived or configured) range: 1 - 4294967295
Referencing LSPs	Name of referencing LSP. If the LSP does not have a path name, then the referencing LSP is displayed as `unnamed path`.
SR-ERO hop count	Number of hops in the segment routing ERO.
Hop 1	Represents the path of the BFD session. If any other LSP is on same path, it has the same BFD session.
Total displayed BFD sessions	Total count of all the BFD sessions.
Tunnel source	Source of the tunnel configuration; for example, `static configuration`.
Ingress telemetry statistics	Ingress telemetry statistics including the sensor name and ID.
Transit telemetry statistics	Transit telemetry statistics including the sensor name and ID.

Table 1593: show spring-traffic-engineering Output Fields *(Continued)*

ERO Valid	Indicates whether the received explicit route object (ERO) is valid or not.

Sample Output

show spring-traffic-engineering lsp name

```
user@host> show spring-traffic-engineering lsp name lsp-name
To         State        LSP Name
                10.1.1.7   Up           to-R1
```

show spring-traffic-engineering lsp detail

```
user@host> show spring-traffic-engineering lsp detail
10.1.1.7
                State: Up
                S-ERO: 10.24.1.1(80001) 10.1.1.3(4509) 10.11.1.2(9875)
                Bandwidth: 100M
```

show spring-traffic-engineering lsp detail (BGP-SR-TE policies based tunnel)

```
user@host> show spring-traffic-engineering lsp detail
Name: tunnel15
Tunnel-source: Static configuration
To: 10.6.6.6
State: Up
  Path: sl-15-primary
  Outgoing interface: NA
  Auto-translate status: Disabled Auto-translate result: N/A
  Compute Status:Disabled , Compute Result:N/A , Compute-Profile Name:N/A
  BFD status: N/A BFD name: N/A
  ERO Valid: true
  SR-ERO hop count: 3
    Hop 1 (Strict):
      NAI: IPv4 Adjacency ID, 0.0.0.0 -> 10.21.4.2
      SID type: None
    Hop 2 (Strict):
```

```
      NAI: None
      SID type: 20-bit label, Value: 400050
    Hop 3 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 400060
  Path: sl-15-backup
  Outgoing interface: ge-0/0/2.0
  Auto-translate status: Disabled Auto-translate result: N/A
  Compute Status:Disabled , Compute Result:N/A , Compute-Profile Name:N/A
  BFD status: N/A BFD name: N/A
  ERO Valid: true
  SR-ERO hop count: 3
    Hop 1 (Strict):
      NAI: IPv4 Adjacency ID, 0.0.0.0 -> 10.21.3.2
      SID type: None
    Hop 2 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 400050
    Hop 3 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 400060
  Path: sl-15-backup
  Outgoing interface: ge-0/0/2.0
  Auto-translate status: Disabled Auto-translate result: N/A
  Compute Status:Disabled , Compute Result:N/A , Compute-Profile Name:N/A
  BFD status: N/A BFD name: N/A
  ERO Valid: true
  SR-ERO hop count: 3
    Hop 1 (Strict):
      NAI: IPv4 Adjacency ID, 0.0.0.0 -> 10.21.3.2
      SID type: None
    Hop 2 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 400050
    Hop 3 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 400060

Name: bgp-srte-0.0.0.0-900-10.7.7.7-111
Tunnel-source: BGP SRTE
To: 10.7.7.7-111 State: Up <c>
Outgoing interface: NA
  Auto-translate status: Disabled Auto-translate result: N/A
```

```
  BFD status: N/A BFD name: N/A
  ERO Valid: true
  SR-ERO hop count: 1
    Hop 1 (Strict):
      NAI: None
      SID type: 32-bit label, Value: 400050, TTL: 10, Exp: 1


Name: bgp-srte-0.0.0.0-900-10.5.5.5-111
Tunnel-source: BGP SRTE
To: 10.5.5.5-111<c>
State: Up
  Outgoing interface: NA
  Auto-translate status: Disabled
  Auto-translate result: N/A
  BFD status: N/A BFD name: N/A
  ERO Valid: true
  SR-ERO hop count: 1
    Hop 1 (Strict):
      NAI: None
      SID type: 32-bit label, Value: 400050, TTL: 10, Exp: 1

Name: bgp-srte-0.0.0.0-900-10.6.6.6-111
Tunnel-source: BGP SRT E
To: 10.6.6.6-111<c>
State: Up
     Outgoing interface: NA
  Auto-translate status: Disabled
  Auto-translate result: N/A
  BFD status: N/A BFD name: N/A
  ERO Valid: true
  SR-ERO hop count: 1
    Hop 1 (Strict):
      NAI: None
      SID type: 32-bit label, Value: 400050, TTL: 10, Exp: 1

Total displayed LSPs: 4 (Up: 4, Down: 0)
```

show spring-traffic-engineering lsp detail (BGP-SR-TE policies based colored tunnels)

For colored tunnels, telemetry details are displayed.

```
user@host> show spring-traffic-engineering lsp detail

Name: bgp-srte-0.0.0.0-900-10.6.6.6-111
Tunnel-source: BGP SRTE
To: 10.6.6.6-111
State: Up Telemetry statistics:
Sensor-name: f4248-10.6.6.6-6f, Id: 3758096396
Sensor-name: 10.6.6.6-6f, Id: 3758096395
  Outgoing interface: NA
  Auto-translate status: Disabled Auto-translate result: N/A
  BFD status: N/A BFD name: N/A
  ERO Valid: true
  SR-ERO hop count: 1
  Hop 1 (Strict):
    NAI: None
    SID type: 32-bit label, Value: 400050, TTL: 10, Exp: 1
```

show spring-traffic-engineering lsp detail (BGP-SRTE filter in tunnel source)

```
user@host> show spring-traffic-engineering lsp detail
```

show spring-traffic-engineering lsp detail (PCE-Delegated LSPs)

```
user@host> show spring-traffic-engineering lsp detail
srte_at_dlg_to_r5
Oct 16 14:39:11
Name: srte_at_dlg_to_r5
Tunnel-source: Static configuration
To: 10.128.14.141
State: Up
  Path: sr_auto_to_r5
  Outgoing interface: NA
  Delegation info:
   Control-status: Externally controlled
   Routing-status: Externally routed
```

```
  Auto-translate status: Disabled Auto-translate result: N/A
  BFD status: N/A BFD name: N/A
```

show spring-traffic-engineering lsp detail (with Auto Derivation)

```
user@host> show spring-traffic-engineering lsp detail
Name: lsp1
  Tunnel-source: Static configuration
  Tunnel Forward Type: SRMPLS
  To: 10.1.1.1
  Te-group-id: 0
  State: Up
  Telemetry statistics:
  Sensor-name: ingress-lsp1, Id: 3758096386
    Path: sl1
    Path Status: NA
    Outgoing interface: NA
    Auto-translate status: Disabled Auto-translate result: N/A
    Compute Status:Enabled , Compute Result:success , Compute-Profile Name:cprof1
    Total number of computed paths: 2
    Segment ID : 128
    Computed-path-index: 1
      BFD status: Up BFD name: V4-srte_bfd_session-2
      BFD remote-discriminator: 184549479 (auto-derived)
      TE metric: 105, IGP metric: 105
      ...

Name: lsp2
  Tunnel-source: Static configuration
  Tunnel Forward Type: SRMPLS
  To: 10.1.1.1
  Te-group-id: 0
  State: Up
  Telemetry statistics:
  Sensor-name: ingress-lsp2, Id: 3758096389
    Path: sl2
    Path Status: NA
    Outgoing interface: NA
    Auto-translate status: Disabled Auto-translate result: N/A
    Compute Status:Enabled , Compute Result:success , Compute-Profile Name:cprof1
    Total number of computed paths: 2
    Segment ID : 128
```

```
    Computed-path-index: 1
      BFD status: Up BFD name: V4-srte_bfd_session-4
      BFD remote-discriminator: 501 (configured)
      TE metric: 105, IGP metric: 105
       ...

Name: lsp3
  Tunnel-source: Static configuration
  Tunnel Forward Type: SRMPLS
  To: 10.1.1.1
  Te-group-id: 0
  State: Up
  Telemetry statistics:
  Sensor-name: ingress-lsp3, Id: 3758096386
    Path: sl2
    Path Status: NA
    Outgoing interface: NA
    Auto-translate status: Disabled Auto-translate result: N/A
    Compute Status:Enabled , Compute Result:success , Compute-Profile Name:cprof1
    Total number of computed paths: 2
    Segment ID : 128
    Computed-path-index: 1
      BFD status: N/A BFD name: N/A
      BFD remote-discriminator: N/A
      TE metric: 105, IGP metric: 105
       ...
```

show spring-traffic-engineering overview

```
user@host> show spring-traffic-engineering overview
Overview of SPRING-TE:
  Route preference: 8
  Number of LSPs: 0 (Up: 0, Down: 0)
  External controllers:
    pccd
```

show spring-traffic-engineering sbfd detail

```
user@host> show spring-traffic-engineering sbfd detail
BFD name: V4-srte_bfd_session-1
```

```
BFD status: Down
Referencing LSPs:
  sr-lsp1:path1
  sr-lsp2:path1
SR-ERO hop count: 2
  Hop 1 (Strict):
    NAI: IPv4 Adjacency ID, 10.1.1.1 -> 10.1.1.2
    SID type: 20-bit label, Value: 299776
  Hop 2 (Strict):
    NAI: IPv4 Adjacency ID, 10.2.0.1 -> 10.2.0.2
      SID type: 20-bit label, Value: 299824


Total displayed BFD sessions: 2 (Up: 2, Down: 0)
```

show spring-traffic-engineering sbfd detail (with Auto-derivation)

```
user@host> show spring-traffic-engineering sbfd detail
BFD name: V4-srte_bfd_session-2
BFD status: Up
BFD remote-discriminator: 184549479
Referencing LSPs:
  delegatd_lsp:sl1
  lsp2:sl2
      SR-ERO hop count: 4
        …

BFD name: V4-srte_bfd_session-4
BFD status: Up
BFD remote-discriminator: 501
Referencing LSPs:
  delegatd_lsp:sl1
  lsp2:sl2
      SR-ERO hop count: 4
        …
```

show spring-traffic-engineering lsp detail name <name>

```
user@host> show spring-traffic-engineering lsp detail name sr_plcy1
Name: sr_plcy1
```

```
Tunnel source: Static configuration
To: 10.1.1.1
State: Up
  Path: sl1
Ingress tele1etry statistics:
      Sensor-Name: i;st;0;f;sr_plcy1;sl1, Id: 3758096390
Transit tele1etry statistics:
      Sensor-Name: t;st;0;f;sr_plcy1;sl1, Id: 3758096391
  Path: sl2
Ingress tele1etry statistics:
      Sensor-Name: i;st;0;f;sr_plcy1;sl2, Id: 3758096390
Transit tele1etry statistics:
      Sensor-Name: t;st;0;f;sr_plcy1;sl2, Id: 3758096391
```

show spring-traffic-engineering lsp name <name> detail (Static route over SR-TE tunnel)

The following output is generated when `tunnel-tracking` statement is configured at the `[edit protocols source-packet-routing]` hierarchy level.

```
user@host> show spring-traffic-engineering lsp name static-lsp1 detail
Name: static-lsp1
Tunnel-source: Static configuration
To: 10.7.7.7
State: Up
Path: Path1
Path Status: Up
Outgoing interface: NA
Auto-translate status: Disabled Auto-translate result: N/A
Compute Status:Disabled , Compute Result:N/A , Compute-Profile Name:N/A
BFD status: N/A BFD name: N/A
ERO Valid: true
SR-ERO hop count: 2
Hop 1 (Strict):
NAI: None
SID type: 20-bit label, Value: 801003
Hop 2 (Strict):
NAI: None
SID type: 20-bit label, Value: 801007
Path: Path2
Path Status: Up
```

```
Outgoing interface: NA
Auto-translate status: Disabled Auto-translate result: N/A
Compute Status:Disabled , Compute Result:N/A , Compute-Profile Name:N/A
BFD status: N/A BFD name: N/A
ERO Valid: true
SR-ERO hop count: 2
Hop 1 (Strict):
NAI: None
SID type: 20-bit label, Value: 801005
Hop 2 (Strict):
NAI: None
SID type: 20-bit label, Value: 801007

Total displayed LSPs: 1 (Up: 1, Down: 0)
```

NOTE: The `Path Status` shows NA when `tunnel-tracking` statement is not configured.

show spring-traffic-engineering lsp color <value>

```
user@host> show spring-traffic-engineering lsp color 111
To              State     LSPname
10.6.6.6-111<c>  Up        bgp-srte-0.0.0.0-900-10.6.6.6-111
10.7.7.7-111<c>  Up        bgp-srte-0.0.0.0-900-10.7.7.7-111
10.5.5.5-111<c>  Up        bgp-srte-0.0.0.0-900-10.5.5.5-111
```

show spring-traffic-engineering lsp destination <address>

```
user@host> show spring-traffic-engineering lsp destination 10.6.6.6
To              State     LSPname
10.6.6.6-111<c>  Up        bgp-srte-0.0.0.0-900-10.6.6.6-111
10.6.6.6         Up        pcep-tunnel1
```

Release Information

Command introduced in Junos OS Release 17.2.

`sbfd` option introduced in Junos OS Release 19.4R1 on all platforms.

`color` and `destination` options introduced in Junos OS 21.2R1.

RELATED DOCUMENTATION

Enable Segment Routing for the Path Computation Element Protocol

show-spring-traffic-statistics

IN THIS SECTION

- show spring traffic-statistics | 26600
- Description | 26600
- Options | 26601
- Additional Information | 26601
- Required Privilege Level | 26601
- Output Fields | 26601
- Sample Output | 26602
- Release Information | 26602

show spring traffic-statistics

```
show spring traffic-statistics
```

Description

Displays the spring sensor-based traffic statistics. The per-interface-per-member-link <ingress> <egress> and per-sid <ingress> <egress> filters are used to view the details of Packets, Bytes, Packet-Rate, and Byte-Rate.

```
show spring traffic-statistics per-interface-per-member-link ingress
show spring traffic-statistics per-interface-per-member-link egress
```

```
show spring traffic-statistics per-sid ingress
show spring traffic-statistics per-sid egress
```

Options

per-interface-per-member-link ingress

egress

per-sid ingress

egress

Additional Information

For Junos OS Evolved, the native subscription path is "../junos/services/segment-routing/sid/egress/usage/", because currently, the OC compliant path for per-sid egress sensor is not supported on EVO platforms.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show spring traffic-statistics` command.

Table 1594: show spring traffic-statistics Output Fields

Field Name	Field Description	Level of Output
Interface/Sid	Displays interface name and segment id.	-
Type	Displays interface type	Ingress/Egress
PacketsBytes	Displays number of packets in bytes.	-
Packet-Rate	Displays number of packets transmitted per second.	-
Byte-Rate	Displays number of bytes transmitted per second.	-

Sample Output

show spring traffic-statistics

```
root@R2-re0>show spring traffic-statistics
```

```
Interface/SID            Type               Packets            Bytes            Packet-
Rate            Byte-Rate
et-0/0/15.0              Intf-Egress        45                 3840
0                      0
et-0/0/3.0               Intf-Ingress       15                 1320
0                      0
500005                   SID-Ingress        30                 2640
0                      0
L-ISIS-7.7.7.7           SID-Egress         15                 1320
0                      0
Next Update expected in 8 seconds.
```

Release Information

Command introduced in Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/ref/statement/sensor-based-stats-edit-protocols-isis-source-packet-routing.html

show ssl-certificates

IN THIS SECTION

- Syntax | 26603
- Description | 26603

- Options | 26603
- Required Privilege Level | 26603
- Output Fields | 26603
- Sample Output | 26605
- Release Information | 26605

Syntax

```
show ssl certificates
```

Description

Display information about the Secure Sockets Layer (SSL) certificates installed on the switch. When you configure PEAP as the authentication protocol for MAC RADIUS authentication, you must load the server-side Secure Sockets Layer (SSL) certificate on the switch. PEAP requires an SSL certificate to create a secure TLS tunnel to protect user authentication, and uses server-side public key certificates to authenticate the server. It then creates an encrypted TLS tunnel between the client and the authentication server. The key for this encryption are transported using the server's public key. The ensuing exchange of authentication information inside the tunnel to authenticate the client is then encrypted and user credentials are safe from eavesdropping.

Options

none Display information about all SSL certificates.

`detail` Display information about SSL certificates for the specified level of output.

Required Privilege Level

view

Output Fields

Table 1595 on page 26604 lists the output fields for the `show ssl-certificates` command. Output fields are listed in the approximate order in which they appear.

Table 1595: show ssl-certificates Output Fields

Field Name	Field Description	Level of Output
`Issuer`	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • `C`—Country of origin. • `ST`—State or province name. • `L`—Locality. • `O`—Organization of origin. • `OU`—Organizational unit. • `CN`—Common name of the authority.	All levels
`Valid from`	Start time when the digital certificate becomes valid.	`detail`
`Valid from`	End time when the digital certificate becomes invalid.	`detail`
`Serial number`	Unique serial number of the digital certificate.	`detail`
`Subject`	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • `C`—Country of origin. • `ST`—State or province name. • `L`—Locality. • `O`—Organization of origin. • `OU`—Organizational unit. • `CN`—Common name of the authority.	`detail`

Sample Output

show ssl-certificates

```
user@root> show ssl-certificates
 Issuer:        /C=IN/ST=KA/L=Blr/O=JNPR/OU=CP/CN=User-Radius/emailAddress=user@juniper.net
```

show ssl-certificates detail

```
user@root> show ssl-certificates detail

    Issuer:        /C=IN/ST=KA/L=Blr/O=JNPR/OU=CP/CN=User-Radius/emailAddress=user@juniper.net
    Valid From:    May 30 17:41:04 2016 GMT
    Valid Till:    May 29 17:41:04 2026 GMT
    Serial Number: 0
    Subject:       /C=IN/ST=KA/L=Blr/O=JNPR/OU=CP/CN=User-Radius/emailAddress=user@juniper.net
```

Release Information

Command introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Configuring PEAP for MAC RADIUS Authentication

show static-subscribers sessions

IN THIS SECTION

- Syntax | 26606
- Description | 26606
- Options | 26606
- Required Privilege Level | 26606

- Output Fields | 26606
- Sample Output | 26607
- Release Information | 26609

Syntax

```
show static-subscribers sessions
<group group-name>
<interface interface-name>
```

Description

Display information about the subscriber sessions for all static subscribers, all static subscribers on an interface group, or a single subscriber on an interface.

Options

group-name	(Optional) Display session information for static subscribers on all interfaces in the specified group.
interface-name	(Optional) Display session information for the static subscriber on the specified in the specified group.

Required Privilege Level

view

Output Fields

Table 1596 on page 26607 lists the output fields for the `show static-subscribers sessions` command. Output fields are listed in the approximate order in which they appear.

Table 1596: show static-subscribers sessions Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the interface.	None specified
`State`	State of the static subscriber session: • `authenticating`—Subscriber is being authenticated. • `activating client`—Client is being activated. • `activating services`—Subscriber services are being activated. • `deactivating client`—Client is being deactivated. • `deactivating services`—Subscriber services are being deactivated. • `initializing`—Process is initializing. • `logged in`—Subscriber is logged in to the interface. • `logged out`—Subscriber is logged out of the interface. • `processing statistics`—Session statistics are being processed. • `terminating session`—Subscriber session is being terminated.	None specified
`Group`	Name of the interface group to which the interface belongs.	None specified
`User Name`	Username used for the static subscriber. Can be the interface name.	None specified

Sample Output

show static-subscribers sessions

```
user@host> show static-subscribers sessions

Static subscriber information:
Interface         State                 Group            User Name
ge-9/1/0.1        logged out            SS1              ge-9-1-0.1
```

```
ge-9/1/0.10       logged out            SS1              ge-9-1-0.10
ge-9/1/0.100      logged out            SS1              ge-9-1-0.100
ge-9/1/0.11       logged out            SS1              ge-9-1-0.11
ge-9/1/0.12       logged out            SS1              ge-9-1-0.12
ge-9/1/0.13       logged out            SS1              ge-9-1-0.13
ge-9/1/0.14       logged out            SS1              ge-9-1-0.14
ge-9/1/0.15       logged out            SS1              ge-9-1-0.15
ge-9/1/0.16       logged out            SS1              ge-9-1-0.16
ge-9/1/0.17       logged out            SS1              ge-9-1-0.17
ge-9/1/0.18       logged out            SS1              ge-9-1-0.18
ge-9/1/0.19       logged out            SS1              ge-9-1-0.19
ge-9/1/0.2        logged out            SS1              ge-9-1-0.2
ge-9/1/0.20       logged out            SS1              ge-9-1-0.20
ge-9/1/0.21       logged out            SS1              ge-9-1-0.21
```

show static-subscribers sessions group

```
user@host> show static-subscribers sessions group boston

Interface         State                 Group            User Name
ge-0/0/1.1        logged in             boston           ge-0/0/1.1
ge-0/0/1.2        logged in             boston           ge-0/0/1.2
```

show static-subscribers sessions interface

```
user@host> show static-subscribers sessions interface ge-0/0/1.1

Interface         State                 Group            User Name
ge-0/0/1.1        logged in             foo              ge-0/0/1.1
```

show static-subscribers sessions interface (Pseudowire Interface)

```
user@host> show static-subscribers sessions interface ps100.4040
Static subscriber information:
```

```
Interface          State                  Group            User Name
ps100.4040         logged in              foo              static-user:4000@jnpr.net
```

show static-subscribers sessions group (Pseudowire Interface)

```
user@host> show static-subscribers sessions group grp-1
Static subscriber information:
Interface          State                  Group            User Name
ps100.4040         logged in              foo              static-user:4000@jnpr.net
ps100.4041         logged in              foo              static-user:4001@jnpr.net
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

JSRC and Subscribers on Static Interfaces

show subscribers

IN THIS SECTION

- Syntax | 26610
- Description | 26610
- Options | 26610
- Required Privilege Level | 26614
- Output Fields | 26614
- Sample Output | 26627
- Sample Output for AGF | 26661
- Release Information | 26662

Syntax

```
show <user-plane> subscribers
<detail | extensive | terse>
<aci-interface-set-name aci-interface-set-name>
<address address>
<agent-circuit-identifier agent-circuit-identifier>
<agent-remote-identifier agent-remote-identifier>
<aggregation-interface-set-name interface-set-name>
<client-type client-type>
<count>
<id session-id <accounting-statistics>>
<interface interface <accounting-statistics>>
<logical-system logical-system>
<mac-address mac-address>
<physical-interface physical-interface-name>
<profile-name profile-name>
<routing-instance routing-instance>
<stacked-vlan-id stacked-vlan-id>
<sub-system (agf | bng)>
<subscriber-state subscriber-state>
<user-name user-name>
<vci vci-identifier>
<vpi vpi-identifier>
<vlan-id vlan-id>
```

Description

Display information for active subscribers on both the broadband network gateway (BNG) and Access Gateway Function (AGF) subsystems. When you have subscribers logged in to both the BNG and the AGF, you can use the `sub-system` option to view the outputs specific to either subsytem.

See "Sample Output for AGF" on page 26661 for subscriber session information about the AGF subsystem.

Options

detail | extensive | terse	(Optional) Display the specified level of output.

aci-interface-set-name

(Optional) Display all the dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. You must use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) control packets.

address

(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.0.2.0). If you specify the IP address as a prefix with a netmask (for example, 192.0.2.0/32), the router displays a message that the IP address is invalid and rejects the command.

agent-circuit-identifier

(Optional) Display all the dynamic subscriber sessions whose ACI value matches the specified string. You can specify either the complete ACI string or a substring. To specify a substring, you must enter characters that form the beginning of the string, followed by an asterisk (*) as a wildcard to substitute for the remainder of the string. The wildcard can be used only at the end of the specified substring; for example:

```
user@host1> show subscribers agent-circuit-identifier substring*
```

The table below lists supported string values for different Junos OS releases.

Table 1597: Supported Substring

Junos OS Release	Substring Support
Junos OS Release 13.3R1	You can specify a substring without a wildcard.
Starting in Junos OS Release 14.1R1	You must specify the complete ACI string; you cannot specify a wildcard.
Starting in Junos OS Releases 15.1R7, 16.1R7, 16.2R3, 17.1R3, 17.2R3, 17.3R3, 17.4R2, 18.1R2, 18.2R1	You can specify a substring, but you must include the wildcard character at the end of the substring.

agent-remote-identifier

(Optional) Display all the dynamic subscriber sessions whose agent remote identifier (ARI) value matches the specified string. You must specify the complete ACI string; you cannot specify a wildcard.

aggregation-interface-set-name *interface-set-name*

(Optional) Display summary information for the specified aggregation node interface set, including interface, VLAN ID, username, and logical system and routing instance (LS:RI).

client-type

(Optional) Display subscribers whose client type matches one of the following client types:

- `dhcp`—Dynamic Host Configuration Protocol (DHCP) clients only.
- `dotlx`—802.1X clients only.
- `essm`—Extensible Subscribers Services Manager (ESSM) clients only.
- `fixed-wireless-access`—Fixed wireless access clients only.
- `fwauth`—FwAuth (authenticated across a firewall) clients only.
- `l2tp`—Layer 2 Tunneling Protocol (L2TP) clients only.
- `mlppp`—Multilink Point-to-Point Protocol (MLPPP) clients only.
- `ppp`—Point-to-Point Protocol (PPP) clients only.
- `pppoe`—Point-to-Point Protocol over Ethernet (PPPoE) clients only.
- `static`—Static clients only.
- `vlan`—VLAN clients only.
- `vlan-oob`—VLAN out-of-band (triggered by by Access Node Control Protocol or ANCP) clients only.
- `vpls-pw`—Virtual private LAN service (VPLS) pseudowire clients only.
- `xauth`—Extended Authentication (XAuth) clients only.

count

(Optional) Display the count of the total subscribers and active subscribers for any specified option. You can use the `count` option alone or in combination with one or more of the following options:

- `address`
- `client-type`
- `interface`
- `logical-system`

- `mac-address`
- `profile-name`
- `routing-instance`
- `stacked-vlan-id`
- `subscriber-state`
- `vlan-id`

id *session-id*
(Optional) Display a specific subscriber session whose session ID matches the specified subscriber ID. You can display the subscriber IDs by using the `show subscribers extensive` command or the `show subscribers interface extensive` command.

id *session-id* accounting-statistics
(Optional) Display accurate subscriber accounting statistics for a subscriber session based on the session ID you specify. You must configure the `actual-transmit-statistics` statement in the dynamic profile for the dynamic logical interface. If you do not configure statement, the CLI displays a value of 0 for the accounting statistics.

interface
(Optional) Display subscribers whose interface matches the specified interface.

***interface* accounting-statistics**
(Optional) Display subscriber accounting statistics for the specified interface. If you do not configure statement, the CLI displays a value of 0 for the accounting statistics.

logical-system
(Optional) Display subscribers whose logical system matches the specified logical system.

mac-address
(Optional) Display subscribers whose MAC address matches the specified MAC address.

physical-interface-name
(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

profile-name
(Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance
(Optional) Display subscribers whose routing instance matches the specified routing instance.

stacked-vlan-id
(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

subsystem (agf | bng)
(Optional) Display information for subscribers on the AGF or BNG subsystem.

> **NOTE**: The subsystem option is only available when both the AGF and BNG subscribers are logged into the router at the same time.

subscriber-state (Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

user-name (M120, M320, and MX Series routers only) (Optional) Display subscriber whose username matches the specified subscriber name.

vci-identifier (MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active Asynchronous Transfer Mode (ATM) subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is 0 through 255.

vpi-identifier (MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is 0 through 65,535.

vlan-id (Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. Among the subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To see the subscribers (among subscribers using a double-tagged VLAN) whose outer VLAN tag matches the specified value, you can use the `stacked-vlan-id` *`stacked-vlan-id`* option to match the outer VLAN tag.

> **NOTE**: Because of display limitation, the logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

Output Fields

Table 1598 on page 26615 lists the output fields for the `show subscribers` command. Output fields are listed in the approximate order in which they appear.

Table 1598: show subscribers Output Fields

Field Name	Field Description
`Interface`	Interface associated with the subscriber. The router or the switch displays the subscribers whose interface matches or begins with the specified interface. The asterisk (*) indicates a continuation of addresses for the same session.
`IP Address/VLAN ID`	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id* No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions, the value is `Tunnel-switched`.
`User Name`	Name of subscriber.
`LS:RI`	Logical system and routing instance associated with the subscriber.
`Type`	Subscriber client type (DHCP, FWA, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
`IP Address`	Subscriber IPv4 address.
`IP Netmask`	Subscriber IP netmask. (MX Series) This field displays 255.255.255.255 by default. For tunneled or terminated PPP subscribers only, this field displays the actual value of Framed-IP-Netmask when the SDB_FRAMED_PROTOCOL attribute in the session database is equal to AUTHD_FRAMED_PROTOCOL_PPP. This occurs in the use case where the L2TP network server (LNS) generates access-internal routes when it receives Framed-IP-Netmask from RADIUS during authorization. When it receives Framed-Pool from RADIUS, the pool mask is ignored and the default /32 mask is used.
`Primary DNS Address`	IP address of the primary Domain Name System (DNS) server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Secondary DNS Address`	IP address of the secondary DNS server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.
`IPv6 Primary DNS Address`	IPv6 address of the primary DNS server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.
`IPv6 Secondary DNS Address`	IPv6 address of the secondary DNS server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.
`Domain name server inet`	IP addresses for the DNS server, displayed in order of configuration. This field is displayed with the `extensive` option only when the addresses are derived from the access profile or the global access configuration.
`Domain name server inet6`	IPv6 addresses for the DNS server, displayed in order of configuration. This field is displayed with the `extensive` option only when the addresses are derived from the access profile or the global access configuration.
`Primary WINS Address`	IP address of the primary Windows Internet Name Service (WINS) server.
`Secondary WINS Address`	IP address of the secondary WINS server.
`IPv6 Address`	Subscriber IPv6 address, or multiple addresses.
`IPv6 Prefix`	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
`IPv6 User Prefix`	IPv6 prefix obtained through Neighbor Discovery Router Advertisement (NDRA).

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`IPv6 Address Pool`	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
`IPv6 Network Prefix Length`	Length of the network portion of the IPv6 address.
`IPv6 Prefix Length`	Length of the subscriber IPv6 prefix.
`Logical System`	Logical system associated with the subscriber.
`Routing Instance`	Routing instance associated with the subscriber.
`Interface`	(Enhanced subscriber management for the MX Series routers) Name of the enhanced subscriber management logical interface, in the form `demux0.`*nnnn* (for example, `demux0.3221225472`), to which access-internal and framed subscriber routes are mapped.
`Interface Type`	Subscriber interface (`Static` or `Dynamic`)

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Interface Set`	Internally generated name of the dynamic agent circuit identifier (ACI) or ATM line interface (ALI) interface set used by the subscriber session. The prefix of the name indicates the string received in DHCP or PPPoE control packets on which the interface set is based. For ALI interface sets, the prefix indicates that the value is configured as a trusted option to identify the subscriber line. The name of the interface set uses one of the following prefixes: • `aci`—ACI; for example, aci-1033-demux0.3221225524. This is the only prefix allowed for ACI interface sets. • `ari`—ARI or agent remote identifier; for example, ari-1033-demux0.3221225524. • `aci+ari`—Both the ACI and the ARI; for example, aci+ari-1033-demux0.3221225524. • `noids`—Neither the ACI nor the ARI were received; for example, noids-1033-demux0.3221225524. **NOTE**: ACI interface sets are configured with the `agent-circuit-identifier` autoconfiguration stanza. ALI interface sets are configured with the `line-identity` autoconfiguration stanza. Besides dynamic ACI and ALI interface sets, this field can be an interface set based on a substring of the ARI string. This occurs when the dynamic profile includes the predefined variable *$junos-pon-id-interface-set-name*, and the profile is applied for a passive optical network (PON). The ARI string is inserted by the optical line terminal (OLT). The final substring in the string, unique for the PON, identifies individual subscriber circuits, and is used as the name of the interface set.
`Interface Set Type`	Interface type of the ACI interface set: `Dynamic`. This is the only ACI interface set type currently supported.
`Interface Set Session ID`	Identifier of the dynamic ACI interface set entry in the session database.
`Underlying Interface`	Name of the underlying interface for the subscriber session.
`Dynamic Profile Name`	Dynamic profile used for the subscriber.

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Dynamic Profile Version`	Version number of the dynamic profile used for the subscriber.
`MAC Address`	MAC address associated with the subscriber.
`State`	Current state of the subscriber session (`Init`, `Configured`, `Active`, `Terminating`, or `Tunneled`).
`L2TP State`	Current state of the L2TP session (`Tunneled` or `Tunnel-switched`). When the value is `Tunnel-switched`, two entries are displayed for the subscriber; the first entry is at the L2TP network server (LNS) interface on the L2TP tunnel switch (LTS) and the second entry is at the L2TP access concentrator (LAC) interface on the LTS.
`Tunnel switch Profile Name`	Name of the L2TP tunnel switch profile that initiates tunnel switching.
`Local IP Address`	IP address of the local gateway (LAC).
`Remote IP Address`	IP address of the remote peer (LNS).
`PFE Flow ID`	Forwarding flow identifier.
`VLAN Id`	VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
`Stacked VLAN Id`	Stacked VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
`RADIUS Accounting ID`	RADIUS accounting ID associated with the subscriber.

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Agent Circuit ID`	For the `dhcp` client type, option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in the hexadecimal format. For the `vlan-oob` client type, the agent circuit ID or access-loop circuit identifier that identifies the subscriber line based on the subscriber-facing DSL access multiplexer (DSLAM) interface on which the subscriber request originates.
`Agent Remote ID`	For the `dhcp` client type, option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format. For the `vlan-oob` client type, the agent remote ID or access-loop remote identifier that identifies the subscriber line based on the network access server (NAS)-facing DSLAM interface on which the subscriber request originates.
`Aggregation Interface-set Name`	Value of the $junos-aggregation-interface-set-name predefined variable. The value is determined based on the following factors:; • When the `hierarchical-access-network-detection` option is configured for the access lines and the value of the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003) received either in the Access Node Control Protocol (ANCP) Port Up message or PPPoE Active Discover Request (PADR) IA tags begins with a # character, then the variable takes the value of the remainder of the string after the # character. • When the `hierarchical-access-network-detection` option is not configured, or if the sting does not begin with the # character, then the variable takes the value specified with the `predefined-variable-defaults` statement.
`Accounting Statistics`	Actual transmitted subscriber accounting statistics by the session ID or the interface. Service accounting statistics are not included. These statistics do not include overhead bytes or dropped packets; they are the accurate statistics used by RADIUS. The statistics are counted when the `actual-transmit-statistics` statement is included in the dynamic profile.
`DHCP Relay IP Address`	IP address used by the DHCP relay agent.

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
ATM VPI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
ATM VCI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.
Login Time	Date and time at which the subscriber logged in.
DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for the DHCPv6 options.
Server DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for the DHCP options.
Server DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for the DHCPv6 options.
DHCPV6 Header	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for the DHCPv6 options.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber in kilobits per second.
IPv4 Input Service Set	Input service set in access dynamic profile.
IPv4 Output Service Set	Output service set in access dynamic profile.
PCEF Profile	Policy and charging enforcement function (PCEF) profile in access dynamic profile.
PCEF Rule/Rulebase	PCC rule or rulebase used in dynamic profile.

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Dynamic configuration`	Values for variables that are passed into the dynamic profile from RADIUS.
`Service activation time`	Time at which the first family in this service became active.
`IPv4 rpf-check Fail Filter Name`	Name of the filter applied by the dynamic profile to the IPv4 packets that fail the reverse-path-forwarding (RPF) check.
`IPv6 rpf-check Fail Filter Name`	Name of the filter applied by the dynamic profile to the IPv6 packets that fail the RPF check.
`DHCP Options`	`len` = number of hex values in the message. The hex values specify the type, length, value (TLV) for the DHCP options, as defined in RFC 2132.
`Session ID`	ID number for a subscriber session.
`Underlying Session ID`	Session ID of the underlying PPPoE interface for the DHCPv6 subscribers on a PPPoE network.
`Service Sessions`	Number of service sessions (that is, a service activated using RADIUS Change of Authorization or CoA) associated with the subscribers.
`Service Session ID`	ID number for a subscriber service session.
`Service Session Name`	Service session profile name.
`Session Timeout (seconds)`	Number of seconds of access provided to the subscriber before the session is automatically terminated.
`Idle Timeout (seconds)`	Number of seconds for which the subscriber can be idle before the session is automatically terminated.

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute. This field is displayed only when the predefined variable $junos-ipv6-address is used in the dynamic profile.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv4 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Input Filter Name	Name assigned to the ADF interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Output Filter Name	Name assigned to the ADF interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`IPv6 Input Filter Name`	Name assigned to the IPv6 input filter (client or service session).
`IPv6 Output Filter Name`	Name assigned to the IPv6 output filter (client or service session).
`IFL Input Filter Name`	Name assigned to the logical interface input filter (client or service session).
`IFL Output Filter Name`	Name assigned to the logical interface output filter (client or service session).
`DSL type`	PPPoE subscriber's access line type reported by the PPPoE intermediate agent in a PPPoE Active Discovery Initiation (PADI) or PPPoE Active Discovery Offer(PADO) packet in the Vendor-Specific-Tags TLV in subattribute DSL-Type (0x0091). The DSL can be one of the following types: `ADSL`, `ADSL2`, `ADSL2+`, `OTHER`, `SDSL`, `VDSL`, or `VDSL2`.
`Frame/Cell Mode`	Mode type of the PPPoE subscriber's access line determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091): • `Cell`—When the DSL line type is one of the following: asymmetric digital subscriber line (ADSL), asymmetric digital subscriber line 2 (ADSL2), or asymmetric digital subscriber line 2 plus (ADSL2+). • `Frame`—When the DSL line type is one of the following: OTHER, SDSL, VDSL, or VDSL2. The value is stored in the subscriber session database.
`Overhead accounting bytes`	Number of bytes added to or subtracted from the actual downstream cell or frame overhead to account for the technology overhead of the DSL line type. The value is determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091). The value is stored in the subscriber session database.

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Actual upstream data rate`	Unadjusted upstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Upstream (0x0081).
`Actual downstream data rate`	Unadjusted downstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Downstream (0x0082).
`Adjusted downstream data rate`	Adjusted downstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.
`Adjusted upstream data rate`	Adjusted upstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database. **NOTE**: This output field is only available on Junos Release 19.1R1 or earlier versions.
`AGF Mode`	Type of access. Adaptive mode indicates an FN-RG.
`Local TEID-U`	Tunnel endpoint identifier (TEID) on the BNG for the GPRS Tunnelling Protocol User Plane (GTP-U) tunnel to the evolved node B (eNodeB). The identifier is allocated by the BNG. A fully qualified local TEID-C consists of this identifier and the `GTPU Tunnel Local IP address` value.
`Local TEID-C`	Tunnel endpoint identifier on the BNG for the GTP-C control plane tunnel to the Mobility management Entity (MME). The identifier is allocated by the BNG. A fully qualified local TEID-C consists of this identifier and the `GTPC Local IP address` value.
`Remote TEID-U`	Tunnel endpoint identifier on the eNodeB for the GTP-U user plane tunnel to the BNG. The identifier is allocated by the eNodeB. A fully qualified remote TEID-U consists of this identifier and the `GTPU Tunnel Remote IP address` value.

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Remote TEID-C`	Tunnel endpoint identifier on the MME for the GPRS tunneling protocol, control (GTP-C) plane tunnel to the BNG. The identifier is allocated by the MME. A fully qualified remote TEID-C consists of this identifier and the `GTPC Remote IP address` value.
`GTPU Tunnel Remote IP address`	IP address of the S1-U interface on the eNodeB for the GPRS tunneling protocol, user plane (GTP-U) tunnel endpoint. A fully qualified remote TEID-U consists of this address and the `Remote TEID-U` value.
`GTPU Tunnel Local IP address`	IP address of the S1-U interface on the BNG for the GTP-U tunnel endpoint. A fully qualified local TEID-U consists of this address and the `Local TEID-U` value
`GTPC Remote IP address`	IP address of the S11 interface on the MME for the GTP-C tunnel endpoint. A fully qualified remote TEID-C consists of this address and the `Remote TEID-C` value.
`GTPC Local IP address`	IP address of the S11 interface on the BNG for the GTP-C tunnel endpoint. A fully qualified local TEID-C consists of this address and the `Local TEID-C` value.
`Access Point Name`	Access point name (APN) for the user equipment. The APN corresponds to the connection and service parameters that the subscriber's mobile device can use for connecting to the carrier's gateway to the Internet.
`Tenant`	Name of the tenant system. You can create multiple tenant system administrators for a tenant system and assign the administrators different permission levels based on your requirements.
`Routing instance`	Name of the routing instance. When a custom routing instance is created for a tenant system, all the interfaces defined in that tenant system are added to that routing instance.

Table 1598: show subscribers Output Fields *(Continued)*

Field Name	Field Description
Dynamic Profile Version Alias	Configured name for a specific variation of a base dynamic profile. The presence of this name presence indicates that the profile configuration is different from that of the base profile. The value is conveyed to the RADIUS server during authentication in the Client-Profile-Name VSA (26–4874–174).

Sample Output

show subscribers (IPv4)

```
user@host> show subscribers
Interface               IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824     10                                    default:default
demux0.1073741824       203.0.113.10        WHOLESALER-CLIENT default:default
demux0.1073741825       203.0.113.3         RETAILER1-CLIENT  test1:retailer1
demux0.1073741826       203.0.113.3         RETAILER2-CLIENT  test1:retailer2
```

show subscribers (IPv6)

```
user@host> show subscribers
Interface          IP Address/VLAN ID   User Name           LS:RI
ge-1/0/0.0         2001:db8:c0:0:0:0/74    WHOLESALER-CLIENT   default:default
     *             2001:db8:1/128          subscriber-25       default:default
```

show subscribers (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers
Interface           IP Address/VLAN ID         User Name                      LS:RI
demux0.1073741834   0x8100.1002 0x8100.1                                 default:default
demux0.1073741835   0x8100.1001 0x8100.1                                 default:default
pp0.1073741836      203.0.113.13             dualstackuser1@example1.com     default:ASP-1
*                   2001:db8:1::/48
*                   2001:db8:1:1::/64
pp0.1073741837      203.0.113.33             dualstackuser2@example1.com     default:ASP-1
```

```
*                   2001:db8:1:2:5::/64
```

show subscribers (Single Session DHCP Dual Stack)

```
user@host> show subscribers

Interface           IP Address/VLAN ID        User Name               LS:RI
demux0.1073741364   192.168.10.10             dual-stack-retail35     default:default
                    2001:db8::100:0:0:0/74                            default:default
                    2001:db8:3ffe:0:4::/64
```

show subscribers (Single Session DHCP Dual Stack detail)

```
user@host> show subscribers id 27 detail
Type: DHCP
User Name: dual-stack-retail33
IP Address: 10.10.0.53
IPv6 Address: 2001:db8:3000:0:0:8003::2
IPv6 Prefix: 2001:db8:3ffe:0:4::/64
Logical System: default
Routing Instance: default
Interface: ae0.3221225472
Interface type: Static
Underlying Interface: ae0.3221225472
Dynamic Profile Name: dhcp-retail-18
MAC Address: 00:00:5E:00:53:02
State: Active
DHCP Relay IP Address: 10.10.0.1
Radius Accounting ID: 27
Session ID: 27
PFE Flow ID: 2
Stacked VLAN Id: 2000
VLAN Id: 1
Login Time: 2014-05-15 10:12:10 PDT
DHCP Options: len 60
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 00 64 01 01 02
00 06 00 04 00 03 00 19 00 03 00 0c 00 00 00 00 00 00 00 00
00 00 00 00 00 19 00 0c 00 00 00 00 00 00 00 00 00 00 00 00
```

show subscribers (LNS on MX Series Routers)

```
user@host> show subscribers
Interface          IP Address/VLAN ID   User Name         LS:RI
si-4/0/0.1         192.0.2.0          user@example.com   default:default
```

show subscribers (L2TP Switched Tunnels)

```
user@host> show subscribers
Interface           IP Address/VLAN ID    User Name              LS:RI
si-2/1/0.1073741842 Tunnel-switched       user@example.com        default:default
si-2/1/0.1073741843 Tunnel-switched       user@example.com        default:default
```

show subscribers aggregation-interface-set-name

```
user@host> show subscribers aggregation-interface-set-name FRA*

Interface              IP Address/VLAN ID              User Name                  LS:RI
ge-1/0/0.3221225472    50                                ancp                       default:isp1-
subscriber
```

show subscribers client-type dhcp detail

```
user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 203.0.113.29
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux
MAC Address: 00:00:5e:00:53:98
State: Active
Radius Accounting ID: user :2304
Login Time: 2009-08-25 14:43:52 PDT
```

```
Type: DHCP
IP Address: 203.0.113.27
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:f3
State: Active
Radius Accounting ID: 1234 :2560
Login Time: 2009-08-25 14:43:56 PDT
```

show subscribers client-type dhcp detail (DHCPv6)

```
user@host> show subscribers client-type dhcp detail
Type: DHCP
User Name: DEFAULTUSER
IPv6 Address: 2001:db8::2
IPv6 Prefix: 2001:db8:1::/64
Logical System: default
Routing Instance: default
Interface: demux0.3221225602
Interface type: Static
Underlying Interface: demux0.3221225602
Dynamic Profile Name: client-profile
MAC Address: 00:00:5E:00:53:01
State: Active
Radius Accounting ID: 142
Session ID: 142
PFE Flow ID: 148
Stacked VLAN Id: 1
VLAN Id: 1
Login Time: 2018-03-29 12:27:38 EDT
DHCP Options: len 56
00 08 00 02 00 00 00 01 00 0e 00 01 00 01 22 4f d0 33 00 11
01 00 00 01 00 03 00 0c 00 00 00 0a 00 04 9d 40 00 07 62 00
00 19 00 0c 00 00 00 0b 00 04 9d 40 00 07 62 00
Server DHCPV6 Options: len 94
00 0a 00 06 11 22 33 44 55 66 00 11 00 09 00 00 0c 4c 00 02
```

```
00 01 aa 00 11 00 20 00 00 0a 4c 00 02 00 02 32 33 00 03 00
03 34 35 36 00 05 00 06 31 32 33 34 35 36 00 06 00 01 31 00
11 00 09 00 00 0b 4c 00 02 00 01 bb 00 11 00 12 00 00 0d e9
00 01 00 03 aa bb cc 00 02 00 03 dd ee cc
DHCPV6 Header: len 4
01 fc e4 96
```

show subscribers client-type dhcp extensive

```
user@host> show subscribers client-type dhcp extensive
Type: DHCP
User Name: user
IP Address: 192.0.2.4
IP Netmask: 255.0.0.0
IPv6 Address: 2001:db8:3::103
IPv6 Prefix: 2001:db8::/68
Domain name server inet6: 2001:db8:1 abcd::2
Logical System: default
Routing Instance: default
Interface: ge-0/0/0.0
Interface type: Static
Underlying Interface: ge-0/0/0.0
MAC Address: 00:00:5e:00:53:01
State: Configured
Radius Accounting ID: 10
Session ID: 10
PFE Flow ID: 2
VLAN Id: 100
Agent Circuit ID: ge-0/0/0:100
Agent Remote ID: ge-0/0/0:100
Login Time: 2017-05-23 12:52:22 IST
DHCPV6 Options: len 69
00 01 00 0e 00 01 00 01 59 23 e3 31 00 10 94 00 00 01 00 08
00 02 00 00 00 19 00 29 00 00 00 00 00 04 9d 40 00 07 62 00
00 1a 00 19 00 09 3a 80 00 27 8d 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
Server DHCP Options: len 13
3a 04 00 00 00 ff 00 3b 04 00 00 0f 00
Server DHCPV6 Options: len 8
00 0a 00 04 ab cd ef ab
DHCPV6 Header: len 4
```

```
01 00 00 04
IP Address Pool: al_pool30
IPv6 Address Pool: ia_na_pool
IPv6 Delegated Address Pool: prefix_delegate_pool
```

show subscribers client-type fixed-wireless-access

```
user@host> show subscribers client-type fixed-wireless-access

Interface             IP Address/VLAN ID             User Name                 LS:RI
ps1.3221225472        192.0.2.10                     505024101215074           default:default
ps1.3221225473        192.0.2.11                     505024101215075           default:default
```

show subscribers client-type fixed-wireless-access detail (Detail)

```
user@host> show subscribers client-type fixed-wireless-access detail
Type: FWA
User Name: 505024101215074
IP Address: 192.0.2.10
IP Netmask: 255.255.0.0
Interface: ps1.3221225472
Interface type: Dynamic
Dynamic Profile Name: fwa-profile
State: Active
Radius Accounting ID: 1
Session ID: 1
PFE Flow ID: 11
Login Time: 2019-04-10 14:10:12 PDT
Local TEID-U: 1
Local TEID-C: 1
Remote TEID-U: 2000000
Remote TEID-C: 1000000
GTPU Tunnel Remote IP Address: 203.0.113.1.3
GTPU Tunnel Local IP Address: 203.0.113.2.5
GTPC Remote IP Address: 203.0.113.1.2
GTPC Local IP Address: 203.0.113.1.1
Access Point Name: user21
```

show subscribers client-type vlan-oob detail

```
user@host> show subscribers client-type vlan-oob detail
Type: VLAN-OOB
User Name: L2WS.line-aci-1.line-ari-1
Logical System: default
Routing Instance: ISP1
Interface: demux0.1073744127
Interface type: Dynamic
Underlying Interface: ge-1/0/0
Dynamic Profile Name: Prof_L2WS
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 1234
Session ID: 77
VLAN Id: 126
Core-Facing Interface: ge-2/1/1
VLAN Map Id: 6
Inner VLAN Map Id: 2001
Agent Circuit ID: line-aci-1
Agent Remote ID: line-ari-1
Login Time: 2013-10-29 14:43:52 EDT
```

show subscribers count

```
user@host> show subscribers count
Total Subscribers: 188, Active Subscribers: 188
```

show subscribers address detail (IPv6)

```
user@host> show subscribers address 203.0.113.137 detail
Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 203.0.113.137
IP Netmask: 255.0.0.0
IPv6 User Prefix: 2001:db8:0:c88::/32
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
```

```
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:00:5e:00:53:53
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: example demux0.8201:6544
Session ID: 6544
Agent Circuit ID: ifl3720
Agent Remote ID: ifl3720
Login Time: 2012-05-21 13:37:27 PDT
Service Sessions: 1
```

show subscribers detail (IPv4)

```
user@host> show subscribers detail
Type: DHCP
IP Address: 203.0.113.29
IP Netmask: 255.255.0.0
Primary DNS Address: 192.0.2.0
Secondary DNS Address: 192.0.2.1
Primary WINS Address: 192.0.2.3
Secondary WINS Address: 192.0.2.4
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:98
State: Active
Radius Accounting ID: example :2304
Idle Timeout (seconds): 600
Login Time: 2009-08-25 14:43:52 PDT
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
33 2d 37 2d 30 37 05 01 06 0f 21 2c
Service Sessions: 2
```

show subscribers detail (IPv6)

```
user@host> show subscribers detail
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2001:db8:ffff:1::/32
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:00:5e:00:53:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
```

show subscribers detail (pseudowire Interface for GRE Tunnel)

```
user@host> show subscribers detail
Interface           IP Address/VLAN ID        User Name           LS:RI
ps0.3221225484      192.0.2.2
ps0.3221225485      192.0.2.3
demux0.3221225486    1                                            default:default
demux0.3221225487    1                                            default:default
demux0.3221225488   198.51.0.1                                    default:default
demux0.3221225489   198.51.0.2                                    default:default
```

show subscribers detail (IPv6 Static Demux Interface)

```
user@host> show subscribers detail
Type: STATIC-INTERFACE
User Name: user@example.com
IPv6 Prefix: 2001:db8:3:4:5:6:7:aa/32
Logical System: default
Routing Instance: default
```

```
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT
```

show subscribers detail (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers detail
Type: L2TP
User Name: user@example.com
IP Address: 203.0.113.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
```

show subscribers detail (L2TP Switched Tunnels)

```
user@host> show subscribers detail
Type: L2TP
User Name: user@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741842
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 203.0.113.51
Remote IP Address: 192.0.2.0
```

```
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST

Type: L2TP
User Name: user@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741843
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 203.0.113.31
Remote IP Address: 192.0.2.1
Session ID: 22
Login Time: 2013-01-18 03:01:14 PST
```

show subscribers detail (Tunneled Subscriber)

```
user@host> show subscribers detail
Type: PPPoE
User Name: user1@example.com
Logical System: default
Routing Instance: default
Interface: pp0.1
State: Active, Tunneled
Radius Accounting ID: 512
```

show subscribers detail (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
```

```
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@example1.com
IP Address: 203.0.113.13
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST

Type: DHCP
IPv6 Prefix: 2001:db8:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
```

show subscribers detail (ACI Interface Set Session)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0
Interface Set: aci-1001-ge-1/0/0.2800
Interface Set Session ID: 0
Underlying Interface: ge-1/0/0.2800
Dynamic Profile Name: aci-vlan-set-profile-2
Dynamic Profile Version: 1
State: Active
Session ID: 1
Agent Circuit ID: aci-ppp-dhcp-20
Login Time: 2012-05-26 01:54:08 PDT
```

show subscribers detail (PPPoE Subscriber Session with ACI Interface Set)

```
user@host> show subscribers detail
Type: PPPoE
User Name: ppphint2
IP Address: 203.0.113.15
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Interface Set: aci-1001-demux0.1073741824
Interface Set Type: Dynamic
Interface Set Session ID: 2
Underlying Interface: demux0.1073741824
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 3
Session ID: 3
Agent Circuit ID: aci-ppp-dhcp-dvlan-50
Login Time: 2012-03-07 13:46:53 PST
```

show subscribers detail (Dynamic Profile Version Alias)

```
user@host> show subscribers detail

Type: PPPoE
User Name: DEFAULTUSER
IP Address: 192.0.2.21
IP Netmask: 255.255.255.255
IPv6 Address: 2001:db8::17
Logical System: default
Routing Instance: default
Interface: pp0.3221225720
Interface type: Dynamic
Underlying Interface: demux0.3221225719
Dynamic Profile Name: pppoe-client-profile
Dynamic Profile Version Alias: profile-version1a
MAC Address: 00:00:5E:00:53:38
State: Active
Radius Accounting ID: 288
Session ID: 288
PFE Flow ID: 344
VLAN Id: 1
Login Time: 2019-09-23 10:40:56 IST
```

show subscribers extensive

```
user@host> show subscribers extensive
Type: DHCP
User Name: uer@host
IP Address: 192.0.2.136
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: ge-0/0/0.0
Interface type: Static
Underlying Interface: ge-0/0/0.0
MAC Address: 00:10:94:00:00:01
State: Active
Radius Accounting ID: 15
Session ID: 15
```

```
PFE Flow ID: 2
VLAN Id: 100
Login Time: 2021-05-24 11:30:07 IST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 31 2f
31 2d 30 2d 30 37 05 01 06 0f 21 2c
DHCP Header: len 44
01 01 06 00 00 00 00 1d 00 00 80 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 10 94 00 00 01 00 00 00 00 00 00
00 00 00 00
IP Address Pool: al_pool30
Access Line Attributes:
  Actual upstream data rate: 19998
  Actual downstream data rate: 79999
  Access loop encapsulation: 01 02 00
```

show subscribers extensive (Aggregation Node Interface Set and DSL Forum Attributes)

```
user@host> show subscribers extensive
Type: VLAN-OOB
User Name: ancp
Logical System: default
Routing Instance: isp1-subscriber
Interface: ge-1/0/0.3221225472
Interface type: Dynamic
Interface Set: FRA-DPU-C-100
Underlying Interface: ge-1/0/0
Core IFL Name: ge-1/0/4.0
Dynamic Profile Name: Prof_L2BSA
State: Active
Radius Accounting ID: 1
Session ID: 1
PFE Flow ID: 13
VLAN Id: 50
VLAN Map Id: 20
Inner VLAN Map Id: 1
Inner VLAN Tag Protocol Id: 0x88a8
Agent Circuit ID: circuit 201
Agent Remote ID: remote-id
Aggregation Interface-set Name: FRA-DPU-C-100
```

```
Login Time: 2018-05-29 08:43:42 EDT
Accounting interval: 72000
Dynamic configuration:
  junos-cos-scheduler-map: 100m
  junos-inner-vlan-tag-protocol-id: 0x88a8
  junos-vlan-map-id: 20

Type: PPPoE
IP Address: 192.85.128.1
IP Netmask: 255.255.255.255
Logical System: default
Routing Instance: default
Interface: pp0.3221225474
Interface type: Dynamic
Interface Set: ge-1/0/0
Underlying Interface: demux0.3221225473
Dynamic Profile Name: pppoe-client-profile-with-cos
MAC Address: 00:10:94:00:00:03
State: Active
Radius Accounting ID: 3
Session ID: 3
PFE Flow ID: 16
Stacked VLAN Id: 50
VLAN Id: 7
Agent Circuit ID: circuit 201
Agent Remote ID: remote-id
Aggregation Interface-set Name: FRA-DPU-C-100
Login Time: 2018-05-29 08:43:45 EDT
IP Address Pool: pool-1
Accounting interval: 72000
DSL type: G.fast
Frame/cell mode: Frame
Overhead accounting bytes: 10
Actual upstream data rate: 100000 kbps
Actual downstream data rate: 200000 kbps
Calculated downstream data rate: 180000 kbps
Calculated upstream data rate: 90000 kbps

Adjusted downstream data rate: 160000 kbps
DSL Line Attributes
  Agent Circuit ID: circuit 201
  Agent Remote ID: remote-id
  Actual upstream data rate: 100000
```

```
  Actual downstream data rate: 200000
  DSL type: G.fast
  Access Aggregation Circuit ID: #FRA-DPU-C-100
  Attribute type: 0xAA, Attribute length: 4
    198 51 100 78
```

show subscribers extensive (Passive Optical Network Circuit Interface Set)

```
user@host> show subscribers client-type dhcp extensive
Type: DHCP
IP Address: 192.0.2.136
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073741842
Interface type: Dynamic
Interface Set: otl01.xyz101-202
Underlying Interface: demux0.1073741841
Dynamic Profile Name: dhcp-profile
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: user :19
Session ID: 19
VLAN Id: 1100
Agent Remote ID: ABCD01234|100M|AAAA01234|otl01.xyz101-202

Login Time: 2017-03-29 10:30:46 PDT
DHCP Options: len 97
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 02 33 04 00 00
17 70 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
32 2d 31 2d 31 37 05 01 06 0f 21 2c 52 2b 02 29 41 42 43 44
30 31 32 33 34 7c 31 30 30 4d 7c 41 41 41 41 30 31 32 33 34
7c 6f 74 6c 30 31 2e 78 79 7a 31 30 31 2d 32 30 32
IP Address Pool: POOL-V4
```

show subscribers extensive (DNS Addresses from Access Profile or Global Configuration)

```
user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com
```

```
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Domain name server inet: 198.51.100.1 198.51.100.2
IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

show subscribers extensive (DNS Addresses from RADIUS)

```
user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
IPv6 Address: 2001:db8::1:11
IPv6 Primary DNS Address: 2001:db8:5001::12
IPv6 Secondary DNS Address: 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
```

```
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

show subscribers extensive (IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration)

```
user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

show subscribers extensive (RPF Check Fail Filter)

```
user@host> show subscribers extensive
...
Type: VLAN
  Logical System: default
  Routing Instance: default
  Interface: ae0.1073741824
  Interface type: Dynamic
  Dynamic Profile Name: vlan-prof
  State: Active
  Session ID: 9
  VLAN Id: 100
  Login Time: 2011-08-26 08:17:00 PDT
  IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
  IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...
```

show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers extensive
Type: L2TP
User Name: user@example.com
IP Address: 203.0.113.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out
```

show subscribers extensive (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@example1.com
IP Address: 203.0.113.13
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2001:db8:2016:1:1::1/64
IPv6 Framed Interface Id: 1:1:2:2
IPv4 Input Filter Name: FILTER-IN-pp0.1073741825-in
IPv4 Output Filter Name: FILTER-OUT-pp0.1073741825-out
IPv6 Input Filter Name: FILTER-IN6-pp0.1073741825-in
IPv6 Output Filter Name: FILTER-OUT6-pp0.1073741825-out

Type: DHCP
IPv6 Prefix: 2001:db8:1::/32
Logical System: default
Routing Instance: ASP-1
```

```
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Network Prefix Length: 48
```

show subscribers extensive (ADF Rules)

```
user@host> show subscribers extensive
...
   Service Session ID: 12
   Service Session Name: SERVICE-PROFILE
   State: Active
   Family: inet
     ADF IPv4 Input Filter Name: __junos_adf_12-demux0.3221225474-inet-in
                         Rule 0: 010101000b0101020b020200201811
                                 from {
                                     source-address 203.0.113.232;
                                     destination-address 198.51.100.0/24;
                                     protocol 17;
                                 }
                                 then {
                                     accept;
                                 }
```

show subscribers extensive (Effective Shaping-Rate)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
```

```
Interface: demux0.1073741837
Interface type: Dynamic
Interface Set: ifset-1
Underlying Interface: ae1
Dynamic Profile Name: svlan-dhcp-test
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.201
VLAN Id: 0x8100.201
Login Time: 2011-11-30 00:18:04 PST
Effective shaping-rate: 31000000k
...
```

show subscribers extensive (PPPoE Subscriber Access Line Rates)

```
user@host> show subscribers extensive
Type: PPPoE
 IP Address: 198.51.100.1
 IP Netmask: 255.255.255.255
 Logical System: default
 Routing Instance: default
 Interface: pp0.3221225475
 Interface type: Dynamic
 Underlying Interface: demux0.3221225474
 Dynamic Profile Name: pppoe-client-profile-with-cos
 MAC Address: 00:00:5e:00:53:02
 State: Active
 Radius Accounting ID: 4
 Session ID: 4
 PFE Flow ID: 14
 Stacked VLAN Id: 40
 VLAN Id: 1
 Agent Circuit ID: circuit0
 Agent Remote ID: remote0
 Login Time: 2017-04-06 15:52:32 PDT

 User Name: DAVE-L2BSA-SERVICE
 Logical System: default
 Routing Instance: isp-1-subscriber
 Interface: ge-1/2/4.3221225472
 Interface type: Dynamic
```

```
  Interface Set: ge-1/2/4
  Underlying Interface: ge-1/2/4
  Core IFL Name: ge-1/3/4.0
  Dynamic Profile Name: L2BSA-88a8-400LL1300VO
  State: Active
 Radius Accounting ID: 1
  Session ID: 1
  PFE Flow ID: 14
  VLAN Id: 13
  VLAN Map Id: 102
  Inner VLAN Map Id: 1
  Agent Circuit ID: circuit-aci-3
  Agent Remote ID: remote49-3
  Login Time: 2017-04-05 16:59:29 EDT
  Service Sessions: 4
  IFL Input Filter Name: L2BSA-CP-400LL1300VO-ge-1/2/4.3221225472-in
  IFL Output Filter Name: L2BSA-CP-400LL1300VO-ge-1/2/4.3221225472-out
  Accounting interval: 900
  DSL type: VDSL
  Frame/Cell Mode: Frame
  Overhead accounting bytes: -10
  Actual upstream data rate: 1024 kbps
  Actual downstream data rate: 4096 kbps
  Adjusted downstream data rate: 3686 kbps
   Dynamic configuration:
    junos-vlan-map-id: 102
    Service Session ID: 5
    Service Session Name: SRL-L1
    State: Active
    Family: inet, inet6
    IFL Input Filter Name: L2BSA-FWF-in-10048-ge-1/2/4.3221225472-in
    IFL Output Filter Name: L2BSA-FWF-out-25088-ge-1/2/4.3221225472-out
    Service Activation time: 2017-04-05 16:59:30 EDT
    Dynamic configuration:
      l2bsa-fwf-in: L2BSA-FWF-in-10048
      l2bsa-fwf-out: L2BSA-FWF-out-25088
      rldown: 25088
      rlup: 10048
```

show subscribers extensive (Subscriber Session Using PCEF Profile)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.3221225517
Interface type: Dynamic
Underlying Interface: ge-1/0/3
Dynamic Profile Name: svlan-dhcp
State: Active
Session ID: 59
PFE Flow ID: 71
Stacked VLAN Id: 0x8100.1
VLAN Id: 0x8100.2
Login Time: 2017-03-28 08:23:08 PDT

Type: DHCP
User Name: pcefuser
IP Address: 192.0.2.26
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225518
Interface type: Dynamic
Underlying Interface: demux0.3221225517
Dynamic Profile Name: dhcp-client-prof
MAC Address: 00:00:5e:00:53:01
State: Active
Radius Accounting ID: 60
Session ID: 60
PFE Flow ID: 73
Stacked VLAN Id: 1
VLAN Id: 2
Login Time: 2017-03-28 08:23:08 PDT
Service Sessions: 1
DHCP Options: len 9
35 01 01 37 04 01 03 3a 3b
IP Address Pool: pool-ipv4
IPv4 Input Service Set: tdf-service-set
IPv4 Output Service Set: tdf-service-set
PCEF Profile: pcef-prof-1
```

```
PCEF Rule/Rulebase: default
Dynamic configuration:
  junos-input-service-filter: svc-filt-1
  junos-input-service-set: tdf-service-set
  junos-output-service-filter: svc-filt-1
  junos-output-service-set: tdf-service-set
  junos-pcef-profile: pcef-prof-1
  junos-pcef-rule: default

   Service Session ID: 61
   Service Session Name: pcef-serv-prof
   State: Active
   Family: inet
   IPv4 Input Service Set: tdf-service-set
   IPv4 Output Service Set: tdf-service-set
   PCEF Profile: pcef-prof-1
   PCEF Rule/Rulebase: limit-fb
   Service Activation time: 2017-03-28 08:31:19 PDT
   Dynamic configuration:
     pcef-prof: pcef-prof-1
     pcef-rule1: limit-fb
     svc-filt: svc-filt-1
     svc-set: tdf-service-set
```

show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set)

```
user@host> show subscribers aci-interface-set-name aci-1003-ge-1/0/0.4001 detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
```

```
User Name: ppphint2
IP Address: 203.0.113.17
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address:
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT
```

show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring)

```
user@host> show subscribers agent-circuit-identifier aci-ppp-vlan detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 203.0.113.17
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
```

```
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:5e:00:53:52
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT
```

show subscribers id accounting-statistics

```
user@host> show subscribers id 601 accounting-statistics
Session ID: 601
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
```

show subscribers interface accounting-statistics

```
user@host> show subscribers interface pp0.3221226949 accounting-statistics
Session ID: 501
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:
Input bytes : 0
```

```
Output bytes : 0
Input packets: 0
Output packets: 0

Session ID: 502
Accounting Statistics:
Input bytes : 87654
Output bytes : 72108
Input packets: 3322
Output packets: 3322
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Session ID: 503
Accounting Statistics:
Input bytes : 156528
Output bytes : 123865
Input packets: 7448
Output packets: 7448
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
```

show subscribers interface extensive

```
user@host> show subscribers interface demux0.1073741826 extensive
Type: VLAN
User Name: user@test.example.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:5e:00:53:04
State: Active
Radius Accounting ID: 12
```

```
Session ID: 12
Stacked VLAN Id: 0x8100.1500
VLAN Id: 0x8100.2902
Login Time: 2011-10-20 16:21:59 EST

Type: DHCP
User Name: user@test.example.com
IP Address: 192.0.2.0
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:5e:00:53:04
State: Active
Radius Accounting ID: 21
Session ID: 21
Login Time: 2011-10-20 16:24:33 EST
Service Sessions: 2

Service Session ID: 25
Service Session Name: SUB-QOS
State: Active

Service Session ID: 26
Service Session Name: service-cb-content
State: Active
IPv4 Input Filter Name: content-cb-in-demux0.1073741826-in
IPv4 Output Filter Name: content-cb-out-demux0.1073741826-out
```

show subscribers logical-system terse

```
user@host> show subscribers logical-system test1 terse
Interface                       IP Address/VLAN ID   User Name         LS:RI
demux0.1073741825           203.0.113.3         RETAILER1-CLIENT  test1:retailer1
demux0.1073741826           203.0.113.4         RETAILER2-CLIENT  test1:retailer2
```

show subscribers physical-interface count

```
user@host> show subscribers physical-interface ge-1/0/0 count
Total subscribers: 3998, Active Subscribers: 3998
```

show subscribers routing-instance inst1 count

```
user@host> show subscribers routing-instance inst1 count
Total Subscribers: 188, Active Subscribers: 183
```

show subscribers stacked-vlan-id detail

```
user@host> show subscribers stacked-vlan-id 101 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id detail (Combined Output)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers user-name detail

```
user@host> show subscribers user-name larry1 detail
Type: DHCP
User Name: larry1
IP Address: 203.0.113.37
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.1
Interface type: Static
Dynamic Profile Name: foo
MAC Address: 00:00:5e:00:53:01
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-11-07 08:25:59 PST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
37 2d 30 2d 30 37 05 01 06 0f 21 2c
```

show subscribers vlan-id

```
user@host> show subscribers vlan-id 100
Interface           IP Address                              User Name
```

```
ge-1/0/0.1073741824
ge-1/2/0.1073741825
```

show subscribers vlan-id detail

```
user@host> show subscribers vlan-id 100 detail
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT
```

show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session)

```
user@host> show subscribers vpi 40 vci 50 extensive
Type: PPPoE
User Name: testuser
IP Address: 203.0.113.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
ATM VPI: 40
ATM VCI: 50
Login Time: 2012-12-03 07:49:26 PST
```

```
IP Address Pool: pool_1
IPv6 Framed Interface Id: 200:65ff:fe23:102
```

show subscribers address detail (Enhanced Subscriber Management)

```
user@host> show subscribers address 203.0.113.111 detail
Type: DHCP
User Name: simple_filters_service
IP Address: 203.0.113.111
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225482
Interface type: Dynamic
Underlying Interface: demux0.3221225472
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:0f
State: Active
Radius Accounting ID: 11
Session ID: 11
PFE Flow ID: 15
Stacked VLAN Id: 210
VLAN Id: 209
Login Time: 2014-03-24 12:53:48 PDT
Service Sessions: 1
DHCP Options: len 3
35 01 01
```

show subscribers extensive (Tenant Systems)

```
user@host:TSYS1> show subscribers extensive
Type: XAUTH
User Name: userX
+    Tenant: TSYS1
    Routing Instance: TSYS1-ri
IP Address: 192.0.2.0
IP Netmask: 203.0.113.0
Primary DNS Address: 198.51.100.0
Secondary DNS Address: 198.51.100.1
Dynamic Profile Name: radius
```

```
State: Active
Session ID: 1
Login Time: 2018-09-18 13:49:00 PDT
```

Sample Output for AGF

The following sample output shows subscribers connected to the AGF:

show subscribers extensive (DHCP on an FN-RG)

```
user@host> show subscribers extensive
Type: DHCP
User Name: USER2
IP Address: 172.16.0.227
Logical System: default
Routing Instance: default
Interface: demux0.3221230587
Interface type: Dynamic
Underlying Interface: demux0.3221230586
Dynamic Profile Name: dhcp-profile
MAC Address: 00:44:46:44:44:44
State: Active
DHCP Relay IP Address: 10.1.0.1
Radius Accounting ID: 5128
Session ID: 5128
PFE Flow ID: 5185
Stacked VLAN Id: 1
VLAN Id: 1
Agent Circuit ID: aci1
Agent Remote ID: ari1
Login Time: 2022-04-26 09:24:56 PDT
Service Sessions: 1
DHCP Options: len 9
35 01 01 37 04 01 03 3a 3b
DHCP Header: len 44
01 01 06 00 10 bd b4 93 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 44 46 44 44 44 00 00 00 00 00 00
00 00 00 00
AGF Mode: Adaptive
Local TEID-U: 10183810
Remote TEID-U: 6354992
```

```
GTPU Tunnel Remote IP Address: 10.0.0.1
GTPU Tunnel Local IP Address: 10.0.0.1
5G-QFI: 1
IPv6 Framed Interface Id: 0:7fff:ffff:eaa0
IPv4 Input Filter Name: RG-LWAC-V4-INPUT-FILTER-01-demux0.3221230587-in
IPv4 Output Filter Name: RG-LWAC-V4-OUTPUT-FILTER-01-demux0.3221230587-out
Access Line Attributes:
  Agent Circuit ID: aci1
  Agent Remote ID: ari1
Dynamic configuration:
  junos-cos-guaranteed-rate: 1000000
  junos-cos-guaranteed-rate-burst: 250000
      junos-cos-scheduler: GOLD
          junos-cos-scheduler-shaping-rate: 200000000
          junos-cos-scheduler-tx: 200000000
  junos-cos-scheduler-map: DATA_VOICE_VIDEO_SMAP_UID1468
  junos-cos-shaping-rate: 200000000
  junos-cos-shaping-rate-burst: 50000
  junos-cos-traffic-control-profile: TRAFFIC_CONTROL_PROFILE
  junos-input-filter: RG-LWAC-V4-INPUT-FILTER-01
  junos-input-ipv6-filter: RG-LWAC-V6-INPUT-FILTER-01
  junos-output-filter: RG-LWAC-V4-OUTPUT-FILTER-01
  junos-output-ipv6-filter: RG-LWAC-V6-OUTPUT-FILTER-01

   Service Session ID: 5129
   Service Session Name: SERVICE-PROFILE-BASIC-POLICER
   State: Active
   Family: inet
   Service session type: Service-Profile
   IPv4 Input Filter Name: CAP-POLICER-demux0.3221230587-in
   IPv4 Output Filter Name: CAP-POLICER-demux0.3221230587-out
   Service Activation time: 2022-04-26 09:24:57 PDT
   Dynamic configuration:
     bandwidth-limit: 200k
     burst-size-limit: 75k
```

Release Information

Command introduced in Junos OS Release 9.3.

`client-type`, `mac-address`, `subscriber-state`, and `extensive` options introduced in Junos OS Release 10.2.

`count` option usage with other options introduced in Junos OS Release 10.2.

Options `aci-interface-set-name` and `agent-circuit-identifier` introduced in Junos OS Release 12.2.

The `physical-interface` and `user-name` options introduced in Junos OS Release 12.3.

Options `vci` and `vpi` introduced in Junos OS Release 12.3R3 and supported in later 12.3R*x* releases.

Options `vci` and `vpi` supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

Enhanced subscriber management supported in Junos OS Release 15.1R3 on MX Series routers.

`accounting-statistics` option added in Junos OS Release 15.1R3 and 17.4R1 on MX Series routers.

`aggregation-interface-set-name` option added in Junos OS Release 18.4R1 on MX Series routers.

`sub-system` added in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers

Verifying and Managing Junos OS Enhanced Subscriber Management

show subscribers summary

IN THIS SECTION

- Syntax | 26664
- Description | 26664
- Options | 26664
- Required Privilege Level | 26665
- Output Fields | 26665
- Sample Output | 26668
- Release Information | 26673

Syntax

```
show subscribers summary
<all>
<detail | extensive | terse>
<count>
<physical-interface physical-interface-name>
<logical-system logical-system pic | port | routing-instance routing-instance | slot>
<sub-system (agf | bng)>
```

Description

Display summary information for subscribers.

Options

none	Display summary information by state and client type for all subscribers.
all	(Optional) Display summary information by state, client type, and logical system and routing instance (LS:RI).
detail | extensive | terse	(Not supported on MX Series routers) (Optional) Display the specified level of output.
count	(Not supported on MX Series routers) (Optional) Display the count of total subscribers and active subscribers for any specified option.
logical-system ***logical-system***	(Optional) Display subscribers whose logical system matches the specified logical system.
physical-interface ***physical-interface-name***	(M120, M320, and MX Series routers only) (Optional) Display the count of subscribers whose physical interface matches the specified physical interface, by subscriber state, client type, and LS:RI.
pic	(M120, M320, and MX Series routers only) (Optional) Display the count of subscribers by PIC number and the total number of subscribers.
port	(M120, M320, and MX Series routers only) (Optional) Display the count of subscribers by port number and the total number of subscribers.
routing-instance ***routing-instance***	(Optional) Display subscribers whose routing instance matches the specified routing instance.

slot (M120, M320, and MX Series routers only) (Optional) Display the count of subscribers by Flexible PIC Concentrator (FPC) slot number and the total number of subscribers.

sub-system (agf | bng) (Optional) Display a count of subscribers using either the Access Gateway Function (AGF) or broadband network gateway (BNG) services.

> **NOTE**: The subsystem option is only available when both the AGF and BNG subscribers are logged into the router at the same time.

> **NOTE**: Because of display limitation, the logical system and routing instance output values are truncated when necessary.

Starting from Junos OS 20.4R1 release, you need a license to use the Extensible Subscribers Services Manager (ESSM) feature.

Required Privilege Level

view

Output Fields

Table 1599 on page 26666 lists the output fields for the `show subscribers summary` command. Output fields are listed in the approximate order in which they appear.

Table 1599: show subscribers summary Output Fields

Field Name	Field Description	Level of Output
`Subscribers by State`	Number of subscribers summarized by state. The summary includes the following information: • Init—Number of subscribers currently in the initialization state. • Configured—Number of configured subscribers. • Active—Number of active subscribers. • Terminating—Number of subscribers currently terminating. • Terminated—Number of terminated subscribers. • Total—Total number of subscribers for all states.	`detail` none
`Subscribers by Client Type`	Number of subscribers summarized by client type. Client types can include DHCP, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN, and VLAN-OOB. This field also displays the total number of subscribers for all client types (Total).	`detailextensive` none
`Subscribers by LS:RI`	Number of subscribers summarized by logical system:routing instance (LS:RI) combination. This field also displays the total number of subscribers for all the LS:RI combinations (Total).	`detail` none
`Subscribers by Connection Type`	Number of subscribers summarized by connection type, `Cross-connected` or `Terminated`.	`extensive`
`Subscribers by Sub-System`	Number of subscribers on the subsystem—AGF or BNG as well as the total number of subscriber.	All levels

Table 1599: show subscribers summary Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Interface`	Interface associated with the subscriber. The router or the switch displays subscribers whose interface matches or begins with the specified interface.
	
	The asterisk (*) indicates a continuation of addresses for the same session.
	
	For aggregated Ethernet interfaces, the output of the `summary (pic | port | slot)` options prefixes the interface name with ae0:.
	
	For pseudowire IFDs, this field displays both the pseudowire and the associated logical tunnel (LT) and the redundant logical tunnel (RLT) anchor interface. For example:
	
	`ps0: lt-2/1/0`
	`ps1: rlt0: lt-4/0/0`
		All levels

`Count`	Count of subscribers displayed for each PIC, port, or slot when those options are specified with the `summary` option. For an aggregated Ethernet configuration, the total subscriber count does not equal the sum of the individual PIC, port, or slot counts, because each subscriber can be in more than one aggregated Ethernet link. Multiple pseudowire interfaces can share a given logical tunnel or redundant logical tunnel anchor interface. Starting in Junos OS Release 18.1R1, the field displays subscribers per individual pseudowire interface. In earlier releases, the field displays the same number of subscribers for all the pseudowire interfaces that share the same tunnel interface as their anchor point.	`detailextensive` none
`Total Subscribers`	Total number of subscribers for all physical interfaces, all PICs, all ports, or all LS:RI slots.	`detailextensive` none
`IP Address/VLAN ID`	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id*	`terse`

Table 1599: show subscribers summary Output Fields *(Continued)*

Field Name	Field Description	Level of Output
User Name	Name of the subscriber.	terse
LS:RI	Logical system and routing instance associated with the subscriber.	terse

Sample Output

show subscribers summary

```
user@host> show subscribers summary

Subscribers by State
   Active: 52194
   Total: 52194

Subscribers by Client Type
   DHCP: 10000
   VLAN: 15997
   VLAN-OOB: 3600
   PPPoE: 15998
   ESSM: 6599
   Total: 52194
```

show subscribers summary all

```
user@host> show subscribers summary all
Subscribers by State
    Init                     3
    Configured             2
    Active                  183
    Terminating             2
    Terminated             1

    TOTAL                   191
```

```
Subscribers by Client Type
    DHCP                     107
    PPP                      76
    VLAN                      8

    TOTAL               191

Subscribers by LS:RI
    default:default      1
    default:ri1              28
    default:ri2              16
    ls1:default             22
    ls1:riA                 38
    ls1:riB                 44
    logsysX:routinstY     42

    TOTAL                    191
```

show subscribers summary physical-interface

```
user@host> show subscribers summary physical-interface ge-1/0/0
Subscribers by State
   Active: 3998
   Total: 3998

Subscribers by Client Type
   DHCP: 3998
   Total: 3998

Subscribers by LS:RI
   default:default: 3998
   Total: 3998
```

show subscribers summary physical-interface pic

```
user@host> show subscribers summary physical-interface ge-0/2/0 pic
Subscribers by State
   Active: 4825
   Total: 4825
```

```
Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary physical-interface port

```
user@host> show subscribers summary physical-interface ge-0/3/0 port
Subscribers by State
   Active: 4825
   Total: 4825

Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary physical-interface slot

```
user@host> show subscribers summary physical-interface ge-2/0/0 slot
Subscribers by State
   Active: 4825
   Total: 4825

Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary pic

```
user@host> show subscribers summary pic
Interface           Count
ge-1/0              1000
ge-1/3              1000

Total Subscribers: 2000
```

show subscribers summary pic (Aggregated Ethernet Interfaces)

```
user@host> show subscribers summary pic
Interface           Count
ae0: ge-1/0         801
ae0: ge-1/3         801

Total Subscribers: 801
```

show subscribers summary port

```
user@host> show subscribers summary port
Interface           Count
ge-5/0/1              201
ge-5/0/2              301

Total Subscribers: 502
```

show subscribers summary port (Pseudowire Interfaces)

```
user@host> show subscribers summary port
ps0: lt-2/1/0 10
ps1: lt-2/1/0 20

Total Subscribers: 30
```

show subscribers summary port extensive

```
user@host>show subscribers summary port extensive
Interface: xe-3/0/3
Port Count: 100
Detail:
Subscribers by Client Type
   PPPoE: 1
   ESSM: 99
Subscribers by Connection Type
   Terminated: 1

Interface: xe-3/1/3
Port Count: 3100
Detail:
Subscribers by Client Type
   PPPoE: 1600
   ESSM: 1100
   VLAN-OOB: 400
Subscribers by Connection Type
   Tunneled: 500
   Terminated: 1100
   Cross-connected: 400

Total Subscribers: 26197
```

show subscribers summary slot

```
user@host> show subscribers summary slot
Interface           Count
ge-1                2000

Total Subscribers: 2000
```

show subscribers summary terse

```
user@host> show subscribers summary terse
Interface                       IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824         100                                     default:default
```

```
demux0.1073741824           203.0.113.10         WHOLESALER-CLIENT default:default
demux0.1073741825           203.0.113.13         RETAILER1-CLIENT  test1:retailer1
demux0.1073741826           203.0.113.213        RETAILER2-CLIENT  test1:retailer2
```

show subscribers summary (agf)

```
user@host> show subscribers summary
Subscribers by Sub-System
   AGF: 3000
   BNG: 0
   Total: 3000

Subscribers by State
   Active: 4000
   Total: 4000

Subscribers by Client Type
   VLAN: 1000
   PPPoE: 3000
   Total: 4000
```

show subscribers summary sub-system agf

```
user@host> show subscribers summary sub-system agf
AGF Subscribers by State
    Init:     3
    Active:  12
    Total:   15

AGF Subscribers by Client Type
     DHCP:           2
     PPPOE:          5
     PPPOE-5G        5
     TOTAL          12
```

Release Information

Command introduced in Junos OS Release 10.2.

`sub-system` added in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

show subscribers | 26609

show synchronous-ethernet esmc statistics

IN THIS SECTION

- Syntax | 26674
- Description | 26674
- Required Privilege Level | 26674
- Output Fields | 26675
- Sample Output | 26675
- Release Information | 26676

Syntax

```
show synchronous-ethernet esmc statistics
<brief interface interface-name>
<detail interface interface-name>
< interface interface-name>
< interface interface-name (brief | detail)>
```

Description

(MX5, MX10, MX40, MX80, MX80-T, MX240, MX480, MX960, MX2010, MX2020 , and PTX Series routers only) Display the Synchronous Ethernet ESMC statistics.

Required Privilege Level

maintenance

Output Fields

Table 1600 on page 26675 lists the output fields for the `show synchronous-ethernet esmc statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1600: show synchronous-ethernet esmc statistics Output Fields

Field Name	Field Description
`Interface Name`	`interface-slot/pic/port`—Displays the name of the interface for which the ESMC statistics are displayed.
`Transmit Count`	`number`—Displays the number of ESMC packets transmitted.
`Receive Count`	`number`—Displays the number of ESMC packets received.

Sample Output

show synchronous-ethernet esmc statistics

```
user@host> show synchronous-ethernet esmc statistics
ESMC statistics:
Interface Name       Transmit Count      Receive Count
ge-1/0/4             3540                0
ge-1/0/2             3539                0
ge-1/2/4             3540                0
```

show synchronous-ethernet esmc statistics detail

```
user@host> show synchronous-ethernet esmc statistics detail
ESMC Statistics:

Interface Name  :  xe-2/0/10
Transmit Count  :  40908             Receive Count        : 40534
Total Drop Count:  336               Ineligible Drop Count: 0
Adjacency Count :  4
```

show synchronous-ethernet esmc statistics interface (PTX)

```
user@host> show synchronous-ethernet esmc statistics interface et-1/1/0:0
ESMC statistics:
Interface Name       Transmit Count      Receive Count
et-1/1/0:0           358                 338
```

Release Information

Command introduced in Junos OS Release 11.2R4.

Command introduced in Junos OS Release 14.2R1 for PTX Series Packet Transport Routers.

RELATED DOCUMENTATION

Synchronous Ethernet Overview

Configuring an External Clock Synchronization Interface for MX Series Routers

request chassis synchronization mode | **14570**

request chassis synchronization switch | **14574**

show synchronous-ethernet global-information | **26679**

show synchronous-ethernet esmc transmit | **26676**

clear synchronous-ethernet esmc statistics | **4216**

show synchronous-ethernet esmc transmit

IN THIS SECTION

- Syntax | **26677**
- Description | **26677**
- Required Privilege Level | **26677**
- Output Fields | **26677**
- Sample Output | **26678**
- show synchronous-ethernet esmc transmit (PTX) | **26678**

Release Information | 26678

Syntax

```
show synchronous-ethernet esmc transmit
<brief interface interface-name>
<detail interface interface-name>
< interface interface-name>
< interface interface-name (brief | detail)>
```

Description

(MX5, MX10, MX40, MX80, MX80-T, MX240, MX480, MX960, MX2010, MX2020, and PTX Series routers only) Display the Synchronous Ethernet ESMC transmit interface details.

Required Privilege Level

maintenance

Output Fields

Table 1601 on page 26677 lists the output fields for the `show synchronous-ethernet esmc transmit detail` command. Output fields are listed in the approximate order in which they appear.

Table 1601: show synchronous-ethernet esmc transmit detail Output Fields

Field Name	Field Description
`Interface name`	`interface-slot/pic/port`—Displays the name of the interface for which the ESMC transmit details are displayed.
`Status`	`string`—Displays the ESMC transmit interface status details.

Sample Output

command-name

```
user@host# show synchronous-ethernet esmc transmit detail
ESMC Transmit interface details:
    Interface name: ge-1/0/4          Status: ESMC Tx (QL SSU-A/SSM 0x4)
    Interface name: ge-1/0/2          Status: ESMC Tx (QL DNU/SSM 0xf)
    Interface name: ge-1/2/4          Status: ESMC Tx (QL SSU-A/SSM 0x4)
```

show synchronous-ethernet esmc transmit

```
user@host> show synchronous-ethernet esmc transmit
ESMC Transmit interfaces:
    xe-2/0/10
```

show synchronous-ethernet esmc transmit (PTX)

command-name

```
user@host> show synchronous-ethernet esmc transmit
ESMC Transmit interfaces:
    et-1/1/0:0
    et-2/0/22:3
```

Release Information

Command introduced in Junos OS Release 11.2R4.

Command introduced in Junos OS Release 14.2R1 for PTX Series Routers.

RELATED DOCUMENTATION

Synchronous Ethernet Overview

Configuring an External Clock Synchronization Interface for MX Series Routers

request chassis synchronization mode | **14570**

request chassis synchronization switch | 14574
show synchronous-ethernet global-information | 26679
show synchronous-ethernet esmc statistics | 26674

show synchronous-ethernet global-information

IN THIS SECTION

- Syntax | 26679
- Description | 26679
- Required Privilege Level | 26679
- Output Fields | 26679
- Sample Output | 26680
- Release Information | 26681

Syntax

```
show synchronous-ethernet global-information
<brief>
```

Description

(MX5, MX10, MX40, MX80, MX80-T, MX240, MX480, MX960, and PTX routers only) Display information about the global configuration for Synchronous Ethernet chassis synchronization.

Required Privilege Level

maintenance

Output Fields

Table 1602 on page 26680 lists the output fields for the `show synchronous-ethernet global-information` command. Output fields are listed in the approximate order in which they appear.

Table 1602: show synchronous-ethernet global-information Output Fields

Field Name	Field Description
`Network option`	`(option-1(EEC1) | option-2(EEC2))`—Displays the network option configuration, either option-1(EEC1) or option-2(EEC2).
`Clock mode`	`(free-run | auto-select)`—Displays the configured mode of operation. The clock source can be either from the free-run local oscillator or from an external qualified clock. The default is auto-select mode.
`QL mode`	`(enable | disable)`—Displays the configured quality level mode configuration. The default is disable.
`Switchover mode`	`(revertive | non-revertive)`—Displays the configured synchronization clock switching mode. The default mode is revertive.
`Config change holdover`	`seconds`—Displays the time interval to wait before selecting the new clock source during a configuration change. The default value is 30 seconds.
`Switchover holdover`	`seconds`—Displays the time interval to wait before selecting the new clock source during switchover. The default value is 30 seconds.
`Reboot holdover`	`seconds`—Displays the time interval to wait before selecting the new clock source during reboot. The default value is 120 seconds.

Sample Output

show synchronous-ethernet global-information (MX)

```
user@host# show synchronous-ethernet global-information
Global Configuration:

Network option         : option-1(EEC1)
Clock mode             : Auto-select
Max transmit quality   : PRC
QL mode                : Enabled
```

```
Clock selection mode   : Receive-QL based
Switchover mode        : Revertive
Config change holdover : 15  seconds
Switchover holdover    : 30  seconds
Reboot holdover        : 60  seconds
RE Status              : Master
Global Wait to Restore : 1   min
```

show synchronous-ethernet global-information (PTX)

```
user@host# show synchronous-ethernet global-information
Global Configuration:

Network option         : option-1(EEC1)
Clock mode             : Auto-select
Max transmit quality   : PRC
QL mode                : Disabled
Clock selection mode   : Config-QL based
Switchover mode        : Revertive
Config change holdover : 15  seconds
Switchover holdover    : 30  seconds
Reboot holdover        : 120 seconds
RE Status              : Master
Global Wait to Restore : 0   min
```

Release Information

Command introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

Synchronous Ethernet Overview

Configuring an External Clock Synchronization Interface for MX Series Routers

request chassis synchronization mode | 14570

show synchronous-ethernet esmc statistics | 26674

show synchronous-ethernet esmc transmit | 26676

show system alarms

IN THIS SECTION

- Syntax | 26682
- Description | 26682
- Options | 26682
- Additional Information | 26682
- Required Privilege Level | 26682
- Sample Output | 26683
- Release Information | 26683

Syntax

```
show system alarms
```

Description

Display active system alarms.

Options

This command has no options.

Additional Information

System alarms are preset. You cannot modify them, although you can configure them to appear automatically in the J-Web user interface or CLI. They include a `configuration` alarm that appears when no rescue configuration alarm is set and a `license` alarm that appears when a software feature is configured but no valid license is configured for the feature.

Required Privilege Level

admin

Sample Output

show system alarms

```
user@host> show system alarms
5 alarms currently active
Alarm time               Class  Description
2012-05-29 16:47:18 UTC  Major  /var partition usage crossed critical threshold
2012-05-29 16:47:18 UTC  Minor  /var partition usage crossed high threshold
2012-05-29 16:47:18 UTC  Major  /root partition usage crossed critical threshold
2012-05-29 16:47:18 UTC  Minor  /root partition usage crossed high threshold
2012-05-29 16:47:18 UTC  Minor  Rescue configuration is not set
```

Release Information

Command introduced in Junos OS Release 11.1.

show system alarms

IN THIS SECTION

- Syntax | 26684
- Description | 26684
- Options | 26684
- Additional Information | 26684
- Required Privilege Level | 26685
- Output Fields | 26685
- Sample Output | 26686
- Release Information | 26689

Syntax

```
show system alarms
```

Description

Display active system alarms. System alarms are software or operating system software related alarms. System alarms indicate a missing rescue configuration or software license, where valid. System alarms are preset. You cannot modify them, although you can configure them to appear automatically in the J-Web user interface or CLI. They include a *configuration* alarm that appears when no rescue configuration alarm is set and a *license* alarm that appears when a software feature is configured and no valid license is configured for the feature. You can also determine when a license will expire from syslog messages that appear starting from four weeks before expiry of the license. On EX6200 switches, an alarm can be triggered by an internal link error.

Chassis alarms indicate a failure on the device or one of its components. For example, chassis alarms include an alarm for fan tray failure on a device, or if an FPC on a device is unreachable. Chassis alarms are preset and cannot be modified.

Options

This command has no options.

Additional Information

The the logic for multiple feature licenses is based on the highest validity among the licenses. Also, for capacity non-cumulative, exclusive type licenses (such as for scale), the logic is based on the highest validity of the license.

> **NOTE**: As of Junos OS Release 17.3R1, the logic for multiple capacity type licenses and when their expiry raises alarms was changed. Before, the behavior had alarms and syslog messages for expiring licenses raised based on the highest validity, which would mislead users in the case of a license expiring earlier than the highest validity license. The newer behavior has the about-to-expire logic based on the first expiring license.

In Junos OS release 11.1 and later, alarms for fans also show the slot number of the malfunctioning fans in the CLI output.

Starting with Junos OS Release 13.2, you can view degraded fabric alarms on a routing matrix based on TX Matrix Plus router with 3D SIBs. The alarm indicates that the source FPC is running with a degraded fabric condition. This alarm is an early warning of a possible fabric null route condition. When the

degraded fabric alarm is raised on the source FPC, you can take remedial action to avoid a fabric null route condition. The degraded fabric alarm is raised on the source FPC if both the following conditions are met:

- The active Packet Forwarding Engine destinations are reachable on one or no active switching planes.
- At least one of the inactive switching planes has a fault that causes the destination Packet Forwarding Engine to become unreachable.

> **NOTE**: On Junos OS Evolved, the `show system alarms` command does not display the error number. Instead, you can use the `show chassis fpc errors detail` and the `show system errors` commands to list the errors that contribute to a failure.

Required Privilege Level

view

Output Fields

Table 1603 on page 26685 lists the output fields for the `show system alarms` command. Output fields are listed in the approximate order in which they appear.

Table 1603: show system alarms Output Fields

Field Name	Field Description
`Alarm time`	Date and time the alarm was first recorded.
`Class`	Severity class for this alarm: `Minor` or `Major`.
`Description`	Information about the alarm.

Sample Output

show system alarms

```
user@host> show system alarms

2 alarms currently active
Alarm time                    Class     Description
2005-02-24 17:29:34 UTC       Minor     IPsec VPN tunneling usage requires a license
2005-02-24 17:29:34 UTC       Minor     Rescue configuration is not sent
```

show system alarms

```
user@host> show system alarms

5 alarms currently active
Alarm time               Class  Description
2012-05-29 16:47:18 UTC  Major  /var partition usage crossed critical threshold
2012-05-29 16:47:18 UTC  Minor  /var partition usage crossed high threshold
2012-05-29 16:47:18 UTC  Major  /root partition usage crossed critical threshold
2012-05-29 16:47:18 UTC  Minor  /root partition usage crossed high threshold
2012-05-29 16:47:18 UTC  Minor  Rescue configuration is not set
```

show system alarms (Fan Tray)

```
user@host> show system alarms

4 alarms currently active
Alarm time               Class  Description
2010-11-11 20:27:38 UTC  Major  Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC  Minor  Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC  Major  Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC  Major  Side Fan Tray 0 Failure
```

show system alarms (QFX Series and OCX Series)

```
user@switch> show system alarms
2 alarms currently active
Alarm time Class Description
2005-02-24 17:29:34 UTC Minor Rescue configuration is not sent
```

show system alarms (EX6200)

```
user@switch> show system alarms
2 alarms currently active
Alarm time               Class  Description
2013-04-05 16:51:41 PDT  Major  FPC 8 internal link errors detected
2013-04-04 18:05:35 PDT  Minor  Rescue configuration is not set
```

show system alarms (SRX5400, SRX5600, and SRX5800)

This command displays the active alarms due to L2 channel error.

```
user@switch> show system alarms
1 alarm currently active
Alarm time      Class Description
2021-01-16 20:42:40 PST Minor NSD l2 channel error on physical interfaces
```

show system alarms (SRX4700)

This command displays the active alarms indicating a failure in writing the security logs to traffic logs due to disk corruption or a read/write error.

```
user@switch> show system alarms
8 alarms currently active
Alarm time               Class  Description
2024-08-19 07:01:43 UTC  Minor  Failed to write in Traffic-Log
```

show system alarms (TX Matrix Plus router with 3D SIBs)

```
user@router> show system alarms

sfc0-re0:
--------------------------------------------------------------------------
2 alarms currently active
Alarm time               Class  Description
2013-05-08 18:13:58 UTC  Major  LCC 0 Major Errors
2013-05-08 17:48:46 UTC  Major  LCC 7 Major Errors

lcc0-re1:
--------------------------------------------------------------------------
1 alarm currently active
Alarm time               Class  Description
2013-05-08 18:19:24 UTC  Major  FPC 1 degraded fabric condition detected

lcc7-re0:
--------------------------------------------------------------------------
1 alarm currently active
Alarm time               Class  Description
2013-05-08 18:19:24 UTC  Major  FPC 7 degraded fabric condition detected
```

show system alarms (Junos OS Evolved with PCI corrected error)

```
user@router> show system alarms
10 alarms currently active
Alarm time               Class  Description
2019-02-01 02:20:09 PST  Major  PCI Corrected error on dev 0000:00:01.0
```

Starting in Junos OS Evolved Release 19.1R1, the alarm string for PCI Corrected error is shown as `PCI Corrected error on dev 0000:00:01.0`. Also, a PCI uncorrectable error does not cause a reboot, but only raises an alarm.

show system alarms (Junos OS Evolved when FEB is absent)

```
user@router> show system alarms
31 alarms currently active
```

```
Alarm time               Class  Description
2023-01-31 11:25:49 UTC  Major  FEB 1 Absent
2023-01-31 10:09:26 UTC  Major  Chassis No Redundant Power
```

Starting in Junos OS Evolved Release 23.1R1, an alarm will be raised when a RCB is present but its mated FEB is absent.

Release Information

Command introduced before Junos OS Release 7.4.

show system anomalies (GNF)

IN THIS SECTION

- Syntax | 26689
- Description | 26689
- Options | 26690
- Required Privilege Level | 26690
- Output Fields | 26690
- Sample Output | 26691
- Release Information | 26691

Syntax

```
show system anomalies (all-anomalies | config | fru | system)
```

Description

Display the incompatibilities between the software version running on the guest network function (GNF) and the version running on the base system (BSYS).

Options

all-anomalies Display the multiversion software incompatibilities from all categories—system, configuration, and FRU.

config Display the feature incompatibilities between software versions.

fru Display the FRU-level incompatibilities between software versions. This can be an incompatibility pertaining to the support for a specific FRU.

system Display the system-level incompatibilities between software versions. These include interprocess communication (IPC) message, CLI, or SNMP incompatibility.

Required Privilege Level

View

Output Fields

Table 1604 on page 26690 lists the output fields for the `show system anomalies` command. Output fields are listed in the approximate order in which they appear.

Table 1604: show system anomalies gnf-id Output Fields

Field Name	Field Description
`Anomaly Type`	Shows the software incompatibility type. The following are the possible values: • SYS—Indicates system-specific incompatibilities. • FRU—Indicates FRU-specific incompatibilities. • CONFIG—Indicates feature-specific incompatibilities.
`Anomaly ID`	Shows the incompatibility ID.

Table 1604: show system anomalies gnf-id Output Fields *(Continued)*

Field Name	Field Description
Default Action	Shows the default actions associated with incompatibilities. The following are the possible values: • WARN—Indicates the presence of a minor incompatibility. It causes a warning. • ABORT—Indicates the presence of a major incompatibility. It causes a termination action.
Class	Indicates of the incompatibility is present in GNF, or BSYS, or both.
Message	Incompatibility description.
FRU ID	Field-replaceable unit (FRU) ID. Applicable in the case of FRU-specific incompatibilities.
Firmware	Firmware of the FRU.

Sample Output

command-name

```
user@router> show system anomalies

Anomaly Type    Anomaly ID   Default Action     Class
Message                               FRU ID    Firmware
-------------------------------------------------------------------------------------------
--------------------------------------------------
SYS              100         WARN               GNF Present   system incompatibility 1
SYS              103         WARN               BOTH Present  system incompatibility 2
SYS              200         WARN               BSYS Present  system incompatibility 3
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

show system anomalies gnf-id | 26692

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | 15481

request server authenticate-peer-server | 14901

show system anomalies gnf-id

IN THIS SECTION

- Syntax | 26692
- Description | 26692
- Options | 26692
- Required Privilege Level | 26693
- Output Fields | 26693
- Sample Output | 26694
- Release Information | 26694

Syntax

```
show system anomalies gnf-id id (all-anomalies | config | fru | system)
```

Description

Display incompatibilities between the software version running on the base system (BSYS) and the software running on a specific guest network function (GNF).

Options

gnf-id *id* Specify the GNF ID for which you want to view the software incompatibilities.

all-anomalies Display the multiversion software incompatibilities from all categories—system, configuration and FRU.

config Display feature-level incompatibilities between software versions.

fru Display the FRU-level incompatibilities between software versions. This can be an incompatibility pertaining to the support for a specific FRU.

system Display the system-level incompatibilities between software versions. These include IPC incompatibility, CLI or SNMP incompatibility.

Required Privilege Level

View

Output Fields

Table 1605 on page 26693 lists the output fields for the `show system anomalies gnf-id` command. Output fields are listed in the approximate order in which they appear.

Table 1605: show system anomalies gnf-id Output Fields

Field Name	Field Description
`Anomaly Type`	Shows the software incompatibility type. The following are the possible values: • SYS—Indicates system-specific incompatibilities. • FRU—Indicates FRU-specific incompatibilities. • CONFIG—Indicates feature-specific incompatibilities.
`Anomaly ID`	Shows the incompatibility ID.
`Default Action`	Shows the default actions associated with incompatibilities. The following are the possible values: • WARN—Indicates the presence of a minor incompatibility. It causes a warning. • ABORT—Indicates the presence of a major incompatibility. It causes a termination action.
`Class`	Indicates of the incompatibility is present in GNF, or BSYS, or both.

Table 1605: show system anomalies gnf-id Output Fields *(Continued)*

Field Name	Field Description
Message	Incompatibility description.
FRU ID	Field-replaceable unit (FRU) ID. Applicable in the case of FRU-specific incompatibilities.

Table 1605: show system anomalies gnf-id Output Fields *(Continued)*

Firmware	Firmware of the FRU.

Sample Output

command-name

```
user@router> show system anomalies gnf-id 1

Anomaly Type    Anomaly ID   Default Action     Class
Message                               FRU ID    Firmware
---------------------------------------------------------------------------------------------
---------------------------------------------------
SYS              100         WARN               GNF Present   system incompatibility 1
SYS              103         WARN               BOTH Present  system incompatibility 2
SYS              200         WARN               BSYS Present  system incompatibility 3
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

show system anomalies (GNF) | **26689**

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | **15481**

request server authenticate-peer-server | **14901**

show system applications (Junos OS Evolved)

IN THIS SECTION

- Syntax | 26695
- Description | 26695
- Options | 26695
- Required Privilege Level | 26696
- Output Fields | 26696
- Sample Output | 26702
- Release Information | 26709

Syntax

```
show system applications
<app app-name>
<brief>
<detail>
<error>
<node node-name>
```

Description

This command displays application summary information in one of the following forms:

- Shows all application summary information for all nodes.
- Shows the application summary information for a specific application.
- Shows the application summary information for a specific node.

Options

app *app-name* (Optional) Specify application name for which you want to display application summary information.

brief	(Optional) Display brief output. This is the default format of display.
detail	(Optional) Display detailed output.
error	(Optional) Display information about errors. You can specify the `app`, `node`, and `detail` options at the same time to further refine the list of errors, and to get more information about each one.
node *node-name*	(Optional) Specify node name for which you want to display application summary information.

Required Privilege Level

view

Output Fields

For a description of the output fields, see Table 1606 on page 26696. Output fields are listed in the approximate order in which they appear.

Table 1606: show system applications Output Fields

Field Name	Description		Level
`Applications Information`	**Application**	Name of the application.	all options, except `error`
	Node	Name of the node the application is running on.	
	App State	State of the application: online, offline, failed, or active.	
	App Weight	A relative weight for multiple instances of the application across multiple nodes. The application instance with the higher weight provides more functionality.	
	App Zookeeper Session	Zookeeper session ID.	

Table 1606: show system applications Output Fields *(Continued)*

Field Name	Description	Level
Object Producer details	**Producer ID** Identifies which production set the object is part of. **Epoch ID** A number that identifies the current process that owns a production set. There can only be one owning process (active producer) that owns a production set at one time. The current owning process has an Epoch ID that is larger than any previous producer. **Production Topic** Hierarchical string that represents the production set. **Producer State** active or standby: • active indicates the application has production rights to modify the state in the production set. • standy means that the application is waiting to get the production right for the production set.	all options, except error
Description	A short description of the application, it also lists the systemd service file used for the application.	detail, without the error option
Loaded	A systemd state that indicates if the application is loaded in the system or not.	detail, without the error option
Run State from OS	A systemd state that indicates if the application is active or not.	detail, without the error option
Main PID	Process identifier (PID) of the application.	detail, without the error option
Command	Command line to launch the application.	detail, without the error option

Table 1606: show system applications Output Fields *(Continued)*

Field Name	Description	Level
`ID`	Name of the application.	`detail`, without the `error` option
`Meta`	Meta data for the application includes the following fields: **`Bin path`** Path to application execution. **`Log file`** Where logs go. **`Working Dir`** Working directory. **`Production Set`** Global or local production set. Values might be `shared` or `local`.	`detail`, without the `error` option

Table 1606: show system applications Output Fields *(Continued)*

Field Name	Description	Level
`Resource`	Resource data for the application includes the following fields: **`all nodes`** Does the application run on all nodes, `true` or `false`. **`Max instances`** How many instances of the application are there. **`Max instances per node`** How many instances of the application per node are there. **`Run on startup`** Is the application launched at bootup, `true` or `false`. **`Node attributes`** Typical node attributes are RE, FPC, MasterRE. You can see the node attributes by using the `show system node-attributes` command. **`Node attribute match`** What is the node attribute required to launch this application on a node? For example, if this field has the output `re`, `Service file: lfmd`, it indicates that the process `lfmd` will be launched on a node that has the attribute RE.	`detail`, without the `error` option
`Failure`	Failure data for the application includes the following fields: **`Alarm color`** Which alarm to be raised on failure, or `none`. **`Alarm ID`** The alarm ID. **`Restart`** Whether to restart the application, `true` or `false`.	`detail`, without the `error` option

Table 1606: show system applications Output Fields *(Continued)*

Field Name	Description	Level
`Upgrade`	Upgrade data for the application includes the following fields: **`Upgrade parallely`** Options are `true` or `false`. **`Node upgrade done notify`** Application requires notification when node upgrade is done. Options are `true` or `false`. **`System upgrade done notify`** Application requires notification when system upgrade is done. Options are `true` or `false`. **`Upgrade restart node`** Options are `true` or `false`. **`Upgrade style`** Option is `stop-start`. **`hw sync`** Application requires notification to sync hardware state during a unified in-service software upgrade (unified ISSU). Options are `true` or `false`. **`wait for sw sync`** Application provides notification after it has synced software state during a unified in-service software upgrade (unified ISSU). Options are `true` or `false`. **`prepare to upgrade final`** Application requires notification to prepare for upgrade. Options are `true` or `false`. **`exits on prep final`** Application exits after processing prepare to upgrade. Options are `true` or `false`. **`Weight`** Application weight.	`detail`, without the `error` option

Table 1606: show system applications Output Fields *(Continued)*

Field Name	Description	Level
App-Exit	App-Exit data for the application includes the following fields: **Restart Supported** True/false. When the application exits, should the application be restarted. **Restart Node** True/false. When the application exits, should the node be rebooted. **Mark node spare** When an application exits, should the node be marked spare.	detail, without the error option
Node	Name of the node the application is running on.	error option, all combinations
Application	Name of the application.	error option, all combinations
Error Count	How many errors are on the node for the stated application.	error detail, either with or without the app or node options
Error Guids	Identifiers for particular error instances.	error detail, either with or without the app or node options
Error Number	Identifier for a particular type of error. For example, number 3 is for the error "No such process".	error detail, either with or without the app or node options
Severity	How severe the error is.	error detail, either with or without the app or node options

Table 1606: show system applications Output Fields *(Continued)*

Field Name	Description	Level
GUID	Identifier for a particular error instance.	error detail, either with or without the app or node options
Error Description	Text describing the error.	error detail, either with or without the app or node options
Error Module	The module causing the error.	error detail, either with or without the app or node options
Error Object Name	Name of the object causing the error.	error detail, either with or without the app or node options
Error Object Parent Name	Name of the object's parent.	error detail, either with or without the app or node options
Error Timestamp	Date and time at which the error occurred.	error detail, either with or without the app or node options

Sample Output

show system applications

```
user@host> show system applications
Applications Information:
Application              : bcmd_evo
Node                     : fpc0
App State               : offline
Object Producer details
```

```
Producer ID             : 0
  Epoch ID                : 0

Applications Information:
Application             : ccdpfe
Node                    : fpc0
App State              : online
Object Producer details
Producer ID             : 576
  Epoch ID                : 65
  Production Topic        : /Root/fpc0/ccdpfe/100143551468101228
  Producer State          : active

Applications Information:
Application             : cmdd
Node                    : fpc0
App State              : online
Object Producer details
Producer ID             : 570
  Epoch ID                : 66
  Production Topic        : /Root/fpc0/cmdd/1099227235289688912
  Producer State          : active

...

Applications Information:
Application             : alarm-mgmtd
Node                    : re0
App State              : online
Object Producer details
Producer ID             : 26
  Epoch ID                : 1
  Production Topic        : /Root/alarm-mgmtd/2988563069668674039
  Producer State          : active

Applications Information:
Application             : alarmd
Node                    : re0
App State              : online
Object Producer details
Producer ID             : 377
  Epoch ID                : 30
  Production Topic        : /Root/alarmd/6512784671716237713
```

```
  Producer State          : active

Applications Information:
Application             : arpd
Node                    : re0
App State              : online
Object Producer details
Producer ID             : 396
  Epoch ID                : 41
  Production Topic        : /Root/arpd/14284058728950342139
  Producer State          : active

...

Applications Information:
Application             : alarm-mgmtd
Node                    : re1
App State              : online
Object Producer details
Producer ID             : 26
  Epoch ID                : 0
  Production Topic        : /Root/alarm-mgmtd/29885630696686740 39
  Producer State          : standby

Applications Information:
Application             : bcmd_evo
Node                    : re1
App State              : offline
Object Producer details
Producer ID             : 0
  Epoch ID                : 0

Applications Information:
Application             : charonctl
Node                    : re1
App State              : online
Object Producer details
Producer ID             : 25
  Epoch ID                : 4
  Production Topic        : /Root/re1/charonctl/10854553120394604032
  Producer State          : active
```

```
...
```

show system applications error

```
{master}
user@host-re0> show system applications error
Error Summary:
Node                : fpc0
Application         : evo-aftmand-bt
Error Count         : 3
Error Guids:
    1249835833539
    1249835833540
    1249835833541

Error Summary:
Node                : re0
Application         : rpdagent
Error Count         : 3
Error Guids:
    1249835833539
    1249835833540
    1249835833541
```

show system applications error node node-name

```
user@host-re0> show system applications error node fpc0
Error Summary:
Node                : fpc0
Application         : evo-aftmand-bt
Error Count         : 3
Error Guids:
    1249835833539
    1249835833540
    1249835833541
```

show system applications error app application-name

```
user@host-re0> show system applications error app rpdagent
Error Summary:
Node                : re0
Application         : rpdagent
Error Count         : 3
Error Guids:
    1249835833539
    1249835833540
    1249835833541
```

show system applications app application-name node node-name

```
user@host> show system applications app alarm-mgmtd node re1
Applications Information:
Application             : alarm-mgmtd
Node                    : re1
App State              : online
Object Producer details
Producer ID             : 26
  Epoch ID                : 0
  Production Topic        : /Root/alarm-mgmtd/2988563069668674039
  Producer State          : standby
```

show system applications app application-name detail

```
user@host> show system applications app cmdd detail
Applications Information:
Application             : cmdd
Node                    : re0
App State               : online ready
App Weight              : 1
App Zookeeper Session   : 1000000934d000d
Object Producer details
Producer ID             : 50331736
  Epoch ID                : 47
  Production Topic        : /Root/re0/cmdd/3158206796014561683
  Producer State          : active
```

```
Description             : cmdd.service - "Command Daemon"
Loaded                  : loaded (/etc/systemd/system/cmdd.service;static;vendor preset:enabled)
Run State from OS       : active (running) (Result: success) since Mon 2018-10-29 05:02:24 PDT
Main PID                : 5814
Command                 : /usr/sbin/cmdd --app-name cmdd -I object_select --shared-objects-mode 3
App Config Info
 ID                     : cmdd
 Meta
  Bin path              : /usr/sbin/cmdd
  Log file              : /var/log
  Working Dir           : /usr/sbin
  Production Set        : local
  Sysman Managed        : true
  Type Evo              : true
 Resource
  All nodes             : true
  Max instances         : 1
  Max instances per node: 1
  App Suite             : default,diags_default
  Run on startup        : true
  Node attributes       :
(Node attribute match  : *, Service file  : cmdd)
 Failure
  Alarm color           : red
  Restart Node          : false
  Mark node spare       : false
 Upgrade
  Upgrade parallely           : true
  Upgrade restart node        : false
  Upgrade style               : stop-start
 App-Exit
  Restart Supported           : true
  Restart Node                : false
  Mark node spare             : false
```

show system applications error app application-name detail

```
user@host-re0> show system applications error app rpdagent detail
Error detail:
Node                            : re0
Application                     : rpdagent
```

```
Error Number                    : 3
Severity                        : FATAL
GUID                            : 1249835833539
Error Description               : No such process
Error Module                    : CHASSIS0:FPC0
Error Object Name               : proto: ipv4 prefix: 10.35.1.1
Error Object Parent Name        : Type:RtTable : default.inet
Error Timestamp                 : Thu Aug 18 03:54:34 2022

Node                            : re0
Application                     : rpdagent
Error Number                    : 3
Severity                        : FATAL
GUID                            : 1249835833540
Error Description               : No such process
Error Module                    : CHASSIS0:FPC0
Error Object Name               : proto: ipv4 prefix: 10.35.1.2
Error Object Parent Name        : Type:RtTable : default.inet
Error Timestamp                 : Thu Aug 18 03:54:34 2022

Node                            : re0
Application                     : rpdagent
Error Number                    : 3
Severity                        : FATAL
GUID                            : 1249835833541
Error Description               : No such process
Error Module                    : CHASSIS0:FPC0
Error Object Name               : proto: ipv4 prefix: 10.35.1.3
Error Object Parent Name        : Type:RtTable : default.inet
Error Timestamp                 : Thu Aug 18 03:54:34 2022
```

show system applications error app application-name node node-name detail

```
{master}
user@host-re0> show system applications error app evo-aftmand-bt node fpc1 detail
Error detail:
Node                            : fpc1
Application                     : evo-aftmand-bt
Error Number                    : 1001
Severity                        : FATAL
GUID                            : 1151051261231
```

```
Error Description              : Generic failure
Error Module                   : FPC:1/pfe/rt
Error Object Name              : proto: ipv4 prefix: 10.32.1.1
Error Object Parent Name       : Type:RtTable : default.inet
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

`error` option added in Junos OS Evolved Release 22.3R1.

RELATED DOCUMENTATION

- request node halt (Junos OS Evolved) | **14708**
- request node (offline | online) (Junos OS Evolved) | **14711**
- request node power-off (Junos OS Evolved) | **14713**
- request node power-on (Junos OS Evolved) | **14715**
- request node reboot (re0 | re1) (Junos OS Evolved) | **14717**

show system audit

IN THIS SECTION

- Syntax | **26710**
- Syntax (EX Series Switch and MX Series Router) | **26710**
- Syntax (TX Matrix Router) | **26710**
- Syntax (TX Matrix Plus Router) | **26710**
- Syntax (QFX Series) | **26711**
- Syntax (OCX Series) | **26711**
- Description | **26711**
- Options | **26711**
- Additional Information | **26713**
- Required Privilege Level | **26713**

Sample Output | 26713
Release Information | 26720

Syntax

```
show system audit
<root-only>
```

Syntax (EX Series Switch and MX Series Router)

```
show system audit
<all-members>
<local>
<member member-id>
<root-only>
```

Syntax (TX Matrix Router)

```
show system audit
<all-lcc | lcc number | scc>
<root-only>
```

Syntax (TX Matrix Plus Router)

```
show system audit
<all-chassis | all-lcc | lcc number | sfc number>
<root-only>
```

Syntax (QFX Series)

```
show system audit
<infrastructure name | interconnect-device name | node-group name | root-only>
```

Syntax (OCX Series)

```
show system audit
<root-only>
```

Description

Display the state and checksum values for file systems.

Options

none Display the state and checksum values for all file systems.

all-chassis (TX Matrix routers and TX Matrix Plus routers only) (Optional) Display file system MD5 hash and permissions information for all of the chassis.

all-lcc (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display file system MD5 hash and permissions information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display file system MD5 hash and permissions information for all T1600 or T4000 routers connected to the TX Matrix Plus router.

all-members (EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on all members of the Virtual Chassis configuration.

lcc *number* (TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display file system MD5 hash and permissions information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display file system MD5 hash and permissions information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

infrastructure *name*
(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for a fabric control Routing Engine or a fabric control Routing Engine.

interconnect-device *name*
(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for the Interconnect device.

local
(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on the local Virtual Chassis member.

member *member-id*
(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

node-group *name*
(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for the Node group

root-only
(Optional) Check only the root (/) file system. On a QFabric system, you can check the root (/) file system on the infrastructure (fabric manager Routing Engine and fabric control Routing Engine), Interconnect device, or Node group.

scc
(TX Matrix routers only) (Optional) Display file system MD5 hash and permissions information for the TX Matrix router (or switch-card chassis).

sfc *number*
(TX Matrix Plus routers only) (Optional) Display file system MD5 hash and permissions information for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with **0**.

Additional Information

To redirect the output to a file, issue the following command:

```
ssh device-name  'show system audit root-only' > output-file
```

If you save the output of the `show system audit root-only` command to a file, you can compare it to subsequent output from the command to determine whether anything has changed.

By default, when you issue the `show system audit` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

admin

Sample Output

show system audit root-only

```
user@host> show system audit root-only
#          user: root
#       machine: host
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.               type=dir nlink=23 size=1024 time=950252640.0
  .cshrc      uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
                md5digest=f414e06fea6bd646244b98e13d6e6226
  .kernel.jkernel.backup \
                mode=0744 size=1934552 time=944688902.0 \
                md5digest=2c343cf0bd9fea8f04f78604feed7aa4
  .profile    uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
                md5digest=55a1e3c6c67789c9d3a1cce1ea39f670
    COPYRIGHT   uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \
                md5digest=7df8bc77dcee71382ea73eb0ec6a9243
```

```
    boot.config mode=0644 size=3 time=945902618.0 \
                md5digest=93d722493ed38477338a1405d7dcbb40
    boot.help   uid=3 gid=7 mode=0444 size=411 time=939182876.0 \
                md5digest=9b7126385734bcae753f4179ab59d8e5
    compat      type=link mode=0777 size=11 time=915149058.0 \
                link=/usr/compat
    kernel      mode=0444 size=1947607 time=950230892.0 \
                md5digest=1a2a8aff2fec678a918ba0d6bf063980
    kernel.avr  uid=1112 size=1947642 time=950252597.0 \
                md5digest=82e1637682d58ec28964dfee7fccb62e
    kernel.config \
                mode=0644 size=0 time=915149058.0 \
                md5digest=d41d8cd98f00b204e9800998ecf8427e
    sys         type=link mode=0777 size=11 time=915149029.0 \
                link=usr/src/sys
```

show system audit lcc (TX Matrix Router)

```
user@host> show system audit lcc 2
lcc2-re0:
--------------------------------------------------------------------------
#          user: root
#       machine: test-lcc2
#          tree: /
#          date: Mon Sep 13 11:55:33 2004

# .
/set type=file uid=0 gid=0 mode=0555 nlink=1 flags=none
.               type=dir nlink=20 size=512 time=1094982121.0
    COPYRIGHT   mode=0644 size=4735 time=986012708.0 \
                md5digest=78396df1404ad742e6eb1be28f0cd63b
    kernel      type=link mode=0700 size=17 time=1090266262.0 \
                link=/packages/jkernel

# ./altconfig
altconfig       type=dir nlink=2 size=512 time=1089801320.0
# ./altconfig
..


# ./altroot
```

```
altroot         type=dir nlink=2 size=512 time=1089801320.0
# ./altroot
..

# ./b
b               type=dir mode=0755 nlink=2 size=512 time=1093961429.0
# ./b
..


# ./bin
/set type=file uid=0 gid=0 mode=0700 nlink=1 flags=none
bin             type=dir mode=0755 nlink=2 size=512 time=1089843059.0
    [           type=link size=28 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/test
    cat         type=link size=27 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/cat
    chmod       type=link size=29 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/chmod
    cp          type=link size=26 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/cp
    csh         type=link size=27 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/csh
    date        type=link size=28 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/date
    dd          type=link size=26 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/dd
    df          type=link size=26 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/df
    echo        type=link size=28 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/echo
    ed          type=link size=26 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/ed
    expr        type=link size=28 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/expr
    hostname    type=link size=32 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/hostname
    kill        type=link size=28 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/kill
    ln          type=link size=26 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/ln
    ls          type=link size=26 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/ls
```

```
    mkdir       type=link size=29 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/mkdir
    mv          type=link size=26 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/mv
    ps          type=link size=26 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/ps
    pwd         type=link size=27 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/pwd
    rcp         type=link size=27 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/rcp
    red         type=link size=26 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/ed
    rm          type=link size=26 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/rm
    rmdir       type=link size=29 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/rmdir
    sh          type=link size=26 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/sh
    sleep       type=link size=29 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/sleep
    stty        type=link size=28 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/stty
    sync        type=link size=28 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/sync
    tcsh        type=link size=27 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/csh
    test        type=link size=28 time=1090266270.0 \
                link=/packages/mnt/jbase/bin/test
# ./bin
..


# ./boot
/set type=file uid=0 gid=0 mode=0444 nlink=1 flags=none
boot            type=dir mode=0555 nlink=3 size=512 time=1095069935.0
    boot0       size=512 time=1094978286.0 \
                md5digest=6f780822dd4ae482a20462b66e542cca
    boot1       mode=0555 size=512 time=1094978294.0 \
                md5digest=8d112b09df342cd0b60fdb9bdcde8e07
    boot2       mode=0555 size=7680 time=1094978294.0 \
                md5digest=28eb58c4068c6b85717e1484f9e028e4
    cdboot      mode=0555 size=165888 time=1094978298.0 \
                md5digest=1474c6b800dfc82ba552d7c36116d07d
```

```
    kgzldr.o    size=5996 time=1094982121.0 \
                md5digest=c53dc948eb07e2ea4eb0413e4c4634a3
    loader      mode=0555 size=163840 time=1094978298.0 \
                md5digest=82d9dc2d31033476bfb61bb7264c4fed
    loader.4th  size=9237 time=986013631.0 \
                md5digest=43144391465ad50267d31e0a320be1de
...
```

show system audit lcc (TX Matrix Plus Router)

```
user@host> show system audit all-chassis

sfc0-re0:
--------------------------------------------------------------------------
#          user: root
#       machine: test
#          tree: /
#          date: Mon May 18 00:13:16 2009

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1 flags=none
.               type=dir nlink=23 size=512 time=1242347096.0
    COPYRIGHT   mode=0644 size=6196 time=1168587741.0 \
                md5digest=bbad415e1c29bbedd9b383537100412c
    kernel      type=link size=17 time=1242347011.0 link=/packages/jkernel
    staging     type=link mode=0777 size=8 time=1242346935.0 link=/var/tmp

# ./.snap
.snap           type=dir mode=0775 nlink=2 size=512 time=1242346922.0
# ./.snap
..


# ./altconfig
altconfig       type=dir mode=0500 nlink=2 size=512 time=1242319843.0
# ./altconfig
..


# ./altroot
altroot         type=dir mode=0500 nlink=2 size=512 time=1242319843.0
```

```
# ./altroot
..


# ./bin
bin             type=dir nlink=2 size=512 time=1242346944.0
    \133        type=link size=28 time=1242346942.0 \
                link=/packages/mnt/jbase/bin/test
    cat         type=link size=27 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/cat
    chflags     type=link size=31 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/chflags
    chmod       type=link size=29 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/chmod
    cp          type=link size=26 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/cp
    csh         type=link size=27 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/csh
    date        type=link size=28 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/date
    dd          type=link size=26 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/dd
    df          type=link size=26 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/df
    echo        type=link size=28 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/echo
    ed          type=link size=26 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/ed
    expr        type=link size=28 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/expr
    hostname    type=link size=32 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/hostname
    kill        type=link size=28 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/kill
    ln          type=link size=26 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/ln
    ls          type=link size=26 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/ls
    mkdir       type=link size=29 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/mkdir
    mv          type=link size=26 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/mv
    pax         type=link size=27 time=1242346944.0 \
```

```
                link=/packages/mnt/jbase/bin/pax
    ps          type=link size=26 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/ps
    pwd         type=link size=27 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/pwd
    rcp         type=link size=27 time=1242346942.0 \
                link=/packages/mnt/jbase/bin/rcp
    red         type=link size=26 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/ed
    rm          type=link size=26 time=1242346942.0 \
                link=/packages/mnt/jbase/bin/rm
    rmdir       type=link size=29 time=1242346942.0 \
                link=/packages/mnt/jbase/bin/rmdir
    sh          type=link size=26 time=1242346942.0 \
                link=/packages/mnt/jbase/bin/sh
    sleep       type=link size=29 time=1242346942.0 \
                link=/packages/mnt/jbase/bin/sleep
    stty        type=link size=28 time=1242346942.0 \
                link=/packages/mnt/jbase/bin/stty
    sync        type=link size=28 time=1242346942.0 \
                link=/packages/mnt/jbase/bin/sync
    tcsh        type=link size=27 time=1242346941.0 \
                link=/packages/mnt/jbase/bin/csh
    test        type=link size=28 time=1242346942.0 \
                link=/packages/mnt/jbase/bin/test
# ./bin
...
```

show system audit root-only (QFX3500 Switch)

```
user@switch> show system audit root-only
#          user: root
#       machine: test
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.               type=dir nlink=23 size=1024 time=950252640.0
  .cshrc      uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
                md5digest=f414e06fea6bd646244b98e13d6e6226
```

```
.kernel.jkernel.backup \
              mode=0744 size=1934552 time=944688902.0 \
              md5digest=2c343cf0bd9fea8f04f78604feed7aa4
.profile    uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
              md5digest=55a1e3c6c67789c9d3a1cce1ea39f670
  COPYRIGHT   uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \
              md5digest=7df8bc77dcee71382ea73eb0ec6a9243
  boot.config mode=0644 size=3 time=945902618.0 \
              md5digest=93d722493ed38477338a1405d7dcbb40
  boot.help   uid=3 gid=7 mode=0444 size=411 time=939182876.0 \
              md5digest=9b7126385734bcae753f4179ab59d8e5
  compat      type=link mode=0777 size=11 time=915149058.0 \
              link=/usr/compat
  kernel      mode=0444 size=1947607 time=950230892.0 \
              md5digest=1a2a8aff2fec678a918ba0d6bf063980
  kernel.avr  uid=1112 size=1947642 time=950252597.0 \
              md5digest=82e1637682d58ec28964dfee7fccb62e
  kernel.config \
              mode=0644 size=0 time=915149058.0 \
              md5digest=d41d8cd98f00b204e9800998ecf8427e
  sys         type=link mode=0777 size=11 time=915149029.0 \
              link=usr/src/sys
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

show system auto-snapshot

IN THIS SECTION

- Syntax | **26721**
- Description | **26721**
- Options | **26721**
- Required Privilege Level | **26721**

- Output Fields | 26721
- Sample Output | 26722
- Release Information | 26722

Syntax

```
show system auto-snapshot
```

Description

This command displays automatic snapshot status information.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1607 on page 26722 describes the output fields for the `show system auto-snapshot` command. Output fields are listed in the approximate order in which they appear.

Table 1607: show system auto-snapshot status Output Fields

Field Name	Field Description
`Auto-snapshot configuration`	Status of the configuration: • Enabled—If the system reboots from the alternate partition, the automatic snapshot feature automatically takes a snapshot of the alternate partition and copies it onto the primary partition. • Disabled—The system does not automatically take a snapshot of the alternate partition. You must use the manual snapshot command, `request system snapshot`, to take a snapshot of one partition and copy it onto the other.
`Auto-snapshot state`	Status of the automatic snapshot procedure: • `Completed`—The automatic snapshot procedure has completed copying the alternate partition to the primary partition and the alarm has been cleared. • `Disabled`—The automatic snapshot procedure is inactive. • `In progress`—The automatic snapshot procedure is in progress. It takes about 10 to 15 minutes to complete, depending upon disk size.

Sample Output

show system auto-snapshot

```
user@switch> show system auto-snapshot

Auto-snapshot Configuration: Enabled
Auto-snapshot State: Disabled
```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring Dual-Root Partitions

show system autoinstallation status

IN THIS SECTION

- Syntax | 26723
- Description | 26723
- Options | 26723
- Required Privilege Level | 26723
- Output Fields | 26723
- Sample Output | 26724
- Release Information | 26725

Syntax

```
show system autoinstallation status
```

Description

This command displays autoinstallation status information (ACX Series routers, and EX Series switches only).

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1608 on page 26724 describes the output fields for the `show system autoinstallation status` command. Output fields are listed in the approximate order in which they appear.

Table 1608: show system autoinstallation status Output Fields

Field Name	Field Description

`Autoinstallation status`	Display autoinstallation status information:
	• `Last committed file`—File last committed for autoinstallation configuration.
	• `Configuration server of last committed file`—IP address or URL of the server configured to retrieve configuration information for the last committed configuration file.
	• `Interface`—Interface configured for autoinstallation.
	  • `Name`—Name of the interface.
	  • `State`—Interface state.
	• `Address acquisition`—Display IP address acquired and protocol used for acquisition upon startup.
	  • `Protocol`—Protocol used for acquisition: BOOTP/DHCP or RARP.
	  • `Acquired address`—IP address acquired from the DHCP server.

Sample Output

show system autoinstallation status

```
user@host> show system autoinstallation status
Autoinstallation status:
Master state: Active
Last committed file: None
Configuration server of last committed file: 10.0.0.0
Interface:
  Name: ge-0/0/1
  State: None
  Address acquisition:
    Protocol: DHCP Client
    Acquired address: None
    Protocol: RARP Client
    Acquired address: None
```

Release Information

Command introduced before Junos OS Release 7.4.

Command supported in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

ACX Series Autoinstallation Overview

Before You Begin Autoinstallation on an ACX Series Universal Metro Router

Autoinstallation Configuration of ACX Series Universal Metro Routers

USB Autoinstallation on ACX Series Routers

Verifying Autoinstallation on ACX Series Universal Metro Routers

autoinstallation | 2753

show system autorecovery state

IN THIS SECTION

- Syntax | 26725
- Description | 26726
- Required Privilege Level | 26726
- Output Fields | 26726
- Sample Output | 26727
- Release Information | 26727

Syntax

```
show system autorecovery state
```

Description

This command perform checks and show status of all autorecovered items.

In devices running FreeBSD Release 12 or later, you cannot back up data with the autorecovery feature. Instead, back up data using snapshots. To learn if your device is running FreeBSD Release 12 or later, issue the `show version` command and look for the `fbsd_builder_stable` string in the module names. If the string includes the number 12 or later, your device is running FreeBSD Release 12 or later.

Required Privilege Level

view

Output Fields

Table 1609 on page 26726 lists the output fields for the `show system autorecovery state` command. Output fields are listed in the approximate order in which they appear.

Table 1609: show system autorecovery state Output Fields

Field Name	Field Description
`File`	The name of the file on which autorecovery checks are performed.
`Slice`	The disk partition on which autorecovery checks are performed.
`Recovery Information`	Indicates whether autorecovery information for the file or slice has been saved.
`Integrity Check`	Displays the status of the file's integrity check (passed or failed).
`Action / Status`	Displays the status of the item, or the action required to be taken for that item.

Sample Output

show system autorecovery state

```
user@host> show system autorecovery state

Configuration:
File             Recovery Information    Integrity Check    Action / Status
rescue.conf.gz   Saved                   Passed             None
Licenses:
File             Recovery Information    Integrity Check    Action / Status
JUNOS282736.lic  Saved                   Passed             None
JUNOS282737.lic  Not Saved               Not checked        Requires save
BSD Labels:
Slice            Recovery Information    Integrity Check    Action / Status
s1               Saved                   Passed             None
s2               Saved                   Passed             None
s3               Saved                   Passed             None
s4               Saved                   Passed             None
```

Release Information

Command introduced in Junos OS Release 15.1X49-D35.

RELATED DOCUMENTATION

request system autorecovery state | 15044

show system boot console

IN THIS SECTION

- Syntax | 26728
- Description | 26728
- Required Privilege Level | 26728

Sample Output | 26728
Release Information | 26728

Syntax

```
show system boot-console
```

Description

Displays boot console information.

Required Privilege Level

view

Sample Output

show system boot-console

```
user@host> show system boot-console
fpc0:
--------------------------------------------------------------------------
Boot console for current boot was Console port
Boot console for next boot is Auxiliary port
```

Release Information

Command introduced in Junos OS Release 21.3R1.

show system boot-messages

IN THIS SECTION

- Syntax | 26729
- Syntax (EX Series Switches) | 26729
- Syntax (TX Matrix Router) | 26730
- Syntax (TX Matrix Plus Router) | 26730
- Syntax (MX Series Router) | 26730
- Syntax (QFX Series) | 26730
- Description | 26730
- Options | 26730
- Additional Information | 26732
- Required Privilege Level | 26732
- Sample Output | 26732
- Release Information | 26740

Syntax

```
show system boot-messages
```

Syntax (EX Series Switches)

```
show system boot-messages
<all-members>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system boot-messages
<all-chassis | all-lcc | lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
show system boot-messages
<all-chassis | all-lcc | lcc number | sfc number>
```

Syntax (MX Series Router)

```
show system boot-messages
<all-members>
<local>
<member member-id>
```

Syntax (QFX Series)

```
show system boot-messages
infrastructure name | interconnect-device name | node-group name
```

Description

This command displays initial messages generated by the system kernel upon startup. These messages are the contents of **/var/run/dmesg.boot**.

Options

none Display all boot time messages.

all-chassis (TX Matrix routers and TX Matrix Plus routers only) (Optional) Display boot time messages for all of the chassis.

all-lcc

(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for all T640 routers connected to a TX Matrix router. On a TX Matrix Plus router, display boot time messages for all connected T1600 or T4000 LCCs.

all-members

(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on all members of the Virtual Chassis configuration.

infrastructure *name*

(QFabric systems only) (Optional) Display boot time messages on the fabric control Routing Engine or fabric manager Routing engines.

interconnect-device *name*

(QFabric systems only) (Optional) Display boot time messages on the Interconnect device.

lcc *number*

(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for a specific T640 router connected to a TX Matrix router. On a TX Matrix Plus router, display boot time messages for a specific router connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on the local Virtual Chassis member.

member *member-id*

(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name*

(QFabric systems only) (Optional) Display boot time messages on the Node group.

scc

(TX Matrix routers only) (Optional) Display boot time messages for the TX Matrix router (or switch-card chassis).

sfc *number* (TX Matrix Plus routers only) (Optional) Display boot time messages for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system boot-messages` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system boot-messages (TX Matrix Router)

```
user@host> show system boot-messages
Copyright (c) 1992-1998 FreeBSD Inc.
Copyright (c) 1996-2000 Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1982, 1986, 1989, 1991, 1993
        The Regents of the University of California.  All rights reserved.

JUNOS 4.1-20000216-Zf8469 #0: 2000-02-16 12:57:28 UTC
    builder@device1.example.com:/p/build/20000216-0905/4.1/release_kernel/sys/compil
e/GENERIC
CPU: Pentium Pro (332.55-MHz 686-class CPU)
  Origin = "GenuineIntel"  Id = 0x66a  Stepping=10
  Features=0x183f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,<b
16>,<b17>,MMX,<b24>>
Teknor CPU Card Recognized
real memory  = 805306368 (786432K bytes)
avail memory = 786280448 (767852K bytes)
Probing for devices on PCI bus 0:
chip0 <generic PCI bridge (vendor=8086 device=7192 subclass=0)> rev 3 class 6000
0 on pci0:0:0
chip1 <Intel 82371AB PCI-ISA bridge> rev 1 class 60100 on pci0:7:0
```

```
chip2 <Intel 82371AB IDE interface> rev 1 class 10180 on pci0:7:1
chip3 <Intel 82371AB USB interface> rev 1 class c0300 int d irq 11 on pci0:7:2
smb0 <Intel 82371AB SMB controller> rev 1 class 68000 on pci0:7:3
pcic0 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int a irq 15 on pci0:13
:0
TI1131 PCI Config Reg: [pci only][FUNC0 pci int]
pcic1 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int b irq 12 on pci0:13
:1
TI1131 PCI Config Reg: [pci only][FUNC1 pci int]
fxp0 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 12 on
 pci0:16:0
chip4 <generic PCI bridge (vendor=1011 device=0022 subclass=4)> rev 4 class 6040
0 on pci0:17:0
fxp1 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on
 pci0:19:0
Probing for devices on PCI bus 1:
mcs0 <Miscellaneous Control Subsystem> rev 12 class ff0000 int a irq 12 on pci1:
13:0
fxp2 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on
 pci1:14:0
Probing for devices on the ISA bus:
sc0 at 0x60-0x6f irq 1 on motherboard
sc0: EGA color <16 virtual consoles, flags=0x0>
ed0 not found at 0x300
ed1 not found at 0x280
ed2 not found at 0x340
psm0 not found at 0x60
sio0 at 0x3f8-0x3ff irq 4 flags 0x20010 on isa
sio0: type 16550A, console
sio1 at 0x3e8-0x3ef irq 5 flags 0x20000 on isa
sio1: type 16550A
sio2 at 0x2f8-0x2ff irq 3 flags 0x20000 on isa
sio2: type 16550A
pcic0 at 0x3e0-0x3e1 on isa
PC-Card ctlr(0) TI PCI-1131 [CardBus bridge mode] (5 mem & 2 I/O windows)
pcic0: slot 0 controller I/O address 0x3e0
npx0 flags 0x1 on motherboard
npx0: INT 16 interface
fdc0: direction bit not set
fdc0: cmd 3 failed at out byte 1 of 3
fdc0 not found at 0x3f0
wdc0 at 0x1f0-0x1f7 irq 14 on isa
wdc0: unit 0 (wd0): <SunDisk SQFXB-80>, single-sector-i/o
```

```
wd0: 76MB (156672 sectors), 612 cyls, 8 heads, 32 S/T, 512 B/S
wdc0: unit 1 (wd1): <IBM-DCXA-210000>
wd1: 8063MB (16514064 sectors), 16383 cyls, 16 heads, 63 S/T, 512 B/S
wdc1 not found at 0x170
wdc2 not found at 0x180
ep0 not found at 0x300
fxp0: Ethernet address 00:a0:a5:12:05:5a
fxp1: Ethernet address 00:a0:a5:12:05:59
fxp2: Ethernet address 02:00:00:00:00:01
swapon: adding /dev/wd1s1b as swap device
Automatic reboot in progress...
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd0s1e: clean, 9233 free (9 frags, 1153 blocks, 0.1% fragmentation)
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd1s1f: clean, 4301055 free (335 frags, 537590 blocks, 0.0% fragmentation)
```

show system boot-messages lcc (TX Matrix Router)

```
user@host> show system boot-messages lcc 2
lcc2-re0:
--------------------------------------------------------------------------
Copyright (c) 1996-2001, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2001 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
        The Regents of the University of California. All rights reserved.
JUNOS 7.0-20040912.0 #0: 2004-09-12 09:16:32 UTC

builder@device1.example.com:/build/benten-b/7.0/20040912.0/obj-i386/sys/compile/JUNIPER
Timecounter "i8254"  frequency 1193182 Hz
Timecounter "TSC"  frequency 601368936 Hz
CPU: Pentium III/Pentium III Xeon/Celeron (601.37-MHz 686-class CPU)
  Origin = "GenuineIntel"  Id = 0x68a  Stepping = 10

Features=0x387f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,PAT,PSE36,PN,MMX,FXSR
,SSE>
real memory  = 2147467264 (2097136K bytes)
sio0: gdb debugging port
avail memory = 2084040704 (2035196K bytes)
Preloaded elf kernel "kernel" at 0xc06d9000.
DEVFS: ready for devices
```

```
Pentium Pro MTRR support enabled
md0: Malloc disk
DRAM Data Integrity Mode: ECC Mode with h/w scrubbing
npx0: <math processor> on motherboard
npx0: INT 16 interface
pcib0: <ServerWorks NB6635 3.0LE host to PCI bridge> on motherboard
pci0: <PCI bus> on pcib0
pcic-pci0: <TI PCI-1410 PCI-CardBus Bridge> irq 15 at device 1.0 on pci0
pcic-pci0: TI12XX PCI Config Reg: [pwr save][pci only]
fxp0: <Intel Embedded 10/100 Ethernet> port 0x1000-0x103f mem
0xfb800000-0xfb81ffff,0xfb820000-0xfb820fff irq 9 at device 3.0 on pci0
fxp1: <Intel Embedded 10/100 Ethernet> port 0x1040-0x107f mem
0xfb840000-0xfb85ffff,0xfb821000-0xfb821fff irq 11 at device 4.0 on pci0
...
```

show system boot-messages (TX Matrix Plus Router)

```
user@host> show system boot-messages
sfc0-re0:
--------------------------------------------------------------------------
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
        The Regents of the University of California. All rights reserved.
JUNOS 9.6B3.3 #0: 2009-06-17 19:52:08 UTC
    builder@device1.example.com:/volume/build/junos/9.6/release/9.6B3.3/obj-i386/bsd/sys/compile/
JUNIPER
MPTable:  Timecounter "i8254" frequency 1193182 Hz quality 0 CPU: Intel(R) Xeon(R) CPU
L5238  @ 2.66GHz (2660.01-MHz 686-class CPU)   Origin = "GenuineIntel"  Id = 0x1067a  Stepping =
10   Features=0xbfebfbff
...
lcc1-re0:
--------------------------------------------------------------------------
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
        The Regents of the University of California. All rights reserved.
JUNOS 9.6-20090617.0 #0: 2009-06-17 04:15:14 UTC
    builder@device1.example.com:/volume/build/junos/9.6/production/20090617.0/obj-i386/bsd/sys/
```

```
compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz quality 0
CPU: Intel(R) Xeon(R) CPU                  @ 1.86GHz (1862.01-MHz 686-class CPU)
  Origin = "GenuineIntel"  Id = 0x1067a  Stepping = 10
  Features=0xbfebfbff
...
```

show system boot-messages (QFX3500 Switch)

```
user@switch> show system boot-messages
getmemsize: msgbufp[size=32768] = 0x81d07fe4

System physical memory distribution:
--------------------------------------------------------------------------------
Total physical memory: 4160749568 (3968 MB)
Physical memory used: 3472883712 (3312 MB)
Physical memory allocated to kernel: 2130706432 (2032 MB)
Physical memory allocated to user BTLB: 1342177280 (1280 MB)
--------------------------------------------------------------------------------

Copyright (c) 1996-2010, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
        The Regents of the University of California. All rights reserved.
JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC
    user@device1.example.com:/c/user/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
WARNING: debug.mpsafenet forced to 0 as ipsec requires Giant
JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC
    user@device1.example.com:/c/user/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
real memory  = 3472883712 (3312MB)
avail memory = 1708171264 (1629MB)
cpuid: 0, btlb_cpumap:0xffffffff8
FreeBSD/SMP: Multiprocessor System Detected: 12 CPUs
ETHERNET SOCKET BRIDGE initialising
Initializing QFX platform properties ..
cpu0 on motherboard
: RMI's XLR CPU Rev. 0.3 with no FPU implemented
        L1 Cache: I size 32kb(32 line), D size 32kb(32 line), eight way.
```

```
        L2 Cache: Size 1024kb, eight way
pic_lbus0: <XLR Local Bus>
pic_lbus0: <XLR Local Bus> on motherboard
Enter qfx control ethernet probe addr:0xc5eeec00
gmac4: <XLR GMAC GE Ethernet> on pic_lbus0
me0: Ethernet address 00:1d:b5:f7:68:40
Enter qfx control ethernet probe addr:0xc5eeeb40
gmac5: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:41
Enter qfx control ethernet probe addr:0xc5eeea80
gmac6: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:42
sio0 on pic_lbus0
Entering sioattach
sio0: type 16550A, console
xls_setup_intr: skip irq 3, xlr regs are set up somewhere else.
gblmem0 on pic_lbus0
ehci0: <RMI XLS USB 2.0 controller> on pic_lbus0
ehci_bus_attach: allocated resource.  tag=1, base=bef24000
xls_ehci_init: endian hardware swapping NOT enabled.
usb0: EHCI version 1.0
usb0 on ehci0
usb0: USB revision 2.0
uhub0: vendor 0x0000 EHCI root hub, class 9/0, rev 2.00/1.00, addr 1
uhub0: 2 ports with 2 removable, self powered
umass0: USB USBFlashDrive, rev 2.00/11.00, addr 2
pcib0: PCIe link 0 up
pcib0: PCIe link 2 up
pcib0: PCIe link 3 up
pcib0: <XLS PCI Host Controller> on pic_lbus0
pci0: <PCI bus> on pcib0
pcib1: <PCI-PCI bridge> at device 0.0 on pci0
pci1: <PCI bus> on pcib1
pci1: <network, ethernet> at device 0.0 (no driver attached)
pcib2: <PCI-PCI bridge> at device 1.0 on pci0
pcib3: <PCI-PCI bridge> at device 2.0 on pci0
pci2: <PCI bus> on pcib3
pci2: <network, ethernet> at device 0.0 (no driver attached)
pcib4: <PCI-PCI bridge> at device 3.0 on pci0
pci3: <PCI bus> on pcib4
pci3: <network, ethernet> at device 0.0 (no driver attached)
cfi device address space at 0xbc000000
cfi0: <AMD/Fujitsu - 8MB> on pic_lbus0
```

```
cfi device address space at 0xbc000000
i2c0: <I2C bus controller> on pic_lbus0
i2c1: <I2C bus controller> on pic_lbus0
qfx_fmn0 on pic_lbus0
pool offset 1503776768
xlr_lbus0: <XLR Local Bus Controller> on motherboard
qfx_bcpld_probe[124]
qfx_bcpld_probe[138]: dev_type=0x0
qfx_bcpld_probe[124]
qfx_bcpld0: QFX BCPLD probe success
qfx_bcpld0qfx_bcpld_attach[174]
qfx_bcpld_attach[207] : bus_space_tag=0x0, bus_space_handle=0xbd900000
qfx_bcpld_probe[124]
qfx_bcpld1: QFX BCPLD probe success
qfx_bcpld1qfx_bcpld_attach[174]
tor_bcpld_slave_attach[1245] : bus_space_tag=0x0, bus_space_handle=0xbda00000
Initializing product: 96 ..
bmeb: bmeb_lib_init done 0xc60a5000, addr 0x809c99a0
bme0:Virtual BME driver initializing
Timecounter "mips" frequency 1200000000 Hz quality 0
Timecounter "xlr_pic_timer" frequency 66666666 Hz quality 1
Timecounters tick every 1.000 msec
Loading the NETPFE fc module
IPsec: Initialized Security Association Processing.
SMP: AP CPU #3 Launched!
SMP: AP CPU #1 Launched!
SMP: AP CPU #2 Launched!
SMP: AP CPU #4 Launched!
SMP: AP CPU #5 Launched!
SMP: AP CPU #7 Launched!
SMP: AP CPU #6 Launched!
SMP: AP CPU #11 Launched!
SMP: AP CPU #10 Launched!
SMP: AP CPU #9 Launched!
SMP: AP CPU #8 Launched!
da0 at umass-sim0 bus 0 target 0 lun 0
da0: <USB USBFlashDrive 1100> Removable Direct Access SCSI-0 device
da0: 40.000MB/s transfers
da0: 3920MB (8028160 512 byte sectors: 255H 63S/T 499C)
Trying to mount root from ufs:/dev/da0s1a
```

show system boot-messages (SRX300)

```
user@host> show system boot-messages |no-more
kld_map_v: 0x8ff80000, kld_map_p: 0x0
Running in PARTITIONED TLB MODE
Copyright (c) 1996-2022, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2007 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
FreeBSD is a registered trademark of The FreeBSD Foundation.
JUNOS 21.2R3-S2.9 #0: 2022-07-16 04:20:05 UTC
    builder@example1.juniper.net:/volume/build/junos/21.2/release/21.2R3-S2.9/obj/bsd6/octeon/
junos/bsd/kernels/JSRXNLE/kernel
can't re-use a leaf (perf_mon)!
can't re-use a leaf (threshold)!
can't re-use a leaf (debug)!
JUNOS 21.2R3-S2.9 #0: 2022-07-16 04:20:05 UTC
    builder@example1.juniper.net:/volume/build/junos/21.2/release/21.2R3-S2.9/obj/bsd6/octeon/
junos/bsd/kernels/JSRXNLE/kernel
real memory  = 4294967296 (4194304K bytes) avail memory = 2286268416 (2180MB)
FreeBSD/SMP: Multiprocessor System Detected: 2 CPUs Security policy loaded: Junos MAC/veriexec
(mac_veriexec) Security policy loaded: JUNOS MAC/pcap (mac_pcap) Security policy loaded: JUNOS
MAC/runasnonroot (mac_runasnonroot) MAC/veriexec fingerprint module loaded: SHA1 MAC/veriexec
fingerprint module loaded: SHA256
netisr_init: forcing maxthreads from 4 to 2
random: <Software, Yarrow> initialized
cpu0 on motherboard
: CAVIUM's OCTEON 70XX/71XX CPU Rev. 0.2 with no FPU implemented
        L1 Cache: I size 78kb(128 line), D size 32kb(128 line), thirty two way.
        L2 Cache: Size 512kb, 4 way
obio0 on motherboard
uart0: <Octeon-16550 channel 0> on obio0
uart0: console (9600,n,8,1)
twsi0 on obio0
set clock 0x49
xhci0: <Cavium Octeon 7xxx xHCI Host Driver> on obio0
usb0: <USB bus for xHCI Controller> on xhci0
usb0: USB revision 3.0
uhub0: vendor 0x0000 XHCI root hub, class 9/0, rev 3.00/1.00, addr 1
uhub0: 2 ports with 2 removable, self powered
xhci1: <Cavium Octeon 7xxx xHCI Host Driver> on obio0
```

```
usb1: <USB bus for xHCI Controller> on xhci1
usb1: USB revision 3.0
uhub1: vendor 0x0000 XHCI root hub, class 9/0, rev 3.00/1.00, addr 1
uhub1: 2 ports with 2 removable, self powered
cpld0 on obio0
pcib0: <Cavium on-chip PCIe HOST bridge> on obio0 Disabling Octeon big bar support
pcib0: Initialized controller
pci0: <PCI bus> on pcib0
pci0: <network, ethernet> at device 0.0 (no driver attached)
pci0: <network, ethernet> at device 0.1 (no driver attached)
gblmem0 on obio0
octpkt0: <Octeon RGMII> on obio0
boot_bus0 on obio0
cfi0: <Macronix MX25L64 - 8MB> on boot_bus0
cfi1: <Macronix MX25L64 - 8MB> on boot_bus0
umass0: ATP Electronics ATP CG eUSB, rev 2.00/11.00, addr 2 Timecounter "mips" frequency
1200000000 Hz quality 0
da0 at umass-sim0 bus 0 target 0 lun 0
da0: <ATP ATP CG eUSB 1100> Fixed Direct Access SCSI-4 device
da0: 40.000MB/s transfers
da0: 7672MB (15712256 512 byte sectors: 255H 63S/T 978C)
random: unblocking device.
hwpmc: OCTEON/4/64/0x1ff<INT,USR,SYS,EDG,THR,REA,WRI,INV,QUA>
Trying to mount root from ufs:/dev/da0s1a
WARNING: / was not properly dismounted
WARNING: / was not properly dismounted
WARNING: R/W mount of /cf/var denied.  Filesystem is not clean - run fsck

LPC bus driver
lpcbus0 on cpld0
tpm0: <Trusted Platform Module> on lpcbus0
tpm: IFX SLB 9660 TT 1.2 rev 0x10
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system buffers

IN THIS SECTION

- Syntax | 26741
- Syntax (EX Series) | 26741
- Syntax (TX Matrix Router) | 26742
- Syntax (TX Matrix Plus Router) | 26742
- Syntax (MX Series Router) | 26742
- Syntax (QFX Series) | 26742
- Description | 26742
- Options | 26743
- Additional Information | 26744
- Required Privilege Level | 26744
- Output Fields | 26744
- Sample Output | 26746
- Release Information | 26750

Syntax

```
show system buffers
```

Syntax (EX Series)

```
show system buffers
<all-members>
```

```
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system buffers
<all-chassis | all-lcc | lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
show system buffers
<all-chassis | all-lcc | lcc number | sfc number>
```

Syntax (MX Series Router)

```
show system buffers
<all-members>
<local>
<member member-id>
```

Syntax (QFX Series)

```
show system buffers
<infrastructure name | interconnect-device name | node-group name | root-only (infrastructure
name | interconnect-device name | node-group name)>
```

Description

Display information about the buffer pool that the Routing Engine uses for local traffic. Local traffic is the routing and management traffic that is exchanged between the Routing Engine and the Packet Forwarding Engine within the router or switch, as well as the routing and management traffic from IP (that is, from OSPF, BGP, SNMP, ping operations, and so on).

Options

none Show all buffer statistics.

all-lcc (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show buffer statistics for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, show buffer statistics for all routers connected to the TX Matrix Plus router.

all-chassis (TX Matrix routers and TX Matrix Plus routers only) (Optional) Show buffer statistics for all of the chassis.

all-members (EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for all members of the Virtual Chassis configuration.

infrastructure *name* (QFabric systems only) (Optional) Show buffer statistics for a fabric control Routing Engine or a fabric control Routing Engine.

interconnect-device *name* (QFabric systems only) (Optional) Show buffer statistics for the Interconnect device.

lcc *number* (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show buffer statistics for a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, show buffer statistics for a specific router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for the local Virtual Chassis member.

member *member-id* (EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for the specified member of the Virtual Chassis configuration. For EX4200 switches,

replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

node-group *name* (QFabric systems only) (Optional) Show buffer statistics for the Node group

sfc (TX Matrix Plus routers only) (Optional) Show buffer statistics for the TX Matrix Plus router. Replace ***number*** with **0**.

Additional Information

By default, when you issue the `show system buffers` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

A special type of memory buffer called a *cluster* is 2 KB in size. For more information, see *The Design and Implementation of the 4.4BSD Operation System* by McKusic, Bostic, Karels, and Quarterman.

Required Privilege Level

view

Output Fields

Table 1610 on page 26744 describes the output fields for the `show system buffers` command. Output fields are listed in the approximate order in which they appear.

Table 1610: show system buffers Output Fields

Field Name	Field Description
mbufs in use	Memory buffers (mbufs) are 128-byte buffers that are used for various purposes inside the kernel. Each memory buffer has a type, and the output itemizes the amount allocated for each type. Types with no memory buffers allocated are not displayed.
mbufs allocated to packet headers	Number of memory buffers currently holding packet headers

Table 1610: show system buffers Output Fields *(Continued)*

Field Name	Field Description
mbufs allocated to control blocks	Number of memory buffers currently holding the state for sockets.
mbufs allocated to send data	Number of memory buffers currently holding socket send data.
mbufs allocated to pfe refill data	Number of memory buffers currently holding Packet Forwarding Engine refill data.
mbufs allocated to fxp data	Number of memory buffers currently holding fxp data.
mbufs allocated to socket names and addresses	Number of memory buffers currently holding addresses for sockets.
mbuf clusters in use	Allocation statistics for memory buffer clusters.
allocated to network	Total amount of memory in use by the networking and interprocess communication (IPC) code.
requests for memory denied	Number of times a memory allocation request within the IPC and networking code failed.
requests for memory delayed	Number of times a memory allocation request within the IPC and networking code was postponed.
calls to protocol drain routines	Number of times a memory allocation request within the IPC and networking code triggered a memory reclamation attempt.

Sample Output

show system buffers

```
user@host> show system buffers
397/893/1290 mbufs in use (current/cache/total)
395/331/726/30000 mbuf clusters in use (current/cache/total/max)
384/256 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
889K/885K/1774K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/5/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

show system buffers scc (TX Matrix Router)

```
user@host> show system buffers scc
213 mbufs in use:
        11 mbufs allocated to packet headers
        26 mbufs allocated to socket names and addresses
        2 mbufs allocated to socket options
        17 mbufs allocated to socket send data
        2 mbufs allocated to pfe data
        155 mbufs allocated to fxp data (rx)
        511 mbufs allocated to <mbuf type 86>
        256 mbufs allocated to <mbuf type 92>
924/1162 mbuf clusters in use
2788 Kbytes allocated to network (75% in use)
0 requests for memory denied
0 requests for memory delayed
0 calls to protocol drain routines
```

show system buffers sfc (TX Matrix Plus Router)

```
user@host> show system buffers sfc 0

sfc0-re0:
--------------------------------------------------------------------------
4363/2807/7170 mbufs in use (current/cache/total)
4358/1968/6326/30000 mbuf clusters in use (current/cache/total/max)
256/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
9806K/4637K/14444K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/10/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

show system buffers all-chassis (TX Matrix Plus Router)

```
user@host> show system buffers all-chassis

sfc0-re0:
--------------------------------------------------------------------------
4363/2807/7170 mbufs in use (current/cache/total)
4358/1968/6326/30000 mbuf clusters in use (current/cache/total/max)
256/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
9806K/4637K/14444K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/10/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
```

```
0 calls to protocol drain routines

lcc0-re0:
------------------------------------------------------------------------
772/2558/3330 mbufs in use (current/cache/total)
772/598/1370/30000 mbuf clusters in use (current/cache/total/max)
768/512 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1737K/1835K/3572K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

lcc1-re0:
------------------------------------------------------------------------
773/2437/3210 mbufs in use (current/cache/total)
773/453/1226/30000 mbuf clusters in use (current/cache/total/max)
768/384 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1739K/1515K/3254K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/7/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

lcc2-re0:
------------------------------------------------------------------------
816/2514/3330 mbufs in use (current/cache/total)
816/554/1370/30000 mbuf clusters in use (current/cache/total/max)
768/512 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
```

```
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1836K/1736K/3572K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
```

show system buffers node-group (QFabric System)

```
user@switch> show system buffers node-group
node1
node-group node1:
-----------------------------------------------------------------------------
2/2698/2700 mbufs in use (current/cache/total)
2/1520/1522/30000 mbuf clusters in use (current/cache/total/max)
0/1280 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
4K/3714K/3719K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/6/6656 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

re0:
-----------------------------------------------------------------------------
516/639/1155 mbufs in use (current/cache/total)
515/147/662/30000 mbuf clusters in use (current/cache/total/max)
512/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1159K/453K/1612K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
```

```
0/4/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

re1:
--------------------------------------------------------------------------
519/771/1290 mbufs in use (current/cache/total)
518/176/694/30000 mbuf clusters in use (current/cache/total/max)
512/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1165K/544K/1710K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sfbufs in use (current/peak/max)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Routing Matrix with a TX Matrix Plus Router Solutions Page

show system certificate

IN THIS SECTION

- Syntax | 26751
- Description | 26751
- Options | 26751
- Required Privilege Level | 26751
- Output Fields | 26752
- Sample Output | 26753
- Release Information | 26754

Syntax

```
show system certificate
<certificate-id>
```

Description

(Encryption interface on M Series, T Series routers, QFX Series, and OCX Series switches only) Display installed certificates signed by the Juniper Networks certificate authority.

Options

none Display all installed certificates signed by the Juniper Networks certificate authority.

certificate-id (Optional) Display the details of a particular certificate.

Required Privilege Level

maintenance

Output Fields

Table 1611 on page 26752 lists the output fields for the `show system certificate` command. Output fields are listed in the approximate order in which they appear.

Table 1611: show system certificate Output Fields

Field Name	Field Description
Certificate identifier	Unique identifier associated with a certificate. The certificate identifier is the common name of the subject.
Issuer **Subject**	Information about the certificate issuer and the distinguished name (**DN**) of the issuer, respectively: • **Organization**—Name of the owner's organization. • **Organizational unit**—Name of the owner's department. • **Country**—Two-character country code in which the owner's system is located. • **State**—State in the USA in which the owner is using the certificate. • **Locality**—City in which the owner's system is located. • **Common name**—Name of the owner of the certificate. • **E-mail address**—E-mail address of the owner of the certificate.
Validity	When a certificate is valid.
Signature algorithm	Encryption algorithm applied to the installed certificate.
Public key algorithm	Encryption algorithm applied to the public key.

Sample Output

show system certificate

```
user@host> show system certificate
Certificate identifier: Dallas-v3
  Issuer:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country:  US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
E-mail address:ca@example.com
  Subject:
  Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
E-mail address:ca@example.com
Validity:
 Not before: Mar 13 03:23:25 2004 GMT
  Not after: Mar 24 03:23:25 2014 GMT
Signature algorithm: sha1WithRSAEncryption
Public key algorithm: dsaEncryption
```

show system certificate (QFX Series)

```
user@host> show system certificate
Certificate identifier: Dallas-v3
  Issuer:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country:  US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
E-mail address:ca@example.com
  Subject:
  Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
E-mail address:ca@example.com
Validity:
 Not before: Mar 13 03:23:25 2004 GMT
  Not after: Mar 24 03:23:25 2014 GMT
Signature algorithm: sha1WithRSAEncryption
Public key algorithm: dsaEncryption
```

Release Information

Command introduced before Junos OS Release 7.4.

show system commit

IN THIS SECTION

- Syntax | 26754
- Description | 26754
- Options | 26754
- Required Privilege Level | 26755
- Output Fields | 26755
- Sample Output | 26756
- Release Information | 26758

Syntax

```
show system commit
<revision | server | synchronize-server pending-jobs | include-configuration-revision>
```

Description

Display the system commit history and pending commit operations.

Options

none Display the last 50 commit operations on the static configuration database, starting with the most recent.

revision (Optional) Display the revision number of the active configuration of the Routing Engine(s).

server (Optional) Display the commit server status.

> **NOTE**: By default, the status of the commit server is “Not running”. The commit server starts running only when a commit job is added to the batch.

synchronize-server pending-jobs (Optional) Display the pending commit synchronize operations for all instances of the ephemeral configuration database on an MX Series Virtual Chassis or a device with dual Routing Engines. This option can only be executed on the primary Routing Engine of the Virtual Chassis primary router or the dual Routing Engine system.

include-configuration-revision (Optional) Display configuration revision information including the revision identifier string for each commit record.

Required Privilege Level

view

Output Fields

Table 1612 on page 26755 describes the output fields for the `show system commit` command. Output fields are listed in the approximate order in which they appear.

Table 1612: show system commit Output Fields

Field Name	Field Description	Level of Output
<number>	Displays the last 50 commit operations listed, most recent to first. The identifier *<number>* designates a configuration created for recovery using the `request system configuration rescue save` command.	none
<time-stamp>	Date and time of the commit operation.	none
<root>/ *<username>*	User who executed the commit operation.	none

Table 1612: show system commit Output Fields *(Continued)*

Field Name	Field Description	Level of Output
<method>	Method used to execute the commit operation: • `CLI`—CLI interactive user performed the commit operation. • `Junos XML protocol`—Junos XML protocol client performed the commit operation. • `synchronize`—The `commit synchronize` command was performed on the other Routing Engine. • `snmp`—An SNMP `set` request caused the commit operation. • `button`—A button on the router or switch was pressed to commit a rescue configuration for recovery. • `autoinstall`—A configuration obtained through autoinstallation was committed. • `other`—When there is no login name associated with the session, the values for user and client default to root and other. For example, during a reboot after package installation, mgd commits the configuration as a system commit, and there is no login associated with the commit.	`none`
<rollback pending>	Identifies whether `commit confirmed` is issued. It is removed once `commit` or `commit check` is issued or `commit confirmed` is rolled back after rollback timeout.	`none`

Sample Output

show system commit

```
user@host> show system commit
0   2003-07-28 19:14:04 PDT by root via other
1   2003-07-25 22:01:36 PDT by user via cli
2   2003-07-25 22:01:32 PDT by user via cli
3   2003-07-25 21:30:13 PDT by root via button
4   2003-07-25 13:46:48 PDT by user via cli
5   2003-07-25 05:33:21 PDT by root via autoinstall
```

```
...
rescue  2002-05-10 15:32:03 PDT by root via other
```

show system commit (At a Particular Time)

```
user@host> show system commit
commit requested by root via cli at Tue May  7 15:59:00 2002
```

show system commit (At the Next Reboot)

```
user@host> show system commit
commit requested by root via cli at reboot
```

show system commit (Rollback Pending)

```
user@host> show system commit
0   2022-05-06 06:33:42 PDT by root via cli commit confirmed, rollback in 10mins
```

show system commit (QFX Series)

```
user@switch> show system commit
0   2011-11-25 19:17:49 PST by root via cli
```

show system commit synchronize-server pending-jobs

```
user@host> show system commit synchronize-server pending-jobs
Job Id                 Commit Synchronize Model       Database                       Version
36                     Asynchronous Commit            ephemeral                      130
37                     Asynchronous Commit            ephemeral                      131
38                     Asynchronous Commit            ephemeral                      132
39                     Asynchronous Commit            ephemeral                      133
40                     Asynchronous Commit            ephemeral                      134
41                     Asynchronous Commit            ephemeral                      135
42                     Asynchronous Commit            ephemeral                      2
```

show system commit include-configuration-revision

```
user@host> show system commit include-configuration-revision
0   2020-08-02 00:42:58 IST by user via cli re0-1596309177-4
1   2020-08-02 00:42:53 IST by user via cli re0-1596309173-3
2   2020-08-02 00:42:50 IST by user via cli re0-1596309170-2
3   2020-08-02 00:42:40 IST by user via other re0-1596309160-1
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `server` introduced in Junos OS Release 12.1 for the PTX Series router.

Option `revision` introduced in Junos OS Release 14.1.

Option `synchronize-server` introduced in Junos OS Release 17.2R1 and Junos OS Evolved Release 22.1R1.

Option `include-configuration-revision` introduced in Junos OS Release 20.4R1 and Junos OS Evolved Release 20.4R1.

RELATED DOCUMENTATION

clear system commit | 4219

show system commit revision

show system commit server queue

IN THIS SECTION

- Syntax | 26759
- Description | 26759
- Options | 26759
- Required Privilege Level | 26759
- Sample Output | 26760
- Release Information | 26763

Syntax

```
show system commit server queue
<id commit-id>
<job-status (all| error| pending| success)>
<patch (none | id commit-id) | (job-status (all | error | pending | success)>
```

Description

Display the status of commit server queue transactions.

> **NOTE**: Only 50 successful commit jobs are stored in the database and displayed in the output. When the fifty-first job is committed, the first job is deleted from the database and is no longer displayed in the output.

Options

id *commit-id*	(Optional) Display the batch commit operation status messages for a specific commit ID.
job-status	(Optional) Display batch commit operation status messages for the following batch commit statuses: • `all`—Status messages for all batch commit operations. • `error`—Status messages for batch commit operations that have errors. • `pending`—Status messages for batch commit operations that are pending. • `success`—Status messages for batch commit operations that are successful.
patch (none | id *commit-id*) | job-status (all |error | pending | success)	(Optional) Display the patch file containing the configuration changes for all batch commit operations, a specific batch commit ID, or a specific job status.

Required Privilege Level

view

Sample Output

show system commit server queue

```
user@host> show system commit server queue

Pending commits:
  none

Completed commits:
  Id: 1000
  Last Modified: Tue Nov  1 22:46:43 2011
  Status: Successfully committed 1000

  Id: 1002
  Last Modified: Tue Nov  1 22:50:35 2011
  Status: Successfully committed 1002

  Id: 1004
  Last Modified: Tue Nov  1 22:51:48 2011
  Status: Successfully committed 1004

  Id: 1007
  Last Modified: Wed Nov  2 01:08:04 2011
  Status: Successfully committed 1007

  Id: 1009
  Last Modified: Wed Nov  2 01:16:45 2011
  Status: Successfully committed 1009

  Id: 1010
  Last Modified: Wed Nov  2 01:19:25 2011
  Status: Successfully committed 1010

  Id: 1011
  Last Modified: Wed Nov  2 01:28:16 2011
  Status: Successfully committed 1011


Error commits:
  Id: 1008
```

```
  Last Modified: Wed Nov  2 01:08:18 2011
  Status: Error while commiting 1008
```

show system commit server queue job-status success

```
user@host> show system commit server queue job-status success
Completed commits:
  Id: 1000
  Last Modified: Tue Nov  1 22:46:43 2011
  Status: Successfully committed 1000

  Id: 1001
  Last Modified: Tue Nov  1 22:47:02 2011
  Status: Successfully committed 1001
```

show system commit server queue patch

```
user@host> show system commit server queue patch
Pending commits:
  none


Completed commits:
  Id: 1000
  Last Modified: Tue Nov  1 22:46:43 2011
  Status: Successfully committed 1000

Patch:
[edit system commit]
+ server {
+     days-to-keep-error-logs 4294967295;
+     traceoptions {
+         file commitd_nov;
+         flag all;
+     }
+ }
  Id: 1002
  Last Modified: Tue Nov  1 22:50:35 2011
  Status: Successfully committed 1002
```

```
Patch:
[edit system commit server]
- days-to-keep-error-logs 4294967295;
  Id: 1004
  Last Modified: Tue Nov  1 22:51:48 2011
  Status: Successfully committed 1004

Patch:
[edit system commit server]
+ days-to-keep-error-logs 4294967295;
  Id: 1007
  Last Modified: Wed Nov  2 01:08:04 2011
  Status: Successfully committed 1007

Patch:
[edit system commit server]
-  days-to-keep-error-logs 4294967295;
+  days-to-keep-error-logs 2;
  Id: 1009
  Last Modified: Wed Nov  2 01:16:45 2011
  Status: Successfully committed 1009

Patch:
[edit]
+ snmp {
+     community abc;
+ }
  Id: 1010
  Last Modified: Wed Nov  2 01:19:25 2011
  Status: Successfully committed 1010

Patch:
[edit system syslog]
  file test { ... }
+ file j {
+     any any;
+ }
  Id: 1011
  Last Modified: Wed Nov  2 01:28:16 2011
  Status: Successfully committed 1011


Error commits:
```

```
  Id: 1008
  Last Modified: Wed Nov  2 01:08:18 2011
  Status: Error while commiting 1008

Patch:
[edit system]
+ radius-server {
+     10.1.1.1 port 222;
+ }
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Batch Commit Server Properties

show system commit server status

IN THIS SECTION

- Syntax | 26763
- Description | 26764
- Options | 26764
- Required Privilege Level | 26764
- Sample Output | 26764
- Release Information | 26764

Syntax

```
show system commit server status
```

Description

Display commit server status.

> **NOTE**: By default, the status of the commit server is “Not running”. The commit server starts running only when a commit job is added to the batch.

Options

This command has no options.

Required Privilege Level

view

Sample Output

show system commit server status (When Server Is Inactive)

```
user@host> show system commit server status
Commit server status : Not running
```

show system commit server status (When Server Is Active)

```
user@R0> show system commit server status

Commit server status : Running
Jobs in process:
  1369 1370 1371
```

Release Information

Command introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Batch Commit Server Properties

show system configuration archival

IN THIS SECTION

- Syntax | 26765
- Description | 26765
- Options | 26765
- Required Privilege Level | 26765
- Sample Output | 26766
- Release Information | 26766

Syntax

```
show system configuration archival
```

Description

Display directory and number of files queued for archival transfer.

NOTE: The `[edit system configuration]` hierarchy is not available on QFabric systems.

Options

This command has no options.

Required Privilege Level

maintenance

Sample Output

show system configuration archival

```
user@host> show system configuration archival

/var/transfer/config/:
total 8
```

Release Information

Introduced in Junos OS Release 7.6.

Command introduced in Junos OS Release 9.0.

show system configuration database status

IN THIS SECTION

- Syntax | 26766
- Description | 26767
- Options | 26767
- Required Privilege Level | 26767
- Output Fields | 26767
- Release Information | 26767

Syntax

```
show system configuration database status
```

Description

Display the users currently editing the configuration.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1 describes the output fields for the `show system configuration database status` command. Output fields are listed in the approximate order in which they appear.

Table 1613: show system configuration database status Output Fields

Field Name	Field Description	Level of Detail
`User`	Userid.	all
`Terminal`	Terminal ID.	all
PID	Process ID.	all
`Start-time`	Amount of time user editing system configuration.	all

Release Information

Command introduced before Junos OS Release 23.3R1.

RELATED DOCUMENTATION

Displaying Users Currently Editing the Configuration

show system configuration database usage

IN THIS SECTION

- Syntax | 26768
- Description | 26768
- Options | 26768
- Required Privilege Level | 26768
- Output Fields | 26769
- Sample Output | 26769
- Release Information | 26769

Syntax

```
show system configuration database usage
<ephemeral-instance (default | instance-name)>
```

Description

Display configuration database disk space usage statistics.

Options

none	Display statistics for the static configuration database.
ephemeral-instance (default | *instance-name*)	(Optional) Display statistics for the specified ephemeral configuration database instance.

Required Privilege Level

maintenance

Output Fields

Table 1614 on page 26769 describes the output fields for the `show system configuration database usage` command. Output fields are listed in the approximate order in which they appear.

Table 1614: show system configuration database usage Output Fields

Field Name	Field Description
`Maximum size of the database`	Display the maximum available space on the disk to store the configuration database
`Current database size on disk`	Display the total space on the disk used by the current configuration database
`Actual database usage`	Display the actual space on the disk used by the current configuration data
`Available database space`	Display the free space available on the disk to store the configuration database

Sample Output

show system configuration database usage

```
user@host> show system configuration database usage
Maximum size of the database: 665.99 MB
Current database size on disk: 1.50 MB
Actual database usage: 1.48 MB
Available database space: 664.51 MB
```

Release Information

Command introduced in Junos OS Release 15.1.

`ephemeral-instance` option introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Overview for Junos OS

show system configuration rescue

IN THIS SECTION

- Syntax | 26770
- Description | 26770
- Options | 26770
- Required Privilege Level | 26770
- Sample Output | 26771
- Release Information | 26771

Syntax

```
show system configuration rescue
```

Description

Display a rescue configuration, if one exists.

NOTE: The `[edit system configuration]` hierarchy is not available on QFabric systems.

Options

This command has no options.

Required Privilege Level

maintenance

Sample Output

show system configuration rescue

```
user@switch> show system configuration rescue
version "7.3"; groups {
    global {
        system {
            host-name router1;
            domain-name customer.net;
            domain-search [ customer.net ];
            backup-router 192.0.2.0;
            name-server {
                192.0.2.11;
                192.0.2.101;
                192.0.2.100;
                192.0.2.10;
            }
            login {
                user user1 {
                    uid 928;
                    class ;
                    shell csh;
                    authentication {
                        encrypted-password "$ABC123"; ## SECRET-DATA
                    }
                }
            }
            services {
                ftp;
                rlogin;
                rsh;
                telnet;
            }
        }
    }
....
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show system configuration archival | 26765

show system configuration revision

IN THIS SECTION

- Syntax | 26772
- Description | 26772
- Options | 26772
- Required Privilege Level | 26773
- Sample Output | 26773
- Release Information | 26775

Syntax

```
show system configuration revision revision-string
<compare revision-string revision-string2>
<rollback-number revision-string >
```

Description

Display the configuration associated to the specified revision.

Options

revision-string	Show the configuration associated with the specified revision identifier string. Use the "show system commit include-configuration-revision" on page 26754 command to view the revision identifier string for each revision.
compare *revision-string revision-string2*	(Optional) Compare two different revisions by entering both revision identifier strings. The output displays the differences between the two configurations.

rollback-number *revision-string* (Optional) Show corresponding rollback number for the given configuration revision identifier string.

Required Privilege Level

view

Sample Output

show system configuration revision

```
user@host> show system configuration revision re0-1604691570-294
## Last changed: 2020-11-06 11:39:28 PST
version 20.4I-20201012.0.0359;
groups {
    re0 {
        system {
            host-name host;
            backup-router 192.0.2.254 destination 0.0.0.0/0;
        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 192.0.2.42/24;
                    }
                }
            }
        }
    }
    re1 {
        system {
            host-name host1;
            backup-router 192.0.2.254 destination 0.0.0.0/0;
        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 192.0.2.42/24;
```

```
                    }
                }
            }
        }
    }
    global {
        system {
            root-authentication {
                encrypted-password "$ABC123"; ## SECRET-DATA
    debugger-on-break;
    switchover-on-routing-crash;
    authentication-order [ password radius ];
    name-server {
        192.0.2.68;
        192.0.2.11;
    }
    radius-server {
        192.0.2.162 secret "$ABC123"; ## SECRET-DATA
        192.0.2.10 secret "$ABC123"; ## SECRET-DATA
    }
    schema {
        openconfig {
            unhide;
        }
    }
}
chassis {
    redundancy {
        graceful-switchover;
    }
    network-services enhanced-ip;
}
routing-options {
    static {
        route 0.0.0.0/0 next-hop 192.0.2.254;
    }
    nonstop-routing;
}
protocols {
    isis {
        interface all;
        interface fxp0.0 {
            disable;
```

```
        }
    }
}
```

Release Information

Command introduced in Junos OS Release 20.4R1 and Junos OS Evolved Release 20.4R1.

show system connections

IN THIS SECTION

- Syntax | 26775
- Syntax (EX Series) | 26776
- Syntax (MX Series Router) | 26776
- Syntax (QFX Series) | 26776
- Syntax (OCX Series) | 26776
- Syntax (Junos OS Evolved) | 26777
- Description | 26777
- Options | 26777
- Additional Information | 26778
- Required Privilege Level | 26778
- Output Fields | 26778
- Sample Output | 26779
- Release Information | 26782

Syntax

```
show system connections
<extensive>
<inet | inet6>
<show-routing-instances>
```

Syntax (EX Series)

```
show system connections
<extensive>
<all-members>
<inet | inet6>
<local>
<member member-id>
<show-routing-instances>
```

Syntax (MX Series Router)

```
show system connections
<extensive>
<all-members>
<inet | inet6>
<local>
<member member-id>
<show-routing-instances>
```

Syntax (QFX Series)

```
show system connections
<extensive>
<inet>
<infrastructure name>
<interconnect-device name>
<node-group name>
<show-routing-instances>
```

Syntax (OCX Series)

```
show system connections
<extensive>
<inet>
<show-routing-instances>
```

Syntax (Junos OS Evolved)

```
show system connections
<inet | inet6>
<node node-name>
```

Description

Display information about the active IP sockets on the Routing Engine. Use this command to verify which servers are active on a system and what connections are currently in progress.

Options

The following table lists all possible options for the `show system connections` command. The options that appear may vary depending on the platform and software release.

none	Display information about all active IP sockets on the Routing Engine.
extensive	(Optional) Display exhaustive system process information, which, for TCP connections, includes the TCP control block and maximum segment size (MSS). This option is useful for debugging TCP connections.
all-members	(Optional) Display system connection activity for all members of the Virtual Chassis configuration.
inet | inet6	(Optional) Display IPv4 connections or IPv6 connections, respectively.
infrastructure *name*	(QFabric systems only) (Optional) Display system connection activity for the fabric control Routing Engines or fabric manager Routing Engines.
interconnect-device *name*	(QFabric systems only) (Optional) Display system connection activity for the Interconnect device.
local	(Optional) Display system connection activity for the local Virtual Chassis member.
member *member-id*	(Optional) Display system connection activity for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.
node *node-name*	(Optional) Display system connection activity for the specified node.

node-group *name* (QFabric systems only) (Optional) Display system connection activity for the Node group.

show-routing-instances (Optional) Display routing instances.

Additional Information

NOTE: The device calculates the TCP MSS value as described in RFC 6691.

Required Privilege Level

view

Output Fields

Table 1615 on page 26778 describes the output fields for the `show system connections` command. Output fields are listed in the approximate order in which they appear.

Table 1615: show system connections Output Fields

Field Name	Field Description
`Proto`	Protocol of the socket: `IP`, `TCP`, or `UDP` for IPv4 or IPv6.
`Recv-Q`	Number of input bytes received by the protocol and waiting to be processed by the application.
`Send-Q`	Number of output bytes sent by the application and waiting to be processed by the protocol.
`Local Address`	Local address and port of the socket, separated by a period. An asterisk (*) indicates that the bound address is the wildcard address. Server sockets typically have the wildcard address and a well-known port bound to them.
`Foreign Address`	Foreign address and port of the socket, separated by a period. An asterisk (*) indicates that the address or port is a wildcard.

Table 1615: show system connections Output Fields *(Continued)*

Field Name	Field Description
`Routing Instance`	(Displayed only when the `show-routing-instance` option is used.) Routing instances associated with active IP sockets on the Routing Engine.
`(state)`	For TCP, the protocol state of the socket.

Sample Output

show system connections

```
user@host> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q  Local Address          Foreign Address        (state)
tcp        0      2  192.0.2.16.513       192.0.2.254.894    ESTABLISHED
tcp        0      0  192.0.2.16.513       192.0.2.195.945    ESTABLISHED
tcp        0      0  *.23                   *.*                    LISTEN
tcp        0      0  *.22                   *.*                    LISTEN
tcp        0      0  *.513                  *.*                    LISTEN
tcp00 *.514      *.*                    LISTEN
tcp 0 0*.21      *.*                    LISTEN
tcp00 *.79     *.*                    LISTEN
tcp 00 *.1023     *.*                    LISTEN
tcp 00 *.111     *.*                    LISTEN
udp00192.192.0.2.1634     192.0.2.249.2049
udp00192.192.0.2.1627     192.0.2.254.2049
udp00192.192.0.2.1371     192.0.2.195.2049
udp00*.*      *.*
udp00*.9999     *.*
udp00 *.161     *.*
udp00192.192.0.2.1039     192.0.2.16.1023
udp00192.192.0.2.1038     192.0.2.16.1023
udp 00 192.0.2.16.1037     192.0.2.16.1023
udp00 192.0.2.16.1036     192.0.2.16.1023
udp00*.1022     *.*
udp00*.1023     *.*
```

```
udp00*.111     *.*
udp00*.*      *.*
```

show system connections extensive

```
user@host> show system connections extensive

Active Internet connections (including servers)
Proto Recv-Q Send-Q  Local Address                                Foreign
Address                                (state)
tcp4       0      6  192.0.2.15.23
192.0.2.138.3013                          ESTABLISHED
   sndsbcc:          6 sndsbmbcnt:        256  sndsbmbmax:     272000
sndsblowat:       2048 sndsbhiwat:      34000
   rcvsbcc:          0 rcvsbmbcnt:          0  rcvsbmbmax:     533120
rcvsblowat:          1 rcvsbhiwat:      66640
   proc id:          0  proc name:
       iss: 2566994072      sndup: 2566994491
    snduna: 2566994491     sndnxt: 2566994494      sndwnd:      64094
    sndmax: 2566994494    sndcwnd:       6589 sndssthresh:       2720
       irs:  236981199      rcvup:  236981325
    rcvnxt:  236981327     rcvadv:  237046862      rcvwnd:      66640
       rtt:  140058623       srtt:      15519        rttv:        908
    rxtcur:       1200   rxtshift:          0       rtseq: 2566994491
    rttmin:       1000  mss:       1360
     flags: SACK_PERMIT [0x2000200]
tcp4       0      0  10.255.165.93.179
10.255.165.203.65141                          ESTABLISHED
   sndsbcc:          0 sndsbmbcnt:          0  sndsbmbmax:     131072
sndsblowat:       2048 sndsbhiwat:      16384
   rcvsbcc:          0 rcvsbmbcnt:          0  rcvsbmbmax:     131072
rcvsblowat:          1 rcvsbhiwat:      16384
   proc id:          0  proc name:
       iss: 2555961065      sndup: 2555995917
    snduna: 2555995917     sndnxt: 2555995917      sndwnd:      16384
    sndmax: 2555995917    sndcwnd:       1000 sndssthresh: 1073725440
       irs: 2123825753      rcvup: 2123860681
    rcvnxt: 2123860681     rcvadv: 2123877065      rcvwnd:      16384
       rtt:          0       srtt:       3309        rttv:         72
    rxtcur:       1200   rxtshift:          0       rtseq: 2555995898
    rttmin:       1000  mss:        500
```

```
     flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x3e0]
tcp4       0      0  10.255.165.203.65141
10.255.165.93.179                            ESTABLISHED
   sndsbcc:          0 sndsbmbcnt:          0  sndsbmbmax:     131072
sndsblowat:       2048 sndsbhiwat:      16384
   rcvsbcc:          0 rcvsbmbcnt:          0  rcvsbmbmax:     131072
rcvsblowat:          1 rcvsbhiwat:      16384
   proc id:       5022  proc name:        rpd
       iss: 2123825753      sndup: 2123860662
    snduna: 2123860681     sndnxt: 2123860681      sndwnd:      16384
    sndmax: 2123860681    sndcwnd:       1000 sndssthresh: 1073725440
       irs: 2555961065      rcvup: 2555995917
    rcvnxt: 2555995917     rcvadv: 2556012301      rcvwnd:      16384
       rtt:          0       srtt:       3279        rttv:         22
    rxtcur:       1200   rxtshift:          0       rtseq: 2123860662
    rttmin:       1000  mss:        500
     flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x100003e0]
tcp4       0      0  10.255.165.203.179
10.255.165.113.52404                         ESTABLISHED
   sndsbcc:          0 sndsbmbcnt:          0  sndsbmbmax:     131072
sndsblowat:       2048 sndsbhiwat:      16384
   rcvsbcc:          0 rcvsbmbcnt:          0  rcvsbmbmax:     131072
rcvsblowat:          1 rcvsbhiwat:      16384
   proc id:          0  proc name:
       iss: 1109297190      sndup: 1109332099
    snduna: 1109332118     sndnxt: 1109332118      sndwnd:      16384
    sndmax: 1109332118    sndcwnd:       1000 sndssthresh: 1073725440
       irs: 1476831634      rcvup: 1476866449
    rcvnxt: 1476866449     rcvadv: 1476882833      rcvwnd:      16384
       rtt:          0       srtt:       3235        rttv:         18
    rxtcur:       1200   rxtshift:          0       rtseq: 1109332099
    rttmin:       1000  mss:        500
     flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x3e0]
```

show system connections show-routing-instances

```
user@host> show system connections show-routing-instances
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q  Local Address          Foreign Address        Routing Instance       (state)
tcp4       0      0  192.0.2.204.23      192.0.2.19.4267      default                ESTABLISHED
tcp4       0      0  192.0.2.204.58540   10.209.7.138.23        default
```

```
ESTABLISHED
tcp4       0      0  192.0.2.204.23      192.0.2.19.1098      default                ESTABLISHED
tcp4       0      0  192.0.2.1.57668      192.0.2.1.179        default                ESTABLISHED
tcp4       0      0  192.0.2.1.179        192.0.2.1.49209      default                ESTABLISHED
tcp4       0      0  192.0.2.1.6234         192.0.2.17.1024        __juniper_private1__
ESTABLISHED
tcp4       0      0  192.0.2.4.9000         192.0.24.59103        __juniper_private1__
ESTABLISHED
tcp4       0      0  1192.0.2.4.59103        192.0.2.4.9000         __juniper_private1__
ESTABLISHED
tcp4       0      0  *.32012                *.*                    __juniper_private1__   LISTEN
tcp4       0      0  *.9000                 *.*                    __juniper_private1__   LISTEN
tcp4       0      0  *.33007                *.*                    __juniper_private2__   LISTEN
tcp46      0      0  *.179                  *.*                    default                LISTEN
tcp4       0      0  *.179                  *.*                    default                LISTEN
tcp4       0      0  *.6154                 *.*                    __juniper_private1__   LISTEN
tcp4       0      0  *.6153                 *.*                    __juniper_private1__   LISTEN
tcp4       0      0  *.7000                 *.*                    __juniper_private1__   LISTEN
tcp4       0      0  *.6152                 *.*                    __juniper_private1__   LISTEN
tcp4       0      0  *.6156                 *.*                    __juniper_private1__   LISTEN
tcp4       0      0  *.33005                *.*                    __juniper_private2__   LISTEN
tcp4       0      0  *.31343                *.*                    __juniper_private1__   LISTEN
tcp4       0      0  *.31341                *.*                    __juniper_private1__   LISTEN
tcp4       0      0  *.32003                *.*                    __juniper_private2__   LISTEN
tcp4       0      0  *.666                  *.*                    __juniper_private1__   LISTEN
tcp4       0      0  *.38                   *.*                    __juniper_private1__   LISTEN
tcp4       0      0  *.3221                 *.*                    default                LISTEN
```

show system connections | find sctp

```
user@host> show system connections | find sctp
Active SCTP associations (including servers)
Proto  Type  Local Address          Foreign Address        (state)
sctp4  1to1  30.30.30.1.50000       20.20.20.2.62324       ESTABLISHED
             20.20.20.1.50000       30.30.30.2.62324
```

Release Information

Command introduced before Junos OS Release 7.4.

Options `extensive` and `show-routing-instance` deprecated in Junos OS Evolved Release 17.3.

`node` option introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

tcp-mss (Internet Options) | 28996

tcp-mss (Protocols BGP) | 28998

show system core-dumps

IN THIS SECTION

- Syntax | 26783
- Syntax (SRX Series) | 26784
- Syntax (Junos OS Evolved) | 26784
- Syntax (EX Series Switches) | 26784
- Syntax (TX Matrix Router) | 26784
- Syntax (TX Matrix Plus Router) | 26784
- Syntax (QFX Series and OCX Series) | 26784
- Description | 26785
- Options | 26785
- Required Privilege Level | 26787
- Output Fields | 26787
- Sample Output | 26790
- Release Information | 26802

Syntax

```
show system core-dumps
<re0>
<re1>
```

```
<routing-engine>
<satellite [fpc-slot-id |device-alias  alias-name]>
```

Syntax (SRX Series)

```
show system core-dumps
```

Syntax (Junos OS Evolved)

```
show system core-dumps
```

Syntax (EX Series Switches)

```
show system core-dumps
<all-members>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system core-dumps
<all-chassis | all-lcc | lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
show system core-dumps
<all-chassis | all-lcc | lcc number | sfc number>
```

Syntax (QFX Series and OCX Series)

```
show system core-dumps
<component (UUID / serial number | all)>
<display-period (hours / minutes / seconds)>
```

```
<display-order>
<kernel-crashinfo component (UUID | serial number)>
<repository (core | log)>
```

Description

Show core files on all routers or switches running Junos OS. You can use the `show system core-dumps` command to show a list of system core files created when the router or switch has failed. This command can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification date, and path and filename. If dual Routing Engines are present, you can view core-dump files for either routing engine or both routing engines together. On a QFabric system, you can view core-dump files on individual QFabric system devices as well as on the entire QFabric system.

For Junos OS, all cores files are stored at **/var/core/re**. For Junos OS Evolved, a core file created during early bootup is stored in **/var/core/re**. But a core file created later in the bootup, for example, after the Routing Engine slot number can be determined, is stored in **/var/core/re0** or **/var/core/re1**. The command `show system core-dumps` continues to show all cores generated.

The core files are placed in the **/var/tmp/corefiles** on the SPC3 cards. Each PIC of the SPC3 card has five core files quota on the RE. When no more than five core files from one SPC3 PIC are on the RE, and the RE hard drive has more than 5 GB capacity, core file from the specific PIC is saved at the time it arrives. When there are already five core files from one SPC3 PIC on the RE, the newly arrived core file from the specific PIC replaces the last core file created by that PIC on the RE. When 5 GB capacity limit is reached, core file is not copied onto the RE. Only a zero sized shadow file with the same file name suffixed by ".shadow" is created. The core file is reached on the specific SPC3 PIC.

On vSRX Virtual Firewall, the srxpfe process might stop if you perform a reboot or an upgrade. You might observe core files whn you execute the show system core-dump command.

Options

none

Display a list of all existing core-dump files.

> **NOTE**: If dual Routing Engines are present, then only the core-dump files for the active Routing Engine are listed. For Junos OS Evolved, core-dump files for all Routing Engines are listed.

all-chassis

(TX Matrix and TX Matrix Plus routers only) (Optional) On a routing matrix based on a TX Matrix router, display system core files for the TX Matrix router switch-

card chassis [SCC] and all the T640 routers [LCCs] connected to the TX Matrix router.

On a routing matrix based on a TX Matrix Plus router, display system core files for the TX Matrix Plus router (switch-fabric chassis [SFC]) and all the T1600 routers [LCCs] connected to the TX Matrix Plus router.

<all-lcc | lcc *number*> (TX Matrix and TX Matrix Plus routers only) (Optional) On a routing matrix based on the TX Matrix router, display core dump files for all T640 routers (line-card chassis [LCCs]) or a specific T640 router [LCC] connected to the TX Matrix router.

On a routing matrix based on the TX Matrix Plus router, display logging information for all T1600 routers (line-card chassis [LCCs]) or a specific T1600 router (LCC) connected to the TX Matrix Plus router. When using the `lcc` *`number`* option, replace *`number`* with a value from `0` through `3`.

> **NOTE**: The `all-chassis` option displays system core files for the SCC or SFC and the LCCs connected to the SCC or SFC in the routing matrix while the `all-lcc` option only displays system core files for the LCCs in the routing matrix.

all-members (EX4200 switches) (Optional) Display system core files on all members of the Virtual Chassis configuration.

component (*UUID* | *serial number* | *all*) (QFabric systems only) (Optional) Display a list of core-dump files located on individual QFabric system device or on the entire QFabric system.

display-order (timestamp-sort | alphanumeric-sort) (QFabric systems only) (Optional) Display list of debug artifacts generated within the specified period—for example, within the last hour, within the last 20 minutes, or within the last 32 seconds—or according to their filename.

display-period (*hours* | *minutes* | *seconds*) (QFabric systems only) (Optional) Display core-dump files generated within the specified period—for example, within the last hour, within the last 20 minutes, or within the last 32 seconds.

kernel-crashinfo component (*UUID* | *serial number*) (QFabric systems only) (Optional) Display kernel crash information from the EEPROM on a QFabric system device.

local (EX4200 switches only) (Optional) Display system core files on the local Virtual Chassis member.

member *member-id*	(EX4200 switches only) (Optional) Display system core files on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.
node *node-name*	(Optional) (Junos OS Evolved only) Display system core files generated on the specified node.
re0	(Dual Routing Engines only) Display the core-dump files on re0.
re1	(Dual Routing Engines only) Display the coredump files on re1.
repository (core | log)	(QFabric systems only) (Optional) Specify either the core or log repository in which to view core-dump files.
routing-engine (backup | both | local | master | other)	(Dual routing engines only) Display a list of core-dump files for either the backup, local, primary, or other routing engine or both routing engines.
satellite [*fpc-slot-id* **| device-alias** *alias-name***]**	(Junos Fusion only)(Optional) Display system core files for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.
scc	(TX Matrix routers only) (Optional) Display system core files on the TX Matrix router (or switch-card chassis).
sfc	(TX Matrix Plus routers only) (Optional) Display system core files on the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level

view

Output Fields

Table 1616 on page 26787 describes the output fields for the `show system core-dumps` command. Output fields are listed in the approximate order in which they appear.

Table 1616: show system core-dumps Output Fields

Field Name	Field Description
Permissions	Read/write permissions for the file named.

Table 1616: show system core-dumps Output Fields *(Continued)*

Field Name	Field Description
Links	Number of links to the file.
Owner	Name of the file owner.
Group	Name of the group with file access.
File size	File size in bytes.
Modified	Last file modification date and time.
Path/filename	File path where the file resides and the filename. (MX Series routers only) When you display the core files for an MX Series Virtual Chassis, the `show system core-dumps` command does not display information about files pertaining to the relayd process.
`Repository scope:`	Repository where core-dump files and log files are stored. The core-dump files are located in the `core` repository, and the log files are located in the `log` repository. The default `Repository scope` is shared since both the `core` and `log` repositories are shared by all of the QFabric system devices.
`Repository head:`	Path to the top-level repository location.
`Repository name:`	Name of the repository: `core` or `log`.
`List of nodes for core repository:`	List of core-dump files associated with a particular QFabric system device located in the core repository.
`Node Group`	Name of the QFabric system device.
`Node Identifier`	UUID or serial number of the QFabric system device.

Table 1616: show system core-dumps Output Fields *(Continued)*

Field Name	Field Description
`Num`	Number of core-dump and log files.
`Model`	Model number of the QFabric system device.
`Usage`	Usage of the repository in megabytes.
`Total usage of core repository:`	Total usage of core-dump files associated with a particular QFabric system device located in the core repository. Usage is specified in megabytes and as a percentage.
`Total usage of log repository:`	Total usage of log files associated with a particular QFabric system device located in the log repository. Usage is specified in megabytes and as a percentage.
`List of nodes for core repository:`	List of core-dump files associated with a particular QFabric system device located in the core repository.
`List of nodes for log repository:`	List of log files associated with a particular QFabric system device located in the log repository.
`Filename`	Name of the core-dump file.
`Date`	Last core-dump file modification date and time.
`Size`	Size of the core-dump file.
`Core filename`	Filename of the core-dump file.
`Process name`	Name of the process that is generating a core-dump file or log file.
`Release`	Junos OS release.

Table 1616: show system core-dumps Output Fields *(Continued)*

Field Name	Field Description
`Build server`	Junos OS build server.
`Build date`	Junos OS build date.
`Stack trace`	Stack trace of the core-dump file.

Sample Output

show system core-dumps

This example shows the command output if core files exist.

```
user@host> show system core-dumps
-rw------- 1 root wheel 268369920 Jun 18 17:59 /var/crash/vmcore.0
-rw-rw---- 1 root field 3371008 Jun 18 17:53 /var/tmp/rpd.core.0
-rw-r--r-- 1 root wheel 27775914 Jun 18 17:59 /var/crash/kernel.0
```

show system core-dumps

This example shows the command output if core files do not exist.

```
user@host> show system core-dumps
/var/crash/*core*: No such file or directory
/var/tmp/*core*: No such file or directory
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
```

show system core-dumps routing-engine both

This example shows the command output if dual Routing Engines are present.

```
user@host> show system core-dumps routing-engine both
re0:
--------------------------------------------------------------------------
/var/crash/*core*: No such file or directory
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory

/var/tmp/cores:
total blocks: 496776
-rw-rw----  1 root  field   11910589 Nov 8  13:20 chassisd.core.0.201311081320
...

-rw-rw----  1 root  field   11737227 Oct 28 14:21 rpd.core-tarball.4.tgz.201310281421.3458162
total files: 10

re1:
--------------------------------------------------------------------------
/var/crash/*core*: No such file or directory
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory

/var/tmp/cores:
total blocks: 3178420
-rw-rw----  1 root  field   19039721 Nov 8  14:29 chassisd.core.0.201311081429.3485600.gz
-rw-rw----  1 root  field   19039793 Nov 8  14:37 chassisd.core.1.201311081437.3485599.gz
..

-rw-rw----  1 root  field   11710113 Oct 17 15:26 rpd.core-tarball.1.tgz.201310171526.3430028
```

show system core-dumps (SRX Series)

```
user@host> show system core-dumps

/var/crash/*core*: No such file or directory
-rw-r--r--  1 nobody wheel   1439949 Apr 24 10:38 /var/tmp/FPC0_PIC0.localhost.J-
UKERN.23421.1556127502.core.tgz
```

```
-rw-r--r--  1 nobody wheel   1435531 Apr 24 10:44 /var/tmp/FPC0_PIC0.localhost.J-
UKERN.24702.1556127821.core.tgz
-rw-r--r--  1 nobody wheel 288761042 Apr 24 10:32 /var/tmp/
FPC0_PIC0.localhost.flowd_spc3.elf.31620.1556126342.core.tgz
-rw-r--r--  1 nobody wheel     35082 Apr 24 10:47 /var/tmp/
FPC0_PIC0.localhost.tnp_hello.20972.1556128038.core.tgz
-rw-r--r--  1 nobody wheel     35367 Apr 24 10:49 /var/tmp/
FPC0_PIC0.localhost.tnp_hello.27233.1556128140.core.tgz
-rw-r--r--  1 nobody wheel     35372 Apr 24 11:32 /var/tmp/
FPC0_PIC1.localhost.tnp_hello.22289.1556130737.core.tgz
-rw-r--r--  1 nobody wheel     35357 Apr 24 10:51 /var/tmp/
FPC0_PIC1.localhost.tnp_hello.22492.1556128268.core.tgz
-rw-r--r--  1 nobody wheel     34812 Apr 24 11:33 /var/tmp/
FPC0_PIC1.localhost.tnp_hello.24235.1556130795.core.tgz
-rw-r--r--  1 nobody wheel     35383 Apr 24 11:18 /var/tmp/
FPC0_PIC1.localhost.tnp_hello.27070.1556129899.core.tgz
-rw-r--r--  1 nobody wheel     34675 Apr 24 11:18 /var/tmp/
FPC0_PIC1.localhost.tnp_hello.31621.1556129928.core.tgz
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
/var/jails/rest-api/tmp/*core*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
total files: 10
```

show system core-dumps (TX Matrix Plus Router)

```
user@host> show system core-dumps
sfc0-re0:
--------------------------------------------------------------------------
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory

/var/crash/cores:
total 8

/var/tmp/cores:
total 1627592
-rw-r--r--  1 root  field  535346090 May 15 07:36 rpd.core-tarball.0.090515.0736.tgz
-rw-r--r--  1 root  field  105632057 May 15 07:37 rpd.core-tarball.1.090515.0737.tgz
-rw-r--r--  1 root  field  101981681 May 15 07:38 rpd.core-tarball.2.090515.0738.tgz
-rw-r--r--  1 root  field   85854573 May 15 07:40 rpd.core-tarball.3.090515.0740.tgz
```

```
-rw-r--r--  1 root  field    4157845 May 15 08:18 rpd.core-tarball.4.090515.0818.tgz

lcc0-re0:
--------------------------------------------------------------------------
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory

/var/crash/cores:
total 8

/var/tmp/cores:
total 12

lcc1-re0:
--------------------------------------------------------------------------
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory

/var/crash/cores:
total 8

/var/tmp/cores:
total 10024
-rw-r--r--  1 root  field    1875794 Apr 22 15:47 chassisd.core-tarball.0.090422.1547.tgz
-rw-r--r--  1 root  field    1894183 Apr 22 19:02 chassisd.core-tarball.0.090422.1902.tgz
-rw-r--r--  1 root  field    1290240 Apr 26 16:01 ksyncd_1558.core.0.090426.1601

lcc2-re0:
--------------------------------------------------------------------------
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory

/var/crash/cores:
total 21124008
-rw-r--r--  1 root  wheel  1022376528 May 2  06:43 core-LCC2-EGFPC7.core.0.090502.0643
-rw-r--r--  1 root  wheel  1022376528 May 2  08:13 core-LCC2-EGFPC7.core.0.090502.0813
-rw-r--r--  1 root  wheel  1022376544 May 5  06:15 core-LCC2-EGFPC7.core.0.090505.0615
-rw-r--r--  1 root  wheel  1022376544 May 6  10:59 core-LCC2-EGFPC7.core.0.090506.1059
-rw-r--r--  1 root  wheel  1022376528 May 2  06:58 core-LCC2-EGFPC7.core.1.090502.0658
-rw-r--r--  1 root  wheel  754271232 May 5  06:33 core-LCC2-EGFPC7.core.1.090505.0633
-rw-r--r--  1 root  wheel  264897536 May 6  11:12 core-LCC2-EGFPC7.core.1.090506.1112
-rw-r--r--  1 root  wheel  1022376528 May 2  07:22 core-LCC2-EGFPC7.core.2.090502.0722
-rw-r--r--  1 root  wheel  163633152 May 5  06:52 core-LCC2-EGFPC7.core.2.090505.0652
```

```
-rw-r--r--  1 root  wheel  171312128 May 6  12:13 core-LCC2-EGFPC7.core.2.090506.1213
-rw-r--r--  1 root  wheel  1022376528 May 2  07:39 core-LCC2-EGFPC7.core.3.090502.0739
-rw-r--r--  1 root  wheel  1022376528 May 2  07:55 core-LCC2-EGFPC7.core.4.090502.0755
-rw-r--r--  1 root  wheel  427277312 May 7  04:47 core-LCC2-STFPC4.core.0.090507.0447
-rw-r--r--  1 root  wheel  419609600 May 7  04:47 core-LCC2-STFPC5.core.0.090507.0447
-rw-r--r--  1 root  wheel  432356352 May 7  04:47 core-LCC2-STFPC6.core.0.090507.0447

/var/tmp/cores:
total 2568
-rw-r--r--  1 root  field    1290240 May 14 14:26 ksyncd_1540.core.0.090514.1426
...
```

show system core-dumps (QFX3500 Switch)

```
user@switch> show system core-dumps
/var/crash/*core*: No such file or directory
-rw-rw----  1 root  field    1545143 Jun 4   2012 /var/tmp/pafxpc.core.0.gz
-rw-rw----  1 root  field    1545146 Jun 4   2012 /var/tmp/pafxpc.core.1.gz
-rw-rw----  1 root  field    1545141 Jun 4   2012 /var/tmp/pafxpc.core.2.gz
-rw-rw----  1 root  field    1545146 Jun 4   2012 /var/tmp/pafxpc.core.3.gz
-rw-rw----  1 root  field    1545142 Jun 5   2012 /var/tmp/pafxpc.core.4.gz
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
total 5
```

show system core-dumps (QFabric Systems)

```
user@switch> show system core-dumps
Repository scope: shared
Repository head: /pbdata/export
List of nodes for core repository: /pbdata/export/rdumps/
Node Group      Node Identifier                          Num  Model        Usage

DG-0            BCF7208D-E44F-E011-802F-4171BAAC781D    0    qfx3100      0M
FM-0            73747cd8-0710-11e1-b6a4-00e081c5297e    0    fx-jvre      0M
DRE-0           77116f18-0710-11e1-a2a0-00e081c5297e    0    fx-jvre      0M
NW-NG-0         BBAK0394                                0    qfx3500      0M
NW-NG-0         cd78871a-0710-11e1-878e-00e081c5297e    0    fx-jvre      0M
NW-NG-0         d0afda1e-0710-11e1-a1d0-00e081c5297e    0    fx-jvre      0M
```

```
FC-0            d31ab7a6-0710-11e1-ad1b-00e081c5297e    0    fx-jvre      0M
FC-1            d4d0f254-0710-11e1-90c3-00e081c5297e    0    fx-jvre      0M
IC-WS001        WS001                                   0    -            -
IC-WS001        WS001/YW3803                            0    qfxc08-3008  0M
IC-WS001        WS001/YN5999                            0    qfxc08-3008  0M
node-device1    BBAK0372                                0    qfx3500      0M
node-device1    EE3093                                  0    qfx3500      0M
Total usage of core repository:0M of 70000M (0.0%)

List of nodes for log repository: /pbdata/export/rlogs/
Node Group      Node Identifier                         Num  Model        Usage

DG-0            BCF7208D-E44F-E011-802F-4171BAAC781D    0    qfx3100      0M
FM-0            73747cd8-0710-11e1-b6a4-00e081c5297e    1    fx-jvre      0M
DRE-0           77116f18-0710-11e1-a2a0-00e081c5297e    1    fx-jvre      0M
NW-NG-0         BBAK0394                                1    qfx3500      0M
NW-NG-0         cd78871a-0710-11e1-878e-00e081c5297e    1    fx-jvre      0M
NW-NG-0         d0afda1e-0710-11e1-a1d0-00e081c5297e    3    fx-jvre      0M
FC-0            d31ab7a6-0710-11e1-ad1b-00e081c5297e    1    fx-jvre      0M
FC-1            d4d0f254-0710-11e1-90c3-00e081c5297e    1    fx-jvre      0M
IC-WS001        WS001                                   0    -            -
IC-WS001        WS001/YN5999                            1    qfxc08-3008  0M
IC-WS001        WS001/YW3803                            1    qfxc08-3008  0M
node-device1    BBAK0372                                1    qfx3500      0M
node-device1    EE3093                                  1    qfx3500      0M
Total usage of log repository:0M of 70000M (0.0%)
```

show system core-dumps component serial number display-order alphanumeric-sort repository core (QFabric Systems)

```
user@switch> show system core-dumps component BBAK8891 display-order alphanumeric-sort
repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
List of core dumps for component BBAK8891
Repository location: /pbdata/export/rdumps/BBAK8891
Filename                                Date                      Size

eswd.core.0.1361.11172011214257.gz      Nov 17 21:43:10 2011      4779553
eswd.core.1.80267.11172011214514.gz     Nov 17 21:45:19 2011      3541648
```

```
eswd.core.2.80682.11172011214535.gz        Nov 17 21:45:43 2011      2156683
vccpd.core.0.1195.11182011151131.gz        Nov 18 15:11:35 2011      375617
Number of core dumps in repository:4
```

show system core-dumps display-period (QFabric Systems)

```
user@switch> show system core-dumps display-period 24h
 show system core-dumps display-period 24h
Repository scope: shared
Repository head: /pbdata/export
List of core dumps at repository: /pbdata/export/rdumps
Delta timespec: Last 24h
Component: BBAK8273
Filename                                   Size                      Date

vccpd.core.0.1195.11182011151131.gz        Nov 18 15:11:35 2011      375794
Component: cedb7b0e-0025-11e1-9a5f-00e081c52990
Filename                                   Size                      Date

vccpd.core.0.1461.11182011151131.gz        Nov 18 15:11:31 2011      120951
Component: ee19c4f8-0025-11e1-aef6-00e081c52990
Filename                                   Size                      Date

vccpd.core.0.1462.11182011151131.gz        Nov 18 15:11:31 2011      109420
Component: BBAK8281
Filename                                   Size                      Date

vccpd.core.0.1196.11182011151131.gz        Nov 18 15:11:36 2011      375373
Component: BBAK8891
Filename                                   Size                      Date

vccpd.core.0.1195.11182011151131.gz        Nov 18 15:11:35 2011      375617
Component: BBAK8276
Filename                                   Size                      Date

vccpd.core.0.1196.11182011151131.gz        Nov 18 15:11:35 2011      375350
Component: BBAK8868
Filename                                   Size                      Date

vccpd.core.0.1196.11182011151130.gz        Nov 18 15:11:34 2011      376211
Component: BBAK8835
```

```
Filename                                    Size                      Date

vccpd.core.0.1195.11182011151130.gz         Nov 18 15:11:35 2011      375700
Component: BBAK8283
Filename                                    Size                      Date

vccpd.core.0.1195.11182011151131.gz         Nov 18 15:11:36 2011      368298
Component: YW3781/YW3781
Filename                                    Size                      Date

vccpd.core.0.1220.11182011151131.gz         Nov 18 15:11:38 2011      380002
Component: 09726be2-0026-11e1-82d9-00e081c52990
Filename                                    Size                      Date

vccpd.core.0.1461.11182011151130.gz         Nov 18 15:11:31 2011      119965
Component: BBAK8309
Filename                                    Size                      Date

vccpd.core.0.1196.11182011151131.gz         Nov 18 15:11:36 2011      378930
Component: 303d476a-0026-11e1-abf4-00e081c52990
Filename                                    Size                      Date

vccpd.core.0.1460.11182011151131.gz         Nov 18 15:11:31 2011      118385
Component: YW3798/YW3798
Filename                                    Size                      Date

vccpd.core.0.1219.11182011151131.gz         Nov 18 15:11:36 2011      380455
List of log dumps at repository: /pbdata/export/rlogs
Delta timespec: Last 24h
Component: BBAK8273
Filename                                    Size                      Date

vccpd.tarball.0.1195.11182011151138.tgz     Nov 18 15:11:39 2011      20415
Component: cedb7b0e-0025-11e1-9a5f-00e081c52990
Filename                                    Size                      Date

vccpd.tarball.0.1461.11182011151131.tgz     Nov 18 15:11:33 2011      19651
Component: ee19c4f8-0025-11e1-aef6-00e081c52990
Filename                                    Size                      Date

vccpd.tarball.0.1462.11182011151133.tgz     Nov 18 15:11:36 2011      24650
Component: BBAK8281
Filename                                    Size                      Date
```

```
vccpd.tarball.0.1196.11182011151137.tgz    Nov 18 15:11:41 2011      19445
Component: BBAK8891
Filename                                   Size                      Date

vccpd.tarball.0.1195.11182011151138.tgz    Nov 18 15:11:41 2011      21916
Component: BBAK8276
Filename                                   Size                      Date

vccpd.tarball.0.1196.11182011151137.tgz    Nov 18 15:11:39 2011      20461
Component: BBAK8868
Filename                                   Size                      Date

vccpd.tarball.0.1196.11182011151137.tgz    Nov 18 15:11:41 2011      21924
Component: BBAK8835
Filename                                   Size                      Date

vccpd.tarball.0.1195.11182011151137.tgz    Nov 18 15:11:39 2011      19424
Component: BBAK8283
Filename                                   Size                      Date

vccpd.tarball.0.1195.11182011151138.tgz    Nov 18 15:11:42 2011      31186
Component: YW3781/YW3781
Filename                                   Size                      Date

vccpd.tarball.0.1220.11182011151141.tgz    Nov 18 15:11:45 2011      27565
Component: 09726be2-0026-11e1-82d9-00e081c52990
Filename                                   Size                      Date

vccpd.tarball.0.1461.11182011151130.tgz    Nov 18 15:11:34 2011      19613
Component: BBAK8309
Filename                                   Size                      Date

vccpd.tarball.0.1196.11182011151138.tgz    Nov 18 15:11:46 2011      50362
Component: 303d476a-0026-11e1-abf4-00e081c52990
Filename                                   Size                      Date

vccpd.tarball.0.1460.11182011151133.tgz    Nov 18 15:11:33 2011      19360
Component: YW3798/YW3798
Filename                                   Size                      Date

vccpd.tarball.0.1219.11182011151140.tgz    Nov 18 15:11:49 2011      24473
```

show system core-dumps kernel-crashinfo component serial number (QFabric Systems)

```
user@switch> show system core-dumps kernel-crashinfo component A0001/YA0197
Node: A0001/YA0197

Information about previous kernel crash:

-- Kernel panic data --

Panic string: kdb_sysctl_panic
System uptime: 3 day 20 hr 59 min 40 sec Kernel crash time: 2011-11-15 Wed 15:25:17 Kernel build
linkstamp: JUNOS 11.3I #0: 2011-11-10 20:42:27 UTC

-- Stacktrace of panicing context --
Processor 1 (crash monarch):
savectx+0x0 (c9552800,80214efc,802a7fbc,c88ad05c) ra 801b93a8 sz 0
kdm_kcore_save_crashinfo+0x254 (c9552800,0,802a7fbc,c88ad05c) ra 801b9f44 sz 784
kdm_kcore_kern_panic_event_handler+0x4b0 (c9552800,0,802a7fbc,c88ad05c) ra 8022a9b8 sz 88
panic+0x1d0 (c9552800,0,4,77fed534) ra 802540c0 sz 56
kdb_sysctl_panic+0x70 (c9552800,0,4,77fed534) ra 80237e58 sz 40 sysctl_root+0x12c
(c9552800,0,4,e8bc5cf8) ra 80238e50 sz 48
userland_sysctl+0x164 (c9552800,0,4,e8bc5cf8) ra 8023956c sz 104
__sysctl+0xe4 (c9552800,0,4,e8bc5cf8) ra 806d62e8 sz 160
trap+0xe1c (c9552800,0,4,e8bc5cf8) ra 80896e68 sz 128
MipsUserGenException+0x1a4 (c9552800,0,4,405cd12c) ra 0 sz 0
pid 82340, process: sysctl

Processor 0:
restoreintr+0x14 (1,81bca820,3,0) ra 806cdc3c sz 0
spinlock_exit+0x30 (1,81bca820,3,0) ra 8025d354 sz 24
sleepq_release+0x64 (1,81bca820,3,0) ra 8025e670 sz 24
sleepq_timeout+0x224 (1,81bca820,3,0) ra 80240294 sz 48
softclock+0x434 (1,81bca820,3,0) ra 802067f8 sz 80
ithread_loop+0x244 (1,81bca820,3,0) ra 80200e28 sz 64 fork_exit+0xc0 (1,81bca820,3,0) ra
80897c28 sz 48
MipsNMIException+0x34 (1,81bca820,3,0) ra 0 sz 0
pid 82340, process: sysctl

Processor 2:
cpu_idle+0x20 (80960000,51bbc,2031df,81bca1b8) ra 80204948 sz 24 idle_proc+0x130
(80960000,51bbc,2031df,81bca1b8) ra 80200e28 sz 56 fork_exit+0xc0
(80960000,51bbc,2031df,81bca1b8) ra 80897c28 sz 48
```

```
MipsNMIException+0x34 (80960000,51bbc,2031df,81bca1b8) ra 0 sz 0
pid 82340, process: sysctl

Processor 3:
cpu_idle+0x20 (80960000,51bbc,2038df,81bca300) ra 80204948 sz 24 idle_proc+0x130
(80960000,51bbc,2038df,81bca300) ra 80200e28 sz 56 fork_exit+0xc0
(80960000,51bbc,2038df,81bca300) ra 80897c28 sz 48
MipsNMIException+0x34 (80960000,51bbc,2038df,81bca300) ra 0 sz 0
pid 82340, process: sysctl

Processor 4:
cpu_idle+0x20 (80960000,51bbc,2037df,81bca448) ra 80204948 sz 24 idle_proc+0x130
(80960000,51bbc,2037df,81bca448) ra 80200e28 sz 56 fork_exit+0xc0
(80960000,51bbc,2037df,81bca448) ra 80897c28 sz 48
MipsNMIException+0x34 (80960000,51bbc,2037df,81bca448) ra 0 sz 0
pid 82340, process: sysctl

Processor 5:
restoreintr+0x14 (1,51bbc,203edf,81bca590) ra 806cdc3c sz 0
spinlock_exit+0x30 (1,51bbc,203edf,81bca590) ra 80204a34 sz 24 idle_proc+0x21c
(1,51bbc,203edf,81bca590) ra 80200e28 sz 56 fork_exit+0xc0 (1,51bbc,203edf,81bca590) ra 80897c28
sz 48
MipsNMIException+0x34 (1,51bbc,203edf,81bca590) ra 0 sz 0
pid 82340, process: sysctl

Processor 6:
cpu_idle+0x20 (80960000,51bbc,205cdf,81bca6d8) ra 80204948 sz 24 idle_proc+0x130
(80960000,51bbc,205cdf,81bca6d8) ra 80200e28 sz 56 fork_exit+0xc0
(80960000,51bbc,205cdf,81bca6d8) ra 80897c28 sz 48
MipsNMIException+0x34 (80960000,51bbc,205cdf,81bca6d8) ra 0 sz 0
pid 82340, process: sysctl

Processor 7:
lockmgr+0x5ac (c97e8484,c8dd9800,0,c8dd9800) ra 8c11c81c sz 48
sal_sem_take+0x134 (c97e8484,c8dd9800,0,c8dd9800) ra 8c351108 sz 56 _bcm_esw_linkscan_thread
+0x45c (c97e8484,c8dd9800,0,c8dd9800) ra 8c11cdb4 sz 104
sal_thread_start_wrap+0x74 (c97e8484,c8dd9800,0,c8dd9800) ra 80200e28 sz 32 fork_exit+0xc0
(c97e8484,c8dd9800,0,c8dd9800) ra 80897c28 sz 48
MipsNMIException+0x34 (c97e8484,c8dd9800,0,c8dd9800) ra 0 sz 0
pid 82340, process: sysctl
-- End of stacktrace --
```

show system core-dumps repository core (QFabric Systems)

```
user@switch> show system core-dumps repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
List of nodes for core repository: /pbdata/export/rdumps/
Node Group      Node Identifier                          Num  Model        Usage

DG-0            BCF7208D-E44F-E011-802F-4171BAAC781D    0    qfx3100      0M
FM-0            73747cd8-0710-11e1-b6a4-00e081c5297e    0    fx-jvre      0M
DRE-0           77116f18-0710-11e1-a2a0-00e081c5297e    0    fx-jvre      0M
NW-NG-0         BBAK0394                                0    qfx3500      0M
NW-NG-0         cd78871a-0710-11e1-878e-00e081c5297e    0    fx-jvre      0M
NW-NG-0         d0afda1e-0710-11e1-a1d0-00e081c5297e    0    fx-jvre      0M
FC-0            d31ab7a6-0710-11e1-ad1b-00e081c5297e    0    fx-jvre      0M
FC-1            d4d0f254-0710-11e1-90c3-00e081c5297e    0    fx-jvre      0M
IC-WS001        WS001                                   0    -            -
IC-WS001        WS001/YW3803                            0    qfxc08-3008  0M
IC-WS001        WS001/YN5999                            0    qfxc08-3008  0M
node-device1    BBAK0372                                0    qfx3500      0M
node-device1    EE3093                                  0    qfx3500      0M
Total usage of core repository:0M of 70000M (0.0%)
```

show system core-dumps repository log (QFabric Systems)

```
user@switch> show system core-dumps repository log
Repository scope: shared
Repository head: /pbdata/export
Repository name: log
List of nodes for log repository: /pbdata/export/rlogs/
Node Group      Node Identifier                          Num  Model        Usage

DG-0            BCF7208D-E44F-E011-802F-4171BAAC781D    0    qfx3100      0M
FM-0            73747cd8-0710-11e1-b6a4-00e081c5297e    1    fx-jvre      0M
DRE-0           77116f18-0710-11e1-a2a0-00e081c5297e    1    fx-jvre      0M
NW-NG-0         BBAK0394                                1    qfx3500      0M
NW-NG-0         cd78871a-0710-11e1-878e-00e081c5297e    1    fx-jvre      0M
NW-NG-0         d0afda1e-0710-11e1-a1d0-00e081c5297e    3    fx-jvre      0M
FC-0            d31ab7a6-0710-11e1-ad1b-00e081c5297e    1    fx-jvre      0M
```

```
FC-1            d4d0f254-0710-11e1-90c3-00e081c5297e    1    fx-jvre      0M
IC-WS001        WS001                                   0    -            -
IC-WS001        WS001/YN5999                            1    qfxc08-3008  0M
IC-WS001        WS001/YW3803                            1    qfxc08-3008  0M
node-device1    BBAK0372                                1    qfx3500      0M
node-device1    EE3093                                  1    qfx3500      0M
Total usage of log repository:0M of 70000M (0.0%)
```

Release Information

Command introduced before Junos OS Release 8.5.

`sfc` option introduced in Junos OS Release 9.6 for the TX Matrix Plus router.

`re0`, `re1`, and `routing-engine` options introduced for dual Routing Engines in Junos OS Release 13.1.

`satellite` option introduced in Junos OS Release 14.2R3.

`core-file-info` option is deprecated in Junos OS Release 16.1R3.

show system core dumps (Junos OS Evolved)

IN THIS SECTION

- Syntax | 26802
- Description | 26803
- Required Privilege Level | 26803
- Output Fields | 26803
- Sample Output | 26803

Syntax

```
show system core-dumps
<node node-name>
```

Description

Show core files on all routers or switches running Junos OS Evolved. You use this command to show a list of system core files created when the device has failed, which can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification date, and path and filename.

> **NOTE**: For Junos OS Evolved, if dual Routing Engines are present, the command lists the core-dump files for both Routing Engines.

Required Privilege Level

view

Output Fields

The command displays a list of core-dump files. If a node does not have any core-dump files, then the command displays just the node name.

Sample Output

show system core-dumps (Dual-Routing Engine System, with Core Dump on only One Routing Engine)

```
user@host> show system core-dumps
re0:
--------------------------------------------------------------------------

re1:
--------------------------------------------------------------------------
-rw-r--r--  1 root  root    10389341 Mar 16 00:38 /var/core/re1/
agentd.re.re1.19293.2021_03_16.00_37_32.tar.gz
total files: 1
```

The output shows that there aren't any core-dump files on node RE0, but node RE1 has one core-dump file.

show system core-dumps (Dual-Routing Engine System, with Core Dumps on both Routing Engines)

```
user@host-re0> show system core-dumps
re0:
--------------------------------------------------------------------------
-rw-r--r--  1 root  root    52340949 Apr 13 11:01 /var/core/fpc0/
hwdfpc.fpc_x86_64.fpc0.14522.2021_04_13.10_59_10.tar.gz
total files: 1

re1:
--------------------------------------------------------------------------
-rw-r--r--  1 root  root    32432932 Apr 13 13:01 /var/core/re1/
imgd.re.re1.11040.2021_04_13.12_59_45.tar.gz
-rw-r--r--  1 root  root   346853497 Apr 8  10:52 /var/core/re1/
rpdagent.re.re1.17935.2021_04_08.10_46_34.tar.gz
-rw-r--r--  1 root  root   369435949 Apr 8  10:58 /var/core/re1/
rpdagent.re.re1.1908.2021_04_08.10_52_22.tar.gz
-rw-r--r--  1 root  root   192094114 Apr 8  11:00 /var/core/re1/
rpdagent.re.re1.5148.2021_04_08.10_56_18.tar.gz
-rw-r--r--  1 root  root   214337055 Apr 8  10:51 /var/core/re1/
rpdagent.re.re1.17935.2021_04_08.10_46_34/rpd-agent_trace.tar.gz
total files: 5
```

show system cpu (JDM)

IN THIS SECTION

- Syntax | 26805
- Description | 26805
- Required Privilege Level | 26805
- Output Fields | 26805
- Sample Output | 26805
- Release Information | 26806

Syntax

```
show system cpu
```

Description

Display the CPUs and their usage status.

Required Privilege Level

View

Output Fields

Table 1617 on page 26805 describes the output fields for the `show system cpu` command. Output fields are listed in the approximate order in which they appear.

Table 1617: show system cpu Output Fields

Field Name	Field Description
VNF	Name of the GNF.
State	GNF status. For example, “Running”.
CPU-Id(s)	List of CPU IDs.
Usage	CPU usage percentage.
Qemu Pid	Qemu process ID.

Sample Output

show system cpu (External Server Model)

```
user@jdm> show system cpu
```

```
VNF CPU Utilization and Allocation Information
-------------------------------------------------------------------------------
                                                                        -------------
VNF                                      CPU-Id(s)               Usage  Qemu
Pid
State
---------------------------------------- ----------------------- ------
--------
 -----------
test                                     4,5,6,7,8,9,10,11       5.0%
32392
 Running

Free CPUs      : 12,13,14,15
Host Isolcpu(s): 2-15
Emulator Pins  : 2-3
```

show system cpu (In-chassis Junos Node Slicing)

```
user@jdm> show system cpu
> VNF CPU Utilization and Allocation Information
 ---------------------------------------------------------------------------------------------
 VNF                                      CPU-Id(s)               Usage  Qemu Pid  State
 ---------------------------------------- ----------------------- ------ --------  -----------
 gnf-bkhl                                 2                       102%   21045     Running
 gnf-core                                 3,4                     101%   22205     Running

 Free CPU-id(s) : 5
 Host Isolcpu(s): 1-5
 Emulator Pins  : 0
```

For in-chassis Junos node slicing, it is normal to see 100 percentage or higher CPU usage in the `Usage` column.

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | 15481

request server authenticate-peer-server | 14901

show system data-models

IN THIS SECTION

- Syntax | 26807
- Description | 26807
- Options | 26807
- Required Privilege Level | 26808
- Sample Output | 26808
- Release Information | 26808

Syntax

```
show system data-models model xpath-cli-command-mapping xpath
```

Description

Display information for the requested data model.

Options

model	Specify the data model for the requested operation. Models include: • `genstate`—YANG data models for operational state.
xpath-cli-command-mapping *path*	For genstate models, retrieve the operational mode command that generates the state data associated with the specified resource path.

Required Privilege Level

view

Sample Output

show system data-models genstate xpath-cli-command-mapping

```
{master}
user@host> show system data-models genstate xpath-cli-command-mapping /genstate/system-
information/os-version

Genstate xpath to CLI command mapping information (sometimes multiple commands might be mapped):

> show system information
```

Release Information

Command introduced in Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

Junos Genstate YANG Data Models

show system directory-usage

IN THIS SECTION

- Syntax | 26809
- Syntax (EX Series) | 26809
- Syntax (TX Matrix Router) | 26809
- Syntax (TX Matrix Plus Router) | 26810
- Syntax (MX Series Router) | 26810

- Syntax (QFX Series and OCX Series) | 26810
- Description | 26810
- Options | 26810
- Required Privilege Level | 26812
- Output Fields | 26812
- Sample Output (Junos OS) | 26812
- Sample Output (Junos OS Evolved) | 26814
- Release Information | 26814

Syntax

```
show system directory-usage
<depth number>
<path>
```

Syntax (EX Series)

```
show system directory-usage
<all-members>
<depth number>
<local>
<member member-id>
<path>
```

Syntax (TX Matrix Router)

```
show system directory-usage
<all-chassis | all-lcc | lcc number | scc>
<depth number>
<path>
```

Syntax (TX Matrix Plus Router)

```
show system directory-usage
<all-chassis | all-lcc | lcc number | sfc number>
<depth number>
<path>
```

Syntax (MX Series Router)

```
show system directory-usage
<all-members>
<depth number>
<local>
<member member-id>
<path>
```

Syntax (QFX Series and OCX Series)

```
show system directory-usage
<depth number>
<path>
<infrastructure name>
<interconnect-device name>
<node-group name>
```

Description

Display directory usage information.

Options

none Display all directory usage information.

all-chassis (TX Matrix routers and TX Matrix Plus routers only) (Optional) Display directory usage information about all the T640 routers (in a routing matrix based on a TX Matrix router). Display directory usage information about all the T1600 or T4000 routers (in a routing matrix based on a TX Matrix Plus router) in the chassis.

all-lcc (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display directory information for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display directory information for all connected T1600 or T4000 LCCs.

all-members (EX4200 switches and MX Series routers only) (Optional) Display directory information for all members of the Virtual Chassis configuration.

depth *number* (Optional) Depth of the directory to traverse. This option is useful when you want to limit the output shown for a large file system.

infrastructure *name* (QFabric systems only) (Optional) Display directory information for the fabric control Routing Engines and fabric manager Routing Engines.

interconnect-device *name* (QFabric systems only) (Optional) Display directory information for the Interconnect device.

node-group *name* (QFabric systems only) (Optional) Display directory information for the Node group.

lcc *number* (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display directory information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display directory information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (EX4200 switches and MX Series routers only) (Optional) Display directory information for the local Virtual Chassis member.

member *member-id* (EX4200 switches and MX Series routers only) (Optional) Display directory information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

path	(Optional) Path or root directory to traverse.
scc	(TX Matrix router only) (Optional) Display directory information for the TX Matrix router (or switch-card chassis).
sfc *number*	(TX Matrix Plus routers only) (Optional) Display directory information for the TX Matrix Plus router. Replace *number* with **0**.

Required Privilege Level

view

Output Fields

Table 1618 on page 26812 describes the output fields for the `show system directory-usage` command. Output fields are listed in the approximate order in which they appear.

Table 1618: show system directory-usage Output Fields

Field Name	Field Description
bytes	Number of bytes used by files in a directory.
directory-name	Name of the directory.

Sample Output (Junos OS)

show system directory-usage scc (TX Matrix Router)

```
user@host> show system directory-usage /var/tmp
scc
          /var/tmp
1.0K      /var/tmp/vi.recover
2.0K      /var/tmp/instmp.tPMk8u
1.0K      /var/tmp/install
          /var/tmp/instmp.GUMpur
4.8M      /var/tmp/instmp.GUMpur/packages
6.4M      /var/tmp/troy1
297M      /var/tmp/dsw
```

```
          /var/tmp/pkg_tmp.2073
83K       /var/tmp/pkg_tmp.2073/bin
          /var/tmp/instmp.oMIDbl
89K       /var/tmp/instmp.oMIDbl/bin
          /var/tmp/instmp.byhMjR
4.6M      /var/tmp/instmp.byhMjR/packages
          /var/tmp/instmp.6fqHf3
1.7M      /var/tmp/instmp.6fqHf3/packages
          /var/tmp/instmp.mljECe
4.6M      /var/tmp/instmp.mljECe/packages
```

show system directory-usage sfc (TX Matrix Plus Router)

```
user@host show system directory-usage /var/tmp
sfc 0
sfc0-re0:
--------------------------------------------------------------------------
          /var/tmp
46K       /var/tmp/gres-tp
          /var/tmp/sec-download
2.0K      /var/tmp/sec-download/sub-download
2.0K      /var/tmp/vi.recover
2.0K      /var/tmp/install
795M      /var/tmp/cores
766K      /var/tmp/pr440594
```

show system directory-usage (QFX Series and OCX Series)

```
user@host> show system directory-usage
          /var/tmp
30K       /var/tmp/gres-tp
2.0K      /var/tmp/rtsdb
2.0K      /var/tmp/vi.recover
2.0K      /var/tmp/install
2.0K      /var/tmp/pics
```

show system directory-usage (Relative Directory)

In this example, **example_dir** is a directory inside **/var/tmp/**, and the user ran this command inside **/var/tmp/**:

```
user@host> show system directory-usage example_dir
          example_dir
4.0K      example_dir/test_dir
```

Sample Output (Junos OS Evolved)

show system directory-usage (Relative Directory)

In this example, **example_dir** is a directory inside **/var/tmp/**, and the user ran this command inside **/var/tmp/**:

```
user@host> show system directory-usage example_dir
          /var/tmp/example_dir
4.0K      /var/tmp/example_dir/test_dir
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system download

IN THIS SECTION

- Syntax | 26815
- Description | 26815
- Options | 26815
- Required Privilege Level | 26815
- Output Fields | 26815
- Sample Output | 26816
- Release Information | 26817

Syntax

```
show system download <download-id>
```

Description

This command displays a brief summary of all the download instances along with their current state and extent of progress. If a `download-id` is provided, the command displays a detailed report of the particular download instance.

Options

- `download-id`—(Optional) The ID number of the download instance.

Required Privilege Level

view

Output Fields

Table 1619 on page 26816 lists the output fields for the `show system download` command. Output fields are listed in the approximate order in which they appear.

Table 1619: show system download Output Fields

Field Name	Field Description
ID	Displays the download identification number.
Status	Displays the state of a particular download.
Start Time	Displays the start time of a particular download.
Progress	Displays the percentage of a download that has been completed.
URL	Displays the URL from which the file was downloaded.

Sample Output

show system download

```
user@host> show system download
Download Status Information:
ID  Status      Start Time       Progress  URL
1   Active      May 4 06:28:36   5%        ftp://ftp-server//tftpboot/1m_file
2   Active      May 4 06:29:07   3%        ftp://ftp-server//tftpboot/5m_file
3   Error       May 4 06:29:22   Unknown   ftp://ftp-server//tftpboot/badfile
4   Completed   May 4 06:29:40   100%      ftp://ftp-server//tftpboot/smallfile
```

show system download 1

```
user@host> show system download 1

Download ID     : 1
Status          : Active
Progress        : 6%
URL             : ftp://ftp-server//tftpboot/1m_file
Local Path      : /var/tmp/1m_file
```

```
Maximum Rate     : 1k
Creation Time    : May 4 06:28:36
Scheduled Time   : May 4 06:28:36
Start Time       : May 4 06:28:37
Error Count      : 0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D35.

RELATED DOCUMENTATION

request system download start | 15076

show system errors

IN THIS SECTION

- Syntax | 26817
- Description | 26818
- Options | 26818
- Required Privilege Level | 26818
- Output Fields | 26818
- Sample Output | 26819
- Release Information | 26821

Syntax

```
show system errors
<cb slot| ccg slot | fan slot | fpc slot | psm slot | re slot | sib slot>
```

Description

Display information about faults in the system. You can display all errors or the errors for one system component. Use this command to understand about faults and their correlation with other events. First, top level root causes are listed, with board level faults followed by component level faults. Next, details for affected faults are listed.

The show output represents five faults, F1 through F5. F4 and F5 are top level faults, where F4 is affected by F1, F2, and F3; and F3 is affected by F1 and F2. The lowest level (leaf) faults, F1, F2, and F5, have no affected events.

> **NOTE**: For Junos OS Evolved, only the QFX5200 supports this command. For all other Junos OS Evolved platforms, use the "show system errors active" on page 26822, "show system errors count" on page 26829, "show system errors error-id" on page 26831, "show system errors fru" on page 26834, or "show system errors inactive" on page 26844 command.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1620 on page 26818 lists the output fields for the `show system errors` command. Output fields are listed in the approximate order in which they appear.

Table 1620: show system errors Output Fields

Field Name	Field Description
`Top level root causes`	Display of the top level faults with board level faults followed by component level faults.
`F`*x*	Fault number F1 to Fn, where F1 is the first fault and n is the last fault generated by the system.

Table 1620: show system errors Output Fields *(Continued)*

Field Name	Field Description
(*module*, *error-id*, *board-name*, *component-name*)	Information about the fault. Component level faults include the component name.
Group	Fault severity level is Fatal, Major, or Minor.
Scope	Affected scope of fault is System, Component, Board, or Link.
Corr-enabled	Correlation with fault is always enabled, Y.
Time	Time in the format yyyy-mm-dd hh:nn:ss.lll TMZ, where nn is minutes, lll is milliseconds, and TMZ is time zone.
Desc	Description of the fault.
Actions	List of errors that caused the fault.
Root-causes	List of faults that caused this error.
Affected	List of faults that correlate to this root cause.
Details for affected errors	Display the affected errors listed in top level faults.

Sample Output

show system errors

```
user@host> show system errors
Top level root-causes:
F4: {pciesw, 1, fpc0} Group: Fatal Scope: Board Corr-enabled: Y
Time: “2017-02-22 16:37:47.456 PST”
Desc: PCIe Switch Fatal AER Errors
Actions: Alarm: FPC_FATAL_ERRORS + FRU_FAULT
```

```
Root-causes: None
Affected:
F3: {hwd, 1, fpc0}
F1: {pechip, 1, fpc0, pechip0}
F2: {pechip, 1, fpc0, pechip3}
F5: {pfchip, 3, sib0, pfchip5} Group: Major Scope: Component Corr-enabled: Y
Time: “2017-02-22 18:37:47.456 PST”
Desc: Midplane link errors
Actions: Alarm: ASIC_FABRIC_LINK_ERRORS
Affected: None
Details for Affected Errors:
F3: {hwd, 1, fpc0}, Group: Fatal Scope: Board Corr-enabled: Y
Time: “2017-02-22 16:37:47.856 PST”
Desc: FPC Fault
Root-causes: F4 : { pciesw, 1, fpc0}
Affected:
F1: {pechip, 1, fpc0, pechip0}
F2: {pechip, 1, fpc0, pechip3}
F1: {pechip, 3, fpc0, pechip0}, Group: Fatal Scope: Component Corr-enabled: Y
Time: “2017-02-22 16:37:48.500 PST”
Desc: PIO Fault
Root-causes:
F4 : {pciesw, 1, fpc0}
Affected: None
F2: {pechip, 3, fpc0, pechip1}, Group: Fatal Scope: Component Corr-enabled: Y
Time: “2017-02-22 16:37:48.600 PST”
Desc: PIO Fault
Root-causes:
F4 : {pciesw, 1, fpc0}
Affected: None
```

show system errors fpc 0

```
user@host> show system errors fpc 0
Top level root-causes:
F4: {pciesw, 1, fpc0} Group: Fatal   Scope: Board  Corr-enabled: Y
  Time:    “2017-02-22 16:37:47.456 PST”
  Desc:    PCIe Switch Fatal AER Errors
    Actions: Alarm: FPC_FATAL_ERRORS + FRU_FAULT
    Root-causes: None
    Affected:
```

```
      F3:    {hwd, 1, fpc0}
      F1:    {pechip, 1, fpc0, pechip0}
      F2:    {pechip, 1, fpc0, pechip3}

Details for Affected Errors:
F3: {hwd, 1, fpc0}, Group: Fatal   Scope: Board  Corr-enabled: Y
    Time: “2017-02-22 16:37:47.856 PST”
  Desc:    FPC Fault
  Root-causes: F4 : { pciesw, 1, fpc0}
  Affected:
      F1:    {pechip, 1, fpc0, pechip0}
      F2:    {pechip, 1, fpc0, pechip3}

F1: {pechip, 3, fpc0, pechip0}, Group: Fatal   Scope: Component  Corr-enabled: Y
    Time: “2017-02-22 16:37:48.500 PST”
    Desc: PIO Fault
  Root-causes:
      F4 :   {pciesw, 1, fpc0}
  Affected: None
F2: {pechip, 3, fpc0, pechip1}, Group: Fatal   Scope: Component  Corr-enabled: Y
    Time: “2017-02-22 16:37:48.600 PST”
    Desc: PIO Fault
  Root-causes:
      F4 :   {pciesw, 1, fpc0}
  Affected: None
```

show system errors sib 0

```
user@host> show system errors sib 0
Top level root-causes:
F5: {pfchip, 3, sib0, pfchip5} Group: Major   Scope: Component  Corr-enabled: Y
  Time:    “2017-02-22 18:37:47.456 PST”
  Desc:    Midplane link errors
    Actions: Alarm: ASIC_FABRIC_LINK_ERRORS
    Affected: None
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

show system errors history | 26841

show system errors active

IN THIS SECTION

- Syntax | 26822
- Description | 26822
- Options | 26823
- Required Privilege Level | 26823
- Output Fields | 26823
- Sample Output | 26824
- Release Information | 26829

Syntax

```
show system errors active
<detail [fru slot-number [scope error-scope] [category error-category]]>
<fru slot-number>
```

Description

Display information collected by the J-Insight fault monitoring feature. Specifically, display summary or detailed information about the active errors based on FRU, error scope, or error category.

NOTE: In PTX Series routers with Junos OS Evolved, the details of the Packet Forwarding Engine errors (reported through CMError), when set and cleared, are moved from the output of `show system errors active` command to the output of `show system errors inactive` command. However, the output of the `show system errors inactive detail` does not contain the details of the active FRU board errors that are cleared.

Options

none	Display a brief summary of the system error information for all applicable FRUs.
category *error-category*	(Optional) Display system error information based on error category. An error category categorizes errors into various subgroups under a specific error scope level. Values include: core, functional, io, memory, processing, storage, and switch.
detail	(Optional) Display detailed system error information.
fru slot-number	(Optional) Display system error information for a specific FRU. For devices running Junos OS, output displays error details for FPC FRUs. For devices running Junos OS Evolved, output displays error details for FPC and other components such as fan, PSM, CB, and chassis.
scope *error-scope*	(Optional) Display system error information based on error scope. An error scope provides a level of classification above error category. Values include: board, pfe, and scope-all.

Required Privilege Level

admin

Output Fields

Table 1621 on page 26823 list the output fields for the `show system errors active` command. Output fields are listed in the approximate order in which they appear.

Table 1621: show system errors active Output Fields

Field Name	Field Description
Error Name	Name of error.
Identifier	Each error is uniquely identified with an error ID that is represented as a Uniform Resource Identifier (URI).
Description	Description of the error.

Table 1621: show system errors active Output Fields *(Continued)*

Field Name	Field Description
State	State of the error. Values are: enabled or disabled.
Scope	Scope classification to which the error belongs. Values include board and pfe.
Category	Category subgroup under the scope level to which the error belongs. Values include: core, functional, io, memory, processing, storage, and switch.
Level	Severity level of the error.
Threshold	Configured threshold value. The associated detection and recovery actions are triggered when this value is exceeded.
Error Limit	The maximum number of times the error is reported.
Support	Support details for the error type.
Occur count	Number of times errors of a specific scope, category, and severity level has occurred.
Clear count	Number of times error instances have been cleared.
Last occurred (ms ago)	Amount of time (in milliseconds) passed since the error last occurred.

Sample Output

show system errors active

For devices running Junos OS, output displays error details for FPC FRUs. For devices running Junos OS Evolved, output displays error details for FPC and other components such as fan, PSM, CB, and chassis.

```
user@host> show system errors active
```

```
System Active Errors Information
CB 0
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
CHASSIS 0
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 5
Active Fatal Errors           : 0
FAN 0
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
FAN 1
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
FAN 2
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
FAN 3
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
FAN 4
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
FPC 0
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
FPC 1
---------------------------------
```

```
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
FPC 2
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
FPC 3
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
FPM 0
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
PDU 0
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
PICS 0
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
PICS 1
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
PSM 0
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
PSM 1
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
```

```
PSM 2
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
PSM 3
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
RE 0
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
SIB 0
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
SIB 1
---------------------------------
Active Minor Errors           : 0
Active Major Errors           : 0
Active Fatal Errors           : 0
```

show system errors active fpc-slot

```
user@host> show system errors active fpc-slot
0
System Active Errors Information
FPC 0
--------------------------------------------------
Active Minor Errors: 0
Active Major Errors: 1
Active Fatal Errors: 0
```

show system errors active detail

```
user@host> show system errors active detail

System Active Errors Detail Information
FPC 7
---------------------------------------------------------------
Error Name : btchip_temp_monitor_pfe_throttled_bandwidth
Identifier : /fpc/7/evo-cda-bt/0/cm/0/btchip/0/btchip_temp_monitor_pfe_throttled_bandwidth
Description : btchip_temp_monitor_pfe_throttled_bandwidth
State : enabled
Scope : pfe
Category : functional
Level : minor
Threshold : 10
Error limit : 30
Occur count : 1
Clear count : 0
Last occurred(ms ago) : 2021-07-07 18:32:43 PDT (211961 ms ago)
```

show system errors active detail (PTX series: PTX10004, PTX10008, and PTX10016)

```
user@host> show system errors active detail
System Active Errors Detail Information
CHASSIS 0
---------------------------------------------------------------
Error Name : fan_tray_removal
Identifier : /chassis/0/hwdre/0/cm/0/fan_tray/Fan Tray 0/fan_tray_removal
Description : Fan_tray_absent
State : disabled
Scope : board
Category : functional
Level : major
Threshold : 1
Error limit : 1
Support : No help info provided
Occur count : 1
Clear count : 0
Last occurred(ms ago) : 339112691
```

Release Information

Command introduced in Junos OS Release 18.2R1.

Command enhanced to include automatic temperature performance throttle and "btchip_temp_monitor_pfe_throttled_bandwidth" option error display in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

- show system errors count | 26829
- show system errors error-id | 26831
- show system errors fru | 26834

show system errors count

IN THIS SECTION

- Syntax | 26829
- Description | 26829
- Options | 26830
- Required Privilege Level | 26830
- Output Fields | 26830
- Sample Output | 26830
- Release Information | 26831

Syntax

```
show system errors count
```

Description

Display information collected by the J-Insight fault monitoring feature. Specifically, display information about the number of detected errors and recovery actions triggered based on error severity level.

Options

This command has no options.

Required Privilege Level

admin

Output Fields

Table 1622 on page 26830 lists the output fields for the `show system errors count` command. Output fields are listed in the approximate order in which they appear.

Table 1622: show system errors count Output Fields

Field Name	Field Description
Level	Severity level of the error. Values are: Minor, Major, or Fatal.
Occurred	Number of times errors of a specific severity level occurred.
Cleared	Number of times errors of a specific severity level were cleared.
Action-Taken	Number of times a recovery action was triggered for a specific severity level.

Sample Output

show system errors count

```
user@host> show system errors count
Level   Occurred   Cleared      Action-Taken
-------------------------------------------------------
 Minor:  0          0             0
 Major:  1          0             1
 Fatal:  0          0             0
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

show system errors active | 26822
show system errors error-id | 26831
show system errors fru | 26834

show system errors error-id

IN THIS SECTION

- Syntax | 26831
- Description | 26831
- Options | 26832
- Additional Information | 26832
- Required Privilege Level | 26832
- Output Fields | 26832
- Sample Output | 26833
- Release Information | 26834

Syntax

```
show system errors error-id error-id-uri
```

Description

Display information collected by the J-Insight fault monitoring feature. Specifically, display information about detected errors based on the error ID Uniform Resource Identifier (URI). For devices running Junos OS Evolved, output displays only errors that have occurred at least once in the system.

Options

This command has no options.

Additional Information

Required Privilege Level

admin

Output Fields

Table 1623 on page 26832 lists the output fields for the `show system errors error-id` command. Output fields are listed in the approximate order in which they appear.

Table 1623: show system errors error-id Output Fields

Field Name	Field Description
Error Name	Name of error.
Identifier	Each error is uniquely identified with an error ID that is represented as a Uniform Resource Identifier (URI).
Description	Description of the error.
State	State of the error. Values are: enabled or disabled.
Scope	Scope classification to which the error belongs. Values include board and pfe.
Category	Category subgroup under the scope level to which the error belongs. Values include: core, functional, io, memory, processing, storage, and switch.
Level	Severity level of the error.
Threshold	Configured threshold value. The associated detection and recovery actions are triggered when this value is exceeded.

Table 1623: show system errors error-id Output Fields *(Continued)*

Field Name	Field Description
Error Limit	The maximum number of times the error is reported.
Support	Support details for the error type.
Occur count	Number of times errors of a specific scope, category, and severity level has occurred.
Clear count	Number of times error instances have been cleared.
Last occurred (ms ago)	Amount of time (in milliseconds) passed since the error last occurred.

Sample Output

show system errors error-id

```
user@host> show system errors error-id "/chassis/0/hwdre/0/cm/0/fan_tray/Fan
Tray 0/fan_tray_removal"
System Errors Detail Information
CHASSIS 0
----------------------------------------------------------------
Error Name            : fan_tray_removal
Identifier            : /chassis/0/hwdre/0/cm/0/fan_tray/Fan Tray 0/fan_tray_removal
Description           : Fan_tray_absent
State                 : enabled
Scope                 : board
Category              : functional
Level                 : major
Threshold             : 1
Error limit           : 1
Support               : No help info provided
Occur count           : 1
Clear count           : 0
Last occurred(ms ago) : 84091182
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

show system errors active | 26822
show system errors count | 26829
show system errors fru | 26834

show system errors fru

IN THIS SECTION

- Syntax | 26834
- Description | 26834
- Options | 26835
- Required Privilege Level | 26835
- Output Fields | 26835
- Sample Output (Junos OS) | 26836
- Sample Output (Junos OS Evolved) | 26839
- Release Information | 26841

Syntax

```
show system errors fru detail [fru slot-number]
```

Description

Display information collected by the J-Insight fault monitoring feature. Specifically, display information about detected errors based on the FRU.

Options

none Display a brief summary of the system error information for the FRU.

detail (Optional) Display detailed system error information.

fru slot-number (Optional) Display system error information for a specific FRU. For devices running Junos OS, output displays error details for FPC FRUs. For devices running Junos OS Evolved, output displays error details for FPC and other components such as fan, PSM, CB, and chassis.

Required Privilege Level

admin

Output Fields

Table 1624 on page 26835 lists the output fields for the `show system errors fru` command. Output fields are listed in the approximate order in which they appear.

Table 1624: show system errors fru Output Fields

Field Name	Field Description
FRU	FRU identification number.
Scope	An error scope provides a level of classification above error category. Error scope values are: pfe and board.
Category	An error category categorizes errors into various subgroups under a specific error scope level. Values include: functional, io, memory, processing, storage, and switch.
Level	Severity level of the error.
Occurred	Number of times errors of a specific scope, category, and severity level has occurred.
Cleared	Number of times errors of a specific scope, category, and severity level were cleared.

Table 1624: show system errors fru Output Fields *(Continued)*

Field Name	Field Description
Threshold	Configured threshold value. The associated detection and recovery actions are triggered when this value is exceeded.
Action-Taken	Number of times a user-configured recovery action was triggered for errors of a specific scope, category, and severity level.
Action	Action that is triggered when the threshold value is exceeded.

Sample Output (Junos OS)

show system errors fru detail

```
user@host> show system errors fru detail
Fru     Scope   Category    Level   Occurred    Cleared     Threshold   Action-Taken    Action
FPC 0
        board
                functional  Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
PFE
                  memory    Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
PFE
                      io    Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
PFE
                 storage    Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
```

```
PFE
                  switch    Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
PFE
                processing  Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
PFE
        pfe
                functional  Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
PFE
                  memory    Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
PFE
                      io    Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
PFE
                 storage    Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
PFE
                  switch    Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
PFE
                processing  Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          DISABLE
PFE
```

show system errors fru detail (MX240, MX480, MX960, MX2008, MX2010, MX2020)

```
user@host> show system errors fru detail
Fru     Scope   Category    Level   Occurred    Cleared     Threshold   Action-Taken    Action
FPC 0
        board
                functional  Minor        0           0           1          0          LOG|CM
ALARM|
                            Major        0           0           1          0          GET
STATE|CM ALARM|
                            Fatal        0           0           1          0          DISABLE
PFE
                  memory    Minor        0           0           1          0          LOG|CM
ALARM|
                            Major        0           0           1          0          GET
STATE|CM ALARM|
                            Fatal        0           0           1          0          DISABLE
PFE
                      io    Minor        0           0           1          0          LOG|CM
ALARM|
                            Major        0           0           1          0          GET
STATE|CM ALARM|
                            Fatal        0           0           1          0          DISABLE
PFE
                 storage    Minor        0           0           1          0          LOG|CM
ALARM|
                            Major        0           0           1          0          GET
STATE|CM ALARM|
                            Fatal        0           0           1          0          DISABLE
PFE
                  switch    Minor        0           0           1          0          LOG|CM
ALARM|
                            Major        0           0           1          0          GET
STATE|CM ALARM|
                            Fatal        0           0           1          0          DISABLE
PFE
                processing  Minor        0           0           1          0          LOG|CM
ALARM|
                            Major        0           0           1          0          GET
STATE|CM ALARM|
                            Fatal        0           0           1          0          DISABLE
PFE
```

```
        pfe
                functional  Minor        0           0            1          0          LOG|CM
ALARM|
                            Major        0           0            1          0            LOG|
RESET PFE
```

Sample Output (Junos OS Evolved)

show system errors fru detail (PTX10003)

The following output has been shortened for clarity. For each part of a FRU, the full output displays any errors in the functional, io, memory, processing, storage, and switch categories, similar to the CB 0 FRU below.

```
user@host> show system errors fru detail

Fru     Scope   Category    Level   Occurred    Cleared     Threshold   Action-Taken    Action
CB 0
        board
                functional  Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          CM ALARM|
RESET
                      io    Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          CM ALARM|
RESET
                  memory    Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          CM ALARM|
RESET
                processing  Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
STATE|CM ALARM|
                            Fatal        0           0            1          0          CM ALARM|
RESET
                 storage    Minor        0           0           10          0            LOG|
                            Major        0           0            1          0          GET
```

```
STATE|CM ALARM|
                          Fatal        0           0            1          0          CM ALARM|
RESET
                  switch  Minor        0           0           10          0            LOG|
                          Major        0           0            1          0          GET
STATE|CM ALARM|
                          Fatal        0           0            1          0          CM ALARM|
RESET
CHASSIS 0
        board
...
FAN 0
        board
...
FAN 1
        board
...
FPC 0
        board

...
        pfe
...
FPC 1
        board
...

        pfe
...
FPM 0
        board
...
PDU 0
        board
...
PICS 0
        board
...
PICS 1
        board
...
PSM 0
        board
```

```
...
PSM 1
        board
...
RE 0
        board
 ...
SIB 0
        board
...
        switch
...
SIB 1
        board
...
        switch
...
```

Release Information

Command introduced in Junos OS Release 18.2R1.

Reset-pfe option added in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

- show system errors active | 26822
- show system errors count | 26829
- show system errors error-id | 26831

show system errors history

IN THIS SECTION

- Syntax | 26842
- Description | 26842

- Options | 26842
- Required Privilege Level | 26842
- Output Fields | 26842
- Sample Output | 26843
- Release Information | 26844

Syntax

```
show chassis errors history
<cb slot| ccg slot | fan slot | fpc slot | psm slot | re slot | sib slot>
```

Description

Display information about cleared faults in the error history buffer. You can display history for all errors or the errors for one system component. The error history is displayed in chronological order and includes a description of each PFE and FCHIP fault and when the fault was raised and cleared.

NOTE: For Junos OS Evolved, only the QFX5200 supports this command. For all other Junos OS Evolved platforms, use the show system errors active, show system errors count, show system errors error-id, or show system errors fru command.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1625 on page 26843 lists the output fields for the `show system errors history` command. Output fields are listed in the approximate order in which they appear.

Table 1625: show system errors history Output Fields

Field Name	Field Description
Fault (*module*, *error-id*, *board-name/component-name*, *PFE-or-FCHIP*)	Information about the fault.
Group	Fault severity level is Fatal, Major, or Minor.
Scope	Affected scope of fault is System, Component, Board, or Link.
Corr-enabled	Correlation with fault is always enabled, Y.
Raised	Time the fault was raised, in the format yyyy-mm-dd hh:nn:ss.lll TMZ, where nn is minutes, lll is milliseconds, and TMZ is time zone.
Desc	Description of the fault.
Cleared	Time the fault was cleared, in the format yyyy-mm-dd hh:nn:ss.lll TMZ, where nn is minutes, lll is milliseconds, and TMZ is time zone.

Sample Output

show system errors history

```
user@host> show system errors history
Fault: {pechip, 1487, /Chassis[0]/Fpc[4], Fchip[0]}
 Group:   Major             Scope: Link           Corr-enabled: Y
 Raised:  2017-04-19 18:18:48.652000 PDT
 Desc:    Fabric Down condition on PFE
 Cleared: 2017-04-19 18:18:49.474975 PDT
 Fault: {pechip, 1487, /Chassis[0]/Fpc[4], Fchip[1]}
 Group:   Major             Scope: Link           Corr-enabled: Y
 Raised:  2017-04-19 18:18:48.653000 PDT
 Desc:    Fabric Down condition on PFE
 Cleared: 2017-04-19 18:18:49.474668 PDT
```

```
  Fault: {pechip, 1487, /Chassis[0]/Fpc[4], Fchip[2]}
  Group:   Major             Scope: Link           Corr-enabled: Y
  Raised:  2017-04-19 18:18:48.654000 PDT
  Desc:    Fabric Down condition on PFE
  Cleared: 2017-04-19 18:18:49.474245 PDT
  Fault: {pechip, 1487, /Chassis[0]/Fpc[4], Fchip[3]}
  Group:   Major             Scope: Link           Corr-enabled: Y
  Raised:  2017-04-19 18:18:48.654000 PDT
  Desc:    Fabric Down condition on PFE
  Cleared: 2017-04-19 18:18:49.210875 PDT
  Fault: {pechip, 1143, /Chassis[0]/Fpc[4], Pfe[0]}
  Group:   Major             Scope: Component      Corr-enabled: N
  Raised:  2017-04-19 18:18:57.533000 PDT
  Desc:    hostif_local_int_wnack0
 Fault: {pechip, 1487, /Chassis[0]/Fpc[4], Fchip[0]}
  Group:   Major             Scope: Link           Corr-enabled: Y
  Raised:  2017-04-19 19:45:20.949000 PDT
  Desc:    Fabric Down condition on PFE
  Cleared: Active
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

show system errors | 26817

show system errors inactive

IN THIS SECTION

- Syntax | 26845
- Description | 26845
- Options | 26845
- Required Privilege Level | 26845

- Output Fields | 26845
- Sample Output | 26847
- Release Information | 26851

Syntax

```
show system errors inactive
<detail>
```

Description

Display information collected by the J-Insight fault monitoring feature. Specifically, display summary or detailed information about the inactive errors in the system. This commands shows the information about errors that had occurred and were then cleared.

Options

none Display a brief summary of the system error information for all applicable FRUs.

detail (Optional) Display detailed system error information.

Required Privilege Level

admin

Output Fields

Table 1626 on page 26846 list the output fields for the `show system errors inactive` command. Output fields are listed in the approximate order in which they appear.

Table 1626: show system errors inactive Output Fields

Field Name	Field Description
Error Name	Name of error.
Identifier	Each error is uniquely identified with an error ID that is represented as a Uniform Resource Identifier (URI).
Description	Description of the error.
State	State of the error. Values are: enabled or disabled.
Scope	Scope classification to which the error belongs. Values include board and pfe.
Category	Category subgroup under the scope level to which the error belongs. Values include: core, functional, io, memory, processing, storage, and switch.
Level	Severity level of the error.
Threshold	Configured threshold value. The associated detection and recovery actions are triggered when this value is exceeded.
Error Limit	The maximum number of times the error is reported.
Support	Support details for the error type.
Occur count	Number of times errors of a specific scope, category, and severity level has occurred.
Clear count	Number of times error instances have been cleared.
Last occurred (ms ago)	Amount of time (in milliseconds) passed since the error last occurred.

Sample Output

show system errors inactive

```
user@host> show system errors inactive
System Inactive Errors Information
CB 0
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
CB 1
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
CHASSIS 0
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
FAN 0
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
FAN 1
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
FPC 0
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
FPC 1
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
FPC 2
```

```
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
FPC 3
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
FPC 4
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
FPC 5
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
FPC 6
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
FPC 7
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
FPM 0
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
PICS 0
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
PSM 0
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
```

```
Inactive Fatal Errors           : 0
PSM 1
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
PSM 2
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
PSM 3
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
PSM 4
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
PSM 5
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
RE 0
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
RE 1
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
SIB 0
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 3
Inactive Fatal Errors           : 0
SIB 1
---------------------------------
```

```
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
SIB 2
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
SIB 3
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
SIB 4
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
SIB 5
---------------------------------
Inactive Minor Errors           : 0
Inactive Major Errors           : 0
Inactive Fatal Errors           : 0
```

show system errors inactive detail

```
user@host> show system errors inactive detail
System Inactive Errors Detail Information
SIB 0
---------------------------------------------------------------
Error Name            : Cell_drop_errors
Identifier            : /sib/0/fabspoked-fchip/0/cm/0/fchip/0/Cell_drop_errors
Description           : Cell drop errors
State                 : enabled
Scope                 : board
Category              : internal
Level                 : major
Threshold             : 10
Error limit           : 0
Support               : No help info provided
Occur count           : 1
```

```
Clear count           : 1
Last occurred(ms ago) : 973151
System Inactive Errors Detail Information
SIB 0
---------------------------------------------------------------
Error Name            : Sib_Asic_PIO_Fault
Identifier            : /sib/0/fabspoked-fchip/0/cm/0/fchip/0/Sib_Asic_PIO_Fault
Description           : Sib Asic PIO Fault
State                 : enabled
Scope                 : switch
Category              : internal
Level                 : major
Threshold             : 10
Error limit           : 0
Support               : No help info provided
Occur count           : 1
Clear count           : 1
Last occurred(ms ago) : 777971
System Inactive Errors Detail Information
SIB 0
---------------------------------------------------------------
Error Name            : sib_link_to_fpc_fault
Identifier            : /sib/0/fabspoked-fchip/0/cm/0/fchip/0/sib_link_to_fpc_fault
Description           : sib_link_to_fpc_fault
State                 : enabled
Scope                 : board
Category              : internal
Level                 : major
Threshold             : 10
Error limit           : 0
Support               : No help info provided
Occur count           : 1
Clear count           : 1
Last occurred(ms ago) : 862333
```

Release Information

Command introduced in Junos OS Evolved Release 19.4R1.

RELATED DOCUMENTATION

show system errors count | 26829
show system errors error-id | 26831
show system errors fru | 26834

show system filesystem encryption status

IN THIS SECTION

- Syntax | 26852
- Description | 26852
- Options | 26852
- Required Privilege Level | 26853
- Output Fields | 26853
- Sample Output | 26853
- Release Information | 26854

Syntax

```
show system filesystem encryption status
<re0 | re1 | routing-engine>
```

Description

View the file-system encryption status.

Options

none	View the file-system encryption status for all Routing Engines.
re0	(Optional) View the file-system encryption status on RE0.

re1 (Optional) View the file-system encryption status on RE1.

routing-engine (Optional) View the file-system encryption status for the specified Routing Engine. Use one of the following options to specify the Routing Engine:

Backup	Backup Routing Engine
Both	Both Routing Engines
Local	Local Routing Engine
Master	Primary Routing Engine
Other	Other Routing Engine

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show system filesystem encryption status` command.

Table 1627: show system filesystem encryption status Output Fields

Field Name	Field Description	Level of Output
Device	Shows the partition disk that is encrypted.	Enabled or disabled

Sample Output

show system filesystem encryption status

```
user@host> show system filesystem encryption status
```

```
Device /dev/sda :
  Partition /dev/sda5 is encrypted.
  LUKS2 UUID: 5eb80d8e-d75f-4637-a90f-cbdd217d7593
  LUKS2 volume opened and mapped to /dev/mapper/luks-sda5
  LVM volumes on /dev/mapper/luks-sda5
    NAME               TYPE    SIZE MOUNTPOINT
    luks-sda5          crypt 181.8G
```

```
    |-jvg_P-jlvmrootrw lvm       8G /.old/tdev/.union/.s
    |-jvg_P-jlvmjunos  lvm      32G /junos
    |-jvg_P-jlvmvm     lvm   141.7G /vm
    |-jvg_P-jlvmspare  lvm       8M /spare
    `-jvg_P-jlvmswap   lvm     108M
  Keyslot: 0 enabled and TPM2.0 token type  <juniper> present
  Keyslot: 1 disabled
  Keyslot: 2 disabled
  Keyslot: 3 disabled
  Keyslot: 4 disabled
  Keyslot: 5 disabled
  Keyslot: 6 disabled
  Keyslot: 7 disabled
Device /dev/sdb :
  Partition /dev/sdb5 is encrypted.
  LUKS2 UUID: e58ae66a-4e3b-451a-a97d-d6b471232060
  LUKS2 volume opened and mapped to /dev/mapper/luks-sdb5
  LVM volumes on /dev/mapper/luks-sdb5
    NAME               TYPE    SIZE MOUNTPOINT
    luks-sdb5          crypt 181.8G
    |-jvg_S-jlvmrootrw lvm       8G
    |-jvg_S-jlvmjunos  lvm      32G
    |-jvg_S-jlvmvm     lvm   141.7G
    |-jvg_S-jlvmspare  lvm       8M
    `-jvg_S-jlvmswap   lvm     108M
  Keyslot: 0 enabled and TPM2.0 token type  <juniper> present
  Keyslot: 1 disabled
  Keyslot: 2 disabled
  Keyslot: 3 disabled
  Keyslot: 4 disabled
  Keyslot: 5 disabled
  Keyslot: 6 disabled
  Keyslot: 7 disabled
```

Release Information

Command introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Encryption with Trusted Platform Module

request system filesystem encryption enable | 15079
request system filesystem encryption keys delete | 15082

show system firmware

IN THIS SECTION

- Syntax | 26855
- Description | 26855
- Options | 26856
- Required Privilege Level | 26857
- Output Fields | 26857
- Sample Output (Junos OS) | 26858
- Release Information | 26865

Syntax

```
show system firmware
<cb number>
<compatibility>
<detail <(fpc slot-number>> | sfb slot-number)> >
<fpc slot-number>
<ftc slot-number>
<pem slot-number>
<routing-engine slot-number>
<sfb slot-number>
```

Description

Display firmware information.

Options

cb *number*

(Optional; MX10004 and MX10008 routers only) Show Control Board firmware version information.

Range: 0 to 1

compatibility

(Optional) Display firmware compatibility information.

detail <(fpc *slot-number*>> **| sfb** *slot-number*)>

(Optional; MX10004 and MX10008 routers only) Show the current, primary, golden, and available firmware version information:

- Current version: Firmware version running on the component
- Primary version: Primary firmware version for the component
- Golden version: Recovery firmware version (recovery image, if component fails to boot from the primary image).
- Available version: Firmware available as part of the Jfirmware package, which is available for upgrade.

Ranges:

- fpc *slot-number*: 0 to 3 (MX10004); 0 to 7 (MX10008)
- sfb *slot-number*: 0 to 5

fpc *slot-number*

(Optional; MX10004 and MX10008 routers only) Show Flexible PIC Concentrator (FPC) firmware version information.

Range: 0 to 3 (MX10004); 0 to 7 (MX10008)

ftc *slot-number*

(Optional; MX10004 and MX10008 routers only) Show Fan Tray Controller (FTC) Board firmware version information.

Range: 0 to 6

pem *slot-number*

(Optional; MX10004 and MX10008 routers only) Show Power Entry Module (PEM) Board firmware version information.

Range: 0 to 2 (MX10004); 0 to 5 (MX10008)

routing-engine *slot-number*

(Optional; MX10004 and MX10008 routers only) Show Routing Engine firmware version information.

Range: 0 to 1

sfb ***slot-number*** (Optional; MX10004 and MX10008 routers with SFB2 boards only) Show Switch Fabric Board firmware version information.

Range: 0 to 5

Required Privilege Level

view

Output Fields

The following table lists all possible output fields for the `show system firmware` command. The fields that are displayed may vary depending on the platform and software release. Output fields are listed in the approximate order in which they appear.

Table 1628: show system firmware Output Fields

Field Name	Field Description	With Command Options
Part	Physical part on the router or switch affected by the firmware.	all
Type	Type of firmware on the router or switch.	all
Tag	Location of the firmware on the interface.	all
Current version	Firmware version on the affected router or switch parts.	all
Primary version	Primary version number for the particular component. If `NA`, a primary firmware package is not available for that component.	`detail`
Golden version	Recovery firmware version for the particular component (recovery image, if the component fails to boot from the primary image). If `NA`, a golden firmware package is not available for that component.	`detail`

Table 1628: show system firmware Output Fields *(Continued)*

Field Name	Field Description	With Command Options
Available version	New versions of firmware for upgrading or downgrading. If `0`, you cannot upgrade the firmware for that component. If `NA`, a firmware package is not available for that component.	all
Status	Firmware condition on the router or switch.	all, except `compatibility` and `detail`
Action	Whether you can upgrade or downgrade, or if no action is available (**none**).This field is displayed only if the `show system firmware` command is used with the `compatibility` option.	`compatibility`

Sample Output (Junos OS)

show system firmware

```
user@host> show system firmware
Part             Type           Tag Current   Available Status
                                    version   version
FPC 0            ROM Monitor 0  0   6.4.10              OK
Routing Engine 0 RE BIOS        0   0                   OK
```

show system firmware compatibility

```
user@host> show system firmware compatibility
Part             Type           Tag Current   Available Action
                                    version   version
FPC 0            ROM Monitor 0  0   6.4.10              None
Routing Engine 0 RE BIOS        0   0                   None
```

show system firmware fpc 0 (MX10004 and MX10008 routers, with line cards LC2101, LC480, LC9600 and Switch Fabric Board SFB2)

```
user@host> show system firmware fpc 0

Part             Type              Tag Current               Available             Status
                                       version               version

FPC 0            ROM Monitor 0     0   9.14.0                0
OK
FPC 0            PCIE Sw(0)        4   1.0.0                 1.0.0
OK
FPC 0            MPCS(0)           2   0.4.0                 0.4.0
OK
FPC 0            I2CS CPLD         6   0.8.0                 0.8.0
OK
FPC 0            BOOT CPLD         5   0.8.0                 0.8.0
OK
FPC 0            ETH-SWITCH        1   1.0.11td4             1.0.11td4
OK
```

show system firmware routing-engine 0 (MX10004 and MX10008 routers, with Switch Fabric Board SFB2)

```
user@host> show system firmware routing-engine 0

Part             Type              Tag Current               Available             Status
                                       version               version

Routing Engine 0 RE BIOS           1   1.15.1                1.15.1
OK
Routing Engine 0 RE FPGA           2   43.0.0                43.0.0
OK
Routing Engine 0 RE SSD1           5   202                   0
OK
Routing Engine 0 RE SSD2           6   202                   0
OK
Routing Engine 0 RE i40e-NVM       10  7.0                   0                     OK
```

show system firmware detail (MX10004 and MX10008 routers, with line cards LC2101, LC480, LC9600 and Switch Fabric Board SFB2)

```
user@host> show system firmware detail

Part             Type             Tag Current             Primary
Golden           Available
                                      Version              version
version          version

FPC 0            ROM Monitor 0     0   9.14.0             NA                   NA
NA
FPC 0            PCIE Sw(0)        4   1.0.0              1.0.0                1.0.0
1.0.0
FPC 0            MPCS(0)           2   0.4.0              0.4.0                0.4.0
0.4.0
FPC 0            I2CS CPLD         6   0.8.0              0.8.0                0.8.0
0.8.0
FPC 0            BOOT CPLD         5   0.8.0              0.8.0                0.8.0
0.8.0
FPC 0            ETH-SWITCH        1   1.0.11td4          1.0.11td4            1.0.11td4
1.0.11td4
FPC 1            BIOS              0   29.1.0             NA                   NA
29.1.0
FPC 1            HM CTRL FPGA      1   0.5.0              NA                   NA
0.5.0
FPC 1            HM MEZZ FPGA      2   0.9.0              NA                   NA
0.9.0
FPC 1            BOOT CPLD         3   74.0.0             74.0.0               NA
74.0.0
FPC 1            I2CS CPLD         4   7.0.0              NA                   NA
7.0.0
FPC 1            PORT CPLD(0)      5   7.0.0              NA                   NA
7.0.0
FPC 1            PORT CPLD(1)      6   7.0.0              NA                   NA
7.0.0
FPC 1            PORT CPLD(2)      7   7.0.0              NA                   NA
7.0.0
FPC 1            ETH-SWITCH SPI    8   2021.3.0           2021.3.0             2021.3.0
2021.3.0
SFB 0            FPGA PRIM         0   0.5.0              0.5.0                0.5.0
0.5.0
```

```
SFB 1            FPGA PRIM         0   0.5.0              0.5.0                0.5.0
0.5.0
SFB 2            FPGA PRIM         0   0.5.0              0.5.0                0.5.0
0.5.0
SFB 3            FPGA PRIM         0   0.5.0              0.5.0                0.5.0
0.5.0
SFB 4            FPGA PRIM         0   0.5.0              0.5.0                0.5.0
0.5.0
SFB 5            FPGA PRIM         0   0.5.0              0.5.0                0.5.0
0.5.0
```

show system firmware (QFX10008 Switch)

```
user@host> show system firmware
Part         Type            Status
CB 0         FPGA            OK
FPC 0        U-Boot             OK
CTRL         FPGA            PROGRAMMING
PORT         FPGA            PROGRAMMING
FPM          FPGA            OK
FTC 0        FPGA            OK
FTC 1        FPGA            OK
SIB 0        FPGA            OK
SIB 1        FPGA            OK
SIB 2        FPGA            OK
SIB 3        FPGA            OK
SIB 4        FPGA            OK
SIB 5        FPGA            OK
```

show system firmware (MX Series Routers with MIC3-100G-DWDM)

```
user@host> show system firmware
Part             Type          Tag  Current   Available Status
                      version   version
FPC 0            ROM Monitor 0  0   10.4.1              OK
FPC 1            ROM Monitor 0  0   10.4.1              OK
FPC 2            ROM Monitor 0  0   10.4.1              OK
PIC 0            CMIC LTC 2/0   1   .0        1.0       OK
FPC 3            ROM Monitor 0  0   10.4.1              OK
FPC 4            ROM Monitor 0  0   13.3.1              OK
```

```
FPC 4            MPCS(0)        2   0.24.0              OK
Routing Engine 0 RE BIOS        0   1.18                OK
Routing Engine 1                0   1.18                OK
```

The current firmware version `.0` does not match the available version `1.0`. This output displays the status before the firmware upgrade.

show system firmware (MX Series Routers with MIC3-100G-DWDM)

```
user@host> show system firmware
Part             Type          Tag Current   Available Status
                       version   version
FPC 0            ROM Monitor 0  0   10.4.1              OK
FPC 1            ROM Monitor 0  0   10.4.1              OK
FPC 2            ROM Monitor 0  0   10.4.1              OK
PIC 0            CMIC LTC 2/0   1   1.0        1.0      UPGRADED SUCCESSFULLY
FPC 3            ROM Monitor 0  0   10.4.1              OK
FPC 4            ROM Monitor 0  0   13.3.1              OK
FPC 4            MPCS(0)        2   0.24.0              OK
Routing Engine 0 RE BIOS        0   1.18                OK
Routing Engine 1                0   1.18                OK
```

The current firmware version `1.0` matches the available version `1.0`. This output displays the status after the firmware upgrade.

show system firmware (MX10004 Router)

The MX10004 router supports up to six SFB2 slots and up to four FPC line cards.

```
user@host> show system firmware
Part             Type              Tag Current               Available             Status
                                       version               version
CB 0             CB FPGA           0   264.0.0               267.0.0
OK
Routing Engine 0 RE BIOS           1   1.14.1                1.15.1
OK
Routing Engine 0 RE FPGA           2   41.0.0                43.0.0
OK
Routing Engine 0 RE SSD1           3   12034                 12050
OK
Routing Engine 0 RE SSD2           4   12034                 12050
```

```
OK
Routing Engine 0 RE i40e-NVM       10  6.1                  0
OK
Routing Engine 1                   0   0.0.0                71.63d
OK
FPC 1            ROM Monitor 0     0   <TBD>                <TBD>
OK
FPC 1            AR CTL FPGA       1   <TBD>                <TBD>                OK
FPC 1            AR MEZZ FPGA      2   <TBD>                <TBD>
OK
FPC 1            Port PLD          3   <TBD>                <TBD>
OK
FPC 1            I2CS CPLD         4   <TBD>                <TBD>
OK
FPC 1            BOOT CPLD         5   <TBD>                <TBD>                OK
FPC 1            PCIE Sw(0)        6   <TBD>                <TBD>
OK
FTC 0            FPGA              0   2.0.0                2.0.0
OK
FTC 1            FPGA              0   2.0.0                2.0.0
OK
PEM 1            PSU AC            1   20.54.0              0
OK
PEM 2            PSU AC            1   20.54.0              0
OK
PEM 4            PSU AC            1   20.54.0              0
OK
PEM 5            PSU AC            1   15.36.0              0
OK
FPM 0            FPGA              0   1.9.0                1.9.0
OK
SPMB 0           ROM Monitor 0     0   ??                   71.63d
OK
SFB 0            FPGA              0   <TBD>                <TBD>
OK
SFB 1            FPGA              0   <TBD>                <TBD>
OK
SFB 2            FPGA              0   <TBD>                <TBD>
OK
SFB 3            FPGA              0   <TBD>                <TBD>
OK
SFB 4            FPGA              0   <TBD>                <TBD>
```

```
OK
SFB 5            FPGA             0   <TBD>                <TBD>                OK
```

show system firmware (PTX10003-80C Routers)

```
user@router> show system firmware

Part             Type             Tag Current   Available Status
                                     version   version
CB 0             MBAM FPGA        0   181011.09 181011.09 OK
FPC 0            PBAM FPGA        0             181011.09 OK
FPC 0            QBAM FPGA        1             181011.0B OK
FPC 1            PBAM FPGA        0   181011.09 181011.09 OK
FPC 1            QBAM FPGA        1   181011.0B 181011.0B OK
PSM 0            Psm MCU AC       0   0116.0036 0209.9205 OK
PSM 1            Psm MCU AC       0   255255.255255 0209.9205
OK

RE 0             UBAM FPGA        0   181023.0F 181023.0F OK
RE 0             i40e             1   6.01      6.01      OK
RE 0             SSD Primary      2   SBR12050  SBR12050  OK
RE 0             SSD Secondary    3   SBR12050  SBR12050  OK
RE 0             BIOS ROM         4   14.1      15.01     OK
```

Starting in Junos OS Evolved Release 19.1R1, on PTX10003-80C and PTX10003-160C routers, the `show system firmware` command displays information based on the accessibility of the device and not the FRU state. The firmware information is cached in the FRU online state if the device is accessible. This cached information is displayed even when the FRU is in error or offline condition.

show system firmware (SRX1600, SRX2300, and SRX4300)

Starting in Junos OS Release 23.4R1, the `show system firmware` command is supported on SRX1600, SRX2300 and SRX4300 devices.

```
user@host> show system firmware
Part             Type             Tag Current               Available            Status
                                     version               version
Routing Engine 0 BIOS             1   0.27.1                0
OK
Routing Engine 0 FPGA             2   1.103                 0
```

```
OK
Routing Engine 0 SSD1              5   22.10.27             0
OK
Routing Engine 0 I210              7   3.30.1908            0
OK
Routing Engine 0 I210-HA0          8   3.30.2047            0
OK
Routing Engine 0 I210-HA1          9   3.30.2047            0
OK
Routing Engine 0 CTRL-FPGA         11  1.43                 0
OK
Routing Engine 0 UCD               12  0.3                  0                    OK
```

Release Information

Command introduced in Junos OS Release 7.4.

Options `cb`, `detail`, `fpc`, `ftc`, `pem`, `routing-engine`, and `sfb` introduced in Junos OS Release 24.2R1.

show system firmware (Junos OS Evolved)

IN THIS SECTION

- Syntax | 26866
- Description | 26866
- Options | 26866
- Required Privilege Level | 26866
- Output Fields | 26866
- Sample Output | 26867
- Release Information | 26868

Syntax

```
show system firmware
```

Description

Display firmware information.

Options

Required Privilege Level

view

Output Fields

The following table lists all possible output fields for the `show system firmware` command. The fields that are displayed may vary depending on the platform and software release. Output fields are listed in alphabetical order.

Table 1 lists the output fields for the `show system firmware` command.

Table 1629: show system firmware Output Fields

Field Name	Field Description
Available version	New versions of firmware for upgrading or downgrading. If `0`, you cannot upgrade the firmware for that component. If `NA`, a firmware package is not available for that component.
Current version	Firmware version on the affected device parts.
Part	Physical part on the device affected by the firmware.
Status	Firmware condition on the device.

Table 1629: show system firmware Output Fields *(Continued)*

Field Name	Field Description
Tag	Location of the firmware on the interface.
Type	Type of firmware on the device.

Sample Output

show system firmware on a PTX10008

```
user@host> show system firmware
```

```
Part            Type            Tag Current              Available            Status
                                    version              version
FPC 0           CTRL FPGA       0   190912.11            190912.11            OK
FPC 0           PORT FPGA       1   15                   15                   OK
FPC 0           WANIO FPGA      2   15                   15                   OK
FPC 0           I2CS FPGA       3   19                   19                   OK
FPC 0           PMB BIOS        4   26.01                28.01                OK
FPC 0           BOOT CPLD       5   17                   17                   OK
FPC 1           CTRL FPGA       0   190912.11            190912.11            OK
FPC 1           PORT FPGA       1   15                   15                   OK
FPC 1           WANIO FPGA      2   15                   15                   OK
FPC 1           I2CS FPGA       3   19                   19                   OK
FPC 1           PMB BIOS        4   26.01                28.01                OK
FPC 1           BOOT CPLD       5   17                   17                   OK
PSM 4           PSM MCU AC      0   aa17.aa17.ab15.ab15.ac27 aa17.aa17.ab15.ab15.ac27 OK
PSM 5           PSM MCU AC      0   aa17.aa17.ab15.ab15.ac27 aa17.aa17.ab15.ab15.ac27 OK
RE 0            CBC FPGA        0   0x10B                0x10B                OK
RE 0            RE FPGA         1   0x2B                 0x2B                 OK
RE 0            SSD Primary     2   SBR12050             SBR12050             OK
RE 0            SSD Secondary   3   SBR12050             SBR12050             OK
RE 0            i40e            4   7.00                 7.00                 OK
RE 0            BIOS ROM        5   14.01                19.01                OK
```

```
RE 1            CBC FPGA          0   0x10B                0x10B                OK
RE 1            RE FPGA           1   0x2B                 0x2B                 OK
RE 1            SSD Primary       2   SBR12050             SBR12050             OK
RE 1            SSD Secondary     3   SBR12050             SBR12050             OK
RE 1            i40e              4   6.01                 7.00                 OK
RE 1            BIOS ROM          5   14.01                19.01                OK
SIB 0           SIB FPGA          0   200311.0F            200831.1D            OK
FTC 0           FTC FPGA          0   0.5                  0.5                  OK
FTC 1           FTC FPGA          0   0.5                  0.5                  OK
```

Release Information

show system health-monitor

IN THIS SECTION

- Syntax | 26868
- Description | 26869
- Options | 26869
- Required Privilege Level | 26869
- Output Fields | 26869
- Sample Output | 26870
- Release Information | 26872

Syntax

```
show system health-monitor
<fpc fpc-slot fpc-slot>
```

Description

Display the J-Insight health monitor results. Starting with Junos OS Release 18.2R1, J-Insight supports health monitoring for FPC FRUs on the MX Series routers.

Before you can generate output with **show system health-monitor**, you will need to perform an initial CLI configuration. Please see J-Insight Device Monitor Basic Configuration for information on how to configure J-Insight.

Options

none	Display information for all FPCs.
fpc fpc-slot *fpc-slot*	(Optional) Display information for a specified FPC.

Required Privilege Level

admin

Output Fields

Table 1630 on page 26869 lists the output fields for the `show system health-monitor` command. Output fields are listed in the approximate order in which they appear.

Table 1630: show system health-monitor Output Fields

Field Name	Field Description
Component	Platform component name.
Health-Parameter	Health parameter name.
Value	Reported health value collected by the health monitor.
Threshold	Default threshold value for the health parameter.
Health-Status	State of the health parameter. Values are: GREEN, YELLOW, or RED.

Table 1630: show system health-monitor Output Fields *(Continued)*

Field Name	Field Description
FPC SLOT	FPC slot number.

Sample Output

show system health-monitor

```
user@host> show system health-monitor
Component      Health-Parameter     Value        Threshold       Health-Status
-------------------------------------------------------------------------------
FPC SLOT: 0

board.0.cpu.0     CPU Load 1 (1  sec)         15              NA          NA
board.0.cpu.0     CPU Load 2 (5  sec)         16              NA          NA
board.0.cpu.0     CPU Load 3 (10 sec)         15              NA          NA
board.0.cpu.0     CPU Load 4 (1  min)         15              NA          NA
board.0.cpu.0     heap_util[Kernel]           11              NA          NA
board.0.cpu.0     heap_util[LAN buffer]       20              NA          NA
board.0.temp.0    Exhaust A                   46 C/114.8 F    75          GREEN
board.0.temp.0    Exhaust B                   59 C/138.2 F    75          GREEN
board.0.temp.0    Intake                      41 C/105.8 F    75          GREEN
board.0.temp.0    LU 0 Chip                   55 C/131 F      NA          NA
board.0.temp.0    LU 0 TSen                   50 C/122 F      NA          NA
board.0.temp.0    LU 1 Chip                   49 C/120.2 F    NA          NA
board.0.temp.0    LU 1 TSen                   50 C/122 F      NA          NA
board.0.temp.0    LU 2 Chip                   57 C/134.6 F    NA          NA
board.0.temp.0    LU 2 TSen                   50 C/122 F      NA          NA
board.0.temp.0    LU 3 Chip                   64 C/147.2 F    NA          NA
board.0.temp.0    LU 3 TSen                   50 C/122 F      NA          NA
board.0.temp.0    PLX Switch Chip             55 C/131 F      NA          NA
board.0.temp.0    PLX Switch TSen             50 C/122 F      NA          NA
board.0.temp.0    XF 0 Chip                   69 C/156.2 F    NA          NA
board.0.temp.0    XF 0 TSen                   50 C/122 F      NA          NA
board.0.temp.0    XM 0 Chip                   58 C/136.4 F    NA          NA
board.0.temp.0    XM 0 TSen                   50 C/122 F      NA          NA
npu.0.fabric.0    PLANE0.dest[0-31]           0x80000003      NA          NA
npu.0.fabric.0    PLANE0.dest[128-159]        0x00000300      NA          NA
```

```
npu.0.fabric.0    PLANE0.dest[160-191]        0x20000000      NA          NA
npu.0.fabric.0    PLANE0.dest[192-223]        0x00000000      NA          NA
npu.0.fabric.0    PLANE0.dest[224-239]        0x00000000      NA          NA
npu.0.fabric.0    PLANE0.dest[32-63]          0x00200000      NA          NA
npu.0.fabric.0    PLANE0.dest[64-95]          0x00000000      NA          NA
npu.0.fabric.0    PLANE0.dest[96-127]         0x00000080      NA          NA
npu.0.fabric.0    PLANE1.dest[0-31]           0x80000003      NA          NA
npu.0.fabric.0    PLANE1.dest[128-159]        0x00000300      NA          NA
npu.0.fabric.0    PLANE1.dest[160-191]        0x20000000      NA          NA
npu.0.fabric.0    PLANE1.dest[192-223]        0x00000000      NA          NA
npu.0.fabric.0    PLANE1.dest[224-239]        0x00000000      NA          NA
npu.0.fabric.0    PLANE1.dest[32-63]          0x00200000      NA          NA
npu.0.fabric.0    PLANE1.dest[64-95]          0x00000000      NA          NA
npu.0.fabric.0    PLANE1.dest[96-127]         0x00000080      NA          NA
npu.0.fabric.0    PLANE2.dest[0-31]           0x80000003      NA          NA
npu.0.fabric.0    PLANE2.dest[128-159]        0x00000300      NA          NA
npu.0.fabric.0    PLANE2.dest[160-191]        0x20000000      NA          NA
npu.0.fabric.0    PLANE2.dest[192-223]        0x00000000      NA          NA
npu.0.fabric.0    PLANE2.dest[224-239]        0x00000000      NA          NA
npu.0.fabric.0    PLANE2.dest[32-63]          0x00200000      NA          NA
npu.0.fabric.0    PLANE2.dest[64-95]          0x00000000      NA          NA
npu.0.fabric.0    PLANE2.dest[96-127]         0x00000080      NA          NA
npu.0.fabric.0    PLANE3.dest[0-31]           0x80000003      NA          NA
npu.0.fabric.0    PLANE3.dest[128-159]        0x00000300      NA          NA
npu.0.fabric.0    PLANE3.dest[160-191]        0x20000000      NA          NA
npu.0.fabric.0    PLANE3.dest[192-223]        0x00000000      NA          NA
npu.0.fabric.0    PLANE3.dest[224-239]        0x00000000      NA          NA
npu.0.fabric.0    PLANE3.dest[32-63]          0x00200000      NA          NA
npu.0.fabric.0    PLANE3.dest[64-95]          0x00000000      NA          NA
npu.0.fabric.0    PLANE3.dest[96-127]         0x00000080      NA          NA
npu.0.fabric.0    PLANE4.dest[0-31]           0x00000000      NA          NA
npu.0.fabric.0    PLANE4.dest[128-159]        0x00000000      NA          NA
npu.0.fabric.0    PLANE4.dest[160-191]        0x00000000      NA          NA
npu.0.fabric.0    PLANE4.dest[192-223]        0x00000000      NA          NA
npu.0.fabric.0    PLANE4.dest[224-239]        0x00000000      NA          NA
npu.0.fabric.0    PLANE4.dest[32-63]          0x00000000      NA          NA
npu.0.fabric.0    PLANE4.dest[64-95]          0x00000000      NA          NA
npu.0.fabric.0    PLANE4.dest[96-127]         0x00000000      NA          NA
npu.0.fabric.0    PLANE5.dest[0-31]           0x00000000      NA          NA
npu.0.fabric.0    PLANE5.dest[128-159]        0x00000000      NA          NA
npu.0.fabric.0    PLANE5.dest[160-191]        0x00000000      NA          NA
npu.0.fabric.0    PLANE5.dest[192-223]        0x00000000      NA          NA
npu.0.fabric.0    PLANE5.dest[224-239]        0x00000000      NA          NA
```

```
npu.0.fabric.0    PLANE5.dest[32-63]          0x00000000      NA          NA
npu.0.fabric.0    PLANE5.dest[64-95]          0x00000000      NA          NA
npu.0.fabric.0    PLANE5.dest[96-127]         0x00000000      NA          NA
npu.0.fabric.0    PLANE6.dest[0-31]           0x00000000      NA          NA
npu.0.fabric.0    PLANE6.dest[128-159]        0x00000000      NA          NA
npu.0.fabric.0    PLANE6.dest[160-191]        0x00000000      NA          NA
npu.0.fabric.0    PLANE6.dest[192-223]        0x00000000      NA          NA
npu.0.fabric.0    PLANE6.dest[224-239]        0x00000000      NA          NA
npu.0.fabric.0    PLANE6.dest[32-63]          0x00000000      NA          NA
npu.0.fabric.0    PLANE6.dest[64-95]          0x00000000      NA          NA
npu.0.fabric.0    PLANE6.dest[96-127]         0x00000000      NA          NA
npu.0.fabric.0    PLANE7.dest[0-31]           0x00000000      NA          NA
npu.0.fabric.0    PLANE7.dest[128-159]        0x00000000      NA          NA
npu.0.fabric.0    PLANE7.dest[160-191]        0x00000000      NA          NA
npu.0.fabric.0    PLANE7.dest[192-223]        0x00000000      NA          NA
npu.0.fabric.0    PLANE7.dest[224-239]        0x00000000      NA          NA
npu.0.fabric.0    PLANE7.dest[32-63]          0x00000000      NA          NA
npu.0.fabric.0    PLANE7.dest[64-95]          0x00000000      NA          NA
npu.0.fabric.0    PLANE7.dest[96-127]         0x00000000      NA          NA
npu.0.memory.0    Counters_EDMEM Utilization  50              NA          NA
npu.0.memory.0    EDMEM Utilization           37              NA          NA
npu.0.memory.0    ENCAPS_EDMEM Utilization    100             NA          NA
npu.0.memory.0    Firewall_EDMEM Utilization  1               NA          NA
npu.0.memory.0    HASH_EDMEM Utilization      100             NA          NA
npu.0.memory.0    HASH_OMEM Utilization       100             NA          NA
npu.0.memory.0    IDMEM Utilization           86              NA          NA
npu.0.memory.0    LMEM_LMEM Utilization       100             NA          NA
npu.0.memory.0    Next_Hop_EDMEM Utilization  65              NA          NA
npu.0.memory.0    OMEM Utilization            1               NA          NA
npu.0.memory.0    UEID_SHARED_SPACE_EDMEM Utilization 1       NA          NA
npu.0.memory.0    UEID_SPACE_EDMEM Utilization 1              NA          NA
npu.0.util.0      EDMEM Avg Load              1               NA          NA
npu.0.util.0      Global Utilization          1               NA          NA
npu.0.util.0      IDMEM Avg Load              1               NA          NA
npu.0.util.0      OMEM Avg Load               0               NA          NA
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

J-Insight Device Monitor Basic Configuration

J-Insight Device Monitor Basic Configuration

J-Insight Device Monitor Basic Configuration

J-Insight Device Monitor Basic Configuration

delete services jinsightd subscribe health-monitor | 5284

set services jinsightd traceoptions | 16894

show system information

IN THIS SECTION

- Syntax | 26873
- Description | 26873
- Options | 26873
- Required Privilege Level | 26874
- Sample Output | 26874
- Release Information | 26874

Syntax

```
show system information
```

Description

Display high-level system information for the device including the model number, device family, Junos OS release, and hostname.

Options

none Display system information for the device.

Required Privilege Level

view

Sample Output

show system information

```
user@host> show system information
Model: mx960
Family: junos
Junos: 17.2R1
Hostname: host
```

Release Information

Command introduced in Junos OS Release 17.2.

show system inventory hardware cpu

IN THIS SECTION

- Syntax | 26874
- Description | 26875
- Required Privilege Level | 26875
- Output Fields | 26875
- Sample Output | 26877
- Release Information | 26880

Syntax

```
show system inventory hardware cpu
```

Description

Display system CPU statistics.

Required Privilege Level

view

Output Fields

Table 1631 on page 26875 lists the output fields for the `show system inventory hardware cpu` command. Output fields are listed in the approximate order in which they appear.

> **NOTE**: Starting in Junos OS Release 22.1R1, the CPU allocations are displayed as a number range (for example, 4-7) instead of individual numbers (for example, 4,5,6,7).

Table 1631: show system inventory hardware cpu Output Fields

Field Name	Field Description
Fields for Inventory CPU Information	
No of CPUs	Total number of CPUs.
No of Logical CPUs/Hyper threads	Total number of hyper threads.
No of CPU sockets	Total number of CPU sockets.
No of Cores per socket	Total number of cores per socket.
No of Hyper threads per core	Total number of hyper threads per core.
CPU Vendor	The name of the CPU vendor.
CPU Model	The CPU model.
CPU Architecture	The type of CPU architecture.
CPU Speed	The CPU speed, in GHz.

Table 1631: show system inventory hardware cpu Output Fields *(Continued)*

Field Name	Field Description
CPU Cache	The size of the CPU cache.
No of Max vCPUs	The maximum number of allowed virtual CPUs.
Fields for CPU Statistics	
`CPU ID`	The CPU ID
`User Time`	The amount of user time, in seconds.
`System Time`	The amount of system time, in seconds.
`Idle Time`	The amount of time spent in idle mode, in seconds.
`Nice Time`	The amount of spent nice time, in seconds.
`I/O Wait Time`	The amount of time spent waiting for input/output (I/O) operations, in seconds.
`Interrupt Service Time`	The amount of interrupt service time, in seconds.
Fields for CPU Pinning Information	
`Virtual Machine`	The name of the virtual machine.
`vCPU`	The ID of virtual CPUs used by the virtual machine.
`CPU`	The ID of CPUs used by the virtual machine.
Fields for Emulator Thread Pinning Information (This section is displayed starting in Junos OS Release 22.1R1.)	
Virtual Machine	The name of the virtual machine.

Table 1631: show system inventory hardware cpu Output Fields *(Continued)*

Field Name	Field Description
CPU	The ID of CPUs used by the virtual machine.
`System Component`	The name of the system component.
`CPUs`	The ID of CPUs used by the system component.
Quota	The CPU quota configured using the `cpu colocation quota` parameter. **NOTE**: This parameter is displayed starting in Junos OS Release 22.1R1.

Sample Output

show system inventory hardware cpu (NFX150)

```
user@host> show system inventory hardware cpu
```

```
Inventory CPU Information
-------------------------
No of CPUs:                       4
No of Logical CPUs/Hyper threads: 4
No of CPU sockets:                1
No of Cores per socket:           4
No of Hyper threads per core:     1
No of Max VCPUs:                  4
CPU Vendor:                       GenuineIntel
CPU Model:                        Intel(R) Atom(TM) CPU C3558 @ 2.20GHz
CPU Architecture:                 x86_64
CPU Speed:                        2.2000 GHz
CPU Cache:                        2048 KB

CPU Statistics (Time in sec)
----------------------------
User Time:              1044
System Time:            90
```

```
Idle Time:              5605
I/O Wait Time:          21
Nice Time:              0
Interrupt Service Time: 0

CPU Pinning Information
---------------------------------------
Virtual Machine             vCPU    CPU
--------------------------- ------- ---
vjunos0                     0       0
testvnf1                    0       2-3
```

show system inventory hardware cpu (NFX250 (NG))

```
user@host> show system inventory hardware cpu
```

```
Inventory CPU Information
-------------------------
No of CPUs:                        4
No of Logical CPUs/Hyper threads: 8
No of CPU sockets:                 1
No of Cores per socket:            4
No of Hyper threads per core:      2
No of Max VCPUs:                   8
CPU Vendor:                        GenuineIntel
CPU Model:                         Intel(R) Pentium(R) CPU D1517 @ 1.60GHz
CPU Architecture:                  x86_64
CPU Speed:                         1.6000 GHz
CPU Cache:                         6144 KB

CPU Statistics (Time in sec)
----------------------------
User Time:              88347
System Time:            2121
Idle Time:              241343
I/O Wait Time:          73
Nice Time:              0
Interrupt Service Time: 0

CPU Pinning Information
```

```
--------------------------------------
Virtual Machine             vCPU    CPU
--------------------------- ------- ---
vjunos0                     0       0
testvnf1                    0       2
testvnf1                    1       6-7

Emulator Thread Pinning Information
-------------------------------
Virtual Machine             CPU
--------------------------- ---
vjunos0                     0
testvnf1                    2-3


System Component                CPUs     Quota
------------------------------- -------- -----
ovs-vswitchd                    0,4
srxpfe                          1,5
jdm
```

show system inventory hardware cpu (NFX350)

```
user@host> show system inventory hardware cpu
```

```
Inventory CPU Information
-------------------------
No of CPUs:                       8
No of Logical CPUs/Hyper threads: 16
No of CPU sockets:                1
No of Cores per socket:           8
No of Hyper threads per core:     2
No of Max VCPUs:                  14
CPU Vendor:                       GenuineIntel
CPU Model:                        Intel(R) Xeon(R) D-2146NT CPU @ 2.30GHz
CPU Architecture:                 x86_64
CPU Speed:                        2.3000 GHz
CPU Cache:                        11264 KB

CPU Statistics (Time in sec)
```

```
----------------------------
User Time:              4738
System Time:            115
Idle Time:              19399
I/O Wait Time:          17
Nice Time:              0
Interrupt Service Time: 0

CPU Pinning Information
---------------------------------------
Virtual Machine             vCPU    CPU
--------------------------- ------- ---
vjunos0                     0       8
testvnf1                    0       3-4
testvnf1                    1       5-7

Emulator Thread Pinning Information
-------------------------------
Virtual Machine             CPU
--------------------------- ---
vjunos0                     8
testvnf1                    11-15


System Component                CPUs     Quota
------------------------------- -------- -----
ovs-vswitchd                    1,8
srxpfe                          0,2
jdm                             8
```

Release Information

Command introduced in Junos OS Release 15.1X53-D40.

show-system-inventory-hardware-memory

IN THIS SECTION

- Syntax | 26881
- Description | 26881
- Required Privilege Level | 26881
- Output Fields | 26881
- Sample Output | 26882
- Sample Output | 26883
- Release Information | 26883

Syntax

```
show system inventory hardware memory
```

Description

Display hardware memory statistics for a disaggregated Junos OS platform.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show system inventory hardware memory` command. Output fields are listed in the approximate order in which they appear.

Table 1632: show system inventory hardware memory Output Fields

Field Name	Field Description
`Total`	The total amount of available virtual memory, in kibibytes (KiBs).
`Used`	The total amount of used virtual memory, in kibibytes (KiBs).
`Available`	The total amount of available virtual memory, in kibibytes (KiBs).
`Free`	The total amount of free virtual memory, in kibibytes (KiBs).
`Percent Used`	The percentage of virtual memory used.
Fields for Swap Memory	
`Total`	The total amount of available swap space memory, in kibibytes (KiBs).
`Used`	The total amount of used swap space memory, in kibibytes (KiBs).
`Free`	The total amount of free swap space memory, in kibibytes (KiBs).
`Percent Used`	The percentage of swap space memory used.

Sample Output

show system inventory hardware memory (NFX250)

```
user@host> show system inventory hardware memory
```

```
Inventory Memory Information
----------------------------
Virtual Memory:
---------------
Total  (KiB): 15949116
```

```
Used   (KiB): 5542256
Available (KiB): 10437928
Free   (KiB): 10406860
Percent Used: 28.6

Swap Memory:
------------
Total  (KiB): 1249996
Used   (KiB): 0
Free   (KiB): 1249996
Percent Used: 0.0
```

Sample Output

show system inventory hardware memory (NFX150)

```
user@host> show system inventory hardware memory
```

```
Inventory Memory Information
----------------------------
Virtual Memory:
---------------
Total       (KiB): 7936420
Used        (KiB): 7880136
Available   (KiB): 1801340
Free        (KiB): 56284
Percent Used     : 77.30

Swap Memory:
------------
Total  (KiB): 0
Used   (KiB): 0
Free   (KiB): 0
Percent Used: 0.00
```

Release Information

Command introduced in Junos OS Release 15.1X53-D40 for the NFX250 Network Services Platform.

Command introduced in Junos OS Release 18.1R1 for NFX150 devices.

show-system-inventory-hardware-optics

IN THIS SECTION

- Syntax | 26884
- Description | 26884
- Required Privilege Level | 26884
- Output Fields | 26884
- Sample Output | 26885
- Release Information | 26885

Syntax

```
show system inventory hardware optics
```

Description

Display port information of all hardware optics installed in the router, including the Xcvr vendor details and serial number.

Required Privilege Level

view

Output Fields

No Link Title lists the output fields for the `show system inventory hardware optics` command. Output fields are listed in the approximate order in which they appear.

Table 1633: show system inventory hardware optics Output Fields

Field Name	Field Description
Port	Name of the port
Cable type	Name of the cable type
Serial Number	Serial number of the optics component. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis.
Xcvr vendor number	Name of the Xcvr vendor
Xcvr vendor part number	Part number of the Xcvr vendor

Sample Output

show system inventory hardware optics (NFX150)

```
user@host> show system inventory hardware optics
```

```
Port information:
                        Serial                        Xcvr vendor
 Port    Cable type      number        Xcvr vendor    part number
 heth-0-4  SFP-SX        AM17062RZ6U    AVAGO       AFBR-5715PZ-JU1
 heth-0-5  SFP+-10G-SR    N73EBHS      FINISAR CORP.   FTLX8574D3BCL-J1
```

Release Information

Command introduced in Junos OS Release 19.4R1.

show-system-inventory-hardware-storage

IN THIS SECTION

- Syntax | 26886
- Description | 26886
- Required Privilege Level | 26886
- Output Fields | 26886
- Sample Output | 26888
- Release Information | 26889

Syntax

```
show system inventory hardware storage
```

Description

Display hardware storage details such as the list of partitions, disk usage per partition, and disk I/O statistics for a disaggregated Junos OS platform.

Required Privilege Level

view

Output Fields

Table 1 lists the output fields for the `show system inventory hardware storage` command. Output fields are listed in the approximate order in which they appear.

Table 1634: show system inventory hardware storage Output Fields

Field Name	Field Description
`Device`	The device path.
`Mount Point`	The mount point of the device path.
`File System`	The file system type.
`Options`	Options available for the device path.
Fields for Disk Usage Information	
`Total`	The total amount of disk usage space, in mebibytes (MiB).
`Used`	The amount of used disk usage space, in mebibytes (MiB).
`Free`	The amount of free disk usage space, in mebibytes (MiB).
`Percentage Used`	The percentage of used disk space.
Fields for Disk I/O Information	
`Read Count`	The number of times the disk has been read.
`Write Count`	The number of times a write operation has happened on the disk.
`Read Bytes`	The number of bytes used in read operations on the disk.
`Write Bytes`	The number of bytes used in write operations on the disk.
`Read Time`	The amount of time the disk has been read, in milliseconds.

Table 1634: show system inventory hardware storage Output Fields *(Continued)*

Field Name	Field Description
Read Time	The amount of time write operations have been performed on the disk, in milliseconds.

Sample Output

show system inventory hardware storage

```
user@host> show system inventory hardware storage
```

```
Device                                    Mount Point      File System  Options
----------------------------------------- ---------------- -----------
----------------------------------
/dev/sda4                                 /                ext4         rw,relatime,data=ordered
/dev/sda5                                 /var             ext4         rw,relatime,data=ordered
/dev/sda2                                 /boot            vfat
rw,noatime,fmask=0022,dmask=0022,codepage=437,iocharset=iso8859-1,shortname=mixed,errors=remount-
ro
/dev/sda1                                 /boot/efi        vfat
rw,noatime,fmask=0022,dmask=0022,codepage=437,iocharset=iso8859-1,shortname=mixed,errors=remount-
ro
/dev/sda3                                 /app_disk        ext4         rw,noatime
/dev/mapper/vg0_vjunos-lv_junos           /junos           ext4
rw,noatime,discard,data=ordered
/dev/mapper/vg0_vjunos-lv_var_third_party /third-party     ext4
rw,noatime,discard,data=ordered

Disk Usage Information
-----------------------------------------------------------------------------------------
Disk                                        Total (MiB) Used (MiB)  Free (MiB)  % Used
------------------------------------------- ----------- ----------- ----------- ------
/dev/sda4                                   1409        822         497         58.0
/dev/sda5                                   5639        201         5128        3.0
/dev/sda2                                   976         569         406         58.0
/dev/sda1                                   189         17          172         9.0
```

```
/dev/sda3                                   1           0           1           1.0
/dev/mapper/vg0_vjunos-lv_junos             9951        3986        5436        40.0
/dev/mapper/vg0_vjunos-lv_var_third_party   94849       5734        84274       6.0

Disk I/O Information
--------------------
 Read Count: 304501
Write Count: 1176577
 Read Bytes: 4199641600
Write Bytes: 16493698560
  Read Time: 45528
 Write Time: 1181938
```

Release Information

Command introduced in Junos OS Release 15.1X53-D40 or the NFX250 Network Services Platform.

show system khms-stats

IN THIS SECTION

- Syntax | 26889
- Description | 26890
- Options | 26890
- Required Privilege Level | 26890
- Sample Output | 26890
- Release Information | 26896

Syntax

```
show system khms-stats
<detail>
```

Description

Display information about interface states and statistics for next hop and memory usage.

Options

detail Display information about kernel health parameters, such as graceful Routing Engine switchover (GRES) and in-service software upgrade (ISSU).

Required Privilege Level

View

Sample Output

show system khms-stats

```
user@host> show system khms-stats
Memory usage by ifstates, routes, nexthops, gencfg blobs
==============================================================================
Memory type  Total no. of states     Memory used  Percentage of Dead states
     Ifstate                885           525KB             0%
      Routes                150            25KB             0%
     Nexthop                163            39KB             0%

Number of delayed unrefs of RT and NH: 0 (Max allowed 40000)

Number of delayed weight unrefs of RT and NH: 0 (Max allowed 400000)

Mbufs/Cluster Usage
=====================================================
852/1938/2790 mbufs in use (current/cache/total)
772/1198/1970/250436 mbuf clusters in use (current/cache/total/max)
769/1002 mbuf+clusters out of packet secondary zone in use (current/cache)
0/19/19/125218 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/37101 9k jumbo clusters in use (current/cache/total/max)
0/0/0/20869 16k jumbo clusters in use (current/cache/total/max)
1757K/2956K/4713K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
```

```
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile

Memory Pages stats
=====================================================
Free pages threshold: 19226
Number of Free pages: 118312
Number of pages on buffer cache: 0

IFL Index space usage
=====================================================
MAX Number of Logical Interfaces(IFL's): 262144
Number of Logical Interfaces in use: 23
Number of Invalidated Logical Interfaces: 0


NH Index Space Allocation
=====================================================
    Index_Space_type             Used        Available
            Reserved               64             1280
             Private               39              665
             Regular               60          1046466
            Extended                0         15728638

Maximum possible number of private nexthops: 704
Percentage of private nexthop index space reserved for IRI:(internal routing) 30
Number of nexthop indices reserved for IRI:(internal routing) 212

Printing gencfg blobs statistics by major type

                                Type   Num Blobs   Data size
                      Proxy gencfg           33           0
                 L2 forwarding blob           1         128
                           COS BLOB          41       36996
                    Feature License           1           2
                           Firewall          21        4044
                            CH Info           1           0
              Libpeer PeerInfo blob           2        4392
                        Resync blob           4          44
                        MACSEC blob           1           8
                        Agentd blob           4        5888
                      RPD EVPN blob           1          52
```

```
                              LFM Blob           1           8
                                 Total         111       51562
```

show system khms-stats detail

```
user@host> show system khms-stats detail
Memory usage by ifstates, routes, nexthops, gencfg blobs
==============================================================================
Memory type  Total no. of states     Memory used  Percentage of Dead states
     Ifstate                885           525KB             0%
      Routes                150            25KB             0%
     Nexthop                163            39KB             0%

Number of delayed unrefs of RT and NH: 0 (Max allowed 40000)

Number of delayed weight unrefs of RT and NH: 0 (Max allowed 400000)

Mbufs/Cluster Usage
=======================================================
852/1938/2790 mbufs in use (current/cache/total)
772/1198/1970/250436 mbuf clusters in use (current/cache/total/max)
769/1002 mbuf+clusters out of packet secondary zone in use (current/cache)
0/19/19/125218 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/37101 9k jumbo clusters in use (current/cache/total/max)
0/0/0/20869 16k jumbo clusters in use (current/cache/total/max)
1757K/2956K/4713K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0 requests for sfbufs denied
0 requests for sfbufs delayed
0 requests for I/O initiated by sendfile

Memory Pages stats
=======================================================
Free pages threshold: 19226
Number of Free pages: 118117
Number of pages on buffer cache: 0

IFL Index space usage
=======================================================
MAX Number of Logical Interfaces(IFL's): 262144
```

```
Number of Logical Interfaces in use: 23
Number of Invalidated Logical Interfaces: 0


Kernel Health Parameters:
------------------------
  Record Time                       : 1568295214   : Thu Sep 12 19:03:34 2019

  Security:
    Veriexec                        : 3            : VERIEXEC_STATE_ACTIVE_AND_LOADED

   Ifstate:
    Alive Ifstate count             : 883
    Dead Ifstate count              : 0
    Delayed Unrefs count            : 0
    Delayed Unrefs max              : 40000
    Stuck Ifstate clients           : 0
    Alive Ifstate clients           : 40
    Dead Ifstate clients            : 0
    Ifstate client limit reached    : 0

  GRES:
    GRES time (mastership acquired) : 0            : NA
    Protocol Slave Connect time     : 0            : NA
    GRES Configured state           : 0            : NOT_CONFIGURED
    Master Kernel GRES Ready        : 0            : GRES_NOT_READY
    Slave Kernel GRES Ready         : 0            : GRES_NOT_READY
    GRES Error state                : 0            : KSYNCD_NO_ERROR
    GRES Other RE Present           : 1            : OTHER_RE_PRESENT
    GRES Other RE Alive             : 1            : OTHER_RE_ALIVE
    GRES Is Protocol Master         : 1            : PROTOCOL_MASTER

  ISSU:
    ISSU Failure stage              : 0            : ISSU_STATE_IDLE
    Current ISSU stage              : 0            : ISSU_STATE_IDLE

  Peer Infra:
    PFEMAN connection drops         : 0
    Spurious PPT wakeups            : 0

  TNP:
    tnp hello drop count            : 0
    tnp fragment drop count         : 1
```

```
VPLS:
  unicast token count             : 0
  unicast token max               : 131072
  flood token count               : 1
  flood token max                 : 327680

TUNNEL:
  rpf tunnelid count              : 0
  rpf tunnelid MAX                : 1048575
  non-rpf tunnelid count          : 0
  non-rpf tunnelid MAX            : 15728640
  looped count                    : 0

MULTICAST:
  iif mismatch error count        : 0
  resolve request error count     : 0

RTSOCK:
  RTSOCK Total Veto               : 0
  RTSOCK Total Error              : 57

Aggregated Devices:
  PS LT Unstacking error count    : 0
  PS IFL l2circuit down count     : 0
  RLT LP Backup link down count   : 0
  AE Unstacking error count       : 0
  AE LP Backup link down count    : 0

TCPIP:
  Maximum Public nexthops         : 1046526
  Used  Public nexthops           : 68
  Maximum Private nexthops        : 704
  Used Private nexthops           : 35
  RTB Clone Route cnt             : 6
  RTB Clone Route Max             : 1000000
  ARP cache iri max               : 200
  ARP cache mgmt max              : 14960
  ARP cache public max            : 59840
  ARP cache iri cnt               : 4
  ARP cache mgmt cnt              : 2
  ARP cache public cnt            : 0
  ARP cache iri drop cnt          : 0
```

```
    ARP cache mgmt drop cnt         : 0
    ARP cache public drop cnt       : 0
    NDP cache iri max               : 200
    NDP cache mgmt max              : 14960
    NDP cache public max            : 59840
    NDP cache iri cnt               : 3
    NDP cache mgmt cnt              : 0
    NDP cache public cnt            : 0
    NDP cache iri drop cnt          : 0
    NDP cache mgmt drop cnt         : 0
    NDP cache public drop cnt       : 0
    Netisr queue ether wm cnt       : 0
    Netisr queue ether drop cnt     : 0
    Netisr queue ether pkt handled  : 524311
    Netisr queue ether pkt queued   : 524311
    Netisr queue ip wm cnt          : 0
    Netisr queue ip drop cnt        : 0
    Netisr queue ip pkt handled     : 451581
    Netisr queue ip pkt queued      : 451581
    Netisr queue ip6 wm cnt         : 0
    Netisr queue ip6 drop cnt       : 0
    Netisr queue ip6 pkt handled    : 2
    Netisr queue ip6 pkt queued     : 2
    Netisr queue arp wm cnt         : 0
    Netisr queue arp drop cnt       : 0
    Netisr queue arp pkt handled    : 50420
    Netisr queue arp pkt queued     : 50420
    TCP DDOS attack count           : 0
    TCP Connection drop count       : 19
    TCP TIME WAIT connection count  : 0

  Non Stop Routing:
    JSR Split failure count         : 0
    JSR Merge failure count         : 0
    JSR PRL Queue full count        : 0

NH Index Space Allocation
=========================================================
    Index_Space_type             Used        Available
             Reserved              64             1280
              Private              39              665
              Regular              60          1046466
            Extended               0         15728638
```

```
Maximum possible number of private nexthops: 704
Percentage of private nexthop index space reserved for IRI:(internal routing) 30
Number of nexthop indices reserved for IRI:(internal routing) 212

Printing gencfg blobs statistics by major type

                              Type   Num Blobs   Data size
                     Proxy gencfg          33           0
                L2 forwarding blob           1         128
                          COS BLOB          41       36996
                   Feature License           1           2
                          Firewall          21        4044
                           CH Info           1           0
             Libpeer PeerInfo blob           2        4392
                       Resync blob           4          44
                       MACSEC blob           1           8
                       Agentd blob           4        5888
                     RPD EVPN blob           1          52
                          LFM Blob           1           8
                             Total         111       51562
```

Release Information

Command introduced in Junos OS 19.3R1.

show system license (View)

IN THIS SECTION

- Syntax | 26897
- Description | 26897
- Options | 26897
- Required Privilege Level | 26897
- Output Fields | 26898
- Sample Output | 26899

- Release Information | **26901**

Syntax

```
show system license
<bandwitdh | installed | fpc | keys | status | usage>
```

Description

Display licenses and information about how licenses are used.

Options

none Display all license information.

bandwidth (Optional) Display a summary of advanced and premium bandwidth used by flex capable FPCs.

installed (Optional) Display installed licenses only.

fpc (Optional) Display premium and advanced bandwidth usage for Flex-Only and flex-capable FPCs.

keys (Optional) Display a list of license keys. Use this information to verify that each expected license key is present.

status (Optional) Display license status for a specified logical system or for all logical systems.

usage (Optional) Display the state of licensed features.

Required Privilege Level

view

Output Fields

Table 1635 on page 26898 lists the output fields for the `show system license` command. Output fields are listed in the approximate order in which they appear.

Table 1635: show system license Output Fields

Field Name	Field Description
`Feature name`	Name assigned to the configured feature. You use this information to verify that all the features for which you installed licenses are present.
`Licenses used`	Number of licenses used by the device. You use this information to verify that the number of licenses used matches the number configured. If a licensed feature is configured, the feature is considered used.
`Licenses installed`	Information about the installed license key: • `License identifier`—Identifier associated with a license key. • `License version`—Version of a license. The version indicates how the license is validated, the type of signature, and the signer of the license key. • `Valid for device`—Device that can use a license key. • `Features`—Feature associated with a license.
`Licenses needed`	Number of licenses required for features being used but not yet properly licensed.
`Expiry`	Time remaining in the grace period before a license is required for a feature being used.
`Logical system license status`	Displays whether a license is enabled for a logical system.

Sample Output

show system license

```
user@host> show system license

License usage:
                                 Licenses     Licenses    Licenses    Expiry
  Feature name                       used    installed      needed
  av_key_kaspersky_engine               1            1           0    2012-03-30 01:00:00 IST
  wf_key_surfcontrol_cpa                0            1           0    2012-03-30 01:00:00 IST
  dynamic-vpn                           0            1           0    permanent
  ax411-wlan-ap                         0            2           0    permanent

Licenses installed:
  License identifier: JUNOS301998
  License version: 2
  Valid for device: AG4909AA0080
  Features:
    av_key_kaspersky_engine - Kaspersky AV
      date-based, 2011-03-30 01:00:00 IST - 2012-03-30 01:00:00 IST

  License identifier: JUNOS302000
  License version: 2
  Valid for device: AG4909AA0080
  Features:
    wf_key_surfcontrol_cpa - Web Filtering
      date-based, 2011-03-30 01:00:00 IST - 2012-03-30 01:00:00 IST
```

show system license installed

```
user@host> show system license installed

License identifier: JUNOS301998
  License version: 2
  Valid for device: AG4909AA0080
  Features:
    av_key_kaspersky_engine - Kaspersky AV
      date-based, 2011-03-30 01:00:00 IST - 2012-03-30 01:00:00 IST
```

```
  License identifier: JUNOS302000
  License version: 2
  Valid for device: AG4909AA0080
  Features:
    wf_key_surfcontrol_cpa - Web Filtering
      date-based, 2011-03-30 01:00:00 IST - 2012-03-30 01:00:00 IST
```

show system license keys

```
user@host> show system license keys

XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
            xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
            xxxxxx xxxxxx xxx
```

show system license usage

```
user@host> show system license usage

                                   Licenses     Licenses    Licenses  Expiry
Feature name                           used        installed      needed
  av_key_kaspersky_engine               1            1           0    2012-03-30 01:00:00 IST
  wf_key_surfcontrol_cpa                0            1           0    2012-03-30 01:00:00 IST
  dynamic-vpn                           0            1           0    permanent
  ax411-wlan-ap                         0            2           0    permanent
```

show system license status logical-system all

```
user@host> show system license status logical-system all
Logical system license status:

logical system name                 license status
root-logical-system                 enabled
LSYS0                               enabled
LSYS1                               enabled
LSYS2                               enabled
```

show system license bandwidth

```
user@host> show system license bandwidth
Advance Bandwidth Usage: 110
Premium Bandwidth Usage: 0
```

show system license fpc

```
user@host> show system license fpc
FPC                 Serial Number
    5                 ABC123
Advance Bandwidth Usage: 100
Premium Bandwidth Usage: 0

Flex-capable linecards:
  FPC                 Serial Number
    1                 ABC124
Advance Bandwidth Usage: 10
Premium Bandwidth Usage: 0
```

Release Information

Command introduced in Junos OS Release 9.5. Logical system status option added in Junos OS Release 11.2. Support for flex FPC and bandwidth licensing added in Junos OS Release 22.1 for MX platforms.

RELATED DOCUMENTATION

Adding New Licenses (CLI Procedure)

show system login lockout

IN THIS SECTION

- Syntax | 26902
- Description | 26902
- Required Privilege Level | 26902
- Output Fields | 26902
- Sample Output | 26903
- Release Information | 26903

Syntax

```
show system login lockout
```

Description

This command displays the usernames locked after unsuccessful login attempts.

Required Privilege Level

view and system

Output Fields

Table 1636 on page 26902 lists the output fields for the `show system login lockout` command. Output fields are listed in the approximate order in which they appear.

Table 1636: show system login lockout

Field Name	Field Description	Level of Output
`User`	Username	All levels

Table 1636: show system login lockout *(Continued)*

Field Name	Field Description	Level of Output
`Lockout start`	Date and time the username was locked	All levels
`Lockout end`	Date and time the username was unlocked	All levels

Sample Output

show system login lockout

```
user@host> show system login lockout

User                 Lockout start           Lockout end
root                 2011-05-11 09:11:15 UTC 2011-05-11 09:13:15 UTC
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

retry-options

clear system login lockout | 4225

show system mac-addresses (JDM)

IN THIS SECTION

- Syntax | 26904
- Description | 26904

- Required Privilege Level | 26904
- Output Fields | 26904
- Sample Output | 26906
- Release Information | 26906

Syntax

```
show system mac-addresses
```

Description

Display information related to the Juniper Device Manager (JDM) management interface MAC address and available MAC address range.

Required Privilege Level

View

Output Fields

Table 1637 on page 26904 describes the output fields for the `show system mac-addresses` command. Output fields are listed in the approximate order in which they appear.

Table 1637: show system mac-addresses Output Fields

Field Name	Field Description
Management Interface MAC Addresses	
Domain	This could be a Juniper Device Manager (JDM) domain or guest network function (GNF) domain.
Interface	Interface name.
Management-MAC	The MAC address of the JDM or GNF.

Table 1637: show system mac-addresses Output Fields *(Continued)*

Field Name	Field Description
Donor-GNF-ID	Displays name of the GNF which is the donor for the JDM management interface (jmgmt0) MAC address. This field is marked as not applicable (NA) for the GNFs listed in the column 'Management Interface MAC Addresses'. **NOTE**: • If there are no licensed GNFs in the system, no donor will be available for the `jmgmt0` MAC address. In that case, a virtual license is used and this field is marked as '-'. • For in-chassis Junos node slicing, `jmgmt0` and `fxp0` MAC addresses come from Juniper's licensed address space available on BSYS, not from a licensed GNF.
VNF MAC Address Pool	
Base-Mac-Address	Starting address of the MAC pool.
Size	Size of the MAC address pool.
Licensed	Indicates if the GNF is licensed or not. The following are the possible options: • Y • N
Donor	Indicates if the GNF is donor for jmgmt0 MAC address. The following are the possible options: • Y • N

Sample Output

show system mac-addresses

```
user@jdm> show system mac-addresses

Management Interface MAC Addresses
------------------------------------------------------------------------------------------
Domain                                        Interface   Management-MAC     Donor-GNF-ID
----------------------------------------      ---------   ------------------  ------------
JDM                                           jmgmt0      7C:E2:CA:FF:36:73   NA
gnf1                                          fxp0        7C:E2:CA:FF:36:74   NA
gnf3                                          fxp0        7C:E2:CA:FF:36:75   NA

MAC Blocks
--------------------------------------------------
GNF-ID  Base-Mac-Address    Size   Licensed  Donor
------  ------------------  -----  --------  -----
1       aa:bb:cc:dd:ee:00   4096   Y         N
3       00:90:69:3a:40:00   4096   N         N
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | **15481**

request server authenticate-peer-server | **14901**

show system max-smbios-version (JDM)

IN THIS SECTION

- Syntax | 26907
- Description | 26907
- Required Privilege Level | 26907
- Sample Output | 26907
- Release Information | 26907

Syntax

```
show system max-smbios-version
```

Description

Display the latest SMBIOS version supported by the Juniper Device Manager (JDM).

Required Privilege Level

View

Sample Output

show system max-smbios-version

```
user@jdm> show system max-smbios-version
Version : v2
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | 15481

request server authenticate-peer-server | 14901

show system memory

IN THIS SECTION

- Syntax | 26908
- Syntax (EX Series Switches) | 26908
- Syntax (MX Series Routers) | 26909
- Syntax (QFX Series) | 26909
- Syntax (Junos OS Evolved) | 26909
- Description | 26909
- Options | 26909
- Required Privilege Level | 26910
- Output Fields | 26910
- Sample Output | 26912
- Release Information | 26918

Syntax

```
show system memory
```

Syntax (EX Series Switches)

```
show system memory
<all-members>
```

```
<local>
<member member-id>
```

Syntax (MX Series Routers)

```
show system memory
<all-members>
<local>
<member member-id>
```

Syntax (QFX Series)

```
show system memory
<infrastructure name>
<interconnect-device name>
<node-group name>
```

Syntax (Junos OS Evolved)

```
show system memory
<node name>
```

Description

Display system-wide memory distribution and usage including the Junos kernel, software processes, and memory disks. Use the `show system memory` command for troubleshooting with Juniper Networks Customer Support.

Options

none Display the system memory distribution and usage information.

all-members (EX4200 switches and MX Series routers only) (Optional) Display system memory distribution and usage information for all members of the Virtual Chassis configuration.

infrastructure *name* (QFabric switches only) (Optional) Display system memory distribution and usage information for the fabric control Routing Engine and fabric manager Routing Engine.

interconnect-device *name* (QFabric switches only) (Optional) Display system memory distribution and usage information for the Interconnect device.

local (EX4200 switches and MX Series routers only) (Optional) Display system memory distribution and usage information for the local Virtual Chassis member.

member *member-id* (EX4200 switches and MX Series routers only) (Optional) Display system memory distribution and usage information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node *name* (Junos OS Evolved) (Optional) Display system memory distribution and usage information for the specified node.

node-group *name* (QFabric switches only) (Optional) Display system memory distribution and usage information for the Node group.

Required Privilege Level

view

Output Fields

Table 1638 on page 26911 lists the output fields for the `show system memory` command. Output fields are listed in the approximate order in which they appear.

Table 1638: show system memory Output Fields

Field Name	Field Description
`System memory usage distribution`	Amount of memory used by the system, represented in kilobytes and as a percentage of the total memory. • Total memory—Actual memory available to the system at the time of system initialization. • Reserved memory—Amount of total memory reserved for the system at the time of system initialization, but not managed by virtual memory sub-system. • Wired memory—Amount of total memory wired in the system that never paged out. • Active memory—Amount of total memory actively in use. • Inactive—Amount of total memory in use, but not referenced recently. • Used memory—Amount of total memory used. • Free—Amount of total memory available for allocation. • Buffer—Amount of memory used as buffers. • Swap—Amount of virtual memory that is swapped out to a block disk. • Available—Estimate of how much memory is available for starting new applications without swapping memory. This estimate takes into account that the system needs some page cache to function well. Not all reclaimable slab is reclaimable because certain items are being used. • Cached—Amount of total memory almost available for allocation.
`Memory disk resident memory`	Memory used by the memory disks, but not accounted for in the kernel map. For example, swap-based memory disk.

Table 1638: show system memory Output Fields *(Continued)*

Field Name	Field Description
VM-Kbytes	Virtual memory used as per all virtual mappings in the kernel or process map and percentage thereof based on maximum addressable virtual memory as per any kernel, system, architecture, or resource limits.
Resident	Real memory used as per actual resident pages of all mappings in the kernel or process map and percentage thereof based on the total real memory available in the system.
Map-name	Name of the map.
PID	Process ID.
Process-name	Name of the process. (MX Series routers only) When you display the memory usage for an MX Series Virtual Chassis, the `show system memory` command does not display information pertaining to the relayd (/usr/sbin/relayd) process.

Sample Output

show system memory

```
user@host> show system memory
System memory usage distribution:
       Total memory: 1035552 Kbytes (100%)
    Reserved memory:   18684 Kbytes (  1%)
       Wired memory:  526072 Kbytes ( 50%)
      Active memory:  154792 Kbytes ( 14%)
    Inactive memory:   64664 Kbytes (  6%)
       Cache memory:   69840 Kbytes (  6%)
        Free memory:  200688 Kbytes ( 19%)
Memory disk resident memory:   19016 Kbytes
VM-Kbytes(  %  ) Resident(  %  ) Map-name
```

```
   570580(54.36)    92704(08.95) kernel
Pid     VM-Kbytes(  %  ) Resident(  %  ) Process-name
   1451      4756(00.15)     1000(00.09) /sbin/pmap
   1411    727012(23.14)    18008(01.73) mgd: (mgd) (root)/dev/ttyd0
---(more)---
```

show system memory | display xml (QFX Series)

```
user@switch> show system-memory | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/12.1I0/junos">
    <system-memory-information>
        <system-memory-summary-information>
            <system-memory-total>3391488</system-memory-total>
            <system-memory-total-percent>100%</system-memory-total-percent>
            <system-memory-reserved>1619892</system-memory-reserved>
            <system-memory-reserved-percent> 47%</system-memory-reserved-percent>
            <system-memory-wired>1592936</system-memory-wired>
            <system-memory-wired-percent> 46%</system-memory-wired-percent>
            <system-memory-active> 177736</system-memory-active>
            <system-memory-active-percent>  5%</system-memory-active-percent>
            <system-memory-inactive>  60480</system-memory-inactive>
            <system-memory-inactive-percent>  1%</system-memory-inactive-percent>
            <system-memory-cache>  39240</system-memory-cache>
            <system-memory-cache-percent>  1%</system-memory-cache-percent>
            <system-memory-free>1244628</system-memory-free>
            <system-memory-free-percent> 36%</system-memory-free-percent>
        </system-memory-summary-information>
        <memory-disk-resident-information>
            <resident-memory>  16424</resident-memory>
        </memory-disk-resident-information>
        <output>
            pmap: unable to load pmap_helper module: No such file or directory
        </output>
        <output>
            pmap: unable to load pmap_helper module: No such file or directory
        </output>
    </system-memory-information>
    <cli>
        <banner></banner>
```

```
    </cli>
</rpc-reply>
```

show system memory (Junos OS Evolved)

```
user@host> show system memory
fpc1:
--------------------------------------------------------------------------
System memory usage distribution:
       Total memory: 16125892 Kbytes (100%)
      Active memory: 2447108 Kbytes (15%)
    Inactive memory: 2099532 Kbytes (13%)
        Free memory: 9329768 Kbytes (57%)
Pid     VM-Kbytes(  %  ) Resident(  %  ) Process-name
      1    159116(    0)     8408(    0) /lib/systemd/systemd
      2         0(    0)        0(    0) [kthreadd]
      3         0(    0)        0(    0) [rcu_gp]
...

re0:
--------------------------------------------------------------------------
System memory usage distribution:
       Total memory: 16125576 Kbytes (100%)
      Active memory: 4935984 Kbytes (30%)
    Inactive memory: 8935956 Kbytes (55%)
        Free memory: 178448 Kbytes (1%)
Pid     VM-Kbytes(  %  ) Resident(  %  ) Process-name
      1    162220(    0)    10780(    0) /sbin/init
      2         0(    0)        0(    0) [kthreadd]
      3         0(    0)        0(    0) [rcu_gp]
...

re1:
--------------------------------------------------------------------------
System memory usage distribution:
       Total memory: 16125560 Kbytes (100%)
      Active memory: 3834540 Kbytes (23%)
    Inactive memory: 5579724 Kbytes (34%)
        Free memory: 5500228 Kbytes (34%)
Pid     VM-Kbytes(  %  ) Resident(  %  ) Process-name
      1    160928(    0)     9860(    0) /sbin/init
```

```
      2         0(    0)        0(    0) [kthreadd]
      3         0(    0)        0(    0) [rcu_gp]
...
```

Starting in Junos OS Evolved Release 23.1R1, the `Buffer memory`, `Swap memory`, and `Available memory` fields have been added to the `show system memory` command.

show system memory (Junos OS Evolved)

```
user@host> show system memory
Total memory: 16062004 Kbytes (100%)
Active memory: 3152044 Kbytes (13%)
Inactive memory: 4751112 Kbytes (61%)
Free memory: 6602668 Kbytes (41%)
Buffer memory: 370760 Kbytes (2%)
Swap memory: 6045820 Kbytes (38%)
Available memory: 12610560 Kbytes (79%)
Cached memory: 5694904 Kbytes (36%)
```

show system memory | display xml (Junos OS Evolved)

```
user@host> show system memory | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/22.1R0/junos">
    <multi-routing-engine-results>
        <multi-routing-engine-item>
            <re-name>fpc1</re-name>
            <system-memory-information>
                <system-memory-summary-information>
                    <system-memory-total>16125892</system-memory-total>
                    <system-memory-total-percent>100%</system-memory-total-percent>
                    <system-memory-used>3885112</system-memory-used>
                    <system-memory-used-percent>24%</system-memory-used-percent>
                    <system-memory-active>2447796</system-memory-active>
                    <system-memory-active-percent>15%</system-memory-active-percent>
                    <system-memory-inactive>2101128</system-memory-inactive>
                    <system-memory-inactive-percent>13%</system-memory-inactive-percent>
                    <system-memory-free>9327316</system-memory-free>
                    <system-memory-free-percent>57%</system-memory-free-percent>
                    <system-memory-buffer>314516</system-memory-buffer>
                    <system-memory-buffer-percent>1%</system-memory-buffer-percent>
```

```
                <system-memory-swap>2598948</system-memory-swap>
                <system-memory-swap-percent>16%</system-memory-swap-percent>
            </system-memory-summary-information>
            <pmap-terse-information xmlns="http://xml.juniper.net/fbsd10/14.2I0/junos-pmap">
                <pmap-terse-summary junos:style="pmap-process-terse-summary">
                    <pid>1</pid>
                    <process-name>/lib/systemd/systemd</process-name>
                    <size>159116</size>
                    <size-percent>0</size-percent>
                    <resident>8408</resident>
                    <resident-percent>0</resident-percent>
                </pmap-terse-summary>
                ...
            </pmap-terse-information>
        </system-memory-information>
    </multi-routing-engine-item>
    <multi-routing-engine-item>
        <re-name>re0</re-name>
        <system-memory-information>
            <system-memory-summary-information>
                <system-memory-total>16125576</system-memory-total>
                <system-memory-total-percent>100%</system-memory-total-percent>
                <system-memory-used>6912492</system-memory-used>
                <system-memory-used-percent>42%</system-memory-used-percent>
                <system-memory-active>4936580</system-memory-active>
                <system-memory-active-percent>30%</system-memory-active-percent>
                <system-memory-inactive>8939976</system-memory-inactive>
                <system-memory-inactive-percent>55%</system-memory-inactive-percent>
                <system-memory-free>170744</system-memory-free>
                <system-memory-free-percent>1%</system-memory-free-percent>
                <system-memory-buffer>524676</system-memory-buffer>
                <system-memory-buffer-percent>3%</system-memory-buffer-percent>
                <system-memory-swap>8517664</system-memory-swap>
                <system-memory-swap-percent>52%</system-memory-swap-percent>
            </system-memory-summary-information>
            <pmap-terse-information xmlns="http://xml.juniper.net/fbsd10/14.2I0/junos-pmap">
                <pmap-terse-summary junos:style="pmap-process-terse-summary">
                    <pid>1</pid>
                    <process-name>/sbin/init</process-name>
                    <size>162220</size>
                    <size-percent>0</size-percent>
                    <resident>10780</resident>
                    <resident-percent>0</resident-percent>
```

```
                    </pmap-terse-summary>
                    ...
                </pmap-terse-information>
            </system-memory-information>
        </multi-routing-engine-item>
        ...
    </multi-routing-engine-results>
```

show system memory (Junos OS Evolved) (shows 'used memory')

```
user@host> show system memory
fpc0:
--------------------------------------------------------------------------
System memory usage distribution:
        Total memory: 16092800 Kbytes (100%)
       Active memory: 367208 Kbytes (2%)
     Inactive memory: 3866136 Kbytes (24%)
         Used memory: 5876340 Kbytes (36%)
         Free memory: 10216460 Kbytes (63%)
       Buffer memory: 17868 Kbytes (0%)
         Swap memory: 0 Kbytes (0%)
    Available memory: 12437340 Kbytes (77%)
       Cached memory: 2510420 Kbytes (15%)
Pid     VM-Kbytes(  %  ) Resident(  %  ) Process-name
      1     94084(    0)    10232(    0) /lib/systemd/systemd
      2         0(    0)        0(    0) [kthreadd]
      3         0(    0)        0(    0) [rcu_gp]
      4         0(    0)        0(    0) [rcu_par_gp]
      5         0(    0)        0(    0) [slub_flushwq]
      6         0(    0)        0(    0) [netns]
      8         0(    0)        0(    0) [kworker/0:0H-events_highpri]
     10         0(    0)        0(    0) [mm_percpu_wq]
     11         0(    0)        0(    0) [rcu_tasks_kthre]

...

re0:
--------------------------------------------------------------------------
System memory usage distribution:
        Total memory: 16026056 Kbytes (100%)
       Active memory: 1075256 Kbytes (6%)
```

```
     Inactive memory: 5483160 Kbytes (34%)
         Used memory: 7606360 Kbytes (47%)
         Free memory: 8419696 Kbytes (52%)
       Buffer memory: 28356 Kbytes (0%)
         Swap memory: 0 Kbytes (0%)
    Available memory: 12931772 Kbytes (80%)
       Cached memory: 4736032 Kbytes (29%)
Pid     VM-Kbytes(  %  ) Resident(  %  ) Process-name
      1     97040(    0)     6688(    0) /sbin/init
      2         0(    0)        0(    0) [kthreadd]
      3         0(    0)        0(    0) [rcu_gp]
      4         0(    0)        0(    0) [rcu_par_gp]
      5         0(    0)        0(    0) [slub_flushwq]
      6         0(    0)        0(    0) [netns]
      8         0(    0)        0(    0) [kworker/0:0H-kblockd]
     10         0(    0)        0(    0) [mm_percpu_wq]
     11         0(    0)        0(    0) [rcu_tasks_kthre]
     12         0(    0)        0(    0) [rcu_tasks_rude_]
     13         0(    0)        0(    0) [rcu_tasks_trace]


...
```

Release Information

Command introduced in Junos OS Release 12.1.

show system memory (JDM)

IN THIS SECTION

- Syntax | 26919
- Description | 26919
- Required Privilege Level | 26919
- Output Fields | 26919
- Sample Output | 26920
- Release Information | 26920

Syntax

```
show system memory
```

Description

Display the memory usage information about the host server, Juniper Device Manager (JDM), and guest network functions (GNF).

Required Privilege Level

View

Output Fields

Table 1639 on page 26919 describes the output fields for the `show system memory` command. Output fields are listed in the approximate order in which they appear.

Table 1639: show system memory Output Fields

Field Name	Field Description
Memory Usage Information	Displays the following memory usage information about host server and JDM: • Total—Total memory. • Used—Used memory. • Free—Available memory.
VNF Memory Information	Displays the following memory information about the GNFs: • Name—GNF name. • Resident—The memory used by the GNFs. • Actual—Actual memory.

Sample Output

show system memory

```
user@jdm> show system memory

Memory Usage Information
---------------------------
       Total  Used   Free
       ------ ------ ------
Host:  251G   14G    219G

JDM :  2.0G   2.0G   0

VNF Memory Information
-----------------------------------------------------------------
Name                                              Actual Resident
------------------------------------------------- ------ --------
gnf1                                              16.0G  13.6G
gnf2                                              16.0G  11.6G
gnf3                                              16.0G  12.6G
gnf4                                              16.0G  12.0G
gnf5                                              16.0G  12.7G
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | **15481**

request server authenticate-peer-server | **14901**

show system monitor memory actions

IN THIS SECTION

- Syntax | 26921
- Description | 26921
- Options | 26921
- Required Privilege Level | 26922
- Output Fields | 26922
- Sample Output | 26922
- Release Information | 26923

Syntax

```
show system monitor memory actions (all | system | process)
```

Description

Displays the history of actions performed when the threshold level for any event is crossed. You can view the history of actions for a process or for a system or both.

Options

all Display the history of actions performed when the threshold level for any event is crossed for the entire system as well as per process.

system Display the history of actions performed when the threshold level for any event is crossed at the system level.

process Display the history of actions performed when the threshold level for any event is crossed at the process level.

Required Privilege Level

admin

Output Fields

Table 1640 on page 26922 lists the output fields for the `show system monitor memory actions` command. Output fields are listed in the approximate order in which they appear.

Table 1640: show system monitor memory action Output Fields

Field Name	Field Description	Level of Output

Category	Specifies if the action is generated at the system level or at the process level.	`all`
		`process`
		`system`

process	Name of the process that has crossed the threshold value and process identification number used to identify the process that crossed the threshold value.	`all` `process`

Action	The actions performed when a threshold is crossed either at the system level or at the process level.	`all`
		`process`
		`system`

Sample Output

show system monitor memory actions all

```
user@host> show system monitor memory actions all

May 08 10:56:00 [process-action]: process: ('xxxxx', 'xxx') action: snmp-trap
May 08 10:56:02 [process-action]: process: ('xxxxx', 'xxx') action: snmp-trap
```

show system monitor memory actions process

```
user@host> show system monitor memory actions process

May 08 10:56:00 [process-action]: process: ('11727', 'dcd') action: snmp-trap
May 08 10:56:02 [process-action]: process: ('11716', 'jkhmd') action: snmp-trap
```

Release Information

Command introduced in Junos OS Release 20.2R1

show system monitor memory events

IN THIS SECTION

- Syntax | 26923
- Description | 26923
- Options | 26924
- Required Privilege Level | 26924
- Output Fields | 26924
- Sample Output | 26925
- Release Information | 26925

Syntax

```
show system monitor memory events (all | system | process)
```

Description

Displays the history of events, whether they are major, minor, or critical. The events can be either for the system or per process. Use this command to view all the recorded events. When the condition is first detected, an event is raised. If the condition persists with the same severity, the event is not

reported multiple times. However, if the condition persists with a changed severity or is detected after a reset, the event is reported again.

Options

all Display the history of events for the entire system as well as per process.

system Display the history of events at the system level.

process Display the history of events per process.

Required Privilege Level

admin

Output Fields

Table 1641 on page 26924 lists the output fields for the `show system monitor memory events` command. Output fields are listed in the approximate order in which they appear.

Table 1641: show system monitor memory events Output Fields

Field Name	Field Description	Level of Output

Category	Specifies if the event is generated at the system level or at the process level.	`all`
		`process`
		`system`

Event	The Event that is raised when a threshold is crossed either at the system level or at the process level.	`all`
		`process`
		`system`

Type	Type of event. Values are: major, minor, and critical.	`all`
		`process`
		`system`

for process	Name of the process that has crossed the threshold value.	`all` `process`

Table 1641: show system monitor memory events Output Fields *(Continued)*

Field Name	Field Description	Level of Output
with pid	Process identification number used to identify the process that crossed the threshold value.	`all` `process`

Sample Output

show system monitor memory events all

```
user@host> show system monitor memory events all

Feb 24 10:08:21 [system] Event: free_page_event Type minor
Feb 24 13:42:13 [process-threshold] Event: critical-event for process: xxxxx with pid xxxxx
Feb 24 17:15:32 [system] Event: free_page_event Type minor
Feb 25 11:47:50 [process-threshold] Event: critical-event for process: xxxxx with pid xxxxx
```

show system monitor memory events process

```
user@host> show system monitor memory events process
Feb 24 10:08:21 [system] Event: free_page_event Type minor
Feb 24 17:15:32 [system] Event: free_page_event Type minor
```

show system monitor memory events system

```
user@host> show system monitor memory events system
Feb 24 13:42:13 [process-threshold] Event: critical-event for process: xxxxx with pid xxxxx
Feb 25 11:47:50 [process-threshold] Event: critical-event for process: xxxxx with pid xxxxx
```

Release Information

Command introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

show system monitor memory status | 26926

show system monitor memory status

IN THIS SECTION

- Syntax | 26926
- Description | 26926
- Options | 26926
- Required Privilege Level | 26927
- Output Fields | 26927
- Sample Output | 26927
- Release Information | 26928

Syntax

```
show system monitor memory status (all | system | process)
```

Description

Displays the current instant snapshot of the health of the device. You can view current health of both the modules: system and process.

Options

all Display the memory usage status of the entire system as well as per process.

system Display the memory usage status at the system level.

process Display the memory usage status per process.

Required Privilege Level

admin

Output Fields

Table 1642 on page 26927 lists the output fields for the `show system monitor memory` command. Output fields are listed in the approximate order in which they appear.

Table 1642: show system monitor memory Output Fields

Field Name	Field Description	Level of Output
No.	Number	`all` `process`
Resource	Resource that is monitored. For instance, Cache and free pages, memory utilization at the kernel level, swap usage, and others.	`all` `system`
processID	The Process identification number of the process that has crossed the threshold value.	`all` `process`
processName	Name of the process that has crossed the threshold value.	`all` `process`
Alarm-Type	The type of alarm generated. Values are: major-event, minor-event, and critical-event.	`all` `process`
Event-Type	The type of event raised. Values are: major-event, minor-event, and critical-event.	`all` `system`

Sample Output

show system monitor memory status all

```
user@host> show system monitor memory status all
-----------------------------------------------
```

```
No.      (processID, processName)  Alarm-Type
-----------------------------------------------
1.       ('xxxxx', 'xxxxx')        critical-event
---------------------------------------------
Resource                      Event-Type
---------------------------------------------
Free pages:                   minor-event
```

NOTE: This is not applicable for PTX3000 Series routers.

show system monitor memory status process

```
user@host> show system monitor memory status process
-----------------------------------------------
No.      (processID, processName)  Alarm-Type
-----------------------------------------------
1.       ('xxxxx', 'xxxxx')        critical-event
```

show system monitor memory status system

```
user@host> show system monitor memory status system
---------------------------------------------
Resource                      Event-Type
---------------------------------------------
Free pages:                   minor-event
```

Release Information

Command introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

show system monitor memory events | 26923

show system name-resolution

IN THIS SECTION

- Syntax | 26929
- Description | 26929
- Options | 26929
- Required Privilege Level | 26929
- Output Fields | 26929
- show system name-resolution | 26930
- Release Information | 26931

Syntax

```
show system name-resolution
```

Description

Display hostname-to-IP-address mappings.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1643 on page 26930 lists the output fields for the `show system name-resolution` command. Output fields are listed in the approximate order in which they appear.

Table 1643: show system name-resolution Output Fields

Field Name	Field Description
Last update	Date and time when the hostname-to-IP address mapping were last resolved.
Refresh interval	Interval for refreshing the cache with the updated hostname-to-IP address mappings.
Addresses	Resolved IP addresses based on the hostname-to-IP address mappings.
Error	Error message displayed if there is a DNS hostname lookup failure.
Last change	Timestamp for the last change in the hostname-to-IP address mappings.

show system name-resolution

command-name

```
user@host> show system name-resolution

Hostname to IP-address mappings:
--------------------------------------------------
 Last update: Mon Sep 29 18:42:21 2008
 Refresh interval: 600 secs
 Host: ntp1
      Addresses:
        3.3.3.11
      Last change: Mon Sep 29 18:42:20 2008
 Host: radauth1
      Error: Host name lookup failure
Last change: Mon Sep 29 18:42:20 2008
Host: radacct1
      Error: Host name lookup failure
Host: snmp1
      Addresses:
        4.4.4.1
        4.4.4.2
      Last change: Mon Sep 29 18:45:20 2008
```

```
Host: sys1
      Addresses:
        192.168.68.69
      Last change: Mon Sep 29 18:42:21 2008
```

Release Information

Command introduced in Junos OS Release 9.6.

show system network (JDM)

IN THIS SECTION

- Syntax | 26931
- Description | 26931
- Required Privilege Level | 26931
- Output Fields | 26932
- Sample Output | 26935
- Release Information | 26936

Syntax

```
show system network
```

Description

Display the statistics information for physical interface, JDM interface, and interfaces per guest network function (GNF).

Required Privilege Level

View

Output Fields

Table 1644 on page 26932 describes the output fields for the `show system network` command. Output fields are listed in the approximate order in which they appear.

Table 1644: show system network Output Fields

Field Name	Field Description
Physical Interfaces	
Name	Name of the physical interface.
Index	Interface index.
MTU	The maximum transmission unit.
Hardware-address	The physical interface hardware address.
Rcvd Bytes	The number of bytes received.
Rcvd Packets	The number of packets received.
Rcvd Error	The packets received with error.
Rcvd Drop	The packets dropped during reception.
Trxd Bytes	The number of bytes transmitted.
Trxd Packets	The number of packets transmitted.
Trxd Error	The packets transmitted with error.
Trxd Drop	The packets dropped during transmission.
Flags	Interface flags.

Table 1644: show system network Output Fields *(Continued)*

Field Name	Field Description
Per VNF Interface Statistics	
Name	List of the interfaces for each GNF.
Source	Interface source.
MAC Address	The MAC address of the interface.
Rcvd Bytes	The number of bytes received.
Rcvd Packets	The number of packets received.
Rcvd Error	The packets received with error.
Rcvd Drop	The packets dropped during reception.
Trxd Bytes	The number of bytes transmitted.
Trxd Packets	The number of packets transmitted.
Trxd Error	The packets transmitted with error.
Trxd Drop	The packets dropped during transmission.
JDM Interface Statistics	
Name	Name of the JDM interface.
Index	Interface index.
MTU	The maximum transmission unit.

Table 1644: show system network Output Fields *(Continued)*

Field Name	Field Description
Hardware-address	The JDM interface hardware address.
Rcvd Bytes	The number of bytes received.
Rcvd Packets	The number of packets received.
Rcvd Error	The packets received with error.
Rcvd Drop	The packets dropped during reception.
Trxd Bytes	The number of bytes transmitted.
Trxd Packets	The number of packets transmitted.
Trxd Error	The packets transmitted with error.
Trxd Drop	The packets dropped during transmission.
Flags	Interface flags.
VNF MAC Address Pool	
Start MAC Address	MAC address start value.
Range	Available MAC address range.

Sample Output

show system network

```
user@jdm> show system network

Physical Interfaces
---------------------------------------------------------------------------------------------------
-------------------------------------------
Name     Index MTU     Hardware-address  Rcvd Bytes   Rcvd Packets Rcvd Error Rcvd Drop Trxd
Bytes   Trxd Packets Trxd Error Trxd Drop Flags
-------- ----- ------- ----------------- ------------ ------------ ---------- ---------
------------ ------------ ---------- --------- ------
enp3s0f1 4     1500    00:25:90:b5:75:51 8787662837   51975964     0          538926
40009223     407379       0          0         BMPRU
ens3f1   7     1500    3c:fd:fe:08:87:02 1019880532   16723722     0          11243028
19265494115  31971968     0          0         BMPRU
ens3f0   3     1500    3c:fd:fe:08:87:00 5951717054   81330473     0          11226877
139135292735 124708008    0          0         BMPRU
enp3s0f2 5     1500    00:25:90:b5:75:52 3343179197   40806691     0          461955
3449064446   12191724     0          0         BMRU

Per VNF Interface Statistics
---------------------------------------------------------------------------------------------------
-------------------------------------------------------
Interface                  Source        MAC Address        Rcvd Bytes   Rcvd packets Rcvd Error
Rcvd Drop Trxd bytes   Trxd Packets Trxd Error Trxd Drop
-------------------------- ------------- -----------------  ------------ ------------ ----------
--------- ------------ ------------ ---------- ---------
VNF name: test
macvtap0                   ens3f0        00:f1:60:3d:20:23  91526        1697         0
0         1254382      16782        0          0
macvtap1                   ens3f1        00:f1:60:3d:20:24  561558       3727         0
0         803396       12958        0          0
vnet1                      bridge_jdm_vm 00:f1:60:3d:20:25  2888         32           0
0         2282         25           0          0
macvtap2                   enp3s0f2      00:f1:60:3d:20:26  152563       1080         0
0         4472700      55664        0          0

JDM Interface Statistics
---------------------------------------------------------------------------------------------------
-----------------------------------------
```

```
Name     Index MTU   Hardware-address  Rcvd Bytes   Rcvd Packets Rcvd Error Rcvd Drop Trxd
Bytes   Trxd Packets Trxd Error Trxd Drop Flags
-------- ----- ----- ----------------- ------------ ------------ ---------- ---------
------------ ------------ ---------- --------- ------
bme1     1433  1500  52:54:00:21:20:2e 502730       4506         0          0
477328       2619         0          0         BMRU
jmgmt0   1439  1500  00:f1:60:3d:20:22 4991675      66429        0          2862
100548       891          0          0         BMRU
bme2     1435  1500  52:54:00:88:b5:dd 2930         33           0          0
3466         39           0          0         ABMRU
cb0.4002 2     1500  00:f1:60:3d:20:20 12204921     209269       0          0
3688591023   195579       0          0         ABMRU
cb1.4002 3     1500  00:f1:60:3d:20:21 161850       3026         0          0
204784       3029         0          0         ABMRU

.......................
```

Starting in Junos OS Release 19.2R1, the `show system network` output does not display information related to MAC address pool. For MAC address information use the command `show system mac-addresses`.

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | **15481**

request server authenticate-peer-server | **14901**

show system node-attributes

IN THIS SECTION

- Syntax | **26937**
- Description | **26937**

- Required Privilege Level | 26937
- Output Fields | 26937
- Sample Output | 26938
- Release Information | 26939

Syntax

```
show system node-attributes
```

Description

Show node attribute configuration. Attributes are properties of the node that determine what kind of application gets launched on the node.

Required Privilege Level

view

Output Fields

For a description of the output fields, see Table 1645 on page 26937. Output fields are listed in the approximate order of appearance.

Table 1645: show system node-attributes Output Fields

Field Name	Field Description
`Attribute`	Name of the node attribute. Typical node attribute names are FPC, GlobalIPower, MasterRE, and RE.
`Configured ratio`	Shows the ratio currently configured of active to spare nodes.
`Active Nodes`	Lists the nodes that are active.

Table 1645: show system node-attributes Output Fields *(Continued)*

Field Name	Field Description
Spare Nodes	Lists the nodes that are standby.

Sample Output

show system node-attributes

```
user@host>show system node-attributes
Attribute: ASICS
   Active Nodes: fpc1

Attribute: BT
   Active Nodes: fpc1

Attribute: FABRIC_CONTROL
   Configured num active: 1
   Active Nodes: re0

Attribute: FABRIC_FCHIP_PARALLEL
   Configured num active: 1
   Active Nodes: re0

Attribute: FABRIC_PFE
   Active Nodes: fpc1

Attribute: FAB_TOKEN
   Configured num active: 1
   Active Nodes: re0

Attribute: FPC
   Active Nodes: fpc1

Attribute: GlobalIPOwner
   Active Nodes: re0

Attribute: MasterRE
   Active Nodes: re0
```

```
Attribute: PIC
   Active Nodes: re0

Attribute: RE
   Configured num active: 1
   Active Nodes: re0

Attribute: TIMINGD_FPC
   Active Nodes: fpc1

Attribute: TIMINGD_RE
   Configured num active: 1
   Active Noes: re0
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

show system nodes

IN THIS SECTION

- Syntax | 26940
- Description | 26940
- Options | 26940
- Required Privilege Level | 26940
- Output Fields | 26940
- Sample Output | 26941
- Release Information | 26941

Syntax

```
show system nodes
<node-name>
```

Description

Display information about active nodes on the system.

Options

node-name (Optional) Name of the node. If no specific node is named, the output shows data for all nodes.

Required Privilege Level

view

Output Fields

Table 1646 on page 26940 lists the output fields for the `show system nodes` command.

Table 1646: show system nodes Output Fields

Field Name	Field Description
`Attributes`	Shows the node attributes by name and whether the attribute is active or spare.
`Node`	Name of the node.
`Node ID`	A unique integer that identifies the node in the cluster.
`Node Nonce`	A unique number that identifies the incarnation of the node. The node nonce changes across node reboot.

Table 1646: show system nodes Output Fields *(Continued)*

Field Name	Field Description
Status	Status of the node, online or offline. The status indicates the operational state of the system. For all show chassis commands dealing with FPCs, PICs, or SIBs, the status indicates the operational state of the hardware.

Sample Output

show system nodes

```
user@host> show system nodes

Node: re0
  Node Id    : 598209226576388
  Node Nonce : 4112
  Status     : online
  Attributes : ASICS (Active), MasterRE (Active), PFE4 (Active), PIC (Active), RE (Active),
GlobalIPOwner (Active)
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

node (System) | 11656

show system ntp threshold

IN THIS SECTION

Syntax | 26942

- Description | 26942
- Options | 26942
- Required Privilege Level | 26942
- Output Fields | 26942
- Sample Output | 26943
- Release Information | 26943

Syntax

```
show system ntp threshold
```

Description

Display the current threshold and reject mode configured information.

Options

Required Privilege Level

view

Output Fields

lists the output fields for the Table 1647 on page 26942 `show system ntp threshold` command. Output fields are listed in the approximate order in which they appear.

Table 1647: show system ntp threshold Output Fields

Field Name	Field Description
`NTP threshold`	Assign a threshold value for Network Time Protocol (NTP) adjustment that is outside of the acceptable NTP update and specify whether to accept or reject NTP synchronization when the proposed time from the NTP server exceeds the configured threshold value.

Table 1647: show system ntp threshold Output Fields *(Continued)*

Field Name	Field Description
`Success Criteria`	Verifies the NTP threshold and provide the status of NTP adjustment mode (accept or reject).

Sample Output

show system ntp threshold

```
user@host> show system ntp threshold
NTP threshold:  400 sec
NTP adjustment reject mode is enabled
Success Criteria: verify threshold and reject mode can appear after user configuration.
```

Release Information

Command introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

- set date ntp | **16888**
- ntp threshold | **11731**
- ntp | **11721**
- *NTP Time Synchronization on SRX Series Devices*

show system packet-forwarding-options hw-resource-monitor monitor-profile

IN THIS SECTION

- Syntax | **26944**

- Description | 26944
- Required Privilege Level | 26944
- Output Fields | 26944
- Sample Output | 26945
- Release Information | 26945

Syntax

```
show system packet-forwarding-options hw-resource-monitor monitor-profile monitor-profile-name
```

Description

Display the parameters for the monitor profile that is used for hardware resource threshold monitoring. You can simplify how utilization information is displayed by mapping resource lists to a monitor profile.The same monitor profile can be used for multiple resource lists, but a single resource list cannot be mapped to multiple monitor profiles. A resource list can only be mapped only to a single monitor profile. Use hardware resource threshold monitoring to display utilization information to help assist with capacity utilization and resource management.

Required Privilege Level

view

Output Fields

"show system packet-forwarding-options hw-resource-monitor monitor-profile" on page 26943 lists the output fields for the `show system packet-forwarding-options hw-resource-monitor monitor-profile` command.

Table 1648: show system packet-forwarding-options hw-resource-monitor monitor-profile Output Fields

Field Name	Field Description
Lower Threshold	The lower threshold of the maximum capacity for the resource.

Table 1648: show system packet-forwarding-options hw-resource-monitor monitor-profile Output Fields *(Continued)*

Field Name	Field Description
Upper Threshold	The upper threshold of the maximum capacity for the resource.
Notification Type	The type of notification issued when the threshold is breached. The value is one of the following: `Alarm`, `None`, or `Syslog`.

Sample Output

show system packet-forwarding-options hw-resource-monitor monitor-profile

```
user@host> show system packet-forwarding-options hw-resource-monitor monitor-profile profile1
```

```
lower-threshold 1;
upper-threshold 2;
notification-type alarm;
```

Release Information

Command introduced in Junos OS Release 24.2R1 for EX4100-48MP, EX4400-24MP, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM.

RELATED DOCUMENTATION

No Link Title

polling-interval (Network Management) | **12752**

resource-list (Network Management) | **15610**

show system packet-forwarding-options hw-resource-utilization-info | **26952**

show system packet-forwarding-options hw-resource-monitor resource-list | **26946**

show system packet-forwarding-options hw-resource-monitor utilization-info | **26948**

show system packet-forwarding-options hw-resource-monitor resource-list

IN THIS SECTION

- Syntax | 26946
- Description | 26946
- Required Privilege Level | 26946
- Output Fields | 26946
- Sample Output | 26948
- Release Information | 26948

Syntax

```
show system packet-forwarding-options hw-resource-monitor resource-list resource-list-name
```

Description

Displays the hardware resources in a resource list.

Required Privilege Level

view

Output Fields

"show system packet-forwarding-options hw-resource-monitor resource-list" on page 26946 lists the output fields for the `show system packet-forwarding-options hw-resource monitor resource-list` command.

Table 1649: show system packet-forwarding-options hw-resource-monitor resource-list Output Fields

Field Name	Field Description
resource- names	The resources contained in the resource list. The value is one of the following: • `all-resources`—All applicable hardware resources. • `ecmp-group`—Equal Cost Multi Path (ECMP) group hardware table. • `ecmp-member`—ECMP member hardware table. • `efp`—TCAM: Egress Field Processor hardware table. • `egress-l3-interface`—Egress layer 3 Interface hardware table. • `host-ipv4`—IPv4 host routes hardware table. • `host-ipv6`—IPv6 host routes hardware table. • `ifp`—TCAM: Ingress field processor hardware table. • `l3-next-hop`—Layer 3 next hop hardware table. • `lpm-ipv4`—Longest prefix match IPv4 32 bits wide hardware table. • `lpm-ipv6-128`—Longest prefix match IPv6 128 bits wide hardware table. • `lpm-ipv6-64`—Longest prefix match IPv6 64 bits wide hardware table. • `mac`—Layer 2 entry hardware table. • `mpls-ingress`—Multiprotocol Label Switching (MPLS) ingress hardware table. • `mpls-swap`—MPLS swap hardware table. • `tunnel`—IP/Layer 3 tunnel hardware table. • `vfp`—TCAM: VLAN field processor hardware table. • `vport`—Virtual port hardware table.

Sample Output

show system packet-forwarding-options hw-resource-monitor resource-list

```
user@host> show system packet-forwarding-options hw-resource-monitor resource-list
```

```
resource-names {
    efp;
    ifp;
    vfp;
}
```

Release Information

Command introduced in Junos OS Release 24.2R1 for EX4100-48MP, EX4400-24MP, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM.

RELATED DOCUMENTATION

No Link Title

polling-interval (Network Management) | 12752

resource-list (Network Management) | 15610

show system packet-forwarding-options hw-resource-utilization-info | 26952

show system packet-forwarding-options hw-resource-monitor monitor-profile | 26943

show system packet-forwarding-options hw-resource-monitor utilization-info | 26948

show system packet-forwarding-options hw-resource-monitor utilization-info

IN THIS SECTION

Syntax | 26949

- Description | 26949
- Required Privilege Level | 26949
- Output Fields | 26949
- Sample Output | 26951
- Release Information | 26952

Syntax

```
show system packet-forwarding-options hw-resource-monitor utilization-info
```

Description

Displays hardware resource utilization information for a resource list.

Required Privilege Level

edit

Output Fields

Table 1650 on page 26949 lists the output fields for the `show system packet-forwarding-options hw-resource-monitor utilization-info` command.

Table 1650: show system packet-forwarding-options hw-resource-monitor utilization-info Output Fields

Field Name	Field Description
HW Resource Name	The resource name. The value is one of the following: `ecmp-group`, `ecmp-member`, `efp`, `egress-l3-interface` , `host-ipv4`, `host-ipv6`, `ifp`, `l3-next-hop`, `lpm-ipv4`, `lpm-ipv6-128`, `lpm-ipv6-64`, `mac`, `mpls-ingress`, `mpls-swap`, `tunnel`, `vfp`, `vport` , and `all-resources`.
Max Capacity	The maximum capacity of the resource in number of packets.

Table 1650: show system packet-forwarding-options hw-resource-monitor utilization-info Output Fields *(Continued)*

Field Name	Field Description
Current Utilization	The current hardware resource utilization as an absolute number.
Current Utilization %	The current utilization of the resource as a percentage against the configured threshold.
Lower Threshold %	The lower threshold percentage of the maximum capacity for the resource.
Upper Threshold %	The upper threshold percentage of the maximum capacity for the resource.
Health	The visual health indicator value. This value is updated as `Green`, `Yellow`, or `Red` when a hardware utilization value is compared against the upper and lower threshold boundaries. Whenever there is a change in the resource health, notifications are raised or cleared as required, depending upon the notification type (`alarm` or `syslog`). • A `Green` resource health band indicates that the hardware utilization is safely within threshold boundaries. • A `Yellow` resource health band indicates that the hardware resource utilization has crossed the lower threshold but is still within the upper threshold. • A `Red` resource health band indicates that the hardware resource utilization has crossed both the lower and upper thresholds.
Notification Type	The notification type used to alert you to changes in resource health. Values are `alarm` or `syslog`. The default is `syslog`. • When the resource health changes from `Green` to `Yellow` or vice-versa, a minor alarm is raised or cleared (as applicable). • When the resource health changes from `Yellow` to `Red` or vice-versa, a major alarm is raised or cleared (applicable only if the notification type is alarm). • When the resource health changes from `Green` to `Yellow`, a syslog WARNING is logged (as applicable). • When the resource health changes from `Yellow` to `Red`, a syslog CRITICAL is logged (applicable only if the notification type is syslog).

Sample Output

show system packet-forwarding-options hw-resource-monitor utilization-info

```
user@host> show system packet-forwarding-options hw-resource-monitor utilization-info
```

```
***** HW Resource Monitoring Information  *****

Polling Interval: 1000 milliseconds (1 seconds)

HW Resource Name    | Max Capacity | Current Utilization | Current Utilization % | Lower
Threshold % | Upper Threshold % | Health | Notification Type
--------------------|--------------|---------------------|-----------------------|---------------
----|-------------------|--------|-----------------
ECMP-GROUP          | 4096         | 0                   | 0                     |
50                | 90                | GREEN  | Sylog
ECMP-MEMBER         | 32768        | 0                   | 0                     |
50                | 90                | GREEN  | Sylog
EFP                 | 2048         | 0                   | 0                     |
50                | 90                | GREEN  | Sylog
EGRESS-L3-INTERFACE | 16384        | 7                   | 1                     |
50                | 90                | GREEN  | Sylog
HOST-IPv4           | 147456       | 21                  | 1                     |
50                | 90                | GREEN  | Sylog
HOST-IPv6           | 73728        | 5                   | 1                     |
50                | 90                | GREEN  | Sylog
IFP                 | 18432        | 182                 | 1                     |
50                | 90                | GREEN  | Sylog
L3-NEXT-HOP         | 65536        | 10                  | 1                     |
50                | 90                | GREEN  | Sylog
LPM-IPv4            | 24576        | 8                   | 1                     |
50                | 90                | GREEN  | Sylog
LPM-IPv6-128        | 2048         | 0                   | 0                     |
50                | 90                | GREEN  | Sylog
LPM-IPv6-64         | 12288        | 4                   | 1                     |
50                | 90                | GREEN  | Sylog
MAC                 | 163840       | 2                   | 1                     |
50                | 90                | GREEN  | Sylog
MPLS-INGRESS        | 16384        | 0                   | 0                     |
50                | 90                | GREEN  | Sylog
```

```
MPLS-SWAP          | 16384        | 0                  | 0                    |
50               | 90               | GREEN  | Sylog
TUNNEL             | 4096         | 0                  | 0                    |
50               | 90               | GREEN  | Sylog
VFP                | 1024         | 0                  | 0                    |
50               | 90               | GREEN  | Sylog
VPORT              | 16384        | 0                  | 0                    |
50               | 90               | GREEN  | Sylog
IFP-DYN-GROUP      | 6144         | 364                | 5                    |
50               | 90               | GREEN  | Sylog
```

Release Information

Command introduced in Junos OS Release 24.2R1 for EX4100-48MP, EX4400-24MP, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM.

RELATED DOCUMENTATION

No Link Title

polling-interval (Network Management) | **12752**

resource-list (Network Management) | **15610**

show system packet-forwarding-options hw-resource-monitor resource-list | **26946**

show system packet-forwarding-options hw-resource-monitor monitor-profile | **26943**

show system packet-forwarding-options hw-resource-utilization-info | **26952**

show system packet-forwarding-options hw-resource-utilization-info

IN THIS SECTION

- Syntax | **26953**
- Description | **26953**
- Required Privilege Level | **26953**
- Output Fields | **26953**
- Sample Output | **26954**

Release Information | 26956

Syntax

```
show system packet-forwarding-options hw-resource-utilization-info
```

Description

Show the maximum capacity and current utilization for all applicable resources. This command can be used to display the utilization of hardware resources even if resource lists are not configured.

Required Privilege Level

edit

Output Fields

Table 1651 on page 26953 lists the output fields for the `show system packet-forwarding-options hw-resource-utilization-info` command.

Table 1651: show system packet-forwarding-options hw-resource-utilization-info Output Fields

Field Name	Field Description
HW Resource Name	The resource name. The values can be: `ecmp-group`, `ecmp-member`, `efp`, `egress-l3-interface` , `host-ipv4`, `host-ipv6`, `ifp`, `l3-next-hop`, `lpm-ipv4`, `lpm-ipv6-128`, `lpm-ipv6-64`, `mac`, `mpls-ingress`, `mpls-swap`, `tunnel`, `vfp`, `vport`, and `all-resources`. For QFX5230-64cd and QFX5130-32cd platforms, the values can be `L3 NextHop`, `Host V6`, `Host V4`, `LPM V4`, `LPM V6`, `VFI`, `Virtual Port`, `Tunnel`, `MC Replication NextHop`, `Overlay Egress Interface`, `Underlay Egress Interface`, `Overlay ECMP Group`, `Underlay ECMP Group`, `Overlay ECMP Member`, `Underlay ECMP Member`, `L2 Mac` , `My Mac`, `VFP`, `IFP`, `EFP`, and `System Filters`.
Max Capacity	The maximum capacity of for all applicable resources.
Current Utilization	The current hardware resource utilization in absolute numbers.

Table 1651: show system packet-forwarding-options hw-resource-utilization-info Output Fields *(Continued)*

Field Name	Field Description
Current Utilization %	The current utilization of the resource as a percentage against the configured threshold.

Sample Output

NOTE: The hardware resources displayed in the output and their corresponding values may vary across releases and platforms based on the supported features.

show system packet-forwarding-options hw-resource-utilization-info

```
user@host> show system packet-forwarding-options hw-resource-utilization-info
```

```
Slot 0


***** HW Resource Maximum Capacity and Current Usage *****

HW Resource Name    | Max Capacity | Current Utilization | Current Utilization %
--------------------|--------------|---------------------|----------------------
ecmp-group          | 4096         | 0                   | 0
ecmp-member         | 32768        | 0                   | 0
efp                 | 2048         | 0                   | 0
egress-l3-interface | 16384        | 7                   | 1
host-ipv4           | 147456       | 21                  | 1
host-ipv6           | 73728        | 5                   | 1
ifp                 | 18432        | 182                 | 1
l3-next-hop         | 65536        | 10                  | 1
lpm-ipv4            | 24576        | 8                   | 1
lpm-ipv6-128        | 2048         | 0                   | 0
lpm-ipv6-64         | 12288        | 4                   | 1
mac                 | 163840       | 2                   | 1
mpls-ingress        | 16384        | 0                   | 0
mpls-swap           | 16384        | 0                   | 0
tunnel              | 4096         | 0                   | 0
```

```
vfp                 | 1024         | 0                   | 0
vport               | 16384        | 0                   | 0
```

Sample outout for QFX5130-32cd

```
user@host> show system packet-forwarding-options hw-resource-utilization-info
```

```
--------------------------------|--------------|---------------------|----------------------
HW Resource Name                | Max Capacity | Current Utilization | Current Utilization %
--------------------------------|--------------|---------------------|----------------------
L3 NextHop                       65536          8                     1
Host V4                          32768          0                     1
Host V6                          16384          6                     1
LPM V4                           737280         22                    1
LPM V6                           737280         7                     1
VFI                              16384          4096                  25
Virtual Port                     16384          2                     1
Tunnel                           1024           0                     1
MC Replication NextHop           16383          2                     1
Overlay Egress Interface         16382          6                     1
Underlay Egress Interface        8191           1                     1
Overlay ECMP Group               4095           0                     1
Underlay ECMP Group              4095           0                     1
Overlay ECMP Member              32767          0                     1
Underlay ECMP Member             32767          0                     1
L2 Mac                           32768          0                     1
My Mac                           1022           6                     1
VFP                              4096           0                     1
IFP                              98304          284                   1
EFP                              8192           0                     1
System Filters                   24575          71                    1
```

Sample outout for QFX5230-64cd

```
user@host> show system packet-forwarding-options hw-resource-utilization-info
```

```
--------------------------------|-------------|--------------------|----------------------
HW Resource Name                | Max Capacity | Current Utilization | Current Utilization %
--------------------------------|-------------|--------------------|----------------------
L3 NextHop                       32768          6                     1
Host V4                          16384          0                     1
Host V6                          8192           6                     1
LPM V4                           737280         13                    1
LPM V6                           737280         6                     1
VFI                              4096           2                     1
Virtual Port                     8192           1                     1
Tunnel                           8191           0                     1
MC Replication NextHop           511            2                     1
Egress Interface                 8192           5                     1
Overlay ECMP Group               4096           0                     1
Underlay ECMP Group              4095           0                     1
Overlay ECMP Member              16384          0                     1
Underlay ECMP Member             32768          0                     1
L2 Mac                           8192           0                     1
My Mac                           511            3                     1
My Tunnel Mac                    127            0                     1
VFP                              4096           0                     1
IFP                              12288          496                   4
EFP                              2048           0                     1
System Filters                   1023           62                    6
```

Release Information

Command introduced in Junos OS Release 24.2R1 for EX4100-48MP, EX4400-24MP, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM.

Command introduced in Junos OS Release 24.4R1 for QFX5220-32CD, QFX5220-128C, QFX5130-32CD, QFX5700, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD.

RELATED DOCUMENTATION

No Link Title

polling-interval (Network Management) | 12752

resource-list (Network Management) | 15610

show system packet-forwarding-options hw-resource-monitor utilization-info | 26948

show system processes

IN THIS SECTION

- Syntax | 26957
- Syntax (EX Series Switches and MX Series Routers) | 26958
- Syntax (QFX Series Switches) | 26958
- Syntax (OCX Series) | 26958
- Syntax (TX Matrix Routers) | 26959
- Syntax (TX Matrix Plus Router) | 26959
- Syntax (Junos OS Evolved) | 26959
- Description | 26959
- Options | 26959
- Additional Information | 26966
- Required Privilege Level | 26967
- Output Fields | 26967
- Sample Output | 26972
- Release Information | 26975

Syntax

```
show system processes
<brief | detail | extensive | summary>
<health (pid process-identifer | process-name process-name)>
<providers>
```

```
<resource-limits (brief | detail) process-name>
<wide>
```

Syntax (EX Series Switches and MX Series Routers)

```
show system processes
<all-members>
<brief | detail | extensive | summary>
<health (pid process-identifer | process-name process-name)>
<local>
<member member-id>
<providers>
<resource-limits (brief | detail) process-name>
<wide>
```

Syntax (QFX Series Switches)

```
show system processes
<all-members>
<brief | detail | extensive | summary>
<health (pid process-identifer | process-name process-name)>
host-processes (brief|detail )
<local>
<member member-id>
<providers>
<resource-limits (brief | detail) process-name>
<wide>
```

Syntax (OCX Series)

```
show system processes
<brief | detail | extensive | summary >
<health (pid process-identifer | process-name process-name)>
host-processes (brief|detail )
<providers>
<resource-limits>
<wide>
```

Syntax (TX Matrix Routers)

```
show system processes
<brief | detail | extensive | summary>
<all-chassis| all-lcc | lcc number | scc>
<wide>
```

Syntax (TX Matrix Plus Router)

```
show system processes
<brief | detail | extensive | summary>
<all-chassis| all-lcc | lcc number | sfc number>
<wide>
```

Syntax (Junos OS Evolved)

```
show system processes
<brief | detail | extensive | summary>
<node node-name>
<wide>
```

Description

Display information about software processes that are running on the router or switch and that have controlling terminals.

Options

none	Display standard information about system processes.
brief | detail | extensive | summary	(Optional) Display the specified level of detail.
adaptive-services	(Optional) Display the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

alarm-control	(Optional) Display the process to configure the system alarm.
all-chassis	(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display standard system process information about all the T640 routers (in a routing matrix based on the TX Matrix router) or all the T1600 or T4000 routers (in a routing matrix based on the TX Matrix Plus router) in the chassis.
all-lcc	(TX Matrix routers and TX Matrix Plus router only) (Optional) Display standard system process information for all T640 routers (or line-card chassis) connected to the TX Matrix router. Display standard system process information for all connected T1600 or T4000 LCCs.
all-members	(EX4200 switches, QFX Series Virtual Chassis, and MX Series routers) (Optional) Display standard system process information for all members of the Virtual Chassis configuration.
ancpd-service	Display the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.
application-identification	Display the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.
audit-process	(Optional) Display the RADIUS accounting process.
auto-configuration	Display the Interface Auto-Configuration process.
bootp	Display the process that enables a router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent. DHCP relaying is disabled.
captive-portal-content-delivery	Display the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.
ce-l2tp-service	(Optional) (M10, M10i, M7i, and MX Series routers only) Display the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.
cfm	Display Ethernet Operations, Administration, and Maintenance (OAM) connectivity fault management (CFM) process, which can be used to monitor the physical link between two switches.

chassis-control	(Optional) Display the chassis management process.
class-of-service	(Optional) Display the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.
clksyncd-service	Display the external clock synchronization process, which uses synchronous Ethernet (SyncE).
craft-control	Display the process for the I/O of the craft interface.
database-replication	(EX Series switches and MX Series routers only) (Optional) Display the database replication process.
datapath-trace-service	Display the packet path tracing process.
dhcp-service	(EX Series switches and MX Series routers only) (Optional) Display the Dynamic Host Configuration Protocol process, which enables a DHCP server to allocate network IP addresses and deliver configuration settings to client hosts without user intervention.
diameter-service	(Optional) Display the diameter process.
disk-monitoring	(Optional) Display the disk monitoring process, which checks the health of the hard disk drive on the Routing Engine.
dynamic-flow-capture	(Optional) Display the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.
ecc-error-logging	(Optional) Display the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.
ethernet-connectivity-fault-management	Display the process that provides IEEE 802.1ag OAM connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.
ethernet-link-fault-management	(EX Series switches and MX Series routers only) (Optional) Display the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.
event-processing	(Optional) Display the event process (eventd).
firewall	(Optional) Display the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

general-authentication-service	(EX Series switches and MX Series routers only) (Optional) Display the general authentication process.
health (pid *process-identifer* | process-name *process-name*)	(Optional) Display process health information, either by process id (PID) or by process name.
host-processes	Display process information of processes running on the host system. (On OCX Series only) The following options are available: • brief | detail—(Optional) Display the specified level of detail.
iccp-service	Display the Inter-Chassis Communication Protocol (ICCP) process.
idp-policy	Display the intrusion detection and prevention (IDP) protocol process.
ilmi	Display the Integrated Local Management Interface (ILMI) protocol process, which provides bidirectional exchange of management information between two ATM interfaces across a physical connection.
inet-process	Display the IP multicast family process.
init	Display the process that initializes the USB modem.
interface-control	(Optional) Display the interface process, which controls the router's or switch's physical interface devices and logical interfaces.
kernel-replication	(Optional) Display the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.
l2-learning	(Optional) Display the Layer 2 address flooding and learning process.
l2cpd-service	Display the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.
lacp	(Optional) Display the Link Aggregation Control Protocol (LACP)process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.
lcc *number*	(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display standard system process information for a specific T640 router

that is connected to the TX Matrix router. On a TX Matrix Plus router, display standard system process information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(EX4200 switches, QFX Series Virtual Chassis, and MX Series routers) (Optional) Display standard system process information for the local Virtual Chassis member.

local-policy-decision-function

Display the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

logical-system-mux

Display the logical router multiplexer process (lrmuxd), which manages the multiple instances of the routing protocols process (rpd) on a machine running logical routers.

mac-validation

Display the MAC validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member *member-id*

(EX4200 switches, QFX Series Virtual Chassis, and MX Series routers) (Optional) Display standard system process information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

mib-process

(Optional) Display the MIB II process, which provides the router's MIB II agent.

mobile-ip

(Optional) Display the Mobile IP process, which configures Junos OS Mobile IP features.

mountd-service	(EX Series switches and MX Series routers only) (Optional) Display the service for NFS mounts requests.
mpls-traceroute	(Optional) Display the MPLS Periodic Traceroute process.
mspd	(Optional) Display the Multiservice process.
multicast-snooping	(EX Series switches and MX Series routers only) (Optional) Display the multicast snooping process, which makes Layer 2 devices such as VLAN switches aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.
named-service	(Optional) Display the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.
neighbor-liveness	Display the process, which specifies the maximum length of time that the router waits for its neighbor to re-establish an LDP session.
nfsd-service	(Optional) Display the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.
ntp	Display the Network Time Protocol (NTP) process, which provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network.
packet-triggered-subscribers	Display the packet-triggered subcribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.
peer-selection-service	(Optional) Display the Peer Selection Service process.
periodic-packet-services	Display the Periodic packet management process, which is responsible for processing a variety of time-sensitive periodic tasks so that other processes can more optimally direct their resources.
pfe	Display the Packet Forwarding Engine management process.
pgcp-service	(Optional) Display the pgcpd service process running on the Routing Engine.
pgm	Display the Pragmatic General Multicast (PGM) protocol process, which enables a reliable transport layer for multicast applications.
pic-services-logging	(Optional) Display the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

ppp	(Optional) Display the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.
ppp-service	Display the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.
pppoe	(Optional) Display the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.
process-monitor	Display the process health monitor process (pmond).
providers	(Optional) Display provider processes.
redundancy-interface-process	(Optional) Display the ASP redundancy process.
remote-operations	(Optional) Display the remote operations process, which provides the ping and traceroute MIBs.
resource-cleanup	Display the resource cleanup process.
resource-limits (brief | detail) *process-name*	(Optional) Display process resource limits.
routing	(Optional) Display the routing protocol process.
sampling	(Optional) Display the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.
sbc-configuration-process	Display the session border controller (SBC) process of the border signaling gateway (BSG).
scc	(TX Matrix routers only) (Optional) Display standard system process information for the TX Matrix router (or switch-card chassis).
sdk-service	Display the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.
secure-neighbor-discovery	(EX Series switches and MX Series routers only) (Optional) Display the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

send	(Optional) Display the Secure Neighbor Discovery Protocol (SEND) process, which provides support for protecting Neighbor Discovery Protocol (NDP) messages.
service-deployment	(Optional) Display the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.
sfc *number*	(TX Matrix Plus routers only) (Optional) Display system process information for the TX Matrix Plus router. Replace *number* with 0.
snmp	Display the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP primary agent.
sonet-aps	Display the SONET Automatic Protection Switching (APS) process, which monitors any SONET interface that participates in APS.
static-subscribers	(Optional) Display the Static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.
tunnel-oamd	(Optional) Display the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.
vrrp	(EX Series switches and MX Series routers only) (Optional) Display the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.
watchdog	Display the watchdog timer process, which enables the watchdog timer when Junos OS encounters a problem.
wide	(Optional) Display process information that might be wider than 80 columns.
node *node-name*	Specify a name if you want to view the system process details for that node. Example: `re0`.

Additional Information

By default, when you issue the `show system processes` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of

the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Output Fields

The following table describes the output fields for the `show system processes` command. Output fields are listed in the approximate order in which they appear.

Table 1652: show system processes Output Fields

Field Name	Field Description	Level of Output
`last pid`	Last process identifier assigned to the process.	`brief extensive summary`
`load averages`	Three load averages followed by the current time.	`brief extensive summary`
`processes`	Number of existing processes and the number of processes in each state (`sleeping`, `running`, `starting`, `zombies`, and `stopped`).	`brief extensive summary`
`CPU`	(For systems running Junos OS with upgraded FreeBSD only) Breakdown of the percent usage on a per-CPU basis into the following categories: `% user`, `% nice`, `% system`, `% interrupt`, `% idle`. **NOTE**: This field shows up in the second frame of output. To see which platforms run Junos OS with upgraded FreeBSD, see Release Information for Junos OS with Upgraded FreeBSD. **NOTE**: On ACX Series routers running Junos OS Evolved, the `show system process` command might report CPU utilization spikes greater than 100%. This kind of CPU utilization is normal behavior, and no user action is required. The CPU utilization spikes represent the sum of individual processor threads and not of the entire system CPU capacity.	`extensive`

Table 1652: show system processes Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Mem`	Information about physical and virtual memory allocation.	`brief extensive summary`
`Active`	Memory allocated and actively used by the program. When the system is under memory pressure, the pageout process reuses memory from the free, cache, inact and, if necessary, active pages. When the pageout process runs, it scans memory to see which pages are good candidates to be unmapped and freed up. Thus, the distinction between `Active` and `Inact` memory is only used by the pageout process to determine which pool of pages to free first at the time of a memory shortage. The pageout process first scans the `Inact` list, and checks whether the pages on this list have been accessed since the time they have been listed here. The pages that have been accessed are moved from the `Inact` list to the `Active` list. On the other hand, pages that have not been accessed become prime candidates to be freed by the pageout process. If the pageout process cannot produce enough free pages from the `Inact` list, pages from the `Active` list get freed up. Because the pageout process runs only when the system is under memory pressure, the pages on the `Inact` list remain untouched – even if they have not been accessed recently – when the amount of `Free` memory is adequate.	`brief extensive summary`
`Inact`	Memory allocated but not recently used or memory freed by the programs. Inactive memory remains mapped in the address space of one or more processes and, therefore, counts toward the RSS value of those processes. Any amount of memory freed by the routing protocol process might still be considered part of the `RES` value. Generally, the kernel delays the migrating of memory out of the `Inact` queue into the `Cache` or `Free` list unless there is a memory shortage.	`brief extensive summary`
`Wired`	Memory that is not eligible to be swapped, usually used for in-kernel memory structures and/or memory physically locked by a process.	`brief extensive summary`

Table 1652: show system processes Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Cache`	Memory that is not associated with any program and does not need to be swapped before being reused.	`brief extensive summary`
`Buf`	Size of memory buffer used to hold data recently called from the disk.	`brief extensive summary`
`Free`	Memory that is not associated with any programs. Memory freed by a process can become `Inactive`, `Cache`, or `Free`, depending on the method used by the process to free the memory.	`brief extensive summary`
`Swap`	Information about physical and virtual memory allocation. **NOTE**: Memory can remain swapped out indefinitely if it is not accessed again. Therefore, the show system process extensive command shows that memory is swapped to disk even though there is plenty of free memory, and such a situation is not unusual.	`brief extensive summary`
`PID`	Process identifier.	`detail extensive summary`
`TT`	Control terminal name.	none `detail`

Table 1652: show system processes Output Fields *(Continued)*

Field Name	Field Description	Level of Output
STAT	Symbolic process state. The state is given by a sequence of letters. The first letter indicates the run state of the process: • D—In disk or other short-term, uninterruptible wait • I—Idle (sleeping longer than about 20 seconds) • R—Runnable • S—Sleeping for less than 20 seconds • T—Stopped • Z—Dead (zombie) • + —The process is in the foreground process group of its control terminal. • <—The process has raised CPU scheduling priority. • >—The process has specified a soft limit on memory requirements and is currently exceeding that limit; such a process is not swapped. • A—The process requested random page replacement. • E—The process is trying to exit. • L—The process has pages locked in core. • N—The process has reduced CPU scheduling priority. • S—The process requested first-in, first-out (FIFO) page replacement. • s—The process is a session leader. • V—The process is temporarily suspended. • W—The process is swapped out. • X—The process is being traced or debugged.	none detail
UID	User identifier.	detail

Table 1652: show system processes Output Fields *(Continued)*

Field Name	Field Description	Level of Output
USERNAME	Process owner.	extensive summary
PPID	Parent process identifier.	detail
CPU	**(D)** Short-term CPU usage. **(E and S)** Raw (unweighted) CPU usage. The value of this field is used to sort the processes in the output.	detail extensive summary
RSS	Resident set size.	detail
WCHAN	Symbolic name of the wait channel.	detail
STARTED	Local time when the process started running.	detail
PRI	Current priority of the process. A lower number indicates a higher priority.	detail extensive summary
NI or NICE	UNIX "niceness" value. A lower number indicates a higher priority.	detail extensive summary
SIZE	Total size of the process (text, data, and stack), in kilobytes.	extensive summary
RES	Current amount of program resident memory, in kilobytes. This is also known as RSS or Resident Set Size. The RES value includes shared library pages used by the process. Any amount of memory freed by the process might still be considered part of the RES value. Generally, the kernel delays the migrating of memory out of the Inact queue into the Cache or Free list unless there is a memory shortage. This can lead to large discrepancies between the values reported by the routing protocol process and the kernel, even after the routing protocol process has freed a large amount of memory.	extensive summary

Table 1652: show system processes Output Fields *(Continued)*

Field Name	Field Description	Level of Output
STATE	Current state of the process (for example, sleep, wait, run, idle, zombie, or stop).	extensive summary
C	CPU number. **NOTE**: There is no such column in output from Junos OS Evolved. To see the CPU number, issue the show system processes wide | detail command and look at the PSR column.	extensive summary
TIME	**(S)** Number of system and user CPU seconds that the process has used. **(None, D, and E)** Total amount of time that the command has been running.	detail extensive summary
WCPU	Weighted CPU usage.	extensive summary
COMMAND	Command that is currently running. (MX Series routers only) When you display the software processes for an MX Series Virtual Chassis, the show system processes command does not display information about the relayd process.	detail extensive summary
THR	Number of threads in the process	extensive

Sample Output

show system processes

```
user@host> show system processes
PID  TT  STAT      TIME COMMAND
    0  ??  DLs    0:00.70  (swapper)
    1  ??  Is     0:00.35 /sbin/init --
    2  ??  DL     0:00.00  (pagedaemon)
```

```
    3  ??  DL     0:00.00  (vmdaemon)
    4  ??  DL     0:42.37  (update)
    5  ??  DL     0:00.00  (if_jnx)
   80  ??  Ss     0:14.66 syslogd -s
   96  ??  Is     0:00.01 portmap
  128  ??  Is     0:02.70 cron
  173  ??  Is     0:02.24 /usr/local/sbin/sshd (sshd1)
  189  ??  S      0:03.80 /sbin/watchdog -t180
  190  ??  I      0:00.03 /usr/sbin/tnetd -N
  191  ??  S      2:24.76 /sbin/ifd -N
  192  ??  S<     0:55.44 /usr/sbin/xntpd -N
  195  ??  S      0:53.11 /usr/sbin/snmpd -N
  196  ??  S      1:15.73 /usr/sbin/mib2d -N
  198  ??  I      0:00.75 /usr/sbin/inetd -N
 2677  ??  I      0:00.01 /usr/sbin/mgd -N
 2712  ??  Ss     0:00.24 rlogind
 2735  ??  R      0:00.00 /bin/ps -ax
 1985  p0- S      0:07.41 ./rpd -N
 2713  p0  Is     0:00.24 -tcsh (tcsh)
 2726  p0  S+     0:00.07 cli
```

show system processes brief

```
user@host> show system processes brief
last pid:   543;  load averages:  0.00,  0.00,  0.00    18:29:47
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free
```

show system processes detail

```
user@host> show system processes detail
  PID   UID  PPID CPU PRI  NI    RSS WCHAN    STARTED TT  STAT        TIME COMMAND
   11     0     0   0 155   0     64 -        Wed16    -  RL    6411:48.16 [idle]
17967     0 17936   0  24   0  96976 select   Wed16    -  S      138:17.79 /usr/sbin/chassisd -N
    0     0     0   0 -16   0    384 swapin   Wed16    -  DLs      0:22.33 [kernel]
    1     0     0   0  20   0    600 wait     Wed16    -  ILs      0:00.11 /sbin/init --
    2     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:30.05 [jfe_job_0_0]
    3     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_1_0]
```

```
    4     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_1_1]
    5     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_1_2]
    6     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_1_3]
    7     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_2_0]
    8     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_3_0]
    9     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_3_1]
   10     0     0   0 -16   0     16 audit_wo Wed16    -  DL       0:00.00 [audit]
   12     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_3_2]
   13     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_3_3]
   14     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_4_0]
   15     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_4_1]
   16     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.00 [jfe_job_4_2]
   17     0     0   0 -16   0     16 jfe_job_ Wed16    -  DL       0:00.92 [jfe_job_4_3]
   18     0     0   0 -64   0    656 -        Wed16    -  WL       4:47.14 [intr]
   19     0     0   0  -8   0     48 -        Wed16    -  DL       0:02.31 [geom]
   20     0     0   0 -16   0     16 crypto_w Wed16    -  DL       0:00.00 [crypto]
   21     0     0   0 -16   0     16 crypto_r Wed16    -  DL       0:00.00 [crypto returns]
   22     0     0   0 -16   0     32 -        Wed16    -  DL       0:00.01 [cam]
```

show system processes extensive

```
user@host> show system processes extensive
Mem: 241M Active, 99M Inact, 78M Wired, 325M Cache, 69M Buf, 1251M Free
Swap: 2048M Total, 2048M Free


  PID USERNAME  THR PRI NICE   SIZE    RES STATE    TIME   WCPU COMMAND
   11 root        1 171   52     0K    12K RUN    807.5H 98.73% idle
   13 root        1 -20 -139     0K    12K WAIT    36:17  0.00% swi7: clock sio
 1499 root        1  96    0  7212K  3040K select  34:01  0.00% license-check
 1621 root        1  96    0 20968K 11216K select  20:25  0.00% mib2d
 1465 root        2   8  -88   115M 11748K nanslp  14:32  0.00% chassisd
 1478 root        1  96    0  6336K  3816K select  11:28  0.00% ppmd
   20 root        1 -68 -187     0K    12K WAIT    10:28  0.00% irq10: em0 em1+++*
 1490 root        1  96    0 11792K  4336K select   9:44  0.00% shm-rtsdbd
 1618 root        1  96    0 39584K  7464K select   8:47  0.00% pfed
 1622 root        1  96    0 15268K 10988K select   6:16  0.00% snmpd
 1466 root        1  96    0  7408K  2896K select   5:44  0.00% alarmd
    7 root        1 -16    0     0K    12K client   5:09  0.00% ifstate notify
 1480 root        1  96    0  5388K  2660K select   4:29  0.00% ksyncd
   12 root        1 -40 -159     0K    12K WAIT     4:15  0.00% swi2: netisr 0
```

```
 1462 root        1  96    0  1836K  1240K select   3:57  0.00% bslockd
   55 root        1 -16    0     0K    12K -        3:44  0.00% schedcpu
 1392 root        1  16    0     0K    12K bcmsem   3:37  0.00% bcmLINK.0
   47 root        1 -16    0     0K    12K psleep   3:25  0.00% vmkmemdaemon
   36 root        1  20    0     0K    12K syncer   2:46  0.00% syncer
 1484 root        1  96    0  7484K  3428K select   2:38  0.00% clksyncd
 1616 root        1  96    0  4848K  2848K select   2:18  0.00% irsd
 1487 root        1  96    0 32800K  6992K select   2:10  0.00% smid
 1623 root        1  96    0 34616K  5464K select   2:01  0.00% dcd
   15 root        1 -16    0     0K    12K -        1:59  0.00% yarrow
   49 root        1 -16    0     0K    12K .        1:51  0.00% ddostasks
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `sfc` introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Enhanced output regarding per CPU usage introduced in Junos OS Release 16.1R3 for Junos OS with upgraded FreeBSD.

RELATED DOCUMENTATION

> *Junos OS Routing Processes*
>
> Routing Matrix with a TX Matrix Plus Router Solutions Page

show system queues

IN THIS SECTION

- Syntax | **26976**
- Syntax (TX Matrix Router) | **26976**
- Syntax (TX Matrix Plus Router) | **26976**
- Syntax (MX Series Router) | **26976**
- Description | **26976**
- Options | **26977**

- Additional Information | 26978
- Required Privilege Level | 26978
- Output Fields | 26978
- Sample Output | 26980
- Release Information | 26980

Syntax

```
show system queues
```

Syntax (TX Matrix Router)

```
show system queues
<all-chassis| all-lcc | lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
show system queues
<all-chassis | all-lcc | lcc number | sfc number>
```

Syntax (MX Series Router)

```
show system queues
<all-members>
<local>
<member member-id>
```

Description

Display queue statistics.

Options

all-chassis (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system queue statistics for all the T640 routers in the chassis that are connected to the TX Matrix router. On a TX Matrix Plus router, display system queue statistics for all the T1600 or T4000 routers in the chassis that are connected the TX Matrix Plus router.

all-lcc (TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system queue statistics for all LCC chassis attached to the TX Matrix or TX Matrix Plus router.

all-members (MX Series routers only) (Optional) Display system queue statistics for all members of the Virtual Chassis configuration.

lcc *number* (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system queue statistics for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system queue statistics for a specific connected router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (MX Series routers only) (Optional) Display system queue statistics for the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display system queue statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

scc (TX Matrix routers only) (Optional) Display queue statistics for the TX Matrix router.

sfc *number* (TX Matrix Plus routers only) (Optional) Display system queue statistics for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system queues` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

maintenance

Output Fields

Table 1653 on page 26979 lists the output fields for the `show system queues` command. Output fields are listed in the approximate order in which they appear.

Table 1653: show system queues Output Fields

Field Name	Field Description
`Output interface`	Interface on the device on which the queue exists: • `fxp0`—Management Ethernet interface. • `fxp1`—Internal Ethernet interface. • `ipip, lsi, tap, mt, mtun, pimd, and pime`—Internally generated interface and not configurable. • `dsc`—Discard interface. • `em`—Management and internal Ethernet interfaces. • `gre`—Internally generated interface that is configurable only as the control channel for Generalized MPLS (GMPLS). • `ge`—Gigabit Ethernet interface. • `xe`—10-Gigabit Ethernet interface. • `lo`—Loopback interface; the Junos OS automatically configures one loopback interface (lo0). • `lsq`—Link services IQ interface. • `lt`—Logical tunnel interface. • `gr, ip, sp`—Services interfaces. • `irb`—integrated routing and bridging interface. • `vtep`—Virtual Tunnel End Point (VTEP). • `ppd and ppe`—Interfaces used to enable a cluster to act as a rendezvous point (RP) or first hop router in the multicast domain.
`bytes`	Number of bytes in the queue.
`max`	Maximum number of bytes allowed in the queue.
`packets`	Number of packets in the queue.

Table 1653: show system queues Output Fields *(Continued)*

Field Name	Field Description
max	Maximum number of packets allowed in the queue.
drops	Number of packets dropped from the queue.

Sample Output

show system queues

```
user@host> show system queues
output interface            bytes          max  packets      max    drops
fxp0                            0      1250000        0     4166        6
fxp1                            0      1250000        0     4166       19
lsi                             0        12500        0       41        0
dsc                             0            0        0        0        0
```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system random-mac-prefix (JDM)

IN THIS SECTION

- Syntax | 26981
- Description | 26981
- Required Privilege Level | 26981
- Sample Output | 26981
- Release Information | 26982

Syntax

```
show system random-mac-prefix
```

Description

Display the random MAC prefix generated by JDM. This pseudo-random MAC prefix forms part of a MAC address associated with an unlicensed guest network function (GNF).

> **NOTE**: Before creating a GNF, you must ensure that you have configured commit synchronization as part of JDM interfaces configuration so that the random MAC prefixes generated by the JDM instances are in sync.

Required Privilege Level

View

Sample Output

show system random-mac-prefix

```
user@jdm> show system random-mac-prefix
Mac-Prefix : 0x1e
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | **15481**

request server authenticate-peer-server | **14901**

show system reboot

IN THIS SECTION

- Syntax | **26982**
- Syntax (EX Series and MX Series) | **26983**
- Syntax (TX Matrix Router) | **26983**
- Syntax (TX Matrix Plus Router) | **26983**
- Syntax (QFX Series and OCX Series) | **26983**
- Syntax (Junos OS Evolved) | **26983**
- Description | **26984**
- Options | **26984**
- Additional Information | **26985**
- Required Privilege Level | **26986**
- Sample Output | **26986**
- Release Information | **26987**

Syntax

```
show system reboot
<both-routing-engines>
```

Syntax (EX Series and MX Series)

```
show system reboot
<all-members>
<both-routing-engines>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system reboot
<all-chassis | all-lcc | lcc number / scc>
<both-routing-engines>
```

Syntax (TX Matrix Plus Router)

```
show system reboot
<all-chassis| all-lcc | lcc number /  sfc number>
<both-routing-engines>
```

Syntax (QFX Series and OCX Series)

```
show system reboot
<both-routing-engines>
<infrastructure name>
<interconnect-device name>
<node-device name>
```

Syntax (Junos OS Evolved)

```
show system reboot
```

Description

Display pending system reboots or halts.

Options

none

Display pending reboots or halts on the active Routing Engine.

For Junos OS Evolved, the `show system reboot` command is applicable to all nodes (Routing Engines and FPCs). There is no system reboot command for a specific Routing Engine. Hence, the `show system reboot` command shows the pending reboot for the system and not for a specific Routing Engine.

all-chassis

(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display halt or reboot request information for all the T640 routers in the chassis that are connected to the TX Matrix router. On a TX Matrix router, display halt or reboot request information for all the T1600 or T4000 routers in the chassis that are connected to the TX Matrix Plus router.

all-members

(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for all members of the Virtual Chassis configuration.

all-lcc

(TX Matrix routers and TX Matrix Plus router only) (Optional) On a TX Matrix router, display system halt or reboot request information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display halt or reboot request information for all connected T1600 or T4000 LCCs.

both-routing-engines

(Systems with multiple Routing Engines) (Optional) Display halt or reboot request information on both Routing Engines.

> **NOTE**: Junos OS Evolved does not support the both-routing-engines option.

infrastructure *name*

(QFabric systems only) (Optional) Display reboot request information on the fabric manager Routing Engines and fabric control Routing Engines.

interconnect-device *name*

(QFabric systems only) (Optional) Display reboot request information on the Interconnect device.

lcc *number*

(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display halt or reboot request information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display halt or reboot

request information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for the local Virtual Chassis member.

member *member-id* (EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node-group *name* (QFabric systems only) (Optional) Display reboot request information on the Node group.

scc (TX Matrix router only) (Optional) Display halt or reboot request information for the TX Matrix router (or switch-card chassis).

sfc (TX Matrix Plus router only) (Optional) Display halt or reboot request information for the TX Matrix Plus router.

Additional Information

By default, when you issue the `show system reboot` command on a TX Matrix or TX Matrix Plus primary Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) primary Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) backup Routing Engines that are connected to it.

For Junos OS Evolved, the `show system reboot` command is applicable to all nodes (Routing Engines and FPCs). There is no system reboot command for a specific Routing Engine. Hence, the `show system reboot` command shows the pending reboot for the system and not for a specific Routing Engine.

Required Privilege Level

maintenance

Sample Output

show system reboot

```
user@host> show system reboot
reboot requested by root at Wed Feb 10 17:40:46 1999
[process id 17885]
```

show system reboot all-lcc (TX Matrix Router)

```
user@host> show system reboot all-lcc
lcc0-re0:
--------------------------------------------------------------------------
No shutdown/reboot scheduled.

lcc2-re0:
--------------------------------------------------------------------------
No shutdown/reboot scheduled.
```

show system reboot sfc (TX Matrix Plus Router)

```
user@host> show system sfc 0
No shutdown/reboot scheduled.
```

show system reboot (QFX3500 Switch)

```
user@switch> show system reboot
No shutdown/reboot scheduled.
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system resource-monitor

IN THIS SECTION

- Syntax | 26987
- Description | 26987
- Options | 26988
- Additional Information | 26988
- Required Privilege Level | 26988
- Output Fields | 26988
- Sample Output | 26989
- Release Information | 26992

Syntax

```
show system resource-monitor ( fpc | summary )
<slot slot-number>
```

Description

Display the utilization of memory resources on the Packet Forwarding Engines for all FPCs or a specific FPC. The filter memory denotes the filter counter memory used for firewall filter counters. The asterisk

(*) displayed next to each of the memory regions denotes the ones for which the configured threshold is being currently exceeded.

Options

slot *slot-number* Display the Junos OS utilization information of memory resources for the specified slot number in which the FPC (or MPC) is installed.

> **NOTE**: On the MX80 router, the MPCs are built into the chassis and are represented in the CLI as either `FPC 0` or `FPC 1`. This command supports only FPC 1 for the MX80 router.

Additional Information

The filter memory denotes the filter counter memory used for firewall filter counters. From the Ukern perspective, MPC5E contains only one Packet Forwarding Engine instance. The `show chassis fabric plane` command output displays the state of fabric plane connections to the Packet Forwarding Engine. Because two Packet Forwarding Engines exist, you notice PFE-0 and PFE-1 in the output.

Required Privilege Level

view

Output Fields

Table 1654 on page 26988 lists the output fields for the `show system resource-monitor fpc` command. Output fields are listed in the approximate order in which they appear.

Table 1654: show system resource-monitor fpc Output Fields

Field Name	Field Description
`Free Heap Memory Watermark`	Configured watermark value for the percentage of free memory space used for ukernel or heap memory to be monitored

Table 1654: show system resource-monitor fpc Output Fields *(Continued)*

Field Name	Field Description
`Free FW Memory Watermark`	Configured watermark value for the percentage of free memory space used for firewall or filter memory to be monitored
`Free NH Memory Watermark`	Configured watermark value for the percentage of free memory space used for next-hop memory to be monitored
`* - watermark reached`	An asterisk (*) displayed beside any of the memory regions denotes the memory types for which the configured threshold is being currently exceeded.
`Slot #`	Slot number in which the line card is installed
`PFE #`	Number or identifier of the Packet Forwarding Engine in the specified line card slot
`Heap % free`	Percentage of free space associated with heap or ukernel memory
`Encap mem % free`	Percentage of free space associated with encapsulation memory
`NH mem % free`	Percentage of free space associated with next-hop memory
`Filter / FW mem % free`	Percentage of free space associated with firewall or filter memory

Sample Output

show system resource-monitor fpc (All Slots)

```
show system resource-monitor fpc
FPC Resource Usage Summary
```

```
Free Heap Mem Watermark         : 20  %
Free NH Mem Watermark           : 20  %
Free Filter Mem Watermark       : 20  %

* - Watermark reached


Slot #         % Heap Free        RTT      Average RTT
     0             89                103     102(30)

                 PFE #      % ENCAP mem Free      % NH mem Free      % FW mem Free
                    0              NA                 78                 99

Slot #         % Heap Free        RTT      Average RTT
     2             88                103     103(30)

                 PFE #      % ENCAP mem Free      % NH mem Free      % FW mem Free
                    0              NA                 80                 99
                    1              NA                 80                 99

Slot #         % Heap Free        RTT      Average RTT
     3             91                  0      --(--)

                 PFE #      % ENCAP mem Free      % NH mem Free      % FW mem Free
                    0              99                 82                 72
                    1              99                 82                 72
```

show system resource-monitor fpc (Specific Slot)

```
show system resource-monitor fpc slot 2
FPC Resource Usage Summary

Free Heap Mem Watermark         : 20  %
Free NH Mem Watermark           : 20  %
Free Filter Mem Watermark       : 20  %

* - Watermark reached


Slot #         % Heap Free        RTT      Average RTT
```

```
     2             88                103     103(30)

                 PFE #      % ENCAP mem Free      % NH mem Free      % FW mem Free
                   0              NA                  80                 99
                   1              NA                  80                 99
```

show system resource-monitor summary

```
user@host> show system resource-monitor summary
 Resource Usage Summary

 Throttle                       : Enabled
 Load Throttle                  : Enabled
 Heap Mem Threshold             : 70
 IFL Counter Threshold          : 95
 Round Trip Delay Threshold(ms) : 1000
 Filter Counter Threshold       : 100
 Expansion Threshold            : 95
 CoS Queue Threshold            : 0
 MFS threshold                  : 70          Used : 0

 Slot +ACM- 0
      Client allowed                   : Yes
      Service allowed                  : Yes
      Heap memory used                 : 294197664       In  : 8
      Average Round-trip Delay(ms)     : 104  (30  )
      Round-trip Delay(ms) : 101

      MAX session rate allowed(+ACU-)  : 100
      Client denied                    : 6
      Service Denied                   : 0
      Performance Denial Client        : 6
      Performance Denial Service       : 0
      IFL Denied                       : 0

       Filter counter memory      IFL counter memory   Expansion memory
       PFE    +ACM-         used         +AHw-       used  +AHw-  +ACU-  used  +AHw-   +ACU-
       0       23301216     29           420320      1       8337184      4

      CoS Queue Utilization
      PFE +ACM-   Scheduler Block +ACM-        Used      +ACU-
```

```
        0                   0                   8140      49
        0                   1                   8140      49
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

- show system resource-monitor summary
- show system resource-monitor ifd-cos-queue-mapping fpc | 26992
- show system resource-monitor subscribers-limit | 26995

show system resource-monitor ifd-cos-queue-mapping fpc

IN THIS SECTION

- Syntax | 26992
- Description | 26993
- Options | 26993
- Required Privilege Level | 26993
- Output Fields | 26993
- Sample Output | 26994
- Platform-Specific Behavior Resource-Monitor | 26994
- Release Information | 26994

Syntax

```
show system resource-monitor ifd-cos-queue-mapping fpc slot-number
```

Description

Display the scheduler usage of each Packet Forwarding Engine and the mapping of scheduler blocks to CoS-capable physical interface (IFDs).

Options

slot-number Display the Junos OS utilization information of memory resources for the specified slot number in which the FPC is installed.

Required Privilege Level

view

Output Fields

Table 1655 on page 26993 lists the output fields for the `show system resource-monitor ifd-cos-queue-mapping fpc` command. Output fields are listed in the approximate order in which they appear.

Table 1655: show system resource-monitor ifd-cos-queue-mapping fpc Output Fields

Field	Field Description
`PFE #`	Number or identifier of the Packet Forwarding Engine in the specified line-card slot.
`Scheduler Block #`	Scheduler block identifier. For each Packet Forwarding Engine there are multiple schedulers.
`used`	Scheduler queues used.
`%`	Percentage of queue resource used.
`Devices`	List of physical interfaces that map the scheduler.

Sample Output

show system resource-monitor ifd-cos-queue-mapping fpc

```
user@host> show system resource-monitor ifd-cos-queue-mapping fpc 0
Feb 07 23:50:59
    PFE #   Scheduler Block #       Used    %             Devices
        0                   0       5972   36     xe-0/0/1  xe-0/0/0
        0                   1       5972   36     xe-0/1/1  xe-0/1/0
```

Platform-Specific Behavior Resource-Monitor

Table 1656: Platform-Specific Behavior

Platform	Behavior
MPC5E Line Card with flexible-queuing-mode enabled	For MPC5E line card with flexible-queuing-mode enabled, queue resources are shared between scheduler block 0 & 1. Resource monitor CLI output displays an equal distribution of the total available and used queues between scheduler blocks. This correctly represents the queue availability to the routing engine.

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

show system resource-monitor fpc

show system resource-monitor subscribers-limit | 26995

show system resource-monitor subscribers-limit

IN THIS SECTION

- Syntax | 26995
- Description | 26995
- Options | 26995
- Required Privilege Level | 26996
- Output Fields | 26996
- Sample Output | 26997
- Release Information | 27001

Syntax

```
show system resource-monitor subscribers-limit
<chassis>
<fpc slot-number>
<pic number>
<port number>
<extensive>
```

Description

Display information about subscriber limits for the specified hardware element, chassis, MPC, MIC, or port by client type. Shows the configured limit, the number of subscribers of the type currently logged in, and the number of subscribers that have been denied login because the limit has been reached. Use the `extensive` option to display information for the specified element and all subordinate elements that have a configured subscriber limit.

Options

extensive (Optional) Display information for the specified hardware element and all subordinate elements that have a configured subscriber limit.

chassis (Optional) Subscriber limit statistics for the chassis.

fpc *slot-number* (Optional) Subscriber limit statistics for MPC in the specified slot.

pic *number* (Optional) Subscriber limit statistics for the specified MIC.

port *number* (Optional) Subscriber limit statistics for the specified port.

Required Privilege Level

view

Output Fields

Table 1657 on page 26996 lists the output fields for the `show system resource-monitor subscribers-limit` command. Output fields are listed in the approximate order in which they appear.

Table 1657: show system resource-monitor subscribers-limit Output Fields

Field Name	Field Description	Level of Output
`fpc`, `pic`, `port`	Hardware element on which a maximum subscriber limit is configured. `fpc` corresponds to MPC. `pic` corresponds to MIC.	All levels
`Client-type`	Type of client for which a maximum subscriber limit is configured on the specified hardware element: `ANY`, `DHCP`, `L2TP`, or `PPPoE`.	All levels
`Configured limit`	Maximum number of subscribers that can be logged in for the client type.	All levels
`Current count`	Current number of subscribers that can log in for the client type.	All levels
`Denied count`	Number of subscribers for the client type that have been denied login because the maximum subscriber limit has been reached.	All levels

Sample Output

show system resource-monitor subscribers-limit (Chassis)

```
user@host> show system resource-monitor subscribers-limit chassis
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

Client-type : any
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0
```

show system resource-monitor subscribers-limit (Chassis Extensive)

```
user@host> show system resource-monitor subscribers-limit chassis extensive
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

Client-type : any
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

    fpc : 1
    Client-type : pppoe
        Configured limit    : 0
        Current count       : 1
        Denied count        : 0

    Client-type : any
        Configured limit    : 0
        Current count       : 1
        Denied count        : 0

        pic : 2
```

```
        Client-type : pppoe
            Configured limit    : 0
            Current count       : 1
            Denied count        : 0

        Client-type : any
            Configured limit    : 0
            Current count       : 1
            Denied count        : 0

            port : 0
            Client-type : pppoe
                Configured limit    : 0
                Current count       : 1
                Denied count        : 0

            Client-type : any
                Configured limit    : 0
                Current count       : 1
                Denied count        : 0
```

show system resource-monitor subscribers-limit (MPC)

```
user@host> show system resource-monitor subscribers-limit fpc 1
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

Client-type : any
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0
```

show system resource-monitor subscribers-limit (MPC Extensive)

```
user@host> show system resource-monitor subscribers-limit fpc 1 extensive
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
```

```
    Denied count        : 0

Client-type : any
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

    pic : 2
    Client-type : pppoe
        Configured limit    : 0
        Current count       : 1
        Denied count        : 0

    Client-type : any
        Configured limit    : 0
        Current count       : 1
        Denied count        : 0

        port : 0
        Client-type : pppoe
            Configured limit    : 0
            Current count       : 1
            Denied count        : 0

        Client-type : any
            Configured limit    : 0
            Current count       : 1
            Denied count        : 0
```

show system resource-monitor subscribers-limit (MIC)

```
user@host> show system resource-monitor subscribers-limit fpc 1 pic 2
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

Client-type : any
    Configured limit    : 0
```

```
    Current count       : 1
    Denied count        : 0
```

show system resource-monitor subscribers-limit (MIC Extensive)

```
user@host> show system resource-monitor subscribers-limit fpc 1 pic 2 extensive
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

Client-type : any
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

    port : 0
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

Client-type : any
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0
```

show system resource-monitor subscribers-limit (Port)

```
user@host> show system resource-monitor subscribers-limit fpc 1 pic 2 port 0
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

Client-type : any
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0
```

show system resource-monitor subscribers-limit (Port Extensive)

```
user@host> show system resource-monitor subscribers-limit fpc 1 pic 2 port 0 extensive
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

Client-type : any
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0
```

Release Information

Command introduced in Junos OS 17.3R1.

RELATED DOCUMENTATION

Limiting Subscribers by Client Type and Hardware Element with Resource Monitor

show system resource-monitor fpc

show system resource-monitor ifd-cos-queue-mapping fpc | **26992**

show system rollback

IN THIS SECTION

- Syntax | **27002**
- Description | **27002**
- Options | **27002**
- Required Privilege Level | **27002**
- Sample Output | **27002**
- Release Information | **27003**

Syntax

```
show system rollback number
<compare number | configuration-revision>
```

Description

This command displays the contents of a previously committed configuration, or the differences between two previously committed configurations.

> **NOTE**: The `show system rollback` command is an operational mode command and cannot be issued with `run` from configuration mode.

Options

number	Number of a configuration to view. The output displays the configuration. The range of values is 0 through 49.
compare *number*	(Optional) Number of another previously committed (rollback) configuration to compare to rollback *number*. The output displays the differences between the two configurations. The range of values is 0 through 49.
configuration-revision	(Optional) Display corresponding configuration revision for this rollback number.

Required Privilege Level

view

Sample Output

show system rollback compare

```
user@host> show system rollback 3 compare 1
[edit]
+  interfaces {
+      ge-1/1/1 {
+          unit 0 {
```

```
+             family inet {
+                 filter {
+                     input mf_plp;
+                 }
+                 address 10.1.1.1/10;
+             }
+         }
+     }
+     ge-1/2/1 {
+         unit 0 {
+             family inet {
+                 filter {
+                     input mf_plp;
+                 }
+                 address 10.1.1.1/10;
+             }
+         }
+     }
+     ge-1/3/0 {
+         unit 0 {
+             family inet {
+                 filter {
+                     input mf_plp;
+                 }
+                 address 10.1.1.1/10;
+             }
+         }
+     }
+}
```

show system rollback configuration-revision

```
user@host> show system rollback 0 configuration-revision
The corresponding configuration revision is: re0-1596379942-3
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `configuration-revision` introduced in Junos OS Release 20.4R1 and Junos OS Evolved Release 20.4R1.

show system schema

IN THIS SECTION

- Syntax | 27004
- Description | 27004
- Options | 27004
- Required Privilege Level | 27005
- Sample Output | 27006
- Release Information | 27008

Syntax

```
show system schema module module output-directory path
<filter [filter1 filter2]>
<format format>
<module-name output-module-name>
<output-file-name path>
<version version>
```

Description

Display the Junos OS schema in the specified format. If you do not specify a format, the device displays the schema in YANG.

Options

filter [*filter1 filter2*] Display the schema only for the specified configuration hierarchies when generating the schema for the `configuration` module. Specify a space-delimited list of filters representing each hierarchy to include. When you use the `filter` option, you must also include the `module-name` option. This option is deprecated starting in Junos OS Release 17.4R1

> **NOTE**: In the filter path, the root element represents the top-level `configuration` element in the configuration hierarchy. For example, to only retrieve the `[edit system services]` hierarchy, set the value of `filter` to `/system/services`.

format *format*

(Optional) Data modeling language of the schema. Specify `yang` to display the schema in YANG format.

- **Default:** yang

output-directory *path*

Specify the directory where the schema files will be saved.

> **NOTE**: Starting in Junos OS Release 19.1R2 and 19.2R1, you must specify the `output-directory` option when requesting any schema files. In earlier releases, you can omit the `output-directory` option when requesting a single module to display the module in standard output.

> **NOTE**: To specify the output file in Junos OS Release 15.1 and earlier releases, use the `output-file-name` option.

output-file-name *path*

(Optional) File to which the output is written. If you do not specify an absolute path, the device places the file in the current working directory, which defaults to the user's home directory in **/var/home**. If you omit this option, the output is sent to standard output. This option is deprecated starting in Junos OS Release 16.1.

module *module*

Module for which to display the schema.

module-name *output-module-name*

(Optional) Name used for the generated module. If you also include the `output-directory` option in the command to direct the output to a file, the filename for the output file uses this module name as the filename base and the format as the file extension. This option is deprecated starting in Junos OS Release 17.4R1.

Required Privilege Level

view

Sample Output

show system schema module all-conf

```
user@host> show system schema module all-conf output-directory /var/tmp/yang
user@host>
```

show system schema module (Junos OS Release 19.1 and earlier)

```
user@host> show system schema module junos-conf-root

/*
 * Copyright (c) 2017 Juniper Networks, Inc.
 * All rights reserved.
 */
 module junos-conf-root {
   namespace "http://yang.juniper.net/junos/conf/root";

   prefix jc;

   import junos-common-types {
     prefix jt;
   }

   organization "Juniper Networks, Inc.";

   contact "yang-support@juniper.net";

   description "Junos YANG module for configuration hierarchies.";

   revision 2017-01-01 {
     description "Junos: 17.4R1.17";
   }

   container configuration {
     config true;
     uses juniper-config;
     list groups {
       key name;
       ordered-by user;
```

```
        description "Configuration groups";
        uses juniper-group;
      }
    }
...
```

show system schema module configuration filter (Junos OS Release 17.4 and earlier)

```
user@host> show system schema module configuration filter [/system /interfaces] module-name
config-sys-int
/*
 * Copyright (c) 2016 Juniper Networks, Inc.
 * All rights reserved.
 */

 module config-sys-int {
  namespace "http://yang.juniper.net/yang/1.1/jc/configuration/junos/17.2R1.13";
  prefix jc;
  import junos-extension {
    prefix junos;
  }
  ...
   container configuration {
     config true;
     uses juniper-config;
     list groups {
       key group_name;
       ordered-by user;
       description "Configuration groups";
       uses juniper-group;
     }
   }
   grouping juniper-config {
     container system {
       description "System parameters";
       uses juniper-system;
     }
     container interfaces {
       description "Interface configuration";
```

```
        uses apply-advanced;
    ...
```

Release Information

Command introduced in Junos OS Release 14.2.

`filter`, `module-name`, and `output-directory` options added in Junos OS Release 16.1R1.

`output-file-name` option deprecated in Junos OS Release 16.1R1.

`filter` and `module-name` options deprecated in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Understanding YANG on Devices Running Junos OS

Understanding Junos OS YANG Modules

YANG Modules Overview

Using Juniper Networks YANG Modules

show system security-profile

IN THIS SECTION

- Syntax | 27009
- Description | 27009
- Options | 27009
- Required Privilege Level | 27010
- Output Fields | 27010
- Sample Output | 27011
- Release Information | 27015

Syntax

```
show system security-profile (all-resource | resource)
detail | terse
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display information about a resource allocated to the logical system in a security profile. For each resource specified, the number used by the logical system and the configured maximum and reserved values are displayed.

The `show system security-profile` command can be used by the primary administrator to display resource information for the primary logical system or user logical system. This command can also be used by the user logical system administrator to display resource information for a user logical system.

Options

Either specify `all-resource` to display information about all resources allocated for the logical system, or specify one of the following resources:

- address-book—Address books.
- appfw-rule-set—Application firewall rule set entries.
- appfw-rule—Application firewall rule entries.
- auth-entry—Firewall authentication entries.
- cpu—CPU utilization.
- flow-gate—Flow gates, also known as pinholes.
- flow-session—Flow sessions.
- icap-redirect-profile—ICAP redirect profile resource information.
- nat-cone-binding—Network Address Translation (NAT) cone bindings.
- nat-destination-pool—NAT destination pools.
- nat-destination-rule—NAT destination rules.

- nat-nopat-address—NAT without port address translations.
- nat-pat-address—NAT with port address translations.
- nat-pat-portnum—NAT source port numbers for port translation
- nat-port-ol-ipnumber—NAT port overloading IP numbers.
- nat-rule-referenced-prefix—NAT rule referenced IP-prefixes.
- nat-source-pool—NAT source pools.
- nat-source-rule—NAT source rules.
- nat-static-rule—NAT static rules.
- policy—Security policies.
- policy-with-count—Security policies with a count.
- scheduler—Schedulers.
- zone—Security zones.

detail | terse—(Optional) Display the specified level of output.

The following options are available only to the primary administrator:

- logical-system—Display resource information for a specified user logical system. Specify `all` to display resource information for all logical systems, including the primary logical system.
- root-logical-system—Display resource information for the primary (root) logical system.
- summary—Display summary information about the resource for all logical systems.
- tenant—Display resource information for a specified tenant system. Specify `all` to display resource information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 1658 on page 27011 lists the output fields for the `show system security-profile` command. Output fields are listed in the approximate order in which they appear.

Table 1658: show system security-profile Output Fields

Field Name	Field Description
`logical system name`	Name of the logical system.
`tenant name`	Name of the tenant system.
`security profile name`	Name of the security profile bound to the logical system.
`usage`	Number of resources that are currently being used by the logical system.
`reserved`	Number of resources that are guaranteed to be available to the logical system.
`maximum`	Number of resources that the logical system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.
`CPU control`	`TRUE` if CPU control is enabled or `FALSE` if CPU control is not enabled.
`CPU control target`	Upper limit for CPU utilization on the device. The default value is 80 percent.
`CPU name`	Central point (CP) or services processing unit (SPU). CP utilization and average utilization of all SPUs is shown. The `detail` option shows CPU utilization on each SPU.
`drop rate`	Packets dropped for CPU control.

Sample Output

show system security-profile all-resource

```
user@host> show system security-profile all-resource

resource                                usage    reserved     maximum
```

```
[logical system name:   root-logical-system]
[security profile name: Default-Profile]
address-book                                 0           0         512
auth-entry                                   0           0  2147483647
cpu on CP                                0.00%       1.00%      80.00%
cpu on SPU                               0.00%       1.00%      80.00%
flow-gate                                    0           0      524288
flow-session                                 2           0     6291456
nat-cone-binding                             0           0       65536
nat-destination-pool                         0           0        4096
nat-destination-rule                         0           0        8192
nat-nopat-address                            0           0     1048576
nat-pat-address                              0           0        2048
nat-port-ol-ipnumber                         0           0           4
nat-rule-referenced-prefix                   0           0     1048576
nat-source-pool                              0           0        2048
nat-source-rule                              0           0        8192
nat-static-rule                              0           0       20480
policy                                       0           0       40000
policy-with-count                            0           0        1024
scheduler                                    0           0          64
zone                                         0           0         512
```

show system security-profile all-resource tenant all

```
user@host> show system security-profile all-resource tenant all

resource                                 usage    reserved     maximum

[logical system or tenant name: tn1]
[security profile name:         SP1]
address-book                                 0           0        2000
appfw-profile                                0           0        2048
appfw-rule                                   0           0      114688
appfw-rule-set                               0           0       57344
auth-entry                                   0           0       50000
cpu on CP                                0.00%       0.00%      80.00%
cpu on SPU                               0.00%       0.00%      80.00%
dslite-softwire-initiator                    0           0      100000
flow-gate                                    0           0      524288
```

```
flow-session                                0           0   119537664
icap-redirect-profile                       0           0          64
nat-cone-binding                            0           0     2097152
nat-destination-pool                        0           0        8192
nat-destination-rule                        0           0       30720
nat-interface-port-ol                       0           0         256
nat-nopat-address                           0           0     4194304
nat-pat-address                             0           0     1048576
nat-pat-portnum                             0           0  2576980378
nat-port-ol-ipnumber                        0           0         128
nat-rule-referenced-prefix                  0           0     1048576
nat-source-pool                             0           0       12288
nat-source-rule                             0           0       30720
nat-static-rule                             0           0       30720
policy                                      0           0       80000
policy-with-count                           0           0        1024
scheduler                                   0           0          64
security-log-stream-number                  1           0           3
sla-policy                                  0           0        1024
zone                                        0           0        2000
```

show system security-profile policy

```
user@host> show system security-profile policy
logical system name   security profile name       usage    reserved     maximum

ls-product-design     ls-design-profile               0          40          50
```

show system security-profile cpu

```
user@host> show system security-profile cpu
CPU control: TRUE
CPU control target: 80.00%
logical system name    profile name    CPU name    usage(%)    reserved(%)    drop rate(%)
root-logical-system    Default-Profile CP             0.00%          1.00%           0.00%
root-logical-system    Default-Profile SPU            0.00%          1.00%           0.00%
```

show system security-profile cpu logical-system all

```
user@host> show system security-profile cpu logical-system all
CPU control: TRUE
CPU control target: 80.00%
logical system name    profile name    CPU name    usage(%)    reserved(%)    drop rate(%)
root-logical-system    Default-Profile CP             0.00%          1.00%           0.00%
root-logical-system    Default-Profile SPU            0.00%          1.00%           0.00%
ls-product-design      ls-design-profile CP           0.00%          0.00%           0.00%
ls-product-design      ls-design-profile SPU          0.00%          0.00%           0.00%
ls-marketing-dept      ls-acct-mrkt-profile CP        0.00%          0.00%           0.00%
ls-marketing-dept      ls-acct-mrkt-profile SPU       0.00%          0.00%           0.00%
logical system name    security profile name    usage    reserved    maximum
root-logical-system    Default-Profile         67108864     0      134217728
lsys1                  profile1                193536     6000     134217728
```

show system security-profile cpu summary

```
user@host> show system security-profile cpu summary
CPU control: TRUE
CPU control target: 80.00%

CPU type                :     CP
global used amount      :  0.00%
global maximum quota    : 80.00%
global available amount : 80.00%
total logical systems   :      3
total security profiles :      3
heaviest usage / user   :  0.00%      / root-logical-system
lightest usage / user   :  0.00%      / root-logical-system

CPU type                :    SPU
global used amount      :  0.00%
global maximum quota    : 80.00%
global available amount : 80.00%
total logical systems   :      3
total security profiles :      3
heaviest usage / user   :  0.00%      / root-logical-system
lightest usage / user   :  0.00%      / root-logical-system
```

show system security-profile nat-pat-portnum

```
user@host> show system security-profile cpu nat-pat-portnum
CPU control: TRUE
CPU control target: 80.00%
logical system name    security profile name        usage(%)    reserved(%)    maximum
root-logical-system    Default-Profile CP             67108864          0           134217728
```

show system security-profile nat-pat-portnum summary

```
user@host> show system security-profile nat-pat-portnum summary
global used amount    :67302400
global maximum quota    :134217728
global available amount   :66915328
total logical systems    :2
total security profiles   :1
heaviest usage / user    :193536 / lsys1
```

show system security-profile icap-redirect-profile logical-system all

```
user@host> show system security-profile icap-redirect-profile logical-system all
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                 2           0          64
LSYS1                        SP1                             1          30          64
LSYS2                        SP2                             1          30          64
```

Release Information

Command introduced in Junos OS Release 11.2.

Support for application firewall added in Junos OS Release 11.3.

Option to display all resources for a logical system added in Junos OS Release 11.

Resource information for ports in source NAT pools with port translation added in Release Junos OS 11.4.

The tenant option is introduced in Junos OS Release 18.3R1.

The icap redirect profile option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

security-profile-resources | 16512

show system security-profile scheduler

IN THIS SECTION

- Syntax | 27016
- Description | 27017
- Options | 27017
- Required Privilege Level | 27017
- Output Fields | 27017
- Sample Output | 27019
- Sample Output | 27019
- Sample Output | 27019
- Sample Output | 27020
- Sample Output | 27020
- Release Information | 27020

Syntax

```
show system security-profile scheduler
detail
logical-system <logical-system-name>
root-logical-system
summary
```

```
tenant <tenant-name>
terse
```

Description

Displays the number of schedulers that the user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems. The existing show command for the security-profile scheduler is enhanced with tenant support.

Options

- `detail`—Displays the detailed output.
- `logical-system`—Displays the name of the logical system.
- `root-logical-system`—Displays root logical system as default.
- `summary`—Displays the summary output.
- `tenant`—Displays the name of the tenant system.
- `terse`—Displays the terse output.

Required Privilege Level

view

Output Fields

Table 1659 on page 27017 lists the output fields for the `show system security-profile scheduler` command. Output fields are listed in the approximate order in which they appear.

Table 1659: show security-profile scheduler Output Fields

Field Name	Field Description
`logical-system tenant name`	Name of the logical system or tenant system name.
`security profile name`	Name of the security profile bound to the logical system or tenant system.

Table 1659: show security-profile scheduler Output Fields *(Continued)*

Field Name	Field Description
usage/used amount	Number of resources that are currently being used.
reserved amount	Number of resources that are guaranteed to be available to the logical system or the tenant system.
maximum quota	Number of resources that the logical system or the tenant system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.
global maximum quota	If a logical system or the tenant system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems or tenant system.
global available amount	Number of resources available across all the logical systems or tenant system.
total logical systems/tenants	Total number of logical systems and tenant systems.
heaviest usage/user	Using the most security log streams with the detailed number.
lightest usage/user	Using the least security log streams with the detailed number.
total security profiles	Total number of the resources configured for the security profile.
global used amount	Number of resources used across all the logical systems or tenant systems.

Sample Output

show system security-profile scheduler

```
root@host> show system security-profile scheduler
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                0           0         256
```

Sample Output

show system security-profile scheduler detail

```
root@host> show system security-profile scheduler detail
logical-system tenant name : root-logical-system
security profile name      : Default-Profile
used amount                : 0
reserved amount            : 0
maximum quota              : 256
```

Sample Output

show system security-profile scheduler summary

```
root@host> show system security-profile scheduler summary
global used amount            : 0
global maximum quota          : 256
global available amount       : 256
total logical systems/tenants : 2
total security profiles       : 2
heaviest usage / user         : 0     / root-logical-system ...(2 logical system & tenants)
lightest usage / user         : 0     / root-logical-system ...(2 logical system & tenants)
```

Sample Output

show system security-profile scheduler terse

```
root@host> show system security-profile scheduler terse
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                0           0         256
```

Sample Output

show system security-profile scheduler tenant all

```
root@host> show system security-profile scheduler tenant all
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                0           0         256
```

Release Information

Statement introduced in Junos OS Release 11.2.

The `tenant` option is added in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

show security policies | 25028

show system security-profile secure-wire

IN THIS SECTION

Syntax | 27021

- Description | 27021
- Options | 27021
- Required Privilege Level | 27022
- Output Fields | 27022
- Sample Output | 27023
- Sample Output | 27023
- Sample Output | 27023
- Sample Output | 27024
- Sample Output | 27024
- Sample Output | 27024
- Release Information | 27025

Syntax

```
show system security-profile secure-wire
<all-logical-systems-tenants>
<detail>
<logical-system (all | logical-system-name)>
<root-logical-system>
<summary>
<tenant (all | tenant-system-name)>
<terse>
```

Description

Display information about secure wire resource allocation.

Options

- none—Display information about all configured secure wire resource allocation.
- all-logical-systems-tenants—Display information about all configured logical system and tenant system.
- detail—Display detailed information about all configured logical system and tenant system.

- logical-system (all | *logical-system-name*)—Perform this operation on all logical systems or on a particular logical system.
- root-logical-system—Perform this operation for primary logical systems.
- summary—Display detailed summary of resource allocation.
- tenant (all | *tenant-system-name*)—Secure wire does not support tenant system.
- terse—Display summary of resource allocation.

Required Privilege Level

view

Output Fields

Table 1660 on page 27022 lists the output fields for the `show system security-profile secure-wire` command in the approximate order in which they appear.

Table 1660: show system security-profile secure-wire Output Fields

Field Name	Field Description
logical-system tenant name	Name of the logical system and tenant system.
security profile name	Name of the security profile name.
usage	Number of usage that are currently being used by the logical system.
reserved	Number of resources that are guaranteed to be available to the logical system.
maximum	Number of resources that the logical system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.

Sample Output

show system security-profile secure-wire

```
user@host> show system security-profile secure-wire

logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                0           0         255
```

Sample Output

show system security-profile secure-wire logical-system LSYS1

```
user@host> show system security-profile secure-wire logical-system LSYS1

logical-system tenant name   security profile name       usage    reserved     maximum

LSYS1                        prof1                            1           1         100
```

Sample Output

show system security-profile secure-wire all-logical-systems-tenants

```
user@host> show system security-profile secure-wire all-logical-systems-tenants

logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                0           0         255
LSYS1                        prof1                          1           1         100
```

Sample Output

show system security-profile secure-wire detail

```
user@host> show system security-profile secure-wire detail

logical-system tenant name : root-logical-system
security profile name      : Default-Profile
used amount                : 0
reserved amount            : 0
maximum quota              : 255
```

Sample Output

show system security-profile secure-wire summary

```
user@host> show system security-profile secure-wire summary

global used amount            : 1
global maximum quota          : 255
global available amount       : 254
total logical systems/tenants : 2
total security profiles       : 1
heaviest usage / user         : 1     / LSYS1
lightest usage / user         : 0     / root-logical-system
```

Sample Output

show system security-profile secure-wire terse

```
user@host> show system security-profile secure-wire terse

logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                 0           0         255
```

Release Information

Command introduced in Junos OS Release 12.3X48-D10.

`logical-system` option introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Understanding Secure Wire on Security Devices

Secure Wire for Logical Systems

show system security-profile security-log-stream-number

IN THIS SECTION

- Syntax | 27025
- Description | 27026
- Options | 27026
- Required Privilege Level | 27026
- Output Fields | 27026
- Sample Output | 27028
- Release Information | 27028

Syntax

```
show system security-profile security-log-stream-number
 logical-system (all |logical-system-name)
 tenant (all |tenant-system-name)
```

Description

Display information about a resource allocated to the logical system or tenant system in a security profile. This command can be used by the primary administrator to display resource information for the primary logical system or user logical system.

Options

logical-system (all | *logical-system-name*)	Display resource information for all logical systems, including the primary logical system or a particular logical system.
tenant (all | *tenant-system-name*)	Display resource information for all tenant systems, including the primary logical system or a particular tenant system.

Required Privilege Level

view

Output Fields

Table 1661 on page 27026 lists the output fields for the `show system security-profile security-log-stream-number logical-system all` command. Output fields are listed in the approximate order in which they appear.

Table 1661: show system security-profile security-log-stream-number logical-system all Output Fields

Field Name	Field Description
`logical system name`	Name of the logical system.
`security profile name`	Name of the security profile bound to the logical system.
`usage`	Number of resources that are currently being used by the logical system.
`reserved`	Number of resources that are guaranteed to be available to the logical system.

Table 1661: show system security-profile security-log-stream-number logical-system all Output Fields *(Continued)*

Field Name	Field Description
maximum	Number of resources that the logical system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.
root-logical-system	Display resource information for the primary (root) logical system.
Default-Profile	Specify the authentication profile to use if no profile is specified.

Table 1662 on page 27027 lists the output fields for the `show system security-profile security-log-stream-number tenant all` command. Output fields are listed in the approximate order in which they appear.

Table 1662: show system security-profile security-log-stream-number tenant all Output Fields

Field Name	Field Description
logical-system tenant name	Name of the tenant system.
security profile name	Name of the security profile bound to the tenant system.
usage	Number of resources that are currently being used by the tenant system.
reserved	Number of resources that are guaranteed to be available to the tenant system.
maximum	Number of resources that the tenant system can use. The maximum resource does not guarantee that the amount specified for the resource in the security profile is available. The maximum resource is not applicable for CPU resources.
root-logical-system	Display resource information for the primary (root) logical system.
Default-Profile	Specify the authentication profile to use if no profile is specified.

Sample Output

show system security-profile security-log-stream-number logical-system all

```
user@host> show system security-profile security-log-stream-number logical-system all
logical system name   security profile name       usage    reserved     maximum
root-logical-system   Default-Profile                1           0           3
LSYS1                 sp1                            0           1           3
LSYS2                 sp2                            1           0           3
```

show system security-profile security-log-stream-number tenant all

```
user@host> show system security-profile security-log-stream-number tenant all
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                1           0           8
LSYS1                        p1                             1           1           2
TN1                          sp2                            0           0           8
```

Release Information

Command introduced in Junos OS Release 18.2R1.

`tenant` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

security-profile-resources | 16512

show system security-profile security-log-stream-number detail

IN THIS SECTION

Syntax | 27029

- Description | 27029
- Options | 27029
- Required Privilege Level | 27029
- Output Fields | 27029
- Sample Output | 27030
- Release Information | 27031

Syntax

```
show system security-profile security-log-stream-number detail
logical-system (all | logical-system-name)
tenant (all | tenant-system-name)
```

Description

Display information about a resource allocated to the logical system or tenant system in a security profile with security log stream number.

Options

logical-system-name (all | *logical-system-name*)	Display resource information for all logical systems, including the primary logical system or a particular logical system.
tenant (all | *tenant-system-name*)	Display resource information for all tenant systems, including the primary logical system or a particular tenant system.

Required Privilege Level

view

Output Fields

Table 1663 on page 27030 lists the output fields for the `show system security-profile security-log-stream-number summary` command. Output fields are listed in the approximate order in which they appear.

Table 1663: show system security-profile security-log-stream-number summary Output Fields

Field Name	Field Description
`logical system name`	Displays the logical system name
`security profile name`	Name of the security profile
`used amount`	Number of resources that are currently being used by the logical system.
reserved amount	Reserved quota that guarantees that the resource amount specified is always available to the logical system.
maximum quota	Maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

Sample Output

show system security-profile security-log-stream-number detail logical-system all

```
user@host> show system security-profile security-log-stream-number detail logical-system all
logical system name     : root-logical-system
security profile name   : Default-Profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8

logical system name     : lsys0
security profile name   : lsys_profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8
```

```
logical system name     : lsys1
security profile name   : lsys_profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8

logical system name     : lsys2
security profile name   : lsys_profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8
```

show system security-profile security-log-stream-number detail tenant all

```
user@host> show system security-profile security-log-stream-number detail tenant all
logical-system tenant name : root-logical-system
security profile name      : Default-Profile
used amount                : 1
reserved amount            : 0
maximum quota              : 8

logical-system tenant name : LSYS1
security profile name      : p1
used amount                : 1
reserved amount            : 1
maximum quota              : 2

logical-system tenant name : TN1
security profile name      : sp2
used amount                : 0
reserved amount            : 0
maximum quota              : 8
```

Release Information

Command introduced in Junos OS Release 18.2R1.

The `tenant` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

security-profile-resources | 16512

show system security-profile security-log-stream-number summary

IN THIS SECTION

- Syntax | 27032
- Description | 27032
- Options | 27032
- Required Privilege Level | 27032
- Output Fields | 27033
- Sample Output | 27033
- Release Information | 27034

Syntax

```
show system security-profile security-log-stream-number summary  (detail |terse)
```

Description

Display summary information about the resource for all logical systems.

Options

detail	Display detailed output.
terse	Display terse output (default).

Required Privilege Level

view

Output Fields

Table 1664 on page 27033 lists the output fields for the `show system security-profile security-log-stream-number summary` command. Output fields are listed in the approximate order in which they appear.

Table 1664: show system security-profile security-log-stream-number summary Output Fields

Field Name	Field Description
`global used amount`	Number of resources that are currently being used by the logical system.
`global maximum quota`	Number of resources that the logical system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.
`global available amount`	Number of resources that are guaranteed to be available to the logical system.
`total logical systems`	Total number of logical systems
`total security profiles`	Total number of resources configured for the security profile
`heaviest usage / user`	Using the most security log streams with the detailed number
`lightest usage / user`	Using the least security log streams with the detailed number

Sample Output

show system security-profile security-log-stream-number summary

```
user@host> show system security-profile security-log-stream-number summary
global used amount      : 0
global maximum quota    : 32
global available amount : 32
total logical systems   : 1
total security profiles : 0
```

```
heaviest usage / user   : 0     / root-logical-system
lightest usage / user   : 0     / root-logical-system
```

show system security-profile security-log-stream-number summary

```
user@host> show system security-profile security-log-stream-number summary
global used amount            : 1
global maximum quota          : 100
global available amount       : 99
total logical systems/tenants : 1
total security profiles       : 2
heaviest usage / user         : 1     / root-logical-system
lightest usage / user         : 1     / TN1
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

security-profile-resources | 16512

show system services dhcp binding

IN THIS SECTION

- Syntax | 27035
- Description | 27035
- Options | 27035
- Required Privilege Level | 27035
- Output Fields | 27035
- Sample Output | 27036
- Release Information | 27038

Syntax

```
show system services dhcp binding
<detail>
<address>
```

Description

(EX Series switches only) Display Dynamic Host Configuration Protocol (DHCP) server client binding information.

Options

none Display brief information about all active client bindings.

detail (Optional) Display detailed information about all active client bindings.

address (Optional) Display detailed client binding information for the specified IP address only.

Required Privilege Level

view and system

Output Fields

Table 1665 on page 27035 describes the output fields for the `show system services dhcp binding` command. Output fields are listed in the approximate order in which they appear.

Table 1665: show system services dhcp binding Output Fields

Field Name	Field Description	Level of Output
`Allocated address`	List of IP addresses the DHCP server has assigned to clients.	All levels
`MAC address`	Corresponding media access control (MAC) hardware address of the client.	All levels

Table 1665: show system services dhcp binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Client identifier`	(*address* option only) Client's unique identifier (represented by an ASCII string or hexadecimal digits). This identifier is used by the DHCP server to index its database of address bindings.	All levels
`Binding Type`	Type of binding assigned to the client. DHCP servers can assign a dynamic binding from a pool of IP addresses or a static binding to one or more specific IP addresses.	All levels
`Lease Expires at`	Time the lease expires or `never` for leases that do not expire.	All levels
`Lease Obtained at`	(*address* option only) Time the client obtained the lease from the DHCP server.	`detail`
`State`	Status of the binding. Bindings can be active or expired.	`detail`
`Pool`	Address pool that contains the IP address assigned to the client.	`detail`
`Request received on`	Interface on which the DHCP message exchange occurs. The IP address pool is configured based on the interface's IP address. If a relay agent is used, its IP address is also displayed.	`detail`
`DHCP options`	User-defined options created for the DHCP server. If no options have been defined, this field is blank.	`detail`

Sample Output

show system services dhcp binding

```
user@host> show system services dhcp binding

Allocated address   MAC address        Binding Type   Lease expires at
```

```
192.168.1.2         00:a0:12:00:12:ab  static         never
192.168.1.3         00:a0:12:00:13:02  dynamic        2004-05-03 13:01:42 PDT
```

show system services dhcp binding address

```
user@host> show system services dhcp binding 192.168.1.3

DHCP binding information:
Allocated address: 192.168.1.3
Mac address: 00:a0:12:00:12:ab
Client identifier
61 63 65 64 2d 30 30 3a 61 30 3a 31 32 3a 30 30aced-00:a0:12:00
3a 31 33 3a 30 32:13:02

Lease information:
    Binding Type dynamic
    Obtained at 2004-05-02 13:01:42 PDT
    Expires at 2004-05-03 13:01:42 PDT
```

show system services dhcp binding address detail

```
user@host> show system services dhcp binding 192.168.1.3 detail
DHCP binding information:
Allocated address        192.168.1.3
MAC address  00:a0:12:00:12:ab
    Pool              192.168.1.0/24
Request received on fe-0/0/0, relayed by 192.168.4.254

Lease information:
    Type              DHCP
    Obtained at       2004-05-02 13:01:42 PDT
    Expires at        2004-05-03 13:01:42 PDT
State active

DHCP options:
    Name: name-server, Value: { 6.6.6.6, 6.6.6.7 }
    Name: domain-name, Value: mydomain.tld
    Code: 19, Type: flag, Value: off
```

```
Code: 40, Type: string, Value: domain.tld
Code: 32, Type: ip-address, Value: 3.3.3.33
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear system services dhcp binding | 4230

show system services dhcp client

IN THIS SECTION

- Syntax | 27038
- Description | 27039
- Options | 27039
- Required Privilege Level | 27039
- Output Fields | 27039
- Sample Output | 27041
- Sample Output | 27042
- Sample Output | 27042
- Release Information | 27043

Syntax

```
show system services dhcp client
<                    interface-name                  >
<statistics>
```

Description

Display information about DHCP clients.

Options

- none—Display DHCP information for all interfaces.
- *interface-name*—(Optional) Display DHCP information for the specified interface.
- statistics—(Optional) Display DHCP client statistics.

Required Privilege Level

view and system

Output Fields

Table 1666 on page 27039 lists the output fields for the show system services dhcp client command. Output fields are listed in the approximate order in which they appear.

Table 1666: show system services dhcp client Output Fields

Field Name	Field Description
Logical Interface Name	Name of the logical interface.
Client Status	State of the client binding.
Vendor Identifier	Vendor ID.
Server Address	IP address of the DHCP server.
Address obtained	IP address obtained from the DHCP server.
Lease Obtained at	Date and time the lease was obtained.

Table 1666: show system services dhcp client Output Fields *(Continued)*

Field Name	Field Description
`Lease Expires in`	(EX Series switches only) Time the current lease expires in (seconds).
`Lease Expires at`	Date and time the lease expires.
`DHCP Options`	• `Name: server-identifier`, `Value:` IP address of the name server. • `Name: device`, `Value:` IP address of the name device. • `Name: domain-name`, `Value:` Name of the domain.
`Packets dropped`	`Total` packets dropped.
`Messages received`	`Number of the following DHCP messages received:` • `DHCPOFFER`—First packet received on a logical interface when DHCP is enabled. • `DHCPACK`—When received from the server, the client sends an ARP request for that address and adds a (ARP response) timer for 4 seconds and stops the earlier timer added for DHCPACK. • `DHCPNAK`—When a DHCPNAK is received instead of DHCPACK, the logical interface sends a DHCPDISCOVER packet.

Table 1666: show system services dhcp client Output Fields *(Continued)*

Field Name	Field Description
Messages sent	Number of the following DHCP messages sent: • DHCPDECLINE—Packet sent when ARP response is received and there is a conflict. The logical interface sends a new DHCPDISCOVER packet. • DHCPDISCOVER—Packet sent on the interface for which the DHCP client is enabled. • DHCPREQUEST—Packet sent to the DHCP server after accepting the DHCPOFFER. After sending the DHCPREQUEST, the device adds a retransmission-interval timer. • DHCPINFORM—Packet sent to the DHCP server for local configuration parameters. • DHCPRELEASE—Packet sent to the DHCP server to relinquish network address and cancel remaining lease. • DHCPRENEW—Packet sent to the DHCP server to renew the address. The next message to be sent will be a DHCPREQUEST message, which will be unicast directly to the server. • DHCPREBIND—Packet sent to any server to renew the address. The next message to be sent will be a DHCPREQUEST message, which will be broadcast.

Sample Output

show system services dhcp client

```
user@host> show system services dhcp client
Logical Interface name         ge-0/0/34.0
        Hardware address        00:1f:12:38:5f:e5
        Client status           bound
        Address obtained        10.0.0.2
        Update server           disabled
        Lease obtained at       2013-12-23 08:11:40 UTC
        Lease expires in        93
        Lease expires at        2013-12-23 08:13:20 UTC
```

```
DHCP options:
    Name: server-identifier, Value: 10.0.0.1
    Code: 1, Type: ip-address, Value: 255.255.255.0
```

Sample Output

show system services dhcp client ge-0/0/34.0

```
user@host> show system services dhcp client ge-0/0/34.0
 Logical Interface name         ge-0/0/34.0
        Hardware address        00:1f:12:38:5f:e5
        Client status           bound
        Address obtained        10.0.0.2
        Update server           disabled
        Lease obtained at       2013-12-23 08:11:40 UTC
        Lease expires in        87
        Lease expires at        2013-12-23 08:13:20 UTC

DHCP options:
    Name: server-identifier, Value: 10.0.0.1
    Code: 1, Type: ip-address, Value: 255.255.255.0
```

Sample Output

show system services dhcp client statistics

```
user@host> show system services dhcp client statistics
Packets dropped:
    Total                      0
Messages received:
    DHCPOFFER                  0
    DHCPACK                    8
    DHCPNAK                    0
Messages sent:
    DHCPDECLINE                0
    DHCPDISCOVER               0
    DHCPREQUEST                1
    DHCPINFORM                 0
```

```
    DHCPRELEASE                0
    DHCPRENEW                  7
    DHCPREBIND                 0
```

Release Information

Command introduced in Junos OS Release 8.5.

Command introduced in Junos OS Release 9.0 for EX Series switches.

RELATED DOCUMENTATION

dhcp (DHCP Client) | 5598

request system services dhcp | 15185

show system services dhcp conflict

IN THIS SECTION

- Syntax | 27043
- Description | 27044
- Options | 27044
- Required Privilege Level | 27044
- Output Fields | 27044
- Sample Output | 27044
- Release Information | 27045

Syntax

```
show system services dhcp conflict
```

Description

(J Series routers only and EX Series switches) Display Dynamic Host Configuration Protocol (DHCP) client-detected conflicts for IP addresses. When a conflict is detected, the DHCP server removes the address from the address pool.

Options

This command has no options.

Required Privilege Level

view and system

Output Fields

Table 1667 on page 27044 describes the output fields for the `show system services dhcp conflict` command. Output fields are listed in the approximate order in which they appear.

Table 1667: show system services dhcp conflict Output Fields

Field Name	Field Description
`Detection time`	Date and time the client detected the conflict.
`Detection method`	How the conflict was detected.
`Address`	IP address where the conflict occurs. The addresses in the conflicts list remain excluded from the pool until you use a `clear system services dhcp conflict` command to manually clear the list.

Sample Output

show system services dhcp conflict

```
user@host> show system services dhcp conflict

Detection time           Detection method  Address
```

```
2004-08-03 19:04:00 PDT  ARP              10.0.0.1
2004-08-04 04:23:12 PDT  Ping             10.0.0.2
2004-08-05 21:06:44 PDT  Client           10.0.0.3
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear system services dhcp conflict | 4232

show system services dhcp global

IN THIS SECTION

- Syntax | 27045
- Description | 27045
- Options | 27046
- Required Privilege Level | 27046
- Output Fields | 27046
- Sample Output | 27046
- Release Information | 27047

Syntax

```
show system services dhcp global
```

Description

(J Series routers and EX Series switches only) Display Dynamic Host Configuration Protocol (DHCP) global configuration options. Global options apply to all scopes and clients served by the DHCP server.

Global options are overridden if specified otherwise in scope or client options. Scope options apply to specific subnets or ranges of addresses. Client options apply to specific clients.

Options

This command has no options.

Required Privilege Level

view and system

Output Fields

Table 1668 on page 27046 describes the output fields for the `show system services dhcp global` command. Output fields are listed in the approximate order in which they appear.

Table 1668: show system services dhcp global Output Fields

Field Name	Field Description
`BOOTP lease length`	Length of lease time assigned to BOOTP clients.
`Default lease time`	Lease time assigned to clients that do not request a specific lease time.
`Minimum lease time`	Minimum time a client retains an IP address lease on the server.
`Maximum lease time`	Maximum time a client can retain an IP address lease on the server.
`DHCP options`	User-defined options created for the DHCP server. If no options have been defined, this field is blank.

Sample Output

show system services dhcp global

```
user@host> show system services dhcp global
```

```
Global settings:
    BOOTP lease length         infinite

DHCP lease times:
    Default lease time         1 hour
    Minimum lease time         2 hours
    Maximum lease time         infinite

DHCP options:
    Name: name-server, Value: { 6.6.6.6, 6.6.6.7 }
    Name: domain-name, Value: mydomain.tld
    Code: 19, Type: flag, Value: off
    Code: 40, Type: string, Value: domain.tld
    Code: 32, Type: ip-address, Value: 3.3.3.33
```

Release Information

Command introduced before Junos OS Release 7.4.

show system services dhcp pool

IN THIS SECTION

- Syntax | 27048
- Description | 27048
- Options | 27048
- Required Privilege Level | 27048
- Output Fields | 27048
- Sample Output | 27050
- Release Information | 27051

Syntax

```
show system services dhcp pool
<detail>
<subnet-address>
```

Description

(J Series routers and EX Series switches only) Display Dynamic Host Configuration Protocol (DHCP) server IP address pools.

Options

none	Display brief information about all IP address pools.
detail	(Optional) Display detailed information.
subnet-address	(Optional) Display information for the specified subnet address.

Required Privilege Level

view and system

Output Fields

Table 1669 on page 27048 describes the output fields for the `show system services dhcp pool` command. Output fields are listed in the approximate order in which they appear.

Table 1669: show system services dhcp pool Output Fields

Field Name	Field Description	Level of Output
`Pool name`	Subnet on which the IP address pool is defined.	None specified
`Low address`	Lowest address in the IP address pool.	None specified
`High address`	Highest address in the IP address pool.	None specified

Table 1669: show system services dhcp pool Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Excluded addresses`	Addresses excluded from the address pool.	None specified
`Subnet`	(*subnet-address* option only) Subnet to which the specified address pool belongs.	None specified
`Address range`	(*subnet-address* option only) Range of IP addresses in the address pool.	None specified
`Addresses assigned`	Number of IP addresses in the pool that are assigned to DHCP clients and the total number of IP addresses in the pool.	`detail`
`Active`	Number of assigned IP addresses in the pool that are active.	`detail`
`Excluded`	Number of assigned IP addresses in the pool that are excluded.	`detail`
`Default lease time`	Lease time assigned to clients that do not request a specific lease time.	`detail`
`Minimum lease time`	Minimum time a client can retain an IP address lease on the server.	`detail`
`Maximum lease time`	Maximum time a client can retain an IP address lease on the server.	`detail`
`DHCP options`	User-defined options created for the DHCP server. If no options have been defined, this field is blank.	`detail`

Sample Output

show system services dhcp pool

```
user@host> show system services dhcp pool

Pool name    Low address   High address   Excluded addresses
192.0.2.0/24 192.0.2.2     192.0.2.254     192.0.2.1
```

show system services dhcp pool subnet-address

```
user@host> show system services dhcp pool 192.0.2.0/24

Pool information:
    Subnet                     192.0.2.0/24
    Address range              192.0.2.2 - 192.0.2.254
    Addresses assigned         2/253
```

show system services dhcp pool subnet-address detail

```
user@host> show system services dhcp pool 192.0.2.0/24 detail

Pool information:
    Subnet                     192.0.2.0/24
    Address range              192.0.2.2 - 192.0.2.254
    Addresses assigned         2/253
      Active: 1, Excluded: 1

DHCP lease times:
    Default lease time         1 hour
    Minimum lease time         2 hours
    Maximum lease time         infinite

DHCP options:
    Name: name-server, Value: { 6.6.6.6, 6.6.6.7 }
    Name: domain-name, Value: mydomain.tld
    Name: router, Value: { 192.0.2.1 }
    Name: server-identifier, Value: 192.0.2.1
```

```
    Code: 19, Type: flag, Value: off
    Code: 40, Type: string, Value: domain.tld
    Code: 32, Type: ip-address, Value: 192.0.2.1
```

Release Information

Command introduced before Junos OS Release 7.4.

show system services dhcp relay-statistics

IN THIS SECTION

- Syntax | 27051
- Description | 27051
- Required Privilege Level | 27051
- Output Fields | 27052
- Sample Output | 27052
- Release Information | 27053

Syntax

```
show system services dhcp relay-statistics
```

Description

Display information about the DHCP relay.

Required Privilege Level

view and system

Output Fields

Table 1670 on page 27052 lists the output fields for the `show system services dhcp relay-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1670: show system services dhcp relay-statistics Output Fields

Field Name	Field Description
`Received packets`	Total DHCP packets received.
`Forwarded packets`	Total DHCP packet forwarded.
`Dropped packets`	Total DHCP packets dropped for the following reasons: • `Due to a missing interface in the relay database`—Number of packets discarded because they did not belong to a configured interface. • `Due to a missing matching routing instance`—Number of packets discarded because they did not belong to a configured routing instance. • `Due to an error during packet read`—Number of packets discarded because of a system read error. • `Due to an error during packet send`—Number of packets that the DHCP relay application could not send. • `Due to an invalid server address`—Number of packets discarded because an invalid server address was specified. • `Due to a missing valid local address`—Number of packets discarded because there was no valid local address. • `Due to a missing route to the server or client`—Number of packets discarded because there were no addresses available for assignment.

Sample Output

show system services dhcp relay-statistics

```
user@host> show system services dhcp relay-statistics
  Received packets: 4
```

```
Forwarded packets: 4
Dropped packets: 4
  Due to missing interface in relay database: 4
  Due to missing matching routing instance: 0
  Due to an error during packet read: 0
  Due to an error during packet send: 0
  Due to invalid server address: 0
  Due to missing valid local address: 0
  Due to missing route to server/client: 0
```

Release Information

Command introduced in Junos OS Release 8.5 .

RELATED DOCUMENTATION

dhcp (DHCP Client) | 5598

show system services dhcp statistics

IN THIS SECTION

- Syntax | 27054
- Description | 27054
- Options | 27054
- Required Privilege Level | 27054
- Output Fields | 27054
- Sample Output | 27056
- Release Information | 27057

Syntax

```
show system services dhcp statistics
```

Description

(J Series routers and EX Series switches only) Display Dynamic Host Configuration Protocol (DHCP) server statistics.

Options

This command has no options.

Required Privilege Level

view and system

Output Fields

Table 1671 on page 27054 describes the output fields for the `show system services dhcp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1671: show system services dhcp statistics Output Fields

Field Name	Field Description
`Default lease time`	Lease time assigned to clients that do not request a specific lease time.
`Minimum lease time`	Minimum time a client can retain an IP address lease on the server.
`Maximum lease time`	Maximum time a client can retain an IP address lease on the server.

Table 1671: show system services dhcp statistics Output Fields *(Continued)*

Field Name	Field Description
`Packets dropped`	Total number of packets dropped and number of packets dropped because of: • Invalid hardware address • Invalid opcode • Invalid server address • No available address • No interface match • No routing instance match • No valid local addresses • Packet too short • Read error • Send error
`Messages received`	Number of the following message types sent from DHCP clients and received by the DHCP server: • BOOTREQUEST • DHCPDECLINE • DHCPDISCOVER • DHCPINFORM • DHCPRELEASE • DHCPREQUEST

Table 1671: show system services dhcp statistics Output Fields *(Continued)*

Field Name	Field Description
`Messages sent`	Number of the following message types sent from the DHCP server to DHCP clients: • BOOTREPLY • DHCPACK • DHCPOFFER • DHCPNAK

Sample Output

show system services dhcp statistics

```
user@host> show system services dhcp statistics

DHCP lease times:
    Default lease time         1 hour
    Minimum lease time         2 hours
    Maximum lease time         infinite

Packets dropped:
    Total                      0
    Bad hardware address       0
    Bad opcode                 0
    Invalid server address     0
    No available addresses     0
    No interface match         0
    No routing instance match  0
    No valid local address     0
    Packet too short           0
    Read error                 0
    Send error                 0

Messages received:
    BOOTREQUEST                0
    DHCPDECLINE                0
```

```
    DHCPDISCOVER              0
    DHCPINFORM                0
    DHCPRELEASE               0
    DHCPREQUEST               0

Messages sent:
    BOOTREPLY                 0
    DHCPACK                   0
    DHCPOFFER                 0
    DHCPNAK                   0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

clear system services dhcp statistics | 4233

show system services dns dns-proxy

IN THIS SECTION

- Syntax | 27058
- Description | 27058
- Options | 27058
- Required Privilege Level | 27058
- Output Fields | 27058
- Sample Output | 27060
- Release Information | 27061

Syntax

```
show system services dns dns-proxy
```

Description

Display domain name system (DNS) proxy information. This option is supported on the SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550M devices.

Options

- `none`—Display DNS proxy statistics information.
- `cache`—(Optional) Display the DNS proxy cache.
- `statistics`—(Optional) Display the DNS proxy statistics.

Required Privilege Level

view

Output Fields

Table 1672 on page 27059 lists the output fields for the `show system services dns-proxy` command. Output fields are listed in the approximate order in which they appear.

Table 1672: show system services dns-proxy

Field Name	Field Description
`DNS proxy statistics`	Display information about the DNS proxy. • `Status`—State of the proxy server as Enabled or disabled. • `Queries received`—Number of DNS queries received by the DNS proxy. • `Responses sent`—Number of DNS responses sent by the DNS proxy. • `Queries forwarded`—Number of DNS queries forwarded by the DNS proxy. • `Negative responses`—Number of negative responses the DNS proxy sent to the DNS client. • `Retry requests`—Number of retries the DNS proxy received from the DNS client. • `Pending requests`—Number of pending queries the DNS proxy has yet to send the DNS client a response for. • `Server failures`—Number of DNS proxy server failures.
`Hostname`	Hostname of the host that has been cached.
`IP address`	IP address of the host.
`Time-to-live`	Length of time before an entry is purged from the DNS cache.
`Type`	Type of DNS Resource Record. For example, A records refer to IPv4 host addresses.
`Class`	Class of DNS. A parameter used to define a DNS Resource Record. For example, IN class is used for Internet domain names.

Sample Output

show system services dns-proxy statistics

```
user@host> show system services dns-proxy statistics
DNS proxy statistics        :
    DNS proxy statistics          :
    Status                     : enabled
    IPV4 Queries received      : 30
    IPV6 Queries received      : 0
    Responses sent             : 30
    Queries forwarded          : 13
    Negative responses         : 23
    Positive responses         : 23
    Retry requests             : 0
    Pending requests           : 0
    Server failures            : 0
    Interfaces                 : fe-0/0/0.0, fe-1/0/1.0
```

show system services dns-proxy cache

```
user@host>show system services dns-proxy cache
Hostname                  Time-to-live      Type        Class   IP address/Hostname
device1.example.com        408               A           IN      207.17.137.229
device2.example.com        408               A           IN      192.0.2.50
device3.example.com        408               A           IN      192.0.2.11
device4.example.com        408               A           IN      10.209.194.131
device5.example.com        408               A           IN      10.10.4.202
device6.example.com        408               A           IN      10.16.0.11
device7.example.com        408               A           IN      192.0.2.100
a.l.example.com            408               A           IN      203.0.113.9
b.l.example.com            408               A           IN      198.51.100.9
maps.l.example.com         408               A           IN      198.51.100.104
c.l.example.com            408               A           IN      198.51.100.91
d.l.example.com            408               A           IN      198.51.100.89
e.l.example.com            408               A           IN      198.51.100.39
g.l.example.com            408               A           IN      198.51.100.69
device8.example.com.       408               A           IN      192.0.2.123
mail.example.com.          408               CNAME       IN      device8.example.com.
```

show system services dns-proxy cache <viewname V1>

```
user@host>show system services dns-proxy cache <viewname V1>
Hostname                  Time-to-live   Type        Class    IP address/Hostname
device1.example.com.      495            A           IN       198.51.100.123
mail.example.com.         495            CNAME       IN       device1.example.com.
```

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

clear system services dns dns-proxy | 4235

dns-proxy | 5862

show system services dynamic-dns

IN THIS SECTION

- Syntax | 27061
- Description | 27062
- Required Privilege Level | 27062
- Output Fields | 27062
- Sample Output | 27062
- Release Information | 27063

Syntax

```
show system services dynamic-dns
```

Description

Display information about dynamic DNS clients. This option is supported on the SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550M devices.

Required Privilege Level

view

Output Fields

Table 1673 on page 27062 lists the output fields for the `show system services dynamic-dns` command. Output fields are listed in the approximate order in which they appear.

Table 1673: show system services dynamic-dns

Field Name	Field Description
`Hostname`	Hostname of the registered client
`Server`	DDNS server name
`Agent`	Name of the DDNS agent
`Last response`	Status of the last response
`Last update`	Date and time of the last update
`Interface`	Name of the interface

Sample Output

show system services dynamic-dns client

```
user@host> show system services dynamic-dns client
Internal hostname      Server                  Last response
device1.example.com         example.org                 success
```

```
device2.example.com.        example.org                failure
```

show system services dynamic-dns client detail

```
user@host>show system services dynamic-dns client detail
Hostname     : device1.example.com
Server       : example.org
Agent        : branch-0.1
Last response: success
Last update  : 2006-08-29 04:02:52 PDT
Interface    : fe-0/0/0.0

Hostname     : device2.example.com
Server       : example.org
Agent        : Branch-0.1
Last response: failure
Last update  : 2006-08-29 04:03:03 PDT
Interface    : fe-0/0/0.0
```

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

dynamic-dns | 6106

show system services outbound-ssh jefflee

IN THIS SECTION

Syntax | 27064

- Description | 27064
- Options | 27064
- Required Privilege Level | 27064
- Output Fields | 27064
- Sample Output | 27065

Syntax

```
show system services outbound-ssh client <name> (brief | detail)
```

Description

Display information about an outbound-ssh client.

Options

This command has no options.

Required Privilege Level

system services

Output Fields

Table 1 on page 27064 lists the output fields for the `show system services outbound-ssh client` command. Output fields are listed in the approximate order in which they appear.

Table 1674: show system services outbound-ssh client Output Fields

Field Name	Field Description
`Client Name`	Name of client.
`Current Application Server`	IP address of application server.

Table 1674: show system services outbound-ssh client Output Fields *(Continued)*

Field Name	Field Description
Proxy In Use	Type of proxy used (DHCP of Static).
Connection Status	Status of connection (Up or Down).
Connection Uptime	Duration of connection.

Sample Output

show system services outbound-ssh client

```
user@host> show system services outbound-ssh client

------------   ------------------------   ------    -----------  ----------
Client         Current Application        Proxy     Connection   Connection
Name           Server                     In Use    Status       Uptime
------------   ------------------------   ------    -----------  ----------
nsm_cont       10.220.0.191:6066          Static    Up           00:01:24
nsm_cont_2     10.220.0.191:6066          DHCP      Down         None
```

show system services outbound-ssh client brief

```
user@host> show system services outbound-ssh client brief

------------   ------------------------   ------    -----------  ----------
Client         Current Application        Proxy     Connection   Connection
Name           Server                     In Use    Status       Uptime
------------   ------------------------   ------    -----------  ----------
nsm_cont       10.220.0.191:6066          Static    Up           00:01:46
nsm_cont_2     10.220.0.191:6066          DHCP      Down         None
```

show system services outbound-ssh client detail

```
user@host> show system services outbound-ssh client detail

------------------------------
Outbound SSH Connection Status
------------------------------

Application ID :  nsm_cont
   Device ID                  :  dsm_new
   Current Application Server :  10.220.0.191:6066
   Proxy Type                 :  Static
   Proxy Server               :  10.209.0.224:8443
   Proxy Authentication       :  Enabled
   Connection Status          :  Up
   Connection Uptime          :  00:01:55

Application ID :  nsm_cont_2
   Device ID                  :  dsm_new_2
   Current Application Server :  10.220.0.191:6066
   Proxy Type                 :  DHCP
   Proxy Authentication       :  Disabled
   Connection Status          :  Down
   Connection Uptime          :  None
```

show system services outbound-ssh client nsm_cont

```
user@host> show system services outbound-ssh client nsm_cont

------------   ------------------------   ------    -----------  ----------
Client         Current Application        Proxy     Connection   Connection
Name           Server                     In Use    Status       Uptime
------------   ------------------------   ------    -----------  ----------
nsm_cont       10.220.0.191:6066          Static    Up           00:00:05
```

show system services outbound-ssh client nsm_cont brief

```
user@host> show system services outbound-ssh client nsm_cont brief

------------   ------------------------   ------    -----------  ----------
Client         Current Application        Proxy     Connection   Connection
Name           Server                     In Use    Status       Uptime
------------   ------------------------   ------    -----------  ----------
nsm_cont       10.220.0.191:6066          Static    Up           00:00:08
```

show system services outbound-ssh client nsm_cont detail

```
user@host> show system services outbound-ssh client nsm_cont detail

------------------------------
Outbound SSH Connection Status
------------------------------

Application ID :  nsm_cont
   Device ID                  :  dsm_new
   Current Application Server :  10.220.0.191:6066
   Proxy Type                 :  Static
   Proxy Server               :  10.209.0.224:8443
   Proxy Authentication       :  Enabled
   Connection Status          :  Up
   Connection Uptime          :  00:00:13
```

Command introduced in Junos OS Release 23.4.

show system services service-deployment

IN THIS SECTION

Syntax | 27068

- Description | 27068
- Options | 27068
- Required Privilege Level | 27068
- Output Fields | 27068
- Sample Output | 27069
- Release Information | 27069

Syntax

```
show system services service-deployment
```

Description

Display information about a Session and Resource Control (SRC) client.

Options

This command has no options.

Required Privilege Level

system

view

Output Fields

Table 1675 on page 27069 lists the output fields for the `show system services service-deployment` command. Output fields are listed in the approximate order in which they appear.

Table 1675: show system services service-deployment Output Fields

Field Name	Field Description
PDT Keepalive settings	Configured PDT keepalive interval, in seconds.
Keepalives sent	Number of keepalives sent.
Notifications sent	Number of notifications sent.
Last update from peer	Time at which the last update from a peer was received.

Sample Output

show system services service-deployment

```
user@host> show system services service-deployment
Connected to 192.0.2.0 port 10288 since 2004-05-03 11:04:34 PDT Keepalive settings: Interval 15
seconds Keepalives sent: 750 Notifications sent: 0 Last update from peer: 00:00:06 ago
```

Release Information

Command introduced before Junos OS Release 7.4.

show system snapshot (Junos OS)

IN THIS SECTION

- Syntax | 27070
- Syntax (EX Series Switches) | 27070
- Description | 27070
- Options | 27070

- Required Privilege Level | 27071
- Output Fields | 27071
- Sample Output | 27071
- Release Information | 27073

Syntax

```
show system snapshot
```

Syntax (EX Series Switches)

```
show system snapshot
<all-members | local | member member-id>
<media (external | internal)>
```

Description

This command displays information about the backup software:

- On the routers, this command display information about the backup software, which is located in the **/altroot**, and **/altconfig** file systems or on the alternate media.
- On the switches, this command display information about the backup of the root file system **(/)** and directories **/altroot**,**/config**, **/var**, and **/var/tmp**, which are located either on an external USB flash drive or in internal flash memory.

  To back up software, use the `request system snapshot` command.

Options

none	Display information about the backup software.
all-members | local | member *member-id*	(EX Series switch Virtual Chassis only) (Optional) Display the snapshot in a Virtual Chassis:

- all-members—Display the snapshot for all members of the Virtual Chassis.
- local—Display the snapshot on the member of the Virtual Chassis that you are currently logged into.
- member *member-id*—Display the snapshot for the specified member of the Virtual Chassis.

media (external | internal) (EX Series switch only) (Optional) Display the destination media location for the snapshot. The `external` option specifies the snapshot on an external mass storage device, such as a USB flash drive. The `internal` option specifies the snapshot on an internal memory source, such as internal flash memory. If no additional options are specified, the command displays the snapshot stored in both slices.

Required Privilege Level

view

Output Fields

Table 1676 on page 27071 lists the output fields for the `show system snapshot` command. Output fields are listed in the approximate order in which they appear.

Table 1676: show system snapshot Output Fields

Field Name	Field Description
`Creation date`	Date and time of the last snapshot.
`JUNOS version on snapshot`	Junos OS release number of individual software packages.

Sample Output

show system snapshot (Router)

```
user@host> show system snapshot
Information for snapshot on hard-disk
Creation date: Oct 5 13:53:29 2005
```

```
JUNOS version on snapshot:
  jbase  : 7.3R2.5
  jcrypto: 7.3R2.5
  jdocs  : 7.3R2.5
  jkernel: 7.3R2.5
  jpfe   : M40-7.3R2.5
  jroute : 7.3R2.5
```

show system snapshot media external (Switch)

```
user@switch> show system snapshot media external
Information for snapshot on       external (/dev/da1s1a) (backup)
Creation date: Mar 19 03:37:18 2012
JUNOS version on snapshot:
  jbase  : ex-12.1I20120111_0048_user
  jcrypto-ex: 12.1I20120111_0048_user
  jdocs-ex: 12.1I20120111_0048_user
  jroute-ex: 12.1I20120111_0048_user
  jswitch-ex: 12.1I20120111_0048_user
  jweb-ex: 12.1I20120111_0048_user
Information for snapshot on       external (/dev/da1s2a) (primary)
Creation date: Mar 19 03:38:25 2012
JUNOS version on snapshot:
  jbase  : ex-12.2I20120305_2240_user
  jcrypto-ex: 12.2I20120305_2240_user
  jdocs-ex: 12.2I20120305_2240_user
  jroute-ex: 12.2I20120305_2240_user
  jswitch-ex: 12.2I20120305_2240_user
  jweb-ex: 12.2I20120305_2240_user
```

show system snapshot media internal (Switch)

```
user@switch> show system snapshot media internal
Information for snapshot on internal (/dev/da0s1a) (backup)
Creation date: Mar 14 05:01:02 2011
JUNOS version on snapshot:
  jbase  : 11.1R1.9
  jcrypto-ex: 11.1R1.9
  jdocs-ex: 11.1R1.9
  jkernel-ex: 11.1R1.9
```

```
  jroute-ex: 11.1R1.9
  jswitch-ex: 11.1R1.9
  jweb-ex: 11.1R1.9
  jpfe-ex42x: 11.1R1.9
Information for snapshot on internal (/dev/da0s2a) (primary)
Creation date: Mar 30 08:46:27 2011
JUNOS version on snapshot:
  jbase  : 11.2-20110330.0
  jcrypto-ex: 11.2-20110330.0
  jdocs-ex: 11.2-20110330.0
  jkernel-ex: 11.2-20110330.0
  jroute-ex: 11.2-20110330.0
  jswitch-ex: 11.2-20110330.0
  jweb-ex: 11.2-20110330.0
  jpfe-ex42x: 11.2-20110330.0
```

Release Information

Command introduced in Junos OS Release 7.6.

Option `slice` deprecated for Junos OS with Upgraded FreeBSD in Junos OS Release 15.1. You can find which platforms run Junos OS with Upgraded FreeBSD here: Release Information for Junos OS with Upgraded FreeBSD.

RELATED DOCUMENTATION

request system snapshot (Junos OS with FreeBSD Prior to Release 10) | **15190**

show system snapshot (Junos OS Evolved)

IN THIS SECTION

- Syntax | **27074**
- Description | **27074**
- Options | **27074**
- Required Privilege Level | **27074**

Output Fields | 27074
Sample Output | 27074

Syntax

```
show system snapshot
```

Description

This command displays information about the backup software—the contents of the **/soft** directory, which includes the running version of Junos OS Evolved. When nodes are synchronized in a cluster, this command shows what versions are available on all nodes in the cluster that contain persistent storage.

To back up the software, use the `request system snapshot` command.

Options

There are no options for this command.

Required Privilege Level

view

Output Fields

When you issue this command, you see a list of the snapshots available on each node.

Sample Output

show system snapshot (Junos OS Evolved)

```
user@host-re0> show system snapshot
-------------------------------
node: re0
-------------------------------
Current snapshot device: /dev/sdb
```

```
Snapshot boot device: sdb
List of installed version(s) in Snapshot boot device sdb:

[1]   < junos-evo-install-ptx-x86-64-20.4-202103151803.0-EVO - [2021-03-16 15:09:46]
[2]     junos-evo-install-ptx-x86-64-20.4-202103111254.0-EVO - [2021-03-16 15:10:32]
[3] ->  junos-evo-install-ptx-x86-64-20.4-202103150459.0-EVO - [2021-03-16 15:07:49]
[4]     junos-evo-install-ptx-x86-64-20.4-202103141559.0-EVO - [2021-03-16 15:11:52]

    '-' running version
    '>' next boot version after upgrade/downgrade
    '<' rollback boot version
------------------------------
node: re1
------------------------------
Current snapshot device: /dev/sdb
Snapshot boot device: sdb
List of installed version(s) in Snapshot boot device sdb:

[1] ->  junos-evo-install-ptx-x86-64-20.4-202103051234.0-EVO - [2021-03-05 01:10:31]

    '-' running version
    '>' next boot version after upgrade/downgrade
    '<' rollback boot version
```

RELATED DOCUMENTATION

request system snapshot (Junos OS Evolved) | **15187**

show system snapshot (Junos OS with Upgraded FreeBSD)

IN THIS SECTION

- Syntax | **27076**
- Description | **27076**
- Required Privilege Level | **27076**
- Output Fields | **27076**

- Sample Output | 27076
- Sample Output | 27077
- Release Information | 27077

Syntax

```
show system snapshot
```

Description

This command displays information about the non-recovery backup software, which is located in the **junos** file system on the hard disk drive or solid-state drive (SSD).

This command displays information about recovery snapshot after the non-recovery information.

To back up software, use the `request system snapshot` command.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. If there are no snapshots available, the command returns null output.

Sample Output

show system snapshot

```
user@host> show system snapshot
Snapshot snap.20141219.122106:
Location: /packages/sets/snap.20141219.122106
Creation date: Dec 19 12:21:06 2014
Junos version: 15.1-20141216_ib_15_1_psd.0
```

Sample Output

show system snapshot (updated as of Junos OS Release 17.2)

```
user@host> show system snapshot
Non-recovery snapshots:

Snapshot snap.20170112.105151:
Location: /packages/sets/snap.20170112.105151
Creation date: Jan 12 10:51:51 2017
Junos version: 17.2I20170112_0239

Snapshot snap.20170112.112307:
Location: /packages/sets/snap.20170112.112307
Creation date: Jan 12 11:23:07 2017
Junos version: 17.2I20170112_0239

Snapshot snap.20170112.112314:
Location: /packages/sets/snap.20170112.112314
Creation date: Jan 12 11:23:14 2017
Junos version: 17.2I20170112_0239

Total non-recovery snapshots:        3

Recovery Snapshots:
Snapshots available on the OAM volume:
recovery.ufs
Date created: Wed Jan 11 15:59:35 PST 2017
Junos version: 17.2I20170111_2242

Total recovery snapshots: 1
```

Release Information

Command introduced starting in Junos OS Release 15.1 for supported platforms. See Feature Explorer.

Output for recovery snapshots provided in Junos Release 17.2 for all the devices using Junos OS with upgraded FreeBSD.

RELATED DOCUMENTATION

request system snapshot (Junos OS with Upgraded FreeBSD) | 15203

request system reboot (Junos OS with Upgraded FreeBSD) | 15154

Release Information for Junos OS with Upgraded FreeBSD

show system snapshot media

IN THIS SECTION

- Syntax | 27078
- Description | 27078
- Options | 27078
- Required Privilege Level | 27079
- Output Fields | 27079
- Sample Output | 27079
- Release Information | 27081

Syntax

```
show system snapshot < media (compact-flash | external | harddisk | internal | usb) >
```

Description

This command displays information about the partitioning scheme present on the media. Information for only one root is displayed for single-root partitioning, whereas information for both roots is displayed for dual-root partitioning.

Options

- compact-flash— Show snapshot information from the CompactFlash card. (Supported on SRX5400, SRX5600, SRX5800)

- external— Show snapshot information from the external CompactFlash card. (Not supported on SRX5000 line devices)
- hard-disk— Show snapshot information from the Hard Disk. (Supported on SRX5400, SRX5600, SRX5800)
- internal— Show snapshot information from internal media. (Not supported on SRX5000 line devices)
- usb— Show snapshot information from device connected to USB port.

Required Privilege Level

View

Output Fields

Table 1677 on page 27079 lists the output fields for the `show system snapshot media` command. Output fields are listed in the approximate order in which they appear.

Table 1677: show system snapshot media Output Fields

Field Name	Field Description
`Creation date`	Date and time of the last snapshot.
`JUNOS version on snapshot`	Junos OS release number of individual software packages.

Sample Output

show system snapshot media compact-flash

```
show system snapshot media compact-flash
Information for snapshot on compact-flash (ad0s1)
Creation date: Aug 21 11:58:14 2017
JUNOS version on snapshot:
  junos  : 12.3X48-D40.5-domestic
```

show system snapshot media external

```
show system snapshot media external
Information for snapshot on      external (/dev/da1s2a) (primary)
Creation date: Apr 9 09:41:16 2018
JUNOS version on snapshot:
  junos  : 12.3X48-D40.5-domestic
Information for snapshot on       external (/dev/da1s1a) (backup)
Creation date: Apr 9 09:41:16 2018
JUNOS version on snapshot:
  junos  : 12.3X48-D40.5-domestic
```

show system snapshot media internal

```
show system snapshot media internal
Information for snapshot on       internal (/dev/da0s1a) (primary)
Creation date: Jan 15 10:43:26 2010
JUNOS version on snapshot:
  junos  : 10.1B3-domestic
Information for snapshot on       internal (/dev/da0s2a) (backup)
Creation date: Jan 15 10:15:32 2010
JUNOS version on snapshot:
  junos  : 10.2-20100112.0-domestic
```

show system snapshot media usb

```
show system snapshot media usb
Information for snapshot on usb (da0s1)
Creation date: Apr 9 08:44:46 2018
JUNOS version on snapshot:
  junos  : 12.3X48-D40.5-domestic
```

show system snapshot media hard-disk

```
show system snapshot media hard-disk
Information for snapshot on hard-disk (ad2s1)
Creation date: Apr 9 16:40:18 2018
```

```
JUNOS version on snapshot:
  junos  : 12.3X48-D40.5-domestic
```

Release Information

Command introduced in Junos OS Release 10.2 .

RELATED DOCUMENTATION

Creating a Snapshot and Using It to Boot an SRX Series device

show system software

IN THIS SECTION

- Syntax | 27082
- Syntax (EX Series Switches) | 27082
- Syntax (TX Matrix Router) | 27082
- Syntax (TX Matrix Plus Router) | 27082
- Syntax (QFX Series) | 27082
- Syntax (Junos OS Evolved) | 27083
- Description | 27083
- Options | 27083
- Required Privilege Level | 27084
- Output Fields | 27084
- Sample Output | 27085
- Release Information | 27086

Syntax

```
show system software
<detail>
```

Syntax (EX Series Switches)

```
show system software
<all-members>
<detail>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system software
<all-chassis | all-lcc | lcc number | scc>
<detail>
```

Syntax (TX Matrix Plus Router)

```
show system software
<all-chassis | all-lcc | lcc number | sfc number>
<detail>
```

Syntax (QFX Series)

```
show system software
<detail>
<infrastructure name>
<interconnect-device name>
<node-group name>
```

Syntax (Junos OS Evolved)

```
show system software
<add-restart>
<external-packages>
<list>
```

Description

Display the Junos OS extensions loaded on your router or switch.

Options

none	Display standard information about all loaded Junos OS extensions.
add-restart	(Junos OS Evolved only) (Optional) Display all console messages from the last in-service software upgrade (ISSU).
all-chassis	(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system software information for all the T640 routers (TX Matrix Router) or all the routers (TX Matrix Plus Router) in the chassis.
all-lcc	(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system software information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system software information for all connected T1600 or T4000 LCCs.
all-members	(EX4200 switches only) (Optional) Display the system software running on all members of the Virtual Chassis configuration.
detail	(Optional) Display detailed information about available Junos OS extensions.
external-packages	(Junos OS Evolved only) (Optional) Display installed third-party packages and the applications inside each package.
infrastructure *name*	(QFabric systems only) (Optional) Display the system software running on the fabric control Routing Engine and the fabric manager Routing Engine.
interconnect-device *name*	(QFabric systems only) (Optional) Display the system software running on the Interconnect device.

lcc *number*

(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system software information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system software information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(EX4200 switches only) (Optional) Display the system software running on the local Virtual Chassis member.

member *member-id*

(EX4200 switches only) (Optional) Display the system software running on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

node-group *name*

(QFabric systems only) (Optional) Display the system software running on the Node group.

scc

(Routing matrix only) (Optional) Display the system software running on a TX Matrix router (or switch-card chassis).

sfc

(TX Matrix Plus routers only) (Optional) Display system software information for the TX Matrix Plus router.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided a list of Junos OS packages installed on the router and their corresponding Junos OS release number.

Sample Output

show system software

```
user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]


Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]
Information for jdocs:

Comment:
JUNOS Online Documentation [7.2R1.7]


Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]


Information for jpfe:

Comment:
JUNOS Packet Forwarding Engine Support (M20/M40) [7.2R1.7]


Information for jroute:

Comment:
JUNOS Routing Software Suite [7.2R1.7]


Information for junos:
```

```
Comment:
JUNOS Base OS boot [7.2R1.7]
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system software add-restart (Junos OS Evolved)

IN THIS SECTION

- Syntax | 27086
- Description | 27086
- Required Privilege Level | 27087
- Output Fields | 27087
- Sample Output | 27087

Syntax

```
show system software
<add-restart>
```

Description

Display all console messages from the last in-service software upgrade (ISSU).

Required Privilege Level

maintenance

Output Fields

When you enter this command, the output shows a list of console messages from the last in-service software upgrade (ISSU).

Sample Output

show system software add-restart (no Log File Found)

```
user@host> show system software add-restart
-------------------------------
node: re0
-------------------------------
Software ISSU upgrade log file not found
```

show system software add-restart (with Log Information)

```
user@host-re0> show system software add-restart
-------------------------------
node: re0
-------------------------------
2021-04-10 22:31:26 Adding software images. This process can take several minutes. Please be
patient...
2021-04-10 22:31:26 Download and Validate in Progress
2021-04-10 22:33:09 re0: Starting upgrade : /var/tmp/sharat/snapshot/ptx-J4.iso
2021-04-10 22:33:09 re0: Upgrade version : junos-evo-install-ptx-x86-64-21.1I20210410155013-EVO-
J4
2021-04-10 22:33:30 re0: Running prechecks for JSU image...
2021-04-10 22:33:30 re0: JSU base release 21.1I20210410155013-EVO matches with Current Release:
21.1I20210410155013-EVO.
2021-04-10 22:33:30 re0: Starting JSU package installation...
2021-04-10 22:33:30 re0: Copying Files from base version junos-evo-install-ptx-
x86-64-21.1I20210410155013-EVO-J3 ...
2021-04-10 22:33:32 re0: Copying Files from JSU image...
2021-04-10 22:33:34 re0: Signing files...
2021-04-10 22:33:36 re0: Running pre-checks for 'junos-evo-install-ptx-
```

```
x86-64-21.1I20210410155013-EVO-J4'
2021-04-10 22:33:58 re0: Pre-checks pass successfully, copying files to software area
2021-04-10 22:34:26 re0: Running post install commands...
2021-04-10 22:34:33 re0: Post install sequence was successful.
2021-04-10 22:34:33 re0: Validating existing configs. See /var/log/validation_config.log for
config validation logs.
2021-04-10 22:35:22 re0: Validation Passed
2021-04-10 22:35:22 re0: Going ahead with Installation
2021-04-10 22:35:43 re0: Boot version is now 'junos-evo-install-ptx-x86-64-21.1I20210410155013-
EVO-J4'
2021-04-10 22:36:00 re1: Running pre-checks for 'junos-evo-install-ptx-
x86-64-21.1I20210410155013-EVO-J4'
2021-04-10 22:36:04 re1: Pre-checks pass successfully, copying files to software area
2021-04-10 22:36:43 re1: Running post install commands...
2021-04-10 22:36:51 re1: Post install sequence was successful.
2021-04-10 22:37:10 re1: Boot version is now 'junos-evo-install-ptx-x86-64-21.1I20210410155013-
EVO-J4'
2021-04-10 22:37:23 re0: Updating all nodes...
2021-04-10 22:38:23 re0: Boot version is now 'junos-evo-install-ptx-x86-64-21.1I20210410155013-
EVO-J3'
2021-04-10 22:38:23 re0: Removing version junos-evo-install-ptx-x86-64-21.1I20210410155013-EVO-
J4...
2021-04-10 22:38:24 re0: One or more remote node(s) failed updating
2021-04-10 22:38:38 re1: Boot version is now 'junos-evo-install-ptx-x86-64-21.1I20210410155013-
EVO-J3'
2021-04-10 22:38:38 re1: Removing version junos-evo-install-ptx-x86-64-21.1I20210410155013-EVO-
J4...
2021-04-10 22:38:51 ERROR: Upgrade failed on current RE. Node:re0 Image: re0:/var/tmp/user/
snapshot/ptx-J4.iso
2021-04-10 22:38:51 Validation and download failed. Aborting upgrade.
```

show system software list

IN THIS SECTION

- Syntax | 27089
- Description | 27089

- Required Privilege Level | 27089
- Output Fields | 27089
- Sample Output | 27090
- Release Information | 27092

Syntax

```
show system software list
```

Description

This command displays all the software versions, including `jfirmware` package versions, in the persistent storage on the Routing Engines in the system and displays the current software version running on the FPCs. FPCs cannot store more than one version, because FPCs do not contain any persistent storage media.

Required Privilege Level

view

Output Fields

For a description of the output fields, see Table 1678 on page 27089. Output fields are listed in the approximate order in which they appear.

Table 1678: show system software list Output Fields

Field Name	Description
node	Name of the node.

Table 1678: show system software list Output Fields *(Continued)*

Field Name	Description
`List of installed version(s)`	Ordered list of software that is or has been installed on the node: • - indicates the running software version. • > indicates the next boot software version, which occurs only after an upgrade or a downgrade. If no upgrade or downgrade has been performed, the > symbol will not appear in the list of installed versions. • < indicates the rollback boot software version if there is one.
`External Software`	List of running third party packages.

Sample Output

show system software list

```
user@host-re0> show system software list
-------------------------------
node: re0
-------------------------------
Active boot device is primary: /dev/vda
List of installed version(s) :

    '-' running version
    '>' next boot version after upgrade/downgrade
    '<' rollback boot version
    '*' deleted JSU version

 >   junos-evo-install-ptx-fixed-x86-64-22.4I20230106124155 - [2023-01-06 12:47:59]
       JET app custom_logger 1.0.1
 -   junos-evo-install-ptx-fixed-x86-64-22.4I20230106121117 - [2023-01-06 12:33:31]
       JET app custom_logger 1.0.1
 <   junos-evo-install-ptx-fixed-x86-64-22.4I20230106120732 - [2023-01-06 12:15:25]
       JET app custom_logger 1.0.1
     junos-evo-install-ptx-fixed-x86-64-22.4I20230106120459 - [2023-01-06 12:08:05]
       JET app custom_logger 1.0.1
```

```
     junos-evo-install-ptx-fixed-x86-64-22.4I20230106105842-EVO - [2023-01-06 11:41:50]
       JET app custom_logger 1.0.1
       JET app custom_multi_app 1.0.1
       JET app multi_app 1.1.1
```

show system software list, with `jfirmware` package information

```
user@host-re0> show system software list
-------------------------------
node: re0
-------------------------------
Active boot device is primary: /dev/vda
List of installed version(s) :

    '-' running version
    '>' next boot version after upgrade/downgrade
    '<' rollback boot version
    '*' deleted JSU version

 -   junos-evo-install-ptx-x86-64-24.4 I202410114083306-EVO - [2024-10-11 09:14:39]
       jfirmware-junos-evo-install-ptx-x86-64-24.4I202410114083306-EVO 1.0.1
 <   junos-evo-install-ptx-x86-64-22.4I20221214083306 - [2022-12-14 08:58:24]

-------------------------------
node: re1
-------------------------------
Active boot device is primary: /dev/vda
List of installed version(s) :

    '-' running version
    '>' next boot version after upgrade/downgrade
    '<' rollback boot version
    '*' deleted JSU version

 -   junos-evo-install-ptx-x86-64-24.4 I202410114083306-EVO - [2024-10-11 09:16:48]
       jfirmware-junos-evo-install-ptx-x86-64-24.4I202410114083306-EVO 1.0.1
 <   junos-evo-install-ptx-x86-64-22.4I20221214083306 - [2022-12-14 08:58:59]
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

Support for `jfirmware` packages introduced in Junos OS Evolved Release 24.4R1.

show system software recovery-package

IN THIS SECTION

- Syntax | 27092
- Description | 27092
- Options | 27092
- Required Privilege Level | 27093
- Output Fields | 27093
- Sample Output | 27093
- Release Information | 27094

Syntax

```
show system software recovery-package
<backup>
<master>
```

Description

Display the software recovery packages loaded on an XRE200 External Routing Engine in an EX8200 Virtual Chassis.

Options

An EX8200 Virtual Chassis supports up to two external Routing Engines for the purpose of external Routing Engine redundancy. If you have configured two external Routing Engines, one functions as the

primary external Routing Engine and the other functions as the backup. If you have not configured two external Routing Engines, the `backup` and `master` options are not relevant.

none Display the software recovery packages available on the external Routing Engine.

backup (Optional) Display the software recovery packages available on the backup external Routing Engine.

master (Optional) Display the software recovery packages available on the primary external Routing Engine.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided a list of software recovery packages, including the software release number, installed on the external Routing Engine.

Sample Output

show system software recovery-package

```
user@external-routing-engine> show system software recovery-package

Member 1 :
---------------------------------------------------------------------------
1. jinstall-ex-xre200-11.1I20110608_1356_pravasp-domestic.tgz
Member 0 :
---------------------------------------------------------------------------
1. jinstall-ex-xre200-11.4I20110608_1356_pravasp-domestic.tgz
```

show system software recovery-package backup

```
user@external-routing-engine> show system software recovery-package backup

Member 0 :
```

```
---------------------------------------------------------------------------
1. jinstall-ex-xre200-11.4I20110608_1356_pravasp-domestic.tgz
```

show system software recovery-package master

```
user@external-routing-engine> show system software recovery-package master

Member 1 :
---------------------------------------------------------------------------
1. jinstall-ex-xre200-11.1I20110608_1356_pravasp-domestic.tgz
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

request system software recovery-package | 15334

show system software restore-point-status

IN THIS SECTION

- Syntax | 27095
- Description | 27095
- Required Privilege Level | 27095
- Output Fields | 27095
- Sample Output | 27096
- Release Information | 27096

Syntax

```
show system software restore-point-status
```

Description

This command displays the status of the restore-point for QFabric system. A restore-point contains both a snapshot of the software and a configuration file.

You can create only one restore-point. When you create a new restore-point, the existing restore-point, if available, is erased.

Required Privilege Level

view

Output Fields

Table 1679 on page 27095 lists the output fields for the `show system software restore-point status` command. Output fields are listed in the approximate order in which they appear.

Table 1679: show system software restore-point status Output Fields

Field Name	Field Description
`Member`	Name of the Director device.
`Creation Time`	Time when the restore-point was created.
`Status`	Status of restore-point creation.
`Restore volume`	Name and path to restore volume used to create the restore-point.

Sample Output

show system software restore-point status

```
user@qfabric> show system software restore-point status
Member        Creation Time            Status        Restore volume
------        -------------            -----        ---------------
dg0        Aug 15 07:42:39 2014        success        /dev/VolGroup00/LogVol03
dg1        Aug 15 07:42:27 2014        success        /dev/VolGroup00/LogVol03
```

Release Information

Command introduced in Junos OS Release 14.1X53-D15.

RELATED DOCUMENTATION

request system software restore-point | 15338

show system software upgrade status

IN THIS SECTION

- Syntax | 27097
- Description | 27097
- Required Privilege Level | 27097
- Output Fields | 27097
- Sample Output | 27098
- Release Information | 27098

Syntax

```
show system software upgrade status
```

Description

(QFabric systems only) Display the status of a software upgrade, including details for both nonstop software upgrades and *component all* style upgrades.

Required Privilege Level

view

Output Fields

Table 1680 on page 27097 lists the output fields for the `show system software upgrade status` command. Output fields are listed in the approximate order in which they appear.

Table 1680: show system software upgrade status Output Fields

Field Name	Field Description
`Timestamp`	Displays the day of the week, month, date, hour, minute, second, and year when you issue the `show system software upgrade status` command. An example of the timestamp format is as follows: Wed Jan 16 22:06:02 2013.
`Software nonstop upgrade on:`	Status of the upgrade: • `FM-0 in progress`—A Director group nonstop software upgrade is in process for the fabric manager Routing Engine. • `NW-NG-0 in progress`—A Node group nonstop software upgrade is in process for the network Node group. • `RSNG in progress`—A Node group nonstop software upgrade is in process for the redundant server Node group. • `all in progress`—A *component all* style upgrade is in process for the entire QFabric system.

Sample Output

show system software upgrade status (Nonstop Software Upgrade)

```
user@qfabric> show system software upgrade status
Wed Jan 16 22:06:02 2013 Software nonstop upgrade on:
                         NW-NG-0 in progress
                         RSNG in progress
```

show system software upgrade status (Component All Upgrade)

```
user@qfabric> show system software upgrade status
Wed Jan 16 22:37:48 2013 Software component upgrade on:
                         all in progress
```

Release Information

Command introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Performing a Nonstop Software Upgrade on the QFabric System

Verifying Nonstop Software Upgrade for QFabric Systems

Upgrading Software on a QFabric System

request system software nonstop-upgrade (QFabric Systems) | **15311**

show system software usb-software-version

IN THIS SECTION

- Syntax | **27099**
- Description | **27099**

- Additional Information | 27099
- Required Privilege Level | 27099
- Output Fields | 27100
- Sample Output | 27100
- Release Information | 27100

Syntax

```
show system software usb-software-version
```

Description

This comamnd displays the version of software present on a standard USB installer key attached to each Director Group (DG) device (QFabric systems only).

Additional Information

When issuing the **show system software usb-software-version** command, the USB installer key must be attached to the DGs.

The format of the USB installer key (including partitions) must conform to the standard specifications of the Juniper-provided USB installer.

A Juniper-provided or Juniper-recommended USB installer device must have the following partitions:

```
Device     Boot      Start         End      Blocks      Id   System
/dev/sdb1   *           1            75      295244      4    FAT16 <32M
/dev/sdb2               76          709     2496058   83  Linux
```

Required Privilege Level

view

Output Fields

Table 1681 on page 27100 lists the output fields for the `show system software usb-software-version` command. Output fields are listed in the approximate order in which they appear.

Table 1681: show system software usb-software-version Output Fields

Field Name	Field Description
`Node`	Node supporting the device.
`Device`	Device on which the software is present.
`Version`	Version of the software present.
`Filename`	Software filename.

Sample Output

show system software usb-software-version

```
user@host> show system software usb-software-version

  NODE     DEVICE
VERSION                                                FILENAME
  ----     ------
-------                                                --------
   dg0   /dev/sdb     14.1-20160516_x141X53_vjqfd.0     jinstall-
qfabric-14.1-20160516_x141X53_vjqfd.0.rpm
   dg1   /dev/sdb     14.1-20160516_x141X53_vjqfd.0     jinstall-
qfabric-14.1-20160516_x141X53_vjqfd.0.rpm
```

Release Information

Command introduced in Junos OS Release 14.1X53-D40.

RELATED DOCUMENTATION

Performing a Nonstop Software Upgrade on the QFabric System

Verifying Nonstop Software Upgrade for QFabric Systems

Upgrading Software on a QFabric System

show system statistics

IN THIS SECTION

- Syntax | 27101
- Syntax (EX Series Switches) | 27101
- Syntax (MX Series Router) | 27102
- Syntax (QFX Series) | 27102
- Description | 27102
- Options | 27102
- Required Privilege Level | 27103
- Sample Output | 27104
- Release Information | 27119

Syntax

```
show system statistics
```

Syntax (EX Series Switches)

```
show system statistics
<all-members>
<local>
<member member-id>
```

Syntax (MX Series Router)

```
show system statistics
<all-members>
<local>
<member member-id>
<extended <ipv4 | ipv6>>
```

Syntax (QFX Series)

```
show system statistics
```

Description

Display system-wide protocol-related statistics.

Options

The following are all possible options for the `show system statistics` command. The available options vary depending on the platform, software release, and operating system (Junos OS and Junos OS Evolved).

NOTE: Use the `show system statistics ?` command to determine the available options.

none Display system statistics for all the following protocols:

- `arp`—Address Resolution Protocol
- `backup` - Statistics of backup JunosVM or RE
- `bridge`—IEEE 802.1 Bridging
- `clns`—Connectionless Network Service
- `esis`—End System-to-Intermediate System
- `ethoamcfm`—Ethernet OAM protocol for connectivity fault management
- `ethoamlfm`—Ethernet OAM protocol for link fault management

- `extended`—System statistics for IPv4 and IPv6 traffic
- `icmp`—Internet Control Message Protocol
- `icmp6`—Internet Control Message Protocol version 6
- `igmp`—Internet Group Management Protocol
- `ip`—Internet Protocol version 4
- `ip6`—Internet Protocol version 6
- `jsr`—Juniper Socket Replication
- `jtd` - Show jtd statistics
- `mpls`—Multiprotocol Label Switching
- `rdp`—Reliable Datagram Protocol
- `tcp`—Transmission Control Protocol
- `tnp`—Trivial Network Protocol
- `ttp`—TNP Tunneling Protocol
- `tudp`—Trivial User Datagram Protocol
- `udp`—User Datagram Protocol
- `vpls`—Virtual Private LAN Service

all-members (EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for all members of the Virtual Chassis configuration.

local (EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for the local Virtual Chassis member.

member *member-id* (EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *`member-id`* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *`member-id`* with a value of 0 or 1.

Required Privilege Level

view

Sample Output

show system statistics

```
user@host> show system statistics
ip:
        3682087 total packets received
        0 bad header checksums
        0 with size smaller than minimum
        0 with data size < data length
        0 with header length < data size
        0 with data length < header length
        0 with incorrect version number
        0 packets destined to dead next hop
        0 fragments received
        0 fragments dropped (dup or out of space)
        0 fragments dropped (queue overflow)
        0 fragments dropped after timeout
        0 fragments dropped due to over limit
        0 packets reassembled ok
        3664774 packets for this host
        17316 packets for unknown/unsupported protocol
        0 packets forwarded
        0 packets not forwardable
        0 redirects sent
        6528 packets sent from this host
        0 packets sent with fabricated ip header
        0 output packets dropped due to no bufs
        0 output packets discarded due to no route
        0 output datagrams fragmented
        0 fragments created
        0 datagrams that can't be fragmented
        0 packets with bad options
        1123 packets with options handled without error
        0 strict source and record route options
        0 loose source and record route options
        0 record route options
        0 timestamp options
        0 timestamp and address options
        0 timestamp and prespecified address options
        0 option packets dropped due to rate limit
        1123 router alert options
```

```
        0 multicast packets dropped (no iflist)
        0 packets dropped (src and int don't match)
icmp:
        0 drops due to rate limit
        0 calls to icmp_error
        0 errors not generated because old message was icmp
        Output histogram:
                echo reply: 75
        0 messages with bad code fields
        0 messages less than the minimum length
        0 messages with bad checksum
        0 messages with bad source address
        0 messages with bad length
        0 echo drops with broadcast or multicast destination address
        0 timestamp drops with broadcast or multicast destination address
        Input histogram:
                echo: 75
                router advertisement: 130
        75 message responses generated
tcp:
        3844 packets sent
                3618 data packets (1055596 bytes)
                0 data packets (0 bytes) retransmitted
                0 resends initiated by MTU discovery
                205 ack-only packets (148 packets delayed)
                0 URG only packets
                0 window probe packets
                0 window update packets
                1079 control packets
        5815 packets received
                3377 acks (for 1055657 bytes)
                24 duplicate acks
                0 acks for unsent data
                2655 packets (15004 bytes) received in-sequence
                1 completely duplicate packet  (0 bytes)
                0 old duplicate packets
                0 packets with some dup. data (0 bytes duped)
                0 out-of-order packets (0 bytes)
                0 packets (0 bytes) of data after window
                0 window probes
                7 window update packets
                0 packets received after close
                0 discarded for bad checksums
```

```
        0 discarded for bad header offset fields
        0 discarded because packet too short
1 connection request
32 connection accepts
0 bad connection attempts
0 listen queue overflows
33 connections established (including accepts)
30 connections closed (including 0 drops)
        27 connections updated cached RTT on close
        27 connections updated cached RTT variance on close
        0 connections updated cached ssthresh on close
0 embryonic connections dropped
3374 segments updated rtt (of 3220 attempts)
0 retransmit timeouts
        0 connections dropped by rexmit timeout
0 persist timeouts
        0 connections dropped by persist timeout
344 keepalive timeouts
        0 keepalive probes sent
        0 connections dropped by keepalive
1096 correct ACK header predictions
1314 correct data packet header predictions
32 syncache entries added
        0 retransmitted
        0 dupsyn
        0 dropped
        32 completed
        0 bucket overflow
        0 cache overflow
        0 reset
        0 stale
        0 aborted
        0 badack
        0 unreach
        0 zone failures
0 cookies sent
0 cookies received
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
1058 RST packets
0 ICMP packets ignored by TCP
```

```
        0 send packets dropped by TCP due to auth errors
        0 rcv packets dropped by TCP due to auth errors
udp:
        3658884 datagrams received
        0 with incomplete header
        0 with bad data length field
        0 with bad checksum
        3657342 dropped due to no socket
        3657342 broadcast/multicast datagrams dropped due to no socket
        0 dropped due to full socket buffers
        0 not for hashed pcb
        4291311496 delivered
        1551 datagrams output
ipsec:
        0 inbound packets processed successfully
        0 inbound packets violated process security policy
        0 inbound packets with no SA available
        0 invalid inbound packets
        0 inbound packets failed due to insufficient memory
        0 inbound packets failed getting SPI
        0 inbound packets failed on AH replay check
        0 inbound packets failed on ESP replay check
        0 inbound AH packets considered authentic
        0 inbound AH packets failed on authentication
        0 inbound ESP packets considered authentic
        0 inbound ESP packets failed on authentication
        0 outbound packets processed successfully
        0 outbound packets violated process security policy
        0 outbound packets with no SA available
        0 invalid outbound packets
        0 outbound packets failed due to insufficient memory
        0 outbound packets with no route
igmp:
        17186 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent
arp:
```

```
        44181302 datagrams received
        2 ARP requests received
        2028 ARP replies received
        3156 resolution requests received
        0 unrestricted proxy requests
        0 received proxy requests
        0 proxy requests not proxied
        0 with bogus interface
        787 with incorrect length
        712 for non-IP protocol
        0 with unsupported op code
        0 with bad protocol address length
        0 with bad hardware address length
        0 with multicast source address
        7611 with multicast target address
        0 with my own hardware address
        14241699 for an address not on the interface
        0 with a broadcast source address
        0 with source address duplicate to mine
        29929250 which were not for me
        0 packets discarded waiting for resolution
        6 packets sent after waiting for resolution
        17812 ARP requests sent
        2 ARP replies sent
        0 requests for memory denied
        0 requests dropped on entry
        0 requests dropped during retry
ip6:
        0 total packets received
        0 with size smaller than minimum
        0 with data size < data length
        0 with bad options
        0 with incorrect version number
        0 fragments received
        0 fragments dropped (dup or out of space)
        0 fragments dropped after timeout
        0 fragments that exceeded limit
        0 packets reassembled ok
        0 packets for this host
        0 packets forwarded
        0 packets not forwardable
        0 redirects sent
        0 packets sent from this host
```

```
        0 packets sent with fabricated ip header
        0 output packets dropped due to no bufs, etc.
        0 output packets discarded due to no route
        0 output datagrams fragmented
        0 fragments created
        0 datagrams that can't be fragmented
        0 packets that violated scope rules
        0 multicast packets which we don't join
        Mbuf statistics:
        0 packets whose headers are not continuous
        0 tunneling packets that can't find gif
        0 packets discarded due to too may headers
        0 failures of source address selection
        0 forward cache hit
        0 forward cache miss
        0 packets destined to dead next hop
        0 option packets dropped due to rate limit
        0 packets dropped (src and int don't match)
        0 packets dropped due to bad protocol
icmp6:
        0 calls to icmp_error
        0 errors not generated because old message was icmp error or so
        0 errors not generated because rate limitation
        0 messages with bad code fields
        0 messages < minimum length
        0 bad checksums
        0 messages with bad length
        Histogram of error messages to be generated:
                0 no route
                0 administratively prohibited
                0 beyond scope
                0 address unreachable
                0 port unreachable
                0 packet too big
                0 time exceed transit
                0 time exceed reassembly
                0 erroneous header field
                0 unrecognized next header
                0 unrecognized option
                0 redirect
                0 unknown
        0 message responses generated
        0 messages with too many ND options
```

```
ipsec6:
        0 inbound packets processed successfully
        0 inbound packets violated process security policy
        0 inbound packets with no SA available
        0 invalid inbound packets
        0 inbound packets failed due to insufficient memory
        0 inbound packets failed getting SPI
        0 inbound packets failed on AH replay check
        0 inbound packets failed on ESP replay check
        0 inbound AH packets considered authentic
        0 inbound AH packets failed on authentication
        0 inbound ESP packets considered authentic
        0 inbound ESP packets failed on authentication
        0 outbound packets processed successfully
        0 outbound packets violated process security policy
        0 outbound packets with no SA available
        0 invalid outbound packets
        0 outbound packets failed due to insufficient memory
        0 outbound packets with no route
clnl:
        0 total packets received
        0 packets delivered
        0 too small
        0 bad header length
        0 bad checksum
        0 bad version
        0 unknown or unsupported protocol
        0 bogus sdl size
        0 no free memory in socket buffer
        0 send packets discarded
        0 sbappend failure
        0 mcopy failure
        0 address fields were not reasonable
        0 segment information forgotten
        0 forwarded packets
        0 total packets sent
        0 output packets discarded
        0 non-forwarded packets
        0 packets fragmented
        0 fragments sent
        0 fragments discarded
        0 fragments timed out
        0 fragmentation prohibited
```

```
        0 packets reconstructed
        0 packets destined to dead nexthop
        0 packets discarded due to no route
        0 Error pdu rate drops
        0 ER pdu generation failure
esis:
        0 total pkts received
        0 total packets consumed by protocol
        0 pdus received with bad checksum
        0 pdus received with bad version number
        0 pdus received with bad type field
        0 short pdus received
        0 bogus sdl size
        0 bad header length
        0 unknown or unsupported protocol
        0 no free memory in socket buffer
        0 send packets discarded
        0 sbappend failure
        0 mcopy failure
        0 ISO family not configured
tnp:
        146776365 unicast packets received
        0 broadcast packets received
        0 fragmented packets received
        0 hello packets dropped
        0 fragments dropped
        0 fragment reassembly queue flushes
        0 hello packets received
        0 control packets received
        49681642 rdp packets received
        337175 udp packets received
        96757548 tunnel packets received
        0 input packets discarded with no protocol
        98397591 unicast packets sent
        0 broadcast packets sent
        0 fragmented packets sent
        0 hello packets dropped
        0 fragments dropped
        0 hello packets sent
        0 control packets sent
        49681642 rdp packets sent
        337175 udp packets sent
        48378774 tunnel packets sent
```

```
        0 packets sent with unknown protocol
rdp:
        49681642 input packets
        0 discards for bad checksum
        0 discards bad sequence number
        0 refused connections
        2031964 acks received
        0 dropped due to full socket buffers
        49692 retransmits
        49681642 output packets
        24815968 acks sent
        28 connects
        0 closes
        22783990 keepalives received
        22783990 keepalives sent
tudp:
        337175 datagrams received
        0 with incomplete header
        0 with bad data length field
        0 with bad checksum
        0 dropped due to no socket
        0 broadcast/multicast datagrams dropped due to no socket
        0 dropped due to full socket buffers
        337175 delivered
        337175 datagrams output
ttp:
        398749 packets sent
        0 packets sent while unconnected
        0 packets sent while interface down
        0 packets sent couldn't get buffer
        0 packets sent couldn't find neighbor
        44696687 L2 packets received
        0 unknown L3 packets received
        3682087 IPv4 L3 packets received
        0 MPLS L3 packets received
        0 MPLS->IPv4 L3 packets received
        0 IPv4->MPLS L3 packets received
        0 VPLS L3 packets received
        0 IPv6 L3 packets received
        0 ARP L3 packets received
        0 CLNP L3 packets received
        0 TNP L3 packets received
        0 NULL L3 packets received
```

```
        0 cyclotron cycle L3 packets received
        0 cyclotron send L3 packets received
        0 packets received while unconnected
        0 packets received from unknown ifl
        0 input packets couldn't get buffer
        0 input packets with bad type
        0 input packets with discard type
        0 Input packets with bad tlv header
        70633 Input packets with bad tlv type
        68877 Input packets dropped based on tlv result
        0 input packets for which rt lookup is bypassed
mpls:
        0 total mpls packets received
        0 packets forwarded
        0 packets dropped
        0 with header too small
        0 after tagging, can't fit link MTU
        0 with IPv4 explicit NULL tag
        0 with IPv4 explicit NULL cksum errors
        0 with router alert tag
        0 lsp ping packets (ttl-expired/router alert)
        0 with ttl expired
        0 with tag encoding error
        0 packets discarded, no route
jsr:
     Handle-inf:o
         0 Handles in use
         0 Handles allocated so far
         0 Handles freed so far
         0 Handles in delayed free state
     IHA:
         0 IHA invalid subtype messages
         0 IHA invalid length messages
         0 IHA invalid version messages
         0 IHA too short messages
         0 IHA invalid dst handle messages
         0 IHA invalid src handle messages
         0 IHA unmatched src handle messages
         0 IHA invalid messages for primary
         0 IHA invalid messages for secondary
         0 IHA invalid messages for current state
         0 IHA messages sent for subtype init
         0 IHA messages rcvd for subytpe init
```

```
        0 IHA messages sent for subtype init
        0 IHA messages rcvd for subytpe init
        0 IHA messages sent for subtype init
        0 IHA messages rcvd for subytpe init
        0 IHA messages sent for subtype init
        0 IHA messages rcvd for subytpe init
        0 IHA messages sent for subtype init
        0 IHA messages rcvd for subytpe init
        0 IHA message timeouts
        0 IHA socket unreplicate messages
    SDRL:
        0 SDRL socket teardowns
        0 SDRL socket teardown failures
        0 SDRL socket unreplicates
        0 SDRL socket unreplicate failures
        0 SDRL external timeouts
        0 SDRL internal timeouts
        0 SDRL ipc messages sent
        0 SDRL ipc send failures
        0 SDRL ipc messages recvd
        0 SDRL ipc messages recvd
        0 SDRL primary replication messages sent
        0 SDRL primary replication message send failures
        0 SDRL primary ack messages received
        0 SDRL primary ack message receive failures
        0 SDRL primary sock replication inits
        0 SDRL primary sock replication init failures
        0 SDRL primary throttle remove messages
        0 SDRL primary throttle remove failures
        0 SDRL primary init handshake messages
        0 SDRL primary init handshake failures
        0 SDRL secondary replication messages received
        0 SDRL secondary replication message receive failures
        0 SDRL secondary replication acks sent
        0 SDRL secondary replication ack send failures
        0 SDRL secondary sock splits
        0 SDRL secondary sock split failures
        0 SDRL secondary sock merges
        0 SDRL secondary sock merge failures
        0 SDRL secondary sockets closed
        0 SDRL secondary rcv snoop fd close failures
        0 SDRL secondary snd snoop fd close failures
        0 SDRL secondary init handshake messages
```

```
        0 SDRL secondary init handshake failures
    PRL:
        0 PRL packets enqueued
        0 PRL packets failed to enqueue
        0 PRL packets dequeued
        0 PRL packets failed to dequeue
        0 PRL queue entry allocations
        0 PRL queue entry frees
        0 calls to layer 4 input handlers
        0 failed calls to layer 4 input handlers
        0 PRL queue drains
        0 PRL replication timeouts
        0 PRL replication messages sent
        0 PRL replication message send failures
        0 PRL acknowledgment messages sent
        0 PRL acknowledgement message send failures
        0 PRL replication messages received
        0 PRL replication message receive failures
        0 PRL acknowledgement messages received
        0 PRL acknowledgement receive failures
        0 PRL messages with bad IPC type
        0 PRL messages with no handler
        2 PRL global state initializations
        1 PRL global state cleanups
        0 PRL per-socket state creations
        0 PRL per-socket state creation failures
        0 PRL per-socket state cleanups
        0 PRL socket closes
        0 PRL socket merges
        0 PRL socket unreplicates
        0 PRL primary socket replication initializations
        0 PRL secondary socket replication initializations
        0 PRL primary socket replication activations
        0 PRL secondary socket replication activations
        0 packets received from peers
        0 PRL packets receive operations from peer failed
        0 PRL buffer pullup failures
        0 new pkts dropped on secondary socket
    PSRM:
        0 PSRM replication timeouts
        0 PSRM replication messages sent
        0 PSRM replication message send failures
        0 PSRM acknowledgment messages sent
```

```
          0 PSRM acknowledgement message send failures
          0 PSRM flow control messages sent
          0 PSRM flow control message send failures
          0 PSRM replication messages received
          0 PSRM replication message receive failures
          0 PSRM acknowledgment messages received
          0 PSRM acknowledgment message receive failures
          0 PSRM flow control messages received
          0 PSRM flow control message receive failures
          0 SRM messages with bad IPC type
          0 PSRM messages with no handler
          2 PSRM global state initializations
          1 PSRM global state cleanups
          0 PSRM per-socket state creations
          0 PSRM per-socket state creation failures
          0 PSRM per-socket state cleanups
          0 PSRM socket closes
          0 PSRM socket merges
          0 PSRM socket unreplicates
          0 PSRM primary socket replication initializations
          0 psrm-secondary-socket-replication-initializations
          0 PSRM primary socket replication activations
          0 secondary socket replication activations
          0 PSRM tcpcb updates
          0 PSRM buffer pullup failures
          73 PSRM tcp timestamp msg recv counters
          0 PSRM tcp timestamp msg recv failures
          0 PSRM tcp timestamp msg send counters
          0 PSRM tcp timestamp msg send failures
     TCP:
          0 TCP out-of-order packets on JSR sockets
vpls:
        0 total packets received
        0 with size smaller than minimum
        0 with incorrect version number
        0 packets for this host
        0 packets with no logical interface
        0 packets with no family
        0 packets with no route table
        0 packets with no auxiliary table
        0 packets with no corefacing entry
        0 packets with no CE-facing entry
        0 mac route learning requests
```

```
        0 mac routes learnt
        0 requests to learn an existing route
        0 learning requests while learning disabled on interface
        0 learning requests over capacity
        0 mac routes moved
        0 requests to move static route
        0 mac route aging requests
        0 mac routes aged
        0 bogus address in aging requests
        0 requests to age static route
        0 requests to re-ageout aged route
        0 requests involving multiple peer FEs
        0 aging acks from PFE
        0 aging non-acks from PFE
        0 aging requests timed out waiting on FEs
        0 aging requests over max-rate
        0 errors finding peer FEs
Sctp:
         195 sinput-packets
         195 datagrams
         9 packets that had data
         9 input SACK chunks
         9 input DATA chunks
         0 duplicate DATA chunks
         27 input HB chunks
         27 HB-ACK chunks
         0 input ECNE chunks
         0 input AUTH chunks
         0 chunks missing AUTH
         0 invalid HMAC ids received
         0 invalid secret ids received
         0 auth failed
         7 fast path receives all one chunk
         0 fast path multi-part data
         209 output packets
         9 output SACKs
         9 output DATA chunks
         0 retransmitted DATA chunks
         0 fast retransmitted DATA chunks
         0 FR's that happened more than once to same chunk
         153 output HB chunks
         0 output ECNE chunks
         0 output AUTH chunks
```

```
        9 ip_output error counter
        0 from middle box
        0 from end host
        0 with data
        0 non-data, non-endhost
        0 non-endhost, bandwidth rep only
        0 not enough for chunk header
        0 not enough data to confirm
        0 where process_chunk_drop said break
        0 failed to find TSN
        0 attempt reverse TSN lookup
        0 e-host confirms zero-rwnd
        0 midbox confirms no space
        0 data did not match TSN
        0 TSN's marked for Fast Retran
        0 iterator timers fired
        1 T3 data time outs
        0 window probe (T3) timers fired
        0 INIT timers fired
        7 sack timers fired
        0 shutdown timers fired
        157 heartbeat timers fired
        0 a cookie timeout fired
        0 an endpoint changed its cookiesecret
        0 PMTU timers fired
        0 shutdown ack timers fired
        0 shutdown guard timers fired
        0 stream reset timers fired
        0 early FR timers fired
        0 an asconf timer fired
        0 auto close timer fired
        0 asoc free timers expired
        1 inp free timers expired
        0 packet shorter than header
        0 checksum error
        0 no endpoint for port
        0 bad v-tag
        0 bad SID
        0 no memory
        0 number of multiple FR in a RTT window
        9 RFC813 allowed sending
        0 RFC813 does not allow sending
        0 times max burst prohibited sending
```

```
        0 look ahead tells us no memory in interface
        0 numbers of window probes sent
        0 times an output error to clamp down on next user send
        1 times sctp_senderrors were caused from a user
        0 number of in data drops due to chunk limit reached
        0 number of in data drops due to rwnd limit reached
        0 times a ECN reduced the cwnd
        195 used express lookup via vtag
        0 collision in express lookup
        0 times the sender ran dry of user data on primary
        0 sacks the slow way
        0 window update only sacks sent
        0 sends with sinfo_flags !=0
        0 unordered sends
        0 sends with EOF flag set
        0 sends with ABORT flag set
        188 times protocol drain called
        0 times we did a protocol drain
        0 times recv was called with peek
        197 cached chunks used
        0 cached stream oq's used
        0 unread messages abandonded by close
        0 send burst avoidance, already max burst inflight to net
        0 send cwnd full avoidance, already max burst inflight to net
        0 number of map array over-runs via fwd-tsn's
```

Release Information

Command introduced before JUNOS Release 7.4.

show system statistics arp

IN THIS SECTION

- Syntax | 27120
- Syntax (EX Series Switches) | 27120
- Syntax (TX Matrix Router) | 27120

- Syntax (TX Matrix Plus Router) | 27120
- Description | 27121
- Options | 27121
- Additional Information | 27122
- Required Privilege Level | 27122
- Sample Output | 27122
- Release Information | 27129

Syntax

```
show system statistics arp
```

Syntax (EX Series Switches)

```
show system statistics arp
<all-members>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics arp
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics arp
<all-chassis | all-lcc | lcc number / sfc number>
```

Description

Display system-wide Address Resolution Protocol (ARP) statistics.

Options

none Display system-wide ARP statistics.

all-chassis (Optional) Display ARP statistics for all the routers in the chassis.

all-lcc (Optional) On a TX Matrix router, display system-wide ARP statistics for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system-wide ARP statistics for all routers connected to the TX Matrix Plus router

all-members (Optional) Display ARP statistics for all members of the Virtual Chassis configuration.

lcc *number* (Optional) On a TX Matrix router, display ARP statistics for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display ARP statistics for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (Optional) Display ARP statistics for the local Virtual Chassis member.

member *member-id* (Optional) Display ARP statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

scc (Optional) Display ARP statistics for the TX Matrix router (or switch-card chassis).

sfc *number* (Optional) Display ARP statistics for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system statistics arp` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics arp

```
user@host> show system statistics arp
arp:
         184710 datagrams received
         2886 ARP requests received
         684 ARP replies received
         0 resolution request  received
         0 unrestricted proxy requests
         0 restricted proxy requests
         0 received proxy requests
         0 unrestricted proxy requests not proxied
         0 restricted proxy requests not proxied
         0 datagrams with bogus interface
         0 datagrams with incorrect length
         0 datagrams for non-IP protocol
         0 datagrams with unsupported op code
         0 datagrams with bad protocol address length
         0 datagrams with bad hardware address length
         0 datagrams with multicast source address
         0 datagrams with multicast target address
         0 datagrams with my own hardware address
         0 datagrams for an address not on the interface
         0 datagrams with a broadcast source address
         0 datagrams with source address duplicate to mine
         181140 datagrams which were not for me
         0 packets discarded waiting for resolution
```

```
        4 packets sent after waiting for resolution
        703 ARP requests sent
        2886 ARP replies sent
        0 requests for memory denied
        0 requests dropped on entry
        0 requests dropped during retry
        0 requests dropped due to interface deletion
        0 requests on unnumbered interfaces
        0 new requests on unnumbered interfaces
        0 replies for from unnumbered interfaces
        0 requests on unnumbered interface with non-subnetted donor
        0 replies from unnumbered interface with non-subnetted donor
```

show system statistics arp (EX Series Switches)

```
user@host> show system statistics arp
arp:
        186423 datagrams received
        88 ARP requests received
        88 ARP replies received
        0 resolution request  received
        0 unrestricted proxy requests
        0 restricted proxy requests
        0 received proxy requests
        0 proxy requests not proxied
        0 restricted proxy requests not proxied
        0 datagrams with bogus interface
        0 datagrams with incorrect length
        0 datagrams for non-IP protocol
        0 datagrams with unsupported op code
        0 datagrams with bad protocol address length
        0 datagrams with bad hardware address length
        0 datagrams with multicast source address
        0 datagrams with multicast source address
        0 datagrams with my own hardware address
        164 datagrams for an address not on the interface
        0 datagrams with a broadcast source address
        0 datagrams with source address duplicate to mine
        186075 datagrams which were not for me
        0 packets discarded waiting for resolution
        0 packets sent after waiting for resolution
```

```
          50 ARP requests sent
          88 ARP replies sent
          0 requests for memory denied
          0 requests dropped on entry
          0 requests dropped during retry
          0 requests dropped due to interface deletion
          0 requests on unnumbered interfaces
          0 new requests on unnumbered interfaces
          0 replies for from unnumbered interfaces
          0 requests on unnumbered interface with non-subnetted donor
          0 replies from unnumbered interface with non-subnetted donor
```

show system statistics arp (TX Matrix Plus Router)

```
user@host> show system statistics arp
sfc0-re0:
--------------------------------------------------------------------------
arp:
        487 datagrams received
        8 ARP requests received
        438 ARP replys received
        438 resolution requests received
        0 unrestricted proxy requests
        0 restricted proxy requests
        0 received proxy requests
        0 proxy requestss not proxied
        0 restricted-proxy requestss not proxied
        0 with bogus interface
        0 with incorrect length
        0 for non-IP protocol
        0 with unsupported op code
        0 with bad protocol address length
        0 with bad hardware address length
        0 with multicast source address
        0 with multicast target address
        0 with my own hardware address
        0 for an address not on the interface
        0 with a broadcast source address
        0 with source address duplicate to mine
        41 which were not for me
        0 packets discarded waiting for resolution
```

```
        438 packets sent after waiting for resolution
        1282 ARP requests sent
        8 ARP replys sent
        0 requests for memory denied
        0 requests dropped on entry
        0 requests dropped during retry
        0 requests dropped due to interface deletion
        0 requests on unnumbered interfaces
        0 new requests on unnumbered interfaces
        0 replies for from unnumbered interfaces
        0 requests on unnumbered interface with non-subnetted donor
        0 replies from unnumbered interface with non-subnetted donor

lcc0-re0:
--------------------------------------------------------------------------
arp:
        19 datagrams received
        0 ARP requests received
        1 ARP reply  received
        0 resolution requests received
        0 unrestricted proxy requests
        0 restricted proxy requests
        0 received proxy requests
        0 proxy requestss not proxied
        0 restricted-proxy requestss not proxied
        0 with bogus interface
        0 with incorrect length
        0 for non-IP protocol
        0 with unsupported op code
        0 with bad protocol address length
        0 with bad hardware address length
        0 with multicast source address
        0 with multicast target address
        0 with my own hardware address
        0 for an address not on the interface
        0 with a broadcast source address
        0 with source address duplicate to mine
        18 which were not for me
        0 packets discarded waiting for resolution
        0 packets sent after waiting for resolution
        8 ARP requests sent
        0 ARP replys sent
        0 requests for memory denied
```

```
        0 requests dropped on entry
        0 requests dropped during retry
        0 requests dropped due to interface deletion
        0 requests on unnumbered interfaces
        0 new requests on unnumbered interfaces
        0 replies for from unnumbered interfaces
        0 requests on unnumbered interface with non-subnetted donor
        0 replies from unnumbered interface with non-subnetted donor

lcc1-re0:
--------------------------------------------------------------------------
arp:
        17 datagrams received
        0 ARP requests received
        1 ARP reply  received
        0 resolution requests received
        0 unrestricted proxy requests
        0 restricted proxy requests
        0 received proxy requests
        0 proxy requestss not proxied
        0 restricted-proxy requestss not proxied
        0 with bogus interface
        0 with incorrect length
        0 for non-IP protocol
        0 with unsupported op code
        0 with bad protocol address length
        0 with bad hardware address length
        0 with multicast source address
        0 with multicast target address
        0 with my own hardware address
        0 for an address not on the interface
        0 with a broadcast source address
        0 with source address duplicate to mine
        16 which were not for me
        0 packets discarded waiting for resolution
        0 packets sent after waiting for resolution
        9 ARP requests sent
        0 ARP replys sent
        0 requests for memory denied
        0 requests dropped on entry
        0 requests dropped during retry
        0 requests dropped due to interface deletion
        0 requests on unnumbered interfaces
```

```
        0 new requests on unnumbered interfaces
        0 replies for from unnumbered interfaces
        0 requests on unnumbered interface with non-subnetted donor
        0 replies from unnumbered interface with non-subnetted donor

lcc2-re0:
--------------------------------------------------------------------------
arp:
        18 datagrams received
        1 ARP request  received
        1 ARP reply  received
        0 resolution requests received
        0 unrestricted proxy requests
        0 restricted proxy requests
        0 received proxy requests
        0 proxy requestss not proxied
        0 restricted-proxy requestss not proxied
        0 with bogus interface
        0 with incorrect length
        0 for non-IP protocol
        0 with unsupported op code
        0 with bad protocol address length
        0 with bad hardware address length
        0 with multicast source address
        0 with multicast target address
        0 with my own hardware address
        0 for an address not on the interface
        0 with a broadcast source address
        0 with source address duplicate to mine
        16 which were not for me
        0 packets discarded waiting for resolution
        0 packets sent after waiting for resolution
        9 ARP requests sent
        1 ARP reply  sent
        0 requests for memory denied
        0 requests dropped on entry
        0 requests dropped during retry
        0 requests dropped due to interface deletion
        0 requests on unnumbered interfaces
        0 new requests on unnumbered interfaces
        0 replies for from unnumbered interfaces
        0 requests on unnumbered interface with non-subnetted donor
        0 replies from unnumbered interface with non-subnetted donor
```

```
lcc3-re0:
--------------------------------------------------------------------------
arp:
        13 datagrams received
        0 ARP requests received
        1 ARP reply  received
        0 resolution requests received
        0 unrestricted proxy requests
        0 restricted proxy requests
        0 received proxy requests
        0 proxy requestss not proxied
        0 restricted-proxy requestss not proxied
        0 with bogus interface
        0 with incorrect length
        0 for non-IP protocol
        0 with unsupported op code
        0 with bad protocol address length
        0 with bad hardware address length
        0 with multicast source address
        0 with multicast target address
        0 with my own hardware address
        0 for an address not on the interface
        0 with a broadcast source address
        0 with source address duplicate to mine
        12 which were not for me
        0 packets discarded waiting for resolution
        0 packets sent after waiting for resolution
        8 ARP requests sent
        0 ARP replys sent
        0 requests for memory denied
        0 requests dropped on entry
        0 requests dropped during retry
        0 requests dropped due to interface deletion
        0 requests on unnumbered interfaces
        0 new requests on unnumbered interfaces
        0 replies for from unnumbered interfaces
        0 requests on unnumbered interface with non-subnetted donor
        0 replies from unnumbered interface with non-subnetted donor
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Routing Matrix with a TX Matrix Plus Router Solutions Page
- Example: Configuring Proxy ARP on an EX Series Switch
- Verifying That Proxy ARP Is Working Correctly

show system statistics icmp

IN THIS SECTION

- Syntax | 27129
- Syntax (EX Series Switches) | 27130
- Syntax (TX Matrix Router) | 27130
- Syntax (TX Matrix Plus Router) | 27130
- Description | 27130
- Options | 27130
- Additional Information | 27131
- Required Privilege Level | 27131
- Sample Output | 27132
- Release Information | 27136

Syntax

```
show system statistics icmp
```

Syntax (EX Series Switches)

```
show system statistics icmp
<all-members>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics icmp
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics icmp
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide Internet Control Message Protocol (ICMP) statistics.

Options

none Display system statistics for ICMP.

all-chassis (Optional) Display system statistics for ICMP for all the routers in the chassis.

all-lcc (Optional) On a TX Matrix router, display system statistics for ICMP for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for ICMP for all connected T1600 or T4000 LCCs.

all-members (Optional) Display ICMP statistics for all members of the Virtual Chassis configuration.

lcc *number* (Optional) On a TX Matrix router, display system statistics for ICMP for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for ICMP for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (Optional) Display ICMP statistics for the local Virtual Chassis member.

member ***member-id*** (Optional) Display ICMP statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

scc (Optional) Display system statistics for ICMP for the TX Matrix router (or switch-card chassis).

sfc ***number*** (Optional) Display system statistics for ICMP for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system statistics icmp` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics icmp

```
user@host> show system statistics icmp
icmp:
         16783 drops due to rate limit
         9998 calls to icmp_error
         0 errors not generated because old message was icmp
         Output Histogram
                 38877 echo reply
                 1 destination unreachable
                 1 routing redirect
                 163 echo
                 5000 time exceeded
                 4996 parameter problem
                 5000 time stamp reply
         0 messages with bad code fields
         0 messages less than the minimum length
         0 messages with bad checksum
         0 messages with bad source address
         20000 messages with bad length
         0 echo drops with broadcast or multicast destinaton address
         0 timestamp drops with broadcast or multicast destination address
         Input Histogram
                 5093 echo reply
                 5000 destination unreachable
                 5000 source quench
                 5000 routing redirect
                 5000 alternate host address
                 38877 echo
                 5000 router advertisement
                 5000 router solicitation
                 5000 time exceeded
                 5000 parameter problem
                 5000 time stamp
                 5000 time stamp reply
                 5000 information request
                 5000 information request reply
                 5000 address mask request
                 5000 address mask reply
                 5000 traceroute
```

```
                 5000 data conversion error
                 5000 mobile host redirect
                 5000 IPv6 where-are-you
                 5000 IPv6 i-am-here
                 5000 mobile registration request
                 5000 mobile registration reply
                 5000 skip
                 5000 photuris
         43877 message responses generated
```

show system statistics icmp (EX Series Switches)

```
user@host> show system statistics icmp
icmp:
         0 drops due to rate limit
         12 calls to icmp_error
         0 errors not generated because old message was icmp
         Output histogram:
                 297 echo reply
                 12 destination unreachable
         0 messages with bad code fields
         0 messages less than the minimum length
         0 messages with bad checksum
         0 messages with bad source address
         0 messages with bad length
         0 echo drops with broadcast or multicast destinaton address
         0 timestamp drops with broadcast or multicast destination address
         Input histogram:
                 297 echo
         297 message responses generated
```

show system statistics icmp (TX Matrix Plus Router)

```
user@host> show system statistics icmp
sfc0-re0:
--------------------------------------------------------------------------
icmp:
        0 drops due to rate limit
        0 calls to icmp_error
        0 errors not generated because old message was icmp
```

```
        Output histogram:
                echo reply: 21
        0 messages with bad code fields
        0 messages less than the minimum length
        0 messages with bad checksum
        0 messages with bad source address
        0 messages with bad length
        0 echo drops with broadcast or multicast destinaton address
        0 timestamp drops with broadcast or multicast destination address
        Input histogram:
                echo: 21
        21 message responses generated

lcc0-re0:
--------------------------------------------------------------------------
icmp:
        0 drops due to rate limit
        1 call  to icmp_error
        0 errors not generated because old message was icmp
        Output histogram:
                echo reply: 24
                destination unreachable: 1
        0 messages with bad code fields
        0 messages less than the minimum length
        0 messages with bad checksum
        0 messages with bad source address
        0 messages with bad length
        0 echo drops with broadcast or multicast destinaton address
        0 timestamp drops with broadcast or multicast destination address
        Input histogram:
                echo: 24
        24 message responses generated

lcc1-re0:
--------------------------------------------------------------------------
icmp:
        0 drops due to rate limit
        0 calls to icmp_error
        0 errors not generated because old message was icmp
        Output histogram:
                echo reply: 23
        0 messages with bad code fields
        0 messages less than the minimum length
```

```
        0 messages with bad checksum
        0 messages with bad source address
        0 messages with bad length
        0 echo drops with broadcast or multicast destinaton address
        0 timestamp drops with broadcast or multicast destination address
        Input histogram:
                echo: 23
        23 message responses generated

lcc2-re0:
--------------------------------------------------------------------------
icmp:
        0 drops due to rate limit
        0 calls to icmp_error
        0 errors not generated because old message was icmp
        Output histogram:
                echo reply: 22
        0 messages with bad code fields
        0 messages less than the minimum length
        0 messages with bad checksum
        0 messages with bad source address
        0 messages with bad length
        0 echo drops with broadcast or multicast destinaton address
        0 timestamp drops with broadcast or multicast destination address
        Input histogram:
                echo: 22
        22 message responses generated

lcc3-re0:
--------------------------------------------------------------------------
icmp:
        0 drops due to rate limit
        0 calls to icmp_error
        0 errors not generated because old message was icmp
        Output histogram:
                echo reply: 22
        0 messages with bad code fields
        0 messages less than the minimum length
        0 messages with bad checksum
        0 messages with bad source address
        0 messages with bad length
        0 echo drops with broadcast or multicast destinaton address
        0 timestamp drops with broadcast or multicast destination address
```

```
        Input histogram:
                echo: 22
        22 message responses generated
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system statistics icmp6

IN THIS SECTION

- Syntax (EX Series Switches) | 27137
- Syntax (MX Series Routers, SRX Series) | 27137
- Syntax (TX Matrix Router) | 27137
- Syntax (TX Matrix Plus Router) | 27137
- Description | 27137
- Options | 27137
- Additional Information | 27138
- Required Privilege Level | 27138
- Sample Output | 27139
- Sample Output | 27142
- Release Information | 27146

Syntax (EX Series Switches)

```
show system statistics icmp6
<all-members>
<local>
<member member-id>
```

Syntax (MX Series Routers, SRX Series)

```
show system statistics icmp6
```

Syntax (TX Matrix Router)

```
show system statistics icmp6
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics icmp6
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide Internet Control Message Protocol for IPv6 (ICMPv6) statistics.

Options

none Display system statistics for ICMPv6.

all-chassis (Optional) Display system statistics for ICMPv6 for all the routers in the chassis.

all-lcc (Optional) On a TX Matrix router, display system statistics for ICMPv6 for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for ICMPv6 for all connected T1600 or T4000 LCCs.

all-members (Optional) Display ICMPv6 statistics for all members of the Virtual Chassis configuration.

lcc *number* (Optional) On a TX Matrix router, display system statistics for ICMPv6 for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for ICMPv6 for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (Optional) Display ICMPv6 statistics for the local Virtual Chassis member.

member *member-id* (Optional) Display ICMPv6 statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

scc (Optional) Display system statistics for ICMPv6 for the TX Matrix router (or switch-card chassis).

sfc *number* (Optional) Display system statistics for ICMPv6 for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system statistics icmp6` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics icmp6 (MX Series Routers)

```
user@host> show system statistics icmp6
icmp6:
         79 Calls to icmp_error
         0 Errors not generated because old message was icmp error
         0 Errors not generated because rate limitation
         Output histogram:
                 79 unreach
                 30 echo
                 163 multicast listener query
                 6 multicast listener report
                 940 neighbor solicitation
                 694184 neighbor advertisement
         0 Messages with bad code fields
         0 Messages < minimum length
         0 Bad checksums
         0 Messages with bad length
         Input histogram:
                 10 echo reply
                 6 multicast listener report
                 693975 neighbor solicitation
         Histogram of error messages to be generated:
                 0 No route
                 0 Administratively prohibited
                 0 Beyond scope
                 79 Address unreachable
                 0 Port unreachable
                 0 Time exceed transit
                 0 Time exceed reassembly
                 0 Erroneous header field
                 0 Unrecognized next header
                 0 Unrecognized option
                 0 Unknown
         0 Message responses generated
         0 Messages with too many ND options
         100000 Max System ND nh cache limit
         79840 Max Public ND nh cache limit
         200 Max IRI ND nh cache limit
         19960 Max Management intf ND nh cache limit
```

```
        79840 Current Public ND nexthops present
        4 Current IRI ND nexthops present
        0 Current Management ND nexthops present
        909266 Total ND nexthops creation failed as limit reached
        909266 Public ND nexthops creation failed as public limit reached
        0 IRI ND nexthops creation failed as iri limit reached
        0 Management ND nexthops creation failed as mgt limit reached
```

show system statistics icmp6 (EX Series Switches)

```
user@host> show system statistics icmp6
icmp6:
        0 Calls to icmp_error
        0 Errors not generated because old message was icmp error
        0 Errors not generated because rate limitation
        0 Messages with bad code fields
        0 Messages < minimum length
        0 Bad checksums
        0 Messages with bad length
                0 No route
                0 Administratively prohibited
                0 Beyond scope
                0 Address unreachable
                0 Port unreachable
                0 packet too big
                0 Time exceed transit
                0 Time exceed reassembly
                0 Erroneous header field
                0 Unrecognized next header
                0 Unrecognized option
                0 redirect
                0 Unknown
        0 Message responses generated
        0 Messages with too many ND options
```

show system statistics icmp6 (SRX Series and vSRX Virtual Firewall 3.0)

```
0 Calls to icmp_error
        0 Errors not generated because old message was icmp error
        0 Errors not generated because rate limitation
```

```
        0 Messages with bad code fields
        0 Messages < minimum length
        0 Bad checksums
        0 Messages with bad length
        Histogram of error messages to be generated:
                0 No route
                0 Administratively prohibited
                0 Beyond scope
                0 Address unreachable
                0 Port unreachable
                0 Time exceed transit
                0 Time exceed reassembly
                0 Erroneous header field
                0 Unrecognized next header
                0 Unrecognized option
                0 Unknown
        0 Message responses generated
        0 Messages with too many ND options
        100000 Max System ND nh cache limit
        79840 Max Public ND nh cache limit
        200 Max IRI ND nh cache limit
        19960 Max Management intf ND nh cache limit
        0 Current Public ND nexthops present
        0 Current IRI ND nexthops present
        0 Current Management ND nexthops present
        0 Total ND nexthops creation failed as limit reached
        0 Public ND nexthops creation failed as public limit reached
        0 IRI ND nexthops creation failed as iri limit reached
        0 Management ND nexthops creation failed as mgt limit reached
        0 interface-restricted ndp proxy requests
        0 interface-restricted dad proxy requests
        0 interface-restricted ndp proxy responses
        0 interface-restricted dad proxy conflicts
        0 interface-restricted dad proxy duplicates
        0 interface-restricted ndp proxy resolve requests
        0 interface-restricted dad proxy resolve requests
        0 interface-restricted dad packets from same node dropped
        0 interface-restricted proxy packets dropped with nomac
        0 interface-unrestricted ndp proxy requests
        0 interface-unrestricted dad proxy requests
        0 interface-unrestricted ndp proxy responses
        0 interface-unrestricted dad proxy conflicts
        0 interface-unrestricted dad proxy duplicates
```

```
          0 interface-unrestricted ndp proxy resolve requests
          0 interface-unrestricted dad proxy resolve requests
          0 interface-unrestricted dad packets from same port dropped
          0 interface-unrestricted proxy packets dropped with nomac
          0 ND hold nexthops dropped on entry by RED mark
          0 ND hold nexthops dropped on timer expire by RED mark
```

Sample Output

show system statistics icmp6 (TX Matrix Plus Router)

```
user@host> show system statistics icmp6
sfc0-re0:
--------------------------------------------------------------------------
icmp6:
        0 calls to icmp_error
        0 errors not generated because old message was icmp error or so
        0 errors not generated because rate limitation
        Output histogram:
                neighbor solicitation: 12
                neighbor advertisment: 4
        0 messages with bad code fields
        0 messages < minimum length
        0 bad checksums
        0 messages with bad length
        Histogram of error messages to be generated:
                0 no route
                0 administratively prohibited
                0 beyond scope
                0 address unreachable
                0 port unreachable
                0 packet too big
                0 time exceed transit
                0 time exceed reassembly
                0 erroneous header field
                0 unrecognized next header
                0 unrecognized option
                0 redirect
                0 unknown
        0 message responses generated
        0 messages with too many ND options
```

```
lcc0-re0:
--------------------------------------------------------------------------
icmp6:
        0 calls to icmp_error
        0 errors not generated because old message was icmp error or so
        0 errors not generated because rate limitation
        Output histogram:
                neighbor solicitation: 12
                neighbor advertisment: 4
        0 messages with bad code fields
        0 messages < minimum length
        0 bad checksums
        0 messages with bad length
        Histogram of error messages to be generated:
                0 no route
                0 administratively prohibited
                0 beyond scope
                0 address unreachable
                0 port unreachable
                0 packet too big
                0 time exceed transit
                0 time exceed reassembly
                0 erroneous header field
                0 unrecognized next header
                0 unrecognized option
                0 redirect
                0 unknown
        0 message responses generated
        0 messages with too many ND options

lcc1-re0:
--------------------------------------------------------------------------
icmp6:
        0 calls to icmp_error
        0 errors not generated because old message was icmp error or so
        0 errors not generated because rate limitation
        Output histogram:
                neighbor solicitation: 12
                neighbor advertisment: 4
        0 messages with bad code fields
        0 messages < minimum length
        0 bad checksums
```

```
        0 messages with bad length
        Input histogram:
                neighbor advertisment: 2
        Histogram of error messages to be generated:
                0 no route
                0 administratively prohibited
                0 beyond scope
                0 address unreachable
                0 port unreachable
                0 packet too big
                0 time exceed transit
                0 time exceed reassembly
                0 erroneous header field
                0 unrecognized next header
                0 unrecognized option
                0 redirect
                0 unknown
        0 message responses generated
        0 messages with too many ND options

lcc2-re0:
--------------------------------------------------------------------------
icmp6:
        0 calls to icmp_error
        0 errors not generated because old message was icmp error or so
        0 errors not generated because rate limitation
        Output histogram:
                neighbor solicitation: 12
                neighbor advertisment: 4
        0 messages with bad code fields
        0 messages < minimum length
        0 bad checksums
        0 messages with bad length
        Input histogram:
                neighbor advertisment: 2
        Histogram of error messages to be generated:
                0 no route
                0 administratively prohibited
                0 beyond scope
                0 address unreachable
                0 port unreachable
                0 packet too big
                0 time exceed transit
```

```
                0 time exceed reassembly
                0 erroneous header field
                0 unrecognized next header
                0 unrecognized option
                0 redirect
                0 unknown
        0 message responses generated
        0 messages with too many ND options

lcc3-re0:
--------------------------------------------------------------------------
icmp6:
        0 calls to icmp_error
        0 errors not generated because old message was icmp error or so
        0 errors not generated because rate limitation
        Output histogram:
                neighbor solicitation: 12
                neighbor advertisment: 4
        0 messages with bad code fields
        0 messages < minimum length
        0 bad checksums
        0 messages with bad length
        Input histogram:
                neighbor advertisment: 2
        Histogram of error messages to be generated:
                0 no route
                0 administratively prohibited
                0 beyond scope
                0 address unreachable
                0 port unreachable
                0 packet too big
                0 time exceed transit
                0 time exceed reassembly
                0 erroneous header field
                0 unrecognized next header
                0 unrecognized option
                0 redirect
                0 unknown
        0 message responses generated
        0 messages with too many ND options
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system statistics igmp

IN THIS SECTION

- Syntax | 27146
- Syntax (EX Series Switches) | 27147
- Syntax (TX Matrix Router) | 27147
- Syntax (TX Matrix Plus Router) | 27147
- Description | 27147
- Options | 27147
- Additional Information | 27148
- Required Privilege Level | 27148
- Sample Output | 27149
- Release Information | 27151

Syntax

```
show system statistics igmp
```

Syntax (EX Series Switches)

```
show system statistics igmp
<all-members>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics igmp
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics igmp
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide Internet Group Management Protocol (IGMP) statistics.

Options

none Display system statistics for IGMP.

all-chassis (Optional) Display system statistics for IGMP for all the routers in the chassis.

all-lcc (Optional) On a TX Matrix router, display system statistics for IGMP for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for all connected T1600 or T4000 LCCs.

all-members (Optional) Display IGMP statistics for all members of the Virtual Chassis configuration.

lcc *number* (Optional) On a TX Matrix router, display system statistics for IGMP for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (Optional) Display IGMP statistics for the local Virtual Chassis member.

member ***member-id*** (Optional) Display IGMP statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

scc (Optional) Display system statistics for IGMP for the TX Matrix router (or switch-card chassis).

sfc ***number*** (Optional) Display system statistics for IGMP for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system statistics igmp` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics igmp

```
user@host> show system statistics igmp
igmp:
        17178 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent
```

show system statistics igmp (EX Series Switches)

```
user@host> show system statistics igmp
igmp:
         0 messages received
         0 messages received with too few bytes
         0 messages received with bad checksum
         0 membership queries received
         0 membership queries received with invalid fields
         0 membership reports received
         0 membership reports received with invalid fields
         0 membership reports received for groups to which we belong
         0 Membership reports sent
```

show system statistics igmp (TX Matrix Plus Router)

```
user@host> show system statistics igmp
sfc0-re0:
--------------------------------------------------------------------------
igmp:
        0 messages received
        0 messages received with too few bytes
```

```
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent

lcc0-re0:
--------------------------------------------------------------------------
igmp:
        0 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent

lcc1-re0:
--------------------------------------------------------------------------
igmp:
        0 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent

lcc2-re0:
--------------------------------------------------------------------------
igmp:
        0 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
```

```
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent

lcc3-re0:
--------------------------------------------------------------------------
igmp:
        0 messages received
        0 messages received with too few bytes
        0 messages received with bad checksum
        0 membership queries received
        0 membership queries received with invalid field(s)
        0 membership reports received
        0 membership reports received with invalid field(s)
        0 membership reports received for groups to which we belong
        0 membership reports sent
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system statistics ip

IN THIS SECTION

- Syntax | 27152
- Syntax (EX Series Switches) | 27152
- Syntax (TX Matrix Router) | 27152
- Syntax (TX Matrix Plus Router) | 27152
- Description | 27152

- Options | 27153
- Additional Information | 27154
- Required Privilege Level | 27154
- Sample Output | 27154
- Release Information | 27162

Syntax

```
show system statistics ip
```

Syntax (EX Series Switches)

```
show system statistics ip
<all-members>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics ip
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics ip
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide IPv4 statistics.

Some of the statistics include, the total number of IPv4 packets received, number of packets destined to dead next hop, number of fragments received and dropped, number of multicast packets dropped, number of Time-triggered Protocol over IP (TTPoIP) packets sent, received, and dropped, and number of raw packets dropped.

Options

none Display system statistics for IPv4.

all-chassis (Optional) Display system statistics for IPv4 for all the routers in the chassis.

all-lcc (Optional) On a TX Matrix router, display system statistics for IPv4 for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IPv4 for all T1600 or T4000 routers connected to the TX Matrix Plus router.

all-members (Optional) Display IPv4 statistics for all members of the Virtual Chassis configuration.

lcc *number* (Optional) On a TX Matrix router, display system statistics for IPv4 for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IPv4 for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (Optional) Display IPv4 statistics for the local Virtual Chassis member.

member *member-id* (Optional) Display IPv4 statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

scc (Optional) Display system statistics for IPv4 for the TX Matrix router (or switch-card chassis).

sfc *number* (Optional) Display system statistics for IPv4 for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system statistics ip` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics ip

```
user@host> show system statistics ip
ip:
         1752658 total packets received
         0 bad header checksums
         0 with size smaller than minimum
         0 with data size < data length
         0 with header length < data size
         0 with data length < header length
         0 with incorrect version number
         0 packets destined to dead next hop
         0 fragments received
         0 fragments dropped (dup or out of space)
         0 fragment sessions dropped (queue overflow)
         0 fragments dropped after timeout
         0 packets reassembled ok
         1709456 packets for this host
         10494 packets for unknown/unsupported protocol
         546 packets forwarded
         0 packets not forwardable
         546 redirects sent
         1340179 packets sent from this host
```

```
          0 packets sent with fabricated ip header
          0 output packets dropped due to no bufs
          0 output packets discarded due to no route
          0 output datagrams fragmented
          0 fragments created
          0 datagrams that can't be fragmented
          0 packets with bad options
          10494 packets with options handled without error
          0 strict source and record route options
          0 loose source and record route options
          0 record route options
          0 timestamp options
          0 timestamp and address options
          0 timestamp and prespecified address options
          0 option packets dropped due to rate limit
          10494 router alert options
          0 multicast packets dropped (no iflist)
          0 packets dropped (src and int don't match)
          0 transit re packets dropped on mgmt i/f
          0 packets used first nexthop in ecmp unilist
          0 incoming ttpoip packets received
          0 incoming ttpoip packets dropped
          0 outgoing TTPoIP packets sent
          0 outgoing TTPoIP packets dropped
          0 raw packets dropped. no space in socket recv buffer
```

show system statistics ip (EX Series Switches)

```
user@host> show system statistics ip
ip:
          74121 total packets received
          0 bad header checksums
          0 with size smaller than minimum
          0 with data size < data length
          0 with header length < data size
          0 with data length < header length
          0 with incorrect version number
          0 packets destined to dead next hop
          0 fragments received
          0 fragments dropped (dup or out of space)
          0 fragment sessions dropped (queue overflow)
```

```
          0 fragments dropped after timeout
          0 packets reassembled ok
          1134061 packets for this host
          0 packets for unknown/unsupported protocol
          40177 packets forwarded
          0 packets not forwardable
          40177 redirects sent
          1122558 packets sent from this host
          0 packets sent with fabricated ip header
          0 output packets dropped due to no bufs
          0 output packets discarded due to no route
          0 output datagrams fragmented
          0 fragments created
          0 datagrams that can't be fragmented
          0 packets with bad options
          0 packets with options handled without error
          0 strict source and record route options
          0 loose source and record route options
          0 record route options
          0 timestamp options
          0 timestamp and address options
          0 timestamp and prespecified address options
          0 option packets dropped due to rate limit
          0 router alert options
          0 multicast packets dropped (no iflist)
          0 packets dropped (src and int don't match)
          0 transit re packets dropped on mgmt i/f
          0 packets used first nexthop in ecmp unilist
          0 incoming ttpoip packets received
          0 incoming ttpoip packets dropped
          0 outgoing TTPoIP packets sent
          0 outgoing TTPoIP packets dropped
```

show system statistics ip (TX Matrix Plus Router)

```
user@host> show system statistics ip
sfc0-re0:
--------------------------------------------------------------------------
ip:
         47695035 total packets received
         0 bad header checksums
```

```
        0 with size smaller than minimum
        0 with data size < data length
        0 with header length < data size
        0 with data length < header length
        0 with incorrect version number
        0 packets destined to dead next hop
        42350 fragments received
        0 fragments dropped (dup or out of space)
        0 fragment sessions dropped (queue overflow)
        0 fragments dropped after timeout
        21175 packets reassembled ok
        47674941 packets for this host
        146 packets for unknown/unsupported protocol
        0 packets forwarded
        0 packets not forwardable
        0 redirects sent
        61304579 packets sent from this host
        8496 packets sent with fabricated ip header
        0 output packets dropped due to no bufs
        0 output packets discarded due to no route
        6746344 output datagrams fragmented
        0 fragments created
        0 datagrams that can't be fragmented
        0 packets with bad options
        2400 packets with options handled without error
        0 strict source and record route options
        0 loose source and record route options
        0 record route options
        0 timestamp options
        0 timestamp and address options
        0 timestamp and prespecified address options
        0 option packets dropped due to rate limit
        2400 router alert options
        0 multicast packets dropped (no iflist)
        0 packets dropped (src and int don't match)
        0 transit re packets dropped on mgmt i/f
        0 packets used first nexthop in ecmp unilist
        12995412 incoming ttpoip packets received
        0 incoming ttpoip packets dropped
        16959177 outgoing TTPoIP packets sent
        0 outgoing TTPoIP packets dropped
        0 raw packets dropped. no space in socket recv buffer
```

```
lcc0-re0:
--------------------------------------------------------------------------
ip:
         12990061 total packets received
         0 bad header checksums
         0 with size smaller than minimum
         0 with data size < data length
         0 with header length < data size
         0 with data length < header length
         0 with incorrect version number
         0 packets destined to dead next hop
         0 fragments received
         0 fragments dropped (dup or out of space)
         0 fragment sessions dropped (queue overflow)
         0 fragments dropped after timeout
         0 packets reassembled ok
         12989979 packets for this host
         82 packets for unknown/unsupported protocol
         0 packets forwarded
         0 packets not forwardable
         0 redirects sent
         9318381 packets sent from this host
         0 packets sent with fabricated ip header
         0 output packets dropped due to no bufs
         0 output packets discarded due to no route
         3440 output datagrams fragmented
         0 fragments created
         0 datagrams that can't be fragmented
         0 packets with bad options
         82 packets with options handled without error
         0 strict source and record route options
         0 loose source and record route options
         0 record route options
         0 timestamp options
         0 timestamp and address options
         0 timestamp and prespecified address options
         0 option packets dropped due to rate limit
         82 router alert options
         0 multicast packets dropped (no iflist)
         0 packets dropped (src and int don't match)
         0 transit re packets dropped on mgmt i/f
         0 packets used first nexthop in ecmp unilist
         0 incoming ttpoip packets received
```

```
          0 incoming ttpoip packets dropped
          548071 outgoing TTPoIP packets sent
          0 outgoing TTPoIP packets dropped
          0 raw packets dropped. no space in socket recv buffer

lcc1-re0:
--------------------------------------------------------------------------
ip:
          12849723 total packets received
          0 bad header checksums
          0 with size smaller than minimum
          0 with data size < data length
          0 with header length < data size
          0 with data length < header length
          0 with incorrect version number
          0 packets destined to dead next hop
          0 fragments received
          0 fragments dropped (dup or out of space)
          0 fragment sessions dropped (queue overflow)
          0 fragments dropped after timeout
          0 packets reassembled ok
          12849641 packets for this host
          82 packets for unknown/unsupported protocol
          0 packets forwarded
          0 packets not forwardable
          0 redirects sent
          7676351 packets sent from this host
          0 packets sent with fabricated ip header
          0 output packets dropped due to no bufs
          0 output packets discarded due to no route
          0 output datagrams fragmented
          0 fragments created
          0 datagrams that can't be fragmented
          0 packets with bad options
          82 packets with options handled without error
          0 strict source and record route options
          0 loose source and record route options
          0 record route options
          0 timestamp options
          0 timestamp and address options
          0 timestamp and prespecified address options
          0 option packets dropped due to rate limit
          82 router alert options
```

```
          0 multicast packets dropped (no iflist)
          0 packets dropped (src and int don't match)
          0 transit re packets dropped on mgmt i/f
          0 packets used first nexthop in ecmp unilist
          0 incoming ttpoip packets received
          0 incoming ttpoip packets dropped
          0 outgoing TTPoIP packets sent
          0 outgoing TTPoIP packets dropped
          0 raw packets dropped. no space in socket recv buffer

lcc2-re0:
--------------------------------------------------------------------------
ip:
          16926850 total packets received
          0 bad header checksums
          0 with size smaller than minimum
          0 with data size < data length
          0 with header length < data size
          0 with data length < header length
          0 with incorrect version number
          0 packets destined to dead next hop
          0 fragments received
          0 fragments dropped (dup or out of space)
          0 fragment sessions dropped (queue overflow)
          0 fragments dropped after timeout
          0 packets reassembled ok
          16926768 packets for this host
          82 packets for unknown/unsupported protocol
          0 packets forwarded
          0 packets not forwardable
          0 redirects sent
          10039747 packets sent from this host
          0 packets sent with fabricated ip header
          0 output packets dropped due to no bufs
          0 output packets discarded due to no route
          0 output datagrams fragmented
          0 fragments created
          0 datagrams that can't be fragmented
          0 packets with bad options
          82 packets with options handled without error
          0 strict source and record route options
          0 loose source and record route options
          0 record route options
```

```
          0 timestamp options
          0 timestamp and address options
          0 timestamp and prespecified address options
          0 option packets dropped due to rate limit
          82 router alert options
          0 multicast packets dropped (no iflist)
          0 packets dropped (src and int don't match)
          0 transit re packets dropped on mgmt i/f
          0 packets used first nexthop in ecmp unilist
          0 incoming ttpoip packets received
          0 incoming ttpoip packets dropped
          0 outgoing TTPoIP packets sent
          0 outgoing TTPoIP packets dropped
          0 raw packets dropped. no space in socket recv buffer

lcc3-re0:
--------------------------------------------------------------------------
ip:
          18025026 total packets received
          0 bad header checksums
          0 with size smaller than minimum
          0 with data size < data length
          0 with header length < data size
          0 with data length < header length
          0 with incorrect version number
          0 packets destined to dead next hop
          0 fragments received
          0 fragments dropped (dup or out of space)
          0 fragment sessions dropped (queue overflow)
          0 fragments dropped after timeout
          0 packets reassembled ok
          18024944 packets for this host
          82 packets for unknown/unsupported protocol
          0 packets forwarded
          0 packets not forwardable
          0 redirects sent
          10456545 packets sent from this host
          0 packets sent with fabricated ip header
          0 output packets dropped due to no bufs
          0 output packets discarded due to no route
          0 output datagrams fragmented
          0 fragments created
          0 datagrams that can't be fragmented
```

```
        0 packets with bad options
        82 packets with options handled without error
        0 strict source and record route options
        0 loose source and record route options
        0 record route options
        0 timestamp options
        0 timestamp and address options
        0 timestamp and prespecified address options
        0 option packets dropped due to rate limit
        82 router alert options
        0 multicast packets dropped (no iflist)
        0 packets dropped (src and int don't match)
        0 transit re packets dropped on mgmt i/f
        0 packets used first nexthop in ecmp unilist
        0 incoming ttpoip packets received
        0 incoming ttpoip packets dropped
        0 outgoing TTPoIP packets sent
        0 outgoing TTPoIP packets dropped
        0 raw packets dropped. no space in socket recv buffer
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system statistics ip6

IN THIS SECTION

- Syntax | **27163**
- Syntax (EX Series Switches) | **27163**

- Syntax (TX Matrix Router) | 27163
- Syntax (TX Matrix Plus Router) | 27163
- Description | 27164
- Options | 27164
- Additional Information | 27165
- Required Privilege Level | 27165
- Sample Output | 27165
- Release Information | 27172

Syntax

```
show system statistics ip6
```

Syntax (EX Series Switches)

```
show system statistics ip6
<all-members>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics ip6
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics ip
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide IPv6 statistics.

Options

none Display system statistics for IPv6.

all-chassis (TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system statistics for IPv6 for all the routers in the chassis.

all-lcc (Optional) On a TX Matrix router, display system statistics for IPv6 for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IPv6 for all connected T1600 or T4000 LCCs.

all-members (Optional) Display IPv6 statistics for all members of the Virtual Chassis configuration.

lcc *number* (Optional) On a TX Matrix router, display system statistics for IPv6 for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IPv6 for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (Optional) Display IPv6 statistics for the local Virtual Chassis member.

member *member-id* (Optional) Display IPv6 statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

scc (Optional) Display system statistics for IPv6 for the TX Matrix router (or switch-card chassis).

sfc *number* (Optional) Display system statistics for IPv6 for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system statistics ip6` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics ip6

```
user@host> show system statistics ip6
ip6:
        0 total packets received
        0 with size smaller than minimum
        0 with data size < data length
        0 with bad options
        0 with incorrect version number
        0 fragments received
        0 fragments dropped (dup or out of space)
        0 fragments dropped after timeout
        0 fragment sessions dropped (queue overflow)
        0 packets reassembled ok
        0 packets for this host
        0 packets forwarded
        0 packets not forwardable
        0 redirects sent
        0 packets sent from this host
        0 packets sent with fabricated ip header
        0 output packets dropped due to no bufs, etc.
        0 output packets discarded due to no route
        0 output datagrams fragmented
```

```
        0 fragments created
        0 datagrams that can't be fragmented
        0 packets that violated scope rules
        0 multicast packets which we don't join
        Mbuf statistics:
        0 packets whose headers are not continuous
        0 tunneling packets that can't find gif
        0 packets discarded due to too may headers
        0 failures of source address selection
        0 forward cache hit
        0 forward cache miss
        0 packets destined to dead next hop
        0 option packets dropped due to rate limit
        0 packets dropped (src and int don't match)
        0 packets dropped due to bad protocol
```

show system statistics ip6 (EX Series Switches)

```
user@host> show system statistics ip6
ip6:
         0 total packets received
         0 packets with size smaller than minimum
         0 packets with data size < data length
         0 packets with bad options
         0 packets with incorrect version number
         0 fragments received
         0 fragments dropped (dup or out of space)
         0 fragments dropped after timeout
         0 fragment sessions dropped (queue overflow)
         0 packets reassembled ok
         0 packets for this host
         0 packets forwarded
         0 packets not forwardable
         0 redirects sent
         0 packets sent from this host
         0 packets sent with fabricated ip header
         0 output packets dropped due to no bufs, etc.
         0 output datagrams fragmented
         0 fragments created
         0 datagrams that can't be fragmented
         0 packets that violated scope rules
```

```
          0 multicast packets which we don't join
          0 packets whose headers are not continuous
          0 tunneling packets that can't find gif
          0 packets discarded due to too may headers
          0 failures of source address selection
          0 forward cache hit
          0 forward cache miss
          0 Packets destined to dead next hop
          0 option packets dropped due to rate limit
          0 Packets dropped (src and int don't match)
          0 packets dropped due to bad protocol
          0 transit re packet(null) dropped on mgmt i/f
```

show system statistics ip6 (TX Matrix Router)

```
user@host> show system statistics ip6
sfc0-re0:
--------------------------------------------------------------------------
ip6:
        0 total packets received
        0 with size smaller than minimum
        0 with data size < data length
        0 with bad options
        0 with incorrect version number
        0 fragments received
        0 fragments dropped (dup or out of space)
        0 fragments dropped after timeout
        0 fragment sessions dropped (queue overflow)
        0 packets reassembled ok
        0 packets for this host
        0 packets forwarded
        0 packets not forwardable
        0 redirects sent
        0 packets sent from this host
        0 packets sent with fabricated ip header
        0 output packets dropped due to no bufs, etc.
        0 output packets discarded due to no route
        0 output datagrams fragmented
        0 fragments created
        0 datagrams that can't be fragmented
        0 packets that violated scope rules
```

```
        0 multicast packets which we don't join
        Mbuf statistics:
        0 packets whose headers are not continuous
        0 tunneling packets that can't find gif
        0 packets discarded due to too may headers
        0 failures of source address selection
        source addresses on an outgoing I/F
                4 link-locals
        source addresses of same scope
                4 link-locals
        0 forward cache hit
        0 forward cache miss
        0 packets destined to dead next hop
        0 option packets dropped due to rate limit
        0 packets dropped (src and int don't match)
        0 packets dropped due to bad protocol
        0 transit re packet(null) dropped on mgmt i/f
        0 packet(null) used first nexthop in ecmp unilist

lcc0-re0:
---------------------------------------------------------------------------
ip6:
        0 total packets received
        0 with size smaller than minimum
        0 with data size < data length
        0 with bad options
        0 with incorrect version number
        0 fragments received
        0 fragments dropped (dup or out of space)
        0 fragments dropped after timeout
        0 fragment sessions dropped (queue overflow)
        0 packets reassembled ok
        0 packets for this host
        0 packets forwarded
        0 packets not forwardable
        0 redirects sent
        0 packets sent from this host
        0 packets sent with fabricated ip header
        0 output packets dropped due to no bufs, etc.
        0 output packets discarded due to no route
        0 output datagrams fragmented
        0 fragments created
        0 datagrams that can't be fragmented
```

```
        0 packets that violated scope rules
        0 multicast packets which we don't join
        Mbuf statistics:
        0 packets whose headers are not continuous
        0 tunneling packets that can't find gif
        0 packets discarded due to too may headers
        0 failures of source address selection
        source addresses on an outgoing I/F
                4 link-locals
        source addresses of same scope
                4 link-locals
        0 forward cache hit
        0 forward cache miss
        0 packets destined to dead next hop
        0 option packets dropped due to rate limit
        0 packets dropped (src and int don't match)
        0 packets dropped due to bad protocol
        0 transit re packet(null) dropped on mgmt i/f
        0 packet(null) used first nexthop in ecmp unilist

lcc1-re0:
--------------------------------------------------------------------------
ip6:
        2 total packets received
        0 with size smaller than minimum
        0 with data size < data length
        0 with bad options
        0 with incorrect version number
        0 fragments received
        0 fragments dropped (dup or out of space)
        0 fragments dropped after timeout
        0 fragment sessions dropped (queue overflow)
        0 packets reassembled ok
        0 packets for this host
        0 packets forwarded
        0 packets not forwardable
        0 redirects sent
        0 packets sent from this host
        0 packets sent with fabricated ip header
        0 output packets dropped due to no bufs, etc.
        0 output packets discarded due to no route
        0 output datagrams fragmented
        0 fragments created
```

```
        0 datagrams that can't be fragmented
        0 packets that violated scope rules
        0 multicast packets which we don't join
        Input histogram:
                ICMP6: 2
        Mbuf statistics:
        0 packets whose headers are not continuous
        0 tunneling packets that can't find gif
        0 packets discarded due to too may headers
        0 failures of source address selection
        source addresses on an outgoing I/F
                4 link-locals
        source addresses of same scope
                4 link-locals
        0 forward cache hit
        0 forward cache miss
        0 packets destined to dead next hop
        0 option packets dropped due to rate limit
        0 packets dropped (src and int don't match)
        0 packets dropped due to bad protocol
        0 transit re packet(null) dropped on mgmt i/f
        0 packet(null) used first nexthop in ecmp unilist

lcc2-re0:
--------------------------------------------------------------------------
ip6:
        2 total packets received
        0 with size smaller than minimum
        0 with data size < data length
        0 with bad options
        0 with incorrect version number
        0 fragments received
        0 fragments dropped (dup or out of space)
        0 fragments dropped after timeout
        0 fragment sessions dropped (queue overflow)
        0 packets reassembled ok
        0 packets for this host
        0 packets forwarded
        0 packets not forwardable
        0 redirects sent
        0 packets sent from this host
        0 packets sent with fabricated ip header
        0 output packets dropped due to no bufs, etc.
```

```
        0 output packets discarded due to no route
        0 output datagrams fragmented
        0 fragments created
        0 datagrams that can't be fragmented
        0 packets that violated scope rules
        0 multicast packets which we don't join
        Input histogram:
                ICMP6: 2
        Mbuf statistics:
        0 packets whose headers are not continuous
        0 tunneling packets that can't find gif
        0 packets discarded due to too may headers
        0 failures of source address selection
        source addresses on an outgoing I/F
                4 link-locals
        source addresses of same scope
                4 link-locals
        0 forward cache hit
        0 forward cache miss
        0 packets destined to dead next hop
        0 option packets dropped due to rate limit
        0 packets dropped (src and int don't match)
        0 packets dropped due to bad protocol
        0 transit re packet(null) dropped on mgmt i/f
        0 packet(null) used first nexthop in ecmp unilist

lcc3-re0:
--------------------------------------------------------------------------
ip6:
        2 total packets received
        0 with size smaller than minimum
        0 with data size < data length
        0 with bad options
        0 with incorrect version number
        0 fragments received
        0 fragments dropped (dup or out of space)
        0 fragments dropped after timeout
        0 fragment sessions dropped (queue overflow)
        0 packets reassembled ok
        0 packets for this host
        0 packets forwarded
        0 packets not forwardable
        0 redirects sent
```

```
        0 packets sent from this host
        0 packets sent with fabricated ip header
        0 output packets dropped due to no bufs, etc.
        0 output packets discarded due to no route
        0 output datagrams fragmented
        0 fragments created
        0 datagrams that can't be fragmented
        0 packets that violated scope rules
        0 multicast packets which we don't join
        Input histogram:
                ICMP6: 2
        Mbuf statistics:
        0 packets whose headers are not continuous
        0 tunneling packets that can't find gif
        0 packets discarded due to too may headers
        0 failures of source address selection
        source addresses on an outgoing I/F
                4 link-locals
        source addresses of same scope
                4 link-locals
        0 forward cache hit
        0 forward cache miss
        0 packets destined to dead next hop
        0 option packets dropped due to rate limit
        0 packets dropped (src and int don't match)
        0 packets dropped due to bad protocol
        0 transit re packet(null) dropped on mgmt i/f
        0 packet(null) used first nexthop in ecmp unilist
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system statistics netisr per-priority-stats

IN THIS SECTION

- Syntax | 27173
- Description | 27173
- Required Privilege Level | 27173
- Sample Output | 27173
- Release Information | 27174

Syntax

```
show system statistics netisr per-priority-stats
```

Description

Displays the management traffic priority level statistics. You can prioritize the internal traffic between the Routing Engine and the Packet Forwarding Engine or other Routing Engine based on the packet classification. This enhancement increases the internal and management traffic performance and provides better packet processing.

Required Privilege Level

view

Sample Output

show system statistics netisr per-priority-stats

```
user@host> show system statistics netisr per-priority-stats
Netisr:
Priority-wise incoming packet statistics
Priority 2:
    0 Received ttp packets
```

```
    0 Dropped ttp packets
    95246 Received non-ttp packets
    38 Dropped non-ttp packets
Priority 3:
    0 Received ttp packets
    0 Dropped ttp packets
    1728 Received non-ttp packets
    0 Dropped non-ttp packets
Priority 4:
    0 Received ttp packets
    0 Dropped ttp packets
    214226 Received non-ttp packets
    0 Dropped non-ttp packets
```

Release Information

Command introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

- priority (Management Traffic) | 13180
- show system statistics netisr qdrops | 27174
- show system statistics netisr socket-drops | 27176

show system statistics netisr qdrops

IN THIS SECTION

- Syntax | 27175
- Description | 27175
- Required Privilege Level | 27175
- Sample Output | 27175
- Release Information | 27176

Syntax

```
show system statistics netisr qdrops
```

Description

Displays the packet queue drops in kernel network stack.

Required Privilege Level

view

Sample Output

show system statistics netisr qdrops

```
user@host> show system statistics netisr qdrops
Netisr:
38  total netisr packet drops
Incoming packet drops in Netisr queues:
    ip:
        0 packet drops
    ip6:
        0 packet drops
    esis:
        0 packet drops
    arp:
        38 ether drops
        38 arp drops
    autoconf:
        0 packet drops
    lacp:
        0 packet drops
    ether:
        0 packet drops
Priority-wise incoming packet drops :
    Priority 0 :
        0 packet drops
    Priority 1 :
```

```
        0 packet drops
    Priority 2 :
        38 ether drops
        38 arp drops
    Priority 3 :
        0 packet drops
    Priority 4 :
        0 packet drops
    Priority 5 :
        0 packet drops
    Priority 6 :
        0 packet drops
    Priority 7 :
        0 packet drops
```

Release Information

Statement introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

- priority (Management Traffic) | 13180
- show system statistics netisr per-priority-stats | 27173
- show system statistics netisr socket-drops | 27176

show system statistics netisr socket-drops

IN THIS SECTION

- Syntax | 27177
- Description | 27177
- Required Privilege Level | 27177
- Sample Output | 27177
- Release Information | 27177

Syntax

```
show system statistics netisr socket-drops
```

Description

Displays the raw and UDP socket drops.

Required Privilege Level

view

Sample Output

show system statistics netisr socket-drops

```
user@host> show system statistics netisr socket-drops
Netisr:
0 raw packets dropped. no space in socket recv buffer
0 dropped due to full udp socket buffers
```

Release Information

Command introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

- priority (Management Traffic) | 13180
- show system statistics netisr per-priority-stats | 27173
- show system statistics netisr qdrops | 27174

show system statistics tcp

IN THIS SECTION

- Syntax | 27178
- Syntax (EX Series Switches) | 27178
- Syntax (TX Matrix Router) | 27178
- Syntax (TX Matrix Plus Router) | 27179
- Description | 27179
- Options | 27179
- Additional Information | 27180
- Required Privilege Level | 27180
- Sample Output | 27180
- Release Information | 27190

Syntax

```
show system statistics tcp
```

Syntax (EX Series Switches)

```
show system statistics tcp
<all-members>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics tcp
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics tcp
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide Transmission Control Protocol (TCP) statistics.

Options

none Display system statistics for TCP.

all-chassis (Optional) Display system statistics for TCP for all the routers in the chassis.

all-lcc (Optional) On a TX Matrix router, display system statistics for TCP for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for TCP for all connected T1600 or T4000 LCCs.

all-members (Optional) Display TCP statistics for all members of the Virtual Chassis configuration.

lcc *number* (Optional) On a TX Matrix router, display system statistics for TCP for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for TCP for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (Optional) Display TCP statistics for the local Virtual Chassis member.

member *member-id*	(Optional) Display TCP statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.
scc	(Optional) Display system statistics for TCP for the TX Matrix router (or switch-card chassis).
sfc *number*	(Optional) Display system statistics for TCP for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system statistics tcp` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics tcp

```
user@host> show system statistics tcp
tcp:
        3844 packets sent
                3618 data packets (1055596 bytes)
                0 data packets (0 bytes) retransmitted
                0 resends initiated by MTU discovery
                205 ack-only packets (148 packets delayed)
                0 URG only packets
                0 window probe packets
                0 window update packets
                1079 control packets
        5815 packets received
                3377 acks (for 1055657 bytes)
                24 duplicate acks
                0 acks for unsent data
                2655 packets (15004 bytes) received in-sequence
```

```
        1 completely duplicate packet  (0 bytes)
        0 old duplicate packets
        0 packets with some dup. data (0 bytes duped)
        0 out-of-order packets (0 bytes)
        0 packets (0 bytes) of data after window
        0 window probes
        7 window update packets
        0 packets received after close
        0 discarded for bad checksums
        0 discarded for bad header offset fields
        0 discarded because packet too short
1 connection request
32 connection accepts
0 bad connection attempts
0 listen queue overflows
33 connections established (including accepts)
30 connections closed (including 0 drops)
        27 connections updated cached RTT on close
        27 connections updated cached RTT variance on close
        0 connections updated cached ssthresh on close
0 embryonic connections dropped
3374 segments updated rtt (of 3220 attempts)
0 retransmit timeouts
        0 connections dropped by rexmit timeout
0 persist timeouts
        0 connections dropped by persist timeout
344 keepalive timeouts
        0 keepalive probes sent
        0 connections dropped by keepalive
1096 correct ACK header predictions
1314 correct data packet header predictions
32 syncache entries added
        0 retransmitted
        0 dupsyn
        0 dropped
        32 completed
        0 bucket overflow
        0 cache overflow
        0 reset
        0 stale
        0 aborted
        0 badack
        0 unreach
```

```
                0 zone failures
        0 cookies sent
        0 cookies received
        0 ACKs sent in response to in-window but not exact RSTs
        0 ACKs sent in response to in-window SYNs on established connections
        0 rcv packets dropped by TCP due to bad address
        0 out-of-sequence segment drops due to insufficient memory
        1058 RST packets
        0 ICMP packets ignored by TCP
        0 send packets dropped by TCP due to auth errors
        0 rcv packets dropped by TCP due to auth errors
```

show system statistics tcp (EX Series Switches)

```
user@host> show system statistics tcp
Tcp:
         572724 packets sent
                 21936 data packets (1887657 bytes)
                 2 data packets retransmitted (20 bytes)
                 0 resends initiated by MTU discovery
                 3724 ack only packets (537 packets delayed)
                 0 URG only packets
                 1 window probe packets
                 1 window update packets
                 1094083 control packets
         1134258 packets received
                 21371 acks(for 1886660 bytes)
                 5870 duplicate acks
                 0 acks for unsent data
                  19908  packets received in-sequence(267794 bytes)
                 3022 completely duplicate packets(0 bytes)
                 0 old duplicate packets
                 4 packets with some duplicate data(4 bytes duped)
                 2 out-of-order packets(2 bytes)
                 0 packets of data after window(0 bytes)
                 0 window probes
                 40 window update packets
                 0 packets received after close
                 0 discarded for bad checksums
                 0 discarded for bad header offset fields
                 0 discarded because packet too short
```

```
        547027 connection requests
        80 connection accepts
        0 bad connection attempts
        0 listen queue overflows
        103 connections established (including accepts)
        547106 connections closed (including 6 drops)
                47 connections updated cached RTT on close
                47 connections updated cached RTT variance on close
                0 connections updated cached ssthresh on close
        547004 embryonic connections dropped
        20862 segments updated rtt(of 567830 attempts)
        2 retransmit timeouts
                0 connections dropped by retransmit timeout
        0 persist timeouts
                0 connections dropped by persist timeout
        3032 keepalive timeouts
                3031 keepalive probes sent
                1 connections dropped by keepalive
        7823 correct ACK header predictions
        12533 correct data packet header predictions
        80 syncache entries added
                0 retransmitted
                0 dupsyn
                4 dropped
                80 completed
                0 bucket overflow
                0 cache overflow
                0 reset
                0 stale
                0 aborted
                0 badack
                0 unreach
                0 zone failures
        0 cookies sent
        0 cookies received
        1 SACK recovery episodes
        1 segment retransmits in SACK recovery episodes
        1 byte retransmits in SACK recovery episodes
        71 SACK options (SACK blocks) received
        1 SACK options (SACK blocks) sent
        0 SACK scoreboard overflow
        0 ACKs sent in response to in-window but not exact RSTs
        0 ACKs sent in response to in-window SYNs on established connections
```

```
         0 rcv packets dropped by TCP due to bad address
         0 out-of-sequence segment drops due to insufficient memory
         547024 RST packets
         0 ICMP packets ignored by TCP
         0 send packets dropped by TCP due to auth errors
         0 rcv packets dropped by TCP due to auth errors
         0 outgoing segments dropped due to policing
```

show system statistics tcp lcc (TX Matrix Router)

```
user@host> show system statistics tcp lcc 2
lcc2-re0:
--------------------------------------------------------------------------
tcp:
        21271 packets sent
                11069 data packets (12044 bytes)
                0 data packets (0 bytes) retransmitted
                0 resends initiated by MTU discovery
                10198 ack-only packets (10194 packets delayed)
                0 URG only packets
                0 window probe packets
                0 window update packets
                4 control packets
        13363 packets received
                11073 acks (for 12044 bytes)
                0 duplicate acks
                0 acks for unsent data
                12895 packets (2400874 bytes) received in-sequence
                0 completely duplicate packets (0 bytes)
                0 old duplicate packets
                0 packets with some dup. data (0 bytes duped)
                0 out-of-order packets (0 bytes)
                0 packets (0 bytes) of data after window
                0 window probes
                0 window update packets
                0 packets received after close
                0 discarded for bad checksums
                0 discarded for bad header offset fields
                0 discarded because packet too short
        4 connection requests
        0 connection accepts
```

```
        0 bad connection attempts
        0 listen queue overflows
        4 connections established (including accepts)
        33 connections closed (including 0 drops)
                0 connections updated cached RTT on close
                0 connections updated cached RTT variance on close
                0 connections updated cached ssthresh on close
        0 embryonic connections dropped
        11073 segments updated rtt (of 11073 attempts)
0 retransmit timeouts
                0 connections dropped by rexmit timeout
        0 persist timeouts
                0 connections dropped by persist timeout
        0 keepalive timeouts
                0 keepalive probes sent
                0 connections dropped by keepalive
        464 correct ACK header predictions
        2172 correct data packet header predictions
        0 ACKs sent in response to in-window but not exact RSTs
        0 ACKs sent in response to in-window SYNs on established connections
        0 out-of-sequence segment drops due to insufficient memory
        0 RST packets
        0 ICMP packets ignored by TCP
```

show system statistics tcp (TX Matrix Plus Router)

```
user@host> show system statistics tcp
sfc0-re0:
--------------------------------------------------------------------------
Tcp:
         10420 packets sent
                 10203 data packets (2374613 bytes)
                 0 data packets retransmitted (0 bytes)
                 0 resends initiated by MTU discovery
                 202 ack only packets (120 packets delayed)
                 0 URG only packets
                 0 window probe packets
                 0 window update packets
                 30 control packets
         16635 packets received
                 9468 acks(for 2374674 bytes)
```

```
        32 duplicate acks
        0 acks for unsent data
        7764  packets received in-sequence(38286 bytes)
        20 completely duplicate packets(0 bytes)
        0 old duplicate packets
        0 packets with some duplicate data(0 bytes duped)
        0 out-of-order packets(0 bytes)
        0 packets of data after window(0 bytes)
        0 window probes
        356 window update packets
        0 packets received after close
        0 discarded for bad checksums
        0 discarded for bad header offset fields
        0 discarded because packet too short
10 connection requests
33 connection accepts
0 bad connection attempts
0 listen queue overflows
34 connections established (including accepts)
50 connections closed (including 0 drops)
        24 connections updated cached RTT on close
        24 connections updated cached RTT variance on close
        0 connections updated cached ssthresh on close
9 embryonic connections dropped
9468 segments updated rtt(of 9256 attempts)
0 retransmit timeouts
        0 connections dropped by retransmit timeout
0 persist timeouts
        0 connections dropped by persist timeout
14 keepalive timeouts
        14 keepalive probes sent
        0 connections dropped by keepalive
6220 correct ACK header predictions
6625 correct data packet header predictions
33 syncache entries added
        0 retransmitted
        0 dupsyn
        0 dropped
        33 completed
        0 bucket overflow
        0 cache overflow
        0 reset
        0 stale
```

```
                0 aborted
                0 badack
                0 unreach
                0 zone failures
        0 cookies sent
        0 cookies received
        0 SACK recovery episodes
        0 segment retransmits in SACK recovery episodes
        0 byte retransmits in SACK recovery episodes
        0 SACK options (SACK blocks) received
        0 SACK options (SACK blocks) sent
        0 SACK scoreboard overflow
        0 ACKs sent in response to in-window but not exact RSTs
        0 ACKs sent in response to in-window SYNs on established connections
        0 rcv packets dropped by TCP due to bad address
        0 out-of-sequence segment drops due to insufficient memory
        15 RST packets
        0 ICMP packets ignored by TCP
        0 send packets dropped by TCP due to auth errors
        0 rcv packets dropped by TCP due to auth errors
        0 outgoing segments dropped due to policing

lcc0-re0:
--------------------------------------------------------------------------
Tcp:
        1306 packets sent
                1251 data packets (161855 bytes)
                0 data packets retransmitted (0 bytes)
                0 resends initiated by MTU discovery
                51 ack only packets (1 packets delayed)
                0 URG only packets
                0 window probe packets
                0 window update packets
                6 control packets
        1397 packets received
                1218 acks(for 161904 bytes)
                2 duplicate acks
                0 acks for unsent data
                612  packets received in-sequence(12495 bytes)
                0 completely duplicate packets(0 bytes)
                0 old duplicate packets
                0 packets with some duplicate data(0 bytes duped)
                0 out-of-order packets(0 bytes)
```

```
                0 packets of data after window(0 bytes)
                0 window probes
                22 window update packets
                0 packets received after close
                0 discarded for bad checksums
                0 discarded for bad header offset fields
                0 discarded because packet too short
        1 connection requests
        24 connection accepts
        0 bad connection attempts
        0 listen queue overflows
        25 connections established (including accepts)
        27 connections closed (including 0 drops)
                24 connections updated cached RTT on close
                24 connections updated cached RTT variance on close
                0 connections updated cached ssthresh on close
        0 embryonic connections dropped
        1218 segments updated rtt(of 1192 attempts)
        0 retransmit timeouts
                0 connections dropped by retransmit timeout
        0 persist timeouts
                0 connections dropped by persist timeout
        0 keepalive timeouts
                0 keepalive probes sent
                0 connections dropped by keepalive
        196 correct ACK header predictions
        119 correct data packet header predictions
        24 syncache entries added
                0 retransmitted
                0 dupsyn
                0 dropped
                24 completed
                0 bucket overflow
                0 cache overflow
                0 reset
                0 stale
                0 aborted
                0 badack
                0 unreach
                0 zone failures
        0 cookies sent
        0 cookies received
        0 SACK recovery episodes
```

```
        0 segment retransmits in SACK recovery episodes
        0 byte retransmits in SACK recovery episodes
        0 SACK options (SACK blocks) received
        0 SACK options (SACK blocks) sent
        0 SACK scoreboard overflow
        0 ACKs sent in response to in-window but not exact RSTs
        0 ACKs sent in response to in-window SYNs on established connections
        0 rcv packets dropped by TCP due to bad address
        0 out-of-sequence segment drops due to insufficient memory
        2 RST packets
        0 ICMP packets ignored by TCP
        0 send packets dropped by TCP due to auth errors
        0 rcv packets dropped by TCP due to auth errors
        0 outgoing segments dropped due to policing

lcc1-re0:
--------------------------------------------------------------------------
Tcp:
        1118 packets sent
                1066 data packets (131896 bytes)
                0 data packets retransmitted (0 bytes)
                0 resends initiated by MTU discovery
                48 ack only packets (2 packets delayed)
                0 URG only packets
                0 window probe packets
                0 window update packets
                6 control packets
        1215 packets received
```

show system statistics tcp (Junos OS Evolved)

```
user@host> show system statistics tcp
Tcp:
        2635574 packets sent
                0 window probe packets
        1124324 packets received
                0 discarded for bad checksums
                0 discarded for bad header offset fields
        3495 connection requests
        2371 bad connection attempts
        0 listen queue overflows
```

```
        1574 connections established (including accepts)
        1 embryonic connections dropped
                22 connections dropped by retransmit timeout
                3677 keepalive probes sent
                8 retransmitted
                17 reset
                0 aborted
        0 cookies sent
        0 cookies received
        199 SACK recovery episodes
        13743 segment retransmits in SACK recovery episodes
        0 out-of-sequence segment drops due to insufficient memory
        2365 RST packets
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system storage

IN THIS SECTION

- Syntax | **27191**
- Syntax (EX Series Switches and MX Series Routers) | **27191**
- Syntax | **27191**
- Syntax (QFX Series) | **27191**
- Syntax (SRX Series) | **27192**
- Syntax (Junos OS Evolved) | **27192**
- Description | **27192**

- Options | 27192
- Required Privilege Level | 27193
- Output Fields | 27193
- Sample Output | 27194
- Release Information | 27197

Syntax

```
show system storage
<detail>
<invoke-on (all-routing-engines | other-routing-engine)>
```

Syntax (EX Series Switches and MX Series Routers)

```
show system storage
<detail>
<all-members>
<local>
<member member-id>
<invoke-on (all-routing-engines | other-routing-engine)>
```

Syntax

Syntax (QFX Series)

```
show system storage
<detail>
<infrastructure name>
<interconnect-device name>
<node-group name>
<invoke-on (all-routing-engines | other-routing-engine)>
```

Syntax (SRX Series)

```
show system storage
<detail>
<partitions>
```

For more information, see "show system storage partitions" on page 27198 (View SRX Series).

Syntax (Junos OS Evolved)

```
show system storage
<detail>
<node node-name>
```

Description

Display statistics about the amount of free disk space in the router's or switch's file systems.

Options

none	Display standard information about the amount of free disk space in the router's or switch's file systems.
detail	(Optional) Display detailed output.
all-members	(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for all members of the Virtual Chassis configuration.
infrastructure *name*	(QFabric systems only) (Optional) Display system storage statistics for the fabric control Routing Engines or fabric manager Routing Engines.
interconnect-device *name*	(QFabric systems only) (Optional) Display system storage statistics for the Interconnect device.
local	(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for the local Virtual Chassis member.
member *member-id*	(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for the specified member of the Virtual Chassis configuration. For

EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node (Junos OS Evolved only) (Optional) Display system storage statistics for the specified node.

node-group *name* (QFabric systems only) (Optional) Display system storage statistics for the Node group.

Required Privilege Level

view

Output Fields

Table 1682 on page 27193 describes the output fields for the `show system storage` command. Output fields are listed in the approximate order in which they appear.

Table 1682: show system storage Output Fields

Field Name	Field Description	Level of Detail
`Filesystem`	Name of the filesystem.	all
`Size`	Size of the filesystem. Size is reported in human readable form (GB or MB, etc.).	standard output
1024-blocks	Size of the filesystem. Size is reported in 1024-bytes (1KB) blocks.	`detail`
`Used`	Amount of space used in the filesystem.	all (see note)
`Avail`	Amount of space available in the filesystem.	all (see note)
`Capacity`	Percentage of the filesystem space that is being used.	all
`Mounted on`	Directory in which the filesystem is mounted.	all

Table 1682: show system storage Output Fields *(Continued)*

Field Name	Field Description	Level of Detail

NOTE: In detailed output, the output is in bytes, whereas in regular output, the size is in human-readable form (like GB or MB, etc.).

Sample Output

show system storage

```
user@host> show system storage
Filesystem              Size       Used      Avail  Capacity   Mounted on
/dev/ad0s1a              77M        37M        34M       52%  /
devfs                    16K        16K         0B      100%  /dev/
/dev/vn0                 12M        12M         0B      100%  /packages/mnt/jbase
/dev/vn1                 39M        39M         0B      100%  /packages/mnt/jkernel-7.2R1.7
/dev/vn2                 12M        12M         0B      100%  /packages/mnt/jpfe-M40-7.2R1.7
/dev/vn3                2.3M       2.3M         0B      100%  /packages/mnt/jdocs-7.2R1.7
/dev/vn4                 14M        14M         0B      100%  /packages/mnt/jroute-7.2R1.7
/dev/vn5                4.5M       4.5M         0B      100%  /packages/mnt/jcrypto-7.2R1.7
mfs:172                 1.5G       4.0K       1.3G        0%  /tmp
/dev/ad0s1e              12M        20K        11M        0%  /config
procfs                  4.0K       4.0K         0B      100%  /proc
/dev/ad1s1f             9.4G       4.9G       3.7G       57%  /var
```

show system storage (SRX Series)

```
user@host> show system storage
Filesystem              Size       Used      Avail  Capacity   Mounted on
/dev/da0s1a             2.4G       369M       1.9G       16%  /
devfs                   1.0K       1.0K         0B      100%  /dev
/dev/md0                 20M        11M       6.7M       63%  /junos
/cf/packages            2.4G       369M       1.9G       16%  /junos/cf/packages
devfs                   1.0K       1.0K         0B      100%  /junos/cf/dev
/dev/md1                1.2G       1.2G         0B      100%  /junos
/cf                      20M        11M       6.7M       63%  /junos/cf
devfs                   1.0K       1.0K         0B      100%  /junos/dev/
```

```
/cf/packages            2.4G       369M       1.9G       16%  /junos/cf/packages1
procfs                  4.0K       4.0K         0B      100%  /proc
/dev/bo0s3e             185M        74K       170M        0%  /config
/dev/bo0s3f             2.1G       1.7G       219M       89%  /cf/var
/dev/md2                1.0G        90M       859M       10%  /mfs
/cf/var/jail            2.1G       1.7G       219M       89%  /jail/var
/cf/var/log             2.1G       1.7G       219M       89%  /jail/var/log
devfs                   1.0K       1.0K         0B      100%  /jail/dev
/dev/md3                1.8M       4.0K       1.7M        0%  /jail/mfs
```

show system storage node

```
user@host> show system storage node re1
Filesystem              Size       Used      Avail  Capacity   Mounted on
/dev/root                44M        44M          0      100%  /pivot
devtmpfs                7.8G          0       7.8G        0%  /dev
/dev/sda2                16G       4.9G        11G       33%  /soft
/dev/sda5               3.0G       122M       2.7G        5%  /etc
/dev/sda6              1000M       1.3M       932M        1%  /config
/dev/sda7                16G       9.4G       5.4G       64%  /var
/dev/sda1               189M       5.3M       170M        4%  /boot
/dev/loop0              1.8G       1.8G          0      100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/ccd-ptx-re64
/dev/loop1               14M        14M          0      100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/perl-5.20.0
/dev/loop2               94M        94M          0      100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/java
/dev/loop3              2.4M       2.4M          0      100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/modules
/dev/loop4              9.4M       9.4M          0      100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/zookeeper
/dev/loop5               57M        57M          0      100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/python-2.7
/dev/loop6               14M        14M          0      100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/python-3.3
/dev/loop7              191M       191M          0      100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/dev
/dev/loop8              3.8M       3.8M          0      100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/jimbase
/dev/loop9              103M       103M          0      100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/osbase
```

```
/dev/loop10              44M        44M          0      100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/initrd
unionfs                 5.2G       2.4G       2.7G       48%  /
/tmp                    7.8G       4.0K       7.8G        1%  /tmp
run                     7.8G       7.1M       7.8G        1%  /run
tmpfs                   7.8G       1.2G       6.7G       15%  /dev/shm
tmpfs                   7.8G          0       7.8G        0%  /sys/fs/cgroup
tmpfs                   1.6G          0       1.6G        0%  /run/user/0
```

show system storage node detail

```
user@host> show system storage node re1 detail
Filesystem       1024-blocks       Used      Avail   Capacity   Mounted on
/dev/root              44376      44376          0       100%  /pivot
devtmpfs             8103560          0    8103560         0%  /dev
/dev/sda2           16513960    5057236   10601480        33%  /soft
/dev/sda5            3055376     124232    2757476         5%  /etc
/dev/sda6            1023892       1308     953772         1%  /config
/dev/sda7           16310696    9809324    5656368        64%  /var
/dev/sda1             193242       5418     173561         4%  /boot
/dev/loop0           1816864    1816864          0       100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/ccd-ptx-re64
/dev/loop1             13432      13432          0       100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/perl-5.20.0
/dev/loop2             95344      95344          0       100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/java
/dev/loop3              2384       2384          0       100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/modules
/dev/loop4              9528       9528          0       100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/zookeeper
/dev/loop5             57816      57816          0       100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/python-2.7
/dev/loop6             13320      13320          0       100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/python-3.3
/dev/loop7            195024     195024          0       100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/dev
/dev/loop8              3872       3872          0       100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/jimbase
/dev/loop9            105272     105272          0       100%  /pivot/data/junos-install-qfx-
x86-64-16.2I20170508115447_evo-builder/osbase
/dev/loop10            44376      44376          0       100%  /pivot/data/junos-install-qfx-
```

```
x86-64-16.2I20170508115447_evo-builder/initrd
unionfs              5412608    2481464    2757476        48%  /
/tmp                 8127388          4    8127384         1%  /tmp
run                  8127388       7216    8120172         1%  /run
tmpfs                8127388    1190096    6937292        15%  /dev/shm
tmpfs                8127388          0    8127388         0%  /sys/fs/cgroup
tmpfs                1625480          0    1625480         0%  /run/user/0
```

show system storage node (SRX4000 line only)

```
user@host> show system storage node 0
node0:
--------------------------------------------------------------------------
Filesystem              Size       Used      Avail  Capacity   Mounted on
/dev/gpt/junos          4.0G       735M       3.0G       20%  /.mount
tmpfs                   5.7G        64K       5.7G        0%  /.mount/tmp
tmpfs                   643M       352K       642M        0%  /.mount/mfs
/dev/gpt/config         1.5G        84K       1.4G        0%  /.mount/config
/dev/gpt/var             57G       131M        52G        0%  /.mount/var
host_corefiles          379G       872M       359G        0%  /.mount/var/crash/corefiles
host_volatile           125G       4.0K       125G        0%  /.mount/var/log/host
host_log                379G       872M       359G        0%  /.mount/var/log/hostlogs
host_traffic_log        379G       872M       359G        0%  /.mount/var/traffic-log
host_local              379G       872M       359G        0%  /.mount/var/db/host
host_aamwd              379G       872M       359G        0%  /.mount/var/db/aamwd
host_geoip              379G       872M       359G        0%  /.mount/var/db/geoip
host_secinteld          379G       872M       359G        0%  /.mount/var/db/secinteld
host_app_disk           9.0G       2.3G       6.2G       27%  /.mount/var/install_disk
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `invoke-on (all-routing-engines | other-routing-engine)` introduced in Junos OS Release 14.1

`node` option introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

show system storage partitions

show system storage partitions

IN THIS SECTION

- Syntax (EX Series) | 27198
- Syntax (SRX Series) | 27198
- Description | 27198
- Options | 27199
- Required Privilege Level | 27199
- Output Fields | 27199
- Sample Output | 27200
- Release Information | 27201

Syntax (EX Series)

```
show system storage partitions
<all-members>
<local>
<member member-id>
```

Syntax (SRX Series)

```
show system storage partitions
```

Description

This command displays information about the disk partitioning scheme.

Options

none Display partition information.

`all-members` (Virtual Chassis systems only) (Optional) Display partition information for all members of the Virtual Chassis.

`local` (Virtual Chassis systems only) (Optional) Display partition information for the local Virtual Chassis member.

`member` *`member-id`* (Virtual Chassis systems only) (Optional) Display partition information for the specified member of the Virtual Chassis configuration.

Required Privilege Level

view

Output Fields

Table 1683 on page 27199 describes the output fields for the `show system storage partitions` command. Output fields are listed in the approximate order in which they appear.

Table 1683: show system storage partitions Output Fields

Field Name	Field Description
Boot Media	Media (internal or external) from which the switch was booted.
Active Partition	Name of the active root partition.
Backup Partition	Name of the backup (alternate) root partition.
Currently booted from	Partition from which the switch was last booted.

Table 1683: show system storage partitions Output Fields *(Continued)*

Field Name	Field Description
Partitions information	Information about partitions on the boot media: • Partition—Partition identifier. • Size—Size of partition. • Mountpoint—Directory on which the partition is mounted.

Sample Output

show system storage partitions (EX Series)

```
user@switch> show system storage partitions
fpc0:
--------------------------------------------------------------------------
Boot Media: internal (da0)
Active Partition: da0s1a
Backup Partition: da0s2a
Currently booted from: active (da0s1a)

Partitions information:
  Partition  Size   Mountpoint
  s1a        184M   /
  s2a        184M   altroot
  s3d        369M   /var/tmp
  s3e        123M   /var
  s4d        62M    /config
  s4e               unused (backup config)
```

show system storage partitions (SRX Series, Dual Root Partitioning)

```
show system storage partitions
Boot Media: internal (da0)
Active Partition: da0s2a
Backup Partition: da0s1a
```

```
Currently booted from: active (da0s2a)

Partitions Information:
  Partition  Size   Mountpoint
      s1a    293M   altroot
      s2a    293M   /
      s3e    24M    /config
      s3f    342M   /var
      s4a    30M    recovery
```

show system storage partitions (SRX Series, Single Root Partitioning)

```
show system storage partitions
Boot Media: internal (da0)
Partitions Information:
  Partition  Size   Mountpoint
      s1a    898M   /
      s1e    24M    /config
      s1f    61M    /var
```

show system storage partitions (SRX Series, USB)

```
show system storage partitions
Boot Media: usb (da1)
Active Partition: da1s1a
Backup Partition: da1s2a
Currently booted from: active (da1s1a)

Partitions Information:
  Partition  Size   Mountpoint
      s1a    293M   /
      s2a    293M   altroot
      s3e    24M    /config
      s3f    342M   /var
      s4a    30M    recovery
```

Release Information

Command introduced in Junos OS Release 15.1X49-D35.

RELATED DOCUMENTATION

Verifying Junos OS and Boot Loader Software Versions on an EX Series Switch

Example: Installing Junos OS on SRX Series Devices Using the Partition Option

[EX] Switch boots from backup root partition after file system corruption occurred on the primary root partition

show system storage satellite

IN THIS SECTION

- Syntax | 27202
- Description | 27202
- Options | 27202
- Required Privilege Level | 27203
- Output Fields | 27203
- Sample Output | 27204
- Release Information | 27205

Syntax

```
show system storage satellite
<fpc-slot>
device-alias <device-alias>
```

Description

Displays the storage usage of satellite devices in Junos Fusion.

Options

none	Display the storage usage for all listed satellite devices.

detail	(Optional) Display detail output.
fpc-slot	(Optional) Display satellite device connection information for the satellite device using the specified FPC slot number only.
device-alias ***device-alias***	(Optional) Display satellite device connection information for the satellite device using the specified device alias only.

Required Privilege Level

view

Output Fields

Table 1684 on page 27203 describes the output fields for the `show system storage` command. Output fields are listed in the approximate order in which they appear.

Table 1684: show system storage Output Fields

Field Name	Field Description
`Filesystem`	Name of the filesystem.
`Size`	Size of the filesystem.
`Used`	Amount of space used in the filesystem.
`Avail`	Amount of space available in the filesystem.
`Use %`	Percentage of the filesystem space that is being used.
`Mounted on`	Directory in which the filesystem is mounted.

Sample Output

show system storage satellite

```
user@host> show system storage satellite
Slot-ID: 125
Filesystem      Size  Used Avail Use% Mounted on
rootfs          665M  287M  330M  47% /
udev            896M   24K  896M   1% /dev
none            896M   24K  896M   1% /dev
tmpfs           936M  812K  935M   1% /run
/dev/sda3       665M  287M  330M  47% /
/dev/sda5       327M   99M  207M  33% /var
tmpfs           936M  812K  935M   1% /run
tmpfs           936M   14M  923M   2% /var/volatile
/dev/sda1       481M  254M  228M  53% /boot
/dev/sda2       259M  2.1M  240M   1% /app_disk
tmpfs           936M  812K  935M   1% /run/named-chroot/var/run/named
tmpfs           936M  812K  935M   1% /run/named-chroot/var/run/bind
none            896M   24K  896M   1% /run/named-chroot/dev/random
none            896M   24K  896M   1% /run/named-chroot/dev/zero
none            896M   24K  896M   1% /run/named-chroot/dev/null

Slot-ID: 134
Filesystem      Size  Used Avail Use% Mounted on
rootfs          665M  287M  330M  47% /
udev            896M   24K  896M   1% /dev
none            896M   24K  896M   1% /dev
tmpfs           936M  812K  935M   1% /run
/dev/sda3       665M  287M  330M  47% /
/dev/sda5       327M   97M  209M  32% /var
tmpfs           936M  812K  935M   1% /run
tmpfs           936M   13M  924M   2% /var/volatile
/dev/sda1       481M  254M  228M  53% /boot
/dev/sda2       259M  2.1M  240M   1% /app_disk
tmpfs           936M  812K  935M   1% /run/named-chroot/var/run/named
tmpfs           936M  812K  935M   1% /run/named-chroot/var/run/bind
none            896M   24K  896M   1% /run/named-chroot/dev/random
none            896M   24K  896M   1% /run/named-chroot/dev/zero
none            896M   24K  896M   1% /run/named-chroot/dev/null
```

Release Information

Command introduced in Junos OS Release 17.1R1.

show system subscriber-management redundancy-state dhcp active-leasequery interface

IN THIS SECTION

- Syntax | 27205
- Description | 27205
- Options | 27206
- Required Privilege Level | 27206
- Output Fields | 27206
- Sample Output | 27207
- Release Information | 27207

Syntax

```
show system subscriber-management redundancy-state dhcp active-leasequery interface interface-name
```

Description

Display the redundancy state of the specified access logical interface based on the redundancy primary role election protocol, VRRP or pseudowire.

For VRRP redundancy, the redundancy state is the same as the VRRP state of the underlying logical interface:

- The redundancy state is `Master` if the logical interface is in VRRP active mode.
- The redundancy state is `Backup` if the logical interface is in VRRP backup mode.

For pseudowire redundancy, the redundancy state is based on the state of the pseudowire interface:

- The redundancy state is `Master` if the logical interface is in the UP state for a pseudowire interface.
- The redundancy state is `Backup` if the logical interface is in the DOWN state for a pseudowire interface.

Options

interface-name	Name of the access logical interface.

Required Privilege Level

view

Output Fields

Table 1685 on page 27206 lists the output fields for the `show system subscriber-management redundancy-state dhcp active-leasequery` command. Output fields are listed in the approximate order in which they appear.

Table 1685: show system subscriber-management redundancy-state dhcp active-leasequery Output Fields

Field Name	Field Description
interface	Name of the access logical interface.
Redundancy State	Redundancy state of the specified interface. • `Master`—Interface is on the current primary BNG. This BNG is handling traffic for all members of the subscriber redundancy group that is on this interface. For VRRP, the interface state is `Master`; for pseudowire, the interface state is `Up`. • Backup—Interface is on the current backup BNG. If the primary BNG fails, this BNG is elected as the new primary. This interface transitions to the `Master` state and handles traffic for all members of the subscriber redundancy group previously handled on the corresponding access interface on the former primary. For VRRP, the interface state is backup; for pseudowire, the interface state is down.

Sample Output

show system subscriber-management redundancy-state dhcp active-leasequery interface (VLAN Underlying Interface)

```
user@host> show system subscriber-management redundancy-state dhcp active-leasequery interface
ge-0/0/0.1

Interface                      Redundancy State
ge-0/0/0.1                     Master
```

show system subscriber-management redundancy-state dhcp active-leasequery interface (Subscriber Demux Interface)

```
user@host> show system subscriber-management redundancy-state dhcp active-leasequery interface
demux0.3221225473

Interface                      Redundancy State
demux0.3221225473              Master
```

show system subscriber-management redundancy-state dhcp active-leasequery interface (Pseudowire Interface)

```
user@host> show system subscriber-management redundancy-state dhcp active-leasequery interface
ps1.0

Interface                      Redundancy State
ps1.0                          Master
```

Release Information

Command introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

M:N Subscriber Redundancy

show system subscriber-management redundancy-state interface

IN THIS SECTION

- Syntax | 27208
- Description | 27208
- Options | 27208
- Required Privilege Level | 27208
- Output Fields | 27208
- Sample Output | 27209
- Release Information | 27211

Syntax

```
show system subscriber-management redundancy-state interface interface-name
```

Description

Displays the status of the subscriber management redundancy service of the pseudowire interface.

Options

interface *interface-name*	Name of the pseudowire interface.

Required Privilege Level

view

Output Fields

Table 1686 on page 27209 lists the output fields for the `show system subscriber-management redundancy-state interface interface-name` command.

Table 1686: show system subscriber-management redundancy-state interface Output Fields

Field Name	Field Description
Forwarding state	Status of the packet forwarding.
Service Activation Programming	Status of the service activation programming.
Standby-mode	Status of the standby mode.

Sample Output

Primary BNG Interface Status

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: hot-standby
Forwarding state: Active
Service Activation Programming: Completed
```

Secondary BNG Interface Status in Normal Operation

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: service-activation-on-failover
Forwarding state: Inactive
Service Activation Programming: Not-applicable
```

Secondary BNG Interface Status Immediately on Failover

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: hot-standby
Forwarding state: Active
Service Activation Programming: In-progress
```

Secondary BNG Interface Status on Failover Completion

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: hot-standby
Forwarding state: Active
Service Activation Programming: Completed
```

Secondary BNG Interface Status Immediately on Failover Reversal

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: service-activation-on-failover
Forwarding state: Inactive
Service Activation Programming: In-rollback
```

Secondary BNG Interface Status Immediately on Failover Reversal Completion

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: service-activation-on-failover
Forwarding state: Inactive
Service Activation Programming: Not-applicable
```

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

N+1 Support for BNG M:N Subscriber Service Redundancy

show system subscriber-management resiliency

IN THIS SECTION

- Syntax | 27212
- Description | 27212
- Options | 27212
- Required Privilege Level | 27212
- Output Fields | 27212
- Sample Output | 27215
- Release Information | 27222

Syntax

```
show system subscriber-management resiliency
<detail>
<extensive>
<summary>
```

Description

Display information that indicates the health and relationship of session database replication between the primary and standby Routing Engines.

Options

detail (Optional) Displays brief information about the shared memory state for the primary and standby Routing Engines.

extensive (Optional) Displays very detailed statistics for the SDB components in shared memory for the primary and standby Routing Engines, enabling you to evaluate the state of replication between the two.

summary (Optional) Displays only an indication of whether the system is okay (replication is normal) or has some unexpected condition.

Required Privilege Level

system

Output Fields

Table 1687 on page 27213 lists the output fields for the `show system subscriber-management resiliency` command. Output fields are listed in the approximate order in which they appear.

Table 1687: show system subscriber-management resiliency Output

Field Name	Field Description	Level
`Overall Status`	Indicates the condition of the system: • `Ok`—The system is functioning normally. • `Not-Ok`—An unexpected condition has been discovered. This status may require investigation by the Juniper Networks Technical Assistance Center (JTAC) to confirm whether anything is wrong and the root cause of the status.	`summary`
`shared memory type`	One of the following types of shared memory objects: • `mmap`—Memory-mapped file that stores the hash or entry data for an MMDB. • `mmap Database (MMDB)`–Memory-mapped database that uses memory-mapped files to store the MMDB hash and entry data. Each MMDB typically stores a type of statistic, such as statistics related logical interfaces, logical interface sets, or subscribers. • `Shared Memory Segment`—An operating system object that is a chunk of contiguous shared memory. `Total`—Number of memory objects of all types.	`detail`
`count`	Number of shared memory instances of a type.	`detail`
`mapped bytes`	Number of bytes mapped into process space.	`detail`
`mmfs`	Memory-mapped file information.	`extensive`
`Name`	File path including the filename of the shared memory object. For MMFs, the filename is the name of its associated MMDB and a suffix to indicate whether it stores hash or data. For MMDBs, the filename indicates the type of statistics stored in the database.	`extensive`
`Current Bytes`	Current total size of the shared memory object.	`extensive`

Table 1687: show system subscriber-management resiliency Output *(Continued)*

Field Name	Field Description	Level
Maximum Bytes	Maximum size of the shared memory object.	extensive
Mapped Bytes	Number of bytes mapped into process space.	extensive
Lock Count	Number of times the shared memory object has been locked by a global, inter-process lock.	extensive
Contention Count	Number of times that a process or thread object waited to lock a shared memory object because a different process or thread already has the lock. This is a global, inter-process lock.	extensive
Lock Wait Secs	How long a process or thread taking a global, inter-process lock waited because a different process or thread already had the lock.	extensive
mmap Count	Number of times that parts of the overall memory mapped data have been mapped.	extensive
Shared Memory Segments	Information about the shared memory segments; each segment is a chunk of contiguous shared memory.	extensive
Size in Bytes	Number of bytes in the shared memory segment.	extensive
MMDBs	Information about the memory-mapped file databases that use memory-mapped files to store data (typically statistics associated with interfaces and subscribers).	extensive
Hash Entries	Number of different hash entries a key could be hashed to in this table.	extensive
PLock Count	Number of times the MMDB shared memory object has been locked by a process-level, intra-process lock.	extensive

Table 1687: show system subscriber-management resiliency Output *(Continued)*

Field Name	Field Description	Level
PLock Contention Count	Number of times that a process or thread object waited to lock a shared memory object because a different process or thread already has the lock. This is a process-level, intra-process lock.	extensive
PLock Wait Secs	How long a process or thread taking a process-level, intra-process lock waited because a different process or thread already had the lock.	extensive

Sample Output

show system subscriber-management resiliency (Summary)

```
user@host> show system subscriber-management resiliency summary
Overall Status: Ok
```

show system subscriber-management resiliency (Detail)

```
user@host> show system subscriber-management resiliency detail
Master:
shared memory type             count mapped bytes
mmap                              43    195027200
mmap Database (MMDB)               9    (in mmap)
Shared Memory Segment              6     39163504
Total                             58    234190704

Standby:
shared memory type             count mapped bytes
mmap                              41    192930048
mmap Database (MMDB)               9    (in mmap)
Shared Memory Segment              6     39163504
Total                             56    232093552
```

show system subscriber-management resiliency (Extensive)

```
user@host> show system subscriber-management resiliency extensive
Master:

   mmfs:
      Name                                                        Current Bytes Maximum Bytes
Mapped bytes   Lock Count Lock Contention Count   Lock Wait Secs   mmap Count
      /mfs/var/smm_accounting-stats-db_hash                              15736832
15736832     15736832           17                     0         0.000000            0
      /mfs/var/smm_accounting-stats-db_data                             1139015680
9112125440      2097152           17                     0         0.000000           18
      /mfs/var/mmcq/mmdb_rep_mmcq                                         1048576
104857600      1048576           25                     1         0.011021            0
      /mfs/var/smm_accounting-ifl-db_hash                                   28672
28672        28672           17                     0         0.000000            0
      /mfs/var/smm_accounting-ifl-db_data                                33554432
536870912      4194304           17                     0         0.000000           18
      /mfs/var/smm_accounting-iflset-db_hash                                28672
28672        28672           17                     0         0.000000            0
      /mfs/var/smm_accounting-iflset-db_data                             33554432
536870912      4194304           17                     0         0.000000           18
      /mfs/var/sdb/shmem/sdb.head                                         7680256
7680256      7680256       384006                     0         0.000000            0
      /mfs/var/sdb/shmem/sdb.lts.data                                   1620049920
8589934592     20971520           41                     0         0.000000           60
      /mfs/var/sdb/shmem/sdb_sts_data                                    51216384
51216384     51216384        20012                     0         0.000000            0
      /mfs/var/sdb/shmem/sdb_intf.db                                       409600
409600       409600            0                     0         0.000000            0
      /mfs/var/sdb/shmem/subscriber_hash                                  2408448
2408448      2408448           21                     0         0.000000            0
      /mfs/var/sdb/shmem/subscriber_data                                 33554432
536870912      2097152           22                     0         0.000000           22
      /mfs/var/sdb/shmem/service_hash                                     2408448
2408448      2408448           21                     0         0.000000            0
      /mfs/var/sdb/shmem/service_data                                    33554432
536870912      2097152           22                     0         0.000000           22
      /mfs/var/sdb/shmem/interface_hash                                     28672
28672        28672           21                     0         0.000000            0
      /mfs/var/sdb/shmem/interface_data                                  33554432
536870912      4194304          109                     0         0.000000           22
```

```
      /mfs/var/sdb/shmem/interface_set_hash                                28672
28672        28672           21                     0         0.000000            0
      /mfs/var/sdb/shmem/interface_set_data                             33554432
536870912      4194304           22                     0         0.000000           22
      /mfs/var/sdb/shmem/mobile_subs_location_hash                       1208320
1208320      1208320           21                     0         0.000000            0
      /mfs/var/sdb/shmem/mobile_subs_location_data                      33554432
536870912      2097152           22                     0         0.000000           22
      /mfs/var/sdb/shmem/mobile_subscriber_hash                          1208320
1208320      1208320           21                     0         0.000000            0
      /mfs/var/sdb/shmem/mobile_subscriber_data                         33554432
536870912      2097152           21                     0         0.000000           22
      /mfs/var/mmq/mmq_queue                                              126976
126976       126976            5                     0         0.000000            0
      /mfs/var/mmq/mmq_heap                                              5120000
5120000      5120000            4                     0         0.000000            0
      /mfs/var/mmcq/sdb_bbe_mmcq                                        25165824
318767104     25165824           21                     0         0.000000            0
      /mfs/var/mmcq/authdRxQueue                                         1048576
20971520      1048576            6                     0         0.000000            0
      /mfs/var/mmcq/pppdRxQueue                                          1048576
20971520      1048576            4                     0         0.000000            0
      /mfs/var/mmcq/bbeStatsdGetCollector                                1048576
20971520      1048576           16                     0         0.000000            0
      /mfs/var/mmdb/mmdb_ack_registry                                       8192
8192         8192          141                     0         0.000000            0
      /mfs/var/mmcq/mmdb_ackq_bbe-statsd                                 1048576
67108864      1048576            2                     0         0.000000            0
      /mfs/var/mmcq/jdchpdAccountingClientApp                            1048576
20971520      1048576            2                     0         0.000000            0
      /mfs/var/ss/domain.0.data                                         16777216
2147483648      4194304          262                     0         0.000000           18
      /mfs/var/tmp/bbe_throttle_control                                     8192
8192         8192            7                     0         0.000000            0
      /mfs/var/mmcq/statsPluginGCClient                                  1048576
20971520      1048576            2                     0         0.000000            0
      /mfs/var/sdb/shmem/sdb_reg_info                                       8192
8192         8192            2                     0         0.000000            0
      /mfs/var/mmcq/sdb_reg_q_bbe-statsd                                16777216
16777216     16777216            2                     0         0.000000            0
      /mfs/var/mmcq/jl2tpdCliRxQ                                         1048576
20971520      1048576            2                     0         0.000000            0
      /mfs/var/mmcq/jl2tpdSnmpRxQ                                        1048576
```

```
20971520      1048576            2                     0         0.000000            0
      /mfs/var/mmcq/authd                                                    1048576
20971520      1048576            2                     0         0.000000            0
      /mfs/var/mmcq/jpppdAccountingClientApp                                 1048576
20971520      1048576            2                     0         0.000000            0
      /mfs/var/mmcq/mmdb_mmcq_0                                              1048576
104857600      1048576           42                     0         0.000000            0
      /mfs/var/ss/domain.0                                                    409600
4294967295       409600      6400000                  3037         0.002642            0

   Shared Memory Segments:
      Name                                                        Size in Bytes
      /mfs/var/shmlog/shmlog                                           39071744
      sdb_rsmon_shared_memory                                             22536
      sdb_rsmon_ae_table                                                   4096
      sdb_rsmon_ps_table                                                  60008
      sdb_rsmon_rlt_table                                                  1024
      sdb_bbe_rep_mailbox                                                  4096

   MMDBs:
      Name                                                          Hash Entries   Lock Count
Lock Contention Count   Lock Wait Secs   PLock Count PLock Contention Count   PLock Wait Secs
      /mfs/var/smm_accounting-stats-db                                    655360
7208990                     0         0.000000       1966111                     0
0.000000
      /mfs/var/smm_accounting-ifl-db                                        1000
11024                     0         0.000000          3025                     0
0.000000
      /mfs/var/smm_accounting-iflset-db                                     1000
11024                     0         0.000000          3025                     0
0.000000
      /mfs/var/sdb/shmem/subscriber                                       100000
1400010                     2         0.043705        400012                     0
0.000000
      /mfs/var/sdb/shmem/service                                          100000
1400010                     0         0.000000        400012                     0
0.000000
      /mfs/var/sdb/shmem/interface                                          1000
14430                     0         0.000000          4427                     0
0.000000
      /mfs/var/sdb/shmem/interface_set                                      1000
14010                     0         0.000000          4012                     0
0.000000
```

```
      /mfs/var/sdb/shmem/mobile_subs_location                                50000
700018                     0         0.000000        200020                     0
0.000000
      /mfs/var/sdb/shmem/mobile_subscriber                                   50000
700010                     0         0.000000        200012                     0
0.000000

   Total Mapped Bytes                                                 234190704

Standby:

   mmfs:
      Name                                                            Current Bytes Maxiumum Bytes
Mapped bytes   Lock Count Lock Contention Count   Lock Wait Secs   mmap Count
      /mfs/var/smm_accounting-stats-db_hash                                15736832
15736832     15736832           13                     0         0.000000            0
      /mfs/var/smm_accounting-stats-db_data                               1139015680
9112125440      2097152           13                     0         0.000000           14
      /mfs/var/mmcq/mmdb_rep_mmcq                                           1048576
104857600      1048576           15                     0         0.000000            0
      /mfs/var/smm_accounting-ifl-db_hash                                     28672
28672        28672           13                     0         0.000000            0
      /mfs/var/smm_accounting-ifl-db_data                                  33554432
536870912      4194304           13                     0         0.000000           14
      /mfs/var/smm_accounting-iflset-db_hash                                  28672
28672        28672           13                     0         0.000000            0
      /mfs/var/smm_accounting-iflset-db_data                               33554432
536870912      4194304           13                     0         0.000000           14
      /mfs/var/sdb/shmem/sdb.head                                           7680256
7680256      7680256       384005                     0         0.000000            0
      /mfs/var/sdb/shmem/sdb.lts.data                                     1620049920
8589934592     20971520           11                     0         0.000000           20
      /mfs/var/sdb/shmem/sdb_sts_data                                      51216384
51216384     51216384        17510                     0         0.000000            0
      /mfs/var/sdb/shmem/sdb_intf.db                                         409600
409600       409600            0                     0         0.000000            0
      /mfs/var/sdb/shmem/subscriber_hash                                    2408448
2408448      2408448            5                     0         0.000000            0
      /mfs/var/sdb/shmem/subscriber_data                                   33554432
536870912      2097152            4                     0         0.000000            6
      /mfs/var/sdb/shmem/service_hash                                       2408448
2408448      2408448            5                     0         0.000000            0
      /mfs/var/sdb/shmem/service_data                                      33554432
```

```
536870912      2097152            4                    0         0.000000            6
      /mfs/var/sdb/shmem/interface_hash                                    28672
28672        28672            5                    0         0.000000            0
      /mfs/var/sdb/shmem/interface_data                                 33554432
536870912      4194304            4                    0         0.000000            6
      /mfs/var/sdb/shmem/interface_set_hash                                28672
28672        28672            5                    0         0.000000            0
      /mfs/var/sdb/shmem/interface_set_data                             33554432
536870912      4194304            4                    0         0.000000            6
      /mfs/var/sdb/shmem/mobile_subs_location_hash                      1208320
1208320      1208320            5                    0         0.000000            0
      /mfs/var/sdb/shmem/mobile_subs_location_data                      33554432
536870912      2097152            4                    0         0.000000            6
      /mfs/var/sdb/shmem/mobile_subscriber_hash                         1208320
1208320      1208320            5                    0         0.000000            0
      /mfs/var/sdb/shmem/mobile_subscriber_data                         33554432
536870912      2097152            4                    0         0.000000            6
      /mfs/var/mmq/mmq_queue                                             126976
126976       126976            4                    0         0.000000            0
      /mfs/var/mmq/mmq_heap                                             5120000
5120000      5120000            3                    0         0.000000            0
      /mfs/var/mmcq/sdb_bbe_mmcq                                        25165824
318767104     25165824           11                    0         0.000000            0
      /mfs/var/mmcq/authdRxQueue                                        1048576
20971520      1048576            6                    0         0.000000            0
      /mfs/var/mmcq/pppdRxQueue                                         1048576
20971520      1048576            2                    0         0.000000            0
      /mfs/var/mmcq/bbeStatsdGetCollector                               1048576
20971520      1048576           14                    0         0.000000            0
      /mfs/var/mmdb/mmdb_ack_registry                                      8192
8192         8192            2                    0         0.000000            0
      /mfs/var/mmcq/mmdb_ackq_bbe-statsd                                1048576
67108864      1048576            2                    0         0.000000            0
      /mfs/var/mmcq/jdchpdAccountingClientApp                           1048576
20971520      1048576            2                    0         0.000000            0
      /mfs/var/ss/domain.0.data                                        16777216
2147483648      4194304          261                    0         0.000000           16
      /mfs/var/tmp/bbe_throttle_control                                    8192
8192         8192            6                    0         0.000000            0
      /mfs/var/mmcq/statsPluginGCClient                                 1048576
20971520      1048576            2                    0         0.000000            0
      /mfs/var/sdb/shmem/sdb_reg_info                                      8192
8192         8192            2                    0         0.000000            0
```

```
      /mfs/var/mmcq/sdb_reg_q_bbe-statsd                                  16777216
16777216     16777216            2                     0         0.000000            0
      /mfs/var/mmcq/jl2tpdCliRxQ                                           1048576
20971520      1048576            2                     0         0.000000            0
      /mfs/var/mmcq/jl2tpdSnmpRxQ                                          1048576
20971520      1048576            2                     0         0.000000            0
      /mfs/var/mmcq/authd                                                  1048576
20971520      1048576            2                     0         0.000000            0
      /mfs/var/ss/domain.0                                                  409600
4294967295       409600      8000000                  4044         0.002962            0

   Shared Memory Segments:
      Name                                                        Size in Bytes
      /mfs/var/shmlog/shmlog                                           39071744
      sdb_rsmon_shared_memory                                             22536
      sdb_rsmon_ae_table                                                   4096
      sdb_rsmon_ps_table                                                  60008
      sdb_rsmon_rlt_table                                                  1024
      sdb_bbe_rep_mailbox                                                  4096

   MMDBs:
      Name                                                          Hash Entries   Lock Count
Lock Contention Count   Lock Wait Secs   PLock Count PLock Contention Count   PLock Wait Secs
      /mfs/var/smm_accounting-stats-db                                    655360
5898264                     0         0.000000       1966105                      0
0.000000
      /mfs/var/smm_accounting-ifl-db                                        1000
9020                     0         0.000000          3021                      0
0.000000
      /mfs/var/smm_accounting-iflset-db                                     1000
9020                     0         0.000000          3021                      0
0.000000
      /mfs/var/sdb/shmem/subscriber                                       100000
300002                     0         0.000000        100003                      0
0.000000
      /mfs/var/sdb/shmem/service                                          100000
300002                     0         0.000000        100003                      0
0.000000
      /mfs/var/sdb/shmem/interface                                          1000
3002                     0         0.000000          1003                      0
0.000000
      /mfs/var/sdb/shmem/interface_set                                      1000
3002                     0         0.000000          1003                      0
```

```
0.000000
      /mfs/var/sdb/shmem/mobile_subs_location                                50000
150002                    0         0.000000         50003                       0
0.000000
      /mfs/var/sdb/shmem/mobile_subscriber                                   50000
150002                    0         0.000000         50003                       0
0.000000

   Total Mapped Bytes                                              232093552
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Junos OS Enhanced Subscriber Management

show system subscriber-management route

IN THIS SECTION

- Syntax | 27223
- Description | 27223
- Options | 27223
- Required Privilege Level | 27224
- Output Fields | 27224
- Sample Output | 27226
- Release Information | 27229

Syntax

```
show system subscriber-management route
<family family>
<incomplete>
<level (brief | detail)>
<next-hop index>
<prefix>
<routing-instance name>>
<route-type type>
<rrt-index index>
<summary>
```

Description

Display information about how routes are mapped to specific enhanced subscriber management interfaces. You can customize the output by including one or more optional filters in the command. With the exception of the `summary` option, all filter options can be combined in a single command.

Options

family *family*	(Optional) Display route mapping information for the specified protocol family: `inet` (IPv4) or `inet6` (IPv6).
incomplete	(Optional) Display route mapping information for incomplete routes that are missing elements required to add the routes to the routing table.
level (brief | detail)	(Optional) Display the specified level of output: `brief` or `detail`.
next-hop *index*	(Optional) Display the next hop associated with the route entry with the specified next-hop index, in the range 1 through 65535.
prefix *address*	(Optional) Use the same prefix and prefix length as the subscriber host address. Output includes attributes that originate in the Famed-Route record of an upstream RADIUS server (Tag, Metric, Preference).
route-type *type*	(Optional) Display route mapping information for the specified route type: `access`, `access-internal`, `kernel`, or `local`.
routing-instance *name*	(Optional) Display route mapping information for the specified routing-instance

rrt-index *index* (Optional) Display mapping information for the specified routing table index, in the range 0 through 65535. An `rtt-index` value of 0 (zero) denotes routes in the default routing table managed by enhanced subscriber management.

summary (Optional) Display summary information about the routes managed by enhanced subscriber management.

Required Privilege Level

view

Output Fields

Table 1688 on page 27224 lists the output fields for the `show system subscriber-management route` command. Output fields are listed in the approximate order in which they appear.

Table 1688: show system subscriber-management route Output Fields

Field Name	Field Description	Level of Output
address	IPv4 or IPv6 address associated with the route entry.	All levels
`Route Type`	One of the following route types: • `Access` • `Access-internal` • `Framed` • `Kernel` • `Local`	All levels
`Interface`	Name of the enhanced subscriber management interface associated with the route entry.	All levels
`Next-hop`	Next-hop associated with the route entry.	All levels

Table 1688: show system subscriber-management route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Tag	Reflects the Tag attribute used in the RADIUS Framed-Route type record.	All levels
Metric	Reflects the Metric attribute used in the RADIUS Framed-Route type record.	All levels
Preference	Reflects the Preference attribute used in the RADIUS Framed-Route type record.	All levels
Rtt-index	Value of the routing table index. A value of 0 (zero) denotes a route in the default routing table managed by enhanced subscriber management.	detail
Bbe index	Value of the interface index for the control plane.	detail
Flow id	Value of the route object index.	detail
Reference Count	Used for internal accounting.	detail
Dirty Flags	Used for internal accounting.	detail
Flags	Used for internal accounting.	detail
Family	One of the following protocol families: • AF_INET—IPv4 • AF_INET6—IPv6	detail

Sample Output

show system subscriber-management route prefix <address>

rtt-index 0

```
user@host> show system subscriber-management route prefix 10.10.0.1/32
Route:  10.10.0.1/32
     Routing-instance:         default:default
     Kernel rt-table id :      0
     Family:                   AF_INET
     Route Type:               Framed
     Protocol Type:            Unspecified
     Interface:                pp0.3221225491
     Interface index:          26
     Internal Interface index: 26
     Route index:              20
     Next-Hop:                 684
     Tag:                      9999
     Metric:                   56
     Preference:               10
     Reference-count:          1
     L2 Address:               00:00:5e:00:53:0b
     Flags:                    0x0
     Dirty Flags:              0x0
```

show system subscriber-management route family route-type rtt-index level brief

The following example displays abbreviated information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show system subscriber-management route family inet6 route-type access rtt-index 0
level brief
2001:db8::/64
     Route Type: Access
     Interface: pp0.3221225479, Next-hop:721
2001:db8:0:0:1::/64
     Route Type: Access
     Interface: pp0.3221225477, Next-hop:721
2001:db8:0:0:2::/64
```

```
     Route Type: Access
     Interface: pp0.3221225478, Next-hop:721
2001:db8:0:0:3::/64
     Route Type: Access
     Interface: pp0.3221225480, Next-hop:721
2001:db8:0:0:4::/64
     Route Type: Access
     Interface: pp0.3221225481, Next-hop:721
2001:db8:2002::/84
     Route Type: Access
     Interface: demux0.3221225492, Next-hop:721
2001:db8:0:0:5::/64
     Route Type: Access
     Interface: pp0.3221225487, Next-hop:721
2001:db8:0:0:6::/64
     Route Type: Access
```

show system subscriber-management route family route-type rtt-index level detail

The following example displays detailed information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show system subscriber-management route family inet6 route-type access rtt-index 0
level detail
2001:db8::/64
     Route Type:      Access
     Interface:       pp0.3221225479
     Next-hop:        721
     Rtt-index:       0
     Bbe index:       9
     Flow id:         1
     Reference Count: 1
     Dirty Flags:     0
     Flags:           0x10082
     Family:          AF_INET6
2001:db8:0:0:1::/64
     Route Type:      Access
     Interface:       pp0.3221225477
     Next-hop:        721
     Rtt-index:       0
     Bbe index:       9
```

```
     Flow id:         1
     Reference Count: 1
     Dirty Flags:     0
     Flags:           0x10082
     Family:          AF_INET6
2001:db8:0:0:2::/64
     Route Type:      Access
     Interface:       pp0.3221225478
     Next-hop:        721
     Rtt-index:       0
     Bbe index:       9
     Flow id:         1
     Reference Count: 1
     Dirty Flags:     0
     Flags:           0x10082
     Family:          AF_INET6
2001:db8:0:0:3::/64
     Route Type:      Access
     Interface:       pp0.3221225480
     Next-hop:        721
     Rtt-index:       0
     Bbe index:       9
     Flow id:         1
     Reference Count: 1
     Dirty Flags:     0
     Flags:           0x10082
     Family:          AF_INET6
```

show system subscriber-management route family route-type rtt-index level brief

The following example displays abbreviated information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show system subscriber-management route family inet6 route-type access rtt-index 0
level brief
2001:db8::/64
     Route Type: Access
     Interface: pp0.3221225479, Next-hop:721
2001:db8:0:0:1::/64
     Route Type: Access
     Interface: pp0.3221225477, Next-hop:721
```

```
2001:db8:0:0:2::/64
     Route Type: Access
     Interface: pp0.3221225478, Next-hop:721
2001:db8:0:0:3::/64
     Route Type: Access
     Interface: pp0.3221225480, Next-hop:721
2001:db8:0:0:4::/64
     Route Type: Access
     Interface: pp0.3221225481, Next-hop:721
2001:db8:2002::/84
     Route Type: Access
     Interface: demux0.3221225492, Next-hop:721
2001:db8:0:0:5::/64
     Route Type: Access
     Interface: pp0.3221225487, Next-hop:721
2001:db8:0:0:6::/64
     Route Type: Access
```

Release Information

Command introduced in Junos OS Release 15.1R3.

Support for passing **Framed-Route** attributes from a RADIUS server to the router was added in Junos OS Release 17.2 on MX Series routers for enhanced subscriber management. This allows the tagged subscriber host routes to be imported to the routing table and advertised by BGP.

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

show system subscriber-management statistics

IN THIS SECTION

- Syntax | 27230
- Description | 27230

- Options | 27230
- Required Privilege Level | 27231
- Output Fields | 27231
- Sample Output | 27232
- Release Information | 27241

Syntax

```
show system subscriber-management statistics
<all>
<dhcp>
<dvlan>
<fixed-wireless-access>
<l2tp>
<ppp>
<pppoe>
```

Description

Display statistics for the specified option. You can customize the output by including one or more optional filters in the command. With the exception of the `extensive` option, all filter options can be combined in a single command.

Options

all	(Optional) Display packet statistics for all protocols.
dhcp	(Optional) Display DHCP packet statistics.
dvlan	(Optional) Display DVLAN packet statistics.
fixed-wireless-access	(Optional) Display fixed wireless access packet statistics.
l2tp	(Optional) Display L2TP packet statistics.
ppp	(Optional) Display PPP packet statistics.

pppoe (Optional) Display PPPoE packet statistics.

Required Privilege Level

view

Output Fields

Table 1689 on page 27231 lists the output fields for the `show system subscriber-management statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1689: show system subscriber-management statistics Output Fields

Field Name	Field Description
`Rx Statistics`	Statistics for packets received.
`Tx Statistics`	Statistics for packets sent.
`Enhanced I/O Statistics`	Statistics for visibility into packet drops from the queue.
`Error Statistics`	Includes connection packets, flow control, and messages and packets sent to and received from the daemon.
`ERA discards`	Event Rate Analyzer discards. For DHCP and PPPoE in advanced subscriber management, ERA packet discard counts are included for Discover, Solicit, and PADI packets .
`Layer 3 Statistics`	Statistics for Layer 3 packets.
`padis`	PPPoE Active Discovery Initiation (PADI) packets. PADI is the first step in the PPPoE establishment protocol.
`padrs`	PPPoE Active Discovery Request packets.
`ppp`	Point-to-Point Protocol packets.

Table 1689: show system subscriber-management statistics Output Fields *(Continued)*

Field Name	Field Description
`router solicitations`	Number of router solicitations sent or received. Router solicitations are sent to prompt all on-link routers to send it router advertisements.
`router advertisements`	Number of router advertisements sent or received.
`route solicit response packet`	Number of router solicitation responses sent or received.

Sample Output

The following examples displays packet statistics accumulated for DHCP, hybrid access, and PPPoE since the last time the session manager was cleared.

show system subscriber-management statistics all

```
user@host> show system subscriber-management statistics all
  user@host> show system subscriber-management statistics all
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
-------------------------------------------------------------
                     Packet Statistics
-------------------------------------------------------------
I/O Statistics:
-------------------------------------------------------------
    Rx Statistics
        packets                         : 784711
    Tx Statistics
        packets                         : 7013122
  Layer 3 Statistics
    Rx Statistics
        packets                         : 356218
    Tx Statistics
        packets                         : 6604660

DHCP Statistics:
```

```
-------------------------------------------------------------
    Rx Statistics
        packets                          : 320008
        ERA discards                     : 6274
    Tx Statistics
        transmit request packets         : 320482
        sent packets                     : 320482
    Error Statistics
    Connection Statistics
        no connection packets            : 0

PPPoE Statistics:
-------------------------------------------------------------
    Rx Statistics
        packets                          : 486165
        padis                            : 36768
        padrs                            : 35421
        ppp packets                      : 341787
        ERA discards                     : 8249
    Tx Statistics
        packets                          : 70842
        send failures                    : 6240
```

show system subscriber-management statistics dhcp

```
user@host> show system subscriber-management statistics dhcp
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
-------------------------------------------------------------
                      Packet Statistics
-------------------------------------------------------------
I/O Statistics:
-------------------------------------------------------------
    Rx Statistics
        packets                          : 784711
    Tx Statistics
        packets                          : 7013122
  Layer 3 Statistics
    Rx Statistics
        packets                          : 356218
    Tx Statistics
```

```
        packets                          : 6604660

DHCP Statistics:
---------------------------------------------------------------
    Rx Statistics
        packets                          : 320008
        ERA discards                     : 6274
    Tx Statistics
        transmit request packets         : 320482
        sent packets                     : 320482
    Error Statistics
    Connection Statistics
        no connection packets            : 0
```

show system subscriber-management statistics dhcp extensive

```
user@host> show system subscriber-management statistics dhcp extensive
Session Manager started @ Mon Dec  6 06:14:27 2021
Session Manager cleared @ Mon Dec  6 06:14:27 2021
---------------------------------------------------------------
                     Packet Statistics
---------------------------------------------------------------
I/O Statistics:
---------------------------------------------------------------
    Rx Statistics
        packets                          : 7986
    Tx Statistics
        packets                          : 15
          l2 inject                      : 15
          l3 inject                      : 0
    Buffer Statistics
        allocations                      : 7990
        frees                            : 7990
        allocation failures              : 0
  Layer 3 Statistics
    Rx Statistics
        packets                          : 29
    Tx Statistics
        packets                          : 15
    PFE Event Statistics
        packets                          : 7957
```

```
-------------------------------------------------------------
Enhanced I/O Statistics:
-------------------------------------------------------------
    bbe_io_rcv l2                        : 0
    bbe_io_rcv l3                        : 29
    bbe_io_rcv l3 v4                     : 17
    bbe_io_rcv l3 v6                     : 12
    bbe_io_rcv l3 unspec                 : 0
    bbe_io_rcv l3 unknown af             : 0
    bbe_io_rcv routed                    : 0
    bbe_io_rcv routed_v4                 : 0
    bbe_io_rcv routed_v6                 : 0
    bbe_io_rcv routed_no_route           : 0
    bbe_io_rcv routed_default            : 0
    bbe_io_rcv resolve v4                : 0
    bbe_io_rcv resolve v6                : 0
    bbe_io_rcv resolve default           : 0
    bbe_io_rcv pfe event                 : 0
    bbe_io_rcv default                   : 0
    -----------------------------------------------------------
    rx inet                              : 17
    rx inet6                             : 12
    rx inet6 icmp6                       : 2
    rx inet igmp                         : 0
    rx inet gre                          : 0
    rx inet agf                          : 0
    rx inet icmp                         : 0
    rx inet udp v4                       : 17
    rx inet udp v6                       : 10
          rx inet udp l2tp               : 0
          rx inet udp jdhcp v4           : 17
          rx inet udp jdhcp v6           : 10
          rx inet udp bfd v4             : 0
          rx inet udp bfd v6             : 0
    rx inet tcp proxy v4                 : 0
    rx inet tcp proxy v6                 : 0
    rx v4 ip frag reassm proc cnt        : 0
    rx v4 ip frag reassm pkt cnt         : 0
    rx v4 ip frag reassm alloc cntx cnt  : 0
    rx v4 ip frag reassm free cntx cnt   : 0
    tx l3 forward                        : 15
    tx udp v4                            : 5
    tx udp v6                            : 6
```

```
    tx tcp proxy client v4              : 0
    tx tcp proxy client v4 error        : 0
    tx tcp proxy v4 drop                : 0
    tx tcp proxy v6 drop                : 0
    tx igmp                             : 0
    tx v6 l3 route forwards             : 0
    tx v6 l3 route drops                : 0
    tx v6 kernel forwards               : 0
    tx v6 kernel forward drops          : 0
    tx v6 l2 forwards                   : 10
    tx v6 l2 drops                      : 0
    tx v4 l3 route forwards             : 0
    tx v4 l3 route drops                : 0
    tx v4 kernel forwards               : 0
    tx v4 kernel forward drops          : 0
    tx v4 l2 forwards                   : 5
    tx v4 l2 drops                      : 0
    tx v4 ip fragment forward           : 0
    tx v4 ip fragment DF drops          : 0
    tx v4 ip fragment drops             : 0
    tx v4 ip fragment failed            : 0
    bbe_ifl_output tx failed            : 0
    bbe_io_send tx failed               : 0
    bbe_io_send tx partial failed       : 0
    io_queue low_prio_packets_dropped   : 0
    io_queue mlow_prio_packets_dropped  : 0
    io_queue med_prio_packets_dropped   : 0
    io_queue high_prio_packets_dropped  : 0
-------------------------------------------------------------
DHCP Statistics:
-------------------------------------------------------------
    Rx Statistics
        packets                         : 27
    Tx Statistics
        transmit request packets        : 11
        sent packets                    : 11
    DHCPv4 Rx Statistics
        total packets                   : 17
        boot request                    : 15
        boot reply                      : 2
        discover                        : 11
        offer                           : 1
        request                         : 3
```

```
        ack                             : 1
        release                         : 1
    DHCPv4 Tx Statistics
        total packets                   : 5
        boot reply                      : 5
        offer                           : 2
        ack                             : 3
    DHCPv6 Rx Statistics
        total packets                   : 10
        solicit                         : 4
        advertise                       : 1
        request                         : 2
        reply                           : 1
        renew                           : 1
        release                         : 1
        relay repl                      : 2
    DHCPv6 Tx Statistics
        total packets                   : 6
        advertise                       : 2
        reply                           : 4
    Error Statistics
    Connection Statistics
        no connection packets           : 0
        connection down events          : 1
        connection up events            : 2
        flow control invoked            : 0
        flow control released           : 0
        packets sent to daemon          : 27
        packets received from daemon    : 11
        messages sent to daemon         : 0
        messages received from daemon   : 1481
        notifies while not connected    : 0
NET Statistics:
---------------------------------------------------------------
  ICMP6 Statistics
    Rx Statistics
        packets:                        : 2
        neighbor advertisements         : 2
    Tx Statistics
        packets:                        : 4
        neighbor solicitations          : 4
Management Statistics:
---------------------------------------------------------------
```

```
    ipdemux                              : 4
    ipdemux add                          : 4
    ipdemux delete                       : 2
    ifl flow                             : 7
    ifl flow adds                        : 6
    ifl flow changes                     : 1
    ifl flow deletes                     : 3
    ip flow                              : 9
    ip flow add                          : 6
    ip flow delete                       : 3
    service                              : 27
    service adds                         : 24
    service deletes                      : 3
Management Config Status:
---------------------------------------------------------------
    gres state enabled state             : 0
    shmlog disabled state                : 0
    shmlog filtering state               : 0
    vc backup member local switch state  : 0
    dscp code point value                : 0x30
```

show system subscriber-management statistics pppoe

```
user@host> show system subscriber-management statistics pppoe
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
---------------------------------------------------------------
                     Packet Statistics
---------------------------------------------------------------
I/O Statistics:
---------------------------------------------------------------
    Rx Statistics
        packets                          : 784711
    Tx Statistics
        packets                          : 7013122
  Layer 3 Statistics
    Rx Statistics
        packets                          : 356218
    Tx Statistics
        packets                          : 6604660
```

```
PPPoE Statistics:
-------------------------------------------------------------
    Rx Statistics
        packets                          : 486165
        padis                            : 36768
        padrs                            : 35421
        ppp packets                      : 341787
        ERA discards                     : 8249
    Tx Statistics
        packets                          : 70842
        send failures                    : 6240
```

show system subscriber-management statistics extensive

```
user@host> show system subscriber-management statistics extensive
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
-------------------------------------------------------------
                      Packet Statistics
-------------------------------------------------------------
I/O Statistics:
-------------------------------------------------------------
    Rx Statistics
        packets                          : 784711
    Tx Statistics
        packets                          : 7013122
    Buffer Statistics
        allocations                      : 7032618
        frees                            : 7032624
        allocation failures              : 0
  Layer 3 Statistics
    Rx Statistics
        packets                          : 356218
    Tx Statistics
        packets                          : 6604660
    PFE Event Statistics
        packets                          : 0

-------------------------------------------------------------
Enhanced I/O Statistics:
-------------------------------------------------------------
```

```
    bbe_io_rcv l2                        : 0
    bbe_io_rcv l3                        : 0
    bbe_io_rcv l3 v4                     : 0

    io low queue drops                   :12
    io mlow queue drops                  :0
    io medium queue drops                :0
    io high queue drops                  :0
```

show system subscriber-management statistics ppp (LCP Vendor-Specific Counters)

```
user@host> show system subscriber-management statistics ppp
Session Manager started @ Thu Feb 11 00:37:43 2020
Session Manager cleared @ Thu Feb 11 00:37:43 2020
-------------------------------------------------------------
                     Packet Statistics
-------------------------------------------------------------
I/O Statistics:
-------------------------------------------------------------
    Rx Statistics
        packets                          : 486783
    Tx Statistics
        packets                          : 144
  Layer 3 Statistics
    Rx Statistics
        packets                          : 8
    Tx Statistics
        packets                          : 0
PPP Statistics:
-------------------------------------------------------------
    Rx Statistics
        network packets                  : 123
        plugin packets                   : 123
        lcp config requests              : 18
        lcp config acks                  : 18
        lcp conf nacks                   : 8
        lcp conf rejects                 : 6
        lcp termination requests         : 4
        lcp termination acks             : 13
        lcp code rejects                    : 2
        lcp vendor-specific acks       : 10
```

```
        pap requests                      : 8
        ipcp requests                     : 27
        ipcp acks                         : 9
        ipv6cp requests                   : 11
        ipv6cp acks                       : 1
    Tx Statistics
        packets                           : 101
        lcp config requests               : 32
        lcp config acks                   : 18
        lcp termination requests          : 13
        lcp termination acks              : 4
        lcp vendor-specific requests   : 10
        pap acks                          : 8
        ipcp requests                     : 9
        ipcp acks                         : 5
        ipcp nacks                        : 9
        ipv6cp requests                   : 1
        ipv6cp acks                       : 1
        ipv6cp nacks                      : 1
NET Statistics:
-------------------------------------------------------------
  ICMP6 Statistics
    Rx Statistics
        packets:                          : 8
        router solicitations              : 8
    Tx Statistics
        packets:                          : 0
```

Release Information

Command introduced in Junos OS Release 15.1R3.

Enhanced I/O Statistics introduced as part of Extensive output in Junos OS Release 15.1R4 on MX Series routers for enhanced subscriber management.

RELATED DOCUMENTATION

Understanding Dropped Packets and Untransmitted Traffic Using show Commands

show system subscriber-management summary

IN THIS SECTION

- Syntax | 27242
- Description | 27242
- Options | 27242
- Required Privilege Level | 27242
- Output Fields | 27242
- Sample Output | 27245
- Release Information | 27246

Syntax

```
show system subscriber-management summary
```

Description

Display complete subscriber management database summary information.

Options

none This command has no options.

Required Privilege Level

view

Output Fields

Table 1690 on page 27243 lists the output fields for the `show system subscriber-management summary` command. Output fields are listed in the approximate order in which they appear.

Table 1690: show system subscriber-management summary Output Fields

Field Name	Field Description
`Graceful Restart`	State of graceful Routing Engine switchover (GRES): • `Enabled` • `Disabled` (Enhanced subscriber management for MX Series routers) The name of this field is `Graceful Switchover`.
`Mastership`	State of the Routing Engine: • `Master` • `Standby`
`Database`	State of the subscriber management database: • `Available` • `Init` • `Not-available`
`Standby`	(Enhanced subscriber management for MX Series routers) State of the standby Routing Engine: • `Connected`—Connected but not synchronized • `Disconnected`—Not connected • `Resync (`*nn*`%)`—Connected and *nn* percent synchronized with the primary Routing Engine • `Synchronized`—Synchronized with the primary Routing Engine
`Disconnection Reason`	Reason why both Routing Engines are disconnected when there is a DRAM mismatch. • Primary/Standy RE DRAM Size Mismatch—Displayed when the amount of memory is different on the primary and standby Routing Engines.

Table 1690: show system subscriber-management summary Output Fields *(Continued)*

Field Name	Field Description
`Chassisd ISSU State`	State of unified ISSU chassis daemon: • `ABORT` • `DAEMON_ISSU_PREPARE` • `DAEMON_ISSU_PREPARE_DONE` • `DAEMON_SWITCHOVER_PREPARE` • `DAEMON_SWITCHOVER_PREPARE_DONE` • `FRU_ISSU` • `FRU_ISSU_DONE` • `IDLE` • `UNKNOWN`
`ISSU State`	State of unified ISSU aggregate daemon: • `ABORT` • `IDLE` • `PREPARE` • `READY` • `SWITCHOVER_PREPARE` • `SWITCHOVER_READY` • `UNKNOWN`
`ISSU Wait`	Amount of time, in seconds, requested by a daemon to perform cleanup. If multiple daemons request time, the displayed value is the highest wait time requested by a daemon.

Sample Output

show system subscriber-management summary

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Chassisd ISSU State  DAEMON_ISSU_PREPARE
    ISSU State           PREPARE
    ISSU Wait            198
```

show system subscriber-management summary (Enhanced Subscriber Management)

```
user@host> show system subscriber-management summary
General:
    Graceful Switchover     Enabled
    Mastership              Master
    Database                Available
    Standby                   Resync (75%)
    Chassisd ISSU State     IDLE
    ISSU State              IDLE
    ISSU Wait               0
```

show system subscriber-management summary (DRAM Size Mismatch Error)

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Standby              Disconnected
<emphasis>    Disconnection Reason    Master/Standy RE DRAM Size Mismatch</emphasis
>
    Chassisd ISSU State  IDLE
    ISSU State           IDLE
    ISSU Wait            0
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show database-replication statistics | 19183

show database-replication summary | 19185

show system switchover

IN THIS SECTION

- Syntax | 27246
- Syntax (MX Series Router) | 27246
- Description | 27247
- Options | 27247
- Additional Information | 27247
- Required Privilege Level | 27247
- Output Fields | 27248
- Sample Output | 27249
- Release Information | 27251

Syntax

```
show system switchover
```

Syntax (MX Series Router)

```
show system switchover
<all-members>
```

```
<local>
<member member-id>
```

Description

Display whether graceful Routing Engine switchover is configured, the state of the kernel replication (ready or synchronizing), any replication errors, and whether the primary and standby Routing Engines are using compatible versions of the kernel database.

> **NOTE**: Issue the `show system switchover` command *only* on the backup Routing Engine. This command is *not* supported on the primary Routing Engine because the `kernel-replication` process daemon does not run on the primary Routing Engine. This process runs only on the backup Routing Engine.

Options

all-members (MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on all members of the Virtual Chassis configuration.

local (MX Series routers only) (Optional) Display graceful Routing Engines switchover information for all Routing Engines on the local Virtual Chassis member.

member *member-id* (MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

Additional Information

If you issue the `show system switchover` command on the active Routing Engine in the primary router of an MX Series Virtual Chassis, the router displays a message that this command is not applicable on this member of the Virtual Chassis.

Required Privilege Level

view

Output Fields

Table 1691 on page 27248 describes the output fields for the `show system switchover` command. Output fields are listed in the approximate order in which they appear.

Table 1691: show system switchover Output Fields

Field Name	Field Description
`Graceful switchover`	Display graceful Routing Engine switchover status: • `On`—Indicates `graceful-switchover` is specified for the `routing-options` configuration command. • `Off`—Indicates `graceful-switchover` is not specified for the `routing-options` configuration command.
`Configuration database`	State of the configuration database: • `Ready`—Configuration database has synchronized. • `Synchronizing`—Configuration database is synchronizing. Displayed when there are updates within the last 5 seconds. • `Synchronize failed`—Configuration database synchronize process failed.
`Kernel database`	State of the kernel database: • `Ready`—Kernel database has synchronized. This message implies that the system is ready for GRES. • `Synchronizing`—Kernel database is synchronizing. Displayed when there are updates within the last 5 seconds. • `Version incompatible`—The primary and standby Routing Engines are running incompatible kernel database versions. • `Replication error`—An error occurred when the state was replicated from the primary Routing Engine. Inspect **Steady State** for possible causes, or notify Juniper Networks customer support.

Table 1691: show system switchover Output Fields *(Continued)*

Field Name	Field Description
`Switchover Status`	Switchover Status: • `Ready`—Message for system being switchover ready. • `Not Ready`—Message for system not being ready for switchover.
`Platform Components`	Platform Components: FEB1: • `Not Online`—The backup FEB2 is not yet online. • `Online`—The backup FEB2 is online. FEB1-PFE0, FEB1- PFE1: • `Ready`—Message for FEB1- PFE0, or FEB1- PFE1 being switchover ready. The backup PFE is ready for switchover (does not include time for routes or nexthops for a scaled configuration). • `Not Online`—The backup PFE is not yet online. • `Sync In Progress`—Two-minute synchronization in progress when information flow data-path is programmed on the backup PFE (does not include time for routes or next-hops programming).

Sample Output

show system switchover (Backup Routing Engine - Ready)

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Status: Ready
```

show system switchover (Backup Routing Engine - Not Ready)

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Status: Not Ready
```

show system switchover (ACX7509)

The switchover status option for ACX7509 is only available on the backup Routing Engine. The show outputs display the system switchover status as "Ready" and "Not Ready".

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Object database: Ready
Applications' ready state: Ready
Switchover Status: Ready
Platform Components: Ready
FEB1 Online
FEB1-PFE0 Ready
FEB1-PFE1 Ready
```

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Object database: Ready
Applications' ready state: Ready
Switchover Status: Not Ready
Platform Components: Not Ready
FEB1 Not Online
FEB1-PFE0 Not Online
FEB1-PFE1 Not Online
```

show system switchover (Backup Routing Engine - Junos OS Evolved)

```
{backup}
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Object database: Ready
Applications' ready state: Ready
Switchover Status: Not Ready
```

Release Information

Command introduced before Junos OS Release 7.4.

Command support added for ACX7509 in 22.1R1 for High Availability Platform Redundancy RCB and FEB switchover.

RELATED DOCUMENTATION

request chassis routing-engine master | 14529

show system unified-services status

IN THIS SECTION

- Syntax | 27252
- Description | 27252
- Required Privilege Level | 27252
- Output Fields | 27252
- Sample Output | 27252
- Release Information | 27252

Syntax

```
show system unified-services status
```

Description

Determine whether Next Gen Services is enabled or disabled on the MX.

Required Privilege Level

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

show system unified-services status

```
user@host> show system unified-services status
```

One of the following four messages appears:

```
Enabled
Unified Services : Upgrade staged , please reboot with 'request system reboot' to enable unified
services.
Disabled
Unified Services : Upgrade staged , please reboot with 'request system reboot' to disable
unified services.
```

Release Information

Command introduced in Junos OS Release 19.3R1.

show system uptime

IN THIS SECTION

- Syntax | 27253
- Syntax (EX Series Switches) | 27253
- Syntax (QFX Series) | 27254
- Syntax (TX Matrix Router) | 27254
- Syntax (TX Matrix Plus Router) | 27254
- Syntax (MX Series Router) | 27254
- Syntax (Junos OS Evolved) | 27254
- Description | 27255
- Options | 27255
- Additional Information | 27257
- Required Privilege Level | 27257
- Output Fields | 27257
- Sample Output | 27258
- Release Information | 27259

Syntax

```
show system uptime
```

Syntax (EX Series Switches)

```
show system uptime
<all-members>
<local>
<member member-id>
```

Syntax (QFX Series)

```
show system uptime
<director-group name>
<infrastructure name>
<interconnect-device name>
<node-group name>
```

Syntax (TX Matrix Router)

```
show system uptime
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system uptime
<detail>
<all-chassis | all-lcc | lcc number | sfc number>
```

Syntax (MX Series Router)

```
show system uptime
<all-members>
<invoke-on>
<local>
<member member-id>
```

Syntax (Junos OS Evolved)

```
show system uptime
<node node-name>
```

Description

Display the current time and information about how long the router or switch, router or switch software, and routing protocols have been running.

> **NOTE**: Time values computed from differences in timestamps can vary due to the insertion or deletion of leap-seconds between them.

The `show system uptime` command is a little different in how it displays output in Junos OS Evolved. The `show system uptime` command by itself shows system-wide uptime information. Use the `show system uptime node` *`node-name`* command to see node-specific uptime information, where *`node-name`* can be `re0 | re1 | fpc0 | all`.

Options

none Show time since the system rebooted and processes started.

all-chassis (TX Matrix routers and TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started on all the routers in the chassis.

all-lcc (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, show time since the system rebooted and processes started for all connected T1600 or T4000 LCCs.

all-members (EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on all members of the Virtual Chassis configuration.

director-group *name* (QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Director group.

infrastructure *name* (QFabric systems only) (Optional) Show time since the system rebooted and processes started on the fabric control Routing Engine and fabric manager Routing Engine.

interconnect-device *name* (QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Interconnect device.

invoke-on (MX Series routers only) (Optional) Display the time since the system rebooted and processes started on the primary Routing Engine, backup Routing Engine, or both, on a router with two Routing Engines.

lcc *number* (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show time since the system rebooted and processes started for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the local Virtual Chassis member.

member *member-id* (EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node *node-name* (Junos OS Evolved only) (Optional) Display uptime information for the specified node. When node option is not used, display uptime information for entire system. From the system's point of view, the system booted time is the time when the system was assembled and is available in **/var/run/system_boot_time**.

node-group *name* (QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Node group.

scc (TX Matrix routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix router (or switch-card chassis).

sfc *number* (TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system uptime` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Output Fields

Table 1692 on page 27257 describes the output fields for the `show system uptime` command. Output fields are listed in the approximate order in which they appear.

Table 1692: show system uptime Output Fields

Field Name	Field Description
`System booted`	(Only for Junos OS Evolved when `node` option is not used) Time system was last booted. From the system's point of view, the system booted time is the time when the system was assembled and is available in **/var/run/system_boot_time**.
`System-wide users`	(Only for Junos OS Evolved when `node` option is not used) Number of system-wide users.
`Current time`	Current system time in UTC.
`Time Source`	Reference time source that the system is locked to.
`System booted`	Date and time when the Routing Engine on the router or switch was last booted and how long it has been running.
`Protocols started`	Date and time when the routing protocols were last started and how long they have been running.

Table 1692: show system uptime Output Fields *(Continued)*

Field Name	Field Description
`Last configured`	Date and time when a configuration was last committed. Also shows the name of the user who issued the last `commit` command.
`time` `and up`	Current time, in the local time zone, and how long the router or switch has been operational.
`users`	Number of users logged in to the router or switch.
`load averages`	Load averages for the last 1 minute, 5 minutes, and 15 minutes.

Sample Output

show system uptime

```
user@host> show system uptime
Current time:      2017-10-13 19:45:47 UTC
Time Source:       NTP CLOCK
System booted:     2017-10-12 20:51:41 UTC (22:54:06 ago)
Protocols started: 2017-10-13 19:33:45 UTC (00:12:02 ago)
Last configured:   2017-10-13 19:33:45 UTC (00:12:02 ago) by abc
12:45PM  up 22:54, 2 users, load averages: 0.07, 0.02, 0.01
```

show system uptime node (for Junos OS Evolved)

```
user@host> show system uptime node re0

Current time: 2017-05-16 16:13:18 PDT
Node booted: 2017-05-10 15:45:29 PDT (6d 00:27 ago)
Last configured: 2017-05-10 15:31:46 PDT (6d 00:41 ago) by root
16:13:18 up 6 days, 27 min, 1 user, load averages: 2.69, 2.58, 2.57
```

show system uptime (QFX Series)

```
user@switch> show system uptime
Current time: 2017-08-27 03:12:30 PDT
Time Source: NTP CLOCK
System booted: 2017-08-13 17:11:54 PDT (1w6d 10:00 ago)
Protocols started: 2017-08-13 17:13:56 PDT (1w6d 09:58 ago)
Last configured: 2017-08-26 05:54:00 PDT (21:18:30 ago) by user
 3:12AM  up 13 days, 10:01, 3 users, load averages: 0.00, 0.00, 0.00
```

show system uptime (Junos OS Evolved)

```
user@router> show system uptime

System booted: 2019-02-20 19:17:52 PST (02:20:33 ago)
System-wide users: 7 users
```

Starting in Junos OS Evolved 19.1R1 release, the `show system uptime` output displays only the `System booted` and `System-wide users` information. The output does not display the details such as `Current time`, `System booted`, `Protocols started`, and `Last configured parameters`. To see node-specific uptime information, use the `node` option.

show system uptime node re0 (Junos OS Evolved)

```
user@router> show system uptime node re0
Current time: 2019-07-09 14:24:34 PDT
Time Source:  NTP CLOCK
Node booted: 2019-07-03 09:59:58 PDT (6d 04:24 ago)
Protocols started: 2019-07-03 10:01:41 PDT (6d 04:22 ago)
Last configured: 2019-07-03 10:01:10 PDT (6d 04:23 ago) by root
 2:24PM  up 6 days, 4:24, 1 user, load averages: 1.25, 0.51, 0.36
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced for the TX Matrix Plus router in JUNOS Release 9.6.

`node` option introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

- 10-Gigabit Ethernet LAN/WAN PIC with XFP (T640 Router)
- Routing Matrix with a TX Matrix Plus Router Solutions Page

show system users

IN THIS SECTION

- Syntax | 27260
- Syntax (TX Matrix Router) | 27260
- Syntax (TX Matrix Plus Router) | 27261
- Syntax (MX Series Router) | 27261
- Description | 27261
- Options | 27261
- Additional Information | 27262
- Required Privilege Level | 27262
- Output Fields | 27263
- Sample Output | 27263
- Release Information | 27266

Syntax

```
show system users
<no-resolve>
```

Syntax (TX Matrix Router)

```
show system users
<all-chassis | all-lcc | lccnumber | scc>
<no-resolve>
```

Syntax (TX Matrix Plus Router)

```
show system users
<detail>
<all-chassis | all-lcc | lcc number | sfc number> <no-resolve>
```

Syntax (MX Series Router)

```
show system users
<all-members>
<local>
<member member-id>
<no-resolve>
```

Description

List information about the users who are currently logged in to the router or switch.

> **NOTE**: The `show system users` command lists the information about administrative users that are logged in to a router or switch using the CLI, J-Web, or an SSH client. The output does not list information about web users or automated users that are logged in from a remote client application using Junos XML APIs, such as NETCONF.

Options

none List information about the users who are currently logged in to the router or switch.

all-chassis (TX Matrix routers and TX Matrix Plus routers only) (Optional) Show users currently logged in to all the routers in the chassis.

all-lcc (TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show users currently logged in to all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, show users currently logged in to all connected T1600 or T4000 LCCs.

all-members (MX Series routers only) (Optional) Display users currently logged in to all members of the Virtual Chassis configuration.

lcc *number*

(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show users currently logged in to a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show users currently logged in to a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(MX Series routers only) (Optional) Display users currently logged in to the local Virtual Chassis member.

member *member-id*

(MX Series routers only) (Optional) Display users currently logged in to the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

no-resolve

(Optional) Do not attempt to resolve IP addresses to hostnames.

scc

(TX Matrix routers only) (Optional) Show users currently logged in to the TX Matrix router (or switch-card chassis).

sfc *number*

(TX Matrix Plus routersonly) (Optional) Show users currently logged in to the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system users` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Output Fields

Table 1693 on page 27263 describes the output fields for the `show system users` command. Output fields are listed in the approximate order in which they appear.

Table 1693: show system users Output Fields

Field Name	Field Description
time and up	Current time, in the local time zone, and how long the router or switch has been operational.
users	Number of users logged in to the router or switch.
load averages	Load averages for the last 1 minute, 5 minutes, and 15 minutes.
USER	Username.
TTY	Terminal through which the user is logged in.
FROM	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.
LOGIN@	Time when the user logged in.
IDLE	How long the user has been idle.
WHAT	Processes that the user is running.

Sample Output

show system users

```
user@host> show system users
 7:30PM  up 4 days,  2:26, 2 users, load averages: 0.07, 0.02, 0.01
USER     TTY FROM              LOGIN@  IDLE WHAT
```

```
root     d0  -                Fri05PM 4days -csh (csh)
blue   p0 level5.company.net 7:30PM     - cli
```

show system users lcc no-resolve (TX Matrix, TX Matrix Plus Router)

```
user@host> show system users lcc 2 no-resolve

lcc2-re0:
--------------------------------------------------------------------------
10:34AM PDT up 1 day,  7:11, 5 users, load averages: 0.03, 0.01, 0.00
USER     TTY      FROM                              LOGIN@  IDLE WHAT
root     d0       -                                3:21AM   7:12 /bin/csh
user1    p0       scc-re0                          10:15AM     - telnet hostA
user1    p1       scc-re0                          10:16AM     - telnet hostA
user1    p2       scc-re0                          10:19AM     - telnet hostA
user1    p3       scc-re0                          10:24AM     - telnet hostA
```

show system users (TX Matrix Plus Router)

```
user@host> show system users
sfc0-re0:
--------------------------------------------------------------------------
 1:41AM  up 26 mins, 3 users, load averages: 0.08, 0.04, 0.03
USER     TTY      FROM                              LOGIN@  IDLE WHAT
user2    p0       10.209.208.123                   1:18AM     21 cli
user2    p1       192.0.2.207                    1:37AM      2 cli
user2    p2       192.0.2.19                     1:40AM      - cli

lcc0-re0:
--------------------------------------------------------------------------
 1:41AM  up 26 mins, 0 users, load averages: 0.00, 0.00, 0.03

lcc1-re0:
--------------------------------------------------------------------------
 1:41AM  up 26 mins, 0 users, load averages: 0.00, 0.02, 0.03

lcc2-re0:
--------------------------------------------------------------------------
 1:41AM  up 26 mins, 0 users, load averages: 0.16, 0.06, 0.02
```

```
lcc3-re0:
--------------------------------------------------------------------------
 1:41AM  up 26 mins, 0 users, load averages: 0.12, 0.04, 0.04

user3@aj> show system users
sfc0-re0:
--------------------------------------------------------------------------
 1:42AM  up 28 mins, 4 users, load averages: 0.02, 0.03, 0.02
USER     TTY      FROM                              LOGIN@  IDLE WHAT
user     p0       device1.example.com              1:18AM     22 cli
user     p1       device2.example.com              1:37AM      - cli
user     p2       device3.example.com              1:40AM      - cli
user     p3       device4.example.com              1:42AM      - -csh (csh)

lcc0-re0:
--------------------------------------------------------------------------
 1:42AM  up 28 mins, 0 users, load averages: 0.02, 0.01, 0.03

lcc1-re0:
--------------------------------------------------------------------------
 1:42AM  up 28 mins, 0 users, load averages: 0.07, 0.04, 0.03

lcc2-re0:
--------------------------------------------------------------------------
 1:42AM  up 27 mins, 0 users, load averages: 0.07, 0.06, 0.02

lcc3-re0:
--------------------------------------------------------------------------
 1:42AM  up 28 mins, 0 users, load averages: 0.05, 0.04, 0.04
```

show system users (QFX Series)

```
user@switch> show system users
USER     TTY      FROM                              LOGIN@  IDLE WHAT
user1   p0       192.0.2.117                    2:54AM     39 -cli (cli)
user2   p1       192.0.2.116                    3:01AM      - -cli (cli)
user3   p2       192.0.2.97                    3:08AM     11 -cli (cli)
```

show system users no-resolve (QFX Series)

```
user@switch> show system users no-resolve
USER     TTY      FROM                              LOGIN@  IDLE WHAT
user1    p0       192.0.2.117                    2:54AM     39 -cli (cli)
user2    p1       192.0.2.116                    3:01AM      - -cli (cli)
user3    p2       192.0.2.97                    3:08AM     11 -cli (cli)
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced for the TX Matrix Plus router in JUNOS OS Release 9.6.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system virtual-memory

IN THIS SECTION

- Syntax | 27267
- Syntax (EX Series) | 27267
- Syntax (TX Matrix Router) | 27267
- Syntax (TX Matrix Plus Router) | 27267
- Syntax (MX Series Router) | 27267
- Syntax (QFX Series) | 27268
- Syntax (Junos OS Evolved) | 27268
- Description | 27268
- Options | 27268
- Additional Information | 27269
- Required Privilege Level | 27270

- Output Fields | 27270
- Sample Output | 27273
- Release Information | 27278

Syntax

```
show system virtual-memory
```

Syntax (EX Series)

```
show system virtual-memory
<all-members>
<local>
<member member-id>
```

Syntax (TX Matrix Router)

```
show system virtual-memory
<all-chassis | all-lcc | lcc number | scc>
```

Syntax (TX Matrix Plus Router)

```
show system virtual-memory
<all-chassis | all-lcc | lcc number | sfc number>
```

Syntax (MX Series Router)

```
show system virtual-memory
<all-members>
<local>
<member member-id>
```

Syntax (QFX Series)

```
show system virtual-memory
<infrastructure name>
<interconnect-device name>
<node-group name>
```

Syntax (Junos OS Evolved)

```
show system virtual-memory
<node node-name>
```

Description

Display the usage of Junos OS kernel memory listed first by size of allocation and then by type of usage. Use the `show system virtual-memory` command for troubleshooting with Juniper Networks Customer Support.

Options

none	Display kernel dynamic memory usage information.
all-chassis	(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display kernel dynamic memory usage information for all chassis.
all-lcc	(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display kernel dynamic memory usage information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display kernel dynamic memory usage information for all connected T1600 or T4000 LCCs.
all-members	(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for all members of the Virtual Chassis configuration.
infrastructure *name*	(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the fabric control Routing Engine and fabric manager Routing Engine.
interconnect-device *name*	(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the Interconnect device.

lcc *number*

(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display kernel dynamic memory usage information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display kernel dynamic memory usage information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local

(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for the local Virtual Chassis member.

member *member-id*

(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

node *node-name*

(Junos OS Evolved only) (Optional) Display the kernel memory information for the specified node. If no node is specified, information is displayed for all nodes.

node-group *name*

(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the Node group.

scc

(TX Matrix routers only) (Optional) Display kernel dynamic memory usage information for the TX Matrix router (or switch-card chassis).

sfc *number*

(TX Matrix Plus routers only) (Optional) Display kernel dynamic memory usage information for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system virtual-memory` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of

the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

> **NOTE**: The `show system virtual-memory` command with the `| display XML` pipe option now displays XML output for the command in the parent tags: `<vmstat-memstat-malloc>`, `<vmstat-memstat-zone>`, `<vmstat-sumstat>`, `<vmstat-intr>`, and `<vmstat-kernel-state>` with each child element as a separate XML tag. In Junos OS Releases 10.1 and earlier, the `| display XML` option for this command does not have an XML API element and the entire output is displayed in a single `<output>` tag element.
>
> `kernel direct memory map` only displays for the 64-bit platform.

Required Privilege Level

view

Output Fields

Table 1694 on page 27270 lists the output fields for the `show system virtual-memory` command. Output fields are listed in the approximate order in which they appear.

Table 1694: show system virtual-memory Output Fields

Field Name	Field Description
Memory statistics by bucket size	
`Size`	Memory block size (bytes). The kernel memory allocator appropriates blocks of memory whose size is exactly a power of 2.
`In Use`	Number of memory blocks of this size that are in use (bytes).
`Free`	Number of memory blocks of this size that are free (bytes).
`Requests`	Number of memory allocation requests made.
`HighWater`	Maximum value the free list can have. Once the system starts reclaiming physical memory, it continues until the free list is increased to this value.

Table 1694: show system virtual-memory Output Fields *(Continued)*

Field Name	Field Description
`Couldfree`	Total number of times that the free elements for a bucket size exceed the high-water mark for that bucket size.
Memory usage type by bucket size	
`Size`	Memory block size (bytes).
`Type(s)`	Kernel modules that are using these memory blocks. For a definition of each type, refer to a FreeBSD book.
Memory statistics by type	
`Type`	Kernel module that is using dynamic memory.
`InUse`	Number of memory blocks used by this type. The number is rounded up.
`MemUse`	Amount of memory in use, in kilobytes (KB).
`HighUse`	Maximum memory ever used by this type.
`Limit`	Maximum memory that can be allocated to this type.
`Requests`	Total number of dynamic memory allocation requests this type has made.
`Type Limit`	Number of times requests were blocked for reaching the maximum limit.
`Kern Limit`	Number of times requests were blocked for the kernel map.
`Size(s)`	Memory block sizes this type is using.
Memory Totals	

Table 1694: show system virtual-memory Output Fields *(Continued)*

Field Name	Field Description
`In Use`	Total kernel dynamic memory in use (bytes, rounded up).
`Free`	Total kernel dynamic memory free (bytes, rounded up).
`Requests`	Total number of memory allocation requests.
`ITEM`	Kernel module that is using memory.
`Size`	Memory block size (bytes).
`Limit`	Maximum memory that can be allocated to this type.
`Used`	Number of memory blocks used by this type. The number is rounded up.
`Free`	Number of memory blocks available to this type.
`Requests`	Total number of memory allocation requests this type has made.
`interrupt`	Timer events and scheduling interruptions.
`total`	Total number of interruptions for each type.
`rate`	Interruption rate.
`Total`	Total for all interruptions.

Sample Output

show system virtual-memory

```
user@host> show system virtual-memory
Memory statistics by bucket size
Size   In Use   Free   Requests  HighWater  Couldfree
  16      906    118     154876    1280          0
  32      455    313     209956     640          0
  64     4412    260      75380     320         20
 128     3200     32      19361     160         81
 256     1510     10       8844      80          4
 512      446      2       5085      40          0
  1K       18      2       5901      20          0
  2K     1128      2       4445      10       1368
  4K      185      1        456       5          0
  8K        5      1       2653       5          0
 16K      181      0        233       5          0
 32K        2      0       1848       5          0
 64K       20      0         22       5          0
128K        5      0          5       5          0
256K        2      0          2       5          0
512K        1      0          1       5          0

Memory usage type by bucket size
Size  Type(s)
  16  uc_devlist, nexusdev, iftable, temp, devbuf, atexit, COS, BPF,
          DEVFS mount, DEVFS node, vnodes, mount, pcb, soname, proc-args, kld,
          MD disk, rman, ATA generic, bus, sysctl, ippool, pfestat, ifstate,
          pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode
  32  atkbddev, dirrem, mkdir, diradd, freefile, freefrag, indirdep,
          bmsafemap, newblk, temp, devbuf, COS, vnodes, cluster_save buffer,
          pcb, soname, proc-args, sigio, kld, Gzip trees, taskqueue, SWAP,
          eventhandler, bus, sysctl, uidinfo, subproc, pgrp, pfestat, itable32,
          ifstate, pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode, rtnexthop
  64  isadev, iftable, MFS node, allocindir, allocdirect, pagedep, temp,
          devbuf, lockf, COS, NULLFS hash, DEVFS name, vnodes,
          cluster_save buffer, vfscache, pcb, soname, proc-args, file,
          AR driver, AD driver, Gzip trees, rman, eventhandler, bus, sysctl,
          subproc, pfestat, pic, ifstate, pfe_ipc, mkey, ifaddr, rtable, ipfw
 128  ZONE, freeblks, inodedep, temp, devbuf, zombie, COS, DEVFS node,
          vnodes, mount, vfscache, pcb, soname, proc-args, ttys, dev_t,
```

```
          timecounter, kld, Gzip trees, ISOFS node, bus, uidinfo, cred,
          session, pic, itable16, ifstate, pfe_ipc, rtable, ifstat, metrics,
          rtnexthop, iffamily
 256  iflogical, iftable, MFS node, FFS node, newblk, temp, devbuf,
          NFS daemon, vnodes, proc-args, kqueue, file desc, Gzip trees, bus,
          subproc, itable16, ifstate, pfe_ipc, sysctl, rtnexthop
 512  UFS mount, temp, devbuf, mount, BIO buffer, ptys, ttys, AR driver,
          Gzip trees, ISOFS mount, msg, ioctlops, ATA generic, bus, proc,
          pfestat, lr, ifstate, pfe_ipc, rtable, ipfw, ifstat, rtnexthop
  1K  iftable, temp, devbuf, NQNFS Lease, kqueue, kld, AD driver,
          Gzip trees, sem, MD disk, bus, ifstate, pfe_ipc, ipfw
  2K  uc_devlist, UFS mount, temp, devbuf, BIO buffer, pcb, AR driver,
          Gzip trees, ioctlops, bus, ipfw, ifstat, rcache
  4K  memdesc, iftable, UFS mount, temp, devbuf, kld, Gzip trees, sem, msg
  8K  temp, devbuf, syncache, Gzip trees
 16K  indirdep, temp, devbuf, shm, msg
 32K  pagedep, kld, Gzip trees
 64K  VM pgdata, devbuf, MSDOSFS mount
128K  UFS ihash, inodedep, NFS hash, kld, ISOFS mount
256K  mbuf, vfscache
512K  SWAP

Memory statistics by type                          Type  Kern
        Type  InUse MemUse HighUse  Limit Requests Limit Limit Size(s)
       isadev    13     1K      1K127753K       13    0     0  64
     atkbddev     2     1K      1K127753K        2    0     0  32
   uc_devlist    24     3K      3K127753K       24    0     0  16,2K
     nexusdev     3     1K      1K127753K        3    0     0  16
      memdesc     1     4K      4K127753K        1    0     0  4K
         mbuf     1   152K    152K127753K        1    0     0  256K
    iflogical     6     2K      2K127753K        6    0     0  256
      iftable    17     9K      9K127753K       18    0     0  16,64,256,1K,4K
         ZONE    15     2K      2K127753K       15    0     0  128
    VM pgdata     1    64K     64K127753K        1    0     0  64K
    UFS mount    12    26K     26K127753K       12    0     0  512,2K,4K
    UFS ihash     1   128K    128K127753K        1    0     0  128K
     MFS node     6     2K      3K127753K       35    0     0  64,256
     FFS node   906   227K    227K127753K     1352    0     0  256
       dirrem     0     0K      4K127753K      500    0     0  32
        mkdir     0     0K      1K127753K       38    0     0  32
       diradd     0     0K      6K127753K      521    0     0  32
     freefile     0     0K      4K127753K      374    0     0  32
     freeblks     0     0K      8K127753K      219    0     0  128
```

```
      freefrag     0     0K      1K127753K      193    0     0  32
    allocindir     0     0K     25K127753K     1518    0     0  64
      indirdep     0     0K     17K127753K       76    0     0  32,16K
   allocdirect     0     0K     10K127753K      760    0     0  64
     bmsafemap     0     0K      1K127753K       72    0     0  32
        newblk     1     1K      1K127753K     2279    0     0  32,256
      inodedep     1   128K    175K127753K     2367    0     0  128,128K
       pagedep     1    32K     33K127753K       47    0     0  64,32K
          temp  1239    92K     96K127753K     8364    0     0  16,32,64K
        devbuf  1413  5527K   5527K127753K     1535    0     0  16,32,64,128,256
         lockf    38     3K      3K127753K     2906    0     0  64
        atexit     1     1K      1K127753K        1    0     0  16
        zombie     0     0K      2K127753K     3850    0     0  128
      NFS hash     1   128K    128K127753K        1    0     0  128K
   NQNFS Lease     1     1K      1K127753K        1    0     0  1K
    NFS daemon     1     1K      1K127753K        1    0     0  256
      syncache     1     8K      8K127753K        1    0     0  8K
           COS   353    44K     44K127753K      353    0     0  16,32,64,128
           BPF   189     3K      3K127753K      189    0     0  16
 MSDOSFS mount     1    64K     64K127753K        1    0     0  64K
   NULLFS hash     1     1K      1K127753K        1    0     0  64
   DEVFS mount     2     1K      1K127753K        2    0     0  16
    DEVFS name   487    31K     31K127753K      487    0     0  64
    DEVFS node   471    58K     58K127753K      479    0     0  16,128
        vnodes    28     7K      7K127753K      429    0     0  16,32,64,128,256
         mount    15     8K      8K127753K       18    0     0  16,128,512
cluster_save buffer     0     0K      1K127753K       55    0     0  32,64
      vfscache  1898   376K    376K127753K     3228    0     0  64,128,256K
    BIO buffer    49    98K    398K127753K      495    0     0  512,2K
           pcb   159    16K     17K127753K      399    0     0  16,32,64,128,2K
        soname    82    10K     10K127753K    42847    0     0  16,32,64,128
     proc-args    57     2K      3K127753K     2105    0     0  16,32,64,128,256
          ptys    32    16K     16K127753K       32    0     0  512
          ttys   254    33K     33K127753K      522    0     0  128,512
        kqueue     5     3K      4K127753K       23    0     0  256,1K
         sigio     1     1K      1K127753K       27    0     0  32
          file   383    24K     24K127753K    16060    0     0  64
     file desc    76    19K     20K127753K     3968    0     0  256
           shm     1    12K     12K127753K        1    0     0  16K
         dev_t   286    36K     36K127753K      286    0     0  128
   timecounter    10     2K      2K127753K       10    0     0  128
           kld    11   117K    122K127753K       34    0     0  16,32,128,1K,4K
     AR driver     1     1K      3K127753K        5    0     0  64,512,2K
```

```
     AD driver     2     2K      3K127753K     2755    0     0  64,1K
    Gzip trees     0     0K     46K127753K   133848    0     0  32,64,128,256
    ISOFS node  1136   142K    142K127753K     1189    0     0  128
   ISOFS mount     9   132K    132K127753K       10    0     0  512,128K
           sem     3     6K      6K127753K        3    0     0  1K,4K
       MD disk     2     2K      2K127753K        2    0     0  16,1K
           msg     4    25K     25K127753K        4    0     0  512,4K,16K
          rman    59     4K      4K127753K      461    0     0  16,64
      ioctlops     0     0K      2K127753K      992    0     0  512,2K
     taskqueue     2     1K      1K127753K        2    0     0  32
          SWAP     2   413K    413K127753K        2    0     0  32,512K
   ATA generic     6     3K      3K127753K        6    0     0  16,512
  eventhandler    17     1K      1K127753K       17    0     0  32,64
           bus   340    30K     31K127753K      794    0     0  16,32,64,128,256
        sysctl     0     0K      1K127753K   130262    0     0  16,32,64
       uidinfo     4     1K      1K127753K       10    0     0  32,128
          cred    22     3K      3K127753K     3450    0     0  128
       subproc   156    10K     10K127753K     7882    0     0  32,64,256
          proc     2     1K      1K127753K        2    0     0  512
       session    12     2K      2K127753K       34    0     0  128
          pgrp    16     1K      1K127753K       45    0     0  32
        ippool     1     1K      1K127753K        1    0     0  16
       pfestat     0     0K      1K127753K    47349    0     0  16,32,64,512
           pic     5     1K      1K127753K        5    0     0  64,128
            lr     1     1K      1K127753K        1    0     0  512
      itable32   110     4K      4K127753K      110    0     0  32
      itable16   161    26K     26K127753K      161    0     0  128,256
       ifstate   694   159K    160K127753K     1735    0     0  16,32,64,128,1K
       pfe_ipc     0     0K      1K127753K    56218    0     0  16,32,64,128,1K
          mkey   250     4K      4K127753K      824    0     0  16,32,64
        ifaddr     9     1K      1K127753K        9    0     0  64
        sysctl     0     0K      1K127753K       30    0     0  256
        rtable    49     6K      6K127753K      307    0     0  16,32,64,128,512
       ifmaddr    22     1K      1K127753K       22    0     0  16,32
          ipfw    23    10K     10K127753K       48    0     0  16,32,64,512,2K
        ifstat   698   805K    805K127753K      698    0     0  128,512,2K
        rcache     4     8K      8K127753K        4    0     0  2K
         rnode    27     1K      1K127753K      285    0     0  16,32
       metrics     1     1K      1K127753K        3    0     0  128
     rtnexthop    57     9K      9K127753K      312    0     0  32,128,256,512
      iffamily    12     2K      2K127753K       12    0     0  128

Memory Totals:  In Use       Free    Requests
```

```
                  9311K        54K      489068

ITEM            SIZE     LIMIT    USED    FREE  REQUESTS
PIPE:            192,        0,      4,     81,     4422
SWAPMETA:        160,    95814,      0,      0,        0
unpcb:           160,        0,    114,     36,      279
ripcb:           192,    25330,      5,     37,        5
syncache:        128,    15359,      0,     64,        5
tcpcb:           576,    25330,     23,     12,       32
udpcb:           192,    25330,     14,     28,      255
socket:          256,    25330,    246,     26,      819
KNOTE:            96,        0,     27,     57,       71
NFSNODE:         352,        0,      0,      0,        0
NFSMOUNT:        544,        0,      0,      0,        0
VNODE:           224,        0,   2778,     43,     2778
NAMEI:          1024,        0,      0,      8,    40725
VMSPACE:         192,        0,     57,     71,     3906
PROC:            448,        0,     73,     17,     3923
DP fakepg:        64,        0,      0,      0,        0
PV ENTRY:         28,   499566,  44530, 152053,  1525141
MAP ENTRY:        48,        0,   1439,    134,   351075
KMAP ENTRY:       48,    35645,    179,    119,    10904
MAP:             108,        0,      7,      3,        7
VM OBJECT:        92,        0,   2575,    109,    66912

   792644 cpu context switches
  9863474 device interrupts
   286510 software interrupts
   390851 traps
  3596829 system calls
       16 kernel threads created
     3880  fork() calls
       27 vfork() calls
        0 rfork() calls
        0 swap pager pageins
        0 swap pager pages paged in
        0 swap pager pageouts
        0 swap pager pages paged out
      380 vnode pager pageins
      395 vnode pager pages paged in
      122 vnode pager pageouts
     1476 vnode pager pages paged out
        0 page daemon wakeups
```

```
         0 pages examined by the page daemon
       101 pages reactivated
    161722 copy-on-write faults
         0 copy-on-write optimized faults
     84623 zero fill pages zeroed
     83063 zero fill pages prezeroed
         7 intransit blocking page faults
    535606 total VM faults taken
         0 pages affected by kernel thread creation
    238254 pages affected by  fork()
      2535 pages affected by vfork()
         0 pages affected by rfork()
    283379 pages freed
         0 pages freed by daemon
    190091 pages freed by exiting processes
     17458 pages active
     29166 pages inactive
         0 pages in VM cache
     10395 pages wired down
    134610 pages free
      4096 bytes per page
    183419 total name lookups
           cache hits (90% pos + 7% neg) system 0% per-directory
           deletions 0%, falsehits 0%, toolong 0%

interrupt                   total       rate
ata0 irq14                 113338          3
mux irq7                   727643         21
fxp1 irq10                1178671         34
sio0 irq4                     833          0
clk irq0                  3439769         99
rtc irq8                  4403221        127
Total                     9863475        286

Kernel direct memory map:
        4423  pages used
     4057340  pages maximum
```

Release Information

Command introduced before Junos OS Release 7.4.

`sfc` option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

`node` option introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

Routing Matrix with a TX Matrix Plus Router Solutions Page

show system visibility cpu

IN THIS SECTION

- Syntax | 27279
- Description | 27279
- Required Privilege Level | 27279
- Output Fields | 27280
- Sample Output | 27281
- Release Information | 27285

Syntax

```
show system visibility cpu
```

Description

Display details such as per CPU statistics, per CPU usage, and CPU pinning for a Junos OS platform.

Required Privilege Level

view

Output Fields

Table 1695 on page 27280 lists the output fields for the `show system visibility cpu` command. Output fields are listed in the approximate order in which they appear.

> **NOTE**: Starting in Junos OS Release 22.1R1, the CPU allocations are displayed as a number range (for example, 4-7) instead of individual numbers (for example, 4,5,6,7).

Table 1695: show system visibility cpu Output Fields

Field Name	Field Description
Fields for CPU Statistics	
`CPU ID`	The CPU ID
`User Time`	The amount of user time, in seconds.
`System Time`	The amount of system time, in seconds.
`Idle Time`	The amount of time spent in idle mode, in seconds.
`Nice Time`	The amount of spent nice time, in seconds.
`I/O Wait Time`	The amount of time spent waiting for input/output (I/O) operations, in seconds.
`Interrupt Service Time`	The amount of interrupt service time, in seconds.
`Service Time`	The amount of service time, in seconds.
Fields for CPU Usages	
`CPU ID`	The CPU ID

Table 1695: show system visibility cpu Output Fields *(Continued)*

Field Name	Field Description
`CPU Usage`	The percentage of CPU used.
Fields for CPU Pinning Information	
`Virtual Machine`	The name of the virtual machine.
`vCPU`	The ID of virtual CPUs used by the virtual machine.
`CPU`	The ID of CPUs used by the virtual machine.
Fields for Emulator Thread Pinning Information (This section is displayed starting in Junos OS Release 22.1R1.)	
Virtual Machine	The name of the virtual machine.
CPU	The ID of CPUs used by the virtual machine.
`System Component`	The name of the system component.
`CPUs`	The ID of CPUs used by the system component.
Quota	The CPU quota configured using the `cpu colocation quota` parameter. **NOTE**: This parameter is displayed starting in Junos OS Release 22.1R1.

Sample Output

show system visibility cpu (NFX150)

```
user@host> show system visibility cpu
CPU Statistics (Time in sec)
--------------------------------------------------------------------------------
CPU Id User Time System Time Idle Time Nice Time IOWait Time Intr. Service Time
------ --------- ----------- --------- --------- ----------- ------------------
```

```
0      542       392         879       0         5           0
1      1747      1           97        0         0           0
2      29        40          1742      0         0           0
3      1361      0           483       0         0           0
4      213       27          1594      0         5           0
5      0         0           1844      0         0           0
6      17        3           1821      0         1           0
7      17        5           1821      0         0           0

CPU Usages
----------------
CPU Id CPU Usage
------ ---------
0      34.600000000000001
1      100.0
2      7.0
3      100.0
4      0.0
5      0.0
6      0.0
7      0.0

CPU Pinning Information
---------------------------------------
Virtual Machine             vCPU    CPU
--------------------------- ------- ---
vjunos0                     0       0

Emulator Thread Pinning Information
-------------------------------
Virtual Machine             CPU
--------------------------- ---
vjunos0                     0


System Component                CPUs     Quota
------------------------------- -------- -----
ovs-vswitchd                    0-1
riot                            0,2
srxpfe                          0,3
jdm                             0
```

show system visibility cpu (NFX250 (NG))

```
user@host> show system visibility cpu
CPU Statistics (Time in sec)
--------------------------------------------------------------------------------
CPU Id User Time System Time Idle Time Nice Time IOWait Time Intr. Service Time
------ --------- ----------- --------- --------- ----------- ------------------
0      40084     1324        270       0         2           0
1      203       395         41009     0         10          0
2      209       21          41435     0         22          0
3      0         0           41680     0         21          0
4      41659     0           64        0         0           0
5      6150      376         33392     0         0           0
6      0         0           41712     0         8           0
7      0         0           41714     0         6           0

CPU Usages
----------------
CPU Id CPU Usage
------ ---------
0      100.0
1      1.0
2      0.0
3      0.0
4      100.0
5      100.0
6      0.0
7      0.0

CPU Pinning Information
---------------------------------------
Virtual Machine             vCPU    CPU
--------------------------- ------- ---
vjunos0                     0       0
testvnf1                    0       2
testvnf1                    1       6-7

Emulator Thread Pinning Information
-------------------------------
Virtual Machine             CPU
--------------------------- ---
vjunos0                     0
```

```
testvnf1                    2-3



System Component                CPUs     Quota
------------------------------- -------- -----
ovs-vswitchd                    0,4
srxpfe                          1,5
jdm                             1
```

show system visibility cpu (NFX350)

```
user@host> show system visibility cpu
CPU Statistics (Time in sec)
--------------------------------------------------------------------------------
CPU Id User Time System Time Idle Time Nice Time IOWait Time Intr. Service Time
------ --------- ----------- --------- --------- ----------- ------------------
0      1463      76          148       0         3           0
1      1616      3           78        0         0           0
2      1300      1           395       0         0           0
3      112       13          1567      0         2           0
4      1         0           1693      0         1           0
5      0         0           1695      0         2           0
6      0         0           1697      0         0           0
7      0         0           1696      0         0           0
8      188       18          1475      0         4           0
11     0         0           1698      0         0           0
12     0         0           1698      0         0           0
13     0         0           1698      0         0           0
14     0         0           1698      0         0           0
15     0         0           1698      0         0           0

CPU Usages
----------------
CPU Id CPU Usage
------ ---------
0      100.0
1      100.0
2      100.0
3      0.0
4      0.0
5      1.0
```

```
6      0.0
7      0.0
8      23.0
11     0.0
12     0.0
13     0.0
14     0.0
15     0.0

CPU Pinning Information
---------------------------------------
Virtual Machine             vCPU    CPU
--------------------------- ------- ---
vjunos0                     0       8
testvnf1                    0       3-4
testvnf1                    1       5-7

Emulator Thread Pinning Information
-------------------------------
Virtual Machine             CPU
--------------------------- ---
vjunos0                     8
testvnf1                    11-15


System Component                CPUs     Quota
------------------------------- -------- -----
ovs-vswitchd                    1,8
srxpfe                          0,2
jdm                             8
```

Release Information

Command introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

show system visibility host | 27286

show system visibility memory | 27304

show system visibility network | 27308

show system visibility vnf | 27316

show system visibility host

IN THIS SECTION

- Syntax | 27286
- Description | 27286
- Required Privilege Level | 27286
- Output Fields | 27286
- Sample Output | 27290
- Release Information | 27295

Syntax

```
show system visibility host
```

Description

Displays details such as the host uptime, number of tasks, CPU statistics, list of disk partitions, disk usage, disk I/O statistics, list of network interfaces, and per port statistics for a Junos OS platform.

Required Privilege Level

view

Output Fields

Table 1696 on page 27287 lists the output fields for the `show system visibility host` command. Output fields are listed in the approximate order in which they appear.

Table 1696: show system visibility host Output Fields

Field Name	Field Description
Field for Host Uptime	
`Uptime`	The time the host has been operational.
Fields for Host Tasks	
`Total`	The total number of tasks.
`Running`	The total number of tasks running.
`Sleeping`	The total number of tasks in sleeping state.
`Stopped`	The total number of tasks that are stopped.
`Zombie`	The total number of zombie processes.
Fields for Host CPU Information	
`User Time`	The amount of user time, in seconds.
`System Time`	The amount of system time, in seconds.
`Idle Time`	The amount of time spent in idle mode, in seconds.
`Nice Time`	The amount of spent nice time, in seconds.
`I/O Wait Time`	The amount of time spent waiting for input/output (I/O) operations, in seconds.
`Interrupt Service Time`	The amount of interrupt service time, in seconds.

Table 1696: show system visibility host Output Fields *(Continued)*

Field Name	Field Description
Fields for Host Disk Partitions	
`Device`	The device path.
`Mount Point`	The mount point of the device path.
`File System`	The file system type.
`Options`	Options available for the device path.
Fields for Host Disk Usage Information	
`Total`	The total amount of disk usage space, in mebibytes (MiB).
`Used`	The amount of used disk usage space, in mebibytes (MiB).
`Free`	The amount of free disk usage space, in mebibytes (MiB).
`Percentage Used`	The percentage of used disk space.
Fields for Host Disk I/O Information	
`Read Count`	The number of times the disk has been read.
`Write Count`	The number of times a write operation has happened on the disk.
`Read Bytes`	The number of bytes used in read operations on the disk.
`Write Bytes`	The number of bytes used in write operations on the disk.
`Read Time`	The amount of time the disk has been read, in milliseconds.

Table 1696: show system visibility host Output Fields *(Continued)*

Field Name	Field Description
`Write Time`	The amount of time write operations have been performed on the disk, in milliseconds.
Fields for List of Host Interfaces	
`Interfaces`	The name of the interface.
`State`	The state of the Host Interface.
`MAC`	The MAC address of the interface.
Fields for List of Host Port Statistics	
`Interface`	The name of the interface.
`Bytes Sent`	The number of bytes sent.
`Bytes Received`	The number of bytes received.
`Packets Sent`	The number of packets sent.
`Packets Received`	The number of packets received.
`Errors In`	The number of errors in.
`Errors Out`	The number of errors out.
`Drops In`	The number of drops in.
`Drops Out`	The number of drops out.

Sample Output

show system visibility host (NFX150)

```
user@host> show system visibility host
Host Uptime
-----------
Uptime: 1 day 23:19:41.21000

Host Tasks
----------
Total:    187
Running:  3
Sleeping: 179
Stopped:  0
Zombie:   5

Host CPU Information (Time in sec)
----------------------------------
User Time:              79359
System Time:            0
Idle Time:              502215
I/O Wait Time:          103
Nice Time:              103724
Interrupt Service Time: 0

Host Disk Partitions
-----------------------------------------------------------------------------------------------
---------
Device                                    Mount Point      File System  Options
----------------------------------------- ---------------- -----------
----------------------------------
/dev/sda2                                 /                ext4
rw,relatime,i_version,data=ordered
/dev/sda1                                 /boot/efi        vfat
rw,noatime,fmask=0022,dmask=0022,codepage=437,iocharset=iso8859-1,shortname=mixed,errors=remount-
ro
/dev/sda7                                 /config          ext4         rw,noatime,data=ordered
/dev/sda8                                 /var/log         ext4         rw,noatime,data=ordered
/dev/sda9                                 /mnt/.share      ext4
rw,noatime,discard,data=ordered
/dev/sda5                                 /junos           ext4
```

```
rw,noatime,discard,data=ordered
/dev/loop0                                /var/tmp         ext4         rw,relatime,data=ordered
/dev/loop1                                /mnt/.share/lshare/jnpr/jlog ext4
rw,relatime,data=ordered
/dev/loop0                                /mnt/.share/lshare/jnpr/jtmp ext4
rw,relatime,data=ordered

Host Disk Usage Information
---------------------------
Total (MiB):     1469
Used  (MiB):     948
Free  (MiB):     429
Percentage Used: 64.5

Host Disk I/O Information
-------------------------
 Read Count: 187083
Write Count: 256206
 Read Bytes: 2290787328
Write Bytes: 3331667456
  Read Time: 33977
 Write Time: 258864

Host Interfaces
----------------------------------------------
Interface          State     MAC
------------------ --------- -----------------
heth-0-1           active    00:00:5e:00:53:8e
heth-0-0           active    00:00:5e:00:53:8d
heth-0-3           active    00:00:5e:00:53:90
heth-0-2           active    00:00:5e:00:53:8f
heth-0-5           inactive  00:00:5e:00:53:92
heth-0-4           inactive  00:00:5e:00:53:91
ctrlbr0            active    00:00:5e:00:53:10
docker0            inactive  00:00:5e:00:53:8c
eth0br             active    00:00:5e:00:53:00
eth1br             inactive  00:00:5e:00:53:67
l3_h_ge_1_0_0      active    00:00:5e:00:53:6d
l3_h_ltectrl       active    00:00:5e:00:53:f1
l3_h_ltedata       active    00:00:5e:00:53:91
lo                 inactive  00:00:00:00:00:00
lte_crtl0          active    00:00:5e:00:53:91
lte_data0          active    00:00:5e:00:53:fc
```

```
ovs-sys-br         inactive  00:00:5e:00:53:4f
ovs-system         inactive  00:00:5e:00:53:1b
sit0               inactive  00:00:00:00
veth00             active    00:00:5e:00:53:79
veth01             active    00:00:5e:00:53:87
veth10             active    00:00:5e:00:53:40
veth11             active    00:00:5e:00:53:65
virbr0             active    00:00:5e:00:53:83
virbr1             active    00:00:5e:00:53:6f

Host Port Statistics
-----------------------------------------------------------------------------------------------
----
Interface Bytes Sent   Bytes Rcvd   Packets Sent Packets Rcvd Errors In Errors Out Drops In
Drops Out
--------- ------------ ------------ ------------ ------------ --------- ---------- --------
---------
l3_h_ge_1_0_0 11025    648          74           8            0         0          0        0
veth10    0            11673        0            82           0         0          12       0
veth11    11673        0            82           0            0         0          0        0
ovs-system 0           0            0            0            0         0          0        0
ovs-sys-br 0           0            0            0            0         0          82       0
vnet0     31080352     10698402     153074       136451       0         0          0        0
vnet1     858553596    712231555    9325949      10546588     0         0          0        0
vnet2     735033102    50689829     4956943      180168       0         0          0        0
vnet3     4428680      602          85168        13           0         0          0        0
eth0      50689829     1077880063   180168       5551593      0         0          6146     0
eth1br    0            0            0            0            0         0          0        0
lte_data0 0            1648         0            14           0         0          0        0
lo        96584        96584        1219         1219         0         0          0        0
lte_crtl0 749623       12570778     22710        22762        0         0          0        0
virbr0-nic 0           0            0            0            0         0          0        0
docker0   0            0            0            0            0         0          0        0
veth01    4558         4743808      53           89402        0         0          0        0
veth00    4743808      4558         89402        53           0         0          8        0
dcapi-tap 0            0            0            0            0         0          0        0
l3_h_ltedata 1648      648          14           8            0         0          0        0
sit0      0            0            0            0            0         0          0        0
flowd_h_mgmt 391536979 448871585    5975703      5507199      0         0          0        0
virbr1    29553905     8096581      137792       128808       0         0          0        0
virbr0    46365        48232        467          540          0         0          0        0
l3_h_ltectrl 12570778  818395       22762        22718        0         0          0        0
jdm-hbme1 4474379      55866        85622        537          0         0          0        0
```

```
jdm-hbme2 813479       1526643      7992         15288        0         0          0        0
eth0br    0            595875398    0            4835907      0         0          222      0
ctrlbr0   408483097    256713674    3800585      4571275      0         0          0        0
heth-0-1  0            5368334      0            89330        0         0          0        0
heth-0-0  0            5366462      0            89349        0         0          0        0
heth-0-3  0            5367002      0            89358        0         0          0        0
heth-0-2  0            5365262      0            89329        0         0          0        0
heth-0-5  0            0            0            0            0         0          0        0
heth-0-4  0            0            0            0            0         0          0        0
```

show system visibility host (NFX250 (NG))

```
user@host> show system visibility host
Host Uptime
-----------
Uptime: 3 days 3:47:05.09000

Host Tasks
----------
Total:    198
Running:  1
Sleeping: 194
Stopped:  0
Zombie:   3

Host CPU Information (Time in sec)
----------------------------------
User Time:              574351
System Time:            0
Idle Time:              2692218
I/O Wait Time:          216
Nice Time:              4609
Interrupt Service Time: 0

Host Disk Partitions
-----------------------------------------------------------------------------------------------
---------
Device                                    Mount Point      File System  Options
----------------------------------------- ---------------- -----------
----------------------------------
/dev/sda2                                 /                ext4
```

```
rw,relatime,i_version,data=ordered
/dev/sda1                                /boot/efi        vfat
rw,noatime,fmask=0022,dmask=0022,codepage=437,iocharset=iso8859-1,shortname=mixed,errors=remount-
ro
/dev/sda7                                /config          ext4         rw,noatime,data=ordered
/dev/sda8                                /var/log         ext4         rw,noatime,data=ordered
/dev/sda9                                /mnt/.share      ext4
rw,noatime,discard,data=ordered
/dev/sda5                                /junos           ext4
rw,noatime,discard,data=ordered
/dev/loop0                               /var/tmp         ext4         rw,relatime,data=ordered

Host Disk Usage Information
---------------------------
Total (MiB):     1469
Used  (MiB):     906
Free  (MiB):     470
Percentage Used: 61.7

Host Disk I/O Information
-------------------------
 Read Count: 245805
Write Count: 333782
 Read Bytes: 2967304704
Write Bytes: 6147921408
  Read Time: 34906
 Write Time: 448918

Host Interfaces
----------------------------------------------
Interface          State     MAC
------------------ --------- -----------------
hsxe0              active    30:7c:5e:4c:78:44
hsxe1              active    30:7c:5e:4c:78:45
ctrlbr0            active    02:00:00:00:00:10
docker0            inactive  02:42:f9:e7:08:5f
eth0br             active    4c:96:14:00:00:00
eth1br             inactive  66:7e:98:6c:9d:a7
l3_h_ge_1_0_0      active    ca:6b:5a:fe:39:2c
lo                 inactive  00:00:00:00:00:00
sit0               inactive  00:00:00:00
virbr0             active    30:7c:5e:4c:78:43
virbr1             active    be:51:f7:ac:03:1b
```

```
Host Port Statistics
-----------------------------------------------------------------------------------------------
----
Interface Bytes Sent   Bytes Rcvd   Packets Sent Packets Rcvd Errors In Errors Out Drops In
Drops Out
--------- ------------ ------------ ------------ ------------ --------- ---------- --------
---------
l3_h_ge_1_0_0 0        648          0            8            0         0          0        0
ovs-sys-br 0           0            0            0            0         0          0        0
vnet0     2573491477   117345734    2448205      1790887      0         0          0        0
vnet1     670930985    585788796    7585078      8400542      0         0          0        0
vnet2     454043208    224389433    2873376      416585       0         0          0        0
vnet3     7129616      9814         137213       231          0         0          0        0
eth0      224389433    464747548    416585       2889060      0         0          9829     0
lo        61305        61305        920          920          0         0          0        0
virbr1    2475291351   90762062     1008399      1774468      0         0          0        0
irb       0            0            0            0            0         0          0        0
hsxe1     0            0            0            0            0         0          0        0
hsxe0     0            0            0            0            0         0          0        0
docker0   0            0            0            0            0         0          0        0
dcapi-tap 0            0            0            0            0         0          0        0
sit0      0            0            0            0            0         0          0        0
flowd_h_mgmt 387545386 426690199    5662328      5294853      0         0          0        0
virbr0-nic 0           0            0            0            0         0          0        0
virbr0    3021873      1067179      4573         6153         0         0          0        0
jdm-hbme1 1785562      33378        34145        404          0         0          0        0
jdm-hbme2 41904        72344        321          323          0         0          0        0
eth0br    0            401858893    0            2755416      0         0          226      0
ctrlbr0   243770080    159923150    2283092      2738720      0         0          0        0
eth1br    0            0            0            0            0         0          0        0
ovs-netdev 0           0            0            0            0         0          0        0
```

Release Information

Command introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

show system visibility cpu | 27279

show system visibility memory | 27304
show system visibility network | 27308
show system visibility vnf | 27316

show system visibility jcp

IN THIS SECTION

- Syntax | 27296
- Description | 27296
- Required Privilege Level | 27296
- Output Fields | 27296
- Sample Output | 27299
- Release Information | 27300

Syntax

```
show system visibility jcp
```

Description

Display details such as CPU statistics, memory usage, internal IP address, list of network interfaces, interface statistics, and the list of disks for Junos VM.

Required Privilege Level

view

Output Fields

Table 1697 on page 27297 lists the output fields for the `show system visibility jcp` command. Output fields are listed in the approximate order in which they appear.

Table 1697: show system visibility jcp Output Fields

Field Name	Field Description
Fields for JCP CPU Statistics	
`CPU Time`	The total CPU time, in seconds.
`System Time`	The total system time, in seconds.
`User Time`	The total user time, in seconds.
Fields for JCP Memory Usage	
`Maximum Memory`	The maximum amount of memory, in kibibytes (KiBs).
`Used Memory`	The total amount of used memory, in kibibytes (KiBs).
Fields for JCP Internal IP Addresses	
`Interface`	The name of the interface.
`Address`	The IP address of the interface.
Fields for JCP Interfaces	
`Interface`	The name of the interface.
`Type`	The type of the interface.
`Source`	The connectivity source.
`Model`	The connectivity model.
`MAC`	The MAC address of the interface.

Table 1697: show system visibility jcp Output Fields *(Continued)*

Field Name	Field Description
Fields for JCP Interfaces Statistics	
`Port`	The name of the port.
`Rcvd Bytes`	The number of bytes received.
`Rcvd Packets`	The number of packets received.
`Rcvd Error`	The number of errors received.
`Rcvd Drop`	The number of drops received.
`Trxd Bytes`	The number of bytes transmitted.
`Trxd Packets`	The number of packets transmitted.
`Trxd Error`	The number of errors transmitted.
`Trxd Drop`	The number of drops transmitted.
Fields for JCP Disk Information	
`Disk`	The type of disk.
`File`	The path to the disk.

Sample Output

show system visibility jcp

```
user@jdm> show system visibility jcp
JCP CPU Statistics (Time in sec)
-----------------------------
CPU Time:    21435526
System Time: 4981660
User Time:   780770

JCP Memory Usage
----------------
Maximum Memory (KiB): 200089
Used Memory (KiB):    200089

JCP Internal IP Addresses
-------------------------
Interface: em2.32768
Address  : 192.168.1.2

JCP Interfaces
--------------------------------------------------------------
Interface Type      Source       Model      MAC
--------- --------- ------------ ---------- -----------------
vnet1     bridge    virbr0       e1000      52:54:00:7b:28:91
iri1      bridge    ctrlbr0      e1000      52:54:00:db:ba:c2
vnet3     bridge    eth0br       e1000      52:54:00:c6:8e:e9
vjunos0_em1 bridge  ovs-sys-br   e1000      52:54:00:e4:4d:2e

JCP Interfaces Statistics
-----------------------------------------------------------------------------------------------
------
Port      Rcvd Bytes   Rcvd Packets Rcvd Error Rcvd Drop Trxd Bytes   Trxd Packets Trxd Error
Trxd Drop
--------- ------------ ------------ ---------- --------- ------------ ------------ ----------
---------
vnet1     107535558    1150988      0          0         57984280     923183       0          0
iri1      384836814    5242053      0          0         349833594    5250766      0          0
vnet3     1341546383   6654730      0          125       1169070      27835        0          0
vjunos0_em1 4656       58           0          0         0            0            0          0
```

```
JCP Disk Information
---------------------------------------------------------------------------
Disk        File
----------- -------------------------------------------------------------
hda         /junos/images/0/vjunos.img
hdb         /junos/images/0/vjunos-data.img
hdc         /junos/images/0/vjunos-config.img
hdd         /junos/images/0/vjunos-platform.img
```

Release Information

Command introduced in Junos OS Release 15.1X53-D40.

RELATED DOCUMENTATION

- show system visibility cpu
- show system visibility host
- show system visibility jdm | 27300
- show system visibility memory
- show system visibility network
- show system visibility storage
- show system visibility vnf

show system visibility jdm

IN THIS SECTION

- Syntax | 27301
- Description | 27301
- Required Privilege Level | 27301
- Output Fields | 27301
- Sample Output | 27303
- Release Information | 27304

Syntax

```
show system visibility jdm
```

Description

Display details such as uptime, number of tasks, CPU usage, disk usage, memory usage, the list of network interfaces, and internal IP address for JDM container.

Required Privilege Level

view

Output Fields

Table 1698 on page 27301 lists the output fields for the `show system visibility jdm` command. Output fields are listed in the approximate order in which they appear.

Table 1698: show system visibility jdm Output Fields

Field Name	Field Description
Fields for JDM CPU Usage	
`Percentage Used`	The percentage of CPU utilisation.
Fields for JDM Disk Usage Information	
`Total`	The total amount of disk usage space, in mebibytes (MiB).
`Used`	The amount of used disk usage space, in mebibytes (MiB).
`Free`	The amount of free disk usage space, in mebibytes (MiB).
`Percentage Used`	The percentage of used disk space.
Fields for JDM Memory Information	

Table 1698: show system visibility jdm Output Fields *(Continued)*

Field Name	Field Description
`Maximum Memory`	The maximum virtual memory used, in kilobytes (KiB).
`Used Memory`	The amount of used virtual memory, in kilobytes (KiB).
`Percentage Used`	The percentage of memory utilisation.
Fields for List of JDM Interfaces	
`Interface`	The name of the interface.
`Type`	The type of interface.
`Source`	The connectivity source.
`MAC`	The MAC address of the interface.
Field for JDM Uptime	
`Uptime`	The time the JDM has been operational.
Fields for JDM Tasks	
`Total`	The total number of tasks.
`Running`	The total number of tasks running.
`Sleeping`	The total number of tasks in sleeping state.
`Stopped`	The total number of tasks that are stopped.
`Zombie`	The number of zombie processes.

Table 1698: show system visibility jdm Output Fields *(Continued)*

Field Name	Field Description
Fields for JDM Internal IP Addresses	
`Interface`	The ID of the interface.
`Address`	The IP address of the interface.

Sample Output

show system visibility jdm

```
user@jdm> show system visibility jdm
JDM CPU Usage Information
-------------------------
Percentage Used:  0.01

JDM Disk Usage Information
--------------------------
Total (MiB):      15829
Used  (MiB):      6275
Free  (MiB):      8728
Percentage Used:  39.6

JDM Memory Usage
-------------------
Maximum Memory (KiB): 1048576
Used Memory (KiB):    135180
Percentage Used:      13.6

JDM Uptime
----------
Uptime: 0 days 1:38:48

JDM Tasks
---------
Total:    14
```

```
Running:  1
Sleeping: 11
Stopped:  0
Zombie:   2

JDM Internal IP Addresses
-------------------------
Interface: bme1
Address  : 192.0.2.254
Interface: bme2
Address  : 192.168.1.254
```

Release Information

Command introduced in Junos OS Release 15.1X53-D40.

RELATED DOCUMENTATION

- show system visibility cpu
- show system visibility host
- show system visibility jcp | 27296
- show system visibility memory
- show system visibility network
- show system visibility storage
- show system visibility vnf

show system visibility memory

IN THIS SECTION

- Syntax | 27305
- Description | 27305
- Required Privilege Level | 27305
- Output Fields | 27305

Sample Output | 27306
Release Information | 27307

Syntax

```
show system visibility memory
```

Description

Display the details about virtual memory and shared memory for a Junos OS platform.

Required Privilege Level

view

Output Fields

Table 1699 on page 27305 lists the output fields for the `show system visibility memory` command. Output fields are listed in the approximate order in which they appear.

Table 1699: show system visibility memory Output Fields

Field Name	Field Description
Fields for Memory Information—Virtual Memory	
`Total`	The total amount of available virtual memory, in kibibytes (KiBs).
`Used`	The total amount of used virtual memory, in kibibytes (KiBs).
`Available`	The total amount of available virtual memory, in kibibytes (KiBs).
`Free`	The total amount of free virtual memory, in kibibytes (KiBs).

Table 1699: show system visibility memory Output Fields *(Continued)*

Field Name	Field Description
Percent Used	The percentage of buffer virtual memory used.
Fields for Memory Information—Swap Memory	
Total	The total amount of available swap memory, in kibibytes (KiBs).
Used	The total amount of used swap memory, in kibibytes (KiBs).
Free	The total amount of free swap memory, in kibibytes (KiBs).
Percent Used	The percentage of buffer swap memory used.

Sample Output

show system visibility memory (NFX150)

```
user@host> show system visibility memory
Memory Information
------------------
Virtual Memory:
---------------
Total       (KiB): 7946732
Used        (KiB): 3292908
Available   (KiB): 5844376
Free        (KiB): 4653824
Percent Used     : 26.50
```

show system visibility memory (NFX250 (NG))

```
user@host> show system visibility memory
Memory Information
```

```
------------------

Virtual Memory:
---------------
Total       (KiB): 15914412
Used        (KiB): 6723092
Available   (KiB): 10250492
Free        (KiB): 9191320
Percent Used     : 35.60

Huge Pages:
------------
Total 1GiB Huge Pages:      2
Free 1GiB Huge Pages:       0
Configured 1GiB Huge Pages: 0
Total 2MiB Huge Pages:      401
Free 2MiB Huge Pages:       1
Configured 2MiB Huge Pages: 0

Hugepages Usage:
-----------------------------------------------------------------------------------------------
---------
Name                                Type                                Used 1G Hugepages  Used 2M
Hugepages
--------------------------------- ---------------------------------- ------------------
------------------
srxpfe                              other process                       1                  400
ovs-vswitchd                        other process                       2                  0
```

Release Information

Command introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

- show system visibility cpu | 27279
- show system visibility host | 27286
- show system visibility network | 27308
- show system visibility vnf | 27316

show system visibility network

IN THIS SECTION

- Syntax | 27308
- Description | 27308
- Required Privilege Level | 27308
- Output Fields | 27308
- Sample Output | 27310
- Release Information | 27315

Syntax

```
show system visibility network
```

Description

Displays details such as the list of MAC addresses assigned to VNF interfaces, the list of internal IP addresses for VNFs, the list of virtual functions used by VNFs, and the list of VNF interfaces for a Junos OS platform.

Required Privilege Level

view

Output Fields

Table 1700 on page 27309 lists the output fields for the `show system visibility network` command. Output fields are listed in the approximate order in which they appear.

Table 1700: show system visibility network Output Fields

Field Name	Field Description
Fields for List of VNF MAC Addresses	
VNF	The name of the VNF.
MAC	The MAC address of the VNF.
Fields for List of VNF Internal IP Addresses	
VNF	The name of the VNF.
IP	The IP address of the VNF.
Fields for List of VNF Virtual Functions	
VNF	The name of the VNF.
PF	The names of the Physical Functions available.
VF	The names of the Virtual Functions available for each Physical Function.
Fields for List of Free Virtual Functions	
PF	The names of the Physical Functions available.
VF	The names of the Virtual Functions available for each Physical Function.
Reserved For	The owner type for the Virtual Functions.
Fields for List of VNF Interfaces	
VNF	The name of the VNF.

Table 1700: show system visibility network Output Fields *(Continued)*

Field Name	Field Description
Interface	The name of the interface.
Type	The type of interface.
Source	The connectivity source.
Model	The connectivity model.
MAC	The MAC address of the VNF.

Sample Output

show system visibility network (NFX150)

```
user@host> show system visibility network
VNF MAC Addresses
-----------------------------------------------------------
VNF                                        MAC
----------------------------------------- -----------------
centos1_ethdef0                            00:00:5E:00:53:9E
centos1_ethdef1                            00:00:5E:00:53:9F
centos1_eth2                               00:00:5E:00:53:A0
centos1_eth3                               00:00:5E:00:53:A1
centos2_ethdef0                            00:00:5E:00:53:A2
centos2_ethdef1                            00:00:5E:00:53:A3
centos2_eth2                               00:00:5E:00:53:A4
centos2_eth3                               00:00:5E:00:53:A5

VNF Internal IP Addresses
-------------------------------------------------------------
VNF                                        IP
----------------------------------------- ---------------
centos1                                    192.0.2.103
centos2                                    192.0.2.102
```

```
VNF Virtual Functions
-----------------------------------------------------------------
VNF                                        PF        VF
----------------------------------------- --------- ------------
l3_ge_1_0_4_vfdef0                         heth-0-1  0000:04:10:0
l2_ge_0_0_0_vfdef0                         heth-0-0  0000:04:10:1
l2_ge_0_0_0_vfdef1                         heth-0-0  0000:04:10:5
l2_ge_0_0_0_vfdef2                         heth-0-0  0000:04:11:1
l2_ge_0_0_0_vfdef3                         heth-0-0  0000:04:11:5
l3_ge_1_0_2_vfdef0                         heth-0-5  0000:07:10:0
l2_ge_0_0_2_vfdef0                         heth-0-2  0000:04:10:3
l2_ge_0_0_2_vfdef1                         heth-0-2  0000:04:10:7
l2_ge_0_0_2_vfdef2                         heth-0-2  0000:04:11:3
l2_ge_0_0_2_vfdef3                         heth-0-2  0000:04:11:7
l3_ge_1_0_1_vfdef0                         heth-0-4  0000:07:10:1
l2_ge_0_0_3_vfdef0                         heth-0-3  0000:04:10:2
l2_ge_0_0_3_vfdef1                         heth-0-3  0000:04:10:6
l2_ge_0_0_3_vfdef2                         heth-0-3  0000:04:11:2
l2_ge_0_0_3_vfdef3                         heth-0-3  0000:04:11:6

Free Virtual Functions
---------------------------------------
PF        VF           Reserved For
--------- ------------ ----------------
heth-0-0  0000:02:10.5 fpc,vnf
heth-0-0  0000:02:11.1 fpc,vnf
heth-0-0  0000:02:11.5 fpc,vnf
heth-0-0  0000:05:10.2 fpc,vnf
heth-0-0  0000:05:10.6 fpc,vnf
heth-0-0  0000:05:10.4 fpc,vnf
heth-0-0  0000:02:10.1 fpc,vnf

VNF Interfaces
-----------------------------------------------------------------------------------------------
VNF             Interface Type   Source             Model      MAC                  VLAN-ID
--------------  ---------------  -----------------  ---------- -------------------- --------
centos2         centos2_vnet6    network default    virtio     00:00:5e:00:53:a2    --
centos2         centos2_vnet7    bridge eth0br      virtio     00:00:5e:00:53:a3    --
centos2         centos2_eth2     bridge ovs-sys-br  virtio     00:00:5e:00:53:a4    199
centos2         centos2_eth3     bridge custom1     virtio     00:00:5e:00:53:a5    --
centos1         centos1_vnet4    network default    virtio     00:00:5e:00:53:9e    --
centos1         centos1_vnet5    bridge eth0br      virtio     00:00:5e:00:53:9f    --
```

```
centos1         centos1_eth2     bridge ovs-sys-br  virtio     00:00:5e:00:53:a0    100
centos1         centos1_eth3     bridge custom1     virtio     00:00:5e:00:53:a1    --

OVS Interfaces
---------------- ------
NAME              MTU
---------------- ------
custom1           1500
centos2_eth3      1500
centos1_eth3      1500
veth11            9200
l3_h_ge_1_0_0     9200
veth01            9200
ovs-sys-br        1500
centos1_eth2      1500
centos2_eth2      1500
```

show system visibility network (NFX250 (NG))

```
user@host> show system visibility network
VNF Virtual Functions
-----------------------------------------------------------------
VNF                                       PF        VF
---------------------------------------- --------- ------------
System_vfdef0                             hsxe0     0000:03:13:6
System_vfdef0                             hsxe1     0000:03:13:7

Free Virtual Functions
--------------------------------------
PF        VF           Reserved For
--------- ------------ ----------------
hsxe0     0000:03:10.0 vnf
hsxe0     0000:03:11.4 vnf
hsxe0     0000:03:10.2 vnf
hsxe0     0000:03:11.6 vnf
hsxe0     0000:03:10.4 vnf
hsxe0     0000:03:11.0 vnf
hsxe0     0000:03:10.6 vnf
hsxe0     0000:03:11.2 vnf
hsxe0     0000:03:12.2 vnf
hsxe0     0000:03:12.0 vnf
```

```
hsxe0     0000:03:13.4 vnf
hsxe0     0000:03:12.6 vnf
hsxe0     0000:03:13.2 vnf
hsxe0     0000:03:12.4 vnf
hsxe0     0000:03:13.0 vnf
hsxe1     0000:03:11.5 vnf
hsxe1     0000:03:10.1 vnf
hsxe1     0000:03:11.7 vnf
hsxe1     0000:03:10.3 vnf
hsxe1     0000:03:11.1 vnf
hsxe1     0000:03:10.5 vnf
hsxe1     0000:03:11.3 vnf
hsxe1     0000:03:10.7 vnf
hsxe1     0000:03:12.3 vnf
hsxe1     0000:03:13.5 vnf
hsxe1     0000:03:12.1 vnf
hsxe1     0000:03:13.3 vnf
hsxe1     0000:03:12.7 vnf
hsxe1     0000:03:13.1 vnf
hsxe1     0000:03:12.5 vnf

OVS Interfaces
------------------------
NAME              MTU
----------------- ------
dpdk1             1500
ovs-sys-br        1500
l3_h_ge_1_0_0     1500
dpdk0             1500
```

show system visibility network (NFX350)

```
user@host> show system visibility network
VNF MAC Addresses
-------------------------------------------------------------
VNF                                          MAC
----------------------------------------- -----------------
new_ethdef0                                  78:4F:9B:2B:2E:4B
new_ethdef1                                  78:4F:9B:2B:2E:4C

VNF Internal IP Addresses
```

```
-------------------------------------------------------
VNF                                      IP
---------------------------------------- ---------------
new                                      192.0.2.100

VNF Virtual Functions
-----------------------------------------------------------------
VNF                                      PF        VF
---------------------------------------- --------- ------------
new                                      hsxe0     0000:b7:02.1
new                                      hsxe0     0000:b7:02.2

Free Virtual Functions
---------------------------------------
PF        VF           Reserved For
--------- ------------ ----------------
hsxe0     0000:b6:02.3 vnf
hsxe0     0000:b6:03.5 fpc
hsxe0     0000:b6:02.1 vnf
hsxe0     0000:b6:02.2 vnf
hsxe0     0000:b6:03.6 fpc
hsxe0     0000:b6:02.4 vnf
hsxe0     0000:b6:02.5 vnf
hsxe0     0000:b6:02.6 vnf
hsxe0     0000:b6:02.7 vnf
hsxe1     0000:b6:06.4 vnf
hsxe1     0000:b6:06.5 vnf
hsxe1     0000:b6:06.6 vnf
hsxe1     0000:b6:06.7 vnf
hsxe1     0000:b6:06.1 vnf
hsxe1     0000:b6:06.2 vnf
hsxe1     0000:b6:06.3 vnf
hsxe1     0000:b6:07.5 fpc
hsxe1     0000:b6:07.6 fpc
hsxe2     0000:b6:0b.5 fpc
hsxe2     0000:b6:0b.6 fpc
hsxe2     0000:b6:0a.7 vnf
hsxe2     0000:b6:0a.6 vnf
hsxe2     0000:b6:0a.5 vnf
hsxe2     0000:b6:0a.4 vnf
hsxe2     0000:b6:0a.3 vnf
hsxe2     0000:b6:0a.2 vnf
hsxe2     0000:b6:0a.1 vnf
```

```
hsxe3     0000:b6:0e.3 vnf
hsxe3     0000:b6:0e.2 vnf
hsxe3     0000:b6:0e.1 vnf
hsxe3     0000:b6:0e.7 vnf
hsxe3     0000:b6:0e.6 vnf
hsxe3     0000:b6:0e.5 vnf
hsxe3     0000:b6:0e.4 vnf
hsxe3     0000:b6:0f.5 fpc
hsxe3     0000:b6:0f.6 fpc

VNF Interfaces
-----------------------------------------------------------------------------------------------
VNF                  Interface Type      Source       Model      MAC                  VLAN-ID
-------------------- --------- --------- ------------ ---------- -------------------- -------
new                  --        hostdev   --           --         78:4f:9b:2b:2e:4b    97
new                  --        hostdev   --           --         78:4f:9b:2b:2e:4c    4000

OVS Interfaces
------------------------
NAME               MTU
----------------- ------
dpdk0              9216
xdsl_eth0          9192
ovs-sys-br         9192
dpdk2              9216
dpdk1              9216
dpdk3              9216
l3_h_ge_1_0_0      9216
```

Release Information

Command introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

show system visibility cpu | 27279

show system visibility host | 27286

show system visibility memory | 27304

show system visibility vnf | 27316

show system visibility vnf

IN THIS SECTION

- Syntax | 27316
- Description | 27316
- Required Privilege Level | 27316
- Output Fields | 27316
- Sample Output | 27320
- Release Information | 27322

Syntax

```
show system visibility vnf vnf name
```

Description

If a VNF name is not specified, this command displays the details of all VNFs present in the system. Details include VNF memory usage, CPU statistics, the list of network interfaces, the list of disk files, per disk usage, per port I/O statistics, and media information, which includes details about CD-ROM and USB storage devices.

If a VNF name is specified, this command displays the details of that particular VNF. Details include VNF memory usage, CPU statistics, the list of network interfaces, the list of disk files, per disk usage, per port I/O statistics, and media information, which includes details about CD-ROM and USB storage devices.

Required Privilege Level

view

Output Fields

Table 1701 on page 27317 lists the output fields for the `show system visibility vnf` command. Output fields are listed in the approximate order in which they appear.

Table 1701: show system visibility vnf Output Fields

Field Name	Field Description
Fields for List of VNFs	
`ID`	ID of the VNF.
`Name`	Name of the VNF.
`State`	State of the VNF.
Fields for VNF Memory Usage	
`Name`	Name of the VNF.
`Maximum Memory`	The maximum amount of memory, in kibibytes (KiBs).
`Used Memory`	The total amount of used memory, in kibibytes (KiBs).
`Used 1G Hugepages`	The total number of 1G hugepages used.
`Used 2M Hugepages`	The total number of 2M hugepages used.
Fields for VNF CPU Stats	
`Name`	Name of the VNF.
`CPU Time`	The total CPU time, in seconds.
`System Time`	The amount of system CPU time, in seconds.
`User Time`	The amount of user CPU time, in seconds.
Fields for List of VNF MAC Addresses	

Table 1701: show system visibility vnf Output Fields *(Continued)*

Field Name	Field Description
VNF	Names of the VNFs.
MAC	MAC addresses of the VNFs.
Fields for List of VNF Internal IP Addresses	
VNF	Names of the VNFs.
IP	Internal IP addresses of the VNFs.
Fields for List of Virtual Functions per VNF	
VNF	Names of the VNFs.
PF	The names of the Physical Functions available.
VF	The names of the Virtual Functions available for each Physical Function.
Fields for the VNF Interfaces	
VNF	The name of the VNF.
Interface	The name of the interface.
Type	The type of interface.
Source	The connectivity source.
Model	The connectivity model.
MAC	The MAC address of the VNF.

Table 1701: show system visibility vnf Output Fields *(Continued)*

Field Name	Field Description
Fields for List of VNF Disk Information	
VNF	The name of the VNF.
Disk	The name of the disk.
File	The path to the disk.
Fields for List of VNF Disk Usage	
VNF	The name of the VNF.
Disk	The name of the disk.
Read Requests	The number of times a read operation has happened on the disk.
Bytes Read	The number of read bytes on the disk.
Write Requests	The number of times a write operation has happened on the disk.
Bytes Written	The number of bytes written on the disk.
Fields for List of VNF Port Statistics	
VNF	The name of the VNF.
Port	The name of the port.
Rcvd Bytes	The number of bytes received.
Rcvd Packets	The number of packets received.

Table 1701: show system visibility vnf Output Fields *(Continued)*

Field Name	Field Description
Rcvd Error	The number of errors received.
Rcvd Drop	The number of drops received.
Trxd Bytes	The number of bytes transferred.
Trxd Packets	The number of packets transferred.
Trxd Error	The number of errors transferred.
Trxd Drop	The number of drops transferred.

Sample Output

show system visibility vnf

```
user@host> show system visibility vnf
List of VNFs
-------------------------------------------------------
ID   Name                                  State
---- ------------------------------------- ---------------
5    centos                                Running
VNF Memory Usage
-----------------------------------------------------------------------------------------------
-------------------
Name                                   Maximum Memory (KiB)  Used Memory (KiB)  Used 1G
Hugepages  Used 2M Hugepages
------------------------------------- --------------------- ------------------
------------------ -----------------
centos                                 2097152               260741
0                  0
VNF CPU Statistics (Time in ms)
-------------------------------------------------------------------------------
```

```
Name                                   CPU Time           System Time  User Time
-------------------------------------- ------------------ ------------ ------------
centos                                 14029              3650         1540
VNF MAC Addresses
-------------------------------------------------------------
VNF                                       MAC
---------------------------------------- -----------------
centos_ethdef0                            E8:B6:C2:CC:66:9B
centos_ethdef1                            E8:B6:C2:CC:66:9C
VNF Internal IP Addresses
-------------------------------------------------------
VNF                                       IP
---------------------------------------- ---------------
centos                                    192.0.2.100
-----------------------------------------------------------------
VNF Virtual Functions
-----------------------------------------------------------------
VNF                                       PF        VF
---------------------------------------- --------- ------------
l2_ge_0_0_0_vfdef0                        heth-0-0  0000:02:10:1
l2_ge_0_0_0_vfdef1                        heth-0-0  0000:02:10:5
l2_ge_0_0_0_vfdef2                        heth-0-0  0000:02:11:1
l2_ge_0_0_0_vfdef3                        heth-0-0  0000:02:11:5
l2_ge_0_0_2_vfdef0                        heth-0-2  0000:02:10:3
l2_ge_0_0_2_vfdef1                        heth-0-2  0000:02:10:7
l2_ge_0_0_2_vfdef2                        heth-0-2  0000:02:11:3
l2_ge_0_0_2_vfdef3                        heth-0-2  0000:02:11:7
l3_ge_1_0_2_vfdef0                        heth-0-5  0000:05:10:0
l2_ge_0_0_1_vfdef0                        heth-0-1  0000:02:10:0
l2_ge_0_0_1_vfdef1                        heth-0-1  0000:02:10:4
l2_ge_0_0_1_vfdef2                        heth-0-1  0000:02:11:0
l2_ge_0_0_1_vfdef3                        heth-0-1  0000:02:11:4
l2_ge_0_0_3_vfdef0                        heth-0-4  0000:05:10:1
l2_ge_0_0_3_vfdef1                        heth-0-4  0000:05:10:3
l2_ge_0_0_3_vfdef2                        heth-0-4  0000:05:10:5
l2_ge_0_0_3_vfdef3                        heth-0-4  0000:05:10:7
l3_ge_1_0_1_vfdef0                        heth-0-3  0000:02:10:2
VNF Interfaces
-----------------------------------------------------------------------------------------------
-------
VNF                  Interface Type          Source       Model      MAC               IPv4-
address
-------------------- --------- ---------     ------------ ---------- -----------------
```

```
---------------------
centos               centos_vnet4 network default         virtio     e8:b6:c2:cc:66:9b --
centos               centos_vnet5 bridge     eth0br           virtio         e8:b6:c2:cc:66:9c
--
VNF Disk Information
-----------------------------------------------------------------------------------------------
--------
VNF                           Disk        File
----------------------------- -----------
---------------------------------------------------------------
centos                        vda         /var/public/centos-linux-1.img
centos                        hda         /var/public/vnf_config_data_vnf0
VNF Disk Usage
---------------------------------------------------------------------------
VNF                  Disk      Read Req   Read Bytes   Write Req  Write Bytes
-------------------- --------- ---------- ------------ ---------- ------------
centos               vda       5382       84654592     2068       4372480
centos               hda       15         37068        0          0
VNF Port Statistics
-----------------------------------------------------------------------------------------------
---------------------------
VNF                  Port      Rcvd Bytes   Rcvd Packets Rcvd Error Rcvd Drop Trxd Bytes   Trxd
Packets Trxd Error Trxd Drop
-------------------- --------- ------------ ------------ ---------- --------- ------------
------------ ---------- ---------
centos               centos_vnet4 572       11           0          0         850
7            0          0
centos               centos_vnet5 21729     258          0          395       0
0            0          0
VNF Media Information
-----------------------------------------------------------------------------------------------
---------------
VNF                           Media Disk         File
----------------------------- ----- ------------
---------------------------------------------------------------
vnf0                          CDROM hda          /var/public/vnf_config_data_vnf0
```

Release Information

Command introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

show system visibility cpu | 27279
show system visibility host | 27286
show system visibility memory | 27304
show system visibility network | 27308

show system yang package

IN THIS SECTION

- Syntax | 27323
- Description | 27323
- Options | 27324
- Required Privilege Level | 27324
- Output Fields | 27324
- Sample Output | 27325
- Release Information | 27325

Syntax

```
show system yang package
<package-name>
```

Description

Display YANG packages that are installed on the device and list the corresponding modules, action scripts, and translation scripts associated with the package. The output also includes the translation status for the package, which determines whether the device invokes the package's translation scripts during a commit operation.

Options

none Display information about all YANG packages that are installed on the device.

package-name Name of a specific YANG package for which to display information.

Required Privilege Level

view

Output Fields

Table 1702 on page 27324 lists the output fields for the `show system yang package` command. Output fields are listed in the approximate order in which they appear.

Table 1702: show system yang package Output Fields

Field Name	Field Description
`Package ID`	Package identifier.
`YANG Module(s)`	List of YANG modules and deviation modules associated with the package ID. Data models defined in the modules are merged with the Junos OS schema.
`Action Script(s)`	List of action scripts associated with the package ID.
`Translation Script(s)`	List of translation scripts associated with the package ID.
`Translation script status`	Specifies whether translation for that package is `enabled` or `disabled`. When translation is disabled, the device does not invoke the translation scripts for that package during a commit operation, and the changes made to the configuration statements and hierarchies added by that package are not translated into Junos OS syntax and committed as transient configuration changes.

Sample Output

show system yang package

```
user@host> show system yang package p1
Package ID            :p1
YANG Module(s)        :if.yang if-aggregate.yang if-show.yang
Action Script(s)      :if-show.py
Translation Script(s) :if.slax
Translation script status is enabled
```

Release Information

Command introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Managing YANG Packages, Modules, and Scripts on Devices Running Junos OS

Understanding the Management of Nonnative YANG Modules on Devices Running Junos OS

Disabling and Enabling YANG Translation Scripts on Devices Running Junos OS

show system ztp

IN THIS SECTION

- Syntax | 27326
- Description | 27326
- Required Privilege Level | 27326
- Output Fields | 27326
- Sample Output | 27330
- Release Information | 27331

Syntax

```
show system ztp
```

Description

This command displays the Zero Touch Provisioning (ZTP) and Secure Zero Touch Provisioning (SZTP) state information.

Required Privilege Level

view

Output Fields

For a description of the output fields, see Table 1703 on page 27326. Output fields are listed in the approximate order in which they appear. The state field can have multiple settings. The rest of the fields are self explanatory based on DHCP arguments provided by the server.

Table 1703: show system ztp Output Fields

Field Name	Description
`ZtpState`	ZTP state field values for starting: • `INITIALIZED`—ZTP is initializing. • `STARTED`—ZTP started running.
	ZTP state field values for image download: • `IMAGE_DOWNLOADING`—ZTP is downloading the next software image. • `IMAGE_DOWNLOADED`—ZTP is finished downloading the next software image. • `RETRY_IMAGE_DOWNLOAD`—ZTP is retrying image download. • `IMAGE_NOT_FOUND`—ZTP could not find the image at the specified location on the server.

Table 1703: show system ztp Output Fields *(Continued)*

Field Name	Description
	ZTP state field values for configuration download: • CONFIG_DOWNLOADING—ZTP is downloading the configuration. • CONFIG_DOWNLOADED—ZTP is finished downloading the configuration. • RETRY_CONFIG_DOWNLOAD—ZTP is retrying configuration download. • CONFIG_NOT_FOUND—ZTP could not find the configuration.
	ZTP state field values for upgrading configuration: • IMAGE_CONFIG_UPGRADING—ZTP got an image and a configuration from the server. • CONFIG_UPGRADING—ZTP is upgrading the configuration.
	ZTP state field values for upgrading image: • RETRY_IMAGE_UPGRADE—ZTP is retrying image upgrade. • IMAGE_CONFIG_UPGRADING—ZTP got an image and a configuration from the server. • IMAGE_UPGRADING—ZTP is downloading the image. • IMAGE_UPGRADED—ZTP is finished upgrading the image.
	ZTP state field values for scripts: • SCRIPT_UPGRADING—ZTP is running the script provided by server. • SCRIPT_UPGRADED—ZTP is finished upgrading the script. • SCRIPT_UPGRADE_SUCCEEDED—ZTP script upgrade finished with success. • SCRIPT_UPGRADE_FAILED—ZTP script upgrade finished with failure status.

Table 1703: show system ztp Output Fields *(Continued)*

Field Name	Description
	ZTP state field values for reboot: • `REBOOTING`—ZTP is rebooting the system. • `REBOOTED`—ZTP is finished rebooting the system.
	ZTP state field values for configuration commit: • `CONFIG_COMMIT_SUCCEEDED`—ZTP succeeded in committing user configuration. • `CONFIG_COMMIT_FAILED`—ZTP user configuration commit failed.
	ZTP state field values for finishing: • `FAILED`—ZTP failed. • `SUCCEEDED`—ZTP succeeded.
ZTPDefaultMode	On secure platforms running Junos OS, the default bootstrapping mode that is currently running. • `Secure`— SZTP is the default bootstrapping mode. • `Unsecure`— ZTP is the default bootstrapping mode.
	On secure platforms running Junos OS Evolved, the default bootstrapping mode that is currently running. • `Secure`— SZTP is the default bootstrapping mode. • `Unsecure`— ZTP is the default bootstrapping mode.
`ZtpInterface`	Name of interface.

Table 1703: show system ztp Output Fields *(Continued)*

Field Name	Description
ZtpPlatformDefault	On secure platforms running Junos OS, the default bootstrapping mode. • `Secure`— SZTP is the default bootstrapping mode. • `Unsecure`— ZTP is the default bootstrapping mode.
	On secure platforms running Junos OS Evolved, the default bootstrapping mode. • `Secure`— SZTP is the default bootstrapping mode. • `Unsecure`— ZTP is the default bootstrapping mode.
`FtpIpAddr`	IP address.
`DefaultRouter`	When the log server, NTP server, or FTP server are on a remote subnet, the value of `DefaultRouter` is used to configure a route to reach the servers.
`LogServers`	ZTP allows specification of a remote log server address. ZTP logs are then streamed to the remote log server.
`NtpServers`	ZTP allows specification of a remote NTP server address.
`TransferMode`	Options for `TransferMode` are ftp, tftp, http, or https
`ImageFileType`	It can be a symbolic link.
`ConfigFileName`	Configuration file name.
`ConfigUrl`	Configuration URL.

Table 1703: show system ztp Output Fields *(Continued)*

Field Name	Description
ConfigStatus	This field specifies whether the configuration file is downloading, is downloaded, or the download is being retried.
ZtpRetryCount	If the ZTP state machine, which applies the image and configuration, fails, the number of retries attempted.
DhcpRetryCount	If the DHCP state machine, which fetches parameters for ZTP from the DHCP server, fails, the number of times it retries.
ZTP State History	Lists the last 10 state transitions by Time (date and time) and Description or which state it was in then.

Sample Output

show system ztp

```
user@host> show system ztp
Attribute          Value
--------------------------
ZtpState           IMAGE_CONFIG_UPGRADING
ZtpInterface       vmb0
FtpIpAddr          10.10.213.1
HostName           sw-123
LogServers         [u'10.10.213.1']
NtpServers         [u'10.10.255.62', u'10.10.255.63']
TransferMode       tftp
ImageFileName      /ZTP_IMAGES/test.iso
ImageFileType      None
ImageUrl           tftp://17.17.213.1//ZTP_IMAGES/junos-evo-scapa.iso
ImageStatus        IMAGE_DOWNLOADED
ConfigFileName     /ZTP_CONFIG/sw-123.cfg
ConfigUrl          tftp://10.10.213.1//ZTP_CONFIG/sw-123.cfg
ConfigStatus       CONFIG_DOWNLOADED
ZtpRetryCount      0
DhcpRetryCount     0
```

```
ZTP State History(last 10 transitions)
   Time                        Description
   Fri Jun  5 22:35:40 2020   Started
   Fri Jun  5 22:36:46 2020   Initialized
   Fri Jun  5 22:37:08 2020   Discovering interfaces
   Fri Jun  5 22:37:31 2020   Querying DHCP Server
   Fri Jun  5 22:37:43 2020   DHCP query succeeded
   Fri Jun  5 22:41:46 2020   Upgrading image and config
```

show system ztp (Default Platform Boot Mode and Default Boot Mode that is currently running on the device)

```
user@host> show system ztp
ZtpPlatformDefault UnSecure
ZtpBootMode UnSecure
```

Release Information

Command introduced in Junos OS Evolved Release 19.1R1.

RELATED DOCUMENTATION

Zero Touch Provisioning

show task

IN THIS SECTION

- Syntax | 27332
- Description | 27332
- Options | 27332
- Required Privilege Level | 27333
- Output Fields | 27333

Sample Output | 27334
Release Information | 27335

Syntax

```
show task
    <logical-system (all | logical-system-name)>
    <task-name>
    io
    logical-system-mux
    memory
    replication
    scheduler-slip-history
    snooping
    summary
```

Description

Display routing protocol tasks on the Routing Engine.

Options

none	Display all routing protocol tasks on the Routing Engine.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
logical-system-mux	Display the logical router multiplexer process (lrmuxd) per-task information.
task-name	(Optional) Display information about running tasks for all tasks whose name begins with this string (for example, `BGP_Group_69_153` and `BGP_Group_70_153` are both displayed when you run the `show task BGP_Group` command).
io	Show i/o statistics for all tasks displayed.
memory	Show memory statistics for all tasks displayed.

replication	Show only replication tasks.
scheduler-slip-history	Show the number of scheduler slips (wherein the scheduler is unable to provide requested CPU time to a process) that have occurred since the last RPD start, including some scheduler slip details.
snooping	Show only snooping tasks.
summary	(Optional) Display summary information about running tasks.

Required Privilege Level

view

Output Fields

Table 1704 on page 27333 describes the output fields for the `show task` command. Output fields are listed in the approximate order in which they appear.

Table 1704: show task Output Fields

Field Name	Field Description
`Pri`	Current priority of the process. A lower number indicates a higher priority.
`Task Name`	Name of the task.
`Pro`	IP protocol number associated with the process.
`Port`	TCP or UDP port number associated with the task.
`So`	Socket number of the task.

Table 1704: show task Output Fields *(Continued)*

Field Name	Field Description
`Flags`	Flags for the task: • `Accept`—Task is waiting for incoming connections. • `Connect`—Task is waiting for a connection to be completed. • `Delete`—Task has been deleted and is being cleaned up. • `LowPrio`— Task will be dispatched to read its socket after other higher-priority tasks.

Sample Output

show task

```
user@host> show task
Pri Task Name                          Pro  Port So Flags
 10 IF
 15 LABEL
 15 ISO
 15 INET                                          7
 20 Aggregate
 20 RT
 30 ICMP                                  1       9
 39 ISIS I/O                                     12
 40 IS-IS                                        10
 40 BGP RT Background                               <LowPrio>
 40 BGP.0.0.0.0+179                          179 15 <Accept LowPrio>
 50 BGP_69.192.168.201.234+179               179 17 <LowPrio>
 50 BGP_70.192.168.201.233+179               179 16 <LowPrio>
 50 BGP_Group_69_153                                <LowPrio>
 50 BGP_Group_70_153                                <LowPrio>
 50 ASPaths
 60 KRT                                  255       1
 60 Redirect
 70 MGMT.local                                   14 <LowPrio>
 70 MGMT_Listen./var/run/rpd_mgmt                13 <Accept LowPrio>
 70 SNMP Subagent./var/run/sub_rpd.sock           8 <LowPrio>
```

```
 40 KRT IO task                                                  {krtio-th}
 40 krtio-th                                                     {krtio-th}
 60 krt solic client                    255       85 <ReadDisableWriteDisable>    {krtio-th}
 13 rsvp-iobagent./var/run/sub_rpd.sock           46 <WriteDisable> {rsvp-io}
 80 jtrace_jthr_task                    255       85                 {TraceThread}
```

show task scheduler-slip-history

```
user@host> show task scheduler-slip-history
Scheduler Slip Report:
Total number of Scheduler Slips: 11
Scheduler Slip Logs(last 64):
Timestamp           Total            User          System
10/13/21 09:19:08     5 sec     0 sec 382 usec 0 sec 0 usec
10/13/21 09:18:55     5 sec     0 sec 575 usec 0 sec 0 usec
10/13/21 09:18:26     5 sec     0 sec 483 usec 0 sec 0 usec
10/13/21 09:18:20     5 sec     0 sec 616 usec 0 sec 0 usec
10/13/21 09:18:14     5 sec     0 sec 471 usec 0 sec 0 usec
10/13/21 09:17:58     5 sec     0 sec 0 usec  0 sec 90 usec
10/13/21 09:17:52     5 sec     0 sec 0 usec  0 sec 6290 usec
10/13/21 09:17:46     5 sec     0 sec 543 usec 0 sec 0 usec
10/13/21 09:17:40     5 sec     0 sec 1368 usec 0 sec 0 usec
10/13/21 09:17:34     5 sec     0 sec 487 usec 0 sec 0 usec
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show task io | 27336

show task logical-system-mux | 27339

show task memory | 27342

show task io

IN THIS SECTION

- Syntax | 27336
- Syntax (EX Series Switches) | 27336
- Description | 27336
- Options | 27336
- Required Privilege Level | 27337
- Output Fields | 27337
- Sample Output | 27337
- Release Information | 27338

Syntax

```
show task io
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show task io
```

Description

Display I/O statistics for routing protocol tasks on the Routing Engine.

Options

none	Display I/O statistics for routing protocol tasks on the Routing Engine.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1705 on page 27337 describes the output fields for the `show task io` command. Output fields are listed in the approximate order in which they appear.

Table 1705: show task io Output Fields

Field Name	Field Description
`Task Name`	Name of the task.
`Reads`	Number of input ready notifications.
`Writes`	Number of output ready notifications.
`Rcvd`	Number of requests to the kernel for input.
`Sent`	Number of requests to the kernel for output.
`Dropped`	Number of sent requests that failed.

Sample Output

show task io

```
user@host> show task io
Task Name                        Reads  Writes    Rcvd    Sent Dropped
 LMP Client                          1       1       0       0       0
 IF                                  0       0       0       0       0
 INET6                               0       0       0       0       0
 INET                                0       0       0       0       0
 ISO                                 0       0       0       0       0
 Memory                              0       0       0       0       0
```

```
  RPD Unix Domain Server./var/ru       0       0       0       0       0
  RPD Unix Domain Server./var/ru       1       0       0       0       0
  RPD Unix Domain Server./var/ru       2       0       0       0       0
  RPD Server.0.0.0.0+666               0       0       0       0       0
  Aggregate                            0       0       0       0       0
  RT                                   0       0       0       0       0
  ICMP                                 0       0       0       0       0
  Router-Advertisement                 0       0       0       0       0
  ICMPv6                               0       0       0       0       0
  IS-IS I/O./var/run/ppmd_contro    1307       1       0       0       0
  l2vpn global task                    0       0       0       0       0
  IS-IS                                0       0       0       0       0
  BFD I/O./var/run/bfdd_control     1307       1       0       0       0
  TED                                  0       0       0       0       0
  ASPaths                              0       0       0       0       0
  Resolve tree 1                       0       0       0       0       0
  KStat                                0       0       0       0       0
  KRT Request                          0       0      63       0       0
  KRT Ifstate                        106       0     295       0       0
  KRT                                  0       0       0       0       0
  Redirect                             0       0       0       0       0
  KRT IO task                          0       0       0       0       0            {krtio-th}
  krtio-th                             0       0       0       0       0            {krtio-th}
  krt solic client                     0       1       0       0       0            {krtio-th}
  rsvp-io                          83826       0  117827  139682       0            {rsvp-io}
  jtrace_jthr_task                     0       0       0       0       0            {TraceThread}

...
```

Release Information

Command introduced before Junos OS Release 7.4.

show task logical-system-mux

IN THIS SECTION

- Syntax | 27339
- Description | 27339
- Options | 27339
- Required Privilege Level | 27340
- Output Fields | 27340
- Sample Output | 27341
- Release Information | 27341

Syntax

```
show task logical-system-mux <task-name>
    <io>
    <memory>
    <replication>
<summary>
```

Description

Display the logical router multiplexer process (lrmuxd) per-task information.

Options

none Display all the logical router multiplexer process (lrmuxd) per-task information.

task-name (Optional) Display information about running tasks for all tasks whose name begins with this string (for example, `lsys_session_task:lr2` and `lsys_session_task:lr1` are both displayed when you run the `show task logical-system-mux lsys` command).

io Show I/O statistics for all tasks displayed.

memory Show memory statistics for all lrmuxd processes displayed.

replication Show only replication tasks.

summary Display summary information about running tasks.

Required Privilege Level

view

Output Fields

Table 1706 on page 27340 describes the output fields for the `show task logical-system-mux` command. Output fields are listed in the approximate order in which they appear.

Table 1706: show task logical-system-mux Output Fields

Field Name	Field Description
`Pri`	Current priority of the process. A lower number indicates a higher priority.
`Task Name`	Name of the task.
`Pro`	IP number associated with the process.
`Port`	TCP or UDP port number associated with the task.
`So`	Socket number of the task.
`Flags`	Flags for the task: • `Accept`—Task is waiting for incoming connections. • `Connect`—Task is waiting for a connection to be completed. • `Delete`—Task has been deleted and is being cleaned up. • `LowPrio`— Task will be dispatched to read its socket after other higher-priority tasks.

Sample Output

show task logical-system-mux

```
user@host> show task logical-system-mux
Pri Task Name                          Pro  Port So Flags
 15 Memory
 40 lsys_session_task:lr2                        14 <WriteDisable>
 40 lsys_session_task:lr1                        11 <WriteDisable>
 40 lrmuxd lsys info task
 60 Mirror Task.128.0.0.6                  63793  9 <WriteDisable>
 70 MGMT.local                                   15 <WriteDisable>
 70 MGMT_Listen./var/run/lrmuxd_mgmt              6 <Accept WriteDisable>
```

show task logical-system-mux io

```
user@host> show task logical-system-mux io
Task Name                          Reads  Writes    Rcvd    Sent Dropped
 Memory                                0       0       0       0       0
 lsys_session_task:lr2                 7       2       0       0       0
 lsys_session_task:lr1                 7       2       0       0       0
 lrmuxd lsys info task                 0       0       0       0       0
 Mirror Task.128.0.0.6              1940       1       0       0       0
 MGMT.local                            0       0       1       0       0
 MGMT_Listen./var/run/lrmuxd_mg       12       0      12       0       0
```

show task logical-system-mux memory

```
user@host> show task logical-system-mux memory
Memory                 Size (kB)  Percentage  When
  Currently In Use:         1483          0%  now
  Maximum Ever Used:        1483          0%  13/03/20 02:28:18
  Available:             1589641        100%  now
```

Release Information

Command introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

show task | 27331
show task io | 27336
show task memory | 27342

show task memory

IN THIS SECTION

- Syntax | 27342
- Syntax (EX Series Switches) | 27342
- Description | 27343
- Options | 27343
- Required Privilege Level | 27343
- Output Fields | 27343
- Sample Output | 27346
- Release Information | 27348

Syntax

```
show task memory
<brief | detail | history | summary>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show task memory
<brief | detail | history | summary>
```

Description

Display memory utilization for routing protocol tasks on the Routing Engine.

> **NOTE**: The `show task memory` command does not display all the memory used by the routing protocol process. This value does not account for the memory used for the `TEXT` and `STACK` segments, or the memory used by the routing protocol process's internal memory manager.

Options

none	Display standard information about memory utilization for routing protocol tasks on the Routing Engine on all logical systems.
brief | detail | history | summary	(Optional) Display the specified level of output. Use the `history` option to display a history of memory utilization information.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1707 on page 27343 describes the output fields for the `show task memory` command. Output fields are listed in the approximate order in which they appear.

Table 1707: show task memory Output Fields

Field Name	Field Description	Level of Output
`Memory Currently In Use`	Memory currently in use. Dynamically allocated memory plus the `DATA` segment memory in kilobytes.	All levels
`Memory Maximum Ever Used`	Maximum memory ever used.	none specified, `brief`, `history`

Table 1707: show task memory Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Memory Available`	Memory currently available. **NOTE**: The maximum currently available memory is displayed incorrectly. On 32-bit Junos OS, the actual available memory is 2,097,152 kilobytes (2147483648 / 1024) but instead it is displayed as 2,147,484 kilobytes (2147483648 / 1000). On 64-bit Junos OS, the actual available memory is 3,145,728 kilobytes (3221225472 / 1024) but instead it is displayed as 3221225 kilobytes (3221225472 / 1000).	none specified, `brief`
`Size (kB)`	Memory capacity in 1000-byte kilobytes.	none specified, `brief`, `history`, `summary`
`Percentage`	Percentage of memory currently available.	none specified, `brief`
`When`	Timestamp.	none specified, `brief`, `history`
`Overall Memory Report`	Memory utilization by memory size: • `Size`—Block size, in bytes. • `TXP`—`T` indicates transient memory, `X` indicates exclusive memory, and `P` indicates full page. • `Allocs`—Number of blocks allocated for named objects. • `Mallocs`—Number of blocks allocated for anonymous objects. • `Alloc Bytes`—Number of blocks allocated times block size. • `MaxAllocs`—Maximum value of`Allocs`. • `MaxBytes`—Maximum value of `Alloc Bytes`. • *`FreeBytes`*—Total number of bytes unused on memory pages for this block size.	`detail`

Table 1707: show task memory Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Allocator Memory Report`	Memory utilization by named objects: • `Size`—Size of the named object in bytes. • `Alloc Size`—Actual memory used by that object in bytes. • `DTXP`—D indicates debug, `T` indicates transient memory, `X` indicates exclusive memory, and `P` indicates full page. • `Alloc Blocks`—Number of named objects allocated. • `AllocBytes`—Number of blocks allocated times block size. • `MaxAlloc Blocks`—Maximum value of `Alloc Blocks`. • `Max Alloc Bytes`—Maximum value of `AllocBytes`.	`detail`
`Malloc Usage Report`	Memory utilization for miscellaneous use: • `Allocs`—Number of allocations. • `Bytes`—Total bytes consumed. • `MaxAllocs`—Maximum value of `Allocs`. • `MaxBytes`—Maximum value of `Bytes`. • `FuncCalls`—Cumulative number of `Allocs`.	`detail`
`Dynamically allocated memory`	Memory allocated dynamically by the system.	`detail`
`Program data+BSS memory`	Program and base station subsystem (BSS) memory.	`detail`
`Page data overhead`	Internal memory overhead.	`detail`

Table 1707: show task memory Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Page directory size`	Internal memory overhead.	`detail`
`Total bytes in use`	Total memory, in bytes, that is currently in use and percentage of available memory (in parentheses).	`detail`

Sample Output

show task memory

```
user@host> show task memory
Memory                  Size (kB)  Percentage  When
  Currently In Use:        29417          3%  now
  Maximum Ever Used:       33882          4%  00/02/11 22:07:03
  Available:              756281        100%  now
```

show task memory detail

```
user@host> show task memory detail
------------------------- Overall Memory Report -------------------------
  Size TP    Allocs   Mallocs  AllocBytes  MaxAllocs    MaxBytes  FreeBytes
     8            -       111         888        112         896       3208
    12           92       149        2892        247        2964       1204
    12 T          -         -           -          5          60          -
    16            7        11         288         23         368       3808
    20          100        33        2660        164        3280       1436
    20 T          -         -           -         40         800          -
    24          162        15        4248        177        4248       3944
    24 T          -         -           -          4          96          -
    28          371         -       10388        372       10416       1900
    32            6        23         928         30         960       3168
...
---------------------------------------------------------------------------
                                   606182                  715302     118810
```

```
  ------------------------ Allocator Memory Report ------------------------
  Name                  Size Alloc DTP    Alloc     Alloc MaxAlloc  MaxAlloc
                             Size       Blocks     Bytes   Blocks     Bytes
  patroot                  8    12           84      1008       87      1044
  sockaddr_un.i802         8    12            2        24        2        24
  cos_nhm_nh               8    12            1        12        1        12
  sockaddr_un.tag          8    12            3        36        4        48
  gw_entry_list            8    12            1        12        1        12
  bgp_riblist_01           8    12            1        12        2        24
  ospf_intf_ev             8    12            -         -        6        72
  krt_remnant_rt           8    12  T         -         -        5        60
...
  ---------------------------------------------------------------------------
                                                   164108            221552

  -------------------------- Malloc Usage Report --------------------------
  Name                       Allocs     Bytes MaxAllocs  MaxBytes  FuncCalls
  MGMT.local                      1         8         1         8          1
  BGP.0.0.0.0+179                 -         -         1         8          2
  BGP RT Background               4     74748         4     74748          4
  SNMP Subagent./var/run/         -        52         1      9172         56
  OSPFv2 I/O./var/run/ppm         1     66536         2     66552       4551
  OSPF                            6     67655         7     67703         68
  KRT                             -         -         1      3784         18
  ASPaths                         3        80         3        80          3
  -- sockaddr --                183      2100       184      2108       1645
  BFD I/O./var/run/bfdd_c         1     65535         2     65551       4555
  RT                             48       872        48       872         48
  Scheduler                      42       628        43       628         88
  --Anonymous--                  56      1100        58      1140        112
  --System--                     82     58364       114     60044       4654
  ---------------------------------------------------------------------------
                                        337678               352398

    Dynamically allocated memory:     765952      Maximum:    765952
          Program data+BSS memory:   1568768      Maximum:   1568768
               Page data overhead:     53248      Maximum:     53248
              Page directory size:      4096      Maximum:      4096
                                   ----------
               Total bytes in use:   2392064 (0% of available memory)
```

Release Information

Command introduced before Junos OS Release 7.4.

show task replication

IN THIS SECTION

- Syntax | 27348
- Description | 27348
- Options | 27349
- Required Privilege Level | 27349
- Output Fields | 27349
- Sample Output | 27350
- Release Information | 27351

Syntax

```
show task replication
```

Description

Displays nonstop active routing (NSR) status. When you issue this command on the primary Routing Engine, the status of nonstop active routing synchronization is also displayed.

CAUTION: If BGP is configured, before attempting nonstop active routing switchover, check the output of `show bgp replication` to confirm that BGP routing table synchronization has completed on the backup Routing Engine. The `complete` status in the output of `show task replication` only indicates that the socket replication has completed and the BGP synchronization is in progress.

To determine whether BGP synchronization is complete, you must check the `Protocol state` and `Synchronization state` fields in the output of `show bgp replication` on the primary Routing Engine. The `Protocol state` must be `idle` and the `Synchronization state` must be `complete`. If you perform NSR switchover before the BGP synchronization has completed, the BGP session might flap.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1708 on page 27349 lists the output fields for the `show task replication` command. Output fields are listed in the approximate order in which they appear.

Table 1708: show task replication Output Fields

Field Name	Field Description
`Stateful replication`	Displays whether or not graceful Routing Engine switchover is configured. The status can be `Enabled` or `Disabled`.
`RE mode`	Displays the Routing Engine on which the command is issued: `Master`, `Backup`, or `Not applicable` (when the router has only one Routing Engine).
`Protocol`	Protocols that are supported by nonstop active routing.
`Synchronization Status`	Nonstop active routing synchronization status for the supported protocols. States are `NotStarted`, `InProgress`, and `Complete`. Synchronization states are shown for each of the supported protocols that are running on the device at that moment.

Sample Output

show task replication (Issued on the Primary Routing Engine)

```
user@host> show task replication
         Stateful Replication: Enabled
         RE mode: Master

     Protocol                Synchronization Status
     OSPF                    NotStarted
     BGP                     Complete
     IS-IS                   NotStarted
     LDP                     Complete
     PIM                     Complete
```

show task replication (Issued on the Backup Routing Engine)

```
user@host> show task replication
         Stateful Replication: Enabled
         RE mode: Backup
```

show task replication (Junos OS Evolved)

In Junos OS Evolved, both the primary and backup Routings have the same CLI output. If you configured any protocol, you should see the synchronization state for the same.

```
user@host> show task replication

         Stateful Replication: Enabled
         RE mode: Master

     Protocol                Synchronization Status
     OSPF                    NotStarted
     BGP                     Complete
     IS-IS                   NotStarted
     LDP                     Complete
     PIM                     Complete
```

Release Information

Command introduced in Junos OS Release 8.5.

Support for logical systems introduced in Junos OS Release 13.3

RELATED DOCUMENTATION

Example: Configuring Nonstop Active Routing on Switches

show tcp-forwarding status

IN THIS SECTION

- Syntax | 27351
- Description | 27351
- Options | 27352
- Required Privilege Level | 27352
- Output Fields | 27352
- Sample Output | 27355
- Release Information | 27357

Syntax

```
show tcp-forwarding status
listening-port port-number listening-address ipv4-listening-address
routing-instance routing-instance-name
```

Description

Display the status of TCP mapping and the current TCP connections for each mapping. You can limit the display to a specific listening port/listening address combination, per routing instance. If you do not specify a routing instance, the default routing instance is assumed.

Options

listening-address *ipv4-listening-address*

IPv4 address that is part of a listening port/listening address combination. The listening address is one on the BNG that external management systems or remote devices must use when attempting to trigger connections on the listening port. You must also specify a listening port.

listening-port *port-number*

Port number that is part of a listening port/listening address combination. The listening port is one that the BNG monitors for connections to be triggered by external management systems or remote devices.

- **Range:** 8000 through 8031

routing-instance *routing-instance-name*

Name of the routing instance for the TCP mapping.

Required Privilege Level

view

Output Fields

Table 1709 on page 27353 lists the output fields for the `show tcp-forwarding status` command. Output fields are listed in the approximate order in which they appear.

Table 1709: show tcp-forwarding status Output Fields

Field Name	Field Description

`Listening on`	Routing instance, IPv4 address, and port that the BNG is monitoring for connection triggers from external management systems or remote devices, in the following format:
	
	[*routing-instance-name*]:*ip-address*:*port-number*
	
	The following status and statistics are displayed for the TCP listening connection:
	
	• `Status`—Current state of the BNG regarding the TCP connection:
	  • `listening`—The connection is active.
	  • `not-listening`—The connection is inactive. This is typically indicative of a misconfiguration, such as an invalid listening address for the routing instance.
	
	• `Total Bytes`—Total number of bytes sent and received on the TCP listening connection.
	  • `Rx`—Received bytes.
	  • `Tx`—Transmitted bytes.

`Forwarding to`	Routing instance, IPv4 address, and port number where the BNG is forwarding data on the TCP connection, in the following format:
	
	[*routing-instance-name*]:*ip-address*:*port-number*
	
	The following statistics are displayed for the TCP port forwarding connection:
	
	• `Total Bytes`—Total number of bytes sent and received on the TCP port forwarding connection.
	  • `Rx`—Received bytes.
	  • `Tx`—Transmitted bytes.

Table 1709: show tcp-forwarding status Output Fields *(Continued)*

Field Name	Field Description
`Allowed Source Prefixes`	IPv4 source prefixes from which the BNG can accept connection requests. Request from all other addresses are rejected. A /32 mask indicates a single acceptable address. The allowed values are compared to the source address in the TCP header from the triggering entity.

`Connections`	Status and statistics for the configured TCP connections.
	
	• `Max`—Configured maximum number of connections allowed on the listening port.
	
	• `Active`—Number of TCP connections that are currently active on the listening port.
	
	• `Source`—List of source address and source port combinations with which the BNG has a current connection. The field value is displayed in the following format:
	
	`source-ip-address`:*`port-number`*
	
	• `Listening` and `Forwarding`—Status of the listening and forwarding connections.
	
	  • `connected`—The connection is active.
	
	  • `disconnected`—The connection is removed. This is a transient state, because the TCP connection pair is removed when disconnected.
	
	• `Bytes`—Number of bytes transmitted and received on the listening or forwarding connections for the source address/port combination.
	
	  • `Rx`—Received bytes.
	
	  • `Tx`—Transmitted bytes.

Sample Output

show tcp-forwarding status

```
user@host> show tcp-forwarding status
Listening on: [default:]203.0.113.50:8000
    Status: listening
    Total Bytes Rx: 0 Tx: 0
  Forwarding to: [default:]192.0.0.2:830
    Total Bytes Rx: 0 Tx: 0
  Allowed Source Prefixes:
    198.51.100.3/32
  Connections Max: 1 Active: 1
    Source: 198.51.100.3:55000
    Listening: connected Bytes Rx: 0 Tx: 0
    Forwarding: connected Bytes Rx: 0 Tx: 0

Listening on: [default:]203.0.113.50:8001
    Status: listening
    Total Bytes Rx: 0 Tx: 0
  Forwarding to: [default:]192.0.0.3:830
    Total Bytes Rx: 0 Tx: 0
  Allowed Source Prefixes:
    198.51.100.3/32
  Connections Max: 1 Active: 1
    Source: 198.51.100.3:55001
      Listening: connected Bytes Rx: 0 Tx: 0
      Forwarding: connected Bytes Rx: 0 Tx: 0

Listening on: [default:]203.0.113.50:8002
    Status: listening
    Total Bytes Rx: 0 Tx: 0
  Forwarding to: [default:]192.0.0.4:830
    Total Bytes Rx: 0 Tx: 0
  Allowed Source Prefixes:
    198.51.100.3/32
  Connections Max: 1 Active: 1
    Source: 198.51.100.3:55002
      Listening: connected Bytes Rx: 0 Tx: 0
      Forwarding: connected Bytes Rx: 0 Tx: 0

Listening on: [default:]203.0.113.50:8003
```

```
    Status: listening
    Total Bytes Rx: 0 Tx: 0
  Forwarding to: [default:]192.0.0.5:830
    Total Bytes Rx: 0 Tx: 0
  Allowed Source Prefixes:
    198.51.100.3/32
  Connections Max: 1 Active: 1
    Source: 198.51.100.3:55003
      Listening: connected Bytes Rx: 0 Tx: 0
      Forwarding: connected Bytes Rx: 0 Tx: 0

Listening on: [default:]203.0.113.50:8020
    Status: listening
    Total Bytes Rx: 0 Tx: 0
  Forwarding to: [default:]198.51.100.1:49
    Total Bytes Rx: 0 Tx: 0
  Allowed Source Prefixes:
    192.0.0.1/24
  Connections Max: 4 Active: 4
    Source: 192.0.0.2:55000
      Listening: connected Bytes Rx: 0 Tx: 0
      Forwarding: connected Bytes Rx: 0 Tx: 0
    Source: 192.0.0.3:55000
      Listening: connected Bytes Rx: 0 Tx: 0
      Forwarding: connected Bytes Rx: 0 Tx: 0
    Source: 192.0.0.4:56022
      Listening: connected Bytes Rx: 0 Tx: 0
      Forwarding: connected Bytes Rx: 0 Tx: 0
    Source: 192.0.0.5:55000
      Listening: connected Bytes Rx: 0 Tx: 0
      Forwarding: connected Bytes Rx: 0 Tx: 0
```

show tcp-forwarding status (Listening Port and Address)

```
user@host> show tcp-forwarding status listening-port 203.0.113.50 listening-address 8002
Listening on: [default:]203.0.113.50:8002
    Status: listening
    Total Bytes Rx: 0 Tx: 0
  Forwarding to: [default:]192.0.0.4:830
    Total Bytes Rx: 0 Tx: 0
  Allowed Source Prefixes:
```

```
    198.51.100.3/32
  Connections Max: 1 Active: 1
    Source: 198.51.100.3:55002
      Listening: connected Bytes Rx: 0 Tx: 0
      Forwarding: connected Bytes Rx: 0 Tx: 0
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

TCP Port Forwarding for Remote Device Management

show ted database

IN THIS SECTION

- Syntax | 27357
- Syntax (EX Series Switches) | 27358
- Description | 27358
- Options | 27358
- Required Privilege Level | 27358
- Output Fields | 27359
- Sample Output | 27363
- Release Information | 27372

Syntax

```
show ted database
<brief | detail | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

```
<system-name>
<topology-id topology bgp-ls-epe>
```

Syntax (EX Series Switches)

```
show ted database
<brief | detail | extensive>
<system-name>
```

Description

Display the entries in the Multiprotocol Label Switching (MPLS) traffic engineering database.

Options

none	Display standard information about all entries in the traffic engineering database.
brief | detail | extensive	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, information is displayed for the master instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
system-name	(Optional) Display traffic engineering database information for a particular system.
topology-id *topology*	Display the topology information. By default, traffic engineering topology information is displayed.

Required Privilege Level

view

Output Fields

Table 1710 on page 27359 describes the output fields for the `show ted database` command. Output fields are listed in the approximate order in which they appear.

Table 1710: show ted database Output Fields

Field Name	Field Description	Level of Output
TED database	Number of nodes and pseudonodes participating in IS-IS and OSPF domain routing.	All levels
ID	Hostname and address of the node that the link is coming from. An address of `.00` indicates that the node is the routing device itself. An address in the range `0.01` through `0.FF` indicates that the node is a pseudonode. If the node contains a router ID, it is displayed in parentheses.	`brief`
NodeID	Hostname and address of the node that the link is coming from. An address of `.00` indicates that the node is the routing device itself. An address in the range `0.01` through `0.FF` indicates that the node is a pseudonode.	`extensive`
Type	Type of node. It can be either `Rtr` (router) or `Net` (pseudonode).	All levels
Age(s)	How long since the node was last refreshed, in seconds.	All levels
LnkIn	Number of nodes pointing toward this node.	All levels
LnkOut	Number of nodes to which this node points.	All levels
Protocol	Protocol that reported the node information: • `IS-IS(1)`—IS-IS Level 1. • `IS-IS(2)`—IS-IS Level 2. • `OSPF (area-number)`—OSPF from the specified area.	All levels

Table 1710: show ted database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
To	Address on the far end of a link.	`detail extensive`
Local	Address of the local interface being used to reach the remote node.	`detail extensive`
Remote	Address of the interface on the remote node.	`detail extensive`
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	`detail extensive`
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	`detail extensive`
Metric	Configured traffic engineering metric.	`extensive`
IGP metric	Configured interior gateway protocol metric.	`extensive`
Static BW	Total interface bandwidth in bps.	`extensive`
Reservable bandwidth	Subscription factor for the interface, which is the percentage of the link bandwidth that can be used for the RSVP reservation process. You configure this by including the `subscription` statement when configuring RSVP.	`extensive`
Flood reflector client, cluster-id *<number>*	Displays flood-reflector related information on the TE links and the cluster ID that the you have connected at the client side.	`extensive`
Flood reflector, cluster-id *<number>*	Displays flood-reflector related information on the TE links and the cluster ID that you have connected in the flood reflector.	`extensive`

Table 1710: show ted database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Available BW [priority]`	(Must include `diffserv-te` statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each priority level. The bandwidth shown is for the entire interface, not for each individual LSP.	`extensive`
`Unidirectional link delay`	Displays unidirectional link delays.	`extensive`
`Minimum unidirectional link delay`	Displays minimum unidirectional link delays.	`extensive`
`Maximum unidirectional link delay`	Displays maximum unidirectional link delays.	`extensive`
`Unidirectional link delay variation`	Displays unidirectional link delay variation.	`extensive`
Diffserv-TE BW Model	Bandwidth constraint model used by the LSPs.	`extensive`
`Available BW [TE-class]`	(Must include the `diffserv-te` statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each traffic engineering class.	`extensive`
`Static BW [CT-class]`	Total interface bandwidth used by an MPLS traffic class, in bps.	`extensive`

Table 1710: show ted database Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Interface Switching Capability Descriptor (*n*)	Information about the interface switching capability descriptor, which is a subtype length value (TLV) of the link TLV. *n* is the index number.
	• `Switching type`—Type of switching to be performed on a particular link:
	  • `PSC-1`—Packet switch-capable 1
	  • `PSC-2`—Packet switch-capable 2
	  • `PSC-3`—Packet switch-capable 3
	  • `PSC-4`—Packet switch-capable 4
	  • `L2SC`—Layer-2-switch-capable
	  • `TDM`—Time-division-multiplexing-capable
	  • `LSC`—Lambda switch-capable
	  • `FSC`—Fiber switch-capable
	• `Encoding type`—Encoding of the LSP being requested:
	  • `Packet`
	  • `Ethernet`
	  • `ANSI/ETSI PDH`
	  • `Reserved`
	  • `SDH /SONET`
	  • `Digital Wrapper`
	  • `Lambda (photonic)`
	  • `Fiber`
	  • `FiberSDH/SONET`
		`extensive`

Table 1710: show ted database Output Fields *(Continued)*

Field Name	Field Description	Level of Output

	• `Maximum LSP BW [priority] bps`—Maximum LSP bandwidth information. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface.
	  • `[n]`—Priority level. The range is from `0` (high) through `7` (low).
	  • *n* `Mbps`—Amount of the maximum bandwidth.
	• `Minimum LSP BW`—Minimum LSP bandwidth in Mbps. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. `Minimum LSP BW` is displayed only when `switching type` is `PSC-1` or `TDM`.
	• `Interface MTU`—Displayed only when `switching type` is `TDM`.
	• `Interface supports standard SONET/SDH`—Displayed only when `switching type` is `TDM`.	

`P2P Adjacency SID`	Displays the IPv4 and IPv6 adjacency segment identifier information.	`extensive`
`LAN Adjacency SID`	Displays the IPv4 and IPv6 LAN adjacency segment identifier information.	`extensive`
`Prefixes`	Displays the IPv4 and IPv6 Prefix SID information.	`extensive`

Sample Output

show ted database brief

```
user@host> show ted database brief
TED database: 12 ISIS nodes 0 INET nodes
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-A.00                     ---    3178     2      0
Router-B.00                     ---    3152     2      0
Router-B.02                     Net     802     0      2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
```

```
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-C.00                      ---    3126     2      0
Router-C.02                      Net      38     0      2 IS-IS(2)
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-D.00                      ---    3144     2      0
Router-D.02                      Net     723     0      2 IS-IS(2)
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-D.03                      Net     607     0      2 IS-IS(2)
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-E.00                      ---    3178     2      0
Router-E.02                      Net     131     0      2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-F.00                      ---    3153     2      0
Router-F.02                      Net     769     0      2 IS-IS(2)
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
```

show ted database detail

```
TED database: 12 ISIS nodes 0 INET nodes
ID                            Type Age(s) LnkIn LnkOut Protocol
```

```
Router-A.00                      ---    2913     2      0
Router-B.00                      ---    2887     2      0
Router-B.02                      Net     537     0      2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-C.00                      ---    2861     2      0
Router-C.02                      Net     597     0      2 IS-IS(2)
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-D.00                      ---    2879     2      0
Router-D.02                      Net     458     0      2 IS-IS(2)
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-D.03                      Net     342     0      2 IS-IS(2)
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-E.00                      ---    2913     2      0
Router-E.02                      Net     640     0      2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-F.00                      ---    2888     2      0
Router-F.02                      Net     504     0      2 IS-IS(2)
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
```

show ted database extensive

The output of the `show ted database extensive` command on the reflector client.

```
user@host> show ted database extensive R21_re0
TED database: 32 ISIS nodes 12 INET nodes
NodeID: R21_re0.00(21.21.21.21)
  Protocol: IS-IS(2)
   21.21.21.21, 128.205.175.165
    To: R28_re0.00(28.28.28.28), Local: 10.21.28.21, Remote: 10.21.28.28
      Local interface index: 71, Remote interface index: 71
      Color: 0 <none>
      Flood reflector client, cluster-id 28
      Metric: 1000
      IGP metric: 1000
      Static BW: 18.4467Ebps
      Reservable BW: 18.4467Ebps
      Available BW [priority] bps:
          [0] 18.4467Ebps  [1] 18.4467Ebps [2] 18.4467Ebps [3] 18.4467Ebps
          [4] 18.4467Ebps  [5] 18.4467Ebps [6] 18.4467Ebps [7] 18.4467Ebps
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 18.4467Ebps  [1] 18.4467Ebps [2] 18.4467Ebps [3] 18.4467Ebps
          [4] 18.4467Ebps  [5] 18.4467Ebps [6] 18.4467Ebps [7] 18.4467Ebps
```

The output of the `show ted database extensive` command on the flood reflector.

```
user@host> show ted database extensive R28_re0
TED database: 32 ISIS nodes 12 INET nodes
NodeID: R28_re0.00(28.28.28.28)
  Type: Rtr, Age: 278 secs, LinkIn: 8, LinkOut: 8
  Protocol: IS-IS(2)
   28.28.28.28, 128.205.174.178
    To: R21_re0.00(21.21.21.21), Local: 10.21.28.28, Remote: 10.21.28.21
      Local interface index: 71, Remote interface index: 71
      Color: 0 <none>
      Flood reflector, cluster-id 28
      Metric: 1000
      IGP metric: 1000
```

```
      Static BW: 18.4467Ebps
      Reservable BW: 18.4467Ebps
      Available BW [priority] bps:
          [0] 18.4467Ebps  [1] 18.4467Ebps [2] 18.4467Ebps [3] 18.4467Ebps
          [4] 18.4467Ebps  [5] 18.4467Ebps [6] 18.4467Ebps [7] 18.4467Ebps
      Unidirectional link delay:
      Minimum unidirectional link delay:
      Maximum unidirectional link delay:
      Unidirectional link delay variation:
     Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 18.4467Ebps  [1] 18.4467Ebps [2] 18.4467Ebps [3] 18.4467Ebps
          [4] 18.4467Ebps  [5] 18.4467Ebps [6] 18.4467Ebps [7] 18.4467Ebps
```

The enhanced output of the `show ted database extensive` command to include IPv6 SR-MPLS attributes. The following sample output shows the traffic-engineering database information specific to one node.

```
user@host> show ted database VMX_R2_re.00 extensive
TED database: 3 ISIS nodes 2 INET nodes 2 INET6 nodes
NodeID: VMX_R2_re.00(2.2.2.2)(2::2)
  Type: Rtr, Age: 216 secs, LinkIn: 1, LinkOut: 1
  Protocol: IS-IS(2)
   2.2.2.2, 128.204.49.202
    To: VMX_R2_re.02, Local: 23.0.0.1, Remote: 0.0.0.0
      IPv6 Local: 23::1, IPv6 Remote: ::
      Local interface index: 333, Remote interface index: 0
      Color: 0 <none>
      Metric: 10
      IGP metric: 10
      Static BW: 1000Mbps
      Reservable BW: 1000Mbps
      Available BW [priority] bps:
          [0] 1000Mbps     [1] 1000Mbps    [2] 1000Mbps    [3] 1000Mbps
          [4] 1000Mbps     [5] 1000Mbps    [6] 1000Mbps    [7] 1000Mbps
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 1000Mbps     [1] 1000Mbps    [2] 1000Mbps    [3] 1000Mbps
          [4] 1000Mbps     [5] 1000Mbps    [6] 1000Mbps    [7] 1000Mbps
```

```
      LAN Adjacency-SID:
        IPV6, SID: 299824, Flags: 0xb0, Weight: 0, Neighbor: VMX_R3_re
        IPV4, SID: 299808, Flags: 0x30, Weight: 0, Neighbor: VMX_R3_re
   Prefixes:
     192.12.2.2/32
      IGP Flags: 0x00
      Prefix-SID:
        SID: 2002, Flags: 0x40, Algo: 0
     128.204.49.202/32
      IGP Flags: 0x00
     192::2/128
      IGP Flags: 0x00
      Prefix-SID:
        SID: 3002, Flags: 0x40, Algo: 0
     23::/64
      Metric: 10, IGP Flags: 0x00
     abcd::128:204:49:202/128
      IGP Flags: 0x00
   SPRING-Capabilities:
     SRGB block [Start: 800000, Range: 4096, Flags: 0xc0]
   SPRING-Algorithms:
     Algo: 0
     Algo: 1
```

show ted database extensive topology-type l3-unicast

```
user@host> show ted database extensive topology-type l3-unicast
SPRING Flex-Algorithms:
     Flex-Algo: 128
        Metric: 2, Calc: 1, Prio: 10
        Flags: 0X80000000, Inc-Any: 0x00000002, Inc-All: 0x00000002, Exclude: 0x00000004
```

show ted database extension

```
user@host> show ted database extension
TED database: 3 ISIS nodes 3 INET nodes 0 INET6 nodes
NodeID: R0.00(100.100.100.0)
  Type: Rtr, Age: 15 secs, LinkIn: 0, LinkOut: 0
  Protocol: IS-IS(2)
   100.100.100.0, 128.221.133.20
```

```
   Prefixes:
     100.100.100.0/32
      Flags: 0x00
      Prefix-SID:
        SID: 100, Flags: 0x40, Algo: 0
     100.100.100.1/32
      Metric: 10, Flags: 0x00
      Prefix-SID:
        SID: 10, Flags: 0xe0, Algo: 0
        SID: 780, Flags: 0xe0, Algo: 129
      Flex-Algo metric:
Algo: 129, Metric: 10
     122.2.2.2/32
      Flags: 0x00
      Prefix-SID:
        SID: 707, Flags: 0x40, Algo: 129
     128.221.133.20/32
      Flags: 0x00
   SRV6 Locators:
     4000::/64, Metric: 0, Flags: 0x0, Algorithm: 129
       4000::7:8:9, 0x00, USP
SPRING-Capabilities:
     SRGB block [Start: 400000, Range: 5000, Flags: 0xc0]
   SPRING-Algorithms:
      Algo: 0
      Algo: 1
      Algo: 129
    SPRING Flex-Algorithms:
        Flex-Algo: 129
             Metric: 0, Calc: 0, Prio: 129
             Flags: 0x02, Inc-Any: 0x00040000, Exclude: 0x00008000, Inc-All: 0x00004000
        Flex-Algo: 130
             Metric: 0, Calc: 0, Prio: 130
             Flags: 0x02, Inc-Any: 0x00040000, Exclude: 0x00008000, Inc-All: 0x00004000
          ............
         ..............
```

show ted database topology-id igp

```
user@host> show ted database topology-id igp
TED database: 3 ISIS nodes 3 INET nodes
```

```
ID                            Type Age(s) LnkIn LnkOut Protocol
Router A.00(128.220.1.2)        Rtr     193     2      2 IS-IS(2)
    To: Router B.00(128.220.18.198), Local: 2.3.0.2, Remote: 2.3.0.1
      Local interface index: 334, Remote interface index: 336
    To: Router B.00(128.220.18.198), Local: 2.3.1.2, Remote: 2.3.1.1
      Local interface index: 333, Remote interface index: 335
ID                            Type Age(s) LnkIn LnkOut Protocol
Router C.00(128.220.1.52)       Rtr     193     2      2 IS-IS(2)
    To: Router B.00(128.220.18.198), Local: 1.2.0.1, Remote: 1.2.0.2
      Local interface index: 335, Remote interface index: 334
    To: Router B.00(128.220.18.198), Local: 1.2.1.1, Remote: 1.2.1.2
      Local interface index: 334, Remote interface index: 333
ID                            Type Age(s) LnkIn LnkOut Protocol
Router B.00(128.220.18.198)     Rtr     193     4      4 IS-IS(2)
    To: Router A.00(128.220.1.2), Local: 2.3.0.1, Remote: 2.3.0.2
      Local interface index: 336, Remote interface index: 334
    To: Router A.00(128.220.1.2), Local: 2.3.1.1, Remote: 2.3.1.2
      Local interface index: 335, Remote interface index: 333
    To: Router C.00(128.220.1.52), Local: 1.2.0.2, Remote: 1.2.0.1
      Local interface index: 334, Remote interface index: 335
    To: Router C.00(128.220.1.52), Local: 1.2.1.2, Remote: 1.2.1.1
      Local interface index: 333, Remote interface index: 334
```

show ted database topology-id bgp-ls-epe extensive

```
user@host> show ted database topology-id bgp-ls-epe extensive
TED database: 0 ISIS nodes 3 INET nodes
NodeID: 4.4.4.4                                << DUT router-id
  Type: Rtr, Age: 270 secs, LinkIn: 0, LinkOut: 3
  Protocol: BGP-LS-EPE(0)                           << Protocol
    To: 5.5.5.5, Local: 50.1.1.1, Remote: 50.1.1.2 << Peer router-id and local and remote
interface used for BGP session.
      Local interface index: 0, Remote interface index: 0
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 0bps         [1] 0bps        [2] 0bps        [3] 0bps
          [4] 0bps         [5] 0bps        [6] 0bps        [7] 0bps
      BGP-Peer-SID:                             << BGP-Peer-SID Information
        SID: 1000007, Type: Node-SID Flags: 0x30, Weight: 0 << BGP-Node-SID
```

```
        SID: 1000002, Type: Set-SID Flags: 0x30, Weight: 0  << BGP-Set-SID
    To: 7.7.7.7, Local: 4.4.4.4, Remote: 7.7.7.7
      Local interface index: 0, Remote interface index: 0
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 0bps         [1] 0bps        [2] 0bps        [3] 0bps
          [4] 0bps         [5] 0bps        [6] 0bps        [7] 0bps
      BGP-Peer-SID:
        SID: 1000006, Type: Node-SID Flags: 0x30, Weight: 0  << BGP-Node-SID
    To: 7.7.7.7, Local: 4.4.4.4, Remote: 7.7.7.7
      Local interface index: 339, Remote interface index: 0
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 0bps         [1] 0bps        [2] 0bps        [3] 0bps
          [4] 0bps         [5] 0bps        [6] 0bps        [7] 0bps
      BGP-Peer-SID:
        SID: 1000005, Type: Adj-SID Flags: 0x30, Weight: 0 << BGP-ADj-SID
NodeID: 5.5.5.5
  Type: Rtr, Age: 270 secs, LinkIn: 1, LinkOut: 0
  Protocol: BGP-LS-EPE(0)
NodeID: 7.7.7.7
  Type: Rtr, Age: 270 secs, LinkIn: 2, LinkOut: 0
  Protocol: BGP-LS-EPE(0)
```

show ted database (IPv6)

```
user@host> show ted database
TED database: 2 ISIS nodes 0 INET nodes 2 INET6 nodes
ID                            Type Age(s) LnkIn LnkOut Protocol
R0_re.00(::1:1:1:1)         Rtr       1     1      1 IS-IS(1)
    To: R1_re.00(::2:2:2:2), Local: 1000:1::1:1, Remote: 1000:1::1:2
      Local interface index: 334, Remote interface index: 335
ID                            Type Age(s) LnkIn LnkOut Protocol
R1_re.00(::2:2:2:2)         Rtr       1    1      3 IS-IS(1)
    To: R0_re.00(::1:1:1:1), Local: 1000:1::1:2, Remote: 1000:1::1:1
      Local interface index: 335, Remote interface index: 334
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` *`instance-name`* option added in Junos OS Release 15.1.

`topology-id` *`topology`* option added in Junos OS Release 17.4R1 for MX Series and PTX Series.

Support for ipv6 introduced on Junos OS Release 20.4R1.

`Flood reflector client, cluster-id` *`<number>`* and `Flood reflector, cluster-id` *`<number>`* fields are introduced in the output field of the `show ted database extensive` command in Junos OS Release 21.2R1.

`Flood reflector client, cluster-id` *`<number>`* and `Flood reflector, cluster-id` *`<number>`* fields are introduced in the output field of the `show ted database extensive` command in Junos OS Evolved Release 22.1R1.

show-ted-ipv4-prefix

IN THIS SECTION

- Syntax | 27372
- Syntax (EX Series Switches) | 27373
- Description | 27373
- Options | 27373
- Required Privilege Level | 27373
- Output Fields | 27373
- Sample Output | 27374
- Release Information | 27374

Syntax

```
show ted ipv4-prefix
<instance instance-name>
<logical-system (all | logical-system-name)>
<topology-type topology bgp-ls-epe | express-segments | l3-unicast | traffic-engineering>
```

Syntax (EX Series Switches)

```
show ted ipv4-prefix
```

Description

Display the entries in the traffic engineering database IPv4 prefix tree.

Options

instance *instance-name* — (Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, information is displayed for the master instance.

logical-system (all | *logical-system-name*) — (Optional) Perform this operation on all logical systems or on a particular logical system.

topology-id *topology bgp-ls-epe | express-segments | l3-unicast | traffic-engineering* — Display the topology information. By default, traffic engineering topology information is displayed.

Required Privilege Level

view

Output Fields

Table 1711 on page 27373 lists the output fields for the `show ted ipv4-prefix` command.

Table 1711: show ted ipv4-prefix Output Fields

Field Name	Field Description
Prefix	IPv4 Prefix information for the node.
Node	Hostname and address of the node. If the node contains a router ID, it is displayed in parentheses. An address of `.00` indicates that the node is the routing device itself. An address in the range `0.01` through `0.FF` indicates that the node is a pseudonode.

Sample Output

show ted ipv4-prefix

```
user@host> show ted ipv4-prefix
Prefix                Node
100.100.100.0/32       R0.00(100.100.100.0)
100.100.100.1/32       R1.00(100.100.100.1)
128.207.200.16/32      R0.00(100.100.100.0)
128.207.201.212/32     R1.00(100.100.100.1)
```

Release Information

Command introduced in Junos OS Release 22.1R1.

RELATED DOCUMENTATION

show-ted-ipv6-prefix | 27374

IPv6 Prefixes and IPv6 Adjacency SIDs MPLS Support in Traffic Engineering Database and BGP Link-State

show-ted-ipv6-prefix

IN THIS SECTION

- Syntax | 27375
- Syntax (EX Series Switches) | 27375
- Description | 27375
- Options | 27375
- Required Privilege Level | 27375
- Output Fields | 27375
- Sample Output | 27376
- Release Information | 27376

Syntax

```
show ted ipv6-prefix
<instance instance-name>
<logical-system (all | logical-system-name)>
<topology-type topology bgp-ls-epe | express-segments | l3-unicast | traffic-engineering>
```

Syntax (EX Series Switches)

```
show ted ipv6-prefix
```

Description

Display the entries in the traffic engineering database IPv6 prefix tree.

Options

instance *instance-name*	(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, information is displayed for the master instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
topology-id *topology bgp-ls-epe | express-segments | l3-unicast | traffic-engineering*	Display the topology information. By default, traffic engineering topology information is displayed.

Required Privilege Level

view

Output Fields

Table 1712 on page 27376 lists the output fields for the `show ted ipv6-prefix` command.

Table 1712: show ted ipv6-prefix Output Fields

Field Name	Field Description
IPv6 Prefix	IPv6 Prefix information for the node.
Node	Hostname and address of the node. If the node contains a router ID, it is displayed in parentheses. An address of `.00` indicates that the node is the routing device itself. An address in the range `0.01` through `0.FF` indicates that the node is a pseudonode.

Sample Output

show ted ipv6-prefix

```
user@host> show ted ipv6-prefix
IPv6 Prefix                Node
1::1/128               R1.00(100.100.100.1)
10::10/128             R0.00(100.100.100.0)
21:0:1::1/128          R0.00(100.100.100.0)
21:0:1::2/128          R1.00(100.100.100.1)
6000::/64              R1.00(100.100.100.1)
abcd::128:207:200:16/128 R0.00(100.100.100.0)
abcd::128:207:201:212/128 R1.00(100.100.100.1)
```

Release Information

Command introduced in Junos OS Release 22.1R1.

RELATED DOCUMENTATION

show-ted-ipv4-prefix | **27372**

IPv6 Prefixes and IPv6 Adjacency SIDs MPLS Support in Traffic Engineering Database and BGP Link-State

show ted link

IN THIS SECTION

- Syntax | 27377
- Syntax (EX Series Switches) | 27377
- Description | 27377
- Options | 27378
- Required Privilege Level | 27378
- Output Fields | 27378
- Sample Output | 27380
- Release Information | 27383

Syntax

```
show ted link
<brief | detail>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show ted link
<brief | detail>
```

Description

Display Multiprotocol Label Switching (MPLS) traffic engineering database link information.

Options

none	Display standard information about traffic engineering database link information.
brief | detail	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, information is displayed for the master instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1713 on page 27378 describes the output fields for the `show ted link` command. Output fields are listed in the approximate order in which they appear.

Table 1713: show ted link Output Fields

Field Name	Field Description	Level of Output
ID	Hostname and address of the node that the link is coming from. An address of `.00` indicates that the node is the routing device itself. An address in the range `0.01` through `0.FF` indicates that the node is a pseudonode.	`brief`
-->ID	Hostname and address of the node that the link is going to. An address of `.00` indicates that the node is the routing device itself. An address in the range `0.01` through `0.FF` indicates that the node is a pseudonode.	`brief`
hostname	Hostname and address of the node that the link is coming from. An address of `.00` indicates that the node is the routing device itself. An address in the range `0.01` through `0.FF` indicates that the node is a pseudonode.	`detail`

Table 1713: show ted link Output Fields *(Continued)*

Field Name	Field Description	Level of Output
hostname	Hostname and address of the node that the link is going to. An address of `.00` indicates that the node is the routing device itself. An address in the range `0.01` through `0.FF` indicates that the node is a pseudonode.	`detail`
Local Path	Number of paths CSPF on the local routing device has placed on the link.	All levels
Metric	Configured traffic engineering metric.	`extensive`
IGP metric	Configured interior gateway protocol metric.	`detail`
LinkDelay	Configured link delay metrics.	`detail`
MinLinkDelay	Configured minimum link delay metrics.	`detail`
MaxLinkDelay	Configured maximum link delay metrics.	`detail`
LinkDelayVariation	Configured link delay variation.	`detail`
Local BW	Amount of bandwidth the local routing device has placed on the link.	All levels
Local	Address of the local interface being used to reach the remote node.	`detail extensive`
Remote	Address of the interface on the remote node.	`detail extensive`
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	`detail`
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	`detail`

Sample Output

show ted link brief

```
user@host> show ted link brief
ID                          ->ID                          LocalPath LocalBW
Router-B.02                     Router-A.00                           0 0bps
Router-B.02                     Router-B.00                           0 0bps
Router-C.02                     Router-B.00                           0 0bps
Router-C.02                     Router-C.00                           0 0bps
Router-D.02                     Router-F.00                           0 0bps
Router-D.02                     Router-D.00                           0 0bps
Router-D.03                     Router-D.00                           0 0bps
Router-D.03                     Router-C.00                           0 0bps
Router-E.02                     Router-A.00                           0 0bps
Router-E.02                     Router-E.00                           0 0bps
Router-F.02                     Router-E.00                           0 0bps
Router-F.02                     Router-F.00                           0 0bps
```

show ted link detail

```
user@host> show ted link detail
Router-B.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: 0bps
  LinkDelay: , MinLinkDelay: , MaxLinkDelay: , LinkDelayVariation:
      Color: 0 <none>
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-B.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 20 AvailBW: 0bps
  LinkDelay: , MinLinkDelay: , MaxLinkDelay: , LinkDelayVariation:
      Color: 0 <none>
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-C.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 40 AvailBW: 0bps
  LinkDelay: , MinLinkDelay: , MaxLinkDelay: , LinkDelayVariation:
```

```
      Color: 0 <none>
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-C.02->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: 0bps
  LinkDelay: , MinLinkDelay: , MaxLinkDelay: , LinkDelayVariation:
      Color: 0 <none>
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-D.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-D.02->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 60 AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-D.03->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-D.03->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-E.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 60 AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-E.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 20 AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-F.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: 0bps
```

```
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-F.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 40 AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
```

The enhanced output of the `show ted link detail` command to include IPv6 SR-MPLS attributes.

```
user@host> show ted link detail
VMX_R3_re.00(3.3.3.3)(3::3)->VMX_R2_re.02, Local: 23.0.0.2, Remote: 0.0.0.0
  IPv6 Local: 23::2, IPv6 Remote: ::
  Local interface index: 334, Remote interface index: 0
  LocalPath: 0, Metric: 10, IGP metric: 10, StaticBW: 1000Mbps, AvailBW: 1000Mbps
      Color: 0 <none>
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  AvailBW [0] 1000Mbps  [1] 1000Mbps  [2] 1000Mbps  [3] 1000Mbps
  AvailBW [4] 1000Mbps  [5] 1000Mbps  [6] 1000Mbps  [7] 1000Mbps
  IPV6 LAN-Adj-SID SID: 17 Flags: 0xb0 Weight: 0 Neighbor: VMX_R2_re
  IPV4 LAN-Adj-SID SID: 16 Flags: 0x30 Weight: 0 Neighbor: VMX_R2_re
VMX_R2_re.00(2.2.2.2)(2::2)->VMX_R2_re.02, Local: 23.0.0.1, Remote: 0.0.0.0
  IPv6 Local: 23::1, IPv6 Remote: ::
  Local interface index: 333, Remote interface index: 0
  LocalPath: 0, Metric: 10, IGP metric: 10, StaticBW: 1000Mbps, AvailBW: 1000Mbps
      Color: 0 <none>
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  AvailBW [0] 1000Mbps  [1] 1000Mbps  [2] 1000Mbps  [3] 1000Mbps
  AvailBW [4] 1000Mbps  [5] 1000Mbps  [6] 1000Mbps  [7] 1000Mbps
  IPV6 LAN-Adj-SID SID: 299824 Flags: 0xb0 Weight: 0 Neighbor: VMX_R3_re
  IPV4 LAN-Adj-SID SID: 299808 Flags: 0x30 Weight: 0 Neighbor: VMX_R3_re
VMX_R2_re.02->VMX_R3_re.00(3.3.3.3)(3::3), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
VMX_R2_re.02->VMX_R2_re.00(2.2.2.2)(2::2), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0
```

```
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
```

show ted link topology-id bgp-ls-epe detail

```
user@host> show ted link topology-id bgp-ls-epe detail
4.4.4.4->5.5.5.5, Local: 50.1.1.1, Remote: 50.1.1.2
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  SID: 1000007 Type: Node-SID Flags: 0x30 Weight: 0
  SID: 1000002 Type: Set-SID  Flags: 0x30 Weight: 0
4.4.4.4->7.7.7.7, Local: 4.4.4.4, Remote: 7.7.7.7
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  SID: 1000006 Type: Node-SID Flags: 0x30 Weight: 0
4.4.4.4->7.7.7.7, Local: 4.4.4.4, Remote: 7.7.7.7
  Local interface index: 339, Remote interface index: 0
  LocalPath: 0, AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  SID: 1000005 Type: Adj-SID Flags: 0x30 Weight: 0
```

show ted link (IPv6)

```
user@host> show ted link
R0_re.00(::1:1:1:1)->R1_re.00(::2:2:2:2), Local: 1000:1::1:1, Remote: 1000:1::1:2
  Local interface index: 334, Remote interface index: 335
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` *`instance-name`* option added in Junos OS Release 15.1.

Support for ipv6 introduced on Junos OS Release 20.4R1.

show ted protocol

IN THIS SECTION

- Syntax | 27384
- Syntax (EX Series Switches) | 27384
- Description | 27384
- Options | 27385
- Required Privilege Level | 27385
- Output Fields | 27385
- Sample Output | 27386
- Release Information | 27386

Syntax

```
show ted protocol
<brief | detail>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show ted protocol
<brief | detail>
```

Description

Display information about the protocols from which the Multiprotocol Label Switching (MPLS) traffic engineering database learned about its nodes.

Options

none	Display standard information about the protocols from which the traffic engineering database learned about its nodes.
brief | detail	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, information is displayed for the master instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1714 on page 27385 describes the output fields for the `show ted protocol` command. Output fields are listed in the approximate order in which they appear.

Table 1714: show ted protocol Output Fields

Field Name	Field Description
`Protocol name`	Protocol that reported the node information: • `IS-IS(1)`—IS-IS Level 1. • `IS-IS(2)`—IS-IS Level 2. • `OSPF` (*area-number*)—OSPF from the specified area.
`Credibility`	If the protocols provide conflicting information about a node, the protocol with the highest credibility value is the one that the traffic engineering database uses.
`Self node`	Address the protocol uses as the local address.

Sample Output

show ted protocol

```
user@host> show ted protocol
Protocol name        Credibility  Self node
IS-IS(2)             2  (highest) corriedale.00(123.456.1.11)
IS-IS(1)             1            corriedale.00(123.456.1.11)
```

command-name

```
user@host> show ted protocol topology-id bgp-ls-epe detail
Protocol name        Credibility  Self node
BGP-LS-EPE(0)        342          200.0.0.4
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance` *`instance-name`* option added in Junos OS Release 15.1.

show ted spring-te-policy

IN THIS SECTION

- Syntax | 27387
- Description | 27387
- Options | 27387
- Required Privilege Level | 27387
- Output Fields | 27387
- Sample Output | 27388
- Release Information | 27392

Syntax

```
show ted spring-te-policy
<brief | detail | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Display information about the TE policy stored in TED.

Options

none	Display standard information about the TE policy stored in TED.
brief | detail | extensive	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, information is displayed for the master instance.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

The Output fields of `show ted spring-te-policy` command are listed in the approximate order in which they appear.

Table 1715: show ted spring-te-policy Output Fields

Field Name	Field Description	Field Description
TED database	Number of nodes and pseudonodes participating in segment routing (SPRING). SPRING-TE	All levels
ID	Hostname and address of the node that the link is coming from.	All levels
Type	Type of node. It can be either Rtr (router) or Net (pseudonode).	All levels
Age(s)	How long since the node was last refreshed, in seconds.	All levels
Protocol	Protocol that reported the node information:	All levels
NodeID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	detail
Color	Fine-grained preference value of a static, aggregate, or generated route.	detail
Endpoint	ID of the system or node that is the destination of the source-routed path.	All levels

Sample Output

show ted spring-te-policy

```
user@host> show ted spring-te-policy
TED database: 1 ISIS nodes 0 INET nodes 0 INET6 nodes
                    ID                            Type Age(s)  Protocol
                    10.4.4.4(10.2.2.5)(abcd::128:49:2:181) Rtr    573 SPRING-TE(0)
    Color: 5120, Endpoint: abcd::17
    Color: 4120, Endpoint: 0.0.0.0
    Color: 5150, Endpoint: abcd::17
    Color: 3120, Endpoint: 17.1.1.1
```

```
    Color: 3121, Endpoint: 17.1.1.1
    Color: 6120, Endpoint: ::
```

show ted spring-te-policy extensive

```
user@host> show ted spring-te-policy extensive
TED database: 1 ISIS nodes 0 INET nodes 0 INET6 nodes
NodeID: 10.49.2.181.00(10.1.1.1)(abcd::128:49:2:181)
  Type: RtrAge: 846 secs
  Protocol: SPRING-TE(0)
    Color: 5120, Endpoint: abcd::17
        Candidate Path Name: static_srte_v6_policy1_1sgl_4segs_r7, Flags: 0x21a
          Origin: 3, Origin Address: 10.1.1.1, AS: 0
          Preference: 100, Priority: 128
          BSID Flags: 0x2, SR-MPLS BSID: 5120, SR-MPLS SPECIFIED BSID: 5120
          Number of Segment Lists: 1
            Segment List: 1, Segment List Flags: 0x1a, MTID: 0, Weight: 1, Algo: 0
            Number of Segments: 4
            Segments:
              SR-MPLS SID: 60210, Segment Flags: 0xf
              SR-MPLS SID: 60220, Segment Flags: 0xf
              SR-MPLS SID: 60250, Segment Flags: 0xf
              SR-MPLS SID: 60270, Segment Flags: 0xf
    Color: 4120, Endpoint: 0.0.0.0
        Candidate Path Name: bgp-srte-99.10.102.1-4120-0.0.0.0-4120, Flags: 0x29c
          Origin: 2, Origin Address: 99.10.102.1, AS: 65000
          Preference: 100, Priority: 128
          BSID Flags: 0x2, SR-MPLS BSID: 4120, SR-MPLS SPECIFIED BSID: 4120
          Number of Segment Lists: 1
            Segment List: 1, Segment List Flags: 0x1a, MTID: 0, Weight: 2, Algo: 0
            Number of Segments: 4
            Segments:
              SR-MPLS SID: 4272072206, Segment Flags: 0xf
              SR-MPLS SID: 4265781006, Segment Flags: 0xf
              SR-MPLS SID: 4263684366, Segment Flags: 0xf
              SR-MPLS SID: 4267944462, Segment Flags: 0xf
        Candidate Path Name: static_bgp_srte_v4_clronly_policy1_1sgl_4segs_r7, Flags: 0x21a
          Origin: 3, Origin Address: 10.1.1.1, AS: 0
          Preference: 100, Priority: 128
          BSID Flags: 0x2, SR-MPLS BSID: 4120, SR-MPLS SPECIFIED BSID: 4120
          Number of Segment Lists: 1
```

```
            Segment List: 1, Segment List Flags: 0x1a, MTID: 0, Weight: 1, Algo: 0
            Number of Segments: 4
            Segments:
              SR-MPLS SID: 60010, Segment Flags: 0xf
              SR-MPLS SID: 60020, Segment Flags: 0xf
              SR-MPLS SID: 60050, Segment Flags: 0xf
              SR-MPLS SID: 60070, Segment Flags: 0xf
Color: 5150, Endpoint: abcd::17
        Candidate Path Name: bgp-srte-99.10.102.1-5150-abcd::17-5150, Flags: 0x29a
          Origin: 2, Origin Address: 99.10.102.1, AS: 65000
          Preference: 100, Priority: 128
          BSID Flags: 0x2, SR-MPLS BSID: 5150, SR-MPLS SPECIFIED BSID: 5150
          Number of Segment Lists: 1
            Segment List: 1, Segment List Flags: 0x1a, MTID: 0, Weight: 2, Algo: 0
            Number of Segments: 4
            Segments:
              SR-MPLS SID: 4263686926, Segment Flags: 0xf
              SR-MPLS SID: 4274172686, Segment Flags: 0xf
              SR-MPLS SID: 4272076046, Segment Flags: 0xf
              SR-MPLS SID: 4276336142, Segment Flags: 0xf
    Color: 3120, Endpoint: 17.1.1.1
        Candidate Path Name: static_srte_v4_policy1_2sgls_4segs_r7, Flags: 0x21a
          Origin: 3, Origin Address: 10.1.1.1, AS: 0
          Preference: 100, Priority: 128
          BSID Flags: 0x2, SR-MPLS BSID: 3120, SR-MPLS SPECIFIED BSID: 3120
          Number of Segment Lists: 2
            Segment List: 1, Segment List Flags: 0x1a, MTID: 0, Weight: 2, Algo: 0
            Number of Segments: 4
            Segments:
              SR-MPLS SID: 60010, Segment Flags: 0xf
              SR-MPLS SID: 60020, Segment Flags: 0xf
              SR-MPLS SID: 60050, Segment Flags: 0xf
              SR-MPLS SID: 60070, Segment Flags: 0xf
            Segment List: 2, Segment List Flags: 0x1a, MTID: 0, Weight: 4, Algo: 0
            Number of Segments: 4
            Segments:
              SR-MPLS SID: 303, Segment Flags: 0xf
              SR-MPLS SID: 60040, Segment Flags: 0xf
              SR-MPLS SID: 60060, Segment Flags: 0xf
              SR-MPLS SID: 60070, Segment Flags: 0xf
    Color: 3121, Endpoint: 17.1.1.1
        Candidate Path Name: static_dynamic_srte_v4_policy2_1sgl_4segs_wo_bsid_r7, Flags: 0x1c
          Origin: 3, Origin Address: 10.1.1.1, AS: 0
```

```
          Preference: 100, Priority: 128
          Number of Segment Lists: 1
            Segment List: 1, Segment List Flags: 0x1a, MTID: 0, Weight: 1, Algo: 0
            Number of Segments: 4
            Segments:
              SR-MPLS SID: 60010, Segment Flags: 0xf
              SR-MPLS SID: 60020, Segment Flags: 0xf
              SR-MPLS SID: 60050, Segment Flags: 0xf
              SR-MPLS SID: 60070, Segment Flags: 0xf
        Candidate Path Name: 17.1.1.1:c31:dt-srte-dynamic_static_srte_v4_policy2_1sgl_4segs_r7,
Flags: 0x3a
          Origin: 3, Origin Address: 10.1.1.1, AS: 0
          Preference: 100, Priority: 128
          Number of Segment Lists: 1
            Segment List: 1, Segment List Flags: 0x1e, MTID: 0, Weight: 1, Algo: 0
            Number of Segments: 4
            Segments:
              SR-MPLS SID: 60030, Segment Flags: 0xf
              SR-MPLS SID: 60040, Segment Flags: 0xf
              SR-MPLS SID: 60060, Segment Flags: 0xf
              SR-MPLS SID: 60070, Segment Descriptor: 17.1.1.1, Segment Flags: 0xd
   Color: 6120, Endpoint: ::
        Candidate Path Name: bgp-srte-99.10.102.1-6120-::-6120, Flags: 0x29c
          Origin: 2, Origin Address: 99.10.102.1, AS: 65000
          Preference: 100, Priority: 128
          BSID Flags: 0x2, SR-MPLS BSID: 6120, SR-MPLS SPECIFIED BSID: 6120
          Number of Segment Lists: 1
            Segment List: 1, Segment List Flags: 0x1a, MTID: 0, Weight: 2, Algo: 0
            Number of Segments: 4
            Segments:
              SR-MPLS SID: 4263686926, Segment Flags: 0xf
              SR-MPLS SID: 4274172686, Segment Flags: 0xf
              SR-MPLS SID: 4272076046, Segment Flags: 0xf
              SR-MPLS SID: 4276336142, Segment Flags: 0xf
        Candidate Path Name: static_bgp_srte_v6_clronly_policy1_1sgl_4segs_r7, Flags: 0x21a
          Origin: 3, Origin Address: 10.1.1.1, AS: 0
          Preference: 100, Priority: 128
          BSID Flags: 0x2, SR-MPLS BSID: 6120, SR-MPLS SPECIFIED BSID: 6120
          Number of Segment Lists: 1
            Segment List: 1, Segment List Flags: 0x1a, MTID: 0, Weight: 1, Algo: 0
            Number of Segments: 4
            Segments:
              SR-MPLS SID: 60210, Segment Flags: 0xf
```

```
                SR-MPLS SID: 60220, Segment Flags: 0xf
                SR-MPLS SID: 60250, Segment Flags: 0xf
                SR-MPLS SID: 60270, Segment Flags: 0xf
```

Release Information

Command introduced in Junos OS Release 22.4R1

show trace

IN THIS SECTION

- Syntax | 27392
- Description | 27393
- Options | 27393
- Required Privilege Level | 27393
- Output Fields | 27393
- Sample Output | 27394
- Release Information | 27396

Syntax

```
show trace
<application app-name>
<live>
<node node-name>
<pid pid-value>
<terse>
<time time-elapsed>
```

Description

Show the trace data from all nodes that is collected on the primary Routing Engine in **/var/log/traces**. All applications are traced at the info level for informational messages. You can refine the traces to show by specifying trace time elapsed, application, process ID, and node.

The options provide you with a way to target the traces you want to see. The ouput will prompt you to use the options, like so:

```
[WARNING] Number of contributing trace folders is 2880 .
[WARNING] This might cause some logs not to be displayed.
[WARNING] Please filter your search using the available knobs (including time)
```

Options

none	Display all traces.
`application` *`app-name`*	(Optional) Display traces for the specified application name.
`live`	(Optional) Enable a mode in which the command remains active and new traces are displayed as they come in.
`node` *`node-name`*	(Optional) Display traces for the specified node name.
`pid` *`pid-value`*	(Optional) Display traces for the specified process ID.
`terse`	(Optional) Display briefer output for traces.
`time` *`time-elapsed`*	(Optional) Display traces for the specified elapsed time. • **Range:** 1 through 840 minutes

Required Privilege Level

view

Output Fields

Table 1716 on page 27394 lists the output fields for the `show trace` command. Output fields are listed in the approximate order in which they appear.

Table 1716: show trace Output Fields

Field Name	Field Description
timestamp	Timestamp field in the following format: YYYY-MM-DD HH:MM:SS.123456789.
node	Node where trace message originated.
system-process	System process where trace message originated.
tracepoint	Tracepoint value of the trace message.
trace-level	Trace level of the trace message.
application	Application where trace message originated.
message-type	Message type of the trace message.
Function	Function name where the trace message was generated.
Message	Message associated with the tracepoint.

Sample Output

show trace

```
root@evovbracklaq_RE0> show trace

[WARNING] Number of contributing trace folders is 2880 .
[WARNING] This might cause some logs not to be displayed.
[WARNING] Please filter your search using the available knobs (including time)


2019-09-26 08:46:29.658883645 re0:aft-sysinfo:14325 libevoinfra_INFO_APP  Function =
"evoapp_init_commons", node_type = "RE", node_slot = 0, node_name = "re0", app_name = "aft-
```

```
sysinfo", app_id = 0
2019-09-26 08:46:29.659055906 re0:aft-sysinfo:14325 libevoinfra_INFO_STR  Function =
"evoapp_init_commons", Message = "Running : /usr/sbin/aft-sysinfo -p /var/pfe --app-name aft-
sysinfo "
2019-09-26 08:46:29.659076131 re0:aft-sysinfo:14325 libevoinfra_INFO_2STR  Function =
"evoapp_init_commons", Message1 = "Object subscription mode", Message2 = "Object Select"
2019-09-26 08:46:29.659755689 re0:aft-sysinfo:14325 libevoinfra_INFO_2STR  Function =
"evoapp_load_dsl", Message1 = "App Lua config not set, using app file", Message2 = "/usr/conf/
evoapp/aft-sysinfo.lua"
2019-09-26 08:46:30.291258500 re0:aft-sysinfo:14325 lltp_info  message = "Setting up ZooClient
for app aft-sysinfo"
2019-09-26 08:46:30.291305775 re0:aft-sysinfo:14325 lltp_info  message = "Connecting to
Zookeeper : attempt 1"
2019-09-26 08:46:30.291422845 re0:aft-sysinfo:14325 lltp_info  message = "Zookeeper address
127.0.0.1:2181"
2019-09-26 08:46:30.291441778 re0:aft-sysinfo:14325 lltp_info  message = "Connecting to
Zookeeper: path 127.0.0.1:2181"
2019-09-26 08:46:30.308878435 re0:aft-sysinfo:14325 lltp_info  message = "Wait for Zookeeper
connection to get established"
2019-09-26 08:46:30.314930581 re0:aft-sysinfo:14325 lltp_info  message = "zookeeperWatcher:
event type ZOO_SESSION_EVENT state ZOO_CONNECTED_STATE path "
2019-09-26 08:46:30.314958284 re0:aft-sysinfo:14325 lltp_info  message = "Saving client id
10000015c6f0068 to aft-sysinfo"
2019-09-26 08:46:30.315649988 re0:aft-sysinfo:14325 lltp_info  message = "Async getConfig
completed path /zookeeper/config rc 0"
2019-09-26 08:46:31.309018911 re0:aft-sysinfo:14325 libevoinfra_INFO_STR  Function =
"evoapp_zoo_init", Message = "Connected to Zookeeper"
2019-09-26 08:46:31.312922419 re0:aft-sysinfo:14325 lltp_info  message = "Get Xapp static config
for aft-sysinfo"
2019-09-26 08:46:31.314689699 re0:aft-sysinfo:14325 libevoinfra_INFO_EVOAPP  Function =
"evoapp_zoo_init", Message = "App managed by SysMan", app_name = "aft-sysinfo", node_name =
"re0", app_version = 0, shared_app_version = 0, node_attr_match = ""
2019-09-26 08:46:31.321618770 re0:aft-sysinfo:14325 lltp_info  message = "Create node local
production set: Path /system/nodes/re0/apps/aft-sysinfo/infra/restart Version "
2019-09-26 08:46:31.321623149 re0:aft-sysinfo:14325 libevoinfra_INFO_EVOAPP  Function =
"evoapp_zoo_init", Message = "Generated app local version", app_name = "aft-sysinfo", node_name
= "re0", app_version = 3698991357035064051, shared_app_version = 0, node_attr_match = ""
2019-09-26 08:46:31.324009193 re0:aft-sysinfo:14325 lltp_info  message = "Create shared
production set: Path /system/apps/aft-sysinfo/restart Version 12377670930056552194"
. . .
```

show trace application live

```
user@host> show trace application cmdd live
2019-09-27 10:57:13.999923130 re0:cmdd:7955 lltp_info  message = "DdxClientConn::DdxClientConn:
name = cmdd, owner = 0x7f097a7d5c00, this = 0x7f097abd4600, stream = 0x7f097a797b80"
2019-09-27 10:57:13.999926928 re0:cmdd:7955 lltp_info  message = "DdxClientConn::start: name =
cmdd, owner = 0x7f097a7d5c00, this = 0x7f097abd4600, stream = 0x7f097a797b80"
```

Release Information

Command introduced in Junos OS Evolved Release 18.3R1.

RELATED DOCUMENTATION

clear trace | 4246

show unified-edge tdf aaa radius client statistics

IN THIS SECTION

- Syntax | 27396
- Description | 27397
- Options | 27397
- Required Privilege Level | 27397
- Output Fields | 27397
- Sample Output | 27399
- Release Information | 27404

Syntax

```
show unified-edge tdf aaa radius client statistics
<brief | detail>
```

```
<fpc-slot fpc-slot>
<gateway gateway>
<client name>
<pic-slot pic-slot>
```

Description

Display the statistics for the accounting packets transmitted and received from the RADIUS client for one or more TDF gateways. If a TDF gateway is not specified, then information for all TDF gateways is displayed.

Options

none	Display statistics for all TDF gateways.
brief | detail	(Optional) Display the specified level of output.
client *name*	(Optional) Display statistics for the specified RADIUS client.
fpc-slot *fpc-slot*	(Optional) Display statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway*	(Optional) Display statistics for the specified TDF gateway.
pic-slot *pic-slot*	(Optional) Display statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

view

Output Fields

Table 1717 on page 27398 lists the output fields for the `show unified-edge tdf aaa radius client statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1717: show unified-edge tdf aaa radius client statistics Output Fields

Field Name	Field Description	Level of Output
`Client`	Name of the RADIUS client.	All levels
`Gateway Name`	Name of the TDF gateway.	All levels
`FPC/PIC`	FPC and PIC slot numbers for which the statistics are displayed.	`detail`
`Accounting Requests`	Number of accounting requests sent to the RADIUS server from the FPC slot and PIC slot. The following information is displayed about each request type: • `Start`—Number of Accounting Start requests sent. • `Stop`—Number of Accounting Stop requests sent. • `Interim`—Number of Accounting Interim-Update requests sent. • `On`—Number of Accounting On requests sent. • `Off`—Number of Accounting Off requests sent.	All levels
`Accounting Responses`	Number of accounting responses sent to the RADIUS server from the FPC slot and PIC slot. The following information is displayed about each request type: • `Start`—Number of Accounting Start responses sent. • `Stop`—Number of Accounting Stop responses sent. • `Interim`—Number of Accounting Interim-Update responses sent. • `On`—Number of Accounting On responses sent. • `Off`—Number of Accounting Off responses sent.	All levels
`Duplicate Requests`	Number of duplicate accounting requests sent to the RADIUS server.	All levels
`Malformed Requests`	Number of malformed accounting requests sent to the RADIUS server.	All levels

Table 1717: show unified-edge tdf aaa radius client statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Bad Authenticators	Number of responses received from the RADIUS server with bad authenticators.	All levels
Unknown Types	Number of unknown type responses (that the TDF gateway does not recognize) received from the RADIUS server.	All levels
Dropped Packets	Number of packets dropped.	All levels

Sample Output

show unified-edge tdf aaa radius client statistics brief

```
user@host> show unified-edge tdf aaa radius client statistics brief
Client: pgwclient
Gateway Name: TDF
    Accounting Requests: 8
        Start: 8
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Accounting Responses: 8
        Start: 8
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0

Client: pgwclient_jrad
```

```
Gateway Name: TDF
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Accounting Responses: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0

Client: pgwclient_jrad1
Gateway Name: TDF
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Accounting Responses: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0
```

show unified-edge tdf aaa radius client statistics detail

```
user@host> show unified-edge tdf aaa radius client statistics detail
Client: pgwclient
Gateway Name: TDF
FPC/PIC: 2/0
    Accounting Requests: 8
        Start: 8
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Accounting Responses: 8
        Start: 8
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0

Client: pgwclient
Gateway Name: TDF
FPC/PIC: 2/1
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Accounting Responses: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
```

```
    Unknown Types: 0
    Dropped Packets: 0

Client: pgwclient_jrad
Gateway Name: TDF
FPC/PIC: 2/0
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Accounting Responses: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0

Client: pgwclient_jrad
Gateway Name: TDF
FPC/PIC: 2/1
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Accounting Responses: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
```

```
    Dropped Packets: 0

Client: pgwclient_jrad1
Gateway Name: TDF
FPC/PIC: 2/0
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Accounting Responses: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0

Client: pgwclient_jrad1
Gateway Name: TDF
FPC/PIC: 2/1
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Accounting Responses: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
```

```
    Unknown Types: 0
    Dropped Packets: 0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf aaa radius client statistics | 4248

show unified-edge tdf aaa statistics | 27425

show unified-edge tdf aaa radius client status

IN THIS SECTION

- Syntax | 27404
- Description | 27405
- Options | 27405
- Required Privilege Level | 27405
- Output Fields | 27405
- Sample Output | 27406
- Release Information | 27406

Syntax

```
show unified-edge tdf aaa radius client status
<fpc-slot fpc-slot>
<gateway gateway>
<client name>
<pic-slot pic-slot>
```

Description

Display the status of the RADIUS client for one or more TDF gateways. If a TDF gateway is not specified, then information for all TDF gateways is displayed.

Options

none	Display RADIUS client status for all TDF gateways.
client *name*	(Optional) Display the status for the specified RADIUS client.
fpc-slot *fpc-slot*	(Optional) Display the status for the specified Flexible PIC Concentrator (FPC).
gateway *gateway*	(Optional) Display the status for the specified TDF gateway.
pic-slot *pic-slot*	(Optional) Display the status for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

view

Output Fields

Table 1718 on page 27405 lists the output fields for the `show unified-edge tdf aaa radius client status` command. Output fields are listed in the approximate order in which they appear.

Table 1718: show unified-edge tdf aaa radius client status Output Fields

Field Name	Field Description
`Client`	Name of the RADIUS client.
`FPC/PIC`	FPC and PIC slot numbers for which the statistics are displayed.
`Address`	IP address of the RADIUS client.

Table 1718: show unified-edge tdf aaa radius client status Output Fields *(Continued)*

Field Name	Field Description
`Last activity`	Day of the week, month, date, time, and year when the last operation occurred on the RADIUS client. The term `No activity` is displayed if no communication occurred between the RADIUS client and the TDF gateway.

Sample Output

show unified-edge tdf aaa radius client status

```
user@host> show unified-edge tdf aaa radius statistics accounting brief
Client       FPC/PIC  Address              Last activity
-----------------------------------------------------------------
pgwclient    2/0      192.0.2.3           Mon Jul 21 11:00:16 2014
pgwclient_j  2/0      198.51.100.2            No activity
pgwclient_j  2/0      203.0.113.1            No activity
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf aaa radius client statistics | 4248

show unified-edge tdf aaa statistics | 27425

show unified-edge tdf aaa radius network-element statistics

IN THIS SECTION

Syntax | 27407

- Description | 27407
- Options | 27407
- Required Privilege Level | 27408
- Output Fields | 27408
- Sample Output | 27409
- Release Information | 27410

Syntax

```
show unified-edge tdf aaa radius network-element statistics
<brief | detail>
<fpc-slot fpc-slot>
<gateway gateway>
<name name>
<pic-slot pic-slot>
```

Description

Display RADIUS network element statistics. If a TDF gateway is not specified, then statistics for all TDF gateways are displayed.

Options

none Display statistics for all TDF gateways.

brief | detail (Optional) Display the specified level of output.

The `brief` option is the default and displays the consolidated statistics for all TDF gateways, and the `detail` option displays the statistics for each Services PIC on the configured TDF gateways.

fpc-slot *fpc-slot* (Optional) Display statistics for the specified Flexible PIC Concentrator (FPC).

gateway *gateway* (Optional) Display statistics for the specified TDF gateway.

name *name* (Optional) Display statistics for the specified network element.

pic-slot *pic-slot* (Optional) Display statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

view

Output Fields

Table 1719 on page 27408 lists the output fields for the `show unified-edge tdf aaa radius network-element statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1719: show unified-edge tdf aaa radius network-element statistics Output Fields

Field Name	Field Description	Level of Output
`Network-element`	Name of the network element to which the statistics belong.	All levels
`FPC/PIC`	FPC and PIC slot numbers for which statistics are displayed.	`detail`
`Requests Attempted`	Number of access and accounting requests that were attempted.	All levels
`Access Requests Sent`	Number of access requests sent.	All levels
`Accounting Requests Sent`	Number of accounting requests sent.	All levels
`Responses Received`	Number of access and accounting response messages received.	All levels
`Request Timeouts`	Number of access and accounting requests to the RADIUS server that timed out.	All levels

Table 1719: show unified-edge tdf aaa radius network-element statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Memory Failures	Number of internal memory allocation failures.	All levels
Invalid State Errors	Number of access requests and accounting requests that were attempted in non-operational state.	All levels
No Radius Server Found	Number of access requests and accounting requests that failed because no more RADIUS servers were available.	All levels
Source Port allocation Errors	Number of access and accounting requests that failed because of source port allocation failure for outgoing RADIUS messages.	All levels
Send Failures	Total number of failed attempts to send access requests and accounting requests.	All levels

Sample Output

show unified-edge tdf aaa radius network-element statistics detail

```
user@host> show unified-edge tdf aaa radius network-element statistics detail
  Network-element: ne1
  FPC/PIC:         5/2
    Requests Attempted:            0
    Access Requests Sent:          0
    Accounting Requests Sent:      0
    Responses Received:            0
    Request Timeouts:              0
    Memory Failures:               0
    Invalid State Errors:          0
```

```
    No Radius Server Found:        0
    Source Port allocation Errors: 0
    Send Failures:                 0

  Network-element: ne2
  FPC/PIC:         5/2
    Requests Attempted:            0
    Access Requests Sent:          0
    Accounting Requests Sent:      0
    Responses Received:            0
    Request Timeouts:              0
    Memory Failures:               0
    Invalid State Errors:          0
    No Radius Server Found:        0
    Source Port allocation Errors: 0
    Send Failures:                 0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf aaa radius network-element statistics | 4250

Understanding Network Elements

show unified-edge tdf aaa radius server statistics

IN THIS SECTION

- Syntax | 27411
- Description | 27411
- Options | 27411
- Required Privilege Level | 27412
- Output Fields | 27412

Sample Output | 27415
Release Information | 27416

Syntax

```
show unified-edge tdf aaa radius server statistics
<brief | detail>
<fpc-slot fpc-slot>
<gateway gateway>
<name name>
<pic-slot pic-slot>
```

Description

Display RADIUS server statistics. If a TDF gateway is not specified, then statistics for all TDF gateways are displayed.

Options

none Display the same output as the `brief` option.

brief | detail (Optional) Display the specified level of output.

The `brief` option is the default and displays the consolidated statistics for all TDF gateways, and the `detail` option displays the statistics for each Services PIC on the configured TDF gateways.

fpc-slot *fpc-slot* (Optional) Display statistics for the specified Flexible PIC Concentrator (FPC).

gateway *gateway* (Optional) Display statistics for the specified TDF gateway.

name *name* (Optional) Display statistics for the specified RADIUS server.

pic-slot *pic-slot* (Optional) Display statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

view

Output Fields

Table 1720 on page 27412 lists the output fields for the `show unified-edge tdf aaa radius server statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1720: show unified-edge tdf aaa radius server statistics Output Fields

Field Name	Field Description	Level of Output
`RADIUS server`	Name of the RADIUS server.	All levels
`Address`	IP address of the RADIUS server.	All levels
`Routing-instance`	Routing-instance of RADIUS server's source address.	`detail`
`Authentication Statistics`		
`Port`	RADIUS server port number to which access requests are sent.	All levels
`FPC/PIC`	FPC and PIC slot numbers for which the statistics are displayed.	`detail`
`Access requests`	Number of access requests sent to the RADIUS server.	All levels
`Access req retransmissions`	Number of access requests retransmitted to the RADIUS server.	All levels
`Access accepts`	Number of access accepts sent by the RADIUS server.	All levels
`Access rejects`	Number of access requests rejected by the RADIUS server.	All levels
`Malformed responses`	Number of malformed access responses received from the RADIUS server.	All levels

Table 1720: show unified-edge tdf aaa radius server statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Bad authenticators`	Number of bad authentication responses received for access-requests.	All levels
`Timeouts`	Number of access requests to the RADIUS server that timed out.	All levels
`Unknown types`	Number of unknown type responses received from the RADIUS server for access requests.	All levels
`Packets dropped`	Number of packets dropped for access requests and responses.	All levels
`Accounting Statistics`		
`Port`	RADIUS server port number to which accounting requests are sent.	All levels
`Accounting requests`	Number of accounting requests sent to the RADIUS server. The following information is displayed about each request type for the `detail` level: • `Start`—Number of accounting start requests sent. • `Stop`—Number of accounting stop requests sent. • `Interim`—Number of accounting interim-update requests sent. • `On`—Number of accounting on requests sent. • `Off`—Number of accounting off requests sent.	All levels
`Accounting req retransmissions`	Number of accounting requests retransmitted to the RADIUS server.	All levels

Table 1720: show unified-edge tdf aaa radius server statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting responses`	Number of accounting responses received from the RADIUS server.	All levels
`Malformed responses`	Number of malformed accounting responses received from the RADIUS server.	All levels
`Bad authenticators`	Number of bad accounting responses received for accounting requests.	All levels
`Timeouts`	Number of accounting requests to the RADIUS server that timed out.	All levels
`Unknown types`	Number of unknown type responses (that the TDF gateway does not recognize) received from the RADIUS server for accounting requests.	All levels
`Packets dropped`	Number of packets dropped for accounting requests and responses.	All levels
`Dynamic Authorization Request Statistics`		
`CoA requests received`	Number of change of authorization (CoA) requests received from the RADIUS server.	All levels
`DM requests received`	Number of Disconnect Message (DM) requests received from the RADIUS server.	All levels
`CoA Acks sent`	Number of CoA acknowledgements sent to the RADIUS server.	All levels
`CoA Nacks sent`	Number of CoA negative acknowledgements sent to the RADIUS server.	All levels

Table 1720: show unified-edge tdf aaa radius server statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
DM Acks sent	Number of DM acknowledgements sent to the RADIUS server.	All levels
DM Nacks sent	Number of DM negative acknowledgements sent to the RADIUS server.	All levels
Dropped	Number of dynamic authorization requests dropped.	All levels

Sample Output

show unified-edge tdf aaa radius server statistics detail

```
user@host> show unified-edge tdf aaa radius server statistics detail
RADIUS server: radius1 (FPC/PIC: 5/2)
Address: 192.0.2.2
Routing-instance: default
Authentication Statistics:
    Port: 1812
    Access requests: 0
    Access req retransmissions: 0
    Access accepts: 0
    Access rejects: 0
    Malformed responses: 0
    Bad authenticators: 0
    Timeouts: 0
    Unknown types: 0
    Packets dropped: 0
Accounting Statistics:
    Port: 1813
    Accounting requests: 0
      Start: 0      Stop: 0      Interim: 0      On: 0      Off: 0
    Accounting req retransmissions: 0
    Accounting responses: 0
    Malformed responses: 0
    Bad authenticators: 0
    Timeouts: 0
```

```
    Unknown types: 0
    Packets dropped: 0
Dynamic Authorization Request Statistics:
    CoA requests received: 0
    DM requests received: 0
    CoA Acks sent: 0
    CoA Nacks sent: 0
    DM Acks sent: 0
    DM Nacks sent: 0
    Dropped: 0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf aaa radius server statistics | 4252

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

show unified-edge tdf aaa radius server status

IN THIS SECTION

- Syntax | 27417
- Description | 27417
- Options | 27417
- Required Privilege Level | 27417
- Output Fields | 27417
- Sample Output | 27419
- Release Information | 27420

Syntax

```
show unified-edge tdf aaa radius server status
<brief | detail>
<fpc-slot fpc-slot>
<gateway gateway>
<name name>
<pic-slot pic-slot>
```

Description

Display RADIUS server status. If a TDF gateway is not specified, then statistics for all TDF gateways are displayed.

Options

none	(Same as brief) Display consolidated statistics for all TDF gateways.
brief | detail	(Optional) Display the specified level of output. The `brief` option is the default.
fpc-slot *fpc-slot*	(Optional) Display statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway*	(Optional) Display statistics for the specified TDF gateway.
pic-slot *pic-slot*	(Optional) Display statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

view

Output Fields

Table 1721 on page 27418 lists the output fields for the `show unified-edge tdf aaa radius server status` command. Output fields are listed in the approximate order in which they appear.

Table 1721: show unified-edge tdf aaa radius server status Output Fields

Field Name	Field Description	Level of Output
`Server`	Name of the RADIUS server.	`brief`
`RADIUS Server`	Name of the RADIUS server.	`detail`
`FPC/PIC`	FPC and PIC slot numbers for which the statistics are displayed.	All levels
`Address`	IP address of the RADIUS server.	All levels
`State`	State of the RADIUS server: `Active` or `Inactive` (dead).	All levels
`Duration`	Duration, in *weeks:days:MM:SS* format, for which the RADIUS server has been in the current state.	All levels
`Previous duration`	Duration, in *HH:MM:SS* format, for which the RADIUS server was in the previous state.	All levels
`Flaps`	Number of times that the RADIUS server transitioned from the active to inactive state.	All levels
`Authentication Information`		
`Pending requests`	Number of access requests waiting for responses from the RADIUS server.	`detail`
`Round trip time (ms)`	Time taken to receive the response from the RADIUS server for access requests. The minimum, maximum, and average round-trip times are also displayed.	`detail`

Table 1721: show unified-edge tdf aaa radius server status Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Accounting Information		
Pending requests	Number of accounting requests waiting for responses from the RADIUS server.	detail
Round trip time (ms)	Time taken to receive the response from the RADIUS server for accounting requests. The minimum, maximum, and average round-trip times are also displayed.	detail

Sample Output

show unified-edge tdf aaa radius server status brief

```
user@host> show unified-edge tdf aaa radius server status brief
             FPC/                                   Previous
Server       PIC     Address          State   Duration      Duration    Flaps
-------------------------------------------------------------------------------
radius1      5/2     192.0.2.2        Active  1w5d 23:12    00:00:00    0
radius2      5/2     198.51.100.100   Active  1w5d 23:12    00:00:00    0
radius3      5/2     203.0.113.100     Active  1w5d 23:12    00:00:00    0
radius4      5/2     203.0.113.100     Active  1w5d 23:12    00:00:00    0
```

show unified-edge tdf aaa radius server status detail

```
user@host> show unified-edge tdf aaa radius server status detail
RADIUS server: pgwcl (FPC/PIC: 4/0)
Address           : 198.51.100.100
State             : Active
Duration          : 1w6d 11:29
Previous Duration : 00:00:00
Flaps             : 0
```

```
Authentication Information:
    Pending requests     : 0
    Round trip time (ms) : 1 (Min: 1 Max: 1 Avg: 1)
Accounting Information:
    Pending requests     : 0
    Round trip time (ms) : 0 (Min: 0 Max: 0 Avg: 0)
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf aaa radius server statistics | 27410

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

show unified-edge tdf aaa radius snoop-segment statistics

IN THIS SECTION

- Syntax | 27420
- Description | 27421
- Options | 27421
- Required Privilege Level | 27421
- Output Fields | 27421
- Sample Output | 27423
- Release Information | 27425

Syntax

```
show unified-edge tdf aaa radius snoop-segment statistics
<brief | detail>
<fpc-slot fpc-slot>
```

```
<gateway gateway>
<pic-slot pic-slot>
<segment snoop-segment-name>
```

Description

Display statistics for snoop segments. If a snoop segment is not specified, then statistics for all snoop segments are displayed.

Options

none	(Same as brief) Display statistics for all snoop segments for all TDF gateways.
brief | detail	(Optional) Display the specified level of output. The `brief` option displays the consolidated statistics for all TDF gateways, and the `detail` option displays the statistics for each Services PIC on the configured TDF gateways.
fpc-slot *fpc-slot*	(Optional) Display statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway*	(Optional) Display statistics for the specified TDF gateway.
pic-slot *pic-slot*	(Optional) Display statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.
segment *snoop-segment-name*	(Optional) Display statistics for the specified snoop segment.

Required Privilege Level

view

Output Fields

Table 1722 on page 27422 lists the output fields for the `show unified-edge tdf aaa radius snoop-segment statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1722: show unified-edge tdf aaa radius snoop-segment statistics Output Fields

Field Name	Field Description	Level of Output
`Snoop-segment`	Name of the snoop-segment for which statistics are displayed.	All levels
`Gateway Name`	Name of the TDF gateway. If the statistics for all TDF gateways are displayed, then `All` is displayed.	All levels
`FPC/PIC`	FPC and PIC slot numbers for which the statistics are displayed.	`detail`
`Accounting Requests`	The following information is displayed for each category: • `Start`—Number of snooped Accounting Start requests. • `Interim`—Number of snooped Accounting Interim-Update requests. • `Stop`—Number of snooped Accounting Stop requests. • `On`—Number of snooped Accounting On requests. • `Off`—Number of snooped Accounting Off requests.	All levels
`Duplicate Requests`	Number of duplicate snooped accounting requests.	All levels
`Malformed Requests`	Number of snooped malformed accounting requests.	All levels
`Bad Authenticators`	Number of snooped acccounting requests with bad authenticators.	All levels
`Unknown Types`	Number of snooped accounting requests of unknown type.	All levels
`Dropped Packets`	Number of snooped packets dropped.	All levels

Sample Output

show unified-edge tdf aaa radius snoop-segment statistics brief

```
user@host> show unified-edge tdf aaa radius snoop-segment statistics brief

Snoop-segment: 123
Gateway Name: TDF
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0

Snoop-segment: dummy
Gateway Name: TDF
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0
```

show unified-edge tdf aaa radius snoop-segment statistics detail

```
user@host> show unified-edge tdf aaa radius snoop-segment statistics detail
Snoop-segment: 123
Gateway Name: TDF
FPC/PIC: 4/0
```

```
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0

Snoop-segment: 123
Gateway Name: TDF
FPC/PIC: 4/1
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0

Snoop-segment: dummy
Gateway Name: TDF
FPC/PIC: 4/0
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0

Snoop-segment: dummy
```

```
Gateway Name: TDF
FPC/PIC: 4/1
    Accounting Requests: 0
        Start: 0
        Stop: 0
        Interim: 0
        On: 0
        Off: 0
    Duplicate Requests: 0
    Malformed Requests: 0
    Bad Authenticators: 0
    Unknown Types: 0
    Dropped Packets: 0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf aaa radius snoop-segment statistics | 4253

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview

show unified-edge tdf aaa statistics

IN THIS SECTION

- Syntax | 27426
- Description | 27426
- Options | 27426
- Required Privilege Level | 27426
- Output Fields | 27427
- Sample Output | 27432

- Release Information | **27436**

Syntax

```
show unified-edge tdf aaa statistics
<brief | detail>
<fpc-slot fpc-slot>
<gateway gateway>
<pic-slot pic-slot>
```

Description

Display global statistics for accounting requests and responses for one or more TDF gateways. If a TDF gateway is not specified, then statistics for all TDF gateways are displayed.

Options

none	(Same as brief) Display statistics for all TDF gateways.
brief | detail	(Optional) Display the specified level of output. The `brief` option displays the consolidated statistics for all TDF gateways, and the `detail` option displays the statistics for each Services PIC on the configured TDF gateways.
fpc-slot *fpc-slot*	(Optional) Display statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway*	(Optional) Display statistics for the specified TDF gateway.
pic-slot *pic-slot*	(Optional) Display statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

view

Output Fields

Table 1723 on page 27427 lists the output fields for the `show unified-edge tdf aaa statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1723: show unified-edge tdf aaa statistics Output Fields

Field Name	Field Description	Level of Output
`Gateway Name`	Name of the TDF gateway. If the statistics for all TDF gateways are displayed, then `All` is displayed.	All levels
`FPC/PIC`	FPC and PIC slot numbers for which the statistics are displayed.	`detail`
`Total Messages`	Total number of all RADIUS requests and responses for the following categories: • `Received` • `Sent` • `Snooped`—Snooped by the MX Series router. • `Forwarded In`—Forwarded into the interface. • `Forwarded Out`—Forwarded out of the interface.	All levels `Forwarded In` and `Forwarded Out` are displayed only with the `detail` option.
`Access Requests`	Number of access requests for the following category: • `Sent`—Sent to the RADIUS server from the FPC slot and PIC slot.	All levels
`Access Responses`	Number of access responses for the following category: • `Received`—Received from the RADIUS server for the FPC slot and PIC slot. The following information is displayed: • `Accept`—Number of access accepts sent by the RADIUS server. • `Reject`—Number of access requests rejected by the RADIUS server.	All levels

Table 1723: show unified-edge tdf aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting Requests`	Number of accounting requests for the following categories: • `Received` • `Sent` • `Snooped`—Snooped by the MX Series router. • `Forwarded In`—Forwarded into the interface. • `Forwarded Out`—Forwarded out of the interface. The following information is displayed for each category: • `Start`—Number of Accounting Start requests. • `Interim`—Number of Accounting Interim-Update requests. • `Stop`—Number of Accounting Stop requests. • `On`—Number of Accounting On requests. • `Off`—Number of Accounting Off requests.	All levels `Forwarded In` and `Forwarded Out` are displayed only with the `detail` option.

Table 1723: show unified-edge tdf aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting Responses`	Number of accounting responses for the following categories: • `Received` • `Sent` • `Snooped`—Snooped by the MX Series router. • `Forwarded In`—Forwarded into the interface. • `Forwarded Out`—Forwarded out of the interface. The following information is displayed for each category: • `Start`—Number of Accounting Start responses. • `Interim`—Number of Accounting Interim-Update responses. • `Stop`—Number of Accounting Stop responses. • `On`—Number of Accounting On responses. • `Off`—Number of Accounting Off responses.	All levels `detail`—Number of responses that are forwarded into the interface and forwarded out of the interface is displayed only with the `detail` option.
`Change of Auth Requests`	Number of change of authorization (CoA) requests for the following categories: • `Received`—Received from the RADIUS server. • `Forwarded In`—Forwarded into the interface. • `Forwarded Out`—Forwarded out of the interface.	All levels `Forwarded In` and `Forwarded Out` are displayed only with the `detail` option.

Table 1723: show unified-edge tdf aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Change of Auth Responses`	Number of CoA responses for the following category: • `Sent`—Sent to the RADIUS server from the FPC slot and PIC slot. • `Forwarded In`—Forwarded into the interface. • `Forwarded Out`—Forwarded out of the interface. The following information is displayed: • `Ack`—Number of CoA acknowledgements sent to the RADIUS server. • `Nack`—Number of CoA negative acknowledgements sent to the RADIUS server.	All levels `Forwarded In` and `Forwarded Out` are displayed only with the `detail` option.
`Disconnect Message Requests`	Number of Disconnect Message requests for the following categories: • `Received`—Received from the RADIUS server. • `Forwarded In`—Forwarded into the interface. • `Forwarded Out`—Forwarded out of the interface.	All levels `Forwarded In` and `Forwarded Out` are displayed only with the `detail` option.
`Disconnect Message Responses`	Number of Disconnect Message responses for the following categories: • `Sent`—Sent to the RADIUS server. • `Forwarded In`—Forwarded into the interface. • `Forwarded Out`—Forwarded out of the interface. The following information is displayed: • `Ack`—Number of Disconnect Message acknowledgements sent to the RADIUS server. • `Nack`—Number of Disconnect Message negative acknowledgements sent to the RADIUS server.	All levels `Forwarded In` and `Forwarded Out` are displayed only with the `detail` option.

Table 1723: show unified-edge tdf aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Duplicates`	Number of duplicate requests received from RADIUS clients.	All levels
`Request Processing errors`	Number of errors that occurred during the processing of accounting requests.	All levels
`Response Processing errors`	Number of errors that occurred during the processing of access and accounting response packets from the RADIUS server.	All levels
`Request Transmit errors`	Number of errors that occurred during the transmission of access and accounting requests.	All levels
`Response Transmit errors`	Number of errors that occurred during the transmission of access and accounting responses to the RADIUS server.	All levels
`Request Enqueue Errors`	Number of errors that occurred while trying to place an access or accounting request packet in the queue.	All levels
`Response Enqueue Errors`	Number of errors that occurred while trying to place an access or accounting response packet in the queue.	All levels
`Request Timeouts`	Number of access and accounting requests to the RADIUS server that timed out.	All levels
`Request Retransmissions`	Number of access and accounting requests that were retransmitted to the RADIUS server because they did not receive a response.	All levels
`Dropped Requests`	Number of accounting requests that were dropped.	All levels
`Dropped Responses`	Number of access or accounting responses from the RADIUS server that were dropped.	All levels

Table 1723: show unified-edge tdf aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Missing TDF Domain`	Number of accounting requests from the GGSN, PGW, or BNG for which the TDF domain corresponding to the subscriber was not available.	All levels
`Missing PCEF profile`	Number of accounting requests from the GGSN, PGW, or BNG for which the PCEF profile corresponding to the subscriber was not available.	All levels
`Server Initiated Request Processing Errors`	Number of processing errors of CoA and Disconnect Message requests from the RADIUS server.	All levels
`Dropped Server Initiated Requests`	Number of CoA and Disconnect Message requests from the RADIUS server that were dropped.	All levels
`Duplicate Server Initiated Requests`	Number of duplicate requests received from RADIUS servers.	All levels
`Cached Reply Sent`	Number of RADIUS cached responses sent for RADIUS accounting request messages from the GGSN, PGW, or BNG. RADIUS replies are stored in the MX Series router response cache.	All levels

Sample Output

show unified-edge tdf aaa statistics brief

```
user@host> show unified-edge tdf aaa statistics brief


Gateway Name: TDF
Messages                     Received        Sent        Snooped
---------------------------------------------------------------------
Total Messages               15              15          0
```

```
Access Requests              0               7           0
Access Responses             7               0           0
    Accept                   7               0           0
    Reject                   0               0           0
Accounting Requests          8               0           0
    Start                    8               0           0
    Interim                  0               0           0
    Stop                     0               0           0
    On                       0               0           0
    Off                      0               0           0
Accounting Responses         0               8           0
    Start                    0               8           0
    Interim                  0               0           0
    Stop                     0               0           0
    On                       0               0           0
    Off                      0               0           0
Change of Auth Requests      0               0           0
Change of Auth Responses     0               0           0
    Ack                      0               0           0
    Nak                      0               0           0
Disconnect Message Requests  0               0           0
Disconnect Message Responses 0               0           0
    Ack                      0               0           0
    Nak                      0               0           0

Duplicates:                                   0
Request Processing Errors:                    0
Response Processing Errors:                   0
Request Transmit Errors :                     0
Response Transmit Errors:                     0
Request Enqueue Errors:                       0
Response Enqueue Errors:                      0
Request Timeouts:                             0
Request Retransmissions:                      0
Missing TDF Domain:                           0
Missing PCEF profile:                         0
Dropped Requests:                             0
Dropped Responses:                            0
Server Initiated Request Processing Errors:   0
Dropped Server Initiated Requests:            0
Duplicate Server Initiated Requests:          0
```

```
Cached Reply Sent:                          0
```

show unified-edge tdf aaa statistics detail

```
user@host> show unified-edge tdf aaa statistics detail


Gateway Name: TDF
FPC/PIC: 2/0
Messages                  Received        Sent        Forwarded In     Forwarded Out   Snooped
-----------------------------------------------------------------------------------------------

Total Messages               2               2           0                   0            0
Access Requests              0               0           0                   0            0
Access Responses             0               0           0                   0            0
    Accept                   0               0           0                   0            0
    Reject                   0               0           0                   0            0
Accounting Requests          2               0           0                   0            0
    Start                    2               0           0                   0            0
    Interim                  0               0           0                   0            0
    Stop                     0               0           0                   0            0
    On                       0               0           0                   0            0
    Off                      0               0           0                   0            0
Accounting Responses         0               2           0                   0            0
    Start                    0               2           0                   0            0
    Interim                  0               0           0                   0            0
    Stop                     0               0           0                   0            0
    On                       0               0           0                   0            0
    Off                      0               0           0                   0            0
Change of Auth Requests      0               0           0                   0            0
Change of Auth Responses     0               0           0                   0            0
    Ack                      0               0           0                   0            0
    Nak                      0               0           0                   0            0
Disconnect Message Requests  0               0           0                   0            0
Disconnect Message Responses 0               0           0                   0            0
    Ack                      0               0           0                   0            0
    Nak                      0               0           0                   0            0

Duplicates:                                  0
Request Processing Errors:                   0
```

```
Response Processing Errors:                    0
Request Transmit Errors :                      0
Response Transmit Errors:                      0
Request Enqueue Errors:                        0
Response Enqueue Errors:                       0
Request Timeouts:                              0
Request Retransmissions:                       0
Missing TDF Domain:                            0
Missing PCEF profile:                          0
Dropped Requests:                              0
Dropped Responses:                             0
Server Initiated Request Processing Errors:    0
Dropped Server Initiated Requests:             0
Duplicate Server Initiated Requests:           0
Cached Reply Sent:                             0

Gateway Name: TDF
FPC/PIC: 2/1
Messages                     Received        Sent        Forwarded In      Forwarded Out   Snooped
-----------------------------------------------------------------------------------------------------
-
Total Messages                  0              0             0                    0            0
Access Requests                 0              0             0                    0            0
Access Responses                0              0             0                    0            0
    Accept                      0              0             0                    0            0
    Reject                      0              0             0                    0            0
Accounting Requests             0              0             0                    0            0
    Start                       0              0             0                    0            0
    Interim                     0              0             0                    0            0
    Stop                        0              0             0                    0            0
    On                          0              0             0                    0            0
    Off                         0              0             0                    0            0
Accounting Responses            0              0             0                    0            0
    Start                       0              0             0                    0            0
    Interim                     0              0             0                    0            0
    Stop                        0              0             0                    0            0
    On                          0              0             0                    0            0
    Off                         0              0             0                    0            0
Change of Auth Requests         0              0             0                    0            0
Change of Auth Responses        0              0             0                    0            0
    Ack                         0              0             0                    0            0
    Nak                         0              0             0                    0            0
Disconnect Message Requests     0              0             0                    0            0
```

```
Disconnect Message Responses   0              0             0                    0            0
    Ack                        0              0             0                    0            0
    Nak                        0              0             0                    0            0

Duplicates:                                 0
Request Processing Errors:                  0
Response Processing Errors:                 0
Request Transmit Errors :                   0
Response Transmit Errors:                   0
Request Enqueue Errors:                     0
Response Enqueue Errors:                    0
Request Timeouts:                           0
Request Retransmissions:                    0
Missing TDF Domain:                         0
Missing PCEF profile:                       0
Number of Dropped Requests:                           0
Dropped Responses:                          0
Server Initiated Request Processing Errors: 0
Dropped Server Initiated Requests:          0
Duplicate Server Initiated Requests:        0
Cached Reply Sent:                          0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- clear unified-edge tdf aaa statistics | **4255**
- show unified-edge tdf aaa radius client statistics | **27396**
- show unified-edge tdf aaa radius server statistics | **27410**
- *Understanding How a RADIUS Server Controls Policy and Charging Control Rules*
- *IP-Based Subscriber Setup Overview*
- *Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview*

show unified-edge tdf address-assignment pool

IN THIS SECTION

- Syntax | 27437
- Description | 27437
- Options | 27437
- Required Privilege Level | 27438
- Output Fields | 27438
- Sample Output | 27440
- Release Information | 27442

Syntax

```
show unified-edge tdf address-assignment pool
<brief | detail | summary>
<fpc-slot fpc-slot>
<gateway gateway-name>
<name pool-name>
<pic-slot pic-slot>
<routing-instance routing-instance>
```

Description

Display information about the address pools for one or more TDF gateways. If a TDF gateway is not specified, then information for all TDF gateways is displayed.

Options

none	(Same as brief) Display address information about the address pools in brief for all TDF gateways.
brief | detail | summary	(Optional) Display the specified level of output.

fpc-slot ***fpc-slot***	(Optional) Display address pool information for the session PIC in the specified FPC slot.
gateway ***gateway-name***	(Optional) Display address pool information for the specified TDF gateway.
name ***pool-name***	(Optional) Display information for the specified address pool.
pic-slot ***pic-slot***	(Optional) Display address pool information for the session PIC in the specified PIC slot.
routing-instance ***routing-instance***	(Optional) Display the address pool information for the specified routing instance.

Required Privilege Level

view

Output Fields

Table 1724 on page 27438 lists the output fields for the `show unified-edge tdf address-assignment pool` command. Output fields are listed in the approximate order in which they appear.

Table 1724: show unified-edge tdf address-assignment pool Output Fields

Field Name	Field Description	Level of Output
`Pool` or `Name`	Name of the address pool.	All levels
`FPC/PIC`	FPC and PIC slots of the session PIC for which the address pool information is displayed.	`detail`
`Total addresses`	Total number of addresses available in the address pool.	`brief` `detail`
`Total`	Total number of addresses available in the address pool.	`summary`

Table 1724: show unified-edge tdf address-assignment pool Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Addresses in use`	Number of addresses that have been allocated.	`brief` `detail`
`In-use`	Number of addresses that have been allocated.	`summary`
`Addresses skipped`	Number of addresses that are excluded from allocation.	`brief` `detail`
`Address usage (percent)`	Percentage of the total addresses used.	`brief` `detail`
`Util (%)`	Percentage of the total addresses used.	`summary`
`Addresses in aging period`	Number of addresses that are currently being released and that cannot be allocated.	`brief` `detail`
`Routing Instance`	Name of the routing instance to which the address pool belongs.	All levels
`Gateway`	TDF gateway to which the session PIC belongs.	`detail`
`Pool Maintenance Mode`	Service mode of the address pool; for example, Operational or Maintenance.	`detail`
`Address chunks`	Number of chunks of IP addresses in the address pool (for the session PIC) that are currently being assigned	`detail`

Table 1724: show unified-edge tdf address-assignment pool Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Total address chunk size`	Total number of addresses in the address chunk (for the session PIC).	`detail`
`Total allocation failures`	Total number of addresses that were not allocated.	`detail`

Sample Output

show unified-edge tdf address-assignment pool brief

```
user@host> show unified-edge tdf address-assignment pool brief


Pool: pool1
  Total addresses:           16777215
  Addresses in use:          1600
  Addresses skipped:         416
  Address usage (percent):   0
  Addresses in aging period: 1600
  Routing instance:          default


Pool: pool2
  Total addresses:           256
  Addresses in use:          254
  Addresses skipped:         2
  Address usage (percent):   99
  Addresses in aging period: 0
  Routing instance:          default

[...output truncated...]
```

show unified-edge tdf address-assignment pool detail

```
user@host> show unified-edge tdf address-assignment pool detail

Pool: pool1 (FPC/PIC: 4/0)
  Pool Maintenance Mode:     Operational
  Total addresses:           16777215
  Addresses in use:          822
  Addresses skipped:         208
  Address usage (percent):   0
  Addresses in aging period: 822
  Routing instance:          default
  Gateway:                   TDF
  Address chunks:            26
  Total address chunk size:  26416
  Total allocation failures: 0

Pool: pool1 (FPC/PIC: 4/1)
  Pool Maintenance Mode:     Operational
  Total addresses:           16777215
  Addresses in use:          778
  Addresses skipped:         208
  Address usage (percent):   0
  Addresses in aging period: 778
  Routing instance:          default
  Gateway:                   TDF
  Address chunks:            26
  Total address chunk size:  26416
  Total allocation failures: 0

Pool: pool2 (FPC/PIC: 4/0)
  Pool Maintenance Mode:     Operational
  Total addresses:           256
  Addresses in use:          0
  Addresses skipped:         0
  Address usage (percent):   0
  Addresses in aging period: 0
  Routing instance:          default
  Gateway:                   TDF
  Address chunks:            0
  Total address chunk size:  0
  Total allocation failures: 0
```

```
[...output truncated...]
```

show unified-edge tdf address-assignment pool summary

```
user@host> show unified-edge tdf address-assignment pool summary
                                                    Util
Name                        Total       In-use      (%)  Routing instance
pool1                       16777215    1600        0    default
pool2                       256         254         99   default
pool3                       256         47          18   default
v4_pool                     16777216    0           0    default
v4_pool1                    16777215    0           0    default
v6_pool                     16777215    0           0    default
v6_pool1                    16777215    0           0    default
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf address-assignment pool | 4257

show unified-edge tdf address-assignment service-mode

IN THIS SECTION

- Syntax | 27443
- Description | 27443
- Options | 27443
- Required Privilege Level | 27443
- Output Fields | 27443

Sample Output | 27444
Release Information | 27445

Syntax

```
show unified-edge tdf address-assignment service-mode
<brief | detail>
<pool pool-name>
<routing-instance routing-instance-name>
```

Description

Display service mode information about address pools.

Options

none	(Same as brief) Display service mode information in brief.
brief | detail	(Optional) Display the specified level of output.
pool *pool-name*	(Optional) Display service mode information for the specified address pool.
routing-instance *routing-instance-name*	(Optional) Display service mode information about the address pools that are part of the specified routing instance.

Required Privilege Level

view

Output Fields

Table 1725 on page 27444 lists the output fields for the `show unified-edge tdf address-assignment service-mode` command. Output fields are listed in the approximate order in which they appear.

Table 1725: show unified-edge tdf address-assignment service-mode Output Fields

Field Name	Field Description	Level of Output
`Maintenance Mode`	Phases applicable when the address pool is in maintenance mode. • `MM - Active Phase`—All the attributes of the address pool can be modified. • `MM - In/Out Phase`—Only the non-maintenance mode attributes of the address pool can be modified.	None specified
`Pool Name`	Name of the address pool.	All levels
`Routing Instance`	Routing instance to which the address pool belongs.	All levels
`Service Mode`	Service mode for the address pool: • `Operational`—Address pool is in operational mode. • `Maintenance`—Address pool is in maintenance mode. • `Maintenance - Active Phase`—All the attributes of the address pool can be modified. • `Maintenance - In/Out Phase`—Only the non-maintenance mode attributes of the address pool can be modified.	All levels

Sample Output

show unified-edge tdf address-assignment service-mode brief

```
user@host> show unified-edge tdf address-assignment service-mode brief
Maintenance Mode
    MM Active Phase - System is ready to accept configuration changes for all
                      attributes of this object and its sub-hierarchies.
    MM In/Out Phase - System is ready to accept configuration changes only for
                      non-maintenance mode attributes of this object and
                      its sub-hierarchies.

Routing-Instance                Pool Name                       Service Mode
```

```
default                         my_pool                         Operational
default                         v6_pool                         Operational
```

show unified-edge tdf address-assignment service-mode detail

```
user@host> show unified-edge tdf address-assignment service-mode detail
Routing Instance: default
Pool Name       : my_pool
Service Mode    : Operational

Routing Instance: default
Pool Name       : v6_pool
Service Mode    : Operational
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Address Pools for Source-IP Filtering of IP-Based Subscribers

show unified-edge tdf address-assignment statistics

IN THIS SECTION

- Syntax | 27446
- Description | 27446
- Options | 27446
- Required Privilege Level | 27446
- Output Fields | 27446

Sample Output | 27447
Release Information | 27448

Syntax

```
show unified-edge tdf address-assignment statistics
<brief | detail>
<fpc-slot fpc-slot>
<gateway gateway-name>
<pic-slot pic-slot>
```

Description

Display address assignment statistics for one or more TDF gateways. If a TDF gateway is not specified, then the consolidated statistics for all TDF gateways are displayed.

Options

none	(Same as brief) Display address assignment statistics in brief for all TDF gateways.
brief | detail	(Optional) Display the specified level of output.
fpc-slot *fpc-slot*	(Optional) Display statistics for the session PIC in the specified FPC slot.
gateway *gateway-name*	(Optional) Display consolidated statistics for the specified TDF gateway.
pic-slot *pic-slot*	(Optional) Display statistics for the session PIC in the specified PIC slot.

Required Privilege Level

view

Output Fields

Table 1726 on page 27447 lists the output fields for the `show unified-edge tdf address-assignment statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1726: show unified-edge tdf address-assignment statistics Output Fields

Field Name	Field Description	Level of Output
FPC/PIC	FPC and PIC slots for which the statistics are displayed.	detail
Gateway	Name of the TDF gateway.	detail
Total address allocations	Total number of addresses allocated.	All levels
Total allocation failures	Total number of address allocations that failed.	All levels
Total address releases	Total number of addresses that were released.	All levels

Sample Output

show unified-edge tdf address-assignment statistics

```
user@host> show unified-edge tdf address-assignment statistics


Address assignment statistics
  Total address allocations: 1101
  Total allocation failures: 0
  Total address releases:    800
```

show unified-edge tdf address-assignment statistics detail

```
user@host> show unified-edge tdf address-assignment statistics detail


Address assignment statistics   (FPC/PIC: 4/0)
  Gateway:                   TDF
  Total address allocations: 416
  Total allocation failures: 0
```

```
  Total address releases:    416

Address assignment statistics   (FPC/PIC: 4/1)
  Gateway:                   TDF
  Total address allocations: 685
  Total allocation failures: 0
  Total address releases:    384
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf address-assignment statistics | 4259

show unified-edge tdf call-admission-control statistics

IN THIS SECTION

- Syntax | 27448
- Description | 27449
- Options | 27449
- Required Privilege Level | 27449
- Output Fields | 27449
- Sample Output | 27451
- Release Information | 27452

Syntax

```
show unified-edge tdf call-admission-control statistics
<detail>
```

```
<fpc-slot fpc-slot>
<gateway gateway-name>
<pic-slot pic-slot>
```

Description

Display call admission control (CAC) statistics for one or more TDF gateways. If a TDF gateway is not specified, then statistics for all TDF gateways are displayed.

> **NOTE**: CAC statistics are not stored on the Routing Engine. When this command is executed, the Routing Engine fetches the statistics from the active session PICs and displays the consolidated statistics for one or more TDF gateways.

Options

none Display CAC statistics for all TDF gateways.

detail (Optional) Display detailed CAC statistics for the specified FPC and PIC slot numbers.

> **NOTE**: The `detail` option is valid only when you specify an FPC and PIC slot number configured on the TDF gateway.

fpc-slot *fpc-slot* (Optional) Display statistics for the session PIC in the specified FPC slot.

pic-slot *pic-slot* (Optional) Display statistics for the session PIC in the specified PIC slot.

gateway *gateway-name* (Optional) Display CAC statistics for the specified TDF gateway.

Required Privilege Level

view

Output Fields

Table 1727 on page 27450 lists the output fields for the `show unified-edge tdf call-admission-control statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1727: show unified-edge tdf call-admission-control statistics Output Fields

Field Name	Field Description	Level of Output
`Gateway: TDF`	Output is displayed for TDF gateways.	`detail` none
`GW CAC Statistics`	Statistical details are displayed at the TDF gateway level.	`detail` none
`Memory High Rejects`	Number of subscribers or PDP contexts that were rejected because the memory load or utilization (at the session PIC level) was high.	`detail` none
`Memory High Redirects`	Number of subscribers or PDP contexts that were redirected because the memory load or utilization (at the session PIC level) was high.	`detail` none
`CPU High Rejects`	Number of subscribers or PDP contexts that were rejected because the CPU load or utilization (at the session PIC level) was high.	`detail` none
`CPU High Redirects`	Number of subscribers or PDP contexts that were redirected because the CPU load or utilization (at the session PIC level) was high.	`detail` none
`Session Reservation Rejects`	Number of sessions that were rejected for reservation of TDF subscribers on a particular TDF gateway or domain.	`detail` none
`Session Reservation Redirects`	Number of sessions that were redirected to a different TDF gateway or domain for reservation of TDF subscribers.	`detail` none
`Gateway Subscriber Count`	Total number of subscribers that are connected to the TDF gateway.	`detail` none

Table 1727: show unified-edge tdf call-admission-control statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`TDF DOMAIN CAC Statistics`	Statistical details are displayed at the TDF domain level.	`detail` none
`Session Reservation Rejects`	Number of sessions that were rejected for reservation of TDF subscribers on a particular TDF gateway or domain.	`detail` none
`Session Reservation Redirects`	Number of sessions that were redirected to a different TDF gateway or domain for reservation of TDF subscribers.	`detail` none

Sample Output

show unified-edge tdf call-admission-control statistics

```
user@host> show unified-edge tdf call-admission-control statistics
Gateway: TDF

GW CAC Statistics:
  Memory High Rejects                        :   0
  Memory High Redirects                      :   0
  CPU High Rejects                           :   0
  CPU High Redirects                         :   0
  Session Reservation Rejects                :   0
  Session Reservation Redirects              :   0
  Gateway Subscriber Count                   :   1

Domain CAC Statistics:
  Session Reservation Rejects                :   0
  Session Reservation Redirects              :   0
```

show unified-edge tdf call-admission-control statistics fpc-slot pic-slot detail

```
user@host> show unified-edge tdf call-admission-control statistics fpc-slot 3 pic-slot 1 detail

Gateway: TDF

GW CAC Statistics:
  Memory High Rejects                        :   0
  Memory High Redirects                      :   0
  CPU High Rejects                           :   0
  CPU High Redirects                         :   0
  Session Reservation Rejects                :   0
  Session Reservation Redirects              :   0
  Gateway Subscriber Count                   :   1

Domain CAC Statistics:
  Session Reservation Rejects                :   0
  Session Reservation Redirects              :   0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf call-admission-control statistics | 4261

show unified-edge tdf call-rate statistics

IN THIS SECTION

- Syntax | 27453
- Description | 27453
- Options | 27453
- Required Privilege Level | 27453

- Output Fields | 27453
- Sample Output | 27455
- Release Information | 27455

Syntax

```
show unified-edge tdf call-rate statistics (domain domain-name | gateway gateway-name)
```

Description

Display call-rate statistics for the specified TDF domain or specified TDF gateway.

Options

domain *domain-name* Display call-rate statistics for the specified TDF domain.

gateway *gateway-name* Display call-rate statistics for the specified TDF gateway.

Required Privilege Level

view

Output Fields

Table 1728 on page 27453 lists the output fields for the `show unified-edge tdf call-rate statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1728: show unified-edge tdf call-rate statistics Output Fields

Field Name	Field Description
`Gateway`	Name of the TDF gateway.

Table 1728: show unified-edge tdf call-rate statistics Output Fields *(Continued)*

Field Name	Field Description
`TDF domain name`	Name of the TDF domain. This is displayed only when the `domain` option is used.
`Record` *n*	Displays statistics for the most recent *n* number of intervals.
`Number of Activations`	Number of successful subscriber logins for this record.
`Number of Deactivations`	Number of subscriber logouts for this record.
`Activations processing time (in ms)`	Average subscriber login activation processing time.
`Subscriber session duration (in mins)`	Average subscriber session duration.
`Statistics collection time`	Time at which the statistics were collected.
`Control Plane Standard Deviation`	Standard deviations for the following: • `Number of Activations`—Number of subscriber logins • `Number of Deactivations`—Number of subscriber logouts • `Activations processing time`—Length of time of subscriber login • `Subscriber session duration`—Length of time of subscriber logout

Sample Output

show unified-edge tdf call-rate statistics gateway

```
user@host> show unified-edge tdf call-rate statistics gateway TDF
Gateway: TDF
Record 1 (Call-rate statistics for the past 1 min):
Control Plane:
        Number of Activations:                  0
        Number of Deactivations:                0
        Activations processing time (in ms):    0
        Subscriber session duration (in mins):  0

Statistics collection time: 2014-03-04 11:45:44 UTC (00:03:06 ago)

Record 2 (Call-rate statistics for the past 2 min):
Control Plane:
        Number of Activations:                  0
        Number of Deactivations:                0
        Activations processing time (in ms):    0
        Subscriber session duration (in mins):  0

Control Plane Standard Deviation:
        Number of Activations:                  0
        Number of Deactivations:                0
        Activations processing time:            0
        Subscriber session duration:            0

Statistics collection time: 2014-03-04 11:44:44 UTC (00:04:06 ago)
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Call-Rate Statistics Collection

show unified-edge tdf diameter network-element statistics

IN THIS SECTION

- Syntax | 27456
- Description | 27456
- Options | 27456
- Required Privilege Level | 27457
- Output Fields | 27457
- Sample Output | 27458
- Release Information | 27459

Syntax

```
show unified-edge tdf diameter network-element statistics
<brief | detail>
<fpc-slot fpc-slot>
<gateway gateway-name>
<network-element-name network-element-name>
<pic-slot pic-slot>
```

Description

Display statistics for network elements for one or more TDF gateways. If a network element is not specified, then statistics for all network elements are displayed. If a TDF gateway is not specified, then statistics for all TDF gateways are displayed.

Options

none	Display statistics for network elements for all TDF gateways.
brief | detail	(Optional) Display the specified level of output. The `brief` output is displayed by default.
fpc-slot *fpc-slot*	(Optional) Display statistics for the specified Flexible PIC Concentrator (FPC).

gateway *gateway-name*	(Optional) Display statistics for the specified TDF gateway.
network-element-name *network-element-name*	(Optional) Display statistics for the specified network element.
pic-slot *pic-slot*	(Optional) Display the statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

view

Output Fields

Table 1729 on page 27457 lists the output fields for the `show unified-edge tdf diameter network-element statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1729: show unified-edge tdf diameter network-element statistics Output Fields

Field Name	Field Description	Level of Output
`Name`	Name of the network element.	All levels
`FPC/PIC`	FPC and PIC slot numbers through which the network element was reached.	`detail`
`Packets Received`	Number of incoming packets.	All levels
`Packets Transmitted`	Number of outgoing packets.	All levels
`Request Timeouts`	Number of request timeouts.	All levels
`Credit Control Request Transmitted`	Number of outgoing Credit-Control-Request messages.	All levels
`Credit Control Answer Received`	Number of incoming Credit-Control-Answer messages.	All levels

Sample Output

show unified-edge tdf diameter network-element statistics

```
user@host> show unified-edge tdf diameter network-element statistics

  Name:   pcrf-dne
     Packets Received :                       0
     Packets Transmitted :                    0
     Request Timeouts :                       0
     Credit Control Request Transmitted :     0
     Credit Control Answer Received :         0

  Name:   ocs-dne
     Packets Received :                       3
     Packets Transmitted :                    4
     Request Timeouts :                       1
     Credit Control Request Transmitted :     4
     Credit Control Answer Received :         3
```

show unified-edge tdf diameter network-element statistics detail

```
user@host> show unified-edge tdf diameter network-element statistics detail

  Name :                                  pcrf-dne
  FPC/PIC :                               0/0
  Packets Received :                      0
  Packets Transmitted :                   0
  Request Timeouts :                      0
  Credit Control Request Transmitted :    0
  Credit Control Answer Received :        0

  FPC/PIC :                               0/1
  Packets Received :                      0
  Packets Transmitted :                   0
  Request Timeouts :                      0
  Credit Control Request Transmitted :    0
  Credit Control Answer Received :        0

  Name :                                  ocs-dne
```

```
FPC/PIC :                                0/0
Packets Received :                       0
Packets Transmitted :                    0
Request Timeouts :                       0
Credit Control Request Transmitted :     0
Credit Control Answer Received :         0

FPC/PIC :                                0/1
Packets Received :                       3
Packets Transmitted :                    4
Request Timeouts :                       1
Credit Control Request Transmitted :     4
Credit Control Answer Received :         3
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf diameter network-element statistics | 4262

show unified-edge tdf diameter network-element status | 27459

show unified-edge tdf diameter network-element status

IN THIS SECTION

- Syntax | 27460
- Description | 27460
- Options | 27460
- Required Privilege Level | 27460
- Output Fields | 27460
- Sample Output | 27461
- Release Information | 27462

Syntax

```
show unified-edge tdf diameter network-element status
<fpc-slot fpc-slot>
<gateway gateway-name>
<network-element-name network-element-name>
<pic-slot pic-slot>
```

Description

Display the status for one or more Diameter network elements. If a network element is not specified, then status for all network elements is displayed. If a TDF gateway is not specified, then status for all TDF gateways is displayed.

Options

none	Display status for all network elements for all TDF gateways.
fpc-slot *fpc-slot*	(Optional) Display the status for the specified Flexible PIC Concentrator (FPC).
gateway *gateway-name*	(Optional) Display the status for the specified TDF gateway.
network-element-name *network-element-name*	(Optional) Display the status for the specified network element.
pic-slot *pic-slot*	(Optional) Display the status for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

view

Output Fields

Table 1730 on page 27461 lists the output fields for the `show unified-edge tdf diameter network-element status` command. Output fields are listed in the approximate order in which they appear.

Table 1730: show unified-edge tdf diameter network-element status Output Fields

Field Name	Field Description
DNE	Name of the network element.
PEER	Name of the peer.
FPC/PIC	FPC and PIC slot numbers through which the network element was reached.
PEER STATE	Current state of the peer. Possible states are: Closed, Closing, I-Open, R-Open, Wait-Conn-Ack, Wait-Conn-Ack/Elect, Wait-I-CEA, and Wait-Returns.
WATCHDOG STATE	Peer watchdog status. • closed—Connection between Diameter peers is terminated. • initial—Connection between Diameter peers is being initialized. • okay—Connection between Diameter peers is established and active.

Sample Output

show unified-edge tdf diameter network-element status

```
user@host> show unified-edge tdf diameter network-element status
DNE : pcrf-dne
    PEER : pcrf
       FPC/PIC          PEER STATE            WATCHDOG STATE
         0/0            Closed                initial
         0/1            Closed                initial
DNE : ocs-dne
    PEER : ocs
       FPC/PIC          PEER STATE            WATCHDOG STATE
         0/0            I-Open                okay
         0/1            I-Open                okay
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf diameter network-element statistics | 27456

show unified-edge tdf diameter pcc-gx statistics

IN THIS SECTION

- Syntax | 27462
- Description | 27462
- Options | 27463
- Required Privilege Level | 27463
- Output Fields | 27463
- Sample Output | 27467
- Release Information | 27469

Syntax

```
show unified-edge tdf diameter pcc-gx statistics
<brief | detail>
<fpc-slot fpc-slot>
<gateway gateway-name>
<pic-slot pic-slot>
```

Description

Display statistics for the Gx application for one or more TDF gateways. If a TDF gateway is not specified, then statistics for all TDF gateways are displayed.

Options

none	Display statistics for the Gx application for all TDF gateways.
brief | detail	(Optional) Display the specified level of output. The `brief` output is displayed by default.
fpc-slot *fpc-slot*	(Optional) Display statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway-name*	(Optional) Display statistics for the specified TDF gateway.
pic-slot *pic-slot*	(Optional) Display statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

view

Output Fields

Table 1731 on page 27463 lists the output fields for the `show unified-edge tdf diameter pcc-gx statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1731: show unified-edge tdf diameter pcc-gx statistics Output Fields

Field Name	Field Description	Level of Output
`Gateway`	Name of the TDF gateway.	All levels
`FPC/PIC`	FPC and PIC slots for which the statistics are displayed.	`detail`
`Total Sessions Established`	Total number of active sessions.	All levels
`Total Sessions Terminated`	Total number of terminated sessions.	All levels
`Internal Errors`	Number of internal errors.	`detail`

Table 1731: show unified-edge tdf diameter pcc-gx statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Total	• Requests—Total number of request messages. • Answers—Total number of answer messages.	none brief
Credit Control Initial	• Requests—Number of initial transfer type Credit-Control-Request (CCR) messages. • Answers—Number of initial transfer type Credit-Control-Answer (CCA) messages.	none brief
Credit Control Update	• Requests—Number of update transfer type CCR messages. • Answers—Number of update transfer type CCA messages.	none brief
Credit Control Terminate	• Requests—Number of terminate transfer type CCR messages. • Answers—Number of terminate transfer type CCA messages.	none brief
Re-Auth	• Requests—Number of Re-Auth-Request (RAR) messages. • Answers—Number of Re-Auth-Answer (RAA) messages.	none brief
Dropped	• Requests—Number of dropped request messages. • Answers—Number of dropped answer messages.	none brief
Requests Transmitted	• Initial—Number of initial transfer type CCR messages sent. • Update—Number of update transfer type CCR messages sent. • Terminate—Number of terminate transfer type CCR messages sent. • Total—Number of CCR messages sent.	detail

Table 1731: show unified-edge tdf diameter pcc-gx statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Request Timeouts`	• `Initial`—Number of initial transfer type CCR messages that timed out. • `Update`—Number of update transfer type CCR messages that timed out. • `Terminate`—Number of terminate transfer type CCR messages that timed out. • `Total`—Number of CCR messages that timed out.	`detail`
`Request Tx Timeouts`	• `Initial`—Number of initial transfer type CCR messages sent that timed out. • `Update`—Number of update transfer type CCR messages sent that timed out. • `Terminate`—Number of terminate transfer type CCR messages sent that timed out. • `Total`—Number of CCR messages sent that timed out.	`detail`
`Request Discarded`	• `Initial`—Number of initial transfer type CCR messages sent that were discarded. • `Update`—Number of update transfer type CCR messages sent that were discarded. • `Terminate`—Number of terminate transfer type CCR messages sent that were discarded. • `Total`—Number of CCR messages sent that were discarded.	`detail`

Table 1731: show unified-edge tdf diameter pcc-gx statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Answers Received`	• `Initial`—Number of initial transfer type CCA messages received. • `Update`—Number of update transfer type CCA messages received. • `Terminate`—Number of terminate transfer type CCA messages received. • `Total`—Number of CCA messages received.	`detail`
`Answers Dropped`	• `Initial`—Number of initial transfer type CCA messages dropped. • `Update`—Number of update transfer type CCA messages dropped. • `Terminate`—Number of terminate transfer type CCA messages dropped. • `Total`—Number of CCA messages dropped.	`detail`
`Answers Parse Errors`	• `Initial`—Number of initial transfer type CCA messages with parse errors. • `Update`—Number of update transfer type CCA messages with parse errors. • `Terminate`—Number of terminate transfer type CCA messages with parse errors. • `Total`—Number of CCA messages with parse errors.	`detail`
`Answers with Invalid AVP(s)`	• `Initial`—Number of initial transfer type CCA messages with invalid AVPs. • `Update`—Number of update transfer type CCA messages with invalid AVPs. • `Terminate`—Number of terminate transfer type CCA messages with invalid AVPs. • `Total`—Number of CCA messages with invalid AVPs.	`detail`

Table 1731: show unified-edge tdf diameter pcc-gx statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Requests Received	Number of RAR messages received.	detail
Requests Dropped	Number of RAR messages dropped.	detail
Requests Parse Errors	Number of RAR messages with parse errors.	detail
Requests with Invalid AVP(s)	Number of RAR messages with invalid AVPs.	detail
Answers Transmitted	Number of RAA messages sent.	detail

Sample Output

show unified-edge tdf diameter pcc-gx statistics

```
user@host> show unified-edge tdf diameter pcc-gx statistics
  Gateway: TDF
    Total Sessions Established:             0
    Total Sessions Terminated:                      0
                                Requests            Answers
    -------------------------------------------------------
    Total                       0                   0
    Credit Control Initial      0                   0
    Credit Control Update       0                   0
    Credit Control Terminate    0                   0
    Re-Auth                     0                   0
    Dropped                     0                   0
```

show unified-edge tdf diameter pcc-gx statistics detail

```
user@host> show unified-edge tdf diameter pcc-gx statistics detail
  Gateway: TDF
  FPC/PIC: 0/0
    Total Sessions Established:             0
    Total Sessions Terminated:  0
    Internal Errors:            0

    Credit Control                  Initial     Update      Terminate       Total
    -------------------------------------------------------------------------------
    Requests Transmitted            0           0           0               0
    Request Timeouts                0           0           0               0
    Request Tx Timeouts             0           0           0               0
    Request Discarded               0           0           0               0
    Answers Received                0           0           0               0
    Answers Dropped                 0           0           0               0
    Answers Parse Errors            0           0           0               0
    Answers with Invalid AVP(s)     0           0           0               0

    Server Requests                 Re-Auth
    --------------------------------------
    Requests Received               0
    Requests Dropped                0
    Requests Parse Errors           0
    Requests with Invalid AVP(s)    0
    Answers Transmitted             0

  Gateway: TDF
  FPC/PIC: 0/1
    Total Sessions:             0
    Total Sessions Terminated:  0
    Internal Errors:            0

    Credit Control                  Initial     Update      Terminate       Total
    -------------------------------------------------------------------------------
    Requests Transmitted            0           0           0               0
    Request Timeouts                0           0           0               0
    Request Tx Timeouts             0           0           0               0
    Request Discarded               0           0           0               0
    Answers Received                0           0           0               0
    Answers Dropped                 0           0           0               0
```

```
    Answers Parse Errors            0           0           0               0
    Answers with Invalid AVP(s)     0           0           0               0

    Server Requests                 Re-Auth
    --------------------------------------
    Requests Received               0
    Requests Dropped                0
    Requests Parse Errors           0
    Requests with Invalid AVP(s)    0
    Answers Transmitted             0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf diameter pcc-gx statistics | 4264

show unified-edge tdf diameter peer statistics

IN THIS SECTION

- Syntax | 27470
- Description | 27470
- Options | 27470
- Required Privilege Level | 27470
- Output Fields | 27470
- Sample Output | 27473
- Release Information | 27476

Syntax

```
show unified-edge tdf diameter peer statistics
<brief | detail>
<fpc-slot fpc-slot>
<gateway gateway-name>
<peer-name peer-name>
<pic-slot pic-slot>
```

Description

Display statistics for Diameter peers for one or more TDF gateways. If a peer is not specified, then statistics for all Diameter peers are displayed. If a TDF gateway is not specified, then statistics for all TDF gateways are displayed.

Options

none	(Same as brief) Display statistics for Diameter peers for all TDF gateways in brief.
brief | detail	(Optional) Display the specified level of output. The `brief` output is displayed by default.
fpc-slot *fpc-slot*	(Optional) Display statistics for the specified Flexible PIC Concentrator (FPC).
gateway *gateway-name*	(Optional) Display statistics for the specified TDF gateway.
peer-name *peer-name*	(Optional) Display statistics for the specified peer.
pic-slot *pic-slot*	(Optional) Display statistics for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

view

Output Fields

Table 1732 on page 27471 lists the output fields for the `show unified-edge tdf diameter peer statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1732: show unified-edge tdf diameter peer statistics Output Fields

Field Name	Field Description	Level of Output
`Peer`	Name of the peer.	All levels
`FPC/PIC`	FPC and PIC slot numbers through which the peer was reached.	`detail`
`Request Timeouts`	Number of request timeouts.	All levels
`Request Retransmissions`	Number of request retransmissions.	All levels
`Connect Failures`	Number of connection failures.	`detail`
`Duplicate Requests`	Number of duplicate requests.	`detail`
`Malformed Messages`	Number of malformed requests.	`detail`
`Dropped Responses`	Number of dropped responses.	`detail`
`Dropped Requests`	Number of dropped requests.	`detail`
`Last Disconnect Cause`	Number of last disconnect cause messages.	`detail`
`Transport Failures`	Number of transport failures.	`detail`
`Unknown Messages`	Number of unknown type errors.	`detail`
`High Watermark Hits`	Number of times the high watermark is reached.	`detail`
`Low Watermark Hits`	Number of times the low watermark is reached.	`detail`
`Device Watchdog Failured`	Number of device watchdog failures.	`detail`

Table 1732: show unified-edge tdf diameter peer statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Capabilities Exchange Failures`	Number of capabilities exchange failures.	`detail`
`Total Messages`	Total number of messages transmitted and received.	All levels
`Credit Control Requests`	Number of Credit-Control-Request messages transmitted and received.	All levels
`Credit Control Answers`	Number of Credit-Control-Answer messages transmitted and received.	All levels
`Re-Auth Requests`	Number of Re-Auth-Request messages transmitted and received.	All levels
`Re-Auth Answers`	Number of Re-Auth-Answer messages transmitted and received.	All levels
`Abort Session Requests`	Number of Abort-Session-Request messages transmitted and received.	All levels
`Abort Session Answers`	Number of Abort-Session-Answer messages transmitted and received.	All levels
`Capability Exchange Requests`	Number of Capabilities-Exchange-Request messages transmitted and received.	All levels
`Capability Exchange Answers`	Number of Capabilities-Exchange-Answer messages transmitted and received.	All levels
`Device Watchdog Requests`	Number of Device-Watchdog-Request messages transmitted and received.	All levels
`Device Watchdog Answers`	Number of Device-Watchdog-Answer messages transmitted and received.	All levels

Table 1732: show unified-edge tdf diameter peer statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Disconnect Peer Requests	Number of Disconnect-Peer-Request messages transmitted and received.	All levels
Disconnect Peer Answers	Number of Disconnect-Peer-Answer messages transmitted and received.	All levels
Permanent Failures	Number of permanent failure result codes transmitted and received.	detail
Protocol Errors	Number of protocol error result codes transmitted and received.	detail
Transient Failures	Number of transient failure result codes transmitted and received.	detail

Sample Output

show unified-edge tdf diameter peer statistics

```
user@host> show unified-edge tdf diameter peer statistics
Peer: ocs
    Request Timeouts:                   1
    Request Retransmissions:            0
    Messages                            Transmitted       Received
    ---------------------------------------------------------------
    Total Messages                      6                 5
    Credit Control Requests             4                 0
    Credit Control Answers              0                 3
    Re-Auth Requests                    0                 0
    Re-Auth Answers                     0                 0
    Abort Session Requests              0                 0
    Abort Session Answers               0                 0
    Capability Exchange Requests        2                 0
    Capability Exchange Answers         0                 2
```

```
    Device Watchdog Requests            0                 0
    Device Watchdog Answers             0                 0
    Disconnect Peer Requests            0                 0
    Disconnect Peer Answers             0                 0
```

show unified-edge tdf diameter peer statistics detail

```
user@host> show unified-edge tdf diameter peer statistics detail
Peer: ocs
    FPC/PIC: 0/0
    Request Timeouts:                   0
    Request Retransmissions:            0
    Connect Failures:                   0
    Duplicate Requests:                 0
    Malformed Messages:                 0
    Dropped Responses:                  0
    Dropped Requests:                   0
    Last Disconnect Cause:              0
    Transport Failures:                 0
    Unknown Messages:                   0
    High Watermark Hits:                0
    Low Watermark Hits:                 0
    Device Watchdog Failured:           0
    Capabilities Exchange Failures:     0

    Messages                            Transmitted       Received
    ---------------------------------------------------------------
    Total Messages                      1                 1
    Credit Control Requests             0                 0
    Credit Control Answers              0                 0
    Re-Auth Requests                    0                 0
    Re-Auth Answers                     0                 0
    Abort Session Requests              0                 0
    Abort Session Answers               0                 0
    Capability Exchange Requests        1                 0
    Capability Exchange Answers         0                 1
    Device Watchdog Requests            0                 0
    Device Watchdog Answers             0                 0
    Disconnect Peer Requests            0                 0
    Disconnect Peer Answers             0                 0
```

```
Result-Code                        Transmitted       Received
-------------------------------------------------------------
Permanent Failures                 0                 0
Protocol Errors                    0                 0
Transient Failures                 0                 0

FPC/PIC: 0/1
Request Timeouts:                  1
Request Retransmissions:           0
Connect Failures:                  0
Duplicate Requests:                0
Malformed Messages:                0
Dropped Responses:                 0
Dropped Requests:                  0
Last Disconnect Cause:             0
Transport Failures:                0
Unknown Messages:                  0
High Watermark Hits:               0
Low Watermark Hits:                0
Device Watchdog Failured:          0
Capabilities Exchange Failures:    0

Messages                           Transmitted       Received
-------------------------------------------------------------
Total Messages                     5                 4
Credit Control Requests            4                 0
Credit Control Answers             0                 3
Re-Auth Requests                   0                 0
Re-Auth Answers                    0                 0
Abort Session Requests             0                 0
Abort Session Answers              0                 0
Capability Exchange Requests       1                 0
Capability Exchange Answers        0                 1
Device Watchdog Requests           0                 0
Device Watchdog Answers            0                 0
Disconnect Peer Requests           0                 0
Disconnect Peer Answers            0                 0

Result-Code                        Transmitted       Received
-------------------------------------------------------------
Permanent Failures                 0                 0
```

```
    Protocol Errors                    0                 0
    Transient Failures                 0                 0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf diameter peer statistics | 4266

show unified-edge tdf diameter peer status | 27476

show unified-edge tdf diameter peer status

IN THIS SECTION

- Syntax | 27476
- Description | 27477
- Options | 27477
- Required Privilege Level | 27477
- Output Fields | 27477
- Sample Output | 27479
- Release Information | 27480

Syntax

```
show unified-edge tdf diameter peer status
<brief | detail>
<fpc-slot fpc-slot>
<gateway gateway-name>
<peer-name peer-name>
<pic-slot pic-slot>
```

Description

Display the status for one or more Diameter peers. If a peer is not specified, then status for all Diameter peers is displayed. If a TDF gateway is not specified, then status for all TDF gateways is displayed.

Options

none	(Same as brief) Display the status for Diameter peers for all TDF gateways in brief.
brief | detail	(Optional) Display the specified level of output. The `brief` output is displayed by default.
fpc-slot *fpc-slot*	(Optional) Display the status for the specified Flexible PIC Concentrator (FPC).
gateway *gateway-name*	(Optional) Display the status for the specified TDF gateway.
peer-name *peer-name*	(Optional) Display the status for the specified peer.
pic-slot *pic-slot*	(Optional) Display the status for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

Required Privilege Level

view

Output Fields

Table 1733 on page 27477 lists the output fields for the `show unified-edge tdf diameter peer status` command. Output fields are listed in the approximate order in which they appear.

Table 1733: show unified-edge tdf diameter peer status Output Fields

Field Name	Field Description	Level of Output
`Name`	Name of the peer. For the `brief` output, the name is truncated if it exceeds 11 characters.	All levels
`FPC/PIC`	FPC and PIC slot numbers through which the peer was reached.	All levels

Table 1733: show unified-edge tdf diameter peer status Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Address`	IP address of the Diameter peer.	`brief` none
`Port`	Port number of the Diameter peer.	`brief` none
`State`	Current state of the Diameter peer. Possible states are: `Closed`, `Closing`, `I-Open`, `R-Open`, `Wait-Conn-Ack`, `Wait-Conn-Ack/Elect`, `Wait-I-CEA`, and `Wait-Returns`. For the `brief` and none output, the state is truncated if it exceeds 11 characters.	All levels
`Duration`	Duration for which the Diameter peer has been in the current state in Coordinated Universal Time (UTC) format (HH:MM:SS).	none `brief`
`State Duration`	Duration for which the Diameter peer has been in the current state in Coordinated Universal Time (UTC) format (HH:MM:SS).	`detail`
`Watchdog`	Peer watchdog status. • `closed`—Connection between Diameter peers is terminated. • `initial`—Connection between Diameter peers is being initialized. • `okay`—Connection between Diameter peers is established and active.	none `brief`
`Watchdog State`	Peer watchdog status. • `closed`—Connection between Diameter peers is terminated. • `initial`—Connection between Diameter peers is being initialized. • `okay`—Connection between Diameter peers is established and active.	`detail`
`Origin Host`	Diameter Origin-Host.	`detail`

Table 1733: show unified-edge tdf diameter peer status Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Origin Realm	Diameter Origin-Realm.	detail
Peer Address	IP address of the Diameter peer.	detail
Peer port	Port number of the Diameter peer.	detail
Source Address	Local source IP address used to connect to the peer.	detail
Source Port	Local source port number used to connect to the peer.	detail

Sample Output

show unified-edge tdf diameter peer status

```
user@host> show unified-edge tdf diameter peer status
Name         FPC/PIC  Address          Port    State        Duration  Watchdog
p_jpkt1      4/0      192.0.2.2        3868    Closed       00:00:00  initial
p_jpkt1      4/1      192.0.2.2        3868    Closed       00:00:00  initial
p_jpkt1      5/0      192.0.2.2        3868    Wait-Conn-A  00:00:00  initial
abcabcabcab  4/0      192.0.2.2        3868    Closed       00:00:00  initial
abcabcabcab  4/1      192.0.2.2        3868    Closed       00:00:00  initial
abcabcabcab  5/0      192.0.2.2        3868    Wait-Conn-A  00:00:00  initial
```

show unified-edge tdf diameter peer status detail

```
user@host> show unified-edge tdf diameter peer status detail
Diameter Peer Status
  Name : ocs
     FPC/PIC               :          0/0
     State                 :          I-Open
     State Duration        :          00:00:00
     Watchdog State        :          okay
```

```
    Origin Host          :          host5
    Origin Realm         :          example.com
    Peer Address         :          198.51.100.2
    Peer port            :          3868
    Source Address       :          203.0.113.1
    Source Port          :          30965
Name : ocs
    FPC/PIC              :          0/1
    State                :          I-Open
    State Duration       :          00:00:00
    Watchdog State       :          okay
    Origin Host          :          host5
    Origin Realm         :          example.com
    Peer Address         :          198.51.100.2
    Peer port            :          3868
    Source Address       :          203.0.113.1
    Source Port          :          30709
Name : pcrf
    FPC/PIC              :          0/0
    State                :          Closed
    State Duration       :          00:00:00
    Watchdog State       :          initial
    Peer Address         :          192.168.1.2
    Peer port            :          3868
    Source Address       :          203.0.113.1
    Source Port          :          0
Name : pcrf
    FPC/PIC              :          0/1
    State                :          Closed
    State Duration       :          00:00:00
    Watchdog State       :          initial
    Peer Address         :          192.168.1.2
    Peer port            :          3868
    Source Address       :          203.0.113.1
    Source Port          :          0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf diameter peer statistics | 27469

show unified-edge tdf domain service-mode

IN THIS SECTION

- Syntax | 27481
- Description | 27481
- Options | 27481
- Required Privilege Level | 27482
- Output Fields | 27482
- Sample Output | 27483
- Release Information | 27483

Syntax

```
show unified-edge tdf domain service-mode
<domain-name tdf-domain-name>
<brief | detail>
<gateway gateway>
```

Description

Display service mode information for a TDF domain for one or more TDF gateways. If a TDF domain is not specified, then the information for all domains for one or more TDF gateways is displayed.

Options

none (Same as brief) Display the TDF domain service mode information in brief for all TDF gateways.

brief | detail	(Optional) Display the specified level of output.
domain-name ***tdf-domain-name***	(Optional) Display service mode information for the specified TDF domain.
gateway ***gateway***	(Optional) Display service mode information for the specified TDF gateway.

Required Privilege Level

view

Output Fields

Table 1734 on page 27482 lists the output fields for the `show unified-edge tdf domain service-mode` command. Output fields are listed in the approximate order in which they appear.

Table 1734: show unified-edge tdf domain service-mode Output Fields

Field Name	Field Description	Level of Output
`Maintenance Mode`	Phases applicable when the address pool is in maintenance mode. • `MM - Active Phase`—All the attributes of the address pool can be modified. • `MM - In/Out Phase`—Only the non-maintenance mode attributes of the address pool can be modified.	None specified
`Gateway Name`	Name of the TDF gateway.	None specified
`Gateway`	Name of the TDF gateway.	`detail`
`TDF domain Name`	Name of the TDF domain.	All levels
`Service Mode`	Service mode for the TDF gateway: • `Operational`—Gateway is in operational mode. • `Maintenance`—Gateway is in maintenance mode.	All levels

Sample Output

show unified-edge tdf domain service-mode brief

```
user@host> show unified-edge tdf domain service-mode brief
Maintenance Mode
    MM Active Phase - System is ready to accept configuration changes for all
                      attributes of this object and its sub-hierarchies.
    MM In/Out Phase - System is ready to accept configuration changes only for
                      non-maintenance mode attributes of this object and
                      its sub-hierarchies.

TDF domain Name               Gateway Name                     Service Mode

jnpr-sunnyvale                  TDF                              Operational
jnpr-toxin                      TDF                              Operational
zoo                             TDF1                             Maintenance - Active Phase
```

show unified-edge tdf domain service-mode detail

```
user@host> show unified-edge tdf domain service-mode detail
Gateway: TDF
TDF domain Name        : jnpr-sunnyvale
Service Mode    : Operational

TDF domain Name        : jnpr-toxin
Service Mode    : Operational
Gateway: TDF1

TDF domain Name        : zoo
Service Mode    : Maintenance - Active Phase
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf service-mode | 27493

show unified-edge tdf domain statistics

IN THIS SECTION

- Syntax | 27484
- Description | 27484
- Options | 27484
- Required Privilege Level | 27485
- Output Fields | 27485
- Sample Output | 27488
- Release Information | 27490

Syntax

```
show unified-edge tdf domain statistics
<domain-name domain-name>
<gateway gateway>
```

Description

Display statistics for one or more domains in a TDF gateway. If a TDF gateway is not specified, then statistics for all TDF gateways are displayed.

Options

none	Display statistics for all domains for all TDF gateways.
domain-name *domain-name*	(Optional) Display the statistics for the specified TDF domain.

The output of the `show unified-edge tdf domain statistics` command is the same as the output of the `show unified-edge tdf statistics` command with the `tdf-domain` option.

gateway *gateway*

(Optional) Display the statistics for the specified TDF gateway.

Required Privilege Level

view

Output Fields

Table 1735 on page 27485 lists the output fields for the `show unified-edge tdf domain statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1735: show unified-edge tdf domain statistics Output Fields

Field Name	Field Description
`Gateway`	Name of the TDF gateway.
`Control Plane Statistics`	
`Subscriber attach attempts`	Number of attempted session establishments and number of successful session establishments (`Success`).
`TDF Time of day initiated update attempts`	Number of attempted activations of rules based on time of day settings and number of successful activations (`Success`).
`TDF initiated subscriber detach attempts`	Number of attempted subscriber session detachments initiated by the TDF.
`PCRF initiated subscriber detach attempts`	Number of attempted subscriber session detachments initiated by the PCRF.

Table 1735: show unified-edge tdf domain statistics Output Fields *(Continued)*

Field Name	Field Description
`Peer initiated subscriber detach attempts`	Number of attempted subscriber session detachments initiated by the peer.
`Subscriber attach failures by cause`	Number of session establishments that failed: • `System failure` • `No resources` • `Policy denied` • `Service denied` • `Others`
`Rejects due to early CAC`	Number of subscriber sessions rejected due to early call admission control (CAC) for the TDF gateway.
`Policy statistics`	
`Subscriber session activation attempts`	Number of subscriber session activations attempted. In addition, the number of successful subscriber session establishments (`Success`) is displayed.
`TDF initiated modification attempts`	Number of session modifications initiated by TDF gateway. In addition, the number session modifications that were successful (`Success`) is displayed.
`PCRF initiated modification attempts`	Number of session modifications initiated by the policy and charging rules function (PCRF). In addition, the number of modifications that were successful (`Success`) is displayed.
`TDF initiated session deactivations`	Number of subscriber session deactivations initiated by the TDF gateway.

Table 1735: show unified-edge tdf domain statistics Output Fields *(Continued)*

Field Name	Field Description
PCRF initiated session deactivations	Number of subscriber session deactivations initiated by the PCRF.
Modification event reason	The number of Gx modifications for each event reason: • Application Start • Application Stop
Failure Statistics	• Session terminations due to unreachable PCRF—Number of sessions terminated because the PCRF was unreachable. • Session terminations due to PCRF restart—Number of sessions terminated because the PCRF was restarted. • Rule Validation Failures—Number of sessions terminated because the validation of rules failed.
PCC Rule Statistics	• Dynamic rule activations—Number of dynamic rule activations and deactivations (Deactivations). • Static rule activations—Number of static rule activations and deactivations (Deactivations). • Dynamic rule modifications—Number of dynamic rule modifications.
PCC Rule Failure Statistics	• Rule update failure—Number of rules that cannot be updated.
ePCC/ADC Rule Statistics	• Dynamic rule activations—Number of dynamic rule activations and deactivations (Deactivations). • Static rule activations—Number of static rule activations and deactivations (Deactivations). • Dynamic rule modifications—Number of dynamic rule modifications.

Table 1735: show unified-edge tdf domain statistics Output Fields *(Continued)*

Field Name	Field Description
ePCC/ADC Rule Failure Statistics	• Rule update failure—Number of rules that cannot be updated.

Sample Output

show unified-edge tdf domain statistics gateway

```
user@host> show unified-edge tdf domain statistics gateway tdf
Gateway: TDF
Control Plane Statistics:
    Subscriber attach attempts:                 0     Success: 0
    TDF Time of day initiated update attempts:  0     Success: 0
    TDF initiated subscriber detach attempts:   0
    PCRF initiated subscriber detach attempts:  0
    Peer initiated subscriber detach attempts:  0
    Subscriber attach failures by cause:
       System failure:      0
       No resources:        0
       Service denied:      0
       Policy denied:       0
       Others:              0
    Rejects due to early CAC: 0
Policy Statistics:
    Subscriber session activation attempts:     0     Success: 0
    TDF initiated modification attempts:        0     Success: 0
    PCRF initiated modification attempts:       0     Success: 0
    TDF initiated session deactivations:        0
    PCRF initiated session deactivations:       0
  Modification Event Reason:
       Application Start:   0
       Application Stop:    0
  Failure Statistics:
       Session terminations due to unreachable PCRF:  0
       Session terminations due to PCRF restart:      0
       Rule validation failures:                      0
  PCC Rule Statistics:
```

```
        Dynamic rule activations:     0      Deactivations:   0
        Static rules activations:     0      Deactivations:   0
        Dynamic rule modifications:   0
  PCC Rule Failure Statistics:
        Rule update failure:                  0
  ePCC/ADC Rule Statistics:
        Dynamic rule activations:     0      Deactivations:   0
        Static rules activations:     0      Deactivations:   0
        Dynamic rule modifications:   0
  ePCC/ADC Rule Failure Statistics:
        Rule update failure:                  0
```

show unified-edge tdf domain statistics domain-name

```
user@host> show unified-edge tdf domain statistics domain-name domain1
domain-name domain1
Gateway: TDF
Control Plane Statistics:
    Subscriber attach attempts:                 0     Success: 0
    TDF Time of day initiated update attempts:  0     Success: 0
    TDF initiated subscriber detach attempts:   0
    PCRF initiated subscriber detach attempts:  0
    Peer initiated subscriber detach attempts:  0
    Subscriber attach failures by cause:
       System failure:      0
       No resources:        0
       Service denied:      0
       Policy denied:       0
       Others:              0
Policy Statistics:
    Subscriber session activation attempts:     0     Success: 0
    TDF initiated modification attempts:        0     Success: 0
    PCRF initiated modification attempts:       0     Success: 0
    TDF initiated session deactivations:        0
    PCRF initiated session deactivations:       0
  Modification Event Reason:
       Application Start:   0
       Application Stop:    0
  Failure Statistics:
       Session terminations due to unreachable PCRF:  0
```

```
        Session terminations due to PCRF restart:      0
        Rule validation failures:                      0
    PCC Rule Statistics:
        Dynamic rule activations:     0      Deactivations:   0
        Static rules activations:     0      Deactivations:   0
        Dynamic rule modifications:   0
    PCC Rule Failure Statistics:
        Rule update failure:                  0
    ePCC/ADC Rule Statistics:
        Dynamic rule activations:     0      Deactivations:   0
        Static rules activations:     0      Deactivations:   0
        Dynamic rule modifications:   0
    ePCC/ADC Rule Failure Statistics:
        Rule update failure:                  0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf statistics | 4268

show unified-edge tdf resource-manager clients

IN THIS SECTION

- Syntax | 27491
- Description | 27491
- Options | 27491
- Required Privilege Level | 27491
- Output Fields | 27491
- Sample Output | 27492
- Release Information | 27493

Syntax

```
show unified-edge tdf resource-manager clients
<gateway gateway>
```

Description

Display information about the resource management clients (the session Dense Port Concentrators [DPCs] and interface DPCs and Modular Port Concentrators [MPCs]) on one or more TDF gateways. If a TDF gateway is not specified, then information for all TDF gateways is displayed.

Options

none	Display information for all TDF gateways.
gateway *gateway-name*	(Optional) Display information for the specified TDF gateway.

Required Privilege Level

view

Output Fields

Table 1736 on page 27491 lists the output fields for the `show unified-edge gateways tdf resource-manager clients` command. Output fields are listed in the approximate order in which they appear.

Table 1736: show unified-edge tdf resource-manager clients Output Fields

Field Name	Field Description
`Client`	Name of the resource manager client slot identified by the FPC and PIC slot numbers; for example, pfe-1/2/0 or ms-/7/0/0.
`State`	Resource manager client state. `In-Service` means that the client can handle session creation requests.

Table 1736: show unified-edge tdf resource-manager clients Output Fields *(Continued)*

Field Name	Field Description
`Role`	Role of the resource manager client slot: • `Primary`—The resource manager client is a primary member. • `Secondary`—The resource manager client is a secondary or backup member.
`Client type`	Type of resource manager client: • `Session PIC`—Session PIC client used for the mobile control plane in the TDF gateway. • `Service PIC`—services PIC used for anchoring services-related subscriber sessions in the TDF gateway.
`Gateway`	Name of the TDF gateway to which the resource manager client belongs.

Sample Output

show unified-edge tdf resource-manager clients

```
user@host> show unified-edge tdf resource-manager clients
Client      State          Redundancy role Client type Gateway
ms-2/0/0    In-Service     Primary         Service-PIC TDF
ms-2/1/0    In-Service     Secondary       Service-PIC TDF
ms-3/0/0    In-Service     Primary         Service-PIC TDF
ms-3/1/0    In-Service     Primary         Service-PIC TDF
ms-5/0/0    In-Service     Primary         Session-PIC TDF
ms-5/1/0    In-Service     Secondary       Session-PIC TDF
```

show unified-edge tdf resource-manager clients gateway

```
user@host> show unified-edge tdf resource-manager clients gateway TDF
Client      State          Redundancy role Client type Gateway
ms-3/0/0    In-Service     Secondary       Session-PIC TDF
ms-3/1/0    In-Service     Primary         Session-PIC TDF
```

```
ms-3/2/0    In-Service     Secondary       Service-PIC TDF
ms-3/3/0    In-Service     Primary         Service-PIC TDF
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- show unified-edge tdf subscribers | 27513
- show unified-edge tdf system interfaces | 27532

show unified-edge tdf service-mode

IN THIS SECTION

- Syntax | 27493
- Description | 27494
- Options | 27494
- Required Privilege Level | 27494
- Output Fields | 27494
- Sample Output | 27495
- Release Information | 27496

Syntax

```
show unified-edge tdf service-mode
<brief | detail>
<domain-name tdf-domain-name>
<gateway gateway-name>
```

Description

Display service mode information for one or more TDF gateways. If a TDF gateway is not specified, then service mode information for all the TDF gateways is displayed.

Options

none	(Same as brief) Display service mode information in brief for all TDF gateways.
brief | detail	(Optional) Display the specified level of output.
tdf-domain ***domain-name***	(Optional) Display service mode information for the specified TDF domain. The output of the `show unified-edge tdf service-mode` command with the `tdf-domain` option is the same as the output of the `show unified-edge tdf domain service-mode` command.
gateway ***gateway-name***	(Optional) Display service mode information for the specified TDF gateway.

Required Privilege Level

view

Output Fields

Table 1737 on page 27494 lists the output fields for the `show unified-edge tdf service-mode` command. Output fields are listed in the approximate order in which they appear.

Table 1737: show unified-edge tdf service-mode Output Fields

Field Name	Field Description	Level of Output
`Maintenance Mode`	Phases applicable when the TDF domain is in maintenance mode. • `MM - Active Phase`—All the attributes of the address pool can be modified. • `MM - In/Out Phase`—Only the non-maintenance mode attributes of the address pool can be modified.	none
`Gateway Name`	Name of the TDF gateway.	none

Table 1737: show unified-edge tdf service-mode Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Service Mode`	Service mode for the TDF gateway: • `Operational`—Gateway is in operational mode. • `Maintenance`—Gateway is in maintenance mode.	All levels

Sample Output

show unified-edge tdf service-mode brief

```
user@host> show unified-edge tdf service-mode brief
Maintenance Mode
    MM Active Phase - System is ready to accept configuration changes for all
                      attributes of this object and its sub-hierarchies.
    MM In/Out Phase - System is ready to accept configuration changes only for
                      non-maintenance mode attributes of this object and
                      its sub-hierarchies.

Gateway Name                    Service Mode

TDF                             Operational
TDF2                            Operational
```

show unified-edge tdf service-mode detail

```
user@host> show unified-edge tdf service-mode detail
Service Mode Status
Gateway Name    : PGW
Service Mode    : Operational
Service Mode Status
Gateway Name    : PGW2
Service Mode    : Operational
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf domain service-mode | 27481

show unified-edge tdf statistics

IN THIS SECTION

- Syntax | 27496
- Description | 27496
- Options | 27497
- Required Privilege Level | 27497
- Output Fields | 27497
- Sample Output | 27504
- Release Information | 27506

Syntax

```
show unified-edge tdf statistics
<domain domain-name>
<gateway gateway>
```

Description

Display statistics for one or more TDF gateways. If a TDF gateway is not specified, then statistics for all TDF gateways are displayed.

Options

none Display statistics for all TDF gateways.

domain *domain-name* (Optional) Display statistics for the specified TDF domain.

The output of the `show unified-edge tdf statistics` command with the `domain` *`domain-name`* option is the same as the output of the `show unified-edge tdf domain statistics` command.

gateway *gateway* (Optional) Display statistics for the specified TDF gateway.

Required Privilege Level

view

Output Fields

Table 1738 on page 27497 lists the output fields for the `show unified-edge tdf statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1738: show unified-edge tdf statistics Output Fields

Field Name	Field Description
`Gateway`	Name of the TDF gateway.
`Control Plane Statistics`	
`Subscriber attach attempts`	Number of attempted session establishments and number of successful session establishments for IP-based subscribers (`Success`).
`Peer initiated subscriber update attempts`	Number of RADIUS client attempts to update the subscriber context of an IP-based subscriber.
`TDF Time of day initiated update attempts`	Number of attempted activations, deactivations, and revalidations of PCC rules and revalidations of the PCEF session for the time-of-day feature, and number of successful attempts (`Success`).

Table 1738: show unified-edge tdf statistics Output Fields *(Continued)*

Field Name	Field Description
`TDF initiated update attempts`	Number of TDF gateway attempts to update an IFL-based subscriber context as a result of access interfaces going up or down, or as a result of access interfaces being added to or deleted from the subscriber configuration.
`TDF initiated subscriber detach attempts`	Number of attempted subscriber session detachments initiated by the TDF gateway.
`Policy Server initiated subscriber detach attempts`	Number of attempted subscriber session detachments initiated by the policy server.
`Peer initiated subscriber detach attempts`	Number of attempted IP-based subscriber session detachments initiated by the RADIUS client. For IFL-based subscribers, 0 is displayed.
`Subscriber attach failures by cause`	Number of session establishments that failed: • `System failure` • `No resources` • `Service denied` • `Policy denied` • `Service PIC NACK` • `Others`
`Rejects due to early CAC`	Number of rejects on the TDF gateway caused by early CAC.
`Subscriber detach by cause`	Number of subscriber detachments for the following cause: • `service PIC NACK`

Table 1738: show unified-edge tdf statistics Output Fields *(Continued)*

Field Name	Field Description
`Policy statistics`	
`Subscriber session activation attempts`	Number of subscriber session activations attempted. In addition, the number of successful subscriber session establishments (`Success`) is displayed.
`TDF initiated modification attempts`	Number of session modifications initiated by TDF gateway. In addition, the number of session modifications that were successful (`Success`) is displayed.
`Policy Server initiated modification attempts`	Number of session modifications initiated by the policy server. In addition, the number of modifications that were successful (`Success`) is displayed.
`TDF initiated session deactivations`	Number of subscriber session deactivations initiated by the TDF gateway.
`Policy Server initiated session deactivations`	Number of subscriber session deactivations initiated by the policy server.
`Modification event reason`	Number of Gx modifications for each event reason: • `Application Start` • `Application Stop` • `Revalidation`—PCEF re-requested PCC rules from the PCRF.

Table 1738: show unified-edge tdf statistics Output Fields *(Continued)*

Field Name	Field Description
Failure Statistics	• Session terminations due to unreachable policy server—Number of sessions terminated because the policy server was unreachable. • Session terminations due to PCRF restart—Number of sessions terminated because the PCRF was restarted. • Rule Validation Failures—Number of sessions terminated because the validation of rules failed.
PCC Rule Statistics	• Dynamic rule activations—Number of dynamic rule activations and deactivations (Deactivations). • Static rules activations—Number of static rule activations and deactivations (Deactivations). • Dynamic rule modifications—Number of dynamic rule modifications.
PCC Rule Failure Statistics	• Rule update failure—Number of rules that cannot be updated.
ePCC/ADC Rule Statistics	• Dynamic rule activations—Number of dynamic rule activations and deactivations (Deactivations). • Static rules activations—Number of static rule activations and deactivations (Deactivations). • Dynamic rule modifications—Number of dynamic rule modifications.
ePCC/ADC Rule Failure Statistics	• Rule update failure—Number of rules that cannot be updated.
Usage Monitoring Statistics	
UMI AVP validation failures	Number of times that decoding fails for any of the grouped AVPs that belong to the Usage Monitoring Information, such as the Monitoring key, Monitoring Level, and Granted Service Unit AVPs.

Table 1738: show unified-edge tdf statistics Output Fields *(Continued)*

Field Name	Field Description
`Session Level`	The following information about usage monitoring at the session level is displayed: • `UM activations`—Number of session-level monitoring keys that the TDF gateway has activated. • `UM update quota attempts`—Number of times the PCRF has attempted to update the quota for a session-level monitoring key. The number of reports that the TDF gateway sent as a result of the update quota attempts is shown in `Stats report sent`. • `UM implicit deactivations`—Number of times that a session-level monitoring key has been implicitly deactivated by the TDF gateway. For example, this happens if a monitoring key does not receive additional quota after a threshold has been reached. • `UM explicit deactivations`—Number of session-level monitoring key deactivations that the TDF gateway has received from the PCRF. The number of reports that the TDF gateway sent as a result of the deactivations is shown in `Stats report sent`. • `Usage report request received`—Number of requests for session-level usage reports that the TDF gateway has received from the PCRF. The number of reports that the TDF gateway sent as a result of the requests is shown in `Stats report sent`. • `UM threshold hit`—Number of times that a threshold for a session-level monitoring key has been reached. The number of reports that the TDF gateway sent as a result of the threshold being reached is shown in `Stats report sent`.

Table 1738: show unified-edge tdf statistics Output Fields *(Continued)*

Field Name	Field Description
`Rule Level`	The following information about usage monitoring at the rule level is displayed: • `UM activations`—Number of rule-level monitoring keys that the TDF gateway has activated. • `UM update quota attempts`—Number of times the PCRF has attempted to update the quota for a rule-level monitoring key. The number of reports that the TDF gateway sent as a result of the update quota attempts is shown in `Stats report sent`. • `UM implicit deactivations`—Number of times that a rule-level monitoring key has been implicitly deactivated by the TDF gateway. For example, this happens if a monitoring key does not receive additional quota after a threshold has been reached. • `UM explicit deactivations`—Number of rule-level monitoring key deactivations that the TDF gateway has received from the PCRF. The number of reports that the TDF gateway sent as a result of the deactivations is shown in `Stats report sent`. • `Usage report request received`—Number of requests for rule-level usage reports that the TDF gateway has received from the PCRF. The number of reports that the TDF gateway sent as a result of the requests is shown in `Stats report sent`. • `UM threshold hit`—Number of times that a threshold for a rule-level monitoring key has been reached. The number of reports that the TDF gateway sent as a result of the threshold being reached is shown in `Stats report sent`. • `UM with no rule reference`—Number of rule-level monitoring keys received by the TDF gateway that had no rule referring to it. These keys are not activated.
`Service plane statistics`	

Table 1738: show unified-edge tdf statistics Output Fields *(Continued)*

Field Name	Field Description
`Subscriber detach attempts (NACK) by cause`	Number of service PIC messages to session PIC indicating that subscriber creation or modification failed for the following causes: • `Memory watermark high threshold hit` • `Memory watermark critical threshold hit` • `Memory alloc failure` • `Subscriber lookup failure` • `Others`
`Data plane statistics`	
`Subscriber Stats`	The following information about packets processed by the data plane for subscribers connected to the TDF domains in the TDF gateway is displayed: • `Uplink`—Statistics for traffic in the uplink direction from the TDF gateway to the PDN (Internet). • `Downlink`—Statistics for traffic in the downlink direction from the PDN (Internet) to the TDF gateway. • `Packets`—Number of packets forwarded in the uplink direction and in the downlink direction. • `Bytes`—Number of bytes forwarded in the uplink direction and in the downlink direction. • `Dropped Packets`—Number of packets dropped in the uplink direction and in the downlink direction. • `Dropped Bytes`—Number of bytes dropped in the uplink direction and in the downlink direction.

Table 1738: show unified-edge tdf statistics Output Fields *(Continued)*

Field Name	Field Description
`Non Subscriber Stats`	The following information about packets processed by the data plane for traffic that does not belong to subscribers connected to the TDF domains in the TDF gateway is displayed: • `Uplink`—Statistics for traffic in the uplink direction from the TDF gateway to the PDN (Internet). • `Downlink`—Statistics for traffic in the downlink direction from the PDN (Internet) to the TDF gateway. • `Packets`—Number of packets sent in the uplink direction and in the downlink direction. • `Bytes`—Number of bytes sent in the uplink direction and in the downlink direction. • `Dropped Packets`—Number of packets dropped in the uplink direction and in the downlink direction. • `Dropped Bytes`—Number of bytes dropped in the uplink direction and in the downlink direction.

Sample Output

show unified-edge tdf statistics

```
user@host> show unified-edge tdf statistics
Gateway: TDF
Control Plane Statistics:
    Subscriber attach attempts:                 0     Success: 0
    Peer initiated subscriber update attempts:  0     Success: 0
    TDF Time of day initiated update attempts:  0     Success: 0
    TDF initiated update attempts:              0     Success: 0
    TDF initiated subscriber detach attempts:   0
    Policy Server initiated subscriber detach attempts:  0
    Peer initiated subscriber detach attempts:  0
    Subscriber attach failures by cause:
       System failure:      0
       No resources:        0
       Service denied:      0
       Policy denied:       0
       Service PIC NACK:    0
       Others:              0
```

```
  Rejects due to early CAC:   0
  Subscriber detach by cause:
     service PIC NACK: 0
Policy Statistics:
  Subscriber session activation attempts:     0     Success: 0
  TDF initiated modification attempts:        0     Success: 0
  Policy Server initiated modification attempts:       0     Success: 0
  TDF initiated session deactivations:        0
  Policy Server initiated session deactivations:       0
Modification Event Reason:
     Application Start:                0
     Application Stop:                 0
     Revalidation:                     0
Failure Statistics:
     Session terminations due to unreachable policy server: 0
     Session terminations due to PCRF restart:              0
     Rule validation failures:                              0
PCC Rule Statistics:
     Dynamic rule activations:     0      Deactivations:   0
     Static rules activations:     0      Deactivations:   0
     Dynamic rule modifications:   0
PCC Rule Failure Statistics:
     Rule update failure:                  0
ePCC/ADC Rule Statistics:
     Dynamic rule activations:     0      Deactivations:   0
     Static rules activations:     0      Deactivations:   0
     Dynamic rule modifications:   0
ePCC/ADC Rule Failure Statistics:
     Rule update failure:                  0
Usage Monitoring Statistics:
     UMI AVP validation failures:   0
  Session Level:
        UM activations:                0
        UM update quota attempts:      0      Stats report sent: 0
        UM implicit deactivations:     0
        UM explicit deactivations:     0      Stats report sent: 0
        Usage report request received: 0      Stats report sent: 0
        UM threshold hit:              0      Stats report sent: 0
  Rule Level:
        UM activations:                0
        UM update quota attempts:      0      Stats report sent: 0
        UM implicit deactivations:     0
        UM explicit deactivations:     0      Stats report sent: 0
```

```
          Usage report request received: 0      Stats report sent: 0
          UM threshold hit:              0      Stats report sent: 0
          UM with no rule reference: 0

Service plane statistics:
   Subscriber detach attempts (NACK) by cause:
     Memory watermark high threshold hit:      0
     Memory watermark critical threshold hit:  0
     Memory alloc failure:                     0
     Subscriber lookup failure:                0
     Others:                                   0

Data plane statistics:

Subscriber Stats:
                      Uplink                Downlink
---------------------------------------------------------------
    Packets          :0                    :0
    Bytes            :0                    :0
    Dropped Packets  :0                    :0
    Dropped Bytes    :0                    :0

Non Subscriber Stats:
                      Uplink                Downlink
---------------------------------------------------------------
    Packets          :0                    :0
    Bytes            :0                    :0
    Dropped Packets  :0                    :0
    Dropped Bytes    :0                    :0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

clear unified-edge tdf statistics | 4268

show unified-edge tdf domain statistics | 27484

IP-Based and IFL-Based TDF Subscribers Overview

show unified-edge tdf status

IN THIS SECTION

- Syntax | 27507
- Description | 27507
- Options | 27507
- Required Privilege Level | 27508
- Output Fields | 27508
- Sample Output | 27511
- Release Information | 27512

Syntax

```
show unified-edge tdf status
<brief | detail | extensive>
<domain domain-name>
<fpc-slot fpc-slot>
<gateway gateway>
<pic-slot pic-slot>
<subscriber-state>
```

Description

Display status information, such as the number of subscribers, active sessions, and so on, for one or more TDF gateways. If a TDF gateway name is not specified, then the status information for all the TDF gateways is displayed.

Options

none	(Same as brief) Display the TDF gateway status information in brief for all TDF gateways.
brief | detail | extensive	(Optional) Display the specified level of output.

domain *domain-name* (Optional) Display the status information for the specified TDF domain.

fpc-slot *fpc-slot* (Optional) Display the status information for the specified FPC slot number.

gateway *gateway* (Optional) Display the status information for the specified TDF gateway name.

pic-slot *pic-slot* (Optional) Display the status information for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.

subscriber-state (Optional) Display the status of the subscribers.

Required Privilege Level

view

Output Fields

Table 1739 on page 27508 lists the output fields for the `show unified-edge tdf status` command. Output fields are listed in the approximate order in which they appear.

Table 1739: show unified-edge tdf status Output Fields

Field Name	Field Description	Level of Output
`Gateway`	Name of the TDF gateway.	All levels
`Established`	Number of established subscribers.	none with the `subscriber-state` option
`Deleting`	Number of subscribers that are being deleted.	none with the `subscriber-state` option

Table 1739: show unified-edge tdf status Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Control Plane`	The following is displayed for the control plane:
	• `Active Subscribers`—Number of subscribers that are active in each of the following categories:
	  • `IP Subscribers`
	  • `IFL Subscribers`
		none
		`brief`

`Service Plane`	The following is displayed for the service plane:
	• `Active Subscribers`—Number of subscribers that are actively using services in each of the following categories:
	  • `IP Subscribers`
	  • `IFL Subscribers`
	• `Active Sessions`—Number of active subscriber sessions.
		none
		`brief`

`CPU Load (%)`	Percentage of the CPU load.	All levels
`Memory Load (%)`	Percentage of the memory load.	All levels
`FPC SLOT`	FPC slot number of the interface for which the status information is displayed.	`detail` `extensive`
`PIC SLOT`	PIC slot number of the FPC for which the status information is displayed.	`detail` `extensive`

Table 1739: show unified-edge tdf status Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Role`	Role of the Packet Forwarding Engine, services PIC, or session PIC on the TDF gateway: • `Standalone` • `Primary`—Primary member. • `Secondary`—Secondary member.	`detail` `extensive`
`Type`	Indicates whether the PIC is a Packet Forwarding Engine, a session PIC, or a services PIC.	`detail` `extensive`

`Active Subscribers`	Number of logged-in subscribers on the TDF gateway in each of the following categories:
	• `IP Subscribers`
	• `IFL Subscribers`
		`brief`
		`detail`
		`extensive`

`Delete Pending Subscribers`	Number of pending subscribers that are being deleted on the TDF gateway in each of the following categories: • `IP Subscribers` • `IFL Subscribers`	`detail` `extensive`
`Active Sessions`	Number of logged-in sessions on the TDF gateway. **NOTE**: Active Sessions count may not match the output of the `show services session count` command. This is due to internal asynchronous message queues.	`detail` `extensive`

Sample Output

show unified-edge tdf status brief

```
user@host> show unified-edge tdf status brief
Gateway: TDF
    TDF gateway status:
    Control Plane:
        Active Subscribers                          :         0
            IP Subscribers                          :         0
            IFL Subscribers                         :         0
    Service Plane:
        Active Subscribers                          :         0
            IP Subscribers                          :         0
            IFL Subscribers                         :         0
        Active Sessions                             :         0
    CPU Load (%)                                    :         0
    Memory Load (%)                                 :         26
```

show unified-edge tdf status detail

```
user@host> show unified-edge tdf status detail
Gateway: TDF

    FPC SLOT: 0   PIC SLOT: 2
    Role                                        :       Primary
    Type                                        :       Session-PIC
    Active Subscribers                          :       0
        IP Subscribers                          :       0
        IFL Subscribers                         :       0
    CPU Load (%)                                :       0
    Memory Load (%)                             :       26

    FPC SLOT: 0   PIC SLOT: 3
    Role                                        :       Primary
    Type                                        :       Service-PIC
    Active Subscribers                          :       0
        IP Subscribers                          :       0
        IFL Subscribers                         :       0
    Delete Pending Subscribers                  :       0
```

```
        IP Subscribers                          :       0
        IFL Subscribers                         :       0
    Active Sessions                             :       0
    CPU Load (%)                                :       1
    Memory Load (%)                             :       12

    FPC SLOT: 1   PIC SLOT: 1
    Role                                        :       Secondary
    Type                                        :       Session-PIC
    Active Subscribers                          :       0
        IP Subscribers                          :       0
        IFL Subscribers                         :       0
    CPU Load (%)                                :       0
    Memory Load (%)                             :       26
```

show unified-edge tdf status subscriber-state

```
user@host> show unified-edge tdf status subscriber-state
show unified-edge tdf status subscriber-state
Gateway: TDF
    Established          :               1
    Deleting             :               0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> show unified-edge tdf aaa statistics | **27425**
>
> *IP-Based and IFL-Based TDF Subscribers Overview*

show unified-edge tdf subscribers

IN THIS SECTION

- Syntax | 27513
- Description | 27513
- Options | 27514
- Required Privilege Level | 27514
- Output Fields | 27515
- Sample Output | 27526
- Release Information | 27531

Syntax

```
show unified-edge tdf subscribers
<brief | detail | extensive>
<business-subscribers>
<data-plane>
<domain domain-name>
<fpc-slot fpc-slot>
<gateway gateway>
<pdn-type (ipv4 | ipv4-v6 | ipv6)>
<pic-slot pic-slot>
<routing-instance routing-instance>
<stuck>
<subscriber-name subscriber-name>
<v4-addr v4-addr>
<v6-addr v6-addr>
```

Description

Display the subscriber information for one or more TDF gateways. If a TDF gateway is not specified, then subscriber information for all TDF gateways is displayed.

Options

none	(Same as brief) Display subscriber information in brief for all TDF gateways.
brief | detail | extensive	(Optional) Display the specified level of output.
business-subscribers	(Optional) Display subscriber information for only enterprise business subscribers (subscribers whose IPv4 prefix length is less than 32).
data-plane	(Optional) Display subscriber information for the data plane.
domain *domain-name*	(Optional) Display subscriber information for the specified TDF domain.
fpc-slot *fpc-slot*	(Optional) Display subscriber information for the specified FPC slot number.
gateway *gateway*	(Optional) Display subscriber information for the specified TDF gateway.
pdn-type (ipv4 | ipv4-v6 | ipv6)	(Optional) Display subscriber information according to the type of Packet Data Network (PDN): IPv4, IPv6, and both IPv4 and IPv6.
pic-slot *pic-slot*	(Optional) Display subscriber information for the specified PIC slot number. You must first specify an FPC slot number before specifying the PIC slot number.
routing-instance *routing-instance*	(Optional) Display subscriber information for the specified routing instance.
stuck	(Optional) Display subscribers for the TDF gateway that are not logged in successfully and are in a blocked state.
subscriber-name *subscriber-name*	(Optional) Display subscriber information for the specified IFL-based subscriber.
v4-addr *v4-addr*	(Optional) Display subscriber information for the specified IPv4 address of the subscriber's user equipment.
v6-addr *v6-addr*	(Optional) Display subscriber information for the specified IPv6 address of the subscriber's user equipment.

Required Privilege Level

view

Output Fields

Table 1740 on page 27515 lists the output fields for the `show unified-edge tdf subscribers` command. Output fields are listed in the approximate order in which they appear.

Table 1740: show unified-edge tdf subscribers Output Fields

Field Name	Field Description	Level of Output
`Gateway`	Name of the TDF gateway.	All levels none
`MSISDN/Name`	MSISDN number of the IP-based subscriber's user equipment.	`brief` none
`V4 Address`	IPv4 address of the IP-based subscriber.	`brief` none
`V6 Address`	IPv6 address, if any, of the IP-based subscriber. Otherwise, `None` is displayed.	`brief` none
`NAS-IP-Address`	IP address to be used for the NAS IP address attribute of the IP-based subscriber when sending the requests to the RADIUS server.	`brief` none
`Domain`	TDF domain, on the TDF gateway, to which the subscriber is attached.	`brief` none
`IFL-Subscriber-Name`	Name of the IFL-based subscriber.	`brief` none
`Subscriber Information`		

Table 1740: show unified-edge tdf subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Subscriber Type`	Type of subscriber: • `IFL`—Interface-based subscriber. • `IP`—IP-based subscriber.	`detail` `extensive`

`IMSI`	IMSI of the IP-based subscriber's user equipment.	`detail`
		`extensive`
		none

`IMEI`	International Mobile Station Equipment Identity (IMEI) of the IP-based subscriber's user equipment.	`detail` `extensive`
`MSISDN/Username`	MSISDN number of the IP-based subscriber's user equipment.	`detail` `extensive`
`Subscriber Name`	Name of the IFL-based subscriber.	`detail` `extensive`
`State`	State of the subscriber session on the signaling plane.	`detail` `extensive`
`Session Duration`	Duration of the PDP session.	`detail` `extensive`
`Domain`	Name of the TDF domain that is used to establish the session on the TDF gateway for the subscriber.	`detail` `extensive`

Table 1740: show unified-edge tdf subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Data VRF`	Name of the data plane VRF.	`detail` `extensive`
`TDF domain Name`	Unique identifier that denotes the TDF domain to be used for the subscriber's session. This setting is applicable only when the domain specified in the Create Session Request message from the subscriber is virtual.	`detail` `extensive`
`NAS-IP-Addr`	IP address to be used for the NAS IP address attribute of the IP-based subscriber when sending the requests to the RADIUS server.	`detail` `extensive`
`APN name`	Name of the APN for the IP-based subscriber that is denoted by a unique identifier.	`detail` `extensive`
`V4 Address`	IPv4 address of the IP-based subscriber.	`detail` `extensive`

`V4 Prefix Length`	IPv4 prefix length of the IP-based subscriber's IPv4 address. This is displayed only if the length is less than 32.	`detail`
		`extensive`
		`data-plane`

`V6 Address`	IPv6 address of the IP-based subscriber, if any. Otherwise, `None` is displayed.	`detail` `extensive`
`V6 Prefix Length`	IPv6 prefix length of the IP-based subscriber's IPv6 address.	`detail`
`Session PIC`	FPC and PIC slots for the session PIC on which the subscriber control session is present.	`detail` `extensive`

Table 1740: show unified-edge tdf subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Service PIC`	FPC and PIC slots for the service PIC on which the subscriber control session is present.	`detail` `extensive`

Table 1740: show unified-edge tdf subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
PCRF Event Triggers	Policy and charging rules function (PCRF) event triggers, if any. If no trigger is configured, None is displayed. The notation used for the event triggers displayed in the output and the corresponding event triggers as per the 3GPP specifications are as follows: • SGSN—SGSN CHANGE (0) • QoS—QOS CHANGE (1) • RAT—RAT CHANGE (2) • TFT—TFT CHANGE (3) • PLMN—PLMN CHANGE (4) • BL—subscriber LOSS (5) • BR—subscriber RECOVERY (6) • IPCAN—IPCAN CHANGE (7) • EAUTH—EXCEEDING AUTH (11) • RAI—RAI CHANGE (12) • ULI—ULI CHANGE (13) • NET—NO EVENT TRIGGERS (14) • OOC—OUT OF CREDIT (15) • ROC—REALLOCATION OF CREDIT (16) • REVALIDATION_TIMEOUT—REVALIDATION TIMEOUT (17) • IP ALLOC—UE_IP_ADDRESS_ALLOCATE (18) • IP RELEASE—UE_IP_ADDRESS_RELEASE (19) • DEFAULT QoS—DEFAULT QoS (20) • GW—AN GW CHANGE (21)	detail extensive

Table 1740: show unified-edge tdf subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• RA—RESOURCE_ALLOCATION (22) • RM—RESOURCE_MODIFICATION (23) • TRACE—PGW TRACE CONTROL (24) • TZ —UE_TZ_CHANGE (25) • TAI—TAI CHANGE (26) • ECGI—ECGI CHANGE (27) • CCE—CHARGING CORRELATION EXCHANGE (28) • AMBR—AMBR CHANGE (29) • UCIC—USR CSG INFO CHANGE (30) • QMF—QoS MODIFICATION FAILURE (31) • UR—USAGE REPORT (33)	
Revalidation due in	Time remaining in days, hours, minutes, and seconds until PCEF session revalidation takes place if the REVALIDATION_TIMEOUT event trigger is armed. Otherwise N/A is displayed.	detail extensive
Idle Timeout	Idle timeout for the session, in minutes.	detail extensive
Subscriber MBR	TDF subscriber maximum bit rate (MBR) for uplink and downlink traffic. Uplink traffic originates from the subscriber towards the public data network (PDN), and downlink traffic comes from the PDN and is destined for the subscriber.	detail extensive
Subscriber burst	TDF subscriber burst size configured for uplink and downlink traffic. Uplink traffic originates from the subscriber towards the PDN, and downlink traffic comes from the PDN and is destined for the subscriber.	detail extensive

Table 1740: show unified-edge tdf subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Access IFL List`	The following is displayed for each interface assigned to an IFL-based subscriber: • `Name`—Name of the interface. • `Index`—Index number of the interface. • `State`—Operational state of the interface: `Active` or `Inactive`.	
`PCC Profile Name`	Name of the PCEF profile that is assigned to the subscriber.	`detail` `extensive`

`Usage Monitoring Information`	The following is displayed for each monitoring key, which corresponds to a data set that is being monitored for the subscriber:
	• `Monitoring Key`—Identifier for the monitoring key.
	• `Level`—Indication of whether the monitoring key applies to particular PCC rules (Rule) or to the entire TDF subscriber session (Session).
	• `Status`—Indication of whether monitoring with the key is active or inactive.
	• `Total Available Quota`—Volume and time quota sent from the PCRF to indicate when a report should be sent to the PCRF. A value of zero indicates that the field is not applicable to the key.
	  • `Input`—Uplink traffic volume quota.
	  • `Output`—Downlink traffic volume quota.
	  • `Total`—Uplink and downlink traffic volume quota.
	  • `Time`—Time quota, in seconds.
		`detail`
		`extensive`

Table 1740: show unified-edge tdf subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output

PCC Rule Information	**NOTE**: Both ePCC rules and PCC rules appear under PCC Rule Information. Fields that apply only to ePCC rules are identified in the description.
	
	The following information for each PCC or ePCC rule is displayed per subscriber:
	
	• Rule Name—Name of the rule. In addition, the following is displayed:
	  • Type—Rrule type: Static or Dynamic.
	  • Associated Rule Base—Rule set with which the rule is associated, if any.
	  • Precedence—Rule precedence, which defines the order in which the policy is applied for incoming or outgoing packets; the lower the number, the higher its precedence.
	  • Activation due in—Day, time, or both at which the rule is scheduled for activation for the subscriber. If activation/deactivation settings have not been applied to the rule, then N/A appears.
	  • Deactivation due in—Day, time, or both at which the rule is scheduled for deactivation for the subscriber. If activation/deactivation settings have not been applied to the rule, then N/A appears.
	  • Status—Rule status: Initialized, Active, Inactive, or Removal Pending.
	  • Application Id—(ePCC rules only) Name of the application identification parameter associated with the rule.
	  • Application Id Base–(ePCC rules only) Name of the base application that serves as the primary application identification service if a group or cluster are configured.
	  • Mute Notification—(ePCC rules only) Whether the Mute-Notification AVP is included in the rule.
	
	• QoS Parameters—The following QoS attributes are displayed for each rule per subscriber:
	  • MBR Uplink (kbps)—Maximum bit rate (MBR) in the uplink direction, in kbps. Identifier.
		detail
		
		extensive

Table 1740: show unified-edge tdf subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output

	• `MBR Downlink (kbps)`—MBR in the downlink direction, in kbps.
	• `Burst size Uplink (bytes)`—TDF domain's default TDF subscriber burst size configured for uplink traffic, in bytes.
	• `Burst size Downlink (bytes)`—TDF domain's default TDF subscriber burst size configured for downlink traffic, in bytes. Uplink traffic originates from the subscriber towards the public data network (PDN), and downlink traffic comes from the PDN and is destined for the subscriber.
	• `Charging Attributes`—The following charging attributes are displayed for each rule per subscriber:
	  • `Rating Group`—Rating group for the rule.
	  • `Service Id`—Service ID for the rule.
	  • `Gating Status`—Whether the flow is enabled or not. One of the following:
	  • `enable-uplink`
	  • `enable-downlink`
	  • `enable-both`
	  • `disable-both`
	  • `AF Charging Id`—Application function record information, which contains an octet string and the charging ID.
	  • `Charging Method`—Charging method for the rule (`none`, `offline`, `offline-online`, or `online`).
	  • `Metering Method`—Charging metering method for the rule:
	  • `Time`—Time based.
	  • `Volume`—Volume based.	

Table 1740: show unified-edge tdf subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Volume-Time`—Both volume and time based. • `None`—No metering. • `Usage Monitoring Key`—Monitoring key that is associated with the rule. • `Services Attributes`—The following information about resource management and steering is displayed for the subscribers connected to the TDF gateway or the TDF domain: • `Steering IP`—IPv4 or IPv6 address for HTTP steering of the packets. • `Keep existing steering`—Whether existing steering is enabled or disabled. • `Service Chain VRF`—Routing instance for steering of packets. Use this to steer traffic to either a local service chain or external service chain. • `Forwarding Class`—Forwarding class that needs to be assigned to the packet. • `HCM ID`—Profile that identifies the HTTP header enrichment rules to apply. This action is restricted to PCC rules that are only matching HTTP-based applications. • `LRF ID`—Unique ID of the Location Retrieval Function • `Filter Attributes`—The following filter attributes are displayed per filter in each rule: • `Remote IP/Mask`—Remote IP address and subnet mask of the filter. • `Protocol`—Protocol configured for the filter. If all protocols are supported for the filter, `Any` is displayed. For the explanation of what the numbers represent, refer to the 3GPP specifications. • `Direction`—Direction in which the filter is applicable (`Downlink`, `Uplink`, or `Both`). • `Local Ports`—Local ports or port range for the filter. `Any` indicates that the filter does not restrict the local ports.	

Table 1740: show unified-edge tdf subscribers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Remote Ports`—Remote ports or port range for the filter. `Any` indicates that the filter does not restrict the remote ports. • `Application Name`—(ePCC rules only) Name of the predefined or custom application signature.	
`Data Plane statistics`	The following information about packets processed by the data plane for subscribers connected to the TDF domains in the TDF gateway: • `Subscriber-Name`—Name of the IFL-based subscriber. • `V4 Address`—IPv4 address of the IP-based subscriber. • `V6 Address`—IPv6 address of the IP-based subscriber. • `V6 Prefix Length`—IPv6 prefix length of the IP-based subscriber's IPv6 address. • `Vrf Id`—Name of the data plane VRF. • `Subscriber Stats`—Total statistics for the subscriber. • `Rule`—Statistics for traffic that was handled by the specified PCC rule. • `Uplink`—Statistics for traffic in the uplink direction from the TDF gateway to the PDN (Internet). • `Downlink`—Statistics for traffic in the downlink direction from the PDN (Internet) to the TDF gateway. • `Sessions`—Number of sessions in the uplink and downlink direction. • `Packets`—Number of packets forwarded in the uplink direction and in the downlink direction. • `Bytes`—Number of bytes forwarded in the uplink direction and in the downlink direction. • `Dropped Packets`—Number of packets dropped in the uplink direction and in the downlink direction. • `Dropped Bytes`—Number of bytes dropped in the uplink direction and in the downlink direction.	`data-plane` option

Sample Output

show unified-edge tdf subscribers (IP-Based Subscriber)

```
user@host> show unified-edge tdf subscribers
Gateway: TDF
   MSISDN/name       V4 Address      V6 Address          NAS-IP-Address  Domain
   att               192.0.2.11     None                198.51.100.123 domain1
```

show unified-edge tdf subscribers (IFL-Based Subscriber)

```
user@host> show unified-edge tdf subscribers
IFL-Subscriber-Name                                   Domain
   ifl-sub-radius-001                                 domain-ifl-radius
   ifl-sub-static-001                                 domain-ifl-static
```

show unified-edge tdf subscribers extensive

```
user@host> show unified-edge tdf subscribers extensive
Gateway: TDF

Subscriber Information:
     Subscriber Type : IFL
     Subscriber Name : IFL1
     State      : Established                     Session Duration: 000065 hrs 50 mins 15 secs
     Domain     : domain1
     Data VRF   : default
     Session PIC: 3 /0  (FPC/PIC)
     Service PIC: 3 /1  (FPC/PIC)
     PCRF Event Triggers : None
     Revalidation due in : N/A
     Subscriber MBR:   Uplink (kbps):  0             Downlink (kbps):  0
     Subscriber burst: Uplink (bytes): 0             Downlink (bytes): 0
     Access IFL List:          Name      (Index)      State
                         ge-1/1/8.0      (362  )      Active
                         ge-1/1/1.0      (361  )      Active
                         ge-1/0/9.0      (360  )      Active
```

```
PCC Profile Name : pcef-prof-static
PCC Rule Information:
Rule Name: google-traffic
   Type      : Static             Associated Rule Base: None
   Precedence: 20                 Status: Active
   Activation due in  : N/A
   Deactivation due in: N/A
 QoS Parameters:
   MBR Uplink (kbps):         0           MBR Downlink (kbps):         0
   Burst size Uplink (bytes): 0           Burst size Downlink (bytes): 0
 Charging Attributes:
   Rating Group: 0          Service ID: 0               Gating Status: enable-both
   AF Charging Id: None     Charging Method: None       Metering Method: None
   Usage Monitoring Key :  NULL
 Logging Rule Name :  r1
 Services Attributes:
   Forwarding Class: best-effort
 Filter Attributes:
   Remote IP/Mask: any/any            Protocol: any Direction: Both
   Local Ports:  any
   Remote Ports: any
   Application Name   : junos:google
   Application Name   : junos:udp
   Application Name   : junos:http
Rule Name: http-traffic
   Type      : Static             Associated Rule Base: None
   Precedence: 30                 Status: Active
   Activation due in  : N/A
   Deactivation due in: N/A
 QoS Parameters:
   MBR Uplink (kbps):         0           MBR Downlink (kbps):         0
   Burst size Uplink (bytes): 0           Burst size Downlink (bytes): 0
 Charging Attributes:
   Rating Group: 0          Service ID: 0               Gating Status: enable-both
   AF Charging Id: None     Charging Method: None       Metering Method: None
   Usage Monitoring Key :  NULL
 Logging Rule Name :  r1
 Services Attributes:
   Forwarding Class: best-effort
 Filter Attributes:
   Remote IP/Mask: any/any            Protocol: any Direction: Both
   Local Ports:  any
   Remote Ports: any
```

```
        Application Name   : junos:http
    Rule Name: all-traffic
        Type      : Static             Associated Rule Base: None
        Precedence: 40                 Status: Active
        Activation due in  : N/A
        Deactivation due in: N/A
      QoS Parameters:
        MBR Uplink (kbps):         0           MBR Downlink (kbps):         0
        Burst size Uplink (bytes): 0           Burst size Downlink (bytes): 0
      Charging Attributes:
        Rating Group: 0          Service ID: 0               Gating Status: enable-both
        AF Charging Id: None     Charging Method: None       Metering Method: None
        Usage Monitoring Key :  NULL
      Logging Rule Name :  r1
      Services Attributes:
        Forwarding Class: best-effort
      Filter Attributes:
        Remote IP/Mask: any/any            Protocol: any Direction: Both
        Local Ports:  any
        Remote Ports: any
```

show unified-edge tdf subscribers detail

```
user@host> show unified-edge tdf subscribers detail
Gateway: TDF

Subscriber Information:
     Subscriber Type : IP

     IMSI       : 988888888888899                MSISDN/Username : 9741488201
     IMEI       : None
     State      : Established                     Session Duration: 000000 hrs 41 mins 04 secs
     Domain     : aaa
     Data VRF   : bng_vrf
     NAS-IP-Addr: 198.51.100.123
     NAS-ID     : dfssw
     APN name   : 3242
     V4 Address : 192.0.2.11
     V6 Address : 2001:db8::                        V6 Prefix Length: 64
     Session PIC: 3 /3  (FPC/PIC)
     Service PIC: 3 /0  (FPC/PIC)
```

```
 PCRF Event Triggers : UR
 Revalidation due in : N/A
 Idle Timeout: 0   min
 Subscriber MBR:   Uplink (kbps):  0             Downlink (kbps):  0
 Subscriber burst: Uplink (bytes): 0             Downlink (bytes): 0
PCC Profile Name : pcef-jpkt-prof-dyn
Usage Monitoring Information:
  Monitoring Key: 302
   Level: Session Status: Active
     Total Available Quota:
       Input: 0       Output: 0
       Total: 20000   Time  : 100
  Monitoring Key: 301
    Level: PCC-Rule Status: Active
      Total Available Quota:
        Input: 0       Output: 0
        Total: 20000   Time  : 0
PCC Rule Information:
Rule Name: Dyn_Rule_1
   Type : Dynamic Associated Rule Base: None
   Precedence: 1 Status: Active
   Activation due in : N/A
   Deactivation due in: N/A
 QoS Parameters:
   MBR Uplink (kbps): 2000 MBR Downlink (kbps): 3000
   Burst size Uplink (bytes): 0 Burst size Downlink (bytes): 0
 Charging Attributes:
   Rating Group: 0 Service ID: 0 Gating Status: enable-uplink
   AF Charging Id: None Charging Method: None Metering Method: None
   Usage Monitoring Key : 301
 Services Attributes:
   Steering VRF Uplink: changed_vrf     Downlink: new_vrf
   HCM ID: hcmtag1
 Filter Attributes:
   Remote IP/Mask: 203.0.113/32 Protocol: 1 Direction: Both
   Local Ports: any
   Remote Ports: any
   Application Id : None
   Application Id Base: None
```

show unified-edge tdf subscribers data-plane

```
user@host> show unified-edge tdf subscribers data-plane
Gateway: TDF
Data plane statistics :

V4 Address:192.0.2.11
V6 Address:2001:db8::                        V6 Prefix Length: 64
Subscriber-Type:IP
Vrf Id: 11
Subscriber Stats:
                      Uplink                Downlink
-----------------------------------------------------------------
    Packets          :0                    :0
    Bytes            :0                    :0
    Dropped Packets  :0                    :0
    Dropped Bytes    :0                    :0


Rule: rule_zynga
                      Uplink                Downlink
-----------------------------------------------------------------
    Sessions         :0                    :0
    Packets          :0                    :0
    Bytes            :0                    :0
    Dropped Packets  :0                    :0
    Dropped Bytes    :0                    :0


Rule: rule_youtube
                      Uplink                Downlink
-----------------------------------------------------------------
    Sessions         :0                    :0
    Packets          :0                    :0
    Bytes            :0                    :0
    Dropped Packets  :0                    :0
    Dropped Bytes    :0                    :0


Rule: rule_amazon
                      Uplink                Downlink
-----------------------------------------------------------------
```

```
    Sessions         :0                    :0
    Packets          :0                    :0
    Bytes            :0                    :0
    Dropped Packets  :0                    :0
    Dropped Bytes    :0                    :0


Rule: rule_monster
                      Uplink                Downlink
-----------------------------------------------------------------
    Sessions         :0                    :0
    Packets          :0                    :0
    Bytes            :0                    :0
    Dropped Packets  :0                    :0
    Dropped Bytes    :0                    :0


Rule: all-traffic-s
                      Uplink                Downlink
-----------------------------------------------------------------
    Sessions         :0                    :0
    Packets          :0                    :0
    Bytes            :0                    :0
    Dropped Packets  :0                    :0
    Dropped Bytes    :0                    :0
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf statistics | **27496**

clear unified-edge tdf subscribers | **4269**

IP-Based and IFL-Based TDF Subscribers Overview

show unified-edge tdf system interfaces

IN THIS SECTION

- Syntax | 27532
- Description | 27532
- Options | 27532
- Required Privilege Level | 27532
- Output Fields | 27533
- Sample Output | 27534
- Release Information | 27534

Syntax

```
show unified-edge tdf system interfaces
<gateway gateway-name>
```

Description

Display information about the aggregated Packet Forwarding Engine and the aggregated multiservices (AMS) interfaces and their states on one or more TDF gateways. If a TDF gateway is not specified, then information for all TDF gateways is displayed.

Options

none	Display information for all TDF gateways.
gateway *gateway-name*	(Optional) Display information for the specified TDF gateway.

Required Privilege Level

view

Output Fields

Table 1741 on page 27533 lists the output fields for the `show unified-edge tdf system interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 1741: show unified-edge tdf system interfaces Output Fields

Field Name	Field Description
`Gateway`	Name of the TDF gateway.
`Interfaces`	Name of the interface: • Aggregated multiservices; for example, ams0 • Aggregated Packet Forwarding Engine; for example, apfe1 • Member of aggregated multiservices; for example, mams-1/0/0 • Multiservices; for example, ms-1/0/0 • Packet Forwarding Engine; for example, pfe-0/1/0
`Members`	For ams and apfe interfaces, the member interfaces that are part of the aggregated interfaces are displayed.
`Operational State`	Whether the interface is operational (`Active`) or not (`Inactive`).
`Redundancy Role`	Redundancy state in which the interface is configured: • `Primary`—Interface is a primary member. • `Secondary`—Interface is a backup to all the primary members. • `Standalone`—Interface has not been configured for redundancy.

Sample Output

show unified-edge tdf system interfaces

```
user@host> show unified-edge tdf system interfaces
Gateway: TDF
    Interfaces          Members        Operational    Redundancy
                                          State          Role
    ms-1/0/0                           Active         Standalone
    ms-1/1/0                           Active         Standalone
    ms-2/0/0                           Active         Standalone
    ms-2/1/0                           Active         Standalone
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- show interfaces anchor-group (Aggregated Packet Forwarding Engine) | 20437
- show interfaces load-balancing (Aggregated Multiservices) | 21024
- show unified-edge tdf resource-manager clients | 27490
- show unified-edge tdf system interfaces service-mode | 27534

show unified-edge tdf system interfaces service-mode

IN THIS SECTION

- Syntax | 27535
- Description | 27535
- Options | 27535
- Required Privilege Level | 27535

- Output Fields | 27535
- Sample Output | 27537
- Release Information | 27537

Syntax

```
show unified-edge tdf system interfaces service-mode
<brief | detail>
<gateway gateway-name>
```

Description

Display the service mode information for the interfaces on one or more TDF gateways. If a TDF gateway is not specified, then information for all TDF gateways is displayed.

Options

none	(Same as brief) Display service mode information for all TDF gateways.
brief | detail	(Optional) Display the specified level of output.
gateway-name *gateway-name*	(Optional) Display service mode information for the specified TDF gateway.

Required Privilege Level

view

Output Fields

Table 1742 on page 27536 lists the output fields for the `show unified-edge tdf system interfaces service-mode` command. Output fields are listed in the approximate order in which they appear.

Table 1742: show unified-edge tdf system interfaces service-mode Output Fields

Field Name	Field Description	Level of Output
`Maintenance Mode`	Phases applicable when the TDF interface is in maintenance mode. • `MM - Active Phase`—All the attributes of the address pool can be modified. • `MM - In/Out Phase`—Only the non-maintenance mode attributes of the address pool can be modified.	None `brief`
`Interface Name`	Name of the interface for which the service mode information is displayed: • Aggregated multiservices; for example, ams0 • Aggregated Packet Forwarding Engine; for example, apfe1 • Multiservices; for example, ms-1/0/0	All levels
`Gateway`	Name of the TDF gateway.	None `brief`
`Gateway Name`	Name of the TDF gateway.	`detail`
`Service Mode Status`	Status of service mode for the TDF gateway.	`detail`
`Service Mode`	Service mode for the TDF gateway. The following service modes are possible: • `Operational`—Gateway is in operational mode. • `Maintenance`—Gateway is in maintenance mode.	All levels

Sample Output

show unified-edge tdf system interfaces service-mode brief

```
user@host> show unified-edge tdf system interfaces service-mode brief
Maintenance Mode
    MM Active Phase - System is ready to accept configuration changes for all
                      attributes of this object and its sub-hierarchies.
    MM In/Out Phase - System is ready to accept configuration changes only for
                      non-maintenance mode attributes of this object and
                      its sub-hierarchies.

Interface Name     Gateway                              Service Mode

ms-2/1/0          TDF                              Operational
ams1               TDF                              Operational
```

show unified-edge tdf system interfaces service-mode detail

```
user@host> show unified-edge tdf system interfaces service-mode detail
Service Mode Status
Interface Name  : ms-2/1/0
Gateway Name    : TDF
Service Mode    : Operational
Service Mode Status
Interface Name  : ams1
Gateway Name    : TDF
Service Mode    : Operational
```

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

show unified-edge tdf system interfaces | 27532

show uplink-failure-detection

IN THIS SECTION

- Syntax | 27538
- Description | 27538
- Options | 27538
- Required Privilege Level | 27538
- Output Fields | 27539
- Sample Output | 27539
- Sample Output | 27540
- Release Information | 27540

Syntax

```
show uplink-failure-detection
<group  group-name>
```

Description

Display information about the uplink-failure-detection group, the member interfaces, and their status.

Options

none	Display information about all groups configured for uplink failure detection.
group *group-name*	(Optional) Display information about the specified group only.

Required Privilege Level

view

Output Fields

Table 1743 on page 27539 lists the output fields for the `show uplink-failure-detection` command. Output fields are listed in the approximate order in which they appear.

Table 1743: show uplink-failure-detection Output Fields

Field Name	Field Description
Group	Name of the group.
Uplink	The uplink interface or interfaces configured as link-to-monitor. **NOTE**: The asterisk (*) indicates that the link is up.
Downlink	The downlink interface or interfaces configured as link-to-disable. **NOTE**: The asterisk (*) indicates that the link is up.
Failure Action	Status of uplink failure detection: • Active—The switch has detected an uplink failure and has brought the downlink down. • Inactive—The uplink or uplinks are up.
Debounce Interval	The amount of time, in seconds, that elapses before the downlink interfaces are brought up after a state change of the uplink interfaces.

Sample Output

show uplink-failure-detection

```
user@switch> show uplink-failure-detection
Group             : group1
Uplink            :  ge-0/0/0*
Downlink          :  ge-0/0/1*
Failure Action    :  Inactive

Group             : group2
Uplink            :  ge-0/0/3.0
```

```
Downlink          :  ge-0/0/4.0
Failure Action    :  Active

Debounce Interval :  20
```

Sample Output

show uplink-failure-detection group g2

```
user@switch> show uplink-failure-detection group g2
Group             : group1
Uplink            :  ge-0/0/0*
Downlink          :  ge-0/0/1*
Failure Action    :  Inactive

Group             : group2
Uplink            :  ge-0/0/3.0
Downlink          :  ge-0/0/4.0
Failure Action    :  Active

Debounce Interval :  20
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Overview of Uplink Failure Detection

Configuring Interfaces for Uplink Failure Detection

Example: Configuring Interfaces for Uplink Failure Detection

show uplink-failure-detection

IN THIS SECTION

- Syntax | 27541
- Description | 27541
- Options | 27541
- Required Privilege Level | 27541
- Output Fields | 27542
- Sample Output | 27542
- Sample Output | 27543
- Release Information | 27543

Syntax

```
show uplink-failure-detection
<group  group-name>
```

Description

Display information about the uplink-failure-detection group, the member interfaces, and their status.

Options

none	Display information about all groups configured for uplink failure detection.
group *group-name*	(Optional) Display information about the specified group only.

Required Privilege Level

view

Output Fields

Table 1744 on page 27542 lists the output fields for the `show uplink-failure-detection` command. Output fields are listed in the approximate order in which they appear.

Table 1744: show uplink-failure-detection Output Fields

Field Name	Field Description
Group	Name of the group.
Uplink	The uplink interface or interfaces configured as link-to-monitor. **NOTE**: The asterisk (*) indicates that the link is up.
Downlink	The downlink interface or interfaces configured as link-to-disable. **NOTE**: The asterisk (*) indicates that the link is up.
Failure Action	Status of uplink failure detection: • Active—The switch has detected an uplink failure and has brought the downlink down. • Inactive—The uplink or uplinks are up.
Debounce Interval	The amount of time, in seconds, that elapses before the downlink interfaces are brought up after a state change of the uplink interfaces.

Sample Output

show uplink-failure-detection

```
user@switch> show uplink-failure-detection
Group             : group1
Uplink            :  ge-0/0/0*
Downlink          :  ge-0/0/1*
Failure Action    :  Inactive

Group             : group2
Uplink            :  ge-0/0/3.0
```

```
Downlink          :  ge-0/0/4.0
Failure Action    :  Active

Debounce Interval :  20
```

Sample Output

show uplink-failure-detection group g2

```
user@switch> show uplink-failure-detection group g2
Group             : group1
Uplink            :  ge-0/0/0*
Downlink          :  ge-0/0/1*
Failure Action    :  Inactive

Group             : group2
Uplink            :  ge-0/0/3.0
Downlink          :  ge-0/0/4.0
Failure Action    :  Active

Debounce Interval :  20
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Interfaces for Uplink Failure Detection (CLI Procedure)

show v4ov6-tunnels

IN THIS SECTION

- Syntax | 27544
- Description | 27544
- Options | 27544
- Required Privilege Level | 27545
- Output Fields | 27545
- Sample Output | 27546
- Release Information | 27546

Syntax

```
show v4ov6-tunnels information
<anchor-pfe>
<anti-spoof-ip>
<fabric>
<logical-system>
<tcnh-index>
<v6-dest>
<v6-source>
```

Description

Display all configured V4oV6 tunnels in the routing protocol process (rpd). If the tunnel composite `NHINDEX` is `0`, then the route is not yet installed in the forwarding information base (FIB) also known as the forwarding table.

Options

none Display dynamic tunnel localization information for the Packet Forwarding Engine tunnel.

anchor-pfe	Filter the V4oV6 tunnels in the routing protocol process based on the anchor-pfe.
anti-spoof-ip	(Optional) Filter the IPv4-over-IPv6 tunnels in the routing protocol process based on the anti-spoof IP address.
fabric	(Optional) Internal fabric state.
logical-system (all | *logical-system-name*	(Optional) Perform this operation on all logical systems or on a particular logical system.
tcnh-index	(Optional) Filter the V4oV6 tunnels in the routing protocol process based on the tcnh-index.
v6-dest	(Optional) Filter the V4oV6 tunnels in the routing protocol process based on the IPv6 destination address.
v6-source	Filter the V4oV6 tunnels in the routing protocol process based on the IPv6 source address.

Required Privilege Level

view

Output Fields

Table 1745 on page 27545 lists the output fields for the `show v4ov6-tunnels` command. Output fields are listed in the approximate order in which they appear.

Table 1745: show v4ov6-tunnels

Field Name	Field Description
`Destination`	Destination IPv6 address of the dynamic tunnel.
`Source`	Source IPv6 address of the dynamic tunnel.
`Antispoof`	Anti-spoof address of the dynamic tunnel.
`Antispoof Status`	Status of the anti-spoof ability.

Table 1745: show v4ov6-tunnels *(Continued)*

Field Name	Field Description
Mtu	Maximum transmission unit configured for the dynamic tunnel.
Anchor pfe	Anchor PFE of the dynamic tunnel.
tcnh	Tunnel composite next hop address.
tcnh-index	Index value of tunnel composite next hop address.
tcnh-refcount	Reference count of the tunnel composite next hop address.

Sample Output

show v4ov6-tunnels information

```
user@host> show v4ov6-tunnels information
Destination                        Source          Antispoof      Antispoof Status
Mtu          Anchor pfe    tcnh       tcnh-index    tcnh-refcount2001:db8:5555::5555
2001:db8:9999::9999      10.4.4.4              DISABLED       1400         pfe-0/0/0
0x9dc0f7c    559            2
2001:db8:3333::3333          2001:db8:7777::7777      10.2.2.2              DISABLED
1400         pfe-0/0/0    0x9dc1858    561           2
2001:db8:4444::4444          2001:db8:8888::8888      10.3.3.3              DISABLED
1400         pfe-0/0/0    0x9dc11c8    560           2
2001:db8:2222::2222          2001:db8:6666::6666      10.1.1.1              DISABLED
1400         pfe-0/0/0    0x9dc1fe4    622            2
```

Release Information

Command introduced in 17.3R1.

RELATED DOCUMENTATION

dynamic-tunnels | 6161

extended-nexthop | 6700

tunnel-attributes | 30194

show dynamic-tunnels pfe-tunnel-localization | 19609

Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP

show validation database

IN THIS SECTION

- Syntax | 27547
- Description | 27548
- Options | 27548
- Required Privilege Level | 27548
- Output Fields | 27548
- Sample Output | 27549
- Release Information | 27550

Syntax

```
show validation database
<name database-name>
<brief | detail>
<instance instance-name>
<logical-system logical-system-name>
<mismatch>
<origin-autonomous-system as-number>
<record ip-prefix>
<session ip-address>
```

Description

Display information about the route validation database when resource public key infrastructure (RPKI) BGP route validation is configured. You can query all route validation records that match a given prefix or origin-autonomous-system. In addition, you can filter the output by a specific RPKI cache session.

Options

none	Display all route validation database entries.
name database-name	(Optional) Display the records of the specified database in that routing-instance or logical-system. If unspecified, the records of all databases are displayed.
brief | detail	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display information about route validation database entries for the specified routing instance. The instance name can be primary for the main instance, or any valid configured instance name or its prefix.
logical-system *logical-system-name*	(Optional) Perform this operation on a particular logical system.
mismatch	(Optional) Filter the output by mismatched origin autonomous systems.
origin-autonomous-system *as-number*	(Optional) Filter the output by mismatched origin autonomous systems. The `mismatch` qualifier is useful for finding conflicting origin-autonomous-system information between RPKI caches. Mismatches might occur during cache reconfiguration.
record *ip-prefix*	(Optional) Filter the output by route validation records that match a given prefix.
session *ip-address*	(Optional) Filter the output by a specific RPKI cache session.

Required Privilege Level

view

Output Fields

Table 1746 on page 27549 describes the output fields for the `show validation database` command. Output fields are listed in the approximate order in which they appear.

Table 1746: show validation database Output Fields

Field Name	Field Description	Level of Output
Prefix	Route validation (RV) record prefix. RV records are received from the cache server and can also be configured statically at the [`edit routing-options validation static`] hierarchy level .	All levels
Origin-AS	Legitimate originator autonomous system (AS).	All levels
Session	IP address of the RPKI cache server.	All levels
State	State of the route validation records. The state can be `valid`, `invalid` or `unknown`.	All levels
Mismatch	Conflicting origin-autonomous-system information between RPKI caches when nonstop active routing (NSR) is configured.	All levels
IPv4 records	Number of IPv4 route validation records.	All levels
IPv6 records	Number of IPv6 route validation records.	All levels

Sample Output

show validation database

```
user@host> show validation database
RV database: default
Prefix             Origin-AS   Session   State   Mismatch
172.16.1.0/24-32               1 10.0.77.1    valid
172.16.2.0/24-32               2 10.0.77.1    valid
172.16.3.0/24-32               3 10.0.77.1    valid
172.16.4.0/24-32               4 10.0.77.1    valid
```

```
IPv4 records: 4
IPv6 records: 0

RV database: foo
Prefix            Origin-AS    Session   State   Mismatch
172.16.5.0/24-32               5 10.0.77.1    valid
172.16.6.0/24-32               6 10.0.77.1 valid
172.16.7.0/24-32               7 10.0.77.1 valid
172.16.8.0/24-32               8 10.0.77.1 valid
72.9.224.0/19-24       26234 192.168.1.100 valid *
72.9.224.0/19-24       3320 192.168.1.200 invalid *
10.0.0.0/8-32             0 internal valid

IPv4 records: 7
IPv6 records: 0

RV database: vrf1.red

Prefix           Origin-AS    Session   State    Mismatch
172.16.5.0/24-32              5 10.0.77.1 valid
172.16.6.0/24-32              6 10.0.77.1 valid
172.16.7.0/24-32              7 10.0.77.1 valid
172.16.8.0/24-32              8 10.0.77.1 valid
72.9.224.0/19-24        26234 192.168.1.100 valid *
72.9.224.0/19-24        3320 192.168.1.200 invalid *
10.0.0.0/8-32              0 internal valid

IPv4 records: 7
IPv6 records: 0
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- Use Case and Benefit of Origin Validation for BGP
- Understanding Origin Validation for BGP
- Example: Configuring Origin Validation for BGP

show validation group

IN THIS SECTION

- Syntax | 27551
- Description | 27551
- Options | 27551
- Required Privilege Level | 27552
- Output Fields | 27552
- Sample Output | 27553
- Release Information | 27553

Syntax

```
show validation group
<instance instance-name>
<logical-system logical-system-name>
```

Description

Display information about route validation redundancy groups.

Options

none	Display information about all route validation groups.
instance *instance-name*	(Optional) Display information about route validation groups for the specified routing instance. The instance name can be primary for the main instance, or any valid configured instance name or its prefix.
logical-system *logical-system-name*	(Optional) Perform this operation on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1747 on page 27552 describes the output fields for the `show validation group` command. Output fields are listed in the approximate order in which they appear.

Table 1747: show validation group Output Fields

Field Name	Field Description
Group	Group name.
Maximum sessions	Number of concurrent sessions for each group. The default is 2. The number is configured with the `max-sessions` statement.
Session	Resource public key infrastructure (RPKI) cache session IP address.
State	State of the connection between the routing device and the cache server. `Up` means that the connection is established. `Connect` means that the connection is not established.
Preference	Each cache server has a preference. Higher preferences are preferred. During a session start or restart, the routing device attempts to start a session with the cache server that has the numerically highest preference. The routing device connects to multiple cache servers in preference order. The default preference is 100. The preference value is configured with the `preference` statement at the [edit routing-options validation group group-name session] hierarchy level.

Sample Output

show validation group

```
user@host> show validation group
master
    Group: test, Maximum sessions: 3
      Session 10.255.255.11, State: Up, Preference: 100
      Session 10.255.255.12, State: Up, Preference: 100
    Group: test2, Maximum sessions: 2
      Session 10.255.255.13, State: Connect, Preference: 100
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- Use Case and Benefit of Origin Validation for BGP
- Understanding Origin Validation for BGP
- Example: Configuring Origin Validation for BGP

show validation replication database

IN THIS SECTION

- Syntax | 27554
- Description | 27554
- Options | 27554
- Required Privilege Level | 27555
- Output Fields | 27555
- Sample Output | 27555
- Release Information | 27556

Syntax

```
show validation replication database
<brief | detail>
<name database-name>
<instance instance-name>
<logical-system logical-system-name>
<origin-autonomous-system as-number>
<record ip-prefix>
<session ip-address>
```

Description

Display the state of the nonstop active routing (NSR) records. The output is the same as the output of the `show validation database` command, except for the `Mismatch` column.

Options

none	Display all route validation database entries.
name database-name	(Optional) Display the records of the specified database in that routing-instance or logical-system. If unspecified, the records of all databases are displayed.
brief | detail	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display information about route validation database entries for the specified routing instance. The instance name can be primary for the main instance, or any valid configured instance name or its prefix.
logical-system *logical-system-name*	(Optional) Perform this operation on a particular logical system.
origin-autonomous-system *as-number*	(Optional) Filter the output by mismatched origin autonomous systems. The `mismatch` qualifier is useful for finding conflicting origin-autonomous-system information between resource public key infrastructure (RPKI) caches. Mismatches might occur during cache reconfiguration.
record *ip-prefix*	(Optional) Filter the output by route validation records that match a given prefix.
session *ip-address*	(Optional) Filter the output by a specific RPKI cache session.

Required Privilege Level

view

Output Fields

Table 1748 on page 27555 describes the output fields for the `show validation replication database` command. Output fields are listed in the approximate order in which they appear.

Table 1748: show validation replication database Output Fields

Field Name	Field Description	Level of Output
Prefix	Route validation (RV) record prefix. RV records are received from the cache server and can also be configured statically at the [`edit routing-options validation static`] hierarchy level.	All levels
Origin-AS	Legitimate originator autonomous system (AS).	All levels
Session	IP address of the RPKI cache server.	All levels
State	State of the route validation records. The state can be `valid` or `invalid`.	All levels
IPv4 records	Number of IPv4 route validation records.	All levels
IPv6 records	Number of IPv6 route validation records.	All levels

Sample Output

show validation replication database

```
user@host> show validation replication database
RV database for instance master
```

```
Prefix                Origin-AS Session         State
172.16.1.0/24-32              1 10.0.77.1       valid
172.16.2.0/24-32              2 10.0.77.1       valid
172.16.3.0/24-32              3 10.0.77.1       valid
172.16.4.0/24-32              4 10.0.77.1       valid
172.16.5.0/24-32              5 10.0.77.1       valid
172.16.6.0/24-32              6 10.0.77.1       valid
172.16.7.0/24-32              7 10.0.77.1       valid
172.16.8.0/24-32              8 10.0.77.1       valid
72.9.224.0/19-24          26234 192.168.1.100   valid
72.9.224.0/19-24           3320 192.168.1.200   invalid
10.0.0.0/8-32                 0 internal        valid

IPv4 records: 14
IPv6 records: 0
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- Use Case and Benefit of Origin Validation for BGP
- Understanding Origin Validation for BGP
- Example: Configuring Origin Validation for BGP

show validation session

IN THIS SECTION

- Syntax | 27557
- Description | 27557
- Options | 27557
- Required Privilege Level | 27557
- Output Fields | 27558

Sample Output | 27562
Release Information | 27562

Syntax

```
show validation session
<brief | detail>
<destination>
<instance instance-name>
<logical-system logical-system-name>
```

Description

Display information about all sessions or a specific session with a resource public key infrastructure (RPKI) cache server.

Options

none	Display information about all sessions.
destination	(Optional) Display information about a specific session.
brief | detail	(Optional) Display the specified level of output.
instance *instance-name*	(Optional) Display information about sessions for the specified routing instance. The instance name can be primary for the main instance, or any valid configured instance name or its prefix.
logical-system *logical-system-name*	(Optional) Perform this operation on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1749 on page 27558 describes the output fields for the `show validation session` command. Output fields are listed in the approximate order in which they appear.

Table 1749: show validation session Output Fields

Field Name	Field Description	Level of Output
Session	IP address of the RPKI cache server. You configure the session and all of its elements with the `session` statement.	All levels

Table 1749: show validation session Output Fields *(Continued)*

Field Name	Field Description	Level of Output

State	State of the connection between the routing device and the cache server.
	• `Up`—The session is up and the validation route payload (VRP) database is synchronized with the RPKI cache server.
	• `Down`—The session is down, In other words, the client router cannot communicate effectively with this RPKI cache server because:
	  • TCP is down
	  • The RPKI cache server is unresponsive. In other words, the RPKI cache server does not send PDUs for a while
	• `Connect`—The connection is not established and the client router retries connecting to the RPKI cache server.
	• `Ex-Full`—The client router requests an immediate full update from the RPKI cache server by sending a “Reset Query”. This state indicates that the router is in the process of receiving cache response PDUs on this session from the cache server with validation records relevant to the query. The “End of Data” PDU indicates the end of responses from the server. When the router receives this PDU, the session moves to the Up state.
	• `Ex-Incr`—The client router requests an incremental update from the RPKI cache server by sending a “Serial Query”. This state indicates that the router is in the process of receiving cache response PDUs on this session from the cache server with validation records relevant to this query. The “End of Data” PDU indicates end of responses from the server. When the router receives this PDU, the session moves to the Up state.
		All levels

Flaps	Number of attempts to establish a session.	None and `brief`

Table 1749: show validation session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Uptime	Length of time that the session has remained established.	None and `brief`
#IPv4/IPv6 records	Number of IPv4 and IPv6 route validation records.	None and `brief`
Session index	Every session has an index number.	`detail`
Group	Name of the group to which the session belongs.	`detail`
Preference	Each cache server has a preference. Higher preferences are preferred. During a session start or restart, the routing device attempts to start a session with the cache server that has the numerically highest preference. The routing device connects to multiple cache servers in preference order. The default preference is 100. The preference is configurable with the `preference` statement.	`detail`
Port	TCP port number for the outgoing connection with the cache server. The well-known RPKI port is TCP port 2222. For a given deployment, an RPKI cache server might listen on some other TCP port number. If so, you can configure the alternative port number with the `port` statement.	`detail`
Refresh time	Liveliness check interval for an RPKI cache server. Every`refresh-time` (seconds), a serial query protocol data unit (PDU) with the last known serial number is transmitted. The `hold-time` must be at least 2 x the `refresh-time`.	`detail`

Table 1749: show validation session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Hold time	Length of time in seconds that the session between the routing device and the cache server is considered operational without any activity. After the hold time expires, the session is dropped. Receiving any PDU from the cache server resets the hold timer. The `hold-time` is 600 seconds, by default, and must be least 2 x the `refresh-time`. If the hold time expires, the session is considered to be down. This, in turn, triggers a session restart event. During a session restart, the routing device attempts to start a session with the cache server that has the numerically highest `preference`.	`detail`
Record Life time	Amount of time that route validation (RV) records learned from a cache server are valid. RV records expire if the session to the cache server goes down and remains down for the `record-lifetime` (seconds).	`detail`
Serial (Full Update)	Number of full serial updates.	`detail`
Serial (Incremental Update)	Number of incremental serial updates.	`detail`
Session flaps	Number of attempts to establish a session.	`detail`
Session uptime	Length of time that the session has remained established.	`detail`
Last PDU received	Time when the most recent PDU was received.	`detail`
IPv4 prefix count	Number of IPv4 sessions.	`detail`
IPv6 prefix count	Number of IPv6 sessions.	`detail`
Authentication key-chain	Configured TCP keychain	`detail`

Sample Output

show validation session brief

```
user@host> show validation session brief
 Session                                State   Flaps     Uptime #IPv4/IPv6 records
  1.3.0.2                                up          2   00:01:37 13/0
  10.255.255.11                          up          3   00:00:01 1/0
  10.255.255.12                          connect     2            64/68
```

show validation session detail

```
user@host> show validation session detail
Session 10.0.77.1, State: up
    Group: test, Preference: 100
    Local IPv4 address: 10.0.77.2, Port: 2222
    Refresh time: 300s
    Session flaps: 14, Last Session flap: 5h13m18s ago
    Hold time: 900s
    Record Life time: 3600s
    Serial (Full Update): 0
    Serial (Incremental Update): 0
    Authentication key-chain: new_auth_key
      Session flaps 2
      Session uptime: 00:48:35
      Last PDU received: 00:03:35
      IPv4 prefix count: 71234
      IPv6 prefix count: 345
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- Use Case and Benefit of Origin Validation for BGP
- Understanding Origin Validation for BGP
- Example: Configuring Origin Validation for BGP

show validation statistics

IN THIS SECTION

- Syntax | 27563
- Description | 27563
- Options | 27563
- Required Privilege Level | 27564
- Output Fields | 27564
- Sample Output | 27565
- Release Information | 27566

Syntax

```
show validation statistics
<instance instance-name>
<logical-system logical-system-name>
```

Description

Display route validation statistics.

Options

none	Display statistics for all routing instances.
instance *instance-name*	(Optional) Display information for the specified routing instance. The instance name can be primary for the main instance, or any valid configured instance name or its prefix.
logical-system *logical-system-name*	(Optional) Perform this operation on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1750 on page 27564 describes the output fields for the `show validation statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1750: show validation statistics Output Fields

Field Name	Field Description
Total RV records	Total number of route validation records.
Total Replication RV records	Total number of replicated route records.
Static records	Number of static records.
Prefix entries	The number of route prefixes.
Origin-AS entries	The number of route entries originated by the local AS.
Memory utilization	The amount of memory consumed by the RV process and its database.
Policy origin-validation requests	Number of queries for validation state of a given instance and prefix.
Valid	Number of valid prefixes reported by the validation query.
Invalid	Number of invalid prefixes reported by the validation query.
Unknown	Number of unknown prefixes reported by the validation query. This means that the prefix is not found in the database.
BGP import policy reevaluation notifications	A change, addition, or deletion of a route validation record triggers a BGP import reevaluation for all exact matching and more specific prefixes.

Table 1750: show validation statistics Output Fields *(Continued)*

Field Name	Field Description
inet.0	Number of IPv4 route validation records that have been added, deleted, or changed.
inet6.0	Number of IPv6 route validation records that have been added, deleted, or changed.
Count of VRP records	Total count of Validation Route Payload (VRP) entries for each address family.
Count of reevaluations	Count of VRP entries added for each address family.
Count of VRP records added	Count of VRP entries withdrawn from each address family.
Count of VRP records withdrawn	Count of reevaluations using this VRP table for per address family.

Sample Output

show validation statistics (with VRP statistics)

```
user@host> show validation statistics
Total RV records: 5644
Total Replication RV records: 443
  Prefix entries: 54
  Origin-AS entries: 2
Memory utilization: 1592 bytes
RV database: default
  RV records in Database: 57
  Origin-AS entries in Database: 32
  Database origin-validation re-evaluation statistics: 766
    Attempts resulting Valid: 32
    Attempts resulting Invalid: 56
    Attempts resulting Unknown: 76
  BGP import policy reevaluation notifications: 87
Policy origin-validation re-evaluation statistics: 566
```

```
  Attempts resulting Valid: 12
  Attempts resulting Invalid: 77
  Attempts resulting Unknown: 87
BGP import policy reevaluation notifications: 23
Count of VRP records: 32000
Count of reevaluations: 24000
Count of VRP records added: 88000
Count of VRP records withdrawn: 56000
```

show validation statistics

```
user@host> show validation statistics

Total RV records: 32000
Total Replication RV records: 32000
  Prefix entries: 8000
  Origin-AS entries: 8000
Memory utilization: 8550016 bytes
RV database: default
  RV records in Database: 32000
  Origin-AS entries in Database: 8000
  Database origin-validation re-evaluation statistics: 54000
    Attempts resulting Valid: 6000
    Attempts resulting Invalid: 12000
    Attempts resulting Unknown: 36000
  BGP import policy reevaluation notifications: 0
    inet.0, 0
    inet6.0, 0
Policy origin-validation re-evaluation statistics: 54000
  Attempts resulting Valid: 6000
  Attempts resulting Invalid: 12000
  Attempts resulting Unknown: 36000
BGP import policy reevaluation notifications: 0
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- Use Case and Benefit of Origin Validation for BGP
- Understanding Origin Validation for BGP
- Example: Configuring Origin Validation for BGP

show version (JDM)

IN THIS SECTION

- Syntax | 27567
- Description | 27567
- Options | 27567
- Required Privilege Level | 27568
- Sample Output | 27568
- Release Information | 27569

Syntax

```
show version (all-servers | server | vnf | all-vnfs | detail | brief)
```

Description

Display the version information about the Juniper Device Manager (JDM).

Options

all-servers Display the version details of the JDM instances on both the servers.

server Display the version details of the JDM instance on one specific server.

- **Range:** `0` through `1`

vnf Display the version details for a particular guest network function (GNF). You need to mention the GNF name in the command. Example: `show version vnf` *`gnf2`*.

all-vnfs Display the version details for all the GNFs.

(detail | brief) Display the specified level of output.

Required Privilege Level

View

Sample Output

show version

```
user@jdm> show version

Hostname: mgb-dvaita-ixr1-jdm
Model: junos_node_slicing
Server slot : 1
JDM package version : 17.4-R1.7
Host Software [Red Hat Enterprise Linux]
JDM container Software [Ubuntu 14.04.1 LTS]
JDM daemon jdmd [Version: 17.4R1.7-secure]
JDM daemon jinventoryd [Version: 17.4R1.7-secure]
JDM daemon jdmmon [Version: 17.4R1.7-secure]
Host daemon jlinkmon [Version: 17.4R1.7-secure]

user@jdm> show version brief

Hostname: mgb-dvaita-ixr1-jdm
Model: junos_node_slicing
Server slot : 1
JDM package version : 17.4-R1.7
Host Software [Red Hat Enterprise Linux]
JDM container Software [Ubuntu 14.04.1 LTS]
JDM daemon jdmd [Version: 17.4R1.7-secure]
JDM daemon jinventoryd [Version: 17.4R1.7-secure]
JDM daemon jdmmon [Version: 17.4R1.7-secure]
Host daemon jlinkmon [Version: 17.4R1.7-secure]

user@jdm> show version detail

Hostname: mgb-dvaita-ixr1-jdm
```

```
Model: junos_node_slicing
Server slot : 1
JDM package version : 17.4-R1.7
Host Software [Red Hat Enterprise Linux]
JDM container Software [Ubuntu 14.04.1 LTS]
JDM daemon jdmd [Version: 17.4R1.7-secure]
JDM daemon jinventoryd [Version: 17.4R1.7-secure]
JDM daemon jdmmon [Version: 17.4R1.7-secure]
Host daemon jlinkmon [Version: 17.4R1.7-secure]
KERNEL 3.10.0-514.el7.x86_64
MGD release 17.4R1.7 built by builder on 2017-11-17 11:29:41 UTC
CLI release 17.4R1.7 built by builder on 2017-11-17 10:53:44 UTC
base-actions-dd release 17.4R1.7 built by builder on 2017-11-17 10:06:17 UTC
jdmd_common-actions-dd release 17.4R1.7 built by builder on 2017-11-17 10:06:09 UTC
jdmd_nv_jdm-actions-dd release 17.4R1.7 built by builder on 2017-11-17 10:06:09 UTC
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | **15481**

request server authenticate-peer-server | **14901**

show version (Junos OS)

IN THIS SECTION

- Syntax | **27570**
- Syntax (EX Series Switches) | **27570**
- Syntax (MX Series Router) | **27570**
- Syntax (QFX Series) | **27570**
- Syntax (SRX Series) | **27571**

- Description | 27571
- Options | 27571
- Required Privilege Level | 27572
- Release Information | 27572

Syntax

```
show version
<brief | detail>
```

Syntax (EX Series Switches)

```
show version
<all-members>
<brief | detail>
<local>
<member member-id>
```

Syntax (MX Series Router)

```
show version
<brief | detail>
<all-members>
<local>
<member member-id>
```

Syntax (QFX Series)

```
show version
<brief | detail>
<component component-name | all>
```

Syntax (SRX Series)

```
show version
<brief | detail>
<node node-id | local | primary>
```

Description

Display the hostname and version information about the software running on the router or switch.

Beginning in Junos OS Release 13.3, the `show version` command output includes the `Junos` field that displays the Junos OS version, including any selective upgrade (JSU) packages, running on the device. This field provides a consistent means of identifying the Junos OS version, rather than extracting that information from the list of installed sub-packages.

Table 1751: Common Package Prefixes for Junos OS

Junos OS Package Prefix	Junos OS Architecture
jinstall-*	Junos OS for M Series, MX Series, and T Series, routers
junos-install-*	Junos OS based on an upgraded FreeBSD kernel instead of older versions of FreeBSD
junos-vmhost-install-*	Junos OS with upgraded FreeBSD on a VM Host

Options

none Display standard information about the hostname and version of the software running on the router or switch.

brief | detail (Optional) Display the specified level of output.

all-members (EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on all members of the Virtual Chassis configuration.

component all (QFabric systems only) (Optional) Display the host name and version information about the software running on all the components on the QFabric system.

component *component-name*	(QFabric systems only) (Optional) Display the host name and version information about the software running on a specific QFabric system component. Replace *component-name* with the name of the QFabric system component. The *component-name* can be the name of a diagnostics Routing Engine, Director group, fabric control Routing Engine, fabric manager Routing Engine, Interconnect device, or Node group.
local	(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on the local Virtual Chassis member.
member *member-id*	(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.
node (all | *node-name)*	(Optional) Display version information for the specified node or all nodes.
primary	(SRX Series only) Display the software version on the primary node.

Required Privilege Level

view

Release Information

Command introduced before Junos OS Release 7.4.

show version (Junos OS Evolved)

IN THIS SECTION

- Syntax | 27573
- Description | 27573
- Options | 27573
- Required Privilege Level | 27573

Sample Output | 27573
Release Information | 27575

Syntax

```
show version
<node (all | node-name)>
```

Description

Display the hostname and the version information about the software running on the router or switch.

The output for the `show version` command for Junos OS Evolved includes a `Junos` field that indicates the installation package name. From the prefix of this package name, you can decode which Junos OS Evolved architecture the device is running.

Options

none	Display standard information about the hostname and version of the software running on the router or switch.
node (all | *node-name*)	(Optional) Display version information for the specified node or all nodes.

Required Privilege Level

view

Sample Output

show version (with Third-party Applications Installed)

```
user@host-re0> show version
Hostname: host-re0
Model: ptx10008
Junos: 22.4R1.11-EVO
```

```
Yocto: 3.0.2
Linux Kernel: 5.2.60-yocto-standard-gae998d995
JUNOS-EVO OS 64-bit [junos-evo-install-ptx-x86-64-22.4R1.11-EVO]
External Software:
JET app acmeMonitor 1.2.3
JET app multi_app 1.1.1
JET app custom_logger 1.0.2
```

show version with a `jfirmware` package present on the device

```
user@host> show version
Hostname: host
Model: ptx10008
Junos: 24.4I20241214083306-EVO
Yocto: 3.0.2
Linux Kernel: 5.2.60-yocto-standard-gae998d995
JUNOS-EVO OS 64-bit [junos-evo-install-ptx-x86-64-24.4I20241214083306-EVO]
Additional Software:
jfirmware-junos-evo-install-ptx-x86-64-24.4I202410114083306-EVO 1.0.1
```

show version node re0 (on a Dual-Routing Engine Device)

```
{master}
user@host-re0> show version node re0
Hostname: host-re0
Model: ptx10008
Junos: 20.4R2.4-EVO
Yocto: 2.2.1
Linux Kernel: 4.8.28-WR2.2.1_standard-g4d37950
JUNOS-EVO OS 64-bit [junos-evo-install-ptx-x86-64-20.4R2.4-EVO]
```

show version node all

```
{master}
user@host-re0> show version node all
re0:
--------------------------------------------------------------------------
Hostname: host-re0
```

```
Model: ptx10008
Junos: 20.4R2.4-EVO
Yocto: 2.2.1
Linux Kernel: 4.8.28-WR2.2.1_standard-g4d37950
JUNOS-EVO OS 64-bit [junos-evo-install-ptx-x86-64-20.4R2.4-EVO]

re1:
--------------------------------------------------------------------------
Hostname: host-re1
Model: ptx10008
Junos: 20.4R2.4-EVO
Yocto: 2.2.1
Linux Kernel: 4.8.28-WR2.2.1_standard-g4d37950
JUNOS-EVO OS 64-bit [junos-evo-install-ptx-x86-64-20.4R2.4-EVO]
```

Release Information

`node` option introduced in Junos OS Evolved Release 18.3R1.

show version vnf

IN THIS SECTION

- Syntax | 27575
- Description | 27576
- Options | 27576
- Required Privilege Level | 27576
- Sample Output | 27576
- Release Information | 27578

Syntax

`show version vnf` *`vnf-name`*

Description

Display the hostname and version information about the specified guest network function (GNF).

Options

vnf-name Name of the GNF for which you want to view the version details.

Required Privilege Level

View

Sample Output

Depending on the platform running Junos OS, you might see different installed sub-packages.

show version vnf

```
user@jdm> show version vnf gnf2

Hostname: gnf2
Model: mx960
Junos: 17.4X48-D10.3
JUNOS OS Kernel 64-bit  [20170913.201739_fbsd-builder_stable_11]
JUNOS OS libs [20170913.201739_fbsd-builder_stable_11]
JUNOS OS runtime [20170913.201739_fbsd-builder_stable_11]
JUNOS OS time zone information [20170913.201739_fbsd-builder_stable_11]
JUNOS network stack and utilities [20170926.111120_builder_junos_174_x48_d10]
JUNOS modules [20170926.111120_builder_junos_174_x48_d10]
JUNOS mx modules [20170926.111120_builder_junos_174_x48_d10]
JUNOS libs [20170926.111120_builder_junos_174_x48_d10]
JUNOS OS libs compat32 [20170913.201739_fbsd-builder_stable_11]
JUNOS OS 32-bit compatibility [20170913.201739_fbsd-builder_stable_11]
JUNOS libs compat32 [20170926.111120_builder_junos_174_x48_d10]
JUNOS runtime [20170926.111120_builder_junos_174_x48_d10]
Junos vmguest package [20170926.111120_builder_junos_174_x48_d10]
JUNOS py extensions [20170926.111120_builder_junos_174_x48_d10]
JUNOS py base [20170926.111120_builder_junos_174_x48_d10]
JUNOS OS vmguest [20170913.201739_fbsd-builder_stable_11]
JUNOS OS crypto [20170913.201739_fbsd-builder_stable_11]
```

```
JUNOS mx libs compat32 [20170926.111120_builder_junos_174_x48_d10]
JUNOS mx runtime [20170926.111120_builder_junos_174_x48_d10]
JUNOS common platform support [20170926.111120_builder_junos_174_x48_d10]
JUNOS mx libs [20170926.111120_builder_junos_174_x48_d10]
JUNOS mtx Data Plane Crypto Support [20170926.111120_builder_junos_174_x48_d10]
JUNOS daemons [20170926.111120_builder_junos_174_x48_d10]
JUNOS mx daemons [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services URL Filter package [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services TLB Service PIC package [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services SSL [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services SOFTWIRE [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services Stateful Firewall [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services RPM [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services PTSP Container package [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services PCEF package [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services NAT [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services Mobile Subscriber Service Container package
[20170926.111120_builder_junos_174_x48_d10]
JUNOS Services MobileNext Software package [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services Logging Report Framework package [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services LL-PDF Container package [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services Jflow Container package [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services Deep Packet Inspection package [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services IPSec [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services IDS [20170926.111120_builder_junos_174_x48_d10]
JUNOS IDP Services [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services HTTP Content Management package [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services Crypto [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services Captive Portal and Content Delivery Container package
[20170926.111120_builder_junos_174_x48_d10]
JUNOS Services COS [20170926.111120_builder_junos_174_x48_d10]
JUNOS AppId Services [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services Application Level Gateways [20170926.111120_builder_junos_174_x48_d10]
JUNOS Services AACL Container package [20170926.111120_builder_junos_174_x48_d10]
JUNOS SDN Software Suite [20170926.111120_builder_junos_174_x48_d10]
JUNOS Extension Toolkit [20170926.111120_builder_junos_174_x48_d10]
JUNOS Packet Forwarding Engine Support (wrlinux) [20170926.111120_builder_junos_174_x48_d10]
JUNOS Packet Forwarding Engine Support (MX/EX92XX Common)
[20170926.111120_builder_junos_174_x48_d10]
JUNOS Packet Forwarding Engine Support (M/T Common) [20170926.111120_builder_junos_174_x48_d10]
JUNOS Packet Forwarding Engine Support (MX Common) [20170926.111120_builder_junos_174_x48_d10]
```

```
JUNOS jfirmware [20170926.111120_builder_junos_174_x48_d10]
JUNOS Online Documentation [20170926.111120_builder_junos_174_x48_d10]
```

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | 15481

request server authenticate-peer-server | 14901

show virtual-chassis

IN THIS SECTION

- Syntax | 27578
- Description | 27579
- Options | 27579
- Required Privilege Level | 27579
- Output Fields | 27579
- Sample Output | 27582
- Release Information | 27586

Syntax

```
show virtual-chassis <status>
```

Description

Display information about all members of the Virtual Chassis or VCF.

Options

none Display information about all Virtual Chassis or VCF member devices.

status Same output as for `show virtual-chassis` without any options.

Required Privilege Level

view

Output Fields

Table 1752 on page 27579 lists the output fields for the `show virtual-chassis` command. Output fields are listed in the approximate order in which they appear.

Table 1752: show virtual-chassis Output Fields

Field Name	Field Description
Fabric ID	(VCF only) Assigned ID used to identify the VCF.
Fabric Mode	(VCF only) Mode of the VCF: Enabled, Disabled, or Mixed.
Preprovisioned Virtual Chassis or **Preprovisioned Virtual Chassis Fabric**	Virtual Chassis or VCF is configured using preprovisioning.
Virtual Chassis ID	Assigned ID that applies to the entire Virtual Chassis or VCF.

Table 1752: show virtual-chassis Output Fields *(Continued)*

Field Name	Field Description
Virtual Chassis Mode	Mode of the Virtual Chassis or VCF. This field indicates support for the Virtual Chassis feature and, if a Virtual Chassis is configured, if it is a mixed or homogenous Virtual Chassis. Values can be: • `Enabled`—The platform supports the Virtual Chassis feature. If a Virtual Chassis is currently configured, this is a homogenous Virtual Chassis (all members are the same type of switch). • `Disabled`—The switch does not support the Virtual Chassis feature. **NOTE**: Switches that support the Virtual Chassis feature do not display this value. Even if a Virtual Chassis is not currently configured, those switches display `Enabled` in this field. • `Mixed`—The platform supports the Virtual Chassis feature, and is configured as a mixed mode Virtual Chassis (members consist of more than one type of switch).
Member ID	Assigned member ID and FPC: • On all EX Series Virtual Chassis except EX8200 Virtual Chassis, and on a VCF, the FPC number refers to the member ID assigned to the switch. • On EX8200 Virtual Chassis, member IDs are numbered 0 through 9. The FPC number indicates the slot number of the line card within the Virtual Chassis. The FPC number on member 0 is always 0 through 15. The FPC number on member 1 is always 16 through 31. The FPC number on member 2 is always 32 through 47; and so on for the members.
Status	For a nonprovisioned configuration: • `Prsnt` for a member that is currently connected to the Virtual Chassis or VCF configuration. • `NotPrsnt` for a member ID that has been assigned but is not currently connected. For a preprovisioned configuration: • `Prsnt` for a member that is specified in the preprovisioned configuration file and is currently connected to the Virtual Chassis or VCF. • `Unprvsnd` for a member that is interconnected with the Virtual Chassis or VCF configuration but is not specified in the preprovisioned configuration file.

Table 1752: show virtual-chassis Output Fields *(Continued)*

Field Name	Field Description
Serial No	Serial number of the member device.
Alias-Name	The user-configured alias of the member device. The **Alias-Name** field appears only if an alias has been configured for at least one device in the Virtual Chassis or VCF. Aliases are configured using the `alias-name` statement in the [`edit virtual-chassis aliases serial-number` *`serial-number`*] hierarchy.
Model	Model number of the member device.
Mastership Priority or **Mstr prio**	Primary-role priority value of the member device.
Role	Role of the member device: primary, backup, or linecard. An asterisk (*) following the Role denotes the member device on which the `show virtual-chassis <status>` command was issued.
Mixed Mode	Mixed mode configuration status: • `Y` for a member device configured in mixed mode. • `N` for a member device not configured in mixed mode. • `NA` for a member device that cannot be configured in mixed mode.
Route Mode	The route mode of the member device: • `F` for fabric (VCF) mode • `VC` for Virtual Chassis mode
Location	Location of the member device. If this field is empty, the location field was not set for the device.

Table 1752: show virtual-chassis Output Fields *(Continued)*

Field Name	Field Description
Neighbor List ID and **Interface**	Member ID of the neighbor member to which this member's Virtual Chassis port (VCP) is connected, and the VCP interface name.

Sample Output

show virtual-chassis (EX2300 multigigabit model—EX2300-24MP and EX2300-48MP—Virtual Chassis)

```
user@switch> show virtual-chassis
Virtual Chassis ID: cdef.789a.bc12
Virtual Chassis Mode: Enabled
                                          Mstr        Mixed Route  Neighbor List
Member ID  Status  Serial No  Model        prio Role    Mode  Mode  ID  Interface
0 (FPC 0)  Prsnt   ...        ex2300-24mp  128  Backup   N    VC   3  vcp-255/1/0
                                                                    1  vcp-255/1/2
1 (FPC 1)  Prsnt   ...        ex2300-48mp  128  Master*  N    VC   2  vcp-255/1/4
                                                                    0  vcp-255/1/1
2 (FPC 2)  Prsnt   ...        ex2300-24mp  128  Linecard N    VC   3  vcp-255/1/0
                                                                    1  vcp-255/1/1
3 (FPC 3)  Prsnt   ...        ex2300-48mp  128  Linecard N    VC   0  vcp-255/1/4
                                                                    2  vcp-255/1/5
```

show virtual-chassis (EX2300 and EX2300 multigigabit model (EX2300-24MP, EX2300-48MP) Virtual Chassis)

```
user@switch> show virtual-chassis status
Preprovisioned Virtual Chassis
Virtual Chassis ID: 9876.5432.abcd
Virtual Chassis Mode: Enabled
                                          Mstr           Mixed Route Neighbor List
Member ID  Status  Serial No  Model        prio  Role      Mode Mode  ID  Interface
0 (FPC 0)  Prsnt   ...        ex2300-24mp  129   Backup      N  VC    1 vcp-255/1/3
                                                                      3 vcp-255/1/0
1 (FPC 1)  Prsnt   ...        ex2300-24p   129   Master*     N  VC    2 vcp-255/1/0
```

```
                                                                    0 vcp-255/1/1
2 (FPC 2)  Prsnt   ...        ex2300-24p     0   Linecard    N  VC    3 vcp-255/1/2
                                                                    1 vcp-255/1/3
3 (FPC 3)  Prsnt   ...        ex2300-48mp    0   Linecard    N  VC    2 vcp-255/1/5
                                                                    0 vcp-255/1/0
```

show virtual-chassis (EX4200 Virtual Chassis)

```
user@switch> show virtual-chassis
Virtual Chassis ID: 00ab.cdef.1234
Virtual Chassis Mode: Enabled
                                        Mastership      Mixed  Neighbor List
Member ID  Status  Serial No  Model      priority   Role  Mode  ID  Interface
0 (FPC 0)  Prsnt   ...        ex4200-24t   249    Master*   N   8   vcp-0
                                                              1   vcp-1
1 (FPC 1)  Prsnt   ...        ex4200-24t   248    Backup    N   0   vcp-0
                                                              2   vcp-1
2 (FPC 2)  Prsnt   ...        ex4200-48p   247   Linecard   N   1   vcp-0
                                                              3   vcp-1
3 (FPC 3)  Prsnt   ...        ex4200-24t   246   Linecard   N   2   vcp-0
                                                              4   vcp-1
4 (FPC 4)  Prsnt   ...        ex4200-48p   245   Linecard   N   3   vcp-0
                                                              5   vcp-1
5 (FPC 5)  Prsnt   ...        ex4200-48t   244   Linecard   N   4   vcp-0
                                                              6   vcp-1
6 (FPC 6)  Prsnt   ...        ex4200-48t   243   Linecard   N   5   vcp-0
                                                              7   vcp-1
7 (FPC 7)  Prsnt   ...        ex4200-24f   242   Linecard   N   6   vcp-0
                                                              8   vcp-1
8 (FPC 8)  Prsnt   ...        ex4200-24f   241   Linecard   N   7   vcp-0
                                                              0   vcp-1
Member ID for next new member: 9 (FPC 9)
```

show virtual-chassis (Mixed EX4300 multigigabit model—EX4300-48MP—Virtual Chassis)

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis
Virtual Chassis ID: abcd.ef00.1234
Virtual Chassis Mode: Mixed
```

```
                                          Mstr         Mixed Route Neighbor List
Member ID  Status   Serial No  Model       prio  Role    Mode  Mode  ID Interface
0 (FPC 0)  Prsnt    ...        ex4300-48mp 129   Master*    Y  VC   1  vcp-255/1/0
                                                                     3  vcp-255/1/3
1 (FPC 1)  Prsnt    ...        ex4300-48mp 129   Backup     Y  VC   0  vcp-255/1/1
                                                                     2  vcp-255/1/3
2 (FPC 2)  Prsnt    ...        ex4300-48mp   0   Linecard   Y  VC   3  vcp-255/1/1
                                                                     1  vcp-255/1/3
3 (FPC 3)  Prsnt    ...        ex4300-48t    0   Linecard   Y  VC   0  vcp-255/1/0
                                                                     2  vcp-255/1/2
```

show virtual-chassis (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis
Virtual Chassis ID: cdcd.1212.efef
Virtual Chassis Mode: Enabled
                                       Mastership          Neighbor List
Member ID       Status Serial No  Model   priority  Role     ID Interface
0 (FPC 0-15)    Prsnt  ...        ex8216      0    Linecard  8  vcp-0/0
                                                             8  vcp-0/1
                                                             1  vcp-4/0/4
1 (FPC 16-31)   Prsnt  ...        ex8208      0    Linecard  8  vcp-0/0
                                                             0  vcp-3/0/4
8 (FPC 128-143) Prsnt  ...        ex-xre    128    Master    9  vcp-1/0
                                                             1  vcp-1/2
                                                             9  vcp-1/3
                                                             0  vcp-2/0
                                                             9  vcp-2/1
                                                             0  vcp-1/1
9 (FPC 144-159) Prsnt  ...        ex-xre    128    Backup*   8  vcp-1/0
                                                             8  vcp-1/2
                                                             8  vcp-1/3
                                                             8  vcp-1/3
```

show virtual-chassis (QFX5110 Virtual Chassis)

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis
Virtual Chassis ID: abab.1212.cdcd
```

```
Virtual Chassis Mode: Enabled
                                          Mstr         Mixed Route  Neighbor List
MemberID  Status Serial No  Model          prio  Role    Mode  Mode  ID Interface
0 (FPC 0) Prsnt  ...        qfx5110-32q    129   Backup     Y  VC    9 vcp-255/0/12
                                                                     1 vcp-255/0/16
1 (FPC 1) Prsn   ...        qfx5110-32q    129   Master*    Y  VC    2 vcp-255/0/16
                                                                     0 vcp-255/0/17
2 (FPC 2) Prsnt  ...        qfx5110-32q      0   Linecard   Y  VC    3 vcp-255/0/11
                                                                     1 vcp-255/0/14
3 (FPC 3) Prsnt  ...        qfx5110-32q      0   Linecard   Y  VC    4 vcp-255/0/14
                                                                     2 vcp-255/0/31
4 (FPC 4) Prsnt  ...        qfx5110-32q      0   Linecard   Y  VC    5 vcp-255/0/10
                                                                     3 vcp-255/0/5
5 (FPC 5) Prsnt  ...        qfx5100e-24q-2p  0   Linecard   Y  VC    6 vcp-255/0/12
                                                                     4 vcp-255/0/20
6 (FPC 6) Prsnt  ...        qfx5110-48s-4c   0   Linecard   Y  VC    7 vcp-255/0/10
                                                                     5 vcp-255/0/49
7 (FPC 7  Prsnt  ...        qfx5100e-48s-6q  0   Linecard   Y  VC    6 vcp-255/0/18
                                                                     8 vcp-255/0/31
8 (FPC 8) Prsnt  ...        qfx5110-48s-4c   0   Linecard   Y  VC    7 vcp-255/0/21
                                                                     9 vcp-255/0/49
9 (FPC 9) Prsnt  ...        qfx5110-32q      0   Linecard   Y  VC    8 vcp-255/0/13
                                                                     0 vcp-255/0/17
```

show virtual-chassis (QFX5200 Virtual Chassis)

```
user@switch> show virtual-chassis
Virtual Chassis ID: abab.1212.cdcd
Virtual Chassis Mode: Enabled
                 Serial                  Mstr           Mixed Route Neighbor List
Member ID Status    No   Model           prio  Role     Mode  Mode  ID Interface
0 (FPC 0) Prsnt   ...    qfx5200-32c-r   128  Master*   N     VC    1 vcp-255/0/19
                                                                     2 vcp-255/0/26
1 (FPC 1) Prsnt   ...    qfx5200-32c-r   128  Linecard  N     VC    0 vcp-255/0/20
                                                                     2 vcp-255/0/16
2 (FPC 2) Prsnt   ...    qfx5200-32c-32q 128  Backup    N     VC    0 vcp-255/0/15
                                                                     1 vcp-255/0/18
```

show virtual-chassis (QFX5100 Virtual Chassis Fabric)

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 0123.abcd.4567
Fabric Mode: Enabled
                   Serial                Mstr        Mixed Route Neighbor List
Member ID Status    No    Model          prio Role   Mode  Mode  ID Interface
0 (FPC 0) Prsnt   ...     qfx5100-48s    129  Master*  N    F     3  vcp-255/1/0
                                                                  2  vcp-255/1/1
                                                                  4  vcp-255/1/2
                                                                  4  vcp-255/1/3
1 (FPC 1) Prsnt   ...     qfx5100-48s    129  Backup   N    F     3  vcp-255/1/0
                                                                  2  vcp-255/1/1
                                                                  4  vcp-255/1/2
                                                                  4  vcp-255/1/3
2 (FPC 2) Prsnt   ...     qfx5100-48s      0  Linecard N    F     1  vcp-255/1/0
                                                                  0  vcp-255/1/1
3 (FPC 3) Prsnt   ...     qfx5100-48s      0  Linecard N    F     1  vcp-255/1/0
                                                                  0  vcp-255/1/1
4 (FPC 4) Prsnt   ...     qfx5100-48s      0  Linecard N    F     1  vcp-255/1/0
                                                                  0  vcp-255/1/1
```

Release Information

Command introduced in Junos OS Release 9.2.

`Fabric ID`, `Fabric Mode`, and `Route Mode` output fields introduced in Junos OS Release 13.2X51-D20.

`Alias-Name` output field introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

- show virtual-chassis active-topology | **27587**
- show virtual-chassis protocol adjacency | **27622**
- show virtual-chassis vc-path | **27676**
- Understanding Mixed EX Series and QFX Series Virtual Chassis
- *Understanding Mixed Virtual Chassis Fabric*
- Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

show virtual-chassis active-topology

IN THIS SECTION

- Syntax | 27587
- Description | 27587
- Options | 27587
- Required Privilege Level | 27588
- Output Fields | 27588
- Sample Output | 27588
- Release Information | 27593

Syntax

```
show virtual-chassis active-topology
<all-members | local | member member-id>
```

Description

Display the active topology of the Virtual Chassis or VCF with next-hop reachability information.

Options

none	Display the active topology of the member switch where you enter this command.
all-members	(Optional) Display the active topology of all members of the Virtual Chassis or VCF.
local	(Optional) Display the active topology of the switch or external Routing Engine where you enter this command.
member *member-id*	(Optional) Display the active topology of the specified member of the Virtual Chassis or VCF.

Required Privilege Level

view

Output Fields

Table 1753 on page 27588 lists the output fields for the `show virtual-chassis active-topology` command. Output fields are listed in the approximate order in which they appear.

Table 1753: show virtual-chassis active-topology Output Fields

Field Name	Field Description
Destination ID	Specifies the member ID of the destination.
Next-hop	Specifies the member ID and Virtual Chassis port (VCP) of the next hop to which packets for the destination ID are forwarded. The next hop can be more than one device in a VCF.

Sample Output

show virtual-chassis active-topology (EX4200 Virtual Chassis)

```
user@switch> show virtual-chassis active-topology

  1                   1(vcp-1)



  2                   1(vcp-1)



  3                   1(vcp-1)



  4                   1(vcp-1)
```

```
  5                    8(vcp-0)  1(vcp-1)


  6                    8(vcp-0)


  7                    8(vcp-0)


  8                    8(vcp-0)
```

show virtual-chassis active-topology (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis active-topology

 member0:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  1                    1(vcp-4/0/4.32768)

  8                    8(vcp-0/0.32768)

  9                    8(vcp-0/0.32768)

member1:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    0(vcp-3/0/4.32768)

  8                    8(vcp-0/0.32768)

  9                    8(vcp-0/0.32768)
```

```
member8:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    0(vcp-1/1.32768)

  1                    1(vcp-1/2.32768)

  9                    9(vcp-2/1.32768)

member9:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    8(vcp-1/2.32768)

  1                    8(vcp-1/2.32768)

  8                    8(vcp-1/2.32768)
```

show virtual-chassis active-topology (Virtual Chassis Fabric)

```
user@device> show virtual-chassis active-topology
fpc0:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  1                    4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  2                    4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  3                    4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  4                    4(vcp-255/0/2.32768)

  5                    5(vcp-255/0/3.32768)

  6                    6(vcp-255/0/1.32768)

fpc1:
```

```
-----------------------------------------------------------------------------
  Destination ID        Next-hop

  0                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  2                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  3                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  4                     4(vcp-255/0/2.32768)

  5                     5(vcp-255/0/3.32768)

  6                     6(vcp-255/0/1.32768)

fpc2:
-----------------------------------------------------------------------------
  Destination ID        Next-hop

  0                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  1                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  3                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  4                     4(vcp-255/0/2.32768)

  5                     5(vcp-255/0/3.32768)

  6                     6(vcp-255/0/1.32768)

fpc3:
-----------------------------------------------------------------------------
  Destination ID        Next-hop

  0                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  1                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  2                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  4                     4(vcp-255/0/2.32768)
```

```
  5                    5(vcp-255/0/3.32768)

  6                    6(vcp-255/0/1.32768)

fpc4:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    0(vcp-255/0/48.32768)

  1                    1(vcp-255/0/49.32768)

  2                    2(vcp-255/0/50.32768)

  3                    3(vcp-255/0/51.32768)

  5                    3(vcp-255/0/51.32768)  2(vcp-255/0/50.32768)  0(vcp-255/0/48.32768)
1(vcp-255/0/49.32768)

  6                    3(vcp-255/0/51.32768)  2(vcp-255/0/50.32768)  0(vcp-255/0/48.32768)
1(vcp-255/0/49.32768)

fpc5:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    0(vcp-255/0/48.32768)

  1                    1(vcp-255/0/49.32768)

  2                    2(vcp-255/0/50.32768)

  3                    3(vcp-255/0/51.32768)

  4                    3(vcp-255/0/51.32768)  2(vcp-255/0/50.32768)  0(vcp-255/0/48.32768)
1(vcp-255/0/49.32768)

  6                    3(vcp-255/0/51.32768)  2(vcp-255/0/50.32768)  0(vcp-255/0/48.32768)
1(vcp-255/0/49.32768)

fpc6:
--------------------------------------------------------------------------
  Destination ID        Next-hop
```

```
  0                    0(vcp-255/0/0.32768)

  1                    1(vcp-255/0/1.32768)

  2                    2(vcp-255/0/2.32768)

  3                    3(vcp-255/0/3.32768)

  4                    3(vcp-255/0/3.32768)  2(vcp-255/0/2.32768)  0(vcp-255/0/0.32768)
1(vcp-255/0/1.32768)

  5                    3(vcp-255/0/3.32768)  2(vcp-255/0/2.32768)  0(vcp-255/0/0.32768)
1(vcp-255/0/1.32768)
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

show virtual-chassis active-topology (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 27594
- Description | 27594
- Options | 27594
- Required Privilege Level | 27594
- Output Fields | 27594
- Sample Output | 27595
- Release Information | 27596

Syntax

```
show virtual-chassis active-topology
<(all-members | local | member member-id)>
```

Description

Display information about neighbor reachability from each member router in an MX Series Virtual Chassis configuration. You can issue the `show virtual-chassis active-topology` command from the console of either member router in the Virtual Chassis.

Options

all-members (Optional) Display neighbor reachability information for both member routers in a Virtual Chassis configuration. This is the default behavior if no options are specified.

local (Optional) Display neighbor reachability information for the member router on which you are issuing the command.

member *member-id* (Optional) Display neighbor reachability information for the specified Virtual Chassis member router. Replace ***member-id*** with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 1754 on page 27594 lists the output fields for the `show virtual-chassis active-topology` command. Output fields are listed in the approximate order in which they appear.

Table 1754: show virtual-chassis active-topology Output Fields

Field Name	Field Description
member*n*	Member ID of the Virtual Chassis member router for which output is displayed.
Destination ID	Member ID of the destination (neighbor) router.

Table 1754: show virtual-chassis active-topology Output Fields *(Continued)*

Field Name	Field Description
Next-hop	Member ID and Virtual Chassis port interface (in the format **vcp-*slot/pic/port.logical-unit-number***) of the next-hop to which the router forwards packets for the destination ID.

Sample Output

show virtual-chassis active-topology all-members

```
{master:member0-re0}
user@host> show virtual-chassis active-topology all-members
member0:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  1                    1(vcp-5/0/0.32768)

member1:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    0(vcp-1/3/0.32768)
```

show virtual-chassis active-topology local

```
{master:member0-re0}
user@host> show virtual-chassis active-topology local

  Destination ID        Next-hop

  1                    1(vcp-5/0/0.32768)
```

show virtual-chassis active-topology member 1

```
{master:member0-re0}
user@host> show virtual-chassis active-topology member 1
member1:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    0(vcp-1/3/0.32768)
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Verifying Neighbor Reachability for Member Routers or Switches in a Virtual Chassis

show virtual-chassis device-topology

IN THIS SECTION

- Syntax | 27597
- Description | 27597
- Options | 27597
- Required Privilege Level | 27597
- Output Fields | 27597
- Sample Output | 27598
- Release Information | 27605

Syntax

```
show virtual-chassis device-topology
<all-members | local | member member-id>
```

Description

Display the device topology—the member and system IDs, the VCP numbers, and device status—for all hardware devices in the Virtual Chassis or VCF.

Options

none	Display the device topology for all members of the Virtual Chassis or VCF.
all-members	(Optional) Display the device topology for all members of the Virtual Chassis or VCF.
local	(Optional) Display the device topology for the switch or external Routing Engine where you enter this command.
member ***member-id***	(Optional) Display the device topology for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 1755 on page 27597 lists the output fields for the `show virtual-chassis device-topology` command. Output fields are listed in the approximate order in which they appear.

Table 1755: show virtual-chassis device-topology Output Fields

Field Name	Field Description
Member	Assigned member ID.

Table 1755: show virtual-chassis device-topology Output Fields *(Continued)*

Field Name	Field Description
Device	Assigned device ID. For an EX8200 Virtual Chassis, the member ID and the device ID are always identical.
Status	The status of the device within the Virtual Chassis or VCF. Outputs include: • **Prsnt**—Device is currently connected to and participating in the Virtual Chassis or VCF. • **NotPrsnt**—Device is assigned but is not currently connected.
System ID	System ID of the device. The system ID of the device is the device's MAC address.
Member (Neighbor List)	Assigned member ID of the neighbor device.
Device (Neighbor List)	Assigned device ID of the neighbor device. For an EX8200 Virtual Chassis, the member ID and the device ID are always identical.
Interface (Neighbor List)	The interface connecting the device to the neighbor.

Sample Output

show virtual-chassis device-topology

```
user@switch> show virtual-chassis device-topology
member0:
--------------------------------------------------------------------------
                                                Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    0021.59f7.d000       8       8    vcp-0/0
                                                1       1    vcp-4/0/1
```

```
   1       1    Prsnt    0026.888d.6800       8       8    vcp-0/0
                                              9       9    vcp-0/1
                                              0       0    vcp-3/0/4
   8       8    Prsnt    0000.4a75.9b7c       9       9    vcp-1/0
                                              0       0    vcp-1/1
                                              1       1    vcp-1/2
   9       9    Prsnt    0000.73e9.9a57       8       8    vcp-1/0
                                              1       1    vcp-1/1

member1:
--------------------------------------------------------------------------
                                                 Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    0021.59f7.d000       8       8    vcp-0/0
                                              1       1    vcp-4/0/1
   1       1    Prsnt    0026.888d.6800       8       8    vcp-0/0
                                              9       9    vcp-0/1
                                              0       0    vcp-3/0/4
   8       8    Prsnt    0000.4a75.9b7c       9       9    vcp-1/0
                                              0       0    vcp-1/1
                                              1       1    vcp-1/2
   9       9    Prsnt    0000.73e9.9a57       8       8    vcp-1/0
                                              1       1    vcp-1/1

member8:
--------------------------------------------------------------------------
                                                 Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    0021.59f7.d000       8       8    vcp-0/0
                                              1       1    vcp-4/0/1
   1       1    Prsnt    0026.888d.6800       8       8    vcp-0/0
                                              9       9    vcp-0/1
                                              0       0    vcp-3/0/4
   8       8    Prsnt    0000.4a75.9b7c       9       9    vcp-1/0
                                              0       0    vcp-1/1
                                              1       1    vcp-1/2
   9       9    Prsnt    0000.73e9.9a57       8       8    vcp-1/0
                                              1       1    vcp-1/1

member9:
--------------------------------------------------------------------------
                                                 Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
```

```
   0       0    Prsnt    0021.59f7.d000       8       8    vcp-0/0
                                              1       1    vcp-4/0/1
   1       1    Prsnt    0026.888d.6800       8       8    vcp-0/0
                                              9       9    vcp-0/1
                                              0       0    vcp-3/0/4
   8       8    Prsnt    0000.4a75.9b7c       9       9    vcp-1/0
                                              0       0    vcp-1/1
                                              1       1    vcp-1/2
   9       9    Prsnt    0000.73e9.9a57       8       8    vcp-1/0
                                              1       1    vcp-1/1
```

show virtual-chassis device-topology (Virtual Chassis Fabric)

```
user@device> show virtual-chassis device-topology
fpc0:
--------------------------------------------------------------------------
                                                Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900       4       4    vcp-255/0/2
                                              5       5    vcp-255/0/3
                                              6       6    vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40       4       4    vcp-255/0/2
                                              5       5    vcp-255/0/3
                                              6       6    vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700       4       4    vcp-255/0/2
                                              5       5    vcp-255/0/3
                                              6       6    vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440       4       4    vcp-255/0/2
                                              5       5    vcp-255/0/3
                                              6       6    vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40       3       3    vcp-255/0/51
                                              2       2    vcp-255/0/50
                                              0       0    vcp-255/0/48
                                              1       1    vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0       3       3    vcp-255/0/51
                                              2       2    vcp-255/0/50
                                              1       1    vcp-255/0/49
                                              0       0    vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00       3       3    vcp-255/0/3
                                              2       2    vcp-255/0/2
                                              0       0    vcp-255/0/0
```

```
                                                 1       1    vcp-255/0/1

fpc1:
--------------------------------------------------------------------------
                                                  Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40       3       3    vcp-255/0/51
                                                 2       2    vcp-255/0/50
                                                 0       0    vcp-255/0/48
                                                 1       1    vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0       3       3    vcp-255/0/51
                                                 2       2    vcp-255/0/50
                                                 1       1    vcp-255/0/49
                                                 0       0    vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00       3       3    vcp-255/0/3
                                                 2       2    vcp-255/0/2
                                                 0       0    vcp-255/0/0
                                                 1       1    vcp-255/0/1

fpc2:
--------------------------------------------------------------------------
                                                  Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
```

```
                                                 6       6    vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440         4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40         3       3    vcp-255/0/51
                                                 2       2    vcp-255/0/50
                                                 0       0    vcp-255/0/48
                                                 1       1    vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0         3       3    vcp-255/0/51
                                                 2       2    vcp-255/0/50
                                                 1       1    vcp-255/0/49
                                                 0       0    vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00         3       3    vcp-255/0/3
                                                 2       2    vcp-255/0/2
                                                 0       0    vcp-255/0/0
                                                 1       1    vcp-255/0/1

fpc3:
--------------------------------------------------------------------------
                                                Neighbor List
Member  Device  Status   System ID          Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900         4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40         4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700         4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440         4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40         3       3    vcp-255/0/51
                                                 2       2    vcp-255/0/50
                                                 0       0    vcp-255/0/48
                                                 1       1    vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0         3       3    vcp-255/0/51
                                                 2       2    vcp-255/0/50
                                                 1       1    vcp-255/0/49
                                                 0       0    vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00         3       3    vcp-255/0/3
                                                 2       2    vcp-255/0/2
```

```
                                                 0       0    vcp-255/0/0
                                                 1       1    vcp-255/0/1

fpc4:
--------------------------------------------------------------------------
                                                   Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40       3       3    vcp-255/0/51
                                                 2       2    vcp-255/0/50
                                                 0       0    vcp-255/0/48
                                                 1       1    vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0       3       3    vcp-255/0/51
                                                 2       2    vcp-255/0/50
                                                 1       1    vcp-255/0/49
                                                 0       0    vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00       3       3    vcp-255/0/3
                                                 2       2    vcp-255/0/2
                                                 0       0    vcp-255/0/0
                                                 1       1    vcp-255/0/1

fpc5:
--------------------------------------------------------------------------
                                                   Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40       4       4    vcp-255/0/2
                                                 5       5    vcp-255/0/3
                                                 6       6    vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700       4       4    vcp-255/0/2
```

```
                                                5       5    vcp-255/0/3
                                                6       6    vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440         4       4    vcp-255/0/2
                                                5       5    vcp-255/0/3
                                                6       6    vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40         3       3    vcp-255/0/51
                                                2       2    vcp-255/0/50
                                                0       0    vcp-255/0/48
                                                1       1    vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0         3       3    vcp-255/0/51
                                                2       2    vcp-255/0/50
                                                1       1    vcp-255/0/49
                                                0       0    vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00         3       3    vcp-255/0/3
                                                2       2    vcp-255/0/2
                                                0       0    vcp-255/0/0
                                                1       1    vcp-255/0/1

fpc6:
--------------------------------------------------------------------------
                                                  Neighbor List
Member  Device  Status   System ID          Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900         4       4    vcp-255/0/2
                                                5       5    vcp-255/0/3
                                                6       6    vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40         4       4    vcp-255/0/2
                                                5       5    vcp-255/0/3
                                                6       6    vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700         4       4    vcp-255/0/2
                                                5       5    vcp-255/0/3
                                                6       6    vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440         4       4    vcp-255/0/2
                                                5       5    vcp-255/0/3
                                                6       6    vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40         3       3    vcp-255/0/51
                                                2       2    vcp-255/0/50
                                                0       0    vcp-255/0/48
                                                1       1    vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0         3       3    vcp-255/0/51
                                                2       2    vcp-255/0/50
                                                1       1    vcp-255/0/49
                                                0       0    vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00         3       3    vcp-255/0/3
```

```
                                        2       2     vcp-255/0/2
                                        0       0     vcp-255/0/0
                                        1       1     vcp-255/0/1
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Understanding Virtual Chassis Port Link Aggregation
- Understanding EX8200 Virtual Chassis Topologies

show virtual-chassis device-topology (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 27605
- Description | 27606
- Options | 27606
- Required Privilege Level | 27606
- Output Fields | 27606
- Sample Output | 27607
- Release Information | 27608

Syntax

```
show virtual-chassis device-topology
<(all-members | local | member member-id)>
```

Description

Display information about neighbor reachability for each hardware device in an MX Series Virtual Chassis configuration. On the MX Series router, there is only one active device for each member router. You can issue the `show virtual-chassis device-topology` command from the console of either member router in the Virtual Chassis.

Options

all-members (Optional) Display neighbor reachability information for the device in both member routers in a Virtual Chassis configuration.

local (Optional) Display neighbor reachability information for the device in the member router on which you are issuing the command. This is the default behavior if no options are specified.

member *member-id* (Optional) Display neighbor reachability information for the device in the specified Virtual Chassis member router. Replace ***member-id*** with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 1756 on page 27606 lists the output fields for the `show virtual-chassis device-topology` command. Output fields are listed in the approximate order in which they appear.

Table 1756: show virtual-chassis device-topology Output Fields

Field Name	Field Description
member*n*	Member ID of the Virtual Chassis member router for which output is displayed.
Member	Identifier assigned to the member router in the preprovisioned Virtual Chassis configuration.

Table 1756: show virtual-chassis device-topology Output Fields *(Continued)*

Field Name	Field Description
Device	Identifier assigned to the device in the member router. Because there is only one active device per member router in an MX Series Virtual Chassis configuration, the values in the Device and Member fields are identical.
Status	Status of the device: • **Prsnt**—Device is currently connected to the Virtual Chassis. • **NotPrsnt**—Device is not currently connected to the Virtual Chassis.
System ID	Unique identifier derived from the device's media access control (MAC) address. The System ID is included in each Virtual Chassis Control Protocol (VCCP) packet to identify the packet owner to all members of the Virtual Chassis.
Neighbor List Member/Device/Interface	Member IDs, Device IDs, and Virtual Chassis port interfaces (in the format **vcp-*slot/pic/port***) to which this device is connected.

Sample Output

show virtual-chassis device-topology all-members

```
{master:member0-re0}
user@host> show virtual-chassis device-topology all-members
member0:
--------------------------------------------------------------------------
                                                Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    b0c6.9abf.6800       1       1    vcp-5/0/0
   1       1    Prsnt    001d.b510.0800       0       0    vcp-1/3/0

member1:
--------------------------------------------------------------------------
                                                Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
```

```
   0       0    Prsnt    b0c6.9abf.6800       1       1    vcp-5/0/0
   1       1    Prsnt    001d.b510.0800       0       0    vcp-1/3/0
```

show virtual-chassis device-topology local

```
{master:member0-re0}
user@host> show virtual-chassis device-topology local

                                                Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    b0c6.9abf.6800       1       1    vcp-5/0/0
   1       1    Prsnt    001d.b510.0800       0       0    vcp-1/3/0
```

show virtual-chassis device-topology member 1

```
{master:member0-re0}
user@host> show virtual-chassis device-topology member 1
member1:
--------------------------------------------------------------------------
                                                Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    b0c6.9abf.6800       1       1    vcp-5/0/0
   1       1    Prsnt    001d.b510.0800       0       0    vcp-1/3/0
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

> *Verifying Neighbor Reachability for Hardware Devices in a Virtual Chassis*

show virtual-chassis fast-failover

IN THIS SECTION

- Syntax | 27609
- Description | 27609
- Required Privilege Level | 27609
- Output Fields | 27609
- Sample Output | 27610
- Release Information | 27610

Syntax

```
show virtual-chassis fast-failover
```

Description

Display information about the fast failover feature in a Virtual Chassis configuration.

Required Privilege Level

view

Output Fields

Table 1757 on page 27609 lists the output fields for the `show virtual-chassis fast-failover` command. Output fields are listed in the approximate order in which they appear.

Table 1757: show virtual-chassis fast-failover Output Fields

Field Name	Field Description
Fast failover on dedicated VCP ports	Indicates fast failover status on dedicated VCPs.

Table 1757: show virtual-chassis fast-failover Output Fields *(Continued)*

Field Name	Field Description
Fast failover on XE uplink VCP ports	Indicates fast failover status on XFP uplink VCPs.
Fast failover on GE uplink VCP ports	Indicates fast failover status on SFP uplink VCPs.

Sample Output

show virtual-chassis fast-failover

```
user@switch1> show virtual-chassis fast-failover
Fast failover on dedicated VCP ports: Enabled
Fast failover on XE uplink VCP ports: Disabled
Fast failover on GE uplink VCP ports: Enabled
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Example: Configuring Fast Failover on Uplink Module VCPs to Reroute Traffic When an EX4200 Virtual Chassis Switch or Intermember Link Fails

Configuring Fast Failover in an EX Series Virtual Chassis

Disabling Fast Failover in an EX Series Virtual Chassis

show virtual-chassis heartbeat (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 27611
- Description | 27611
- Options | 27611
- Required Privilege Level | 27612
- Output Fields | 27612
- Sample Output | 27614
- Release Information | 27616

Syntax

```
show virtual-chassis heartbeat
<(brief | detail)>
<(all-members | local | member member-id)>
```

Description

Display the state of one or both member routers in an MX Series Virtual Chassis configuration when an IP-based heartbeat packet connection is configured. You can issue the `show virtual-chassis heartbeat` command from the console of either member router in the Virtual Chassis.

Options

none Display the heartbeat state of both member routers in an MX Series Virtual Chassis.

brief | detail (Optional) Display the specified level of output. Using the `brief` option is equivalent to issuing the command with no options (the default). The `detail` option provides more output than the `brief` option.

all-members (Optional) Display the heartbeat state of both member routers in an MX Series Virtual Chassis. This is the default behavior if no options are specified.

local (Optional) Display the heartbeat state of the member router from which you are issuing the command.

member *member-id* (Optional) Display the heartbeat state of the specified member router. Replace *member-id* with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 1758 on page 27612 lists the output fields for the `show virtual-chassis heartbeat` command. Output fields are listed in the approximate order in which they appear.

Table 1758: show virtual-chassis heartbeat Output Fields

Field Name	Field Description	Level of Output
`member`*n*	Member ID of the Virtual Chassis member router for which output is displayed.	All levels
`Local`	IP address of the local member router that forms the heartbeat connection.	All levels
`Remote`	IP address of the remote member router that forms the heartbeat connection.	All levels

Table 1758: show virtual-chassis heartbeat Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`State`	State of the heartbeat connection: • `Alive`—Primary Routing Engine in the specified member router is connected and has received a heartbeat response message. • `Connected`—Primary Routing Engine in the specified member router is connected and awaiting a heartbeat response message. • `Conn-Pending`—Primary Routing Engine in the specified member router is connecting to the other member router. • `Detected`—Primary Routing Engine in the specified member router has successfully exchanged a heartbeat connection message with the other member router when an adjacency disruption or split occurs in the Virtual Chassis. The `Detected` state persists until the heartbeat connection is reset, or until the Virtual Chassis forms again and a primary router (protocol primary) and backup router (protocol backup) are elected. • `Idle`—Heartbeat connection is enabled but not yet connected. • `Inactive`—Heartbeat connection is not active. To activate the connection, you must configure the heartbeat address and, optionally, the heartbeat timeout value.	All levels
`Time`	Date and time, in the format *yyyy-mm-dd hh:mm:ss*, of the last connection state change. If you have cleared or never collected timestamps, this field displays `Never` and the time interval displays `0` (zero).	All levels

Table 1758: show virtual-chassis heartbeat Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Heartbeat Statistics`	Detailed statistics for the heartbeat connection: • `Heartbeats sent`—Number of heartbeat request messages sent. • `Heartbeats received`—Number of heartbeat response messages received. • `Heartbeats lost/missed`—Calculated difference between `Heartbeats sent` value and `Heartbeats received` value. • `Last time sent`—Date and time, in the format *yyyy-mm-dd hh*:*mm*:*ss*, and time since, in the format (*hh*:*mm*:*ss* `ago`), the last heartbeat request message was sent. If you have cleared or never collected timestamps, this field displays `Never` and the time interval displays `0` (zero). • `Last time received`—Date and time, in the format *yyyy-mm-dd hh*:*mm*:*ss*, and time since, in the format (*hh*:*mm*:*ss* `ago`), the last heartbeat request response message was received. If you have cleared or never collected timestamps, this field displays `Never` and the time interval displays `0` (zero). • `Maximum latency`—Maximum number of seconds that elapse on the local member router between transmission of a heartbeat request message and receipt of a heartbeat response message. • `Minimum latency`—Minimum number of seconds that elapse on the local member router between transmission of a heartbeat request message and receipt of a heartbeat response message.	`detail`

Sample Output

show virtual-chassis heartbeat

```
{master:member0-re0}
user@host> show virtual-chassis heartbeat
member0:
--------------------------------------------------------------------------
Local        Remote          State       Time
```

```
10.4.2.210   10.4.2.100      Alive       2014-01-14 14:46:52 PDT

member1:
--------------------------------------------------------------------------
Local        Remote          State       Time
10.4.2.100   10.4.2.210      Alive       2014-01-14 14:46:53 PDT
```

show virtual-chassis heartbeat detail

```
{master:member0-re0}
user@host> show virtual-chassis heartbeat detail
member0:
--------------------------------------------------------------------------
Local        Remote          State       Time
10.4.2.210   10.4.2.100      Alive       2014-01-14 14:46:52 PDT

Heartbeat statistics
    Heartbeats sent: 8812
    Heartbeats received: 8806
    Heartbeats lost/missed: 6
    Last time sent: 2014-01-14 14:28:00 PST (00:00:01 ago)
    Last time received: 2014-01-14 14:28:00 PST (00:00:01 ago)
    Maximum latency (secs): 1
    Minimum latency (secs): 0

member1:
--------------------------------------------------------------------------
Local        Remote          State       Time
10.4.2.100   10.4.2.210      Alive       2014-01-14 14:46:53 PDT

Heartbeat statistics
    Heartbeats sent: 8812
    Heartbeats received: 8806
    Heartbeats lost/missed: 6
    Last time sent: 2014-01-14 14:28:00 PST (00:00:01 ago)
    Maximum latency (secs): 1
    Minimum latency (secs): 0
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Verifying and Managing the Virtual Chassis Heartbeat Connection

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in the Same Subnet

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in Different Subnets

show virtual-chassis login

IN THIS SECTION

- Syntax | 27616
- Description | 27616
- Required Privilege Level | 27617
- Sample Output | 27617
- Release Information | 27617

Syntax

```
show virtual-chassis login
```

Description

Supply the address of the host that logged into the Virtual Chassis or VCF, or identify the location of the member switch that redirected the current session to a different member switch.

You might need this information for tracing or troubleshooting purposes.

Required Privilege Level

view

Sample Output

show virtual-chassis login (Direct Login to the Primary Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host 192.0.2.3
```

show virtual-chassis login (Backup Console Session Redirected to the Primary Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host backup
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

request session member | 15011

Understanding Global Management of a Virtual Chassis

show virtual-chassis mode

IN THIS SECTION

- Syntax | 27618
- Description | 27618
- Options | 27618

- Required Privilege Level | 27618
- Output Fields | 27619
- Sample Output | 27620
- Sample Output | 27620
- Sample Output | 27621
- Release Information | 27621

Syntax

```
show virtual-chassis mode
<all-members | local | member member-id>
```

Description

Display the Virtual Chassis or Virtual Chassis Fabric (VCF) mixed mode status.

Options

none	Display the Virtual Chassis or VCF mixed mode status for the device where you enter the command.
all-members	(Optional) Display the Virtual Chassis or VCF mixed mode status for all member devices in the Virtual Chassis or VCF.
local	(Optional) Display the Virtual Chassis or VCF mixed mode status for the device where you enter the command.
member ***member-id***	(Optional) Display the Virtual Chassis or VCF mixed mode status for the specified member device.

Required Privilege Level

view

Output Fields

Table 1759 on page 27619 lists the output fields for the `show virtual-chassis mode` command.

Table 1759: show virtual-chassis mode Output Fields

Field Name	Field Description
Mixed Mode	Specifies the mixed mode status of the member switch. Mixed mode is either **Enabled** or **Disabled**.
Current mode	Shows the current mixed and fabric mode settings running on the member device or devices. (EX4400 switches only) Shows whether the switch has network port mode enabled to use the default Virtual Chassis ports (VCPs) as network ports instead of as VCPs. A device reboot is required to change the device mode. The **Current mode** and **Future mode after reboot** are different when the mode has been changed but the device has not been rebooted. Outputs include: • **Fabric with mixed devices**—Fabric mode and mixed mode are enabled. • **Fabric with similar devices**—Fabric mode is enabled and mixed mode is disabled. • **Virtual Chassis with mixed devices**—Fabric mode is disabled and mixed mode is enabled. • **Virtual Chassis with similar devices**—Fabric mode is disabled and mixed mode is disabled. • (EX4400 switches only) **network-port mode**—Network port mode is enabled, which means the default VCPs on the switch function as network ports rather than as VCPs.

Table 1759: show virtual-chassis mode Output Fields *(Continued)*

Field Name	Field Description
Future mode after reboot	Shows the mixed and fabric mode settings that you have requested to be enabled on the member device or devices the next time the device is rebooted. A device reboot is required to change the fabric or mixed mode. The **Current mode** and **Future mode after reboot** are different only when the mode has been changed but the device has not been rebooted. Outputs include: • **Fabric with mixed devices**—Fabric mode and mixed mode are enabled. • **Fabric with similar devices**—Fabric mode is enabled and mixed mode is disabled. • **Virtual Chassis with mixed devices**—Fabric mode is disabled and mixed mode is enabled. • **Virtual Chassis with similar devices**—Fabric mode is disabled and mixed mode is disabled. • (EX4400 switches only) **network-port mode**—Network port mode is enabled so the default VCPs on the switch will function as network ports rather than as VCPs.

Sample Output

show virtual-chassis mode (EX4200 in a non-mixed Virtual Chassis)

```
user@switch> show virtual-chassis mode

fpc0:
--------------------------------------------------------------------------
Mixed Mode: Disabled
```

Sample Output

show virtual-chassis mode (EX4400 with network port mode enabled)

```
user@switch> show virtual-chassis mode

fpc0:
--------------------------------------------------------------------------
```

```
Current mode : network-port mode
Future mode after reboot: network-port mode
```

Sample Output

show virtual-chassis mode (QFX5100 in a non-mixed VCF)

```
user@switch> show virtual-chassis mode

fpc0:
--------------------------------------------------------------------------
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc1:
--------------------------------------------------------------------------
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc2:
--------------------------------------------------------------------------
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc3:
--------------------------------------------------------------------------
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc4:
--------------------------------------------------------------------------
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices
```

Release Information

Command introduced in Junos OS Release 11.1.

Current mode and **Future mode after reboot** fields introduced in Junos OS Release 13.2X51-D20.

RELATED DOCUMENTATION

request virtual-chassis mode | 15455

Verifying the Virtual Chassis Fabric Mode Settings

Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)

show virtual-chassis protocol adjacency

IN THIS SECTION

- Syntax | 27622
- Description | 27622
- Options | 27623
- Required Privilege Level | 27623
- Output Fields | 27623
- Sample Output | 27624
- Release Information | 27627

Syntax

```
show virtual-chassis protocol adjacency
<brief | detail | extensive>
<all-members | local | member member-id>
<system-id>
```

Description

Display the Virtual Chassis Control Protocol (VCCP) adjacency statistics in the Virtual Chassis or VCF for all hardware devices.

Options

none	Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.
brief | detail | extensive	(Optional) Display the specified level of output. Using the `brief` option is equivalent to entering the command with no options (the default). The `detail` and `extensive` options provide identical displays.
all-members	(Optional) Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.
local	(Optional) Display VCCP adjacency statistics for the switch or external Routing Engine on which this command is entered.
member *member-id*	(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.
system-id	(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 1760 on page 27623 lists the output fields for the `show virtual-chassis protocol adjacency` command. Output fields are listed in the approximate order in which they appear.

Table 1760: show virtual-chassis protocol adjacency Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the Virtual Chassis port (VCP) interface.	All levels
`System`	The MAC address of the device on the receiving side of the VCP link.	All levels

Table 1760: show virtual-chassis protocol adjacency Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the link. Outputs include: • Up—The link is up. • Down—The link is down. • New—The link is new. • One-way—The link is transmitting traffic in one direction. • Initializing—The link is initializing. • Rejected—The link is rejected.	All levels
Hold, Expires in	Remaining holdtime of the adjacency.	All levels
Priority	Priority to become the designated intermediary system.	detail
Up/Down Transitions	Count of adjacency status transition changes from up to down or down to up.	detail
Last transition	Time of the last up/down transition.	detail

Sample Output

show virtual-chassis protocol adjacency

```
user@switch> show virtual-chassis protocol adjacency
member0:
--------------------------------------------------------------------------
Interface             System         State        Hold (secs)
vcp-0/0.32768         0000.4a75.9b7c Up                 57
vcp-0/1.32768         0000.4a75.9b7c Up                 59
vcp-4/0/1.32768       0026.888d.6800 Up                 57

member1:
```

```
---------------------------------------------------------------------------
Interface             System         State        Hold (secs)
vcp-0/0.32768         0000.4a75.9b7c Up                  58
vcp-0/1.32768         0000.73e9.9a57 Up                  59
vcp-3/0/4.32768       0021.59f7.d000 Up                  58

member8:
---------------------------------------------------------------------------
Interface             System         State        Hold (secs)
vcp-1/0.32768         0000.73e9.9a57 Up                  58
vcp-1/1.32768         0021.59f7.d000 Up                  58
vcp-1/2.32768         0026.888d.6800 Up                  59
vcp-2/0.32768         0021.59f7.d000 Up                  59

member9:
---------------------------------------------------------------------------
Interface             System         State        Hold (secs)
vcp-1/0.32768         0000.4a75.9b7c Up                  58
vcp-1/1.32768         0026.888d.6800 Up                  59
```

show virtual-chassis protocol adjacency detail

```
user@switch> show virtual-chassis protocol adjacency detail
member0:
---------------------------------------------------------------------------

0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago

0000.4a75.9b7c
  interface-name: vcp-0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago

0026.888d.6800
  interface-name: vcp-4/0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago

member1:
---------------------------------------------------------------------------
```

```
0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

0000.73e9.9a57
  interface-name: vcp-0/1.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago

0021.59f7.d000
  interface-name: vcp-3/0/4.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago

member8:
--------------------------------------------------------------------------

0000.73e9.9a57
  interface-name: vcp-1/0.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

0021.59f7.d000
  interface-name: vcp-1/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

0026.888d.6800
  interface-name: vcp-1/2.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

0021.59f7.d000
  interface-name: vcp-2/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

member9:
--------------------------------------------------------------------------

0000.4a75.9b7c
  interface-name: vcp-1/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

0026.888d.6800
  interface-name: vcp-1/1.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Understanding Virtual Chassis Port Link Aggregation
- Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis

show virtual-chassis protocol adjacency (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 27627
- Description | 27628
- Options | 27628
- Required Privilege Level | 27628
- Output Fields | 27628
- Sample Output | 27631
- Release Information | 27632

Syntax

```
show virtual-chassis protocol adjacency
<(brief | detail | extensive)>
<(all-members | local | member member-id)>
<system-id>
```

Description

Display the entries (neighbors) in the Virtual Chassis Control Protocol (VCCP) adjacency database for an MX Series Virtual Chassis configuration. You can issue the `show virtual-chassis protocol adjacency` command from the console of either member router in the Virtual Chassis.

Options

brief | detail | extensive (Optional) Display the specified level of output. Using the `brief` option is equivalent to issuing the command with no options (the default). The `detail` option provides more output than the `brief` option. The `extensive` option provides complete output and is most useful for customer support personnel.

all-members (Optional) Display the VCCP adjacency database for both member routers in a Virtual Chassis. This is the default behavior if no options are specified.

local (Optional) Display the VCCP adjacency database for the member router on which you are issuing the command.

member *member-id* (Optional) Display the VCCP adjacency database for the specified member router. Replace *`member-id`* with the value 0 or 1.

system-id (Optional) Display the VCCP adjacency database for the device with the specified system ID.

Required Privilege Level

view

Output Fields

Table 1761 on page 27628 lists the output fields for the `show virtual-chassis protocol adjacency` command. Output fields are listed in the approximate order in which they appear.

Table 1761: show virtual-chassis protocol adjacency Output Fields

Field Name	Field Description	Level of Output
`member`*n*	Member ID of the Virtual Chassis member router for which output is displayed.	All levels

Table 1761: show virtual-chassis protocol adjacency Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface`	Name of the Virtual Chassis port interface, in the format `vcp-`*`slot`*`/`*`pic`*`/`*`port`*`.`*`logical-unit-number`*.	`brief`
`System`	System ID of the device associated with the Virtual Chassis port interface. The System ID is derived from the device's media access control (MAC) address.	`brief`
`State`	State of the adjacency: • `Up`—The adjacency is up. • `Down`—The adjacency is down.	All levels
`Hold (secs)`	Remaining hold time of the adjacency, in seconds.	`brief`
`system-id`	System ID of the device associated with the Virtual Chassis port interface. The System ID is derived from the device's media access control (MAC) address.	`detail`, `extensive`
`interface-name`	Name of the Virtual Chassis port interface, in the format `vcp-`*`slot`*`/`*`pic`*`/`*`port`*`.`*`logical-unit-number`*.	`detail`, `extensive`
`Expires in`	Number of seconds until the adjacency expires.	`detail`, `extensive`
`Priority`	Priority to become the designated intermediate system.	`detail`, `extensive`
`Up/Down transitions`	Count of adjacency status changes from `Up` to `Down` or from `Down` to `Up`.	`detail`, `extensive`
`Last transition`	Time of the last `Up/Down` transition.	`detail`, `extensive`

Table 1761: show virtual-chassis protocol adjacency Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Transition log`	List of recent adjacency transitions, including:
	• `When`—Date and time at which a VCCP adjacency transition occurred.
	• `State`—Current state of the VCCP adjacency:
	  • `Up`—Adjacency is up and operational.
	  • `Down`—Adjacency is down and not available.
	  • `Rejected`—Adjacency has been rejected.
	• `Event`—Type of transition that occurred:
	  • `Seenself`—Possible routing loop has been detected.
	  • `Interface down`—Virtual Chassis port interface has gone down and is no longer available.
	  • `Error`—Adjacency error.
	• `Down reason`—Reason that a VCCP adjacency is down:
	  • `3-Way Handshake Failed`—Connection establishment failed.
	  • `Address Mismatch`—Address mismatch caused link failure.
	  • `Aged Out`—Link expired.
	  • `ISO Area Mismatch`—VCCP area mismatch caused link failure.
	  • `Bad Hello`—Unacceptable hello message caused link failure.
	  • `BFD Session Down`—Bidirectional failure detection caused link failure.
	  • `Interface Disabled`—Virtual Chassis port interface is disabled.
	  • `Interface Down`—Virtual Chassis port interface is unavailable.
	  • `Interface Level Disabled`—VCCP level is disabled.
		`extensive`

Table 1761: show virtual-chassis protocol adjacency Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Level Changed`—VCCP level has changed on the adjacency. • `Level Mismatch`—Levels on adjacency are not compatible. • `MPLS LSP Down`—Label-switched path (LSP) is unavailable. • `MT Topology Changed`—VCCP topology has changed. • `MT Topology Mismatch`—VCCP topology is mismatched. • `Remote System ID Changed`—Adjacency peer system ID changed. • `Protocol Shutdown`—VCCP is disabled. • `CLI Command`—Adjacency brought down by user. • `Unknown`—Unknown.	

Sample Output

show virtual-chassis protocol adjacency all-members brief

```
{master:member0-re0}
user@host> show virtual-chassis protocol adjacency all-members brief
member0:
--------------------------------------------------------------------------
Interface             System         State        Hold (secs)
vcp-5/0/0.32768       001d.b510.0800 Up                 57

member1:
--------------------------------------------------------------------------
Interface             System         State        Hold (secs)
vcp-1/3/0.32768       b0c6.9abf.6800 Up                 58
```

show virtual-chassis protocol adjacency member 0 detail

```
{master:member0-re0}
user@host> show virtual-chassis protocol adjacency member 0 detail
member0:
--------------------------------------------------------------------------

001d.b510.0800
  interface-name: vcp-5/0/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 18:50:41 ago
```

show virtual-chassis protocol adjacency member 0 detail 001d.b510.0800

```
{master:member0-re0}
user@host> show virtual-chassis protocol adjacency member 0 detail 001d.b510.0800
member0:
--------------------------------------------------------------------------

001d.b510.0800
  interface-name: vcp-5/0/0.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 18:52:08 ago
```

show virtual-chassis protocol adjacency local extensive

```
{master:member0-re0}
user@host> show virtual-chassis protocol adjacency local extensive


001d.b510.0800
  interface-name: vcp-5/0/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 18:52:40 ago
  Transition log:
  When                  State        Event           Down reason
  Mon Sep 20 17:26:44   Up           Seenself
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Viewing Information in the Virtual Chassis Control Protocol Adjacency Database

show virtual-chassis protocol database

IN THIS SECTION

- Syntax | 27633
- Description | 27633
- Options | 27633
- Required Privilege Level | 27634
- Output Fields | 27634
- Sample Output | 27635
- Release Information | 27638

Syntax

```
show virtual-chassis protocol database
<brief | detail | extensive>
<all-members | local | member member-id>
```

Description

Display the Virtual Chassis Control Protocol (VCCP) database statistics for all hardware devices within the Virtual Chassis or VCF.

Options

none Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.

brief | detail | extensive (Optional) Display the specified level of output. Using the `brief` option is equivalent to entering the command with no options (the default). The `detail` option provides more output than the `brief` option. The `extensive` option provides all output and is most useful for customer support personnel.

all-members (Optional) Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.

local (Optional) Display VCCP database statistics for the switch or external Routing Engine where you enter this command.

member *member-id* (Optional) Display VCCP database statistics for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 1762 on page 27634 lists the output fields for the `show virtual-chassis protocol database` command. Output fields are listed in the approximate order in which they appear.

Table 1762: show virtual-chassis protocol database Output Fields

Field Name	Field Description	Level of Output
`LSP ID`	Link-state protocol (LSP) data unit identifier.	All levels
`Sequence`	Sequence number of the LSP.	All levels
`Checksum`	Checksum value of the LSP.	All levels
`Lifetime`	Remaining lifetime of the LSP, in seconds.	All levels
`Neighbor`	MAC address of the neighbor on the advertising system.	detail
`Interface`	Virtual Chassis port (VCP) interface name.	detail

Table 1762: show virtual-chassis protocol database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Metric`	Metric of the prefix or neighbor.	detail

The `extensive` output was omitted from this list. The `extensive` output is useful for customer support personnel only.

Sample Output

show virtual-chassis protocol database

```
user@switch> show virtual-chassis protocol database
member0:
--------------------------------------------------------------------------
LSP ID                      Sequence Checksum Lifetime
0000.4a75.9b7c.00-00         0x1dd80   0xc2e3      116
0000.73e9.9a57.00-00          0xf361   0x27e8      113
0021.59f7.d000.00-00         0x16882   0x3993      118
0026.888d.6800.00-00         0x1691f   0x82b7      116
  4 LSPs

member1:
--------------------------------------------------------------------------
LSP ID                      Sequence Checksum Lifetime
0000.4a75.9b7c.00-00         0x1dd80   0xc2e3      116
0000.73e9.9a57.00-00          0xf361   0x27e8      114
0021.59f7.d000.00-00         0x16883    0x289      116
0026.888d.6800.00-00         0x1691f   0x82b7      118
  4 LSPs

member8:
--------------------------------------------------------------------------
LSP ID                      Sequence Checksum Lifetime
0000.4a75.9b7c.00-00         0x1dd80   0xc2e3      118
0000.73e9.9a57.00-00          0xf361   0x27e8      114
0021.59f7.d000.00-00         0x16883    0x289      116
0026.888d.6800.00-00         0x16920   0xa335      116
  4 LSPs
```

```
member9:
--------------------------------------------------------------------------
LSP ID                      Sequence Checksum Lifetime
0000.4a75.9b7c.00-00         0x1dd80   0xc2e3      116
0000.73e9.9a57.00-00          0xf361   0x27e8      116
0021.59f7.d000.00-00         0x16883    0x289      114
0026.888d.6800.00-00         0x16920   0xa335      116
  4 LSPs
```

show virtual-chassis protocol database detail

```
user@switch> show virtual-chassis protocol database detail
member0:
--------------------------------------------------------------------------

0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-1/1.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/2.32768  Metric:     150

0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 114 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/1.32768  Metric:     150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 118 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-4/0/1.32768 Metric:      15

0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-0/1.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-3/0/4.32768 Metric:      15

member1:
--------------------------------------------------------------------------

0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-1/1.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/2.32768  Metric:     150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 116 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/1.32768  Metric:     150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-4/0/1.32768 Metric:      15

0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 117 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-0/1.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-3/0/4.32768 Metric:      15

member8:
------------------------------------------------------------------------

0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 118 secs
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-1/1.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/2.32768  Metric:     150

0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 115 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/1.32768  Metric:     150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-4/0/1.32768 Metric:      15

0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-0/1.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-3/0/4.32768 Metric:      15

member9:
------------------------------------------------------------------------

0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 116 secs
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-1/1.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/2.32768  Metric:     150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 117 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/1.32768  Metric:     150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 113 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-4/0/1.32768 Metric:      15

0026.888d.6800.00-00 Sequence: 0x1694f, Checksum: 0xa61a, Lifetime: 116 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-0/1.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-3/0/4.32768 Metric:      15
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis
- Understanding Virtual Chassis Components

show virtual-chassis protocol database (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 27639
- Description | 27639
- Options | 27639
- Required Privilege Level | 27639
- Output Fields | 27640
- Sample Output | 27643
- Release Information | 27646

Syntax

```
show virtual-chassis protocol database
<(brief | detail | extensive)>
<(all-members | local | member member-id)>
<system-id>
```

Description

Display the entries in the Virtual Chassis Control Protocol (VCCP) link-state database for an MX Series Virtual Chassis configuration. The VCCP link-state database contains information about protocol data unit (PDU) packets. You can issue the `show virtual-chassis protocol database` command from the console of either member router in the Virtual Chassis.

Options

brief | detail | extensive	(Optional) Display the specified level of output. Using the **brief** option is equivalent to issuing the command with no options (the default). The **detail** option provides more output than the **brief** option. The **extensive** option provides complete output and is most useful for customer support personnel.
all-members	(Optional) Display the VCCP link-state database for both member routers in a Virtual Chassis. This is the default behavior if no options are specified.
local	(Optional) Display the VCCP link-state database for the member router on which you are issuing the command.
member ***member-id***	(Optional) Display the VCCP link-state database for the specified member router. Replace ***member-id*** with the value 0 or 1.
system-id	(Optional) Display the VCCP link-state database for the neighbor with the specified system ID.

Required Privilege Level

view

Output Fields

Table 1763 on page 27640 lists the output fields for the `show virtual-chassis protocol database` command. Output fields are listed in the approximate order in which they appear.

Table 1763: show virtual-chassis protocol database Output Fields

Field Name	Field Description	Level of Output
member*n*	Member ID of the Virtual Chassis member router for which output is displayed.	All levels
LSP ID	Link-state PDU (LSP) identifier.	All levels
Sequence	Sequence number of the link-state PDU.	All levels
Checksum	Checksum value of the link-state PDU.	All levels
Lifetime	Remaining lifetime of the link-state PDU, in seconds.	All levels
***number* LSPs**	Total number of link-state PDUs in the specified link-state database.	none, **brief**
Neighbor, Neighbor Info	Media access control (MAC) address of the neighbor on the advertising system.	**detail**, **extensive**
Interface	Name of the Virtual Chassis port interface, in the format **vcp-*slot/pic/port.logical-unit-number***.	**detail**, **extensive**
Metric	Metric value of the prefix or neighbor.	**detail**, **extensive**

Table 1763: show virtual-chassis protocol database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Header	Link-state PDU (LSP) packet header: • **LSP ID**—LSP identifier in the header. • **Length**—Header length, in bytes. • **Allocated length**—Length available for the header, in bytes. • **Remaining lifetime**—Remaining lifetime of the link-state PDU, in seconds. • **Interface**—The interface from which the LSP is received. • **Estimated free bytes**—Estimated number of available bytes in the LSP. • **Actual free bytes**—Actual number of available bytes in the LSP. • **Aging timer expires in**—Remaining lifetime of the LSP, in seconds.	**extensive**
Packet	Link-state PDU (LSP) packet: • **LSP ID**—Identifier for the link-state packet. • **Length**—Packet length, in bytes. • **Lifetime**—Remaining lifetime, in seconds. • **Checksum**—Checksum of the link-state PDU. • **Sequence**—Sequence number of the link-state PDU. This number increments whenever the link-state PDU is updated. • **Fixed length**—Set length for the packet, in bytes. • **Version**—Protocol version. • **Sysid length**—Length of the system ID, in bytes. The value 0 represents 6 bytes. • **Packet type**—Protocol data unit (PDU) type of the LSP. • **SW version**—Junos OS Release number.	**extensive**

Table 1763: show virtual-chassis protocol database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
TLVs	Link-state PDU (LSP) type, length, and value (TLV): • **Member ID**—Identifier configured for the Virtual Chassis member router. • **VC ID**—Identifier assigned to the Virtual Chassis member router. • **Flags**—Internal flags that keep track of the member state for the purpose of primary-role election in the Virtual Chassis. • **Priority**—Priority value associated with the role assigned to a member router in the preprovisioned Virtual Chassis configuration. For example, the priority value for the **routing-engine** role is 129. The priority value is used for primary-role election in the Virtual Chassis. • **System ID**—System ID of the device associated with the Virtual Chassis port interface. The System ID is derived from the device's media access control (MAC) address. • **Device ID**—Identifier for the device; usually the same as the Member ID. • **Neighbor Info**—System ID, Virtual Chassis port interface, and metric value for VCCP neighbor. • **Topology Info**—System ID of the VCCP neighbor. • **IRI Addr Info**—Internal routing interface (IRI) IP address, which is reserved for internal communication. • **Master Info**—System ID of the primary router in the Virtual Chassis. • **Backup Info**—System ID of the backup router in the Virtual Chassis. • **Stable State Info**—Internal state information used for primary-role election in the Virtual Chassis. • **Member Info**—System ID, Member ID, and role of each member router in the Virtual Chassis.	**extensive**

Table 1763: show virtual-chassis protocol database Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• **Provision Info**—Member ID and chassis serial number specified for each member router in the preprovisioned configuration for an MX Series Virtual Chassis. • **Unknown TLV**—Type and length of TLVs with unsupported content received on this device.	
***number* queued**	Number of link-state PDUs queued on the specified Virtual Chassis port interface.	**extensive**

Sample Output

show virtual-chassis protocol database all-members brief

```
{master:member1-re0}
user@host> show virtual-chassis protocol database all-members brief
member0:
--------------------------------------------------------------------------
LSP ID                      Sequence Checksum Lifetime
001d.b510.0800.00-00           0x9eb   0xb8f1      115
b0c6.9abf.6800.00-00           0x9ee   0x8f35      116
  2 LSPs

member1:
--------------------------------------------------------------------------
LSP ID                      Sequence Checksum Lifetime
001d.b510.0800.00-00           0x9eb   0xb8f1      117
b0c6.9abf.6800.00-00           0x9ee   0x8f35      114
  2 LSPs
```

show virtual-chassis protocol database member 0 detail

```
{master:member1-re0}
user@host> show virtual-chassis protocol database member 0 detail
member0:
```

```
---------------------------------------------------------------------------

001d.b510.0800.00-00 Sequence: 0x9f5, Checksum: 0x5b2b, Lifetime: 116 secs
   Neighbor: b0c6.9abf.6800.00  Interface: vcp-1/3/0.32768 Metric:      15

b0c6.9abf.6800.00-00 Sequence: 0x9f8, Checksum: 0x326e, Lifetime: 117 secs
   Neighbor: 001d.b510.0800.00  Interface: vcp-5/0/0.32768 Metric:      15
```

show virtual-chassis protocol database member 0 b0c6.9abf.6800 detail

```
{master:member1-re0}
user@host> show virtual-chassis protocol database member 0 b0c6.9abf.6800 detail
member0:
---------------------------------------------------------------------------

b0c6.9abf.6800.00-00 Sequence: 0xa06, Checksum: 0x925b, Lifetime: 117 secs
   Neighbor: 001d.b510.0800.00  Interface: vcp-5/0/0.32768 Metric:      15
```

show virtual-chassis protocol database member 0 extensive

```
{master:member1-re0}
user@host> show virtual-chassis protocol database member 0 extensive
member0:
---------------------------------------------------------------------------

001d.b510.0800.00-00 Sequence: 0xa09, Checksum: 0xa696, Lifetime: 116 secs
   Neighbor: b0c6.9abf.6800.00  Interface: vcp-1/3/0.32768 Metric:      15

  Header: LSP ID: 001d.b510.0800.00-00, Length: 804 bytes
    Allocated length: 804 bytes,
    Remaining lifetime: 116 secs, Interface: 64
    Estimated free bytes: 0, Actual free bytes: 0
    Aging timer expires in: 116 secs

  Packet: LSP ID: 001d.b510.0800.00-00, Length: 804 bytes, Lifetime : 118 secs
    Checksum: 0xa696, Sequence: 0xa09,
    Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 18, SW version: 11.1

  TLVs:
```

```
    Node Info: Member ID: 1, VC ID: 5a6a.e747.8511, Flags: 3, Priority: 129
       System ID: 001d.b510.0800, Device ID: 1
    Unknown TLV, Type: 0, Length: 0
    ...
    Unknown TLV, Type: 0, Length: 0
    Unknown TLV, Type: 1, Length: 1
    Neighbor Info: b0c6.9abf.6800.00, Interface: vcp-1/3/0.32768, Metric: 15
    Topology Info: System ID: 001d.b510.0800,
    Topology Info: System ID: b0c6.9abf.6800,
    IRI Addr Info: IP Address: 128.0.0.1,
    IRI Addr Info: IP Address: 128.0.0.4,
    IRI Addr Info: IP Address: 128.0.0.5,
    IRI Addr Info: IP Address: 128.0.0.6,
    IRI Addr Info: IP Address: 128.0.0.17,
    Master Info: System ID: 001d.b510.0800
    Backup Info: System ID: b0c6.9abf.6800
    Stable State Info: Master ID: 001d.b510.0800, Backup ID: b0c6.9abf.6800
    Member Info: System ID: b0c6.9abf.6800, Member ID: 0 Member role: Backup
       System ID: b0c6.9abf.6800, Device ID: 0
    Member Info: System ID: 001d.b510.0800, Member ID: 1 Member role: Master
       System ID: 001d.b510.0800, Device ID: 1
    Provision Info: Member ID: 1 Serial Number: JN10C78D1AFC,
    Provision Info: Member ID: 0 Serial Number: JN115FDADAFB,
    Unknown TLV, Type: 24, Length: 1
    Unknown TLV, Type: 28, Length: 56

  1 queued :
  Send PSN on vcp-5/0/0.32768 for 00:00:01

b0c6.9abf.6800.00-00 Sequence: 0xa0d, Checksum: 0x82d2, Lifetime: 118 secs
   Neighbor: 001d.b510.0800.00  Interface: vcp-5/0/0.32768 Metric:      15

  Header: LSP ID: b0c6.9abf.6800.00-00, Length: 808 bytes
    Allocated length: 1400 bytes,
    Remaining lifetime: 118 secs, Interface: 0
    Estimated free bytes: 546, Actual free bytes: 592
    Aging timer expires in: 118 secs

  Packet: LSP ID: b0c6.9abf.6800.00-00, Length: 808 bytes, Lifetime : 118 secs
    Checksum: 0x82d2, Sequence: 0xa0d,
    Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 18, SW version: 11.1
```

```
TLVs:
  Node Info: Member ID: 0, VC ID: 5a6a.e747.8511, Flags: 5, Priority: 129
     System ID: b0c6.9abf.6800, Device ID: 0
  Unknown TLV, Type: 0, Length: 0
  ...
  Unknown TLV, Type: 0, Length: 0
  Unknown TLV, Type: 1, Length: 1
  Neighbor Info: 001d.b510.0800.00, Interface: vcp-5/0/0.32768, Metric: 15
  Topology Info: System ID: 001d.b510.0800,
  Topology Info: System ID: b0c6.9abf.6800,
  IRI Addr Info: IP Address: 128.0.0.1,
  IRI Addr Info: IP Address: 128.0.0.4,
  IRI Addr Info: IP Address: 128.0.0.5,
  IRI Addr Info: IP Address: 128.0.0.6,
  IRI Addr Info: IP Address: 128.0.0.17,
  IRI Addr Info: IP Address: 128.0.0.21,
  Master Info: System ID: 001d.b510.0800
  Backup Info: System ID: b0c6.9abf.6800
  Stable State Info: Master ID: 001d.b510.0800, Backup ID: b0c6.9abf.6800
  Member Info: System ID: b0c6.9abf.6800, Member ID: 0 Member role: Backup
     System ID: b0c6.9abf.6800, Device ID: 0
  Member Info: System ID: 001d.b510.0800, Member ID: 1 Member role: Master
     System ID: 001d.b510.0800, Device ID: 1
  Provision Info: Member ID: 1 Serial Number: JN10C78D1AFC,
  Provision Info: Member ID: 0 Serial Number: JN115FDADAFB,
  Unknown TLV, Type: 24, Length: 1
  Unknown TLV, Type: 28, Length: 56

1 queued :
Retransmit on vcp-5/0/0.32768 for 00:00:01
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Viewing Information in the Virtual Chassis Control Protocol Link-State Database

show virtual-chassis protocol interface

IN THIS SECTION

- Syntax | 27647
- Description | 27647
- Options | 27647
- Required Privilege Level | 27648
- Output Fields | 27648
- Sample Output | 27648
- Release Information | 27649

Syntax

```
show virtual-chassis protocol interface
<brief | detail>
<all-members | local | member member-id>
<interface-name>
```

Description

Display information about Virtual Chassis Control Protocol (VCCP) statistics for VCCP-enabled interfaces within the Virtual Chassis or VCF.

Options

none Display the VCCP interface statistics in brief form for all members of the Virtual Chassis or VCF.

brief | detail (Optional) Display the specified level of output. Using the `brief` option is equivalent to entering the command with no options (the default). The `detail` option provides more output than the `brief` option.

all-members (Optional) Display VCCP interface statistics for all members of the Virtual Chassis or VCF.

interface-name	(Optional) Display VCCP interface statistics for the specified interface.
local	(Optional) Display VCCP interface statistics for the switch or external Routing Engine where you enter this command.
member *member-id*	(Optional) Display VCCP interface statistics for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 1764 on page 27648 lists the output fields for the `show virtual-chassis protocol interface` command. Output fields are listed in the approximate order in which they appear.

Table 1764: show virtual-chassis protocol interface Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the VCP.	All levels
`State`	State of the link. Outputs include: • `Up`—The link is up. • `Down`—The link is down.	All levels
`Metric`	Metric of the prefix or neighbor.	All levels

Sample Output

show virtual-chassis protocol interface

```
user@switch> show virtual-chassis protocol interface
member0:
--------------------------------------------------------------------------
IS-IS interface database:
```

```
Interface             State         Metric
vcp-0/0.32768         Up             150
vcp-0/1.32768         Up             150
vcp-4/0/1.32768       Up             15
vcp-4/0/7.32768       Down           15

member1:
--------------------------------------------------------------------------
IS-IS interface database:
Interface             State         Metric
vcp-0/0.32768         Up             150
vcp-0/1.32768         Up             150
vcp-3/0/4.32768       Up             15

member8:
--------------------------------------------------------------------------
IS-IS interface database:
Interface             State         Metric
vcp-0/0.32768         Down           150
vcp-1/0.32768         Up             150
vcp-1/1.32768         Up             150
vcp-1/2.32768         Up             150
vcp-1/3.32768         Down           150
vcp-2/0.32768         Up             150
vcp-2/1.32768         Down           150
vcp-2/2.32768         Down           150
vcp-2/3.32768         Down           150

member9:
--------------------------------------------------------------------------
IS-IS interface database:
Interface             State         Metric
vcp-0/0.32768         Down           150
vcp-1/0.32768         Up             150
vcp-1/1.32768         Up             150
vcp-1/2.32768         Down           150
vcp-1/3.32768         Down           150
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Understanding EX Series Virtual Chassis
- Understanding QFX Series Virtual Chassis
- Understanding Virtual Chassis Ports in an EX8200 Virtual Chassis
- Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis

show virtual-chassis protocol interface (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 27650
- Description | 27650
- Options | 27651
- Required Privilege Level | 27651
- Output Fields | 27651
- Sample Output | 27653
- Release Information | 27654

Syntax

```
show virtual-chassis protocol interface
<(brief | detail)>
<interface-name>
<(all-members | local | member member-id)>
```

Description

Display Virtual Chassis Control Protocol (VCCP) information about Virtual Chassis port interfaces in an MX Series Virtual Chassis. You can issue the `show virtual-chassis protocol interface` command from the console of either member router in the Virtual Chassis.

Options

brief | detail (Optional) Display the specified level of output. Using the **brief** option is equivalent to issuing the command with no options (the default). The **detail** option provides more output than the **brief** option.

all-members (Optional) Display VCCP information about Virtual Chassis port interfaces for both member routers in a Virtual Chassis. This is the default behavior if no options are specified.

interface-name (Optional) Display VCCP information about Virtual Chassis port interfaces for the specified Virtual Chassis port, in the format **vcp-*slot/ pic/ port.logical-unit-number***.

local (Optional) Display VCCP information about Virtual Chassis port interfaces for the member router on which you are issuing the command.

member *member-id* (Optional) Display VCCP information about Virtual Chassis port interfaces for the specified member router. Replace ***member-id*** with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 1765 on page 27651 lists the output fields for the `show virtual-chassis protocol interface` command. Output fields are listed in the approximate order in which they appear.

Table 1765: show virtual-chassis protocol interface Output Fields

Field Name	Field Description	Level of Output
member*n*	Member ID of the Virtual Chassis member router for which output is displayed.	All levels
Interface	Name of the Virtual Chassis port interface, in the format **vcp-*slot/ pic/ port.logical-unit-number***.	**brief**

Table 1765: show virtual-chassis protocol interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the Virtual Chassis port interface: • **Up**—The interface is up. • **Down**—The interface is down.	**brief**
Metric	Metric value for this Virtual Chassis port interface.	All levels
vcp-*slot/ pic/ port. logical-unit-number*	Name of the Virtual Chassis port interface.	**detail**
Index	Interface index number assigned by the Junos OS software.	**detail**
State	Internal implementation information.	**detail**
LSP interval	Interval, in milliseconds, between link-state protocol data units (PDUs) sent from the interface.	**detail**
***type* Hello padding**	Type of hello padding: • **Adaptive**—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state type, length, and value (TLV). If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. • **Loose**—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. • **Strict**—Padding is performed on all interface types and for all adjacency states, and is continuous.	**detail**

Table 1765: show virtual-chassis protocol interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Adjacencies	Number of adjacencies established on this Virtual Chassis port interface.	**detail**
Hello(s)	Hello interval for the Virtual Chassis port interface.	**detail**
Hold(s)	Hold time for the Virtual Chassis port interface.	**detail**

Sample Output

show virtual-chassis protocol interface brief all-members

```
{master:member1-re0}
user@host> show virtual-chassis protocol interface brief all-members
member0:
--------------------------------------------------------------------------
IS-IS interface database:
Interface             State         Metric
vcp-5/0/0.32768       Up             15

member1:
--------------------------------------------------------------------------
IS-IS interface database:
Interface             State         Metric
vcp-1/3/0.32768       Up             15
```

show virtual-chassis protocol interface detail all-members

```
{master:member1-re0}
user@host> show virtual-chassis protocol interface detail all-members
member0:
--------------------------------------------------------------------------

IS-IS interface database:
```

```
vcp-5/0/0.32768
  Index: 64, State: 0x46
  LSP interval: 100 ms, Loose Hello padding
  Adjacencies Metric Hello (s) Hold (s)n           1     15         3     60

member1:
--------------------------------------------------------------------------

IS-IS interface database:
vcp-1/3/0.32768
  Index: 64, State: 0x86
  LSP interval: 100 ms, Loose Hello padding
  Adjacencies Metric Hello (s) Hold (s)n           1     15         3     60
```

show virtual-chassis protocol interface detail local

```
{master:member1-re0}
user@host> show virtual-chassis protocol interface detail local


IS-IS interface database:
vcp-1/3/0.32768
  Index: 64, State: 0x46
  LSP interval: 100 ms, Loose Hello padding
  Adjacencies Metric Hello (s) Hold (s)n           1     15         3     60
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

> *Viewing Information About Virtual Chassis Port Interfaces in the Virtual Chassis Control Protocol Database*

show virtual-chassis protocol route

IN THIS SECTION

- Syntax | 27655
- Description | 27655
- Options | 27655
- Required Privilege Level | 27656
- Output Fields | 27656
- Sample Output | 27657
- Release Information | 27659

Syntax

```
show virtual-chassis protocol route
<all-members | local | member member-id>
<destination-id>
```

Description

Display the unicast and multicast Virtual Chassis Control Protocol (VCCP) routing tables within the Virtual Chassis or VCF.

Options

none	Display the unicast and multicast routing tables for all members of the Virtual Chassis.
all-members	(Optional) Display the unicast and multicast routing tables for all members of the Virtual Chassis or VCF.
destination-id	(Optional) Display the unicast and multicast routing tables to the specified destination member ID for each member of the Virtual Chassis or VCF.

local	(Optional) Display the unicast and multicast routing tables on the device where you enter this command.
member ***member-id***	(Optional) Display the unicast and multicast routing tables for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 1766 on page 27656 lists the output fields for the `show virtual-chassis protocol route` command. Output fields are listed in the approximate order in which they appear.

Table 1766: show virtual-chassis protocol route Output Fields

Field Name	Field Description
`Dev`	MAC address of the member storing the VCCP routing table.
`Version`	Version of the shortest-path-first algorithm that generated the routing table.
`System ID`	MAC address of the device.
`Version`	Version of the shortest-path-first (SPF) algorithm that generated the route.
`Metric`	The metric number to get to that device.
`Interface`	Name of the Virtual Chassis port (VCP) interface connecting the devices.
`Via`	MAC address of the next-hop device, if applicable.

Sample Output

show virtual-chassis protocol route

```
user@switch> show virtual-chassis protocol route
member0:
--------------------------------------------------------------------------

Dev 0021.59f7.d000 ucast routing table            Current version: 21
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          21      150 vcp-0/1.32768 0000.4a75.9b7c
0000.73e9.9a57          21      165 vcp-4/0/1.32768 0026.888d.6800
0021.59f7.d000          21        0
0026.888d.6800          21       15 vcp-4/0/1.32768 0026.888d.6800

Dev 0021.59f7.d000 mcast routing table            Current version: 21
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          21
0000.73e9.9a57          21
0021.59f7.d000          21          vcp-4/0/1.32768
                                    vcp-0/1.32768
0026.888d.6800          21

member1:
--------------------------------------------------------------------------

Dev 0026.888d.6800 ucast routing table            Current version: 25
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          25      150 vcp-0/0.32768 0000.4a75.9b7c
0000.73e9.9a57          25      150 vcp-0/1.32768 0000.73e9.9a57
0021.59f7.d000          25       15 vcp-3/0/4.32768 0021.59f7.d000
0026.888d.6800          25        0

Dev 0026.888d.6800 mcast routing table            Current version: 25
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          25
0000.73e9.9a57          25          vcp-3/0/4.32768
0021.59f7.d000          25          vcp-0/1.32768
```

```
0026.888d.6800          25          vcp-3/0/4.32768
                                    vcp-0/0.32768
                                    vcp-0/1.32768

member8:
--------------------------------------------------------------------------

Dev 0000.4a75.9b7c ucast routing table             Current version: 39
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          39        0
0000.73e9.9a57          39      150 vcp-1/0.32768 0000.73e9.9a57
0021.59f7.d000          39      150 vcp-2/0.32768 0021.59f7.d000
0026.888d.6800          39      150 vcp-1/2.32768 0026.888d.6800

Dev 0000.4a75.9b7c mcast routing table             Current version: 39
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          39          vcp-1/0.32768
                                    vcp-2/0.32768
                                    vcp-1/2.32768
0000.73e9.9a57          39
0021.59f7.d000          39
0026.888d.6800          39

member9:
--------------------------------------------------------------------------

Dev 0000.73e9.9a57 ucast routing table             Current version: 31
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          31      150 vcp-1/0.32768 0000.4a75.9b7c
0000.73e9.9a57          31        0
0021.59f7.d000          31      165 vcp-1/1.32768 0026.888d.6800
0026.888d.6800          31      150 vcp-1/1.32768 0026.888d.6800

Dev 0000.73e9.9a57 mcast routing table             Current version: 31
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          31
0000.73e9.9a57          31          vcp-1/0.32768
                                    vcp-1/1.32768
```

```
0021.59f7.d000          31
0026.888d.6800          31
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Understanding EX Series Virtual Chassis
- Understanding QFX Series Virtual Chassis
- Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis

show virtual-chassis protocol route (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 27659
- Description | 27660
- Options | 27660
- Required Privilege Level | 27660
- Output Fields | 27660
- Sample Output | 27661
- Release Information | 27663

Syntax

```
show virtual-chassis protocol route
<destination-id>
<(all-members | local | member member-id)>
```

Description

Display the Virtual Chassis Control Protocol (VCCP) unicast and multicast routing tables for an MX Series Virtual Chassis. You can issue the `show virtual-chassis protocol route` command from the console of either member router in the Virtual Chassis.

Options

all-members (Optional) Display the VCCP unicast and multicast routing tables for both member routers in a Virtual Chassis configuration. This is the default behavior if no options are specified.

destination-id (Optional) Display the VCCP unicast and multicast routing tables to the destination with the specified system ID.

local (Optional) Display the VCCP unicast and multicast routing tables for the member router on which you are issuing the command.

member *member-id* (Optional) Display the VCCP unicast and multicast routing tables for the specified member router. Replace ***member-id*** with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 1767 on page 27660 lists the output fields for the `show virtual-chassis protocol route` command. Output fields are listed in the approximate order in which they appear.

Table 1767: show virtual-chassis protocol route Output Fields

Field Name	Field Description
member*n*	Member ID of the Virtual Chassis member router for which output is displayed.
Dev	System ID of the device (member router) that stores the VCCP routing tables. The System ID is derived from the router's media access control (MAC) address.

Table 1767: show virtual-chassis protocol route Output Fields *(Continued)*

Field Name	Field Description
ucast routing table	VCCP unicast routing table.
mcast routing table	VCCP multicast routing table.
Current version	Version of the shortest-path-first (SPF) algorithm that generated the VCCP unicast or multicast routing table.
System ID	System ID of the device, derived from the device's MAC address.
Version	Version of the SPF algorithm that generated this route in the VCCP unicast or multicast routing table.
Metric	Metric value required to reach this device.
Interface	Name of the Virtual Chassis port interface (in the format **vcp-*slot/pic/port.logical-unit-number***) that interconnects the devices.
Via	MAC address of the next-hop device, if applicable.

Sample Output

show virtual-chassis protocol route all-members

```
{master:member1-re0}
user@host> show virtual-chassis protocol route all-members
member0:
--------------------------------------------------------------------------

Dev b0c6.9abf.6800 ucast routing table            Current version: 17
----------------
System ID          Version   Metric Interface     Via
001d.b510.0800          17       15 vcp-5/0/0.32768 001d.b510.0800
b0c6.9abf.6800          17        0
```

```
Dev b0c6.9abf.6800 mcast routing table            Current version: 17
----------------
System ID          Version   Metric Interface     Via
001d.b510.0800          17
b0c6.9abf.6800          17          vcp-5/0/0.32768

member1:
--------------------------------------------------------------------------

Dev 001d.b510.0800 ucast routing table            Current version: 17
----------------
System ID          Version   Metric Interface     Via
001d.b510.0800          17        0
b0c6.9abf.6800          17       15 vcp-1/3/0.32768 b0c6.9abf.6800

Dev 001d.b510.0800 mcast routing table            Current version: 17
----------------
System ID          Version   Metric Interface     Via
001d.b510.0800          17          vcp-1/3/0.32768
b0c6.9abf.6800          17
```

show virtual-chassis protocol route member 0 001d.b510.0800 (For Specific Member ID and Destination ID)

```
{master:member1-re0}
user@host> show virtual-chassis protocol route member 0 001d.b510.0800
member0:
--------------------------------------------------------------------------

Dev b0c6.9abf.6800 ucast routing table            Current version: 17
----------------
System ID          Version   Metric Interface     Via
001d.b510.0800          17       15 vcp-5/0/0.32768 001d.b510.0800
b0c6.9abf.6800          17        0

Dev b0c6.9abf.6800 mcast routing table            Current version: 17
----------------
System ID          Version   Metric Interface     Via
```

```
001d.b510.0800          17
b0c6.9abf.6800          17          vcp-5/0/0.32768
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Viewing Virtual Chassis Control Protocol Routing Tables

show virtual-chassis protocol statistics

IN THIS SECTION

- Syntax | 27663
- Description | 27664
- Options | 27664
- Required Privilege Level | 27664
- Output Fields | 27664
- Sample Output | 27665
- Release Information | 27667

Syntax

```
show virtual-chassis protocol statistics
<all-members | local | member member-id>
<interface-name>
```

Description

Display the Virtual Chassis Control Protocol (VCCP) statistics for all hardware devices within the Virtual Chassis or VCF.

Options

none	Display VCCP statistics for all members of the Virtual Chassis or VCF.
all-members	(Optional) Display VCCP statistics for all members of the Virtual Chassis or VCF.
interface-name	(Optional) Display VCCP statistics for the specified interface.
local	(Optional) Display VCCP statistics for the switch or external Routing Engine where you run this command.
member *member-id*	(Optional) Display VCCP statistics for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 1768 on page 27664 lists the output fields for the `show virtual-chassis protocol interface` command. Output fields are listed in the approximate order in which they appear.

Table 1768: show virtual-chassis protocol statistics Output Fields

Field Name	Field Description
`PDU type`	Protocol data unit type.
`Received`	Number of PDUs received since VCCP started or since the statistics were set to zero.
`Processed`	Number of PDUs received minus the number of PDUs dropped.
`Drops`	Number of PDUs dropped.

Table 1768: show virtual-chassis protocol statistics Output Fields *(Continued)*

Field Name	Field Description
Sent	Number of PDUs transmitted since VCCP started or since the statistics were set to zero.
Rexmit	Number of PDUs retransmitted since VCCP started or since the statistics were set to zero.
Total Packets Received	Number of PDUs received since VCCP started or since the statistics were set to zero.
Total Packets Sent	Number of PDUs sent since VCCP started or since the statistics were set to zero.
LSP queue length	Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.
SPF runs	Number of shortest-path-first (SPF) calculations that have been performed.
Fragments Rebuilt	Number of link-state PDU fragments that the local system has computed.
LSP Regenerations	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software determines that a link-state PDU must be removed from the network.

Sample Output

show virtual-chassis protocol statistics

```
user@switch> show virtual-chassis protocol statistics
member0:
--------------------------------------------------------------------------

IS-IS statistics for 0021.59f7.d000:
```

```
PDU type       Received    Processed        Drops         Sent       Rexmit
LSP                8166         8166            0         4551            0
HELLO              1659         1659            0         1693            0
CSNP                  2            2            0            3            0
PSNP               1909         1909            0         2293            0
Unknown               0            0            0            0            0
Totals            11736        11736            0         8540            0

Total packets received: 11736 Sent: 8540

LSP queue length: 0 Drops: 0
SPF runs: 9
Fragments rebuilt: 1640
LSP regenerations: 1
Purges initiated: 0

member1:
---------------------------------------------------------------------------

IS-IS statistics for 0026.888d.6800:
PDU type       Received    Processed        Drops         Sent       Rexmit
LSP               10909        10909            0        12088            0
HELLO              1877         1877            0         2251            0
CSNP                  3            3            0            3            0
PSNP               3846         3846            0         3732            0
Unknown               0            0            0            0            0
Totals            16635        16635            0        18074            0

Total packets received: 16635 Sent: 18074

LSP queue length: 0 Drops: 0
SPF runs: 13
Fragments rebuilt: 1871
LSP regenerations: 2
Purges initiated: 0

member8:
---------------------------------------------------------------------------

IS-IS statistics for 0000.4a75.9b7c:
PDU type       Received    Processed        Drops         Sent       Rexmit
LSP                7935         7935            0        14865            0
HELLO              2695         2695            0         7124            0
```

```
CSNP                  4            4            0            4            0
PSNP               4398         4398            0         3666            0
Unknown               0            0            0            0            0
Totals            15032        15032            0        25659            0

Total packets received: 15032 Sent: 25659

LSP queue length: 0 Drops: 0
SPF runs: 26
Fragments rebuilt: 2666
LSP regenerations: 4
Purges initiated: 0

member9:
--------------------------------------------------------------------------

IS-IS statistics for 0000.73e9.9a57:
PDU type       Received    Processed        Drops         Sent       Rexmit
LSP               10800        10800            0         6327            0
HELLO              1492         1492            0         2356            0
CSNP                  2            2            0            2            0
PSNP               2683         2683            0         3149            0
Unknown               0            0            0            0            0
Totals            14977        14977            0        11834            0

Total packets received: 14977 Sent: 11834

LSP queue length: 0 Drops: 0
SPF runs: 19
Fragments rebuilt: 1510
LSP regenerations: 6
Purges initiated: 0
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Understanding EX Series Virtual Chassis

Understanding QFX Series Virtual Chassis

Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis

show virtual-chassis protocol statistics (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 27668
- Description | 27668
- Options | 27668
- Required Privilege Level | 27669
- Output Fields | 27669
- Sample Output | 27670
- Release Information | 27672

Syntax

```
show virtual-chassis protocol statistics
<interface-name>
<(all-members | local | member member-id)>
```

Description

Display Virtual Chassis Control Protocol (VCCP) statistics for one or both member routers, or for a specified Virtual Chassis port interface, in an MX Series Virtual Chassis. You can issue the `show virtual-chassis protocol statistics` command from the console of either member router in the Virtual Chassis.

Options

all-members (Optional) Display VCCP statistics for both member routers in a Virtual Chassis configuration. This is the default behavior if no options are specified.

interface-name	(Optional) Display VCCP statistics for the specified Virtual Chassis port interface, in the format **vcp-***slot/pic/port.logical-unit-number*.
local	(Optional) Display VCCP statistics for the member router on which you are issuing the command.
member *member-id*	(Optional) Display VCCP statistics for the specified member router. Replace *member-id* with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 1769 on page 27669 lists the output fields for the `show virtual-chassis protocol statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1769: show virtual-chassis protocol statistics Output Fields

Field Name	Field Description
member*n*	Member ID of the Virtual Chassis member router for which output is displayed.
PDU type	Type of protocol data unit (PDU).
Received	Number of PDUs received since VCCP started or since the statistics were set to zero.
Processed	Number of PDUs received minus the number of PDUs dropped.
Drops	Number of PDUs dropped.
Sent	Number of PDUs transmitted since VCCP started or since the statistics were set to zero.
Rexmit	Number of PDUs retransmitted since VCCP started or since the statistics were set to zero.

Table 1769: show virtual-chassis protocol statistics Output Fields *(Continued)*

Field Name	Field Description
Total packets received	Total number of PDUs received since VCCP started or since the statistics were set to zero.
Sent	Total number of PDUs transmitted since VCCP started or since the statistics were set to zero.
LSP queue length	Number of link-state PDUs waiting in the queue to be processed.
Drops	Number of link-state PDUs dropped.
SPF runs	Number of shortest-path-first (SPF) calculations performed.
Fragments rebuilt	Number of link-state PDU fragments computed by the local system.
LSP regenerations	Number of regenerated link-state PDUs. A link-state PDU is regenerated when the PDU nears the end of its lifetime and has not changed.
Purges initiated	Number of purges initiated by the software. A purge is initiated when the software determines that it must remove a link-state PDU from the network.

Sample Output

show virtual-chassis protocol statistics all-members

```
{master:member1-re0}
user@host> show virtual-chassis protocol statistics all-members
member0:
--------------------------------------------------------------------------

IS-IS statistics for b0c6.9abf.6800:
PDU type       Received    Processed        Drops         Sent       Rexmit
LSP                2937         2937            0         2934            0
HELLO              2913         2913            0         2922            0
```

```
CSNP                  1            1            0            1            0
PSNP               2916         2916            0         2925            0
Unknown               0            0            0            0            0
Totals             8767         8767            0         8782            0

Total packets received: 8767 Sent: 8782

LSP queue length: 0 Drops: 0
SPF runs: 17
Fragments rebuilt: 2955
LSP regenerations: 14
Purges initiated: 0

member1:
--------------------------------------------------------------------------

IS-IS statistics for 001d.b510.0800:
PDU type       Received    Processed        Drops         Sent       Rexmit
LSP                2934         2934            0         2937            0
HELLO              2922         2922            0         2914            0
CSNP                  1            1            0            1            0
PSNP               2925         2925            0         2916            0
Unknown               0            0            0            0            0
Totals             8782         8782            0         8768            0

Total packets received: 8782 Sent: 8768

LSP queue length: 0 Drops: 0
SPF runs: 17
Fragments rebuilt: 2953
LSP regenerations: 11
Purges initiated: 0
```

show virtual-chassis protocol statistics vcp-1/3/0.32768 member 1 (For Specific Virtual Chassis Port Interface and Member ID)

```
{master:member1-re0}
user@host> show virtual-chassis protocol statistics vcp-1/3/0.32768 member 1
member1:
--------------------------------------------------------------------------
vcp-1/3/0.32768
```

```
IS-IS statistics for 001d.b510.0800:
PDU type       Received    Processed        Drops         Sent       Rexmit
LSP                3013         3013            0         3016            0
HELLO              3001         3001            0         2993            0
CSNP                  1            1            0            1            0
PSNP               3003         3003            0         2994            0
Unknown               0            0            0            0            0
Totals             9018         9018            0         9004            0

Total packets received: 9018 Sent: 9004
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Viewing Virtual Chassis Control Protocol Statistics for Member Devices and Virtual Chassis Ports

show virtual-chassis status (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 27673
- Description | 27673
- Required Privilege Level | 27673
- Output Fields | 27673
- Sample Output | 27675
- Release Information | 27675

Syntax

```
show virtual-chassis status
```

Description

Display information about the status of both member routers in an MX Series Virtual Chassis configuration. You can issue the `show virtual-chassis status` command from the console of either member router in the Virtual Chassis.

Required Privilege Level

view

Output Fields

Table 1770 on page 27673 lists the output fields for the `show virtual-chassis status` command. Output fields are listed in the approximate order in which they appear.

Table 1770: show virtual-chassis status Output Fields

Field Name	Field Description
`Virtual Chassis ID`	Assigned ID that applies to the entire Virtual Chassis configuration.
`Member ID`	Member ID assigned in the preprovisioned Virtual Chassis configuration, and the Flexible PIC Concentrator (FPC) slot range, including offset, for each member router in the Virtual Chassis.

Table 1770: show virtual-chassis status Output Fields *(Continued)*

Field Name	Field Description
`Status`	State of the member router: • `Heartbt`—Router has used heartbeat packet connection to maintain primary role roles during an adjacency disruption or split in the Virtual Chassis configuration. • `Mismatch`—Virtual Chassis not formed due to a mismatch in the slot count value of the member routers. In a Virtual Chassis with an MX2020 member router and either an MX960 or MX2010 member router, a mismatch status occurs if you do not set the slot count value of the MX960 router or MX2010 router to 20 to match the slot count of the MX2020 router. • `Prsnt`—Router is currently connected to the Virtual Chassis. • `NotPrsnt`—Router is not currently connected to the Virtual Chassis.
`Serial No`	Serial number of the member router.
`Model`	Model number of the member router.
`Mastership priority`	Metric used by the Virtual Chassis software for the primary-role election algorithm. This value is assigned by the software and is not configurable in the current release.
`Role`	Role of the member router in the Virtual Chassis: `Master` or `Backup`. The asterisk (*) following the Role denotes the router on which the `show virtual-chassis status` command was issued.
`Neighbor List ID Interface`	Member IDs and Virtual Chassis port interfaces (in the format `vcp-`*slot*/*pic*/*port*) to which this member router is connected.

Sample Output

show virtual-chassis status

```
{master:member1-re0}
user@host> show virtual-chassis status
Preprovisioned Virtual Chassis
Virtual Chassis ID: 5a6a.e747.8511
                                                Mastership          Neighbor List
Member ID       Status   Serial No    Model    priority    Role      ID  Interface
0 (FPC   0- 11) Prsnt    JN115FDADAFB mx480           129  Backup     1  vcp-5/0/0
1 (FPC  12- 23) Prsnt    JN10C78D1AFC mx240           129  Master*    0  vcp-1/3/0
```

show virtual-chassis status (Mismatch Status)

```
{master:member1-re0}
user@host> show virtual-chassis status
Preprovisioned Virtual Chassis
Virtual Chassis ID: 4b43.17d9.1dcf
                                                Mastership            Neighbor List
Member ID       Status   Serial No    Model    priority    Role      ID  Interface
0 (FPC   0- 19) Mismatch JN11CA3F9AFK mx2010
1 (FPC  20- 39) Prsnt    JN122D1DAAFJ mx2020          129  Master*
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Verifying the Status of Virtual Chassis Member Routers or Switches

Verifying and Managing the Virtual Chassis Heartbeat Connection

Example: Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in the Same Subnet

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in Different Subnets

show virtual-chassis vc-path

IN THIS SECTION

- Syntax | 27676
- Description | 27676
- Options | 27676
- Required Privilege Level | 27677
- Output Fields | 27677
- Sample Output | 27678
- Release Information | 27680

Syntax

```
show virtual-chassis vc-path source-interface interface-name destination-interface interface-name
```

Description

Show the forwarding path a packet takes when going from a source interface to a destination interface in a Virtual Chassis or VCF configuration.

Starting in Junos OS Releases 14.1X53-D40, 15.1R5, and 16.1R3, when used in VCF configurations, this command displays additional information about next hops, including when the forwarding path has multiple possible next hops within the VCF.

Options

source-interface *interface-name*	Name of the interface from which the packet originates in the Virtual Chassis or VCF

destination-interface *interface-name*	Name of the interface to which the packet is being delivered in the Virtual Chassis or VCF

Required Privilege Level

view

Output Fields

Table 1771 on page 27677 lists the output fields for the `show virtual-chassis vc-path` command. Output fields are listed in the approximate order in which they appear.

Some output field names and display order differ between the output for a Virtual Chassis and for a VCF. When this command is used to display forwarding paths in a VCF, additional fields are included in the output to show details of multiple possible next hops. The differences are described in Table 1771 on page 27677 and shown in sample output for each mode.

Table 1771: show virtual-chassis vc-path Output Fields

Field Name	Field Description
`Hop`	The hop number along the path between the source and destination interfaces. The first hop entry (`Hop 0`) is the packet's source, intermediate hop information represents transitions through the members within the Virtual Chassis or VCF, and the last hop entry represents arrival at the packet's destination. For VCF output, each `Hop` entry also shows information about multiple possible next hops towards the destination. See `Next-hop PFE`, `Interface`, and `Bandwidth` output field descriptions for details.
`Member` (Virtual Chassis) `Member-ID` (VCF)	The Virtual Chassis or VCF member ID of the switch that contains the Packet Forwarding Engine for each hop through which the packet passes.

Table 1771: show virtual-chassis vc-path Output Fields *(Continued)*

Field Name	Field Description
PFE-Device(Virtual Chassis) PFE(VCF)	The number of the Packet Forwarding Engine in each Virtual Chassis or VCF member through which a packet passes. For Virtual Chassis output, the Packet Forwarding Engine in each row is the next hop of the preceding Packet Forwarding Engine, including intermediate transitions through members within the Virtual Chassis. VCF output is expanded on multiple rows to show more information about multiple possible next hops—each hop entry is followed by one or more Next-hop PFE output fields with interface name and bandwidth information for each possible next hop.
Next-hop PFE (VCF only)	One or more possible next-hop Packet Forwarding Engine numbers for VCF member Member-ID. Each Next-hop entry includes the next-hops's interface name and bandwidth.
Interface	The name of the interface through which the Packet Forwarding Engines are connected. The interface for the first hop (Hop 0) is always the source interface. VCF output provides more information about multiple possible next hops for each hop entry, listing the Interface and the interface Bandwidth for each possible Next-hop PFE.
Bandwidth (VCF only)	The bandwidth (in Gbps) of the next-hop interface for the associated Next-hop PFE entry.

Sample Output

show virtual-chassis vc-path source-interface destination-interface (Virtual Chassis)

```
user@switch> show virtual-chassis vc-path source-interface ge-0/0/0 destination-interface
ge-1/0/1
vc-path from ge-0/0/0 to ge-1/0/1
Hop      Member     PFE-Device      Interface
0         0          1              ge-0/0/0
1         0          0              internal-1/24
2         1          3              vcp-0
3         1          4              ge-1/0/1
```

show virtual-chassis vc-path source-interface destination-interface (Virtual Chassis Fabric)

This example shows `vc-path` command output for two source-to-destination paths in the following Virtual Chassis Fabric displayed by the `show virtual-chassis` command:

```
user@switch> show virtual-chassis

Preprovisioned Virtual Chassis Fabric
Fabric ID: ec56.0915.8595
Fabric Mode: Enabled
                                          Mstr         Mixed Route Neighbor List
Member ID Status Serial No    Model       prio  Role    Mode Mode  ID  Interface
1 (FPC 1) Prsnt  ...          qfx5100-24q-aa 129   Backup     N  F    7  vcp-255/0/7
                                                                      8  vcp-255/0/8
                                                                      12 vcp-255/0/12
                                                                      3  vcp-255/0/3
                                                                      4  vcp-255/0/4
                                                                      5  vcp-255/0/5
                                                                      6  vcp-255/0/6
                                                                      10 vcp-255/0/10
                                                                      11 vcp-255/0/11
                                                                      9  vcp-255/0/9
2 (FPC 2) Prsnt  ...          qfx5100-24q-aa 129   Master*    N  F    7  vcp-255/0/7
                                                                      8  vcp-255/0/8
                                                                      12 vcp-255/0/12
                                                                      3  vcp-255/0/3
                                                                      4  vcp-255/0/4
                                                                      5  vcp-255/0/5
                                                                      6  vcp-255/0/6
                                                                      10 vcp-255/0/10
                                                                      11 vcp-255/0/11
                                                                      9  vcp-255/0/9
3 (FPC 3)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                      2  vcp-255/0/2
4 (FPC 4)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                      2  vcp-255/0/2
5 (FPC 5)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                      2  vcp-255/0/2
6 (FPC 6)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                      2  vcp-255/0/2
7 (FPC 7)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                      2  vcp-255/0/2
```

```
8 (FPC 8)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                       2  vcp-255/0/2
9 (FPC 9)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                       2  vcp-255/0/2
10 (FPC 10)Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                       2  vcp-255/0/2
11 (FPC 11)Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                       2  vcp-255/0/2
12 (FPC 12)Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                       2  vcp-255/0/2


user@switch> show virtual-chassis vc-path source-interface xe-12/0/25 destination-interface
xe-5/0/25
Fabric forwarding path from xe-12/0/25 (PFE 12) to xe-5/0/25 (PFE 5)
  Hop   0  Member-ID  12  PFE   12
    Next-hop PFE   1
      Interface vcp-255/0/1.32768  Bandwidth   40
    Next-hop PFE   2
      Interface vcp-255/0/2.32768  Bandwidth   40
  Hop   1  Member-ID   1  PFE    1
    Next-hop PFE   5
      Interface vcp-255/0/5.32768  Bandwidth   40
  Hop   1  Member-ID   2  PFE    2
    Next-hop PFE   5
      Interface vcp-255/0/5.32768  Bandwidth   40
  Hop   2  Member-ID   5  PFE    5


user@switch> show virtual-chassis vc-path source-interface xe-12/0/25 destination-interface
xe-1/0/25
Fabric forwarding path from xe-12/0/25 (PFE 12) to xe-1/0/25 (PFE 1)
  Hop   0  Member-ID  12  PFE   12
    Next-hop PFE   1
      Interface vcp-255/0/1.32768  Bandwidth   40
  Hop   1  Member-ID   1  PFE    1
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis
- Understanding EX Series Virtual Chassis
- Understanding QFX Series Virtual Chassis
- EX8200 Virtual Chassis Overview

show virtual-chassis vc-port

IN THIS SECTION

- Syntax | 27681
- Description | 27681
- Options | 27682
- Required Privilege Level | 27682
- Output Fields | 27682
- Sample Output | 27684
- Release Information | 27688

Syntax

```
show virtual-chassis vc-port
<all-members | local | member member-id>
```

Description

Display the status of the Virtual Chassis ports (VCPs), including dedicated VCPs, default-configured VCPs, and uplink ports configured as VCPs, if present.

Options

none Display the operational status of all VCPs of the member switch where you enter the command.

all-members (Optional) Display the operational status of all VCPs on all members of the Virtual Chassis or VCF.

local (Optional) Display the operational status of the switch or external Routing Engine where you enter this command.

member *member-id* (Optional) Display the operational status of all VCPs for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

view

Output Fields

Table 1772 on page 27682 lists the output fields for the `show virtual-chassis vc-port` command. Output fields are listed in the approximate order in which they appear.

Table 1772: show virtual-chassis vc-port Output Fields

Field Name	Field Description
member name	Device or member number in the Virtual Chassis, represented as `fpc`*number*, `member`*number*, or local-re (the primary member).

Table 1772: show virtual-chassis vc-port Output Fields *(Continued)*

Field Name	Field Description
`Interface or PIC/Port`	VCP name. • The dedicated VCPs in an EX4200 or EX4500 Virtual Chassis are `vcp-0` and `vcp-1`. The dedicated VCPs in an EX4550 Virtual Chassis are `VCP-1/0`, `VCP-1/1`, `VCP-2/0`, and `VCP-2/1`. • Optical ports set as VCPs are named `1/0` and `1/1`, representing the PIC number and the port number. • The native VCP (port 0) on an XRE200 External Routing Engine in an EX8200 Virtual Chassis is named `vcp-0`. • The VCPs on each Virtual Chassis Control Interface (VCCI) module in an XRE200 External Routing Engine are named using the `vcp-`*slot-number*/*port-number* convention; for instance, `vcp-1/0`. • The VCPs on EX8200 member switches are named using the `vcp-`*slot-number*/*pic-number*/*interface-number* convention; for instance, `vcp-3/0/2`. • A `255` as the first number in your port number indicates that your VCP is part of a Link Aggregation group (LAG) bundle. For example, `vcp-255/1/0` indicates that the dedicated VCP named `vcp-1/0` is part of a LAG bundle, and `vcp-255/1/0` represents an uplink port previously named `xe-0/1/0` that is now part of a VCP LAG bundle.
`Type`	Type of VCP: • `Dedicated`—Rear panel dedicated VCP on an EX4200 or EX4300 multigigabit model switch, a Virtual Chassis module port on EX4500 or EX4550 switch, or any VCP link connected to an XRE200 External Routing Engine in an EX8200 Virtual Chassis. • `Configured`—Optical port configured as a VCP, which includes default-configured VCPs (set in the default factory configuration) or those that are set by a user. • `Auto-Configured`—Optical port automatically converted into a VCP. See Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port or Configuring an EX4650 or a QFX Series Virtual Chassis for information about configuring VCPs, and Automatic Virtual Chassis Port (VCP) Conversion for details on how a port is automatically converted into a VCP.

Table 1772: show virtual-chassis vc-port Output Fields *(Continued)*

Field Name	Field Description
`Trunk ID`	A positive-number ID assigned to a link aggregation group (LAG) formed by the Virtual Chassis. The trunk ID value is -1 if no trunk is formed. A LAG between uplink VCPs requires that the link speed be the same on connected interfaces and that at least two VCPs on one member be connected to at least two VCPs on the other member in an EX4200 or EX4500 Virtual Chassis. Dedicated VCP LAGs are assigned trunk IDs 1 and 2. Trunk IDs for LAGs formed with uplink VCPs therefore have values of 3 or greater. The trunk ID value changes if the link-adjacency state between LAG members changes; trunk membership is then allocated or deallocated.
`Status`	Interface status: • `absent`—Interface is not a VCP link. • `down`—VCP link is down. • `up`—VCP link is up.
`Speed (mbps)`	Speed of the interface in megabits per second.
`Neighbor ID/ Interface`	The Virtual Chassis member ID and interface of a VCP on a member that is connected to the interface or PIC/Port field in the same row as this interface.

Sample Output

show virtual-chassis vc-port (EX4200 Virtual Chassis)

```
user@switch> show virtual-chassis vc-port

fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status    Speed        Neighbor
or                             ID              (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up        32000        1   vcp-1
```

```
vcp-1       Dedicated           2    Up        32000        0   vcp-0
1/0         Auto-Configured     3    Up        1000         2   vcp-255/1/0
1/0         Auto-Configured     3    Up        1000         2   vcp-255/1/1
```

show virtual-chassis vc-port (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis vc-port
member0:
--------------------------------------------------------------------------
Interface       Type              Trunk  Status    Speed        Neighbor
or                                 ID              (mbps)       ID  Interface
Slot/PIC/Port
vcp-0/0         Dedicated          -1    Up        1000         8   vcp-1/1
vcp-0/1         Dedicated          -1    Up        1000         8   vcp-2/0
4/0/4           Configured         -1    Up        10000        1   vcp-3/0/4
4/0/7           Configured         -1    Down      10000
4/0/3           Configured               Absent
4/0/2           Configured               Absent
4/0/5           Configured               Absent
4/0/6           Configured               Absent
4/0/1           Configured               Absent
4/0/0           Configured               Absent

member1:
--------------------------------------------------------------------------
Interface       Type              Trunk  Status    Speed        Neighbor
or                                 ID                 (mbps)    ID  Interface
Slot/PIC/Port
vcp-0/0         Dedicated          -1    Up        1000         8   vcp-1/2
3/0/0           Configured         -1    Down      10000
3/0/1           Configured         -1    Down      10000
3/0/4           Configured         -1    Up        10000        0   vcp-4/0/4
3/0/5           Configured               Absent
4/0/5           Configured               Absent
4/0/4           Configured               Absent

member8:
--------------------------------------------------------------------------
Interface       Type              Trunk  Status    Speed        Neighbor
or                                 ID              (mbps)       ID  Interface
Slot/PIC/Port
```

```
vcp-0/0         Dedicated          -1    Down      1000
vcp-1/0         Dedicated          -1    Up        1000         9   vcp-1/0
vcp-1/1         Dedicated          -1    Up        1000         0   vcp-0/0
vcp-1/2         Dedicated          -1    Up        1000         1   vcp-0/0
vcp-1/3         Dedicated          -1    Up        1000         9   vcp-1/3
vcp-2/0         Dedicated          -1    Up        1000         0   vcp-0/1
vcp-2/1         Dedicated          -1    Up        1000         9   vcp-1/2
vcp-2/2         Dedicated          -1    Down      1000
vcp-2/3         Dedicated          -1    Down         1000

member9:
--------------------------------------------------------------------------
Interface       Type              Trunk  Status    Speed        Neighbor
or                                 ID              (mbps)       ID  Interface
Slot/PIC/Port
vcp-0/0         Dedicated          -1    Disabled  1000
vcp-1/0         Dedicated          -1    Up        1000         8   vcp-1/0
vcp-1/1         Dedicated          -1    Down      1000
vcp-1/2         Dedicated          -1    Up        1000         8   vcp-2/1
vcp-1/3         Dedicated          -1    Up        1000         8   vcp-1/3
```

show virtual-chassis vc-port all-members

```
user@switch> show virtual-chassis vc-port all-members



fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status  Speed    Neighbor
or                             ID            (mbps)   ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up      32000     1   vcp-1
vcp-1       Dedicated           2    Up      32000     0   vcp-0
1/0         Auto-Configured     3    Up      1000      2   vcp-255/1/0
1/1         Auto-Configured     3    Up      1000      2   vcp-255/1/1

fpc1:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status  Speed    Neighbor
or                             ID            (mbps)   ID  Interface
PIC / Port
```

```
vcp-0       Dedicated           1    Up      32000     0   vcp-1
vcp-1       Dedicated           2    Up      32000     0   vcp-0
1/0         Auto-Configured    -1    Up      1000      3   vcp-255/1/0

fpc2:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status   Speed    Neighbor
or                             ID             (mbps)   ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up      32000     3   vcp-1
vcp-1       Dedicated           2    Up      32000     3   vcp-0
1/0         Auto-Configured     3    Up      1000      0   vcp-255/1/0
1/1         Auto-Configured     3    Up      1000      0   vcp-255/1/1

fpc3:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status   Speed    Neighbor
or                             ID             (mbps)   ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up      32000     2   vcp-0
vcp-1       Dedicated           2    Up      32000     2   vcp-1
1/0         Auto-Configured    -1    Up      1000      1   vcp-255/1/0
```

show virtual-chassis vc-port (QFX5120-32C Virtual Chassis)

```
user@switch> show virtual-chassis vc-port

localre:
-
Interface   Type              Trunk  Status  Speed     Neighbor
or                             ID            (mbps)    ID  Interface
PIC / Port
0/13        Auto-Configured     6    Up      100000    1   vcp-255/0/13
0/23        Auto-Configured     6    Up      100000    1   vcp-255/0/23
0/26        Auto-Configured     5    Up      40000     1   vcp-255/0/26
0/27        Configured          5    Up      40000     1   vcp-255/0/27
0/6         Configured          5    Up      40000     1   vcp-255/0/6
0/8         Auto-Configured     5    Up      40000     1   vcp-255/0/8
0/9         Auto-Configured     6    Up      100000    1   vcp-255/0/9

fpc1:
```

```
-
Interface   Type              Trunk  Status  Speed     Neighbor
or                             ID            (mbps)    ID  Interface
PIC / Port
0/13        Auto-Configured     6    Up      100000    0   vcp-255/0/13
0/23        Auto-Configured     6    Up      100000    0   vcp-255/0/23
0/26        Auto-Configured     5    Up      40000     0   vcp-255/0/26
0/27        Configured          5    Up      40000     0   vcp-255/0/27
0/6         Configured          5    Up      40000     0   vcp-255/0/6
0/8         Auto-Configured     5    Up      40000     0   vcp-255/0/8
0/9         Auto-Configured     6    Up      100000    0   vcp-255/0/9
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show virtual-chassis vc-port statistics | 27713

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

show virtual-chassis vc-port diagnostics optics

IN THIS SECTION

- Syntax | 27689
- Description | 27689
- Options | 27689
- Required Privilege Level | 27689
- Output Fields | 27690
- Sample Output | 27693
- Release Information | 27708

Syntax

```
show virtual-chassis vc-port diagnostics optics
<all-members | local | member member-id>
<interface-name>
```

Description

Display diagnostics data and alarms for Ethernet optical transceivers installed in ports configured as Virtual Chassis Ports (VCPs) in an EX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that a transceiver is not operating properly. DOM information can be used to diagnose why a transceiver is not working.

On some EX Series switches, the `request virtual-chassis vc-port diagnostics optics` command must be entered to run a diagnostic scan before you can gather the `show virtual-chassis vc-port diagnostics optics` output.

Options

none	Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.
all-members	(Optional) Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.
interface-name	(Optional) Display diagnostics information for the transceiver installed in a specified VCP.
local	(Optional) Display diagnostics information for transceivers installed in VCPs on the switch or external Routing Engine where you enter this command.
member ***member-id***	(Optional) Display diagnostics information for transceivers installed in VCPs on a specified member of a Virtual Chassis or VCF.

Required Privilege Level

view

Output Fields

Table 1773 on page 27690 lists the output fields for the `show virtual-chassis vc-port diagnostics optics` command. Output fields are listed in the approximate order in which they appear.

Table 1773: show virtual-chassis vc-port diagnostics optics Output Fields

Field Name	Field Description
FPC	Displays the FPC slot number.
Virtual chassis port	Displays the name of the VCP.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in Volts.
Receiver signal average optical power	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is *On* or *Off*.
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is *On* or *Off*.
Laser bias current high warning	Displays whether the laser bias power setting high warning is *On* or *Off*.
Laser bias current low warning	Displays whether the laser bias power setting low warning is *On* or *Off*.
Laser output power high alarm	Displays whether the laser output power high alarm is *On* or *Off*.

Table 1773: show virtual-chassis vc-port diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Laser output power low alarm	Displays whether the laser output power low alarm is *On* or *Off*.
Laser output power high warning	Displays whether the laser output power high warning is *On* or *Off*.
Laser output power low warning	Displays whether the laser output power low warning is *On* or *Off*.
Module temperature high alarm	Displays whether the module temperature high alarm is *On* or *Off*.
Module temperature low alarm	Displays whether the module temperature low alarm is *On* or *Off*.
Module temperature high warning	Displays whether the module temperature high warning is *On* or *Off*.
Module temperature low warning	Displays whether the module temperature low warning is *On* or *Off*.
Module voltage high alarm	Displays whether the module voltage high alarm is *On* or *Off*.
Module voltage low alarm	Displays whether the module voltage low alarm is *On* or *Off*.
Module voltage high warning	Displays whether the module voltage high warning is *On* or *Off*.
Module voltage low warning	Displays whether the module voltage low warning is *On* or *Off*.
Laser rx power high alarm	Displays whether the receive laser power high alarm is *On* or *Off*.
Laser rx power low alarm	Displays whether the receive laser power low alarm is *On* or *Off*.
Laser rx power high warning	Displays whether the receive laser power high warning is *On* or *Off*.
Laser rx power low warning	Displays whether the receive laser power low warning is *On* or *Off*.

Table 1773: show virtual-chassis vc-port diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.

Table 1773: show virtual-chassis vc-port diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.
Module voltage high alarm threshold	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser rx power low warning.

Sample Output

show virtual-chassis vc-port diagnostics optics

```
user@switch> show virtual-chassis vc-port diagnostics optics
fpc0:
```

```
---------------------------------------------------------------------------
Virtual chassis port: vcp-0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-1
    Optical diagnostics                       :  N/A

fpc1:
---------------------------------------------------------------------------
Virtual chassis port: vcp-0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-1
    Optical diagnostics                       :  N/A

fpc2:
---------------------------------------------------------------------------
Virtual chassis port: vcp-2/0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-2/1
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/14
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/15
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/24
    Laser bias current                        :  4.130 mA
    Laser output power                        :  0.2450 mW / -6.11 dBm
    Module temperature                        :  32 degrees C / 90 degrees F
    Module voltage                            :  3.3530 V
    Receiver signal average optical power     :  0.0971 mW / -10.13 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
```

```
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  14.998 mA
    Laser bias current low alarm threshold    :  0.998 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  1.198 mA
    Laser output power high alarm threshold   :  0.7940 mW / -1.00 dBm
    Laser output power low alarm threshold    :  0.0790 mW / -11.02 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold  :  0.0990 mW / -10.04 dBm
    Module temperature high alarm threshold   :  85 degrees C / 185 degrees F
    Module temperature low alarm threshold    :  -10 degrees C / 14 degrees F
    Module temperature high warning threshold :  80 degrees C / 176 degrees F
    Module temperature low warning threshold  :  -5 degrees C / 23 degrees F
    Module voltage high alarm threshold       :  3.600 V
    Module voltage low alarm threshold        :  3.000 V
    Module voltage high warning threshold     :  3.499 V
    Module voltage low warning threshold      :  3.099 V
    Laser rx power high alarm threshold       :  1.5848 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  1.2589 mW / 1.00 dBm
    Laser rx power low warning threshold      :  0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.428 mA
    Laser output power                        :  0.4760 mW / -3.22 dBm
    Module temperature                        :  28 degrees C / 83 degrees F
    Module voltage                            :  3.3440 V
    Receiver signal average optical power     :  0.4002 mW / -3.98 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
```

```
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm

fpc3:
--------------------------------------------------------------------------
Virtual chassis port: vcp-255/0/2
    Laser bias current                        :  7.876 mA
    Laser output power                        :  0.5330 mW / -2.73 dBm
    Module temperature                        :  26 degrees C / 78 degrees F
    Module voltage                            :  3.3060 V
    Receiver signal average optical power     :  0.4885 mW / -3.11 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
```

```
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  14.500 mA
    Laser bias current low alarm threshold    :  3.500 mA
    Laser bias current high warning threshold :  14.500 mA
    Laser bias current low warning threshold  :  3.500 mA
    Laser output power high alarm threshold   :  1.8620 mW / 2.70 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7410 mW / -1.30 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.9952 mW / 3.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.052 mA
    Laser output power                        :  0.5030 mW / -2.98 dBm
    Module temperature                        :  24 degrees C / 75 degrees F
    Module voltage                            :  3.2890 V
    Receiver signal average optical power     :  0.5028 mW / -2.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
```

```
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/4
    Laser bias current                        :  7.978 mA
    Laser output power                        :  0.5460 mW / -2.63 dBm
    Module temperature                        :  24 degrees C / 76 degrees F
    Module voltage                            :  3.3060 V
```

```
    Receiver signal average optical power     :  0.6305 mW / -2.00 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  14.500 mA
    Laser bias current low alarm threshold    :  3.500 mA
    Laser bias current high warning threshold :  14.500 mA
    Laser bias current low warning threshold  :  3.500 mA
    Laser output power high alarm threshold   :  1.8620 mW / 2.70 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7410 mW / -1.30 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.9952 mW / 3.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
fpc4:
--------------------------------------------------------------------------
```

```
Virtual chassis port: vcp-0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-1
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/4
    Laser bias current                        :  7.860 mA
    Laser output power                        :  0.5370 mW / -2.70 dBm
    Module temperature                        :  24 degrees C / 75 degrees F
    Module voltage                            :  3.2920 V
    Receiver signal average optical power     :  0.6271 mW / -2.03 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  14.500 mA
    Laser bias current low alarm threshold    :  3.500 mA
    Laser bias current high warning threshold :  14.500 mA
    Laser bias current low warning threshold  :  3.500 mA
    Laser output power high alarm threshold   :  1.8620 mW / 2.70 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7410 mW / -1.30 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
```

```
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.9952 mW / 3.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
```

show virtual-chassis vc-port diagnostics optics (interface-name)

```
user@external-routing-engine> show virtual-chassis vc-port diagnostics optics vcp-255/0/3
fpc0:
--------------------------------------------------------------------------

fpc1:
--------------------------------------------------------------------------

fpc2:
--------------------------------------------------------------------------
Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.448 mA
    Laser output power                        :  0.4770 mW / -3.21 dBm
    Module temperature                        :  28 degrees C / 82 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.3973 mW / -4.01 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
```

```
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm

fpc3:
--------------------------------------------------------------------------
Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.040 mA
    Laser output power                        :  0.5020 mW / -2.99 dBm
    Module temperature                        :  24 degrees C / 74 degrees F
    Module voltage                            :  3.2870 V
    Receiver signal average optical power     :  0.5073 mW / -2.95 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
```

```
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm

fpc4:
--------------------------------------------------------------------------
```

show virtual-chassis vc-port diagnostics optics local

```
user@switch> show virtual-chassis vc-port diagnostics optics local
Virtual chassis port: vcp-2/0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-2/1
    Optical diagnostics                       :  N/A
```

```
Virtual chassis port: vcp-255/0/14
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/15
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/24
    Laser bias current                        :  4.130 mA
    Laser output power                        :  0.2450 mW / -6.11 dBm
    Module temperature                        :  32 degrees C / 90 degrees F
    Module voltage                            :  3.3530 V
    Receiver signal average optical power     :  0.0961 mW / -10.17 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  14.998 mA
    Laser bias current low alarm threshold    :  0.998 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  1.198 mA
    Laser output power high alarm threshold   :  0.7940 mW / -1.00 dBm
    Laser output power low alarm threshold    :  0.0790 mW / -11.02 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold  :  0.0990 mW / -10.04 dBm
    Module temperature high alarm threshold   :  85 degrees C / 185 degrees F
    Module temperature low alarm threshold    :  -10 degrees C / 14 degrees F
    Module temperature high warning threshold :  80 degrees C / 176 degrees F
    Module temperature low warning threshold  :  -5 degrees C / 23 degrees F
    Module voltage high alarm threshold       :  3.600 V
```

```
    Module voltage low alarm threshold        :  3.000 V
    Module voltage high warning threshold     :  3.499 V
    Module voltage low warning threshold      :  3.099 V
    Laser rx power high alarm threshold       :  1.5848 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  1.2589 mW / 1.00 dBm
    Laser rx power low warning threshold      :  0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.426 mA
    Laser output power                        :  0.4760 mW / -3.22 dBm
    Module temperature                        :  28 degrees C / 83 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.3955 mW / -4.03 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
```

```
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
```

show virtual-chassis vc-port diagnostics optics (member member-id)

```
user@switch> show virtual-chassis vc-port diagnostics optics member 2
fpc2:
--------------------------------------------------------------------------
Virtual chassis port: vcp-2/0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-2/1
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/14
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/15
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/24
    Laser bias current                        :  4.130 mA
    Laser output power                        :  0.2450 mW / -6.11 dBm
    Module temperature                        :  31 degrees C / 88 degrees F
    Module voltage                            :  3.3530 V
    Receiver signal average optical power     :  0.0961 mW / -10.17 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
```

```
    Module temperature high warning          :  Off
    Module temperature low warning           :  Off
    Module voltage high alarm                :  Off
    Module voltage low alarm                 :  Off
    Module voltage high warning              :  Off
    Module voltage low warning               :  Off
    Laser rx power high alarm                :  Off
    Laser rx power low alarm                 :  Off
    Laser rx power high warning              :  Off
    Laser rx power low warning               :  Off
    Laser bias current high alarm threshold  :  14.998 mA
    Laser bias current low alarm threshold   :  0.998 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  1.198 mA
    Laser output power high alarm threshold  :  0.7940 mW / -1.00 dBm
    Laser output power low alarm threshold   :  0.0790 mW / -11.02 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold  :  0.0990 mW / -10.04 dBm
    Module temperature high alarm threshold  :  85 degrees C / 185 degrees F
    Module temperature low alarm threshold   :  -10 degrees C / 14 degrees F
    Module temperature high warning threshold :  80 degrees C / 176 degrees F
    Module temperature low warning threshold  :  -5 degrees C / 23 degrees F
    Module voltage high alarm threshold      :  3.600 V
    Module voltage low alarm threshold       :  3.000 V
    Module voltage high warning threshold    :  3.499 V
    Module voltage low warning threshold     :  3.099 V
    Laser rx power high alarm threshold      :  1.5848 mW / 2.00 dBm
    Laser rx power low alarm threshold       :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold    :  1.2589 mW / 1.00 dBm
    Laser rx power low warning threshold     :  0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
    Laser bias current                       :  5.418 mA
    Laser output power                       :  0.4770 mW / -3.21 dBm
    Module temperature                       :  28 degrees C / 83 degrees F
    Module voltage                           :  3.3450 V
    Receiver signal average optical power    :  0.3964 mW / -4.02 dBm
    Laser bias current high alarm            :  Off
    Laser bias current low alarm             :  Off
    Laser bias current high warning          :  Off
    Laser bias current low warning           :  Off
    Laser output power high alarm            :  Off
    Laser output power low alarm             :  Off
    Laser output power high warning          :  Off
```

```
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

show virtual-chassis vc-port | 27681

Install a Transceiver

Remove a Transceiver

Junos OS Ethernet Interfaces Configuration Guide

show virtual-chassis vc-port (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 27709
- Description | 27709
- Options | 27709
- Required Privilege Level | 27710
- Output Fields | 27710
- Sample Output | 27711
- Release Information | 27713

Syntax

```
show virtual-chassis vc-port
<(all-members | local | member member-id)>
```

Description

Display the operational status of the Virtual Chassis ports for both member routers or for a specified member router in an MX Series Virtual Chassis configuration. You can issue the `show virtual-chassis vc-port` command from the console of either member router in the Virtual Chassis.

Options

all-members (Optional) Display the operational status of the Virtual Chassis ports for both member routers in a Virtual Chassis configuration.

local (Optional) Display the operational status of the Virtual Chassis ports on the member router on which you are issuing the command. This is the default behavior if no options are specified.

member *member-id* (Optional) Display the operational status of the Virtual Chassis ports on the specified member router. Replace *member-id* with the value 0 or 1.

Required Privilege Level

view

Output Fields

Table 1774 on page 27710 lists the output fields for the `show virtual-chassis vc-port` command. Output fields are listed in the approximate order in which they appear.

Table 1774: show virtual-chassis vc-port Output Fields

Field Name	Field Description
`member`*n*	Member ID of the Virtual Chassis member router for which output is displayed.
`Interface or Slot/PIC/Port`	Location, in the format *slot*/*pic*/*port*, of the Virtual Chassis ports configured on the member router.
`Type`	Type of Virtual Chassis port. `Configured` indicates that the Virtual Chassis port is properly configured.
`Trunk ID`	Trunk ID value assigned to a link aggregation group (LAG) formed by the Virtual Chassis. A positive number indicates that a trunk exists. The value `-1` indicates that a trunk is not present.

Table 1774: show virtual-chassis vc-port Output Fields *(Continued)*

Field Name	Field Description
`Status`	State of the Virtual Chassis port interface: • `Up`—The port interface is up. • `Up (Resync-not-set)`—The port interface is up but the packet forwarding engine is not ready for forwarding. • `Down`—The port interface is down. • `Absent`—The port interface is absent.
`Speed (mbps)`	Speed, in megabits per second, of the Virtual Chassis port interface.
`Neighbor ID Interface`	Member IDs and Virtual Chassis port interfaces (in `vcp-`*`slot`*`/`*`pic`*`/`*`port`* format) that are connected to this member router.

Sample Output

show virtual-chassis vc-port all-members

```
{master:member1-re0}
user@host> show virtual-chassis vc-port all-members
member0:
--------------------------------------------------------------------------
Interface       Type              Trunk  Status       Speed        Neighbor
or                                 ID                (mbps)       ID  Interface
Slot/PIC/Port
5/0/0           Configured         -1    Up           10000        1   vcp-1/3/0

member1:
--------------------------------------------------------------------------
Interface       Type              Trunk  Status       Speed        Neighbor
or                                 ID                (mbps)       ID  Interface
Slot/PIC/Port
1/3/0           Configured         -1    Up           10000        0   vcp-5/0/0
```

show virtual-chassis vc-port local

```
{master:member1-re0}
user@host> show virtual-chassis vc-port local

Interface       Type              Trunk  Status       Speed        Neighbor
or                                 ID                 (mbps)       ID  Interface
Slot/PIC/Port
1/3/0           Configured         -1    Up           10000        0   vcp-5/0/0
```

show virtual-chassis vc-port member 0

```
{master:member1-re0}
user@host> show virtual-chassis vc-port member 0
member0:
--------------------------------------------------------------------------
Interface       Type              Trunk  Status       Speed        Neighbor
or                                 ID                 (mbps)       ID  Interface
Slot/PIC/Port
5/0/0           Configured         -1    Up           10000        1   vcp-1/3/0
```

show virtual-chassis vc-port (Packet Forwarding Engine is not Ready)

```
{master:member1-re0}
user@host> show virtual-chassis vc-port
member1:
--------------------------------------------------------------------------
Interface      Type          Trunk     Status
Speed     Neighbor
or                           ID                                                    (mbps)
ID  Interface
Slot/PIC/Port
3/0/2          Configured     -1        Up                                          10000
1   vcp-1/1/4
3/0/3          Configured     -1        Up (Resync-not-set)      10000     1   vcp-1/1/0
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Verifying the Operation of Virtual Chassis Ports

Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis

Example: Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis

show virtual-chassis vc-port statistics

IN THIS SECTION

- Syntax | 27713
- Description | 27714
- Options | 27714
- Required Privilege Level | 27714
- Output Fields | 27714
- Sample Output | 27718
- Release Information | 27722

Syntax

```
show virtual-chassis vc-port statistics
<all-members | local | member member-id>
<brief | detail | extensive >
<interface-name>
```

Description

Display the traffic statistics collected on Virtual Chassis ports (VCPs).

Options

none	Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF.
brief | detail | extensive	(Optional) Display the specified level of output. Using the `brief` option is equivalent to entering the command with no options (the default). The `detail` and `extensive` options provide identical displays.
all-members	(Optional) Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF
interface-name	(Optional) Display traffic statistics for the specified VCP.
local	(Optional) Display traffic statistics for VCPs on the switch or external Routing Engine where you enter this command.
member *member-id*	(Optional) Display traffic statistics for VCPs on the specified member of a Virtual Chassis or VCF.

Required Privilege Level

view

Output Fields

Table 1775 on page 27714 lists the output fields for the `show virtual-chassis vc-port statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1775: show virtual-chassis vc-port statistics Output Fields

Field Name	Field Description	Level of Output
`fpc`*`number`*	(All Virtual Chassis except EX8200 Virtual Chassis. VCF) ID of the Virtual Chassis member. The FPC number is the same as the member ID.	All levels

Table 1775: show virtual-chassis vc-port statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
member *number*	(EX8200 Virtual Chassis only) Member ID of the Virtual Chassis member.	All levels
Interface	VCP name.	brief
Input Octets/Packets	Number of octets and packets received on the VCP.	brief, member, none
Output Octets/ Packets	Number of octets and packets transmitted on the VCP.	brief, member, none
master: *number*	Member ID of the primary Routing Engine.	All levels
Port	VCP for which RX (Receive) statistics,TX (Transmit) statistics, or both are reported by the VCP subsystem during a sampling interval—since the statistics counter was last cleared.	detail, extensive
Total octets	Total number of octets received and transmitted on the VCP.	detail, extensive
Total packets	Total number of packets received and transmitted on the VCP.	detail, extensive
Unicast packets	Number of unicast packets received and transmitted on the VCP.	detail, extensive
Broadcast packets	Number of broadcast packets received and transmitted on the VCP.	detail, extensive
Multicast packets	Number of multicast packets received and transmitted on the VCP.	detail, extensive

Table 1775: show virtual-chassis vc-port statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`MAC control frames`	Number of media access control (MAC) control frames received and transmitted on the VCP.	`detail`, `extensive`
`CRC alignment errors`	Number of packets received on the VCP that had a length—excluding framing bits, but including frame check sequence (FCS) octets—of between 64 and 1518 octets, inclusive, and had one of the following errors: • Invalid FCS with an integral number of octets (FCS error) • Invalid FCS with a nonintegral number of octets (alignment error)	`detail`, `extensive`
`Oversize packets`	Number of packets received on the VCP that were longer than 1518 octets (excluding framing bits, but including FCS octets) but were otherwise well formed.	`detail`, `extensive`
`Undersize packets`	Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits but including FCS octets) and were otherwise well formed..	`detail`, `extensive`
`Jabber packets`	Number of packets received on the VCP that were longer than 1518 octets—excluding framing bits, but including FCS octets—and that had either an FCS error or an alignment error. **NOTE**: This definition of *jabber* is different from the definition in IEEE-802.3 section 8.2.1.5 (10Base5) and section 10.3.1.4 (10Base2). These documents define *jabber* as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.	`detail`, `extensive`

Table 1775: show virtual-chassis vc-port statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Fragments received`	Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.	`detail`, `extensive`
`Ifout errors`	Number of outbound packets received on the VCP that could not be transmitted because of errors.	`detail`, `extensive`
`Packet drop events`	Number of outbound packets received on the VCP that were dropped, rather than being encapsulated and sent out of the switch as fragments. The packet drop counter is incremented if a temporary shortage of packet memory causes packet fragmentation to fail.	`detail`, `extensive`
`64 octets frames`	Number of packets received on the VCP (including invalid packets) that were 64 octets in length (excluding framing bits, but including FCS octets).	`detail`, `extensive`
`65-127 octets frames`	Number of packets received on the VCP (including invalid packets) that were between 65 and 127 octets in length, inclusive (excluding framing bits, but including FCS octets).	`detail`, `extensive`
`128-255 octets frames`	Number of packets received on the VCP (including invalid packets) that were between 128 and 255 octets in length, inclusive (excluding framing bits, but including FCS octets).	`detail`, `extensive`
`256-511 octets frames`	Number of packets received on the VCP (including invalid packets) that were between 256 and 511 octets in length, inclusive (excluding framing bits, but including FCS octets).	`detail`, `extensive`

Table 1775: show virtual-chassis vc-port statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`512-1023 octets frames`	Number of packets received on the VCP (including invalid packets) that were between 512 and 1023 octets in length, inclusive (excluding framing bits, but including FCS octets).	`detail`, `extensive`
`1024-1518 octets frames`	Number of packets received on the VCP (including invalid packets) that were between 1024 and 1518 octets in length, inclusive (excluding framing bits, but including FCS octets).	`detail`, `extensive`
`Rate packets per second`	Number of packets per second received and transmitted on the VCP.	`detail`, `extensive`
`Rate bytes per second`	Number of bytes per second received and transmitted on the VCP.	`detail`, `extensive`

Sample Output

show virtual-chassis vc-port statistics

```
user@switch> show virtual-chassis vc-port statistics
fpc0:
--------------------------------------------------------------------------
Interface              Input Octets/Packets          Output Octets/Packets
internal-0/24          0           / 0               0           / 0
internal-0/25          0           / 0               0           / 0
internal-1/26          0           / 0               0           / 0
internal-1/27          0           / 0               0           / 0
vcp-0                  0           / 0               0           / 0
vcp-1                  0           / 0               0           / 0
internal-0/26          0           / 0               0           / 0
internal-0/27          0           / 0               0           / 0
internal-1/24          0           / 0               0           / 0
internal-1/25          0           / 0               0           / 0
```

```
{master:0}
```

show virtual-chassis vc-port statistics (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis vc-port statistics
member0:
--------------------------------------------------------------------------
Interface               Input Octets/Packets          Output Octets/Packets
vcp-4/0/4               43171238    / 48152           47687133    / 51891
vcp-4/0/7               0           / 0               0           / 0

member1:
--------------------------------------------------------------------------
Interface               Input Octets/Packets          Output Octets/Packets
vcp-3/0/0               0           / 0               0           / 0
vcp-3/0/1               0           / 0               0           / 0
vcp-3/0/4               47695376    / 51899           43180556    / 48160

member8:
--------------------------------------------------------------------------

member9:
--------------------------------------------------------------------------
```

show virtual-chassis vc-port statistics brief

```
user@switch> show virtual-chassis vc-port statistics brief
fpc0:
--------------------------------------------------------------------------
Interface               Input Octets/Packets          Output Octets/Packets
internal-0/24           0           / 0               0           / 0
internal-0/25           0           / 0               0           / 0
internal-1/26           0           / 0               0           / 0
internal-1/27           0           / 0               0           / 0
vcp-0                   0           / 0               0           / 0
vcp-1                   0           / 0               0           / 0
internal-0/26           0           / 0               0           / 0
internal-0/27           0           / 0               0           / 0
internal-1/24           0           / 0               0           / 0
```

```
internal-1/25          0           / 0               0           / 0

{master:0}
```

show virtual-chassis vc-port statistics extensive

```
user@switch> show virtual-chassis vc-port statistics extensive
fpc0:
--------------------------------------------------------------------------

                            RX                      TX

Port: internal-0/24
Total octets:               0                       0
Total packets:              0                       0
Unicast packets:            0                       0
Broadcast packets:          0                       0
Multicast packets:          0                       0
MAC control frames:         0                       0
CRC alignment errors:       0
Oversize packets:           0
Undersize packets:          0
Jabber packets:             0
Fragments received:         0
Ifout errors:               0
Packet drop events:         0
64        octets frames:    0
65-127    octets frames:    0
128-255   octets frames:    0
256-511   octets frames:    0
512-1023  octets frames:    0
1024-1518 octets frames:    0
Rate packets per second:    0                       0
Rate bytes per second:      0                       0

...

Port: vcp-0
Total octets:               0                       0
Total packets:              0                       0
Unicast packets:            0                       0
```

```
Broadcast packets:         0                       0
Multicast packets:         0                       0
MAC control frames:        0                       0
CRC alignment errors:      0
Oversize packets:          0
Undersize packets:         0
Jabber packets:            0
Fragments received:        0
Ifout errors:              0
Packet drop events:        0
64        octets frames:   0
65-127    octets frames:   0
128-255   octets frames:   0
256-511   octets frames:   0
512-1023  octets frames:   0
1024-1518 octets frames:   0
Rate packets per second:   0                       0
Rate bytes per second:     0                       0

Port: vcp-1
Total octets:              0                       0
Total packets:             0                       0
Unicast packets:           0                       0
Broadcast packets:         0                       0
Multicast packets:         0                       0
MAC control frames:        0                       0
CRC alignment errors:      0
Oversize packets:          0
Undersize packets:         0
Jabber packets:            0
Fragments received:        0
Ifout errors:              0
Packet drop events:        0
64        octets frames:   0
65-127    octets frames:   0
128-255   octets frames:   0
256-511   octets frames:   0
512-1023  octets frames:   0
1024-1518 octets frames:   0
Rate packets per second:   0                       0
Rate bytes per second:     0                       0

...
```

```
{master:0}
```

show virtual-chassis vc-port statistics member 0

```
user@switch>show virtual-chassis vc-port statistics member 0
fpc0:
--------------------------------------------------------------------------
Interface              Input Octets/Packets          Output Octets/Packets
internal-0/24          0           / 0               0           / 0
internal-0/25          0           / 0               0           / 0
internal-1/26          0           / 0               0           / 0
internal-1/27          0           / 0               0           / 0
vcp-0                  0           / 0               0           / 0
vcp-1                  0           / 0               0           / 0
internal-0/26          0           / 0               0           / 0
internal-0/27          0           / 0               0           / 0
internal-1/24          0           / 0               0           / 0
internal-1/25          0           / 0               0           / 0

{master:0}
```

Release Information

Command introduced in Junos OS Release 9.0.

The options `all-members`, `brief`, `detail`, `extensive`, and `local` were added in Junos OS Release 9.3 for EX Series switches.

RELATED DOCUMENTATION

- clear virtual-chassis vc-port statistics | **4283**
- show virtual-chassis vc-port | **27681**
- Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis
- Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis

show virtual-network-functions

IN THIS SECTION

- Syntax | 27723
- Description | 27723
- Options | 27723
- Required Privilege Level | 27724
- Output Fields | 27724
- Sample Output | 27725
- Sample Output | 27726
- Sample Output | 27726
- Sample Output | 27727
- Release Information | 27728

Syntax

```
show virtual-network-functions
show virtual-network-functions vnf-name
show virtual-network-functions ipsec-nm
show virtual-network-functions vjunos0
```

Description

Display Virtual Network Function (VNF) information.

Options

vnf-name	(Optional) Display information for a specific VNF.
ipsec-nm	(Optional) Display information of the system VNF IPsec.
vjunos0	(Optional) Display information of the system VNF vjunos0.

brief (Optional) Display brief output.

> **NOTE**: This is the default option.

detail (Optional) Display detailed output.

Required Privilege Level

view

Output Fields

Table 1776 on page 27724 describes the output fields for the `show virtual-network-functions` command. Output fields are listed in the approximate order in which they appear.

Table 1776: show virtual-network functions Output Fields

Field Name	Field Description
ID	ID of the VNF
Name	Name of the VNF
State	Status of the VNF. Possible values are Running, Shutdown, or Undefined.
Liveliness	Indicates whether or not the IP address of the VNF is reachable.
IP Address	IP address of the VNF
VCPUs	Number of virtual CPUs
Maximum Memory	Maximum amount of memory available to the VNF
Used Memory	Amount of memory used by the VNF

Table 1776: show virtual-network functions Output Fields *(Continued)*

Field Name	Field Description
Used 1G Hugepages	The number of used 1G hugepages.
Used 2M Hugepages	The number of used 2M hugepages

Sample Output

show virtual-network-functions vjunos0

```
user@host> show virtual-network-functions vjunos0
ID       Name                                            State      Liveliness
--------------------------------------------------------------------------------
2        vjunos0                                         Running    alive
```

show virtual-network-functions vjunos0 detail

```
user@host> show virtual-network-functions vjunos0 detail
Virtual Network Function Information
------------------------------------

Id:                 3
Name:               vjunos0
State:              Running
Liveliness:         Up
IP Address:         192.0.2.2
VCPUs:              1
Maximum Memory:     1000448 KiB
Used Memory:        1000448 KiB
Used 1G Hugepages:  0
Used 2M Hugepages:  0

     Block Devices
     ----------------------------
     Target   | Source
```

```
    ----------------------------
    vda      | /junos/images/0/vjunos.img
    vdb      | /junos/images/0/vjunos-data.img
    hdc      | /junos/images/shared/swap-disk.img
    vdd      | /junos/images/0/vjunos-platform.img
    vde      | /junos/images/0/metadata-usb-re.img
```

Sample Output

show virtual-network-functions vsrx

```
user@host> show virtual-network-functions vsrx
ID       Name                                            State      Liveliness
--------------------------------------------------------------------------------
67       vsrx                                            Running    alive
```

Sample Output

show virtual-network-functions vmx

```
user@host> show virtual-network-functions vmx
ID       Name                                            State      Liveliness
--------------------------------------------------------------------------------
3       vjunos0                                           Running   alive
10      vm    x
Running        alive
11341     jdm                                                            Running        alive
```

show virtual-network-functions

```
user@host> show virtual-network-functions
ID       Name                                            State      Liveliness
--------------------------------------------------------------------------------
3        vjunos0                                         Running    alive
1        LTE-VM                                          Running    down
```

Sample Output

show virtual-network-functions detail

```
user@host> show virtual-network-functions detail
Virtual Network Function Information
------------------------------------

Id:                 3
Name:               vjunos0
State:              Running
Liveliness:         Up
IP Address:         192.0.2.2
VCPUs:              1
Maximum Memory:     1000448 KiB
Used Memory:        1000448 KiB
Used 1G Hugepages:  0
Used 2M Hugepages:  0

     Block Devices
     ----------------------------
     Target   | Source
     ----------------------------
     vda      | /junos/images/0/vjunos.img
     vdb      | /junos/images/0/vjunos-data.img
     hdc      | /junos/images/shared/swap-disk.img
     vdd      | /junos/images/0/vjunos-platform.img
     vde      | /junos/images/0/metadata-usb-re.img

Virtual Network Function Information
------------------------------------

Id:                 1
Name:               LTE-VM
State:              Running
Liveliness:         Down
IP Address:         -
VCPUs:              2
Maximum Memory:     122880 KiB
Used Memory:        122880 KiB
Used 1G Hugepages:  0
Used 2M Hugepages:  0
```

```
Block Devices
----------------------------
Target    | Source
----------------------------
vdb       | /usr/share/juniper/LTE/lte_vm.latest
```

Release Information

Command introduced in Junos OS Release 15.1X53-D45.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions (Junos Node Slicing)

IN THIS SECTION

- Syntax | 27729
- Description | 27729
- Options | 27729
- Required Privilege Level | 27729
- Output Fields | 27729
- Sample Output | 27731
- Sample Output | 27732
- Sample Output | 27732
- Release Information | 27738

Syntax

```
show virtual-network-functions (all-servers | server| slot)
<vnf-name>
<vnf-name detail>
<vnf-name vnf-xml-profile>
```

Description

Display the list of guest network functions (GNFs) along with their IDs, status and availability.

Options

all-servers Display the details of the GNFs on both the servers.

server Display the details of the GNFs on one specific server. Applicable value is `0` or `1`.

vnf-name Display additional details of a particular GNF. You can use the `detail` option to view the detailed output. For example, `show virtual-network-functions gnf1 detail`.

vnf-xml-profile Displays the xml profile of a GNF.

Required Privilege Level

View

Output Fields

Table 1777 on page 27729 lists the output fields for the `show virtual-network-functions` command.

Table 1777: show virtual network functions Output Fields

Field Name	Field Description
`ID`	The ID associated with the VNF.

Table 1777: show virtual network functions Output Fields *(Continued)*

Field Name	Field Description
`Name`	Name of the VNF.
`State`	Status of the VNF. • `Running`—The VNF is online and running. • `Shut off`—The VNF is in the shut down state.
`Liveness`	Indicates the availability of VNF. • `Up` • `Down`
`Name`	Name of the VNF.
`IP Address`	IP address of the VNF.
`Status`	Status of a particular VNF. • `Running`—The VNF is online and running. • `Shut off`—The VNF is in the shut down state.
`Cores`	Number of cores in the VNF.
`Memory`	The VNF memory.
`Resource Template`	The resource template associated with a VNF.
`Qemu Process id`	Qemu process ID.

Table 1777: show virtual network functions Output Fields *(Continued)*

VNF CPU Utilization and Allocation Information	Shows the GNF CPU utilization details. See also: "show system cpu (JDM)" on page 26804.
VNF Memory Information	Displays the following memory information about the GNFs: • Name—GNF name. • Resident—The memory used by the GNFs. • Actual—Actual memory.
VNF Storage Information	Displays the following guest network function (GNF) storage information: • Directories—Names of the directories. • Size—Total storage size. • Used—Storage used.
VNF Interfaces Statistics	Shows the GNF interface statistics information. See also: "show system network (JDM)" on page 26931.
VNF Network Information	Shows the list of Physical Interfaces, Virtual Interfaces and MAC addresses.

Sample Output

show virtual-network-functions

```
user@jdm> show virtual-network-functions


ID        Name                                              State      Liveness
--------------------------------------------------------------------------------
1         bittern-gnf-a                                     Running    Up
2         bittern-gnf-b                                     Running    Up
3         bittern-gnf-c                                     Running    Up
4         bittern-gnf-d                                     Running    Up
```

```
5        bittern-gnf-e                                    Running    Up
```

Sample Output

show virtual-network-functions (for a specific VNF)

```
user@jdm> show virtual-network-functions gnf1


Virtual Machine Information
---------------------------
Name:                gnf1
IP Address:          192.168.2.1
Status:              Running
Liveness:            up
Cores:               4
Memory:              32GB
Resource Template:   4core-32g
Qemu Process id:     10891
```

Sample Output

show virtual-network-functions <vnf-name> detail

```
user@jdm> show virtual-network-functions gnf1 detail


VNF Information
---------------------------
ID                   1
Name:                gnf1
Status:              Running
Liveness:            up
IP Address:          192.168.2.1
Cores:               2
Memory:              16GB
Resource Template:   2core-16g
Qemu Process id:     20478
```

```
SMBIOS version:     v1

VNF Uptime: 206:09.58


VNF CPU Utilization and Allocation Information
-------------------------------------------------------------------------------
GNF                                      CPU-Id(s)               Usage  Qemu Pid
---------------------------------------- ----------------------- -----  --------
gnf1                                     10,11                   6.1%   20478


VNF Memory Information
-----------------------------------------------------------------
Name                                                Actual Resident
------------------------------------------------- ------ --------
gnf1                                                16.0G  15.5G


VNF Storage Information
-------------------------------------------------------
Directory                                    Size   Used
------------------------------------------- ------ ------
/vm-primary/gnf1                             50.2G  4.2G


VNF Interfaces Statistics
-----------------------------------------------------------------------------------------------------
------------------------
Interface                   Rcvd Bytes    Rcvd packets Rcvd Error Rcvd Drop Trxd bytes   Trxd
Packets Trxd Error Trxd Drop
-------------------------- ------------  ------------ ---------- --------- ------------
------------ ---------- ---------
macvtap0                    19077011      335687       0          0         261601960
3545272      0          0
macvtap1                    99369778      689729       0          0         172763580
2786344      0          0
vnet1                       24686         527          0          0         11620
250          0          0
macvtap2                    428385        7405         0          0         861020752
10813152     0          0
```

```
VNF Network Information
--------------------------------------------------------------------------------
Virtual Interface            Physical Interface    MAC
--------------------------   ------------------    ----------------------------
macvtap0                     ens3f0                00:85:49:54:cd:30
macvtap1                     ens3f1                00:85:49:54:cd:31
vnet1                        bridge_jdm_vm         00:85:49:54:cd:32
macvtap2                     enp3s0f2              00:85:49:54:cd:33
```

show virtual-network-functions <vnf-name> vnf-xml-profile

```
user@jdm> show virtual-network-functions gnf1 vnf-xml-profile

<domain type='kvm' id='1'>
  <name>gnf1</name>
  <uuid>a1c62c35-910f-4354-991c-7ad26c1b32e7</uuid>
  <memory unit='KiB'>33554432</memory>
  <currentMemory unit='KiB'>33554432</currentMemory>
  <vcpu placement='static'>4</vcpu>
  <cputune>
    <vcpupin vcpu='0' cpuset='4'/>
    <vcpupin vcpu='1' cpuset='5'/>
    <vcpupin vcpu='2' cpuset='6'/>
    <vcpupin vcpu='3' cpuset='7'/>
    <emulatorpin cpuset='2-3'/>
  </cputune>
  <resource>
    <partition>/machine</partition>
  </resource>
  <sysinfo type='smbios'>
    <system>
      <entry name='manufacturer'>JUNIPER</entry>
      <entry name='product'>RE-GNF</entry>
      <entry name='version'>v2-mx480-gnf1-re0-vtnet0-vtnet1-vtnet3-vtnet2-
external-0x009069136000:0x0800</entry>
    </system>
  </sysinfo>
  <os>
    <type arch='x86_64' machine='pc-i440fx-rhel7.0.0'>hvm</type>
    <boot dev='hd'/>
```

```
    <smbios mode='sysinfo'/>
  </os>
  <features>
    <acpi/>
    <apic/>
  </features>
  <clock offset='utc'/>
  <on_poweroff>destroy</on_poweroff>
  <on_reboot>restart</on_reboot>
  <on_crash>restart</on_crash>
  <devices>
    <emulator>/usr/libexec/qemu-kvm</emulator>
    <disk type='file' device='disk'>
      <driver name='qemu' type='raw' cache='directsync' io='native'/>
      <source file='/vm-primary/gnf1/gnf1.img'/>
      <backingStore/>
      <target dev='hda' bus='virtio'/>
      <alias name='virtio-disk0'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x07' function='0x0'/>
    </disk>
    <disk type='file' device='disk'>
      <driver name='qemu' type='raw' cache='directsync' io='native'/>
      <source file='/vm-primary/gnf1/var-config.img'/>
      <backingStore/>
      <target dev='vdb' bus='virtio'/>
      <alias name='virtio-disk1'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x08' function='0x0'/>
    </disk>
    <disk type='file' device='disk'>
      <driver name='qemu' type='raw' cache='directsync' discard='unmap'/>
      <source file='/vm-primary/gnf1/swap-disk.img'/>
      <backingStore/>
      <target dev='hdc' bus='virtio'/>
      <alias name='virtio-disk2'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x09' function='0x0'/>
    </disk>
    <disk type='file' device='disk'>
      <driver name='qemu' type='raw' cache='directsync' discard='unmap'/>
      <source file='/vm-primary/gnf1/aux-disk.img'/>
      <backingStore/>
      <target dev='hdb' bus='ide'/>
      <alias name='ide0-0-1'/>
      <address type='drive' controller='0' bus='0' target='0' unit='1'/>
```

```
    </disk>
    <disk type='file' device='disk'>
      <driver name='qemu' type='raw' cache='directsync'/>
      <source file='/vm-primary/gnf1/usb-disk.img'/>
      <backingStore/>
      <target dev='sda' bus='usb'/>
      <alias name='usb-disk0'/>
      <address type='usb' bus='0' port='1'/>
    </disk>
    <controller type='usb' index='0'>
      <alias name='usb'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x01' function='0x2'/>
    </controller>
    <controller type='pci' index='0' model='pci-root'>
      <alias name='pci.0'/>
    </controller>
    <controller type='ide' index='0'>
      <alias name='ide'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x01' function='0x1'/>
    </controller>
    <interface type='direct' trustGuestRxFilters='yes'>
      <mac address='52:54:00:09:ab:83'/>
      <source dev='ens2f0' mode='vepa'/>
      <target dev='macvtap0'/>
      <model type='virtio'/>
      <alias name='net0'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x03' function='0x0'/>
    </interface>
    <interface type='direct' trustGuestRxFilters='yes'>
      <mac address='52:54:00:45:e3:ba'/>
      <source dev='ens2f1' mode='vepa'/>
      <target dev='macvtap1'/>
      <model type='virtio'/>
      <alias name='net1'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0'/>
    </interface>
    <interface type='bridge'>
      <mac address='52:54:00:ee:73:93'/>
      <source bridge='bridge_jdm_vm'/>
      <target dev='vnet1'/>
      <model type='virtio'/>
      <alias name='net2'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x05' function='0x0'/>
```

```
    </interface>
    <interface type='direct'>
      <mac address='00:90:69:13:7f:00'/>
      <source dev='eno2' mode='bridge'/>
      <bandwidth>
        <inbound average='125000' peak='125000' burst='256'/>
        <outbound average='125000' peak='125000' burst='256'/>
      </bandwidth>
      <target dev='macvtap2'/>
      <model type='virtio'/>
      <driver name='qemu'/>
      <alias name='net3'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x06' function='0x0'/>
    </interface>
    <serial type='pty'>
      <source path='/dev/pts/6'/>
      <target port='0'/>
      <alias name='serial0'/>
    </serial>
    <console type='pty' tty='/dev/pts/6'>
      <source path='/dev/pts/6'/>
      <target type='serial' port='0'/>
      <alias name='serial0'/>
    </console>
    <input type='mouse' bus='ps2'>
      <alias name='input0'/>
    </input>
    <input type='keyboard' bus='ps2'>
      <alias name='input1'/>
    </input>
    <memballoon model='virtio'>
      <alias name='balloon0'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x0a' function='0x0'/>
    </memballoon>
  </devices>
  <seclabel type='none' model='none'/>
  <seclabel type='dynamic' model='dac' relabel='yes'>
    <label>+107:+107</label>
    <imagelabel>+107:+107</imagelabel>
  </seclabel>
</domain>
```

Release Information

Command introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

request virtual-network-functions | **15481**

request server authenticate-peer-server | **14901**

show vlans

IN THIS SECTION

- Syntax (EX Series and QFX Series Switches) | **27738**
- Syntax (EX Series with ELS Switches and MX Routers) | **27739**
- Syntax (SRX Series Firewalls) | **27739**
- Description | **27739**
- Options | **27740**
- Required Privilege Level | **27741**
- Output Fields | **27741**
- Sample Output (EX Series and QFX Series Switches) | **27745**
- Sample Output: EX Series with ELS Switches and MX Routers | **27758**
- Sample Output (SRX Series Firewalls) | **27765**
- Release Information | **27767**

Syntax (EX Series and QFX Series Switches)

```
show vlans
<brief | detail | extensive>
<dot1q-tunneling>
<management-vlan>
```

```
<sort-by (tag | name)>
<vlan-range-name>
<summary>
<vlan-name>
<vlan-range-name>
```

Syntax (EX Series with ELS Switches and MX Routers)

```
show vlans
<brief | detail | extensive>
<instance instance-name>
<logical-system logical-system-name>
<operational>
<vlan-name)>
<interface interface-name>
```

Syntax (SRX Series Firewalls)

```
show vlans
<brief | detail | extensive>
<interface interface-name>
<logical-system (logical-system | all)>
<operational>
```

Description

Display information about VLANs configured on bridged Ethernet interfaces. For interfaces configured to support a VoIP VLAN and a data VLAN, the `show vlans` command displays both tagged and untagged membership for those VLANs.

NOTE: When a series of VLANs is created using the `vlan-range` statement, such VLAN names are preceded and followed by a double underscore. For example, a series of VLANs using the VLAN range 1 through 3 and the base VLAN name **marketing** would be displayed as **__marketing_1__**, **__marketing_2__**, and **__marketing_3__**.

> **NOTE**: To display an 802.1X supplicant successfully authenticated in multiple-supplicant mode with dynamic VLAN movement, use the `show vlans vlan-name extensive` operational mode command, where ***vlan-name*** is the dynamic VLAN.

Options

For EX Series and QFX Series switches:

none	Display information for all VLANs. VLAN information is displayed by VLAN name in ascending order.
brief | detail | extensive	(Optional) Display the specified level of output.
dot1q-tunneling	(Optional) Display VLANs with the Q-in-Q tunneling feature enabled.
management-vlan	(Optional) Display management VLANs.
sort-by (tag | name)	(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.
vlan-range-name	(Optional) Display VLANs in ascending order of VLAN range names.
summary	(Optional) Display the total number of VLANs and counts of VLANs by type—for example, the number of dynamic, 802.1Q-tagged, and Q-in-Q tunneled VLANs.
vlan-range-name	(Optional) Display information for the specified VLAN range. To display information for all members of the VLAN range, specify the base VLAN name—for example, **employee** for a VLAN range that includes **__employee_1__** through **__employee_10__**.

For EX Series with ELS Switches and MX Routers:

none	Display information for all VLANs.
brief | detail | extensive	(Optional) Display the specified level of output.
instance ***instance-name***	(Optional) Display information for the specified routing instance.
logical-system ***logical-system-name***	(Optional) Display Ethernet-switching statistics information for the specified logical system.
operational	(Optional) Display information for the operational routing instances.

vlan-name)	(Optional) Display information about the specified VLAN.
interface *interface-name*	(Optional) Display information about the specified interface.

For SRX Series Firewalls:

none	Display information for all VLANs.
brief | detail | extensive	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display information about a specific interface.
logical system	(Optional) Display name of the logical system or all.
operational	(Optional) Display information for the operational switching instances.

Required Privilege Level

view

Output Fields

Table 1778 on page 27741 lists the output fields for the `show vlans` command. Output fields are listed in the approximate order in which they appear.

Table 1778: show vlans Output Fields

Field Name	Field Description	Level of Output
Name	Name of a VLAN.	none, **brief**
Tag	802.1Q tag applied to this VLAN. If **none** is displayed, no tag is applied.	All levels
Interfaces	Interface associated with learned MAC addresses or **All-members** option (flood entry). An asterisk (*****) beside the interface indicates that the interface is **UP**.	All levels
Address	IP address.	none, **brief**

Table 1778: show vlans Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Ports Active /Total	Number of interfaces associated with a VLAN: **Active** indicates interfaces that are **UP**, and **Total** indicates interfaces that are active and inactive.	**brief**
VLAN	Name of a VLAN.	**detail**, **extensive**
Admin state	State of the interface. Values are: **enabled**—The interface is turned on, and the physical link is operational and can pass packets.	**detail**,**extensive**
MAC learning Status	Indicates if MAC learning is disabled.	**detail**, **extensive**
Description	Description for the VLAN.	**detail**,**extensive**
Primary IP	Primary IP address associated with a VLAN.	**detail**
Number of interfaces	Number of interfaces associated with a VLAN. Both the total number of interfaces and the number of active interfaces associated with a VLAN are displayed.	**detail**, **extensive**
STP	Spanning tree associated with a VLAN.	**detail**,**extensive**
Tagged interfaces	Tagged interfaces with which a VLAN is associated.	**detail**,**extensive**
Untagged interfaces	Untagged interfaces with which a VLAN is associated.	**detail**. **extensive**
Dot1q Tunneling Status	Indicates if Q-in-Q tunneling is enabled.	**extensive**
Customer VLAN ranges	List of customer VLAN (C-VLAN) ranges associated with this service VLAN (S-VLAN).	**extensive**

Table 1778: show vlans Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Private VLAN Mode	The private VLAN mode for this VLAN. Values include **Primary**, **Isolated**, and **Community**.	**extensive**
Primary VLAN	Primary VLAN tag for this secondary VLAN.	**extensive**
Internal Index	VLAN index internal to Junos OS software.	**extensive**
Origin	Manner in which the VLAN was created: **static** or **learn**.	**extensive**
Protocol	Port-based VLAN or MAC-based VLAN. MAC-based protocol is displayed when VLAN assignment is done either statically or dynamically through 802.1X,	**extensive**
IP addresses	IP address associated with a VLAN.	**extensive**
Number of MAC entries	For MAC-based VLANs created either statically or dynamically, the MAC addresses associated with an interface.	**extensive**
Number of mapping rules	Number of mapping rules for Q-in-Q tunneling (**Push**) and VLAN translation (**Swap**).	
Secondary VLANs	Secondary VLANs associated with a primary VLAN.	**extensive**
Isolated VLANs	Isolated VLANs associated with a primary VLAN.	**extensive**
Community VLANs	Community VLANs associated with a primary VLAN.	**extensive**

Table 1778: show vlans Output Fields *(Continued)*

Field Name	Field Description	Level of Output
VLANs summary	VLAN counts: • **Total**—Total number of VLANs on the switch. • **Configured VLANs**—Number of VLANs that are based on user-configured settings. • **Internal VLANs**—Number of VLANs created by the system with no explicit configuration or protocol—for example, the **default** VLAN and the VLAN created when a trunk interface is not configured with native VLAN membership. • **Temporary VLANs**—Number of VLANs from the previous configuration that the system retains for a limited time after restart. Temporary VLANs are converted into one of the other types of VLAN, or are removed from the system if the current configuration does not require them.	All levels

Dot1q VLANs summary	802.1Q VLAN counts:
	• **Total**—Total number of 802.1Q-tagged and untagged VLANs on the switch.
	• **Tagged VLANs**—Number of 802.1Q-tagged VLANs.
	• **Untagged VLANs**—Number of untagged 802.1Q VLANs.
	• **Private VLAN**—Counts of the following kinds of 802.1Q private VLANs (PVLANs):
	  • **Primary VLANs**—Number of primary forwarding private VLANs.
	  • **Community VLANs**—Number of community transporting and forwarding private VLANs.
	  • **Isolated VLANs**—Number of isolated receiving and forwarding private VLANs.
	  • **Inter-switch-isolated VLANs**—Number of inter-switch isolated receiving and forwarding private VLANs.
		All levels

Table 1778: show vlans Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Dot1q Tunneled VLANs summary	Q-in-Q-tunneled VLAN counts: • **Total**—Total number of Q-in-Q-tunneled VLANs on the switch. • **Private VLAN**—Counts of primary, community, and isolated Q-in-Q-tunneled private VLANs (PVLANs).	All levels

Sample Output (EX Series and QFX Series Switches)

show vlans (EX Series and QFX Series)

```
user@switch>  show vlans

Name            Tag     Interfaces
default         None
                        xe-0/0/34.0, xe-0/0/33.0, xe-0/0/32.0, xe-0/0/31.0,
                        xe-0/0/30.0, xe-0/0/29.0, xe-0/0/28.0, xe-0/0/27.0,
                        xe-0/0/26.0, xe-0/0/25.0, xe-0/0/19.0, xe-0/0/18.0,
                        xe-0/0/17.0, xe-0/0/16.0, xe-0/0/15.0, xe-0/0/14.0,
                        xe-0/0/13.0, xe-0/0/11.0, xe-0/0/9.0, xe-0/0/8.0,
                        xe-0/0/3.0, xe-0/0/2.0, xe-0/0/1.0
v0001           1
                        xe-0/0/24.0, xe-0/0/23.0, xe-0/0/22.0, xe-0/0/21.0
v0002           2
                        None
v0003           3
                        None
v0004           4
                        None
v0005           5
                        None
```

show vlans (Private VLANs on EX and QFX Series)

```
user@switch>  show vlans

Name           Tag     Interfaces
__pvlan_pvlan_xe-0/0/46.0__
                       xe-0/0/44.0*, xe-0/0/46.0*
c1
                       xe-0/0/4.0*, xe-0/0/44.0*
c2
                       xe-0/0/28.0*, xe-0/0/44.0*
default
                       None
pvlan          500
                       xe-0/0/4.0*, xe-0/0/28.0*, xe-0/0/44.0*, xe-0/0/46.0*
```

show vlans brief (EX and QFX Series)

```
user@switch> show vlans brief
                                          Ports
Name           Tag     Address            Active/Total
default        None                       0/23
v0001          1                          0/4
v0002          2                          0/0
v0003          3                          0/0
v0004          4                          0/0
v0005          5                          0/0
v0006          6                          0/0
v0007          7                          0/0
v0008          8                          0/0
v0009          9                          0/0
v0010          10                         0/2
v0011          11                         0/0
v0012          12                         0/0
v0013          13                         0/0
v0014          14                         0/0
v0015          15                         0/0
v0016          16                         0/0
```

show vlans detail (EX Series and QFX Series)

```
user@switch> show vlans detail
VLAN: default, Tag: Untagged, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 23 (Active = 0)
  STP: None, RTG: None
  Untagged interfaces: xe-0/0/34.0, xe-0/0/33.0, xe-0/0/32.0, xe-0/0/31.0,
  xe-0/0/30.0, xe-0/0/29.0, xe-0/0/28.0, xe-0/0/27.0, xe-0/0/26.0,
  xe-0/0/25.0, xe-0/0/19.0, xe-0/0/18.0, xe-0/0/17.0, xe-0/0/16.0,
  xe-0/0/15.0, xe-0/0/14.0, xe-0/0/13.0, xe-0/0/11.0, xe-0/0/9.0, xe-0/0/8.0,
  xe-0/0/3.0, xe-0/0/2.0, xe-0/0/1.0,
  Tagged interfaces: None

VLAN: v0001, Tag: 802.1Q Tag 1, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 4 (Active = 0)
  Dot1q Tunneling Status: Enabled
  STP: None, RTG: None
  Untagged interfaces: None
  Tagged interfaces: xe-0/0/24.0, xe-0/0/23.0, xe-0/0/22.0, xe-0/0/21.0,

VLAN: v0002, Tag: 802.1Q Tag 2, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 0 (Active = 0)
  STP: None, RTG: None
  Untagged interfaces: None
  Tagged interfaces: None

VLAN: v0003, Tag: 802.1Q Tag 3, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 0 (Active = 0)
  STP: None, RTG: None
  Untagged interfaces: None
  Tagged interfaces: None

VLAN: vlan4000, 802.1Q Tag: Untagged, Admin State: Enabled
MAC learning Status: Disabled
Number of interfaces: 0 (Active = 0)
```

show vlans extensive (for a PVLAN spanning multiple switches)

```
user@switch> show vlans extensive
VLAN: COM1, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 100, Internal index: 3, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  1 (Active = 1)
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk
      ge-0/0/7.0*, untagged, access

VLAN: __pvlan_primary_ge-0/0/0.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  1 (Active = 1)
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk
      ge-0/0/0.0*, untagged, access


VLAN: __pvlan_primary_ge-0/0/2.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  1 (Active = 0)
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk
      ge-0/0/2.0, untagged, access


VLAN: __pvlan_primary_isiv__, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 50, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  0 (Active = 0)
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
```

```
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk


VLAN: community2, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 20, Internal index: 8, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  2 (Active = 2)
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk
      ge-0/0/1.0*, untagged, access
      ge-1/0/6.0*, untagged, access


VLAN: primary, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 10, Internal index: 2, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged  5 (Active = 4)
      ge-0/0/20.0*, tagged, trunk
      ge-0/0/22.0*, tagged, trunk, pvlan-trunk
      ge-0/0/23.0*, tagged, trunk, pvlan-trunk
      ge-0/0/0.0*, untagged, access
      ge-0/0/1.0*, untagged, access
      ge-0/0/2.0, untagged, access
      ge-0/0/7.0*, untagged, access
      ge-1/0/6.0*, untagged, access

Secondary VLANs: Isolated 2, Community  2, Inter-switch-isolated  1
  Isolated VLANs :
      __pvlan_primary_ge-0/0/0.0__
      __pvlan_primary_ge-0/0/2.0__
  Community VLANs :
      COM1
      community2
  Inter-switch-isolated VLAN :
      __pvlan_primary_isiv__
```

show vlans extensive (Port-Based on EX Series and QFX Series)

```
user@switch> show vlans extensive
VLAN: default, created at Mon Feb  4 12:13:47 2008
  Tag: None, Internal index: 0, Admin state: Enabled, Origin: static
  Description: None
  Customer VLAN ranges:
                     1-4100
    Protocol: Port based
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 0 (Active = 0), Untagged 23 (Active = 0)
    xe-0/0/34.0 (untagged, access)
    xe-0/0/33.0 (untagged, access)
    xe-0/0/32.0 (untagged, access)
    xe-0/0/31.0 (untagged, access)
    xe-0/0/30.0 (untagged, access)
    xe-0/0/29.0 (untagged, access)
    xe-0/0/28.0 (untagged, access)
    xe-0/0/27.0 (untagged, access)
    xe-0/0/26.0 (untagged, access)
    xe-0/0/25.0 (untagged, access)
    xe-0/0/19.0 (untagged, access)
    xe-0/0/18.0 (untagged, access)
    xe-0/0/17.0 (untagged, access)
    xe-0/0/16.0 (untagged, access)
    xe-0/0/15.0 (untagged, access)
    xe-0/0/14.0 (untagged, access)
    xe-0/0/13.0 (untagged, access)
    xe-0/0/11.0 (untagged, access)
    xe-0/0/9.0 (untagged, access)
    xe-0/0/8.0 (untagged, access)
    xe-0/0/3.0 (untagged, access)
    xe-0/0/2.0 (untagged, access)
    xe-0/0/1.0 (untagged, access)

Secondary VLANs: Isolated 1, Community  1
  Isolated VLANs :
      __pvlan_pvlan_xe-0/0/3.0__
  Community VLANs :
      comm1
```

```
VLAN: v0001, created at Mon Feb  4 12:13:47 2008
  Tag: 1, Internal index: 1, Admin state: Enabled, Origin: static
  Description: None
  Protocol: Port based, Layer 3 interface: None
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 4 (Active = 0), Untagged 0 (Active = 0)
    xe-0/0/24.0 (tagged, trunk)
    xe-0/0/23.0 (tagged, trunk)
    xe-0/0/22.0 (tagged, trunk)
    xe-0/0/21.0 (tagged, trunk)

VLAN: v0002, created at Mon Feb  4 12:13:47 2008
  Tag: 2, Internal index: 2, Admin state: Enabled, Origin: static
  Description: None
  Protocol: Port based, Layer 3 interface: None
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
    None

VLAN: v0003, created at Mon Feb  4 12:13:47 2008
  Tag: 3, Internal index: 3, Admin state: Enabled, Origin: static
  Description: None
  Protocol: Port based, Layer 3 interface: None
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
    None
```

show vlans extensive (MAC-based)

```
user@switch> show vlans extensive
VLAN: default, Created at: Thu May 15 13:43:09 2008
Internal index: 3, Admin State: Enabled, Origin: Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged  2 (Active = 2)
      ge-0/0/0.0*, untagged, access
      ge-0/0/14.0*, untagged, access

VLAN: vlan_dyn, Created at: Thu May 15 13:43:09 2008
```

```
Internal index: 4, Admin State: Enabled, Origin: Static
Protocol: Port Mode
Number of interfaces: Tagged 0 (Active = 0), Untagged  0 (Active = 0)
Protocol: MAC Based
Number of MAC entries: 6
      ge-0/0/0.0*
                  00:00:00:00:00:02 (untagged)
                  00:00:00:00:00:03 (untagged)
                  00:00:00:00:00:04 (untagged)
                  00:00:00:00:00:05 (untagged)
                  00:00:00:00:00:06 (untagged)
                  00:00:00:00:00:07 (untagged)
```

show vlans (Q-in-Q Tunneling on EX Series and QFX Series)

```
user@switch> show vlans dot1q-tunneling
Name           Tag     Interfaces
sv100          100     xe-0/0/4.0*, xe-0/0/15.0*
```

show vlans extensive (Q-in-Q Tunneling on EX Series and QFX Series)

```
user@switch> show vlans sv100 extensive
VLAN: sv100, Created at: Sat Sep 10 12:53:52 2011
802.1Q Tag: 100, Internal index: 2, Admin State: Enabled, Origin: Static
Dot1q Tunneling Status: Enabled
Customer VLAN ranges:
                     10-20
                     40-50
Protocol: Port Mode
Number of interfaces: Tagged 1 (Active = 1), Untagged  0 (Active = 0)
      ge-0/0/0.0, tagged, trunk

Number of mapping rules:
      Push 1 (Active = 0), Policy 0 (Active = 0), Swap 0 (Active = 0)

      xe-0/0/3.0*, 300, push
```

show vlans extensive (Q-in-Q Tunneling and L2TP on EX Series and QFX Series)

```
user@switch> show vlans v1 extensive
VLAN: v1, Created at: Fri Mar 2 05:07:38 2012
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Dot1q Tunneling status: Enabled
Layer2 Protocol Tunneling status: Enabled
```

show vlans sort-by tag (EX Series and QFX Series)

```
user@switch> show vlans sort-by tag
Name           Tag     Interfaces
default                       None
__vlan-x_1__   1
                       None
__vlan-x_2__   2
                       None
__vlan-x_3__   3
                       None
__vlan-x_4__   4
                       None
__vlan-x_5__   5
                       None
__vlan-x_6__   6
                       None
__vlan-x_7__   7
                       None
__vlan-x_8__   8
                       None
__vlan-x_9__   9
                       None
__vlan-x_10__  10
                       None
__vlan-x_11__  11
                       None
__vlan-x_12__  12
                       None
__vlan-x_13__  13
                       None
__vlan-x_14__  14
```

```
                        None
__vlan-x_15__  15
                        None
__vlan-x_16__  16
                        None
__vlan-x_17__  17
                        None
__vlan-x_18__  18
                        None
__vlan-x_19__  19
                        None
__vlan-x_20__  20
                        None
```

show vlans sort-by name (EX Series and QFX Series)

```
user@switch> show vlans sort-by employee

Name             Tag   Interfaces

__employee_120__ 120
                       xe-0/0/22.0*
__employee_121__ 121
                       xe-0/0/22.0*
__employee_122__ 122
                       xe-0/0/22.0*
__employee_123__ 123
                       xe-0/0/22.0*
__employee_124__ 124
                       xe-0/0/22.0*
__employee_125__ 125
                       xe-0/0/22.0*
__employee_126__ 126
                       xe-0/0/22.0*
__employee_127__ 127
                       xe-0/0/22.0*
__employee_128__ 128
                       xe-0/0/22.0*
__employee_129__ 129
                       xe-0/0/22.0*
```

```
__employee_130__ 130
                        xe-0/0/22.0*
```

show vlans tag (EX Series and QFX Series)

```
user@switch> show vlans employee

Name             Tag   Interfaces

__employee_120__ 120
                        xe-0/0/22.0*
__employee_121__ 121
                        xe-0/0/22.0*
__employee_122__ 122
                        xe-0/0/22.0*
__employee_123__ 123
                        xe-0/0/22.0*
__employee_124__ 124
                        xe-0/0/22.0*
__employee_125__ 125
                        xe-0/0/22.0*
__employee_126__ 126
                        xe-0/0/22.0*
__employee_127__ 127
                        xe-0/0/22.0*
__employee_128__ 128
                        xe-0/0/22.0*
__employee_129__ 129
                        xe-0/0/22.0*
__employee_130__ 130
                        xe-0/0/22.0*
```

show vlans sort-by tag (EX Series)

```
user@switch> show vlans sort-by tag
Name           Tag     Interfaces
default                None
__vlan-x_1__   1
                       None
__vlan-x_2__   2
```

```
                        None
__vlan-x_3__   3
                        None
__vlan-x_4__   4
                        None
__vlan-x_5__   5
                        None
__vlan-x_6__   6
                        None
__vlan-x_7__   7
                        None
__vlan-x_8__   8
                        None
__vlan-x_9__   9
                        None
__vlan-x_10__  10
                        None
__vlan-x_11__  11
                        None
__vlan-x_12__  12
                        None
__vlan-x_13__  13
                        None
__vlan-x_14__  14
                        None
__vlan-x_15__  15
                        None
__vlan-x_16__  16
                        None
__vlan-x_17__  17
                        None
__vlan-x_18__  18
                        None
__vlan-x_19__  19
                        None
__vlan-x_20__  20
                        None
```

show vlans employee (vlan-range-name)

```
user@switch> show vlans employee

Name             Tag   Interfaces

__employee_120__ 120
                       ge-0/0/22.0*
__employee_121__ 121
                       ge-0/0/22.0*
__employee_122__ 122
                       ge-0/0/22.0*
__employee_123__ 123
                       ge-0/0/22.0*
__employee_124__ 124
                       ge-0/0/22.0*
__employee_125__ 125
                       ge-0/0/22.0*
__employee_126__ 126
                       ge-0/0/22.0*
__employee_127__ 127
                       ge-0/0/22.0*
__employee_128__ 128
                       ge-0/0/22.0*
__employee_129__ 129
                       ge-0/0/22.0*
__employee_130__ 130
                       ge-0/0/22.0*
```

show vlans summary (EX Series)

```
user@switch> show vlans summary
VLANs summary:
         Total: 8,   Configured VLANs: 5
         Internal VLANs: 1,  Temporary VLANs: 0

     Dot1q VLANs summary:
         Total: 8,  Tagged VLANs: 2, Untagged VLANs: 6
         Private VLAN:
             Primary VLANs: 2,  Community VLANs: 2, Isolated VLANs: 3
```

```
     Dot1q Tunneled VLANs summary:
         Total: 0
         Private VLAN:
             Primary VLANs: 0,  Community VLANs: 0, Isolated VLANs: 0

     Dynamic VLANs:
        Total: 2,  Dot1x: 2, MVRP: 0
```

Sample Output: EX Series with ELS Switches and MX Routers

show vlans brief (EX Series Switch)

```
user@switch> show vlans brief

Routing instance        VLAN name             Tag          Interfaces
default-switch          c1                    20
                                                           ge-0/0/0.0*
                                                           ge-1/0/0.0*
                                                           ge-2/0/0.0*
default-switch          c2                    30
                                                           ge-0/0/0.0*
                                                           ge-2/0/0.0*
default-switch          default               1

default-switch          iso                   10
                                                           ge-0/0/1.0*
default-switch          iso1                  50
                                                           ge-0/0/0.0*
                                                           ge-2/0/0.0*
default-switch          pri                   100
                                                           ge-0/0/0.0*
                                                           ge-1/0/0.0*
                                                           ge-2/0/0.0*
```

show vlans brief (MX Routers)

```
user@host> show vlans brief

Routing instance        VLAN name             Tag         Interfaces
VPLS-1                  __VPLS-1__               all
                                                            ae1.0
VPLS-2                  __VPLS-2__               all
                                                            ae3.0
                                                            ge-3/1/2.0
                                                            vt-3/3/10.1048576
default-switch          VLAN1000                 1000
                                                            ae26.0
default-switch          VLAN101                  101
                                                            ae20.0
default-switch          VLAN102                  102
                                                            ae20.0
default-switch          VLAN103                  103
                                                            ae20.0
default-switch          VLAN104                  104
                                                            ae20.0
default-switch          VLAN105                  105
                                                            ae20.0
default-switch          VLAN106                  106
                                                            ae20.0
default-switch          VLAN107                  107
                                                            ae20.0
default-switch          VLAN108                  108
                                                            ae20.0
[...output truncated...]
```

show vlans detail (EX Series Switch)

```
user@switch> show vlans detail

Routing instance: default-switch
  VLAN Name: c1                             State: Active
Tag: 20
PVLAN type : Community
Internal index: 16, Generation Index: 21, Origin: Static
```

```
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: c2                             State: Active
Tag: 30
PVLAN type : Community
Internal index: 17, Generation Index: 22, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 2    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: default                        State: Active
Tag: 1
Internal index: 5, Generation Index: 5, Origin: Static
MAC aging time: 300 seconds
Number of interfaces: Tagged 0    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso                            State: Active
Tag: 10
Internal index: 14, Generation Index: 19, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/1.0*,untagged,access
Number of interfaces: Tagged 0    , Untagged 1
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso1                           State: Active
Tag: 50
PVLAN type : Isolated
Internal index: 15, Generation Index: 20, Origin: Static
```

```
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 2    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: pri                            State: Active
Tag: 100
PVLAN type : Primary
Isolated VLAN :
vlan-id : 50 vlan name : iso1
Community VLAN :
vlan-id : 20 vlan name : c1
vlan-id : 30 vlan name : c2
Internal index: 9, Generation Index: 14, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 0
```

show vlans detail (MX Routers)

```
user@host> show vlans detail
Routing instance: VPLS-1
  VLAN Name: __VPLS-1__                     State: Active
Tag: all
Internal index: 2, Generation Index:  , Origin: Dynamic
Interfaces:
    ae1.0,tagged
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 0

Routing instance: VPLS-2
  VLAN Name: __VPLS-2__                     State: Active
Tag: all
```

```
Internal index: 3, Generation Index:  , Origin: Dynamic
Interfaces:
    ae3.0,tagged
    ge-3/1/2.0,tagged
    vt-3/3/10.1048576,tagged
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 4

Routing instance: default-switch
  VLAN Name: VLAN1000                       State: Active
Tag: 1000
Internal index: 4, Generation Index: 1, Origin: Static
Layer 3 interface: irb.1000
Interfaces:
    ae26.0,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: VLAN101                        State: Active
Tag: 101
Internal index: 5, Generation Index: 2, Origin: Static
Layer 3 interface: irb.101
Interfaces:
    ae20.0,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 1

Routing instance: default-switch
  VLAN Name: VLAN102                        State: Active
Tag: 102
Internal index: 6, Generation Index: 3, Origin: Static
Layer 3 interface: irb.102
Interfaces:
    ae20.0,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 1
[...output truncated...]
```

show vlans extensive (EX Series Switch)

```
user@switch> show vlans extensive

Routing instance: default-switch
  VLAN Name: c1                             State: Active
Tag: 20
PVLAN type : Community
Internal index: 16, Generation Index: 21, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: c2                             State: Active
Tag: 30
PVLAN type : Community
Internal index: 17, Generation Index: 22, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 2    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: default                        State: Active
Tag: 1
Internal index: 5, Generation Index: 5, Origin: Static
MAC aging time: 300 seconds
Number of interfaces: Tagged 0    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso                            State: Active
Tag: 10
Internal index: 14, Generation Index: 19, Origin: Static
MAC aging time: 300 seconds
```

```
Interfaces:
    ge-0/0/1.0*,untagged,access
Number of interfaces: Tagged 0    , Untagged 1
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso1                           State: Active
Tag: 50
PVLAN type : Isolated
Internal index: 15, Generation Index: 20, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 2    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: pri                            State: Active
Tag: 100
PVLAN type : Primary
Isolated VLAN :
vlan-id : 50 vlan name : iso1
Community VLAN :
vlan-id : 20 vlan name : c1
vlan-id : 30 vlan name : c2
Internal index: 9, Generation Index: 14, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 0
```

show vlans extensive (MX Routers)

```
user@host> show vlans extensive

Routing instance: default-switch
```

```
  VLAN Name: VLAN_10                        State: Active
Tag: 10
Internal index: 2, Generation Index: 1, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-1/0/3.0*,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: VLAN_20                        State: Active
Tag: 20
Internal index: 3, Generation Index: 2, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-1/0/3.0*,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 0
```

Sample Output (SRX Series Firewalls)

show vlans (SRX Series Firewalls)

```
user@host> show vlans
Routing instance        VLAN name             Tag          Interfaces
default-switch          vlan-22                  22

default-switch          vlan-333                 333
                                                             ge-0/0/3.0*
                                                             ge-0/0/4.0*
default-switch          default               1

default-switch          vlan100               100
                                                             ge-0/0/1.0*
```

show vlans brief (SRX Series Firewalls)

```
user@host> show vlans brief
Routing instance        VLAN name             Tag          Interfaces
default-switch          vlan-22                  22

default-switch          vlan-333                 333
                                                             ge-0/0/3.0*
                                                             ge-0/0/4.0*
default-switch          default               1

default-switch          vlan100               100
                                                             ge-0/0/1.0*
```

show vlans detail (SRX Series Firewalls)

```
user@host> show vlans detail
Routing instance: default-switch
  VLAN Name: vlan-22                          State: Active
Tag: 22
Internal index: 2, Generation Index: 1, Origin: Static
MAC aging time: 300 seconds
VXLAN Enabled : No
Number of interfaces: Tagged 0    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: vlan-333                         State: Active
Tag: 333
Internal index: 3, Generation Index: 2, Origin: Static
MAC aging time: 300 seconds
VXLAN Enabled : No
Interfaces:
    ge-0/0/3.0*,tagged,trunk
    ge-0/0/4.0*,tagged,trunk
Number of interfaces: Tagged 2    , Untagged 0
Total MAC count: 0
```

Release Information

Command introduced in Junos OS Release 8.4.

Option **dot1q-tunneling** added in Junos OS Release 12.1 for the QFX Series.

Option `interface` introduced in Junos OS Release 13.2X50-D10 (ELS).

RELATED DOCUMENTATION

Example: Setting Up Basic Bridging and a VLAN on Switches

Example: Setting Up Bridging with Multiple VLANs

Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch

Example: Setting Up Bridging with Multiple VLANs for EX Series Switches

Example: Configuring a Private VLAN on a Single EX Series Switch

Example: Configuring a Private VLAN Spanning Multiple EX Series Switches

Example: Setting Up Q-in-Q Tunneling on EX Series Switches

Understanding Bridging and VLANs on Switches

show ethernet-switching interfaces | **19641**

show vlans (JDM)

IN THIS SECTION

- Syntax | **27768**
- Description | **27768**
- Options | **27768**
- Required Privilege Level | **27768**
- Output Fields | **27768**
- Sample Output | **27769**
- Release Information | **27769**

Syntax

```
show vlans vlan-name
```

Description

Display the details about the VLANs.

Options

vlan-name	Display information for a specific VLAN.
brief | detail | extensive | terse	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 1779 on page 27768 describes the output fields for the `show virtual-network-functions` command. Output fields are listed in the approximate order in which they appear.

Table 1779: show virtual-network functions Output Fields

Field Name	Field Description
vlan-name	Display information for a specified VLAN
brief	Display brief output
detail	Display detailed output
extensive	Display extensive output
instance	Display information for a specified instance

Table 1779: show virtual-network functions Output Fields *(Continued)*

Field Name	Field Description
interface	Name of interface for which to display table
logical-system	Name of logical systems

Sample Output

show vlans

```
root@jdm> show vlans

Routing instance        VLAN name             Tag          Interfaces
host-os                 vlan100               100
                                                           vsrx1_eth6.0
                                                           vsrx2_eth6.0
host-os                 vlan200-202-vlan-0200 200
                                                           vsrx1_eth7.0
                                                           vsrx2_eth7.0
host-os                 vlan200-202-vlan-0202 202
                                                           vsrx1_eth7.0
                                                           vsrx2_eth7.0
```

Release Information

Command introduced in Junos OS Release 15.1X53-D40.

show vlans evpn nd-table

IN THIS SECTION

Syntax | 27770

- Description | 27770
- Options | 27770
- Required Privilege Level | 27771
- Output Fields | 27771
- Sample Output | 27772
- Release Information | 27773

Syntax

```
show vlans evpn nd-table
<brief | detail | extensive>
<count>
<instance-name>
<mac-address>
<vlan-name>
```

Description

Display information about INET entries associated with MAC addresses learned through Network Discovery Protocol (NDP).

Options

none	Display information for all INET entries.
brief | detail | extensive	(Optional) Display the specified level of output.
count	(Optional) Display the number of INET addresses learned in a routing instance.
instance	(Optional) Display information for a specified instance.
mac-adress	(Optional) Display information for a specified MAC address.
vlan-name (all)	(Optional) Display information for a specified VLAN or for all VLANs.

Required Privilege Level

view

Output Fields

Table 1780 on page 27771 lists the output fields for the `show vlans evpn nd-table` command. Output fields are listed in the approximate order in which they appear.

Table 1780: show vlans evpn nd-table Output Fields

Field Name	Field Description	Level of Output
INET address	The INET address related to the INET entries that are added to the NDP table.	All levels
MAC address	MAC addresses learned through NDP.	**brief**, **detail**, **extensive**, **instance**, **mac-addressvlan-name**,,
Logical Interface	Logical interface associated with the routing instance in which the NDP INET address is learned.	**brief**, **instance**, **mac-addressvlan-name**,,
Routing instance	Routing instance in which the NDP INET address is learned.	all levels
Bridging domain	Bridging domain in which the NDP INET address is learned.	all levels
Learning interface	Interface on which the NDP INET address is learned.	**detail**, **extensive**
Count	Indicates the number of NDP INET addresses learned in a routing instance in a bridge domain.	**count**

Sample Output

show vlans evpn nd-table brief

```
user@switch> show vlans evpn nd-table brief
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
8002::2             00:05:86:a0:d5:00  irb.0         evpn1      vlan10
```

show vlans evpn nd-table detail

```
user@switch> show vlans evpn nd-table detail
INET address: 8002::2
MAC address: 00:05:86:a0:d5:00
  Routing instance: evpn1
   Bridging domain: vlan10
   Learning interface: irb.0
```

show vlans evpn nd-table extensive

```
user@switch> show vlans evpn nd-table extensive
INET address: 8002::2
MAC address: 00:05:86:a0:d5:00
  Routing instance: evpn1
   Bridging domain: vlan10
   Learning interface: irb.0
```

show vlans evpn nd-table count

```
user@switch> show vlans evpn nd-table count
1 ND INET addresses learned in routing instance evpn1 bridge domain vlan10
```

show vlans evpn nd-table instance evpn1

```
user@switch> show vlans evpn nd-table instance
evpn1
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
8002::2
            00:05:86:a0:d5:00
  irb.0         evpn1        vlan10
```

show vlans evpn nd-table instance 00:05:86:90:bd:f0 (MAC address)

```
user@switch> show vlans evpn nd-table
INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
8002::2
            00:05:86:a0:d5:00
  irb.0         evpn1        __evpn1__
```

show vlans evpn nd-table vlan-name vlan10

```
user@switch>show vlans evpn nd-table vlan-name
vlan10

INET                MAC                Logical       Routing      Bridging
address             address            interface     instance     domain
8002::2            00:05:86:a0:d5:00
  irb.0         evpn1        vlan10
```

Release Information

Command introduced in Junos OS Release 17.3R1.

> **NOTE**: Starting with Junos OS Release 17.4R2, the `show vlans evpn nd-table` command is replaced by the `show ethernet-switching evpn nd-table` command.

show vlans vlan-groups

IN THIS SECTION

- Syntax | 27774
- Description | 27774
- Options | 27774
- Required Privilege Level | 27774
- Output Fields | 27774
- Sample Output | 27775
- Release Information | 27775

Syntax

```
show vlans vlan-groups <vlan-group-name>
```

Description

Display the results of the vlan-groups configuration on the switch.

Options

vlan-group-name (Optional) Information about a specified VLAN group name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Output Fields

Table 1 on page 27775 lists the output fields for the `show vlans vlan-groups` command.

Table 1781: show vlans vlan-groups Output Fields

Field Name	Field Description
vlan-group-name	(Optional) Information about a specified VLAN group name

Sample Output

show vlans vlan-groups

```
user@host> show vlans vlan-groups
Vlan Group Name:          Last vlan used:    Associated Vlan Ids:
group1                              0                             30, 40-50
```

Release Information

Statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

vlan-groups | 30792

802.1X Authentication

show vmhost bridge

IN THIS SECTION

- Syntax | 27776
- Syntax | 27776
- Description | 27776
- Options | 27776
- Required Privilege Level | 27777

Sample Output | 27777
Release Information | 27777

Syntax

```
show vmhost bridge
<invoke-on>
<re0 | re1>
<routing engine>
```

Syntax

Description

Display bridge table information. The bridge table provides information about the interfaces used for communication between host and guest operating systems.

Options

invoke-on (Optional) Display the bridge table information of Routing Engines on a device that has dual or redundant Routing Engines. You can use the `all-routing-engine` option to display the bridge table information of all the Routing Engines or the `other-routing-engine` option to display the bridge table information of the other Routing Engine. For example, If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, display bridge table information about the Routing Engine in slot 0 (`re0`) or the Routing Engine in slot 1 (`re1`).

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

routing-engine (Optional) Specify the Routing Engine for which the bridge information is to be displayed. The following options are available:

- `backup`—Backup Routing Engine.
- `both`—Primary and the backup Routing Engines.
- `local`— Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

view

Sample Output

show vmhost bridge

```
user@host> show vmhost bridge
Compute cluster: rainier-re-cc
Compute node: rainier-re-cn
Bridge Table
 ============
bridge name              bridge id            STP enabled     interfaces
   jnpr-int-br          8000.bee5a8cfdb9a       no              tap1
   virbr0               8000.52540051f94b       yes             virbr0-nic
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

show vmhost connections

IN THIS SECTION

- Syntax | 27778
- Description | 27778
- Options | 27778
- Required Privilege Level | 27778
- Output Fields | 27779
- Sample Output | 27779
- Release Information | 27780

Syntax

```
show vmhost connections
```

Description

Display the details for the cross-connect connections. The NFX150 and NFX250 (NG) supports VLAN PUSH, POP, and SWAP operations.

Options

name	Display the details of a specific connection.
down	Display the details of connections that are not operational.
up	Display the details of connections that are operational.
up-down	Display the details of both operational and non-operational connections.

Required Privilege Level

view

Output Fields

Table 1782 on page 27779 lists the output fields for the `show vmhost connections` command. Output fields are listed in the approximate order in which they appear.

Table 1782: show vmhost connections Output Fields

Field Name	Field Description
Connection	Displays the type of the cross-connect.
Function	Displays the name of the virtual network function.
Interface	Specifies an interface on which the connection is established.
Status	Displays the status of the connection.

Sample Output

show vmhost connections

```
user@host> show vmhost connections
Connection                  Function                    Interface Vlan  Status
--------------------------------------------------------------------------------

phy_cc                      system                      sxe0     200   up
                            centos1                     eth2      500

push_pop_cc                 centos1                     eth2      none  down
                            centos2                     eth3      none

swap_cc                     centos1                     eth2      300   up
                            centos2                     eth2      400

vlan_cc                     centos1                     eth2      100   up
                            centos2                     eth2      100
```

Release Information

Command introduced in Junos OS Release 18.1R1.

show vmhost control-plane

IN THIS SECTION

- Syntax | 27780
- Description | 27780
- Required Privilege Level | 27780
- Sample Output | 27780
- Release Information | 27782

Syntax

```
show vmhost control-plane
```

Description

Display the status of the JCP, JDM, Layer 2 dataplane, Layer 3 dataplane, and LTE.

Required Privilege Level

view

Sample Output

NOTE: Starting in Junos OS Release 22.1R1, the output of the `show vmhost control-plane` command is enhanced to display the following information:

- Any system component that is currently offline. If you disabled a system component in the currently active custom mode configuration, then the status of that component is displayed as OFFLINE.
- Status and state of the NFV backplane

show vmhost control-plane (NFX150)

```
user@host> show vmhost control-plane
Vmhost Control Plane Information
--------------------------------
Name                        | State              | Status
-------------------------------------------------------
Junos Control Plane           RUNNING              OK
Juniper Device Manager        RUNNING              OK
Layer 2 Infrastructure        NOT_RUNNING          OFFLINE
Layer 3 Infrastructure        RUNNING              OK
NFV Backplane                 RUNNING              OK
```

show vmhost control-plane (NFX250 (NG))

```
user@host> show vmhost control-plane
Vmhost Control Plane Information
--------------------------------
Name                        | State              | Status
-------------------------------------------------------
Junos Control Plane           RUNNING              OK
Juniper Device Manager        RUNNING              OK
Layer 2 Infrastructure        RUNNING              OK
Layer 3 Infrastructure        RUNNING              OK
NFV Backplane                 RUNNING              OK
```

show vmhost control-plane (NFX350)

```
user@host> show vmhost control-planeVmhost Control Plane Information
--------------------------------
```

```
Name                        | State              | Status
-------------------------------------------------------
Junos Control Plane          RUNNING              OK
Juniper Device Manager       RUNNING              OK
Layer 2 Infrastructure       RUNNING              OK
Layer 3 Infrastructure       RUNNING              OK
NFV Backplane                RUNNING              OK
```

Release Information

Command introduced in Junos OS Release 18.1R1.

show vmhost crash

IN THIS SECTION

- Syntax | 27782
- Description | 27782
- Required Privilege Level | 27783
- Sample Output | 27783
- show vmhost crash (NFX Series) | 27783
- Release Information | 27783

Syntax

```
show vmhost crash
```

Description

Display the number of times the host OS crashed. The crash dumps are available at **/var/crash**.

Required Privilege Level

view

Sample Output

show vmhost crash

```
user@host> show vmhost crash
Compute cluster: rainier-re-cc

  Compute node: rainier-re-cn

   Crash Info
   ==========
   total 0
```

show vmhost crash (NFX Series)

show vmhost crash

```
user@host> show vmhost crash

-rw-r--r-- 1 root root 306773 Mar 22 10:41 local-node.srxpfe.7439.1521715280.core.tgz
-rw-r--r-- 1 root root 307058 Mar 22 10:42 local-node.srxpfe.8184.1521715324.core.tgz
-rw-r--r-- 1 root root 306999 Mar 22 10:42 local-node.srxpfe.8918.1521715357.core.tgz
-rw-r--r-- 1 root root 315121 Apr 18 05:35 localhost.dummy_flowdapp.3037.1524029709.core.tgz
-rw-r--r-- 1 root root 315033 Apr 18 05:17 localhost.dummy_flowdapp.3432.1524028674.core.tgz
-rw-r--r-- 1 root root 315088 Apr 13 18:11 localhost.dummy_flowdapp.3435.1523643106.core.tgz
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

show system core-dumps | 26783

show vmhost crash (NFX Series)

IN THIS SECTION

- Syntax | 27784
- Description | 27784
- Required Privilege Level | 27784
- Sample Output | 27784
- Release Information | 27785

Syntax

```
show vmhost crash
```

Description

Display host OS crash information.

Required Privilege Level

view

Sample Output

show vmhost crash

```
user@host> show vmhost crash

-rw-r--r-- 1 root root 306773 Mar 22 10:41 local-node.srxpfe.7439.1521715280.core.tgz
```

```
-rw-r--r-- 1 root root 307058 Mar 22 10:42 local-node.srxpfe.8184.1521715324.core.tgz
-rw-r--r-- 1 root root 306999 Mar 22 10:42 local-node.srxpfe.8918.1521715357.core.tgz
-rw-r--r-- 1 root root 315121 Apr 18 05:35 localhost.dummy_flowdapp.3037.1524029709.core.tgz
-rw-r--r-- 1 root root 315033 Apr 18 05:17 localhost.dummy_flowdapp.3432.1524028674.core.tgz
-rw-r--r-- 1 root root 315088 Apr 13 18:11 localhost.dummy_flowdapp.3435.1523643106.core.tgz
```

Release Information

Command introduced in Junos OS Release 18.1R1.

show vmhost forwarding-options analyzer

IN THIS SECTION

- Syntax | 27785
- Description | 27785
- Options | 27786
- Required Privilege Level | 27786
- Output Fields | 27786
- Sample Output | 27786
- Release Information | 27787

Syntax

```
show vmhost forwarding-options analyzer analyzer-name
```

Description

Displays information about the VNF analyzers that are configured for port mirroring on a Junos OS platform.

Options

analyzer-name Displays the details of a specific analyzer on the device.

Required Privilege Level

view

Output Fields

Table 1783 on page 27786 lists the output fields for the `show vmhost forwarding-options analyzer` command. Output fields are listed in the approximate order in which they appear.

Table 1783: show vmhost forwarding-options analyzer Output Fields

Field Name	Field Description
Analyzer name	Displays the name of the analyzer instance.
Egress monitored interfaces	Displays interfaces for which the traffic leaving the interfaces is mirrored.
Output interface	Specifies an interface to which mirrored packets are sent.
Ingress monitored interfaces	Displays interfaces for which the traffic entering the interfaces is mirrored.

Sample Output

show vmhost forwarding-options analyzer

```
user@host> show vmhost forwarding-options analyzer
Analyzer name                    : mon1
Egress monitored interfaces      : vnf-name1:eth2
Output interface                 : analyzer1:eth2

Analyzer name                    : mon2
Ingress monitored interfaces     : vnf-name2:eth2
Output interface                 : analyzer1:eth3
```

Release Information

Command introduced in Junos OS Release 18.1R1.

show vmhost hard-disk-test

IN THIS SECTION

- Syntax | 27787
- Description | 27787
- Options | 27787
- Required Privilege Level | 27788
- Sample Output | 27788
- Release Information | 27789

Syntax

```
show vmhost hard-disk -test { disk disk-name | status}
```

Description

Display memory and diagnostics monitoring test status on the solid-state drive (SSD). The test output provides information about the various attributes of the SSD that help to monitor the status of the hard disk memory. This command should be used only after initiating the disk test with the `request vmhost hard-disk-test` command.

Options

disk *disk-name*	Display the name of the SSD.
status	Display the status of the test.

Required Privilege Level

maintenance

Sample Output

show vmhost hard-disk-test

```
user@host> show vmhost hard-disk-test status disk /dev/sda
smartctl 5.42 2014-07-28 r3460 [x86_64-linux-3.10.79-ovp-rt74-WR6.0.0.20_preempt-rt] (local
build)
Copyright (C) 2002-11 by Bruce Allen, http://smartmontools.sourceforge.net

=== START OF INFORMATION SECTION ===
Model Family:     UNIGEN SATA SSD
Device Model:     SATA SSD
Serial Number:    3AF607410C3800117282
Firmware Version: S9FM01.3
User Capacity:    64,023,257,088 bytes [64.0 GB]
Sector Size:      512 bytes logical/physical
Device is:        In smartctl database [for details use: -P show]
ATA Version is:   8
ATA Standard is:  ACS-3 (revision not indicated)
Local Time is:    Sun Jan  8 08:02:22 2017 UTC
SMART support is: Available - device has SMART capability.
SMART support is: Enabled

=== START OF READ SMART DATA SECTION ===
SMART overall-health self-assessment test result: PASSED

General SMART Values:
Offline data collection status:  (0x00) Offline data collection activity
                                        was never started.
                                        Auto Offline Data Collection: Disabled.
Self-test execution status:      (   0) The previous self-test routine completed
                                        without error or no self-test has ever
                                        been run.
Total time to complete Offline
data collection:                (   30) seconds.
Offline data collection
capabilities:                    (0x1b) SMART execute Offline immediate.
```

```
                                        Auto Offline data collection on/off support.
                                        Suspend Offline collection upon new
...
...
...
```

Release Information

Command introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

request vmhost hard-disk-test | 15495

show vmhost hardware

IN THIS SECTION

- Syntax | 27789
- Description | 27790
- Options | 27790
- Required Privilege Level | 27790
- Sample Output | 27790
- Release Information | 27791

Syntax

```
show vmhost hardware
<invoke-on>
<re0 | re1>
<routing engine>
```

Description

Display details of RAM and solid-state drives (SSDs) installed in the Routing Engine.

Options

none (Optional) Display information about hardware.

invoke-on (Optional) Display the details of RAM and solid-state drives (SSDs) installed on a device that has dual Routing Engines. You can use the `all-routing-engine` option to display the hardware information of all the Routing Engines or the `other-routing-engine` option to display the hardware information of the other Routing Engine. For example, if you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, display hardware information about the Routing Engine in slot 0 (`re0`) or the Routing Engine in slot 1 (`re1`).

routing-engine (Optional) Specify the Routing Engine for which the details of the installed RAM and solid-state drives (SSDs) is to be displayed. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Primary and backup Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`— If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

view

Sample Output

show vmhost hardware

```
user@host> show vmhost hardware
Compute cluster: rainier-re-cc

  Compute node: rainier-re-cn
   Hardware inventory:

   Item    Capacity  Part number           Serial number         Description
   DIMM 0  16384 MB  36ADS2G72PZ-2G1A1     0x0CF49320            DDR4 2133 MHz
   DIMM 1  16384 MB  36ADS2G72PZ-2G1A1     0x0CF4934C            DDR4 2133 MHz
   DIMM 2  16384 MB  36ADS2G72PZ-2G1A1     0x0CF49329            DDR4 2133 MHz
   DIMM 3  16384 MB  36ADS2G72PZ-2G1A1     0x0CF49352            DDR4 2133 MHz
   Disk1   50.0  GB  StorFly-VSF202CC050G  P1T13003443810130041  SLIM SATA SSD
   Disk2   50.0  GB  StorFly-VSF202CC050G  P1T13003443810130012  SLIM SATA SSD
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

show chassis hardware | 18430

show vmhost information

IN THIS SECTION

- Syntax | 27792
- Description | 27792
- Options | 27792
- Required Privilege Level | 27793

Sample Output | 27793
Release Information | 27793

Syntax

```
show vmhost information
<invoke-on>
<re0 | re1>
<routing engine>
```

Description

Display information about the host—such as IP address of the host Routing Engine, host OS version, model number or name of the Routing Engine, and so on.

Options

invoke-on (Optional) Display information about the host on a device that has dual Routing Engines. You can use the `all-routing-engine` option to display information about the host of all the Routing Engines or the `other-routing-engine` option to display the information about the host of the other Routing Engine. If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine .

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, display information about the host of Routing Engine in slot 0 (`re0`) or the Routing Engine in slot 1 (`re1`).

routing-engine (Optional) Specify the Routing Engine for which the information about the host is to be displayed. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Primary and the backup Routing Engines.

- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

view

Sample Output

show vmhost information

```
user@host> show vmhost information
Compute cluster: rainier-re-cc
  Compute node       Model           Kernel release                 Machine    Management IP
  rainier-re-cn      RAINIER         3.10.79-ovp-rt74-WR6.0.0.20_preempt-rt x86_64 192.168.1.2/24
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

show vmhost jdm status (In-Chassis Model)

IN THIS SECTION

- Syntax | 27794
- Description | 27794
- Required Privilege Level | 27794
- Sample Output | 27794

Release Information | 27794

Syntax

```
show vmhost jdm status
```

Description

Displays the Juniper Device Manager (JDM) status. This command is applicable to in-chassis Junos node slicing.

Required Privilege Level

View

Sample Output

show vmhost jdm status

```
user@router> show vmhost jdm status

JDM Information
---------------------------
Package    :  jns-jdm-vmhost-19.1-B2.x86_64
Status     :  Running
PID        :  3088
Free Space :  62967 (MiB)
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

show virtual-network-functions (Junos Node Slicing) | 27728

request server authenticate-peer-server | 14901

show vmhost logs

IN THIS SECTION

- Syntax | 27795
- Description | 27795
- Options | 27795
- Required Privilege Level | 27796
- Sample Output | 27796
- Release Information | 27796

Syntax

```
show vmhost logs
<re0 | re1>
<routing engine>
```

Description

Display trace logs information of the host OS.

Options

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, display trace logs information of the host os running on the Routing Engine in slot 0 (`re0`) or the Routing Engine in slot 1 (`re1`).

routing-engine (Optional) Specify the Routing Engine for which the trace logs information of the host OS is to be displayed. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Primary and the backup Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

view

Sample Output

show vmhost logs

```
show vmhost logs
```

Release Information

Command introduced in Junos OS Release 15.1F3.

> **NOTE**: PTX3000 router supports the Routing and Control Board, RCBPTX.

show vmhost management-if

IN THIS SECTION

- Syntax | 27797
- Description | 27797
- Required Privilege Level | 27797
- Sample Output | 27797
- Release Information | 27798

Syntax

```
show vmhost management-if
```

Description

Display the administrative status, speed and operational mode of the host interface eth0, which serves as a management interface.

Required Privilege Level

view

Sample Output

show vmhost management-if

```
user@host> show vmhost management-if
Administrative status: Up
Link status: Up
Link speed: 1000Mb/s
Link operational mode: Full
```

Release Information

Command introduced in Junos OS Release 15.1F6.

> **NOTE**: The command is supported on the routers with RE-MX-X6, RE-MX-X8, and RE-PTX-X8 Routing Engines only.

show vmhost memory

IN THIS SECTION

- Syntax | 27798
- Description | 27798
- Required Privilege Level | 27798
- Output Fields | 27798
- Sample Output | 27799
- Release Information | 27799

Syntax

```
show vmhost memory
```

Description

Display the memory information for the host OS.

Required Privilege Level

view

Output Fields

Sample Output

show vmhost memory

```
user@host> show vmhost memory
Memory Controller Information
-----------------------------

Id :MC0
correctable-error                   :0
uncorrectable-error                 :0
```

Release Information

Command introduced in Junos OS Release 18.1R1.

show vmhost mode

IN THIS SECTION

- Syntax | 27800
- Description | 27800
- Required Privilege Level | 27800
- Sample Output (NFX150) | 27800
- Sample Output (NFX250 (NG)) | 27804
- Sample Output (NFX350) | 27807
- Release Information | 27810

Syntax

```
show vmhost mode
show vmhost mode mode-name
```

Description

The `show vmhost mode` command displays the CPU and memory allocations for various components in the current mode of the device. The `show vmhost mode` *`mode-name`* command displays the CPU and memory allocations for various components for a specific mode of the device.

> **NOTE**: Starting in Junos OS Release 22.1R1, the output of the `show vmhost mode` command displays the maximum memory that can be configured and the currently configured memory for VNFs. In addition, the CPU allocations are displayed as a number range (for example, 4-7) instead of individual numbers (for example, 4,5,6,7).

Required Privilege Level

view

Sample Output (NFX150)

show vmhost mode

```
user@host> show vmhost mode

Mode:
--------
Current Mode: compute

CPU Allocations:
Name                                Configured                              Used
-------------------------------------------------------------------------------------------------
---------------------
Junos Control Plane                 0                                       0,2
Juniper Device Manager              0                                       0
LTE                                 0                                       -
```

```
NFV Backplane Control Path          0                                       0
NFV Backplane Data Path             -                                       -
Layer 2 Control Path                0                                       0
Layer 2 Data Path                   1                                       0-1
Layer 3 Control Path                0                                       0
Layer 3 Data Path                   1                                       1
CPUs available for VNFs             2-3                                     2-3
CPUs turned off                     -                                       -

Memory Allocations:
Name                                Configured                              Used
-----------------------------------------------------------------------------------------------
---------------------
Junos Control Plane (mB)            2048                                    1895
NFV Backplane 1G hugepages          -                                       0
NFV Backplane 2M hugepages          -                                       0
Layer 2 1G hugepages                -                                       0
Layer 2 2M hugepages                128                                     128
Layer 3 1G hugepages                1                                       1
Layer 3 2M hugepages                401                                     400
VNF max memory limit (gB) (approx)  1                                       1.000
```

show vmhost mode hybrid

```
user@host> show vmhost mode hybrid

Mode: hybrid

CPU Allocations:
Name                                Configured
-------------------------------------------------------------------------------
Junos Control Plane                 0
LTE                                 0
Juniper Device Manager              0
NFV Backplane Control Path          0
NFV Backplane Data Path             -
Layer 2 Control Path                0
Layer 2 Data Path                   1
Layer 3 Control Path                0
Layer 3 Data Path                   2
CPUs turned off                     -
```

```
CPUs available for VNFs              3

Memory Allocations:
Name                                 Configured
---------------------------------------------------------------------------
Junos Control Plane (mB)             2048
NFV Backplane 1G hugepages           -
NFV Backplane 2M hugepages           -
Layer 2 1G hugepages                 -
Layer 2 2M hugepages                 128
Layer 3 1G hugepages                 1
Layer 3 2M hugepages                 401
VNF max memory limit (gB) (approx)   1
```

show vmhost mode throughput

```
user@host> show vmhost mode throughput

Mode: throughput

CPU Allocations:
Name                                 Configured
---------------------------------------------------------------------------
Junos Control Plane                  0
LTE                                  0
Juniper Device Manager               0
NFV Backplane Control Path           0
NFV Backplane Data Path              -
Layer 2 Control Path                 0
Layer 2 Data Path                    1
Layer 3 Control Path                 0
Layer 3 Data Path                    2-3
CPUs turned off                      -
CPUs available for VNFs              -

Memory Allocations:
Name                                 Configured
---------------------------------------------------------------------------
Junos Control Plane (mB)             2048
NFV Backplane 1G hugepages           -
NFV Backplane 2M hugepages           -
```

```
Layer 2 1G hugepages                -
Layer 2 2M hugepages                128
Layer 3 1G hugepages                1
Layer 3 2M hugepages                401
VNF max memory limit (gB) (approx)  -
```

show vmhost mode flex

```
user@host> show vmhost mode flex

Mode: flex

CPU Allocations:
Name                                Configured
---------------------------------------------------------------------------
Junos Control Plane                 0
LTE                                 0
Juniper Device Manager              0
NFV Backplane Control Path          2
NFV Backplane Data Path             -
Layer 2 Control Path                0
Layer 2 Data Path                   1
Layer 3 Control Path                0
Layer 3 Data Path                   2
CPUs turned off                     -
CPUs available for VNFs             3

Memory Allocations:
Name                                Configured
---------------------------------------------------------------------------
Junos Control Plane (mB)            1536
NFV Backplane 1G hugepages          -
NFV Backplane 2M hugepages          -
Layer 2 1G hugepages                -
Layer 2 2M hugepages                128
Layer 3 1G hugepages                -
Layer 3 2M hugepages                282
VNF max memory limit (gB) (approx)  2
```

Sample Output (NFX250 (NG))

show vmhost mode

```
user@host> show vmhost mode

Mode:
--------
Current Mode: compute

CPU Allocations:
Name                                Configured                              Used
-----------------------------------------------------------------------------------------------
---------------------
Junos Control Plane                 0                                       0,2
Juniper Device Manager              1                                       1
LTE                                 0                                       -
NFV Backplane Control Path          0                                       0
NFV Backplane Data Path             4                                       4
Layer 2 Control Path                -                                       -
Layer 2 Data Path                   -                                       -
Layer 3 Control Path                1                                       1
Layer 3 Data Path                   5                                       5
CPUs available for VNFs             2-3,6-7                                 2-3,6
CPUs turned off                     -                                       -

Memory Allocations:
Name                                Configured                              Used
-----------------------------------------------------------------------------------------------
---------------------
Junos Control Plane (mB)            2048                                    2020
NFV Backplane 1G hugepages          1                                       2
NFV Backplane 2M hugepages          -                                       0
Layer 2 1G hugepages                -                                       -
Layer 2 2M hugepages                -                                       -
Layer 3 1G hugepages                1                                       1
Layer 3 2M hugepages                1376                                    1375
VNF max memory limit (gB) (approx)  6                                       2.000
```

show vmhost mode hybrid

```
user@host> show vmhost mode hybrid

Mode: hybrid

CPU Allocations:
Name                                Configured
---------------------------------------------------------------------------
Junos Control Plane                 0
LTE                                 0
Juniper Device Manager              1
NFV Backplane Control Path          0
NFV Backplane Data Path             4
Layer 2 Control Path                -
Layer 2 Data Path                   -
Layer 3 Control Path                1
Layer 3 Data Path                   2,5
CPUs turned off                     6
CPUs available for VNFs             3,7

Memory Allocations:
Name                                Configured
---------------------------------------------------------------------------
Junos Control Plane (mB)            2048
NFV Backplane 1G hugepages          1
NFV Backplane 2M hugepages          -
Layer 2 1G hugepages                -
Layer 2 2M hugepages                -
Layer 3 1G hugepages                1
Layer 3 2M hugepages                1376
VNF max memory limit (gB) (approx)  6
```

show vmhost mode throughput

```
user@host> show vmhost mode throughput

Mode: throughput

CPU Allocations:
```

```
Name                                Configured
--------------------------------------------------------------------------
Junos Control Plane                 4
LTE                                 4
Juniper Device Manager              4
NFV Backplane Control Path          4
NFV Backplane Data Path             -
Layer 2 Control Path                -
Layer 2 Data Path                   -
Layer 3 Control Path                0
Layer 3 Data Path                   2-3
CPUs turned off                     1,5-7
CPUs available for VNFs             -

Memory Allocations:
Name                                Configured
--------------------------------------------------------------------------
Junos Control Plane (mB)            2048
NFV Backplane 1G hugepages          -
NFV Backplane 2M hugepages          -
Layer 2 1G hugepages                -
Layer 2 2M hugepages                -
Layer 3 1G hugepages                1
Layer 3 2M hugepages                1376
VNF max memory limit (gB) (approx)  -
```

show vmhost mode flex

```
user@host> show vmhost mode flex

Mode: flex

CPU Allocations:
Name                                Configured
--------------------------------------------------------------------------
Junos Control Plane                 0-1
LTE                                 0
Juniper Device Manager              1
NFV Backplane Control Path          7
NFV Backplane Data Path             6
```

```
Layer 2 Control Path                -
Layer 2 Data Path                   -
Layer 3 Control Path                1
Layer 3 Data Path                   7
CPUs turned off                     -
CPUs available for VNFs             2-5,8-11

Memory Allocations:
Name                                Configured
---------------------------------------------------------------------------
Junos Control Plane (mB)            1536
NFV Backplane 1G hugepages          1
NFV Backplane 2M hugepages          -
Layer 2 1G hugepages                -
Layer 2 2M hugepages                -
Layer 3 1G hugepages                -
Layer 3 2M hugepages                282
VNF max memory limit (gB) (approx)  24
```

Sample Output (NFX350)

show vmhost mode

```
user@host> show vmhost mode

Mode:
--------
Current Mode: compute

CPU Allocations:
Name                                Configured                              Used
-------------------------------------------------------------------------------------------------
---------------------
Junos Control Plane                 8                                       3,8
Juniper Device Manager              8                                       8
LTE                                 8                                       -
NFV Backplane Control Path          8                                       8
NFV Backplane Data Path             1                                       1
Layer 2 Control Path                -                                       -
Layer 2 Data Path                   -                                       -
Layer 3 Control Path                0                                       0
```

```
Layer 3 Data Path                    2                                       2
CPUs available for VNFs              3-7,11-15                               4-5,11,15
CPUs turned off                      9-10                                    -

Memory Allocations:
Name                                 Configured                              Used
-----------------------------------------------------------------------------------------------
---------------------
Junos Control Plane (mB)             2048                                    1926
NFV Backplane 1G hugepages           4                                       8
NFV Backplane 2M hugepages           -                                       0
Layer 2 1G hugepages                 -                                       -
Layer 2 2M hugepages                 -                                       -
Layer 3 1G hugepages                 4                                       4
Layer 3 2M hugepages                 5633                                    5377
VNF max memory limit (gB) (approx)   7                                       2.000
```

show vmhost mode hybrid

```
user@host> show vmhost mode hybrid

Mode: hybrid

CPU Allocations:
Name                                 Configured
-------------------------------------------------------------------------------
Junos Control Plane                  8
LTE                                  8
Juniper Device Manager               8
NFV Backplane Control Path           8
NFV Backplane Data Path              1
Layer 2 Control Path                 -
Layer 2 Data Path                    -
Layer 3 Control Path                 0
Layer 3 Data Path                    2-3
CPUs turned off                      9-11
CPUs available for VNFs              4-7,12-15

Memory Allocations:
Name                                 Configured
-------------------------------------------------------------------------------
```

```
Junos Control Plane (mB)             2048
NFV Backplane 1G hugepages           4
NFV Backplane 2M hugepages           -
Layer 2 1G hugepages                 -
Layer 2 2M hugepages                 -
Layer 3 1G hugepages                 4
Layer 3 2M hugepages                 5633
VNF max memory limit (gB) (approx)   7
```

show vmhost mode throughput

```
user@host> show vmhost mode throughput

Mode: throughput

CPU Allocations:
Name                                 Configured
---------------------------------------------------------------------------
Junos Control Plane                  8
LTE                                  8
Juniper Device Manager               8
NFV Backplane Control Path           8
NFV Backplane Data Path              -
Layer 2 Control Path                 -
Layer 2 Data Path                    -
Layer 3 Control Path                 0
Layer 3 Data Path                    2-7
CPUs turned off                      1,9-15
CPUs available for VNFs              -

Memory Allocations:
Name                                 Configured
---------------------------------------------------------------------------
Junos Control Plane (mB)             2048
NFV Backplane 1G hugepages           -
NFV Backplane 2M hugepages           -
Layer 2 1G hugepages                 -
Layer 2 2M hugepages                 -
Layer 3 1G hugepages                 4
```

```
Layer 3 2M hugepages                 5633
VNF max memory limit (gB) (approx)   -
```

show vmhost mode flex

```
user@host> show vmhost mode flex

Mode: flex
CPU Allocations:
Name                                 Configured
-----------------------------------------------------------------------------
Junos Control Plane                  0,1
LTE                                  0
Juniper Device Manager               1
NFV Backplane Control Path           1
NFV Backplane Data Path              4
Layer 2 Control Path                 -
Layer 2 Data Path                    -
Layer 3 Control Path                 1
Layer 3 Data Path                    5
CPUs turned off                      -
CPUs available for VNFs              2,3,6,7



Memory Allocations:
Name                                 Configured
-----------------------------------------------------------------------------
Junos Control Plane (mB)             1536
NFV Backplane 1G hugepages           1
NFV Backplane 2M hugepages           -
Layer 2 1G hugepages                 -
Layer 2 2M hugepages                 -
Layer 3 1G hugepages                 1
Layer 3 2M hugepages                 200
```

Release Information

Command introduced in Junos OS Release 19.1R1.

show vmhost netstat

IN THIS SECTION

- Syntax | 27811
- Description | 27811
- Options | 27811
- Required Privilege Level | 27812
- Sample Output | 27812
- Release Information | 27813

Syntax

```
show vmhost netstat
<invoke-on>
<re0 | re1>
<routing engine>
```

Description

Display network statistics information for the host OS. The statistics contains information related to the interfaces used for the communication between the host and the guest, such as the IP address of the destination, IP address of the gateway, mask, flags, and so on.

Options

invoke-on (Optional) Display the network statistics for the host OS on a device that has dual Routing Engines. You can use the `all-routing-engine` option to display the network statistics information for the host OS running on all the Routing Engines or the `other-routing-engine` option to display the network statistics information for the host OS running on the other Routing Engine. For example, If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine .

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, display the network statistics information for the host OS running on the Routing Engine in slot 0 (`re0`) or the Routing Engine in slot 1 (`re1`).

routing-engine (Optional) Specify the Routing Engine for which the network statistics information for the host OS is to be displayed. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Primary and the backup Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

view

Sample Output

show vmhost netstat

```
user@host> show vmhost netstat
Compute cluster: rainier-re-cc

  Compute node: rainier-re-cn

   Netstat
   =======
   Kernel IP routing table
   Destination     Gateway         Genmask         Flags   MSS Window  irtt Iface
   0.0.0.0         10.216.63.254   0.0.0.0         UG        0 0          0 eth0
   10.216.48.0     0.0.0.0         255.255.240.0   U         0 0          0 eth0
```

```
192.168.1.0     0.0.0.0         255.255.255.0   U         0 0          0 jnpr-int-br
192.168.122.0   0.0.0.0         255.255.255.0   U         0 0          0 virbr0
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

show vmhost network external-network

IN THIS SECTION

- Syntax | 27813
- Description | 27813
- Required Privilege Level | 27813
- Sample Output | 27814
- Release Information | 27814

Syntax

```
show vmhost network external-network
```

Description

Display current configured and created external interfaces on the hypervisor.

Required Privilege Level

view

Sample Output

show vmhost network external-network

```
user@host> show vmhost network external-network
```

```
Interface : hext-ib.10
Type : in-band
Link State : RUNNING
PF : heth-0-0, VF : 0000:02:10.1
MAC : c8:e7:f0:c9:86:45, VLAN : 10, MTU : 1500
IPv4 : 10.10.10.10/24, Gateway : 10.10.10.254
    Rx-packets :       20
    Rx-errors  :       0
    Rx-drops   :       0
    Tx-packets :       69
    Tx-drops   :       0
    Tx-errors  :       0

Interface : hext-ob
Type : out-of-band
Link State : RUNNING
MAC : c8:e7:f0:c9:86:44, MTU : 1500
IPv4 : 10.204.97.175/20, Gateway : 10.204.111.254
    Rx-packets :       277166
    Rx-errors  :       0
    Rx-drops   :       0
    Tx-packets :       65541
    Tx-drops   :       0
    Tx-errors  :       0
```

Release Information

Command introduced in Junos OS Release 24.2R1 for NFX150, NFX250, and NFX350 devices.

show vmhost processes

IN THIS SECTION

- Syntax | 27815
- Description | 27815
- Options | 27815
- Required Privilege Level | 27816
- Sample Output | 27816
- Release Information | 27817

Syntax

```
show vmhost processes
<invoke-on>
<re0 | re1>
<routing engine>
```

Description

Display information about the host processes that are running on the device.

Options

invoke-on (Optional) Display information about the host processes that are running on a device with dual Routing Engines. You can use the `all-routing-engine` option to display information about the host processes running on all the Routing Engines or the `other-routing-engine` option to display information about the host processes running on the other Routing Engine. For example, If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine .

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, display information about the host processes running on the Routing Engine in slot 0 (re0) or on the Routing Engine in slot 1 (re1).

routing-engine (Optional) Specify the Routing Engine for which the information about the host processes is to be displayed. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Primary and the backup Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

view

Sample Output

show vmhost processes

```
user@host> show vmhost processes
Compute cluster: rainier-re-cc

  Compute node: rainier-re-cn
   UID        PID  PPID  C STIME TTY          TIME CMD
   root         1     0  0 21:56 ?        00:00:02 init [3]
   root         2     0  0 21:56 ?        00:00:00 [kthreadd]
   root         3     2  0 21:56 ?        00:00:04 [ksoftirqd/0]
   root         5     2  0 21:56 ?        00:00:00 [kworker/0:0H]
   root         7     2  0 21:56 ?        00:00:00 [posixcputmr/0]
   root         8     2  0 21:56 ?        00:00:00 [kclksetdelayd]
   root         9     2  0 21:56 ?        00:00:00 [rcub/0]
   root        10     2  0 21:56 ?        00:00:04 [rcu_preempt]
   root        11     2  0 21:56 ?        00:00:00 [rcu_sched]
   root        12     2  0 21:56 ?        00:00:00 [rcu_bh]
   root        13     2  0 21:56 ?        00:00:03 [rcuc/0]
```

```
root        14     2  0 21:56 ?        00:00:00 [kcmosdelayd]
root        15     2  0 21:56 ?        00:00:00 [migration/0]
root        16     2  0 21:56 ?        00:00:00 [migration/1]
root        17     2  0 21:56 ?        00:00:03 [rcuc/1]
root        18     2  0 21:56 ?        00:00:04 [ksoftirqd/1]
root        19     2  0 21:56 ?        00:00:00 [posixcputmr/1]
root        20     2  0 21:56 ?        00:00:00 [kworker/1:0]
root        21     2  0 21:56 ?        00:00:00 [kworker/1:0H]
root        22     2  0 21:56 ?        00:00:00 [migration/2]
root        23     2  0 21:56 ?        00:00:10 [rcuc/2]
root        24     2  0 21:56 ?        00:00:02 [ksoftirqd/2]
root        25     2  0 21:56 ?        00:00:00 [posixcputmr/2]
root        26     2  0 21:56 ?        00:00:00 [kworker/2:0]
root        27     2  0 21:56 ?        00:00:00 [kworker/2:0H]
root        28     2  0 21:56 ?        00:00:00 [migration/3]
root        29     2  0 21:56 ?        00:00:01 [rcuc/3]
root        30     2  0 21:56 ?        00:00:01 [ksoftirqd/3]
root        31     2  0 21:56 ?        00:00:00 [posixcputmr/3]
root        32     2  0 21:56 ?        00:00:00 [kworker/3:0]
root        33     2  0 21:56 ?        00:00:00 [kworker/3:0H]
root        34     2  0 21:56 ?        00:00:00 [migration/4]
root        35     2  0 21:56 ?        00:00:01 [rcuc/4]
root        36     2  0 21:56 ?        00:00:01 [ksoftirqd/4]
root        37     2  0 21:56 ?        00:00:00 [posixcputmr/4]
root        38     2  0 21:56 ?        00:00:00 [kworker/4:0]
root        39     2  0 21:56 ?        00:00:00 [kworker/4:0H]
root        40     2  0 21:56 ?        00:00:00 [migration/5]
root        41     2  0 21:56 ?        00:00:01 [rcuc/5]
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

show system processes

show vmhost resource-usage

IN THIS SECTION

- Syntax | 27818
- Description | 27818
- Options | 27818
- Required Privilege Level | 27819
- Sample Output | 27819
- Release Information | 27820

Syntax

```
show vmhost resource-usage
<invoke-on>
<re0 | re1>
<routing engine>
```

Description

Display the current usage of solid-state drive (SSD), RAM, and CPU resources of the host OS.

Options

invoke-on (Optional) Display information about resources used by the host OS running on a device that has dual Routing Engines. You can use the `all-routing-engine` option to display information about resources used by the host OS on all the Routing Engines or the `other-routing-engine` option to display information about resources used by the host OS on the other Routing Engine. For example, If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, display information about resources used by the host OS on the Routing Engine in slot 0 (re0) or on the Routing Engine in slot 1 (re1).

routing-engine (Optional) Specify the Routing Engine for which the information about resources used by the host OS is to be displayed. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Primary and the backup Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

view

Sample Output

show vmhost resource-usage

```
user@host> show vmhost resource-usage
Compute cluster: rainier-re-cc

  Compute node: rainier-re-cn
   CPU Usage
   =========
   23:15:09   CPU  %usr  %nice    %sys %iowait   %irq   %soft  %steal  %guest   %idle
   23:15:09   all  0.36   0.00    1.16    0.07   0.00    0.01   0.00   74.07   24.34
   23:15:09   0    1.42   0.00    4.46    0.28   0.00    0.03   0.00    0.00   93.81
   23:15:09   1    1.43   0.00    3.87    0.30   0.00    0.03   0.00    0.00   94.38
   23:15:09   2    0.02   0.00    0.42    0.00   0.00    0.01   0.00   98.33    1.22
   23:15:09   3    0.00   0.00    0.14    0.00   0.00    0.00   0.00   98.65    1.21
   23:15:09   4    0.00   0.00    0.09    0.00   0.00    0.00   0.00   98.71    1.19
   23:15:09   5    0.00   0.00    0.10    0.00   0.00    0.00   0.00   98.71    1.19
   23:15:09   6    0.00   0.00    0.11    0.00   0.00    0.00   0.00   98.70    1.19
   23:15:09   7    0.00   0.00    0.12    0.00   0.00    0.00   0.00   98.68    1.19
```

```
Memory Usage
============
             total       used       free     shared    buffers     cached
Mem:         63851      51388      12462          0         35        123
Swap:            0          0          0

Disk Usage
==========
Filesystem                    Size  Used Avail Use% Mounted on
rootfs                        3.3G  127M  3.0G   5% /
/dev/sda4                     2.0G  1.7G  166M  91% /.old
tmpfs                          32G  1.2M   32G   1% /.old/run
none                           32G  4.0K   32G   1% /.old/dev
tmpfs                          32G     0   32G   0% /.old/tmp
tmpfs                          32G     0   32G   0% /.old/tdev
/dev/mapper/jvg_P-jlvmrootrw  3.3G  127M  3.0G   5% /.old/tdev/.union/.s
unionfs                       3.3G  127M  3.0G   5% /
none                           32G  4.0K   32G   1% /dev
tmpfs                          32G  180K   32G   1% /run
tmpfs                          32G  8.0K   32G   1% /var/volatile
/dev/mapper/jvg_P-jlvmjunos    32G   13G   18G  43% /junos
/dev/mapper/jvg_P-jlvmvm      6.1G  2.7G  3.1G  47% /vm
/dev/mapper/jvg_P-jlvmspare   287M  2.1M  266M   1% /spare
cgroup                         32G     0   32G   0% /sys/fs/cgroup
unionfs                       3.3G  127M  3.0G   5% /run/named-chroot/etc/bind
tmpfs                          32G  180K   32G   1% /run/named-chroot/var/run/named
tmpfs                          32G  180K   32G   1% /run/named-chroot/var/run/bind
unionfs                       3.3G  127M  3.0G   5% /run/named-chroot/var/cache/bind
unionfs                       3.3G  127M  3.0G   5% /run/named-chroot/etc/localtime
none                           32G  4.0K   32G   1% /run/named-chroot/dev/random
none                           32G  4.0K   32G   1% /run/named-chroot/dev/zero
none                           32G  4.0K   32G   1% /run/named-chroot/dev/null
```

It is normal to see 100 percentage or close to 100 percentage of CPU usage in the `%guest` column in the `CPU Usage` section.

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

show vmhost snapshot

IN THIS SECTION

- Syntax | 27821
- Description | 27821
- Options | 27821
- Required Privilege Level | 27822
- Sample Output | 27822
- Release Information | 27823

Syntax

```
show vmhost snapshot
<invoke-on>
<re0 | re1>
<routing engine>
```

Description

Display snapshot details including Linux host kernel version, software version, and other package version details for both the sets of software in the backup disk.

Options

invoke-on (Optional) Display the host snapshot information of Routing Engines on a device that has dual Routing Engines. You can use the `all-routing-engine` option to display the host snapshot information of all the Routing Engines or the `other-routing-engine` option to display the host snapshot information of the other Routing Engine. For example, If you issue the command

from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, display host snapshot information about the Routing Engine in slot 0 (`re0`) or the Routing Engine in slot 1 (`re1`).

routing-engine (Optional) Specify the Routing Engine for which the host snapshot details is to be displayed. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Primary and backup Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

view

Sample Output

show vmhost snapshot

```
user@host> show vmhost snapshot
UEFI    Version: NGRE_v00.53.00.01

Secondary Disk, Snapshot Time: Tue Dec  8 19:49:09 UTC 2015

Version: set p
VMHost Version: 2.897
VMHost Root: vmhost-x86_64-15.1I20151203_0011_rbu-builder
VMHost Core: vmhost-core-x86_64-15.1I20151203_0011_rbu-builder
```

```
kernel: 3.10.79-ovp-rt74-WR6.0.0.20_preempt-rt
Junos Disk: junos-install-x86-64-15.1F-20151204.0

Version: set b
VMHost Version: 2.897
VMHost Root: vmhost-x86_64-15.1I20151203_0011_rbu-builder
VMHost Core: vmhost-core-x86_64-15.1I20151203_0011_rbu-builder
kernel: 3.10.79-ovp-rt74-WR6.0.0.20_preempt-rt
Junos Disk: junos-install-x86-64-15.1F-20151204.0
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

request vmhost software rollback | 15556

request vmhost snapshot | 15539

show vmhost snmp mib walk

IN THIS SECTION

- Syntax | 27824
- Description | 27824
- Required Privilege Level | 27824
- Sample Output | 27824
- Release Information | 27825

Syntax

```
show vmhost snmp mib walk mib-name
```

Description

The `show vmhost snmp mib walk` *`mib-name`* command displays the libvirt SNMP object values that are associated with the libvirt MIB.

NOTE: To view information for other MIBs, use the `show snmp mib walk` command.

Required Privilege Level

view

Sample Output

show vmhost snmp mib walk

```
root@host> show vmhost snmp mib walk LIBVIRT-MIB::libvirtMIB
LIBVIRT-MIB::libvirtGuestName[STRING: 33913919-8dde-4689-b54b-81beb25d84ae] = STRING: "centos1"
LIBVIRT-MIB::libvirtGuestName[STRING: aabbccdd-eeff-42e3-0-564a554e4f53] = STRING: "vjunos0"
LIBVIRT-MIB::libvirtGuestState[STRING: 33913919-8dde-4689-b54b-81beb25d84ae] = INTEGER:
running(1)
LIBVIRT-MIB::libvirtGuestState[STRING: aabbccdd-eeff-42e3-0-564a554e4f53] = INTEGER: running(1)
LIBVIRT-MIB::libvirtGuestCpuCount[STRING: 33913919-8dde-4689-b54b-81beb25d84ae] = Gauge32: 1
LIBVIRT-MIB::libvirtGuestCpuCount[STRING: aabbccdd-eeff-42e3-0-564a554e4f53] = Gauge32: 1
LIBVIRT-MIB::libvirtGuestMemoryCurrent[STRING: 33913919-8dde-4689-b54b-81beb25d84ae] = Gauge32:
512
LIBVIRT-MIB::libvirtGuestMemoryCurrent[STRING: aabbccdd-eeff-42e3-0-564a554e4f53] = Gauge32: 1954
LIBVIRT-MIB::libvirtGuestMemoryLimit[STRING: 33913919-8dde-4689-b54b-81beb25d84ae] = Gauge32: 512
LIBVIRT-MIB::libvirtGuestMemoryLimit[STRING: aabbccdd-eeff-42e3-0-564a554e4f53] = Gauge32: 1954
LIBVIRT-MIB::libvirtGuestCpuTime[STRING: 33913919-8dde-4689-b54b-81beb25d84ae] = Counter64:
12300000000
LIBVIRT-MIB::libvirtGuestCpuTime[STRING: aabbccdd-eeff-42e3-0-564a554e4f53] = Counter64:
734900000000
LIBVIRT-MIB::libvirtGuestRowStatus[STRING: 33913919-8dde-4689-b54b-81beb25d84ae] = INTEGER:
active(1)
```

```
LIBVIRT-MIB::libvirtGuestRowStatus[STRING: aabbccdd-eeff-42e3-0-564a554e4f53] = INTEGER:
active(1)
```

Release Information

Command introduced in Junos OS Release 21.4R1.

show vmhost status

IN THIS SECTION

- Syntax | 27825
- Description | 27825
- Options | 27826
- Required Privilege Level | 27826
- Sample Output | 27826
- Release Information | 27827

Syntax

```
show vmhost status
<invoke-on>
<re0 | re1>
<routing engine>
```

Description

Display information about the status of communication between the host OS and the guest OS. The following status outputs are displayed:

- `Online`—Communication between the host OS and the guest OS is good.
- `Offline`—Communication with the host is lost. Any state other than `Online` is considered as `Offline`.

Options

invoke-on (Optional) Display the status of communication between the host OS and the guest OS running on a router with dual Routing Engines. You can use the `all-routing-engine` option to display the status of host-to-guest communication on all the Routing Engines or the `other-routing-engine` option to display the status of host-to-guest communication on the other Routing Engine. For example, If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, display the status of communication between the host OS and the guest OS on the Routing Engine in slot 0 (`re0`) or on the Routing Engine in slot 1 (`re1`).

routing-engine (Optional) Specify the Routing Engine for which the status of communication between the host OS and the guest OS is to be displayed. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Primary and backup Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

view

Sample Output

show vmhost status

```
user@host> show vmhost status
Compute cluster: rainier-re-cc
```

```
Compute Node: rainier-re-cn, Online
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

show vmhost status (NFX Series)

IN THIS SECTION

- Syntax | 27827
- Description | 27827
- Required Privilege Level | 27827
- Sample Output | 27828
- Release Information | 27828

Syntax

```
show vmhost status
```

Description

Display the virtualization status and status of all the CPUs.

Required Privilege Level

view

Sample Output

show vmhost status

```
user@host> show vmhost status
 Virtualization status :
 -------------------------------
 kvm_status      : ok
 libvirt_status  : ok
 qemu_status     : ok

 CPU Status [Since Boot Time]:
 ---------------------------------------------------------------------------------------
 CPU        %usr   %nice    %sys   %iowait  %irq   %soft  %steal  %guest  %gnice  %idle

 Load Avg  : 4.04   0.00     4.74   0.01     0.00   0.01   0.00    0.30    0.00    90.90
 cpu0      : 8.26   0.00     15.91  0.06     0.00   0.06   0.00    2.47    0.00    73.23
 cpu1      : 24.73  0.00     22.95  0.00     0.00   0.00   0.00    0.00    0.00    52.32
 cpu2      : 0.00   0.00     0.01   0.00     0.00   0.00   0.00    0.02    0.00    99.97
 cpu3      : 0.00   0.00     0.00   0.00     0.00   0.00   0.00    0.00    0.00    100.00
 cpu4      : 0.00   0.00     0.00   0.00     0.00   0.02   0.00    0.00    0.00    99.98
 cpu5      : 0.00   0.00     0.00   0.00     0.00   0.00   0.00    0.00    0.00    100.00
 cpu6      : 0.00   0.00     0.00   0.00     0.00   0.00   0.00    0.00    0.00    100.00
 cpu7      : 0.00   0.00     0.00   0.00     0.00   0.00   0.00    0.00    0.00    100.00

 Device:  tps    kB_read/s    kB_wrtn/s    kB_read    kB_wrtn
 ---------------------------------------------------------
 sda      2.15   7.60         30.04        4057951    16046703
```

Release Information

Command introduced in Junos OS Release 18.1R1.

show vmhost status (In-Chassis Model)

IN THIS SECTION

- Syntax | 27829
- Description | 27829
- Options | 27829
- Required Privilege Level | 27830
- Sample Output | 27830
- Release Information | 27830

Syntax

```
show vmhost status (re0 | re1)
<routing-engine (backup | both | local | master | other)>
<invoke-on  (all-routing-engines | other-routing-engine)>
```

Description

Displays the VM host status. This command is applicable to in-chassis Junos node slicing. When the router is operating in the in-chassis Junos node slicing mode, the `vJunos Resource Status` is shown as `Compact`. When the router is not operating in the in-chassis Junos node slicing mode, the `vJunos Resource Status` is shown as `Normal`.

Options

re0 Displays Routing Engine 0 (re0) VM host status.

re1 Displays Routing Engine 1 (re1) VM host status.

routing-engine Displays the VM host status of a specific routing engine. You have the following sub-options:

- backup

- primary
- both
- local
- other

invoke-on Displays the VM host status of the remote Routing Engines or all Routing Engines. You have the following options:

- all-routing-engines
- other-routing-engine

Required Privilege Level

View

Sample Output

show vmhost jdm status re0

```
user@router> show vmhost status re0

bsys-re0:
--------------------------------------------------------------------------

Compute cluster: rainier-re-cc
  Compute Node: rainier-re-cn, Online

vJunos Resource Status: Compact
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Generic Guidelines for Using JDM Server Commands

show virtual-network-functions (Junos Node Slicing) | 27728

request server authenticate-peer-server | 14901

show vmhost storage

IN THIS SECTION

- Syntax | 27831
- Description | 27831
- Required Privilege Level | 27831
- Sample Output | 27832
- Sample Output | 27833
- Sample Output | 27833
- Release Information | 27837

Syntax

```
show vmhost storage
```

Description

Display the vmhost storage information.

Required Privilege Level

view

Sample Output

show vmhost storage (NFX150)

```
user@host> show vmhost storage
Vmhost Storage Information
------------------------------
Storage Name          : sda
SSD Description       : Internal disk 1
SSD Model Number      : SFSA100GQ1AA4TO-C-LB-216-JUN
SSD Serial Number     : 000060124205B1000099
SSD Firmware Version  : SBR13025

  ID      Storage S.M.A.R.T attribute        Raw value

  1       Raw_Read_Error_Rate                0
  5       Reallocated_Sector_Ct              0
  9       Power_On_Hours                     20792
  12      Power_Cycle_Count                  66
  160     Uncorrectable_Sector_Count         0
  161     Spare_Blocks                       568
  163     Number_of_Initial_Invalid_Blocks   18
  164     Total_Erase_Count                  163038
  165     Maximum_Erase_Count                160
  166     Minimum_Erase_Count                34
  167     Average_Erase_Count                78
  168     Maximum_Specified_Erase_Count      3000
  169     Power-On_UECC_Count                54
  192     Power-Off_Retract_Count            568
  193     Dynamic_Remaps                     0
  194     Temperature_Celsius                32
  195     Hardware_ECC_Recovered             1345461
  196     Reallocated_Event_Count            0
  198     Offline_Uncorrectable              0
  199     UDMA_CRC_Error_Count               0
  215     TRIM_Count                         71048
  235     Total_Flash_LBAs_Written           289438408
  237     Total_Flash_LBAs_Written_Expanded  0
  241     Total_LBAs_Written                 13595913833
  242     Total_LBAs_Read                    6786635984
  243     Total_Host_LBAs_Written_Expanded   0
  244     Total_Host_LBAs_Read_Expanded      0
```

```
  248     SSD_Remaining_Life                 98
  249     Spare_Blocks_Remaining_Life        100
```

Sample Output

show vmhost storage (NFX250 NextGen)

```
user@host> show vmhost storage
Vmhost Storage Information
------------------------------
Storage Name          : sda
SSD Description       : Internal disk 1
SSD Model Number      : StorFly VSFBM6CC100G-JUN
SSD Serial Number     : P1T13004007308160267
SSD Firmware Version  : 1130-000

  ID      Storage S.M.A.R.T attribute        Raw value

  1       Raw_Read_Error_Rate                0
  9       Power_On_Hours                     1
  12      Power_Cycle_Count                  37
  192     Power-Off_Retract_Count            28
  194     Temperature_Celsius                40
  199     UDMA_CRC_Error_Count               0
  160     Uncorrectable_Sector_Count         0
  161     Spare_Blocks                       100
  241     Total_LBAs_Written                 30678
  242     Total_LBAs_Read                    7542
  169     Power-On_UECC_Count                100
  248     SSD_Remaining_Life                 100
  249     Spare_Blocks_Remaining_Life        100
```

Sample Output

show vmhost storage (NFX350)

```
user@host> show vmhost storage
Vmhost Storage Information
------------------------------
```

```
Storage Name          : sda
SSD Description       : Internal disk 1
SSD Model Number      : SFSA050GM3AA2TO-C-LB-34A-JUN
SSD Serial Number     : 000060154396B1000059
SSD Firmware Version  : SBR12050

  ID      Storage S.M.A.R.T attribute        Raw value

  1       Raw_Read_Error_Rate                1
  5       Reallocated_Sector_Ct              0
  9       Power_On_Hours                     8467
  12      Power_Cycle_Count                  120
  160     Uncorrectable_Sector_Count         0
  161     Spare_Blocks                       277
  163     Number_of_Initial_Invalid_Blocks   15
  164     Total_Erase_Count                  113168
  165     Maximum_Erase_Count                146
  166     Minimum_Erase_Count                47
  167     Average_Erase_Count                108
  168     Maximum_Specified_Erase_Count      3000
  169     Power-On_UECC_Count                85
  192     Power-Off_Retract_Count            277
  193     Dynamic_Remaps                     0
  194     Temperature_Celsius                42
  195     Hardware_ECC_Recovered             2092
  196     Reallocated_Event_Count            0
  198     Offline_Uncorrectable              0
  199     UDMA_CRC_Error_Count               1
  215     TRIM_Count                         20355
  235     Total_Flash_LBAs_Written           110143092
  237     Total_Flash_LBAs_Written_Expanded  0
  241     Total_LBAs_Written                 9943202407
  242     Total_LBAs_Read                    6158124561
  243     Total_Host_LBAs_Written_Expanded   0
  244     Total_Host_LBAs_Read_Expanded      0
  248     SSD_Remaining_Life                 97
  249     Spare_Blocks_Remaining_Life        100

Vmhost Storage Information
------------------------------
Storage Name          : sdb
SSD Description       : Public disk 0
SSD Model Number      : SFSA800GM3AA8TO-C-OC-626-JUN
```

```
SSD Serial Number     : 000060154239B1000059
SSD Firmware Version  : SBR13056
External SSD State    : INITIALIZED
External SSD Slot     : SSD0
Public Directory Path : /var/public-disk0

  ID      Storage S.M.A.R.T attribute        Raw value

  1       Raw_Read_Error_Rate                0
  5       Reallocated_Sector_Ct              0
  9       Power_On_Hours                     7604
  12      Power_Cycle_Count                  98
  160     Uncorrectable_Sector_Count         0
  161     Spare_Blocks                       1068
  163     Number_of_Initial_Invalid_Blocks   98
  164     Total_Erase_Count                  15715
  165     Maximum_Erase_Count                43
  166     Minimum_Erase_Count                0
  167     Average_Erase_Count                3
  168     Maximum_Specified_Erase_Count      3000
  169     Power-On_UECC_Count                28
  192     Power-Off_Retract_Count            1068
  193     Dynamic_Remaps                     971
  194     Temperature_Celsius                37
  195     Hardware_ECC_Recovered             18110
  196     Reallocated_Event_Count            0
  198     Offline_Uncorrectable              0
  199     UDMA_CRC_Error_Count               0
  215     TRIM_Count                         343556
  235     Total_Flash_LBAs_Written           81364321
  237     Total_Flash_LBAs_Written_Expanded  0
  241     Total_LBAs_Written                 5041956446
  242     Total_LBAs_Read                    3934034061
  243     Total_Host_LBAs_Written_Expanded   0
  244     Total_Host_LBAs_Read_Expanded      0
  248     SSD_Remaining_Life                 100
  249     Spare_Blocks_Remaining_Life        100

Vmhost Storage Information
------------------------------
Storage Name          : sdc
SSD Description       : Internal disk 2
SSD Model Number      : SFSA050GM3AA2TO-C-LB-34A-JUN
```

```
SSD Serial Number     : 000060154396B1000058
SSD Firmware Version  : SBR12050

  ID      Storage S.M.A.R.T attribute        Raw value

  1       Raw_Read_Error_Rate                0
  5       Reallocated_Sector_Ct              0
  9       Power_On_Hours                     8467
  12      Power_Cycle_Count                  122
  160     Uncorrectable_Sector_Count         0
  161     Spare_Blocks                       275
  163     Number_of_Initial_Invalid_Blocks   17
  164     Total_Erase_Count                  7492
  165     Maximum_Erase_Count                19
  166     Minimum_Erase_Count                0
  167     Average_Erase_Count                7
  168     Maximum_Specified_Erase_Count      3000
  169     Power-On_UECC_Count                30
  192     Power-Off_Retract_Count            275
  193     Dynamic_Remaps                     0
  194     Temperature_Celsius                42
  195     Hardware_ECC_Recovered             207
  196     Reallocated_Event_Count            0
  198     Offline_Uncorrectable              0
  199     UDMA_CRC_Error_Count               0
  215     TRIM_Count                         3843
  235     Total_Flash_LBAs_Written           4950046
  237     Total_Flash_LBAs_Written_Expanded  0
  241     Total_LBAs_Written                 532128913
  242     Total_LBAs_Read                    291859128
  243     Total_Host_LBAs_Written_Expanded   0
  244     Total_Host_LBAs_Read_Expanded      0
  248     SSD_Remaining_Life                 100
  249     Spare_Blocks_Remaining_Life        100

Vmhost Storage Information
------------------------------
Storage Name          : sdd
SSD Description       : Public disk 1
SSD Model Number      : M.2 (S80) 3MG2-P
SSD Serial Number     : B0021811130190037
SSD Firmware Version  : M271112J
External SSD State    : ADDED
```

```
External SSD Slot     : SSD1
Public Directory Path : /var/public-disk1

  ID      Storage S.M.A.R.T attribute        Raw value

  1       Raw_Read_Error_Rate                0
  5       Reallocated_Sector_Ct              0
  9       Power_On_Hours                     8
  12      Power_Cycle_Count                  137
  160     Uncorrectable_Sector_Count         0
  163     Number_of_Initial_Invalid_Blocks   78
  164     Total_Erase_Count                  1001
  165     Maximum_Erase_Count                2
  166     Minimum_Erase_Count                0
  167     Average_Erase_Count                0
  168     Maximum_Specified_Erase_Count      3000
  175     Program_Fail_Count_Chip            0
  176     Erase_Fail_Count_Chip              0
  177     Wear_Leveling_Count                0
  178     Used_Rsvd_Blk_Cnt_Chip             0
  181     Program_Fail_Cnt_Total             0
  182     Erase_Fail_Count_Total             0
  192     Power-Off_Retract_Count            12
  194     Temperature_Celsius                42
  195     Hardware_ECC_Recovered             7633
  196     Reallocated_Event_Count            0
  197     Current_Pending_Sector             0
  198     Offline_Uncorrectable              0
  199     UDMA_CRC_Error_Count               0
  232     Available_Reservd_Space            100
  241     Total_LBAs_Written                 4154
  242     Total_LBAs_Read                    183
  245     Unknown_Attribute                  8008
  248     SSD_Remaining_Life                 100
  249     Spare_Blocks_Remaining_Life        100
```

Release Information

Command introduced in Junos OS Release 18.1R1.

show vmhost uptime

IN THIS SECTION

- Syntax | 27838
- Description | 27838
- Options | 27838
- Required Privilege Level | 27839
- Sample Output | 27839
- Release Information | 27839

Syntax

```
show vmhost uptime
<invoke-on>
<re0 | re1>
<routing engine>
```

Description

Display the current time and information such as how long the host OS has been running, number of users, average load, and reason for the last reboot that occurred.

Options

invoke-on (Optional) Display the uptime information about the host on a device with dual Routing Engines. You can use the `all-routing-engine` option to display the uptime information about the host on all the Routing Engines or the `other-routing-engine` option to display the uptime information about the host on the other Routing Engine. For example, If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, display the uptime information about the hoston the Routing Engine in slot 0 (`re0`) or on the Routing Engine in slot 1 (`re1`).

routing-engine (Optional) Specify the Routing Engine for which the uptime information about the host is to be displayed. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Primary and backup Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

view

Sample Output

show vmhost uptime

```
show vmhost uptime
```

Release Information

Command introduced in Junos OS Release 15.1F3.

> **NOTE**: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

show chassis routing-engine | 18703

show vmhost uptime (NFX Series)

IN THIS SECTION

- Syntax | 27840
- Description | 27840
- Required Privilege Level | 27840
- Reboot Reason Codes | 27840
- Sample Output | 27841
- Release Information | 27841

Syntax

```
show vmhost uptime
```

Description

Display the current time and information such as how long the host OS has been running, number of users, average load, and the last reboot reason.

Required Privilege Level

view

Reboot Reason Codes

Vmhost last reboot reason: 0x20	power cycle
Vmhost last reboot reason: 0x04	reset button
Vmhost last reboot reason: 0x01	cold reset
Vmhost last reboot reason: 0x80	hypervisor reboot
Vmhost last reboot reason: 0x40	watchdog reset

Sample Output

show vmhost uptime

```
user@host> show vmhost uptime
 Vmhost Current time: 2020-02-05 10:04:09+00:00
 Vmhost Uptime:
       10:04:09 up 7 days, 21:43, 0 users, load average: 1.33, 1.26, 1.19
 Vmhost last reboot reason: 0x20
```

In the output message, the `vmhost last reboot reason` field provides the reboot reason code. To understand various reboot reason codes and its description, see "Reboot Reason Codes" on page 27840.

Release Information

Command introduced in Junos OS Release 18.1R1.

show vmhost version

IN THIS SECTION

- Syntax | 27841
- Description | 27842
- Options | 27842
- Required Privilege Level | 27843
- Sample Output | 27843
- Release Information | 27844

Syntax

```
show vmhost version
        invoke-on;
        (re0 | re1);
         routing engine {
```

```
            backup;
            both;
            local;
            master;
            other;
}
```

Description

Display host version information including Linux host kernel version, host software version, and other package version details for both the sets of software in the primary disk.

Options

invoke-on (Optional) Display the version information of the host running on a device with dual Routing Engines. You can use the `all-routing-engine` option to display the version information of the host software running on all the Routing Engines or the `other-routing-engine` option to display the version information of the host software running on other Routing Engine. For example, If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

re0 | re1 (Optional) On devices that support dual or redundant Routing Engines, display version information of the host software on the Routing Engine in slot 0 (re0) or on the Routing Engine in slot 1 (re1).

routing-engine (Optional) Specify the Routing Engine for which the version information of the host software is to be displayed. The following options are available:

> **NOTE**: The QFX10002-60C and PTX10002-60C devices do not have primary and backup routing engines.

- `backup`—Backup Routing Engine.
- `both`—Primary and backup Routing Engines.
- `local`—Routing Engine on the local Virtual Chassis member.
- `master`—Primary Routing Engine.
- `other`—If you issue the command from the primary Routing Engine, the backup Routing Engine is referred to as the other Routing Engine.

Required Privilege Level

view

Sample Output

show vmhost version

```
user@hostshow vmhost version
Current root details,         Device sda, Label: jrootp_P, Partition: sda3
Current boot disk: Primary
Current root set: p
UEFI Version: CICL_P_BUG1_00.23.01

Primary Disk, Upgrade Time: Wed Mar  1 16:35:37 PST 2023

Version: set p
VMHost Version: 7.2648
VMHost Root: vmhost-x86_64-22.2R3-20230129_2258_builder
VMHost Core: vmhost-core-x86-64-22.2R3.9
kernel: 5.2.60-rt15-LTS19
Junos Disk: junos-install-mx-x86-64-22.2R3.9

Version: set b
VMHost Version: 7.2648
VMHost Root: vmhost-x86_64-22.2R3-20230129_2258_builder
VMHost Core: vmhost-core-x86-64-22.2R3.10
kernel: 5.2.60-rt15-LTS19
Junos Disk: junos-install-mx-x86-64-22.2R3.10
```

show vmhost version

```
user@host> show vmhost version | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/20.3I0/junos">
    <output>
        Current root details,           Device sda, Label: jrootp_P, Partition: sda3
        Current boot disk: Primary
        Current root set: p
        UEFI    Version: NGRE_v00.53.00.01
```

```
        Primary Disk, Upgrade Time: Mon May 11 22:20:39 PDT 2020

        Version: set p
        VMHost Version: 4.2222
        VMHost Root: vmhost-x86_64-20.2Iwrlsdk9_20200427_1232.ad9e120_rbu-builder
        VMHost Core: vmhost-core-x86-64-20.3I-20200507_dev_common.0.2113
        kernel: 4.8.28-rt10-WR9.0.0.24_ovp
        Junos Disk: junos-install-mx-x86-64-20.3I-20200507_dev_common.0.2113

        Version: set b
        VMHost Version: 4.2222
        VMHost Root: vmhost-x86_64-20.2Iwrlsdk9_20200427_1232.ad9e120_rbu-builder
        VMHost Core: vmhost-core-x86-64-20.3I20200508_0827_sshanmugasun
        kernel: 4.8.28-rt10-WR9.0.0.24_ovp
        Junos Disk: junos-install-mx-x86-64-20.3I20200508_0827_sshanmugasun

    </output>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

Release Information

Command introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

show vmhost version (NFX Series)

IN THIS SECTION

- Syntax | 27845
- Description | 27845
- Required Privilege Level | 27845

- Sample Output | 27845
- Sample Output | 27846
- Sample Output | 27846
- Release Information | 27847

Syntax

```
show vmhost version detail
```

Description

Display host version information including Linux host kernel version and host software version.

Required Privilege Level

view

Sample Output

show vmhost version (NFX150)

```
user@host> show vmhost version detail
Partition set     : primary
Software version  : 20.3I-20200601_dev_common.0.0613
                    Host kernel release  : 4.1.27-rt30-WR8.0.0.30_ovp
                    Host kernel version  : #1 SMP Fri Jun  1 22:42:16 IST 2019

Partition set             : primary
Software version          : 20.3I-20200601_dev_common.0.0613
Installed/Upgraded at     : Wed Jun  3 12:58:46 UTC 2020
Status                    : Boot success

Partition set             : alternate
Software version          : 20.3I-20200404_dev_common.0.0613
```

```
Installed/Upgraded at     : Fri May 5 05:33:45 UTC 2020
Status                    : Boot success
```

Sample Output

show vmhost version (NFX250 NextGen)

```
user@host> show vmhost version detail
Partition set     : primary
Software version  : 20.3I-20200518_dev_common.0.2122
                    Host kernel release  : 4.1.27-rt30-WR8.0.0.30_ovp
                    Host kernel version  : #1 SMP Fri Dec 27 22:42:16 IST 2019

Partition set             : primary
Software version          : 20.3I-20200518_dev_common.0.2122
Installed/Upgraded at     : Wed May 20 10:11:27 UTC 2020
Status                    : Boot success

Partition set             : alternate
Software version          : 20.3I-20200601_dev_common.0.0613
Installed/Upgraded at     : Thu Jun  4 12:50:37 UTC 1970
Status                    : Boot success
```

Sample Output

show vmhost version (NFX350)

```
user@host> show vmhost version detail
Partition set     : alternate
Software version  : 20.3I-20200601_dev_common.0.0613
                    Host kernel release  : 4.1.27-rt30-WR8.0.0.30_ovp
                    Host kernel version  : #1 SMP Fri Dec 27 22:42:16 IST 2019

Partition set             : primary
Software version          : 20.3I-20200527_dev_common.0.1016
Installed/Upgraded at     : Mon Jun  1 20:02:10 UTC 2020
Status                    : Boot success

Partition set             : alternate
```

```
Software version          : 20.3I-20200601_dev_common.0.0613
Installed/Upgraded at     : Mon Jun  1 08:17:51 UTC 2020
Status                    : Boot success

Partition set             : second primary
Software version          : 20.3I-20200527_dev_common.0.1016
Installed/Upgraded at     : Mon May 28 09:05:30 UTC 2020
Status                    : Boot success

Partition set             : second alternate
Software version          : 20.3I-20200527_dev_common.0.1016
Installed/Upgraded at     : Mon May 28 09:05:34 UTC 2020
Status                    : Boot success
```

Release Information

Command introduced in Junos OS Release 18.1R1.

show vmhost vlans

IN THIS SECTION

- Syntax | 27847
- Description | 27848
- Options | 27848
- Required Privilege Level | 27848
- Output Fields | 27848
- Sample Output | 27849
- Release Information | 27849

Syntax

```
show vmhost vlans
```

Description

Display details about the vmhost VLANs.

Options

vlan-name	Display information for a specified VLAN.
brief | detail | extensive	Display the specified level of output.
instance	Display information for a specified instance.
interface	Name of interface for which the table is displayed.
logical-system	Name of logical system.

Required Privilege Level

view

Output Fields

Table 1784 on page 27848 describes the output fields for the"show vmhost forwarding-options analyzer" on page 27785"show vmhost vlans" on page 27847 `show vmhost vlans` command. Output fields are listed in the approximate order in which they appear.

Table 1784: show vmhost vlans Output Fields

Field Name	Field Description
vlan-name	Display information for a specified VLAN
brief	Display brief output
detail	Display detailed output
extensive	Display extensive output
instance	Display information for a specified instance

Table 1784: show vmhost vlans Output Fields *(Continued)*

Field Name	Field Description
interface	Name of interface for which to display table
logical-system	Name of logical system

Sample Output

show vmhost vlans

```
root@host> show vmhost vlans

Routing instance        VLAN name             Tag          Interfaces
vmhost                  test-1                56
                                                           centos1_eth2.0
----
```

Release Information

Command introduced in Junos OS Release 18.1R1.

show vpls connections

IN THIS SECTION

- Syntax | 27850
- Description | 27850
- Options | 27850
- Required Privilege Level | 27851
- Output Fields | 27851
- Sample Output | 27858

- Release Information | **27865**

Syntax

```
show vpls connections
<brief | extensive>
<down | up | up-down>
<history>
<instance instance-name local-site local-site-name remote-site remote-site-name>
<instance-history>
<logical-system (all | logical-system-name)>
<status>
<summary>
```

Description

(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) connection information.

Options

none	Display information about all VPLS connections for all routing instances.
brief | extensive	(Optional) Display the specified level of output.
down | up | up-down	(Optional) Display nonoperational, operational, or both types of connections.
history	(Optional) Display information about connection history.
instance *instance-name*	(Optional) Display the VPLS connections for the specified routing instance only.
instance-history	(Optional) Display information about connection history for a particular instance.

local-site *local-site-name*	(Optional) Display the VPLS connections for the specified local site name or ID only.
remote-site *remote-site-name*	(Optional) Display the VPLS connections for the specified remote site name or ID only. Label block size information is always shown as 0 when using this option.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.
status	(Optional) Display information about the connection and interface status.
summary	(Optional) Display summary of all VPLS connections information.

Required Privilege Level

view

Output Fields

Table 1785 on page 27851 lists the output fields for the `show vpls connections` command. Output fields are listed in the approximate order in which they appear.

Table 1785: show vpls connections Output Fields

Field Name	Field Description
`Instance`	Name of the VPLS instance.
`Local site`	Name of the local site.
`VPLS-id`	Identifier for the VPLS site.
`Number of local interfaces`	Number of interfaces configured for the local site.
`Number of local interfaces up`	Number of interfaces configured for the local site that are currently up.

Table 1785: show vpls connections Output Fields *(Continued)*

Field Name	Field Description
`IRB interface present`	Indicates whether or not an integrated routing and bridging (IRB) interface is present (`yes` or `no`).
`Intf`	List of all of the interfaces configured for the local site. The types of interfaces can include VPLS virtual loopback tunnel interfaces and label-switched interfaces. Any interface that supports VPLS could be listed here. Virtual loopback tunnel interfaces are displayed using the `vt-`*`fpc`*`/`*`pic`*`/`*`port`*`.`*`nnnnn` format. Label-switched interfaces are displayed using the `lsi.`*`nnnnn` format. In both cases, *`nnnnn`* is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain. Each interface might include the following information: • Identification as a VPLS interface • Name of the associated VPLS routing instance • Local site number • Remote site number • VPLS neighbor address • VPLS identifier
`Interface flags`	Flag associated with the interface. Can include the following: • `VC-Down`—The virtual circuit associated with this interface is down.
`Label-base`	First label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router.
`Offset`	Displays the VPLS Edge (VE) block offset in the Layer 2 VPN NLRI. The VE block offset is used to identify a label block from which a particular label value is selected to setup a pseduowire for a remote site. The block offset value itself indicates the starting VE ID that maps to the label base contained in the VPLS NLRI advertisement.

Table 1785: show vpls connections Output Fields *(Continued)*

Field Name	Field Description
`Size`	Label block size. A configurable value that represents the number of label blocks required to cover all the pseudowires for the remote peer. Acceptable configuration values are: `2`, `4`, `8` and `16`. The default value is `2`. A value of `0` will be displayed when using the `remote-site` option.
`Range`	Label block range. A value that keeps track of the numbers of remote sites discovered within each label block.
`Preference`	Preference value advertised for a VPLS site. When multiple PE routers are assigned the same VE ID for multihoming, you might need to specify that a particular PE router acts as the designated forwarder by configuring the site preference value. The site preference indicates the degree of preference for a particular customer site. The site preference is one of the tie-breaking criteria used in a designated forwarder election.
`status-vector`	Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of `0` indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of `1` indicates either one or both are down.
`connection-site`	Name of the connection site.
`Neighbor`	IP address and VPLS identifier for the VPLS neighbor. If multiple pseudowires have been configured, the IP address will also show the PW-specific "vpls-id-list" on page 30929, for example, 203.0.113.144 (vpls-id 200).
`Type`	Type of connection: `loc` (local) or `rmt` (remote).

Table 1785: show vpls connections Output Fields *(Continued)*

Field Name	Field Description

St	Status of the VPLS connection (corresponds with Legend for Connection Status):
	• EI—The local VPLS interface is configured with an encapsulation that is not supported.
	• EM—The encapsulation type received on this VPLS connection from the neighbor does not match the local VPLS connection interface encapsulation type.
	• VC-Dn—The virtual circuit is currently down.
	• CM—The two routers do not agree on a control word, which causes a control word mismatch.
	• CN—The virtual circuit is not provisioned properly.
	• OR—The label associated with the virtual circuit is out of range.
	• OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit.
	• LD—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established.
	• RD—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established.
	• LN—The local site has lost path selection to the remote site and therefore no pseduowires can be established from this local site.
	• RN—The remote site has lost path selection to a local site or other remote site and therefore no pseudowires are established to this remote site.
	In a multihoming configuration, one multihomed PE site displays the state LN, and the other multihomed PE site displays the state RN in the following circumstances:
	  • The multihomed links are both configured to be the backup site.
	  • The two multihomed PE routers have the same site ID, but have a peering relationship with a route reflector (RR) that has a different site ID.
	• XX—The VPLS connection is down for an unn reason. This is a programming error.

Table 1785: show vpls connections Output Fields *(Continued)*

Field Name	Field Description
	• MM—The MTU for the local site and the remote site do not match. • BK—The router is using a backup connection. • PF—Profile parse failure. • RS—The remote site is in a standby state. • NC—The interface encapsulation is not configured as an appropriate CCC, TCC, or VPLS encapsulation. • WE—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the VPLS routing instance. • NP—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the desired type, or the interface might be configured in a different routing instance. • -->—Only the outbound connection is up. • <-—Only the inbound connection is up. • Up—The VPLS connection is operational. • Dn—The VPLS connection is down. • CF—The router cannot find enough bandwidth to the remote router to satisfy the VPLS connection bandwidth requirement. • SC—The local site identifier matches the remote site identifier. No pseudowire can be established between these two sites. You should configure different values for the local and remote site identifiers. • LM—The local site identifier is not the minimum designated, meaning it is not the lowest. There is another local site with a lower site identifier. Pseudowires are not being established to this local site. and the associated local site identifier is not being used to distribute VPLS label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. • RM—The remote site identifier is not the minimum designated, meaning it is not the lowest. There is another remote site connected to the same PE router which has

Table 1785: show vpls connections Output Fields *(Continued)*

Field Name	Field Description
	lower site identifier. The PE router cannot established a pseudowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic can continue to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. • IL—The incoming packets for the VPLS connection have no MPLS label. • MI—The configured mesh group identifier is in use by another system in the network. • ST—The router has switched to a standby connection. • PB—Profile busy. • SN—The VPLS neighbor is static.
Time last up	Time connection was last in the Up condition.
# Up trans	Number of transitions from Down to Up condition.
Status	Status of the (local or remote circuit) local interface: • Up—Operational • Dn—Down • NP—Not present • DS—Disabled • WE—Wrong encapsulation • UN—Uninitialized
Encapsulation	Type of encapsulation: VPLS.
Remote PE	Address of the remote provider edge router.

Table 1785: show vpls connections Output Fields *(Continued)*

Field Name	Field Description
`Negotiated control-word`	Whether a control word has been negotiated: `Yes` or `No`.
`Incoming label`	Name of the incoming label.
`Outgoing label`	Name of the outgoing label.
`Negotiated PW status TLV`	Indicates whether or not the pseudowire status TLV has been negotiated for the VPLS connection.
`Local interface`	Provides the following information about the local interface configured for the VPLS neighbor: • Name of the local interface • `Status`—Interface status (`Up` or `Down`) • `Encapsulation`—Interface encapsulation (for example, `ETHERNET`) • `Description`—Includes the VPLS instance name, the VPLS neighbor address, and the VPLS identifier
`Time`	Date and time of VPLS connection event.
`Event`	Type of event.
`Interface/Lbl/PE`	Interface, label, or PE router.
`Connection History`	Each entry can include the date, time, year, and the connection event. Connection events include any of a variety of events related to VPLS connections, such as route changes, label updates, and interfaces going down or coming up.

Sample Output

show vpls connections

```
user@host> show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -< -- only outbound connection is up
CN -- circuit not provisioned    >- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unn connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy

Legend for interface status
Up -- operational
Dn -- down

Instance: vpls-1
  Local site: 1 (11)
    Number of local interfaces: 1
    Number of local interfaces up: 1
    IRB interface present: no
    lt-1/3/0.10496
    vt-1/3/0.1048588    1         Intf - vpls vpls-1 local site 11 remote site 1
    vt-1/2/0.1048591    2         Intf - vpls vpls-1 local site 11 remote site 2
    vt-1/2/0.1048585    3         Intf - vpls vpls-1 local site 11 remote site 3
    vt-1/2/0.1048587    4         Intf - vpls vpls-1 local site 11 remote site 4
    vt-1/2/0.1048589    5         Intf - vpls vpls-1 local site 11 remote site 5
    vt-1/3/0.1048586    6         Intf - vpls vpls-1 local site 11 remote site 6
    vt-1/3/0.1048590    7         Intf - vpls vpls-1 local site 11 remote site 7
    vt-1/3/0.1048584    8         Intf - vpls vpls-1 local site 11 remote site 8
```

```
    Label-base        Offset     Size  Range     Preference
    800256            1          16     16        100
    Timer Values:
        Startup wait time: 120 seconds
        New site wait-time: 20 seconds
        Collision detect time: 30 seconds
        Reclaim wait time: 748 milliseconds
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Apr 28 13:28:24 2009           2
      Remote PE: 192.0.2.1, Negotiated control-word: No
      Incoming label: 800256, Outgoing label: 800026
      Local interface: vt-1/3/0.1048588, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls vpls-1 local site 11 remote site 1
    Connection History:
        Apr 28 13:28:24 2009  status update timer
        Apr 28 13:28:24 2009  PE route down
        Apr 28 13:24:27 2009  status update timer
        Apr 28 13:24:27 2009  loc intf up             vt-1/3/0.1048588
        Apr 28 13:24:27 2009  PE route changed
        Apr 28 13:24:27 2009  Out lbl Update                    800026
        Apr 28 13:24:27 2009  In lbl Update                     800256
        Apr 28 13:24:27 2009  loc intf down
    2                         rmt   Up     Apr 28 13:28:24 2009           2
      Remote PE: 192.0.2.71, Negotiated control-word: No
      Incoming label: 800257, Outgoing label: 800034
      Local interface: vt-1/2/0.1048591, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls vpls-1 local site 11 remote site 2
    Connection History:
        Apr 28 13:28:24 2009  status update timer
        Apr 28 13:28:24 2009  PE route down
        Apr 28 13:24:28 2009  status update timer
        Apr 28 13:24:28 2009  loc intf up             vt-1/2/0.1048591
        Apr 28 13:24:28 2009  PE route changed
        Apr 28 13:24:28 2009  Out lbl Update                    800034
        Apr 28 13:24:28 2009  In lbl Update                     800257
        Apr 28 13:24:28 2009  loc intf down
```

show vpls connections (with FEC128 and FEC129 in the same routing-instance)

```
user@host> show vpls connections
Instance: fec129
```

```
  L2vpn-id: 1:1
  Local-id: 203.0.113.0
FEC129-VPLS State:
  Mesh-group connections: __ves__
    Remote-id                 Type  St     Time last up          # Up trans
    203.0.3.3                  rmt   Up     Sep 19 09:59:56 2017           1
      Remote PE: 203.0.3.3, Negotiated control-word: No
      Incoming label: 262155, Outgoing label: 262164
      Negotiated PW status TLV: No
      Local interface: lsi.1048844, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls fec129 local-id 10.4.4.4 remote-id 203.0.3.3 neighbor 203.0.3.3
      Flow Label Transmit: No, Flow Label Receive: No
    203.0.2.2                  rmt   Up     Sep 19 09:59:52 2017           1
      Remote PE: 203.0.2.2, Negotiated control-word: No
      Incoming label: 262154, Outgoing label: 262157
      Negotiated PW status TLV: No
      Local interface: lsi.1048846, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls fec129 local-id 10.4.4.4 remote-id 203.0.2.2 neighbor 203.0.2.2
      Flow Label Transmit: No, Flow Label Receive: No
    203.0.1.1                  rmt   Up     Sep 19 09:59:48 2017           1
      Remote PE: 203.0.1.1, Negotiated control-word: No
      Incoming label: 262156, Outgoing label: 262157
      Negotiated PW status TLV: No
      Local interface: lsi.1048845, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls fec129 local-id 10.4.4.4 remote-id 203.0.1.1 neighbor 203.0.1.1
      Flow Label Transmit: No, Flow Label Receive: No

LDP-VPLS State
   Mesh-group connections: MG1
     Neighbor                  Type  St     Time last up          # Up trans
     203.0.6.6(vpls-id 1)       rmt   Up     Sep 17 19:17:11 2017           1
       Remote PE: 203.0.6.6, Negotiated control-word: No
       Incoming label: 262423, Outgoing label: 262145
       Negotiated PW status TLV: No
       Local interface: lsi.1049859, Status: Up, Encapsulation: ETHERNET
         Description: Intf - vpls bgp-vpls neighbor 203.0.6.6 vpls-id 1
       Flow Label Transmit: No, Flow Label Receive: No
     203.0.7.7(vpls-id 1)       rmt   Up     Sep 17 19:17:04 2017           1
       Remote PE: 203.0.7.7, Negotiated control-word: No
       Incoming label: 262424, Outgoing label: 262145
       Negotiated PW status TLV: No
       Local interface: lsi.1049857, Status: Up, Encapsulation: ETHERNET
         Description: Intf - vpls bgp-vpls neighbor 203.0.7.7 vpls-id 1
```

```
      Flow Label Transmit: No, Flow Label Receive: No
  Mesh-group connections: MG2
    Neighbor                  Type  St     Time last up          # Up trans
    203.0.5.5(vpls-id 1)       rmt   Up     Sep 17 19:17:00 2017           1
      Remote PE: 203.0.5.5, Negotiated control-word: No
      Incoming label: 262425, Outgoing label: 299872
      Negotiated PW status TLV: No
      Local interface: lsi.1049856, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls bgp-vpls neighbor 203.0.5.5 vpls-id 1
      Flow Label Transmit: No, Flow Label Receive: No
```

show vpls connections (with multiple pseudowires)

```
user@host> show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: vpls
```

```
  VPLS-id: 100
  Mesh-group connections: __ves__
    Neighbor                  Type  St     Time last up          # Up trans
    10.255.114.3 (vpls-id 100)  rmt   Up     Apr 11 23:38:38 2013           1
      Remote PE: 10.255.114.3, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262145
      Negotiated PW status TLV: No
      Local interface: lsi.1049090, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls h-vpls neighbor 10.255.114.3 vpls-id 100

  Mesh-group connections: spokes
    Neighbor                  Type  St     Time last up          # Up trans
    10.255.114.4 (vpls-id 200)  rmt   Up     Apr 11 23:39:25 2013           1
      Remote PE: 10.255.114.4, Negotiated control-word: No
      Incoming label: 262148, Outgoing label: 304224
      Negotiated PW status TLV: Yes
      local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
      Local interface: lsi.1049091, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls h-vpls neighbor 10.255.114.4 vpls-id 200
      10.255.114.4 (vpls-id 201)  rmt   Up     Apr 11 23:39:25 2013           1
        Remote PE: 10.255.114.4, Negotiated control-word: No
        Incoming label: 262149, Outgoing label: 304225
        Negotiated PW status TLV: Yes
        local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
        Local interface: lsi.1049096, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls h-vpls neighbor 10.255.114.4 vpls-id 201
```

show vpls connections extensive (Static VPLS Neighbors)

```
user@host> show vpls connections extensive instance red
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range              Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
```

```
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unn connection status  IL -- no incoming label
MM -- MTU mismatch              MI -- Mesh-Group ID not available
BK -- Backup connection                ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby       SN -- Static Neighbor

Legend for interface status
Up -- operational
Dn -- down

Instance: static
  VPLS-id: 1
   Number of local interfaces: 1
   Number of local interfaces up: 1
   ge-0/0/5.0
   lsi.1049344                  Intf - vpls static neighbor 10.255.114.3 vpls-id 1
   Neighbor                 Type  St     Time last up          # Up trans
   10.255.114.3(vpls-id 1)(SN) rmt Up     Mar  4 08:48:41 2010           1
     Remote PE: 10.255.114.3, Negotiated control-word: No
     Incoming label: 29696, Outgoing label: 29697
     Negotiated PW status TLV: No
     Local interface: lsi.1049344, Status: Up, Encapsulation: ETHERNET
       Description: Intf - vpls static neighbor 10.255.114.3 vpls-id 1
   Connection History:
       Mar  4 08:48:41 2010  status update timer
       Mar  4 08:48:41 2010  PE route changed
       Mar  4 08:48:41 2010  Out lbl Update                    29697
       Mar  4 08:48:41 2010  In lbl Update                     29696
       Mar  4 08:48:41 2010  loc intf up                 lsi.1049344
```

command-name

```
user@PE1> show vpls connections extensive (Multihoming with FEC 129)
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
```

```
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: green
  L2vpn-id: 100:100
  Local-id: 192.0.2.2
    Number of local interfaces: 2
    Number of local interfaces up: 2
    ge-0/3/1.0
    ge-0/3/3.0
    lsi.101711873                 Intf - vpls green local-id 192.0.2.2 remote-id 192.0.2.4
neighbor 192.0.2.4
    Remote-id                 Type  St     Time last up          # Up trans
    192.0.2.4                   rmt   Up     Jan 31 13:49:52 2012           1
      Remote PE: 192.0.2.4, Negotiated control-word: No
      Incoming label: 262146, Outgoing label: 262146
      Local interface: lsi.101711873, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls green local-id 192.0.2.2 remote-id 192.0.2.4 neighbor 192.0.2.4
    Connection History:
        Jan 31 13:49:52 2012  status update timer
        Jan 31 13:49:52 2012  PE route changed
        Jan 31 13:49:52 2012  Out lbl Update                     262146
        Jan 31 13:49:52 2012  In lbl Update                      262146
        Jan 31 13:49:52 2012  loc intf up                lsi.101711873
  Multi-home:
```

```
    Local-site                Id     Pref   State
    test                      1      100    Up
      Number of interfaces: 1
      Number of interfaces up: 1
        ge-0/3/1.0
      Received multi-homing advertisements:
        Remote-PE         Pref   flag   Description
        192.0.2.4           100    0x0
```

Release Information

Command introduced before Junos OS Release 7.4.

`instance-history` option introduced in Junos OS Release 12.3R2.

show vpls flood event-queue

IN THIS SECTION

- Syntax | 27865
- Description | 27866
- Options | 27866
- Required Privilege Level | 27866
- Output Fields | 27866
- Sample Output | 27867
- Release Information | 27867

Syntax

```
show vpls flood event-queue
```

Description

Display the pending events in the VPLS flood queue.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 1786 on page 27866 lists the output fields for the `show vpls flood event-queue` command. Output fields are listed in the approximate order in which they appear.

Table 1786: show vpls flood event-queue Output Fields

Field Name	Field Description
Current Pending Event	Provides information on the current event in the VPLS flood event queue.
Name	Name of the event.
Owner Name	Name of the interface associated with the flood event.
Pending Op	Pending operation for the event.
Last Error	Name of the last error encountered.
Number of Retries	Number of attempts made to update the event queue.
Pending Event List	List of the events awaiting processing.
Event Name	Name of the event.

Table 1786: show vpls flood event-queue Output Fields *(Continued)*

Field Name	Field Description
Pending Op	Pending operation for the event.
Event Identifier	Name of the interface associated with the flood event.

Sample Output

show vpls flood event-queue

```
user@host> show vpls flood event-queue
Current Pending Event
    Name:          Flood Nexthop
    Owner Name:ge-4/3/0.0
    Pending Op: ADD
    Last Error:ENOMEM
    Number of Retries:3
    Pending Event List:
    Event Name      Pending Op      Event Identifier
    Flood Nexthop   ADD             ge-4/3/0.0
    Flood Route     ADD             ge-4/3/0.0
```

Release Information

Command introduced in Junos OS Release 8.0.

show vpls flood instance

IN THIS SECTION

- Syntax | 27868
- Description | 27868

- Options | 27868
- Required Privilege Level | 27868
- Output Fields | 27869
- Sample Output | 27869
- Release Information | 27870

Syntax

```
show vpls flood instance
<brief | detail | extensive>
 <instance-name>
<logical-system logical-system-name>
```

Description

Display VPLS information related to the flood process.

Options

none	Display VPLS information related to the flood process for all routing instances.
brief | detail | extensive	(Optional) Display the specified level of output.
instance-name	(Optional) Display VPLS information related to the flood process for the specified routing instance.
logical-system *logical-system-name*	(Optional) Display VPLS information related to the flood process for the specified logical system.

Required Privilege Level

view

Output Fields

Table 1787 on page 27869 lists the output fields for the `show vpls flood instance` command. Output fields are listed in the approximate order in which they appear.

Table 1787: show vpls flood instance Output Fields

Field Name	Field Description
Logical system	Name of the logical system.
Name	Name of the VPLS routing instance.
CEs	Number of CE routers connected to the VPLS instance.
VEs	Number of VE routers connected to the VPLS instance.
Flood routes	List of all flood routes associated with the VPLS instance.
Prefix	Prefix for the route.
Type	Type of route.
Owner	VPLS routing instance or interface associated with the route.
Nhtype	Next-hop type. For example, **flood** for a flood route.
Nhindex	Next-hop index number for the route.

Sample Output

show vpls flood instance

```
user@host> show vpls flood instance

Logical system: __example_ls1__
```

```
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  default   ALL_CE_FLOOD  green                flood           383
  0x47/16   CE_FLOOD      fe-1/2/1.0           flood           388
```

show vpls flood instance logical-system-name

```
user@host:__example_ls1__> show vpls flood instance example_ls1


Logical system: __example_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  default   ALL_CE_FLOOD  green                flood           383
  0x47/16   CE_FLOOD      fe-1/2/1.0           flood           388
```

show vpls flood instance detail

```
user@host:__example_ls1__> show vpls flood instance detail


Logical system: __example_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  default   ALL_CE_FLOOD  green                flood           383
  0x47/16   CE_FLOOD      fe-1/2/1.0           flood           388
```

Release Information

Command introduced in Junos OS Release 8.0.

show vpls flood route

IN THIS SECTION

- Syntax | 27871
- Description | 27871
- Options | 27871
- Required Privilege Level | 27872
- Output Fields | 27872
- Sample Output | 27873
- Release Information | 27874

Syntax

```
show vpls flood route
(all-ce-flood instance-name instance-name <logical-system-name logical-system-name> |
ce-flood interface interface-name)
```

Description

Display VPLS route information related to the flood process for either the specified routing instance or the specified interface.

Options

all-ce-flood	Display the flood next-hop route for all customer edge routers for traffic coming from the core of the network.
ce-flood interface *interface-name*	Display the flood next-hop route for traffic coming from the specified customer edge interface.
instance-name *instance-name*	Display the flood routes for the specified instance.

logical-system-name *logical-system-name*

(Optional) Specify the logical system whose flood routes you want to display. You can only specify the default logical system name for VPLS. The default logical system name is **__example_ls1__**\ (the name must be entered in the command with the underscore characters).

Required Privilege Level

view

Output Fields

Table 1788 on page 27872 lists the output for the `show vpls flood route` command. Output fields are listed in the approximate order in which they appear.

Table 1788: show vpls flood route Output Fields

Field Name	Field Description
Flood route prefix	Prefix for the flood route.
Flood route type	Type of flood route (either **CE_FLOOD** or **ALL_CE_FLOOD**).
Flood route owner	VPLS routing instance or interface associated with the flood route.
Nexthop type	Next-hop type. For example, **flood** for a flood route.
Nexthop index	Next-hop index number for the route.
Interfaces flooding to	Interfaces to which VPLS routes are being flooded.
Name	Name of the interface.
Type	Type of VPLS router (**CE** or **VE**).
Nh type	Next-hop type.

Table 1788: show vpls flood route Output Fields *(Continued)*

Field Name	Field Description
Index	Index number for the flood route.

Sample Output

show vpls flood route all-ce-flood

```
user@host:__example_ls1__> show vpls flood route all-ce-flood logical-system-name
__example_ls1__instance-name green


  Flood route prefix: default
  Flood route type: ALL_CE_FLOOD
  Flood route owner: green
  Nexthop type: flood
  Nexthop index: 383
    Interfaces Flooding to:
    Name             Type          NhType          Index
    fe-1/2/1.0       CE
```

show vpls flood route ce-flood

```
user@host:__example_ls1__> show vpls flood route ce-flood interface fe-1/2/1.0


  Flood route prefix: 0x47/16
  Flood route type: CE_FLOOD
  Flood route owner: fe-1/2/1.0
  Nexthop type: flood
  Nexthop index: 388
    Interfaces Flooding to:
    Name             Type          NhType          Index
    lsi.49152        VE             indr            262142
```

Release Information

Command introduced in Junos OS Release 8.0.

show vpls mac-move-action

IN THIS SECTION

- Syntax | 27874
- Description | 27874
- Options | 27874
- Required Privilege Level | 27875
- Output Fields | 27875
- Sample Output | 27875
- Release Information | 27876

Syntax

```
show vpls mac-move-action instance
<instance instance-name>
```

Description

Display the learning interfaces (IFLs) disabled due to a MAC move in a VPLS instance.

Options

none	Display all the IFLs disabled due to MAC moves.
`instance` ***instance-name***	(Optional) Display the IFLs disabled due to a MAC move in a VPLS instance.

Required Privilege Level

view

Output Fields

Table 1789 on page 27875 describes the output fields for the **show vpls mac-move-action** command. Output fields are listed in the approximate order in which they appear.

Table 1789: show vpls mac-move-action Output Fields

Field Name	Field Description
`Instance`	Instance name
`Local Interface`	Local interface disabled due to MAC move.
`Index`	Index number.
`Algorithm Used`	The following algorithm can be in use: • Base IFL — Primary approach • Statistical — Secondary approach
`Time rec`	Time recorded in hh:mm:sec.
`Recovery Timer`	Enable the interface after a set time.

Sample Output

show vpls mac-move-action (MX Series)

```
user@host> show vpls mac-move-action

        Instance: vpls_1
       Local interface: xe-1/3/1.600, Index: 341
     Algorithm used : Statistical
```

```
     Time rec       : 02:30:35
     Recovery timer : Yes
```

show vpls mac-move-action instance instance-name

```
user@host> show vpls mac-move-action instance vpls_1

        Instance: vpls_1
       Local interface: xe-1/3/1.600, Index: 341
     Algorithm used : Base IFL
     Time rec       : 02:29:35
     Recovery timer : Yes
```

Release Information

Command introduced in Junos OS Release 14.2.

show vpls mac-table

IN THIS SECTION

- Syntax | 27877
- Description | 27877
- Options | 27877
- Required Privilege Level | 27878
- Output Fields | 27878
- Sample Output | 27879
- Release Information | 27884

Syntax

```
show vpls mac-table
<age>
<brief | detail | extensive | summary>
<bridge-domain bridge-domain-name>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
<mac-address>
<vlan-id vlan-id-number>
```

Description

Display learned virtual private LAN service (VPLS) media access control (MAC) address information.

Options

none	Display all learned VPLS MAC address information.
age	(Optional) Display age of a single mac-address.
brief | detail | extensive | summary	(Optional) Display the specified level of output.
bridge-domain *bridge-domain-name*	(Optional) Display learned VPLS MAC addresses for the specified bridge domain.
instance *instance-name*	(Optional) Display learned VPLS MAC addresses for the specified instance.
interface *interface-name*	(Optional) Display learned VPLS MAC addresses for the specified instance.
logical-system (all | *logical-system-name*)	(Optional) Display learned VPLS MAC addresses for all logical systems or for the specified logical system.
mac-address	(Optional) Display the specified learned VPLS MAC address information..
vlan-id *vlan-id-number*	(Optional) Display learned VPLS MAC addresses for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 1790 on page 27878 describes the output fields for the `show vpls mac-table` command. Output fields are listed in the approximate order in which they appear.

Table 1790: show vpls mac-table Output fields

Field Name	Field Description
Age	Age of a single mac-address.
Routing instance	Name of the routing instance.
Bridging domain	Name of the bridging domain.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: • `S`—Static MAC address configured. • `D`—Dynamic MAC address learned. • `SE`—MAC accounting is enabled. • `NM`—Nonconfigured MAC.
Logical interface	Name of the logical interface.
MAC count	Number of MAC addresses learned on a specific routing instance or interface.
Learning interface	Logical interface or logical Label Switched Interface (LSI) the address is learned on.
Base learning interface	Base learning interface of the MAC address. This field is introduced in Junos OS Release 14.2.

Table 1790: show vpls mac-table Output fields *(Continued)*

Field Name	Field Description
Learn VLAN ID/ VLAN	VLAN ID of the routing instance or bridge domain in which the MAC address was learned.
VXLAN ID/VXLAN	VXLAN Network Identifier (VNI)
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

Sample Output

show vpls mac-table

```
user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_ldp1
VLAN : 223
   MAC                 MAC      Logical
   address             flags    interface
   00:00:5e:00:53:5d   D        ge-0/2/5.400

MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)
```

```
Routing instance : vpls_red
VLAN : 401
   MAC                 MAC      Logical
   address             flags    interface
   00:00:5e:00:53:12   D        lsi.1051138
   00:00:5e:00:53:f0   D        lsi.1051138
```

show vpls mac-table (with Layer 2 Services over GRE Interfaces)

```
user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls_4site:1000
 Bridging domain : __vpls_4site:1000__,   MAC                 MAC      Logical
   address             flags    interface
   00:00:5e:00:53:f4   D,SE     ge-4/2/0.1000
   00:00:5e:00:53:33   D,SE     lsi.1052004
   00:00:5e:00:53:32   D,SE     lsi.1048840
   00:00:5e:00:53:14   D,SE     lsi.1052005
   00:00:5e:00:53:f7   D,SE     gr-1/2/10.10
```

show vpls mac-table (with VXLAN enabled)

```
user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls_4site:1000
 Bridging domain : __vpls_4site:1000__, VLAN : 4094,4093
  VXLAN: Id : 300, Multicast group: 233.252.0.1
   MAC                 MAC      Logical
   address             flags    interface
   00:00:5e:00:53:f4   D,SE     ge-4/2/0.1000
   00:00:5e:00:53:33   D,SE     lsi.1052004
   00:00:5e:00:53:32   D,SE     lsi.1048840
   00:00:5e:00:53:14   D,SE     lsi.1052005
```

```
   00:00:5e:00:53:f7   D,SE     vtep.1052010
   00:00:5e:00:53:3f   D,SE     vtep.1052011
```

show vpls mac-table age (for GE interface)

```
user@host> show vpls mac-table age 00:00:5e:00:53:1a instance vpls_instance_1
MAC Entry Age information
Current Age: 4 seconds
```

show vpls mac-table age (for AE interface)

```
user@host> show vpls mac-table age 000:00:5e:00:53:1a instance vpls_instance_1
MAC Entry Age information
Current Age on FPC1: 102 seconds
Current Age on FPC2: 94 seconds
```

show vpls mac-table count

```
user@host> show vpls mac-table count
0 MAC address learned in routing instance __example_private1__

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    lc-0/0/0.32769                   0
    lc-0/1/0.32769                   0
    lc-0/2/0.32769                   0
    lc-2/0/0.32769                   0
    lc-0/3/0.32769                   0
    lc-2/1/0.32769                   0
    lc-9/0/0.32769                   0
    lc-11/0/0.32769                  0
    lc-2/2/0.32769                   0
    lc-9/1/0.32769                   0
    lc-11/1/0.32769                  0
    lc-2/3/0.32769                   0
    lc-9/2/0.32769                   0
    lc-11/2/0.32769                  0
    lc-11/3/0.32769                  0
```

```
    lc-9/3/0.32769                  0

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID            MAC count
                0                    0

1 MAC address learned in routing instance vpls_ldp1

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    lsi.1051137                      0
    ge-0/2/5.400                     1

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID            MAC count
                0                    1

1 MAC address learned in routing instance vpls_red

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    ge-0/2/5.300                     1

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID            MAC count
                0                    1
```

show vpls mac-table detail

```
user@host> show vpls mac-table detail
MAC address: 00:00:5e:00:53:5d
  Routing instance: vpls_ldp1
   Learning interface: ge-0/2/5.400
   Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
   Epoch: 0                            Sequence number: 1
   Learning mask: 0x1                  IPC generation: 0

MAC address: 00:00:5e:00:53:5d
  Routing instance: vpls_red
   Learning interface: ge-0/2/5.300
   Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
```

```
    Epoch: 0                            Sequence number: 1
    Learning mask: 0x1                  IPC generation: 0
```

show vpls mac-table extensive

```
user@host> show vpls mac-table extensive

MAC address: 00:00:5e:00:53:00
  Routing instance: vpls_1
   Bridging domain: __vpls_1__, VLAN : NA
   Learning interface: lsi.1049165
   Base learning interface: lsi.1049165
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 1
   Learning mask: 0x00000001

MAC address: 00:00:5e:00:53:01
  Routing instance: vpls_1
   Bridging domain: __vpls_1__, VLAN : NA
   Learning interface: lsi.1049165
   Base learning interface: lsi.1049165
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 1
   Learning mask: 0x00000001

MAC address: 00:00:5e:00:53:02
  Routing instance: vpls_1
   Bridging domain: __vpls_1__, VLAN : NA
   Learning interface: lsi.1049165
   Base learning interface: lsi.1049165
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 1
   Learning mask: 0x00000001

MAC address: 00:00:5e:00:53:03
  Routing instance: vpls_1
   Bridging domain: __vpls_1__, VLAN : NA
   Learning interface: lsi.1049165
   Base learning interface: lsi.1049165
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
```

```
    Epoch: 0                          Sequence number: 1
    Learning mask: 0x00000001
```

Release Information

Command introduced in Junos OS Release 8.5.

show vpls statistics

IN THIS SECTION

- Syntax | 27884
- Description | 27884
- Options | 27885
- Required Privilege Level | 27885
- Output Fields | 27885
- Sample Output | 27886
- Release Information | 27887

Syntax

```
show vpls statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) statistics.

Options

none	Display VPLS statistics for all routing instances.
instance *instance-name*	(Optional) Display VPLS statistics for a specific VPLS routing instance only.
logical-system (all | *logical-system-name*)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1791 on page 27885 lists the output fields for the `show vpls statistics` command. Output fields are listed in the approximate order in which they appear.

Table 1791: show vpls statistics Output Fields

Field Name	Field Description
Instance	Name of the VPLS instance.
Local interface	Name of the local VPLS virtual loopback tunnel interface, **vt-*fpc/pic/port.nnnnn***, where ***nnnnn*** is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain.
Index	Number associated with the next hop.
Remote provider edge router	Address of the remote provider edge router.
Multicast packets	Number of multicast packets received.
Multicast bytes	Number of multicast bytes received.

Table 1791: show vpls statistics Output Fields *(Continued)*

Field Name	Field Description
Flood packets	Number of VPLS flood packets received.
Flood bytes	Number of VPLS flood bytes received.
Current MAC count	Number of MAC addresses learned by the interface and the configured maximum limit on the number of MAC addresses that can be learned.

Sample Output

show vpls statistics

```
user@host> show vpls statistics

VPLS statistics:

Instance: green

   Local interface: fe-2/2/1.0, Index: 69
     Multicast packets:                     1
     Multicast bytes  :                    60
     Flooded packets  :                    18
     Flooded bytes    :                  2556
     Current MAC count:                     1

   Local interface: lt-0/3/0.2, Index: 72
     Multicast packets:                     3
     Multicast bytes  :                   153
     Flooded packets  :                     1
     Flooded bytes    :                    51
     Current MAC count:                     1

   Local interface: lsi.32769, Index: 75
     Current MAC count:                     0

   Local interface: lsi.32771, Index: 77
```

```
   Remote PE: 10.255.14.222
     Current MAC count:                    2

Instance: red

   Local interface: vt-0/3/0.32768, Index: 74
     Multicast packets:                    0
     Multicast bytes  :                    0
     Flooded packets  :                    0
     Flooded bytes    :                    0
     Current MAC count:                    0

   Local interface: vt-0/3/0.32770, Index: 76
     Multicast packets:                    0
     Multicast bytes  :                    0
     Flooded packets  :                    0
     Flooded bytes    :                    0
     Current MAC count:                    0
```

show vpls statistics instance

```
user@host> show vpls statistics instance red


Layer-2 VPN Statistics:
Instance: red

   Local interface: vt-3/2/0.32768, Index: 73
   Remote provider edge router: 10.255.17.35
     Multicast packets:                    0
     Multicast bytes  :                    0
     Flood packets    :                    0
     Flood bytes      :                    0
     Current MAC count:                    1 (Limit 20)
```

Release Information

Command introduced before Junos OS Release 7.4.

show vrrp

IN THIS SECTION

- Syntax | 27888
- Description | 27888
- Options | 27888
- Required Privilege Level | 27889
- Output Fields | 27889
- Sample Output | 27899
- Release Information | 27905

Syntax

```
show vrrp
<brief | detail | extensive | summary>
<interface interface-name <group number>>
<logical-system logical-system-name >
<nsr>
```

Description

Display status information about Virtual Router Redundancy Protocol (VRRP) groups.

Options

none	(Same as `brief`) Display brief status information about all VRRP interfaces.
`brief | detail | extensive | summary`	(Optional) Display the specified level of output.
`interface` *`interface-name`* `<group` *`number`*`>`	(Optional) Display information and status about the specified VRRP interface and, optionally, the group number.
`logical-system` *`logical-system-name`*	(Optional) Perform this operation on a particular logical system.

nsr	(Optional) Display state replication information when graceful Routing Engine switchover (GRES) with nonstop active routing (NSR) is configured. Use only on the backup Routing Engine.

Required Privilege Level

view

Output Fields

Table 1792 on page 27889 lists the output fields for the `show vrrp` command. Output fields are listed in the approximate order in which they appear

Table 1792: show vrrp Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the logical interface.	`brief` `extensive` none `summary`
`Interface index`	Physical interface index number, which reflects its initialization sequence.	`extensive`
`Groups`	Total number of VRRP groups configured on the interface.	`extensive`
`Active`	Total number of VRRP groups that are active (that is, whose interface state is either up or down).	`extensive`

Table 1792: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface VRRP PDU statistics`	Non-errored statistics for the logical interface: • `Advertisement sent`—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted. • `Advertisement received`—Number of VRRP advertisement PDUs received by the interface. • `Packets received`—Number of VRRP packets received for VRRP groups on the interface. • `No group match received`—Number of VRRP packets received for VRRP groups that do not exist on the interface.	`extensive`
`Interface VRRP PDU error statistics`	Errored statistics for the logical interface: • `Invalid IPAH next type received`—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets. • `Invalid VRRP ttl value received`—Number of packets received whose IP time- to-live (TTL) value is not 255. • `Invalid VRRP version received`—Number of packets received whose VRRP version is not 2. • `Invalid VRRP pdu type received`—Number of packets received whose VRRP PDU type is not 1. • `Invalid VRRP authentication type received`—Number of packets received whose VRRP authentication is not none, simple, or md5. • `Invalid VRRP IP count received`—Number of packets received whose VRRP IP count exceeds 8. • `Invalid VRRP checksum received`—Number of packets received whose VRRP checksum does not match the calculated one.	`extensive`

Table 1792: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Physical interface`	Name of the physical interface.	`detail extensive`
`Unit`	Logical unit number.	All levels
`Address`	Address of the physical interface.	`brief detail extensive` none
`Index`	Physical interface index number, which reflects its initialization sequence.	`detail extensive`
`SNMP ifIndex`	SNMP index number for the physical interface.	`detail extensive`
`VRRP-Traps`	Status of VRRP traps: `Enabled` or `Disabled`.	`detail extensive`
`VRRP-Version`	VRRP version: `2` or `3`.	`detail extensive`
`Type and Address`	Identifier for the address and the address itself: • `lcl`—Configured local interface address. • `mas`—Address of the primary virtual router. This address is displayed only when the local interface is acting as a backup router. • `vip`—Configured virtual IP addresses.	`brief` none `summary`

Table 1792: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface state/Int state/State`	State of the physical interface: • `down`—The device is present and the link is unavailable. • `not present`—The interface is configured, but no physical device is present. • `unknown`—The VRRP process has not had time to query the kernel about the state of the interface. • `up`—The device is present and the link is established.	`brief` `extensive` none `summary`
`Group`	VRRP group number.	`brief` `extensive` none `summary`
`State`	The state of the interface on which VRRP is running: • `backup`—The interface is acting as the backup router interface. • `bringup`—VRRP is just starting and the physical device is not yet present. • `idle`—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • `init`—VRRP is initializing. • `master`—The interface is acting as the primary router interface. • `master(ISSU)`—The primary router interface is going through a unified in-service software upgrade. • `transition`—The interface is changing between being the backup and being the primary router.	`extensive`

Table 1792: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`VRRP Mode`	If the interface inherits its state and configuration from the active VRRP group, or if it is part of the active VRRP group. • Active—Part of the active VRRP group • Inherit—Inherits state and configuration from the active VRRP group.	`detail extensive`
`Priority`	Configured VRRP priority for the interface.	`detail extensive`
`Advertisement interval`	Configured VRRP advertisement interval.	`detail extensive`
`Authentication type`	Configured VRRP authentication type: `none`, `simple`, or `md5`.	`detail extensive`
`Advertisement Threshold`	A value from 1 through 15, used for setting the time when a peer should be considered down. • The time a peer is considered down is equal to the advertisement-threshold multiplied by the advertisement-interval. • (advertisement-threshold *advertisement-interval) = Peer down.	`detail extensive`
`Computed Send Rate`	How many protocol data units (PDUs) are generated per second. Based on the number of instances and the advertisement interval.	`detail extensive`
`Preempt`	Whether preemption is allowed on the interface: `yes` or `no`.	`detail extensive`
`Accept-data mode`	Whether the interface is configured to accept packets destined for the virtual IP address: `yes` or `no`.	`detail extensive`

Table 1792: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`VIP count`	Number of virtual IP addresses that have been configured on the interface.	`detail extensive`
`VIP`	List of virtual IP addresses configured on the interface.	`detail extensive`
`Advertisement timer`	How long, in seconds, until the advertisement timer expires.	`detail extensive`
`Master router`	IP address of the interface that is acting as the primary. If the VRRP interface is down, the output is `N/A`.	`detail extensive`
`Virtual router uptime`	How long, in seconds, that the virtual router has been up.	`detail extensive`
`Master router uptime`	How long, in seconds, that the primary route has been up.	`detail extensive`
`Virtual MAC`	MAC address associated with the virtual IP address.	`detail extensive`
`Tracking`	Whether tracking is `enabled` or `disabled`.	`detail extensive`
`Current priority`	Current operational priority for being the VRRP primary.	`detail extensive`
`Configured priority`	Configured base priority for being the VRRP primary.	`detail extensive`
`Priority hold-time`	Minimum time interval, in seconds, between successive changes to the current priority. `Disabled` indicates no minimum interval.	`detail extensive`
`Remaining-time`	(`track` option only) Displays the time remaining in the priority hold-time interval.	detail
`Interface tracking`	Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.	`detail extensive`

Table 1792: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface/Tracked interface/Track Int`	Name of the tracked interface.	`detail extensive`
`Int state/Interface state/State`	Current operational state of the tracked interface: `up` or `down`.	`detail extensive`
`Int speed/Speed`	Current operational speed, in bits per second, of the tracked interface.	`detail extensive`
`Incurred priority cost`	Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.	`detail extensive`
`Threshold`	Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred. An entry of `down` means that the corresponding priority cost is incurred when the interface is down.	`detail extensive`
`Route tracking`	Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.	`detail extensive`
`Route count`	The number of routes being tracked.	`detail extensive`
`Route`	The IP address of the route being tracked.	`detail extensive`
`VRF name`	The VPN routing and forwarding (VRF) routing instance that the tracked route is in.	`detail extensive`
`Route state`	The state of the route being tracked: `up`, `down`, or `unknown`.	`detail extensive`

Table 1792: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Priority cost`	Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.	`detail extensive`
`Active`	Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.	`detail extensive`
`Group VRRP PDU statistics`	Number of VRRP advertisements sent and received by the group.	`extensive`
`Group VRRP PDU error statistics`	Errored statistics for the VRRP group: • `Bad authentication type received`—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be `none`, `simple`, or `md5` and must be the same for all routers in the VRRP group. • `Bad password received`—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group • `Bad MD5 digest received`—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router. • `Bad advertisement timer received`—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group. • `Bad VIP count received`—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance. • `Bad VIPADDR received`—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance.	`extensive`

Table 1792: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Group state transition statistics`	State transition statistics for the VRRP group: • `Idle to master transitions`—Number of times that the VRRP instance transitioned from the idle state to the primary state. • `Idle to backup transitions`—Number of times that the VRRP instance transitioned from the idle state to the backup state. • `Backup to master transitions`—Number of times that the VRRP instance transitioned from the backup state to the primary state. • `Master to backup transitions`—Number of times that the VRRP instance transitioned from the primary state to the backup state.	`extensive`

Table 1792: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`VR state`	The state of the VRRP: • `backup`—The interface is acting as the backup router interface. • `bringup`—VRRP is just starting, and the physical device is not yet present. • `idle`—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • `init`—VRRP is initializing. • `master`—The interface is acting as the primary router interface. • `transition`—The interface is changing between being the backup and being the primary router. **NOTE**: When show vrrp nsr is used on the backup Routing Engine, it displays the current VRRP state on the primary Routing Engine, which is the future VRRP state for the backup Routing Engine. Do not use on the primary Routing Engine.	`brief` none `summary`

Table 1792: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
NSR	VRRP nonstop active routing is enabled for the configured VRRP group: yes or no. **NOTE**: A yes value means that the new primary Routing Engine will immediately start with the VRRP State value from the original primary Routing Engine. A no value means that the VRRP session will: • Start afresh. • Go through asilent startup period. • Move to a backup state. • Wait for the D Timer to run out before becoming the primary (only if the primary has not been configured already).	brief none
RPD-NSR	The routing options have been set to nonstop active routing: yes or no.	brief none
Timer	VRRP timer information: • A—How long, in seconds, until the advertisement timer expires. • D—How long, in seconds, until the Primary is Down timer expires.	brief none

Sample Output

show vrrp

```
user@host> show vrrp
Interface            State        Group       VR state    Timer            Type     Address
fe-0/0/0.121         up               1        master      A 1.052        lcl      fec0::12:1:1:1
                                                                      vip      fe80::12:1:1:99
                                                                      vip      fec0::12:1:1:99
```

```
fe-0/0/2.131     up                  1         master      A 0.364     lcl      fec0::13:1:1:1
                                                                     vip      fe80::13:1:1:99
                                                                     vip      fec0::13:1:1:99
```

show vrrp brief

The output for the `show vrrp brief` command is identical to that for the `show vrrp` command. For sample output, see "show vrrp" on page 27899.

show vrrp detail (IPv6)

```
user@host> show vrrp detail
Physical interface: fe-0/0/0, Unit: 121, Vlan-id: 212, Address: fec0::12:1:1:1/120
  Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master, VRRP Mode: Active
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: fe80::12:1:1:99, fec0::12:1:1:99
  Advertisement timer: 1.121s, Master router: fe80::12:1:1:1
  Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled

Physical interface: fe-0/0/2, Unit: 131, Vlan-id: 213, Address: fec0::13:1:1:1/120
  Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: fe80::13:1:1:99, fec0::13:1:1:99
  Advertisement timer: 0.327s, Master router: fe80::13:1:1:1
  Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
```

show vrrp detail (Route Track)

```
user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 1, Vlan-id: 1, Address: 101.1.1.1/24
  Index: 324, SNMP ifIndex: 623, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 1, State: master(ISSU), VRRP Mode: Active
```

```
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 101.1.1.3
  Advertisement Timer: 0.469s, Master router: 101.1.1.1
  Virtual router uptime: 00:02:10, Master router uptime: 00:02:05
  Virtual Mac: 00:00:5e:00:01:01
  Tracking: disabled
```

show vrrp detail (Route Track)

```
user@host> show vrrp detail
Physical interface: ge-1/2/0, Unit: 0, Address: 30.30.30.30/24
  Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 100, State: master
  Priority: 150, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100
  Advertisement timer: 1.218s, Master router: 30.30.30.30
  Virtual router uptime: 00:04:28, Master router uptime: 00:00:13
  Virtual MAC: 00:00:5e:00:01:64
  Tracking: enabled
    Current priority: 150, Configured priority: 150
    Priority hold-time: disabled
    Interface tracking: disabled
    Route tracking: enabled, Route count: 1
      Route                VRF name       Route state   Priority cost
      192.168.40.0/22      default        up            30
```

show vrrp extensive

```
user@host> show vrrp extensive
Interface: ge-2/0/0.0, Interface index :65539, Groups: 1, Active :1
  Interface VRRP PDU statistics
    Advertisement sent                       :0
    Advertisement received                   :0
    Packets received                         :0
    No group match received                  :0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received          :0
    Invalid VRRP TTL value received          :0
    Invalid VRRP version received            :0
```

```
    Invalid VRRP PDU type received           :0
    Invalid VRRP authentication type received:0
    Invalid VRRP IP count received           :0
    Invalid VRRP checksum received           :0

Physical interface: ge-2/0/0, Unit: 0, Address: 10.10.10.1/24
  Index: 65539, SNMP ifIndex: 648, VRRP-Traps: enabled, VRRP-Version: 3
  Interface state: up, Group: 1, State: backup, VRRP Mode: Active
  Priority: 100, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 10.10.10.2
  Dead timer: 3.078s, Master priority: 0, Master router: 10.10.10.1
  Virtual router uptime: 00:00:04
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent                       :0
    Advertisement received                   :0
  Group VRRP PDU error statistics
    Bad authentication Type received         :0
    Bad password received                    :0
    Bad MD5 digest received                  :0
    Bad advertisement timer received         :0
    Bad VIP count received                   :0
    Bad VIPADDR received                     :0
  Group state transition statistics
    Idle to master transitions               :0
    Idle to backup transitions               :1
    Backup to master transitions             :0
    Master to backup transitions             :0
```

show vrrp interface

```
user@host> show vrrp interface ge-0/0/0.1
Interface: ge-0/0/0.1, Interface index :324, Groups: 2, Active :2
  Interface VRRP PDU statistics
    Advertisement sent                       :39
    Advertisement received                   :0
    Packets received                         :0
    No group match received                  :0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received          :0
```

```
    Invalid VRRP TTL value received          :0
    Invalid VRRP version received            :0
    Invalid VRRP PDU type received           :0
    Invalid VRRP authentication type received:0
    Invalid VRRP IP count received           :0
    Invalid VRRP checksum received           :0

Physical interface: ge-0/0/0, Unit: 1, Vlan-id: 1, Address: 101.1.1.1/24
  Index: 324, SNMP ifIndex: 623, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 1, State: master(ISSU), VRRP Mode: Active
  Advertisement threshold: 3, Computed send rate: 0
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 101.1.1.3
  Advertisement Timer: 0.619s, Master router: 101.1.1.1
  Virtual router uptime: 00:00:22, Master router uptime: 00:00:17
  Virtual Mac: 00:00:5e:00:01:01
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent                       :20
    Advertisement received                   :0
  Group VRRP PDU error statistics
    Bad authentication Type received         :0
    Bad password received                    :0
    Bad MD5 digest received                  :0
    Bad advertisement timer received         :0
    Bad VIP count received                   :0
    Bad VIPADDR received                     :0
  Group state transition statistics
    Idle to master transitions               :0
    Idle to backup transitions               :1
    Backup to master transitions             :1
    Master to backup transitions             :0
Interface: fe-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                       :          205
    Advertisement received                   :            0
    Packets received                         :            0
    No group match received                  :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received          :            0
    Invalid VRRP TTL value received          :            0
    Invalid VRRP version received            :            0
```

```
    Invalid VRRP PDU type received           :            0
    Invalid VRRP authentication type received:            0
    Invalid VRRP IP count received           :            0
    Invalid VRRP checksum received           :            0
```

show vrrp nsr

This command is similar to `show vrrp`. Here, the **VR state** column displays the current VRRP state on the primary Routing Engine, which is the future VRRP state for the backup Routing Engine. Do not use on the primary Routing Engine.

NSR is yes if VRRP nonstop active routing is enabled for the configured VRRP group.

RPD-NSR is yes if the routing options have been set to nonstop active routing.

```
user@host>show vrrp nsr
Interface  State  Group  VR state  VR Mode  Type   NSR  RPD-NSR  Address
ge-1/0/1.0  up    1       master    Active    lcl  yes    yes    10.0.0.1
                                              vip               10.0.0.3
ge-1/0/1.0  up    2       master    Active    lcl  yes    yes    20.0.0.1
                                              vip               20.0.0.3
ge-1/0/1.0  up    3       master    Active    lcl  yes    yes    30.0.0.1
                                              vip               30.0.0.3
ge-1/0/1.0  up    4       master    Active    lcl  yes    yes    40.0.0.1
                                              vip               40.0.0.3
ge-1/0/1.0  up    5       master    Active    lcl  yes    yes    50.0.0.1
                                              vip               50.0.0.3
ge-1/0/1.0  up    1       master    Active    lcl  yes    yes    1000::1
                                              vip               fe80::200:5eff:fe00:1
                                              vip               1000::3
ge-1/0/1.0  up    2       master    Active    lcl  yes    yes    2000::1
                                              vip               fe80::200:5eff:fe00:2
                                              vip               2000::3
ge-1/0/1.0  up    3       master    Active    lcl  yes    yes    3000::1
                                              vip               fe80::200:5eff:fe00:3
                                              vip               3000::3
ge-1/0/1.0  up    4       master    Active    lcl  yes    yes    4000::1
                                              vip               fe80::200:5eff:fe00:4
                                              vip               4000::3
ge-1/0/1.0  up    5       master    Active    lcl  yes    yes    5000::1
```

```
                                                vip                 fe80::200:5eff:fe00:5
                                                vip                 5000::3
```

show vrrp summary

```
user@host> show vrrp summary
Interface     State       Group   VR state      Type    Address
ge-4/2/0.0    up              1   backup        lcl     10.57.0.2
                                                vip     10.57.0.100
```

Release Information

Command introduced before Junos OS Release 7.4.

`nsr` option added in Junos OS Release 13.2.

RELATED DOCUMENTATION

show vrrp track | 27920

clear vrrp | 4290

show vrrp

IN THIS SECTION

- Syntax | 27906
- Description | 27906
- Options | 27906
- Required Privilege Level | 27906
- Output Fields | 27906
- Sample Output | 27914
- Release Information | 27920

Syntax

```
show vrrp
<brief | detail | extensive | summary>
<interface interface-name>
<logical-system logical-system-name>
<nsr>
<track track-interfaces>
```

Description

Display information and status about VRRP groups.

Options

none	(Same as brief) Display brief status information about all VRRP interfaces.
brief | detail | extensive | summary	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display information and status about the specified VRRP interface.
logical-system	(Optional) Display information and status about the specified logical system.
nsr	(Optional) Display information and status about the primary routing engine.
track *track-interfaces*	(Optional) Display information and status about VRRP track interfaces.

Required Privilege Level

view

Output Fields

Table 1793 on page 27907 lists the output fields for the `show vrrp` command. Output fields are listed in the approximate order in which they appear.

Table 1793: show vrrp Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the logical interface.	**none, brief, extensive, summary**
Interface index	Physical interface index number, which reflects its initialization sequence.	**extensive**
Groups	Total number of VRRP groups configured on the interface.	**extensive**
Active	Total number of VRRP groups that are active (that is, whose interface state is either up or down).	**extensive**
Interface VRRP PDU statistics	Nonerrored statistics for the logical interface: • **Advertisement sent**Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted. • **Advertisement received**Number of VRRP advertisement PDUs received by the interface. • **Packets received**Number of VRRP packets received for VRRP groups on the interface. • **No group match received**Number of VRRP packets received for VRRP groups that do not exist on the interface.	**extensive**

Table 1793: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Interface VRRP PDU error statistics	Errored statistics for the logical interface: • **Invalid IPAH next type received**Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets. • **Invalid VRRP ttl value received**Number of packets received whose IP time- to-live (TTL) value is not255. • **Invalid VRRP version received**Number of packets received whose VRRP version is not 2. • **Invalid VRRP pdu type received**Number of packets received whose VRRP PDU type is not 1. • **Invalid VRRP authentication type received**Number of packets received whose VRRP authentication is not none, simple, or md5. • **Invalid VRRP IP count received**Number of packets received whose VRRP IP count exceeds 8. • **Invalid VRRP checksum received**Number of packets received whose VRRP checksum does not match the calculated value.	**extensive**
Physical interface	Name of the physical interface.	**detail, extensive**
Unit	Logical unit number.	All levels
Address	Address of the physical interface.	**none, brief, detail, extensive**
Index	Physical interface index number, which reflects its initialization sequence.	**detail, extensive**
SNMP ifIndex	SNMP index number for the physical interface.	**detail, extensive**
VRRP-Traps	Status of VRRP traps: **Enabled** or **Disabled**.	**detail, extensive**

Table 1793: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Type and Address	Identifier for the address and the address itself: • **lcl**Configured local interface address. • **mas**Address of the primary device. This address is displayed only when the local interface is acting as a backup device. • **vip**Configured virtual IP addresses.	**none, brief, summary**
Interface state or **Int state**	State of the physical interface: • **down**The device is present and the link is unavailable. • **not present**The interface is configured, but no physical device is present. • **unknown**The VRRP process has not had time to query the kernel about the state of the interface. • **up**The device is present and the link is established.	**none, brief, extensive, summary**
Group	VRRP group number.	**none, brief, extensive, summary**
State	VRRP state: • **backup**The interface is acting as the backup device interface. • **bringup**VRRP is just starting, and the physical device is not yet present. • **idle**VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • **initializing**VRRP is initializing. • **master**The interface is acting as the primary device interface. • **transition**The interface is changing between being the backup and being the primary device.	**extensive**

Table 1793: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Priority	Configured VRRP priority for the interface.	**detail, extensive**
Advertisement interval	Configured VRRP advertisement interval.	**detail, extensive**
Authentication type	Configured VRRP authentication type: **none**, **simple**, or **md5**.	**detail, extensive**
Preempt	Whether preemption is allowed on the interface: **yes** or **no**.	**detail, extensive**
Accept-data mode	Whether the interface is configured to accept packets destined for the virtual IP address: **yes** or **no**.	**detail, extensive**
VIP count	Number of virtual IP addresses that have been configured on the interface.	**detail, extensive**
VIP	List of virtual IP addresses configured on the interface.	**detail, extensive**
Advertisement timer	Time until the advertisement timer expires.	**detail, extensive**
Master router	IP address of the interface that is acting as the primary. If the VRRP interface is down, the output is **N/A**.	**detail, extensive**
Virtual router uptime	Time that the virtual device has been up.	**detail, extensive**
Master router uptime	Time that the primary device has been up.	**detail, extensive**
Virtual MAC	MAC address associated with the virtual IP address.	**detail, extensive**
Tracking	Whether tracking is **enabled** or **disabled**.	**detail, extensive**

Table 1793: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Current priority	Current operational priority for being the VRRP primary.	**detail, extensive**
Configured priority	Configured base priority for being the VRRP primary.	**detail, extensive**
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority. **Disabled** indicates no minimum interval.	**detail, extensive**
Remaining-time	(**track** option only) Displays the time remaining in the priority hold-time interval.	**detail**
Interface tracking	Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.	**detail extensive**
Interface/ Tracked interface	Name of the tracked interface.	**detail extensive**
Int state/ Interface state	Current operational state of the tracked interface: **up** or **down**.	**detail, extensive**
Int speed/Speed	Current operational speed, in bits per second, of the tracked interface.	**detail, extensive**
Incurred priority cost	Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.	**detail, extensive**
Threshold	Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred. An entry of **down** means that the corresponding priority cost is incurred when the interface is down.	**detail, extensive**

Table 1793: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Route tracking	Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.	**detail, extensive**
Route count	The number of routes being tracked.	**detail, extensive**
Route	The IP address of the route being tracked.	**detail, extensive**
VRF name	The VPN routing and forwarding (VRF) routing instance that the tracked route is in.	**detail, extensive**
Route state	The state of the route being tracked: **up**, **down**, or **unknown**.	**detail, extensive**
Priority cost	Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.	**detail, extensive**
Active	Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.	**detail, extensive**
Group VRRP PDU statistics	Number of VRRP advertisements sent and received by the group.	**extensive**

Table 1793: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Group VRRP PDU error statistics	Errored statistics for the VRRP group: • **Bad authentication type received**Number of VRRP PDUs received with an invalid authentication type. The received authentication can be **none**, **simple**, or **md5** and must be the same for all devices in the VRRP group. • **Bad password received**Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all devices in the VRRP group • **Bad MD5 digest received**Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the device. • **Bad advertisement timer received**Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the devices in the VRRP group. • **Bad VIP count received**Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance. • **Bad VIPADDR received**Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance.	**extensive**
Group state transition statistics	State transition statistics for the VRRP group: • **Idle to master transitions**Number of times that the VRRP instance transitioned from the idle state to the primary state. • **Idle to backup transitions**Number of times that the VRRP instance transitioned from the idle state to the backup state. • **Backup to master transitions**Number of times that the VRRP instance transitioned from the backup state to the primary state. • **Master to backup transitions**Number of times that the VRRP instance transitioned from the primary state to the backup state.	**extensive**

Table 1793: show vrrp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Vlan-id	ID of Vlan	**detail**
VR state	VRRP information: • **backup**—The interface is acting as the backup device interface. • **bringup**—VRRP is just starting, and the physical device is not yet present. • **idle**—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • **initializing**—VRRP is initializing. • **master**—The interface is acting as the primary device interface. • **transition**—The interface is changing between being the backup and being the primary device.	**none, brief**
Timer	VRRP timer information: • **A**—Time, in seconds, until the advertisement timer expires. • **D**—Time, in seconds, until the Primary is Dead timer expires.	**none, brief**

Sample Output

show vrrp

```
user@host> show vrrp
Interface            State        Group       VR state    Timer            Type     Address
ge-0/0/0.121         up               1        master      A 1.052        lcl
ge2001:db8::12:1:1:1
                                                                     vip
ge2001:db8:0:1:12:1:1:99
                                                                     vip
ge2001:db8::12:1:1:99
```

```
ge-0/0/2.131     up                  1         master      A 0.364     lcl
ge2001:db8::13:1:1:1
                                                                 vip
ge2001:db8:0:1:13:1:1:99
                                                                 vip
ge2001:db8::13:1:1:99
```

show vrrp brief

The output for the `show vrrp brief` command is identical to that for the `show vrrp` command. For sample output, see "show vrrp" on page 27914.

show vrrp detail (IPv6)

```
user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: ge2001:db8::12:1:1:1/120
  Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge2001:db8:0:1:12:1:1:99,
ge2001:db8::12:1:1:99
  Advertisement timer: 1.121s, Master router: ge2001:db8:0:1:12:1:1:1
  Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: ge2001:db8::13:1:1:1/120
  Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge2001:db8:0:1:13:1:1:99,
ge2001:db8::13:1:1:99
  Advertisement timer: 0.327s, Master router: ge2001:db8:0:1:13:1:1:1
  Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
```

show vrrp detail (Route Track)

```
user@host> show vrrp detail
Physical interface: ge-1/1/0, Unit: 0, Address: 192.0.2.30/24
  Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled
  Interface state: up, Group: 100, State: master
  Priority: 150, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 192.0.2.100
  Advertisement timer: 1.218s, Master router: 192.0.2.30
  Virtual router uptime: 00:04:28, Master router uptime: 00:00:13
  Virtual MAC: 00:00:5e:00:01:64
  Tracking: enabled
    Current priority: 150, Configured priority: 150
    Priority hold-time: disabled
    Interface tracking: disabled
    Route tracking: enabled, Route count: 1
      Route                VRF name       Route state   Priority cost
      198.51.100.0/24      default        up            30
```

show vrrp extensive

```
user@host> show vrrp extensive
Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                      :          188
    Advertisement received                  :            0
    Packets received                        :            0
    No group match received                 :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received         :            0
    Invalid VRRP TTL value received         :            0
    Invalid VRRP version received           :            0
    Invalid VRRP PDU type received          :            0
    Invalid VRRP authentication type received:           0
    Invalid VRRP IP count received          :            0
    Invalid VRRP checksum received          :            0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: ge2001:db8::12:1:1:1/120
  Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
```

```
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge2001:db8:0:1:12:1:1:99,
ge2001:db8::12:1:1:99
  Advertisement timer: 1.034s, Master router: ge2001:db8:0:1:12:1:1:1
  Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent              :          188
    Advertisement received          :            0
  Group VRRP PDU error statistics
    Bad authentication type received:            0
    Bad password received           :            0
    Bad MD5 digest received         :            0
    Bad advertisement timer received:            0
    Bad VIP count received          :            0
    Bad VIPADDR received            :            0
  Group state transition statistics
    Idle to master transitions      :            0
    Idle to backup transitions      :            1
    Backup to master transitions    :            1
    Master to backup transitions    :            0

Interface: ge-0/0/2.131, Interface index: 69, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                        :          186
    Advertisement received                    :            0
    Packets received                          :            0
    No group match received                   :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received           :            0
    Invalid VRRP TTL value received           :            0
    Invalid VRRP version received             :            0
    Invalid VRRP PDU type received            :            0
    Invalid VRRP authentication type received:             0
    Invalid VRRP IP count received            :            0
    Invalid VRRP checksum received            :            0

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: ge2001:db8::13:1:1:1/120
  Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge2001:db8:0:1:13:1:1:99,
```

```
ge2001:db8::13:1:1:99
  Advertisement timer: 0.396s, Master router: ge2001:db8:0:1:13:1:1:1
  Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent              :          186
    Advertisement received          :            0
  Group VRRP PDU error statistics
    Bad authentication type received:            0
    Bad password received           :            0
    Bad MD5 digest received         :            0
    Bad advertisement timer received:            0
    Bad VIP count received          :            0
    Bad VIPADDR received            :            0
  Group state transition statistics
    Idle to master transitions      :            0
    Idle to backup transitions      :            1
    Backup to master transitions    :            1
     Master to backup transitions    :            0
```

show vrrp interface

```
user@host> show vrrp interface
Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                       :          205
    Advertisement received                   :            0
    Packets received                         :            0
    No group match received                  :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received          :            0
    Invalid VRRP TTL value received          :            0
    Invalid VRRP version received            :            0
    Invalid VRRP PDU type received           :            0
    Invalid VRRP authentication type received:            0
    Invalid VRRP IP count received           :            0
    Invalid VRRP checksum received           :            0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: ge2001:db8::12:1:1:1/120
  Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
```

```
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge2001:db8:0:1:12:1:1:99,
gec2001:db8::12:1:1:99
  Advertisement timer: 0.789s, Master router: ge2001:db8:0:1:12:1:1:1
  Virtual router uptime: 00:04:26, Master router uptime: 00:04:20
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent              :          205
    Advertisement received          :            0
  Group VRRP PDU error statistics
    Bad authentication type received:            0
    Bad password received           :            0
    Bad MD5 digest received         :            0
    Bad advertisement timer received:            0
    Bad VIP count received          :            0
    Bad VIPADDR received            :            0
  Group state transition statistics
    Idle to master transitions      :            0
    Idle to backup transitions      :            1
    Backup to master transitions    :            1
    Master to backup transitions    :            0
```

show vrrp summary

```
user@host> show vrrp summary
Interface     State       Group   VR state      Type   Address
ge-3/2/6.0    up              1   backup        lcl    10.57.0.2
                                                vip    10.57.0.100
```

show vrrp track detail

```
user@host> show vrrp track detail
Tracked interface: ae1.211
  State: up, Speed: 400m
  Incurred priority cost: 0
  Threshold   Priority cost   Active
  400m        10
  300m        60
```

```
  200m        110
  100m        160
  down        190
  Tracking VRRP interface: ae0.210, Group: 1
    VR State: master
    Current priority: 200, Configured priority: 200
    Priority hold-time: disabled,    Remaining-time: 50.351
```

show vrrp track summary

```
user@host> show vrrp track summary
Track if      State     Speed       VRRP if   Group   VR State   Current priority
ae1.211       up        400m        ae0.210       1      master     200
```

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

- Configuring VRRP for IPv6 (CLI Procedure)
- *Understanding VRRP on SRX Series Devices*
- *Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces*

show vrrp track

IN THIS SECTION

- Syntax | 27921
- Description | 27921
- Options | 27921
- Required Privilege Level | 27921
- Output Fields | 27922

Sample Output | 27924
Release Information | 27926

Syntax

```
show vrrp track
<all | interfaces | routes>
<detail | summary>
<logical-system logical-system-name>
```

Description

Display status information about Virtual Router Redundancy Protocol (VRRP) tracked routes and tracked interfaces.

Options

none
(Same as `summary`) Display summarized status information of tracked routes and tracked interfaces.

`all | interfaces | routes`
(Optional) These options display the following information:

- `all`—Output is the same as for the `show vrrp track` command.
- `interfaces`—Show summary of VRRP tracked interfaces.
- `routes`—Show summary of VRRP tracked routes

`detail | summary`
(Optional) Display detailed or summarized information.

`logical-system` *`logical-system-name`*
(Optional) Perform this operation on a particular logical system.

Required Privilege Level

view

Output Fields

Table 1794 on page 27922 lists the output fields for the `show vrrp track` command. Output fields are listed in the approximate order in which they appear.

Table 1794: show vrrp track Output Fields

Fields	Description	Level
`Tracked interface/ Track Int`	Name of the tracked interface.	`detail` or `summary`
`State`	Current operational state of the tracked interface: `up` or `down`.	`detail` or `summary`
`Speed`	Current operational speed, in bits per second, of the tracked interface.	`detail` or `summary`
`Incurred priority cost`	Operational priority cost incurred resulting from the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority cost.	`detail`
`VRRP Int/Tracking VRRP interface`	Name of the VRRP interface.	`detail` or `summary`
`Group`	VRRP group number.	`detail` or `summary`

Table 1794: show vrrp track Output Fields *(Continued)*

Fields	Description	Level
VR state	The state of the VRRP: • backup—The interface is acting as the backup router interface. • bringup—VRRP is just starting, and the physical device is not yet present. • idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • init—VRRP is initializing. • master—The interface is acting as the primary router interface. • transition—The interface is changing between being the backup and being the primary router. **NOTE**: When the show vrrp nsr command is used on the backup Routing Engine, it displays the current VRRP state on the primary Routing Engine, which is the future VRRP state for the backup Routing Engine. Do not use the show vrrp nsr command on the primary Routing Engine.	detail or summary
Current priority	Current operational priority for being the VRRP primary.	detail or summary
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority cost. Disabled indicates no minimum interval.	detail
Track route	IP address of route.	detail or summary
State	State of route. Possible values are unknown, up, and down.	detail or summary
Cost	Priority cost. When the route state is not up, the cost will be deducted from the configured priority of the VRRP session.	detail or summary

Table 1794: show vrrp track Output Fields *(Continued)*

Fields	Description	Level
Interface	Name of the logical interface (for example, ge-0/0/1.0) on which the corresponding VRRP session is configured.	detail or summary
Cfg	Configured priority.	detail or summary
Run	Current (or running) priority cost.	detail or summary

Sample Output

show vrrp track summary

```
user@host> show vrrp track summary

Track Int   State         Speed   VRRP Int   Group   VR State      Current prio
ge-0/0/2.0  up               1g   ge-0/0/1.0     1   master                  80
ge-0/0/8.0  up               1g   ge-0/0/1.0     1   master                  80

Track route         State       Cost    Interface  Group   Cfg   Run   VR State
44.44.44.0/24       unknown       10    ge-0/0/1.0     1   100    80   master
55.55.55.0/24       unknown       10    ge-0/0/1.0     1   100    80   master
```

show vrrp track detail

```
user@host> show vrrp track detail

Tracked interface: ge-0/0/2.0
  State: up, Speed: 1g
  Incurred priority cost: 0
  Tracking VRRP interface: ge-0/0/1.0, Group: 1
    VR State: master
    Current priority: 80, Configured priority: 100
    Priority hold-time: disabled
```

```
Tracked interface: ge-0/0/8.0
  State: up, Speed: 1g
  Incurred priority cost: 0
  Tracking VRRP interface: ge-0/0/1.0, Group: 1
    VR State: master
    Current priority: 80, Configured priority: 100
    Priority hold-time: disabled

Track route         State       Cost    Interface  Group   Cfg   Run   VR State
44.44.44.0/24       unknown       10    ge-0/0/1.0     1   100    80   master
55.55.55.0/24       unknown       10    ge-0/0/1.0     1   100    80   master
```

show vrrp track interfaces summary

```
user@host> show vrrp track interfaces summary
Track Int   State         Speed   VRRP Int   Group   VR State      Current prio
ge-0/0/2.0  up               1g   ge-0/0/1.0     1   master                  80
ge-0/0/8.0  up               1g   ge-0/0/1.0     1   master                  80
```

show vrrp track interfaces detail

```
user@host> show vrrp track interfaces detail
Tracked interface: ge-0/0/2.0
  State: up, Speed: 1g
  Incurred priority cost: 0
  Tracking VRRP interface: ge-0/0/1.0, Group: 1
    VR State: master
    Current priority: 80, Configured priority: 100
    Priority hold-time: disabled

Tracked interface: ge-0/0/8.0
  State: up, Speed: 1g
  Incurred priority cost: 0
  Tracking VRRP interface: ge-0/0/1.0, Group: 1
    VR State: master
    Current priority: 80, Configured priority: 100
    Priority hold-time: disabled
```

show vrrp track routes summary

```
user@host> show vrrp track routes summary

Track route         State       Cost    Interface  Group   Cfg   Run   VR State
44.44.44.0/24       unknown       10    ge-1/0/0.0     1   100    60   bringup
55.55.55.0/24       unknown       10    ge-1/0/0.0     1   100    60   bringup
```

show vrrp track routes detail

The output for `show vrrp track routes detail` is the same as that for `show vrrp track routes summary`.

Release Information

Command introduced before Junos OS Release 7.4.

`all` and `routes` options added in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Logical Interface to Be Tracked for a VRRP Group

Configuring a Route to Be Tracked for a VRRP Group

show vrrp | 27888

show wlan access-points

IN THIS SECTION

- Syntax | 27927
- Description | 27927
- Options | 27927
- Required Privilege Level | 27928
- Output Fields | 27928
- Sample Output | 27930

Release Information | 27934

Syntax

Syntax (SRX Series Firewalls)

```
show wlan
<detail>
<virtual-access-points>
<client-associations>
<neighbors>
<summary>
<radio>
<ha>
```

Description

Display information about wireless and virtual access point WLANs configured on the wireless LAN interface wl-*x*/0/0.

Options

detail	(Optional) Display the specified level of output.
virtual access points	Display Virtual access Points (VAPs) status and statistics.
client-associations	Client association number of the specified access point.
neighbors	List neighboring access point information including the MAC address, WPA, band, channel and SSID values.
summary	Display the access point configuration summary as output.
radio	Display access point radio information.
ha	Display wireless access point status in chassis cluster mode. In chassis cluster mode, only one wirelesss interface is active and the wireless client can choose the required

secure method to establish the wireless connection. After the wireless interface failover, wireless client retries the secure process to establish the new wireless connection.

In WIFI mPIM HA mode, there is one WAP mPIM card on each node. The WAP mPIM cards on both nodes have the same WAP configuration such as SSID, channel, and bandwidth. WAP0 is active on node0 and WAP1 is inactive on node1. Users are connected to WAP0 initially. When there is WAP0 failure, WAP1 is changed to active and the users are connected to WAP1 automatically.

The `show wlan access-points ha status` command is supported on primary routing engine. This command displays information from WAP mPIM card on both nodes.

Required Privilege Level

view

Output Fields

Table 1795 on page 27928 lists the output fields for the `show wlan access-points` command. Output fields are listed in the approximate order in which they appear.

Table 1795: `show wlan access-points` output fields

Field Name	Field Description	Level of Output
Access point	Name of the wireless access point.	All levels.
Type	Internal.	`detail`
Location	Location of the access point.	`detail`
Serial Number	Serial number of the Mini-PIM.	`detail`
Firmware Version	Firmware version of the Mini-PIM.	`detail`
Alternate Version	Alternate firmware version of the Mini-PIM.	`detail`

Table 1795: `show wlan access-points` **output fields** ***(Continued)***

Field Name	Field Description	Level of Output
Country	Country code.	`detail`
Access Interface	Name of the wireless LAN interface on the WI-Fi Mini-PIM.	`detail`
Packet Capture	Status of the traffic flow.	`detail`
Capture Interface	Interface captured on radio.	`detail`
Capture File	Capture file name.	`detail`
Capture Duration	Duration of capture in seconds.	`detail`
Capture File Size	File size of capture in kilobytes.	`detail`
MAC Address	MAC address of the Ethernet port.	`detail`
IPv4 Address	IPv4 address of the access point.	`detail`
Mode	Authentication mode on the radio.	`detail`
Channel	Channel bandwidth on the radio.	`detail` `neighbors`
VAP	Name of the virtual access point.	
SSID	Network name of the virtual access point.	`detail` `neighbors`
VLAN ID	VLAN ID associated with the access point.	`detail`

Table 1795: `show wlan access-points` output fields *(Continued)*

Field Name	Field Description	Level of Output
Traffic Statistics	Number and rate of bytes and packets received and transmitted on the wireless LAN interface. • Input bytes—Number of bytes received on the wireless LAN interface. • Output bytes—Number of bytes transmitted on the wireless LAN interface. • Input packets—Number of packets received on the wireless LAN interface. • Output packets—Number of packets transmitted on the wireless LAN interface.	`detail`
Client number	Client number on Radios 1 and 2.	`client-associations`
MAC Privacy	Client MAC address as per the virtual access.	`neighbors`
WPA	Status of the security authentication method.	`neighbors`
Band	Status of the band.	`neighbors`

Sample Output

show wlan access-points (SRX Series Firewalls)

```
user@router> show wlan access-points
Active access points information
Access-Point     Type   Interface    Radio-mode/Channel/Bandwidth
bj340d-ha        Int    wl-2/0/0     Off/Off/Off, Off/Off/Off
```

show wlan access-points (SRX Series Firewalls)

```
user@router> show wlan access-points e09-22-ha-ap
node0:
--------------------------------------------------------------------------

Active access point information

Access Point        : e09-22-ha-ap
Type                : Internal
Access Interface    : wl-2/0/0
Packet Capture      : Off
Radio1              : Mode: IEEE 802.11a/n/ac, Channel: 153, Bandwidth: 40
Radio2              : Mode: Off, Channel: Off, Bandwidth: Off

node1:
--------------------------------------------------------------------------

Active access point information

Access Point        : e09-22-ha-ap
Type                : Internal
Access Interface    : wl-7/0/0
Packet Capture      : Off
Radio1              : Mode: IEEE 802.11a/n/ac, Channel: 153, Bandwidth: 40
Radio2              : Mode: Off, Channel: Off, Bandwidth: Off
```

show wlan access-points detail (SRX Series Firewalls)

```
user@router> show wlan access-points bj340d-ha detail
Active access point detail information
Access Point          : bj340d-ha
Type                  : Internal
Location              : Default Location
Serial Number         : EV0519AF0022
Firmware Version      : v1.2.2
Alternate Version     : v1.1.8
Country               : US
Access Interface      : wl-2/0/0
System Time           : Thu Oct 31 12:11:40 UTC 2019
```

```
Packet Capture        : Off
Ethernet Port:
MAC Address           : 94:f7:ad:2c:7b:87

Radio1:
                        Status            : Off
                        MAC Address       : 94:f7:ad:2c:7b:89
                        Temperature       : 0
                        Mode              : Off
                        Channel           : Off
                        Bandwidth         : Off
                        Transmit Power    : Off

Radio2:
                        Status            : Off
                        MAC Address       : 94:f7:ad:2c:7b:88
                        Temperature       : 0
                        Mode              : Off
                        Channel           : Off
                        Bandwidth         : Off
                        Transmit Power    : Off
```

show wlan access-points virtual access-points (SRX Series Firewalls)

```
user@host> show wlan access-points bj340d-ha virtual-access-points
Virtual access points information
Access Point           : bj340d-ha
Radio1:
                    VAP0:
                    SSID                : juniper_ap_0
                    MAC Address         : 00:11:22:33:44:55
                    VLAN ID             : 5
                    Traffic Statistics  :
                    Input Bytes         : 37979930
                    Output Bytes        : 54231321
                    Input Packets       : 210737
                    Output Packets      : 298451

                    VAP1:
                    SSID                : juniper_ap_3
                    MAC Address         : 00:11:22:33:44:58
```

```
                    VLAN ID             : 15
                    Traffic Statistics  :
                    Input Bytes         : 83421130
                    Output Bytes        : 21221311
                    Input Packets       : 510837
                    Output Packets      : 238451

Radio2:
                    VAP0:
                    SSID                : juniper_ap_7
                    MAC Address         : 00:11:22:33:44:75
                    VLAN ID             : 5
                    Traffic Statistics  :
                    Input Bytes         : 67079335
                    Output Bytes        : 34932321
                    Input Packets       : 415737
                    Output Packets      : 308451

                    VAP1:
                    SSID                : juniper_ap_3
                    MAC Address         : 00:11:22:33:44:78
                    VLAN ID             : 49
                    Traffic Statistics  :
                    Input Bytes         : 37979930
                    Output Bytes        : 54231321
                    Input Packets       : 512837
                    Output Packets      : 901151
```

show wlan access-points client-associations summary (SRX Series Firewalls)

```
user@host> show wlan access-points bj340d-ha client associations summary
Access point client associations summary
Access point                        : bj340d-ha

Client number on radio 1 (5.0 GHz)  : 30
Client number on radio 2 (2.4 GHz)  : 20
Total client number on access point : 50
```

show wlan access-points neighbors (SRX Series Firewalls)

```
user@host> show wlan access-points bj340d-ha neighbors
Access point neighbors information
Access point         : bj340d-ha

MAC                  Privacy   WPA    Band    Channel     SSID
00:11:22:33:44:55    Off       Off    2.4     10          xyz-ap
00:12:23:45:56:67    On        On     5       100         abc-ap
```

show wlan access-points ha status (SRX Series Firewalls)

```
user@host> show wlan access-points ha status

Acess-point                                Interface            HA-status

e09-22-ha-ap                             wl-2/0/0                 Inactive
e09-22-ha-ap                             wl-7/0/0                 Active
```

Release Information

Command introduced in Junos OS Release 19.4R1.

`ha` option is introduced in the Junos OS Release 20.3R1 for SRX Series Firewalls.

RELATED DOCUMENTATION

Wi-Fi Mini-Physical Interface Module Overview

Configure Wi-Fi Mini-PIM

wlan | **31137**

shutdown (BPDU Block)

IN THIS SECTION

- Syntax | 27935
- Hierarchy Level | 27935
- Description | 27935
- Default | 27935
- Required Privilege Level | 27935
- Release Information | 27936

Syntax

```
shutdown;
```

Hierarchy Level

```
[edit ethernet-switching-options bpdu-block interface  (all | [interface-name])]
```

Description

Shut down all or specified interfaces to prevent spanning-tree protocol BPDUs (for STP, MSTP, RSTP, and VSTP) from entering the interfaces on which BPDU protection is configured.

Default

Not enabled

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Understanding BPDU Protection for STP, RSTP, and MSTP

Unblocking an Interface on non-ELS EX Series Switches That Receives BPDUs in Error (CLI Procedure)

shutdown (Protocols BGP)

IN THIS SECTION

- Syntax | 27936
- Hierarchy Level | 27936
- Description | 27937
- Options | 27937
- Required Privilege Level | 27937
- Release Information | 27937

Syntax

```
shutdown {
    notify-message notify-message;
}
```

Hierarchy Level

```
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
```

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address]
```

Description

Extended BGP administrative shutdown communication. This shutdown communication feature allows configuration, encode, and sending of additional text message along with BGP cease notification to the peer. This feature is useful during maintenance or unstable window.

Options

notify-message	Notification message

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

graceful-shutdown (Protocols BGP) | 7784

shutdown-threshold

IN THIS SECTION

- Syntax | 27938
- Hierarchy Level | 27938
- Description | 27938

- Default | 27938
- Options | 27938
- Required Privilege Level | 27939
- Release Information | 27939

Syntax

```
shutdown-threshold number;
```

Hierarchy Level

```
[edit vlans vlan-name dot1q-tunneling layer2-protocol-tunneling (all | protocol-name)]
```

Description

Specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the interface is disabled. Once an interface is disabled, you must explicitly reenable it using the `clear ethernet-switching layer2-protocol-tunneling error` command. Otherwise, the interface remains disabled.

The shutdown threshold value must be greater than or equal to the drop threshold value. If the shutdown threshold value is less than the drop threshold value, the drop threshold value has no effect.

You can specify a shutdown threshold value without specifying a drop threshold value.

Default

No shutdown threshold is specified.

Options

number—Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the interface is disabled.

- **Range:** 1 through 1000

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

drop-threshold | 5978

Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

Configuring Layer 2 Protocol Tunneling on EX Series Switches Without ELS Support

sib

IN THIS SECTION

- Syntax | 27939
- Hierarchy Level | 27940
- Description | 27940
- Options | 27940
- Required Privilege Level | 27940
- Release Information | 27940

Syntax

```
sib {
    minimum number;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the minimum number of SIBs on an M320 router. With this configuration, SIB absent alarms are generated only if the SIB count falls below the minimum specified.

Options

number—Minimum number of SIBs on the router.

- **Range:** 0 through 3

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Junos OS to Support Entry-Level Configuration on an M320 Router with a Minimum Number of SIBs and PIMs

pem (MX Series Devices, ACX7024 Devices) | **12399**

signal-bandwidth

IN THIS SECTION

- Syntax | 27941
- Hierarchy Level | 27941
- Description | 27941
- Options | 27941
- Required Privilege Level | 27941
- Release Information | 27942

Syntax

```
signal-bandwidth type;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name lsp-
attributes],
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
```

Description

Specify the bandwidth encoding of the signal used for path computation and admission control.

Options

type—Configure the type of bandwidth encoding used on the LSP. It can be any of the following values: `10gigether`, `ds1`, `ds3`, `e1`, `e3`, `ethernet`, `fastether`, `gigether`, `stm-1`, `stm-4`, `stm-16`, `stm-64`, `stm-256`, `sts-1`, `vt1-5`, or `vt2`.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Signal Bandwidth Type

signal-degrade

IN THIS SECTION

- Syntax | 27942
- Hierarchy Level | 27943
- Description | 27943
- Default | 27943
- Options | 27943
- Required Privilege Level | 27943
- Release Information | 27943

Syntax

```
signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
    interval value;
}
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify bit error rate (BER) signal degradation thresholds and time interval for raising and clearing alarms for optical transport network (OTN) links.

Default

If you omit the `signal-degrade` statement, the default threshold values are used.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

10-Gigabit Ethernet OTN Options Configuration Overview

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

signal-degrade-monitor-enable

IN THIS SECTION

- Syntax | 27944
- Hierarchy Level | 27944
- Description | 27944
- Default | 27945
- Options | 27945
- Required Privilege Level | 27945
- Release Information | 27945

Syntax

```
(signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

Description

Enable or disable pre-forward error correction (FEC) bit error rate (BER) monitoring.

With pre-FEC BER monitoring enabled, when the configured pre-FEC BER signal degrade threshold is reached, the PIC stops forwarding packets to the remote interface and raises an interface alarm. Ingress packets continue to be processed. If pre-FEC BER monitoring is used with MPLS fast reroute or another link protection method, then traffic is rerouted to a different interface.

You can also configure backward fast reroute to insert local pre-FEC BER status into transmitted OTN frames, notifying the remote interface of signal degradation. The remote interface can use the information to reroute traffic to a different interface. If you use pre-FEC BER monitoring together with backward fast reroute, then notification of signal degradation and rerouting of traffic occurs in less time

than that required through a Layer 3 protocol. To configure backward fast reroute, include the `backward-frr-enable` statement at the same hierarchy level.

> **NOTE**: When you configure pre-FEC BER signal degrade monitoring, we recommend that you configure both the `signal-degrade-monitor-enable` and `backward-frr-enable` statements.

You can also configure the pre-FEC BER thresholds that raise or clear a signal degrade alarm and the time interval for the thresholds. If the BER thresholds and interval are not configured, the default values are used. Include the `ber-threshold-signal-degrade` *value*, `ber-threshold-clear` *value*, and `interval` *value* statements at the [`edit interfaces` *interface-name* `otn-options signal-degrade`] hierarchy level to configure the BER thresholds and time interval. See *Understanding Pre-FEC BER Monitoring and BER Thresholds* for more information about pre-FEC BER monitoring and determining BER threshold settings.

Default

By default, pre-FEC BER signal degrade monitoring is disabled.

Options

`signal-degrade-monitor-enable`—Enable pre-FEC BER signal degrade monitoring.

`no-signal-degrade-monitor-enable`—Do not enable pre-FEC BER signal degrade monitoring.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding Pre-FEC BER Monitoring and BER Thresholds

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

signal-type

IN THIS SECTION

- Syntax | 27946
- Hierarchy Level | 27946
- Description | 27946
- Options | 27946
- Required Privilege Level | 27947
- Release Information | 27947

Syntax

```
signal-type (1hz | 5mhz | 10mhz | 2048mhz | e1 | t1);
```

Hierarchy Level

```
[edit chassis synchronization interfaces external]
[edit chassis synchronization interfaces (external-0/0 | external-1/0)]
```

Description

Configure the frequency for the provided grandmaster clock.

Options

1mhz Set the signal with a clock frequency of 1 MHz.

5mhz Set the signal with a clock frequency of 5 MHz.

10mhz Set the signal with a clock frequency of 10 MHz.

2048khz Set the signal as an E1 unframed 2048 kHz G.703 signal.

e1 Set the signal as an E1-coded 2048-kHz signal on a 120-ohm balanced line.

t1 Set the signal as a T1-coded 1.544-MHz signal on a 100-ohm balanced line.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

signal-type (BITS Interfaces)

IN THIS SECTION

- Syntax | 27947
- Hierarchy Level | 27948
- Description | 27948
- Options | 27948
- Required Privilege Level | 27948
- Release Information | 27948

Syntax

```
signal-type (2048KHZ | e1 | t1)
```

Hierarchy Level

```
[edit chassis synchronization interfaces bits]
```

Description

Configure the input and output signal type to the same value for the building-integrated timing supply (BITS) interface.

Options

2048khz Configure the BITS interface to output or accept as input the 2.048 MHz square wave.

e1 Configure the BITS interface to output or accept as input the E1 framed signal.

t1 Configure the BITS interface to output or accept as input the T1 framed signal.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

External Clock Synchronization Overview for ACX Series Routers

Configuring External Clock Synchronization for ACX Series Routers

Synchronous Ethernet Overview

show chassis synchronization | **18879**

e1-options (BITS Interfaces Signal Type) | **6176**

t1-options (BITS Interfaces Signal Type) | **28887**

signaling

IN THIS SECTION

- Syntax | 27949
- Hierarchy Level | 27949
- Description | 27949
- Required Privilege Level | 27950
- Release Information | 27950

Syntax

```
signaling;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family inet-mdt],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family inet-mvpn],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family inet-mdt],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family inet-mvpn],
[edit routing-instances routing-instance-name protocols bgp family inet-mdt],
[edit routing-instances routing-instance-name protocols bgp family inet-mvpn],
[edit routing-instances routing-instance-name protocols bgp group group-name family inet-mdt],
[edit routing-instances routing-instance-name protocols bgp group group-name family inet-mvpn]
```

Description

Enable signaling in BGP. For multicast distribution tree (MDT) subaddress family identifier (SAFI) NLRI signaling, configure signaling under the **inet-mdt** family. For multiprotocol BGP (MBGP) intra-AS NLRI signaling, configure signaling under the **inet-mvpn** family.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs

signature

IN THIS SECTION

- Syntax | 27950
- Hierarchy Level | 27951
- Description | 27951
- Options | 27951
- Required Privilege Level | 27951
- Release Information | 27951

Syntax

```
signature name {
    member name {
        context context;
        direction (any | client-to-server | server-to-client);
        pattern pattern;
    }
```

```
    port-range [ port-range ... ];
}
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type]
```

Description

Application signature for pattern matching. A unique application signature identifier. Must be a unique name with a maximum length of 63 characters.

You need to define an application signature to match the pattern by defining a unique application signature identifier, application signature member identifier, connection direction of the packets, and set the context to be matched. You also need to specify port range for TCP or UDP.

Options

member Member name for a custom application signature. Custom signatures can contain multiple members that define attributes for an application. (The supported member name range is m01 through m15.)

port-range Port range. This option is applicable for TCP-based or UDP-based applications only.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

Custom Application Signatures for Application Identification

signature (Application Identification)

IN THIS SECTION

- Syntax | 27952
- Hierarchy Level | 27952
- Description | 27953
- Required Privilege Level | 27953
- Release Information | 27953

Syntax

```
signature l4-l7-signature-name {
    chain-order
    member member-name {
        check-bytes  max-bytes-to-check;
        context context;
        pattern pattern;
        direction (any | client-to-server | server-to-client);
    }
    order order;
    order-priority (high | low);
    port-range {
        tcp [port-range];
        udp [port-range];
    }
protocol (http | ssl | tcp | udp);
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type]
```

Description

Configure an application signature for pattern matching.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Releases 19.3R2 and 19.4R1 on MX Series MX240, MX480 and MX960 using the MX-SPC3 services card.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

Signature Hash Algorithm (Security IKE)

IN THIS SECTION

- Syntax | 27954
- Hierarchy Level | 27954
- Description | 27954
- Options | 27954
- Required Privilege Level | 27954
- Release Information | 27954

Syntax

```
signature-hash-algorithm (sha1 | sha256 | sha384 | sha512);
```

Hierarchy Level

```
[edit security ike proposal proposal-name]
```

Description

Defines the digital signature hash algorithms that must match with the received signature hash algorithms in the hierarchical order for IKEv2. If your device doesn't find a match, negotiation with the IKE peer fails. This configuration ensures that the peers negotiate a specific set of signature hash algorithms. Note that you cannot use the configuration statement with IKEv1.

If not configured, the default value is all the supported hash algorithms.

Options

sha1 Select SHA-1 signature hash algorithm.

sha256 Select SHA 256-bit signature hash algorithm.

sha384 Select SHA 384-bit signature hash algorithm.

sha512 Select SHA 512-bit signature hash algorithm.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.4R1.

RELATED DOCUMENTATION

Signature Authentication in IKEv2

signature (Security IDP)

IN THIS SECTION

- Syntax | 27955
- Hierarchy Level | 27961
- Description | 27961
- Options | 27961
- Required Privilege Level | 27961
- Release Information | 27962

Syntax

```
signature {
    context context-name;
    direction (any | client-to-server | server-to-client);
    negate;
    pattern signature-pattern;
    pattern-pcre signature-pattern-pcre;
    protocol (Security IDP Signature Attack) {
        iicmp (Security IDP Signature Attack) {
            checksum-validate {
                match (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
            code {
                match (equal | greater-than | less-than | not-equal);
                value code-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
```

```
            }
            identification {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            sequence-number {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
            type {
                match (equal | greater-than | less-than | not-equal);
                value type-value;
            }
        }
        icmpv6 {
            checksum-validate {
                match (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
            code {
                match (equal | greater-than | less-than | not-equal);
                value code-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
            }
            identification {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            sequence-number {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
            type {
                match (equal | greater-than | less-than | not-equal);
                value type-value;
            }
        }
        ipv4 (Security IDP Signature Attack) {
            checksum-validate {
                match (equal | greater-than | less-than | not-equal);
```

```
                value checksum-value;
            }
            destination {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            identification (Security ICMP Headers) {
                match (equal | greater-than | less-than | not-equal);
                value identification-value;
            }
            ihl {
                match (equal | greater-than | less-than | not-equal);
                value ihl-value;
            }
            ip-flags {
                (df | no-df);
                (mf | no-mf);
                (rb | no-rb);
            }
            protocol {
                match (equal | greater-than | less-than | not-equal);
                value transport-layer-protocol-id;
            }
            source {
                match (equal | greater-than | less-than | not-equal);
                value ip-address-or-hostname;
            }
            tos {
                match (equal | greater-than | less-than | not-equal);
                value type-of-service-in-decimal;
            }
            total-length {
                match (equal | greater-than | less-than | not-equal);
                value total-length-of-ip-datagram;
            }
            ttl {
                match (equal | greater-than | less-than | not-equal);
                value time-to-live;
            }
        }
        ipv6 {
            destination {
                match (equal | greater-than | less-than | not-equal);
```

```
            value ip-address-or-hostname;
        }
        extension-header {
            destination-option {
                home-address {
                    match (equal | greater-than | less-than | not-equal);
                    value header-value;
                }
                    option-type {
                        match (equal | greater-than | less-than | not-equal);
                        value header-value;
                    }
            }
            routing-header {
                header-type {
                    match (equal | greater-than | less-than | not-equal);
                    value header-value;
                    }
                }
        }
        flow-label {
            match (equal | greater-than | less-than | not-equal);
            value flow-label-value;
        }
        hop-limit {
            match (equal | greater-than | less-than | not-equal);
            value hop-limit-value;
        }
        next-header {
            match (equal | greater-than | less-than | not-equal);
            value next-header-value;
        }
        payload-length {
            match (equal | greater-than | less-than | not-equal);
            value payload-length-value;
        }
        source {
            match (equal | greater-than | less-than | not-equal);
            value ip-address-or-hostname;
        }
        traffic-class {
            match (equal | greater-than | less-than | not-equal);
            value traffic-class-value;
```

```
            }
        tcp (Security IDP Signature Attack) {
            ack-number {
                match (equal | greater-than | less-than | not-equal);
                value acknowledgement-number;
            }
            checksum-validate {
                match (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value tcp-data-length;
            }
            destination-port {
                match (equal | greater-than | less-than | not-equal);
                value destination-port;
            }
            header-length {
                match (equal | greater-than | less-than | not-equal);
                value header-length;
            }
            mss {
                match (equal | greater-than | less-than | not-equal);
                value maximum-segment-size;
            }
            option {
                match (equal | greater-than | less-than | not-equal);
                value tcp-option;
            }
            reserved {
                match (equal | greater-than | less-than | not-equal);
                value reserved-value;
            }
            sequence-number {
                match (equal | greater-than | less-than | not-equal);
                value sequence-number;
            }
            source-port {
                match (equal | greater-than | less-than | not-equal);
                value source-port;
            }
            tcp-flags {
```

```
                (ack | no-ack);
                (fin | no-fin);
                (psh | no-psh);
                (r1 | no-r1);
                (r2 | no-r2);
                (rst | no-rst);
                (syn | no-syn);
                (urg | no-urg);
            }
            urgent-pointer {
                match (equal | greater-than | less-than | not-equal);
                value urgent-pointer;
            }
            window-scale {
                match (equal | greater-than | less-than | not-equal);
                value window-scale-factor;
            }
            window-size {
                match (equal | greater-than | less-than | not-equal);
                value window-size;
            }
        }
        udp (Security IDP Signature Attack) {
            checksum-validate {
                match (equal | greater-than | less-than | not-equal);
                value checksum-value;
            }
            data-length {
                match (equal | greater-than | less-than | not-equal);
                value data-length;
            }
            destination-port {
                match (equal | greater-than | less-than | not-equal);
                value destination-port;
            }
            source-port {
                match (equal | greater-than | less-than | not-equal);
                value source-port;
            }
        }
    }
    protocol-binding {
        application application-name;
```

```
        icmp;
        icmpv6;
        ip {
            protocol-number transport-layer-protocol-number;
        }
        ipv6 {
            protocol-number transport-layer-protocol-number;
        }
        rpc {
            program-number rpc-program-number;
        }
        tcp {
            minimum-port port-number <maximum-port port-number>;
        }
        udp {
            minimum-port port-number <maximum-port port-number>;
        }
    }
    regexp regular-expression;
    shellcode (all | intel | no-shellcode | sparc);
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type]
```

Description

IDP uses stateful signatures to detect attacks. Stateful signatures are more specific than regular signatures. With stateful signatures, IDP can look for the specific protocol or service used to perpetrate the attack.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

simple-filter

IN THIS SECTION

- Syntax | 27962
- Hierarchy Level | 27963
- Description | 27963
- Options | 27963
- Required Privilege Level | 27963
- Release Information | 27964

Syntax

```
simple-filter filter-name {
    term term-name {
        from {
            match-conditions;
        }
        then {
            forwarding-class class-name;
            loss-priority (high | low | medium);
        }
    }
}
```

Hierarchy Level

```
[edit firewall family inet],
[edit logical-systems logical-system-name firewall family inet]
```

Description

Configure simple filters.

Options

filter-name	Name that identifies the simple filter. The name must be a non-reserved string of not more than 64 characters. No special characters are restricted. To include spaces in the name, enclose them in quotation marks (“ ”).
from	Match packet fields to values. If the `from` option is not included, all packets are considered to match and the actions and action modifiers in the `then` statement are taken.
match-conditions	One or more conditions to use to make a match.
term *term-name*	Define a simple-filter term. The name that identifies the term can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose them in quotation marks (“ ”).
then	Actions to take on matching packets. If the `then` option is not included and a packet matches all the conditions in the `from` statement, the packet is accepted.

NOTE: Only `forwarding-class` and `loss-priority` are valid actions in a simple filter configuration.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Logical systems support introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

- simple-filter (Applying to an Interface) | 27964
- *Simple Filter Overview*
- *How Simple Filters Evaluate Packets*
- *Guidelines for Configuring Simple Filters*
- *Guidelines for Applying Simple Filters*

simple-filter (Applying to an Interface)

IN THIS SECTION

- Syntax | 27964
- Hierarchy Level | 27965
- Description | 27965
- Options | 27965
- Required Privilege Level | 27965
- Release Information | 27965

Syntax

```
simple-filter {
    input filter-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet]
```

Description

Apply a simple filter to an interface. You can apply simple filters to the family `inet` only, and only in the input direction.

Options

`input` *`filter-name`*—Name of one filter to evaluate when packets are received on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring Multifield Classifiers

filter (Applying to an Interface) | 7125

simple-filter (Interfaces)

IN THIS SECTION

Syntax | 27966

- Hierarchy Level | 27966
- Description | 27966
- Options | 27966
- Required Privilege Level | 27966
- Release Information | 27966

Syntax

```
simple-filter;
```

Hierarchy Level

```
[edit interfaces interfaces-name unit logical-unit-number family family-name]
```

Description

Apply a simple filter to an interface. You can apply simple filters on ingress interfaces only.

Options

input *filter-name*: Name of one filter to evaluate when packets are received on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

simple-mail-client-service

IN THIS SECTION

- Syntax | 27967
- Hierarchy Level | 27967
- Description | 27967
- Options | 27967
- Required Privilege Level | 27968
- Release Information | 27968

Syntax

```
simple-mail-client-service {
    command binary-file-path;
    disable;
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the SMTP client process.

Options

- `command` *`binary-file-path`*—Path to the binary process.

- `disable`—Disable the SMTP client process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

single-home (Junos Fusion)

IN THIS SECTION

- Syntax | 27968
- Hierarchy Level | 27968
- Description | 27969
- Options | 27969
- Required Privilege Level | 27969
- Release Information | 27969

Syntax

```
single-home {
    satellite [slot-id |slot-id-range | all];
}
```

Hierarchy Level

```
[edit chassis satellite-management]
```

Description

Specify that the links connecting the satellite device to the aggregation device are connected to the aggregation device only.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

Statement introdcued in Junos OS Release 16.1R1 for Junos Fusion Enterprise.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

single-rate

IN THIS SECTION

- Syntax | 27970
- Hierarchy Level | 27970
- Description | 27970
- Required Privilege Level | 27970
- Release Information | 27970

Syntax

```
single-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    excess-burst-size bytes;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name],
[edit firewall three-color-policer policer-name],
[edit logical-systems logical-system-name firewall three-color-policer policer-name]
```

Description

Configure a single-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).

Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the EBS are assigned medium-high loss priority (yellow). Packets that exceed the EBS are assigned high loss priority (red).

Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

Support at the [edit dynamic-profiles ... three-color-policer *name*] hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Three-Color Policer Configuration Overview
- color-aware | 4428
- color-blind | 4429
- two-rate | 30259

sip

IN THIS SECTION

- Syntax | 27971
- Hierarchy Level | 27972
- Description | 27972
- Default | 27972
- Required Privilege Level | 27972
- Release Information | 27972

Syntax

```
sip {
    video {
        dscp (alias | bits);
        forwarding-class class-name;
    }
    voice {
        dscp (alias | bits);
        forwarding-class class-name;
```

```
    }
}
```

Hierarchy Level

```
[edit services cos application-profile profile-name]
```

Description

Set the appropriate **dscp** and **forwarding-class** value for SIP traffic.

Default

By default, the system will not alter the DSCP or forwarding class for SIP traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring CoS Rules

sip (Application Profile)

IN THIS SECTION

- Syntax | 27973
- Hierarchy Level | 27973
- Description | 27973
- Default | 27974
- Required Privilege Level | 27974
- Release Information | 27974

Syntax

```
sip {
    video {
        dscp (alias | bits);
        forwarding-class class-name;
    }
    voice {
        dscp (alias | bits);
        forwarding-class class-name;
    }
}
```

Hierarchy Level

```
[edit services cos application-profile profile-name]
```

Description

Set the appropriate `dscp` and `forwarding-class` value for SIP traffic.

Default

By default, the system will not alter the DSCP or forwarding class for SIP traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

ftp (Services CoS) | 7628

sip-call-hold-timeout

IN THIS SECTION

- Syntax | 27974
- Hierarchy Level | 27975
- Description | 27975
- Options | 27975
- Required Privilege Level | 27975
- Release Information | 27975

Syntax

```
sip-call-hold-timeout seconds;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Timeout period for SIP calls placed on hold, in seconds.

Options

seconds—Length of time the application holds a SIP call open before it times out.

- **Default:** 7200 seconds
- **Range:** 0 through 36,000 seconds (10 hours)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

ALG Descriptions

Configuring SIP

Examples: Configuring Application Protocols

Verifying the Output of ALG Sessions

sip-password

IN THIS SECTION

- Syntax | 27976
- Hierarchy Level | 27976
- Description | 27976
- Options | 27976
- Required Privilege Level | 27976
- Release Information | 27977

Syntax

```
sip-password simple-ip-password;
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options profiles profile-name]
```

Description

Configure the password provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.

Options

simple-ip-password—Password.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

sip (Security)

IN THIS SECTION

- Syntax | 27977
- Hierarchy Level | 27978
- Description | 27978
- Options | 27979
- Required Privilege Level | 27979
- Release Information | 27979

Syntax

```
sip {
    application-screen {
        protect {
            deny (all | varname);
            timeout varseconds;
        }
        unknown-message {
            permit-nat-applied;
            permit-routed;
        }
    }
    bulk-call-mode;
    c-timeout varminutes;
```

```
    disable;
    dscp-rewrite {
        code-point string;
    }
    varname distribute-tofpc {
        pic varname;
        varname;
    }
    dscp-rewrite {
        code-point varcode-point;
    }
    enable-call-distribution;
    hide-via-headers disable;
    inactive-media-timeout varseconds;
    maximum-call-duration varminutes;
    retain-hold-resource;
    t1-interval varmilliseconds;
    t4-interval varseconds;
    traceoptions {
        flag {
            all {
                <detail | extensive | terse>
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit tenants name security alg]
```

Description

Configure sip to specify the Session Initiation Protocol (SIP) ALG on the device. Session Initiation Protocol (SIP) is a signaling protocol for initiating, modifying, and terminating multimedia sessions over the internet. By default, the SIP ALG is disabled for SRX5600 and SRX5800 devices. Use the CLI command `delete security alg alg-name disable` to enable the SIP ALG. IPv6 is supported on the SIP ALG

along with Network Address Translation Protocol Translation (NAT-PT) mode and NAT64 address translation. The type of the destination-ip-address is changed from IPv4 address to IP prefix to support all kinds of IP addresses, and correspondingly a prefix is supported to allow multiple IP addresses.

Options

`disable`—Disable the SIP ALG. Use the `set security alg` *`alg-name`* `disable` CLI command to disable the SIP ALG.

`distribution-ip`—Configure IPv4 or IPv6 distribution IP addresses. You can configure upto 20 IP addresses.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

The `distribution-ip` option introduced in Junos OS Release 18.1R1 for SRX5400, SRX5600, and SRX5800 devices.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding the SIP ALG

sip-server

IN THIS SECTION

- Syntax | 27980
- Hierarchy Level | 27980

- Description | 27980
- Options | 27980
- Required Privilege Level | 27980
- Release Information | 27981

Syntax

```
sip-server [address | name];
```

Hierarchy Level

```
[edit system services dhcp],
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]
```

Description

Configure Session Initiation Protocol (SIP) server addresses or names for DHCP servers.

Options

address—IPv4 address of the SIP server. To configure multiple SIP servers, include multiple ***address*** options. This address must be accessible by all clients served within a specified range of addresses (based on an address pool or static binding).

name—Fully qualified domain name of the SIP server. To configure multiple SIP servers, include multiple ***name*** options. This domain name must be accessible by all clients served within a specified range of addresses (based on an address pool or static binding).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configure DHCP SIP Server

Configuring a DHCP Server on Switches

sip (Services CoS Next Gen Services)

IN THIS SECTION

- Syntax | 27981
- Hierarchy Level | 27982
- Description | 27982
- Options | 27982
- Required Privilege Level | 27982
- Release Information | 27982

Syntax

```
sip {
    data {
        dscp (alias | bits);
        forwarding-class class-name;
    }
}
```

Hierarchy Level

```
[edit services cos application-profile profile-name]
```

Description

Configure CoS actions for SIP traffic in an application profile. The application profile can then be used in CoS rule actions.

Options

dscp (*alias* | *bits*) Either a code point alias or a DSCP bit value to apply to the SIP packets.

forwarding-class *class-name* Forwarding class name to apply to the SIP packets. The choices are:

- assured-forwarding
- best-effort
- expedited-forwarding
- network-control

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Class of Service for Services PICs (Next Gen Services)

sip-user-id

IN THIS SECTION

- Syntax | 27983
- Hierarchy Level | 27983
- Description | 27983
- Options | 27983
- Required Privilege Level | 27983
- Release Information | 27984

Syntax

```
sip-user-id simple-ip-user-id;
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options profiles profile-name]
```

Description

Configure the username provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.

Options

simple-ip-user-id—Username.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

site-identifier (Layer 2 Circuits)

IN THIS SECTION

- Syntax | 27984
- Hierarchy Level | 27984
- Description | 27984
- Options | 27985
- Required Privilege Level | 27985
- Release Information | 27985

Syntax

```
site-identifier identifier;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn site site-name],
[edit routing-instances routing-instance-name protocols l2vpn site site-name]
```

Description

Specify the numerical identifier for the local Layer 2 VPN site.

Options

identifier—The numerical identifier for the Layer 2 VPN site, which can be any number from 1 through 65,534.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Site

Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)

site-identifier (VPLS)

IN THIS SECTION

- Syntax | 27986
- Hierarchy Level | 27986
- Description | 27986
- Options | 27986
- Required Privilege Level | 27986
- Release Information | 27986

Syntax

```
site-identifier identifier;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
site site-name],
[edit routing-instances routing-instance-name protocols vpls site site-name]
```

Description

Specify the numerical identifier for the local VPLS site.

Options

identifier—Specify the numerical identifier for the local VPLS site. The identifier must be an unsigned 16-bit number greater than zero.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the VPLS Site Name and Site Identifier

site (Layer 2 Circuits)

IN THIS SECTION

- Syntax | 27987
- Hierarchy Level | 27987
- Description | 27988
- Options | 27988
- Required Privilege Level | 27988
- Release Information | 27988

Syntax

```
site site-name {
    hot-standby;
    site-identifier identifier;
    site-preference preference-value {
        backup;
        primary;
    }
    interface interface-name {
        description text;
        remote-site-id remote-site-ID;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn],
[edit routing-instances routing-instance-name protocols l2vpn]
```

Description

Specify the site name, site identifier, and interfaces connecting to the site. Allows you to configure a remote site ID for remote sites.

Options

`hot-standby`—Turn on the protector behavior for the site. This keeps backup pseudowire in continuous standby mode and ready for traffic forwarding.

`site-identifier` *`identifier`*—Numerical identifier for the site used as a default reference for the remote site ID.

`site-name`—Name of the site.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`hot-standby` option introduced in Junos OS Release 14.2 for MX Series routers.

RELATED DOCUMENTATION

Configuring the Site

Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)

site-preference

IN THIS SECTION

- Syntax | 27989
- Hierarchy Level | 27989
- Description | 27989
- Options | 27990
- Required Privilege Level | 27990
- Release Information | 27990

Syntax

```
site-preference preference-value {
    backup;
    primary;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn site site-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
site site-name],
[edit routing-instances routing-instance-name protocols l2vpn site site-name],
[edit routing-instances routing-instance-name protocols vpls site site-name]
```

Description

Specify the preference value advertised for a particular Layer 2 VPN or VPLS site. The site preference value is encoded in the BGP local preference attribute. When a PE router receives multiple advertisements with the same VE identifier, the advertisement with the highest local preference value is preferred. You can use this statement to enable multihoming for Layer 2 VPNs and VPLS.

Options

preference-value—Specify the preference value advertised for a Layer 2 VPN or VPLS site.

- **Range:** 1 through 65,535

backup—Set the preference value to 1.

primary—Set the preference value to 65,535.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for Layer 2 VPNs in Junos OS Release 9.1.

RELATED DOCUMENTATION

> *Configuring a Site Preference and Layer 2 VPN Multihoming*
>
> *Configuring the VPLS Site Preference*

site-range

IN THIS SECTION

- Syntax | 27991
- Hierarchy Level | 27991
- Description | 27991
- Options | 27991
- Required Privilege Level | 27991
- Release Information | 27991

Syntax

```
site-range number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Specify an upper limit on the maximum site identifier that can be accepted to allow a pseudowire to be brought up. Pseudowires cannot be established to sites with site identifiers greater than the configured site range. If you issue the `show vpls connections` command, such sites are displayed as OR (out of range).

Options

number Maximum number of site identifiers. We recommend using the default value.

- **Range:** 1 through 65,534
- **Default:** 65,534

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Site Range

site-reputation-action

IN THIS SECTION

- Syntax | 27992
- Hierarchy Level | 27992
- Description | 27992
- Options | 27993
- Required Privilege Level | 27993
- Release Information | 27993

Syntax

```
site-reputation-action {
    harmful   (block | log-and-permit | permit | quarantine);
    fairly-safe   (block | log-and-permit | permit | quarantine);
    moderately-safe  (block | log-and-permit | permit | quarantine);
    suspicious  (block | log-and-permit | permit | quarantine);
    very-safe   (block | log-and-permit | permit | quarantine);
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name category
category-name]
```

Description

Specify the action to be taken depending on the site reputation returned for all types of URLs whether it is categorized or uncategorized.

> **NOTE**: Starting with Junos OS Release 17.4R1, the reputation base scores are configurable. Users can apply global reputation values, provided by the Websense ThreatSeeker Cloud (TSC). For the non-category URLs, the global reputation value is used to perform filtering,

Options

fairly-safe Permit, log-and-permit, block, or quarantine a request if a site-reputation of 70 through 79 is returned.

harmful Permit, log-and-permit, block, or quarantine a request if a site-reputation of zero through 59 is returned.

moderately-safe Permit, log-and-permit, block, or quarantine a request if a site-reputation of 80 through 89 is returned.

suspicious Permit, log-and-permit, block, or quarantine a request if a site-reputation of 60 through 69 is returned.

very-safe Permit, log-and-permit, block, or quarantine a request if a site-reputation of 90 through 100 is returned.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

site (VPLS Multihoming for FEC 128)

IN THIS SECTION

- Syntax | 27994
- Hierarchy Level | 27994
- Description | 27995
- Options | 27995
- Required Privilege Level | 27995
- Release Information | 27995

Syntax

```
site site-name {
    mac-pinning;
    active-interface (any | primary interface-name);
    best-site;
    interface interface-name {
        interface-mac-limit limit;
    }
    mesh-group mesh-group-name;
    multi-homing;
    site-identifier identifier;
    site-preference preference-value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Specify the site name and site identifier for a site. Allows you to configure a remote site ID for remote sites.

Options

site-name—Name of the site.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

site (VPLS Multihoming for FEC 129)

IN THIS SECTION

- Syntax | 27996
- Hierarchy Level | 27996
- Description | 27996
- Options | 27996
- Required Privilege Level | 27996
- Release Information | 27997

Syntax

```
site site-name {
    active-interface interface-name {
        any;
        primary interface-name;
    }
    identifier identifier;
    interface interface-name {
        preference preference-value;
    }
    peer-active;
    preference (preference-value | backup | primary);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls multi-homing],
[edit routing-instances instance-name protocols vpls multi-homing]
```

Description

For VPLS autodiscovery (FEC 129), specify the parameters for a VPLS site that is multihomed to two or more provider edge (PE) routers.

> **NOTE**: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Options

site-name—Name of the site.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Configuring VPLS Multihoming (FEC 129)

size

IN THIS SECTION

- Syntax | 27997
- Hierarchy Level | 27997
- Description | 27998
- Options | 27998
- Required Privilege Level | 27998
- Release Information | 27998

Syntax

```
size bytes;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family (inet |inet6 |mpls) output file],
[edit forwarding-options sampling traceoptions file]
```

Description

Specify the maximum size of each file containing sample or log data. The file size is limited by the number of files to be created and the available hard disk space.

When a traffic sampling file named `sampling-file` reaches the maximum size, it is renamed `sampling-file.0`. When the `sampling-file` again reaches its maximum size, `sampling-file.0` is renamed `sampling-file.1` and `sampling-file` is renamed `sampling-file.0`. This renaming scheme continues until the maximum number of traffic sampling files is reached. Then the oldest traffic sampling file is overwritten.

Options

bytes—Maximum size of each traffic sampling file or trace log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

- **Syntax:** *x*`k` to specify KB, *x*`m` to specify MB, or *x*`g` to specify GB
- **Range:** 10 KB through the maximum file size supported on your router
- **Default:** 1 MB for sampling data; 128 KB for log information

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

Configuring Traffic Sampling on MX, M and T Series Routers

size (VNF)

IN THIS SECTION

- Syntax | 27999
- Hierarchy Level | 27999
- Description | 27999
- Required Privilege Level | 27999
- Release Information | 27999

Syntax

```
size size;
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Size of the memory in kilobytes.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D40.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

size (Interface)

IN THIS SECTION

- Syntax | 28000
- Hierarchy Level | 28000
- Description | 28000
- Options | 28001
- Required Privilege Level | 28001
- Release Information | 28001

Syntax

```
size bytes;
```

Hierarchy Level

```
[edit accounting-options file filename]
```

Description

Specify attributes of an accounting-data log file.

Options

bytes—Maximum size of each log file, in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). When a log file (for example, `profilelog`) reaches its maximum size, it is renamed `profilelog.0`, then `profilelog.1`, and so on, until the maximum number of log files is reached. Then the oldest log file is overwritten. If you do not specify a size, the file is closed, archived, and renamed when the time specified for the transfer interval is exceeded.

- **Syntax:** *x* to specify bytes, *x*k to specify KB, *x*m to specify MB, *x*g to specify GB
- **Range:** 256 KB through 1 GB

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Accounting Options, Source Class Usage and Destination Class Usage Options

size (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | 28002
- Hierarchy Level | 28002
- Description | 28002
- Options | 28002
- Required Privilege Level | 28002
- Release Information | 28002

Syntax

```
size size;
```

Hierarchy Level

```
[edit services rtlog traceoptions file]
```

Description

Maximum trace file size

Options

size Maximum trace file size

- **Default:** 128k
- **Range:** through

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

size (Sampling and Traceoptions)

IN THIS SECTION

- Syntax | 28003
- Hierarchy Level | 28003
- Description | 28003
- Options | 28004
- Required Privilege Level | 28004
- Release Information | 28004

Syntax

```
size bytes;
```

Hierarchy Level

```
[edit forwarding-options helperstraceoptions file],
[edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family family-name output file],
[edit forwarding-options sampling traceoptions file]
```

Description

Specify the maximum size of each file containing sample or log data. The file size is limited by the number of files to be created and the available hard disk space.

When a traffic sampling file named **sampling-file** reaches the maximum size, it is renamed **sampling-file.0**. When the **sampling-file** file again reaches its maximum size, **sampling-file.0** is renamed **sampling-file.1** and **sampling-file** is renamed **sampling-file.0**. This renaming scheme continues until the maximum number of traffic sampling files is reached. Then the oldest traffic sampling file is overwritten.

Options

bytes—Maximum size of each traffic sampling file or trace log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your router
- **Default:** 1 MB for sampling data; 128 KB for log information

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in a File

Tracing Traffic-Sampling Operations

size (Security Web Filtering Cache)

IN THIS SECTION

- Syntax | 28005
- Hierarchy Level | 28005
- Description | 28005
- Options | 28005
- Required Privilege Level | 28005
- Release Information | 28005

Syntax

```
size value;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated cache]
[edit security utm feature-profile web-filtering juniper-enhanced cache]
```

Description

Set the cache size parameters for Web filtering.

Options

- **Range:** 0 through 4096 kilobytes.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The surf-control feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

size (Services)

IN THIS SECTION

- Syntax | 28006
- Hierarchy Level | 28006
- Description | 28006
- Options | 28006
- Required Privilege Level | 28006
- Release Information | 28007

Syntax

```
size size;
```

Hierarchy Level

```
[edit services ssl traceoptions file file-name]
```

Description

Specify the maximum trace file size. This statement is supported on the SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall.

Options

`size`—Specify the maximum trace file size.

Range: 10,240 to 1,073,741,824.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

Firewall User Authentication Overview

size (System)

IN THIS SECTION

- Syntax | 28007
- Hierarchy Level | 28007
- Description | 28008
- Options | 28008
- Required Privilege Level | 28008
- Release Information | 28008

Syntax

```
size size;
```

Hierarchy Level

```
[edit system syslog archive],
[edit system syslog file filename archive]
```

Description

Configure the maximum amount of data that the Junos OS logging utility writes to a log file ***logfile*** before archiving it (closing it, compressing it, and changing its name to ***logfile*.0.gz**). The utility then opens and writes to a new file called ***logfile***. For information about the number of archive files that the utility creates in this way, see "files" on page 7108.

Options

size—Maximum size of each system log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

- **Syntax:** *x*k to specify the number of kilobytes, *x*m for the number of megabytes, or *x*g for the number of gigabytes
- **Range:** 64 KB through 1 GB
- **Default:** 1 MB for MX Series routers the QFX Series, and the OCX Series

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Specify Log File Size, Number, and Archiving Properties

System Log Explorer

files (syslog) | **7108**

skew-timer-disable

IN THIS SECTION

- Syntax | 28009
- Hierarchy Level | 28009
- Description | 28009
- Default | 28009
- Required Privilege Level | 28010
- Release Information | 28010

Syntax

```
skew-timer-disable;
```

Hierarchy Level

```
[edit protocols vrrp]
```

Description

Disable the skew timer, thereby reducing the time required to transition from the backup state to the primary state.

NOTE: The `skew-timer-disable` statement is used when there is only one primary router and one backup router in the network.

Default

By default, the skew timer is enabled for all the VRRP groups.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Improving the Convergence Time for VRRP

Configuring VRRP to Improve Convergence Time

skip-arp-iteration (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 28010
- Junos OS Hierarchy Level | 28011
- Junos OS Evolved Hierarchy Level | 28011
- Description | 28011
- Required Privilege Level | 28011
- Release Information | 28011

Syntax

```
skip-arp-iteration;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Disable the Address Resolution Protocol (ARP) test iteration for IPv4 or inet services during a benchmarking test. This parameter is valid only for an `inet` family. An ARP test iteration is a 3-second iteration that is run for all `inet` tests. The results of this iteration are disregarded in the test result calculations. The ARP test iteration is performed by sending test frames for 3 seconds to ensure that all devices on the path to the destination add ARP entries in the cache of the corresponding devices. This parameter is not applicable for CCC and bridge families.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

sla-iterator-profile

IN THIS SECTION

- Syntax | 28012
- Hierarchy Level | 28012
- Description | 28012
- Options | 28013
- Required Privilege Level | 28013
- Release Information | 28013

Syntax

```
sla-iterator-profile profile-name {
    data-tlv-size size;
    iteration-count count-value;
    priority priority-value;
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-
association ma-name mep mep-id remote-mep remote-mep-id]
```

Description

Configure a remote MEP with an iterator profile and specify the options.

Y.1731 performance monitoring (PM) over Aggregated Ethernet Interfaces is not supported on EX4300 switches.

Options

profile-name—Name of the iterator profile configured for a remote MEP. For more information about configuring a remote MEP with an iterator profile, see *Configuring a Remote MEP with an Iterator Profile*.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring an Iterator Profile

Configuring a Remote MEP with an Iterator Profile

Verifying the Configuration of an Iterator Profile

Managing Iterator Statistics

sla-iterator-profiles | **28013**

sla-iterator-profiles

IN THIS SECTION

- Syntax | **28014**
- Hierarchy Level | **28014**
- Description | **28014**
- Options | **28014**
- Required Privilege Level | **28015**

Release Information | 28015

Syntax

```
sla-iterator-profiles {
    profile-name {
        avg-fd-twoway-threshold;
            avg-ifdv-twoway-threshold;
            avg-flr-forward-threshold;
            avg-flr-backward-threshold;
            calculation-weight {
            delay delay-weight;
            delay-variation delay-variation-weight;
        }
        cycle-time milliseconds;
        flap-trap-monitor seconds
        iteration-period iteration-period-value;
        measurement-type (loss | statistical-frame-loss | two-way-delay);
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
```

Description

Configure an iterator application and specify the iterator profile options.

Options

profile-name—Name of the iterator profile. For more information about configuring the iterator profile, see *Configuring an Iterator Profile*.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring an Iterator Profile

Configuring a Remote MEP with an Iterator Profile

Verifying the Configuration of an Iterator Profile

Managing Iterator Statistics

sla-iterator-profiles (OAM CFM)

IN THIS SECTION

- Syntax | 28015
- Hierarchy Level | 28016
- Description | 28016
- Options | 28016
- Required Privilege Level | 28016
- Release Information | 28016

Syntax

```
sla-iterator-profiles {
    profile-name {
        calculation-weight {
```

```
            delay delay-value;
            delay-variation delay-variation-value;
        }
        cycle-time cycle-time-value;
        iteration-period iteration-period-value;
        measurement-type two-way-delay;
        passive;
    }
}
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
```

Description

Configure an iterator application and specify the iterator profile options.

Options

profile-name—Name of the iterator profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configure an Iterator Profile on a Switch (CLI Procedure)

sla-options

IN THIS SECTION

- Syntax | 28017
- Hierarchy Level | 28017
- Description | 28017
- Options | 28017
- Required Privilege Level | 28018
- Release Information | 28018

Syntax

```
sla-options {
    dynamic-load-balancer;
    log (disabled | syslog);
    max-passive-probe-limit {
        interval refresh-interval;
        number-of-probes probes;
    }
}
```

Hierarchy Level

[edit security advance-policy-based-routing]

Description

Specify global SLA options for Application Quality of Experience (AppQoE).

Options

dynamic-load-balancer Enable the session-based load balancing functionality. When you enable session-based load balancing option, the application traffic distribution depends on the `weight` value that

you configure at the [`edit security advance-policy-based-routing interface interface-name`] hierarchy level.

Load balancing is disabled by default.

log

Choose the logging mechanism for AppQoE. The log type configuration by default uses `syslog` to support application-level logging. You need to configure the log type as `disabled` to turn off the AppQoE logging.

disabled Disable logging.

syslog Enable logging.

- **Default:** syslog

max-passive-probe-limit

Set maximum passive probe limits.

interval Interval within which to send passive probes (in ms).

- **Default:** 500
- **Range:** 500 through 1000

number-of-probes Number of passive probes to be sent.

- **Default:** 5000
- **Range:** 10 through 5000

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

`dynamic-load-balancer` option added in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

interface (advance-policy-based-routing) | 8833

Application Quality of Experience

Advanced Policy-Based Routing

sla-rule

IN THIS SECTION

- Syntax | 28019
- Syntax (Prior Junos OS Release 21.2R1) | 28020
- Hierarchy Level | 28021
- Description | 28021
- Options | 28021
- Required Privilege Level | 28027
- Release Information | 28027

Syntax

```
sla-rule sla-rule-name {
    active-probe-params {
        probe-params-name;
    }
    link-preferences {
            preferred-tag preferred-tag;
            affinity strict;
    }
    metrics-profile {
        metric-profile-name;
    }
    passive-probe-params {
        sampling-percentage {
            percentage;
        }
```

```
        sampling-period {
            period;
        }
        type {
            book-ended;
        }
    switch-idle-time {
        period;
    }
    type {
        saas;
    }
    violation-count {
        count;
    }
}
```

Syntax (Prior Junos OS Release 21.2R1)

```
sla-rule sla-rule-name {
    active-probe-params {
        probe-params-name;
    }
    link-type-affinity strict;
    metrics-profile {
        metric-profile-name;
    }
    passive-probe-params {
        sampling-percentage {
            percentage;
        }
        sampling-period {
            period;
        }
        type {
            book-ended;
        }
    preferred-link-type (Any | IP | MPLS);
    sla-export-factor {
        value;
    }
```

```
    switch-idle-time {
        period;
    }
    type {
        saas;
    }
    violation-count {
        count;
    }
}
```

Hierarchy Level

[edit security advance-policy-based-routing]

Description

Configure an SLA rule.

An SLA rule includes all information required to measure the SLA and to identify whether any SLA violation has occurred or not. It contains the complete probe profiles, time interval which the profiles need to be sent, preferred SLA configuration, and so on.

When you configure an APBR rule, you must associate the corresponding SLA rule for the application.

The presence of SLA rule in the APBR configuration triggers the AppQoE functionality; If there are no SLA profiles available, APBR operates without AppQoE.

Options

active-probe-params ***probe-params-name*** Name of the active probe parameter. Associate the active probe parameter with the SLA rule.

affinity Define link type affinity. For strict affinity, AppQoE ensures that the selected path is always from the preferred tag. For the default affinity (loose), if there is no SLA meeting link in the preferred tag, AppQoE selects a link outside the preferred tag that meets the SLA requirements.

- **Default:** loose

If the link-type affinity is not configured as strict, then an overlay link with the highest SLA score is selected. If there are no links meeting the SLA requirements and belonging to the preferred link type available, then AppQoE selects a link outside of the preferred link type that meets the SLA requirements. However, path selection mechanism continues to check for a preferred link meeting the SLA requirements. Once the preferred link type meeting the SLA requirement is available, then the application traffic is switched back to that link.

link-type-affinity (strict)

(optional) Configure the link-type affinity as strict for the preferred link type. For strict affinity, AppQoE ensures that the path selected is always of the preferred link type. When the default affinity (loose) is configured and if there are no SLA meeting links belonging to the preferred link type available, then AppQoE selects a link outside the preferred link type that meets the SLA requirements.

Following table provides details on AppQoE pat selection mechanism based on affinity type.

Table 1796: Enhanced AppQoE Path Selection

Affinity Type	SLA Metrics Type	Link Selection Preference		
		If there is a preferred tag	If there are multiple preferred tags	If there is no preferred tag
Strict	With SLA metrics	AppQoE selects the links with the preferred tag.	AppQoE selects the link with highest path score. If there are multiple links with same path score, then the link selection happens as per the priority.	To configure affinity as strict, the user need to configure a preferred tag.
Strict	Without SLA metrics	AppQoE selects the links with the preferred tag.	AppQoE selects the link based on the priority.	To configure affinity as strict, the user need to configure a preferred tag.

Table 1796: Enhanced AppQoE Path Selection *(Continued)*

Affinity Type	SLA Metrics Type	Link Selection Preference		
		If there is a preferred tag	If there are multiple preferred tags	If there is no preferred tag
Loose	With SLA metrics	AppQoE selects the links with the preferred tag.	AppQoE selects the link with highest path score. If there are multiple links with same path score, then the link selection happens as per the priority.	AppQoE selects the link outside the preferred tag with best path score. The link selection switches back to preferred link when a link with a preferred tag is available. If there are multiple links with same the preferred tag after switchback, AppQoE selects the preferred link with highest path score. If there are multiple links with same path score, then the link selection happens as per the priority.

Table 1796: Enhanced AppQoE Path Selection *(Continued)*

Affinity Type	SLA Metrics Type	Link Selection Preference		
		If there is a preferred tag	If there are multiple preferred tags	If there is no preferred tag
Loose	Without SLA metrics	AppQoE selects the links with the preferred tag.	AppQoE selects the link based on the priority.	AppQoE selects non preferred link with highest priority. If there are multiple non preferred links with similar priority, then the link selection happens randomly.

Prior to Junos OS Release 21.2R1, during traffic flow, switching from a link with a lower priority value that meets the SLA requirements to a link with a higher priority is not supported when the link with higher priority meets the SLA requirements.

When a currently selected link and the highest priority link are part of a indirect next-hop table (belonging to a different routing instance), the application traffic does not switch to the highest priority link. It switches to the highest priority link only when the SLA manager is triggered. The SLA manager is triggered for the following reasons:

- When no path for application traffic is selected,
- SLA is violated partially
- A non-preferred link is available as best path
- A non-highest priority link is available as the best path
- Change in next-hop ID
- Configuration changes
- Current best path link is down
- SLA violation is detected

metrics-profile *profile-name*

Metric profile name. The SLA rule contains metric profiles that provide the acceptable threshold. If the violation goes beyond the threshold, an alternate path is identified and then traffic is rerouted.

passive-probe-params

Passive probe parameter name. Passive probes are installed on links within the network, and they monitor all the traffic that flows through those links.. This option is not supported for SaaS applications.

preferred-link-type (IP | MPLS | Any)

Select an MPLS or Internet link as the preferred path. If you do not select `IP` or `MPLS`, the preferred link type `Any` is selected when the link-type affinity is configured as loose (default link type affinity). Configuring the link type as `Any` when the link-type affinity is configured as `strict` is not supported.

- **Default:** Any

preferred-tag

Define a custom tag such as ISP1, ISP2, WAN1, IP, MPLS, and so on as per the requirement. The security device does not support the preferred tag as `Any` when you configure affinity as `strict`.

The preferred tag is case sensitive. For example, the security device considers the prefered tag `IP` and `ip` as different tags.

- **Default:** `any`

Custom Link Tag

SLA link preference for security device interfaces support custom link tags. If there are multiple links that meet SLA, you can set link preference by configuring the `preferred-tag` and `affinity` options at the [`edit security advance-policy-based-routing sla-rule` *`sla-rule-name`* `link-preferences`] hierarchy level.

Application Traffic Switch to the Higher Priority Link of the Preferred Tags

During the traffic flow, AppQoE allows the application traffic to switch from a lower priority link to a higher priority link that meets SLA requirements.

Non-SLA metrics Based Deployments

We support AppQoE for application-based traffic or non-application based traffic even when the SLA metrics are not available.

Overlay Interface Attribute Preference

In case of conflicting attributes within an overlay or between an overlay and corresponding underlay interfaces, the highest level overlay link attributes take

precedence. Link attribute `gr` interface take precedence over corresponding `st` interface and `st` interface takes precedence over `ge` or `xe` underlay interfaces.

sla-export-factor *value*

Set interval to report passive probe report metrics at the application level.

Example: When you configure the `sla-export-factor` as 5, passive probe results are exported once at the end of the 5th, 10th, and 15th probe interval. You can use a passive probe report to report any data that remains unreported in the probe interval at the end of a session.

With application-level summarization, each probe candidate session must send data to central location where the metrics are aggregated. The data thus aggregated is sent out once the configured SLA export factor is met.

- **Range:** 5 through 1000
- **Default:** 500

switch-idle-time *period*

Path switch idle time in seconds. This is the period during which no subsequent switching of application traffic path occurs until the switch idle time expires. This idle time starts when application traffic switches the path.

- **Range:** 5 through 300 seconds

type Define SLA rule type.

saas Select SLA rule type sa Software as a Service (SaaS).

To perform AppQoE for SaaS application, you must define an SLA rule as SaaS and use policy-based APBR profile. You must specify SaaS destination server URL in address book configuration.

To configure AppQoE for SaaS applications:

1. Define the SLA rule type as SaaS (`set security advance-policy-based-routing sla-rule sla1 type saas`).
2. Include SaaS server details in the address book (`set security address-book global address` *`address-book`* `dns-name` *`saas-server-url`* `ipv4-only`).
3. Attach the SLA rule to the policy-based APBR profile.

violation-count *number*

Indicates the number of violations that must occur in a sampling-period for a given session before a link is marked as having violated the SLA.

- **Range:** 1 through 32 seconds
- **Default:** 5

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

The options `preferred-link-type` and `link-type-affinity` are introduced in Junos OS Release 18.4R1.

The options `sla-export-factor` and `violation-count` are moved to [edit security advance-policy-based-routing sla-rule] hierarchy in Junos OS Release 19.2R1.

The options `preferred-tag` and `affinity` are introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

passive-probe-params | **12216**

active-probe-params | **1813**

Application Quality of Experience

Advanced Policy-Based Routing

sla-tracking

IN THIS SECTION

- Syntax | **28028**
- Hierarchy Level | **28028**
- Description | **28028**
- Default | **28029**

- Options | 28029
- Required Privilege Level | 28030
- Release Information | 28030

Syntax

```
sla-tracking {
    route destination-prefix {
        metric metric;
        next-hop next-hop {;
           rpm {
               owner name;
               test test-name;
           }
           twamp {
               connection name;
               test test-name;
           }
        }
        preference preference;
        tag tag;
    }
 }
```

Hierarchy Level

```
[edit routing-options]
```

Description

(Junos OS Evolved only) You configure this configuration statement when you want to install static routes only if they meet the criteria of a given service level agreement (SLA). You configure an RPM or TWAMP test instance with the criteria for the SLA at either the `[edit services monitoring rpm]` or the `[edit services monitoring twamp]` hierarchy levels. You then configure this statement to couple the SLA-tracked

static route with the configured RPM or TWAMP test instance. The system then installs or removes the static routes according to the results of the given RPM or TWAMP test. When installed, these routes are automatically given a preference of 1, and so are preferred over static routes that may already exist with the same prefix. You use the command `show route sla-tracking` to see the results of the tests.

If the test result is "success," then all SLA-tracked static routes that match the probe owner (RPM), connection name (TWAMP), and the test name of the successful test are added to the routing table and are considered to be in an active state. If the test result is "failure," then all SLA-tracked static routes that match the probe owner (RPM), connection name (TWAMP), and test name of the failed test are placed into an inactive state.

Static route tracking supports both IPv4 and IPv6 routes. SLA-tracked static routes are configured individually; wildcards, ranges, and regular expressions are not supported.

Default

See individual options for any defaults.

Options

metric *metric*

(Optional) Set the route metric for the destination prefix. The route with the lowest metric is active in the routing table.

- **Default:** 1
- **Range:** 1 through 16.

next-hop *next-hop*

(Required) Configure a next-hop address. Must be an IPv4 or IPv6 address. You can configure up to 16 multiple paths (next-hops) for any given SLA-tracked static route.

preference *preference*

(Optional) Set the route preference value for the destination prefix. The route with the lowest preference value is active in the routing table. Qualified next hops are not supported with SLA-tracked static routes. Hence, the setting of the `preference`, `metric`, and `tag` options apply only to the SLA-tracked static route and not to the related next hops.

- **Default:** 1

route *destination-prefix*

(Required) Configure an SLA-tracked static route. Must be an IPv4 or IPv6 destination prefix. The maximum allowed number of SLA-tracked routes is 10000.

rpm owner *name* **test** *test-name*

Associate the path for an SLA-tracked static route with a particular RPM owner and test. You must also configure an RPM test owner and test name using the `owner` statement at the `[edit services monitoring rpm]` hierarchy level.

twamp connection ***name*** **test** ***test-name***	Associate the path for an SLA-tracked static route with a particular TWAMP connection and test. You must also configure a TWAMP control connection and test session using the `control-connection` statement at the `[edit services monitoring twamp client]` hierarchy level.
tag ***tag***	(Optional) Set the route tag value for the destination prefix. The route with the lowest tag value is active in the routing table. • **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Support introduced in Junos OS Evolved Release 24.4R1.

slaac-snooping

IN THIS SECTION

- Syntax | 28030
- Hierarchy Level | 28031
- Description | 28031
- Options | 28032
- Required Privilege Level | 28032
- Release Information | 28032

Syntax

```
slaac-snooping {
    interface  (interface-name | all) {
        auto-dad {
```

```
            retries retry-count;
            retrans-interval seconds;
        }
        mark-interface {
            trusted;
        }
        max-allowed-contentions {
            count integer;
            duration seconds;
        }
    }
    link-local {
        expiry interval seconds;
    }
    vlans (vlan-name | all);
}
```

Hierarchy Level

```
[edit forwarding-options access-security]
```

Description

Configure IPv6 stateless address auto-configuration (SLAAC) snooping. SLAAC enables an IPv6 client to generate its own addresses using a combination of locally-available information and information advertised by routers through Neighbor Discovery Protocol (NDP). NDP messages are unsecured, which makes SLAAC susceptible to attacks that involve the spoofing (or forging) of link-layer addresses. IPv6 clients using SLAAC for dynamic address assignment are validated against the SLAAC snooping binding table before being allowed access to the network.

SLAAC snooping is similar to DHCP snooping, in that it snoops packets to build a table of IP-MAC address bindings. SLAAC snooping extracts address information from DAD packets exchanged during the SLAAC process to build the SLAAC snooping table. The address bindings in this table are used to inspect and validate NDP/IP packets sent by IPv6 clients using SLAAC.

NOTE: You must configure SLAAC snooping to allow IPv6 clients using SLAAC access to the network.

The remaining statements are explained separately. See CLI Explorer.

Options

`link-local expiry interval` *`seconds`*	Configure the expiration period for a link-local address learned by SLAAC. When the lease for the address expires, the snooping device sends a DAD message with the client address as the target. If the client is still reachable, the lease is renewed. • **Default:** 86400 seconds • **Range:** 60 to 86400 seconds
`vlans (`*`vlan-name`* `| all)`	Configure SLAAC snooping on a specific VLAN or on all VLANs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

IPv6 Stateless Address Auto-configuration (SLAAC) Snooping

slamon-services

IN THIS SECTION

- Syntax | 28033
- Hierarchy Level | 28033
- Description | 28033
- Options | 28033
- Required Privilege Level | 28033
- Release Information | 28033

Syntax

```
slamon-services rfc2544;
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

(MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, and MX10003 (with the LC2103 card) routers only) Enable service-level agreement (SLA) monitoring services support for RFC 2544-based benchmarking tests on MX Series routers with the MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFP), MPC3E (MX-MPC3E-3D), MPC3E-NG (MX-MPC3E-3D-NG), MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE), MPC5E (MPC5E-40G10G, MPC5EQ-40G10G, MPC5E-100G10G, and MPC5EQ-100G10G), MPC6E (MX2K-MPC6E), MPC7E (MPC7E-MRATE or MPC7E-10G), MX2K-MPC8E or MX2K-MPC9E, MPC10E (MPC10E-15C-MRATE or MPC10E-10C-MRATE), and MX2K-MPC11E line cards that are hosting test interfaces. For aggregated interfaces, enable support for RFC 2544-based benchmarking tests on all MPCs hosting child links. A system log is generated when you enable support for RFC 2544-based benchmarking tests on unsupported MPCs.

> **NOTE**: On MX104 and MX80 Series routers that have a single fixed FPC, this configuration is not required.

Options

`rfc2544`—Enable support for RFC 2544-based benchmarking tests.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

Configuring an RFC 2544-Based Benchmarking Test

Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers

slave

IN THIS SECTION

- MX Series | **28034**
- QFX Series | **28035**
- ACX Series (ACX5448, ACX5448-D, and ACX5448-M) | **28036**
- Hierarchy Level | **28036**
- Description | **28036**
- Options | **28036**
- Required Privilege Level | **28038**
- Release Information | **28038**

MX Series

```
slave {
    announce-interval announce-interval-value
    announce-timeout announce-timeout-value;
    delay-request delay-request-value;
    frequency-only;
    hybrid
    interface interface-name {
        unicast-mode {
            transport ipv4;
            clock-source ip-address {
                local-ip-address local-ip-address {
            }
        }
```

```
        multicast-mode {
           { hybrid
            transport 802.3 link-local;
        }

     primary-asymmetry <primary-asymmetry-value>
     secondary-asymmetry <secondary-asymmetry-value>
    }
     primary <primary-interface>
    secondary <secondary-interface>
   }
    sync-interval interval;
}
```

QFX Series

```
slave {
    interface interface-name {
        unicast-mode {
            transport ipv4;
            clock-client ip-address {
                local-ip-address local-ip-address;
            }
        }
    }
    multicast-mode {
        transport (ipv4 | ieee-802.3)
        local-ip-address local-ip-address;
        local-priority local-ip-address;
    }
max-announce-interval max-announce-interval;
max-delay-response-interval max-delay-response-interval;
max-sync-interval max-sync-interval;
min-announce-interval min-announce-interval;
min-delay-response-interval min-delay-response-interval;
min-sync-interval min-sync-interval;
sync-interval sync-interval;
}
```

ACX Series (ACX5448, ACX5448-D, and ACX5448-M)

```
slave {
    announce-interval announce-interval-value
    announce-timeout announce-timeout-value;
    delay-request delay-request-value;
    grant-duration;
    interface interface-name {
        unicast-mode {
            transport [ipv4 | ipv6];
            clock-source ip-address {
                local-ip-address local-ip-address {
            }
        }
    }
    sync-interval interval;
}
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

Configure the timeReceiver with parameters.

> **NOTE**: Multicast mode is not supported on the QFX Series, ACX5448, ACX5448-D, and ACX5448-M devices.

Options

announce-interval

Configure the logarithmic mean interval for the announce messages to be sent by the timeTransmitter.

- **Range:** 0 through 3 seconds
- **Default:** 1

announce-timeout	Specify the number of announce messages a timeReceiver must miss before an announce-timeout is declared. • **Range:** 2 through 10 seconds • **Default:**
delay-request	Configure the logarithmic mean interval in seconds between the delay request messages sent by the timeReceiver to the timeTransmitter. Specify the grant duration value for the announce messages. • **Range:** 60 through 7200 seconds • **Default:** 300 seconds
frequency-only	Configure frequency synchronization.
hybrid	Configure the timing and synchronization feature to operate in Sync-E assisted PTP mode of operation.
interface *<interface-name>*	Interface on which to respond to upstream PTP timeTransmitter, or aggregated Ethernet ID. For example, interface ae0.0.
unicast-mode	Configure the timeReceiver in unicast mode.
clock-source	Configure the IP address of the timeTransmitter.
multicast-mode	Configure multicast transmission of PTP packets between the timeTransmitter node and the timeReceiver node.
transport 802.3	Configure Ethernet as the encapsulation type for transport of PTP packets.
sync-interval	Configure the logarithmic mean interval for sync interval messages to be sent by the timeTransmitter.
unicast-mode	Configure the timeTransmitter in unicast mode.
clock-client	Configure the IP address of the timeReceiver
transport	Configure the encapsulation type for PTP packet transport.
timeTransmitter-asymmetry *<timeTransmitter-asymmetry-value>*	Configure asymmetry for timeTransmitter link of the aggregated Ethernet bundle. (MX10004 Device)

secondary-asymmetry *<secondary-asymmetry-value>*	Configure asymmetry for secondary link of the aggregated Ethernet bundle. (MX10004 Device)
primary *<primary-interface>*	Configure primary interface name for the aggregated Ethernet bundle. (MX10004 Device)
secondary *<secondary-interface>*	Configure primary interface name for the aggregated Ethernet bundle. (MX10004 Device)
local-ip-address	Configure the IP address of the interface acting as the timeReceiver or the timeTransmitter.
local-priority	Configure a clock's local priority to be used as a tie-breaker in the dataset comparison algorithm, if all other attributes of the datasets being compared are equal.
max-announce-interval	Configure the maximum log mean interval between announce messages.
max-delay-response-interval	Configure the maximum log mean interval between delay-response messages.
max-sync-interval	Configure the maximum log mean interval between synchronization messages.
min-announce-interval	Configure the minimum log mean interval between announce messages.
min-delay-response-interval	Configure the minimum log mean interval between delay-response messages.
min-sync-interval	Configure the minimum log mean interval between synchronization messages.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

Configuring the Precision Time Protocol G.8275.2 Enhanced Profile (Telecom Profile)

slice (CoS Interfaces)

IN THIS SECTION

- Syntax | 28039
- Hierarchy Level | 28039
- Description | 28040
- Options | 28040
- Required Privilege Level | 28040
- Release Information | 28041

Syntax

```
slice name {
 output-traffic-control-profile profile-name;
}
```

Hierarchy Level

```
[edit class-of-service interfaces],
[edit dynamic-profiles name class-of-service interfaces]
```

Description

Assign output traffic control profiles to slices to provide hierarchical CoS for slices.

At the core of the slice per-hop behavior (PHB) in the data plane is the hierarchical class of service (CoS) queuing. In the network slicing architecture, each slice has a designated queue set consisting of queues corresponding to traditional behavioral aggregate (BA) classifiers. Similar to hierarchical queuing for logical interfaces, you can define hierarchical CoS (HCoS) for transport slices. You can configure up to three levels of HCoS for transport slices:

- Physical or aggregated Ethernet interface (Level 1)
- Individual slices (Level 2)
- Eight queues per slice (Level 3)

NOTE:

- You cannot configure hierarchical CoS for both slices and logical interfaces on the same physical or aggregated Ethernet interface.
- You cannot configure hierarchical CoS for slices on a device if the device also has "policy map" on page 12686 configured.
- You must define the slice name at [`edit services network-slicing`].
- You must enable `enhanced-ip` at [`edit chassis network-services`].
- You must enable `hierarchical-scheduler` at [`edit interfaces` *`interface-name`*].

Options

name	Name of the network slice
output-traffic-control-profile	Name of the traffic control profile

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3R1 for MX Series routers.

RELATED DOCUMENTATION

Hierarchical Class of Service for Network Slicing

show interfaces queue | 21208

slice (export routing policy action)

IN THIS SECTION

- Syntax | 28041
- Hierarchy Level | 28041
- Description | 28042
- Required Privilege Level | 28042
- Release Information | 28042

Syntax

```
slice slice-name
```

Hierarchy Level

```
[edit policy-options policy-statement policy-name term term-name then]
```

Description

Enables a matched route to be marked with the slice identifier corresponding to the slice name specified in the Class of Service configuration. The `slice` export routing policy action is supported for routes that also undergo color or transport-class based route resolution.

Required Privilege Level

- routing—To view this statement in the configuration
- routing-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS 22.2R1.

slice (firewall filter action)

IN THIS SECTION

- Syntax | 28042
- Hierarchy Level | 28042
- Description | 28043
- Required Privilege Level | 28043
- Release Information | 28043

Syntax

```
slice slice-name
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name then]
```

Description

Mark packets that pass the match conditions of the rule with the slice identifier corresponding to the slice name specified in the Class of Service configuration.

NOTE: *family-name* can be `inet`, `inet6`, `mpls`, or `any`.

Required Privilege Level

- firewall—To view this statement in the configuration
- firewall-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS 22.2R1.

slice (firewall filter match condition)

IN THIS SECTION

- Syntax | 28043
- Hierarchy Level | 28044
- Description | 28044
- Required Privilege Level | 28044
- Release Information | 28044

Syntax

```
slice slice-name
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name from]
```

Description

Match the slice identifier of the packet.

> **NOTE**: *family-name* can be `inet`, `inet6`, `mpls`, or `any`.

Required Privilege Level

- firewall—To view this statement in the configuration
- firewall-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS 22.2R1.

slow-pfe-alarm

IN THIS SECTION

- Syntax | 28045
- Hierarchy Level | 28045
- Description | 28045
- Required Privilege Level | 28045
- Release Information | 28045

Syntax

```
slow-pfe-alarm;
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable the slow Packet Forwarding Engine alarm on a M Series, MX Series, or a T Series router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1R6, 12.2R5, 12.3R3, 13.1R2, and 13.2R1.

RELATED DOCUMENTATION

Configuring Slow Packet Forwarding Engine Alarm

smart-check

IN THIS SECTION

- Syntax | 28046
- Hierarchy Level | 28046
- Description | 28046

- Options | 28046
- Required Privilege Level | 28047
- Release Information | 28047

Syntax

```
smart-check {
    ssd-series ssd-series-name {
        id id-num {
            id-threshold id-threshold-value;
            id-value (norm | raw);
            id-flag (high | low);
        }
    }
}
```

Hierarchy Level

```
[edit chassis routing-engine disk]
```

Description

Configure the device to perform health monitoring on the Routing Engine disk and raise alarm when an error occurs. You can specify threshold value for each smart attribute of the solid-state drive (SSD). The device raises an alarm when the threshold is breached. You can view the alarms by using the command `show chassis alarms`.

Options

disk	Configure health check parameters for the SSD.
ssd-series *ssd-series-name*	Specify the SSD series name.

id *id-num*

Specify the ID of the SSD smart attribute for which you are configuring the health check. You can view the smart attributes and their IDs of an SSD by using the command `show vmhost hard-disk-test status disk /dev/sda`.

- **Range:** 1 - 255

id-threshold

Specify the SSD smart attribute threshold, breaching which will result in the device logging a health event or raising an alarm.

id-value

Choose the attribute value type from the following options:

- norm—Use this option to instruct the device to consider normalized value for the specified smart attribute.
- raw—Use this option to instruct the device to consider raw value for the specified smart attribute.

id-flag

Choose the attribute value flag from the following options:

- high—The software considers the threshold to be breached if the value is greater than smart threshold.
- low—The software considers the threshold to be breached if the value is less than smart threshold.

smart-check

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Managing Errors

smart-optimize-timer

IN THIS SECTION

- Syntax | 28048
- Hierarchy Level | 28048
- Description | 28048
- Default | 28049
- Options | 28049
- Required Privilege Level | 28049
- Release Information | 28049

Syntax

```
smart-optimize-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Enable the smart optimization timer. When you enable the smart optimization timer on a router, the Junos OS operates on the assumption that the original LSP path is preferable to any alternate or secondary path. When you enable the smart optimization timer and an LSP fails and its traffic is switched to an alternate path, the smart optimization timer starts and waits 3 minutes (this time is configurable). After 3 minutes have passed, the LSP is switched back to the original path. If the original path fails again and the LSP is switched to an alternate path again, the router waits 1 hour before attempting to switch the LSP back to its original path.

NOTE: This feature is not applicable for LSPs for which `no-cspf` is configured.

If you want to disable the smart optimizer, you can set it to zero. The `smart-optimize-timer` value in seconds indicates the time before which the LSP is switched back to its primary path in case the primary path becomes available. Otherwise, the time to wait is controlled by the `optimize-timer` , which is usually set to a high value. Some ISPs have the `optimize-timer` set to once a day. Sometimes after the smart optimizer causes the LSP to be placed back on its primary path, the primary path goes down again within 60 minutes. When this happens, the `smart-optimize-timer` is disabled automatically, and the `optimize-timer` (regular path optimization) goes into effect. This is to protect against a flapping link being used.

Default

The smart optimization timer is enabled by default.

Options

seconds—(Optional) Specify the number of seconds to wait before switching an LSP back to its original path. If you do not specify the number of seconds, the default value is used.

- **Range:** 0 through 65,535 seconds
- **Default:** 180 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Smart Optimize Timer for LSPs

Optimizing Signaled LSPs

optimize-aggressive | **11855**

optimize-timer (Protocols MPLS) | **11859**

smg-service (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 28050
- Hierarchy Level | 28050
- Description | 28050
- Required Privilege Level | 28051
- Release Information | 28051

Syntax

```
smg-service {
    failover other-routing-engine;
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag <disable>;
        level level;
        no-remote-trace
    }
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Configure system services, including tracing operations and Routing Engine failover, for the main enhanced subscriber management session management process, `smg-service`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

trace—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

Configuring the Subscriber Management Database Trace Log Filename

smtp-profile (Security Content Security Policy Antispam)

IN THIS SECTION

- Syntax | 28051
- Hierarchy Level | 28052
- Description | 28052
- Required Privilege Level | 28052
- Release Information | 28052

Syntax

```
smtp-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name anti-spam]
```

Description

Configure a Content Security policy for the antispam SMTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

smtp-profile (Security Content Security Policy Antivirus)

IN THIS SECTION

- Syntax | 28053
- Hierarchy Level | 28053
- Description | 28053
- Required Privilege Level | 28053
- Release Information | 28053

Syntax

```
smtp-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name anti-virus]
```

Description

Configure a Content Security policy for the antivirus SMTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

smtp-profile (Security Content Security Policy Content Filtering)

IN THIS SECTION

- Syntax | 28054
- Hierarchy Level | 28054
- Description | 28054
- Required Privilege Level | 28054

Release Information | 28054

Syntax

```
smtp-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name content-filtering]
```

Description

Configure a Content Security policy for the content-filtering SMTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

snapshot

IN THIS SECTION

- Syntax | 28055
- Hierarchy Level | 28055
- Description | 28056
- Options | 28056
- Required Privilege Level | 28056
- Release Information | 28056

Syntax

```
snapshot {
   configuration;
   no-configuration;
   recovery {
      automatic;
      configuration-only;
      configuration-preference;
      no-automatic;
      on-upgrade;
   }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure system snapshot options. The `configuration` and `no-configuration` options are used for creating non-recovery snapshots and `recovery` is used for recovery snapshot.

This statement provides flexibility and control over the snapshot process.

Options

configuration Save configuration in snapshot. When you configure the `configuration` option and issue the `request system snapshot` command, it stores the configuration files as part of the non-recovery snapshot.

no-configuration When you configure the `no-configuration` option and issue the `request system snapshot` command, it does not store the configuration files in the non-recovery snapshot.

recovery This option is used for system snapshot recovery management. This option allows you to create recovery snapshots for the purpose of system recovery.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

System-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

snext

IN THIS SECTION

- Syntax | 28057
- Hierarchy Level | 28057
- Description | 28057

- Options | 28057
- Required Privilege Level | 28057
- Release Information | 28057

Syntax

```
snext margin;
```

Hierarchy Level

```
[edit interfaces interface-name shdsl-options snr-margin],
[edit logical-systems logical-system-name interfaces interface-name shdsl-options snr-margin]
```

Description

For J Series Services Routers only, configure self-near-end crosstalk (SNEXT) signal-to-noise ratio (SNR) margin for a SHDSL line. When configured, the line trains at higher than SNEXT threshold. The SNR margin is the difference between the desired SNR and the actual SNR.

Options

margin—Desired SNEXT margin. Possible values are disabled or a margin between –10dB and 10 dB.

- **Default:** `disabled`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

snmp

IN THIS SECTION

- Syntax: MX, M, T, ACX, PTX Series Routers and EX Series Switches | 28058
- Syntax: QFX Series Switches, QFabric, OCX1100 and EX4600 | 28058
- Options | 28063
- Hierarchy Level | 28063
- Description | 28064
- Required Privilege Level | 28064
- Release Information | 28064

Syntax: MX, M, T, ACX, PTX Series Routers and EX Series Switches

```
snmp { ... }
```

Syntax: QFX Series Switches, QFabric, OCX1100 and EX4600

```
snmp {
    client-list client-list-name {
        ip-addresses;
    }
     community community-name {
        authorization authorization;
        client-list-name client-list-name;
        clients {
            address restrict;
        }
        logical-system logical-system-name {
            routing-instance routing-instance-name {
                clients {
                addresses;
                }
            }
        }
```

```
        routing-instance routing-instance-name {
            clients {
            addresses;
            }
        }
        view view-name;
    }
    contact contact;
    description description;
    filter-duplicates;
    filter-interfaces;
    health-monitor {
        falling-threshold integer;
        interval seconds;
        rising-threshold integer;
    }
    interface [ interface-names ];
    location location;
    name name;
    nonvolatile {
        commit-delay seconds;
    }
    rmon {
        alarm index {
            description description;
            falling-event-index index;
            falling-threshold integer;
            No Link Title seconds;
            interval seconds;
            No Link Title;
            No Link Title index;
            rising-threshold integer;
            No Link Title (absolute-value | delta-value);
            No Link Title (falling-alarm | rising-alarm | rising-or-falling alarm);
            syslog-subtag syslog-subtag;
            variable oid-variable;
        }
        No Link Title index {
            community community-name;
            description description;
            type type;
        }
        history history-index {
```

```
            bucket-size number;
            interface interface-name;
            interval seconds;
            owner (RMON) owner-name;
        }
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
regular-expression>;
        flag flag;
    }
    trap-group group-name {
        categories {
            category;
        }
        destination-port port-number;
        routing-instance routing-instance-name;
        targets {
            address;
        }
        version (all | v1 | v2);
    }
    trap-options {
        agent-address outgoing-interface;
        source-address address;
    }
    No Link Title {
        notify name {
            tag tag-name;
            type trap;
        }
        notify-filter profile-name {
            oid object-identifier (include | exclude);
        }
        snmp-community community-index {
            community-name community-name;
            security-name security-name;
            tag tag-name;
        }
        target-address target-address-name {
            address address;
            address-mask address-mask;
            logical-system logical-system;
```

```
            port port-number;
            retry-count number;
            routing-instance routing-instance-name;
            tag-list tag-list;
            target-parameters target-parameters-name;
            timeout seconds;
        }
        target-parameters target-parameters-name {
            notify-filter profile-name;
            parameters {
                message-processing-model (v1 | v2c | V3);
                security-level (authentication | none | privacy);
                security-model (usm | v1 | v2c);
                security-name security-name;
            }
        }
        usm {
            local-engine {
                user username {
                    authentication-sha {
                        authentication-password authentication-password;
                    }
                    authentication-md5 {
                        authentication-password authentication-password;
                    }
                    authentication-none;
                    privacy-aes128 {
                        privacy-password privacy-password;
                    }
                    privacy-des {
                        privacy-password privacy-password;
                    }
                    privacy-3des {
                        privacy-password privacy-password;
                    }
                    privacy-none;
                }
            }
            remote-engine engine-id {
                user username {
                    authentication-sha {
                        authentication-password authentication-password;
                    }
```

```
                    authentication-md5 {
                        authentication-password authentication-password;
                    }
                    authentication-none;
                    privacy-aes128 {
                        privacy-password privacy-password;
                    }
                    privacy-des {
                        privacy-password privacy-password;
                    }
                    privacy-3des {
                        privacy-password privacy-password;
                    }
                    privacy-none {
                        privacy-password privacy-password;
                    }
                }
            }
        }
        vacm {
            access {
                group group-name {
                    (default-context-prefix | context-prefix context-prefix) {
                        security-model (any | usm | v1 | v2c) {
                            security-level (authentication | none | privacy) {
                                notify-view view-name;
                                read-view view-name;
                                write-view view-name;
                            }
                        }
                    }
                }
            }
            security-to-group {
                security-model (usm | v1 | v2c) {
                    security-name security-name {
                        group group-name;
                    }
                }
            }
        }
    }
    view view-name {
```

```
        oid object-identifier (include | exclude);
    }
}
}
```

Options

`notify` *`name`* is the name assigned to the notification. Each notify entry name must be unique.

`tag` *`tag-name`* defines the target addresses that are sent this notification. The notification is sent to all target addresses that have this tag in their tag list. The *`tag-name`* is not included in the notification.

`type inform` is the type of notification.

`target-address` *`target-address-name`* identifies the target address. The target address defines a management application's address and parameters that are used to respond to informs.

`timeout` *`seconds`* is the number of seconds to wait for an acknowledgment. If no acknowledgment is received within the timeout period, the inform is retransmitted. The default timeout is `15` seconds.

`retry-count` *`number`* is the maximum number of times an inform is transmitted if no acknowledgment is received. The default is `3`. If no acknowledgment is received after the inform is transmitted the maximum number of times, the inform message is discarded.

`message-processing-model` defines which version of SNMP to use when SNMP notifications are generated. Informs require a `v3` message processing model.

`security-model` defines the security model to use when SNMP notifications are generated. Informs require a `usm` security model.

`security-level` specifies whether the inform is authenticated and encrypted before it is sent. For the `usm` security model, the security level must be one of the following:

- `authentication`—Provides authentication but no encryption.
- `privacy`—Provides authentication and encryption.

`security-name` identifies the username that is used when generating the inform.

Hierarchy Level

```
[edit]
```

Description

Configure SNMP.

SNMP modules cannot have the slash (/) character or the @ character in the name.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

snmp-command

IN THIS SECTION

- Syntax | 28064
- Hierarchy Level | 28065
- Description | 28065
- Options | 28065
- Required Privilege Level | 28065
- Release Information | 28065

Syntax

```
snmp-command command;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

SNMP command format.

Options

command—Supported commands are SNMP `get`, `get-next`, `set`, and `trap`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *ALG Descriptions*
- *Configuring an SNMP Command for Packet Matching*
- *Examples: Configuring Application Protocols*
- *Verifying the Output of ALG Sessions*

snmp-community

IN THIS SECTION

Syntax | 28066

- Hierarchy Level | 28066
- Description | 28066
- Options | 28066
- Required Privilege Level | 28066
- Release Information | 28067

Syntax

```
snmp-community community-index {
    community-name community-name;
    security-name security-name;
    tag tag-name;
}
```

Hierarchy Level

```
[edit snmp v3]
```

Description

Configure the SNMP community.

Options

community-index—(Optional) String that identifies an SNMP community.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

snmp-options

IN THIS SECTION

- Syntax | 28067
- Hierarchy Level | 28067
- Description | 28067
- Options | 28068
- Required Privilege Level | 28068
- Release Information | 28068

Syntax

```
snmp-options backward-traps-only-from-established;
```

Hierarchy Level

```
[edit protocols bgp]
```

Description

Configure the SNMP options and customize the behavior of BGP-related MIBs. By default, SNMP generates traps when a session moves from a higher state to a lower state. For example, from `Active` to `Idle` or from `Established` to `Idle`. This can result in many uninteresting traps, especially when there are a large number of unconfigured BGP sessions that toggle continuously between`Active` and `Idle` state. Set the `backward-traps-only-from-established` option to eliminate the uninteresting traps, and receive backward trap notifications only when transitioning away from the `Established` state. This can substantially reduce the number of traps sent.

Options

backward-traps-only-from-established Limit the generation of traps for backward transitions to session states that are moving from `Established` state only.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Example: Viewing BGP Trace Files on Logical Systems

snmp (SNMP Scripts)

IN THIS SECTION

- Syntax | 28068
- Hierarchy Level | 28069
- Description | 28069
- Required Privilege Level | 28070
- Release Information | 28070

Syntax

```
snmp {
    file filename {
```

```
            cert-file path;
            checksum hash-algorithm hash-value;
            oid id {
                priority;
            }
            python-script-user username;
            refresh;
            refresh-from;
            routing-instance routing-instance-name;
            source;
            start {
                start-options {
                    mem-factor percentage;
                }
            }
        }
        max-datasize size;
        refresh;
        refresh-from;
        start {
            start-options {
                mem-factor percentage;
            }
        }
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
```

Hierarchy Level

```
[edit system scripts]
```

Description

For Junos OS SNMP scripts, configure the SNMP scripting mechanism.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Storing and Enabling Scripts

snmp-trap-thresholds

IN THIS SECTION

- Syntax | 28070
- Hierarchy Level | 28071
- Description | 28071
- Options | 28071
- Required Privilege Level | 28072
- Release Information | 28072

Syntax

```
snmp-trap-thresholds {
    flows high high-threshold | low low-threshold;
    nat-address-port high high-threshold | low low-threshold;
    }
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Configures SNMP flow thresholds for all flows for a service set or flows for all NAT pools configured for a service set..

Options

The remaining options are described separately.

`flows high` *`high-threshold`*—Configure the upper limit for all flows on the service set. The limit is expressed as a percentage of `max-flows` configured for the service set. When the number of active flows exceeds this limit, an SNMP trap is set.

- **Default:** 90 percent of `max-flows`

`flows low` *`low-threshold`*—Configure the lower limit for all flows on the service set . The limit is expressed as a percentage of `max-flows` configured for the service set. When the number of active flows falls below this limit, an SNMP trap is set.

- **Default:** 70 percent of `max-flows`

`nat-address-port high` *`high-threshold`*—Configure the upper limit for flows for all NAT pools on the service set. The limit is expressed as a percentage of `max-flows` configured for the service set. When the number of active flows exceeds this limit, an SNMP trap is set.

- **Default:** 90 percent of `max-flows`

`nat-address-port low` *`low-threshold`*—Configure the lower limit for flows. The limit is expressed as a percentage of `max-flows` configured for the service set. When the number of active flows falls below this limit, an SNMP trap is set.

- **Default:** 80 percent of `max-flows`

> **NOTE**: SNMP traps that are generated when you modify the threshold value for flows of NAT address pools in a service set (by using the `snmp-trap-thresholds nat-address-port (high` *`high-threshold`* `| low` *`low-threshold`*) statement) are not effective in the PIC. Only the initial threshold value that is set is effective on the PIC and subsequent changes to the threshold value are not reflected on the PIC. As a workaround, for the configuration

changes under the `[edit services nat pool nat-pool-name]` hierarchy level, you must deactivate and activate the relevant service-set to enable the updated configuration to become effective. Otherwise, you must reboot the PIC for the updated threshold value of to take effect.

NOTE: Until Junos OS Release 14.1, when the NAT pool utilization exceeded the high threshold value configured, an SNMP trap was sent. However, a similar SNMP trap was not triggered when the NAT pool utilization fell below the configured lower limit or threshold. Because NMS systems are being used to monitor and set alarm for threshold values, the absence of an SNMP trap when the low threshold value was reached caused NMS to retain an active alarm in the alarms list. As a result, starting with Release 14.2R1, an SNMP trap is generated when the NAT pool utilization reaches the lower threshold, thereby causing the alarm in NMS to be reset.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Configuring Service Set Limitations

snmp-trap-thresholds (Next Gen Services)

IN THIS SECTION

- Syntax | **28073**
- Hierarchy Level | **28073**

- Description | 28073
- Options | 28073
- Required Privilege Level | 28073
- Release Information | 28074

Syntax

```
snmp-trap-thresholds {
    flow high percent low percent;
    nat-address-port high percent low percent;
    session high percent low percent;
}
```

Hierarchy Level

```
[edit services service-set]
```

Description

Define snmp traps for Next Gen Services service sets.

Options

session Specify the low and high session threshold limits for generating SNMP traps.

The default for high = 90%.

The default for low = 70%.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

snmp-value-match-msmic (Services NAT Options)

IN THIS SECTION

- Syntax | 28074
- Hierarchy Level | 28074
- Description | 28074
- Required Privilege Level | 28074
- Release Information | 28075

Syntax

```
snmp-value-match-msmic;
```

Hierarchy Level

```
[edit services service-set service-set-name nat-options]
```

Description

Match the values for MS-MIC-specific objects in the jnxNatObjects MIB table with the values for MS-DPC objects. Use the `deactivate services service-set *service-set-name* nat-options snmp-value-match-msmic` configuration mode command to disable this feature. By default, this feature is disabled.

Required Privilege Level

control—To add this statement to the configuration.

Release Information

Statement introduced with Junos OS Release 13.3R7

RELATED DOCUMENTATION

> *Configuring Service Rules*
>
> *Troubleshooting the Mismatch of jnxNatObjects Values for MS-DPC and MS-MIC*

snoop-pseudowires

IN THIS SECTION

- Syntax | 28075
- Hierarchy Level | 28075
- Description | 28076
- Required Privilege Level | 28076
- Release Information | 28076

Syntax

```
snoop-pseudowires;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name igmp-snooping-options]
[edit logical-systems logical-system -name routing-instances routing-instance-nameigmp-snooping-
options]
```

Description

The default IGMP snooping implementation for a VPLS instance adds each pseudowire interface to its **oif** list. It includes traffic from the ingress PE that is sent to egress PE even if there is no interest. The `snoop-pseudowires` option prevents multicast traffic from traversing the pseudowire (to egress PEs) unless there are IGMP receivers for the traffic. In other words, multicast traffic is forwarded only to VPLS core interfaces that are router interfaces, or that are IGMP receivers. In addition to the benefit of sending traffic to only interested PEs, `snoop-pseudowires` also optimizes a common path between PE-P routers wherever possible (so if two PEs connect via the same P router, only one copy of packet is sent; the packet would be replicated only on P routers for which the path is divergent).

> **NOTE**: Note that this option can only be enabled when `instance-type` is *vpls*. The `snoop-pseudowires` option cannot be enabled if `use-p2mp-lsp` is enabled for `igmp-snooping-options`.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

instance-type | **8773**

Example: Configuring IGMP Snooping

snoop-segment (TDF Domain Selection)

IN THIS SECTION

- Syntax | **28077**
- Hierarchy Level | **28077**

- Description | 28077
- Options | 28077
- Required Privilege Level | 28077
- Release Information | 28077

Syntax

```
snoop-segment snoop-segment-name;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from]
```

Description

Specify the snoop segment that matches the RADIUS request.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

Options

snoop-segment-name	Name of the snoop segment.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers

snoop-segments (RADIUS)

IN THIS SECTION

- Syntax | 28078
- Hierarchy Level | 28079
- Description | 28079
- Options | 28079
- Required Privilege Level | 28079
- Release Information | 28079

Syntax

```
snoop-segments snoop-segment-name {
    destination-ip-address destination-address;
    <destination-port destination-port;>
    <request-cache-timeout timeout;>
    shared-secret secret;
    source-interface source-interface;
    <source-ip-address source-address;>
}
```

Hierarchy Level

```
[edit access radius]
```

Description

Specify which accounting messages to snoop. You must specify at least the destination IP address for the accounting messages, the shared secret, and the source interface.

Options

snoop-segment-name Name for the snoop segment. The maximum length is 32 characters.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview

IP-Based Subscriber Setup Overview

snoop-segments (TDF Gateway)

IN THIS SECTION

- Syntax | 28080
- Hierarchy Level | 28080
- Description | 28080
- Options | 28080
- Required Privilege Level | 28080
- Release Information | 28081

Syntax

```
snoop-segments [snoop-segment-name];
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name aaa]
```

Description

Specify one or more snoop segments that control RADIUS accounting request snooping for the TDF gateway. The snoop segments must already be configured at the `[edit access radius]` hierarchy level.

Options

snoop-segment-name	Name of a snoop segment.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview

IP-Based Subscriber Setup Overview

snr-margin

IN THIS SECTION

- Syntax | **28081**
- Hierarchy Level | **28081**
- Description | **28082**
- Required Privilege Level | **28082**
- Release Information | **28082**

Syntax

```
snr-margin {
    snext margin;
}
```

Hierarchy Level

```
[edit interfaces interface-name shdsl-options],
[edit logical-systems logical-system-name interfaces interface-name shdsl-options]
```

Description

For J Series Services Routers only, configure the SHDSL signal-to-noise ratio (SNR) margin. The SNR margin is the difference between the desired SNR and the actual SNR. Configuring the SNR creates a more stable SHDSL connection by making the line train at a SNR margin higher than the threshold. If any external noise below the threshold is applied to the line, the line remains stable.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

sockets

IN THIS SECTION

- Syntax | 28082
- Hierarchy Level | 28083
- Description | 28083
- Required Privilege Level | 28083
- Release Information | 28083

Syntax

```
sockets value;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name]
```

Description

Enter the number of sockets used for communicating between the client and server. The default is 1.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

soft-gre

IN THIS SECTION

- Syntax | 28084
- Hierarchy Level | 28084
- Description | 28084
- Options | 28084
- Required Privilege Level | 28085
- Release Information | 28085

Syntax

```
soft-gre group-name {
    destination-networks [prefix] {
    dynamic-profile profile-name;
    service-interface psn;
    source-address wag-ip-address;
    <tunnel-idle-timeout seconds>;
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the conditions for enabling dynamic-bridged generic routing encapsulation (GRE) tunnel creation on the MX Series router Wi-Fi access gateway (WAG).

NOTE: Configuration of multiple dynamic tunnel groups is supported.

Options

destination-networks [*prefix*]
Use the specified IP subnets from which soft-GRE connection requests from the customer can be processed.

group-name
Name of the dynamic GRE tunnel group.

dynamic-profile *profile-name*
Name of the dynamic profile that creates the tunnel.

NOTE: To support VLAN autosensing on a GRE tunnel, you must also specify the `auto-configure` options at the `[edit dynamic-profile profile-name interfaces unit]` hierarchy level. These options include a reference to the dynamic profile that creates VLANs.

service-interface ps*n*
Use the specified pseudowire subscriber interface device (IFD) on which the tunnels are built.

source-address *wag-ip-address*
Use the specified source IP address of the GRE tunnels for the WAG. This is the IP address on which incoming GRE traffic must be received by the MX Series router.

tunnel-idle-timeout *seconds*
(Optional) Use the specified number of seconds that a GRE tunnel remains up after the last subscriber session on the tunnel has ended. If set to 0, then no idle timeout is set, and the tunnel remains up for an unlimited period of time.

- **Range:** 0 through 65535
- **Default:** 120

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Wi-Fi Access Gateway Deployment Model Overview

Supported Access Models for Dynamic-Bridged GRE Tunnels on the Wi-Fi Access Gateway

Configuring Conditions for Enabling Dynamic-Bridged GRE Tunnel Creation

show services soft-gre tunnel | **26186**

soft-limit

IN THIS SECTION

- Syntax | **28086**
- Hierarchy Level | **28086**
- Description | **28086**
- Options | **28086**

Required Privilege Level | 28086
Release Information | 28086

Syntax

```
soft-limit bandwidth;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Specify a bandwidth threshold at which congestion notifications are sent to each control source of the criteria that point to this content destination. If the control source is configured with the `syslog` statement, a log message also be generated.

Options

bandwidth—Soft limit threshold, in bits per second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring the Content Destination

soft-limit-clear

IN THIS SECTION

- Syntax | 28087
- Hierarchy Level | 28087
- Description | 28087
- Options | 28087
- Required Privilege Level | 28087
- Release Information | 28088

Syntax

```
soft-limit-clear bandwidth;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Specify a bandwidth threshold at which the latch set by the soft-limit threshold is cleared.

Options

bandwidth—Soft-limit clear threshold, in bits per second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring the Content Destination

soft-limit | 28085

soft-preemption (Protocols MPLS)

IN THIS SECTION

- Syntax | 28088
- Hierarchy Level | 28088
- Description | 28088
- Required Privilege Level | 28089
- Release Information | 28089

Syntax

```
soft-preemption;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Attempt to establish a new path for a preempted LSP before tearing it down.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring MPLS Soft Preemption

soft-preemption (Protocols RSVP)

IN THIS SECTION

- Syntax | 28089
- Hierarchy Level | 28090
- Description | 28090
- Options | 28090
- Required Privilege Level | 28090
- Release Information | 28090

Syntax

```
soft-preemption {
    cleanup-timer seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp preemption],
[edit protocols rsvp preemption]
```

Description

Enable soft preemption to attempt to establish a new path for a preempted LSP before tearing it down.

Options

`cleanup-timer`—A value of 0 disables soft preemption.

- **Range:** 0 through 10800 seconds
- **Default:** 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring MPLS Soft Preemption

software-upgrade

IN THIS SECTION

Syntax | 28091

- Hierarchy Level | 28091
- Description | 28091
- Required Privilege Level | 28091
- Release Information | 28091

Syntax

```
software-upgrade
```

Hierarchy Level

```
[edit chassis high-availability]
```

Description

Enable software upgrade mode. Use this configuration option when you want to upgrade your security device in Multinode High Availability. This statement marks SRG status as offline and diverts the transit traffic to the other node.

Once you complete the software upgrade, delete the statement on both nodes using the command `delete chassis high-availability software-upgrade`.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Software Upgrade in Multinode High Availability

softwire-concentrator

IN THIS SECTION

- Syntax | 28092
- Hierarchy Level | 28092
- Description | 28093
- Options | 28093
- Required Privilege Level | 28093
- Release Information | 28093

Syntax

```
softwire-concentrator {
    ds-lite ds-lite-softwire-concentrator {
        auto-update-mtu;
        flow-limit flow-limit | session-limit-per-prefix session-limit-per-prefix;
        mtu-v6 bytes;
        softwire-address address;
    }
    map-e
    v6rd v6rd-softwire-concentator {
        ipv4-prefix ipv4-prefix;
        v6rd-prefix ipv6-prefix;
        mtu-v4 mtu-v4;
    }
}
```

Hierarchy Level

```
[edit services softwire]
```

Description

Configure settings for a softwire concentrator.

Softwires are supported on the MS-DPC, MS-100, MS-400, and MS-500 line cards. Starting in Junos OS release 17.4R1, softwires for DS-Lite are supported on MX Series routers with MS-MPCs and MS-MICs.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`map-e` option introduced in Junos OS Release 18.2R1 for MX Series Routers with MPC and MIC interfaces.

RELATED DOCUMENTATION

Configuring a DS-Lite Softwire Concentrator

Configuring a 6rd Softwire Concentrator

softwire-name (Next Gen Services)

IN THIS SECTION

- Syntax | 28094
- Hierarchy Level | 28094
- Description | 28094

- Options | 28094
- Required Privilege Level | 28095
- Release Information | 28095

Syntax

```
softwire-name v6rd-softwire-concentrator {
    ipv4-prefix ipv4-prefix;
    mtu-v4 number-of-bytes;
    softwire-concentrator address;
    softwire-type v6rd;
    v6rd-prefix v6rd-prefix
}
```

Hierarchy Level

```
[edit services softwires]
```

Description

Configure a 6rd softwire concentrator. A 6rd softwire allows an IPv6 end user to send traffic over an IPv4 network to reach an IPv6 network. The softwire concentrator decapsulates IPv6 packets that were encapsulated in IPv4 packets by a software initiator at the customer edge WAN, and forwards the packets for IPv6 routing.

Options

ipv4-prefix *ipv4-prefix*

IPv4 prefix of the customer edge (CE) network.

mtu-v4 *number-of-bytes*

The size, in bytes, of the maximum transmission unit for IPv6 packets encapsulated in IPv4. Compute this as the maximum expected IPv4 packet size plus 20. Packets that are larger than the configured value are dropped.

- **Range:** 576 through 9192

softwire-concentrator ***address***	IPv4 address of a softwire concentrator. This is an IPv4 address independent of any interface and on a different prefix.
softwire-name ***v6rd-softwire-concentrator***	Name of the softwire concentrator.
softwire-type v6rd	Sets softwire concentrator type to 6rd.
v6rd-prefix ***v6rd-prefix***	IPv6 prefix for the 6rd domain.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

6rd Softwires in Next Gen Services

softwire-options

IN THIS SECTION

- Syntax | 28096
- Hierarchy Level | 28096
- Description | 28096
- Options | 28096
- Required Privilege Level | 28096
- Release Information | 28096

Syntax

```
softwire-options {
    dslite-ipv6-prefix-length dslite-ipv6-prefix-length ;
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the IPv6 prefix length associated with a subscriber's basic broadband bridging device that is subject to a limited number of sessions.

This feature is supported on MX Series routers equipped with MS-DPCs. Starting in Junos OS Release 18.2R1, this option is also supported on MS-MPCs and MS-MICs.

Options

dslite-ipv6-prefix-length Subnet prefix representing the size of the subnet subject to session limitation.

- **Values:** 56, 64, 96, 128
- **Default:** 0—no limitation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Support added in Junos OS 20.2R1 for Next Gen Services on MX240, MX480, and MX960 routers.

RELATED DOCUMENTATION

DS-Lite Per Subnet Limitation Overview

softwire-rules

IN THIS SECTION

- Syntax | 28097
- Hierarchy Level | 28097
- Description | 28097
- Options | 28098
- Required Privilege Level | 28098
- Release Information | 28098

Syntax

```
(softwire-rule rule-name | softwire-rule-sets rule-set-name);
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the DS-Lite or 6rd rules or rule set included in this service set. You can configure multiple rules; however, you can only configure one rule set for each service set.

Softwire rules are supported on the MS-DPC, MS-100, MS-400, and MS-500 line cards. Starting in Junos OS release 17.4R1, softwire rules for DS-Lite are supported on MX Series routers with MS-MPCs and MS-MICs.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

rule-set-name—Identifier for the set of rules to be included.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Service Rules

softwire-types (Next Gen Services)

IN THIS SECTION

- Syntax | 28099
- Hierarchy Level | 28099
- Description | 28099
- Options | 28099
- Required Privilege Level | 28101
- Release Information | 28101

Syntax

```
softwire-types {
    ds-lite ds-lite-softwire-concentrator {
        auto-update-mtu;
        flow-limit flow-limit | session-limit-per-prefix session-limit-per-prefix;
        mtu-v6 bytes;
        softwire-address address;
    }
    map-e
    v6rd v6rd-softwire-concentator {
        ipv4-prefix ipv4-prefix;
        v6rd-prefix ipv6-prefix;
        mtu-v4 mtu-v4;
    }
}
```

Hierarchy Level

```
[edit services softwires]
```

Description

Configure ds-lite, 6rd and MAP-E softwire objects.

Options

The following options are available for each type of softwire:

ds-lite Specify options for DS-Lite softwites.

v6rd Specify options for v6rd softwites.

map-e Specify options for map-e softwires.

auto-update-mtu This option is not currently supported.

copy-dscp Copy DSCP information to IPv4 headers during decapsulation.

flow-limit —Maximum number of IPv4 flows per softwire.

ipv4-prefix IPv4 prefix of the customer edge (CE) network

mtu-v4 Maximum transmission unit (MTU), in bytes (576 through 9192), for IPv6 packets encapsulated into IPv4. If the final length is greater than the configured value, the IPv4 packet is dropped. This option is mandatory except for DS-Lite softwires since it depends on other network parameters under administrator control.

mtu-v6 Maximum transmission unit when encapsulating IPv4 packets into IPv6. If the final length is greater than the MTU, the IPv6 packet is fragmented. This option is mandatory since it depends on other network parameters under administrator control.

session-limit-per-prefix Maximum number of sessions per B4 subnet prefix.

softwire-concentrator Specify the IP address of the softwire concentrator.

softwire-type Sets softwire concentrator type to 6rd.

- **Values:** v6rd

v6rd-prefix IPv6 prefix for the 6rd domain.

For map-e softwires:

Options for MAP-E rules:

name Name of the MAP-E softwire domain name.

br-address Specify the Border Relay (BR) device unicast IPv6 address as the softwire concentrator IPV6 address.

version 3(Optional) Configure version number to distinguish between currently supported version of the Internet draft draft-ietf-softwire-map-03 (expires on July 28, 2013), *Mapping of Address and Port with Encapsulation (MAP)* and the latest available version.

rule Specify the name of Map-E the rule.

v4-reassembly | v6-reassembly (Optional) Enable IPv4 and IPv6 reassembly for MAP-E.

disable-auto-route Disable auto-routes and enable static routes to facilitate ECMP load balancing.

> **NOTE**: When you enable the `disable-auto-route` option, you must configure static routes.

ipv4-prefix Configure rule for IPv4 prefix of the MAP-E domain.

ipv6-prefix Configure rule for IPv6 prefix of the MAP-E domain.

ea-bits-length Configure rule for Embedded Address (EA) length for the MAP-E domain.

- **Range:** 0 through 48

psid-length Configure Port Set ID (PSID) length value for the MAP-E domain.

> **NOTE**:
>
> - If the sum of `v4-prefix-len` and `ea-bits-len` is less than 32, then the `psid-len` must be equal to the difference between 32 and the sum total of `v4-prefix-len` and `ea-bits-len`.

- **Range:** 0 through 16

psid-offset (Optional) Configure PSID offset value for the MAP-E domain.

- **Default:** 4
- **Range:** 0 through 16

mtu-v6 (Optional) Specify the Maximum transmission unit (MTU) for the MAP-E softwire tunnel.

- **Default:** 9192
- **Range:** 1280 through 9192

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2 for Next Gen Services on MX240, MX480 and MX960.

softwires

IN THIS SECTION

- Syntax | 28102
- Hierarchy Level | 28103
- Description | 28103
- Options | 28103
- Required Privilege Level | 28103
- Release Information | 28104

Syntax

```
softwires {
    rule-set name {
        match-direction (input | output);
        rule name {
            then {
                v6rd v6rd;
            }
        }
    }
    softwire-name name {
        ipv4-prefix ipv4-prefix;
        mtu-v4 mtu-v4;
        mtu-v6 mtu-v6;
        post-fragmentation(off | on);
        pre-fragmentation(clear-df-bit | off | on);

        softwire-concentrator softwire-concentrator;
        softwire-type (IPv4-in-IPv6 | v6rd);
        v6rd-prefix v6rd-prefix;
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
```

```
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit logical-systems name security],
[edit security]
[edit services]
```

Description

Configure softwires for IPv6 dual-stack lite (DS-Lite). DS-Lite allows migration to an IPv6 access network without changing end-user software. IPv4 users can continue to access IPv4 internet content using their current hardware, while IPv6 users are able to access IPv6 content.

Options

map-e Name of the map-e rules configuration.

softwire-name Name of the softwire configuration.

traceoptions Traceoptions for Network Security DS-Lite.

mtu-v6 Specify the configured MTU version 6 name for the softwire tunnel.

post-fragmentation You can enable or disable the post-fragmentations to fragment the IPv6 packet. By default the post-fragmentation is off.

pre-fragmentation You can enable or disable the pre-fragmentation or clear the df bit in the IP packet.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 12.1.

softwires (Next Gen Services)

IN THIS SECTION

- Syntax | 28104
- Hierarchy Level | 28105
- Description | 28105
- Required Privilege Level | 28105
- Release Information | 28105

Syntax

```
softwires {
                        rule-set                    name {
        match-direction (input | output);
        rule name {
            then {
                (ds-lite ds-lite | map-e map-e | v6rd v6rd);
            }
        }
    }
                        softwire-name                   name {
    }
    softwire-types {
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace;
```

```
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure softwire feature

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 20.2 for Next Gen Services.

softwires-rule-set (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 28106
- Hierarchy Level | 28106
- Description | 28106
- Required Privilege Level | 28106
- Release Information | 28106

Syntax

```
softwires-rule-set softwire-rule-set-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the softwire rule-set that contains the rule to be used with the service set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

6rd Softwires in Next Gen Services

solicit-router-advertisement-unicast

IN THIS SECTION

- Syntax | 28107
- Hierarchy Level | 28107
- Description | 28107

- Required Privilege Level | 28107
- Release Information | 28107

Syntax

```
solicit-router-advertisement-unicast;
```

Hierarchy Level

```
[edit protocols router-advertisement interface interface-name]
```

Description

Configure devices to send router advertisements as unicast in response to the router solicitation message sent by IPv6 routers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R1 onwards.

RELATED DOCUMENTATION

> https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-routing/config-guide-routing-neighbor-discovery.html

sonet

IN THIS SECTION

- Syntax | 28108
- Hierarchy Level | 28108
- Description | 28108
- Options | 28108
- Required Privilege Level | 28108
- Release Information | 28109

Syntax

```
sonet {
    device-count number;
}
```

Hierarchy Level

```
[edit chassis aggregated-devices]
```

Description

Configure properties for SONET/SDH aggregated devices on the router.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Junos OS for Supporting Aggregated Devices

sophos-engine

IN THIS SECTION

- Syntax | 28109
- Hierarchy Level | 28110
- Description | 28110
- Options | 28110
- Required Privilege Level | 28110
- Release Information | 28111

Syntax

```
sophos-engine {
    file-reputation-action {
        potentially-unwanted-applications (block | log-and-permit | permit);
        unknown (block | log-and-permit | permit);
    }
    server {
        host host;
        port port;
        proxy-profile proxy-profile;
        routing-instance routing-instance;
        source-address source-address;
        ssl-profile ssl-profile;
```

```
                    proxy-profile proxy-profile;
                }
            }
```

Hierarchy Level

```
[edit security utm default-configuration anti-virus],
[edit logical-systems name security utm default-configuration anti-virus],
[edit logical-systems name tenants name security utm default-configuration anti-virus],
[edit tenants name security utm default-configuration anti-virus]
```

Description

Configure the Content Security Sophos antivirus feature.

Options

file-reputation-action	Specify the Sophos file reputation action. File reputation based actions provides more control over security.
potentially-unwanted-applications (block | log-and-permit | permit)	Specify Sophos file reputation actions such as block, log and permit, and permit for the potentially unwanted applications. • **Default:** log-and-permit
proxy-profile *proxy-profile*	Specify the proxy profile name.
unknown (block | log-and-permit | permit)	Specify Sophos file reputation actions such as block, log and permit, and permit for the unknown applications. • **Default:** log-and-permit

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration anti-virus]` hierarchy level is introduced in Junos OS Release 18.2R1.

The `file-reputation-action` and `proxy-profile` options are introduced in Junos OS Release 24.4R1.

source (GRE Encapsulation)

IN THIS SECTION

- Syntax | 28111
- Hierarchy Level | 28111
- Description | 28112
- Default | 28112
- Options | 28112
- Required Privilege Level | 28112
- Release Information | 28112

Syntax

```
source {
    address address;
}
```

Hierarchy Level

```
[edit interfaces fti number unit logical-unit-number tunnel encapsulation ( gre | ipip | udp )]
```

Description

Specify the source address of the tunnel.

Default

Source address is not required. Not configuring a source address indicates that the tunnel will terminate on the destination address.

> **NOTE**: Starting in Junos OS Evolved Release 20.3R1, when you do not configure the "tunnel-termination" on page 30245, configuring the source address for the tunnel encapsulation is mandatory.

Options

address Interface address prefix of the source address.

Required Privilege Level

interface

Release Information

Statement introduced for GRE tunnels in Junos OS Evolved Release 20.2R1.

Statement introduced for IPIP tunnels in Junos OS Evolved Release 20.4R1.

Statement introduced for UDP tunnels in Junos OS Evolved Release 21.1R1.

source (Interfaces)

IN THIS SECTION

- Syntax | **28113**
- Hierarchy Level (EX, NFX, OCX1100 and QFX Series) | **28113**
- Hierarchy Level (M-series and T-series) | **28113**

- Description | 28113
- Default | 28113
- Options | 28113
- Required Privilege Level | 28114
- Release Information | 28114

Syntax

```
source source-address;
```

Hierarchy Level (EX, NFX, OCX1100 and QFX Series)

```
[edit interfaces interface-name unit logical-unit-number tunnel]
```

Hierarchy Level (M-series and T-series)

```
[edit interfaces interface-name unit logical-unit-number tunnel address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
tunnel address]
```

Description

Specify the source address of the tunnel.

Default

If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.

Options

source-address—Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *Tunnel Services Overview*
- multicast-only
- primary (Address on Interface) | **13130**
- Junos OS Services Interfaces Library for Routing Devices

source-active-advertisement

IN THIS SECTION

- Syntax | **28114**
- Hierarchy Level | **28115**
- Description | **28115**
- Required Privilege Level | **28115**
- Release Information | **28115**

Syntax

```
source-active-advertisement {
    dampen minutes;
    min-rate seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system--name protocols mvpn mvpn-mode spt-only],
[edit logical-systems logical-system--name routing-instances instance-name protocols mvpn mvpn-
mode spt-only],
[edit routing-instances protocols mvpn mvpn-mode spt-only],
[edit routing-instances instance-name protocols mvpn mvpn-mode spt-only]
```

Description

Attributes associated with advertising Source-Active A-D routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs

source-address (Services IPSec VPN)

IN THIS SECTION

- Syntax | 28116
- Hierarchy Level | 28116
- Description | 28116
- Options | 28116

- Required Privilege Level | 28116
- Release Information | 28116

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name from]
```

Description

Specify the source address for rule matching.

Options

address—Source IP address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

source-address (Dynamic Tunnel)

IN THIS SECTION

- Syntax | 28117
- Hierarchy Level | 28117
- Description | 28117
- Options | 28117
- Required Privilege Level | 28118
- Release Information | 28118

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options dynamic-tunnels tunnel-name],
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name],
[edit routing-instances routing-instance-name routing-options dynamic-tunnels tunnel-name],
[edit routing-options dynamic-tunnels tunnel-name]
```

Description

Configure the tunnel source address.

Options

address—Name of the source address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Dynamic Tunnels Overview

source-address (IGMP and MLD Snooping)

IN THIS SECTION

- Syntax | 28118
- Hierarchy Level | 28119
- Description | 28119
- Options | 28119
- Required Privilege Level | 28119
- Release Information | 28119

Syntax

```
source-address ip-address;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping proxy],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id proxy],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
snooping proxy],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
snooping vlan vlan-id proxy]
```

Description

Specify the IP address to use as the source for IGMP snooping or MLD snooping reports in proxy mode. Reports are sent with address 0.0.0.0 as the source address unless there is a source address configured. You can also use this statement to configure the source address to use for IGMP snooping or MLD snooping queries.

Options

ip-address—IP address to use as the source for proxy-mode IGMP snooping or MLD snooping reports.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping

source-address

IN THIS SECTION

- Syntax | 28120
- Hierarchy Level | 28120
- Description | 28120
- Required Privilege Level | 28120
- Release Information | 28121

Syntax

```
source-address ip-address;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name from],
[edit firewall family family-name filter filter-name term term-name from ip-version ip-version]
```

Description

IP source address field, which is the address of the source node sending the packet. You cannot specify the `source-address` and `address` match conditions in the same term.

For IPv4, the source-address field is 32 bits in length. The filter description syntax supports either a mask value that can be noncontiguous, such as 10.0.0.10/255.0.0.255, or prefix notation such as 10.0.0.0/8. Simple filters do not support noncontiguous mask values.

For IPv6, the source-address field is 128 bits in length. The filter description syntax supports the text representations for IPv6 addresses that are described in RFC 2373, *IP Version 6 Addressing Architecture*.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1R1.

RELATED DOCUMENTATION

family (Firewall) | 6867

family (Firewall Filter) | 6869

Firewall Filter Match Conditions Based on Address Fields

Firewall Filter Match Conditions for IPv4 Traffic

Firewall Filter Match Conditions for IPv6 Traffic

Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic

Guidelines for Configuring Simple Filters

source-address (Content Security)

IN THIS SECTION

- Syntax | 28121
- Hierarchy Level | 28122
- Description | 28122
- Options | 28122
- Required Privilege Level | 28122
- Release Information | 28122

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit security utm default-configuration web-filtering juniper-enhanced server]
[edit security utm default-configuration web-filtering websense-redirect server]
[edit security utm default-configuration anti-virus sophos-engine server]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name server]
```

Description

You can configure source address for Enhanced Web Filtering (EWF) cloud service, websense redirect policy service, antivirus scan service, and antispam scan service. Configuring a source address for these Content Security services enhance the network disaster recovery. Since the antivirus and antispam scan services shares the same Domain Name Service (DNS) to forward the scan request, they share the same source address configuration under Sophos antivirus engine server.

The source address configuration is optional. If you do not configure the source address, the Content Security find the route and get the source address.

Options

address	Source address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

source-address (AACL)

IN THIS SECTION

- Syntax | 28123
- Hierarchy Level | 28123
- Description | 28123
- Options | 28123
- Required Privilege Level | 28123
- Release Information | 28124

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit services aacl rule rule-name term term-name from]
```

Description

Specify the source address for rule matching.

Options

address—Source IPv4 or IPv6 address or prefix value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

IPv6 support introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Match Conditions in AACL Rules

source-address (Access LDAP)

IN THIS SECTION

- Syntax | 28124
- Hierarchy Level | 28124
- Description | 28125
- Options | 28125
- Required Privilege Level | 28125
- Release Information | 28125

Syntax

```
source-address source-address;
```

Hierarchy Level

```
[edit access ldap-server server-address],
[edit access profile profile-name ldap-server server-address]
```

Description

Configure a source address for each configured LDAP server. Each LDAP request sent to a LDAP server uses the specified source address.

Options

source-address—Valid IP address configured on one of the device interfaces.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Example: Configuring RADIUS and LDAP User Authentication

source-address (Access RADIUS)

IN THIS SECTION

- Syntax | 28126
- Hierarchy Level | 28126
- Description | 28126
- Options | 28126
- Required Privilege Level | 28126
- Release Information | 28126

Syntax

```
source-address source-address;
```

Hierarchy Level

```
[edit access radius-server server-address],
[edit access profile profile-name radius-server server-address]
```

Description

Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address.

Options

source-address—Valid IP address configured on one of the device interfaces.

Required Privilege Level

secret—To view this statement in the configuration.

secret-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

Firewall User Authentication Overview

source-address-excluded

IN THIS SECTION

- Syntax | 28127
- Hierarchy Level | 28127
- Description | 28127
- Required Privilege Level | 28127
- Release Information | 28127

Syntax

```
source-address-excluded;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy  policy-name match]
```

Description

Exclude the source address(es) from the policy.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Security Policies Overview

source-address-filter

IN THIS SECTION

- Syntax | 28128
- Hierarchy Level | 28128
- Description | 28128
- Options | 28129
- Required Privilege Level | 28129
- Release Information | 28129

Syntax

```
source-address-filter {
    mac-address;
}
```

Hierarchy Level

```
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name fastether-options],
[edit interfaces interface-name gigether-options]
```

Description

For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), specify the MAC addresses from which the interface can receive packets. For this

statement to have any effect, you must include the `source-filtering` statement in the configuration to enable source address filtering.

Options

mac-address—MAC address filter. You can specify the MAC address as *nn*:*nn*:*nn*:*nn*:*nn*:*nn* or *nnnn*.*nnnn*.*nnnn*, where *n* is a decimal digit. To specify more than one address, include multiple *mac-address* options in the `source-address-filter` statement.

If you enable the VRRP on a Fast Ethernet or Gigabit Ethernet interface, as described in VRRP and VRRP for IPv6 Overview, and if you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the `source-address-filter` statement. MAC addresses ranging from `00:00:5e:00:01:00` through `00:00:5e:00:01:ff` are reserved for VRRP, as defined in RFC 3768, *Virtual Router Redundancy Protocol*. When you configure the VRRP group, the group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

On untagged Gigabit Ethernet interfaces, you should not configure the `source-address-filter` statement and the `accept-source-mac` statement simultaneously. On tagged Gigabit Ethernet interfaces, you should not configure the `source-address-filter` statement and the `accept-source-mac` statement with an identical MAC address specified in both filters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring MAC Address Filtering for Ethernet Interfaces

Configuring MAC Address Filtering on PTX Series Packet Transport Routers

source-filtering | **28208**

source-address-filter (Interfaces)

IN THIS SECTION

- Syntax | 28130
- Hierarchy Level | 28130
- Description | 28130
- Options | 28131
- Required Privilege Level | 28131
- Release Information | 28131

Syntax

```
source-address-filter mac-address;
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options]
```

Description

For redundant Ethernet interfaces, specify the MAC addresses from which the interface can receive packets. For this statement to have any effect, you must include the `source-filtering` statement in the configuration to enable source address filtering.

Be sure to update the MAC address if the remote Ethernet card is replaced. Replacing the interface card changes the MAC address. Otherwise, the interface cannot receive packets from the new card.

NOTE:

- Software based MAC limiting is supported on SRX300, SRX320, and SRX340 devices.

  A maximum of 32 devices are supported per device.

Options

mac-address —MAC address filter. You can specify the MAC address as six hexadecimal bytes in one of the following formats: *nn:nn:nn:nn:nn:nn* (for example, `00:11:22:33:44:55`) or *nnnn:nnnn:nnnn* (for example, `0011.2233.4455`). You can configure up to 64 source addresses. To specify more than one address, include multiple *mac-address* options in the `source-address-filter` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

source-address-filter-list

IN THIS SECTION

- source-address-filter-list (Create List) | **28132**
- Hierarchy Level | **28132**
- source-address-filter-list (Use list) | **28132**
- Hierarchy Level | **28132**
- Description | **28132**
- Required Privilege Level | **28133**
- Release Information | **28133**

source-address-filter-list (Create List)

```
source-address-filter-list source-address-filter-list-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options],
[edit policy-options]
```

source-address-filter-list (Use list)

```
source-address-filter-list source-address-filter-list-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-statement policy-statement-name term term-name
from],
[edit policy-options policy-statement policy-statement-name term term-name from]
```

Description

The source address filter list is a user configured list of individual source address filters, typically used to match an incoming route address to unicast source addresses in Multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments, that you create at the `[edit policy-options]` hierarchy level. Each item in the list consists of a complete source address filter statement, made up of a source-prefix address, a match-type, and an optional action.

For example:

```
[edit]
user@router# show policy-options source-address-filter-list saf-test-list
203.0.113.0/26 exact;
192.0.2.0/24 longer accept;
198.51.100.8/29 exact reject;
```

Once configured, `source-address-filter-list` is used by referencing its *`source-address-filter-list-name`* in a `policy-statement` at the `[edit policy-options policy-statement` *`policy-statement-name`* `term` *`term-name`* `from]` hierarchy level. Source address filter lists can be used in conjunction with other `source-address-filter` statements.

For example:

```
[edit]
user@router# show policy-options policy-statement test-saf-list-statement
from {
source-address-filter 198.51.100.16/29 exact accept;
source-address-filter-list saf-test-list;
}
then reject;
```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement first introduced in Junos OS Release 16.1 on M Series, MX Series, and T Series.

source-address (Forwarding Options)

IN THIS SECTION

- Syntax | 28134
- Hierarchy Level | 28134
- Description | 28134
- Options | 28134
- Required Privilege Level | 28134
- Release Information | 28134

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit forwarding-options accounting name outputinterface interface-name],
[edit forwarding-options monitoring namefamilyfamily inet output interface interface-name],
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls | vpls) output
interface interface-name],
[edit forwarding-options sampling family (inet |inet6 |mpls) output interface interface-name],
[edit forwarding-options sampling  instance instance-name family inet output inline-jflow]
```

Description

Specify the source address for monitored packets.

Options

address—Interface source address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for IPv6 addresses for the `source-address` *address* statement within sampling instances introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Configuring Discard Accounting

Configuring Flow Monitoring

source-address-giaddr

IN THIS SECTION

- Syntax | 28135
- Hierarchy Level | 28135
- Description | 28135
- Required Privilege Level | 28136
- Release Information | 28136

Syntax

```
source-address-giaddr;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface interface-name]
```

Description

Configure the gateway IP address (giaddr) as the source IP address of the switch for relayed DHCP packets when the switch is used as the DHCP relay agent.

When this statement is entered in the [**edit forwarding-options helpers bootp**] hierarchy, the gateway IP address is configured as the source IP address of the switch for relayed DHCP packets exiting all interfaces on the switch.

When this statement is entered in the [**edit forwarding-options helpers bootp interface** *interface-name*] hierarchy, the gateway IP address is configured as the source IP address of the switch for relayed DHCP packets exiting the specified interface of the switch.

In Junos OS Release 10.1 for EX Series switches and later releases, the IP address of the interface that the DHCP packet exits on the switch acting as a DHCP relay agent is used as the source IP address for relayed DHCP packets by default.

In Junos OS Releases 9.6 and 10.0 for EX Series switches, the gateway IP address of the switch is always used as the source IP address for relayed DHCP packets when the switch is used as the DHCP relay agent.

In Junos OS Releases 9.3 through 9.5 for EX Series switches, the IP address of the interface that the DHCP packet exits on the switch acting as a DHCP relay agent is always used as the source IP address for relayed DHCP packets.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

DHCP/BOOTP Relay for Switches Overview

source-address (IDS MS-DPC)

IN THIS SECTION

- Syntax | 28137
- Hierarchy Level | 28137
- Description | 28137
- Options | 28137
- Required Privilege Level | 28137
- Release Information | 28137

Syntax

```
source-address (address | any-unicast) <except>;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name from]
```

Description

Specify the source address for rule matching when using the MS-DPC.

Options

address—Source IPv4 or IPv6 address or prefix value.

`any-unicast`—Any unicast packet.

`except`—(Optional) Exempt the specified address, prefix, or unicast packets from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Match Conditions in IDS Rules

source-address (LRF Profile)

IN THIS SECTION

- Syntax | 28138
- Hierarchy Level | 28138
- Description | 28138
- Options | 28138
- Required Privilege Level | 28138
- Release Information | 28139

Syntax

```
source-address source-address;
```

Hierarchy Level

```
[edit services lrf profile profile-name collector collector-name]
```

Description

Configure the source address to be used when exporting data to the collector.

Options

source-address	IP address to be used as the source address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

source-address-name (NAT Next Gen Services)

IN THIS SECTION

- Syntax | 28139
- Hierarchy Level | 28139
- Description | 28140
- Required Privilege Level | 28140
- Release Information | 28140

Syntax

```
source-address-name address-name;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name match],
[edit services nat source rule-set rule-set rule rule-name match]
```

Description

Specify the name of the range of source addresses that the packet must match for the NAT rule to take effect. The range of addresses is configured with the `address` statement at the `[edit services address-book global]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

source-address-name (Security Destination NAT)

IN THIS SECTION

- Syntax | 28140
- Hierarchy Level | 28141
- Description | 28141
- Options | 28141
- Required Privilege Level | 28141
- Release Information | 28141

Syntax

```
source-address-name [address-name];
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name rule rule-name match]
```

Description

Specify a source address name to match the rule. You can configure multiple address names.

Options

address-name—Source address name.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

source-address-name (Security Source NAT)

IN THIS SECTION

- Syntax | 28142
- Hierarchy Level | 28142
- Description | 28142
- Options | 28142
- Required Privilege Level | 28142
- Release Information | 28142

Syntax

```
source-address-name [address-name];
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name match]
```

Description

Specify a source address name to match the rule. You can configure multiple address names.

Options

address-name—Source address name.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

source-address-name (Security Static NAT Rule Set)

IN THIS SECTION

- Syntax | 28143
- Hierarchy Level | 28143
- Description | 28143
- Options | 28143

- Required Privilege Level | **28143**
- Release Information | **28143**

Syntax

```
source-address-name [address-name];
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name match]
```

Description

Specify a source address name to match the rule. Up to 8 address names are supported.

Options

address-name—Source address name.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

source-address (NAT Next Gen Services)

IN THIS SECTION

- Syntax | 28144
- Hierarchy Level | 28144
- Description | 28144
- Options | 28144
- Required Privilege Level | 28144
- Release Information | 28145

Syntax

```
source-address (address | any-unicast);
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name match],
[edit services nat source rule-set rule-set rule rule-name match]
```

Description

Specify the source address that the packet must match for the NAT rule to take effect.

Options

address	A specific address that must be matched.
any-unicast	Any unicast source address results in a match.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

source-address (Next Gen Services Service-Set Remote System Logging)

IN THIS SECTION

- Syntax | 28145
- Hierarchy Level | 28145
- Description | 28145
- Required Privilege Level | 28146
- Release Information | 28146

Syntax

```
source-address address;
```

Hierarchy Level

```
edit services service-set name syslog
```

Description

Specify the IP address of the source for Next Gen Services system log messages.

BEST PRACTICE: The syslog source address can be any arbitrary IP address. It does not have to be an IP address that is assigned to the device. Rather, this IP address is used on

the syslog collector to identify the syslog source. The best practice is to configure the source address as the IP address of the interface that the traffic is sent out on.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

stream (Next Gen Services Service-Set Remote System Logging) | **28631**

source-address (PCP)

IN THIS SECTION

- Syntax | **28147**
- Hierarchy Level | **28147**
- Description | **28147**
- Options | **28147**
- Required Privilege Level | **28147**
- Release Information | **28147**

Syntax

```
source-address address <except>;
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name from]
```

Description

Specify the source address that must be matched for the PCP rule. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

address Destination address or prefix value.

except (Optional) Prevent the specified address or prefix from matching the PCP rule.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

source-address (RADIUS, System Logging, or TACACS+)

IN THIS SECTION

- Syntax | 28148
- Hierarchy Level | 28148
- Description | 28148
- Options | 28148
- Required Privilege Level | 28149
- Release Information | 28149

Syntax

```
source-address source-address;
```

Hierarchy Level

```
[edit system accounting destination radius server server-address],
[edit system accounting destination tacplus server server-address],
[edit system radius-server server-address],
[edit system syslog],
[edit system tacplus-server server-address]
```

Description

Specify a source address for each configured TACACS+ server, RADIUS server, or the source address to record in system log messages that are directed to a remote machine.

Options

source-address—Valid IP address configured on one of the switch interfaces. For system logging, the address is recorded as the message source in messages sent to the remote machines specified in all `host hostname` statements at the `[edit system syslog]` hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring RADIUS Authentication (QFX Series or OCX Series)

Synchronizing and Coordinating Time Distribution Using NTP

Specifying an Alternative Source Address for System Log Messages Directed to a Remote Destination

source-address-range

IN THIS SECTION

- Syntax | 28149
- Hierarchy Level | 28150
- Description | 28150
- Options | 28150
- Required Privilege Level | 28150
- Release Information | 28150

Syntax

```
source-address-range low minimum-value high maximum-value;
```

Hierarchy Level

```
[edit services aacl rule rule-name term term-name from]
```

Description

Specify the source address range for rule matching.

Options

minimum-value—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

IPv6 support introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Match Conditions in AACL Rules

source-address-range (IDS MS-DPC)

IN THIS SECTION

- Syntax | 28151
- Hierarchy Level | 28151

- Description | 28151
- Options | 28151
- Required Privilege Level | 28151
- Release Information | 28151

Syntax

```
source-address-range low minimum-value high maximum-value <except>;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name from]
```

Description

Specify the source address range for rule matching when using the MS-DPC.

Options

minimum-value—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

`except`—(Optional) Exempt the specified address range from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

minimum-value and *maximum-value* options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Configuring Match Conditions in IDS Rules*

source-address-range (PCP)

IN THIS SECTION

- Syntax | 28152
- Hierarchy Level | 28152
- Description | 28152
- Options | 28153
- Required Privilege Level | 28153
- Release Information | 28153

Syntax

```
source-address-range high maximum-value low minimum-value <except>;
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name from]
```

Description

Specify the source address range that must be matched for the PCP rule. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

maximum-value	Upper boundary for the address range.
minimum-value	Lower boundary for the address range.
except	(Optional) Prevent the specified address range from matching the PCP rule.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

source-address-range (Services NAT)

IN THIS SECTION

- Syntax | 28154
- Hierarchy Level | 28154
- Description | 28154
- Options | 28154
- Required Privilege Level | 28154
- Release Information | 28154

Syntax

```
source-address-range low minimum-value high maximum-value <except>;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name from]
```

Description

Specify the source address range for rule matching.

Options

minimum-value—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

except—(Optional) Prevent the specified address range from being translated.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

minimum-value and ***maximum-value*** options enhanced to support IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

source-address-range (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 28155
- Hierarchy Level | 28155
- Description | 28155
- Options | 28155
- Required Privilege Level | 28155
- Release Information | 28156

Syntax

```
source-address-range low minimum-value high maximum-value <except>;
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name term term-name from]
```

Description

Source address range for rule matching.

Options

minimum-value—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

`except`—(Optional) Exclude the specified address, prefix, or unicast packets from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

minimum-value and *maximum-value* options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Match Conditions in Stateful Firewall Rules

source-address (Routing Options)

IN THIS SECTION

- Syntax | 28156
- Hierarchy Level | 28156
- Description | 28157
- Options | 28157
- Required Privilege Level | 28157
- Release Information | 28157

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options
dynamic-tunnels tunnel-name],
```

```
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name],
[edit routing-instances routing-instance-name routing-options dynamic-tunnels tunnel-name],
[edit routing-options dynamic-tunnels tunnel-name]
```

Description

Specify the source address for the generic routing encapsulation (GRE) tunnels. The source address specifies the address used as the source for the local tunnel endpoint. This address can be any local address on the router, typically the router ID or the loopback address.

Options

address—Name of the source address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring GRE Tunnels for Layer 3 VPNs

source-address (RPM)

IN THIS SECTION

- Syntax | **28158**
- Junos OS Hierarchy Level | **28158**
- Junos OS Evolved Hierarchy Level | **28158**

- Description | 28158
- Options | 28159
- Required Privilege Level | 28159
- Release Information | 28159

Syntax

```
source-address address;
```

Junos OS Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the source IP address used for probes. If the source IP address is not one of the router's or switch's assigned addresses, the packet uses the outgoing interface's address as its source. Prior to Junos OS Evolved Release 23.1R1, only IPv4 addresses are supported for Junos OS Evolved.

The following addresses cannot be used for the source IP address used for probes:

- 0.0.0.0
- 127.0.0.0/8 (loopback)
- 224.0.0.0/4 (multicast)
- 255.255.255.255 (broadcast)
- ::/128
- ::1/128 (loopback)

- fc00::/7 (unique-local)
- 2001:db8::/32 (documentation prefix)
- 2002::/16 (6to4)
- 5f00::/8 (6bone)
- 2001:10::/28 (ORCHID)
- 2001::/32 (Teredo)
- ::/0 (default route)
- ff00::/8 (multicast)
- :ffff:0:0/96 (IPv4-mapped addresses)

Options

address—Valid IP address. You cannot configure this statement if the address configured for the `target` statement is an IPv6 link-local address.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

Support for IPv6 addresses added in Junos OS Evolved 23.1R1.

RELATED DOCUMENTATION

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

source-address (Security Destination NAT)

IN THIS SECTION

- Syntax | 28160
- Hierarchy Level | 28160
- Description | 28160
- Options | 28160
- Required Privilege Level | 28160
- Release Information | 28161

Syntax

```
source-address [ip-address];
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name rule rule-name match]
```

Description

Specify source address to match the rule. You can configure multiple addresses or subnets.

Options

ip-address —Source address or a subnet.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

source-address (Security IDP)

IN THIS SECTION

- Syntax | 28161
- Hierarchy Level | 28161
- Description | 28161
- Options | 28161
- Required Privilege Level | 28162

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit security idp security-package]
```

Description

Sets the source address to be used for sending download requests.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

source-address (Security IDP Policy)

IN THIS SECTION

- Syntax | 28162
- Hierarchy Level | 28162
- Description | 28162
- Options | 28163
- Required Privilege Level | 28163
- Release Information | 28163

Syntax

```
source-address ([address-name] | any | any-ipv4 | any-ipv6);
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match]
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match]
```

Description

Specify a source IP address or IP address set object to be used as the match source address object. The default value is any.

Options

- *address-name*—IP address or IP address set object.
- *any*—Specify any IPv4 or IPv6 address.
- *any-ipv4*—Specify any IPv4 address.
- *any-ipv6*—Specify any IPv6 address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

source-address (Security IDP Sensor Configuration)

IN THIS SECTION

- Syntax | 28163
- Hierarchy Level | 28164
- Description | 28164
- Required Privilege Level | 28164
- Release Information | 28164

Syntax

```
source-address ip-address;
```

Hierarchy Level

```
[edit security idp sensor-configuration packet-log]
```

Description

Configure the source IP address for the carrier UDP packet.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

source-address (Security Policies)

IN THIS SECTION

- Syntax | 28164
- Hierarchy Level | 28165
- Description | 28165
- Options | 28165
- Required Privilege Level | 28165
- Release Information | 28165

Syntax

```
source-address {
    [address];
```

```
    any;
    any-ipv4;
    any-ipv6;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name match]
```

```
[edit security policies global policy policy-name match]
```

Description

Define the matching criteria. You can specify one or more IP addresses, address sets, or wildcard addresses. You can specify wildcards `any`, `any-ipv4`, or `any-ipv6`.

Options

address—IP addresses, address sets, or wildcard addresses (represented as A.B.C.D/wildcard-mask). You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

source-address (Security Source NAT)

IN THIS SECTION

- Syntax | 28166
- Hierarchy Level | 28166
- Description | 28166
- Options | 28166
- Required Privilege Level | 28166
- Release Information | 28167

Syntax

```
source-address [ip-address];
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name match]
```

Description

Specify source address to match the rule. You can configure multiple addresses or subnets.

Options

ip-address—Source address or a subnet.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

source-address (Security Static NAT Rule Set)

IN THIS SECTION

- Syntax | 28167
- Hierarchy Level | 28167
- Description | 28167
- Options | 28167
- Required Privilege Level | 28168
- Release Information | 28168

Syntax

```
source-address [ip-address];
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name match]
```

Description

Specify the source address to match the rule. Up to 8 addresses are supported.

Options

ip-address —Source address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

source-address (Service Sets)

IN THIS SECTION

- Syntax | 28168
- Hierarchy Level | 28168
- Description | 28169
- Options | 28169
- Required Privilege Level | 28169
- Release Information | 28169

Syntax

```
source-address source-address
```

Hierarchy Level

```
[edit services service-set service-set-name syslog host hostname]
```

Description

Specify a source address to record in system log messages that are directed to a remote machine specified in the *hostname* statement.

> **NOTE**: The supported interfaces are ms, rms, and mams interfaces. If you do not specify the interface parameter, the command loops on all supported interfaces.

Options

source-address	A valid IP address, which is recorded as the message source in messages sent to the remote machines specified in the `host` *hostname* statement

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

Configuring System Logging for Service Sets

host (service-set) | **8169**

service-set (Services) | **16750**

source-address (Services CoS)

IN THIS SECTION

- Syntax | **28170**

- Hierarchy Level | 28170
- Description | 28170
- Options | 28170
- Required Privilege Level | 28170
- Release Information | 28170

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit services cos rule rule-name term term-name from]
```

Description

Source address for rule matching.

Options

address—Source IPv4 or IPv6 address or prefix value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

Configuring Match Conditions In CoS Rules

source-address (Services NAT)

IN THIS SECTION

- Syntax | 28171
- Hierarchy Level | 28171
- Description | 28171
- Options | 28171
- Required Privilege Level | 28172
- Release Information | 28172

Syntax

```
source-address (address | any-unicast) <except>;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name from]
```

Description

Specify the source address for rule matching.

Options

address—Source IPv4 or IPv6 address or prefix value.

any-unicast—Any unicast packet.

except—(Optional) Prevent the specified address or unicast packets from being translated.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

any-unicast and **except** options introduced in Junos OS Release 7.6.

address option enhanced to support IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

source-address (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 28172
- Hierarchy Level | 28173
- Description | 28173
- Options | 28173
- Required Privilege Level | 28173
- Release Information | 28173

Syntax

```
source-address (address | any-ipv4 | any-ipv6 | any-unicast ) <except>;
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name term term-name from]
```

Description

Source address for rule matching.

Options

address—Source IPv4 or IPv6 address or prefix value.

`any-ipv4`—Any IPv4 packet.

`any-ipv6`—Any IPv6 packet.

`any-unicast`—Any unicast packet.

`except`—(Optional) Exclude the specified address, prefix, IPv4, IPv6, or unicast packets from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`any-unicast` and `except` options introduced in Junos OS Release 7.6.

address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Match Conditions in Stateful Firewall Rules

source-address (SNMP)

IN THIS SECTION

- Syntax | 28174
- Hierarchy Level | 28174
- Description | 28174
- Options | 28174
- Required Privilege Level | 28175
- Release Information | 28175

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit snmp trap-options]
```

Description

Set the source address of every SNMP trap packet sent by this router to a single address regardless of the outgoing interface. If the source address is not specified, the default is to use the address of the outgoing interface as the source address.

Options

address—Source address of SNMP traps. You can configure the source address of trap packets two ways: `lo0` or a valid IPv4 address configured on one of the router interfaces. The value `lo0` indicates that the source address of all SNMP trap packets is set to the lowest loopback address configured at interface `lo0`.

- **Default:** Disabled. (The source address is the address of the outgoing interface.)

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

source-address (SRC Software)

IN THIS SECTION

- Syntax | 28175
- Hierarchy Level | 28175
- Description | 28175
- Options | 28176
- Required Privilege Level | 28176
- Release Information | 28176

Syntax

```
source-address source-address;
```

Hierarchy Level

```
[edit system services service-deployment]
```

Description

Enable Junos OS to work with the Session and Resource Control (SRC) software.

Options

source-address— Local IPv4 address to be used as source address for traffic to the SRC server. The source address restricts traffic within the out-of-band network.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Junos OS to Work with SRC Software

source-address (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 28176
- Hierarchy Level | 28177
- Description | 28177
- Options | 28177
- Required Privilege Level | 28177
- Release Information | 28177

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit services radius-flow-tap policy policy-name inet drop-policy rule-name  from],
[edit services radius-flow-tap policy policy-name inet6 drop-policy rule-name from]
```

Description

Specify source IP address or prefix value from which to inherit configuration data for radius-flow-tap policy rule mapping.

Options

address— IPv4 or IPv6 address for the radius-flow-tap policy.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Subscriber Secure Policy Overview

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

source-address (System Logging)

IN THIS SECTION

- Syntax | 28178
- Hierarchy Level | 28178

- Description | 28178
- Options | 28178
- Required Privilege Level | 28178
- Release Information | 28179

Syntax

```
source-address source-address <routing-instance routing-instance-name>;
```

Hierarchy Level

```
[edit system syslog]
```

Description

Specify a source address to record in system log messages that are directed to a remote machine.

Options

source-address—A valid IP address configured on one of the device interfaces. For system logging, the address is recorded as the message source in messages sent to the remote machines specified in all `host hostname` statements at the `[edit system syslog]` hierarchy level, but not for messages directed to the other Routing Engine.

`routing-instance routing-instance-name`—(Optional) The routing instance name in which the source address is defined.

- **Default:** The primary address of the interface

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

source-address (TWAMP)

IN THIS SECTION

- Syntax | 28179
- Hierarchy Level | 28179
- Description | 28179
- Default | 28180
- Options | 28180
- Required Privilege Level | 28180
- Release Information | 28181

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit services rpm twamp client control-connection connection-name test-session session-name
```

Description

Specify the source IP address used for probes. If the source IP address is not one of the router's or switch's assigned addresses, the packet uses the outgoing interface's address as its source.

The following IPv4 addresses cannot be used as the source IP address used for probes:

- 0.0.0.0

- 127.0.0.0/8 (loopback)
- 224.0.0.0/4 (multicast)
- 255.255.255.255 (broadcast)

The following IPv6 addresses cannot be used as the source IP address used for probes:

- ::/128
- ::1/128 (loopback)
- fc00::/7 (unique-local)
- 2001:db8::/32 (documentation prefix)
- 2002::/16 (6to4)
- 5f00::/8 (6bone)
- 2001:10::/28 (ORCHID)
- 2001::/32 (Teredo)
- ::/0 (default route)
- ff00::/8 (multicast)
- :ffff:0:0/96 (IPv4-mapped addresses)

Default

The source address is selected based on the available default route.

Options

address Valid IPv4 or IPv6 address. You can only specify an IPv6 address as the source address if the source address corresponds to a TWAMP Light IPv6 link-local target address.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

source-addresses

IN THIS SECTION

- Syntax | 28181
- Hierarchy Level | 28181
- Description | 28181
- Options | 28181
- Required Privilege Level | 28182
- Release Information | 28182

Syntax

```
source-addresses [ addresses ];
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

List the IP addresses from which the control source can send control protocol requests to the Juniper Networks router.

Options

address—Allowed IP source address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Control Source

source (Application Identification)

IN THIS SECTION

- Syntax | 28182
- Hierarchy Level | 28182
- Description | 28183
- Options | 28183
- Required Privilege Level | 28183
- Release Information | 28183

Syntax

```
source ip ip-address-prefix;
```

Hierarchy Level

```
[edit services application-identification application application-name address-mapping]
```

Description

Specify the source IP address for address mapping-based application identification.

Options

ip-address-prefix	IP address and prefix for matching.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

Application Identification Overview

source-attachment-identifier (Protocols VPWS)

IN THIS SECTION

- Syntax | 28184
- Hierarchy Level | 28184
- Description | 28184
- Options | 28184
- Required Privilege Level | 28184
- Release Information | 28185

Syntax

```
source-attachment-identifier identifier;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn site site-name],
[edit routing-instances routing-instance-name protocols l2vpn site site-name]
```

Description

For FEC 129, specify the VPWS source attachment identifier. The point-to-point nature of VPWS requires that you specify the source access individual identifier (SAII) and the target access individual identifier (TAII). This SAII-TAII pair defines a unique pseudowire between two PE devices.

Auto-discovery routes are used by BGP to allow auto-discovery of remote source access individual identifiers (SAIIs) (the sources of the point-to-point pseudowires). One auto-discovery route is advertised for each source attachment identifier (SAI) in the instance or mesh group.

The SAII is specified with the `source-attachment-identifier` statement within the FEC 129 VPWS routing instance. You configure the source attachment identifier and the interfaces to associate with that source attachment identifier. Under each interface, you can configure the TAII with the `target-attachment-identifier` statement. If the configured target identifier matches a source identifier advertised by a remote PE device by way of a BGP auto-discovery message, the pseudowire between that source-target pair is signaled. If there is no match between an advertised source identifier and the configured target identifier, the pseudowire is not established.

Options

identifier—The numerical identifier for the Layer 2 VPN site.

- **Range:** 1 through 4,292,967,295

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

> *Example: Configuring FEC 129 BGP Autodiscovery for VPWS*
>
> target-attachment-identifier (Protocols VPWS) | **28947**

source-bmac

IN THIS SECTION

- Syntax | **28185**
- Hierarchy Level | **28185**
- Description | **28186**
- Options | **28186**
- Required Privilege Level | **28186**
- Release Information | **28186**

Syntax

```
source-bmac <mac-address> <length> [auto];
```

Hierarchy Level

```
[edit interfaces interface-name unit unit esi],
[edit routing-instances routing-instance-name pbb-options],
[edit routing-instances routing-instance-name pbb-options service-groups service-group-name],
[edit routing-instances routing-instance-name pbb-options service-groups service-group-name isid
isid-name]
```

Description

Set a shared source Provider Backbone Bridging (PBB) MAC (B-MAC) address for all service identifiers (ISIDs) under a specified routing instance. If a source B-MAC address is derived automatically from the aggregated Ethernet system ID, you must specify the `auto` keyword.

For PBB-Ethernet VPN (EVPN) functionality, one service group (a group of ISIDs) can correspond to a single-homed, single-Active, or Active-Active scenario. You cannot mix and match service groups. For a shared B-MAC model, you configure one source B-MAC per service group using the `routing-instances` *`routing-instance-name`* `pbb-options service-groups` *`service-group-name`* hierarchy.

To specify a source B-MAC per ISID, use the `routing-instances` *`routing-instance-name`* `pbb-options service-groups` *`service-group-name`* `isid` *`isid-name`* hierarchy. If you do not specify the source B-MAC per ISID, then the B-MAC inherits from the B-MAC configuration located under the `pbb-options` setting.

The support of the `source-bmac` statement at the `[edit interfaces` *`interface-name`* `unit` *`unit`* `esi]` hierarchy level allows configuring of the Ethernet Segment Identifier (ESI) and source B-MAC address per logical interface, in addition to the existing support for physical devices.

Options

Length

Length of the B-MAC address.

- **Default:** 48B

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support at the `[edit interfaces` *`interface-name`* `unit` *`unit`* `esi]` hierarchy level introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

routing-instances (Multiple Routing Entities) | **15976**

source (Bridge Domains)

IN THIS SECTION

- Syntax | 28187
- Hierarchy Level | 28187
- Description | 28187
- Options | 28187
- Required Privilege Level | 28188
- Release Information | 28188

Syntax

```
source ip-address;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name static
group],
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name static
group],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
snooping interface interface-name static group],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols vlan
vlan-id igmp-snooping interface interface-name static group]
```

Description

Statically define multicast group source addresses on an interface.

Options

ip-address—IP address to use as the source for the group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping

source (Chassis Synchronization)

IN THIS SECTION

- Syntax | 28188
- Hierarchy Level | 28189
- Description | 28189
- Options | 28189
- Required Privilege Level | 28189
- Release Information | 28189

Syntax

```
source {
    (bits | gps) {
        priority number;
        quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
        request request (force-switch | lockout);
    }
    interfaces interface-name {
```

```
        priority number;
        quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
        request request (force-switch | lockout);
        wait-to-restore minutes;
    }
}
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Configure a clock source used by the clock selection process.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

> *External Clock Synchronization Overview for ACX Series Routers*
>
> *Configuring External Clock Synchronization for ACX Series Routers*
>
> *Synchronous Ethernet Overview*
>
> *show chassis synchronization*

source-checking

IN THIS SECTION

- Syntax | 28190
- Hierarchy Level | 28190
- Description | 28190
- Required Privilege Level | 28191
- Release Information | 28191

Syntax

```
source-checking;
```

Hierarchy Level

```
[edit forwarding-options family inet6]
```

Description

(MX Series 5G Universal Routing Platforms Only) Discard IPv6 packets when the source address type is unspecified, loopback, multicast, or link-local unicast.

RFC 4291, *IP Version 6 Addressing Architecture*, refers to four address types that require special treatment when they are used as source addresses. The four address types are:

- Unspecified--The IPv6 unspecified address is 0:0:0:0:0:0:0:0, which can be notated as ::/128.
- Loopack--The IPv6 loopback address is 0:0:0:0:0:0:0:1, which can be notated as ::1/128.
- Multicast--IPv6 multicast addresses use the prefix FF00::/8.
- Link-Local--IPv6 link-local addresses have the prefix FE80::/10.

The loopback and multicast addresses must never be used as a source address in IPv6 packets. The unspecified and link-local addresses can be used as source addresses but routers must never forward

packets that have these addresses as source addresses. Typically, packets that contain unspecified or link-local addresses as source addresses are delivered to the local host. If the destination is not the local host, then the packet must not be forwarded. Configuring this statement filters or discards IPv6 packets of these four address types.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Applying Forwarding Table Filters

source-class-usage

IN THIS SECTION

- Syntax | 28192
- Hierarchy Level | 28192
- Description | 28192
- Options | 28192
- Required Privilege Level | 28192
- Release Information | 28192

Syntax

```
source-class-usage {
    direction;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet accounting],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet accounting],
[edit routing-instances routing-instance-name vrf-table-label]
```

Description

Enable packet counters on an interface that count packets that arrive from specific prefixes on the provider core router and are destined for specific prefixes on the customer edge router.

Options

direction can be one of the following:

`input`—Configure at least one expected ingress point.

`output`—Configure at least one expected egress point.

`input output`—On a single interface, configure at least one expected ingress point and one expect egress point.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `vrf-table-label` statement added in Junos OS Release 9.3.

RELATED DOCUMENTATION

- Enabling Source Class and Destination Class Usage
- accounting
- destination-class-usage
- Junos OS Services Interfaces Library for Routing Devices
- vrf-table-label | 30966

source-classes

IN THIS SECTION

- Syntax | 28193
- Hierarchy Level | 28193
- Description | 28193
- Options | 28194
- Required Privilege Level | 28194
- Release Information | 28194

Syntax

```
source-classes {
    source-class-name;
}
```

Hierarchy Level

```
[edit accounting-options class-usage-profile profile-name]
```

Description

Specify the source classes for which statistics are collected.

Options

source-class-name—Name of the source class to include in the class usage profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Accounting Options, Source Class Usage and Destination Class Usage Options

source (Commit Scripts)

IN THIS SECTION

- Syntax | 28194
- Hierarchy Level | 28195
- Description | 28195
- Options | 28195
- Required Privilege Level | 28195
- Release Information | 28195

Syntax

```
source url;
```

Hierarchy Level

```
[edit system scripts commit file filename]
```

Description

Specify the location of the master source file for a commit script. When you issue the `set refresh` configuration mode command at the same hierarchy level, the local copy of the script is overwritten by the version stored at the specified URL. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

> **NOTE**: On the QFabric system, commit scripts are stored in the **/pbdata/mgd_shared/*partition-ip*/var/db/scripts/op/** directory on the Director device.

Options

url The source specified as an HTTPS URL, HTTP URL, FTP URL, or scp-style remote file specification.

If you specify an HTTPS URL, you must also configure the `cert-file` statement at the same hierarchy level. Specify the path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate when you refresh the script.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configure and Use a Master Source Location for a Script

Example: Configure and Refresh from the Master Source for a Script

Overview of Updating Scripts from a Remote Source

cert-file (Scripts) | **3252**

refresh (Commit Scripts) | **14234**

source-count (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 28196
- Hierarchy Level | 28196
- Description | 28196
- Options | 28196
- Required Privilege Level | 28197
- Release Information | 28197

Syntax

```
source-count number;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name static group multicast-group-address source]
```

Description

Configure the number of multicast source addresses that should be accepted for each static group created on dynamic interfaces.

Options

number—Number of source addresses.

- **Default:** 1
- **Range:** 1 through 1024

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Enabling MLD Static Group Membership

source-count (Protocols IGMP)

IN THIS SECTION

- Syntax | 28197
- Hierarchy Level | 28198
- Description | 28198
- Options | 28198
- Required Privilege Level | 28198
- Release Information | 28198

Syntax

```
source-count number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name static group
multicast-group-address source],
[edit protocols igmp interface interface-name static group multicast-group-address source]
```

Description

Configure the number of multicast source addresses that should be accepted for each static group created.

Options

number—Number of source addresses.

- **Default:** 1
- **Range:** 1 through 1024

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Enabling IGMP Static Group Membership

source-count (Protocols MLD)

IN THIS SECTION

- Syntax | 28199
- Hierarchy Level | 28199
- Description | 28199
- Options | 28199
- Required Privilege Level | 28200
- Release Information | 28200

Syntax

```
source-count number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name static group
multicast-group-address source],
[edit protocols mld interface interface-name static group multicast-group-address source]
```

Description

Configure the number of multicast source addresses that should be accepted for each static group created.

Options

number—Number of source addresses.

- **Default:** 1
- **Range:** 1 through 1024

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Enabling MLD Static Group Membership

source (Distributed IGMP)

IN THIS SECTION

- Syntax | 28200
- Hierarchy Level | 28200
- Description | 28201
- Options | 28201
- Required Privilege Level | 28201
- Release Information | 28201

Syntax

```
source source-address <distributed>;
```

Hierarchy Level

```
[edit protocols pim static group multicast-group-address]
```

Description

Specify an IP unicast source address for a multicast group being statically configured on an interface.

Options

distributed (Optional) Enable a static join for multiple multicast address groups so that all Packet Forwarding Engines receive traffic, but preprovision only one multicast group.

source-address Specific IP unicast source address for a multicast group.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X50.

RELATED DOCUMENTATION

Enabling Distributed IGMP

Junos OS Multicast Protocols User Guide

Junos OS Multicast Protocols User Guide

source (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | 28202
- Hierarchy Level | 28202
- Description | 28202
- Options | 28202

Required Privilege Level | 28202
Release Information | 28202

Syntax

```
source source;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name static]
```

Description

Specify the IP version 4 (IPv4) unicast address to send data on an interface.

Options

source—IPv4 unicast address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Enabling IGMP Static Group Membership

source (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 28203
- Hierarchy Level | 28203
- Description | 28203
- Options | 28203
- Required Privilege Level | 28204
- Release Information | 28204

Syntax

```
source ip-address {
    source-count number;
    source-increment increment;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name static group multicast-group-address]
```

Description

IP version 6 (IPv6) unicast source address for the multicast group being configured on a dynamic interface.

Options

ip-address—One or more IPv6 unicast addresses.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Enabling MLD Static Group Membership

source-end-user-profile

IN THIS SECTION

- Syntax | 28204
- Hierarchy Level | 28205
- Description | 28205
- Options | 28205
- Required Privilege Level | 28205
- Release Information | 28205

Syntax

```
source-end-user-profile device-identity-profile-name;
```

Hierarchy Level

```
[edit security policies from-zone from-zone to-zone to-zone policy policy-name match]
```

Description

The `source-end-user-profile` field in a security policy enables you to specify a device identity profile that identifies the traffic source based on the device from which the traffic issued. The security policy action is applied to traffic issuing from a device if the device matches the attributes specified in the profile and it matches the rest of the security policy parameters.

The device identity profile feature provides a solution for cases in which you cannot or do not want to use the user identity to control access to network resources. The device identity feature allows you to use the identity of a device and its attributes to control access to network resources instead of the identity of the user of that device.

You might want to control network access based on the device identity for various reasons. For example, you might allow your users to use their own devices (BYOD) to access network resources and you do not want to use captive portal authentication. Also, some companies might have older switches that do not support 802.1, or they might not have a Network Access Control (NAC) system.

Options

device-identity-profile-name	Device identity profile that specifies characteristics that can apply to one or more devices.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

Understanding Access Control to Network Resources Based on Device Identity Information

Understanding the Device Identity Attributes and Profiles for the Integrated User Firewall Device Identity Authentication Feature

Understanding the Device Identity Authentication Table and Its Entries

source (Event Scripts)

IN THIS SECTION

- Syntax | 28206
- Hierarchy Level | 28206
- Description | 28206
- Options | 28207
- Required Privilege Level | 28207
- Release Information | 28207

Syntax

```
source url;
```

Hierarchy Level

```
[edit event-options event-script file filename]
```

Description

Specify the location of the master source file for an event script. When you issue the `set refresh` configuration mode command at the same hierarchy level, the local copy of the script is overwritten by the version stored at the specified URL. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

Options

url Master source file for an event script specified as an HTTPS URL, HTTP URL, FTP URL, or scp-style remote file specification.

If you specify an HTTPS URL, you must also configure the `cert-file` statement at the same hierarchy level. Specify the path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate when you refresh the script.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configure and Use a Master Source Location for a Script

Example: Configure and Refresh from the Master Source for a Script

cert-file (Scripts) | **3252**

refresh (Event Scripts) | **14236**

source-except

IN THIS SECTION

- Syntax | **28208**
- Hierarchy Level | **28208**
- Description | **28208**
- Options | **28208**
- Required Privilege Level | **28208**
- Release Information | **28208**

Syntax

```
source-except [address-name];
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match]
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match]
```

Description

Specify a source IP address or IP address set object to specify all source address objects except the specified address objects. The default value is any.

Options

address-name—IP address or IP address set object.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

source-filtering

IN THIS SECTION

- Syntax | 28209
- Hierarchy Level | 28209

- Description | 28209
- Default | 28209
- Required Privilege Level | 28209
- Release Information | 28210

Syntax

```
(source-filtering | no-source-filtering);
```

Hierarchy Level

```
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name fastether-options],
[edit interfaces interface-name gigether-options]
```

Description

For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and Gigabit Ethernet IQ interfaces only, enable the filtering of MAC source addresses, which blocks all incoming packets to that interface. To allow the interface to receive packets from specific MAC addresses, include the `source-address-filter` statement.

If the remote Ethernet card is changed, the interface is no longer able to receive packets from the new card because it has a different MAC address.

Default

Source address filtering is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring MAC Address Filtering for Ethernet Interfaces

Configuring MAC Address Filtering on PTX Series Packet Transport Routers

accept-source-mac | **1647**

source-address-filter | **28128**

source-filtering (Interfaces)

IN THIS SECTION

- Syntax | **28210**
- Hierarchy Level | **28210**
- Description | **28211**
- Required Privilege Level | **28211**
- Release Information | **28211**

Syntax

```
(source-filtering | no-source-filtering);
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options]
```

Description

For redundant Ethernet interfaces, enable the filtering of MAC source addresses, which blocks all incoming packets to that interface. To allow the interface to receive packets from specific MAC addresses, include the `source-address-filter` statement.

If the remote Ethernet card is changed, the interface cannot receive packets from the new card because it has a different MAC address.

By default, source address filtering is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

source-gateway (Tunnel Profile)

IN THIS SECTION

- Syntax | 28212
- Hierarchy Level | 28212
- Description | 28212
- Required Privilege Level | 28212
- Release Information | 28212

Syntax

```
source-gateway {
    address client-ip-address;
    gateway-name client-name;
}
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Description

Specify the IP address and hostname of the source gateway at the local L2TP tunnel endpoint, the LAC.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

source-hostname

IN THIS SECTION

- Syntax | 28213
- Hierarchy Level | 28213
- Description | 28213
- Required Privilege Level | 28213
- Release Information | 28214

Syntax

```
source-hostname source-hostname;
```

Hierarchy Level

```
[edit interfaces name unit name tunnel]
```

Description

Define the hostname of the tunnel source for the IP tunnel interface.

After the DNS resolves the IP address for the hostname, the resolved address is used as tunnel source address. The ipip interface (IFL) is not created unless the hostname is resolved.

.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 for SRX Series Firewalls (SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, SRX5800, and vSRX3.0)

Statement introduced in Junos OS Release 23.4R1 for SRX Series Firewalls (SRX38, SRX300, SRX320, SRX340, and SRX345)

RELATED DOCUMENTATION

Tunnel Services Overview

tunnel | **30190**

source-id

IN THIS SECTION

- Syntax | **28214**
- Hierarchy Level | **28215**
- Description | **28215**
- Options | **28215**
- Required Privilege Level | **28215**
- Release Information | **28215**

Syntax

```
source-id source-id;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

For version 9 flows, a 32-bit value that identifies the Exporter Observation Domain is called the source ID. NetFlow collectors use the combination of the source IP address and the source ID field to separate different export streams originating from the same exporter.

MX and QFX Series: Typically, this attribute is used to limit the scope of other information elements. If the observation domain is not unique, the collector cannot uniquely identify an version 9 device. If you configure the same source ID for different template types, such as for IPv4 and IPv6, it does not impact flow monitoring because the actual or the base source ID is transmitted in the flow. The actual source ID is derived from the value you configure and plus other parameters such as the slot number, lookup chip (LU) instance, and Packet Forwarding Engine instance. Such a method of computation of the observation domain ID ensures that this ID is not the same for two version 9 devices.

PTX Series: When you configure the source ID, the software attaches the ID to a particular template type. If you configure the same source ID for two different template types, such as for IPv4 and IPv6, this does not impact flow monitoring, because the configured ID is not what is being sent. The value sent in the packets is derived from that configured value and the FPC slot value. This method ensures two version 9 devices can never have the same value of observation domain ID.

Options

source-id—Identifier for the source for version 9 flows.

- **Range:** 0 through 255

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Evolved 22.4R1.

RELATED DOCUMENTATION

Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows

Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows

source-identity

IN THIS SECTION

- Syntax | 28216
- Hierarchy Level | 28217
- Description | 28217
- Options | 28217
- Required Privilege Level | 28218
- Release Information | 28218

Syntax

```
source-identity {
    [user-or-role-name];
    any;
    authenticated-user;
    unauthenticated-user;
    unknown-user;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name match]
```

```
[edit security policies global policy policy-name match]
```

```
[edit security advance-policy-based-routing from-zone zone-name policy policy-name match]
```

Description

Identifies users and roles to be used as match criteria for a policy. If a value other than `any` is specified as match criteria for a policy within a zone pair, the traffic is matched to table entries to retrieve associated user and roles before policy lookup occurs. Users and roles are retrieved from the local authentication table or from a UIT pushed to the SRX Series Firewall from an access control service when a user is authenticated.

Options

The following entries specify the source identities that match a policy:

user-or-role-name A list of specific users and roles.

- **Range:** 0 through 39 characters

> **NOTE**: SRX Series Firewalls truncate imported roles to 39 characters. You need to ensure that all of your roles are 39 characters or less.

any Any user or role, as well as the keywords authenticated-user, unauthenticated-user, and unknown-user.

authenticated-user All users and roles that have been authenticated.

unauthenticated-user Any user or role that does not have an IP-address mapped to authentication sources and the authentication source is up and running.

unknown-user Any user or role that does not have an IP address mapped to authentication sources, because the authentication source is disconnected from the SRX Series Firewall. In this case, users are unable to be authenticated due to an authentication server disconnection, such as a power outage.

Unknown-user must be configured for non-domain users to be able to authenticate and log in.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1. Statement updated in Junos OS Release 12.1X44-D10. Statement is supported in [edit security advance-policy-based-routing from-zone *zone-name* policy *policy-name* match] hierarchy in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Understanding User Role Firewalls

Understanding the User Identification Table

Security Policies Overview

source-identity-log (Security)

IN THIS SECTION

- Syntax | 28219
- Hierarchy Level | 28219
- Description | 28219
- Required Privilege Level | 28219
- Release Information | 28219

Syntax

```
source-identity-log
```

Hierarchy Level

```
[edit security zones security-zone zone-name]
```

Description

Specify the source-identity-log parameter as part of the configuration for a zone to enable it to trigger user identity logging when that zone is used as the source zone (from-zone) in a security policy. If a zone is configured for zone-based user identity logging and it is used as the source zone in a security policy, the system logs the user identity of any user who belongs to that zone and whose traffic matches the security policy's terms.

A zone configured for zone-based user identity logging is reusable. That is, you can use it as the source zone in any security policy.

For zone-based user identity logging to occur, you must have configured the session initialization (session-init) and the session termination (session-close) events as actions for the security policy.

Zone-based user identity logging allows you to broaden the scope of users whose identities are recorded in the session log. The source-identity security policy tuple writes the user or group name to log, but it restricts application of the security policy to the specified user or user group.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Understanding How to Include User Identity Information in the Session Log File Based on the Source Zone

Example: Configuring Integrated User Firewall to Write User Identity to the Session Log Based On the Source Zone

Overview of Integrated User Firewall

Example: Configuring Integrated User Firewall on SRX Series

source-increment (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 28220
- Hierarchy Level | 28220
- Description | 28220
- Options | 28221
- Required Privilege Level | 28221
- Release Information | 28221

Syntax

```
source-increment increment;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name static group multicast-group-address source]
```

Description

Configure the number of times the address should be incremented for each static group created on the dynamic interface. The increment is specified in a format similar to an IPv6 address.

Options

increment—Number of times the source address should be incremented.

- **Default:** ::1
- **Range:** ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff:

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Enabling MLD Static Group Membership

source-increment (Protocols IGMP)

IN THIS SECTION

- Syntax | 28222
- Hierarchy Level | 28222
- Description | 28222
- Options | 28222
- Required Privilege Level | 28222
- Release Information | 28222

Syntax

```
source-increment number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name static group
multicast-group-address source],
[edit protocols igmp interface interface-name static group multicast-group-address source]
```

Description

Configure the number of times the multicast source address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.

Options

increment—Number of times the source address should be incremented.

- **Default:** 0.0.0.1
- **Range:** 0.0.0.1 through 255.255.255.255

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Enabling IGMP Static Group Membership

source-increment (Protocols MLD)

IN THIS SECTION

- Syntax | 28223
- Hierarchy Level | 28223
- Description | 28223
- Options | 28223
- Required Privilege Level | 28224
- Release Information | 28224

Syntax

```
source-increment number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name static group
multicast-group-address source],
[edit protocols mld interface interface-name static group multicast-group-address source]
```

Description

Configure the number of times the address should be incremented for each static group created. The increment is specified in a format similar to an IPv6 address.

Options

increment—Number of times the source address should be incremented.

- **Default:** ::1
- **Range:** ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff:

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Enabling MLD Static Group Membership

source-interface (RADIUS Client)

IN THIS SECTION

- Syntax | 28224
- Hierarchy Level | 28224
- Description | 28225
- Options | 28225
- Required Privilege Level | 28225
- Release Information | 28225

Syntax

```
source-interface interface ipv4-address address;
```

Hierarchy Level

```
[edit access radius clients client-name]
```

Description

Configure the MX Series router interface and IPv4 address that receive RADIUS requests from the GGSN, PGW, or BNG RADIUS client.

Options

interface	Name of the interface.
address	IPv4 address on the MX Series router.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers

source-interface

IN THIS SECTION

- Syntax | 28226
- Hierarchy Level | 28226
- Description | 28226
- Required Privilege Level | 28226
- Release Information | 28226

Syntax

```
source-interface source-interface;
```

Hierarchy Level

```
[edit interfaces name unit name tunnel]
```

Description

Define the tunnel source interface. IP address from source-interface is used as tunnel source address.

When you configure IP adddress for the tunnel source-interface using a CLI, the IP tunnel derives that IP address and uses it as tunnel source address.

When you configure the tunnel source-interface to derive IP address from the DHCP server, IP tunnel derives the address once the source interface gets the address from the DHCP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 for SRX Series Firewalls (SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, SRX5800, and vSRX3.0)

Statement introduced in Junos OS Release 23.4R1 for SRX Series Firewalls (SRX38, SRX300, SRX320, SRX340, and SRX345)

RELATED DOCUMENTATION

Tunnel Services Overview

tunnel | **30190**

source-interface (RADIUS Server)

IN THIS SECTION

- Syntax | 28227
- Hierarchy Level | 28227
- Description | 28227
- Options | 28227
- Required Privilege Level | 28228
- Release Information | 28228

Syntax

```
source-interface interface [ipv4-address address];
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Specify the source interface and one or more IPv4 addresses on the MX Series router that receive RADIUS requests from which the RADIUS requests are sent to the RADIUS server.

Options

interface Source interface that sends the RADIUS requests.

address Source IPv4 address that sends the RADIUS requests. You can specify multiple source IPv4 addresses.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

source-interface (RADIUS Snoop Segment)

IN THIS SECTION

- Syntax | 28228
- Hierarchy Level | 28229
- Description | 28229
- Options | 28229
- Required Privilege Level | 28229
- Release Information | 28229

Syntax

```
source-interface source-interface;
```

Hierarchy Level

```
[edit access radius snoop-segments segment-name]
```

Description

Specify the interface of the MX Series router that receives accounting packets from the access network to be snooped.

Options

source-interface	Name of the interface.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview

IP-Based Subscriber Setup Overview

source-interface-set-at-login

IN THIS SECTION

Syntax | 28230

- Hierarchy Level | 28230
- Description | 28230
- Options | 28230
- Required Privilege Level | 28230
- Release Information | 28231

Syntax

```
source-interface-set-at-login {
    svlan;
}
```

Hierarchy Level

```
[edit protocols ppp-service]
```

Description

Enable PPP to generate the name of an interface set for use during subscriber login. The AAA process then sends the name by means of the Juniper Networks VSA QoS-Set-Name (26-130) to the RADIUS server in the Access-Request message. You can use this statement to discriminate between business and residential subscribers in heterogeneous networks. The interface set name is created for business subscribers. The RADIUS server then returns the VSA in Access-Accept messages only for business subscribers.

Options

svlan Specifies that PPP constructs the name of an interface from the physical interface name and the outer (SVLAN) VLAN tag.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Automatic Creation of Business Subscriber Interface Sets

source interfaces

IN THIS SECTION

- Syntax | 28231
- Hierarchy Level | 28232
- Description | 28232
- Options | 28232
- Required Privilege Level | 28233
- Release Information | 28233

Syntax

```
source interfaces (external | interface-name)
    hold-off-time time;
    priority number;
    quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
    request (force-switch | lockout);
    wait-to-restore minutes;
    aggregated-ether aggregated-ether-group-name;
}
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Configure the clock source that is used by the clock selection process on an interface.

- (SCB only) Specify the timeTransmitter clock source as the external-a interface and the secondary clock source as the external-b interface. The clock source is chosen using the clock selection process.
- (SCBE only) Specify the external interface to select the external clock source.
- (SCBE2 only) Specify the external-0/0 interface or the external-1/0 interface to select the external clock source.

Options

time Amount of time in milliseconds that a signal is held before being passed to the clock selection process.

Range: 300–1800 milliseconds

Default: 1000 milliseconds

number Set the priority level of the clock source.

Range: 1 through 5

`prc`—Primary reference clock

`prs` —Primary reference source—Stratum 1

`sec`—SDH equipment clock

`smc`—SONET minimum clock

`ssu-a`—Synchronization supply unit A

`ssu-b`—Synchronization supply unit B

`st2`—Stratum 2

`st3`—Stratum 3

`st3e`—Stratum 3E

st4—Stratum 4

force-switch Force switching to a clock source, provided the clock source is enabled and not locked out. Only one configured source may be force-switched.

lockout Clock source is not considered by the selection process. Lockout may be configured for any source.

minutes Set the time for the port signal to be up before the port is opened to receive and transmit ESMC messages.

Range: 0 through 12 minutes

Default: 5 minutes

aggregated-ether-group-name aggregated ethernet group name

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

source-ip

IN THIS SECTION

- Syntax (EX Series) | 28234
- Syntax (MX Series, QFX Series, OCX Series, and PTX Series) | 28234

- Hierarchy Level | 28234
- Description | 28234
- Required Privilege Level | 28234
- Release Information | 28234

Syntax (EX Series)

```
source-ip (ip4-address | ipv6-address);
```

Syntax (MX Series, QFX Series, OCX Series, and PTX Series)

```
source-ip {
    ipv4-address;
    inet6 ipv6-address;
}
```

Hierarchy Level

```
[edit protocols sflow]
```

Description

Configure the IP address to be used for the sFlow datagrams. By configuring an IP address, you ensure that the IP address is not dynamic. You can specify an IPv4 or IPv6 address or both.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

IPv6 addresses are supported in Junos OS Release 19.1R1 for MX Series, PTX Series, and QFX Series devices

RELATED DOCUMENTATION

Configuring sFlow Technology for Network Monitoring (CLI Procedure)

source-ip-address-list

IN THIS SECTION

- Syntax | 28235
- Hierarchy Level | 28235
- Description | 28236
- Options | 28236
- Required Privilege Level | 28236
- Release Information | 28236

Syntax

```
source-ip-address-list address-list-name;
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard policy policy-name discard
(forwarding-options)]
[edit forwarding-options access-security router-advertisement-guard policy policy-name accept (Router
Advertisement Guard Policy) match-list]
```

Description

Configure a list of IPv6 addresses for an IPv6 Router Advertisement (RA) guard policy to validate the source IPv6 address of an incoming RA message against the IPv6 addresses in this list. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.

You can use a list of IPv6 addresses for validating RA messages as part of either an accept policy or a discard policy. Before you can include a list in an RA policy, you must configure the list name at the [`edit policy-options prefix-list`] hierarchy level. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped. When RA guard is enabled by using a discard policy, any RA messages that match the conditions are dropped, and RA messages that do not match the conditions are forwarded.

Options

address-list-name Configure a list of IPv6 addresses to use in an RA guard policy. The policy is used to validate the source of an incoming RA message by comparing the IPv6 address of the RA message to the IPv6 addresses contained in the list.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

source-ip-address (RADIUS Snoop Segment)

IN THIS SECTION

- Syntax | 28237
- Hierarchy Level | 28237
- Description | 28237
- Options | 28237
- Required Privilege Level | 28237
- Release Information | 28238

Syntax

```
source-ip-address source-address;
```

Hierarchy Level

```
[edit access radius snoop-segments segment-name]
```

Description

Specify the source IP address of accounting requests from a GGSN, PGW, or BNG to snoop. If you do not enter a source IP address, accounting requests from any IP address can be snooped.

Options

source-address	Source IPv4 address.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview

IP-Based Subscriber Setup Overview

source-ip-based

IN THIS SECTION

- Syntax | 28238
- Hierarchy Level | 28238
- Description | 28239
- Options | 28239
- Required Privilege Level | 28239
- Release Information | 28239

Syntax

```
source-ip-based number;
```

Hierarchy Level

```
[edit security screen ids-option screen-name limit-session]
```

Description

Limit the number of concurrent sessions the device can initiate from a single source IP address.

Options

number —Maximum number of concurrent sessions that can be initiated from a source IP address.

- **Range:** 1 through 1,000,000

For SRX Series Firewalls, the applicable range is 1 through 8,000,000.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

source-ip-change (Forwarding Options)

IN THIS SECTION

- Syntax | 28240
- Hierarchy Level | 28240
- Description | 28240
- Required Privilege Level | 28240
- Release Information | 28240

Syntax

```
source-ip-change;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay]
```

Description

For Dynamic Host Configuration Protocol (DHCP) client request forwarding, enable source IP change for the device to use address of egress interface as source IP address.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement changed from `vpn` to `source-ip-change` in Junos OS Release 15.1X49-D130 and later releases.

RELATED DOCUMENTATION

DHCP Overview

source-ip (Flow Monitoring Logs for NAT)

IN THIS SECTION

- Syntax | 28241
- Hierarchy Level | 28241

- Description | 28241
- Options | 28241
- Required Privilege Level | 28241
- Release Information | 28241

Syntax

```
source-ip address;
```

Hierarchy Level

```
[edit services jflow-log collector collector-name]
```

Description

Specify the source IPv4 address of the services PIC interface to be used for generation of flow monitoring log messages in flow monitoring template format for NAT events.

Options

address—Valid IPv4 address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

source-ip (Security Alarms)

IN THIS SECTION

- Syntax | 28242
- Hierarchy Level | 28242
- Description | 28243
- Options | 28243
- Required Privilege Level | 28243
- Release Information | 28243

Syntax

```
source-ip {
    duration interval;
    size count;
    threshold value;
}
```

Hierarchy Level

```
[edit security alarms potential-violation policy]
```

Description

Configure alarms for a number of policy violations by a source network identifier within a specified time period.

Options

- `duration` *`interval`*—Indicate the duration of counters.
- **Range:** 1 through 3600 seconds.
- **Default:** 1 second.
- `size` *`count`*—Indicate the number of source IP addresses for which policy violation checks can be done concurrently.
- **Range:** 1 through 10240.
- **Default:** 1024.
- `threshold` *`value`*—Indicate the maximum number of source IP address matches required to raise an alarm.
- **Range:** 1 through 4294967295.
- **Default:** 1000.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Security Policies Overview

source-ipv4-address (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 28244
- Junos OS Hierarchy Level | 28244
- Junos OS Evolved Hierarchy Level | 28244
- Description | 28244
- Options | 28245
- Required Privilege Level | 28245
- Release Information | 28245

Syntax

```
source-ipv4-address address;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the source IPv4 address to be used in generated test frames. This parameter is optional for both `ccc` and `inet` families. If you do not configure the source IPv4 address for an `inet` family, the source address of the interface is used to transmit the test frames.

Options

address Valid IPv4 address.

- **Default:** If you do not configure the source IPv4 address for a ccc family, default value of 192.168.1.10 is used.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

rfc2544-benchmarking | **15759**

source (JET Scripts)

IN THIS SECTION

- Syntax | **28246**
- Hierarchy Level | **28246**
- Description | **28246**
- Options | **28246**
- Required Privilege Level | **28246**
- Release Information | **28246**

Syntax

```
source url;
```

Hierarchy Level

```
[edit system extensions extension-service application file filename]
```

Description

Specify the location of the master source file for a JET script. When you issue the `set refresh` configuration mode command at the same hierarchy level, the local copy of the script is overwritten by the version stored at the specified URL. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

NOTE: JET scripts are stored in the **/var/db/scripts/jet** directory.

Options

url Master source file for a JET script specified as an HTTPS URL, HTTP URL, FTP URL, or scp-style remote file specification.

If you specify an HTTPS URL, you must also configure the `cert-file` statement at the same hierarchy level. Specify the path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate when you refresh the script.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

- Configuring and Using a Master Source Location for a Script
- Example: Configuring and Refreshing from the Master Source for a Script

source-mac-address

IN THIS SECTION

- Syntax | 28247
- Junos OS Hierarchy Level | 28247
- Junos OS Evolved Hierarchy Level | 28247
- Description | 28248
- Options | 28248
- Required Privilege Level | 28248
- Release Information | 28248

Syntax

```
source-mac-address mac-address;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the source MAC address used in generated test frames. This parameter is mandatory for a `bridge`, `ethernet-switching`, or `vpls` family and is optional for a `ccc` family.

Options

mac-address

Source MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn*.*nnnn*.*nnnn* or *nn*:*nn*:*nn*:*nn*:*nn*:*nn*, for example, `0000:5e00:5355` or `00:00:5e:00:53:55`.

Default: 0x00:0x60:0x67:0x71:0xC6:0x62

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

rfc2544-benchmarking | **15759**

rfc2544 | **15757**

RFC 2544-Based Benchmarking Tests Overview

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

Configuring an RFC 2544-Based Benchmarking Test

source-mac-address-list

IN THIS SECTION

- Syntax | 28249
- Hierarchy Level | 28249
- Description | 28249
- Options | 28250
- Required Privilege Level | 28250
- Release Information | 28250

Syntax

```
source-mac-address-list address-list-name;
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard policy policy-name discard
(forwarding-options)]
[edit forwarding-options access-security router-advertisement-guard policy policy-name accept (Router
Advertisement Guard Policy) match-list]
```

Description

Configure a list of MAC addresses for an IPv6 Router Advertisement (RA) guard policy to validate the source MAC address of an incoming RA message against the MAC addresses in this list. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.

You can use a list of MAC address for validating RA messages as part of either an accept policy or a discard policy. Before you can include a list in an RA policy, you much configure the list name at the [`edit policy-options mac-list`] hierarchy level. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not

match the conditions are dropped. When RA guard is enabled by using a discard policy, any RA messages that match the conditions are dropped, and RA messages that do not match the conditions are forwarded.

Options

address-list-name Configure the RA guard policy to match the MAC source address of an incoming RA message to a MAC address contained in the list.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

source-mode

IN THIS SECTION

- Syntax | **28251**
- Syntax (ACX710) | **28251**
- Hierarchy Level | **28251**
- Description | **28251**
- Options | **28251**
- Required Privilege Level | **28251**

Release Information | 28252

Syntax

```
source-mode (chassis | line);
```

Syntax (ACX710)

```
source-mode (chassis | ptp-clock);
```

Hierarchy Level

```
[edit chassis synchronization output interfaces external]
[edit chassis synchronization output interfaces (external-0/0 | external-1/0)]
[edit chassis synchronization interfaces bits]
```

`[edit chassis synchronization interfaces bits]` hierarchy added for ACX710 in Junos OS Release 20.2R1.

Description

Configure a clock source for output for the configured BITS interface.

Options

chassis Set the chassis clock for output.

line Set the best line clock source for output.

ptp-clock (ACX710 only) Set the PTP clock for output.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

`ptp-clock` option introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

synchronization (ACX Series) | 28787

synchronization (MX Series) | 28795

source (Multicast VLAN Registration)

IN THIS SECTION

- Syntax | 28252
- Hierarchy Level | 28253
- Description | 28253
- Default | 28253
- Options | 28253
- Required Privilege Level | 28253
- Release Information | 28253

Syntax

```
source {
    groups group-prefix;
}
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan  vlan-name data-forwarding]
```

Description

Configure a VLAN to be a multicast source VLAN (MVLAN), and specify the IP address range of the multicast source groups.

To configure a data-forwarding VLAN as an MVLAN, you also configure one or more multicast receiver VLANs (MVR receiver VLANs) with hosts that might be interested in receiving traffic on the MVLAN for the specified multicast groups. You can configure a VLAN as either an MVLAN or MVR receiver VLAN, but not both at the same time.

> **NOTE**: On EX4300 and EX4300 multigigabit switches, you can configure up to 10 MVLANs, and up to a total of 4K MVR receiver VLANs and MVLANs together. On EX2300 and EX3400, you can configure up to 5 MVLANs and the remaining configurable VLANs can be MVR receiver VLANs.

The remaining statement is explained separately. See CLI Explorer.

Default

Disabled

Options

groups *group-prefix* — IP address range of the source groups. Each MVLAN must have exactly one `groups` statement. If there are multiple MVLANs on the switch, their group ranges must be unique.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring Multicast VLAN Registration on EX Series Switches Without ELS

Configuring Multicast VLAN Registration on EX Series Switches

source-nat

IN THIS SECTION

- Syntax | 28254
- Hierarchy Level | 28255
- Description | 28255
- Options | 28255
- Required Privilege Level | 28255
- Release Information | 28255

Syntax

```
source-nat {
    interface {
        persistent-nat {
            address-mapping;
            inactivity-timeout seconds;
            max-session-number value;
            permit (any-remote-host | target-host | target-host-port);
        }
    }
    off;
    pool <pool-name>;
         persistent-nat {
             address-mapping;
            inactivity-timeout seconds;
            max-session-number number;
            permit (any-remote-host | target-host | target-host-port);
    }
```

```
    rule-session-count-alarm (clear-threshold value | raise-threshold value);
}
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name  then]
```

Description

Specify the action of the source NAT rule.

Options

- `off`—Do not perform the source NAT operation.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security — To view this statement in the configuration.

security-control— To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10.

source (Op Scripts)

IN THIS SECTION

- Syntax | 28256
- Hierarchy Level | 28256
- Description | 28256
- Options | 28256

Required Privilege Level | 28256
Release Information | 28257

Syntax

```
source url;
```

Hierarchy Level

```
[edit system scripts op file filename]
```

Description

Specify the location of the master source file for an op script. When you issue the `set refresh` configuration mode command at the same hierarchy level, the local copy of the script is overwritten by the version stored at the specified URL. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.

NOTE: On the QFabric system, commit scripts are stored in the **/pbdata/mgd_shared/*partition-ip*/var/db/scripts/op/** directory on the Director device.

Options

url Master source file for an op script specified as an HTTPS URL, HTTP URL, FTP URL, or scp-style remote file specification.

If you specify an HTTPS URL, you must also configure the `cert-file` statement at the same hierarchy level. Specify the path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate when you refresh the script.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configure and Use a Master Source Location for a Script

Example: Configure and Refresh from the Master Source for a Script

cert-file (Scripts) | **3252**

refresh (Op Scripts) | **14251**

source-packet-routing (Segment Routing Traffic Engineering)

IN THIS SECTION

- Syntax | **28257**
- Hierarchy Level | **28258**
- Description | **28258**
- Options | **28258**
- Required Privilege Level | **28260**
- Release Information | **28260**

Syntax

```
source-packet-routing {
    inherit-label-nexthops;
    lsp-external-controller name;
    maximum-segment-list-depth maximum-segment-list-depth
    maximum-srv6-te-segment-list-depth maximum-srv6-segment-list-depth;
    preference preference;
    preserve-nexthop-hierarchy;
    retry-timer seconds;
    sbfd-frr;
```

```
    segment-list;
    source-routing-path name;
    sr-preference sr-preference;
    sr-preference-override sr-preference-override;
    telemetry;
    srv6;
    traceoptions;
    traffic-engineering (SR-TE import policy)
    tunnel-tracking;
}
```

Hierarchy Level

```
[edit logical-systems name protocols],
[edit protocols]
```

Description

Enable Source Packet Routing in Networking (SPRING) or segment routing for traffic engineering packets for protocols that support segment routing. For example, you can configure segment routing capability for a BGP address family to steer traffic according to network requirements.

To enable segment routing traffic engineering for a BGP family, configure `segment-routing-te` at the `[edit protocols bgp family inet|inet6]` and `[edit protocols bgp group group-name neighbor neighborfamily inet|inet6]` hierarchy levels. Also, configure `sr-preference-override` at the `[edit protocols bgp family inet|inet6 source-packet-routing]` hierarchy level. This allows BGP to support segment routing policies and to steer traffic based on the segment list.

NOTE: `sr-preference-override` is not applicable for static SR-TE SRv6 LSPs.

A policy action `color: color-mode:color-value` is configured at the `[edit policy-options community namemembers]` hierarchy level to categorize segment routing traffic. The color value is attached to BGP communities when exporting prefixes from inet-unicast and inet6-unicast address families, which associates a BGP community with a segment routing policy path definition.

Options

inherit-label-nexthops Inherit label next hops for first hop in all the segment lists.

When the `inherit-label-nexthops` is configured at the `[edit protocols source-packet-routing]` hierarchy, it is applied globally, and takes precedence over individual `segment-list` level configuration.

For dynamic non-colored static LSPs, that is the PCEP-driven segment routing LSPs, the `inherit-label-nexthops` statement must be enabled globally, as the segment-level configuration is not applied.

lsp-external-controller *name*

Specify an external path computing entity, typically the network controller.

maximum-segment-list-depth

Specify a maximum segment list depth for segment routing traffic engineering policies.

- **Default:** The default value is 5.
- **Range:** 1 through 16.

maximum-srv6-segment-list-depth

Specify a maximum srv6 segment list depth for segment routing traffic engineering policies.

- **Default:** For reduced SRH, the default value is 6.
- **Range:** 1 through 6.
- **Default:** For non-reduced SRH, the default value is 5.
- **Range:** 1 through 5.

> **NOTE**: When `no-reduced-srh` is configured under the `[edit routing-options source-packet-routing srv6]` hierarchy level, ensure that the segment-list has a maximum of 5 SIDs, else configure `maximum-srv6-segment-list-depth` to have a commit constraint.

preference

Route preference for segment routing traffic engineering routes.

retry-timer

Number of seconds to wait before trying to reconnect to the auto-translate service if the previous attempt failed. Connections may fail if there is no link, or node, for a given IP address, or if the link or node does not have a SID.

- **Default:** 30 seconds
- **Range:** 1 through 600

sbfd-frr

Enable seamless BFD (S-BFD) fast re-route (FRR), a feature designed to enhance the resilience of SR-TE tunnels by monitoring end-to-end paths and initiating local repair

mechanisms when failures are detected, rerouting traffic to alternate paths to minimize disruption.

sr-preference Specify a preference for static segment routing policies. BGP chooses a segment routing policy with a higher value over policies with lower values.

- **Range:** `0` through `4,294,967,295`

sr-preference-override Specify a preference override for static segment routing policies that BGP uses to select a segment routing policy. BGP chooses a policy with a higher value over policies with lower values.

- **Default:** The default value is `100`
- **Range:** `0` through`4,294,967,295`

traffic-engineering Enable import of all Spring-TE policies into TED

tunnel-tracking When the `tunnel-tracking` option is enabled, it does a first-hop validation for SR-TE tunnels according to their route resolution. Status of each path will also be displayed when this option is enabled.

srv6 Specify SRv6 support for SR-TE policies. Enable the `srv6` option to enable the SRv6 data plane support for statically configured SR-TE policies.

preserve-nexthop-hierarchy Specify preserve nexthop hierarchy support for SR-TE routes.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1 for the MX Series and PTX Series with FPC-PTX-P1-A.

`telemetry` option introduced in Junos OS Release 18.3R1 for the MX Series and PTX Series with FPC-PTX-P1-A.

`inherit-label-nexthops` and `retry-timer` statements introduced in Junos OS Release 19.1R1 on MX Series routers.

`tunnel-tracking` option introduced in Junos OS Release 21.2R1.

`srv6` option introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

- show spring-traffic-engineering | 26587
- extended-nexthop-color
- segment-list | 16521
- source-routing-path | 28316
- sr-preference-override
- *Segment Routing Traffic Engineering at BGP Ingress Peer Overview*
- *Understanding Static Segment Routing LSP in MPLS Networks*
- *Understanding SR-TE Policy for SRv6 Tunnel*

source-packet-routing (Junos Telemetry Interface)

IN THIS SECTION

- Syntax | 28262
- Hierarchy Level | 28262
- Description | 28262
- Default | 28263
- Options | 28263
- Required Privilege Level | 28264
- Release Information | 28264

Syntax

```
source-packet-routing {
    telemetry {
        statistics {
            per-source per-segment-list
            no-transit;
            no-ingress;
            }
        }
     source-routing-path {
         use-for-shortcut;
         }
    }
```

Hierarchy Level

```
[edit protocols]
```

Description

Enable BGP and statically configured Segment Routing Traffic Engineering (SR-TE) traffic statistics sensor support for Junos telemetry interface (JTI).

Export JTI statistics using either remote procedure call (gRPC) services or UDP native sensors to stream statistics. The following resource paths are supported. For UDP native sensors:

- **/junos/services/segment-routing/traffic-engineering/ingress/usage/**
- **/junos/services/segment-routing/traffic-engineering/transit/usage**

For gRPC streaming:

- /mpls/signaling-protocols/segment-routing/

Export SR-TE per Label Switched Path (LSP) route statistics using JTI and gRPC services. Using JTI and gRPC services. You can stream SR-TE telemetry statistics for uncolored SR-TE policies statistics to an outside collector. Ingress statistics include statistics for all traffic steered by means of an SR-TE LSP. Transit statistics include statistics for traffic to the Binding SID (BSID) of the SR-TE policy.

> **NOTE**: Enabling the parameter `per-segment-list` will disable tunnel sensors. For example, configuring `per-souorce per-segment-list` will disable the sensor `/junos/services/segment-routing/traffic-engineering/tunnel/ingress/usage/`.

To enable these statistics, include the **per-source per-segment-list** option at the [`edit protocols source-packet-routing telemetry statistics`] hierarchy level. When configuring:

- If the statement `set protocols source-packet-routing telemetry statistics no-ingress` is issued, ingress sensors are not created.
- If the statement `set protocols source-packet-routing telemetry statistics no-transit` is issued, transit sensors are not created. Otherwise, if BSID is configured for a tunnel, transit statistics are created.

The following resource paths (sensors) are supported:

- **/junos/services/segment-routing/traffic-engineering/tunnel/lsp/ingress/usage/**
- **/junos/services/segment-routing/traffic-engineering/tunnel/lsp/transit/usage/**

For exporting statistics using UDP native sensors, configure parameters at the [edit services analytics] hierarchy level. To provision sensors to export data through gRPC streaming, use the telemetrySubscribe RPC to specify telemetry parameters.

Default

Disabled.

Options

statistics Create sensors for both the SR-TE policy nexthop and the binding SID that are installed in the forwarding plane. For the SR-TE policy nexthop, the sensors collect traffic statistics steered by all routes that use the SR-TE policy as a nexthop. For the binding SID, the sensors collect statistics on labeled traffic that is steered by the binding-SID route.

per-source per-segment-list Create sensors to export SR-TE per LSP route statistics for uncolored SR-TE policies. Ingress statistics include statistics for all traffic steered by means of an SR-TE LSP. Transit statistics include statistics for traffic to the BSID of the SR-TE policy.

no-transit Enable sensors only for SR-TE policy nexthops. The sensor will collect statistics on all steering routes that use the SR-TE policy as a nexthop.

no-ingress	Enable sensors only for Binding-SID transit routes.
source-routing-path	Specify the name of the source routing path.
use-for-shortcut	Enable to allow the LSP to be used as a shortcut tunnel.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.3R1.

Option **per-source per-segment-list** is introduced in Junos OS Release 20,1R1 for MX Series and PTX Series routers.

Statement supported in Junos OS Release 20.2R1 for MX240. MX480, MX960, MX2010, and MX2020 with MPC-10E or MPC-11E routers.

Statement supported in Junos OS Evolved Release 21.4R1 for PTX10001-36MR, PTX10004, PTX10008, and PTX10016 routers.

RELATED DOCUMENTATION

Understanding OpenConfig and gRPC on Junos Telemetry Interface

Configure a NETCONF Proxy Telemetry Sensor in Junos

sensor (Junos Telemetry Interface) | **16564**

statistics | **28591**

telemetry | **29054**

source-packet-routing (routing-options)

IN THIS SECTION

- Syntax | **28265**
- Hierarchy Level | **28265**

- Description | 28265
- Required Privilege Level | 28265
- Release Information | 28265

Syntax

```
mapping-server-entry mapping-server-entry;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options],
[edit logical-systems name routing-options],
[edit routing-instances name routing-options],
[edit routing-options]
```

Description

Configure interoperability between islands of devices supporting only segment routing and LDP in an LDP network domain where there is gradual deployment of segment routing.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.2R1.

source-packet-routing (Protocols IS-IS)

IN THIS SECTION

- Syntax | 28266
- Hierarchy Level | 28267
- Description | 28268
- Default | 28268
- Options | 28268
- Required Privilege Level | 28269
- Release Information | 28269

Syntax

```
source-packet-routing {
    adjacency-segment {
        hold-time hold-time;
    }
    disable;
    explicit-null;
    node-segment {
        index-range index range;
        ipv4-index index;
        ipv6-index index;
    }
    flex-algorithm name {
        color color;
        definition;
    }
    srgb {
        start-label start-label;
        index-range index range;
    }
    no-binding-sid-leaking;
    no-strict-spf
    srv6 {
```

```
        locator name {
            end-sid sid
                flavor {
                    psp;
                    usp;
                    usd;
                }
            }
        }
    udp-tunneling {
        encapsulation;
   }
use-flex-algorithm-metric-always;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
[edit logical-systems logical-system-name protocols isis level level-number],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols isis
level level-number],
[edit protocols isis level level-number],
[edit routing-instances routing-instance-name protocols isis level level-number]
[edit logical-systems logical-system-name protocols ospf],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf],
[edit protocols ospf],
[edit routing-instances routing-instance-name protocols ospf]
[edit logical-systems logical-system-name protocols ospf area-number],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area-number],
[edit protocols ospf area-number],
[edit routing-instances routing-instance-name protocols ospf area-number]
```

Description

Enable source packet routing in networking (SPRING) feature on IS-IS levels or OSPF areas.

Source packet routing for IPv6 is supported only for IS-IS. The `adjacency-segment`, `explicit-null`, and `srgb` statements are supported only for IS-IS.

The remaining statements are explained separately.

Default

Disabled.

Options

adjacency-segment <hold-time *hold-time*>

Configure attributes for adjacency segments in source packet routing in networking (SPRING), or configure segment routing (SR) to ensure that the adjacency segment identifiers are retained during adjacency or link flaps. The adjacency segments are not released immediately and are retained for the configured hold time duration.

- hold-time *hold-time*—(Optional) Duration, in milliseconds, to retain adjacency segments after isolating from an interface. The range is 180,000 through 900,000, with the default being 300,000 milliseconds for IS-IS interfaces and 180,000 milliseconds for OSPF interfaces.

disable

Disable source packet routing from a specific level.

explicit-null

Configure E and P bits in all prefix segment identifier (SID) advertisements.

no-strict-spf

Disable strict spf algo 1 advertisement.

node-segment

Enable source packet routing in networking (SPRING) at all levels. SPRING or segment routing is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to determine the actual path it should take.

Provisioning the IPv4 and IPv6 node segment index is allowed per routing-instance, and will NOT be allowed per IS-IS level. Node segment index is attached to the IPv4 and IPv6 router-id, if the router-ids are configured on the loopback interface. Else, lowest IP address on the loopback is chosen to attach the node-sid.

index-range *index range*

Range of node segment indices allowed. The range is 32 through 16384, and the default is 4096.

ipv4-index *index*

IPv4 node segment index. The range is 0 through 199999.

Starting with Junos OS Release 17.2, the maximum index for IPv4 node segment index is 199999.

ipv6-index *index*

IPv6 node segment index. The range is 0 through 199999.

Starting with Junos OS Release 17.2, the maximum index for IPv6 node segment index is 199999.

udp-tunneling

Configure `udp-tunneling` option to enable SR-o-UDP feature. When this option is enabled, during route computation, IS-IS looks for routes whose next hop neighbor does not support segment routing. IS-IS creates UDP tunnel next hops to every route that does not support segment routing. The source address of UDP tunnel is set to self router-id and destination address of the UDP tunnel is set to the prefix for which route computation is being performed. The statement is introduced only at [edit protocols isis] hierarchy level.

encapsulation

The encapsulation statement enables SR-over-UDP encapsulation, thus creates a tunnel to traverse an SR traffic across a non-SR core.

> **NOTE**: Configure a filter-based de-encapsulation on MX series router and an FTI-based de-encapsulation on PTX Series routers.

use-flex-algorithm-metric-always

Specify to indicate flex-algo path computations to always use the metric advertised in FAPM during path computation. By default, flex-algo path computation will use FAPM for readvertised prefixes, for which the R bit in prefix SID is considered. In a multi-vendor environment, if this check needs to be relaxed, you can use this option. Path computations will still use FAPM if the winner FAD in that level is set to M flag.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2R1.

Support for `explicit-null` and `adjacency-segment` statements introduced in Junos OS Release 17.2R1 for MX Series routers, PTX Series routers, QFX5100 switches and QFX10000 switches.

The udp-tunneling option is introduced in Junos OS Evolved Release 21.2R1.

RELATED DOCUMENTATION

Understanding Source Packet Routing in Networking (SPRING)

Understanding Adjacency Segments, Anycast Segments, and Configurable SRGB in SPRING

Configuring Anycast and Prefix segments in SPRING for IS-IS Protocol

source-packet-routing (Protocols OSPF)

IN THIS SECTION

- Syntax | 28270
- Hierarchy Level | 28271
- Description | 28271
- Default | 28271
- Options | 28271
- Required Privilege Level | 28273
- Release Information | 28273

Syntax

```
source-packet-routing {
    adjacency-segment {
        hold-time hold-time;
    }
    disable;
    explicit-null;
    node-segment {
        index-range index range;
        ipv4-index index;
```

```
    }
   flex-algorithm name {
        color color;
        definition;
    }
    sensor-based-stats;
 strict-asla-based-flex-algorithm;
    srgb {
        start-label start-label;
        index-range index range;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf],
[edit protocols ospf],
[edit routing-instances routing-instance-name protocols ospf]
```

Description

Configures source packet routing in networking (SPRING) feature.

Default

Disabled on all the levels.

Options

adjacency-segment <hold-time *hold-time*>

Configure attributes for adjacency segments in source packet routing in networking (SPRING), or configure segment routing (SR) to ensure that the adjacency segment identifiers are retained during adjacency or link flaps. The adjacency segments are not released immediately and are retained for the configured hold time duration.

- hold-time *hold-time*—(Optional) Duration, in milliseconds, to retain adjacency segments after isolating from an interface. The range is 180,000 through 900,000,

with the default being 300,000 milliseconds for IS-IS interfaces and 180,000 milliseconds for OSPF interfaces.

disable Disable source packet routing from a specific level.

explicit-null Configure E and P bits in all prefix segment identifier (SID) advertisements.

node-segment Enable source packet routing in networking (SPRING) at all levels. SPRING or segment routing is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to determine the actual path it should take.

> **NOTE**: Provisioning the IPv4 and IPv6 node segment index is allowed per routing-instance, and will NOT be allowed per IS-IS level. Node segment index is attached to the IPv4 and IPv6 router-id, if the router-ids are configured on the loopback interface. Else, lowest IP address on the loopback is chosen to attach the node-sid.

index-range *index range* Range of node segment indices allowed. The range is 32 through 16384, and the default is 4096.

ipv4-index *index* IPv4 node segment index. The range is 0 through 199999.

> **NOTE**: Starting with Junos OS Release 17.2, the maximum index for IPv4 node segment index is 199999.

ipv6-index *index* IPv6 node segment index. The range is 0 through 199999.

> **NOTE**: Starting with Junos OS Release 17.2, the maximum index for IPv6 node segment index is 199999.

strict-asla-based-flex-algorithm Specify to indicate flexible algorithm path computations to use only the links that advertise relevant te-attributes through Application-Specific Link Attributes (ASLA).

The option `source-packet-routing strict-asla-based-flex-algorithm` has to be applied to all the flexible algorithms running on the devices in the network to avoid routing loops.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

The `strict-asla-based-flex-algorithm` option introduced in Junos OS and Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

OSPF Overview

source-packet-routing (Routing Instances Protocols EVPN)

IN THIS SECTION

- Syntax | 28273
- Hierarchy Level | 28274
- Description | 28274
- Options | 28274
- Required Privilege Level | 28274
- Release Information | 28274

Syntax

```
source-packet-routing {
 srv6 locator loc-name (end-dt2-sid | micro-dt2-sid);
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols evpn],
[edit routing-instances name protocols evpn]
```

Description

Enable IPv6 segment routing for EVPN E-LAN services.

Options

locator *loc-name* Required. Name of the SRv6 locator (defined at `[edit routing-options source-packet-routing srv6]`).

end-dt2-sid Segment Routing Header-based End.DT2 SIDs.

micro-dt2-sid Micro-SID based End.DT2 SIDs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4 R1 and Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

encapsulation (protocols evpn) | 6284

source (PIM RPF Selection)

IN THIS SECTION

- Syntax | 28275
- Hierarchy Level | 28275
- Description | 28275
- Options | 28275
- Required Privilege Level | 28276
- Release Information | 28276

Syntax

```
source source-address {
    next-hop next-hop-address;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address],
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-
list-addresses]
```

Description

Configure the source address for the PIM group.

Options

source-address—Specific source address for the PIM group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring PIM RPF Selection

source-pool

IN THIS SECTION

- Syntax | 28276
- Hierarchy Level | 28276
- Description | 28277
- Required Privilege Level | 28277
- Release Information | 28277

Syntax

```
source-pool nat-pool-name;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated]
```

Description

Specify the source address pool for translated traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

source-port

IN THIS SECTION

- Syntax | 28277
- Hierarchy Level | 28278
- Description | 28278
- Options | 28278
- Required Privilege Level | 28278
- Release Information | 28278

Syntax

```
source-port port-number;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Source port identifier.

Options

port-value—Identifier for the port. For a complete list, see *Configuring Source and Destination Ports*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- *ALG Descriptions*
- *Configuring Application Properties*
- *Configuring Source and Destination Ports*
- *Verifying the Output of ALG Sessions*

source-port (firewall)

IN THIS SECTION

Syntax | 28279

- Hierarchy Level | 28279
- Description | 28279
- Options | 28279
- Required Privilege Level | 28279
- Release Information | 28279

Syntax

```
source-port <source-port>;
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name from ip-version ip-version
protocol (tcp|udp)]
```

Description

Configure the source port of the Layer 4 header.

Options

source-port The source port of the Layer 4 header.

- **Range:** 0 through 65,535

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1R1.

RELATED DOCUMENTATION

Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic

source-port (Applications)

IN THIS SECTION

- Syntax | 28280
- Hierarchy Level | 28280
- Description | 28280
- Options | 28280
- Required Privilege Level | 28281
- Release Information | 28281

Syntax

```
source-port port-number;
```

Hierarchy Level

```
[edit applications application  application-name  ]
[edit applications application  application-name  term  term-name]
```

Description

Specify the source port identifier.

Options

port-number—Identifier for the port. You can use a numeric value or one of the text synonyms listed in "destination-port (Applications)" on page 5465.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policy Application Sets Overview

source-port (Security Forwarding Options)

IN THIS SECTION

- Syntax | 28281
- Hierarchy Level | 28281
- Description | 28282
- Required Privilege Level | 28286
- Release Information | 28286

Syntax

```
source-port port-number;
```

Hierarchy Level

```
[edit security forwarding-options mirror-filter filter-name]
```

Description

Specify a Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port number to be matched for mirroring. You can specify a numeric value or one of the text synonyms listed in Table 1797 on page 28282.

Table 1797: Ports Supported by Services Interfaces

Port Name	Corresponding Port Number
afs	1483
bgp	179
biff	512
bootpc	68
bootps	67
cmd	514
cvspserver	2401
dhcp	67
domain	53
eklogin	2105
ekshell	2106
excc	512
finger	79

Table 1797: Ports Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
ftp	21
ftp-data	20
http	80
https	443
ident	113
imap	143
kerberos-sec	88
klogin	543
kpasswd	761
krb-prop	754
krbupdate	760
kshell	544
ldap	389
ldp	646
login	513

Table 1797: Ports Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
mobileip-agent	434
mobilip-mn	435
msdp	639
netbios-dgm	138
netbios-ns	137
netbios-ssn	139
nfsd	2049
nntp	119
ntalk	518
ntp	123
pop3	110
pptp	1723
printer	515
radacct	1813
radius	1812

Table 1797: Ports Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
rip	520
rkinit	2108
smtp	25
snmp	161
snmp-trap	162
snpp	444
socks	1080
ssh	22
sunrpc	111
syslog	514
tacacs	49
tacacs-ds	65
talk	517
telnet	23
tftp	69

Table 1797: Ports Supported by Services Interfaces *(Continued)*

Port Name	Corresponding Port Number
`timed`	525
`who`	513
`xdmcp`	177

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

mirror-filter (Security Forwarding Options) | **10768**

show security forwarding-options mirror-filter | **24363**

source-port (Security IDP)

IN THIS SECTION

- Syntax | **28287**
- Hierarchy Level | **28287**
- Description | **28287**
- Options | **28287**

Required Privilege Level | 28287
Release Information | 28288

Syntax

```
source-port {
    match (equal | greater-than | less-than | not-equal);
    value source-port;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol udp],

[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify the port number on the attacking device.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`source-port`*—Port number on the attacking device.
- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

source-port (Security Source NAT Rule Set)

IN THIS SECTION

- Syntax | 28288
- Hierarchy Level | 28288
- Description | 28288
- Options | 28288
- Required Privilege Level | 28289
- Release Information | 28289

Syntax

```
source-port (port-or-low <to high>);
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name match]
```

Description

Specify the port number or port range for a source rule. Up to 8 ports or port ranges are supported.

Options

port Specify a port number.

low	Specify the lower limit of the port range.
<to *high*>	Specify the upper limit of the port range.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

source-port (Security Static NAT Rule Set)

IN THIS SECTION

- Syntax | 28289
- Hierarchy Level | 28289
- Description | 28290
- Options | 28290
- Required Privilege Level | 28290
- Release Information | 28290

Syntax

```
source-port (port or low <to high>);
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name match]
```

Description

Specify the port or port range for a source rule. Up to 8 ports or port ranges are supported.

Options

port	Specify a port number.
low	Specify the lower limit of the port range.
<to *high*>	Specify the upper limit of the port range.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

source-port (Subscriber Secure Policy)

IN THIS SECTION

- Syntax | 28291
- Hierarchy Level | 28291
- Description | 28291
- Options | 28291
- Required Privilege Level | 28291
- Release Information | 28291

Syntax

```
source-port port-number;
```

Hierarchy Level

```
[edit services radius-flow-tap policy policy-name inet drop-policy rule-name from],
[edit services radius-flow-tap policy policy-name inet6 drop-policy rule-name from]
```

Description

Specify the match source port for the radius-flow-tap policy.

Options

port-number— Number of the IPv4 or IPv6 source port for the radius-flow-tap policy.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

> *Subscriber Secure Policy Overview*
>
> *Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview*

source-prefix (IDS)

IN THIS SECTION

- Syntax | 28292
- Hierarchy Level | 28292
- Description | 28292
- Options | 28292
- Required Privilege Level | 28293
- Release Information | 28293

Syntax

```
source-prefix prefix-value;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then aggregation]
```

Description

Specify the prefix length for source IPv4 address aggregation for the IDS rule. This applies session limits to an aggregation of all attacks from within a subnet of the specified length.

For example, if you configure a value of 24 for `source-prefix`, then attacks from 10.1.1.2 and 10.1.1.3 are counted as attacks from the 10.1.1/24 subnet. However, if a single host on a subnet generates a large number of network probing or flooding attacks, the flows for the entire subnet might be stopped.

Options

prefix-value—Integer value.

- **Range:** 1 through 32

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IDS Rules on an MS-DPC

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

source-prefix-ipv6 (IDS)

IN THIS SECTION

- Syntax | 28293
- Hierarchy Level | 28294
- Description | 28294
- Options | 28294
- Required Privilege Level | 28294
- Release Information | 28294

Syntax

```
source-prefix-ipv6 prefix-value;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then aggregation]
```

Description

Specify the prefix length for source IPv6 address aggregation for the IDS rule. This applies session limits to an aggregation of all attacks from within a subnet of the specified length.

For example, if you configure a value of 64 for `source-prefix-ipv6`, then attacks from 2001:db8:1234:72a2::2 and 2001:db8:1234:72a2::3 are counted as attacks from the 2001:db8:1234:72a2::/64 subnet. However, if a single host on a subnet generates a large number of network probing or flooding attacks, the flows for the entire subnet might be stopped.

Options

`prefix-value`—Integer value.

- **Range:** 1 through 128

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Configuring IDS Rules on an MS-DPC*
>
> *Configuring Protection Against Network Attacks on an MS-MPC*
>
> *Understanding IDS on an MS-MPC*

source-prefix-list (PCP)

IN THIS SECTION

- Syntax | 28295
- Hierarchy Level | 28295
- Description | 28295
- Options | 28295
- Required Privilege Level | 28296
- Release Information | 28296

Syntax

```
source-prefix-list list-name <except>;
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name from]
```

Description

Specify the source prefix list that must be matched for the PCP rule. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

list-name Source prefix list.

except (Optional) Prevent the specified prefix list from matching the PCP rule.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

source-prefix-list (Services AACL)

IN THIS SECTION

- Syntax | 28296
- Hierarchy Level | 28296
- Description | 28297
- Options | 28297
- Required Privilege Level | 28297
- Release Information | 28297

Syntax

```
source-prefix-list list-name;
```

Hierarchy Level

```
[edit services aacl rule rule-name term term-name from]
```

Description

Specify the source prefix list for rule matching. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level.

Options

`list-name`—Source prefix list.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Match Conditions in AACL Rules

source-prefix-list (Services CoS)

IN THIS SECTION

- Syntax | 28298
- Hierarchy Level | 28298
- Description | 28298
- Options | 28298
- Required Privilege Level | 28298
- Release Information | 28298

Syntax

```
source-prefix-list list-name <except>;
```

Hierarchy Level

```
[edit services cos rule rule-name term term-name from]
```

Description

Specify the source prefix list for rule matching. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level.

Options

`list-name`—Destination prefix list.

`except`—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

> *Configuring CoS Rules*
>
> https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-policy/config-guide-policy.html

source-prefix-list (Services IDS)

IN THIS SECTION

- Syntax | 28299
- Hierarchy Level | 28299
- Description | 28299
- Options | 28299
- Required Privilege Level | 28299
- Release Information | 28300

Syntax

```
source-prefix-list list-name <except>;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name from]
```

Description

Specify the source prefix list for rule matching. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level.

Options

`list-name`—Destination prefix list.

`except`—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

> *Configuring Match Conditions in IDS Rules*
>
> https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-policy/config-guide-policy.html

source-prefix-list (Services NAT)

IN THIS SECTION

- Syntax | 28300
- Hierarchy Level | 28300
- Description | 28301
- Options | 28301
- Required Privilege Level | 28301
- Release Information | 28301

Syntax

```
source-prefix-list list-name <except>;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name from]
```

Description

Specify the source prefix list for rule matching. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level.

Options

list-name—Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

> *Network Address Translation Rules Overview*
>
> https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-policy/config-guide-policy.html

source-prefix-list (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 28302
- Hierarchy Level | 28302
- Description | 28302
- Options | 28302
- Required Privilege Level | 28302

- Release Information | **28302**

Syntax

```
source-prefix-list list-name <except>;
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name term term-name from]
```

Description

Specify the source prefix list for rule matching. You configure the prefix list by including the `prefix-list` statement at the `[edit policy-options]` hierarchy level.

Options

`list-name`—Destination prefix list.

`except`—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

> *Configuring Match Conditions in Stateful Firewall Rules*

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-policy/config-guide-policy.html

source-prefix (Security Forwarding Options)

IN THIS SECTION

- Syntax | 28303
- Hierarchy Level | 28303
- Description | 28303
- Required Privilege Level | 28303
- Release Information | 28304

Syntax

```
source-prefix source-prefix;
```

Hierarchy Level

```
[edit security forwarding-options mirror-filter filter-name]
```

Description

Specify the source IP prefix or address to be matched for mirroring.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

RELATED DOCUMENTATION

mirror-filter (Security Forwarding Options) | 10768

show security forwarding-options mirror-filter | 24363

source-prefix (Services NAT)

IN THIS SECTION

- Syntax | 28304
- Hierarchy Level | 28304
- Description | 28304
- Options | 28305
- Required Privilege Level | 28305
- Release Information | 28305

Syntax

```
source-prefix source-prefix;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated]
```

Description

Specify the source prefix for translated traffic.

Options

source-prefix—IPv4 or IPv6 source prefix value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

source-prefix option enhanced to support IPv6 addresses in Junos OS Release 8.5.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

source (Protocols IGMP)

IN THIS SECTION

- Syntax | 28306
- Hierarchy Level | 28306
- Description | 28306
- Options | 28306
- Required Privilege Level | 28306
- Release Information | 28306

Syntax

```
source ip-address {
    source-count number;
    source-increment increment;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name static group
multicast-group-address],
[edit protocols igmp interface interface-name static group multicast-group-address]
```

Description

Specify the IP version 4 (IPv4) unicast source address for the multicast group being statically configured on an interface.

Options

ip-address—IPv4 unicast address.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling IGMP Static Group Membership

source (Protocols MLD)

IN THIS SECTION

- Syntax | 28307
- Hierarchy Level | 28307
- Description | 28307
- Options | 28307
- Required Privilege Level | 28308
- Release Information | 28308

Syntax

```
source ip-address {
    source-count number;
    source-increment increment;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name static group
multicast-group-address],
[edit protocols mld interface interface-name static group multicast-group-address]
```

Description

IP version 6 (IPv6) unicast source address for the multicast group being statically configured on an interface.

Options

ip-address — One or more IPv6 unicast addresses.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling MLD Static Group Membership

source (Protocols MSDP)

IN THIS SECTION

- Syntax | 28308
- Hierarchy Level | 28309
- Description | 28309
- Default | 28309
- Options | 28309
- Required Privilege Level | 28309
- Release Information | 28309

Syntax

```
source ip-address</prefix-length> {
    active-source-limit {
        maximum number;
        threshold number;
```

```
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
msdp],
[edit protocols msdp],
[edit routing-instances routing-instance-name protocols msdp]
```

Description

Limit the number of active source messages the routing device accepts from sources in this address range.

Default

If you do not include this statement, the routing device accepts any number of MSDP active source messages.

Options

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP with Active Source Limits and Mesh Groups

source-route-option (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 28310
- Hierarchy Level | 28310
- Description | 28310
- Required Privilege Level | 28310
- Release Information | 28310

Syntax

```
source-route-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop IPv4 packets that have either the IP option of 3 (Loose Source Routing) or the IP option of 9 (Strict Source Routing).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

source-routing

IN THIS SECTION

- Syntax | 28311
- Hierarchy Level | 28311
- Description | 28311
- Default | 28312
- Required Privilege Level | 28312
- Release Information | 28312

Syntax

```
source-routing {
    (ip | ipv6)
}
```

Hierarchy Level

```
[edit routing-options]
```

Description

Enable source routing.

Source routing allows a sender of a packet to partially or completely specify the route the packet takes through the network. In contrast, in non-source routing protocols, routers in the network determine the path based on the packet's destination.

> **NOTE**: We recommend that you not use source routing. Instead, we recommend that you use policy-based routing or filter-based forwarding to route packets based on source addresses.

Default

Disabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement for IPv6 introduced in Junos OS Release 8.2.

Statement for IPv4 introduced in Junos OS Release 8.5.

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Filter-Based Forwarding on the Source Address

source (Routing Instances)

IN THIS SECTION

- Syntax | 28313
- Hierarchy Level | 28313
- Description | 28313
- Options | 28313
- Required Privilege Level | 28313

Release Information | 28314

Syntax

```
source source-address {
    rate threshold-rate;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
mdt threshold group group-address],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel family intet | inet6mdt threshold group group-address],
[edit routing-instances routing-instance-name protocols pim mdt threshold group group-address],
[edit routing-instances routing-instance-name provider-tunnel family intet | inet6mdt threshold
group group-address]
```

Description

Establish a threshold to trigger the automatic creation of a data MDT for the specified unicast address or prefix of the source of multicast information.

Options

source-address—Explicit unicast address of the multicast source.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4. In Junos OS Release 17.3R1, the `mdt` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default MDT in Rosen 7, and data MDT for Rosen 6 and Rosen 7. The `provider-tunnel mdt` hierarchy is now hidden for backward compatibility with existing scripts.

RELATED DOCUMENTATION

> *Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode*
>
> *Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode*

source (Routing Instances Provider Tunnel Selective)

IN THIS SECTION

- Syntax | **28314**
- Hierarchy Level | **28315**
- Description | **28315**
- Options | **28315**
- Required Privilege Level | **28315**
- Release Information | **28315**

Syntax

```
source source-address {
    ldp-p2mp;
    pim-ssm {
        group-range multicast-prefix;
    }
    rsvp-te {
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
```

```
        static-lsp lsp-name;
    }
    threshold-rate number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group address],
[edit routing-instances routing-instance-name provider-tunnel selective group address]
```

Description

Specify the IP address for the multicast source. This statement is a part of the point-to-multipoint LSP and PIM-SSM GRE selective provider tunnel configuration for MBGP MVPNs.

Options

source-address—IP address for the multicast source.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

Configuring PIM-SSM GRE Selective Provider Tunnels

source-routing-path

IN THIS SECTION

- Syntax | 28316
- Hierarchy Level | 28317
- Description | 28317
- Options | 28317
- Required Privilege Level | 28319
- Release Information | 28319

Syntax

```
source-routing-path name {
    binding-sid binding-sid;
    color color;
    lsp-external-controller;
    metric value;
    no-ingress;
    primary name {
        compute {
            compute-profile-name;
        }
        lsp-external-controller;
        weight weight;
    }
    secondary name {
        compute {
            compute-profile-name;
        }
        weight weight;
    }
    sr-preference sr-preference;
    te-group-id ID;
    to to;
```

```
    srv6;
}
```

Hierarchy Level

```
[edit logical-systems name protocols source-packet-routing],
[edit protocols source-packet-routing]
```

Description

Configure a source-routing label switched path (LSP) for steering traffic at an ingress router. Specify a binding segment identifier from the static label range. Configure other parameters such as color, weight, preference, and segment routing preference for traffic engineering.

Starting with Junos OS Release 18.1R1, compute static noncolored segment routing LSPs for protocol SPRING-TE in an MPLS network. Configure parameters such as destination address, binding SIDs, primary segment, secondary segment, metric, and preference. These segment routing LSPs do not have a color associated with them. If an ingress route is not required for a non-colored segment routing LSP then the ingress route installation in inet.3 table can be disabled.

Options

name

Specify a name to identify a source routing path.

> **NOTE**: The combined length of tunnel-name and LSP name must not exceed 53 characters for per-path telemetry to work.

binding-sid

(Optional) Specify the binding label to enable transit functionality for this tunnel. For a non-colored static segment routing LSP, the binding SID label of protocol SPRING-TE have a default preference value of 8 and a metric of 1.

- **Range:** `16` through `1,048,576`

color

(Colored segment routing LSPs only) Specify a color identifier for the tunnel endpoint. For noncolored segment routing LSPs, you do not have to configure the color parameter.

lsp-external-controller

Enable external path computing capability for the device. See "lsp-external-controller" on page 10187 for more information.

metric

Specify metric for routes downloaded for the non-colored static segment routing tunnel.

- **Default:** 1,000,000 through 1,048,575
- **Range:** 1 through 16,777,215 (for BGP)

no-ingress

Disable ingress route that is not required for the non-colored static segment routing tunnel

primary

Specify a primary segment list for the configured source routing path.

The non-colored static segment routing LSP can have a maximum of 8 primary paths. incase of multiple operational primary paths, the PFE distributes the traffic over the paths based on the weight configured on the paths. If none of the paths have weights configured then the weights default to 1 making it an ECMP path. the paths become weighted ECMP if at least one of the paths have a non-zero weight. In both cases , when one or some of the paths fail, the PFE automatically re-balances the traffic over the remaining paths resulting in path protection.

compute *compute-profle-name*

Enable computation for the specified compute-profile. Note that if you include this option in the configuration, you don't need to explicitly configure the segment list. The segment list is dynamically created based on the specified compute-profile name.

weight *weight_value*

Specify a percentage of the bandwidth with respect to the sum of weights of all paths for the primary segment list. If forwarding interfaces are also configured with weighted ECMP, then Junos OS applies hierarchical weighted ECMP. If the weight percentage is not configured, then only IGP weights are applied on the forwarding interfaces.

secondary

Specify a secondary segment list for the configured non-colored static segment routing LSP.

sr-preference

Configure a preference for segment routing routes for traffic engineering. BGP chooses a higher preference over a lower preference value.

- **Range:** `0` through `4,294,967,295`

te-group-id

Specify the Te group ID (1..128).

to

Specify the IP address of the tunnel end-point

srv6

Specify the tunnel type SRv6.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

The `metric`, `no-ingress`, and `secondary` statements are introduced in Junos OS Release 18.1R1.

`lsp-external-controller` statement introduced in Junos OS Release 20.1R1.

`srv6` option introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

- extended-nexthop-color
- segment-list | **16521**
- source-packet-routing (Segment Routing Traffic Engineering) | **28257**
- sr-preference-override
- *Segment Routing Traffic Engineering at BGP Ingress Peer Overview*
- *Enable Segment Routing for the Path Computation Element Protocol*
- compute-profile | **4497**

source-routing-path-template

IN THIS SECTION

- Syntax | **28320**
- Hierarchy Level | **28320**
- Description | **28320**
- Options | **28320**
- Required Privilege Level | **28321**
- Release Information | **28322**

Syntax

```
source-routing-path-template name {
    bfd-liveness-detection;
    ldp-tunneling;
    metric metric;
    no-ingress;
    primary name {
        compute {
            compute-profile-name;
        }
        lsp-external-controller;
        weight weight;
    }
    secondary name;
    compute {
        compute-profile-name;
    }
    sr-preference sr-preference;
}
```

Hierarchy Level

```
[edit protocols source-packet-routing]
```

Description

Configure a source-routing path template for dynamic creation of segment routing label-switched paths (LSPs).

Options

name Name of the source routing path.

bfd-liveness-detection Configure Bidirectional Forwarding Detection (BFD) options for PCE-initiated segment routing LSP template. See "bfd-liveness-detection" on page 2955 for more information.

ldp-tunneling Allow LDP to use this LSP for tunneling. This configuration can be applied to PCE-initiated segment routing LSPs only.

lsp-external-controller Enable external path computing capability for the device. See "lsp-external-controller" on page 10187 for more information.

metric Metric for routes downloaded for this tunnel.

- **Range:** 1 through 16,777,215

no-ingress Disable ingress functionality for this tunnel.

primary *name* Configure a named identifier for the segment-list that describes the primary segment routing path along which the packet is to be routed. This segment list must have the `dynamic` statement enabled for dynamic creation of segment routing LSPs.

compute *compute-profle-name* Enable computation for the specified compute-profile.

weight *weight* (Optional) Specify the balance factor for this segment list in SR-TE tunnel.

If the forwarding interfaces have weights assigned by IGP, then hierarchical weighted ECMP is applied. When weight is not specified, only the IGP weights are applied on the forwarding interfaces.

secondary *name* Configure a named identifier for the segment-list that describes the secondary segment routing path along which the packet is to be routed. This segment list must have the `dynamic` statement enabled for dynamic creation of segment routing LSPs.

sr-preference Segment routing preference for the segment routing traffic engineered (SPRING-TE) routes. A greater value indicates higher preference.

The preference value of the routes programmed for the segment routing LSP is inherited from the `preference` *`value`* statement at the `[edit protocols source-packet-routing` hierarchy level. When this value is not configured, the default preference value of 8 is used.

- **Range:** 0 through 4,294,967,295

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.2R1.

`bfd-liveness-detection` and `ldp-tunneling` options introduced in Junos OS Release 19.4R1.

`lsp-external-controller` statement introduced in Junos OS Release 20.1R1.

`compute` statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

> *Understanding Static Segment Routing LSP in MPLS Networks*
>
> source-routing-path-template-map

source (Security Source NAT)

IN THIS SECTION

- Syntax | 28322
- Hierarchy Level | 28325
- Description | 28325
- Options | 28325
- Options | 28325
- Required Privilege Level | 28325
- Release Information | 28326

Syntax

```
source {
    address-persistent;
    interface (port-overloading off | port-overloading-factor number | enhanced-port-overloading-
algorithm);
    pool pool-name {
```

```
        address ip-address {
             to ip-address;
        }
        address-persistent subscriber ipv6-prefix-length prefix-length;
        address-pooling (paired | no-paired);
        address-shared;
        description text;
        host-address-base ip-address;
        overflow-pool (interface | pool-name);
        pool-utilization-alarm (clear-threshold value | raise-threshold value);
        port {
            block-allocation {
                active-block-timeout timeout-interval;
                block-size block-size;
                interim-logging-interval timeout-interval;
                last-block-recycle-timeout timeout-interval;
                log disable;
                maximum-blocks-per-host maximum-block-number
            }
            deterministic {
                block-size block-size;
                host {
                    address ip-address;
                    address-name address-name;
            }
            no-translation;
            port-overloading-factor number;
            range {
                port-low <to port-high>;
                to port-high;
                twin-port port-low <to port-high>;
            }
        }
        routing-instance routing-instance-name;
    }
    pool-default-port-range lower-port-range to upper-port-range;
    pool-default-twin-port-range lower-port-range to upper-port-range;
    pool-utilization-alarm (clear-threshold value | raise-threshold value);
    port-randomization disable;
    port-round-robin disable;
    port-scaling-enlargement;
    rule-set rule-set-name {
        description text;
```

```
        from {
            interface [interface-name];
            routing-instance [routing-instance-name];
            zone [zone-name];
        }
        rule rule-name {
            description text;
            match {
                application {
                    [application];
                    any;
                }
                (destination-address <ip-address> | destination-address-name <address-name>);
                destination-port (port-or-low <to high>);
                protocol [protocol-name-or-number];
                source-address [ip-address];
                source-address-name [address-name];
                source-port (port-or-low <to high>)
            }
            then source-nat
                    interface {
                        persistent-nat {
                            address-mapping;
                            inactivity-timeout seconds;
                            max-session-number value;
                            permit (any-remote-host | target-host | target-host-port);
                        }
                    }
                    off;
                    pool <pool-name>
                         persistent-nat {
                             address-mapping;
                            inactivity-timeout seconds;
                            max-session-number number;
                            permit (any-remote-host | target-host | target-host-port);
                         }
                    rule-session-count-alarm (clear-threshold value | raise-threshold value);
                }
            }
        }
        to {
            interface [interface-name];
            routing-instance [routing-instance-name];
```

```
            zone [zone-name];
        }
    }
}
```

Hierarchy Level

```
[edit security nat source pool pool-name port]
```

Description

Configure source NAT, which allows you to configure the following:

- Translate source IP address or addresses to the egress interface's IP address.
- Translate a range of source IP addresses to another range of IP addresses. This mapping is dynamic and without PAT.
- Translate a range of source IP addresses to another range of IP addresses. This mapping is dynamic and with PAT.
- Translate a range of source IP addresses to another range of IP addresses. This mapping is one-to-one, static, and without PAT.

Options

enhanced-port-overloading-algorithm	The hash algorithm uses destination IP and port to calculate the port overloading index.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. The `description` option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10. Statement modified in Junos OS Release 12.3X48-D10. Statement modified in Junos OS Release 15.1X49-D60. The `enhanced-port-overloading-algorithm` option added in Junos OS Release 22.4R1.

source (SNMP Scripts)

IN THIS SECTION

- Syntax | 28326
- Hierarchy Level | 28326
- Description | 28326
- Options | 28327
- Required Privilege Level | 28327
- Release Information | 28327

Syntax

```
source url;
```

Hierarchy Level

```
[edit system scripts snmp file filename]
```

Description

Specify the location of the master source file for an SNMP script. When you issue the `set refresh` configuration mode command at the same hierarchy level, the local copy of the script is overwritten by the version stored at the specified URL. If the `load-scripts-from-flash` statement is configured, the device refreshes the scripts on the flash drive instead of the hard drive.

Options

url The source specified as an HTTPS URL, HTTP URL, FTP URL, or scp-style remote file specification.

If you specify an HTTPS URL, you must also configure the `cert-file` statement at the same hierarchy level. Specify the path to the certificate (Root CA or self-signed) that the device will use to validate the server's certificate when you refresh the script.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configure and Use a Master Source Location for a Script

Overview of Updating Scripts from a Remote Source

cert-file (Scripts) | **3252**

refresh (SNMP Scripts) | **14255**

source (Source-Specific Multicast)

IN THIS SECTION

- Syntax | **28328**
- Hierarchy Level | **28328**
- Description | **28328**
- Options | **28328**
- Required Privilege Level | **28328**
- Release Information | **28328**

Syntax

```
source [ addresses ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast ssm-map ssm-map-name],
[edit logical-systems logical-system-name routing-options multicast ssm-map ssm-map-name],
[edit routing-instances routing-instance-name routing-options multicast ssm-map ssm-map-name],
[edit routing-options multicast ssm-map ssm-map-name]
```

Description

Specify IPv4 or IPv6 source addresses for an SSM map.

Options

addresses—IPv4 or IPv6 source addresses.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To view this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring SSM Mapping

source-threshold

IN THIS SECTION

- Syntax | 28329
- Hierarchy Level | 28329
- Description | 28329
- Options | 28329
- Required Privilege Level | 28330
- Release Information | 28330

Syntax

```
source-threshold number;
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp syn-flood]
```

Description

Specify the number of SYN segments that the device can receive per second from a single source IP address (regardless of the destination IP address and port number) before the device begins dropping connection requests from that source.

Options

number —Number of SYN segments to be received per second before the device starts dropping connection requests.

- **Range:** 4 through 500,000 per second
- **Default:** 4000 per second

NOTE: For SRX Series Firewalls the applicable range is 4 through 1,000,000 per second.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

source-tree (MBGP MVPN)

IN THIS SECTION

- Syntax | 28330
- Hierarchy Level | 28331
- Description | 28331
- Required Privilege Level | 28331
- Release Information | 28331

Syntax

```
source-tree;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvpn
static-umh],
[edit routing-instances routing-instance-name protocols mvpn static-umh]
```

Description

Specify that a statically selected upstream multicast hop (UMH) only affects type 7 (S,G) routes.

The `source-tree` option is mandatory. Type 6 routes are sent toward the rendezvous point (RP), and use the dynamic UMH selection that is configured with the `unicast-umh-election` statement, or the default method of highest IP address is used if `unicast-umh-election` is not configured.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

Example: Configuring Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

sender-based-rpf (MBGP MVPN) | **16552**

static-umh (MBGP MVPN) | **28585**

unicast-umh-election | **30353**

source-udp-port (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 28332
- Junos OS Hierarchy Level | 28332
- Junos OS Evolved Hierarchy Level | 28332
- Description | 28332
- Options | 28333
- Required Privilege Level | 28333
- Release Information | 28333

Syntax

```
source-udp-port port-number;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the UDP port of the source to be used in the UDP header for the generated frames. If you do not specify the UDP port, the default value of 4041 is used.

Options

port-number Source UDP port number for the test frames

- **Default:** 4041

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

rfc2544-benchmarking | **15759**

source-based-fpolicy-statementorwarding

IN THIS SECTION

- Syntax | **28334**
- Hierarchy Level | **28334**
- Description | **28334**
- Required Privilege Level | **28337**
- Release Information | **28337**

Syntax

```
source-based-forwarding next-hop next-table forwarding table name
```

Hierarchy Level

```
[edit policy-options policy-statement policy name term term name then]
```

Description

Source based forwarding is used to provide value added services where paths through the network are selected based on the packet's source address. In some deployments, the destination node of the packet, identified by the protocol next hop, stays the same, but the path through the network may vary.

In some deployments, the destination of the packet changes too. In this scenario, separate forwarding tables are used for traffic forwarding, and traffic is steered into these tables based on the packet's source address.

PTX routers provide a mechanism to conditionally route a packet with a specific source IP address via a specific forwarding table. When a packet arrives at the router interface, the source IP address information is retrieved from the packet, and a source lookup is performed in the forwarding table using the packet's source IP address. If a match is found in the forwarding table, the matching route in the forwarding table has the details on which forwarding table to lookup the destination IP address for this packet. Subsequently, the destination IP address and next-hop is used from the designated forwarding table to forward the packet.

This type of conditional traffic steering is done only after initial configuration. You initially configure policy options and within the policy options, under a policy statement, for a match criteria, specify the forwarding table for this route, from where the destination IP address and next-hop is retrieved to forward the packet. In the following example, a packet with BGP community *c1* is matched in the *from* clause, and the matching packet's forwarding table specified in the *then* clause.

```
policy-options {

policy-statement p1 {

    term t1 {

        from community c1;
```

```
        then {

            source-based-forwarding {

                 next-hop {
                              next-table foo.inet.0;
                         }

            }

        }

    }

}

}
```

You then include the policy in a forwarding table export policy.

```
routing-options {
 forwarding-table {
 export pl;
 }
}
```

You have to enable the interface to do a source lookup when the packet arrives at the interface - you use the `source-lookup` configuration statement to enable the interface to do a source lookup. You can enable `source-lookup` per interface, per routing protocol family, for INET or INET6. You cannot use `source-lookup` along with Source Class Usage (SCU). In the following example, `source-lookup` is enabled for an interface at the `[edit interfaces]` hierarchy of the CLI.

```
et-0/0/1 {
 unit 0 {
 family inet {
 source-lookup;
             }
```

```
        }
}
```

NOTE:

- Only the table prefixes - inet.0, inet6.0 and instance tables (including L3vVPN prefixes) are supported.
- The `source-lookup` configuration statement cannot be used alongside source-class-usage configuration statement (`set interfaces` *`interface name`* `unit` *`unit number`* `family` *`routing protocol family source-class-usage`*). But because `souce-lookup` is enabled per family for INET or INET 6, there can be a situation where INET has `source-lookup` enabled and INET6 has source-class-usage enabled - this use case is supported.
- After configuring the policy, issue `show policy` *`policy name`* to display the configured policy, like in the following example.

```
{master}[edit]
host@router# run show policy p1
Policy p1: [RESOLVED]
 Term unnamed:
   from
       route filter:
             10.0.0.8/32 orlonger
    then
       source-based-forwarding
              next-hop next-table: VRF1.inet.0
```

- You can use the `show route extensive` command to display whether a source-lookup for the route is being performed. In the following output, the `Opaque Data` field displays Source Based Forwarding details.

```
{master}[edit]
root@R11_re0# run show route 10.0.0.8/32 extensive
inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
10.0.0.8/32 (1 entry, 1 announced)
TSI:
```

```
            KRT in-kernel 10.0.0.8/32 -> {172.16.1.2}
            Opaque data client: KRT
            Opaque data: Source Based Forwarding:
             Table :VRF1.inet.0
             Nexthop Index :599
             KRT Table Kernel ID :8
            Address: 0x4e813d5d840
            Opaque-data reference count: 2
             *Static Preference: 5
             Next hop type: Router, Next hop index: 590
             Address: 0x4e813f9a61c
             Next-hop reference count: 4, Next-hop session id: 320
             Kernel Table Id: 0
             Next hop: 172.16.1.2 via ge-0/0/0.0, selected
             Session Id: 320
            State: <Active Int Ext OpaqueData>
             Age: 9:01
             Validation State: unverified
             Task: RT
             Announcement bits (1): 0-KRT
             AS path: I
             Thread: junos-main
```

Required Privilege Level

routing

routing-control

interface

interface-control

Release Information

Statement introduced in Junos OS Evolved Release 24.4R1

source-lookup

IN THIS SECTION

- Syntax | 28338
- Hierarchy Level | 28338
- Description | 28338
- Required Privilege Level | 28338
- Release Information | 28339

Syntax

```
source-lookup
```

Hierarchy Level

```
[edit interfaces interface name unit logical unit number family routing protocol family]
```

Description

On PTX platforms, enable the interface to perform a source lookup in the forwarding table using the packet's IP source address. `source-lookup` is used for `source-based-forwrding`.

> **NOTE**: The `source-lookup` configuration statement cannot be used alongside source-class-usage configuration statement (`set interfaces` *`interface name`* `unit` *`unit number`* `family` *`routing protocol family source-class-usage`*). But because `souce-lookup` is enabled per family for INET or INET 6, there can be a situation where INET has `source-lookup` enabled and INET6 has source-class-usage enabled - this use case is supported.

Required Privilege Level

routing

routing-control

interface

interface-control

Release Information

Statement introduced in Junos OS Evolved Release 24.4R1

source-routing

IN THIS SECTION

- Syntax | 28339
- Hierarchy Level | 28339
- Description | 28340
- Options | 28340
- Required Privilege Level | 28340
- Release Information | 28340

Syntax

```
source-routing {
    forward (srmpls-nodes <label hop1...n> | srv6-nodes <IPv6 SID hop1...n>);
    return (srmpls-nodes <label hop1...n> | srv6-nodes <IPv6 SID hop1...n>);
}
```

Hierarchy Level

```
[edit services monitoring twamp client control-connection name test-session name]
```

Description

TWAMP supports segment routing as defined in RFC 8402, which broadly specifies the segment routing architecture. We support two types of segment routing for TWAMP probes:

- **SR-MPLS:** Uses a list of labels, each representing a segment end node.
- **SRv6:** Uses a type 4 IPv6 routing header introduced in RFC 8754, with each segment end node represented as an IPv6 address or IPv6 segment identifier (SID).

You can specify the list of SR-MPLS or SRv6 segments for a TWAMP probe to reach the reflector. In addition, you can specify the same information for the return path from the reflector to the client. This return path information is embedded in the probe itself by using the extensions proposed in *Simple TWAMP (STAMP) Extensions for Segment Routing Networks*, draft-ietf-ippm-stamp-srpm, namely the return path TLV and its return segment list sub-TLVs, as appropriate depending on the segment routing type. The TWAMP probes are timestamped in either the Routing Engine or the Packet Forwarding Engine.

You do not have to configure anything additional on the TWAMP server to activate this feature. You can specify only the forward path, only the return path, or both paths. If configuring both paths, the forward and return paths should be of the same type.

Options

forward (srmpls-nodes <label *hop1...n*> | srv6-nodes <IPv6 SID *hop1...n*>)	Specify the forward path.
return (srmpls-nodes <label *hop1...n*> | srv6-nodes <IPv6 SID *hop1...n*>)	Specify the return path.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 24.4R1.

spam-action

IN THIS SECTION

- Syntax | 28341
- Hierarchy Level | 28341
- Description | 28341
- Options | 28341
- Required Privilege Level | 28342
- Release Information | 28342

Syntax

```
spam-action (block | tag-header | tag-subject);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-spam sbl profile profile-name]
```

Description

Configure the action to be taken by the device when spam is detected.

Options

- `block`— Block e-mail.
- `tag-header`—Tag header of e-mail.
- `tag-subject`—Tag subject of e-mail.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Example: Configuring Server-Based Antispam Filtering

Example: Configuring Local List Antispam Filtering

speed (24-port and 12-port 10 Gigabit Ethernet PIC)

IN THIS SECTION

- Syntax | 28342
- Hierarchy Level | 28343
- Description | 28343
- Options | 28343
- Required Privilege Level | 28343
- Release Information | 28344

Syntax

speed 1G | 10G

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number]
[edit chassis fpc slot-number pic pic-number port port-number]
[edit chassis lcc number fpc slot-number pic pic-number mixed-rate-mode](Routing Matrix)
```

Description

Configure the port speed on the following interface modules:

- PF-24XGE-SFPP or the PF-12XGE-SFPP PIC on a T4000 standalone router or on an LCC in a TX Matrix Plus routing matrix with 3D SIBs

  NOTE: To change the port speed from 10 Gbps to 1 Gbps on PF-24XGE-SFPP and PF-12XGE-SFPP PICs, SFP optics is required.

- P1-PTX-24-10GE-SFPP PIC on the PTX3000 router
- P1-PTX-24-10GE-SFPP PIC with the FPC2-PTX-P1A on the PTX5000 router

Dual-rate support for the P1-PTX-24-10GE-SFPP enables you to switch all port speeds to either 1 Gbps or 10 Gbps. The default is 10 Gpbs. All ports are configured to the same speed; there is no mixed-rate-mode capability. Changing the port speed causes the PIC to reboot.

To return all ports to the 10-Gbps port speed, use the `delete chassis fpc fpc-number pic pic-number speed 1G` statement. To check the port speed, use the `show interfaces` command.

NOTE: For the 1-Gbps port speed on the P1-PTX-24-10GE-SFPP PIC, you can use either the SFP-1GE-SX or the SFP-1GE-LX transceiver.

Options

1 G	1 Gbps
10 G	10 Gbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Modes of Operation

Configuring Mixed-Rate Mode Operation

mixed-rate-mode | **10774**

speed

IN THIS SECTION

- Syntax | **28344**
- Hierarchy Level | **28344**
- Description | **28345**
- Default | **28345**
- Options | **28345**
- Required Privilege Level | **28347**
- Release Information | **28347**

Syntax

```
speed (oc3-stm1 | oc12-stm4 | oc48-stm16 | 10G | 25G | 40G | 50G | 100G | 200G | 400G | 800G);
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number port port-number]
```

Description

Configure the port speed for the ports on a line card or a device with a built-in FPC.

Default

`oc3-stm1`

Options

100g	Supported ports operate at 100-Gbps speed.
10g	Supported ports operate at 10-Gbps speed.
1g	Supported ports operate at 1-Gbps speed.
200g	Supported ports operate at 200-Gbps speed.
25g	Supported ports operate at 25-Gbps speed.
40g	Supported ports operate at 40-Gbps speed.
50g	Supported ports operate at 50-Gbps speed.
400g	Supported ports operate at 400-Gbps speed.
800g	Supported ports operate at 800-Gbps speed.
oc12-stm4	Supported ports operate at OC12 or STM4 speed.
oc3-stm1	Supported ports operate at OC3 or STM1 speed.
oc48-stm16	Supported ports operate at OC48 or STM16 speed.

NOTE: You can configure the `oc12-stm4`, `oc3-stm1`, and `oc48-stm16` port speed options for SONET/SDH OC3/STM1 (Multi-Rate) MICs. However, for Channelized SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP and ATM MICs, you can configure only the `oc12-stm4` and `oc3-stm1` port speed options.

(MX Series with MPCs and ATM MICs with SFP) To configure up to OC12 CBR bandwidth speed per VC on an ATM MIC with SFP (MIC-3D-8OC3-2OC12-ATM), specify `oc12-stm4` as the speed for the specified port. You can configure the `oc12-stm4` port speed option only for ports 0 and 4 on an ATM MIC. If you

configure the `oc12-stm4` port speed option for port 0, then ports 1, 2, and 3 are disabled. Similarly, if you configure the `oc12-stm4` port speed for port 4, then ports 5, 6, and 7 are disabled.

(MX Series with MPC7E-MRATE) To configure 100-Gbps, 10-Gbps, or 40-Gbps speed per port on an MPC7E (Multi-Rate) MPC, specify `100G`, `10G`, or `40G`, respectively, as the speed for the specified port. You can configure `10G` and `40G` port speed options on all the six ports of PIC 0 and PIC 1 of an MPC7E-MRATE MPC. However, you can configure the `100G` port speed option only for ports 2 and 5 of PIC 0 and PIC 1 of an MPC7E-MRATE MPC.

(MX10003 routers with MX10003 MPC) To configure 100 Gbps, 10 Gbps, and 40 Gbps speed on all supported ports, specify 100GbE, 10GbE, or 40GbE, respectively, as the speed for the specified PIC. All the six ports of the fixed port PIC support 10-Gbps and 40-Gbps speeds. All the 12 ports of the Multi-rate MIC support 100-Gbps, 10-Gbps and 40-Gbps speeds. For more information, seeMX10003 MPC (Multi-Rate) and Supported Active Physical Ports for Configuring Port Speed to Prevent Oversubscription on MX10003 MPC.

(MX204 Routers) The MX204 has four rate-selectable ports (referred to a PIC 0 ports) that can be configured as 100-Gigabit Ethernet ports or 40-Gigabit Ethernet port, or each port can be configured as four 10-Gigabit Ethernet ports (by using a breakout cable). The MX204 also has eight 10-Gigabit Ethernet ports (referred to as PIC 1 ports). For more information, see MX204 Router Overview and Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on MX204 Router

(PTX1000 routers) The PTX1000 routers has 72 network ports that support the QSFP+ transceivers and can be configured as channelized 4x10-Gigabit Ethernet interfaces by default (for a maximum of 288 10-Gigabit Ethernet ports). When you configure the speed of a port, you must reset the FPC for the configuration to take effect. If you do not reset the FPC after modifying the port speed configuration, the router displays the following error message:

```
pic_need_skip_port_speed_change: pic-0/0: Port speed changed while previous port speed transition is pending. This PIC needs to be restarted. Please take this PIC or FPC offline and online via CLI.
```

(PTX10001-36MR routers) See *Port Speed on PTX10001-36MR Router Overview* to learn about multiple port speeds supported on PTX10001-36MR router, guidelines, and how to configure the port speed.

(PTX10003-80C and PTX10003-160C) To configure 100-Gbps, 10-Gbps, or 40-Gbps speed per port on the PTX10003-80C and PTX10003-160C line cards, specify `100g`, `10g`, or `40g`, respectively, as the speed for the specified port. Use the `number-of-sub-ports` statement to configure the number of channels for a particular port if the optics are used in a channelized mode. Use the `number-of-sub-ports` statement with the `speed` statement to configure a supported for the different channels based on the optics used.

(ACX5448-D) To configure 100-Gbps, 10-Gbps, or 40-Gbps speed per port on the ACX5448-D router, specify `100g`, `10g`, or `40g`, respectively, as the speed for the specified port.

Reboot the PIC or FPC after you change an existing port speed configuration for the new configuration to take effect. See *Port Speed* for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support for PTX3000 and PTX5000 introduced in Junos OS Release 15.1F3 and 16.1R2.

Support for MPC7E (Multi-Rate) introduced in Junos OS Release 15.1F4.

Support for PTX1000 introduced in Junos OS Release 16.1X65.

Support for MX10003 MPC introduced in Junos OS Release 17.3R1.

Support for MX204 routers introduced in Junos OS Release 17.4R1.

Speed option 10Gbps, 40Gbps, and 100Gbps introduced in Junos OS Evolved Release 19.1R1 for PTX10003-80C, PTX10003-160C routers.

Speed options 100-Gbps, 40-Gbps, 25-Gbps, and 10-Gbps introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

Speed options 800-Gbps introduced in Junos OS Release 23.4R2 for PTX10002-36QDD routers.

RELATED DOCUMENTATION

Configuring Port Speed on Multi-Rate MICs

speed (Gigabit Ethernet interface - MX Series) | 28373

speed (Chassis Cluster)

IN THIS SECTION

- Syntax | 28348
- Description | 28348
- Options | 28349

Required Privilege Level | 28349
Release Information | 28349

Syntax

```
set chassis cluster control-port speed (1g |10g);
```

Description

The SRX4600 supports three different PIC types—8-port 10-Gigabit Ethernet PIC, 4-port 40-Gigabit or 100-Gigabit Ethernet PIC, and 4-port 10-Gigabit Ethernet PIC (in a chassis cluster). Out of the four ports on the 10-Gigabit Ethernet PIC in a chassis cluster, two ports are fabric ports and the other two ports are chassis cluster control ports. The two fabric ports do not support 1-Gbps speed. Only the two control ports of the chassis cluster support a port speed of 1 Gbps. When you replace a node through return merchandise authorization (RMA), set the speed and reboot the node to activate the control-link.

On chassis cluster control interfaces, you can configure the operating speed of the 4-port 10-Gigabit Ethernet PIC from default 10-Gbps port speed to 1-Gbps port speed. You must reboot the device for the changed configuration to take effect.

The chassis cluster control interfaces do not support multiple speeds.

Following are the list of optics supported on SRX4600:

- SRX-SFP-1GE-LX
- SRX-SFP-1GE-LX-ET
- SRX-SFP-1GE-SX
- SRX-SFP-1GE-SX-ET

Autonegotiation is automatically disabled when 1-Gbps speed is configured on the interfaces.

NOTE:

- The interface name for any xe interface remains same after converting its speed from 10G to 1G.

- To view the speed configured for the interface, execute the `show interfaces extensive` command. The `Speed Configuration` field's value of `1G` or `AUTO` in the command output indicates whether the current operation speed of the interface is 1 Gbps or the default 10 Gbps, respectively.

Options

- 1g — Link speed of 1 Gbps
- 10g — Link speed of 10 Gbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1 for SRX4600.

RELATED DOCUMENTATION

speed (Gigabit Ethernet interface) | 28372

speed (Ethernet)

IN THIS SECTION

- Syntax (ACX Series, EX Series, MX Series) | 28350
- Syntax (ACX5448, ACX710) | 28350
- Syntax (ACX5448-D) | 28350
- Syntax (ACX5448-M) | 28350
- Syntax (ACX7100) | 28351
- Syntax (ACX7509) | 28351

- Syntax (EX Series) | **28351**
- Syntax (EX2300) | **28351**
- Syntax (EX4300) | **28351**
- Syntax (EX4600, OCX1100, QFX Series) | **28351**
- Syntax (QFX5100-48T) | **28351**
- Syntax (QFX5130-32CD, QFX5130E-32CD, QFX5700) | **28352**
- Hierarchy Level (ACX Series, EX Series, MX Series) | **28352**
- Hierarchy Level (EX Series) | **28352**
- Hierarchy Level (ACX5448, ACX5448-D, ACX710, ACX7100, EX2300, EX4300, EX4600, OCX Series, QFX Series, QFX5100-48T, QFX5130-32CD, QFX5700) | **28352**
- Description | **28352**
- Options | **28353**
- Required Privilege Level | **28369**
- Release Information | **28369**

Syntax (ACX Series, EX Series, MX Series)

```
speed (10m | 100m | 10g | 1g | 2.5g |5g | auto | auto-10m-100m);
```

Syntax (ACX5448, ACX710)

```
speed (10g | 25g | 40g | 100g |100m | 1g | auto);
```

Syntax (ACX5448-D)

```
speed (10g | 25g | 40g | 100g  | 100m | 1g | auto);
```

Syntax (ACX5448-M)

```
speed (10g | 25g | 40g | 100g  | 1g | auto);
```

Syntax (ACX7100)

```
speed (10g | 25g | 40g | 100g |200g | 400g);
```

Syntax (ACX7509)

```
speed (1g | 10g | 25g | 40g | 100g | 400g);
```

These are supported speeds of ACX7509-FPC-20Y, ACX7509-FPC-16C, and ACX7509-FPC-4CD line cards of ACX7509.

Syntax (EX Series)

```
speed (auto-negotiation | speed) ;
```

Syntax (EX2300)

```
speed (10m | 100m | 1g | 2.5g );
```

Syntax (EX4300)

```
speed (10m | 100m | 1g | 2.5g | 5g | 10g);
```

Syntax (EX4600, OCX1100, QFX Series)

```
speed (10g | 1g | 100m);
```

Syntax (QFX5100-48T)

```
speed (10g | 1g | 100m | auto);
```

Syntax (QFX5130-32CD, QFX5130E-32CD, QFX5700)

```
speed (10g | 25g | 40g | 50g | 100g | 200g | 400g);
```

Hierarchy Level (ACX Series, EX Series, MX Series)

```
[edit interfaces interface-name],
[edit interfaces ge-pim/0/0 switch-options switch-port port-number]
```

Hierarchy Level (EX Series)

```
[edit interfaces interface-name ether-options]
```

Hierarchy Level (ACX5448, ACX5448-D, ACX710, ACX7100, EX2300, EX4300, EX4600, OCX Series, QFX Series, QFX5100-48T, QFX5130-32CD, QFX5700)

```
[edit interfaces interface-name]
```

Description

Configure the interface speed. This statement applies to the following interfaces:

- Management Ethernet interface (`fxp0` or `em0`)
- Fast Ethernet 12-port and 48-port PICs
- Built-in Fast Ethernet port on the FIC (M7i router)
- Combo Line Rate DPCs and Tri-Rate Ethernet Copper interfaces on MX Series routers
- Gigabit Ethernet interfaces on EX Series switches

If you enable autonegotiation, then the device automatically negotiates the speed based on the speed of the other end of the link. Table 1798 on page 28353 describes the autonegotiation option available for different platforms:

Table 1798: Autonegotiation Options

Autonegotiation Option	Description
`auto-negotiation`	Automatically negotiate the speed based on the speed of the other end of the link.
`auto`	Automatically negotiate the speed (10 Mbps, 100 Mbps, or 1 Gbps) based on the speed of the other end of the link.
`auto-10m-100m`	Automatically negotiate the speed (10 Mbps or 100 Mbps) based on the speed of the other end of the link.

See *Speed and Autonegotiation* for more details.

Options

Routers support autonegotiation by default. To enable or disable autonegotiation, see "autonegotiation (Routers)" on page 2719. Table 1799 on page 28353 summarizes the speed and autonegotiation supported on different routing platforms:

Table 1799: Speed and Autonegotiation support on Routers

Router	Speed	Autonegotiation Supported (Yes/No) ?
ACX Series	• 10M • 100M • 1G • 2.5G • 5G • 10G • auto • auto-10m-100m	Yes

Table 1799: Speed and Autonegotiation support on Routers *(Continued)*

Router	Speed	Autonegotiation Supported (Yes/No) ?
ACX5448, ACX710	• 10G • 25G • 40G • 100G • 100M • 1G • auto	Yes
ACX5448-D	• 10G • 25G • 40G • 100G • 100M • 1G • auto	Yes
ACX5448-M	• 10G • 25G • 40G • 100G • 1G • auto	Yes

Table 1799: Speed and Autonegotiation support on Routers *(Continued)*

Router	Speed	Autonegotiation Supported (Yes/No) ?
ACX7100	• 10G • 25G • 40G • 100G • 200G • 400G	No
ACX7509-FPC-20Y	• 1G • 10G • 25G	Yes
ACX7509-FPC-16C	• 40G • 100G	No
ACX7509-FPC-4CD	• 400G	No

Table 1799: Speed and Autonegotiation support on Routers *(Continued)*

Router	Speed	Autonegotiation Supported (Yes/No) ?
MX Series	• 10M • 100M • 1G • 2.5G • 5G • 10G • auto • auto-10m-100m	Yes

Switches support autonegotiation by default. To enable or disable autonegotiation, see "auto-negotiation (Switches)" on page 2718. Table 1800 on page 28356 summarizes the speed and autonegotiation supported on different switching platforms:

Table 1800: Speed and Autonegotiation Support on Switches

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
EX2300-12 T or P, EX2300-24/48 T or P, and EX2300-24/48 MP	RJ-45 Interfaces • 10 Mbps • 100 Mbps • 1 Gbps • 2.5 Gbps (MP model only) • Auto (default)	Yes	NA	NA

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
	• 4 x built-in fiber SFP+ interfaces (EX2300-24MP) • 6 x built-in fiber SFP+ interfaces (EX2300-48MP)	No	Yes	No
EX3400-24/48 T or P	RJ-45 Interfaces • 10 Mbps • 100 Mbps • 1 Gbps • Auto (default)	Yes	NA	NA
	4 x built-in SFP+ interfaces (Network default)	No	Yes	No
	2 x built-in QSFP+ interfaces (VCP default) • 40 Gbps	Yes	Yes	No

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
EX4300-24/48 T or P and EX4300-48 MP	RJ-45 interfaces • 10 Mbps (EX4300-24/48 T or P only) • 100 Mbps • 1 Gbps • 2.5 Gbps (MP models only; 24 maximum) • 5 Gbps (MP models only; 24 maximum) • 10 Gbps (MP models only; 24 maximum) • Auto (default)	Yes	NA	NA
	4 x built-in VCP (default) QSFP+ interfaces • 40 Gbps	No	Yes	No
EX4300-32F	32 x 1 Gbps (SFP ports)	Yes	Yes	No
	8 x 1/10 Gbps (SFP+ ports)	No	Yes	No

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
	2 x 40 Gbps (QSFP+ ports)	No **NOTE**: If 40G is a DAC and used as network port, AN is enabled. For all other cases, AN is not enabled.	Yes	No
EX4300 Uplink Modules				
EX-UM-4X4SFP (Non-MP Model)	4 x SFP+ interfaces	No	Yes	No
EX-UM-4SFPP-MR (MP Model)	4 x SFP+ interfaces	No	Yes	No
EX-UM-2QSFP-MR (MP Model)	2 x QSFP+ interfaces • 40 Gbps • 100 Gbps • Channelized into 4x10 Gbps • Channelized into 4x25 Gbps	No	Yes	No
EX4100-F 12 P/T (No Macsec)	RJ-45 interfaces (PIC0) • 10 Mbps • 100 Mbps • 1000 Mbps	Yes	NA	NA

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
	RJ-45 interfaces (PIC2) • 100 Mbps • 1000 Mbps • 1 Gbps • 2.5 Gbps • 5 Gpps • 10 Gbps	Yes	NA	NA
	Built-in SFP+ interfaces (PIC1) (VCP default) • 1 Gbps (Network mode only) • 10 Gbps	No	Yes	No

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
EX4100 24/48 T or P and EX4100-24/48 MP **NOTE**: MACsec is supported only in Multigigabit ports and 4x10 GbE uplink module.	RJ-45 interfaces • 10 Mbps • 100 Mbps • 1 Gbps • 2.5 Gbps (MP model only) • 5 Gbps (24MPmodel only) • 10 Gbps (24MP model only) • Auto	Yes	NA	NA
	4 x built-in VCP SFP28 Interfaces (VCP Default) • 1 Gbps (only in network mode) • 10 Gbps • 25 Gbps (default)	No	Yes	No
	4 x built-in SFP+ interfaces (Network default) • 100 Mbps • 1 Gbps • 10 Gbps	No	Yes	No

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
EX4100-F 24/48 T or P (no MACsec)	RJ-45 interfaces • 10 Mbps • 100 Mbps • 1 Gbps • Auto	Yes	NA	NA
	Built-in VCP SFP+ interfaces • 1 Gbps (only in network mode) • 10 Gbps	No	Yes	No
	Fiber SFP+ Interfaces • 100 Mbps • 1 Gbps • 10 Gbps	No	Yes	No

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
EX4400 24/48 T or P and EX4400-24/48 MP (has MACsec support on all interfaces except built-in 100 GbE ports)	RJ-45 interfaces • 10 Mbps ((EX4400 24/48 T or P only) • 100 Mbps • 1 Gbps • 2.5 Gbps (MP Model only) • 5 Gbps (MP Model only) • 10 Gbps (MP Model only) • Auto	Yes	NA	NA

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
	2 x 100 GbE QSFP28 interfaces (No MACsec; by default VCP) • 100 Gbps • 40 Gbps • Channelized 4 x 10 Gbps; supported for 40 GbE when changed to Network Interface from VCP. • Channelized 4x25 Gbps; supported for 100 GbE when changed to Network Interface from VCP	NA	Yes	Yes
EX4400-48 F	• 10 Mbps (SFP-T on 0-35 ports) • 100 Mbps (SFP-T on 0-35 ports) • 36x1 Gbps (SFP) • 12x1/10 Gbps (SFP+)	Yes (for 36 PIC0 ports)	Yes	Yes (22.3R1 onwards) **NOTE**: HGoE is not supported on PIC0 ports.

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
	2 x 100 GbE QSFP28 interfaces (No MACsec; by default VCP) • 40 Gbps • 100 Gbps • Channelized 4 x 10 Gbps; supported for 40 Gbps when changed to Network Interface from VCP. • Channelized 4x25 Gbps; supported for 100 Gbps when changed to Network Interface from VCP	No	Yes	Yes (22.3R1 onwards)
EX4400-24X	• 24x1/10 GbE (SFP+)	No	Yes	NA

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
	2 x 100 GbE QSFP28 interfaces (Supports MACsec; by default VCP) • 100 Gbps • 40 Gbps • Channelized 4 x 10 Gbps; supported for 40 Gbps when changed to Network Interface from VCP. • Channelized 4x25 GbE; supported for 100 GbE when changed to Network Interface from VCP	No	Yes	Yes (23.1R1 onwards)
EX4400-EM-4S (No MACsec Support)	4 x SFP+ interfaces	No	No	NA
EX4400-EM-4Y (has MACsec Support)	4 x SFP28 Network interfaces • 1 Gbps (23.1R1 onwards) • 10 Gbps • 25 Gbps (default)	No	No	Yes (22.3R1 onwards)

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
EX4400-EM-1C (has MACsec Support)	1 x QSFP28 Network Interface • 100 Gbps • 40 Gbps • Channelized 4 x 10 Gbps supported for 40 Gbps (supports HGoE) • Channelized 4x25 Gbps supported for 100 Gbps (supports HGoE)	No	Yes	Yes (23.1R1 onwards)
EX4600-40F	24 x 1/10 Gbps (SFP+)	No	Yes	No
	4 x 40 Gbps (QSFP+)	No	Yes	No
OCX1100, QFX Series, QFabric	• 100 Mbps • 1 Gbps • 10 Gbps	No		No
QFX5100-48T	• 100 Mbps • 1 Gbps • 10 Gbps • Auto	Yes		No

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
QFX5110-48S (connected with QFX-SFP-1GE-T transceiver)	• 100 Mbps • 1 Gbps • 10 Gbps • 40 Gbps • 100 Gbps • Auto	Yes		No
QFX5110-48S and QFX5100-48S (connected with JNP-SFPP-10GE-T transceiver)	• 100 Mbps • 1 Gbps • 10 Gbps	Yes		No
QFX5120-48Y (connected with JNP-SFPP-10GE-T transceiver)	• 1 Gbps • 10 Gbps	No		No
QFX5130-32CD, QFX5130E-32CD	• 400 Gbps • 200 Gbps • 100 Gbps • 50 Gbps • 40 Gbps • 25 Gbps • 10 Gbps	No		No

Table 1800: Speed and Autonegotiation Support on Switches *(Continued)*

Switch	Speed	Autonegotiation Support (Yes/No/NA)?	Autosense Support (Yes/No/NA)	HGoE Support (Yes/No/NA)?
QFX5700	• 400 Gbps • 200 Gbps • 100 Gbps • 50 Gbps • 40 Gbps • 25 Gbps • 10 Gbps	No		No

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Speed option 2.5Gbps introduced in Junos OS Release 18.1R2 for EX2300 switch.

Speed option 10Gbps and 5Gbps introduced in Junos OS Release 18.2R1 for EX4300 switch.

Speed option 1-Gbps is introduced in Junos OS Release 19.1R1 on the 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module on EX4300-48MP switches.

Speed options 100-Mbps, 1-Gbps, and auto is introduced in Junos OS Releases 18.4R1S2, 18.4R2, and 19.2R1 and later for ACX5448 Universal Metro Routers.

Speed option 10Gbps, 40Gbps, and 100Gbps introduced in Junos OS Evolved Release 19.1R1 for PTX10003-80C, PTX10003-160C routers.

Speed options 100-Gbps, 40-Gbps, 25-Gbps, and 10-Gbps introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

RELATED DOCUMENTATION

Speed and Autonegotiation

Port Settings

Configuring the Interface Speed

Configuring the Interface Speed on Ethernet Interfaces

Configuring Gigabit Ethernet Autonegotiation

Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support

auto-negotiation (QFX Series) | **2718**

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

Junos OS Network Interfaces Library for Routing Devices

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

Configuring Gigabit Ethernet Interfaces (J-Web Procedure)

Junos OS Ethernet Interfaces Configuration Guide

speed (Fibre Channel Interfaces)

IN THIS SECTION

- Syntax | **28371**
- Hierarchy Level | **28371**
- Description | **28371**
- Options | **28371**
- Required Privilege Level | **28371**
- Release Information | **28371**

Syntax

```
speed (auto-negotiation | 2g | 4g | 8g);
```

Hierarchy Level

```
[edit interfaces interface-name                fibrechannel-options]
```

Description

Configure FC interface speed.

Options

auto-negotiation—Automatically negotiate interface speed to match the speed of the attached link (2 Gbps, 4 Gbps, 8 Gbps).

2g—2 Gbps link speed

4g—4 Gbps link speed

8g—8 Gbps link speed

- **Default: auto-negotiation**

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show fibre-channel interfaces | **19980**

Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric

Converting an Ethernet Interface To a Fibre Channel Interface

Configuring a Physical Fibre Channel Interface

Understanding Interfaces on an FCoE-FC Gateway

speed (Gigabit Ethernet interface)

IN THIS SECTION

- Syntax | 28372
- Hierarchy Level | 28372
- Description | 28372
- Options | 28373
- Required Privilege Level | 28373
- Release Information | 28373

Syntax

```
speed (1g |10g);
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options]
```

Description

Configure the operating speed of the 8-port 10-Gigabit Ethernet PIC from default 10-Gbps port speed to 1-Gbps port speed. Each of the interfaces in the 8-port 10-Gigabit Ethernet PIC can be independently configured to 1Gbps or 10Gbps speeds. For information about platforms support, see hardware compatibility tool (HCT).

Autonegotiation is automatically disabled when 1-Gbps speed is configured on the interfaces.

On 1/10-Gbps capable Gigabit Ethernet SFP interfaces, the duplex is always full and the speed matches that of the inserted optic. These interfaces support either 1-Gbps or 10-Gbps SFP optics. For SRX Series Firewalls, the display and configuration is always xe- only, even if a 1G optic is inserted. The xe- value is used to denote that the interface is 10G capable. If a 1G optic is used, show commands for the interface will display the correct speed, but the config will always show as xe-. If a speed configuration is changed, you cannot change it again in the next 180 seconds. The interface link might drop down, if you try to change the speed configuration again within 180 seconds of the first speed configuration change. The 8x10-Gbps ports supports multiple port speeds, that is, some ports operates at 10G speed and some at 1G speed. To view the speed configured for the interface, execute the `show interfaces extensive` command. The `Speed Configuration` field's value of `1G` or `AUTO` in the command output indicates whether the current operation speed of the interface is 1 Gbps or the default 10 Gbps, respectively.

Options

- 1g — Link speed of 1 Gbps
- 10g — Link speed of 10 Gbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

speed (Chassis Cluster) | 28347

speed (Gigabit Ethernet interface - MX Series)

IN THIS SECTION

Syntax | 28374

- Hierarchy Level | 28374
- Description | 28374
- Default | 28375
- Options | 28375
- Required Privilege Level | 28376
- Release Information | 28376

Syntax

```
speed (1G | 10G);
```

Hierarchy Level

```
[edit interfaces intf-name gigether-options]
```

Description

(MX10003, MX204, MX10008, and MX10016 routers) the 10-Gbps port can operate in 1-Gbps mode also. When a port operates in 10-Gbps speed, you can change the operating speed to 1-Gpbs using the configuration `speed 1G` in this configuration statement. After you commit this configuration, the operating speed of the 10-Gbps port changes to 1-Gbps speed.

On fixed-port PIC and non-MACsec MIC of MX10003 router, you can configure any or all 10-Gbps port operating in 4X10-Gbps speed to operate in 1-Gbps speed. On MX204 routers, you can configure the 4X10-Gbps port on one of the fixed-port PICs to operate in 1-Gbps mode. And on the other fixed-port PIC, you can configure the 10-Gbps port to 1-Gbps speed. On MX10003 and MX204 routers, 1-Gbps speed is supported with `speed 1g` configuration. You can configure auto-negotiation by using the command `set interfaces interface-name gigether-options auto-negotiation`. For more information, see "auto-negotiation" on page 2719.

On JNP10K-LC2101 MPC, you can configure one or all 10-Gbps ports operating in 4x10-Gbps speed to operate in 1-Gbps speed. Also, autonegotiation is supported when the interface speed is configured for 1-Gbps speed.

Starting in Junos OS Release 20.4R1, you can configure 1-Gbps speed on all 40 10-Gbps Ethernet ports of the MPC7E-10G line cards.

As of Junos OS Release 23.1R1, the MX304 supports 1G fibre optic speeds in addition to the previously supported speeds of 400G, 100G, 40G, 25G, and 10G.

NOTE:

- On the MX10003 router, the MACsec MIC does not provide 1-Gbps speed. If you attempt to change the operating speed to 1-Gbps, syslog displays that this feature is not supported on the MACsec MIC.
- (MX10003, MX204, MX10008, and MX10016 routers) Rate selectability at PIC level and port level does not support 1-Gbps speed.
- (MX10003, MX204, MX10008, and MX10016 routers) The interface name prefix is **xe**.
- (MX10003, MX204, MX10008, and MX10016 routers) After configuring 1-Gbps speed, the protocol continues to advertise the bandwidth as 10-Gigabit Ethernet.
- On MX10003 and MX204 routers, Link Aggregation Group (LAG) is supported on 10-Gbps speed only. It is not supported on 1-Gbps speed.
- On MX304 for 1G interface, the maximum MTU supported is 2000 bytes. Hence, the maximum MRU supported is 2008 bytes.
- On MX304, 1G interface does not support MACsec in this release.

To view the speed configured for the interface, execute the " show interfaces extensive" on page 20676 command. The `Speed Configuration` output parameter in the command output indicates the current operation speed of the interface. If the interface is configured with 1-Gbps speed, then `Speed Configuration` displays 1G; if the interface is configured with 10-Gbps speed, Speed Configuration displays AUTO.

Default

`10G`

Options

`1G`—Supported ports operate at 1-Gbps speed.

`10G`—Supported ports operate at 10-Gbps speed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Port Speed

speed (Interfaces)

IN THIS SECTION

- Syntax | 28376
- Hierarchy Level | 28376
- Description | 28377
- Options | 28377
- Required Privilege Level | 28377
- Release Information | 28377

Syntax

```
speed (100m |10m | 1g);
```

Hierarchy Level

```
[edit interfaces interface-name speed]
```

Description

Configure the operating speed for the 2-Port 10 Gigabit Ethernet XPIM.

Options

- 100m — Link speed of 100 Mbps
- 10g — Link speed of 10 Gbps
- 10m — Link speed of 10 Mbps
- 1g — Link speed of 1 Gbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Ethernet Interfaces

Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface

speed (MX Series DPC)

IN THIS SECTION

- Syntax | 28378
- Hierarchy Level | 28378
- Description | 28378
- Options | 28378

- Required Privilege Level | 28378
- Release Information | 28379

Syntax

```
speed (auto | 1Gbps | 100Mbps | 10Mbps);
```

Hierarchy Level

```
[edit interfaces ge-fpc/pic/port]
```

Description

On MX Series routers with Combo Line Rate DPCs and Tri-Rate Copper SFPs you can set auto negotiation of speed. To specify the auto negotiation speed, use the `speed (auto | 1Gbps | 100Mbps | 10Mbps)` statement under the `[edit interface ge-/fpc/pic/port]` hierarchy level. The `auto` option will attempt to automatically match the rate of the connected interface. To set port speed negotiation to a specific rate, set the port speed to `1Gbps`, `100Mbps`, or `10Mbps`.

> **NOTE**: If the negotiated speed and the interface speed do not match, the link will not be brought up. Half duplex mode is not supported.

Options

You can specify the speed as either `auto` (autonegotiate), `10Mbps` (10 Mbps), `100Mbps` (100 Mbps), or `1Gbps` (1 Gbps).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Autonegotiation

no-auto-mdix | 11409

spf-options (Protocols IS-IS)

IN THIS SECTION

- Syntax | 28379
- Hierarchy Level | 28379
- Description | 28380
- Options | 28380
- Required Privilege Level | 28381
- Release Information | 28381

Syntax

```
spf-options {
    delay milliseconds;
    holddown milliseconds;
    rapid-runs number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
```

```
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Configure options for running the shortest-path-first (SPF) algorithm.

Running the SPF algorithm is usually the beginning of a series of larger system-wide events. For example, the SPF algorithm can lead to interior gateway protocol (IGP) prefix changes, which then lead to BGP nexthop resolution changes. Consider what happens if there are rapid link changes in the network. The local routing device can become overwhelmed. This is why it sometimes makes sense to throttle the scheduling of the SPF algorithm.

You can configure the following SPF options:

- The delay in the time between the detection of a topology change and when the SPF algorithm actually runs.
- The maximum number of times that the SPF algorithm can run in succession before the hold-down timer begins.
- The time to hold down, or wait, before running another SPF calculation after the SPF algorithm has run in succession the configured maximum number of times.

If the network stabilizes during the hold-down period and the SPF algorithm does not need to run again, the system reverts to the configured values for the `delay` and `rapid-runs` statements.

Options

`delay` *`milliseconds`*—Time interval between the detection of a topology change and when the SPF algorithm runs.

- **Range:** 50 through 1000 milliseconds
- **Default:** 200 milliseconds

`holddown` *`milliseconds`*—Time interval to hold down, or wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession.

- **Range:** 2000 through 10,000 milliseconds
- **Default:** 5000 milliseconds

`rapid-runs` *number*—Maximum number of times the SPF algorithm can run in succession. After the maximum is reached, the holddown interval begins.

- **Range:** 1 through 5
- **Default:** 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Loop-Free Alternate Routes for IS-IS

Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN

spf-options (Protocols OSPF)

IN THIS SECTION

- Syntax | 28382
- Hierarchy Level | 28382
- Description | 28382
- Options | 28383
- Required Privilege Level | 28383
- Release Information | 28383

Syntax

```
spf-options {
    delay milliseconds;
    holddown milliseconds;
    rapid-runs number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf topology (default | ipv4-multicast |
name)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
topology (default | ipv4-multicast | name)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf topology (default | ipv4-multicast | name)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf topology (default | ipv4-multicast |
name)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Configure options for running the shortest-path-first (SPF) algorithm. You can configure the following:

- A delay for when to run the SPF algorithm after a network topology change is detected.
- The maximum number of times the SPF algorithm can run in succession.

- A hold-down interval after the SPF algorithm runs the maximum number of times. If the network stabilizes during the holddown period and the SPF algorithm does not need to run again, the system reverts to the configured values for the **delay** and `rapid-runs` statements.

Running the SPF algorithm is usually the beginning of a series of larger system-wide events. For example, the SPF algorithm can lead to interior gateway protocol (IGP) prefix changes, which then lead to BGP nexthop resolution changes. Consider what happens if there are rapid link changes in the network. The local routing device can become overwhelmed. This is why it sometimes makes sense to throttle the scheduling of the SPF algorithm.

Options

delay ***milliseconds***—Time interval between the detection of a topology change and when the SPF algorithm runs.

- **Range:** 50 through 8000 milliseconds
- **Default:** 200 milliseconds

holddown ***milliseconds***—Time interval to hold down, or to wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession.

- **Range:** 2000 through 20,000 milliseconds
- **Default:** 5000 milliseconds

rapid-runs ***number***—Maximum number of times the SPF algorithm can run in succession. After the maximum is reached, the hold down interval begins.

- **Range:** 1 through 10
- **Default:** 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Support for Multitopology Routing introduced in Junos OS Release 9.0.

Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring SPF Algorithm Options for OSPF

Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths

Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Understanding Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Understanding Multitopology Routing in Conjunction with PIM

spi

IN THIS SECTION

- Syntax | **28384**
- Hierarchy Level | **28385**
- Description | **28385**
- Options | **28385**
- Required Privilege Level | **28385**
- Release Information | **28385**

Syntax

```
spi spi-value;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]
```

Description

Configure the SPI for an SA.

Options

spi-value—An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).

- **Range:** 256 through 16,639

> **NOTE**: Use the auxiliary SPI when you configure the `protocol` statement to use the `bundle` option.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Security Associations

spi (Junos-FIPS Software)

IN THIS SECTION

- Syntax | 28386
- Hierarchy Level | 28386
- Description | 28386
- Options | 28386
- Required Privilege Level | 28386
- Release Information | 28387

Syntax

```
spi spi-value;
```

Hierarchy Level

```
[edit security ipsec internal security-association manual direction]
```

Description

The security parameter index (SPI) value used for the internal Routing Engine-to-Routing Engine IPsec security association (SA) configuration.

Options

spi-value—Integer to use for this SPI.

- **Range:** 256 through 16,639

Required Privilege Level

Crypto Officer—To add and view this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring IPsec for Enabling Internal Communications Between Routing Engines for Junos OS in FIPS Mode*

spi (Junos OS)

IN THIS SECTION

- Syntax | 28387
- Hierarchy Level | 28387
- Description | 28388
- Options | 28388
- Required Privilege Level | 28388
- Release Information | 28388

Syntax

```
spi spi-value;
```

Hierarchy Level

```
[edit security ipsec security-association sa-name manual direction
(inbound | outbound | bi-directional)]
```

Description

Configure the security parameter index (SPI) for a security association (SA).

Options

spi-value—An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).

- **Range:** 256 through 16,639

> **NOTE**: Use the auxiliary SPI when you configure the `protocol` statement to use the `bundle` option.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

spid1

IN THIS SECTION

- Syntax | 28389
- Hierarchy Level | 28389
- Description | 28389
- Options | 28389
- Required Privilege Level | 28389
- Release Information | 28389

Syntax

```
spid1 spid1-string;
```

Hierarchy Level

```
[edit interfaces br-pim/0/port isdn-options]
```

Description

Configure the Service Profile Identifier (SPID).

Options

spid1-string—Numeric SPID.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

spid2

IN THIS SECTION

- Syntax | 28390
- Hierarchy Level | 28390
- Description | 28390
- Options | 28390

Required Privilege Level | 28390
Release Information | 28390

Syntax

```
spid2 spid2-string;
```

Hierarchy Level

```
[edit interfaces br-pim/0/port isdn-options]
```

Description

Configure an additional SPID.

Options

spid2-string—Numeric SPID.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

splitting-merging

IN THIS SECTION

- Syntax | 28391
- Hierarchy Level | 28391
- Description | 28391
- Options | 28392
- Required Privilege Level | 28393
- Release Information | 28393

Syntax

```
splitting-merging {
    maximum-member-lsps number;
    maximum-signaling-bandwidth bps;
    merging-bandwidth bps;
    minimum-member-lsps number;
    minimum-signaling-bandwidth bps;
    normalization;
    sampling;
    splitting-bandwidth bps;
    splitting-merging-threshold percent;
    no-normalize-same-members;
}
```

Hierarchy Level

```
[edit protocols mpls container-label-switched-path lsp-name]
```

Description

Perform splitting and merging.

Options

maximum-member-lsps *number*

Number of label-switched paths (LSPs) that a container LSP can have as member LSPs at maximum.

- **Default:** 1

maximum-signaling-bandwidth *bandwidth*

Amount of bandwidth in bits per second (bps) that can be signaled for an LSP at maximum after normalization. When `maximum-signaling-bandwidth` is not configured, the value is derived from the `splitting-bandwidth`.

When auto-bandwidth adjustment is done between two normalization events, per LSP auto-bandwidth configuration and thresholds are used instead of the `splitting-bandwidth`.

- **Default:** 1 bps

merging-bandwidth *bandwidth*

Amount of bandwidth in bits per second (bps) that is used for merging during normalization.

- **Default:** 1 bps

minimum-member-lsps *number*

Number of LSPd that a container LSP can have as member LSPs at minimum.

- **Default:** 64

minimum-signaling-bandwidth *bandwidth*

Amount of bandwidth in bits per second (bps) that can be signaled for an LSP at minimum after normalization. When `minimum-signaling-bandwidth` is not configured, the value is derived from the `merging-bandwidth`.

When auto-bandwidth adjustment is done between two normalization events, per LSP auto-bandwidth configuration and thresholds are used instead of the `merging-bandwidth`.

- **Default:** 1 bps

splitting-bandwidth *bandwidth*

Amount of bandwidth in bits per second (bps) that can be used for splitting during normalization.

- **Default:** 1 bps

splitting-merging-threshold *percent*

Percentage changes in aggregate bandwidth relevant for splitting and merging.

- **Default:** 0%

no-normalize-same-members

Do not normalize when we are not splitting or merging member LSPs .

If normalization does not need to create or delete any member LSPs, we re-signal the member LSPs with equal bandwidth. Enabling `no-normalize-same-members` statement, we avoid such normalizations. .

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced for QFX Series switches in Junos OS Release 15.1X53-D30.

RELATED DOCUMENTATION

container-label-switched-path | 4577

spring-te

IN THIS SECTION

- Syntax | 28394
- Hierarchy Level | 28394
- Description | 28394
- Options | 28394
- Required Privilege Level | 28395

Syntax

```
spring-te {
             apply-groups
             apply-groups-except
             groups
             color
             discriminator
             endpoint
             origin-addr
             origin-as
             protocol-origin
                   }
```

Hierarchy Level

```
[edit policy-options policy-statement name from traffic-engineering]
```

Description

Specify SPRING-TE policy related parameters.

Options

color Color of the SPRING-TE policy candidate path

discriminator Discriminator of the SPRING-TE policy candidate path

endpoint Endpoint of the SPRING-TE policy candidate path

origin-addr Address of the node originating the SPRING-TE policy candidate path

origin-as AS number of the node creating the SPRING-TE policy candidate path

protocol-origin Protocol originating the SPRING-TE policy candidate path

Required Privilege Level

routing

RELATED DOCUMENTATION

Link-State Distribution using BGP

spring-te (Dynamic Tunnels)

IN THIS SECTION

- Syntax | 28395
- Hierarchy Level | 28395
- Description | 28396
- Options | 28396
- Required Privilege Level | 28396
- Release Information | 28396

Syntax

```
spring-te {
    destination-networks name;
    source-routing-path-template name {
        (color [ color ... ] | color-any);
    }
}
```

Hierarchy Level

```
[edit routing-options dynamic-tunnels]
```

Description

Enable next hop-base dynamic tunnel mode.

Options

source-routing-path-template *template-name*
Configure template color mapping for segment routing traffic-engineered (SPRING-TE) dynamic LSP parameters.

color
Specify set of color list to be mapped to corresponding SPRING-TE template.

- When enabled, all templates should have color objects defined.
- All destinations are assumed to be colored as well.
- A color can be mapped to only one template at a given point in time.
- Both the `color` and `color-any` statements can coexist. When the two statements are enabled together, preference is given to a specific color template than color-any.
- A colored and non-colored destination cannot co-exist in the same SR-TE configuration.

color-any
Enables mapping of any color to corresponding SPRING-TE template. Only one color-any template can be configured for one SR-TE LSP.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Understanding Static Segment Routing LSP in MPLS Networks

spring-te-lsp-next-hop

IN THIS SECTION

- Syntax | 28397
- Hierarchy Level | 28397
- Description | 28397
- Options | 28398
- Required Privilege Level | 28398
- Release Information | 28398

Syntax

```
spring-te-lsp-next-hop{
    <nexthop>
    lsp-source {
        bgp;
        dtm;
        pcep;
        static;
        }
}
```

Hierarchy Level

```
[edit logical-systems name routing-options static route destination],
[edit logical-systems name routing-options rib rib name static route destination]
[edit routing-options static route destination],
[edit routing-options rib rib name static route destination]
```

Description

The configuration allows you to resolve static route over SR-TE tunnels.

Options

nexthop—Name of the SR-TE tunnel to reach destination address.

lsp-source—Source of tunnel creation and the following source types are currently supported. If the source is not configured, by default, it takes the next hop as static. So you must configure the source type if LSP is dtm, bgp, or pcep.

- **bgp**
- **dtm**
- **pcep**
- **static**

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

static (Routing Options) | 28570

spring-traffic-engineering-edit-protocols

IN THIS SECTION

Syntax | 28399
Hierarchy Level | 28399
Description | 28399
Options | 28399

Required Privilege Level | 28399
Release Information | 28399

Syntax

```
spring-traffic-engineering {
lsp-external-controller controller-name;
preference;
traceoptions;
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Enable traffic engineering (TE) for Source Packet Routing in Networking (SPRING).

Options

preference (Optional) Configure the preference value for the traffic engineered label-switched paths (LSPs). Default: 8

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2 on MX Series routers.

spt-only

IN THIS SECTION

- Syntax | 28400
- Hierarchy Level | 28400
- Description | 28400
- Required Privilege Level | 28400
- Release Information | 28400

Syntax

```
spt-only;
```

Hierarchy Level

```
[edit logical-systems profile-name routing-instances instance-name protocols mvpn mvpn-mode],
[edit routing-instances instance-name protocols mvpn mvpn-mode]
```

Description

Set the MVPN mode to learn about active multicast sources using multicast VPN source-active routes. This is the default mode.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs

spt-threshold

IN THIS SECTION

- Syntax | 28401
- Hierarchy Level | 28401
- Description | 28401
- Required Privilege Level | 28402
- Release Information | 28402

Syntax

```
spt-threshold {
    infinity [ policy-names ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Set the SPT threshold to infinity for a source-group address pair. Last-hop multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or an SPT rooted at the source. By default, last-hop routing devices transition

to a direct SPT to the source. You can configure this routing device to set the SPT transition value to infinity to prevent this transition for any source-group address pair.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Example: Configuring the PIM SPT Threshold Policy

sql

IN THIS SECTION

- Syntax | 28402
- Hierarchy Level | 28403
- Description | 28403
- Options | 28403
- Required Privilege Level | 28403
- Release Information | 28403

Syntax

```
sql {
    disable;
```

```
    traceoptions {
        flag {
            all <extensive>;
        }
    }
}
```

Hierarchy Level

```
[edit security alg]
[edit tenants tenant-name security alg]
```

Description

Specify the Oracle SQL ALG on the device.

Options

- `disable`—Disable the SQL ALG. By default, the SQL ALG is enabled.
- `traceoptions`—Configure SQL ALG tracing options.
  - `flag`—Trace operation to perform.
    - `all`—Trace all events.
    - *`extensive`*—Display extensive amount of data.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding the SQLNET ALG

squelch

IN THIS SECTION

- Syntax (ACX710: ESMC Hierarchy Level) | 28404
- Syntax (ACX710: Output Hierarchy Level) | 28404
- Hierarchy Level | 28405
- Description | 28405
- Options | 28405
- Required Privilege Level | 28406
- Release Information | 28406

Syntax (ACX710: ESMC Hierarchy Level)

```
squelch {
    method ql-dnu;
    quality-level quality-level;
}
```

Syntax (ACX710: Output Hierarchy Level)

```
squelch {
    method (ais | ql-dnu | rai | shut);
}
```

Hierarchy Level

```
[edit chassis synchronization esmc-transmit interfaces]
[edit chassis synchronization output interfaces bits]
```

Description

Squelch the transmitting signal.

Options

method Specify the method to squelch the transmitting signal. Choose from the following options:

- ais—Transmits an alarm indication signal (AIS).
- ql-dnu—Transmits a Synchronous Digital Hierarchy (SDH) synchronization status message (SSM) or a Synchronous Optical Networks (SONET) SSM.
- rai—Transmits a remote alarm indication (RAI).
- shut—Shuts down the output.

quality-level Specify the quality level below which `squelch` is activated. The available quality levels are given in Table 1801 on page 28405.

Table 1801: Quality Levels

Quality Level	Description
prc	Timing quality of a primary reference clock (option-1 only).
prs	Clock traceable to a primary reference source (option-2 only).
sec	Timing quality of an SDH equipment clock (option-1 only).
smc	Clock traceable to a self-timed SONET (option-2 only).
ssu-a	Timing quality of a type I or IV timeReceiver clock (option-1 only).

Table 1801: Quality Levels *(Continued)*

Quality Level	Description
ssu-b	Timing quality of a type VI timeReceiver clock (option-1 only).
st2	Clock traceable to Stratum 2 (option-2 only).
st3	Clock traceable to Stratum 3 (option-2 only).
st3e	Clock traceable to Stratum 3E (option-2 only).
st4	Clock traceable to Stratum 4 free-run (option-2 only).
stu	Clock traceable to an unknown quality (option-2 only).
tnc	Clock traceable to a transit node clock (option-2 only).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

show chassis synchronization | 18879

srgb (Protocols BGP)

IN THIS SECTION

- Syntax | 28407
- Hierarchy Level | 28407
- Description | 28407
- Options | 28407
- Required Privilege Level | 28408
- Release Information | 28408

Syntax

```
srgb index-range index-range start-label start-label;
```

Hierarchy Level

```
[edit fabric protocols bgp source-packet-routing],
[edit logical-systems name protocols bgp source-packet-routing],
[edit logical-systems name routing-instances name protocols bgp source-packet-routing],
[edit logical-systems name tenants name routing-instances name protocols bgp source-packet-
routing],
[edit protocols bgp source-packet-routing],
[edit routing-instances name protocols bgp source-packet-routing],
[edit tenants name routing-instances name protocols bgp source-packet-routing]
```

Description

Set the SRGB global block in SPRING

Options

index-range Index to the SRGB start label block

- **Range:** 16 through 1048575

start-label Start range for SRGB label block

- **Range:** 16 through 1048575

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

advertise-prefix-sid | 2022

srgb (Protocols IS-IS)

IN THIS SECTION

- Syntax | 28408
- Hierarchy Level | 28409
- Description | 28409
- Options | 28409
- Required Privilege Level | 28409
- Release Information | 28409

Syntax

```
srgb {
    start-label start-label-value;
```

```
    index-range index range-value;
}
```

Hierarchy Level

```
[edit logical-systems name protocols isis source-packet-routing],
[edit logical-systems name routing-instances name protocols isis source-packet-routing],
[edit protocols isis source-packet-routing],
[edit routing-instances name protocols isis source-packet-routing]
```

Description

Configure the segment routing global block (SRGB) in source packet routing in networking (SPRING) or segment routing (SR). The SRGB label range is based on the start label and the index range. The value of the start label indicates the start of the label range, and the value of the index range along with the value of the start label indicate the end of the label range.

Options

index-range *index-range-value*

Index range of the SRGB label block.

- **Range:** 32 through 1,048,559

start-label *start-label-value*

Start of the SRGB label block.

- **Range:** 16 through 1,048,575

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

Statement introduced in Junos OS Release 17.3 for QFX5110 and QFX5200 switches.

RELATED DOCUMENTATION

Understanding Adjacency Segments, Anycast Segments, and Configurable SRGB in SPRING

Configuring Segment Routing Global Blocks Label Ranges in SPRING for IS-IS Protocol

Example: Configuring Segment Routing Global Blocks in SPRING for IS-IS to Increase Network Speed

source-packet-routing (Protocols IS-IS) | 28266

srgb-label-range

IN THIS SECTION

- Syntax | 28410
- Hierarchy Level | 28410
- Description | 28410
- Options | 28411
- Required Privilege Level | 28411
- Release Information | 28411

Syntax

```
srgb-label-range <range-start> <range-end>
```

Hierarchy Level

```
[edit protocols mpls label-range],
[edit protocols ospf source-packet-routing srgb]
```

Description

SRGB configured under `mpls label-range` is termed as global SRGB. The MPLS label range is based on the start range and the end range. The value of the start range indicates the start of the label range, and the value of the end range along with the value of the start range indicate the end of the label range.

Options

<range-start> <range-end>	Start range and end range of the global SRGB label block.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

source-packet-routing (Protocols IS-IS) | 28266

srlg

IN THIS SECTION

- Syntax | 28411
- Hierarchy Level | 28412
- Description | 28412
- Options | 28412
- Required Privilege Level | 28412
- Release Information | 28412

Syntax

```
srlg {
    srlg-name {
```

```
        srlg-cost srlg-cost;
        srlg-value srlg-value;
    }
}
```

Hierarchy Level

```
[edit routing-options],
[edit logical-systems logical-system-name routing-options]
[edit protocols mpls interface interface-name]
```

Description

Configure Shared Risk Link Group (SRLG) parameters.

Options

srlg-cost *srlg-cost* Specify a cost for the SRLG ranging from **1** through **65535**.

srlg-value *srlg-value* Specify a Group ID for the SRLG ranging from **1** through **4294967295**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring SRLG

srlg-cost

IN THIS SECTION

- Syntax | 28413
- Hierarchy Level | 28413
- Description | 28413
- Required Privilege Level | 28413
- Release Information | 28413

Syntax

```
srlg-cost srlg-cost;
```

Hierarchy Level

```
[edit routing-options srlg],
[edit logical-systems logical-system-name routing-options srlg]
```

Description

Specify a cost for the Shared Risk Link Group (SRLG) ranging from **1** through **65535**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring SRLG

srlg-value

IN THIS SECTION

- Syntax | 28414
- Hierarchy Level | 28414
- Description | 28414
- Required Privilege Level | 28414
- Release Information | 28415

Syntax

```
srlg-value srlg-value;
```

Hierarchy Level

```
[edit routing-options srlg],
[edit logical-systems logical-system-name routing-options srlg]
```

Description

Specify a Group ID for the Shared Risk Link Group (SRLG) ranging from **1** through **4294967295**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring SRLG

srv6 (Routing Options)

IN THIS SECTION

- Syntax | 28415
- Hierarchy Level | 28416
- Description | 28416
- Options | 28416
- Required Privilege Level | 28416
- Release Information | 28416

Syntax

```
srv6 {
  block block-name;
  locator locator-name;
  no-reduced-srh;
  transit-srh-insert;
  traceoptions;
  oam;
}
```

Hierarchy Level

```
[edit routing-instances name routing-options source-packet-routing],
[edit routing-options source-packet-routing],
[edit tenants name routing-instances name routing-options source-packet-routing]
```

Description

Enable SRv6 source packet routing in an IPv6 network and configure the SRv6 options. SRv6 uses IPv6 addresses in extended headers instead of labels. These IPv6 addresses are stacked up in a routing extension header.

Options

no-reduced-srh Configure non-reduced Segment Routing Header (SRH) mode of tunnel encapsulation if you want to preserve the entire SID list in the SRH. In an SRv6 domain, an SRv6 segment can have multiple SRHs in an IPv6 header. Note that multiple SRHs can significantly raise the encapsulation overhead that can sometimes be larger than the actual packet payload.

- **Default:** Reduced SRH is enabled for SRv6 tunnel encapsulation. In a reduced SRH mode the device does not preserve the entire SID list in the SRH and thereby reduces the encapsulation overhead of an SRv6 packet.

transit-srh-insert Enable the SRv6 SRH insert mode on transit devices. IS-IS inserts an SRH as the next header in the original IPv6 packet header and modifies the next header according to the value of SRH. The IPv6 destination address is replaced with the IPv6 address of the first SID in the segment list and the original IPv6 destination address is carried in the SRH header as the last segment in the list.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

locator | 9997

Understanding SRv6 Network Programming in IS-IS Networks

Example: Configuring SRv6 Network Programming in IS-IS Networks

srv6 (Protocols BGP)

IN THIS SECTION

- Syntax | 28417
- Hierarchy Level | 28417
- Description | 28418
- Options | 28418
- Required Privilege Level | 28418
- Release Information | 28418

Syntax

```
srv6 {
     locator locator-name {
               end-dt4-sid end-dt4-sid | end-dt6-sid end-dt6-sid | end-dt46-sid end-dt46-sid |
non-default;
               micro-dt4-sid micro-sid-value | micro-dt6-sid micro-sid-value | micro-dt46-sid
micro-sid-value | non-default;
     }
```

Hierarchy Level

```
[edit fabric protocols bgp source-packet-routing],
[edit logical-systems name protocols bgp source-packet-routing],
[edit logical-systems name routing-instances name protocols bgp source-packet-routing],
[edit logical-systems name tenants name routing-instances name protocols bgp source-packet-
```

```
routing],
[edit protocols bgp source-packet-routing],
[edit routing-instances name protocols bgp source-packet-routing],
[edit tenants name routing-instances name protocols bgp source-packet-routing]
```

Description

Enable IPv6 segment routing capability in BGP. You can also configure BGP-based Layer 3 services over the SRv6 core network with BGP as the control pane and SRv6 as the dataplane.

Options

locator *locator-name*	Specify a name to identify the locator.
end-dt4-sid *end-dt4-sid*	(optional) Specify the SID to enable IPv4 transport over SRv6 core network.
end-dt6-sid *end-dt6-sid*	(optional) Specify the SID to enable IPv6 transport over SRv6 core network.
end-dt46-sid *end-dt46-sid*	(optional) Specify the SID to enable endpoint decapsulation of outer IPV6 header and lookup of the inner IPV4 or IPv6 packets in specific route table,
micro-dt4-sid *micro-sid-value*	(optional) Specify a static micro SID value 16-bit or 32-bit. If you do not specify the micro SID value for IPv4 prefixes, BGP allocates a dynamic micro SID value.
micro-dt6-sid *micro-sid-value*	(optional) Specify a static micro SID value 16-bit or 32-bit. If you do not specify the micro SID value for IPv6 prefixes, BGP allocates a dynamic micro SID value.
micro-dt46-sid *micro-sid-value*	(optional) Specify a static micro SID value 16-bit or 32-bit. If you do not specify the micro SID value for IPv4 and IPv6 prefixes, BGP allocates a dynamic micro SID value.
non-default	Use this option to configure multiple service SIDs. Specify non-default SIDs through a BGP routing policy. Note that you can use this option for service prefixes only.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

`end-dt46-sid` option introduced in Junos OS Release 22.1R1

`micro-dt4-sid micro-sid-value`, `micro-dt6-sid micro-sid-value`, `micro-dt46-sid micro-sid-value`, and `non-default`; options introduced in Junos OS Evolved Release 23.2R1 and Junos OS Release 23.4R1.

RELATED DOCUMENTATION

advertise-srv6-service | 2024

accept-srv6-service | 1649

srv6-sid

IN THIS SECTION

- Syntax | 28419
- Hierarchy Level | 28419
- Description | 28420
- Options | 28420
- Required Privilege Level | 28420
- Release Information | 28420

Syntax

```
srv6-sid {
    as as;
    router-id router-id;
    prefix prefix;
    system-id system-id;
}
```

Hierarchy Level

```
[edit logical-systems name policy-options policy-statement name from traffic-engineering],
```

```
[edit policy-options policy-statement name from traffic-engineering]
```

Description

Configure filter options for a traffic engineering policy. You can specify additional parameters, such as autonomous system (AS), prefix, router ID, and system ID for filtering SRV6 SID related traffic. SRv6 enables IPv6 segment routing capability in BGP.

Options

as *as*	Specify an AS to filter traffic
router-id *router-id*	Specify an IP prefix to match the router-ID against.
prefix	Specify an IPv6 prefix to match against.
system-id *system-id*	Specify an ISO address for the node.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

Link-State Distribution using SRv6

ssb

IN THIS SECTION

- Syntax | 28421
- Hierarchy Level | 28421
- Description | 28421
- Default | 28421
- Options | 28421
- Required Privilege Level | 28422
- Release Information | 28422

Syntax

```
ssb slot-number (always | preferred);
```

Hierarchy Level

```
[edit chassis redundancy]
```

Description

On M20 routers, configure which System and Switch Board (SSB) is the primary and which is the backup.

Default

By default, the SSB in slot 0 is the primary and the SSB in slot 1 is the backup.

Options

slot-number—Specify which slot is the primary and which is the backup.

always—Define this SSB as the sole device.

preferred—Define this SSB as the preferred device of at least two.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring SSB Redundancy on the M20 Router

ssh

IN THIS SECTION

- Syntax | 28423
- Syntax (EX Series Switch and the QFX Series) | 28423
- Description | 28423
- Options | 28424
- Additional Information | 28424
- Required Privilege Level | 28425
- Output Fields | 28425
- Sample Output | 28425
- Release Information | 28425

Syntax

```
ssh host
<bypass-routing>
<verbose>
<extensive>
<inet | inet6>
<interface interface-name>
<logical-system logical-system-name>
<tenant tenant-name>
<routing-instance routing-instance-name>
<source address>
<v2>
<port port-number>
```

Syntax (EX Series Switch and the QFX Series)

```
ssh host
<bypass-routing>
<verbose>
<extensive>
<inet | inet6>
<interface interface-name>
<routing-instance routing-instance-name>
<source address>
<v2>
<port port-number>
```

Description

Use the SSH program to open a connection between a local router or switch and a remote system and execute commands on the remote system. You can issue the `ssh` command from the Junos OS CLI to log in to a remote system or from a remote system to log in to the local router or switch. When executing this command, you include one or more CLI commands by enclosing them in quotation marks and separating the commands with semicolons:

```
ssh address 'cli-command1 ; cli-command2 '
```

Options

host	Name or address of the remote system.
bypass-routing	(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.
verbose	(Optional) Display verbose messages during the SSH attempt.
extensive	Display an extensive level of messaging during the SSH attempt.
inet | inet6	(Optional) Create an IPv4 or IPv6 connection, respectively.
interface *interface-name*	(Junos OS only) (Optional) Interface name for the SSH session. (This option does not work when *default-address-selection* is configured at the `[edit system]` hierarchy level, because this configuration uses the loopback interface as the source address for all locally generated IP packets.)
logical-system *logical-system-name*	(Optional) Name of a particular logical system for the SSH attempt.
tenant *tenant-name*	(Optional) Name of a particular tenant system for the SSH attempt.
routing-instance *routing-instance-name*	(Optional) Name of the routing instance for the SSH attempt.
source *address*	(Optional) Source address of the SSH connection.
v2	(Optional) Use SSH version 2 when connecting to a remote host.
port *port-number*	(Optional) Specify a port number for the SSH connection.

Additional Information

To configure an SSH (version 2) key for your user account, include the `authentication dsa-rsa` statement at the `[edit system login user *user-name*]` hierarchy level.

You can limit the number of times a user can attempt to enter a password while logging in through SSH. To specify the number of times a user can attempt to enter a password to log in through SSH, include the `retry-options` statement at the `[edit system login]` hierarchy level.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

ssh

```
user@switch> ssh user
Host key not found from the list of known hosts.
Are you sure you want to continue connecting (yes/no)? yes

Host ?user' added to the list of known hosts.
user@device's password:
Last login: Sun Jun 21 10:43:42 1998 from junos-router
% ...
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

The command `tenant` option is introduced in Junos OS Release 19.2R1 for SRX Series.

RELATED DOCUMENTATION

Configuring SSH Host Keys for Secure Copying of Data

ssh-known-hosts

IN THIS SECTION

- Syntax | 28426
- Hierarchy Level | 28426
- Description | 28426
- Options | 28427
- Required Privilege Level | 28427
- Release Information | 28427

Syntax

```
ssh-known-hosts {
    fetch-from-server server;
    host hostname {
        dsa-key key;
        ecdsa-sha2-nistp256-key key;
        ecdsa-sha2-nistp384-key key;
        ecdsa-sha2-nistp521-key key;
        ed25519-key key;
        rsa-key key;

    }
    load-key-file filename;
}
```

Hierarchy Level

```
[edit security]
```

Description

Configure SSH support for known hosts and for administering SSH host key updates.

Options

fetch-from-server *server* — Retrieve SSH public host key information from the specified server. Specify by server name or IP address.

host *host-name* — Hostname of the SSH known host entry. This option has the following suboptions:

- `dsa-key` *key*—Base64-encoded Digital Signature Algorithm (DSA) key for SSH version 2.
- `ecdsa-sha2-nistp256-key` *key*—Base64-encoded ECDSA-SHA2-NIST256 key.
- `ecdsa-sha2-nistp384-key` *key*—Base64-encoded ECDSA-SHA2-NIST384 key.
- `ecdsa-sha2-nistp521-key` *key*—Base64-encoded ECDSA-SHA2-NIST521 key.
- `ed25519-key` *key*—Base64-encoded ED25519 key.
- `rsa-key` *key*—Base64-encoded public key algorithm that supports encryption and digital signatures for SSH version 1 and SSH version 2.

load-key-file *filename* — Import SSH host key information from the named file. If the file is in a directory other than the home directory of the device, specify pathname as well. The default filename is `/var/tmp/ssh-known-hosts`.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configuring SSH Host Keys for Secure Copying of Data

ssh (NETCONF)

IN THIS SECTION

- Syntax | 28428
- Hierarchy Level | 28428
- Description | 28428
- Options | 28429
- Required Privilege Level | 28429
- Release Information | 28429

Syntax

```
ssh {
    client-alive-count-max number;
    client-alive-interval seconds;
    connection-limit limit;
    port port-number;
    rate-limit limit;
}
```

Hierarchy Level

```
[edit system services netconf]
```

Description

Enable access to the NETCONF SSH subsystem using the default port number 830, as specified by RFC 4742.

Options

client-alive-count-max *number*

(Optional) Threshold of client-alive responses that can be missed before the sshd process disconnects the client, thereby terminating the NETCONF session. Use this statement in conjunction with the `client-alive-interval` statement to disconnect unresponsive NETCONF clients.

- Default: 3
- Range: 0 through 255

client-alive-inteval *seconds*

(Optional) Timeout interval in seconds, after which, if no data has been received from the client, the sshd process sends a message through the encrypted channel to request a response from the client. Use this statement in conjunction with the `client-alive-count-max` statement to disconnect unresponsive NETCONF clients. This option applies to SSH protocol version 2 only.

- Default: 0 seconds
- Range: 0 through 65535 seconds

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

`client-alive-count-max` and `client-alive-interval` options added in Junos OS Release 21.1R1 and Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

connection-limit | **4537**

netconf | **11290**

port (NETCONF) | **12844**

rate-limit | **14063**

ssh (System Services)

IN THIS SECTION

- Syntax | 28430
- Hierarchy Level | 28431
- Description | 28431
- Options | 28431
- Required Privilege Level | 28437
- Release Information | 28437

Syntax

```
ssh {
    access-disable-external;
    authentication-order  [method 1 method2...];
    authorized-keys-command authorized-keys-command;
    authorized-keys-command-user authorized-keys-command-user;
    (authorized-principals principal-names | authorized-principals-command program-path)
    authorized-principals-file filename
    ciphers [ cipher-1 cipher-2 cipher-3 ...];
    client-alive-count-max number;
    client-alive-interval seconds;
    connection-limit limit;
    fingerprint-hash (md5 | sha2-256);
    host-certificate-file filename
    hostkey-algorithm (algorithm | no-algorithm);
    key-exchange [algorithm1 algorithm2...];
    log-key-changes log-key-changes;
    macs [algorithm1 algorithm2...];
    max-pre-authentication-packets number;
    max-sessions-per-connection number;
    no-challenge-response;
    no-password-authentication;
    no-passwords;
    no-public-keys;
    allow-tcp-forwarding;
```

```
    port port-number;
    protocol-version [v2];
    rate-limit number;
    rekey {
        data-limit bytes;
        time-limit minutes;
    }
    root-login (allow | deny | deny-password);
    sftp-server;
    trusted-user-ca-key-file filename
}
```

Hierarchy Level

```
[edit system services]
```

Description

Allow SSH requests from remote systems to access the local device.

Options

access-disable-external Disable external SSH access without disabling internal SSH access.

allow-tcp-forwarding Enable a user to create an SSH tunnel over a CLI session to a disaggregated Junos OS platform by using SSH.

Starting in Junos OS Release 22.2R1, we've disabled the TCP forwarding feature by default to enhance security. To enable the TCP forwarding feature, you can configure the `allow-tcp-forwarding` statement at the [`edit system services ssh`] hierarchy level. In addition, we've deprecated the `tcp-forwarding` and `no-tcp-forwarding` statements at the [`edit system services ssh`] hierarchy level.

authentication-order [*method1 method2...*] Configure the order in which the software tries different user authentication methods when attempting to authenticate a user. For each login attempt, the software tries the authentication methods in order, starting with the first one, until the password matches.

- **Default:** If you do not include the `authentication-order` statement, users are verified based on their configured passwords.
- **Syntax:** Specify one or more of the following authentication methods listed in the order in which they must be tried:
  - `ldaps`—Use LDAP authentication services.
  - `password`—Use the password configured for the user with the `authentication` statement at the `[edit system login user]` hierarchy level.
  - `radius`—Use RADIUS authentication services.
  - `tacplus`—Use TACACS+ authentication services.

authorized-keys-command

Specify a command string to be used to look up the user's public keys.

authorized-keys-command-user

Specify the user under whose account the authorized-keys-command is run.

authorized-principals *principal-names*

Specify a list of principals that can be accepted for authenticaiton. Principals added through this command are supplemental to the principals added with the `authorized-principals-file` command.

> **NOTE**: The `authorized-principals` and `authorized-principals-command` options are mutually exclusive.

authorized-principals-file *filename*

Configure the `AuthorizedPrincipals` file at **/var/etc**, for SSH certificate-based authentication. This file contains a list of names, one of which must appear in the certificate for it to be accepted for authentication.

authorized-principals-command *program-path*

Specify a program to be used for generating the list of allowed certificate principals found in the `AuthorizedPrincipals` file for SSH certificate-based authentication.

> **NOTE**: The `authorized-principals-command` and `authorized-principals` options are mutually exclusive.

ciphers [*cipher-1 cipher-2 cipher-3* **...]**

Specify the set of ciphers the SSH server can use to perform encryption and decryption functions.

- **Values:** Specify one or more of the following ciphers:

- `aes128-ctr`—128-bit AES in counter mode.
- `aes192-ctr`—192-bit AES in counter mode.
- `aes256-ctr`—256-bit AES in counter mode.

client-alive-count-max *number*

Configure the number of client alive messages that can be sent without sshd receiving any messages back from the client. If this threshold is reached while client alive messages are being sent, sshd will disconnect the client, terminating the session. Client alive messages are sent through the encrypted channel. Use in conjunction with the client-alive-interval statement to disconnect unresponsive SSH clients.

- **Default:** 3 messages
- **Range:** 0 through 255 messages

client-alive-interval *seconds*

Configure a timeout interval in seconds, after which if no data has been received from the client, sshd will send a message through the encrypted channel to request a response from the client. This option applies to SSH protocol version 2 only. Use in conjunction with the client-alive-count-max statement to disconnect unresponsive SSH clients.

- **Default:** 0 seconds
- **Range:** 1 through 65535 seconds

fingerprint-hash (md5 | sha2-256)

Specify the hash algorithm used by the SSH server when it displays key fingerprints.

> **NOTE**: The FIPS image does not permit the use of MD5 fingerprints. On systems in FIPS mode, `sha2-256` is the only available option.

- **Values:** Specify one of the following:
  - md5—Enable the SSH server to use the MD5 algorithm.
  - sha2-256—Enable the SSH server to use the sha2-256 algorithm.
- **Default:** sha2-256

host-certificate-file *filename*

Configure the `HostCertificate` file at **/etc/ssh/sshd_config** for SSH certificate-based authentication. This file contains the signed host certificate.

log-key-changes
log-key-changes

Enable Junos OS to log the authorized SSH keys. When the `log-key-changes` statement is configured and committed, Junos OS logs the changes to the set of authorized SSH keys for each user (including the keys that were added or removed). Junos OS logs the differences since the last time the `log-key-changes` statement was configured. If the `log-key-changes` statement was never configured, then Junos OS logs all the authorized SSH keys.

- **Default:** Junos OS logs all the authorized SSH keys.

macs [*algorithm1 algorithm2...*]

Specify the set of message authentication code (MAC) algorithms that the SSH server can use to authenticate messages.

> **NOTE**: The *macs* configuration statement represents a set. Therefore, it must be configured as follows:
>
> `user@host#set system services ssh macs [hmac-md5 hmac-sha1]`

- **Values:** Specify one or more of the following MAC algorithms to authenticate messages:
  - `hmac-md5`—Hash-based MAC using Message-Digest 5 (MD5)
  - `hmac-md5-96`—96-bits of hash-based MAC using MD5
  - `hmac-md5-96-etm@openssh.com`—96-bits of hash-based Encrypt-then-MAC using MD5
  - `hmac-md5-etm@openssh.com`—Hash-based Encrypt-then-MAC using MMD5
  - `hmac-sha1`—Hash-based MAC using secure hash algorithm-1 (SHA-1)
  - `hmac-sha1-96`—96-bits of hash-based MAC using SHA-1
  - `hmac-sha1-96-etm@openssh.com`—96-bits of hash-based Encrypt-then-MAC using SHA-1
  - `hmac-sha1-etm@openssh.com`—Hash-based Encrypt-then-MAC using SHA-1
  - `hmac-sha2-256`—256-bits of hash-based MAC using secure hash algorithm-2 (SHA-2)
  - `hmac-sha2-256-etm@openssh.com`—Hash-based Encrypt-then-Mac using SHA-2

- `hmac-sha2-512`—512-bits of hash-based MAC using SHA-2
- `hmac-sha2-512-etm@openssh.com`—Hash-based Encrypt-then-Mac using SHA-2
- `umac-128-etm@openssh.com`—Encrypt-then-MAC using UMAC-128 algorithm specified in RFC4418
- `umac-128@openssh.com`—UMAC-128 algorithm specified in RFC4418
- `umac-64-etm@openssh.com`—Encrypt-then-MAC using UMAC-64 algorithm specified in RFC4418
- `umac-64@openssh.com`—UMAC-64 algorithm specified in RFC4418

max-pre-authentication-packets *number*

Define the maximum number of pre-authentication SSH packets that the SSH server will accept prior to user authentication.

- **Range:** 20 through 2147483647 packets
- **Default:** 128 packets

max-sessions-per-connection *number*

Specify the maximum number of ssh sessions allowed per single SSH connection.

- **Range:** 1 through 65535 sessions
- **Default:** 10 sessions

no-challenge-response

Disable SSH challenge-response-based authentication methods.

> **NOTE**: Configuring this statement under the `[edit system services ssh]` hierarchy affects both the **SSH** login service and the **NETCONF** over SSH service.

no-password-authentication

Disable SSH password-based authentication methods.

> **NOTE**: Configuring this statement under the `[edit system services ssh]` hierarchy affects both the **SSH** login service and the **NETCONF** over SSH service.

no-passwords

Disable both password-based and challenge-response-based authentication for SSH.

> **NOTE**: Configuring this statement under the `[edit system services ssh]` hierarchy affects both the **SSH** login service and the **NETCONF** over SSH service.

no-public-keys

Disable public key authentication system wide. If you specify the no-public-keys statement at the [edit system login user *user-name* authentication] hierarchy level, you disable public key authentication for a specific user.

port *port-number*

Specify the port number on which to accept incoming SSH connections.

- **Default:** 22
- **Range:** 1 through 65535

protocol-version [v2]

Specify the Secure Shell (SSH) protocol version.

Starting in Junos OS Release 19.3R1 and Junos OS Release 18.3R3, on all SRX Series devices, we've removed the nonsecure SSH protocol version 1 (`v1`) option from the `[edit system services ssh protocol-version]` hierarchy level. You can use the SSH protocol version 2 (`v2`) as the default option to remotely manage systems and applications. With the `v1` option deprecated, Junos OS is compatible with OpenSSH 7.4 and later versions.

Junos OS releases before 19.3R1 and 18.3R3 continue to support the `v1` option to remotely manage systems and applications.

- **Default:** v2—SSH protocol version 2 is the default, introduced in Junos OS Release 11.4.

rate-limit *number*

Configure the maximum number of connection attempts per minute, per protocol (either IPv6 or IPv4) on an access service. For example, a rate limit of 10 allows 10 IPv6 SSH session connection attempts per minute and 10 IPv4 SSH session connection attempts per minute.

- **Range:** 1 through 250 connections
- **Default:** 150 connections

rekey

Specify limits before the session keys are renegotiated.

data-limit *bytes* Specify the data limit before renegotiating the session keys.

time-limit *minutes* Specify the time limit before renegotiating the session keys.

- **Range:** 1 through 1440 minutes

root-login (allow | deny | deny-password)

Control user access through SSH.

- allow—Allow users to log in to the device as root through SSH.
- deny—Disable users from logging in to the device as root through SSH.
- deny-password—Allow users to log in to the device as root through SSH when the authentication method (for example, RSA authentication) does not require a password.
- **Default:** `deny-password` is the default for most systems.

  Starting in Junos release 17.4R1 for MX Series routers, the default for root-login is **deny**. In previous Junos OS releases, the default setting for the MX240, MX480, MX960, MX2010 and MX2020 was **allow**.

sftp-server

Globally enable incoming SSH File Transfer Protocol (SFTP) connections. By configuring the `sftp-server` statement, you enable authorized devices to connect to the device through SFTP. If the `sftp-server` statement is not present in the configuration, then SFTP is globally disabled and no devices can connect to the device through SFTP.

trusted-user-ca-key-file *filename*

Configure the `TrustedUserCAKey` file at **/etc/ssh/sshd_config** for SSH certificate-based authentication. This file contains the public keys of an SSH certificate.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`ciphers`, `hostkey-algorithm`, `key-exchange`, and `macs` statements introduced in Junos OS Release 11.2.

`max-sessions-per-connection` and `no-tcp-forwarding` statements introduced in Junos OS Release 11.4.

SHA-2 options introduced in Junos OS Release 12.1.

Support for the curve25519-sha256 option on the `key-exchange` statement added in Junos OS Release 12.1X47-D10.

`client-alive-interval` and `client-alive-count-max` statements introduced in Junos OS Release 12.2.

`max-pre-authentication-packets` statement introduced in Junos OS Release 12.3X48-D10.

`no-passwords` statement introduced in Junos OS Release 13.3.

`no-public-keys` statement introduced in Junos OS release 15.1.

`tcp-forwarding` statement introduced in Junos OS Release 15.1X53-D50 for the NFX250 Network Services Platform.

`fingerprint-hash` statement introduced in Junos OS Release 16.1.

`log-key-changes` statement introduced in Junos OS Release 17.4R1.

`sftp-server` statement introduced in Junos OS Release 19.1R1.

`no-challenge-response` and `no-password-authentication` statements introduced in Junos OS Release 19.4R1.

Option `ldaps` introduced in Junos OS Release 20.2R1.

`allow-tcp-forwarding` option added in Junos OS Release 22.2R1.

`access-disable-external` option added in Junos OS Release 22.2R3.

The `athorized-prinicpals`, `authorized-principals-command`, and the `authorized-principals-file` options added in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

- Configuring SSH Service for Remote Access to the Router or Switch
- Junos OS User Authentication Methods
- *Configuring clear-text or SSL Service for Junos XML Protocol Client Applications*
- Configuring SSH Service for Remote Access to the Disaggregated Junos OS Platform

ssl

IN THIS SECTION

- Syntax | 28439
- Hierarchy Level | 28439
- Description | 28439
- Options | 28440
- Required Privilege Level | 28441
- Release Information | 28441

Syntax

```
ssl {
    address ip-address;
    hot-reloading;
    local-certificate local-certificate;
    mutual-authentication {
        certificate-authority certificate-authority-profile-name;
        client-certificate-request (no-certificate | request-certificate | request-certificate-
and-verify | require-certificate | require-certificate-and-verify);
    }
    port port;
    use-pki;
}
```

Hierarchy Level

```
[edit system services extension-service request-response grpc]
```

Description

Configure API connection settings based on Secure Sockets Layer (SSL) technology.

Options

address *ip-address*

Specify the IP address to listen for incoming connections. If you use the default IP address 0.0.0.0, the JET service process (jsd) listens on the IP address in the default routing instance.

- **Default:** 0.0.0.0

mutual-authentication

Enable bidirectional authentication. Use this option, in conjunction with `client-certificate-request` and `certificate-authority` *`profile-name`* to configure client authentication using SSL-based certificates.

client-certificate-request

Specify the requirements for a client certificate.

`no-certificate`—Client certificate is not requested.

> **NOTE**: We strongly recommend that you use this option in a test environment only.

`request-certificate`—Request certificate from client but do not verify.

`request-certificate-and-verify`—Request certificate from client and verify if provided.

`require-certificate`—Client certificate is mandatory, but do not verify.

`require-certificate-and-verify`—Client certificate is mandatory, and certificate is verified.

- **Default:** no-certificate

> **NOTE**: You can specify only one value for a client certificate.

hot-reloading

Enable persistent gRPC sessions across SSL certificate updates between a network management system or collector and a network device. If this feature is not enabled, when a certificate is updated between the network device and remote management system, all existing gRPC sessions are terminated.

certificate-authority *profile-name*

Specify the name of a certificate-authorirty profile configured at the [`edit security pki ca-profile`] hierarchy level. This profile is used to validate the certificate provided by the client.

port *port*

Specify the port number to accept incoming connections.

> **NOTE**: For gRPC connections used to stream telemetry data, the required port number is 32767.

- **Range:** 1 through 65535
- **Default:** 9090

`use-pki` Use the Public Key Infrastructure (PKI) database on the device for gRPC-based operations that require certificates.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`mutual-authentication`, `client-certificate-request`, and `certificate-authority` options added in Junos OS Release 17.4R1.

`hot-reloading` option added in Junos OS Release 20.4R1.

`use-pki` option added in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

grpc | 7920

JET Service Process Overview

Configuring Request-Response Service for JET Applications

ssl-encryption

IN THIS SECTION

- Syntax | 28442
- Hierarchy Level | 28442
- Description | 28442
- Options | 28442
- Required Privilege Level | 28443
- Release Information | 28443

Syntax

```
ssl-encryption (any | no | yes);
```

Hierarchy Level

```
[edit security application-firewall rule-sets rule-set-name rule rule-name match]
```

Description

Distinguishes between encrypted and unencrypted SSL traffic as match criteria for the rule. In application firewall usage, this option lets you specify different actions for encrypted and unencrypted SSL traffic.

Starting in Junos OS Release 18.2R1 application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration.

Options

- `any`—Matches both encrypted and unencrypted SSL traffic.

- `no`—Matches unencrypted SSL traffic only.
- `yes`—Matches encrypted SSL traffic only.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

ssl-inspection

IN THIS SECTION

- Syntax | 28443
- Hierarchy Level | 28444
- Description | 28444
- Options | 28444
- Required Privilege Level | 28445
- Release Information | 28445

Syntax

```
ssl-inspection {
    cache-prune-chunk-size number;
    key-protection;
```

```
    maximum-cache-size number;
    session-id-cache-timeout seconds;
    sessions number;
}
```

Hierarchy Level

```
[edit security idp sensor-configuration]
```

Description

Inspect HTTP traffic encrypted in SSL protocol. SSL inspection is disabled by default. It is enabled if you configure SSL inspection.

With the Intrusion Detection and Prevention (IDP) Secure Sockets Layer (SSL) decryption feature, SRX Series Firewalls load configured RSA private keys to memory and use them to establish SSL session keys to decrypt data. IDP is required to decrypt the RSA keys and to check the integrity before performing normal encryption or decryption operations using the keys.

Options

The remaining statements are explained separately. See CLI Explorer.

cache-prune-chunk-size

Number of cache entries to delete when pruning SSL session ID cache.

- **Syntax:** *cache-prune-chunk-size*—Number of cache entries to delete when pruning SSL session ID cache.
- **Range:** 1 through 100,000
- **Default:** 100

key-protection

Enabling key protection provides improved security. When key protection is enabled, persistent keys are encrypted when not in use.

Enabling or disabling of this option requires rebooting the device.

Enable secure key handling. This option is off by default.

maximum-cache-size

Maximum SSL session ID cache size.

- **Syntax:** *maximum-cache-size*—Maximum number of SSL session ID cache size.

- **Range:** 1 through 5,000,000 sessions
- **Default:** 100,000

session-id-cache-timeout

Sets the timeout value for an IDP session ID cache (range: 1 through 7200 seconds).

- **Syntax:** *maximum-cache-size*—Maximum number of SSL session ID cache size.

sessions

Maximum number of SSL sessions for inspection. This limit is per Services Processing Unit (SPU).

- **Syntax:** *number*—Number of SSL session to inspect.
- **Range:** 1 through 100,000
- **Default:** 10,000

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statements introduced in Junos OS Release 9.3.

Options `cache-prune-chunk-size` and `maximum-cache-size` introduced in Junos OS Release 10.2.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D100	Starting from 15.1X49, the IDP SSL Inspection feature is deprecated. Juniper recommends use of SSL Proxy feature.

ssl-proxy (Application Services)

IN THIS SECTION

- Syntax | 28446
- Hierarchy Level | 28446
- Description | 28446
- Options | 28446
- Required Privilege Level | 28447
- Release Information | 28447

Syntax

```
ssl-proxy {
    profile-name profile-name
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
application-services]
```

Description

Enable SSL proxy and identify the name of the SSL proxy profile to be used. This option is supported on SRX340, SRX345, SRX550M, SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices.

Options

profile-name	SSL proxy profile.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

ssl (Services)

IN THIS SECTION

- Syntax | 28447
- Hierarchy Level | 28449
- Description | 28449
- Options | 28450
- Required Privilege Level | 28450
- Release Information | 28450

Syntax

```
ssl {
    initiation {
        profile profile-name {
            actions {
                crl {
                    disable;
                    if-not-present (allow | drop);
```

```
                ignore-hold-instruction-code;
            }
            ignore-server-auth-failure;
        }
        client-certificate;
        custom-ciphers [cipher];
        enable-flow-tracing;
        enable-session-cache;
        preferred-ciphers (custom | medium | strong | weak);
        protocol-version (all | tls1  | tls11 | tls12);
        trusted-ca (all | [ca-profile] );
    }
}
proxy {
    global-config {
        session-cache-timeout seconds;
    }
    profile profile-name {
        actions {
            crl {
                disable;
                if-not-present (allow | drop);
                ignore-hold-instruction-code;
            }
            disable-session-resumption;
            ignore-server-auth-failure;
            log {
                all;
                errors;
                info;
                sessions-allowed;
                sessions-dropped;
                sessions-ignored;
                sessions-whitelisted;
                warning;
            }
            renegotiation {
                (allow | allow-secure | drop);
            }
        }
        custom-ciphers [cipher];
        enable-flow-tracing;
         preferred-ciphers (custom | medium | strong | weak);
```

```
            root-ca root-certificate;
            trusted-ca (all | [ca-profile] );
            whitelist [global-address-book-addresses];
        }
    }
    termination {
        profile profile-name {
            custom-ciphers [cipher];
            enable-flow-tracing;
            enable-session-cache;
            preferred-ciphers (custom | medium | strong | weak);
            protocol-version (all | tls1  | tls11 | tls12);
            server-certificate certificate-identifier;
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            (no-world-readable | world-readable);
            size maximum-file-size;
        }
        flag flag;
        level [brief | detail | extensive | verbose];
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Specify the configuration for Secure Socket Layer (SSL) support service. This statement is supported on the SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. The `crl` statement is supported from 15.1X49-D30. The `protocol-version` statement is updated to include `tls11` and `tls12` from Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

Firewall User Authentication Overview

ssl-termination-profile

IN THIS SECTION

- Syntax | 28451
- Hierarchy Level | 28451
- Description | 28451
- Options | 28451
- Required Privilege Level | 28451
- Release Information | 28451

Syntax

```
ssl-termination-profile profile-name;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication pass-through]
```

Description

Specify the SSL termination profile used for SSL offloading.

Options

profile-name Specify the name of the SSL termination profile used to the SSL offload.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Security Policies Overview

ssm-groups

IN THIS SECTION

- Syntax | 28452
- Hierarchy Level | 28452
- Description | 28452
- Options | 28453
- Required Privilege Level | 28453
- Release Information | 28453

Syntax

```
ssm-groups [ ip-addresses ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Description

Configure source-specific multicast (SSM) groups.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the `ssm-groups` statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the `ssm-groups` statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.

IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.

Options

ip-addresses—List of one or more additional SSM group addresses separated by a space.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring Source-Specific Multicast Groups with Any-Source Override

ssm-map (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | 28454
- Hierarchy Level | 28454
- Description | 28454
- Options | 28454
- Required Privilege Level | 28455
- Release Information | 28455

Syntax

```
ssm-map ssm-map-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name]
```

Description

Apply an SSM map to a dynamic IGMP interface. SSM mapping translates IGMPv1 and IGMPv2 membership reports to an IGMPv3 report, which enables hosts running IGMPv1 or IGMPv2 to participate in SSM. The SSM map associates an SSM policy that matches group addresses to be translated with the source addresses where the group addresses are found.

Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the `ssm-map` statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the statement is removed from the CLI. Use the `ssm-map-policy` instead to associate the policy with the dynamic IGMP interface for all releases with enhanced subscriber management.

If you upgrade from an earlier release that does not support enhanced subscriber management (any release earlier than Junos OS Release 15.1R4) with a configuration that includes the `ssm-map` statement, the results vary depending on the release to which you are upgrading:

- Upgrade to Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, or later releases in those branches—The `ssm-map` configuration is allowed and does not cause the upgrade to fail. However, the configuration has no effect and subscribers cannot log in.
- Upgrade to Junos OS Release 17.3R1 or later releases—The upgrade fails because the `ssm-map` configuration is not allowed. If you perform the upgrade without validation (`no-validate`), the upgrade passes and the `ssm-map` configuration is accepted, but it has no effect.

BEST PRACTICE: Delete the `ssm-map` configuration before you upgrade.

Options

ssm-map-name—Name of SSM map.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Statement became non-functional in Junos OS Release 15.1R4.

Statement deprecated in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1.

Statement removed from CLI in Junos OS Release 17.3R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
16.1R5	Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the `ssm-map` statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the statement is removed from the CLI.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Source-Specific Multicast Groups Overview

ssm-map (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 28456
- Hierarchy Level | 28456
- Description | 28456

- Options | 28457
- Required Privilege Level | 28457
- Release Information | 28457

Syntax

```
ssm-map ssm-map-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Description

Apply an SSM map to a dynamic MLD interface. SSM mapping translates MLDv1 membership reports to an MLDv2 report, which enables hosts running MLDv1 to participate in SSM. The SSM map associates an SSM policy that matches group addresses to be translated with the source addresses where the group addresses are found.

Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the `ssm-map` statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the statement is removed from the CLI. Use the `ssm-map-policy` instead to associate the policy with the dynamic MLD interface for all releases with enhanced subscriber management.

If you upgrade from an earlier release that does not support enhanced subscriber management (any release earlier than Junos OS Release 15.1R4) with a configuration that includes the `ssm-map` statement, the results vary depending on the release to which you are upgrading:

- Upgrade to Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, or later releases in those branches—The `ssm-map` configuration is allowed and does not cause the upgrade to fail. However, the configuration has no effect and subscribers cannot log in.

- Upgrade to Junos OS Release 17.3R1 or later releases—The upgrade fails because the `ssm-map` configuration is not allowed. If you perform the upgrade without validation (`no-validate`), the upgrade passes and the `ssm-map` configuration is accepted, but it has no effect.

BEST PRACTICE: Delete the `ssm-map` configuration before you upgrade.

Options

`ssm-map-name`—Name of SSM map.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

Statement became non-functional in Junos OS Release 15.1R4.

Statement deprecated in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1.

Statement removed from CLI in Junos OS Release 17.3R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
16.1R5	Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the `ssm-map` statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the statement is removed from the CLI.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Example: Configuring SSM Mapping

ssm-map-policy (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | 28458
- Hierarchy Level | 28458
- Description | 28458
- Required Privilege Level | 28459
- Release Information | 28459

Syntax

```
ssm-map-policy ssm-map-policy-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name]
```

Description

Apply an SSM map policy to a dynamic IGMP interface. SSM mapping translates IGMPv1 and IGMPv2 membership reports to an IGMPv3 report, which enables hosts running IGMPv1 or IGMPv2 to participate in SSM. The map policy associates the group addresses to be translated with the source addresses where the group addresses are found. You configure the SSM map policy with the `policy-statement` statement at the `[edit policy-options]` hierarchy level.

For statically-configured IGMP interfaces, use the `ssm-map-policy (IGMP)` statement.

Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1, use this statement instead of the `ssm-map` statement to associate the policy with the dynamic IGMP interface for all releases with enhanced subscriber management. Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the `ssm-map` statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the statement is removed from the CLI.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R5.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
16.1R5	Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1, use this statement instead of the `ssm-map` statement to associate the policy with the dynamic IGMP interface for all releases with enhanced subscriber management.

RELATED DOCUMENTATION

Configuring SSM Mapping for Dynamic IGMP and MLD

Dynamic IGMP Configuration Overview

Example: Configuring SSM Maps for Different Groups to Different Sources

ssm-map-policy (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 28460
- Hierarchy Level | 28460
- Description | 28460
- Required Privilege Level | 28460
- Release Information | 28460

Syntax

```
ssm-map-policy ssm-map-policy-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Description

Apply an SSM map policy to a dynamic MLD interface. SSM mapping translates MLDv1 membership reports to an MLDv2 report, which enables hosts running MLDv1 to participate in SSM. The map policy associates the group addresses to be translated with the source addresses where the group addresses are found. You configure the SSM map policy with the `policy-statement` statement at the `[edit policy-options]` hierarchy level.

For statically-configured MLD interfaces, use the `ssm-map-policy (MLD)` statement.

Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1, use this statement instead of the `ssm-map` statement to associate the policy with the dynamic MLD interface for all releases with enhanced subscriber management. Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the `ssm-map` statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the `ssm-map` statement is removed from the CLI.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R5.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
16.1R5	Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1, use this statement instead of the `ssm-map` statement to associate the policy with the dynamic MLD interface for all releases with enhanced subscriber management.

RELATED DOCUMENTATION

Configuring SSM Mapping for Dynamic IGMP and MLD

Dynamic MLD Configuration Overview

Example: Configuring SSM Maps for Different Groups to Different Sources

ssm-map-policy (IGMP)

IN THIS SECTION

- Syntax | **28461**
- Hierarchy Level | **28462**
- Description | **28462**
- Options | **28462**
- Required Privilege Level | **28462**
- Release Information | **28462**

Syntax

```
ssm-map-policy ssm-map-policy-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Description

Apply an SSM map policy to a statically configured IGMP interface.

For dynamically-configured IGMP interfaces, use the `ssm-map-policy (Dynamic IGMP Interface)` statement.

Options

ssm-map-policy-name—Name of SSM map policy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring SSM Maps for Different Groups to Different Sources

ssm-map-policy (MLD)

IN THIS SECTION

- Syntax | 28463
- Hierarchy Level | 28463

- Description | 28463
- Options | 28463
- Required Privilege Level | 28463
- Release Information | 28463

Syntax

```
ssm-map-policy ssm-map-policy-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

Apply an SSM map policy to a statically configured MLD interface.

For dynamically-configured MLD interfaces, use the `ssm-map-policy (Dynamic MLD Interface)` statement.

Options

ssm-map-policy-name—Name of SSM map policy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Example: Configuring SSM Maps for Different Groups to Different Sources

ssm-map (Protocols IGMP)

IN THIS SECTION

- Syntax | 28464
- Hierarchy Level | 28464
- Description | 28464
- Options | 28464
- Required Privilege Level | 28465
- Release Information | 28465

Syntax

```
ssm-map ssm-map-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Description

Apply an SSM map to an IGMP interface.

Options

ssm-map-name—Name of SSM map.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring SSM Mapping

ssm-map (Protocols IGMP AMT)

IN THIS SECTION

- Syntax | 28465
- Hierarchy Level | 28466
- Description | 28466
- Options | 28466
- Required Privilege Level | 28466
- Release Information | 28466

Syntax

```
ssm-map ssm-map-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp amt relay defaults],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols igmp
amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances routing-instance-name protocols igmp amt relay defaults]
```

Description

Apply a source-specific multicast (SSM) map to all Automatic Multicast Tunneling (AMT) interfaces.

Options

ssm-map-name—Name of the SSM map.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Default IGMP Parameters for AMT Interfaces

ssm-map (Protocols MLD)

IN THIS SECTION

Syntax | 28467

- Hierarchy Level | 28467
- Description | 28467
- Options | 28467
- Required Privilege Level | 28467
- Release Information | 28467

Syntax

```
ssm-map ssm-map-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

Apply an SSM map to an MLD interface.

Options

ssm-map-name—Name of SSM map.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring SSM Mapping

ssm-map (Routing Options Multicast)

IN THIS SECTION

- Syntax | 28468
- Hierarchy Level | 28468
- Description | 28468
- Options | 28469
- Required Privilege Level | 28469
- Release Information | 28469

Syntax

```
ssm-map ssm-map-name {
    policy [ policy-names ];
    source [ addresses ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Description

Configure SSM mapping.

Options

ssm-map-name—Name of the SSM map.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring SSM Mapping

stacked-interface-set (Dynamic Profiles)

IN THIS SECTION

- Syntax | 28470
- Hierarchy Level | 28470
- Description | 28470
- Options | 28470
- Required Privilege Level | 28471
- Release Information | 28471

Syntax

```
stacked-interface-set {
    interface-set-name interface-set-name {
        interface-set-name interface-set-name;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces]
```

Description

For MX Series routers with MPCs that support 5-level hierarchical scheduling, define an interface set over interface set hierarchy. The child interface set (L3) contains subscriber logical interfaces, and the parent interface set (L2) contains one or more interface sets as members.

> **NOTE**: For both interface-sets in the stack (at L3 and L2), a CoS traffic-control-profile (TCP) must be assigned to each.

Options

interface-set-name—Name of the interface set to be configured or one of the following Junos OS predefined variables:

- $junos-aggregation-interface-set-name—L2 interface-set representing a logical intermediate node (DPU-C or PON tree) in the access network.
- $junos-default-interface-set-name—An intermediate, default variable used to resolve the L3 (child) interface-set name based on whether subscriber access is via a logical intermediate aggregation node
- $junos-interface-set-name—Predefined variable that, when used, is replaced with the interface-set obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.
- $junos-phy-ifd-underlying-intf-set-name—A default, topology-based interface-set derived from the predefined variable, $junos-phy-ifd-underlying-intf-set-name by appending “-underlying” to conserve

L2 CoS nodes. It is a variation of the $junos-phy-ifd-interface-set-name variable used as a default L2 interface-set for 4-level hierarchy or default L3 interface-set for 5-level hierarchy.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

CoS for Interface Sets of Subscribers Overview

interface-set (Dynamic Profiles) | 9037

stacked-vlan-ranges

IN THIS SECTION

- Syntax | 28471
- Hierarchy Level | 28472
- Description | 28472
- Required Privilege Level | 28472
- Release Information | 28473

Syntax

```
stacked-vlan-ranges {
    access-profile profile-name;
    authentication {
        packet-types  [packet-types];
```

```
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-18
            option-37
            option-82;
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | inet);
        access-profile vlan-dynamic-profile-name;
        ranges (any | low-tag-high-tag),(any | low-tag-high-tag);
    }
    override;
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure]
```

Description

Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

Configuring Interfaces to Support Both Single and Stacked VLANs

stacked-vlan-ranges (RADIUS Options)

IN THIS SECTION

- Syntax | 28473
- Hierarchy Level | 28473
- Description | 28474
- Options | 28474
- Required Privilege Level | 28474
- Release Information | 28474

Syntax

```
stacked-vlan-ranges (any | low-outer-tag-high-outer-tag),any;
```

Hierarchy Level

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
```

Description

Configure the stacked VLAN (S-VLAN) range of subscribers to which the named NAS-Port options definition applies.

> **NOTE**: You can configure a maximum of 16 NAS-Port options definitions per physical interface. Each definition can include a maximum of 32 VLAN ranges or 32 S-VLAN ranges, but cannot include a combination of VLAN ranges and S-VLAN ranges.

Options

any Entire S-VLAN range representing all S-VLAN IDs. The inner tag (S-VLAN ID) of the S-VLAN range must be configured as `any` to represent all inner VLAN ID tags.

low-outer-tag Outer VLAN ID tag representing the lower limit of the S-VLAN range.

- **Range:** 1 through 4094

high-outer-tag Outer VLAN ID tag representing the upper limit of the S-VLAN range.

- **Range:** 1 through 4094

> **NOTE**: To specify a single outer VLAN ID tag, set *`low-outer-tag`* and *`high-outer-tag`* to the same value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

> *RADIUS NAS Port Attributes and Options*
>
> *RADIUS NAS Port Attributes and Options*

stacked-vlan-tagging

IN THIS SECTION

- Syntax | 28475
- Hierarchy Level | 28475
- Description | 28475
- Required Privilege Level | 28475
- Release Information | 28476

Syntax

```
stacked-vlan-tagging;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

For Gigabit Ethernet IQ interfaces, Gigabit Ethernet, 10-Gigabit Ethernet LAN/WAN PIC, and 100-Gigabit Ethernet Type 5 PIC with CFP, enable stacked VLAN tagging for all logical interfaces on the physical interface.

For pseudowire subscriber interfaces, enable stacked VLAN tagging for logical interfaces on the pseudowire service.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview

stale-labels-holddown-period

IN THIS SECTION

- Syntax | 28476
- Hierarchy Level | 28476
- Description | 28477
- Options | 28477
- Required Privilege Level | 28477
- Release Information | 28477

Syntax

```
stale-labels-holddown-period stale-labels-holddown-period;
```

Hierarchy Level

```
[edit fabric protocols bgp],
[edit protocols bgp],
[edit routing-instances name protocols bgp]
```

Description

Configure a time duration for which the BGP allocated MPLS labels are retained after they go stale. Holding the stale labels without deleting them immediately allows the upstream peers to receive and install the new labeled advertisements. Note that these MPLS stale labels are retained even when MPLS fast reroute (FRR) is not configured.

> **NOTE**: During this hold down period you cannot use the label space for new label allocation, which momentarily reduces the supported label scale. Therefore, we recommend setting an appropriate minimum value as the hold down time period for stale labels.

Options

stale-labels-holddown-period

Specify a time duration in seconds for which the MPLS labels allocated by BGP are kept after they go stale.

- **Range:** 1 through 600

Required Privilege Level

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

stale-routes-time | 28478

stale-routes-time

IN THIS SECTION

- Syntax | 28478
- Hierarchy Level | 28478
- Description | 28478
- Options | 28478
- Required Privilege Level | 28479
- Release Information | 28479

Syntax

```
stale-routes-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-routing-name protocols bgp graceful-restart],
[edit logical-systems logical-routing-name routing-instances routing-instance-name protocols bgp
graceful-restart],
[edit protocols bgp graceful-restart],
[edit routing-instances routing-instance-name protocols bgp graceful-restart]
```

Description

Specify the maximum time that stale routes are kept during a restart. The `stale-routes-time` statement allows you to set the length of time the routing device waits to receive messages from restarting neighbors before declaring them down.

Options

seconds—Time the router device waits to receive messages from restarting neighbors before declaring them down.

- **Range:** 1 through 600 seconds
- **Default:** 300 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Configuring Graceful Restart for QFabric Systems

restart-time (BGP Graceful Restart) | **15663**

stale-routes-time (Fabric Control)

IN THIS SECTION

- Syntax | **28480**
- Hierarchy Level | **28480**
- Description | **28480**
- Options | **28480**
- Required Privilege Level | **28480**
- Release Information | **28480**

Syntax

```
stale-routes-time seconds;
```

Hierarchy Level

```
[edit fabric protocols fabric-control graceful-restart]
```

Description

Set the length of time that the fabric control Routing Engine waits to receive messages from devices before declaring them down. Configure a stale routes time of 1800 seconds if the number of VLAN members (vmembers) exceeds 32k.

Options

seconds Amount of time that the fabric control Routing Engine waits to receive messages from other devices before declaring them down.

- **Default:** 900 seconds
- **Range:** 900 to 1800 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X52-D10.

RELATED DOCUMENTATION

> Understanding Bridging and VLANs on Switches
>
> *Understanding Routing Engines in the QFabric System*

stale-time (Long-Lived Graceful Restart for BGP Restarter)

IN THIS SECTION

- Syntax | 28481
- Hierarchy Level | 28481
- Description | 28482
- Options | 28483
- Required Privilege Level | 28483
- Release Information | 28483

Syntax

```
stale-time interval;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family (l2vpn | route-target | inet)
(labeled-unicast | flow) graceful-restart long-lived restarter],
[edit logical-systems logical-system-name protocols bgp group group-name family (l2vpn | route-
target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
(l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived
restarter],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart
long-lived restarter],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family (l2vpn | route-target | inet) (labeled-unicast | flow)
graceful-restart long-lived restarter],
[edit routing-instances routing-instance-name protocols bgp family (l2vpn | route-target | inet)
(labeled-unicast | flow) graceful-restart long-lived restarter],
[edit routing-instances routing-instance-name protocols bgp group group-name family (l2vpn |
```

```
route-target | inet) (labeled-unicast | flow) graceful-restart long-lived restarter],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-restart long-lived
restarter],
[edit protocols bgp family (l2vpn | route-target | inet) (labeled-unicast | flow) graceful-
restart long-lived restarter],
[edit protocols bgp group group-name family (l2vpn | route-target | inet) (labeled-unicast |
flow) graceful-restart long-lived restarter],
[edit protocols bgp group group-name neighbor address family (l2vpn | route-target | inet)
(labeled-unicast | flow) graceful-restart long-lived restarter]
```

NOTE: Each routing table is identified by the protocol family or address family indicator (AFI) and a subsequent address family identifier (SAFI). The AFI parameter can be one of the `(l2vpn | inet | route-target)` protocols and the SAFI parameter can be either of the `(flow | labeled-unicast)` protocols for inet family and one of the `(auto-discovery-mspw | auto-discovery-only | signaling)` protcols for L2VPN family..

Configuring LLGR does not require that BGP graceful restart also be configured. The long-lived-graceful-restart section is visible only for families l2vpn, inet labeled-unicast, inet flow and route-target. It is prohibited for inet-mvpn, inet6-mvpn and inet-mdt. It is hidden for other families.

Description

Specify the period of time for which stale routes must be preserved by using le the long-lived graceful restart capability for BGP sessions on the restarting router. When LLGR restarter is enabled or disabled for a family or the stale- time is changed, the session is reset so that the new capability can be sent to the neighbor. You can configure the stale period for each address family at the logical system or routing instance level.

The stanzas in the per-family graceful-restart long-lived restarter configuration section enables LLGR restarter mode negotiation for BGP globally, or for a group or neighbor. The values are inherited by groups from the global configuration, and by neighbors from the group configuration. The disable attribute is used to override configuration inherited from a higher level. It does not disable LLGR receiver mode; you must disable LLGR receiver mode explicitly for all families as necessary.

In addition, times can also be configured using the following notation: <hours>:<minutes>:<seconds> For example, 12:00:00 specifies twelve hours. The hours and minutes are optional.

The two notations can be combined, for example, 2w1d 12:00:02 specifies two weeks, one day, twelve hours and two seconds (1339202 seconds). (Note that the CLI requires double-quotes around a value like this with spaces in it.) Expressed in this notation, the maximum stale time is 27w5d 04:20:15 (27 weeks, 5 days, 4 hours, 20 minutes and 15 seconds). While the show configuration command displays

the actually configured values, when the associated timers are displayed in run-time show commands such as `show bgp neighbor`, the values are normalized, such as 1d36h becoming 2d 12:00:00. The full rules for displaying normalized LLGR times depend on the `clear bgp neighbor neighbor-address gracefully` command configuration.

Options

interval Period as a measure of the number of weeks, days, hours, minutes, and seconds for which stale routes must be maintained when long-lived graceful restart mechanism is enabled on the restarter router.

The range of values for stale-time is from 1 to 16777215 (2^24 – 1) seconds. The value is a simple integer giving the number of seconds by default, but it can also be specified using the following notation:

```
[<weeks>w][<days>d][<hours>h][<minutes>m][<seconds>s]
```

For example, you can specify 27 days as 27d, 648h, 38880m or 2332800s. 90 minutes can be configured as 1h30m, 90m or 5400s. The specified number of days is multiplied by 86400, the number of hours by 3600 and the number of minutes by 60; these are added to the seconds to get the total. A combined format of days and hours, in different time period units, such as 1d36h are permitted, as long as the specified total does not exceed the maximum stale time.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for BGP

Junos OS High Availability User Guide

stamp (Sampling)

IN THIS SECTION

- Syntax | 28484
- Hierarchy Level | 28484
- Description | 28484
- Options | 28484
- Required Privilege Level | 28484
- Release Information | 28485

Syntax

```
(stamp | no-stamp);
```

Hierarchy Level

```
[edit forwarding-options sampling family (inet |inet6 |mpls) output file]
```

Description

Include a timestamp with each line in the output file.

Options

`no-stamp`—Do not include timestamps. This is the default.

`stamp`—Include a timestamp with each line of packet sampling information.

- **Default:** No timestamp is included.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

Collecting Traffic Sampling Output in a File

standby

IN THIS SECTION

- Syntax | 28485
- Hierarchy Level | 28485
- Description | 28486
- Required Privilege Level | 28486
- Release Information | 28486

Syntax

```
standby;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
```

```
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

Have the path remain up at all times to provide instant switchover if connectivity problems occur.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Hot Standby of Secondary Paths for LSPs

standby

IN THIS SECTION

- Syntax | 28487
- Hierarchy Level | 28487
- Description | 28487
- Required Privilege Level | 28487
- Release Information | 28487

Syntax

```
standby;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

Enable the path to remain up at all times to provide instant switchover if connectivity problems occur.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Hot Standby of Secondary Paths for LSPs

Configuring Path Protection in an MPLS Network (CLI Procedure)

standby-path-creation-delay

IN THIS SECTION

- Syntax | 28488
- Hierarchy Level | 28488
- Description | 28488
- Options | 28489
- Required Privilege Level | 28489
- Release Information | 28489

Syntax

```
standby-path-creation-delay <seconds>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Configure the time interval after which a standby path is created, when a new ECMP interface or neighbor is added to the network.

In the absence of this statement, ECMP joins are redistributed as soon as a new ECMP interface or neighbor is added to the network.

Options

<seconds> Time interval after which a standby path is created, when a new ECMP interface or neighbor is added to the network. Range is from 1 through 300.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring PIM Make-Before-Break Join Load Balancing

Configuring PIM Join Load Balancing

clear pim join-distribution | **3809**

join-load-balance | **9431**

idle-standby-path-switchover-delay | **8301**

standby (Protocols Layer 2 Circuit)

IN THIS SECTION

- Syntax | **28490**
- Hierarchy Level | **28490**
- Description | **28490**
- Required Privilege Level | **28490**
- Release Information | **28490**

Syntax

```
standby;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name end-interface interface interface-name backup-neighbor address],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name backup-neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor address backup-neighbor address],
[edit protocols l2circuit neighbor address interface interface-name backup-neighbor address],
[edit routing-instances routing-instance-name protocols vpls neighbor address backup-neighbor
address]
```

Description

Configure the pseudowire to the specified backup neighbor as the standby. When you configure this statement, traffic flows over both the active and standby pseudowires to the backup device (either a CE device or PE router). The backup device drops the traffic from the standby pseudowire, unless the active pseudowire fails. If the active pseudowire fails, the backup device automatically switches to the standby pseudowire.

The `standby` statement is quite similar to the `hot-standby` statement introduced in Junos OS Release 12.3. The `hot-standby` statement allows for a faster forwarding-path switchover during transition periods, as compared to what is allowed by the `standby` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Hierarchy levels associated with the `backup-neighbor` statement added in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Layer 2 Circuits over Both RSVP and LDP LSPs

Configuring Pseudowire Redundancy on the PE Router

Example: Configuring Layer 2 Circuit Switching Protection

standby (Routing Policy Condition)

IN THIS SECTION

- Syntax | 28491
- Hierarchy Level | 28491
- Description | 28491
- Required Privilege Level | 28492
- Release Information | 28492

Syntax

```
standby;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options condition if-route-exists address-family
ccc],
[edit policy-options condition if-route-exists address-family ccc]
```

Description

Specify that the route must be in standby state to be considered a match.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring Pseudowire Redundancy in a Mobile Backhaul Scenario

start-date

IN THIS SECTION

- Syntax | 28492
- Hierarchy Level | 28492
- Description | 28493
- Options | 28493
- Required Privilege Level | 28493
- Release Information | 28493

Syntax

```
start-date date-time;
```

Hierarchy Level

```
[edit schedulers scheduler scheduler-name]
```

Description

Specify the time, day, month, and year that the schedule starts.

Specifying the year is optional. If no year is specified, the schedule applies to the current year and all subsequent years. If the year is specified in either the `start-date` or `stop-date` statement, that year is used for both statements.

Options

date-time —Use the format [*yyyy* -] *mm* - *dd* . *hh* . *mm* to specify the year, month, day, hour, and minutes.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

start-log

IN THIS SECTION

- Syntax | 28494
- Hierarchy Level | 28494
- Description | 28494
- Options | 28494
- Required Privilege Level | 28494
- Release Information | 28494

Syntax

```
start-log value;
```

Hierarchy Level

```
[edit security idp sensor-configuration log suppression]
```

Description

Specify how many instances of a specific event must occur before log suppression begins.

Options

value—Log suppression begins after how many occurrences.

- **Range:** 1 through 128
- **Default:** 1

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

start-measurement

IN THIS SECTION

Syntax | 28495

- Hierarchy Level | 28495
- Description | 28495
- Default | 28495
- Options | 28495
- Required Privilege Level | 28495
- Release Information | 28496

Syntax

```
(no-start-measurement | start-measurement);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

Description

Start or do not start a delay measurement (DM) session.

Default

By default, do not start a DM session.

Options

no-start-measurement	Do not start a DM session.
start-measurement	Start a DM session.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

start (Scripts)

IN THIS SECTION

- Syntax | 28496
- Hierarchy Level | 28497
- Description | 28497
- Default | 28497
- Options | 28497
- Required Privilege Level | 28498
- Release Information | 28498

Syntax

```
start {
    start-options {
        mem-factor percentage;
    }
}
```

Hierarchy Level

```
[edit event-options event-script],
[edit event-options event-script file filename],
[edit system extensions extension-service application],
[edit system extensions extension-service application file filename],
[edit system scripts op],
[edit system scripts op file filename],
[edit system scripts snmp],
[edit system scripts snmp file filename]
```

Description

Configure options to control script execution.

> **NOTE**: Junos devices do not enforce script start options for remote op scripts, which are executed using the `op url` command. Script start options are enforced only for local op scripts.

Default

If you do not configure start options for an individual script or script type, the device does not perform any system resource checks before executing a script.

Options

mem-factor *percentage*

System memory usage threshold, specified as a percentage, above which a script does not execute.

You can configure the statement for op, event, SNMP, and Juniper Extension Toolkit (JET) scripts. You can apply the statement to an individual script or to all scripts of a given script type. If you define different thresholds for the script type and an individual script of the same type, the configuration for the individual script takes precedence.

If a device prevents the execution of a script, the cscript process records system log messages in the log file. The CLI does not issue any warning or error that the script was halted.

- Range: 1 through 100

> **NOTE**: For JET applications, the `mem-factor` and `daemonize` statements are mutually exclusive.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1.

Statement introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Configure Script Start Options

start shell

IN THIS SECTION

- Syntax | 28499
- Description | 28499
- Options | 28499
- Additional Information | 28499
- Required Privilege Level | 28500
- Output Fields | 28500
- Sample Output | 28500
- Release Information | 28500

Syntax

```
start shell (csh | sh)
<user username>
```

Description

Exit from the CLI environment and create a UNIX-level shell. To return to the CLI, type `exit` from the shell.

Juniper Networks does not provide support for operations in the shell.

NOTE:

- To issue this command, the user must have the required login access privileges configured by including the `permissions` statement at the `[edit system login class` *`class-name`*`]` hierarchy level.
- UNIX wheel group membership or permissions are no longer required to issue this command.

Options

csh	Create a UNIX C shell.
sh	Create a UNIX Bourne shell.
user *username*	(Optional) Start the shell as another user.

Additional Information

When you are in the shell, the shell prompt has the following format:

```
username@hostname%
```

An example of the prompt is:

```
root@host%
```

Required Privilege Level

shell or maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

start shell csh

```
user@host> start shell csh
%

exit
%

username@hostname% start shell sh
%

exit
user@host>
```

Release Information

Command introduced before Junos OS Release 7.4.

start-time

IN THIS SECTION

- Syntax | 28501
- Hierarchy Level | 28501
- Description | 28501

- Options | 28501
- Required Privilege Level | 28501
- Release Information | 28501

Syntax

```
start-time start-time;
```

Hierarchy Level

```
[edit security utm default-configuration anti-virus avira-engine pattern-update]
[edit security utm default-configuration anti-virus sophos-engine pattern-update]
```

Description

Specify the time that the device automatically starts downloading the updated signature database from the specified URL.

Options

start-time—Time in MM-DD.hh:mm format.

Required Privilege Level

security— To view this statement in the configuration.

security-control— To add this statement to the configuration.

Release Information

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for Avira engine added in Junos OS Release 18.4R1.

start-time (Accounting)

IN THIS SECTION

- Syntax | 28502
- Hierarchy Level | 28502
- Description | 28502
- Options | 28502
- Required Privilege Level | 28502
- Release Information | 28503

Syntax

```
start-time time;
```

Hierarchy Level

```
[edit accounting-options file filename]
```

Description

Specify the start time for transfer of an accounting-data log file.

Options

time—Start time for file transfer.

- **Syntax:** *YYYY-MM-DD.hh:mm*

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring Accounting Options, Source Class Usage and Destination Class Usage Options

start-time (Schedulers)

IN THIS SECTION

- Syntax | 28503
- Hierarchy Level | 28503
- Description | 28504
- Options | 28504
- Required Privilege Level | 28504
- Release Information | 28504

Syntax

```
start-time hh:mm;
```

Hierarchy Level

```
[edit schedulers scheduler  scheduler-name  daily],
[edit schedulers scheduler  scheduler-name  friday],
[edit schedulers scheduler  scheduler-name monday],
[edit schedulers scheduler  scheduler-name saturday],
[edit schedulers scheduler  scheduler-name  sunday],
[edit schedulers scheduler  scheduler-name   tuesday],
```

```
[edit schedulers scheduler  scheduler-name  thursday],
[edit schedulers scheduler  scheduler-name  wednesday]
```

Description

Specify the time that a schedule starts for a specified day.

If you specify a starting time for a daily schedule with the `daily` statement and also include the `friday`, `monday`, `saturday`, `sunday`, `tuesday`, `thursday`, and `wednesday` statements in the schedule, the starting time specified for a specific day (for example, Friday using the `friday` statement) overrides the starting time set with the `daily` statement.

Options

time —Use the 24-hour format (*hh*:*mm*:*ss*) to specify the hours, minutes, and seconds.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

start-time (Security IDP)

IN THIS SECTION

- Syntax | 28505
- Hierarchy Level | 28505

- Description | 28505
- Options | 28505
- Required Privilege Level | 28505
- Release Information | 28505

Syntax

```
start-time start-time;
```

Hierarchy Level

```
[edit security idp security-package automatic]
```

Description

Specify the time that the device automatically starts downloading the updated signature database from the specified URL.

Options

start-time—Time in MM-DD.hh:mm format.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

starts-with (DHCP Relay Agent Option)

IN THIS SECTION

- Syntax | 28506
- Hierarchy Level | 28506
- Description | 28509
- Options | 28509
- Required Privilege Level | 28509
- Release Information | 28510

Syntax

```
equals {
    ascii name {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
    hexadecimal name {
        drop drop;
        forward-only forward-only;
        local-server-group local-server-group;
    }
}
```

Hierarchy Level

```
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay group name relay-option option-77],
[edit bridge-domains name forwarding-options dhcp-relay relay-option],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-60],
[edit bridge-domains name forwarding-options dhcp-relay relay-option option-77],
[edit forwarding-options dhcp-relay group name relay-option],
```

```
[edit forwarding-options dhcp-relay group name relay-option option-60],
[edit forwarding-options dhcp-relay group name relay-option option-77],
[edit forwarding-options dhcp-relay relay-option],
[edit forwarding-options dhcp-relay relay-option option-60],
[edit forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit logical-systems name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay group name relay-option option-77],
[edit logical-systems name forwarding-options dhcp-relay relay-option],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems name forwarding-options dhcp-relay relay-option option-77],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option option-60],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay group name relay-option option-77],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option option-60],
[edit logical-systems name routing-instances name bridge-domains name forwarding-options dhcp-
relay relay-option option-77],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option option-60],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay group name relay-
option option-77],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option],
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option
option-60],
```

```
[edit logical-systems name routing-instances name forwarding-options dhcp-relay relay-option
option-77],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option option-60],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay group
name relay-option option-77],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option option-60],
[edit logical-systems name routing-instances name vlans name forwarding-options dhcp-relay relay-
option option-77],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option
option-60],
[edit logical-systems name vlans name forwarding-options dhcp-relay group name relay-option
option-77],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option option-60],
[edit logical-systems name vlans name forwarding-options dhcp-relay relay-option option-77],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay group name relay-
option option-77],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-60],
[edit routing-instances name bridge-domains name forwarding-options dhcp-relay relay-option
option-77],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay group name relay-option option-77],
[edit routing-instances name forwarding-options dhcp-relay relay-option],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name forwarding-options dhcp-relay relay-option option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay group name relay-option
```

```
option-77],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-60],
[edit routing-instances name vlans name forwarding-options dhcp-relay relay-option option-77],
[edit vlans name forwarding-options dhcp-relay group name relay-option],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-60],
[edit vlans name forwarding-options dhcp-relay group name relay-option option-77],
[edit vlans name forwarding-options dhcp-relay relay-option],
[edit vlans name forwarding-options dhcp-relay relay-option option-60],
[edit vlans name forwarding-options dhcp-relay relay-option option-77]
```

Description

Configure a partial match criteria used with the DHCP relay agent selective processing feature. DHCP relay agent compares the configured partial match string with the option-specific string received in DHCP client packets. If there is an partial left-to-right match, DHCP performs the action you define for the match criteria.

The option-specific string in the DHCP client packets can contain a superset of the specified ASCII or hexadecimal match string, provided that the leftmost characters of the option-specific string entirely match the characters in the configured match string.

You can configure an unlimited number of match strings. If you have multiple partial match configurations, the longest match rule applies. For example, DHCP relay agent matches the string “test123” before it matches the string “test”. Match strings do not support wildcard attributes.

The `local-server-group` option is not supported for DHCPv6 relay agent.

Options

ascii-string ASCII string of 1 through 255 alphanumeric characters.

hexadecimal-string Hexadecimal string of 1 through 255 hexadecimal characters (0 through 9, a through f, A through F).

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Statement updated in Junos OS Release 17.4.

RELATED DOCUMENTATION

DHCP Options and Selective Traffic Processing

DHCP Options and Selective Traffic Processing

starts-with (Event Policy)

IN THIS SECTION

- Syntax | **28510**
- Hierarchy Level | **28510**
- Description | **28511**
- Options | **28511**
- Required Privilege Level | **28511**
- Release Information | **28511**

Syntax

```
event1.attribute-name starts-with event2.attribute-name;
```

Hierarchy Level

```
[edit event-options policy policy-name attributes-match event1.attribute-name]
```

Description

Configure an event policy to trigger only if the specified attribute of *event1* starts with the specified attribute of *event2*.

Options

event1.attribute-name	Attribute of one event.
event2.attribute-name	Attribute of another event.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Use Correlated Events to Trigger an Event Policy

startup-alarm

IN THIS SECTION

- Syntax | 28512
- Hierarchy Level | 28512
- Description | 28512
- Options | 28512
- Required Privilege Level | 28512
- Release Information | 28512

Syntax

```
startup-alarm (falling-alarm | rising-alarm | rising-or-falling-alarm);
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Set an initial alarm that is sent after the configured SNMP RMON alarm becomes active.

Options

`falling-alarm`—Generated if the first sample after the alarm entry becomes active is less than or equal to the falling threshold.

`rising-alarm`—Generated if the first sample after the alarm entry becomes active is greater than or equal to the rising threshold.

`rising-or-falling-alarm`—Generated if the first sample after the alarm entry becomes active satisfies either of the corresponding thresholds.

- **Default:** `rising-or-falling-alarm`

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Remote Nework Monitoring (RMON)

startup-silent-period

IN THIS SECTION

- Syntax | 28513
- Hierarchy Level | 28513
- Description | 28513
- Options | 28513
- Required Privilege Level | 28514
- Release Information | 28514

Syntax

```
startup-silent-period seconds;
```

Hierarchy Level

```
[edit protocols vrrp]
```

Description

Instruct the system to ignore the Primary Down Event when an interface transitions from the down state to the up state. This statement is used to avoid incorrect error alarms caused by the delay or interruption of incoming Virtual Router Redundancy Protocol (VRRP) advertisement packets during the interface startup phase.

Options

seconds—Number of seconds for the startup period.

- **Default:** 4 seconds
- **Range:** 1 through 2000 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Startup Period for VRRP Operations

stateful (Router Advertisement Guard)

IN THIS SECTION

- Syntax | 28514
- Hierarchy Level | 28515
- Description | 28515
- Default | 28515
- Required Privilege Level | 28515
- Release Information | 28515

Syntax

```
(stateful | stateless);
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard interface  (interface-name |
interface-range-name)]
[edit forwarding-options access-security router-advertisement-guard vlans  vlan-name]
```

Description

Configure stateful IPv6 Router Advertisement (RA) guard. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. RA guard performs checks on incoming RA messages to make sure that they are sent from legitimate routers. If the sender of the RA message cannot be validated, the RA message is dropped.

Stateful RA guard enables the switch to learn about the sources of RA messages for a certain period of time. During this period, when the switch is known to be in the learning state, the information contained in attributes of received RA messages is stored and compared to the policy. When the learning period ends, the switch has a record of which interfaces are attached to links with valid IPv6 routers. If there is no valid IPv6 router attached to the interface, the switch dynamically transitions the interface from the learning state into the blocking state. Subsequent RA messages received after the transition to blocking state are dropped. If there is a valid IPv6 router attached to the interface, the interface transitions to the forwarding state, and subsequent RA messages that can be validated against the configured policy are forwarded.

You can enable stateful RA guard on an interface or on a VLAN. When you enable stateful RA guard, the initial state is `Off`. You initiate the learning state by issuing the `request access-security router-advertisement-guard-learn` command.

Default

RA guard is stateless by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

stateful (Precision Time Protocol)

IN THIS SECTION

- Syntax | 28516
- Hierarchy Level | 28517
- Description | 28517
- Options | 28517
- Required Privilege Level | 28517
- Release Information | 28517

Syntax

```
stateful {
    interface interface-name {
        multicast-mode {
            local-priority
            transport 802.3 ( link-local ) ;
            asymmetry number;
            primary-asymmetry number;
            secondary-asymmetry number
        }
        passive-port-monitor {
            measure-only;
            phase-threshold <offset-threshold-value>;
        }
        not-master;
    }
}
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

Configure the stateful port or bidirectional port with parameters. A stateful port is useful when you configure PTP over Ethernet for multicast mode of transmission of PTP traffic. A stateful port configuration for PTP over Ethernet and default profile is supported only on boundary clock mode and not on ordinary clock mode.

For PTP over Ethernet, you can configure a port to function as both a timeReceiver port and a timeTransmitter port. This type of port is called a stateful port, or a bidirectional port. Such a stateful port enables the transfer of frequency for synchronization services, in addition to time and phase alignment, when PTP functionality is not hop-by-hop and you have provisioned timeTransmitter and timeReceiver roles or interfaces.

The `passive-port-monitor` and `not-master` are only available for PTX10004 and PTX10008.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

Precision Time Protocol Overview

Guidelines for Configuring PTP over Ethernet

Configuring PTP Multicast Primary and Member Ports for Ethernet Encapsulation

Configuring PTP Dynamic Ports for Ethernet Encapsulation

stateful-firewall-rule-set (Next Gen Services)

IN THIS SECTION

- Syntax | 28518
- Hierarchy Level | 28518
- Description | 28518
- Options | 28518
- Required Privilege Level | 28519
- Release Information | 28519

Syntax

```
stateful-firewall-rule-set {
    stateful-firewall-rule [rule-name];
]
```

Hierarchy Level

```
[edit services policies]
```

Description

Specify a set of stateful firewall rules, which are processed in the order in which they appear in the rule set configuration. Once a stateful firewall rule in the rule set matches a flow, that rule is applied and no other rules in the rule set are processed˙.

Options

stateful-firewall-rule [*rule-name*] Names of the stateful firewall rules that belong to the rule set. A stateful firewall rule is configured at the `[edit services policies]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Stateful Firewalls for Next Gen Services

stateful-firewall-rule-sets (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 28519
- Hierarchy Level | 28519
- Description | 28520
- Required Privilege Level | 28520
- Release Information | 28520

Syntax

```
stateful-firewall-rule-sets [rule-set-name];
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the stateful firewall rule sets to be used with the service set. A stateful firewall rule set is configured at the `[edit services policies]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Stateful Firewalls for Next Gen Services

stateful-firewall-rules

IN THIS SECTION

- Syntax | 28520
- Hierarchy Level | 28521
- Description | 28521
- Options | 28521
- Required Privilege Level | 28521
- Release Information | 28521

Syntax

```
(stateful-firewall-rules rule-names | stateful-firewall-rule-sets rule-set-name);
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the stateful firewall rules or rule set included in this service set. You can configure multiple rules, but only one rule set for each service.

Options

rule-name—Identifier for the collection of terms that make up this rule.

rule-set-name—Identifier for the set of rules to be included.

Required Privilege Level

System—To view this statement in the configuration.

System-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Service Rules

stateful-firewall-rules (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 28522
- Hierarchy Level | 28522

- Description | 28522
- Required Privilege Level | 28522
- Release Information | 28522

Syntax

```
stateful-firewall-rules [rule-name];
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the stateful firewall rules to be used with the service set. A stateful firewall rule is configured at the `[edit services policies]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Stateful Firewalls for Next Gen Services

stateful-nat64

IN THIS SECTION

- Syntax | 28523
- Hierarchy Level | 28523
- Description | 28523
- Options | 28524
- Required Privilege Level | 28524
- Release Information | 28524

Syntax

```
stateful-nat64 {
    clear-dont-fragment-bit;
    disable-h323-ras;
    ipv6-mtu;

}
```

Hierarchy Level

```
[edit services service-set service-set-name nat-options]
```

Description

Set parameters for stateful NAT64 operation.

NOTE: These parameters do not change the operation of other types of NAT.

Options

`clear-dont-fragment-bit` Clear the DF (don't fragment) bit in a translated IPv4 packet if its packet size is less than 1280 bytes. If the packet is greater than or equal to 1280 bytes, the DF bit is not cleared.

`disable-h323-ras` Disable the H323 and RAS ALG for NAT64.

`ipv6-mtu` Configure path Maximum transmission unit (MTU) of IPv6 network. Range is 1280 through 9192 bytes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced with Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Service Rules

clear-dont-fragment-bit (Services NAT Options) | **3561**

stateless

IN THIS SECTION

- Syntax | **28525**
- Hierarchy Level | **28525**
- Description | **28525**
- Default | **28525**
- Required Privilege Level | **28525**
- Release Information | **28525**

Syntax

```
(stateful | stateless);
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard interface  (interface-name |
interface-range-name)]
[edit forwarding-options access-security router-advertisement-guard vlans  vlan-name]
```

Description

Configure stateless IPv6 Router Advertisement (RA) guard. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in a policy.

You can configure RA guard to be stateless or stateful. If stateless RA guard is enabled, the switch examines incoming RA messages and filters each message on the basis of whether it matches the conditions configured in the policy. After the switch has validated the content of the RA message, it forwards the RA message to its destination; otherwise, the RA message is dropped. For example, an interface can be statically configured to forward RA messages only from predefined sources.

You can enable stateless RA guard on an interface or on a VLAN.

Default

RA guard is stateless by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

Understanding IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

static (802.1X)

IN THIS SECTION

- Syntax | 28526
- Hierarchy Level | 28526
- Description | 28526
- Options | 28527
- Required Privilege Level | 28527
- Release Information | 28527

Syntax

```
static mac-address {
    bridge-domain-assignment bridge-domain-assignment;
        interface [interface-names];
    vlan-assignment (vlan-id |vlan-name );
}
```

Hierarchy Level

```
[edit protocols dot1x                 authenticator]
```

Description

Configure MAC addresses to exclude from 802.1X authentication. The static MAC list provides an authentication bypass mechanism for supplicants connecting to a port, permitting devices such as printers that are not 802.1X-enabled to be connected to the network on 802.1X-enabled ports.

Using this 802.1X authentication-bypass mechanism, the supplicant connected to the MAC address is assumed to be successfully authenticated and the port is opened for it. No further authentication is done for the supplicant.

You can optionally configure the VLAN so that the supplicant is moved to or the interfaces on which the MAC address can gain access from.

Options

mac-address	The MAC address of the device for which 802.1X authentication should be bypassed and the device permitted access to the port.
bridge-domain-assignment *bridge-domain-assignment*	(MX Series only) Specify the bridge-domain name or 802.1q tag identifier for the MAC address that should be allowed to bypass RADIUS authentication.
interface [*interface-names*]	Specify a list of interfaces on which the specified MAC addresses are allowed to bypass RADIUS authentication and allowed to connect to the LAN without authentication.
vlan-assignment (*vlan-id* | *vlan-name*	(EX, QFX, and SRX Series only) Specify the VLAN 802.1q tag identifier or VLAN name associated with the list of MAC addresses that should be allowed to bypass RADIUS authentication.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show dot1x static-mac-address | **19577**

Example: Configuring Static MAC Bypass of 802.1X and MAC RADIUS Authentication on an EX Series Switch

Configuring 802.1X Interface Settings (CLI Procedure)

Configuring 802.1X Authentication (J-Web Procedure)

Understanding Authentication on Switches

static-binding

IN THIS SECTION

- Syntax | 28528
- Hierarchy Level | 28528
- Description | 28528
- Options | 28529
- Required Privilege Level | 28529
- Release Information | 28529

Syntax

```
static-binding mac-address {
    client-identifier (ascii client-id | hexadecimal client-id);
    fixed-address {
        address;
    }
    host-name client-hostname;
}
```

Hierarchy Level

```
[edit system services dhcp],
[edit system services dhcp]
```

Description

For J Series Services routers and EX Series switches only. Set static bindings for DHCP clients. A static binding is a mapping between a fixed IP address and the client's MAC address or client identifier.

Options

mac-address—The MAC address of the client. This is a hardware address that uniquely identifies a client on the network.

`fixed-address` *address*—Fixed IP address assigned to the client. Typically a client has one address assigned, but you can assign more.

`host-name` *client-hostname*—Hostname of the client requesting the DHCP server. The name can include the local domain name. Otherwise, the name is resolved based on the `domain-name` statement.

`client-identifier (ascii` *client-id* `| hexadecimal` *client-id*`)`—Used by the DHCP server to index the database of address bindings. The client identifier is an ASCII string or hexadecimal number and can include a type-value pair as specified in RFC 1700, *Assigned Numbers*. Either a client identifier or the client's MAC address must be configured to uniquely identify the client on the network.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

static (Bridge Domains)

IN THIS SECTION

- Syntax | 28530
- Hierarchy Level | 28530
- Description | 28530
- Required Privilege Level | 28530
- Release Information | 28530

Syntax

```
static {
    group multicast-group-address {
        source ip-address;
    }
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping interface interface-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id interface interface-name]
```

Description

Define static multicast groups on an interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring IGMP Snooping

static (Distributed IGMP)

IN THIS SECTION

- Syntax | 28531
- Hierarchy Level | 28531
- Description | 28531
- Options | 28532
- Required Privilege Level | 28532
- Release Information | 28532

Syntax

```
static {
    <distributed>;
    group multicast-group-address{
        <distributed>;
        source  source- address<distributed>;
    }
}
```

Hierarchy Level

```
[edit protocols pim]
```

Description

Configure static source and group (S, G) addresses when distributed IGMP is enabled. Reduces the first join delay time and brings multicast traffic to the last-hop router. Specified (S, G) addresses join statically without waiting for the first join.

Options

distributed (Optional) Enable static joins for specified (S,G) addresses and preprovision all of them so that all distributed IGMP Packet Forwarding Engines receive traffic.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X50.

RELATED DOCUMENTATION

Enabling Distributed IGMP

Junos OS Multicast Protocols User Guide

Junos OS Multicast Protocols User Guide

static (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | 28533
- Hierarchy Level | 28533
- Description | 28533
- Options | 28533
- Required Privilege Level | 28533
- Release Information | 28533

Syntax

```
static {
    group group;
    group group {
        source source;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name]
```

Description

Test multicast forwarding on an interface without a receiver host.

Options

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Enabling IGMP Static Group Membership

static (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 28534
- Hierarchy Level | 28534
- Description | 28534
- Required Privilege Level | 28535
- Release Information | 28535

Syntax

```
static {
    group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
            source-count number;
            source-increment increment;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Description

Test multicast forwarding on an interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Enabling MLD Static Group Membership

static-host-mapping

IN THIS SECTION

- Syntax | 28535
- Hierarchy Level | 28536
- Description | 28536
- Default | 28536
- Options | 28536
- Required Privilege Level | 28536
- Release Information | 28537

Syntax

```
static-host-mapping {
    hostname {
        alias [ aliases ];
        inet [ addresses ];
```

```
        inet6 [ addresses];
        sysid system-identifier;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

(Optional) Statically map a hostname to one or more IP addresses and aliases, and configure an International Organization for Standardization (ISO) system identifier (system ID).

Default

If you do not statically map the hostname, the mapping is generated dynamically, based on the system configuration. For instance, if you omit the `static-host-mapping` *`hostname`* `sysid` statement, the IS-IS system ID is dynamically generated from the host portion of the ISO address configured on the loopback interface (lo0) and is mapped to the `host-name` statement configured at the `[edit system]` hierarchy level.

Options

alias ***alias***	Alias for the hostname.
hostname	Fully qualified hostname.
inet ***address***	IP address. You can specify one or more IP addresses for the host.
sysid ***system-identifier***	ISO system identifier (system ID). This is the 6-byte portion of the Intermediate System-to-Intermediate System (IS-IS) network service access point (NSAP). We recommend that you use the host's IP address represented in binary-coded decimal (BCD) format. For example, the IP address 208.197.169.18 is 2081.9716.9018 in BCD.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1. for the QFX Series.

RELATED DOCUMENTATION

No Link Title

static (IGMP Snooping)

IN THIS SECTION

- Syntax | 28537
- Hierarchy Level | 28537
- Description | 28538
- Default | 28538
- Required Privilege Level | 28538
- Release Information | 28538

Syntax

```
static {
    group ip-address;
}
```

Hierarchy Level

```
[edit protocols igmp-snooping vlan  (all | vlan-name) interface  interface-name
```

Description

Statically define multicast groups for IGMP snooping on a Layer 2 (L2) interface.

You might use this statement to configure an L2 interface to statically join a particular multicast group in cases such as the following:

- In a network that doesn't include Layer 3 (L3) multicast. You can't establish a static IGMP state at L3, but you can create a static join state on an interface for multicast forwarding at L2.
- To enable static IGMP join states with more granularity in a network with L2 and L3 multicast elements.

  For example, consider that you have Layer 2 interfaces in a VLAN associated with an IRB interface for Layer 3 multicast routing. You have two options to establish a static join state on an interface at the two different levels:

  - Assign a static group membership on an individual Layer 2 interface in the VLAN associated with the IRB interface:

    In this case, the device creates IGMP and PIM join states for the static group at L3 on the IRB interface. However, with IGMP snooping enabled and the group membership established only on that individual L2 interface, the device forwards traffic for the group only on that interface.

  - Assign a static group membership only on the IRB interface itself at L3:

    In this case, without enabling IGMP snooping, the device creates creates IGMP and PIM join states for the group at L3 on the IRB interface, pulls multicast traffic, and floods it on all L2 ports in the VLAN. With IGMP snooping enabled, the device pulls multicast traffic for the group, but doesn't forward it on the VLAN (because you didn't configure any particular L2 ports to join that multicast group).

The remaining statement is explained separately. See CLI Explorer.

Default

No multicast groups are statically defined.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

show igmp snooping membership | 20165

show igmp-snooping vlans | 20198

static-interface

IN THIS SECTION

- Syntax | 28539
- Hierarchy Level | 28539
- Description | 28539
- Options | 28540
- Required Privilege Level | 28540
- Release Information | 28540

Syntax

```
static-interface interface-name;
```

Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci
circuit-id-string ari remote-id-string]
```

Description

Reserve the specified static PPPoE interface for use only by the PPPoE client with matching agent circuit identifier (ACI) and agent remote identifier (ARI) information. You can specify only one static interface per ACI/ARI pair configured for a named service entry, `empty` service entry, or `any` service entry in the PPPoE service name table.

The static interface associated with an ACI/ARI pair takes precedence over the general pool of static interfaces associated with the PPPoE underlying interface.

If you include the `static-interface` statement in the configuration, you cannot also include either the `dynamic-profile` statement or the `routing-instance` statement. The `dynamic-profile`, `routing-instance`, and `static-interface` statements are mutually exclusive for ACI/ARI pair configurations.

Options

`interface-name`—Name of the static PPPoE interface reserved for use by the PPPoE client with matching ACI/ARI information. Specify the interface in the format `pp0.`*`logical`*, where *`logical`* is a logical unit number from 0 through 16385 for static interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client

static-ip

IN THIS SECTION

- Syntax | 28541
- Hierarchy Level | 28541
- Description | 28541
- Options | 28542

Required Privilege Level | 28542
Release Information | 28542

Syntax

```
static-ip ip-addresses {
    vlan  vlan-name;
    mac mac-address;
}
```

Hierarchy Level

- For platforms with ELS:

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name interface interface-name]
```

- For platforms without ELS:

```
[edit ethernet-switching-options secure-access-port interface (all | interface-name)]
```

Description

Configure a static IP address to MAC address (IP-MAC) binding to be added to the DHCP snooping database.

NOTE: The VLAN is specified at the higher hierarchy level when `static-ip` is configured at `[edit vlans vlan-name forwarding-options dhcp-security group group-name interface interface-name]`.

Options

ip-address—Static IP address assigned to a device connected on the specified interface.

`mac` *mac-address*—Static MAC address assigned to a device connected on the specified interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Hierarchy level `[edit vlans` *vlan-name* `forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

RELATED DOCUMENTATION

Configuring Static DHCP IP Addresses for DHCP snooping (non-ELS)

Configuring Static DHCP IP Addresses for DHCP snooping (ELS)

static-ip (MX Series)

IN THIS SECTION

- Syntax | 28543
- Hierarchy Level | 28543
- Description | 28543
- Options | 28543
- Required Privilege Level | 28543
- Release Information | 28543

Syntax

```
static-ip ip-address mac mac-address;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security group group-name interface
interface-name]
```

Description

Configure a static IP address to MAC address (IP-MAC) binding record to be added to the DHCP snooping database.

Options

ip-address—Static IP address assigned to a device connected on the specified interface.

mac-address—Static MAC address assigned to a device connected on the specified interface.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Juos OS Release 14.1.

RELATED DOCUMENTATION

> *Configuring Static DHCP IP Addresses for DHCP snooping (MX routers)*

static-ipv6

IN THIS SECTION

- Syntax | 28544
- Hierarchy Level | 28544
- Description | 28544
- Options | 28544
- Required Privilege Level | 28545
- Release Information | 28545

Syntax

```
static-ipv6 ip-address {
    mac mac-address;
}
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name interface interface-name],
[edit ethernet-switching-options secure-access-port interface interface-name]
```

Description

Configure a static IP-MAC binding to be added to the DHCPv6 snooping database.

Options

ip-address—Static IPv6 address assigned to a device connected on the specified interface.

`mac` *mac-address*—Static MAC address assigned to a device connected on the specified interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D20.

Support at the `[edit ethernet-switching-options secure-access-port interface interface-name]` hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

RELATED DOCUMENTATION

Configuring Static DHCP IP Addresses for DHCP snooping (ELS)

static-label-switched-path

IN THIS SECTION

- Syntax | 28545
- Hierarchy Level | 28546
- Description | 28547
- Options | 28547
- Required Privilege Level | 28547
- Release Information | 28547

Syntax

```
static-label-switched-path lsp-name {
    bypass bypass-name {
        bandwidth bps;
        description string;
        next-hop (address | interface-name | address/interface-name);
```

```
        next-table string
        push out-label;
        to address;
    }
    ingress {
        bandwidth bps;
        class-of-service cos-value;
        description string;
        install {
            destination-prefix <active>;
        }
        link-protection bypass-name name;
        metric metric;
        next-hop (address | interface-name | address/interface-name);
        node-protection bypass-name name next-next-label label;
        no-install-to-address;
        policing {
            filter filter-name;
            no-auto-policing;
        }
        preference preference;
        push out-label;
        to address;
    }
    transit incoming-label {
        bandwidth bps;
        description string;
        link-protection bypass-name name;
        next-hop (address | interface-name | address/interface-name);
        node-protection bypass-name name next-next-label label;
        pop;
        swap out-label;
    }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Configure a static LSP.

Options

lsp-name—Name of the path.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring Static LSPs

static-lsp

IN THIS SECTION

- Syntax | 28548
- Hierarchy Level | 28548
- Description | 28548
- Required Privilege Level | 28549
- Release Information | 28549

Syntax

```
static-lsp lsp-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel rsvp-te],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group address source source-address rsvp-te],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group address wildcard-source rsvp-te],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet wildcard-source rsvp-te],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet6 wildcard-source rsvp-te],
[edit routing-instances routing-instance-name provider-tunnel rsvp-te],
[edit routing-instances routing-instance-name provider-tunnel selective group address source
source-address rsvp-te],
[edit routing-instances routing-instance-name provider-tunnel selective group address wildcard-
source rsvp-te],
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet
wildcard-source rsvp-te],
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet6
wildcard-source rsvp-te]
```

Description

Specify the name of the static point-to-multipoint (P2MP) LSP used for a specific MBGP MVPN; static P2MP LSP cannot be shared by multiple VPNs. Use this statement to specify the static LSP for both inclusive and selective point-to-multipoint LSPs.

Use a static P2MP LSP when you know all the egress PE router endpoints (receiver nodes) and you want to avoid the setup delay incurred by dynamically created P2MP LSPs (configured with the `label-switched-path-template`). These static LSPs are signaled before the MVPN requires or uses them, consequently avoiding any signaling latency and minimizing traffic loss due to latency.

If you add new endpoints after the static P2MP LSP is established, you must update the configuration on the ingress PE router. In contrast, a dynamic P2MP LSP learns new endpoints without any configuration changes.

BEST PRACTICE: Multiple multicast flows can share the same static P2MP LSP; this is the preferred configuration when the set of egress PE router endpoints on the LSP are all interested in the same set of multicast flows. When the set of relevant flows is different between endpoints, we recommend that you create a new static P2MP LSP to associate endpoints with flows of interest.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

- Point-to-Multipoint LSPs Overview
- Configuring Static LSPs
- *Configuring Point-to-Multipoint LSPs for an MBGP MVPN*
- Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems

static-mac

IN THIS SECTION

- Syntax | 28550
- Hierarchy Level | 28550
- Description | 28550
- Options | 28550
- Required Privilege Level | 28551
- Release Information | 28551

Syntax

```
static-mac mac-address;
```

```
static-mac mac-address {
    vlan-id number;
}
```

Hierarchy Level

```
[edit vlans vlan-name switch-options interface interface-name]
```

```
[edit bridge-domains bridge-domain-name bridge-options interface interface-name],
```

```
[edit logical-systems logical-system-name bridge-domains bridge-domain-name bridge-options
interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name bridge-options interface interface-name],
```

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options
interface interface-name],
[edit routing-instances routing-instance-name protocols evpn interface interface-name]
```

Description

Configure a static MAC address for a logical interface in a bridge domain or VLAN.

The `vlan-id` option can be specified for `static-macs` only if `vlan-id all` is configured for the bridging domain or VLAN.

Options

mac-address—MAC address

`vlan-id` *number*—(Optional) VLAN identifier to associate with static MAC address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement modified in Junos OS Release 9.5.

Support for logical systems added in Junos OS Release 9.6.

[edit vlans *vlan-name* switch-options interface *interface name*] hierarchy level introduced in Junos OS Release 12.3R2 for EX Series switches.

Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers. The `vlan-id` option is not available for EVPNs.

[edit vlans *vlan-name* switch-options interface *interface name*] hierarchy level introduced in Junos OS Release 13.2 for the QFX Series.

RELATED DOCUMENTATION

- *Configuring EVPN Routing Instances*
- Understanding Layer 2 Learning and Forwarding for Bridge Domains
- Layer 2 Learning and Forwarding for VLANs Overview
- Adding a Static MAC Address Entry to the Ethernet Switching Table on a Switch with ELS Support

static-nat

IN THIS SECTION

- Syntax | **28552**
- Hierarchy Level | **28552**

- Description | 28553
- Options | 28553
- Required Privilege Level | 28553
- Release Information | 28553

Syntax

```
static-nat {
    inet {
        routing-instance (default | routing-instance-name);
    }
    nptv6-prefix {
        address-prefix;
            routing-instance routing-instance-name;
    nptv6-prefix-name {
        address-prefix-name;
            routing-instance routing-instance-name;
    [
    prefix {
        address-prefix;
        routing-instance (default | routing-instance-name);
    }
    prefix-name {
        address-prefix-name;
        routing-instance (default | routing-instance-name);
    }
    rule-session-count-alarm (clear threshold value | raise threshold value);
}
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name then]
```

Description

Specify the translated address of the static NAT rule.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.3X48-D25.

static (Origin Validation for BGP)

IN THIS SECTION

- Syntax | 28553
- Hierarchy Level | 28554
- Description | 28554
- Required Privilege Level | 28554
- Release Information | 28554

Syntax

```
static {
    record destination {
        maximum-length prefix-length {
            origin-autonomous-system as-number {
```

```
                validation-state (invalid | valid);
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation],
[edit logical-systems logical-system-name routing-options validation],
[edit routing-instances instance-name routing-options validation],
[edit routing-options validation]
```

Description

Configure a static route validation (RV) record.

RV records are received from the RPKI cache server using the protocol defined in Internet draft draft-ietf-sidr-rpki-rtr-19, *The RPKI/Router Protocol*, and can also be configured statically, as shown here.

Static records are useful for overwriting the information received from an RPKI cache server.

An RV record matches any route whose prefix matches the RV prefix "record" on page 14138, whose prefix length does not exceed the "maximum-length" on page 10531 given in the RV record, and whose origin AS equals the "origin-autonomous-system" on page 11932 number given in the RV record.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

static-policy-control

IN THIS SECTION

- Syntax | 28555
- Hierarchy Level | 28555
- Description | 28556
- Required Privilege Level | 28556
- Release Information | 28556

Syntax

```
static-policy-control {
    pcc-rules {
        [rule-name precedence number <time-of-day-profile profile-name>];
    }
    pcc-rulebases {
        [rulebase-name <time-of-day-profile profile-name>];
    }
}
```

Hierarchy Level

```
[edit unified-edge pcef profiles profile-name],
[edit services pcef profiles profile-name]
```

Description

Configure static policy control for the policy and charging control (PCC) rules or PCC rulebase in a policy and charging enforcement function (PCEF) profile. You can configure a maximum of 32 PCC rules in a PCEF profile. There is no limit to the number of PCC rulebases you can configure in a PCEF profile.

> **NOTE**: For Junos OS Subscriber Aware, you can configure only one of the following statements in a PCEF profile: `aaa-policy-control`, `static-policy-control`, or `dynamic-policy-control`. For Junos OS Subscriber Management, you can configure only `static-policy-control`.

If you are using Junos OS Subscriber Aware, configure static policy control at the [`edit unified-edge pcef profiles` *`profile-name`*] hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure static policy control at the [`edit services pcef profiles` *`profile-name`*] hierarchy level.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the [`edit services pcef profiles` *`profile-name`*] hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Static Policies

Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management

static (Protocols IGMP)

IN THIS SECTION

- Syntax | 28557
- Hierarchy Level | 28558
- Description | 28558
- Required Privilege Level | 28558
- Release Information | 28558

Syntax

```
static {
    group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
            source-count number;
            source-increment increment;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
```

Description

Test multicast forwarding on an interface without a receiver host.

The `static` statement simulates IGMP joins on a routing device statically on an interface without any IGMP hosts. It is supported for both IGMPv2 and IGMPv3 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.

> **NOTE**: To prevent joining too many groups accidentally, the `static` statement is not supported with the `interface all` statement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling IGMP Static Group Membership

static (Protocols Layer 2 Circuit)

IN THIS SECTION

- Syntax | 28559
- Hierarchy Level | 28559
- Description | 28559
- Options | 28560
- Required Privilege Level | 28560
- Release Information | 28560

Syntax

```
static {
    incoming-label label;
    outgoing-label label;
    send-oam;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name l2circuit neighbor address interface interface-name],
[edit logical-systems logical-system-name l2circuit neighbor address interface interface-name
backup-neighbor neighbor],
[edit protocols l2circuit neighbor address interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name backup-neighbor neighbor]
```

Description

Configures static Layer 2 circuit pseudowires. Static pseudowires are designed for networks that do not support LDP or do not have LDP enabled. You configure a static pseudowire by configuring static values for the in and out labels needed to enable a pseudowire connection.

Options

incoming-label (Optional for PTX Series Packet Transport Routers only) Configure the Layer 2 circuit incoming static pseudowire label.

- **Range:** 1000000 through 1048575

outgoing-label (Optional for PTX Series Packet Transport Routers only) Configure the Layer 2 circuit outgoing static pseudowire label.

- **Range:** 16 through 1048575

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Static Layer 2 Circuits

static (Protocols MLD)

IN THIS SECTION

- Syntax | **28561**
- Hierarchy Level | **28561**
- Description | **28561**
- Required Privilege Level | **28561**
- Release Information | **28562**

Syntax

```
static {
    group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
            source-count number;
            source-increment increment;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

Test multicast forwarding on an interface.

The **static** statement simulates MLD joins on a routing device statically on an interface without any MLD hosts. It is supported for both MLDv1 and MLDv2 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.

> **NOTE**: To prevent joining too many groups accidentally, the `static` statement is not supported with the `interface all` statement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling MLD Static Group Membership

static (Protocols PIM)

IN THIS SECTION

- Syntax | 28562
- Hierarchy Level | 28563
- Description | 28563
- Required Privilege Level | 28563
- Release Information | 28563

Syntax

```
static {
    address address {
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
        override;
        version version;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim rp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp],
[edit protocols pim rp],
[edit routing-instances routing-instance-name protocols pim rp]
```

Description

Configure static RP addresses. The default static RP address is 224.0.0.0/4. To configure other addresses, include one or more `address` statements. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.

For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Static PIM RP Address on the Non-RP Routing Device

static (Protocols VPLS)

IN THIS SECTION

- Syntax | 28564
- Hierarchy Level | 28564
- Description | 28565
- Options | 28565
- Required Privilege Level | 28565
- Release Information | 28565

Syntax

```
static {
    incoming-label label;
    outgoing-label label;
    control-word;
    no-control-word;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor address backup-neighbor address],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name neighbor
address],
[edit routing-instances routing-instance-name protocols vpls neighbor address],
[edit routing-instances routing-instance-name protocols vpls neighbor address backup-neighbor
address]
```

Description

Specifies a static pseudowire for a VPLS domain. By configuring static pseudowires for the VPLS domain, you do not need to configure the LDP or BGP protocols that would normally be used for signaling. Static pseudowires require that you configure a set of in and out labels for each pseudowire configured for the VPLS domain. You can configure both static and dynamic neighbors within the same VPLS routing instance. You can also configure a static pseudowire for a backup neighbor (if you configure the neighbor as static the backup must also be static) and for a mesh group.

Options

`incoming-label` *`label`*—You must configure an incoming label for the static pseudowire.

- **Range:** 29,696 through 41,983 and 1,000,000 through 1,048,575

`outgoing-label` *`label`*—You must configure an outgoing label for the static pseudowire.

- **Range:** 16 through 1,048,575

`control-word`—Set `control-word` to request that other routers insert a control word between the label stack and the MPLS payload.

`no-control-word`—Set `no-control-word` to request that other routers not insert a control word between the label stack and the MPLS payload. This is the default setting for BGP VPLS.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Options `control-word` and `no-control-word` are introduced in Junos OS Release 20.3R1.

NOTE: Options `control-word` and `no-control-word` are not introduced at the [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls neighbor *address* backup-neighbor *address*] and [edit routing-instances *routing-instance-name* protocols vpls neighbor *address* backup-neighbor *address*] hierarchy level.

RELATED DOCUMENTATION

Configuring Static Pseudowires for VPLS

control-word | 4644

no-control-word | 11426

static-remote-vtep-list

IN THIS SECTION

- Syntax | 28566
- Hierarchy Level | 28566
- Description | 28566
- Options | 28570
- Required Privilege Level | 28570
- Release Information | 28570

Syntax

```
static-remote-vtep-list [ remote-vtep-loopback-address ... ];
```

Hierarchy Level

```
[edit bridge-domains name vxlan],
[edit routing-instances name bridge-domains name vxlan],
[edit routing-instances name vlans name vxlan],
[edit vlans name vxlan]
```

Description

Statically configure a list of one or more remote virtual tunnel endpoints (VTEPs) on an Layer 2 (L2) VXLAN gateway device, and map the list to a particular bridge domain or VLAN.

You can configure a static VXLAN configuration at the VLAN level or bridge domain level as follows:

- In the default switch instance, supported on:
  - MX Series routers—Bridge domain and VLAN levels.
  - QFX Series switches—VLAN level.
- In a configured `virtual-switch` L2 routing instance, supported on:
  - MX Series routers—Bridge domain and VLAN levels.
  - The PTX10000 line of routers (running Junos OS Evolved)—VLAN level.

When configuring a static remote VTEP list, keep the following in mind:

- When specifying remote VTEPs at the bridge domain level in a routing instance, you must also use the `remote-vtep-list` statement to specify the same VTEPs at the global level in the same routing instance.

  When specifying remote VTEPs at the VLAN level in the default switching instance, you must also use the `remote-vtep-list` statement to specify the same VTEPs at the global level in the default switching instance.

  To illustrate this point, note the following in sample configurations 1 and 2:

  - Routing instance `rt1` and the default switching instance include the VTEPs with the loopback addresses of 10.1.1.1 and 10.1.1.2.
  - Bridge domain `bd1` and `VLAN-1` include the VTEPs with the loopback addresses of 10.1.1.1 and 10.1.1.2.
  - Bridge domain `bd2` and `VLAN-2` include the VTEP with the loopback address of 10.1.1.1.
  - Bridge domain `bd3` and `VLAN-3` don't include a static remote VTEP list. As a result, `bd3` and `VLAN-3` inherit all VTEPs specified in routing instance `rt1` and the default switching instance (the VTEPs with the loopback addresses of 10.1.1.1 and 10.1.1.2), respectively.

> **NOTE**: If you configure a global list of VTEPs in a routing instance using the `remote-vtep-list` statement, but you don't also configure the `static-remote-vtep-list` statement for a VLAN in the routing instance, then that VLAN inherits all the globally configured remote VTEPs.
>
> Also, on platforms that support EVPN with VXLAN encapsulation, you can have EVPN-VXLAN and static VXLAN configurations on the same device, as long as the two configurations use different sets of remote peer VTEPs.

- To replicate and flood broadcast, unknown unicast, and multicast (BUM) traffic, you must specify the `ingress-node-replication` configuration statement at the `[edit vlans name vxlan]`, `[edit bridge-domains name vxlan]`, or `[edit routing-instances name bridge-domains name vxlan]` hierarchy level. In sample configuration 1, `ingress-node-replication` is configured for all bridge domains in routing instance `rt1`. In sample configuration 2, `ingress-node-replication` is configured for all VLANs in the default switching instance.

  Note that this ingress node replication configuration restricts the BUM traffic flood domain to only those VTEPs mapped to a particular bridge domain or VLAN. For example, for sample configurations 1 and 2, the BUM traffic flood domain is restricted as described for each bridge domain and VLAN, respectively:

  - `bd1` and `bd3`; `VLAN-1` and `VLAN-3`—BUM traffic is flooded to the VTEPs with the loopback addresses of 10.1.1.1 and 10.1.1.2.
  - `bd2`; `VLAN-2`—BUM traffic is flooded to the VTEP with the loopback address of 10.1.1.1.

> **NOTE**: To ensure that BUM traffic is flooded to only those VTEPs mapped to a particular VLAN, we strongly recommend that the static VTEP lists on all QFX5*XXX* switches are symmetric. We provide this guidance to avoid a situation where, for example:
>
> - Switch 1's list includes switches 2 and 3,
> - Switch 2's list includes switches 1 and 3, and
> - Switch 3's list includes only switch 2.
>
> Because of this asymmetric configuration, when switch 1 sends BUM traffic to switch 2 and switch 3 on its list, Switch 3 forwards the BUM traffic to its local ports in the VLAN.

Sample Configuration 1 From MX Router (Routing Instance of Type virtual-switch)

```
routing-instances {
    rt1 {
        vtep-source-interface lo0.0;
        remote-vtep-list [ 10.1.1.1 10.1.1.2 ];
        instance-type virtual-switch;
        bridge-domains {
            bd1 {
                vlan-id 101;
                vxlan {
                    vni 101;
                    ingress-node-replication;
                    static-remote-vtep-list [ 10.1.1.1 10.1.1.2 ];
                }
```

```
            }
            bd2 {
                vlan-id 102;
                vxlan {
                    vni 102;
                    ingress-node-replication;
                    static-remote-vtep-list [ 10.1.1.1 ];
                }
            }
            bd3 {
                vlan-id 103;
                vxlan {
                    vni 103;
                    ingress-node-replication;
                }
            }
        }
    }
}
```

Sample Configuration 2 From QFX5XXX Switch (Default Switching Instance)

```
switch-options {
    vtep-source-interface lo0.0;
    remote-vtep-list [ 10.1.1.1 10.1.1.2 ];
}
vlans {
    VLAN-1 {
        vlan-id 10001;
        interface xe-0/0/0:0.1001;
        vxlan {
            vni 10001;
            ingress-node-replication;
            static-remote-vtep-list [ 10.1.1.1 10.1.1.2 ];
        }
    }
    VLAN-2 {
        vlan-id 10002;
        interface xe-0/0/0:0.1002;
        vxlan {
            vni 10002;
            ingress-node-replication;
```

```
            static-remote-vtep-list [ 10.1.1.1 ];
        }
    }
    VLAN-3 {
        vlan-id 10003;
        interface xe-0/0/0:0.1003;
        vxlan {
            vni 10003;
            ingress-node-replication;
        }
    }
}
```

Options

static-remote-vtep-list	Specify the loopback address of one or more remote VTEPs in square brackets ([]). You do not need to include a delimiter such as a comma between names.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Static VXLAN

static (Routing Options)

IN THIS SECTION

Syntax | **28571**

- Hierarchy Level | 28572
- Description | 28573
- Options | 28573
- Required Privilege Level | 28576
- Release Information | 28576

Syntax

```
static {
    defaults {
        static-options;
    }
    rib-group group-name;
    route destination-prefix {
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            description description;
            local-address ip-address;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-receive-ttl number;
            multiplier number;
            neighbor address;
            no-adaptation;
            transmit-interval {
                threshold milliseconds;
                minimum-interval milliseconds;
            }
            version (1 | automatic);
        }
        next-hop address;
```

```
        next-hop options;
        qualified-next-hop address {
            bfd-liveness-detection {
                authentication {
                    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 | meticulous-keyed-
sha-1 | simple-password);
                    key-chain key-chain-name;
                    loose-check;
                }
                detection-time {
                    threshold milliseconds;
                }
                holddown-interval milliseconds;
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                version (1 | automatic);
            }
            metric metric;
            preference preference;
        }
        static-options;
        spring-te-lsp-next-hop
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-options rib routing-table-name],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options],
[edit routing-options rib routing-table-name]
```

Description

Configure static routes to be installed in the routing table. You can specify any number of routes within a single `static` statement, and you can specify any number of **static** options in the configuration.

Options

defaults—(Optional) Specify global static route options. These options only set default attributes inherited by all newly created static routes. These are treated as global defaults and apply to all the static routes you configure in the `static` statement.

> **NOTE**: Specifying the global static route options does not create default routes. These options only set default attributes inherited by all newly created static routes.

route—Configure individual static routes. In this part of the `static` statement, you optionally can configure static route options. These options apply to the individual destination only and override any options you configured in the **defaults** part of the `static` statement.

- ***destination-prefix/ prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

When you configure an individual static route in the **route** part of the `static` statement, specify the destination of the route (in **route** ***destination-prefix***) in one of the following ways:

- ***network/ mask-length***, where ***network*** is the network portion of the IP address and ***mask-length*** is the destination prefix length.
- **default** if this is the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.

  > **NOTE**: IPv4 packets with a destination of **0.0.0.0** (the obsoleted limited broadcast address) and IPv6 packets with a destination of **0::0** are discarded by default. To forward traffic destined to these addresses, you can add a static route to **0.0.0.0/32** for IPv4 or **0::0/128** for IPv6.

- ***nsap-prefix***—***nsap-prefix*** is the network service access point (NSAP) address for ISO.
- **next-hop** ***address***—Reach the next-hop routing device by specifying an IP address, an interface name, or an ISO network entity title (NET).

  IPv4 or IPv6 address of the next hop to the destination, specified as:

  - IPv4 or IPv6 address of the next hop

- Interface name (for point-to-point interfaces only)
- *address* or *interface-name* to specify an IP address of a multipoint interface or an interface name of a point-to-point interface.

> **NOTE**: If an interface becomes unavailable, all configured static routes on that interface are withdrawn from the routing table.

> **NOTE**: Load balancing is not supported on management and internal Ethernet (**fxo**) interfaces because this type of interface cannot handle the routing process. On **fxp** interfaces, you cannot configure multiple next hops and enable load balancing.

next-hop options—Additional information for how to manage forwarding of packets to the next hop.

- **discard**—Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP (or ICMPv6) unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.
- *iso-net*—Reach the next-hop routing device by specifying an ISO NSAP.
- **next-table** *routing-table-name*—Name of the next routing table to the destination.

  If you use the **next-table** action, the configuration must include a term qualifier that specifies a different table than the one specified in the **next-table** action. In other words, the term qualifier in the `from` statement must exclude the table in the **next-table** action. In the following example, the first term contains **rib vrf-customer2.inet.0** as a matching condition. The action specifies a next-hop in a different routing table, **vrf-customer1.inet.0**. The second term does the opposite by using **rib vrf-customer1.inet.0** in the match condition and **vrf-customer2.inet.0** In the **next-table** action.

```
term 1 {
    from {
        protocol bgp;
        rib vrf-customer2.inet.0;
        community customer;
    }
    then {
        next-hop next-table vrf-customer1.inet.0;
    }
}
term 2 {
    from {
```

```
        protocol bgp;
        rib vrf-customer1.inet.0;
        community customer;
    }
    then {
        next-hop next-table vrf-customer2.inet.0;
    }
}
```

> **NOTE**: Within a routing instance, you cannot configure a static route with the `next-table inet.0` statement if any static route in the main routing instance is already configured with the `next-table` statement to point to the **inet.0** routing table of the routing instance. For example, if you configure on the main routing instance a static route **192.168.88.88/32** with the `next-table test.inet.0` statement and the routing instance test is also configured with a static route **192.168.88.88/32** with the `next-table inet.0` statement, the commit operation fails. Instead, you must configure a routing table group both on the main instance and on the routing instance, which enables you to install the static route into both routing tables.

> **NOTE**: With the fix for PR 1383419, configuring a static route with a next-table action on QFX5xxx/EX4xxx switches results in software forwarding when the prefix length is less than 16 bits for IPv4 or 64 bits for IPv6. To avoid possible performance impact you should configure multiple static routes with prefixes equal to or longer than 16 and 64 bits for IPv6 and IPv6, respectively

- **receive**—Install a route for this next-hop destination into the routing table.

  The **receive** option forces the packet to be sent to the Routing Engine.

  The **receive** option can be useful in the following cases:

  - For receiving MPLS packets destined to a VRF instance's loopback address
  - For receiving packets on a link's subnet address, with zeros in the host portion of the address

- **reject**—Do not forward packets addressed to this destination. Instead, drop the packets, send ICMP (or ICMPv6) unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

static-options—(Optional under **route**) Additional information about static routes, which is included with the route when it is installed in the routing table.

You can specify one or more of the following in ***static-options***. Each of the options is explained separately.

- **(active | passive)**;
- **as-path <*as-path*> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator *as-number in-address*>**;
- **community [*community-ids*]**;
- **(install | no-install)**;
- **(metric | metric2 | metric3 | metric4) *value* <type *type*>**;
- **(preference | preference2 | color | color2) *preference* <type *type*>**;
- **(readvertise | no-readvertise)**;
- **(resolve | no-resolve)**;
- **(retain | no-retain)**;
- **tag *metric* type *number***;

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for BFD authentication introduced in Junos 9.6.

Support for BFD authentication introduced in Junos 9.6 for EX Series switches.

RELATED DOCUMENTATION

> *Understand Basic Static Routing*
>
> *Example: Configure IPv4 Static Routing for a Stub Network*
>
> *Example: Configure IPv6 Static Routing for a Stub Network*

static (Security NAT)

IN THIS SECTION

- Syntax | 28577
- Hierarchy Level | 28578
- Description | 28578
- Options | 28578
- Required Privilege Level | 28578
- Release Information | 28579

Syntax

```
static {
    rule-set rule-set-name {
        description text;
        from {
            interface [interface-name];
            routing-instance [routing-instance-name];
            zone [zone-name];
        }
        rule rule-name {
            description text;
            match {
                (destination-address <ip-address> | destination-address-name <address-name>);
                destination-port (port-or-low | <to high>);
                source-address [ip-address];
                source-address-name [ip-address-name];
                source-port (port-or-low <to high>);
            }
            then {
                static-nat {
                    inet {
                        routing-instance (routing-instance-name| default);
                    }
                    nptv6-prefix {
                        address-prefix;
```

```
                        routing-instance routing-instance-name;
                    nptv6-prefix-name {
                        address-prefix-name;
                            routing-instance routing-instance-name;
                    [
                    prefix {
                        address-prefix;
                            mapped-port lower-port-range to upper-port-range;
                        routing-instance (routing-instance-name| default);
                    }
                    prefix {
                        address-prefix-name;
                        mapped-port lower-port-range to upper-port-range;
                        routing-instance (routing-instance-name| default);
                    }
                    rule-session-count-alarm (clear-threshold value | raise-threshold value);
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit security nat]
```

Description

Configure static NAT.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3. The `description` option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.3X48-D25.

static-subscriber-application

IN THIS SECTION

- Syntax | 28579
- Hierarchy Level | 28579
- Description | 28579
- Additional Information | 28580
- Required Privilege Level | 28580
- Release Information | 28580

Syntax

```
static-subscriber-application;
```

Hierarchy Level

```
[edit services service-set name service-set-options]
```

Description

Apply advanced services, such as deep packet inspection, to static subscribers configured under a dynamic profile. Static subscribers are supported on GE, AE, and PS logical interfaces over LT interfaces.

To configure advanced services for static subscribers, in addition to the `static-subscriber-application` statement, you must also configure the:

- `access-profile`, `dynamic-profile`, `authentication`, and the `interface` statements at the `[edit system services static-subscribers group]` hierarchy level. You configure a demux interface for the group.
- Demux interface and the GE, AE, or PS interface at the `[edit interfaces]` hierarchy level.
- `profile` statement at both the `[edit services hcm]` and `[edit services pcef]` hierarchy levels.
- `service-set` statement at the `[edit services]` hierarchy level. Be sure to configure the `hcm-profile`, `pcef-profile`, and the `interface-service` statements. For the `interface-service` statement, configure the `ms-` or `ams-` interface name.
- `service-filter` statement at the `[edit firewall family (inet|inet6)]` hierarchy level.
- `dynamic-profiles` statement at the `[edit]` hierarchy level, making sure to include the service-set and service-filter for each family on an interface, and configure the PCEF profile name on the `pcef` statement at the `[edit dynamic-profiles` *`name`* `interfaces` *`name`* `unit` *`number`*`]` hierarchy level.

Additional Information

To verify that the services have been applied to the static subscriber, use these operational-mode commands:

- `show services pcef subscriber (terse|extensive)`
- `show services pcef pic` *`fpc-shot-number pic-slot-number`*
- `show subscriber extensive`

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

static-subscribers

IN THIS SECTION

- Syntax | 28581
- Hierarchy Level | 28581
- Description | 28581
- Options | 28582
- Required Privilege Level | 28582
- Release Information | 28582

Syntax

```
static-subscribers {
    disable;
    traceoptions {
        file filename<files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Disable static subscribers processes or configure tracing operations for static subscriber processes

Options

`disable`—Disable the static subscribers process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

static-subscribers (Dynamic Service Provisioning)

IN THIS SECTION

- Syntax | 28582
- Hierarchy Level | 28583
- Description | 28584
- Required Privilege Level | 28584
- Release Information | 28584

Syntax

```
static-subscribers {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            delimiter delimiter-character;
            domain-name domain-name;
            interface;
            logical-system-name;
```

```
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dynamic-profile profile-name {
        aggregate-clients  (merge | replace);
    }
    group group-name {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                delimiter delimiter-character;
                domain-name domain-name;
                interface;
                logical-system-name;
                routing-instance-name;
                 user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
        }
        interface  interface-name <exclude> <upto upto-interface-name>;
    }
    service-profileservice-profile-name
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name system services],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services],
[edit routing-instances routing-instances-name system services],
[edit system services]
```

Description

Configure and associate subscribers with statically configured interfaces for dynamic service provisioning.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

JSRC and Subscribers on Static Interfaces

static-tei-val

IN THIS SECTION

- Syntax | 28585
- Hierarchy Level | 28585
- Description | 28585
- Options | 28585
- Required Privilege Level | 28585
- Release Information | 28585

Syntax

```
static-tei-val value;
```

Hierarchy Level

```
[edit interfaces br-pim/0/port isdn-options]
```

Description

For J Series Services Routers only. Statically configure the Terminal Endpoint Identifier (TEI) value. The TEI value represents any ISDN-capable device attached to an ISDN network that is the terminal endpoint. TEIs are used to distinguish between several different devices using the same ISDN links.

Options

value—Value between 0 through 63.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

static-umh (MBGP MVPN)

IN THIS SECTION

- Syntax | 28586
- Hierarchy Level | 28586

- Description | 28586
- Required Privilege Level | 28587
- Release Information | 28587

Syntax

```
static-umh {
    primary address;
    backup address;
    source-tree;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
mvpn],
[edit routing-instances routing-instance-name protocols mvpn]
```

Description

In a BGP multicast VPN (MVPN) with RSVP-TE point-to-multipoint provider tunnels, statically set the upstream multicast hop (UMH), instead of using one of the dynamic methods to choose the UMH routers, such as that described in `unicast-umh-election`.

The `static-umh` statement causes all type 7 (S,G) routes to use the configured primary and backup upstream multicast hops. If these UMHs are not available, no UMH is selected. If the primary is not available, but the backup UMH is available, the backup is used as the UMH.

The `static-umh` statement only affects type 7 (S,G) routes. Type 6 routes are sent toward the rendezvous point (RP), and use the dynamic UMH selection that is configured with the `unicast-umh-election` statement, or the default method of highest IP address is used if `unicast-umh-election` is not configured.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

Example: Configuring Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

sender-based-rpf (MBGP MVPN) | **16552**

unicast-umh-election | **30353**

static-mac-ip

IN THIS SECTION

- Syntax | **28587**
- Hierarchy Level | **28588**
- Description | **28588**
- Options | **28588**
- Required Privilege Level | **28589**
- Release Information | **28589**

Syntax

```
static-mac-ip ip-address {
```

```
                mac-address [ mac-address ... ];
                router;
            }
```

Hierarchy Level

```
[edit routing-instances name vlans name switch-options interface],
[edit vlans name switch-options interface]
```

```
[edit logical-systems name routing-instances name bridge-domains name bridge-options interface],
[edit routing-instances name bridge-domains name bridge-options interface],
```

```
[edit logical-systems name routing-instances name protocols evpn interface],
[edit logical-systems name tenants name routing-instances name protocols evpn interface],
[edit routing-instances name protocols evpn interface],
```

Description

Configure a static IP address to MAC address (MAC-IP) binding. You can configure an IP address to one or more MAC address entries.

NOTE: You can configure a maximum of eight MAC addresses per static IP address.

Options

ip-address	Static IP address assigned to bind to the MAC address of the device.
mac-address [mac-address ...]	The MAC address of the switch-options interface.
router	Router flag for IPv6 addresses.

Required Privilege Level

[none specified]

Release Information

Statement introduced in Junos OS Release 23.4R1.

static-multihoming-peer

IN THIS SECTION

- Syntax | 28589
- Hierarchy Level | 28589
- Description | 28590
- Options | 28590
- Required Privilege Level | 28590
- Release Information | 28590

Syntax

```
static-multihoming-peer (static-multihoming-peer-IP-address | [ static-multihoming-peer-IP-
address ... ]);
```

Hierarchy Level

```
[edit logical-systems name protocols evpn],
[edit protocols evpn]
```

Description

Statically identify a device's multihoming peer leaf devices in an Ethernet VPN (EVPN) network running enhanced optimized intersubnet multicast (OISM).

OISM leaf devices are multihoming peers when they share an Ethernet segment (ES) for a multihomed customer edge (CE) device or attached host. When enhanced OISM leaf devices have static information about their multihoming peers, the multihoming peer leaf devices can avoid multicast traffic loss when their peer devices go up and down.

We support this statement with OISM in enhanced OISM mode only.

> **NOTE**: In Junos OS Release 24.2R1, this statement was named `multihoming-peer-gateways` to support multiple multihoming feature use cases.
>
> Starting in Junos OS and Junos OS Evolved Release 24.4R1, instead we provide this `static-multihoming-peer` statement for OISM multihoming peer leaf devices. You won't see the `multihoming-peer-gateways` statement in the Junos OS CLI anymore.

Options

(*static-multihoming-peer-IP-address* | [*static-multihoming-peer-IP-address* ...])—One or more multihoming peer device loopback IPv4 addresses. Specify a single address, or list multiple peer device addresses within square brackets.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS and Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

Statically Identify Multihoming Peers With Enhanced OISM To Improve Convergence

statistics

IN THIS SECTION

- Syntax | 28591
- Hierarchy Level | 28591
- Description | 28591
- Options | 28591
- Required Privilege Level | 28592
- Release Information | 28592

Syntax

```
statistics {
    no-ingress;
    no-transit;
}
```

Hierarchy Level

```
[edit logical-systems name protocols source-packet-routing telemetry],
[edit protocols source-packet-routing telemetry]
```

Description

Enable traffic-statistics collection on SR-TE policies through sensors for both SR-TE policy next hop and binding segment identifier (SID) installed in the forwarding table, by using JVISION traffic sensors in Junos data plane to stream out traffic statistics on static segment routing policies and their corresponding binding SID routes.

Options

- **Default:** By default statistics collection is disabled for static SR-TE routes.

no-ingress (Optional) Create sensors for binding SID transit routes only.

no-transit (Optional) Create sensors for SR-TE policy next hops and collect statistics on all steering routes that use the SR-TE policy as nexthop.

> **NOTE**: `no-ingress` and `no-transit` are mutually exclusive. You cannot disable traffic collection at both destination route and binding SID route.If you enable statistics collection and do not configure one of the two options then the sensors collect traffic statistics steered by all routes that use the SR-TE policy as a nexthop and the labeled traffic steered by the binding SID that are installed in the forwarding table.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

telemetry | **29054**

source-packet-routing (Segment Routing Traffic Engineering) | **28257**

Segment Routing Traffic Engineering at BGP Ingress Peer Overview

statistics (Access Profile)

IN THIS SECTION

- Syntax | **28593**
- Hierarchy Level | **28593**
- Description | **28593**

- Options | 28593
- Required Privilege Level | 28593
- Release Information | 28593

Syntax

```
statistics (time | volume-time);
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router or switch to collect time statistics, or both volume and time statistics, for the sessions being managed by AAA.

Options

`time`—Collect uptime statistics only.

`volume-time`—Collect both volume and uptime statistics.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`volume-time` option added in Junos OS Release 9.4.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

statistics (Protocols MPLS)

IN THIS SECTION

- Syntax | 28594
- Hierarchy Level | 28594
- Description | 28595
- Options | 28595
- Required Privilege Level | 28596
- Release Information | 28596

Syntax

```
statistics {
    auto-bandwidth (MPLS Statistics);
    file filename <files number> <size size> <world-readable | no-world-readable>;
    interval seconds;
    no-transit-statistics;
    traffic-class-statistics;
    transit-statistics-polling;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Enable MPLS statistics collection and reporting.

Options

`file` *`filename`*—(Optional) Name of the file to receive the output. We recommend that you place MPLS tracing output in the file mpls-stat in the /var/log directory.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *file* reaches its maximum size, it is renamed *file*.0, then *file*.1, and so on, until the maximum number of files is reached. Then, the oldest file is overwritten.

- **Range:** 2 or more
- **Default:** 2 files

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

`interval` *`seconds`*—Interval at which to periodically collect statistics.

- **Range:** 1 through 65,535
- **Default:** 300 seconds

`no-world-readable`—(Optional) Prevent users from reading the log file.

`size` *`size`*—(Optional) Maximum size of each file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a file named *file* reaches this size, it is renamed *file*.0. When the *file* again reaches its maximum size, *file*.0 is renamed *file*.1 and *file* is renamed *file*.0. This renaming scheme continues until the maximum number of files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of files with the `files` option.

`world-readable`—(Optional) Enable users to read the log file.

- **Syntax:** Syntax: *x*`k` to specify KB, *x*`m` to specify MB, or *x*`g` to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

`traffic-class-statistics`—(Optional) Create counters that maintain data traffic statistics per traffic class at the ingress of all types of LSPs and egress of ultimate hop popping (UHP) point-to-point LSPs. These counters are not created by default and are required to be configured to perform traffic-class-scoped loss measurement.

`transit-statistics-polling`—(Optional) Enable the polling and display of MPLS statistics for LSPs transiting the router. By default, RSVP does not periodically poll for transit LSP statistics. You cannot configure this statement and the `no-transit-statistics` statement at the same time.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`traffic-class-statistics` option introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring MPLS to Gather Statistics

Configuring Automatic Bandwidth Allocation for LSPs

statistics-service

IN THIS SECTION

- Syntax | 28597
- Hierarchy Level | 28597
- Description | 28597
- Options | 28597
- Required Privilege Level | 28597
- Release Information | 28597

Syntax

```
statistics-service {
    command binary-file-path;
    disable;
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the Packet Forwarding Engine (PFE) statistics service management process.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the Packet Forwarding Engine (PFE) statistics service management process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

statistics (Service Accounting)

IN THIS SECTION

- Syntax | 28598
- Hierarchy Level | 28598
- Description | 28598
- Options | 28598
- Required Privilege Level | 28599
- Release Information | 28599

Syntax

```
statistics (time | volume-time);
```

Hierarchy Level

```
[edit access profile profile-name service accounting]
```

Description

Configure the router to collect time statistics, or both volume and time statistics, for the service accounting sessions being managed by AAA.

Options

`time`—Collect uptime statistics only.

`volume-time`—Collect both volume and uptime statistics.

NOTE: You must not configure the `volume-time` option if the `local` option is configured at the `[edit access profile profile-name service accounting-order]` hierarchy level, because the

`local` configuration automatically determines that both volume and time statistics are collected for the service accounting sessions. If you configure the `volume-time` option in this case, an error is generated when you commit the configuration.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R1.

RELATED DOCUMENTATION

RADIUS Accounting for Subscriber Access

RADIUS Accounting for Subscriber Access

RADIUS Accounting for Subscriber Access

statistics (Services)

IN THIS SECTION

- Syntax | 28600
- Hierarchy Level | 28600
- Description | 28600
- Options | 28600
- Required Privilege Level | 28600
- Release Information | 28600

Syntax

```
statistics {
    interval interval-number;
}
```

Hierarchy Level

```
[edit services application-identification]
```

Description

Specify the interval, in minutes, for statistics collection.

Options

`interval` *interval-number*—Length of time, in minutes, that application statistics are collected.

- **Range:** 1 through 1440 minutes
- **Default:** 1 minute

> **NOTE**: For SRX Series Firewalls, the maximum number of interval periods for which statistics are stored is 8.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Onbox Application Identification Statistics

stats-aggregation-interval (programmable-rpd)

IN THIS SECTION

- Syntax | 28601
- Hierarchy Level | 28601
- Description | 28601
- Options | 28601
- Required Privilege Level | 28602
- Release Information | 28602

Syntax

```
stats-aggregation-interval seconds;
```

Hierarchy Level

```
[edit routing-options programmable-rpd]
```

Description

Configure the aggregation window for programmable routing protocol process (prpd) statistics.

If the OpenConfig telemetry collector is receiving periodic streaming every 1 minute, then we recommend that you configure a 60-second `stats-aggregation-interval` in order for the collector to not lose any data. This is because, after every aggregation period, the previous data is reset.

Options

seconds Statistics aggregation time in seconds.

- **Range:** 60 through 3600 seconds
- **Default:** 60 seconds

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

show programmable-rpd clients | 23287

stats-sampling-interval (programmable-rpd)

IN THIS SECTION

- Syntax | 28602
- Hierarchy Level | 28602
- Description | 28603
- Options | 28603
- Required Privilege Level | 28603
- Release Information | 28603

Syntax

```
stats-sampling-interval seconds;
```

Hierarchy Level

```
[edit routing-options programmable-rpd]
```

Description

Configure the sampling interval at which to collect programmable routing protocol process (prpd) statistics.

Options

seconds Interval in seconds.

- **Range:** 30 through 3600 seconds
- **Default:** 30 seconds

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

show programmable-rpd clients | 23287

status

IN THIS SECTION

- Syntax | 28604
- Description | 28604
- Options | 28604
- Required Privilege Level | 28604
- Release Information | 28604

Syntax

```
status
```

Description

Display the users currently editing the configuration.

Options

This command has no options.

Required Privilege Level

configure—Enters status configuration mode.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Displaying Users Currently Editing the Configuration

status-control

IN THIS SECTION

- Syntax | 28605
- Hierarchy Level | 28605
- Description | 28605
- Required Privilege Level | 28605
- Release Information | 28605

Syntax

```
status-control (active | standby);
```

Hierarchy Level

```
[edit interfaces aggregated-ether-options mc-ae]
```

Description

Specify whether a peer hosting a multichassis link aggregation group (MC-LAG) is primary or secondary. Primary is considered active, and secondary is considered standby.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

steering

IN THIS SECTION

- Syntax | 28606
- Hierarchy Level | 28606
- Description | 28606
- Required Privilege Level | 28606
- Release Information | 28607

Syntax

```
steering {
    keep-existing-steering;
    path {
        ipv4-address ipv4-address;
        ipv6-address ipv6-address;
    }
    routing-instance {
        downlink downlink-vrf-name;
        uplink uplink-vrf-name;
    }
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],
[edit services pcef pcc-action-profiles profile-name]
```

Description

Specify the method that a PCC action profile uses for steering traffic

If you are using Junos OS Subscriber Aware, configure steering at the `[edit unified-edge pcef pcc-action-profiles profile-name]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the PCC action profile at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-action-profiles profile-name]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Understanding Application-Aware Policy Control for Subscriber Management

Configuring Policy and Charging Control Action Profiles for Subscriber Management

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

step-percent (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 28608
- Junos OS Hierarchy Level | 28608
- Junos OS Evolved Hierarchy Level | 28608
- Description | 28608
- Options | 28608
- Required Privilege Level | 28608
- Release Information | 28608

Syntax

```
step-percent percent;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking profilestest-profile profile-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 test-profile profile-name]
```

Description

Specify the step percentage for frame-loss tests. This parameter is not applicable for the other types of tests. If you do not configure this parameter, the default step-percent is 10 percent.

Options

percent Step percent for frame-loss tests.

- **Default:** 10 percent
- **Range:** 1 through 100 percent

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

stickydr

IN THIS SECTION

- Syntax | **28609**
- Hierarchy Level | **28609**
- Description | **28609**
- Required Privilege Level | **28610**
- Release Information | **28610**

Syntax

```
stickydr
```

Hierarchy Level

```
[edit protocols pim interface interface-name]
[edit routing-instances instance-name protocols pim interface interface-name]
```

Description

The `stickydr` feature protects against traffic loss as can happen when the designated router (DR) changes once a new router joins the LAN and/or following an interface down event, or device upgrade. Set `stickydr` on all the last hop devices in the LAN, and it will assign one DR special priority (that is, `0xfffffffe`, the second highest priority) irrespective of existing DR election logic (DM priority and IP address of PIM neighbors). The sticky DR priority remains with the device until it is explicitly transferred to another eligible device on the LAN.

This feature is especially useful for countering DR elections cases wherein a new interface on the LAN appears, immediately wins the DR election, and even before it has received an IGMP join from host, starts pulling traffic from the upstream router.

Consider the example of a new device with higher DM priority and/or IP address that joins the LAN. Instead of immediately ceding DR status to the new interface, an existing device with a lower IP address and/or lower priority can remain the DR and receive IGMP joins and send PIM joins upstream. When the new device (with higher priority or IP address) appears, it detects the sticky DR and joins as a non-DR. No traffic is lost because of a DR transition.

Another example is when a DR interface goes down. If the devices in the LAN were configured for `stickydr`, a new DR election amongst the remaining PIM routers will take place as usual, and as per the RFC, but the election winner will inherit the "sticky" property of the down DR when wins. The sticky status will persist even if another device with higher priority joins the LAN. Later, when the previous DR comes back up, it's DR status is not resumed.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding Designated Routers

Configuring Basic PIM Settings

Configuring a Designated Router for PIM

stop-date

IN THIS SECTION

Syntax | **28611**

- Hierarchy Level | 28611
- Description | 28611
- Options | 28611
- Required Privilege Level | 28611
- Release Information | 28611

Syntax

```
stop-date date-time;
```

Hierarchy Level

```
[edit schedulers scheduler scheduler-name]
```

Description

Specify the time, day, month, and year that the schedule ends.

Specifying the year is optional. If no year is specified, the schedule applies to the current year and all subsequent years. If the year is specified in either the `start-date` or `stop-date` statement, that year is used for both statements.

Options

`date-time` —Use the format [*`yyyy`* -] *`mm`* - *`dd`* . *`hh`* . *`mm`* to specify the year, month, day, hour, and minutes.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

stop-time

IN THIS SECTION

- Syntax | 28612
- Hierarchy Level | 28612
- Description | 28612
- Options | 28613
- Required Privilege Level | 28613
- Release Information | 28613

Syntax

```
stop-time hh:mm;
```

Hierarchy Level

```
[edit schedulers scheduler  scheduler-name  daily],
[edit schedulers scheduler  scheduler-name  friday],
[edit schedulers scheduler  scheduler-name monday],
[edit schedulers scheduler  scheduler-name saturday],
[edit schedulers scheduler  scheduler-name  sunday],
[edit schedulers scheduler  scheduler-name   tuesday],
[edit schedulers scheduler  scheduler-name  thursday],
[edit schedulers scheduler  scheduler-name  wednesday]
```

Description

Specify the time that a schedule stops for a specified day.

If you specify a stop time for a daily schedule with the `daily` statement and also include the the `friday`, `monday`, `saturday`, `sunday`, `tuesday`, `thursday`, and `wednesday` statements in the schedule, the stop time specified for a specific day (for example, Friday using the `friday` statement) overrides the stop time set with the `daily` statement.

Options

time —Use the 24-hour format (*hh*:*mm*:*ss*) to specify the hours, minutes, and seconds.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release.

RELATED DOCUMENTATION

Security Policies Overview

storage

IN THIS SECTION

- Syntax | 28614
- Hierarchy Level | 28614
- Description | 28614
- Options | 28614
- Required Privilege Level | 28615
- Release Information | 28615

Syntax

```
storage device-name {
    type {
        cdrom {
            source {
                file filename;
            }
        }
        disk {
            bus-type [ide | virtio];
                file-type [qcow2 | raw];
                source {
                file filename;
            }
        }
        usb {
            source {
                file filename;
            }
        }
    }
}
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Configure storage parameters on VNFs.

Options

storage ***device-name***	Name of the storage device. For example, hda, hdb, sdb, or vdb.
type ***disk-type***	Type of disk. For example, cdrom, usb, or disk.

source *file-name*

Path of the source file of the storage device.

> **NOTE**: The source file to be attached to the VNF must be located in the **/var/third-party/** directory and must have read and write permissions for all.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D40.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

store

IN THIS SECTION

- Syntax | 28616
- Hierarchy Level | 28616
- Description | 28616
- Default | 28617
- Options | 28617
- Required Privilege Level | 28618
- Release Information | 28618

Syntax

```
store {
   database;
   file name {
      files number;
      (no-world-readable | world-readable);
      size size;
   }
   fwding-file name {
      files number;
      (no-world-readable | world-readable);
      size size;
   }
}
```

Hierarchy Level

```
[edit system resiliency]
```

Description

Store exception data in a file or database.

The on-box collector is disabled by default.

To enable the on-box collector for forwarding exceptions and store the exception data in a file, configure the following:

```
user@host# set system resiliency exceptions forwarding
user@host# set system resiliency store fwding-file file-name
user@host# set system resiliency store fwding-file size size
```

For forwarding exceptions, you must also configure inline-monitoring services to create the IPFIX records used to carry the exception data; see *Understand Juniper Resiliency Interface*.

To enable the on-box collector for routing and kernel exceptions and store the telemetry key-value-pair exception data in a file, configure the following:

```
user@host# set system resiliency exceptions routing
user@host# set system resiliency exceptions os
user@host# set system resiliency store file file-name
user@host# set system resiliency store size size
```

Default

database

Options

database

Store exceptions data in the on-box SQLite database, in the **/var/db** directory. You can copy this database to a remote server and issue SQL commands using the SQLite command line interface.

Starting in Release 22.2R1, you can only configure one place to store your exception data, either in the database or in a file on the device. You cannot store data in both places. If you configure both the `database` statement and the `file` or `fwding-file` statement, you won't be able to commit the configuration.

file *file-name*

Use the specified file to receive the telemetry exceptions data. All files are placed in the directory **/var/log**. The files are also archived in the same way as trace logs.

Starting in Release 22.2R1, you can only configure one place to store your exception data, either in the database or in a file on the device. You cannot store data in both places. If you configure both the `database` statement and the `file` statement, you won't be able to commit the configuration.

files *files*

(Optional) Use the specified maximum number of files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Default:** 3 files
- **Range:** 2 through 1000 files

fwding-file *file-name*

Use the specified file to receive the forwarding IPFIX exceptions data. All files are placed in the directory **/var/log**. The files are also archived in the same way as trace logs.

Starting in Release 22.2R1, you can only configure one place to store your exception data, either in the database or in a file on the device. You cannot store data in both places. If you configure both the `database` statement and the `fwding-file` statement, you won't be able to commit the configuration.

no-world-readable (Default) Disable unrestricted file access. This means the exceptions file can be accessed only by the user who configured the exceptions operation.

world-readable (Optional) Enable unrestricted file access.

size *size* (Optional) Use the specified maximum size of each file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of files with the `files` option.

- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 for MX Series routers.

Statement introduced in Junos Evolved OS Release 22.2R1 for PTX Series routers.

RELATED DOCUMENTATION

Juniper Resiliency Interface

storm-control (Video Monitoring)

IN THIS SECTION

- Syntax | 28619
- Hierarchy Level | 28619
- Description | 28619
- Options | 28619
- Required Privilege Level | 28620
- Release Information | 28620

Syntax

```
storm-control {
    count number;
    interval number;
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the count and the interval to control the flooding of SNMP traps per flow.

Options

count *number* Use the specified maximum number of SNMP traps generated in the configured interval.

interval *number* Use the specified minimum time period, in seconds, between the generation of successive traps.

- **Default:** The default count value is 1.

  The default interval is 1 second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

> *Understanding Inline Video Monitoring on MX Series Routers*
>
> alarms (Video Monitoring) | 2125

storm-control

IN THIS SECTION

- Syntax | 28620
- Hierarchy Level | 28621
- Description | 28621
- Required Privilege Level | 28621
- Release Information | 28621

Syntax

```
storm-control storm-control-profile;
```

Hierarchy Level

```
[edit interfaces interface-name unit number family ethernet-switching],
[edit interfaces interface-name unit number family bridge]
[edit interfaces interface-name ether-options ethernet-switch-profile]
[edit logical-systems name  interfaces interface-name unit number family bridge]
```

Description

Bind a storm control profile to a given interface.

On switches running ELS software, storm control is enabled by default on all switch interfaces at a level of 80 percent of the combined broadcast and unknown unicast streams. (For the equivalent statement for platforms running non-ELS software, see "storm-control" on page 28622.)

> **NOTE**: If you configure storm control on an aggregated Ethernet interface, the storm-control level is applies to each member interface individually. For example, if the aggregated interface has two members and you configure a storm-control level of 20 kbps, Junos will not detect a storm if one or both of the member interfaces receives traffic at 15 kbps because in neither of these cases does an individual member receive traffic at a rate greater than the configured storm-control level. In this example, Junos detects a storm only if at least one member interface receives traffic at greater than 20 Kbps.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

Starting in Junos OS release 17.4R1 for MX Series routers, support for storm control was added for logical systems.

RELATED DOCUMENTATION

- Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches

Understanding Storm Control

storm-control (switch)

IN THIS SECTION

- Syntax | 28622
- Hierarchy Level | 28622
- Description | 28623
- Default | 28623
- Required Privilege Level | 28623
- Release Information | 28623

Syntax

```
storm-control {
    action-shutdown;
    interface (all |  interface-name) {
        bandwidth bandwidth;
        level level;
        multicast;
        no-broadcast;
        no-multicast;
        no-registered-multicast;
        no-unknown-unicast;
        no-unregistered-multicast;
    }
}
```

Hierarchy Level

```
[edit ethernet-switching-options]
```

Description

Configure storm control on the switch.

Apply storm control to all interfaces or to the specified interfaces on switches running non-ELS software. (For the equivalent statement for switches running ELS software, see "storm-control" on page 28620.

The remaining statements are explained separately. See CLI Explorer.

Default

On switches running non-ELS software, storm control is disabled by default on all switch interfaces. If you enable storm control and do not specify a storm control level, the default level is 80 percent of the available bandwidth for ingress traffic.You can change the storm control level by configuring it as a specific bandwidth value.

When you configure storm control bandwidth on an aggregated Ethernet interface, each member of the aggregated interface is assigned that bandwidth. For example, if you configure 7000000 Kbps on aggregated interface `ae1`, and `ae1` has two members, `xe-2:0/0/0` and `xe-2:0/0/1`, each member is allowed a bandwidth level of 7000000 Kbps. Thus, the storm control bandwidth on `ae1` could be as much as 14000000 Kbps of combined broadcast and unknown unicast traffic.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

> *Understanding Storm Control*
>
> *Example: Using Storm Control to Prevent Network Outages (non-ELS)*

storm-control-profiles

IN THIS SECTION

- Syntax | 28624
- Hierarchy Level | 28624
- Description | 28624
- Required Privilege Level | 28625
- Release Information | 28625

Syntax

```
storm-control-profiles profile-name {
    action-shutdown;
    all {
        bandwidth-level;
        bandwidth-percentage;
        no-broadcast;
        no-multicast;
        no-registered-multicast;
        no-unknown-unicast;
        no-unregistered-multicast;
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
[edit logical-systems name forwarding-options]
```

Description

Configure a storm control profile on a switch or router. Storm control is used to prevent network outages that are caused by broadcast traffic storms. Storm control enables the switching device to

monitor traffic levels and to drop broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level or storm control bandwidth—is exceeded, thus preventing packets from proliferating and degrading the LAN.

NOTE: The name of the storm control profile can contain no more than 127 characters.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

Starting in Junos OS release 17.4R1 for MX Series routers, support for storm control was added for logical systems.

RELATED DOCUMENTATION

Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches

Understanding Storm Control

stp

IN THIS SECTION

- Syntax | **28626**
- Syntax (EX Series) | **28626**
- Hierarchy Level | **28627**
- Description | **28627**
- Default | **28627**
- Required Privilege Level | **28628**

- Release Information | **28628**

Syntax

```
stp   {
                          disable;
                        bridge-priority                     priority;
                         forward-delay                     seconds;
                         hello-time                       seconds;
                             interface (all | interface-name) {
                                  disable;
                                     bpdu-timeout-action {
                                          block;
                                          alarm;
        }
                                 cost                          cost;
                                  edge;
                                 mode                         mode;
                                  no-root-port;
                                 priority                         priority;
    }
                         max-age                      seconds;
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
           <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Syntax (EX Series)

```
stp   {
                          bpdu-block-on-edge                  ;
                        bridge-priority                     priority;
                         disable;
                         forward-delay                     seconds;
                         hello-time                       seconds;
```

```
                                interface (all | interface-name) {
        arp-on-stp;
                                        bpdu-timeout-action {
            block;
                                            log;
        }
                                    cost                            cost;
                                    disable;
                                    mode                            mode;
                                    no-root-port;
                                    priority                            priority;
    }
                            max-age                         seconds;
                                traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable | no-world-
readable>;
        flag flag;
    }
}
```

Hierarchy Level

```
[edit protocols],
[edit protocols rstp force-version]
```

Description

When you explicitly configure STP, a switch uses the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP (defined in the IEEE 802.1D 1998 specification).

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: Although the edge configuration statement appears in the CLI on the switch, this statement has no effect on the switch operation if you configure it.

Default

STP is disabled; by default, RSTP is enabled on all Ethernet switching ports.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Overview of Spanning-Tree Protocols
- *Understanding STP*
- *Configuring STP on EX Series Switches (CLI Procedure)*
- Configuring STP
- show spanning-tree bridge | **26548**
- show spanning-tree interface | **26555**

stream (Logical Systems and Tenant Systems)

IN THIS SECTION

- Syntax | **28629**
- Hierarchy Level | **28629**
- Description | **28629**
- Options | **28629**
- Required Privilege Level | **28630**
- Release Information | **28631**

Syntax

```
stream name {
    category name;
    filter name;
    host {
        ipaddr;
        port port;
        routing-instance routing-instance;
    }
    rate-limit rate;
}
```

Hierarchy Level

```
[edit logical-systems name security log]
[edit tenants name security log]
```

Description

Define security log stream options for a logical and tenant system. When the logging mode is set to `stream`, security logs generated in the data plane are streamed out a revenue traffic port directly to a remote server. All the categories can be configured for sending specific category logs to different log servers for stream mode log forwarding.

Options

category Type of logged events.

- Values:
  - aamw—AAMW events are logged.
  - alg—Alg events are logged.
  - all—All events are logged.
  - appqos—Appqos events are logged.

- content-security—Content security events are logged.
- flow—Flow events are logged.
- fw-auth—Fw-auth events are logged.
- gtp—Gtp events are logged.
- idp—Idp events are logged.
- ipsec—Ipsec events are logged.
- nat—Nat events are logged.
- pst-ds-lite—Pst-ds-lite events are logged.
- rtlog—Rtlog events are logged.
- screen—Screen events are logged.
- sctp—Sctp events are logged.
- secintel—Secintel events are logged.

`filter` Selects the filter to filter the logs to be logged.

- Values:
  - `threat-attack` — Threat attack security events are logged.

`host` Destination to send security logs.

- *ip-address*—Specify IP address of the host.
- `port` *port-number*—Specify host port number.
- `routing-instance` *routing-instance*—Specify the routing instance.

`rate-limit` Rate-limit for security logs.

- **Range:** 1 through 65535
- **Default:** 65535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1 for logical Systems.

Statement introduced in Junos OS Release 18.3R1.

The `routing-instance` option introduced in Junos OS Release 18.3R1 for tenant systems.

RELATED DOCUMENTATION

Troubleshooting Logical Systems

stream (Next Gen Services Service-Set Remote System Logging)

IN THIS SECTION

- Syntax | 28631
- Hierarchy Level | 28631
- Description | 28632
- Options | 28632
- Required Privilege Level | 28632
- Release Information | 28632

Syntax

```
stream stream-name (severity debug | category screen | format sd-syslog | host);
```

Hierarchy Level

```
edit services service-set name syslog
```

Description

Specify the name of the stream to the remote log server.

NOTE: Each remote server requires a unique stream name.

Options

severity debug

category screen

format sd-syslog

host

Required Privilege Level

system

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

stream (Next Gen Services Service-Set Remote System Logging) | **28631**

stream-option (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 28633
- Hierarchy Level | 28633
- Description | 28633
- Required Privilege Level | 28633
- Release Information | 28633

Syntax

```
stream-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop IPv4 packets that have the IP option of 8 (Stream ID).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

stream-protection (Multicast-Only Fast Reroute)

IN THIS SECTION

- Syntax | 28634
- Hierarchy Level | 28634
- Description | 28635
- Required Privilege Level | 28635
- Release Information | 28635

Syntax

```
stream-protection {
    mofrr-asm-starg;
    mofrr-disjoint-upstream-only;
    mofrr-no-backup-join;
    mofrr-primary-path-selection-by-routing;
    policy policy-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Description

Enable multicast-only fast reroute (MoFRR) on a routing or switching device. MoFRR minimizes packet loss in a network when there is a link failure.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Multicast-Only Fast Reroute

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain

stream (Security Log)

IN THIS SECTION

- Syntax | 28636
- Hierarchy Level | 28637
- Description | 28637
- Options | 28637
- Required Privilege Level | 28639
- Release Information | 28640

Syntax

```
stream stream-name {
    category (all | content-security | fw-auth | screen | alg | nat | flow | sctp | gtp | ipsec
| idp | rtlog |pst-ds-lite | appqos |secintel |aamw);
    file {
        allow-duplicates;
        name file-name;
        size maximum-file-size;
        rotation maximum-file-rotation-number;
    filter {
        threat-attack;
        }
    format (binary | sd-syslog | syslog | welf);
    host {
        ip-address;
        port port-number;
        routing-instanceinstance-name;
    }
    radius {
            ipaddr {
                backup backup;
                primary primary;
            }
            port port number;
            secret secret;
    }
    rate-limit {
        log-rate;
    }
    severity (alert | critical | debug | emergency | error | info | notice | warning);
    source-address {
        ip-address;
    }
    time-format (year | millisecond);
    transport {
        protocol (tcp | tls | udp);
        tcp-connections tcp-connections;
        tls-profile tls-profile;
    }
}
```

Hierarchy Level

```
[edit security log]
[edit logical-systems name security log]
[edit tenants tenant-name security log]
```

Description

Define the security log steam settings.

Options

stream Every stream can configure file or host.

- Values:
  - `category`— Type of events that may be logged.
    - `all`— All events are logged
    - `content-security` — Content security events are logged
    - `fw-auth`— Fw-auth events are logged.
    - `screen`— Screen events are logged.
    - `alg`— Alg events are logged.
    - `nat`— Nat events are logged.
    - `flow`— Flow events are logged.
    - `sctp`— Sctp events are logged.
    - `gtp`— Gtp events are logged.
    - `ipsec`— IPsec events are logged.
    - `idp`— Idp events are logged.
    - `rtlog`— Rtlog events are logged.

  - `pst-ds-lite`— Pst-ds-lite events are logged.
  - `appqos`— Appqos events are logged.
  - `secintel`— AAMW events are logged.
- `file`— Security log file options for logs in local file.
  - `allow-duplicates`— To disable log consolidation.
  - `file-name`— Name of local log file.
  - `file-size`— Specify the local log file size in megabytes.
    - **Range:** 10M through 50M.
  - `rotation`— Configure the max file number for rotation.
    - **Range:** 2 through 19.
- `filter` *`threat-attack`*— Selects the filter to filter the threat attack security events to be logged
- `format` *`(binary | sd-syslog | syslog)`*— Specify the log stream format in binary or sd-syslog or syslog formats.
- `host` *`ip-address`*— Destination to send security logs.
  - `port`— Specify the host port number.
  - `routing-instance`— Specify the routing instance name.
- `radius`— Configure RADIUS accounting for security events by specifying the RADIUS accounting server.
  - `ipaddr` — Specify the IP address of the RADIUS server for accounting messages.
    - `primary` — IP address of the Primary RADIUS server. Configure the primary IP address to enable `radius` stream.
    - `backup` — IP address of the backup RADIUS server.
- `rate-limit` *`rate`*— Specify the rate limit for security logs.
  - The range is 1 through 65535 logs per second and the default value is 65535

- `severity`— Specify the severity threshold for security logs.
  - `alert`— Specify the conditions that require immediate attention.
  - `critical`— Specify the critical conditions.
  - `debug`— Specify the information normally used in debugging.
  - `emergency`— Specify the conditions that cause security functions to stop.
  - `error`— Specify the general error conditions.
  - `info`— Specify the Information about normal security operations.
  - `notice`— Specify the non error conditions that are of interest.
  - `warning`— Specify the general warning conditions.
- `source-address`— Specify the source address to the stream log.
- `time-format` *(year | millisecond)*— Specify the year, the millisecond, or both in the timestamp.
- `transport`— Set the security log transport settings.
  - `protocol` *(tcp | tls | udp)*— Specify the security log transport protocol for the device. Values: tcp, tls, and udp.
  - `tcp-connections`— Specify the number of tcp connections per stream. Values: 1 through 5.

    NOTE: `tcp-connections` is not supported on logical and tenant systems.
  - `tls-profile`— Specify the tls profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

The [edit `logical-systems` *name* security log] and [edit `tenants` *tenant-name* security log] hierarchy levels introduced in Junos OS Release 19.1R1.

`time-format`, `source-address`, and `transport` options are introduced in Junos OS Release 20.2R1.

`radius` option is introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Example: Configuring AppTrack

category (Security Logging) | **3224**

streaming-server (Junos Telemetry Interface)

IN THIS SECTION

- Syntax | **28640**
- Hierarchy Level | **28641**
- Description | **28641**
- Options | **28641**
- Required Privilege Level | **28642**
- Release Information | **28642**

Syntax

```
streaming-server streaming-server-name {
    remote-address ip-address;
    remote-port number;
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

For Junos telemetry interface, configure the parameters of the server that collects exported data streamed by a monitored system resource. You can configure more than one streaming server. To collect data, you must associate a configured server with one or more configured sensors. The sensor configuration defines the parameters to monitor a specific system resource. To configure a sensor, include the `sensor sensor-name` statement at the `[edit services analytics]` hierarchy level.

To configure the server that collects data, you must also configure a destination IP address and a destination port. Junos telemetry interface relies on neighbor reachability information to deliver packets to the destination address. That means that all policies, such as filtering, that apply to the packets for that destination also apply to the exported packets.

NOTE: Starting with Junos OS Release 15.1F6. you can also associate more than one server with a specific sensor configuration, which enables you to transmit streamed data for the same sensor to more than one server.

NOTE: Junos telemetry interface was introduced in Junos OS Release 15.1F3 on MX Series routers with interfaces configured on MPC1 through MPC6E and on PTX Series routers with interfaces configured on FPC3. Starting in Junos OS Release 15.1F5, Junos Telemetry Interface is also supported on MPC7E, MPC8E, and MPC9E on MX Series routers.

Starting with Junos OS Release 16.1R3, FPC1 and FPC2 on PTX Series routers are also supported.

Options

streaming-server-name

Specify a name for the server configured to collect data streamed through Junos Telemetry Interface. You can configure multiple streaming servers. To associate as many as four server names with a sensor configuration, include each name at the `[edit services analytics sensor sensor-name streaming server [ streaming-server-names ] ]` hierarchy level. If you specify more than one streaming server, you must enclose the names in brackets.

remote-address *ip-address*	Specify the destination address of the streaming server for exported packets.
remote-port *number*	Specify a port number for the destination address of the streaming server for exported packets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F3.

RELATED DOCUMENTATION

export-profile (Junos Telemetry Interface) | 6669

streaming-servers

IN THIS SECTION

- Syntax | 28643
- Hierarchy Level | 28643
- Description | 28643
- Options | 28643
- Required Privilege Level | 28644
- Release Information | 28644

Syntax

```
streaming-servers {
    address ip-address {
        port number {
            stream-format format;
            stream-type type
        }
    }
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Configure remote servers to receive streaming output for the network analytics monitoring of traffic and queue statistics. The streaming function supports TCP connections only, and sends records separated by a newline character.

> **NOTE**: Before you use the remote server to receive streaming data, you must set up the TCP server software to process records that are separated by the newline character (\n).

You can configure multiple servers and multiple ports on each server to receive the streaming data. You can configure different streaming data types and formats for different ports on a server, but you can configure only one streaming type and one format for each port on a server.

Options

address ***ip-address***	IP address of the remote server receiving the streaming data.
port ***number***	Port number of the remote server receiving the streaming data.
stream-format ***format***	Format of the streaming data being sent to a server. Only one format can be sent to each port on a server.

- Values:
  - **csv**—Comma-separated Values (CSV). Records sent in this format contain a “q” for a queue statistics, or a “t” for a traffic statistics.
  - **json**—JavaScript Object Notification (JSON). Records sent in this format contain “queue-statistics” or “traffic-statistics” in the “record type” field.
  - **tsv**—Tab-separated Values (TSV). Records sent in this format contain a “q” for a queue statistics, or a “t” for a traffic statistics.

stream-type *type*

Type of streaming data sent to a port. You can specify different types of streaming data to be sent to different ports on the same server.

- Values:
  - **queue-statistics**
  - **traffic-statistics**

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

strict-asla-based-flex-algorithm

IN THIS SECTION

- Syntax | 28645
- Hierarchy Level | 28645
- Release Information | 28645
- Description | 28646
- Required Privilege Level | 28646

Syntax

```
strict-asla-based-flex-algorithm;
```

Hierarchy Level

```
[edit logical-systems name protocols isis source-packet-routing],
[edit logical-systems name routing-instances name protocols isis source-packet-routing],
[edit logical-systems name tenants name routing-instances name protocols isis source-packet-routing],
[edit protocols isis source-packet-routing],
[edit routing-instances name protocols isis source-packet-routing],
[edit tenants name routing-instances name protocols isis source-packet-routing]
```

Release Information

Statement introduced in Junos OS 21.3R1.

Description

Configure IS-IS to ignore links that do not have TE Attributes advertised in Flex-Algo specific Application-Specific Link Attribute (ASLA) sub-TLVs. This is applicable only to the following TE attribute sub-TLVs:

- Min unidirectional link delay
- TE default metric
- Administrative group

> **NOTE**: If you do not enable strict ASLA-based flexible algorithm, but have routers in your network that do not advertise ASLA sub-TLVs, the route computation of the flexible algorithm could be incorrect and might cause network issues such as routing loops.

Advertising Delay-Metric in Flex Algo ASLA

You can advertise delay-metric inside flex algo Application Specific Link Attributes (ASLA). This functionality is essential in a multi-vendor deployment running latency-based flex algo.

Required Privilege Level

routing

RELATED DOCUMENTATION

How to Configure Flexible Algorithms in IS-IS for Segment Routing Traffic Engineering

strict (DHCP Local Server)

IN THIS SECTION

Syntax | 28647
Hierarchy Level | 28647

- Description | 28647
- Default | 28647
- Required Privilege Level | 28648
- Release Information | 28648

Syntax

```
strict;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Specify whether the server denies a client to bind when the client does not indicate that it accepts reconfigure messages. This feature is available only for DHCPv6.

Default

Accept solicit messages from clients that do not support reconfiguration and permit them to bind.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview

Preventing Binding of Clients That Do Not Support Reconfigure Messages

strict-dual-isis (Protocols IS-IS)

IN THIS SECTION

- Syntax | 28648
- Hierarchy Level | 28649
- Description | 28649
- Default | 28649
- Options | 28649
- Required Privilege Level | 28649
- Release Information | 28650

Syntax

```
strict-dual-isis {
    holddown delay-value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit protocols isis]
```

Description

Configure strict dual stack (IPv4 and IPv6) support for IS-IS.

The `strict-dual-isis` option was introduced in Junos release 18.RR1-S6. This option avoids protocol flap when both IPv4 and IPv6 are configured on the same interface or in the same ISIS instance. The option works by introducing a holdown delay to allow both IPv4 ARP and IPv6 ND to complete before the interface is advertised. The default hold time is 15 seconds. Use the `holddown` argument to specify the desired delay in the range of 1 to 65535 seconds.

To configure IS-IS for strict dual stack operation include the following statements:

```
isis {
    strict-dual-isis {
       holdown delay-value;
    }
}
```

Default

Strict dual stack operation for IS-IS is not enabled.

Options

holddown Specify the delay before ISIS advertises IPv4 and IPv6 reachability, either for the entire ISIS instance or under a specific interface. For details on interface level application, see "interface (Protocols IS-IS)" on page 8966

By default the `holdown` delay is 15 seconds. Users can configure a value in the range of 1 to 65535 seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 18.4R1-S6.

RELATED DOCUMENTATION

interface (Protocols IS-IS) | 8966

strict-priority-scheduler

IN THIS SECTION

- Syntax | 28650
- Hierarchy Level | 28650
- Description | 28651
- Required Privilege Level | 28651
- Release Information | 28651

Syntax

```
strict-priority-scheduler;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name]
```

Description

On PTX Series routers only, enable *strict-priority scheduling* for all scheduler maps within the traffic control profile.

Physical interfaces on PTX Series routers support two mutually exclusive scheduling processes:

- *normal scheduling* (default mode) — A queue's transmit-rate is used to determine whether it is operating within the *guaranteed region* or in the *excess region*.

  Within the *guaranteed region* (transmit-rates credits are positive), the scheduler uses the transmit rates to decide the bandwidth allocation. Queues that are at priority-level `low` or higher and have transmit-rate credits are serviced first by priority order and then within a priority level using packet round robin algorithm.

- *strict-priority scheduling* — Queues are processed in strict-priority order. There is no concept of guaranteed region and excess region. The packet scheduler is always operating in the guaranteed region, with the exception of priority `low`, which is always assigned to the excess priority level. The configured `transmit-rate` does not affect how the queue is serviced because packets are processed in order of queue priority. Among queues that are configured `low` priority, if `excess-rate` weights are configured, they are used by the hardware to perform WRR. Queues that are mapped to the same hardware priority or that have the same configured priority other than `low` are serviced in a packet round-robin fashion.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Example: Configuring Strict-Priority Scheduling on a PTX Series Router

Understanding Scheduling on PTX Series Routers

Understanding CoS CLI Configuration Statements on PTX Series Routers

How Schedulers Define Output Queue Properties

traffic-control-profiles | **30009**

strict-source-route-option (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 28652
- Hierarchy Level | 28652
- Description | 28652
- Required Privilege Level | 28652
- Release Information | 28652

Syntax

```
strict-source-route-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop IPv4 packets that have the IP option of 9 (Strict Source Routing).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

strict-syn-check

IN THIS SECTION

- Syntax | 28653
- Hierarchy Level | 28653
- Description | 28653
- Required Privilege Level | 28653
- Release Information | 28654

Syntax

```
strict-syn-check;
```

Hierarchy Level

```
[edit security flow tcp-session]
```

Description

Enable the strict three-way handshake check for the TCP session. It enhances security by dropping data packets before the three-way handshake is done. By default, strict-syn-check is disabled.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

strict-targeted-hellos

IN THIS SECTION

- Syntax | 28654
- Hierarchy Level | 28654
- Description | 28655
- Required Privilege Level | 28655
- Release Information | 28655

Syntax

```
strict-targeted-hellos;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling Strict Targeted Hello Messages for LDP

string

IN THIS SECTION

- Syntax | 28655
- Hierarchy Level | 28656
- Description | 28656
- Required Privilege Level | 28656
- Release Information | 28656

Syntax

```
string {
    equals;
```

```
    has-prefix;
    has-suffix;
    matches;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format]
```

Description

Specify the custom AVP attribute's format as a string and the value to match for the incoming RADIUS request from the IP-based subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

strip-domain (Domain Map)

IN THIS SECTION

- Syntax | 28657
- Hierarchy Level | 28657
- Description | 28657
- Required Privilege Level | 28657
- Release Information | 28657

Syntax

```
strip-domain;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Remove the domain name from the username before continuing with any AAA services specified in a domain map.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Mapping Subscriber Domains to Access and Session Options

Mapping Subscriber Domains to Access and Session Options

strip-nexthop

IN THIS SECTION

- Syntax | 28658
- Hierarchy Level | 28658
- Description | 28659
- Required Privilege Level | 28659
- Release Information | 28659

Syntax

```
strip-nexthop;
```

Hierarchy Level

```
[edit protocols bgp family (inet | inet-vpn | inet6 | inet6-vpn) flow],
[edit protocols bgp group group-name family (inet | inet-vpn | inet6 | inet6-vpn) flow],
[edit protocols bgp group group-name neighbor address family (inet | inet-vpn | inet6 | inet6-
vpn flow)],
[edit routing-instances routing-instance-name protocols bgp family (inet | inet-vpn | inet6 |
inet6-vpn) flow],
[edit routing-instances routing-instance-name protocols bgp group group-name family (inet | inet-
vpn | inet6 | inet6-vpn) flow],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (inet | inet-vpn | inet6 | inet6-vpn) flow]
```

Description

Prevents BGP from advertising flow route updates with real nexthop address even when the route is present in the local routing table. When `strip-nexthop` is not configured, Junos advertises the real next hop to its neighbors in order to interoperate with devices that have the capability to accept and advertise real BGP nexthops. You can either propagate the received next hop to an EBGP peer in the case of a route reflector or define a policy to advertise self next hop.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Example: Enabling BGP to Carry Flow-Specification Routes

strip-user-name (Access Profile)

IN THIS SECTION

- Syntax | 28660
- Hierarchy Level | 28660
- Description | 28660
- Required Privilege Level | 28660
- Release Information | 28660

Syntax

```
strip-user-name {
    delimiter delimiter;
    parse-direction (left-to-right | right-to-left);
}
```

Hierarchy Level

```
[edit access profile profile-name session-options]
```

Description

Configure the details of username stripping for a subscriber access profile. This configuration determines how a portion of a subscriber login string is identified and discarded, leaving the remainder for subsequent use as a modified username by an external AAA server for session authentication and accounting. For example, the modified username might appear in RADIUS Access-Request, Acct-Start, and Acct-Stop messages, as well as RADIUS-initiated disconnect requests and change of authorization (CoA) requests.

You can use the `aaa-options` *`aaa-options-name`* statement at the `[edit access]` hierarchy level to define options that specify the LS:RI context for AAA authorization and configuration for a specific subscriber or a set of subscribers, overriding any such configuration within a domain map.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access

Configuring Username Modification for Subscriber Sessions

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

strip-username (Domain Map)

IN THIS SECTION

- Syntax | 28661
- Hierarchy Level | 28661
- Description | 28661
- Options | 28662
- Required Privilege Level | 28662
- Release Information | 28662

Syntax

```
strip-username (left-to-right | right-to-left)
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Strip the user portion of the username for authentication to simplify off-chassis provisioning. Parsing direction can be left-to-right (default) or right-to-left.

Options

left-to-right Strip the user portion of the username from left to right.

right-to-left Strip the user portion of the username from right to left.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Changing the Username and Password to Simplify Off-Chassis Provisioning

strip-as-path

IN THIS SECTION

- Syntax | 28662
- Hierarchy Level | 28663
- Description | 28663
- Required Privilege Level | 28663
- Release Information | 28663

Syntax

```
strip-as-path;
```

Hierarchy Level

```
[edit policy-options policy-statement policy-name term term-name then]
```

Description

Removes the incoming autonomous system (AS) path as part of the import policy for a BGP session and replaces the received autonomous system (AS) path with the receiving router's local AS number for the receiving session. Note that the local AS number may be different from the number configured under `autonomous system` in the `[edit routing-options]` hierarchy.

If you need to normalize externally injected routes, you can use this policy option for the incoming autonomous system (AS) path so that it may be used similarly to routes that originate solely within the fabric. The new `strip-as-path` policy option has no impact on the BGP export policy.

You can configure the `strip-as-path` option under `policy-options then` clause:

```
set policy-options policy-statement do-strip term a then strip-as-path
```

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3.

RELATED DOCUMENTATION

Autonomous Systems for BGP Sessions

structured-data

IN THIS SECTION

- Syntax | 28664
- Hierarchy Level | 28664
- Description | 28664
- Required Privilege Level | 28665
- Release Information | 28665

Syntax

```
structured-data {
    brief;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name system syslog file filename],
[edit system syslog file filename]
```

Description

Write system log messages to the log file in structured-data format, which complies with Internet draft draft-ietf-syslog-protocol-23, *The syslog Protocol* (http://tools.ietf.org/html/draft-ietf-syslog-protocol-23).

Starting in Junos OS 19.2R1 release and later, after configuring this statement, you cannot configure any other statements that specify the format for messages written to the file. Commit error appears (the **explicit-priority** statement at the **[edit system syslog file *filename*]** hierarchy level and the **time-format** statement at the **[edit system syslog]** hierarchy level) if you try to configure any other statement after configuring this statement. In releases prior to Junos OS 19.2R1 release, there is no commit error when you run the `request system software validate` command before an upgrade to validate the software package as a pre-requisiste.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Log Messages in Structured-Data Format

explicit-priority | 6657

time-format | 29356

stub

IN THIS SECTION

- Syntax | 28665
- Hierarchy Level | 28666
- Description | 28666
- Options | 28666
- Required Privilege Level | 28666
- Release Information | 28667

Syntax

```
stub <default-metric metric> <(no-summaries | summaries)>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3) area area-id],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Specify that this area not be flooded with AS external link-state advertisements (LSAs). You must include the `stub` statement when configuring all routing devices that are in the stub area.

The backbone cannot be configured as a stub area.

You cannot configure an area to be both a stub area and a not-so-stubby area (NSSA).

Options

default-metric *metric*	Metric value, ranging from 1 through 16,777,215.
no-summaries	(Optional) Do not advertise routes into the stub area. If you include the `default-metric` option, only the default route is advertised.
summaries	(Optional) Flood summary LSAs into the stub area.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Understanding OSPF Areas

Example: Configuring OSPF Stub and Totally Stubby Areas

nssa | **11717**

stub-network

IN THIS SECTION

- Syntax | **28667**
- Hierarchy Level | **28667**
- Description | **28668**
- Required Privilege Level | **28668**
- Release Information | **28668**

Syntax

```
stub-network;
```

Hierarchy Level

```
[edit logical-systems name protocols ospf overload],
[edit logical-systems name routing-instances name protocols ospf overload],
```

```
[edit protocols ospf overload],
[edit routing-instances name protocols ospf overload]
```

Description

Advertise Stub Network with maximum metric.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2.

RELATED DOCUMENTATION

Understanding OSPF Overload Function

allow-route-leaking | 2202

intra-area-prefix | 9224

as-external | 2486

sub-domain

IN THIS SECTION

- Syntax | 28669
- Hierarchy Level | 28669
- Description | 28669
- Options | 28672
- Required Privilege Level | 28673
- Release Information | 28673

Syntax

```
sub-domain name {
 aaa-logical-system name {
      aaa-routing-instance (default | name);
 }
 target-logical-system name {
     target-routing-instance (default | name);
 }
 access-profile access-profile;
 address-pool address-pool;
 dynamic-profile dynamic-profile;
 override-chap-password override-chap-password;
 override-password override-password;
 qualifier {
  vlan-id-list [ vlan-id-list ... ];
 }
 strip-domain;
 strip-username (left-to-right | right-to-left);
 tunnel-profile tunnel-profile;
 using-user-password;
}
```

Hierarchy Level

```
[edit access domain map]
```

Description

You can configure subdomains under a domain map. Subdomain allows you to select different access-profile for users within the same domain different VLAN ID or VLAN ID range. Subdomain gives the flexibility to differentiate the users in a domain. It provides different services per the profile configuration.

The options available at the domain map level are also available at the subdomain. But the options defined for a subdomain are independent from other subdomains and the options at the domain map level.

Characteristics of a Subdomain

- Subdomain configuration within a domain takes the higher precedence than the domain map level configuration.
- The qualifier option is mandatory to define a subdomain.
- Qualifiers (VLAN ID) cannot have any overlap within a domain.
- You can configure maximum of 16 subdomains within a domain.

Here're some examples of different types of subdomain configurations and their behavior.

1. Access profile configuration in a domain with VLAN ID qualifier using subdomain.

```
[edit access domain map abc.com]
user@host# set access-profile A
user@host# set sub-domain sub1 qualifier vlan-id-list 1
user@host# set sub-domain sub1 access-profile B
user@host# set sub-domain sub2 qualifier vlan-id-list [ 11 20 29-35 100-199 ]
user@host# set sub-domain sub2 access-profile C
user@host# set sub-domain sub3 qualifier vlan-id-list 300-399
user@host# set sub-domain sub3 access-profile D
```

   This configuration creates different access profiles using the VLAN ID qualifier and subdomain.

2. Access profile and strip domain configuration through domain map.

```
[edit access domain map abc.com]
user@host# set access-profile A
user@host# set strip-domain
user@host# set sub-domain sub1 qualifier vlan-id-list 1
user@host# set sub-domain sub1 access-profile B
user@host# set sub-domain sub1 strip-domain
user@host# set sub-domain sub2 qualifier vlan-id-list [ 11 20 29-35 100-199 ]
user@host# set sub-domain sub2 access-profile C
user@host# set sub-domain sub2 strip-domain
user@host# set sub-domain sub3 qualifier vlan-id-list 300-399
user@host# set sub-domain sub3 access-profile D
user@host# set sub-domain sub3 strip-domain
```

   The intention of this configuration is to use the domain map for an access profile selection, then strip the domain name. For each subdomain you can assign different access profile depending on the VLAN ID, but the `strip-domain` kept common for all cases. Though it appears a repeat configuration in

subdomain, since the subdomain is totally independent, it gives better flexibility when it comes to assign modifiers selectively.

3. Domain map and subdomain configuration for completely independent attribute selection.

```
[edit access domain map abc.com]
user@host# set access-profile A
user@host# set strip-domain
user@host# set sub-domain sub1 qualifier vlan-id-list 1
user@host# set sub-domain sub1 access-profile B
user@host# set sub-domain sub1 strip-domain
user@host# set sub-domain sub2 qualifier vlan-id-list [ 11 20 29-35 100-199 ]
user@host# set sub-domain sub2 access-profile C
user@host# set sub-domain sub2 address-pool
user@host# set sub-domain sub3 qualifier vlan-id-list 900-399
user@host# set sub-domain sub3 access-profile D
user@host# set sub-domain sub3 dynamic-profile vlan-profile-9xx
```

For the same domain `abc.com`, each subdomain and unqualified domain map (top level) is independently defining its actions. VLAN ID is the qualifier for subdomain, which takes precedence and overrides the unqualified attributes with independent set available in the qualified subdomain.

4. Configure an empty subdomain.

```
[edit access domain map abc.com]
user@host# set access-profile A
user@host# set strip-domain
user@host# set sub-domain sub1 qualifier vlan-id-list [ 100 200-299 400-450 ]
```

This configuration creates an empty subdomain with a set of VLAN ranges. This configuration is an example to exclude the users of the same domain depending on their VLAN ID.

Any user login qualifying for the subdomain match do not apply any options. All the other non-matching users in the domain get the options from un-qualified top level domain map.

5. Invalid subdomain configuration (overlapping VLAN ID ranges in subdomains).

```
[edit access domain map abc.com]
user@host# set access-profile A
user@host# set strip-domain
user@host# set sub-domain sub1 qualifier vlan-id-list [ 100 200-299 400-450 ]
user@host# set sub-domain sub1 access-profile B
```

```
user@host# set sub-domain sub1 strip-domain
user@host# set sub-domain sub2 qualifier vlan-id-list [ 250-300 ]
user@host# set sub-domain sub2 access-profile C
user@host# set sub-domain sub2 strip-domain
```

This configuration gets rejected during commit. The qualifiers within the same domain map cannot have any overlap.

Example of an error message while trying such invalid subdomain configuration commit:

```
root@host# commit
2021-02-03 22:50:39.730422 IST: Running FIPS Self-tests
Veriexec is not enforced, FIPS mode not available
2021-02-03 22:50:39.768595 IST: FIPS Self-tests Skipped
[edit access domain map abc.com sub-domain sub2 qualifier]
  'vlan-id-list 250-300'
    Range 250-300 overlaps with range 200-299 in another sub-domain under same domain
error: configuration check-out failed
```

Options

sub-domain *name*
Name of a subdomain.

aaa-logical-system
Logical system used for applying AAA services.

aaa-routing-instance
Routing instance used for applying AAA services.

- default—Default routing instance.
- name—Name of the routing instance you want to configure.

target-routing-instance
Specify the routing instance of the subscriber context.

- default—Default routing instance.
- name—Name of the routing instance you want to configure.

access-profile *profile-name*
Name of an access profile.

address-pool
Specify the address pool used to assign addresses to subscribers associated with the domain map.

dynamic-profile	Specify the dynamic profile that is used for subscriber sessions associated with the subdomain.
override-chap-password	Use this CHAP password for authentication.
override-password	Use this password for authentication.
strip-domain	Enable domain name stripping from the username.
strip-username	Enable user name stripping from the username. • left-to-right—Strip to first domain delimiter on the left. • right-to-left—Strip to first domain delimiter on the right.
tunnel-profile	Specify the tunnel profile that provides definitions for tunnels associated with the subdomain.
using-user-password	Send overridden CHAP-Password using User-Password.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

Domain Mapping Overview

sub-prefix-length

IN THIS SECTION

- Syntax | 28674
- Hierarchy Level | 28674
- Description | 28674
- Required Privilege Level | 28674
- Release Information | 28675

Syntax

```
sub-prefix-length sub-prefix-length;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcpv6-client prefix-
delegating]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet6 dhcpv6-client prefix-delegating]
[edit tenants tenant-name interfaces interface-name unit logical-unit-number family inet6 dhcpv6-
client prefix-delegating]
```

Description

Allows you to configure DHCPv6 client sub prefix length. The DHCPv6 client separates the delegated prefix according to sub-prefix lengths. If the delegated prefix is not enough for all interfaces, the client sends out a syslog message.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30 and in Junos OS Release 15.1X49-D100.

The `logical-systems` and `tenants` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

preferred-prefix-length | 13059

subnet-address (Easy EVPN LAG Configuration)

IN THIS SECTION

- Syntax | 28675
- Hierarchy Level | 28676
- Description | 28676
- Options | 28677
- Required Privilege Level | 28680
- Release Information | 28680

Syntax

```
subnet-address {
    inet (ipv4-subnet-addr | [ipv4-subnet-addr1 ipv4-subnet-addr2 ...]);
    inet6 (ipv6-subnet-addr | [ipv6-subnet-addr1 ipv6-subnet-addr2 ...]);
    no-irb-address-auto-derive;
    virtual-gateway-v4-address ipv4-virtual-gateway-address;
    virtual-gateway-v6-address ipv6-virtual-gateway-address;
}
```

Hierarchy Level

```
[edit services evpn evpn-vxlan irb irb-instance]
```

Description

Specify IRB instance subnet address parameters as part of a generated Ethernet VPN (EVPN) fabric configuration.

You use statements at the `[edit services evpn]` hierarchy level to simplify configuring an EVPN fabric with Ethernet segment identifier (ESI) link aggregation groups (LAGs) in the topology. The feature you enable with statements in this hierarchy is an easy EVPN LAG configuration feature, also called EZ-LAG.

You only need to configure a few statements to provide the parameters required for an EVPN core fabric with multihomed or single-homed customer edge (CE) devices (or end server devices). When you commit the `[edit services evpn]` configuration statements, the device uses a commit script to generate the corresponding standard EVPN fabric configuration statements.

The commit script uses the parameters you specify at the `[edit services evpn evpn-vxlan irb irb-instance subnet-address]` hierarchy level to generate IPv4 and/or IPv6 address configurations for an IRB instance as follows:

- IRB interface subnet address.

  NOTE: For IPv6 data traffic, you can configure a global IPv6 address and an IPv6 link-local address for the IRB interface.

- Virtual gateway address

  NOTE: By default, the commit script sets the virtual gateway MAC address for the device (which applies to all IRB instances) as follows:

  - IPv4 virtual gateway MAC address—00:00:5e:00:01:01
  - IPv6 virtual gateway MAC address—00:00:5e:00:02:01

  You can override the default virtual gateway MAC addresses using the following options at the `[edit services evpn global-parameters]` hierarchy level:

- `virtual-gateway (v4-mac | v6-mac)` *`virtual-gateway-mac-address`*—To set different virtual gateway MAC addresses for IPv4 traffic or IPv6 traffic
- `virtual-gateway-mac` *`virtual-gateway-mac-address`*—To set the same virtual gateway MAC address for both IPv4 traffic and IPv6 traffic.

Starting in Junos OS and Junos OS Evolved Release 23.4R1, with the `subnet-address inet` and `subnet-address inet6` options, you can specify multiple IRB subnet addresses using a list syntax [*addr1addr2* ...]. You can still use the original syntax with a single subnet address (no list brackets), and use multiple statements with single subnet addresses to specify multiple IRB addresses.

Also as of those releases, in the generated configuration, the commit script includes `router-advertisement` statements at the `[edit protocols]` hierarchy level as follows to configure the device with default values to send router advertisement packets for a configured EZ-LAG IRB instance:

- With an anycast gateway EZ-LAG configuration:

```
set protocols router-advertisement interface irb.unit max-advertisement-interval 30
set protocols router-advertisement interface irb.unit min-advertisement-interval 10
```

- With a virtual gateway address EZ-LAG configuration:

```
set protocols router-advertisement interface irb.unit max-advertisement-interval 30
set protocols router-advertisement interface irb.unit min-advertisement-interval 10
set protocols router-advertisement interface irb.unit virtual-router-only
```

Options

inet ***address(es)***

Specify an IPv4 subnet address or multiple IPv4 subnet addresses for the IRB interface—you can specify a single address value *`ipv4-subnet-addr`* or an address list `[` *`ipv4-subnet-addr1 ipv4-subnet-addr2 …`* `]`.

You must configure this option or the `inet6` option (or both options).

inet6 ***address(es)***

Specify an IPv6 subnet address or multiple IPv6 subnet addresses for the IRB interface—you can specify a single address value *`ipv6-subnet-addr`* or an address list `[` *`ipv6-subnet-addr1 ipv6-subnet-addr2 …`* `]`.

You must configure this option or the `inet` option (or both options).

For IPv6 data traffic on an IRB interface, you can use either multiple statements with single addresses or one statement with an address list to specify the IPv6 global subnet address and the IPv6 link-local address (prefix fe80::) of the IRB interface. For example, if you include the following statements for an IRB interface IPv6 subnet address and an IPv6 link-local address for logical unit 10 in your EZ-LAG configuration:

```
set services evpn evpn-vxlan irb irb_10 subnet-address inet6
2001:db8::10:a:1:1/112
set services evpn evpn-vxlan irb irb_10 subnet-address inet6 fe80::10:a:1:1/112
```

or:

```
set services evpn evpn-vxlan irb irb_10 subnet-address inet6
[ 2001:db8::10:a:1:1/112 fe80::10:a:1:1/112 ]
```

then the commit script adds the following statements in the generated configuration:

```
set interfaces irb unit 10 family inet6 address 2001:db8::10:a:1:1/112
set interfaces irb unit 10 family inet6 address fe80::10:a:1:1/112
```

The commit script also includes default `router-advertisement` statements in the generated configuration for that IRB interface. For example, if you have an EZ-LAG virtual gateway configuration, the generated configuration also includes the `virtual-router-only` option, which enables sending router advertisement packets only for the link-local virtual gateway address:

```
set protocols router-advertisement interface irb.10 max-advertisement-interval 30
set protocols router-advertisement interface irb.10 min-advertisement-interval 10
set protocols router-advertisement interface irb.10 virtual-router-only
```

See *Routing IPv6 Data Traffic through an EVPN-VXLAN Network with an IPv4 Underlay* and *IPv6 Neighbor Discovery* for more on EVPN-VXLAN IPv6 data traffic and setting IPv6 link-local addresses.

no-irb-address-auto-derive

Instruct the commit script to not derive IRB interface subnet addresses in the generated configuration for this IRB instance.

For configuration simplicity, with this feature you can set the same IRB interface subnet address on connected PE devices using this `subnet-address` option at the `[edit services`

evpn evpn-vxlan irb *irb-instance*] hierarchy level. By default, the commit script derives different device IRB subnet addresses from the subnet-address value you configure for peer ID 1, based on the peer ID of the device on which it's running.

When you set no-irb-address-auto-derive, instead of the default behavior, the commit script uses exactly the IPv4 or IPv6 subnet addresses you specify with the subnet-address options on each device.

> **NOTE**: This option applies individually to the specified IRB instance. You can alternatively set this option globally at the [edit services evpn global-parameters] hierarchy level. In that case, the option takes effect for all IRB instances on the device that you configure with the irb *irb-instance* statement at the [edit services evpn evpn-vxlan] hierarchy level.

If you don't set no-irb-address-auto-derive (either per IRB instance or globally), by default the commit script derives the subnet address based on the provided *subnet-address* value on each device. In this case, the commit process doesn't do a commit check to enforce setting the same subnet address across both devices.

> **NOTE**: We strongly recommend that you make sure you set subnet-address to the same base value on both peer PE devices in the easy EVPN LAG configuration. That way the default address derivation works as expected for the IRB interface subnet addresses on each peer PE device.

virtual-gateway-v4-address *ipv4-virtual-gateway-address*

Specify an IPv4 virtual gateway address for the IRB interface.

If you don't include this option to set a specific IPv4 virtual gateway address, the device derives a virtual gateway IPv4 address for the IRB interface from the address you specify with the inet option and the highest configurable IPv4 address in the subnet range.

For example, if you configure subnet address 10.1.1.1/24, the commit script generates IRB interface configuration statements with an IPv4 virtual gateway address 10.1.1.254.

virtual-gateway-v6-address *ipv6-virtual-gateway-address*

Specify an IPv6 virtual gateway address for the IRB interface.

If you don't include this option to set a specific IPv6 virtual gateway address, the device derives a virtual gateway IPv6 address for the IRB interface from the address you specify with the inet6 option and the highest configurable IPv6 address in the subnet range.

For example, if you configure subnet address `2001:db8::10:1:1:1/112` , the commit script generates IRB interface configuration statements with an IPv6 virtual gateway address `2001:db8::10:1:1:fffe`.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 and Junos OS Evolved Release 23.2R1.

RELATED DOCUMENTATION

Easy EVPN LAG (EZ-LAG) Configuration

subscriber (Access Profile)

IN THIS SECTION

- Syntax | 28680
- Hierarchy Level | 28681
- Description | 28681
- Options | 28682
- Required Privilege Level | 28682
- Release Information | 28682

Syntax

```
subscriber username {
    delegated-pool delegated-pool-name;
```

```
    framed-ip-address ipv4-address;
    framed-ipv6-pool ipv6-pool-name;
    framed-pool ipv4-pool-name;
    password password;
    target-logical-system logical-system-name<(target-routing-instance (default | routing-
instance-name)>;
    target-routing-instance (default | routing-instance-name);
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Enable local authentication for subscribers by configuring a password to match the subscriber. Local authentication can take the form of either user password authentication or Challenge Handshake Authentication Protocol' (CHAP) authentication. For user password authentication, the configured password is used to verify the subscriber's login password. For CHAP authentication, the configured password acts as the challenge secret to verify the subscriber's challenge password and challenge response credential.

> **NOTE**: Local authentication and authorization also requires the `password` option to be configured as an `authentication-order` method for the access profile.

You can also optionally configure several attributes, such as an address, address pool, logical system, or routing instance, to be authorized locally for the subscriber when authentication is successful.

Local authentication supports all subscriber types that are currently supported by subscriber management and services on MX Series routers.

Local authentication is useful when you do not want to use external authentication servers. The associated local authorization similarly is useful when you do not want to use external authorization servers. Another use case might be when you are migrating a network from E Series routers running JunosE software to MX Series routers running Junos OS. You may also want to configure local authentication and authorization as a backup for RADIUS authentication.

If you do not configure an address or address pool for local authorization, address assignment is based on network matching or the first address pool assigned to the routing instance.

> **NOTE**: Local authentication and authorization supports a chassis-wide maximum of 100 subscribers. If subscribers are configured in access profiles where `authentication-order password` is not configured, local authentication does not occur, but these subscriber count against the system limit of 100 subscribers for local authentication.

Options

delegated-pool *delegated-pool-name*	(Optional) Specify the name of an address pool used to locally allocate a delegated IPv6 prefix for the subscriber. Corresponds to RADIUS standard attribute Delegated-IPv6-Prefix (123).
framed-ip-address *ipv4-address*	(Optional) Specify the IP address to be configured for the subscriber. Corresponds to RADIUS standard attribute Framed-IP-Address (8).
framed-ipv6-pool *ipv6-pool-name*	(Optional) Specify the name of an address pool used to assign a router advertisement IPv6 prefix or a DHCPv6 IA_NA/128 address for the subscriber. Corresponds to RADIUS standard attribute Framed-IPv6-Pool (100).
framed-pool *ipv4-pool-name*	(Optional) Specify the name of an address pool used to assign an IPv4 address for the subscriber. Corresponds to RADIUS standard attribute Framed-Pool (88).
password *password*	Specify the password used to authenticate the subscriber locally. Corresponds to RADIUS standard attributes User-Password (2) or CHAP-Password (3).
target-logical-system *logical-system-name*	(Optional) Specify the name of the logical system assigned to the subscriber.
target-routing-instance (default | *routing-instance-name***)**	(Optional) Specify the name of the routing instance assigned to the subscriber; either the default routing instance or a nondefault routing instance.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Example: Configuring CHAP Authentication with RADIUS

RADIUS Authentication and Accounting Basic Configuration

Configuring Local Authentication and Authorization for Subscribers

subscriber-address

IN THIS SECTION

- Syntax | 28683
- Hierarchy Level | 28683
- Description | 28684
- Required Privilege Level | 28684
- Release Information | 28684

Syntax

```
subscriber-address {
    inet {
        pool pool-name;
    }
    inet6 {
        pool pool-name;
    }
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
```

Description

Specify the address pool that contains the source IP addresses for IP-based subscriber packets that can undergo TDF processing.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

subscriber-aware-services

IN THIS SECTION

- Syntax | 28685
- Hierarchy Level | 28685
- Description | 28685
- Required Privilege Level | 28685
- Release Information | 28685

Syntax

```
subscriber-aware-services;
```

Hierarchy Level

```
[edit chassis fpc name pic name]
```

Description

Enable subscriber-aware services.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 20.2R1 for Next Gen Services on MX240, MX480 and MX960.

subscriber-awareness (Service Set Options)

IN THIS SECTION

- Syntax | 28686
- Hierarchy Level | 28686
- Description | 28686
- Default | 28686
- Required Privilege Level | 28686
- Release Information | 28686

Syntax

```
subscriber-awareness;
```

Hierarchy Level

```
[edit services service-set service-set-options]
```

Description

Enable subscriber awareness on the service set.

Default

If you do not include the `subscriber-awareness` statement, then subscriber-aware services cannot be provided.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> *Applying Services to Subscriber-Aware Traffic with a Service Set*

subscriber-context (AAA Options)

IN THIS SECTION

- Syntax | 28687
- Hierarchy Level | 28687
- Description | 28687
- Options | 28687
- Required Privilege Level | 28688
- Release Information | 28688

Syntax

```
subscriber-context subscriber-context-name;
```

Hierarchy Level

```
[edit access aaa-options aaa-options-name]
```

Description

Specify the logical-system:routing-instance (LS:RI) in which the subscriber interface is placed. For example, this may correspond to the LAC-facing interface on the LNS that is accessed by all requests from a subscriber residence.

NOTE: Only the default (primary) logical system is supported.

Options

subscriber-context-name	Name of the logical-system:routing-instance in which the subscriber interface is placed.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access

Configuring Username Modification for Subscriber Sessions

L2TP LNS Inline Service Interfaces

L2TP LNS Inline Service Interfaces

subscriber (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 28688
- Hierarchy Level | 28689
- Description | 28689
- Options | 28689
- Required Privilege Level | 28690
- Release Information | 28690

Syntax

```
subscriber (flow-bandwidth | flow-control-mode | flow-detection-mode)
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type flow-level-bandwidth],
[edit system ddos-protection protocols protocol-group packet-type flow-level-control],
[edit system ddos-protection protocols protocol-group packet-type flow-level-detection]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure flow bandwidth, flow control mode, or flow detection mode at the subscriber flow aggregation level for the packet type.

Options

flow-bandwidth—Specify the bandwidth for the flow at the subscriber level. Available only at the `[edit system ddos-protection protocols protocol-group packet-type flow-level-bandwidth]` hierarchy level.

- **Default:** 100 packets per second
- **Range:** 1 through 10,000 packets per second

flow-control-mode—Specify how traffic in the detected flow is controlled at the subscriber level. Available only at the `[edit system ddos-protection protocols protocol-group packet-type flow-level-control]` hierarchy level.

> **NOTE**: The configuration at this level overrides the global configuration using the `flow-level-control` statement at the `[edit system ddos-protection global]` hierarchy level.

- `drop`—Drop all traffic in flow.
- `keep`—Keep all traffic in flow.
- `police`—Police the traffic to within its allowed bandwidth.
- **Default:** `drop`

flow-detection-mode—Specify how flow detection operates at the subscriber level when a policer has been violated. Available only at the `[edit system ddos-protection protocols protocol-group packet-type flow-level-detection]` hierarchy level.

> **NOTE**: The configuration at this level overrides the global configuration using the `flow-detection-mode` statement at the `[edit system ddos-protection global]` hierarchy level.

- `automatic`—Search flows at the subscriber level only when a DDoS policer is being violated and only until it is established that the flow causing the violation is not at this level. When the suspicious flow is not at this level, then the search moves to a coarser level of flow aggregation (logical interface). Flows at the subscriber level are subsequently not searched again until the policer is no longer violated at the coarser level.
- `off`—Disable flow detection at the subscriber level so that flows are never searched at this level.
- `on`—Search flows at the subscriber level, even when no DDoS protection policer is currently being violated. Monitoring continues at this level regardless of whether a suspect flow is identified at this level.
- **Default:** `automatic`

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring the Maximum Flow Bandwidth at Each Flow Aggregation Level

Configuring How Traffic in a Culprit Flow Is Controlled at Each Flow Aggregation Level

Configuring How Flow Detection Operates at Each Flow Aggregation Level

Setting Up and Using Flow Detection

subscriber-exclude-prefix

IN THIS SECTION

- Syntax | 28691
- Hierarchy Level | 28691
- Description | 28692
- Required Privilege Level | 28692
- Release Information | 28692

Syntax

```
subscriber-exclude-prefix {
    apply-groups [group-names];
    apply-groups-except [group-names];
    family {
        inet {
            apply-groups [group-names];
            apply-groups-except [group-names];
            network address net-mask;
        }
        inet6 {
            apply-groups [group-names];
            apply-groups-except [group-names];
            network address net-mask;
        }
    }
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
```

Description

Specify the network prefix of source IP addresses for uplink packets and destination IP addresses for downlink packets that do not undergo TDF processing.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain*
- *Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain*
- *Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain*
- *IP-Based and IFL-Based TDF Subscribers Overview*
- *IP-Based Subscriber Setup Overview*

subscriber-extension

IN THIS SECTION

- Syntax | 28693
- Hierarchy Level | 28693
- Description | 28693
- Required Privilege Level | 28693
- Release Information | 28693

Syntax

```
subscriber-extension;
```

Hierarchy Level

```
[edit security nat source]
```

Description

Enable subscriber NAT extension to treat B4 (softwire initiator) as subscriber in the Dual-Stack Lite (DS-Lite). This configuration is necessary to generate RADIUS logs for NAT PBA usage on per subscriber basis.

Enabling the `subscriber-extension` feature and performing a reboot is mandatory in order to enable RADIUS accounting logging.

> **NOTE**: For each subscriber, the initial message received to the RADIUS accounting server is considered as the start of the session and the last message received is considered as the stop of the session.
>
> Any messages received to the accounting server periodically with information about the specific session are considered as interim updates for that subscriber.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

subscriber-leave-timer

IN THIS SECTION

- Syntax | 28694
- Hierarchy Level | 28694
- Description | 28694
- Options | 28694
- Required Privilege Level | 28695
- Release Information | 28695

Syntax

```
subscriber-leave-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast interface interface-name],
[edit logical-systems logical-system-name routing-options multicast interface interface-name],
[edit routing-instances routing-instance-name routing-options multicast interface interface-name],
[edit routing-options multicast interface interface-name]
```

Description

Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message.

Options

seconds—Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message. Specifying a value of 0 results in an immediate update. This is the same as if the statement were not configured.

- **Range:** 0 through 30
- **Default:** 0 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

subscriber-management-helper

IN THIS SECTION

- Syntax | 28695
- Hierarchy Level | 28696
- Description | 28696
- Options | 28696
- Required Privilege Level | 28696
- Release Information | 28696

Syntax

```
subscriber-management-helper {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the subscriber management helper process.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the subscriber management helper process.
- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

subscriber-management

IN THIS SECTION

- Syntax | 28697
- Hierarchy Level | 28698
- Description | 28698
- Required Privilege Level | 28699
- Release Information | 28699

Syntax

```
subscriber-management {
    enable;
    enforce-strict-scale-limit-license;
    gres-route-flush-delay;
    }
    overrides {
        event {
            catastrophic-failure {
                reboot (master | standby);
            }
        }
        force-show-arp-resolve;
        interfaces {
            family (inet | inet6) {
                layer2-liveness-detection;
                ipoe-dynamic-arp-enable;
                receive-gratuitous-arp;
            }
        }
        no-unsolicited-ra;
        ra-initial-interval-max seconds;
        ra-initial-interval-min seconds;
        shmlog {
            disable;
            file filename <files maximum-no-files> <size maximum-file-size——>;
```

```
            filtering enable;
            log-name {
                all;
                logname {
                    <brief | detail | extensive | none | terse>;
                    <file-logging |no-file-logging>;
                }
                }
            log-type (debug | info | notice);
        |
    }
    redundancy {
        interface name {
            local-inet-address v4-address;
            local-inet6-address v6-address;
            shared-key string;
            virtual-inet-address virtual-v4-address;
            virtual-inet6-address virtual-v6-address;
        }
        no-advertise-routes-on-backup;
        protocol {
            pseudo-wire;
            vrrp;
        }
    }
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure global services for subscriber management, such as maintaining subscribers, tracing operations, and enabling enhanced subscriber management.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- *Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events*
- *Tracing Subscriber Management Database Events for Troubleshooting*
- *Junos OS Enhanced Subscriber Management*
- *Configuring Junos OS Enhanced Subscriber Management*
- DHCP Liveness Detection Using ARP and Neighbor Discovery Packets
- *Minimize Traffic Loss Due to Stale Route Removal After a Graceful Routing Engine Switchover*
- *How to Configure M:N Subscriber Redundancy with VRRP and DHCP Binding Synchronization*

subscriber-packet-idle-timeout

IN THIS SECTION

- Syntax | 28700
- Hierarchy Level | 28700
- Description | 28700
- Options | 28700
- Required Privilege Level | 28700

Syntax

```
subscriber-packet-idle-timeout subscriber-packet-idle-timeout;
```

Hierarchy Level

```
[edit system services packet-triggered-subscribers]
```

Description

The subscriber packet idle timeout for packet triggered subscribers.

Options

subscriber-packet-idle-timeout—Maximum idle time.

- **Range:** 15 through 1440 minutes.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Packet-Triggered Subscribers Services Overview

subscriber-profile

IN THIS SECTION

Syntax | 28701

Hierarchy Level | 28701

- Description | 28701
- Options | 28701
- Required Privilege Level | 28702
- Release Information | 28702

Syntax

```
subscriber-profile profile-name {
    enable service-name {
        concurrent-data-sessions max-session-number;
    }
    disable service-name;
    max-data-sessions-per-subscriber {
        limit max-sub-sessions;
        exceed-action {
            drop;
            syslog;
        }
    }
}
```

Hierarchy Level

```
[edit services service-set services-set-name]
```

Description

Specify the subscriber profile name. A subscriber profile specifies which services should be enabled and which services should be disabled for traffic belonging to a subscriber bound to a particular subscriber profile. A subscriber is bound to a minimum of one subscriber profile at any given time.

Options

profile-name—Name of the profile.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

subscriber-type (TDF Domain)

IN THIS SECTION

- Syntax | 28702
- Hierarchy Level | 28702
- Description | 28703
- Options | 28703
- Required Privilege Level | 28703
- Release Information | 28703

Syntax

```
subscriber-type (ip | ifl);
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
```

Description

Configure the type of subscriber that this domain is applied to — an IP-based subscriber or an IFL-based (interface-based) subscriber. If you do not include this statement, `subscriber-type ip` is used.

Options

ip (Default) Apply the TDF domain to IP-based subscribers, for which a RADIUS accounting request is sent to the MX Series router. An individual subscriber session is created for each unique source IP address.

ifl Apply the TDF domain to IFL-based subscribers, which are defined by a set of interfaces. One subscriber session is created for all traffic that is received on those interfaces.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IFL-Based Subscriber Properties with a TDF Domain

subscriber-info

IN THIS SECTION

- Syntax | 28704
- Hierarchy Level | 28704
- Description | 28704
- Options | 28705
- Required Privilege Level | 28705
- Release Information | 28705

Syntax

```
subscriber-info {
    service-statistics;
    subscriber-statistics;
}
```

Hierarchy Level

```
[edit dynamic-profiles name telemetry]
```

Description

The `subscriber-info` statement enables the streaming of subscriber metadata to an external telemetry server when the subscriber logs into or logs out of the device under test (DUT). Subscriber metadata includes subscriber-related information such as username, profile name, underlying interface name, customer VLAN (C-VLAN), and service VLAN (S-VLAN) tags. The subscriber-info is a mandatory config under the telemetry stanza.

Options

service-statistics	Enables service accounting statistics telemetry. Use this option to allow periodic streaming of firewall filter statistics for IPv4 and IPv6 filters to an external telemetry server.
subscriber-statistics	Enables periodic streaming of subscriber accounting statistics such as input packets/octets and output packets/octets to an external telemetry server.

Required Privilege Level

edit - To set the statement in the configuration

Release Information

Statement introduced in Junos OS Release 24.3R1.

RELATED DOCUMENTATION

No Link Title

subscriber-pool-utilization-alarm

IN THIS SECTION

- Syntax | 28706
- Hierarchy Level | 28706
- Description | 28706
- Options | 28706
- Required Privilege Level | 28707
- Release Information | 28707

Syntax

```
subscriber-pool-utilization-alarm
<clear-threshold clear-threshold alarm value>
<dampening-interval dampening-interval value>
<log-type <snmp-trap><syslog>>
<raise-threshold raise-threshold alarm value>;
```

Hierarchy Level

```
[edit security nat source],
[edit services nat source],
```

Description

You can configure threshold limits to receive notifications when port or port block usage exceeds the configured thresholds.

Options

clear-threshold

Clear threshold for subscriber pool utilization alarm.

- **Range:** 40 through 100

dampening-interval

Configure dampening interval for subscriber pool alarm.

- **Range:** 5 through 60

log-type

Define log type. By default, `snmp-trap` and `syslog` are enabled.

- snmp-trap—Enable SNMP trap for subscriber pool utilization.
- syslog—Enable system log for subscriber pool utilization.

raise-threshold

Raise threshold for subscriber pool utilization alarm.

- **Range:** 50 through 100

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration

Release Information

Statement introduced in Junos OS Release 24.4R1.

subscribers-limit (Resource Monitor)

IN THIS SECTION

- Syntax | 28707
- Hierarchy Level | 28708
- Description | 28708
- Required Privilege Level | 28708
- Release Information | 28708

Syntax

```
subscribers-limit {
    client-type (any | dhcp | l2tp | pppoe) {
        chassis {
            limit limit;
        }
        fpc slot-number {
            limit limit;
            pic number {
                limit limit;
                port number {
                    limit limit;
                }
            }
```

```
        }
    }
}
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Configure the maximum number of subscribers of a specified client type allowed to be logged in on the chassis, per MPC, per MIC, and per port. When that number is reached, subsequent logins are denied until the current number of subscribers drops below the maximum allowed.

Limit the number of subscribers allowed to log in per chassis, MPC, MIC, or port.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Limiting Subscribers by Client Type and Hardware Element with Resource Monitor

Resource Monitoring for Subscriber Management and Services

subscription

IN THIS SECTION

- Syntax | 28709
- Hierarchy Level | 28709
- Description | 28709
- Options | 28710
- Required Privilege Level | 28710
- Release Information | 28710

Syntax

```
subscription percentage {
    ct0 percentage;
    ct1 percentage;
    ct2 percentage;
    ct3 percentage;
    piority priority-value percentage;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-
protection],
[edit protocols rsvp interface interface-name],
[edit protocols rsvp interface interface-name link-protection]
```

Description

Configure the amount of bandwidth subscribed to a class type (when you have enabled Differentiated Services) or bypass LSP (when you have enabled link protection). `subscription` is the percentage of the link bandwidth that can be used for the RSVP reservation process.

Options

`ct`*`number percentage`*—Percentage of the class-type bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type. You can specify bandwidth subscriptions for class types 0 through 3. This option is not available for bypass LSPs.

- **Range:** 0 through 65,000
- **Default:** 100 percent

`percentage`—Percentage of the class-type or bypass LSP bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type or bypass LSP.

- **Range:** 0 through 65,000
- **Default:** 100 percent

`priority` *`priority-value percentage`*—Priority for which subscription percent is being configured.

- **Range:** 0 through 7
- **Default:** 100 percent

`percentage`—Subscription percent for the specific priority.

- **Range:** 0 through 65,000
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`priority` option introduced in Junos OS Release 20.4R1 for MX960, MX2008, MX2010, MX2020, and PTX Series.

RELATED DOCUMENTATION

> *Configuring the Bandwidth Subscription Percentage for LSPs*

Configuring the Amount of Bandwidth Subscribed for Bypass LSPs

subscription-id

IN THIS SECTION

- Syntax | 28711
- Hierarchy Level | 28712
- Description | 28712
- Required Privilege Level | 28712
- Release Information | 28712

Syntax

```
subscription-id {
    constant value;
    subscription-id-options {
        entry-name {
            id-components {
                use-class;
                use-imsi;
                use-msisdn;
                use-nai;
                use-nas-port;
                use-nas-port-id;
                use-realm;
                use-username;
            }
        }
    }
    use-class {
        regex "value";
        pattern "pattern";
        subscription-id-type (imsi | msisdn | nai | private | sip-uri);
```

```
    }
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
```

Description

Specify how the Subscription-Id is constructed for the Diameter credit control request (CCR) message that is sent from the TDF to the PCRF for IP-based subscribers belonging to the TDF domain.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

subscription-id-data-include (PCRF Partition)

IN THIS SECTION

- Syntax | 28713
- Hierarchy Level | 28713
- Description | 28713
- Options | 28714
- Required Privilege Level | 28715
- Release Information | 28715

Syntax

```
subscription-id-data-include {
    base-interface-name;
    delimiter delimiter-character;
    domain-name name;
    interface-name;
    mac-address;
    nas-port-id;
    origin-host;
    origin-realm;
    user-name;
    user-prefix prefix;
    vlan-tags;
}
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the subscriber data to use in a Policy and Charging Rules Function (PCRF).

Options

base-interface-name
Use the specified physical or underlying interface name. If both the underlying interface and the client application perform authentication, authorization, and provisioning, the identification attributes in the server requests enable the AAA or PCRF server to make an association between the two entities.

delimiter *delimiter-character*
Use the specified character used as the delimiter between the concatenated components of the subscription-id-data-include. You cannot use the semicolon (;) as a delimiter.

- **Default:** @

domain-name *name*
Use the specified domain name that is concatenated with the subscription-id-data-include during the subscriber identification process.

interface-name
Use the specified interface name that is concatenated with the subscription-id-data-include during the subscriber identification process; for example, demux0.

mac-address
Use the specified client hardware address (chaddr) from the incoming packet that is concatenated with the subscription-id-data-include during the subscriber identification process.

nas-port-id
Use the specified NAS-Port-ID (RADIUS attribute 87), which identifies the physical interface that subscriber management uses to identify subscribers that is concatenated with the subscription-id-data-include during the subscriber identification process. By default, the NAS-Port-ID includes the interface-description value that describes the physical interface.

origin-host
Use the specified name of the host that originates the Diameter message that is concatenated with the subscription-id-data-include during the subscriber identification process. Supplied as the value of Origin-Host AVP for all messages sent by the Diameter master instance.

origin-realm
Use the specified realm of the host that originates the Diameter message that is concatenated with the subscription-id-data-include during the subscriber identification process. Supplied as the value of Origin-Realm AVP for all messages sent by the Diameter master instance.

user-name
Use the specified subscriber username that is concatenated with the subscription-id-data-include during the subscriber identification process.

user-prefix *prefix*
Use the specified user prefix that is concatenated with the subscription-id-data-include during the subscriber identification process.

vlan-tags Include the subscriber session VLAN tags in the subscription ID. Both single-tagged and double-tagged VLANs are supported: The tags are added in the format *outer-vlan-tag-inner-vlan-tag*. The outer tag is always included; the inner tag is included for double-tagged VLANs.

Use this option instead of the `interface-name` option when you the outer VLAN tag is unique across the system and you do not need the underlying physical interface name to be part of the format.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

`vlan-tags` option added in Junos OS Release 18.3R1 on MX Series routers.

RELATED DOCUMENTATION

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

subscription-id-options

IN THIS SECTION

- Syntax | **28716**
- Hierarchy Level | **28716**

- Description | 28716
- Options | 28717
- Required Privilege Level | 28717
- Release Information | 28717

Syntax

```
subscription-id-options {
    [entry-name] {
        id-components {
            use-class;
            use-imsi;
            use-msisdn;
            use-nai;
            use-nas-port;
            use-nas-port-id;
            use-realm;
            use-username;
        }
    }
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber subscription-id]
```

Description

Specify a method for constructing the Subscription-Id for the RADIUS credit control request (CCR) message that is sent from the TDF to the PCRF for IP-based subscribers belonging to the TDF domain. To specify multiple methods, include the *entry-name* option multiple times.

Options

entry-name Identifier for the Subscription-Id construction method.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

IP-Based Subscriber Setup Overview

subscription-id-type (Class Attribute)

IN THIS SECTION

- Syntax | 28718
- Hierarchy Level | 28718
- Description | 28718
- Options | 28718
- Required Privilege Level | 28718
- Release Information | 28718

Syntax

```
subscription-id-type (imsi | msisdn | nai | private | sip-uri);
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name subscription-id use-class]
```

Description

Configure the subscription ID type when the Class attribute is used for the subscription ID.

Options

imsi Use the IMSI subscriber type.

msisdn Use the MSISDN (E164) subscriber type.

nai Use the NAI subscriber type.

private Use the Private subscriber type.

sip-uri Use the SIP URI name subscriber type.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

subscription-id-type (PCRF Partition)

IN THIS SECTION

- Syntax | 28719
- Hierarchy Level | 28719
- Description | 28719
- Options | 28720
- Required Privilege Level | 28720
- Release Information | 28720

Syntax

```
subscription-id-type number;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Identify the subscriber. You can identify the subscriber using several IANA predefined values, or define your own. The Policy and Charging Rule Function (PCRF) uses the subscription-id-type during a subscriber login session using for CCR-GX-I and CCA-GX-I messages.

BEST PRACTICE: For wireline service, we recommend that you define your own customer-specific value.

If you do not define or include any data in the `subscription-id-data-include` statement, then the value defaults to reserved. For some applications, the PCRF may be configured to identify subscribers using another method, and instead use the subscription-id-type for protocol compliance. If the dynamic-subscription-id is provided by the AAA server, then the PCRF uses the dynamic-subscription-id as the identifier regardless of what you defined for the local configuration.

Options

number Identifier for the type of subscriber for a PCRF partition.

- **Default:** 4 (END_USER_PRIVATE)
- **Range:** 1 through 2147483647
- **Values:** 1 (END_USER_IMSI)
  2 (END_USER_SIP_URI)
  3 (END_USER_NAI)
  4 (END_USER_PRIVATE)

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

success

IN THIS SECTION

Syntax | 28721

- Hierarchy Level | 28721
- Description | 28721
- Options | 28721
- Required Privilege Level | 28721
- Release Information | 28722

Syntax

```
success string;
```

Hierarchy Level

```
[edit access firewall-authentication pass-through default-profile          name        (ftp |
http | telnet) banner],
[edit access firewall-authentication web-authentication]
```

Description

Specify the banner (message) that users see when trying to connect using FTP, HTTP, or Telnet after successful authentication.

Options

string—Banner text. Maximum length of the message text is 250 characters. Enclose the banner text within spaces or special characters—for example, quotation marks (” “).

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Understanding Pass-Through Authentication

sunrpc

IN THIS SECTION

- Syntax | 28722
- Hierarchy Level | 28723
- Description | 28723
- Options | 28723
- Required Privilege Level | 28723
- Release Information | 28723

Syntax

```
sunrpc {
    disable;
    map-entry-timeout map-entry-timeout;
    traceoptions {
        flag {
            all {
                extensive
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Specify the Sun Microsystems remote procedure call (RPC) ALG on the device. Sun Microsystems Remote Procedure Call (Sun RPC)—also known as Open Network Computing Remote Procedure Call (ONC RPC)—provides a way for a program running on one host to call procedures in a program running on another host.

Options

- `disable`—Disable the Sun RPC ALG. By default, the Sun RPC ALG is enabled.
- `map-entry-timeout` *`map-entry-timeout`*—Specify the Sun RPC ALG mapping entry timeout value in hours.

  Range: 1 through 72 hours.

  Default: 32 hours.
- `traceoptions`—Configure the Sun RPC ALG tracing options.
  - `flag`—Trace operation to perform.
    - `all`—Trace all events.
    - *`extensive`*—Display extensive amount of data.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding Sun RPC ALGs

supplementary-blackout-timer (Host Fast Reroute)

IN THIS SECTION

- Syntax | 28724
- Hierarchy Level | 28724
- Description | 28724
- Default | 28725
- Options | 28725
- Required Privilege Level | 28726
- Release Information | 28726

Syntax

```
supplementary-blackout-timer minutes;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
interface interface-name],
[edit routing-instances instance-name routing-options interface interface-name]
```

Description

Override the global supplementary blackout timer for the specified host fast-reroute (HFRR) profile.

When you configure HFRR, an optional ARP prefix limit sets a maximum for the number of ARP routes and, therefore FRR routes created for each HFRR profile in the routing table. This limit prevents ARP attacks from exhausting the virtual memory on the routing devices.

There are two configuration statements (`global-arp-prefix-limit` and `arp-prefix-limit`) that set the ARP prefix limit, one at the global `[edit routing-options host-fast-reroute]` hierarchy level and the other at the `[edit routing-instances` *`instance-name`* `routing-options interface` *`interface-name`*`]` hierarchy level, respectively. The global `global-arp-prefix-limit` statement sets a default ARP prefix limit for all HFRR profiles configured on the routing device. The `arp-prefix-limit` statement overrides the `global-arp-prefix-limit` for that HFRR profile for that protected interface.

Warning system log messages begin when the ARP routes in an HFRR profile reaches 80% of the configured limit. When the number crosses the 100% threshold, the HFRR profile is deactivated. When this happens, all ARP/FRR routes are deleted from the routing table. FRR routes are deleted from forwarding table as well.

After the HFRR profile is deactivated, a blackout timer is started. The timeout value of this timer is the ARP cache timeout (kernel timeout) + the supplementary blackout timer.

There are global and per-HFRR CLI statements (`global-supplementary-blackout-timer` and `supplementary-blackout-timer`) to configure the supplementary blackout timer. The global value is at the `[edit routing-options host-fast-reroute]` hierarchy level and applies to all HFRR profiles on the routing device. The value for the routing-instance interface is at the `[edit routing-instances` *`instance-name`* `routing-options interface` *`interface-name`*`]` hierarchy level, and overrides the global value for that HFRR profile only.

When the blackout timer expires, the HFRR profile is reactivated, and the Junos OS relearns the ARP routes and re-creates the HFRR routes. If the ARP prefix limit is not exceeded again, the HFRR routes will be up.

If an HFRR profile is in the deactivated state, a reevaluation of the ARP state is preformed during every commit operation or whenever the routing process (rpd) is restarted with the `restart routing` command.

Default

If you omit the-`supplementary-blackout-timer` statement, the `global-supplementary-blackout-timer` takes effect for all HFRR profiles on the device. If you omit both of these statements, the blackout timeout value is set to the ARP cache timeout value, which by default is 20 minutes and is configurable with the `aging-timer` statement at the `[edit system arp]` hierarchy level.

Options

minutes Number of minutes, in addition to the ARP cache timeout value, before an HFRR profile is reactivated after the ARP prefix limit is exceeded.

- **Range:** 1 through 15 minutes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Configuring Link Protection with Host Fast Reroute

supplicant

IN THIS SECTION

- Syntax | 28726
- Hierarchy Level | 28727
- Description | 28727
- Options | 28727
- Required Privilege Level | 28727
- Release Information | 28728

Syntax

```
supplicant  {
    interface interface-name {
        local-certificate certificate-id;
        authentication-method eap-tls;
```

```
        user-id user-id;
    }
}
```

Hierarchy Level

```
[edit protocols dot1x]
```

Description

Configure the interface as a supplicant for 802.1X authentication. You can configure an interface to act as both an authenticator and as a supplicant on a link connecting switches or routers. This can be a switch-to-switch, switch-to-router, or router-to-router link. This enables the devices to authenticate each other, which is required to secure the link using MACsec in dynamic connectivity association key (CAK) mode.

MACsec in dynamic CAK mode relies on certificate-based validation using Extensible Authentication Protocol – Transport Layer Security (EAP-TLS). You must configure the supplicant interface to use EAP-TLS and assign a digital certificate to the interface.

Supplicant interfaces do not support captive portal or MAC RADIUS authentication.

Options

local-certificate *certificate-id*	Specify the the local certificate for the supplicant interface when the local device has multiple loaded certificates.
authentication-method (eap-tls | eap-md5)	Configure the authentication method for the supplicant. To support MACsec in dynamic CAK mode, you must configure EAP-TLS authentication.
user-id *user-id*	Configure the user ID.
password *password*	Configure the password.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

suppress-host-routes-from-dci-to-dc

IN THIS SECTION

- Syntax | 28728
- Hierarchy Level | 28728
- Description | 28728
- Required Privilege Level | 28729
- Release Information | 28729

Syntax

```
suppress-host-routes-from-dci-to-dc
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols evpn interconnect]
```

Description

You can suppress Ethernet VPN (EVPN) Type 5 host route advertisements re-originated across the data center interconnect (DCI), to the local data center (DC). Setting this statement achieves better scaling and performance on leaf devices by preventing host route advertisements for non-existing subnets on the local DC. Host routes corresponding to existing subnets stretched from the DCI gateway (GW) to the local DC will be advertised.

The DCI GW device treats the following Type 5 routes as host routes:

- IPv4 routes with a /32 prefix

- IPv6 routes with a /128 prefix

The DCI GW will continue to advertise non-host Type 5 routes to the DC by default.

Consider the following points before deploying this configuration:

- You configure this option on the DCI GW devices.
- You must configure integrated routing and bridging (IRB) interfaceds on DCI GW devices in edge-routed bridging (ERB) mode.
- Setting this option applies to every DCI-learned host route, and could suppress necessary route advertisements. You can selectively suppress Type 5 route advertisements with an `export` policy at the `[edit routing-instances` *`routing-instance-name`* `protocols evpn ip-prefix-routes]` hierarchy, instead of using this option. Setting this option doesn't affect locally learned host route advertisements.
- Setting this option doesn't affect IRB's configured with the `disable` statement. IRB's associated both with a Layer 3 context and EVPN bridge domain are still validated when a subnet is stretched from DCI GW to DC.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

Statement introduced in Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

interconnect | 8797

suppress-interface-leaf

IN THIS SECTION

- Syntax | 28730
- Hierarchy Level | 28730
- Description | 28730
- Required Privilege Level | 28730
- Release Information | 28731

Syntax

```
suppress-interface-leaf;
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure a single source of data for the following operational state sensors:

- **/interfaces/interface/state/high-speed**
- **/interfaces/interface/state/oper-status**

You can use this statement to suppress telemetry streaming of these statistics from the packet forwarding engine (PFE). This might be required for collectors that require a single source of data for each sensor.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.2R1.

suppression

IN THIS SECTION

- Syntax | 28731
- Hierarchy Level | 28731
- Description | 28732
- Options | 28732
- Required Privilege Level | 28733
- Release Information | 28733

Syntax

```
suppression {
    disable;
    (include-destination-address | no-include-destination-address);
    max-logs-operate value;
    max-time-report value;
    start-log value;
}
```

Hierarchy Level

```
[edit security idp sensor-configuration log]
```

Description

Log suppression reduces the number of logs by displaying a single record for multiple occurrences of the same event. Log suppression can negatively impact sensor performance if the reporting interval is set too high. By default this feature is enabled.

Options

disable —Disable log suppression.

include-destination-address When log suppression is enabled, multiple occurrences of events with the same source, service, and matching attack object generate a single log record with a count of occurrences. If you enable this option, log suppression will only combine log records for events with a matching source as well. The IDP Sensor does not consider destination when determining matching events for log suppression. By default this setting is disabled.

max-logs-operate When log suppression is enabled, IDP must cache log records so that it can identify when multiple occurrences of the same event occur. This setting specifies how many log records are tracked simultaneously by IDP.

- **Syntax:** *value*—Maximum number of log records are tracked by IDP.
- **Range:** 256 through 65,536 records
- **Default:** 16,384 records

max-time-report When log suppression is enabled, IDP maintains a count of multiple occurrences of the same event. After the specified number of seconds has passed, IDP writes a single log entry containing the count of occurrences.

- **Syntax:** *value*—Time after which IDP writes a single log entry containing the count of occurrences.
- **Range:** 1 through 60 seconds
- **Default:** 5 seconds

start-log Specify how many instances of a specific event must occur before log suppression begins.

- **Syntax:** *value*—Log suppression begins after how many occurrences.
- **Range:** 1 through 128

- **Default:** 1

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

surf-control-integrated

IN THIS SECTION

- Syntax | 28733
- Hierarchy Level | 28734
- Description | 28734
- Options | 28734
- Required Privilege Level | 28734
- Release Information | 28734

Syntax

```
surf-control-integrated {
    cache {
        size value;
        timeout value;
    }
    profile profile-name {
            category customurl-list name {
                action  (block | log-and-permit | permit);
            }
```

```
        custom-block-message value;
        default (block | log-and-permit | permit);
        fallback-settings {
            default  (block | log-and-permit);
            server-connectivity   (block | log-and-permit);
            timeout    (block | log-and-permit);
            too-many-requests  (block | log-and-permit);
            }
            timeout value;
        }
    }
    server {
        host host-name;
        port number;
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering]
```

Description

Configure the Content Security web-filtering integrated feature.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The surf-control- integrated feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The surf-control- integrated feature is not supported from Junos OS Release 15.1X49-D10 onwards.

suspend-for

IN THIS SECTION

- Syntax | 28735
- Hierarchy Level | 28735
- Description | 28736
- Default | 28736
- Required Privilege Level | 28736
- Release Information | 28736

Syntax

```
suspend-for;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name mka]
```

Description

Configure the `suspend-for` statement to maintain non-stop MACsec service during graceful routing engine switchover (GRES).

The MACsec Key Agreement (MKA) protocol maintains the MACsec session between two nodes on a point-to-point MACsec link. The MKA protocol works at the control plane level between the two nodes. When you configure the `suspend-for` statement on the local node, in the event of an RE switchover, it sends a request to the remote peer node to suspend the MACsec session at the control plane. The suspension lasts for 120 seconds.

At the data plane, traffic continues to traverse the point-to-point link during suspension. The secure association key (SAK) that was programmed prior to suspension remains in use until the switchover is complete. After the switchover, the key server generates a new SAK to secure the link. The key server will continue to periodically create and share a SAK over the link for as long as MACsec is enabled.

Default

You must configure the `suspend-for` statement to enable this feature. It is not enabled by default. If this statement is not configured, in the event of an RE switchover, the MACsec session is terminated, resulting in traffic loss during the switchover.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

suspend-on-request

IN THIS SECTION

- Syntax | 28737
- Hierarchy Level | 28737
- Description | 28737
- Default | 28738
- Required Privilege Level | 28738
- Release Information | 28738

Syntax

```
suspend-on-request;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name mka]
```

Description

Configure the `suspend-on-request` statement to maintain non-stop MACsec service during graceful routing engine switchover (GRES).

The MACsec Key Agreement (MKA) protocol maintains the MACsec session between two nodes on a point-to-point MACsec link. The MKA protocol works at the control plane level between the two nodes. With GRES for MACsec is enabled, when one node initiates an RE switchover, it sends an MKA hello packet with a suspension request to the peer node. The peer node suspends the MACsec session at the control plane.

At the data plane, traffic continues to traverse the point-to-point link during suspension. The secure association key (SAK) that was programmed prior to suspension remains in use until the switchover is complete. After the switchover, the key server generates a new SAK to secure the link. The key server will continue to periodically create and share a SAK over the link for as long as MACsec is enabled.

You must configure the `suspend-on-request` statement on the key server so that it can accept a request for suspension from the peer node. If this statement is not configured, a suspension request sent to the key server results in termination of the MACsec session, which can cause traffic loss during the switchover.

Default

You must configure the `suspend-on-request` statement to enable a node on a MACsec link to accept requests for suspension of MACsec sessions. Suspension requests are not accepted by default.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

swap (protocols mpls)

IN THIS SECTION

- Syntax | 28739
- Hierarchy Level | 28739
- Description | 28739
- Options | 28739
- Required Privilege Level | 28739
- Release Information | 28739

Syntax

```
swap out-label;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
transit incoming-label],
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Description

Remove the label at the top of the label stack and replace it with the specified label. Manually assigned incoming labels can have values from 1,000,000 through 1,048,575. This statement is used to configure static LSPs at transit routers.

Options

out-label—Manually assigned outgoing label value.

- **Range:** 0 through 1,048,575
- **Default:** If you do not define the *out-label* option, the original label value remains unchanged.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

pop (protocols mpls) | **12797**

push (protocols mpls) | **13873**

swap

IN THIS SECTION

- Syntax | 28740
- Hierarchy Level | 28740
- Description | 28740
- Required Privilege Level | 28741
- Release Information | 28741

Syntax

```
swap;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

Description

Specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.

On MX Series routers, you can enter this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet using Gigabit Ethernet IQ interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP. On EX Series switches, you can enter this

statement on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Rewriting the VLAN Tag on Tagged Frames

Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support

swap-by-poppush

IN THIS SECTION

- Syntax | 28741
- Hierarchy Level | 28742
- Description | 28742
- Default | 28742
- Required Privilege Level | 28742
- Release Information | 28742

Syntax

```
swap-by-poppush;
```

Hierarchy Level

```
[edit interfaces  interface-name unit logical-unit-number]
```

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the VLAN rewrite operation to replace a VLAN tag. Pop original tag, then push an entirely new tag. The swap operation is performed as pop followed by push.

Default

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2

swap (Dynamic VLANs)

IN THIS SECTION

- Syntax | 28743
- Hierarchy Level | 28743
- Description | 28743
- Required Privilege Level | 28743
- Release Information | 28743

Syntax

```
swap;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number input-vlan-map],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map]
```

Description

For dynamic VLAN interfaces, specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

> Rewriting the VLAN Tag on Tagged Frames
>
> *Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution*

swap-push

IN THIS SECTION

- Syntax | 28744
- Hierarchy Level | 28744
- Description | 28744
- Required Privilege Level | 28745
- Release Information | 28745

Syntax

```
swap-push;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

Description

Specify the VLAN rewrite operation to replace the outer VLAN tag of the frame with a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Rewriting a VLAN Tag and Adding a New Tag

swap-swap

IN THIS SECTION

- Syntax | 28745
- Hierarchy Level | 28745
- Description | 28746
- Required Privilege Level | 28746
- Release Information | 28746

Syntax

```
swap-swap;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
```

```
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

Description

Specify the VLAN rewrite operation to replace both the inner and the outer VLAN tags of the frame with a user-specified VLAN tag value.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and for 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Rewriting the Inner and Outer VLAN Tags

switch-away-lsps

IN THIS SECTION

- Syntax | 28747
- Hierarchy Level | 28747
- Description | 28747
- Required Privilege Level | 28747
- Release Information | 28747

Syntax

```
switch-away-lsps;
```

Hierarchy Level

```
[edit logical-systems logical-systems-name protocols mpls interface interface-name],
[edit protocols mpls interface interface-name]
```

Description

(MX Series routers only) Enable you to switch an LSP away from a network node using a bypass LSP. This feature could be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic. Configure this statement only after you have configured and committed the `always-mark-connection-protection-tlv` statement.

The `always-mark-connection-protection-tlv` statement marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. When you configure the `switch-away-lsps` statement, traffic is switched to the bypass LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Switching LSPs Away from a Network Node

switch-options (Interfaces)

IN THIS SECTION

- Syntax | 28748
- Hierarchy Level | 28748
- Description | 28748
- Required Privilege Level | 28748
- Release Information | 28749

Syntax

```
switch-options {
    switch-port port-number {
        (auto-negotiation | no-auto-negotiation);
        speed (10m | 100m | 1g);
        link-mode (full-duplex | half-duplex);
    }
}
```

Hierarchy Level

```
[edit interfaces ge-pim/0/0]
```

Description

Configuration of the physical port characteristics is done under the single physical interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

switch-options

IN THIS SECTION

- Syntax | 28749
- Hierarchy Level | 28750
- Description | 28750
- Options | 28750
- Required Privilege Level | 28750
- Release Information | 28750

Syntax

```
switch-options {
    interface interface-name {
        mac-pinning;
        mac-learning-priority value packet-action action;
        interface-mac-limit limit;
    }
    interface-mac-limit limit {
        packet-action drop;
    }
    mac-statistics;
    mac-table-size limit {
        packet-action drop;
    }
    no-mac-learning;
    route-distinguisher (as-number:id | ip-address:id);
     service-id number; number;
    vrf-target {
        community;
```

```
        auto
        import community-name;
        export community-name;
    }
    vrf-import[ policy-names ];
    vrf-export[ policy-names ];
}
```

Hierarchy Level

```
[edit],
[edit logical-systems logical-system-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Configure Layer 2 learning and forwarding properties for a set of bridge domains.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support for logical systems added in Junos OS Release 9.6.

14.1x53-D10

RELATED DOCUMENTATION

Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports

switch-options (VLANs)

IN THIS SECTION

- Syntax (EX Series, MX Series, QFX Series and NFX Series) | **28751**
- Syntax (SRX Series) | **28752**
- EX Series, MX Series, QFX Series and NFX Series | **28752**
- SRX Series | **28752**
- Description | **28753**
- Required Privilege Level | **28753**
- Release Information | **28753**

Syntax (EX Series, MX Series, QFX Series and NFX Series)

```
switch-options {
    interface interface-name {
        interface-mac-limit limit {
            packet-action drop;
        }
        mac-pinning
        no-mac-learning;
        static-mac static-mac-address {
            vlan-id number;
        }
    }
    interface-mac-limit limit {
        packet-action drop;
    }
    mac-statistics;
    mac-ip-table-size number;
```

```
    mac-table-size limit {
        packet-action drop;
    }
    no-mac-learning;
    service-id number;
    vtep-source-interface
}
```

Syntax (SRX Series)

```
switch-options {
    interface interface-name {
        encapsulation-type;
        ignore-encapsulation-mismatch;
        pseudowire-status-tlv;
        static-mac mac-address {
            vlan-id vlan-id;
        }
    }
    mac-table-aging-time seconds;
    mac-table-size {
        number;
        packet-action drop;
    }
}
```

EX Series, MX Series, QFX Series and NFX Series

```
[edit ],
[edit logical-systems logical-system-name routing-instances routing-instance-name vlans vlan-
name],
[edit routing-instances routing-instance-name vlans vlan-name],
[edit vlans vlan-name]
```

SRX Series

```
[edit vlans vlans-name]
```

Description

Configure Layer 2 learning and forwarding properties for a VLAN or a virtual switch.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.5.

Statement introduced in Junos OS Release 12.3R2.

Statement (mac-pinning) introduced in Junos OS 16.2 for MX Series routers.

`mac-ip-table-size` statement introduced in Junos OS 17.4 Release for MX Series routers and EX9200 switches.

RELATED DOCUMENTATION

Ethernet Switching and Layer 2 Transparent Mode Overview

switch-port

IN THIS SECTION

- Syntax | 28754
- Hierarchy Level | 28754
- Description | 28754
- Default | 28754
- Options | 28754
- Required Privilege Level | 28754

Release Information | 28755

Syntax

```
switch-port port-number {
    (auto-negotiation | no-auto-negotiation);
    speed (10m | 100m | 1g);
    link-mode (full-duplex | half-duplex);
}
```

Hierarchy Level

```
[edit interfaces ge-pim/0/0 switch-options]
```

Description

Configuration of the physical port characteristics, done under the single physical interface.

Default

Autonegotiation is enabled by default. If the link speed and duplex are also configured, the interfaces use the values configured as the desired values in the negotiation.

Options

port-number—Ports are numbered 0 through 5 on the 6-port Gigabit Ethernet uPIM, 0 through 7 on the 8-port Gigabit Ethernet uPIM, and 0 through 15 on the 16-port Gigabit Ethernet uPIM.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

switch (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 28755
- Hierarchy Level | 28756
- Description | 28756
- Options | 28756
- Required Privilege Level | 28756
- Release Information | 28756

Syntax

```
switch switch-name {
    controller {
        address address;
        id id;
        protocol tcp {
            port port;
        }
        role equal;
    }
    default-action (drop | packet-in);
    interfaces {
        interface-name port-id port;
    }
    purge-flow-timer seconds;
}
```

Hierarchy Level

```
[edit protocols openflow]
```

Description

Configure an OpenFlow virtual switch.

Options

switch-name User-configured identifier for the OpenFlow virtual switch. The identifier can include a maximum of 60 characters.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Support for OpenFlow on Devices Running Junos OS

controller (Protocols OpenFlow) | **4654**

default-action (Protocols OpenFlow) | **5179**

interfaces (Protocols OpenFlow) | **9153**

openflow (Protocols OpenFlow) | **11842**

switch-type

IN THIS SECTION

- Syntax | 28757
- Hierarchy Level | 28757
- Description | 28757
- Options | 28757
- Required Privilege Level | 28758
- Release Information | 28758

Syntax

```
switch-type (att5e | etsi | ni1 | ntdms-100)
```

Hierarchy Level

```
[edit interfaces br-pim/0/port isdn-options]
```

Description

For J Series Services Routers only. Configure the ISDN variant supported.

Options

`att5e`—AT&T switch variant.

`etsi`—European Telecommunications Standards Institute switch variant.

`ni1`—National ISDN 1 switch variant.

`ntdms-100`—Northern Telecom DMS-100.

`ntt`—NTT Group switch for Japan.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

switching-type

IN THIS SECTION

- Syntax | 28758
- Hierarchy Level | 28758
- Description | 28759
- Default | 28759
- Required Privilege Level | 28759
- Release Information | 28759

Syntax

```
switching-type (fiber | lambda | psc-1 | tdm);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name lsp-
attributes],
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
```

Description

Specify the switching method for the LSP. The switching method can be one of the following values:

- `fiber`—Fiber switching
- `lambda`—Lambda switching
- `psc-1`—Packet switching
- `tdm`—Time-division multiplexing (TDM) switching

Default

`psc-1`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring MPLS LSPs for GMPLS

switchover-delay

IN THIS SECTION

- Syntax | 28760
- Hierarchy Level | 28760
- Description | 28760

- Options | 28760
- Required Privilege Level | 28760
- Release Information | 28761

Syntax

```
switchover-delay milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
neighbor address],
[edit protocols l2circuit neighbor address interface interface-name],
[edit routing-instances routing-instance-name protocols vpls neighbor address]
```

Description

After the primary pseudowire goes down, specifies the delay (in milliseconds) to wait before the backup pseudowire takes over. You configure this statement for each backup neighbor configuration to adjust the switchover time after a failure is detected.

Options

milliseconds—Specify the time to wait before switching to the backup pseudowire after the primary pseudowire fails.

- **Default:** 10,000 milliseconds
- **Range:** 0 through 180,000 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring the Switchover Delay for the Pseudowires

switchover-mode

IN THIS SECTION

- Syntax | 28761
- Hierarchy Level | 28761
- Description | 28762
- Default | 28762
- Options | 28762
- Required Privilege Level | 28762
- Release Information | 28762

Syntax

```
switchover-mode (revertive | non-revertive);
```

Hierarchy Level

```
[edit chassis synchronization]
```

Description

Specify whether the device should switch to a higher-quality clock source or use the current clock source only.

Default

The default is revertive mode.

Options

`revertive` In revertive mode (the default), the system automatically switches from a lower to a higher quality clock source whenever the higher clock source becomes available.

`non-revertive` Set this option so that the router continues to use the current clock source as long as it is valid.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2R4.

RELATED DOCUMENTATION

- synchronization (ACX Series) | 28787
- synchronization (MX Series) | 28795
- synchronization (PTX Series) | 28802
- quality-level
- show chassis synchronization | 18879

switchover-on-routing-crash

IN THIS SECTION

- Syntax | 28763
- Hierarchy Level | 28763
- Description | 28763
- Required Privilege Level | 28763
- Release Information | 28764

Syntax

```
switchover-on-routing-crash;
```

Hierarchy Level

```
[edit system]
```

Description

Prevent loss of traffic in the case of NSR being configured. With the `switchover-on-routing-crash` configuration statement enabled, when rpd on the primary Routing Engine crashes with NSR configured, the Routing Engine will switch over immediately to the backup Routing Engine to preserve protocol state and adjacencies. Prior to having this statement, if NSR was configured and rpd on the primary Routing Engine crashed, it would cause network impact (protocol neighbor and adjacency drops and traffic loss).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Configuring Nonstop Active Routing

sxl-retry

IN THIS SECTION

- Syntax | 28764
- Hierarchy Level | 28764
- Description | 28765
- Options | 28765
- Required Privilege Level | 28765
- Release Information | 28765

Syntax

```
sxl-retry value;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus sophos-engine]
```

Description

Configure the number of retry attempts to the remote Sophos Extensible List (SXL) server when a request timeout occurs.

Starting in Junos OS Release 23.1R1, we've deprecated the `sxl-retry` statement at the [`edit security utm default-configuration anti-virus sophos-engine`] and [`security utm feature-profile anti-virus sophos-engine`] hierarchy levels.

Options

`value` —Number of retries.

- **Range:** 0 through 5
- **Default:** 1

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The [`edit security utm default-configuration`] hierarchy level is introduced in Junos OS Release 18.2R1.

sxl-timeout

IN THIS SECTION

- Syntax | 28766
- Hierarchy Level | 28766
- Description | 28766
- Options | 28766
- Required Privilege Level | 28766

Release Information | 28766

Syntax

```
sxl-timeout seconds;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus sophos-engine]
```

Description

Configure the timeout value for responses to a Sophos checksum or URI query.

Starting in Junos OS Release 23.1R1, we've deprecated the `sxl-timeout` statement at the [`edit security utm default-configuration anti-virus sophos-engine`] and [`security utm feature-profile anti-virus sophos-engine`] hierarchy levels.

Options

seconds —Number of seconds before timeout occurs.

- **Range:** 1 through 5 seconds
- **Default:** 2 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The [edit security utm default-configuration] hierarchy level is introduced in Junos OS Release 18.2R1.

symbol-period

IN THIS SECTION

- Syntax | 28767
- Hierarchy Level | 28767
- Description | 28767
- Options | 28768
- Required Privilege Level | 28768
- Release Information | 28768

Syntax

```
symbol-period count;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile name event link-event-rate],
[edit protocols oam ethernet link-fault-management interface name No Link Title]
```

Description

Configure the threshold for sending symbol period events or taking the action specified in the action profile.

A symbol error is any symbol code error on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period window. The default period window is the number of symbols that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.

Options

count—Threshold count in seconds for symbol period events.

- **Range:** 0 through 100 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configure Ethernet OAM Link Fault Management

symmetric-hash

IN THIS SECTION

- Syntax | 28769
- Hierarchy Level | 28769
- Description | 28769
- Options | 28769
- Required Privilege Level | 28769
- Release Information | 28769

Syntax

```
symmetric-hash {
    complement;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic slot-number hash-key family inet],
[edit chassis fpc slot-number pic slot-number hash-key family multiservice]
```

Description

(MX Series 5G Universal Routing Platforms only) Configure the symmetric hash or symmetric hash complement at the PIC level for configuring symmetrical load balancing on an 802.3ad Link Aggregation Group.

Options

complement Include the complement of the symmetric hash in the hash key.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

> *Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers*

syn-ack-ack-proxy

IN THIS SECTION

- Syntax | 28770
- Hierarchy Level | 28770
- Description | 28770
- Options | 28770
- Required Privilege Level | 28771
- Release Information | 28771

Syntax

```
syn-ack-ack-proxy; {
    threshold number,
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp]
```

Description

Prevent the SYN-ACK-ACK attack, which occurs when the attacker establishes multiple telnet sessions without allowing each session to terminate. This behavior consumes all open slots, generating a denial-of-service (DoS) condition.

Options

threshold *number* — Number of connections from any single IP address.

- **Range:** 1 through 250,000
- **Default:** 512

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5; support.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

syn-ack-ack-proxy (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 28771
- Hierarchy Level | 28772
- Description | 28772
- Options | 28772
- Required Privilege Level | 28772
- Release Information | 28772

Syntax

```
syn-ack-ack-proxy {
    threshold number;
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Description

Configure the maximum number of connections from an IP address that can be opened without being completed. Once this threshold has been reached, further connection requests are rejected. In the SYN-ACK-ACK attack, the session table can fill up, resulting in the device rejecting legitimate connection requests.

Options

threshold *number* Maximum number of uncompleted connections from any single IP address.

- **Range:** 1 through 250,000
- **Default:** 512

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

syn-check-required

IN THIS SECTION

- Syntax | 28773
- Hierarchy Level | 28773
- Description | 28773
- Required Privilege Level | 28773
- Release Information | 28773

Syntax

```
syn-check-required;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name
policy policy-name then permit tcp-options]
```

Description

Enable sync check per policy. The syn-check-required value overrides the global value no-syn-check.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Attack Detection and Prevention Overview
- Example: Configuring Multiple Screening Options

syn-cookie (IDS MS-DPC)

IN THIS SECTION

- Syntax | 28774
- Hierarchy Level | 28774
- Description | 28774
- Options | 28775
- Required Privilege Level | 28775
- Release Information | 28775

Syntax

```
syn-cookie {
    mss value;
    threshold rate;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Enable SYN-cookie defenses against SYN attacks when using the MS-DPC. By default, SYN-cookie techniques are not applied.

When SYN cookie is enabled on Junos OS and becomes the TCP-negotiating proxy for the destination server, it replies to each incoming SYN segment with a SYN/ACK containing an encrypted cookie as its initial sequence number (ISN). The cookie is an MD5 hash of the original source address and port number, destination address and port number, and ISN from the original SYN packet. After sending the cookie, Junos OS drops the original SYN packet and deletes the calculated cookie from memory. If there is no response to the packet containing the cookie, the attack is noted as an active SYN attack and is effectively stopped.

If the initiating host responds with a TCP packet containing the cookie +1 in the TCP ACK field, Junos OS extracts the cookie, subtracts 1 from the value, and recomputes the cookie to validate that it is a legitimate ACK. If it is legitimate, Junos OS starts the TCP proxy process by setting up a session and sending a SYN to the server containing the source information from the original SYN. When Junos OS receives a SYN/ACK from the server, it sends ACKs to the server and to the initiation host. At this point the connection is established and the host and server are able to communicate directly.

Options

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in IDS Rules

syn-fin

IN THIS SECTION

- Syntax | 28776
- Hierarchy Level | 28776
- Description | 28776
- Required Privilege Level | 28776
- Release Information | 28776

Syntax

```
syn-fin;
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp]
```

Description

Enable detection of an illegal combination of flags that attackers can use to consume sessions on the target device, thus resulting in a denial-of-service (DoS) condition.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

syn-fin (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 28777
- Hierarchy Level | 28777
- Description | 28777
- Required Privilege Level | 28777
- Release Information | 28778

Syntax

```
syn-fin;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Description

Identify and drop packets that have both the SYN and FIN flags set, which can cause unpredictable behavior.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

syn-flood

IN THIS SECTION

- Syntax | 28778
- Hierarchy Level | 28779
- Description | 28779
- Options | 28779
- Required Privilege Level | 28779
- Release Information | 28779

Syntax

```
syn-flood {
    alarm-threshold number;
    attack-threshold number;
    destination-threshold number;
    source-threshold number;
    timeout seconds;
    white-list name {
        destination-address destination-address;
        source-address source-address;
    }
}
```

Hierarchy Level

```
[edit security screen ids-option          screen-name        tcp]
```

Description

Configure detection and prevention of SYN flood attacks. Such attacks occur when the connecting host continuously sends TCP SYN requests without replying to the corresponding ACK responses.

> **NOTE**: On all SRX Series Firewalls, the TCP synchronization flood alarm threshold value does not indicate the number of packets dropped, however the value does show the packet information after the alarm threshold has been reached.
>
> The synchronization cookie or proxy never drops packets; therefore the `alarm-without-drop` (not `drop`) action is shown in the system log.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

syn-flood-protection-mode

IN THIS SECTION

- Syntax | 28780
- Hierarchy Level | 28780

- Description | 28780
- Options | 28780
- Required Privilege Level | 28780
- Release Information | 28780

Syntax

```
syn-flood-protection-mode (syn-cookie | syn-proxy);
```

Hierarchy Level

```
[edit security flow]
```

Description

Enable SYN cookie or SYN proxy defenses against SYN attacks. SYN flood protection mode is enabled globally on the device and is activated when the configured `syn-flood attack-threshold value` is exceeded.

Options

- `syn-cookie`—Uses a cryptographic hash to generate a unique Initial Sequence Number (ISN). This is enabled by default.
- `syn-proxy`—Uses a proxy to handle the SYN attack.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5; support.

syn-frag

IN THIS SECTION

- Syntax | 28781
- Hierarchy Level | 28781
- Description | 28781
- Required Privilege Level | 28781
- Release Information | 28782

Syntax

```
syn-frag;
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp]
```

Description

Enable detection of a SYN fragment attack and drops any packet fragments used for the attack. A SYN fragment attack floods the target host with SYN packet fragments. The host caches these fragments, waiting for the remaining fragments to arrive so it can reassemble them. The flood of connections that cannot be completed eventually fills the host's memory buffer. No further connections are possible, and damage to the host's operating system can occur.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

syn-frag (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 28782
- Hierarchy Level | 28782
- Description | 28782
- Required Privilege Level | 28782
- Release Information | 28783

Syntax

```
syn-frag;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Description

Identify and drop SYN packet fragments. In TCP SYN fragment attacks, the target caches SYN fragments, waiting for the remaining fragments to arrive so it can reassemble them and complete the connection. A flood of SYN fragments eventually fills the host's memory buffer, preventing valid traffic connections.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

sync-active-path-bandwidth

IN THIS SECTION

- Syntax | 28783
- Hierarchy Level | 28783
- Description | 28784
- Default | 28784
- Required Privilege Level | 28784
- Release Information | 28784

Syntax

```
sync-active-path-bandwidth;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mplslabel-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-
bandwidth],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
```

```
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Description

When you have a primary and a secondary path configuration, specify that a path needs to be signaled with the active-path bandwidth when the auto-bandwidth adjustment happens and that the secondary path synchronizes the bandwidth reservations to that of the primary path.

When a primary path fails, bandwidth reservations are made by the secondary path on the links that it uses. If you include the `sync-active-path-bandwidth` statement, the secondary path releases the bandwidth it has reserved and adjusts its bandwidth after the primary path begins carrying traffic.

For example, suppose the active path is a secondary path with a reserved bandwidth of 10 GB as a result of the automatic bandwidth adjustment. Then suppose there is a switchover from the secondary path to the primary path. After some time the primary path reserves 5 GB as a result of a new automatic adjustment. Without the `sync-active-path-bandwidth` statement, the secondary path does not release the 10 GB after a switchover occurs. That bandwidth is wasted. If the `sync-active-path-bandwidth` is included in the configuration, the secondary path adjusts its bandwidth to 5 GB along with the primary path.

Default

When you have a primary and a secondary path configuration, and the primary path fails, bandwidth reservations are made by the secondary path on the links that it uses. When the primary path comes back and the traffic switches over, the secondary path does not release its bandwidth reservations.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Disabling Constrained-Path LSP Computation

Configuring Explicit-Path LSPs

sync-reset

IN THIS SECTION

- Syntax | 28785
- Hierarchy Level | 28785
- Description | 28785
- Default | 28786
- Options | 28786
- Required Privilege Level | 28786
- Release Information | 28786

Syntax

```
sync-reset (disable | enable);
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options lacp],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-ether-options lacp],
[edit interfaces name aggregated-ether-options lacp]
```

Description

For redundant Ethernet interface link aggregation group links, you can configure the minimum number of physical child links on the primary node in the redundant Ethernet interface that must be working. If the minimum number of operating child links falls below the configured value, the interface is marked

down even if some of the child interfaces are working. LACP marks these operating child interfaces or links that are working, as out of synchronization. This enables a peer switch, that does not have the minimum link configuration, to mark its interface as down as well. The peer switch can be a Juniper Switch or any other third party switch. As a result, the interface is marked as down on both the switches until the number of operating child links is more than the configured value.

For an aggregated ethernet interface, you cannot configure all three configuration options, `bfd-liveness-detection`, `minimum-links`, and `sync-reset` at the same time.

Default

(ACX Series routers) Disabled.

(MX Series routers) Disabled.

(PTX Series routers) Disabled.

(QFX Series switches) Disabled.

(EX Series switches) Disabled.

Options

disable	Disable synchronization reset.
enable	Enable synchronization reset.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

minimum-links | **10741**

show lacp interfaces | **21829**

synchronization (ACX Series)

IN THIS SECTION

- Syntax | 28787
- Syntax (ACX710) | 28788
- Syntax (ACX7100-48L and ACX7100-32C) | 28789
- Hierarchy Level | 28791
- Description | 28791
- Options | 28791
- Required Privilege Level | 28791
- Release Information | 28791

Syntax

```
synchronization {
    clock-mode (auto-select | free-run);
    esmc-transmit {
        interfaces {
            (all | interface-name);
        }
    }
    global-wait-to-restore minutes;
    hold-interval {
        configuration-change seconds;
        restart seconds;
        switchover seconds;
    }
    interfaces bits {
        signal-type (e1 | t1);
        e1-options {
            framing (g704 | g704-no-cro4);
        }
        t1-options {
            framing (esf | sf);
        }
```

```
    }
    network-option (option-1 | option-2);
    quality-mode-enable;
    selection-mode (configured-quality | received-quality);
    source {
        interfaces interface-name {
            priority number;
            quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu |
tnc);
            request (force-switch | lockout);
            wait-to-restore minutes;
        }
    }
    switchover-mode (revertive | non-revertive);
}
```

Syntax (ACX710)

```
synchronization {
    clock-mode (auto-select | free-run);
    esmc-transmit {
            interfaces {
    (all | interface-name) {
    squelch {
    method ql-dnu;
    quality-level quality-level;
    }
    }
    }
    }
        global-wait-to-restore minutes;
    hold-interval {
    configuration-change seconds;
    restart seconds;
    switchover seconds;
    }
        interfaces bits {
    signal-type (e1 | t1);
    e1-options {
    framing (g704 | g704-no-cro4);
    }
```

```
    t1-options {
    framing (esf | sf);
    }
    }
        minimum-quality (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu |
tnc);
    network-option (option-1 | option-2);
        output interfaces {
    bits {
                source-mode (chassis | ptp-clock);
    squelch {
    method (ais | ql-dnu | rai | shut);
    }
    ssm option-2  (GEN-1 | GEN-2);
    }
    }
        quality-mode-enable;
    selection-mode (configured-quality | received-quality);
    source {
    interfaces interface-name {
    priority number;
    quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
    request (force-switch | lockout);
    wait-to-restore minutes;
    }
    }
    switchover-mode (revertive | non-revertive);
    }
```

Syntax (ACX7100-48L and ACX7100-32C)

```
synchronization {
    clock-mode(auto-select | free-run);
    enable-extended-ql-tlv;
    esmc-transmit {
       interfaces (all | interface-name) {
       squelch {
          method(ais | ql-dnu | rai | shut);
          quality-level (eEEC | ePRC | ePRTC | prc | prs | PRTC | sec | smc | ssu-a | ssu-b |
st2 | st3 | st3e | st4 | stu | tnc);
          }
```

```
        }
    }
    global-wait-to-restore minutes;
    hold-interval {
        configuration-change seconds;
        restart seconds;
        switchover seconds;
        }
    loop-prevention master-side;
    max-transmit-quality-level(eEEC | ePRC | ePRTC | prc | prs | PRTC | sec | ssu-a | ssu-b | st2
| st3e | stu | tnc);
    minimum-quality (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
    network-option(option-1 | option-2);
    output {
       interfaces interface-name {
       holdover-mode-disable;
       minimum-quality(eEEC | ePRC | ePRTC | prc | prs | PRTC | sec | smc | ssu-a | ssu-b | st2
| st3 | st3e | st4 | stu | tnc);
       source-mode(chassis | equipment-clock | line | ptp-clock);
       squelch {
        method(ais | ql-dnu | rai | shut);
        }
       }
       ssm {

    }
    quality-mode-enable;
    selection-mode (configured-quality | received-quality);
    source {
        interfaces interface-name {
        hold-off-time milliseconds;
        no-ssm-support;
        priority priority;
        quality-level(eEEC | ePRC | ePRTC | prc | prs | PRTC | sec | smc | ssu-a | ssu-b | st2 |
st3 | st3e | st4 | stu | tnc);
        request(force-switch | lockout);
        wait-to-restore min;
        }
    }
    switchover-mode(non-revertive | revertive);
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure Synchronous Ethernet parameters.

> **NOTE:**
> Synchronous Ethernet

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

- *External Clock Synchronization Overview for ACX Series Routers*
- *Configuring External Clock Synchronization for ACX Series Routers*
- *Synchronous Ethernet Overview*
- show chassis synchronization | **18879**

synchronization (M Series and T Series)

IN THIS SECTION

- Syntax | 28792
- Hierarchy Level | 28792
- Description | 28793
- Options | 28793
- Required Privilege Level | 28794
- Release Information | 28794

Syntax

```
synchronization {
    primary (external-a | external-b | fpc-slot-number | gps-0-10mhz | gps-0-5mhz | gps-1-10mhz
| gps-1-5mhz | bits-a | bits-b);
    secondary (external-a | external-b | fpc-slot-number | gps-0-10mhz | gps-0-5mhz |
gps-1-10mhz | gps-1-5mhz | bits-a | bits-b );
    signal-type (t1 | e1);
    switching-mode (revertive | non-revertive);
    transmitter-enable;
    validation-interval seconds;
    y-cable-line-termination;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

(M320, M40e, M120, T320, T640, and T1600 routers and PTX Series Packet Transport Routers only) Configure an external synchronization interface to synchronize the internal Stratum 3 clock to an external source, and then synchronize the chassis interface clock to that source.

Options

`primary`—First external timing source specified in the configuration hierarchy. This statement has the following suboptions:

- `external-a`—Use `external-a` as the timeTransmitter clock synchronization source.
- `external-b`—Use `external-b` as the timeTransmitter clock synchronization source.
- `fpc-`*`slot-number`*—Use `fpc-`*`slot-number`* as the timeTransmitter clock synchronization source. For the PTX5000 Packet Transport Router, replace *slot-number* with a value from `0` through `7`.
- `gps-0-10mhz`—Use `gps-0-10mhz` as the timeTransmitter clock synchronization source.
- `gps-0-5mhz`—Use `gps-0-5mhz` as the timeTransmitter clock synchronization source.
- `gps-1-10mhz`—Use `gps-1-10mhz` as the timeTransmitter clock synchronization source.
- `gps-1-5mhz`—Use `gps-1-5mhz` as the timeTransmitter clock synchronization source.
- `bits-a`—Use `bits-a` as the timeTransmitter clock synchronization source.
- `bits-b`—Use `bits-b` as the timeTransmitter clock synchronization source.

`secondary`—Second external timing source specified in the configuration hierarchy.

- `external-a`—Use `external-a` as the secondary clock synchronization source.
- `external-b`—Use `external-b` as the secondary clock synchronization source.
- `fpc-`*`slot-number`*—Use `fpc-`*`slot-number`* as the secondary clock synchronization source. For the PTX5000 Packet Transport Router, replace *slot-number* with a value from `0` to `7`.
- `gps-0-10mhz`—Use`gps-0-10mhz` as the secondary clock synchronization source.
- `gps-0-5mhz`—Use `gps-0-5mhz` as the secondary clock synchronization source.
- `gps-1-10mhz`—Use `gps-1-10mhz` as the secondary clock synchronization source.
- `gps-1-5mhz`—Use `gps-1-5mhz` as the secondary clock synchronization source.
- `bits-a`—Use `bits-a` as the secondary clock synchronization source.

- `bits-b`—Use `bits-b` as the secondary clock synchronization source.

`signal-type`—Specify the line encoding mode for interfaces: either `t1` or `e1`. For the M40e router, only the `t1 signal-type` mode is supported.

- **Default:** `t1`

`switching-mode`—Specify `revertive` if a lower-priority synchronization can be switched to a valid, higher-priority synchronization.

- **Default:** `non-revertive`

`transmitter-enable`—(M320 routers only) Control whether the diagnostic timing signal is transmitted.

`validation-interval`—Validate the synchronized deviation. If revertive switching is enabled and a higher-priority clock is validated, the clock module is directed to the higher-priority clock, and all configured and active synchronizations are validated. The validation timer resumes after the current validation interval expires. This feature is not supported on PTX Series Packet Transport Routers.

- **Range:** (M320, M40e, T320, T640, T1600 routers) 90 through 86,400 seconds.

  (M120 routers) 30 through 86,400 seconds.

- **Default:** (M320, M40e, T320, T640, T1600 routers) 90 seconds. (M120 routers) 30 seconds

`y-cable-line-termination`—(M320 routers only) Specify that a single signal be wired to both Control Boards (CBs) using a Y-cable.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

> *Configuring Junos OS to Support an External Clock Synchronization Interface for M Series, MX Series, and T Series Routers*

synchronization (MX Series)

IN THIS SECTION

- Syntax | 28795
- Hierarchy Level | 28797
- Description | 28797
- Options | 28797
- Required Privilege Level | 28801
- Release Information | 28801

Syntax

```
synchronization {
    clock-mode (auto-select | free-run);
    esmc-transmit {
        interfaces (all | interface-name);
    }
    hold-interval {
        configuration-change seconds;
        restart seconds;
        switchover seconds;
    }
    interfaces {
        external | (external-0/0 | external-1/0) {
            e1-options {
                framing (g704 | g704-no-crc4);
                line-encoding (ami | hdb3);
                sabit number;
            }
            pulse-per-second-enable;
            signal-type (2048khz | e1 | t1);
            t1-options {
                framing (esf | sf);
                line-encoding (ami | b8zs);
            }
        }
```

```
        bits {
            e1-options {
                framing (g704 | g704-no-crc4);
                sabit number;
                }
            }
        }
    }
    max-transmit-quality-level (prc | prs | sec | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu |
tnc);
    network-option (option-1 | option-2);
    output interfaces {
         (external | (external-0/0 | external-1/0)) {
            holdover-mode-disable;
            minimum-quality (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu
| tnc);
            source-mode (chassis | line);
            tx-dnu-to-line-source-enable;
            wander-filter-disable;
        }
    }
    port {
        auxiliary {
            client {
                time-of-day-format {
                    ascii string;
                }
            }
        }
    }
    quality-mode-enable;
    selection-mode (configured-quality | received-quality);
    source {
        interfaces (interface-name | (bits |external | (external-0/0 | external-1/0)) {
            priority number;
            quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu |
tnc);
            request (force-switch | lockout);
        }
    }
    switchover-mode (revertive | non-revertive);
}
```

Hierarchy Level

```
[edit chassis]
```

Description

(MX5-T, MX10-T, MX40-T, MX80-T, MX240, MX480, MX960 , and MX10003 routers) Configure Synchronous Ethernet parameters. For configuration details, see *Configuring Clock Synchronization Interface on MX Series Routers*.

(MX240, MX480, MX960, and MX2020 routers with SCBE or SCBE2) Configure centralized clocking parameters.

- Starting in Junos 12.2, configure distribution of the selected chassis clock source to downstream network elements through supported line interfaces.
- Starting in Junos 12.3, configure an external building-integrated timing supply (BITS) timing source. You can also configure the selected chassis clock, or an incoming Synchronous Ethernet or PTP line source for transmission out the external interface.
- Starting in Junos 16.1 for the MX104 routers, SSM is supported on the BITS interface and can be configured at the `[edit chassis synchronization interfaces bits e1-options]` hierarchy level.

For configuration details, see *Example: Configuring Centralized Clocking on the Enhanced MX Switch Control Board*.

NOTE:

- Unified ISSU is not supported when clock synchronization is configured for Synchronous Ethernet.
- The commit operation fails if you configure bits as source, without first configuring bits interface.

Options

`clock-mode (auto-select | free-run)`—Specify the mode of operation to select the clock source either from a free-run local oscillator or from an external qualified clock. On MX240, MX480, and MX960 routers with enhanced MPCs, the free-run clock is provided by a local oscillator. On other MX Series routers, the free-run clock is provided by the SCB. The default setting is auto-select mode.

`hold-interval (configuration-change | restart | switchover)` *seconds*—Specify the chassis synchronization hold interval type and clock selection wait time:

- `configuration-change`—Clock select wait time after change in configuration. The range is 15 through 60 seconds.
- `restart`— Clock select wait time after reboot. The range is 60 through 180 seconds. The default is 120 seconds.
- `switchover`— Switchover wait time after clock recovery. The range is 30 to 60 seconds. The default is 30 seconds.

`interfaces (external | (external-0/0 | external-1/0))`—Configure the external interface for operating with a connected external device. This interface can be configured as a clock source, which then becomes a candidate for selection as the chassis clock source by the clock source selection algorithm.

- `signal-type (1mhz | 5mhz | 10mhz | 2048khz | e1 | t1)`—Specify the external interface signal type:
  1. `1mhz`—Set the signal with a clock frequency of 1 MHz.
  2. `5mhz`—Set the signal with a clock frequency of 5 MHz.
  3. `10mhz`—Set the signal with a clock frequency of 10 MHz.
  4. `2048khz`—Set the signal with a clock frequency of 2048 kHz.
  5. `e1`—Set the signal as an E1-coded 2048 kHz signal on a 120-ohm balanced line.
  6. `t1`—Set the signal as a T1-coded 1.544 MHz signal on a 100-ohm balanced line.
- `e1-options`—Specify the E1 options:
  1. `framing (g704 | g704-no-crc4)`—Specify the framing format:
     - `g704`—G.704 framing format for E1 interfaces
     - `g704-no-crc4`—G.704 framing with no CRC4 for E1 interfaces
  2. `line-encoding (ami | hdb3)`—Specify the line encoding:
     - `ami`— Alternate mark inversion (AMI)
     - `hdb3`—High-d Density bipolar 3 code (HDB3)
  3. `sabit` *number*—Specify the San synchronization status bit used for exchanging SSN quality. The value can be 4, 5, 6, 7, or 8. The default is 4.
- `t1-options`—Specify the T1 options:
  1. `framing (esf | sf)`—Specify the framing format:
     - `esf`—Extended superframe (ESF)

- `sf`—Superframe (SF)

2. `line-encoding (ami | b8zs)`—Specify the line encoding:
   - `ami`— Alternate mark inversion (AMI)
   - `b8zs`—8-bit zero suppression, bipolar with 8-zero substitution (B8ZS)

`network-option (option-1 | option-2)`—Specify the synchronization networking:

- `option-1`— EEC-1 maps to G.813 option 1 clock
- `option-2`—EEC-2 maps to G.812 type IV clock

`quality-mode-enable`—Specify the clock selection, quality level, and priority setting. The quality level parameter for a Synchronous Ethernet interface is optional when quality mode is enabled and the selection mode is set to `received-quality`. The default quality level for a Synchronous Ethernet interface is based on the value of `network-option`: `option-1` selects `SEC` and `option-2` selects `ST3`. Table 1802 on page 28799 shows whether SSM quality level is supported for a given external interface signal type and framing. The default setting is disabled.

Table 1802: SSM-Quality Level Support by Signal Type and Framing

Signal Type	Framing	SSM Quality Level Supported
E1	G.704	yes
E1	G.704 no CRC4	no
T1	ESF	yes
T1	SF	no
2048 KHz	Not applicable	no

`selection-mode (configured-quality | received-quality)`—Specify whether the clock source selection should use the configured or received ESMC or SSM quality level for a qualifying interface. In both selection modes, the interface qualifies for clock source selection only when the received ESMC or SSM quality level on the interface is equal to or greater than the configured ESMC or SSM quality level for the interface.

> **NOTE**: For the `selection-mode` statement configuration to take effect, you must set the `quality-mode-enable` statement at the `[edit chassis synchronization]` hierarchy level.

1. `configured-quality`—The clock source selection algorithm uses the ESMC or SSM quality level configured for a qualifying interface.
2. `received-quality`—The clock source selection algorithm uses the ESMC or SSM quality level received on the qualifying interface.

`source (external-a | external-b | interfaces (`*`interface-name`*`) | external | (external-0/0 | external-1/0))`—Specify clock sources.

(SCB only) The timeTransmitter clock source is external-a interface, the secondary clock source is external-b interface. The clock source is chosen using the clock selection process.

(SCBE only) Specify the external interface to select the external clock source.

(SCBE2 only) Specify the external-0/0 interface or external-1/0 interface to select the external clock source.

- `priority` *`number`*—Specify a priority level from 1 to 5. When not specified, `external-a` has higher default priority than external-b interface , and external-b interface has higher default priority than other Gigabit Ethernet or 10-Gigabit Ethernet clock sources, which have the lowest default priority. Any priority you configure is higher than any default priority.
- `quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc)`—Specify the `quality-level` option based on the configured `network-option`. For quality level details, see "quality-level" on page 13917.

> **NOTE**: Starting with Junos OS Release 12.2R1, the `quality-level` parameter for a Synchronous Ethernet interface is optional when quality mode is enabled and the selection mode is set to received quality. The default quality level for a Synchronous Ethernet interface is `SEC` for the `option-1` network type and `ST3` for the `option-2` network type.
>
> Both option I and option II SSM quality levels are supported:
>
> - For option-1, quality level must be configured for external clocks (`external-a` or `external-b`) whether or not quality level is enabled.
> - For option-2, the default quality level for external clocks is QL_STU whether or not quality level is enabled.

- Quality level is set to DNU for network-option 1 and set to DUS for network-option 2, if quality-level is not configured and no ESMC messages are received.
- On the selected active source (primary or secondary, whichever is active), even if ESMC transmit is not enabled, a DNU ESMC will be sent out if `network-option` is 1, and DUS ESMC will be sent out if `network-option` is 2. This is applicable only for sources of Ethernet interface type to avoid source looping.

- `request force-switch`—Force a switch to this source if the source is enabled and not locked out. You can configure only one source to be force-switched.
- `request lockout`—You can configure lockout for any source. When configured, this source is not considered by the clock selection process.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Statement introduced in Junos OS Release 17.3R1 for MX10003 router.

Options `interfaces`, `output`, and `source interfaces external` introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Synchronous Ethernet Overview

show chassis synchronization (MX Series Routers) | **18894**

Configuring Clock Synchronization Interface on MX Series Routers

Example: Configuring Framing Mode for Synchronous Ethernet on MX Series Routers with 10-Gigabit Ethernet MIC

Example: Configuring Centralized Clocking on the Enhanced MX Switch Control Board

Example: Configuring Centralized Clocking on an MX2020

request chassis synchronization mode | **14570**

Clock Sources for PTX Series Packet Transport Routers

synchronization (PTX Series)

IN THIS SECTION

- Syntax | 28802
- Hierarchy Level | 28803
- Description | 28803
- Options | 28803
- Required Privilege Level | 28803
- Release Information | 28803

Syntax

```
synchronization {
    clock-mode (auto-select | free-run);
    esmc-transmit {
        interfaces (all | interface-name);
    }
    global-wait-to-restore minutes;
    hold-interval {
        configuration-change seconds;
        restart seconds;
        switchover seconds;
    }
    interfaces (bits-a | bits-b | gps-0 | gps-1) (pulse-per-second-enable | signal-type (5mhz|
10mhz|e1|t1));
    max-transmit-quality-level (prc | prs | sec | ssu-a | ssu-b | st2 | st3e | stu | tnc);
    network-option (option-1 | option-2);
    primary (fpc-slot-number | gps-0 | gps-1 | bits-a | bits-b);
    quality-mode-enable;
    secondary (fpc-slot-number | gps-0 | gps-1 | bits-a | bits-b);
    selection-mode (configured-quality | received-quality);
    source (bits-a | bits-b | gps-0 | gps-1) {
        priority number;
        quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu |
tnc);
        request (force-switch | lockout);
```

```
    }
    switchover-mode (revertive | non-revertive);
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure Synchronous Ethernet parameters.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X48.

RELATED DOCUMENTATION

- *Synchronous Ethernet Overview*
- *Synchronizing Internal Stratum 3 Clock to External Clock Sources on PTX Series Routers*
- request chassis synchronization switch | **14574**
- *Clock Sources for PTX Series Packet Transport Routers*
- *Understanding Clock Synchronization*

synchronize

IN THIS SECTION

- Syntax | 28804
- Hierarchy Level | 28804
- Description | 28804
- Options | 28805
- Required Privilege Level | 28806
- Release Information | 28806

Syntax

```
synchronize;
```

Hierarchy Level

```
[edit system commit]
```

Description

For devices with multiple Routing Engines only. Configure the `commit` command to automatically perform a `commit synchronize` action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the `commit` command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engine. Each Routing Engine then performs a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.

NOTE: If you configure the `commit synchronize` statement at the `[edit system]` hierarchy level and issue a `commit` in the primary Routing Engine, the primary configuration is automatically synchronized with the backup. However, if the backup Routing Engine is down when you issue the `commit`, the Junos OS displays a warning and commits the

candidate configuration in the primary Routing Engine. When the backup Routing Engine comes up, its configuration will automatically be synchronized with the primary. A newly inserted backup Routing Engine automatically synchronizes its configuration with the primary Routing Engine configuration.

NOTE: When you configure nonstop active routing (NSR), you must configure the `commit synchronize` statement. Otherwise, the commit operation fails.

NOTE: Starting in Junos OS Release 20.2R1, when the `commit synchronize` statement is configured and the backup Routing Engine synchronizes its configuration with the primary Routing Engine, for example, when it is newly inserted, brought back online, or during a change in primary role, it also synchronizes the ephemeral configuration database.

On the TX Matrix router, synchronization only occurs between the Routing Engines within the same chassis. When synchronization is complete, the new configuration is then distributed to the Routing Engines on the T640 routers. That is, the primary Routing Engine on the TX Matrix router distributes the configuration to the primary Routing Engine on each T640 router. Likewise, the backup Routing Engine on the TX Matrix router distributes the configuration to the backup Routing Engine on each T640 router.

On the TX Matrix Plus router, synchronization only occurs between the Routing Engines within the switch-fabric chassis and when synchronization is complete, the new configuration is then distributed to the Routing Engines on the line-card chassis (LCC). That is, the primary Routing Engine on the TX Matrix Plus router distributes the configuration to the primary Routing Engine on each LCC. Likewise, the backup Routing Engine on the TX Matrix Plus router distributes the configuration to the backup Routing Engine on each LCC.

In EX Series Virtual Chassis configurations:

- On EX4200 switches in Virtual Chassis, synchronization occurs between the switch in the primary role and the switch in the backup role.
- On EX8200 switches in a Virtual Chassis, synchronization occurs only between the primary and backup XRE200 External Routing Engines.

Options

and-quit (Optional) Quit configuration mode if the commit synchronization succeeds.

at (Optional) Time at which to activate configuration changes.

comment (Optional) Write a message to the commit log.

force (Optional) Force a commit synchronization on the other Routing Engine (ignore warnings).

scripts (Optional) Push scripts to the other Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Synchronizing the Routing Engine Configuration

Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically

synchronize (Scripts)

IN THIS SECTION

- Syntax | 28807
- Hierarchy Level | 28807
- Description | 28807
- Required Privilege Level | 28807
- Release Information | 28808

Syntax

```
synchronize;
```

Hierarchy Level

```
[edit system scripts]
```

Description

Synchronize all commit, event, lib, op, and SNMP scripts from the requesting Routing Engine to the responding Routing Engine when executing a `commit synchronize` operation. The device synchronizes all scripts regardless of whether they are enabled in the configuration or have been updated since the last synchronization.

When you issue the `commit synchronize` command, the device performs a commit check on the requesting Routing Engine, synchronizes all scripts to the other Routing Engine, synchronizes, performs a commit check, and commits the configuration on the responding Routing Engine, and finally commits the configuration on the requesting Routing Engine. If the commit check operation fails for the requesting Routing Engine, the process stops, and the scripts are not copied to the responding Routing Engine. If the commit check or commit operation fails for the responding Routing Engine, the scripts are still synchronized since the synchronization occurs prior to the commit check operation on the responding Routing Engine.

If the `load-scripts-from-flash` statement is configured for the requesting Routing Engine, the device synchronizes the scripts from flash memory on the requesting Routing Engine to flash memory on the responding Routing Engine. Otherwise, the device synchronizes the scripts from the hard disk on the requesting Routing Engine to the hard disk on the responding Routing Engine.

NOTE: Configuring the `synchronize` statement causes the device to synchronize all scripts even if they have not been updated since the last synchronization. If the device has a large number of scripts that are infrequently updated, it might be more suitable to synchronize scripts either manually using the `request system scripts synchronize` operational mode command or on a per-commit basis using the `commit synchronize scripts` command.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding Script Synchronization Between Routing Engines

Synchronizing Scripts Between Routing Engines

Example: Synchronizing Scripts Between Routing Engines

commit | **4436**

request system scripts synchronize | **15182**

synchronous-ethernet-mapping

IN THIS SECTION

- Syntax | **28808**
- Hierarchy Level | **28809**
- Description | **28809**
- Required Privilege Level | **28809**
- Release Information | **28809**

Syntax

```
synchronous-ethernet-mapping {
    clock-source ip-address {
        interface interface1-name;
        interface interface2-name;
    }
}
```

Hierarchy Level

```
[edit protocols ptp slave hybrid]
```

Description

Configure the Synchronous Ethernet mapping for hybrid mode.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2R2.

RELATED DOCUMENTATION

- *Configuring Hybrid Mode and ESMC Quality Level Mapping*
- *Example: Configuring Hybrid Mode and ESMC Quality Level Mapping*
- *Understanding Hybrid Mode*
- *Precision Time Protocol Overview*
- *Synchronous Ethernet Overview*

profile-type (PTP)

IN THIS SECTION

- Syntax (ACX Series) | **28810**
- Syntax (MX Series) | **28810**

- Syntax (PTX Series) | 28810
- Syntax (QFX Series) | 28810
- Hierarchy Level | 28810
- Description | 28811
- Options | 28812
- Required Privilege Level | 28812
- Release Information | 28813

Syntax (ACX Series)

```
profile-type (g.8275.1 | g.8275.2 | g.8275.2.enh)
```

Syntax (MX Series)

```
profile-type (g.8275.1 | g.8275.1.enh)
```

Syntax (PTX Series)

```
profile-type (aes67 | aes67-smpte | enterprise-profile | g.8275.1 | g.8275.2 | g.8275.2.enh |
ieee-2008 | ieee1588v2-custom | smpte)
```

Syntax (QFX Series)

```
profile-type (aes67 | aes67-smpte | enterprise-profile | g.8275.2.enh | smpte)
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

On the MX Series, configure the G.8275.1 or the enhanced G.8275.1 PTP profile for applications that require accurate phase and time synchronization. This profile supports the architecture defined in ITU-T G.8275 to enable the distribution of phase and time with full timing support and is based on the second version of PTP defined in IEEE 1588.

> **NOTE**: When the enhanced G.8275.1 PTP profile is configured on either MX240, MX480, MX960, MX2010, MX2020, or MX2008 with MPC2E-NG, MPC3E-NG, MPC5E, or MPC6E line cards with MICS (MIC-3D-20GE-SFP-E), and if PTPoIPv4 BC (Boundary-Clock) stream is configured on the MIC, then the PTPoIPv4 traffic that pass through the PIC gets time stamped one-way and sees a huge correction field update on the pass through traffic. This can cause PTP sync issues on the devices consuming this traffic.

On QFX Series switches that support the enterprise-profile feature, you can configure the enterprise profile, which supports IEEE 1588 PTPv2 transport over multicast IPv4. If you do not specify a profile, the IEEE 1588 default profile is enabled by default.

On QFX Series switches that support the G.8275.2 enhanced profile feature, you can configure the G.8275.2 enhanced profile, which supports telecom applications that require accurate phase and time synchronization for phase alignment and time of day synchronization over a wide area network. This profile supports PTP over IPv4 unicast, ordinary and boundary clocks, and unicast negotiation.

On QFX Series switches that support the media profile, you can configure the SMPTE, AES67, and the AES67+SMPTE profiles to support video applications for capture (for example, cameras), video edit, and playback to be used in professional broadcast environments. The standard allows multiple video sources to stay in synchronization across various equipment by providing time and frequency synchronization to all devices. This profile supports PTP over IPv4 multicast and ordinary and boundary clocks.

ACX 710 devices support G.8275.2 profile with node type T-BC-P (BC).

ACX5448, ACX5448-D, and ACX5448-M devices support ACX G.8275.2 enhanced profile. You can configure the G.8275.2 enhanced profile, which supports telecom applications that require accurate phase and time synchronization for phase alignment and time of day synchronization over a wide area network. This profile supports PTP over IPv4 and IPV6 unicast traffic.

ACX710, ACX2100, ACX2200, ACX7100-48L, and ACX7100-32C devices support G.8275.1 telecome profile. This profile supports the architecture defined in ITU-T G.8275 to enable the distribution of phase and time with full timing support. This profile requires all devices in the network to operate in combined or hybrid modes, which means that PTP and Synchronous Ethernet are enabled on all devices.

PTX10008 device supports G.8275.1 telecome profile. This profile supports the architecture defined in ITU-T G.8275 to enable the distribution of phase and time with full timing support. This profile requires

all devices in the network to operate in combined or hybrid modes, which means that PTP and Synchronous Ethernet are enabled on all devices.

Options

aes67	(QFX Series only) Enable the AES67 PTP profile.
aes67-smpte	(QFX Series only) Enable the AES67+SMPTE PTP profile.
enterprise-profile	(QFX Series only) Enable the enterprise profile. The enterprise profile supports IEEE 1588 PTPv2 transport over multicast IPv4. When the enterprise profile is enabled, no other profiles can be enabled. Also, unicast negotiation is disabled when you enable the enterprise profile.
g.8275.1	(ACX710, ACX2100, ACX2200, ACX7100-48L, ACX7100-32C, MX Series, and PTX10008 only) Enable the G.8275.1 PTP profile.
g.8275.1.enh	(MX Series, PTX series) Enable the enhanced G.8275.1 PTP profile. This profile supports PTP over IPv4.
g.8275.2	(ACX710 only) Enable the G.8275.2 PTP profile.
g.8275.2.enh	(ACX5448, ACX5448-D, ACX5448-M, and QFX Series only) Enable the enhanced G.8275.2 PTP profile.
ieee-2008	(PTX Series) Enable the IEEE 2008 profile type (default).
ieee1588v2-custom	(PTX Series) Enable the IEEE1588v2 with Multicast.
smpte	(QFX Series only) Enable the SMPTE PTP profile.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1R1.

PTP profile-type options added to QFX5130-48C devices in Junos OS Evolved Release 24.2R1

RELATED DOCUMENTATION

Precision Time Protocol Overview

Configuring the Precision Time Protocol Enterprise Profile

Configuring the Precision Time Protocol G.8275.2 Enhanced Profile (Telecom Profile)

syntonized-e2e-transparent

IN THIS SECTION

- Syntax | 28813
- Hierarchy Level | 28813
- Description | 28814
- Options | 28814
- Required Privilege Level | 28814
- Release Information | 28814

Syntax

```
syntonized-e2e-transparent;
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

Configure the syntonized-end-to-end (E2E) transparent clock for Precision Time Protocol (PTP). In syntonized transparent clock, the transparent clock requires physical layer frequency based on the ITU-T G.8262/.1 standard. To set syntonized transparent clock, you can enable Synchronous Ethernet configuration with Ethernet Synchronization Message Channel (ESMC) along with PTP transparent clock configuration. Syntonization supports frequency synchronization but does not support phase or time synchronization.

The transparent clock measures the residence time of PTP packets as the packets pass through the router. Transparent clock adds the residence time into the correction field of the PTP packet. The timeReceiver or boundary clocks determine this resident time while they receive the PTP packet from the upstream transparent clock router. The timeReceiver clock estimates and removes these delays from the offset computation and reduce the packet's jitter effects.

Transparent clock functionality is supported on PTP over Ethernet, IPv4, IPv6, unicast, and multicast. Transparent clock functionality is enabled globally and might be required in scenarios in which the interface and the packets are received and transmitted is unknown.

ACX5048 and ACX5096 routers do not support PTP over IPv6 for transparent clocks.

Options

There are no options.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1 EVO.

RELATED DOCUMENTATION

Understanding Transparent Clocks in Precision Time Protocol

show ptp global-information | **23354**

syslog

IN THIS SECTION

- Syntax | 28815
- Hierarchy Level | 28815
- Description | 28815
- Required Privilege Level | 28815
- Release Information | 28816

Syntax

```
(syslog | no-syslog);
```

Hierarchy Level

```
[edit interfaces mo-fpc/pic/port multiservice-options]
```

Description

System logging is enabled by default. The system log information of the Monitoring Services PIC is passed to the kernel for logging in the **/var/log** directory.

- `syslog`—Enable PIC system logging.
- `no-syslog`—Disable PIC system logging.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Flow Monitoring

syslog (Services IPsec VPN)

IN THIS SECTION

- Syntax | 28816
- Hierarchy Level | 28816
- Description | 28816
- Required Privilege Level | 28817
- Release Information | 28817

Syntax

```
syslog;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Enable system logging. The system log information for the Adaptive Services or Multiservices Physical Interface Card (PIC) is passed to the kernel for logging in the **/var/log** directory.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

syslog (Chassis)

IN THIS SECTION

- Syntax | 28817
- Hierarchy Level | 28818
- Description | 28818
- Options | 28818
- Required Privilege Level | 28818
- Release Information | 28819

Syntax

```
syslog {
    facility {
        severity;
        destination destination;
    }
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-
provider]
```

Description

Enable PIC system logging to record or view system log messages on a specific PIC. The system log information is passed to the kernel for logging in the **/var/log** directory.

Options

facility—Group of messages that are either generated by the same software process or concern a similar condition or activity. Possible values include the following: `daemon`, `external`, `kernel`, and `pfe`.

severity—Classification of effect on functioning. Possible values are the following options:

- `any`—Include all severity levels.
- `none`—Disable logging of the associated facility to a destination.
- `emergency`—System panic or other condition that causes the routing platform to stop functioning.
- `alert`—Conditions that require immediate correction, such as a corrupted system database.
- `critical`—Critical conditions, such as hard errors.
- `error`—Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels.
- `warning`—Conditions that warrant monitoring.
- `notice`—Conditions that are not errors but might warrant special handling.
- `info`—Events or nonerror conditions of interest.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Options `daemon` and `kernel` (for *`facility`*) introduced in Junos OS Release 9.5.

syslog (IDS MS-DPC)

IN THIS SECTION

- Syntax | 28819
- Hierarchy Level | 28819
- Description | 28819
- Required Privilege Level | 28819
- Release Information | 28820

Syntax

```
syslog;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then logging]
```

Description

Enable system logging when using the MS-DPC. The system log information from the MS-DPC is passed to the kernel for logging in the **/var/log** directory.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in IDS Rules

syslog (Interfaces)

IN THIS SECTION

- Syntax | 28820
- Hierarchy Level | 28821
- Description | 28821
- Options | 28821
- Required Privilege Level | 28821
- Release Information | 28821

Syntax

```
syslog {
    host hostname {

 facility-override facility-name;
        log-prefix prefix-number;
        services priority-level;
port
    }
message-rate-limit messages-per-second;
}
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Description

For adaptive services interfaces, configure generation of system log messages for the service set. System log information is passed to the kernel for logging in the `/var/log` directory. Any values configured in the service set definition override these values.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Services Interfaces Library for Routing Devices

syslog-local-system-timestamp

IN THIS SECTION

- Syntax | 28822
- Hierarchy Level | 28822
- Description | 28822

- Default | 28822
- Required Privilege Level | 28822
- Release Information | 28822

Syntax

```
syslog-local-system-timestamp;
```

Hierarchy Level

For routers with MS-MPC, you can configure the statement at the following hierarchy level:

```
[edit interfaces ms-interface|ams-interface services-options
```

Description

Display the system timestamp in the syslog by overriding by overriding the default UTC timestamp to local system timestamp format.

Default

When the statement is not enabled, the syslog displays the system timestamp in UTC format.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1.

RELATED DOCUMENTATION

syslog (Interfaces) | 28820

syslog (Next Gen Services Service-Set System Logging)

IN THIS SECTION

- Syntax | 28823
- Hierarchy Level | 28823
- Description | 28823
- Options | 28823
- Required Privilege Level | 28824
- Release Information | 28824

Syntax

```
syslog ;
```

Hierarchy Level

```
[edit services service-set name]
```

Description

Configure the filename Next Gen Services system logs.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

- *Understanding Next Gen Services CGNAT Global System Logging*
- *Enabling Global System Logging for Next Gen Services*
- *Configuring System Logging to One or More Remote Servers for Next Gen Services*
- *Configuring Local System Logging for Next Gen Services*

syslog (OAM Action)

IN THIS SECTION

- Syntax | 28824
- Hierarchy Level | 28825
- Description | 28825
- Required Privilege Level | 28825
- Release Information | 28825

Syntax

```
syslog;
```

Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile action]
```

Description

Generate a syslog message for the Ethernet Operation, Administration, and Management (OAM) event.

Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Specifying the Actions to Be Taken for Link-Fault Management Events

Configuring Ethernet OAM Link Fault Management

syslog (QFabric System)

IN THIS SECTION

- Syntax | 28826
- Hierarchy Level | 28826

- Description | 28827
- Required Privilege Level | 28827
- Release Information | 28827

Syntax

```
syslog {
    file filename {
        archive {
            size maximum-file-size;
        }
        explicit-priority;
        facility severity;
        match "regular-expression";
        structured-data;
    }
    filter all {
        facility severity;
        match "regular-expression";
    }
    host hostname {
        explicit-priority;
        facility severity;
        facility-override facility;
        log-prefix string;
        match "regular-expression";
        structured-data;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure system log messages for the QFabric system.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Understanding the Implementation of System Log Messages on the QFabric System

Directing System Log Messages to a Remote Machine

syslog (Services CoS)

IN THIS SECTION

- Syntax | 28828
- Hierarchy Level | 28828
- Description | 28828
- Required Privilege Level | 28828
- Release Information | 28828

Syntax

```
syslog;
```

Hierarchy Level

```
[edit services cos rule rule-name term term-name then],
[edit services cos rule rule-name term term-name then reverse]
```

Description

Enable system logging. The system log information from the Multiservices and Services PICs is passed to the kernel for logging in the **/var/log** directory. This setting overrides any `syslog` statement setting included in the service set or interface default configuration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

Configuring Actions in CoS Rules

syslog (Services L2TP)

IN THIS SECTION

- Syntax | 28829
- Hierarchy Level | 28829
- Description | 28829
- Options | 28830
- Required Privilege Level | 28830
- Release Information | 28830

Syntax

```
syslog {
    host hostname {
        services severity-level;
        facility-override facility-name;
        log-prefix prefix-value;
    }
}
```

Hierarchy Level

```
[edit services l2tp tunnel-group group-name]
```

Description

Configure the generation of system log messages for L2TP services. System log information is passed to the kernel for logging in the **/var/log/l2tpd** directory.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring System Logging of L2TP Tunnel Activity

syslog (Services NAT)

IN THIS SECTION

- Syntax | 28830
- Hierarchy Level | 28831
- Description | 28831
- Required Privilege Level | 28831
- Release Information | 28831

Syntax

```
syslog;
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then]
```

Description

Enable system logging. The system log information from the Multiservices PIC is passed to the kernel for logging in the **/var/log** directory.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

syslog (Services Service Set)

IN THIS SECTION

- Syntax | 28832
- Hierarchy Level | 28832
- Description | 28832
- Options | 28833
- Required Privilege Level | 28833
- Release Information | 28833

Syntax

```
syslog {
    host hostname {
        class {
            alg-logs;
            deterministic-nat-configuration-log;
            ids-logs;
            nat-logs;
            packet-logs;
            pcp-logs;
            session-logs <open | close>;
            stateful-firewall-logs;
            urlf-logs;
        }
        facility-override facility-name;
        interface-service prefix-value;
        log-prefix prefix-value;
        port port-number;
        services severity-level;
        source-address source-address;
    }
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Configure generation of system log messages for the service set. The system log information is passed to the kernel for logging in the **/var/log** directory. These settings override the values defined at the `[edit interfaces interface-name services-options]` hierarchy level; for more information on configuring those values, see *Configuring System Logging for Services Interfaces*.

NOTE: Starting with Junos OS release 14.1X55, 14.2R5, 15.1R3, and 16.1R1, for multiservices (ms-) interfaces, you cannot configure system logging for PCP and ALGs by

including the pcp-logs and alg-logs statements at the [edit services service-set service-set-name syslog host hostname class] hierarchy level. An error message is displayed if you attempt to commit a configuration that contains the pcp-logs and alg-logs options to define system logging for PCP and ALGs for ms- interfaces.

Options

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring System Logging for Service Sets

syslog (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 28834
- Hierarchy Level | 28834
- Description | 28834
- Required Privilege Level | 28834
- Release Information | 28834

Syntax

```
syslog;
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name term term-name then]
```

Description

Enable system logging. The system log information from the Adaptive Services or Multiservices PIC is passed to the kernel for logging in the **/var/log** directory. This setting overrides any `syslog` statement setting included in the service set or interface default configuration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in Stateful Firewall Rules

syslog-subtag

IN THIS SECTION

Syntax | 28835

- Hierarchy Level | 28835
- Description | 28835
- Options | 28835
- Required Privilege Level | 28835
- Release Information | 28835

Syntax

```
syslog-subtag syslog-subtag;
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Add a tag to the system log message.

Options

`syslog-subtag` *`syslog-subtag`*—Tag of not more than 80 uppercase characters to be added to syslog messages.

- **Default:** None

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Remote Nework Monitoring (RMON)

RMON MIB Event, Alarm, Log, and History Control Tables

syslog (System)

IN THIS SECTION

- Syntax | 28836
- Hierarchy Level | 28838
- Description | 28838
- Options | 28838
- Required Privilege Level | 28839
- Release Information | 28839

Syntax

```
syslog {
    allow-duplicates;
    alternate-format;
    archive {
        (binary-data| no-binary-data);
        files number;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    console {
        facility severity;
        application-status (any | error);
    }
    file filename {
        allow-duplicates;
        archive {
            (binary-data| no-binary-data);
```

```
            files number;
            size maximum-file-size;
            start-time "YYYY-MM-DD.hh:mm";
            transfer-interval minutes;
            (world-readable | no-world-readable);
        }
        archive-sites [url <password password>];
        facility severity;
        explicit-priority;
        match "regular-expression";
        match-strings [match-strings...];
        structured-data {
            brief;
        }
    }
    grpc-replay;
    host (hostname | other-routing-engine | scc-master) {
        facility severity;
        allow-duplicates;
        exclude-hostname;
        explicit-priority;
        facility-override facility;
        log-prefix string;
        match "regular-expression";
        match-strings [match-strings...];
        routing-instance routing-instance-name;
        source-address source-address;
        structured-data {
            brief;
        }
        port port-number;
        tlsdetails {
            local-certificate;
            mutual-authentication {
                certificate-authority;
            }
            trusted-ca-group [trusted-ca-group-name];
        }
        transport;
    }
    log-rotate-frequency frequency;
    routing-instance routing-instance-name;
    server {
```

```
        routing-instances (routing-instance-name | all | default) {
            disable;
        }
    }
    shell;
    source-address source-address;
    time-format (millisecond | year |  year millisecond);
    user (username | *) {
        allow-duplicates;
        facility severity;
        match "regular-expression";
        match-strings [match-strings...];
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the system log messages types to send to different destinations such as files, remote destinations, user terminals, or the system console.

The remaining statements are explained separately.

Options

archive	Define parameters for archiving log messages.
allow-duplicates	Allow the repeated log messages for all targets.
alternate-format	Configure to append node name with daemon name.
console	Send log messages of a specified class and severity to the console.
file	Send log messages to a named file.
host	Remote location to be notified of specific log messages.
log-rotate-frequency	Configure the interval for checking logfile size and archiving messages.

server	Enable a syslog server for compute nodes and VMs in an App Engine.
shell	Configure to log command executed from the shell.
source-address	Include a specified address as the source address for log messages.
time-format	Additional information to include in the system log time stamp. • millisecond — Include milliseconds in time stamp. • year — Include year in time stamp.
user	Notify a specific user of the log event.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Option `alternate-format` introduced in Junos OS Evolved Release 20.4R2.

Option `shell` introduced in Junos OS Evolved Release 23.1R1.

`application-status` option introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Overview of System Logging

System Log Explorer

system (Analytics Resource)

IN THIS SECTION

- Syntax | 28840
- Hierarchy Level | 28840
- Description | 28840
- Options | 28841
- Required Privilege Level | 28841
- Release Information | 28841

Syntax

```
system {
    polling-interval {
        queue-monitoring interval;
        traffic-monitoring interval;
    }
    resource-profile profile-name;
}
```

Hierarchy Level

```
[edit services analytics resource]
```

Description

Apply a network analytics resource profile to a system for which you wish to enable queue or traffic monitoring. The resource profile is a template that specifies the parameters for queue and traffic monitoring, as well as for the depth and latency thresholds.

> **NOTE**: The `system` statement in the `[edit services analytics resource]` hierarchy is not available in Junos OS Releases prior to 13.2X51-D15.

Options

polling-interval

Configure the polling interval for queue and traffic monitoring:

queue-monitoring *polling-interval*

Configure the queue monitoring interval in milliseconds.

- **Range:** 1 to 1000 milliseconds (1 millisecond to 1 second) on devices other than EX4300 switches.

  8 to 1000 milliseconds (8 milliseconds to 1 second) on EX4300 switches.

traffic-monitoring *polling-interval*

Configure the traffic monitoring interval in seconds.

- **Range:** 1 to 300 seconds (1 second to 5 minutes) on devices other than EX4300 switches.

  5 to 300 seconds (5 seconds to 5 minutes) on EX 4300 switches.

resource-profile *profile-name*

Name of a resource profile containing the global analytics parameters that have been configured for the system. Information contained in a resource profile includes the configuration of queue and traffic monitoring (whether enabled or disabled), and values for the depth and latency thresholds (if applicable).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

system-defaults

IN THIS SECTION

- Syntax | 28842
- Hierarchy Level | 28842
- Description | 28842
- Options | 28843
- Required Privilege Level | 28843
- Release Information | 28843

Syntax

```
system-defaults {
    classifiers exp {
        classifier-name;
    }
    traffic-control-profiles {
        interface-set-input profile-name;
        interface-set-output profile-name;
    }
}
```

Hierarchy Level

```
[edit class-of-service],
[edit dynamic-profiles name class-of-service]
```

Description

Define system defaults for CoS parameters. You can define a global system default EXP classifier, a global default input traffic control profile (TCP), and/or a global default output TCP.

If you configure a global system default EXP classifier, then all interfaces use that EXP classifier to classify MPLS traffic.

On devices that have a default EXP classifier, if you do not configure a global system default classifier, interfaces configured as `family mpls` use the default EXP classifier.

On devices that do not have a default EXP classifier, if you do not configure a global system default EXP classifier, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier are applied to the interface, MPLS traffic is treated as best-effort traffic using the IEEE 802.1 default untrusted classifier. DSCP classifiers are not applied to MPLS traffic. Because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier.

Options

`classifier-name`—Name of the behavior aggregate (BA) `exp` classifier or *default* for the default `exp` classifier.

`profile-name`—For `interface-set-input`, the default input TCP for dynamic interface-sets. For `interface-set-output`, the default output TCP for dynamic interface-sets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

`traffic-control-profiles` option added in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

classifiers (Definition) | 3388

traffic-control-profiles | 30009

system (Export Profiles)

IN THIS SECTION

- Syntax | 28844
- Hierarchy Level | 28844
- Description | 28844
- Options | 28845
- Required Privilege Level | 28845
- Release Information | 28845

Syntax

```
system {
    information;
    status {
        queue;
        traffic;
    }
}
```

Hierarchy Level

```
[edit services analytics export-profiles]
```

Description

Configure a system-wide export profile for streaming network analytics data to remote servers. Each profile is a template that defines the type of data being streamed for that system.

NOTE: The `system` statement is not available in Junos OS Releases prior to 13.2X51-D15.

Options

information Information about the system, including boot time, model, serial number, maximum number of ports, collector information, and interface list.

status System status information to be streamed.

- Values:
  - **queue**
  - **traffic**

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

system-id (Interfaces)

IN THIS SECTION

- Syntax | 28846
- Hierarchy Level | 28846
- Description | 28846
- Required Privilege Level | 28846
- Release Information | 28846

Syntax

```
system-id system-id;
```

Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp]
```

Description

Define the LACP system identifier at the aggregated Ethernet interface level.

The user-defined system identifier in LACP enables two ports from two separate routers (M Series or MX Series routers) to act as though they were part of the same aggregate group.

The system identifier is a 48-bit (6-byte) globally unique field. It is used in combination with a 16-bit system-priority value, which results in a unique LACP system identifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2R1

RELATED DOCUMENTATION

Configuring LACP for Aggregated Ethernet Interfaces

system-id

IN THIS SECTION

- Syntax | 28847
- Hierarchy Level | 28847
- Description | 28847
- Options | 28848
- Required Privilege Level | 28848
- Release Information | 28848

Syntax

```
system-id system-id-value {
    ip-address(es);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp)],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id],
[edit protocols (mstp | rstp)],
[edit protocols vstp vlan vlan-id],
[edit routing-instances routing-instance-name protocols (mstp | rstp)],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id]
```

Description

Determine the system identifier value for bridges in a VPLS multihomed Layer 2 ring with MPLS infrastructure.

Options

system-id-value—System identifier in the format ***nnnnnn.nnnnnn***, where ***n*** = any digit from 0 through 9.

- **Range:** Any valid value
- **Default:** None

ip-address(es)—Valid IP host addresses in the format ***ip-address*/32**.

- **Range:** Any valid IP address
- **Default:** None

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Understanding VPLS Multihoming

system-id (Junos Fusion)

IN THIS SECTION

- Syntax | 28849
- Hierarchy Level | 28849
- Description | 28849
- Default | 28850
- Options | 28850
- Required Privilege Level | 28850

Release Information | 28850

Syntax

```
system-id mac-address;
```

Hierarchy Level

```
[edit chassis satellite-management fpc slot-id]
```

Description

Assign the specified FPC identifier to the satellite device based on the satellite device's MAC address.

For instance, if you wanted the satellite device using MAC address **01:02:03:AA:BB:CC** to be assigned FPC identifier 101, enter the **set chassis satellite-management fpc 101 system-id 01:02:03:AA:BB:CC** statement.

A Junos Fusion provides two methods of assigning an FPC identifier: Unique ID-based FPC identification and connectivity-based FPC identification. Unique ID-based FPC identification maps an FPC slot ID to a satellite device's MAC address or serial number, while connectivity-based FPC identification maps an FPC slot ID to a cascade port. This statement is used to assign an FPC ID using unique ID-based FPC identification by mapping the FPC slot ID to the satellite device's MAC address.

In a Junos Fusion, each satellite device must be mapped to an FPC identifier (FPC ID). The FPC ID is used for Junos Fusion configuration, monitoring, and maintenance. Interface names—which are identified using the *type-fpc* / *pic* / *port* format—use the FPC ID as the *fpc* variable when the satellite device is participating in a Junos Fusion. For instance, built-in port 2—a gigabit Ethernet interface on a satellite device that is using 101 as its FPC ID— uses ge-101/0/2 as its interface name.

For additional information on the role of FPC slot IDs in a Junos Fusion, see *Understanding Junos Fusion Provider Edge Components* or Understanding Junos Fusion Enterprise Components.

If the serial number that is configured using this statement does not match the serial number of the satellite device, the device is not converted into a satellite device.

If the MAC address that is configured using this statement does not match the MAC address of the satellite device, the device is not converted into a satellite device.

Default

No FPC slot IDs are associated with satellite devices, by default.

Options

mac-address Specifies the MAC address of the satellite device.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

system-id (Junos Fusion Satellite Device Cluster)

IN THIS SECTION

- Syntax | 28851
- Hierarchy Level | 28851
- Description | 28851
- Default | 28851
- Options | 28851
- Required Privilege Level | 28851
- Release Information | 28851

Syntax

```
system-id mac-address;
```

Hierarchy Level

```
[edit chassis satellite-management cluster cluster-name fpc slot-id ]
```

Description

Assign the specified FPC identifier to the satellite device in the satellite device cluster based on the satellite device's MAC address.

For instance, if you wanted the satellite device using MAC address **01:02:03:AA:BB:CC** in the satellite device cluster named `building-1` to be assigned FPC identifier 101, enter the **set chassis satellite-management cluster building-1 fpc 101 system-id 01:02:03:AA:BB:CC** statement.

If the MAC address that is configured using this statement does not match the MAC address of the satellite device, the device is not converted into a satellite device.

Default

No FPC slot ID numbers are associated with satellite devices, by default.

Options

mac-address	Specifies the MAC address of the satellite device.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

system login (Junos Node Slicing)

IN THIS SECTION

- Syntax | 28852
- Hierarchy Level | 28852
- Description | 28852
- Options | 28853
- Required Privilege Level | 28853
- Release Information | 28853

Syntax

```
system login {
    user username ;
    class class-name;
    idle-timeout idle-timeout;
}
```

Hierarchy Level

[edit]

Description

Create a non-root user in JDM for Junos node slicing. The non-root user accounts function only inside JDM, not on the host server.

Options

user *username* Username of the account.

class *class-name* Predefined login classes that JDM supports for non-root users.

- super-user
- operator
- read-only
- unauthorized

idle-timeout Maximum idle time before logout. We recommend that you set an idle time-out. Otherwise, if a user forgets to log out of a VNF console session, another user can log in to it without providing the access credentials.

- **Range:** 1 through 60 minutes

Required Privilege Level

Release Information

Statement introduced in Junos OS Release 18.3 R1.

RELATED DOCUMENTATION

Setting Up Junos Node Slicing

Setting Up Junos Node Slicing

system (Memory)

IN THIS SECTION

Syntax | 28854

- Hierarchy Level | 28854
- Description | 28854
- Options | 28855
- Required Privilege Level | 28855
- Release Information | 28855

Syntax

```
system {
    frequency (minutes | hours | days) value;
    minor-event {
        threshold swap-usage value ;
        action snmp-trap;
    }
    major-event {
        threshold swap-usage value ;
        action [snmp-trap | live-vmcore];
    }
    critical-event {
        threshold swap-usage value ;
        action [snmp-trap | live-vmcore | (RE-restart | RE-switchover)];
    }
}
```

Hierarchy Level

```
[edit system monitor memory]
```

Description

Configure system level memory monitoring.

Options

frequency Specify the frequency of system level memory monitoring. You can specify the frequency in days, hours or minutes. Possible values: 1 through 65535.

minor-event Indicates a non-critical event that, if left unchecked, may cause an interruption in service or degradation in performance but does not significantly affect the device.

major-event Indicates an event that, if left unchecked, may result in significant loss of service or interruption of traffic.

critical-event Indicates an event that, if left unchecked, may result in serious loss of service or interruption of traffic. In case a critical event is triggered, you can either restart the Routing Engine or configure the primary Routing Engine to switchover gracefully to a backup Routing Engine without interruption in packet forwarding.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1.

RELATED DOCUMENTATION

System Log Overview

System Log Explorer

system-priority (chassis)

IN THIS SECTION

Syntax | 28856

- Hierarchy Level | 28856
- Description | 28856
- Options | 28856
- Required Privilege Level | 28857
- Release Information | 28857

Syntax

```
system-priority priority;
```

Hierarchy Level

```
[edit chassis aggregated-devices ethernet lacp]
```

Description

Define LACP system priority for aggregated Ethernet interfaces at the global (chassis) level.

The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored. In priority comparisons, numerically lower values have higher priority. Therefore, the system with the numerically lower value (higher priority value) for LACP system priority becomes the controlling system. If both devices have the same LACP system priority (for example, they are both configured with the default setting of 127), the device MAC address determines which switch is in control.

Options

priority—Priority for the aggregated Ethernet system. A smaller value indicates a higher priority.

- **Range:** 0 through 65535
- **Default:** 127

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Junos OS for Supporting Aggregated Devices

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

system-priority (interfaces)

IN THIS SECTION

- Syntax | 28857
- Hierarchy Level | 28858
- Description | 28858
- Options | 28858
- Required Privilege Level | 28858
- Release Information | 28858

Syntax

```
system-priority priority;
```

Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp]
```

Description

Define LACP system priority at the aggregated Ethernet interface level. This system priority value takes precedence over a system priority value configured at the global `[edit chassis]` hierarchy level.

The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored. In priority comparisons, numerically lower values have higher priority. Therefore, the system with the numerically lower value (higher priority value) for LACP system priority becomes the controlling system. If both devices have the same LACP system priority (for example, they are both configured with the default setting of 127), the device MAC address determines which switch is in control.

Options

priority—Priority for the aggregated Ethernet system. A smaller value indicates a higher priority.

- **Range:** 0 through 65535
- **Default:** 127

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

system-services (Security Zones Host Inbound Traffic)

IN THIS SECTION

- Syntax | 28859
- Hierarchy Level | 28859
- Description | 28859
- Options | 28860
- Required Privilege Level | 28861
- Release Information | 28861

Syntax

```
system-services {
    (service-name | all <service-name except>);
}
```

Hierarchy Level

```
[edit security zones security-zone zone-name host-inbound-traffic]
```

Description

Specify the types of incoming system service traffic that can reach the device for all interfaces in a security zone. By default, a security zone has all system services disabled. You can allow the inbound system services traffic in one of the following ways:

- Allow system services individually.
- Allow all system services.
- Allow all system services with the exception of the specified services.

Options

service-name— Name of system service traffic that can reach the device.

- `all`—Traffic from the defined system services available on the Routing Engine (RE). Use the *except* option to disallow specific system services. Enabling all the system services does not override any interface-specific configuration under a particular zone.
- `any-service`—All system services on an entire port range including the system services that are not defined.
- `bootp`—Traffic destined to BOOTP and DHCP relay agents
- `dhcp`—DHCP requests
- `dhcpv6`—DHCP requests for IPv6
- `dns`—DNS services
- `finger`—Finger traffic
- `ftp`—FTP traffic
- `http`—J-Web or clear-text Web authentication traffic
- `https`—J-Web or Web authentication traffic over Secure Sockets Layer (SSL)
- `ident-reset`—Access that has been blocked by an unacknowledged identification request
- `ike`—Internet Key Exchange (IKE) traffic
- `lsping`—Label-switched path (LSP) ping service
- `netconf`—NETCONF service
- `ntp`—Network Time Protocol (NTP) traffic
- `ping`—ICMP echo request responses
- `r2cp`—Radio-to-Router Control Protocol traffic
- `reverse-ssh`—Reverse SSH traffic
- `reverse-telnet`—Reverse Telnet traffic
- `rlogin`—Incoming `rlogin` (remote login) traffic
- `rpm`—Real-time performance monitoring (RPM) traffic

- `rsh`—Remote shell (`rsh`) traffic
- `snmp`—SNMP traffic (UDP port 161)
- `snmp-trap`—SNMP traps (UDP port 162)
- `ssh`—SSH traffic
- `telnet`—Telnet traffic
- `tftp`—TFTP services
- `traceroute`—Traceroute traffic (UDP port 33434)
- `xnm-clear-text`—Junos XML protocol traffic for all specified interfaces
- `xnm-ssl`— Junos XML protocol-over-SSL traffic for all specified interfaces

service-name **except—** (Optional) Allow all inbound service traffic, except the specified service traffic types, to reach the device. In the following example, the configuration allows all system service traffic, with the exception of FTP and HTTP, to reach the device:

```
[edit]
user@host# set security zones security-zone trust host-inbound-traffic system-
services all
user@host# set security zones security-zone trust host-inbound-traffic system-
services ftp except
user@host# set security zones security-zone trust host-inbound-traffic system-
services http except
```

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Zones Overview

Supported System Services for Host Inbound Traffic

system-services (Security Zones Interfaces)

IN THIS SECTION

- Syntax | 28862
- Hierarchy Level | 28862
- Description | 28862
- Options | 28863
- Required Privilege Level | 28864
- Release Information | 28864

Syntax

```
system-services service-name {
               except;
}
```

Hierarchy Level

```
[edit security zones security-zone zone-name interfaces interface-name host-inbound-traffic]
```

Description

Specify the types of traffic that can reach the device on a particular interface.

Options

- *service-name* —Service for which traffic is allowed. The following services are supported:
  - `all`—Enable all possible system services available on the Routing Engine (RE).
  - `any-service`—Enable services on entire port range.
  - `bootp`—Enable traffic destined to BOOTP and DHCP relay agents.
  - `dhcp`—Enable incoming DHCP requests.
  - `dhcpv6`—Enable incoming DHCP requests for IPv6.
  - `dns`—Enable incoming DNS services.
  - `finger`—Enable incoming finger traffic.
  - `ftp`—Enable incoming FTP traffic.
  - `http`—Enable incoming J-Web or clear-text Web authentication traffic.
  - `https`—Enable incoming J-Web or Web authentication traffic over Secure Sockets Layer (SSL).
  - `ident-reset`—Enable the access that has been blocked by an unacknowledged identification request.
  - `ike`—Enable Internet Key Exchange traffic.
  - `netconf SSH`—Enable incoming NetScreen Security Manager (NSM) traffic over SSH.
  - `ntp`—Enable incoming Network Time Protocol (NTP) traffic.
  - `ping`—Allow the device to respond to ICMP echo requests.
  - `r2cp`—Enable incoming Radio Router Control Protocol traffic.
  - `reverse-ssh`—Reverse SSH traffic.
  - `reverse-telnet`—Reverse Telnet traffic.
  - `rlogin`—Enable incoming `rlogin` (remote login) traffic.
  - `rpm`—Enable incoming real-time performance monitoring (RPM) traffic.
  - `rsh`—Enable incoming Remote Shell (`rsh`) traffic.
  - `snmp`—Enable incoming SNMP traffic (UDP port 161).
  - `snmp-trap`—Enable incoming SNMP traps (UDP port 162).

- `ssh`—Enable incoming SSH traffic.
- `telnet`—Enable incoming Telnet traffic.
- `tftp`—Enable TFTP services.
- `traceroute`—Enable incoming traceroute traffic (UDP port 33434).
- `xnm-clear-text`—Enable incoming Junos XML protocol traffic for all specified interfaces.
- `xnm-ssl`— Enable incoming Junos XML protocol-over-SSL traffic for all specified interfaces.

- `except`—(Optional) except can only be used if all has been defined.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Zones Overview

Supported System Services for Host Inbound Traffic

T

IN THIS SECTION

- t1-interval | **28885**
- t1-options | **28886**
- t1-options (BITS Interfaces Signal Type) | **28887**
- t310 | **28889**

- t391 | **28890**
- t392 | **28892**
- t4-interval | **28893**
- table | **28895**
- tacplus | **28896**
- tacplus-options | **28898**
- tacplus-server | **28902**
- tag | **28904**
- tag (Access) | **28905**
- tag-attribute (HTTP Header Enrichment) | **28907**
- tag-attribute (HTTP Header Enrichment Tag Rule) | **28908**
- tag (Configuring Notification Targets) | **28910**
- tag-group | **28911**
- tag-header (HTTP Header Enrichment) | **28913**
- tag (HTTP Header Enrichment) | **28914**
- tag-list | **28916**
- tag-operation (HTTP Header Enrichment) | **28917**
- tag-protocol-id (Dynamic VLANs) | **28919**
- tag-protocol-id (TPID to Rewrite) | **28920**
- tag-protocol-id TPID(TPIDs Expected to Be Sent or Received) | **28922**
- tag (Routing Options) | **28924**
- tag-rule (HTTP Header Enrichment) | **28926**
- tag-rule (Profiles for HTTP Header Enrichment) | **28928**
- tag-rule-set (HTTP Header Enrichment) | **28930**
- tag-rule-sets (Service Set) | **28931**
- tag-rules (Service Set) | **28932**
- tag-separator (HTTP Header Enrichment) | **28934**
- tag-value (HTTP Header Enrichment) | **28935**
- tag2 (Dynamic Access Routes) | **28937**
- tags (Application Identification) | **28938**
- talk | **28940**
- target | **28941**

- target-address (TWAMP) | **28944**
- target-address (SNMP) | **28946**
- target-attachment-identifier (Protocols VPWS) | **28947**
- target-logical-system (Domain Map) | **28949**
- target-parameters | **28951**
- target-routing-instance (Domain Map) | **28952**
- target (Routing Instances MVPN) | **28954**
- targeted-broadcast | **28956**
- targeted-distribution | **28958**
- targeted-distribution (Dynamic Demux Interfaces over Aggregated Ethernet) | **28959**
- targeted-distribution (Static Interfaces over Aggregated Ethernet) | **28961**
- targeted-hello | **28962**
- targeted-options (Grouping Subscribers by Bandwidth Usage) | **28964**
- targeted-options (Manual Targeting) | **28966**
- targeted-options (PS Interface) | **28968**
- targets | **28970**
- targets (Unified Edge) | **28971**
- tca | **28973**
- tcp | **28977**
- tcp-aggressive-transmission | **28978**
- tcp-encap | **28980**
- tcp-fast-open | **28982**
- tcp-flags | **28983**
- tcp-forwarding (Processes) | **28985**
- tcp-forwarding (Remote Device Management) | **28987**
- tcp-initial-timeout | **28989**
- tcp-keepcnt | **28991**
- tcp-keepidle | **28992**
- tcp-keepintvl | **28994**
- tcp-log | **28995**
- tcp-mss (Internet Options) | **28996**
- tcp-mss (Protocols BGP) | **28998**

- tcp-mss (Security Flow) | **29000**
- tcp-mss (Services) | **29002**
- tcp-no-flag | **29004**
- tcp-no-flag (IDS Screen Next Gen Services) | **29005**
- tcp-non-syn | **29006**
- tcp-options (Security Policies) | **29008**
- tcp-rst | **29012**
- tcp (Security IDP Protocol Binding) | **29013**
- tcp (Security IDP Signature Attack) | **29014**
- tcp (Security Screen) | **29017**
- tcp-seq-ooo-window (Security Forwarding Options) | **29019**
- tcp-session | **29020**
- tcp-session-install-interval (Security Forwarding Options) | **29022**
- tcp-session (Service Set Next Gen Services) | **29023**
- tcp-sweep | **29025**
- tcp-syn-defense (IDS MS-MPC) | **29027**
- tcp-syn-fragment-check (IDS MS-MPC) | **29028**
- tcp-tickles | **29029**
- tcp-tickles (Service Set Next Gen Services) | **29030**
- tcp-winnuke-check (IDS MS-MPC) | **29031**
- tcp-mss (Interfaces) | **29033**
- tcpao-auth-mismatch | **29034**
- tdf-interface | **29036**
- tdf (Unified Edge) | **29037**
- tdm-options (Interfaces) | **29039**
- te-class-matrix | **29041**
- te-link | **29043**
- te-policy | **29044**
- tear-drop (IDS Screen Next Gen Services) | **29046**
- tei-option | **29047**
- telemetries | **29048**
- telemetries (PoE) | **29050**

- telemetry | **29051**
- telemetry | **29054**
- telnet | **29056**
- telnet (Access) | **29059**
- telnet (Services) | **29060**
- temperature-sensor | **29063**
- template | **29071**
- template (Services URL Filter) | **29072**
- template (Flow Monitoring IPFIX Version) | **29075**
- template (Flow Monitoring Version 9) | **29077**
- template (Forwarding Options) | **29079**
- template (Forwarding Options Version IPFIX) | **29080**
- template-id | **29081**
- template (Inline Monitoring) | **29084**
- template (LRF Profile) | **29087**
- template (LRF Rule) | **29089**
- template-profile (Flow Monitoring Logs for NAT) | **29090**
- template(Protocols MVPN) | **29092**
- template-refresh-rate | **29094**
- template-refresh-rate (Services) | **29096**
- template-tx-interval (LRF Profile) | **29097**
- template-type (Flow Monitoring Logs for NAT) | **29098**
- template-type (LRF Profile) | **29100**
- templates | **29102**
- term (Application Aware Access List) | **29106**
- term (IP Security VPN) | **29107**
- term (Applications) | **29110**
- term (candidate-uplink-policy) | **29111**
- term (Captive Portal Content Delivery) | **29113**
- term (Dynamic Profiles) | **29115**
- term (Firewall Filter) | **29117**
- term (HTTP Header Enrichment) | **29120**

- term (IDS MS-DPC) | **29123**
- term (IDS MS-MPC) | **29125**
- term (PCP) | **29128**
- term (satellite-policies) | **29130**
- term (Services CoS) | **29132**
- term (Services NAT) | **29134**
- term (Services Stateful Firewall) | **29136**
- term (Simple Filter) | **29137**
- term (TDF Domain Selection) | **29139**
- term (URL Filter) | **29143**
- terminal | **29146**
- terminate-code | **29147**
- terminate (PPPoE Service Name Tables) | **29149**
- termination (Services) | **29151**
- test | **29152**
- test aaa authd-lite user | **29156**
- test aaa dhcp user | **29161**
- test aaa ppp user | **29169**
- test access profile | **29178**
- test access radius-server | **29184**
- test-agent (Paragon Active Assurance) | **29187**
- test configuration | **29190**
- test-count | **29192**
- test-finish-wait-duration (RFC 2544 Benchmarking) | **29193**
- test interface restart-auto-negotiation | **29195**
- test-interface (RFC 2544 Benchmarking) | **29196**
- test-interval | **29198**
- test-iterator-duration (RFC 2544 Benchmarking) | **29200**
- test-iterator-pass-threshold (RFC 2544 Benchmarking) | **29201**
- test msdp | **29202**
- test-name | **29204**
- test-only-mode | **29208**

- test policy | **29209**
- test-profile (RFC 2544 Benchmarking) | **29211**
- test (Security IDP) | **29214**
- test security utm anti-spam | **29215**
- test security utm enhanced-web-filtering url-check | **29219**
- test security utm web-filtering profile | **29224**
- test services l2tp tunnel | **29227**
- test services monitoring rfc2544 | **29230**
- test services rpm rfc2544-benchmarking test | **29234**
- test-session (Junos OS) | **29237**
- test-session (Junos OS Evolved) | **29239**
- test-type (RFC 2544 Benchmarking) | **29244**
- tests | **29247**
- tftp | **29251**
- tftp (Security ALG) | **29252**
- tftp-server | **29254**
- then (IP Security VPN) | **29255**
- then (Hierarchical Policer) | **29257**
- then (Captive Portal Content Delivery) | **29259**
- then (Event Policy) | **29261**
- then (Filters) | **29264**
- then (Firewall Filters) | **29266**
- then (HTTP Header Enrichment) | **29267**
- then (IDS MS-DPC) | **29269**
- then (IDS MS-MPC) | **29271**
- then (LRF rule) | **29274**
- then (PCC Rules) | **29276**
- then (PCP) | **29278**
- then (Policer Action) | **29279**
- then (Policers) | **29281**
- then (Security Application Firewall) | **29282**
- then (Security Destination NAT) | **29284**

- then (Security IDP Policy) | **29285**
- then (Security Policies) | **29287**
- then (Security Source NAT) | **29292**
- then (Security Static NAT) | **29293**
- then (Services CoS) | **29295**
- then (Services CoS Next Gen Services) | **29296**
- then (Services NAT) | **29299**
- then (Services Stateful Firewall) | **29300**
- then (Stateful Firewall Rule Next Gen Services) | **29302**
- then (TDF Domain Selection) | **29303**
- thermal-health-check | **29305**
- three-color-policer | **29309**
- three-color-policer (Applying) | **29311**
- three-color-policer (Configuring) (EX) | **29313**
- three-color-policer (Configuring) | **29314**
- threshold | **29317**
- threshold (Bridge Domains) | **29319**
- threshold-count | **29320**
- threshold (Detection Time) | **29322**
- threshold (detection-time) | **29323**
- threshold (Interfaces) | **29325**
- threshold-logging-interval | **29326**
- threshold (MSDP Active Source Messages) | **29327**
- threshold (Multicast Forwarding Cache) | **29329**
- threshold (PIM BFD Detection Time) | **29331**
- threshold (PIM BFD Transmit Interval) | **29333**
- threshold (PIM Entries) | **29334**
- threshold-rate | **29336**
- threshold (Routing Instances) | **29338**
- threshold (Services IPsec) | **29340**
- threshold (Services Logging and SYN-Cookie Defenses) | **29341**
- threshold-time | **29343**

- threshold (transmit-interval) | **29344**
- thresholds | **29346**
- thresholds (Junos OS) | **29348**
- thresholds (Junos OS Evolved) | **29350**
- time | **29352**
- time-binding | **29354**
- time-format | **29356**
- time-interval (Dampen Options) | **29358**
- time-limit (LRF Rule) | **29359**
- time-wait-state | **29361**
- time-zone | **29362**
- timeout (SNMP) | **29365**
- timeout | **29367**
- timeout (Access LDAP) | **29368**
- timeout (Access RADIUS) | **29370**
- timeout-action | **29371**
- timeout-active-flows (DDoS Flow Detection) | **29372**
- timeout (DHCP Local Server) | **29374**
- timeout (Diameter Network Element) | **29376**
- timeout (Flow Maps) | **29377**
- timeout-grace (Access) | **29379**
- timeout (Multicast) | **29380**
- timeout (RADIUS Server) | **29382**
- timeout (Security Antivirus Fallback Options) | **29383**
- timeout (Security Antivirus Fallback Options Sophos Engine) | **29385**
- timeout (Security Antivirus Scan Options) | **29386**
- timeout (Security GTP) | **29388**
- timeout (Security Screen) | **29389**
- timeout (Security Web Filtering) | **29391**
- timeout (Security Web Filtering Cache) | **29392**
- timeout (Security Web Filtering Fallback Settings) | **29393**
- timeout (Services) | **29395**

- timestamp | **29396**
- timestamp-format (RFC 2544 Benchmarking) | **29398**
- timestamp-option (IDS Screen Next Gen Services) | **29399**
- tls | **29400**
- tls-min-version | **29402**
- tls (NETCONF) | **29403**
- tls-peer-name | **29405**
- tls-timeout | **29406**
- tls-type | **29407**
- tlsdetails | **29409**
- tlv-filter | **29410**
- tlv-select | **29414**
- tm | **29417**
- tm-polarity | **29418**
- to | **29419**
- to (Security Source NAT) | **29421**
- to-zone (Security IDP Policy) | **29422**
- to-zone (Security Policies) | **29423**
- to-zone (Security Policies Global) | **29427**
- token (DHCP Local Server) | **29428**
- too-many-requests (Security Antivirus Fallback Options) | **29430**
- too-many-requests (Security Antivirus Fallback Options Sophos Engine) | **29432**
- too-many-requests (Security Web Filtering Fallback Settings) | **29433**
- top | **29435**
- topologies (Multitopology Routing) | **29436**
- topologies (Protocols IS-IS) | **29438**
- topology (Filter-Based Forwarding) | **29440**
- topology-id | **29441**
- topology (Multitopology Routing) | **29443**
- topology (OSPF) | **29445**
- topology (OSPF Interface) | **29446**
- topology (Protocols BGP) | **29449**

- tos-reflect (L2TP) | **29451**
- total-memory | **29452**
- trace | **29454**
- trace (DHCP Local Server) | **29456**
- trace (DHCP Relay Agent) | **29458**
- trace (EVPN-VXLAN) | **29459**
- traceoptions (routing-options) | **29462**
- traceoptions (Ethernet Switching) | **29465**
- traceoptions (High Availability) | **29471**
- traceoptions (802.1X) | **29474**
- traceoptions (Access) | **29477**
- traceoptions (Access Port Security) | **29479**
- traceoptions (Active Directory Access) | **29483**
- traceoptions (Address-Assignment Pool) | **29485**
- traceoptions (advanced policy-based routing) | **29488**
- traceoptions (Analytics) | **29491**
- traceoptions (Analytics Agent) | **29492**
- traceoptions (ANCP) | **29494**
- traceoptions (App Engine Virtual Machine Management Service) | **29497**
- traceoptions (Batch Commits) | **29499**
- traceoptions (Captive Portal Content Delivery) | **29501**
- traceoptions (CASB) | **29503**
- traceoptions (Chassis Cluster) | **29506**
- traceoptions (Class of Service) | **29508**
- traceoptions (Commit, Op, and SNMP Scripts) | **29511**
- traceoptions (DDoS) | **29515**
- traceoptions (DHCP) | **29517**
- traceoptions (DHCP Server) | **29521**
- traceoptions (Diameter Base Protocol for Subscriber Aware Policy Control) | **29524**
- traceoptions (Diameter Base Protocol) | **29527**
- traceoptions (DNS, Port, and TFTP Packet Forwarding) | **29530**
- traceoptions (dynamic-address) | **29533**

- traceoptions (dynamic-application) | **29535**
- traceoptions (Dynamic Flow Capture) | **29538**
- traceoptions (Egress Protection) | **29539**
- traceoptions (Enhanced Subscriber Management) | **29541**
- traceoptions (Event Options) | **29544**
- traceoptions (Event Scripts) | **29547**
- traceoptions (Extensible Subscriber Services Manager) | **29550**
- traceoptions (Extensions) | **29552**
- traceoptions (FC-2 Fibre Channel) | **29554**
- traceoptions (Fibre Channel) | **29557**
- traceoptions (FIP Protocol Fibre Channel) | **29560**
- traceoptions (flow) | **29563**
- traceoptions (Forwarding Options) | **29565**
- traceoptions (General Authentication Service) | **29567**
- traceoptions (Health Check Monitoring) | **29569**
- traceoptions (IGMP Snooping) | **29573**
- traceoptions (Individual Interfaces) | **29576**
- traceoptions (Inline Monitoring) | **29585**
- traceoptions (Interfaces) | **29587**
- traceoptions (Juniper Secure Connect) | **29589**
- traceoptions (LACP) | **29591**
- traceoptions (Layer 2 Learning) | **29593**
- traceoptions (LLDP) | **29596**
- traceoptions (MACsec) | **29599**
- traceoptions (MACsec interfaces) | **29602**
- traceoptions (Multicast Snooping Options) | **29604**
- traceoptions (MVRP) | **29607**
- traceoptions (MX Series Virtual Chassis) | **29609**
- traceoptions (NETCONF and Junos XML Protocol) | **29613**
- traceoptions (NETCONF TLS) | **29616**
- traceoptions (Next Gen Services Global System Logging) | **29619**
- traceoptions (Next Gen Services Service-Set Flow) | **29620**

- traceoptions (Next Gen Services Softwires) | **29624**
- traceoptions (Origin Validation for BGP) | **29626**
- traceoptions (Outbound SSH) | **29627**
- traceoptions (OVSDB) | **29630**
- traceoptions (PCE) | **29633**
- traceoptions (PCEF) | **29635**
- traceoptions (PIM Snooping) | **29638**
- traceoptions (Port Mirroring and Traffic Sampling) | **29640**
- traceoptions (PPP Process) | **29641**
- traceoptions (PPPoE) | **29645**
- traceoptions (Process Monitor) | **29648**
- traceoptions (programmable-rpd) | **29651**
- traceoptions (Protocols) | **29654**
- traceoptions (Protocols AMT) | **29656**
- traceoptions (Protocols BFD) | **29659**
- traceoptions (Protocols BGP) | **29662**
- traceoptions (Protocols BMP) | **29666**
- traceoptions (Protocols DVMRP) | **29669**
- traceoptions (Protocols EVPN) | **29672**
- traceoptions (Protocols IGMP) | **29675**
- traceoptions (Protocols IGMP Snooping) | **29678**
- traceoptions (Protocols IPv6 Neighbor Discovery) | **29681**
- traceoptions (Protocols IS-IS) | **29684**
- traceoptions (Protocols Layer 2 Circuit) | **29688**
- traceoptions (Protocols Layer 2 VPN) | **29690**
- traceoptions (Protocols LDP) | **29693**
- traceoptions (Protocols Link Management) | **29697**
- traceoptions (Protocols MPLS) | **29699**
- traceoptions (Protocols MSDP) | **29703**
- traceoptions (Protocols MVPN) | **29706**
- traceoptions (Protocols OpenFlow) | **29710**
- traceoptions (Protocols OSPF) | **29713**

- traceoptions (Protocols PCEP) | **29717**
- traceoptions (Protocols PIM) | **29719**
- traceoptions (Protocols PPP Service) | **29723**
- traceoptions (Protocols RIP) | **29727**
- traceoptions (Protocols RIPng) | **29730**
- traceoptions (Protocols Router Discovery) | **29733**
- traceoptions (Protocols RSVP) | **29736**
- traceoptions (Protocols Secure Neighbor Discovery) | **29739**
- traceoptions (Protocols Spring-TE) | **29742**
- traceoptions (Protocols VPLS) | **29744**
- traceoptions (Protocols VRRP) | **29747**
- traceoptions (Proxy Fibre Channel) | **29750**
- traceoptions (Resiliency) | **29752**
- traceoptions (Resource Cleanup) | **29755**
- traceoptions (Resource Monitor) | **29758**
- traceoptions (REST API) | **29760**
- traceoptions (RPM) | **29761**
- traceoptions (SBC Configuration Process) | **29764**
- traceoptions (Security) | **29767**
- traceoptions (Security ALG) | **29770**
- traceoptions (Security Antispam) | **29773**
- traceoptions (Security Antivirus) | **29774**
- traceoptions (Security Application Firewall) | **29776**
- traceoptions (Security Application Proxy) | **29778**
- traceoptions (Security Content Filtering) | **29781**
- traceoptions (Security Content Security) | **29782**
- traceoptions (Security Datapath Debug) | **29783**
- traceoptions (Security Dynamic VPN) | **29786**
- traceoptions (Security Firewall Authentication) | **29788**
- traceoptions (Security Flow) | **29789**
- traceoptions (Security Group VPN) | **29795**
- traceoptions (Security GTP) | **29798**

- traceoptions (Security H323 ALG) | **29801**
- traceoptions (Security IDP) | **29803**
- traceoptions (Security IKE) | **29806**
- traceoptions (Security IPsec) | **29809**
- traceoptions (Security Log) | **29811**
- traceoptions (Security MGCP ALG) | **29814**
- traceoptions (Security NAT) | **29815**
- traceoptions (Security PKI) | **29818**
- traceoptions (Security Policies) | **29821**
- traceoptions (Security SCCP ALG) | **29823**
- traceoptions (Security Screen) | **29825**
- traceoptions (Security SCTP) | **29828**
- traceoptions (Security SIP ALG) | **29830**
- traceoptions (Security User Identification) | **29832**
- traceoptions (Security Web Filtering) | **29835**
- traceoptions (Services) | **29836**
- traceoptions (Services Application Identification) | **29839**
- traceoptions (Services IPsec VPN) | **29842**
- traceoptions (Services L2TP) | **29844**
- traceoptions (Services Logging) | **29849**
- traceoptions (Services Redundancy Daemon) | **29851**
- traceoptions (Services SSL) | **29854**
- traceoptions (Services UAC) | **29857**
- traceoptions (Services User Identification) | **29859**
- traceoptions (sFlow Technology) | **29862**
- traceoptions (SMTP) | **29865**
- traceoptions (SNMP) | **29866**
- traceoptions (Spanning Tree) | **29869**
- traceoptions (Static Subscribers) | **29873**
- traceoptions (Subscriber Management) | **29876**
- traceoptions (Subscriber Session Database Replication) | **29878**
- traceoptions (System Accounting) | **29881**

- traceoptions (System License) | **29883**
- traceoptions (System Monitor) | **29885**
- traceoptions (System Services DNS) | **29888**
- traceoptions (TCP Encapsulation) | **29891**
- traceoptions (TCP Port Forwarding) | **29894**
- traceoptions (TDF Gateway) | **29896**
- traceoptions (Traffic Load Balancer) | **29899**
- traceoptions (Uplink Failure Detection) | **29903**
- traceoptions (Virtual Chassis) | **29905**
- traceroute | **29909**
- traceroute clns | **29917**
- traceroute ethernet | **29919**
- traceroute fabric unicast-flow | **29924**
- traceroute monitor | **29927**
- traceroute mpls bgp | **29930**
- traceroute mpls ldp | **29935**
- traceroute mpls rsvp | **29940**
- traceroute mpls segment-routing isis | **29946**
- traceroute mpls segment-routing ospf | **29953**
- traceroute mpls segment-routing spring-te | **29959**
- traceroute mpls segment-routing static egress-ip | **29970**
- traceroute overlay | **29976**
- traceroute srv6 | **29982**
- tracing | **29987**
- track-igp-metric | **29989**
- track-igp-metric (LSP) | **29990**
- track (VRRP) | **29992**
- tracking-area | **29994**
- traffic-class (delay) | **29996**
- traffic-class (loss) | **29998**
- traffic-class (loss-delay) | **30000**
- traffic-class-map | **30003**

- traffic-class-map (Apply to Interface) | **30006**
- traffic-class (Tunnels) | **30008**
- traffic-control-profiles | **30009**
- traffic-control-profiles (Dynamic CoS Definition) | **30012**
- traffic-engineering | **30014**
- traffic-engineering (policy-options) | **30016**
- traffic-engineering (OSPF) | **30018**
- traffic-engineering (Passive TE Mode) | **30021**
- traffic-engineering (Protocols BGP) | **30023**
- traffic-engineering (Protocols IS-IS) | **30025**
- traffic-engineering (Protocols MPLS) | **30031**
- traffic-engineering (SR-TE import policy) | **30032**
- traffic-load-balance (Traffic Load Balancer) | **30034**
- traffic-manager | **30036**
- traffic-options | **30040**
- traffic-selector | **30042**
- traffic-statistics (Services Analytics) | **30045**
- traffic-statistics (Protocols IS-IS) | **30048**
- traffic-statistics-labeled-path | **30049**
- traffic-statistics (Protocols BGP) | **30051**
- traffic-statistics (Protocols LDP) | **30053**
- transaction-timeout | **30055**
- transfer | **30056**
- transfer-delay (Event Policy) | **30058**
- transfer-interval | **30059**
- transfer-log-archive | **30061**
- transient-filter | **30062**
- transit | **30064**
- transit (Chained Composite Next Hops) | **30066**
- transit-lsp-association | **30069**
- translated | **30070**
- translated-port | **30072**

- translation | **30073**
- translation-table | **30075**
- translation-type | **30077**
- translation-vni | **30080**
- transmit-clock | **30083**
- transmit-failure-threshold (RFC 2544 Benchmarking) | **30084**
- transmit-interval | **30085**
- transmit-interval (BFD Liveness Detection) | **30087**
- transmit-interval (LAG) | **30090**
- transmit-interval (Liveness Detection) | **30092**
- transmit-interval (MACsec) | **30093**
- transmit-interval (MACsec for MX Series) | **30095**
- transmit-interval (PIM BFD Liveness Detection) | **30096**
- transmit-rate | **30098**
- transmit-rate (Dynamic Schedulers) | **30103**
- transmit-rate (EX Series Switches) | **30105**
- transmit-rate (Schedulers) | **30107**
- transparent | **30110**
- transport (Services) | **30111**
- transport (System Syslog) | **30112**
- transport (Protocols BGP) | **30113**
- transport-address | **30118**
- transport-class | **30120**
- transport (Diameter Base Protocol) | **30123**
- transport (Diameter Peer) | **30124**
- transport-monitoring | **30126**
- transport (Next Gen Services Syslog Message Security) | **30127**
- transport (Security Log) | **30129**
- transport-type | **30131**
- transport-class | **30132**
- trap | **30133**
- trap-group | **30135**

- trap-options | **30136**
- trap (Security IKE) | **30139**
- trap (Security PKI) | **30140**
- traps | **30142**
- traps (RPM and TWAMP) | **30144**
- tri-color | **30147**
- tri-color (EX Series Switches) | **30149**
- trickling | **30152**
- trigger | **30154**
- trigger (CoS) | **30159**
- trigger (DHCP Local Server) | **30160**
- trigger (Event Policy) | **30162**
- trigger-link-failure | **30164**
- trigger-type (LRF Profile) | **30165**
- trio-flow-offload | **30167**
- trust-option-82 | **30168**
- trusted | **30170**
- trusted-ca | **30171**
- trusted-ca (Services) | **30172**
- trusted (DHCP Security) | **30174**
- trusted-keys (DNSSEC) | **30175**
- tti | **30176**
- ttl | **30178**
- ttl (Dynamic Flow Capture) | **30179**
- ttl (RPM probe) | **30180**
- ttl-threshold | **30183**
- tunable-name | **30184**
- tunable-value | **30186**
- tunables | **30187**
- tunnel | **30189**
- tunnel | **30190**
- tunnel | **30192**

- tunnel-attributes | **30194**
- tunnel-devices (Protocols AMT) | **30196**
- tunnel-devices (Tunnel-Capable PICs) | **30198**
- tunnel-end-point | **30200**
- tunnel (fti) | **30204**
- tunnel-group | **30206**
- tunnel-inspection | **30208**
- tunnel (L2TP) | **30210**
- tunnel-limit (Protocols AMT) | **30212**
- tunnel-limit (Routing Instances) | **30213**
- tunnel-limit (Routing Instances Provider Tunnel Selective) | **30215**
- tunnel-mtu | **30216**
- tunnel-mtu (Services IPsec VPN) | **30218**
- tunnel-mtu (Services Service Set) | **30219**
- tunnel-observation | **30221**
- tunnel-port | **30223**
- tunnel-port (GRE) | **30224**
- tunnel-profile (Domain Map) | **30225**
- tunnel-profile (L2TP Tunnel Switching) | **30226**
- tunnel-profile (Tunnel Profile) | **30228**
- tunnel-queuing | **30230**
- tunnel (Security Policies) | **30231**
- tunnel (Security Screen) | **30232**
- tunnel-services (Chassis) | **30234**
- tunnel-services (Routing Instances VPLS) | **30236**
- tunnel-services (RSVP) | **30238**
- tunnel-source | **30239**
- tunnel-source-protocol (Protocols OSPF) | **30240**
- tunnel-switch-profile (L2TP Tunnel Switching, Application) | **30242**
- tunnel-switch-profile (L2TP Tunnel Switching, Definition) | **30243**
- tunnel-termination | **30245**
- tunnel-timeout | **30247**

- tunnel (Tunnel Profile) | **30248**
- twamp | **30250**
- twamp (Security ALG) | **30256**
- twamp-server | **30258**
- two-rate | **30259**
- tx-address-change (L2TP LAC) | **30261**
- tx-buffers | **30263**
- tx-connect-speed-method (L2TP LAC) | **30265**
- tx-dnu-to-line-source-enable | **30268**
- tx-duration | **30269**
- tx-enable | **30271**
- tx-power | **30272**
- type | **30273**
- type (Application Identification) | **30275**
- type (ICMP Mapping for Application Identification) | **30276**
- type (Protocols BGP) | **30278**
- type (RMON) | **30279**
- type (RMON Notification) | **30281**
- type (Security Antivirus Feature Profile) | **30282**
- type (Security Content Filtering Notification Options) | **30284**
- type (Security Fallback Block) | **30285**
- type (Security IDP Dynamic Attack Group) | **30287**
- type (Security IDP ICMP Headers) | **30288**
- type (Security Virus Detection) | **30289**
- type (Security Web Filtering) | **30291**
- type (SNMP v3) | **30293**
- type (Tunnel Profile) | **30294**

t1-interval

IN THIS SECTION

- Syntax | 28885
- Hierarchy Level | 28885
- Description | 28885
- Options | 28885
- Required Privilege Level | 28886
- Release Information | 28886

Syntax

```
t1-interval value-in-milliseconds;
```

Hierarchy Level

```
[edit logical-systems name security alg sip],
[edit logical-systems name tenants name security alg sip],
[edit security alg sip],
[edit services alg sip],
[edit tenants name security alg sip]
```

Description

Configure t1-interval parameter to specify the maximum round-trip time (RTT) (in milliseconds) allowed for Session Initiation Protocol (SIP) transactions.

Options

value-in-milliseconds—Maximum round-trip time (RTT) allowed measured in milliseconds.

- **Range:** 500 through 5000 milliseconds

- **Default:** 500 milliseconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sip (Security) | 27977

t1-options

IN THIS SECTION

- Syntax | 28886
- Hierarchy Level | 28887
- Description | 28887
- Options | 28887
- Required Privilege Level | 28887
- Release Information | 28887

Syntax

```
t1-options {
    framing (esf | sf);
    line-encoding (ami | b8zs);
}
```

Hierarchy Level

```
[edit chassis synchronization interfaces external]
[edit chassis synchronization interfaces (external-0/0 | external-1/0) e1-options]
```

Description

Configure the T1 interface options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

t1-options (BITS Interfaces Signal Type)

IN THIS SECTION

- Syntax | 28888
- Hierarchy Level | 28888
- Description | 28888

- Options | 28888
- Required Privilege Level | 28888
- Release Information | 28888

Syntax

```
t1-options {
    framing (esf | sf);
}
```

Hierarchy Level

```
[edit chassis synchronization interfaces bits]
```

Description

Specify options for the T1 signal type for the building-integrated timing supply (BITS) interface. Valid only if `t1` is the `signal-type` configured at the [`edit chassis synchronization source interfaces` *`interface-name`* `bits signal-type`] hierarchy.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

External Clock Synchronization Overview for ACX Series Routers

Configuring External Clock Synchronization for ACX Series Routers

Synchronous Ethernet Overview

show chassis synchronization | **18879**

t310

IN THIS SECTION

- Syntax | **28889**
- Hierarchy Level | **28889**
- Description | **28889**
- Options | **28890**
- Required Privilege Level | **28890**
- Release Information | **28890**

Syntax

```
t310-value seconds;
```

Hierarchy Level

```
[edit interfaces br-pim/0/port isdn-options]
```

Description

For ISDN interfaces, configure the Q.931-specific timer for T310, in seconds. The Q.931 protocol is involved in the setup and termination of connections.

Options

seconds—Timer value, in seconds.

- **Range:** 1 through 65,536 seconds
- **Default:** 10 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

t391

IN THIS SECTION

- Syntax | 28890
- Hierarchy Level | 28891
- Description | 28891
- Options | 28891
- Required Privilege Level | 28891
- Release Information | 28891

Syntax

```
t391 number;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

For link services interfaces only, set the Frame Relay link integrity polling interval.

Options

number—Link integrity polling interval.

- **Range:** 5 through 30 seconds
- **Default:** 10 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Keepalives on Link Services Physical Interfaces

n391 | **11163**

n392 | **11164**

t392 | **28892**

n393 | **11166**

t392

IN THIS SECTION

- Syntax | 28892
- Hierarchy Level | 28892
- Description | 28892
- Options | 28892
- Required Privilege Level | 28892
- Release Information | 28893

Syntax

```
t392 number;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

For link services interfaces only, set the Frame Relay polling verification interval.

Options

number—Polling verification interval.

- **Range:** 5 through 30 seconds
- **Default:** 15 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Keepalives on Link Services Physical Interfaces

n391 | **11163**

n392 | **11164**

n393 | **11166**

t391 | **28890**

t4-interval

IN THIS SECTION

- Syntax | **28893**
- Hierarchy Level | **28894**
- Description | **28894**
- Options | **28894**
- Required Privilege Level | **28894**
- Release Information | **28894**

Syntax

```
t4-interval value-in-seconds;
```

Hierarchy Level

```
[edit logical-systems name security alg sip],
[edit logical-systems name tenants name security alg sip],
[edit security alg sip],
[edit services alg sip],
[edit tenants name security alg sip]
```

Description

Configure t4-interval parameter to specify the maximum length of time (in seconds) that the network can take to clear messages between client and server Session Initiation Protocol (SIP) transactions.

Options

value-in-seconds—Maximum number of seconds that the network can take to clear messages between client and server transactions.

- **Range:** 5 through 10 seconds
- **Default:** 5 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sip (Security) | 27977

table

IN THIS SECTION

- Syntax | 28895
- Hierarchy Level | 28895
- Description | 28895
- Options | 28896
- Required Privilege Level | 28896
- Release Information | 28896

Syntax

```
table table-name;
```

Hierarchy Level

```
[edit dynamic policy-options condition],
[edit logical-systems logical-system-name policy-options condition if-route-exists ],
[edit logical-systems logical-system-name policy-options condition if-route-exists address-family
ccc],
[edit logical-systems logical-system-name policy-options condition if-route-exists address-family
inet],
[edit policy-options condition if-route-exists],
[edit policy-options condition if-route-exists address-family ccc],
[edit policy-options condition if-route-exists address-family inet]
```

Description

Specify a routing table in which the route must exist for the condition to be met and the route to be considered a match.

Options

`table` *`table-name`*—Routing table name, such as inet.0.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.

Support for the address families introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Conditional Advertisement and Import Policy (Routing Table) with certain match conditions

Example: Configuring Pseudowire Redundancy in a Mobile Backhaul Scenario

dynamic-db | **6105**

tacplus

IN THIS SECTION

- Syntax | **28897**
- Hierarchy Level | **28897**
- Description | **28897**
- Options | **28897**
- Required Privilege Level | **28897**

Release Information | 28898

Syntax

```
tacplus {
    server {
        server-address {
            port port-number;
            routing-instance routing-instance;
            secret password;
            single-connection;
            source-address address;
            timeout seconds;
        }
    }
}
```

Hierarchy Level

```
[edit system accounting destination]
```

Description

Configure the Terminal Access Controller Access Control System Plus (TACACS+) servers.

Options

server-address—Address of the TACACS++ authentication server.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`routing-instance` option introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Configuring TACACS+ System Accounting

tacplus-options

IN THIS SECTION

- Syntax | **28898**
- Hierarchy Level | **28899**
- Description | **28899**
- Options | **28899**
- Required Privilege Level | **28901**
- Release Information | **28901**

Syntax

```
tacplus-options {
    (exclude-cmd-attribute | no-cmd-attribute-value);
    authorization-time-interval minutes;
    enhanced-accounting;
    (strict-authorization | no-strict-authorization);
    service-name service-name;
    timestamp-and-timezone;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure TACACS+ options for authentication and accounting.

Options

authorization-time-interval *minutes*

Configure the time interval at which the authorization profile that is configured on the TACACS+ server is fetched by the Junos OS device during a TACACS+ authentication session. The TACACS+ server sends the authorization profile once by default after the user is successfully authenticated, and the authorization profile is stored locally on the Junos OS device. The authorization-time-interval option enables the Junos OS device to periodically check the authorization profile configured remotely on the TACACS+ server at the configured time interval.

If there is a change in the remote authorization profile, the device fetches the authorization profile from the TACACS+ server and the authorization profile configured locally under the [edit system login class class-name] hierarchy. The device refreshes the authorization profile stored locally by combining the remote and locally-configured authorization profiles. This ensures that any changes made to the authorization profile configuration on the TACACS+ server are reflected on the Junos OS device without the user having to restart the authentication process.

To enable the periodic refresh of the authorization profile, you must set the authorization time interval at which the Junos OS device fetches the authorization profile configuration from the TACACS+ server and refreshes the authorization profile stored locally. The time interval can be configured directly on the TACACS+ server or locally on the Junos OS device using the CLI. Use the following guidelines to determine which time interval configuration takes precedence:

- If there is no time interval configured on the TACACS+ server for periodic refresh, the Junos OS device does not receive the time interval value in the authorization response. In this case, the value configured locally on the Junos OS device will take effect.
- If the time interval is configured on the TACACS+ server and there is no authorization time interval configured locally on the Junos OS device, the value configured on the TACACS+ server will take effect.

- If the periodic refresh time interval is configured on the TACACS+ server and also locally on the Junos OS device, the value configured on the TACACS+ server will take precedence.
- If there is no periodic refresh time interval configured on the TACACS+ server and there is no authorization time interval configured on the Junos OS device, there will be no periodic refresh.
- If the periodic refresh time interval configured on the TACACS+ server is out of range or invalid, the authorization time interval value configured locally will take effect.
- If the periodic refresh time interval configured on the TACACS+ server is out of range or invalid and there is no authorization time interval configured locally, there will be no periodic refresh.

After the periodic authorization time interval is set, if the user changes the interval before the authorization request is sent from the Junos OS device, the updated interval takes effect after the next immediate periodic refresh.

- **Default:** If the authorization time interval is not configured, the authorization profile is not refreshed during a TACACS+ authentication session.
- **Range:** 15 through 1440 minutes

enhanced-accounting

Configure the audit of TACACS+ authentication events, such as access method, remote port, and access privileges.

exclude-cmd-attribute

Exclude the `cmd` attribute value completely from start and stop accounting records to enable logging of accounting records in the correct log file on a TACACS+ server.

no-cmd-attribute-value

Set the `cmd` attribute value to an empty string in the TACACS+ accounting start and stop requests to enable logging of accounting records in the correct log file on a TACACS+ server.

no-strict-authorization

Don't deny login if the authorization request fails. When a user is logging in, Junos OS issues two TACACS+ requests—first the authentication request followed by the authorization request.

- **Default:** By default, when the authorization request is rejected by the TACACS+ server, Junos OS ignores this and allows full access to the user. Specifying `no-strict-authorization` restores this default behavior.

service-name *service-name*

Name of the authentication service used when you configure multiple TACACS+ servers to use the same authentication service.

- **Default:** `junos-exec`

strict-authorization

Deny login if the authorization request fails. When a user is logging in, Junos OS issues two TACACS+ requests—first the authentication request followed by the authorization request. When the `strict-authorization` option is specified, Junos OS denies access to the user even when the TACACS+ authorization request fails.

- **Default:** By default, when the authorization request is rejected by the TACACS+ server, Junos OS ignores this and allows full access to the user.

timestamp-and-timezone

Include this statement if you want start time, stop time, and time zone attributes included in the start and stop accounting records.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`no-cmd-attribute-value` and `exclude-cmd-attribute` options introduced in Junos OS Release 9.3.

`timestamp-and-timezone` option introduced in Junos OS Release 12.2.

`strict-authorization` and `no-strict-authorization` options introduced in Junos OS Release 13.3 for EX Series, M Series, MX Series, PTX Series, and T Series. Support for these options was added to Junos OS Evolved starting in Release 21.1.

`enhanced-accounting` option introduced in Junos OS Release 14.1.

`authorization-time-interval` option introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

Configuring Periodic Refresh of the TACACS+ Authorization Profile

Configuring TACACS+ Authentication

Configuring TACACS+ System Accounting

Determine the Authentication Order for LDAPS, RADIUS, TACACS+, and Password Authentication

tacplus-server

IN THIS SECTION

- Syntax | 28902
- Hierarchy Level | 28902
- Description | 28902
- Options | 28902
- Required Privilege Level | 28904
- Release Information | 28904

Syntax

```
tacplus-server server-address {
    port port-number;
    routing-instance routing-instance;
    secret password;
    single-connection;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the IPv4 or IPv6 TACACS+ server.

Options

server-address	Address of the IPv4 or IPv6 TACACS+ authentication server.

> **NOTE**: Wildcard characters cannot be used in the TACACS+ server address or source address. This is because the TACACS+ server and source can accept both IPv4 and IPv6 addresses and, if you use wildcard characters for these addresses, Junos OS cannot validate mismatching server and source address families.

port *port-number*

Configure the port number on which to contact the TACACS+ authentication server.

- **Default:** 49

routing-instance *routing-instance*

Configure the routing instance name for the management routing instance, that is `mgmt_junos`. Configuring this parameter along with the `management-instance` statement enables authentication processes (for example, RADIUS and TACACS+) to use the non-default management routing instance for packet traffic.

> **NOTE**: You must also define the mgmt_junos routing instance under the `[edit routing-instances]` hierarchy level.

If you no not configure the mgmt_junos instance under the `[edit routing-instances]` hierarchy level and configure it only under `tacplus-server` or `radius-server`, the commit will fail.

secret *password*

Configure the password to use with the RADIUS or TACACS+ server. The secret password used by the local router or switch must match that used by the server. The password can include spaces included in quotation marks.

> **NOTE**: To ensure better security, we recommend you configure the TACACS+ secret password with a minimum of 14 characters.

single-connection

Optimize attempts to connect to a TACACS+ server. The software maintains one open TCP connection to the server for multiple requests rather than opening a connection for each connection attempt.

source-address *source-address*

Specify a source address for each configured TACACS+ server to record in system log messages that are directed to a remote machine. Configure a valid IP address on one of the device interfaces. For system logging, the address is recorded as the message source in messages sent to the remote machines specified in all `host` *`hostname`* statements at the `[edit system syslog]` hierarchy level.

- **Default:** The primary address of the interface.

timeout *seconds*

The amount of time that the local device waits to receive a response from a TACACS+ server.

- **Default:** 3 seconds
- **Range:** 1 through 90 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`routing-instance` option introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Configuring TACACS+ Authentication

tag

IN THIS SECTION

- Syntax | 28905
- Hierarchy Level | 28905
- Description | 28905
- Options | 28905
- Required Privilege Level | 28905
- Release Information | 28905

Syntax

```
tag tag-name;
```

Hierarchy Level

```
[edit snmp v3 notify name],
[edit snmp v3 snmp-community community-index]
```

Description

Configure a set of targets to receive traps or informs (for IPv4 packets only).

Options

tag-name—Identifies the address of managers that are allowed to use a community string.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

tag (Access)

IN THIS SECTION

- Syntax | 28906
- Hierarchy Level | 28906
- Description | 28906

- Options | 28906
- Required Privilege Level | 28906
- Release Information | 28907

Syntax

```
tag route-tag;
```

Hierarchy Level

```
[edit routing-options access route ip-prefix</prefix-length>]
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
access route prefix],
[edit dynamic-profiles profile-name routing-instances  $junos-routing-instance routing-options
rib routing-table-name access route prefix],
[edit dynamic-profiles profile-name routing-options access route prefix]
```

Description

Dynamically configure the tag for an access route.

Options

route-tag—Either the specific tag you want to assign to the access route or either of the following tag variables:

- `$junos-framed-route-tag`—Tag assigned to an IPv4 access route; the variable is dynamically replaced with the preference value (Subattribute 6) from the RADIUS Framed-Route attribute [22].
- `$junos-framed-route-ipv6-tag`—Tag assigned to an IPv6 access route; the variable is dynamically replaced with the preference value (Subattribute 6) from the RADIUS Framed-IPv6-Route attribute [99].

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Support at the [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*] and [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*] hierarchy levels introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Routes for DHCP and PPP Subscriber Access Networks

tag-attribute (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 28907
- Hierarchy Level | 28907
- Description | 28908
- Options | 28908
- Required Privilege Level | 28908
- Release Information | 28908

Syntax

```
tag-attribute [tag-attr-name];
```

Hierarchy Level

```
[edit services hcm]
```

Description

Specify one or more tag attributes that can be used in tag rules for HTTP header enrichment.

These attributes are stored in the subscriber database for subscribers. After these attributes are configured, they can be used in the tag rules. HTTP tag rules can be configured to choose one or more of these attributes to insert in the HTTP header.

Options

tag-attr-name—Tag attribute. To specify multiple attributes at one time, include the attributes in square brackets ([]). The supported attributes are `apn`, `ggsnipv4`, `ggsnipv6`, `imei`, `imsi`, `ipv4addr`, `ipv6addr`, and `msisdn`.

- **Range:** 1 through 63 alphanumeric characters

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

Junos Web Aware HTTP Header Enrichment Overview

tag-attribute (HTTP Header Enrichment Tag Rule)

IN THIS SECTION

Syntax | 28909

Hierarchy Level | 28909

- Description | 28909
- Options | 28909
- Required Privilege Level | 28909
- Release Information | 28910

Syntax

```
tag-attribute [tag-attr-name];
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
```

Description

Specify one or more tag attributes (for the tag header and separator) to insert into the HTTP header.

> **NOTE**: The tag attribute specified here must already be defined at the `[edit services hcm]` hierarchy level.

Options

tag-attr-name—Tag attribute. To specify multiple attributes at one time, include the attributes in square brackets ([]). The supported attributes are `apn`, `ggsnipv4`, `ggsnipv6`, `imei`, `imsi`, `ipv4addr`, `ipv6addr`, and `msisdn`.

- **Range:** 1 through 63 alphanumeric characters

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

tag (Configuring Notification Targets)

IN THIS SECTION

- Syntax | 28910
- Hierarchy Level | 28910
- Description | 28910
- Options | 28911
- Required Privilege Level | 28911
- Release Information | 28911

Syntax

```
tag tag-name;
```

Hierarchy Level

```
[edit snmp v3 notify name]
```

Description

Configure a set of target addresses to receive SNMP traps or informs (for IPv4 packets only).

Options

tag-name—Define the target addresses to which an SNMP notification is sent. Target addresses containing the same tag in their tag list are sent the same notification. The *tag-name* is not included in the notification.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

tag-group

IN THIS SECTION

- Syntax | 28911
- Hierarchy Level | 28912
- Description | 28912
- Options | 28912
- Required Privilege Level | 28912
- Release Information | 28912

Syntax

```
tag-group name {
      application-tags [ application-tags... ]
      }
```

Hierarchy Level

```
[edit services application-identification application-group application-group-name
```

Description

Specify a unique tag group identifier to create your own application group based on the application attributes. Application group allows you to associate a related applications under a single name for simplified, consistent reuse when using any application services.

For example: If you want to group application traffic for web or remote access, you can specify respective tags as follows:

```
user@host# set services application-identification application-group A1 tag-group TG-1
application-tags [web remote_access]
```

You can use the application groups as match condition when defining security policies.

Options

name A unique tag group identifier

application-tags Application tag to configure the application group.

(Supported tags available in the latest version of application signature) - basic, networking, email, webmail, im_mc, chat, audio_chat, video_chat, file_transfer, social_network, voip, web, web_search, web_ecom, web_sites, mobile, p2p, file_mngt, db, enterprise, update, gaming, aaa, remote_access, mm_streaming, vpn_tun, cdn, news_portal, classified_ads, advertising, analytics, adult_content, anonymizer, blog, forum, standardized, scada, iot, cloud_services, dea_mail, ebook_reader, transportation, cryptocurrency, map_service, crowdfunding, speedtest, payment_service

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Relesae 21.1R1.

RELATED DOCUMENTATION

Configure Packet Capture For Unknown Application Traffic

tag-header (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 28913
- Hierarchy Level | 28913
- Description | 28913
- Options | 28913
- Required Privilege Level | 28914
- Release Information | 28914

Syntax

```
tag-header header;
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
```

Description

Specify the tag header for the tag to be inserted into the HTTP header. This is a required configuration.

You can configure a total of 16 unique tag headers for all the tag rules you configure.

Options

header—Tag header.

- **Values:** You *cannot* use the following values: `accept`, `accept-charset`, `accept-encoding`, `accept-language`, `authorization`, `expect`, `host`, `if-match`, `if-modified-since`, `if-none-match`, `if-range`, `if-unmodified-since`, `max-forwards`, `proxy-authorization`, `referer`, `user-agent`, or `x-moz`. These header values are reserved; you cannot configure them.
- **Range:** 1 through 63 alphanumeric characters

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

Configuring HCM Profiles and Assigning Tag Rules

Junos Web Aware HTTP Header Enrichment Overview

tag (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | **28915**
- Hierarchy Level | **28915**
- Description | **28915**
- Options | **28915**
- Required Privilege Level | **28915**
- Release Information | **28916**

Syntax

```
tag tag-name {
    encrypt {
        hash algorithm;
        prefix hash-prefix;
    }
    ipv4-mask ipv4-mask;
    ipv6-mask ipv6-mask;
    ipv4-or-value ipv4-or-value;
    ipv6-or-value ipv6-or-value;
    tag-attribute tag-attr-name;
    tag-header header;
    tag-separator separator;
}
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number then]
```

Description

Configure the tags to be applied to the HTTP headers. If you configure a tag, you must include the `tag-header` statement.

Options

tag-name	Name of the tag.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

Configuring HCM Profiles and Assigning Tag Rules

Junos Web Aware HTTP Header Enrichment Overview

tag-list

IN THIS SECTION

- Syntax | 28916
- Hierarchy Level | 28916
- Description | 28917
- Options | 28917
- Required Privilege Level | 28917
- Release Information | 28917

Syntax

```
tag-list tag-list;
```

Hierarchy Level

```
[edit snmp v3 target-address target-address-name]
```

Description

Configure an SNMP tag list used to select target addresses.

Options

tag-list—Define sets of target addresses (tags). To specify more than one tag, specify the tag names as a space-separated list enclosed within double quotes.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

tag-operation (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 28917
- Hierarchy Level | 28918
- Description | 28918
- Options | 28918
- Required Privilege Level | 28918
- Release Information | 28918

Syntax

```
tag-operation (add | delete | modify);
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-name then tag tag-name]
```

Description

Specify the operation to be performed on the specified tag of the tag rule set.

NOTE: The tag rule must already be defined at the `[edit services hcm]` hierarchy level.

Options

add Add the specified tag with previously existing tag in the tag rule set.

delete Delete the specified tag from the tag rule set.

modify Modify the existing tag in the tag rule set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm (HTTP Header Enrichment) | 8005

tag-protocol-id (Dynamic VLANs)

IN THIS SECTION

- Syntax | 28919
- Hierarchy Level | 28919
- Description | 28919
- Default | 28919
- Options | 28920
- Required Privilege Level | 28920
- Release Information | 28920

Syntax

```
tag-protocol-id tpids;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number input-vlan-map],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map]
```

Description

For dynamic VLAN interfaces, configure the outer TPID value. All TPIDs you include in input and output VLAN maps must be among those you specify at the `[edit interfaces interface-name gigether-options ethernet-switch-profile tag-protocol-id [ tpids ]]` hierarchy level.

Default

If the `tag-protocol-id` statement is not configured, the TPID value is 0x8100.

Options

tpids—TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal format.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Inner and Outer TPIDs and VLAN IDs

tag-protocol-id (TPID to Rewrite)

IN THIS SECTION

- Syntax | 28920
- Hierarchy Level | 28921
- Description | 28921
- Default | 28921
- Required Privilege Level | 28921
- Release Information | 28921

Syntax

```
tag-protocol-id tpid;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, configure the outer TPID value. All TPIDs you include in input and output VLAN maps must be among those you specify at the `[edit interfaces interface-name gigether-options ethernet-switch-profile tag-protocol-id [ tpids ]]` hierarchy level.

For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers the default TPID value (`0x8100`) is supported.

Default

If the `tag-protocol-id` statement is not configured, the TPID value is 0x8100.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Inner and Outer TPIDs and VLAN IDs

tag-protocol-id TPID(TPIDs Expected to Be Sent or Received)

IN THIS SECTION

- Syntax | 28922
- Hierarchy Level | 28922
- Description | 28922
- Options | 28923
- Required Privilege Level | 28923
- Release Information | 28923

Syntax

```
tag-protocol-id [tpids];
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile],
[edit interfaces interface-name aggregated-ether-options ethernet-switch-profile],
[edit interfaces interface-name aggregated-ether-options ethernet-switch-profile],
[edit interfaces interface-name ether-options ethernet-switch-profile]
```

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, and the built-in Gigabit Ethernet port on the M7i router), define the TPIDs expected to be sent or received on a particular VLAN. For each Gigabit Ethernet port, you can configure up to eight TPIDs using the `tag-protocol-id` statement; but only the first four TPIDs are supported on IQ2 and IQ2-E interfaces.

For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers only the default TPID value (`0x8100`) is supported.

For Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches, define the TPIDs expected to be sent or received on a particular VLAN. The default TPID value is `0x8100`. Other supported values are `0x88a8`, `0x9100`, and `0x9200`.

Starting in Junos OS Release 22.4R1, the QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 switches support up to four TPID configurations per VLAN. The system supports the following values for the TPID configuration: 0x8100, 0x9200, 0x88a8, and 0x9100. A TPID needs to be defined in the command-line interface for aggregated ethernet interfaces and the ether interfaces for the device, before it can be sent or received.

> **NOTE**: Routing Engine can handle only the default TPID `0x8100`. The traffic destined to the Routing Engine is discarded if the TPID is not `0x8100`.

Options

tpids—TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

- Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames
- Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support

tag (Routing Options)

IN THIS SECTION

- Syntax | 28924
- Hierarchy Level | 28924
- Description | 28925
- Default | 28925
- Options | 28925
- Required Privilege Level | 28925
- Release Information | 28925

Syntax

```
tag metric type number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options (aggregate | generate | static) (defaults | route)],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name (aggregate | generate | static) (defaults | route)],
[edit logical-systems logical-system-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit logical-systems logical-system-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-instances routing-instance-name routing-options (aggregate | generate | static)
(defaults | route)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name (aggregate |
generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib routing-table-name (aggregate | generate | static) (defaults | route)]
```

Description

Associate a tag with a static, aggregate, or generated route.

Default

No tag strings are associated with routes.

Options

metric—Tag metric.

- **Range:** 0 through 4,294,967,295

type *number*—Tag type.

- **Range:** 1 through 16

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation

aggregate (Routing) | **2070**

generate | **7672**

static (Routing Options) | **28570**

tag-rule (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 28926
- Hierarchy Level | 28927
- Description | 28927
- Options | 28927
- Required Privilege Level | 28928
- Release Information | 28928

Syntax

```
tag-rule rule-name {
    term term-number {
        from {
            destination-address {
                (any-ipv4 | any-ipv4 except);
                (any-ipv6 | any-ipv6 except);
                (any-unicast | any-unicast except);
                (prefix | prefix except);
            }
            destination-address-range {
                high address low address <except>;
            }
            destination-port-range {
                high port-number low port-number;
            }
            destination-ports value;
            destination-prefix-list {
                (prefix-name | prefix-name except);
            }
        }
        then {
            count;
            tag tag-name {
                encrypt {
```

```
                    hash algorithm;
                    prefix hash-prefix;
                }
                ipv4-mask ipv4-mask;
                ipv6-mask ipv6-mask;
                ipv4-or-value ipv4-or-value;
                ipv6-or-value ipv6-or-value;
                tag-attribute tag-attr-name;
                tag-header header;
                tag-separator separator;
            }
        }
    }
}
```

Hierarchy Level

```
[edit services hcm]
```

Description

Configure the tag rules that enrich HTTP headers with the appropriate tags.

You must configure at least one term for a tag rule, but you can configure multiple terms. Terms are evaluated in the order they are configured. If a data packet matches all the criteria in the `from` statement in any of the terms, then the actions specified in the `then` statement are applied. If the `from` statement does not identify any criteria, then all traffic matches. After a term matches a data packet, further terms are not evaluated. If no terms match, then the HTTP header is not enriched.

For subscriber traffic under static policy control, a tag rule is used if it is included in the HCM profile specified in a PCC rule that the traffic matches. For subscribers under dynamic policy control, a message from the PCRF identifies the configured HCM profile and tag rules to use for HTTP header enrichment.

Options

`rule-name`—Name of the tag rule.

- **Range:** 1 through 63 alphanumeric characters

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

Configuring HCM Profiles and Assigning Tag Rules

Junos Web Aware HTTP Header Enrichment Overview

tag-rule (Profiles for HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 28928
- Hierarchy Level | 28929
- Description | 28929
- Options | 28929
- Required Privilege Level | 28929
- Release Information | 28929

Syntax

```
tag-rule rule-name;
```

Hierarchy Level

```
[edit services hcm profile profile-name]
```

Description

Specify the tag rule to be associated with the HCM profile.

> **NOTE**: The tag rule specified here must already be defined at the `[edit services hcm]` hierarchy level.

Options

rule-name Name of the tag rule.

- **Range:** 1 through 63 alphanumeric characters

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

Configuring HCM Profiles and Assigning Tag Rules

tag-rule-set (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 28930
- Hierarchy Level | 28930
- Description | 28930
- Options | 28930
- Required Privilege Level | 28931
- Release Information | 28931

Syntax

```
tag-rule-set rule-set-name {
    [rule rule-name];
}
```

Hierarchy Level

```
[edit services hcm]
```

Description

Configure the tag rule set for HTTP header enrichment so that you can group multiple configured tag rules into one tag rule set.

Options

rule-set-name	Name of the tag rule set.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm (HTTP Header Enrichment) | 8005

tag-rule-sets (Service Set)

IN THIS SECTION

- Syntax | 28931
- Hierarchy Level | 28932
- Description | 28932
- Options | 28932
- Required Privilege Level | 28932
- Release Information | 28932

Syntax

```
tag-rule-sets rule-set-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the HTTP header enrichment tag rule set included in the service set. You can configure only one tag rule set. If you specify a tag rule set, you cannot specify a tag rule.

Options

rule-set-name	Name of the tag rule set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Applying Services to Subscriber-Aware Traffic with a Service Set

Configuring HTTP Header Enrichment Overview

tag-rules (Service Set)

IN THIS SECTION

- Syntax | 28933
- Hierarchy Level | 28933

- Description | 28933
- Options | 28933
- Required Privilege Level | 28933
- Release Information | 28933

Syntax

```
[tag-rules rule-name;]
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify one or more HTTP header enrichment tag rules to include in the service set. You can configure multiple tag rules. If you specify any tag rules, you cannot specify a tag rule set.

Options

rule-name Name of the tag rule.

- **Range:** 1 through 63 alphanumeric characters

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Applying Services to Subscriber-Aware Traffic with a Service Set

Configuring HTTP Header Enrichment Overview

tag-separator (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 28934
- Hierarchy Level | 28934
- Description | 28934
- Options | 28934
- Required Privilege Level | 28935
- Release Information | 28935

Syntax

```
tag-separator separator;
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
```

Description

Specify the tag separator for the tag to be inserted into the HTTP header.

Options

separator Tag separator. You may use a forward slash (/) or pipe (|).

- **Default:** / (forward slash)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

Configuring HCM Profiles and Assigning Tag Rules

Junos Web Aware HTTP Header Enrichment Overview

tag-value (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 28936
- Hierarchy Level | 28936
- Description | 28936
- Options | 28936
- Required Privilege Level | 28936
- Release Information | 28936

Syntax

```
tag-value value;
```

Hierarchy Level

```
[edit services hcm tag-rule then tag]
```

Description

Specify the tag value for the specified tag name.

NOTE: The tag rule must already be defined at the `[edit services hcm]` hierarchy level.

Options

value String of up to 16 alphanumeric characters

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

hcm (HTTP Header Enrichment) | 8005

tag2 (Dynamic Access Routes)

IN THIS SECTION

- Syntax | 28937
- Hierarchy Level | 28937
- Description | 28937
- Options | 28937
- Required Privilege Level | 28938
- Release Information | 28938

Syntax

```
tag2 route-tag2;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances $junos-routing-instance routing-options
access route prefix],
[edit dynamic-profiles profile-name routing-options access route prefix]
```

Description

Options

route-tag2 One of the following values the specific tag2 value you want to assign to the access route or the following predefined variable:

- A specific tag 2 value for the specified access route prefix.

- `$junos-framed-route-tag2`—Tag2 value assigned to an IPv4 access route. The value is dynamically replaced with the preference value (subattribute 6) from the RADIUS Framed-Route attribute [22]. You configure this variable only when the access route prefix is derived from the $junos-framed-route-ip-address-prefix predefined variable; this value is (subattribute 1) of the RADIUS Framed-Route attribute [22].

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

PPP Subscriber Access Networks Overview

tags (Application Identification)

IN THIS SECTION

- Syntax | 28939
- Hierarchy Level | 28939
- Description | 28939
- Options | 28939
- Required Privilege Level | 28939
- Release Information | 28939

Syntax

```
tags tag-name tag-value;
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Description

Specify an application tag that provides general information about the application, such as associated risk factors, technology, and the type of traffic. The tag consists of a user-defined name and value.

Options

tag-name	Name for the tag, which is a textual string.
tag-value	Value for the tag.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> *Application Identification Overview*
>
> *Configuring Custom Application Signatures*

talk

IN THIS SECTION

- Syntax | 28940
- Hierarchy Level | 28940
- Description | 28941
- Options | 28941
- Required Privilege Level | 28941
- Release Information | 28941

Syntax

```
talk {
    disable;
    traceoptions {
        flag {
            all {
                extensive
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Configure `talk` to specify the TALK program ALG on the device. TALK ALG is a visual communication program used for interactive communication between two users.

Options

- `disable`—Disable the TALK program ALG. By default, the TALK program ALG is enabled.
- `traceoptions`—Configure TALK program ALG tracing options.
  - `flag`—Trace operation to perform.
    - `all`—Trace all events.
    - *`extensive`*—Display extensive amount of data.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding the TALK ALG

target

IN THIS SECTION

Syntax (Junos OS) | **28942**

- Syntax (Junos OS Evolved) | 28942
- Junos OS Hierarchy Level | 28942
- Junos OS Evolved Hierarchy Level | 28942
- Description | 28942
- Options | 28943
- Required Privilege Level | 28943
- Release Information | 28944

Syntax (Junos OS)

```
target (address IPv4-address | inet6-address IPv6-address | inet6-url IPv6-URL| url IPv4-URL)
```

Syntax (Junos OS Evolved)

```
target (url URL | address) local-link IPv6-link-local-interface-name;
```

Junos OS Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

(Required) Specify the destination IPv4 or IPv6 address or URL used for the probes.

Options

address (Junos OS Evolved) For all RPM probe types other than the HTTP probes, use the specified IPv4 or IPv6 address for the target host. (Prior to Junos OS Evolved Release 23.1R1, only IPv4 addresses are supported for RPM probes.) If you also configure the `link-local` option, you must configure an IPv6 address in the link-local prefix fe80::/8 range.

address ***IPv4-address*** (Junos OS) For all RPM probe types other than the HTTP probes, use the specified IPv4 address for the target host.

> **NOTE**: Starting with Junos OS Release 14.2R2, the RPM client router (the router or switch that originates the RPM probes) can send probe packets to the RPM probe server (the device that receives the RPM probes) that contains an IPv6 address.

inet6-address ***IPv6-address*** (Junos OS) For all RPM probe types other than the HTTP probes, use the specified IPv6 address for the target host.

inet6-url ***IPv6-URL*** (Junos OS) For HTTP probe types, use the specified fully formed URL that includes `http://` in the URL address. You can specify an IPv6 address of a host in the URL to denote the destination or server to which the RPM probes must be sent.

local-link ***IPv6-link-local-interface-name*** (Junos OS Evolved) For test sessions with IPv6 target addresses, configure the link-local logical interface name for the egress interface. If you configure this option, you cannot configure the `source-address` statement for the test.

url ***URL*** For HTTP probe types, use the specified fully formed URL that includes `http://` in the URL address. You can specify an IPv4 address of a host in the URL to denote the destination or server to which the RPM probes must be sent. For Junos OS Evolved, you can also use this option to specify an IPv6 address of a host in the URL; for Junos OS, you use the `inet6-url` option instead to specify the same.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Packet Transport routers.

Statement introduced in Junos OS Release 18.1 for QFX Series switches.

Statement introduced in Junos OS Evolved Release 20.1R1.

Support for IPv6 addresses introduced in Junos OS Evolved Release 23.1R1.

Support for TCP and HTTP probe types and URLs introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

target-address (TWAMP)

IN THIS SECTION

- Syntax | 28944
- Junos OS Hierarchy Levels | 28945
- Description | 28945
- Options | 28945
- Required Privilege Level | 28945
- Release Information | 28945

Syntax

```
target-address address local-link IPv6-link-local-interface-name;
```

Junos OS Hierarchy Levels

```
[edit services rpm twamp client control-connection control-client-name]
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Description

(Required) Specify the destination IPv4 or IPv6 address used for Junos OS probes. For specifying the target address for TWAMP control connections and tests, see "control-connection (Junos OS Evolved)" on page 4627.

Options

address

For TWAMP managed control connections and test sessions, use the specified IPv4 or IPv6 address for the target host. MX Series and PTX Series Junos OS TWAMP Light test sessions support both IPv4 and IPv6 addresses.

> **NOTE**: Starting with Junos OS Release 14.2R2, the RPM client router (the router or switch that originates the RPM probes) can send probe packets to the RPM probe server (the device that receives the RPM probes) that contains an IPv6 address.

local-link *IPv6-link-local-interface-name*

For TWAMP Light test sessions with IPv6 target addresses (where the `control-type light` statement is configured for the control connection), configure the link-local logical interface name for the egress interface.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Packet Transport routers.

Support at the [`edit services rpm twamp client control-connection` *`control-client-name`*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Release 18.1 for QFX Series switches.

Support for IPv6 addresses for TWAMP Light test sessions introduced in Junos OS Release 21.3R1 for MX Series and PTX1000, PTX3000, and PTX5000 routers.

`local-link` option for IPv6 TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

target-address (SNMP)

IN THIS SECTION

- Syntax | 28946
- Hierarchy Level | 28947
- Description | 28947
- Options | 28947
- Required Privilege Level | 28947
- Release Information | 28947

Syntax

```
target-address target-address-name {
    address (SNMP) address;
    address-mask address-mask;
    logical-system (SNMP) logical-system;
    port port-number;
    retry-count (SNMP) number;
    routing-instance (SNMPv3) instance;
    tag-list tag-list;
```

```
    target-parameters target-parameters-name;
    timeout (SNMP) seconds;
}
```

Hierarchy Level

```
[edit snmp v3]
```

Description

Configure the address of an SNMP management application and the parameters to be used in sending notifications.

Options

target-address-name—String that identifies the target address.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

target-attachment-identifier (Protocols VPWS)

IN THIS SECTION

- Syntax | 28948
- Hierarchy Level | 28948

- Description | 28948
- Options | 28948
- Required Privilege Level | 28949
- Release Information | 28949

Syntax

```
target-attachment-identifier identifier;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols l2vpn site
site-name interface interface-name],
[edit routing-instances instance-name protocols l2vpn site site-name interface interface-name]
```

Description

For FEC 129, specify the VPWS target attachment identifier. The point-to-point nature of VPWS requires that you specify the source access individual identifier (SAII) and the target access individual identifier (TAII). This SAII-TAII pair defines a unique pseudowire between two PE devices.

Auto-discovery routes are used by BGP to allow auto-discovery of SAIIs (the sources of the point-to-point pseudowires). One auto-discovery route is advertised for each source attachment identifier (SAI) in the instance or mesh group.

The SAII is specified with the `source-attachment-identifier` statement within the FEC 129 VPWS routing instance. You configure the source attachment identifier and the interfaces to associate with that source attachment identifier. Under each interface, you can configure the TAII with the `target-attachment-identifier` statement. If the configured target identifier matches a source identifier advertised by a remote PE device by way of a BGP auto-discovery message, the pseudowire between that source-target pair is signaled. If there is no match between an advertised source identifier and the configured target identifier, the pseudowire is not established.

Options

identifier—The numerical identifier for the Layer 2 VPN site.

- **Range:** 1 through 4,292,967,295

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring FEC 129 BGP Autodiscovery for VPWS

source-attachment-identifier (Protocols VPWS) | **28183**

target-logical-system (Domain Map)

IN THIS SECTION

- Syntax | **28950**
- Hierarchy Level | **28950**
- Description | **28950**
- Default | **28950**
- Options | **28950**
- Required Privilege Level | **28950**
- Release Information | **28951**

Syntax

```
target-logical-system logical-system-name {
    target-routing-instance routing-instance-name;
}
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Configure a non-default logical system and optionally a non-default routing instance for the subscriber's interface in a domain map.

You use the `target-routing-instance` statement at the `[edit access domain map domain-map-name]` hierarchy level to configure a non-default routing instance for the default logical system.

NOTE: Subscriber management is supported in the default logical system only.

Default

Default logical system for the subscriber..

Options

logical-system-name—Name of the logical system.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Mapping Subscriber Domains to Access and Session Options

target-parameters

IN THIS SECTION

- Syntax | 28951
- Hierarchy Level | 28952
- Description | 28952
- Required Privilege Level | 28952
- Release Information | 28952

Syntax

At the `[edit snmp v3]` hierarchy level:

```
target-parameters target-parameters-name {
    profile-name;
    parameters {
        message-processing-model (v1 | v2c | V3);
        security-level (authentication | none | privacy);
        security-model (usm | v1 | v2c);
        security-name security-name;
    }
}
```

At the [edit snmp v3 target-address *target-address-name*] hierarchy level:

```
target-parameters target-parameters-name;
```

Hierarchy Level

```
[edit snmp v3]
[edit snmp v3 target-address target-address-name]
```

Description

Configure the message processing and security parameters for sending notifications to a particular management target. The target parameters are configured at the [edit snmp v3] hierarchy level. The remaining statements at this level are explained separately.

Then apply the target parameters configured at the [edit snmp v3 target-parameters *target-parameters-name*] hierarchy level to the target address configuration at the [edit snmp v3] hierarchy level.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

target-routing-instance (Domain Map)

IN THIS SECTION

- Syntax | 28953
- Hierarchy Level | 28953
- Description | 28953

- Default | 28953
- Options | 28953
- Required Privilege Level | 28954
- Release Information | 28954

Syntax

```
target-routing-instance (routing-instance-name | default);
```

Hierarchy Level

```
[edit access domain map domain-map-name],
[edit access domain map domain-map-name target-logical-system logical-system-name]
```

Description

Configure the routing instance of the subscriber context.

> **NOTE**: Subscriber management is supported in the default logical system only. The `target-logical-system` statement, which appears in the CLI, is not supported in current Junos OS releases.

Default

For dynamic LNS sessions, the routing instance of the peer (LAC facing) interface. For all other sessions, the default logical system/routing instance context.

Options

routing-instance-name—Name of the routing instance.

`default`—The default routing instance.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`default` option added in Junos OS Release 13.3.

RELATED DOCUMENTATION

Mapping Subscriber Domains to Access and Session Options

Mapping Subscriber Domains to Access and Session Options

target (Routing Instances MVPN)

IN THIS SECTION

- Syntax | 28954
- Hierarchy Level | 28955
- Description | 28955
- Options | 28955
- Required Privilege Level | 28955
- Release Information | 28955

Syntax

```
target target-value {
    receiver target-value;
```

```
    sender target-value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvpn
route-target import-target],
[edit routing-instances routing-instance-name protocols mvpn route-target import-target]
```

Description

Specify the target value when importing sender and receiver site routes.

Options

target-value—Specify the target value when importing sender and receiver site routes.

receiver—Specify the target community used when importing receiver site routes.

sender—Specify the target community used when importing sender site routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring VRF Route Targets for Routing Instances for an MBGP MVPN

targeted-broadcast

IN THIS SECTION

- Syntax (EX Series, MX Series, QFX Series, ACX Series, SRX Series) | 28956
- Syntax (OCX1100, EX4600, NFX Series) | 28956
- Hierarchy Level (EX Series, MX Series, ACX Series) | 28956
- Hierarchy Level (QFX Series, OCX1100, EX4600, NFX Series) | 28957
- Hierarchy Level (SRX Series) | 28957
- Description | 28957
- Options | 28957
- Default | 28957
- Required Privilege Level | 28957
- Release Information | 28957

Syntax (EX Series, MX Series, QFX Series, ACX Series, SRX Series)

```
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
```

Syntax (OCX1100, EX4600, NFX Series)

```
targeted-broadcast;
```

Hierarchy Level (EX Series, MX Series, ACX Series)

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet]
```

Hierarchy Level (QFX Series, OCX1100, EX4600, NFX Series)

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit interfaces interface-range interface-range-name unit logical-unit-number family inet]
```

Hierarchy Level (SRX Series)

```
[edit interfaces interface-name unit name family inet]
```

Description

Specify the IP packets destined for a Layer 3 broadcast address to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.

The remaining statements are explained separately. See CLI Explorer.

Options

forward-and-send-to-re Allow packets to be forwarded and sent to Routing Engine.

forward-only Allow packets only to be forwarded.

Default

When this statement is not included, broadcast packets are sent to the Routing Engine only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Statement introduced in Junos OS Release 20.4R1 for the SRX Series.

RELATED DOCUMENTATION

Configuring Targeted Broadcast

Understanding Targeted Broadcast

targeted-distribution

IN THIS SECTION

- Syntax | 28958
- Hierarchy Level | 28958
- Description | 28959
- Options | 28959
- Required Privilege Level | 28959
- Release Information | 28959

Syntax

```
targeted-distribution primary-list primary-list |backup-list backup-list;
```

Hierarchy Level

```
[edit logical-systems name interfaces name unit ]
```

Description

Configure egress data for a member link in an aggregated Ethernet bundle. Specify a distribution list as primary list and a different distribution list as backup list. A backup list is provisioned in the event the primary list goes down.

Options

primary-list (Optional) Specify the role of the distribution list as primary. Member links of the aggregated Ethernet are assigned membership to the distribution list.

backup-list (Optional) Specify the role of the distribution list as backup. Member links of the aggregated Ethernet are assigned membership to the distribution list.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

- Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links
- distribution-list

targeted-distribution (Dynamic Demux Interfaces over Aggregated Ethernet)

IN THIS SECTION

Syntax | 28960

- Hierarchy Level | 28960
- Description | 28960
- Required Privilege Level | 28960
- Release Information | 28960

Syntax

```
targeted-distribution;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number]
```

Description

Configure egress data for a dynamic logical interface to be sent across a single member link in an aggregated Ethernet bundle. A backup link is provisioned and CoS scheduling resources are switched to the backup link in the event that the primary assigned link goes down. The aggregated Ethernet interface must be configured without link protection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces

targeted-distribution (Static Interfaces over Aggregated Ethernet)

IN THIS SECTION

- Syntax | 28961
- Hierarchy Level | 28961
- Description | 28961
- Default | 28961
- Options | 28962
- Required Privilege Level | 28962
- Release Information | 28962

Syntax

```
targeted-distribution (replicate | scale);
```

Hierarchy Level

```
[edit interfaces demux0 unit logical-unit-number],
[edit interfaces pp0 unit logical-unit-number]
```

Description

Configure egress data for a logical interface to be sent across a single member link in an aggregated Ethernet bundle. A backup link is provisioned and CoS scheduling resources are switched to the backup link in the event that the primary assigned link goes down. The aggregated Ethernet interface must be configured without link protection.

Default

By default, if you do not include the targeted-scheduler statement, scheduler parameters are applied to the targeted links in the scale mode.

Options

`replicate`—Scheduler parameters are copied to each of the targeted aggregated interface links.

`scale`—Scheduler parameters are scaled based on number of targeted links and applied each of the aggregated interface targeted links.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

- Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces
- Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface
- Targeted Traffic Distribution on Aggregated Ethernet Interfaces in a Virtual Chassis
- Configuring Module Redundancy for a Virtual Chassis
- Configuring Chassis Redundancy for a Virtual Chassis
- *CoS for PPPoE Subscriber Interfaces Overview*

targeted-hello

IN THIS SECTION

- Syntax | 28963
- Hierarchy Level | 28963
- Description | 28963
- Options | 28963
- Required Privilege Level | 28963

Release Information | 28963

Syntax

```
targeted-hello {
    hello-interval seconds;
    hold-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Specify the LDP timer and LDP hold time for targeted hellos.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring the LDP Timer for Hello Messages

Configuring the Delay Before LDP Neighbors Are Considered Down

targeted-options (Grouping Subscribers by Bandwidth Usage)

IN THIS SECTION

- Syntax | 28964
- Hierarchy Level | 28964
- Description | 28965
- Options | 28965
- Required Privilege Level | 28965
- Release Information | 28966

Syntax

```
targeted-options {
    backup backup;
    group group;
    primary primary;
    weight ($junos-interface-target-weight | weight-value);
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit logical-unit-number],
[edit dynamic-profiles name logical-systems name interfaces name unit logical-unit-number],
[edit interfaces name unit logical-unit-number]
```

Description

Configure primary and backup links, group similar subscribers, and specify a subscriber weight for manual targeting to distribute subscribers across aggregated Ethernet member links.

Options

backup (Optional) Specify a backup member link per subscriber when you configure manual targeting.

group (Optional) Assign a group name for subscribers with similar bandwidth usage. Subscribers that are configured for targeted distribution without a group name are added to the `default` group and distributed evenly across member links. Grouping of subscribers is supported only for static subscribers.

- **Default:** default

primary Specify a primary member link per subscriber when you configure manual targeting. You must always configure a primary link when you configure manual targeting.

weight ($junos-interface-target-weight | *weight-value*) Specify the weight for targeted subscribers like PPPoe, demux, and conventional VLANs based on factors such as customer preferences, class of service (CoS), or bandwidth requirement. Member links for logical interfaces of aggregated Ethernet logical interfaces are assigned based on the value of the weight . When a new VLAN is added to the same aggregated Ethernet bundle, then the primary member link selected for targeting is the one with the minimum primary load and the backup link selected for targeting is the one with the minimum overall load.

The $junos-interface-target-weight predefined variable is supported for dynamic configuration only. When you configure this predefined variable, the weight value is sourced from VSA 26-213 in the RADIUS Access-Accept message when a dynamic subscriber is authenticated.

- **Range:** 1 through 1000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`weight` option added in Junos OS Release 17.3 for MX Series and MX Virtual Chassis.

`$junos-interface-target-weight` option added in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

- Understanding Support for Targeted Distribution of Logical Interface Sets of Static VLANs over Aggregated Ethernet Logical Interfaces
- *Using RADIUS-Sourced Weights for Interface and Interface Set Targeted Distribution*
- *RADIUS-Sourced Weights for Interface and Interface Set Targeted Distribution*

targeted-options (Manual Targeting)

IN THIS SECTION

- Syntax | **28966**
- Hierarchy Level | **28967**
- Description | **28967**
- Options | **28967**
- Required Privilege Level | **28968**
- Release Information | **28968**

Syntax

```
targeted-options {
    (logical-interface-chassis-redundancy | logical-interface-fpc-redundancy);
    rebalance-periodic {
        interval interval;
        start-time start-time;
    }
```

```
    rebalance-subscriber-granularity rebalance-subscriber-granularity;
    type (auto | manual);
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-ether-options],
[edit interfaces name aggregated-ether-options]
```

Description

Configure manual targeting or auto-targeting.

Specify the type of targeting to be used for targeted distribution. Specify the targeting option as `manual` for conventional VLAN targeting. By default, the targeting option is `auto`.

Options

type Configure manual targeting type as `manual` or `auto`.

Specify the type of targeting to be used for targeted distribution.

- Values:
  - `auto`—Configure targeted-distribution without specific primary and backup links.
  - `manual`—Configure targeted distribution with specific member links as primary and backup for a subscriber. When you configure manual targeting, you must always configure a primary link. Configuring a backup link is optional. You specify the primary and backup links for a subscriber in the individual interface configuration. You configure primary and backup links by using the "targeted-options" on page 28964 statement at the `[edit interfaces name unit]` hierarchy level.

    Manual targeting enhances the distribution of targeted VLANs or subscribers across member links of an aggregated Ethernet bundle by making it bandwidth-aware.

    Use `manual` keyword to enforce manual targeting on conventional VLANs.
- **Default:** `auto`—By default, targeted option is set to `auto`.

rebalance-subscriber-granularity Rebalancing takes place when the member links have more subscribers than the configured value on the rebalance-subscriber-granularity option.

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Ethernet Interfaces Overview

targeted-options (Grouping Subscribers by Bandwidth Usage) | **28964**

Targeted Traffic Distribution on Aggregated Ethernet Interfaces in a Virtual Chassis

Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links

distribution-list

targeted-distribution | **28958**

targeted-options (PS Interface)

IN THIS SECTION

- Syntax | **28969**
- Hierarchy Level | **28969**
- Description | **28969**
- Required Privilege Level | **28969**
- Release Information | **28969**

Syntax

```
targeted-options {

    logical-interface-chassis-redundancy;
    logical-interface-fpc-redundancy;
    rebalance-periodic {
        interval interval;
        start-time start-time;
    }
    rebalance-subscriber-granularity;
    single-targeted-link;
    type;
}
```

Hierarchy Level

```
[edit interfaces ps name]
```

Description

Configure the option to achieve targeted distribution only for subscriber-interface on pseudowire (PS) IFDs when the PS IFD anchor point is a redundant logical tunnel (RLT), which has multiple LT links in active-active mode (not in link-protection mode).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

MPLS Pseudowire Subscriber Logical Interfaces

targeted-distribution

targets

IN THIS SECTION

- Syntax | 28970
- Hierarchy Level | 28970
- Description | 28970
- Options | 28970
- Required Privilege Level | 28971
- Release Information | 28971

Syntax

```
targets {
    address;
}
```

Hierarchy Level

```
[edit snmp trap-group group-name]
```

Description

Configure one or more systems to receive SNMP traps.

Options

address—IPv4 or IPv6 address of the system to receive traps. You must specify an address, not a hostname.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

targets (Unified Edge)

IN THIS SECTION

- Syntax | 28971
- Hierarchy Level | 28972
- Description | 28972
- Options | 28972
- Required Privilege Level | 28972
- Release Information | 28973

Syntax

```
targets {
    target-name {
        <destination-host hostname>;
        destination-realm realm-name;
        network-element element-name;
        priority priority-value;
    }
}
```

Hierarchy Level

```
[edit unified-edge diameter-profiles gx-profile profile-name]
```

Description

Configure the targets for this Diameter profile.

Options

target-name Name of the target.

destination-host *hostname* (Optional) Use the name of the destination host associated with this target.

destination-realm *realm-name* Use the name of the destination realm associated with this target.

network-element *element-name* Use the name of the network element.

> **NOTE**: The Diameter network element must be previously configured at the `[edit access diameter network-element]` hierarchy level.

- **Range:** 1 through 32 characters

priority *priority-value* Use the specified priority for the target within the Diameter profile. A value with a lower number has a higher priority. For load balancing, configure the targets with the same priority.

> **NOTE**: Failover handling depends on how the policy for the application is configured. For example, switching between the primary and secondary online charging servers set with the appropriate priority can occur only when the failover handling policy is configured to do so.

- **Range:** 1 through 65,535

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

gx-profile | 7954

tca

IN THIS SECTION

- Syntax | 28973
- Hierarchy Level | 28974
- Description | 28974
- Default | 28974
- Options | 28974
- Required Privilege Level | 28976
- Release Information | 28976

Syntax

```
tca tca-identifier (enable-tca | no-enable-tca) (threshold number | threshold-24hrs number)
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

```
[edit interfaces interface-name otn-options]
```

Description

TCAs can give the management system an early indication as to the state of the associated entity when it crosses a certain threshold. TCAs can be set for both minimum and maximum values for gauges and only maximum values for counters. The timely detection of TCAs is essential to proactively manage the interface. TCAs are not an indication of a fault, but rather an indication that the entity may be close to a fault. You can choose which TCAs you want monitored by enabling the TCA. You can either keep the default threshold settings or change the settings.

Enable threshold crossing alerts (TCAs) for the following:

- Optical channel data unit (ODU)
- Optical channel transport unit (OTU)
- Laser power
- Laser temperature

Default

By default, TCAs are not enabled.

Options

tca-identifier —At the otn-options hierarchy level, it can be one of the following:

- `odu-tca-bbe`—ODU background block error threshold-crossing defect trigger
- `odu-tca-bbe-fe`—ODU far-end background block error threshold-crossing defect trigger
- `odu-tca-es`—ODU errored seconds threshold-crossing defect trigger
- `odu-tca-es-fe`—ODU far-end errored seconds threshold-crossing defect trigger
- `odu-tca-ses`—ODU severely errored seconds threshold-crossing defect trigger

- `odu-tca-ses-fe`—ODU far-end severely errored seconds threshold-crossing defect trigger
- `odu-tca-uas`—ODU unavailable seconds threshold-crossing defect trigger
- `odu-tca-uas-fe`—ODU far-end unavailable seconds threshold-crossing defect trigger
- `otu-tca-bbe`—OTU background block error threshold-crossing defect trigger
- `otu-tca-bbe-fe`—OTU far-end background block error threshold-crossing defect trigger
- `otu-tca-es`—OTU errored seconds threshold-crossing defect trigger
- `otu-tca-es-fe`—OTU far-end errored seconds threshold-crossing defect trigger
- `otu-tca-fec-ber`—OTU forward error correction bit error rate threshold-crossing defect trigger
- `otu-tca-ses`—OTU severely errored seconds threshold-crossing defect trigger
- `otu-tca-ses-fe`—OTU far-end severely errored seconds threshold-crossing defect trigger
- `otu-tca-uas`—OTU unavailable seconds threshold-crossing defect trigger
- `otu-tca-uas-fe`—OTU far-end unavailable seconds threshold-crossing defect trigger

tca-identifier —At the optics-options hierarchy level, it can be one of the following:

- `carrier-frequency-offset-high-tca`—Carrier frequency high threshold setting trigger
- `carrier-frequency-offset-low-tca`—Carrier frequency low threshold setting trigger
- `fec-ber`—Optics Errored Seconds Threshold crossing defect trigger
- `fec-corrected-errors-high-tca`—FEC Corrected Error High Threshold crossing defect trigger
- `fec-ucorrected-words-high-tca`—FEC Uncorrected Words High Threshold crossing defect trigger
- `laser-frequency-error-high-tca`—Laser frequency error high TCA
- `laser-frequency-error-low-tca`—Laser frequency error low TCA
- `pam-histogram-high-tca`—PAM Histogram high TCA
- `residual-isi-high-tca`—Residual ISI high TCA
- `residual-isi-low-tca`—Residual ISI low TCA
- `rx-power-high-tca`—Rx power high threshold setting trigger
- `rx-power-low-tca`—Rx power low threshold setting trigger
- `snr-low-tca`—SNR low TCA

- `tec-current-high-tca`—TEC Current high TCA
- `tec-current-low-tca`—TEC Current low TCA
- `temperature-high-tca`—Temperature high threshold setting trigger
- `temperature-low-tca`—Temperature low threshold setting trigger
- `tx-power-high-tca`—Tx power high threshold setting trigger
- `tx-power-low-tca`—Tx power low threshold setting trigger

enable-tca | no-enable-tca—To enable or disable the threshold crossing alert.

threshold | threshold-24hrs:

- threshold *number*—Set the 15-minute interval threshold.
- threshold-24hrs *number*—Set the 24-hour interval threshold.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

optics-options | **11847**

tcp

IN THIS SECTION

- Junos OS Syntax | 28977
- Junos OS Evolved Syntax | 28977
- Junos OS Hierarchy Level | 28977
- Junos OS Evolved Hierarchy Level | 28978
- Description | 28978
- Required Privilege Level | 28978
- Release Information | 28978

Junos OS Syntax

```
tcp {
    destination-interface interface-name;
    port port;
}
```

Junos OS Evolved Syntax

```
tcp {
    port port;
}
```

Junos OS Hierarchy Level

```
[edit services rpm probe-server]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm probe-server]
```

Description

Specify the port information for the TCP server.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved 23.4R1.

RELATED DOCUMENTATION

Configuring RPM Receiver Servers

tcp-aggressive-transmission

IN THIS SECTION

- Syntax | 28979
- Hierarchy Level | 28979
- Description | 28979
- Required Privilege Level | 28979
- Release Information | 28979

Syntax

```
tcp-aggressive-transmission;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor neighbor-
address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor neighbor-address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor neighbor_address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor neighbor-
address]
```

Description

Enables a TCP socket option for the affected BGP sessions, which prioritizes pure ACKs and does not exponentially back-off retransmission for couple of retransmissions.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

tcp-encap

IN THIS SECTION

- Syntax | 28980
- Hierarchy Level | 28980
- Description | 28981
- Options | 28981
- Required Privilege Level | 28981
- Release Information | 28981

Syntax

```
tcp-encap {
    profile profile-name;
        ssl-profile ssl-profile-name;
        log ;
    }
    traceoptions {
        file filename {
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag (all | configuration | session | tunnel);
        level (all | error | info | notice | verbose | warning);
        no-remote-trace’
    }
```

Hierarchy Level

```
[edit security]
```

Description

Specify TCP encapsulation operations for a remote access client to a remote access gateway on an SRX Series Firewall to support IPsec messages encapsulated within a TCP connection.

Options

profile *profile-name*

Configure a TCP encapsulation profile for a remote access client to a remote access gateway on an SRX Series Firewall to define the data encapsulation operation.

- **ssl-profile *ssl-profile-name***

  Specify the SSL termination profile that is configured at the [`edit services ssl termination profile`] hierarchy level. This parameter is required for NCP Exclusive Remote Access Client of Full SSL Session.

- **log**

  Enable logging for remote access client connections.

traceoptions

Configure TCP encapsulation tracing options.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D80.

Support for the `ssl-profile` option added in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

gateway (Security IKE) | 7663

Local User Authentication Using Pre-shared Key (CLI Procedure)

tcp-fast-open

IN THIS SECTION

- Syntax | 28982
- Hierarchy Level | 28982
- Description | 28982
- Default | 28982
- Options | 28983
- Required Privilege Level | 28983
- Release Information | 28983

Syntax

```
tcp-fast-open  {
    disabled;
    drop;
}
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Description

Specify how TCP Fast Open (TFO) enabled packets are to be handled.

Default

By default, all TFO packets are forwarded by the service PIC.

Options

`disabled` Any TCP packet with the TFO option present has the TFO option stripped from the TCP header of the packet, and the rest of the packet is forwarded as is. The benefit of stripping the header over dropping the packet is that the client does not have to wait for the retransmission timer to go off and then retransmit the SYN packet without the TFO option.

`drop` Any TCP packet with the TFO option present is dropped.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Exchanging Data More Efficiently Using TCP Fast Open

tcp-flags

IN THIS SECTION

- Syntax | 28984
- Hierarchy Level | 28984
- Description | 28984
- Options | 28984
- Required Privilege Level | 28985
- Release Information | 28985

Syntax

```
tcp-flags {
    (ack | no-ack);
    (fin | no-fin);
    (psh | no-psh);
    (r1 | no-r1);
    (r2 | no-r2);
    (rst | no-rst);
    (syn | no-syn);
    (urg | no-urg);
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify that IDP looks for a pattern match whether or not the TCP flag is set.

Options

- `ack` | `no-ack`—When set, the acknowledgment flag acknowledges receipt of a packet.
- `fin` | `no-fin`—When set, the final flag indicates that the packet transfer is complete and the connection can be closed.
- `psh` | `no-psh`—When set, the push flag indicates that the receiver should push all data in the current sequence to the destination application (identified by the port number) without waiting for the remaining packets in the sequence.
- `r1` | `no-r1`—When set, indicates that the R1 retransmission threshold has been reached.
- `r2` | `no-r2`—When set, indicates that the R2 retransmission threshold has been reached.
- `rst` | `no-rst`—When set, the reset flag resets the TCP connection, discarding all packets in an existing sequence.
- `syn` | `no-syn`—When set, indicates that the sending device is asking for a three-way handshake to initialize communications.

- `urg` | `no-urg`—When set, the urgent flag indicates that the packet data is urgent.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

tcp-forwarding (Processes)

IN THIS SECTION

- Syntax | 28985
- Hierarchy Level | 28986
- Description | 28986
- Options | 28986
- Required Privilege Level | 28986
- Release Information | 28986

Syntax

```
tcp-forwarding {
    disable;
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
```

```
    }
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Disable the TCP port forwarding process or configure tracing operations for TCP port forwarding events.

Options

disable Disable the TCP port forwarding process, tcpfwdd.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Tracing TCP Port Forwarding Events for Troubleshooting

TCP Port Forwarding for Remote Device Management

tcp-forwarding (Remote Device Management)

IN THIS SECTION

- Syntax | 28987
- Hierarchy Level | 28987
- Description | 28988
- Options | 28988
- Required Privilege Level | 28989
- Release Information | 28989

Syntax

```
tcp-forwarding {
    listening-port port-number listening-address ipv4-listening-address {
        allowed-source ipv4-prefix;
        forwarding-address ipv4-forwarding-address;
        forwarding-port forwarding-port-number;
        max-connections number;
    }

}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services],
[edit logical-systems logical-system-name system services],
[edit routing-instances routing-instance-name system services],
[edit system services]
```

Description

Configure the mapping between the TCP listening address/listening port combination on the BNG and the TCP port forwarding address/port combination where the BNG forwards the incoming data stream. TCP port forwarding is used when the BNG, together with one or more access nodes, is treated as a single addressable point of management by an external management system. The TCP port forwarding connections enable the BNG to demultiplex and multiplex management requests exchanged between the access nodes and the management system.

Options

allowed-source *ipv4-prefix*

(Optional) Restrict the IPv4 prefixes from which TCP connections are accepted on the listening port. The `allowed-source` value is compared to the source address in the TCP header from the triggering entity. When you do not configure an allowed source, TCP connections are accepted from any source prefix.

You can use a /32 IPv4 mask to specify a single address as the source or you can use other masks to specify an IPv4 subnet as the source. You can configure an unlimited number of prefixes for each listening port. To configure multiple sources, you must include the statement multiple times, once for each additional source prefix.

> **NOTE**: You can also configure an unlimited number of allowed-source prefixes across the system.

forwarding-address *ipv4-forwarding-address*

Specify the IPv4 address to which MX BNG must open the second connection of the TCP pair after it opens the first connection triggered on the listening port/listening address combination. All packets received on one connection of the TCP pair are transmitted on the peer (second) connection. This address is used with the forwarding port to open the peer connection.

forwarding-port *forwarding-port-number*

Specify the TCP port of the peer (second) connection of the TCP pair. This port is used with the forwarding address to open the peer connection.

- **Range:** 1 through 65,535

listening-address *ipv4-listening-address*

Specify a particular IPv4 address on the BNG that a triggering entity (an external management or provisioning system or a remote device) must use when attempting to trigger connections on the listening port. You must configure a unique combination of listening port and listening address for each TCP mapping.

listening-port *port-number*

Specify the TCP port that the BNG monitors for connections to be triggered by a remote device or an external management or provisioning system.

- **Range:** 8000 through 8031

max-connections *number*

(Optional) Set a limit on the number of simultaneous TCP connections that the BNG allows on a single listening port. Connection requests received after this limit is reached are rejected.

> **NOTE**: In addition to this per-listening port limit, the system-wide limit for TCP connections is 128 (64 pairs) across all routing instances and listening ports.

- **Range:** 1 through 16
- **Default:** 1

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

TCP Port Forwarding for Remote Device Management

tcp-initial-timeout

IN THIS SECTION

- Syntax | 28990
- Hierarchy Level | 28990
- Description | 28990

- Options | 28990
- Required Privilege Level | 28990
- Release Information | 28991

Syntax

```
tcp-initial-timeout seconds;
```

Hierarchy Level

```
[edit security flow tcp-session]
```

Description

Define the length of time (in seconds) that the device keeps an initial TCP session in the session table before dropping it, or until the device receives a FIN (no more data) or RST (reset) packet. The FIN flag indicates the end of data transmission to finish a TCP connection.

Options

seconds—Number of seconds that the device keeps an initial TCP session in the session table before dropping it.

- **Range:** 4 through 300 seconds
- **Default:** 20 seconds

> **NOTE**: The minimum value you can configure for TCP session initialization is 4 seconds. The default value is 20 seconds; if required you can set the TCP session initialization value to less than 20 seconds.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

tcp-keepcnt

IN THIS SECTION

- Syntax | 28991
- Hierarchy Level | 28991
- Description | 28992
- Default | 28992
- Required Privilege Level | 28992
- Release Information | 28992

Syntax

```
tcp-keepcnt number;
```

Hierarchy Level

```
[edit services rpm twamp (client control-connection control-client-name |server)]
```

Description

Number of unacknowledged probes to send before considering the connection dead and notifying the application layer. The range is 1 through 50.

Default

The default number of TCP KEEPALIVEs sent is 6.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

Understanding Two-Way Active Measurement Protocol on Routers

tcp-keepidle

IN THIS SECTION

- Syntax | 28993
- Hierarchy Level | 28993
- Description | 28993
- Default | 28993
- Required Privilege Level | 28993
- Release Information | 28993

Syntax

```
tcp-keepidle seconds;
```

Hierarchy Level

```
[edit services rpm twamp (client control-connection control-client-name |server)]
```

Description

Time interval between the last data packet sent and the first keepalive probe. The range is 1 through 600 seconds.

Default

The default value is 120 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

> *Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches*
>
> *Understanding Two-Way Active Measurement Protocol on Routers*

tcp-keepintvl

IN THIS SECTION

- Syntax | 28994
- Hierarchy Level | 28994
- Description | 28994
- Default | 28994
- Required Privilege Level | 28994
- Release Information | 28995

Syntax

```
tcp-keepintvl seconds;
```

Hierarchy Level

```
[edit services rpm twamp (client control-connection control-client-name |server)]
```

Description

Time interval between successive keepalive probes. The range is 1 second through 600 seconds.

Default

The default value is 5 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

Understanding Two-Way Active Measurement Protocol on Routers

tcp-log

IN THIS SECTION

- Syntax | 28995
- Hierarchy Level | 28995
- Description | 28995
- Required Privilege Level | 28996

Syntax

```
set services service-set ss1 syslog host server-IP tcp-log
```

Hierarchy Level

```
[edit services service-set]
```

Description

Configure TCP/TLS for logging syslog events and notifications.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

tcp-mss (Internet Options)

IN THIS SECTION

- Syntax | 28996
- Hierarchy Level | 28996
- Description | 28996
- Options | 28997
- Required Privilege Level | 28997
- Release Information | 28998

Syntax

```
tcp-mss bytes;
```

Hierarchy Level

```
[edit system internet-options]
```

Description

Enable and specify the TCP maximum segment size (MSS) for IPV4 traffic in bytes.

When the device receives a TCP packet with the SYN bit and MSS option set, it examines the MSS value. If the MSS option specified in the packet is larger than the MSS you configured with the `tcp-mss` statement, the router replaces the MSS value in the packet with the lower value you configured.

> **NOTE**: This configuration statement is for transit packets. To adjust the TCP MSS for host-originated packets, see "tcp-mss (Protocols BGP)" on page 28998.

There are multiple factors that define the MSS value for TCP packets, which are reflected in the MSS value displayed in the output of the `show system connections` command:

- The MSS value offered by the peer in the SYN packet.
- Rounding the MSS off to the nearest multiple of 2048.
- The MTU value of the interface.
- The configured path MTU value.
- Whether TCP sessions are not directly connected and path MTU discovery is disabled.
- Whether the TCP sessions are on a directly-connected network.

This statement enables you to specify the MSS size in TCP SYN packets used during session establishment. Decreasing the MSS size helps to limit packet fragmentation and to protect against packet loss that can occur when a packet must be fragmented to meet the MTU size but the packet's DF (don't fragment) bit is set.

> **NOTE**: SRX Series Firewalls running in packet mode with MPLS do not support TCP MSS.

Options

bytes TCP MSS value for SYN packets with a higher MSS value set.

- **Range:** 64 through 65535 bytes.
- **Default:** TCP MSS is disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. Use Feature Explorer to confirm platform and release support for this feature.

RELATED DOCUMENTATION

Configure TCP Options

tcp-mss (Protocols BGP)

IN THIS SECTION

- Syntax | 28998
- Hierarchy Level | 28998
- Description | 28999
- Options | 28999
- Required Privilege Level | 29000
- Release Information | 29000

Syntax

```
tcp-mss segment-size;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor neighbor-name],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor neighbor-
name]
```

Description

Configure the maximum segment size (MSS) for the TCP connection for BGP neighbors. You can set the MSS globally or for a particular neighbor. If the router receives a TCP packet with the SYN bit and the MSS option set, and the MSS option specified in the packet is larger than the MSS value specified by the `tcp-mss` statement, the router replaces the MSS value in the packet with the lower value specified by the `tcp-mss` statement.

The MSS is only valid in increments of 2 KB. The value used is based on the value set, but is rounded down to the nearest multiple of 2048.

> **NOTE**: This configuration statement is for host-originated packets. To configure the TCP MSS value for transit traffic on a particular interface, see tcp-mss (Interfaces) .

There are multiple factors that define the MSS value for TCP packets, which are reflected in the MSS value displayed in the output of the `show system connections` command:

- The MSS value offered by the peer in the SYN packet.
- Rounding the MSS off to the nearest multiple of 2048.
- The MTU value of the interface.
- The configured path MTU value.
- Whether TCP sessions are not directly connected and path MTU discovery is disabled.
- Whether the TCP sessions are on a directly-connected network.

Options

segment-size	MSS for the TCP connection.

- **Range:** 1 through 4096

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Example: Limiting TCP Segment Size for BGP

tcp-mss (Security Flow)

IN THIS SECTION

- Syntax | 29000
- Hierarchy Level | 29001
- Description | 29001
- Options | 29002
- Required Privilege Level | 29002
- Release Information | 29002

Syntax

```
tcp-mss {
    all-tcp mss value;
    gre-in {
        mss value;
```

```
    }
    gre-out {
        mss value;
    }
    ipsec-vpn {
        mss value;
    }
}
```

Hierarchy Level

```
[edit security flow]
```

Description

Configure TCP maximum segment size (TCP MSS) for the following packet types:

- All TCP packets for network traffic.
- GRE packets entering the IPsec VPN tunnel.
- GRE packets exiting the IPsec VPN tunnel.
- TCP packets entering the IPsec VPN tunnel.

If all the four TCP MSS options are configured simultaneously, then the order of preference is as follows:

- If TCP packet enters an IPsec VPN tunnel, then an ipsec-vpn mss value has high priority over all-tcp mss value, hence ipsec-vpn mss value is set.
- If TCP packet enters GRE , then gre-in mss value overrides all-tcp mss value, hence gre-in mss value is set.
- If TCP packet exits GRE, then all-tcp mss value overrides gre-in mss value, hence all-tcp mss value is set.

TCP MSS may not override in GRE over IPsec (GREoIPsec) scenarios. Consequently, there may be more fragmentation in the network as GREoIPsec traffic isn't modified for TCP MSS. To ensure that TCP MSS works with GREoIPsec, set the priority of MSS applied to the TCP traffic in the following order (highest to lowest):

1. `gre-in`
2. `gre-out` based on direction of the GREoIPSec TCP traffic

3. `ipsec-vpn` for GREoIPsec
4. IPsec traffic
5. `all-tcp` for all the TCP traffic.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

- all-tcp | 2159
- gre-in | 7793
- gre-out | 7794
- ipsec-vpn (Security Flow)
- *Understanding Traffic Processing on Security Devices*

tcp-mss (Services)

IN THIS SECTION

- Syntax | 29003
- Hierarchy Level | 29003
- Description | 29003
- Options | 29003

Required Privilege Level | 29003
Release Information | 29003

Syntax

```
tcp-mss number;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the TCP Maximum Segment Size (MSS) allowed for the service set.

Options

number—MSS value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Service Set Limitations

tcp-no-flag

IN THIS SECTION

- Syntax | 29004
- Hierarchy Level | 29004
- Description | 29004
- Required Privilege Level | 29004
- Release Information | 29004

Syntax

```
tcp-no-flag;
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp]
```

Description

Enable the device to drop illegal TCP packets with a missing or malformed flag field.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

tcp-no-flag (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 29005
- Hierarchy Level | 29005
- Description | 29005
- Required Privilege Level | 29005
- Release Information | 29005

Syntax

```
tcp-no-flag;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Description

Identify and drop TCP packets that have no flag fields set. A TCP no flag attack can cause unpredictable behavior on the target.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

tcp-non-syn

IN THIS SECTION

- Syntax | 29006
- Hierarchy Level | 29006
- Description | 29006
- Options | 29007
- Required Privilege Level | 29007
- Release Information | 29007

Syntax

```
tcp-non-syn  {
    drop-flow;
    drop-flow-send-rst;
}
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Description

Specify how the first non-SYN TCP packet is processed on services PICs. When a services PIC receives the first non-SYN TCP packet for processing, the packet is dropped.

Options

`drop-flow` When a services PIC receives the first non-SYN TCP packet for processing, the packet is dropped.

A drop flow created on the services PIC ensures that subsequent non-SYN TCP packets with the same 5-tuple information (source and destination addresses, protocol, and source and destination ports) are dropped. If this statement is not configured, a session is created when a packet hits the services set and matches the stateful firewall rule even if the packet is a non-SYN packet.

`drop-flow-send-rst` When a services PIC receives the first non-SYN TCP packet for processing, the packet is dropped and a reset packet is sent to originator to ensure that no further packets are generated.

A drop flow created on the services PIC ensures that subsequent non-SYN TCP packets with the same 5-tuple information (source and destination addresses, protocol, and source and destination ports) are dropped. If this statement is not configured, a session is created when a packet hits the services set and matches the stateful firewall rule even if the packet is a non-SYN packet.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R2.

RELATED DOCUMENTATION

Configuring Service Sets to be Applied to Services Interfaces

tcp-options (Security Policies)

IN THIS SECTION

- Syntax | 29008
- Hierarchy Level | 29008
- Hierarchy Level | 29008
- Description | 29009
- Options | 29010
- Required Privilege Level | 29011
- Release Information | 29011

Syntax

```
tcp-options {
    initial-tcp-mss mss-value;
    reverse-tcp-mss mss-value;
    sequence-check-required;
    syn-check-required;
    window-scale;
}
```

Hierarchy Level

Hierarchy Level

```
[edit logical-systems logical-systems-name security policies explicit-proxy profile name policy
policy-name from then permit]
[edit logical-systems logical-systems-name security policies-zone  from-zone zone-name to-zone
zone-name policy policy-name then permit]
[edit logical-systems logical-systems-name security policies global policy policy-name then
permit]
```

```
[edit logical-systems logical-systems-name  security policies policy-set policy-policy-set-name
policy policy-name then permit]
[edit logical-systems logical-systems-name security policies policy-set policy-set-name policy
policy-name then permit]

[edit logical-systems logical-systems-name tenants tenants-tenants-systems-name security policies
explicit-proxy profile policy-name policy policy-name from then]
[edit logical-systems logical-systems-name tenants tenants-systems-name security policies-zone
from-zone zone-name to-zone zone-name policy policy-name then permit]
[edit logical-systems logical-systems-name tenants tenants-tenants-systems-name security
policies  global policy policy-name then permit]
[edit logical-systems logical-systems-name tenants tenants-tenants-systems-name security policies
policy-setset policy-set-name policy policy-name then permit]
[edit security policies explicit-proxy profile policy-set-policy-name policy policy-name then
permit]
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit]
[edit security policies policiesglobal from-zone zone-name to-zone zone-name policy policy-name
then permit  then permit]
[edit security policies policies policy-set policy-policy-set-name policy policy-name then
permit]
[edit tenants tenants-systems-name securitypolicies policies policyexplicit-setproxy profile
policy-set-name policy policy-name then permit]
[edit tenants tenants-systems-name security policies from-zone zone-name to-zone zone-name policy
policy-name then permit]
[edit tenants tenants-systems-name security policies global policy policy-name then permit]
[edit tenants tenants-systems-name security policies policy-set policy-set-name policy policy-
name then permit]
```

Description

Specify the TCP options for each policy. You can configure sync and sequence checks for each policy based on your requirements, and, because each policy has two directions, you can configure a TCP MSS value for both directions or for just one direction. To configure per-policy TCP options, you must turn off the respective global options. Otherwise, the commit check will fail.

NOTE: TCP maximum segment size (MSS) does not have a default value. The value is determined either from policy configuration or from flow configuration.

Example:

[set security flow tcp-mss all-tcp mss *mss-value*]

The priority for setting the MSS is given to the policy-level configuration. If there are no configurations at both the policy and flow levels, then the MSS indicated by the respective endpoint will remain as it is.

See tcp-mss (Security Flow) for details.

NOTE: TCP maximum segment size (MSS) does not have a default value. The value is determined either from policy configuration or from flow configuration.

Example:

[set security flow tcp-mss all-tcp mss <*mss-value*>]

The priority for setting the MSS is given to the policy-level configuration. If there are no configurations at both the policy and flow levels, then the MSS indicated by the respective endpoint will remain as it is.

See tcp-mss (Security Flow) for details.

Options

initial-tcp-mss

Configure the TCP maximum segment size (MSS) for packets that arrive at the ingress interface (initial direction), match a specific policy, and for which a session is created. The value you configure overrides the TCP MSS value in the incoming packet when the value in the packet is higher than the one you specify.

The `initial-tcp-mss` value per policy takes precedence over a global `tcp-mss` value (`all-tcp`, `ipsec-vpn`, `gre-in`, `gre-out`), if one is configured. However, when the `syn-flood-protection-mode syn-proxy` statement at the `[edit security flow]` hierarchy level is used to enable SYN proxy defenses against SYN attacks, the TCP MSS value is not overriden.

Because each policy has two directions, you can configure a value for both directions or for just one direction. To configure a TCP MSS value for the reverse session, use the `reverse-tcp-mss` option.

- **Range:** 64 through 65535

reverse-tcp-mss

Configure the TCP maximum segment size (MSS) for packets that match a specific policy and travel in the reverse direction of a session. The value you configure replaces the TCP MSS value when the value in the packet is higher than the one you specify.

The `reverse-tcp-mss` value per policy takes precedence over a global `tcp-mss` value (`all-tcp`, `ipsec-vpn`, `gre-in`, `gre-out`), if one is configured. However, when the `syn-flood-protection-mode syn-proxy` statement at the `[edit security flow]` hierarchy level is used to enable SYN proxy defenses against SYN attacks, the TCP MSS value is not overridden.

Because each policy has two directions, you can configure a value for both directions or for just one direction. To configure the TCP MSS value for the initial session, use the `initial-tcp-mss` option.

- **Range:** 64 through 65535

sequence-check-required Enable sequence check per policy. The sequence-check-required value overrides the global value no-sequence-check.

syn-check-required Enable sync check per policy. The syn-check-required value overrides the global value no-syn-check.

window-scale Enable window-scale per policy.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The `tcp-options`, `sequence-check-required`, and `syn-check-required` statements are introduced in Junos OS Release 10.4.

The `initial-tcp-mss` and `reverse-tcp-mss` statements are introduced in Junos OS Release 12.3X48-D20.

The `window-scale` statement is introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

tcp-mss (Security Flow)

tcp-rst

IN THIS SECTION

- Syntax | 29012
- Hierarchy Level | 29012
- Description | 29012
- Default | 29012
- Required Privilege Level | 29012
- Release Information | 29013

Syntax

```
tcp-rst;
```

Hierarchy Level

```
[edit security zones security-zone zone-name]
```

Description

Enable the device to send a TCP segment with the RST (reset) flag set to 1 (one) in response to a TCP segment with any flag other than SYN set and that does not belong to an existing session.

During flow first path process, a TCP RST packet is sent to the traffic originator if the TCP packet trying to create the flow session is not a SYN packet.

Default

By default, `tcp-rst` is disabled.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring the TCP-Reset Parameter

tcp (Security IDP Protocol Binding)

IN THIS SECTION

- Syntax | 29013
- Hierarchy Level | 29014
- Description | 29014
- Options | 29014
- Required Privilege Level | 29014
- Release Information | 29014

Syntax

```
tcp {
    minimum-port port-number <maximum-port port-number>;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain protocol-binding],

[edit security idp custom-attack attack-name attack-type signature protocol-binding]
```

Description

Allow IDP to match the attack for specified TCP ports.

Options

`minimum-port` *`port-number`*—Minimum port in the port range.

- **Range:** 0 through 65,535

`maximum-port` *`port-number`*—Maximum port in the port range.

- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

tcp (Security IDP Signature Attack)

IN THIS SECTION

- Syntax | 29015
- Hierarchy Level | 29016
- Description | 29016

- Options | 29016
- Required Privilege Level | 29017
- Release Information | 29017

Syntax

```
tcp {
    ack-number {
        match (equal | greater-than | less-than | not-equal);
        value acknowledgement-number;
    }
    data-length {
        match (equal | greater-than | less-than | not-equal);
        value tcp-data-length;
    }
    destination-port (Security Signature Attack) {
        match (equal | greater-than | less-than | not-equal);
        value destination-port;
    }
    header-length {
        match (equal | greater-than | less-than | not-equal);
        value header-length;
    }
    mss (Security IDP) {
        match (equal | greater-than | less-than | not-equal);
        value maximum-segment-size;
    }
    option (Security IDP) {
        match (equal | greater-than | less-than | not-equal);
        value tcp-option;
    }
    reserved (Security IDP Custom Attack) {
        match (equal | greater-than | less-than | not-equal);
        value reserved-value;
    }
    sequence-number (Security IDP ICMP Headers) {
        match (equal | greater-than | less-than | not-equal);
        value sequence-number;
```

```
    }
    source-port (Security IDP) {
        match (equal | greater-than | less-than | not-equal);
        value source-port;
    }
    tcp-flags {
        (ack | no-ack);
        (fin | no-fin);
        (psh | no-psh);
        (r1 | no-r1);
        (r2 | no-r2);
        (rst | no-rst);
        (syn | no-syn);
        (urg | no-urg);
    }
    urgent-pointer {
        match (equal | greater-than | less-than | not-equal);
        value urgent-pointer;
    }
    window-scale {
        match (equal | greater-than | less-than | not-equal);
        value window-scale-factor;
    }
    window-size {
        match (equal | greater-than | less-than | not-equal);
        value window-size;
    }
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol]
```

Description

Allow IDP to match the TCP header information for the signature attack.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

tcp (Security Screen)

IN THIS SECTION

- Syntax | 29017
- Hierarchy Level | 29018
- Description | 29018
- Options | 29018
- Required Privilege Level | 29018
- Release Information | 29019

Syntax

```
tcp {
    fin-no-ack;
    land;
    port-scan {
        threshold number;
    }
    syn-ack-ack-proxy {
        threshold number;
    }
    syn-fin;
    syn-flood {
        alarm-threshold number;
```

```
        attack-threshold number;
        destination-threshold number;
        source-threshold number;
        timeout seconds;
        white-list name {
            destination-address destination-address;
            source-address source-address;
        }
    }
    syn-frag;
    tcp-no-flag;
    tcp-sweep {
        threshold threshold number;
    }
    winnuke;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name]
```

Description

Configure TCP-layer intrusion detection service (IDS) options.

> **NOTE**: On all SRX Series Firewalls, the TCP synchronization flood alarm threshold value does not indicate the number of packets dropped, however the value does show the packet information after the alarm threshold has been reached.
>
> The synchronization cookie or proxy never drops packets; therefore the `alarm-without-drop` (not `drop`) action is shown in the system log.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

tcp-seq-ooo-window (Security Forwarding Options)

IN THIS SECTION

- Syntax | 29019
- Hierarchy Level | 29019
- Description | 29019
- Required Privilege Level | 29020
- Release Information | 29020

Syntax

```
tcp-seq-ooo-window Kbytes;
```

Hierarchy Level

```
[edit security forwarding-options services-offload]
```

Description

Specify the TCP sequence out of order window threshold (Kbytes) for Service offload. The configured threshold range is from 0 Kbytes to 1023 Kbytes

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1

RELATED DOCUMENTATION

TCP Sessions

tcp-session

IN THIS SECTION

- Syntax | 29020
- Hierarchy Level | 29021
- Description | 29021
- Options | 29022
- Required Privilege Level | 29022
- Release Information | 29022

Syntax

```
tcp-session {
    aggressive-tcp-ack;
    fin-invalidate-session;
    maximum-window(128K | 1M | 256K | 512K | 64K);
    no-sequence-check;
    no-syn-check;
    no-syn-check-in-tunnel;
```

```
    rst-invalidate-session;
    rst-sequence-check;
    strict-syn-check;
    tcp-fast-open;
    tcp-initial-timeout seconds;
    tcp-selective-ack;
    time-wait-state {
        (session-ageout | session-timeout seconds);
        }
}
```

Hierarchy Level

```
[edit security flow]
```

Description

Configure TCP session attributes:

- Aggressive TCP ACK
- FIN invalidate session—To Immediately end the session on receipt of fin (FIN) segment.
- Maximum window—Maximum TCP proxy scaled receive window, default 256K bytes.
- TCP sequence number checking.
- TCP SYN bit checking.
- Reset (RST) checking.
- Initial TCP session timeout—The minimum value you can configure for TCP session initialization is 4 seconds. The default value is 20 seconds; if required you can set the TCP session initialization value to less than 20 seconds.
- Strict TCP SYN checking.
- TCP fast open (FSO) for TCP proxy session.
- TCP session timeout for time-wait state.
- TCP selective ack for TCP proxy session.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

tcp-session-install-interval (Security Forwarding Options)

IN THIS SECTION

- Syntax | 29022
- Hierarchy Level | 29023
- Description | 29023
- Required Privilege Level | 29023
- Release Information | 29023

Syntax

```
tcp-session-install-interval mseconds;
```

Hierarchy Level

```
[edit security forwarding-options services-offload]
```

Description

Specify the timeout value for TCP install NP session (mseconds). The configured interval range is from 50 msec to 30,000 msec.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1

RELATED DOCUMENTATION

TCP Sessions

tcp-session (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 29024
- Hierarchy Level | 29024
- Description | 29024
- Options | 29024
- Required Privilege Level | 29025

Release Information | 29025

Syntax

```
tcp-session {
    inactivity-asymm-tcp-timeout ;
    inactivity-tcp-timeout ;
    open-timeout ;
    tcp-fast-open ;
    tcp-mss ;
    tcp-non-syn ;
    tcp-tickles ;
}
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Description

Configure the TCP options for the service set.

Options

close-timeout	Timeout period for TCP session tear-down (2. through 300 seconds)
ignore-errors	Ignore anomalies or errors for TCP
"inactivity-asymm-tcp-timeout " on page 8500	
"tcp-tickles" on page 29030	Number of TCP keep-alive packets to be sent for bidirectional TCP flows
inactivity-tcp-timeout	Inactivity timeout period for TCP established sessions
open-timeout	Timeout period for TCP session establishment (seconds)

tcp-fast-open	Tcp-fast-Open enabled packets will be handled accordingly
tcp-mss	Enable the limit on TCP Max. Seg. Size in SYN packets
tcp-non-syn	Deny session creation on receiving first non SYN packet

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

tcp-sweep

IN THIS SECTION

- Syntax | 29025
- Hierarchy Level | 29026
- Description | 29026
- Options | 29026
- Required Privilege Level | 29026
- Release Information | 29026

Syntax

```
tcp-sweep {
    threshold number;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp]
```

Description

Configure the device to detect and prevent TCP sweep attack. In a TCP sweep attack, an attacker sends TCP SYN packets to the target device as part of the TCP handshake. If the device responds to those packets, the attacker gets an indication that a port in the target device is open, which makes the port vulnerable to attack. If a remote host sends TCP packets to 10 addresses in 0.005 seconds (5000 microseconds), then the device flags this as a TCP sweep attack.

If the `alarm-without-drop` option is not set, the device rejects the eleventh and all further TCP packets from that host for the remainder of the specified threshold period.

Options

`threshold` *`number`*—Maximum number of microseconds during which up to 10 TCP SYN packets from the same host are allowed into the device. More than 10 requests from a host during this period triggers TCP Sweep attack response on the router during the remainder of the second.

- **Range:** 1000 through 1,000,000 microseconds
- **Default:** 5000 microseconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

tcp-syn-defense (IDS MS-MPC)

IN THIS SECTION

- Syntax | 29027
- Hierarchy Level | 29027
- Description | 29027
- Required Privilege Level | 29027
- Release Information | 29027

Syntax

```
tcp-syn-defense;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Close unestablished TCP connections and send a TCP RST to the end host to clear the TCP states on it when the `open-timeout` value at the `[edit interfaces interface-name service-options]` hierarchy level expires. This provides protection against TCP SYN flooding attacks. This statement can only be used in IDS rules assigned to a service set on an MS-MPC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

tcp-syn-fragment-check (IDS MS-MPC)

IN THIS SECTION

- Syntax | 29028
- Hierarchy Level | 29028
- Description | 29028
- Required Privilege Level | 29028
- Release Information | 29029

Syntax

```
tcp-syn-fragment-check;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Identify and drop TCP SYN packets that are IP fragments. This statement can only be used in IDS rules assigned to a service set on an MS-MPC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

tcp-tickles

IN THIS SECTION

- Syntax | 29029
- Hierarchy Level | 29029
- Description | 29029
- Options | 29030
- Required Privilege Level | 29030
- Release Information | 29030

Syntax

```
tcp-tickles tcp-tickles;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Description

Define the maximum number of keep-alive messages sent before a TCP session is allowed to timeout.

Options

tcp-tickles—Number of keep-alive messages.

- **Range:** 0 through 30
- **Default:** 4

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Default Timeout Settings for Services Interfaces

tcp-tickles (Service Set Next Gen Services)

IN THIS SECTION

- Syntax | 29031
- Hierarchy Level | 29031
- Description | 29031
- Required Privilege Level | 29031
- Release Information | 29031

Syntax

```
tcp-tickles tcp-tickles;
```

Hierarchy Level

```
[edit service-set service-set-name service-set-options tcp-session]
```

Description

Define the maximum number of keep-alive messages sent before a TCP session is allowed to timeout.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

tcp-winnuke-check (IDS MS-MPC)

IN THIS SECTION

- Syntax | 29032
- Hierarchy Level | 29032
- Description | 29032
- Required Privilege Level | 29032
- Release Information | 29032

Syntax

```
tcp-winnuke-check;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then]
```

Description

Identify and drop TCP segments that are destined for port 139 and have the urgent (URG) flag set, which provides protection against WinNuke attacks. This statement can only be used in IDS rules assigned to a service set on an MS-MPC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> *Configuring Protection Against Network Attacks on an MS-MPC*
>
> *Understanding IDS on an MS-MPC*

tcp-mss (Interfaces)

IN THIS SECTION

- Syntax | 29033
- Hierarchy Level | 29033
- Description | 29033
- Options | 29034
- Required Privilege Level | 29034
- Release Information | 29034

Syntax

```
tcp-mss tcp-mss;
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit name family inet],
[edit dynamic-profiles name interfaces name unit name family inet6],
[edit dynamic-profiles name logical-systems name interfaces name unit name family inet],
[edit dynamic-profiles name logical-systems name interfaces name unit name family inet6],
[edit interfaces name unit name family inet],
[edit interfaces name unit name family inet6]
```

Description

Enable and specify the TCP maximum segment size (MSS) in bytes for transit packets in the specified protocol family. The TCP MSS value must be lower than the MTU of the interface.

When the device receives a TCP packet with the SYN bit and MSS option set, it examines the MSS value. If the MSS option specified in the packet is larger than the MSS you configured with the `tcp-mss` statement, the router replaces the MSS value in the packet with the lower value you configured.

You cannot configure TCP MSS on service interfaces (ms, vms, or ams).

> **NOTE**: This configuration statement is for transit packets. To adjust the TCP MSS for host-originated packets, see "tcp-mss (Protocols BGP)" on page 28998.

Options

bytes TCP MSS value for SYN packets with a higher MSS value set.

- **Range:** 64 through 65535 bytes.
- **Default:** TCP MSS is disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Configure TCP Options

tcpao-auth-mismatch

IN THIS SECTION

- Syntax | 29035
- Hierarchy Level | 29035
- Description | 29035
- Options | 29035
- Required Privilege Level | 29035

Release Information | 29036

Syntax

```
tcpao-auth-mismatch {
    allow-without-tcpao;
}
```

Hierarchy Level

```
[edit fabric protocols bgp]
[edit logical-systems name protocols bgp]
[edit logical-systems name routing-instances name protocols bgp]
[edit logical-systems name tenants name routing-instances name protocols bgp]
[edit routing-instances name protocols bgp]
[edit tenants name routing-instances name protocols bgp]
[edit protocols bgp group group neighbor neighbor]
[edit protocols bgp]
[edit protocols ldp session session]
```

Description

Continue without TCP-AO if any one TCP endpoint does not have TCP-AO configured.

Options

allow-without-tcpao Allow the connection establishment without TCP-AO.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

TCP Authentication Option (TCP-AO)

tdf-interface

IN THIS SECTION

- Syntax | 29036
- Hierarchy Level | 29036
- Description | 29036
- Options | 29037
- Required Privilege Level | 29037
- Release Information | 29037

Syntax

```
tdf-interface mif.number;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
```

Description

Specify the TDF interface that the TDF domain uses. A TDF interface is different from other types of interfaces, and is associated with the TDF service set that is used for the TDF subscriber.

> **NOTE**: The TDF interface must have been previously configured at the `[edit interfaces]` hierarchy level.
>
> The TDF interface and the access-facing interfaces in the TDF domain must be included in the same VRF routing instance.

Options

mif.*number* Use the specified TDF interface unit number.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain*
- *Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain*
- *IP-Based Subscriber Setup Overview*
- *Configuring TDF Interface to Access Interface Associations in VRFs*
- *Configuring TDF Interface to Access Interface Associations in VRFs*
- *Configuring a TDF Logical Interface*

tdf (Unified Edge)

IN THIS SECTION

Syntax | 29038

- Hierarchy Level | 29038
- Description | 29038
- Options | 29038
- Required Privilege Level | 29038
- Release Information | 29038

Syntax

```
tdf gateway-name;
```

Hierarchy Level

```
[edit unified-edge gateways]
```

Description

Specify the name to be used for the traffic detection function (TDF) gateway.

Options

gateway-name

Name of the gateway.

- **Range:** 1 through 16 characters.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

TDF Gateway Service PICs and Session PICs for Subscriber-Aware Traffic Treatment

tdm-options (Interfaces)

IN THIS SECTION

- Syntax | 29039
- Hierarchy Level | 29039
- Description | 29040
- Options | 29040
- Required Privilege Level | 29040
- Release Information | 29040

Syntax

```
tdm-options {
    ces-psn-channel ...
        ces-psn-port-dmac-check-enable ;
        iwf-params...
        sfp-type sfptype;
        tdm-in-loop ;
    tdm-out-loop ;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure TDM options to enable the smart SFP to encapsulate PDH (Plesiochronous Digital Hierarchy) and SDH (synchronous Digital Hierarchy) traffic over packet switched networks (PSNs).

Options

ces-psn-port-dmac-check-enable Enables checking of the destination MAC address of the incoming packets on the receiving SFP. If you have configured the destination MAC address using the `dmac-address` option, use this option to verify the MAC address on the receiving SFP. If you have enabled MAC address verification and the MAC address does not match, the packet is discarded by the smart SFP.

sfp-type *sfp-type* Specify the smart SFP type you want to configure on the interface. Values: T1, E1, DS3, STM1, STM4, and STM16.

tdm-in-loop Enables looping back of the input path (Rx) of TDM traffic on the SFP TDM port.

tdm-out-loop Enables looping back of the output path (Tx) of TDM traffic on the SFP TDM port.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4.

RELATED DOCUMENTATION

show interfaces smart-sfp-statistics | 21372

show interfaces smart-sfp-defects | 21359

te-class-matrix

IN THIS SECTION

- Syntax | 29041
- Hierarchy Level | 29041
- Description | 29041
- Default | 29042
- Options | 29042
- Required Privilege Level | 29042
- Release Information | 29042

Syntax

```
te-class-matrix {
    te number {
        priority priority;
        traffic-class {
            ct number priority priority;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls diffserv-te],
[edit protocols mpls diffserv-te]
```

Description

Specify the traffic engineering class matrix for a multiclass LSP or a DiffServ-aware traffic engineering LSP.

Default

The default traffic engineering class matrix is:

```
te-class-matrix {
    te0 traffic-class ct0 priority 7;
    te1 traffic-class ct1 priority 7;
    te2 traffic-class ct2 priority 7;
    te3 traffic-class ct3 priority 7;
    te4 traffic-class ct0 priority 0;
    te5 traffic-class ct1 priority 0;
    te6 traffic-class ct2 priority 0;
    te7 traffic-class ct3 priority 0;
}
```

If you define any of the traffic engineering classes, all the default values are dropped.

Options

`ct`*`number`*—Specify the number of the class type. It can be one of four values: `ct0`, `ct1`, `ct2`, or `ct3`.

`priority` *`priority`*—Specify the priority of the class type. It can be one of eight values from 0 through 7.

`te`*`number`*—Specify the number of the traffic engineering class. It can be one of eight values: `te0`, `te1`, `te2`, `te3`, `te4`, `te5`, `te6`, or `te7`. You must configure the traffic engineering classes in order, starting with `te0`.

`traffic-class`—Specify the traffic class for the traffic engineering class.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Traffic Engineering Classes

te-link

IN THIS SECTION

- Syntax | 29043
- Hierarchy Level | 29043
- Description | 29044
- Options | 29044
- Required Privilege Level | 29044
- Release Information | 29044

Syntax

```
te-link te-link-name {
    disable;
    ethernet-vlan;
    interface interface-name {
        disable;
        local-address ip-address;
        remote-address ip-address;
        remote-id id-number;
    }
    local-address ip-address;
    remote-address ip-address;
    remote-id id-number;
}
```

Hierarchy Level

```
[edit protocols link-management],
[edit protocols link-management peer peer-name]
```

Description

Represent a collection of physical ports or time slots. Assign a traffic engineering link to the specified network peer.

Options

te-link-name—Name of the collection of physical ports or the name of the time slots.

`disable`—Disable the traffic engineering link or an interface to a traffic engineering link.

The remaining statements are explained separately..

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`ethernet-vlan` option introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring LMP Traffic Engineering Links

te-policy

IN THIS SECTION

- Syntax | 29045
- Hierarchy Level | 29045
- Description | 29045
- Options | 29045

- Required Privilege Level | **29046**

Syntax

```
te-policy {
                apply-groups
                apply-groups-except
                groups
                node-as
                router-id
                spring-te
                |
                    }
```

Hierarchy Level

```
[edit policy-options policy-statement name from traffic-engineering]
```

Description

Specify traffic engineering related parameters.

Options

node-as Node AS number

router-id Router ID to match against the node added by TE policy

spring-te SPRING-TE policy related parameters

Required Privilege Level

routing

RELATED DOCUMENTATION

Link-State Distribution using BGP

tear-drop (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 29046
- Hierarchy Level | 29046
- Description | 29046
- Required Privilege Level | 29047
- Release Information | 29047

Syntax

```
tear-drop;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop fragmented IP packets that overlap, which protects against teardrop attacks. In teardrop attacks, the target machine uses up its resources as it attempts to reassemble the packets, and then it can no longer process valid traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

tei-option

IN THIS SECTION

- Syntax | 29047
- Hierarchy Level | 29047
- Description | 29048
- Options | 29048
- Required Privilege Level | 29048
- Release Information | 29048

Syntax

```
tei-option (first-call | power-up);
```

Hierarchy Level

```
[edit interfaces br-pim/0/portisdn-options]
```

Description

For ISDN interfaces, configure when the Terminal Endpoint Identifier (TEI) negotiates with the ISDN provider.

Options

`first-call`—Activation does not occur until the call setup is sent.

`power-up`—Activation occurs when the Services Router is powered on.

- **Default:** `power-up`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

telemetries

IN THIS SECTION

- Syntax | 29049
- Hierarchy Level | 29049
- Description | 29049
- Default | 29049
- Required Privilege Level | 29049
- Release Information | 29049

Syntax

```
telemetries {
    disable;
    duration hours;
    interval minutes;
}
```

Hierarchy Level

```
[edit poe interface (all | interface-name)]
```

Description

Enable the logging of power consumption of a PoE interface over time.

If you want to log the power consumption of a PoE interface, you must explicitly specify the `telemetries` statement. When you commit the configuration, logging begins, with data being collected at the specified intervals. Logging stops at the end of the specified duration. If you do not specify the **duration** and `interval` statements, data is collected at the default interval of five minute intervals every hour.

The remaining statements are explained separately. See CLI Explorer.

Default

Logging of power consumption is disabled.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Metro Routers.

RELATED DOCUMENTATION

Configuring PoE Interfaces on EX Series Switches

show poe telemetries | 23189

telemetries (PoE)

IN THIS SECTION

- Syntax | 29050
- Hierarchy Level | 29050
- Description | 29050
- Default | 29051
- Required Privilege Level | 29051
- Release Information | 29051

Syntax

```
telemetries {
    disable;
    duration hours;
    interval minutes;
}
```

Hierarchy Level

```
[edit poe interface (all | interface-name)]
```

Description

Allow logging of per-port PoE power consumption. The telemetries section must be explicitly specified to enable logging. If left unspecified, telemetries is disabled by default.

Default

If the telemetries statement is specified, logging is enabled with the default values for interval and duration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

telemetry

IN THIS SECTION

- Syntax | 29051
- Hierarchy Level | 29052
- Description | 29052
- Options | 29053
- Required Privilege Level | 29053
- Release Information | 29054

Syntax

```
telemetry {
    subscriber-statistics;
    service-statistics;
    }
    interfaces junos-interface-name {
        unit "$junos-interface-unit" {
            family [inet | inet6] {
```

```
                filter {
                    input "inputFilter";
                    output "outputFilter";
                }
            }
        }
    }
    queue-statistics {
        interface $junos-interface-name {
            refresh rate;
            queues queue-set;
        }
        interface-set $junos-interface-set-name {
            refresh rate;
            queues queue-set;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Enable telemetry collection of subscriber statistics and queue statistics.

Include the **subscriber-statistics** statement to enable the export of subscriber statistics through telemetry. When this statement is configured, you must also include the **actual-transit-statistics** statement at the [`edit dynamic-profiles` *`profile-name`* `interfaces` *`interface-name`* `unit` *`unit-name`*} hierarchy level to enable subscriber-statistics.

Include the **queue-statistics** statement to instruct the statistics infrastructure to collect queue statistics for dynamic interfaces or interface-sets queue-statistics and enable export via Junos Telelemetry Interface (JTI).

The profile variable `$junos-interface-name`" and "`$junos-interface-set-name`" are generated from the corresponding device, unit and interface-set elements in the interfaces stanza at profile instantiation time. Using these derived variables is a convenient way to configure telemetry behavior for the interface or interface-set without the need to mimic the specific configuration in the interfaces stanza.

After telemetry for these statistics is enabled, they are eligible for export through a collector subscription.

For information about subscribing to the statistics through an external collector, see *Configure a NETCONF Proxy Telemetry Sensor in Junos*. For information about supported sensors for subscriber statistics and queue statistics, see *Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)*.

Options

subscriber-statistics

Enable the export of interface meta-data and export of subscriber accurate statistics. When this statement is configured, you must also include the **actual-transit-statistics** statement at the [`edit dynamic-profiles` *`profile-name`* `interfaces` *`interface-name`* `unit` *`unit-name`*] hierarchy level.

queue statistics interface "$junos-interface-name"

Enable the export of interface meta-data and interface queue statistics. The profile variable `$junos-interface-name"` is generated from the corresponding device, unit and interface elements in the interfaces stanza at profile instantiation time.

queue statistics interface-set "$junos-interface-set-name"

Enable the export of interface-set meta-data and interface-set queue statistics. The profile variable `"$junos-interface-set-name"` is generated from the corresponding device, unit and interface-set elements in the interfaces stanza at profile instantiation time.

refresh *rate*

Override the default internal queue statistics collection interval. If dynamic interfaces and interface-sets are created as a result of multiple dynamic profiles, each with their own refresh intervals, the smallest interval for each object type (interface or interface-set) is used to poll queue statistics for that object type. The default is 900 seconds.

Range: 300 seconds (5 minutes) to 86,400 seconds (24 hours)

queue "*queue-set*"

Specify the set of queues for which queue-statistics will be exported. The queue set is a comma delimited string of integers. The default is all queues (`0,1,2,3,4,5,6,7`) are eligible for export.

Range: 0 to 7

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 18.4.

RELATED DOCUMENTATION

Enable Export of Subscriber Statistics and Queue Statistics

Understanding Enabling Export of Subscriber Statistics and Queue Statistics for Dynamic Interfaces and Interface-Sets

Guidelines for Exporting Subscriber Statistics and Queue Statistics for Dynamic Interfaces and Interface-Sets

Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)

telemetry

IN THIS SECTION

- Syntax | **29054**
- Hierarchy Level | **29055**
- Description | **29055**
- Options | **29055**
- Required Privilege Level | **29055**
- Release Information | **29055**

Syntax

```
telemetry {
    statistics{
        no-ingress;
        no-transit;
    }
}
```

Hierarchy Level

```
[edit logical-systems name protocols source-packet-routing],
[edit protocols source-packet-routing]
```

Description

Enable telemetry on segment routing traffic engineering policies at ingress nodes. Junos OS creates the following sensors to collect traffic statistics for segment routing:

- BGP-SRTE policies for ingress routes, sensors are attached to nexthops in inet{6}color.0 table.
- Static SRTE policies for ingress routes, sensors are attached to nexthops in inet{6}color.0 table.
- Transit routes for BGP-SRTE policies in mpls.0 table.
- Transit routes for static SRTE policies in mpls.0 table.

> **NOTE**: If both static SRTE policy and BGP SRTE policy are to present for the same destination and color then only one of them is made active and BSID of the corresponding active policy is programmed to the mpls.0 table. In this case the sensor is attached to BSID of the active policy only.

Options

statistics Enable collection of traffic statistics on segment routing traffic engineering policies.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

statistics | **28591**

source-packet-routing (Segment Routing Traffic Engineering) | 28257

Segment Routing Traffic Engineering at BGP Ingress Peer Overview

telnet

IN THIS SECTION

- Syntax | 29056
- Syntax (EX Series Switches) | 29056
- Syntax (Junos OS Evolved) | 29057
- Description | 29057
- Options | 29057
- Additional Information | 29058
- Required Privilege Level | 29058
- Output Fields | 29058
- Sample Output | 29058
- Release Information | 29059

Syntax

```
telnet host
<8bit>
<inet | inet6>
<port port-number>
<routing-instance routing-instance-name>
<logical-system logical-system-name>
<tenant tenant-name>
```

Syntax (EX Series Switches)

```
telnet host
<8bit>
```

```
<bypass-routing>
<inet | inet6>
<interface interface-name>
<no-resolve>
<port port-number>
<routing-instance routing-instance-name>
<source source-address>
```

Syntax (Junos OS Evolved)

```
telnet host
<8bit>
<inet | inet6>
<port port-number>
<routing-instance routing-instance-name>
```

Description

Open a telnet session to a remote system. Type Ctrl+] to escape from the telnet session to the telnet command level, and then type `quit` to exit from telnet.

NOTE: For Junos OS Evolved, use the `routing-instance mgmt_junos` option to access a remote system through the management interface.

Options

host	Name or address of the remote system.
8bit	(Optional) Use an 8-bit data path.
bypass-routing	(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.
inet | inet6	(Optional) Open an IPv4 or IPv6 session, respectively.
interface *interface-name*	(Optional) Interface name for the telnet session. (This option does not work when `default-address-selection` is configured at the `[edit system]`

	hierarchy level, because this configuration uses the loopback interface as the source address for all locally generated IP packets.)
logical-system *logical-system-name*	(Optional) Name of a particular logical system for the telnet attempt.
tenant *tenant-name*	(Optional) Name of a particular tenant system for the telnet attempt.
no-resolve	(Optional) This option is not supported for Junos OS Evolved Release 18.3R1. Do not attempt to determine the hostname that corresponds to the IP address.
port *port-number*	(Optional) Port number or service name on the remote system.
routing-instance *routing-instance-name*	(Optional) Name of the routing instance for the telnet attempt.
source *source-address*	(Optional) This option is not supported for Junos OS Evolved Release 18.3R1. Source address of the telnet connection.

Additional Information

You can limit the number of times a user can attempt to enter a password while logging in through telnet. To specify the number of times a user can attempt to enter a password to log in through telnet, include the "retry-options" on page 15712 statement at the [`edit system login`] hierarchy level.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

telnet

```
user@host> telnet 192.154.1.254
Trying 192.154.169.254...
Connected to level5.company.net.
Escape character is '^]'.
ttypa
```

```
login:
```

Release Information

Command introduced before Junos OS Release 7.4.

The following options are deprecated in Junos OS Evolved Release 18.3R1: `bypass-routing`, `interface`, `no-resolve`, and `source`.

The command `tenant` option is introduced in Junos OS Release 19.2R1 for SRX Series.

telnet (Access)

IN THIS SECTION

- Syntax | 29059
- Hierarchy Level | 29060
- Description | 29060
- Options | 29060
- Required Privilege Level | 29060
- Release Information | 29060

Syntax

```
telnet {
    banner {
        fail string;
        login string;
        success string;
    }
}
```

Hierarchy Level

```
[edit access firewall-authentication pass-through]
```

Description

Configure banners for Telnet login prompt, successful authentication, and failed authentication.

Options

The remaining statements are explained separately.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Understanding Pass-Through Authentication

telnet (Services)

IN THIS SECTION

- Syntax | **29061**
- Hierarchy Level | **29061**
- Description | **29061**
- Options | **29061**

Required Privilege Level | 29062
Release Information | 29062

Syntax

```
telnet {
    authentication-order [authentication-methods];
    connection-limit limit;
    rate-limit limit;
}
```

Hierarchy Level

```
[edit system services]
```

Description

Provide Telnet connections from remote systems to the local device.

Options

authentication-order [*authentication-methods*]

Configure the order in which the software tries different user authentication methods when attempting to authenticate a user. For each login attempt, the software tries the authentication methods in order, starting with the first one, until the password matches.

- **Values:** Specify one or more of the following authentication methods listed in the order in which they must be tried:
  - `ldaps`—Use LDAP authentication services.
  - `password`—Use the password configured for the user with the `authentication` statement at the `[edit system login user]` hierarchy level.

- `radius`—Use RADIUS authentication services.
- `tacplus`—Use TACACS+ authentication services.

- **Default:** If you do not include the `authentication-order` statement, users are verified based on their configured passwords.

connection-limit ***limit***

Configure the maximum number of connections sessions for the telnet service per protocol (either IPv6 or IPv4).

> **NOTE**: The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

- **Range:** 1 through 250 connections
- **Default:** 75 connections

rate-limit ***limit***

Configure the maximum number of connections attempts per minute, per protocol (either IPv6 or IPv4) on an access service. For example, a rate limit of 10 allows 10 IPv6 telnet session connection attempts per minute and 10 IPv4 telnet session connection attempts per minute.

- **Range:** 1 through 250 connections
- **Default:** 150 connections

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Option `ldaps` introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Configuring Telnet Service for Remote Access to a Router or Switch

temperature-sensor

IN THIS SECTION

- Syntax | 29063
- Hierarchy Level | 29063
- Description | 29064
- Options | 29064
- Required Privilege Level | 29070
- Release Information | 29070

Syntax

```
temperature-sensor sensor-name {
    temperature-threshold {
        fans-on-full-speed fans-full-speed-temp;
        fans-on-full-speed-if-failed-fan full-speed-failed-fan-temp;
        fans-on-intermediate-speed fans-intermediate-speed-temp;
        fans-to-normal-speed fans-normal-speed-temp;
        fire-shutdown fire-shutdown-temp;
        red-alarm red-alarm-temp;
        red-alarm-if-failed-fan red-alarm-failed-fan-temp;
        yellow-alarm yellow-alarm-temp;
        yellow-alarm-if-failed-fan yellow-alarm-failed-fan-temp;
    }
}
```

Hierarchy Level

```
[edit chassis cb 0]
[edit chassis fpc fpc_number]
```

Description

Modifies the temperature thresholds at which fan speeds change. This command is used to thermally manage high-power optics such as 400G ZR and ZR-M in Junos OS and Junos OS Evolved releases. This command is supported on PTX10K with PTX10K-LC1201-36CD line card, PTX10001-36MR, ACX7100-32C, ACX7100-48L, MX10K with MX10K-LC9600, and MX304 with MX304-LMIC16 line card.

> **NOTE**: Optics EM Policy incorporates periodically polled temperature reading of pluggable optical modules in the system's automated control of fan speed behavior. This eliminates the need for manual fan speed provisioning for high power optics (ex. 400G ZR/ZR+) using this command. Please check Junos OS Release notes and Junos OS Evolved Release Notes to verify if your device supports Optics EM Policy.

Options

Options	Description
temperature sensor *sensor-name*	Specify name of the temperature sensor in quotes if the name includes spaces. Use the *sensor-name* as given below: • **"Temp Sensor PSM Inlet"** for PTX10001-36MR devices. • **"Intake A Temp Sensor"** for PTX10K devices with PTX10K-LC1201-36CD line card. • **"Main PCBLeft"** and **"Main PCBRight"** for ACX7100-32C and ACX7100-48L devices. • **"Intake A Temp Sensor"** for MX10K devices with MX10K-LC9600 line card. • **"Intake_A"** for MX304 devices with MX304-LMIC16 line card. • **xcvr-*fpc-slot-number/ pic-number/ port-number*** for devices with optic modules.
temperature threshold	Specify the temperature threshold values.
fans-on-full-speed *fans-full-speed-temp*	Temperature threshold based on chassis-cooling requirements for fans on full speed.

(Continued)

Options	Description
fans-on-full-speed-if-failed-fan *full-speed-failed-fan-temp*	Temperature threshold based on chassis-cooling requirements if fans are at full speed and any fan fails.
fans-on-intermediate-speed *fans-intermediate-speed-temp*	Temperature threshold based on chassis-cooling requirements when fans are at intermediate speed.
fans-to-normal-speed *fans-normal-speed-temp*	Temperature threshold based on chassis-cooling requirements for fans at normal speed.
fire-shutdown *fire-shutdown-temp*	Temperature threshold beyond which Junos triggers a fire shutdown.
red-alarm *red-alarm-temp*	Temperature threshold beyond which Junos triggers a red alarm alert.
red-alarm-if-failed-fan *red-alarm-failed-fan-temp*	Temperature threshold beyond which Junos triggers a red alarm when any fan fails.
yellow-alarm *yellow-alarm-temp*	Temperature threshold beyond which Junos triggers a yellow alarm alert.
yellow-alarm-if-failed-fan *yellow-alarm-failed-fan-temp*	Temperature threshold beyond which Junos triggers a yellow alarm when any fan fails.

NOTE: When configuring temperature thresholds, you must maintain the following order:

`fans-to-normal-speed-temp < fans-on-intermediate-speed-temp < fans-full-speed-temp <= yellow-alarm-temp < red-alarm-temp < fire-shutdown-temp`

The configuration snippets below provide recommended values for the temperature thresholds for different platforms in celsius.

PTX10001-36MR

```
[edit chassis cb 0]
temperature-sensor "Temp Sensor PSM Inlet" {
    temperature-threshold {
        fans-on-full-speed 30;
        fans-on-full-speed-if-failed-fan 28;
        fans-on-intermediate-speed 20;
        fans-to-normal-speed 15;
        fire-shutdown 80;
        red-alarm 75;
        red-alarm-if-failed-fan 72;
        yellow-alarm 65;
        yellow-alarm-if-failed-fan 62;
    }
}
```

PTX10K-LC1201-36CD line card for PTX10004, PTX10008, and PTX10016

```
[edit chassis cb 0]
temperature-sensor "Intake A Temp Sensor" {
    temperature-threshold {
        fans-on-full-speed 3;
        fans-on-full-speed-if-failed-fan 0;
        fans-on-intermediate-speed 2;
        fans-to-normal-speed 1;
        fire-shutdown 80;
        red-alarm 75;
        red-alarm-if-failed-fan 72;
        yellow-alarm 65;
        yellow-alarm-if-failed-fan 62;
    }
}
```

ACX7100-48L

JCO400-QDD-ZR, JCO400-QDD-ZR-M, QDD-400G-ZR-M, QDD-400G-ZR-M-HP Optic Modules

```
[edit chassis fpc fpc_number]
temperature-sensor xcvr-0/0/1 {
    temperature-threshold {
        fans-on-full-speed 38;
        fans-on-full-speed-if-failed-fan 38;
        fans-on-intermediate-speed 26;
        fans-to-normal-speed 21;
        fire-shutdown 90;
        red-alarm 80;
        red-alarm-if-failed-fan 80;
        yellow-alarm 70;
        yellow-alarm-if-failed-fan 70;
    }
}
temperature-sensor xcvr-0/0/2 {
    temperature-threshold {
        fans-on-full-speed 38;
        fans-on-full-speed-if-failed-fan 38;
        fans-on-intermediate-speed 26;
        fans-to-normal-speed 21;
        fire-shutdown 90;
        red-alarm 80;
        red-alarm-if-failed-fan 80;
        yellow-alarm 70;
        yellow-alarm-if-failed-fan 70;
    }
}
```

ACX7100-32C

JCO400-QDD-ZR, JCO400-QDD-ZR-M, QDD-400G-ZR, QDD-400G-ZR-M, QDD-400G-ZR-M-HP Optic Modules

```
[edit chassis fpc fpc_number]
temperature-sensor xcvr-0/0/1 {
    temperature-threshold {
        fans-on-full-speed 30;
```

```
        fans-on-full-speed-if-failed-fan 30;
        fans-on-intermediate-speed 21;
        fans-to-normal-speed 16;
        fire-shutdown 100;
        red-alarm 90;
        red-alarm-if-failed-fan 90;
        yellow-alarm 80;
        yellow-alarm-if-failed-fan 80;
    }
}
temperature-sensor xcvr-0/0/2 {
    temperature-threshold {
        fans-on-full-speed 34;
        fans-on-full-speed-if-failed-fan 34;
        fans-on-intermediate-speed 25;
        fans-to-normal-speed 20;
        fire-shutdown 100;
        red-alarm 90;
        red-alarm-if-failed-fan 90;
        yellow-alarm 80;
        yellow-alarm-if-failed-fan 80;
    }
}
```

MX10K-LC9600 Line Card for MX10004 and MX10008

```
[edit chassis fpc fpc_number]
temperature-sensor "Intake-A Temp Sensor" {
    temperature-threshold {
        fans-on-full-speed 3;
        fans-on-full-speed-if-failed-fan 0;
        fans-on-intermediate-speed 2;
        fans-to-normal-speed 1;
        fire-shutdown 80;
        red-alarm 75;
        red-alarm-if-failed-fan 72;
        yellow-alarm 65;
        yellow-alarm-if-failed-fan 62;
    }
}
```

NOTE: The *fpc_number* is the FPC where MX10K-LC9600 line card is installed.

MX304-LMIC16 Line Card for MX304

```
[edit chassis cb 0]
temperature-sensor Intake_A {
    temperature-threshold {
        fans-to-normal-speed 5;
        fans-on-intermediate-speed 15;
        fans-on-full-speed 20;
        fans-on-full-speed-if-failed-fan 17;
        yellow-alarm 50;
        yellow-alarm-if-failed-fan 50;
        red-alarm 70;
        red-alarm-if-failed-fan 70;
        fire-shutdown 75;
    }
}
```

Verify Configured Thresholds

You can verify the configured thresholds using the `show chassis temperature-thresholds` command:

```
user@host> show chassis temperature-thresholds
                                        Fan speed      Yellow alarm      Red alarm      Fire
Shutdown
                                       (degrees C)      (degrees C)     (degrees C)
(degrees C)
Item                                  Normal  High   Normal  Bad fan   Normal  Bad fan     Normal
Routing Engine 0 Temp Sensor Exhaust Left   65   70      80       78       85       82         90
Routing Engine 0 Temp Sensor Exhaust Middle   65   70    80       78       85       82         90
Routing Engine 0 Temp Sensor Exhaust Right   65   70     80       78       85       82         90
Routing Engine 0 CPU Temperature          72    77       88       85       95       92         98
CB 0 Temp Sensor Left 1                   51    55       70       68       75       72         80
CB 0 Temp Sensor Left 2                   75    80       90       87       95       92         98
CB 0 Temp Sensor PSM Inlet                39    43       65       62       75       72         80
CB 0 Temp Sensor Exhaust                  90    95      112      110      116      114        120
CB 0 Temp Sensor Exhaust Bottom           70    75       85       82       90       87         95
```

```
CB 0 Inlet-1                              75    80       90       87       95       92        98
CB 0 Temp Sensor Exhaust Right 1          75    80       90       87       95       92        98
CB 0 Temp Sensor Exhaust Right 2          75    80       90       87       95       92        98
FPC 0 BT-0 HBM-0 Temperature              80    85      103      100      105      103       110
FPC 0 BT-0 HBM-1 Temperature              80    85      103      100      105      103       110
FPC 0 BT-0 TVsense Temp Max               95   100      120      117      124      120       130
FPC 0 BT-1 HBM-0 Temperature              80    85      103      100      105      103       110
FPC 0 BT-1 HBM-1 Temperature              80    85      103      100      105      103       110
FPC 0 BT-1 TVsense Temp Max               95   100      120      117      124      120       130
FPC 0 BT-2 HBM-0 Temperature              80    85      103      100      105      103       110
FPC 0 BT-2 HBM-1 Temperature              80    85      103      100      105      103       110
FPC 0 BT-2 TVsense Temp Max               95   100      120      117      124      120       130
FPC 0 WAN RT0.0                           82    87      107      104      110      107       115
FPC 0 WAN RT0.1                           82    87      107      104      110      107       115
FPC 0 WAN RT1.0                           82    87      107      104      110      107       115
FPC 0 WAN RT1.1                           82    87      107      104      110      107       115
FPC 0 WAN RT2.0                           82    87      107      104      110      107       115
FPC 0 WAN RT2.1                           82    87      107      104      110      107       115
FPC 0 BT-0 Temp sensor 0                  95   100      120      117      124      120       130
FPC 0 BT-0 Temp sensor 1                  95   100      120      117      124      120       130
FPC 0 BT-1 Temp sensor 0                  95   100      120      117      124      120       130
FPC 0 BT-1 Temp sensor 1                  95   100      120      117      124      120       130
FPC 0 BT-2 Temp sensor 0                  95   100      120      117      124      120       130
FPC 0 BT-2 Temp sensor 1                  95   100      120      117      124      120       130
SIB 0 ASIC Temp, ZF0                      80    85      115      112      121      117       125
FPC 0 xcvr-0/0/1                          47    52       70       70       74       74        75
FPC 0 xcvr-0/0/4                          50    55       73       73       76       76        77
FPC 0 xcvr-0/0/5                          47    52       70       70       75       75        76
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Hardware Compatibility Tool

template

IN THIS SECTION

- Syntax | 29071
- Hierarchy Level | 29071
- Description | 29071
- Required Privilege Level | 29071
- Release Information | 29072

Syntax

```
template;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path p2mp-lsp-template-name],
[edit protocols mpls label-switched-path p2mp-lsp-template-name]
```

Description

Specify a template for the dynamically generated point-to-multipoint LSPs used for VPLS flooding.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring Dynamic Point-to-Multipoint Flooding LSPs with a Preconfigured Template

template (Services URL Filter)

IN THIS SECTION

- Syntax | 29072
- Hierarchy Level | 29073
- Description | 29073
- Options | 29073
- Required Privilege Level | 29074
- Release Information | 29075

Syntax

```
template template-name {
    client-interfaces [ client-interface-name1 client-interface-name2 ];
    disable-url-filtering;
    dns-resolution-interval minutes;
    dns-resolution-rate seconds;
    dns-retries number;
    dns-routing-instance dns-routing-instance-name;
    dns-server [ ip-address1 ip-address2 ip-address3 ];
    dns-source-interface loopback-interface-name;
    routing-instance routing-instance-name;
    server-interfaces  [ server-interface-name1 server-interface-name2 ];
    term term-name {
        from {
```

```
            src-ip-prefix [prefix1 prefix2];
            dest-port   [port1 port2];
        }
        then {
            accept;
            custom-page custom-page;
            http-status-code http-status-code;
            redirect-url redirect-url;
            tcp-reset;
        }
    }
    url-filter-database filename
}
```

Hierarchy Level

```
[edit services url-filter profile profile-name]
```

Description

Configure a URL filter template.

> **NOTE**: Starting in Junos OS Release 18.3R1, the `template` statement is deprecated and has been replaced by the `url-filter-template` statement. The `template` statement is supported for backward compatibility.

Options

template-name	Name of the template.
client-interfaces [*client-interface-name1 client-interface-name2* **]**	The list of client-facing logical interfaces (uplink) on which the URL filtering is configured. This option is mandatory.
disable-url-filtering	Disables the filtering of HTTP traffic that contains an embedded IP address (for example, http:/10.1.1.1) belonging to a blocklisted domain name in the URL filter database.
dns-resolution-interval *minutes*	DNS resolution time interval in minutes.

- **Default:** 1440
- **Range:** 60 through 1440 minutes.

`dns-resolution-rate` *seconds*

Number of DNS queries per second sent out from the system before initiating further DNS queries.

- **Default:** 50
- **Range:** 50 through 100.

`dns-retries` *number*

Number of retries for a DNS query in case query fails or times out.

- **Default:** 3
- **Range:** 1 through 5.

`dns-routing-instance` *dns-routing-instance-name*

The VRF on which the DNS server is reachable. This option is mandatory. You can use the default routing instance inet.0 or a defined routing instance.

`dns-server [` *ip-address1 ip-address2 ip-address3* `]`

One or more IP (IPv4 or IPv6) addresses of DNS servers to which the DNS queries are sent out. This option is mandatory.

`dns-source-interface` *loopback-interface-name*

The loopback interface for which source IP address is picked for sending DNS queries. This option is mandatory.

`routing-instance` *routing-instance-name*

The VRF on which URL filtering feature is configured. This option is mandatory. You can use the default routing instance inet.0 or a defined routing instance.

`server-interfaces [` *server-interface-name1 server-interface-name2* `]`

Server-facing interfaces to which traffic is destined. This option is mandatory.

The list of server-facing logical interfaces (downlink) on which the URL filtering is configured. This option is mandatory.

`url-filter-database` *filename*

The filename of the URL filter database. The file should be placed in the **/var/db/url-filterd** directory, but indicate just the filename here and not the full path.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

`disable-url-filtering` option introduced in Junos OS Release 17.2R2 and 17.4R1.

RELATED DOCUMENTATION

Configuring URL Filtering

template (Flow Monitoring IPFIX Version)

IN THIS SECTION

- Syntax | 29075
- Hierarchy Level | 29076
- Description | 29076
- Options | 29076
- Required Privilege Level | 29076
- Release Information | 29077

Syntax

```
template template-name {
    data-record-fields {
        source-prefix-as-path count;
        destination-prefix-as-path count;
        bgp-source-standard-community count;
        bgp-destination-standard-community count;
        bgp-source-extended-community count;
        bgp-destination-extended-community count;
        bgp-source-large-community count;
        bgp-destination-large-community count;
    }
```

```
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    flow-key {
        flow-direction;
        vlan-id;
        output-interface;
    }
    (bridge-template|ipv4-template | ipv6-template | mpls-ipv4-template | mpls-ipvx-template | vpls-
template);
    multi-bgp-path;
    nexthop-learning (enable |disable);
    observation-domain-id id;
    option-refresh-rate packets packets seconds seconds;
    options-template-id id;
    template-id id;
    template-refresh-rate packets packets seconds seconds;
    tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
}
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix]
```

Description

Specify the IPFIX output template properties to support flow monitoring.

Options

template-name	Name of the IPFIX template.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

`data-record-fields` option introduced in Junos OS Evolved Release 21.4R1.

`multi-bgp-path` option added in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices

template (Flow Monitoring Version 9)

IN THIS SECTION

- Syntax | 29077
- Hierarchy Level | 29078
- Description | 29078
- Options | 29078
- Required Privilege Level | 29078
- Release Information | 29078

Syntax

```
template template-name {
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    flow-key {
        flow-direction;
        vlan-id;
        output-interface;
    }
```

```
    (bridge-template|ipv4-template | ipv6-template | mpls-template |vpls-template label-position
[ positions ] | mpls-ipv4-template label-position [ positions ] | mpls-ipvx-template);
    multi-bgp-path;
    option-refresh-rate packets packets seconds seconds;
    options-template-id id;
    peer-as-billing-template;
    source-id id;
    template-id id;
    template-refresh-rate packets packets seconds seconds;
    tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
}
```

Hierarchy Level

```
[edit services flow-monitoring version9]
```

Description

Specify the version 9 output template properties to support flow monitoring.

Options

template-name—Name of the version 9 template.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

`multi-bgp-path` option added in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

> *Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates*

template (Forwarding Options)

IN THIS SECTION

- Syntax | 29079
- Hierarchy Level | 29079
- Description | 29079
- Options | 29079
- Required Privilege Level | 29080
- Release Information | 29080

Syntax

```
template template-name;
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server hostname version9],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server hostname version9]
```

Description

Specify flow monitoring version 9 template to be used for output of sampling records.

Options

template-name—Name of the version 9 template.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

> *Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates*

template (Forwarding Options Version IPFIX)

IN THIS SECTION

- Syntax | 29080
- Hierarchy Level | 29081
- Description | 29081
- Required Privilege Level | 29081
- Release Information | 29081

Syntax

```
template;
```

Hierarchy Level

```
[edit forwarding-options sampling instance family (inet | inet6 | mpls | vpls) output flow-
server hostname version-ipfix]
```

Description

Specify flow monitoring version IPFIX properties to apply to output sampling records.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices

template-id

IN THIS SECTION

- Syntax | 29082
- Hierarchy Level | 29082
- Description | 29082
- Options | 29083
- Required Privilege Level | 29083

Release Information | 29083

Syntax

```
template-id id;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Define a template ID to be used for flow aggregation of version 9 and IPFIX flows.

MX and QFX Series

The template ID can be a value in the range of 1024 through 65535. If you do not configure a value for the template ID, a default value is assumed for this ID, which is different for the various address families. If you configure the same template ID value for different address families, such a setting is not processed properly and might cause unexpected behavior. For example, if you configure the same template ID value for both IPv4 and IPv6, the collector validates the export data based on the template ID value that it last receives. In this case, if IPv6 is configured after IPv4, the value is effective for IPv6 and the default value is used for IPv4.

PTX Series

If you choose to configure the template ID, the range is 1024 to 65520. If you do not configure this ID, the default value that are set is in the range 256-511 and is different for each template.

You can configure the `template-id` statement for family `inet`, `inet6`, and `mpls` only.

You must not configure the same IDs for different templates (option or data template). The operating system does not check to ensure that you do not configure the same ID value for different templates. If you configure the same ID value, such a setting is not processed properly and might cause unexpected behavior.

The template ID range [configured `template-id` value + 20) is reserved and you must not configure any another ID in this range. The difference between configured template IDs across families should be at least 20; for example, if `template-id 1056` is configured for family `inet`, you should not configure template IDs in the range of 1056 to 1075 for any other families.

For Junos OS, if you change the template ID, all flows are inactively timed out. New flows are learned afresh.

For Junos OS Evolved, if you change the template ID, this change does not impact the flows.

Options

id—Unique identifier for the template to be used for version 9 or IPFIX flows. The range 0-255 is reserved and is used for template sets, options template sets, and for future expansion. If you do not configure a template ID, the software generates a template ID for you.

- **Range:** 1024 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows

Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows

template (Inline Monitoring)

IN THIS SECTION

- Syntax (Junos OS) | 29084
- Syntax (Junos OS Evolved) | 29085
- Hierarchy Level | 29085
- Description | 29085
- Options | 29085
- Required Privilege Level | 29086
- Release Information | 29086

Syntax (Junos OS)

```
template name {
    flow-inactive-timeout seconds;
    flow-monitoring {
        counter-profile profile-identifier;
        flow-rate kbps burst-size bytes;
        flow-limit number;
        sampling-rate bytes;
        sampling-profile profile-name;
        security-enable;
    }
    observation-domain-id observation-domain-id;
    option-template-id option-template-id;
    option-template-refresh-rate seconds;
    primary-data-record-fields name;
    template-id template-id;
    template-refresh-rate seconds;
    template-type (ipv4-template | ipv6-template);
}
```

Syntax (Junos OS Evolved)

```
template name {
    observation-domain-id observation-domain-id;
    option-template-refresh-rate seconds;
    primary-data-record-fields name;
    template-refresh-rate seconds;
}
```

Hierarchy Level

```
[edit services inline-monitoring]
```

Description

Configure template for inline monitoring services.

Options

name	Name of the template.
flow-inactive-timeout ***seconds***	(EX switches only) Configure the inactive-timeout period for a flow, in seconds, for flow-based telemetry. Once there is no active traffic for a flow, the flow is aged out after the configured inactive-timeout period.
flow-monitoring	(EX switches only) Configure flow-based telemetry. See flow-monitoring (Inline Monitoring Services) for more information.
observation-domain-id ***observation-domain-id***	Significant one byte of observation domain ID used to uniquely identify the exporting process. The other three bytes are system generated and are unique within the chassis. • **Default:** 0 • **Range:** 1 through 255
option-template-id ***option-template-id***	(Junos OS only) Option template ID. For Junos OS Evolved, the system generates the option template ID.

- **Default:** 640
- **Range:** 1024 through 65535

option-template-refresh-rate *seconds*

Option refresh rate in seconds.

- **Default:** 600 seconds
- **Range:** 10 through 600 seconds

primary-data-record-fields *name*

Configure which IPFIX information elements (IEs) to include in the primary data record for Juniper Resiliency Interface. See "primary-data-record-fields" on page 13134 for more information.

template-id *template-id*

(Junos OS only) Template ID. For Junos OS Evolved, the system generates the template ID.

- **Default:** 384
- **Range (EX4400):** 1024 through 65535
- **Range (EX4100 and EX4100-F):** 32768 through 32831

template-refresh-rate *seconds*

Refresh rate in seconds.

- **Default:** 600 seconds
- **Range:** 10 through 600 seconds

template-type (ipv4-template | ipv6-template)

(QFX5120 only, Junos OS only) For flow-based telemetry for VXLANs, specify that the template is for either IPv4 or IPv6 traffic.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

`flow-inactive-timeout` and `flow-monitoring` options introduced in Junos OS Release 21.1R1.

`primary-data-record-fields` option introduced in Junos OS Release 21.2R1.

Statement introduced in Junos OS Evolved Release 22.2R1 for PTX Series routers.

`template-type` option introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Inline Monitoring Services

Flow-Based Telemetry

Flow-Based Telemetry for VXLANs

template (LRF Profile)

IN THIS SECTION

- Syntax | 29087
- Hierarchy Level | 29088
- Description | 29088
- Options | 29088
- Required Privilege Level | 29088
- Release Information | 29088

Syntax

```
template template-name {
    format ipfix;
    template-tx-interval tx-time;
    template-type template-type;
    trigger-type (session-close | volume);
}
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Description

Configure a template, which specifies a set of data to be transmitted. This template can be specified in LRF rules.

Options

template-name

Name for the template.

- **Range:** Up to 32 characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

template (LRF Rule)

IN THIS SECTION

- Syntax | 29089
- Hierarchy Level | 29089
- Description | 29089
- Options | 29089
- Required Privilege Level | 29089
- Release Information | 29090

Syntax

```
template template-name;
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name then report]
```

Description

Specify the template that identifies the type of data to report if the LRF rule is matched.

Options

template-name Name of the template that is used. The referenced template must be configured.

- **Range:** Up to 32 characters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

template-profile (Flow Monitoring Logs for NAT)

IN THIS SECTION

- Syntax | 29090
- Hierarchy Level | 29091
- Description | 29091
- Options | 29091
- Required Privilege Level | 29091
- Release Information | 29091

Syntax

```
template-profile template-profile-name;
```

Hierarchy Level

```
[edit services jflow-log],
[edit services service-set service-set-name jflow-log]
```

Description

Specify the name of the flow template profile to be used for generating flow monitoring format messages for NAT events and for transmitting them to the collector. You can define a template profile for the Jflow service by using this statement at the `[edit services jflow-log]` hierarchy level, and associate the template profile with a service set by using this statement at the `[edit services service-set` *`service-set-name`* `jflow-log]` hierarchy level.

Options

`template-profile-name`—Name of the flow template profile for NAT events. The name can be up to 32 alphanumeric characters in length. Allowed characters are [a-zA-Z0-9_].

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

template(Protocols MVPN)

IN THIS SECTION

- Syntax | 29092
- Hierarchy Level | 29093
- Description | 29093
- Required Privilege Level | 29093
- Release Information | 29093

Syntax

```
template template-name {
    all-regions {
        incoming;
        ingress-replication {
            create-new-ucast-tunnel;
            label-switched-path {
                label-switched-path-template (Multicast) {
                    (default-template | lsp-template-name);
                }
            }
        }
        ldp-p2mp;
        rsvp-te {
            label-switched-path-template (Multicast) {
                (default-template | lsp-template-name);
            }
            static-lsp static-lsp;
        }
    }
    region region-name{
        incoming;
        ingress-replication {
            create-new-ucast-tunnel;
            label-switched-path {
                label-switched-path-template (Multicast){
```

```
                    (default-template | lsp-template-name);
                }
            }
        }
        ldp-p2mp;
        rsvp-te {
            label-switched-path-template (Multicast) {
                (default-template | lsp-template-name);
            }
            static-lsp static-lsp;
        }
    }
}
```

Hierarchy Level

```
[edit protocols mvpn inter-region-template]
```

Description

A template is configured on a transit area boundary router (ABR) to define the tunnel type for each region. If a template is not configured on the ABR then the ABR simply reflects the routes without performing any inter-region segmentation.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

> *Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview*

Example: Configuring Segmented Inter-Area P2MP LSP

Configuring Segmented Inter-Area P2MP LSP

all-regions | **2157**

inter-region | **8789**

inter-region-template | **8793**

inter-region-segmented | **8791**

region | **14262**

template-refresh-rate

IN THIS SECTION

- Syntax | **29094**
- Hierarchy Level | **29094**
- Description | **29095**
- Options | **29095**
- Required Privilege Level | **29095**
- Release Information | **29095**

Syntax

```
template-refresh-rate packets packets seconds seconds;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name],
```

Description

Specify the frequency at which the flow generator sends updates about template definitions to the flow collector. Specify either the number of packets or the number of seconds.

Options

packets—Refresh rate, in number of packets.

- **Range:** 1 through 480,000
- **Default:** 4800

seconds—Refresh rate, in number of seconds.

- **Range:** 10 through 600
- **Default (Junos OS):** 600
- **Default (Junos OS Evolved):** 60

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level added in Junos OS Release 10.2.

Support at the `[edit services flow-monitoring version9 template template-name]` hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.

RELATED DOCUMENTATION

Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates

Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices

template-refresh-rate (Services)

IN THIS SECTION

- Syntax | 29096
- Hierarchy Level | 29096
- Description | 29096
- Options | 29096
- Required Privilege Level | 29096
- Release Information | 29097

Syntax

```
template-refresh-rate;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

Specify the template refresh rate.

Options

- `packets`—Specify the number of packets. The range is from 1 through 480,000.
- `seconds`—Specify the number of seconds. The range is from 10 through 600.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding Interfaces

template-tx-interval (LRF Profile)

IN THIS SECTION

- Syntax | 29097
- Hierarchy Level | 29097
- Description | 29097
- Options | 29098
- Required Privilege Level | 29098
- Release Information | 29098

Syntax

```
template-tx-interval tx-time;
```

Hierarchy Level

```
[edit services lrf profile profile-name template template-name]
```

Description

Configure the interval at which to retransmit the template to the collector.

Options

tx-time Time interval in seconds.

- **Default:** 60
- **Range:** 10 through 600

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

template-type (Flow Monitoring Logs for NAT)

IN THIS SECTION

- Syntax | 29099
- Hierarchy Level | 29099
- Description | 29099
- Options | 29099
- Required Privilege Level | 29099

Release Information | 29099

Syntax

```
template-type nat;
```

Hierarchy Level

```
[edit services jflow-log template-profile template-profile-name]
```

Description

Specify the type of service for which flow template profiles, in version or IPFIX format, must be used for generating flow monitoring format messages for NAT events and for transmitting them to the collector. Currently, you can configure only NAT events or services for generation of log messages in flow monitoring format.

Options

`nat`—Use flow template profiles for generation of flow monitoring logs for NAT events.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

template-type (LRF Profile)

IN THIS SECTION

- Syntax | 29100
- Hierarchy Level | 29100
- Description | 29100
- Options | 29101
- Required Privilege Level | 29102
- Release Information | 29102

Syntax

```
template-type template-type;
```

Hierarchy Level

```
[edit services lrf profile profile-name template template-name]
```

Description

Configure the template types for the template, which specify the data fields to include. You must configure at least one type, and you can configure multiple types.

If Next Gen Services is enabled, then the template types `dns`, `ifl-subscriber`, `ipv4-extended`, `ipv6-extended`, `mobile-subscriber`, `video`, and `wireline-subscriber` are not available.

Options

template-type

Template type. You must configure at least one of the following types, and you can configure multiple types:

- `device-data`—Use data fields specific to the device collecting the logging feed.
- `dns`—(Not available if Next Gen Services is enabled) Use the DNS response time data field.
- `flow-id`—Use the Flow ID data field.
- `http`—Use data fields for the HTTP metadata from header fields.
- `ifl-subscriber`—(Not available if Next Gen Services is enabled) Use data fields specific to interface-based subscribers.
- `ipflow`—Use data fields for the uplink and downlink octets and bytes.
- `ipflow-extended`—Use data fields for the service set name, routing instance, and payload timestamps.
- `ipflow-tcp`—Use data fields for TCP-related timestamps.
- `ipflow-tcp-ts`—Use IBM-specific data fields for TCP-related timestamps.When configuring a `ipflow-tcp-ts` template, configure `vendor-support ibm` at the `[edit services lrf profile` *`profile-name`*`]` hierarchy level to avoid a commit warning.
- `ipflow-ts`—Use data fields for the flow start and end timestamps.
- `ipv4`—Use data fields for the basic source and destination IPv4 information.
- `ipv4-extended`—(Not available if Next Gen Services is enabled) Use data fields for the elements of IPv4 extended fields.
- `ipv6`—Use data fields for the basic source and destination IPv6 information.
- `ipv6-extended`—(Not available if Next Gen Services is enabled) Use data fields for the elements of IPv6 extended fields.
- `l7-app`—Use data fields for the Layer 7 application.
- `mobile-subscriber`—(Not available if Next Gen Services is enabled) Use data fields specific to mobile subscribers.
- `pcc`—Use the PCC rule name data field.

- `status-code-dist`—Use data fields for the HTTP or DNS status codes.
- `subscriber-data`—Use data fields for Generic Subscriber information that can be included with wireless (mobile) subscribers or wireline subscribers.
- `transport-layer`—Use data fields for the transport layer.
- `video`—(Not available if Next Gen Services is enabled) Use data fields for video traffic.
- `wireline-subscriber`—(Not available if Next Gen Services is enabled) Use the UserName data field for wireline subscribers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

templates

IN THIS SECTION

- Syntax | **29103**
- Hierarchy Level | **29104**
- Description | **29104**

- Options | 29104
- Required Privilege Level | 29105
- Release Information | 29105

Syntax

```
templates {
    template-name {
        interval-duration interval-duration;
        inactive-timeout inactive-timeout;
        rate {
            (layer3 layer3-packets-per-second | media media-bits-per-second);
        }
        delay-factor {
            ;
            threshold {
                (info | warning | critical) delay-factor-threshold;
            }
        }
        media-loss-rate {
            disable;
            threshold {
                (info | warning | critical) percentage mlr-percentage | packet-count mlr-packet-
count);
            }
        }
        media-rate-variation {
            disable;
            threshold {
                (info | warning | critical) mrv-variation;
            }
        }
        media-packets-count-in-layer3 media-packets-count-in-layer3;
        media-packet-size media-packet-size;
    }
}
```

Hierarchy Level

```
[edit services video-monitoring]
```

Description

Configure the media delivery index template containing the measurement parameters for video monitoring.

Options

delay-factor	Define delay factor syslog threshold levels.
delay-factor-threshold	Delay factor threshold in milliseconds. When the threshold is exceeded, a syslog message is generated. • **Default:** 0—Do not generate syslogs. • **Range:** 0 though 65,535 milliseconds
disable	Disable logging for the threshold.
inactive-timeout	Number of seconds of flow inactivity after which time media delivery index statistics collection for a flow is terminated. • **Range:** 30 through 300 seconds
info | warning | critical	Level of syslog message generated when a threshold is exceeded.
interval-duration	Number of seconds after which time media delivery index flow monitoring statistics for the interval are reported. • **Range:** 1 through 50
layer3-packets-per-second	Layer 3 packet rate in IP packets per second. • **Range:** 0 though 4,294,967,295 pps
media-bits-per-second	Media bit rate for the stream in bits per second.
media-loss-rate	Define media loss rate syslog threshold levels.
media-packets-count-in-layer-3	Number of media packets in an IP packet.

- **Range:** 1 through 32

media-packet-size Size of media packet in bits.

- **Default:** 188
- **Range:** 1 through 2048

media-rate-variation Define delay factor syslog threshold levels.

mlr-packet-count Media loss rate threshold expressed as the number of packets dropped. When the threshold is exceeded, a syslog message is generated.

mlr-percentage Media loss rate threshold expressed as the percentage of total packets dropped. When the threshold is exceeded, a syslog message is generated.

- **Range:** 0 through 100

mrv-variation Media rate variation threshold. The variation is the ratio of actual media rate to the configured media rate, expressed as a percentage.

template-name Name of the template containing media delivery index measurement criteria. The template can be assigned to an interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring Inline Video Monitoring on MX Series Routers

term (Application Aware Access List)

IN THIS SECTION

- Syntax | 29106
- Hierarchy Level | 29107
- Description | 29107
- Options | 29107
- Required Privilege Level | 29107
- Release Information | 29107

Syntax

```
term term-name {
    from {
        application-group-any;
        application-groups [ application-group-names ];
        application-unknown;
        applications [ application-names ];
        destination-address address <any-unicast>;
        destination-address-range low minimum-value high maximum-value;
        destination-prefix-list list-name;
        nested-application-unknown;
        source-address address <any-unicast>;
        source-address-range low minimum-value high maximum-value;
        source-prefix-list list-name;
    }
    then {
        (accept | discard);
        count (application | application-group | application-group-any | nested-application |
none);
        forwarding-class class-name;
        policer policer-name;
    }
}
```

Hierarchy Level

```
[edit services aacl rule rule-name]
```

Description

Define the AACL term properties.

Options

term-name—Identifier for the term.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring AACL Rules

term (IP Security VPN)

IN THIS SECTION

- Syntax | 29108
- Hierarchy Level | 29109
- Description | 29109

- Options | 29109
- Required Privilege Level | 29109
- Release Information | 29109

Syntax

```
term term-name {
    from {
        destination-address address;
        ipsec-inside-interface interface-name;
        source-address address;
    }
    then {
        anti-replay-window-size bits;
        backup-remote-gateway address;
        clear-dont-fragment-bit;
        dynamic {
            ike-policy policy-name;
            ipsec-policy policy-name;
        }
        initiate-dead-peer-detection;
        manual {
            direction (inbound | outbound | bidirectional) {
                authentication {
                    algorithm (hmac-sha-256);
                    key (ascii-text key | hexadecimal key);
                }
                auxiliary-spi spi-value;
                encryption {
                    algorithm algorithm;
                    key (ascii-text key | hexadecimal key);
                }
                protocol (bundle | esp);
                spi spi-value;
            }
        }
        no-anti-replay;
        remote-gateway address;
```

```
        syslog;
        tunnel-mtu bytes;
    }
}
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name]
```

Description

Define the IPsec term properties.

Options

term-name—Identifier for the term.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

term (Applications)

IN THIS SECTION

- Syntax | 29110
- Hierarchy Level | 29110
- Description | 29110
- Options | 29111
- Required Privilege Level | 29111
- Release Information | 29111

Syntax

```
term term-name {
   alg application;
   destination-port port-identifier;
   icmp-code value;
   icmp-type value;
   icmp6-code value;
   icmp6-type value;
   inactivity-timeout (seconds | never);
   protocol number;
   rpc-program-number number;
   source-port port-number;
   uuid hex-value;
}
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Define individual application protocols.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policy Applications Overview

term (candidate-uplink-policy)

IN THIS SECTION

- Syntax | 29111
- Hierarchy Level | 29112
- Description | 29112
- Options | 29112
- Required Privilege Level | 29112
- Release Information | 29113

Syntax

```
term term-name{
    from {
        product-model model-name;
```

```
        <holddown holddown-time>;
        <minimum-links number-of-links>;
        <uplink-port-group uplink-port-group-name>;
    }
}
```

Hierarchy Level

```
[edit policy-options satellite-policies candidate-uplink-port-profile policy-name]
```

Description

Create and configure a term in a candidate uplink satellite policy within a Junos Fusion.

A term in a candidate uplink port policy in a Junos Fusion is used to apply an uplink failure detection configuration to certain satellite devices, based on their product model only. The more complex options that are available for other policies in Junos OS—such as the terms available for routing policies—are not available for candidate uplink port policies.

Options

term-name The user-defined name of the term.

A *term* is a named structure in which match conditions and configuration statements are defined. A candidate uplink policy policy can contain multiple terms. The terms are evaluated in the order they appear in the configuration. The first term that matches a satellite device is applied to that satellite device.

from The statements under the `from` statement define the satellite device model match criteria and uplink failure detection configuration for the term. Each term can contain only one `from` statement.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

term (Captive Portal Content Delivery)

IN THIS SECTION

- Syntax | 29113
- Hierarchy Level | 29114
- Description | 29114
- Options | 29114
- Required Privilege Level | 29114
- Release Information | 29114

Syntax

```
term term-name{
    from {
        destination-address address <except>;
    }
    then {
        accept;
        insert tag-name tag-name tag-value tag-value;
        redirect url;
        rewrite destination-address address <destination-port port-number>;
        syslog;
```

```
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule rule-name],
[edit services captive-portal-content-delivery rule rule-name]
```

Description

Define the term match and action properties for the captive portal content delivery rule. Use the statement at the `[edit services...]` hierarchy level for static CPCD. Use the statement at the `[edit dynamic-profiles profile-name services...]` hierarchy level for converged services CPCD.

Options

term-name—Identifier for the term.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the `[edit dynamic-profiles profile-name services captive-portal-content-delivery rule rule-name]` hierarchy level added in Junos OS Release 17.2R1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

HTTP Redirect Service Overview

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

Configuring Routing Engine-Based, Static HTTP Redirect Services

Configuring Routing Engine-Based, Converged HTTP Redirect Services

term (Dynamic Profiles)

IN THIS SECTION

- Syntax | **29115**
- Hierarchy Level | **29116**
- Description | **29116**
- Options | **29116**
- Required Privilege Level | **29116**
- Release Information | **29116**

Syntax

```
term term-name {
    from {
        match-conditions;
    }
    then {
        action;
        action-modifiers;
    }
    only-at-create;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name fast-update-filter filter-name],
[edit dynamic-profiles profile-name firewall family family-name  filter filter-name]
```

Description

Define terms for fast update filters.

Options

action—(Optional) An action to take if conditions match. If you do not specify an action, the packets that match the conditions in the `from` statement are accepted.

action-modifiers—(Optional) One or more actions to perform on a packet.

`from`—(Optional) Match packet fields to values. If not included, all packets are considered to match and the actions and action modifiers in the `then` statement are taken.

match-conditions—One or more conditions to make a match.

`only-at-create`—(Optional) Specify that the term is added only when the fast update filter is first created. No subsequent changes can be made to the term in the filter. Use this option only for terms that do not include subscriber-specific data in their match conditions, such as common or default terms (for example, counting the default drop packets).

term-name—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

`then`—(Optional) Actions to take on matching packets. If not included and a packet matches all the conditions in the `from` statement, the packet is accepted.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support at the [edit dynamic-profiles ... filter *filter-name*] hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- *Configuring Fast Update Filters*
- *Configuring Terms for Fast Update Filters*
- *Fast Update Filter Match Conditions*
- *Fast Update Filter Actions and Action Modifiers*
- *Parameterized Filters Overview*
- *Firewall Filter Terminating and Nonterminating Actions for Protocol-Independent Traffic in Dynamic Service Profiles*
- *Parameterized Filter Nonterminating and Terminating Actions and Modifiers*
- *Parameterized Filter Match Conditions for IPv4 Traffic*
- *Parameterized Filter Match Conditions for IPv6 Traffic*

term (Firewall Filter)

IN THIS SECTION

- Syntax | **29117**
- Hierarchy Level | **29118**
- Description | **29118**
- Options | **29118**
- Required Privilege Level | **29120**
- Release Information | **29120**

Syntax

```
term term-name {
    from {
        match-conditions;
```

```
        vxlan {
            vni vni-id
                flags value mask-in-hex value
                reserved1 value
                reserved2 value
            }
        ip-version ipv4 {
            match-conditions-mpls-ipv4-address;
                protocol (tcp | udp) {
                    match conditions-mpls-ipv4-port;
                }
        }
    }
    then {
        actions;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name filter filter-name],
[edit firewall family family-name filter filter-name],
[edit firewall family family-name service-filter filter-name],
[edit firewall family family-name simple-filter filter-name],
[edit logical-systems logical-system-name firewall family family-name filter filter-name],
[edit logical-systems logical-system-name firewall family family-name service-filter filter-name],
[edit logical-systems logical-system-name firewall family family-name simple-filter filter-name]
```

Description

Define a firewall filter term.

Options

actions—(Optional) Actions to perform on the packet if conditions match. You can specify one *terminating action* supported for the specified filter type. If you do not specify a terminating action, the packets that match the conditions in the `from` statement are accepted by default. As an option, you can specify one or more *nonterminating actions* supported for the specified filter type.

filter-name—(Optional) For `family` *family-name* `filter` *filter-name* only, reference another standard stateless firewall filter from within this term.

`from`—(Optional) Match packet fields to values. If not included, all packets are considered to match and the actions and action modifiers in the `then` statement are taken.

match-conditions—One or more conditions to use to make a match on a packet.

match-conditions-mpls-ipv4-address—(MPLS-tagged IPv4 traffic only) One or more IP address match conditions to match on the IPv4 packet header. Supports network-based service in a core network with IPv4 packets as an inner payload of an MPLS packet with labels stacked up to five deep.

match-conditions-mpls-ipv4-port—(MPLS-tagged IPv4 traffic only) One or more UDP or TCP port match conditions to use to match a packet in an MPLS flow. Supports network-based service in a core network with IPv4 packets as an inner payload of an MPLS packet with labels stacked up to five deep.

vxlan—(Optional) Match packets belonging to a particular VXLAN Network Identifier (VNI).

term-name—Name that identifies the term. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

`then`—(Optional) Actions to take on matching packets. If not included and a packet matches all the conditions in the `from` statement, the packet is accepted.

The Firewall Filer Match Conditions for the different protocols are explained separately:

- *Firewall Filter Match Conditions for IPv4 Traffic*
- *Firewall Filter Match Conditions for IPv6 Traffic*
- *Firewall Filter Match Conditions for MPLS Traffic*
- *Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic*
- *Firewall Filter Match Conditions for VPLS Traffic*
- *Firewall Filter Match Conditions for Protocol-Independent Traffic*
- *Firewall Filter Match Conditions for Protocol-Independent Traffic in Dynamic Service Profiles*
- *Firewall Filter Match Conditions Based on Numbers or Text Aliases*
- *Firewall Filter Match Conditions Based on Bit-Field Values*
- *Firewall Filter Match Conditions Based on Address Fields*
- *Firewall Filter Match Conditions Based on Address Classes*
- *Firewall Filter Match Conditions for Layer 2 Bridging Traffic*

- *Firewall Filter Match Conditions for Layer 2 CCC Traffic*

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`filter` option introduced in Junos OS Release 7.6.

Logical systems support introduced in Junos OS Release 9.3.

`ip-version ipv4` support introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

- *Guidelines for Configuring Firewall Filters*
- *Configuring Multifield Classifiers*
- *Guidelines for Configuring Simple Filters*
- *Guidelines for Configuring and Applying Firewall Filters in Logical Systems*

term (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | 29121
- Hierarchy Level | 29122
- Description | 29122
- Options | 29122
- Required Privilege Level | 29122
- Release Information | 29122

Syntax

```
term term-number {
    from {
        destination-address {
            (any-ipv4 | any-ipv4 except);
            (any-ipv6 | any-ipv6 except);
            (any-unicast | any-unicast except);
            (prefix | prefix except);
        }
        destination-address-range {
            high address low address <except>;
        }
        destination-port-range {
            high port-number low port-number;
        }
        destination-ports value;
        destination-prefix-list {
            (prefix-name | prefix-name except);
        }
    }
    then {
        count;
        tag tag-name {
            encrypt {
                hash algorithm;
                prefix hash-prefix;
            }
            ipv4-mask ipv4-mask;
            ipv6-mask ipv6-mask;
            ipv4-or-value ipv4-or-value;
            ipv6-or-value ipv6-or-value;
            tag-attribute tag-attr-name;
            tag-header header;
            tag-separator separator;
        }
    }
}
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name]
```

Description

Configure a term in a tag rule, which is used to enrich HTTP headers with the appropriate tags. You must configure at least one term for a tag rule, but you can configure multiple terms. Terms are evaluated in the order they are configured. If a data packet matches all the criteria in the `from` statement in any of the terms, then the actions specified in the `then` statement are applied. If the `from` statement does not identify any criteria, then all traffic matches. After a term matches a data packet, further terms are not evaluated. If no terms match, then the HTTP header is not enriched.

For subscriber traffic under static policy control, a tag rule is used if it is included in the HCM profile specified in a PCC rule that the traffic matches. For subscribers under dynamic policy control, a message from the PCRF identifies the configured HCM profile and tag rules to use for HTTP header enrichment.

Options

term-number

Number for the term.

- **Range:** 1 through 32,767

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview

Configuring Tag Rules

Configuring HCM Profiles and Assigning Tag Rules

Junos Web Aware HTTP Header Enrichment Overview

term (IDS MS-DPC)

IN THIS SECTION

- Syntax | 29123
- Hierarchy Level | 29124
- Description | 29124
- Options | 29124
- Required Privilege Level | 29125
- Release Information | 29125

Syntax

```
term term-name {
    from {
        application-sets set-name;
        applications [ application-names ];
        destination-address (address | any-unicast) <except>;
        destination-address-range low minimum-value high maximum-value <except>;
        source-address (address | any-unicast) <except>;
        source-address-range low minimum-value high maximum-value <except>;
    }
    then {
        aggregation (IDS) {
            destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
            source-prefix prefix-value | source-prefix-ipv6 prefix-value;
        }
        (force-entry | ignore-entry);
        logging {
            syslog;
            threshold rate;
        }
```

```
        session-limit {
            by-destination (IDS MS-DPC) {
                hold-time seconds;
                maximum number;
                packets number;
                rate number;
            }
            by-pair (IDS MS-DPC) {
                hold-time seconds;
                maximum number;
                packets number;
                rate number;
            }
            by-source (IDS MS-DPC) {
                hold-time seconds;
                maximum number;
                packets number;
                rate number;
            }
        }
        syn-cookie {
            mss value;
            threshold rate;
        }
    }
}
```

Hierarchy Level

```
[edit services ids rule rule-name]
```

Description

Define the IDS term properties when using the MS-DPC.

Options

term-name—Identifier for the term.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IDS Rules on an MS-DPC

term (IDS MS-MPC)

IN THIS SECTION

- Syntax | 29125
- Hierarchy Level | 29127
- Description | 29127
- Required Privilege Level | 29128
- Release Information | 29128

Syntax

```
term {
    then {
        aggregation (IDS) {
            destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
            source-prefix prefix-value | source-prefix-ipv6 prefix-value;
        }
        allow-ip-options {
            any;
```

```
            loose-source-route;
            route-record;
            route-alert;
            security;
            stream-id;
            strict-source-route;
            timestamp;
        }
        allow-ipv6-extension-header {
            any;
            ah;
            dstopts;
            esp;
            fragment;
            hop-by-hop;
            mobility;
            routing;
        }
        icmp-fragment-check;
        icmp-large-packet-check;
        land-attack-check (ip-only | ip-port);
        session-limit {
            by-destination {
                by-protocol {
                    icmp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                    tcp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                    udp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                }
                maximum number;
                packets number;
                rate number;
```

```
            }
            by-source {
                by-protocol {
                    icmp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                    tcp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                    udp {
                        maximum number;
                        packets number;
                        rate number;
                    }
                }
                maximum number;
                packets number;
                rate number;
            }
        }
        tcp-syn-defense;
        tcp-syn-fragment-check;
        tcp-winnuke-check;
    }
}
```

Hierarchy Level

```
[edit services ids rule rule-name]
```

Description

Configure the network attack prevention actions for an IDS rule for a service set on an MS-MPC.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> *Configuring Protection Against Network Attacks on an MS-MPC*
>
> *Understanding IDS on an MS-MPC*

term (PCP)

IN THIS SECTION

- Syntax | 29128
- Hierarchy Level | 29129
- Description | 29129
- Options | 29129
- Required Privilege Level | 29129
- Release Information | 29129

Syntax

```
term term-name {
    from {
        application-sets set-name;
        applications [ application-name ];
        destination-address address <except>;
        destination-address-range high maximum-value low minimum-value <except>;
```

```
        destination-port  high maximum-value low minimum-value;
        destination-prefix-list list-name <except>;
        source-address address <except>;
        source-address-range high maximum-value low minimum-value <except>;
        source-prefix-list list-name <except>;
    }
    then {
        pcp-server server-name;
    }
}
```

Hierarchy Level

```
[edit services pcp rule rule-name]
```

Description

Define the PCP rule term properties. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

The remaining statements are explained separately.

Options

term-name

Identifier for the term.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

term (satellite-policies)

IN THIS SECTION

- Syntax | 29130
- Hierarchy Level | 29130
- Description | 29130
- Options | 29131
- Required Privilege Level | 29131
- Release Information | 29131

Syntax

```
term term-name{
    from {
        product-modelmodel-name;
    }
}
```

Hierarchy Level

```
[edit policy-options satellite-policies candidate-uplink-policy policy-name],
[edit policy-options satellite-policies environment-monitoring-policy policy-name],
[edit policy-options satellite-policies forwarding-policy policy-name]
```

Description

Create and configure a term in a candidate uplink satellite policy, an environment monitoring satellite policy, or a forwarding policy satellite policy, within a satellite policy.

A term in a candidate uplink policy, an environment monitoring policy, or a forwarding policy for satellite devices is used to apply the policy to certain satellite devices only. The more complex options that are available for other policies in Junos OS—such as the terms available for routing policies—are not available for satellite policies.

The actions of a candidate uplink satellite policy, an environment monitoring satellite policy, or forwarding-policy satellite policy are defined at the [`edit policy-options satellite-policies candidate-uplink-policy` *`policy-name`*, [`edit policy-options satellite-policies environment-monitoring-policy` *`policy-name`*], and [`edit policy-options satellite-policies forwarding-policy` *`policy-name`*hierarchy levels.

Options

term *term-name*

Specifies the user-defined name of the term.

A *term* is a named structure in which match conditions and actions are defined. A candidate uplink policy or environment monitoring policy can contain multiple terms. The term name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

from

(Optional) Defines the match criteria for the satellite policy.

The only match criteria available for a satellite policy is the product model.

If you omit the `from` option, the satellite policy is applied globally.

product-model *model-name*

Defines the product model of the satellite device that uses the satellite policy.

If you want the satellite policy to apply to all EX4300 switches in the satellite device role, enter `EX4300*` as the *model-name*.

If you want the satellite policy to apply to all QFX5100 switches in the satellite device role, enter `QFX5100*` as the *model-name*.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

Statement introduced in Junos OS Release 16.1R1 for Junos Fusion Enterprise.

RELATED DOCUMENTATION

Configuring Satellite Device Alarm Handling Using an Environment Monitoring Satellite Policy in a Junos Fusion

Understanding Satellite Policies in a Junos Fusion

term (Services CoS)

IN THIS SECTION

- Syntax | **29132**
- Hierarchy Level | **29133**
- Description | **29133**
- Options | **29133**
- Required Privilege Level | **29133**
- Release Information | **29133**

Syntax

```
term term-name {
    from {
        application-sets set-name;
        applications [ application-names ];
        destination-address address;
        destination-prefix-list list-name <except>;
        source-address address;
        source-prefix-list list-name <except>;
    }
    then {
        application-profile profile-name;
        dscp (alias | bits);
        forwarding-class class-name;
        syslog;
        reflexive; | revert; | reverse {
            application-profile profile-name;
```

```
            dscp (alias | bits);
            forwarding-class class-name;
            syslog;
        }
    }
}
```

Hierarchy Level

```
[edit services cos rule rule-name]
```

Description

Define the CoS term properties.

Options

term-name—Identifier for the term.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring CoS Rules

Configuring CoS Rules on Services PICs

term (Services NAT)

IN THIS SECTION

- Syntax | 29134
- Hierarchy Level | 29135
- Description | 29135
- Options | 29135
- Required Privilege Level | 29135
- Release Information | 29135

Syntax

```
term term-name {
    from {
        application-sets set-name;
        applications [ application-names ];
        destination-address (address | any-unicast) <except>;
        destination-address-range low minimum-value high maximum-value <except>;
        destination-port range high maximum-value low minimum-value;
        source-address (address | any-unicast) <except>;
        source-address-range low minimum-value high maximum-value <except>;
    }
    then {
        no-translation;
        port-forwarding-mappings map-name;
        translated {
            address-pooling paired;
            clat-prefix clat-prefix;
            destination-pool nat-pool-name;
            destination-prefix destination-prefix;
            dns-alg-pool dns-alg-pool;
            dns-alg-prefix dns-alg-prefix;
            filtering-type endpoint-independent;
            mapping-type endpoint-independent;
            source-pool nat-pool-name;
            source-prefix source-prefix;
```

```
            translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | deterministic-
napt44 | deterministic-napt64 |dynamic-nat44 | napt-44 | napt-66 | napt-pt | stateful-nat464 |
stateful-nat64 | twice-basic-nat-44 | twice-dynamic-nat-44 | twice-napt-44);
            }
        }
        syslog;
    }
}
```

Hierarchy Level

```
[edit services nat rule rule-name]
```

Description

Define the NAT term properties.

Options

term-name	Identifier for the term.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

term (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 29136
- Hierarchy Level | 29136
- Description | 29137
- Options | 29137
- Required Privilege Level | 29137
- Release Information | 29137

Syntax

```
term term-name {
    from {
        application-sets set-name;
        applications [ application-names ];
        destination-address (address | any-ipv4 | any-ipv6 | any-unicast) <except>;
        destination-address-range low minimum-value high maximum-value <except>;
        destination-prefix-list list-name <except>;
        source-address (address | any-ipv4 | any-ipv6 | any-unicast) <except>;
        source-address-range low minimum-value high maximum-value <except>;
        source-prefix-list list-name <except>;
    }
    then {
        (accept | discard | reject);
        syslog;
    }
}
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name]
```

Description

Define the stateful firewall term properties.

Options

term-name—Identifier for the term.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Stateful Firewall Rules

term (Simple Filter)

IN THIS SECTION

- Syntax | 29138
- Hierarchy Level | 29138
- Description | 29138
- Options | 29138
- Required Privilege Level | 29138
- Release Information | 29139

Syntax

```
term term-name {
    from {
        match-conditions;
    }
    then {
        forwarding-class class-name;
        loss-priority (high | low | medium);
    }
}
```

Hierarchy Level

```
[edit firewall family inet simple-filter filter-name]
```

Description

Define a simple filter term.

Options

`from`—Match packet fields to values. If the `from` option is not included, all packets are considered to match and the actions and action modifiers in the `then` statement are taken.

`match-conditions`—One or more conditions to use to make a match.

`term-name`—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

`then`—Actions to take on matching packets. If the `then` option is not included and a packet matches all the conditions in the `from` statement, the packet is accepted. For CoS, only the actions listed are allowed. These statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Multifield Classification

Simple Filter Overview

Firewall Filter Match Conditions for IPv4 Traffic

Firewall Filter Match Conditions for IPv6 Traffic

term (TDF Domain Selection)

IN THIS SECTION

- Syntax | 29139
- Hierarchy Level | 29142
- Description | 29142
- Options | 29142
- Required Privilege Level | 29143
- Release Information | 29143

Syntax

```
term term-name {
    from {
        3gpp-imsi {
            equals value;
            has-prefix value;
            has-suffix value;
            matches value;
        }
        attribute name {
```

```
            code numeric-code;
            vendor-id vendor-id;
            format {
                integer {
                    equals {
                        value;
                    }
                    greater-than value;
                    less-than value;
                }
                string {
                    equals {
                        value;
                    }
                    has-prefix{
                        value;
                    }
                    has-suffix {
                        value;
                    }
                    matches {
                        value;
                    }
                }
                time {
                    equals {
                        value;
                    }
                    greater-than value;
                    less-than value;
                }
                v4address {
                    equals {
                        value;
                    }
                }
                v6address {
                    equals {
                        value;
                    }
                }
                v6prefix {
                    equals {
```

```
                        value;
                    }
                }
            }
        }
        called-station-id {
            equals value;
            matches value;
        }
        calling-station-id {
            equals value;
            matches value;
        }
        class {
            equals value;
            has-prefix value;
            has-suffix value;
            matches value;
        }
        client client-name;
        framed-ip-address {
            equals value;
        }
        framed-ipv6-prefix {
            equals value;
        }
        nas-ip-address {
            equals value;
        }
        snoop-segment snoop-segment-name;
        user-name {
            equals value;
            has-prefix value;
            has-suffix value;
            matches value;
        }
    }
    then {
        domain tdf-domain-name;
        pcef-profile pcef-profile-name;
    }
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection]
```

Description

Configure a term that can be used to select the TDF domain for an IP-based subscriber. You can configure multiple terms (up to 10 terms) for the TDF domain selection, and each term is applied in the order in which it is configured. You can specify multiple match conditions within the `from` statement of a term, and all of the conditions have to match. If the incoming RADIUS request from the subscriber matches the criteria in a term, then the TDF domain specified in the `then` statement of the term is used to create the TDF subscriber session.

A term can also be used to select a PCEF profile for a an IP-based subscriber. Setting up a term so that it is used to select a profile is required if the TDF domain selected for a subscriber does not specify a PCEF profile or you want to allow different members of the same TDF domain to have different PCEF profiles.

After a term matches and a TDF domain is selected, further terms are not evaluated when the PCEF profile is specified in either the `then` statement or in the selected TDF domain. If a PCEF profile is not specified in either the `then` statement or in the selected TDF domain, further terms are evaluated to find a PCEF profile for the subscriber.

If no TDF domain is selected for a subscriber, then a TDF subscriber session is not created.

> **NOTE**: The TDF domain must have been previously configured at the `[edit unified-edge gateways tdf gateway-name domains]` hierarchy level.
>
> The PCEF profile must have been previously configured at the `[edit unified-edge pcef]` hierarchy level.

Options

term-name Identifier for the term.

- **Range:** 1 through 50 alphanumeric characters.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: You must configure at least one term.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

term (URL Filter)

IN THIS SECTION

- Syntax | **29144**
- Hierarchy Level (starting in Junos OS Release 18.3R1) | **29144**
- Hierarchy Level (before Junos OS Release 18.3R1) | **29144**
- Description | **29144**
- Options | **29144**
- Required Privilege Level | **29145**
- Release Information | **29145**

Syntax

```
term term-name {
    from {
        src-ip-prefix [prefix1 prefix2];
        dest-port   [port1 port2];
    }
    then {
        accept;
        custom-page custom-page;
        http-status-code http-status-code;
        redirect-url redirect-url;
        tcp-reset;
    }
}
```

Hierarchy Level (starting in Junos OS Release 18.3R1)

```
[edit services web-filter profile profile-name url-filter-template template-name]
```

Hierarchy Level (before Junos OS Release 18.3R1)

```
[edit services url-filter profile profile-name template template-name]
```

Description

Define a URL filtering term. A term is a set of match criteria with actions to be taken if the match criteria is met. You must configure `term` to configure URL filtering.

Options

term-name	Name of the term.
from	Define match criteria.

The `from` statement is optional. If you omit the `from` statement, all source IP prefixes and all destination ports are considered to match. All such combinations then take the configured actions of the term.

Only one term in a template can have an optional `from` statement. If you omit more than one `from` statement per template, you will get the following error message on commit:

```
URLFD_CONFIG_FAILURE: Configuration not valid:
Cannot have two wild card terms in template template1
error: configuration check-out failed
```

Similarly, no two templates within a profile can have a term without a `from` statement.

dest-ports Destination port list specification.

- **Range:** 1 through 65535

src-ip-prefix Source IP prefix list specification.

then Specify one of the following actions to be taken if the `from` condition is matched:

accept Accept the traffic and allow it to flow as normal.

custom-page *custom-page* Custom-page string.

http-status-code *http-status-code* HTTP status code value.

- **Range:** 400 through 599

redirect-url *redirect-url* URL to redirect traffic to.

tcp-reset Reset TCP.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

terminal

IN THIS SECTION

- Syntax | 29146
- Hierarchy Level | 29146
- Description | 29146
- Required Privilege Level | 29146
- Release Information | 29146

Syntax

```
terminal;
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name]
```

Description

Set or unset a terminal rule flag. The device stops matching rules for a session when a terminal rule is matched.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

terminate-code

IN THIS SECTION

- Syntax | 29147
- Hierarchy Level | 29147
- Description | 29147
- Options | 29148
- Required Privilege Level | 29149
- Release Information | 29149

Syntax

```
terminate-code (aaa (deny | service-shutdown | shutdown) | dhcp | l2tp | ppp | vlan) term-reason
radius value;
```

Hierarchy Level

```
[edit access]
```

Description

Customize the mapping between a termination cause (the internal termination identifier) and a numerical code value for the cause that is reported in the RADIUS Acct-Terminate-Cause attribute (49).

When a RADIUS Acct-Stop message is issued as a result of the termination of a subscriber or service session, the RADIUS Acct-Terminate-Cause attribute (49) reports the cause or reason for the termination. This attribute is included only in RADIUS Acct-Stop messages. The termination cause is conveyed as a code value in the attribute. *RFC 2866, RADIUS Accounting,* defines the standard mapping between 18 code values and termination causes.

Junos OS defines a set of internal termination cause codes for AAA, DHCP, L2TP, PPP, and VLAN subscriber and service session failures. By default, these internal cause codes are mapped to the RFC-defined code values. When a subscriber or service session is terminated, the router logs a message for

the internal termination cause and logs another message for the RADIUS Acct-Terminate-Cause attribute. You can use the logged information to help monitor and troubleshoot the events.

Because there are many different Junos OS internal identifiers for termination causes and only 18 supported, RFC-defined standard code values, by default a given code value can map to multiple identifiers. Instead of using the default code values, you can use the `terminate-code` statement to map any of the internally defined termination causes to any 32-bit number (1 through 4,294,967,295). The flexibility of customized mapping greatly increases the possibilities for fine-grained analytics and failure tracking.

Options

aaa Map internal identifiers for AAA-specific termination causes to a numerical value.

deny Limit selection of termination causes to those associated with denial of subscriber access.

dhcp Map internal identifiers for DHCP-specific termination causes to a numerical value.

l2tp Map internal identifiers for L2TP-specific termination causes to a numerical value.

ppp Map internal identifiers for PPP-specific termination causes to a numerical value.

radius ***value*** Number that represents the termination cause in the RADIUS Acct-Terminate-Cause attribute (49).

- **Range:** 1 through 4,294,967,295

service-shutdown Limit selection of termination causes to those associated with established service sessions independent of the parent subscriber session.

shutdown Limit selection of termination causes to those associated with established subscriber sessions.

vlan Map internal identifiers for VLAN-specific termination causes to a numerical value.

term-reason Internal identifier for the termination causes defined for the specified protocol type.

For protocol-specific termination causes, see the following topics:

- *AAA Termination Causes and Code Values*
- *DHCP Termination Causes and Code Values*
- *L2TP Termination Causes and Code Values*

- *PPP Termination Causes and Code Values*
- *VLAN Termination Causes and Code Values*

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

`vlan` option added in Junos OS Release 16.1.

RELATED DOCUMENTATION

Session Termination Causes and RADIUS Termination Cause Codes

Session Termination Causes and RADIUS Termination Cause Codes

terminate (PPPoE Service Name Tables)

IN THIS SECTION

- Syntax | 29150
- Hierarchy Level | 29150
- Description | 29150
- Required Privilege Level | 29150
- Release Information | 29150

Syntax

```
terminate;
```

Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name],
[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci
circuit-id-string ari remote-id-string]
```

Description

Direct the router to immediately respond to a PPPoE Active Discovery Initiation (PADI) control packet received from a PPPoE client by sending the client a PPPoE Active Discovery Offer (PADO) packet. The PADO packet contains the name of the access concentrator (router) that can service the client request. The `terminate` action is the default action for a named service entry, `empty` service entry, `any` service entry, or agent circuit identifier/agent remote identifier (ACI/ARI) pair in a PPPoE service name table.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at `[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci circuit-id-string ari remote-id-string]` hierarchy level introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

termination (Services)

IN THIS SECTION

- Syntax | 29151
- Hierarchy Level | 29151
- Description | 29151
- Options | 29152
- Required Privilege Level | 29152
- Release Information | 29152

Syntax

```
termination {
    profile name {
        custom-ciphers;
        enable-flow-tracing   enable-flow-tracing;
        enable-session-cache enable-session-cache;
        preferred-ciphers (custom | medium | strong | weak);
        protocol-version (all | ssl3 | tls1 | tls11 | tls12);
        server-certificate  server-certificate;
        trusted-ca  ;
    }
}
```

Hierarchy Level

```
[edit services ssl]
```

Description

Specify the configuration for Secure Socket Layer (SSL) termination support service.

Following types of SSL profiles are supported on SRX Series to secure connections based on the role of the SRX Series Firewall:

- SSL initiation: The SRX Series Firewall, acting as an SSL proxy client, initiates and maintains SSL sessions between itself and an SSL server. SRX Series Firewall receives unencrypted data from an HTTP client, and encrypts and transmits the data as ciphertext to the SSL server.
- SSL termination: The SRX Series Firewall, actings as an SSL proxy server, terminates the SSL session from the client and then establishing a new SSL connection to the server. The SRX Series Firewall decrypts the data and then sends the data as un-encrypted request to the other servers (HTTP server).

The SSL proxy profile will be applied to the security policy as application services.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. The `protocol-version` statement is updated to include `tls11` and `tls12` from Junos OS Release 15.1X49-D30.

test

IN THIS SECTION

- Junos OS | **29153**
- Junos OS Evolved | **29154**
- Junos OS Hierarchy Level | **29155**
- Junos OS Evolved Hierarchy Level | **29155**
- Description | **29155**

- Options | 29155
- Required Privilege Level | 29155
- Release Information | 29155

Junos OS

```
test test-name {
    data-fill data;
    data-size size;
    destination-interface interface-name;
    destination-port port;
    dscp-code-points dscp-bits;
    hardware-timestamp;
    history-size size;
    inet6-options {
        source-address address;
    }
    moving-average-size number;
    next-hop next-hop;
    one-way-hardware-timestamp;
    probe-count count;
    probe-interval seconds;
    probe-type type;
    routing-instance instance-name;
    rpm-scale {
        destination {
            interface interface-name.logical-unit-number;
            subunit-cnt subunit-cnt;
        }
        source {
            address-base ipv4-address-base;
            count ipv4-count;
            step ipv4-step;
        }
        source-inet6 {
            address-base ipv6-address-base;
            count ipv6-count;
            step ipv6-step;
```

```
        }
        target {
            address-base ipv4-address-base;
            count ipv4-count;
            step ipv4-step;
        }
        target-inet6 {
            address-base ipv6-address-base;
            count ipv6-count;
            step ipv6-step;
        }
        tests-count tests-count;
    }
    source-address address;
    target (url url | address address);
    test-interval interval;
    thresholds thresholds;
    traps traps;
    ttl hop-count
}
```

Junos OS Evolved

```
test test-name {
    data-fill data;
    data-size size;
    destination-port port;
    dscp-code-points dscp-bits;
    history-size size;
    moving-average-size number;
    offload-type {
        none;
        pfe-timestamp;
    }
    probe-count count;
    probe-interval seconds;
    probe-type type;
    routing-instance instance-name;
    source-address address;
    target (url url | address address);
    test-interval interval;
```

```
    thresholds thresholds;
    traps traps;
    ttl hop-count
}
```

Junos OS Hierarchy Level

```
[edit services rpm probe owner]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name]
```

Description

Specify the range of probes over which the standard deviation, average, and jitter are calculated. The test name combined with the owner name represent a single RPM configuration instance.

Options

`test` *`name`*—You can configure any name up to 32 characters in length. For Junos OS Evolved, if the owner name is one of the pre-defined names, then four pre-defined test names are available to choose from. For example, if the owner name is the pre-defined name `icmp-evo`, then these four pre-defined test names are available to configure: `icmp-evo-1`, `icmp-evo-2`, `icmp-evo-3`, and `icmp-evo-4`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`inet6-options` option added in Junos OS Release 14.1R4 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

`traps` option introduced in Junos OS Evolved Release 21.2R1.

RELATED DOCUMENTATION

Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, and MX Series Routers

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

test aaa authd-lite user

IN THIS SECTION

- Syntax | 29156
- Description | 29156
- Options | 29157
- Required Privilege Level | 29157
- Output Fields | 29157
- Sample Output | 29158
- Release Information | 29161

Syntax

```
test aaa authd-lite user username password password profile access-profile-name
<port nas-port>
<zero-stats>
```

Description

Verify authd-lite subscriber access authentication, accounting, and address allocation configuration.

The `test aaa` command supports all RADIUS-sourced attributes, both IETF standard attributes and Juniper Networks VSAs. Received attributes are displayed in the output. For information about standard RADIUS attributes, see *Standard and Vendor-Specific RADIUS Attributes*. For information about Juniper Networks VSAs, see *Standard and Vendor-Specific RADIUS Attributes*.

Starting in Junos OS Release 19.3R1, the XML output format has changed. Each RADIUS server attribute name has an associated attribute value. Each of these pairs is now enclosed by the <radius-server-data> tag. The new tag makes it easier to recognize the name/value pairs, both for operators and API clients. You may have to change any scripts that use the XML output to work properly with the new format.

Options

username	Specify the subscriber username to test.
password ***password***	Specify the password associated with the username.
profile ***access-profile-name***	Specify the access profile associated with the subscriber.
port ***nas-port***	(Optional) Specify the NAS port used for the test.
zero-stats	(Optional) Specify that no accounting statistics are set for this test.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. For information about output fields related to authentication, accounting, and subscriber-specific information, see the **show network-access aaa statistics**, **show network-access aaa statistics authentication**, **show network-access aaa subscribers**, and **show subscribers** commands.

The **test** command does not support volume-time accounting. If volume-time accounting is configured for the test subscriber, the **test** command replaces the statistics with time-only accounting statistics.

This command displays only attributes that are supported by Junos OS; these attributes appear even when their values are not set. The Virtual Router Name (LS:RI) field matches the Juniper Networks Virtual-Router VSA (26-1), if present; otherwise the field displays default:default. The displayed value for all other attributes that are not received is `<not set>`.

Sample Output

test aaa authd-lite user

The following example tests the configuration for authd-lite subscriber user1bt with a password of $ABC123 and an access profile of employee12, and displays the resulting output:

```
user@host> test aaa authd-lite user user1bt password $ABC123 profile employee12
    Authentication Grant
    ************User Attributes***********
         User Name -                              user1bt
         Framed IPv6 Prefix -                     ::/0
         Framed IPv6 Pool -                       NULL
         Nas IPv6 Address -                       ::
         NDRA IPv6 Prefix -                       NULL
         Login IPv6 Host -                        ::
         Framed Interface Id -                    0:0:0:0
         Delegated IPv6 Prefix -                  ::/0
         NDRA IPv6 Pool -                         NULL
         User Password -                          $ABC123
         Nas Ip Address -                         0.0.0.0
         NAS Port -                               0
         Service Type-                            0
         Framed IP Address -                      0.0.0.0
         Framed IP Netmask -                      0.0.0.0
         Filter Id -                              NULL
         Framed MTU -                             0
         Reply Message -                          NULL
         Framed Route-                           <not set>
         Framed MTU -                             0
         Class -                                  SBR2CL
         Virtual Router Name (LS:RI)              default:default
         Primary DNS IP Address -                 0.0.0.0
         Secondary DNS IP Address -               0.0.0.0
         Primary WINS IP Address -                0.0.0.0
         Secondary WINS IP Address -              0.0.0.0
         Ingress Statistics -                     disabled
         Egress Statistics -                      disabled
         Ingress Policy Name -                   <not set>
         Engress Policy Name -                   <not set>
         IGMP Enable -                             disabled
         Redirect VR Name (LS:RI)                 default:default
```

```
        Service Bundle                          <not set>
        Framed Ip Route Tag                     <not set>
        LI Action                                0
        LI Interception Identifier               0
        LI Mediation Device IP Address           0.0.0.0
        LI_Mediation_Device_Port_Number          0
        Activate Service                         NULL
        Deactivate Service                       NULL
        Service Statistics                       0
        Ignore_DF_Bit -                          disabled
        IGMP Access Group Name                  <not set>
        IGMP Access Source Group_Name -         <not set>
        MLD Access Group Name                   <not set>
        MLD Access Source Group Name            <not set>
        MLD Version -                            MLD Version not set
        IGMP Version                             IGMP Version not set
        IGMP Immediate Leave -                  <not set>
        MLD Immediate Leave -                   <not set>
        IPv6_Ingress_Policy_Name -              <not set>
        IPv6_Egress_Policy_Name -               <not set>
        Cos_Parameter_Type -                    <not set>
        Service Interim Acct Interval            0
        Max Clients Per Interface                0
        Cos_Scheduler_Pmt_Type                  <not set>
        Session Timeout                          599999940
        NAS Port Type                            0
        Framed Pool                              NULL
        Idle Timeout                             0
    Acct-start sent
    Acct-start succeeded
    Pausing 10 seconds
    Interim-Acct sent
    Acct-interim succeeded
    Pausing 10 seconds
    Acct-stop sent
    Acct-stop succeeded
    Logging out subscriber
    Test complete. Exiting
```

test aaa authd-lite user (XML Output, Old Format)

The following example shows an excerpt of sample XML output in the old format:

```
user@host>test aaa authd-lite user user45@test.net password $ABC123 profile test | display xml

<rpc-reply xmlns:junos="namespace-URL"
    <aaa-test-result>
        <aaa-test-status>Authentication Grant</aaa-test-status>
        <aaa-test-status>************User Attributes************</aaa-test-status>
        <radius-server-attribute-name>User Name -</radius-server-attribute-name>
        <radius-server-attribute-value>user45@test.net</radius-server-attribute-value>
        <radius-server-attribute-name>Framed IPv6 Prefix -</radius-server-attribute-name>
        <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        <radius-server-attribute-name>Framed IPv6 Pool -</radius-server-attribute-name>
        <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        <radius-server-attribute-name>NDRA IPv6 Prefix -</radius-server-attribute-name>
        <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
...
        <aaa-test-status>Test complete. Exiting</aaa-test-status>
    </aaa-test-result>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

test aaa authd-lite user (XML Output, New Format)

The following example shows an excerpt of sample XML output in the new format:

```
user@host>test aaa authd-lite user user45@test.net password $ABC123 | display xml

<rpc-reply xmlns:junos="namespace-URL">
    <aaa-test-result>
        <aaa-test-status>Authentication Grant</aaa-test-status>
        <aaa-test-status>************User Attributes************</aaa-test-status>
        <radius-server-data>
            <radius-server-attribute-name>User Name -</radius-server-attribute-name>
            <radius-server-attribute-value>user45@test.net</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
```

```
            <radius-server-attribute-name>Framed IPv6 Prefix -</radius-server-attribute-name>
            <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Framed IPv6 Pool -</radius-server-attribute-name>
            <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>NDRA IPv6 Prefix -</radius-server-attribute-name>
            <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        </radius-server-data>
...
        <aaa-test-status>Test complete. Exiting</aaa-test-status>
    </aaa-test-result>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

AAA Testing and Troubleshooting

Understanding External Authentication Servers

test aaa dhcp user

IN THIS SECTION

- Syntax | 29162
- Description | 29162
- Options | 29163

- Required Privilege Level | 29164
- Output Fields | 29164
- Sample Output | 29165
- Release Information | 29168

Syntax

```
test aaa dhcp user username
<agent-remote-id ari>
<logical-system logical-system-name>
<mac-address mac-address>
<no-address-request>
<option-82 option-82>
<password password>
<profile access-profile-name>
<routing-instance routing-instance-name>
<service-type service-type>
<source-address source-address>
<terminate-code code-value>
```

Description

Verify Dynamic Host Configuration Protocol (DHCP) subscriber access authentication, accounting, and address allocation configuration by creating a test pseudo session.

NOTE: The `test aaa` command supports all RADIUS-sourced attributes, both IETF standard attributes and Juniper Networks VSAs. Received attributes are displayed in the output. For information about standard RADIUS attributes, see *Standard and Vendor-Specific RADIUS Attributes*. For information about Juniper Networks VSAs, see *Standard and Vendor-Specific RADIUS Attributes*.

NOTE: Starting in Junos OS Release 19.3R1, the XML output format has changed. Each RADIUS server attribute name has an associated attribute value. Each of these pairs is now enclosed by the <radius-server-data> tag. The new tag makes it easier to recognize

the name/value pairs, both for operators and API clients. You may have to change any scripts that use the XML output to work properly with the new format.

Options

username — Subscriber username to test.

agent-remote-id *ari* — (Optional) Value of the DSL Forum Agent-Remote-Id (VSA 26–2).

logical-system *logical-system-name* — (Optional) Logical system in which the subscriber is authenticated. This is the logical system in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25).

mac-address *mac-address* — (Optional) MAC address of the DHCP client.

no-address-request — (Optional) Request is sent for authentication without address allocation. Use for Layer 2-only scenarios where no address allocation request is needed.

> **NOTE**: The `test aaa dhcp user` command tries to allocate an IPv4 address even when the subscriber is supposed to get only an IPv6 address. If that behavior is undesirable, include the `no-address-request option` when you issue the command.

option-82 *option-82* — (Optional) DHCP relay agent information option (option-82) value.

password *password* — (Optional) Password associated with the username.

profile *access-profile-name* — (Optional) Access profile associated with the subscriber.

routing-instance *routing-instance-name* — (Optional) Routing instance in which the subscriber is authenticated. This is the routing instance in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25). In the case of VSA 26-25, the subscriber is re-authenticated in the subscriber context.

service-type *service-type* — (Optional) Value of the Service Type RADIUS attribute [6] that is associated with the test user; either a number in the range 1 through 255 or one of the following strings that corresponds to an RFC-defined service type; the numbers are the values that are carried in the RADIUS attribute to specify the service:

administrative (6)	callback-nas-prompt (9)
authenticate-only (8)	framed (2)
call-check (10)	login (1)
callback-admin (11)	nas-prompt (7)
callback-framed (4)	outbound (5)
callback-login (3)	–

source-address *source-address* (Optional) IP address of the outgoing interface.

terminate-code *code-value* (Optional) Code associated with the subscriber termination.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. For information about output fields related to authentication, accounting, and subscriber-specific information, see the **show network-access aaa statistics**, **show network-access aaa statistics authentication**, **show network-access aaa subscribers**, and **show subscribers** commands.

The **test** command does not support volume-time accounting. If volume-time accounting is configured for the test subscriber, the **test** command replaces the statistics with time-only accounting statistics.

This command displays only attributes that are supported by Junos OS; these attributes appear even when their values are not set. The Virtual Router Name (LS:RI) field matches the Juniper Networks Virtual-Router VSA (26-1), if present; otherwise the field displays default:default. The displayed value for all other attributes that are not received is `<not set>`.

Sample Output

test aaa dhcp user

The following example tests the configuration for DHCP subscriber user1DB and password $ABC123, and displays the resulting output:

```
user@host> test aaa dhcp user user1DB@test.net password $ABC123
Authentication Grant
    ************User Attributes***********
         User Name -                              user1DB@test.net
         Client IP Address -                      192.168.1.1
         Client IP Netmask -                      255.255.0.0
         Virtual Router Name (LS:RI)-             default:default

         Agent Remote Id -                        NULL
         Reply Message -                          NULL
         Primary DNS IP Address -                 0.0.0.0
         Secondary DNS IP Address -               0.0.0.0
         Primary WINS IP Address -                0.0.0.0
         Secondary WINS IP Address -              0.0.0.0
         Primary DNS IPv6 Address  -              ::
         Secondary DNS IPv6 Address  -            ::
         Framed Pool -                           <not set>
         Service Type -                           0
         DHCP Guided Relay Server -               0
         Class Attribute -                        TEST
         Client IPv6 Address -                    ::
         Client IPv6 Mask -                       null
         Framed IPv6 Prefix -                     ::/0
         Framed IPv6 Pool -                      <not-set>
         NDRA IPv6 Prefix -                      <not-set>
         Login IPv6 Host -                        ::
         Framed Interface Id -                    0:0:0:0
         Delegated IPv6 Prefix -                  ::/0
         Delegated IPv6 Pool -                   <not-set>
         User Password -                          $ABC123
         CHAP Password -                          NULL
         Mac Address -                            00:00:5E:00:53:ab
         Idle Timeout -                           600
         Session Timeout -                        6000
         Service Name (1) -                       cos-service(video_sch, nc_sch)
```

```
Service Statistics (1) -                1
Service Acct Interim (1) -              600
Service Activation Type (1) -           1
Service Name (2) -                      filter-service(in_filter, out_filter)
Service Statistics (2) -                2
Service Acct Interim (2) -              900
Service Activation Type (2) -           1
Cos shaping rate -                      100m
Filter Id -                            <not set>
Framed MTU -                            (null)
Framed Route -                         <not set>
Ingress Policy Name -                  <not set>
Egress Policy Name -                   <not set>
IGMP Enable -                           disabled
Redirect VR Name (LS:RI)-               default:default
Service Bundle -                        Null
Framed Ip Route Tag -                  <not set>
Ignore DF Bit -                         disabled
IGMP Access Group Name -               <not set>
IGMP Access Source Group Name -        <not set>
MLD Access Group Name -                <not set>
MLD Access Source Group Name -         <not set>
IGMP Version -                         <not set>
MLD Version -                          <not set>
IGMP Immediate Leave -                 <not set>
MLD Immediate Leave -                  <not set>
IPv6 Ingress Policy Name -             <not set>
IPv6 Egress Policy Name -              <not set>
Dynamic Profile -                      <not set>
Acct Session ID -                       1
Acct Interim Interval -                 750
Acct Type -                             1
Ingress Statistics -                    disabled
Egress Statistics -                     disabled
Chargeable user identity -              0
NAS Port Id -                           -0/0/0.0
NAS Port -                              4095
NAS Port Type -                         15
Framed Protocol -                       1
IPv4 ADF Rule -                         010100
IPv4 ADF Rule -                         010101
IPv6 ADF Rule -                         030100
IPv6 ADF Rule -                         030101
```

```
    ****Pausing 10 seconds before disconnecting the test user*********
    Logging out subscriber
         Terminate Id -                          <not set>
    Test complete. Exiting
```

test aaa dhcp user (XML Output, Old Format)

The following example shows an excerpt of sample XML output in the old format:

```
user@host>test aaa dhcp user user45@test.net password $ABC123 profile test | display xml

<rpc-reply xmlns:junos="namespace-URL"
    <aaa-test-result>
        <aaa-test-status>Authentication Grant</aaa-test-status>
        <aaa-test-status>************User Attributes***********</aaa-test-status>
        <radius-server-attribute-name>User Name -</radius-server-attribute-name>
        <radius-server-attribute-value>user45@test.net</radius-server-attribute-value>
        <radius-server-attribute-name>Virtual Router Name (LS:RI) -</radius-server-attribute-
name>
        <radius-server-attribute-value>default:default</radius-server-attribute-value>
        <radius-server-attribute-name>Client IP Address -</radius-server-attribute-name>
        <radius-server-attribute-value>198.51.100.7</radius-server-attribute-value>
        <radius-server-attribute-name>Client IP Netmask -</radius-server-attribute-name>
        <radius-server-attribute-value>255.255.255.255</radius-server-attribute-value>

...
        <aaa-test-status>Test complete. Exiting</aaa-test-status>
    </aaa-test-result>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

test aaa dhcp user (XML Output, New Format)

The following example shows an excerpt of sample XML output in the new format:

```
user@host>test aaa dhcp user user45@test.net password $ABC123 | display xml

<rpc-reply xmlns:junos="namespace-URL">
```

```
    <aaa-test-result>
        <aaa-test-status>Authentication Grant</aaa-test-status>
        <aaa-test-status>************User Attributes************</aaa-test-status>
        <radius-server-data>
            <radius-server-attribute-name>User Name -</radius-server-attribute-name>
            <radius-server-attribute-value>user45@test.net</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Virtual Router Name (LS:RI) -</radius-server-attribute-
name>
            <radius-server-attribute-value>default:default</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Client IP Address -</radius-server-attribute-name>
            <radius-server-attribute-value>198.51.100.7</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Client IP Netmask -</radius-server-attribute-name>
            <radius-server-attribute-value>255.255.255.255</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>

...
        <aaa-test-status>Test complete. Exiting</aaa-test-status>
    </aaa-test-result>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

Release Information

Command introduced in Junos OS Release 11.2.

Option `terminate-code` added in Junos OS Release 11.4.

Option `agent-remote-id` added in Junos OS Release 14.1.

Options `no-address-request` and `service-type` added in Junos OS Release 16.1.

RELATED DOCUMENTATION

AAA Testing and Troubleshooting

test aaa ppp user

IN THIS SECTION

- Syntax | 29169
- Description | 29169
- Options | 29170
- Required Privilege Level | 29172
- Output Fields | 29172
- Sample Output | 29172
- Release Information | 29178

Syntax

```
test aaa ppp user username
<agent-remote-id ari>
<logical-system logical-system-name>
<no-address-request>
<password password>
<profile access-profile-name>
<routing-instance routing-instance-name>
<service-type service-type>
<terminate-code code-value>
```

Description

Verify Point-to-Point Protocol (PPP) subscriber access authentication, accounting, and address allocation configuration by creating a test pseudo session.

NOTE: The `test aaa` command supports all RADIUS-sourced attributes, both IETF standard attributes and Juniper Networks VSAs. Received attributes are displayed in the output. For information about standard RADIUS attributes, see *Standard and Vendor-Specific RADIUS Attributes*. For information about Juniper Networks VSAs, see *Standard and Vendor-Specific RADIUS Attributes*.

NOTE: Starting in Junos OS Release 19.3R1, the XML output format has changed. Each RADIUS server attribute name has an associated attribute value. Each of these pairs is now enclosed by the <radius-server-data> tag. The new tag makes it easier to recognize the name/value pairs, both for operators and API clients. You may have to change any scripts that use the XML output to work properly with the new format.

Options

username Subscriber username to test.

agent-remote-id ***ari*** (Optional) Value of the DSL Forum Agent-Remote-Id (VSA 26–2).

logical-system ***logical-system-name*** (Optional) Logical system in which the subscriber is authenticated. This is the logical system in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25).

no-address-request (Optional) Request is sent for authentication without address allocation. Use for Layer 2-only scenarios where no address allocation request is needed.

> **NOTE**: The `test aaa ppp user` command tries to allocate an IPv4 address even when the subscriber is supposed to get only an IPv6 address. If that behavior is undesirable, include the `no-address-request option` when you issue the command.

password ***password*** (Optional) Password associated with the username.

profile ***access-profile-name*** (Optional) Access profile associated with the subscriber.

> **NOTE**: The system logically treats this profile as a client-level configuration. An access profile configured in a domain map takes

precedence over client-level configurations. If you have configured one or more domain maps, the username for the user under test is evaluated against the domain maps the same as any other subscriber.

For example, the username can exactly match a domain map or partially match a wildcard domain map. If it matches neither of those, then it matches the `default` domain map if it is configured. If the username has no domain or realm ,then it matches the `none` domain map, if it is configured.

The consequence is that if the test user matches any configured domain map, then an access profile configured in that map is used for the test in preference to an access profile that you specify with the `test` command.

See Specifying an Access Profile in a Domain Map for more information about domain maps and access profiles.

routing-instance *routing-instance-name*

(Optional) Routing instance in which the subscriber is authenticated. This is the routing instance in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25). In the case of VSA 26-25, the subscriber is re-authenticated in the subscriber context.

service-type *service-type*

(Optional) Value of the Service Type RADIUS attribute [6] that is associated with the test user; either a number in the range 1 through 255 or one of the following strings that corresponds to an RFC-defined service type; the numbers are the values that are carried in the RADIUS attribute to specify the service:

administrative (6)	callback-nas-prompt (9)
authenticate-only (8)	framed (2)
call-check (10)	login (1)
callback-admin (11)	nas-prompt (7)
callback-framed (4)	outbound (5)
callback-login (3)	–

terminate-code *code-value*	(Optional) Code associated with the subscriber termination.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. For information about output fields related to authentication, accounting, and subscriber-specific information, see the **show network-access aaa statistics**, **show network-access aaa statistics authentication**, **show network-access aaa subscribers**, and **show subscribers** commands.

The **test** command does not support volume-time accounting. If volume-time accounting is configured for the test subscriber, the **test** command replaces the statistics with time-only accounting statistics.

This command displays only attributes that are supported by Junos OS; these attributes appear even when their values are not set. The Virtual Router Name (LS:RI) field matches the Juniper Networks Virtual-Router VSA (26-1), if present; otherwise the field displays default:default. The displayed value for all other attributes that are not received is `<not set>`.

Sample Output

test aaa ppp user

The following example tests the configuration for PPP subscriber user98BEDC and password $ABC123, and displays the resulting output:

```
user@host> test aaa ppp user user98BEDC@test.net password $ABC123
Authentication Grant
    ************User Attributes***********
         User Name -                              user98BEDC@test.net
         Client IP Address -                      192.168.1.1
         Client IP Netmask -                      255.255.0.0
         Virtual Router Name (LS:RI) -            default:default
         Agent Remote Id -                        NULL
         Reply Message -                          NULL
         Primary DNS IP Address -                 0.0.0.0
         Secondary DNS IP Address -               0.0.0.0
         Primary WINS IP Address -                0.0.0.0
```

```
Secondary WINS IP Address -              0.0.0.0
Primary DNS IPv6 Address  -              ::
Secondary DNS IPv6 Address  -            ::
Framed Pool -                           <not set>
Class Attribute -                        TEST
Service Type -                           0
Client IPv6 Address -                    ::
Client IPv6 Mask -                       null
Framed IPv6 Prefix -                     ::/0
Framed IPv6 Pool -                      <not-set>
NDRA IPv6 Prefix -                      <not-set
Login IPv6 Host -                        ::
Framed Interface Id -                    0:0:0:0
Delegated IPv6 Prefix -                  ::/0
Delegated IPv6 Pool -                   <not-set>
User Password -                          $ABC123
CHAP Password -                          NULL
Mac Address -                            00:00:5E:00:53:ab
Idle Timeout -                           600
Session Timeout -                        6000
Service Name (1) -                       cos-service(video_sch, nc_sch)
Service Statistics (1) -                 1
Service Acct Interim (1) -               600
Service Activation Type (1) -            1
Service Name (2) -                       filter-service(in_filter, out_filter)
Service Statistics (2) -                 2
Service Acct Interim (2) -               900
Service Activation Type (2) -            1
Cos shaping rate -                       100m
Filter Id -                             <not set>
Framed MTU -                             (null)
Framed Route -                          <not set>
Ingress Policy Name -                   <not set>
Egress Policy Name -                    <not set>
IGMP Enable -                            disabled
Redirect VR Name (LS:RI) -               default
Service Bundle -                         Null
Framed Ip Route Tag -                   <not set>
Ignore DF Bit -                          disabled
IGMP Access Group Name -                <not set>
IGMP Access Source Group Name -         <not set>
MLD Access Group Name -                 <not set>
MLD Access Source Group Name -          <not set>
```

```
        IGMP Version -                          <not set>
        MLD Version -                           <not set>
        IGMP Immediate Leave -                  <not set>
        MLD Immediate Leave -                   <not set>
        IPv6 Ingress Policy Name -              <not set>
        IPv6 Egress Policy Name -               <not set>
        Dynamic Profile -                       <not set>
        Acct Session ID -                        1
        Acct Interim Interval -                  750
        Acct Type -                              1
        Chargeable user identity -               0
        NAS Port Id -                            -0/0/0.0
        NAS Port -                               4095
        NAS Port Type -                          15
        Framed Protocol -                        1
        IPv4 ADF Rule -                          010100
        IPv4 ADF Rule -                          010101
        IPv6 ADF Rule -                          030100
        IPv6 ADF Rule -                          030101
    ****Pausing 10 seconds before disconnecting the test user*********
    Logging out subscriber
        Terminate Id -                          <not set>
    Test complete. Exiting
```

test aaa ppp user (tunneled user)

The following example tests the configuration for PPP tunneled subscriber accounting14, with password $ABC123 and access profile finance-b, and displays the resulting output:

```
user@host> test aaa ppp user accounting14 password $ABC123 14 profile finance-b
    Authentication Grant with Tunnel Attributes
    ************Tunnel Attributes***********
        ****Tunnel Definiton -                  1
            Tunnel Medium           -           1
            Tunnel Type             -           3
            Tunnel Max Sessions     -           100
            Tunnel Server Endpoint  -           192.0.2.4
            Tunnel Client Endpoint  -           198.51.100.5
            Tunnel Server AuthId    -           rt1
            Tunnel Client AuthId    -           ts1
            Tunnel Password         -           radius
```

```
            Tunnel Assignment Id    -            til
            Tunnel Logical System   -
            Tunnel Routing Instance -
    ****Pausing 10 seconds before disconnecting the test user*********
    Logging out subscriber
         Terminate Id -                          l2tp session-receive-cdn-avp-bad-hidden
    Test complete. Exiting
```

test aaa ppp user (authentication failure)

The following example shows sample output when the authentication grant fails due to an invalid password:

```
user@host>test aaa ppp user user45@test.net password $ABC123123
 Authentication Deny
    Reason : Access Denied
    Received Attributes :
         User Name -                             user45@test.net
         Client IP Address -                     0.0.0.0
         Client IP Netmask -                     0.0.0.0
         Virtual Router Name (LS:RI)-            default
         Agent Remote Id -                       NULL
         Reply Message -                         NULL
         Primary DNS IP Address -                0.0.0.0
         Secondary DNS IP Address -              0.0.0.0
         Primary WINS IP Address -               0.0.0.0
         Secondary WINS IP Address -             0.0.0.0
         Primary DNS IPv6 Address  -             ::
         Secondary DNS IPv6 Address  -           ::
         Framed Pool -                           not set
         Class Attribute -                       not set
         Service Type -                          0
         Client IPv6 Address -                   ::
         Client IPv6 Mask -                      null
         Framed IPv6 Prefix -                    ::/0
         Framed IPv6 Pool -                      not-set
         NDRA IPv6 Prefix -                      not-set
         Login IPv6 Host -                       ::
         Framed Interface Id -                   0:0:0:0
         Delegated IPv6 Prefix -                 ::/0
         Delegated IPv6 Pool -                   not-set
```

```
        User Password -                         $ABC123123
        CHAP Password -                         NULL
        Mac Address -                           00:00:5E:00:53:ab
        Filter Id -                             not set
        Framed MTU -                            (null)
        Framed Route -                          not set
        Ingress Policy Name -                   not set
        Egress Policy Name -                    not set
        IGMP Enable-                            disabled
        Redirect VR Name (LS:RI)-               default
        Service Bundle -                        Null
        Framed Ip Route Tag -                   not set
        Ignore DF Bit -                         disabled
        IGMP Access Group Name -                not set
        IGMP Access Source Group Name -         not set
        MLD Access Group Name -                 not set
        MLD Access Source Group Name -          not set
        IGMP Version -                          not set
        MLD Version -                           not set
        IGMP Immediate Leave -                  not set
        MLD Immediate Leave -                   not set
        IPv6 Ingress Policy Name -              not set
        IPv6 Egress Policy Name -               not set
        Acct Session ID -                       12
        Acct Interim Interval -                 0
        Acct Type -                             0                        Chargeable user
identity -              0
        NAS Port Id -                           -0/0/0.0
        NAS Port -                              4095
        NAS Port Type -                         15
        Framed Protocol -                       0
    Test complete. Exiting
```

test aaa ppp user (XML Output, Old Format)

The following example shows an excerpt of sample XML output in the old format:

```
user@host>test aaa ppp user user45@test.net password $ABC123 | display xml

<rpc-reply xmlns:junos="namespace-URL"
    <aaa-test-result>
```

```
        <aaa-test-status>Authentication Grant</aaa-test-status>
        <aaa-test-status>************User Attributes***********</aaa-test-status>
        <radius-server-attribute-name>User Name -</radius-server-attribute-name>
        <radius-server-attribute-value>user45@test.net</radius-server-attribute-value>
        <radius-server-attribute-name>Virtual Router Name (LS:RI) -</radius-server-attribute-
name>
        <radius-server-attribute-value>default:default</radius-server-attribute-value>
        <radius-server-attribute-name>Service Type -</radius-server-attribute-name>
        <radius-server-attribute-value>Framed</radius-server-attribute-value>
        <radius-server-attribute-name>Agent Remote Id -</radius-server-attribute-name>
        <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
...
  <aaa-test-status>Test complete. Exiting</aaa-test-status>
    </aaa-test-result>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

test aaa ppp user (XML Output, New Format)

The following example shows an excerpt of sample XML output in the new format:

```
user@host>test aaa ppp user user45@test.net password $ABC123 | display xml

<rpc-reply xmlns:junos="namespace-URL">
    <aaa-test-result>
        <aaa-test-status>Authentication Grant</aaa-test-status>
        <aaa-test-status>************User Attributes***********</aaa-test-status>
        <radius-server-data>
            <radius-server-attribute-name>User Name -</radius-server-attribute-name>
            <radius-server-attribute-value>user45@test.net</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Virtual Router Name (LS:RI) -</radius-server-attribute-
name>
            <radius-server-attribute-value>default:default</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Service Type -</radius-server-attribute-name>
            <radius-server-attribute-value>Framed</radius-server-attribute-value>
```

```
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Agent Remote Id -</radius-server-attribute-name>
            <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        </radius-server-data>
...
        <aaa-test-status>Test complete. Exiting</aaa-test-status>
    </aaa-test-result>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

Release Information

Command introduced in Junos OS Release 11.2.

Option `terminate-code` added in Junos OS Release 11.4.

Option `agent-remote-id` added in Junos OS Release 14.1.

Options `no-address-request` and `service-type` added in Junos OS Release 16.1.

RELATED DOCUMENTATION

AAA Testing and Troubleshooting

test access profile

IN THIS SECTION

- Syntax | **29179**
- Description | **29179**
- Options | **29179**
- Required Privilege Level | **29179**
- Output Fields | **29179**

Sample Output | 29181
Release Information | 29184

Syntax

```
test access profile profile-name user username password password <detail>
```

Description

Specify a profile to use to get information from a RADIUS server, which includes all the information from the `test access radius-server` command.

Options

detail (Optional) Show the RADIUS attributes returned by the server.

profile-name Access profile name configured.

password Password for the username.

username User name to be authenticated to the RADIUS server.

Required Privilege Level

view

Output Fields

Table 1803 on page 29180 lists the output fields for the `test access profile` command. Output fields are listed in the approximate order in which they appear.

Table 1803: test access profile Output Fields

Field Name	Field Description
`Profile Name`	Name of the configured access profile.
`Client Username`	The user name authenticated by the RADIUS server.
`Client Password`	The user password authenticated by the RADIUS server.
`Num Servers`	Number of RADIUS servers in the configured access profile.
`Server List`	List of RADIUS servers in the configure access profile.
`IP Address`	The IP address of the RADIUS server authenticated.
`UDP Port`	The RADIUS server port utilized during the authentication test.
`Source Address`	The source IP address of the client making the RADIUS request. If no address is shown, it defaults to the address of the outgoing interface.
`Timeout`	The RADIUS server timeout period.
`Retry Count`	The number of authentication attempts allowed by the RADIUS server.
`Secret`	The shared secret used for authentication with the RADIUS server.
`Status`	The test result status (Accepted or Rejected) and the number of retransmits utilized during authentication.
`Attempts`	The number of authentication attempts on the RADIUS server.
`Attribute List`	The list of returned RADIUS attributes, sorted by the attribute name, and including parameter length and value. See your RADIUS server documentation for attribute descriptions.

Table 1803: test access profile Output Fields *(Continued)*

Field Name	Field Description
`(Attribute) Name`	The name of the attribute.
`(Attribute) Length`	The attribute length in bytes.
`(Attribute) Value`	The attribute value.

Sample Output

test access profile

The following example uses the `test access profile` command to access and display basic information about the RADIUS server(s) shown in the resulting output:

```
user@host> test access profile alpha user TEST password TEST
  user@host> test access profile alpha user TEST password TEST
  Test Radius Profile Access
      Profile Name        : alpha
      Client Username     : TEST
      Client Password     : TEST
      Num Servers         : 5
                            Server List
                 UDP      Source                  Retry
    IP Address   Port     Address         Timeout Count Secret        Status      Attempts
  10.1.1.1        1812     10.10.10.10     2       1     TEST          Timeout     2
  10.1.3.4        1812     Default         1       2     TEST          Timeout     3
  192.168.10.10   1812     Default         3       3     TEST          Accepted    1
```

test access profile detail

The following example uses the `test access profile detail` command to access and display detailed information about the RADIUS server(s) shown in the resulting output:

```
user@host> test access profile alpha user TEST password TEST detail
  user@host> test access profile alpha user TEST password TEST detail
  Test Radius Profile Access Detailed
      Profile Name        : alpha
      Client Username     : TEST
      Client Password     : TEST
      Num Servers         : 5
                      Radius Server List

         IP Address           : 10.1.3.4
         UDP Port             : 1812
         Source Address       : 192.168.10.10
         Timeout              : 2
         Retry Count          : 1
         Secret               : TEST
         Status               : Timeout
         Attempts             : 2

         IP Address           : 10.1.3.5
         UDP Port             : 1812
         Source Address       : Default
         Timeout              : 1
         Retry Count          : 2
         Secret               : TEST
         Status               : Timeout
         Attempts             : 3

         IP Address           : 192.168.10.10
         UDP Port             : 1812
         Source Address       : Default
         Timeout              : 3
         Retry Count          : 3
         Secret               : TEST
         Status               : Accepted
         Attempts             : 1

                               Attribute List
```

```
Name                                Length  Value
Class                               52      SBR2CLÍ
Acct-Interim-Interval               4       5
Callback-Id                         12      123-456-789
Callback-Number                     13      555-555-1212
Class                               15      Class information
Filter-Id                           4       999
Filter-Id                           6       12345
Framed-Compression                  4       0
Framed-IP-Address                   4       10:2:3:4
Framed-IP-Netmask                   4       255:255:255:255
Framed-IPv6-Route                   15      2001:db8:10:2:3:4:5:6

Framed-MTU                          4       1024
Framed-Pool                         9       pool sbr
Framed-Protocol                     4       1
Framed-Route                        8       iproute
Framed-Routing                      4       0
Vendor-Specific                     11      583
Idle-Timeout                        4       3
Vendor-Specific                     10      a4c
Vendor-Specific                     14      a4c
Login-IP-Host                       4       10:1:1:1
Login-LAT-Group                     10      lat group
Login-LAT-Node                      9       lat node
Login-LAT-Port                      9       lat port
Login-LAT-Service                   12      lat service
Login-Service                       4       0
Login-TCP-Port                      4       1812
Vendor-Specific                     10      137
Vendor-Specific                     38      137
Vendor-Specific                     10      137
Vendor-Specific                     9       137
Vendor-Specific                     16      137
Vendor-Specific                     10      137
Vendor-Specific                     10      137
Vendor-Specific                     10      137
Vendor-Specific                     9       137
Vendor-Specific                     10      137
Vendor-Specific                     10      137
Vendor-Specific                     10      137
Vendor-Specific                     10      137
Password-Retry                      4       3
```

```
    Port-Limit                          4       100
    Prompt                              4
    Reply-Message                       18      Radius Server SB
    Service-Type                        4       2
    Session-Timeout                     4       10
    Termination-Action                  4       1
    Tunnel-Assignment-ID                4
    Tunnel-Client-Auth-ID               6
    Tunnel-Client-Endpoint              4
    Tunnel-Password                     19
    Tunnel-Type                         4       12
    MS BAP Usage                        4       0
    MS-CHAP MPPE-Keys                   32      -1234567890
    MS-CHAP2 Success                    3       123456789
    MS Filter                           10      ms-filter
    MS Link Drop Time Limit             4       5
    MS Link Utilization Threshold       4       6
    MS MPPE Encryption Policy           4       1
    MS MPPE Encryption Types            3       -556677889
    MS Primary DNS Server               4       10:1:1:1
    MS Primary NBNS Server              4       10:2:2:2
    MS Secondary DNS Server             4       10:3:3:3
    MS Secondary NBNS Server            4       10:4:4:4
```

Release Information

Command introduced in Junos OS Release 9.1.

test access radius-server

IN THIS SECTION

- Syntax | 29185
- Description | 29185
- Options | 29185
- Required Privilege Level | 29185

- Output Fields | 29186
- Sample Output | 29187
- Release Information | 29187

Syntax

```
test access radius-server address user username password password secret secret
<authentication-port port>
<retry number>
<source-address address>
<timeout number>
```

Description

Verify RADIUS server authentication parameters.

Options

address	RADIUS server under test IP address.
password	Password for the user.
secret	Secret shared with the RADIUS server.
user	User name to be authenticated to the RADIUS server.
authentication-port	(Optional) RADIUS server authentication port number (1through 65535).
retry	(Optional) Retry attempts (1through 10).
source-address	(Optional) Use an alternate address as the source address.
timeout	(Optional) Request timeout period (1through 90 seconds).

Required Privilege Level

view

Output Fields

Table 1804 on page 29186 lists the output fields for the `test access radius-server` command. Output fields are listed in the approximate order in which they appear.

Table 1804: test access radius-server Output Fields

Field Name	Field Description
`Server`	The IP address of the RADIUS server authenticated.
`UDP port`	The RADIUS server port utilized during the authentication test.
`Source IP Address`	“Default” is shown if the IP address is the same as that of the RADIUS server. Alternatively, an IP address specified for authentication is shown.
`Server timeout`	The RADIUS server timeout period.
`Sever retry count`	The number of authentication attempts allowed by the RADIUS server.
`Secret`	The shared secret used for authentication with the RADIUS server.
`Client Username`	The user name authenticated by the RADIUS server.
`Client Password`	The user password authenticated by the RADIUS server.
`Status`	The test result status (`Accepted` or `Rejected`) and the number of retransmits utilized during authentication.

Sample Output

test access radius-server user password secret

The following example command tests RADIUS authentication with a specific server (172.28.30.95), user (JOHNDOE), secret (No1Knows), and password (JohnPass); and displays the resulting output:

```
user@host> test access radius-server 172.28.30.95 user JOHNDOE password JohnPass secret No1Knows
Test Radius Server Access
       Server              : 172.28.30.95
       UDP port            : 1812
       Source IP Address   : Default
       Server timeout      : 3
       Sever retry count   : 3
       Secret              : No1Knows
       Client Username     : JOHNDOE
       Client Password     : JohnPass
       Status              : Accepted, retransmits: 0
```

Release Information

Command introduced in Junos OS Release 9.1.

test-agent (Paragon Active Assurance)

IN THIS SECTION

- Syntax | 29188
- Hierarchy Level | 29188
- Description | 29188
- Default | 29189
- Options | 29189
- Required Privilege Level | 29190
- Release Information | 29190

Syntax

```
test-agent {
   cc-account user-account-name;
   cc-host DNS-name;
   cc-insecure;
   cc-password plain-text-password;
   cc-port port-number;
   cc-user "user@domain";
   ta-name test-agent-name;
   ta-version version-number;
}
```

Hierarchy Level

```
[edit services paa]
```

Description

Installs and runs an instance of the Paragon Active Assurance (PAA) test agent on the node. Once this configuration is committed, the node downloads the PAA test agent Docker image from the PAA Control Center. The node then starts the Docker daemon and loads the PAA test agent Docker image into a Docker container. The loaded software sets up the `nlsd` configuration required by the PAA test agent and then starts the PAA test agent.

Install the Paragon Active Assurance (PAA) test agent. Once the configuration for the test agent is committed, the system fetches the PAA test agent software image from the PAA Control Center for you and installs the software into a Docker container.

Before configuring this statement, make sure you have an HTTPS connection to the PAA Control Center (port 443 is open) and either port 6800 or a port of your choosing is open to connect the test agent with the PAA Control Center. In the PAA Control Center, make sure your SSL certificate is correct on both ports, and both are either secure or self-signed.

If the test agent version is prior to 4.2.0, you must configure 127.0.0.1 as the loopback address (lo0.0) on the router. For more information about loopback addresses and how to configure them, see Loopback Interfaces (Junos OS Evolved). If the test agent version is 4.2.0 or newer, you do not need to configure a loopback address.

We support installation on two VRFs, `vrf0` and `mgmt_junos`. By default, the operating system installs the PAA test agent on the `vrf0` VRF. If configuration for `mgmt_junos` is present, the operating system installs the PAA test agent on the `mgmt_junos` VRF.

If you already have a committed configuration for the PAA test agent, and you make and commit some changes to the options below, the PAA test agent is re-installed using the new changes to the configuration.

Default

For `cc-port`, 6800. For `ta-name`, the node name.

Options

cc-account *user-account-name*	(Required) Your PAA Control Center user account.
cc-host *DNS-name*	(Required) IP address or DNS host name of the PAA Control Center.
cc-insecure	(Optional) Allow an insecure connection between the test agent and the PAA Control Center. In other words, disable SSL certificate checking and use self-signed SSL certificates.
cc-password *plain-text-password*	(Required) Your PAA Control Center password. The plain-text password is encrypted and is shown as such in the configuration. But the password is decrypted before being passed along to the test agent, which uses it when authenticating towards the Control Center. The test agent sends the password itself, but inside a secure SSL connection.
cc-port *port-number*	(Optional) Port over which the router can communicate with the PAA Control Center. Make sure the port you specify here (or port 6800 if taking the default) is open. • Range: 1-65535 • Default: 6800
cc-user *"user@domain"*	(Required) The email address registered with the PAA Control Center.
ta-name *test-agent-name*	(Optional) The name you want to use to refer to this PAA test agent. The default is the router's hostname. The `ta-name` is used to identify the test agent in the PAA Control Center, and PAA won't allow two test agents with the same name to run at the same time.

ta-version *version-number*	(Optional) The PAA test agent software version you would like to install. Check the PAA Control Center for the versions available. If not specified, the operating system installs the latest PAA test agent software version from the Control Center.

Required Privilege Level

configure

Release Information

Statement introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

Install the Paragon Active Assurance (PAA) Test Agent

test configuration

IN THIS SECTION

- Syntax | 29191
- Description | 29191
- Options | 29191
- Required Privilege Level | 29191
- Output Fields | 29191
- Sample Output | 29191
- Release Information | 29192

Syntax

```
test configuration filename
syntax-only
```

Description

Verify that the syntax of a configuration file is correct. If the configuration contains any syntax or commit check errors, a message is displayed to indicate the line number and column number in which the error was found. When using the *filename* option, this command only accepts text files.

Options

filename	Name of the configuration file. This file must be a text file and no other type.
syntax-only	(Optional) Check the syntax of a partial configuration file, without checking for commit errors.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

test configuration

```
user@host> test configuration terminal
[Type ^D to end input]
system {
host-name host;
test1;
login;
}
terminal:3:(8) syntax error: test
```

```
  [edit system]
    'test;'
      syntax error
terminal:4:(11) statement must contain additional statements: ;
  [edit system login]
    'login ;'
      statement must contain additional statements
configuration syntax failed
```

Release Information

Command introduced before Junos OS Release 7.4.

`syntax-only` option introduced in Junos OS Release 12.1.

test-count

IN THIS SECTION

- Syntax | **29192**
- Hierarchy Level | **29193**
- Description | **29193**
- Options | **29193**
- Required Privilege Level | **29193**
- Release Information | **29193**

Syntax

```
test-count test_count-number;
```

Hierarchy Level

```
[edit services rpm twamp client control-connection control-client-name]
```

Description

Specify the total number of test session iterations. The range is 0 through 4294967290.

Options

- **Default:** The default value is 0.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

test-finish-wait-duration (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 29194
- Hierarchy Level | 29194
- Description | 29194
- Required Privilege Level | 29194
- Release Information | 29194

Syntax

```
test-finish-wait-duration;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Number of seconds to wait after transmitting the last frame and before concluding that the test as complete. Once the test is complete, the frames received after that are not considered for the result computation. Use this parameter if the latency introduced by the network under test is high in the normal conditions. Default value is 1 second.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

test interface restart-auto-negotiation

IN THIS SECTION

- Syntax | 29195
- Description | 29195
- Options | 29195
- Required Privilege Level | 29195
- Output Fields | 29195
- Sample Output | 29196
- Release Information | 29196

Syntax

```
test interface restart-auto-negotiation interface-name
```

Description

Restarts auto-negotiation on a Fast Ethernet or Gigabit Ethernet interface.

Options

interface-name	Interface name: **fe-***fpc/pic/port* or **ge-***fpc/pic/port*.

Required Privilege Level

view

Output Fields

Use the `show interfaces extensive` command to see the state for auto-negotiation.

Sample Output

test interface restart-auto-negotiation

```
user@host> test interface restart-auto-negotiation fe-1/0/0
```

Release Information

Command introduced in Junos OS Release 7.6.

test-interface (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 29196
- Junos OS Hierarchy Level | 29196
- Junos OS Evolved Hierarchy Level | 29197
- Description | 29197
- Options | 29197
- Required Privilege Level | 29197
- Release Information | 29197

Syntax

```
test-interface interface-name;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the logical interface on which the RFC 2544-based benchmarking test is run.

For Junos OS, if you configure an `inet` family and the test mode to initiate and terminate test frames on the same device, the interface you configure is not effective. Instead, the test is run on the egress logical interface that is determined using route lookup on the specified destination IPv4 address. If you configure an `inet` family and the test mode to reflect the frames back on the sender from the other end, the logical interface is used as the interface to enable the reflection service (reflection is performed on the packets entering the specified interface). If you not configure the logical interface for reflection test mode, a lookup is performed on the source IPv4 address to determine the interface that hosts the address.

For Junos OS Evolved, you can configure this statement for only family `ccc` or `ethernet-switching`. You cannot configure this statement for a Layer 3 reflector (family `inet`).

Options

`interface-name` Name of the logical interface on which the test needs to be run.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

RFC 2544-Based Benchmarking Tests Overview

rfc2544-benchmarking | **15759**

rfc2544 | **15757**

test-interval

IN THIS SECTION

- Syntax | **29198**
- Junos OS Hierarchy Levels | **29198**
- Junos OS Evolved Hierarchy Level | **29198**
- Description | **29199**
- Options | **29199**
- Required Privilege Level | **29199**
- Release Information | **29199**

Syntax

```
test-interval seconds;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the time to wait between tests, in seconds. A test interval of 0 seconds causes the RPM test to stop after one iteration.

Options

seconds—Number of seconds to wait between tests.

- **Range:** [edit services rpm bgp], [edit services rpm probe *owner* test *test-name*], and [edit services monitoring rpm owner *name* test *name*] hierarchy levels: 0 through 86,400
- **Range:** [edit services rpm twamp client control-connection *control-client-name*] hierarchy level: 1 through 255
- **Default:** 1

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.

Support at the [edit services rpm twamp client control-connection *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

Configuring BGP Neighbor Discovery Through RPM

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

Understanding Two-Way Active Measurement Protocol on Routers

test-iterator-duration (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 29200
- Junos OS Hierarchy Level | 29200
- Junos OS Evolved Hierarchy Level | 29200
- Description | 29200
- Options | 29201
- Required Privilege Level | 29201
- Release Information | 29201

Syntax

```
test-iterator-duration seconds;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the duration of each iteration in seconds. If you configure this value, the default value of each iteration depends on the type of test being run. For throughput, back-to-back-frames and frame-loss types of tests, the default value is 20 seconds. For latency tests, the default value is 120 seconds for Junos OS and 20 seconds for Junos OS Evolved.

Options

seconds Number of seconds for which each test iteration is run.

- **Range:** 10 through 120 seconds

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

rfc2544

test-iterator-pass-threshold (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | **29202**
- Hierarchy Level | **29202**
- Description | **29202**
- Required Privilege Level | **29202**
- Release Information | **29202**

Syntax

```
test-iterator-pass-threshold;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specify the pass threshold of each iteration.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

test msdp

IN THIS SECTION

Syntax | **29203**

- Description | 29203
- Options | 29203
- Required Privilege Level | 29203
- Output Fields | 29204
- Sample Output | 29204
- Release Information | 29204

Syntax

```
test msdp (dependent-peers prefix | rpf-peer originator)
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Description

Find Multicast Source Discovery Protocol (MSDP) peers.

Options

dependent-peers *prefix*	Find downstream dependent MSDP peers.
rpf-peer *originator*	Find the MSDP reverse-path-forwarding (RPF) peer for the originator.
instance *instance-name*	(Optional) Find MDSP peers for the specified routing instance.
logical-system (all | *logical-system-name)*	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

test msdp dependent-peers

```
user@host> test msdp dependent-peers 10.0.0.1/24
```

Release Information

Command introduced before Junos OS Release 7.4.

test-name

IN THIS SECTION

- Junos OS Syntax (Except SRX300 and SRX550HM) | **29204**
- Junos OS Syntax (SRX300 and SRX550HM) | **29206**
- Junos OS Evolved Syntax | **29206**
- Junos OS Hierarchy Level | **29207**
- Junos OS Evolved Hierarchy Level | **29207**
- Description | **29207**
- Options | **29207**
- Required Privilege Level | **29207**
- Release Information | **29207**

Junos OS Syntax (Except SRX300 and SRX550HM)

```
test-name test-name {
    destination-ipv4-address address;
```

```
    destination-mac-address destination-mac-address;
    destination-udp-port  port-number;
    direction (egress | ingress);
    disable-signature-check;
    dscp-code-points dscp-code-points;
    family (bridge| inet | ccc | vpls);
    forwarding-class forwarding-class;
    halt-on-prefix-down;
    in-service;
    ip-swap;
    ivlan-cfi ivlan-cfi;
    ivlan-id ivlan-id;
    ivlan-priority ivlan-priority;
    mode (initiate-and-terminate | reflect);
    outer-tag-protocol-id outer-tag-protocol-id;
    ovlan-cfi ovlan-cfi;
    ovlan-id ovlan-id;
    ovlan-priority  ovlan-priority;
    packet-loss-priority (high | low | medium-high);
    receive-failure-threshold receive-failure-threshold;
    reflect-etype reflect-etype;
    reflect-mode (mac-rewrite | mac-swap | no-mac-swap);
    service-type (elan | eline);
    skip-arp-iteration;
    source-ipv4-address address;
    source-mac-address source-mac-address;
    source-udp-port port-number;
                test-finish-wait-duration test-finish-wait-duration;
    test-interface interface-name;
    test-iterator-duration test-iterator-duration;
    test-iterator-pass-threshold test-iterator-pass-threshold;
    test-profile test-profile;
    timestamp-format (microseconds | nanoseconds);
    transmit-failure-threshold transmit-failure-threshold;
    udp-tcp-port-swap;
    vlan-cfi vlan-cfi;
    vlan-id vlan-id;
    vlan-priority vlan-priority;
}
```

Junos OS Syntax (SRX300 and SRX550HM)

```
test-name test-name {
    destination-ipv4-address address;
    destination-udp-port  port-number;
    disable-signature-check;
    family inet;
    mode reflect;
    source-ipv4-address address;
    source-udp-port port-number;
    test-interface interface-name;
}
```

Junos OS Evolved Syntax

```
test-name test-name {
            destination-ipv4-address address;
            destination-mac-address mac-address;
            destination-udp-port port-number;
            direction (egress | ingress);
            disable-signature-check;
            dscp-code-points codepoint;
            family (ccc | ethernet-switching |inet);
            halt-on-prefix-down;
            in-service;
            ip-swap;
            mode (initiate-and-terminate | reflect);
            reflect-etype;
            reflect-mode (mac-swap | no-mac-swap);
            service-type service-type;
            skip-arp-iteration;
            source-ipv4-address address;
            source-mac-address mac-address;
            source-udp-port port-number;
            test-interface interface-name;
            test-iterator-duration seconds;
            test-profile profile-name;
            udp-tcp-port-swap;
}
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests]
```

Description

Define the name of the RFC 2544-based benchmarking test. For each unique test name that you configure, you can specify a test profile, which contains the settings for a test and its type, and also a test interface, which contains the settings for test packets that are sent and received on the selected interface.

Options

test-name Test name. The name can be up to 32 characters in length.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

RFC 2544-Based Benchmarking Tests for ACX Routers Overview

rfc2544-benchmarking | 15759
rfc2544 | 15757

test-only-mode

IN THIS SECTION

- Syntax | 29208
- Hierarchy Level | 29208
- Description | 29208
- Required Privilege Level | 29209
- Release Information | 29209

Syntax

```
test-only-mode (true | false):
```

Hierarchy Level

```
[edit services unified-access-control ]
```

Description

Configure the device in test-only mode to log access decisions from the IC Series device without actually enforcing the decisions. When configured in test-only mode, the SRX Series Firewall enables all UAC traffic to go through so you can test the implementation without impeding traffic.

Use this statement when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

test policy

IN THIS SECTION

- Syntax | 29209
- Description | 29209
- Options | 29210
- Additional Information | 29210
- Required Privilege Level | 29210
- Output Fields | 29210
- Sample Output | 29210
- Release Information | 29211

Syntax

```
test policy policy-name prefix
```

Description

Test a policy configuration to determine which prefixes match routes in the routing table.

> **NOTE**: If you are using the `test policy` command on a logical system, you must first set the CLI to the logical system context. For example, if you want to test a routing policy that is configured on logical system R2, first run the `set cli logical-system R2` command.

Options

policy-name	Name of a policy.
prefix	Destination prefix to match.

Additional Information

All prefixes in the default unicast routing table (inet.0) that match prefixes that are the same as or longer than the specific prefix are processed by the `from` clause in the specified policy. All prefixes accepted by the policy are displayed. The `test policy` command evaluates a policy differently from the BGP import process. When testing a policy that contains an `interface` match condition in the `from` clause, the `test policy` command uses the match condition. In contrast, BGP does not use the `interface` match condition when evaluating the policy against routes learned from internal BGP (IBGP) or external BGP (EGBP) multihop peers.

When testing a policy, you can see the length of time (in microseconds) required to evaluate the policy and the number of times it has been executed by running the `show policy` *`policy-name`* `statistics` command.

Required Privilege Level

view

Output Fields

For information about output fields, see the output field tables for the `show route` command, the `show route detail` command, the `show route extensive` command, or the `show route terse` command.

Sample Output

test policy

```
user@host> test policy test-statics 172.16.0.1/8
inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
```

```
Prefixes passing policy:

172.16.3.0/8          *[BGP/170] 16:22:46, localpref 100, from 10.255.255.41
                    AS Path: 50888 I
                     > to 10.11.4.32 via en0.2, label-switched-path l2
172.16.3.1/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
172.16.3.2/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
172.16.3.3/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
172.16.3.4/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
Policy test-statics: 5 prefixes accepted, 0 prefixes rejected
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Understanding Routing Policy Tests
- show policy | 23199
- show route | 23512
- show route detail | 23573
- show route extensive | 23619
- show route terse | 23824

test-profile (RFC 2544 Benchmarking)

IN THIS SECTION

- Junos OS Syntax | 29212
- Junos OS Evolved Syntax | 29212
- Junos OS Hierarchy Level | 29212

- Junos OS Evolved Hierarchy Level | **29212**
- Description | **29213**
- Options | **29213**
- Required Privilege Level | **29213**
- Release Information | **29213**

Junos OS Syntax

```
test-profile profile-name;
```

Junos OS Evolved Syntax

```
[edit services monitoring rfc2544]
test-profile profile-name {
    bandwidth bandwidth;
    packet-size size;
    step-percent percent;
    test-type (throughput | latency | frame-loss | back-back-frames);
}
[edit services monitoring rfc2544 tests test-name test-name]
test-profile profile-name;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
[edit services rpm rfc2544-benchmarking profiles]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544]
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the name of the test profile to be associated with a particular test name. This parameter is required when the test mode is configured as initiate-and-terminate. This parameter is disregarded when the test mode is configured as reflection. A reflection service does not use the parameters specified in the test profile.

For Junos OS Evolved, when configured at the [edit services monitoring rfc2544] hierarchy level, configures a test profile, which specifies the type of test and the manner in which it must be performed. You then use the `test-profile` statement at the [edit services monitoring rfc2544 tests test-name *test-name*] hierarchy level to associate the test profile name with a particular test name. You can associate the same test profile with multiple test names.

Options

profile-name Name of the test profile. The name can be up 32 characters in length. The name must start with a letter. Allowed characters are [a-zA-Z0-9_]

Required Privilege Level

system—To view this statement in the configuration.

system -control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.

Statement introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

test (Security IDP)

IN THIS SECTION

- Syntax | 29214
- Hierarchy Level | 29214
- Description | 29214
- Options | 29214
- Required Privilege Level | 29214
- Release Information | 29215

Syntax

```
test test-condition;
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type anomaly]
```

Description

Specify protocol anomaly condition to be checked.

Options

test-condition—Name of the anomaly test condition.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

test security utm anti-spam

IN THIS SECTION

- Syntax | 29215
- Description | 29215
- Options | 29215
- Required Privilege Level | 29216
- Output Fields | 29216
- Sample Output | 29216
- Release Information | 29219

Syntax

```
test security utm anti-spam ip-check <test-IP>
<test-IP> test security utm anti-spam.
```

Description

Use this command to check if the IP (Internet Protocol) is a spam source or it is a configuration problem when the device doesn't block the spam. The anti-spam feature requires internet connectivity with the Spam Block List (SBL) server. Domain Name Service (DNS) must be available to access the SBL server.

Options

test-IP	Supports both IPv4 and IPv6.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Table 1805 on page 29216 lists the output fields for the `test security utm anti-spam` command. Output fields are listed in the approximate order in which they appear.

Table 1805: test security utm anti-spam

Field Name	Field Description
`SBL Server`	Spam block list (SBL) contains the list of the disallowed IPs.
`DNS Server`	The firewall performs SBL lookups through the Domain Name Service (DNS) protocol.

Sample Output

test security utm anti-spam

The following examples test IP 1.0.0.1 in SBL server msgsecurity.juniper.net through DNS server 172.29.151.60:

```
user@host> test security utm anti-spam ip-check 1.0.0.1

UTM anti-spam IP-check result:

    SBL Server:
        msgsecurity.juniper.net

    DNS Server:
        172.29.151.60

       Test result for IP check:
        SPAM (Match sbl server blacklist)
```

command-name

The following examples test IP 1.0.0.2 in SBL server msgsecurity.juniper.net through DNS server 172.29.151.60:

```
user@host>test security utm anti-spam ip-check 1.0.0.2

    UTM anti-spam IP-check result:

    SBL Server:
        msgsecurity.juniper.net

    DNS Server:
        172.29.151.60

       Test result for IP check:
           NON SPAM (No match)
```

command-name

The following examples test IP 1.0.0.2 in SBL server msgsecurity.juniper.net through DNS server 172.29.151.60:

```
user@host>test security utm anti-spam ip-check 1.0.0.2

UTM anti-spam IP-check result:

    SBL Server:
        msgsecurity.juniper.net

    DNS Server:
        172.29.151.60

       Test result for IP check:
        NON SPAM (DNS error)
```

command-name

The following examples test IP 1.0.0.3 in SBL server msgsecurity.juniper.net through DNS server 172.29.151.66:

```
user@host>test security utm anti-spam ip-check 1.0.0.3

UTM anti-spam IP-check result:

    SBL Server:
        msgsecurity.juniper.net

    DNS Server:
        172.29.151.66

       Test result for IP check:
           NON SPAM (Timeout error)
```

command-name

The following examples test IP 1.0.0.4 in SBL server msgsecurity.juniper.net through DNS server 172.29.151.66:

```
user@host>test security utm anti-spam ip-check 1.0.0.4

UTM anti-spam IP-check result:

    SBL Server:
        msgsecurity.juniper.net

    DNS Server:
        172.29.151.66

       Test result for IP check:
           NON SPAM (Anti-Spam is not enable)
```

command-name

The following examples test IP 1.0.0.4 in SBL server msgsecurity.juniper.net through DNS server 172.29.151.66:

```
user@host>test security utm anti-spam ip-check 1.0.0.4

UTM anti-spam IP-check result:

    SBL Server:
        msgsecurity.juniper.net

    DNS Server:
        172.29.151.66

       Test result for IP check:
          No license
```

Release Information

Command introduced in Junos OS Release 20.2.

RELATED DOCUMENTATION

anti-spam (Security Content Security Policy) | 2265

test security utm enhanced-web-filtering url-check

IN THIS SECTION

- Syntax (Prior to Junos OS Release 23.3R1) | 29220
- Syntax (Starting in Junos OS Release 23.3R1) | 29220
- Description | 29220
- Options | 29220
- Required Privilege Level | 29220

- Output Fields | 29220
- Sample Output | 29221
- Release Information | 29223

Syntax (Prior to Junos OS Release 23.3R1)

```
test security utm enhanced-web-filtering url-check
test-url Enhanced-web-filtering threat-check test URL
```

Syntax (Starting in Junos OS Release 23.3R1)

```
test security utm web-filtering url-check
```

Description

Use this test command to send the test-string to web-filtering server (example: Websense ThreatSeeker Cloud) to test the category and reputation of the test-string.

Starting in Junos OS Release 23.3R1, you can use the `test security utm web-filtering url-check` test command to check the category and reputation of a URL.

Options

<test-url> Enhanced-web-filtering threat-check test URL. The maximum length of test URL is 249.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Table 1806 on page 29221 lists the output fields for the `test security utm enhanced-web-filtering url-check` command. Output fields are listed in the approximate order in which they appear.

Table 1806: test security utm enhanced-web-filtering

Field Name	Field Description
`Enhance web-filtering server`	Enhanced Web Filtering (EWF) supports HTTP methods.

Sample Output

test security utm enhanced-web-filtering

The following example sends URL www.google.com to EWF rp.cloud.threatseeker.com check the category and reputation:

```
user@host> test security utm enhanced-web-filtering url-check www.google.com

UTM enhanced-web-filtering URL-check result:

    Enhance web-filtering server:
        rp.cloud.threatseeker.com

    Test result for URL check:
            Category name: Enhanced_Search_Engines_and_Portals
            Reputation: 90
            Threat level: very-safe
```

command-name

The following example sends URL www.aaa.com to EWF rp.cloud.threatseeker.com check the category and reputation:

```
user@host>test security utm enhanced-web-filtering url-check www.aaa.com

    UTM enhanced-web-filtering URL-check result:

        Enhance web-filtering server:
         rp.cloud.threatseeker.com

     Test result for URL check:
```

```
            URL does not match on remote server
```

command-name

The following example sends URL www.bbb.com to EWF rp.cloud.threatseeker.com check the category and reputation:

```
user@host>test security utm enhanced-web-filtering url-check www.bbb.com

UTM enhanced-web-filtering URL-check result:

       Enhance web-filtering server:
           rp.cloud.threatseeker.com

    Test result for URL check:
           UTM enhanced web filtering URL check test failed
```

command-name

The following example sends URL www.bbb.com to EWF rp.cloud.threatseeker.com check the category and reputation:

```
user@host>test security utm enhanced-web-filtering url-check www.bbb.com

UTM enhanced-web-filtering URL-check result:

       Enhance web-filtering server:
           rp.cloud.threatseeker.com

    Test result for URL check:
           No license
```

command-name

The following example sends URL www.google.com to EWF rp.cloud.threatseeker.com check the category and reputation:

```
user@host>test security utm enhanced-web-filtering url-check www.google.com

UTM enhanced-web-filtering URL-check result:

    Enhanced web-filtering server:
        rp.cloud.threatseeker.com

    Test result for URL check:
        EWF query timeout
```

test security utm web-filtering url-check

The following example sends URL www.yahoo.com to check the category and reputation:

```
user@host>test security utm web-filtering url-check www.yahoo.com
UTM Web-filtering URL-check result:

    Web-filtering server:
        signatures.juniper.net

    Test result for URL check:
        Type: NG-Juniper
        Category name: NG_Search_Engines_and_Portals
        Reputation: 2
        Reputation level: very safe
```

Release Information

Command introduced in Junos OS Release 20.2.

Command `test security utm web-filtering url-check` introduced in Junos OS Release 23.3R1.

RELATED DOCUMENTATION

Enhanced Web Filtering

test security utm web-filtering profile

IN THIS SECTION

- Syntax | 29224
- Description | 29224
- Options | 29224
- Required Privilege Level | 29224
- Sample Output | 29225
- Release Information | 29227

Syntax

```
test security utm web-filtering profile <profile-name>
<test-url> Web-filtering test URL
```

Description

Use this command to check if the test-string matches the specific profile in the Web filtering server.

Options

test-url Enhanced Web filtering threat-check test URL. The maximum length of test URL is 249.

Required Privilege Level

view

Sample Output

test security utm web-filtering profile

The following example sends profile *sportclass* and test url *www.google.com* to the Web filtering server to check whether the test string matches the specific profile:

```
user@host>  test security utm web-filtering profile sportclass www.google.com

UTM web-filtering profile test:

    Test result:       Match EWF category
    Execute action:    Permit
       Match category:    Enhanced_Search_Engines_and_Portals
```

command-name

The following example sends profile *dangerclass* and test url *www.gun-clubs.com* to the Web filtering server to check whether the test string matches the specific profile:

```
user@host> test security utm web-filtering profile dangerclass www.gun-clubs.com

    UTM web-filtering profile test:

    Test result:       Match custom category
    Execute action:    Permit
       Match category:    custom_category
```

command-name

The following example sends profile *dangerclass* and test url *www.gun-clubs.com* to the Web filtering server to check whether the test string matches the specific profile:

```
user@host> test security utm web-filtering profile dangerclass www.gun-clubs.com

UTM web-filtering profile test:
```

```
    Test result:       Hit global reputation action
    Execute action:    Permit
       Match category:    N/A
```

command-name

The following example sends profile *dangerclass* and test url *www.gun-clubs.com* to the Web filtering server to check whether the test string matches the specific profile:

```
user@host> test security utm web-filtering profile dangerclass www.gun-clubs.com

UTM web-filtering profile test:

        Test result: Web filtering engine is not ready, please try again later
```

command-name

The following example sends profile *dangerclass* and test url *www.gun-clubs.com* to the Web filtering server to check whether the test string matches the specific profile:

```
user@host> test security utm web-filtering profile dangerclass www.gun-clubs.com

UTM web-filtering profile test:

        Test result: Can't find webfilter profile dangerclass
```

command-name

The following example sends profile *dangerclass* and test url *www.gun-clubs.com* to the Web filtering server to check whether the test string matches the specific profile:

```
user@host> test security utm web-filtering profile dangerclass www.gun-clubs.com

UTM web-filtering profile test:
```

```
    Test result:       No license
    Execute action:    Permit
       Match category:    N/A
```

command-name

The following example sends profile *junos-wf-enhanced-default* and test *www.google.com* to the Web filtering server to check whether the test string matches the specific profile:

```
user@host> test security utm web-filtering profile junos-wf-enhanced-default www.google.com

UTM web-filtering profile test:

    Test result:       Fallback (query timeout)
    Execute action:    Log and permit
    Match category:    N/A
```

Release Information

Command introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Enhanced Web Filtering

test services l2tp tunnel

IN THIS SECTION

- Syntax | 29228
- Description | 29228
- Options | 29228

- Required Privilege Level | 29228
- Output Fields | 29228
- Sample Output | 29229
- Release Information | 29230

Syntax

```
test services l2tp tunnel user user-name
<password user-password>
<tunnel-name name>
```

Description

(MX Series routers only) Test and verify Layer 2 Tunneling Protocol (L2TP) tunnel configurations from the L2TP access concentrator (LAC). The test determines whether the user can be authenticated and tunneled according to the L2TP configuration. The establishment of all tunnels associated with the user is tested. You can optionally specify a particular tunnel to test for the user.

Options

user *user-name*	Name of the user under test. You must use an existing configured username, although it can be created solely for testing a tunnel configuration.
password *user-password*	(Optional) Authentication password for the specified user. If you omit this option, the test generates a dummy password—*testpass*—for the user.
tunnel-name *name*	(Optional) Name of a tunnel to test.

Required Privilege Level

view

Output Fields

Table 1807 on page 29229 lists the output fields for the `test services l2tp tunnel` command. Output fields are listed in the approximate order in which they appear.

Table 1807: test services l2tp tunnel Output Fields

Field Name	Field Description
Tunnel-name	Name of the tunnel as configured in the local tunnel profile.
Tunnel-peer	IP address of the tunnel's remote peer (the L2TP network server [LNS]).
Logical-System	Logical system in which the tunnel is created.
Routing-Instance	Routing instance in which the tunnel is created.
Status	Status of the tunnel.

Sample Output

test services l2tp tunnel user (User authentication fails)

```
user@host> test services l2tp tunnel user testuser@example.com
Subscriber: testuser@example.com, authentication failed
```

test services l2tp tunnel user (Multiple tunnels tested)

```
user@host> test services l2tp tunnel user testuser@example.com
Subscriber: testuser@example.com, authentication success, l2tp tunneled
  Tunnel-name   Tunnel-peer   Logical-System  Routing-Instance Status
  test1tunnel    192.0.2.3        default         default         Up
  test2tunnel    198.51.100.243   default         default        Peer unresponsive
  test3tunnel    198.51.100.251   default           test          Up
```

test services l2tp tunnel user tunnel-name (Specific tunnel tested)

```
user@host> test services l2tp tunnel user testuser@example.com tunnel-name test1tunnel
Subscriber: testuser@example.com, authentication success, l2tp tunneled
```

```
Tunnel-name   Tunnel-peer   Logical-System  Routing-Instance Status
test1tunnel    192.0.2.3        default         default         Up
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

test services monitoring rfc2544

IN THIS SECTION

- Syntax | **29230**
- Description | **29231**
- Options | **29231**
- Additional Information | **29231**
- Required Privilege Level | **29231**
- Sample Output | **29231**
- Release Information | **29234**

Syntax

```
test services monitoring rfc2544
(test test-name | test-id id)
<routing-instance <routing-instance-name>>
(clear-counters | start | stop)
```

Description

Start or stop an RFC 2544-based benchmarking test, or clear the statistical counters for a test. You can start or stop a specific test name, or you can stop a test based on its test identifier. When you trigger an RFC 2544-based benchmarking test, it passes through a series of states. These states are displayed in the Test state field of the `show services monitoring rfc2544` command.

> **NOTE**: The RFC 2544 test is stopped at the initiator automatically after the test successfully completes all of the test steps. You need not explicitly enter the `test services monitoring rfc2544-benchmarking (test test-name | test-id test-id) stop` command. However, at the reflector, you must explicitly enter this command to stop the test after the test is completed at the initiator.

Options

(test *test-name* | test-id *id*)	Name of the test to be started or stopped or the unique identifier of the test to be stopped.
routing-instance *routing-instance-name*	(Optional) Name of the routing instance for the test.
(clear-counters | start | stop)	Clear the statistics associated with a test, or start or stop a test.

Additional Information

Required Privilege Level

view

Sample Output

test services monitoring rfc2544 test *test-name* start

```
user@host> test services monitoring rfc2544 test mytest start
```

If successful, the identifier of the started test is displayed:

```
Test mytestid 1 started
```

If unsuccessful, an error message is displayed:

```
Error starting test mytest: error-description
```

Some examples:

- If you try to start a test that is already running:

  ```
  Error starting test mytest: previous session with id 1 still active
  ```

- If you try to start an INET reflector test that has an incorrect address configured on the `destination-ipv4-address` statement:

  ```
  Error starting test mytest: inet address '192.0.2.2' not local for default routing-instance
  ```

- If you try to start a test for a test name that has not been configured on the `test-name` statement:

  ```
  Error starting test mytest: configuration entry for test 'mytest' not found
  ```

test services monitoring rfc2544 test *test-name* routing-instance *instance-name* start

```
user@host> test services monitoring rfc2544 test mytest routing-instance lg01 start
```

If successful, the identifier of the started test is displayed:

```
Test mytest id 1 started
```

If unsuccessful, an error message is displayed:

```
Error starting test mytest: error-description
```

Some examples:

- If you try to start an INET reflector test in a routing instance that has an address configured on the `destination-ipv4-address` statement that is not in the given routing-instance:

```
Error starting test mytest: inet address '192.0.2.2' not local in routing-instance lg01
```

- If you try to start a test for a routing instance that has not been configured on the `routing-instance` statement:

```
Error starting test mytest: routing instance 'lg01' not found
```

test services monitoring rfc2544 test-id *test-id* stop

```
user@host> test services monitoring rfc2544 test-id 1 stop
```

If successful, the identifier of the started test is displayed:

```
Test test-name id 1 stopped
```

If unsuccessful, an error message is displayed:

```
Error stopping test with id 1: error-description
```

For example, if the test identifier you specified when you issued the command doesn't exist:

```
Error stopping test test-name: no active sessions found
```

test services monitoring rfc2544 test *test-name* clear-counters

```
user@host> test services monitoring rfc2544 test mytest clear-counters
```

If successful, a message is displayed describing which counters were cleared.

```
Test mytest id test-id clearing counters
```

If unsuccessful, an error message is displayed:

```
Error clearing counters for test mytest: error-description
```

For example, if you try to clear counters for a test that does not have active sessions:

```
Error clearing counters for test mytest: no active sessions found
```

Release Information

Command introduced for Junos OS Evolved Release 21.1R1.

test services rpm rfc2544-benchmarking test

IN THIS SECTION

- Syntax (ACX Series) | **29234**
- Syntax (MX104 Router) | **29235**
- Syntax (SRX300 and SRX550HM) | **29235**
- Description | **29235**
- Options | **29236**
- Additional Information | **29236**
- Required Privilege Level | **29236**
- Output Fields | **29236**
- Sample Output | **29236**
- Release Information | **29237**

Syntax (ACX Series)

```
test services rpm rfc2544-benchmarking test
<clear-counters>
<routing-instance routing-instance-name>
```

```
<test-name>
<test-id>
<start | stop>
```

Syntax (MX104 Router)

```
test services rpm rfc2544-benchmarking test
<test-name>
<test-id>
<start | stop>
```

Syntax (SRX300 and SRX550HM)

```
test services rpm rfc2544-benchmarking test
<clear-counters>
<test-name>
<test-id>
<start | stop>
```

Description

Start or stop an RFC 2544-based benchmarking test. You can start or stop all of the test names that are defined on a device, or start or stop a specific test name. You can also stop a test based on its test identifier. You can also clear the statistical counters associated with the test. When you trigger an RFC 2544-based benchmarking test, it passes through a series of states. These states are displayed in the Test state field in the brief or displayed output information of the `show services rpm rfc2544-benchmarking` command.

> **NOTE**: The RFC 2544 test is stopped at the initiator automatically after the test successfully completes all of the test steps. You need not explicitly enter the `test services rpm rfc2544-benchmarking test <test-name / test-id> stop` command. However, at the reflector, you must explicitly enter this command to stop the test after the test is completed at the initiator.

When a Layer 2 circuit pseudowire is not up, you cannot start the RFC 2544-based benchmarking test in reflection test mode by entering the `test services rpm rfc2544-benchmarking test test-name start` command. If you attempt to start the reflection test mode, a message is displayed explaining the reason for the failure to commence the test.

Options

start	Start the RFC 2544-based benchmarking test
stop	Terminate the RFC 2544-based benchmarking test
clear-counters	(ACX Series routers, SRX300, and SRX550HM devices only) Clear the statistics associated with the benchmarking test that was run.
routing-instance	(ACX Series routers only) (Optional) Name of the routing instance for the test.
test-name	(Optional) Name of the benchmarking test that must be started or stopped.
test-id	(Optional) Unique identifier of the test that must be stopped. You can stop a test based on the test identifier. You can use the *test-id* option with only the `test services rpm rfc2544-benchmarking stop` command.

Additional Information

The test session is supported in out-of-service mode for the underlying service. You must not transmit any traffic to the UNI port, configured as a generator or a reflector, that is being tested during the duration of the test.

Required Privilege Level

view

Output Fields

To display the results of the benchmarking test, use the `show services rpm rfc2544-benchmarking test` command.

Sample Output

test services rpm rfc2544-benchmarking test start

```
user@host> test services rpm rfc2544-benchmarking test test1 start
Test "test1" id 56 started
```

The response specifies that a test has been started with test id 56. The test ID can be further used in `show` commands to view test output.

Release Information

Command introduced in Junos OS Release 12.3X52.

RELATED DOCUMENTATION

> *Configuring an RFC 2544-Based Benchmarking Test*
>
> *Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers*
>
> rfc2544-benchmarking | **15759**

test-session (Junos OS)

IN THIS SECTION

- Syntax | **29237**
- Hierarchy Level | **29238**
- Description | **29238**
- Options | **29238**
- Required Privilege Level | **29238**
- Release Information | **29238**

Syntax

```
test-session session-name {
    data-fill-with zeros data;
    data-size size;
    destination-port (RPM and TWAMP) port;
    dscp-code-points dscp-bits;
    probe-count count;
    probe-interval seconds;
```

```
    source-address source-address;
    target-address (TWAMP) target-address local-link IPv6-link-local-interface-name;
}
```

Hierarchy Level

```
[edit services rpm twamp client control-connection connection-name]
```

Description

Specify the test session details that includes the session name, the contents of the test packet, the data size, the probe details, and the target destination details.

Options

session-name	Name of the session.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

`destination-port` option added in Junos OS Release 21.1R1.

`source-address` option and the `local-link` sub-option of the `target-address` option for TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

test-session (Junos OS Evolved)

IN THIS SECTION

- Syntax | 29239
- Hierarchy Level | 29240
- Description | 29240
- Options | 29240
- Required Privilege Level | 29244
- Release Information | 29244

Syntax

```
test-session name {
    data-size   data-size;
    destination-port    destination-port;
    dscp-code-points dscp-code-points;
    history-size history-size;
     moving-average-size moving-average-size;
    offload-type (none | inline-timestamp | pfe-timestamp);
    probe-count  probe-count;
    probe-interval  seconds;
    source-address source-address;
    source-routing {
        forward (srmpls-nodes <label hop1...n> | srv6-nodes <IPv6 SID hop1...n>);
        return (srmpls-nodes <label hop1...n> | srv6-nodes <IPv6 SID hop1...n>);
    }
    target target-address local-link IPv6-link-local-interface-name;
    thresholds {
        control-failure (on | off);
        successive-loss number;
        total-loss number;
        threshold-type (microseconds | average);
    }
    traps {
                    egress-jitter-exceeded
                    egress-time-exceeded;
```

```
                    ingress-jitter-exceeded;
                    ingress-time-exceeded;
                    jitter-exceeded;
                    probe-failure;
                    rtt-exceeded;
                    test-completion;
                    test-failure;
    }
    ttl hop-count;
}
```

Hierarchy Level

```
[edit services monitoring twamp client control-connection name]
```

Description

Specify the TWAMP test session details, including the session name, the contents of the test packet, the data size, the probe details, and the target destination. You must configure at least one test session.

Options

name Name of the test session.

data-size *bytes* Specify the size of the data portion of the test packet, in bytes.

- **Range:** 0 through 65400 bytes
- **Default:** 0

destination-port *port* Specify the User Datagram Protocol (UDP) port number to which a probe is sent. You must have the control-type option for the control-connection statement set to `light` to be able to configure a destination port for a particular test session. If you set the control-type option to `managed`, the destination port is negotiated for you between the client and the server, and so you cannot configure the destination port for a particular test session.

- **Range:** You can specify port 862, or any port from 49152 through 65535.
- **Default:** 862 (IANA port for TWAMP)

dscp-code-points ***dscp-bits***

Configure the value of the Differentiated Services (DiffServ) field within the IP header. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern.

- **Values:** Configure a valid 6-bit pattern; for example, `001111`, or one of the following DSCP aliases:
  - `af11`—Default: 001010
  - `af12`—Default: 001100
  - `af13`—Default: 001110
  - `af21`—Default: 010010
  - `af22`—Default: 010100
  - `af23` —Default: 010110
  - `af31` —Default: 011010
  - `af32` —Default: 011100
  - `af33` —Default: 011110
  - `af41` —Default: 100010
  - `af42` —Default:100100
  - `af43` —Default:100110
  - `be`—Default: 000000
  - `cs1`—Default: 001000
  - `cs2`—Default: 010000
  - `cs3`—Default: 011000
  - `cs4`—Default: 100000
  - `cs5`—Default: 101000
  - `cs6`—Default: 110000
  - `cs7`—Default: 111000
  - `ef`—Default: 101110

- `nc1`—Default: 110000
- `nc2`—Default: 111000

history-size *number*

Specify the number of stored history entries. The value configured for the history-size option must be equal to or larger than the value configured for the probe-count option, or the configuration will not commit.

- **Range:** 0 to 255
- **Default:** 50

moving-average-size *number*

Enable statistical calculations to be performed across a configurable number of the most recent samples.

- **Range:** 0 through 512
- **Default:** 0 (disable)

offload-type (none | inline-timestamp | pfe-timestamp)

Enable timestamping of TWAMP probe messages in the Routing Engine, inline in the hardware, or Packet Forwarding Engine host processor. Prior to Junos OS Evolved 22.3R1, for IPv6 traffic, you must configure `none`, as only Routing Engine timestamping is supported. Starting in Junos OS Evolved 22.3R1, you can configure the `pfe-timestamping` option for IPv6 traffic.

- **none**—Timestamping in the Routing Engine.
- **inline-timestamping**—Timestamping inline, in the hardware at the generator, and reflection in the hardware at the reflector. When this option is configured for an ACX router, the ACX can act as either a server or client, but not both. Also for an ACX router, `inline-timestamping` is not supported on PM50 ports if FEC74 is enabled on that port. You can configure the `inline-timestamping` option for up to 256 test sessions.
- **pfe-timestamping**—Timestamping in the Packet Forwarding Engine.
- **Default:** pfe-timestamp

probe-count *number*

Specify the number of probes within a test. The value configured for the probe-count option must be smaller than the value configured for the history-size option, or the configuration will not commit.

- **Range:** 1 through 255
- **Default:** 1

probe-interval *seconds*

Specify the number of seconds to wait before sending the next probe.

- **Range:** 1 through 255 seconds
- **Default:** 1

source-address *address*

Specify the IPv4 or IPv6 source address used for test probes. If the source address is not one of the device's assigned addresses, the probe uses the outgoing interface's address as its source.

You cannot use IPv6 link-local addresses or the following addresses as the source address for probes:

- 0.0.0.0
- 127.0.0.0/8 (loopback)
- 224.0.0.0/4 (multicast)
- 255.255.255.255 (broadcast)

When an IPv6 link-local address is configured with the `target` *`target-address`* `local-link` *`IPv6-link-local-interface-name`* option, you cannot configure the `source-address` *`address`* option. Instead, the operating system sets this option for you.

target *target-address* **local-link** *IPv6-link-local-interface-name*

Specify the IPv4 or IPv6 address for the target destination of the probe. Prior to Junos OS Evolved 22.3R1, you cannot use an IPv6 link-local address for the target destination of a probe. Starting in Junos OS Evolved 22.3R1, for TWAMP Light test sessions with IPv6 target addresses (where the `control-type light` statement is configured for the control connection), configure the link-local logical interface name for the egress interface. You cannot configure IPv6 link-local addresses for TWAMP managed test sessions (where the `control-type managed` statement is configured for the control connection) .

ttl *hop-count*

Specify the maximum number of hops a TWAMP probe can travel. You configure the ttl option when necessary to restrict the scope of a probe, so that a probe does not unintentionally monitor an alternative path to the destination (such as may occur after a BGP re-routing). You can also use the ttl option to monitor direct reachability by specifying a TTL of 1. Probes that exceed the number configured are discarded.

- **Range:** 1 through 255 hops
- **Default:** 64 hops

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.3R1.

`traps` option introduced in Junos OS Evolved Release 21.3R1.

IPv6 address support introduced in Junos OS Evolved Release 21.4R1.

The `local-link` sub-option of the `target` option for TWAMP Light test sessions introduced in Junos OS Evolved Release 22.3R1.

The `inline-timestamping` suboption of the `offload-type` option introduced in Junos OS Evolved 22.4R1.

`source-routing` option introduced in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

test-type (RFC 2544 Benchmarking)

IN THIS SECTION

- Junos OS Syntax | **29245**
- Junos OS Evolved Syntax | **29245**
- Junos OS Hierarchy Level | **29245**
- Junos OS Evolved Hierarchy Level | **29245**
- Description | **29245**
- Options | **29245**
- Required Privilege Level | **29246**
- Release Information | **29246**

Junos OS Syntax

```
test-type (throughput | latency | frame-loss | back-back-frames);
```

Junos OS Evolved Syntax

```
test-type (throughput | latency | frame-loss | back-to-back);
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking profiles test-profile profile-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 test-profile profile-name]
```

Description

RFC 2544 defines four main test types. You can configure and perform a test for a certain service, such as IPv4 or Ethernet, and analyze the results of the test to examine the various SLA parameters of the service. The test packets traverse through the same path as the regular service traffic. These tests are repeated for different frame sizes before calculating the measurement. (For example: 64, 128, 256, 512, 1024, 1280 and 1518 bytes.)

Configure the type of RFC 2544-based benchmarking test to be performed. Because of the ability of these tests to measure throughput, bursty frames, frame loss, and latency, this mechanism is also used to diagnose and examine Ethernet-based networks.

Options

throughput Measure the maximum rate at which none of the offered or transmitted frames are dropped by the device on which the test is performed. You use this test to measure the rate of successful transmission against the theoretical rate of the Ethernet media used for testing.

`latency` Measure the time interval between the arrival of the last bit of the input frame at the input port and the output of the first bit of the frame on the output port. The test is repeated 20 times and the average latency is recorded as the measure for the frame size under test.

`frame-loss` Measure the percentage of frames that must have been forwarded by a network device under steady state (constant) load conditions, but were not forwarded due to lack of resources. The test is repeated for the maximum rate of the Ethernet media (100%) and further lower granular rates by steps, which are configurable with the `step-percent` statement. The test is repeated to find the rate where there are no frames lost for at least 2 iterations of this test. The same test is repeated for various stream sizes supported for the tested traffic rate.

`back-back-frames` **(Junos OS) or** `back-to-back` **(Junos OS Evolved)** Measure the number of frames that are forwarded by the device on which the test is performed when a burst of frames with minimum inter-frame gaps is sent to that device from another source device. The minimum burst length must be at least 2 seconds. If the transmitted frames are equal to the frame count of the burst, then the test is repeated by increasing the number of frames in the burst. The back to back value is then determined as the frame count with the maximum burst size transmitted by the device under test with no loss.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.

Statement introduced in Junos OS Evolved 23.4R1.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

rfc2544

tests

IN THIS SECTION

- Junos OS Syntax (Except SRX300 and SRX550HM) | 29247
- Junos OS Syntax (SRX300 and SRX550HM) | 29248
- Junos OS Evolved Syntax | 29249
- Junos OS Hierarchy Level | 29249
- Junos OS Evolved Hierarchy Level | 29249
- Description | 29250
- Required Privilege Level | 29250
- Release Information | 29250

Junos OS Syntax (Except SRX300 and SRX550HM)

```
tests {
    test-name test-name {
        destination-ipv4-address address;
        destination-mac-address destination-mac-address;
        destination-udp-port  port-number;
        direction (egress | ingress);
        disable-signature-check;
        dscp-code-points dscp-code-points;
        family (bridge| inet | ccc | vpls);
        forwarding-class forwarding-class;
        halt-on-prefix-down;
        in-service;
        ip-swap;
        ivlan-cfi ivlan-cfi;
        ivlan-id ivlan-id;
        ivlan-priority ivlan-priority;
        mode (initiate-and-terminate | reflect);
        outer-tag-protocol-id outer-tag-protocol-id;
        ovlan-cfi ovlan-cfi;
        ovlan-id ovlan-id;
        ovlan-priority  ovlan-priority;
```

```
        packet-loss-priority (high | low | medium-high);
        receive-failure-threshold receive-failure-threshold;
        reflect-etype reflect-etype;
        reflect-mode (mac-rewrite | mac-swap | no-mac-swap);
        service-type (elan | eline);
        skip-arp-iteration;
        source-ipv4-address address;
        source-mac-address source-mac-address;
        source-udp-port port-number;
                test-finish-wait-duration test-finish-wait-duration;
        test-interface interface-name;
        test-iterator-duration test-iterator-duration;
        test-iterator-pass-threshold test-iterator-pass-threshold;
        test-profile test-profile;
        timestamp-format (microseconds | nanoseconds);
        transmit-failure-threshold transmit-failure-threshold;
        udp-tcp-port-swap;
        vlan-cfi vlan-cfi;
        vlan-id vlan-id;
        vlan-priority vlan-priority;
        }
}
```

Junos OS Syntax (SRX300 and SRX550HM)

```
tests {
    test-name test-name {
        destination-ipv4-address address;
        destination-udp-port  port-number;
        disable-signature-check;
        family inet;
        mode reflect;
        source-ipv4-address address;
        source-udp-port port-number;
        test-interface interface-name;

    }
}
```

Junos OS Evolved Syntax

```
tests {
    test-name test-name {
            destination-ipv4-address address;
            destination-mac-address mac-address;
            destination-udp-port port-number;
            direction (egress | ingress);
            disable-signature-check;
            dscp-code-points codepoint;
            family (ccc | ethernet-switching |inet);
            halt-on-prefix-down;
            in-service;
            ip-swap;
            mode (initiate-and-terminate | reflect);
            reflect-etype;
            reflect-mode (mac-swap | no-mac-swap);
            service-type service-type;
            skip-arp-iteration;
            source-ipv4-address address;
            source-mac-address mac-address;
            source-udp-port port-number;
            test-interface interface-name;
            test-iterator-duration seconds;
            test-profile profile-name;
            udp-tcp-port-swap;
     }
}
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544]
```

Description

Specify the attributes of the test iteration, such as the address family (type of service, IPv4 or Ethernet), the logical interface, test duration, and test packet size, that are used for a benchmarking test to be run. The test name combined with the test profile represent a single real-time performance monitoring (RPM) configuration instance.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Evolved Release 21.1R1.

Support for the `destination-mac-address`, `direction`, `disable-signature-check`, `in-service`, `ip-swap`, `reflect-etype`, `reflect-mode`, `service-type`, `source-mac-address`, and `udp-tcp-port-swap` configuration statements added in Junos OS Evolved 22.4R1.

RELATED DOCUMENTATION

Configuring an RFC 2544-Based Benchmarking Test

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers

RFC 2544-Based Benchmarking Tests for ACX Routers Overview

rfc2544-benchmarking | **15759**

rfc2544 | **15757**

tftp

IN THIS SECTION

- Syntax | 29251
- Hierarchy Level | 29251
- Description | 29251
- Required Privilege Level | 29252
- Release Information | 29252

Syntax

```
tftp {
    description text-description;
    interface interface-name {
        broadcast;
        description text-description;
        no-listen;
        server address <logical-system logical-system-name> <routing-instance routing-instance-name>;
    }
    server address <logical-system logical-system-name> <routing-instance routing-instance-name>;
}
```

Hierarchy Level

```
[edit forwarding-options helpers]
```

Description

Enable TFTP request packet forwarding.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding

tftp (Security ALG)

IN THIS SECTION

- Syntax | 29252
- Hierarchy Level | 29253
- Description | 29253
- Options | 29253
- Required Privilege Level | 29253
- Release Information | 29253

Syntax

```
tftp {
    disable;
    traceoptions {
        flag {
            all {
                extensive
            }
```

```
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Configure the Trivial File Transfer Protocol (TFTP) ALG on the device. TFTP is a simple protocol used for files transfer (RFC 1350). By default, the TFTP ALG is disabled for NFX Series and SRX Series Firewalls.

Options

- `disable`—Disable the TFTP ALG. By default, the TFTP ALG is enabled.
- `traceoptions`—Configure TFTP ALG tracing options.
  - `flag`—Trace operation to perform.
    - `all`—Trace all events.
    - *`extensive`*—Display extensive amount of data.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding the TFTP ALG

tftp-server

IN THIS SECTION

- Syntax | 29254
- Hierarchy Level | 29254
- Description | 29254
- Options | 29255
- Required Privilege Level | 29255
- Release Information | 29255

Syntax

```
tftp-server {
    connection-limit connections;
    rate-limit connections per minute;
}
```

Hierarchy Level

```
[edit system services]
```

Description

Enable TFTP file transfers

Options

connection-limit

Limit number of simultaneous connections (connections)

- **Default:** 75
- **Range:** 1 through 250

rate-limit

Limit incoming connection rate (connections per minute)

- **Default:** 150
- **Range:** 1 through 250

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

then (IP Security VPN)

IN THIS SECTION

- Syntax | 29256
- Hierarchy Level | 29256
- Description | 29256
- Options | 29257
- Required Privilege Level | 29257
- Release Information | 29257

Syntax

```
then {
    anti-replay-window-size bits;
    backup-remote-gateway address;
    clear-dont-fragment-bit;
    dynamic {
        ike-policy policy-name;
        ipsec-policy policy-name;
    }
    initiate-dead-peer-detection;
    manual {
        direction (inbound | outbound | bidirectional) {
            authentication {
                algorithm (hmac-sha-256);
                key (ascii-text key | hexadecimal key);
            }
            auxiliary-spi spi-value;
            encryption {
                algorithm algorithm;
                key (ascii-text key | hexadecimal key);
            }
            protocol (bundle | esp);
            spi spi-value;
        }
    }
    no-anti-replay;
    remote-gateway address;
    syslog;
    tunnel-mtu bytes;
}
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name]
```

Description

Define the IPsec term actions.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPsec Rules

then (Hierarchical Policer)

IN THIS SECTION

- Syntax | 29258
- Hierarchy Level | 29258
- Description | 29258
- Options | 29258
- Required Privilege Level | 29258
- Release Information | 29258

Syntax

```
then {
    discard;
}
```

Hierarchy Level

```
[edit firewall hierarchical-policer aggregate],
[edit firewall hierarchical-policer premium]
```

Description

On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, discard packets when a specified bandwidth or burst limits for an aggregate level of a hierarchical policer is reached.

Options

`discard`—Discard packets if condition is met.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Applying Policers

Junos OS Class of Service User Guide for Routing Devices

then (Captive Portal Content Delivery)

IN THIS SECTION

- Syntax | 29259
- Hierarchy Level | 29259
- Description | 29259
- Options | 29260
- Required Privilege Level | 29261
- Release Information | 29261

Syntax

```
then {
    accept;
    insert tag-name tag-name tag-value tag-value;
    redirect url;
    rewrite destination-address address <destination-port port-number>;
    syslog;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule rule-name term
term-name],
[edit services captive-portal-content-delivery rule rule-name term term-name]
```

Description

Define the term actions and any optional action modifiers for the captive portal content delivery rule. Use the statement at the `[edit services...]` hierarchy level for static CPCD. Use the statement at the `[edit dynamic-profiles profile-name services...]` hierarchy level for converged services CPCD.

Options

action Actions to accept, redirect, or rewrite packets and all subsequent packets in flows that match the rules.

- `accept`—Accept the packets and all subsequent packets in flows that match the rules.
- `insert`—Insert a tag with a specified value that is added to the HTTP packet header when the term is matched. You must insert the following action modifiers.
  - `tag` *`tag-name`*—Name of the tag inserted in the packet header, with a maximum size of 127 characters. The tag name is case-sensitive, such that `tag ABCD` and `tag abcd` are processed as different names:
  - `tag-value` *`tag-value`*—Value of the specified tag, with a maximum size of 127 characters. You can specify a customer value or use one of four predefined options that causes the associated value to be inserted:
    - *custom*—Value defined by user, with a maximum size of 127 characters.
    - `hostname`—Hostname of the router.
    - `subscriber-ip`—Subscriber's IPv4 address.
    - `subscriber-ipv6`—Subscriber's IPv6 address.
    - `subscriber-mac-addr`—MAC address of the subscriber.
- `redirect`—Redirect the packet and all subsequent packets in flows that match the rules. You can optionally configure the following action modifier:
  - *`url`*— URL destination for the redirected packet. The URL must begin with `http://` or `https://`.

rewrite Rewrite the packet and all subsequent packets in flows that match the rules. You can optionally configure one or both of the following action modifiers:

- `destination-address` *`address`*—Destination address for the rewritten packet.
- `destination-port` *`port-number`*—(Optional) Destination port for the rewritten packet.

syslog Log information about the packet to a system log file.

action Actions to accept, redirect, or rewrite packets and all subsequent packets in flows that match the rules.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the `[edit dynamic-profiles profile-name services captive-portal-content-delivery rule rule-name term term-name]` hierarchy level added in Junos OS Release 17.2R1.

`insert` option added in Junos OS Release 19.1R1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

- *HTTP Redirect Service Overview*
- Firewall Filter Match Conditions Based on Address Fields
- *Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services*
- *Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services*
- *Configuring Routing Engine-Based, Static HTTP Redirect Services*
- *Configuring Routing Engine-Based, Converged HTTP Redirect Services*

then (Event Policy)

IN THIS SECTION

- Syntax | 29262
- Hierarchy Level | 29263
- Description | 29263
- Required Privilege Level | 29263
- Release Information | 29263

Syntax

```
then {
    change-configuration {
        commands {
            "command";
        }
        commit-options {
            check <synchronize>;
            force;
            log "comment-string";
            synchronize;
        }
        retry count number interval seconds;
        user-name username;
    }
    event-script filename {
        arguments {
            argument-name argument-value;
        }
        destination destination-name {
            retry-count count retry-interval seconds;
            transfer-delay seconds;
        }
        output-filename filename;
        output-format (text | xml);
        user-name username;
    }
    execute-commands {
        commands {
            "command";
        }
        destination destination-name {
            retry-count count retry-interval seconds;
            transfer-delay seconds;
        }
        output-filename filename;
        output-format (text | xml);
        user-name username;
    }
    ignore;
    priority-override {
```

```
        facility facility-type;
        severity severity-level;
    }
    raise-trap;
    upload filename (filename | committed) destination destination-name {
        retry-count count retry-interval seconds;
        transfer-delay seconds;
        user-name username;
    }
}
```

Hierarchy Level

```
[edit event-options policy policy-name]
```

Description

Define event policy actions that are executed when an event occurs that triggers the policy. For each event policy, you can configure multiple actions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

- *Configure an Event Policy to Upload Files*
- *Configure an Event Policy to Execute Operational Mode Commands*
- *Execute Event Scripts in an Event Policy*
- *Configure Event Policies to Ignore an Event*
- *Overview of Using Event Policies to Raise SNMP Traps*

then (Filters)

IN THIS SECTION

- Syntax | 29264
- Hierarchy Level | 29264
- Description | 29264
- Options | 29264
- Required Privilege Level | 29265
- Release Information | 29265

Syntax

```
then {
        action;
        action-modifiers;
}
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name]
```

Description

Configure a firewall filter action.

Options

action Actions to accept, discard, or forward packets that match all conditions specified in a filter term.

Starting in Junos OS Release 18.4R1, two new actions – `port-mirror` and `port-mirror-instance` – are added for all match conditions, which enable selective port mirroring of MPLS traffic to a mirrored destination.

The `port-mirror` action enables port mirroring globally on the device, which applies to all Packet Forwarding Engines (PFEs) and associated interfaces.

The `port-mirror-instance` action enables you to customize each instance with different properties for input sampling and port mirroring output destinations, instead of having to use a single system-wide configuration for port mirroring.

NOTE:

- You can configure only two port mirroring instances per Flexible PIC Concentrator (FPC) by including the instance `port-mirror-instance-name` statement at the `[edit forwarding-options port-mirror]` hierarchy level. You can then associate individual port mirroring instances with an FPC, PIC, or (Forwarding Engine Board (FEB) depending on the device hardware.
- For both `port-mirror` and `port-mirror-instance` actions, the output interface must be enabled with Layer 2 family and not family MPLS (Layer 3) for the selective port mirroring feature to work.

action-modifiers Additional actions to analyze, classify, count, or police packets that match all conditions specified in a filter term.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Firewall Filter Match Conditions and Actions (QFX5100, QFX5110, QFX5120, QFX5200, EX4600, EX4650)

Configuring Firewall Filters

Understanding Firewall Filter Match Conditions

Understanding IP-Based Filtering and Selective Port Mirroring of MPLS Traffic

then (Firewall Filters)

IN THIS SECTION

- Syntax | 29266
- Hierarchy Level | 29266
- Description | 29266
- Options | 29266
- Required Privilege Level | 29267
- Release Information | 29267

Syntax

```
then {
        action;
        action-modifiers;
}
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name]
```

Description

Configure a filter action.

Options

action Action to accept, discard, or forward packets that match all match conditions specified in a filter term.

action-modifiers Additional actions to analyze, classify, count, or police packets that match all conditions specified in a filter term.

The `inline-monitoring-instance` action maps the instance for inline monitoring services to the firewall term. You must attach the inline monitoring filter under `inet` or `inet6` family of the logical unit of the interface for applying inline monitoring on the ingress or egress direction. Starting in Junos OS Release 21.1, the `inline-monitoring-instance` option is also supported by the `any`, `bridge`, `ccc`, `mpls`, and `vpls` families.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`inline-monitoring-instance` option added in Junos OS Release 19.4R1 on MX Series routers with MPCs excluding MPC10E and MPC11E linecards.

Added additional family support to the `inline-monitoring-instance` option in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches

Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches

Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device

Configuring Firewall Filters (CLI Procedure)

Understanding Firewall Filter Match Conditions

Understanding Inline Monitoring Services

then (HTTP Header Enrichment)

IN THIS SECTION

- Syntax | **29268**
- Hierarchy Level | **29268**

- Description | 29268
- Required Privilege Level | 29269
- Release Information | 29269

Syntax

```
then {
    count;
    tag tag-name {
        encrypt {
            hash algorithm;
            prefix hash-prefix;
        }
        ipv4-mask ipv4-mask;
        ipv6-mask ipv6-mask;
        ipv4-or-value ipv4-or-value;
        ipv6-or-value ipv6-or-value;
        tag-attribute tag-attr-name;
        tag-header header;
        tag-separator separator;
    }
}
```

Hierarchy Level

```
[edit services hcm tag-rule rule-name term term-number]
```

Description

Specify the actions to be taken if the criteria for the tag rule are matched. For subscribers under static policy control, the matching conditions for a tag rule are determined by the PCC rule that uses the HCM profile specifying the tag rule. For subscribers under dynamic policy control, a message from the PCRF identifies the configured HCM profile to use for HTTP header enrichment.

> **NOTE**: You must configure this statement and include at least one action to be taken for the tag rule term.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

- *Configuring HTTP Header Enrichment Overview*
- *Configuring Tag Rules*
- *Configuring HCM Profiles and Assigning Tag Rules*
- *Junos Web Aware HTTP Header Enrichment Overview*

then (IDS MS-DPC)

IN THIS SECTION

- Syntax | 29270
- Hierarchy Level | 29271
- Description | 29271
- Options | 29271
- Required Privilege Level | 29271
- Release Information | 29271

Syntax

```
then {
    aggregation (IDS) {
        destination-prefix prefix-number | destination-prefix-ipv6 prefix-value;
        source-prefix prefix-number | source-prefix-ipv6 prefix-value;
    }
    (force-entry | ignore-entry);
    logging {
        syslog;
        threshold rate;
    }
    session-limit {
        by-destination (IDS MS-DPC) {
            hold-time seconds;
            maximum number;
            packets number;
            rate number;
        }
        by-pair (IDS MS-DPC) {
            hold-time seconds;
            maximum number;
            packets number;
            rate number;
        }
        by-source (IDS MS-DPC) {
            hold-time seconds;
            maximum number;
            packets number;
            rate number;
        }
    }
    syn-cookie {
        mss value;
        threshold rate;
    }
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name]
```

Description

Define the IDS term actions when using the MS-DPC.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IDS Rules on an MS-DPC

then (IDS MS-MPC)

IN THIS SECTION

- Syntax | 29272
- Hierarchy Level | 29274
- Description | 29274
- Required Privilege Level | 29274

- Release Information | 29274

Syntax

```
then {
    aggregation (IDS) {
        destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
        source-prefix prefix-value | source-prefix-ipv6 prefix-value;
    }
    allow-ip-options {
        any;
        loose-source-route;
        route-record;
        route-alert;
        security;
        stream-id;
        strict-source-route;
        timestamp;
    }
    allow-ipv6-extension-header {
        any;
        ah;
        dstopts;
        esp;
        fragment;
        hop-by-hop;
        mobility;
        routing;
    }
    icmp-fragment-check;
    icmp-large-packet-check;
    land-attack-check (ip-only | ip-port);
    session-limit {
        by-destination {
            by-protocol {
                icmp {
                    maximum number;
                    packets number;
                    rate number;
```

```
                }
                tcp {
                    maximum number;
                    packets number;
                    rate number;
                }
                udp {
                    maximum number;
                    packets number;
                    rate number;
                }
            }
            maximum number;
            packets number;
            rate number;
        }
        by-source {
            by-protocol {
                icmp {
                    maximum number;
                    packets number;
                    rate number;
                }
                tcp {
                    maximum number;
                    packets number;
                    rate number;
                }
                udp {
                    maximum number;
                    packets number;
                    rate number;
                }
            }
            maximum number;
            packets number;
            rate number;
        }
    }
    tcp-syn-defense;
    tcp-syn-fragment-check;
```

```
    tcp-winnuke-check;
}
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name ]
```

Description

Configure the network attack prevention actions for an IDS rule for a service set on an MS-MPC.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Protection Against Network Attacks on an MS-MPC

Understanding IDS on an MS-MPC

then (LRF rule)

IN THIS SECTION

- Syntax | 29275
- Hierarchy Level | 29275

- Description | 29275
- Required Privilege Level | 29275
- Release Information | 29275

Syntax

```
then {
    report {
        collector collector-name;
        template template-name;
        time-limit time-interval;
        volume-limit volume;
    }
}
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name]
```

Description

Configure the actions to take if the LRF rule is matched.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

then (PCC Rules)

IN THIS SECTION

- Syntax | 29276
- Hierarchy Level | 29276
- Description | 29277
- Required Privilege Level | 29277
- Release Information | 29277

Syntax

```
then {
    pcc-action-profile  profile-name;
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rule-name],
[edit services pcef pcc-rules rule-name]
```

Description

Specify the policy and charging control (PCC) action profile for a PCC rule. The PCC action profile specifies the actions to apply to subscriber traffic that matches any of the `from` statements in the PCC rule. A PCC rule configuration must include the `then` statement and a PCC action profile. The referenced PCC action profile must be configured.

If you are using Junos OS Subscriber Aware, specify the name of the PCC action profile at the `[edit unified-edge pcef pcc-rules rule-name]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the PCC action profile at the `[edit services pcef pcc-rules rule-name]` hierarchy level.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-rules rule-name]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Rules

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

then (PCP)

IN THIS SECTION

- Syntax | 29278
- Hierarchy Level | 29278
- Description | 29278
- Required Privilege Level | 29279
- Release Information | 29279

Syntax

```
then {
    pcp-server server-name;
}
```

Hierarchy Level

```
[edit services pcp rule rule-name term term-name]
```

Description

Specify the PCP server to handle the traffic that matches the PCP rule term.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2R1.

RELATED DOCUMENTATION

Configuring Port Control Protocol

then (Policer Action)

IN THIS SECTION

- Syntax | 29279
- Hierarchy Level | 29280
- Description | 29280
- Options | 29280
- Required Privilege Level | 29280
- Release Information | 29280

Syntax

```
then {
                    policer-action;
}
```

Hierarchy Level

```
[edit firewall policer policer-name]
[edit logical-systems logical-system-name firewall policer policer-name]
```

Description

Configure a policer action.

Options

policer-action—Actions to take are:

- `discard`—Discard traffic that exceeds the rate limits defined by the policer.
- `forwarding-class` *class-name*—Classify traffic that exceeds the rate limits defined by the policer.
- `loss-priority`—Set the loss priority for traffic that exceeds the rate limits defined by the policer.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall -control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches
- Configuring Policers to Control Traffic Rates (CLI Procedure)
- Configuring Firewall Filters (CLI Procedure)
- Understanding the Use of Policers in Firewall Filters
- Basic Single-Rate Two-Color Policers

then (Policers)

IN THIS SECTION

- Syntax | 29281
- Hierarchy Level | 29281
- Description | 29281
- Options | 29281
- Required Privilege Level | 29282
- Release Information | 29282

Syntax

```
then {
    policer-action;
}
```

Hierarchy Level

```
[edit firewall policer policer-name]
```

Description

Configure a policer action.

Options

policer-action—Allowed policer actions are **discard**, **loss-priority high**, and **loss-priority low**. **discard** causes the system to drop traffic that exceeds the rate limits defined by the policer. Use **loss-priority high** to allow the system to forward matching traffic in some cases.

NOTE: If you specify a policer in an egress firewall filter, the only supported action is **discard**.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Two-Color and Three-Color Policers to Control Traffic Rates

Configuring Firewall Filters

Overview of Policers

then (Security Application Firewall)

IN THIS SECTION

- Syntax | 29283
- Hierarchy Level | 29283
- Description | 29283
- Options | 29283
- Required Privilege Level | 29284
- Release Information | 29284

Syntax

```
then {
    (deny [block-message] | permit | reject [block-message]);
}
```

Hierarchy Level

```
[edit security application-firewall rule-set rule-set-name rule rule-name]
```

Description

Specify the action to be performed when traffic matches the associated match criteria.

Note that an application firewall is applied after a session has already been created by the security firewall. When traffic is rejected or denied by an application firewall, therefore, logs contain a session open message, a session reject or deny message, and a session close message.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Options

- `deny`—Block the traffic at the firewall. The device drops the packet. By default, no message is returned to the sender.
  - `block-message`—(Optional) In application firewall rules, provide information to the user regarding blocked traffic. Depending on the content of the `profile` option for this rule set, including the `block-message` option displays a default message or customized message, or redirects the user for denied HTTP or HTTPS traffic. All other traffic is dropped silently.
- `permit`—Permit traffic at the firewall.
- `reject`—Block the traffic at the firewall. For TCP traffic, by default the device drops the packet and returns a TCP reset (RST) message to the source host. For UDP and other protocol traffic, by default the device drops the packet and returns an ICMP “destination unreachable, port unreachable” message to both the client and the server.

- `block-message`—(Optional) In application firewall rules, provide information to the user regarding blocked traffic. Depending on the content of the `profile` option for this rule set, including the `block-message` option displays a default message or customized message, or redirects the user for rejected HTTP or HTTPS traffic. All other traffic is dropped as specified in the default action for the `reject` option.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1. Statement updated in Junos OS Release 12.1X44-D10 with the `reject` option. The `block-message` option added in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Example: Configuring an Application Group for Application Firewall

then (Security Destination NAT)

IN THIS SECTION

- Syntax | 29285
- Hierarchy Level | 29285
- Description | 29285
- Options | 29285
- Required Privilege Level | 29285
- Release Information | 29285

Syntax

```
then {
    destination-nat (off | pool pool-name | rule-session-count-alarm (clear-threshold value |
raise-threshold value));
}
```

Hierarchy Level

```
[edit security nat destination rule-set rule-set-name rule rule-name]
```

Description

Specify the action to be performed when traffic matches the destination NAT rule criteria.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10.

then (Security IDP Policy)

IN THIS SECTION

Syntax | 29286

- Hierarchy Level | 29287
- Description | 29287
- Options | 29287
- Required Privilege Level | 29287
- Release Information | 29287

Syntax

```
then {
    action (Security Rulebase IPS) {
        class-of-service {
            dscp-code-point number;
             forwarding-class forwarding-class;
        }
        (close-client | close-client-and-server | close-server |drop-connection | drop-packet |
ignore-connection | mark-diffserv value | no-action | recommended);
    }
    ip-action (Security IDP Rulebase IPS) {
        (ip-block | ip-close | ip-notify);
        log;
        log-create;
        refresh-timeout;
        target (destination-address | service | source-address | source-zone | source-zone-
address | zone-service);
        timeout seconds;
    }
    notification {
        log-attacks {
            alert;
        }
        packet-log (Security IDP Policy) {
            post-attack number;
            post-attack-timeout seconds;
            pre-attack number;
        }
    }
```

```
    severity (critical | info | major | minor | warning);
}
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-ips rule rule-name]
```

Description

Specify the action to be performed when traffic matches the defined criteria.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

then (Security Policies)

IN THIS SECTION

- threshold-logging-interval | 29290

Syntax

```
then {
                        count {
        alarm {
            per-minute-threshold number;
            per-second-threshold number;
        }
    }
    deny;
    log {
        session-close;
        session-init;
    }
    permit {
        application-services {
            application-firewall {
                rule-set rule-set-name;
            }
            application-traffic-control {
                rule-set rule-set-name;
            }
            gprs-gtp-profile profile-name;
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
                profile-name profile-name;
            }
            uac-policy {
                captive-portal captive-portal;
            }
            utm-policy policy-name;
        }
        destination-address {
            drop-translated;
            drop-untranslated;
        }
        firewall-authentication {
            pass-through {
                access-profile profile-name;
                client-match user-or-group-name;
```

```
                ssl-termination-profile profile-name;
                web-redirect;
                web-redirect-to-https;
            }
            user-firewall {
                access-profile profile-name;
                domain domain-name
                ssl-termination-profile profile-name;
            }
            web-authentication {
                client-match user-or-group-name;
            }
        }
        services-offload;
        tcp-options {
            initial-tcp-mss mss-value;
            reverse-tcp-mss mss-value;
            sequence-check-required;
            syn-check-required;
        }
        tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
        }
    }
    reject;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name]
```

Description

Specify the policy action to be performed when packets match the defined criteria.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for the `services-offload` option added in Junos OS Release 11.4. Support for the `ssl-termination-profile` and `web-redirect-to-https` options added in Junos OS Release 12.1X44-D10. Support for the `user-firewall` option added in Junos OS Release 12.1X45-D10. Support for the `initial-tcp-mss` and `reverse-tcp-mss` options added in Junos OS Release 12.3X48-D20.

threshold-logging-interval

IN THIS SECTION

- Syntax | 29290
- Hierarchy Level | 29290
- Description | 29291
- Options | 29291
- Required Privilege Level | 29291
- Release Information | 29291

Syntax

```
threshold-logging-interval <minutes>
```

Hierarchy Level

```
[edit tenants tenant name security idp sensor-configuration packet-log]
[edit logical system logical system name security idp sensor-configuration packet-log]
```

Description

The minimum time interval in minutes between log messages for maximum sessions or memory reached.

Options

minutes-Interval to generate syslog messages when configured packet-log total memory or max-sessions is reached.

- **Range** 1 to 60
- **Default** 15

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.3

RELATED DOCUMENTATION

> packet-log (Security IDP Sensor Configuration) | **12153**

RELATED DOCUMENTATION

> *Security Policies Overview*
>
> *Understanding Security Policy Rules*
>
> *Understanding Security Policy Elements*

then (Security Source NAT)

IN THIS SECTION

- Syntax | 29292
- Hierarchy Level | 29293
- Description | 29293
- Options | 29293
- Required Privilege Level | 29293
- Release Information | 29293

Syntax

```
then source-nat;
        interface {
            persistent-nat {
                address-mapping;
                inactivity-timeout seconds;
                max-session-number value;
                permit (any-remote-host | target-host | target-host-port);
            }
        }
        off;
        pool <pool-name>;
             persistent-nat {
                 address-mapping;
                inactivity-timeout seconds;
                max-session-number number;
                permit (any-remote-host | target-host | target-host-port);
             }
        rule-session-count-alarm (clear-threshold value | raise-threshold value);
    }
}
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name rule rule-name]
```

Description

Specify the action to be performed when traffic matches the source NAT rule criteria.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10.

then (Security Static NAT)

IN THIS SECTION

- Syntax | 29294
- Hierarchy Level | 29294
- Description | 29294
- Options | 29294
- Required Privilege Level | 29295
- Release Information | 29295

Syntax

```
then {
    static-nat {
        inet {
            routing-instance (default | routing-instance-name);
        }
        nptv6-prefix {
            address-prefix;
                routing-instance routing-instance-name;
        nptv6-prefix-name {
            address-prefix-name;
                routing-instance routing-instance-name;
        [
        prefix {
            address-prefix;
            mapped-port lower-port-range to upper-port-range;
            routing-instance (default | routing-instance-name);
        }
        prefix-name {
            address-prefix-name;
            mapped-port lower-port-range to upper-port-range;
            routing-instance (default | routing-instance-name);
        }
    rule-session-count-alarm (clear-threshold value | raise-threshold value);
}
```

Hierarchy Level

```
[edit security nat static rule-set rule-set-name rule rule-name]
```

Description

Specify the action to be performed when traffic matches the static NAT rule criteria.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.3X48-D25.

then (Services CoS)

IN THIS SECTION

- Syntax | 29295
- Hierarchy Level | 29296
- Description | 29296
- Required Privilege Level | 29296
- Release Information | 29296

Syntax

```
then {
    application-profileprofile-name;
    dscp (alias | bits);
    forwarding-class class-name;
    syslog;
    reflexive; | revert; | reverse {
        application-profile profile-name;
        dscp (alias | bits);
        forwarding-class class-name;
        syslog;
    }
}
```

Hierarchy Level

```
[edit services cos rule rule-name term term-name]
```

Description

Define the CoS term actions.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

Configuring Actions in CoS Rules

then (Services CoS Next Gen Services)

IN THIS SECTION

- Syntax | 29297
- Hierarchy Level | 29297
- Description | 29297
- Options | 29297
- Required Privilege Level | 29298

Release Information | 29298

Syntax

```
then {
    application-profile profile-name;
    dscp (alias | bits);
    forwarding-class class-name;
    reflexive; | revert; | reverse {
        application-profile profile-name;
        dscp (alias | bits);
        forwarding-class class-name;
    }
}
```

Hierarchy Level

```
[edit services cos rule rule-name policy policy-name]
```

Description

Specify the Differentiated Services (DiffServ) code point (DSCP) marking and forwarding-class assignments for packets that are processed by a service set and that match the conditions of the policy in a services CoS rule.

The service set that the CoS rule is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

application-profile *profile-name*	The application profile that sets the CoS actions for FTP and SIP traffic.
dscp (*alias* | *bits*)	Either a code point alias or a DSCP bit value to apply to the packet.

forwarding-class *class-name*

Forwarding class name to apply to the packet. The choices are:

- assured-forwarding
- best-effort
- expedited-forwarding
- network-control

reflexive

Applies the CoS rule policy actions to flows in the reverse direction as well as to flows in the matching direction.

revert

Stores the DSCP and forwarding class of a packet that is received in the match direction of the rule and then applies that DSCP and forwarding class to packets that are received in the reverse direction of the same session.

reverse

Specifies actions to apply to flows in the reverse direction of the matching direction.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Class of Service for Services PICs (Next Gen Services)

then (Services NAT)

IN THIS SECTION

- Syntax | 29299
- Hierarchy Level | 29300
- Description | 29300
- Required Privilege Level | 29300
- Release Information | 29300

Syntax

```
then {
    no-translation;
    port-forwarding-mappings map-name;
    translated {
        address-pooling paired;
        clat-prefix clat-prefix;
        destination-pool nat-pool-name;
        destination-prefix destination-prefix;
        dns-alg-pool dns-alg-pool;
        dns-alg-prefix dns-alg-prefix;
        filtering-type endpoint-independent;
        mapping-type endpoint-independent;
        source-pool nat-pool-name;
        source-prefix source-prefix;
        translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | deterministic-napt44 |
deterministic-napt64 | dnat-44 | dynamic-nat44 | napt-44 | napt-66 | napt-pt | stateful-nat464 |
stateful-nat64 | twice-basic-nat-44 | twice-dynamic-nat-44 | twice-napt-44);
        }
    }
    syslog;
}
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name]
```

Description

Define the NAT term actions.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

then (Services Stateful Firewall)

IN THIS SECTION

- Syntax | 29301
- Hierarchy Level | 29301
- Description | 29301
- Options | 29301
- Required Privilege Level | 29301
- Release Information | 29301

Syntax

```
then {
    (accept <skip-ids>| discard | reject);
    syslog;
}
```

Hierarchy Level

```
[edit services stateful-firewall rule rule-name term term-name]
```

Description

Define the stateful firewall term actions. You can configure the router to accept, discard, or reject the targeted traffic. The other actions are optional.

Options

`accept`—Accept the traffic and send it on to its destination.

`accept skip-ids`—The packet is accepted and sent on to its destination, but IDS rule processing configured on an MS-MPC or MS-MIC is skipped.

`discard`—Do not accept traffic or process it further.

`reject`—Do not accept the traffic and return a rejection message. Rejected traffic can be logged or sampled.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`skip-ids` option added in Junos OS Release 17.1 on MS-MPC and MS-MIC on MX Series.

RELATED DOCUMENTATION

Configuring Actions in Stateful Firewall Rules

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-policy/config-guide-policy.html

Configuring Protection Against Network Attacks on an MS-MPC

then (Stateful Firewall Rule Next Gen Services)

IN THIS SECTION

- Syntax | 29302
- Hierarchy Level | 29302
- Description | 29303
- Options | 29303
- Required Privilege Level | 29303
- Release Information | 29303

Syntax

```
then {
    count;
    deny;
    permit;
    reject;
}
```

Hierarchy Level

```
[edit services policies stateful-firewall-rule rule-name policy policy-name]
```

Description

Specify the actions for a stateful firewall rule policy. The policy actions are applied to flows that meet the policy's matching properties.

Options

count Enables a count, in bytes or kilobytes, of all network traffic the policy allows to pass.

deny Drop the packets.

permit Accept the packets and send them to their destination.

reject Drop the packets. For TCP traffic, send a TCP reset (RST) segment to the source host. For UDP traffic, send an ICMP `destination unreachable, port unreachable` message (type 3, code 3) to the source host.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Stateful Firewalls for Next Gen Services

then (TDF Domain Selection)

IN THIS SECTION

Syntax | 29304

- Hierarchy Level | 29304
- Description | 29304
- Required Privilege Level | 29304
- Release Information | 29305

Syntax

```
then {
    domain tdf-domain-name;
    pcef-profile pcef-profile-name;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
```

Description

Specify the TDF domain or the PCEF profile to be selected when the criteria specified in the domain selection statement match.

> **NOTE**: This statement is required even if you have not specified any match criteria. The TDF domain must have been previously configured at the `[edit unified-edge gateways tdf gateway-name domains]` hierarchy level.
>
> The PCEF profile must have been previously configured at the `[edit unified-edge pcef]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

thermal-health-check

IN THIS SECTION

- Syntax (Junos OS) | **29305**
- Syntax (Junos OS Evolved) | **29306**
- Hierarchy Level | **29306**
- Description | **29306**
- Options | **29307**
- Required Privilege Level | **29309**
- Release Information | **29309**

Syntax (Junos OS)

```
thermal-health-check {
      fet-failure-check {
          action-onfail (auto-shutdown | none);
      }
      action-onfail (auto-shutdown | none);
      shutdown-timer value;
      power-threshold value;
```

```
      power-threshold-percent value;
}
```

Syntax (Junos OS Evolved)

```
thermal-health-check {
      action-onfail (auto-shutdown | none);
      power-threshold-percent value;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable thermal health check, and configure an action that will respond to the detection of a thermal health event such as power leakage on PTX5K, MX10K, PTX10K and QFX10K devices. The thermal health check feature monitors the PSM power output and FRU power consumption every minute. When the PSM power output exceeds the FRU power consumption by a default or configured threshold for three consecutive iterations, it assumes a thermal health event and takes an action based on user configuration.

NOTE: The default threshold for QFX10002 devices is 100 W and for other devices is 600 W.

The default action is set to none. You can use the following command to shut down all PSMs when a thermal health check fails:

```
set chassis thermal-health-check action-onfail auto-shutdown shutdown-timer value (power-
threshold | power-threshold-percent) value
```

NOTE: On a Junos OS device, you can configure only `power-threshold` or `power-threshold-percentage` at a time since they cannot coexist.

You can configure `power-threshold-percent` for systems that are connected to either second or third generation power supplies.

To ensure accurate operation, the thermal health check feature enables the system to shut down the PSMs with a load under 20%. You can expect a margin of error in total FRU input and PSM output power readings compared to actual hardware values.

You must enable the PSM watchdog feature along with thermal-health check to shutdown the system in case a thermal health event causes Junos to go down. Please note that PEM firmware upgrade is required for the thermal health check and PSM watchdog feature.

You can enable the `fet-failure-check` option to monitor a failing power supply due to a Field-effect Transistor (FET) failure and take corrective action. You can choose to shutdown a reporting PSM if a redundant power supply is available, raise an alarm and log the events when a risk of thermal event is determined.

NOTE: The `fet-failure-check` option is supported on MX10K and PTX10K devices that run on Junos OS.

Options

fet-failure-check Enable FET failure detection, and configure an action to be taken upon FET failure.

action-onfail Choose an action to be performed on detection of a thermal health event. The following options are available:

- `auto-shutdown`—The software shuts down all PSMs when a thermal health event is detected. It turn off power to the router to avoid further damage or fire.
- `none`—The software raises a major alarm when a thermal health event is detected.

shutdown-timer *value* Set the timer in seconds to shutdown the PSMs (Range: 10sec to 15mins). This setting is necessary if the automatic shutdown did not occur during a thermal health event due to a software freeze. The default shutdown timer is 900 seconds, and you have the option to reconfigure it. Please note that if the configured timer value is shorter than the default, then you need to deactivate it before rebooting the system. The reboot delay could cause the timer to expire. The shutdown-timer option is not available on Junos OS Evolved. Instead, the timeout value configured for the PSM watchdog feature serves as the shutdown-timer.

power-threshold *value* Set power threshold value in watts. The default value for power-threshold is 600. Avoid using the default value for systems connected to second or third generation power supplies.

power-threshold-percent *value*

Set power threshold value in percent of Total System Power Output (Range: 4 to 10 percent). This value indicates the percentage difference between the parameters: Total System Power Output and Total System FRU Input Power. On Junos OS, you must set a value to configure power-threshold-percentage because the system doesn't provide a default. On Junos OS Evolved, power-threshold-percent value has a default value of 8.

Table 1808: Thermal health check alarms

Alarm	Description	Remedy	Severity
`Thermal Check Failed: Exceeded threshold value`	Appears when thermal health check fails upon exceeding the threshold value. If you have configured auto-shutdown, then the system will shut down.	If the action is set to `none` following a thermal health check failure, the alarm clears once the failure is resolved. If the action is set to `auto-shutdown`, you must power-cycle the chassis to recover.	Major
`Thermal Check did not meet conditions`	Appears when the load of any active PSM is below 20%. System will shut down these PSMs, provided N+2 redundancy criteria is met. If the redundancy criteria is not met, the thermal health check feature will raise this alarm because it cannot shut down these PSMs. On Junos OS Evolved, the PSM state becomes `Spare` under this condition.	Shut down the less loaded PSMs to increase the load of active PSMs.	Minor

Table 1808: Thermal health check alarms *(Continued)*

Alarm	Description	Remedy	Severity
`PEM %d shutdown - Insufficient system load` (Junos OS) `PSM %d shutdown - Insufficient system load` (Junos OS Evolved)	Appears at PSM-level when the load of any active PSM is below 20%. System shuts down these PSMs, provided N+2 redundancy criteria is met.	No remedy required for this alarm. Maintaining a load above 20% for all PSMs is recommended to ensure accurate thermal health check feature functionality and shutdown of unnecessary PSMs.	Minor

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1R1.

`fet-failure-check` option introduced in Junos OS Release 21.2R1.

three-color-policer

IN THIS SECTION

- Syntax | 29310
- Hierarchy Level | 29310
- Description | 29310
- Options | 29310
- Required Privilege Level | 29311
- Release Information | 29311

Syntax

```
three-color-policer policer-name {
    action {
        loss-priority high then discard;
    }
    single-rate {
        (color-aware | color-blind);
        committed-information-rate bps;
        committed-burst-size bytes;
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-information-rate bps;
        committed-burst-size bytes;
        peak-information-rate bps;
        peak-burst-size bytes;
    }
}
```

Hierarchy Level

```
[edit firewall],
[edit logical-systems logical-system-name firewall]
```

Description

Configure a three-color policer.

Options

policer-name—Name of the three-color policer. Use this name when you apply the policer to an interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Two-Color and Three-Color Policers to Control Traffic Rates

Overview of Policers

three-color-policer (Applying)

IN THIS SECTION

- Syntax | 29311
- Hierarchy Level | 29312
- Description | 29312
- Options | 29312
- Required Privilege Level | 29312
- Release Information | 29312

Syntax

```
three-color-policer {
    (single-rate | two-rate) policer-name;
}
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name then]
[edit logical-systems logical-system-name firewall family family-name filter filter-name term term-name then]
```

Description

Apply a tricolor marking policer.

Options

`single-rate`—Named tricolor policer is a single-rate policer.

`two-rate`—Named tricolor policer is a two-rate policer.

`policer-name`—Name of a tricolor policer.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

`single-rate` statement added in Junos OS Release 8.2.

Logical systems support introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring and Applying Tricolor Marking Policers

Firewall Filter Nonterminating Actions

Three-Color Policer Configuration Overview

three-color-policer (Configuring) (EX)

IN THIS SECTION

- Syntax | 29313
- Hierarchy Level | 29314
- Description | 29314
- Options | 29314
- Required Privilege Level | 29314
- Release Information | 29314

Syntax

```
three-color-policer policer-name {
    action {
        loss-priority high then discard ;
    }
    single-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        peak-burst-size bytes;
        peak-information-rate bps;
    }
}
```

Hierarchy Level

```
[edit firewall]
```

Description

Configure a three-color policer.

Options

policer-name—Name of the three-color policer. Reference this name when you apply the policer to an interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Tricolor Marking Policers

three-color-policer (Configuring)

IN THIS SECTION

- Syntax | 29315
- Hierarchy Level | 29315

- Description | 29316
- Options | 29316
- Required Privilege Level | 29316
- Release Information | 29316

Syntax

```
three-color-policer policer-name | uid {
    action {
        loss-priority high then discard;
    }
    filter-specific;
    logical-interface-policer;
    physical-interface-policer;
    shared-bandwidth-policer;
    single-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        peak-burst-size bytes;
        peak-information-rate bps;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall],
[edit firewall],
[edit logical-systems logical-system-name firewall]
```

Description

Configure a three-color policer in static firewall filters or dynamic firewall filters in a dynamic client profile or a dynamic service profile.

Options

policer-name—Name of the three-color policer. Reference this name when you apply the policer to an interface.

uid—When you configure a policer at the `[edit dynamic-profiles]` hierarchy level, you must assign a variable UID as the policer name.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The `action` and `single-rate` statements added in Junos OS Release 8.2.

Logical systems support introduced in Junos OS Release 9.3.

Support at the `[edit dynamic-profiles ... firewall]` hierarchy level introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

- Configuring and Applying Tricolor Marking Policers
- Three-Color Policer Configuration Guidelines
- Basic Single-Rate Three-Color Policers
- Basic Two-Rate Three-Color Policers
- Two-Color and Three-Color Logical Interface Policers
- Two-Color and Three-Color Physical Interface Policers
- Two-Color and Three-Color Policers at Layer 2

threshold

IN THIS SECTION

- Syntax | 29317
- Hierarchy Level | 29317
- Description | 29317
- Default | 29318
- Options | 29318
- Required Privilege Level | 29318
- Release Information | 29318

Syntax

```
threshold threshold-value;
```

Hierarchy Level

[edit chassis "fpc" on page 7516 *slot-number* "error " on page 6441 "fatal" on page 6950]

[edit chassis "fpc" on page 7516 *slot-number* "error " on page 6441 "major" on page 10290

[edit chassis "fpc" on page 7516 *slot-number* "error " on page 6441 "minor" on page 10760]

[edit chassis "error " on page 6441 "fatal" on page 6950]

[edit chassis "error " on page 6441 "major" on page 10290]

[edit chassis "error " on page 6441 "minor" on page 10760]

Description

Configure the threshold value at which to take action. If the severity level of the error is fatal, the action is carried out only once when the total number of errors crosses the threshold value. If the severity level of the error is major, the action is carried out once after the occurrence crosses the threshold. If the severity level is minor, the action is carried out as many times as the value specified by the threshold.

For example, when the severity level is minor, and you have configured the threshold value as 10, the action is carried out after the tenth occurrence.

> **NOTE**: You can set the threshold value to 0 for errors with severity level as minor. This implies that no action is taken for that error. You cannot set the threshold value to 0 for errors with severity level as major or fatal.

Default

By default, the error count for fatal and major actions is 1. The default error count for minor actions is 10.

Options

number Specify the threshold of error counts at which to take action.

- **Range:** 0 through 4,294,967,295

Required Privilege Level

interface	To view this statement in the configuration.
interface-control	To add this statement to the configuration.

Release Information

Statement introduced for PTX Series routers in Junos OS Release 13.3.

Statement introduced for MX240, MX480 MX960, and MX2020 routers in Junos OS Release 14.2.

RELATED DOCUMENTATION

Fabric Resiliency

show chassis fpc errors | **18387**

Managing Errors

threshold (Bridge Domains)

IN THIS SECTION

- Syntax | 29319
- Hierarchy Level | 29319
- Description | 29319
- Options | 29319
- Required Privilege Level | 29320
- Release Information | 29320

Syntax

```
threshold suppress value <reuse value>;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name multicast-snooping-options forwarding-cache],
[edit logical-systems logical-system-name routing-instances routing-instance-name multicast-
snooping-options forwarding-cache],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name multicast-snooping-options forwarding-cache],
[edit routing-instances routing-instance-name multicast-snooping-options forwarding-cache],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name multicast-snooping-
options forwarding-cache]
```

Description

Configure the suppression and reuse thresholds for multicast snooping forwarding cache limits.

Options

suppress *value*—Value to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number must be greater than the reuse value.

- **Range:** 1 through 200,000

reuse ***value***—(Optional) Value to begin creating new multicast forwarding cache entries. If configured, this number must be less than the suppress value.

- **Range:** 1 through 200,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Example: Configuring Multicast Snooping

threshold-count

IN THIS SECTION

- Syntax | **29321**
- Hierarchy Level | **29321**
- Description | **29321**
- Default | **29321**
- Options | **29321**
- Required Privilege Level | **29321**
- Release Information | **29321**

Syntax

```
threshold-count count;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

(MX Series routers only) Configure the threshold count value for MAC move reports.

Default

50

Options

count—Number of MAC moves needed in the notification time to generate a MAC move report.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring MAC Move Parameters

threshold (Detection Time)

IN THIS SECTION

- Syntax | 29322
- Hierarchy Level | 29322
- Description | 29322
- Options | 29322
- Required Privilege Level | 29323
- Release Information | 29323

Syntax

```
threshold milliseconds;
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection detection-time]
```

Description

Specify the threshold for the adaptation of the detection time for a Bidirectional Forwarding Detection (BFD) session. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

> **NOTE**: The threshold time must be greater than or equal to the **minimum-interval** or the **minimum-receive-interval** values.

Options

milliseconds— Value for the detection time adaptation threshold.

- **Range:** 1 through 255,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

threshold (detection-time)

IN THIS SECTION

- Syntax | 29323
- Hierarchy Level | 29323
- Description | 29324
- Options | 29324
- Required Privilege Level | 29324
- Release Information | 29324

Syntax

```
threshold milliseconds;
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method bfd detection-time],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd detection-time],
[edit system services dhcp-local-server group group-name liveness-detection method bfd detection-
time],
```

```
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd
detection-time],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd detection-
time],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd
detection-time]
```

Description

Specify the threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

> **NOTE**: The threshold time must be greater than or equal to the `minimum-interval` or the `minimum-receive-interval`.

Options

milliseconds— Value for the detection time adaptation threshold.

- **Range:** 1 through 255,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients
- Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

threshold (Interfaces)

IN THIS SECTION

- Syntax | 29325
- Hierarchy Level | 29325
- Description | 29325
- Required Privilege Level | 29325
- Release Information | 29325

Syntax

```
threshold <value>;
```

Hierarchy Level

```
[edit interfaces interface-name radio-router]
```

Description

This option controls the percentage of bandwidth change required for routing updates.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

threshold-logging-interval

IN THIS SECTION

- Syntax | 29326
- Hierarchy Level | 29326
- Description | 29326
- Options | 29326
- Required Privilege Level | 29327
- Release Information | 29327

Syntax

```
threshold-logging-interval <minutes>
```

Hierarchy Level

```
[edit tenants tenant name security idp sensor-configuration packet-log]
[edit logical system logical system name security idp sensor-configuration packet-log]
```

Description

The minimum time interval in minutes between log messages for maximum sessions or memory reached.

Options

minutes-Interval to generate syslog messages when configured packet-log total memory or max-sessions is reached.

- **Range** 1 to 60
- **Default** 15

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.3

RELATED DOCUMENTATION

packet-log (Security IDP Sensor Configuration) | 12153

threshold (MSDP Active Source Messages)

IN THIS SECTION

- Syntax | 29327
- Hierarchy Level | 29328
- Description | 29328
- Options | 29328
- Required Privilege Level | 29328
- Release Information | 29328

Syntax

```
threshold number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp active-source-limit],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
active-source-limit],
[edit protocols msdp active-source-limit],
[edit routing-instances routing-instance-name protocols msdp active-source-limit]
```

Description

Configure the random early detection (RED) threshold for MSDP active source messages. This number must be less than the configured or default maximum.

Options

number—RED threshold for active source messages.

- **Range:** 1 through 1,000,000
- **Default:** 24,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring MSDP with Active Source Limits and Mesh Groups

maximum (MSDP Active Source Messages) | **10536**

threshold (Multicast Forwarding Cache)

IN THIS SECTION

- Syntax | 29329
- Hierarchy Level | 29329
- Description | 29330
- Options | 29330
- Required Privilege Level | 29330
- Release Information | 29331

Syntax

```
threshold {
    log-warning value;
    suppress value;
    reuse value;
    mvpn-rpt-suppress value;
    mvpn-rpt-reuse value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast forwarding-cache],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast forwarding-cache family (inet | inet6)],
[edit logical-systems logical-system-name routing-options multicast forwarding-cache],
[edit logical-systems logical-system-name routing-options multicast forwarding-cache family (inet
| inet6)],
[edit routing-instances routing-instance-name routing-options multicast forwarding-cache],
[edit routing-instances routing-instance-name routing-options multicast forwarding-cache (inet |
inet6)],
```

```
[edit routing-options multicast forwarding-cache],
[edit routing-options multicast forwarding-cache family (inet | inet6)]
```

Description

Configure the suppression, reuse, and warning log message thresholds for multicast forwarding cache limits. You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the `threshold` statement globally for the multicast forwarding cache or including the `family` statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.

When general forwarding-cache suppression is active, the multicast forwarding-cache prevents forwarding traffic on the shared RP tree (RPT). At the same time, MVPN (*,G) forwarding states are not created for new RPT c-mcast entires, and , (*,G) installed by BGP-MVPN protocol are deleted. When general forwarding-cache suppression ends, BGP-MVPN (*,G) entries are re-added in the RIB and restored to the FIB (up to the MVPN (*,G) limit).

When MVPN RPT suppression is active, for all PE routers in excess of the threshold (including RP PEs), MVPN will not add new (*,G) forwarding entries to the forwarding-cache. Changes are visible once the entries in the current forwarding-cache have timed out or are deleted.

To use `mvpn-rpt-suppress` and/or `mvpn-rpt-reuse`, you must first configure the general `suppress` threshold. If `suppress` is configured but `mvpn-rpt-suppress` is not, both `mvpn-rpt-suppress` and `mvpn-rpt-reuse` will inherit *and use* the value set for the general `suppress`.

Options

reuse or **mvpn-rpt-reuse** *value* (Optional) Value at which to begin creating new multicast forwarding cache entries. If configured, this number should be less than the **suppress** value.

- **Range:** 1 through 200,000

suppress or **mvpn-rpt-suppress** *value* —Value at which to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number should be greater than the **reuse** value.

- **Range:** 1 through 200,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Examples: Configuring the Multicast Forwarding Cache

show multicast forwarding-cache statistics | 22283

threshold (PIM BFD Detection Time)

IN THIS SECTION

- Syntax | 29331
- Hierarchy Level | 29331
- Description | 29332
- Options | 29332
- Required Privilege Level | 29332
- Release Information | 29332

Syntax

```
threshold milliseconds;
```

Hierarchy Level

```
[edit protocols pim interface interface-name bfd-liveness-detection  detection-time],
[edit routing-instances routing-instance-name protocols pim interface interface-name bfd-
liveness-detection detection-time]
```

Description

Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

> **NOTE**: The threshold value must be equal to or greater than the transmit interval.
> The threshold time must be equal to or greater than the value specified in the **minimum-interval** or the `minimum-receive-interval` statement.

Options

milliseconds—Value for the detection time adaptation threshold.

- **Range:** 1 through 255,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Support for BFD authentication introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD for PIM

bfd-liveness-detection (Protocols PIM) | **2945**

detection-time (BFD for PIM) | **5543**

minimum-interval (PIM BFD Liveness Detection) | **10738**

minimum-receive-interval (Protocols PIM) | **10751**

threshold (PIM BFD Transmit Interval)

IN THIS SECTION

- Syntax | 29333
- Hierarchy Level | 29333
- Description | 29333
- Options | 29333
- Required Privilege Level | 29334
- Release Information | 29334

Syntax

```
threshold milliseconds;
```

Hierarchy Level

```
[edit protocols pim interface interface-name bfd-liveness-detection transmit-interval],
[edit routing-instances routing-instance-name protocols pim interface interface-name bfd-
liveness-detection transmit-interval]
```

Description

Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.

Options

milliseconds—Value for the transmit interval adaptation threshold.

- **Range:** 0 through 4,294,967,295 ($2^{32} - 1$)

> **NOTE**: The threshold value specified in the `threshold` statement must be greater than the value specified in the `minimum-interval` statement for the `transmit-interval` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Configuring BFD for PIM

bfd-liveness-detection (Protocols PIM) | 2945

threshold (PIM Entries)

IN THIS SECTION

- Syntax | 29335
- Hierarchy Level | 29335
- Description | 29336
- Options | 29336
- Required Privilege Level | 29336
- Release Information | 29336

Syntax

```
threshold value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim sglimit],
[edit logical-systems logical-system-name protocols pim sglimit family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
sglimit],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
sglimit family],
[edit protocols pim sglimit],
[edit protocols pim sglimit family],
[edit routing-instances routing-instance-name protocols pim sglimit],
[edit routing-instances routing-instance-name protocols pim sglimit family],
[edit logical-systems logical-system-name protocols pim rp group-rp-mapping],
[edit logical-systems logical-system-name protocols pim rp group-rp-mapping family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp group-rp-mapping],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp group-rp-mapping family],
[edit protocols pim rp group-rp-mapping],
[edit protocols pim rp group-rp-mapping family],
[edit routing-instances routing-instance-name protocols pim rp group-rp-mapping],
[edit routing-instances routing-instance-name protocols pim rp group-rp-mapping family],
[edit logical-systems logical-system-name protocols pim rp register-limit],
[edit logical-systems logical-system-name protocols pim rp register-limit family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp register-limit],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp register-limit family],
[edit protocols pim rp register-limit],
[edit protocols pim rp register-limit family],
[edit routing-instances routing-instance-name protocols pim rp register-limit],
[edit routing-instances routing-instance-name protocols pim rp register-limit family]
```

Description

Configure a threshold at which a warning message is logged when a certain number of PIM entries have been received by the device.

Options

value—Threshold at which a warning message is logged. This is a percentage of the maximum number of entries accepted by the device as defined with the `maximum` statement. You can apply this threshold to incoming PIM join messages, PIM register messages, and group-to-RP mappings.

For example, if you configure a maximum number of 1,000 incoming group-to-RP mappings, and you configure a threshold value of 90 percent, warning messages are logged in the system log when the device receives 900 group-to-RP mappings. The same formula applies to incoming PIM join messages and PIM register messages if configured with both the **maximum** ***limit*** and the `threshold` *`value`* statements.

- **Default:** 1 through 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

clear pim join | 3806

threshold-rate

IN THIS SECTION

Syntax | 29337

- Hierarchy Level | 29337
- Description | 29337
- Options | 29338
- Required Privilege Level | 29338
- Release Information | 29338

Syntax

```
threshold-rate kbps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group address source source-address],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group group-address wildcard-source],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet wildcard-source],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet6 wildcard-source],
[edit routing-instances routing-instance-name provider-tunnel selective group address source
source-address]
[edit routing-instances routing-instance-name provider-tunnel selective group address wildcard-
source]
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet
wildcard-source],
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet6
wildcard-source]
```

Description

Specify the data threshold required before a new tunnel is created for a dynamic selective point-to-multipoint LSP. This statement is part of the configuration for point-to-multipoint LSPs for MBGP MVPNs and PIM-SSM GRE or RSVP-TE selective provider tunnels.

Options

number—Specify the data threshold required before a new tunnel is created.

- **Range:** 0 through 1,000,000 kilobits per second. Specifying 0 is equivalent to not including the statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

Configuring PIM-SSM GRE Selective Provider Tunnels

threshold (Routing Instances)

IN THIS SECTION

- Syntax | 29339
- Hierarchy Level | 29339
- Description | 29339
- Required Privilege Level | 29339
- Release Information | 29339

Syntax

```
threshold {
    group group-address {
        source source-address {
            rate threshold-rate;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
mdt],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel family inet | inet6mdt],
[edit routing-instances routing-instance-name protocols pim mdt],
[edit routing-instances routing-instance-name provider-tunnel family inet | inet6mdt]
```

Description

Establish a threshold to trigger the automatic creation of a data MDT.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4. In Junos OS Release 17.3R1, the `mdt` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default MDT in Rosen 7, and data MDT for Rosen 6 and Rosen 7. The `provider-tunnel mdt` hierarchy is now hidden for backward compatibility with existing scripts.

RELATED DOCUMENTATION

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

threshold (Services IPsec)

IN THIS SECTION

- Syntax | 29340
- Hierarchy Level | 29340
- Description | 29340
- Options | 29341
- Required Privilege Level | 29341
- Release Information | 29341

Syntax

```
threshold number;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then dead-peer-detection]
```

Description

Specify the maximum number of unsuccessful dead peer detection (DPD) requests to be sent before the peer is considered unavailable. The `threshold` value is used for IKEv1 security associations (SAs). Starting in Junos OS Release 17.2R1, the `threshold` value is also applicable to IKEv2 SAs. In Junos OS Release 17.1 and earlier, the `threshold` option is not applicable to IKEv2 SAs, which use the default value.

Options

number Maximum number of unsuccessful DPD requests to be sent.

- **Range:** 1 through 10
- **Default:** 3

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

IKEv2 support introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Configuring IPsec Rules

threshold (Services Logging and SYN-Cookie Defenses)

IN THIS SECTION

- Syntax | 29342
- Hierarchy Level | 29342
- Description | 29342
- Options | 29342
- Required Privilege Level | 29342
- Release Information | 29342

Syntax

```
threshold rate;
```

Hierarchy Level

```
[edit services ids rule rule-name term term-name then logging],
[edit services ids rule rule-name term term-name then syn-cookie]
```

Description

Specify the threshold for logging or applying SYN-cookie defenses when using the MS-DPC.

Options

rate—Logging threshold number of events per second.

rate—SYN-cookie defense number of SYN attacks per second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Actions in IDS Rules

threshold-time

IN THIS SECTION

- Syntax | 29343
- Hierarchy Level | 29343
- Description | 29343
- Default | 29343
- Options | 29343
- Required Privilege Level | 29344
- Release Information | 29344

Syntax

```
threshold-time seconds;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

(MX Series routers only) Configure the threshold time value for MAC move reports when the MAC address moves at least a specified number of times (threshold count) in the configured interval.

Default

1 second

Options

seconds—Timer threshold before MAC move reports are generated.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring MAC Move Parameters

threshold (transmit-interval)

IN THIS SECTION

- Syntax | 29344
- Hierarchy Level | 29345
- Description | 29345
- Options | 29345
- Required Privilege Level | 29345
- Release Information | 29345

Syntax

```
threshold milliseconds;
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server group group-name liveness-detection method bfd transmit-
interval],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd
transmit-interval],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd transmit-
interval],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd
transmit-interval]
```

Description

Specify the threshold for detecting the adaptation of the transmit interval. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.

Options

milliseconds — Threshold value.

- **Range:** 0 through 4,294,967,295 ($2^{32} - 1$)

> **NOTE**: The threshold value specified in the `threshold` statement must be greater than the value specified in the `minimum-interval` statement for the `transmit-interval` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients
- Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

thresholds

IN THIS SECTION

- Syntax | 29346
- Hierarchy Level | 29346
- Description | 29347
- Options | 29347
- Required Privilege Level | 29348
- Release Information | 29348

Syntax

```
thresholds {
    clear clear-value;
    interval interval-value;
    set set-value;
    warning-clear warning-clear-value;
    warning-set warning-set-value;
}
```

Hierarchy Level

```
[edit interfaces interfaces-name link-degrade-monitor]
```

Description

Configure the BER threshold values (such as set and clear thresholds) at which different corrective actions must be triggered on a degraded interface.

Options

clear *clear-value*
The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted. You can configure this value in the 1E–*n* format, where 1 is the mantissa (remains constant) and *n* is the exponent. For example, a threshold value of IE–3 refers to the BER threshold value of $1x10^{-3}$. The supported exponent range is 1 through 16, and the default value is 12.

interval *interval-value*
The number of consecutive link degrade events that are considered before any corrective action is taken. The supported value range for the interval is 1 through 256, and the default interval is 10.

set *set-value*
The BER threshold value at which the link is considered degraded and a corrective action, specified by the user, is triggered. You can configure this value in the 1E–*n* format, where 1 is the mantissa (remains constant) and *n* is the exponent. For example, a threshold value of IE–3 refers to the BER threshold value of $1x10^{-3}$. The supported exponent range is 1 through 16, and the default value is 7.

warning clear *warning-clear-value*
The link clear warning threshold. Every time this threshold value is reached, a system message is logged to indicate that the link degrade condition has been cleared on the interface. You can configure this value in the 1E–*n* format, where 1 is the mantissa (remains constant) and *n* is the exponent. For example, a threshold value of IE–3 refers to the BER threshold value of $1x10^{-3}$. The supported exponent range is 1 through 16, and the default value is 11.

warning set *warning-set-value*
The link degrade warning threshold. Every time this threshold value is reached, a system message is logged to indicate that a link degrade has occurred on the interface. You can configure this value in the 1E–*n* format, where 1 is the mantissa (remains constant) and *n* is the exponent. For example, a threshold value of IE–3 refers to the BER threshold value of $1x10^{-3}$. The supported exponent range is 1 through 16, and the default value is 9.

NOTE: The lower the BER with high confidence level, the longer it takes to estimate it. In such cases, a few packet drops might be noticed (based on the bit error distribution) before a link degrade event is detected.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Link Degrade Monitoring Overview

link-degrade-monitor | 9784

recovery | 14146

request interface link-degrade-recover | 14642

thresholds (Junos OS)

IN THIS SECTION

- Syntax | 29348
- Hierarchy Level | 29349
- Description | 29349
- Options | 29349
- Required Privilege Level | 29350
- Release Information | 29350

Syntax

```
thresholds thresholds;
```

Hierarchy Level

```
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name]
```

Description

Specify thresholds used for the probes. A system log message is generated when the configured threshold is exceeded. Likewise, an SNMP trap (if configured) is generated when a threshold is exceeded.

> **NOTE**: If you configure a value of zero using the *thresholds* option for a certain probe parameter, the generation of SNMP traps is disabled for the corresponding probe attribute. For example, if you specify the `set thresholds jitter-egress 0` statement, it denotes that traps are not triggered when the jitter in egress time threshold is met or exceeded.

Options

thresholds—Specify one or more threshold measurements. The following options are supported:

- `egress-time`—Measures maximum source-to-destination time per probe.
- `ingress-time`—Measures maximum destination-to-source time per probe.
- `jitter-egress`—Measures maximum source-to-destination jitter per test.
- `jitter-ingress`—Measures maximum destination-to- source jitter per test.
- `jitter-rtt`—Measures maximum jitter per test, from 0 through 60,000,000 microseconds.
- `rtt`—Measures maximum round-trip time per probe, in microseconds.
- `std-dev-egress`—Measures maximum source-to-destination standard deviation per test.
- `std-dev-ingress`—Measures maximum destination-to-source standard deviation per test.
- `std-dev-rtt`—Measures maximum standard deviation per test, in microseconds.
- `successive-loss`—Measures successive probe loss count, indicating probe failure.
- `total-loss`—Measures total probe loss count indicating test failure, from 0 through 15. The default for this threshold is 1.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [`edit services rpm twamp client control-connection` *`control-client-name`*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

RELATED DOCUMENTATION

> *Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches*
>
> *Understanding Two-Way Active Measurement Protocol on Routers*

thresholds (Junos OS Evolved)

IN THIS SECTION

- Syntax | 29350
- Hierarchy Level | 29351
- Description | 29351
- Options | 29351
- Required Privilege Level | 29352
- Release Information | 29352

Syntax

```
thresholds {
    control-failure (on | off);
    successive-loss number;
```

```
    total-loss number;
    threshold-type (microseconds | average);
}
```

Hierarchy Level

```
[edit services monitoring rpm owner name test name]
[edit services monitoring twamp client control-connection name test-session name]
```

Description

Specify the thresholds used for real-time performance monitoring (RPM) and Two-Way Active Measurement Protocol (TWAMP) probes. A system log message is generated when the configured threshold is exceeded.

Options

control-failure (on | off)

Specify whether control failure should be reported as a test failure.

- **Default:** on

successive-loss *number*

Measures the number of packets lost in succession. Set this number to specify how many packets should be lost in succession before indicating a probe failure.

- **Range:** 0 through 15; setting this option to zero disables this counter.
- **Default:** 1, which indicates that the test should fail with any packet loss.

total-loss *number*

Measures total number of packets lost per test. Set this number to specify how many packets should be lost in total before indicating a probe failure.

- **Range:** 0 through 15; setting this option to zero disables this counter.
- **Default:** 1, which indicates that the test should fail with any packet loss.

***threshold-type* (*microseconds* | average)**

Specify one or more measurement options and set thresholds for them. By default, these options set the maximum threshold. To configure the option to use the average measurement and not the maximum, configure the average option on each configured threshold type. To set the threshold type back to maximum, either delete the average option from the configuration or configure the threshold type again without the average option.

- **Range:** 0 through 60,000,000 microseconds; setting this option to zero disables the configured threshold type.
- **Values:** You can configure these threshold types:
  - egress-time—Source-to-destination time per probe.
  - ingress-time—Destination-to-source time per probe.
  - jitter-egress—Source-to-destination jitter per test.
  - jitter-ingress—Destination-to- source jitter per test.
  - jitter-rtt—Round-trip jitter per test.
  - rtt—Round-trip time per probe.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 20.1R1 for RPM.

Statement introduced in Junos OS Evolved 20.3R1 for TWAMP.

RELATED DOCUMENTATION

Understanding Two-Way Active Measurement Protocol on Routers

time

IN THIS SECTION

Syntax | 29353
Hierarchy Level | 29353

- Description | 29353
- Required Privilege Level | 29353
- Release Information | 29354

Syntax

```
time {
    equals {
        value;
    }
    greater-than value;
    less-than value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format]
```

Description

Specify the custom AVP attribute's format as time and the value to match for the incoming RADIUS request from the IP-based subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

time-binding

IN THIS SECTION

- Syntax | 29354
- Hierarchy Level | 29355
- Description | 29355
- Options | 29355
- Required Privilege Level | 29355
- Release Information | 29356

Syntax

```
time-binding {
    count count-value;
    scope (destination | peer | source);
    interval time-interval;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name]
```

Description

Allow IDP to detect a sequence of the same attacks over a period of time.

Options

count *count-value*

Specify the number of times that IDP detects the attack within the specified scope before triggering an event.

interval *time-interval*

Specify the maximum time interval between any two instances of a time-binding custom attack.

- **Syntax:** 00m-00s
- **Default:** 60 seconds
- **Range:** 0 minutes and 0 seconds to 60 minutes and 0 seconds

scope (destination | peer | source)

Specify whether the counting of the attack is from the same source IP address, the same destination IP address, or a peer.

- Values:

  destination IDP detects attacks to a given destination IP address for the specified number of times, regardless of the source IP address.

  peer IDP detects attacks between source and destination IP addresses of the sessions for the specified number of times.

  source IDP detects attacks from a given source IP address for the specified number of times, regardless of the destination IP address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

`Interval` option introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Time Bindings

time-format

IN THIS SECTION

- Syntax | 29356
- Hierarchy Level | 29356
- Description | 29357
- Options | 29357
- Required Privilege Level | 29357
- Release Information | 29357

Syntax

```
time-format (year | millisecond | year millisecond);
```

Hierarchy Level

```
[edit system syslog]
[edit security log]
```

Description

Include the year, the millisecond, or both, in the timestamp on every standard-format system log message. The additional information is included for messages directed to each destination configured by a `file`, `console`, or `user` statement at the `[edit system syslog]` hierarchy level, but not to destinations configured by a `host` statement.

However, in a Junos OS specific implementation using the FreeBSD console for a syslog destination, additional time information is available in system log messages for that destination when you configure the destination through a `host` statement.

By default, the timestamp specifies the month, date, hour, minute, and second when the message was logged—for example, `Aug 21 12:36:30`.

The following example illustrates the format for a timestamp that includes both the millisecond (`401`) and the year (`2006`):

```
Aug 21 12:36:30.401 2006
```

However, the timestamp for traceoption messages is specified in milliseconds by default, and is independent of the `[edit system syslog time-format]` statement.

> **NOTE**: When the `structured-data` statement is included at the `[edit system syslog file filename]` hierarchy level, this statement is ignored for the file.

Options

`millisecond`—Include the millisecond in the timestamp.

`year`—Include the year in the timestamp.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Include the Year or Millisecond in Timestamps

structured-data | 28664

time-interval (Dampen Options)

IN THIS SECTION

- Syntax | 29358
- Hierarchy Level | 29358
- Description | 29358
- Options | 29359
- Required Privilege Level | 29359
- Release Information | 29359

Syntax

```
time-interval microseconds;
```

Hierarchy Level

```
[edit event-options event-script dampen dampen-options],
[edit event-options event-script file filename dampen dampen-options],
[edit system scripts commit dampen dampen-options],
[edit system scripts op dampen dampen-options],
[edit system scripts op file filename dampen dampen-options]
```

Description

Specify the length of time during which to pause the execution of one or more scripts.

Script dampening is initiated when CPU use exceeds the value specified using the `cpu-factor` statement. Junos OS then executes the number of lines of script specified using the `line-interval` statement and then pauses script execution for the time specified using the `time-interval` statement. If you do not configure values for these statements, the defaults are used.

Options

`microseconds` Specify the time in microseconds to pause script execution.

- **Range:** 1 through 4,294,967,295 microseconds
- **Default:** 100 microseconds

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Dampen Script Execution

time-limit (LRF Rule)

IN THIS SECTION

- Syntax | 29360
- Hierarchy Level | 29360
- Description | 29360
- Options | 29360
- Required Privilege Level | 29360
- Release Information | 29360

Syntax

```
time-limit time-interval;
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name then report]
```

Description

Configure the time limit to be used for reporting. The template that the LRF rule is using must have `trigger-type time` configured.

Options

time-interval

The time limit in seconds.

- **Range:** 60 through 1800
- **Default:** 300

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

time-wait-state

IN THIS SECTION

- Syntax | 29361
- Hierarchy Level | 29361
- Description | 29361
- Options | 29361
- Required Privilege Level | 29362
- Release Information | 29362

Syntax

```
time-wait-state {
    (session-ageout | session-timeout seconds);
}
```

Hierarchy Level

```
[edit security flow tcp-session]
```

Description

Defines the length of time (in seconds) that the device keeps the defined TCP session in the session table. The default is 150 seconds.

Options

- `session-ageout`—Set a TCP session to age out, using the service based timeout value.
- `session-timeout` *seconds*—Set the session timeout value allowed before the device ages out a session from its session table.

  Range: 2 through 600 seconds

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

time-zone

IN THIS SECTION

- Syntax | 29362
- Hierarchy Level | 29363
- Description | 29363
- Default | 29363
- Options | 29363
- Required Privilege Level | 29365
- Release Information | 29365

Syntax

```
time-zone (GMT hour-offset | time-zone);
```

Hierarchy Level

```
[edit system]
```

Description

Set the local time zone. To have the time zone change take effect for all processes running on the router or switch, you must reboot the router or switch.

Default

UTC

Options

`GMT` *`hour-offset`*—Set the time zone relative to UTC time.

- **Range:** –14 through +12
- **Default:** 0

`time-zone`—Specify the time zone as `UTC`, which is the default time zone, or as a string such as PDT (Pacific Daylight Time), or use one of the following continents and major cities:

```
Africa/Abidjan, Africa/Accra, Africa/Addis_Ababa, Africa/Algiers, Africa/Asmera, Africa/Bamako,
Africa/Bangui, Africa/Banjul, Africa/Bissau, Africa/Blantyre, Africa/Brazzaville, Africa/
Bujumbura, Africa/Cairo, Africa/Casablanca, Africa/Ceuta, Africa/Conakry, Africa/Dakar, Africa/
Dar_es_Salaam, Africa/Djibouti, Africa/Douala, Africa/El_Aaiun, Africa/Freetown, Africa/
Gaborone, Africa/Harare, Africa/Johannesburg, Africa/Kampala, Africa/Khartoum, Africa/Kigali,
Africa/Kinshasa, Africa/Lagos, Africa/Libreville, Africa/Lome, Africa/Luanda, Africa/Lubumbashi,
Africa/Lusaka, Africa/Malabo, Africa/Maputo, Africa/Maseru, Africa/Mbabane, Africa/Mogadishu,
Africa/Monrovia, Africa/Nairobi, Africa/Ndjamena, Africa/Niamey, Africa/Nouakchott, Africa/
Ouagadougou, Africa/Porto-Novo, Africa/Sao_Tome, Africa/Timbuktu, Africa/Tripoli, Africa/Tunis,
Africa/Windhoek
America/Adak, America/Anchorage, America/Anguilla, America/Antigua, America/Aruba, America/
Asuncion, America/Barbados, America/Belize, America/Bogota, America/Boise, America/Buenos_Aires,
America/Caracas, America/Catamarca, America/Cayenne, America/Cayman, America/Chicago, America/
Cordoba, America/Costa_Rica, America/Cuiaba, America/Curacao, America/Dawson, America/
Dawson_Creek, America/Denver, America/Detroit, America/Dominica, America/Edmonton, America/
El_Salvador, America/Ensenada, America/Fortaleza, America/Glace_Bay, America/Godthab, America/
Goose_Bay, America/Grand_Turk, America/Grenada, America/Guadeloupe, America/Guatemala, America/
```

```
Guayaquil, America/Guyana, America/Halifax, America/Havana, America/Indiana/Knox, America/
Indiana/Marengo, America/Indiana/Vevay, America/Indianapolis, America/Inuvik, America/Iqaluit,
America/Jamaica, America/Jujuy, America/Juneau, America/La_Paz, America/Lima, America/
Los_Angeles, America/Louisville, America/Maceio, America/Managua, America/Manaus, America/
Martinique, America/Mazatlan, America/Mendoza, America/Menominee, America/Mexico_City, America/
Miquelon, America/Montevideo, America/Montreal, America/Montserrat, America/Nassau, America/
New_York, America/Nipigon, America/Nome, America/Noronha, America/Panama, America/Pangnirtung,
America/Paramaribo, America/Phoenix, America/Port-au-Prince, America/Port_of_Spain, America/
Porto_Acre, America/Puerto_Rico, America/Rainy_River, America/Rankin_Inlet, America/Regina,
America/Rosario, America/Santiago, America/Santo_Domingo, America/Sao_Paulo, America/
Scoresbysund, America/Shiprock, America/St_Johns, America/St_Kitts, America/St_Lucia, America/
St_Thomas, America/St_Vincent, America/Swift_Current, America/Tegucigalpa, America/Thule,
America/Thunder_Bay, America/Tijuana, America/Tortola, America/Vancouver, America/Whitehorse,
America/Winnipeg, America/Yakutat, America/Yellowknife
Antarctica/Casey, Antarctica/DumontDUrville, Antarctica/Mawson, Antarctica/McMurdo, Antarctica/
Palmer, Antarctica/South_Pole
Arctic/Longyearbyen
Asia/Aden, Asia/Alma-Ata, Asia/Amman, Asia/Anadyr, Asia/Aqtau, Asia/Aqtobe, Asia/Ashkhabad, Asia/
Baghdad, Asia/Bahrain, Asia/Baku, Asia/Bangkok, Asia/Beirut, Asia/Bishkek, Asia/Brunei, Asia/
Chungking, Asia/Colombo, Asia/Dacca, Asia/Damascus, Asia/Dubai, Asia/Dushanbe, Asia/Gaza, Asia/
Harbin, Asia/Hong_Kong, Asia/Irkutsk, Asia/Ishigaki, Asia/Jakarta, Asia/Jayapura, Asia/
Jerusalem, Asia/Kabul, Asia/Kamchatka, Asia/Karachi, Asia/Kashgar, Asia/Katmandu, Asia/Kolkata ,
Asia/Krasnoyarsk, Asia/Kuala_Lumpur, Asia/Kuching, Asia/Kuwait, Asia/Macao, Asia/Magadan, Asia/
Manila, Asia/Muscat, Asia/Nicosia, Asia/Novosibirsk, Asia/Omsk, Asia/Phnom_Penh, Asia/Pyongyang,
Asia/Qatar, Asia/Rangoon, Asia/Riyadh, Asia/Saigon, Asia/Seoul, Asia/Shanghai, Asia/Singapore,
Asia/Taipei, Asia/Tashkent, Asia/Tbilisi, Asia/Tehran, Asia/Thimbu, Asia/Tokyo, Asia/
Ujung_Pandang, Asia/Ulan_Bator, Asia/Urumqi, Asia/Vientiane, Asia/Vladivostok, Asia/Yakutsk,
Asia/Yekaterinburg, Asia/Yerevan
Atlantic/Azores, Atlantic/Bermuda, Atlantic/Canary, Atlantic/Cape_Verde, Atlantic/Faeroe,
Atlantic/Jan_Mayen, Atlantic/Madeira, Atlantic/Reykjavik, Atlantic/South_Georgia, Atlantic/
St_Helena, Atlantic/Stanley
Australia/Adelaide, Australia/Brisbane, Australia/Broken_Hill, Australia/Darwin, Australia/
Hobart, Australia/Lindeman, Australia/Lord_Howe, Australia/Melbourne, Australia/Perth, Australia/
Sydney
Europe/Amsterdam, Europe/Andorra, Europe/Athens, Europe/Belfast, Europe/Belgrade, Europe/Berlin,
Europe/Bratislava, Europe/Brussels, Europe/Bucharest, Europe/Budapest, Europe/Chisinau, Europe/
Copenhagen, Europe/Dublin, Europe/Gibraltar, Europe/Helsinki, Europe/Istanbul, Europe/
Kaliningrad, Europe/Kiev, Europe/Lisbon, Europe/Ljubljana, Europe/London, Europe/Luxembourg,
Europe/Madrid, Europe/Malta, Europe/Minsk, Europe/Monaco, Europe/Moscow, Europe/Oslo, Europe/
Paris, Europe/Prague, Europe/Riga, Europe/Rome, Europe/Samara, Europe/San_Marino, Europe/
Sarajevo, Europe/Simferopol, Europe/Skopje, Europe/Sofia, Europe/Stockholm, Europe/Tallinn,
Europe/Tirane, Europe/Vaduz, Europe/Vatican, Europe/Vienna, Europe/Vilnius, Europe/Warsaw,
Europe/Zagreb, Europe/Zurich
```

```
Indian/Antananarivo, Indian/Chagos, Indian/Christmas, Indian/Cocos, Indian/Comoro, Indian/
Kerguelen, Indian/Mahe, Indian/Maldives, Indian/Mauritius, Indian/Mayotte, Indian/Reunion
Pacific/Apia, Pacific/Auckland, Pacific/Chatham, Pacific/Easter, Pacific/Efate, Pacific/
Enderbury, Pacific/Fakaofo, Pacific/Fiji, Pacific/Funafuti, Pacific/Galapagos, Pacific/Gambier,
Pacific/Guadalcanal, Pacific/Guam, Pacific/Honolulu, Pacific/Johnston, Pacific/Kiritimati,
Pacific/Kosrae, Pacific/Kwajalein, Pacific/Majuro, Pacific/Marquesas, Pacific/Midway, Pacific/
Nauru, Pacific/Niue, Pacific/Norfolk, Pacific/Noumea, Pacific/Pago_Pago, Pacific/Palau, Pacific/
Pitcairn, Pacific/Ponape, Pacific/Port_Moresby, Pacific/Rarotonga, Pacific/Saipan, Pacific/
Tahiti, Pacific/Tarawa, Pacific/Tongatapu, Pacific/Truk, Pacific/Wake, Pacific/Wallis, Pacific/
Yap
```

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`GMT` *`hour-offset`* option added in Junos OS Release 7.4.

RELATED DOCUMENTATION

Modifying the Default Time Zone for a Router or Switch Running Junos OS

timeout (SNMP)

IN THIS SECTION

- Syntax | 29366
- Hierarchy Level | 29366
- Description | 29366
- Options | 29366
- Required Privilege Level | 29366

Release Information | 29366

Syntax

```
timeout seconds;
```

Hierarchy Level

```
[edit snmp v3 target-address target-address-name]
```

Description

Configure the timeout period (in seconds) for SNMP informs.

Options

seconds—Number of seconds to wait for an inform acknowledgment. If no acknowledgment is received within the timeout period, the inform is retransmitted.

- **Default:** 15

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

retry-count (SNMP) | 15697

timeout

IN THIS SECTION

- Syntax | 29367
- Hierarchy Level | 29367
- Description | 29367
- Default | 29367
- Options | 29368
- Required Privilege Level | 29368
- Release Information | 29368

Syntax

```
timeout seconds;
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port dhcp-snooping-file],
[edit ethernet-switching-options secure-access-port dhcpv6-snooping-file]
```

Description

Specify a timeout value for remote read and write operations. This value determines the amount of time that the switch waits for a remote system to respond when the DHCP snooping database is stored on the remote FTP site.

Default

None

Options

seconds—Value in seconds.

- **Range:** 10 through 3600.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Support at the `[edit ethernet-switching-options secure-access-port dhcpv6-snooping-file]` hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

RELATED DOCUMENTATION

Configuring Persistent Bindings in the DHCP or DHCPv6 (non-ELS)

Understanding DHCP Snooping (non-ELS)

timeout (Access LDAP)

IN THIS SECTION

- Syntax | 29369
- Hierarchy Level | 29369
- Description | 29369
- Options | 29369
- Required Privilege Level | 29369
- Release Information | 29369

Syntax

```
timeout seconds;
```

Hierarchy Level

```
[edit access ldap-server server-address]
[edit access profile profile-name ldap-server server-address]
```

Description

Configure the amount of time that the local device waits to receive a response from an LDAP server.

Options

seconds—Amount of time to wait.

- **Range:** 1 through 90 seconds
- **Default:** 3 seconds

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

> *Example: Configuring RADIUS and LDAP User Authentication*

timeout (Access RADIUS)

IN THIS SECTION

- Syntax | 29370
- Hierarchy Level | 29370
- Description | 29370
- Options | 29370
- Required Privilege Level | 29371
- Release Information | 29371

Syntax

```
timeout seconds;
```

Hierarchy Level

```
[edit access radius-server server-address]
[edit access profile profile-name radius-server server-address]
```

Description

Configure the amount of time that the local device waits to receive a response from a RADIUS server.

Options

seconds—Amount of time to wait.

- **Range:** 1 through 90 seconds
- **Default:** 3 seconds

Required Privilege Level

secret—To view this statement in the configuration.

secret-control—To add this statement to the configuration.

Release Information

Statement modified in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Firewall User Authentication Overview

timeout-action

IN THIS SECTION

- Syntax | 29371
- Hierarchy Level | 29371
- Description | 29372
- Options | 29372
- Required Privilege Level | 29372
- Release Information | 29372

Syntax

```
timeout-action (close | no-change | open):
```

Hierarchy Level

```
[edit services unified-access-control ]
```

Description

Specify what the SRX Series Firewall should do when a timeout occurs and the device cannot connect to an Infranet Enforcer.

Use this statement when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series UAC Appliance.

Options

- `close`—Close existing sessions and block any further traffic. This is the default option.
- `no-change`—Preserve existing sessions and require authentication for new sessions.
- `open`—Preserve existing sessions and allow new sessions access.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

interval (Services) | 9221

timeout (Services) | 29395

timeout-active-flows (DDoS Flow Detection)

IN THIS SECTION

Syntax | 29373

- Hierarchy Level | 29373
- Description | 29373
- Required Privilege Level | 29373
- Release Information | 29373

Syntax

```
timeout-active-flows;
```

Hierarchy Level

```
[edit system ddos-protection protocols protocol-group packet-type]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Enable culprit flows for the packet type to time out according to the timeout period. The culprit flow is suppressed for the duration of the timeout period. When the period expires, the flow times out and is released from suppression.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Configuring the Timeout Period for a Culprit Flow

Setting Up and Using Flow Detection

timeout (DHCP Local Server)

IN THIS SECTION

- Syntax | 29374
- Hierarchy Level | 29374
- Description | 29375
- Options | 29375
- Required Privilege Level | 29375
- Release Information | 29375

Syntax

```
timeout timeout-value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
reconfigure],
```

```
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit system services dhcp-local-server reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server group group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Configure the initial value in seconds between attempts to reconfigure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces.

Options

timeout-value—Initial retry timeout value.

- **Range:** 1 through 10 seconds
- **Default:** 2 seconds

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview

Configuring Dynamic Reconfiguration Attempts for DHCP Clients

timeout (Diameter Network Element)

IN THIS SECTION

- Syntax | 29376
- Hierarchy Level | 29376
- Description | 29376
- Options | 29376
- Required Privilege Level | 29377
- Release Information | 29377

Syntax

```
timeout seconds;
```

Hierarchy Level

```
[edit access diameter network-element element-name peer peer-name]
```

Description

Configure the amount of time to wait for a response from this peer before transmitting the request to another peer.

Options

seconds Amount of time to wait before transmitting the request.

- **Range:** 1 through 100 seconds
- **Default:** 4 seconds

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

timeout (Flow Maps)

IN THIS SECTION

- Syntax | 29377
- Hierarchy Level | 29378
- Description | 29378
- Options | 29378
- Required Privilege Level | 29378
- Release Information | 29378

Syntax

```
timeout (never non-discard-entry-only | minutes);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast flow-map flow-map-name],
[edit logical-systems logical-system-name routing-options multicast flow-map flow-map-name],
[edit routing-instances routing-instance-name routing-options multicast flow-map flow-map-name],
[edit routing-options multicast flow-map flow-map-name]
```

Description

Configure the timeout value for multicast forwarding cache entries associated with the flow map.

Options

minutes—Length of time that the forwarding cache entry remains active.

- **Range:** 1 through 720

never non-discard-entry-only—Specify that the forwarding cache entry always remain active. If you omit the **non-discard-entry-only** option, all multicast forwarding entries, including those in forwarding and pruned states, are kept forever. If you include the **non-discard-entry-only** option, entries with forwarding states are kept forever, and entries with pruned states time out.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

timeout-grace (Access)

IN THIS SECTION

- Syntax | 29379
- Hierarchy Level | 29379
- Description | 29379
- Options | 29379
- Required Privilege Level | 29380
- Release Information | 29380

Syntax

```
timeout grace seconds;
```

Hierarchy Level

```
[edit access radius-options]
```

Description

Configure a grace period during which a RADIUS server that times out after failing to respond to an authentication request is not considered to be down (if other servers are available) or unreachable (if it is the only configured server). The server is marked as down or unreachable only when it does not successfully respond to an authentication request during the grace period and times out again after the period expires.

You can use long grace periods to give unresponsive servers more opportunities to respond before abandoning them. You can use short grace periods to cause the router to declare unresponsive servers down and direct requests to available servers sooner.

Options

seconds—Duration of the server timeout grace period.

- **Range:** 0 through 30
- **Default:** 10

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

RADIUS Servers and Parameters for Subscriber Access

timeout (Multicast)

IN THIS SECTION

- Syntax | **29380**
- Hierarchy Level | **29381**
- Description | **29381**
- Options | **29381**
- Required Privilege Level | **29381**
- Release Information | **29381**

Syntax

```
timeout minutes <family (inet | inet6)>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast forwarding-cache],
[edit logical-systems logical-system-name routing-options multicast forwarding-cache],
[edit routing-instances routing-instance-name routing-options multicast forwarding-cache],
[edit routing-options multicast forwarding-cache]
```

Description

Configure the timeout value for multicast forwarding cache entries. In general, you should regularly refresh the forwarding cache so it does not fill up with old entries and thus prevent newer, higher-priority, entries from being added.

Options

minutes—Length of time that the forwarding cache limit remains active.

- **Range:** 1 through 720

family (inet | inet6)—(Optional) Apply the configured timeout to either IPv4 or IPv6 multicast forwarding cache entries. Configuring the `timeout` statement globally for the multicast forwarding cache or including the `family` statement to configure the timeout value for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.

- **Default:** Six minutes. By default, the configured timeout applies to both IPv4 and IPv6 multicast forwarding cache entries.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

> *Example: Configuring the Multicast Forwarding Cache*

timeout (RADIUS Server)

IN THIS SECTION

- Syntax | 29382
- Hierarchy Level | 29382
- Description | 29382
- Options | 29382
- Required Privilege Level | 29383
- Release Information | 29383

Syntax

```
timeout seconds;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure the amount of time that the MX Series router waits to receive a response from a RADIUS server before retrying the request.

Options

seconds

Number of seconds to wait.

- **Range:** 1 through 90
- **Default:** 3

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring RADIUS Servers

Understanding Network Elements

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

timeout (Security Antivirus Fallback Options)

IN THIS SECTION

- Syntax | 29383
- Hierarchy Level | 29384
- Description | 29384
- Options | 29384
- Required Privilege Level | 29384
- Release Information | 29384

Syntax

```
timeout (block | log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name fallback-
options]
```

Description

Scanning a complex file could consume resources and time. If the time it is taking to scan exceeds the timeout setting in the antivirus profile, the processing is terminated and the content is either passed or blocked without completing the virus checking. The decision is made based on the timeout fallback option. The default action is block.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

timeout (Security Antivirus Fallback Options Sophos Engine)

IN THIS SECTION

- Syntax | 29385
- Hierarchy Level | 29385
- Description | 29386
- Options | 29386
- Required Privilege Level | 29386
- Release Information | 29386

Syntax

```
default (block | log-and-permit | permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name fallback-
options]
```

Description

Scanning a complex file could consume resources and time. If the time it is taking to scan exceeds the timeout setting in the antivirus profile, the processing is terminated and the content is either passed or blocked without completing the virus checking. The decision is made based on the timeout fallback option.

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic
- `permit`—Permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

timeout (Security Antivirus Scan Options)

IN THIS SECTION

- Syntax | 29387
- Hierarchy Level | 29387
- Description | 29387
- Required Privilege Level | 29387
- Release Information | 29387

Syntax

```
timeout value;
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name scan-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name scan-options]
[edit security utm default-configuration anti-virus scan-options]
```

Description

The scanning timeout value includes the time frame from when the scan request is generated to when the scan result is returned by the scan engine. The time range can be 1 to 1800 seconds. By default, it is 180 seconds.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for Avira engine added in Junos OS Release 18.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

timeout (Security GTP)

IN THIS SECTION

- Syntax | 29388
- Hierarchy Level | 29388
- Description | 29388
- Options | 29389
- Required Privilege Level | 29389
- Release Information | 29389

Syntax

```
timeout value;
```

Hierarchy Level

```
[edit security gtp profile profile-name]
```

Description

Use the timeout option to set the tunnel timeout value in hours. The default is 36 hours. If a device detects no activity in a tunnel for a specified period, it removes the tunnel from the state table.

From Junos OS Release 20.4R1 onwards, the [`edit security gprs gtp profile` *`profile-name`*] hierarchy level is replaced by [`edit security gtp profile` *`profile-name`*].

Options

- **Range:** 1 through 1,000 hours.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. Support.

RELATED DOCUMENTATION

Understanding GTPv2 Tunnel Cleanup

Example: Setting the Timeout Value for GTPv2 Tunnels

timeout (Security Screen)

IN THIS SECTION

- Syntax | 29390
- Hierarchy Level | 29390
- Description | 29390
- Options | 29390
- Required Privilege Level | 29390
- Release Information | 29390

Syntax

```
timeout seconds;
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp syn-flood]
```

Description

Specify the maximum length of time before a half-completed connection is dropped from the queue. You can decrease the timeout value until you see any connections dropped during normal traffic conditions.

Options

seconds —Time interval before a half-completed connection is dropped from the queue.

- **Range:** 1 through 50 seconds
- **Default:** 20 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Attack Detection and Prevention Overview*
>
> *Example: Configuring Multiple Screening Options*

timeout (Security Web Filtering)

IN THIS SECTION

- Syntax | 29391
- Hierarchy Level | 29391
- Description | 29391
- Required Privilege Level | 29391
- Release Information | 29391

Syntax

```
timeout value;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name]
```

Description

Enter a timeout limit for requests. Once this limit is reached, fail mode settings are applied. The default here is 15 seconds.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The [edit security utm default-configuration] hierarchy level is introduced in Junos OS Release 18.2R1.

timeout (Security Web Filtering Cache)

IN THIS SECTION

- Syntax | 29392
- Hierarchy Level | 29392
- Description | 29392
- Options | 29392
- Required Privilege Level | 29393
- Release Information | 29393

Syntax

```
timeout value;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated cache]
[edit security utm feature-profile web-filtering juniper-enhanced cache]
```

Description

Set the cache timeout parameters for surf-control-integrated web filtering (24 hours is the default and the maximum allowed life span of cached items).

Options

- **Range:** 1 through 1800 minutes.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The surf-control-integrated feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The surf-control-integrated feature is not supported from Junos OS Release 15.1X49-D10 onwards.

timeout (Security Web Filtering Fallback Settings)

IN THIS SECTION

- Syntax | 29394
- Hierarchy Level | 29394
- Description | 29394
- Options | 29394
- Required Privilege Level | 29394
- Release Information | 29394

Syntax

```
timeout (block | log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated profile profile-name
fallback-settings]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name fallback-
settings]
[edit security utm feature-profile web-filtering juniper-enhanced profile profile-name fallback-
settings]
```

Description

Fallback settings tell the system how to handle errors.

Options

- log-and-permit—Log the error and permit the traffic
- block—Log the error and deny the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The surf-control-integrated feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The surf-control-integrated feature is not supported from Junos OS Release 15.1X49-D10 onwards.

timeout (Services)

IN THIS SECTION

- Syntax | 29395
- Hierarchy Level | 29395
- Description | 29395
- Required Privilege Level | 29396
- Release Information | 29396

Syntax

```
timeout seconds;
```

Hierarchy Level

```
[edit services unified-access-control ]
```

Description

Specify the value, in seconds, that the SRX Series Firewall should wait to get a heartbeat response from an IC Series UAC Appliance (default is 300). If the SRX Series Firewall does not receive it in the specified time, it takes the action specified by the `timeout-action` configuration statement. It also tries again to make a connection to the IC Series appliance. After the second failed attempt, the SRX Series Firewall fails

over to the next IC Series appliance in the cluster. The SRX Series Firewall continues trying to reach IC Series appliances in the cluster until a connection is established.

Use this statement when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series appliance. When working with a cluster of IC Series appliances, the Junos OS Enforcer connects to one at a time, failing over to other IC Series appliances in the cluster as required.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

interval (Services) | 9221

timeout-action | 29371

timestamp

IN THIS SECTION

- Syntax | 29397
- Hierarchy Level | 29397
- Description | 29397
- Options | 29397
- Required Privilege Level | 29397
- Release Information | 29398

Syntax

```
timestamp {
    clock-drift value;
    known-peer-window seconds;
    new-peer-window seconds;
}
```

Hierarchy Level

```
[edit protocols neighbor-discovery secure]
```

Description

Configure timestamp options, which are used to ensure that solicitation and redirect messages are not being replayed.

Options

clock-drift *value*—Specify the allowable drift in time between the synchronization of peers. For ***value***, specify a fractional value of 100.

- **Default:** 0.01

known-peer-window *seconds*—Specify the expected interval in seconds between Secure Neighbor Discovery messages from an established peer. A message from a known peer that arrives after the specified interval is discarded.

- **Default:** 1 second

new-peer-window *seconds*—Specify the maximum allowable time in seconds between the timestamp of a Secure Neighbor Discovery message from a new peer and when it can be accepted.

- **Default:** 300 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Secure IPv6 Neighbor Discovery

timestamp-format (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 29398
- Hierarchy Level | 29398
- Description | 29398
- Required Privilege Level | 29399
- Release Information | 29399

Syntax

```
timestamp-format;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specify the time stamp format.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

timestamp-option (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | **29399**
- Hierarchy Level | **29400**
- Description | **29400**
- Required Privilege Level | **29400**
- Release Information | **29400**

Syntax

```
timestamp-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop IPv4 packets that have the IP option of 4 (Internet timestamp).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

tls

IN THIS SECTION

- Syntax | 29401
- Hierarchy Level | 29401
- Description | 29401
- Release Information | 29401

Syntax

```
tls
```

Hierarchy Level

```
[edit system syslog host host transport]
```

Description

Configure this Transport Layer Security (TLS) protocol to forward syslog messages securely.

This configuration uses certificates to authenticate peers. Before configuring this statement, make sure that your device has CA certificate and the PKI. To know more about CA certificate and PKI, refer Self-Signed Digital Certificates.

Encapsulating syslog over TLS allows you to:

- Validate the remote destination (syslog server) before transmitting any sensitive syslog information. (Authentication)
- Encrypt the syslog during the transport. (Encryption)
- Verify that the data has not been modified or tampered with (Integrity).

The remaining statements are explained separately. See CLI Explorer.

Release Information

Statement introduced in Junos OS Release 21.2R1.

Statement introduced in Junos OS Evolved Release 23.4R1.

tls-min-version

IN THIS SECTION

- Syntax | 29402
- Hierarchy Level | 29402
- Description | 29402
- Options | 29402
- Required Privilege Level | 29403
- Release Information | 29403

Syntax

```
tls-min-version (v1.1 | v1.2);
```

Hierarchy Level

```
[edit access profile profile-name ldap-server ip-address]
```

Description

Configure Transport Layer Security (TLS) version to limit the lowest supported versions of TLS that are enabled for SSL connections.

Options

v1.1 Accept TLS version 1.1. This enhanced version of TLS provides protection against cipher-block chaining (CBC) attacks.

v1.2 Accept TLS version 1.2. This enhanced version of TLS provides improved flexibility for negotiation of cryptographic algorithms.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

> *Firewall User Authentication Overview*
>
> *Example: Configuring Integrated User Firewall on SRX Series*

tls (NETCONF)

IN THIS SECTION

- Syntax | 29403
- Hierarchy Level | 29404
- Description | 29404
- Options | 29404
- Required Privilege Level | 29404
- Release Information | 29405

Syntax

```
tls {
    client-identity client-id {
        fingerprint fingerprint;
        map-type (san-dirname-cn | specified);
        username username;
    }
```

```
    default-client-identity {
        map-type (san-dirname-cn | specified);
        username username;
    }
    local-certificate local-certificate;
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system services netconf]
```

Description

Enable NETCONF sessions over Transport Layer Security (TLS) with mutual X.509 certificate-based authentication. To enable NETCONF sessions over TLS, you must configure the `local-certificate` statement and either a `client-identity` statement or the `default-client-identity` statement.

Junos devices support TLS version 1.2 for NETCONF sessions over TLS. The TLS server listens for incoming NETCONF-over-TLS connections on TCP port 6513.

Options

`local-certificate` *`local-certificate`*	TLS server's local certificate ID, which must be loaded into the Junos public key infrastructure (PKI).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

Statement introduced in Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

NETCONF Sessions over Transport Layer Security (TLS)

tls-peer-name

IN THIS SECTION

- Syntax | 29405
- Hierarchy Level | 29405
- Description | 29405
- Required Privilege Level | 29406
- Release Information | 29406

Syntax

```
tls-peer-name peer-host-name;
```

Hierarchy Level

```
[edit access profile profile-name ldap-server ip-address]
```

Description

Configure the peer hostname to be authenticated.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Obtaining Username and Role Information Through Firewall Authentication

LDAP Functionality in Integrated User Firewall

tls-timeout

IN THIS SECTION

- Syntax | 29406
- Hierarchy Level | 29407
- Description | 29407
- Required Privilege Level | 29407
- Release Information | 29407

Syntax

```
tls-timeout seconds;
```

Hierarchy Level

```
[edit access profile profile-name ldap-server ip-address]
```

Description

Specify timeout value on the Transport Layer Security (TLS) handshake. The TLS handshake is responsible for the encryption keys exchange necessary to establish secure sessions between client and server.

Range: 3 through 90 seconds.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Obtaining Username and Role Information Through Firewall Authentication

LDAP Functionality in Integrated User Firewall

tls-type

IN THIS SECTION

- Syntax | 29408
- Hierarchy Level | 29408
- Description | 29408

- Options | 29408
- Required Privilege Level | 29408
- Release Information | 29409

Syntax

```
tls-type {
    start-tls;
}
```

Hierarchy Level

```
[edit access profile profile-name ldap-server ip-address]
```

Description

Configure Lightweight Directory Access Protocol (LDAP) over Secure Sockets Layer/Transport Layer Security (SSL/TLS) for secure communication. Transport Layer Security StartTLS extension for LDAP is used for the firewall user authentication and the integrated user firewall authentication for obtaining username and role information through firewall authentication. StartTTLS allows protocol data transfers between the LDAP server and client over the TLS layer after successful negotiation between the peers. StartTLS upgrades an existing insecure LDAP connection to a secure Secure Sockets Layer/Transport Layer Security (SSL/TLS) connection.

Options

- **start-tls**—Configure LDAP over StartTLS. The StartTLS communications occurs over TCP port 389.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

RELATED DOCUMENTATION

Firewall User Authentication Overview

Obtaining Username and Role Information Through Firewall Authentication

LDAP Functionality in Integrated User Firewall

tlsdetails

IN THIS SECTION

- Syntax | 29409
- Hierarchy Level | 29409
- Description | 29410
- Options | 29410
- Release Information | 29410

Syntax

```
tlsdetails {
       local-certificate certificate-id;
       trusted-ca-group trusted-ca-group-name {
            ca-profiles ca-profile-name};
      }
```

Hierarchy Level

```
[edit system syslog host host]
```

Description

Configure the TLS authentication parameter.

Options

local-certificate ***certificate-id***	Use this option to specify the local certificate ID for authentication. The local certificate contains the public key and identify information for the device.
certificate-authority ***certificate-authority***	Use this option to specify Certificate Authority (CA) for authentication.
trusted-ca-group ***trusted-ca-group-name***	Use this option to specify trusted CA group for authentication.
ca-profiles ***ca-profile-name***	CA profile contains the name and URL of the CA or RA, as well as retry-timer settings.

Release Information

Statement introduced in Junos OS Release 21.2R1 for EX and MX Series routers.

Statement introduced in Junos OS Release 21.3R1 for SRX Series and vSRX Virtual Firewall devices.

RELATED DOCUMENTATION

Configuring a Trusted CA Profile Group

tlv-filter

IN THIS SECTION

- Syntax | 29411
- Hierarchy Level | 29411
- Description | 29411
- Default | 29412

- Options | 29412
- Required Privilege Level | 29413
- Release Information | 29413

Syntax

```
tlv-filter tlv-name;
```

Hierarchy Level

```
[edit protocols lldp],
[edit protocols lldp-med],
[edit protocols lldp interface interface-name],
[edit protocols lldp-med interface interface-name]
```

Description

Select the type, length, and value (TLV) messages that should not be advertised by the Link Layer Discovery Protocol (LLDP) or LLDP Media Endpoint Discovery (LLDP-MED) protocol. LLDP-capable devices transmit information in type, length, and value (TLV) messages to neighbor devices. Device information can include information such as chassis and port identification and system name and system capabilities.

In multi-vendor networks, it might not be desirable to send TLV messages because they can contain sensitive information about a network device. You can configure LLDP or LLDP-MED to disable any non-mandatory TLV message. (These mandatory TLVs are always advertised: chassis-id, port-id, and time-to-live.)

When you configure the `tlv-filter` statement, you specify the TLVs that you want to disable. This is useful when you want to allow most, but not all, TLVs.

You can also disable TLVs using the `tlv-select` statement. When you configure the tlv-select statement, you specify the TLVs that you want to be advertised by LLDP or LLDP-MED. All other non-mandatory TLVs are disabled.

> **NOTE**: The `tlv-select` and `tlv-filter` statements are mutually exclusive and cannot be used on the same configuration stanza at the same time.

Default

All TLVs for LLDP and LLDP-MED are enabled by default.

Options

tlv-name Specify the non-mandatory TLV message that you want to disable.

- **Values:** You can disable the following TLV messages for LLDP:
  - `cloud-reach-event-tlv`—Use to have cloud-ready EX4400 Series switches stop sending cloud-reach-events. By default, once an EX4400 connects to the Juniper Mist Cloud Service via its phone-home routine, subsequent TLV messages are automatically halted.
  - `jnpr-chassis-serial`—The chassis serial number.
  - `jnpr-vcp`—Juniper virtual chassis port.
  - `link-aggregation`—Advertises whether the port is aggregated and its aggregated port ID.
  - `mac-phy-config-status`—Advertises information about the physical interface, such as autonegotiation status and support and MAU (medium attachment unit) type.
  - `management-address`— The IP management address of the local system.
  - `maximum-frame-size`—The maximum transmission unit (MTU) of the interface sending LLDP frames.
  - `port-description`—The user-configured port description.
  - `port-vid`—Indicates the port VLAN ID that will be associated with an untagged or priority tagged data frame received on the VLAN port.
  - `power-vi-mdi`—Advertises MDI (media dependent interface) power support, PSE (power sourcing equipment) power pair, and power class information.

- `system-capabilities`—The primary function performed by the system. The capabilities that system supports are defined; for example, bridge or router. This information cannot be configured, but is based on the model of the product.
- `system-description`—The user-configured name of the local system.
- `system-name`—The user-configured name of the local system.
- `vlan-name`—Indicates the assigned name of any VLAN at the device.

You can disable the following TLV messages for LLDP-MED:

- `ext-power-via-mdi`—The power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port.
- `location-id`—The physical location of the endpoint.
- `med-capabilities`—The primary function of the port.
- `network-policy`—The port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier, application types, such as voice or streaming video, 802.1Q VLAN tagging, and 802.1p priority bits and Diffserv code points.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3.

RELATED DOCUMENTATION

show lldp | **21942**

Configuring LLDP (CLI Procedure)

Understanding LLDP and LLDP-MED on EX Series Switches

tlv-select

IN THIS SECTION

- Syntax | 29414
- Hierarchy Level | 29414
- Description | 29414
- Default | 29415
- Options | 29415
- Required Privilege Level | 29416
- Release Information | 29416

Syntax

```
tlv-select tlv-name;
```

Hierarchy Level

```
[edit protocols lldp],
[edit protocols lldp-med],
[edit protocols lldp interface interface-name],
[edit protocols lldp-med interface interface-name]
```

Description

Select the type, length, and value (TLV) messages that should be advertised by Link Layer Discovery Protocol (LLDP) or LLDP Media Endpoint Discovery (LLDP-MED) protocol. LLDP-capable devices transmit information in type, length, and value (TLV) messages to neighbor devices. Device information can include information such as chassis and port identification and system name and system capabilities.

In multi-vendor networks, it might not be desirable to send TLV messages because they can contain sensitive information about a network device. You can configure LLDP or LLDP-MED to choose which non-mandatory TLV messages to advertise. (These mandatory TLVs are always advertised: chassis-id, port-id, and time-to-live.)

When you configure the tlv-select statement, you specify the TLVs that LLDP or LLDP-MED should advertise. All other non-mandatory TLVs are disabled. This is useful when you want to disable most, but not all, TLVs.

You can also disable TLVs using the tlv-filter statement. When you configure the tlv-filter statement, you specify the TLVs that should be disabled.

> **NOTE**: The `tlv-select` and `tlv-filter` statements are mutually exclusive and cannot be used on the same configuration stanza at the same time.

Default

All TLVs for LLDP and LLDP-MED are enabled by default.

Options

tlv-name Specify the non-mandatory TLV message that you want to advertise.

- **Values:** You can advertise the following TLV messages for LLDP:
  - `jnpr-chassis-serial`—The chassis serial number.
  - `jnpr-vcp`—Juniper virtual chassis port.
  - `link-aggregation`—Advertises whether the port is aggregated and its aggregated port ID.
  - `mac-phy-config-status`—Advertises information about the physical interface, such as autonegotiation status and support and MAU (medium attachment unit) type.
  - `management-address`— The IP management address of the local system.
  - `maximum-frame-size`—The maximum transmission unit (MTU) of the interface sending LLDP frames.
  - `port-description`—The user-configured port description.
  - `port-vid`—Indicates the port VLAN ID that will be associated with an untagged or priority tagged data frame received on the VLAN port.
  - `power-vi-mdi`—Advertises MDI (media dependent interface) power support, PSE (power sourcing equipment) power pair, and power class information.

- `system-capabilities`—The primary function performed by the system. The capabilities that system supports are defined; for example, bridge or router. This information cannot be configured, but is based on the model of the product.
- `system-description`—The user-configured name of the local system.
- `system-name`—The user-configured name of the local system.
- `vlan-name`—Indicates the assigned name of any VLAN at the device.

You can advertise the following TLV messages for LLDP-MED:

- `ext-power-via-mdi`—The power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port.
- `location-id`—The physical location of the endpoint.
- `med-capabilities`—The primary function of the port.
- `network-policy`—The port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier, application types, such as voice or streaming video, 802.1Q VLAN tagging, and 802.1p priority bits and Diffserv code points.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3.

RELATED DOCUMENTATION

show lldp | **21942**

Configuring LLDP (CLI Procedure)

Understanding LLDP and LLDP-MED on EX Series Switches

tm

IN THIS SECTION

- Syntax | 29417
- Hierarchy Level | 29417
- Description | 29417
- Options | 29417
- Required Privilege Level | 29418
- Release Information | 29418

Syntax

```
tm (ignore | normal | require);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options dce-options],
[edit interfaces interface-name serial-options dte-options]
```

Description

For EIA-530 interfaces only, configure the from-DCE signal, test-mode (TM).

Options

`ignore`—The from-DCE signal is ignored.

`normal`—Normal TM signal handling as defined by the TIA/EIA Standard 530.

`require`—The from-DCE signal must be asserted.

- **Default:** `normal`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Signal Handling

tm-polarity

IN THIS SECTION

- Syntax | 29418
- Hierarchy Level | 29418
- Description | 29419
- Options | 29419
- Required Privilege Level | 29419
- Release Information | 29419

Syntax

```
tm-polarity (negative | positive);
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

Configure TM signal polarity.

Options

`negative`—Negative signal polarity.

`positive`—Positive signal polarity.

- **Default:** `positive`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Serial Signal Polarities

to

IN THIS SECTION

- Syntax | 29420
- Hierarchy Level | 29420
- Description | 29420
- Options | 29420
- Required Privilege Level | 29420
- Release Information | 29420

Syntax

```
to address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
bypass],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name
ingress],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name bypass],
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Description

Specify the egress router of a dynamic LSP.

Options

address—IPv4 or IPv6 address of the egress router.

> **NOTE**: IPv6 static LSPs are not supported at the `[edit protocols mpls static-label-switched-path lsp-name ingress]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for IPv6 addresses in static LSP configurations are provided in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Configuring the Egress Router Address for LSPs

to (Security Source NAT)

IN THIS SECTION

- Syntax | 29421
- Hierarchy Level | 29421
- Description | 29421
- Options | 29422
- Required Privilege Level | 29422
- Release Information | 29422

Syntax

```
to {
    interface [interface-name];
    routing-instance [routing-instance-name];
    zone [zone-name];
}
```

Hierarchy Level

```
[edit security nat source rule-set rule-set-name]
```

Description

Specify the destination of the packet among the routing instance, interface, or zone.

Options

- `interface [interface-name]`—Name of the interface.
- `routing-instance [routing-instance-name]`—Name of the routing instance.
- `zone [zone-name]`—Name of the zone.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

to-zone (Security IDP Policy)

IN THIS SECTION

- Syntax | 29422
- Hierarchy Level | 29423
- Description | 29423
- Options | 29423
- Required Privilege Level | 29423
- Release Information | 29423

Syntax

```
to-zone (zone-name | any);
```

Hierarchy Level

```
[edit security idp idp-policy policy-name rulebase-exempt rule rule-name match]
[edit security idp idp-policy policy-name rulebase-ips rule rule-name match]
```

Description

Specify a destination zone to be associated with the security policy. The default value is any.

Options

zone-name—Name of the destination zone object.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

to-zone (Security Policies)

IN THIS SECTION

- Syntax | 29424
- Hierarchy Level | 29426
- Description | 29426
- Options | 29426
- Required Privilege Level | 29426
- Release Information | 29426

Syntax

```
to-zone zone-name {
    policy policy-name {
        description description;
        match {
            application {
                [application];
                any;
            }
            destination-address {
                [address];
                any;
                any-ipv4;
                any-ipv6;
            }
            source-address {
                [address];
                any;
                any-ipv4;
                any-ipv6;
            }
            source-identity {
                [role-name];
                any;
                authenticated-user;
                unauthenticated-user;
                unknown-user;
            }
        }
        scheduler-name scheduler-name;
        then {
                                count {
                alarm {
                    per-minute-threshold number;
                    per-second-threshold number;
                }
            }
            deny;
            log {
                session-close;
                session-init;
```

```
            }
            permit {
                application-services {
                    application-firewall {
                        rule-set rule-set-name;
                    }
                    application-traffic-control {
                        rule-set rule-set-name;
                    }
                    gprs-gtp-profile profile-name;
                    gprs-sctp-profile profile-name;
                    idp;
                    redirect-wx | reverse-redirect-wx;
                    ssl-proxy {
                        profile-name profile-name;
                    }
                    uac-policy {
                        captive-portal captive-portal;
                    }
                    utm-policy policy-name;
                }
                destination-address {
                    drop-translated;
                    drop-untranslated;
                }
                firewall-authentication {
                    pass-through {
                        access-profile profile-name;
                        client-match user-or-group-name;
                        ssl-termination-profile profile-name;
                        web-redirect;
                        web-redirect-to-https;
                    }
                    web-authentication {
                        client-match user-or-group-name;
                    }
                }
                services-offload;
                tcp-options {
                    sequence-check-required;
                    syn-check-required;
                }
                tunnel {
```

```
                    ipsec-group-vpn group-vpn;
                    ipsec-vpn vpn-name;
                    pair-policy pair-policy;
                }
            }
            reject;
        }
    }
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name]
```

Description

Specify a destination zone to be associated with the security policy.

Options

- `zone-name`—Name of the destination zone object.
- `junos-host`—Default security zone for self-traffic of the device.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for the `services-offload` and `junos-host` options added in Junos OS Release 11.4. Support for the `source-identity` option added in Junos OS Release 12.1. Support for the `ssl-termination-profile` and `web-redirect-to-https` options added in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

to-zone (Security Policies Global)

IN THIS SECTION

- Syntax | 29427
- Hierarchy Level | 29427
- Description | 29427
- Required Privilege Level | 29428
- Release Information | 29428

Syntax

```
to-zone {
    [zone-name];
    any;
}
```

Hierarchy Level

```
[edit security policies global policy policy-name match]
```

Description

Identify a single destination zone or multiple destination zones to be used as a match criteria for a policy. You must configure specific zones or default to any zone, but you cannot have both in a global policy.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

Global Policy Overview

Example: Configuring a Global Policy with No Zone Restrictions

Example: Configuring a Global Policy with Multiple Zones

token (DHCP Local Server)

IN THIS SECTION

- Syntax | 29428
- Hierarchy Level | 29429
- Description | 29429
- Options | 29429
- Required Privilege Level | 29430
- Release Information | 29430

Syntax

```
token token-value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure],
[edit system services dhcp-local-server reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server group group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Configure a plain-text token for all DHCP clients or only the DHCP clients serviced by the specified group of interfaces. The token enables rudimentary entity authentication to protect against inadvertently instantiated DHCP servers. A null token (empty string) indicates that the configuration token functionality is not enabled. A group configuration takes precedence over a DHCP local server configuration. For more information about tokens, see RFC 3118, *Authentication for DHCP Messages*, section 4.

Options

token-value—Plain-text alphanumeric string.

- **Default:** null (empty string)

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

> *Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview*
>
> *Configuring a Token for DHCP Local Server Authentication*

too-many-requests (Security Antivirus Fallback Options)

IN THIS SECTION

- Syntax | **29430**
- Hierarchy Level | **29431**
- Description | **29431**
- Options | **29431**
- Required Privilege Level | **29431**
- Release Information | **29431**

Syntax

```
too-many-requests (block | log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name fallback-options]
```

Description

If the total number of messages received concurrently exceeds 4000, the content is either passed or blocked depending on the too-many-request fallback option. The default action is block. (The allowed request limit is not configurable.)

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

too-many-requests (Security Antivirus Fallback Options Sophos Engine)

IN THIS SECTION

- Syntax | 29432
- Hierarchy Level | 29432
- Description | 29432
- Options | 29432
- Required Privilege Level | 29433
- Release Information | 29433

Syntax

```
default (block | log-and-permit | permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name fallback-
options]
```

Description

If the total number of messages received concurrently exceeds the device limits, the content is either passed or blocked depending on the too-many-request fallback option. (The allowed request limit is not configurable.)

Options

- `block`—Log the error and deny the traffic
- `log-and-permit`—Log the error and permit the traffic
- `permit`—Permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

too-many-requests (Security Web Filtering Fallback Settings)

IN THIS SECTION

- Syntax | 29433
- Hierarchy Level | 29433
- Description | 29434
- Options | 29434
- Required Privilege Level | 29434
- Release Information | 29434

Syntax

```
too-many-requests (block | log-and-permit);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering surf-control-integrated profile profile-name
fallback-settings]
[edit security utm feature-profile web-filtering websense-redirect profile profile-name fallback-
settings]
```

```
[edit security utm feature-profile web-filtering  juniper-enhanced profile profile-name fallback-
settings]
```

Description

If the total number of messages received concurrently exceeds the device limits, the content is either passed or blocked depending on the too-many-request fallback option. The default action is BLOCK. (The allowed request limit is not configurable.)

Options

- block—Log the error and deny the traffic
- log-and-permit—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The surf-control-integrated feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The surf-control-integrated feature is not supported from Junos OS Release 15.1X49-D10 onwards.

top

IN THIS SECTION

- Syntax | 29435
- Description | 29435
- Options | 29435
- Required Privilege Level | 29435
- Release Information | 29435

Syntax

```
top <configuration-command>
```

Description

Return to the top level of configuration command mode, which is indicated by the `[edit]` banner.

Options

configuration-command (Optional) Issue configuration mode commands from the top of the hierarchy.

Required Privilege Level

configure—To enter configuration mode.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Displaying the Current Configuration

exit | 6638
up | 30436

topologies (Multitopology Routing)

IN THIS SECTION

- Syntax | 29436
- Hierarchy Level | 29436
- Description | 29437
- Options | 29437
- Required Privilege Level | 29437
- Release Information | 29437

Syntax

```
topologies {
    family (inet | inet6) {
        topology topology-name;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Description

Configure a topology for multitopology routing. Each topology creates a new routing table that is populated with direct routes from the topology.

Options

family—Configure the type of family address type.

inet—IPv4

inet6—IPv6

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths

Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Understanding Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Understanding Multitopology Routing in Conjunction with PIM

topologies (Protocols IS-IS)

IN THIS SECTION

- Syntax | 29438
- Hierarchy Level | 29438
- Description | 29438
- Default | 29439
- Options | 29439
- Required Privilege Level | 29439
- Release Information | 29439

Syntax

```
topologies {
    ipv4-multicast;
    ipv6-multicast;
    ipv6-unicast;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Configure alternate IS-IS multicast or unicast topologies.

The remaining statements are explained separately. See CLI Explorer.

> **NOTE**: The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.

Default

By default, multicast and unicast topologies are disabled.

Options

ipv4-multicast Configure alternate IPv4 multicast topologies.

ipv6-multicast Configure alternate IPv6 multicast topologies.

ipv6-unicast Configure alternate IPv6 unicast topologies. This statement causes IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to inet6.0.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

Example: Configuring IS-IS Multicast Topology

IS-IS Multicast Topologies Overview

Understanding IS-IS IPv4 and IPv6 Unicast Topologies

topology (Filter-Based Forwarding)

IN THIS SECTION

- Syntax | 29440
- Hierarchy Level | 29440
- Description | 29440
- Options | 29441
- Required Privilege Level | 29441
- Release Information | 29441

Syntax

```
topology topology-name;
```

Hierarchy Level

```
[edit firewall family (inet | inet6) filter filter-name term term-name then],
[edit firewall family (inet | inet6) filter filter-name term term-name then logical-system
logical-system-name],
[edit firewall family (inet | inet6) filter filter-name term term-name then logical-system
logical-system-name routing-instance routing-instance-name],
[edit firewall family (inet | inet6) filter filter-name term term-name then routing-instance
routing-instance-name]
```

Description

Configure a topology for filter-based forwarding for multitopology routing. The firewall filter you apply to the ingress interface is used to look up traffic against the configured topology, and, if a route matches the conditions you configure for the term, the route is accepted and added to the routing table for the specific topology.

There are multiple ways to configure a topology for filter-based forwarding, depending on the type of instance or logical system you want to specify for the forwarding class.

NOTE: The options for logical system and routing instance precede the `topology` statement with the `then` statement.

Options

topology-name—Name of a topology against which you want to match traffic.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Routing Policies, Firewall Filters, and Traffic Policers User Guide

Understanding Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

topology-id

IN THIS SECTION

- Syntax | 29442
- Hierarchy Level | 29442
- Description | 29442
- Default | 29442
- Options | 29442

- Required Privilege Level | 29442
- Release Information | 29443

Syntax

```
topology-id number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf topology name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
topology name],
[edit protocols ospf topology name],
[edit routing-instances routing-instance-name protocols ospf topology name]
```

Description

Configure a topology identifier for a topology enabled for OSPF.

Default

The default identifier for the default topology is 0, and the default identifier for the topology for IPv4 multicast traffic is 1. These identifiers are predefined and cannot be modified.

Options

number—The integer value used to identify the topology.

- **Range:** 32 through 127

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths

Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Understanding Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Understanding Multitopology Routing in Conjunction with PIM

topology (OSPF) | 29445

topology (Multitopology Routing)

IN THIS SECTION

- Syntax | 29443
- Hierarchy Level | 29444
- Description | 29444
- Options | 29444
- Required Privilege Level | 29444
- Release Information | 29444

Syntax

```
topology topology-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options topologies family (inet | inet6)],
[edit logical-systems logical-system-name routing-options topologies family (inet | inet6)],
[edit routing-instances routing-instance-name routing-options topologies family (inet | inet6)],
[edit routing-options topologies family (inet | inet6)]
```

Description

Configure the name of a topology configured to run multitopology routing.

Options

topology-name—Name of the topology. Include a string value that describes the type of traffic, such as voice or video. For IPv4 multicast traffic, include **ipv4-multicast** as the name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths

Understanding Multitopology Routing in Conjunction with PIM

Understanding Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

topology (OSPF)

IN THIS SECTION

- Syntax | 29445
- Hierarchy Level | 29445
- Description | 29445
- Options | 29446
- Required Privilege Level | 29446
- Release Information | 29446

Syntax

```
topology (default | ipv4-multicast | name) {
    spf-options {
        delay milliseconds;
        holddown milliseconds;
        rapid-runs number;
    }
    topology-id number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf],
[edit protocols ospf],
[edit routing-instances routing-instance-name protocols ospf]
```

Description

Enable a topology for OSPF multitopology routing. You must first configure one or more topologies under the `[edit routing-options]` hierarchy level.

Options

default—Name of the default topology. This topology is automatically created, and all routes that correspond to it are automatically added to the **inet.0** routing table. You can modify certain default parameters, such as for the SPF algorithm.

ipv4-multicast—Name of the topology for IPv4 multicast traffic.

name—Name of a topology you configured at the `[edit routing-options]` hierarchy level to create a topology for a specific type of traffic, such as voice or video.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths
- Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic
- Understanding Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic
- Understanding Multitopology Routing in Conjunction with PIM

topology (OSPF Interface)

IN THIS SECTION

- Syntax | 29447
- Hierarchy Level | 29447

- Description | 29447
- Default | 29448
- Options | 29448
- Required Privilege Level | 29448
- Release Information | 29448

Syntax

```
topology (ipv4-multicast | name) {
    metric metric;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area area-id interface interface-name],
[edit protocols ospf area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols ospf area area-id interface interface-
name]
```

Description

Configure interface-specific properties for multitopology OSPF, including topology-specific metric values for an interface.

All OSPF interfaces have a cost, which is a routing metric that is used in the link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. The default value for the OSPF metric for an interface is 1. You can modify the default value for an OSPF interface and configure a topology-specific metric for that interface. The topology-specific metric applies to routes advertised from the interface that belong only to that topology.

Default

The default value of the topology metric is the same as the default metric value calculated by OSPF or the value configured for the OSPF metric.

Options

ipv4-multicast—Name of the topology for IPv4 multicast traffic.

name—Name of a topology created under the `[edit routing-options]` hierarchy level.

metric *metric*—Cost of a route from an OSPF interface. You can specify a metric value for a topology that is different from the value specified for the interface.

- **Range:** 1 through 65,535
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic
- Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths
- Understanding Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic
- Understanding Multitopology Routing in Conjunction with PIM

topology (Protocols BGP)

IN THIS SECTION

- Syntax | 29449
- Hierarchy Level | 29449
- Description | 29450
- Options | 29450
- Required Privilege Level | 29450
- Release Information | 29451

Syntax

```
topology name {
    community {
        target identifier;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family (inet | inet6) unicast],
[edit logical-systems logical-system-name protocols bgp group group-name family (inet | inet6)
unicast],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family
(inet | inet6) unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family (inet | inet6) unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family (inet | inet6) unicast],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family (inet | inet6) unicast],
[edit protocols bgp family (inet | inet6) unicast],
[edit protocols bgp group group-name family (inet | inet6) unicast],
[edit protocols bgp group group-name neighbor address family (inet | inet6) unicast],
```

```
[edit routing-instances routing-instance-name protocols bgp family (inet | inet6) unicast],
[edit routing-instances routing-instance-name protocols bgp group group-name family (inet |
inet6) unicast],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family (inet | inet6)]
```

Description

Enable a topology for BGP multitopology routing. You must first configure one or more topologies under the `[edit routing-options]` hierarchy level.

Apply the community tags to identify the application topologies by configuring a routing topology name and BGP community value.

In Junos OS, multitopology support for BGP is based on the community value in a BGP route. This configuration determines the association between a topology and one or more community values and populates the topology routing tables. Arriving BGP updates that have a matching community value are replicated in the associated topology routing table. You decide which BGP community values are associated with a given topology.

For example, you can create a configuration that causes BGP updates that are received with community value `target:40:40` to be added into topology routing table `voice.inet.0` (in addition to the default routing table `inet.0`). Likewise, you configuration can specify that updates that are received with community value `target:50:50` are added into topology routing table `video.inet.0` (in addition to the default routing table `inet.0`).

Options

name Name of a topology you configured at the `[edit routing-options]` hierarchy level to create a topology for a specific type of traffic, such as voice or video.

`community` Configure the community to identify the multitopology routes. BGP uses the target community identifier to install the routes it learns in the appropriate multitopology routing tables.

- **Syntax:** target *identifier*—Configure the destination to which the route is going.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths

Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Understanding Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic

Understanding Multitopology Routing in Conjunction with PIM

tos-reflect (L2TP)

IN THIS SECTION

- Syntax | 29451
- Hierarchy Level | 29451
- Description | 29452
- Required Privilege Level | 29452
- Release Information | 29452

Syntax

```
tos-reflect;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Configure the LNS to reflect the IP ToS value in the inner IP header to the outer IP header. When CoS rewrite rules are also configured ([`edit class-of-service interfaces` *`interface-name`* `unit` *`logical-unit-number`* `rewrite-rules`]), the rewrite is performed only on the inner IP ToS; the outer IP TOS value is not affected.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Dynamic CoS for an L2TP LNS Inline Service

total-memory

IN THIS SECTION

- Syntax | 29453
- Hierarchy Level | 29453
- Description | 29453
- Options | 29453
- Required Privilege Level | 29453
- Release Information | 29453

Syntax

```
total-memory <percentage>;
per-logical-domain <number>
```

Hierarchy Level

```
[edit security idp sensor-configuration packet-log]
```

Description

Configure the maximum amount of memory to be allocated to packet capture for the device. This value is expressed as a percentage of the memory available on the device. The total memory for a device will differ depending on its operating mode.

Options

- *percentage*—Amount of packet capture memory expressed as a percentage of total memory for the device mode.
- **Range:** 1 to 100 percent
- **Default:** 10

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

trace

IN THIS SECTION

- Syntax | 29454
- Hierarchy Level | 29454
- Description | 29455
- Options | 29455
- Required Privilege Level | 29456
- Release Information | 29456

Syntax

```
trace application {
    application-name {
        node {
            node-name {
                level level;
                group group-name {
                    enabled (off | on);
                    tracepoint tracepoint-name {
                        enabled (off | on);
                    }
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Set system-wide tracing options for Junos OS Evolved. Tracing is a process about getting information regarding a program's execution. A programmer adds specific code to a program, called a *tracepoint*, which when enabled triggers a log event, which gives the programmer visibility into how the program is executing.

Options

application-name Name of the application for which tracing should be applied. You can see a list of the applications for Junos OS Evolved by issuing the `show system applications` command. If you want to enable tracing for all applications, use the value `all`.

enabled (on | off) Enter `on` to enable or `off` to to disable. This option has two locations in the hierarchy, one for enabling or disabling an entire group, and one for enabling or disabling specific tracepoints.

group *group-name* The name of the trace group for which the configuration should be applied.

level *level* Trace level of the application. The following levels are available, listed in the order of highest severity first:

- emergency—Trace emergency messages (system is unusable).
- alert—Trace alert messages (take immediate action).
- critical—Trace critical conditions.
- error—Trace error conditions.
- warning—Trace warning conditions.
- notice—Trace normal but significant conditions.
- info—Trace informational messages.
- debug—Trace debug messages.

You can configure the application’s level here. The program determines the level of the tracepoints. If the application's level is at least as inclusive as (that is, the same as or of higher severity) the tracepoint's level, then the trace event is able to fire. If the application's level is less inclusive (of lower severity) than the tracepoint's level, then the trace event is not able to fire.

The default trace level is `info`.

node *node-name* Name of a node. Node names can be re0, re1, fpc0, fpc1, and so on, depending on the system.

tracepoint *tracepoint-name* Name of a tracepoint. A tracepoint is an event in the program, which can give the developer visibility into how the program is executing.

Required Privilege Level

admin

Release Information

Statement introduced in Junos OS Evolved 18.2R1.

trace (DHCP Local Server)

IN THIS SECTION

- Syntax | 29456
- Hierarchy Level | 29456
- Description | 29457
- Required Privilege Level | 29457
- Release Information | 29457

Syntax

```
trace;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name interface interface-name],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name interface interface-name],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name interface interface-name],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name interface interface-name],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group
group-name interface interface-name],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
interface interface-name],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit system services dhcp-local-server group group-name interface interface-name]
```

Description

Enable trace operations for a group of interfaces or for a specific interface within a group.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Tracing Extended DHCP Operations

Tracing Extended DHCP Operations for Specific Interfaces

trace (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 29458
- Hierarchy Level | 29458
- Description | 29459
- Required Privilege Level | 29459
- Release Information | 29459

Syntax

```
trace;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay group group-name  interface-name],
[edit  logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
interface interface-name],
[edit  logical-systems logical-system-name forwarding-options dhcp-relay group group-name
interface interface-name],
[edit  logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit  logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name interface interface-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name interface interface-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name]
```

Description

Enable trace operations for a group of interfaces or for a specific interface within a group. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

EX Series switches do not support DHCPv6.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring DHCP Relay Agent

DHCP Monitoring and Management

trace (EVPN-VXLAN)

IN THIS SECTION

- Syntax | 29460
- Hierarchy Level | 29460
- Description | 29460
- Options | 29461
- Required Privilege Level | 29461
- Release Information | 29461

Syntax

```
trace {
evpn-vxlan;
}
```

Hierarchy Level

```
[edit services]
```

Description

Enable EVPN-VXLAN tracing in multiple modules and hierarchies. The `trace` command invokes a built-in commit script to generate the tracing configurations required for troubleshooting EVPN-VXLAN. The commit script for this feature, `services_evpn_commit_script.py`, is included with supported platforms.

> **NOTE**: Configure EVPN-VXLAN tracing and enable the commit-script using the following commands.
>
> ```
> set services trace evpn-vxlan
> set system scripts commit allow-transients
> set system scripts commit file services_evpn_commit_script.py
> ```

See *Commit Script Overview* and *Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts* for more information about commit scripts.

When you commit the configuration, the commit script generates the following EVPN-VXLAN tracing configurations.

RPD/routing-instances tracing for each routing-instance configured.

```
set routing-instances routing-instance-name protocols evpn traceoptions file routing-instance-name.log
set routing-instances routing-instance-name protocols evpn traceoptions file size 50m
```

```
set routing-instances routing-instance-name protocols evpn traceoptions file files 50
set routing-instances routing-instance-name protocols evpn traceoptions flag all detail
```

BGP tracing for each BGP group configured.

```
set protocols bgp group bgp-group-name traceoptions file bgp-group-name.log
set protocols bgp group bgp-group-name traceoptions file size 50m
set protocols bgp group bgp-group-name traceoptions file files 50
set protocols bgp group bgp-group-name traceoptions flag all
```

L2ALD tracing for protocols l2-learning.

```
set protocols l2-learning traceoptions file l2ald.log
set protocols l2-learning traceoptions file size 50m
set protocols l2-learning traceoptions file files 50
set protocols l2-learning traceoptions level all
set protocols l2-learning traceoptions flag all
```

DHCP tracing when DHCP is configured.

```
set system processes dhcp-service traceoptions file dhcp_service.log
set system processes dhcp-service traceoptions file size 50m
set system processes dhcp-service traceoptions file files 50
set system processes dhcp-service traceoptions level all
set system processes dhcp-service traceoptions flag all
```

Options

evpn-vxlan Enable evpn-vxlan tracing

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1 and Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

show | display commit-scripts | 19527

Remove a Persistent or Transient Configuration Change Using Commit Scripts

traceoptions (routing-options)

IN THIS SECTION

- Syntax | 29462
- Hierarchy Level | 29462
- Description | 29463
- Default | 29463
- Options | 29463
- Required Privilege Level | 29465
- Release Information | 29465

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast],
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options],
```

```
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options],
[edit routing-options flow],
[edit routing-options multicast]
```

Description

Define tracing operations that track all routing protocol functionality in the routing device. You can also trace clock synchronization information using this command.

To specify more than one tracing operation, include multiple `flag` statements.

> **NOTE**: On Junos OS Evolved, `traceoptions` is disabled for op, event, and commit scripts. Instead, Junos OS Evolved enables default tracking and trace messages that are logged under `/var/log/traces`.

Default

If you do not include this statement, no global tracing operations are performed.

Options

Values:

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place global routing protocol tracing output in the file **routing-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 10 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization**—Nonstop active routing synchronization
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.

nsr-synchronization and **nsr-packet** flags for BFD sessions added in Junos OS Release 8.5.

nsr-synchronization flag for RIP and RIPng added in Junos OS Release 9.0.

nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.

nsr-synchronization flag for PIM added in Junos OS Release 9.3.

nsr-synchronization flag for MPLS added in Junos OS Release 10.1.

nsr-synchronization flag for MSDP added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Tracing Global Routing Protocol Operations

traceoptions (Ethernet Switching)

IN THIS SECTION

- Ethernet Switching Options | **29466**
- Ethernet Ring Protection | **29466**
- Edge Virtual Bridging | **29466**
- Hierarchy Level | **29467**
- Description | **29467**
- Default | **29467**
- Options | **29467**

Required Privilege Level | 29471
Release Information | 29471

Ethernet Switching Options

```
traceoptions {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
    flag flag <disable>;
}
```

Ethernet Ring Protection

```
traceoptions {
    file filename <no-stamp> <world-readable | no-world-readable> <replace> <size size>;
    flag flag;
}
```

Edge Virtual Bridging

```
traceoptions {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit ethernet-switching-options]
```

```
[edit protocols protection-group]
```

```
[edit protocols edge-virtual-bridging]
```

Description

Define global tracing operations for access security features on Ethernet switches.

Configure trace options for the protection group.

Define global tracing operations for edge virtual bridging (EVB) features on Ethernet switches.

Default

The Ethernet Switching Options **traceoptions** feature is disabled by default.

Edge Virtual Bridging tracing operations are disabled by default.

Ethernet Ring Protection trace options are not set by default. On some EX Series switches, logging of basic ERPS state transitions is set by default. You can configure trace options on those switches to obtain more details than are provided by the default log. See *Understanding Ethernet Ring Protection Switching Functionality* for additional information about default logging of the basic state transitions.

Options

For Ethernet Switching Options:

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename*** —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files ***number*** —(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached (***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes), at which point the

oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000
- **Default:** 3 files

flag ***flag*** —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **access-security**—Trace access security events.
- **all**—All tracing operations.
- **analyzer**—Trace analyzer events.
- **config-internal**—Trace internal configuration operations.
- **filter**—Trace filter transaction events.
- **forwarding-database**—Trace forwarding database events.
- **general**—Trace general events.
- **interface**—Trace interface events.
- **krt**—Trace communications over routing sockets.
- **lib**—Trace library calls.
- **nexthop**—Trace next-hop events.
- **normal**—Trace normal events.
- **parse**—Trace reading of the configuration.
- **regex-parse**—Trace regular-expression parsing operations.
- **rtg**—Trace redundant trunk group events.
- **state**—Trace state transitions.
- **stp**—Trace spanning-tree events.
- **task**—Trace Ethernet-switching task processing.
- **timer**—Trace Ethernet-switching timer processing.
- **unknown-unicast-forwarding**—Trace unknown unicast forwarding events.
- **vlan**—Trace VLAN events.

no-stamp—(Optional) Do not timestamp the trace file.

- **Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Restrict file access to the user who created the file.

replace—(Optional) Replace an existing trace file if there is one rather than appending to it.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size*** —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes

- **Range:** 10 KB through 1 gigabyte

- **Default:** 128 KB

world-readable—(Optional) Enable unrestricted file access.

For Ethernet Ring Protection:

file ***filename***—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**. You can include the following file options:

- **no-stamp**—(Optional) Do not timestamp trace file.

- **no-world-readable**—(Optional) Do not allow any user to read the log file.

- **replace**—(Optional) Replace the trace file rather than appending to it.

- **size**—(Optional) Maximum trace file size (10240..4294967295).

- **world-readable**—(Optional) Allow any user to read the log file.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. The following flags are available for both EX Switches and MX Series routers:

- **all**—Trace all

- **config** —Trace configuration messages

- **debug**—Trace debug messages

- **events**—Trace events to the protocol state machine

- **normal**—Trace normal messages

- **pdu** —Trace RAPS PDU reception and transmission
- **periodic-packet-management**—Trace periodic packet management state and events
- **state-machine**—Trace RAPS state machine
- **timers**—Trace protocol timers

For Edge Virtual Bridging:

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached (***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000
- **Default:** 3 files

no-world-readable—(Optional) Restrict file access to the user who created the file.

replace—(Optional) Replace an existing trace file if there is one rather than appending output to it.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes
- **Range:** 10 KB through 1 gigabyte
- **Default:** 128 KB

world-readable—(Optional) Enable unrestricted file access.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—Trace everything.
- **ecp**—Trace Edge Control Protocol (ECP) events.
- **evb-tlv**— Trace EVB type, length, and value (TLV) events.

- **parse**—Trace configuration parsing.
- **policy**—Trace policy events.
- **vdp**—Trace Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) events.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

> **NOTE**: The `traceoptions` statement is not supported on the QFX3000 QFabric system.

RELATED DOCUMENTATION

- Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Edge Virtual Bridging*

traceoptions (High Availability)

IN THIS SECTION

- Syntax | 29472
- Hierarchy Level | 29472
- Description | 29472

- Options | 29472
- Required Privilege Level | 29474
- Release Information | 29474

Syntax

```
traceoptions {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
    flag name;
    level (alert | all | critical | debug | emergency | error | info | notice | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit chassis high-availability]
```

Description

Set Multinode High Availability traceoptions

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks

file *name*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place global routing protocol tracing output in the file **routing-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000 files
- **Default:** 2 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

no-world-readable—Restrict users from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

  If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

level—Set a level of tracing.

- Values:
  - alert—Match alert conditions
  - all—Match all levels
  - critical—Match critical conditions
  - debug—Match debug messages
  - emergency—Match emergency conditions
  - error—Match error conditions
  - info—Match informational messages
  - notice—Match notice level messages
  - warning—Match warning messages

no-remote-trace—Disable remote tracing.

Required Privilege Level

trace

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Multinode High Availability

traceoptions (802.1X)

IN THIS SECTION

- Syntax | 29474
- Hierarchy Level | 29475
- Description | 29475
- Default | 29475
- Options | 29475
- Required Privilege Level | 29476
- Release Information | 29476

Syntax

```
traceoptions  {
    file filename <files number> <size size> <world-readable | no-world-readable> <match regex>;
    flag flag;
}
```

Hierarchy Level

```
[edit protocols dot1x]
```

Description

Define tracing operations for the 802.1X protocol.

Default

Tracing operations are disabled.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size by using the **size** option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- `all`—All tracing operations.
- `config-internal`—Trace internal configuration operations.
- `dot1x-event`—(Switches with ELS only) Trace 802.1x events.
- `dot1x-debug`—(Switches without ELS) Trace 802.1x events.
- `dot1x-ipc`—(Switches with ELS only) Trace IPC interactions.
- `eapol`—Trace EAPOL packets transmitted and received.
- `esw-if`—(Switches without ELS) Trace ESW interactions.
- `general`—Trace general operations.

- `normal`—Trace normal operations.
- `parse`—Trace reading of the configuration.
- `regex-parse`—Trace regular-expression parsing operations.
- `state`—Trace protocol state changes.
- `task`—Trace protocol task operations.
- `timer`—Trace protocol timer operations.
- `vlan`—Trace VLAN transactions.

match *regex*—(Optional) Refine the output to include lines that contain the regular expression.

no-world-readable—(Optional) Restrict file access to the user who created the file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files with the **files** option, you also must specify a maximum file size.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for the `dot1x-event` and `dot1x-ipc` options introduced in Junos OS Release 13.2X50 for EX Series switches.

RELATED DOCUMENTATION

show lldp | 21942

Configuring 802.1X Interface Settings (CLI Procedure)

802.1X for Switches Overview

traceoptions (Access)

IN THIS SECTION

- Syntax | 29477
- Hierarchy Level | 29477
- Description | 29478
- Options | 29478
- Required Privilege Level | 29479
- Release Information | 29479

Syntax

```
traceoptions {
    file filename {
        files number;
        match regular-expression;
        size maximum-file-size;
        <world-readable | no-world-readable>;
    }
    flag flag;
}
```

Hierarchy Level

```
[edit access firewall-authentication]
```

Description

Define Routing Engine firewall authentication tracing options.

Options

- `file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.
- `files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file.0`*, then *`trace-file.1`*, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.
- If you specify a maximum number of files, you also must specify a maximum file size with the size option and a filename.
- **Range:** 2 through 1000 files
- **Default:** 10 files
- `match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.
- `size`*`maximum-file-size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file.0`*. When the `trace-file` again reaches its maximum size, *`trace-file.0`* is renamed *`trace-file.1`* and *`trace-file`* is renamed *`trace-file.0`*. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.
- If you specify a maximum file size, you also must specify a maximum number of trace files with the files option and filename.
- **Syntax:** *`x`* `k` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB
- `world-readable | no-world-readable`—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.
- `flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags.
- `all`—All tracing operations

- `authentication`—Trace authentication events
- `configuration`—Trace configuration events
- `setup`—Trace setup of firewall authentication service

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

Firewall User Authentication Overview

traceoptions (Access Port Security)

IN THIS SECTION

- Syntax | 29480
- Hierarchy Level | 29480
- Description | 29480
- Default | 29480
- Options | 29480
- Required Privilege Level | 29482
- Release Information | 29482

Syntax

```
traceoptions {
    file (file-name | files files | match match | no-world-readable | size size | world-
readable);
    flag ( all | asynch | chassis-scheduler | cos-adjustment | dynamic | hardware-database |
init | parse | performance-monitor | process | restart | route-socket | show | snmp | util);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit ethernet-switching-options],
[edit class-of-service]
```

Description

Define global tracing operations for access security features on Ethernet switches.

Default

The **traceoptions** feature is disabled by default.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename*** —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached (***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000
- **Default:** 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **access-security**—Trace access security events.
- **all**—All tracing operations.
- **config-internals**—Trace internal configuration operations.
- **forwarding-database**—Trace forwarding database and next-hop events.
- **general**—Trace general events.
- **interface**—Trace interface events.
- **ip-source-guard**—Trace IP source guard events.
- **krt**—Trace communications over routing sockets.
- **lib**—Trace library calls.
- **normal**—Trace normal events.
- **parse**—Trace reading of the configuration.
- **regex-parse**—Trace regular-expression parsing operations.
- **rtg**—Trace redundant trunk group events.
- **state**—Trace state transitions.
- **stp**—Trace spanning-tree events.
- **task**—Trace Ethernet-switching task processing.
- **timer**—Trace Ethernet-switching timer processing.
- **vlan**—Trace VLAN events.

no-stamp—(Optional) Do not timestamp the trace file.

- **Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Restrict file access to the user who created the file.

replace—(Optional) Replace an existing trace file if there is one rather than appending to it.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes
- **Range:** 10 KB through 1 gigabyte
- **Default:** 128 KB

world-readable—(Optional) Enable unrestricted file access.

no-remote-trace—(Optional) Disable remote tracing.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

- Port Security Features
- Interfaces Overview for Switches
- Understanding IP Source Guard for Port Security on Switches
- Understanding Redundant Trunk Links (Legacy RTG Configuration)
- Understanding STP
- Understanding Bridging and VLANs on Switches

traceoptions (Active Directory Access)

IN THIS SECTION

- Syntax | 29483
- Hierarchy Level | 29484
- Description | 29484
- Options | 29484
- Required Privilege Level | 29485
- Release Information | 29485

Syntax

```
traceoptions {
    file filename ;
    flag {
        active-directory-authentication;
        all;
        configuration;
        db;
        ip-user-mapping;
        ip-user-probe;
        ipc;
        user-group-mapping;
        wmic;
    }
    level {
        all
        error
        info
        notice
        verbose
        warning
    }
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services user-identification active-directory-access]
```

Description

Define Active Directory trace options for the integrated user firewall feature.

Options

file *filename*
Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

flag
Trace the operation or operations to perform on the integrated user firewall. To specify more than one trace operation, include multiple flag statements.

- **active-directory-authentication** — Trace the building of and modifications to the Active Directory authentication table.
- **all** — Trace everything.
- **configuration** — Trace configuration events.
- **db** — Trace the database.
- **ip-user-mapping** — Trace the ip-user-mapping module.
- **ip-user-probe** — Trace PC client probing.
- **ipc** — Trace communication events with the Packet Forwarding Engine.
- **user-group-mapping** — Trace the process of getting user-to-group-mapping.
- **wmic** — Trace the Windows Management Instrumentation Client process.

level
Level of trace operation to perform.

- **all** — Match all levels.
- **error** — Match error conditions.

info	Match informational messages.
notice	Match conditions that should be handled specially.
verbose	Match verbose messages.
warning	Match warning messages.

no-remote-trace Disallow tracing from a remote device.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

active-directory-access | **1802**

user-identification (Services) | **30560**

Overview of Integrated User Firewall

traceoptions (Address-Assignment Pool)

IN THIS SECTION

- Syntax | **29486**
- Hierarchy Level | **29486**
- Description | **29486**
- Options | **29486**
- Required Privilege Level | **29487**

Release Information | 29488

Syntax

```
traceoptions {
    file filename {
        files number;
        size maximum-file-size;
        match regex;
        (world-readable | no-world-readable);
    }
    flag address-assignment;
    flag all;
    flag configuration;
    flag framework;
    flag ldap;
    flag local-authentication;
    flag radius;
}
```

Hierarchy Level

```
[edit system processes general-authentication-service]
```

Description

Configure tracing options.

Options

file *filename*—Name of the file that receives the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `address-assignment`—All address-assignment events
- `all`—All tracing operations
- `configuration`—Configuration events
- `framework`—Authentication framework events
- `ldap`—LDAP authentication events
- `local-authentication`—Local authentication events
- `radius`—RADIUS authentication events

`match` *`regex`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-world-readable`—(Optional) Restrict access to the originator of the trace operation only.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and filename.

- **Syntax:** `xk` to specify KB, `xm` to specify MB, or `xg` to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Flag for tracing address-assignment pool operations introduced in Junos OS Release 9.0.

`option-name` option introduced in Junos OS Release 8.3.

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Address-Assignment Pool Configuration Overview

traceoptions (advanced policy-based routing)

IN THIS SECTION

- Syntax | 29488
- Hierarchy Level | 29489
- Description | 29489
- Options | 29489
- Required Privilege Level | 29490
- Release Information | 29490

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
```

```
        no-remote-trace;
    }
```

Hierarchy Level

[edit security advance-policy-based-routing]

Description

Configure tracing operations for advanced policy-based routing.

Options

- `file`—Configure the trace file options.
  - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *`number`*—Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
  - `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

    Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

    Range: 10 KB through 1 GB

Default: 128 KB

- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace with all flags enabled
  - `compilation`—Trace rule set compilation events
  - `configuration`—Trace configuration events
  - `ipc`—Trace process inter communication events
  - `lookup`—Trace rule set lookup events
- `no-remote-trace`—Set remote tracing as disabled.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Example: Configuring Advanced Policy-Based Routing for Application-Aware Traffic Management Solution

Understanding Advanced Policy-Based Routing

traceoptions (Analytics)

IN THIS SECTION

- Syntax | 29491
- Hierarchy Level | 29491
- Description | 29491
- Options | 29491
- Required Privilege Level | 29492
- Release Information | 29492

Syntax

```
traceoptions {
    file filename;
    files number-of-files;
    size size;
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Configure traceoptions for the network analytics daemon (analyticsd) running on the Routing Engine.

Options

file *filename*

Specify a filename for storing the traceoptions data. The file is stored in the **/var/log/** directory of your device.

If you do not specify a filename, the data is not stored in a file.

files *number-of-files*

Specify the number of files to store locally. After the number files witsh the maximum file size is reached, the system starts over and writes the data to the first file.

- **Range:** 2 to 1,000 files.

size *size*

Configure the file size in megabytes (MB).

- **Syntax:** `xm` to specify MB.
- **Range:** 10 to 4095 MB

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

traceoptions (Analytics Agent)

IN THIS SECTION

- Syntax | 29493
- Hierarchy Level | 29493
- Description | 29493
- Options | 29493
- Required Privilege Level | 29494
- Release Information | 29494

Syntax

```
traceoptions {
    file filename;
    flag (debug | error | info | trace);
}
```

Hierarchy Level

```
[edit services analytics agent]
```

Description

Configure tracing operations for Network Telemetry Framework (NTF) agent. You can specify the name of the file where the NTF agent log messages are stored. You can also specify a severity level for messages to be logged. The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify `trace`. As levels become more restrictive, fewer messages are logged.

> **NOTE**: Although the syntax uses the keyword `flag`, its function in this statement corresponds to the `level` keyword used for other `traceoptions` statements.

Options

file *filename*

Name of the file to receive the output of the tracing operation. The file is stored in the **/var/log/** directory of your device.

- **Default:** ntf-agent

flag (debug | error | info | trace)

Specify the severity level for messages to be logged. The order of severity, from most to least severe is as follows:

error > info > debug > trace

- `debug`—Match debug messages.
- `error`—Match error messages. This is the most restrictive level.

- `info`—Match informational messages.
- `trace`—Match all messages.
- **Default:** error

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

IPFIX Mediation on the BNG

Configuring NTF Agent

traceoptions (ANCP)

IN THIS SECTION

- Syntax | 29495
- Hierarchy Level | 29495
- Description | 29495
- Options | 29495
- Required Privilege Level | 29497
- Release Information | 29497

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag <disable>;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols ancp]
```

Description

Define tracing operations for ANCP agent processes.

Options

`file` *`filename`*— Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. Include the `disable` option after a flag to disable tracing for that flag. You can include the following flags:

- `all`—Trace all operations.
- `config`—Trace configuration events.
- `cos`—Trace class-of-service events.
- `general`—Trace general flow.

- `packet`—Trace ANCP packet transmit and receive operations.
- `process`—Trace process internals.
- `protocol`—Trace protocol events.
- `restart`—Trace process restart flow
- `routing-socket`—Trace routing socket events.
- `session`—Trace connection events and flow.
- `startup`—Trace ANCP startup events and flow.
- `subscriber`—Trace subscriber events.
- `timer`—Trace timer processing.

`level`—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions reqcliring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Tracing ANCP Events for Troubleshooting

traceoptions (App Engine Virtual Machine Management Service)

IN THIS SECTION

- Syntax | 29497
- Hierarchy Level | 29498
- Description | 29498
- Default | 29498
- Options | 29498
- Required Privilege Level | 29498
- Release Information | 29498

Syntax

```
traceoptions {
    file filename <files number> | match | no-world-readable |<size size> <world-readable >;
    flag (all |ccif | configuration | heartbeat | init | miscellaneous | platform | pxe |
routing-instances | snmp);
```

```
    level (all | error| info | notice | verbose | warning)
}
```

Hierarchy Level

```
[edit system processes app-engine-virtual-machine-management-service]
```

Description

Enable traceoptions for the app-engine virtual machine management service system process.

Default

Traceoptions are not enabled.

Options

file	Trace file information.
flag	Perform defined tracing operation.
level	Set traceoptions level.
no-remote-trace	Disable remote tracing.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system-control

Release Information

Statement introduced in Junos OS Release 15.1F3.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

processes

level (App Engine Virtual Machine Management Service) | 9719

flag (App Engine Virtual Machine Management Service) | 7204

file (App Engine Virtual Machine Management Service) | 7030

traceoptions (Batch Commits)

IN THIS SECTION

- Syntax | 29499
- Hierarchy Level | 29499
- Description | 29500
- Options | 29500
- Required Privilege Level | 29500
- Release Information | 29500

Syntax

```
traceoptions {
    file filename;
    files number;
    flag (all | batch | commit-server | configuration);
    size maximum-file-size;
    (world-readable | no-world-readable);
}
```

Hierarchy Level

```
[edit system commit server],
[edit system commit synchronize server]
```

Description

For software batch commits, configure tracing operations.

Options

file *name*

Name of the file to receive the output of the tracing operation.

> **NOTE**: If you configure traceoptions and do not explicitly specify a filename for logging the events, the batch commit events are logged in the `commitd` file (var/log/commitd) by default.

files *number*

Maximum number of trace files.

flag *flag*

Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—All tracing operations flags.
- `batch`—Tracing operations for batch events.
- `commit-server`—Tracing operations for commit server events.
- `configuration`—Tracing operations for the reading of configuration.

size

Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

world-readable | no-world-readable

readable—Grant all users permission to read archived log files, or restrict the permission only to the root user and users who have the Junos OS maintenance permission.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Batch Commit Server Properties

traceoptions (Captive Portal Content Delivery)

IN THIS SECTION

- Syntax | 29501
- Hierarchy Level | 29501
- Description | 29501
- Options | 29502
- Required Privilege Level | 29503
- Release Information | 29503

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services captive-portal-content-delivery]
```

Description

Define tracing operations for captive-portal-content-delivery processes.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. Ensure that filenames are unique for each logical system or routing instance in which Mobile IP is configured.

> **NOTE**: Global messages (common to all logical systems and routing instances) are always saved in `/var/log/mipd`. Messages that are specific to a logical system or routing instance are never saved in `/var/log/mipd`. If you do not configure a trace filename for a logical system or routing instance, then nothing is traced for that entity.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** `xk` to specify KB, `xm` to specify MB, or `xg` to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.
- `clicommand`—Trace CLI command operations.
- `configuration`—Trace home agent state machine operations.
- `general`—Trace general operations.
- `gres`—Trace graceful routing switchover operations.
- `ipc`—Trace Inter-Process Communication (IPC) messages between the PIC and the Routing Engine.
- `rtsock`—Trace routing socket operations.

- `rules`—Trace rules operations.
- `ssets`—Trace service sets operations.
- `statistics`—Trace statistics operations.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

HTTP Redirect Service Overview

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

Configuring Routing Engine-Based, Static HTTP Redirect Services

Configuring Routing Engine-Based, Converged HTTP Redirect Services

traceoptions (CASB)

IN THIS SECTION

- Syntax | 29504
- Hierarchy Level | 29504
- Description | 29504
- Options | 29504
- Required Privilege Level | 29505
- Release Information | 29506

Syntax

```
    traceoptions {
        file (file-name | files files | match match | no-world-readable | size size | world-
readable);
        flag ( all | flow | memory | pfeman);
        level;
        no-remote-trace;
}
```

Hierarchy Level

```
[edit security casb]
```

Description

Configure Cloud Access Security Broker (CASB) traceoptions.

Options

file Configure the trace file options.

- `filename`—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.

- `files number`—Maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, it is renamed to `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

  If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

  Range: 2 through 1000 files

  Default: 10 files

- `match regular-expression`—Refine the output to include lines that contain the regular expression.
- `size maximum-file-size`—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named `trace-file` reaches this size, it is renamed `trace-file.0`. When the `trace-file` again reaches its maximum size, `trace-file.0` is renamed `trace-file.1` and `trace-file` is renamed `trace-file.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

  Syntax: `x K` to specify KB, `x m` to specify MB, or `x g` to specify GB

  Range: 10 KB through 1 GB

  Default: 128 KB
- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option

flag Set the level of debugging the output option.

- all—Trace with all flags enabled.
- flow—Trace flow information.
- memory—Trace memory usage.
- pfeman—Trace PFE information.

level Set the level of debugging.

no-remote-trace Set remote tracing as disabled.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

traceoptions (Chassis Cluster)

IN THIS SECTION

- Syntax | 29506
- Hierarchy Level | 29507
- Description | 29507
- Options | 29507
- Required Privilege Level | 29508
- Release Information | 29508

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        (world-readable | no-world-readable);
        size maximum-file-size;
    }
    flag flag;
    level {
        (alert | all | critical | debug | emergency | error | info | notice | warning);
    }
    no-remote-trace;
}
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Define chassis cluster redundancy process tracing operations.

A minimum traceoptions configuration must include both a target file and a flag. The target file determines where the trace output is recorded. The flag defines what type of data is collected.

Options

- `file` *`filename`* —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.
- `files` *`number`* —(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`* `.0`, then *`trace-file.1`* , and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.
- If you specify a maximum number of files, you also must specify a maximum file size with the size option and a filename.
- **Range:** 2 through 1000 files
- **Default:** 10 files
- `match` *`regular-expression`* —(Optional) Refine the output to include lines that contain the regular expression.
- `size` *`maximum-file-size`* —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`* `.0`. When the `trace-file` again reaches its maximum size, *`trace-file`* `.0` is renamed *`trace-file`* `.1` and *`trace-file`* is renamed *`trace-file`* `.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.
- If you specify a maximum file size, you also must specify a maximum number of trace files with the files option and filename.
- **Syntax:** *`x`* `k` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB
- **Range:** 0 KB through 1 GB
- **Default:** 128 KB

- `world-readable | no-world-readable`—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.
- `flag`—Trace operation or operations to perform on chassis cluster redundancy processes. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace all the events
    - `configuration`—Trace configuration events
    - `routing-socket`—Trace logging of rtsock activity
    - `snmp`—Trace SNMP events

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.5.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

traceoptions (Class of Service)

IN THIS SECTION

Syntax | 29509
Hierarchy Level | 29509
Description | 29509
Default | 29509
Options | 29509

Required Privilege Level | 29511
Release Information | 29511

Syntax

```
traceoptions {
    file filename <size size> <files number>
    <world-readable | no-world-readable>;
    flag flag <flag-modifier>;
    no-remote-trace
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Set class-of-service (CoS) tracing options.

NOTE: The `traceoptions` statement is not supported on the QFabric system.

Default

Traceoptions is disabled.

Options

file *filename*
Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the `/var/log/` directory.

files *number*
(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0. The traceoption output continues in a second

trace file named *trace-file*.1. When *trace-file*.1 reaches its maximum size, output continues in a third file named *trace-file*.2, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the size option.

- **Range:** 2 through 1000 files
- **Default:** 1 trace file

flag

Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements:

- `all`—Trace all operations.
- `asynch`—Trace asynchronous configuration processing.
- `chassis-scheduler`—Trace chassis stream scheduler processing.
- `cos-adjustment`—Trace CoS rate adjustments.
- `dynamic`—Trace dynamic CoS functions.
- `hardware-database`—Trace the chassis hardware database related processing.
- `init`—Trace initialization events.
- `performance-monitor`—Trace performance monitor countors.
- `process`—Trace configuration processing.
- `restart`—Trace restart processing.
- `route-socket`—Trace route-socket events.
- `show`—Trace show command servicing.
- `snmp`—Trace SNMP-related processing.
- `util`—Trace utilities.

The following are the global tracing options:

- `all`—Perform all tracing operations
- `parse`—Trace parser processing.

`no-remote-trace`

(Optional) Disable remote tracing.

no-world-readable (Optional) Prevent any user from reading the log file.

size *size* (Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0. Incoming tracefile data is logged in the now empty *trace-file*. When *trace-file* again reaches its maximum size, *trace-file*.0 is renamed *trace-file*.1 and *trace-file* is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the files option.

- **Syntax:** *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB
- **Range:** 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)
- **Default:** 1 MB

world-readable (Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

traceoptions (Commit, Op, and SNMP Scripts)

IN THIS SECTION

- Syntax | 29512
- Hierarchy Level | 29512
- Description | 29512

- Default | 29512
- Options | 29513
- Required Privilege Level | 29514
- Release Information | 29514

Syntax

```
traceoptions {
    file <filename> <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system scripts commit],
[edit system scripts op],
[edit system scripts snmp]
```

Description

Define tracing operations for commit, op, and SNMP scripts.

Default

If you do not include this statement, the default operation during script execution is to log important events in the default trace file for that script type, because the traceoptions `events` flag is turned on by default, regardless of the configuration settings. When the file reaches 128 kilobytes (KB), it is renamed and compressed.

> **NOTE**: For op scripts, you can only customize tracing operations for local op scripts. For remote op scripts (op scripts executed using the `op url` command), Junos OS always logs the default trace events in the default op script log file.

Options

file ***filename***

Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**. By default, commit script process tracing output is placed in the file **cscript.log**, op script process tracing is placed in the file **op-script.log**, and SNMP script tracing is placed in the file **snmp.log**. If you include the `file` statement, you must specify a filename. To retain the default, you can specify **cscript.log**, **op-script.log**, or **snmp.log** as the filename.

- Default:
  - Commit scripts: **/var/log/cscript.log**
  - Op scripts: **/var/log/op-script.log**
  - SNMP scripts: **/var/log/snmp.log**

files ***number***

(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed and compressed to ***trace-file*.0.gz**. When ***trace-file*** again reaches its maximum size, ***trace-file*.0.gz** is renamed ***trace-file*.1.gz** and ***trace-file*** is renamed and compressed to ***trace-file*.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- Range: 2 through 1000
- Default: 10 files

flag ***flag***

Tracing operation to perform. The `events` flag is turned on by default, regardless of the configuration settings. You can include the following flags. To specify more than one tracing operation, include multiple `flag` statements.

- `all`—Log all operations
- `events`—Log important events, including errors, warnings, progress messages, and script processing events
- `input`—Log script input data
- `offline`—Generate data for offline development
- `output`—Log script output data
- `rpc`—Log script RPCs

- `xslt`—Log the XSLT library
- Default: `events`

no-world-readable Restrict file access to owner. This is the default.

size *size* (Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed and compressed to ***trace-file*.0.gz**. When ***trace-file*** again reaches its maximum size, ***trace-file*.0.gz** is renamed ***trace-file*.1.gz** and ***trace-file*** is renamed and compressed to ***trace-file*.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

- Syntax: *size* to specify bytes, *size*k to specify KB, *size*m to specify MB, or *size*g to specify GB
- Range: 10 KB through 1 GB
- Default: 128 KB

world-readable Enable unrestricted file access.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing Commit Script Processing

Tracing Op Script Processing

Tracing SNMP Script Processing

Dampening Script Execution

traceoptions (DDoS)

IN THIS SECTION

- Syntax | 29515
- Hierarchy Level | 29515
- Description | 29515
- Options | 29515
- Required Privilege Level | 29517
- Release Information | 29517

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system ddos-protection]
```

Description

Define tracing operations for DDoS protection processes.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the filename within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.
- `config`—Trace processing of the DDoS configuration at an extensive level.
- `events`—Trace jddosd event processing; currently only exit events are traced.
- `gres`—Trace messages exchanged with the kernel and jddosd process that could affect graceful Routing Engine switchover (GRES).
- `init`—Trace jddosd initialization.
- `ipc`—Trace interface interprocess communication (IPC) messages.
- `memory`—Trace memory management code. This flag is not currently supported.
- `protocol`—Trace DDoS protocol state processing. Only the violation state is currently traced.
- `rtsock`—Trace messages exchanged with the kernel and jddosd process.
- `signal`—Trace system signals that are passed to jddosd, such as SIGTERM.
- `socket`—Trace socket messages that are passed to jddosd from the Packet Forwarding Engine.
- `state`—Trace state machine events. This flag is not currently supported.
- `timer`—Trace jddosd timer events.
- `ui`—Trace user interface processing. This flag is not currently supported.

`level`—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages.

- `warning`—Match warning messages.

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10,240 through 1,073,741,824

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Tracing Control Plane DDoS Protection Operations

traceoptions (DHCP)

IN THIS SECTION

Syntax | 29518

- Hierarchy Level | 29518
- Description | 29518
- Options | 29519
- Required Privilege Level | 29520
- Release Information | 29521

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
    level  (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system processes dhcp-service]
[edit security dynamic-address]
```

Description

Define global tracing operations for extended DHCP local server and extended DHCP relay agent processes.

This statement replaces the deprecated `traceoptions` statements at the `[edit forwarding-options dhcp-relay]` and `[edit system services dhcp-local-server]` hierarchy levels.

> **NOTE**: Traceoptions does not differentiate between a logical system and tenant system, and can be configured under the root logical system.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements:

- `all`—Trace all events.
- `auth`—Trace authentication events.
- `database`—Trace database events.
- `fwd`—Trace firewall process events.
- `general`—Trace miscellaneous events.
- `ha`—Trace high availability-related events.
- `era`—Enables logging.
- `interface`—Trace interface operations.
- `io`—Trace I/O operations.
- `liveness-detection`—Trace liveness detection operations.
- `packet`—Trace packet and option decoding operations.
- `performance`—Trace performance measurement operations.
- `profile`—Trace profile operations.
- `rpd`—Trace routing protocol process events.
- `rtsock`—Trace routing socket operations.
- `security-persistence`—Trace security persistence events.
- `session-db`—Trace session database events.
- `state`—Trace changes in state.

- `statistics`—Trace baseline statistics.
- `ui`—Trace user interface operations.

`level`—Level of tracing to perform; also known as severity level. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- `all`—Match messages of all levels.
- `error`—Match error messages.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure `verbose`, messages at all higher levels are traced. Therefore, the result is the same as when you configure `all`.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access, allowing only the user `root` and users who have the Junos OS `maintenance` permission to access the trace files.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (*`maximum-file-size`*`k`), megabytes (*`maximum-file-size`*`m`), or gigabytes (*`maximum-file-size`*`g`). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Range:** 10,240 through 1,073,741,824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

DHCP Monitoring and Management

traceoptions (DHCP Server)

IN THIS SECTION

- Syntax | 29521
- Hierarchy Level | 29521
- Description | 29522
- Options | 29522
- Required Privilege Level | 29524
- Release Information | 29524

Syntax

```
traceoptions {
    file filename <files number> <match regex> <size size> <world-readable | no-world-readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit system services dhcp]
```

Description

Define tracing operations for DHCP processes.

Options

file ***filename***—Name of the file that receives the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—All tracing operations
- `binding`—Trace binding operations
- `config`—Log reading of configuration
- `conflict`—Trace user-detected conflicts for IP addresses
- `event`—Trace important events
- `ifdb`—Trace interface database operations
- `io`— Trace I/O operations
- `lease`—Trace lease operations
- `main`—Trace main loop operations
- `misc`— Trace miscellaneous operations
- `packet`—Trace DHCP packets
- `options`—Trace DHCP options
- `pool`—Trace address pool operations

- `protocol`—Trace protocol operations
- `rtsock`—Trace routing socket operations
- `scope`—Trace scope operations
- `signal`—Trace DHCP signal operations
- `trace`—All tracing operations
- `ui`—Trace user interface operations

`match` *`regex`*—(Optional) Refine the output to include lines that contain the regular expression.

- `all`—All tracing operations
- `binding`—Trace binding operations
- `config`— Log reading of configuration
- `conflict`—Trace user-detected conflicts for IP addresses
- `event`—Trace important events
- `ifdb`— Trace interface database operations
- `io`—Trace I/O operations
- `lease`—Trace lease operations
- `main`—Trace main loop operations
- `match` *`regex`*— Refine the output to include lines that contain the regular expression.
- `misc`—Trace miscellaneous operations
- `packet`—Trace DHCP packets
- `options`—Trace DHCP options
- `pool`—Trace address pool operations
- `protocol`—Trace protocol operations
- `rtsock`—Trace routing socket operations
- `scope`—Trace scope operations
- `signal`—Trace DHCP signal operations
- `trace`—All tracing operations

- `ui`—Trace user interface operations

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and filename.

- **Syntax:** *x*`k` to specify KB, *x*`m` to specify MB, or *x*`g` to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

traceoptions (Diameter Base Protocol for Subscriber Aware Policy Control)

IN THIS SECTION

- Syntax | 29525
- Hierarchy Level | 29525
- Description | 29525

- Options | 29525
- Required Privilege Level | 29527
- Release Information | 29527

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
    <peer peer-name>;
}
```

Hierarchy Level

```
[edit access diameter]
```

Description

Define tracing options for Diameter peers.

Options

file *filename* — Use the specified file to receive the output of the tracing operation. Enclose the filename within quotation marks. All files are placed in the directory `/var/log`.

files *number* — (Optional) Create the specified maximum number of trace files before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

flag *flag*

Use the specified tracing operation. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.
- `receive`—Trace received packets.
- `receive-detail`—Trace received packets in detail.
- `send`—Trace transmitted packets.
- `send-detail`—Trace transmitted packets in detail.
- `state`—Trace Diameter peer state changes.
- `timeout`—Trace timeout events.

level

Use the specified level of tracing. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

match *regular-expression*

(Optional) Refine the output to include lines that contain the specified regular expression.

no-remote-trace

Disable remote tracing.

no-world-readable

(Optional) Disable unrestricted file access.

peer *peer-name*

(Optional) Trace packets sent to or received from the specified peer. The specified peer must be defined at the `[edit access diameter peer]` hierarchy level.

size *maximum-file-size*

(Optional) Use the specified maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *size*k to specify KB, *size*m to specify MB, or *size*g to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable` (Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

traceoptions (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 29528
- Hierarchy Level | 29528
- Description | 29528
- Options | 29528
- Required Privilege Level | 29529
- Release Information | 29530

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system processes diameter-service]
```

Description

Define tracing options for Diameter processes.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the filename within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations
- `application`—Trace Diameter application interface events
- `configuration`—Trace configuration events
- `daemon`—Trace Diameter daemon level events
- `diameter-instance`—Trace Diameter instance events

- `dne`—Trace Diameter network element events
- `framework`—Trace Diameter framework events
- `memory-management`—Trace memory management events
- `messages`—Trace Diameter messages
- `node`—Trace Diameter node events
- `peer`—Trace Diameter peer events

`level`—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Tracing Diameter Base Protocol Processes for Subscriber Access

traceoptions (DNS, Port, and TFTP Packet Forwarding)

IN THIS SECTION

- Syntax | 29530
- Hierarchy Level | 29531
- Description | 29531
- Default | 29531
- Options | 29531
- Required Privilege Level | 29532
- Release Information | 29532

Syntax

```
traceoptions {
    file filename  <files number> <match regular-expression> <size bytes> <world-readable | no-world-readable>;
    flag flag;
    level level;
    <no-remote-trace>;
}
```

Hierarchy Level

```
[edit forwarding-options helpers]
```

Description

Configure tracing operations for BOOTP, DNS, TFTP, or custom UDP port packet forwarding.

Default

If you do not include this statement, no tracing operations are performed.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" "). All files are placed in a file named **fud** in the directory **/var/log**. If you include the `file` statement, you must specify a filename.

files ***number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

- **Range:** 2 through 1000
- **Default:** 3 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- **address**—Trace address management events
- **all**—Trace all events
- **bootp**—Trace BOOTP or DHCP services events
- **config**—Trace configuration events
- **domain**—Trace DNS service events
- **ifdb**—Trace interface database operations
- **io**—Trace I/O operations

- **main**—Trace main loop events
- **port**—Trace arbitrary protocol events
- **rtsock**—Trace routing socket operations
- **tftp**—Trace TFTP service events
- **trace**—Trace tracing operations
- **ui**—Trace user interface operations
- **util**—Trace miscellaneous utility operations

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—(Optional) Disable remote tracing globally or for a specific tracing operation.

no-world-readable—(Optional) Restrict file access to the owner.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** file again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 0 bytes through 4,294,967,295 KB
- **Default:** 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement standardized and **match** option introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Tracing BOOTP, DNS, and TFTP Forwarding Operations

traceoptions (dynamic-address)

IN THIS SECTION

- Syntax | 29533
- Hierarchy Level | 29534
- Description | 29534
- Options | 29534
- Required Privilege Level | 29534
- Release Information | 29534

Syntax

```
traceoptions {
 file< filename><files files><match match><size size><(world-readable | no-world-readable)>;
 flag name;

 level(all | error | info | notice | verbose | warning);
 no-remote-trace;
}
```

Hierarchy Level

```
[edit security dynamic-address]
```

Description

Security dynamic address tracing options

Options

level Level of debugging output

- all—Match all levels
- error—Match error conditions
- info—Match informational messages
- notice—Match conditions that should be handled specially
- verbose—Match verbose messages
- warning—Match warning messages

no-remote-trace Disable remote tracing

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D25 on SRX Series Firewalls and on cSRX Container Firewall in Junos OS Release 22.2R1.

traceoptions (dynamic-application)

IN THIS SECTION

- Syntax | 29535
- Hierarchy Level | 29535
- Description | 29535
- Options | 29536
- Required Privilege Level | 29537
- Release Information | 29537

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security dynamic-application]
[edit security dynamic-address]
```

Description

Configure dynamic application tracing options.

NOTE: NOTE: Traceoptions does not differentiate between a logical system and tenant system, and can be configured under the root logical system.

Options

file Trace file information

- Values:

  filename Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.

  files *number* Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed as *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

  If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a *filename*.

  - **Range:** 2 through 1000
  - **Default:** 3 files

  match *regular-expression* Refine the output to include lines that contain the regular expression.

  size *maximum-file-size* Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  - **Range:** 10 KB through 1 GB
  - **Default:** 128 KB

  world-readable|no- By default, the log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any

world-readable	user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

flag Configure tracing parameters

- Values:

all	Trace with all flags enabled.
configuration	Trace configuration events.
compilation	Trace compilation events.
ipc	Trace process intercommunication events.
lookup	Trace lookup events.

no-remote-trace Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Unified Security Policies

Unified Security Policies

traceoptions (Dynamic Flow Capture)

IN THIS SECTION

- Syntax | 29538
- Hierarchy Level | 29538
- Description | 29538
- Options | 29538
- Required Privilege Level | 29539
- Release Information | 29539

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | non-world-readable>;
}
```

Hierarchy Level

```
[edit services dynamic-flow-capture]
```

Description

Enable and define tracing options for dynamic flow capture events.

Options

file *filename* Use the specified file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number* (Optional) Use the specified maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is

overwritten. Note that if you specify a maximum number for files, you must also specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files.
- **Default:** 10 files.

no-world-readable (Optional) Restrict access to the file.

world-readable (Optional) Enable free access to the file.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Junos Capture Vision

traceoptions (Egress Protection)

IN THIS SECTION

- Syntax | 29540
- Hierarchy Level | 29540
- Description | 29540
- Options | 29540
- Required Privilege Level | 29541
- Release Information | 29541

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit protocols mpls egress-protection]
```

Description

Configure tracing operations for egress protection.

Options

filename Name of the file to receive the output of the tracing operation.

files *number* (Optional) Maximum number of trace files. If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000
- **Default:** 2 files

flag Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

- `all`—Trace all operations
- `error`—Trace error conditions
- `route`—Trace route transitions
- `state`—Trace state transitions

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4R3.

RELATED DOCUMENTATION

Example: Configuring MPLS Egress Protection for Layer 3 VPN Services

traceoptions (Enhanced Subscriber Management)

IN THIS SECTION

- Syntax | 29541
- Hierarchy Level | 29542
- Description | 29542
- Options | 29542
- Required Privilege Level | 29544
- Release Information | 29544

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag <disable>;
    level level;
```

```
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system processes smg-service]
```

Description

Define tracing operations for the main enhanced subscriber management session management process.

Options

`file` *`filename`*— Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All trace logs are stored in the directory `/var/log/bbesmgd`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 5 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. Include the `disable` option after a flag to disable tracing for that flag. You can include the following flags:

- `all`—Trace all operations.
- `auth`—Trace general authentication events.
- `autoconf`—Trace autoconfiguration events.
- `config`—Trace configuration events.
- `demux`—Trace demultiplexing (demux) events.
- `dhcp`—Trace DHCP Server events.
- `dprof`—Trace dynamic profile events.
- `interface`—Trace interface events.
- `io`—Trace packet socket events.

- `l2tp`—Trace L2TP events.
- `main`—Trace main operations.
- `net`—Trace network processing events.
- `ppp`—Trace PPP events.
- `pppoe`—Trace PPPoE events.
- `rpd`—Trace routing protocol process events.
- `rtsock`—Trace routing socket events.
- `service`—Trace service events.
- `session`—Trace connection events and flow.
- `stats`—Trace statistics events.
- `ucac`—Trace universal call admission control events.
- `ui`—Trace user interface events.
- `vbf`—Trace variable-based forwarding events.

`level` *`level`*—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions reqcliring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *size*k to specify KB, *size*m to specify MB, or *size*g to specify GB
- **Range:** 10,240 through 1,073,741,824 bytes
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1R3.

RELATED DOCUMENTATION

Configuring Junos OS Enhanced Subscriber Management

Configuring the Subscriber Management Database Trace Log Filename

traceoptions (Event Options)

IN THIS SECTION

- Syntax (Junos OS) | 29545
- Hierarchy Level | 29545
- Description | 29545
- Default | 29545
- Options | 29545
- Required Privilege Level | 29547
- Release Information | 29547

Syntax (Junos OS)

```
traceoptions {
    file <filename> <files number> <match regular-expression> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit event-options]
```

Description

Define tracing operations for event policies on devices running Junos OS.

> Junos OS Evolved uses a different tracing infrastructure, which collects trace data from all applications on all nodes by default. For more information, see *Trace Event Policy Processing on Devices Running Junos OS Evolved*.

Default

If you do not include this statement, no event-policy-specific tracing operations are performed.

Options

file *filename*

Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**. By default, event policy process tracing output is placed in the file **eventd**. If you include the `file` statement, you must specify a filename. To retain the default, you can specify **eventd** as the filename.

- **Default:** /var/log/eventd

files *number*

(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed and compressed to ***trace-file*.0.gz**. When ***trace-file*** again reaches its maximum size, ***trace-file*.0.gz** is renamed ***trace-file*.1.gz** and ***trace-file*** is

renamed and compressed to ***trace-file*.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000
- **Default:** 3 files

flag *flag*

Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Log all operations
- `configuration`—Log reading of configuration at the `[edit event-options]` hierarchy level
- `events`—Log eventd processing
- `database`—Log events involving storage and retrieval in events database
- `server`—Log communication with processes that are generating events
- `timer-events`—Log internally generated events

match *regular-expression*

(Optional) Refine the output to include lines that contain the regular expression.

no-world-readable

Restrict file access to owner. This is the default.

size *size*

(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed and compressed to ***trace-file*.0.gz**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0.gz** is renamed ***trace-file*.1.gz** and ***trace-file*** is renamed and compressed to ***trace-file*.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and filename.

- **Syntax:** *size* to specify bytes, *size*k to specify KB, *size*m to specify MB, or *size*g to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

world-readable (Optional) Enable unrestricted file access.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Trace Event Policy Processing on Devices Running Junos OS

traceoptions (Event Scripts)

IN THIS SECTION

- Syntax | 29547
- Hierarchy Level | 29548
- Description | 29548
- Default | 29548
- Options | 29548
- Required Privilege Level | 29549
- Release Information | 29550

Syntax

```
traceoptions {
    file <filename> <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
```

```
    no-remote-trace;
}
```

Hierarchy Level

```
[edit event-options event-script]
```

Description

Define tracing operations for event scripts.

Default

If you do not include this statement, the default operation during script execution is to log important events in the default trace file **/var/log/escript.log**, because the traceoptions `events` flag is turned on by default, regardless of the configuration settings. When the file reaches 128 kilobytes (KB), it is renamed and compressed.

Options

file ***filename***

Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**. By default, event script process tracing output is placed in the file **escript.log**. If you include the `file` statement, you must specify a filename. To retain the default, you can specify **escript.log** as the filename.

- **Default: /var/log/escript.log**

files ***number***

(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed and compressed to ***trace-file*.0.gz**. When ***trace-file*** again reaches its maximum size, ***trace-file*.0.gz** is renamed ***trace-file*.1.gz** and ***trace-file*** is renamed and compressed to ***trace-file*.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000
- **Default:** 10 files

flag *flag*

Tracing operation to perform. The `events` flag is turned on by default, regardless of the configuration settings. You can include the following flags. To specify more than one tracing operation, include multiple `flag` statements.

- `all`—Log all operations
- `events`—Log important events, including errors, warnings, progress messages, and script processing events
- `input`—Log event script input data
- `offline`—Generate data for offline development
- `output`—Log event script output data
- `rpc`—Log event script RPCs
- `xslt`—Log the XSLT library
- **Default:** `events`

no-world-readable

Restrict file access to owner. This is the default.

size *size*

(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed and compressed to ***trace-file*.0.gz**. When ***trace-file*** again reaches its maximum size, ***trace-file*.0.gz** is renamed ***trace-file*.1.gz** and ***trace-file*** is renamed and compressed to ***trace-file*.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

- **Syntax:** *size* to specify bytes, *size*k to specify KB, *size*m to specify MB, or *size*g to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

world-readable

Enable unrestricted file access.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Trace Event Script Processing on Devices Running Junos OS

Dampen Script Execution

traceoptions (Extensible Subscriber Services Manager)

IN THIS SECTION

- Syntax | 29550
- Hierarchy Level | 29550
- Description | 29551
- Options | 29551
- Required Privilege Level | 29551
- Release Information | 29551

Syntax

```
traceoptions file file-name flag authentication | flag configuration | flag dictionary | flag
fsm | flag statistics | flag kernel | flag dynamic | flag database | flag op-script | flag
general | flag all
```

Hierarchy Level

```
[edit system processes extensible-subscriber-services]
```

Description

Configure the trace options for Extensible Subscriber Services Manager. essmd supports Junos OS daemon trace options.

Options

`traceoptions`	Configure the trace options for Extensible Subscriber Services Manager.
`file-name`	Name of the trace file.
`flag all`	Configure trace settings for all operations.
`flag authentication`	Configure trace settings for authentication operations.
`flag configuration`	Configure trace settings for configuration operations.
`flag database`	Configure trace settings for database operations.
`flag dictionary`	Configure trace settings for dictionary operations.
`flag dynamic`	Configure trace settings for dynamic profile operations.
`flag fsm`	Configure trace settings for finite state machine operations.
`flag general`	Configure trace settings for general operations.
`flag kernel`	Configure trace settings for kernel state-change operations.
`flag op-script`	Configure trace settings for op script operations.
`flag statistics`	Configure trace settings for statistics-collection operations.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

traceoptions (Extensions)

IN THIS SECTION

- Syntax | 29552
- Hierarchy Level | 29552
- Description | 29552
- Options | 29553
- Required Privilege Level | 29554
- Release Information | 29554

Syntax

```
traceoptions {
    file <filename> <files number> <match regex> <size size> <world-readable | no-world-
readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system extensions extension-service application],
[edit system extensions extension-service application file script-name]
```

Description

Trace options for extension-service applications. You can set the `traceoptions` statement for the `application` level or for the `file` *`script-name`* level or for both. If the `traceoptions` statement is set for both the `application` and `file` *`script-name`* levels, the latter will have a higher priority.

> **NOTE**: Global traceoptions for daemonized applications do not take effect if the daemonized application and global traceoptions are committed separately

Options

`file` Indicate trace file information.

- *`filename`* Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the **/var/log/** directory.
- `files` *`number`* (Optional) Specify maximum number of trace files.
  - **Range:** 2 through 1000
  - **Default:** 3
- `size` *`size`* (Optional) Specify the maximum size of each trace file. When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`. The traceoption output continues in a second trace file named *`trace-file`*`.1`. When *`trace-file`*`.1` reaches its maximum size, output continues in a third file named *`trace-file`*`.2`, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.
  - **Range:** 10240 through 1073741824
  - **Default:** 128k
- `world-readable | no-world-readable` (Optional). Grant all users permission to read log files, or restrict the permission only to the root user and users who have Junos OS maintenance permission.

`flag` *`flag`* Specify the tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements:

- `all` Trace all operations.
- `config` Trace important events.
- `general` Trace script input data.
- `grpc` Trace grpc server events.

`notification`	Trace notification events.
`routing-socket`	Trace routing socket calls.
`timeouts`	Trace timeouts.
`timer`	Trace internal timer events.

`no-remote-trace` Disable remote tracing. This option is valid only when [system tracing] is configured.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

application (Extensions) | 2316

traceoptions (FC-2 Fibre Channel)

IN THIS SECTION

- Syntax | 29555
- Hierarchy Level | 29555
- Description | 29555
- Default | 29555
- Options | 29555
- Required Privilege Level | 29557
- Release Information | 29557

Syntax

```
traceoptions {
    file filename <replace> <size size> <files number> <no-stamp>;
    <world-readable | no-world-readable>;
    flag flag <flag-modifier>;
}
```

Hierarchy Level

```
[edit fc-fabrics fabric-name fc2]
```

Description

Set FC-2 protocol tracing options.

NOTE: The `traceoptions` statement is not supported on the QFabric system.

Default

Traceoptions is disabled.

Options

file ***name***—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the **/var/log/** directory.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**. The traceoption output continues in a second trace file named ***trace-file*.1**. When ***trace-file*.1** reaches its maximum size, output continues in a third file named ***trace-file*.2**, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 1 trace file

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements:

- **all** —Trace all operations.
- **error**—Trace all error events
- **normal**—Trace all normal events.
- **Default:** If you do not specify the **normal** option, only unusual or abnormal operations are traced.
- **parse**—Trace configuration parsing.
- **rx-frame**—(Optional) Trace received frames.
- **rx-frame-header**—(Optional) Trace received frame headers.
- **tx-frame**—(Optional) Trace transmitted frames.
- **tx-frame-header**—(Optional) Trace transmitted frame headers.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

- **Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Prevent any user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*** **.0**. Incoming tracefile data is logged in the now empty ***trace-file***. When ***trace-file*** again reaches its maximum size, ***trace-file*****.0** is renamed ***trace-file*****.1** and ***trace-file*** is renamed ***trace-file*****.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*****k** to specify KB, ***x*****m** to specify MB, or ***x*****g** to specify GB
- **Range:** 10 KB through the maximum file size of 4 GB (the maximum is lower if 4 GB is not supported on your system)
- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

traceoptions (Fibre Channel)

IN THIS SECTION

- Syntax | 29557
- Hierarchy Level | 29557
- Description | 29558
- Default | 29558
- Options | 29558
- Required Privilege Level | 29560
- Release Information | 29560

Syntax

```
traceoptions {
    file filename <replace> <size size> <files number> <no-stamp>
    <world-readable | no-world-readable>;
    flag flag <flag-modifier>;
}
```

Hierarchy Level

```
[edit fc-options]
```

Description

Set FC protocol tracing options.

> **NOTE**: The `traceoptions` statement is not supported on the QFabric system.

Default

Traceoptions is disabled.

Options

file ***name***—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the **/var/log/** directory.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**. The traceoption output continues in a second trace file named ***trace-file*.1**. When ***trace-file*.1** reaches its maximum size, output continues in a third file named ***trace-file*.2**, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 1 trace file

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements:

- **all**—Trace all operations.
- **fabric**—Trace virtual fabric events.
- **fc2**—Trace the FC2 (network layer protocols) events.
- **fip**—Trace the Fibre Channel over Ethernet (FCoE) Initialization Protocol events.
- **flogi** —Trace the fabric login server events.
- **forwarding-database**—Trace the forwarding database and next-hop events.
- **interface**—Trace the interface events.
- **krt**—Trace the communication over the routing socket.

- **lib**—Trace library calls.
- **lif**—Trace Fibre Channel logical interface (fc-lif) events.
- **vswitch**—Trace virtual switch events.

The following are the global tracing options:

- **all**—All trace operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events.
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **parse**—Trace configuration parsing.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

- **Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Prevent any user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*** **.0.** Incoming tracefile data is logged in the now empty ***trace-file***. When ***trace-file*** again reaches its maximum size, ***trace-file*****.0** is renamed ***trace-file*****.1** and ***trace-file*** is renamed ***trace-file*****.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*****k** to specify KB, ***x*****m** to specify MB, or ***x*****g** to specify GB
- **Range:** 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)

- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

traceoptions (FIP Protocol Fibre Channel)

IN THIS SECTION

- Syntax | 29560
- Hierarchy Level | 29561
- Description | 29561
- Default | 29561
- Options | 29561
- Required Privilege Level | 29562
- Release Information | 29562

Syntax

```
traceoptions {
    file filename <replace> <size size> <files number> <no-stamp>
    <world-readable | no-world-readable>;
    flag flag <flag-modifier>
}
```

Hierarchy Level

```
[edit fc-fabrics fabric-name protocols fip]
```

Description

Set proxy FC protocol tracing options.

NOTE: The `traceoptions` statement is not supported on the QFabric system.

Default

Traceoptions is disabled.

Options

file ***name***—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the **/var/log/** directory.

files ***number*** —(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**. The traceoption output continues in a second trace file named ***trace-file*.1**. When ***trace-file*.1** reaches its maximum size, output continues in a third file named ***trace-file*.2**, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the size option.

- **Range:** 2 through 1000 files
- **Default:** 1 trace file

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements:

- **all**—Trace all operations.
- **error**—Trace all error events
- **normal**—Trace all normal events.
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **packet**—Trace packet decoding operations

- **parse**—Trace configuration parsing.
- **state**—Trace state transitions.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

- **Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Prevent any user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0.** Incoming tracefile data is logged in the now empty ***trace-file***. When ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)
- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

traceoptions (flow)

IN THIS SECTION

- Syntax | 29563
- Hierarchy Level | 29564
- Description | 29564
- Options | 29564
- Required Privilege Level | 29565
- Release Information | 29565

Syntax

```
    traceoptions {
        file< filename><files files><match match><size size><(world-readable | no-world-
readable)>;
        flag name;

        local-debug-buf {
            size size;
        }
        no-remote-trace;
        packet-filter name {
            conn-tag conn-tag;
            destination-port(afs | bgp | biff | bootpc | bootps | cmd | cvspserver | dhcp |
domain | eklogin | ekshell | exec | finger | ftp | ftp-data | http | https | ident | imap |
kerberos-sec | klogin | kpasswd | krb-prop | krbupdate | kshell | ldap | ldp | login | mobileip-
agent | mobilip-mn | msdp | netbios-dgm | netbios-ns | netbios-ssn | nfsd | nntp | ntalk | ntp |
pop3 | pptp | printer | radacct | radius | range | rip | rkinit | smtp | snmp | snmptrap | snpp
| socks | ssh | sunrpc | syslog | tacacs | tacacs-ds | talk | telnet | tftp | timed | who |
xdmcp | zephyr-clt | zephyr-hm | zephyr-srv);
            destination-prefix destination-prefix;
            interface interface;


            protocol(ah | egp | esp | gre | icmp | icmp6 | igmp | ipip | number | ospf | pim |
```

```
rsvp | sctp | tcp | udp);
            source-port(afs | bgp | biff | bootpc | bootps | cmd | cvspserver | dhcp | domain |
eklogin | ekshell | exec | finger | ftp | ftp-data | http | https | ident | imap | kerberos-sec
| klogin | kpasswd | krb-prop | krbupdate | kshell | ldap | ldp | login | mobileip-agent |
mobilip-mn | msdp | netbios-dgm | netbios-ns | netbios-ssn | nfsd | nntp | ntalk | ntp | pop3 |
pptp | printer | radacct | radius | range | rip | rkinit | smtp | snmp | snmptrap | snpp | socks
| ssh | sunrpc | syslog | tacacs | tacacs-ds | talk | telnet | tftp | timed | who | xdmcp |
zephyr-clt | zephyr-hm | zephyr-srv);
            source-prefix source-prefix;
        }
        rate-limit rate-limit;
        root-override;
        trace-level {
            (brief | detail | minimal);
        }
    }
```

Hierarchy Level

```
[edit logical-systems name security flow]
```

Description

Starting in Junos OS Release 24.2R1, we support `local-debug-buf` option to configure number of lines for the flow local trace buffer, and the maximum lines allowed is 40,000. A new flag `packet-track` has been introduced to get information about the packet being handled in SRX data path, The flag `packet-track` should be configured in root logical system. To track information about the packets in non root logical system use the option `root-override` along with flag `packet-track`.

Options

no-remote-trace Disable remote tracing

rate-limit Limit the incoming rate of trace messages

- **Default:** 0
- **Range:** 0 through 4294967295

root-override Allow collect flow trace in root from all logical-systems and tenants

trace-level Flow trace level

- brief—Brief messages including notice, in addition to minimal messages
- detail—Detail messages including info and debug, in addition to brief messages
- minimal—Significant messages including warning, error, criticality, alert and emergency.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 24.2R1.

traceoptions (Forwarding Options)

IN THIS SECTION

- Syntax | 29566
- Hierarchy Level | 29566
- Description | 29566
- Required Privilege Level | 29566
- Release Information | 29566

Syntax

```
traceoptions {
    no-remote-trace;
    file filename <files number> <size bytes> <match expression> <world-readable | no-world-
readable>;
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring],
[edit forwarding-options sampling]
```

Description

Configure traffic sampling tracing operations.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing Traffic Sampling Operations

traceoptions (General Authentication Service)

IN THIS SECTION

- Syntax | 29567
- Hierarchy Level | 29567
- Description | 29567
- Options | 29568
- Required Privilege Level | 29569
- Release Information | 29569

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    filter {
        user user@domain;
    }
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system processes general-authentication-service]
```

Description

Configure tracing options for the general authentication service.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`filter`—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifier simplifies troubleshooting in a scaled environment.

- `user` *`user@domain`*—Username of a subscriber. Optionally use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms.

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `address-assignment`—Trace address-assignment pool events
- `all`—Trace all tracing operations
- `configuration`—Trace configuration events
- `framework`—Trace authentication framework events
- `gx-plus`—Trace Gx-Plus events
- `jsrc`—Trace JSRC events
- `ldap`—Trace LDAP authentication events
- `local-authentication`—Trace local authentication events
- `radius`—Trace RADIUS authentication events
- `user-access`—Trace user access events, such as login, logout, and authenticate.

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB),

megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *size*`k` to specify KB, *size*`m` to specify MB, or *size*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Tracing General Authentication Service Processes

traceoptions (Health Check Monitoring)

IN THIS SECTION

- Syntax | 29570
- Hierarchy Level | 29570
- Description | 29570
- Options | 29570
- Required Privilege Level | 29572
- Release Information | 29573

Syntax

```
traceoptions {
    file file-name ;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    monitor monitor-object-name {
        group-name group-name;
        real-services-name real-service-name;
    }
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services network-monitoring]
```

Description

Specify tracing options for the traffic load balancer health check monitoring function.

Options

file *file-name*

Use the specified name of the file to receive the output of the tracing operation.

flag *flag*

Specify which operations you want to trace. To specify more than one operation, include multiple flag statements. Table 1809 on page 29570 species the various values you can set this option to.

Table 1809: Trace Flags

Flag	Support on MS-MPC and SPC3 Cards	Description
all	MS-MPC and SPC3	Trace all operations.
all-real-services	SPC3	Trace all real services.

Table 1809: Trace Flags *(Continued)*

Flag	Support on MS-MPC and SPC3 Cards	Description
config	MS-MPC	Trace traffic load balancer configuration events.
connect	MS-MPC	Trace traffic load balancer ipc events.
database	MS-MPC and SPC3	Trace database events.
file-descriptor-queue	MS-MPC	Trace file descriptor queue events.
inter-thread	MS-MPC	Trace inter-thread communication events.
filter	MS-MPC	Trace traffic load balancer filter programming events.
health	MS-MPC	Trace traffic load balancer health events.
messages	MS-MPC and SPC3	Trace normal events.
normal	MS-MPC	Trace normal events.
operational-commands	MS-MPC	Trace traffic load balancer show events.
parse	MS-MPC	Trace traffic load balancer parse events.
probe	MS-MPC and SPC3	Trace probe events.
probe-infra	MS-MPC and SPC3	Trace probe infra events.
route	MS-MPC	Trace traffic load balancer route events.

Table 1809: Trace Flags *(Continued)*

Flag	Support on MS-MPC and SPC3 Cards	Description
snmp	MS-MPC	Trace traffic load balancer SNMP events.
statistics	MS-MPC	Trace traffic load balancer statistics events.
system	MS-MPC	Trace traffic load balancer system events.

group-name *group-name*
Specify which server group is to be traced.

level
Use the specified level of tracing. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages
- `notice`—Match conditions that must be handled specially.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

monitor *monitor-object-name*
Name of a monitoring object that contains server group or real service.

no-remote-trace
Disable remote tracing.

real-services-name *real-service-name*
Specify which real service is to be traced.

Required Privilege Level

trace and interface—To view this statement in the configuration.

trace-control and interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring TLB

traceoptions (IGMP Snooping)

IN THIS SECTION

- Syntax | 29573
- Hierarchy Level | 29573
- Description | 29574
- Default | 29574
- Options | 29574
- Required Privilege Level | 29575
- Release Information | 29576

Syntax

```
traceoptions  {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
    flag flag <flag-modifier>;
}
```

Hierarchy Level

```
[edit protocols igmp-snooping]
```

Description

Define tracing operations for IGMP snooping.

Default

The **traceoptions** feature is disabled by default.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files, including the active trace file. When a trace file reaches its maximum size, its contents are archived into a compressed file named ***filename*.0** and the trace file is emptied. When the trace file reaches its maximum size again, the ***filename*.0** archive file is renamed ***filename*.1** and a new ***filename*.0** archive file is created from the contents of the trace file. This process continues until the maximum number of trace files is reached, at which point the system starts overwriting the oldest archive file each time the trace file is archived. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000
- **Default:** 10 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.
- **general**—Trace general IGMP snooping protocol events.
- **krt**—Trace communication over routing socket.
- **leave**—Trace leave group messages (IGMPv2 and IGMPv3 only).
- **nexthop**—Trace nexthop-related events.
- **normal**—Trace normal IGMP snooping protocol events. If you do not specify this flag, only unusual or abnormal operations are traced.
- **packets**—Trace all IGMP packets.
- **policy**—Trace policy processing.
- **query**—Trace IGMP membership query messages.

- **report**—Trace membership report messages.
- **route**—Trace routing information.
- **state**—Trace IGMP state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.
- **vlan**—Trace VLAN-related events.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers per flag:

- **detail**—Provide detailed trace information
- **disable**—Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

no-stamp—(Optional) Omit the timestamp at the beginning of each line in the trace file.

no-world-readable—(Optional) Restrict file access to the user who created the file.

replace—(Optional) Replace an existing trace file if there is one. If you do not include this option, tracing output is appended to an existing trace file.

size ***size*** —(Optional) Maximum size of each trace file, in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is zipped and renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum size, you also must specify a maximum number of files with the **files** option.

- **Syntax:** ***x*** to specify bytes, ***x*****k** to specify KB, ***x*****m** to specify MB, or ***x*****g** to specify GB
- **Range:** 10240 through 4294967295 bytes
- **Default:** 128 KB

world-readable—(Optional) Allow unrestricted file access.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

traceoptions (Individual Interfaces)

IN THIS SECTION

- Syntax (Individual interfaces with PTX Series, EX Series, ACX Series) | **29576**
- Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series) | **29577**
- Syntax (OAMLFM with EX Series, QFX Series, NFX Series) | **29577**
- Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series) | **29577**
- Hierarchy Level (Individual interfaces with PTX Series, EX Series, ACX Series, QFX Series, OCX1100, EX4600, NFX Series) | **29577**
- Hierarchy Level (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series) | **29577**
- Description | **29578**
- Default | **29578**
- Options | **29578**
- Required Privilege Level | **29584**
- Release Information | **29584**

Syntax (Individual interfaces with PTX Series, EX Series, ACX Series)

```
traceoptions {
    file filename <files name> <size size> <world-readable | no-world-readable>;
    flag flag;
    match;
}
```

Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series)

```
traceoptions {
    flag flag;
}
```

Syntax (OAMLFM with EX Series, QFX Series, NFX Series)

```
traceoptions {
    file filename <files number> <match regex> <size size> <world-readable | no-world-readable>;
    flag flag ;
    no-remote-trace;
}
```

Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series)

```
traceoptions {
    file <filename> <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
    flag flag <disable>;
    no-remote-trace;
}
```

Hierarchy Level (Individual interfaces with PTX Series, EX Series, ACX Series, QFX Series, OCX1100, EX4600, NFX Series)

```
[edit interfaces interface-name]
```

Hierarchy Level (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series)

```
[edit interfaces]
```

Description

Define tracing operations for individual interfaces.

To specify more than one tracing operation, include multiple `flag` statements.

The interfaces `traceoptions` statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system `syslog` file in the directory */var/log/dcd* .

On EX Series, QFX Series, and NFX Series platforms, configure tracing options the link fault management.

On ACX Series, SRX Series, MX Series, M Series, and T Series platforms define tracing operations for the interface process (dcd).

Default

If you do not include this statement, no interface-specific tracing operations are performed.

Options

Table 1810 on page 29578 lists options for traceoption command for the following platforms:

Table 1810: Options for traceoptions

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
`file` *filename*	—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory */var/log/dcd* . By default, interface process tracing output is placed in the file.		—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory */var/log/dcd* .	—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory */var/log/dcd* . By default, interface process tracing output is placed in the file `dcd`.

Table 1810: Options for traceoptions *(Continued)*

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
files *number*	—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.		—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the size option.	—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option. *Range:* 2 through 1000 *Default:* 3 files

Table 1810: Options for traceoptions *(Continued)*

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
flag	—Tracing operation to perform. To specify more than one tracing operation, include multiple *flag* statements. The following are the interface-specific tracing options. • `all`—All interface tracing operations • `event`—Interface events • `ipc`—Interface interprocess communication (IPC) messages • `media`—Interface media changes • `q921`—Trace ISDN Q.921 frames • `q931`—Trace ISDN Q.931 frames	—Tracing operation to perform. To specify more than one tracing operation, include multiple *flag* statements. The following are the interface-specific tracing options. • `all`—All interface tracing operations • `event`—Interface events • `ipc`—Interface interprocess communication (IPC) messages • `media`—Interface media changes • `q921`—Trace ISDN Q.921 frames • `q931`—Trace ISDN Q.931 frames	—Tracing operation to perform. To specify more than one tracing operation, include multiple *flag* statements. You can include the following flags: • `action-profile`—Trace action profile invocation events. • `all`—Trace all events. • `configuration`—Trace configuration events. • `protocol`—Trace protocol processing events. • `routing socket`—Trace routing socket events.	—Tracing operation to perform. To specify more than one tracing operation, include multiple *flag* statements. You can include the following flags: • `all` • `change-events`—Log changes that produce configuration events • `config-states`—Log the configuration state machine changes • `kernel`—Log configuration IPC messages to kernel • `kernel-detail`—Log details of configuration messages to kernel

Table 1810: Options for traceoptions *(Continued)*

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
`match`	—(Optional) Regular expression for lines to be traced.		—(Optional) Refine the output to log only those lines that match the given regular expression.	

Table 1810: Options for traceoptions *(Continued)*

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
size *size*	—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named `trace-file` reaches this size, it is renamed `trace-file.0`. When the `trace-file` again reaches its maximum size, `trace-file.0` is renamed `trace-file.1` and `trace-file` is renamed `trace-file.0`. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.		—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the files option. *Syntax:* *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB *Range:* 10 KB through 1 GB *Default:* 128 KB *Default:* If you do not include this option, tracing output is appended	—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the *`trace-file`* again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option. *Syntax:* *x*k to specify kilobytes, *x*m to specify megabytes, or *x*g to specify gigabytes

Table 1810: Options for traceoptions *(Continued)*

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
			to an existing trace file.	*Range:* 10 KB through the maximum file size supported on your router *Default:* 1 MB
`no-world-readable`	—(Optional) Prevent any user from reading the log file.		—(Optional) Restrict file access to the user who created the file.	—(Optional) Disallow any user to read the log file.
`world-readable`	—(Optional) Allow any user to read the log file.		—(Optional) Enable unrestricted file access.	—(Optional) Allow any user to read the log file.
`disable`				—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.
no-remote-trace			—(Optional) Disable the remote trace.	-

Table 1810: Options for traceoptions *(Continued)*

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
`match` *regex*				—(Optional) Refine the output to include only those lines that match the given regular expression.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

- Tracing Operations of an Individual Router Interface
- Tracing Operations of an Individual Router or Switch Interface
- *Example: Configuring Ethernet OAM Link Fault Management*
- *Configuring Ethernet OAM Link Fault Management*
- *Tracing Operations of the Interface Process*

traceoptions (Inline Monitoring)

IN THIS SECTION

- Syntax | 29585
- Hierarchy Level | 29585
- Description | 29585
- Options | 29585
- Additional Information | 29586
- Required Privilege Level | 29587
- Release Information | 29587

Syntax

```
traceoptions {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services inline-monitoring]
```

Description

(Junos OS only) Configure traceoptions for the inline monitoring process.

Options

file ***file-name*** Use the specified file to receive the output of the tracing operation. All files are placed in the directory **/var/log**.

files *files*
(Optional) Use the specified maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Default:** 3 files
- **Range:** 2 through 1000 files

match *match*
(Optional) Use the specified regular expression to refine the output to include lines that contain the regular expression.

no-remote-trace
Disable remote tracing.

no-world-readable
(Default) Disable unrestricted file access. This means the log file can be accessed only by the user who configured the tracing operation.

world-readable
(Optional) Enable unrestricted file access.

size *size*
(Optional) Use the specified maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

Additional Information

Junos OS Evolved uses a new tracing architecture. All running applications create trace information, with multiple instances of the same application having their own trace information. In Junos OS, you enable tracing operations by configuring the traceoptions statement at the specific hierarchy level you want to trace. Junos OS Evolved, on the other hand, uses an application-based model, and thus trace messages are logged, viewed, and configured by application. As a result, Junos OS Evolved does not support the `traceoptions` statement at many of the hierarchy levels that Junos OS supports.

In Junos OS Evolved, you do not view trace files directly, and you should never add, edit, or remove trace files under the **/var/log/traces** directory because this can corrupt the traces. Instead, you use the `show trace application` *`application-name`* `node` *`node-name`* command to read and decode trace messages stored in the trace files. All running applications on Junos OS Evolved create trace information at the `info` level by default.

Inline monitoring services are governed by the `imond` application. For Junos OS Evolved, to configure traces for a severity other than `info` for the `imond` application, include the `application imond node` *`node-name`* `level` *`severity`* statement at the [edit system trace] hierarchy level.

> **NOTE**: For general monitoring and troubleshooting of devices running Junos OS or Junos OS Evolved, we recommend using standard tools such as CLI show commands, system log messages, SNMP, and telemetry data. You should avoid using trace messages for general debugging purposes and long-term solutions because they are subject to change without notice.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 19.4R1.

RELATED DOCUMENTATION

Understanding Inline Monitoring Services

traceoptions (Interfaces)

IN THIS SECTION

- Syntax | 29588
- Hierarchy Level | 29588
- Description | 29588
- Options | 29588
- Required Privilege Level | 29588
- Release Information | 29588

Syntax

```
traceoptions
```

Hierarchy Level

```
[edit interfaces interface-name traceoptions]
```

Description

Define tracing operations for individual interfaces. To specify more than one tracing operation, include multiple flag statements.

Options

flag - Tracing parameters

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

traceoptions (Juniper Secure Connect)

IN THIS SECTION

- Syntax | 29589
- Hierarchy Level | 29589
- Description | 29589
- Options | 29589
- Required Privilege Level | 29590
- Release Information | 29590

Syntax

```
traceoptions {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
    flag name;
    level (brief | detail | extensive | verbose);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security remote-access]
```

Description

Configure remote access tracing operations for Juniper Secure Connect. By default, messages are written to **/var/log/*file-name*** file. The default file name if not configured is ravpn_trace.

Options

filename Name of file in which to write trace information.

files Maximum number of trace files.

- **Default:** 3
- **Range:** 2 through 1000

match Regular expression for lines to be logged.

no-world-readable Don't allow any user to read the log file.

size Maximum trace file size.

- **Default:** 128k
- **Range:** 10 KB through the maximum file size supported on your system.

world-readable Allow any user to read the log file.

flag *flag*—Tracing operation to perform Tracing operation to perform.

- Values:
  - all—Trace everything
  - cli-configuration—Trace CLI configuration events

level *level* Tracing level. The following values are supported:

- brief—Brief debugging output
- detail—Detailed debugging output
- extensive—Extensive debugging output
- verbose—Verbose debugging output

no-remote-trace Disable remote tracing.

Required Privilege Level

trace

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Juniper Secure Connect Administrator Guide

traceoptions (LACP)

IN THIS SECTION

- Syntax | 29591
- Hierarchy Level | 29591
- Description | 29591
- Default | 29592
- Options | 29592
- Required Privilege Level | 29593
- Release Information | 29593

Syntax

```
traceoptions {
    file <filename> <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols lacp]
```

Description

Define tracing operations for the LACP protocol.

Default

If you do not include this statement, no LACP protocol tracing operations are performed.

Options

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, interface process tracing output is placed in the file lacpd.

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the size option.

- **Range:** 2 through 1000
- **Default:** 3 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- all—All LACP tracing operations
- configuration—Configuration code
- packet—Packets sent and received
- process—LACP process events
- protocol—LACP protocol state machine
- routing-socket—Routing socket events
- startup—Process startup events

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file*.0. When the *trace-file* again reaches its maximum size, *trace-file*.0 is renamed *trace-file*.1 and *trace-file* is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option:

- **Syntax:** *x*`k` to specify kilobytes, *x*`m` to specify megabytes, or *x*`g` to specify gigabytes
- **Range:** 10 KB through the maximum file size supported on your router
- **Default:** 1 MB

`world-readable`—(Optional) Allow any user to read the log file.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Configuring LACP for Aggregated Ethernet Interfaces

traceoptions (Layer 2 Learning)

IN THIS SECTION

- Syntax | 29594
- Hierarchy Level | 29594
- Description | 29594
- Default | 29594
- Options | 29594
- Required Privilege Level | 29596
- Release Information | 29596

Syntax

```
traceoptions  {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag (detail | disable | receive | send);
    in-memory-debug;
    level;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

Define tracing operations for Layer 2 learning.

Default

The `traceoptions` feature is disabled by default.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

You can specify the following options:

- `no-world-readable`—(Optional) Restrict file access to the user who created the file.
- `size` *`size`* —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `files` option. Use *x*`k` to specify KB, *x*`m` to specify MB, or *x*`g` to specify gigabytes.
- `world-readable`—(Optional) Enable unrestricted file access.

`flag` *`flag`* —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- `all`—All tracing operations.
- `bmac-next-hop`—Trace backbone MAC next hop operations.
- `bridge-bmac-next-hop`—Trace backbone MAC next hop bridge operations.
- `bridging-interface`—Trace interface bridge operations.
- `bridging-domain`—Trace bridging domain operations.
- `configuration`—Trace configuration operations.
- `flood-next-hop`—Trace flood next hop operations.
- `initialization`—Trace initialization operations.
- `interface-device`—Trace interface device operations.
- `interface-family`—Trace interface family operations.
- `interface-logical`—Trace logical interface operations.
- `ipc`—Trace inter-process communications operations.
- `irb`—Trace integrated routing and bridging operations.
- `isid`—Trace i-tagged service ID operations.
- `kack`—Trace kernel-acknowledgment.
- `learning-domain`—Trace learning domain operations.
- `logical-system`—Trace logical system operations.
- `mac-learning`—Trace MAC address learning.
- `mc-ae`—Trace multichassis aggregated Ethernet interface operations.
- `redundant-trunk-group`—Trace redundant trunk group operations.
- `routing-instance`—Trace routing instance operations.
- `routing-socket`—Trace routing socket operations.
- `storm-control`—Trace storm control operations.
- `unknown-unicast-forwarding`—Trace unknown unicast forwarding events.

- `vpls-ping`—Trace Virtual Private VLAN Service (VPLS) ping operations.

`in-memory-debug`—Enable trace parameters in the memory.

`level`—Specify level of debugging output.

`no-remote-trace`—Disable remote tracing.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

traceoptions (LLDP)

IN THIS SECTION

- Syntax | 29596
- Hierarchy Level | 29597
- Description | 29597
- Default | 29597
- Options | 29597
- Required Privilege Level | 29599
- Release Information | 29599

Syntax

```
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable | no-world-readable>;
    flag flag <disable>;
}
```

Hierarchy Level

```
[edit protocols lldp],
[edit routing-instances routing-instance-name protocols lldp]
```

Description

Define tracing operations for the Link Layer Discovery Protocol (LLDP). You can trace messages under LLDP for LLDP and physical topology SNMP MIBs.

NOTE: The traceoptions statement is not supported on the QFX3000 QFabric system.

Default

The default LLDP protocol-level trace options are inherited from the global `traceoptions` statement.

Options

disable (Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

file *filename* Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory `/var/log`. We recommend that you place spanning-tree protocol tracing output in the file `/var/log/stp-log`.

files *number* (Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 1 trace file only

flag Specify a tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

- **Values:** The following are the LLDP-specific tracing options:
  - `all`—Trace all operations.
  - `configuration`—Log configuration events.
  - `interface`—Trace interface update events.
  - `packet`—Trace packet events.
  - `protocol`—Trace protocol information.
  - `rtsock`—Trace socket events.
  - `snmp`—Trace SNMP configuration operations.
  - `vlan`—Trace VLAN update events.

  The following are the global tracing options:

  - `all`—All tracing operations.
  - `config-internal`—Trace configuration internals.
  - `general`—Trace general events.
  - `normal`—All normal events. This is the default. If you do not specify this option, only unusual or abnormal operations are traced.
  - `parse`—Trace configuration parsing.
  - `policy`—Trace policy operations and actions.
  - `regex-parse`—Trace regular-expression parsing.
  - `route`—Trace routing table changes.
  - `state`—Trace state transitions.
  - `task`—Trace protocol task processing.
  - `timer`—Trace protocol task timer processing.

no-world-readable (Optional) Prevent any user from reading the log file. This is the default. If you do not include this option, tracing output is appended to an existing trace file.

size *maximum-file-size* (Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the *`trace-`*

file again reaches its maximum size, *trace-file*.0 is renamed *trace-file*.1 and *trace-file* is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the files option.

- **Syntax:** *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

world-readable (Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Configuring LLDP-MED (CLI Procedure)
- Understanding LLDP and LLDP-MED on EX Series Switches
- Understanding LLDP
- Tracing LLDP Operations

traceoptions (MACsec)

IN THIS SECTION

Syntax | 29600

- Hierarchy Level | 29600
- Description | 29600
- Default | 29600
- Options | 29601
- Required Privilege Level | 29601
- Release Information | 29601

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit security macsec]
```

Description

Define tracing operations at the MACsec level. Tracing operations provide support for debugging protocol-level issues. MACsec is an industry-standard security technology that provides secure communication for almost all types of traffic on Ethernet links. To specify more than one tracing operation, include multiple `flag` statements.

The interfaces `traceoptions` statement does not support a separate trace file. The logging is done by the kernel, so the tracing information is placed in the `syslog` file in the directory `/var/log/dcd`.

Default

If you do not include this statement, no tracing operations are performed.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. By default, interface process tracing output is placed in the directory. If you do not specify the name of the trace file, all files are placed in the directory `/var/log/dcd`.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches the maximum value, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. Values range from 2 through 1000.

flag *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. The following are the tracing operation options:

`all`	Trace all operations.
`config`	Trace configuration messages.
`debug`	Trace debug messages.
`normal`	Trace normal messages.

`no-world-readable`—(Optional) Prevent any user from reading the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named `trace-file` reaches this size, it is renamed `trace-file.0`. When `trace-file` again reaches its maximum size, `trace-file.0` is renamed `trace-file.1` and `trace-file` is renamed `trace-file.0`. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

`world-readable`—(Optional) Allow any user to read the log file.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

traceoptions (MACsec interfaces)

IN THIS SECTION

- Syntax | 29602
- Hierarchy Level | 29602
- Description | 29602
- Default | 29603
- Options | 29603
- Required Privilege Level | 29604
- Release Information | 29604

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit security macsec interfaces interface-name]
```

Description

Define tracing operations for individual MACsec interfaces. Tracing operations provide support for debugging protocol-level issues. MACsec is an industry-standard security technology that provides secure communication for almost all types of traffic on Ethernet links. To specify more than one tracing operation, include multiple `flag` statements.

The interfaces `traceoptions` statement does not support a separate trace file. The logging is done by the kernel, so the tracing information is placed in the system `syslog` file in the directory `/var/log/dcd`.

> **NOTE**: Interface level tracing options cannot be enabled when the connectivity association is configured on the sub-interfaces.

Default

If you do not include this statement, no tracing operations are performed.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. By default, interface process tracing output is placed in the directory. If you do not specify the name of the tracefile, all files are placed in the directory `/var/log/dcd`.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named trace-file reaches the maximum value, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. Values range from 2 through 1000.

flag *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. The following are the tracing operation options:

`all` Trace all operations.

`keys` Trace key creation or generation information.

`mka-packets` Trace MACsec Key Agreement (MKA) protocol input and output packet information.

`normal` Trace all normal events and messages.

`state` Trace MKA protocol state information.

`to-secy` Trace MKA to security entity state change information.

`no-world-readable`—(Optional) Prevent any user from reading the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named `trace-file` reaches this size, it is renamed `trace-file.0`. When `trace-file` again reaches its maximum size, `trace-file.0` is renamed `trace-file.1` and `trace-file` is renamed `trace-file.0`. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

`world-readable`—(Optional) Allow any user to read the log file.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

traceoptions (Multicast Snooping Options)

IN THIS SECTION

- Syntax | 29604
- Hierarchy Level | 29605
- Description | 29605
- Default | 29605
- Options | 29605
- Required Privilege Level | 29606
- Release Information | 29606

Syntax

```
traceoptions {
    file filename<files number> <size size> <world-readable | no-world-readable>;
    flag flag <disable>;
}
```

Hierarchy Level

```
[edit multicast-snooping-options]
```

Description

Set multicast snooping tracing options.

Default

Tracing operations are disabled.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***name***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place multicast snooping tracing output in the file **/var/log/multicast-snooping-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 1 trace file only

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

The following are the tracing options:

- **all**—All tracing operations
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events.

- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.
- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Multicast Snooping

Example: Configuring Multicast Snooping

Enabling Bulk Updates for Multicast Snooping

Example: Configuring Multicast Snooping

traceoptions (MVRP)

IN THIS SECTION

- Syntax | 29607
- Hierarchy Level | 29607
- Description | 29608
- Default | 29608
- Options | 29608
- Required Privilege Level | 29609
- Release Information | 29609

Syntax

```
traceoptions {
    file name <size size> <files number> <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mvrp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
mvrp] (for virtual switch instance type),
[edit protocols mvrp],
[edit routing-instances routing-instance-name protocols mvrp] (for virtual switch instance type)
```

Description

For Multiple VLAN Registration Protocol (MVRP), configure tracing options.

Default

Traceoptions is disabled.

Options

disable (Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename* Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. To include the `file` statement, you must specify a filename. Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place MVRP tracing output in the file **/var/log/mvrp-log**.

files *number* (Optional) Maximum number of trace files, in the range from 2 through 1000. The default is 1 trace file. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

flag *flag* Specify which tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- **all**—Enable all trace options flags.
- **error**—Trace all failure conditions.
- **events**—Trace process state change and cleanup events.
- **pdu**—Trace RAPS PDU reception and transmission.
- **socket**—Trace socket activity.
- **state-machine**—Trace information about the state machine.
- **timers**—Trace protocol timers.

no-world-readable (Optional) Prevent any user from reading the log file.

size *size* (Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0.** When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0** . This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten. The file size range is from 10240 through 4294967295. The default file size is 1 MB.

world-readable (Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers

Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration

Verifying That MVRP Is Working Correctly

Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

traceoptions (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 29610
- Hierarchy Level | 29610
- Description | 29610
- Default | 29610
- Options | 29610

Required Privilege Level | 29612
Release Information | 29612

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression> <no-stamp> <replace> <size size>
<world-readable | no-world-readable>;
    flag flag <detail> <disable> <receive> <send>;
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Define tracing operations for the MX Series Virtual Chassis configuration.

Default

Tracing operations are disabled.

Options

detail—(Optional) Generate detailed trace information for a flag.

disable—(Optional) Disable a flag.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files. You can specify 2 through 1000 trace files. The default is 3 trace files.

> **NOTE**: When you specify the maximum number of trace files, you must also specify the maximum trace file size with the **size** *size* option.

When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.
- **auto-configuration**—Trace Virtual Chassis ports that have been automatically configured.
- **csn**—Trace Virtual Chassis complete sequence number (CSN) packets.
- **error**—Trace Virtual Chassis errored packets.
- **graceful-restart**—Trace Virtual Chassis graceful restart events.
- **heartbeat**—Trace Virtual Chassis heartbeat packets.
- **hello**—Trace Virtual Chassis hello packets.
- **krt**—Trace Virtual Chassis kernel routing table (KRT) events.
- **lsp**—Trace Virtual Chassis link-state packets.
- **lsp-generation**—Trace Virtual Chassis link-state packet generation.
- **me**—Trace Virtual Chassis primary-role election (ME) events.
- **normal**—Trace normal events.
- **packets**—Trace Virtual Chassis packets.
- **parse**—Trace reading of the configuration.
- **psn**—Trace partial sequence number (PSN) packets.
- **route**—Trace Virtual Chassis routing information.
- **spf**—Trace Virtual Chassis shortest-path-first (SPF) events.
- **state**—Trace Virtual Chassis state transitions.
- **task**—Trace Virtual Chassis task operations.

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-stamp—(Optional) Do not place a timestamp on any trace file.

no-world-readable—(Optional) Restrict file access to the user who created the file.

receive—(Optional) Trace received packets.

replace—(Optional) Replace a trace file instead of appending information to it.

send—(Optional) Trace transmitted packets.

size ***size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can indicate the unit of measure with ***size*** **k** to specify KB, ***size*** **m** to specify MB, or ***size*** **g** to specify GB. You can specify a trace file size of 10240 through 1073741824 bytes.

> **NOTE**: If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** ***number*** option.

vccpd— (Optional) Include VCCPD messages in the trace files.

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Preprovisioned Member Information for a Virtual Chassis

Tracing Virtual Chassis Operations for MX Series 5G Universal Routing Platforms

Example: Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis

traceoptions (NETCONF and Junos XML Protocol)

IN THIS SECTION

- Syntax | 29613
- Hierarchy Level | 29613
- Description | 29613
- Default | 29614
- Options | 29614
- Required Privilege Level | 29615
- Release Information | 29615

Syntax

```
traceoptions {
    file <filename> <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
    flag flag;
    no-remote-trace;
    on-demand;
}
```

Hierarchy Level

```
[edit system services netconf]
```

Description

Define tracing operations for NETCONF and Junos XML protocol sessions.

NOTE: Starting in Junos OS Release 16.1, when you enable tracing operations at the `[edit system services netconf traceoptions]` hierarchy, Junos OS enables tracing operations

for both NETCONF and Junos XML protocol sessions and adds the `[NETCONF]` and `[JUNOScript]` tags to the log file entries to distinguish the type of session. Prior to Junos OS Release 16.1, only NETCONF session data was logged, and the `[NETCONF]` tag was omitted.

Default

If you do not include this statement, NETCONF and Junos XML protocol-specific tracing operations are not performed.

Options

`file` *`filename`* Name of the file in which to write trace information. All files are placed in the **/var/log** directory.

- **Default: /var/log/netconf**

`files` *`number`* (Optional) Maximum number of trace files.

When a trace file named ***trace-file*** reaches its maximum size, it is renamed and compressed to ***trace-file*.0.gz**. When ***trace-file*** again reaches its maximum size, ***trace-file*.0.gz** is renamed ***trace-file*.1.gz**, and ***trace-file*** is renamed and compressed to ***trace-file*.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000 files
- **Default:** 10 files

`flag` *`flag`* Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Log all incoming and outgoing data from NETCONF and Junos XML protocol sessions.
- `debug`—Log debug level information. Using the `flag all` option is recommended.
- `incoming`—Log all incoming data from NETCONF and Junos XML protocol sessions.
- `outgoing`—Log all outgoing data from NETCONF and Junos XML protocol sessions.

match *regular-expression* (Optional) Refine the output to include only those lines that match the regular expression.

no-remote-trace (Optional) Disable remote tracing.

no-world-readable (Optional) Disable unrestricted file access, which restricts file access to the owner. This is the default.

on-demand (Optional) Enable on-demand tracing, which requires that you start and stop tracing operations from within the NETCONF or Junos XML protocol session. If configured, tracing operations are performed for a session only when requested through the `<request-netconf-trace>` operation.

Within a session, issue the `<request-netconf-trace><start/></request-netconf-trace>` RPC to start tracing operations for that session, and issue the `<request-netconf-trace><stop/></request-netconf-trace>` RPC to stop tracing operations for that session.

size *size* (Optional) Maximum size of each trace file in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). If you don't specify a unit, the default is bytes. If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

- **Syntax:** *size* to specify bytes, *size*k to specify KB, *size*m to specify MB, or *size*g to specify GB
- **Range:** 10240 through 1073741824 bytes
- **Default:** 128 KB

world-readable (Optional) Enable unrestricted file access.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

Support for Junos XML protocol sessions added in Junos OS Release 16.1.

Option `flag debug` introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

NETCONF and Junos XML Protocol Tracing Operations Overview

Example: Tracing NETCONF and Junos XML Protocol Session Operations

netconf | 11290

traceoptions (NETCONF TLS)

IN THIS SECTION

- Syntax | 29616
- Hierarchy Level | 29616
- Description | 29617
- Default | 29617
- Options | 29617
- Required Privilege Level | 29618
- Release Information | 29619

Syntax

```
traceoptions {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
    flag name;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system services netconf tls]
```

Description

Enable trace options for NETCONF sessions that use the Transport Layer Security (TLS) protocol.

Default

If you do not include this statement, NETCONF-over-TLS-specific tracing operations are not performed.

Options

`file` *`filename`*

Name of the file to receive the output of the tracing operation. All files are placed in the **/var/log** directory.

- **Default: /var/log/netconf-tls**

`files` *`number`*

(Optional) Maximum number of trace files.

When a trace file named ***trace-file*** reaches its maximum size, it is renamed and compressed to ***trace-file*.0.gz**, then ***trace-file*.1.gz** and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000 files
- **Default:** 3 files

`flag` *`flag`*

Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

- Values:
  - all—Log all communication
  - app—Log the application data in plain text
  - general—Log tls-proxyd process-related messages
  - pki—Log PKI-related messages
  - plugin—Log plugin messages

level

Level of debugging output.

- Values:
  - all—Match all levels
  - error—Match error conditions
  - info—Match informational messages
  - notice—Match conditions that should be handled specially
  - verbose—Match verbose messages
  - warning—Match warning messages
- **Default:** error

`match` *`regular-expression`* (Optional) Refine the output to include only those lines that match the regular expression.

`no-remote-trace` Disable remote tracing.

`no-world-readable` (Optional) Disable unrestricted file access, which restricts file access to the owner. This is the default.

`size` *`size`* (Optional) Maximum trace file size in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB).

If you don't specify a unit, the default is bytes. If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

- **Syntax:** *`size`* to specify bytes, *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10,240 through 1,073,741,824 bytes
- **Default:** 128 KB

`world-readable` (Optional) Enable unrestricted file access.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

Statement introduced in Junos OS Evolved Release 21.4R1.

traceoptions (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | 29619
- Hierarchy Level | 29619
- Description | 29620
- Options | 29620
- Required Privilege Level | 29620
- Release Information | 29620

Syntax

```
traceoptions {
    apply-groups group-names;
    apply-groups-except group-names;
    flag name;
    file filename,
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services rtlog]
```

Description

Specify the trace information you want to include in the system log messages.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system

Release Information

Support introduced in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

traceoptions (Next Gen Services Service-Set Flow)

IN THIS SECTION

- Syntax | 29621
- Hierarchy Level | 29621
- Description | 29621
- Options | 29621
- Required Privilege Level | 29624
- Release Information | 29624

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
    packet-filter filter-name {
        conn-tag session-conn
        destination-port port-identifier;
        destination-prefix address;
        interface interface-name;
        protocol protocol-identifier;
        source-port port-identifier;
        source-prefix address;
    }
    rate-limit messages-per-second;
    trace-level (brief | detail | error);
}
```

Hierarchy Level

```
[edit services service-set name flow]
```

Description

Configure flow tracing options for a service-set.

Options

file Configure the trace file options.

filename

Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.

files *number*

Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000 files
- **Default:** 10 files

match *regular-expression*

Refine the output to include lines that contain the regular expression.

size *maximum-file-size*

Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

- **Range:** 0 KB through 1 GB
- **Default:** 128 KB

world-readable | no-world-readable

By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

flag

Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.

all

Trace with all flags enabled

basic-datapath	Trace basic packet flow activity
fragmentation	Trace IP fragmentation and reassembly events
high-availability	Trace flow high-availability information
host-traffic	Trace flow host traffic information
multicast	Trace multicast flow information
route	Trace route lookup information
session	Trace session creation and deletion events
session-scan	Trace session scan information
tcp-basic	Trace TCP packet flow information
tunnel	Trace tunnel information

no-remote-trace Set remote tracing as disabled.

packet-filter *filter-name* Packet filter to enable during the tracing operation. Configure the filtering options.

destination-port *port-identifier*	Match TCP/UDP destination port
destination-prefix *address*	Destination IP address prefix
interface *interface-name*	Logical interface
protocol *protocol-identifier*	Match IP protocol type
source-port *port-identifier*	Match TCP/UDP source port
source-prefix *address*	Source IP address prefix

rate-limit *messages-per-second* Limit the incoming rate of trace messages.

trace-level Set the level for trace logging. This option is available only when the flag is set.

brief Trace key flow information, such as message types sent between SPU and central point, policy match, and packet drop reasons.

detail Trace extensive flow information, such as detailed information about sessions and fragments. Detail is the default level.

error Trace error information, such as system failure, unknown message type, and packet drop.

Required Privilege Level

trace—To view this in the configuration.

trace-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

traceoptions (Next Gen Services Softwires)

IN THIS SECTION

- Syntax | 29624
- Hierarchy Level | 29625
- Description | 29625
- Options | 29625
- Required Privilege Level | 29626
- Release Information | 29626

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
```

```
        (no-world-readable | world-readable);
        size maximum-file-size;
    }
    flag (all | configuration | flow);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security softwires]
```

Description

Configure softwire tracing options.

Options

- **file**—Configure trace file information.
  - ***filename***—Name of the file to which to write the trace information.
  - **files *number***—Maximum number of trace files.

    Range: 2 through 1000 files
  - **match *regular-expression***—Regular expression for lines to be logged.
  - **no-world-readable | world-readable**—Allow or deny any user to read the log file.
  - **size *maximum-file-size***—Maximum trace file size.

    Range: 10,240 to 1,073,741,824 bytes
- **flag**—Specify events to trace.
  - **all**—Trace all events
  - **configuration**—Trace configuration events
  - **flow**—Trace flow events
- **no-remote-trace**—Disable remote tracing.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Release 12.1 of Junos OS.

traceoptions (Origin Validation for BGP)

IN THIS SECTION

- Syntax | 29626
- Hierarchy Level | 29626
- Description | 29627
- Required Privilege Level | 29627
- Release Information | 29627

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation],
[edit logical-systems logical-system-name routing-instances instance-name routing-options
validation group group-name session address],
[edit logical-systems logical-system-name routing-options validation],
```

```
[edit logical-systems logical-system-name routing-options validation group group-name session
address],
[edit routing-instances instance-name routing-options validation],
[edit routing-instances instance-name routing-options validation group group-name session
address],
[edit routing-options validation]
[edit routing-options validation group group-name session address]
```

Description

Configure tracing operations for resource public key infrastructure (RPKI) BGP route validation.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

traceoptions (Outbound SSH)

IN THIS SECTION

- Syntax | 29628
- Hierarchy Level | 29628

- Description | 29628
- Options | 29628
- Required Privilege Level | 29629
- Release Information | 29629

Syntax

```
traceoptions {
    file <filename> <files number> <match regular-expression> <size size> <(world-readable | no-
world-readable)>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system services outbound-ssh]
```

Description

Set the trace options for the outbound SSH service. By default, tracing operations are disabled.

Options

- `file`—Configure the trace file information.
  - `filename`—(Optional) By default, the name of the file is the name of the process being traced. Use this option to override the default file name and specify a file to receive the output of the tracing operation. Enclose the name within quotation marks. All trace files are placed in the directory `/var/log`.
  - `files` *`number`*—(Optional) Specify the maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a file name with the `filename` option.

    Range: 2 through 1000 files

Default: 3 files

- `match` *`regular-expression`*—(Optional) When configured, the system adds only those lines to the trace file that match the regular expression. For example, if the regular expression is set to `=error`, the system only adds lines to the trace file that include the string `error`.

- `size` *`maximum-file-size`*—(Optional) Specify the maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named `trace-file` reaches this size, it is renamed `trace-file.0`. When `trace-file` again reaches its maximum size, `trace-file.0` is renamed `trace-file.1` and `trace-file` is renamed `trace-file.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum file size, you also must specify a maximum number of trace files by using the `files` option and a filename by using the `file` option.

  Syntax: x K to specify KB, x m to specify MB, or x g to specify GB

  Range: 10 KB through 1 GB

  Default: 128 KB

- `world-readable | no-world-readable`—(Optional) By default, access to the trace files is restricted to the user who configured the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, configure the `no-world-readable` option.

- `flag`—Specify the tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

  - `all`—Trace all events.

  - `configuration`—Trace configuration events.

  - `connectivity`—Trace TCP connection handling between the management application and the device.

- `no-remote-trace`—(Optional) Disable remote tracing and logging operations that track normal operations, error conditions, and packets that are generated by or passed through the device.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Displaying Log and Trace Files

traceoptions (OVSDB)

IN THIS SECTION

- Syntax | 29630
- Hierarchy Level | 29630
- Description | 29630
- Default | 29631
- Options | 29631
- Required Privilege Level | 29632
- Release Information | 29632

Syntax

```
traceoptions {
    file <filename> <files number> <match regular-expression> <no-world-readable | world-
readable> <size size>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols ovsdb]
```

Description

Define tracing operations for the *Open vSwitch Database (OVSDB)* management protocol, which is supported on Juniper Networks devices.

Default

If you do not include this statement, OVSDB-specific tracing operations are not performed.

Options

file *filename*

Name of file in which the system places the output of the tracing operations. By default, the system places all files in the **/var/log** directory.

- **Default: /var/log/vgd**

files *number*

(Optional) Maximum number of trace files. When a trace file reaches the size specified by the `size` option, the filename is appended with 0 and compressed. For example, a trace file named **trace-file.gz** would be renamed **trace-file.0.gz**. When **trace-file.0.gz** reaches the specified size, it is renamed **trace-file.1.gz** and its contents are compressed to **trace-file.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000 files
- **Default:** 10 files

flag *flag*

Tracing operation to perform. You can include one or more of the following flags:

all All OVSDB events.

configuration OVSDB configuration events.

core OVSDB core events.

function OVSDB function events.

interface OVSDB interface events.

l2-client OVSDB Layer 2 client events.

netconf-client (QFX Series switches only) Events for the dynamic configuration of Virtual Extensible LANs (VXLANs).

ovs-client OVSDB client events.

match *regular-expression*
(Optional) Only log lines that match the regular expression.

no-remote-trace
(Optional) Disable tracing and logging operations that track normal operations, error conditions, and packets that are generated by or passed through the Juniper Networks device.

no-world-readable
Restrict access to the trace files to the owner.

- **Default:** no-world-readable

size *size*
(Optional) Maximum size of each trace file in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). If you do not specify a unit, the default is bytes. If you specify a maximum file size, you also must specify a maximum number of trace files by using the `files` option and a filename by using the `file` option.

- **Syntax:** *size* to specify bytes, *size*k to specify KB, *size*m to specify MB, or *size*g to specify GB.
- **Range:** 10,240 through 1,073,741,824 bytes
- **Default:** 128 KB

world-readable
Enable any user to access the trace files.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

traceoptions (PCE)

IN THIS SECTION

- Syntax | 29633
- Hierarchy Level | 29633
- Description | 29633
- Options | 29633
- Required Privilege Level | 29634

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag (all | pcep);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Description

Configure the Path Computation Element Protocol (PCEP) tracing options.

Options

filename Name of the file to receive the output of the tracing operation. All files are placed in the directory `/var/log`.

`files` *number* (Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000 files
- **Default:** 2 files.

  If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size.

flag Area of path computation client process (pccd) to enable debugging output.

- all—Trace all areas of PCD code.
- pcep—Trace Path Computation Element Protocol (PCEP) operations.

no-remote-trace (Optional) Disable remote tracing options.

no-world-readable (Optional) Allow only certain users to read the log file.

`size` *size* (Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches this size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB.
- **Range:** 10 KB through the maximum file size supported on your system.
- **Default:** 1 MB.

  If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of files.

world-readable (Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

pce | 12303

traceoptions (PCEF)

IN THIS SECTION

- Syntax | 29635
- Hierarchy Level | 29635
- Description | 29635
- Options | 29635
- Required Privilege Level | 29637
- Release Information | 29637

Syntax

```
traceoptions {
    file file-name <files number> <no-word-readable | world-readable> <size size>;
    flag flag;
    level (all | critical | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit unified-edge pcef]
```

Description

Specify tracing options for policy and charging enforcement functions (PCEF).

Options

file *file-name* Use the specified name of the file to receive the output of the tracing operation.

files *number*

(Optional) Use the specified maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000
- **Default:** 3 files

flag *flag*

Specify which operations are to be traced. To specify more than one operation, include multiple flag statements.

> **BEST PRACTICE**: You might want to enable traceoptions only when you want to debug specific charging operations. Enabling the traceoption flags might have an impact on the system performance.

- `all`—Trace all operations.
- `config`—Trace configuration events.
- `debug`—Trace debug internal events.
- `fsm`—Trace finite state machine events.
- `general`—Trace general events that do not fit in any specific traces.
- `high-availability`—Trace high-availability events.
- `init`—Trace initialization events.
- `tftmgr`—Trace tftmgr events.

level

Use the specified level of tracing. You can specify any of the following levels:

- `all`—Match all levels.
- `critical`—Match critical conditions.
- `error`—Match error conditions.
- `info`—Match informational messages
- `notice`—Match conditions that must be handled specially.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

no-remote-trace (Optional) Disable remote tracing.

no-world-readable (Optional) Disable unrestricted file access.

size *size* (Optional) Use the specified maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the size option.

- **Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB.
- **Range:** 10,240 through 1,073,741,824 bytes
- **Default:** 128 KB

word-readable (Optional) Enable unrestricted file access.

Required Privilege Level

trace and unified-edge—To view this statement in the configuration.

trace-control and unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Tracing for PCEF Operations

traceoptions (PIM Snooping)

IN THIS SECTION

- Syntax | 29638
- Hierarchy Level | 29638
- Description | 29638
- Default | 29638
- Options | 29639
- Required Privilege Level | 29640
- Release Information | 29640

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit routing-instances <instance-name> protocols pim-snooping],
[edit logical-systems <logical-system-name> routing-instances <instance-name> protocols pim-
snooping]
```

Description

Define tracing operations for PIM snooping.

Default

The `traceoptions` feature is disabled by default.

The default PIM trace options are those inherited from the routing protocol's `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

PIM Snooping Tracing Flags:

- `all`—All tracing operations.
- `general`—Trace general PIM snooping events.
- `hello`—Trace hello packets.
- `join`—Trace join messages.
- `normal`—Trace normal PIM snooping events. If you do not specify this flag, only unusual or abnormal operations are traced.
- `packets`—Trace all PIM packets.
- `policy`—Trace policy processing.
- `prune`—Trace prune messages.
- `route`—Trace routing information.
- `state`—Trace PIM state transitions.
- `task`—Trace PIM protocol task processing.
- `timer`—Trace PIM protocol timer processing.

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers per flag:

- `detail`—Provide detailed trace information.
- `disable`—Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.
- `receive`—Packets being received.
- `send`—Packets being transmitted.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

PIM Snooping for VPLS

traceoptions (Port Mirroring and Traffic Sampling)

IN THIS SECTION

- Syntax | 29640
- Hierarchy Level | 29641
- Description | 29641
- Required Privilege Level | 29641
- Release Information | 29641

Syntax

```
traceoptions {
    file filename  <files number> <size bytes> <world-readable | no-world-readable>;
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring],
[edit forwarding-options sampling]
```

Description

Configure traffic sampling tracing operations.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing Traffic-Sampling Operations

traceoptions (PPP Process)

IN THIS SECTION

- Syntax | 29642
- Hierarchy Level | 29642
- Description | 29642
- Default | 29642
- Options | 29642

Required Privilege Level | 29644
Release Information | 29644

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
    flag flag;
    level severity-level;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols ppp]
```

Description

Define tracing operations for the PPP process.

To specify more than one tracing operation, include multiple `flag` statements.

You cannot specify a separate trace tile. Tracing information is placed in the system `syslog` file in the directory `/var/log/pppd`.

Default

If you do not include this statement, no PPPD-specific tracing operations are performed.

Options

`filename`—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, commit script process tracing output is placed in the file `ppd`. If you include the `file` statement, you must specify a filename. To retain the default, you can specify `eventd` as the filename.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000
- **Default:** 3 files

`disable`—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

`flag`—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. The following are the PPPD-specific tracing options.

- `access`—Access code
- `address-pool`—Address pool code
- `all`—All areas of code
- `auth`—Authentication code
- `chap`—Challenge Handshake Authentication Protocol (CHAP) code
- `config`—Configuration code
- `ifdb`—Interface database code
- `lcp`—LCP state machine code
- `memory`—Memory management code
- `message`—Message processing code
- `mlppp`—Trace MLPPP code
- `ncp`—NCP state machine code
- `pap`—Password Authentication Protocol (PAP) code
- `ppp`—PPP protocol processing code
- `radius`—RADIUS processing code
- `rtsock`—Routing socket code
- `session`—Session management code

- `signal`—Signal handling code
- `timer`—Timer code
- `ui`—User interface code

`match` *`regex`*—(Optional) Refine the output to include only those lines that match the given regular expression.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the *`trace-file`* again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and filename.

- **Syntax:** *`x`*`k` to specify KB, *`x`*`m` to specify MB, or *`x`*`g` to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

`non-world-readable`—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. Specify `non-world-readable` to reset the default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configuring PPP

traceoptions (PPPoE)

IN THIS SECTION

- Syntax | 29645
- Hierarchy Level | 29645
- Description | 29645
- Options | 29646
- Required Privilege Level | 29648
- Release Information | 29648

Syntax

```
traceoptions {
    file <filename> <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    filter {
        aci regular-expression;
        ari regular-expresion;
        service-name regular-expresion;
        underlying-interface interface-name;
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols pppoe]
```

Description

Define tracing operations for PPPoE processes.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`disable`—Disable this trace flag.

`filter`—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.

> **BEST PRACTICE**: Due to the complexity of agent circuit identifiers and agent remote identifiers, we recommend that you do not try an exact match when filtering on these options. For service names, searching on the exact name is appropriate, but you can also use a regular expression with that option.

- `aci` *`regular-expression`*—Regular expression to match the agent circuit identifier provided by PPPoE client.
- `ari` *`regular-expression`*—Regular expression to match the agent remote identifier provided by PPPoE client.
- `service` *`regular-expression`*—Regular expression to match the name of PPPoE service.
- `underlying-interface` *`interface-name`*—Name of a PPPoE underlying interface. You cannot use a regular expression for this filter option.

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.
- `config`—Trace configuration events.
- `events`—Trace events.
- `gres`—Trace GRES events.
- `init`—Trace initialization events.
- `interface-db`—Trace interface database operations.

- `memory`—Trace memory processing events.
- `protocol`—Trace protocol events.
- `rtsock`—Trace routing socket events.
- `session-db`—Trace connection events and flow.
- `signal`—Trace signal operations.
- `state`—Trace state handling events.
- `timer`—Trace timer processing.
- `ui`—Trace user interface processing.

`level`—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Option `filter` introduced in Junos OS Release 12.3

RELATED DOCUMENTATION

Configuring PPPoE Service Name Tables

Tracing PPPoE Operations

traceoptions (Process Monitor)

IN THIS SECTION

- Syntax | 29648
- Hierarchy Level | 29649
- Description | 29649
- Options | 29649
- Required Privilege Level | 29650
- Release Information | 29651

Syntax

```
traceoptions {
    file filename files number match regex size size (world-readable | no-world-readable);
    flag flag;
    level level;
```

```
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system processes process-monitor]
```

Description

Enable tracing options for the process health monitor process (pmond).

> **NOTE**: Starting with Junos OS Release 15.1R2, the pmond process is enabled by default on the Routing Engines of MX Series routers, even when no service interfaces are configured.

Options

file *filename*
: Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. To include the `file` statement, you must specify a filename.

files *number*
: (Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

flag *flag*
: Specify which tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

  - `all`—Enable all trace options flags.
  - `events`—Trace process state change and cleanup events.
  - `gencfg`—Trace GENCFG blobs recorded for cleanup.
  - `module`—Trace module code.
  - `sysvsem`—Trace SYSV semaphores recorded for cleanup.

- `sysvshm`—Trace SYSV shared memory segments recorded for cleanup.
- `tracking`—Trace tracking code.
- `ui`—Trace user interface operations.

level *level* Specify the level of debugging output:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match conditions that warrant special handling (but are not errors).
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

match *regex* (Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace Disable remote tracing.

size *size* (Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` *number* option.

- **Syntax:** `xk` to specify KB, `xm` to specify MB, or `xg` to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

world-readable | no-world-readable (Optional). Grant all users permission to read log files, or restrict the permission only to the `root` user and users who have the Junos `maintenance` permission.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

traceoptions (programmable-rpd)

IN THIS SECTION

- Syntax | 29651
- Hierarchy Level | 29651
- Description | 29651
- Default | 29652
- Options | 29652
- Required Privilege Level | 29653
- Release Information | 29653

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <disable>;
}
```

Hierarchy Level

```
[edit routing-options programmable-rpd]
```

Description

Starts logging traces related to the programmable routing protocol process (prpd). Use the `traceoptions` statement, along with related show commands, to help debug client-server interactions, identify the

flow of control, and detect errors, get client-level information and statistics. You can filter traces according to the flag(s) you have enabled.

Default

If you do not include this statement, no tracing operations are performed.

Options

file *filename*
Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*
(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the size option.

- **Range:** 2 through 1000 files
- **Default:** 10 files

no-world-readable
(Optional) Prevent any user from reading the log file.

size *size*
(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file*.0. When the *trace-file* again reaches its maximum size, *trace-file*.0 is renamed *trace-file*.1 and *trace-file* is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the files option.

- **Syntax:** *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB
- **Range:** 1024 to 4,294,967,295 bytes
- **Default:** 128 KB

world-readable
(Optional) Allow any user to read the log file.

flag *flag*
Specifies the tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The options are:

- `all`—All tracing operations.
- `client`—Client events.
- `general`—All normal operations and routing table changes (a combination of the `normal` and `route` trace operations).
- `japi`—JAPI-related changes.
- `normal`—All normal operations.
- `policy`—Routing policy operations and actions.
- `route`—Routing table changes.
- `routing-interface`—Routing interface changes.
- `state`—State transitions.
- `task`—Interface transactions and processing.
- `te-path-compute`—Traffic engineering path computation changes.
- `timer`—Timer usage.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

purge-timeout (programmable-rpd) | 13870

show programmable-rpd clients | 23287

show route | 23512

traceoptions (Protocols)

IN THIS SECTION

- Syntax | 29654
- Hierarchy Level | 29654
- Description | 29654
- Default | 29654
- Options | 29655
- Required Privilege Level | 29655
- Release Information | 29656

Syntax

```
traceoptions {
    file name <size size> <files number> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit protocols isis],
[edit protocols (ospf | ospf3)]
```

Description

Define tracing operations that graceful restart functionality in the router or switch.

To specify more than one tracing operation, include multiple `flag` statements.

Default

If you do not include this statement, no global tracing operations are performed.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***name***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place global routing protocol tracing output in the file **routing-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000 files
- **Default:** 2 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. The nonstop active routing tracing option is:

- **graceful-restart**—Tracing operations for nonstop active routing

no-world-readable—Restrict users from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

graceful-restart flag for IS-IS and OSPF/OSPFv3 added in Junos OS Release 8.4.

RELATED DOCUMENTATION

Tracking Graceful Restart Events

traceoptions (Protocols AMT)

IN THIS SECTION

- Syntax | 29656
- Hierarchy Level | 29657
- Description | 29657
- Options | 29657
- Required Privilege Level | 29659
- Release Information | 29659

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols amt],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols amt],
[edit protocols amt],
[edit routing-instances routing-instance-name protocols amt]
```

Description

Configure Automatic Multicast Tunneling (AMT) tracing options.

To specify more than one tracing operation, include multiple `flag` statements.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the file **igmp-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size.

- **Range:** 2 through 1000 files
- **Default:** 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

AMT Tracing Flags

- **errors**—All error conditions
- **packets**—All AMT packets
- **tunnels**—All AMT tunnel-related information

Global Tracing Flags

- **all**—All tracing operations
- **normal**—All normal operations
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

- **Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of trace files.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

traceoptions (Protocols BFD)

IN THIS SECTION

- Syntax | 29659
- Hierarchy Level | 29660
- Description | 29660
- Default | 29660
- Options | 29660
- Required Privilege Level | 29661
- Release Information | 29662

Syntax

```
traceoptions {
    file name <size size> <files number> <world-readable | no-world-readable>;
```

```
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit protocols bfd]
```

Description

Define tracing operations that track unified in-service software upgrade (ISSU) functionality in the router.

To specify more than one tracing operation, include multiple `flag` statements.

Default

If you do not include this statement, no global tracing operations are performed.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***name***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place global routing protocol tracing output in the file **routing-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000 files
- **Default:** 2 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

flag ***flag***—Tracing operation to perform. The tracing options are as follows:

- **adjacency**—Trace adjacency messages.
- **all**—Trace everything.

- **error**—Trace all errors.
- **events**—Trace all events.
- **issu**—Trace ISSU packet activity.
- **nsr-packet**—Trace packet activity of NSR.
- **nsr-synchronization**—Trace NSR synchronization events.
- **packet**—Trace all packets.
- **pipe**—Trace pipe messages.
- **pipe-detail**—Trace pipe messages in detail.
- **ppm-packet**—Trace packet activity by periodic packet management.
- **state**—Trace state transitions.
- **timer**—Trace timer processing.

no-world-readable—Restrict users from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

issu flag for BFD added in Junos OS Release 9.1.

RELATED DOCUMENTATION

Managing and Tracing BFD Sessions During Unified ISSU Procedures

traceoptions (Protocols BGP)

IN THIS SECTION

- Syntax | 29662
- Hierarchy Level | 29663
- Description | 29663
- Default | 29663
- Options | 29663
- Required Privilege Level | 29665
- Release Information | 29665

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Description

Configure BGP protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.

NOTE: The `traceoptions` statement is not supported on QFabric systems.

Default

The default BGP protocol-level tracing options are inherited from the routing protocols `traceoptions` statement included at the `[edit routing-options]` hierarchy level. The default group-level trace options are inherited from the BGP protocol-level `traceoptions` statement. The default peer-level trace options are inherited from the group-level `traceoptions` statement.

Options

`disable`—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

`file` *`name`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place BGP tracing output in the file **bgp-log**.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*.0** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 10 files

`flag`—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

BGP Tracing Flags

- `4byte-as`—4-byte AS events.
- `bfd`—BFD protocol events.
- `damping`—Damping operations.
- `graceful-restart`—Graceful restart events.
- `ipsec`—Events related to IPsec support for BGP peers when the IPsec is configured.
- `keepalive`—BGP keepalive messages. If you enable the the BGP `update` flag only, received keepalive messages do not generate a trace message.
- `nsr-synchronization`—Nonstop routing synchronization events.
- `open`—Open packets. These packets are sent between peers when they are establishing a connection.
- `packets`—All BGP protocol packets.
- `refresh`—BGP refresh packets.
- `update`—Update packets. These packets provide routing updates to BGP systems. If you enable only this flag, received keepalive messages do not generate a trace message. Use the `keepalive` flag to generate a trace message for keepalive messages.

Global Tracing Flags

- `all`—All tracing operations
- `general`—A combination of the `normal` and `route` trace operations
- `normal`—All normal operations
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- `policy`—Policy operations and actions

- `route`—Routing table changes
- `state`—State transitions
- `task`—Routing protocol task processing
- `timer`—Routing protocol timer processing

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- `detail`—Provide detailed trace information.
- `filter`—Provide filter trace information. Applies only to `route`, `damping`, and `update` tracing flags.
- `receive`—Trace the packets being received.
- `send`—Trace the packets being transmitted.

`no-world-readable`—(Optional) Prevent any user from reading the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`x`*`k` to specify KB, *`x`*`m` to specify MB, or *`x`*`g` to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

`world-readable`—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`4byte-as` statement introduced in Junos OS Release 9.2.

`4byte-as` statement introduced in Junos OS Release 9.2 for EX Series switches.

`ipsec` statement introduced in Junos OS Evolved Release 21.3R1.

RELATED DOCUMENTATION

- log-updown | 10042
- *Tracing Nonstop Active Routing Synchronization Events*
- *Understanding Trace Operations for BGP Protocol Traffic*
- *Configuring OSPF Refresh and Flooding Reduction in Stable Topologies*

traceoptions (Protocols BMP)

IN THIS SECTION

- Syntax | 29666
- Hierarchy Level | 29666
- Description | 29667
- Options | 29667
- Required Privilege Level | 29668
- Release Information | 29668

Syntax

```
traceoptions {
    file file-name <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options bmp],
[edit logical-systems logical-system-name routing-options bmp station station-name],
```

```
[edit routing-options bmp],
[edit routing-options bmp station station-name]
```

Description

Configure tracing options for BMP monitoring. To specify more than one tracing operation, include multiple flag statements.

Options

`file` *`file-name`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place BMP tracing output in the file **bmp-log**.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*.0** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 10 files

`flag`—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

- `all`—Trace all BMP monitoring operations.
- `down`—Down messages.
- `error`—Error conditions.
- `event`—Major events, station establishment, errors, and events.
- `general`—General events.
- `normal`—Normal events.
- `packets`—All messages.
- `policy`—Policy processing.
- `route`—Routing information.
- `route-monitoring`—Route monitoring messages.

- `state`—State transitions.
- `statistics`—Statistics messages.
- `task`—Routing protocol task processing.
- `timer`—Routing protocol timer processing.
- `up`—Up messages.
- `write`—Writing of messages.

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- `detail`—Provide detailed trace information.
- `disable`—Disable the tracing flag.
- `receive`—Trace the packets being received.
- `send`—Trace the packets being transmitted.

`no-world-readable`—(Optional) Prevent any user from reading the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`x`*`k` to specify KB, *`x`*`m` to specify MB, or *`x`*`g` to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

`world-readable`—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D15.

RELATED DOCUMENTATION

Tracing BMP Operations

Understanding Trace Operations for BGP Protocol Traffic

Configuring OSPF Refresh and Flooding Reduction in Stable Topologies

traceoptions (Protocols DVMRP)

IN THIS SECTION

- Syntax | 29669
- Hierarchy Level | 29669
- Description | 29670
- Default | 29670
- Options | 29670
- Required Privilege Level | 29672
- Release Information | 29672

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols dvmrp],
[edit protocols dvmrp]
```

Description

Configure DVMRP tracing options.

To specify more than one tracing operation, include multiple `flag` statements.

Default

The default DVMRP trace options are those inherited from the routing protocols `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the **dvmrp-log** file.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size.

- **Range:** 2 through 1000 files
- **Default:** 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

DVMRP Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **graft**—Graft messages
- **neighbor**—Neighbor probe messages
- **normal**—All normal operations
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **packets**—All DVMRP packets
- **poison**—Poison-route-reverse packets
- **probe**—Probe packets
- **prune**—Prune messages
- **report**—DVMRP route report packets
- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

- **Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of trace files.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

> **NOTE**: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing DVMRP Protocol Traffic

traceoptions (Protocols EVPN)

IN THIS SECTION

- Syntax | 29673
- Hierarchy Level | 29673
- Description | 29673
- Options | 29673
- Required Privilege Level | 29675
- Release Information | 29675

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols]
[edit routing-instances routing-instance-name protocols evpn]
```

Description

Trace traffic flowing through an EVPN routing instance.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" ").

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches the maximum size as specified by the `size` option, it is renamed *`trace-file`*`.0`. When *`trace-file`* again reaches the maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 2 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can specify the following tracing flags:

- `all`—All EVPN tracing options
- `error`—Error conditions
- `gbptagdb`—GBP tag database related messages. This option is not available for logical systems.

- `general`—General events
- `mac-database`—MAC route database in the EVPN routing instance
- `nlri`—EVPN advertisements received or sent by means of BGP
- `normal`—Normal events
- `oam`—OAM messages
- `policy`—Policy processing
- `route`—Routing information
- `state`—State transitions
- `task`—Routing protocol task processing
- `timer`—Routing protocol timer processing
- `topology`—EVPN topology changes caused by reconfiguration or advertisements received from other provider edge (PE) routers using BGP

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify the following modifiers:

- `detail`—Provide detailed trace information.
- `disable`—Disable this trace flag.
- `receive`—Trace received packets.
- `send`—Trace sent packets.

`no-world-readable`—Do not allow any user to read the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When *`trace-file`* again reaches the maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files (as specified by the `files` option) is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`x`*`k` to specify kilobytes, *`x`*`m` to specify megabytes, or *`x`*`g` to specify gigabytes
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

`world-readable`—Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2 .

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Tracing EVPN Traffic and Operations

traceoptions (Protocols IGMP)

IN THIS SECTION

- Syntax | 29675
- Hierarchy Level | 29676
- Description | 29676
- Default | 29676
- Options | 29676
- Required Privilege Level | 29678
- Release Information | 29678

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
```

```
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp],
[edit protocols igmp]
```

Description

Configure IGMP tracing options.

To specify more than one tracing operation, include multiple `flag` statements.

To trace the paths of multicast packets, use the `mtrace` command.

Default

The default IGMP trace options are those inherited from the routing protocols `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the file **igmp-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size.

- **Range:** 2 through 1000 files
- **Default:** 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

IGMP Tracing Flags

- **leave**—Leave group messages (for IGMP version 2 only).
- **mtrace**—Mtrace packets. Use the `mtrace` command to troubleshoot the software.
- **packets**—All IGMP packets.
- **query**—IGMP membership query messages, including general and group-specific queries.
- **report**—Membership report messages.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

- **Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of trace files.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing IGMP Protocol Traffic

traceoptions (Protocols IGMP Snooping)

IN THIS SECTION

- Syntax | 29679
- Hierarchy Level | 29679
- Description | 29679

- Default | 29679
- Options | 29680
- Required Privilege Level | 29681
- Release Information | 29681

Syntax

```
traceoptions  {
    file filename <files number> <size size> <world-readable | no-world-readable> ;
    flag flag (detail | disable | receive | send);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name bridge-domains domain-name protocols igmp-snooping],
[edit logical-systems logical-system-name routing-instances instance-name bridge-domains domain-
name protocols igmp-snooping],
[edit logical-systems logical-system-name routing-instances instance-name protocols igmp-
snooping],
[edit bridge-domains domain-name protocols igmp-snooping],
[edit routing-instances instance-name bridge-domains domain-name protocols igmp-snooping],
[edit routing-instances instance-name protocols igmp-snooping]
[edit protocols igmp-snooping vlan]
```

Description

Define tracing operations for IGMP snooping.

Default

The **traceoptions** feature is disabled by default.

Enabling **traceoptions** under the `igmp-snooping` hierarchy enables **igmp-snooping**. Correspondingly, disabling **traceoptions** under the `igmp-snooping` hierarchy disables **igmp-snooping**.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached (***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000
- **Default:** 3 files

flag ***flag*** —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.
- **client-notification**—Trace notifications.
- **general**—Trace general IGMP snooping protocol events.
- **group**—Trace group operations.
- **host-notification**—Trace host notifications.
- **leave**—Trace leave group messages (IGMPv2 only).
- **normal**—Trace normal IGMP snooping protocol events.
- **packets**—Trace all IGMP packets.
- **policy**—Trace policy processing.
- **query**—Trace IGMP membership query messages.
- **report**—Trace membership report messages.
- **route**—Trace routing information.
- **state**—Trace IGMP state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.

no-world-readable—(Optional) Restrict file access to the user who created the file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes
- **Range:** 10 KB through 1 gigabytes
- **Default:** 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring IGMP Snooping Trace Operations

Configuring IGMP Snooping

Example: Configuring IGMP Snooping on SRX Series Devices

IGMP Snooping Overview

igmp-snooping | **8387**

traceoptions (Protocols IPv6 Neighbor Discovery)

IN THIS SECTION

Syntax | **29682**

Hierarchy Level | **29682**

- Description | 29682
- Default | 29682
- Options | 29683
- Required Privilege Level | 29684
- Release Information | 29684

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement],
[edit protocols router-advertisement]
```

Description

For IPv6 neighbor discovery, specify router advertisement protocol-level tracing options.

Trace IPv6 Neighbor Discovery protocol traffic to help debug Neighbor Discovery protocol issues.

Global tracing options are inherited from the configuration set by the `traceoptions` statement at the `[edit routing-options]` hierarchy level. You can override the following global trace options for the IPv6 Neighbor Discovery protocol using the `traceoptions flag` statement included at the `[edit protocols router-advertisement]` hierarchy level:

Default

The default trace options are inherited from the global `traceoptions` statement.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place router advertisement tracing output in the file **/var/log/router-advertisement-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 10 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

- **all**—All tracing operations

> **NOTE**: Use the trace flag **all** with caution as this may cause the CPU to become very busy.

- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations.
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—IPv6 interface transactions and processing
- **timer**—IPv6 neighbor discovery protocol timer processing

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming

scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

traceoptions (Protocols IS-IS)

IN THIS SECTION

- Syntax | 29685
- Hierarchy Level | 29685
- Description | 29685
- Default | 29685
- Options | 29685
- Required Privilege Level | 29687
- Release Information | 29687

Syntax

```
traceoptions {
    file name <size size> <files number> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
isis],
[edit protocols isis],
[edit routing-instances routing-instance-name protocols isis]
```

Description

Configure IS-IS protocol-level tracing options. To specify more than one tracing operation, include multiple `flag` statements.

NOTE: The `traceoptions` statement is not supported on QFabric systems.

Default

The default IS-IS protocol-level tracing options are those inherited from the routing protocols `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

Options

`disable`—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

`file` *name*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks (“ ”). All files are placed in the directory `/var/log`. We recommend that you place IS-IS tracing output in the file `isis-log`.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 10 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one flag, include multiple `flag` statements.

`IS-IS Protocol-Specific Tracing Flags`

- `csn`—Complete sequence number PDU (CSNP) packets
- `error`—Errored IS-IS packets
- `graceful-restart`—Graceful restart operation
- `hello`—Hello packets
- `layer2-map`—Mapped ARP or neighbor discovery next hops in the kernel
- `ldp-synchronization`—Synchronization between IS-IS and LDP
- `lsp`—Link-state PDUs
- `lsp-generation`—Link-state PDU generation packets
- `packets`—All IS-IS protocol packets
- `psn`—Partial sequence number PDU (PSNP) packets
- `spf`—Shortest-path-first calculations
- `srv6`—Segment routing for IPv6.

`Global Tracing Flags`

- `all`—All tracing operations
- `general`—A combination of the `normal` and `route` trace operations
- `normal`—All normal operations, including adjacency changes
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- `policy`—Policy operations and actions

- `route`—Routing table changes
- `state`—State transitions
- `task`—Routing protocol task processing
- `timer`—Routing protocol timer processing

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- `detail`—Provide detailed trace information.
- `receive`—Trace the packets being received.
- `send`—Trace the packets being transmitted.
-

`no-world-readable`—(Optional) Prevent any user from reading the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the *`trace-file`* again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`x`*`k` to specify KB, *`x`*`m` to specify MB, or *`x`*`g` to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

`world-readable`—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces

Example: Enabling Packet Checksums on IS-IS Interfaces for Error Checking

Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces

Understanding Link-State PDU Throttling for IS-IS Interfaces

Understanding Checksums on IS-IS Interfaces for Error Checking

Understanding the Transmission Frequency for CSNPs on IS-IS Interfaces

traceoptions (Protocols Layer 2 Circuit)

IN THIS SECTION

- Syntax | 29688
- Hierarchy Level | 29688
- Description | 29689
- Options | 29689
- Required Privilege Level | 29690
- Release Information | 29690

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit],
[edit protocols l2circuit]
```

Description

Trace traffic flowing through a Layer 2 circuit.

Options

`disable`—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" ").

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 2 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

- `connections`—Layer 2 circuit connections (events and state changes)
- `error`—Error conditions
- `fec`—Layer 2 circuit advertisements received or sent by means of LDP
- `topology`—Layer 2 circuit topology changes caused by reconfiguration or advertisements received from other PE routers

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify the `detail` modifier if you want to provide detailed trace information.

`no-world-readable`—(Optional) Do not allow any user to read the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file*.0. When the *trace-file* again reaches its maximum size, *trace-file*.0 is renamed *trace-file*.1 and *trace-file* is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *x*`k` to specify kilobytes, *x*`m` to specify megabytes, or *x*`g` to specify gigabytes
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

`world-readable`—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing Layer 2 Circuit Operations

traceoptions (Protocols Layer 2 VPN)

IN THIS SECTION

- Syntax | 29691
- Hierarchy Level | 29691
- Description | 29691
- Options | 29691
- Required Privilege Level | 29693
- Release Information | 29693

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn],
[edit routing-instances routing-instance-name protocols l2vpn]
```

Description

Trace traffic flowing through a Layer 2 VPN.

Options

`disable`—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" ").

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 2 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

- `all`—All Layer 2 VPN tracing options
- `connections`—Layer 2 connections (events and state changes)

- `error`—Error conditions
- `general`—General events
- `nlri`—Layer 2 advertisements received or sent by means of the BGP
- `normal`—Normal events
- `policy`—Policy processing
- `route`—Routing information
- `state`—State transitions
- `task`—Routing protocol task processing
- `timer`—Routing protocol timer processing
- `topology`—Layer 2 VPN topology changes caused by reconfiguration or advertisements received from other PE routers using BGP

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify the following modifier:

- `detail`—Provide detailed trace information
- `receive`—Trace received packets
- `send`—Trace transmitted packets

`no-world-readable`—(Optional) Prevents any user from reading the trace file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file*.0. When *trace-file* again reaches its maximum size, *trace-file*.0 is renamed *trace-file*.1 and *trace-file* is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *x*`k` to specify kilobytes, *x*`m` to specify megabytes, or *x*`g` to specify gigabytes
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

`world-readable`—(Optional) Allow any user to read the trace file.

- **Default:** The default is `no-world-readable`.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing Layer 2 VPN Traffic and Operations

traceoptions (Protocols LDP)

IN THIS SECTION

- Syntax | 29693
- Hierarchy Level | 29694
- Description | 29694
- Default | 29694
- Options | 29694
- Required Privilege Level | 29696
- Release Information | 29696

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Specify LDP protocol-level trace options.

Default

The default LDP protocol-level trace options are inherited from the routing protocols `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

Options

`disable`—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **ldp-log**. We recommend that you place LDP tracing output in the file **ldp-log**.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000
- **Default:** 2 files

If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size.

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

- `address`—Operation of address and address withdrawal messages
- `binding`—Label-binding operations

- `error`—Error conditions
- `event`—Protocol events
- `initialization`—Operation of initialization messages
- `label`—Operation of label request, label map, label withdrawal, and label release messages
- `notification`—Operation of notification messages
- `nsr-synchronization`— Nonstop active routing synchronization events
- `p2mp-nsr-synchronization`—Point-to-multipoint nonstop active routing synchronization events
- `packets`—Equivalent to setting `address`, `initialization`, `label`, `notification`, and `periodic` flags (see also the `filter` flag modifier)
- `path`—Label-switched path operations
- `periodic`—Operation of hello and keepalive messages
- `route`—Operation of route messages
- `state`—Protocol state transitions

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- `detail`—Provide detailed trace information.
- `disable`—Disable this trace flag.
- `filter`—Filter to apply to this flag. The `filter` flag modifier can be applied only to the `route`, `path`, and `binding` flags. This flag modifier has the following options:
  - `match-on`—Match on argument specified. The `match-on` option has the following suboptions:
    - `address`—Filter based on the source and destination addresses of packets. Available for the `packets` flag option only.
    - `fec`—Filter based on the FEC associated with the traced object.
  - `policy` *`policy-name`*—Specify the filter policy.
- `receive`—Packets being received.
- `send`—Packets being transmitted.

`no-world-readable`—(Optional) Prevent all users from reading the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the *`trace-file`* again reaches this size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** *`x`*`k` to specify KB, *`x`*`m` to specify MB, or *`x`*`g` to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of files.

`world-readable`—(Optional) Enable any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`match-on address` option for the `filter` flag modifier added in Junos OS Release 10.4.

`nsr-synchronization` and `p2mp-nsr-synchronization` operations for `flag` statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Tracing LDP Protocol Traffic

Junos OS Network Management Administration Guide for Routing Devices

traceoptions (Protocols Link Management)

IN THIS SECTION

- Syntax | 29697
- Hierarchy Level | 29697
- Description | 29697
- Options | 29697
- Required Privilege Level | 29699
- Release Information | 29699

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management],
[edit protocols link-management]
```

Description

Trace options for the LMP protocol.

Options

`disable`—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

`filename`—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000
- **Default:** 2 files

If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size.

`flag`—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

- `all`—Trace all available operations
- `hello-packets`—Trace hello packets on any LMP control channel
- `init`—Output from the initialization messages
- `packets`—Trace all packets other than hello packets on any LMP control channel
- `parse`—Operation of the parser
- `process`—Operation of the general configuration
- `route-socket`—Operation of route socket events
- `routing`—Operation of the routing protocols
- `server`—Server processing operations
- `show`—`show` command servicing operations
- `state`—Trace state transitions of the LMP control channels and traffic engineering links

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- `detail`—Provide detailed trace information
- `receive`—Packets being received
- `send`—Packets being transmitted

`no-world-readable`—(Optional) Prevent all users from reading the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the *`trace-file`* again reaches this size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This

renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of files.

`world-readable`—(Optional) Enable log file access for all users.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for `hello-packets`, `packets`, and `state` flags added in Junos OS Release 8.1.

RELATED DOCUMENTATION

Tracing LMP Traffic

Junos OS Network Management Administration Guide for Routing Devices

traceoptions (Protocols MPLS)

IN THIS SECTION

- Syntax | 29700
- Hierarchy Level | 29700
- Description | 29700

- Default | 29700
- Options | 29700
- Required Privilege Level | 29702
- Release Information | 29702

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name]
```

Description

Configure MPLS tracing options at the protocol level or for a label-switched path.

To specify more than one tracing operation, include multiple `flag` statements.

Default

The default MPLS protocol-level tracing options are inherited from the routing protocols `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

Options

filename—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**. We recommend that you place MPLS tracing output in the file **mpls-log**.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000
- **Default:** 2 files

If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size.

`flag`—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

MPLS Tracing Flags

- `all`—Trace all operations
- `autobw-state`—Automatic bandwidth events.
- `connection`—All circuit cross-connect (CCC) activity
- `connection-detail`—Detailed CCC activity
- `cspf`—CSPF computations
- `cspf-link`—Links visited during CSPF computations
- `cspf-node`—Nodes visited during CSPF computations
- `error`—MPLS error packets
- `graceful-restart`—Trace MPLS graceful restart events
- `lsp-history`—Trace LSP history events
- `lsping`—Trace lsping packets and return codes
- `nsr-synchronization`—Trace NSR synchronization events
- `nsr-synchronization-detail`—Trace NSR synchronization events in detail
- `state`—All LSP state transitions
- `static`—Trace static label-switched path
- `ted-export`—Trace leaking of entries from `lsdist.0` table into the traffic engineering database
- `ted-import`—Trace leaking traffic engineering database entries into the `lsdist.0` table

- `timer`—Timer usage

`no-world-readable`—(Optional) Allow only certain users to read the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** *x*`k` to specify KB, *x*`m` to specify MB, or *x*`g` to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of files.

`world-readable`—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

`ted-export` option introduced in Junos OS Release 14.2.

`ted-import` option introduced in Junos OS Release 14.2.

`lsp-history` option added in Junos OS Release 15.1.

RELATED DOCUMENTATION

Tracing MPLS and LSP Packets and Operations

traceoptions (Protocols MSDP)

IN THIS SECTION

- Syntax | 29703
- Hierarchy Level | 29703
- Description | 29704
- Default | 29704
- Options | 29704
- Required Privilege Level | 29706
- Release Information | 29706

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols msdp],
[edit logical-systems logical-system-name protocols msdp group group-name],
[edit logical-systems logical-system-name protocols msdp group group-name peer address],
[edit logical-systems logical-system-name protocols msdp peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
msdp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
group group-name peer address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp
peer address],
[edit protocols msdp],
[edit protocols msdp group group-name],
```

```
[edit protocols msdp group group-name peer address],
[edit protocols msdp peer address],
[edit routing-instances routing-instance-name protocols msdp],
[edit routing-instances routing-instance-name protocols msdp group group-name],
[edit routing-instances routing-instance-name protocols msdp group group-name peer address],
[edit routing-instances routing-instance-name protocols msdp peer address]
```

Description

Configure MSDP tracing options.

To specify more than one tracing operation, include multiple `flag` statements.

Default

The default MSDP trace options are those inherited from the routing protocol's `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the **msdp-log** file.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size.

- **Range:** 2 through 1000 files
- **Default:** 2 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

MSDP Tracing Flags

- **keepalive—**Keepalive messages

- **packets**—All MSDP packets
- **route**—MSDP changes to the routing table
- **source-active**—Source-active packets
- **source-active-request**—Source-active request packets
- **source-active-response**—Source-active response packets

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

- **Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow any user to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again

reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of trace files.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing MSDP Protocol Traffic

traceoptions (Protocols MVPN)

IN THIS SECTION

- Syntax | 29707
- Hierarchy Level | 29707
- Description | 29707
- Options | 29707
- Required Privilege Level | 29709

Release Information | 29709

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mvpn],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
mvpn],
[edit protocols mvpn],
[edit routing-instances routing-instance-name protocols mvpn]
```

Description

Trace traffic flowing through a Multicast BGP (*MBGP*) MVPN.

Options

`disable`—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" ").

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When *`trace-file`* again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 2 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can specify any of the following flags:

- `all`—All multicast VPN tracing options
- `cmcast-join`—Multicast VPN C-multicast join routes
- `error`—Error conditions
- `general`—General events
- `inter-as-ad`—Multicast VPN inter-AS automatic discovery routes
- `intra-as-ad`—Multicast VPN intra-AS automatic discovery routes
- `leaf-ad`—Multicast VPN leaf automatic discovery routes
- `mdt-safi-ad`—Multicast VPN *MDT SAFI* automatic discovery routes
- `nlri`—Multicast VPN advertisements received or sent by means of the BGP
- `normal`—Normal events
- `policy`—Policy processing
- `route`—Routing information
- `source-active`—Multicast VPN source active routes
- `spmsi-ad`—Multicast VPN SPMSI auto discovery active routes
- `state`—State transitions
- `task`—Routing protocol task processing
- `timer`—Routing protocol timer processing
- `tunnel`—Provider tunnel events
- `umh`—Upstream multicast hop (UMH) events

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify the following modifiers:

- `detail`—Provide detailed trace information
- `disable`—Disable the tracing flag

- `receive`—Trace received packets
- `send`—Trace sent packets

`no-world-readable`—Do not allow any user to read the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When *`trace-file`* again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`x`*`k` to specify kilobytes, *`x`*`m` to specify megabytes, or *`x`*`g` to specify gigabytes
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

`world-readable`—Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support at the `[edit protocols mvpn]` hierarchy level introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Tracing MBGP MVPN Traffic and Operations

traceoptions (Protocols OpenFlow)

IN THIS SECTION

- Syntax | 29710
- Hierarchy Level | 29710
- Description | 29710
- Default | 29710
- Options | 29711
- Required Privilege Level | 29712
- Release Information | 29712

Syntax

```
traceoptions {
    file <filename> <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols openflow]
```

Description

Define tracing operations for OpenFlow.

Default

If you do not include this statement, no OpenFlow-specific tracing operations are performed.

Options

file ***filename***

Name of the file to receive the output of the tracing operation. All files are placed in the **/var/log** directory.

- **Default: /var/log/ofd**

files ***number***

(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed and compressed to ***trace-file*.0.gz**. When ***trace-file*** again reaches its maximum size, ***trace-file*.0.gz** is renamed ***trace-file*.1.gz**, and ***trace-file*** is renamed and compressed to ***trace-file*.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. After this, the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size by using the `size` option and also a filename.

- **Range:** 2 through 1000 files
- **Default:** 10 files

flag ***flag***

Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- `all`—All OpenFlow events.
- `barrier`—OpenFlow barrier events.
- `configuration`—OpenFlow configuration events.
- `filter`—OpenFlow filter events.
- `flow`—OpenFlow flow events.
- `function`—OpenFlow entry and exit events.
- `group`—(Appears only for Juniper Networks devices running OpenFlow v1.3.1 or later) OpenFlow group events.
- `interface`—OpenFlow interface events.
- `nh`—OpenFlow next-hop events.
- `packet-io`—OpenFlow packet in and packet out events.
- `packets`—OpenFlow packet events.
- `statistics`—OpenFlow statistics request and reply events.

- `switch`—OpenFlow switch events including controller connection messages and keepalives, and packets sent to and received from the controller.

match *regular-expression*
(Optional) Log only those lines that match the regular expression.

no-remote-trace
(Optional) Disable tracing and logging operations that track normal operations, error conditions, and packets that are generated by or have passed through the Juniper Networks device.

no-world-readable
(Optional) Disable unrestricted file access, which restricts file access to the owner. This is the default.

size *size*
(Optional) Maximum size of each trace file in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). If you do not specify a unit, the default is bytes. If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

- **Syntax:** *size* to specify bytes, *size*`k` to specify KB, *size*`m` to specify MB, or *size*`g` to specify GB
- **Range:** 10,240 through 1,073,741,824 bytes
- **Default:** 128 KB

world-readable
(Optional) Enable unrestricted file access.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Understanding Support for OpenFlow on Devices Running Junos OS

openflow (Protocols OpenFlow) | **11842**

traceoptions (Protocols OSPF)

IN THIS SECTION

- Syntax | 29713
- Hierarchy Level | 29713
- Description | 29714
- Default | 29714
- Options | 29714
- Required Privilege Level | 29716
- Release Information | 29716

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast)]
```

Description

Configure OSPF protocol-level tracing options.

To specify more than one tracing operation, include multiple `flag` statements.

> **NOTE**: The `traceoptions` statement is not supported on QFabric systems.

Default

The default OSPF protocol-level tracing options are those inherited from the routing protocols `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place OSPF tracing output in the file **ospf-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.
If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 10 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

OSPF Tracing Flags

- **database-description**—Database description packets, which are used in synchronizing the OSPF and OSPFv3 topological database.
- **error**—OSPF and OSPFv3 error packets.
- **event**—OSPF and OSPFv3 state transitions.
- **flooding**—Link-state flooding packets.

- **graceful-restart**—Graceful-restart events.
- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable.
- **ldp-synchronization**—Synchronization events between OSPF and LDP.
- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database.
- **lsa-analysis**—Link-state analysis. Specific to the Juniper Networks implementation of OSPF, Junos OS performs LSA analysis before running the shortest-path-first (SPF) algorithm. LSA analysis helps to speed the calculations performed by the SPF algorithm.
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database.
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **on-demand**—Trace demand circuit extensions.
- **packet-dump**—Content of selected packet types.
- **packets**—All OSPF packets.
- **restart-signaling**—(OSPFv2 only) Restart-signaling graceful restart events.
- **spf**—Shortest-path-first (SPF) calculations.

Global Tracing Flags

- **all**—All tracing operations.
- **general**—A combination of the **normal** and **route** trace operations.
- **normal**—All normal operations. If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions.
- **route**—Routing table changes.
- **state**—State transitions.
- **task**—Routing protocol task processing.
- **timer**—Routing protocol timer processing.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.
If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Example: Tracing OSPF Protocol Traffic

traceoptions (Protocols PCEP)

IN THIS SECTION

- Syntax | 29717
- Hierarchy Level | 29717
- Description | 29717
- Options | 29717
- Required Privilege Level | 29719

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols pcep]
```

Description

Configure the Path Computation Element Protocol (PCEP) tracing options.

Options

filename Name of the file to receive the output of the tracing operation. All files are placed in the directory `/var/log`.

`files` *number* (Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000 files
- **Default:** 2 files.

  If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size.

flag Area of path computation client process (pccd) to enable debugging output.

PCEP Tracing Flags

- *all*—Trace all areas of PCCD code
- *pccd-config*—All configuration parsing operations
- *pccd-core*—PCCD core operations
- *pccd-functions*—PCCD function entries and outs
- *pccd-main*—PCCD main module
- *pccd-rpd*—PCCD communication with RPD
- *pccd-ui*—PCCD user interface handling

no-remote-trace (Optional) Disable remote tracing options.

no-world-readable (Optional) Allow only certain users to read the log file.

`size` *size* (Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches this size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB.
- **Range:** 10 KB through the maximum file size supported on your system.
- **Default:** 1 MB.

  If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of files.

world-readable (Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

pcep | 12320

traceoptions (Protocols PIM)

IN THIS SECTION

- Syntax | 29719
- Hierarchy Level | 29720
- Description | 29720
- Default | 29720
- Options | 29720
- Required Privilege Level | 29722
- Release Information | 29722

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Configure PIM tracing options.

To specify more than one tracing operation, include multiple `flag` statements.

Default

The default PIM trace options are those inherited from the routing protocol's `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the **pim-log** file.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size.

- **Range:** 2 through 1000 files
- **Default:** 2 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

PIM Tracing Flags

- **assert**—Assert messages

- **bidirectional-df-election**—Bidirectional PIM designated-forwarder (DF) election events
- **bootstrap**—Bootstrap messages
- **cache**—Packets in the PIM sparse mode routing cache
- **graft**—Graft and graft acknowledgment messages
- **hello**—Hello packets
- **join**—Join messages
- **mt**—Multicast tunnel messages
- **nsr-synchronization**—Nonstop active routing (NSR) synchronization messages
- **packets**—All PIM packets
- **prune**—Prune messages
- **register**—Register and register stop messages
- **rp**—Candidate RP advertisements
- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

- **Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of trace files.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

- **Range:** 0 KB through the maximum file size supported on your system

- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring PIM Trace Options

Tracing DVMRP Protocol Traffic

Tracing MSDP Protocol Traffic

Configuring PIM Trace Options

traceoptions (Protocols PPP Service)

IN THIS SECTION

- Syntax | 29723
- Hierarchy Level | 29723
- Description | 29724
- Options | 29724
- Required Privilege Level | 29726
- Release Information | 29726

Syntax

```
traceoptions {
    file <filename> <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    filter {
        aci regular-expression;
        ari regular-expresion;
        service-name regular-expresion;
        underlying-interface interface-name;
        user user@domain;
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols ppp-service]
```

Description

Define tracing operations for PPP service processes.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`disable`—Disable this trace flag.

`filter`—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.

> **BEST PRACTICE**: Due to the complexity of agent circuit identifiers and agent remote identifiers, we recommend that you do not try an exact match when filtering on these options. For service names, searching on the exact name is appropriate, but you can also use a regular expression with that option.

- `aci` *`regular-expression`*—Regular expression to match the agent circuit identifier provided by PPP client.
- `ari` *`regular-expression`*—Regular expression to match the agent remote identifier provided by PPP client.
- `service` *`regular-expression`*—Regular expression to match the name of PPPoE service.
- `underlying-interface` *`interface-name`*—Name of a PPP underlying interface. You cannot use a regular expression for this filter option.
- `user` *`user@domain`*—Username of a subscriber. Optionally use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms.

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `accounting-statistics`—Trace accounting statistics events.
- `all`—Trace all operations.

- `authentication`—Trace authentication events.
- `chap`—Trace CHAP events.
- `events`—Trace interface events.
- `gres`—Trace GRES events.
- `init`—Trace daemon initialization events.
- `interface-db`—Trace interface database events.
- `lcp`—Trace LCP state machine events.
- `memory`—Trace memory processing events.
- `ncp`—Trace NCP state machine events.
- `packet-error`—Trace packet error events.
- `pap`—Trace PAP events.
- `parse`—Trace parsing events.
- `profile`—Trace libdynamic profile events.
- `receive-packets`—Trace received PPP packets.
- `routing-process`—Trace routing process interactions.
- `rtp`—Trace real-time priority events.
- `rtsock`—Trace routing socket events.
- `session-db`—Trace session database interactions.
- `smi-services-sentry`—Trace SMI services requests and retries.
- `states`—Trace state machine events.
- `transmit-packets`—Trace transmitted PPP packets.
- `tunnel`—Trace L2TP tunneling events.

`level`—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.

- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

`user` option added in Junos OS Release 14.1.

RELATED DOCUMENTATION

Tracing PPP Service Events for Troubleshooting

traceoptions (Protocols RIP)

IN THIS SECTION

- Syntax | 29727
- Hierarchy Level | 29727
- Description | 29727
- Default | 29728
- Options | 29728
- Required Privilege Level | 29729
- Release Information | 29730

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit protocols rip],
[edit routing-instances routing-instance-name protocols rip]
```

Description

Set RIP protocol-level tracing options.

NOTE: The `traceoptions` statement is not supported on QFabric systems.

Default

The default RIP protocol-level trace options are inherited from the global `traceoptions` statement.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIP tracing output in the file **/var/log/rip-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 10 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

RIP Tracing Options

- **auth**—RIP authentication
- **error**—RIP error packets
- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets such as request, poll, and poll entry packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations

- **normal**—All normal operations
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Tracing RIP Protocol Traffic

traceoptions (Protocols RIPng)

IN THIS SECTION

- Syntax | 29730
- Hierarchy Level | 29730
- Description | 29731
- Default | 29731
- Options | 29731
- Required Privilege Level | 29733
- Release Information | 29733

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng],
```

```
[edit protocols ripng],
[edit routing-instances routing-instance-name protocols ripng]
```

Description

Set RIPng protocol-level tracing options.

Default

The default RIPng protocol-level trace options are inherited from the global `traceoptions` statement.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIPng tracing output in the file **/var/log/ripng-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 10 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

RIPng Tracing Options

- **error**—RIPng error packets
- **expiration**—RIPng route expiration processing
- **holddown**—RIPng hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIPng packets
- **request**—RIPng information packets such as request, poll, and poll entry packets
- **trigger**—RIPng triggered updates

- **update**—RIPng update packets

Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

no-world-readable—(Optional) Do not allow any user to read the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Tracing RIPng Protocol Traffic

traceoptions (Protocols Router Discovery)

IN THIS SECTION

- Syntax | 29733
- Hierarchy Level | 29734
- Description | 29734
- Default | 29734
- Options | 29734
- Required Privilege Level | 29735
- Release Information | 29736

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
```

```
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-discovery],
[edit protocols router-discovery]
```

Description

Configure ICMP protocol-level tracing options.

To specify more than one tracing operation, include multiple `flag` statements.

Default

The default ICMP protocol-level tracing options are inherited from the routing protocols `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place ICMP tracing output in the file **icmp-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 2 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

These are the global tracing options:

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

traceoptions (Protocols RSVP)

IN THIS SECTION

- Syntax | 29736
- Hierarchy Level | 29737
- Description | 29737
- Default | 29737
- Options | 29737
- Required Privilege Level | 29739
- Release Information | 29739

Syntax

```
traceoptions {
    enhanced-frr ;
    file filename  <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Enable RSVP-level trace options.

Default

The default RSVP-level trace options are those inherited from the routing protocols `traceoptions` statement included at the `[edit routing-options]` hierarchy level.

Options

`disable`—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

`enhanced-frr` —(Optional) Enable this option to trace internal events and state changes related to the FRR facility protection enhancements associated with the increased RSVP scaling introduced in Junos OS Release 16.1.

`filename`—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place RSVP tracing output in the file **rsvp-log**.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000
- **Default:** 2 files

If you specify a maximum number of files, you must also include the `size` statement to specify the maximum file size.

`flag`—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

- `all`—All tracing operations
- `error`—All detected error conditions

- `event`—RSVP-related events. New trace messages are added in the file that you create.
- `io-event [detail] [disable]` —Enable tracing of events that occur within the RSVP I/O task; can only be configured for the master routing instance. The trace output is generally independent of routing instance.
- `io-packets [detail] [disable] [receive] [send]` —Enable tracing of messages as they are received from the network or as they are sent out. This flag can be configured independently for each routing instance. Both bundled and individual messages are identified. Use *detail* to show all objects contained in the message. Use *send* and *receive* to limit tracing to outgoing or incoming packets.
- `lmp`—RSVP-LMP interactions
- `packets`—All RSVP packets
- `path`—All path messages
- `pathtear`—PathTear messages
- `resv`—Resv messages
- `resvtear`—ResvTear messages
- `route`—Routing information
- `state`—Session state transitions, including when RSVP-signaled LSPs come up and go down.

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- `detail`—Provide detailed trace information
- `receive`—Packets being received
- `send`—Packets being transmitted

`no-world-readable`—(Optional) Enable only certain users to read the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** *x*`k` to specify KB, *x*`m` to specify MB, or *x*`g` to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of files.

`world-readable`—(Optional) Enable any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing RSVP Protocol Traffic

traceoptions (Protocols Secure Neighbor Discovery)

IN THIS SECTION

- Syntax | 29739
- Hierarchy Level | 29740
- Description | 29740
- Options | 29740
- Required Privilege Level | 29741
- Release Information | 29741

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable | no-
```

```
world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols neighbor-discovery secure]
```

Description

Configure tracing operations for Secure Neighbor Discovery events. To specify more than one tracing operation, include multiple `flag` statements.

Options

file ***filename***—Name of the file to receive the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1** and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.
If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 10 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

Secure Neighbor Discovery Tracing Options

- **configuration**—All configuration events.
- **cryptographic-address**—Cryptographically generated address events.
- **protocol**—All protocol processing events.
- **rsa**—RSA events.

Global Tracing Options

- **all**—All tracing operations.

You can specify one or more of following flag modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

match ***regular-expression***—(Optional) Specify a regular expression to match the output of the trace file you want to log.

no-remote-trace—Disable remote tracing globally or for a specific tracing operation.

no-world-readable—(Optional) Prevent any user from reading this log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1**, and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

world-readable—(Optional) Allow any user to read this log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

> *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*
>
> *Understanding Secure IPv6 Neighbor Discovery*

Understanding IPv6 Neighbor Discovery

traceoptions (Protocols Spring-TE)

IN THIS SECTION

- Syntax | 29742
- Hierarchy Level | 29742
- Description | 29742
- Options | 29743
- Required Privilege Level | 29744

Syntax

```
traceoptions {
file filename <files number> <size size> <world-readable | no-world-readable>;
    flag (all | bfd | compute | controller | dtm | general | interface | route | state |
telemetry-statistics | translation) {
        detail;
    }
```

Hierarchy Level

```
[edit protocols source-packet-routing]
```

Description

Configure the segment routing traffic-engineered (SPRING-TE) protocol tracing options. To specify more than one tracing operation, include multiple flag statements.

Options

filename Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**.

files *number* (Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000 files
- **Default:** 2 files.

  If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

flag Area of path computation client process (pccd) to enable debugging output.

SPRING-TE Tracing Flags

- `all`—Trace everything.
- `bfd`—Bfd-related activity.
- `compute`—Compute profile and computed lsp related activity.
- `controller`—Controller-related activity
- `dtm`—Dtm-related activity
- `general`—General activity
- `interface`—Interface-related activity
- `route`—Route-related activity
- `state`—LSP state-related activity
- `telemetry-statistics`—Telemetry statistics related activity
- `translation`—Translation-related activity.

detail Provide detailed trace information.

filename Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**.

no-remote-trace (Optional) Disable remote tracing options.

no-world-readable (Optional) Allow only certain users to read the log file.

`size` *size* (Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches this size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB.
- **Range:** 10 KB through the maximum file size supported on your system.
- **Default:** 1 MB.

  If you specify a maximum file size, you must also include the `files` statement to specify the maximum number of files.

world-readable (Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

source-packet-routing (Segment Routing Traffic Engineering) | **28257**

traceoptions (Protocols VPLS)

IN THIS SECTION

Syntax | **29745**
Hierarchy Level | **29745**

- Description | 29745
- Options | 29745
- Required Privilege Level | 29747
- Release Information | 29747

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Trace traffic flowing through a VPLS routing instance.

Options

`disable`—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as `all`.

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" ").

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When *`trace-file`* again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 2 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can specify the following tracing flags:

- `all`—All VPLS tracing options
- `connections`—VPLS connections (events and state changes)
- `error`—Error conditions
- `nlri`—VPLS advertisements received or sent by means of the BGP
- `route`—Routing information
- `topology`—VPLS topology changes caused by reconfiguration or advertisements received from other provider edge (PE) routers using BGP

`flag-modifier`—(Optional) Modifier for the tracing flag. You can specify the following modifiers:

- `detail`—Provide detailed trace information.
- `disable`—Disable the tracing flag.
- `receive`—Trace received packets.
- `send`—Trace sent packets.

`no-world-readable`—Do not allow any user to read the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When *`trace-file`* again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`x`*`k` to specify kilobytes, *`x`*`m` to specify megabytes, or *`x`*`g` to specify gigabytes
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

`world-readable`—Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing VPLS Traffic and Operations

traceoptions (Protocols VRRP)

IN THIS SECTION

- Syntax | 29747
- Hierarchy Level | 29748
- Description | 29748
- Default | 29748
- Options | 29748
- Required Privilege Level | 29749
- Release Information | 29749

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression> <microsecond-stamp> <size size>
<world-readable | no-world-readable>;
    flag flag;
```

```
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols vrrp]
```

Description

Define tracing operations for the Virtual Router Redundancy Protocol (VRRP) process.

To specify more than one tracing operation, include multiple `flag` statements.

By default, VRRP logs the error, dcd configuration, and routing socket events in a file in the directory **/var/log**.

Default

If you do not include this statement, no VRRP-specific tracing operations are performed.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, VRRP tracing output is placed in the file **vrrpd**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. When the maximum number is reached, the oldest trace file is overwritten.

- **Range:** 0 through 4,294,967,296 files
- **Default:** 3 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. These are the VRRP-specific tracing options:

- **all**—All VRRP tracing operations
- **database**—Database changes

- **general**—General events
- **interfaces**—Interface changes
- **normal**—Normal events
- **packets**—Packets sent and received
- **state**—State transitions
- **timer**—Timer events

match ***regular-expression***—(Optional) Refine the output to include only those lines that match the given regular expression.

microsecond-stamp—(Optional) Provide a timestamp with microsecond granularity.

no-world-readable—(Optional) Restrict users from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes, megabytes, or gigabytes. When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0.** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your routing platform
- **Default:** 1 MB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—(Optional) Allow users to read the log file.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing VRRP Operations

traceoptions (Proxy Fibre Channel)

IN THIS SECTION

- Syntax | 29750
- Hierarchy Level | 29750
- Description | 29750
- Default | 29751
- Options | 29751
- Required Privilege Level | 29752
- Release Information | 29752

Syntax

```
traceoptions {
    file filename <replace> <size size> <files number> <no-stamp>
    <world-readable | no-world-readable>;
    flag flag <flag-modifier>
}
```

Hierarchy Level

```
[edit fc-fabrics fabric-name proxy]
```

Description

Set proxy FC protocol tracing options.

NOTE: The `traceoptions` statement is not supported on the QFabric system.

Default

Traceoptions is disabled.

Options

file ***name***—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the **/var/log/** directory.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**. The traceoption output continues in a second trace file named ***trace-file*.1**. When ***trace-file*.1** reaches its maximum size, output continues in a third file named ***trace-file*.2**, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 1 trace file

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements:

- **all**—Trace all operations.
- **error**—Trace all error events.
- **interface**—Trace the interface events.
- **normal**—Trace all normal events.
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **packet**—Trace packet decoding operations
- **parse**—Trace configuration parsing.
- **state**—Trace state transitions.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

- **Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Prevent any user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**. Incoming tracefile data is logged in the now empty ***trace-file***. When ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

- **Range:** 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)

- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

traceoptions (Resiliency)

IN THIS SECTION

Syntax | 29753

- Hierarchy Level | 29753
- Description | 29753
- Options | 29753
- Additional Information | 29754
- Required Privilege Level | 29755
- Release Information | 29755

Syntax

```
traceoptions {
   file name {
      files number;
      match;
      (no-world-readable | world-readable);
      size size;
   }
   flag flag;
   no-remote-trace;
   }
```

Hierarchy Level

```
[edit system resiliency]
```

Description

(Junos OS only) Configure trace options for the Juniper Resiliency Interface.

Options

file *file-name* (Required) Use the specified file to receive the output of the tracing operation. All files are placed in the directory **/var/log**.

files *files*

(Optional) Use the specified maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Default:** 3 files
- **Range:** 2 through 1000 files

flag *flag*

Use the specified tracing operation. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `error`—Trace errored packets.
- `state`—Trace state transitions.
- `telemetry`—Trace telemetry state machine events.

match *match*

(Optional) Use the specified regular expression to refine the output to include lines that contain the regular expression.

no-remote-trace

(Optional) Disable remote tracing.

no-world-readable

(Default) Disable unrestricted file access. This means the log file can be accessed only by the user who configured the tracing operation.

world-readable

(Optional) Enable unrestricted file access.

size *size*

(Optional) Use the specified maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

Additional Information

Junos OS Evolved uses a new tracing architecture. All running applications create trace information, with multiple instances of the same application having their own trace information. In Junos OS, you enable tracing operations by configuring the traceoptions statement at the specific hierarchy level you want to trace. Junos OS Evolved, on the other hand, uses an application-based model, and thus trace messages are logged, viewed, and configured by application. As a result, Junos OS Evolved does not support the `traceoptions` statement at many of the hierarchy levels that Junos OS supports.

In Junos OS Evolved, you do not view trace files directly, and you should never add, edit, or remove trace files under the **/var/log/traces** directory because this can corrupt the traces. Instead, you use the `show trace application` *`application-name`* `node` *`node-name`* command to read and decode trace messages stored in the trace files. All running applications on Junos OS Evolved create trace information at the `info` level by default.

The Juniper Resiliency Interface is governed by the `rpdtmd` application. For Junos OS Evolved, to configure traces for a severity other than `info` for the `rpdtmd` application, include the `application rpdtmd node` *`node-name`* `level` *`severity`* statement at the [edit system trace] hierarchy level.

> **NOTE**: For general monitoring and troubleshooting of devices running Junos OS or Junos OS Evolved, we recommend using standard tools such as CLI show commands, system log messages, SNMP, and telemetry data. You should avoid using trace messages for general debugging purposes and long-term solutions because they are subject to change without notice.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Juniper Resiliency Interface

traceoptions (Resource Cleanup)

IN THIS SECTION

- Syntax | 29756
- Hierarchy Level | 29756

- Description | 29756
- Options | 29756
- Required Privilege Level | 29758
- Release Information | 29758

Syntax

```
traceoptions {
    file filename files number match regex size size (world-readable | no-world-readable);
    flag flag;
    level level;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system processes resource-cleanup]
```

Description

Enable debugging tracing for resource cleanup process.

Options

file *filename*
: Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. To include the `file statement`, you must specify a filename.

files *number*
: (Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

flag *flag* Specify which tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Enable all trace options flags.
- `events`—Trace process state change and cleanup events.
- `gencfg`-—Trace GENCFG blobs recorded for cleanup.
- `module`—Trace module code.
- `sysvsem`—Trace SYSV semaphores recorded for cleanup.
- `sysvshm`—Trace SYSV shared memory segments recorded for cleanup.
- `tracking`—Trace tracking code.
- `ui`—Trace user interface operations.

level *level* Specify the level of debugging output:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match conditions that warrant special handling (but are not errors).
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

match *regex* (Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace Disable remote tracing.

size *size* (Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` *number* option.

- **Syntax:** `xk` to specify KB, `xm` to specify MB, or `xg` to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

world-readable | no-world-readable	(Optional). Grant all users permission to read log files, or restrict the permission only to the `root` user and users who have the Junos `maintenance` permission.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

traceoptions (Resource Monitor)

IN THIS SECTION

- Syntax | 29758
- Hierarchy Level | 29759
- Description | 29759
- Options | 29759
- Required Privilege Level | 29760
- Release Information | 29760

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Define tracing operations for the memory resource utilization processes.

Options

file ***filename***

Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**.

- **Default: rmopd**

files ***number***

(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

match ***regular-expression***

(Optional) Refine the output to include lines that contain the regular expression.

size ***maximum-file-size***

(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

world-readable

(Optional) Enable unrestricted file access.

no-world-readable

(Default) Disable unrestricted file access. This means the log file can be accessed only by the user who configured the tracing operation.

flag ***flag***

Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

traceoptions (REST API)

IN THIS SECTION

- Syntax | 29760
- Hierarchy Level | 29760
- Description | 29761
- Options | 29761
- Required Privilege Level | 29761
- Release Information | 29761

Syntax

```
traceoptions {
    flag flag;
}
```

Hierarchy Level

```
[edit system services rest]
```

Description

Define tracing operations for the REST API service.

Options

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, specify `all`. REST API tracing options include:

- `all`—All tracing operations. A combination of the juise and lighttpd tracing operations.
- `juise`—Trace juise operations. Trace information is captured in **/var/chroot/rest-api/var/log/juise**.
- `lighttpd`—Trace lighttpd operations. Trace information is captured in **/var/chroot/rest-api/var/log/lighttpd**.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

traceoptions (RPM)

IN THIS SECTION

- Syntax | 29762
- Hierarchy Level | 29762
- Description | 29762
- Options | 29762
- Additional Information | 29763
- Required Privilege Level | 29764
- Release Information | 29764

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit services rpm]
```

Description

(Junos OS only) Define tracing operations for RPM processes.

Options

file ***filename***

Use the specified file to receive the output of the tracing operation. All files are placed in the directory **/var/log**.

- **Default: rmopd**

files ***number***

(Optional) Use the specified maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

match ***regular-expression***

(Optional) Use the specified regular expression to refine the output to include lines that contain the regular expression.

size ***maximum-file-size***

(Optional) Use the specified maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

world-readable (Optional) Enable unrestricted file access.

no-world-readable (Default) Disable unrestricted file access. This means the log file can be accessed only by the user who configured the tracing operation.

flag *flag* Use the specified tracing operation. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.
- `configuration`—Trace configuration events.
- `error`—Trace events related to catastrophic errors in daemon.
- `ipc`—Trace IPC events.
- `ppm`—Trace ppm events.
- `rpd`—Trace rpd events.
- `statistics`—Trace statistics.

Additional Information

Junos OS Evolved uses a new tracing architecture. All running applications create trace information, with multiple instances of the same application having their own trace information. In Junos OS, you enable tracing operations by configuring the traceoptions statement at the specific hierarchy level you want to trace. Junos OS Evolved, on the other hand, uses an application-based model, and thus trace messages are logged, viewed, and configured by application. As a result, Junos OS Evolved does not support the `traceoptions` statement at many of the hierarchy levels that Junos OS supports.

In Junos OS Evolved, you do not view trace files directly, and you should never add, edit, or remove trace files under the **/var/log/traces** directory because this can corrupt the traces. Instead, you use the `show trace application` *`application-name`* `node` *`node-name`* command to read and decode trace messages stored in the trace files. All running applications on Junos OS Evolved create trace information at the `info` level by default.

RPM is governed by the `rmopd` application. For Junos OS Evolved, to configure traces for a severity other than `info` for the `rmopd` application, include the `application rmopd node` *`node-name`* `level` *`severity`* statement at the [edit system trace] hierarchy level.

> **NOTE**: For general monitoring and troubleshooting of devices running Junos OS or Junos OS Evolved, we recommend using standard tools such as CLI show commands, system log messages, SNMP, and telemetry data. You should avoid using trace messages for general debugging purposes and long-term solutions because they are subject to change without notice.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

`rpd` flag added in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Tracing RPM Operations on MX, M, T and ACX Series Routers

traceoptions (SBC Configuration Process)

IN THIS SECTION

- Syntax | 29765
- Hierarchy Level | 29765
- Description | 29765
- Options | 29765
- Required Privilege Level | 29766
- Release Information | 29767

Syntax

```
traceoptions {
    file filename <files number> <match regex> <size size> <world-readable | no-world-readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit system processes sbc-configuration-process]
```

Description

Configure trace options for the session border controller (SBC) process of the border signaling gateway (BSG).

Options

file *filename*—Name of the file that receives the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory **/var/log**. You can include the following file options:

- **files *number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum number of files, you must also specify a maximum file size with the **size** option and a filename.

- **Range:** 2 through 1000

- **Default:** 3 files

- `match` *regex*—(Optional) Refine the output to include lines that contain the regular expression.

- `no-world-readable`—(Optional) Disable unrestricted file access.

- `size` *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the trace-file again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached.

Then the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the files option and filename.

- **Syntax:** *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB.
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB
- world-readable—(Optional) Enable unrestricted file access.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- all *trace-level*—Trace all SBC process operations.
- common *trace-level*—Trace common events.
- configuration *trace-level*—Trace configuration events.
- device-monitor *trace-level*—Trace device monitor events.
- ipc *trace-level*—Trace IPC events.
- memory-pool *trace-level*—Trace memory pool events.
- *trace-level*—Trace level options are related to the severity of the event being traced. When you choose a trace level, messages at that level and higher levels are captured. Enter one of the following trace levels as the *trace-level*:
  - debug—Log all code flow of control.
  - error—Log failures with a short-term effect.
  - info—Log summary for normal operations, such as the policy decisions made for a call.
  - trace—Log program trace START and EXIT macros.
  - warning—Log failure recovery events or failure of an external entity.
- ui *trace-level*—Trace user interface operations.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

traceoptions (Security)

IN THIS SECTION

- Syntax | 29767
- Hierarchy Level | 29768
- Description | 29768
- Options | 29768
- Required Privilege Level | 29770
- Release Information | 29770

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    level level
    no-remote-trace;
    rate-limit messages-per-second;
}
```

Hierarchy Level

```
[edit security]
[edit services ipsec-vpn]
```

Trace options can be configured at either the `[edit security]` or the `[edit services ipsec-vpn]` hierarchy level, but not at both levels.

Description

Configure security tracing options.

To specify more than one trace option, include multiple `flag` statements. Trace option output is recorded in the **/var/log/kmd** file.

NOTE: The `traceoptions`statement is not supported on QFabric systems.

Options

- `file`—Configure the trace file options.
  - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *`number`*—Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
  - `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

Syntax: *x* `K` to specify KB, *x* `m` to specify MB, or *x* `g` to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace all security events.
  - `certificates`—Trace certificate events.
  - `database`—Trace database events.
  - `general`—Trace general events.
  - `ike`—Trace IKE module processing.
  - `parse`—Trace configuration processing.
  - `policy-manager`—Trace policy manager processing.
  - `routing-socket`—Trace routing socket messages.
  - `timer`—Trace internal timer events.
- `level` *level*—(Optional) Set traceoptions level.
  - `all`—match all levels.
  - `error`—Match error conditions.
  - `info`—Match informational messages.
  - `notice`—Match conditions that should be handled specially.
  - `verbose`—Match verbose messages.
  - `warning`—Match warning messages.
- `no-remote-trace`—Set remote tracing as disabled.

- `rate-limit` *`messages-per-second`*—Limit the incoming rate of trace messages.

Required Privilege Level

trace—To view this in the configuration.

trace-control—To add this to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

traceoptions (Security ALG)

IN THIS SECTION

- Syntax | 29770
- Hierarchy Level | 29771
- Description | 29771
- Options | 29771
- Required Privilege Level | 29772
- Release Information | 29772

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
```

```
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    level (brief | detail | extensive | verbose);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit logical-systems name security alg],
[edit logical-systems name tenants name security alg],
[edit security alg],
[edit services alg],
[edit tenants name security alg]
```

Description

Configure ALG tracing options.

Options

- `file`—Configure the trace file options.
  - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *`number`*—Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
  - `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`.

When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `level`—Set the level of debugging the output option.
  - `brief`—Match brief messages
  - `detail`—Match detail messages.
  - `extensive`—Match extensive messages.
  - `verbose`—Match verbose messages.
- `no-remote-trace`—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

alg | 2130

traceoptions (Security Antispam)

IN THIS SECTION

- Syntax | 29773
- Hierarchy Level | 29773
- Description | 29773
- Options | 29773
- Required Privilege Level | 29774
- Release Information | 29774

Syntax

```
traceoptions flag flag;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-spam]
```

Description

Define tracing operations for Content Security antispam features.

Options

- *flag*:
  - **all**—Enable all antispam trace flags.
  - **manager** —Trace antispam manager information.
  - **sbl**—Trace SBL server information.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

traceoptions (Security Antivirus)

IN THIS SECTION

- Syntax | 29774
- Hierarchy Level | 29774
- Description | 29775
- Options | 29775
- Required Privilege Level | 29775
- Release Information | 29775

Syntax

```
traceoptions flag flag;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus]
```

Description

Define tracing operations for Content Security antivirus features.

Options

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Enable trace all antivirus trace options.
  - `basic`—Trace antivirus module generic basic information.
  - `detail`—Trace antivirus module generic detail information.
  - `engine`—Trace scan engine information.
  - `event`—Trace communication events between routing engine side processes.
  - `ipc`—Trace communication events with Packet Forwarding Engine.
  - `manager`—Trace antivirus manager process activities.
  - `pattern`—Trace detail information of pattern loading.
  - `sendmail`—Trace mail notifying process activities.
  - `statistics`—Trace statistics information.
  - `updater`—Trace pattern updater process activities.
  - `worker`—Trace antivirus worker process activities.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

traceoptions (Security Application Firewall)

IN THIS SECTION

- Syntax | 29776
- Hierarchy Level | 29776
- Description | 29776
- Options | 29777
- Required Privilege Level | 29778
- Release Information | 29778

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security application-firewall]
[edit tenants tenant-name security application-firewall]
```

Description

Configure trace options for the application firewall.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Options

- `file`—Configure the trace file options.
  - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *`number`*—Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
  - `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

    Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

    Range: 10 KB through 1 GB

    Default: 128 KB
  - `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace with all flags enabled
  - `compilation`—Trace rule set compilation events
  - `configuration`—Trace configuration events
  - `ipc`—Trace process inter communication events
  - `lookup`—Trace rule set lookup events
- `no-remote-trace`—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The statement `set tenant` *`tenant-name`* `security application-firewall` is introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Application Firewall Overview

traceoptions (Security Application Proxy)

IN THIS SECTION

- Syntax | 29779
- Hierarchy Level | 29779
- Description | 29779

- Options | 29779
- Required Privilege Level | 29780
- Release Information | 29780

Syntax

```
traceoptions {
    flag flag;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm application-proxy]
[edit logical-systems logical-system-name security]
```

Description

Configure tracing options for application proxy.

Options

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `abort`—Trace terminated sessions for application proxy.
  - `all`—Trace with all flags enabled.
  - `anti-virus`—Trace anti-virus information.
  - `application-objects`—Trace application-proxy objects information.
  - `basic`—Trace application-proxy related basic information.
  - `buffer`— Trace application-proxy data buffer information.

- `connection-rating`—Trace connection rating information.
- `detail`—Trace application-proxy related detailed information.
- `express-anti-virus`—Trace anti-virus express engine information.
- `ftp-control`—Trace FTP control connection information.
- `ftp-data`—Trace FTP data connection information.
- `http`—Trace HTTP protocol information.
- `imap`—Trace IMAP protocol information.
- `memory`—Trace memory usage.
- `mime`—Trace MIME parser information.
- `parser`— Trace protocol parser information.
- `pfe`—Trace communication with PFE.
- `pop3`—Trace POP3 protocol information.
- `queue`—Trace queue information.
- `regex-engine`—Trace Pattern Match Engine (PME) information.
- `smtp`—Trace SMTP protocol information.
- `sophos-anti-virus`—Trace anti-virus sophos engine information.
- `tcp`—Trace TCP level information.
- `timer`—Trace timer processing.
- `utm-realtime`—Trace application-proxy realtime-thread information

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

The logical system option is introduced in Junos OS Release 18.3R1.

traceoptions (Security Content Filtering)

IN THIS SECTION

- Syntax | 29781
- Hierarchy Level | 29781
- Description | 29781
- Options | 29781
- Required Privilege Level | 29782
- Release Information | 29782

Syntax

```
traceoptions flag flag;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering]
```

Description

Define tracing options for content filtering features.

Options

- *flag*:
  - **all**—Enable all content filtering trace flags.
  - **basic** —Trace content filtering basic information.

- **detail**—Trace content filtering detailed information.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

traceoptions (Security Content Security)

IN THIS SECTION

- Syntax | 29782
- Hierarchy Level | 29782
- Description | 29783
- Options | 29783
- Required Privilege Level | 29783
- Release Information | 29783

Syntax

```
traceoptions flag flag;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm]
```

Description

Define tracing operations for Content Security features.

Options

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Enable trace for all Content Security trace options.
  - `cli`—Trace CLI configuration activity and command changes.
  - `daemon`—Trace daemon information.
  - `ipc`—Trace communication events with Packet Forwarding Engine (PFE).
  - `pfe`—Trace PFE information.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

traceoptions (Security Datapath Debug)

IN THIS SECTION

- Syntax | 29784
- Hierarchy Level | 29784
- Description | 29784
- Options | 29784

Required Privilege Level | 29785
Release Information | 29785

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security datapath-debug]
```

Description

Sets the trace options for datapath-debug.

> **NOTE**: Data path debugging is supported on SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800.

Options

- `file`—Configure the trace file options.
  - `filename`—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.

- `files` *`number`*—Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed to **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

  If you specify a maximum number of files, you also must specify a maximum file size with the size option and a filename.

  Range: 2 through 1000 files

  Default: 10 files

- `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.

- `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum file size, you also must specify a maximum number of trace files with the files option and a filename.

  Syntax: x K to specify KB, x m to specify MB, or x g to specify GB

  Range: 10 KB through 1 GB

  Default: 128 KB

- `world-readable | no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option

- `no-remote-trace`—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 9.6.

traceoptions (Security Dynamic VPN)

IN THIS SECTION

- Syntax | 29786
- Hierarchy Level | 29786
- Description | 29786
- Options | 29787
- Required Privilege Level | 29787
- Release Information | 29787

Syntax

```
traceoptions {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
    flag {
        all;
    }
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security dynamic-vpn]
```

Description

Configure dynamic VPN tracing options. This feature is supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

In Junos OS Release 21.4R1, we’ve deprecated the `traceoptions` configuration statement and we might remove it completely in a future release.

Options

file Configure the trace file options.

`file` *`filename`*—Name of the file to receive the output of the tracing operation.

flag Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.

- Values:
  - `all`—Enable all tracing operations

level Level of debugging output

- Values:
  - all—Match all levels
  - error—Match error conditions
  - info—Match informational messages
  - notice—Match conditions that should be handled specially
  - verbose—Match verbose messages
  - warning—Match warning messages

no-remote-trace Disable remote tracing

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Dynamic VPN Overview

traceoptions (Security Firewall Authentication)

IN THIS SECTION

- Syntax | 29788
- Hierarchy Level | 29788
- Description | 29788
- Options | 29788
- Required Privilege Level | 29789
- Release Information | 29789

Syntax

```
traceoptions {
    flag {
        all <detail | extensive | terse>;
        authentication <detail | extensive | terse>;
        proxy <detail | extensive | terse>;
    }
}
```

Hierarchy Level

```
[edit security firewall-authentication]
```

Description

Define data-plane firewall authentication tracing options.

Options

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.

  - `all`—Enable all tracing operations
  - `authentication`—Trace data-plane firewall authentication events
  - `proxy`—Trace data-plane firewall authentication proxy events
- `detail`—Display moderate amount of data in trace.
- `extensive`—Display extensive amount of data in trace.
- `terse`—Display minimum amount of data in trace.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Firewall User Authentication Overview

traceoptions (Security Flow)

IN THIS SECTION

- Syntax | 29790
- Hierarchy Level | 29791
- Description | 29791
- Options | 29791
- Required Privilege Level | 29794
- Release Information | 29794

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag {
        all;
        basic-datapath;
        fragmentation;
        high-availability;
        host-traffic;
        jexec;
        multicast;
        route;
        session;
        session-scan;
        tcp-basic;
        tunnel;
    }
    no-remote-trace;
    packet-filter filter-name {
        conn-tag session-conn
        destination-port port-identifier;
        destination-prefix address;
        interface interface-name;
        protocol protocol-identifier;
        source-port port-identifier;
        source-prefix address;
    }
    rate-limit messages-per-second;
    trace-level (brief | detail | minimal);
}
root-override
```

Hierarchy Level

```
[edit security flow]
[edit logical-systems logical-system-name security flow]
[edit tenants tenant-name security flow]
```

Description

Configure flow tracing options.

Options

file

Configure the trace file options.

filename

Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.

files *number*

Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000 files
- **Default:** 10 files

match *regular-expression*

Refine the output to include lines that contain the regular expression.

size *maximum-file-size*

Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

Syntax: *x* `K` to specify KB, *x* `m` to specify MB, or *x* `g` to specify GB

- **Range:** 0 KB through 1 GB
- **Default:** 128 KB

world-readable | no-world-readable

By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

flag

Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.

all	Trace with all flags enabled
basic-datapath	Trace basic packet flow activity
fragmentation	Trace IP fragmentation and reassembly events
high-availability	Trace flow high-availability information
host-traffic	Trace flow host traffic information
jexec	Trace packet footprints
multicast	Trace multicast flow information
route	Trace route lookup information
session	Trace session creation and deletion events
session-scan	Trace session scan information
tcp-basic	Trace TCP packet flow information
tunnel	Trace tunnel information

no-remote-trace

Set remote tracing as disabled.

packet-filter ***filter-name***

Packet filter to enable during the tracing operation. Configure the filtering options.

destination-port *port-identifier*	Match TCP/UDP destination port
destination-prefix *address*	Destination IP address prefix
interface *interface-name*	Logical interface
protocol *protocol-identifier*	Match IP protocol type
source-port *port-identifier*	Match TCP/UDP source port
source-prefix *address*	Source IP address prefix

rate-limit *messages-per-second*
Limit the incoming rate of trace messages.

trace-level
Set the level for trace logging. This option is available only when the flag is set.

- **brief** Trace key flow information, such as message types sent between SPU and central point, policy match, and packet drop reasons.
- **detail** Trace extensive flow information, such as detailed information about sessions and fragments. Detail is the default level.
- **minimal** Trace messages including notice, warning, error, criticality, alert, and emergency.

root-override
Root administrator can debug the packets flow traces from all the logical systems and tenant systems for debug by enabling the `root-override` option.

Table 1811 on page 29793 describes tracing of packet footprints with and without root override options.

Table 1811: Logical System or Tenant System Tracing of Packet Footprints With and Without Root Override Options

Trace Configuration	With Root Override	Without Root Override
Logical systems and tenant systems with own flow trace configuration	The flow trace of logical systems and tenant systems are recorded to the root.	The flow trace of logical systems and tenant systems are recorded to the configured logical systems and tenant systems.

Table 1811: Logical System or Tenant System Tracing of Packet Footprints With and Without Root Override Options *(Continued)*

Trace Configuration	With Root Override	Without Root Override
Logical systems and tenant systems without own flow trace configuration	The flow trace of logical systems and tenant systems are recorded to root.	The flow trace of logical systems and tenant systems are not recorded anywhere

When the packets are switched by the VPLS switch in root logical system, the `jexec` trace packet footprints might not be complete in logical systems and tenant systems.

If a logical system or tenant system does not configure the `security flow traceoptions`, and there is no root-override configured under root either, the packet traces will not be recorded.

To get complete packet footprints for logical systems and tenant systems related `flow` and `jexec` trace packet footprints, we recommend that you configure the required flow trace with `root-override`.

Required Privilege Level

trace—To view this in the configuration.

trace-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Statement updated in Junos OS Release 12.1X46-D10 with the `trace-level` option and additional flags.

Statement updated in Junos OS Release 15.1X49-D70 with the addition of the `conn-tag` filter parameter.

Support at the following hierarchy levels introduced in Junos OS Release 19.4R1: `[edit logical-systems` *`logical-system-name`* `security flow]`, and `[edit tenants` *`tenant-name`* `security flow]`.

Statement `jexec`, `minimal`, and `root-override` options are introduced in Junos OS Release 20.4R1. The option `error` is now changed to `minimal`.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

traceoptions (Security Group VPN)

IN THIS SECTION

- Syntax | 29795
- Hierarchy Level | 29796
- Description | 29796
- Options | 29796
- Required Privilege Level | 29798
- Release Information | 29798

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag (all | certificates | config | database | general | high-availability | ike | next-
hop-tunnels | parse | policy-manager | routing-socket | thread | timer);
    gateway-filter {
        local-address ip-address;
        remote-address ip-address;
    }
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security group-vpn member ike]
[edit security group-vpn server]
```

Description

Configure group VPN tracing options to aid in troubleshooting the IKE or server issues. This helps troubleshoot one or multiple tunnels negotiation by standard tracefile configuration. Tracing allows the user to view the detailed packet exchange and the negotiation information. Group VPNv2 is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600 devices and vSRX Virtual Firewall instances.

Options

- `file`—Configure the trace file options.
  - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.
  - `files` *`number`*—Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
  - `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and filename.

    Syntax: *`x`* `k` to specify KB, *`x`*`m` to specify MB, or *`x`*`g` to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

- `world-readable | no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace all activity.
  - `certificates`—Trace certificate-related activity.
  - `config`—Trace configuration activity.
  - `database`—Trace SA-related database activity.
  - `general`—Trace general activity.
  - `high-availability`—Trace high-availability operations.
  - `ike`—Trace IKE protocol activity.
  - `next-hop-tunnels`—Trace next-hop tunnel operations.
  - `parse`—Trace configuration processing.
  - `policy-manager`—Trace IKE callback activity.
  - `routing-socket`—Trace routing socket activity.
  - `thread`—Trace thread processing.
  - `timer`—Trace timer activity.
- `gateway-filter`—Configure debugging for the tunnel between the group VPN server and a group member. This option is configured on a group VPN server or member.
  - `local-address`—When configured on a server, the IP address of the group VPN server. When configured on a member, the IP address of the group VPN member.
  - `remote-address`—When configured on a server, the IP address of the group VPN member. When configured on a member, the IP address of the group VPN server.
- `level`—Set the level of debugging.
  - `all`—Match all levels.

- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match conditions that should be handled specifically.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

- `no-remote-trace`—Disable remote tracing.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2. Support for `gateway-filter` option for the [`edit security group-vpn member ike`] hierarchy level added in Junos OS Release 15.1X49-D30 for vSRX Virtual Firewall.

RELATED DOCUMENTATION

Group VPNv2 Overview

traceoptions (Security GTP)

IN THIS SECTION

- Syntax | 29799
- Hierarchy Level | 29799
- Description | 29799
- Options | 29799
- Required Privilege Level | 29801
- Release Information | 29801

Syntax

```
traceoptions {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
    flag name;
    no-remote-trace;
    trace-level {
        (error | info | notice | verbose | warning);
    }
}
```

Hierarchy Level

```
[edit security gtp]
```

Description

Use the traceoptions function to enable the device to identify and log the contents of GTP-U or GTP-C messages based on IMSI prefixes or Mobile Station-Integrated Services Data Network (MS-ISDN) identification.

From Junos OS Release 20.4R1 onwards, the `[edit security gprs gtp]` hierarchy level is replaced by `[edit security gtp]`.

Options

- `file`—Configure the trace file options.
  - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *`number`*—Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

Default: 10 files

- `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
- `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

  Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

  Range: 10 KB through 1 GB

  Default: 128 KB

- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - **all**—Trace everything.
  - **chassis-cluster**—Trace chassis cluster events.
  - **configuration**—Trace configuration events.
  - **flow**—Trace flow events.
  - **parser**—Trace parser events.
- `no-remote-trace`—Set remote tracing as disabled.

trace-level GTP trace level

- Values:
  - error—Match error conditions
  - info—Match informational messages

- notice—Match conditions that should be handled specially
- verbose—Match verbose messages
- warning—Match warning messages

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0. Support.

RELATED DOCUMENTATION

Example: Configure the Validity of IP Address in GTP Messages

traceoptions (Security H323 ALG)

IN THIS SECTION

- Syntax | 29802
- Hierarchy Level | 29802
- Description | 29802
- Options | 29802
- Required Privilege Level | 29803
- Release Information | 29803

Syntax

```
traceoptions {
    flag
        all {
            extensive
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg h323],
[edit logical-systems name tenants name security alg h323],
[edit security alg h323],
[edit services alg h323],
[edit tenants name security alg h323]
```

Description

Configure H.323 tracing options. H.323 is a control-signaling protocol. H.323 endpoints use H.323 protocol to exchange messages.

Options

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace with all flags enabled.
  - `chassis-cluster`—Trace chassis cluster information.
  - `h225-asn1`—Trace H.225 ASN.1 processing activity.
  - `h245`—Trace H.245 processing activity.
  - `h245-asn1`—Trace H.245 ASN.1 processing activity.
  - `q931`—Trace Q.931 processing activity.
  - `ras`—Trace remote access server (RAS) processing activity.

- `ras-asn1`—Trace RAS ASN.1 processing activity.

- `detail`—Display moderate amount of data in trace.
- `extensive`—Display extensive amount of data in trace.
- `terse`—Display minimum amount of data in trace.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

h323 | 7961

traceoptions (Security IDP)

IN THIS SECTION

- Syntax | 29804
- Hierarchy Level | 29804
- Description | 29804
- Options | 29804
- Required Privilege Level | 29805
- Release Information | 29806

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag all;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security idp]
```

Description

Configure IDP tracing options.

Options

- `file`—Configure the trace file options.
  - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - `files` *`number`*—Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed to **trace-file.0** then **trace-file.1** and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files

- `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
- `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When `trace-file.0` again reaches its maximum size, **trace-file.1** is renamed **trace-file.2** and **trace-file.0** is renamed **trace-file.1**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

  Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

  Range: 10 KB through 1 GB

  Default: 128 KB
- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Trace operation to perform.
  - `all`—Trace with all flags enabled
- `level`—Set the level of debugging the output option.
  - `all`—Match all levels
  - `error`—Match error conditions
  - `info`—Match informational messages
  - `notice`—Match conditions that should be handled specially
  - `verbose`—Match verbose messages
  - `warning`—Match warning messages
- `no-remote-trace`—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

traceoptions (Security IKE)

IN THIS SECTION

- Syntax | 29806
- Hierarchy Level | 29807
- Description | 29807
- Options | 29807
- Required Privilege Level | 29809
- Release Information | 29809

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }

    level (critical | error | terse | warning | detail);
    flag flag (all | certificates | config | database | general | high-availability | ike | next-
hop-tunnels | parse | policy-manager | routing-socket | thread | timer);
    no-remote-trace;
    rate-limit messages-per-second;
}
```

Hierarchy Level

```
[edit security ike]
```

Description

Configure IKE tracing options to aid in troubleshooting the IKE issues. This helps troubleshoot one or multiple tunnels negotiation by standard tracefile configuration. IKE tracing allows the user to view the detailed packet exchange and the negotiation information in Phase 1 and Phase 2. IKE tracing is not enabled by default. By default , all IKE or IPsec negotiations are logged into /var/log/kmd. But user can also specify customized file name while configuring the IKE traceoptions.

Options

- file—Configure the trace file options.
  - *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

    Default: kmd

  - `files` *number*—Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed to *trace-file*`.0`, then *trace-file*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files

  - `match` *regular-expression*—Refine the output to include lines that contain the regular expression.

  - `size` *maximum-file-size*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file*`.0`. When the `trace-file` again reaches its maximum size, *trace-file*`.0` is renamed *trace-file*`.1` and *trace-file* is renamed *trace-file*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and filename.

    Syntax: *x*`k` to specify KB, *x*`m` to specify MB, or *x*`g` to specify GB

Range: 10 KB through 1 GB

Default: 1024 KB

- `world-readable | no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `level`—Specify the log levels.
  - `critical`—Log single point failures which needs your immediate attention
  - `error`—Log fatal application errors
  - `terse`—Log syslog messages
  - `warning`—Log recoverable errors
  - `detail`—Log all operational information
- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace all iked process modules activity
  - `certificates`—Trace certificate-related activity
  - `config`—Trace configuration download processing
  - `database`—Trace VPN-related database activity
  - `general`—Trace general activity
  - `high-availability`—Trace high-availability operations
  - `ike`—Trace IKE protocol activity
  - `next-hop-tunnels`—Trace next-hop tunnels operations
  - `parse`—Trace VPN parsing activity
  - `policy-manager`—Trace iked callback activity
  - `routing-socket`—Trace routing socket activity
  - `thread`—Trace thread processing
  - `timer`—Trace timer activity

By default, the `flag` statement is not set. You need to explicitly configure the `flag` statement to perform trace operation.

- `no-remote-trace`—Set remote tracing as disabled.
- `rate-limit` *`messages-per-second`*—Configure the incoming rate of trace messages.

  Range: 0 through 4,294,967,295

  Default: 0

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

`level` options introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

IPsec Overview

ike (Security) | 8432

traceoptions (Security IPsec)

IN THIS SECTION

- Syntax | 29810
- Hierarchy Level | 29810
- Description | 29810
- Options | 29810
- Required Privilege Level | 29810

Release Information | 29811

Syntax

```
traceoptions {
    flag flag;
}
```

Hierarchy Level

```
[edit security ipsec]
```

Description

Configure IPsec tracing options. Trace operations track IPsec events and record them in a log file in the /var/log directory.

Trace operations are written to the trace file **/var/log/kmd**.

Options

- `flag`—To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace with all flags enabled
  - `next-hop-tunnel-binding`—Trace next-hop tunnel binding events
  - `packet-drops`—Trace packet drop activity
  - `packet-processing`—Trace data packet processing events
  - `security-associations`—Trace security association (SA) management events

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

IPsec Overview

traceoptions (Security Log)

IN THIS SECTION

- Syntax | 29811
- Hierarchy Level | 29812
- Description | 29812
- Options | 29812
- Required Privilege Level | 29813
- Release Information | 29813

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag (all | configuration | hpl | report | source);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security log]
[edit logical-systems name security log]
[edit tenants tenant-name security log]
```

Description

Configure security log tracing options.

Options

- `file`—Configure the trace file options.
  - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *`number`*—Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
  - `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

    Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

    Range: 10 KB through 1 GB

    Default: 128 KB

  - `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace with all flags enabled
  - `configuration`—Trace configuration events
  - `hpl`— Trace HPL logging
  - `report`— Trace report
  - `source`—Communicate with security log forwarder
- `no-remote-trace`—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

The [edit `logical-systems` *name* security log] and [edit `tenants` *tenant-name* security log] hierarchy levels introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

syslog (System) | **28836**

traceoptions (Security MGCP ALG)

IN THIS SECTION

- Syntax | 29814
- Hierarchy Level | 29814
- Description | 29814
- Options | 29815
- Required Privilege Level | 29815
- Release Information | 29815

Syntax

```
traceoptions {
    flag
        all {
            extensive
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg mgcp],
[edit logical-systems name tenants name security alg mgcp],
[edit security alg mgcp],
[edit services alg mgcp],
[edit tenants name security alg mgcp]
```

Description

Configure Media Gateway Control Protocol (MGCP) tracing options. MGCP is a text-based Application Layer protocol. Use MGCP ALG for call set up and call control.

Options

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace with all flags enabled.
  - `call`—Trace call processing activity.
  - `chassis-cluster`—Trace chassis cluster information.
  - `decode`—Trace decoder operations activity.
  - `error`—Trace processing errors activity.
  - `nat`—Trace Network Address Translation (NAT) processing activity.
  - `packet`—Trace MGCP protocol packet processing activity.
  - `rm`—Trace MGCP Resource Management (Resmgr) functions activity.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

mgcp | 10693

traceoptions (Security NAT)

IN THIS SECTION

Syntax | 29816

- Hierarchy Level | 29816
- Description | 29816
- Options | 29816
- Required Privilege Level | 29818
- Release Information | 29818

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security nat]
```

Description

Configure NAT tracing options.

Options

- `file`—Configure the trace file options.
    - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.

- `files` *`number`*—Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

  If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

  Range: 2 through 1000 files

  Default: 10 files

- `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.

- `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

  Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

  Range: 10 KB through 1 GB

  Default: 128 KB

- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.

  - `all`—Trace with all flags enabled

  - `destination-nat-pfe`—Trace destination NAT events on PFE-ukernel side

  - `destination-nat-re`—Trace destination NAT events on Routing Engine (RE) side

  - `destination-nat-rt`—Trace destination NAT events on Packet Forwarding Engine real-time (PFE-RT) side

  - `source-nat-pfe`—Trace source NAT events on PFE-ukernel side

  - `source-nat-re`—Trace source NAT events on RE side

  - `source-nat-rt`—Trace source NAT events on PFE-RT side

- `static-nat-pfe`—Trace static NAT events on PFE-ukernel side
- `static-nat-re`—Trace static NAT events on RE side
- `static-nat-rt`—Trace static NAT events on PFE-RT side

- `no-remote-trace`—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

traceoptions (Security PKI)

IN THIS SECTION

- Syntax | 29818
- Hierarchy Level | 29819
- Description | 29819
- Options | 29819
- Required Privilege Level | 29820
- Release Information | 29820

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
```

```
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag {
        all;
        certificate-verification;
        online-crl-check;
    }
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security pki]
```

Description

Configure public key infrastructure (PKI) tracing options. To specify more than one trace option, include multiple flag statements. Trace option output is recorded in the /var/log/pkid file.

Options

- `file`—Configure the trace file options.
  - *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *number*—Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed to *trace-file*`.0`, then *trace-file*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - `match` *regular-expression*—Refine the output to include lines that contain the regular expression.

- `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

  Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

  Range: 10 KB through 1 GB

  Default: 128 KB

- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.

  - `all`—Trace with all flags enabled

  - `certificate-verification`—Trace PKI certificate verification events

  - `online-crl-check`—Trace PKI online certificate revocation list (CRL) events

- `no-remote-trace`—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

RELATED DOCUMENTATION

Basic Elements of PKI in Junos OS

traceoptions (Security Policies)

IN THIS SECTION

- Syntax | 29821
- Hierarchy Level | 29821
- Description | 29821
- Options | 29822
- Required Privilege Level | 29823
- Release Information | 29823

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security policies]
```

Description

Configure policy tracing options.

Options

- `file`—Configure the trace file options.
  - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *`number`*—Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
  - `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

    Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

    Range: 10 KB through 1 GB

    Default: 128 KB
  - `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.
- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace with all flags enabled
  - `configuration`—Trace configuration events
  - `compilation`—Trace policy compilation events

- `ipc`—Trace process inter communication events
- `lookup`—Trace policy lookup events
- `routing-socket`—Trace routing socket events
- `rules`—Trace policy rules-related events

- `no-remote-trace`—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Security Policies Overview

traceoptions (Security SCCP ALG)

IN THIS SECTION

- Syntax | 29824
- Hierarchy Level | 29824
- Description | 29824
- Options | 29824
- Required Privilege Level | 29825
- Release Information | 29825

Syntax

```
traceoptions {
    flag
        all {
            extensive
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg sccp],
[edit logical-systems name tenants name security alg sccp],
[edit security alg sccp],
[edit services alg sccp],
[edit tenants name security alg sccp]
```

Description

Configure Skinny Client Control Protocol (SCCP) tracing options. Skinny Client Control Protocol (SCCP) is a Cisco proprietary protocol for call signaling.

Options

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace with all flags enabled.
  - `call`—Trace call processing activity.
  - `chassis-cluster`—Trace chassis cluster information.
  - `cli`—Trace CLI configuration activity and command changes.
  - `decode`—Trace decoder operations activity.
  - `error`—Trace processing errors activity.
  - `init`—Enable tracing for SCCP initialization errors

- `nat`—Enable tracing for SCCP NAT processing
- `rm`——Trace SCCP Resource Management (Resmgr) functions activity.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sccp | 16341

traceoptions (Security Screen)

IN THIS SECTION

- Syntax | 29825
- Hierarchy Level | 29826
- Description | 29826
- Options | 29826
- Required Privilege Level | 29827
- Release Information | 29827

Syntax

```
traceoptions {
    file {
        filename;
```

```
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security screen]
```

Description

Configure screen tracing options.

To specify more than one tracing option, include multiple `flag` statements.

Options

- `file`—Configure the trace file options.
  - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *`number`*—Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
  - `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-`*

file is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the files option and a filename.

Syntax: *x* K to specify KB, *x* m to specify MB, or *x* g to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

- world-readable | no-world-readable—By default, log files can be accessed only by the user who configures the tracing operation. The world-readable option enables any user to read the file. To explicitly set the default behavior, use the no-world-readable option.

- flag—Trace operation to perform. To specify more than one trace operation, include multiple flag statements.
  - all—Trace all screen events
  - configuration—Trace screen configuration events
  - flow—Trace flow events
- no-remote-trace—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

traceoptions (Security SCTP)

IN THIS SECTION

- Syntax | 29828
- Hierarchy Level | 29828
- Description | 29828
- Options | 29829
- Required Privilege Level | 29830
- Release Information | 29830

Syntax

```
traceoptions {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
    flag name;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security sctp]
```

Description

Use the traceoptions function to set the trace options for Stream Control Transmission Protocol (SCTP).

From Junos OS Release 20.4R1 onwards, the `[edit security gprs sctp]` hierarchy level is replaced by `[edit security sctp]`.

Options

file

Configure the trace file options.

filename

Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.

files *number*

Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000 files.
- **Default:** 10 files.

match *match*

Refine the output to include lines that contain the regular expression.

size *size*

Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

- **Range:** 10 KB through 1 GB.
- **Default:** 128 KB.

world-readable | no-world-readable

By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

flag

Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.

- all—Trace everything.

- chassis-cluster—Trace chassis cluster events.
- configuration—Trace configuration events.
- detail—Trace information used for debugging.
- flow—Trace flow events.
- parser—Trace parser events.

no-remote-trace Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2. The flag statement `detail` introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding Stream Control Transmission Protocol

SCTP Configuration

traceoptions (Security SIP ALG)

IN THIS SECTION

- Syntax | 29831
- Hierarchy Level | 29831
- Description | 29831
- Options | 29831

Required Privilege Level | 29832
Release Information | 29832

Syntax

```
traceoptions {
    flag
        all {
            extensive
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name security alg sip],
[edit logical-systems name tenants name security alg sip],
[edit security alg sip],
[edit services alg sip],
[edit tenants name security alg sip]
```

Description

Configure Session Initiation Protocol (SIP) tracing options. Session Initiation Protocol (SIP) is a signaling protocol for initiating, modifying, and terminating multimedia sessions over the internet.

Options

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - `all`—Trace with all flags enabled.
  - `call`—Trace call processing activity.
  - `chassis-cluster`—Trace chassis cluster information.

  - `nat`—Trace SIP Network Address Translation (NAT) processing activity.
  - `parser`—Trace SIP parser operations.
  - `rm`—Trace SIP Resource Management (Resmgr) functions activity.
- detail—Display moderate amount of data in trace.
- extensive—Display extensive amount of data in trace.
- terse—Display minimum amount of data in trace.
-

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sip (Security) | 27977

traceoptions (Security User Identification)

IN THIS SECTION

- Syntax | 29833
- Hierarchy Level | 29833
- Description | 29833
- Options | 29833
- Required Privilege Level | 29834
- Release Information | 29834

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security user-identification]
```

Description

Configure flow tracing options.

Options

- `file`—Configure the trace file options.
  - *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *number*—Maximum number of trace files. When a trace file named *trace-file* its maximum size, it is renamed to *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

    Default: 10 files

- `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
- `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

  Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

  Range: 10 KB through 1 GB

  Default: 128 KB
- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Trace operation to perform.
  - `all`—Trace with all flags enabled
  - `no-remote-trace`—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding User Role Firewalls

traceoptions (Security Web Filtering)

IN THIS SECTION

- Syntax | 29835
- Hierarchy Level | 29835
- Description | 29835
- Options | 29835
- Required Privilege Level | 29836
- Release Information | 29836

Syntax

```
traceoptions flag flag;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering]
```

Description

Define tracing operations for individual Web filtering modules. To specify more than one tracing operation, include multiple flag statements.

Options

- *flag*:
  - **all**—Enable all Web filtering trace flags.
  - **basic** —Trace basic information on the Web filtering module.
  - **cache**—Enable Web filtering flags for the Web filtering cache maintained on the Web filtering module.

- **enhanced**—Enable Web filtering flags for processing through Enhanced Web Filtering.
- **heartbeat**—Trace connectivity information with Web filter server.
- **ipc**—Trace Web filtering IPC messages.
- **packet**—Trace packet information from session management.
- **profile**—Trace profile configuration information.
- **requests**—Trace requests sent to Web filter server.
- **response**—Trace response received from Web filter server.
- **session manager**—Trace session management information.
- **socket**—Trace the communication socket with Web filter server.
- **timer**—Trace aging information for requests sent to server.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

traceoptions (Services)

IN THIS SECTION

- Syntax | 29837
- Hierarchy Level | 29837
- Description | 29837
- Options | 29837

Required Privilege Level | 29839
Release Information | 29839

Syntax

```
traceoptions {
    file <filename> <files number> <match regex> <size size> <world-readable | no-world-
readable>;
    flag flag;
    level <error>;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system services extension-service]
```

Description

Define tracing operations for the JET service process (jsd).

Options

file Indicate trace file information.

- ***filename*** Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the **/var/log/** directory.
- **files *number*** (Optional) Specify the maximum number of trace files.
  - **Range:** 2 through 1000
  - **Default:** 10
- **match *regex*** Specify the regular expression for lines to be logged.

size *size*

(Optional) Specify the maximum size of each trace file. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0. The traceoption output continues in a second trace file named *trace-file*.1. When *trace-file*.1 reaches its maximum size, output continues in a third file named *trace-file*.2, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

- **Range:** 10,240 through 1,073,741,824 bytes
- **Default:** 1000k

world-readable | no-world-readable

(Optional). Grant all users permission to read log files, or restrict the permission only to the root user and users who have Junos OS maintenance permission.

flag *flag*

Specify the tracing operation to perform. To specify more than one tracing operation, include multiple flag statements:

- all—Trace everything.
- config—Trace configuration events.
- general—Trace general events.
- grpc—Trace grpc server events.
- libgrpc-debug—(Junos OS Release 19.3R1 only) Trace all lib grpc-related events.
- libgrpc-errors—(Junos OS Release 19.3R1 only) Trace lib grpc errors.
- libgrpc-info—(Junos OS Release 19.3R1 only) Trace lib grpc info and errors.
- notification—Trace notification events.
- routing-socket—Trace routing socket calls
- timeouts—Trace timeouts.
- timer—Trace internal timer events.

NOTE: The libgrpc trace flags are only supported in Junos OS Release 19.3R1. They allow you to see events from the grpc libraries in the jsd trace log. Within the grpc core there are multiple libraries (for example, iomgr, compression, and profiling).

level Set the trace log verbosity. Use the `error` option to only show error events.

> **NOTE**: This `error` option does not apply to the libgrpc trace flags.

no-remote-trace Disable remote tracing.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`level` option introduced in Junos OS Release 20.2R1.

Trace flag options `libgrpc-debug`, `libgrpc-errors`, and `libgrpc-info` are supported in Junos OS Release 19.3R1 only.

traceoptions (Services Application Identification)

IN THIS SECTION

- Syntax | 29840
- Hierarchy Level | 29840
- Description | 29840
- Options | 29840
- Required Privilege Level | 29841
- Release Information | 29842

Syntax

```
traceoptions {
    file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag all;
    level (all | error | info | notice | verbose | warning)
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services application-identification]
[edit services icap-redirect]
```

Description

Configure tracing operations for application identification services.

Options

- `file`—Configure the trace file options.
  - *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
  - `files` *`number`*—Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed to *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

    Range: 2 through 1000 files

Default: 10 files

- `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.
- `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches this size, it is renamed *`trace-file`*`.0`. When the `trace-file` again reaches its maximum size, *`trace-file`*`.0` is renamed *`trace-file`*`.1` and *`trace-file`* is renamed *`trace-file`*`.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

  Syntax: *`x`* `K` to specify KB, *`x`* `m` to specify MB, or *`x`* `g` to specify GB

  Range: 10 KB through 1 GB

  Default: 128 KB

- `world-readable` | `no-world-readable`—By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.

  `all`—Trace with all flags enabled.

- `level`—Set the level of debugging the output option.
  - `all`—Match all levels.
  - `error`—Match error conditions.
  - `info`—Match informational messages.
  - `notice`—Match conditions that should be handled specially
  - `verbose`—Match verbose messages.
  - `warning`—Match warning messages.
- `no-remote-trace`—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Application Identification Techniques

traceoptions (Services IPsec VPN)

IN THIS SECTION

- Syntax | 29842
- Hierarchy Level | 29842
- Description | 29843
- Options | 29843
- Required Privilege Level | 29844
- Release Information | 29844

Syntax

```
traceoptions {
    file <filename> <files number> <match regular-expression> <size bytes> <world-readable | no-
world-readable>;
    flag flag;
    level level;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services ipsec-vpn]
```

Description

Configure IPsec tracing operations. By default, messages are written to **/var/log/kmd**.

Options

`files` *`number`*—Maximum number of trace data files.

- **Range:** 2 through 1000

`flag` *`flag`*—Tracing operation to perform:

- `all`—Trace everything.
- `certificates`—Trace certificates that apply to the IPsec service set.
- `database`—Trace security associations database events.
- `general`—Trace general events.
- `ike`—Trace IKE module processing.
- `parse`—Trace configuration processing.
- `policy-manager`—Trace policy manager processing.
- `routing-socket`—Trace routing socket messages.
- `snmp`—Trace SNMP operations.
- `timer`—Trace internal timer events.

`level` *`level`*—Key management process (kmd) tracing level. The following values are supported:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`–Match informational messages.
- `notice`—Match conditions that should be handled specially.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

`size` *`bytes`*—Maximum trace file size.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

`level` option added in Junos OS Release 10.0.

traceoptions (Services L2TP)

IN THIS SECTION

- Syntax | 29844
- Hierarchy Level | 29845
- Description | 29845
- Options | 29845
- Required Privilege Level | 29848
- Release Information | 29848

Syntax

```
traceoptions {
    debug-level level;
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    filter {
        protocol name;
        user user@domain;
        user-name username;
    }
    flag flag;
```

```
    interfaces interface-name {
        debug-level level;
        flag flag;
    }
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Define tracing operations for L2TP processes.

Options

`debug-level` *`level`*—Trace level for PPP, L2TP, RADIUS, and UDP; this option does not apply to L2TP on MX Series routers:

- `detail`—Trace detailed debug information.
- `error`—Trace error information.
- `packet-dump`—Trace packet decoding information.

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`filter`—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.

- `protocol` *`name`*—One of the following protocols; this option does not apply to L2TP on MX Series routers:

  - `l2tp`
  - `ppp`
  - `radius`
  - `udp`

- `user` *user@domain*—Username of a subscriber; this option does not apply to L2TP on M Series routers. Optionally use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms.
- `user-name` *username*—Username of a subscriber; this option does not apply to L2TP on MX Series routers.

`flag` *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.
- `configuration`—Trace configuration events.
- `events`—Trace interface events.
- `general`—Trace general events.
- `gres`—Trace GRES events.
- `init`—Trace daemon initialization.
- `ipc-rx`—Trace IPC receive events.
- `ipc-tx`—Trace IPC transmit events.
- `memory`—Trace memory management code.
- `message`—Trace message processing code.
- `packet-error`—Trace packet error events.
- `parse`—Trace parsing events.
- `protocol`—Trace L2TP events.
- `receive-packets`—Trace received L2TP packets.
- `routing-process`—Trace routing process interactions.
- `routing-socket`—Trace routing socket events.

- `session-db`—Trace session database interactions.
- `states`—Trace state machine events.
- `timer`—Trace timer events.
- `transmit-packets`—Trace transmitted L2TP packets.
- `tunnel`—Trace tunnel events.

`interfaces` *`interface-name`*—Apply L2TP traceoptions to a specific services interface. This option does not apply to L2TP on MX Series routers.

- `debug-level` *`level`*—Trace level for the interface; this option does not apply to L2TP on MX Series routers:
  - `detail`—Trace detailed debug information.
  - `error`—Trace error information.
  - `extensive`—Trace all PIC debug information.
- `flag` *`flag`*—Tracing operation to perform for the interface. This option does not apply to L2TP on MX Series routers. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:
  - `all`—Trace everything.
  - `ipc`—Trace L2TP Inter-Process Communication (IPC) messages between the PIC and the Routing Engine.
  - `packet-dump`—Dump each packet content based on debug level.
  - `protocol`—Trace L2TP, PPP, and multilink handling.
  - `system`—Trace packet processing on the PIC.

`level`—Specify level of tracing to perform. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- `all`—Match messages of all levels.
- `error`—Match error messages.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.

- `verbose`—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure `verbose`, messages at all higher levels are traced. Therefore, the result is the same as when you configure `all`.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing L2TP Operations

Tracing L2TP Events for Troubleshooting

traceoptions (Services Logging)

IN THIS SECTION

- Syntax | 29849
- Hierarchy Level | 29849
- Description | 29849
- Options | 29849
- Required Privilege Level | 29851
- Release Information | 29851

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services adaptive-services-pics],
[edit services logging]
```

Description

Configure Adaptive Services or Multiservices PIC tracing operations. The messages are output to **/var/log/serviced**.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform:

- `all`—Trace everything.
- `command-queued`—Trace command enqueue events.
- `config`—Trace configuration events.
- `handshake`—Trace handshake events.
- `init`—Trace initialization events.
- `interfaces`—Trace interface events.
- `mib`—Trace GGSN SNMP MIB events.
- `removed-client`—Trace client cleanup events.
- `show`—Trace CLI command servicing.

`match` *`regex`*—(Optional) Match output to a defined regular expression (regex).

- **Default:** If you do not include this option, the trace operation output includes all lines relevant to the logged events.

`no-world-readable`—(Optional) Prevent any user from reading the log file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`x`*`k` to specify KB, *`x`*`m` to specify MB, or *`x`*`g` to specify GB

- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

`world-readable`—(Optional) Allow any user to read the log file.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`file` option added in Release 8.0.

RELATED DOCUMENTATION

Tracing Services PIC Operations

traceoptions (Services Redundancy Daemon)

IN THIS SECTION

- Syntax | 29852
- Hierarchy Level | 29852
- Description | 29852
- Options | 29852
- Required Privilege Level | 29854
- Release Information | 29854

Syntax

```
traceoptions {
    file file-name <files number> <no-word-readable | world-readable> <size size>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services redundancy-set]
```

Description

Specify the level of redundancy system events to be traced.

Options

file *file-name*
Name of the file to receive the output of the tracing operation.

files *number*
(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000 files
- **Default:** 3 files

flag *flag*
Specify which operations are to be traced. To specify more than one operation, include multiple flag statements.

- `all`—Trace everything.
- `config`—Trace services redundancy configuration events.
- `connect`—Trace services redundancy ipc events.
- `error`—Trace services redundancy errors.

- `database`—Trace services database events.
- `normal`—Trace normal events.
- `opcmd`—Trace services redundancy opcmd events.
- `parse`—Trace services redundancy parse events.
- `route`—Trace services redundancy route events.
- `snmp`—Trace services redundancy snmp events.
- `state`—Trace services redundancy set state-machine.
- `switchover`—Trace switchover events.
- `system`—Trace services redundancy system events.

level

Use the specified level of tracing. You can specify any of the following levels:

- `all`—Match all levels.
- `critical`—Match critical conditions.
- `error`—Match error conditions.
- `info`—Match informational messages
- `notice`—Match conditions that must be handled specially.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

no-remote-trace

(Optional) Disable remote tracing.

no-world-readable

(Optional) Disable unrestricted file access.

size *size*

(Optional) Use the maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the size option.

- **Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB.

- **Range:** 10,240 through 1,073,741,824 bytes.
- **Default:** 128 KB

word-readable (Optional) Enable unrestricted file access.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring the Service Redundancy Daemon

traceoptions (Services SSL)

IN THIS SECTION

- Syntax | 29854
- Hierarchy Level | 29855
- Description | 29855
- Options | 29855
- Required Privilege Level | 29856
- Release Information | 29857

Syntax

```
traceoptions {
    file {
```

```
        filename;
        files number;
        match regular-expression;
        size (Services) maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    level [brief | detail | extensive | verbose];
    no-remote-trace;
    packet-filter {
        destination-ip;
        destination-port;
        source-ip;
        source-port;
    }
}
```

Hierarchy Level

```
[edit services ssl]
```

Description

Specify the trace file information.

Debug tracing on both Routing Engine and the Packet Forwarding Engine can be enabled for SSL proxy by using `[edit services ssl traceoptions]` command.

Options

- *file-name*—Specify the name of file in which to write trace information.
  - `files`—Specify the maximum number of trace files. Range: 2 to 1000.
  - `match`—Specify the regular expression for lines to be logged. This statement is supported on the SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall.
  - `no-world-readable size`—Do not allow any user to read the log file.
  - `size`—Specify the maximum trace file size. Range: 10,240 to 1,073,741,824.

  - `world-readable`—Allow any user to read the log file.

- `flag`—Trace operation to perform. To specify more than one trace operation, include multiple `flag` statements.
  - *all*—Trace all the parameters.
  - *cli-configuration*—Trace CLI configuration events.
  - *initiation*—Trace initiation service events.
  - *proxy*—Trace proxy service events.
  - *selected-profile*—Trace events for profiles with `enable-flow-tracing` set.
  - *termination*—Trace termination service events.
- `level`—Set the level of debugging the output option.
  - `brief`—Match brief messages.
  - `detail`—Match detail messages.
  - `extensive`—Match extensive messages.
  - `verbose`—Match verbose messages.
- `no-remote-trace`—Set remote tracing as disabled.
- `packet-filter`—Set packet filter to capture the traffic details.
  - `destination-ip` *ipvaddress*—Specify a destination IP address.

    Range—1 through 65535
  - `destination-port` *port-number*—Specify a destination port.
  - `source-ip` *ip-address*—Specify a source IP address.
  - `source-port` *port-number*—Specify a source IP port.

    Range—1 through 65535

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10. This statement is supported on the SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall. Junos OS Release 19.3R1 introduces `packet-filter` statement.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

Firewall User Authentication Overview

traceoptions (Services UAC)

IN THIS SECTION

- Syntax | 29857
- Hierarchy Level | 29858
- Description | 29858
- Options | 29858
- Required Privilege Level | 29858
- Release Information | 29859

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
```

```
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services unified-access-control ]
```

Description

Define Unified Access Control (UAC) tracing options.

Use this statement when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series UAC Appliance.

Options

`flag`—Trace operation to perform. To specify more than one trace option, include multiple flag statements.

- `all`—Trace with all flags enabled
- `config`—Trace configuration information for all UAC-related configurations. This includes all configuration controlled through the `unified-access-control` statements at the `edit services` hierarchy level. It also includes other standard Junos OS configurations required for UAC enforcement such as zones, policies, and interfaces.
- `connect`—Trace communications between the Junos OS Enforcer and the IC Series appliance, including SSL handshakes and timeouts.
- `ipc`—Trace interprocess communications. Use this option to trace communications between the Routing Engine (RE) and the UACD enforcement plugin inside the Packet Forwarding Engine (PFE).

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Understanding Unified Access Control

Acquiring User Role Information from an Active Directory Authentication Server

traceoptions (Services User Identification)

IN THIS SECTION

- Syntax | 29859
- Hierarchy Level | 29860
- Description | 29860
- Options | 29860
- Required Privilege Level | 29861
- Release Information | 29862

Syntax

```
traceoptions {
    file filename files files match match size size(world-readable | no-world-readable);
    flag name;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services user-identification authentication-source (Services User Identification ClearPass)
aruba-clearpass]
```

Description

Specify the name of the trace log file and its characteristics. Messages about the behavior of the authentication source are written to this log file. Aruba ClearPass Policy Manager (CPPM) is the authentication source for the SRX Series Firewall integrated ClearPass authentication and enforcement feature.

file- Configure the name of the trace log file and its characteristics to which messages for the behavior of the authentication source are logged. For the SRX Series Firewall integrated ClearPass authentication and enforcement feature, the authentication source is the Aruba ClearPass Policy Manager (CPPM).

flag- Configure the tracing operation. To perform more than one tracing operation, include multiple *flag* statements. Messages about the behavior of the authentication source are written to this flag. Aruba ClearPass Policy Manager (CPPM) is the authentication source for the SRX Series Firewall integrated ClearPass authentication and enforcement feature.

level- Configure the level of messages that are written to trace log file about authentication source behavior.

For the integrated ClearPass authentication enforcement feature, the authentication source is Aruba ClearPass.

Options

file Trace file information.

filename Name of the log file.

files Specifies the maximum number of trace files.

- **Default:** 3
- **Range:** 2 through 1000

match Specifies a regular expression that determines which lines are logged.

no-world-readable Denies users the ability to read the log file.

size Specifies the trace file maximum file size.

- **Default:** 128k
- **Range:** 10,240 through 1,073,741,824

world-readable Allows users to read the log file.

flag Tracing parameters.

- Values:
  - all—Trace Aruba ClearPass all modules
  - clearpass-authentication—Trace Aruba ClearPass auth table management module
  - configuration—Trace Aruba ClearPass configuration
  - dispatcher—Trace dispatcher module
  - ipc—Trace ipc
  - user-query—Trace user-query module

level Level of debugging output

- Values:
  - all—Matches all levels.
  - error—Matches error conditions.
  - info—Matches informational messages.
  - notice—Matches conditions that should be handled specially.
  - verbose—Matches verbose messages.
  - warning—Matches warning messages.

no-remote-trace Disable remote tracing

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

1. services—To view this statement in the configuration.

2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

traceoptions (sFlow Technology)

IN THIS SECTION

- Syntax | 29862
- Hierarchy Level | 29862
- Description | 29863
- Default | 29863
- Options | 29863
- Required Privilege Level | 29864
- Release Information | 29864

Syntax

```
traceoptions {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit protocols sflow]
```

Description

Define tracing operations for sFlow technology.

Default

The traceoptions feature is disabled.

> **NOTE**: In Junos OS Evolved, the `traceoptions` functions have been moved to the trace application command. Use trace application commands instead. See Junos OS Evolved Differences and No Link Title for details.

Options

file *filename*
Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Output files are located in the **/var/log/** directory.

files *number*
(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**. Incoming trace file data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify the maximum number of files, you must also specify the maximum file size using the `size` option.

- **Range:** 2 through 1000 files
- **Default:** 1 trace file

flag *flag*
Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements.

- all—Trace all sFlow monitoring events.
- client-server—Trace sFlow monitoring client-server events.
- configuration—Trace sFlow monitoring configuration events.
- interface—Trace sFlow monitoring interface events.
- rtsock—Trace routing socket code events.

no-stamp (Optional) Do not place timestamp information at the beginning of each line in the trace file.

no-world-readable (Optional) Prevent any user from reading the trace file.

replace (Optional) Replace an existing trace file if there is one.

size *size* (Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**. Incoming trace file data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the files option.

- **Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB
- **Range:** 10 KB through the maximum file size of 4 GB
- **Default:** 128 KB

world-readable (Optional) Allow any user to read the trace file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

Overview of sFlow Technology

traceoptions (SMTP)

IN THIS SECTION

- Syntax | 29865
- Hierarchy Level | 29865
- Description | 29865
- Options | 29865
- Required Privilege Level | 29866
- Release Information | 29866

Syntax

```
traceoptions {
    flag {
        all;
        configuration;
        IPC;
        protocol-exchange;
        send-request;
    }
}
```

Hierarchy Level

```
[edit smtp]
```

Description

Set the Simple Mail Transfer Protocol (SMTP) traceoptions.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 10.0.

RELATED DOCUMENTATION

Content Security | 4589

traceoptions (SNMP)

IN THIS SECTION

- Syntax | 29866
- Hierarchy Level | 29867
- Description | 29867
- Options | 29867
- Required Privilege Level | 29868
- Release Information | 29869

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Track the activities of SNMP agents on the switch and record the information in log files.

> **NOTE**: The `traceoptions` statement is not supported on the QFabric system.

The output of the tracing operations is placed into log files in the `/var/log` directory. Each log file is named after the SNMP agent that generates it. Currently, the following logs are created in the `/var/log` directory when the `traceoptions` statement is used:

- chassisd
- craftd
- ilmid
- mib2d
- rmopd
- serviced
- snmpd

Options

`file` *`filename`*—By default, the name of the log file that records trace output is the name of the process being traced (for example, `mib2d` or `snmpd`). Use this option to specify another name.

`files` *`number`*—(Optional) Maximum number of trace files per SNMP subagent. When a trace file (for example, `snmpd`) reaches its maximum size, it is archived by being renamed to `snmpd.0`. The previous `snmpd.1` is renamed to `snmpd.2`, and so on. The oldest archived file is deleted.

- **Range:** 2 through 1000 files
- **Default:** 10 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements:

- `all`—Log all SNMP events.
- `configuration`—Log reading of configuration at the `[edit snmp]` hierarchy level.
- `database`—Log events involving storage and retrieval in the events database.
- `events`—Log important events.
- `general`—Log general events.
- `interface-stats`—Log physical and logical interface statistics.
- `nonvolatile-sets`—Log nonvolatile SNMP set request handling.
- `pdu`—Log SNMP request and response packets.
- `policy`—Log policy processing.
- `protocol-timeouts`—Log SNMP response timeouts.
- `routing-socket`—Log routing socket calls.
- `server`—Log communication with processes that are generating events.
- `subagent`—Log subagent restarts.
- `timer-events`—Log internally generated events.
- `varbind-error`—Log variable binding errors.

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`size` *`size`*—(Optional) Maximum size, in kilobytes (KB), of each trace file before it is closed and archived.

- **Range:** 10 KB through 1 GB
- **Default:** 1000 KB

`world-readable | no-world-readable`—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`file filename` option added in Junos OS Release 8.1.

`world-readable | no-world-readable` option added in Junos OS Release 8.1.

`match regular-expression` option added in Junos OS Release 8.1.

RELATED DOCUMENTATION

Trace SNMP Activity on a Device Running Junos OS

traceoptions (Spanning Tree)

IN THIS SECTION

- Syntax | 29869
- Hierarchy Level | 29870
- Description | 29870
- Default | 29870
- Options | 29870
- Required Privilege Level | 29872
- Release Information | 29873

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp | vstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp | vstp)],
```

```
[edit protocols (mstp | rstp | vstp | vstp vlan vlan-id)],
```

```
[edit routing-instances routing-instance-name protocols (mstp | rstp | vstp)]
```

Description

Sets protocol-level tracing options for for spanning-tree protocols MPLS, MVRP, STP, RSTP, MSTP, and VSTP.

Default

The default STP protocol-level trace options are inherited from the global `traceoptions` statement. Traceoptions is disabled.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place STP tracing output in the file **/var/log/stp-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 1 trace file only

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. The following are the STP-specific tracing options:

- **all**—Trace all operations.
- **all-failures**—Trace all failure conditions.
- **bpdu**—Trace BPDU reception and transmission.
- **bridge-detection-state-machine**—Trace the bridge detection state machine.
- `error`—Trace all failure conditions.
- **events**—Trace events of the protocol state machine.
- `pdu`—Trace PDUs that were received and sent.
- **port-information-state-machine**—Trace the port information state machine.
- **port-migration-state-machine**—Trace the port migration state machine.
- **port-receive-state-machine**—Trace the port receive state machine.
- **port-role-transit-state-machine**—Trace the port role transit state machine.
- **port-role-select-state-machine**—Trace the port role selection state machine.
- **port-state-transit-state-machine**—Trace the port state transit state machine.
- **port-transmit-state-machine**—Trace the port transmit state machine.
- **ppmd**—Trace the state and events for the ppmd process.
- `socket`—Trace socket activity.
- `state-machine`—Trace state machine information.
- **state-machine-variables**—Trace when the state machine variables change.
- **timers**—Trace protocol timers.
- **topology-change-state-machine**—Trace the topology change state machine.

The following are the global tracing options:

- **all**—All tracing operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.

- **normal**—All normal events.
- **Default:** If you do not specify this option, only unusual or abnormal operations are traced.
- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.
- **regex-parse**—Trace regular-expression parsing.
- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

no-world-readable—(Optional) Prevent any user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

- **Default:** If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the files option.

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding RSTP

Understanding STP

Understanding VSTP

Understanding Multiple VLAN Registration Protocol (MVRP)

Understanding Spanning-Tree Protocol Trace Options

Configuring Tracing Spanning-Tree Operations

Understanding MSTP

Example: Configuring Network Regions for VLANs with MSTP on Switches

Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

Example: Tracing Spanning-Tree Protocol Operations

show spanning-tree bridge | **26548**

show spanning-tree interface | **26555**

traceoptions (Static Subscribers)

IN THIS SECTION

- Syntax | **29874**
- Hierarchy Level | **29874**
- Description | **29874**
- Options | **29874**
- Required Privilege Level | **29876**
- Release Information | **29876**

Syntax

```
traceoptions {
    file filename<files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name system processes static-subscribers],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
processes static-subscribers],
[edit routing-instances routing-instances-name system processes static-subscribers],
[edit system processes static-subscribers]
```

Description

Define tracing operations for static subscriber processes.

Options

`file` *`filename`*— Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.
- `authentication`—Trace authentication events.

- `configuration`—Trace configuration events.
- `database`—Trace database events.
- `general`—Trace general events.
- `gres`—Trace GRES events.
- `profile`—Trace dynamic profile events.
- `rtsock`—Trace routing socket events.
- `statistics`—Trace statistics events.
- `subscriber`—Trace subscriber events.

`level`—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—(Optional) Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Tracing Static Subscriber Events for Troubleshooting

traceoptions (Subscriber Management)

IN THIS SECTION

- Syntax | 29876
- Hierarchy Level | 29877
- Description | 29877
- Options | 29877
- Required Privilege Level | 29878
- Release Information | 29878

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit system services subscriber-management]
```

Description

Define tracing operations for subscriber management interface processes.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the filename within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.
- `database`—Trace database events.
- `general`—Trace general events.
- `issu`—Trace unified ISSU events.
- `server`—Trace server events.
- `session-db`—Trace session database interactions.
- `ui`—Trace user interface events.

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *size*k to specify KB, *size*m to specify MB, or *size*g to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Tracing Subscriber Management Database Events for Troubleshooting

traceoptions (Subscriber Session Database Replication)

IN THIS SECTION

- Syntax | 29879
- Hierarchy Level | 29879
- Description | 29879
- Options | 29879
- Required Privilege Level | 29880
- Release Information | 29880

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system services database-replication]
```

Description

Define tracing operations for subscriber management session database replication processes.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.
- `database`—Trace database events.
- `general`—Trace general flow.
- `mirror`—Trace mirroring events.
- `replication`—Trace database replication events.

- `server`—Trace server events.
- `session-db`—Trace session database interactions.
- `ui`—Trace user interface events.

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

> *Tracing Subscriber Management Session Database Replication Events for Troubleshooting*

traceoptions (System Accounting)

IN THIS SECTION

- Syntax | 29881
- Hierarchy Level | 29881
- Description | 29881
- Default | 29881
- Options | 29882
- Required Privilege Level | 29883
- Release Information | 29883

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag (all| config | events | radius | tacplus);
    no-remote-trace
}
```

Hierarchy Level

```
[edit system accounting]]
```

Description

Define tracing operations for System Accounting.

Default

Trace options are not enabled by default.

Options

`file` *`filename`*—Name of the file in which Junos OS stores the accounting logs. By default, this is created under the /var/log directory.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file reaches the size specified by the size option, the filename is appended with 0 and compressed. For example, when trace file named trace-file-log reaches size *<size>*, it is renamed and compressed to trace-file-log.0.gz. When trace-file-log reaches size *<size>* or the second time, the trace-file-log.0.gz is renamed to trace-file-log.1.gz and trace-file-log is renamed and compressed to trace-file-log.0.gz. This renaming shceme ensures that the older logs to have a greater index number. When number of trace files reach *<number>* then the oldest file is deleted.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` option and a filename.

- **Range:** 2 through 1000 files
- **Default:** 10

`flag` *`flag`*—Tracing operation to perform. You can include one or more of the following flags:

- `all`—Trace all operations.
- `config`—Trace configuration processing.
- `events`—Trace accounting events and their processing.
- `radius`—Trace RADIUS processing.
- `tacplus`—Trace TACPLUS processing.

`no-remote-trace`—(Optional) Disable tracing and logging operations that track normal operations, error conditions, and packets that are generated by or passed through the Juniper Networks device.

`no-world-readable`—Restrict access to the trace files to the owner.

- **Default:** no-world-readable

`size` *`size`*—(Optional) Maximum size of each trace file in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). If you do not specify a unit, the default is bytes. If you specify a maximum file size, you also must specify a maximum number of trace files by using the files option and a filename by using the file option.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** size to specify bytes, sizek to specify KB, sizem to specify MB, or sizeg to specify GB.

- **Range:** 10 KB through 1 MB
- **Default:** 128 KB

`world-readable`—Enable any user to access the trace files.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

`tacplus` option introduced in Junos OS Release 15.1.

traceoptions (System License)

IN THIS SECTION

- Syntax | 29883
- Hierarchy Level | 29884
- Description | 29884
- Options | 29884
- Required Privilege Level | 29885
- Release Information | 29885

Syntax

```
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
```

```
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit system license]
```

Description

Set trace options for licenses.

Options

file Configure the trace file information.

- ***filename*** Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
- **files *number*** Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum number of files, you also must specify a maximum file size with the `size maximum file-size` option.

  - **Range:** 2 through 1000 files
  - **Default:** 10 files
- **match *regular-expression*** Refine the output to include lines that contain the regular expression.
- **size *size*** Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files number` option.

- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

world-readable | no-world-readable By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

flag *flag* Specify which tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags.

- `all`—Trace all operations.
- `config`—Trace license configuration processing.
- `events`—Trace licensing events and their processing.

no-remote-trace Disable the remote tracing.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

traceoptions (System Monitor)

IN THIS SECTION

- Syntax | 29886
- Hierarchy Level | 29886
- Description | 29886
- Options | 29886

Required Privilege Level | 29887
Release Information | 29887

Syntax

```
traceoptions {
    file {
        files files;
        size size;
    }
    level [ all | debug | info| warning| error | critical ];
}
```

Hierarchy Level

```
[edit system monitor]
```

Description

Configure the system memory monitoring trace options on your PTX and MX Series routers. The memory monitoring system on your device monitors anomalies with respect to the memory utilization of Junos applications and overall system in the control plane of your device. Apart from monitoring the system memory, you can also configure the device to monitor the per-process memory footprint as well to help detect possible leaks and take configured action.

Options

file *filename*
Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. The files are available at the **/var/log/jsysmon.log** .

files *number*
(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the size option.

- **Range:** 2 through 1000 files
- **Default:** 10 files

Level Specify the level for trace logging.

all	Trace all messages.
critical	Trace critical messages.
debug	Trace debug messages.
error	Trace error messages.
info	Trace information messages.
warning	Trace warning messages.

- **Default:** error

Size *size* (Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the files option.

- **Syntax:** xk to specify KB, xm to specify MB, and xg to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 128 KB

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

System Log Overview

traceoptions (System Services DNS)

IN THIS SECTION

- Syntax | 29888
- Hierarchy Level | 29888
- Description | 29888
- Options | 29889
- Required Privilege Level | 29890
- Release Information | 29890

Syntax

```
traceoptions {
    category {
        category-type;
    }
    file;
}
```

Hierarchy Level

```
[edit system services dns]
```

Description

Defines tracing options for DNS services.

Options

`category`—Specifies the logging category. See Table 1812 on page 29889 for the different logging categories and their descriptions.

`file`—Trace file information.

Table 1812: Category Names

Category Name	Description
client	Processing of client requests
config	Configuration file parsing and processing
database	Messages relating to the databases
default	Categories for which there is no specific configuration
delegation-only	Delegation only
dispatch	Dispatching of incoming packets to the server
dnssec	DNSSEC and TSIG protocol processing
edns-disabled	Log query using plain DNS
general	General information
lame-servers	Lame servers
network	Network options
notify	NOTIFY protocol
queries	DNS query resolver

Table 1812: Category Names *(Continued)*

Category Name	Description
resolver	DNS resolution security
security	Approval and denial of requests
unmatched	Unable to determine the class for messages named
update	Dynamic updates
update-security	Approval and denial of update requests
xfer-in	Zone transfers that the server is receiving xfer-out
xfer-out	Zone transfers that the server is sending

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

DNS Overview

traceoptions (TCP Encapsulation)

IN THIS SECTION

- Syntax | 29891
- Hierarchy Level | 29891
- Description | 29891
- Options | 29892
- Required Privilege Level | 29893
- Release Information | 29893

Syntax

```
traceoptions {
    file filename {
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag (all | configuration | session | tunnel);
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit security tcp-encap]
```

Description

Configure TCP encapsulation tracing options.

Options

file Configure the trace file options.

- *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.

- files *number*—Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed to *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

  If you specify a maximum number of files, you also must specify a maximum file size with the size option and a filename.

  Range: 2 through 1000 files

  Default: 10 files

- match *regular-expression*—Refine the output to include lines that contain the regular expression.

- size *maximum-file-size*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0. When *trace-file*.0 reaches its maximum size, it is renamed *trace-file*.1 and *trace-file* is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum file size, you also must specify a maximum number of trace files with the files option and filename.

  Syntax: *x* k to specify KB, *x* m to specify MB, or *x* g to specify GB

  Range: 10 KB through 1 GB

  Default: 128 KB

- world-readable | no-world-readable—By default, log files can be accessed only by the user who configures the tracing operation. The world-readable option enables any user to read the file. To explicitly set the default behavior, use the no-world-readable option.

flag Trace operation to perform. To specify more than one trace operation, include multiple flag statements.

- all—Trace all activity.

- configuration—Trace configuration events.

- `session`—Trace session related events.
- `tunnel`—Trace tunnel events.

level Set the level of debugging.

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match conditions that should be handled specifically.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

no-remote-trace Disable remote tracing.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D80.

RELATED DOCUMENTATION

Understanding SSL Remote Access VPNs with NCP Exclusive Remote Access Client

tcp-encap | **28980**

traceoptions (TCP Port Forwarding)

IN THIS SECTION

- Syntax | 29894
- Hierarchy Level | 29894
- Description | 29894
- Options | 29894
- Required Privilege Level | 29896
- Release Information | 29896

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
}
```

Hierarchy Level

```
[edit system processes tcp-forwarding]
```

Description

Define tracing operations for TCP port forwarding processes.

Options

`file` *filename*—Name of the file to receive the output of the tracing operation. Enclose the filename within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all events
- `configuration`—Trace configuration events
- `connection`—Trace TCP connection events
- `init`—Trace TCP port forwarding initialization events

`level`—Level of tracing to perform; also known as severity level. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error messages.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure `verbose`, messages at all higher levels are traced. Therefore, the result is the same as when you configure `all`.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—(Optional) Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB

- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Tracing TCP Port Forwarding Events for Troubleshooting

TCP Port Forwarding for Remote Device Management

traceoptions (TDF Gateway)

IN THIS SECTION

- Syntax | 29897
- Hierarchy Level | 29897
- Description | 29897
- Options | 29897
- Required Privilege Level | 29899
- Release Information | 29899

Syntax

```
traceoptions {
    file file-name <files number> <no-word-readable | world-readable> <size size>;
    flag flag;
    level (all | critical | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name]
```

Description

Specify tracing options for the TDF gateway.

Options

file *file-name*

Use the specified name of the file to receive the output of the tracing operation.

files *number*

(Optional) Use the specified maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000 files
- **Default:** 3 files

flag *flag*

Specify which operations are to be traced. To specify more than one operation, include multiple flag statements.

- `all`—Trace everything.
- `bulkjob`—Trace events that are handled by bulk jobs in order to prevent system overload.
- `config`—Trace configuration events.

- `cos-cac`—Trace class of service (CoS) and call admission control (CAC) events.
- `ctxt`—Trace user equipment, Packet Data Network (PDN), or bearer context events.
- `fsm`—Trace mobile subscriber finite state machine (FSM) events.
- `gtpu`—Trace GPRS tunneling protocol, user plane (GTP-U) events.
- `ha`—Trace high availability events.
- `init`—Trace initialization events.
- `pfem`—Trace Packet Forwarding Engine Manager events.
- `stats`—Trace `stats` events. This flag is used internally by Juniper Networks engineers.
- `waitq`—Trace `waitq` events. This flag is used internally by Juniper Networks engineers.

level

Use the specified level of tracing. You can specify any of the following levels:

- `all`—Match all levels.
- `critical`—Match critical conditions.
- `error`—Match error conditions.
- `info`—Match informational messages
- `notice`—Match conditions that must be handled specially.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

no-remote-trace

(Optional) Disable remote tracing.

no-world-readable

(Optional) Disable unrestricted file access.

size *size*

(Optional) Use the specified maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the size option.

- **Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB.
- **Range:** 10,240 through 1,073,741,824 bytes.
- **Default:** 128 KB

word-readable (Optional) Enable unrestricted file access.

Required Privilege Level

trace and unified-edge—To view this statement in the configuration.

trace-control and unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring Tracing for TDF Gateway

traceoptions (Traffic Load Balancer)

IN THIS SECTION

- Syntax | 29900
- Hierarchy Level | 29900
- Description | 29900
- Options | 29900
- Required Privilege Level | 29902
- Release Information | 29903

Syntax

```
traceoptions {
    file file-name <files number> <no-word-readable | world-readable> <size size>;
    flag flag;
    level (all | critical | error | info | notice | verbose | warning);
    monitor monitor-object-name {
        instance-name instance-name;
        virtual-svc-name virtual-service-name;
    }
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services traffic-load-balance]
```

Description

Configure tracing options for the traffic load balancer.

Options

For Next Gen Services on the MX-SPC3 services card, set the *monitor-object-name* to either:

file *file-name*
Name of the file to receive the output of the tracing operation.

files *number*
(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

- **Range:** 2 through 1000 files
- **Default:** 3 files

flag *flag*
Specify which operations you want to trace from Table 1813 on page 29901. To specify more than one operation, include multiple flag statements.

Table 1813: Trace Flags

Flag	Support on MS-MPC and MX-SPC3 Cards	Description
all	MS-MPC and MX-SPC3	Trace all operations.
all-real-services	MS-MPC and MX-SPC3	Trace all real services.
database	MS-MPC and MX-SPC3	Trace database events.
file-descriptor-queue	MS-MPC and MX-SPC3	Trace file descriptor queue events.
inter-thread	MS-MPC and MX-SPC3	Trace inter-thread communication events.
messages	MS-MPC and MX-SPC3	Trace normal events.
probe	MS-MPC and MX-SPC3	Trace probe events.
probe-infra	MS-MPC and MX-SPC3	Trace probe infra events.

instance-name *instance-name*

(Optional) Name of the TLB instance to monitor.

level

Use the specified level of tracing. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.

- `notice`—Match conditions that must be handled specially.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

These trace levels are available for both the MS-MPC and MX-SPC3 services cards unless otherwise specified.

monitor *monitor-object-name* — Name of a monitoring object that contains an instance name or virtual service name.

no-remote-trace — (Optional) Disable remote tracing.

no-world-readable — (Optional) Disable unrestricted file access.

group-name — Name of the group.

real-services-name — Name of the real service

size *size* — (Optional) Use the maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the size option.

- **Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB.
- **Range:** 10,240 through 1,073,741,824 bytes.
- **Default:** 128 KB

virtual-svc-name *virtual-service-name* — (Optional) Name of the virtual service to monitor.

word-readable — (Optional) Enable unrestricted file access.

Required Privilege Level

trace and interface—To view this statement in the configuration.

trace-control and interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

`instance-name` and `virtual-service-name` options added in Junos OS Release 16.1R6 and 18.2R1 on MX Series.

Support for Next Gen Services MX-SPC3 services card add in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

> *Traffic Load Balancer Overview*
>
> *Configuring TLB*

traceoptions (Uplink Failure Detection)

IN THIS SECTION

- Syntax | 29903
- Hierarchy Level | 29904
- Description | 29904
- Default | 29904
- Options | 29904
- Required Privilege Level | 29905
- Release Information | 29905

Syntax

```
traceoptions  {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit protocols uplink-failure-detection]
```

Description

Define tracing operations for uplink failure detection.

Default

The `traceoptions` feature is disabled by default.

Options

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached (*x*`k` to specify KB, *x*`m` to specify MB, or *x*`g` to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- `all`—Trace everything.
- `dcd`—Trace ufdd interaction with dcd.
- `groups`—Trace uplink-failure-detection group handling.
- `interface`—Trace interface notifcation handlers of ufdd.
- `parse`—Trace configuration parsing.

`no-stamp`—(Optional) Do not place a timestamp on any trace file.

`no-world-readable`—(Optional) Restricted file access to the user who created the file.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `files` option.

- **Syntax:** *`x`*`k` to specify KB, *`x`*`m` to specify MB, or *`x`*`g` to specify gigabytes
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Interfaces for Uplink Failure Detection (CLI Procedure)

traceoptions (Virtual Chassis)

IN THIS SECTION

- Syntax | 29906
- Hierarchy Level | 29906
- Description | 29906
- Default | 29906
- Options | 29906
- Required Privilege Level | 29908

Release Information | 29908

Syntax

```
traceoptions  {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
    flag flag <detail> <disable> <receive> <send>;
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Define tracing operations for the Virtual Chassis or VCF.

Each Virtual Chassis or VCF member is updated with configured tracing options, and the log file is stored locally on each device.

NOTE: In a large-scale VCF, use system logging and tracing with discretion, and only for troubleshooting. These operations place an extra load on the primary Routing Engine device, which can impact VCF convergence time and stability. All tracing options should be disabled during normal VCF operation. To troubleshoot particular problems, selectively enable tracing options, and disable them again after collecting the desired information.

Default

Tracing operations are disabled.

Options

`detail`—(Optional) Generate detailed trace information for a flag.

> **NOTE**: Enable tracing at the detail level only while troubleshooting a particular issue, and disable it again for normal system operation.

`disable`—(Optional) Disable a flag.

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- `all`—All tracing operations.

  > **TIP**: The `all` flag displays a subset of logs that are generally useful in debugging issues, and for more detailed information, you can use options `all detail`. However, to avoid significantly impacting VCF stability, use this high level of tracing only for troubleshooting for a short period of time, and not during normal VCF operation.

- `auto-configuration`—Trace Virtual Chassis ports (VCPs) that have been automatically configured.
- `csn`—Trace Virtual Chassis complete sequence number (CSN) packets.
- `error`—Trace Virtual Chassis errored packets.
- `hello`—Trace Virtual Chassis hello packets.
- `krt`—Trace Virtual Chassis KRT events.
- `lsp`—Trace Virtual Chassis link-state packets.
- `lsp-generation`—Trace Virtual Chassis link-state packet generation.
- `me`—Trace Virtual Chassis ME events.
- `normal`—Trace normal events.

- `packets`—Trace Virtual Chassis packets.
- `parse`—Trace reading of the configuration.
- `psn`—Trace partial sequence number (PSN) packets.
- `route`—Trace Virtual Chassis routing information.
- `spf`—Trace Virtual Chassis SPF events.
- `state`—Trace Virtual Chassis state transitions.
- `task`—Trace Virtual Chassis task operations.

`no-stamp`—(Optional) Do not place a timestamp on any trace file.

`no-world-readable`—(Optional) Restrict file access to the user who created the file.

`receive`—(Optional) Trace received packets.

`replace`—(Optional) Replace a trace file rather than appending information to it.

`send`—(Optional) Trace transmitted packets.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `files` option.

- **Syntax:** *`x`*`k` to specify KB, *`x`*`m` to specify MB, or *`x`*`g` to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Option `detail` added in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

- Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis
- Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member
- Verifying That Virtual Chassis Ports Are Operational
- Troubleshooting an EX Series Virtual Chassis
- *Troubleshooting Virtual Chassis Fabric*

traceroute

IN THIS SECTION

- Syntax | **29909**
- Syntax (QFX Series and OCX Series) | **29910**
- Description | **29911**
- Options | **29911**
- Required Privilege Level | **29913**
- Output Fields | **29913**
- Sample Output | **29913**

Syntax

```
traceroute
<host>
<as-number-lookup>
<bypass-routing>
<ce-ip ip address>
<clns>
<ethernet host>
<extension>
<gateway address>
<inet>
<inet6>
<interface interface-name>
```

```
<logical-system logical-system-name>
<monitor host>
<mpls (ldp FEC address | rsvp label-switched-path-name | bgp FEC address)>
<next-hop next-hop address>
<no-resolve host>
<overlay>
<port>
<propagate-ttl host>
<routing-instance routing-instance-name>
<tenant tenant-name>
<source source-address>
<tos value>
<ttl value>
<wait seconds>
```

Syntax (QFX Series and OCX Series)

```
traceroute
<host>
<as-number-lookup>
<bypass-routing>
<gateway address>
<inet>
<inet6>
<interface interface-name>
<monitor host>
<mpls (ldp FEC address | rsvp label-switched-path-name | bgp FEC address)>
<next-hop next-hop address>
<no-resolve host>
<overlay>
<port>
<propagate-ttl host>
<routing-instance routing-instance-name>
<source source-address>
<tos value>
<ttl value>
<wait seconds>
```

Description

Display the route that packets take to a specified network host. Use `traceroute` as a debugging tool to locate points of failure in a network.

> **NOTE**: The Junos OS Evolved `traceroute` command parses data in the same way as the Linux `traceroute` command, so the output is different compared to Junos OS. See the "Sample Output" on page 29913 section for more details.

Options

host	IP address or name of remote host.
as-number-lookup	(Optional) Display the autonomous system (AS) number of each intermediate hop on the path from the host to the destination.
bypass-routing	(Optional) Bypass the normal routing tables and send requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to display a route to a local system through an interface that has no route through it.
ce-ip *IPv4 or IPv6 address*	(MX Series routers with MPC and MIC interfaces only) (Optional) Check the route to a customer edge (CE) IP address in a virtual private LAN service (VPLS) and Ethernet VPN (EVPN) network.
clns	(Optional) Trace the route belonging to Connectionless Network Service (CLNS).
ethernet *host*	(Optional) Trace the route to an Ethernet host using unicast MAC address.
extension	(Optional) Trace ICMP extensions
gateway *address*	(Optional) Address of a router or switch through which the route transits.
inet | inet6	(Optional) Trace the route belonging to IPv4 or IPv6, respectively.
interface *interface-name*	(Junos OS only) (Optional) Name of the interface over which to send packets.
logical-system (all | *logical-system-name***)**	(Optional) This option is not supported for Junos OS Evolved Release 18.3R1. Perform this operation on all logical systems or on a particular logical system.

tenant *tenant-name*	(Optional) Name of a particular tenant system for traceroute attempt.
monitor *host*	(Optional) Display real-time monitoring information for the specified host.
mpls (ldp *FEC address* **| rsvp** *label-switched-path name***)**	(Optional) See `traceroute mpls ldp` and `traceroute mpls rsvp`.
next-hop	The next-hop through which to send packets to a destination.
no-resolve	(Optional) Do not attempt to determine the hostname that corresponds to the IP address.
overlay	(Optional) Traceroute overlay path.
port	(Optional) Base port number to use in traceroute probes.
propagate-ttl	(Optional) On the PE routing device, use this option to view locally generated Routing Engine transit traffic. This is applicable for MPLS L3VPN traffic only. Use for troubleshooting, when you want to view hop-by-hop information from the local provider router to the remote provider router, when TTL decrementing is disabled on the core network using the `no-proagate-ttl` configuration statement. **NOTE**: Using `propagate-ttl` with `traceroute` on the CE router does not show hop-by-hop information.
routing-instance *routing-instance-name*	(Optional) Name of the routing instance for the traceroute attempt.
source *source-address*	(Optional) Source address of the outgoing traceroute packets.
tos *value*	(Optional) Value to include in the IP type-of-service (ToS) field. The range of values is `0` through `255`.
ttl *value*	(Optional) Maximum time-to-live value to include in the traceroute request. The range of values is `0` through `128`.
wait *seconds*	(Optional) Maximum time to wait for a response to the traceroute request.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

network

Output Fields

Table 1814 on page 29913 describes the output fields for the `traceroute` command. Output fields are listed in the approximate order in which they appear.

Table 1814: traceroute Output Fields

Field Name	Field Description
`traceroute to`	IP address of the receiver.
`hops max`	Maximum number of hops allowed.
`byte packets`	Size of packets being sent.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
`Round trip time`	Average round-trip time, in milliseconds (ms).

Sample Output

traceroute

```
user@host> traceroute santacruz
traceroute to host1.example.com (10.156.169.254), 30 hops max, 40 byte packets
 1  blue23 (10.168.1.254)    2.370 ms  2.853 ms  0.367 ms
 2  red14  (10.168.255.250)  0.778 ms  2.937 ms  0.446 ms
```

```
 3  yellow (10.156.169.254)  7.737 ms  89.905 ms  0.834 ms
```

traceroute as-number-lookup host

```
user@host> traceroute as-number-lookup 10.100.1.1
traceroute to 10.100.1.1 (10.100.1.1), 30 hops max, 40 byte packets
 1  10.39.1.1 (10.39.1.1)  0.779 ms  0.728 ms  0.562 ms
 2  10.39.1.6 (10.39.1.6) [AS  32]  0.657 ms  0.611 ms  0.617 ms
 3  10.100.1.1 (10.100.1.1) [AS  10, 40, 50]  0.880 ms  0.808 ms  0.774 ms


user@host> traceroute as-number-lookup 1::1
traceroute6 to 1::1 (1::1) from 2001:b8::7, 64 hops max, 12 byte packets
user@host> traceroute 2001:b8::7 as-number-lookup
traceroute6 to 2001:b8::7 (2001:b8::7) from 2001:db8::9, 64 hops max, 12 byte packets
 1  2001:db8::10 (2001:db8::10) [AS 18] 0.657 ms  17.319 ms  0.504 ms
 2  2001:b8::7 (2001:b8::7)  0.949 ms  0.930 ms  0.739 ms
```

traceroute no-resolve

```
user@host> traceroute santacruz no-resolve
traceroute to host1.example.com (10.156.169.254), 30 hops max, 40 byte packets
 1  10.168.1.254  0.458 ms  0.370 ms  0.365 ms
 2  10.168.255.250  0.474 ms  0.450 ms  0.444 ms
 3  10.156.169.254  0.931 ms  0.876 ms  0.862 ms
```

traceroute propogate-ttl

```
user@host> traceroute propagate-ttl 100.200.2.2 routing-instance VPN-A
traceroute to 100.200.2.2 (100.200.2.2) from 1.1.0.2, 30 hops max, 40 byte packets
 1  1.2.0.2 (1.2.0.2)  2.456 ms  1.753 ms  1.672 ms
     MPLS Label=299776 CoS=0 TTL=1 S=0
     MPLS Label=299792 CoS=0 TTL=1 S=1
 2  1.3.0.2 (1.3.0.2)  1.213 ms  1.225 ms  1.166 ms
     MPLS Label=299792 CoS=0 TTL=1 S=1
```

```
 3  100.200.2.2 (100.200.2.2)  1.422 ms  1.521 ms  1.443 ms
```

traceroute (Between CE Routers, Layer 3 VPN)

```
user@host> traceroute vpn09
traceroute to host2.example.com (10.255.14.179), 30 hops max, 40
      byte packets
       1  10.39.10.21 (10.39.10.21)  0.598 ms  0.500 ms  0.461 ms
       2  10.39.1.13 (10.39.1.13)  0.796 ms  0.775 ms  0.806 ms
           MPLS Label=100006 CoS=0 TTL=1 S=1
       3  host2.example.com (10.255.14.179)  0.783 ms  0.716 ms  0.686
```

traceroute (Through an MPLS LSP)

```
user@host> traceroute mpls1
traceroute to 10.168.1.224 (10.168.1.224), 30 hops max, 40 byte packets
 1  mpls1-sr0.company.net (10.168.200.101)  0.555 ms  0.393 ms  0.367 ms
     MPLS Label=1024 CoS=0 TTL=1
 2  mpls5-lo0.company.net (10.168.1.224)  0.420 ms  0.394 ms  0.401 ms
```

traceroute routing-instance no-resolve (Through an MPLS LSP)

```
user@host> traceroute routing-instance VRF-1 198.51.100.1 no-resolve
traceroute to 198.51.100.1 (198.51.100.1), 30 hops max, 40 byte packets
 1  198.51.100.20  20.243 ms  13.256 ms  24.194 ms
     MPLS Label=299792 CoS=0 TTL=1 S=0
     MPLS Label=16 CoS=0 TTL=1 S=1
 2  198.51.100.21  14.126 ms  13.090 ms  29.082 ms
     MPLS Label=16 CoS=0 TTL=1 S=0
     MPLS Label=16 CoS=0 TTL=2 S=1
 3  198.51.100.22  16.419 ms  11.564 ms  17.068 ms
     MPLS Label=16 CoS=0 TTL=1 S=1
 4  198.51.100.1  12.794 ms  12.939 ms  17.123 ms
```

traceroute (Junos OS Evolved, Through an MPLS LSP)

The Junos OS Evolved `traceroute` command parses MPLS data in the same way as the Linux `traceroute` command: L=label, E=exp_use, S=stack_bottom, and T=TTL. In the example below, `T=1/L=16` indicates the TTL with label 16.

```
user@host> traceroute 192.0.2.50 ttl 255
traceroute to 192.0.2.50 (192.0.2.50), 255 hops max, 60 byte packets
 1  192.0.2.60 (192.0.2.60)  13.565 ms  11.696 ms  11.448 ms
 2  192.0.2.61 (192.0.2.61) <MPLS:L=17,E=0,S=0,T=1/L=16,E=0,S=1,T=1>  34.034 ms  31.538 ms
27.697 ms
 3  192.0.2.62 (192.0.2.62) <MPLS:L=299776,E=0,S=0,T=1/L=16,E=0,S=1,T=2>  23.174 ms  24.393 ms
21.009 ms
 4  192.0.2.63 (192.0.2.63)  24.553 ms  19.698 ms  25.648 ms
 5  192.0.2.50 (192.0.2.50)  33.322 ms  29.514 ms  24.706 ms
```

traceroute no-resolve extension (QFX5100)

```
user@host> traceroute 48.1.1.2 no-resolve extension
traceroute to 10.255.255.255 (10.255.255.255), 30 hops max, 52 byte packets
1  10.168.1.254  7.776 ms  8.108 ms  8.131 ms
    rx intf ae1.0(560) 10.168.1.254 mtu 1500
    rx sub-ip intf xe-0/0/18:0(803) 10.168.1.254 mtu 1500
    fwd intf ae2.0(562) 10.255.255.254 mtu 1500
    next-hop intf ae2.0(562) 10.255.255.255 mtu 1500
    !
2  10.255.255.255  7.122 ms  8.599 ms  8.267 ms
    rx intf ae2.0(556) 10.255.255.255 mtu 1500
    !
```

RELATED DOCUMENTATION

traceroute monitor | 29927

traceroute clns

IN THIS SECTION

- Syntax | 29917
- Description | 29917
- Options | 29917
- Required Privilege Level | 29918
- Output Fields | 29918
- Sample Output | 29918
- Release Information | 29919

Syntax

```
traceroute clns <host>
routing-instance <name>
source <source-address>
ttl <value>
wait <seconds>
```

Description

Trace the route belonging to Connectionless Network Service (CLNS).

Options

host	IP address or name of remote host.
routing-instance ***name***	Name of the routing instance for a traceroute attempt.
source ***source address***	Source address to be used in outgoing packets.
ttl ***value***	CLNP maximum time-to-live value. The range of values is 1 through 255.

wait *seconds* Number of seconds to wait for a response. The range of values is 1 second through 1 day.

Required Privilege Level

network

Output Fields

Table 1815 on page 29918 describes the output fields for the `traceroute clns` command. Output fields are listed in the approximate order in which they appear.

Table 1815: Traceroute clns Output Fields

Field Name	Field Description
source	Source address used in outgoing packets.
traceroute to	IP address of the receiver.
hops max	Maximum number of hops allowed.
byte packets	Size of packets being sent.

Sample Output

traceroute clns

```
user@host>traceroute clns <ISO address of the destination> source <ISO address of the source>
traceroute clns 49.0005.80ff.f800.0000.0108.0001.0102.5522.4145.00 source
47.0005.80ff.f800.0000.0108.0001.0102.5522.4143.00
clnstraceroute to 49.0005.80ff.f800.0000.0108.0001.0102.5522.4145.00  from
47.0005.80ff.f800.0000.0108.0001.0102.5522.4143.00, 30 hops max, 17 byte packets
 1  47.0005.80ff.f800.0000.0108.0001.0010.0010.0010.00  7.080 ms  5.579 ms  5.882 ms
 2  * * *
 3  * * *
```

```
 4  * * *
 5  49.0005.80ff.f800.0000.0108.0001.0102.5522.4145.00  11.325 ms  7.704 ms  5.261 ms
```

Release Information

Command introduced in Junos OS Release 8.0

traceroute ethernet

IN THIS SECTION

- Syntax | 29919
- Description | 29920
- Options | 29920
- Required Privilege Level | 29920
- Output Fields | 29920
- Sample Output | 29922
- Release Information | 29923

Syntax

```
traceroute ethernet
local-mep mep-id
maintenance-association ma-name
maintenance-domain md-name
<ttl value>
<wait seconds>
mac-address | mep-id
<detail>
```

Description

Triggers the linktrace protocol to trace the route between two maintenance points. The result of the traceroute protocol is stored in the path database. To display the path database, use the `show oam ethernet connectivity-fault-management path-database` command.

Before using the traceroute command, you can verify the remote MEP's MAC address using the `show oam ethernet connectivity-fault-management path-database` command.

Options

local-mep *mep-id*	(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.
detail	(Optional) Provide detailed information of the responder hostname, ingress port name, egress port name, TTL, and relay action.
mac-address	Destination unicast MAC address of the remote maintenance point.
mep-id	MEP identifier of the remote maintenance point. The range of values is 1 through 8191.
maintenance-association *ma-name*	Specifies an existing maintenance association from the set of configured maintenance associations.
maintenance-domain *md-name*	Specifies an existing maintenance domain from the set of configured maintenance domains.
ttl *value*	Number of hops to use in the linktrace request. The range is 1 to 255 hops. The default is 4.
wait *seconds*	(Optional) Maximum time to wait for a response to the traceroute request. The range is 1 to 255 seconds. The default is 5.

Required Privilege Level

network

Output Fields

Table 1816 on page 29921 lists the output fields for the `traceroute ethernet` command. Output fields are listed in the approximate order in which they appear.

Table 1816: traceroute ethernet Output Fields

Field Name	Field Description
`Linktrace to`	MAC address of the destination maintenance point.
`Interface`	Local interface used to send the linktrace message (LTM).
`Maintenance Domain`	Maintenance domain specified in the traceroute command.
`Level`	Maintenance domain level configured.
`Maintenance Association`	Maintenance association specified in the traceroute command.
`Local Mep`	The local maintenance end point identifier.
`Transaction Identifier`	4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all Maintenance Domains. Use the transaction identifier to match an incoming linktrace response (LTR), with a previously sent LTM.
`Hop`	Sequential hop count of the linktrace path.
`TTL`	Number of hops remaining in the linktrace message. The time to live (TTL) is decremented at each hop.
`Source MAC address`	MAC address of the 802.1ag node responding to the LTM or the source MAC address of the LTR.
`Next-hop MAC address`	MAC address of the egress interface of the node to which the LTM is forwarded or the next-hop MAC address derived from the next egress identifier in the Egress-ID TLV of the LTR PDU.

Table 1816: traceroute ethernet Output Fields *(Continued)*

Field Name	Field Description
`Responder Hostname`	The hostname of the responding router. A valid hostname is received only when the responding system is a Juniper Networks router.
`Ingress port name`	The port name for ingress connections.
`Egress port name`	The port name for egress connections.
`Flags`	The configurable flags can include: • H— Hardware only, incoming LT frame has hardware bit set. • T— Terminal MEP, responder is a terminating MEP. • F— FWD yes, LTM frame is relayed further.
`Relay Action`	The associated relay action. Relay action can be one of the following: • RlyHit— Relay hit; target MAC address matches the MP mac address. • RlyFDB— Relay FDB; output port decided by consulting forwarding database. • RlyMPDB— Relay MIP; output port decided by consulting MIP database.

Sample Output

traceroute ethernet

```
user@host> traceroute ethernet maintenance-domain md1 maintenance-association ma1
00:01:02:03:04:05
 Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
       Maintenance Domain: MD1, Level: 7
       Maintenance Association: MA1, Local Mep: 1

       Hop     TTL      Source MAC address        Next hop MAC address
       Transaction Identifier:100001
```

```
        1       63       00:00:aa:aa:aa:aa         00:00:ab:ab:ab:ab
        2       62       00:00:bb:bb:bb:bb         00:00:bc:bc:bc:bc
        3       61       00:00:cc:cc:cc:cc         00:00:cd:cd:cd:cd
        4       60       00:01:02:03:04:05         00:00:00:00:00:00
```

traceroute ethernet detail

```
user@host> run traceroute ethernet maintenance-domain md6 maintenance-association ma6 mep 101
detail
Linktrace to 00:00:5E:00:53:CC, Interface : ge-1/0/0.1
Maintenance Domain: md6, Level: 6
Maintenance Association: ma6, Local Mep: 201
Transaction Identifier: 2077547465

Legend for RelayAction:
RlyHit  -- Relay hit, Target MAC address matches the MP mac address
RlyFDB  -- Relay FDB, output port decided by consulting FDB database
RlyMPDB -- Relay MIP, output port decided by consulting MIP database

Legend for Flags:
H -- Hardware only,incoming LT frame has hardware bit set
T -- Terminal MEP, responder is a terminating MEP
F -- FWD yes, LTM frame is relayed further

TTL   Responder Hostname    Ingress port name     Egress port name       RelayAction
Responder Service     Ingress MAC address   Egress MAC address     Flags
62    host1                  ge-1/0/0.1            ge-2/3/0.1             RlyFDB
br1                   00:00:5E:00:53:00     00:00:5E:00:53:A0      HF-
63    host2                 ge-2/3/0.1            ge-1/0/0.1              RlyFDB
br1                   00:00:5E:00:53:AA     00:00:5E:00:53:A2      HF-
61    host3                 ge-1/0/0.1            --:--                   RlyHit
br1                   00:00:5E:00:53:B0           --:--            H-T
```

Release Information

Command introduced in Junos OS Release 9.0.

`mep-id` option introduced in Junos OS Release 9.1.

`local-mep` option introduced in Junos OS Release 15.1

traceroute fabric unicast-flow

IN THIS SECTION

- Syntax | 29924
- Description | 29924
- Options | 29925
- Required Privilege Level | 29925
- Output Fields | 29925
- Sample Output | 29926
- Sample Output | 29926
- Release Information | 29927

Syntax

```
traceroute fabric unicast-flow flow-spec-name flow-specification-name fma fma-name
<source-fmep fmep-name>
<dest-fmep fmep-name>
<forced-fte-interface fte-interface>
<count count>
<verbose>
```

Description

Trace the path taken by a specific unicast flow (ping) operation across a VLAN on the QFabric system from a source maintenance endpoint (FMEP) to a destination FMEP. The source and destination FMEPs may be on the same Node device, different Node devices connected to the same Interconnect device, or different Node devices connected to different Interconnect devices.

The flow path is the sequence of Packet Forwarding Engine forwarding hops along which the protocol data unit (PDU) travels. The hop-by-hop sequence and number of hops are reported in terms of fabric maintenance intermediate points (FMIPs), which are interfaces on the Packet Forwarding Engine of the Interconnect device. An FMIP sends a response to the source FMEP when the traceroute PDU is received.

Options

dest-fmep-id *source-fmep-id*	(Optional) ID of the destination FMEP of the traceroute operation.
flow-spec-name *flow-specification-name*	Name of the flow specification that defines the parameters of the traceroute operation.
fma-name *fma-name*	Name of the FMA that associates a VLAN with the FMEPs.
forced-fte-interface *fte-interface*	(Optional) Option to force the fabric ping operation to use the specified FTE interface to inject the PDU into the fabric instead of using the internal forwarding path lookup table to determine the FTE interface.
source-fmep-id *source-fmep-id*	(Optional) ID of the FMEP that is the source of the traceroute operation.
verbose	(Optional) Detailed version of the output display.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Output Fields

Table 1817 on page 29925 lists the output fields for the `traceroute fabric-unicast-flow` command. Output fields are listed in the approximate order in which they appear.

Table 1817: traceroute fabric unicast-flow Output Fields

Field Name	Field Description
`Fabric flow traceroute between source` *`source-fmep-name`* `destination` *`destination-fmep-name`*	Source and destination FMEP names.
`Using fabric flow-specification`	Name of the flow specification.
`Ethernet frame-size`	Ethernet frame size.

Table 1817: traceroute fabric unicast-flow Output Fields *(Continued)*

Field Name	Field Description
`Source-MAC`	MAC address of the source FMEP interface.
`Destination-MAC`	MAC address of the destination FMEP interface.
`Ethertype`	EtherType protocol.
`Received response from fmip-id, interface`	FMIP identifier and interface name of the FMIP that received the PDU and responded to the source FMEP.
`Received total` *`number`* `responses`	Number of responses that are received after the traceroute operation has finished.

Sample Output

traceroute fabric unicast-flow

```
user@host> traceroute fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2 flow-
spec-name fspec1
Received response from fmip-id 128.40.0.8, interface A0003:0/1/2
Received response from fmip-id 128.128.0.8, interface A0003:0/0/8
Received response from fmep-id 2
Received total 3 responses
```

Sample Output

traceroute fabric unicast-flow (Verbose)

```
user@host> traceroute fabric unicast-flow fma-name fma1 source-fmep-id 1 dest-fmep-id 2 flow-
spec-name fspec1 verbose
Fabric flow traceroute between source ED1494 destination ED1497
Using fabric flow-specification: fspec1
Ethernet frame-size: 256
```

```
Source-MAC: 0:EF:7:99:85:10
Destination-MAC: 0:D3:A6:11:39:21
Ethertype: 11564
Received response from fmip-id 128.40.0.8, interface A0003:0/1/2
Received response from fmip-id 128.128.0.8, interface A0003:0/0/8
Received response from fmep-id 2
Received total 3 responses
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute monitor

IN THIS SECTION

- Syntax | **29928**
- Description | **29928**
- Required Privilege Level | **29928**
- Options | **29928**
- Output Fields | **29929**
- Sample Output | **29929**
- Release Information | **29930**

Syntax

```
traceroute monitor host
<count value>
<inet | inet 6>
<interval seconds>
<no resolve>
<size value>
<source source-address>
<summary>
```

Description

Display live monitoring of each hop in the route that packets take to a specified network host. Use as a debugging tool to locate points of failure in a network.

> **NOTE**: This command utilizes the global routing table. It does not support individual routing-instances.

Required Privilege Level

network

Options

host	IP address or name of remote host.
count *value*	Number of ping requests, in packets, to send in summary mode. The default value is 10.
inet | inet 6	(Optional) Trace the route belonging to IPv4 or IPv6.
interval *seconds*	(Optional) Number of seconds to wait before sending ping requests. The default value is 1.
no resolve	(Optional) Do not attempt to display addresses symbolically.
size *value*	(Optional) Receive the specified number of bytes for each packet. The ragne is 0 - 65468 bytes. The default value is 64.

source ***source-address***	(Optional) Source address of the outgoing ping packets.
summary	(Optional) Generate and display a summary of live monitoring of each hop on the route that packets take to a specified network host.

Output Fields

Table 1 on page 29913 lists the output fields for the `traceroute monitor` command.

Table 1818: traceroute monitor Output Fields

Field Name	Field Description
Host	Hostname or IP address of the router at each hop.
Loss%	Percent of packet loss. The number of ping responses divided by the number of ping requests, specified as a percentage.
Snt	Number of ping requests sent to the router at this hop.
Last	Most recent round-trip time, in milliseconds, to the router at this hop.
Avg	Average round-trip time, in milliseconds, to the router at this hop.
Best	Shortest round-trip time, in milliseconds, to the router at this hop.
Wrst	Longest round-trip time, in milliseconds, to the router at this hop.
StDev	Standard deviation of round-trip times, in milliseconds, to the router at this hop.

Sample Output

traceroute monitor

```
user@host> traceroute monitor 192.168.1.10
                   Loss   Snt  Last  Avg   Best   Wrst  StDev
```

```
Host
1. 192.168.1.52    0.0%   17   0.7   1.0   0.6    5.4   1.2
2. example.net     0.0%   17   0.6   1.0   0.6    6.6   1.4
3. example2.net    0.0%   17   0.6   0.6   0.6    0.6   0.0
```

Release Information

Command introduced in Junos OS Release 8.0.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

RELATED DOCUMENTATION

traceroute | 29909

traceroute mpls bgp

IN THIS SECTION

- Syntax | 29930
- Description | 29931
- Options | 29931
- Required Privilege Level | 29932
- Output Fields | 29933
- Sample Output | 29934
- Release Information | 29935

Syntax

```
traceroute mpls bgp fec
<destination destination-address>
<detail>
```

```
<exp exp>
<fanout fanout-number>
<logical-system logical system-name>
<no-resolve>
<paths paths-number>
<pipe-mode>
<retries retries-number>
<routing-instance routing-instance-name>
<source source-address>
<ttl value>
<wait seconds>
```

Description

Trace route to a remote host for an MPLS label-switched path (LSP) signaled by the Border Gateway Protocol (BGP). Use `traceroute mpls bgp` as a debugging tool to locate MPLS BGP forwarding issues in a network. (Currently supported for IPv4 packets only.)

To use the `traceroute mpls bgp` command, make sure you have BGP routes with MPLS labels.

NOTE: The `traceroute mpls bgp` *fec* command only supports single paths.

Options

fec	Specify the IP address and optional prefix of the forwarding equivalence class (FEC). Suppose you are at PE1, use would want to use the IP address of PE2 to trace the BGP path to that router.
destination *destination-address*	(Optional) Specify the destination address to use when sending probes.
detail	(Optional) Display detailed output.
exp *exp*	(Optional) Specify the class of service to use when sending probes. • **Range:** 0 through 7 • **Default:** 7
fanout *fanout-number*	(Optional) Specify the maximum number of next hops to search per node. • **Range:** 1 through 16

- **Default:** 16

`logical-system` *`logical-system-name`*	(Optional) Specify the name of the logical system for the traceroute attempt.
`no-resolve`	(Optional) Specify not to resolve the hostname that corresponds to the IP address.
`paths` *`paths-number`*	(Optional) Specify the number of paths to search. • **Range:** 1 through 255 • **Default:** 16
`pipe-mode`	(Optional) Specify to trace only the outermost FEC.
`retries` *`retries-number`*	(Optional) Specify the number of times to resend probe values. • **Range:** 1 through 9 • **Default:** 3
`routing-instance` *`routing-instance-name`*	(Optional) Specify the name of the routing instance for the trace route attempt.
`source` *`source-address`*	(Optional) Specify the source address of the outgoing traceroute packets.
`ttl` *`value`*	(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds. • **Range:** 1 through 125 • **Default:** 64
`wait` *seconds*	(Optional) Specify the number of seconds to wait before resending a probe. • **Range:** 5 though 15 • **Default:** 10

Required Privilege Level

network

Output Fields

Table 1819 on page 29933 describes the output fields for the `traceroute mpls bgp` *fec* command and the `traceroute mpls bgp` *fec* `detail` command. Output fields are listed in the approximate order in which they appear.

Table 1819: traceroute mpls bgp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the `traceroute mpls bgp` *fec* command.	All levels
ttl	Time to live value of the labeled packet.	None
Label	Outgoing label used for forwarding the packet along the label-switched paths.	None
Protocol	Signaling protocol used. For this command, it is BGP.	None
Address	Address of the next hop.	None
Previous Hop	Address of the previous hop. Previous hop address of the first hop is `null`.	None
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	None
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	`detail`
Parent	Address of the previous hop. Parent value for the first hop is `null`.	`detail`
Return Code	Return code for reporting the result of processing the echo request by the receiver.	`detail`
Response time	Time for the echo request to reach the receiver.	`detail`

Table 1819: traceroute mpls bgp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is `none`.	`detail`
Label Stack	Label stack used to forward the packet.	`detail`

Sample Output

traceroute mpls bgp

```
user@host> traceroute mpls bgp fec
    Probe options: retries 3, exp 7
ttl Label  Protocol Address  Previous Hop  Probe Status Fec-Stack-Sent Fec-Change
1   299824  LDP     81.1.2.2   (null)      Success     LDP, BGP        PUSH-RSVP
2   299825  RSVP    81.2.3.3   81.1.2.2    Success     RSVP, LDP, BGP  (null)
3   299826  RSVP    81.3.4.4   81.2.3.3    Egress      RSVP, LDP, BGP  POP-RSVP
3   299826  LDP     81.3.4.4   81.2.3.3    Success     LDP, BGP        (null)
4   299827  LDP     81.4.5.5   81.3.4.4    Egress      LDP, BGP        POP-LDP
4   299827  BGP     81.4.5.5   81.3.4.4    Egress      BGP             (null)
```

traceroute mpls bgp detail

```
user@host> traceroute mpls bgp fec detail
  Probe options: retries 3, exp 7
  Hop 2.2.1.81.rev.sfr.net (81.1.2.2) Depth 1
    Probe status: Success
    Parent: (null)
    Return code: Label switched at stack-depth 1
    Sender timestamp: 2013-03-22 05:55:19 PDT 822.99 msec
    Receiver timestamp: 2013-03-22 05:55:19 PDT 856.05 msec
    Response time: 33.06 msec
    MTU: Unknown
    Multipath type: IP bitmask
      Address Range 1: 127.0.0.64 ~ 127.0.0.127
    Label Stack:
```

```
      Label 1 Value 299824 Protocol LDP
      Label 2 Value 299276 Protocol BGP
   Fec-Stack-Sent: LDP, BGP
   Fec-Change:
     Operation: PUSH   Protocol RSVP
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

ping mpls bgp | 12547

traceroute mpls ldp

IN THIS SECTION

- Syntax | 29935
- Description | 29936
- Options | 29936
- Required Privilege Level | 29938
- Output Fields | 29938
- Sample Output | 29939
- Release Information | 29940

Syntax

```
traceroute mpls <ldp> fec
<destination ip-address>
<detail>
<exp exp>
<fanout fanout-number>
```

```
<logical-system logical-system-name>
<no-resolve>
<paths maximum-paths>
<pipe-mode>
<retries retries-number>
<routing-instance routing-instance-name>
<source ip-address>
<ttl value>
<update>
<wait seconds>
```

Description

Trace route to a remote host for an MPLS label-switched path signaled by the LDP. Use `traceroute mpls ldp` as a debugging tool to locate MPLS label-switched path forwarding issues in a network. (Currently supported for IPv4 packets only.)

Options

fec
Specify the IP address and optional prefix of the forwarding equivalence class (FEC).

destination *ip-address*
(Optional) Specify the destination address to use when sending probes.

- **Values:** The destination IP address must be within the 127.0.0.0/8 IP address space for Operation, Administration, and Maintenance (OAM) packets.

detail
(Optional) Display detailed output.

exp *exp*
(Optional) Specify the class-of-service to use when sending probes.

- **Range:** 0 through 7
- **Default:** 7

fanout *fanout-number*
(Optional) Specify the maximum number of nexthops to search per node.

- **Range:** 1 through 16
- **Default:** 16

logical-system
(Optional) Specify the name of the logical system for the traceroute attempt.

no-resolve
(Optional) Specify not to resolve the hostname that corresponds to the IP address.

paths *maximum-paths*

(Optional) Specify the maximum number of paths to search.

- **Range:** 1 through 255
- **Default:** 16

pipe-mode

(Optional) Specify to trace only nodes that understand LDP FEC.

In an interoperation with other vendor devices or devices running Junos OS Release that do not support tracing of hierarchical LSPs as described in RFC 6424, continuous non-complaint probe status is displayed in the `traceroute mpls ldp` command output. To avoid this LDP loop creation, use the `pipe-mode` option with the `traceroute mpls ldp` *fec* command.

> **NOTE**: Even after using the `traceroute mpls ldp` *fec* `pipe-mode` command, one or more intermediate transit nodes that do not understand LDP FEC can return non-complaint probe status in the command output.

retries *retries-number*

(Optional) Specify the number of times to resend probe values.

- **Range:** 1 through 9
- **Default:** 3

routing-instance *routing-instance-name*

(Optional) Specify the name of the routing instance for the traceroute attempt.

source *source-address*

(Optional) Specify the source address of the outgoing traceroute packets.

ttl *value*

(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds.

- **Range:** 1 through 125
- **Default:** 64

update

(Optional) Update database contents with traceroute results.

wait *seconds*

(Optional) Specify the number of seconds to wait before resending a probe.

- **Range:** 5 through 15
- **Default:** 10

Required Privilege Level

network

Output Fields

Table 1820 on page 29938 describes the output fields for the `traceroute mpls ldp` *fec* command and the `traceroute mpls ldp` *fec* `detail` commands. Output fields are listed in the approximate order in which they appear.

Table 1820: traceroute mpls ldp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the `traceroute mpls ldp` *fec* command.	all levels
ttl	Time to live value of the labeled packet.	none specified
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified
Protocol	Signaling protocol used. For this command, it is LDP.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is **null**.	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	**detail**
Parent	Address of the previous hop. Parent value for the first hop is **null**.	**detail**
Return Code	Return code for reporting the result of processing the echo request by the receiver.	**detail**

Table 1820: traceroute mpls ldp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Response time	Time for the echo request to reach the receiver.	**detail**
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is **none**.	**detail**
Label Stack	Label stack used to forward the packet.	**detail**

Sample Output

traceroute mpls ldp

```
user@router> traceroute mpls ldp 4.4.4.4

  Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   100016  LDP         24.24.24.1       (null)           Success
    2   100000  LDP         20.20.20.2       24.24.24.1       Success
    3        3  LDP         22.22.22.4       20.20.20.2       Egress

Path 1 via fe-0/3/3.101 destination 127.0.0.64
```

traceroute mpls ldp detail

```
user@router> traceroute mpls ldp 4.4.4.4 detail

  Probe Options: ttl 64, retries 3, wait 10, paths 3, exp 7
  Hop 24.24.24.1 Depth 1
    Parent (null)
    Return code: Label switched at stack-depth 1
    Response time 165.93 msec
    Multipath type: IP bitmask
      Address Range 1: 127.0.0.0 ~ 127.0.3.255
    Label Stack:
```

```
      Label 1 Value 100032 Protocol LDP

  Hop 20.20.20.2 Depth 2
    Parent 24.24.24.1
    Return code: Upstream interface index unknown label-switched at stack-depth 1
    Response time 19.05 msec
    Multipath type: IP bitmask
      Address Range 1: 127.0.0.0 ~ 127.0.3.255
    Label Stack:
      Label 1 Value 100000 Protocol LDP

  Hop 22.22.22.4 Depth 3
    Parent 20.20.20.2
    Return code: Egress-ok at stack-depth 1
    Response time 0.79 msec
    Multipath type: None
    Label Stack:
      Label 1 Value 3 Protocol LDP
```

Release Information

Command introduced in Junos OS Release 8.4.

traceroute mpls rsvp

IN THIS SECTION

- Syntax | 29941
- Description | 29941
- Options | 29941
- Required Privilege Level | 29942
- Output Fields | 29942
- Sample Output | 29944
- Release Information | 29946

Syntax

```
traceroute mpls <rsvp> lsp-name
<detail>
<egress>
<exp>
<logical-system>
<multipoint>
<no-resolve>
<retries>
<source source-address>
<ttl>
```

Description

Trace route to a remote host for an MPLS LSP signaled by RSVP. Use `traceroute mpls rsvp` as a debugging tool to locate MPLS label-switched path (LSP) forwarding issues in a network. (Currently supported for IPv4 packets only.)

Options

lsp-name	Specify the name of the LSP to be traced.
detail	(Optional) Display detailed output.
egress	(Optional) Request that a specific point-to-multipoint egress node reply to the trace route. The trace route would follow the associated sub-LSP to the egress node.
exp	(Optional) Specify the class of service to use when sending probes. The range of values is 0 through 7. The default value is 7.
logical-system	(Optional) Specify the name of the logical system for the traceroute attempt.
multipoint	(Optional) Perform a trace route on a point-to-multipoint LSP.
no-resolve	(Optional) Specify not to resolve the hostname that corresponds to the IP address.
retries	(Optional) Specify the number of times to resend probe. The range of values is 1 through 9. The default value is 3.

source ***source-address***	(Optional) Specify the source address of the outgoing traceroute packets.
ttl	(Optional) Specify the number of hops to follow before forcing the trace route to quit.

Required Privilege Level

network

Output Fields

Table 1821 on page 29942 describes the output fields for the `traceroute mpls rsvp` *`lsp-name`* and `traceroute mpls rsvp` *`lsp-name`* `detail` commands. Output fields are listed in the approximate order in which they appear.

Table 1821: traceroute mpls rsvp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the `traceroute mpls rsvp` *`lsp-name`* command.	all levels
ttl	Time-to-live value of the labeled packet.	none specified
Label	MPLS label used to forward the packets along the LSP.	none specified
Protocol	Signaling protocol used. For this command, it is RSVP-TE.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null.	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths). Displays `Success` if the trace to a hop is successful or `Egress` if the trace has reached the last router on the path.	none specified

Table 1821: traceroute mpls rsvp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	`detail`
Parent	Address of the previous hop. Parent value for the first hop is null.	`detail`
Return Code	Return code for reporting the result of processing the echo request by the receiver.	`detail`
Sender timestamp	Displays the timestamp when the MPLS echo request is sent to the next hop.	`detail`
Receiver timestamp	Timestamp when the echo request from the previous hop is received and acknowledged with an echo response by the next hop.	`detail`
Response time	Time for the echo request to reach the receiver.	`detail`
MTU	Size of the largest packet that includes the label stack forwarded to the next hop.	`detail`
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none.	`detail`
Label stack	Label stack used to forward the packet.	`detail`
Path	Displays the sub-lsp path number for this traceroute, the interface used, and the destination address.	all levels

Sample Output

traceroute mpls rsvp

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta

     Probe options: retries 3, exp 7

     ttl    Label  Protocol    Address          Previous Hop     Probe Status
       1   299792  RSVP-TE        192.168.1.2      (null)           Success
       2   299803  RSVP-TE        192.168.2.3      192.168.1.2      Success
       3        3  RSVP-TE        192.168.3.4      192.168.2.3      Egress

     Path 1 via ge-0/0/0.1 destination 127.0.0.64
```

traceroute mpls rsvp detail

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta detail
Probe options: retries 3, exp 7

   Hop 192.168.1.2 Depth 1
     Probe status: Success
     Parent: (null)
     Return code: Label-switched at stack-depth 1
     Sender timestamp: 2008-04-17 09:35:27 EDT 400.88 msec
     Receiver timestamp: 2008-04-17 09:35:27 EDT 427.87 msec
     Response time: 26.99 msec
     MTU: Unknown
     Multipath type: IP bitmask
       Address Range 1: 127.0.0.64 ~ 127.0.0.127
     Label Stack:
       Label 1 Value 299792 Protocol RSVP-TE

Hop 192.168.2.3 Depth 2
     Probe status: Success
     Parent: 192.168.1.2
     Return code: Upstream interface index unknown label-switched at stack-depth 1
     Sender timestamp: 2008-04-17 09:35:27 EDT 522.13 msec
     Receiver timestamp: 2008-04-17 09:35:27 EDT 548.69 msec
     Response time: 26.55 msec
```

```
    MTU: 1518
    Multipath type: IP bitmask
      Address Range 1: 127.0.0.64 ~ 127.0.0.127
    Label Stack:
      Label 1 Value 299803 Protocol RSVP-TE
```

traceroute mpls rsvp multipoint (branch node for sub-LSPs)

The following traceroute output is for a point-to-multipoint LSP where the penultimate node is a branch node for the sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
  Probe options: retries 3, exp 7

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   300000  RSVP-TE     81.1.2.2         (null)           Success
    2   299968  RSVP-TE     81.2.3.3         81.1.2.2         Success
    3   299952  RSVP-TE     81.3.4.4         81.2.3.3         Success
    4   299920  RSVP-TE     81.4.6.6         81.3.4.4         Egress

  Path 1 via lt-1/2/0.102 destination 127.0.0.64

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    4   299920  RSVP-TE     81.4.5.5         81.3.4.4         Egress

  Path 2 via lt-1/2/0.102 destination 127.0.0.64
```

traceroute mpls rsvp multipoint (single-hop sub-LSPs)

The following traceroute output is for a point-to-multipoint LSP with multiple single-hop sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
  Probe options: retries 3, exp 7

   ttl    Label  Protocol    Address          Previous Hop     Probe Status
     1        0  RSVP-TE     81.1.2.2         (null)           Egress

   Path 1 via lt-1/2/0.102 destination 127.0.0.64

   ttl    Label  Protocol    Address          Previous Hop     Probe Status
```

```
  1        0  RSVP-TE     81.1.8.8         (null)           Egress

Path 2 via lt-1/2/0.108 destination 127.0.0.64

ttl    Label  Protocol    Address          Previous Hop     Probe Status
  1        0  RSVP-TE     81.1.9.9         (null)           Egress

Path 3 via lt-1/2/0.109 destination 127.0.0.64
```

Release Information

Command introduced in Junos OS Release 9.2.

`egress`, `multipoint`, and `ttl` options added in Junos OS Release 11.2.

traceroute mpls segment-routing isis

IN THIS SECTION

- Syntax | 29946
- Description | 29947
- Options | 29947
- Required Privilege Level | 29948
- Output Fields | 29949
- Sample Output | 29950
- Release Information | 29953

Syntax

```
traceroute mpls segment-routing isis fec
algorithm algorithm-id
<destination ip-address>
<detail>
<exp exp>
```

```
<fanout fanout-number>
<logical-system logical-system-name>
<no-resolve>
<paths maximum-paths>
<pipe-mode>
<retries retries-number>
<routing-instance routing-instance-name>
<source ip-address>
<ttl value>
<update>
<wait seconds>
```

Description

Trace route to a remote host for a segment routing label-switched path added by the ISIS protocol. Use `traceroute mpls segment-routing isis` as a debugging tool to locate MPLS label-switched path forwarding issues in a network.

Options

algorithm ***algorithm-id*** (Optional) Specify the flexible algorithm identifier. The range of values is `0` through `255`. The default value is `0` which means no algorithm is configured. The *algorithm-id* in conjunction with the prefix is used to distinguish labelled SR paths.

fec Specify the IP address and optional prefix of the forwarding equivalence class (FEC).

destination ***ip-address*** (Optional) Specify the destination address to use when sending probes.

- **Values:** The destination IP address must be within the 127.0.0.0/8 IP address space for Operation, Administration, and Maintenance (OAM) packets.

detail (Optional) Display detailed output.

exp ***exp*** (Optional) Specify the class-of-service to use when sending probes.

- **Range:** 0 through 7
- **Default:** 7

fanout ***fanout-number*** (Optional) Specify the maximum number of nexthops to search per node.

- **Range:** 1 through 16
- **Default:** 16

logical-system (Optional) Specify the name of the logical system for the traceroute attempt.

no-resolve (Optional) Specify not to resolve the hostname that corresponds to the IP address.

paths *maximum-paths* (Optional) Specify the maximum number of paths to search.

- **Range:** 1 through 255
- **Default:** 16

retries *retries-number* (Optional) Specify the number of times to resend probe values.

- **Range:** 1 through 9
- **Default:** 3

routing-instance *routing-instance-name* (Optional) Specify the name of the routing instance for the traceroute attempt.

source *source-address* (Optional) Specify the source address of the outgoing traceroute packets.

ttl *value* (Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds.

- **Range:** 1 through 125
- **Default:** 64

wait *seconds* (Optional) Specify the number of seconds to wait before resending a probe.

- **Range:** 5 through 15
- **Default:** 10

Required Privilege Level

network

Output Fields

Table 1822 on page 29949 describes the output fields for the `traceroute mpls segment-routing isis` *`fec`* command and the `traceroute mpls segment-routing isis` *`fec`* `detail` commands. Output fields are listed in the approximate order in which they appear.

Table 1822: traceroute mpls ldp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the `traceroute mpls ldp` *`fec`* command.	all levels
ttl	Time to live value of the labeled packet.	none specified
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified
Protocol	Signaling protocol used. For this command, it is LDP.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is **null**.	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	**detail**
Parent	Address of the previous hop. Parent value for the first hop is **null**.	**detail**
Return Code	Return code for reporting the result of processing the echo request by the receiver.	**detail**
Response time	Time for the echo request to reach the receiver.	**detail**

Table 1822: traceroute mpls ldp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is **none**.	**detail**
Label Stack	Label stack used to forward the packet.	**detail**

Sample Output

traceroute mpls segment-routing isis

```
user@router> traceroute mpls segment-routing isis 10.4.4.4

Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   402006  ISIS        12.1.1.2         (null)           Success
  FEC-Stack-Sent: ISIS
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2   402006  ISIS        23.1.1.2         12.1.1.2         Success
  FEC-Stack-Sent: ISIS
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    3   402006  ISIS        34.1.1.2         23.1.1.2         Success
  FEC-Stack-Sent: ISIS
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    4   402006  ISIS        45.1.1.2         34.1.1.2         Success
  FEC-Stack-Sent: ISIS
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    5        3  ISIS        56.1.1.2         45.1.1.2         Egress
  FEC-Stack-Sent: ISIS

  Path 1 via ge-0/0/0.0 destination 127.0.0.64

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    3   402006  ISIS        34.2.1.2         23.1.1.2         Success
  FEC-Stack-Sent: ISIS
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
```

```
  4   402006  ISIS        45.1.1.2         34.2.1.2         Success
FEC-Stack-Sent: ISIS
ttl    Label  Protocol    Address          Previous Hop     Probe Status
  5        3  ISIS        56.1.1.2         45.1.1.2         Egress
FEC-Stack-Sent: ISIS

Path 2 via ge-0/0/0.0 destination 127.0.0.65

ttl    Label  Protocol    Address          Previous Hop     Probe Status
  3   402006  ISIS        34.5.1.2         23.1.1.2         Success
FEC-Stack-Sent: ISIS
ttl    Label  Protocol    Address          Previous Hop     Probe Status
  4   402006  ISIS        45.2.1.2         34.5.1.2         Success
FEC-Stack-Sent: ISIS
ttl    Label  Protocol    Address          Previous Hop     Probe Status
  5        3  ISIS        56.1.1.2         45.2.1.2         Egress
FEC-Stack-Sent: ISIS

Path 3 via ge-0/0/0.0 destination 127.0.0.69
```

traceroute mpls segment-routing isis (Labelled IS-IS flexible-algorithm IPv4 traceroute)

```
user@router> traceroute mpls segment-routing isis 10.4.4.4 algorithm 130

  Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   401306  ISIS        21.0.3.2         (null)           Success
  FEC-Stack-Sent: ISIS
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2        3  ISIS        20.34.67.6       21.0.3.2         Egress
  FEC-Stack-Sent: ISIS

  Path 1 via ge-0/0/3.0 destination 127.0.0.64
```

traceroute mpls segment-routing isis (Labelled IS-IS flexible-algorithm IPv6 traceroute)

```
user@router> traceroute mpls segment-routing isis 10:4:4::4 algorithm 130

Probe options: retries 3, exp 7

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   402406  ISIS        d:0:3::3         (null)           Success
  FEC-Stack-Sent: ISIS IPv6
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2        3  ISIS        d:34:67::6       d:0:3::3         Egress
  FEC-Stack-Sent: ISIS IPv6

  Path 1 via ge-0/0/3.0 destination ::ffff:127.0.0.64

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   402406  ISIS        d:0:4::4         (null)           Success
  FEC-Stack-Sent: ISIS IPv6
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2        3  ISIS        d:4:6::6         d:0:4::4         Egress
  FEC-Stack-Sent: ISIS IPv6

  Path 2 via ge-0/0/6.0 destination ::ffff:127.0.1.64
```

traceroute mpls segment-routing isis (Labelled IS-IS flexible-algorithm IPv4 traceroute, Ingress ECMP)

```
user@host> traceroute mpls segment-routing isis 10.4.4.4 algorithm 130
  Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   401406  ISIS        21.0.3.2         (null)           Success
  FEC-Stack-Sent: ISIS
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2        3  ISIS        20.34.67.6       21.0.3.2         Egress
  FEC-Stack-Sent: ISIS

  Path 1 via ge-0/0/3.0 destination 127.0.0.64
```

```
ttl    Label  Protocol    Address          Previous Hop     Probe Status
  1   401406  ISIS        21.0.4.2         (null)           Success
FEC-Stack-Sent: ISIS
ttl    Label  Protocol    Address          Previous Hop     Probe Status
  2        3  ISIS        21.4.6.2         21.0.4.2         Egress
FEC-Stack-Sent: ISIS

Path 2 via ge-0/0/6.0 destination 127.0.1.64
```

Release Information

Command introduced in Junos OS Release 19.1R1.

`algorithm` option introduced in Junos OS Release 23.1R1 and Junos OS Evolved Release 23.1R1.

traceroute mpls segment-routing ospf

IN THIS SECTION

- Syntax | 29953
- Description | 29954
- Options | 29954
- Required Privilege Level | 29955
- Output Fields | 29956
- Sample Output | 29957
- Release Information | 29958

Syntax

```
traceroute mpls segment-routing ospf fec
algorithm algorithm-id
<destination ip-address>
<detail>
```

```
<exp exp>
<fanout fanout-number>
<logical-system logical-system-name>
<no-resolve>
<paths maximum-paths>
<pipe-mode>
<retries retries-number>
<routing-instance routing-instance-name>
<source ip-address>
<ttl value>
<update>
<wait seconds>
```

Description

Trace route to a remote host for a segment routing label-switched path added by the ISIS protocol. Use `traceroute mpls segment-routing isisp` as a debugging tool to locate MPLS label-switched path forwarding issues in a network. (Currently supported for IPv4 packets only.)

Options

algorithm ***algorithm-id*** (Optional) Specify the flexible algorithm identifier. The range of values is `0` through `255`. The default value is `0` which means no algorithm is configured. The *algorithm-id* in conjunction with the prefix is used to distinguish labelled SR paths.

fec Specify the IP address and optional prefix of the forwarding equivalence class (FEC).

destination ***ip-address*** (Optional) Specify the destination address to use when sending probes.

- **Values:** The destination IP address must be within the 127.0.0.0/8 IP address space for Operation, Administration, and Maintenance (OAM) packets.

detail (Optional) Display detailed output.

exp ***exp*** (Optional) Specify the class-of-service to use when sending probes.

- **Range:** 0 through 7
- **Default:** 7

fanout ***fanout-number*** (Optional) Specify the maximum number of nexthops to search per node.

- **Range:** 1 through 16
- **Default:** 16

logical-system (Optional) Specify the name of the logical system for the traceroute attempt.

no-resolve (Optional) Specify not to resolve the hostname that corresponds to the IP address.

paths *maximum-paths* (Optional) Specify the maximum number of paths to search.

- **Range:** 1 through 255
- **Default:** 16

retries *retries-number* (Optional) Specify the number of times to resend probe values.

- **Range:** 1 through 9
- **Default:** 3

routing-instance *routing-instance-name* (Optional) Specify the name of the routing instance for the traceroute attempt.

source *source-address* (Optional) Specify the source address of the outgoing traceroute packets.

ttl *value* (Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds.

- **Range:** 1 through 125
- **Default:** 64

wait *seconds* (Optional) Specify the number of seconds to wait before resending a probe.

- **Range:** 5 through 15
- **Default:** 10

Required Privilege Level

network

Output Fields

Table 1823 on page 29956 describes the output fields for the `traceroute mpls segment-routing isis` *`fec`* command and the `traceroute mpls segment-routing isis` *`fec`* `detail` commands. Output fields are listed in the approximate order in which they appear.

Table 1823: traceroute mpls ldp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the `traceroute mpls ldp` *`fec`* command.	all levels
ttl	Time to live value of the labeled packet.	none specified
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified
Protocol	Signaling protocol used. For this command, it is LDP.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is **null**.	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	**detail**
Parent	Address of the previous hop. Parent value for the first hop is **null**.	**detail**
Return Code	Return code for reporting the result of processing the echo request by the receiver.	**detail**
Response time	Time for the echo request to reach the receiver.	**detail**

Table 1823: traceroute mpls ldp Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is **none**.	**detail**
Label Stack	Label stack used to forward the packet.	**detail**

Sample Output

traceroute mpls segment-routing ospf

```
user@router> traceroute mpls ospf 4.4.4.4

Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   402006  OSPF        12.1.1.2         (null)           Success
  FEC-Stack-Sent: OSP
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2   402006  OSPF        23.1.1.2         12.1.1.2         Success
  FEC-Stack-Sent: OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    3   402006  OSPF        34.1.1.2         23.1.1.2         Success
  FEC-Stack-Sent: OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    4   402006  OSPF        45.1.1.2         34.1.1.2         Success
  FEC-Stack-Sent: OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    5        3  OSPF        56.1.1.2         45.1.1.2         Egress
  FEC-Stack-Sent: OSPF

  Path 1 via ge-0/0/0.0 destination 127.0.0.64

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    3   402006  OSPF        34.2.1.2         23.1.1.2         Success
  FEC-Stack-Sent: OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
```

```
  4   402006  OSPF        45.1.1.2         34.2.1.2         Success
FEC-Stack-Sent: OSPF
ttl    Label  Protocol    Address          Previous Hop     Probe Status
  5        3  OSPF        56.1.1.2         45.1.1.2         Egress
FEC-Stack-Sent: OSPF

Path 2 via ge-0/0/0.0 destination 127.0.0.65

ttl    Label  Protocol    Address          Previous Hop     Probe Status
  3   402006  OSPF        34.5.1.2         23.1.1.2         Success
FEC-Stack-Sent: OSPF
ttl    Label  Protocol    Address          Previous Hop     Probe Status
  4   402006  OSPF        45.2.1.2         34.5.1.2         Success
FEC-Stack-Sent: OSPF
ttl    Label  Protocol    Address          Previous Hop     Probe Status
  5        3  OSPF        56.1.1.2         45.2.1.2         Egress
FEC-Stack-Sent: OSPF

Path 3 via ge-0/0/0.0 destination 127.0.0.69
```

traceroute mpls segment-routing ospf (Labelled OSPF flexible-algorithm IPv4 traceroute)

```
user@router> traceroute mpls segment-routing ospf 10.4.4.4 algorithm 140

  Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   401306  OSPF        21.0.3.2         (null)           Success
  FEC-Stack-Sent: OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2        3  OSPF        20.34.67.6       21.0.3.2         Egress
  FEC-Stack-Sent: OSPF

  Path 1 via ge-0/0/3.0 destination 127.0.0.64
```

Release Information

Command introduced in Junos OS Release 19.1R1.

algorithm option introduced in Junos OS Release 23.1R1 and Junos OS Evolved Release 23.1R1.

traceroute mpls segment-routing spring-te

IN THIS SECTION

- Syntax | 29959
- Description | 29959
- Options | 29960
- Additional Information | 29962
- Required Privilege Level | 29963
- Output Fields | 29963
- Sample Output | 29965
- Release Information | 29969

Syntax

```
traceroute mpls segment-routing spring-te
<egress-ip ( active | color |  detail | exp | logical-system | no-resolve | retries |  routing-
instance | secondary | segment-list | skip-fec-validation  | source | ttl | tunnel-source) >
<label-stack ( detail | egress | labels | logical-system | nexthop-address | nexthop-interface |
no-resolve) >
<source-routing-path ( lsp-name | active | color | detail | egress-ip | exp | logical-system |
no-resolve | retries | routing-instance | secondary | segment-list | skip-fec-validation |
source | tunnel-source) >
```

Description

Trace route to a remote host for a segment routing label-switched path added by segment routing traffic-engineered (SPRING-TE) tunnel. Use traceroute mpls segment-routing spring-te as a debugging tool to locate MPLS label-switched path forwarding issues in a network.

Options

egress-ip Specify to trace to or install IP address to use when sending probes.

active	Specify to use the forwarding path or nexthops from the RIB table.
color	Specify the color identifier for the tunnel end-point. • **Range:** 1 through 4294967295
detail	(Optional) Display detailed output.
exp *exp*	(Optional) Specify the class of service to use when sending probes. • **Range:** 0 through 7
logical-system *logical-system-name*	(Optional) Specify the name of the logical system for the traceroute attempt.
no-resolve	Specify not to attempt to print addresses symbolically.
retries *retries-number*	(Optional) Specify the number of times to resend probe values. • **Range:** 1 through 9
routing-instance *routing-instance-name*	(Optional) Specify the name of the routing instance for the trace route attempt.
secondary	Specify to use configured secondary segment list for the given segment routing path.
segment-list	Specify the segment list to use when sending probes.
skip-fec-validation	Specify to skip forwarding equivalence class (FEC) validation or use NIL FEC.
source *source-address*	(Optional) Specify the source address of the outgoing traceroute packets.
ttl *value*	(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds. • **Range:** 1 through 255
tunnel-source	Specify the source protocol used to create tunnel.

label-stack

Specify the label stack for traceroute packets. This option works only if you do a minimum configuration of `source-packet-routing` at the `[edit protocols]` hierarchy level.

detail	(Optional) Display detailed output.
egress	Specify the egress IP address.
labels	Specify the labels in label stack. • **Range:** 16 through 1048575
logical-system *logical-system-name*	(Optional) Specify the name of the logical system for the traceroute attempt.
nexthop-address *address*	Specify the nexthop IP address for the traceroute packet.
`nexthop-interface` *`interface`*	Specify the outgoing interface for the traceroute packet.
no-resolve	Specify not to attempt to print addresses symbolically.

source-routing-path

Specify the trace source path routing to use when sending probes.

lsp-name	Specify the source path routing name.
active	Specify to use the forwarding path or nexthops from the RIB table.
color	Specify the color identifier for the tunnel end-point. • **Range:** 1 through 4294967295
detail	(Optional) Display detailed output.
egress-ip	Specify to trace to or install IP address to use when sending probes.
exp *exp*	(Optional) Specify the class of service to use when sending probes. • **Range:** 0 through 7
logical-system *logical-system-name*	(Optional) Specify the name of the logical system for the traceroute attempt.
no-resolve	Specify not to attempt to print addresses symbolically.
retries *retries-number*	(Optional) Specify the number of times to resend probe values.

- **Range:** 1 through 9

secondary Specify to use configured secondary segment list for the given segment routing path.

segment-list Specify the segment list to use when sending probes.

skip-fec-validation Specify to skip forwarding equivalence class (FEC) validation or use NIL FEC.

source ***source-address*** (Optional) Specify the source address of the outgoing traceroute packets.

ttl ***value*** (Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds.

- **Range:** 1 through 255

Additional Information

NOTE: With `segment-list` option, valid `tunnel-source` (source protocol used to create tunnel) is only static.

NOTE: `segment-list` option is not present when `tunnel-source` is BGP-SR-TE since BGP-SR-TE does not have named segment list.

NOTE: With `source-routing-path` name, valid `tunnel-source` are PCEP and static.

NOTE: With `tunnel-source` as PCEP, `secondary` option is not valid as PCEP always sends one primary LSP.

NOTE: FEC validation is performed for `traceroute mpls segment-routing spring-te` by default.

NOTE: FEC validation is supported when segment-routing traffic-engineering (SR-TE) tunnel has IGP labels (OSPF and IS-IS) only. For all the other labels (for example, static, ldp, rsvp, etc.), you can use `skip-fec-validation` option.

NOTE: For a combination of `egress-ip` and `color` options, valid `tunnel-source` are static and BGP-SR-TE.

NOTE: `traceroute mpls segment-routing spring-te` command with `label-stack` option is supported in operational mode only. It is not supported in configuration mode.

NOTE: Parallel SID, binding SID, LDP over SR-TE is not supported.

NOTE: In case of `traceroute mpls segment-routing spring-te`, the PFE supports only 16 ECMP paths. If there more than 16 ECMP paths to the SR-TE LSP, traceroute is done up to a maximum of 16 hops.

Required Privilege Level

network

Output Fields

Table 1824 on page 29963 describes the output fields for the `traceroute mpls segment-routing spring-te` command. Output fields are listed in the approximate order in which they appear.

Table 1824: traceroute mpls spring-te Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the `traceroute mpls segment-routing spring-te` command.	all levels
ttl	Time to live value of the labeled packet.	none specified

Table 1824: traceroute mpls spring-te Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified
Protocol	Signaling protocol used. For this command, it is LDP.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is **null**.	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	**detail**
Parent	Address of the previous hop. Parent value for the first hop is **null**.	**detail**
Return Code	Return code for reporting the result of processing the echo request by the receiver.	**detail**
Response time	Time for the echo request to reach the receiver.	**detail**
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is **none**.	**detail**
Label Stack	Label stack used to forward the packet.	**detail**

Sample Output

traceroute mpls segment-routing spring-te

```
user@host> traceroute mpls segment-routing spring-te
Possible completions:
  source-routing-path Source Path routing to use when sending probes
  egress-ip                  To/Install IP address to use when sending probes
  label-stack                Label stack for traceroute packets
```

traceroute mpls segment-routing spring-te (source-routing-path) (No FEC validation)

```
user@host> traceroute mpls segment-routing spring-te source-routing-path path_1 active skip-fec-
validation
Probe options: retries 3, exp 7
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   900041  Static      23.1.1.2         (null)           Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  Static      34.1.1.2         23.1.1.2         Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2        3  Static      45.1.1.2         34.1.1.2         Egress
  FEC-Stack-Sent: SPRING-TE

  Path 1 via ge-0/0/1.0 destination 127.0.0.64

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   900041  Static      23.2.1.2         (null)           Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  Static      34.4.1.2         23.2.1.2         Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2        3  Static      45.1.1.2         34.4.1.2         Egress
  FEC-Stack-Sent: SPRING-TE

  Path 2 via ge-0/0/2.0 destination 127.0.1.64
```

traceroute mpls segment-routing spring-te (source-routing-path) (FEC validation)

```
user@host> traceroute mpls segment-routing spring-te source-routing-path path_1 active
Probe options: retries 3, exp 7
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   800005  OSPF        82.2.3.3         (null)           Success
  FEC-Stack-Sent: OSPF,OSPF,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  OSPF        81.3.5.5         82.2.3.3         Egress
  FEC-Stack-Sent: OSPF,OSPF,OSPF FEC-Change-Recieved: POP-OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  OSPF        81.3.5.5         82.2.3.3         Success
  FEC-Stack-Sent: OSPF,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  OSPF        81.5.6.6         81.3.5.5         Egress
  FEC-Stack-Sent: OSPF,OSPF FEC-Change-Recieved: POP-OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  OSPF        81.5.6.6         81.3.5.5         Success
  FEC-Stack-Sent: OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2        3  OSPF        81.6.7.7         81.5.6.6         Egress
  FEC-Stack-Sent: OSPF

  Path 1 via ge-0/0/6.0 destination 127.0.0.64

ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   800005  OSPF        81.2.4.4         (null)           Success
  FEC-Stack-Sent: OSPF,OSPF,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  OSPF        81.4.5.5         81.2.4.4         Egress
  FEC-Stack-Sent: OSPF,OSPF,OSPF FEC-Change-Recieved: POP-OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  OSPF        81.4.5.5         81.2.4.4         Success
  FEC-Stack-Sent: OSPF,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  OSPF        81.5.6.6         81.4.5.5         Egress
  FEC-Stack-Sent: OSPF,OSPF FEC-Change-Recieved: POP-OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  OSPF        81.5.6.6         81.4.5.5         Success
  FEC-Stack-Sent: OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2        3  OSPF        81.6.7.7         81.5.6.6         Egress
```

```
  FEC-Stack-Sent: OSPF

  Path 2 via ge-0/1/0.0 destination 127.0.1.64
```

traceroute mpls segment-routing spring-te (egress-ip) (No FEC Validation)

```
user@host> traceroute mpls segment-routing spring-te egress-ip 5.5.5.5 color 6 skip-fec-
validation
Probe options: retries 3, exp 7
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   900041  Static      23.1.1.2         (null)           Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  Static      34.1.1.2         23.1.1.2         Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2        3  Static      45.1.1.2         34.1.1.2         Egress
  FEC-Stack-Sent: SPRING-TE

  Path 1 via ge-0/0/1.0 destination 127.0.0.64

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   900041  Static      23.2.1.2         (null)           Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  Static      34.4.1.2         23.2.1.2         Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2        3  Static      45.1.1.2         34.4.1.2         Egress
  FEC-Stack-Sent: SPRING-TE

  Path 2 via ge-0/0/2.0 destination 127.0.1.64
```

traceroute mpls segment-routing spring-te (egress-ip) (FEC Validation)

```
user@host> traceroute mpls segment-routing spring-te egress-ip 81.7.7.7 color 5 secondary
Probe options: retries 3, exp 7

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   804101  ISIS        82.2.3.3         (null)           Success
```

```
  FEC-Stack-Sent: ISIS,ISIS,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  ISIS        81.3.5.5         82.2.3.3         Egress
  FEC-Stack-Sent: ISIS,ISIS,OSPF FEC-Change-Recieved: POP-ISIS
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  ISIS        81.3.5.5         82.2.3.3         Success
  FEC-Stack-Sent: ISIS,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  ISIS        81.4.5.4         81.3.5.5         Success
  FEC-Stack-Sent: ISIS,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2   800007  ISIS        81.4.5.5         81.4.5.4         Success
  FEC-Stack-Sent: ISIS,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    3   800007  ISIS        81.5.6.6         81.4.5.5         Success
  FEC-Stack-Sent: ISIS,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    4        3  ISIS        81.6.7.7         81.5.6.6         Egress
  FEC-Stack-Sent: ISIS,OSPF FEC-Change-Recieved: POP-ISIS
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    4        3  ISIS        81.6.7.7         81.5.6.6         Egress
  FEC-Stack-Sent: OSPF

  Path 1 via ge-0/0/6.0 destination 127.0.0.64

ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   804101  ISIS        81.2.4.4         (null)           Success
  FEC-Stack-Sent: ISIS,ISIS,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  ISIS        81.4.5.5         81.2.4.4         Egress
  FEC-Stack-Sent: ISIS,ISIS,OSPF FEC-Change-Recieved: POP-ISIS
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  ISIS        81.4.5.5         81.2.4.4         Success
  FEC-Stack-Sent: ISIS,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  ISIS        81.4.5.4         81.4.5.5         Success
  FEC-Stack-Sent: ISIS,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2   800007  ISIS        81.4.5.5         81.4.5.4         Success
  FEC-Stack-Sent: ISIS,OSPF
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    3   800007  ISIS        81.5.6.6         81.4.5.5         Success
  FEC-Stack-Sent: ISIS,OSPF
```

```
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    4        3  ISIS        81.6.7.7         81.5.6.6         Egress
  FEC-Stack-Sent: ISIS,OSPF FEC-Change-Recieved: POP-ISIS
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    4        3  ISIS        81.6.7.7         81.5.6.6         Egress
  FEC-Stack-Sent: OSPF
  Path 2 via ge-0/1/0.0 destination 127.0.1.64
```

traceroute mpls segment-routing spring-te (label-stack)

```
user@host> traceroute mpls segment-routing spring-te label-stack labels 900006 labels 900041
nexthop-address 12.1.1.2 nexthop-interface ge-0/0/0.0 egress 6.6.6.6
Probe options: retries 3, exp 7
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   900041  Static      12.1.1.2         (null)           Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2   900041  Static      23.1.1.2         12.1.1.2         Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  Static      34.1.1.2         23.1.1.2         Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2   900006  Static      45.1.1.2         34.1.1.2         Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    3        3  Static      56.1.1.2         45.1.1.2         Egress
  FEC-Stack-Sent: SPRING-TE

  Path 1 via ge-0/0/0.0 destination 127.0.0.64
```

Release Information

Command introduced in Junos OS Release 20.2R1.

traceroute mpls segment-routing static egress-ip

IN THIS SECTION

- Syntax | 29970
- Description | 29970
- Options | 29970
- Additional Information | 29972
- Required Privilege Level | 29973
- Output Fields | 29973
- Sample Output | 29974
- Release Information | 29976

Syntax

```
traceroute mpls segment-routing static egress-ip
<egress-ip ( destination |  detail | exp | fanout | logical-system | no-resolve | retries |
routing-instance | source | ttl | wait) >
```

```
traceroute mpls segment-routing spring-te label-stack
<label-stack ( destination | detail | egress | exp | fanout | labels | logical-system | nexthop-
address | nexthop-interface | no-resolve | paths | retries | source | ttl | wait) >
```

Description

Use `traceroute mpls segment-routing static` as a debugging tool to locate MPLS label-switched path forwarding issues for prpd programmed network.

Options

egress-ip Specify to trace to or install IP address to use when sending probes.

destination Specify the destination address to use when sending probes.

detail	(Optional) Display detailed output.
exp *exp*	(Optional) Specify the class of service to use when sending probes. • **Range:** 0 through 7
fanout*fanout*	(Optional) Specify the maximum number of nexthops to search per node. • **Range:** 1 through 64
logical-system *logical-system-name*	(Optional) Specify the name of the logical system for the traceroute attempt.
no-resolve	(Optional) Specify not to attempt to print addresses symbolically.
paths *paths*	(Optional) Specify maximum number of paths to traverse. • **Range:** 1 through 255
retries *retries-number*	(Optional) Specify the number of times to resend probe values. • **Range:** 1 through 9
source *source-address*	(Optional) Specify the source address of the outgoing traceroute packets.
ttl *value*	(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds. • **Range:** 1 through 255
wait*wait*	(Optional) Specify the number of seconds to wait before resending the outgoing traceroute packets. • **Range:** 5 through 15

label-stack Specify the label stack for traceroute packets.

detail	(Optional) Display detailed output.
egress	(Optional) Specify the egress IP address.
labels	(Optional) Specify the labels in label stack. • **Range:** 16 through 1048575

logical-system *logical-system-name*	(Optional) Specify the name of the logical system for the traceroute attempt.
nexthop-address *address*	(Optional) Specify the nexthop IP address for the traceroute packet.
nexthop-interface *interface*	(Optional) Specify the outgoing interface for the traceroute packet.
no-resolve	(Optional) Specify not to attempt to print addresses symbolically.
retries *retries-number*	(Optional) Specify the number of times to resend probe values. • **Range:** 1 through 9
source *source-address*	(Optional) Specify the source address of the outgoing traceroute packets.
ttl *value*	(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds. • **Range:** 1 through 255
wait *value*	(Optional) Specify the number of seconds to wait before resending probe. • **Range:** 5 through 15

Additional Information

NOTE: NIL FEC validation is supported. The FEC validation for SWAN ping and traceroute command is not supported.

NOTE: `traceroute mpls segment-routing spring-te` command with `label-stack` option is supported in operational mode only. It is not supported in configuration mode.

NOTE: Traceroute with label-stack option can be executed only once at a time

NOTE: In case of `traceroute mpls segment-routing spring-te`, the Packet Forwarding Engine supports only 16 ECMP paths. If there more than 16 ECMP paths to the LSP, you can perform traceroute to a maximum of 16 paths.

NOTE: Traceroute with `label-stack` option displays FEC stack sent as SPRING-TE. SPRING-TE implicitly means NIL FEC.

Required Privilege Level

network

Output Fields

The following table describes the output fields for the `traceroute mpls segment-routing static` command. Output fields are listed in the approximate order in which they appear.

Table 1825: traceroute mpls segment-routing static Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the `traceroute mpls segment-routing static` command.	all levels
ttl	Time to live value of the labeled packet.	none specified
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified
Protocol	Signaling protocol used. For this command, it is static.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is **null**.	none specified

Table 1825: traceroute mpls segment-routing static Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Previous Hop	Address of the previous hop in the label-switched path.	**detail**
FEC-Stack-Sent	FEC stack used to forward the packet.	**detail**

Sample Output

traceroute mpls segment-routing static ?

```
user@host> traceroute mpls segment-routing static ?
Possible completions:
     egress-ip                  To/Install IP address to use when sending probes
```

traceroute mpls segment-routing static (egress-ip)

```
user@host> traceroute mpls segment-routing static egress-ip ::100.100.100.4
          Probe options: retries 3, exp 7

          ttl    Label      Protocol    Address          Previous Hop     Probe Status
          1   600040  Static         21.1.2::2          (null)
Success
          FEC-Stack-Sent: NIL
          ttl    Label      Protocol    Address          Previous Hop     Probe Status
          2   600040   Static        21.2.6::2          21.1.2::2             Success
          FEC-Stack-Sent: NIL
          ttl    Label      Protocol    Address          Previous Hop     Probe Status
          3        3        Static     21.3.4.::2         21.2.3::2
Egress
          FEC-Stack-Sent: NIL
```

```
          Path 1 via ge-0/0/2.0 destination ::ffff:127.0.0.64

          ttl    Label     Protocol    Address          Previous Hop     Probe Status
          1   600040  Static        21.1.5::2          (null)
Success
          FEC-Stack-Sent: NIL
          ttl    Label      Protocol    Address          Previous Hop     Probe Status
          2   600040   Static        21.3.5::1          21.1.5::2              Success
          FEC-Stack-Sent: NIL
          ttl    Label      Protocol    Address          Previous Hop     Probe Status
          3        3        Static     21.3.4::2          21.3.5::1
Egress
          FEC-Stack-Sent: NIL

          Path 2 via ge-0/0/7.0 destination ::ffff:127.0.1.64
```

traceroute mpls segment-routing spring-te with label-stack, nexthop-address, nexthop-interface as input

```
user@host>traceroute mpls segment-routing spring-te label-stack labels 600040 nexthop-address
21.1.2.2 nexthop-interface ge-0/0/2.0
      Warning: License key missing; requires 'sr' license

      Probe options: retries 3, exp 7

      ttl    Label     Protocol    Address          Previous Hop     Probe Status
      1     600040    Static      21.1.2:2          (null)               Success
      FEC-Stack-Sent: SPRING-TE
      ttl    Label     Protocol    Address          Previous Hop     Probe Status
      2     600040    Static      21.2.6:2          21.1.2:2              Success
      FEC-Stack-Sent: SPRING-TE
      ttl    Label     Protocol    Address          Previous Hop     Probe Status
      3        3      Static      21.4.6:1          21.2.6:2              Egress
      FEC-Stack-Sent: SPRING-TE

      Path 1 via ge-0/0/2.0 destination 127.0.0.64
```

traceroute mpls segment-routing spring-te (label-stack)

```
user@host> traceroute mpls segment-routing spring-te label-stack labels 900006 labels 900041
nexthop-address 12.1.1.2 nexthop-interface ge-0/0/0.0
Probe options: retries 3, exp 7
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   900041  Static      12.1.1.2         (null)           Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2   900041  Static      23.1.1.2         12.1.1.2         Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1        3  Static      34.1.1.2         23.1.1.2         Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    2   900006  Static      45.1.1.2         34.1.1.2         Success
  FEC-Stack-Sent: SPRING-TE
  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    3        3  Static      56.1.1.2         45.1.1.2         Egress
  FEC-Stack-Sent: SPRING-TE

  Path 1 via ge-0/0/0.0 destination 127.0.0.64
```

Release Information

Command introduced in Junos OS Release 22.4R1.

traceroute overlay

IN THIS SECTION

- Syntax | 29977
- Description | 29977
- Options | 29978
- Required Privilege Level | 29980
- Output Fields | 29980

Sample Output | 29981
Release Information | 29982

Syntax

```
traceroute overlay
<tunnel-type>
vni vni
tunnel-src source-ip-address
tunnel-dst destination-ip-address
<mac mac-address>
<ttl value>
<hash source-udp-port>
<hash-input-interface input-interface>
<hash-source-mac source-mac-address>
<hash-destination-mac destination-mac-address>
<hash-source-address source-IP-address>
<hash-destination-address destination-IP-address>
<hash-vlan vlan-id>
<hash-protocol protocol-id>
<hash-source-port source-layer4-port>
<hash-destination-port destination-layer4-port>
```

Description

Display the route that packets take between two Virtual Extensible LAN (VXLAN) tunnel endpoints (VTEPs) and within the context of a VXLAN overlay segment. Use `traceroute overlay` as an isolation and debugging tool to locate points of failure within an overlay VXLAN tunnel. The output is useful for diagnosing a point of failure in the path from the device to the destination host, and for addressing network traffic latency and throughput problems.

NOTE: The `traceroute overlay` command is not supported for IPv6.

NOTE: The `traceroute overlay` command is not supported when there are multiple virtual-switch routing instances.

Options

tunnel-type

(Optional) Specify the overlay tunnel type used in a virtualized environment such as: VXLAN, Network Virtualization using Generic Routing Encapsulation (NVGRE), MPLS over User Datagram Protocol (UDP), and MPLS over General Routing Encapsulation (GRE) tunnels.

> **NOTE**: Only VXLAN overlay tunnel types are supported.

vni *vni*

Specify the VNI of the VXLAN overlay segment.

- **Range:** 1 through 16,777,215

tunnel-src *source-ip-address*

Specify the IP address of the source entity at end of the tunnel, such as the source VTEP.

tunnel-dst *destination-ip-address*

Specify the IP address of the destination entity at the end of the tunnel, such as the remote VTEP.

mac *mac-address*

(Optional) Include the physical or hardware address on the end host you are trying to reach.

ttl *value*

(Optional) Time-to-live (TTL) value to include as the maximum number of hops in the traceroute request.

For MX Series routers, the range of values is **0** through **255**. The default value is **255**.

- **Range:** (QFX Series, EX9200 switches) 1 through 255
- **Default:** 255

hash *source-udp-port*

(Optional) Specify the UDP source port. The range of values is **1** through **65,535**. Usage of this parameter will override all other hash-* parameters used to calculate the UDP source port.

hash-source-mac *source-mac-address*

(Optional) Specify the MAC address of the source end host.

> **NOTE**: The hash parameters provide values that correspond to a particular data flow that the `traceroute overlay` command debugs. Based on the values that you specify, the system calculates a VXLAN UDP header source port hash, which is included in the VXLAN UDP header of the overlay ping and traceroute packets.

Including the calculated hash in the VXLAN header enables the overlay ping and traceroute packets to emulate data packets in the flow that you are troubleshooting.

When using the hash parameters, we recommend that you specify a value for each parameter. The exception to this guideline is the hash-vlan parameter, which you do not have to use if the source endpoint is not a member of a VLAN. This practice ensures that the overlay ping and traceroute processes are successful and that the output for each command is accurate. If you do not specify a value for one or more of the hash parameters, the system sends an OAM request that might include incorrect hash values and generates a warning message.

Hash computation supports TCP and UDP protocols only.

hash-destination-mac *destination-mac-address*
(Optional) Specify the MAC address of the destination end host.

hash-source-address *source-IP-address*
(Optional) Specify the IP address of the source end host.

hash-destination-address *destination-IP-address*
(Optional) Specify the IP address of the destination end host.

hash-vlan *vlan-id*
(Optional, QFX Series switches only) Specify the VLAN ID of the end host.

- **Range:** 1 through 4094

hash-input-interface *interface-name*
(Optional, QFX Series switches only) Specify the ingress interface of the flow on the Juniper Networks device.

hash-protocol *protocol-id*
(Optional) Specify the TCP/UDP IP protocol ID of the end host.

- **Range:** 1 through 255

hash-source-port *source-layer4-port*
(Optional) Specify the Layer 4 source port of the end host.

- **Range:** 1 through 65,535

hash-destination-port *destination-layer4-port*
(Optional) Specify the Layer 4 destination port of the end host.

- **Range:** 1 through 65,535

Required Privilege Level

network

Output Fields

Use the `traceroute overlay` command to determine overlay segments within a VXLAN tunnel. The output is useful for diagnosing a point of failure in the path from the device to the destination host, and for addressing network traffic latency and throughput problems.

Table 1826 on page 29980 lists the output fields for the `traceroute overlay` command. Output fields are listed in the approximate order in which they appear.

Table 1826: traceroute overlay Output Fields

Field Name	Field Description
vni	The VNI of the VXLAN overlay segment.
tunnel src ip	The IP address of the source end of the tunnel.
tunnel dst ip	The IP address of the destination end of the tunnel.
mac address	The physical or hardware address of the end host you are trying to reach.
ttl	TTL value for the maximum number of hops in the traceroute request.
hash-parameters	The hash parameters provide the input interface, source MAC address, destination MAC address, source IP address, destination IP address, and the VLAN ID of the two end hosts within an overlay segment. Hash parameters enable platform-specific hash computation to use as the source port in the outer UDP header.
ttl	Number of hops remaining in the traceroute message. The TTL is decremented at each hop.
Address	The sending IPv4 address.

Table 1826: traceroute overlay Output Fields *(Continued)*

Field Name	Field Description
Sender Timestamp	Timestamp in microseconds when hop was sent.
Receiver Timestamp	Timestamp in microseconds when hop was received.
Response Time	Time in microseconds for traceroute to respond.

Sample Output

run traceroute overlay

```
user@host> traceroute overlay tunnel-type vxlan vni 100 tunnel-src 192.0.2.10 tunnel-dst
192.0.2.20 mac 00:00:5E:00:53:cc hash-source-mac 00:00:5E:00:53:aa hash-destination-mac
00:00:5E:00:53:cc hash-source-address 198.51.100.1 hash-destination-address 198.51.100.3 hash-
vlan 150 hash-input-interface xe-0/0/2 hash-protocol 17 hash-source-port 4456 hash-destination-
port 4540
traceroute-overlay protocol vxlan

        vni 100
        tunnel src ip 192.0.2.10
        tunnel dst ip 192.0.2.20
        mac address 00:00:5E:00:53:cc
        ttl 255

        hash-parameters:
                input-ifd-idx 653
                end-host smac 00:00:5E:00:53:aa
                end-host dmac 00:00:5E:00:53:cc
                end-host src ip 198.51.100.1
                end-host dst ip 198.51.100.3
                end-host protocol 17
                end-host l4-src-port 4456
                end-host l4-dst-port 4540
                end-host vlan 150

ttl  Address    Sender Timestamp               Receiver Timestamp             Response Time
```

```
1   10.1.0.1   09-25 00:56:17 PDT.663 msecs            *                      10 msecs
2   192.0.2.20 09-25 00:56:17 PDT.684 msecs   09-25 00:56:17 PDT.689 msecs   11 msecs

Overlay-segment  present at RVTEP 192.0.2.20



    End-System  Present
```

Release Information

Command introduced in Junos OS Release 14.1X53-D30.

RELATED DOCUMENTATION

Understanding Overlay ping and traceroute Packet Support

Example: Troubleshoot a VXLAN Overlay Network with Overlay Ping and Overlay Traceroute on QFX Series Switches

ping overlay | 12596

traceroute srv6

IN THIS SECTION

- Syntax | 29983
- Description | 29983
- Options | 29983
- Required Privilege Level | 29984
- Output Fields | 29984
- Sample Output | 29986
- Release Information | 29987

Syntax

```
routing-options {
source-packet-routing {
srv6 {
+ oam {
+ traceoptions {
file;
flag {
all;
configuration;
database;
network;
pipe;
rpc-packet-details;
traceroute;
)
}
}
}
}
}
```

Description

An SRv6 segment is a 128-bit value Segment Identifier (SID) often used as a shorter reference for SRv6 Segment. Traceroute to a SID is used for hop-by-hop fault localization and to trace path to a SID.

Options

logical-system (Optional) Specify the name of the logical system for the traceroute attempt.

retries ***retries-number*** (Optional) Specify the number of times to resend probe values.

- **Range:** 1 through 9
- **Default:** 3

routing-instance *routing-instance-name*
(Optional) Specify the name of the routing instance for the traceroute attempt.

source *source-address*
(Optional) Specify the source address of the outgoing traceroute packets.

ttl *value*
(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds.

- **Range:** 1 through 125
- **Default:** 64

wait *seconds*
(Optional) Specify the number of seconds to wait before resending a probe.

- **Range:** 5 through 15
- **Default:** 10

Required Privilege Level

network

Output Fields

"traceroute srv6" on page 29982 describes the output fields for the `traceroute srv6 isis` command and the `traceroute srv6 spring-te` commands. Output fields are listed in the approximate order in which they appear.

Table 1827: Output Fields for traceroute srv6 isis

Field Name	Field Description
host	IPv6 loopback address or SRv6 SID of remote host.
logical-system	Name of logical system.
probe-icmp	Use the ICMP echo request for probes.
retries	Number of times to resend the probe.

Table 1827: Output Fields for traceroute srv6 isis *(Continued)*

Field Name	Field Description
routing-instance	Name of routing instance for traceroute attempt.
source	Source address used when you send probes.
tos	Value for IPv6 type-of-service.
ttl	Time to live value of the labeled packet.
wait	Number of seconds to wait before resending probe.

Table 1828: Output Fields for traceroute srv6 te sids-stack

Field Name	Field Description
logical-system	Name of logical system.
nexthop-address	Nexthop IP address for the traceroute packet.
nexthop-interface	Outgoing interface for the traceroute packet.
no-resolve	Do not attempt to print the addresses symbolically.
probe-icmp	Use the ICMP echo request for probes.
retries	Number of times to resend the probe.
routing-instance	Name of routing instance for traceroute attempt.
sids	SRv6 SIDs in srv6-stack.
source	Source address used when you send probes.

Table 1828: Output Fields for traceroute srv6 te sids-stack *(Continued)*

Field Name	Field Description
tos	Value for IPv6 type-of-service.
ttl	Time to live value of the labeled packet.
wait	Number of seconds to wait before resending probe.

Sample Output

Output for ISIS SRv6

```
user@router> traceroute srv6 isis 3000::1
traceroute: to destination 3000::1, 64 hops max, retries 3
  1   ::21.0.1.2   7.153 ms   4.191 ms   2.652 ms
       2   ::21.1.2.2   2.652 ms   7.153 ms   4.191 ms
  3   ::21.2.3.2   7.323 ms   2.652 ms   7.153 ms
```

Output for IPv6/SRv6

```
user@router> traceroute srv6 spring-te sids-stack sids [ 5000::1 3000::1 ] nexthop-interface
ge-0/0/0.0 nexthop-address ::21.0.1.2

    traceroute: to destination 5000::1, 64 hops max, retries 3
 1   ::21.0.1.2   3.532 ms   4.191 ms   2.652 ms
     DA: 3000::1
     SRH:( 5000::1, 3000::1; SL=1)
2   ::21.1.2.2   9.573 ms    3.532 ms   4.191 ms
     DA: 3000::1
     SRH:( 5000::1, 3000::1; SL=1)
3   3333::3   5.253 ms   3.532 ms   4.191 ms
     DA: 3000::1
```

```
     SRH:( 5000::1, 3000::1; SL=1)
4   ::21.3.5.2   6.800 ms   4.191 ms   2.652 ms
     DA: 5000::1
```

Release Information

Command introduced in Junos OS Release 22.3R1.

tracing

IN THIS SECTION

- Syntax | 29987
- Hierarchy Level | 29987
- Description | 29988
- Options | 29988
- Required Privilege Level | 29988
- Release Information | 29988

Syntax

```
tracing {
    destination-override syslog host ip-address;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the router to enable remote tracing to a specified host IP address. The default setting is disabled.

The following processes are supported:

- chassisd—Chassis-control process
- eventd—Event-processing process
- cosd—Class-of-service process
- spd—Adaptive-services process

You can use the `no-remote-trace` statement, under the `[edit system process-name traceoptions]` hierarchy, to disable remote tracing.

Options

`destination-override syslog host` *`ip-address`*—Overrides the global config under `system tracing` and has no effect if `system tracing` is not configured.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Junos OS Tracing and Logging Operations

destination-override | 5454

no-remote-trace

track-igp-metric

IN THIS SECTION

- Syntax | 29989
- Hierarchy Level | 29989
- Description | 29989
- Required Privilege Level | 29989
- Release Information | 29990

Syntax

```
track-igp-metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 18.4R1 on MX5, MX10, MX40, MX104, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, and MX10008 routers.

Statement introduced in Junos OS Release 18.4R1 on PTX1000, PTX3000, PTX5000, PTX10002-60C, PTX10008, and PTX10016 routers.

RELATED DOCUMENTATION

Configuring LDP to Use the IGP Route Metric

track-igp-metric (LSP)

IN THIS SECTION

- Syntax | **29990**
- Hierarchy Level | **29991**
- Description | **29991**
- Options | **29991**
- Required Privilege Level | **29991**
- Release Information | **29991**

Syntax

```
track-igp-metric <install-v4-prefixes> <install-v6-prefixes>;
```

Hierarchy Level

The hierarchy level for `track-igp-metric` globally enabled for all LSPs:

```
[edit protocols mpls]
```

The hierarchy level for `track-igp-metric` at the per LSP level:

```
[edit protocols mpls label-switched-path pathname]
```

Description

Track IGP metric for LSP install prefixes

Options

install-v4-prefixes	Track IGP metric for IPV4 prefixes.
install-v6-prefixes	Track IGP metric for IPV6 prefixes.

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

- Install Prefix IGP Overview

track (VRRP)

IN THIS SECTION

- Syntax | 29992
- Syntax for cRPD | 29992
- Hierarchy Level | 29993
- Description | 29993
- Options | 29993
- Required Privilege Level | 29993
- Release Information | 29994

Syntax

```
track {
    interface interface-name {
        bandwidth-threshold bits-per-second priority-cost priority;
        priority-cost priority;
    }
    priority-hold-time seconds;
    route prefix/prefix-length routing-instance instance-name priority-cost priority;
}
```

Syntax for cRPD

```
track {
    interface interface-name
    notify-backup {script-name name of script;
    username  username}
    notify-fault {script-name name of script;
    username  username}
    notify-master {script-name name of script;
    username  username}
    notify-script {Notify script for VRRP group}
    track-script {Script instance;
```

```
    keepalived_check fall/interval/rise script-name name of script/time-out;
        weight cost/reverse;}
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-
group group-id],
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-
group group-id],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet address address vrrp-group group-id],
[edit logical-systems logical-system-name  interfaces interface-name unit logical-unit-number
family inet6 address address vrrp-inet6-group group-id]
```

Description

Enable logical interface tracking, route tracking, or both, for a Virtual Router Redundancy Protocol (VRRP) group.

Options

The remaining statements are described separately.

notify-backup	Notify script to BACKUP transition
notify-fault	Notify script to FAULT transition
notify-master	Notify script to MASTER transition
notify-script	Notify script for VRRP group
track-script	Tracking script for VRRP group

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`priority-hold-time` statement added in Junos OS Release 8.1.

`route` statement added in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring a Logical Interface to Be Tracked for a VRRP Group

Configuring a Route to Be Tracked for a VRRP Group

tracking-area

IN THIS SECTION

- Syntax | 29994
- Hierarchy Level | 29995
- Description | 29995
- Options | 29995
- Required Privilege Level | 29995
- Release Information | 29995

Syntax

```
tracking-area code {
 plmn name {
  s-nssai name {
   sd sd;
   sst(embb | miot | urllc | v2x);
  }
```

```
  }
 }
```

Hierarchy Level

```
[edit services agf]
```

Description

Tracking area for the AMF.

Options

name Tracking area code

- **Range:** 0 through 16777215

plmn Public land mobile network name

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3R1.

traffic-class (delay)

IN THIS SECTION

- Syntax | 29996
- Hierarchy Level | 29996
- Description | 29996
- Options | 29997
- Required Privilege Level | 29997
- Release Information | 29998

Syntax

```
traffic-class tc-value {
    average-sample-size sample size;
    padding-size size;
    query-interval milliseconds;
    rtt-delay-threshold rtt threshold value;
    twcd-delay-threshold twcd threshold value;
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier delay],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier delay],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring
querier delay],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring
querier delay]
```

Description

Configure traffic class specific options.

Specify the traffic classes for which loss measurement has to be performed. This parameter takes one of the *tc-all|tc-0|tc-1|tc-2|tc-3|tc-4|tc-5|tc-6|tc-7|tc-none* traffic-class values. For each traffic class, you can configure the respective parameters.

To enable traffic-class parameters, configure the `traffic-class-statistics` configuration statement under the `[edit protocol mpls statistic]` hierarchy level.

Options

average-sample-size *sample size* (Optional) Specify the number of samples used for calculating the average of various metrics.

- **Default:** 5
- **Range:** 1 through 30

padding-size *size* (Optional) Specify the delay-measurement message length, which is used to calculate the delay experienced by messages of different sizes.

- **Default:** 0
- **Range:** 1 through 1500

query-interval *milliseconds* Specify the minimum transmit interval, which signifies how often the loss measurement message is generated from the querier.

- **Default:** 10 seconds
- **Range:** 1000 through 4294967295 milliseconds

rtt-delay-threshold *rtt threshold value* Specify the round-trip delay threshold value.

- **Range:** 1 through 4294967295 microseconds

twcd-delay-threshold *twcd threshold value* Specify the two-way channel delay threshold value.

- **Range:** 1 through 4294967295 microseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Pro-Active Loss and Delay Measurements

On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview

performance-monitoring (Protocols MPLS) | **12434**

traffic-class (loss)

IN THIS SECTION

- Syntax | **29998**
- Hierarchy Level | **29999**
- Description | **29999**
- Options | **29999**
- Required Privilege Level | **30000**
- Release Information | **30000**

Syntax

```
traffic-class tc-value {
    average-sample-size sample size;
    loss-threshold loss threshold value;
    loss-threshold-window number of samples for loss threshold;
    measurement-quantity bytes|packets;
    query-interval milliseconds;
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier loss],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier loss],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring
querier loss],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring
querier loss]
```

Description

Configure traffic class specific options.

Specify the traffic classes for which loss measurement has to be performed. This parameter takes one of the *tc-all|tc-0|tc-1|tc-2|tc-3|tc-4|tc-5|tc-6|tc-7|tc-none* traffic-class values. For each traffic class, you can configure the respective parameters.

To enable traffic-class parameters, configure the `traffic-class-statistics` configuration statement under the `[edit protocol mpls statistic]` hierarchy level.

Options

average-sample-size *sample size* — (Optional) Specify the number of samples used for calculating the average of various metrics.

- **Default:** 5
- **Range:** 1 through 30

loss-threshold *loss threshold value* — Specify the threshold value that will be used with `loss-threshold-window` to calculate the loss within specified window size.

- **Range:** 1 through 4294967295

loss-threshold-window *number of samples for loss threshold* — Specify the number of samples used for loss threshold calculation.

- **Range:** 1 through 30

measurement-quantity *bytes|packets* — (Optional) Specify whether packet or byte loss is being measured at the querier.

- **Default:** packets

query-interval ***milliseconds*** Specify the minimum transmit interval, which signifies how often the loss measurement message is generated from the querier.

- **Default:** 10 seconds
- **Range:** 1000 through 4294967295 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Pro-Active Loss and Delay Measurements

On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview

performance-monitoring (Protocols MPLS) | **12434**

traffic-class (loss-delay)

IN THIS SECTION

- Syntax | **30001**
- Hierarchy Level | **30001**
- Description | **30001**
- Options | **30001**
- Required Privilege Level | **30002**
- Release Information | **30003**

Syntax

```
traffic-class tc-value {
    average-sample-size sample size;
    loss-threshold loss threshold value;
    loss-threshold-window number of samples for loss threshold;
    measurement-quantity bytes|packets;
    padding-size size;
    query-interval milliseconds;
    rtt-delay-threshold rtt threshold value;
    twcd-delay-threshold twcd threshold value;
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier loss-delay],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier loss-delay],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring
querier loss-delay],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring
querier loss-delay]
```

Description

Configure traffic class specific options.

Specify the traffic classes for which loss measurement has to be performed. This parameter takes one of the *tc-all|tc-0|tc-1|tc-2|tc-3|tc-4|tc-5|tc-6|tc-7|tc-none* traffic-class values. For each traffic class, you can configure the respective parameters.

To enable traffic-class parameters, configure the `traffic-class-statistics` configuration statement under the `[edit protocol mpls statistic]` hierarchy level.

Options

average-sample-size ***sample size***

(Optional) Specify the number of samples used for calculating the average of various metrics.

- **Default:** 5

- **Range:** 1 through 30

loss-threshold ***loss threshold value*** Specify the threshold value that will be used with `loss-threshold-window` to calculate loss within specified window size.

- **Range:** 1 through 4294967295

loss-threshold-window ***number of samples for loss threshold*** Specify the number of samples used for loss threshold calculation.

- **Range:** 1 through 30

measurement-quantity ***bytes/packets*** (Optional) Specify whether packet or byte loss is being measured at the querier.

- **Default:** packets

padding-size ***size*** (Optional) Specify the delay-measurement message length, which is used to calculate the delay experienced by messages of different sizes.

- **Default:** 0
- **Range:** 1 through 1500

query-interval ***milliseconds*** Specify the minimum transmit interval, which signifies how often the loss measurement message is generated from the querier.

- **Default:** 10 seconds
- **Range:** 1000 through 4294967295 milliseconds

rtt-delay-threshold ***rtt threshold value*** Specify the round-trip delay threshold value.

- **Range:** 1 through 4294967295 microseconds

twcd-delay-threshold ***twcd threshold value*** Specify the two-way channel delay threshold value.

- **Range:** 1 through 4294967295 microseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Pro-Active Loss and Delay Measurements

On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview

performance-monitoring (Protocols MPLS) | **12434**

traffic-class-map

IN THIS SECTION

- Syntax | **30003**
- Hierarchy Level | **30004**
- Description | **30004**
- Options | **30005**
- Required Privilege Level | **30005**
- Release Information | **30005**

Syntax

```
traffic-class-map {
    dscp traffic-class-map-name{
        traffic-class {
            real-time code-points value;
            network-control code-points value;
            best-effort code-points value;
        }
    }
    exp traffic-class-map-name{
        traffic-class {
```

```
            real-time code-points value;
            network-control code-points value;
            best-effort code-points value;
        }
    }
    ieee-802.1 traffic-class-map-name{
        traffic-class {
            real-time code-points value;
            network-control code-points value;
            best-effort code-points value;
        }
    }
    ieee-802.1ad traffic-class-map-name{
        traffic-class {
            real-time code-points value;
            network-control code-points value;
            best-effort code-points value;
        }
    }
    inet precedence traffic-class-map-name{
        traffic-class {
            real-time code-points value;
            network-control code-points value;
            best-effort code-points value;
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Define the traffic class map for a packet on the basis of the CoS code points.

Traffic class map is a user-configurable input priority map that helps prioritize and classify traffic entering a Packet Forwarding Engine during ingress oversubscription. You can define traffic class maps for a packet on the basis of the following CoS code points:

- Differentiated Services code point (DSCP) for IP DiffServ

- IP precedence bits
- MPLS EXP bits
- IEEE 802.1 bits
- IEEE 802.1ad drop eligible indicator (DEI) bits

You can associate the code point values to one of the following traffic classes:

- Real-time
- Network control
- Best-effort

Options

`dscp`—DSCP traffic class map. Applies to both IPv4 and IPv6 traffic.

`exp`—MPLS EXP traffic class map.

`ieee-802.1`—IEEE 802.1 traffic class map.

`ieee-802.1ad`—IEEE 802.1ad traffic class map.

`inet-precedence`—IPv4 precedence traffic class map.

`traffic-class-map-name`—Name of the input priority map.

`traffic-class best-effort`—Map code points to a best-effort traffic class.

`traffic-class network-control`—Map code points to a network control traffic class.

`traffic-class real-time`—Map code points to a real-time traffic class.

`code-points` *`[code-point-alias | code-point-value]`*—List of code point aliases and bit strings.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

show class-of-service forwarding-table traffic-class-map | 19051

show class-of-service traffic-class-map | 19142

Managing Ingress Oversubscription at the PFE

Configuring Traffic Class Maps to Manage Ingress Oversubscription

Example: Configuring Traffic Class Maps

traffic-class-map (Apply to Interface)

IN THIS SECTION

- Syntax | 30006
- Hierarchy Level | 30006
- Description | 30007
- Options | 30007
- Required Privilege Level | 30007
- Release Information | 30007

Syntax

```
traffic-class-map {
    (dscp | inet precedence) traffic-class-map-name;
    exp traffic-class-map-name;
    (ieee-802.1p | ieee-802.1ad) traffic-class-map-name <vlan-tag (inner | outer)>;
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name ]
```

Description

Associate traffic class maps with an interface.

Options

`dscp`—DSCP traffic class map; applies to both IPv4 and IPv6 traffic.

`exp`—MPLS EXP traffic class map.

`ieee-802.1`—IEEE 802.1 traffic class map.

`ieee-802.1ad`—IEEE 802.1ad traffic class map.

`inet-precedence`—IPv4 precedence traffic class map.

`traffic-class-map-name`—Name of the input priority map.

`vlan-tag inner`—Traffic class map based on the inner VLAN tag to be associated with an interface. By default, the outer VLAN tag is associated with an interface.

`vlan-tag outer`—Traffic class map based on the outer VLAN tag to be associated with an interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

show class-of-service forwarding-table traffic-class-map | **19051**

show class-of-service traffic-class-map | **19142**

Managing Ingress Oversubscription at the PFE

Configuring Traffic Class Maps to Manage Ingress Oversubscription

Example: Configuring Traffic Class Maps

traffic-class (Tunnels)

IN THIS SECTION

- Syntax | 30008
- Hierarchy Level | 30008
- Description | 30008
- Options | 30009
- Required Privilege Level | 30009
- Release Information | 30009

Syntax

```
traffic-class traffic-class-value;
```

Hierarchy Level

```
[edit interfaces gre-interface name unit unit number  tunnel ]
[edit logical-systems logical-system-name interfaces gre-interface-name unit unit number  tunnel]
```

Description

Configure a static ToS/Traffic Class value in the IPv4/IPv6 header, respectively, of a GRE tunnel. This can be useful, for example, when you want to apply a different classification to port-mirrored traffic than the source traffic.

This setting overrides ToS reflection set by copying the inner IP header's ToS value to the outer IP packet header.

NOTE: If rewrite rules are configured on the egress WAN interface, those rewrite rules will overwrite this setting. Therefore the `traffic-class` setting only makes sense when no rewrite rules are configured.

Options

traffic-class-value

The value represents the entire 8-bit differentiated services (DS) field in the IP header, and should be chosen based on the desired DSCP/IP precedence value. For example, if a DSCP value of `111000` is desired, then configure the `traffic-class` value to be `224` (corresponding to `111000 00`).

- **Range:** 0-255

Required Privilege Level

interface

Release Information

Statement introduced before Junos OS Release 11.2R1.

RELATED DOCUMENTATION

CoS for Tunnels Overview

Configuring CoS for GRE and IP-IP Tunnels

Understanding Port Mirroring

traffic-control-profiles

IN THIS SECTION

- Syntax | 30010
- Hierarchy Level | 30011
- Description | 30011
- Options | 30011
- Required Privilege Level | 30012
- Release Information | 30012

Syntax

EX Series (Except EX4600), M Series, MX Series, PTX Series, T Series

```
traffic-control-profiles profile-name {
    adjust-minimum rate;
    atm-service (cbr | rtvbr | nrtvbr);
    delay-buffer-rate (percent percentage | rate);
    excess-rate (percent percentage  | proportion value );
    excess-rate-high (percent percentage | proportion value);
    excess-rate-low (percent percentage | proportion value);
    guaranteed-rate (percent percentage | rate) <burst-size bytes>;
    max-burst-size cells;
    overhead-accounting (frame-mode | cell-mode | frame-mode-bytes | cell-mode-bytes) <bytes (byte-
value)>;
    peak-rate rate;
    scheduler-map map-name;
    shaping-rate (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
    strict-priority-scheduler;
    sustained-rate rate;
}
```

QFX Series including QFabric, OCX OCX1100, EX4600, NFX Series

```
traffic-control-profiles profile-name {
    guaranteed-rate (rate| percent percentage);
    scheduler-map map-name;
    shaping-rate (rate| percent percentage);
}
```

ACX Series

```
traffic-control-profiles profile-name {
    atm-service (cbr | nrtvbr | rtvbr);
    delay-buffer-rate cps;
    max-burst-size max-burst-size;
    peak-rate peak-rate;
    sustained-rate sustained-rate;
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

ACX Series Routers

Configure traffic-shaping profiles.

NOTE: For CoS on ACX6360-OR, see the documentation for the PTX1000.

EX Series (Except EX4600), M Series, MX Series, T Series, and PTX Series Routers

For Gigabit Ethernet IQ, Channelized IQ PICs, FRF.15 and FRF.16 LSQ interfaces, Enhanced Queuing (EQ) DPCs, and PTX Series routers only, configure traffic shaping and scheduling profiles. For Enhanced EQ PICs, EQ DPCs, and PTX Series routers only, you can include the `excess-rate` statement.

QFX Series QFabric, OCX1100, EX4600, NFX Series

Configure traffic shaping and scheduling profiles for forwarding class sets (priority groups) to implement enhanced transmission selection (ETS) or for logical interfaces.

Options

profile-name—Name of the traffic-control profile. This name is also used to specify an output traffic control profile.

The remaining statements are explained separately. See CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement was introduced in Junos OS Release 7.6 (EX series, M series, MX series, T series, and PTX series devices).

Statement was introduced in Junos OS Release 11.1 for the QFX Series.

Statement was introduced in Junos OS Release 12.3 for ACX series routers.

Statement was introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

RELATED DOCUMENTATION

- Oversubscribing Interface Bandwidth
- Understanding Scheduling on PTX Series Routers
- Example: Configuring CoS Hierarchical Port Scheduling (ETS)
- Example: Configuring Traffic Control Profiles (Priority Group Scheduling)
- Example: Configuring Forwarding Class Sets
- Assigning CoS Components to Interfaces
- output-traffic-control-profile
- Understanding CoS Traffic Control Profiles

traffic-control-profiles (Dynamic CoS Definition)

IN THIS SECTION

- Syntax | 30013
- Hierarchy Level | 30013
- Description | 30014
- Options | 30014

- Required Privilege Level | **30014**
- Release Information | **30014**

Syntax

```
traffic-control-profiles profile-name {
    adjust-minimum rate;
    delay-buffer-rate (percent percentage | rate);
    excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
    excess-rate-high (percent percentage | proportion value);
    excess-rate-low (percent percentage | proportion value);
    guaranteed-rate (percent percentage | rate) <burst-size  bytes>;
    max-burst-size cells;
    overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
    peak-rate rate;
    scheduler-map map-name;
    shaping-rate (percent percentage | rate | predefined-variable) <burst-size  bytes>;
    shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
    sustained-rate rate;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service]
```

Description

Configure traffic shaping and scheduling profiles for use in a dynamic client profile or a dynamic service profile.

Options

profile-name—Name of the traffic-control profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Traffic Scheduling and Shaping for Subscriber Access

Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers

traffic-engineering

IN THIS SECTION

- Syntax | 30015
- Hierarchy Level | 30015
- Description | 30015
- Default | 30015

- Required Privilege Level | 30015
- Release Information | 30015

Syntax

```
traffic-engineering {
    disable;
}
```

Hierarchy Level

```
[edit protocols ospf | isis]
```

Description

Enable the traffic engineering features of the specified routing protocol.

Default

Traffic engineering is disabled.

Starting in Junos OS release 15.1, traffic engineering is enabled by default whenever the IS-IS protocol is enabled. You can disable it by including the `disable` statement at the `[edit protocols isis traffic-engineering]` hierarchy level. For the EX3300, EX4200, EX4500, EX4550, EX8200 and XRE200, you can disable traffic engineering starting in Junos OS release 15.1R7.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Example: Configuring MPLS on EX8200 and EX4500 Switches

Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect

Configuring MPLS on Provider Edge Switches Using IP-Over-MPLS

Configuring MPLS on EX8200 and EX4500 Provider Switches

Configuring an OSPF Network (J-Web Procedure)

MPLS Applications User Guide

traffic-engineering (policy-options)

IN THIS SECTION

- Syntax | **30016**
- Hierarchy Level | **30017**
- Description | **30017**
- Options | **30017**
- Required Privilege Level | **30017**

Syntax

```
traffic-engineering {
                   apply-groups
                   apply-groups-except
                   groups
                   identifier
                   ipv4-prefix
                   ipv6-prefix
                   link
                   node
                   protocol
                   srv6-sid
                   te-policy
                    |
```

```
                    }
```

Hierarchy Level

```
[edit policy-options policy-statement name from]
```

Description

Specify traffic-engineering related parameters.

Options

identifier BGP-TE identifier (2..18446744073709551615)

ipv4-prefix IPv4 prefix-related parameters

ipv6-prefix IPv6 prefix related parameters

link Link-related parameters

node Node-related parameters

protocol Protocol that originated the entry

srv6-sid SRv6 SID related parameters

te-policy Traffic engineering related parameters

Required Privilege Level

routing

RELATED DOCUMENTATION

Link-State Distribution using BGP

traffic-engineering (OSPF)

IN THIS SECTION

- Syntax | 30018
- Hierarchy Level | 30019
- Description | 30019
- Default | 30019
- Options | 30019
- Required Privilege Level | 30020
- Release Information | 30020

Syntax

```
traffic-engineering {
    <advertise-unnumbered-interfaces>;
    <credibility-protocol-preference>;
    ignore-lsp-metrics;
    multicast-rpf-routes;
    no-topology;
    l3-unicast-topology;
    shortcuts {
        lsp-metric-into-summary;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit protocols (ospf | ospf3)]
```

Description

Enable the OSPF traffic engineering features.

Default

Traffic engineering support is disabled.

Options

`advertise-unnumbered-interfaces`—(Optional) (OSPFv2 only) Include the link-local identifier in the link-local traffic-engineering link-state advertisement. This statement must be included on both ends of an unnumbered link to allow an ingress LER to update the link in its traffic engineering database and use it for CSPF calculations. The link-local identifier is then used by RSVP to signal unnumbered interfaces as defined in RFC 3477.

`credibility-protocol-preference`—(Optional) (OSPFv2 only) Use the configured preference value for OSPF routes to calculate the traffic engineering database credibility value used to select IGP routes. Use this statement to override the default behavior, in which the traffic engineering database prefers IS-IS routes even if OSPF routes are configured with a lower, that is, preferred, preference value. For example, OSPF routes have a default preference value of 10, whereas IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502, whereas IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.

`ignore-lsp-metrics`—(Optional) Ignore RSVP LSP metrics in OSPF traffic engineering shortcut calculations.

`multicast-rpf-routes`—(Optional) (OSPFv2 only) Install routes for multicast RPF checks into the inet.2 routing table. The inet.2 routing table consists of unicast routes used for multicast RPF lookup. RPF is an antispoofing mechanism used to check whether the packet is coming in on an interface that is also sending data back to the packet source.

> **NOTE**: You must enable OSPF traffic engineering shortcuts to use the `multicast-rpf-routes` statement. You must not allow LSP advertisements into OSPF when configuring the `multicast-rpf-routes` statement.

`no-topology`—(Optional) (OSPFv2 only) Disable the dissemination of the link-state topology information.

`l3-unicast-topology`—Download IGP topology information into the traffic engineering database (TED). In Junos OS, the IGPs install topology information into a database called the traffic engineering database. The traffic engineering database contains the aggregated topology information. The IGP routes are installed by the traffic engineering database on behalf of the corresponding IGP into a user-visible routing table called lsdist.0, subject to route policies.

The remaining statements are explained separately. See CLI Explorer.

> **CAUTION**: When the OSPF traffic engineering configuration is considerably modified, the routing table entries are deleted and the routing table is recreated. Changes to configuration that can cause this behavior include enabling or disabling:
>
> - Traffic engineering shortcuts
> - IGP shortcuts
> - LDP tunneling
> - Multiprotocol LSP
> - Advertise summary metrics
> - Multicast RPF routes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`multicast-rpf-routes` option introduced in Junos OS Release 7.5.

`advertise-unnumbered-interfaces` option introduced in Junos OS Release 8.5.

Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4.

Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4 for EX Series switches.

`credibility-protocol-preference` statement introduced in Junos OS Release 9.4.

`credibility-protocol-preference` statement introduced in Junos OS Release 9.4 for EX Series switches.

Support for `l3-unicast-topology` statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Example: Enabling OSPF Traffic Engineering Support

traffic-engineering (Passive TE Mode)

IN THIS SECTION

- Syntax | 30021
- Hierarchy Level | 30022
- Description | 30022
- Default | 30022
- Options | 30022
- Required Privilege Level | 30023
- Release Information | 30023

Syntax

```
traffic-engineering {
    remote-node-id address;
    remote-node-router-id address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface
interface-name passive],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast |
ipv6-multicast) area area-id interface interface-name passive],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) area area-id interface interface-name passive],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-
name passive],
[edit protocols (ospf | ospf3) area area-id interface interface-name passive],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id
interface interface-name passive],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface
interface-name passive],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-
multicast | ipv6-multicast) area area-id interface interface-name passive]
```

Description

Configure an interface in OSPF passive traffic engineering mode to enable dynamic discovery of OSPF AS boundary routers.

Default

OSPF passive traffic-engineering mode is disabled.

Options

remote-node-id *address*

IP address at the far end of the inter-AS link.

remote-node-router-id *address*

Router ID at the far end of the inter-AS link.

> **NOTE**: The `remote-node-router-id` *address* option does not apply under the [`edit routing-instances` *routing-instance-name*] and [`edit protocols ospf3 area` *area-id*] hierarchy levels.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

Support for the `realm` statement introduced in Junos OS Release 9.2.

Support for the `realm` statement introduced in Junos OS Release 9.2 for EX Series switches.

`remote-node-router-id` *address* option introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Example: Configuring OSPF Passive Traffic Engineering Mode

Junos OS MPLS Applications User Guide

traffic-engineering (Protocols BGP)

IN THIS SECTION

- Syntax | 30024
- Hierarchy Level | 30024
- Description | 30024
- Options | 30024
- Required Privilege Level | 30025
- Release Information | 30025

Syntax

```
traffic-engineering {
    unicast;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family],
[edit logical-systems logical-system-name protocols bgp group group-name family],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name neighbor address family],
[edit protocols bgp family],
[edit protocols bgp group group-name family],
[edit protocols bgp group group-name neighbor address family],
[edit routing-instances routing-instance-name protocols bgp family],
[edit routing-instances routing-instance-name protocols bgp group group-name family],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address
family]
```

Description

Enable traffic engineering address family. This generates a multiprotocol address family indicator (AFI) and a subsequent address family identifier (SAFI) to be negotiated with the BGP peers.

The BGP network layer reachability information (NLRI) information is exchanged between the peers only when the traffic engineering AFI and SAFI are shared between them. If the peers do not agree on the use of the AFI and SAFI, the connection between the peers is terminated.

Options

unicast	Include BGP-TE NLRI.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced on QFX5100 switches in Junos OS Release 15.1

Statement introduced on QFX10000 switches in Junos OS Release 17.1.

RELATED DOCUMENTATION

> *Example: Configuring Link State Distribution Using BGP*

traffic-engineering (Protocols IS-IS)

IN THIS SECTION

- Syntax | 30025
- Hierarchy Level | 30026
- Description | 30026
- Default | 30027
- Options | 30028
- Required Privilege Level | 30030
- Release Information | 30030

Syntax

```
traffic-engineering {
    disable;
    credibility-protocol-preference;
```

```
    family;
    ignore-lsp-metrics;
    ipv6
    l3-unicast-topology ;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis],
[edit protocols isis]
```

Description

Configure traffic engineering properties for IS-IS.

IS-IS always performs shortest-path-first (SPF) calculations to determine next hops. For prefixes reachable through a particular next hop, IS-IS places that next hop for that prefix in the inet.0 routing table. In addition, for routers running MPLS, IS-IS installs the prefix for IPv4 routes in the inet.3 routing table as well. The inet.3 table, which is present on the ingress router, contains the host address of each MPLS label-switched path (LSP) egress router. BGP uses this routing table to resolve next-hop addresses.

If you enable IS-IS traffic engineering shortcuts and if there is a label-switched path to a point along the path to that prefix, IS-IS installs the prefix in the inet.3 routing table and uses the LSP as a next hop. The net result is that for BGP egress routers for which there is no LSP, BGP automatically uses an LSP along the path to reach the egress router.

In Junos OS Release 9.3 and later, IS-IS traffic engineering shortcuts support IPv6 routes. LSPs to be used for shortcuts continue to be signaled using IPv4. However, by default, shortcut routes calculated through IPv6 routes are added to the inet6.3 routing table. The default behavior is for only BGP to use LSPs in its calculations. If you configure MPLS so that both BGP and interior gateway protocols use LSPs for forwarding traffic, shortcut routes calculated through IPv6 are added to the inet6.0 routing table. IS-IS ensures that the IPv6 routes running over the IPv4 MPLS LSP are correctly de-encapsulated at the tunnel egress by pushing an extra IPv6 explicit null label between the IPv6 payload and the IPv4 transport label.

RSVP LSPs with a higher preference than IS-IS routes are not considered during the computation of traffic engineering shortcuts.

To configure IS-IS so that it uses LSPs as shortcuts when installing information in the inet.3 or inet6.3 routing table, include the following statements:

```
family inet {
    shortcuts {
        multicast-rpf-routes;
    }
}
family inet6 {
    shortcuts;
}
```

For IPv4 traffic, include the `inet` statement. For IPv6 traffic, include the `inet6` statement.

To configure IPv4 MPLS or IPv6 MPLS shortcuts explicitly for segment routing, include the `inet-mpls` statement for IPv4 MPLS traffic and the `inet6-mpls` statement for IPv6 MPLS traffic.

To configure load balancing across multiple LSPs, include the `multipath` statement.

When traffic engineering shortcuts are used, RSVP first looks at the `metric2` value, which is derived from the IGP cost. After this, RSVP considers the LSP metric value. So, if a certain path changes for an LSP and the cost changes, not all LSPs are used to load- balance the network.

When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) combinations can be lost during route calculation. To override this default behavior for load balancing, include the `lsp-equal-cost` statement to retain the equal cost path information in the routing table.

```
multipath {
    lsp-equal-cost;
}
```

Because the inet.3 routing table is present only on ingress routers, you can configure LSP shortcuts only on these routers.

Default

IS-IS traffic engineering support is enabled.

By default, IS-IS supports traffic engineering by exchanging basic information with the traffic engineering database. To disable this support, and to disable IS-IS shortcuts if they are configured, include the `disable` statement.

Options

credibility-protocol-preference

Specify that IS-IS should use the configured protocol preference for IGP routes to determine the traffic engineering database credibility value. By default, the traffic engineering database prefers IS-IS routes even when the routes of another IGP are configured with a lower, that is, more preferred value. Use this statement to override this default behavior.

The traffic engineering database assigns a credibility value to each IGP and prefers the routes of the IGP with the highest credibility value. In Junos OS Release 9.4 and later, you can configure IS-IS to take protocol preference into account to determine the traffic engineering database credibility value. When protocol preference is used to determine the credibility value, IS-IS routes are not automatically preferred by the traffic engineering database, depending on your configuration. For example, OSPF routes have a default preference value of 10, whereas IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502, whereas IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.

> **NOTE**: This feature is also supported for OSPFv2.

disable

At the `[edit protocols isis traffic-engineering]` hierarchy level, disable IS-IS support for traffic engineering. By default, IS-IS support for traffic engineering is enabled.

ignore-lsp-metrics

Ignore the metrics for RSVP label-switched paths (LSPs) in IS-IS traffic engineering shortcut calculations or when you configure LDP over RSVP LSPs. If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops. Ignoring the metric of RSVP LSPs avoids mutual dependency between IS-IS and RSVP, eliminating the time period when the RSVP metric used for tunneling traffic is not up to date.

ipv6

Configure IS-IS traffic engineering to store IPv6 information in the traffic engineering database (TED) in addition to IPv4 addresses. BGP-LS distributes this information as routes from the TED to the lsdist.0 routing table using the TED import policies. These routes are advertised to BGP-TE peers as network layer reachability information (NLRI) with IPv6 router ID type, length, and value (TLV).

multipath

Enable load balancing for multiple LSPs.

lsp-equal-cost Configure LSPs to be retained as equal cost paths for load balancing when a better route metric is added to the routing table.

When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) path combinations can be lost during route calculation. To override this default IS-IS behavior, include the `lsp-equal-cost` statement for load balancing, so that the equal cost path information is retained in the routing table.

l3-unicast-topology Download IGP topology information into the traffic engineering database (TED). In Junos OS, the IGPs install topology information into a database called the traffic engineering database. The traffic engineering database contains the aggregated topology information. The IGP routes are installed by the traffic engineering database on behalf of the corresponding IGP into a user-visible routing table called lsdist.0, subject to route policies.

> **NOTE**: In earlier Junos releases the `l3-unicast-topology` option was referred to as `igp-topology`. The `igp-topology` syntax is deprecated.

shortcuts Configure IS-IS to use MPLS label-switched paths (LSPs) as next hops if possible when installing routing information into the inet.3 or inet6.3 routing table. Internal gateway protocol (IGP) shortcuts allow the IGP to install prefixes in inet.3 or inet6.3. It is only necessary to enable IGP shortcuts on the ingress router because that is the router performing the shortest-path-first (SPF) calculations.

It is important to understand how IGP shortcuts affect the protocol and routing table relationship. The IGP performs SPF calculations to subnets downstream of LSP egress points, but the results of these calculations are entered into the inet.3 table only. At the same time, the IGP performs its traditional SPF calculations and enters the results of these calculations into the inet.0 table. The result is that although the IGP is making entries into the inet.3 table, BGP is still the only protocol with visibility into that table for the purposes of route resolution. Therefore, forwarding to AS-internal destinations still uses the inet.0 IGP routes, and the LSPs are only used for BGP next-hop resolution. If you want the LSPs to be used for IGP next-hop resolution, you must configure `traffic-engineering bgp-igp`.

multicast-rpf-routes Install unicast IPv4 routes into the multicast routing table (inet.2) for multicast reverse-path-forwarding (RPF) checks. Traffic engineering shortcuts must be enabled. IPv4 multicast topology must not be enabled. Label-switched paths (LSPs) must not be advertised into IS-IS.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `family` statement introduced in Junos OS Release 9.3.

Support for the `credibility-protocol-preference` statement introduced in Junos OS Release 9.4.

Support for the `multipath` statement introduced in Junos OS Release 9.6.

Support for the `lsp-equal-cost` statement introduced in Junos OS Release 9.6.

Support for `inet-mpls` and `inet6-mpls` statements introduced in Junos OS Release 17.2 for MX Series, PTX Series, QFX5100 switches, and QFX10000 line switches.

Support for `inet-mpls` and `inet6-mpls` statements introduced in Junos OS Release 17.3 for QFX5110 and QFX5200 switches.

Support for `igp-topology` statement introduced in Junos OS Release 17.4R1 for MX Series, and PTX Series. The `igp-topology` syntax is now deprecated. Current Junos releases use the `l3-unicast-topology` statement.

`ipv6` option introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Example: Enabling OSPF Traffic Engineering Support

Example: Enabling IS-IS Traffic Engineering Support

Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts

traffic-engineering (Protocols MPLS)

IN THIS SECTION

- Syntax | 30031
- Hierarchy Level | 30031
- Description | 30031
- Default | 30031
- Options | 30031
- Required Privilege Level | 30032
- Release Information | 30032

Syntax

```
traffic-engineering (bgp | bgp-igp | bgp-igp-both-ribs | mpls-forwarding);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Description

Select whether MPLS performs traffic engineering on BGP destinations only or on both BGP and IGP destinations. Affects only LSPs originating from this routing device, not transit or egress LSPs.

Default

`bgp`

Options

`bgp`—On BGP destinations only. Ingress routes are installed in the inet.3 routing table.

`bgp-igp`—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 routing table. If IGP shortcuts are enabled, the shortcut routes are automatically installed in the inet.0 routing table.

`bgp-igp-both-ribs`—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 and inet.3 routing tables. This option is used to support VPNs.

`mpls-forwarding`—On both BGP and IGP destinations. Use ingress routes for forwarding only, not for routing.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Traffic Engineering for LSPs

Configuring MPLS on Provider Edge Switches Using IP-Over-MPLS

traffic-engineering (SR-TE import policy)

IN THIS SECTION

- Syntax | 30033
- Hierarchy Level | 30033
- Description | 30033
- Options | 30033
- Required Privilege Level | 30033
- Release Information | 30033

Syntax

```
traffic-engineering {
         database;
}
```

Hierarchy Level

```
[edit protocols source-packet-routing
        (Segment Routing Traffic Engineering)]
```

Description

Enable export of all SPRING-TE policies into TED.

Options

database Enable export of all Spring-TE policies into the Traffic engineering database.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1.

traffic-load-balance (Traffic Load Balancer)

IN THIS SECTION

- Syntax | 30034
- Hierarchy Level | 30035
- Description | 30035
- Required Privilege Level | 30035
- Release Information | 30036

Syntax

```
traffic-load-balance {
    instance instance-name {
        client-interface client-interface;
        client-vrf client-vrf;
        group group-name {
            health-check-interface-subunit health-check-interface-subunit;
            network-monitoring-profile [profile-name1, <profile-name2>];
            real-service-rejoin-options no-auto-rejoin;
            real-services [server-list];
            <routing-instance routing-instance>;
        }
        interface interface-name;
        real-service real-service  {
            address server-ip-address;
            admin-down;
        }
        server-inet-bypass-filter server-inet-bypass-filter ;
        server-inet6-bypass-filter server-inet6-bypass-filter ;
        server-interface server-interface;
        server-vrf server-vrf;
        traceoptions {
            file file-name <files number> <no-word-readable | world-readable> <size size>;
            flag flag;
            level (all | critical | error | info | notice | verbose | warning);
            monitor {
```

```
                instance-name instance-name;
                virtual-svc-name virtual-service-name;
            }
            no-remote-trace;
        }
        virtual-service virtual-service-name {
            address virtual-ip-address;
            group group-name;
            load-balance-method {
                hash {
                    hash-key method;
                }
                random;
            }
            mode ( layer2-direct-server-return | direct-server-return | translated );
            <routing-instance routing-instance-name>;
            <routing-metric route-metric>;
            server-interface server-interface;
            service service-name {
                protocol (udp | tcp);
                server-listening-port port;
                virtual-port virtual-port;
            }
        }
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure traffic load balancer options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Traffic Load Balancer Overview

Configuring TLB

traffic-manager

IN THIS SECTION

- Syntax (MX Series, PTX Series) Configure Queue Monitoring | **30036**
- Syntax (QFX Series) | **30037**
- Syntax (vSRX Virtual Firewall) | **30038**
- Hierarchy Level | **30038**
- Description | **30038**
- Options | **30038**
- Required Privilege Level | **30040**
- Release Information | **30040**

Syntax (MX Series, PTX Series) Configure Queue Monitoring

```
traffic-manager {
    egress-shaping-overhead number;
    ingress-shaping-overhead number;
    mode {
```

```
        egress-only;
        ingress-and-egress;
        session-shaping;
    }
    enhanced-priority-mode;
    no-enhanced-priority-mode;
    packet-timestamp {
        enable;
    }
    queue-threshold {
        fabric-queue {
            priority high|low {
                threshold threshold-percentage;
            }
        }
        wan-queue {
            priority high|medium-high|medium-low|low {
                threshold threshold-percentage;
            }
        }
    }
}
```

Syntax (QFX Series)

```
traffic-manager {
    buffer-monitor-enable;
    packet-timestamp {
        enable;
    }
    queue-threshold {
        fabric-queue {
            priority high|low {
                threshold threshold-percentage;
            }
        }
        wan-queue {
            priority high|medium-high|medium-low|low {
                threshold threshold-percentage;
            }
        }
```

```
    }
}
```

Syntax (vSRX Virtual Firewall)

```
traffic-manager {
    egress-shaping-overhead number;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number],
[edit chassis fpc slot-number pic pic-number],
[edit chassis fpc slot-number pic pic-number port port-number],
[edit chassis lcc number fpc slot-number pic pic-number]
```

Description

Enable CoS queuing, scheduling, and shaping.

> **NOTE**: Committing changes to `traffic-manager` automatically restarts any necessary components (PICs, DPCs, or FPCs).

Options

buffer-monitor-enable

Enable port buffer monitoring. Buffer utilization data is collected in one-second intervals and compared with the data from the previous interval. The larger value is kept to keep track of peak buffer occupancy for each queue or priority group.

queue-threshold

Enable monitoring of Fabric and WAN queues. When the `fabric-queue` statement is configured, an SNMP trap is generated whenever the fabric power utilization exceeds the configured threshold value.

When `wan-queue` is configured, an SNMP trap is generated whenever the WAN queue depth exceeds the configured threshold value.

egress-shaping-

When traffic management (queueing and scheduling) is configured on the egress side, the number of CoS shaping overhead bytes to add to the packets on the egress interface.

overhead *number*

Replace *number* with a value from -63 through 192 bytes.

For vSRX Virtual Firewall, replace *number* with a value from -62 through 192 bytes.

> **NOTE**: The L2 headers (DA/SA + VLAN tags) are automatically a part of the shaping calculation.

ingress-shaping-overhead *number*

When L2TP session shaping is configured, the number of CoS shaping overhead bytes to add to the packets on the ingress side of the L2TP tunnel to determine the shaped session packet length.

When session shaping is not configured and traffic management (queueing and scheduling) is configured on the ingress side, the number of CoS shaping overhead bytes to add to the packets on the ingress interface.

Replace *number* with a value from –63 through 192 bytes.

mode

Configure CoS traffic manager mode of operation. This option has the following suboptions:

- egress-only—Enable CoS queuing and scheduling on the egress side for the PIC that houses the interface. This is the default mode.

  > **NOTE**: If ingress packet drops are observed at a high rate for a PIC, configure the traffic-manager statement to work in the egress-only mode.

- ingress-and-egress—Enable CoS queueing and scheduling on both the egress and ingress sides for the PIC.

  > **NOTE**:
  >
  > - For EQ DPCs, you must configure the traffic-manager statement with ingress-and-egress mode to enable ingress CoS on the EQ DPC.
  > - EQ DPCs have 250 ms of buffering, with only egress queueing (default mode). When ingress-and-egress is configured, the buffer is partitioned as 50 ms for the ingress direction and 200 ms for the egress direction.

enhanced-priority-mode

Enable the enhanced priority mode. When you enable the enhanced priority mode, the scheduler supports four additional per-priority shaping rates and two additional excess priorities at the interface and interface set level. The four additional per-priority shaping

rates are: Guaranteed Strict-high, Guaranteed Medium-low, Excess medium-high, and Excess medium-low. The two additional excess priorities are: Excess-rate Medium- high and Excess-rate Medium-low. This is the default mode for PTX Series routers.

`no-enhanced-priority-mode`

Disable the enhanced priority mode. This is the default mode for MX Series routers.

> **NOTE**: The line card reboots when you enable or disable the enhanced priority mode feature.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Enabling a Timestamp for Ingress and Egress Queue Packets

show interfaces queue | 21208

traffic-options

IN THIS SECTION

- Syntax | 30041
- Hierarchy Level | 30041
- Description | 30041
- Options | 30041
- Required Privilege Level | 30041
- Release Information | 30041

Syntax

```
traffic-options {
    sessions-per-client {
        limit value;
        over-limit (block | log-and-permit);
    }
}
```

Hierarchy Level

```
[edit security utm utm-policy policy-name]
```

Description

In an attempt to consume all available resources and hinder the ability of the device, a malicious user might generate a large amount of traffic all at once. To prevent such activity from succeeding, you can impose a session throttle.

Starting in Junos OS Release 18.2R1 , we've deprecated the `traffic-options` statement at the `[edit security utm utm-policy policy-name]` hierarchy level.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

traffic-selector

IN THIS SECTION

- Syntax | 30042
- Hierarchy Level | 30042
- Description | 30043
- Options | 30043
- Required Privilege Level | 30044
- Release Information | 30045

Syntax

```
traffic-selector traffic-selector-name {
    local-ip ip-address/netmask;
    remote-ip ip-address/netmask;
    preference pref_value;
    protocol protocol_name/protocol_id;
    source-port low-high;
    destination-port low-high;
    metric metric_value;
    description description_value;
    term term_name {
    local-ip ip-address/netmask;
    remote-ip ip-address/netmask;
    protocol protocol_name/protocol_id;
    source-port low-high;
    destination-port low-high;
                        }
}
```

Hierarchy Level

```
[edit security ipsec vpn vpn-name]
```

Description

A traffic selector is an agreement between IKE peers to permit traffic through a tunnel, if the traffic matches a specified pair of local IP address range, remote IP address range, source port range, destination port range, and protocol. This functionality is supported only for IKEv2.

In the Junos OS Releases earlier to 21.1R1, we support one pair of local IP prefix and remote IP prefix per IPsec tunnel for traffic filtering through IPsec tunnel. From Junos OS Release 21.1R1 onwards, you can configure multiple sets of local IP prefix, remote IP prefix, source port range, destination port range, and protocol for traffic selection.

This means, multiple sets of IP address ranges, port ranges, and protocols can be part of same traffic selector as defined in RFC 7296. In this functionality, concept of term is introduced within the traffic-selectors. Each term defines a set of local IP range, remote IP range, source port range, destination port range, and protocol. All the terms combined will be part of single IPsec SA. The terms in a single traffic selector can have both IPv4 and IPv6 address. Hence a single IPsec SA has both IPv4 and IPv6 as both local and remote IP addresses. A maximum of 200 terms are supported in each traffic selector.

When you configure multiple traffic selectors, each traffic selector leads to a separate negotiation that results in the multiple IPsec tunnels. But, if you configure multiple terms under one traffic selector, this configuration results in single IPsec SA negotiation with multiple IP prefixes, ports, and protocols.

It is mandatory to configure atleast one local IP prefix and one remote IP prefix for a traffic selector. Other parameters are optional.

If multiple traffic selectors have overlapping routes, a tie breaker of routing metric is used for the forwarding decision.

To install the required Junos package for supporting this functionality on your SRX Series Firewall, use the command `request system software add optional://junos-ike.tgz`.

For backward compatibility, we support configuring IP prefixes directly under the [`edit security ipsec vpn` *`vpn-name`* `traffic-selector` *`traffic-selector-name`*] hierarchy.

Use [`edit security ipsec vpn` *`vpn-name`* `traffic-selector` *`traffic-selector-name`* `term` *`term-name`*] hierarchy level to configure multiple sets of IP address ranges, port ranges, and protocols for the same traffic selector as defined in RFC 7296.

You should not configure same values for different traffic selectors for the same IKE gateway. This is not a valid traffic selector configuration. If you configure multiple traffic selectors with the same values, then depending on the peer configuration there might be unintended high CPU utilization.

Options

local-ip ***ip-address/netmask*** A local IP address or a local subnetwork protected by the local VPN device.

remote-ip *ip-address/netmask*

A remote IP address or a remote subnetwork protected by the peer VPN device.

preference *pref_value*

Local preference value of the traffic selector for a particular `ipsec vpn` *vpn-name* that overrides the value specified at global scope.

- **Range:** 0-4294967295.
- **Default:** 5.

term *term_name*

Define a set of local IP range, remote IP range, source port range, destination port range, and protocol. All the terms combined will be part of single IPsec SA. A maximum of 200 terms are supported in each traffic selector. It is optional to configure this parameter.

protocol *protocol_name/protocol_id*

Transport protocol list for a traffic selector for an IPsec tunnel. It is optional to configure this parameter. In case protocol is not configured, then 'any' protocol is assumed to be configured.

- **Range:** Protocol id can range from 0 to 255.

source-port *low-high*

Source port range from lower to higher range port numbers. It is optional to configure this parameter. If no port is configured but only protocol is configured, port 'any' will be assumed for source port ranges for that protocol.

- **Range:** 1 to 65535

destination-port *low-high*

Destination port range from lower to higher range port numbers. It is optional to configure this parameter. If no port is configured but only protocol is configured, port 'any' will be assumed for destination port ranges for that protocol.

- **Range:** 1 to 65535

metric *metric_value*

Tie breaker when multiple traffic selectors have overlapping routes, to decide the most preferred path. It is optional to configure this parameter.

description *description_value*

Traffic selector description. It is optional to configure this parameter. It is optional to configure this parameter.

- **Range:** 0 to 80 characters

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D10.

`term`, `protocol`, `source-port`, `destination-port`, `metric`, and `description` options introduced in Junos OS Release 21.1R1.

`preference` *`pref_value`* option introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

IPsec Overview

vpn (Security) | 30946

ipsec-traffic-selector | 9321

traffic-statistics (Services Analytics)

IN THIS SECTION

- Syntax | 30045
- Hierarchy Level | 30046
- Description | 30046
- Default | 30046
- Options | 30046
- Required Privilege Level | 30047
- Release Information | 30047

Syntax

```
traffic-statistics {
    file filename {
        files number-of-files;
```

```
        size size;
    }
    interval interval;
}
```

Hierarchy Level

```
[edit services analytics]
```

Description

Enable the logging of traffic statistics in a local file. This statement does not enable traffic statistics monitoring.

To enable the monitoring of traffic statistics, configure the `traffic-statistics` configuration statement at the `[edit services analytics interfaces]` hierarchy level.

Default

This feature is disabled by default.

Options

file *filename*

Specify a filename for storing the traffic statistics in the JavaScript Object Notification (JSON) format. The file is stored in the **/var/log/** directory of your device.

If you do not specify a filename, the data is not stored in a file.

files *number-of-files*

Specify the number of files to store locally. After the number files with the maximum file size is reached, the system starts over and writes the data to the first file.

- **Range:** 2 to 1,000 files.

interval *interval*

Configure the polling interval in seconds.

> **NOTE**: You can configure the polling interval for traffic statistics globally for all interfaces only. Due to limitations and variations in the hardware capability of different devices, you might see a difference in value between the actual interval and configured interval.

- Range:
  - Junos OS Release 13.2X51-D10 or later—2 to 300 seconds (2 seconds to 5 minutes)
  - Junos OS Release 13.2X50-D15—1 to 300 seconds (1 second to 5 minutes)

> **NOTE**: In Junos OS Release 13.2X51-D10 or later, if you configured an interval of less than 2 seconds, the following warning messages appear during the commit process:
>
> `Traffic statistics polling interval can not be less than 2 seconds`, and
>
> `Setting Traffic statistics polling interval to 2 seconds`.
>
> These messages do not stop the commit operation, but the interval is automatically set to 2 seconds.

- Default:
  - Junos OS Release 13.2X50-D15—1 second
  - Junos OS Release 13.2X51-D10 or later—2 seconds

size *size* Configure the file size in megabytes (MB).

- **Syntax:** `xm` to specify MB.
- **Range:** 10 to 4095 MB

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Network Analytics Overview

traffic-statistics (Protocols IS-IS)

IN THIS SECTION

- Syntax | 30048
- Hierarchy Level | 30048
- Description | 30048
- Options | 30049
- Required Privilege Level | 30049
- Release Information | 30049

Syntax

```
traffic-statistics auto-bandwidth auto-bandwidth statistics-granularity per-interface per-
interface;
```

Hierarchy Level

```
[edit logical-systems name protocols isis source-packet-routing],
[edit logical-systems name routing-instances name protocols isis source-packet-routing],
[edit protocols isis source-packet-routing],
[edit routing-instances name protocols isis source-packet-routing]
```

Description

Enable traffic measuring on the SPRING link. RSVP-TE needs to measure the bandwidth utilized by SPRING and estimate the available bandwidth at its disposal for traffic engineering. Traffic statistics are collected at a configured interval. The average of the measured SPRING bandwidth is reported to RSVP only when a set threshold is breached. RSVP then passes on this information to the IGP for reallocation of bandwidth.

Options

auto-bandwidth ***auto-bandwidth***	Specify the name of auto bandwidth.
statistics-granularity ***per-interface***	Specify the interface where you need to measure traffic statistics of SPRING traffic.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

auto-bandwidth | 2682

update-threshold-max-reservable | 30458

traffic-statistics-labeled-path

IN THIS SECTION

- Syntax | 30050
- Hierarchy Level | 30050
- Description | 30050
- Options | 30050
- Required Privilege Level | 30051
- Release Information | 30051

Syntax

```
traffic-statistics-labeled-path {
    file filename <files files>     <size size> <(world-readable | no-world-readable)>;
    interval seconds;
}
```

Hierarchy Level

```
[edit logical-systems name protocols bgp],
[edit logical-systems name routing-instances name protocols bgp],
[edit protocols bgp],
[edit routing-instances name protocols bgp]
```

Description

Enable collection of periodic ingress labeled statistics for BGP label-switched paths in a network configured with segment routing. The traffic statistics are collected at a specified time interval and saved in a specified file. To collect traffic statistics of labeled unicast IPv4 and IPv6 families for specific BGP groups, configure the `labeled-path` option under the family or BGP group level at `[edit logical-systems logical-system-name protocols bgp group group-name family (inet | inet6) labeled-unicast]` hierarchy level.

Options

file *filename*
Specify a filename to collect traffic statistics.

interval
Specify a time interval in seconds to collect traffic statistics.

- **Range:** `60` through `65535` seconds.

size
(Optional) Specify the size of the file that records traffic statistics.

- **Range:** `10240` through `4294967295`

world readable | world unreadable
(Optional) Specify whether the traffic statistics log file is accessible to everyone or not.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Enabling Traffic Statistics Collection for BGP Labeled Unicast

show bgp group traffic-statistics | 17171

traffic-statistics (Protocols BGP)

IN THIS SECTION

- Syntax | 30051
- Hierarchy Level | 30052
- Description | 30052
- Options | 30052
- Required Privilege Level | 30052
- Release Information | 30052

Syntax

```
traffic-statistics {
    labeled-path
    file filename <world-readable | no-world-readable>;
    interval seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family (inet | inet6) labeled-unicast],
[edit logical-systems logical-system-name protocols bgp group group-name family (inet | inet6)
labeled-unicast],
[edit protocols bgp family (inet | inet6) labeled-unicast],
[edit protocols bgp group group-name family (inet | inet6) labeled-unicast]
```

Description

Enable the collection of traffic statistics for interprovider or carrier-of-carriers VPNs.

Options

`file` *`filename`*—Specify a filename for the BGP labeled–unicast traffic statistics file. If you do not specify a filename, statistics are still collected but can only be viewed by using the `show bgp group traffic statistics` *`group-name`* command.

`interval` *`seconds`*—Specify how often BGP labeled-unicast traffic statistics are collected.

`labeled-path`—Specify this option to collect labeled path traffic statistics of ingress BGP nodes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`labeled-path` introduced in Junos OS Release 18.1R1 for the MX Series.

RELATED DOCUMENTATION

- Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics
- MPLS Feature Support on QFX Series and EX4600 Switches
- *Understanding Interprovider and Carrier-of-Carriers VPNs*

traffic-statistics (Protocols LDP)

IN THIS SECTION

- Syntax | 30053
- Hierarchy Level | 30053
- Description | 30053
- Options | 30053
- Required Privilege Level | 30054
- Release Information | 30054

Syntax

```
traffic-statistics {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    interval seconds;
    no-penultimate-hop;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Description

LDP traffic statistics display the amount of traffic passed through a router for a particular FEC.

Options

`file` *`filename`*—Name of the file to receive the output of the LDP statistics operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of LDP statistics files. When a statistics file named *`ldp-stat`* reaches its maximum size, it is renamed *`ldp-stat`*`.0`, then *`ldp-stat`*`.1`, and so on, until the maximum number of LDP statistics files is reached. Then the oldest file is overwritten.

- **Range:** 2 through 1000
- **Default:** 2 files

If you specify a maximum number of files, you also must include the `size` statement to specify the maximum file size.

`interval` *`seconds`*—(Optional) Specify the interval at which the statistics are polled and written to the file.

- **Default:** 300 seconds (5 minutes)

`no-penultimate-hop`—(Optional) Do not collect traffic statistics on the penultimate hop router.

`no-world-readable`—(Optional) Prevent all users from reading the log file.

`size` *`size`*—(Optional) Maximum size of each statistics file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a statistics file named *`ldp-stat`* reaches this size, it is renamed *`ldp-stat`*`.0`. When *`ldp-stat`* again reaches this size, *`ldp-stat`*`.0` is renamed *`ldp-stat`*`.1` and *`ldp-stat`* is renamed *`ldp-stat`*`.0`. This renaming scheme continues until the maximum number of statistics files is reached. Then the oldest statistics file is overwritten.

- **Syntax:** *`x`*`k` to specify KB, *`x`*`m` to specify MB, or *`x`*`g` to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

If you specify a maximum file size, you also must also include the `files` statement to specify the maximum number of files.

`world-readable`—(Optional) Enable log file access for all users.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Collecting LDP Statistics

transaction-timeout

IN THIS SECTION

- Syntax | 30055
- Hierarchy Level | 30055
- Description | 30055
- Options | 30056
- Required Privilege Level | 30056
- Release Information | 30056

Syntax

```
transaction-timeout value-in-seconds;
```

Hierarchy Level

```
[edit logical-systems name security alg mgcp],
[edit logical-systems name tenants name security alg mgcp],
[edit security alg mgcp],
[edit services alg mgcp],
[edit tenants name security alg mgcp]
```

Description

Specify the timeout value for Media Gateway Control Protocol (MGCP) transactions. If the timeout value exceeds, the MGCP transactions stops.

Options

value-in-seconds—Timeout value.

- **Range:** 3 through 50 seconds
- **Default:** 30 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

mgcp | 10693

transfer

IN THIS SECTION

- Syntax | 30057
- Hierarchy Level | 30057
- Description | 30057
- Options | 30057
- Required Privilege Level | 30057
- Release Information | 30057

Syntax

```
transfer {
    record-level number;
    timeout seconds;
}
```

Hierarchy Level

```
[edit services flow-collector file-specification variant variant-number]
```

Description

Specify when to send the flow collection file. The file is sent when either of the two conditions is met.

Options

`record-level` *`number`*—Use the specified number of flow collection files collected.

`timeout` *`seconds`*—Use the specified timeout duration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring File Formats

transfer-delay (Event Policy)

IN THIS SECTION

- Syntax | 30058
- Hierarchy Level | 30058
- Description | 30058
- Default | 30059
- Options | 30059
- Required Privilege Level | 30059
- Release Information | 30059

Syntax

```
transfer-delay seconds;
```

Hierarchy Level

```
[edit event-options destinations destination-name],
[edit event-options policy policy-name then event-script filename destination destination-name],
[edit event-options policy policy-name then execute-commands destination destination-name],
[edit event-options policy policy-name then upload filename (filename | committed)
destination destination-name]
```

Description

Configure a delay before transferring files. This allows the files to be completely generated before the upload starts. If you configure a transfer delay at the `[edit event-options destination destination-name]` hierarchy level and at one of the `[edit event-options policy policy-name then ...]` hierarchy levels, the resulting delay is the sum of the two delays.

Default

If you do not include this statement, there is no transfer delay.

Options

seconds Duration of the delay before files are uploaded.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

Support at the `[edit event-options policy policy-name then ...]` hierarchy levels introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configure the Delay Before Files Are Uploaded

Define Destinations for Event Policy File Archiving

transfer-interval

IN THIS SECTION

- Syntax | 30060
- Hierarchy Level | 30060
- Description | 30060
- Options | 30060
- Required Privilege Level | 30060
- Release Information | 30060

Syntax

```
transfer-interval minutes;
```

Hierarchy Level

```
[edit accounting-options file filename]
```

Description

Specify the length of time the file remains open and receives new statistics before it is closed and transferred to an archive site.

Options

minutes—Time the file remains open and receives new statistics before it is closed and transferred to an archive site.

- **Range:** 5 through 2880 minutes
- **Default:** 30 minutes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Accounting Options, Source Class Usage and Destination Class Usage Options

transfer-log-archive

IN THIS SECTION

- Syntax | 30061
- Hierarchy Level | 30061
- Description | 30061
- Required Privilege Level | 30062
- Release Information | 30062

Syntax

```
transfer-log-archive {
    archive-sites {
        ftp:url {
            password "password";
            username username;
        }
    }
    filename-prefix prefix;
    maximum-age minutes;
}
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure the filename prefix, maximum age, and destination FTP server for log files containing the transfer activity history for a flow collector interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Transfer Logs

transient-filter

IN THIS SECTION

- Syntax | 30062
- Hierarchy Level | 30062
- Description | 30063
- Required Privilege Level | 30063
- Release Information | 30063

Syntax

```
transient-filter transient filter name
```

Hierarchy Level

```
[edit firewall family firewall family name filter parent filter name]
```

Description

Configure a transient firewall filter for a parent firewall filter.

When the parent firewall filter is modified, the make-before-break (MBB) process of the system reprograms the modified parent firewall filter in the hardware. Depending on the size of the parent firewall filter, this process may take time. In the interim, the transient firewall filter is used in lieu of the parent firewall filter. The transient filter is expected to have a smaller memory footprint. After the MBB operation completes, the parent firewall filter resumes operation.

A transient firewall filter can only be paired with one parent firewall filter. The transient firewall filter is implicitly bound to the same interface as its parent firewall filter. A transient firewall filter cannot be configured as an interface specific firewall filter. Transient firewall filters can also be configured for fast-lookup parent firewall filters.

A transient firewall filter should be configured under the same firewall filter family as its parent firewall filter. During the modification of a parent firewall filter, any firewall filter statistics that is generated, will be attributed to the transient firewall filter. After the modification completes, the firewall filter statistics are retained and made available to the modified parent firewall filter when it resumes operation.

Deletion of the parent firewall filter will not delete its transient firewall filter. Instead, the transient firewall filter is dissociated with the parent firewall filter, and treated as any regular firewall filter, and its counters will be reset. A transient firewall filter cannot be deleted if it is associated with a parent firewall filter. A parent firewall filter with a transient firewall filter can be bound on ports, VLANs, and layer 3 interfaces in ingress or egress directions.

If a firewall filter is part of a filter-chain, input-list or output-list, it cannot be made a transient firewall filter. However, the parent firewall filter can be part of a firewall filter chain, input-list or output-list. Transient firewall filters are supported on loopback interfaces in the ingress direction. Transient filters are not supported on management interfaces. A transient firewall filter can only be a non-interface specific filter, whereas its parent firewall filter can be interface specific or non-interface specific firewall filter. If firewall filters of an input-list or output-list have associated transient firewall filters, then a consolidated transient firewall filter is created.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 24.1R1

transit

IN THIS SECTION

- Syntax | 30064
- Hierarchy Level | 30064
- Description | 30065
- Options | 30065
- Required Privilege Level | 30065
- Release Information | 30065

Syntax

```
transit incoming-label {
    bandwidth bps;
    description string;
    link-protection bypass-name name;
    member-interface member-interface;
    next-hop (address | interface-name | address/interface-name);
    node-protection bypass-name name next-next-label label;
    pop;
    stitch {
        bandwidth bps;
        description string;
        link-protection bypass-name name;
        next-hop (address | interface-name | address/interface-name);
        node-protection bypass-name name next-next-label label;
    }
    swap out-label;
}
```

Hierarchy Level

```
[edit logical-systems name protocols mpls static-label-switched-path name],
[edit logical-systems name routing-instances name protocols  mpls static-label-switched-path name],
```

```
[edit logical-systems name tenants name routing-instances name protocols mpls static-label-switched-path name],
[edit protocols mpls static-label-switched-path name],
[edit routing-instances name protocols mpls static-label-switched-path name],
[edit tenants name routing-instances name protocols mpls static-label-switched-path name]
```

Description

Configure a transit static LSP.

> **NOTE**: When configuring transit static LSPs with label operation as `stitch`, the configured next-hop can only be a valid IP address and not an interface name.

Options

incoming-label Incoming label value.

- **Values:**
  - Default: 1000000 through 1048575
  - Range: 16 through 1048575

member-interface Aggregated Ethernet (AE) member interface name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`switch` option introduced in Junos OS Release 14.1X53-D25.

`member-interface` option introduced in Junos OS Release 17.4R1 for the MX Series and PTX 5000.

RELATED DOCUMENTATION

Configuring Static LSPs

Configuring Static Adjacency Segment Identifier for Aggregate Ethernet Member Links Using Single-Hop Static LSP

transit (Chained Composite Next Hops)

IN THIS SECTION

- Syntax | 30066
- Hierarchy Level | 30067
- Description | 30067
- Default | 30067
- Options | 30067
- Required Privilege Level | 30068
- Release Information | 30068

Syntax

```
transit {
    (all | no-all);
    (l2vpn | no-l2vpn);
    (l3vpn | no-l3vpn);
    (labeled-bgp | no-labeled-bgp);
    (ldp | no-ldp);
    (ldp-p2mp | no-ldp-p2mp);
    lsp-statistics-from-route;
    (rsvp | no-rsvp);
    (rsvp-p2mp | no-rsvp-p2mp);
    (static | no-static);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-next-hop],
[edit routing-options forwarding-table chained-composite-next-hop]
```

NOTE: The `[edit logical-systems]` hierarchy level is not supported on the QFX10000 switches.

Description

Allows you to configure the chained composite next hops transit configuration options for devices handling transit traffic in the network. This statement and the associated functionality is available only on PTX Packet Transport Routers, QFX10000 switches, and MX204, MX480, MX960, MX10003, and VMX routers.

Default

All of the `transit` statement options are enabled on PTX transport routers and QFX10000 switches. However, you can disable any of the statements with a **no** option.

Starting in Junos OS Release 14.1, the `transit l3vpn` statement is enabled by default on PTX Series Packet Transport Routers only.

Options

all | no-all Enable or disable chained composite next-hops for all of the possible packet transit protocols and applications. The `all | no-all` statements do not apply to the `lsp-statistics-from-route` statement.

l2vpn | no-l2vpn Enable or disable chained composite next-hops for Layer 2 VPNs.

l3vpn | no-l3vpn Enable or disable chained composite next-hops for Layer 3 VPNs.

labeled-bgp | no-labeled-bgp Enable or disable chained composite next-hops for labeled BGP.

ldp | no-ldp Enable or disable chained composite next-hops for LDP.

ldp-p2mp | no-ldp-p2mp Enable or disable chained composite next-hops for LDP-signaled P2MP LSPs.

> **NOTE**: The `ldp-p2mp` and `rsvp-p2mp` statements are not supported on MX series routers.
>
> On an MX series router with redundant Routing Engines and enhanced-ip mode configuration, enabling the `ldp-p2mp` and `rsvp-p2mp` statements under the `[edit routing-options forwarding-table chained-composite-next-hop transit]` hierarchy level causes ping from the current primary logical system to fail at the time of a Routing Engine switchover.

lsp-statistics-from-route Enable LSP statistics collection from the route.

rsvp | no-rsvp Enable or disable chained composite next-hops for RSVP.

rsvp-p2mp | no-rsvp-p2mp Enable or disable chained composite next-hops for RSVP-signaled P2MP LSPs.

> **NOTE**: The `ldp-p2mp` and `rsvp-p2mp` statements are not supported on MX series routers.
>
> On an MX series router with redundant Routing Engines and enhanced-ip mode configuration, enabling the `rsvp-p2mp` and `ldp-p2mp` statements under the `[edit routing-options forwarding-table chained-composite-next-hop transit]` hierarchy level causes ping from the current primary logical system to fail at the time of a Routing Engine switchover.

static | no-static Chained composite next hops are enabled for transit static LSPs by default. You can also disable this functionality for transit static LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs

chained-composite-next-hop | 3277

transit-lsp-association

IN THIS SECTION

- Syntax | 30069
- Hierarchy Level | 30069
- Description | 30069
- Options | 30070
- Required Privilege Level | 30070
- Release Information | 30070

Syntax

```
transit-lsp-association transit-association-lsp-group-name {
    from-1     address-of-associated-lsp-1;
    from-2     address-of-associated-lsp-2;
    lsp-name-1 name-of-associated-lsp-1;
    lsp-name-2 name-of-associated-lsp-2;
}
```

Hierarchy Level

```
[edit protocols mpls]
```

Description

Associate two label-swiched paths (LSPs) at a transit node to configure a path for sending and receiving GAL and G-Ach messages for MPLS-TP OAM.

Options

transit-association-lsp-group-name	Name of the transit association LSP group.
from-1 *address-of-associated-lsp-1*	Address of the first associated LSP.
from-2 *address-of-associated-lsp-2*	Address of the second associated LSP.
lsp-name-1 *name-of-associated-lsp-1*	Name of the first associated LSP.
lsp-name-2 *name-of-associated-lsp-1*	Name of the second associated LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring the MPLS Transport Profile for OAM

translated

IN THIS SECTION

- Syntax | 30071
- Hierarchy Level | 30071
- Description | 30071
- Options | 30071
- Required Privilege Level | 30071
- Release Information | 30072

Syntax

```
translated {
    address-pooling paired;
    clat-prefix clat-prefix;
    cclat-ipv6-prefix-length
    destination-pool nat-pool-name;
    destination-prefix destination-prefix;
    dns-alg-pool dns-alg-pool;
    dns-alg-prefix dns-alg-prefix;
    filtering-type endpoint-independent;
    mapping-type endpoint-independent;
    overload-pool overload-pool-name;
    overload-prefix;
    source-pool nat-pool-name;
    translation-type (basic-nat-pt | basic-nat44 | basic-nat66 |deterministic-napt44 |
deterministic-napt64 | dnat-44 | dynamic-nat44 | napt-44 | napt-66 | napt-pt | stateful-nat464 |
stateful-nat64 | twice-basic-nat-44 | twice-dynamic-nat-44 | twice-napt-44)

    }
}
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then]
```

Description

Define properties for translated traffic.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

translated-port

IN THIS SECTION

- Syntax | 30072
- Hierarchy Level | 30072
- Description | 30072
- Options | 30073
- Required Privilege Level | 30073
- Release Information | 30073

Syntax

```
translated-port port id;
```

Hierarchy Level

```
[edit services nat port-forwarding map-name]
```

Description

Specify the port to which all traffic will be translated.

The `translated-port` statement is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, the `translated-port` statement is also supported on the MS-MPC and MS-MIC.

Options

port id The port number to which traffic will be translated.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring Port Forwarding for Static Destination Address Translation

Configuring Port Forwarding Without Destination Address Translation

translation

IN THIS SECTION

- Syntax | 30074
- Hierarchy Level | 30074
- Description | 30074
- Options | 30074
- Required Privilege Level | 30074
- Release Information | 30074

Syntax

```
translation {
    max-datasize;
}
```

Hierarchy Level

```
[edit system scripts]
```

Description

Configure options for translation scripts.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

> *Storing and Enabling Scripts*

translation-table

IN THIS SECTION

- Syntax | 30075
- Hierarchy Level | 30075
- Description | 30075
- Default | 30076
- Options | 30076
- Required Privilege Level | 30076
- Release Information | 30076

Syntax

```
translation-table {
    (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp | to-inet-precedence-from-
inet-precedence) table-name {
        to-code-point value from-code-points (* | [ values ]);
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

For an Enhanced IQ PIC or Multiservices PIC, specify the input translation tables. You must also apply the translation table to a logical interface on the Enhanced IQ PIC or Multiservices PIC.

Default

If you do not include this statement, the ToS bit values in received packet headers are not changed by the PIC.

Options

`to-dscp-from-dscp`—(Optional) Translate incoming IPv4 DSCP values to new values. You must also configure and apply a DSCP classifier.

`to-dscp-ipv6-from-dscp-ipv6`—(Optional) Translate incoming IPv6 DSCP values to new values. You must also configure and apply an IPv6 DSCP classifier.

`to-inet-precedence-from-inet-precedence`—(Optional) Translate incoming INET precedence values to new values.

`to-exp-from-exp`—(Optional) Translate incoming MPLS EXP values to new values.

`table-name`—The name of the translation table.

`value`—The bit string to which to translate the incoming bit value.

`value(s)`—The bit string(s) from which the incoming bit value(s) are translated.

`*`—(Optional) This translation matches all bit patterns not explicitly listed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Support on Multiservices PIC added in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring ToS Translation Tables

Multiservices PIC ToS Translation

translation-type

IN THIS SECTION

- Syntax | 30077
- Hierarchy Level | 30077
- Description | 30077
- Options | 30077
- Required Privilege Level | 30079
- Release Information | 30079

Syntax

```
translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | nat-44 | deterministic-napt44 |
deterministic-napt64 | dnat-44 | dynamic-nat44 | napt-44 | napt-66 | napt-pt | nptv6 | stateful-
nat464 | stateful-nat64 | twice-basic-nat-44 | twice-dynamic-nat-44 | twice-napt-44)
```

Hierarchy Level

```
[edit services nat rule rule-name term term-name then translated]
```

Description

Specify the NAT translation types. To identify the interface cards and Junos OS releases that support each translations type, see *Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card*.

Options

basic-nat44 Translate the source address statically (IPv4 to IPv4).

basic-nat66 Translate the source address statically (IPv6 to IPv6).

basic-nat-pt Translate the addresses of IPv6 hosts as they originate sessions to the IPv4 hosts in the external domain. The `basic-nat-pt` option is always implemented with DNS ALG.

deterministic-napt44 Translate as deterministic NAPT44.

deterministic-napt64 Translate as deterministic NAPT64.

dnat-44 Translate the destination address statically (IPv4 to IPv4).

dynamic-nat44 Translate only the source address by dynamically choosing the NAT address from the source address pool.

napt-44 Translate the transport identifier of the IPv4 private network to a single IPv4 external address.

napt-66 Translate the transport identifier of the IPv6 private network to a single IPv6 external address.

napt-pt Bind addresses in an IPv6 network with addresses in an IPv4 network and vice versa to provide transparent routing for the datagrams traversing between the address realms.

nptv6 Translate the source address prefix in a stateless manner (IPv6 to IPv6).

stateful-nat464 Implement 464XLAT Provider-Side Translater (PLAT) address translation for source IP addresses and IPv6 prefix removal translation for destination IPv4 addresses.

stateful-nat64 Implement dynamic address and port translation for source IP addresses (IPv6-to-IPv4) and prefix removal translation for the destination IP addresses (IPv6-to-IPv4).

twice-basic-nat-44 Translate the source and destination addresses statically (IPv4 to IPv4).

> **NOTE**: Starting with Junos OS Release 15.1R1, the twice NAT functionality (`twice-basic-nat-44`, `twice-dynamic-nat-44`, and `twice-dynamic-napt-44` options) is supported on MX Series routers with MS-MPCs and MS-MICs.

twice-dynamic-nat-44 Translate the source address by dynamically choosing the NAT address from the source address pool. Translate the destination address statically.

twice-dynamic-napt-44 Translate the transport identifier of the IPv4 private network to a single IPv4 external address. Translate the destination address statically.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The following options introduced in Junos OS Release 11.2, replacing all previous options.

- `basic-nat44`
- `basic-nat66`
- `basic-nat-pt`
- `dnat-44`
- `dynamic-nat44`
- `napt-44`
- `napt-66`
- `napt-pt`
- `stateful-nat64`

`deterministic-napt44` option introduced in Junos OS Release 12.1.

`deterministic-napt64` option introduced in Junos OS Release 17.4R1.

`nptv6` option introduced in Junos OS Release 15.1

`stateful-nat464` option introduced in Junos OS Release 17.1R1.

`twice-basic-nat-44` option introduced in Junos OS Release 11.4.

`twice-dynamic-nat-44` option introduced in Junos OS Release 11.4.

`twice-napt-44` option introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Network Address Translation Rules Overview

translation-vni

IN THIS SECTION

- Syntax | 30080
- Hierarchy Level | 30080
- Description | 30080
- Options | 30082
- Required Privilege Level | 30082
- Release Information | 30082

Syntax

```
translation-vni vni (1..16777214);
```

Hierarchy Level

```
[edit routing-instances routing-instance-name vlans vlan-name vxlan]
[edit routing-instances routing-instance-name vxlan]
[edit vlans vlan-name vxlan]
```

Description

Configure a VXLAN network identifier (VNI) to which a Juniper Networks switch will translate a VNI in a packet that it is forwarding. The Juniper Networks switch translates the VNI while it is forwarding traffic in the following seamless EVPN-VXLAN stitching use cases:

- From one EVPN-VXLAN point of delivery (POD) to another within a data center.

- From one EVPN-VXLAN data center to another. This use case is also known as data center interconnect (DCI).

In these use cases, a Juniper Networks device that supports the seamless EVPN-VXLAN stitching feature can serve as either a spine or super spine device that interconnects the PODs or data centers through an EVPN-VXLAN WAN network.

When configuring the interconnection, you can set up a single routing instance of type `virtual-switch` or `evpn` on each spine or super spine device. Or, you can use the default switching instance.

The type of instance that you use determines the hierarchy level at which you must include the `translation-vni` configuration statement. Table 1829 on page 30081 outlines the instance types and the hierarchy levels at which you must include `translation-vni` for each type.

Table 1829: Interconnection Instance Types and Hierarchy Levels at Which to Include translation-vni Configuration Statement

Instance Types	Hierarchy Levels
Routing instance of type `virtual-switch`	`[edit routing-instances` *`routing-instance-name`* `vlans` *`vlan-name`* `vxlan]`
Routing instance of type `evpn`	`[edit routing-instances` *`routing-instance-name`* `vxlan]`
Default switching instance	`[edit vlans` *`vlan-name`* `vxlan]`

The Juniper Networks switch translates the VNI only if you have included the translated VNI in the interconnected VNI list. For information about configuring this list, see "interconnected-vni-list" on page 8809.

To understand how VNI translation works, Data Centers 1 and 2 are interconnected through an EVPN-VXLAN WAN network. In this topology, Super Spine 1 connects Data Center 1 and the WAN network, and Super Spine 2 connects Data Center 2 and the WAN network.

A routing instance named `evpn-vxlan` is configured on both Super Spines 1 and 2. In this routing instance, VLAN 100 is configured on each super spine device. However, the VNI mapped to VLAN 100 is different on each super spine device.

Table 1830 on page 30082 provides a summary of the VNI-related attributes configured on each super spine device.

Table 1830: Summary of VNI-Related Configurations on Super Spines 1 and 2

VNI-Related Configurable Attributes	Super Spine 1 evpn-vxlan Routing Instance Configuration	Super Spine 2 evpn-vxlan Routing Instance Configuration
VLAN ID	100	100
VNI	100	300
Translated VNI	200	200
Interconnected VNI List	Includes VNI 200	Includes VNI 200

As shown in Table 1830 on page 30082, Super Spine 1 maps VNI 100 to VLAN 100, and Super Spine 2 maps VNI 300 to VLAN 100. Both super spine devices include VNI 200 in their interconnected VNI lists. As a result, when the super spine devices forward VLAN 100 traffic to the WAN network, the VNI is translated to 200.

Options

vni Numeric value that represents the VNI to which the VNI in a packet is translated.

- **Range:** 1through 16777214

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

interconnect | 8797

transmit-clock

IN THIS SECTION

- Syntax | 30083
- Hierarchy Level | 30083
- Description | 30083
- Options | 30083
- Required Privilege Level | 30083
- Release Information | 30084

Syntax

```
transmit-clock invert;
```

Hierarchy Level

```
[edit interfaces interface-name serial-options]
```

Description

Configure the transmit clock signal.

Options

`invert`—Shift the clock phase 180 degrees.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Serial Clocking Mode

transmit-failure-threshold (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 30084
- Hierarchy Level | 30084
- Description | 30084
- Required Privilege Level | 30085
- Release Information | 30085

Syntax

```
transmit-failure-threshold;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specifies the failure threshold value of the transmit test frames.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

RFC 2544-Based Benchmarking Tests Overview

Configuring RFC 2544-Based Benchmarking Tests

rfc2544-benchmarking | **15759**

transmit-interval

IN THIS SECTION

- Syntax | **30085**
- Hierarchy Level | **30086**
- Description | **30086**
- Required Privilege Level | **30086**
- Release Information | **30086**

Syntax

```
transmit-interval {
    threshold milliseconds;
    minimum-interval milliseconds;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options
dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group group-name liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd]
```

Description

Configure the Bidirectional Forwarding Detection (BFD) transmit interval.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients
- Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

transmit-interval (BFD Liveness Detection)

IN THIS SECTION

- Syntax | 30087
- BGP | 30087
- EVPN, L2VPN, VPLS | 30088
- Description | 30088
- Options | 30089
- Required Privilege Level | 30089
- Release Information | 30089

Syntax

```
transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
}
```

BGP

```
[edit logical-systems name protocols bgp bfd-liveness-detection],
[edit logical-systems name protocols bgp group bfd-liveness-detection],
[edit logical-systems name protocols bgp group name neighbor address bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols bgp bfd-liveness-detection],
[edit logical-systems name routing-instances name  protocols bgpgroup bfd-liveness-detection],
[edit logical-systems name routing-instances name  protocols bgpgroup neighbor address bfd-liveness-
detection],
[edit logical-systems name tenants name routing-instances name protocols bgp bfd-liveness-
detection],
[edit logical-systems name tenants name routing-instances name protocols bgpgroup bfd-liveness-
detection],
[edit logical-systems name tenants name routing-instances name protocols bgpgroup neighbor
address bfd-liveness-detection],
[edit protocols bgp bfd-liveness-detection],
```

```
[edit protocols bgp group bfd-liveness-detection],
[edit protocols bgp group neighbor address bfd-liveness-detection],
[edit routing-instances name protocols bgp bfd-liveness-detection],
[edit routing-instances name protocols bgp group bfd-liveness-detection],
[edit routing-instances name protocols bgp group neighbor address bfd-liveness-detection],
[edit tenants name routing-instances name protocols bgp bfd-liveness-detection]
[edit tenants name routing-instances name protocols bgp group bfd-liveness-detection]
[edit tenants name routing-instances name protocols bgp groupneighbor address bfd-liveness-detection]
```

EVPN, L2VPN, VPLS

```
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls)  oam bfd-liveness-
detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-
id oam bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-
group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)oam
bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn |
vpls)neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-
group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-
detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-name neighbor
neighbor-id oam bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-
detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam
bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-
name neighbor neighbor-id oam bfd-liveness-detection],
```

Description

Specify the transmit interval for the `bfd-liveness-detection` statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum time that it requires between packets sent from its peer; the receive interval is not negotiated

between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

Options

minimum-interval *milliseconds*

Configure the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement at this hierarchy level, you can configure the minimum transmit interval using the `minimum-interval` statement at the `bfd-liveness-detection` hierarchy level.

> **NOTE**: The threshold value specified in the `threshold` statement must be greater than the value specified in the `minimum-interval` statement for the `transmit-interval` statement.

- **Range:** 1 through 255,000 milliseconds

threshold *milliseconds*

Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.

- **Range:** 0 through 4,294,967,295 (2^{32} – 1)

> **NOTE**: The threshold value specified in the `threshold` statement must be greater than the value specified in the `minimum-interval` statement for the `transmit-interval` statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Support for BFD authentication introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD for Layer 2 VPN and VPLS

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

bfd-liveness-detection (BGP)

bfd-liveness-detection (Layer 2 VPN and VPLS) | **2929**

transmit-interval (LAG)

IN THIS SECTION

- Syntax | **30090**
- Hierarchy Level | **30090**
- Description | **30091**
- Options | **30091**
- Required Privilege Level | **30091**
- Release Information | **30091**

Syntax

```
transmit-interval {
    minimum-interval  milliseconds;
    threshold  milliseconds;
}
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options bfd-liveness-detection]
```

Description

Configure the minimum interval and the threshold for transmission of BFD packets for aggregated Ethernet interfaces.

Options

minimum-interval *milliseconds*

Specify the minimum time interval between two transmissions of packets.

- **Range:** 1 through 255,000

threshold *milliseconds*

Specify the maximum interval between transmission of packets. If the transmit interval is greater than this value, the device triggers a trap.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

authentication (LAG) | **2622**

bfd-liveness-detection (LAG) | **2926**

detection-time (LAG) | **5549**

Configuring Micro BFD Sessions for LAG

Example: Configuring Independent Micro BFD Sessions for LAG

Understanding Independent Micro BFD Sessions for LAG

transmit-interval (Liveness Detection)

IN THIS SECTION

- Syntax | 30092
- Hierarchy Level | 30092
- Description | 30092
- Required Privilege Level | 30093
- Release Information | 30093

Syntax

```
transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
}
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection]
```

Description

Configure the Bidirectional Forwarding Detection (BFD) transmit interval. The negotiated transmit interval for a peer is the interval between the sending of BFD liveness detection requests to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

transmit-interval (MACsec)

IN THIS SECTION

- Syntax | 30093
- Hierarchy Level | 30093
- Description | 30094
- Default | 30094
- Options | 30094
- Required Privilege Level | 30094
- Release Information | 30094

Syntax

```
transmit-interval interval;
```

Hierarchy Level

```
[edit security macsec connectivity-association mka]
```

Description

Specifies the transmit interval for MACsec Key Agreement (MKA) protocol data units (PDUs).

The MKA transmit interval setting sets the frequency for how often the MKA PDU is sent to the directly connected device to maintain MACsec on a point-to-point Ethernet link. A lower *interval* increases bandwidth overhead on the link; a higher *interval* optimizes the MKA protocol data unit exchange process.

The transmit interval settings must be identical on both ends of the link when MACsec using static connectivity association key (CAK) security mode is enabled.

We recommend increasing the interval to 6000 ms in high-traffic load environments.

Default

The default transmit interval is 10000 milliseconds (10 seconds).

NOTE: Configuring aggressive transmit interval will lead to broken chassis cluster.

Options

interval Specifies the transmit interval, in milliseconds.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D60.

RELATED DOCUMENTATION

Understanding Media Access Control Security (MACsec)

Configuring Media Access Control Security (MACsec)

macsec | **10278**

transmit-interval (MACsec for MX Series)

IN THIS SECTION

- Syntax | 30095
- Hierarchy Level | 30095
- Description | 30095
- Default | 30096
- Options | 30096
- Required Privilege Level | 30096
- Release Information | 30096

Syntax

```
transmit-interval interval;
```

Hierarchy Level

```
[edit security macsec connectivity-association connectivity-association-name mka]
```

Description

Specifies the transmit interval for MACsec Key Agreement (MKA) protocol data units (PDUs).

The MKA transmit interval setting sets the frequency for how often the MKA PDU is sent to the directly connected device to maintain MACsec on a point-to-point Ethernet link. A lower *interval* increases bandwidth overhead on the link; a higher *interval* optimizes the MKA protocol data unit exchange process.

The transmit interval settings must be identical on both ends of the link when MACsec using static connectivity association key (CAK) security mode is enabled.

We recommend increasing the interval to 6000 milliseconds in high-traffic and large scale configuration load environments.

Default

The default transmit interval is 2000 milliseconds.

Options

interval Specifies the transmit interval, in milliseconds.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Media Access Control Security (MACsec) on Routers

transmit-interval (PIM BFD Liveness Detection)

IN THIS SECTION

- Syntax | 30097
- Hierarchy Level | 30097
- Description | 30097
- Required Privilege Level | 30097
- Release Information | 30097

Syntax

```
transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
}
```

Hierarchy Level

```
[edit protocols pim interface interface-name bfd-liveness-detection],
[edit routing-instances routing-instance-name protocols pim interface interface-name bfd-
liveness-detection]
```

Description

Specify the transmit interval for the `bfd-liveness-detection` statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

Support for BFD authentication introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring BFD for PIM

bfd-liveness-detection (Protocols PIM) | 2945
threshold (PIM BFD Transmit Interval) | 29333
minimum-interval (PIM BFD Transmit Interval) | 10739
minimum-receive-interval (Protocols PIM) | 10751

transmit-rate

IN THIS SECTION

- QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | 30098
- QFX10000 Switches | 30098
- Hierarchy Level | 30098
- Description | 30099
- Default | 30100
- Options | 30102
- Required Privilege Level | 30102
- Release Information | 30102

QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems

```
transmit-rate (rate | percent percentage);
```

QFX10000 Switches

```
transmit-rate (rate | percent percentage) <exact>;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, the transmit rate specifies the minimum guaranteed transmission rate or percentage for a queue (forwarding class) scheduler. The queue transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share on switches that support enhanced transmission selection (ETS) hierarchical scheduling.

On QFX10000 switches, the transmit rate specifies the minimum guaranteed transmission rate or percentage for a queue (forwarding class) scheduler. The queue transmit rate also determines the amount of excess (extra) port bandwidth the queue can share if you do not explicitly configure an excess rate in the scheduler. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share on switches that support enhanced transmission selection (ETS) hierarchical scheduling.

On QFX10000 switch strict-high priority queues, the transmit rate limits the amount of traffic the switch treats as strict-high priority traffic. Traffic up to the transmit rate receives strict-high priority treatment. The switch treats traffic that exceeds the transmit rate as best-effort traffic that receives an excess bandwidth sharing weight of “1”; you cannot configure an excess rate on a strict-high priority queue, and unlike queues with other scheduling priorities, the switch does not use the transmit rate to determine extra bandwidth sharing for strict-high priority queues.

CAUTION: We strongly recommend that you configure a transmit rate on strict-high priority queues to limit the amount of traffic the switch treats as strict-high priority traffic on those queues. This is especially important if you configure more than one strict-high priority queue on a port. To prevent a strict-high priority queue from starving the other queues on a port, we recommend that you always configure a transmit rate, even if you only configure one strict-high priority queue.

NOTE: For ETS, the `transmit-rate` setting works only if you also configure the "guaranteed-rate" on page 7940 in the traffic control profile that is attached to the forwarding class set to which the queue belongs. If you do not configure the guaranteed rate, the minimum guaranteed rate for individual queues that you set using the `transmit-rate` statement does not work. The sum of all queue transmit rates in a forwarding class set should not exceed the traffic control profile guaranteed rate.

NOTE: On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, you cannot configure a transmit rate for a strict-high priority queue. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has a strict-high priority queue. To prevent strict-high priority queues from

consuming all of the available bandwidth on these switches, we recommend that you configure a shaping rate to set a maximum amount of bandwidth for strict-high priority queues.

NOTE: For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.

Default

On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, if you do not configure the transmit rate, the default scheduler transmission rate and buffer size percentages for queues 0 through 11 are:

Table 1831: Default Transmit Rates for QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems

Queue Number	Default Minimum Guaranteed Bandwidth (Transmit Rate)
0 (best-effort)	5 %
1	0
2	0
3 (fcoe)	35 %
4 (no-loss)	35 %
5	0
6	0
7 (network control)	5 %

Table 1831: Default Transmit Rates for QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems *(Continued)*

Queue Number	Default Minimum Guaranteed Bandwidth (Transmit Rate)
8 (mcast)	20 %
9	0
10	0
11	0

On QFX10000 switches, if you do not configure the transmit rate, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are:

Table 1832: Default Transmit Rates for QFX10000 Switches

Queue Number	Default Minimum Guaranteed Bandwidth (Transmit Rate)
0 (best-effort)	15 %
1	0
2	0
3 (fcoe)	35 %
4 (no-loss)	35 %
5	0
6	0
7 (network control)	15 %

Configure schedulers if you want to change the minimum guaranteed bandwidth and other queue characteristics.

Options

rate—Minimum transmission rate for the queue, in bps. You can specify a value in bits-per-second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), m (1,000,000), or `g` (1,000,000,000).

- **Range:** 1000 through 10,000,000,000 bps on 10-Gigabit interfaces, 1000 through 40,000,000,000 bps on 40-Gigabit interfaces.

`percent` *percentage*—Minimum percentage of transmission capacity allocated to the queue.

- **Range:** 1 through 100 percent

`exact`—(QFX10000 switches only) Shape queues that are not strict-high priority queues to the transmit rate so that the transmit rate is the maximum bandwidth limit. Traffic that exceeds the exact transmit rate is dropped. You cannot set an excess rate on queues configured as `transmit-rate` *(rate | percentage)* `exact` because the purpose of setting an exact transmit rate is to set a maximum bandwidth (shaping rate) on the traffic.

> **NOTE**: On QFX10000 switches, oversubscribing all 8 queues configured with the `transmit rate exact` (shaping) statement at the `[edit class-of-service schedulers` *scheduler-name*`]` hierarchy level might result in less than 100 percent utilization of port bandwidth.

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Exact option introduced in Junos OS Release 15.1X53-D10 for the QFX Series.

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Queue Schedulers

Example: Configuring Queue Schedulers for Port Scheduling

Understanding CoS Output Queue Schedulers

Understanding CoS Port Schedulers on QFX Switches

transmit-rate (Dynamic Schedulers)

IN THIS SECTION

- Syntax | 30103
- Hierarchy Level | 30103
- Description | 30103
- Default | 30104
- Options | 30104
- Required Privilege Level | 30104
- Release Information | 30104

Syntax

```
transmit-rate (rate | percent percentage | remainder | percent percentage $junos-cos-scheduler-
tx) <exact | rate-limit>;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
```

Description

Specify the transmit rate or percentage for a scheduler in a dynamic profile.

Default

If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

Options

rate—Transmission rate, in bps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 3200 through 6,400,000,000,000 bps

`percent` *percentage*—Percentage of transmission capacity. A percentage of zero drops all packets in the queue.

- **Range:** 0 through 100 percent

`remainder`—Use remaining rate available.

`$junos-cos-scheduler-tx`—Junos predefined variable that is replaced with the transmission rate obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

`exact`—(Optional) Enforce the exact transmission rate. Under sustained congestion, a rate-controlled queue that goes into negative credit fills up and eventually drops packets. Make sure this value never exceeds the rate-controlled amount.

`rate-limit`—(Optional) Limit the transmission rate to the rate-controlled amount during congestion. In contrast to the `exact` option, when there is no congestion, the scheduler with the `rate-limit` option shares unused bandwidth above the rate-controlled amount.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

The `$junos-cos-scheduler-tx` predefined variable introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Configuring Schedulers in a Dynamic Profile for Subscriber Access

scheduler (Dynamic Scheduler Maps) | **16344**

transmit-rate (EX Series Switches)

IN THIS SECTION

- Syntax | **30105**
- Hierarchy Level | **30105**
- Description | **30105**
- Default | **30106**
- Options | **30106**
- Required Privilege Level | **30106**
- Release Information | **30106**

Syntax

```
transmit-rate (rate | percent percentage | remainder);
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Specify the transmit rate or percentage for a scheduler.

Default

If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent.

Options

rate Transmission rate, in bps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

- **Range:** 3200 through 160,000,000,000 bps

  (EX4300 switches only) 8000 through 160,000,000,000 bps

percent ***percentage*** Percentage of transmission capacity. A percentage of zero drops all packets in the queue.

- **Range:** 0 through 100 percent

remainder Remaining rate available

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches

Defining CoS Schedulers and Scheduler Maps (CLI Procedure)

Defining CoS Schedulers (J-Web Procedure)

Understanding CoS Schedulers

transmit-rate (Schedulers)

IN THIS SECTION

- Syntax | 30107
- Hierarchy Level | 30107
- Description | 30107
- Default | 30107
- Options | 30108
- Required Privilege Level | 30109
- Release Information | 30109

Syntax

```
transmit-rate (rate | percent percentage | remainder percentage) <exact | rate-limit>;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Specify the transmit rate or percentage for a scheduler.

Default

If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent, respectively.

NOTE: On PTX10008 and PTX10016 routers, the default scheduler transmission rate is set to 25 percent for strict-high queues if you do not set the transmission rate for the schedulers for strict-high queues.

Options

exact (Optional) Enforce the exact transmission rate. Under sustained congestion, a rate-controlled queue that goes into negative credit fills up and eventually drops packets. This value should never exceed the rate-controlled amount. For PTX Series routers, this option is allowed only on the non-strict-high (high, medium-high, medium-low, or low) queues.

percent *percentage* Percentage of transmission capacity. A percentage of zero drops all packets in the queue unless additional bandwidth is available from other queues.

- **Range:** 0 through 100 percent for M, MX and T Series routers and EX Series switches; 1 through 100 percent for PTX Series routers

> **NOTE**:
>
> - On M Series Multiservice Edge Routers, for interfaces configured on 4-port E1 and 4-port T1 PICs only, you can configure a *percentage* value only from 11 through 100. These two PICs do not support transmission rates less than 11 percent.
> - The configuration of the `transmit-rate percent 0 exact` statement at the [`edit class-of-service schedules` *`scheduler-name`*] hierarchy is ineffective on T4000 routers with Type 5 FPC.
> - On MIC and MPC interfaces on MX Series routers, when the transmit rate is configured as a percentage and **exact** or **rate-limit** is enabled on a queue, the shaping rate of the parent node is used to compute the transmit rate. If **exact** or **rate-limit** is not configured, the guaranteed rate of the parent node is used to compute the transmit rate.
> - On PTX Series routers, unconfigured interfaces are equivalent to `percent 0`. This means the system offers no guaranteed rate on the interface, and the queue will always be scheduled in the excess priority.

rate Transmission rate, in bps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 3200 through 6,400,000,000,000 bps

> **NOTE**: For all MX Series interfaces, the rate can be from 65,535 through 6,400,000,000,000 bps.

rate-limit (Optional) Limit the transmission rate to the rate-controlled amount by applying a policing action to the queue. Packets are hard-dropped when traffic exceeds the specified maximum transmission rate.

> **NOTE**: For PTX Series routers, this option is allowed only on the strict-high queue. We recommend that you configure rate limit on strict-high queues because the other queues may not meet their guaranteed bandwidths. The `rate-limit` option cannot rate limit the queue if strict-priority scheduling is configured with the "strict-priority-scheduler" on page 28650 statement.

> **NOTE**: The configuration of the `rate-limit` statement is supported on T4000 routers only with a Type 5 FPC.

remainder *percentage* Use the remaining rate available.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`rate-limit` option introduced in Junos OS Release 8.3. Applied to the Multiservices PICs in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Schedulers

Configuring Scheduler Transmission Rate

Understanding Scheduling on PTX Series Routers

transparent

IN THIS SECTION

- Syntax | 30110
- Hierarchy Level | 30110
- Description | 30110
- Required Privilege Level | 30110
- Release Information | 30110

Syntax

```
transparent;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers ieee802.1
vlan-tag]
```

Description

Packet classification based on the transparent VLAN tag.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2

transport (Services)

IN THIS SECTION

- Syntax | 30111
- Hierarchy Level | 30111
- Description | 30111
- Options | 30111
- Required Privilege Level | 30112
- Release Information | 30112

Syntax

```
transport [ transport-protocols ];
```

Hierarchy Level

```
[edit services nat pool nat-pool-name pgcp]
```

Description

Configure the BGF to select a NAT pool based on transport protocol type.

Options

[*transport-protocol*]—One or more transport protocols.

- **Values: rtp-avp**, **tcp**, **udp**
- **Syntax:** One or more protocols. If you specify more than one protocol, you must enclose all protocols in brackets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

transport (System Syslog)

IN THIS SECTION

- Syntax | 30112
- Hierarchy Level | 30112
- Description | 30113
- Options | 30113
- Release Information | 30113

Syntax

```
transport(tcp | tls | udp);
```

Hierarchy Level

```
[edit system syslog host host-name]
```

Description

Configure the system syslog transport option.

Options

tcp Configure TCP as the transport protocol.

tls Configure TLS as the transport protocol. TLS uses certificates to authenticate peers. Before configuring this statement, make sure that your device has CA certificate and the PKI.

udp Configure UDP as the transport protocol.

Default: UDP

Release Information

Statement introduced in Junos OS Release 21.2R1 for EX Series switches and MX Series routers.

Statement introduced in Junos OS Release 21.3R1 for SRX Series and vSRX Virtual Firewall devices.

RELATED DOCUMENTATION

Configure Syslog over TLS

transport (Protocols BGP)

IN THIS SECTION

- Syntax | 30114
- Hierarchy Level | 30116
- Description | 30117
- Options | 30117
- Required Privilege Level | 30118
- Release Information | 30118

Syntax

```
transport {
 accepted-prefix-limit {
  drop-excess< limit-threshold>;
  hide-excess< limit-threshold>;
  maximum maximum;
  teardown<idle-timeout(forever |  timeout)>< limit-threshold>;
 }
 accept-srv6-service;
 add-path {
  receive;
  send {


   include-backup-path include-backup-path;
   multipath;
   path-count path-count;
   path-selection-mode {
    (all-paths | equal-cost-paths);
   }
   prefix-policy [ prefix-policy ... ];
  }
 }
 advertise-srv6-service;
 aggregate-label {
  community community;
 }
 aigp {
  disable;
 }
 damping;
 defer-initial-multipath-build {
  maximum-delay maximum-delay;
 }


 delay-route-advertisements {
  always-wait-for-krt-drain;
  maximum-delay {
   route-age route-age;
```

```
      routing-uptime routing-uptime;
     }
     minimum-delay {
      inbound-convergence inbound-convergence;
      routing-uptime routing-uptime;
     }
    }





    extended-nexthop;
    extended-nexthop-color;
    extended-nexthop-tunnel;


    graceful-restart {
     forwarding-state-bit(from-fib | set);
     long-lived {
      extended-route-retention {
       (disable | enable);
       retention-policy [ retention-policy ... ];
       retention-time retention-time;
      }


      restarter {
       (disable | enable);
       stale-time stale-time;
      }
     }
    }
    local-ipv4-address local-ipv4-address;
    loops loops;
    lu-export {
     policy [ policy ... ];
    }
    nexthop-resolution {
     no-resolution;
     preserve-nexthop-hierarchy;
    }
    no-install;
    output-queue-priority(expedited | priority priority);
```

```
    per-group-label;
    per-prefix-label;
    prefix-limit {
     drop-excess< limit-threshold>;
     hide-excess< limit-threshold>;
     maximum maximum;
     teardown<idle-timeout(forever |  timeout)>< limit-threshold>;
    }
    protection;
    rib-group ribgroup-name

;
    route-age-bgp-view;
    route-refresh-priority(expedited | priority priority);
    traffic-statistics {
     file filename<files files>



<size size><(world-readable | no-world-readable)>;
     interval interval;
     labeled-path;
    }
    withdraw-priority(expedited | priority priority);
   }
```

Hierarchy Level

```
[edit fabric protocols bgp family inet6],
[edit fabric protocols bgp group name family inet6],
[edit fabric protocols bgp group name neighbor name family inet6],
[edit logical-systems name protocols bgp family inet6],
[edit logical-systems name protocols bgp group name family inet6],
[edit logical-systems name protocols bgp group name neighbor name family inet6],
[edit logical-systems name routing-instances name protocols bgp family inet6],
[edit logical-systems name routing-instances name protocols bgp group name family inet6],
[edit logical-systems name routing-instances name protocols bgp group name neighbor name family
inet6],
[edit logical-systems name tenants name routing-instances name protocols bgp family inet6],
[edit logical-systems name tenants name routing-instances name protocols bgp group name family
```

```
inet6],
[edit logical-systems name tenants name routing-instances name protocols bgp group name neighbor
name family inet6],
[edit protocols bgp family inet6],
[edit protocols bgp group name family inet6],
[edit protocols bgp group name neighbor name family inet6],
[edit routing-instances name protocols bgp family inet6],
[edit routing-instances name protocols bgp group name family inet6],
[edit routing-instances name protocols bgp group name neighbor name family inet6],
[edit tenants name routing-instances name protocols bgp family inet6],
[edit tenants name routing-instances name protocols bgp group name family inet6],
[edit tenants name routing-instances name protocols bgp group name neighbor name family inet6]
```

Description

Enable Classful Transport (CT) NLRI for inet6 family in BGP networks.

Options

accept-srv6-service Accept SRv6 service.

advertise-srv6-service Advertise SRv6 service.

aigp Allow sending and receiving of AIGP attribute.

- disable—Disable sending and receiving of AIGP attribute

damping Enable route flap damping.

extended-nexthop Extended nexthop encoding.

extended-nexthop-color Resolve using extended color nexthop.

extended-nexthop-tunnel Use BGP tunnel attribute.

local-ipv4-address Configure a local IPv4 address.

nexthop-resolution Configure nexthop resolution properties:

- no-resolution—Consider nexthop good without resolution attempt

- preserve-nexthop-hierarchy—Attempt preserving resolved nexthop chain in forwarding

no-install Dont install received routes in forwarding.

per-group-label Advertise prefixes with unique labels per group.

per-prefix-label Allocate a unique label to each advertised prefix.

protection Compute backup path for active nexthop failure.

route-age-bgp-view Maintain BGP route's age based on update messages only.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

name (Transport-class) | 11181

transport-class | 30120

transport-address

IN THIS SECTION

- Syntax | 30119
- Hierarchy Level | 30119
- Description | 30119

- Default | 30119
- Options | 30119
- Required Privilege Level | 30120
- Release Information | 30120

Syntax

```
transport-address (address | interface | router-id);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name protocols ldp interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit protocols ldp interface interface-name],
[edit routing-instances routing-instance-name protocols ldp interface interface-name]
```

Description

Enables you to configure the IP address used to specify the TCP session for the LDP session. Routers must first establish a TCP session between one another before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.

Default

`router-id`

Options

address Use the specified IP address as the transport address for the LDP session.

interface Use the first IP address on the interface as the transport address for any LDP sessions to neighbors that can be reached over that interface. You cannot specify the `interface` option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the `router-id` option.

router-id Use router identifier as the transport address. Unless otherwise configured, the router identifier is the loopback address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

address option introduced in Junos OS Release 19.1R1 for all platforms.

RELATED DOCUMENTATION

Specifying the Transport Address Used by LDP

transport-class

IN THIS SECTION

- Syntax | 30121
- Hierarchy Level | 30122
- Release Information | 30122
- Description | 30122
- Options | 30122

- Required Privilege Level | **30123**

Syntax

```
transport-class {
    auto-create;
    fallback {
        none;
    }
    name (Transport-class) name {
        color color;
        fallback {
            none;
        }
        tunnel-egress {
            end-point endpoint-address {
                import [ import-policy ... ];
            }
        }
    }
    any-class {
        color color;
        fallback {
            none;
        }
    }
    best-effort {
        color color;
        fallback {
            none;
        }
    }
}
```

Hierarchy Level

```
[edit routing-options]
```

Release Information

Statement introduced in Junos OS Release 21.2R1.

Description

Specify the transport layer options.

Options

any-class Specify the predefined transport class that collects tunnels of all transport classes.

- `color` *`color`*—This is a 32bit color value mapping to any transport class. This auto derives the mapping community.

auto-create Auto create on color discovery.

best-effort Specify the best effort transport class.

- `color` *`color`*—This is a 32bit color value mapping to any transport class. This auto derives the mapping community.

fallback Service routes resolution fallback. Service routes typically resolves over inet.3.

- `none`—Resolve associated service-routes using the specified transport class tunnel only.

name Specify the transport layer that collects tunnels with similar characteristics.

- *`name`*—Specify the transport class name.
- `color` *`color`*—This is a 32bit color value mapping to any transport class. This auto derives the mapping community.
- `tunnel-egress`—Specify the endpoint origination of the egress tunnel.
  - `end-point` *`endpoint-addr`*—Configure the egress tunnel endpoint by specifying the address.
    - `import [`*`import ...`*`]`—Apply the import policy to set the attributes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing

RELATED DOCUMENTATION

use-transport-class | 30541

MPLS Traffic Engineering Configuration

transport (Diameter Base Protocol)

IN THIS SECTION

- Syntax | 30123
- Hierarchy Level | 30124
- Description | 30124
- Options | 30124
- Required Privilege Level | 30124
- Release Information | 30124

Syntax

```
transport transport-name {
    address;
    logical-system logical-system-name <routing-instance routing-instance-name >;
    routing-instance routing-instance-name
}
```

Hierarchy Level

```
[edit diameter]
```

Description

Configure the Diameter instance and the local IP address for the Diameter local transport connection.

Options

transport-name—Name of the transport.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Diameter Base Protocol

Diameter Base Protocol

transport (Diameter Peer)

IN THIS SECTION

Syntax | 30125

- Hierarchy Level | 30125
- Description | 30125
- Default | 30125
- Options | 30125
- Required Privilege Level | 30125
- Release Information | 30126

Syntax

```
transport transport-name;
```

Hierarchy Level

```
[edit diameter peer peer-name No Link Title]
```

Description

Specify the transport layer connection to be used for establishing active connections to the peer.

Default

The transport is defined in the default logical system and master routing instance.

Options

transport-name—Name of the transport.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Diameter Base Protocol

Diameter Base Protocol

transport-monitoring

IN THIS SECTION

- Syntax | 30126
- Hierarchy Level | 30126
- Description | 30126
- Required Privilege Level | 30127
- Release Information | 30127

Syntax

```
transport-monitoring;
```

Hierarchy Level

```
[edit interfaces]
```

Description

Monitor the performance and state of packet transport for OTN and optics modules. The following statistics are monitored:

- Packet transport for ninety-six 15–minute intervals for the current 24 hours.
- Cumulative data of the current 24 hours.
- Cumulative data of the previous 24 hours.

If this statement is configured, transport monitoring related information is shown in the output of `show interface transport` command and corresponding MIBs are available. If this option is disabled, an error is shown in the output and corresponding MIBs are not available.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

show interfaces transport pm | 21509

transport (Next Gen Services Syslog Message Security)

IN THIS SECTION

- Syntax | 30128
- Hierarchy Level | 30128
- Description | 30128
- Options | 30128
- Required Privilege Level | 30128
- Release Information | 30128

Syntax

```
transport;
```

Hierarchy Level

```
[edit services service-set name  syslog
```

Description

Specify the category for which you want to collect local logs.

Options

apply-groups	Groups from which to inherit configuration data
apply-groups-except	Don't inherit configuration data from these groups
protocol	Set security log transport protocol for the device. You can set the protocol to TCP, TLS or UDP
tcp-connections	Set tcp connection number per-stream (1..5)
tls-profile	If you are using the TLS protocol as the security log transport, specify the TLS profile name to use.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

transport (Security Log)

IN THIS SECTION

- Syntax | 30129
- Hierarchy Level | 30129
- Description | 30129
- Options | 30130
- Required Privilege Level | 30130
- Release Information | 30130

Syntax

```
transport {
    protocol (udp | tcp | tls);
    tls-profile tls-profile-name;
    tcp-connections tcp-connections;
}
```

Hierarchy Level

```
[edit security log]
[edit logical-systems name security log]
[edit tenants tenant-name security log]
```

Description

Configure security log transport options.

Options

protocol

Specify the type of transport protocol to be used to log the data.

- **UDP**—Set the transport protocol to UDP.
- **TCP**—Set the transport protocol to TCP.
- **TLS**—Set the transport protocol to TLS.
- **Default:** UDP.

tls-profile ***tls-profile-name***

Specify the TLS profile name.

tcp-connections ***tcp-connections***

Specify the number of TCP connections per SPU.

- **Range:** 1 through 5.
- **Default:** 1.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X46-D25.

The [edit `logical-systems` *name* security log] and [edit `tenants` *tenant-name* security log] hierarchy levels introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Understanding AppTrack

transport-type

IN THIS SECTION

- Syntax | 30131
- Hierarchy Level | 30131
- Description | 30131
- Options | 30131
- Required Privilege Level | 30131
- Release Information | 30132

Syntax

```
transport-type type;
```

Hierarchy Level

```
[edit services hosted-services server-profile server-profile-name]
```

Description

Configure the transport type.

Options

type—Transport type.

- **Range:** GRE, TCP, or UDP

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

transport-class

IN THIS SECTION

- Syntax | 30132
- Hierarchy Level | 30132
- Description | 30133
- Release Information | 30133

Syntax

```
transport-class color;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name provider-tunnel ingress-replication],
[edit logical-systems name tenants name routing-instances name provider-tunnel ingress-
replication],
[edit routing-instances name provider-tunnel ingress-replication],
[edit tenants name routing-instances name provider-tunnel ingress-replication]
```

Description

By setting the transport-class for a specific LSP, you determine which RSVP tunnels are placed into corresponding colored inet.3 tables.

> **NOTE**: A maximum of 3 transport-class color tables are supported.
>
> A maximum of 9000 LSPs are supported per transport-class.
>
> A maximum of 9000 ingress replication provider tunnels are supported with 3 colored tables.
>
> A maximum of 3 VRFs can be mapped to 3 colored tables.

Release Information

Statement introduced in Junos OS Release 24.4R1.

trap

IN THIS SECTION

- Syntax | 30133
- Hierarchy Level | 30134
- Description | 30134
- Options | 30134
- Required Privilege Level | 30134
- Release Information | 30134

Syntax

```
trap {
    interval trap interval;
}
```

Hierarchy Level

```
[edit security screen trap],
[edit tenants tenant-name security screen]
```

Description

Traps are unsolicited messages sent from an SNMP agent to remote network management systems or trap receivers. Many enterprises use SNMP traps as part of a fault-monitoring solution, in addition to system logging. In Junos OS, SNMP traps are not forwarded by default. You can use the SNMP traps by configuring a `trap-group`.

You can create and name a group of one or more types of SNMP traps and then define which systems receive the group of SNMP traps. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name.

Options

interval

Configure the trap interval.

- **Range:** 1 through 3600 seconds
- **Default:** 2 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D20.

The `tenant` option is introduced in Junos OS Release 18.3R1.

trap-group

IN THIS SECTION

- Syntax: MX, M, T, ACX and PTX Series routers, OCX1100, EX and QFX Series Switches and SRX and vSRX Virtual Firewall firewalls | 30135
- Syntax: QFX and EX Series Switches and OCX1100 | 30135
- Hierarchy Level | 30136
- Description | 30136
- Options | 30136
- Required Privilege Level | 30136
- Release Information | 30136

Syntax: MX, M, T, ACX and PTX Series routers, OCX1100, EX and QFX Series Switches and SRX and vSRX Virtual Firewall firewalls

```
trap-group group-name {
    categories {
        category;
    }
    destination-port port-number;
    routing-instance instance;
    targets {
        address;
    }
    version (all | v1 | v2);
}
```

Syntax: QFX and EX Series Switches and OCX1100

```
trap-group group-name {
    categories {
        category;
    }
    destination-port port-number;
```

```
    targets {
        address;
    }
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Create a named group of hosts to receive the specified trap notifications. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name. At least one trap group must be configured for SNMP traps to be sent.

Options

group-name—Name of the trap group. If the name includes spaces, enclose it in quotation marks (" ").

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

trap-options

IN THIS SECTION

Syntax: MX, M, T, ACX and PTX Series Routers, EX Series Switches and SRX Firewalls | 30137

- QFX Series Switches, EX4600, OCX1100 | **30137**
- Hierarchy Level | **30138**
- Description | **30138**
- Default | **30138**
- Options | **30138**
- Required Privilege Level | **30138**
- Release Information | **30138**

Syntax: MX, M, T, ACX and PTX Series Routers, EX Series Switches and SRX Firewalls

```
trap-options {
    agent-address outgoing-interface;
    context-oid;
    enterprise-oid;
    logical-system logical-system-name {
        routing-instance routing-instance-name {
            source-address address;
        }
    }
    routing-instance routing-instance-name {
        source-address address;
    }
}
```

QFX Series Switches, EX4600, OCX1100

```
trap-options {
    agent-address outgoing-interface;
    source-address address;
    routing-instance routing-instance-name {
        source-address address;
    }
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router or switch to a single address, regardless of the outgoing interface. In addition, you can set the agent address of each SNMPv1 trap. For more information about the contents of SNMPv1 traps, see RFC 1157.

Default

Disabled

Options

context-oid Add context oid in varbind of all traps originating from non-default logical system and routing instance. If your network management system is not able to handle prefixes such as `<routing-instance name>@<trap-group-name>` or `<logical-system name>/<routing-instance name>@<trap-group-name>`, setting the `context-oid` configuration statement will send only the `<trap-group-name>` and add `<logical-system name>/<routing-instance name>` as an additional varbind.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`context-oid` option introduced in Junos OS Release 17.1.

trap (Security IKE)

IN THIS SECTION

- Syntax | 30139
- Hierarchy Level | 30139
- Description | 30139
- Default | 30139
- Required Privilege Level | 30140
- Release Information | 30140

Syntax

```
trap {
    ipsec-tunnel-down ;
    peer-down ;
}
```

Hierarchy Level

```
[edit security ike trap]
```

Description

ipsec-tunnel-down— By enabling the `ipsec-tunnel-down` option, the traps are generated for IPsec tunnel going down only when the associated peer IKE SA is up.

peer-down—By enabling the `peer-down` option, the traps are generated when peer goes down. Traps for the corresponding IPsec tunnels going down for the peer are not generated.

Default

- `peer-down`—Disabled

- `ipsec-tunnel-down`—Disabled

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1R1

RELATED DOCUMENTATION

Configure the Certificate Expiration Trap

Enable Peer Down and IPsec Tunnel Down Traps

show security ipsec statistics | **24744**

trap (Security PKI)

IN THIS SECTION

- Syntax | **30141**
- Hierarchy Level | **30141**
- Description | **30141**
- Default | **30141**
- Options | **30141**
- Required Privilege Level | **30141**
- Release Information | **30142**

Syntax

```
trap {
    all certificates number-of-days;
    ca-identity number-of-days;
    certificate-id number-of-days;
}
```

Hierarchy Level

```
[edit security pki]
```

Description

To configure certificate expiration trap and configure the number of days before to generate the trap.

how

Default

- `all certificates`—60 days
- `ca-identity`—60 days
- `certificate-id` —60 days

Options

`all certificates` *`number-of-days`* —Number of days to generate traps for all certificates. You can configure from *1 to 100* days.

`ca-identity` *`number-of-days`* —Number of days to generate traps for specified CA certificate. You can configure from *1 to 100* days.

`certificate-id` *`number-of-days`* —Number of days to generate for specified local certificate. You can configure from *1 to 100* days.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1R1

RELATED DOCUMENTATION

Configure the Certificate Expiration Trap

Enable Peer Down and IPsec Tunnel Down Traps

show security ipsec statistics | **24744**

traps

IN THIS SECTION

- Syntax | **30142**
- Hierarchy Level (ACX Series, MX Series, T Series, M Series, SRX Series, EX Series) | **30142**
- Hierarchy Level (QFX Series, EX4600) | **30143**
- Description | **30143**
- Required Privilege Level | **30143**
- Release Information | **30143**

Syntax

```
(traps | no-traps);
```

Hierarchy Level (ACX Series, MX Series, T Series, M Series, SRX Series, EX Series)

```
[edit dynamic-profiles profile-name interfaces interface-name],
[edit interfaces interface-name],
[edit interfaces interface-name  unit logical-unit-number],
```

```
[edit interfaces interface-range name],
[edit logical-systems logical-system-name interfaces interface-name  unit logical-unit-number]
```

Hierarchy Level (QFX Series, EX4600)

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-range interface-range-name]
```

Description

Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

(Enhanced subscriber management for MX Series routers) To enable SNMP notifications, you must first configure the `interface-mib` statement at the `[edit dynamic-profiles profile-name interfaces interface-name]` hierarchy level. If `interface-mib` is not configured, the `traps` statement has no effect.

BEST PRACTICE: To achieve maximum performance when enhanced subscriber management is enabled, we recommend that you *not* enable SNMP notifications on all dynamic subscriber interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Support at the `[edit dynamic-profiles profile-name interfaces interface-name]` hierarchy level introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

RELATED DOCUMENTATION

- Enabling or Disabling SNMP Notifications on Physical Interfaces
- Enabling or Disabling SNMP Notifications on Logical Interfaces

traps (RPM and TWAMP)

IN THIS SECTION

- Syntax | 30144
- Junos OS Hierarchy Level | 30144
- Junos OS Evolved Hierarchy Level | 30145
- Description | 30145
- Options | 30145
- Additional Information | 30146
- Required Privilege Level | 30147
- Release Information | 30147

Syntax

```
traps traps;
```

Junos OS Hierarchy Level

```
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name]
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner test test-name]
[edit services monitoring twamp client control-connection control-client-name]
[edit services monitoring twamp client control-connection control-client-name test-session session-name]
```

Description

Set the trap bit to generate SNMP traps for probes. Traps are sent if the configured threshold is met or exceeded. For TWAMP, you can configure both control-connection and test-session traps for managed clients, but only test-session traps for light clients. See SNMP MIB Explorer to learn more about the SNMP MIB objects that Juniper supports.

Default: no traps are generated

Options

traps Specify one or more SNMP traps. The options are configured per test session, except for the `control-connection-closed` and the `test-iteration-done` options. These options are configured only for TWAMP, and only per control connection. The following options are supported:

- `control-connection-closed`—(TWAMP control-connection only) Generate traps when the control connection is closed.
- `egress-jitter-exceeded`—Generate traps when the jitter in egress time threshold is met or exceeded. (MIB object jnxPingEgressJitterThresholdExceeded)
- `egress-std-dev-exceeded`—(Junos OS only) Generate traps when the egress time standard deviation threshold is met or exceeded. (MIB object jnxPingEgressStdDevThresholdExceeded)
- `egress-time-exceeded`—Generate traps when the maximum egress time threshold is met or exceeded. (MIB object jnxPingEgressThresholdExceeded)
- `ingress-jitter-exceeded`—Generate traps when the jitter in ingress time threshold is met or exceeded. (MIB object jnxPingIngressJitterThresholdExceeded)
- `ingress-std-dev-exceeded`—(Junos OS only) Generate traps when the ingress time standard deviation threshold is met or exceeded. (MIB object jnxPingIngressStddevThresholdExceeded)
- `ingress-time-exceeded`—Generate traps when the maximum ingress time threshold is met or exceeded. (MIB object jnxPingIngressThresholdExceeded)

- `jitter-exceeded`—Generate traps when the jitter in round-trip time threshold is met or exceeded. (MIB object jnxPingRttJitterThresholdExceeded)
- `max-rtt-threshold`—(Junos OS only) Generate traps when the maximum round trip time threshold at the end of the test is met or exceeded. (MIB object jnxPingMaxRttThresholdExceeded)
- `probe-failure`—Generate traps when successive probe loss thresholds are crossed. (MIB object pingProbeFailed)
- `rtt-exceeded`—Generate traps when the maximum round-trip time threshold is met or exceeded. (MIB object jnxPingRttThresholdExceeded)
- `std-dev-exceeded`—(Junos OS only) Generate traps when the round-trip time standard deviation threshold is met or exceeded. (MIB object jnxPingRttStdDevThresholdExceeded)
- `test-completion`—Generate traps when a test is completed. (MIB object pingTestCompleted)
- `test-failure`—Generate traps when the total probe loss threshold is met or exceeded. (MIB object pingTestFailed)
- `test-iteration-done`—(TWAMP control-connection only) Generate traps when all test sessions under control connections complete one test iteration.

NOTE: To generate RPM traps, you must configure the `remote-operations` SNMP trap category by including the `categories` statement at the `[edit snmp trap-group` *`trap-group-name`*`]` hierarchy level.

Additional Information

(Junos OS Evolved only) The Juniper PING MIB does not support the following notifications, because the corresponding thresholds are not configurable for Junos OS Evolved and hence the traps are never generated:

- jnxPingRttStdDevThresholdExceeded
- jnxPingEgressStdDevThresholdExceeded
- jnxPingIngressStddevThresholdExceeded
- jnxPingMaxRttThresholdExceeded

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [`edit services rpm twamp client control-connection` *`control-client-name`*] and [`edit services rpm twamp client control-connection` *`control-client-name`* `test-session` *`session-name`*] hierarchy levels introduced in Junos OS Release 15.1 for MX Series routers.

Support at the [`edit services monitoring rpm owner test` *`test-name`*] hierarchy level introduced in Junos OS Evolved Release 21.2R1 for PTX Series routers.

Support at the [edit services monitoring twamp client control-connection *control-client-name*] and [edit services monitoring twamp client control-connection *control-client-name* test-session *session-name*] hierarchy levels introduced in Junos OS Evolved Release 21.3R1 for PTX Series routers.

RELATED DOCUMENTATION

- thresholds (Junos OS) | **29348**
- thresholds (Junos OS Evolved) | **29350**
- *Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches*
- *Understanding Two-Way Active Measurement Protocol on Routers*

tri-color

IN THIS SECTION

- Syntax | **30148**
- Hierarchy Level | **30148**
- Description | **30148**
- Default | **30148**
- Required Privilege Level | **30148**

Release Information | 30148

Syntax

```
tri-color;
```

Hierarchy Level

```
[edit class-of-service]
```

Description

For IPv4 packets on M320, MX Series, T Series routers with Enhanced II Flexible PIC Concentrators (FPCs), and EX Series switches, enable tricolor marking (TCM), as defined in RFC 2698.

Default

If you do not include this statement, tricolor marking is not enabled and the medium packet loss priority (PLP) is not configurable.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring and Applying Tricolor Marking Policers

tri-color (EX Series Switches)

IN THIS SECTION

- Syntax | 30149
- Hierarchy Level | 30152
- Description | 30152
- Options | 30152
- Required Privilege Level | 30152
- Release Information | 30152

Syntax

```
tri-color {
    classifiers {
        (dscp | dscp-ipv6 |exp |ieee-802.1 | inet-precedence) classifier-name {
            forwarding-class (Forwarding Policy) (class-name | assured-forwarding | best-effort |
expedited-forwarding | network-control) {
                loss-priority (high | low | medium-high | medium-low) {
                    code-points [aliases] [6 bit-patterns];
                }
            }
            import (classifier-name | default);
        }
    }
    code-point-aliases {
        (dscp | dscp-ipv6 | exp |ieee-802.1 | inet-precedence) {
            alias-name bits;
        }
    }
    drop-profiles {
        profile-name {
            fill-level percentage drop-probability percentage;
            interpolate {
                drop-probability [values];
                fill-level [values]
            }
```

```
        }
    }
    forwarding-classes {
        class class-name {
            priority (high | low);
            queue-num queue-number;
        }
        queue queue-number;
    }
    host-outbound-traffic {
        forwarding-class class-name;
        dscp-code-point value;
    }
    interfaces {
        interface-name  {
            congestion-notification-profile profile-name {
                input {
                    ieee-802.1 {
                        code-point up-bits pfc;
                        }
                    }
                }
            }
            scheduler-map map-name;
            shaping-rate;
            unit (logical-unit-number | * ) {
                classifiers {
                    (dscp | dscp-ipv6|ieee-802.1 | inet-precedence) (classifier-name | default);
                }
                forwarding-class (Forwarding Policy) (class-name | assured-forwarding | best-effort
| expedited-forwarding | network-control) ; {
                    rewrite-rules {
                        (dscp | dscp-ipv6|ieee-802.1 | inet-precedence) (rewrite-rule-name |
default);
                    }
                    classifiers {
                        (dscp | dscp-ipv6|ieee-802.1 | inet-precedence) classifier-name;
                    }
                }
                rewrite-rules {
                    (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) rewrite-rule-name;
                }
            }
```

```
        }
    }
    rewrite-rules {
        (dscp | dscp-ipv6 | exp |ieee-802.1 | inet-precedence) rewrite-rule-name {
            forwarding-class (Forwarding Policy) (class-name | assured-forwarding | best-effort |
expedited-forwarding | network-control) {
                loss-priority (high | low | medium-high | medium-low) code-point (alias | bits);
            }
            import (rewrite-rule-name | default);
        }
    }
    scheduler-maps {
        map-name {
            forwarding-class (Forwarding Policy) (class-name | assured-forwarding | best-effort |
expedited-forwarding | network-control) {
                scheduler scheduler-name;
            }
        }
    }
    schedulers {
        scheduler-name  {
            buffer-size (exact | percent percentage | remainder | temporal);
            drop-profile-map {
                loss-priority (high | low);
                protocol any;
            }
            priority  (low | strict-high);
            shaping-rate (rate | percent percentage);
            transmit-rate (EX Series Switches) (rate | percent percentage | remainder);
        }
    }
    shared-buffer {
        percent percentage;
    }
    traceoptions {
        file (file-name | files files | match match | no-world-readable | size size | world-
readable);
        flag ( all | asynch | chassis-scheduler | cos-adjustment | dynamic | hardware-database |
init | parse | performance-monitor | process | restart | route-socket | show | snmp | util);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Enable tricolor marking.

Options

percent *percentage*—Size of the shared buffer as a percentage of the buffer allocated to the shared buffer pool.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS 11.4 .

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches

Understanding Junos OS CoS Components for EX Series Switches

trickling

IN THIS SECTION

- Syntax | **30153**
- Hierarchy Level | **30153**

- Description | 30153
- Options | 30153
- Required Privilege Level | 30154
- Release Information | 30154

Syntax

```
trickling {
    timeout value;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name]
```

Description

HTTP trickling is a mechanism used to prevent the HTTP client or server from timing-out during a file transfer or during antivirus scanning. HTTP Trickling is time-based and there is only one parameter to configure for this feature, which is the timeout Interval. By default, trickling is disabled.

> **WARNING**: When you enable the trickling option, keep in mind that trickling might send part of a file to the client during its antivirus scan. It is therefore possible that some of the content could be received by the client before the file has been fully scanned.

Options

value—Timeout interval in seconds.

- **Range:** 0 through 600 seconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Statement updated for Sophos support in Junos OS Release 11.1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

trigger

IN THIS SECTION

- Syntax | 30155
- Hierarchy Level | 30155
- Description | 30155
- Default | 30155
- Options | 30155
- Required Privilege Level | 30158
- Release Information | 30158

Syntax

```
trigger trigger-identifier (hold-time hold-time-value | ignore);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify defect triggers.

Default

By default, triggers are ignored.

Options

trigger-identifier—(For M Series, MX Series, SRX Series, and T Series routers only) Trigger identifier. It can be one of the following:

- `oc-lof`—Optical channel Loss of Frame defect trigger.
- `oc-lom`—Optical channel Loss of Multiframe defect trigger.
- `oc-los`—Optical channel Loss of Signal defect trigger.
- `oc-wavelength-lock`—Optical channel Wavelength Lock defect trigger.
- `odu-ais`—Optical channel data unit (ODU) Alarm Indication Signal defect trigger.
- `odu-bbe-th`—ODU Background Block Error Threshold defect trigger.
- `odu-bdi`—ODU Backward Defect Indication defect trigger.
- `odu-bei`—(MX Series routers only) ODU Backward Error Indication defect trigger.
- `odu-es-th`—ODU Errored Seconds Threshold defect trigger.
- `odu-iae`—(MX Series routers only) ODU Incoming Alignment Error defect trigger.
- `odu-lck`—ODU Locked defect trigger.

- `odu-oci`—ODU Open Connection Indication defect trigger.
- `odu-sd`—ODU Signal Degrade defect trigger.
- `odu-ses-th`—ODU Severely Errored Seconds Threshold defect trigger.
- `odu-tca-es`—(MX Series routers only) ODU Errored Seconds Threshold crossing defect trigger.
- `odu-tca-ses`—(MX Series routers only) ODU Severely Errored Seconds Threshold crossing defect trigger.
- `odu-tca-uas`—(MX Series routers only) ODU Unavailable Seconds Threshold crossing defect trigger.
- `odu-ttim`—ODU Trail Trace Identifier Mismatch defect trigger.
- `opu-ptim`—(MX Series routers only) Payload Type Identifier Mismatch defect trigger.
- `odu-uas-th`—ODU Unavailable Seconds Threshold defect trigger.
- `opu-ptm`—Optical Channel Payload (OPU) Payload Type Mismatch defect trigger.
- `otu-ais`—Optical Channel Transport Unit (OTU) Alarm Indication Signal defect trigger.
- `otu-bbe-th`—OTU Background Block Error Threshold defect trigger.
- `otu-bdi`—OTU Backward Defect Indication defect trigger.
- `otu-es-th`—OTU Errored Seconds Threshold defect trigger.
- `otu-fec-deg`—OTU FEC Degrade defect trigger.
- `otu-fec-exe`—OTU FEC Excessive Error defect trigger.
- `otu-iae`—OTU Incoming Alignment defect trigger.
- `otu-sd`—OTU Signal Degrade defect trigger.
- `otu-ses-th`—OTU Severely Errored Seconds Threshold defect trigger.
- `otu-tca-es`—(MX Series routers only) OTU Errored Seconds Threshold crossing defect trigger.
- `otu-tca-ses`—(MX Series routers only) OTU Severely Errored Seconds Threshold crossing defect trigger.
- `otu-tca-uas`—(MX Series routers only) OTU Unavailable Seconds Threshold crossing defect trigger.
- `otu-ttim`—OTU Trail Trace Identifier Mismatch defect trigger.
- `otu-uas-th`—OTU Unavailable Seconds Threshold defect trigger.

trigger-identifier—(For PTX Series routers only) Trigger identifier. It can be one of the following:

- `oc-lof`—Optical channel Loss of Frame defect trigger.
- `oc-lom`—Optical channel Loss of Multiframe defect trigger.
- `oc-los`—Optical channel Loss of Signal defect trigger.
- `oc-tsf`—Optical channel TOE security functionality (TSF) defect trigger.
- `oc-wavelength-lock`—Optical channel Wavelength Lock defect trigger.
- `odu-ais`—ODU Alarm Indication Signal defect trigger.
- `odu-bdi`—ODU Backward Defect Indication defect trigger.
- `odu-bei`—ODU Backward Error Indication defect trigger.
- `odu-iae`—ODU IAE defect trigger.
- `odu-lck`—ODU Locked defect trigger.
- `odu-oci`—ODU Open Connection Indication defect trigger.
- `odu-sd`—ODU Signal Degrade defect trigger.
- `odu-tca-bbe`—ODU Background Block Error Threshold crossing defect trigger.
- `odu-tca-bbe-fe`—ODU far-end Background Block Error (BEI) Threshold crossing defect trigger.
- `odu-tca-es`—ODU Errored Seconds Threshold crossing defect trigger.
- `odu-tca-es-fe`—ODU far-end Errored Seconds Threshold crossing defect trigger.
- `odu-tca-ses`—ODU Severely Errored Seconds Threshold crossing defect trigger.
- `odu-tca-ses-fe`—ODU far-end Severely Errored Seconds Threshold crossing defect trigger.
- `odu-tca-uas`—ODU Unavailable Seconds Threshold crossing defect trigger.
- `odu-tca-uas-fe`—ODU far-end Unavailable Seconds Threshold crossing defect trigger.
- `odu-ttim`—ODU Trail Trace Identifier Mismatch defect trigger.
- `opu-ptim`—Payload Type Identifier Mismatch defect trigger.
- `otu-ais`—OTU Alarm Indication Signal defect trigger.
- `otu-bdi`—OTU Backward Defect Indication defect trigger.
- `otu-fec-deg`—OTU FEC Degrade defect trigger.
- `otu-fec-exe`—OTU FEC Excessive Error defect trigger.

- `otu-iae`—OTU Incoming Alignment defect trigger.
- `otu-sd`—OTU Signal Degrade defect trigger.
- `otu-tca-bbe`—OTU Background Block Error Threshold crossing defect trigger.
- `otu-tca-bbe-fe`—OTU far-end Background Block Error (BEI) Threshold crossing defect trigger.
- `otu-tca-es`—OTU Errored Seconds Threshold crossing defect trigger.
- `otu-tca-es-fe`—OTU far-end Errored Seconds Threshold crossing defect trigger.
- `otu-tca-ses`—OTU Severely Errored Seconds Threshold crossing defect trigger.
- `otu-tca-ses-fe`—OTU far-end Severely Errored Seconds Threshold crossing defect trigger.
- `otu-tca-uas`—OTU Unavailable Seconds Threshold crossing defect trigger.
- `otu-tca-uas-fe`—OTU far-end Unavailable Seconds Threshold crossing defect trigger.
- `otu-ttim`—OTU Trail Trace Identifier Mismatch defect trigger.

hold-time *hold-time-value*—Hold time value. It can be one of the following:

- `down`—Delay before marking interface down when defect occurs (1..65534 milliseconds).
- `up`—Delay before marking interface up when defect is absent (1..65534 milliseconds).

> **NOTE**: The trigger hold time value alone does not mark an interface to be up when the defect is absent or mark an interface to be down when the defect occurs. The hold time value only impacts the alarm reporting time. To mark an interface up or down, you must also configure the physical interface hold time at the `[edit interfaces` *`interface-name`* `hierarchy level]`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

10-Gigabit Ethernet OTN Options Configuration Overview

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

trigger (CoS)

IN THIS SECTION

- Syntax | 30159
- Hierarchy Level | 30159
- Description | 30159
- Options | 30159
- Required Privilege Level | 30160
- Release Information | 30160

Syntax

```
trigger  type  shaping-rate (percent  percentage | rate);
```

Hierarchy Level

```
[edit class-of-service adaptive-shapers adaptive-shaper-name]
```

Description

Specify a trigger type and its associated rate.

Options

type—The type of trigger. Currently, the trigger type can be **becn** only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

shaping-rate (CoS Adaptive Shapers) | 16932

trigger (DHCP Local Server)

IN THIS SECTION

- Syntax | 30160
- Hierarchy Level | 30161
- Description | 30161
- Required Privilege Level | 30161
- Release Information | 30162

Syntax

```
trigger {
    radius-disconnect;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name reconfigure],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
reconfigure],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-
name reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
reconfigure],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group
group-name reconfigure],
[edit system services dhcp-local-server reconfigure],
[edit system services dhcp-local-server dhcpv6 reconfigure],
[edit system services dhcp-local-server group group-name reconfigure],
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
```

Description

Configure behavior in response to a trigger for all DHCP clients or only the DHCP clients serviced by the specified group of interfaces.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

Support at the `[edit ... dhcpv6 ...]` hierarchy levels introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic Reconfiguration of Extended Local Server Clients Overview

Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect

radius-disconnect (DHCP Local Server) | **14003**

trigger (Event Policy)

IN THIS SECTION

- Syntax | **30162**
- Hierarchy Level | **30162**
- Description | **30163**
- Default | **30163**
- Options | **30163**
- Required Privilege Level | **30163**
- Release Information | **30163**

Syntax

```
trigger (on | after | until) event-count;
```

Hierarchy Level

```
[edit event-options policy policy-name within seconds]
```

Description

Configure an event policy to trigger when an event or set of events occurs *event-count* times within a specified time period.

Default

If you do not include this statement, the policy is executed on receipt of the first configured event.

Options

after *event-count*	The policy is executed when the number of matching events received equals *event-count* + 1.
on *event-count*	The policy is executed when the number of matching events received equals *event-count*.
until *event-count*	The policy is executed each time a matching event is received and stops being executed when the number of matching events received equals *event-count*.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Trigger an Event Policy Based on Event Count

trigger-link-failure

IN THIS SECTION

- Syntax | 30164
- Hierarchy Level | 30164
- Description | 30164
- Options | 30164
- Required Privilege Level | 30164
- Release Information | 30165

Syntax

```
trigger-link-failure interface-name;
```

Hierarchy Level

```
[edit interfaces lsq-fpc/pic/port lsq-failure-options]
```

Description

List of SONET interfaces connected to the LSQ interface that can implement Automatic Protection Switching (APS) if the Link Services IQ PIC fails.

Options

interface-name—Name of SONET interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Association between LSQ and SONET Interfaces

trigger-type (LRF Profile)

IN THIS SECTION

- Syntax | 30165
- Hierarchy Level | 30165
- Description | 30166
- Default | 30166
- Options | 30166
- Required Privilege Level | 30166
- Release Information | 30166

Syntax

```
trigger-type  (session-close | volume);
```

Hierarchy Level

```
[edit services lrf profile profile-name template template-name]
```

Description

Configure the type of trigger that causes the generation of data records and transmission to the collector. You can only configure one type of trigger.

Default

If you do not include the `trigger-type` statement, the default trigger is `session-close`.

Options

session-close Use the closing of the data session to cause the generation of data records and transmission to the collector.

volume Use a data volume limit to cause the generation of data records and transmission to the collector. The data volume limit value is configured in the LRF rule.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

trio-flow-offload

IN THIS SECTION

- Syntax | 30167
- Hierarchy Level | 30167
- Description | 30167
- Options | 30167
- Required Privilege Level | 30168
- Release Information | 30168

Syntax

```
trio-flow-offload minimum-bytes minimum-bytes;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Description

Enable any plug-in or daemon on a PIC to generate a request to off-load flows to the Packet Forwarding Engine. This command is available on MX Series routers with Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs).

NOTE: This feature is not supported for Broadband Edge subscribers (given that service PIC off load is not available with aggregate Ethernet (AE).

Options

minimum-bytes—Minimum number of bytes that trigger offloading. When this option is omitted, offloading is triggered when both the forward and reverse flows of the session have begun, meaning that at least one packet has flowed in each direction.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Flow Offloading on MX Series Routers

trust-option-82

IN THIS SECTION

- Syntax | 30168
- Hierarchy Level | 30168
- Description | 30169
- Required Privilege Level | 30169
- Release Information | 30169

Syntax

```
trust-option-82;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group group-name overrides],
[edit logical-systems logical-system-name forwarding-options dhcp-relay overrides],
```

```
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay overrides],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
overrides],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
interface interface-name overrides]
```

Description

Enable processing of DHCP client packets that have a gateway IP address (giaddr) of 0 (zero) and contain option 82 information. By default, the DHCP relay agent treats such packets as if they originated at an untrusted source, and drops them without further processing.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Enable Processing of Untrusted Packets So Option 82 Information Can Be Used

Overriding the Default DHCP Relay Configuration Settings

trusted

IN THIS SECTION

- Syntax | 30170
- Hierarchy Level | 30170
- Description | 30170
- Required Privilege Level | 30170
- Release Information | 30170

Syntax

```
trusted;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security group group-name overrides]
```

Description

Allow DHCP responses from the specified interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Enabling a Trusted DHCP Server (MX Series Routers)

Understanding and Using Trusted DHCP Servers

trusted-ca

IN THIS SECTION

- Syntax | 30171
- Hierarchy Level | 30171
- Description | 30171
- Options | 30172
- Required Privilege Level | 30172
- Release Information | 30172

Syntax

```
trusted-ca ca-profile-name;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
[edit system ntp nts]
```

Description

Identify one or more trusted IPsec certification authorities for system services and Network Time Security (NTS) using Network Time Protocol (NTP)

Use `set system ntp nts trusted-ca trusted-ca-group NTS-ca-group` to specify ca-profiles by name, that are trusted by the NTS clients. You can specify a trusted ca-group (defined under PKI) or ca-profile by name.

This configuration is optional and if it is not specified then, NTP trusts all loaded ca-profiles for NTS. Only clients can configure trusted-ca options.

Options

ca-profile-name—Name of certification authority profile, which is configured at the `[edit security pki]` hierarchy level.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Configuring IPsec Service Sets

trusted-ca (Services)

IN THIS SECTION

- Syntax | 30173
- Hierarchy Level | 30173
- Description | 30173
- Options | 30173
- Required Privilege Level | 30173
- Release Information | 30173

Syntax

```
trusted-ca (all | [ca-profile] );
```

Hierarchy Level

```
[edit services ssl proxy profile profile-name]
[edit services ssl termination profile profile-name]
[edit services ssl initiation profile profile-name]
```

Description

Specify the list of trusted certificate authority profiles. This statement is supported on the SRX1500, SRX5400, SRX5600, and SRX5800 devices, and vSRX Virtual Firewall.

Options

- *trusted-ca-name*—Specify the certificate authority profile name.
- `all`—Select all certificate authority profiles.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

> *Configuring SSL Forward Proxy*
>
> Firewall User Authentication Overview

trusted (DHCP Security)

IN THIS SECTION

- Syntax | 30174
- Hierarchy Level | 30174
- Description | 30174
- Required Privilege Level | 30174
- Release Information | 30175

Syntax

```
trusted;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name overrides]
```

Description

Specify that the interface in this group is trusted. DHCP snooping and DHCPv6 snooping do not apply to the trusted interface, even if the VLAN is enabled for DHCP or DHCPv6 snooping. Likewise, DAI, IP source guard, IPv6 source guard, and IPv6 neighbor discovery inspection—even if they are enabled for the VLAN—do not apply to the interface that is configured with the `overrides` and the `trusted` options. Access interfaces are untrusted by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Enabling a Trusted DHCP Server (ELS)

Understanding DHCP Snooping (non-ELS)

trusted-keys (DNSSEC)

IN THIS SECTION

- Syntax | 30175
- Hierarchy Level | 30175
- Description | 30176
- Options | 30176
- Required Privilege Level | 30176
- Release Information | 30176

Syntax

```
trusted-keys {
    (key dns-key | load-key-file url);
}
```

Hierarchy Level

```
[edit system services dns dnssec]
```

Description

Configure trusted keys in the DNS server.

Options

key ***dns-key*** Specify a DNS trusted key.

load-key-file ***url*** Specifies the URL of the file that contains the DNS trusted keys.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

DNSSEC Overview

Example: Configuring Secure Domains and Trusted Keys for DNSSEC

tti

IN THIS SECTION

- Syntax | **30177**
- Hierarchy Level | **30177**
- Description | **30177**
- Options | **30177**
- Required Privilege Level | **30177**
- Release Information | **30177**

Syntax

```
tti tti-identifier;
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Description

Specify trace identifier options.

Options

tti-identifier—Trace identifier. It can be one of the following:

- `odu-dapi`—Optical Channel Data Unit (ODU) Destination Access Point Identifier.
- `odu-expected-receive-dapi`—ODU Expected Receive Destination Access Point Identifier.
- `odu-expected-receive-sapi`—ODU Expected Receive Source Access Point Identifier.
- `odu-sapi`—ODU Source Access Point Identifier.
- `otu-dapi`—Optical Channel Transport Unit (OTU) Destination Access Point Identifier.
- `otu-expected-receive-dapi`—OTU Expected Receive Destination Access Point Identifier.
- `otu-expected-receive-sapi`—OTU Expected Receive Source Access Point Identifier.
- `otu-sapi`—OTU Source Access Point Identifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

10-Gigabit Ethernet OTN Options Configuration Overview

100-Gigabit Ethernet OTN Options Configuration Overview

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

ttl

IN THIS SECTION

- Syntax | 30178
- Hierarchy Level | 30178
- Description | 30178
- Options | 30178
- Required Privilege Level | 30179
- Release Information | 30179

Syntax

```
ttl value;
```

Hierarchy Level

```
[edit interfaces interface-name unit number tunnel]
```

Description

Set the time-to-live value bit in the header of the outer IP packet.

Options

value—Time-to-live value.

- **Range:** 0 through 255
- **Default:** 64

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tunnel Services Overview

ttl (Dynamic Flow Capture)

IN THIS SECTION

- Syntax | 30179
- Hierarchy Level | 30180
- Description | 30180
- Options | 30180
- Required Privilege Level | 30180
- Release Information | 30180

Syntax

```
ttl hops;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Configure the time-to-live (TTL) value for the IP-IP header.

Options

hops—TTL value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Content Destination

ttl (RPM probe)

IN THIS SECTION

- Syntax | 30181
- Junos OS Hierarchy Level | 30181
- Junos OS Evolved Hierarchy Level | 30181
- Description | 30181

- Options | 30182
- Required Privilege Level | 30182
- Release Information | 30183

Syntax

```
ttl hop-count;
```

Junos OS Hierarchy Level

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name ]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the maximum hop count (TTL) for all types of probes (IPv4 and IPv6) in real-time performance monitoring (RPM) and Two-Way Active Management Protocol (TWAMP). This can be useful when it is necessary to restrict the scope of the RPM probes so a probe does not unintentionally monitor an alternative path to the destination (such as may occur after a BGP re-routing). Another example is to monitor direct reachability by specifying a TTL of 1. Probes that exceed the number set for TTL are discarded.

The TTL configuration is supported on Routing Engine-based RPM, Multiservices Modular PIC Concentrator (MS-MPC) and Multiservices Modular Interfaces Card (MS-MIC)-based RPM, and Two-Way Active Management Protocol (TWAMP).

You can set a TTL value for the probe types listed below.

Software time stamping:

- icmp-ping, and icmp-ping-timestamp

- icmp6-ping (Junos OS only)
- udp-ping, and udp-ping-timestamp
- tcp-ping (Junos OS only)
- http-get, and http-metadata-get (Junos OS only)
- BGP neighbor monitoring using TCP/UDP (Junos OS only)

Hardware time stamping:

- icmp-ping and icmp-ping-timestamp
- udp-ping and udp-ping-timestamp

MS-MPC-PIC based probes (delegate):

- icmp-ping, icmp-ping-timestamp, and icmp6-ping (Junos OS only),

MS-MPC-PIC hardware timestamp:

- icmp-ping and icmp-ping-timestamp
- udp-ping and udp-ping-timestamp

TWAMP probe:

- inline TWAMP client

Options

hop-count—Prior to Junos OS Release 18.2R1, for RPM, the RPM client always sent a TTL of 64 to the RPM server under the IPv4 or IPv6 header. For TWAMP clients, the TTL was 255 sent in the IPv4 header. In Junos OS Release 18.2R1 and later, you can specify the TTL you want for RPM and TWAMP probes.

- **Range:** 1 through 255 for both RPM and TWAMP
- **Default:** 64

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

ttl-threshold

IN THIS SECTION

- Syntax | 30183
- Hierarchy Level | 30183
- Description | 30184
- Options | 30184
- Required Privilege Level | 30184
- Release Information | 30184

Syntax

```
ttl-threshold number;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Specify the traceroute time-to-live (TTL) threshold value. This value sets the acceptable level of network penetration for trace routing.

Options

number—TTL threshold value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

ALG Descriptions

Configuring the TTL Threshold

Examples: Configuring Application Protocols

Verifying the Output of ALG Sessions

tunable-name

IN THIS SECTION

- Syntax | 30185
- Hierarchy Level | 30185
- Description | 30185
- Options | 30185
- Required Privilege Level | 30185

Release Information | 30185

Syntax

```
tunable-name  tunable-name {
    tunable-value protocol-value;
}
```

Hierarchy Level

```
[edit security idp sensor-configuration detector protocol-name protocol-name]
```

Description

Specify the name of the tunable parameter to enable or disable the protocol detector for each of the service. By default, the protocol decoders for all services are enabled.

Options

tunable-name—Name of the specific tunable parameter.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. Support.

tunable-value

IN THIS SECTION

- Syntax | 30186
- Hierarchy Level | 30186
- Description | 30186
- Options | 30186
- Required Privilege Level | 30186
- Release Information | 30187

Syntax

```
tunable-value protocol-value;
```

Hierarchy Level

```
[edit security idp sensor-configuration detector protocol-name protocol-name tunable-name
tunable-name]
```

Description

Specify the value of the tunable parameter to enable or disable the protocol detector for each of the services.

Options

tunable-value—Integer representing a selected option for the switch specified in `tunable-name`. The range of values depends on the options defined for the specified switch.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2. Support.

tunables

IN THIS SECTION

- Syntax | 30187
- Hierarchy Level | 30187
- Description | 30187
- Options | 30188
- Required Privilege Level | 30188
- Release Information | 30188

Syntax

```
tunables {
    drop-on-zone-mismatch;
    enable-logging;
    max-route-change value;
    no-l3l4-app-lookup
}
```

Hierarchy Level

[edit security advance-policy-based-routing]

Description

Configure the advanced policy-based (APBR) routing options to streamline the traffic handling.

You can streamline the traffic handling with APBR such as limiting the number of times a route can change for a session, terminating the session if there is a mismatch between zones when APBR is being applied in the middle of the session, and enabling logging to record events that occur on the device.

Fine-tuning the APBR configuration is required to avoid the possible issues such as excessive transitions due to route changes.

Options

drop-on-zone-mismatch Terminate the session instead of instead of allowing traffic to traverse through the same route bypassing APBR.

enable-logging Enable logging to record events that occur on the device for APBR-related operations.

max-route-change *value* Configure the threshold for limiting the number of times a route can change for a session.

- **Range:** 0-5
- **Default:** 1

no-l3l4-app-lookup Disable Layer 3 and Layer 4 custom application lookup in APBR. Starting in Junos OS Release 20.2R1, APBR supports Layer 3 and Layer 4 custom applications. You can manually disable Layer 3 and Layer 4 custom application lookup by including the `no-l3l4-app-lookup option`.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D110.

RELATED DOCUMENTATION

Example: Configuring Advanced Policy-Based Routing Policies

Advanced Policy-Based Routing

tunnel

IN THIS SECTION

- Syntax | 30189
- Hierarchy Level | 30189
- Description | 30189
- Required Privilege Level | 30189
- Release Information | 30190

Syntax

```
tunnel {
    destination destination-address;
    source source-address;
    ttl number;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Generic Routing Encapsulation Tunneling

tunnel

IN THIS SECTION

- Syntax | 30190
- Syntax (SRX Series) | 30191
- Hierarchy Level | 30192
- Description | 30192
- Required Privilege Level | 30192
- Release Information | 30192

Syntax

```
tunnel {
    backup-destination destination-address;
    destination destination-address;
    encapsulation {
        gre
            destination {
                addressaddress;
            }
            key key;
            source {
                addressaddress;
            }
            tunnel-routing-instance {
```

```
                routing-instance routing-instance;
            }

        }
        ipip {
            destination {
                address address;
            }
            source {
                address address;
                interface interface;
            }
        }
    }
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;
}
```

Syntax (SRX Series)

```
tunnel {
     destination;
destination-hostname;
    do-not-fragment;
    encap-type;
    flow-label;
    no-path-mtu-discovery;
    path-mtu-discovery;
 routing-instance {
        destination routing-instance-name;
    }
            source;
source-hostname;
           source-interface ;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic or virtual private networks (VPNs).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Encryption Interfaces

Tunnel Services Overview

Junos OS VPNs Library for Routing Devices

tunnel

IN THIS SECTION

- Syntax | 30193
- Hierarchy Level | 30193
- Description | 30193

- Options | 30193
- Required Privilege Level | 30194
- Release Information | 30194

Syntax

```
tunnel {
    allow-fragmentation;
    backup-destination address;
    destination destination-address;
    do-not-fragment;
    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    traffic-class traffic-class-value;
    ttl number;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic or virtual private networks (VPNs).

The remaining statements are explained separately. See CLI Explorer.

Options

`encapsulation ipip ( source | address)` —(Optional) Configure IP-IP encapsulation type. To configure an IP-in-IP tunnel at an FTI, use the `encapsulation ipip` option at the [`edit interfaces interface-name unit logical-`

`unit-number tunnel`] hierarchy level. The `source` and `destination` options support IPv4 addresses. You can also specify interface name instead of source address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Encryption Interfaces

tunnel-attributes

IN THIS SECTION

- Syntax | 30194
- Hierarchy Level | 30195
- Description | 30195
- Options | 30195
- Required Privilege Level | 30196
- Release Information | 30196

Syntax

```
tunnel-attributes name{
    dynamic-tunnel-source-prefix dynamic-tunnel-source-prefix;
    dynamic-tunnel-type V4oV6;
    dynamic-tunnel-mtu dynamic-tunnel-mtu;
```

```
    dynamic-tunnel-reassembly;
    dynamic-tunnel-anchor-pfe dynamic-tunnel-anchor-pfe
    dynamic-tunnel-anti-spoof (off | on);
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options dynamic-tunnels],
[edit logical-systems name routing-options dynamic-tunnels],
[edit routing-instances name routing-options dynamic-tunnels],
[edit routing-options dynamic-tunnels]
```

Description

Define dynamic tunnel attributes for transporting IPv4 traffic over an IPv6 network. You can configure multiple tunnels and specify different attributes for each tunnel. Service providers with IPv6 infrastructure can configure individual tunnels for each customer to route IPv4 traffic.

Options

name

Specify a dynamic tunnel name. You can have multiple tunnels configured with different attributes.

dynamic-tunnel-anchor-pfe

Specify a dynamic tunnel anchor PFE name of format pfe-x/y/z.

dynamic-tunnel-anti-spoof

Enable or disable anti-spoofing check.

- Values:
  - `off`—Disable anti-spoofing check
  - `on`—Enable anti-spoofing check.
  - **Default:** Dynamic tunnel anti-spoof is enabled by default.

dynamic-tunnel-mtu

Specify the dynamic tunnel maximum transmission unit (MTU) value

- **Range:** `296` through `9232`

dynamic-tunnel-source-prefix

Specify the tunnel source address

dynamic-tunnel-type

Specify the tunnel type

- Values:
  - `V4oV6`— Tunnel type is IPv4 over IPv6

dynamic-tunnel-reassembly Enable IPV6 fragment reassembly for forwarding Ipv4 traffic

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

dynamic-tunnels | 6161

extended-nexthop | 6700

Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP

tunnel-devices (Protocols AMT)

IN THIS SECTION

- Syntax | 30197
- Hierarchy Level | 30197
- Description | 30197
- Default | 30197
- Options | 30197
- Required Privilege Level | 30198
- Release Information | 30198

Syntax

```
tunnel-devices [ ud-fpc/pic/port ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols amt relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols amt
relay],
[edit protocols amt relay],
[edit routing-instances routing-instance-name protocols amt relay]
```

Description

List one or more tunnel-capable Automatic Multicast Tunneling (AMT) PICs to be used for creating multicast tunnel (**ud**) interfaces. Creating an AMT PIC list enables you to control the load-balancing implementation.

Tunnel-capable PICs include DPC and MPC.

The physical position of the PIC in the routing device determines the multicast tunnel interface name.

Default

Multicast tunnel interfaces are created on all available tunnel-capable AMT PICs, based on a round-robin algorithm.

Options

ud-*fpc/pic/port*—Interface that is automatically generated when a tunnel-capable PIC is installed in the routing device.

> **NOTE**: Each **tunnel-devices** statement keyword is optional. By default, all configured tunnel devices are used. The keyword selects the subset of configured tunnel devices.
>
> Tunnel devices must be configured on MX Series routers. They are not automatically available like M Series routers that have dedicated PICs. On MX Series routers, the tunnel device port is the next highest number after the physical ports – a PIC created

with the `tunnel-services` statement at the `[edit chassis fpc slot-number pic number]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

Example: Configuring the AMT Protocol

tunnel-devices (Tunnel-Capable PICs)

IN THIS SECTION

- Syntax | 30199
- Hierarchy Level | 30199
- Description | 30199
- Default | 30199
- Options | 30199
- Required Privilege Level | 30200
- Release Information | 30200

Syntax

```
tunnel-devices [ mt-fpc/pic/port ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols pim],
[edit routing-instances instance-name protocols pim]
```

Description

List one or more tunnel-capable PICs to be used for creating multicast tunnel (**mt**) interfaces. Creating a PIC list enables you to control the load-balancing implementation.

Tunnel-capable PICs include:

- Adaptive Services PIC
- Multiservices PIC or Multiservices DPC
- Tunnel Services PIC
- On MX Series routers, a PIC created with the `tunnel-services` statement at the `[edit chassis fpc slot-number pic number]` hierarchy level.

The physical position of the PIC in the routing device determines the multicast tunnel interface name. For example, if you have an Adaptive Services PIC installed in FPC slot 0 and PIC slot 0, the corresponding multicast tunnel interface name is **mt-0/0/0**. The same is true for Tunnel Services PICs, Multiservices PICs, and Multiservices DPCs.

Default

Multicast tunnel interfaces are created on all available tunnel-capable PICs, based on a round-robin algorithm.

Options

mt-*fpc/pic/port*—Interface that is automatically generated when a tunnel-capable PIC is installed in the routing device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Load Balancing Multicast Tunnel Interfaces Among Available PICs

tunnel-end-point

IN THIS SECTION

- Syntax | 30200
- Hierarchy Level | 30201
- Description | 30201
- Options | 30202
- Required Privilege Level | 30204
- Release Information | 30204

Syntax

```
tunnel-end-point tunnel-name {
    gre {
        key authentication key;
    }
    ipv4 {
        destination-address destination-host-address;
        source-address source-host-address;
```

```
    }
    ipv6 {
        destination-address destination-host-address;
        source-address source-host-address;
    }
}
```

Hierarchy Level

```
[edit firewall]
```

Description

The `tunnel-end-point` command enables line-rate, filter-based, GRE tunneling of IPv4 and IPv6 payloads across IPv4 networks for MX Series routers running Trio-based FPCs (including MX80, MX104 and MX204). Filter-based tunneling encapsulates the original passenger protocol packet in an outer packet header. For example, for filter-based tunneling across IPv4 networks, the header adds 24 bytes or 28 bytes of overhead, including 20 bytes of IPv4 header. Either IPv4 or IPv6 traffic can be the transport protocol. For outgoing packets that match the configured filter term, the original packet are encapsulated inside an IP+GRE header as specified by the tunnel definition. IP lookup is performed on the outer header, and the packets are forwarded accordingly.

The route lookup for GRE encapsulated traffic is supported on the default routing instance only. GRE encapsulation is not supported for logical systems, or for MPLS traffic.

When an subnet range is configured for either the IPv4 or IPv6 option, traffic between hosts in the range is load balanced.

Note that the device must be enabled for `enhanced-mode` to support the use of GRE tunnel templates, which allows you to define tunnel attributes.

To use the feature with PTX Series routers, install a PTX Series router as an *encapsulator*, that is, an ingress PE router where you can reference a tunnel template name in an type `inet` or `inet6` ingress firewall filter by configuring the encapsulate terminating action.

Configure the tunnel end points as shown here (for IPv4):

- `set firewall tunnel-end-point tunnel-name ipv4 source-address source-host-address`
- `set firewall tunnel-end-point tunnel-name ipv4 destination-address destination-host-address`
- `set firewall tunnel-end-point tunnel-name ipv4 destination-address destination-host-address`

Note that the maximum number of /25 IPv4 or /123 IPv6 subnets allowed for a tunnel-endpoint destination addresses is 64.

An interface-specific encapsulating output filter action is also required. It triggers the Packet Forwarding Engine to use information in the specified tunnel template to encapsulate matching packets and forward the resulting GRE packets. GRE encapsulation is supported only for outgoing IPv4 unicast and IPv6 unicast traffic.

Configure the tunnel end points as shown here (for IPv4):

- `set firewall family inet filter filter-name term term-name then encapsulate tunnel-name`

For the GRE decapsulation with PTX Series routers, use a PTX3000 or PTX5000 router with third generation FPCs that is running Junos OS Release 16.1R2 or later and configure the firewall filter at the hierarchy level shown here:

- `set firewall family inet filter filter-name term term-name then decapsulate gre`

Egress sampling is supported on GRE encapsulated packets, but note that output filter match conditions only work on the contents of ingress packets.

The route lookup for GRE encapsulated traffic is supported on the default routing instance only. GRE encapsulation is supported for ingress IPv4 unicast and IPv6 unicast traffic. It is not supported for logical systems, or for MPLS traffic.

When defining the tunnel end point, or the prefix list, be sure to specify the /32 route. Multiple tunnel end point source-address are not supported.

- A maximum of 1024 tunnel templates is supported. You can configure or change up to 512 tunnel templates at a time
- A maximum of 64 tunnel end point destination addresses are supported in a given tunnel template. When more than one destination IP address exists, the one used for the outer header is based on a hash that is computed on the input packet from the input interface.

Options

gre

- Encapsulation protocol. Options include `gre` and `gre-in-udp`.

  You must also specify whether the tunnel is IPv4 or IPv6. A IPv4 example using GRE follows.

```
tunnel-end-point name {
    ipv4 {
        source-address 10.255.1.1;
```

```
        destination-address 10.255.2.0/25;
    }
    gre {
        key 9;
    }
}
```

- key *number*—An integer value that uniquely identifies a GRE IPv4 tunnel if multiple traffic flows share the same *source-address* and *destination-address* pair. Range: 1 through 0xFFFFFFFF (4,294,967,295 decimal). If a tunnel definition specifies GRE IPv4 tunneling using a key, the system includes the key in the GRE header whenever a Packet Forwarding Engine is instructed to use that tunnel definition to encapsulate a packet.

gre-in-udp For MX Series routers; specify if the tunnel is **gre-in-udp**.

- *destination-port number*—An integer value that uniquely identifies the UDP destination port. Range: 1 through 65535.
- *key number*—An integer value that uniquely identifies the **gre-in-udp** tunnel if multiple traffic flows share the same destination-port and source-port. Range: 1 through 0xFFFFFFFF (4,294,967,295 decimal). If you include a key in the tunnel definition to encapsulate packets, the key used is the one in the GRE header.
- *source-port number*—An integer value that uniquely identifies the UDP source port. Range: 1 through 65535.

ipv4 The IP network protocol used to transport encapsulated passenger protocol packets; IPv4 transports IPv4 packets encapsulated using filter-based GRE. The default prefix length is 32; the supported range is from 25 to 32. When specified, traffic is load-balanced to the hosts on this subnet.

ipv6 The IP network protocol used to transport encapsulated passenger protocol packets; IPv6 transports IPv6 packets encapsulated using filter-based GRE. The default prefix length is 128; the supported range is from 121 to 128. When specified, traffic is load-balanced to the hosts on this subnet.

source-address IP address of the encapsulator (the local ingress PE router). Multiple tunnel end point source-address are not supported.

destination-address IP address or address range of the decapsulator (the remote egress PE router). For both IPv4 and IPv6, a maximum of 64 /25 IPv4 or /123 IPv6 subnets can be configured for the end point destination address.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

- *Components of Filter-Based Tunneling Across IPv4 Networks*
- *Understanding Filter-Based Tunneling Across IPv4 Networks*
- *Firewall Filter Terminating Actions*
- interface-group (Decapsulate GRE) | **8892**
- enhanced-mode (Network Services)

tunnel (fti)

IN THIS SECTION

- Syntax | **30204**
- Hierarchy Level | **30205**
- Description | **30205**
- Required Privilege Level | **30205**
- Release Information | **30205**

Syntax

```
tunnel {
    encapsulation {
            ( gre  | ipip | udp );
```

```
            destination {
                address address;
            }
            key key; key;
            source {
                address address;
            }
            tunnel-routing-instance {
                routing-instance routing-instance;
            }
            tunnel-termination;
            virtual-loop;
    }
}
```

Hierarchy Level

```
[edit interfaces fti number unit logical-unit-number]
```

Description

Configure a tunnel on an FTI interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.2R1.

Statement introduced for IPIP tunnels in Junos OS Evolved Release 20.4R1.

Statement introduced for UDP tunnels in Junos OS Evolved Release 21.1R1.

`virtual-loop` option introduced in Junos OS Evolved Release 24.4R1.

tunnel-group

IN THIS SECTION

- Syntax | 30206
- Hierarchy Level | 30207
- Description | 30207
- Options | 30207
- Required Privilege Level | 30207
- Release Information | 30207

Syntax

```
tunnel-group group-name {
    aaa-access-profile profile-name;
    dynamic-profile profile-name;
    hello-interval seconds;
    hide-avps;
    l2tp-access-profile profile-name;
    local-gateway address {
        address address;
        gateway-name gateway-name;
    }
    maximum-send-window packets;
    maximum-sessions number;
    ppp-access-profile profile-name;
    receive-window packets;
    retransmit-interval seconds;
    service-device-pool pool-name;
    service-interface interface-name;
    service-profile profile-name(parameter)&profile-name;
    syslog {
        host hostname {
            services severity-level;
            facility-override facility-name;
            log-prefix prefix-value;
        }
```

```
    }
    tos-reflect;
    tunnel-switch-profile profile-name;
    tunnel-timeout seconds;
}
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Specify the L2TP tunnel properties.

> **NOTE**: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

group-name—Identifier for the tunnel group.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for MX Series routers introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring L2TP Tunnel Groups

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

tunnel-inspection

IN THIS SECTION

- Syntax | 30208
- Hierarchy Level | 30209
- Description | 30209
- Options | 30209
- Required Privilege Level | 30210
- Release Information | 30210

Syntax

```
tunnel-inspection
{
   inspection-profile profile-name
    {
      vxlan vxlan-tunnel-name
      {
        policy-set  policy-set;
        vni vni-name;
       }
     }
   traceoptions (Security)
    {
      file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
      flag <name>;
      no-remote-trace;
```

```
        }

      vni vni-name
      {
        (vni-id [ vni-id-value ... ] |
         vni-range range-low
            {
               to range-high
            });
      }
    }
```

Hierarchy Level

```
[edit security]
```

Description

Configure security inspection for EVPN-VXLAN tunnel traffic. Configure an outer policy for the outer header and an inner policy for the inner header.

When a packet matches a security policy, the firewall decapsulates the packet to get the inner header. The tunnel inspection profile is applied to the permitted traffic. With the inner packet content and the applied tunnel inspection profile, the device performs a policy lookup and then performs stateful inspection of the inner session traffic.

Options

- `inspection-profile` *`profile-name`*—Configure a tunnel inspection profile to connect the outer policy and inner policy.
- `trace-option`—Trace options for Network Security tunnel-inspection. Configure trace option for tunnel inspection.
- `vni` *`vni-name`*—Define VXLAN network identifier (VNI). VXLAN network identifier (VNI) used to uniquely identify the VXLAN.
- `no-remote-trace`—Disable remote tracing

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

Security—To view this statement in the configuration.

Security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

CLI Explorer

tunnel (L2TP)

IN THIS SECTION

- Syntax | 30210
- Hierarchy Level | 30211
- Description | 30211
- Required Privilege Level | 30211
- Release Information | 30211

Syntax

```
tunnel {
    assignment-id-format (assignment-id | client-server-id);
    failover-resync (failover-protocol | silent-failover);
    idle-timeout seconds;
    maximum-sessions number;
    minimum-retransmission-timeout;
    name name {
        address ip-address {
```

```
            drain;
            routing-instance routing-instance-name {
                drain;
            }
        }
        drain;
    }
    nas-port-method;
    retransmission-count-established count;
    retransmission-count-not-established count;
    rx-window-size packets;
     tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port | reject | reject-ip-
address | reject-udp-port);
}
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Configure L2TP tunnel characteristics.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

`rx-window-size` option added in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP LAC Subscriber Configuration

L2TP LNS Inline Service Interfaces

tunnel-limit (Protocols AMT)

IN THIS SECTION

- Syntax | 30212
- Hierarchy Level | 30212
- Description | 30212
- Options | 30213
- Required Privilege Level | 30213
- Release Information | 30213

Syntax

```
tunnel-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols amt relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols amt
relay],
[edit protocols amt relay],
[edit routing-instances routing-instance-name protocols amt relay]
```

Description

Limit the number of Automatic Multicast Tunneling (AMT) data tunnels created. The system might reach a dynamic upper limit of tunnels of all types before the static AMT limit is reached.

Options

number—Maximum number of data AMTs that can be created on the system.

- **Range:** 0 through 4294967295
- **Default:** 1 tunnel

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring the AMT Protocol

tunnel-limit (Routing Instances)

IN THIS SECTION

- Syntax | 30214
- Hierarchy Level | 30214
- Description | 30214
- Options | 30214
- Required Privilege Level | 30214
- Release Information | 30214

Syntax

```
tunnel-limit limit;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
mdt],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel family inet | inet6mdt],
[edit routing-instances routing-instance-name protocols pim mdt],
[edit routing-instances routing-instance-name provider-tunnel family inet | inet6 mdt]
```

Description

Limit the number of data MDTs created in this VRF instance. If the limit is 0, then no data MDTs are created for this VRF instance.

Options

limit—Maximum number of data MDTs for this VRF instance.

- **Range: 0** through **1024**
- **Default: 0** (No data MDTs are created for this VRF instance.)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4. In Junos OS Release 17.3R1, the `mdt` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default MDT in Rosen 7, and data MDT for Rosen 6 and Rosen 7. The `provider-tunnel mdt` hierarchy is now hidden for backward compatibility with existing scripts.

RELATED DOCUMENTATION

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

tunnel-limit (Routing Instances Provider Tunnel Selective)

IN THIS SECTION

- Syntax | 30215
- Hierarchy Level | 30215
- Description | 30215
- Options | 30216
- Required Privilege Level | 30216
- Release Information | 30216

Syntax

```
tunnel-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective],
[edit routing-instances routing-instance-name provider-tunnel selective]
```

Description

Specify a limit on the number of selective tunnels that can be created for an LSP. This limit can be applied to the following types of selective tunnels:

- Ingress replication tunnels

- LDP-signaled LSP
- LDP point-to-multipoint LSP
- PIM-SSM provider tunnel
- RSVP-signaled LSP
- RSVP-signaled point-to-multipoint LSP

Options

number—Specify the tunnel limit.

- **Range:** 0 through 1024

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

selective | **16529**

wildcard-source (Selective Provider Tunnels) | **31120**

tunnel-mtu

IN THIS SECTION

- Syntax | **30217**
- Hierarchy Level | **30217**

- Description | 30217
- Range | 30217
- Required Privilege Level | 30217
- Release Information | 30217

Syntax

```
tunnel-mtu;
```

Hierarchy Level

```
[edit security ipsec vpn hub-to-spoke-vpn tunnel-mtu tunnel-mtu]
```

Description

Tunnel MTU is the maximum size of transmit packet for IPsec tunnels. The minimum Tunnel MTU you can configure for IPv6 is 1390.

Range

Range - The packet size of the minimum tunnel value ranges from 256 to 9192 .

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.3R1

RELATED DOCUMENTATION

Next Gen Services Overview

tunnel-mtu (Services IPsec VPN)

IN THIS SECTION

- Syntax | 30218
- Hierarchy Level | 30218
- Description | 30218
- Options | 30218
- Required Privilege Level | 30219
- Release Information | 30219

Syntax

```
tunnel-mtu bytes;
```

Hierarchy Level

```
[edit services ipsec-vpn rule rule-name term term-name then]
```

Description

Maximum transmission unit (MTU) size for IPsec tunnels. This defines the maximum size of an IP packet, including the IPsec overhead.

Options

bytes—MTU size.

- **Default:** 1500 bytes

- **Range:** 256 through 9192 bytes

> **NOTE**: Clear the **IPsec SA** in **tunnel-mtu** to accomodate Jumbo frames larger than 1500 bytes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

tunnel-mtu (Services Service Set)

IN THIS SECTION

- Syntax | 30219
- Hierarchy Level | 30220
- Description | 30220
- Options | 30220
- Required Privilege Level | 30220
- Release Information | 30220

Syntax

```
tunnel-mtu bytes;
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Maximum transmission unit (MTU) size for IPsec tunnels. This statement is useful for dynamic endpoint tunnels for which you cannot configure the `tunnel-mtu` statement at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level.

For static IPsec tunnels, this statement sets the tunnel MTU value for all the tunnels within this service set. If you need a specific value for a particular tunnel, then set the `tunnel-mtu` statement at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level.

> **NOTE**: The `tunnel-mtu` setting at the `[edit services ipsec-vpn rule rule-name term term-name then]` hierarchy level overrides the value specified at the `[edit services service-set service-set-name ipsec-vpn-options]` hierarchy level.

Options

bytes—MTU size.

- **Default:** 1500 bytes
- **Range:** 256 through 9192 bytes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

mtu (Multilink and Link Services Logical Interface) | **11025**

Configuring IPsec Service Sets

Configuring IPsec Rules

tunnel-observation

IN THIS SECTION

- Syntax | 30221
- Hierarchy Level | 30221
- Description | 30221
- Options | 30222
- Required Privilege Level | 30222
- Release Information | 30222

Syntax

```
tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Specify the types of MPLS flows on which to enable inline flow monitoring. If you do not configure a `tunnel-observation` type, then plain MPLS flow records are created.

You can configure multiple values for `tunnel-observation`. Flows are created for only the deepest match. For example, if you configure both `ipv4` and `mpls-over-udp` and the traffic type is MPLS-over-UDP, flows are created for MPLS-over-UDP. If you configure `ipv4` but *do not* configure `mpls-over-udp` and the traffic type is MPLS-over-UDP, flows are created for MPLS-IPv4.

If the MPLS traffic type does *not* match any of the `tunnel-observation` values, then plain MPLS flows are created.

If you do not configure `tunnel-observation`, plain MPLS flows are created.

If the `tunnel-observation` statement is added or deleted, or if the configured value is changed, all flows related to the old template will be deleted and replaced by new flows using the changed template.

Options

ipv4 Enable flow monitoring for MPLS-IPv4 traffic. You must also configure `mpls-template` at the `[edit services flow-monitoring (version9 | version-ipfix) template` *`template-name`*`]` hierarchy level.

ipv6 Enable flow monitoring for MPLS-IPv6 traffic. You must also configure `mpls-template` at the `[edit services flow-monitoring (version9 | version-ipfix) template` *`template-name`*`]` hierarchy level. If you are running inline flow monitoring on a Lookup (LU) card on an MX Series router, you must also configure `use-extended-flow-memory` at the `[edit chassis fpc` *`slot-number`* `inline-services]` hierarchy level to create MPLS-IPv6 flow records.

mpls-over-udp (PTX Series only) Enable flow monitoring for MPLS-over-UDP traffic. Monitoring looks past the tunnel header to report the inner payload of the packets. For an MPLS-over-UDP flow that is carried between IPv4 endpoints, you must also configure `ipv4-template` at the `[edit services flow-monitoring (version9 | version-ipfix) template` *`template-name`*`]` hierarchy level. For an MPLS-over-UDP flow that is encapsulated in an RSVP-TE LSP, you must also configure `mpls-ipvx-template` in Junos OS Release 18.1 or `mpls-template` starting in Junos OS 18.2R1 at the `[edit services flow-monitoring (version 9 | version-ipfix) template` *`template-name`*`]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

`ipv4` and `ipv6` options added in Junos OS Release 18.2R1.

Statement introduced in Junos OS Release 18.4R1 on MX Series routers.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring on PTX Series Routers

Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

tunnel-port

IN THIS SECTION

- Syntax | 30223
- Hierarchy Level | 30223
- Description | 30223
- Required Privilege Level | 30223
- Release Information | 30224

Syntax

```
tunnel-port port-number tunnel-services;
```

Hierarchy Level

```
[edit chassis fpc slot pic pic-number]
```

Description

Configure the port number for generic routing encapsulation (GRE) tunneling.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Generic Routing Encapsulation Tunneling

tunnel-port (GRE)

IN THIS SECTION

- Syntax | 30224
- Hierarchy Level | 30224
- Description | 30224
- Required Privilege Level | 30225
- Release Information | 30225

Syntax

```
tunnel-port port-number
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number tunnel-services]
```

Description

For MX Series 5G Universal Routing Platforms, configure the port number of the tunnel. The range is 0 through 3.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

tunnel-services (Chassis) | 30234

bandwidth (Tunnel Services) | 2880

tunnel-profile (Domain Map)

IN THIS SECTION

- Syntax | 30225
- Hierarchy Level | 30226
- Description | 30226
- Options | 30226
- Required Privilege Level | 30226
- Release Information | 30226

Syntax

```
tunnel-profile profile-name;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Tunnel profile that provides definitions for tunnels associated with the domain map.

Options

profile-name—Name of tunnel profile.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Mapping Subscriber Domains to Access and Session Options

Configuring a Tunnel Profile for Subscriber Access

tunnel-profile (L2TP Tunnel Switching)

IN THIS SECTION

- Syntax | 30227
- Hierarchy Level | 30227
- Description | 30227

- Options | 30227
- Required Privilege Level | 30227
- Release Information | 30227

Syntax

```
tunnel-profile profile-name;
```

Hierarchy Level

```
[edit access tunnel-switch-profile profile-name]
```

Description

Specify the name of an L2TP tunnel profile that defines the tunnel to which PPP subscriber traffic is switched.

Options

profile-name—Unique name that identifies the tunnel profile; configured with the `tunnel-profile` statement at the `[edit access]` hierarchy level.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

tunnel-profile (Tunnel Profile)

IN THIS SECTION

- Syntax | 30228
- Hierarchy Level | 30229
- Description | 30229
- Options | 30229
- Required Privilege Level | 30229
- Release Information | 30229

Syntax

```
tunnel-profile profile-name {
    tunnel tunnel-id {
        identification name;
        logical-system logical-system-name;
        max-sessions number;
        medium type;
        nas-port-method cisco-avp;
        preference number;
        remote-gateway {
            address server-ip-address;
            gateway-name server-name;
        }
        routing-instance routing-instance-name;
        secret password;
        source-gateway {
            address client-ip-address;
            gateway-name client-name;
        }
        type tunnel-type;
```

```
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Define a tunnel profile for subscriber access.

Options

profile-name—Unique name that identifies the tunnel profile. The profile can be referenced from within a domain map or by the RADIUS Tunnel-Group VSA [26-64].

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

tunnel-queuing

IN THIS SECTION

- Syntax | 30230
- Hierarchy Level | 30230
- Description | 30230
- Required Privilege Level | 30230
- Release Information | 30230

Syntax

```
tunnel-queuing;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number]
```

Description

Enable class-of-service (CoS) queuing for generic routing encapsulation (GRE) and IP-IP tunnels.

NOTE: The `tunnel-queuing` option is not supported in chassis cluster mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Release 9.0 of Junos OS.

tunnel (Security Policies)

IN THIS SECTION

- Syntax | 30231
- Hierarchy Level | 30231
- Description | 30231
- Options | 30231
- Required Privilege Level | 30232
- Release Information | 30232

Syntax

```
tunnel {
    ipsec-group-vpn group-vpn;
    ipsec-vpn vpn-name;
    pair-policy pair-policy;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit]
```

Description

Encapsulate outgoing IP packets and decapsulate incoming IP packets.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. The `ipsec-group-vpn` option added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Security Policies Overview

Understanding Security Policy Rules

Understanding Security Policy Elements

tunnel (Security Screen)

IN THIS SECTION

- Syntax | 30232
- Hierarchy Level | 30233
- Description | 30233
- Options | 30233
- Required Privilege Level | 30233
- Release Information | 30234

Syntax

```
tunnel {
    gre {
        gre-4in4;
```

```
        gre-4in6;
        gre-6in4;
        gre-6in6;
    }
    ip-in-udp {
        teredo;
    }
    ipip {
        ipip-4in4;
        ipip-4in6;
        ipip-6in4;
        ipip-6in6;
        ipip-6over4;
        ipip-6to4relay;
        isatap;
        dslite;
    }
    bad-inner-header;
}
```

Hierarchy Level

```
[edit security screen ids-option ids-option-name ip]
```

Description

Enable IP tunnel IDS options.

Options

The remaining options are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D10.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

tunnel-services (Chassis)

IN THIS SECTION

- Syntax | 30234
- Hierarchy Level | 30234
- Description | 30235
- Options | 30235
- Required Privilege Level | 30235
- Release Information | 30235

Syntax

```
tunnel-services {
    bandwidth bandwidth-value;
    tunnel-only;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number]
```

Description

For MX Series 5G Universal Routing Platforms, configure the amount of bandwidth for tunnel services.

For ACX Series routers, configure the amount of bandwidth for tunnel services. Only bandwidths of 1 Gbps and 10 Gbps are supported for ACX routers.

For M7i, M10i, M120, M320, T Series and TX Matrix routers with IQ2 PICs and IQ2E PICs, configure support for per unit scheduling for GRE tunnels. You can specify the IQ2 and IQ2E PICs to work exclusively in tunnel mode or as a regular PIC. The default setting uses IQ2 and IQ2E PICs as a regular PIC. If you do not configure the `tunnel-only` option, the IQ2 and IQ2 PICs operate as regular PICs. For M7i, M10i, M120, M320, T Series and TX Matrix routers with IQ2 PICs and IQ2E PICs, you can use the `tunnel-only` option to specify that an IQ2 or IQ2E PIC work in tunnel mode only.

> **NOTE**: Bandwidth rates of 20 gigabits per second and 40 gigabits per second require use of an MX Series router with the 100-Gigabit Ethernet Modular Port Concentrator (MPC) and the 100-Gigabit CFP MIC.

> **NOTE**: On MX80 routers and MX Series routers with Trio-based FPCs, when ingress queuing is enabled for a PIC, tunnel services and inline services are not supported on the same PIC.

Options

`tunnel-only` (Optional)—For M7i, M10i, M120, M320, T Series and TX Matrix routers with IQ2 PICs and IQ2E PICs, specify that an IQ2 or IQ2E PIC work in tunnel mode only.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

RELATED DOCUMENTATION

Example: Configuring Tunnel Interfaces on a Gigabit Ethernet 40-Port DPC

Example: Configuring Tunnel Interfaces on a 10-Gigabit Ethernet 4-Port DPC

Example: Configuring Tunnel Interfaces on the MPC3E

bandwidth (Tunnel Services) | 2880

tunnel-services (Routing Instances VPLS)

IN THIS SECTION

- Syntax | 30236
- Hierarchy Level | 30236
- Description | 30237
- Options | 30237
- Required Privilege Level | 30237
- Release Information | 30237

Syntax

```
tunnel-services {
    devices device-names;
    primary primary-device-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Specify that traffic for particular VPLS routing instances be forwarded to specific virtual tunnel (VT) interfaces, allowing you to load-balance VPLS traffic among all the available VT interfaces on the router.

> **NOTE**: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Options

devices ***device-names***—Specify the VT interfaces acceptable for use by the VPLS routing instance. If you do not configure this option, all VT interfaces available to the router can be used for de-encapsulating traffic for this instance.

primary ***primary-device-name***—Specify the primary VT interface to be used by the VPLS routing instance. The VT interface specified is used to de-encapsulate all VPLS traffic from the MPLS core network for this routing instance. If the VT interface specified is unavailable, then one of the other acceptable VT interfaces is used for handling the VPLS traffic. If you do not configure this option, any acceptable VT interface can be used to de-encapsulate VPLS traffic from the core.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Specifying the VT Interfaces Used by VPLS Routing Instances

tunnel-services (RSVP)

IN THIS SECTION

- Syntax | 30238
- Hierarchy Level | 30238
- Description | 30238
- Default | 30238
- Options | 30239
- Required Privilege Level | 30239
- Release Information | 30239

Syntax

```
tunnel-services {
    devices device-names;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Description

Enable ultimate-hop popping on point-to-multipoint LSPs. The Junos OS selects one of the available virtual tunnel (VT) interfaces to de-encapsulate the egress traffic. By default, the selection process is performed automatically.

Default

Ultimate-hop popping is disabled.

Options

`devices` *`device-names`*—Specify which VT interfaces are used to handle the RSVP traffic.

- **Range:** 0 to 8 devices

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

> *Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs*

tunnel-source

IN THIS SECTION

- Syntax | 30239
- Hierarchy Level | 30240
- Description | 30240
- Required Privilege Level | 30240
- Release Information | 30240

Syntax

```
tunnel-source address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel  family inet / inet6pim-ssm],
[edit routing-instances routing-instance-name provider-tunnel  family inet / inet6pim-ssm]
```

Description

Configure the source address for the provider space multipoint generic router encapsulation (mGRE) tunnel. This statement enables a VPN tunnel source for Rosen 6 or Rosen 7 multicast VPNs. .

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

In Junos OS Release 17.3R1, the `pim-ssm` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` hierarchies as part of an upgrade to add IPv6 support for default multicast distribution tree (MDT) in Rosen 7, and data MDT for Rosen 6 and Rosen 7.

RELATED DOCUMENTATION

group-address (Routing Instances Tunnel Group) | 7802

tunnel-source-protocol (Protocols OSPF)

IN THIS SECTION

- Syntax | 30241
- Hierarchy Level | 30241

- Description | 30241
- Options | 30241
- Required Privilege Level | 30242
- Release Information | 30242

Syntax

```
tunnel-source-protocol {
    rsvp;
    spring-te;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) traffic-engineering],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(ospf | ospf3) traffic-engineering],
[edit protocols (ospf | ospf3) traffic-engineering],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)traffic-engineering]
```

Description

Enable SR-TE tunnel support.

Options

rsvp Pick label-switched paths from rsvp.

- **Range:** 1 through 255

spring-te Pick label-switched paths from spring-te.

- **Range:** 1 through 255

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

Tunneling LDP over SR-TE

tunnel-switch-profile (L2TP Tunnel Switching, Application)

IN THIS SECTION

- Syntax | 30242
- Hierarchy Level | 30243
- Description | 30243
- Options | 30243
- Required Privilege Level | 30243
- Release Information | 30243

Syntax

```
tunnel-switch-profile profile-name;
```

Hierarchy Level

```
[edit access domain map domain-map-name],
[edit services l2tp],
[edit services l2tp tunnel-group group-name]
```

Description

Specify a tunnel switch profile that determines whether packets in an L2TP session from a LAC are switched to another session that has a different destination LNS.

Options

profile-name—Unique name that identifies the tunnel switch profile.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

tunnel-switch-profile (L2TP Tunnel Switching, Definition)

IN THIS SECTION

- Syntax | 30244
- Hierarchy Level | 30244

- Description | 30244
- Options | 30244
- Required Privilege Level | 30245
- Release Information | 30245

Syntax

```
tunnel-switch-profile profile-name {
    avp {
        bearer-type action;
        calling-number action;
        cisco-nas-port-info action;
    }
    tunnel-profile profile-name;
}
```

Hierarchy Level

```
[edit access]
```

Description

Define a tunnel switch profile for subscriber access; specify actions to take for L2TP AVPs in the switched packets and the profile for the tunnel to which the PPP traffic is switched.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

profile-name—Unique name that identifies the tunnel switch profile. The profile can be applied as a default or referenced from within a domain map, a tunnel group, or by the RADIUS Tunnel-Group VSA [26-64].

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

L2TP Tunnel Switching For Multiple-Domain Networks

tunnel-termination

IN THIS SECTION

- Syntax | 30245
- Hierarchy Level | 30246
- Description | 30246
- Required Privilege Level | 30246
- Release Information | 30246

Syntax

```
tunnel-termination;
```

Hierarchy Level

```
[edit interfaces fti number unit logical-unit-number tunnel encapsulation  ( gre | ipip | udp )]
[edit interfaces interface-name family inet]
[edit forwarding-options]
```

Description

Enables the tunnel to terminate on the tunnel underlay interface configured at [edit interfaces *interface-name*"family" on page 6859 inet] hierarchy level.

On a flexible tunnel interface (FTI), enables tunnel termination for GRE, IPIP , or UDP decapsulation.

> **NOTE**: Starting in Junos OS Evolved Release 20.3R1, this statement is optional. If the statement is configured, the router can perform decapsulation of packets only. When the statement is not configured, you can configure encapsulation and decapsulation on the FTI.

> **NOTE**: For the Junos OS Evolved Release 20.1R2, you have to configure the `tunnel-termination` option for IP-IP decapsulation mandatorily. In Junos OS Evolved Release 20.4R1, when tunnel-termination is configured, the router functions in decap-only mode and the understanding of source address and destination address is swapped.

When configured at the [edit forwarding-options] hierarchy level, the `tunnel-termination` option enables the termination of tunnels on all interfaces at the global level. This option is disabled by default. It is mandatory to enable this option for EVPN-VXLAN deployments.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Evolved Release 20.2R1.

Statement introduced for IPIP tunnels in Junos OS Evolved Release 20.4R1.

Statement introduced for UDP tunnels in Junos OS Evolved Release 21.1R1.

Statement introduced at the [edit forwarding-options] hierarchy level in Junos OS Evolved Release 21.4R2.

Statement introduced in Junos OS Evolved 24.1R1 for ACX Series routers (GRE tunnels)

tunnel-timeout

IN THIS SECTION

- Syntax | 30247
- Hierarchy Level | 30247
- Description | 30247
- Options | 30247
- Required Privilege Level | 30248
- Release Information | 30248

Syntax

```
tunnel-timeout seconds;
```

Hierarchy Level

```
[edit services l2tp tunnel-group name]
```

Description

Specify the maximum downtime for an L2TP tunnel, after which the tunnel is terminated because the connection is presumed to have been lost.

Options

seconds—Interval after which the tunnel is terminated if no data can be sent.

- **Default:** 120 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Timers for L2TP Tunnels

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

tunnel (Tunnel Profile)

IN THIS SECTION

- Syntax | 30248
- Hierarchy Level | 30249
- Description | 30249
- Options | 30249
- Required Privilege Level | 30249
- Release Information | 30250

Syntax

```
tunnel tunnel-id {
    identification name;
    logical-system logical-system-name;
    max-sessions number;
```

```
    medium type;
    nas-port-method cisco-avp;
    preference number;
    remote-gateway {
        address server-ip-address;
        gateway-name server-name;
    }
    routing-instance routing-instance-name;
    secret password;
    source-gateway {
        address client-ip-address;
        gateway-name client-name;
    }
    type tunnel-type;
}
```

Hierarchy Level

```
[edit access tunnel-profile profile-name]
```

Description

Define the attributes of a tunnel for the tunnel profile. You can define up to 31 tunnels for each tunnel profile.

Options

tunnel-id—Unique integer that identifies a tunnel defined within a profile. For a subscriber, RADIUS attributes and VSAs can supply or override the attributes defined here for the tunnel.

- **Range:** 1 through 31

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

twamp

IN THIS SECTION

- Syntax (Junos OS) | **30250**
- Syntax (Junos OS Evolved) | **30253**
- Hierarchy Level (Junos OS) | **30254**
- Hierarchy Level (Junos OS Evolved) | **30255**
- Description | **30255**
- Required Privilege Level | **30255**
- Release Information | **30255**

Syntax (Junos OS)

```
twamp {
    client {
        control-connection  name {
            authentication-mode none;
            control-type (managed | light);
            destination-interface  destination-interface;
            destination-port   destination-port;
            history-size   history-size;
            moving-average-size  moving-average-size;
            persistent-results;
            routing-instance  routing-instance;
            target-address  target-address;
```

```
            tcp-keepcnt   count;
            tcp-keepidle  seconds;
            tcp-keepintvl  seconds;
            test-count test-count;
            test-interval   seconds;
            traps {
                control-connection-closed;
                test-iteration-done;
            }
            test-session name {
                data-fill-with zeros;
                data-size  data-size;
                destination-port   destination-port;
                dscp-code-points (RPM) dscp-code-points;
                probe-count probe-count;
                probe-interval seconds;
                source-address source-address;
                target-address target-address local-link IPv6-link-local-interface-name;
                thresholds {
                    egress-time microseconds;
                    ingress-time microseconds;
                    jitter-egress microseconds;
                    jitter-ingress microseconds;
                    jitter-rtt microseconds;
                    max-rtt microseconds;
                    rtt microseconds;
                    std-dev-egress microseconds;
                    std-dev-ingress microseconds;
                    std-dev-rtt microseconds;
                    successive-loss successive-loss;
                    total-loss total-loss;
                }
                traps {
                    egress-jitter-exceeded;
                    egress-std-dev-exceeded;
                    egress-time-exceeded;
                    ingress-jitter-exceeded;
                    ingress-std-dev-exceeded;
                    ingress-time-exceeded;
                    jitter-exceeded;
                    max-rtt-exceeded;
                    probe-failure;
                    rtt-exceeded;
```

```
                    std-dev-exceeded;
                    test-completion;
                    test-failure;
                }
                ttl hop-count;
            }
        }
    }
    post-cli-implicit-firewall;
    server {
        authentication-key-chain name {
            key-id name {
                secret secret;
            }
        }
        authentication-mode <authenticated > <control-only-encrypted> <encrypted >
<none>;
        client-list <list-name>  {
            address address <routing-instance [instance-name...]>;
        }
        light {
            port port-number;
        }
        max-connection-duration  hours;
        maximum-connections  maximum-connections;
        maximum-connections-per-client  maximum-connections-per-client;
        maximum-sessions  maximum-sessions;
        maximum-sessions-per-connection  maximum-sessions-per-connection;
        port  port;
        routing-instance-list  name {
            port port;
        }
        server-inactivity-timeout  minutes;
        tcp-keepcnt  count;
        tcp-keepidle  seconds;
        tcp-keepintvl seconds;
    }
}
```

Syntax (Junos OS Evolved)

```
twamp {
    client {
        control-connection  name {
            control-type (managed | light);
            destination-port  destination-port;
            routing-instance routing-instance;
            source-address source-address;
            target target-address;
            test-start (auto | manual);
            test-interval seconds;
            traps {
                control-connection-closed;
                test-iteration-done;
            }
            test-session name {
                data-size data-size;
                destination-port destination-port;
                dscp-code-points dscp-code-points;
                history-size history-size;
                moving-average-size moving-average-size;
                offload-type (none | inline-timestamping | pfe-timestamp);
                probe-count  probe-count;
                probe-interval  seconds;
                source-address source-address;
                source-routing {
                    forward (srmpls-nodes <label hop1...n> | srv6-nodes <IPv6 SID hop1...n>);
                    return (srmpls-nodes <label hop1...n> | srv6-nodes <IPv6 SID hop1...n>);
                }
                target target-address local-link IPv6-link-local-interface-
name ;
                thresholds {
                    control-failure (on | off);
                    successive-loss number;
                    total-loss number;
                    threshold-type (microseconds | average);
                }
                traps {
                    egress-jitter-exceeded
                    egress-time-exceeded;
                    ingress-jitter-exceeded;
```

```
                    ingress-time-exceeded;
                    jitter-exceeded;
                    probe-failure;
                    rtt-exceeded;
                    test-completion;
                    test-failure;
                }
                ttl hop-count;
            }
        }
    }
    server {
        managed {
            client-limit limit;
            client-list list-name {
                address address <routing-instance [instance-name...]>;
            }
            control-inactivity-timeout seconds;
            control-per-client-limit limit;
            control-maximum-duration seconds;
            offload-type (none | inline-timestamp | pfe-timestamp);
            port  port;
            stateful-sequence;
            test-per-client-limit limit;
        }
        light {
            offload-type (none | inline-timestamp | pfe-timestamp);
            port  port;
            stateful-sequence cache-size size;
        }
    }
}
```

Hierarchy Level (Junos OS)

```
[edit services rpm]
```

Hierarchy Level (Junos OS Evolved)

```
[edit services monitoring]
```

Description

Configure the Two-Way Active Measurement Protocol (TWAMP) client or server settings on all M Series and T Series routers that support Multiservices PICs (running in either Layer 2 or Layer 3 mode), MX Series routers, ACX routers, EX4300 Switches, PTX Series routers, EX9200 Series switches, and QFX10000 Series switches.

TWAMP is an open protocol for measurement of two-way metrics. The host that initiates the TCP connection takes the roles of the control-client and (in the two-host implementation) the session-sender. Such a device is also called the TWAMP client. The host that acknowledges the TCP connection accepts the roles of a server and (in the two-host implementation) and the session-reflector. Such a device is also called the TWAMP server. The TWAMP-Test messages are exchanged between the session-sender and the session-reflector, and the TWAMP-Control messages are exchanged between the control-client and the server.

The following addresses cannot be used for the `client-list` source IP address used for probes:

- 0.0.0.0
- 127.0.0.0/8 (loopback)
- 224.0.0.0/4 (multicast)
- 255.255.255.255 (broadcast)

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Junos OS Evolved syntax and hierarchy level introduced in Junos OS Evolved 20.3R1 for PTX10003 routers.

Support for the rest of the PTX10000 Series routers added in Junos OS Evolved Release 21.1R1.

`control-type` and `light` options added in Junos OS Release 21.1R1.

`traps` option added in Junos OS Evolved Release 21.3R1.

`source-address` option and the `local-link` sub-option of the `target-address` option for TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

The `local-link` sub-option of the `target` option for TWAMP Light test sessions introduced in Junos OS Evolved Release 22.3R1.

The `inline-timestamping` suboption of the `offload-type` option introduced in Junos OS Evolved Release 22.4R1.

`offload-type` and `stateful-sequence` options for light and managed servers introduced in Junos OS Evolved Release 23.4R1.

`source-routing` option introduced in Junos OS Evolved Release 24.4R1.

RELATED DOCUMENTATION

> *Configure TWAMP on ACX, MX, M, T, and PTX Series Routers EX4300 Series, EX9200 Series, and QFX10000 Series Switches*

twamp (Security ALG)

IN THIS SECTION

- Syntax | **30257**
- Hierarchy Level | **30257**
- Description | **30257**
- Options | **30257**
- Required Privilege Level | **30257**
- Release Information | **30258**

Syntax

```
alg {
    twamp {
        traceoptions {
            flag {
                all extensive;
            }
        }
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Specify the Two-Way Active Measurement Protocol (TWAMP) Application Layer Gateway (ALG) on the device. The TWAMP ALG extracts the IP addresses and the port details from messages between the control client and the TWAMP server and performs IP translation and gate opening to permits the test sessions. The TWAMP ALG also performs the message sanity check. By default, the TWAMP ALG is disabled. To enable, you need to configure the TWAMP ALG at `[edit security alg twamp]` hierarchy level.

Options

twamp Configure the TWAMP ALG.

traceoptions Define the tracing operations for the TWAMP ALG. The TWAMP ALG trace options helps to monitor traffic flow into and out of the NFX Series and SRX Series devices.

flag Configure the TWAMP ALG trace flags.

all Trace with all flags enabled.

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

twamp-server

IN THIS SECTION

- Syntax | 30258
- Hierarchy Level | 30258
- Description | 30258
- Required Privilege Level | 30258
- Release Information | 30259

Syntax

```
twamp-server;
```

Hierarchy Level

```
[edit interfaces sp-fpc/pic/port unit logical-unit-number]
```

Description

Specify the service PIC logical interface to provide the TWAMP service.

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring TWAMP on MX, M, T, and PTX Series Routers and EX4300 Series Switches

two-rate

IN THIS SECTION

- Syntax | 30259
- Hierarchy Level | 30260
- Description | 30260
- Required Privilege Level | 30260
- Release Information | 30260

Syntax

```
two-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    peak-information-rate bps;
    peak-burst-size bytes;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name],
[edit firewall three-color-policer policer-name],
[edit logical-systems logical-system-name firewall three-color-policer policer-name]
```

Description

Configure a two-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), peak information rate (PIR), and peak burst size (PBS).

Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the PIR or the PBS are assigned medium-high loss priority (yellow). Packets that exceed the PIR and the PBS are assigned high loss priority (red).

Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Logical systems support introduced in Junos OS Release 9.3.

Support at the `[edit dynamic-profiles ... three-color-policer name` hierarchy levels introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Three-Color Policer Configuration Overview

color-aware | 4428

color-blind | 4429

single-rate | 27969

tx-address-change (L2TP LAC)

IN THIS SECTION

- Syntax | 30261
- Hierarchy Level | 30261
- Description | 30261
- Default | 30262
- Options | 30262
- Required Privilege Level | 30262
- Release Information | 30262

Syntax

```
tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port | reject | reject-ip-address | reject-udp-port);
```

Hierarchy Level

```
[edit services l2tp tunnel]
```

Description

Configure whether the LAC accepts, ignores, or rejects requests from an LNS to change the destination IP address, UDP port, or both.

When configured to ignore change requests, the LAC continues to send packets to the original address or port, but accepts packets from the new address or port.

When configured to reject change requests, the LAC sends a StopCCN message to the original address or port and then terminates the connection to that LNS. This method has precedence over the ignore configuration.

Default

The LAC accepts IP address or UDP port change requests from the LNS.

Options

accept	Accept all change requests for IP address or UDP port.
ignore	Ignore all change requests for IP address or UDP port.
ignore-ip-address	Ignore a change request for IP address, but accept a change request for UDP port.
ignore-udp-port	Ignore a change request for UDP port, but accept a change request for IP address.
reject	Reject all change requests for IP address or UDP port.
reject-ip-address	Reject a change request for IP address, but accept a change request for UDP port.
reject-udp-port	Reject a change request for UDP port, but accept a change request for IP address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

`reject`, `reject-ip-address`, and `reject-udp-port` options added in Junos OS Release 16.1.

RELATED DOCUMENTATION

L2TP LAC Subscriber Configuration

tx-buffers

IN THIS SECTION

- Syntax | 30263
- Hierarchy Level | 30263
- Description | 30263
- Default | 30264
- Options | 30264
- Required Privilege Level | 30264
- Release Information | 30264

Syntax

```
tx-buffers (on | off);
```

Hierarchy Level

```
[edit interfaces interface-name ether-options configured-flow-control]
```

Description

Enable or disable an interface to respond to received Ethernet PAUSE messages. If you enable the transmit buffers to respond to PAUSE messages, when the interface receives a PAUSE message from the connected peer, the interface stops transmitting frames on the entire link. When the receive buffer on the connected peer empties below a certain threshold, the peer interface sends a message to the paused interface to resume sending frames.

Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the `tx-buffers` statement with the `rx-buffers` statement to configure asymmetric Ethernet PAUSE on an interface. (Use the `flow-control` statement to enable symmetric PAUSE and the `no-flow-control` statement to disable symmetric PAUSE on an interface. Symmetric flow control and

asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)

> **NOTE**: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).
>
> Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

Default

Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.

Options

`on | off`—Enable or disable an interface to respond to an Ethernet PAUSE message.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

- flow-control | **7267**
- rx-buffers | **16164**
- *Configuring CoS Asymmetric Ethernet PAUSE Flow Control*
- *Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control*
- *Understanding CoS Flow Control (Ethernet PAUSE and PFC)*

tx-connect-speed-method (L2TP LAC)

IN THIS SECTION

- Syntax | 30265
- Hierarchy Level | 30265
- Description | 30265
- Options | 30266
- Required Privilege Level | 30267
- Release Information | 30268

Syntax

```
tx-connect-speed-method method;
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Specify the method that determines how to derive the connect speed values sent from the LAC to the LNS.

When the session is being established, the speeds are included in the Incoming-Call-Connected (ICCN) message. The transmit speed is conveyed in AVP 24 (Tx-Connect-Speed) and the receive speed is conveyed in AVP 38 (Rx-Connect-Speed). Both values are in bits per seconds (bps). The LAC typically uses the `static` or `pppoe-ia-tags` method, because values for other configured methods are not available when the session is being established.

When connect speed updates are configured, the LAC sends the updated values for each session to the LNS in Connect-Speed-Update-Notification (CSUN) messages. The updated speeds are conveyed in the Connect Speed Update AVP (97).

When connection speed values are not available from the configured method, the LAC falls back to another source for the values. See *L2TP Subscriber Access Lines and Connection Speeds* for tables describing the LAC fallback behavior by Junos OS release.

Options

method Method used to derive the connection speed values.

- `actual`—(Junos OS Releases 15.1, 16.1, 16.2, 17.1) The speed is derived from the CoS effective shaping rate that is enforced on the level 3 node based on local policy. In the supported releases, `actual` is the default method and has the highest preference among all configured methods.

  This method is not available starting in Junos OS Release 17.2. However, it is configurable in the Tunnel-Tx-Speed-Method VSA (26-94). If you do so, it is translated to the `service-profile` method.

- `ancp`—The speed is derived from the configured ANCP value for the underlying interface. This value results from a user-defined percentage correction to the values received from the access node; this is configured per subscriber access line. The percentage accounts for encapsulation differences between, the router, the access loop, and the Layer 1 transport overhead. The initial rate sent to the LNS is the ANCP value reported at the time the ICCN is sent. The ANCP value is not available for the ICCN message and falls over to another method. You can change the configured correction after a subscriber has logged in, but those changes do not affect the actual rate used by the LNS for that subscriber.

- `none`—This method prevents the LAC from sending AVP 24 and AVP 38 to the LNS. This option also overrides the Juniper Networks RADIUS VSAs, Tx-Connect-Speed (26-162) and Rx-Connect-Speed (26-163).

- `pppoe-ia-tags`—The speed is derived from the PPPoE IA tags received by the LAC from the DSLAM. This speed value is transmitted when a subscriber logs in and it cannot subsequently be changed. The value of Actual-Data-Rate-Downstream (VSA 26-129) is used for AVP 24. The value of Actual-Data-Rate-Upstream (VSA 26-130) is used for AVP 38; it is sent only when the values differ.

  > **NOTE**: This speed derived from the IA tags does not apply to subscribers that are already logged in; it is effective only for subscribers that log in after this setting has been saved.

- `service-profile`—(Junos OS Releases 17.2 and higher) The downstream (Tx) speed is derived from the actual CoS that is enforced on the L3 node based on local policy. The upstream (Rx) speed is the value configured in the dynamic service profile; no adjustment is made to this value. The service-profile value is not available for the ICCN message and falls over to another method.

  The `service-profile` method is supported only when the `effective shaping-rate` statement is included at the `[edit chassis]` hierarchy level. If it is not, the commit check fails. If the method is received from RADIUS in VSA 26-94, a system log message is generated instead, because no commit check is performed in this case.

- `static`—(Junos OS Releases 13.3, 14.1, and 14.2; Junos OS Releases 17.2 and higher) The speed is derived from the configured static Layer 2 speed. For Ethernet VLANs, this is the recommended (advisory) shaping rate configured on the PPPoE logical interface underlying the subscriber interface. If the advisory shaping rate is not configured on the underlying interface, then the actual speed of the underlying physical port is used. In the supported releases, `static` is the default method.

  In Junos OS Releases 15.1, 16.1, 16.2, and 17.2, the `static` method is configurable for backward compatibility, but it is not supported. If you configure this method in the CLI or in the Tunnel-Tx-Speed-Method VSA (26-94), the LAC falls back to the port speed of the subscriber access interface.

  > **NOTE**: For ge and xe interfaces, the port speed value is set to 1,000,000,000 and for ae interfaces, the port speed value is set to 0. The value is sent in both AVP 24 and AVP 38.

- Default:
  - `static` (Starting in Junos OS Release 17.2)
  - `actual` (Junos OS Releases 15.1, 16.1, 16.2, 17.1)
  - `static` (Junos OS Releases 13.3, 14.1, and 14.2)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Options **ancp**, **pppoe-ia-tag**, and **static** introduced in Junos OS Release 13.1.

Option `static` deprecated in Junos OS Release 15.1.

Options **actual** and **none** added in Junos OS Release 15.1.

Option `actual` deprecated in Junos OS Release 17.2.

Option `service-profile` added in Junos OS Release 17.2.

Option `static` undeprecated in Junos OS Release 17.2.

RELATED DOCUMENTATION

L2TP Subscriber Access Lines and Connection Speeds

L2TP Subscriber Access Lines and Connection Speeds

tx-dnu-to-line-source-enable

IN THIS SECTION

- Syntax | 30268
- Hierarchy Level | 30269
- Description | 30269
- Required Privilege Level | 30269
- Release Information | 30269

Syntax

```
tx-dnu-to-line-source-enable;
```

Hierarchy Level

```
[edit chassis synchronization output interfaces external]
[edit chassis synchronization output interfaces (external-0/0 | external-1/0)]
```

Description

Configure the transmitting quality level to DNU or DUS on the line source interface that has been selected as the external output source.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

tx-duration

IN THIS SECTION

- Syntax | 30270
- Hierarchy Level | 30270
- Description | 30270
- Options | 30270
- Required Privilege Level | 30270
- Release Information | 30270

Syntax

```
expected-defect {
tx-duration ;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

Description

The expected duration for which the peer MEP should suppress the LoC alarms.

Options

Minimum value The minimum value at which the peer MEP should suppress the LoC alarms is 120 *seconds*.

Minimum value The maximum acceptable value at which the peer MEP should suppress the LoC alarms is 3600 *seconds*.

Default 900 *seconds*.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

connectivity-fault-management | 4556

show oam ethernet connectivity-fault-management mep-database | 22620

tx-enable

IN THIS SECTION

- Syntax | 30271
- Hierarchy Level | 30271
- Description | 30271
- Default | 30271
- Required Privilege Level | 30271
- Release Information | 30272

Syntax

```
expected-defect {
tx-enable ;
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

Description

Enable the ethernet expected defect (ETH-ED) function to control if EDM transmission need to be triggered on ISSU.

The remaining statements are explained separately. See CLI Explorer.

Default

The MEP does not generate EDM PDUs by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

connectivity-fault-management | 4556

show oam ethernet connectivity-fault-management mep-database | 22620

tx-power

IN THIS SECTION

- Syntax | 30272
- Hierarchy Level | 30272
- Description | 30273
- Default | 30273
- Options | 30273
- Required Privilege Level | 30273
- Release Information | 30273

Syntax

```
tx-power dbm;
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Description

Transmit laser output power (dBm).

Default

If you don't specify a value, the default transmit laser output power is –2 dBm.

Options

dbm—Transmit power value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

optics-options | **11847**

100-Gigabit Ethernet OTN Options Configuration Overview

type

IN THIS SECTION

- Syntax | **30274**
- Hierarchy Level | **30274**
- Description | **30274**
- Options | **30274**

Required Privilege Level | 30274
Release Information | 30274

Syntax

```
type {
    linux-container | virtual-machine;
    }
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Type of the VNF.

Options

linux-container	The VNF type is Linux container.
virtual-machine	The VNF type is virtual machine.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D40.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

type (Application Identification)

IN THIS SECTION

- Syntax | 30275
- Hierarchy Level | 30275
- Description | 30275
- Options | 30276
- Required Privilege Level | 30276
- Release Information | 30276

Syntax

```
type type;
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Description

Specify the type of application, such as FTP or HTTP.

Options

type Type of application such as FTP or HTTP.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

type (ICMP Mapping for Application Identification)

IN THIS SECTION

- Syntax | 30277
- Hierarchy Level | 30277
- Description | 30277
- Options | 30277
- Required Privilege Level | 30277
- Release Information | 30277

Syntax

```
type icmp-type;
```

Hierarchy Level

```
[edit services application-identification application application-name icmp-mapping]
```

Description

Match an ICMP type value to create a custom application signature.

Options

value ICMP code value.

- **Range:** 0 through 254

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Application Identification Overview

Configuring Custom Application Signatures

type (Protocols BGP)

IN THIS SECTION

- Syntax | 30278
- Hierarchy Level | 30278
- Description | 30278
- Options | 30278
- Required Privilege Level | 30279
- Release Information | 30279

Syntax

```
type type;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp
group group-name],
[edit protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name]
```

Description

Specify the type of BGP peer group.

When configuring a BGP group, you can indicate whether the group is an IBGP group or an EBGP group. All peers in an IBGP group are in the same AS, while peers in an EBGP group are in different ASs and normally share a subnet.

Options

type—Type of group:

- `external`—External group, which allows inter-AS BGP routing
- `internal`—Internal group, which allows intra-AS BGP routing
- **Default:** If you do not specify the BGP group type or assign a `peer-as`, then Junos OS assigns peer group type `external` by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Configuring External BGP Point-to-Point Peer Sessions

Example: Configuring Internal BGP Peer Sessions

BGP User Guide

type (RMON)

IN THIS SECTION

- Syntax | 30280
- Hierarchy Level | 30280
- Description | 30280
- Options | 30280
- Required Privilege Level | 30280
- Release Information | 30280

Syntax

```
type type;
```

Hierarchy Level

```
[edit snmp rmon event index]
```

Description

Type of notification generated when a threshold is crossed.

Options

type—Type of notification:

- `log`—Add an entry to `logTable`.
- `log-and-trap`—Send an SNMP trap and make a log entry.
- `none`—No notifications are sent.
- `snmptrap`—Send an SNMP trap.
- **Default:** `log-and-trap`

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configure an RMON Event Entry and Its Attributes

type (RMON Notification)

IN THIS SECTION

- Syntax | 30281
- Hierarchy Level | 30281
- Description | 30281
- Default | 30281
- Options | 30281
- Required Privilege Level | 30282
- Release Information | 30282

Syntax

```
type type;
```

Hierarchy Level

```
[edit snmp rmon event index]
```

Description

Configure the type of notification generated when a rising or falling threshold is crossed.

Default

log-and-trap

Options

type—Type of notification. It can be one of the following:

- **log**—Add an entry to the **logTable** object.

- **log-and-trap**—Send an SNMP trap and add a log entry.
- **none**—No notifications are sent.
- **snmptrap**—Send an SNMP trap.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

RMON MIB Event, Alarm, Log, and History Control Tables

RMON Overview

Junos OS Network Management Configuration Guide

type (Security Antivirus Feature Profile)

IN THIS SECTION

- Syntax | 30283
- Hierarchy Level | 30283
- Description | 30283
- Required Privilege Level | 30283
- Release Information | 30283

Syntax

```
type (juniper-express-engine | kaspersky-lab-engine | sophos-engine);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus]
```

Description

Set the antivirus engine that will be used on the device. You can only have one engine type running and you must restart the device if you change engines.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Statement updated for Sophos in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

type (Security Content Filtering Notification Options)

IN THIS SECTION

- Syntax | 30284
- Hierarchy Level | 30284
- Description | 30284
- Options | 30284
- Required Privilege Level | 30285
- Release Information | 30285

Syntax

```
type (message | protocol-only);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile content-filtering profile profile-name notification-options]
```

Description

When content is blocked, the client generally still receives a successful response code but with modified content (file replacement) containing a warning message. But with protocol-only notifications, a protocol-specific error code might be returned to the client.

Options

- `message`—Send a generic notification.
- `protocol-only`—Send a protocol-specific notification.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

type (Security Fallback Block)

IN THIS SECTION

- Syntax | 30285
- Hierarchy Level | 30285
- Description | 30286
- Options | 30286
- Required Privilege Level | 30286
- Release Information | 30286

Syntax

```
type (message | protocol-only);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options fallback-block]
```

```
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options fallback-block]
```

Description

You can configure notifications for both fallback blocking and fallback nonblocking actions. With protocol-only notifications, a protocol-specific error code may be returned to the client.

Options

- message—Send a generic notification.
- protocol-only—Send a protocol-specific notification.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1 .

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

type (Security IDP Dynamic Attack Group)

IN THIS SECTION

- Syntax | 30287
- Hierarchy Level | 30287
- Description | 30287
- Options | 30287
- Required Privilege Level | 30287
- Release Information | 30288

Syntax

```
type {
    values [anomaly signature];
}
```

Hierarchy Level

```
[edit security idp dynamic-attack-group dynamic-attack-group-name filters]
```

Description

Specify an attack type filter to add attack objects based on the type of attack object (signature or protocol anomaly).

Options

`values`—Name of the attack type filter.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

type (Security IDP ICMP Headers)

IN THIS SECTION

- Syntax | 30288
- Hierarchy Level | 30288
- Description | 30288
- Options | 30289
- Required Privilege Level | 30289
- Release Information | 30289

Syntax

```
type {
    match (equal | greater-than | less-than | not-equal);
    value type-value;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol icmp]
```

Description

Specify the primary code that identifies the function of the request/reply.

Options

`match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.

`value` *`type-value`*—Match a decimal value.

- **Range:** 0 through 255

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

type (Security Virus Detection)

IN THIS SECTION

- Syntax | 30289
- Hierarchy Level | 30290
- Description | 30290
- Options | 30290
- Required Privilege Level | 30290
- Release Information | 30290

Syntax

```
type (message | protocol-only);
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options virus-detection]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
notification-options virus-detection]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options virus-detection]
```

Description

When content is blocked because a virus is found or a scan error occurs, the client generally still receives a successful response code but with modified content (file replacement) containing a warning message. But with protocol-only notifications, a protocol-specific error code might be returned to the client.

Options

- message—Send a generic notification.
- protocol-only—Send a protocol-specific notification.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

type (Security Web Filtering)

IN THIS SECTION

- Syntax | 30291
- Hierarchy Level | 30291
- Description | 30291
- Options | 30292
- Required Privilege Level | 30292
- Release Information | 30292

Syntax

```
type (juniper-enhanced | juniper-local | surf-control-integrated | websense-redirect) | ng-juniper;
```

Hierarchy Level

```
[edit security utm default-configuration web-filtering]
[edit security utm feature-profile web-filtering]
```

Description

Define the type of Web filtering solution or URL filtering solution used by the device.

Options

- **juniper-enhanced**—Enable Enhanced Web Filtering on the device.
- **juniper-local** —Enable Juniper Networks local URL filtering on the device.
- **surf-control-integrated**—Enable integrated Web filtering on the device.
- **websense-redirect**—Redirect the URL to the Websense server.
- **ng-juniper**—Enable Juniper NextGen Web filtering on the device.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The surf-control-integrated feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, command introduced in Junos OS Release 9.5.

Command introduced in Junos OS Release 11.4 .

The `[edit security utm default-configuration web-filtering]` hierarchy level is introduced in Junos OS Release 18.2R1.

The ng-juniper option is introduced in Junos OS Release 23.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The surf-control-integrated feature is not supported from Junos OS Release 15.1X49-D10 onwards.

type (SNMP v3)

IN THIS SECTION

- Syntax | 30293
- Hierarchy Level | 30293
- Description | 30293
- Options | 30293
- Required Privilege Level | 30293
- Release Information | 30294

Syntax

```
type (inform | trap);
```

Hierarchy Level

```
[edit snmp v3 notify name]
```

Description

Configure the type of SNMP notification.

Options

`inform`—Defines the type of notification as an inform. SNMP informs are confirmed notifications.

`trap`—Defines the type of notification as a trap. SNMP traps are unconfirmed notifications.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`inform` option added in Junos OS Release 7.4.

type (Tunnel Profile)

IN THIS SECTION

- Syntax | 30294
- Hierarchy Level | 30294
- Description | 30294
- Default | 30295
- Options | 30295
- Required Privilege Level | 30295
- Release Information | 30295

Syntax

```
type tunnel-type;
```

Hierarchy Level

```
[edit access tunnel-profile profile-name tunnel tunnel-id]
```

Description

Specify the tunnel type (protocol).

Default

`l2tp`

Options

`tunnel-type`—Tunnel protocol type. The only value currently available is `l2tp`.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

U

IN THIS SECTION

- uac-policy (Application Services) | **30301**
- uac-service | **30303**
- ucode-imem-remap | **30304**
- udp | **30306**
- udp-encapsulation | **30307**
- udp (FTI) | **30309**
- udp-port | **30310**
- udp (Security IDP Protocol Binding) | **30312**

- udp (Security IDP Signature Attack) | **30313**
- udp (Security Screen) | **30314**
- udp-sweep | **30316**
- udp-tcp-port-swap | **30318**
- uid (Dynamic Profiles) | **30319**
- uid-reference | **30320**
- ukern-trace | **30322**
- ultimate-hop-popping | **30323**
- ultrahigh-power | **30325**
- unattended-boot | **30326**
- unconditional-src-learn | **30328**
- unconfigured-peer-graceful-restart | **30329**
- underlying-interface | **30330**
- underlying-interface (ANCP) | **30332**
- underlying-interface (demux0) | **30333**
- underlying-interface (Dynamic PPPoE) | **30335**
- unframed | no-unframed (Interfaces) | **30337**
- unicast-ethernet | **30338**
- unicast-ethernet-ipv4 | **30340**
- unicast-in-lpm | **30342**
- unicast-mode | **30344**
- unicast-negotiation | **30347**
- unicast-reverse-path | **30348**
- unicast (Route Target Community) | **30350**
- unicast-stream-limit (Protocols AMT) | **30351**
- unicast-umh-election | **30353**
- unicast (Virtual Tunnel in Routing Instances) | **30354**
- unicast-tunnel-name-regular-expression | **30355**
- unidirectional | **30356**
- unidirectional-session-refreshing | **30358**
- unified-access-control (Services) | **30359**
- unified-policy-explicit-match | **30361**

- unified-policy max-lookups | **30362**
- unique-nas-port (Access) | **30364**
- unit (Sampling) | **30366**
- unit (QFX Series) | **30367**
- unit (interfaces) | **30370**
- unit (Aggregated Multiservices) | **30380**
- unit (CoS Interfaces) | **30382**
- unit (CoS Interfaces) | **30384**
- unit (Dynamic Demux Interface) | **30385**
- unit (Dynamic Interface Sets) | **30388**
- unit (Dynamic PPPoE) | **30390**
- unit (Dynamic Profiles Standard Interface) | **30392**
- unit (Dynamic Traffic Shaping) | **30397**
- unit (Interfaces Services) | **30399**
- unit (Interfaces) | **30401**
- unit (Link-Multilink) | **30402**
- unit (Tunnel) | **30404**
- unit (TDF Interface) | **30406**
- unit (Voice Services) | **30408**
- unknown-message (Security H323 ALG) | **30410**
- unknown-message (Security MGCP ALG) | **30411**
- unknown-message (Security SCCP ALG) | **30413**
- unknown-message (Security SIP ALG) | **30415**
- unknown-protocol (IDS Screen Next Gen Services) | **30416**
- unknown-unicast-forwarding | **30417**
- unnumbered-address (Demux) | **30419**
- unnumbered-address (Dynamic PPPoE) | **30421**
- unnumbered-address (Dynamic Profiles) | **30422**
- unnumbered-address (Ethernet) | **30425**
- unnumbered-address (PPP) | **30426**
- unprotect | **30428**
- unreachable-vtep-aging-timer | **30429**

- untagged | 30431
- untrusted | 30432
- untrusted (VLAN) | 30433
- Unused | 30435
- up | 30436
- update | 30438
- update-interval | 30439
- update-interval (Protocols RIP) | 30441
- update-interval (Protocols RIPng) | 30443
- update-interval (Service Accounting) | 30444
- update-rate-limit | 30446
- update-response-timeout (PCRF Partition) | 30447
- update-router-advertisement | 30449
- update-server | 30450
- update-server (dhcp-client) | 30452
- update-server (dhcpv6-client) | 30453
- update-threading | 30454
- update-threshold | 30456
- update-threshold-max-reservable | 30458
- upgrade-group | 30460
- upgrade-groups (Junos Fusion) | 30461
- upgrade-mode | 30464
- uplink-failure-detection | 30465
- uplink-failure-detection (Junos Fusion) | 30467
- uplink-port-group (Junos Fusion) | 30469
- upload (Event Policy) | 30471
- upload-profile (Security Antivirus FTP) | 30472
- upload-profile (Security Content Filtering FTP) | 30473
- upstream-interface | 30474
- upstream-label | 30476
- upstream-mode | 30478
- upstream-rate (Traffic Shaping) | 30479

- upstream-rewrite | **30481**
- urgent-pointer | **30483**
- uri-check | **30484**
- url | **30485**
- url (Security) | **30487**
- url-blacklist | **30488**
- url-feed | **30489**
- url-filter | **30492**
- url-filter-profile | **30495**
- url-filter-template | **30496**
- url-pattern | **30499**
- url (Security) | **30503**
- url (Security Antivirus) | **30504**
- url (Security IDP) | **30506**
- url-whitelist | **30507**
- url-whitelist | **30508**
- usb | **30509**
- usb-control | **30510**
- use-active-child-mac-on-reth | **30511**
- use-actual-mac-on-physical-interfaces | **30512**
- use-class (Class Attribute) | **30514**
- use-extended-flow-memory | **30515**
- use-for-post-convergence-lfa | **30517**
- use-imported-time-zones | **30518**
- use-interface | **30519**
- use-interface-description (MX) | **30520**
- use-interface-description | **30522**
- use-interface-index | **30525**
- use-interface-name | **30526**
- use-option-82 | **30528**
- use-p2mp-lsp | **30530**
- use-post-convergence-lfa | **30531**

- use-post-convergence-lfa (Protocols OSPF) | **30533**
- use-primary (DHCP Local Server) | **30535**
- use-primary (DHCP Relay Agent) | **30537**
- use-string | **30539**
- use-transport-class | **30541**
- use-underlying-interface-mac | **30542**
- use-vlan-id | **30544**
- use-vlan-id (DHCP Relay Agent) | **30546**
- user (Access) | **30547**
- user-defined-option-82 | **30551**
- user-firewall | **30553**
- user-group-mapping | **30555**
- user-id | **30557**
- user-identification | **30558**
- user-identification (Services) | **30560**
- user-name | **30568**
- user-name (Event Policy) | **30569**
- user-name-include (OCS Partition) | **30571**
- user-password (PCEF Profile) | **30574**
- user-planes | **30575**
- user-prefix | **30577**
- user-prefix (DHCP Local Server) | **30578**
- user-prefix (DHCP Relay Agent) | **30580**
- user-prefix (Static Subscribers) | **30582**
- user-query (Services User Identification) | **30584**
- user (SNMP v3) | **30588**
- user (System Logging) | **30589**
- user (System Services) | **30591**
- username-include (Demux) | **30593**
- username-include (DHCP Local Server) | **30595**
- username-include (DHCP Relay Agent) | **30597**
- username-include (Interfaces) | **30600**

- username-include (Local Authentication) | **30602**
- username-include (Static Subscribers) | **30604**
- username (Security Antivirus) | **30606**
- username (Services) | **30607**
- using-user-password | **30609**
- usm | **30610**
- utc-timestamp | **30612**
- utm default-configuration | **30614**
- uuid | **30621**
- uuid (Applications) | **30623**

uac-policy (Application Services)

IN THIS SECTION

- Syntax | **30301**
- Hierarchy Level | **30302**
- Description | **30302**
- Options | **30302**
- Required Privilege Level | **30302**
- Release Information | **30302**

Syntax

```
uac-policy {
    captive-portal captive-portal;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
application-services]
```

Description

Enable Unified Access Control (UAC) for the security policy. This statement is required when you are configuring the SRX Series Firewall to act as a Junos OS Enforcer in a UAC deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series UAC Appliance .

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.4.

RELATED DOCUMENTATION

> *Understanding User Role Firewalls*
>
> *Example: Configuring a User Role Firewall on an SRX Series Device*

uac-service

IN THIS SECTION

- Syntax | 30303
- Hierarchy Level | 30303
- Description | 30303
- Options | 30303
- Required Privilege Level | 30304
- Release Information | 30304

Syntax

```
uac-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the unified access control daemon process.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the unified access control daemon process.

- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Firewall User Authentication Overview

ucode-imem-remap

IN THIS SECTION

- Syntax | 30305
- Hierarchy Level | 30305
- Description | 30305
- Required Privilege Level | 30305
- Release Information | 30305

Syntax

```
ucode-imem-remap;
```

Hierarchy Level

```
[edit chassis feb slot number]
```

Description

M120 routers with a single type-1 FPC mapped to an FEB support a microcode remap feature to resolve microcode overflow resulting in bad PIC combinations.

You can enable the microcode remap by using the `ucode-imem-remap` statement at the [**edit chassis feb slot *number***] hierarchy level. The default microcode map will continue to be available if the `ucode-imem-remap` statement is not configured.

> **NOTE**: On M120 routers, the FEB is automatically restarted once the `ucode-imem-remap` statement is configured and committed.

Required Privilege Level

interfaces—To view this statement in the configuration.

interfaces-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4R2.

udp

IN THIS SECTION

- Junos OS | 30306
- Junos OS Evolved | 30306
- Junos OS Hierarchy Level | 30306
- Junos OS Evolved Hierarchy Level | 30307
- Description | 30307
- Required Privilege Level | 30307
- Release Information | 30307

Junos OS

```
udp {
    destination-interface interface-name;
    port port;
}
```

Junos OS Evolved

```
udp {
    port port;
}
```

Junos OS Hierarchy Level

```
[edit services rpm probe-server]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm probe-server]
```

Description

Enable UDP requests for the RPM probe server.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

> **NOTE**: The `destination-interface` statement is not supported on PTX Series routers or for Junos OS Evolved.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

Configuring RPM Receiver Servers

udp-encapsulation

IN THIS SECTION

Syntax | 30308

- Hierarchy Level | 30308
- Description | 30308
- Options | 30308
- Required Privilege Level | 30309
- Release Information | 30309

Syntax

```
udp-encapsulation {
    <udp-dest-port destination-port>;
}
```

Hierarchy Level

```
[edit services service-set service-set-name ipsec-vpn-options]
```

Description

Enable multiple path forwarding of IPsec traffic by adding a UDP header to the IPsec encapsulation of packets. Doing this increases the throughput of IPsec traffic. If you do not enable UDP encapsulation, all the IPsec traffic follows a single forward path rather than using multiple available paths.

Options

udp-dest-port ***destination-port***

(Optional) Use the specified UDP destination port for the UDP header that is appended to the ESP encapsulation.

- **Range:** 1025 through 65536. Do not use 4500.
- **Default:** If you do not include the `udp-dest-port` statement, the default UDP destination port is 4565.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring IPsec Service Sets

IPsec Multipath Forwarding with UDP Encapsulation

udp (FTI)

IN THIS SECTION

- Syntax | 30309
- Hierarchy Level | 30310
- Description | 30310
- Required Privilege Level | 30310
- Release Information | 30310

Syntax

```
udp {
    destination {
        address ipv4_address;
    }
    source {
        address ipv4_address;
        interface interface;
```

```
    }
}
```

Hierarchy Level

```
[edit interfaces name unit name tunnel encapsulation]
```

Description

Enable UDP encapsulation with IPv4 on FTI logical interfaces.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Flexible Tunnel Interfaces Overview

udp-port

IN THIS SECTION

- Syntax | 30311
- Hierarchy Level | 30311
- Description | 30311
- Options | 30311
- Required Privilege Level | 30311
- Release Information | 30311

Syntax

```
udp-port port-number;
```

Hierarchy Level

```
[edit protocols sflow collector]
```

Description

Configure the UDP port for a remote collector for sFlow network traffic monitoring. The device sends sFlow UDP datagrams to the collector for analysis.

Options

port-number—UDP port number for this collector.

- **Default:** 6343

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

> *Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches*
>
> *Configuring sFlow Technology for Network Monitoring (CLI Procedure)*
>
> *Example: Configuring sFlow Technology to Monitor Network Traffic on MX Series Routers*

udp (Security IDP Protocol Binding)

IN THIS SECTION

- Syntax | 30312
- Hierarchy Level | 30312
- Description | 30312
- Options | 30312
- Required Privilege Level | 30313
- Release Information | 30313

Syntax

```
udp {
    minimum-port port-number <maximum-port port-number>;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type chain protocol-binding],

[edit security idp custom-attack attack-name attack-type signature protocol-binding]
```

Description

Allow IDP to match the attack for specified UDP ports.

Options

- `minimum-port` *`port-number`*—Minimum port in the port range.
- **Range:** 0 through 65,535
- `maximum-port` *`port-number`*—Maximum port in the port range.

- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

udp (Security IDP Signature Attack)

IN THIS SECTION

- Syntax | 30313
- Hierarchy Level | 30314
- Description | 30314
- Options | 30314
- Required Privilege Level | 30314
- Release Information | 30314

Syntax

```
udp {
    data-length {
        match (equal | greater-than | less-than | not-equal);
        value data-length;
    }
    destination-port (Security Signature Attack) {
        match (equal | greater-than | less-than | not-equal);
        value destination-port;
    }
    source-port (Security IDP) {
```

```
        match (equal | greater-than | less-than | not-equal);
        value source-port;
    }
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol]
```

Description

Allow IDP to match the UDP header information for the signature attack.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

udp (Security Screen)

IN THIS SECTION

- Syntax | 30315
- Hierarchy Level | 30315
- Description | 30315
- Options | 30315

Required Privilege Level | 30315
Release Information | 30316

Syntax

```
udp {
    flood {
        threshold number;
    }
    port-scan {
        threshold number;
    }
    udp-sweep {
        threshold threshold number;
    }
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name]
```

Description

Specify the number of packets allowed per second to the same destination IP address/port pair. When the number of packets exceeds this value within any 1-second period, the device generates an alarm and drops subsequent packets for the remainder of that second.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

udp-sweep

IN THIS SECTION

- Syntax | 30316
- Hierarchy Level | 30316
- Description | 30317
- Options | 30317
- Required Privilege Level | 30317
- Release Information | 30317

Syntax

```
udp-sweep {
    threshold number;
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name udp]
```

Description

Configure the device to detect and prevent UDP sweep attack. In a UDP sweep attack, an attacker sends UDP packets to the target device. If the device responds to those packets, the attacker gets an indication that a port in the target device is open, which makes the port vulnerable to attack. If a remote host sends UDP packets to 10 addresses in 0.005 seconds (5000 microseconds), then the device flags this as an UDP sweep attack.

If the `alarm-without-drop` option is not set, the device rejects the eleventh and all further UDP packets from that host for the remainder of the specified threshold period.

Options

`threshold` *`number`*—Maximum number of microseconds during which up to 10 UDP packets from the same host are allowed into the device. More than 10 requests from a host during this period triggers an UDP Sweep attack response on the device during the remainder of the second.

- **Range:** 1000 through 1,000,000 microseconds
- **Default:** 5000 microseconds

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

> *Attack Detection and Prevention Overview*
>
> *Example: Configuring Multiple Screening Options*

udp-tcp-port-swap

IN THIS SECTION

- Syntax | 30318
- Junos OS Hierarchy Level | 30318
- Junos OS Evolved Hierarchy Level | 30318
- Description | 30318
- Required Privilege Level | 30318
- Release Information | 30319

Syntax

```
udp-tcp-port-swap;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Swap the source and destination UDP ports in the test packets. Only UDP port swap and UDP over IPv4 traffic is supported.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

- rfc2544-benchmarking | **15759**
- rfc2544 | **15757**
- *RFC 2544-Based Benchmarking Tests Overview*
- *Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers*
- *Configuring an RFC 2544-Based Benchmarking Test*

uid (Dynamic Profiles)

IN THIS SECTION

- Syntax | **30319**
- Hierarchy Level | **30320**
- Description | **30320**
- Required Privilege Level | **30320**
- Release Information | **30320**

Syntax

```
uid;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name variables variable-name]
```

Description

Configure a unique ID for parameterized filters in a dynamic profile created for services. The values that the system uses for these variables are applied when the subscriber authenticates.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Dynamic Variables Overview

uid-reference

IN THIS SECTION

- Syntax | 30321
- Hierarchy Level | 30321
- Description | 30321
- Required Privilege Level | 30321
- Release Information | 30321

Syntax

```
uid-reference;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name variables variable-name]
```

Description

When you configure a unique ID (UID) variable, include this statement to specify that the value for the UID is supplied by RADIUS when the subscriber authenticates.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Unique Identifiers for Firewall Variables

Configuring Unique Identifiers for Parameterized Filters

Dynamic Variables Overview

ukern-trace

IN THIS SECTION

- Syntax | 30322
- Hierarchy Level | 30322
- Description | 30322
- Options | 30323
- Required Privilege Level | 30323
- Release Information | 30323

Syntax

```
ukern-trace {
    log {
        app-type dfw;
        logging (off | on);
    }
}
```

Hierarchy Level

```
[edit chassis fpc ],
[edit chassis lcc              name fpc  ],
[edit chassis member name fpc ]
```

Description

Enable or disable logging of all debugging firewall (DFW) ukern-trace logs on the specified FPC slot. The CLI change takes effect immediately and persists after the FPC slot reboots.

Options

app-type Name of application with ukern-trace logs. The only supported app-type is dfw.

logging Turn ukern-trace logging on or off. By default logging is on.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 17.4

RELATED DOCUMENTATION

System Log Overview

ultimate-hop-popping

IN THIS SECTION

- Syntax | 30323
- Hierarchy Level | 30324
- Description | 30324
- Default | 30324
- Required Privilege Level | 30324
- Release Information | 30325

Syntax

```
ultimate-hop-popping;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path label-switched-path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path label-switched-path-name]
```

Description

Enable ultimate-hop popping on LSPs. Configure this statement on the device at the LSP ingress. In ultimate-hop popping, the MPLS label is popped from the IP packet at the PE router. The IP address is checked in a second address lookup (also at the PE router), and then the packet is forwarded to its destination.

Be aware of the following platform requirements and restrictions:

- UHP LSPs using VT interfaces—Supported on all M Series, MX Series, T Series, and TX Matrix routers.
- UHP LSPs using LSI interfaces—Supported on MX 3D Series routers only.
- UHP LSP requirements for the egress PE device—For M Series and T Series routers, a VT interface is needed.
- UHP LSPs and Layer 3 VPNs—UHP LSPs are supported for Layer 3 VPNs configured on MX 3D Series routers only.
- UHP LSPs and VPLS—UHP LSPs are supported for VPLS configured on MX 3D Series routers only. You must configure the "no-tunnel-services" on page 11605 statement at the `[edit routing-instances routing-instance-name protocols vpls]` hierarchy level.

Default

Ultimate-hop popping is disabled by default on LSPs. Penultimate-hop popping is the default behavior. In penultimate-hop popping, the final MPLS label is popped from the IP packet at the last provider router in the network before being forwarded to the PE router. The PE router receives the packet and checks the IP address, and then the packet is forwarded to its destination.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring Ultimate-Hop Popping for LSPs

explicit-null (Protocols MPLS) | 6655

ultrahigh-power

IN THIS SECTION

- Syntax | 30325
- Hierarchy Level | 30325
- Description | 30326
- Default | 30326
- Required Privilege Level | 30326
- Release Information | 30326

Syntax

```
ultrahigh-power;
```

Hierarchy Level

```
[edit poe interface (all | interface-name)]
```

Description

Configure a four-pair PoE port to deliver up to 95 W of power. The four-pair PoE standard allows for power to be delivered over all four pairs of wires in a standard CAT5/6 RJ-45 Ethernet cable.

When you configure the four-pair PoE port with the `ultrahigh-power` option, the port provides up to 95 W of power, even if the powered device does not consume that amount of power.

Default

By default the PoE port operates in two-pair mode and provides power up to 30 W in accordance with the IEEE 802.3at PoE+ standard.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Configuring PoE Interfaces on EX Series Switches

unattended-boot

IN THIS SECTION

- Syntax | 30327
- Hierarchy Level | 30327
- Description | 30327
- Default | 30327
- Required Privilege Level | 30327

Release Information | 30327

Syntax

```
unattended-boot;
```

Hierarchy Level

```
[edit system]
```

Description

Set the switch to unattended mode for U-Boot to prevent unauthorized access to the system before the JUNOS OS login prompt appears. In unattended mode, access to the loader CLI is blocked, as well as recovery mechanisms such as password recovery by using single-user mode and booting the switch by using a USB flash drive. In order to access the CLI in U-Boot mode, the user must enter a boot-loader password that has been previously configured.

NOTE: If the root password is lost while the switch is in unattended mode, the switch must be reset to the factory default configuration using the LCD panel. For more information see *Reverting to the Default Factory Configuration for the EX Series Switch*.

Default

Unattended mode is not enabled by default.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X51-D20.

RELATED DOCUMENTATION

Using Unattended Mode for U-Boot to Prevent Unauthorized Access

boot-loader-authentication | 3008

unconditional-src-learn

IN THIS SECTION

- Syntax | 30328
- Hierarchy Level | 30328
- Description | 30328
- Required Privilege Level | 30329
- Release Information | 30329

Syntax

```
unconditional-src-learn;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

Enables the router to learn IP addresses from nonvalidated sources when proxy Address Resolution Protocol (ARP) is configured.

CAUTION: By default, the router learns IP addresses from validated sources only. When this statement is configured and proxy ARP is enabled on an unnumbered interface, the router responds to ARP requests from any IP address, which might lead to exploitable

information disclosure. An attacker can poison the ARP cache and create a fake forwarding table entry for an IP address, effectively creating a denial of service for that subscriber or interface. Therefore, exercise caution when configuring this statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4R16.

RELATED DOCUMENTATION

Understanding Proxy ARP

Example: Configuring Proxy ARP on an EX Series Switch

unconfigured-peer-graceful-restart

IN THIS SECTION

- Syntax | 30329
- Hierarchy Level | 30330
- Description | 30330
- Required Privilege Level | 30330
- Release Information | 30330

Syntax

```
unconfigured-peer-graceful-restart;
```

Hierarchy Level

```
[edit protocols bgp]
```

Description

When `set protocols bgp group` *`group-name`* `allow network` is configured to accept dynamic BGP sessions, `unconfigured-peer-graceful-restart` statement should be configured to avoid traffic drop during graceful restart or graceful Routing Engine switchover.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

underlying-interface

IN THIS SECTION

- Syntax | 30331
- Hierarchy Level | 30331
- Description | 30331
- Options | 30331
- Required Privilege Level | 30331
- Release Information | 30332

Syntax

```
underlying-interface interface-name;
```

Hierarchy Level

```
[edit interfaces pp0 unit logical-unit-number pppoe-options],
[edit interfaces demux0 unit logical-unit-number demux-options],
[edit logical-systems logical-system-name interfaces demux0 unit logical-unit-number demux-
options],
[edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number pppoe-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name interfaces
demux0 unit logical-unit-number demux-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name interfaces pp0
unit logical-unit-number pppoe-options]
```

Description

Configure the interface on which PPP over Ethernet is running.

For demux interfaces, configure the underlying interface on which the demultiplexing (demux) interface is running.

Options

interface-name—Name of the interface on which PPP over Ethernet or demux is running. For example, `at-0/0/1.0` (ATM VC), `fe-1/0/1.0` (Fast Ethernet interface), `ge-2/0/0.0` (Gigabit Ethernet interface), `ae1.0` (for IP demux on an aggregated Ethernet interface), or `ae1` (for VLAN demux on an aggregated Ethernet interface).

> **NOTE**: Demux interfaces are currently supported on Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet interfaces, or aggregated Ethernet devices.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for aggregated Ethernet added in Junos OS Release 9.4.

RELATED DOCUMENTATION

- Configuring an IP Demultiplexing Interface
- Configuring a VLAN Demultiplexing Interface
- Configuring the PPPoE Underlying Interface

underlying-interface (ANCP)

IN THIS SECTION

- Syntax | 30332
- Hierarchy Level | 30332
- Description | 30333
- Options | 30333
- Required Privilege Level | 30333
- Release Information | 30333

Syntax

```
underlying-interface underlying-interface-name;
```

Hierarchy Level

```
[edit protocols ancp interfaces interface-set interface-set-name]
```

Description

Configure the underlying interface on which the VLAN demux interface is running. The VLAN demux interface is the underlying interface for the PPPoE sessions controlled by ANCP.

Options

underlying-interface-name	Name of the underlying interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

ANCP Agent Neighbors and Operations

ANCP Agent Neighbors and Operations

underlying-interface (demux0)

IN THIS SECTION

- Syntax | 30334
- Hierarchy Level | 30334
- Description | 30334
- Options | 30334
- Required Privilege Level | 30334
- Release Information | 30335

Syntax

```
underlying-interface underlying-interface-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 interface-name unit unit logical-unit-number
demux-options]
```

Description

Configure the underlying interface on which the demultiplexing (demux) interface is running.

> **CAUTION**: Before you make any changes to the underlying interface for a demux0 interface, you must ensure that no subscribers are currently present on that underlying interface. If any subscribers are present, you must remove them before you make changes.

Options

underlying-interface-name—Either the specific name of the interface on which the DHCP discover packet arrives or one of the following interface variables:

- `$junos-underlying-interface` when configuring dynamic IP demux interfaces.
- `$junos-interface-ifd-name` when configuring dynamic VLAN demux interfaces.

The variable is used to specify the underlying interface when a new demux interface is dynamically created. The variable is dynamically replaced with the underlying interface that DHCP supplies when the subscriber logs in.

> **NOTE**: Logical demux interfaces are currently supported on Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Support for aggregated Ethernet introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles

Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

Junos OS Network Interfaces Library for Routing Devices

underlying-interface (Dynamic PPPoE)

IN THIS SECTION

- Syntax | 30335
- Hierarchy Level | 30336
- Description | 30336
- Options | 30336
- Required Privilege Level | 30336
- Release Information | 30336

Syntax

```
underlying-interface interface-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" pppoe-options]
```

Description

In a dynamic profile, configure the underlying interface on which the router creates the dynamic PPPoE logical interface.

Options

interface-name—Variable used to specify the name of the underlying interface on which the PPPoE logical interface is dynamically created. In the `underlying-interface` *interface-name* statement for dynamic PPPoE logical interfaces, you must use the predefined variable `$junos-underlying-interface` in place of *interface-name*. When the router creates the dynamic PPPoE interface, the `$junos-underlying-interface` predefined variable is dynamically replaced with the name of the underlying interface supplied by the network when the subscriber logs in.

> **CAUTION**: Before you make any changes to the underlying interface for a demux0 interface, you must ensure that no subscribers are currently present on that underlying interface. If any subscribers are present, you must remove them before you make changes.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *Configuring a PPPoE Dynamic Profile*
>
> *Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview*

unframed | no-unframed (Interfaces)

IN THIS SECTION

- Syntax | 30337
- Hierarchy Level | 30337
- Description | 30337
- Required Privilege Level | 30337
- Release Information | 30337

Syntax

```
(unframed | no-unframed);
```

Hierarchy Level

```
[edit interfaces interface-name  t3-options]
```

Description

Enable or disable framing for the T3 interface on a 1-Port Clear Channel DS3/E3 GPIM on an SRX Series Firewall. By default, unframed mode is enabled. Select no-unframed to enable framing. Select unframed to return to the default mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring a T3 Interface

unicast-ethernet

IN THIS SECTION

- Syntax | 30338
- Hierarchy Level | 30338
- Description | 30338
- Options | 30339
- Required Privilege Level | 30339
- Release Information | 30339

Syntax

```
unicast-ethernet {
    destination-mac destination-mac-address destination-mac-mask destination-mac-mask;
    ethertype ethertype;
    source-mac source-mac-address source-mac-mask source-mac-mask;
}
```

Hierarchy Level

```
[edit protocols oam fabric flow-specs flow-specification-name]
```

Description

Configure unicast Ethernet parameters for a flow specification used in internal fabric monitoring.

Options

destination-mac *destination-mac-address*

(Optional) MAC address of the destination fabric maintenance endpoint (FMEP) that is configured in the flow specification.

destination-mac-mask *destination-mac-mask*

(Optional) MAC address range mask for the destination FMEP that is configured in the flow specification. You cannot configure the `destination-mac-mask` parameter unless you also configure the `destination-mac` parameter.

The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination MAC address parameter. The mask must be contiguous.

- **Values:** Maximum of 6 contiguous masked bits, or a mask range of 64

ethertype *ethertype*

Valid Ethernet type other than the IPv4 Ethertype.

- **Values:** 2 bytes in decimal or hexadecimal notation

source-mac *source-mac-address*

(Optional) MAC address of the source FMEP or interface that is configured in the flow specification.

source-mac-mask *source-mac-mask*

(Optional) MAC address range mask for the source FMEP that is configured in the flow specification. You cannot configure the `source-mac-mask` parameter unless you also configure the `source-mac` parameter.

The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination MAC address parameter. The mask must be contiguous.

- **Values:** Maximum of 6 contiguous masked bits, or a mask range of 64

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | **29924**

unicast-ethernet-ipv4

IN THIS SECTION

- Syntax | **30340**
- Hierarchy Level | **30341**
- Description | **30341**
- Options | **30341**
- Required Privilege Level | **30342**
- Release Information | **30342**

Syntax

```
unicast-ethernet-ipv4 {
    source-ip source-ip-address source-ip-mask source-ip-mask;
    destination-ip destination-ip-address destination-ip-mask destination-ip-mask;
    ip-proto protocol;
    source-l4-port source-l4-port-number;
    destination-l4-port destination-l4-port-number;
}
```

Hierarchy Level

```
[edit protocols oam fabric flow-specs flow-specification-name]
```

Description

Configure unicast parameters for the Ethernet IPv4 protocol type in the internal fabric monitoring flow specifications.

Options

destination-ip *destination-ip-address*

Unicast IPv4 address of the destination fabric maintenance endpoint (FMEP) that is configured in the fabric maintenance association (FMA).

- **Syntax:** Dotted decimal notation without the prefix

destination-l4-port *destination-l4-port-number*

L4 TCP or UDP port for the destination fabric maintenance endpoint (FMEP).

- **Values:** 2-byte decimal or hexadecimal notation

ip-proto *protocol*

IPv4 protocol type.

- **Values:** 1-byte value in decimal or hexadecimal notation

destination-ip-mask *destination-ip-mask*

IPv4 address range mask for the destination FMEP that is configured in the flow specification.

The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination IPv4 address parameter. The mask must be contiguous.

- **Syntax:** Dotted decimal notation without the prefix
- **Values:** Maximum of 6 contiguous masked bits, or a mask range of 64

source-ip *source-ip-address*

Unicast IPv4 address of the source fabric maintenance endpoint (FMEP) that is configured in the fabric maintenance association (FMA).

- **Syntax:** Dotted decimal notation without the prefix

source-l4-port *source-l4-port-number*

TCP or UDP port for the source fabric maintenance endpoint (FMEP).

- **Values:** 2-byte decimal or hexadecimal notation

source-ip-mask
source-ip-mask

IPv4 address range mask for the source FMEP that is configured in the flow specification.

The masked bits denote the range that needs to be generated. The other bits obtain their values from the corresponding source or destination IPv4 address parameter. The mask must be contiguous.

- **Syntax:** Dotted decimal notation without the prefix
- **Values:** Maximum of 6 contiguous masked bits, or a mask range of 64

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Overview of Internal Fabric Monitoring

Configuring a Fabric Maintenance Association

Configuring Flow Specifications

fabric (OAM) | **6761**

ping fabric multicast-flow | **12539**

ping fabric unicast-flow | **12542**

traceroute fabric unicast-flow | **29924**

unicast-in-lpm

IN THIS SECTION

Syntax | **30343**

- Hierarchy Level | 30343
- Description | 30343
- Required Privilege Level | 30344
- Release Information | 30344

Syntax

```
unicast-in-lpm;
```

Hierarchy Level

```
[edit chassis forwarding-options lpm-profile]
```

Description

For the Unified Forwarding Table feature, specify to store all unicast IPv4 and IPv6 entries with prefixes with lengths equal to or less than 64 in the table for longest prefix match (LPM) entries, thereby freeing up space in the Layer 3 host table. Only unicast entries can be moved to the LPM table. Multicast entries must be stored in the Layer 3 host table.

You can also configure this statement in conjunction with the "prefix-65-127-disable" on page 13064 statement, which allocates no memory for IPv6 prefixes with lengths in the range /65 through /127. Together, these two statements allocate more space for unicast IPv4 and IPv6 entries with prefix lengths equal to or less than 64.

> **NOTE**: On QFX5110 switches running Junos OS Release 18.1R1 and higher, the `unicast-in-lpm` and `prefix-65-127-disable` statements cannot be configured at the same time.
>
> This statement is not supported on QFX5200 switches.

> **NOTE**: This statement is supported only on the `lpm-profile`.
>
> This statement is not supported on QFX5200 switches.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1x53-D30.

RELATED DOCUMENTATION

Understanding the Unified Forwarding Table

unicast-mode

IN THIS SECTION

- Syntax | 30344
- ACX710 | 30345
- Hierarchy Level | 30345
- Description | 30345
- Options | 30345
- Required Privilege Level | 30346
- Release Information | 30346

Syntax

```
unicast-mode {
    clock-source ip-address {
        local-ip-address local-ip-address;
            asymmetry  number;
        }
    transport ipv4;
```

```
    clock-ipv6-source  master-ipv6-address {
        local-ip-address local-ip-address;
        }
    transport ipv6;
}
```

ACX710

```
unicast-mode {
    announce-timeout announce-timeout;
    grantor;
    local-priority local-priority;
    master-only;
    remote-ip-address ip-address {
        local-ip-address;
        announce-interval announce-interval;
        asymmetry asymmetry;
        delay-request delay-request;
        dscp dscp;
        grant-duration grant-duration;
        sync-interval sync-interval;
    }
    transport (ipv4 | ipv6);
}
```

Hierarchy Level

```
[edit protocols ptp client interface interface-name]
```

Description

Configure the timeReceiver in unicast mode. You can set this option when PTP unicast mode of messaging is needed.

Options

announce-timeout-value The announce timeout value for announce interval messages.

- **Range:** 3 through 10
- **Default:** 3

grantor Specify that the selected interface port can be grantor. The grantor attribute enables the port to allow for the grant-port property and can provide PTP message services. It allows the port to accept and handle unicast negotiation requests.

Priority value Priority value.

- **Range:** 1 through 255
- **Default:** 128

master-only Specify the timeTransmitter only option. When you set the master-only option, it is not possible to set the interface port as a timeReceiver clock.

transport Encapsulation for PTP packet transport

- Values:
  - ipv4—Use IPv4 as transport
  - ipv6—Use IPv6 as transport

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

IEEE 1588v2 Precision Timing Protocol (PTP)

unicast-negotiation

IN THIS SECTION

- Syntax | 30347
- Hierarchy Level | 30347
- Description | 30347
- Required Privilege Level | 30347
- Release Information | 30348

Syntax

```
unicast-negotiation;
```

Hierarchy Level

```
[edit protocols ptp]
```

Description

Configure unicast negotiation. Unicast negotiation is a method by which the announce, synchronization, and delay response packet rates are negotiated between the timeTransmitter and the timeReceiver before a PTP session is established.

> **NOTE**:
> When unicast negotiation is enabled, you cannot commit any packet rate–related configuration.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Configuring Precision Time Protocol

Example: Configuring Precision Time Protocol

Precision Time Protocol Overview

IEEE 1588v2 Precision Timing Protocol (PTP)

unicast-reverse-path

IN THIS SECTION

- Syntax | 30348
- Hierarchy Level | 30348
- Description | 30349
- Options | 30349
- Required Privilege Level | 30349
- Release Information | 30349

Syntax

```
unicast-reverse-path (active-paths | feasible-paths);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table],
[edit routing-instances routing-instance-name instance-type name routing-options forwarding-
```

```
table],
[edit routing-options forwarding-table]
```

Description

Control the operation of unicast reverse-path-forwarding check. This statement enables the RPF check to be used when routing is asymmetrical.

Options

`active-paths`—Consider only active paths during the unicast reverse-path check.

`feasible-paths`—Consider all feasible paths during the unicast reverse-path check.

- **Default:** If you omit the `unicast-reverse-path` statement, only the active paths to a particular destination are considered.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances added in Junos OS Release 8.3.

> **NOTE**: This feature is not supported on the EX4300 switch, even though it is available on the device.

RELATED DOCUMENTATION

> *Example: Configuring Unicast RPF (On a Router)*
>
> *Enabling Unicast Reverse-Path Forwarding Check for VPNs*

unicast (Route Target Community)

IN THIS SECTION

- Syntax | 30350
- Hierarchy Level | 30350
- Description | 30350
- Options | 30350
- Required Privilege Level | 30351
- Release Information | 30351

Syntax

```
unicast {
    receiver;
    sender;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mvpn
route-target import-target],
[edit routing-instances routing-instance-name protocols mvpn route-target import-target]
```

Description

Specify the same target community configured for unicast.

Options

`receiver`—Specify the unicast target community used when importing receiver site routes.

`sender`—Specify the unicast target community used when importing sender site routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring VRF Route Targets for Routing Instances for an MBGP MVPN

unicast-stream-limit (Protocols AMT)

IN THIS SECTION

- Syntax | 30351
- Hierarchy Level | 30352
- Description | 30352
- Options | 30352
- Required Privilege Level | 30352
- Release Information | 30352

Syntax

```
unicast-stream-limit ;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols amt relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols amt
relay],
[edit protocols amt relay],
[edit routing-instances routing-instance-name protocols amt relay]
```

Description

Set the upper limit for unicast streams (s,g intf).

Options

number—Maximum number of data unicast streams that can be created on the system.

- **Range:** 0 through 4294967295
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

> *Configuring the AMT Protocol*

unicast-umh-election

IN THIS SECTION

- Syntax | 30353
- Hierarchy Level | 30353
- Description | 30353
- Required Privilege Level | 30353
- Release Information | 30353

Syntax

```
unicast-umh-election;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols mvpn]
```

Description

Configure a router to use the unicast route preference to determine the single forwarder election.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN

mvpn (NG-MVPN) | **11143**

unicast (Virtual Tunnel in Routing Instances)

IN THIS SECTION

- Syntax | **30354**
- Hierarchy Level | **30354**
- Description | **30354**
- Default | **30355**
- Required Privilege Level | **30355**
- Release Information | **30355**

Syntax

```
unicast;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name interface vt-
fpc/pic/port.unit-number],
[edit routing-instances routing-instance-name interface vt-fpc/pic/port.unit-number]
```

Description

In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure the virtual tunnel (VT) interface to be used for unicast traffic only.

Default

If you omit this statement, the VT interface can be used for both multicast and unicast traffic.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs

Example: Configuring MBGP MVPN Extranets

unicast-tunnel-name-regular-expression

IN THIS SECTION

- Syntax | 30355
- Hierarchy Level | 30356
- Description | 30356
- Options | 30356
- Release Information | 30356

Syntax

```
unicast-tunnel-name-regular-expression {
    preferred;
```

```
        regular-expression;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name provider-tunnel ingress-replication],
[edit logical-systems name tenants name routing-instances name provider-tunnel ingress-
replication],
[edit routing-instances name provider-tunnel ingress-replication],
[edit tenants name routing-instances name provider-tunnel ingress-replication]
```

Description

Configure a regular expression for a unicast tunnel to match specific RSVP tunnel LSP names.

> **NOTE**: Regular expressions support 1000 existing P2P LSPs using the default inet.3 table.

Options

preferred Fall-back to default behavior of selecting the active LSP.

regular-expression Regular expression to match unicast tunnels to RSVP tunnel LSP names.

Release Information

Statement introduced in Junos OS Release 24.4R1.

unidirectional

IN THIS SECTION

Syntax | 30357

- Hierarchy Level | 30357
- Description | 30357
- Default | 30357
- Required Privilege Level | 30357
- Release Information | 30358

Syntax

```
unidirectional;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit logical-systems logical-system-name interfaces interface-name]
```

Description

Create two new, unidirectional (transmit-only and receive-only) physical interfaces subordinate to the original parent interface. Unidirectional links are currently supported only on 10-Gigabit Ethernet interfaces on the following hardware:

- 4-port 10-Gigabit Ethernet DPC on the MX960 router
- 10–Gigabit Ethernet IQ2 PIC and 10–Gigabit Ethernet IQ2E PIC on the T Series router

Default

Disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

- Understanding Unidirectional Traffic Flow on Physical Interfaces
- Enabling Unidirectional Traffic Flow on Physical Interfaces

unidirectional-session-refreshing

IN THIS SECTION

- Syntax | 30358
- Hierarchy Level | 30358
- Description | 30359
- Required Privilege Level | 30359
- Release Information | 30359

Syntax

```
unidirectional-session-refreshing;
```

Hierarchy Level

```
[edit security zones security-zone]
```

Description

Use unidirectional session refreshing on a zone to enable two options for a session. Refresh a session by any packet from any directions. This is a default behavior. Refresh a session by only the packets in the initial direction.

> **NOTE**: USR (Unidirectional Session Refresh) is skipped for PPTP and IKE ALG data sessions in cases where USR is enabled on only one security zone, but data traffic originates from the opposite zone. For instance, if USR is enabled on the "Trust" zone but disabled on the "Untrust" zone, then for a PPTP ALG data session originating from the "Untrust" zone (where traffic flows unidirectionally from public to private), USR will be skipped.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

- security-zone | **16516**
- show security zones | **25232**

unified-access-control (Services)

IN THIS SECTION

- Syntax | **30360**
- Hierarchy Level | **30361**
- Description | **30361**
- Required Privilege Level | **30361**

- Release Information | **30361**

Syntax

```
unified-access-control {
    captive-portal redirect-policy-name{
        redirect-traffic (all | unauthenticated);
        redirect-url redirect-url;
    }
    certificate-verification [ optional | required | warning ];
    infranet-contoller host-name {
        address ip-address;
        ca-profile [ca-profile];
        interface interface-name;
        password password;
        port port-number;
        server-certificate-subject subject;
    }
    interval seconds;
    test-only-mode;
    timeout seconds;
    timeout-action (close | no-change | open);
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            (no-world-readable | world-readable);
            size maximum-file-size;
        }
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Use this statement to configure the SRX Series Firewall to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series Firewall enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

unified-policy-explicit-match

IN THIS SECTION

- Syntax | 30361
- Hierarchy Level | 30362
- Description | 30362
- Required Privilege Level | 30362
- Release Information | 30362

Syntax

```
unified-policy-explicit-match;
```

Hierarchy Level

```
[edit security policies]
```

Description

Enable explicit match in a unified policy.

- In a unified policy, when you enable explicit match, the session is dropped if a dependent application does not have any matching policy.
- In a unified policy, when explicit match is not configured, then implicit match is applied. That is, if the dependent application does not have a matching policy, then the security policy that is configured for the base application is applied.

By default, the unified policies enforce implicit rules on dependent applications.

Required Privilege Level

security— To view this statement in the configuration.

security-control— To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Unified Security Policies

Unified Security Policies

unified-policy max-lookups

IN THIS SECTION

Syntax | 30363

- Hierarchy Level | 30363
- Description | 30363
- Default | 30363
- Required Privilege Level | 30364
- Release Information | 30364

Syntax

```
unified-policy max-lookups max-lookups-number;
```

Hierarchy Level

```
[edit logical-systems name security policies],
[edit logical-systems name tenants name security policies],
[edit security policies]
[edit tenants tenant-name security policies]
```

Description

Configure maximum number of policy lookups, when micro-applications reach the *transaction final* state, during the dynamic application classification process.

Application classification with micro-applications does not reach the *final match* state because the micro-application keeps changing for the session. When the micro-application changes from one transaction to another, a policy lookup is attempted unlimited number of times. The *final match* state for the application is required for the policy lookup and for processing the policy. You can limit the number of transactions for the micro application by configuring the `unified-policy max-lookups` statement.

Setting the value of `unified-policy max-lookups` as 0 indicates unlimited number of policy lookups.

Range: 0 through 25

Default

The default value for `unified-policy max-lookups` is 2.

Required Privilege Level

flow-tap

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Configure Micro-Applications in Unified Policies

unique-nas-port (Access)

IN THIS SECTION

- Syntax | 30364
- Hierarchy Level | 30365
- Description | 30365
- Options | 30365
- Required Privilege Level | 30365
- Release Information | 30365

Syntax

```
unique-nas-port {
    chassis-id chassis-id;
    chassis-id-width chassis-id-width;
}
```

Hierarchy Level

```
[edit access radius-options]
```

Description

Configure the router to provide a unique value for the NAS-Port attribute (RADIUS attribute 5) of each subscriber. You can configure a NAS-Port value that is unique within the router only, or unique across the different MX routers in the network.

Options

chassis-id Value for the `chassis-id` part of NAS-Port attribute. To ensure that NAS-Port values are unique across all MX series routers in the network, you must specify a unique `chassis-id` for each MX router.

- **Range:** 0-127 bits

chassis-id-width Number of bits used to uniquely identify the chassis across the MX series routers in the network. All routers must use the same `chassis-id-width`. If you do not configure a `chassis-id-width`, the resulting NAS-Port attribute is unique within the router only.

- **Range:** 1-7 bits

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

RADIUS NAS Port Attributes and Options

RADIUS Servers and Parameters for Subscriber Access

unit (Sampling)

IN THIS SECTION

- Syntax | 30366
- Hierarchy Level | 30366
- Description | 30367
- Options | 30367
- Required Privilege Level | 30367
- Release Information | 30367

Syntax

```
unit logical-unit-number {
    family inet {
        address address {
            destination destination-address;
        }
        filter {
            group filter-group-number;
            input filter-name;
            output filter-name;
        }
        sampling direction;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 16,384

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Monitoring, Sampling, and Collection Services Interfaces User Guide

unit (QFX Series)

IN THIS SECTION

- Syntax | 30368
- Hierarchy Level | 30368
- Description | 30369
- Default | 30369
- Options | 30369

Required Privilege Level | 30369
Release Information | 30369

Syntax

The `ethernet-switching` and `fibre-channel` statements and all of their substatements are not supported on OCX Series switches.

```
unit logical-unit-number {
    family {
        ethernet-switching {
            filter input filter-name;
            filter output filter-name;
            native-vlan-id vlan-id;
            port-mode mode;
            vlan {
                members [ (all | names | vlan-ids) ];
            }
        }
        fibre-channel {
            port-mode (f-port | np-port);
        }
        inet {
            address address {
                primary;
            }
            filter input filter-name;
            filter output filter-name;
            primary;
            targeted-broadcast;
        }
```

Hierarchy Level

```
[edit interfaces (QFX Series) interface-name],
[edit interfaces (QFX Series) interface-range interface-range-name]
```

Description

> **NOTE**: The `ethernet-switching` and `fibre-channel` statements and all of their substatements are not supported on OCX Series switches.

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Default

You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 16,384

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

> *Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches*
>
> *Configuring Link Aggregation*
>
> Junos OS Network Interfaces Library for Routing Devices

unit (interfaces)

IN THIS SECTION

- Syntax | 30370
- Hierarchy Level | 30379
- Description | 30379
- Options | 30379
- Required Privilege Level | 30380
- Release Information | 30380

Syntax

```
unit logical-unit-number {
    accept-source-mac {
        mac-address mac-address {
            policer {
                input cos-policer-name;
                output cos-policer-name;
            }
        }
    }
    accounting-profile name;
    advisory-options {
        downstream-rate rate;
        upstream-rate rate;
    }
    allow-any-vci;
    atm-scheduler-map (map-name | default);
    auto-configure {
        agent-circuit-identifier {
            dynamic-profile profile-name;
        }
        line-identity {
            include {
                accept-no-ids;
                circuit-id;
```

```
                remote-id;

            }
            dynamic-profile profile-name;
        }
    }
    backup-options {
        interface interface-name;
    }
    bandwidth rate;
    cell-bundle-size cells;
    clear-dont-fragment-bit;
    compression {
        rtp {
            maximum-contexts number <force>;
            f-max-period number;
            queues [queue-numbers];
            port {
                minimum port-number;
                maximum port-number;
            }
        }
    }
    compression-device interface-name;
    copy-tos-to-outer-ip-header;
    demux {
        inet {
            address-source address;
            auto-configure {
                address-ranges {
                    authentication {
                        password password-string;
                        username-include {
                            auth-server-realm realm-string;
                            delimiter delimiter-character;
                            domain-name domain-name;
                            interface-name;
                            source-address;
                            user-prefix user-prefix-string;
                        }
                    }
                    dynamic-profile profile-name {
                        network ip-address {
```

```
                        range name {
                            low lower-limit;
                            high upper-limit;
                        }
                    }
                }
            }

        }
    }
    inet6 {
        address-source address;
        auto-configure {
            address-ranges {
                authentication {
                    password password-string;
                    username-include {
                        auth-server-realm realm-string;
                        delimiter delimiter-character;
                        domain-name domain-name;
                        interface-name;
                        source-address;
                        user-prefix user-prefix-string;
                    }
                }
                dynamic-profile profile-name {
                    network ip-address {
                        range name {
                            low lower-limit;
                            high upper-limit;
                        }
                    }
                }
            }

        }
    }
}
demux-destination family;
demux-source family;
demux-options {
    underlying-interface interface-name;
}
```

```
    description text;
    etree-ac-role (leaf | root);
    interface {
        l2tp-interface-id name;
        (dedicated | shared);
    }
    dialer-options {
        activation-delay seconds;
        callback;
        callback-wait-period time;
        deactivation-delay seconds;
        dial-string [dial-string-numbers];
        idle-timeout seconds;
        incoming-map {
            caller caller-id | accept-all;
            initial-route-check seconds;
            load-interval seconds;
            load-threshold percent;
            pool pool-name;
            redial-delay time;
            watch-list {
                [routes];
            }
        }
    }
    disable;
    disable-mlppp-inner-ppp-pfc;
    dlci dlci-identifier;
    drop-timeout milliseconds;
    dynamic-call-admission-control {
        activation-priority priority;
        bearer-bandwidth-limit kilobits-per-second;
    }
    encapsulation type;
    epd-threshold cells plp1 cells;
    family family-name {
        ... the family subhierarchy appears after the main [edit interfaces interface-name unit
logical-unit-number] hierarchy ...
    }
    fragment-threshold bytes;
    host-prefix-only;
    inner-vlan-id-range start start-id end end-id;
    input-vlan-map {
```

```
        (pop | pop-pop | pop-swap | push | push-push | swap |
        swap-push | swap-swap);
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        tag-protocol-id tpid;
        vlan-id number;
    }
    interleave-fragments;
    inverse-arp;
    layer2-policer {
        input-policer policer-name;
        input-three-color policer-name;
        output-policer policer-name;
        output-three-color policer-name;
    }
    link-layer-overhead percent;
    minimum-links number;
    mrru bytes;
    multicast-dlci dlci-identifier;
    multicast-vci vpi-identifier.vci-identifier;
    multilink-max-classes number;
    multipoint;
    oam-liveness {
        up-count cells;
        down-count cells;
    }
    oam-period (disable | seconds);
    output-vlan-map {
        (pop | pop-pop | pop-swap | push | push-push | swap |
        swap-push | swap-swap);
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        tag-protocol-id tpid;
    }
    passive-monitor-mode;
    peer-unit unit-number;
    plp-to-clp;
    point-to-point;
    ppp-options {
        mru size;
        mtu (size | use-lower-layer);
        chap {
            access-profile name;
```

```
            default-chap-secret name;
            local-name name;
            passive;
        }
        compression {
            acfc;
            pfc;
        }
        dynamic-profile profile-name;
        ipcp-suggest-dns-option;
        lcp-restart-timer milliseconds;
        loopback-clear-timer seconds;
        ncp-restart-timer milliseconds;
        pap {
            access-profile name;
            default-pap-password password;
            local-name name;
            local-password password;
            passive;
        }
    }
    pppoe-options {
        access-concentrator name;
        auto-reconnect seconds;
        (client | server);
        service-name name;
        underlying-interface interface-name;
    }
    pppoe-underlying-options {
        access-concentrator name;
        direct-connect;
        dynamic-profile profile-name;
        max-sessions number;
    }
    proxy-arp;
    service-domain (inside | outside);
    shaping {
        (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
burst length);
        queue-length number;
    }
    short-sequence;
    targeted-distribution;
```

```
    transmit-weight number;
    (traps | no-traps);
    trunk-bandwidth rate;
    trunk-id number;
    tunnel {
        backup-destination address;
        destination address;
        key number;
        routing-instance {
            destination routing-instance-name;
        }
        source source-address;
        ttl number;
    }
    vci vpi-identifier.vci-identifier;
    vci-range start start-vci end end-vci;
    vpi vpi-identifier;
    vlan-id number;
    vlan-id-range number-number;
    vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
    family family {
        accounting {
            destination-class-usage;
            source-class-usage {
                (input | output | input output);
            }
        }
        access-concentrator name;
        address address {
            ... the address subhierarchy appears after the main [edit interfaces interface-name
unit logical-unit-number family family-name] hierarchy ...
        }
        bundle interface-name;
        core-facing;
        demux-destination {
            destination-prefix;
        }
        demux-source {
            source-prefix;
        }
        direct-connect;
        duplicate-protection;
        dynamic-profile profile-name;
```

```
filter {
    group filter-group-number;
    input filter-name;
    input-list [filter-names];
    output filter-name;
    output-list [filter-names];
}
interface-mode (access | trunk);
ipsec-sa sa-name;
keep-address-and-control;
mac-validate (loose | strict);
max-sessions number;
mtu bytes;
multicast-only;
no-redirects;
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
primary;
protocols [inet iso mpls];
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check {
    fail-filter filter-name
    mode loose;
}
sampling {
    input;
    output;
}
service {
    input {
        post-service-filter filter-name;
        service-set service-set-name <service-filter filter-name>;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
```

```
        service-name-table table-name
        targeted-options {
            backup backup;
            group group;
            primary primary;
            weight ($junos-interface-target-weight | weight-value);
        }
        (translate-discard-eligible | no-translate-discard-eligible);
        (translate-fecn-and-becn | no-translate-fecn-and-becn);
        translate-plp-control-word-de;
        unnumbered-address interface-name destination address destination-profile profile-name;
        vlan-id number;
        vlan-id-list [number number-number];
        address address {
            arp ip-address (mac | multicast-mac) mac-address <publish>;
            broadcast address;
            destination address;
            destination-profile name;
            eui-64;
            primary-only;
            multipoint-destination address {
                dlci dlci-identifier;
                epd-threshold cells <plp1 cells>;
                inverse-arp;
                oam-liveness {
                    up-count cells;
                    down-count cells;
                }
                oam-period (disable | seconds);
                shaping {
                    (cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length
peak rate sustained rate);
                    queue-length number;
                }
                vci vpi-identifier.vci-identifier;
            }
            preferred;
            primary;
            (vrrp-group | vrrp-inet6-group) group-number {
                (accept-data | no-accept-data);
                advertise-interval seconds;
                authentication-type authentication;
                authentication-key key;
```

```
                fast-interval milliseconds;
                (preempt | no-preempt) {
                    hold-time seconds;
                }
                priority number;
                track {
                    interface interface-name {
                        bandwidth-threshold bits-per-second priority-cost number;
                    }
                    priority-hold-time seconds;
                    route ip-address/prefix-length routing-instance instance-name priority-
cost cost;
                }
                virtual-address [addresses];
                virtual-link-local-address ipv6-address;
                vrrp-inherit-from {
                    active-interface interface-name;
                    active-group group-number;
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit logical-systems logical-system-name interfaces interface-name],
[edit interfaces interface-set interface-set-name interface interface-name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 1,073,741,823 for demux, PPPoE, and pseudowire static interfaces. 0 through 16,385 for all other static interface types.

`etree-ac-role (leaf | root)`—To configure an interface as either leaf or root.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Range increased for static pseudowire interfaces to 1,073,741,823 in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Configuring Logical Interface Properties

Junos OS Services Interfaces Library for Routing Devices

unit (Aggregated Multiservices)

IN THIS SECTION

- Syntax | 30381
- Hierarchy Level | 30381
- Description | 30381
- Options | 30381
- Required Privilege Level | 30381
- Release Information | 30381

Syntax

```
unit interface-unit-number {
    family family;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure the logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

The remaining statements are explained separately. See CLI Explorer.

Options

interface-unit-number—Number of the logical unit.

> **NOTE**: Unit 0 is reserved and cannot be configured under the aggregated Multiservices interface (`ams`).

- **Range:** 1 through 16,384

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

interfaces (Aggregated Multiservices) | 9115

unit (CoS Interfaces)

IN THIS SECTION

- Syntax | 30382
- Hierarchy Level | 30383
- Description | 30383
- Options | 30383
- Required Privilege Level | 30383
- Release Information | 30383

Syntax

```
unit logical-unit-number {
    classifiers {
        type (classifier-name | default) family (mpls | all);
    }
    forwarding-class class-name;
    fragmentation-map map-name;
    input-traffic-control-profile profile-name shared-instance instance-name;
    output-traffic-control-profile profile-name shared-instance instance-name;
    per-session-scheduler;
    policy-map policy-map-name;
    policy-map-marking;
    rewrite-rules {
        dscp (rewrite-name | default);
        dscp-ipv6 (rewrite-name | default);
        exp (rewrite-name | default)protocol protocol-types;
        exp-push-push-push default;
```

```
        exp-swap-push-push default;
        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
        inet-precedence (rewrite-name | default);
    }
    scheduler-map map-name;
    shaping-rate rate;
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name]
```

Description

Configure class of service properties on a logical interface on a physical port.

> **NOTE**: Available options can vary by platform and software release (Junos OS or Junos OS Evolved). Use Feature Explorer to confirm platform and release support for specific features.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 16,384

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic

Configuring Rewrite Rules

Assigning CoS Components to Interfaces

unit (CoS Interfaces)

IN THIS SECTION

- Syntax | 30384
- Hierarchy Level | 30384
- Description | 30385
- Options | 30385
- Required Privilege Level | 30385
- Release Information | 30385

Syntax

```
unit logical-unit-number {
    classifiers {
        (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
    }
    forwarding-class class-name;
    rewrite-rules {
        (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
    }
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to use the physical device.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 16,385

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Assigning CoS Components to Interfaces

unit (Dynamic Demux Interface)

IN THIS SECTION

- Syntax | 30386
- Hierarchy Level | 30387
- Description | 30387
- Options | 30387
- Required Privilege Level | 30387

- Release Information | 30387

Syntax

```
unit logical-unit-number {
    demux-options {
        underlying-interface interface-name
    }
    family family {
        access-concentrator name;
        address address;
        demux-source {
            source-address;
        }
        direct-connect;
        duplicate-protection;
        dynamic-profile profile-name;
        filter {
            input filter-name;
            output filter-name;
        }
        mac-validate (loose | strict):
        max-sessions number;
        max-sessions-vsa-ignore;
        rpf-check {
            fail-filter filter-name;
            mode loose;
        }
        service-name-table table-name;
        short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>;
        unnumbered-address interface-name <preferred-source-address address>;
    }
    filter {
        input filter-name;
        output filter-name;
    }
}
vlan-id number;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0]
```

Description

Configure a dynamic logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Either the specific unit number of the interface or the unit number variable (`$junos-interface-unit`). The variable is used to specify the unit of the interface when a new demux interface is dynamically created. The static unit number variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

> *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*

unit (Dynamic Interface Sets)

IN THIS SECTION

- Syntax | 30388
- Hierarchy Level | 30388
- Description | 30388
- Options | 30388
- Required Privilege Level | 30389
- Release Information | 30389

Syntax

```
unit logical-unit-number {
    advisory-options {
        downstream-rate rate;
        upstream-rate rate;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-set interface-set-name interface interface-name]
```

Description

Apply the logical interface unit to the interface set.

Options

logical-unit-number—One of the following options:

- `$junos-underlying-interface-unit`—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- `$junos-interface-unit`—For dynamic demux and dynamic PPPoE interfaces, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP or PPP when it accesses the subscriber network.
- *value*—Specific unit number of the interface you want to assign to the dynamic-profile
- **Range:** 0 through 1,073,741,823.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Dynamic VLAN Subscriber Interfaces Based on Agent Circuit Identifier Information

Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile

Configuring an Interface Set of Subscribers in a Dynamic Profile

Agent Circuit Identifier-Based Dynamic VLANs Overview

Guidelines for Configuring Dynamic CoS for Subscriber Access

unit (Dynamic PPPoE)

IN THIS SECTION

- Syntax | 30390
- Hierarchy Level | 30391
- Description | 30391
- Options | 30392
- Required Privilege Level | 30392
- Release Information | 30392

Syntax

```
unit logical-unit-number {
    keepalives interval seconds;
    no-keepalives;
    pppoe-options {
        underlying-interface interface-name;
        server;
    }
    ppp-options {
        aaa-options aaa-options-name;
        authentication [ authentication-protocols ];
        mru size;
        mtu (size | use-lower-layer);
        chap {
            challenge-length minimum minimum-length maximum maximum-length;
        }
        ignore-magic-number-mismatch;
        initiate-ncp (ip | ipv6 | dual-stack-passive)
        ipcp-suggest-dns-option;
        mru size;
        mtu (size | use-lower-layer);
        on-demand-ip-address;
        pap;
        peer-ip-address-optional;
    }
```

```
    family inet {
        unnumbered-address interface-name;
        address address;
        service {
            input {
                service-set service-set-name {
                    service-filter filter-name;
                }
                post-service-filter filter-name;
            }
            output {
                service-set service-set-name {
                    service-filter filter-name;
                }
            }
        }
        filter {
            input filter-name {
                precedence precedence;
            }
            output filter-name {
                precedence precedence;
            }
        }
    }
    filter {
        input filter-name;
        output filter-name;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0]
```

Description

In a dynamic profile, configure a logical unit number for the dynamic PPPoE logical interface. You must configure a logical interface to be able to use the router.

Options

logical-unit-number—Variable used to specify the unit number when the PPPoE logical interface is dynamically created. In the `unit` *logical-unit-number* statement for dynamic PPPoE logical interfaces, you must use the predefined variable `$junos-interface-unit` in place of *logical-unit-number*. The `$junos-interface-unit` predefined variable is dynamically replaced with the unit number supplied by the router when the subscriber logs in.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring a PPPoE Dynamic Profile

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

unit (Dynamic Profiles Standard Interface)

IN THIS SECTION

- Syntax | 30393
- Hierarchy Level | 30396
- Description | 30396
- Options | 30396
- Required Privilege Level | 30396
- Release Information | 30397

Syntax

```
unit logical-unit-number {
    actual-transit-statistics;
    auto-configure {
        agent-circuit-identifier {
            dynamic-profile profile-name;
        }
        line-identity {
            include {
                accept-no-ids;
                circuit-id;
                remote-id;

            }
            dynamic-profile profile-name;
        }
    }
    dial-options {
        ipsec-interface-id name;
        l2tp-interface-id name;
        (shared | dedicated);
    }
    encapsulation  (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux | atm-
mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-vc-mux |
ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp |
ethernet | frame-relay-ccc | frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type | frame-
relay-ether-type-tcc | multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether | ppp-
over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls);
    family family {
        address address;
        demux-destination,
        filter {
            adf {
                counter;
                input-precedence precedence;
                not-mandatory;
                output-precedence precedence;
                rule rule-value;
            }
            input filter-name {
                precedence precedence;
```

```
                shared-name filter-shared-name;
            }
            output filter-name {
                precedence precedence;
                shared-name filter-shared-name;
            }
        }
        max-sessions number;
        max-sessions-vsa-ignore;
        rpf-check {
            fail-filter filter-name;
            mode loose;
        }
        service {
             input {
                 service-set service-set-name {
                     service-filter filter-name;
                }
                post-service-filter filter-name;
            }
            input-vlan-map {
                inner-tag-protocol-id tpid;
                inner-vlan-id number;
                (push | swap);
                tag-protocol-id tpid;
                vlan-id number;
            }
            output {
                 service-set service-set-name {
                     service-filter filter-name;
                }
            }
            output-vlan-map {
                inner-tag-protocol-id tpid;
                inner-vlan-id number;
                (pop | swap);
                tag-protocol-id tpid;
                vlan-id number;
            }
        }
        service-name-table table-name
        short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>;
        unnumbered-address interface-name <preferred-source-address address>;
```

```
    }
    filter {
        input filter-name {
            shared-name  filter-shared-name;
        }
        output filter-name {
            shared-name  filter-shared-name;
        }
    }
    host-prefix-only;
    keepalives {
        interval seconds;
        }
    ppp-options {
        aaa-options aaa-options-name;
        authentication [ authentication-protocols ];
        chap {
            challenge-length minimum minimum-length maximum maximum-length;
            local-name name;
        }
        ignore-magic-number-mismatch;
        initiate-ncp (dual-stack-passive | ipv6 | ip)
        ipcp-suggest-dns-option;
        mru size;
        mtu (size | use-lower-layer);
        on-demand-ip-address;
        pap;
        peer-ip-address-optional;
        local-authentication {
            password password;
            username-include {
                circuit-id;
                delimiter character;
                domain-name name;
                mac-address;
                remote-id;
            }
        }
    }
    service {
        pcef  pcef-profile-name {
            activate rule-name | activate-all;
        }
```

```
    }
    targeted-options {
        backup backup;
        group group;
        primary primary;
        weight ($junos-interface-target-weight | weight-value);
    }
    vlan-id number;
    vlan-tags outer  [tpid].vlan-id [inner [tpid].vlan-id];
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—The specific unit number of the interface you want to assign to the dynamic profile, or one of the following predefined variables:

- `$junos-underlying-interface-unit`—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- `$junos-interface-unit`—The unit number variable on a dynamic underlying VLAN interface for which you want to enable the creation of dynamic VLAN subscriber interfaces based on the ACI.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information

Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information

Agent Circuit Identifier-Based Dynamic VLANs Overview

unit (Dynamic Traffic Shaping)

IN THIS SECTION

- Syntax | 30397
- Hierarchy Level | 30398
- Description | 30398
- Options | 30398
- Required Privilege Level | 30398
- Release Information | 30399

Syntax

```
unit logical-unit-number {
    classifiers {
        type (classifier-name | default);
    }
    output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
    report-ingress-shaping-rate bps;
    rewrite-rules {
        dscp (rewrite-name | default);
        dscp-ipv6 (rewrite-name | default);
        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
```

```
        inet-precedence (rewrite-name | default);
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name],
[edit dynamic-profiles profile-name interfaces interface-set interface-set-name interface interface-
name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—One of the following options:

- `$junos-underlying-interface-unit`—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- `$junos-interface-unit`—For dynamic demux and dynamic PPPoE interfaces, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP or PPP when it accesses the subscriber network.
- *value*—Specific unit number of the interface you want to assign to the dynamic-profile
- **Range:** 0 through 16385. For demux and PPPoE interfaces, the unit numbers can range from 0 through 1,073,741,823.

The remaining statements are explained separately. The `classifiers`, `output-traffic-control-profile`, and `rewrite-rules` statements are not supported for interface sets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the `[edit dynamic-profiles profile-name class-of-service interfaces interface-set interface-set-name]` hierarchy level introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile

Configuring an Interface Set of Subscribers in a Dynamic Profile

unit (Interfaces Services)

IN THIS SECTION

- Syntax | 30399
- Hierarchy Level | 30400
- Description | 30400
- Options | 30400
- Required Privilege Level | 30400
- Release Information | 30400

Syntax

```
unit logical-unit-number {
    family inet {
        address address {
        }
        service {
            input {
                [ service-set service-set-name <service-filter filter-name> ];
                post-service-filter filter-name;
```

```
            }
            output {
                [ service-set service-set-name <service-filter filter-name> ];
            }
        }
        service-domain (inside | outside);
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name ]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 16,384

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Network Interfaces Library for Routing Devices

unit (Interfaces)

IN THIS SECTION

- Syntax | 30401
- Hierarchy Level | 30401
- Description | 30402
- Options | 30402
- Required Privilege Level | 30402
- Release Information | 30402

Syntax

```
unit logical-unit-number {
    family inet {
        ipsec-sa sa-name;
    }
    tunnel {
        backup-destination destination-address;
        destination destination-address;
        routing-instance {
            destination routing-instance-name;
        }
        source source-address;
        ttl number;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 16,384

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Encryption Interfaces

Junos OS Network Interfaces Library for Routing Devices

Junos OS Network Interfaces Library for Routing Devices

unit (Link-Multilink)

IN THIS SECTION

- Syntax | 30403
- Hierarchy Level | 30403
- Description | 30403

- Options | 30404
- Required Privilege Level | 30404
- Release Information | 30404

Syntax

```
unit logical-unit-number {
    disable-mlppp-inner-ppp-pfc;
    dlci dlci-identifier;
    drop-timeout milliseconds;
    encapsulation type;
    fragment-threshold bytes;
    interleave-fragments;
    minimum-links number;
    mrru bytes;
    multicast-dlci dlci-identifier;
    short-sequence;
    family family {
        address address {
            destination address;
        }
        bundle (ml-fpc/pic/port | ls-fpc/pic/port);
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 16,384

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Link and Multilink Services Interfaces User Guide for Routing Devices
- Junos OS Network Interfaces Library for Routing Devices

unit (Tunnel)

IN THIS SECTION

- Syntax | 30405
- Hierarchy Level | 30405
- Description | 30405
- Options | 30405
- Required Privilege Level | 30406
- Release Information | 30406

Syntax

```
unit logical-unit-number {
    peer-unit unit-number;
    reassemble-packets;
    tunnel {
        allow-fragmentation;
        backup-destination address;
        destination destination-address;
        do-not-fragment;
        key number;
        routing-instance {
            destination routing-instance-name;
        }
        source source-address;
        ttl number;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit logical-systems logical-system-name interfaces interface-name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 16,384

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Network Interfaces Library for Routing Devices

unit (TDF Interface)

IN THIS SECTION

- Syntax | 30406
- Hierarchy Level | 30407
- Description | 30407
- Options | 30407
- Required Privilege Level | 30407
- Release Information | 30407

Syntax

```
unit interface-unit-number {
    family family-name;
}
```

Hierarchy Level

```
[edit interfaces mif]
```

Description

Configure the logical interface on the TDF interface. You must configure a logical interface to be able to use the TDF interface.

Options

interface-unit-number

Number of the logical unit.

- **Range:** 0 through 16,384

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a TDF Logical Interface

unit (Voice Services)

IN THIS SECTION

- Syntax | 30408
- Hierarchy Level | 30409
- Description | 30409
- Options | 30409
- Required Privilege Level | 30409
- Release Information | 30409

Syntax

```
unit logical-unit-number {
    compression {
        rtp {
            f-max-period number;
            maximum-contexts number <force>;
            port {
                minimum port-number;
                maximum port-number;
            }
            queues [ queue-numbers ];
        }
    }
    compression-device interface-name;
    encapsulation type;
    family family {
        address address {
            ...
        }
        bundle (lsq-fpc/pic/port | ...);
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name ]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 16,384

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> Junos OS Network Interfaces Library for Routing Devices
>
> *Configuring Services Interfaces for Voice Services*

unknown-message (Security H323 ALG)

IN THIS SECTION

- Syntax | 30410
- Hierarchy Level | 30410
- Description | 30410
- Options | 30411
- Required Privilege Level | 30411
- Release Information | 30411

Syntax

```
unknown-message {
    permit-nat-applied;
    permit-routed;
}
```

Hierarchy Level

```
[edit logical-systems name security alg h323 application-screen],
[edit logical-systems name tenants name security alg h323 application-screen],
[edit security alg h323 application-screen],
[edit services alg h323 application-screen],
[edit tenants name security alg h323 application-screen]
```

Description

Specify how the SRX Series Firewall handles the unidentified H.323 messages. By default, SRX Series Firewall drops unknown messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement is useful to resolve interoperability issues with disparate vendor equipment. By permitting unknown H.323 (unsupported) messages, you can get your network operational. Later, you can analyze your VoIP traffic to determine why the device dropped the messages.

This statement applies only to received packets identified as supported VoIP packets. SRX Series Firewall always drops unidentified packets and passes the identified packets without processing.

Options

- `permit-nat-applied`—Permits unknown messages to pass if the session is in NAT mode.
- `permit-routed`—Permit unknown messages on routed packets. Sessions in Transparent mode are treated as Route mode.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

h323 | 7961

unknown-message (Security MGCP ALG)

IN THIS SECTION

- Syntax | 30412
- Hierarchy Level | 30412
- Description | 30412
- Options | 30412
- Required Privilege Level | 30412
- Release Information | 30413

Syntax

```
unknown-message {
    permit-nat-applied;
    permit-routed;
}
```

Hierarchy Level

```
[edit logical-systems name security alg mgcp application-screen],
[edit logical-systems name tenants name security alg mgcp application-screen],
[edit security alg mgcp application-screen],
[edit services alg mgcp application-screen],
[edit tenants name security alg mgcp application-screen]
```

Description

Specify how SRX Series Firewall handles unidentified Media Gateway Control Protocol (MGCP) messages. The default is to drop unknown (unsupported) messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement is useful to resolve interoperability issues with disparate vendor equipment. You can permit unknown MGCP (unsupported) messages to get your network operational. Later, you can analyze your VoIP traffic to determine why some messages were dropped.

This statement applies only to received packets identified as supported VoIP packets. Unidentified packets are always dropped. If a packet is identified as a supported protocol, SRX Series Firewall forwards the message without processing.

Options

- `permit-nat-applied`—Permits unknown messages to pass if the session is in NAT mode.
- `permit-routed`—Permit unknown messages on routed packets. Sessions in Transparent mode are treated as Route mode.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

mgcp | 10693

unknown-message (Security SCCP ALG)

IN THIS SECTION

- Syntax | 30413
- Hierarchy Level | 30413
- Description | 30414
- Options | 30414
- Required Privilege Level | 30414
- Release Information | 30414

Syntax

```
unknown-message {
    permit-nat-applied;
    permit-routed;
}
```

Hierarchy Level

```
[edit logical-systems name security alg sccp application-screen],
[edit logical-systems name tenants name security alg sccp application-screen],
[edit security alg sccp application-screen],
```

```
[edit services alg sccp application-screen],
[edit tenants name security alg sccp application-screen]
```

Description

Specify how SRX Series Firewall handles unidentified Skinny Client Control Protocol (SCCP) messages. The default is to drop unknown (unsupported) messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement is useful to resolve interoperability issues with disparate vendor equipment. You can permit unknown MGCP (unsupported) messages to get your network operational. Later, you can analyze your VoIP traffic to determine why some messages were dropped.

This statement applies only to received packets identified as supported VoIP packets. Unidentified packets are always dropped. If a packet is identified as a supported protocol, SRX Series Firewall forwards the message without processing.

Options

- `permit-nat-applied`—Permits unknown messages to pass if the session is in NAT mode.
- `permit-routed`—Permit unknown messages on routed packets. Sessions in Transparent mode are treated as Route mode.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sccp | 16341

unknown-message (Security SIP ALG)

IN THIS SECTION

- Syntax | 30415
- Hierarchy Level | 30415
- Description | 30415
- Options | 30416
- Required Privilege Level | 30416
- Release Information | 30416

Syntax

```
unknown-message {
    permit-nat-applied;
    permit-routed;
}
```

Hierarchy Level

```
[edit logical-systems name security alg sip (Security)  application-screen (Security SIP)],
[edit logical-systems name tenants name security alg sip (Security)  application-screen (Security
SIP)],
[edit security alg sip (Security)  application-screen (Security SIP)],
[edit services alg sip (Security)  application-screen (Security SIP)],
[edit tenants name security alg sip (Security)  application-screen (Security SIP)]
```

Description

Specify how SRX Series Firewall handles unidentified Session Initiation Protocol (SIP) messages. The default is to drop unknown (unsupported) messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement is useful for resolving interoperability issues with disparate vendor equipment. You can permit unknown

SIP (unsupported) messages to get your network operational. Later, you can analyze your VoIP traffic to determine why some messages were dropped.

This statement applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol, the message is forwarded without processing.

Options

- `permit-nat-applied`—Permits unknown messages to pass if the session is in NAT mode.
- `permit-routed`—Permit unknown messages on routed packets. Sessions in Transparent mode are treated as Route mode.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

sip (Security) | 27977

unknown-protocol (IDS Screen Next Gen Services)

IN THIS SECTION

- Syntax | 30417
- Hierarchy Level | 30417
- Description | 30417
- Required Privilege Level | 30417
- Release Information | 30417

Syntax

```
unknown-protocol;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Description

Identify and drop IP frames with protocol numbers greater than 137 for IPv4 and 139 for IPv6, which protects against IP unknown protocol attacks.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

unknown-unicast-forwarding

IN THIS SECTION

- Syntax | 30418
- Hierarchy Level | 30418

- Description | 30418
- Default | 30418
- Required Privilege Level | 30419
- Release Information | 30419

Syntax

```
unknown-unicast-forwarding {
    vlan (all | vlan-name){
        interface interface-name;
    }
}
```

Hierarchy Level

```
[edit ethernet-switching-options],
[edit switch-options]
```

Description

Configure the switch to forward all unknown unicast packets in a VLAN or on all VLANs to a particular interface.

> **NOTE**: Before you can configure unknown unicast forwarding within a VLAN, you must first configure that VLAN.

The remaining statements are explained separately. See CLI Explorer.

Default

Unknown unicast packets are flooded to all interfaces that belong to the same VLAN.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show ethernet-switching table | 19700
show vlans | 27738

unnumbered-address (Demux)

IN THIS SECTION

- Syntax | 30419
- Hierarchy Level | 30420
- Description | 30420
- Options | 30420
- Required Privilege Level | 30420
- Release Information | 30420

Syntax

```
unnumbered-address interface-name <preferred-source-address address>;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family
inet]
```

Description

For IP demultiplexing interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered interface enables IP processing on the interface without assigning an explicit IP address to the interface.

Options

interface-name—Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface.

The `preferred-source-address` statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

`preferred-source-address` option introduced in Junos OS Release 9.0.

IP demultiplexing interfaces supported in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring an Unnumbered Interface

address | **1846**

unnumbered-address (Dynamic PPPoE)

IN THIS SECTION

- Syntax | 30421
- Hierarchy Level | 30421
- Description | 30421
- Options | 30421
- Required Privilege Level | 30421
- Release Information | 30422

Syntax

```
unnumbered-address interface-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family inet]
```

Description

For dynamic PPPoE interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface.

Options

interface-name—Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring a PPPoE Dynamic Profile

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

unnumbered-address (Dynamic Profiles)

IN THIS SECTION

- Syntax | 30422
- Hierarchy Level | 30422
- Description | 30423
- Options | 30423
- Required Privilege Level | 30424
- Release Information | 30424

Syntax

```
unnumbered-address interface-name <preferred-source-address address>;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family],
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family]
```

Description

For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface. To configure unnumbered address dynamically, include the `$junos-loopback-interface-address` predefined variable.

You can configure unnumbered address support on Ethernet interfaces for IPv4 and IPv6 address families.

Options

interface-name

Name of the interface from which the local address is derived. The specified interface must have a logical unit number, a configured IP address, and must not be an unnumbered interface. This value can be a specific interface name or the `$junos-loopback-interface` predefined variable.

When defining the `unnumbered-address` statement using a static interface, keep the following in mind:

- If you choose to include the `routing-instance` statement at the `[edit dynamic-profiles]` hierarchy level, that statement must be configured with a dynamic value by using the `$junos-routing-instance` predefined variable. In addition, whatever static unnumbered interface you specify must belong to that routing instance; otherwise, the profile instantiation fails.
- If you choose to not include the `routing-instance` statement at the `[edit dynamic-profiles]` hierarchy level, the unnumbered-address statement uses the default routing instance. The use of the default routing instance requires that the unnumbered interface be configured statically and that it reside in the default routing instance.

> **NOTE**: When you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the `$junos-routing-instance` predefined variable, you must not configure a preferred source address, whether with the `$junos-preferred-source-address` predefined variable, the `$junos-preferred-source-ipv6-address` predefined variable, or the `preferred-source-address` statement. Configuring the preferred source address in this circumstance causes a commit failure.

When defining the `unnumbered-address` statement using the `$junos-loopback-interface` predefined variable, keep the following in mind:

- To use the `$junos-loopback-interface` predefined variable, the dynamic profile must also contain the `routing-instance` statement configured with the `$junos-routing-instance` predefined variable at the [edit dynamic-profiles] hierarchy level.
- The applied loopback interface is based on the dynamically obtained routing instance of the subscriber.

address (Optional) Secondary IP address of the donor interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network. This value can be a static IP address, the `$junos-preferred-source-address` predefined variable for the inet family, or the `$junos-preferred-source-ipv6-address` predefined variable for the inet6 family.

When defining the `preferred-source-address` value using a static IP address, keep the following in mind:

- The unnumbered interface must be statically configured.
- The IP address specified as the `preferred-source-address` must be configured in the specified unnumbered interface.

When defining the `preferred-source-address` value using the `$junos-preferred-source-address` or the `$junos-preferred-source-ipv6-address` predefined variables, keep the following in mind:

- You must configure the `unnumbered-address` statement using the `$junos-loopback-interface` predefined variable.
- You must configure the `routing-instance` statement using the `$junos-routing-instance` predefined variable at the `[edit dynamic-profiles]` hierarchy level.
- The preferred source address chosen is based on the dynamically applied loopback address which is in turn derived from the dynamically obtained routing instance of the subscriber. The configured loopback address with the closest network match to the user IP address is selected as the preferred source address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support for the `$junos-preferred-source-address` and `$junos-preferred-source-ipv6-address` predefined variables introduced in Junos OS Release 9.6.

Support for the `$junos-loopback-interface` predefined variable introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Dynamic Profiles Overview

unnumbered-address (Ethernet)

IN THIS SECTION

- Syntax | 30425
- Hierarchy Level | 30425
- Description | 30426
- Options | 30426
- Required Privilege Level | 30426
- Release Information | 30426

Syntax

```
unnumbered-address interface-name <preferred-source-address address>;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family
family]
```

Description

For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered Ethernet interface enables IP processing on the interface without assigning an explicit IP address to the interface.

Options

interface-name—Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface.

The `preferred-source-address` statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

`preferred-source-address` option introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- Configuring an Unnumbered Interfaceaddress
- address
- Overview for Junos OS

unnumbered-address (PPP)

IN THIS SECTION

- Syntax | 30427
- Hierarchy Level | 30427

- Description | 30427
- Options | 30427
- Required Privilege Level | 30427
- Release Information | 30428

Syntax

```
unnumbered-address interface-name destination address destination-profile profile-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet]
```

Description

For interfaces with PPP encapsulation, enable the local address to be derived from the specified interface.

Options

interface-name—Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring IPCP Options for Interfaces with PPP Encapsulation

unprotect

IN THIS SECTION

- Syntax | 30428
- Description | 30428
- Options | 30428
- Required Privilege Level | 30429
- Release Information | 30429

Syntax

```
unprotect (hierarchy | statement | identifier)
```

Description

Unprotect a protected hierarchy, configuration statement, or an identifier, so that it can be modified or deleted.

Options

hierarchy (Optional) Unprotect a specific protected hierarchy.

statement (Optional) Unprotect a specific protected statement.

identifier (Optional) Unprotect a specific protected identifier.

Required Privilege Level

configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

- protect | 13404
- top | 29435
- up | 30436
- *Displaying the Current Configuration*
- *Example: Protecting the Configuration from Modification or Deletion*

unreachable-vtep-aging-timer

IN THIS SECTION

- Syntax | 30430
- Hierarchy Level | 30430
- Description | 30430
- Required Privilege Level | 30430
- Release Information | 30430

Syntax

```
unreachable-vtep-aging-timer [300-1800];
```

Hierarchy Level

```
[edit vlans vlan-name vxlan]
```

Description

Configure the system to age out the address for the remote virtual tunnel endpoint (VTEP) if all the MAC addresses learned from that VTEP age out. The address for the remote VTEP expires the configured number of seconds after the last learned media access control (MAC) address expires.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Understanding VXLANs

Manually Configuring VXLANs on QFX Series and EX4600 Switches

Examples: Manually Configuring VXLANs on QFX Series and EX4600 Switches

untagged

IN THIS SECTION

- Syntax | 30431
- Hierarchy Level | 30431
- Description | 30431
- Required Privilege Level | 30431
- Release Information | 30431

Syntax

```
untagged;
```

Hierarchy Level

```
[edit interfaces ps0]
```

Description

Specify that the router supports untagged traffic on pseudowire subscriber interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

MPLS Pseudowire Subscriber Logical Interfaces

untrusted

IN THIS SECTION

- Syntax | 30432
- Hierarchy Level | 30432
- Description | 30432
- Required Privilege Level | 30432
- Release Information | 30433

Syntax

```
untrusted;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security group group-name overrides]
```

Description

Allow DHCP responses from the specified interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Enabling a Trusted DHCP Server (MX Series Routers)

Understanding and Using Trusted DHCP Servers

untrusted (VLAN)

IN THIS SECTION

- Syntax | 30433
- Hierarchy Level | 30433
- Description | 30433
- Required Privilege Level | 30434
- Release Information | 30434

Syntax

```
untrusted;
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name overrides]
```

Description

Configure a trunk interface as untrusted for DHCP security. Trunk interfaces are trusted by default and all packets are allowed. You can override this default behavior and set a trunk interface as untrusted in

order to support DHCP security features on the interface. DHCP snooping, DHCPv6 snooping, dynamic ARP inspection (DAI), and IPv6 neighbor discovery inspection are supported on trunk ports in untrusted mode.

> **NOTE**: On EX Series switches other than the EX9200, IP source guard and IPv6 source guard are not supported on untrusted trunk ports.

Configuring a trunk port as untrusted is useful in deployments where multiple DHCP clients are aggregated onto one interface on the access device. In this scenario, the interface is configured as a trunk interface with one or more VLANs. A DHCP client attached to a trunk interface might start acting as a DHCP server. Trusted ports allow DHCP servers to provide IP addresses and other information to requesting devices, which makes the network vulnerable to a rogue DHCP server attack.

An unauthorized DHCP server might also assign itself as the default gateway device for the network. An attacker can then sniff the network traffic and perpetrate a man-in-the-middle attack—that is, it misdirects traffic intended for a legitimate network device to a device of its choice. To mitigate this problem, you can configure the interface to which the unauthorized server is connected as untrusted, which blocks all ingress DHCP server messages from that interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Understanding and Using Trusted DHCP Servers

Configuring a Trunk Interface as Untrusted for DHCP Security (CLI Procedure)

Unused

IN THIS SECTION

- Syntax | 30435
- Hierarchy Level | 30435
- Description | 30435
- Options | 30436
- Required Privilege Level | 30436
- Release Information | 30436

Syntax

```
unused;
```

Hierarchy Level

```
[edit chassis fpc slot-number  pic pic-number port port-number]
[edit interface interface-name]
```

Description

(PTX10003-80C and PTX10003-160C routers) While setting port 0 to 400-Gigabit Ethernet mode (using QSFP56-DD-400GBASE-LR8 optics), the total bandwidth (speed x `number-of-subports`) of port 1 has to be less than 100G and port 2 has to configured as 'unused'. When using port 4 as 400G, port 3 has to be configured with total bandwidth and port 2 has to be configured 'unused'. Similarly, with port 5, 9 using 400G, port 6, 8 respectively has to be configured for less than 100G and port 7 should be configured as 'unused'. That is, when a port is configured in 400-Gigabit ethernet mode, you cannot configure speed of the adjacent port to be more then 100-Gbps, and the middle port (2 between 0~4 or 7 between 5~9) must be set to unused. To set a particular port to unused port, use `unused` CLI command.

You cannot use port 1,3,6 and/or 7 of the same logical PIC.

(JNP10K-LC1201 on PTX10008 routers) To disable the specified physical port in a PIC. If you disable a port, no interfaces (channelized or non-channelized) are created for the port, irrespective of the port profile configuration configured on the port.

The `edit chassis` hierarchy is available in 20.2R1 Release on PTX10001-36MR. From 20.3R1 onwards you must use only the `edit interface` *`interface-name`* hierarchy.

Options

None

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 19.3R1.

Hierarchy introduced in Junos OS Evolved Release 20.1R1 for JNP10K-LC1201 line card on PTX10008 routers.

RELATED DOCUMENTATION

- speed (Ethernet) | **28349**
- number-of-sub-ports (interface) | **11740**
- PTX10003 Port Panel
- Port Speed on PTX10003

up

IN THIS SECTION

- Syntax | **30437**
- Description | **30437**

- Options | 30437
- Required Privilege Level | 30437
- Release Information | 30437

Syntax

```
up <number> <configuration-command>
```

Description

Move up one level in the statement hierarchy.

Options

none	Move up one level in the configuration hierarchy.
configuration-command	(Optional) Issue configuration mode commands from a location higher in the hierarchy.
number	(Optional) Move up the specified number of levels in the configuration hierarchy.

Required Privilege Level

configure—To enter configuration mode.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Displaying the Current Configuration

exit | 6638

top | 29435

update

IN THIS SECTION

- Syntax | 30438
- Description | 30438
- Options | 30438
- Required Privilege Level | 30438
- Release Information | 30439

Syntax

```
update
```

Description

Update private candidate configuration with a copy of the most recently committed configuration, including your private changes.

NOTE: The `update` command is available only when you are in configure private mode.

Options

This command has no options.

Required Privilege Level

configure—To enter configuration mode.

Release Information

Command introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Updating the configure private Configuration

update-interval

IN THIS SECTION

- Syntax | 30439
- Hierarchy Level | 30439
- Description | 30440
- Default | 30440
- Options | 30440
- Required Privilege Level | 30440
- Release Information | 30440

Syntax

```
update-interval minutes;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Enable interim accounting updates and configure the amount of time that the router or switch waits before sending a new accounting update.

Interim accounting updates are included in the exchange of messages between the client and the accounting server. In RADIUS accounting, the client is the network access server (NAS), which can be the router or switch. The NAS sends Accounting-Request messages to the server, which acknowledges receipt of the requests with Accounting-Response messages. Interim accounting updates are sent in Accounting-Request packets with the Acct-Status-Type attribute set to Interim-Update.

When a user is authenticated, the authentication server issues an Access-Accept message in response to a successful Access-Request message. The interval between interim updates can be configured directly on the server using the Acct-Interim-Interval attribute of the Access-Accept message. However, if the update interval is configured on the NAS using `update-interval`, then the system prefers the attributes returned by RADIUS and overrides the locally configured values.

> **NOTE**: All information in an interim update message is cumulative from the beginning of the session, not from the last interim update message.

Default

No interim updates are sent from the client to the accounting server.

Options

minutes—Amount of time between updates, in minutes. All values are rounded to the next higher multiple of 10. For example, the values 811 through 819 are all accepted by the CLI, but are all rounded up to 820.

- **Range:** 10 through 1440 minutes

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

ANCP Agent and AAA

update-interval (Protocols RIP)

IN THIS SECTION

- Syntax | 30441
- Hierarchy Level | 30441
- Description | 30442
- Options | 30442
- Required Privilege Level | 30442
- Release Information | 30442

Syntax

```
update-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rip],
[edit logical-systems logical-system-name protocols rip group group-name],
[edit logical-systems logical-system-name protocols rip group group-name neighbor neighbor-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols rip
group group-name neighbor neighbor-name],
[edit protocols rip],
[edit protocols rip group group-name],
[edit protocols rip  group group-name neighbor neighbor-name],
[edit routing-instances routing-instance-name protocols rip],
```

```
[edit routing-instances routing-instance-name protocols rip group group-name],
[edit routing-instances routing-instance-name protocols rip  group group-name neighbor neighbor-
name]
```

Description

Configure the interval at which routes learned by RIP are sent to neighbors. This timer controls the interval between routing updates. This timer is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can happen if all routing devices update their neighbors simultaneously.

Options

seconds—Estimated time to wait before making updates to the routing table.

- **Range:** 10 through 60 seconds
- **Default:** 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

Example: Configuring RIP Timers

update-interval (Protocols RIPng)

IN THIS SECTION

- Syntax | 30443
- Hierarchy Level | 30443
- Description | 30443
- Options | 30443
- Required Privilege Level | 30444
- Release Information | 30444

Syntax

```
update-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ripng],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
ripng],
[edit protocols ripng],
[edit routing-instances routing-instance-name protocols ripng]
```

Description

Configure the interval at which routes learned by RIPng are sent to neighbors.

Options

seconds—Estimated time to wait before making updates to the routing table.

- **Range:** 10 through 60 seconds
- **Default:** 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

Support for routing instances introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring RIP Timers

update-interval (Service Accounting)

IN THIS SECTION

- Syntax | 30444
- Hierarchy Level | 30445
- Description | 30445
- Default | 30445
- Options | 30445
- Required Privilege Level | 30445
- Release Information | 30445

Syntax

```
update-interval minutes;
```

Hierarchy Level

```
[edit access profile profile-name service accounting]
```

Description

Enable service interim accounting updates and configure the amount of time that the router waits before sending a new service accounting update.

> **NOTE**: An update interval value configured in the RADIUS attribute, Service-Interim-Acct-Interval (VSA 26-140), takes precedence over the value configured with this statement. In turn, the value configured with this statement takes precedence over the "interval" on page 9211 value configured at the `[edit accounting-options flat-file-profile profile-name]` hierarchy level.

Default

No updates

Options

minutes—Amount of time between updates, in minutes. All values are rounded up to the next higher multiple of 10. For example, the values 811 through 819 are all accepted by the CLI, but are all rounded up to 820.

- **Range:** 10 through 1440 minutes

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R1.

RELATED DOCUMENTATION

RADIUS Accounting for Subscriber Access

RADIUS Accounting for Subscriber Access

RADIUS Accounting for Subscriber Access

Configuring Service Accounting in Local Flat Files

update-rate-limit

IN THIS SECTION

- Syntax | 30446
- Hierarchy Level | 30446
- Description | 30446
- Options | 30447
- Required Privilege Level | 30447
- Release Information | 30447

Syntax

```
update-rate-limit updates-per-minute;
```

Hierarchy Level

```
[edit protocols pcep]
```

Description

Specify the number of updates per minute that the Path Computation Client (PCC) can receive at maximum. Updates above this limit are ignored by the PCC.

Options

updates-per-minute Number of updates per minute that the PCC can receive at maximum.

- **Range:** 1 through 16384
- **Default:** 0 (disabled or no limit)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

pcep | 12320

update-response-timeout (PCRF Partition)

IN THIS SECTION

- Syntax | 30448
- Hierarchy Level | 30448
- Description | 30448
- Options | 30448
- Required Privilege Level | 30448

Release Information | 30449

Syntax

```
update-response-timeout seconds;
```

Hierarchy Level

```
[edit access pcrf partition partition-name]
```

Description

Configure the amount of time in seconds before a Policy and Charging Rule Function (PCRF) partition stops attempting to send an updated rule report response using a CCR-GX-U message. Use this statement if you added or deleted rules in either the subscriber login using a CCA-GX-I message, or the reauthorization request using a RAR-GX-U message, and configured rule reporting,

The last rule report may be preempted by the subscriber logging out too soon. Use this statement to control the amount of time to allow the system to send the last rule report.

Options

seconds Number of seconds to wait before a PCRF partition stops attempting to send an updated rule report response using a CCR-GX-U message.

- **Default:** 300
- **Range:** 0 through 86,400 seconds (24 hours)

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

update-router-advertisement

IN THIS SECTION

- Syntax | 30449
- Hierarchy Level | 30449
- Description | 30450
- Options | 30450
- Required Privilege Level | 30450
- Release Information | 30450

Syntax

```
update-router-advertisement (interface interface-name);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcpv6-client]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
```

```
family inet6 dhcpv6-client]
[edit tenants tenant-name interfaces interface-name unit logical-unit-number family inet6 dhcpv6-
client]
```

Description

Specify the interface used to delegate prefixes.

Options

interface *interface-name*	Interface on which to delegate prefixes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

The `logical-systems` and `tenant` options are introduced in Junos OS Release 18.4R1.

update-server

IN THIS SECTION

- Syntax | 30451
- Hierarchy Level | 30451
- Description | 30451
- Required Privilege Level | 30451
- Release Information | 30451

Syntax

```
update-server;
```

Hierarchy Level

```
[edit interfaces interface-name unit  logical-unit-number  family inet dhcp]
```

Description

Propagate TCP/IP settings learned from an external DHCP server to the DHCP server running on the switch, router, or device.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring a DHCP Client

Example: Configuring the Device as a DHCP Client

interfaces (EX Series switches) | **9135**

unit

family

update-server (dhcp-client)

IN THIS SECTION

- Syntax | 30452
- Hierarchy Level | 30452
- Description | 30452
- Required Privilege Level | 30452
- Release Information | 30452

Syntax

```
update-server;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcp-client]
```

Description

Propagate DHCP options to a local DHCP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

Starting with Junos OS Release 17.3R1, on all SRX Series Firewalls and vSRX Virtual Firewall instances, the CLI option `dhcp-clinet` at `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy is

changed to dhcp to keep consistent with other Junos platforms. There is no change in the functionality with the changed CLI option dhcp.

update-server (dhcpv6-client)

IN THIS SECTION

- Syntax | 30453
- Hierarchy Level | 30453
- Description | 30453
- Required Privilege Level | 30453
- Release Information | 30454

Syntax

```
update-server;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcpv6-client]
```

Description

Propagate TCP/IP settings to the DHCPv6 server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

update-threading

IN THIS SECTION

- Syntax | 30454
- Hierarchy Level | 30454
- Description | 30454
- Options | 30455
- Required Privilege Level | 30455
- Release Information | 30456

Syntax

```
update-threading {
    number-of-threads;
     group-split-size;
        }
    }
```

Hierarchy Level

```
[edit system processes routing bgp]
```

Description

Enables BGP update thread. If you configure update-threading on a routing engine, RPD creates update threads.

NOTE:

- Sharding requires Update/IO thread. If `rib-sharding` is configured and `update-threading` is not configured, the commit check fails.
- RPD would restart automatically when rib-sharding or update-threading configuration is changed.

Options

`number of threads`—the number of update threads created. If you configure update-threading on a routing engine, RPD creates update threads. By default, the number of update threads created is same as the number of CPU cores on the routing engine. Optionally, you can specify the number-of-threads you want to create.

`group-split-size`—the large groups are divided into fragments based on the configuration of the `group-split-size` (The size here is the number of peers). The feature is enabled only after a `group-split-size` number is provided, until then the feature remains disabled (since the split size would be 0). The objective of the feature is to share the load with different threads since multiple threads are alway available to process. This `group-split-size` statement defines the size that we are using to split the peer group. The preferred size value for this statement is 0 to 2000.

You can adjust the `group-split-size` value according to the requirement. There are two general tips for finding a reasonable group-split-size:

- Do not make it too large so that the system is not evenly distributed. For example, when we have one group of 1000 peers and 10 available threads, setting split-size to 100 is better than 200 because it allows the system to utilize all 10 threads for 100 peers per thread compared to only 5 will be used for 200 peers per thread.
- Do not make it too small so that it takes the unnecessary overhead of splitting groups. For example, if we have 10 groups (each group contains 103 peers) and 10 available threads, setting split-size to 103 will be better than 100 because splitting a group of 103 to 100 and 3 will give us unnecessary overhead.
- **Range:** 1-128 threads

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

`group-split-size` statement introduced in Junos OS Release 21.4R1.

Support for NSR with sharding introduced in Junos OS Release 22.2.

update-threshold

IN THIS SECTION

- Syntax | 30456
- Hierarchy Level | 30456
- Description | 30457
- Options | 30457
- Required Privilege Level | 30457
- Release Information | 30457

Syntax

```
update-threshold {
    threshold-percent;
     adaptive
    threshold-value threshold-value;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit protocols rsvp interface interface-name]
```

Description

Adjust the threshold at which a change in bandwidth triggers an interior gateway protocol (IGP) update.

Options

threshold-percent

Specify the percentage change in reserved bandwidth to trigger IGP update.

- **Default:** 10 percent
- **Range:** 0.001 through 20 percent

threshold-value ***threshold-value***

Specify the change in reserved bandwidth to trigger IGP update. (is capped at 20% of link bandwidth).

If the threshold-value is configured to greater than 20% of bandwidth on that link, the threshold-value is capped at 20% of bandwidth.

For instance, if bandwidth on an interface is 1Gbps, and the `threshold-value` is configured greater than 200Mbps, the `threshold-value` is capped at 200Mbps. The *threshold-percent* is displayed as 20.000% and the `threshold-value` as 200Mbps.

> **NOTE**: The two options, *threshold-percent* and `threshold-value`, are mutually exclusive. You can configure only one option at a given point in time to generate an IGP update for lower bandwidth reservations. As a result, when one option is configured, the other option is calculated and displayed on the CLI.
> For instance, on a link of 1Gbps, if the *threshold-percent* is configured to 5%, the `threshold-value` is calculated and displayed as 50Mbps. Similarly, if the `threshold-value` is configured to 50m, then the *threshold-percent* is calculated and displayed as 5%.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`threshold-value` option introduced in Junos OS Release 19.4R1 on all platforms.

RELATED DOCUMENTATION

Configuring the RSVP Update Threshold on an Interface

update-threshold-max-reservable

IN THIS SECTION

- Syntax | 30458
- Hierarchy Level | 30458
- Description | 30458
- Default | 30459
- Options | 30459
- Required Privilege Level | 30459
- Release Information | 30459

Syntax

```
update-threshold-max-reservable update-threshold-max-reservable;
```

Hierarchy Level

```
[edit logical-systems name protocols rsvp interface],
[edit logical-systems name routing-instances name protocols rsvp interface],
[edit protocols rsvp interface],
[edit routing-instances name protocols rsvp interface]
```

Description

Configure a threshold in bits per second. When SPRING bandwidth utilization exceeds this threshold, Resource Reservation Protocol (RSVP) reports this event to IGP and preempts the labeled switched paths (LSPs). If this threshold is not exceeded, RSVP does not make any changes to the bandwidth allocation.

However, if the `local-bw-override-threshold` statement is configured along with the `update-threshold-max-reservable` statement, RSVP modifies local unreserved bandwidths even if the threshold is not exceeded. This change does not report the new maximum reservable bandwidth and local unreserved bandwidth to IGP until the threshold configuration crosses.

Note that the changes are not applied by modifying the `subscription` *`percentage`* at the `[edit protocols rsvp interface` *`interface-name`*`]` hierarchy level.

Default

The default value is `0`.

Options

update-threshold-max-reservable

Specify a threshold value for maximum SPRING bandwidth utilization in bits per second.

- **Range:** 0 through 100000

percent *percent*

Specify a bandwidth utilization change threshold in percentage beyond which RSVP triggers IGP updates. When the change in bandwidth utilization from previously advertised value is more than the percentage configured, RSVP informs IGP about the max reservable bandwidth change.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

`percent` option introduced in Junos OS Release 17.4R1 for MX Series and PTX Series.

RELATED DOCUMENTATION

auto-bandwidth | **2682**

traffic-statistics (Protocols IS-IS) | **30048**

local-bw-override-threshold | **9925**

upgrade-group

IN THIS SECTION

- Syntax | 30460
- Hierarchy Level | 30460
- Description | 30460
- Options | 30461
- Required Privilege Level | 30461
- Release Information | 30461

Syntax

```
upgrade-group group-name {
    fpcs (slot-number | [list-of-slot-numbers]);
    member member-id {
        fpcs (slot-number | [list-of-slot-numbers]);
    }
}
```

Hierarchy Level

```
[edit chassis nssu]
```

Description

Assign a name to a line-card upgrade group being created for nonstop software upgrade (NSSU).

To reduce the time an NSSU takes, you can configure line-card upgrade groups on an EX6200 or EX8200 switch with redundant Routing Engines; EX8200 Virtual Chassis; EX4650 Virtual Chassis; QFX3500, QFX3600, and QFX5100 Virtual Chassis; and QFX5100 Virtual Chassis Fabric (VCF).

NSSU upgrades the groups in the order that they appear in the configuration (in other words, in the order you configure them). If you do not define any line-card upgrade groups, NSSU upgrades line cards or members of a Virtual Chassis or VCF one at a time in ascending order by slot or member number.

Options

group-name	Name of the upgrade group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade

Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches

Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade

upgrade-groups (Junos Fusion)

IN THIS SECTION

- Syntax | 30462
- Hierarchy Level | 30462
- Description | 30462
- Default | 30463
- Options | 30463
- Required Privilege Level | 30463
- Release Information | 30463

Syntax

```
upgrade-groups upgrade-group-name {
    satellite [slot-id | range | all];
}
```

Hierarchy Level

```
[edit chassis satellite-management]
```

Description

Create and name a satellite software upgrade group for a Junos Fusion, or specify an existing satellite software upgrade group to configure.

A satellite software upgrade group is a group of satellite devices that are designated to upgrade to the same satellite software version using the same satellite software package. One Junos Fusion can contain multiple software upgrade groups, and multiple software upgrade groups should be configured in most Junos Fusions to avoid network downtimes during satellite software installations.

The two most common methods of installing satellite software in a Junos Fusion—autoconverting a device into a satellite device when it is cabled into an aggregation device and manually converting a device that is cabled into an aggregation device into a satellite device—require a configured satellite software upgrade group.

Software upgrade groups are configured and managed from the aggregation device.

To associate a satellite software package with a satellite software upgrade group, use the **request system software add *package-name* upgrade-group *upgrade-group-name*** command.

This statement is entered on an aggregation device in a Junos Fusion. Software upgrade groups are configured and managed from the aggregation device.

The software upgrade group configurations must match exactly—including the same *package-name* and *upgrade-group-name*—in every Junos Fusion with dual aggregation devices to avoid satellite device downtime.

All satellite devices in a satellite device cluster are associated with a single satellite software upgrade group, which is automatically created when a satellite device cluster becomes part of a Junos Fusion. The satellite software upgrade group is named after the satellite device cluster name, and ensures that all satellite devices in the cluster run the same version of satellite software. See Understanding Software

in a Junos Fusion Enterprise for additional information on software management for a satellite device cluster.

Default

No satellite software upgrade groups are present, by default.

A satellite software upgrade group with the name of the satellite device cluster is created automatically when a satellite device cluster is created.

Options

upgrade-group-name Specifies the user-defined name for the satellite software upgrade group.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

Statement introdcued in Junos OS Release 16.1R1 for Junos Fusion Enterprise.

RELATED DOCUMENTATION

Configuring or Expanding a Junos Fusion Enterprise

Configuring Junos Fusion Provider Edge

upgrade-mode

IN THIS SECTION

- Syntax | 30464
- Hierarchy Level | 30464
- Description | 30464
- Options | 30464
- Required Privilege Level | 30465
- Release Information | 30465

Syntax

```
upgrade-mode (3d-fabric | default | t4000);
```

Hierarchy Level

```
[edit chassis fabric],
[edit chassis member name fabric]
```

Description

Configures upgrade mode for SIBs and forces them to operate in the same mode until the upgrade is complete. Upgrade mode is used so that two different types of SIBs can be installed in an operational router or routing matrix. After you upgrade the SIBs delete the `upgrade-mode` statement from the configuration. See the hardware installation guide for your router for more information about upgrading SIBs in an operational router or routing matrix.

Options

3d-fabric Enables the TX Matrix Plus router to upgrade to a TX Matrix Plus router with 3D SIBs. On the SFC, enables setting proper support for mixed SIBs (TXP-F13 SIB and TXP-F13-3D SIB). On the T640 or T1600 or T4000 routers in a routing matrix enables support for mixed SIBs

(TXP-T1600 SIB and TXP-3D-LCC SIBs on the T1600 router and SIB-I-T4000 and TXP-3D-LCC SIBs on the T4000 router).

default Enables support for mixed SIBs when upgrading SIBs in the PTX3000 and PTX5000 routers.

t4000 Enables support for mixed SIBs when upgrading to SIB-I-T4000 SIBs in a T640 or T1600.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

- TX Matrix Router and T640 Router Configuration Overview
- Configuring Junos OS to Upgrade a T4000 Router to LCC 0 of a TX Matrix Plus Router with 3D SIBs
- Configuring Junos OS to Upgrade the T1600 Router Chassis to LCC 0 of a TX Matrix Plus Router with 3D SIBs
- Upgrading an Operational Standalone Router and Integrating It into a TX Matrix Plus Routing Matrix with 3D SIBs
- Upgrading to SIB3-SFF-PTX SIBs in an Operational PTX3000
- Upgrading the FPCs in an Operational PTX5000

uplink-failure-detection

IN THIS SECTION

- Syntax | 30466
- Hierarchy Level | 30466
- Description | 30466

Required Privilege Level | 30466
Release Information | 30467

Syntax

```
uplink-failure-detection {
    action {
         log;
    }
    group group-name {
        debounce-interval seconds;
        link-to-monitor interface-name;
        link-to-disable interface-name;
    }
    traceoptions  {
        file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
        flag flag);
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Description

Configure uplink and downlink interfaces in a group to monitor uplink failures and to propagate uplink failures to the downlink interfaces.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Configuring Interfaces for Uplink Failure Detection (CLI Procedure)

uplink-failure-detection (Junos Fusion)

IN THIS SECTION

- Syntax | 30467
- Hierarchy Level | 30468
- Description | 30468
- Default | 30468
- Options | 30468
- Required Privilege Level | 30468
- Release Information | 30469

Syntax

```
uplink-failure-detection {
    <candidate-uplink-policy policy-name>;
}
```

Hierarchy Level

```
[edit chassis satellite-management]
[edit chassis satellite-management fpc slot-id]
```

Description

Enables uplink failure detection in a Junos Fusion.

The uplink failure detection feature on a Junos Fusion enables satellite devices to detect link failures on the uplink interfaces used to connect to aggregation devices. When uplink failure detection detects uplink failure on a satellite device, all of the device's extended ports (which connect to host devices) are shut down. Shutting down the extended ports allows downstream host devices to more quickly identify and adapt to the outage. For example, when a host device is connected to two satellite devices, and uplink failure detection shuts down the extended ports on one satellite device, the host device can more quickly recognize the uplink failure and redirect traffic through the other, active satellite device.

You can configure uplink failure detection in a Junos Fusion for a single satellite device using the **fpc *slot-id*** option. If uplink failure detection is enabled without specifying the **fpc *slot-id*** option, uplink failure detection is enabled for all cascade ports on the aggregation device.

If you enable uplink failure detection without the `candidate-uplink-policy` substatement, the default uplink failure detection settings are applied. To configure non-default uplink failure detection settings, include the `candidate-uplink-policy` substatement. Candidate uplink policies are configured under `[edit policy-options satellite-policies candidate-uplink-port-policy]`.

Default

Uplink failure detection is disabled.

Options

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

uplink-port-group (Junos Fusion)

IN THIS SECTION

- Syntax | 30469
- Hierarchy Level | 30469
- Description | 30470
- Default | 30470
- Options | 30470
- Required Privilege Level | 30470
- Release Information | 30470

Syntax

```
uplink-port-group group-name;
```

Hierarchy Level

```
[edit policy-options satellite-policies candidate-uplink-port-profile profile-name]
[edit policy-options satellite-policies candidate-uplink-port-profile profile-name term term-
name from]
```

Description

In Junos Fusion configuration, assign an uplink port group to a candidate uplink port policy.

An uplink port group defines a set of candidate uplink ports that are assigned to satellite devices to which the candidate uplink port group is assigned.

> **CAUTION**: The physically connected uplink ports on a satellite device must be defined as candidate uplink ports in the Junos Fusion configuration. If the uplink ports on a satellite device are not configured as candidate uplink ports, uplink failure detection cannot be enabled on the device, and a system log message is generated.

Uplink port groups are defined under `[edit policy-options satellite-policies port-group-alias]`.

Default

Each satellite device model has a set of default uplink ports (see *Overview of Uplink Failure Detection on a Junos Fusion*). You only need to assign an uplink port group to a satellite device if you do not want to use the default uplink ports.

Options

group-name	The name of the port group to assign.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

Configuring Junos Fusion Provider Edge

Configuring or Expanding a Junos Fusion Enterprise

upload (Event Policy)

IN THIS SECTION

- Syntax | 30471
- Hierarchy Level | 30471
- Description | 30471
- Options | 30471
- Required Privilege Level | 30472
- Release Information | 30472

Syntax

```
upload filename  (filename | committed) destination destination-name {
    retry-count count retry-interval seconds;
    transfer-delay seconds;
    user-name username;
}
```

Hierarchy Level

```
[edit event-options policy policy-name then]
```

Description

Configure an event policy to upload the committed configuration file or other specified files to a destination when triggered.

Options

destination *destination-name*

Name of the destination for the uploaded file. The destination must be defined in the `destinations` statement at the `[edit event-options]` hierarchy level.

filename (*filename* | committed) Name of the file to upload. Specify either the word `committed` to upload the most recently committed configuration file, or the filename of another file.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

`committed` option for `filename` statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

destinations (Event Policy) | 5533

Configure an Event Policy to Upload Files

upload-profile (Security Antivirus FTP)

IN THIS SECTION

- Syntax | 30472
- Hierarchy Level | 30473
- Description | 30473
- Required Privilege Level | 30473
- Release Information | 30473

Syntax

```
upload-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name anti-virus ftp]
```

Description

Configure a Content Security policy for the antivirus FTP (upload) protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

upload-profile (Security Content Filtering FTP)

IN THIS SECTION

- Syntax | 30474
- Hierarchy Level | 30474
- Description | 30474
- Required Privilege Level | 30474
- Release Information | 30474

Syntax

```
upload-profile profile-name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name content-filtering ftp]
```

Description

Configure a Content Security policy for the content-filtering FTP (upload) protocol and attach this policy to a security profile to implement it.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

upstream-interface

IN THIS SECTION

- Syntax | 30475
- Hierarchy Level | 30475
- Description | 30475
- Options | 30475

Required Privilege Level | 30476
Release Information | 30476

Syntax

```
upstream-interface [ interface-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast pim-to-igmp-proxy],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multicast pim-to-mld-proxy],
[edit logical-systems logical-system-name routing-options multicast pim-to-igmp-proxy],
[edit logical-systems logical-system-name routing-options multicast pim-to-mld-proxy],
[edit routing-instances routing-instance-name routing-options multicast pim-to-igmp-proxy],
[edit routing-instances routing-instance-name routing-options multicast pim-to-mld-proxy],
[edit routing-options multicast pim-to-igmp-proxy],
[edit routing-options multicast pim-to-mld-proxy]
```

Description

Configure at least one, but not more than two, upstream interfaces on the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain. The RP routing device translates PIM join or prune messages into corresponding IGMP report or leave messages (if you include the `pim-to-igmp-proxy` statement), or into corresponding MLD report or leave messages (if you include the `pim-to-mld-proxy` statement). The routing device then proxies the IGMP or MLD report or leave messages to one or both upstream interfaces to forward IPv4 multicast traffic (for IGMP) or IPv6 multicast traffic (for MLD) across the PIM domains.

Options

interface-names—Names of one or two upstream interfaces to which the RP routing device proxies IGMP or MLD report or leave messages for transmission of multicast traffic across PIM domains. You can specify a maximum of two upstream interfaces on the RP routing device. To configure a set of two

upstream interfaces, specify the full interface names, including all physical and logical address components, within square brackets ([ ]).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring PIM-to-IGMP Message Translation

Configuring PIM-to-MLD Message Translation

upstream-label

IN THIS SECTION

- Syntax | 30477
- Hierarchy Level | 30477
- Description | 30477
- Options | 30477
- Required Privilege Level | 30477
- Release Information | 30477

Syntax

```
upstream-label {
    vlan-id vlan-id;
}
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
```

Description

Specify the upstream label for the bidirectional label-switched path (LSP).

Options

vlan-id *vlan-id* VLAN ID to be used for the Generalized MPLS (GMPLS) VLAN LSP at the ingress provider edge (PE) to customer edge (CE) link.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring MPLS LSPs for GMPLS

upstream-mode

IN THIS SECTION

- Syntax | 30478
- Hierarchy Level | 30478
- Description | 30478
- Required Privilege Level | 30478
- Release Information | 30479

Syntax

```
upstream-mode;
```

Hierarchy Level

```
[edit logical-systems name protocols router-advertisement interface],
[edit protocols router-advertisement interface]
```

Description

Configure the interface as upstream interface for RA proxy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.1R1

upstream-rate (Traffic Shaping)

IN THIS SECTION

- Syntax | 30479
- Hierarchy Level | 30479
- Description | 30479
- Options | 30480
- Required Privilege Level | 30480
- Release Information | 30480

Syntax

```
upstream-rate rate;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-interface-
unit advisory-options],
[edit dynamic-profiles profile-name interfaces interface-set $junos-interface-set-name interface
$junos-interface-ifd-name advisory-options],
[edit interfaces demux0 unit logical-unit-number advisory-options],
[edit interfaces interface-name logical-unit-number advisory-options]
```

Description

Specify a recommended shaping rate to be applied to upstream traffic on an interface.

For ANCP interfaces, this configured rate is used as the default value for the Juniper VSA Upstream-Calculated-Qos-Rate (26-142) when the router has not received and processed the attributes from the access node.

For L2TP, the rate is configured on an underlying PPPoE logical interface for a subscriber on an MX Series router acting as a LAC. When the subscriber is tunneled, this rate, referred to as speed for L2TP, is sent to the LNS in the ICCN message as AVP 38.

Options

rate—Traffic rate in bits per second.

- **Range:** 1000 through 4,294,967,295 bits per second

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

Support at the `[edit interfaces demux0 ...]` hierarchy level introduced in Junos OS Release 12.2.

Support at the `[edit dynamic-profiles ...]` hierarchy level introduced in Junos OS Release 13.1.

RELATED DOCUMENTATION

ANCP Agent Traffic Shaping and CoS

ANCP Agent Neighbors and Operations

Configuring the Method to Derive the LAC Connection Speeds Sent to the LNS

upstream-rewrite

IN THIS SECTION

- Syntax | 30481
- Hierarchy Level | 30481
- Description | 30481
- Options | 30482
- Required Privilege Level | 30482
- Release Information | 30482

Syntax

```
upstream-rewrite {
 dscp dscp;
 dscp-n3 dscp-n3;
 forwarding-class forwarding-class;
 loss-priority(high | low);
}
```

Hierarchy Level

```
[edit services agf fiveqi-map]
```

Description

Rewrite QoS attributes for upstream traffic in AGF in the 5QI table.

Options

dscp

Differentiated services code point (DSCP) value on the IP packet containing user data sent to the UPF on N3 Interface. DSCP is a 6-bit value represented by decimal value between 0 to 63.

- **Range:** 0 through 63

dscp-n3

DSCP value of the GTP-U encapsulated packet sent to the UPF on N3 Interface. DSCP is a 6-bit value represented by decimal value between 0 to 63.

- **Range:** 0 through 63

forwarding-class

Name of forwarding class. Groups packets for transmission and assigns packets to output queues.

loss-priority

Apply the packet loss priority (PLP) for dropping packets.

- high—High probability of being dropped at times of congestion
- low—Low probability of being dropped at times of congestion

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 23.2R1.

RELATED DOCUMENTATION

CoS for Subscriber Access Overview

fiveqi-map | **7199**

show services agf fiveqi remarking | **25331**

urgent-pointer

IN THIS SECTION

- Syntax | 30483
- Hierarchy Level | 30483
- Description | 30483
- Options | 30483
- Required Privilege Level | 30484
- Release Information | 30484

Syntax

```
urgent-pointer {
    match (equal | greater-than | less-than | not-equal);
    value urgent-pointer;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify the data in the packet is urgent; the URG flag must be set to activate this field.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`urgent-pointer`*—Match the value of the urgent pointer.
- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

uri-check

IN THIS SECTION

- Syntax | 30484
- Hierarchy Level | 30484
- Description | 30484
- Required Privilege Level | 30485
- Release Information | 30485

Syntax

```
uri-check;
```

Hierarchy Level

```
[edit security utm default-configuration anti-virus scan-options]
```

Description

Perform Sophos antivirus Uniform Resource Identifier (URI) checking. URI checking is a way of analyzing URI content in HTTP traffic against a remote Sophos database to identify malware or malicious content. URI checking is on by default.

NOTE: Starting in Junos OS release 18.4R1, the URI checking is off by default.

You can disable Sophos antivirus URI checking with the `no-uri-check` statement.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

url

IN THIS SECTION

- Syntax | 30485
- Hierarchy Level | 30486
- Description | 30486
- Options | 30486
- Required Privilege Level | 30486
- Release Information | 30486

Syntax

```
url url-name;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name redirect],
[edit services pcef pcc-action-profiles profile-name redirect]
```

Description

Specify the URL name that you want a PCC action profile to use for performing HTTP redirection. If you configure this, the PCC action profile can only be used in PCC rules that match only HTTP-based applications and all flows.

If you are using Junos OS Subscriber Aware, specify the URL name at the `[edit unified-edge pcef pcc-action-profiles profile-name redirect]` hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the URL name at the `[edit services pcef pcc-action-profiles profile-name redirect]` hierarchy level.

Options

url-name	URL for the HTTP redirect.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support at the `[edit services pcef pcc-action-profiles profile-name redirect]` hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

Configuring Policy and Charging Control Action Profiles for Subscriber Management

url (Security)

IN THIS SECTION

- Syntax | 30487
- Hierarchy Level | 30487
- Description | 30487
- Required Privilege Level | 30488
- Release Information | 30488

Syntax

```
url url;
```

Hierarchy Level

```
[edit security dynamic-address feed-server feed-server-name]
```

Description

URL of the bundle file for feed server configuration. Use a browser to verify the URL validity. Examples of URLs are shown below:

- `example.url.com/my_file.tgz`

- `example.url.com/my_file.bundle`

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Dynamic Address Groups in Security Policies

url-blacklist

IN THIS SECTION

- Syntax | 30488
- Hierarchy Level | 30489
- Description | 30489
- Options | 30489
- Required Privilege Level | 30489
- Release Information | 30489

Syntax

```
url-blacklist listname;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering]
```

Description

This is a global blocklist category, blocking content for Web filtering.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

url-feed

IN THIS SECTION

- Syntax | 30490
- Hierarchy Level | 30490
- Description | 30490
- Options | 30492
- Required Privilege Level | 30492
- Release Information | 30492

Syntax

```
url-feed name {
 interval interval;
 no-autoupdate;
 proxy-profile proxy-profile;
 routing-instance routing-instance;
 url url;
}
```

Hierarchy Level

```
[edit logical-systems name security utm custom-objects],
[edit logical-systems name tenants name security utm custom-objects],
[edit security utm custom-objects],
[edit tenants name security utm custom-objects]
```

Description

Configure URL feed object. The URL feed object reduces your effort to add multiple URLs into a single URL pattern. The URL feed enables you to create a URL pattern automatically.

To create a URL pattern automatically, you need to create a URL feed file on the HTTP server. You should add the URLs that need to be added in the URL pattern to this file. When you configure the URL feed on your security device, the system downloads the file from the HTTP server and creates a URL pattern as follows:

1. The system validates the list of URLs based on the existing validation criteria used for URL pattern.
2. The system creates a URL pattern with `_feed` as suffix to the URL feed name.

   For example, if the URL feed name is `juniper`, then the system creates the URL pattern name as `juniper_feed`.
3. The system adds only the URLs that have a valid syntax and ignores the URLs with invalid syntax.
4. The system periodically downloads the file from the HTTP server based on the `update-interval` time you configure. There may be five minutes delay in the timer functionality.

5. The system maintains a hash to check for any updates to the URL list in the downloaded file. The system skips the file validation and URL pattern updates if the hash matches with the previous file.
6. The system does not automatically download the file and update the URL pattern if you configure `no-autoupdate`.

Limitations of URL Feed

- Files on the HTTP server should be in plain text format.
- File should contain only one URL in one line.
- Number of URLs in a file should not exceed the limit of the corresponding platform. See Table 1833 on page 30491.
- The URL length should not be more than 249 characters.
- The URL should not have invalid syntax.
- Deletion of a URL feed does not delete corresponding URL pattern and corresponding custom URL category.
- The system prioritizes the URL pattern that created automatically using URL feed over the URL pattern that you created manually using CLI.
- The file size should not be more than 512 KB.

Table 1833: Number of URLs Supported in One URL Pattern

Platforms	Number of URLs Supported
cSRX	1000
vSRX	1000
SRX300, SRX320, SRX340, SRX345, and SRX380	1000
SRX1500, SRX1600,SRX4200, SRX4600, and SPC3	1000
SRX5400, SRX5600, and SRX5800	1500

Options

url-feed *name*	Configure the name of the URL feed object
interval	Interval in hours for auto-updating list • **Default:** 24 • **Range:** 1 through 168
no-autoupdate	Disable automatic update for URL feed pattern
proxy-profile	Proxy profile
routing-instance	Routing instance name
url	Configure the HTTPS server file URL that contains the list of URLs

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

show security utm custom-objects url-feed status | 25207

url-filter

IN THIS SECTION

Syntax | 30493
Hierarchy Level | 30494

- Description | 30494
- Options | 30494
- Required Privilege Level | 30494
- Release Information | 30494

Syntax

```
url-filter {
    profile  profile-name {
        template template-name {
            client-interfaces [ client-interface-name1 client-interface-name2 ];
            disable-url-filtering;
            dns-resolution-interval minutes;
            dns-resolution-rate seconds;
            dns-retries number;
            dns-routing-instance dns-routing-instance-name;
            dns-server [ ip-address1 ip-address2 ip-address3 ];
            dns-source-interface loopback-interface-name;
            routing-instance routing-instance-name;
            server-interfaces  [ server-interface-name1 server-interface-name2 ];
            term term-name {
                from {
                    src-ip-prefix [prefix1 prefix2];
                    dest-port   [port1 port2];
                }
                then {
                    accept;
                    custom-page custom-page;
                    http-status-code http-status-code;
                    redirect-url redirect-url;
                    tcp-reset;
                }
            }
            url-filter-database filename
        }
        url-filter-database filename;
```

```
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure URL filtering service.

> **NOTE**: Starting in Junos OS Release 18.3R1, the `url-filter` statement is deprecated and has been replaced by the `web-filter` statement. The `url-filter` statement is supported for backward compatibility.

Options

`url-filter-database` ***`filename`*** Specify the filename of the URL filter database. This option is mandatory.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Configuring URL Filtering

URL Filtering Overview

url-filter-profile

IN THIS SECTION

- Syntax | 30495
- Hierarchy Level | 30495
- Description | 30495
- Options | 30496
- Required Privilege Level | 30496
- Release Information | 30496

Syntax

```
url-filter-profile profile-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the URL filter profile that the service set uses. The URL filter profile specifies how to filter access to disallowed URLs, and is configured at the `[edit services url-filter]` hierarchy level.

NOTE: You must also configure the `next-hop-service` statement with this statement.

NOTE: Starting in Junos OS Release 18.3R1, the `url-filter-profile` statement is deprecated and has been replaced by the `web-filter-profile` statement. The `url-filter-profile` statement is supported for backward compatibility.

Options

profile-name	Name of the URL filter profile.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

Configuring URL Filtering

URL Filtering Overview

url-filter | 30492

url-filter-template

IN THIS SECTION

- Syntax | 30497
- Hierarchy Level | 30497
- Description | 30497
- Options | 30498
- Required Privilege Level | 30499
- Release Information | 30499

Syntax

```
url-filter-template template-name {
    client-interfaces [ client-interface-name1 client-interface-name2 ];
    disable-url-filtering;
    dns-resolution-interval minutes;
    dns-resolution-rate seconds;
    dns-retries number;
    dns-routing-instance dns-routing-instance-name;
    dns-server [ ip-address1 ip-address2 ip-address3 ];
    dns-source-interface loopback-interface-name;
    routing-instance routing-instance-name;
    security-intelligence-policy
    server-interfaces  [ server-interface-name1 server-interface-name2 ];
    term term-name {
        from {
            src-ip-prefix [prefix1 prefix2];
            dest-port   [port1 port2];
        }
        then {
            accept;
            custom-page custom-page;
            http-status-code http-status-code;
            redirect-url redirect-url;
            tcp-reset;
        }
    }
    url-filter-database filename
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name]
```

Description

Configure a URL filter template.

Options

template-name	Name of the URL filter template.
client-interfaces [*client-interface-name1 client-interface-name2*]	The list of client-facing logical interfaces (uplink) on which the URL filtering is configured. This option is mandatory.
disable-url-filtering	Disables the filtering of HTTP traffic that contains an embedded IP address (for example, http:/10.1.1.1) belonging to a disallowed domain name in the URL filter database.
dns-resolution-interval *minutes*	DNS resolution time interval in minutes. • **Default:** 1440 • **Range:** 60 through 1440 minutes.
dns-resolution-rate *seconds*	Number of DNS queries per second sent out from the system before initiating further DNS queries. • **Default:** 50 • **Range:** 50 through 100.
dns-retries *number*	Number of retries for a DNS query in case query fails or times out. • **Default:** 3 • **Range:** 1 through 5.
dns-routing-instance *dns-routing-instance-name*	The VRF on which the DNS server is reachable. This option is mandatory. You can use the default routing instance inet.0 or a defined routing instance.
dns-server [*ip-address1 ip-address2 ip-address3*]	One or more IP (IPv4 or IPv6) addresses of DNS servers to which the DNS queries are sent out. This option is mandatory.
dns-source-interface *loopback-interface-name*	The loopback interface for which source IP address is picked for sending DNS queries. This option is mandatory.
routing-instance *routing-instance-name*	The VRF on which URL filtering feature is configured. This option is mandatory. You can use the default routing instance inet.0 or a defined routing instance.
server-interfaces [*server-interface-name1 server-interface-name2*]	Server-facing interfaces to which traffic is destined. This option is mandatory.

	The list of server-facing logical interfaces (downlink) on which the URL filtering is configured. This option is mandatory.
`url-filter-database` *`filename`*	The filename of the URL filter database. The file should be placed in the **/var/db/url-filterd** directory, but indicate just the filename here and not the full path.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series.

Statement introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

Configuring URL Filtering

url-pattern

IN THIS SECTION

- Syntax | 30500
- Hierarchy Level | 30500
- Description | 30500
- Options | 30502
- Required Privilege Level | 30502
- Release Information | 30502

Syntax

```
url-pattern object-name {
    value [value];
}
```

Hierarchy Level

```
[edit security utm custom-objects]
```

Description

Use URL pattern lists to create custom URL category lists. These are lists of patterns that bypass scanning.

> **WARNING**: Custom category does not take precedence over predefined categories when it has the same name as one of the predefined categories. We do not recommend having a custom category name be the same as the predefined category name.

Starting in Junos OS Release 20.4R1, the URL filtering supports the regular expression format given in Table 1834 on page 30500.

Table 1834: Regular Expression Format

Syntax	Pattern Format	Description	Example

Asterisk (*) in domain name	Pattern = [*].sub-domain..sub-domain	Asterisk should be at the head only.
		
		Match 0-N words in domain name.
			`*.juniper.net`
			
			`*.net`

Caret (^) in domain name	Pattern = [^]..[^].sub-domain.[^]..[^].sub-domain.[^]..[^]	Caret could be at any place.
		
		Match one word in domain name.
			`^.google.^.^`
			
			`a.^.b`

Table 1834: Regular Expression Format *(Continued)*

Syntax	Pattern Format	Description	Example

Question Mark (?) in domain name	Pattern = sub-domain.sub-domain.sub-do[?][?]	Question mark should be at the tail only.
		
		Match one character in domain name.
			`dev.local?`
			
			`www.yahoo.??`

All wildcard	Pattern = *, or *.*, or *.*.*	Special pattern. Match all URLs.	Same as pattern.

Prefix in URL path	Pattern = <domain-name>/[prefix]	Match the longest prefix in the URL path.	`<domain>/watch?`
			
			`<domain>/news/`

Keywords in URL path	Pattern = <domain-name>/[prefix][*token][*token][*token]	Match keywords in URL path.
		
		Support 0-3 tokens.
			`<domain>/pub/*/crypto/*.gz`
			
			`<domain>/video?*key1*key`

The system validate the URL patterns when you commit the configuration. If you configure an invalid pattern, the system shows a popup warning with the first bad character in (). For example, the pattern `a.*.com` shows a warning message `Pattern a.(*).com not supported`.

Starting in Junos OS Release 20.4R1, we have introduced a golden match for multiple URL patterns. When you configure multiple patterns for one domain, sometimes a specific URL could match multiple patterns, then the URL filtering module selects the best out of these patterns, which is the golden match. For the selection of the golden match, the URL filtering module prioritizes the URL pattern in the following sequence:

1. Select the best domain name pattern to match.
2. Within the same domain name pattern, select the longest prefix match of the URL path.
3. Within the same domain name pattern and same URL path prefix, keyword match will take the priority.

For example, if you configure the following four patterns:

1. Pattern 1 = finance.abc.com/gb/chinainternet/

2. Pattern 2 = finance.abc.^/gb/chinamkt/
3. Pattern 3 = finance.abc.^/gb/chinamkt/*.shtml
4. Pattern 4 = *.abc.com/gb/chinamkt/chinamkt_cn

For the URL https://finance.abc.com/gb/chinamkt/chinamkt_cn/sinacn/2020-03-29/doc-ifzuwpay8845719.shtml, the golden match section will be as follows:

1. URL filtering module considers all the four patterns as a potential match in the domain name match stage, and the priority order is:

   finance.abc.com > finance.abc.^ > *.abc.com.

2. Within the same domain name pattern, the URL filtering module considers the longest prefix match of the URL path.

   - Pattern 1 does not match the URL path and the URL filtering module skips pattern 1.
   - Pattern 2 and pattern 3 have the same prefix in the URL path. Hence, the keywords match controls the golden match selection. Finally, the URL filtering module prefers pattern 3 that has the longest keywords match as the golden match.

Options

- *object-name*—Name of the URL list object.
- `value` *value*—Value of the URL list object. You can configure multiple values separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Security Policies User Guide for Security Devices

url (Security)

IN THIS SECTION

- Syntax | 30503
- Hierarchy Level | 30503
- Description | 30503
- Options | 30503
- Required Privilege Level | 30504
- Release Information | 30504

Syntax

```
url url-name;
```

Hierarchy Level

```
[edit security pki ca-profile ca-profile-name enrollment],
[edit security pki ca-profile ca-profile-name revocation-check crl]
```

Description

(Supported on Adaptive Services (AS) and MultiServices PICs only.) Specify the certificate authority (CA) URL to use in requesting digital certificates or the URL for the Lightweight Directory Access Protocol (LDAP) location from which the certificate revocation list (CRL) is retrieved.

Options

url-name—Location of the CA to which enrollment requests are sent or LDAP location of the CRL. With Simple Certificate Enrollment Protocol (SCEP), you enroll CA certificates with the `request security pki ca-certificate enroll` command and specify the CA profile. There is no separate command to enroll CA certificates with CMPv2.

The format of the URL is protocol http.

Required Privilege Level

admin—To view the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Specifying an Enrollment URL

Specifying an LDAP URL

crl (Adaptive Services Interface) | **4718**

url (Security Antivirus)

IN THIS SECTION

- Syntax | **30504**
- Hierarchy Level | **30505**
- Description | **30505**
- Required Privilege Level | **30505**
- Release Information | **30505**

Syntax

```
url url;
```

Hierarchy Level

```
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update]
[edit security utm feature-profile anti-virus sophos-engine pattern-update]
[edit security utm default-configuration anti-virus avira-engine pattern-update]
```

Description

Specify the URL for the pattern database. You should not change the default URL unless you are experiencing problems with it and have called for support.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5 for Juniper Express engine and Kaspersky Lab engine. Support for Sophos engine added in Junos OS Release 11.1.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Support for Avira engine added in Junos OS Release 18.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

url (Security IDP)

IN THIS SECTION

- Syntax | 30506
- Hierarchy Level | 30506
- Description | 30506
- Required Privilege Level | 30506
- Release Information | 30506

Syntax

```
url url-name;
```

Hierarchy Level

```
[edit security idp security-package]
```

Description

Specify the URL to automatically download the updated signature database.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

url-whitelist

IN THIS SECTION

- Syntax | 30507
- Hierarchy Level | 30507
- Description | 30507
- Options | 30507
- Required Privilege Level | 30507
- Release Information | 30508

Syntax

```
url-whitelist listname;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering]
```

Description

A URL allowlist is a unique custom list that you define in which all the URLs or IP addresses in that list for a specified category are always bypassed for filtering

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

url-whitelist

IN THIS SECTION

- Syntax | 30508
- Hierarchy Level | 30508
- Description | 30508
- Required Privilege Level | 30508
- Release Information | 30509

Syntax

```
url-whitelist listname;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus]
```

Description

A URL allowlist is a unique custom list that you define in which all the URLs or IP addresses in that list for a specified category are always bypassed for scanning.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

usb

IN THIS SECTION

- Syntax | 30509
- Hierarchy Level | 30509
- Description | 30509
- Options | 30510
- Required Privilege Level | 30510
- Release Information | 30510

Syntax

```
usb {
    disable;
}
```

Hierarchy Level

```
[edit system autoinstallation]
[edit chassis]
```

Description

Disable the USB autoinstallation process.

Options

`disable`—Disable the process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Autoinstallation of Configuration Files

usb-control

IN THIS SECTION

- Syntax | 30510
- Hierarchy Level | 30511
- Description | 30511
- Options | 30511
- Required Privilege Level | 30511
- Release Information | 30511

Syntax

```
usb-control {
    command binary-file-path;
```

```
    disable;
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the universal serial bus (USB) supervise process.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the universal serial bus (USB) supervise process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

use-active-child-mac-on-reth

IN THIS SECTION

- Syntax | 30512
- Hierarchy Level | 30512
- Description | 30512
- Required Privilege Level | 30512

Release Information | 30512

Syntax

```
use-active-child-mac-on-reth;
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Copy child's physical MAC address on RETH parent's current MAC address on KVM based systems.

NOTE: The commands `use-active-child-mac-on-reth` and `use-actual-mac-on-physical-interfaces` need to be configured together and the device rebooted for the feature to work.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 19.4R1.

use-actual-mac-on-physical-interfaces

IN THIS SECTION

Syntax | 30513

- Hierarchy Level | 30513
- Description | 30513
- Required Privilege Level | 30513
- Release Information | 30513

Syntax

```
use-actual-mac-on-physical-interfaces;
```

Hierarchy Level

```
[edit chassis cluster]
```

Description

Configure this command so that SR-IOV uses hypervisor's provided MAC address for physical interfaces on KVM based systems.

NOTE: The commands `use-active-child-mac-on-reth` and `use-actual-mac-on-physical-interfaces` need to be configured together and the device rebooted for the feature to work.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 19.4R1.

use-class (Class Attribute)

IN THIS SECTION

- Syntax | 30514
- Hierarchy Level | 30514
- Description | 30514
- Required Privilege Level | 30514
- Release Information | 30515

Syntax

```
use-class {
    regex "value";
    pattern "pattern";
    subscription-id-type (imsi | msisdn | nai | private | sip-uri);
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domains domain-name subscription-id]
```

Description

Configure a regular expression to parse the Class attribute contents, specify characters to insert between the resulting regular expression groups, and specify the subscription ID type if you configured `subscription-id-options` *`entry-name`* `use-class` under `[edit unified-edge gateways tdf` *`gateway-name`* `domains` *`domain-name`* `subscription-id]`.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

use-extended-flow-memory

IN THIS SECTION

- Syntax | 30515
- Hierarchy Level | 30515
- Description | 30516
- Required Privilege Level | 30516
- Release Information | 30516

Syntax

```
use-extended-flow-memory;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services]
```

Description

Configure the service to extended flow memory. This service provides more scale in flows for inline services sampling.

The new configuration `set chassis fpc slot slot-number inline-services use-extended-flow-memory` allows you to configure table to operate in side band mode with side band memory. This configuration is applicable only on a Lookup (LU) platform. It is not applicable for XL line card because XL has dedicated DMEM memory to hold 64M flow entries.

If you are configuring inline flow monitoring of MPLS-IPv6 flows on an LU platform, you must configure `use-extended-flow-memory` to get MPLS-IPv6 flow records. If you do not configure `use-extended-flow-memory` on an LU platform, plain MPLS flow records are created.

> **NOTE**: This configuration is supported only on LU platforms. The LU platform line cards are MPC1E, MPC2E, MPC3E, MPC4E, and MPC 3D 16x10GE. Exceptions to the MPC2E and MPC3E series line card are MPC2E-NG and MPC3E-NG that use XL platform.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers

use-for-post-convergence-lfa

IN THIS SECTION

- Syntax | 30517
- Hierarchy Level | 30517
- Description | 30517
- Required Privilege Level | 30517
- Release Information | 30517

Syntax

```
use-for-post-convergence-lfa;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options fate-sharing group group-name],
[edit logical-systems name routing-options fate-sharing group group-name],
[edit routing-instances name routing-options fate-sharing group group-name],
[edit routing-options fate-sharing group group-name]
```

Description

Use this fate-sharing group as a constraint for the backup path computed by topology-independent loop-free alternate (TI-LFA).

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 17.4R1 for MX Series, PTX Series, and QFX Series.

RELATED DOCUMENTATION

Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS

Configuring Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS

post-convergence-lfa | 12903

use-post-convergence-lfa | 30531

node-protection | 11651

use-imported-time-zones

IN THIS SECTION

- Syntax | 30518
- Hierarchy Level | 30518
- Description | 30518
- Required Privilege Level | 30519
- Release Information | 30519

Syntax

```
use-imported-time-zones;
```

Hierarchy Level

```
[edit system]
```

Description

Configure a custom time zone from a locally generated time-zone database.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Updating the IANA Time Zone Database on Junos OS Devices

use-interface

IN THIS SECTION

- Syntax | 30519
- Hierarchy Level | 30519
- Description | 30520
- Required Privilege Level | 30520
- Release Information | 30520

Syntax

```
use-interface-description {logical |device};
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcp-client client-
identifier]
```

Description

The description configured at the physical or logical interface level is used for client identification.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

use-interface-description (MX)

IN THIS SECTION

- Syntax | 30520
- For Platforms with Enhanced Layer 2 Software (ELS) | 30521
- For Platforms Without ELS | 30521
- For MX Series Platforms | 30521
- Description | 30521
- Options | 30521
- Required Privilege Level | 30522
- Release Information | 30522

Syntax

```
use-interface-description (device | logical);
```

For Platforms with Enhanced Layer 2 Software (ELS)

```
[edit vlans vlan-name forwarding-options dhcp-security option-82 circuit-id]
```

For Platforms Without ELS

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name) dhcp-option82 circuit-id],
[edit forwarding-options helpers bootp dhcp-option82 circuit-id],
[edit forwarding-options helpers bootp interface interface-name dhcp-option82 circuit-id],
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name) dhcp-option82 remote-id],
[edit forwarding-options helpers bootp dhcp-option82 remote-id],
[edit forwarding-options helpers bootp interface interface-name dhcp-option82 remote-id]
```

For MX Series Platforms

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security option-82circuit-id]
```

Description

Use the interface description rather than the interface name (which is the default value) in the circuit ID or remote ID value in the DHCP option 82 information.

The textual description is configured using the `description` statement at the `[edit interfaces interface-name]` hierarchy level. If you specify that the textual description be used and no description is configured for the interface, DHCP relay defaults to using the Layer 2 interface name. When you use the interface description rather than the interface name, the interface description has to be specified under interface unit ("set interfaces ge-0/0/0 unit 0 description "client"). If you do not do this, then the interface name is used.

Options

device Use the device interface description. Only available for MX Series platform configuration.

logical Use the logical interface description. Only available for MX Series platform configuration.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

Hierarchy level `[edit bridge-domains bridge domain name forwarding-options dhcp-security]` introduced in Junos OS Release 14.1 for the MX Series.

RELATED DOCUMENTATION

Example: Setting Up DHCP Option 82 Using the Same VLAN

Example: Setting Up DHCP Option 82

http://tools.ietf.org/html/rfc3046

use-interface-description

IN THIS SECTION

- Syntax | 30523
- Hierarchy Level | 30523
- Description | 30523
- Options | 30524
- Required Privilege Level | 30524
- Release Information | 30524

Syntax

```
use-interface-description (logical | device);
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name (relay-agent-interface-id | relay-
agent-remote-id)],
[edit forwarding-options dhcp-relay relay-option-82 (circuit-id | remote-id)],
[edit forwarding-options dhcp-relay group group-name relay-option-82 (circuit-id | remote-id)],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 (relay-agent-
interface-id | relay-agent-remote-id)],
[edit logical-systems logical-system-name forwarding-options dhcp-relay relay-option-82 (circuit-
id | remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6  (relay-agent-
interface-id | relay-agent-remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-82
(circuit-id | remote-id)],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-18],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-37]
```

Description

Use the textual interface description instead of the interface identifier in the DHCP base option 82 Agent Circuit ID (suboption 1) or Agent Remote ID (suboption 2) information, or in the DHCPv6 option 18 (Relay Agent Interface ID) or option 37 (Relay Agent Remote ID) information in DHCP packets that the DHCP relay agent sends to a DHCP server.

> **NOTE**: For integrated routing and bridging (IRB) interfaces, the option 82 field must be able to uniquely identify the incoming interface based on either the Agent Circuit ID or Agent Remote ID. You can modify the information in the textual interface description to match the raw IFD (physical interface without a subunit) name and configure the option 82 field to use the interface description.

The textual description is configured using the `description` statement at the `[edit interfaces interface-name]` hierarchy level. If you specify that the textual description be used and no description is configured for the interface, DHCP relay defaults to using the Layer 2 interface name. When you use the interface description rather than the interface name, the interface description has to be specified under interface

unit ("set interfaces ge-0/0/0 unit 0 description "client"). If you do not do this, then the interface name is used.

In the case of integrated routing and bridging (IRB) interfaces, the textual description of the Layer 2 interface is used instead of the IRB interface. If there is no description configured, the Layer 2 logical interface name is used. To include the IRB interface description instead of the Layer 2 interface description, configure the `use-interface-description` and the `no-vlan-interface-name` statements. If no description is configured for the IRB interface, DHCP relay defaults to using the IRB interface name.

> **NOTE**: The `use-interface-description` statement is mutually exclusive with the `use-vlan-id` statement.

If you specify the textual interface description, rather than accepting the default syntax, the identification is for packets returned from the server, and only for instances where that identification would be required by the DHCP relay, such as a stateless pass-through.

> **NOTE**: By default, DHCP relay accepts a maximum of 253 ASCII characters. If the textual interface description exceeds 253 characters, DHCP relay drops the packet, which results in the DHCP client failing to bind.

Options

`logical`—Use the textual description that is configured for the logical interface.

`device`—Use the textual description that is configured for the device interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... relay-agent-remote-id]` and `[edit ... remote-id]` hierarchy levels introduced in Junos OS Release 14.1.

Support at the `[edit vlans vlan-name dhcp-security dhcpv6-options option-18]` and `[edit vlans vlan-name dhcp-security dhcpv6-options option-37]` hierarchy levels introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

RELATED DOCUMENTATION

Using DHCP Option 82 Information

Using DHCP Relay Agent Option 82 Information

Configuring DHCPv6 Relay Agent Options

use-interface-index

IN THIS SECTION

- Syntax | 30525
- Hierarchy Level | 30525
- Description | 30526
- Options | 30526
- Required Privilege Level | 30526
- Release Information | 30526

Syntax

```
use-interface-index (logical | device);
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-18],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-37]
```

Description

Use the index number of the interface instead of the interface name in the DHCPv6 option 18 (Relay Agent Interface-ID) or option 37 (Relay Agent Remote-ID). These options are used by a relay agent to insert information in DHCPv6 requests before the relay agent forwards them to a DHCPv6 server.

Options

`logical`—Use the textual description that is configured for the logical interface.

`device`—Use the textual description that is configured for the device interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Including a Textual Description in DHCP Options

Using DHCP Relay Agent Option 82 Information

Configuring DHCPv6 Relay Agent Options

use-interface-name

IN THIS SECTION

- Syntax | 30527
- Hierarchy Level | 30527
- Description | 30527
- Options | 30527

Required Privilege Level | 30527
Release Information | 30527

Syntax

```
use-interface-name (logical | device);
```

Hierarchy Level

```
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-18],
[edit vlans vlan-name forwarding-options dhcp-security dhcpv6-options option-37]
```

Description

Configure DHCPv6 option 18 (Relay Agent Interface-ID) or option 37 (Relay Agent Remote-ID) to use the interface name to identify the port identity of the DHCP client to the DHCP server.

Options

`logical`—Use the name that is configured for the logical interface.

`device`—Use the name that is configured for the device interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Including a Textual Description in DHCP Options

Using DHCP Relay Agent Option 82 Information

Configuring DHCPv6 Relay Agent Options

use-option-82

IN THIS SECTION

- Syntax | 30528
- Hierarchy Level | 30528
- Description | 30529
- Options | 30529
- Required Privilege Level | 30529
- Release Information | 30529

Syntax

```
use-option-82 <strict>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 (relay-agent-interface-id | relay-agent-remote-id)],
[edit forwarding-options dhcp-relay dhcpv6 group group-name (relay-agent-interface-id | relay-agent-
remote-id)],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 ... (relay-agent-
interface-id | relay-agent-remote-id)],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 ... (relay-agent-interface-id | relay-agent-remote-id)],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 ... (relay-
agent-interface-id |relay-agent-remote-id)]
```

Description

For the `relay-agent-interface-id` statement, specify that in dual-stack environments, the DHCPv6 relay agent uses an Interface-ID option (option 18) that is based on the DHCPv4 relay agent information option (option 82). When you include this statement, the DHCPv6 relay agent checks for the option 82 Agent Circuit-ID suboption (suboption 1) and inserts it into the outgoing RELAY-FORW message.

For the `relay-agent-remote-id` statement, specify that in dual-stack environments, the DHCPv6 relay agent uses a Remote-ID option (option 37) that is based on the DHCPv4 relay agent information option (option 82). When you include this statement, the DHCPv6 relay agent checks for the option 82 Remote-ID suboption (suboption 2) and inserts it into the outgoing RELAY-FORW message.

Options

strict (Optional) The router drops Solicit messages that do not include a DHCPv4 Remote-ID (option 82 suboption 2). If you do not specify the `strict` keyword, the router sends the RELAY-FORW message without adding option 37.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

Support at the `[...relay-agent-remote-id]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Using DHCP Relay Agent Option 82 Information

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets

use-p2mp-lsp

IN THIS SECTION

- Syntax | 30530
- Hierarchy Level | 30530
- Description | 30530
- Required Privilege Level | 30530
- Release Information | 30531

Syntax

```
    igmp-snooping-options {
        use-p2mp-lsp;
    }
}
```

Hierarchy Level

```
[edit routing-instances instance name igmp-snooping-options]
```

Description

Point-to-multipoint LSP for IGMP snooping enables multicast data traffic in the core to take the point-to-multipoint path. The effect is a reduction in the amount of traffic generated on the PE router when sending multicast packets for multiple VPLS sessions because it avoids the need to send multiple parallel streams when forwarding multicast traffic to PE routers participating in the VPLS. Note that the options configured for IGMP snooping are applied on a per-routing-instance so all IGMP snooping routes in the same instance will use the same mode, point to multipoint or pseudowire.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Configuring Point-to-Multipoint LSP with IGMP Snooping

show igmp snooping options | **20179**

multicast-snooping-options | **11087**

use-post-convergence-lfa

IN THIS SECTION

- Syntax | **30531**
- Hierarchy Level | **30531**
- Description | **30532**
- Options | **30532**
- Required Privilege Level | **30532**
- Release Information | **30533**

Syntax

```
use-post-convergence-lfa <maximum-backup-paths maximum-backup-paths> <maximum-labels maximum-
labels> ;
```

Hierarchy Level

```
[edit logical-systems name protocols isis backup-spf-options],
[edit logical-systems name routing-instances name protocols isis backup-spf-options],
```

```
[edit protocols isis backup-spf-options],
[edit routing-instances name protocols isis backup-spf-options]
```

Description

Calculate post-convergence MPLS fast reroute (FRR) backup next hops for the IS-IS protocol using segment routing (SR). Topology-independent loop-free alternate (TI-LFA) provides protection against link failure, node failure, and fate-sharing failures. By default, Junos OS provides link protection. Junos OS allows you to control the maximum number of equal-cost multipath (ECMP) backup paths installed for a given destination. Junos OS also allows you to control the maximum number of labels in the installed backup paths. Configure the `use-source-packet-routing` statement at `[edit protocols isis backup-spf-options]` hierarchy level to allow the backup paths to be available for inet.0 routing table along with inet.3 routing table.

Options

maximum-backup-paths

Set the maximum number of equal-cost post-convergence backup paths to be installed. For remote LFA, the IS-IS protocol installs only one backup path. The maximum backup path range allows you to install multiple backup paths

- **Default:** 1
- **Range:** `1-8`

maximum-labels

Set the maximum number of labels used to construct a post-convergence backup path. If the backup path for a particular prefix requires more number of labels than the configured maximum labels then the backup path for that particular prefix is not installed.

> **NOTE**: If the `maximum-labels` option is not configured, then the number of labels that gets pushed stops at 3.

- **Default:** 3
- **Range:** `2-5`

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 17.4R1 for MX Series, PTX Series, and QFX Series.

RELATED DOCUMENTATION

Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS

Configuring Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS

post-convergence-lfa | **12903**

use-for-post-convergence-lfa | **30517**

node-protection | **11651**

use-post-convergence-lfa (Protocols OSPF)

IN THIS SECTION

- Syntax | **30533**
- Hierarchy Level | **30534**
- Description | **30534**
- Options | **30534**
- Required Privilege Level | **30534**
- Release Information | **30535**

Syntax

```
use-post-convergence-lfa <maximum-backup-paths maximum-backup-paths> <maximum-labels maximum-labels> ;
```

Hierarchy Level

```
[edit logical-systems name protocols ospf backup-spf-options],
[edit logical-systems name routing-instances name protocols ospf backup-spf-options],
[edit protocols ospf backup-spf-options],
[edit routing-instances name protocols ospf backup-spf-options]
```

Description

Calculate post-convergence MPLS fast reroute (FRR) backup next hops for the OSPF protocol using segment routing (SR). Junos OS allows you to control the maximum number of equal-cost multipath (ECMP) backup paths installed for a given destination. Junos OS also allows you to control the maximum number of labels in the installed backup paths. Configure the `use-source-packet-routing` statement at `[edit protocols ospf backup-spf-options]` hierarchy level to allow the backup paths to be available for inet.0 routing table along with inet.3 routing table.

Options

maximum-backup-paths Set the maximum number of equal-cost post-convergence backup paths to be installed.

- **Default:** 1
- **Range:** `1-8`

maximum-labels Set the maximum number of labels used to construct a post-convergence backup path. If the backup path for a particular prefix requires more labels than the configured maximum labels, then the backup path for that particular prefix is not installed.

> **NOTE**: If the `maximum-labels` option is not configured, then the maximum number of labels used to construct a post-convergence backup path is 3.

- **Default:** 3
- **Range:** `2-5`

Required Privilege Level

routing

Release Information

Statement introduced in Junos OS Release 18.2R1 for MX Series, PTX Series, and QFX Series.

RELATED DOCUMENTATION

post-convergence-lfa (Protocols OSPF) | 12904

Topology-Independent Loop-Free Alternate with Segment Routing for OSPF

use-primary (DHCP Local Server)

IN THIS SECTION

- Syntax | 30535
- Hierarchy Level | 30535
- Description | 30536
- Options | 30536
- Required Privilege Level | 30536
- Release Information | 30536

Syntax

```
use-primary primary-profile-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server group group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name system services dhcp-local-server dynamic-profile
profile-name],
```

```
[edit logical-systems logical-system-name system services dhcp-local-server group group-name
dynamic-profile profile-name],
[edit routing-instances routing-instance-name system services dhcp-local-server dynamic-profile
profile-name],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name
dynamic-profile profile-name],
[edit system services dhcp-local-server dynamic-profile profile-name],
[edit system services dhcp-local-server group group-name dynamic-profile profile-name]
```

Description

Specify the dynamic profile to configure as the primary dynamic profile. The primary dynamic profile is instantiated when the first subscriber or DHCP client logs in. Subsequent subscribers (or clients) are not assigned the primary dynamic profile; instead, they are assigned the dynamic profile specified for the interface. When the first subscriber (or client) logs out, the next subscriber (or client) that logs in is assigned the primary dynamic profile.

Options

primary-profile-name—Name of the dynamic profile to configure as the primary dynamic profile

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

use-primary (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 30537
- Hierarchy Level | 30537
- Description | 30538
- Options | 30538
- Required Privilege Level | 30538
- Release Information | 30538

Syntax

```
use-primary primary-profile-name;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dhcpv6 group  group-name dynamic-profile profile-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name dynamic-profile profile-
name],
[edit forwarding-options dhcp-relay group  group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 dynamic-profile
profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dynamic-profile profile-
name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
dynamic-profile profile-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name dynamic-
profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dynamic-profile profile-name],
```

```
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name dynamic-profile profile-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name dynamic-profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 dynamic-
profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dynamic-profile
profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group  group-
name dynamic-profile profile-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group  group-name
dynamic-profile profile-name]
```

Description

Specify the dynamic profile to configure as the primary dynamic profile. The primary dynamic profile is instantiated when the first subscriber logs in. Subsequent subscribers are not assigned the primary dynamic profile; instead, they are assigned the dynamic profile specified for the interface. When the first subscriber logs out, the next subscriber that logs in is assigned the primary dynamic profile.

Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

EX Series switches do not support DHCPv6.

Options

primary-profile-name—Name of the dynamic profile to configure as the primary dynamic profile

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

use-string

IN THIS SECTION

- Syntax | 30539
- For Platforms with Enhanced Layer 2 Software (ELS) | 30539
- For Platforms Without ELS | 30539
- For MX Series Platforms | 30540
- Description | 30540
- Options | 30540
- Required Privilege Level | 30540
- Release Information | 30540

Syntax

```
use-string string;
```

For Platforms with Enhanced Layer 2 Software (ELS)

```
[edit vlans vlan-name forwarding-options dhcp-security option-82 remote-id]
```

For Platforms Without ELS

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-
name) dhcp-option82 remote-id],
[edit forwarding-options helpers bootp dhcp-option82 remote-id] ,
```

```
[edit forwarding-options helpers bootp interface interface-name dhcp-option82
remote-id]
```

For MX Series Platforms

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security option-82 circuit-id]
```

Description

Use a string rather than the MAC address of the host system (the default) in the remote ID value in the DHCP option 82 information.

Options

string—Character string used as the remote ID value.

Range: 1–255 characters

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

Hierarchy level `[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]` introduced in Junos OS Release 14.1 for the MX Series.

RELATED DOCUMENTATION

Configuring DHCP Option 82 on a Router with Bridge Domain

Example: Setting Up DHCP Option 82

http://tools.ietf.org/html/rfc3046

use-transport-class

IN THIS SECTION

- Syntax | 30541
- Hierarchy Level | 30541
- Release Information | 30542
- Description | 30542
- Options | 30542
- Required Privilege Level | 30542

Syntax

```
use-transport-class {
fib-install;
inet3-install;
}
```

Hierarchy Level

```
[edit logical-systems name protocols source-packet-routing],
[edit logical-systems name routing-instances name routing-options flex-algorithm],
[edit logical-systems name routing-options flex-algorithm],
[edit logical-systems name tenants name routing-instances name routing-options flex-algorithm],
[edit protocols source-packet-routing],
[edit routing-instances name routing-options flex-algorithm],
[edit routing-options flex-algorithm],
[edit tenants name routing-instances name routing-options flex-algorithm]
```

Release Information

Statement introduced in Junos OS Release 21.2R1.

Description

Enable transport class on SR-TE colored policies to Install ingress routes via transport class RIBs

Options

fib-install	Install SR-TE routes in FIB
inet3-install	Install SR-TE routes in inet{6}.3 table

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

BGP Classful Transport Planes Overview

Color-Based Traffic Engineering Configuration

use-underlying-interface-mac

IN THIS SECTION

- Syntax | 30543
- Hierarchy Level | 30543
- Description | 30543
- Required Privilege Level | 30543
- Release Information | 30543

Syntax

```
use-underlying-interface-mac
```

Hierarchy Level

```
[edit system demux-options]
```

Description

Configure the software to use the same link-local address for both interfaces when you are using Router Advertisement for IPv6 subscribers on dynamic demux interfaces that run over underlying static demux interfaces. In this case, the link-local address for the underlying interface should be based the MAC address of the underlying interface.

This statement causes the system to assign an address using the 64-bit Extended Unique Identifier (EUI-64) as described in RFC 2373.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

Dual Stack for PPPoE Access Networks Using NDRA

use-vlan-id

IN THIS SECTION

- Syntax | 30544
- For Platforms with Enhanced Layer 2 Software (ELS) | 30544
- For MX Series Platforms | 30544
- Description | 30544
- Required Privilege Level | 30545
- Release Information | 30545

Syntax

```
use-vlan-id;
```

For Platforms with Enhanced Layer 2 Software (ELS)

```
[edit forwarding-options helpers bootp dhcp-option82-circuit-id]
[edit forwarding-options helpers bootp interface  interface-name dhcp-option82-circuit-id]
```

For MX Series Platforms

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security option-82 circuit-id]
```

Description

Use the VLAN ID rather than the VLAN name (the default) in the circuit ID or remote ID value in the DHCP option 82 information.

NOTE: The `use-vlan-id` statement is mutually exclusive with the `use-interface-description` and `no-vlan-interface-name` statements.

The `use-vlan-id` statement only applies to interfaces in a bridge domain. The format of the Agent Circuit ID or Agent Remote ID information for Fast Ethernet or Gigabit Ethernet interfaces is as follows:

```
(fe | ge)-fpc/pic/port.subunit:svlan_id-vlan_id
```

NOTE: The *subunit* is required and used to differentiate the interface for remote systems, and *svlan_id-vlan_id* represents the VLANs associated with the bridge domain.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

Hierarchy level `[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]` introduced in Junos OS Release 14.1 for the MX Series.

NOTE: The EX Series switches that support the `use-vlan-id` statement are the EX4300, EX4600, and EX9200 switches.

RELATED DOCUMENTATION

- Example: Setting Up DHCP Option 82 Using the Same VLAN
- Example: Setting Up DHCP Option 82
- http://tools.ietf.org/html/rfc3046.

use-vlan-id (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 30546
- Hierarchy Level | 30546
- Description | 30547
- Required Privilege Level | 30547
- Release Information | 30547

Syntax

```
use-vlan-id;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 ...],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dual-stack-group dual-
stack-group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dual-stack-group dual-stack-group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dual-stack-group ]
```

Description

Specify that the VLAN is identified by the VLAN ID, rather than the VLAN name, when you configure the DHCPv6 relay agent to include either of the following in packets it sends to a DHCPv6 server:

- DHCPv6 Interface-ID (option 18)
- DHCPv6 Remote-ID (option 37)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

DHCPv6 Relay Agent

Extended DHCP Relay Agent Overview

user (Access)

IN THIS SECTION

- Syntax | 30548
- Hierarchy Level | 30548
- Description | 30548
- Options | 30548
- Required Privilege Level | 30550
- Release Information | 30550

Syntax

```
user user-name {
    authentication {
        (encrypted-password encrypted-password | plain-text-password);
        load-key-file filename;
        no-public-keys;
        ssh-ecdsa name {
            from host-list;
        }
        ssh-ed25519 name {
            from host-list;
        }
        ssh-rsa name {
            from host-list;
        }
    }
    class class-name;
    cli {
        prompt prompt;
    }
    full-name complete-name;
    uid uid;
}
```

Hierarchy Level

```
[edit system login]
```

Description

Configure access permission for individual users. Starting in Junos OS Release 18.3, the `ssh-dsa` hostkey algorithm is deprecated— rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration.

Options

authentication Specify one or more authentication methods that a user can use to log in to the router or switch. You can assign multiple authentication methods to a single user.

encrypted-password

Message Digest 5 (MD5) or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.

> **CAUTION**: Do not use the `encrypted-password` option unless the password is *already* encrypted, and you are entering the encrypted version of the password.
>
> If you accidentally configure the `encrypted-password` statement with a plain-text password or with blank quotation marks (" "), you will not be able to log in to the device as this user.

- **Range:** You cannot configure a blank password for `encrypted-password` using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

plain-text-password

Specify a plain text password for the user through a password prompt. The plain text password will be automatically encrypted as soon as you configure it. Plain-text passwords are hidden and marked as `## SECRET-DATA` in the configuration.

load-key-file

Load SSH public keys from a previously-generated named file at a specified URL location or local path. The file contains one or more SSH keys that are copied into the configuration when the command is issued.

no-public-keys

Disables ssh public key authentication for the user specified. If the no-public-keys statement is specified at the [edit system services ssh] hierarchy level, public key authentication is disabled for all users on the device.

ssh-ecdsa *public-key*

SSH version 2 authentication. Specify the ECDSA public key. You can specify one or more public keys for each user.

from *host-list* Specify a pattern-list of allowed hosts.

ssh-ed25519 *public-key*

SSH version 2 authentication. Specify the ED25519 public key. You can specify one or more public keys for each user.

from *host-list* Specify a pattern-list of allowed hosts.

ssh-rsa *public-key* — SSH version 2 authentication. Specify the RSA public key. You can specify one or more public keys for each user.

- **from** *host-list* — Specify a pattern-list of allowed hosts.

class *class-name* — Assign a user to a login class. You must assign each user to a login class. Specify one of the classes defined at the `[edit system login class]` hierarchy level.

cli — Set the CLI prompt specified for a specified login user or specified login class. The prompt set for the login user has precedence.

prompt *prompt* — Specify the prompt string you want to see displayed in the CLI prompt.

full-name *complete-name* — Specify the user's complete name. If the name contains spaces, enclose it in quotation marks. Do not include colons or commas.

uid *uid-value* — Numeric identifier associated with the user account, either assigned by an administrator or assigned automatically when you commit the user configuration. It is used by applications that request numeric identifiers, such as some RADIUS queries, or secure applications, such as flow-tap monitoring. This value must be unique on the router or switch.

- **Default:** If you do not assign a UID to a user, the software assigns one when you commit the configuration, preferring the lowest available number.
- **Range:** 100 through 64000

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement `no-public-keys` introduced in Junos OS Release 15.1.

Statement `cli` introduced in Junos OS 17.3.

RELATED DOCUMENTATION

Configuring Junos OS User Accounts by Using a Configuration Group

set cli prompt | 16877

class (Defining Login Classes) | 3347

root-authentication

user-defined-option-82

IN THIS SECTION

- Syntax | 30551
- Hierarchy Level | 30551
- Description | 30552
- Required Privilege Level | 30552
- Release Information | 30552

Syntax

```
user-defined-option-82 string;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides user-
defined-option82 string],
[edit forwarding-options dhcp-relay group group-name overrides user-defined-option82 string],
[edit forwarding-options dhcp-relay overrides user-defined-option82 string]
```

Description

Configure a custom text string to use as the interface description in the DHCP option 82 Agent Circuit ID (suboption 1) information. This text string is defined independently of the interface description that is configured using the `description` statement at the `[edit interfaces interface-name]` hierarchy level.

The custom text string is configured using the `user-defined-option-82` statement at the following hierarchy levels:

- To configure a custom string on an interface level:

```
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides user-
defined-option82 string]
```

- To configure a custom string at the group level:

```
[edit forwarding-options dhcp-relay group group-name overrides user-defined-option82 string]
```

- To configure a custom string globally:

```
[edit forwarding-options dhcp-relay overrides user-defined-option82 string]
```

You can define a custom string up to 251 characters in length. To include the custom string in the DHCP option 82 Agent Circuit ID, you must configure the `user-defined` statement at the `[edit forwarding-options dhcp-relay group group-name relay-option-82 circuit-id]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
21.2R1	Starting with Junos OS Release 21.2R1, QFX Series switches support the `user-defined-option-82` CLI statement in a stateless DHCP relay configuration. You can configure stateless DHCP relay using the `forward-only` CLI statement at the [`edit forwarding-options dhcp-relay`] hierarchy level.

RELATED DOCUMENTATION

Including a Textual Description in DHCP Options

Using DHCP Relay Agent Option 82 Information

user-firewall

IN THIS SECTION

- Syntax | 30553
- Hierarchy Level | 30554
- Description | 30554
- Options | 30554
- Required Privilege Level | 30554
- Release Information | 30554

Syntax

```
user-firewall {
    access-profile profile-name;
    domain domain-name
    ssl-termination-profile profile-name;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication]
```

Description

Configure user role firewall authentication, and map the source IP address to the username and its associated roles (groups). The mapped data is written to the firewall authentication table for later retrieval by the user role firewall. The user role firewall uses the username and role information to determine whether to permit or deny a user's session or traffic.

Options

access-profile *profile-name*	Specify the name of the access profile to be used for authentication.
domain *domain-name*	Specify the name of the domain where firewall authentication occurs in the event that the Windows Management Instrumentation client (WMIC) is not available to get IP-to-user mapping for the integrated user firewall feature. The maximum length is 65 bytes.
ssl-termination-profile *profile-name*	For HTTPS traffic, specify the name of the SSL termination profile used for SSL offloading.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10. Support for the **domain** keyword added in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

- Overview of Integrated User Firewall
- Using Firewall Authentication as an Alternative to WMIC

Understanding User Role Firewalls

Example: Configuring a User Role Firewall on an SRX Series Device

user-group-mapping

IN THIS SECTION

- Syntax | 30555
- Hierarchy Level | 30556
- Description | 30556
- Options | 30556
- Required Privilege Level | 30557
- Release Information | 30557

Syntax

```
user-group-mapping {
    ldap {
        address ip-address {
            port port;
        }
        authentication-algorithm {
            simple;
        }
        base base;
        ssl;
        user username {
            password password;
        }
    }
}
```

Hierarchy Level

```
[edit services user-identification active-directory-access domain]
```

Description

Configure the SRX Series Firewall to connect to an LDAP server, so that the server can provide the SRX Series with user-to-group mappings. These mappings are used to implement the integrated user firewall feature. The domain controller acts as the LDAP server in typical customer scenarios.

Most of this statement is optional, because the default communication method is LDAP and most arguments have default values. Only the LDAP keyword and the base are required.

Options

ldap Required. LDAP is the protocol used to access the LDAP server to get user-to-group mappings.

- **address *ip-address*** Optional. Specify the IP address of the LDAP server. If no address is specified, the system uses one of the configured Active Directory domain controllers.
  - **port *port*** Optional. Specify the port number of the LDAP server. If no port number is specified, the system uses port 389 for plaintext or port 636 for encrypted text.
- **authentication-algorithm** Optional. Specify the algorithm used while the SRX Series communicates with the LDAP server.
  - **simple** Configure simple (plaintext) authentication method.
- **base *base*** Required. LDAP base distinguished name (DN).
- **ssl** Optional. Enable Secure Sockets Layer (SSL) to ensure secure transmission with the LDAP server. Disabled by default, which means that the password is sent in plaintext.
- **user *username*** Optional. Username of the LDAP account. If no username is specified, the system will use the configured domain controller's username.

password *password*	Optional. Specify the password for the account. If no password is specified, the system uses the configured domain controller's password.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

- active-directory-access | 1802
- clear services user-identification active-directory-access | 4194
- show services user-identification active-directory-access statistics | 26389
- show services user-identification active-directory-access user-group-mapping | 26393
- traceoptions (Active Directory Access) | 29483
- user-identification (Services) | 30560
- *LDAP Functionality in Integrated User Firewall*

user-id

IN THIS SECTION

- Syntax | 30558
- Hierarchy Level | 30558
- Description | 30558
- Required Privilege Level | 30558

- Release Information | 30558

Syntax

```
user-id {ascii ascii hexadecimal hexadecimal};
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcp-client client-
identifier]
```

Description

Specify an ASCII or hexadecimal user ID for the Dynamic Host Configuration Protocol (DHCP) client.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

user-identification

IN THIS SECTION

- Syntax | 30559
- Hierarchy Level | 30559

- Description | 30559
- Options | 30560
- Required Privilege Level | 30560
- Release Information | 30560

Syntax

```
user-identification {
    authentication-source {
        firewall-authentication (disable | priority priority);
        local-authentication-table (disable | priority priority);
        unified-access-control (disable | priority priority);
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

Identifies one or more tables to be used as the source for user role information.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1. Statement updated in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding User Role Firewalls

Understanding the User Identification Table

user-identification (Services)

IN THIS SECTION

- Syntax | 30561
- Hierarchy Level | 30566
- Description | 30566
- Options | 30566
- Required Privilege Level | 30567
- Release Information | 30567

Syntax

```
user-identification {
    active-directory-access {
        domain domain-name {
            user username;
                password password;
            domain-controller domain-controller-name {
                address domain-controller-address;
            }
            ip-user-mapping {
                discovery-method {
                    wmi {
                        event-log-scanning-interval seconds;
                        initial-event-log-timespan hours;
                    }
                }
            }
            user-group-mapping {
                ldap {
                    address ip-address {
                        port port;
                    }
                    authentication-algorithm {
                        simple;
                    }
                    base base;
                    ssl;
                    user username {
                        password password;
                    }
                }
            }
        }
        authentication-entry-timeout minutes;
        filter {
            include address;
            exclude address;
        }
        no-on-demand-probe;
        wmi-timeout seconds;
        traceoptions {
```

```
            file file;
            flag {
                active-directory-authentication;
                all;
                configuration;
                db;
                ip-user-mapping;
                ip-user-probe;
                ipc;
                user-group-mapping;
                wmic;
            }
            level {
                all;
                error;
                info;
                notice;
                verbose;
                warning;
            }
            no-remote-trace;
        }
    device-information {
        authentication-source;
        end-user-profile profile-name profile-name {
            domain-name domain-name {
                attribute device-category {
                    string string-value;
                }
                attribute device-identity {
                    string string-value;
                }
                attribute device-vendor {
                    string string-value;
                }
                attribute device-type {
                    string string-value;
                }
                attribute device-os {
                    string string-value;
                }
                attribute device-os-version {
                    string string-value;
```

```
                }
            }
        }
        identity-context-profile profile-name profile-name {
            domain-name domain-name {
                attribute device-category {
                    string string-value;
                }
                attribute device-identity {
                    string string-value;
                }
                attribute device-os {
                    string string-value;
                }
                attribute device-os-version {
                    string string-value;
                }
                attribute device-type {
                    string string-value;
                }
                attribute device-vendor {
                    string string-value;
                }
            }
        }
        traceoptions;
    }
    identity-management {
        authentication-entry-timeout minutes;
        batch-query {
            items-per-batch items-per-batch;
            query-interval seconds;
        }
        connection {
            connect-method (http | https);
            port port;
            primary {
                address address;
                ca-profile ca-profile;
                client-id client-id;
                client-secret client-secret;
                interface interface-name;
                routing-instance routing-instance -name;
```

```
            source source-address;
        }
        query-api query-api;
        secondary {
            address address;
            ca-profile ca-profile;
            client-id client-id;
            client-secret client-secret;
            interface interface-name;
            routing-instance routing-instance -name;
            source source-address;
        }
        server-validity-check-interval;
        token-api token-api;
    }
    filter {
        domain name;
        exclude-ip {
            address-book book-name;
            address-set address-set;
        }
        include-ip {
            address-book book-name;
            address-set address-set;
        }
    }
    invalid-authentication-entry-timeout minutes;
    ip-query {
        no-ip-query;
        query-delay-time seconds;
    }
    jims-validator {
        port port;
    }
    session-batch-query {
        items-per-batch;
        query-interval seconds;
    }
    session-query {
        no-session-query;
        query-delay-time seconds;
    }
    traceoptions {
```

```
            file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
            flag name;
            level (all | error | info | notice | verbose | warning);
            no-remote-trace;
        }
    }
        logical-domain-identity-management {
            active {
                authentication-entry-timeout minutes;
                filter {
                    domain name;
                    exclude-ip {
                        address-book book-name;
                        address-set address-set;
                    }
                    include-ip {
                        address-book book-name;
                        address-set address-set;
                    }
                }
                invalid-authentication-entry-timeout minutes;
                ip-query {
                    query-delay-time seconds;
                }
                query-server name {
                    batch-query {
                        items-per-batch items-per-batch;
                        query-interval seconds;
                    }
                    connection {
                        connect-method (http | https);
                        port port;
                        primary {
                            address address;
                            ca-certificate ca-certificate;
                            client-id client-id;
                            client-secret client-secret;
                        }
                        query-api query-api;
                        secondary {
                            address address;
                            ca-certificate ca-certificate;
```

```
                            client-id client-id;
                            client-secret client-secret;
                        }
                        token-api token-api;
                    }
                }
            }
            traceoptions {
                file <filename> <files files> <match match> <size size> <(world-readable | no-
world-readable)>;
                flag name;
                level (all | error | info | notice | verbose | warning);
                no-remote-trace;
            }
        }
    }
```

Hierarchy Level

```
[edit services]
```

Description

Configure the integrated user firewall feature, including access to the Active Directory domain and domain controller, IP address-to-user mapping, and user-to-group mapping. One or two Active Directories are allowed under one domain. The IP address-to-user mapping and user-to-group mapping are configured per domain.

Options

authentication-entry-timeout *minutes*

Timeout interval starting from the Active Directory/domain controller login time, the last active session, or the last successful probe. A setting of 0 means the authentication does not need a timeout. We recommend that you configure a setting of 0 when you disable on-demand-probe to prevent someone from accessing the Internet without logging in again.

- **Range:** 10 through 1440 minutes
- **Default:** 30 minutes

filter Optional. Range of IP addresses that needs to be monitored or not monitored.

- **include *address*** Include IP address or range. Maximum of 20 addresses.
- **exclude *address*** Exclude IP address or range. Maximum of 20 addresses.

no-on-demand-probe Do not use traffic to discover user. Default is disabled.

wmi-timeout *seconds* (Optional) Configures the number of seconds that the domain PC has to respond to the SRX Series device's query through WMI/DCOM.

- If the PC responds within that timeframe to the WMI query, the SRX creates an authentication entry for this PC.
- If the PC does not respond within that timeframe, the WMI query failed. In the case of a failed query, if the SRX had an authentication entry about the queried PC before the WMI query, that authentication entry is deleted. If the SRX had no authentication entry before the WMI query, the SRX does not create an authentication entry.
- **Range:** 3 through 120 seconds
- **Default:** 10 seconds

identity-context-profile Configure identity context profile in firewall policy to control network access.

identity-management Configure identity management to collect identity information.

logical-domain-identity-management Configures the logical domain identity management.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D10.

`logical-domain-identity-management` option introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

active-directory-access | 1802

traceoptions (Active Directory Access) | 29483

user-name

IN THIS SECTION

- Syntax | 30568
- Hierarchy Level | 30568
- Description | 30568
- Required Privilege Level | 30569
- Release Information | 30569

Syntax

```
user-name {
    equals value;
    has-prefix value;
    has-suffix value;
    matches value;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from]
```

Description

Specify the RADIUS AVP User-Name for the incoming RADIUS request from the subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers*
- *Understanding Selection of Properties for an IP-Based TDF Subscriber*
- *Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber*
- *IP-Based Subscriber Setup Overview*

user-name (Event Policy)

IN THIS SECTION

- Syntax | 30570
- Hierarchy Level | 30570
- Description | 30570
- Default | 30570
- Options | 30570
- Required Privilege Level | 30571
- Release Information | 30571

Syntax

```
user-name username;
```

Hierarchy Level

```
[edit event-options policy policy-name then change-configuration],
[edit event-options policy policy-name then event-script filename],
[edit event-options policy policy-name then execute-commands],
[edit event-options policy policy-name then upload filename (filename | committed)
destination destination-name]
```

Description

Associate a user with an action in an event policy. The event policy action is executed under the privileges of the associated user.

> **NOTE**: When you configure the `user-name` statement for event scripts, it only applies to SLAX and XSLT scripts. This statement has no effect when configured for Python event scripts.
>
> To execute a local Python event script under the access privileges of a specific user, configure the `python-script-user` *`username`* statement at the `[edit event-options event-script file` *`filename`*`]` hierarchy level. If you do not configure the `python-script-user` statement for a Python event script, Junos OS executes the script using the access privileges of the generic, unprivileged system account `nobody`.

Default

If you do not associate a user with an action, the action is executed as user `root`.

Options

username Username that is configured at the `[edit system login]` hierarchy level.

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support at the `[edit event-options policy policy-name then change-configuration]` hierarchy level introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

> *Change the User Privilege Level for an Event Policy Action*
>
> python-script-user (Event and SNMP Scripts) | **13884**

user-name-include (OCS Partition)

IN THIS SECTION

- Syntax | **30571**
- Hierarchy Level | **30572**
- Description | **30572**
- Options | **30572**
- Required Privilege Level | **30573**
- Release Information | **30573**

Syntax

```
user-name-include{
    base-interface-name;
    delimiter delimiter-character;
    domain-name my-domain;
    interface-name;
```

```
    mac-address;
    nas-port-id;
    origin-host;
    origin-realm;
    user-name;
    user-prefix pref;
}
```

Hierarchy Level

```
[edit access ocs partition partition-name]
```

Description

Configure the username identification to be used in an Online Charging Function (OCS) partition. You can configure a data string to include the following options.

Options

base-interface-name

Use the physical or underlying interface name. If both the underlying interface and the client application perform authentication, authorization, and provisioning, the identification attributes in the server requests enable the AAA or PCRF server to make an association between the two entities.

delimiter *delimiter-character*

Use the specified character as the delimiter between the concatenated components of the `user-name-include` statement. You cannot use the semicolon (;) as a delimiter.

- **Default:** @

domain-name *my-domain*

Use the specified domain name that is concatenated with the `user-name-include` statement during the subscriber identification process.

interface-name

Use the interface name that is concatenated with the `user-name-include` statement during the subscriber identification process; for example, demux0.

mac-address

Use the client hardware address (chaddr) from the incoming packet that is concatenated with the `user-name-include` statement during the subscriber identification process.

nas-port-id Use the NAS-Port-ID (RADIUS attribute 87), which identifies the physical interface that subscriber management uses to identify users that is concatenated with the `user-name-include` statement during the subscriber identification process. By default, the NAS-Port-ID includes the interface-description value that describes the physical interface.

origin-host Use the name of the host that originates the Diameter message that is concatenated with the `user-name-include` statement during the subscriber identification process. Supplied as the value of Origin-Host AVP for all messages sent by the Diameter master instance.

origin-realm Use the realm of the host that originates the Diameter message that is concatenated with the `user-name-include` statement during the subscriber identification process. Supplied as the value of Origin-Realm AVP for all messages sent by the Diameter master instance.

user-name (Included in the `user-name-include` statement by default) Use the username that is concatenated with the `user-name-include` statement during the subscriber identification process.

user-prefix *prefix* Use the specified user prefix that is concatenated with the `user-name-include` statement during the subscriber identification process.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.2.

RELATED DOCUMENTATION

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

3GPP Policy and Charging Control for Wireline Provisioning and Accounting

user-password (PCEF Profile)

IN THIS SECTION

- Syntax | 30574
- Hierarchy Level | 30574
- Description | 30574
- Options | 30574
- Required Privilege Level | 30574
- Release Information | 30575

Syntax

```
user-password password;
```

Hierarchy Level

```
[edit unified-edge  pcef profiles profile-name aaa-policy-control]
```

Description

Configure the user password for subscribers assigned to the parent PCEF profile.

Options

password Password for access requests to the RADIUS server.

- **Range:** 1 through 32 characters

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Policies That a RADIUS Server Controls

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

user-planes

IN THIS SECTION

- Syntax | 30575
- Hierarchy Level | 30576
- Description | 30576
- Options | 30576
- Required Privilege Level | 30576
- Release Information | 30576

Syntax

```
user-planes name {
 colocated-user-plane {
  ip-endpoint-address ip-endpoint-address;
 }
 data-network-name data-network-name;
 ip-address ip-address;
```

```
  routing-instance routing-instance;
 }
```

Hierarchy Level

```
[edit services agf]
```

Description

Configure a tunnel to the User Plane Function (UPF).

Options

name	User plane name
colocated-user-plane	Specify there is a colocated UPF. The `ip-endpoint-address` *ip-endpoint-address* is the IP address of the UPF.
data-network-name	Data network name Name of the data network
ip-address *ip-endpoint-address*	User plane IP address
routing-instance	User plane routing instance

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3R1.

user-prefix

IN THIS SECTION

- Syntax | 30577
- Hierarchy Level | 30577
- Description | 30577
- Options | 30577
- Required Privilege Level | 30577
- Release Information | 30578

Syntax

```
user-prefix user-prefix-string;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication username-include]
```

Description

Specify the user prefix that is concatenated with the username during the subscriber authentication process.

Options

user-prefix-string—The user prefix string.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

user-prefix (DHCP Local Server)

IN THIS SECTION

- Syntax | 30578
- Hierarchy Level | 30578
- Description | 30579
- Options | 30579
- Required Privilege Level | 30580
- Release Information | 30580

Syntax

```
user-prefix user-prefix-string;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 authentication ],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
```

```
services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server  authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6  authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6  group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Description

Specify the user prefix that is concatenated with the username during the subscriber authentication or DHCP client authentication process.

Options

user-prefix-string—User prefix string.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

RELATED DOCUMENTATION

Specifying Authentication Support

user-prefix (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 30580
- Hierarchy Level | 30581
- Description | 30581
- Options | 30582
- Required Privilege Level | 30582
- Release Information | 30582

Syntax

```
user-prefix user-prefix-string;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit forwarding-options dhcp-relay dual-stack-grouop dual-stack-group-name authentication
username-include],
[edit forwarding-options dhcp-relay group group-name authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay authentication username-
include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 authentication
username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name
authentication username-include],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name
authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay group group-name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay authentication
username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6
authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-
name authentication username-include],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name
authentication username-include]
```

Description

Specify the user prefix that is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the `[edit ... dhcpv6]` hierarchy levels to configure DHCPv6 support.

Options

user-prefix-string—User prefix string.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group dual-stack-group-name]` hierarchy level introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

DHCP with External Authentication Server

user-prefix (Static Subscribers)

IN THIS SECTION

- Syntax | 30583
- Hierarchy Level | 30583
- Description | 30583
- Options | 30583
- Required Privilege Level | 30583
- Release Information | 30584

Syntax

```
user-prefix user-prefix-string;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication username-include],
[edit logical-systems logical-system-name system services static-subscribers  authentication
username-include],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers
authentication username-include],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication username-include],
[edit system services static-subscribers authentication username-include],
[edit system services static-subscribers group group-name authentication username-include]
```

Description

Specify that a string is included as the beginning of the username created for all static subscribers or for the static subscribers in a specified group. The group version of the statement takes precedence over the global version. The username is also sent to RADIUS in the Access-Request message.

Options

user-prefix-string—String that begins the username. The string can include the following characters: a through z, A through Z, 0 through 9, “-“, or “.”.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

JSRC and Subscribers on Static Interfaces

JSRC and Subscribers on Static Interfaces

JSRC and Subscribers on Static Interfaces

user-query (Services User Identification)

IN THIS SECTION

- Syntax | 30584
- Hierarchy Level | 30585
- Description | 30585
- Options | 30586
- Required Privilege Level | 30588
- Release Information | 30588

Syntax

```
user-query {
    ca-certificate ca-certificate;
    client-id client-id;
    client-secret client-secret;
    delay-query-time seconds;
    query-api query-api;
    token-api token-api;
    web-server (Services) {
        server-name;
        address address;
```

```
        connect-method (http | https);
        port port;
    }
}
```

Hierarchy Level

```
[edit services user-identification authentication-source (Services User Identification ClearPass)
aruba-clearpass]
```

Description

ca-certificate- Configures the Integrated ClearPass Authentication and Enforcement feature user query function configuration. User query enables the SRX Series Firewall to query the ClearPass Policy Manager (CPPM) for authentication and identity information for an individual user under certain circumstance when it does not receive that information from the CPPM though the Web API POST requests.

client-id- Configures the client ID that the SRX Series Firewall requires to obtain an access token for the Integrated ClearPass Authentication and Enforcement user query function. The client ID must be consistent with the API client configured on the CPPM.

client-secret- Configures the client secret used with the client ID that the SRX Series Firewall requires to obtain an access token for the Integrated ClearPass Authentication and Enforcement user query function. The client secret must be consistent with the client secret configured on the CPPM.

delay-query-time- If the CPPM does not send to the SRX Series Firewall authentication and identity information for a particular user, then the SRX Series Firewall can request that information for the user if you configure the user query function.

query-api - Configure `query-api` to specify the path of the URL that the SRX Series Firewall uses to query the ClearPass Policy Manager (CPPM) webserver for authentication and identity information for an individual user. For the SRX Series Firewall to be able to make a request, you must have configured it to obtain an access token.

token-api - Configure the token API that is used in generating the URL for acquiring an access token. The token API is combined with the connection method and the IP address of the ClearPass webserver to produce the complete URL used for acquiring an access token.

Options

ca-certificate

Specify the certificate file that the SRX Series Firewall uses to verify the Clearpass server's certificate for the SSL connection that is used for the user query function. As the ClearPass administrator, you must export the server's certificate from the CPPM and import it to the SRX Series Firewall. Afterward, you must configure the ca-certificate path and the certificate filename on the SRX Series Firewall. Here is an example:

```
'/var/tmp/RADIUSServerCertificate.crt'
```

client-id

The ClearPass endpoint API requires use of OAuth (RFC 6749) to authenticate and authorize the SRX Series Firewall access. The SRX Series Firewall uses the Client Credentials grant type access token, which is one of the two types that ClearPass supports.

If it is configured, the user query function allows the SRX Series Firewall to query the CPPM for authentication and identity information about individual users when it does not receive this information from the CPPM through the SRX Series Web API process (webapi).

client-secret

Client secret for OAuth2 grant.

delay-query-time

Delay time to send user query (0~60sec) (seconds). The amount of time for the SRX Series Firewall to delay before sending queries to the Aruba ClearPass Policy Manager (CPPM) for authentication and identity information for individual users.

Delays can occur from when the CPPM initially posts user authentication information to the SRX Series Firewall to when the SRX Series Firewall updates its ClearPass authentication table with that information. In its transit, the user identity information must first pass through the CPPM device's control plane and the control plane of the SRX Series Firewall.

During that period, traffic might arrive at the SRX Series Firewall that is generated by an access request from a user whose authentication and identity information is in transit from the CPPM to the SRX Series Firewall. Rather than allow the SRX Series Firewall to respond automatically by sending a user query request *immediately*, you can set the delay time parameter specifying in seconds how long the SRX Series Firewall should wait before sending the request.

After the delay timeout expires, the SRX Series Firewall sends the query to the CPPM and creates a pending entry for the user in the Routing Engine authentication table. During this period, any arriving traffic matches the default policy whose action on the traffic you can configure.

- **Default:** 15
- **Range:** 0 through 60

query-api The integrated ClearPass authentication and enforcement user query function supplements the Web API process (webapi) by allowing the SRX Series Firewall to obtain from the CPPM authentication information for an individual user whose information does not already exist in the SRX Series ClearPass authentication table.

Consider the following `query-api` example:

```
api/v1/insight/endpoint/ip/$IP$
```

The SRX Series Firewall generates the complete URL for the user query request by combining the query-api string with the connection method (HTTPS) and the CPPM webserver IP address ({$server}).

```
https://{$server}/api/v1/insight/endpoint/ip/$IP$
```

In this example, the SRX Series Firewall replaces the variables with the following values resulting in a specific URL request for the individual user:

```
https://203.0.113.76/api/v1/insight/endpoint/ip/192.0.2.98
```

Under normal circumstances, the ClearPass webserver sends user authentication information to the SRX Series Firewall in POST request messages and the SRX Series Firewall writes that information to its ClearPass authentication table. When the SRX Series Firewall receives an access request from a user, it searches its ClearPass authentication table for an entry for that user.

It can happen that the SRX Series Firewall might not have received authentication for a user from the CPPM because the user has not yet been authenticated by the CPPM. For example, the user might have joined the network through an access layer not on a managed switch or WLAN. When the CPPM receives the user query from the SRX Series Firewall, it authenticates the user and returns the authentication information to the device.

token-api API of acquiring token for OAuth2 authentication.

For example, if the token API is oauth, the connection method is HTTPS, and the IP address of the ClearPass webserver is 192.0.2.199, the complete URL for acquiring an

access token would be https://192.0.2.199/api/oauth. This is a required parameter. There is no default value.

The SRX Series Firewall user query function requires an access token to be able to query the ClearPass webserver. If the user query function is configured, the SRX Series Firewall can request from the ClearPass webserver user authentication and identity information for an individual user.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services	To view this statement in the configuration.
services-control	To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

RELATED DOCUMENTATION

Configure Integrated ClearPass Authentication and Enforcement

user (SNMP v3)

IN THIS SECTION

- Syntax | 30589
- Hierarchy Level | 30589
- Description | 30589
- Options | 30589
- Required Privilege Level | 30589
- Release Information | 30589

Syntax

```
user username;
```

Hierarchy Level

```
[edit snmp v3 usm local-engine],
[edit snmp v3 usm remote-engine engine-id]
```

Description

Specify a user associated with an SNMPv3 group on a local or remote SNMP engine.

Options

username—SNMPv3 user-based security model (USM) username.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

user (System Logging)

IN THIS SECTION

- Syntax | 30590
- Hierarchy Level | 30590
- Description | 30590

- Options | 30590
- Required Privilege Level | 30591
- Release Information | 30591

Syntax

```
user (username | *) {
    facility severity;
    match "regular-expression";
    match-strings string-name;
}
```

Hierarchy Level

```
[edit system syslog]
```

Description

Configure the logging of system messages to user terminals.

Options

***** (the asterisk)—Log messages to the terminal sessions of all users who are currently logged in.

facility—Class of messages to log. To specify multiple classes, include multiple ***facility severity*** statements. For a list of the facilities, see *System Logging Facilities and Message Severity Levels*.

severity—Severity of the messages that belong to the facility specified by the paired ***facility*** name. Messages with severities the specified level and higher are logged. For a list of the severities, see *System Logging Facilities and Message Severity Levels*.

username—Junos OS login name of the user whose terminal session is to receive system log messages. To log messages to more than one user's terminal session, include more than one **user** statement.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Direct System Log Messages to a User Terminal

System Logging Facilities and Message Severity Levels

user (System Services)

IN THIS SECTION

- Syntax | 30591
- Hierarchy Level | 30592
- Description | 30592
- Options | 30592
- Required Privilege Level | 30592
- Release Information | 30592

Syntax

```
user {
    user-name;
    password password;
}
```

Hierarchy Level

```
[edit system services webapi (System Services)]
```

Description

Configure the Web API process (webapi) username and password for the account.

The Web API process, acting as an HTTP server, allows the Aruba ClearPass Policy Manager (CPPM), acting as a client, to send POST request messages to HTTP server. The CPPM, which is the authentication source for this feature, sends to the SRX Series Firewall user authentication and identity information.

Options

user-name Specify the user name for the integrated ClearPass authentication and enforcement feature Web API process (webapi) user.

password Specify the password for the integrated ClearPass authentication and enforcement feature Web API process (webapi) user.

- **Range:** 1 through 128 characters.

Required Privilege Level

1. system—To view this statement in the configuration.
2. system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

username-include (Demux)

IN THIS SECTION

- Syntax | 30593
- Hierarchy Level | 30593
- Description | 30593
- Options | 30594
- Required Privilege Level | 30594
- Release Information | 30594

Syntax

```
username-include {
    auth-server-realm realm-string;
    delimiter delimiter-character;
    domain-name domain-name;
    interface-name;
    source-address;
    user-prefix user-prefix-string;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet auto-configure address-ranges
authentication]
[edit interfaces interface-name unit unit-number demux inet6 auto-configure address-ranges
authentication]
```

Description

Specify the information included in the username created for the demultiplexing (demux) interface options. The remaining statement is explained separately.

Options

auth-server-realm *auth-server-realm-string*	Specify the authentication server name string used for the demultiplexing (demux) interface options.
delimiter *delimiter-character*	Specify the character used as the delimiter between the concatenated components of the username. You cannot use the semicolon (;) as a delimiter. The delimiter defines the boundary between the part of the original username that is kept and the part that is discarded.
domain-name *domain-name-string*	Specify the domain name formatted string that is concatenated with the username during the subscriber authentication process.
interface-name	Append the interface name for the demultiplexing (demux) interface options to the username string used for authentication.
source-address	Set the source address for every user name for the demultiplexing (demux) interface options.
user-prefix *user-prefix-string*	Specify the user prefix that is concatenated with the username for the demultiplexing (demux) interface options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

IP Demultiplexing Interfaces on Packet-Triggered Subscriber Services

username-include (DHCP Local Server)

IN THIS SECTION

- Syntax | 30595
- Hierarchy Level | 30596
- Description | 30596
- Options | 30596
- Required Privilege Level | 30597
- Release Information | 30597

Syntax

```
username-include {
    circuit-type;
    client-id <exclude-headers> <use-automatic-ascii-hex-encoding>;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interfaces-description (device-interface | logical-interface);
    interface-name ;
    logical-system-name;
    mac-address;
    option-60;
    option-82 <circuit-id> <remote-id>;
    relay-agent-interface-id;
    relay-agent-remote-id;
    relay-agent-subscriber-id;
    routing-instance-name;
    user-prefix user-prefix-string;
    vlan-tags;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...],
[edit system services dhcp-local-server authentication],
[edit system services dhcp-local-server group group-name authentication]
[edit system services dhcp-local-server dhcpv6 authentication],
[edit system services dhcp-local-server dhcpv6 group group-name authentication],
[edit system services dhcp-local-server dual-stack-group group-name authentication]
```

Description

Configure the username that the router or switch passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router (or switch) accesses the local authentication service only and does not use external authentication services, such as RADIUS.

The following statements are not supported in the DHCPv6 hierarchy levels:

- **option-60**
- **option-82**

The following statements are supported in the DHCPv6 hierarchy levels only:

- **relay-agent-interface-id**
- **relay-agent-remote-id**
- **relay-agent-subscriber-id**

Options

vlan-tags Include the subscriber session VLAN tags in the username for interactions with an external authority. Both single-tagged and double-tagged VLANs are supported: The tags are added in the format *outer-vlan-tag-inner-vlan-tag*. The outer tag is always included; the inner tag is included for double-tagged VLANs.

Use this option instead of the `interface-name` option when the outer VLAN tag is unique across the system and you do not need the underlying physical interface name to be part of the format.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

`vlan-tags` option added in Junos OS Release 18.3R1 on MX Series routers.

RELATED DOCUMENTATION

Specifying Authentication Support

Creating Unique Usernames for DHCP Clients

username-include (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 30598
- Hierarchy Level | 30598
- Description | 30598
- Options | 30599
- Required Privilege Level | 30599
- Release Information | 30599

Syntax

```
username-include {
    circuit-type;
    client-id <exclude-headers> <use-automatic-ascii-hex-encoding>;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-description (device-interface | logical-interface);
    interface-name;
    logical-system-name;
    mac-address;
    option-60;
    option-82 <circuit-id> <remote-id>;
    relay-agent-interface-id;
    relay-agent-remote-id;
    relay-agent-subscriber-id;
    routing-instance-name;
    user-prefix user-prefix-string;
    vlan-tags;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay authentication],
[edit forwarding-options dhcp-relay group group-name authentication],
[edit forwarding-options dhcp-relay dhcpv6 authentication],
[edit forwarding-options dhcp-relay dhcpv6 group group-name authentication],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name authentication],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Configure the username that the router (or switch) passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router (or switch) accesses the local authentication service only and does not use external

authentication services, such as RADIUS. Use the statement at the `[edit...dhcpv6]` hierarchy levels to configure DHCPv6 support.

The following statements are not supported in the DHCPv6 hierarchy levels:

- **option-60**
- **option-82**

The following statements are supported in the DHCPv6 hierarchy levels only:

- **relay-agent-interface-id**
- **relay-agent-remote-id**
- **relay-agent-subscriber-id**

Options

vlan-tags Include the subscriber session VLAN tags in the username for interactions with an external authority. Both single-tagged and double-tagged VLANs are supported: The tags are added in the format *`outer-vlan-tag-inner-vlan-tag`*. The outer tag is always included; the inner tag is included for double-tagged VLANs.

Use this option instead of the `interface-name` option when the outer VLAN tag is unique across the system and you do not need the underlying physical interface name to be part of the format.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.1.

Support at the `[edit ... dhcpv6]` hierarchy levels introduced in Junos OS Release 11.4.

Support at the `[edit ... dual-stack-group` *`dual-stack-group-name`*`]` hierarchy level introduced in Junos OS Release 15.1.

`vlan-tags` option added in Junos OS Release 18.3R1 on MX Series routers.

RELATED DOCUMENTATION

Creating Unique Usernames for DHCP Clients

Specifying Authentication Support

username-include (Interfaces)

IN THIS SECTION

- Syntax | 30600
- Hierarchy Level | 30601
- Description | 30601
- Options | 30601
- Required Privilege Level | 30601
- Release Information | 30602

Syntax

```
username-include {
    circuit-id;
    circuit-type;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-name;
    mac-address;
    option-18;
    option-37;
    option-82 <circuit-id> <remote-id>;
    radius-realm radius-realm-string;
    remote-id;
    user-prefix user-prefix-string;
```

```
    vlan-tags;
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication]
```

Description

Configure the username that the router passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router accesses the local authentication service only and does not use external authentication services, such as RADIUS.

The username takes the format *user-prefix mac-address circuit-type circuit-id remote-id option-82 interface-name domain-name radius-realm*. By default, each component is separated by a period (.), but you can specify a different delimiter with the `delimiter` statement.

Options

vlan-tags Include the subscriber session VLAN tags in the username for interactions with an external authority. Both single-tagged and double-tagged VLANs are supported: The tags are added in the format *outer-vlan-tag-inner-vlan-tag*. The outer tag is always included; the inner tag is included for double-tagged VLANs.

You can use this option instead of the `interface-name` option when the outer VLAN tag is unique across the system and you do not need the underlying physical interface name to be part of the format.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

`vlan-tags` option added in Junos OS Release 18.3R1 on MX Series routers.

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

Using DHCP Option 82 Suboptions in Authentication Usernames for Autosense VLANs

Using DHCP Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs

Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs

username-include (Local Authentication)

IN THIS SECTION

- Syntax | **30602**
- Hierarchy Level | **30603**
- Description | **30603**
- Options | **30603**
- Required Privilege Level | **30603**
- Release Information | **30604**

Syntax

```
username-include {
    circuit-id;
    delimiter character;
    domain-name name;
    mac-address;
    remote-id;
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces $junos-interface-ifd-name unit $junos-interface-unit ppp-options local-authentication]
```

Description

Configure a local username that authd can use to request authentication for terminated PPP subscribers. This enables the external RADIUS server to pass implementation-specific configuration for successfully authenticated subscribers. Local authentication supports CPEs that do not negotiate authentication protocols in the same dynamic profile as CPEs that use only PAP or CHAP authentication.

The username takes the following format when you use the default delimiter:

```
mac-address.circuit-id.remote-id@domain-name
```

Options

circuit-id Include the agent circuit identifier (ACI) in the local username.

delimiter *character* Specify the character that separates components that make up the concatenated username.

- **Default:** Period (.)

domain-name *name* Specify the domain name that ends the local username created for the subscribers. The username is sent to RADIUS in the Access-Request message. The string can include the following characters: a through z, A through Z, 0 through 9, “-“, or “.”.

mac-address Include the MAC address from the client PDU in the local username.

remote-id Include the agent remote identifier (ARI) in the local username.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

PPP Subscriber Access Networks Overview

username-include (Static Subscribers)

IN THIS SECTION

- Syntax | 30604
- Hierarchy Level | 30605
- Description | 30605
- Options | 30605
- Required Privilege Level | 30605
- Release Information | 30606

Syntax

```
username-include {
    delimiter delimiter-character;
    domain-name domain-name;
    interface;
    logical-system-name;
    routing-instance-name;
    user-prefix user-prefix-string;
    vlan-tags;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers  authentication],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name  authentication],
[edit logical-systems logical-system-name system services static-subscribers  authentication],
[edit logical-systems logical-system-name system services static-subscribers group group-name
authentication],
[edit routing-instances routing-instances-name system services static-subscribers
authentication],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name  authentication],
[edit system services static-subscribers authentication],
[edit system services static-subscribers group group-name authentication]
```

Description

Specify the information included in the username created for all static subscribers or for static subscribers in a specified group. The group version of the statement takes precedence over the global version. The username is also sent to RADIUS in the Access-Request message.

Options

delimiter *delimiter-character*

Specify the delimiter character that separates the different fields configured for the username. You can specify only a single character as the delimiter.

- **Default:** The default delimiter is a period (.).

vlan-tags

Include the VLAN tags (VLAN IDs) for the interface. For single-tagged VLANs, the field is encoded as *outer-tag*. For stacked VLANs, the field is encoded as *outer-tag-inner-tag*. For IP demux interfaces configured for static subscribers, the VLAN tags configured on the underlying interface are used.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

`delimiter` and `vlan-tags` options added in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

JSRC and Subscribers on Static Interfaces

JSRC and Subscribers on Static Interfaces

JSRC and Subscribers on Static Interfaces

username (Security Antivirus)

IN THIS SECTION

- Syntax | 30606
- Hierarchy Level | 30606
- Description | 30607
- Required Privilege Level | 30607
- Release Information | 30607

Syntax

```
username name;
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update proxy]
```

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update proxy]
[edit security utm feature-profile anti-virus sophos-engine pattern-update proxy]
```

Description

Set the username for the proxy server.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 11.2.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Express and Kaspersky Antivirus feature is not supported from Junos OS Release 15.1X49-D10 onwards.

username (Services)

IN THIS SECTION

- Syntax | 30608
- Hierarchy Level | 30608
- Description | 30608
- Options | 30608

Required Privilege Level | 30608
Release Information | 30608

Syntax

```
username user-name;
```

Hierarchy Level

```
[edit services flow-collector transfer-log-archive archive-sites]
```

Description

Specify the username for the transfer log server.

Options

user-name—FTP server username.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Transfer Logs

using-user-password

IN THIS SECTION

- Syntax | 30609
- Hierarchy Level | 30609
- Release Information | 30609
- Description | 30609
- Required Privilege Level | 30610

Syntax

```
using-user-password password;
```

Hierarchy Level

```
[edit access domain map]
```

Release Information

Statement introduced in Junos OS Release 21.3R1.

Description

Enable the use of an override CHAP password. The override CHAP password replaces the CHAP challenge response that the PPP client sends to the server for authentication.

You must configure the override CHAP password using the `override-chap-password` statement at the [`edit access domain map`] hierarchy level.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement in the configuration.

RELATED DOCUMENTATION

Changing the Username and Password to Simplify Off-Chassis Provisioning

usm

IN THIS SECTION

- Syntax | 30610
- Hierarchy Level | 30612
- Description | 30612
- Required Privilege Level | 30612
- Release Information | 30612

Syntax

```
usm {
    local-engine {
        user username {
            authentication-md5 {
                authentication-password authentication-password;
            }
            authentication-none;
            authentication-sha {
                authentication-password authentication-password;
            }
            privacy-aes128 {
                privacy-password privacy-password;
            }
```

```
            privacy-des {
                privacy-password privacy-password;
            }
            privacy-3des {
                privacy-password privacy-password;
            }
            privacy-none {
                privacy-password privacy-password;
            }
        }
        remote-engine engine-id {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-none;
                authentication-sha {
                    authentication-password authentication-password;
                }
                authentication-sha224 {
                    authentication-password authentication-password;
                }
                authentication-sha256 {
                    authentication-password authentication-password;
                }
                authentication-sha348 {
                    authentication-password authentication-password;
                }
                authentication-sha512 {
                    authentication-password authentication-password;
                }
                privacy-aes128 {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-3des {
                    privacy-password privacy-password;
                }
                privacy-none {
                    privacy-password privacy-password;
                }
```

```
            }
        }
    }
}
```

Hierarchy Level

```
[edit snmp v3]
```

Description

Configure user-based security model (USM) information.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

utc-timestamp

IN THIS SECTION

- Syntax | 30613
- Hierarchy Level | 30613
- Description | 30613
- Default | 30613
- Required Privilege Level | 30613
- Release Information | 30613

Syntax

```
utc-timestamp;
```

Hierarchy Level

For routers with the MX-SPC3 services card, you can configure the statement at the following hierarchy levels:

```
[edit services service-set service-set-name syslog],
[edit interfaces vms-interface|ams-interface services-options]
```

Description

Display the system timestamp in the syslog by overriding the default local system timestamp in syslog to UTC format.

Default

When the statement is not enabled, the syslog displays the local system timestamp in syslog.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced for routers with MX-SPC3 cards in Junos OS Release 21.1.

NOTE: In Junos OS Release 21.1R1, the statement is implemented at the at the

```
[edit interfaces vms-interface|ams-interface services-options
```

RELATED DOCUMENTATION

syslog (Interfaces) | 28820

utm default-configuration

IN THIS SECTION

- Syntax | 30614
- Hierarchy Level | 30620
- Description | 30620
- Options | 30620
- Required Privilege Level | 30621
- Release Information | 30621

Syntax

```
utm {
    default-configuration {
        anti-spam {
            address-blacklist;
            address-whitelist;
            sbl {
                custom-tag-string;
                (sbl-default-server | no-sbl-default-server);
                spam-action (block | tag-header | tag-subject);
            }
            traceoptions {
                flag name;
            }
            type (anti-spam-none | sbl);
        }
        anti-virus {
            mime-whitelist {
                exception;
                list;
```

```
            }
            sophos-engine {
                 fallback-options {
                    content-size (block | log-and-permit | permit);
                    default (block | log-and-permit | permit);
                    engine-not-ready (block | log-and-permit | permit);
                    out-of-resources (block | log-and-permit | permit);
                    timeout (block | log-and-permit | permit);
                    too-many-requests (block | log-and-permit | permit);
                }
                    forwarding-mode {
                        hold;
                        inline-tap;
                    }
                notification-options {
                    fallback-block {
                        custom-message;
                        custom-message-subject;
                        (notify-mail-sender | no-notify-mail-sender);
                        type (message | protocol-only);
                    }
                    fallback-non-block {
                        custom-message;
                        custom-message-subject;
                        (notify-mail-recipient | no-notify-mail-recipient);
                    }
                    virus-detection {
                        custom-message;
                        custom-message-subject;
                        (notify-mail-sender | no-notify-mail-sender);
                        type (message | protocol-only);
                    }
                }
                pattern-update {
                    email-notify {
                        admin-email;
                        custom-message;
                        custom-message-subject;
                    }
                    interval;
                    no-autoupdate;
                    proxy {
                        password;
```

```
                    port;
                    server;
                    username;
                }
                routing-instance;
                url;
            }
            scan-options {
                content-size-limit;
                timeout seconds;
                (uri-check | no-uri-check);
            }
            server {
                ip;
                routing-instance;
            }
            sxl-retry;
            sxl-timeout seconds;
            trickling timeout;
        }
        traceoptions {
            flag name;
        }
        url-whitelist;
    }
    content-filtering {
        block-command;
        block-content-type {
            activex;
            exe;
            http-cookie;
            java-applet;
            zip;
        }
        block-extension;
        block-mime {
            exception;
            list;
        }
        notification-options {
            custom-message;
            (notify-mail-sender | no-notify-mail-sender);
             seclog;
```

```
            type (message | protocol-only);
               } permit-command;
            traceoptions {
                flag name;
            }
            rule-set rule-set-name { /* New provision to add to default rules */
          rule rule-name {
          }
        }
type (content-filtering-none | local);
        }
        web-filtering {
            http-persist;
            http-reassemble;
            juniper-enhanced {
                base-filter;
                block-message {
                    type custom-redirect-url;
                    url;
                }
                cache {
                    size kilobytes;
                    timeout minutes;
                }
                category name {
                    action (block | log-and-permit | permit | quarantine);
                    custom-message;
                }
                custom-block-message;
                default (block | log-and-permit | permit | quarantine);
                fallback-settings {
                    default (block | log-and-permit);
                    server-connectivity (block | log-and-permit);
                    timeout (block | log-and-permit);
                    too-many-requests (block | log-and-permit);
                }
                no-safe-search;
                quarantine-custom-message;
                quarantine-message {
                    type custom-redirect-url;
                    url;
                }
                reputation {
```

```
                reputation-fairly-safe;
                reputation-moderately-safe;
                reputation-suspicious;
                reputation-very-safe;
            }
            server {
                host;
                port;
                routing-instance;
            }
            site-reputation-action {
                fairly-safe (block | log-and-permit | permit | quarantine);
                harmful (block | log-and-permit | permit | quarantine);
                moderately-safe (block | log-and-permit | permit | quarantine);
                suspicious (block | log-and-permit | permit | quarantine);
                very-safe (block | log-and-permit | permit | quarantine);
            }
            timeout seconds;
        }
        juniper-local {
            block-message {
                type custom-redirect-url;
                url;
            }
            category name {
                action (block | log-and-permit | permit | quarantine);
                custom-message;
            }
            custom-block-message;
            default (block | log-and-permit | permit);
            fallback-settings {
                default (block | log-and-permit);
                server-connectivity (block | log-and-permit);
                timeout (block | log-and-permit);
                too-many-requests (block | log-and-permit);
            }
            quarantine-custom-message;
            quarantine-message {
                type custom-redirect-url;
                url;
            }
            timeout seconds;
        }
```

```
            traceoptions {
                flag name;
            }
            url-blacklist;
            url-whitelist;
            websense-redirect {
                account;
                block-message {
                    type custom-redirect-url;
                    url;
                }
                category name {
                    action (block | log-and-permit | permit | quarantine);
                    custom-message;
                }
                custom-block-message;
                fallback-settings {
                    default (block | log-and-permit);
                    server-connectivity (block | log-and-permit);
                    timeout (block | log-and-permit);
                    too-many-requests (block | log-and-permit);
                }
                quarantine-custom-message;
                quarantine-message {
                    type custom-redirect-url;
                    url;
                }
                server {
                    host;
                    port;
                    routing-instance;
                }
                sockets;
                timeout seconds;
            }
        }
    }
    application-proxy;
    custom-objects;
    feature-profile;
    traceoptions;
    utm-policy junos-default-utm-policy;
```

```
    }
}
```

Hierarchy Level

```
[edit security utm]
```

Description

The Content Security default configuration is used in two scenarios.

- **Content Security default configuration for unified policies**—For security policies that enable Content Security with no custom Content Security policy defined, the default Content Security policy will be used.
- **Content Security default configuration for existing Content Security policies**—For existing security policies that have a Content Security policy enabled, the default Content Security policy will NOT be used.

Options

default-configuration	Global default Content Security configurations.
anti-spam	Configure the default Content Security configuration for antispam feature profile.
anti-virus	Configure the default Content Security configuration for antivirus feature profile.
content-filtering	Configure the default Content Security configuration for content filtering feature profile.
web-filtering	Configure the default Content Security configuration for Web filtering feature profile.
utm-policy	Configure a Content Security policy for antivirus, antispam, content filtering, traffic options, and Web filtering protocols and attach this policy to a security profile to implement it.
traceoptions	Define tracing operations for Content Security features.
feature-profile	Configure Content Security features, antivirus, antispam, content filtering, and Web filtering by creating feature profiles.

application-proxy Application proxy settings.

custom-objects Configure custom objects before configuring Content Security feature-profile features. Custom category does not take precedence over predefined categories when it has the same name as one of the predefined categories. It is not recommended to have a custom category name be the same as the predefined category name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Starting in Junos OS Release 21.4R1, the rule-set and rules configurations introduced under the `[edit security utm utm-policy <utm-policy-name> content-filtering]` hierarchy level can be used from `[edit security utm default-configuration content-filtering` hierarchy.

Content filtering options based on mime-type, content-type, and protocol command is not supported. After you upgrade to Junos OS Release 21.4R1, previously existing file extension based content filtering options under the `[edit security utm utm-policy <utm-policy-name> content-filtering]` hierarchy are no more available for configuration.

Junos OS Release 21.4R1 allows you to use legacy functionality if you don't want to migrate to this modern functionality. You will be allowed to use the legacy configurations but all the legacy configuration knobs are deprecated and are hidden. Also, you will receive system logs and error message warnings when you use all the legacy deprecated knobs.

Statement introduced in Junos OS Release 18.2R1.

uuid

IN THIS SECTION

Syntax | 30622

- Hierarchy Level | 30622
- Description | 30622
- Options | 30622
- Required Privilege Level | 30622
- Release Information | 30622

Syntax

```
uuid hex-value;
```

Hierarchy Level

```
[edit applications application application-name]
```

Description

Specify the Universal Unique Identifier (UUID) for DCE RPC objects.

Options

hex-value—Hexadecimal value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

ALG Descriptions

Configuring a Universal Unique Identifier

Examples: Configuring Application Protocols

Verifying the Output of ALG Sessions

uuid (Applications)

IN THIS SECTION

- Syntax | 30623
- Hierarchy Level | 30623
- Description | 30623
- Options | 30624
- Required Privilege Level | 30624
- Release Information | 30624

Syntax

```
uuid hex-value;
```

Hierarchy Level

```
[edit applications application  application-name  term  term-name]
```

Description

Specify the Universal Unique Identifier (UUID) for objects. DCE RPC services are mainly used by Microsoft applications. The ALG uses well-known TCP port 135 for port mapping services and uses the Universal Unique Identifier (UUID) instead of the program number to identify protocols. The main application-based DCE RPC is the Microsoft Exchange Protocol.

Support for stateful firewall and NAT services requires that you configure the DCE RPC port map ALG on TCP port 135. The DCE RPC ALG uses the TCP protocol with application-specific UUIDs.

Options

`uuid` *`hex-value`* —Specify the universal unique identifier (UUID) for objects.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 8.5.

V

IN THIS SECTION

- v3 | **30631**
- v4address | **30635**
- v4ov6 | **30636**
- v6address | **30638**
- v6prefix | **30639**
- v6rd | **30641**
- vacm | **30642**
- valid-lifetime | **30644**
- valid-lifetime (Dynamic Router Advertisement) | **30645**
- validation (Origin Validation for BGP) | **30647**
- validity-period | **30650**
- variable | **30651**
- variable | **30653**

- variables (Dynamic Service Profiles) | **30654**
- variant | **30655**
- vbr rate | **30657**
- vc-port | **30658**
- vci | **30660**
- vcp-no-hold-time | **30662**
- vcp-snmp-statistics | **30664**
- vdsl-bytes | **30665**
- vdsl-overhead-adjust | **30667**
- vdsl-profile | **30669**
- vdsl2-bytes | **30671**
- vdsl2-overhead-adjust | **30673**
- vendor | **30675**
- vendor-id (DHCP Client) | **30676**
- vendor-id (dhcp-option82) | **30677**
- vendor-id (Unified Edge) | **30679**
- vendor-id (AAA Profile) | **30681**
- vendor-id (Interfaces) | **30682**
- vendor-option | **30684**
- vendor-specific (DHCP Relay Agent) | **30686**
- vendor-support | **30687**
- verify-path | **30688**
- version-3 | **30691**
- version (Flow Aggregation) | **30692**
- version-alias (Dynamic Profiles) | **30694**
- version (BFD) | **30695**
- version (BFD Liveness Detection) | **30697**
- version (Dynamic IGMP Interface) | **30699**
- version (Dynamic MLD Interface) | **30701**
- version (Flow Monitoring Logs for NAT) | **30702**
- version (IGMP Snooping) | **30704**
- version (IKE) | **30706**

- version-ipfix (Forwarding Options) | **30707**
- version-ipfix (Services) | **30708**
- version (Liveness Detection) | **30710**
- version (MLD Snooping) | **30712**
- version (Protocols BFD) | **30714**
- version (Protocols IGMP) | **30715**
- version (Protocols IGMP AMT) | **30716**
- version (Protocols MLD) | **30718**
- version (Protocols PIM) | **30719**
- version (SNMP) | **30721**
- version9 (Flow Monitoring) | **30722**
- version9 (Forwarding Options) | **30724**
- versioning | **30725**
- video | **30727**
- video (Application Profile) | **30728**
- video-monitoring | **30729**
- view (Configuring a MIB View) | **30733**
- view (SNMP Community) | **30734**
- violation-action (DHCP Local Server and DHCP Relay Agent) | **30735**
- violation-report-rate (DDoS Flow Detection) | **30737**
- virtual-address | **30739**
- virtual-channel-group (CoS Interfaces) | **30740**
- virtual-channel-groups | **30742**
- virtual-channels | **30743**
- virtual-chassis | **30745**
- virtual-chassis (MX Series Virtual Chassis) | **30747**
- virtual-circuit-id | **30749**
- virtual-control-channel | **30751**
- virtual-cpu | **30752**
- virtual-gateway-address | **30754**
- virtual-gateway-v4-mac | **30755**
- virtual-gateway-v6-mac | **30757**

- virtual-inet6-address | **30760**
- virtual-ip | **30761**
- virtual-link | **30763**
- virtual-link-local-address | **30765**
- virtual-mac | **30766**
- virtual-network-functions | **30768**
- virtual-network-functions (Junos node slicing) | **30773**
- virtual-service (Traffic Load Balancer) | **30774**
- virtual-switch | **30777**
- virus-detection (Security Antivirus) | **30778**
- vjunos0 | **30780**
- vlan | **30782**
- vlan (Access Port Security) | **30785**
- vlan (Bridge Domains) | **30788**
- vlan (DHCP Bindings on Access Ports) | **30790**
- vlan-group | **30791**
- vlan-groups | **30792**
- vlan-id (Fast Ethernet) | **30794**
- vlan-id | **30795**
- vlan-id (Dynamic Profiles) | **30801**
- vlan-id (Dynamic VLANs) | **30802**
- vlan-id inner-all | **30804**
- vlan-id-list | **30805**
- vlan-id-list (Ethernet VLAN Circuit) | **30808**
- vlan-id-list (Interface in Bridge Domain) | **30810**
- vlan-id-list (Interface in VPLS) | **30812**
- vlan-id (Logical Port in Bridge Domain) | **30813**
- vlan-id-range | **30815**
- vlan-id (routing instance) | **30817**
- vlan-id-start | **30818**
- vlan-id (VLAN ID to Be Bound to a Logical Interface) | **30820**
- vlan (IGMP Snooping) | **30821**

- vlan-isolation | **30826**
- vlan members (VLANs) | **30827**
- vlan (Mirroring) | **30829**
- vlan (MLD Snooping) | **30830**
- vlan-model | **30833**
- vlan (MSTP) | **30835**
- vlan (PIM Snooping) | **30837**
- vlan (Port Mirroring) | **30839**
- vlan (Port Mirroring) | **30841**
- vlan-prune | **30842**
- vlan-range | **30843**
- vlan-ranges | **30845**
- vlan-ranges (RADIUS Options) | **30847**
- vlan-rewrite | **30849**
- vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) | **30850**
- vlan (Secure Access Port) | **30852**
- vlan (Static IP) | **30854**
- vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) | **30855**
- vlan-tag | **30857**
- vlan-tag (Dynamic Classifiers) | **30858**
- vlan-tag (Dynamic Rewrite Rules) | **30860**
- vlan-tagging | **30861**
- vlan-tagging (Dynamic) | **30864**
- vlan-tags | **30865**
- vlan-tags (Dynamic Profiles) | **30867**
- vlan-tags (Dual-Tagged Logical Interface) | **30869**
- vlan-tags-outer | **30872**
- vlan-tags (Stacked VLAN Tags) | **30873**
- vlan (Unknown Unicast Forwarding) | **30876**
- vlan (VSTP) | **30878**
- vlans | **30880**
- vlans (RA Guard) | **30894**

- vmhost | **30895**
- vmhost management-if disable | **30899**
- vmhost management-if link-mode | **30900**
- vmhost management-if speed | **30902**
- vmhost resize vjunos compact | **30904**
- vnf-license-supplement | **30905**
- vnf-name | **30906**
- vni | **30909**
- vni(Interfaces) | **30911**
- vni-options | **30912**
- vnid (EVPN) | **30914**
- voice | **30915**
- voice (Application Profile) | **30917**
- voip | **30918**
- voip-mac-exclusive | **30920**
- volume-limit (LRF Rule) | **30921**
- vpi (Define Virtual Path) | **30923**
- vpls-flow-table-size | **30924**
- vpls-flush-on-topology-change | **30926**
- vpls-id | **30927**
- vpls-id-list (protocols vpls mesh-group) | **30929**
- vpls (Interfaces) | **30931**
- vpls-mac-move | **30932**
- vpls (Routing Instance) | **30933**
- vpls-template | **30937**
- vpn-apply-export | **30938**
- vpn-group-address | **30939**
- vpn-label | **30941**
- vpn-localization | **30942**
- vpn-monitor | **30943**
- vpn (Security) | **30946**
- vpn-unequal-cost | **30951**

- vpn-global-import | **30953**
- vpws-service-id | **30954**
- vrf-advertise-selective | **30957**
- vrf-export | **30958**
- vrf-import | **30960**
- vrf-mtu-check | **30962**
- vrf-name (Duplicate Accounting) | **30963**
- vrf-propagate-ttl | **30965**
- vrf-table-label | **30966**
- vrf-target | **30968**
- vrrp | **30971**
- vrrp-group | **30973**
- vrrp-inet6-group | **30976**
- vrrp-inherit-from | **30978**
- vsi-discovery | **30980**
- vsi-policy | **30981**
- vstp | **30983**
- vtep-source-interface | **30988**
- vulnerability-type | **30990**
- vxlan | **30991**
- vxlan-disable-copy-tos-decap | **30993**
- vxlan-disable-copy-tos-encap | **30995**
- vxlan-extended | **30998**
- vxlan-gbp-l2-profile | **30999**
- vxlan-gbp-l3-profile | **31001**
- vxlan-gbp-profile | **31003**
- vxlan-gpe (FTI) | **31005**
- vxlan-routing | **31007**
- vxlan-stitching | **31009**
- vxlan-tos-copy-filter | **31011**

v3

IN THIS SECTION

- Syntax | 30631
- Hierarchy Level | 30633
- Options | 30634
- Description | 30634
- Required Privilege Level | 30635
- Release Information | 30635

Syntax

```
v3 {
    notify name {
        tag tag-name;
        type trap;
    }
    notify-filter profile-name {
        oid object-identifier (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        logical-system (SNMP) logical-system;
        port port-number;
        retry-count number;
        routing-instance instance;
        tag-list tag-list;
        target-parameters target-parameters-name;
        timeout seconds;
    }
```

```
    target-parameters target-parameters-name {
        notify-filter profile-name;
        parameters {
            message-processing-model (v1 | v2c | V3);
            security-level (authentication | none | privacy);
            security-model (usm | v1 | v2c);
            security-name security-name;
        }
    }
    usm {
        local-engine {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-sha {
                    authentication-password authentication-password;
                }
                authentication-none;
                privacy-aes128 {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-none;
            }
        }
        remote-engine engine-id {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-sha {
                    authentication-password authentication-password;
                }
                authentication-none;
                privacy-aes128 {
                    privacy-password privacy-password;
                }
```

```
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-3des {
                    privacy-password privacy-password;
                }
                privacy-none {
                    privacy-password privacy-password;
                }
            }
        }
    }
    vacm {
        access {
            group group-name {
                (default-context-prefix | context-prefix context-prefix){
                    security-model (any | usm | v1 | v2c) {
                        security-level (authentication | none | privacy) {
                            notify-view view-name;
                            read-view view-name;
                            write-view view-name;
                        }
                    }
                }
            }
        }
        security-to-group {
            security-model (usm | v1 | v2c) {
                security-name security-name {
                    group group-name;
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit snmp]
```

Options

username is the name that identifies the SNMPv3 user.

target-parameters-name is the name associated with the message processing and security parameters that are used in sending notifications to a particular management target.

When you configure SNMP trap notifications for subscriber secure policy on MX Series routers, you must configure the parameters as follows:

- Message-processing model: `v3`
- Security level: `privacy`
- Security model: `usm`

target-address-name is the string that identifies the target address.

tag-name defines the target addresses to which this notification is sent. This notification is sent to all the target-addresses that have this tag in their tag list. The *tag-name* is not included in the notification.

`trap` is the type of notification.

name is the name assigned to the notification.

profile-name is the name assigned to the notify filter.

oid is the object identifier. All MIB objects represented by this statement have the specified OID as a prefix. You can specify it either by a sequence of dotted integers or by a subtree name.

- `include`—Include the subtree of MIB objects represented by the specified OID.
- `exclude`—Exclude the subtree of MIB objects represented by the specified OID.

authentication-password is the password used to generate the key used for authentication.

SNMPv3 has specific requirements when you create plain-text passwords on a router or switch:

- The password must be at least eight characters long.
- The password can include alphabetic, numeric, and special characters, but it cannot include control characters.

Description

Configure SNMPv3.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

v4address

IN THIS SECTION

- Syntax | 30635
- Hierarchy Level | 30635
- Description | 30636
- Required Privilege Level | 30636
- Release Information | 30636

Syntax

```
v4address {
    equals;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format]
```

Description

Specify the custom AVP attribute's format as an IPv4 address and the value to match for the incoming RADIUS request from the IP-based subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers*
- *Understanding Selection of Properties for an IP-Based TDF Subscriber*
- *Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber*
- *IP-Based Subscriber Setup Overview*

v4ov6

IN THIS SECTION

- Syntax | 30637
- Hierarchy Level | 30637
- Description | 30637
- Options | 30637

Required Privilege Level | 30637

Syntax

```
v4ov6 ipv6-anycast-source-duplication;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options dynamic-tunnels],
[edit logical-systems name routing-options dynamic-tunnels],
[edit logical-systems name tenants name routing-instances name routing-options dynamic-tunnels],
[edit routing-instances name routing-options dynamic-tunnels],
[edit routing-options dynamic-tunnels],
[edit tenants name routing-instances name routing-options dynamic-tunnels]
```

Description

Enable dynamic tunnel V4oV6 mode

Options

ipv6-anycast-source-duplication Enable full resolved nh base source-tunnel

Required Privilege Level

routing

v6address

IN THIS SECTION

- Syntax | 30638
- Hierarchy Level | 30638
- Description | 30638
- Required Privilege Level | 30638
- Release Information | 30639

Syntax

```
v6address {
    equals;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format]
```

Description

Specify the custom AVP attribute's format as an IPv6 address and the value to match for the incoming RADIUS request from the IP-based subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

v6prefix

IN THIS SECTION

- Syntax | 30639
- Hierarchy Level | 30640
- Description | 30640
- Required Privilege Level | 30640
- Release Information | 30640

Syntax

```
v6prefix {
    apply-groups [group-names];
    apply-groups-except [group-names];
    equals;
}
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name
format]
```

Description

Specify the custom AVP attribute's format as an IPv6 address prefix and the value to match for the incoming RADIUS request from the IP-based subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

- *Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers*
- *Understanding Selection of Properties for an IP-Based TDF Subscriber*
- *Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber*
- *IP-Based Subscriber Setup Overview*

v6rd

IN THIS SECTION

- Syntax | 30641
- Hierarchy Level | 30641
- Description | 30641
- Options | 30642
- Required Privilege Level | 30642
- Release Information | 30642

Syntax

```
v6rd v6rd-softwire-concentator {
    ipv4-prefix ipv4-prefix;
    v6rd-prefix ipv6-prefix;
    mtu-v4 mtu-v4;
    softwire-address ipv4-address;
}
```

Hierarchy Level

```
[edit services softwire softwire-concentrator]
[edit services softwires softwire-types
```

Description

Configure settings for a 6rd concentrator used to process IPv6 packets encapsulated in IPv4 packets.

The `v6rd` statement is supported only on the MS-DPC, MS-100, MS-400, and MS-500 line cards. The `v6rd` statement is *not* supported on MS-MPCs and MS-MICs.

Options

ipv4-prefix—IPv4 prefix of the customer edge (CE) network

ipv6-prefix—IPv6 prefix of the 6rd domain.

mtu-v4— Maximum transmission unit (MTU), in bytes (576 through 9192), for IPv6 packets encapsulated into IPv4. If the final length is greater than the configured value, the IPv4 packet will be dropped.

address—IPv4 address of a softwire concentrator. This is an IPv4 address independent of any interface and on a different prefix.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support added in Junos OS release 20.2R1 for the v6rd concentrator at the `[edit services softwires softwire-types` edit hierarchy for Next Gen Services on MX240, MX480, and MX860 routers.

RELATED DOCUMENTATION

Configuring a 6rd Softwire Concentrator

vacm

IN THIS SECTION

- Syntax | 30643
- Hierarchy Level | 30643
- Description | 30643
- Required Privilege Level | 30644

- Release Information | **30644**

Syntax

```
vacm {
    access {
        group group-name {
            (default-context-prefix | context-prefix context-prefix){
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
    security-to-group {
        security-model (usm | v1 | v2c);
        security-name security-name {
            group group-name;
        }
    }
}
```

Hierarchy Level

```
[edit snmp v3]
```

Description

Configure view-based access control model (VACM) information.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

valid-lifetime

IN THIS SECTION

- Syntax | 30644
- Hierarchy Level | 30644
- Description | 30645
- Options | 30645
- Required Privilege Level | 30645
- Release Information | 30645

Syntax

```
valid-lifetime seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-
name prefix prefix],
[edit protocols router-advertisement interface interface-name prefix prefix]
```

Description

Specify how long the prefix remains valid for onlink determination.

Options

seconds—Valid lifetime, in seconds. If you set the valid lifetime to **0xffffffff**, the lifetime is infinite.

- **Default:** 2,592,000 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

preferred-lifetime | 13056

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

valid-lifetime (Dynamic Router Advertisement)

IN THIS SECTION

- Syntax | 30646
- Hierarchy Level | 30646
- Description | 30646
- Options | 30646
- Required Privilege Level | 30646
- Release Information | 30646

Syntax

```
valid-lifetime seconds;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols router-advertisement interface interface-name
prefix prefix]
```

Description

Specify how long the prefix remains valid for onlink determination.

Options

seconds—Valid lifetime, in seconds. If you set the valid lifetime to `0xffffffff`, the lifetime is infinite.

- **Default:** 2,592,000 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

> *IPv6 WAN Link Addressing with NDRA*
>
> *Dynamic Router Advertisement Configuration Overview*

validation (Origin Validation for BGP)

IN THIS SECTION

- Syntax | 30647
- Hierarchy Level | 30648
- Description | 30648
- Options | 30648
- Required Privilege Level | 30649
- Release Information | 30649

Syntax

```
validation {
    traceoptions {
        file filename <files number> <size size> <(world-readable | no-world-readable)>;
        flag flag flag-modifier;
        }
    notification-rib test;
    group group-name {
        max-sessions number;
        session server-ip-address {
            traceoptions {
                file filename <files number> <size size> <(world-readable | no-world-readable)>;
                flag flag flag-modifier;
            }
            refresh-time number;
            hold-time number;
            record-lifetime number;
            preference number;
            port number;
            local-address local-ip-address;
        }
    }
    static {
        record record-destination {
            maximum-length prefix-length {
```

```
                origin-autonomous-system asn-number {
                    validation-state (invalid | valid);
                }
            }
        }
    }
    database database-name {
        static {
            record destination {
                maximum-length prefix-length {
                    origin-autonomous-system as-number {
                        validation-state valid;
                    }
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name routing-options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances instance-name routing-options],
[edit routing-options]
```

Description

Configure resource public key infrastructure (RPKI) BGP route validation.

Options

notification-rib [*notification-rib* ...]; Specify the routing tables that are notified when the validation state changes.

When validation policies are used for BGP peers in routing instances, the policy engine tries to lookup within the local Validated Route Payload (VRP) database of that particular routing instance. If the RPKI session is not in that routing instance, then it falls back to the VRP database of the default routing instance. Modification of VRP records in the fall back VRP database will not be picked up by the non-

default routing instance. As a result, routing table entries of the default and non-default routing instances will differ.

To offset this, configure the `notification-rib` option in the default routing instance to ensure that the modifications to the default routing instance trigger a re-evaluation of the routing tables for the specified routing instance.

Junos OS Release 22.3R1 introduced support for named validation databases in addition to the default VRP database.

To specify a named validation database, use the `validation-state (invalid | valid)` option at the `[edit routing-options validation database` *`database-name`* `static record` *`destination`* `maximum-length` *`prefix-length`* `origin-autonomous-system` *`as-number]`* hierarchy level.

To specify a target route-validation database for a validation session, use the database `database-name` option at the `[edit routing-options validation group` *`group-name`* `session]` hierarchy level.

As a result, we no longer need to configure notification-ribs explicitly. They are created internally to track which VRP databases are being consulted by which routing-instances.

The remaining statements are explained separately. See CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Use Case and Benefit of Origin Validation for BGP

Understanding Origin Validation for BGP

Example: Configuring Origin Validation for BGP

validity-period

IN THIS SECTION

- Syntax | 30650
- Hierarchy Level | 30650
- Description | 30650
- Options | 30650
- Required Privilege Level | 30651
- Release Information | 30651

Syntax

```
validity-period days;
```

Hierarchy Level

```
[edit security pki auto-re-enrollment certificate-id]
```

Description

Certificate validity period, in days, from the enrollment start date. If not specified, the issuing certificate authority (CA) sets this time as per its own policy. The start time is when auto-reenrollment is initiated.

Options

days—Number of days that the certificate is valid.

- **Range:** 1 through 4095 days
- **Default:** Per CA policy

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Auto-Reenrollment of a Router Certificate

auto-re-enrollment | 2727

variable

IN THIS SECTION

- Syntax | 30651
- Hierarchy Level | 30652
- Description | 30652
- Options | 30652
- Required Privilege Level | 30652
- Release Information | 30652

Syntax

```
variable oid-variable;
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Set the object identifier (OID) of the MIB object (also called variable) to be monitored when you configure an SNMP RMON alarm. If the value of the monitored variable exceeds the configured rising threshold or falling threshold, an alarm is triggered and a corresponding event may be generated.

Options

oid-variable—OID of the MIB variable that is being monitored. The OID can be a dotted decimal (for example, `1.3.6.1.2.1.2.1.2.2.1.10.1`). Alternatively, use the MIB object name (for example, `ifInOctets.1`).

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

RMON MIB Event, Alarm, Log, and History Control Tables

RMON Overview

Junos OS Network Management Configuration Guide

variable

IN THIS SECTION

- Syntax | 30653
- Hierarchy Level | 30653
- Description | 30653
- Options | 30653
- Required Privilege Level | 30653
- Release Information | 30654

Syntax

```
variable oid-variable;
```

Hierarchy Level

```
[edit snmp rmon alarm index]
```

Description

Object identifier (OID) of MIB variable to be monitored.

Options

oid-variable—OID of the MIB variable that is being monitored. The OID can be a dotted decimal (for example, `1.3.6.1.2.1.2.1.2.2.1.10.1`). Alternatively, use the MIB object name (for example, `ifInOctets.1`).

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Remote Nework Monitoring (RMON)

variables (Dynamic Service Profiles)

IN THIS SECTION

- Syntax | 30654
- Hierarchy Level | 30654
- Description | 30655
- Options | 30655
- Required Privilege Level | 30655
- Release Information | 30655

Syntax

```
variables variable-name {
    default-value default-value;
    equals expression;
    mandatory;
    uid;
    uid-reference;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Configure user-defined variables in a dynamic service profile. The values that the system uses for these variables are provided by the RADIUS server and applied when the subscriber authenticates. You can configure default values that are used when RADIUS does not return a value. Alternatively, you can specify that the profile fails if RADIUS does not return a value for a variable.

NOTE: Do not use this statement in a dynamic client profile.

Options

variable-name—Name of the variable.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Dynamic Variables Overview

User-Defined Variables in Dynamic Profiles

variant

IN THIS SECTION

Syntax | 30656

- Hierarchy Level | 30656
- Description | 30656
- Required Privilege Level | 30656
- Release Information | 30656

Syntax

```
variant variant-number {
    data-format format;
    name-format format;
    transfer {
        record-level number;
        timeout seconds;
    }
}
```

Hierarchy Level

```
[edit services flow-collector file-specification]
```

Description

Configure a variant of the file format.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring File Formats

vbr rate

IN THIS SECTION

- Syntax | 30657
- Hierarchy Level | 30657
- Description | 30657
- Options | 30657
- Required Privilege Level | 30658
- Release Information | 30658

Syntax

```
vbr rate;
```

Hierarchy Level

```
[edit interfaces interface-name atm-options vpi vpi-identifier shaping]
```

Description

For ATM encapsulation only, define a variable bit rate bandwidth utilization in the traffic-shaping profile.

Options

- Burst Size–The maximum burst size that can be sent at the peak rate.
- Peak Rate–The maximum instantaneous rate at which the user will transmit.

- Sustained Rate–The average rate as measured over a long interval.
- CDVT–Cell Delay Variation Tolerance in microseconds (range: 1 – 9999).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Interfaces

vc-port

IN THIS SECTION

- Syntax | 30659
- Hierarchy Level | 30659
- Description | 30659
- Default | 30659
- Options | 30659
- Required Privilege Level | 30659
- Release Information | 30660

Syntax

```
vc-port {
    lag-hash (packet-based | source-port-based);
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Enable hashing of link aggregation group (LAG) network traffic over a dedicated trunk port within a Virtual Chassis, and select how the traffic within the dedicated trunk port is hashed.

BEST PRACTICE: Do not configure this statement unless you have a compelling reason to configure it. Configuration of this statement is optional and is only useful in a few types of network setups.

Default

source-port-based

Options

lag-hash Select how to direct LAG traffic through the dedicated trunk port using either of the following options:

- **packet-based** Hashes all incoming LAG network traffic on the dedicated trunk port based on the packet.
- **source-port-based** Hashes all incoming LAG network traffic on the dedicated trunk port based on the source.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

vci

IN THIS SECTION

- Syntax | 30660
- Hierarchy Level | 30660
- Description | 30661
- Options | 30661
- Required Privilege Level | 30661
- Release Information | 30661

Syntax

```
vci vpi-identifier.vci-identifier;
```

Hierarchy Level

```
[edit interfaces at-fpc/pic/port unit logical-unit-number],
[edit interfaces at-fpc/pic/port unit logical-unit-number family family address address
multipoint-destination address],
[edit logical-systems logical-system-name interfaces at-fpc/pic/port unit logical-unit-number],
[edit logical-systems logical-system-name interfaces at-fpc/pic/port unit logical-unit-number
family family address address multipoint-destination address]
```

Description

For ATM point-to-point logical interfaces only, configure the virtual circuit identifier (VCI) and virtual path identifier (VPI).

To configure a VPI for a point-to-multipoint interface, specify the VPI in the multipoint-destination statement.

VCIs 0 through 31 are reserved for specific ATM values designated by the ATM Forum.

Options

vci-identifier—ATM virtual circuit identifier. Unless you configure the interface to use promiscuous mode, this value cannot exceed the highest-numbered VC configured for the interface with the `maximum-vcs` option of the `vpi` statement.

- **Range:** 0 through 4089 or 0 through 65,535 with promiscuous mode, with VCIs 0 through 31 reserved.

vpi-identifier—ATM virtual path identifier.

- **Range:** 0 through 255
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring a Point-to-Point ATM1 or ATM2 IQ Connection

Applying Scheduler Maps to Logical ATM Interfaces

vcp-no-hold-time

IN THIS SECTION

- Syntax | 30662
- Hierarchy Level | 30662
- Description | 30662
- Default | 30663
- Required Privilege Level | 30663
- Release Information | 30663

Syntax

```
vcp-no-hold-time;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Disable the Virtual Chassis port (VCP) holddown timer for all VCPs in the Virtual Chassis or Virtual Chassis Fabric (VCF).

The VCP holddown timer is an internal mechanism that delays a Virtual Chassis reconvergence for several seconds when a VCP becomes inactive. The purpose of this delay is to provide the VCP time to return online without having to reconverge the Virtual Chassis to adjust to the inactive VCP. All traffic to the VCP is dropped while the VCP is inactive. If the VCP remains down for a time that exceeds the VCP holddown timer, a Virtual Chassis reconvergence occurs.

This statement disables the holddown timer only in an EX4300 Virtual Chassis or a mixed Virtual Chassis that contains EX4300 switches in releases in the Junos OS 13.2X50 release train. In releases after that, the `vcp-no-hold-time` option is no longer needed and has no effect because the holddown timer is replaced by a planned PFE restart for actions that affect Virtual Chassis reconvergence. Switches and releases that don't support the holddown timer might allow you to configure this statement, but the

configuration posts a warning message saying the statement has no effect. The option will be deprecated in an upcoming release and will no longer appear in the CLI.

When this statement is enabled, the VCP holddown timer is disabled and the Virtual Chassis reconvergence occurs when a VCP becomes inactive. The period of time where traffic is dropped waiting for the VCP to return online is avoided.

We recommend enabling this statement after a Virtual Chassis is operational. We recommend disabling this statement when you are adding or removing member switches from your Virtual Chassis.

The VCP holddown timer cannot be viewed and is not user-configurable. You can only control whether the VCP holddown timer is enabled or disabled by configuring this statement.

> **NOTE**: In an EX4300 Virtual Chassis running a Junos OS 13.2X50 release, you should enable the `vcp-no-hold-time` statement before performing a software upgrade using NSSU. If you do not enable the `vcp-no-hold-time` statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see *Understanding Split and Merge in a Virtual Chassis*.

Default

The VCP holddown timer is enabled by default on all devices that support this statement.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

- *Understanding EX Series Virtual Chassis*
- *Understanding QFX Series Virtual Chassis*
- *Understanding Virtual Chassis Components*

vcp-snmp-statistics

IN THIS SECTION

- Syntax | 30664
- Hierarchy Level | 30664
- Description | 30664
- Default | 30664
- Required Privilege Level | 30665
- Release Information | 30665

Syntax

```
vcp-snmp-statistics;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Enable SNMP monitoring of the Virtual Chassis ports (VCPs) for all VCPs in the Virtual Chassis or Virtual Chassis Fabric (VCF).

When this statement is enabled, SNMP gathers statistics on the Junos VCP MIBs. You can retrieve the statistics gathered by SNMP for these MIBs by using the show snmp mib command with the `walk` and `ascii` options and specifying `jnxVirtualChassisPortInPkts`.

Default

SNMP is disabled by default on devices running Junos OS.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D40 and 15.1R5 only.

Statement deprecated in Junos OS Releases after 14.1X53 and 15.1.

RELATED DOCUMENTATION

Understanding EX Series Virtual Chassis

Understanding QFX Series Virtual Chassis

Understanding Virtual Chassis Components

vdsl-bytes

IN THIS SECTION

- Syntax | 30665
- Hierarchy Level | 30666
- Description | 30666
- Options | 30666
- Required Privilege Level | 30666
- Release Information | 30666

Syntax

```
vdsl-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of frame overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for a VDSL access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

> **NOTE**: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `vdsl-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream frame overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Neighbors and Operations

vdsl-overhead-adjust

IN THIS SECTION

- Syntax | 30667
- Hierarchy Level | 30668
- Description | 30668
- Options | 30668
- Required Privilege Level | 30668
- Release Information | 30668

Syntax

```
vdsl-overhead-adjust percentage;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the actual downstream rate for a VDSL access line reported in the ANCP Port Up message by multiplying the rate by the specified percentage. The ANCP agent reports the adjusted rate to CoS.

> **NOTE**: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `vdsl-overhead-adjust` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

percentage Percentage by which to multiply the rate.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Neighbors and Operations

vdsl-profile

IN THIS SECTION

- Syntax | 30669
- Hierarchy Level | 30670
- Description | 30670
- Options | 30670
- Required Privilege Level | 30670
- Release Information | 30670

Syntax

```
vdsl-profile
```

Hierarchy Level

```
[edit interfaces interface-name vdsl-options]
```

Description

Configure the type of VDSL2 profiles. A profile is a table that contains a list of preconfigured VDSL2 settings.

Options

- Auto (default)
- 8a
- 8b
- 8c
- 8d
- 12a
- 12b
- 17a

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

VDSL2 Interface Support on SRX Series Devices

vdsl2-bytes

IN THIS SECTION

- Syntax | 30671
- Hierarchy Level | 30671
- Description | 30671
- Options | 30672
- Required Privilege Level | 30672
- Release Information | 30672

Syntax

```
vdsl2-bytes bytes;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the number of frame overhead bytes by the specified number of bytes in the actual downstream rate reported in the ANCP Port Up message for a VDSL2 access line. The ANCP agent reports the adjusted value to CoS. The adjusted value accounts for the traffic encapsulation overhead.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `vdsl2-bytes` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

bytes Number of bytes added to or subtracted from the actual downstream frame overhead.

- **Range:** -100 through 100 bytes
- **Default:** 0 bytes

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-overhead-bytes` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Neighbors and Operations

vdsl2-overhead-adjust

IN THIS SECTION

- Syntax | 30673
- Hierarchy Level | 30673
- Description | 30673
- Options | 30674
- Required Privilege Level | 30674
- Release Information | 30674

Syntax

```
vdsl2-overhead-adjust percentage;
```

Hierarchy Level

```
[edit protocols ancp qos-adjust]
```

Description

Adjust the actual downstream rate for a VDSL2 access line received in the ANCP Port Up message by multiplying the rate by the specified percentage. The ANCP agent reports the adjusted rate to CoS.

NOTE: Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level. The ANCP agent ignores the `vdsl2-overhead-adjust` statement if it is present. This means that when you upgrade from Junos OS Release 17.3 or earlier with an existing configuration, you must reconfigure your adjustment with the `access-line` statement.

Options

percentage Percentage by which to multiply the rate.

- **Range:** 80 through 100 percent
- **Default:** 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement deprecated in Junos OS Release 17.4R1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, this statement is deprecated and replaced by the `vdsl2-overhead-adjust` option of the `access-line` statement at the `[edit system]` hierarchy level.

RELATED DOCUMENTATION

ANCP Agent Traffic Shaping and CoS

ANCP Agent Traffic Shaping and CoS

ANCP Agent Neighbors and Operations

vendor

IN THIS SECTION

- Syntax | 30675
- Hierarchy Level | 30675
- Description | 30675
- Options | 30675
- Required Privilege Level | 30676
- Release Information | 30676

Syntax

```
vendor name {
    product-name product-name;
}
```

Hierarchy Level

```
[edit security idp   dynamic-attack-group              name              filters]
```

Description

A vendor has single or multiple products. An attack can be applicable to some of the products of vendors. This filter can be used to group attacks specific to the product of a vendor.

Options

name	Values for vendor field
product-name	Values for product field

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 18.2R1.

vendor-id (DHCP Client)

IN THIS SECTION

- Syntax | 30676
- Hierarchy Level | 30676
- Description | 30676
- Options | 30677
- Required Privilege Level | 30677
- Release Information | 30677

Syntax

```
vendor-id vendor-id;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family dhcp-client]
```

Description

Configure a vendor class ID for the Dynamic Host Configuration Protocol (DHCP) client.

Options

vendor-id	Vendor class ID.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

Starting with Junos OS Release 17.3R1, on all SRX Series Firewalls and vSRX Virtual Firewall instances, the CLI option `dhcp-clinet` at `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy is changed to `dhcp` to keep consistent with other Junos platforms. There is no change in the functionality with the changed CLI option `dhcp`.

vendor-id (dhcp-option82)

IN THIS SECTION

- Syntax | 30678
- For Platforms with Enhanced Layer 2 Software (ELS) | 30678
- For Platforms Without ELS | 30678
- For MX Series Platforms | 30678
- Description | 30678
- Default | 30678
- Options | 30678
- Required Privilege Level | 30679
- Release Information | 30679

Syntax

```
vendor-id <string>;
```

For Platforms with Enhanced Layer 2 Software (ELS)

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]
```

For Platforms Without ELS

```
[edit ethernet-switching-options secure-access-port vlan (all | vlan-name) dhcp-
option82],
[edit forwarding-options helpers bootp dhcp-option82],
[edit forwarding-options helpers bootp interface interface-name dhcp-option82]
```

For MX Series Platforms

```
[edit bridge-domains bridge-domain-name forwarding-options dhcp-security option-82]
```

Description

Insert a vendor ID in the DHCP option 82 information in a DHCP request packet header before forwarding or relaying the request to a DHCP server.

Default

If `vendor-id` is not explicitly configured for DHCP option 82, then no vendor ID is set.

Options

string—(Optional) A single string that designates the vendor ID.

Range: 1–255 characters

Default: If you specify `vendor-id` with no *string* value, then the default vendor ID `Juniper Networks` is configured.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See Using the Enhanced Layer 2 Software CLI for information about ELS.)

Hierarchy level `[edit bridge-domains bridge-domain-name forwarding-options dhcp-security]` introduced in Junos OS Release 14.1 for the MX Series.

RELATED DOCUMENTATION

Configuring DHCP Option 82 on a Router with Bridge Domain

Example: Setting Up DHCP Option 82 Using the Same VLAN

Example: Setting Up DHCP Option 82 on a VLAN

Setting Up DHCP Option 82 on the Switch with No Relay (ELS)

vendor-id (Unified Edge)

IN THIS SECTION

- Syntax | 30680
- Hierarchy Level | 30680
- Description | 30680
- Options | 30680
- Required Privilege Level | 30680
- Release Information | 30680

Syntax

```
vendor-id vendor-id;
```

Hierarchy Level

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute name]
```

Description

Specify the custom attribute's AVP vendor identification number for the incoming RADIUS request from the subscriber.

After this criterion and the other match criteria specified for the TDF domain or PCEF profile selection term are matched, the specified TDF domain or PCEF profile is selected.

Options

vendor-id AVP vendor identification number.

- **Range:** 0 through 65,534.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

Understanding Selection of Properties for an IP-Based TDF Subscriber

Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

IP-Based Subscriber Setup Overview

vendor-id (AAA Profile)

IN THIS SECTION

- Syntax | 30681
- Hierarchy Level | 30681
- Description | 30681
- Options | 30682
- Required Privilege Level | 30682
- Release Information | 30682

Syntax

```
vendor-id vendor-id;
```

Hierarchy Level

```
[edit unified-edge aaa profiles aaa-profile-name radius policy activation-attribute],
[edit unified-edge aaa profiles aaa-profile-name radius policy deactivation-attribute]
```

Description

Configure the vendor identification when a vendor-specific RADIUS attribute is used to carry the policy and charging control (PCC) rulebase name for rulebase activations or deactivations. By default, the rulebase name is carried in the ERX-Service-Activate Juniper vendor-specific attribute (VSA) for activations and in the ERX-Service-Deactivate Juniper VSA for deactivations.

Options

vendor-id Vendor identification number for the RADIUS AVP.

Required Privilege Level

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Configuring an AAA Profile

Understanding AAA Profiles

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

vendor-id (Interfaces)

IN THIS SECTION

- Syntax | 30683
- Hierarchy Level | 30683
- Description | 30683
- Options | 30683
- Required Privilege Level | 30683
- Release Information | 30683

Syntax

```
vendor-id          vendor-id        ;
```

Hierarchy Level

```
[edit interfaces  interface-name unit  logical-unit-number family  family-name dhcp]
```

Description

Configure a vendor class ID for the Dynamic Host Configuration Protocol (DHCP) client.

Options

vendor-id —vendor class ID.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Interfaces

vendor-option

IN THIS SECTION

- Syntax | 30684
- Hierarchy Level | 30684
- Description | 30685
- Options | 30685
- Required Privilege Level | 30685
- Release Information | 30685

Syntax

```
vendor-option {
    (default-relay-server-group server-group-name | default-local-server-group local-server-group-name | drop);
    (equals | starts-with) (ascii match-string | hexadecimal match-hex) {
        (drop | local-server-group local-server-group-name | relay-server-group server-group-name);
    }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-60],
[edit forwarding-options dhcp-relay group group-name relay-option-60],
[edit logical-systems logical-system-name forwarding-options dhcp-relay relay-option-60],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name relay-option-60],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-60],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay group group-name relay-option-60],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay relay-option-60],
```

```
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name relay-
option-60]
```

Description

Configure the match criteria when you use the DHCP vendor class identifier option (option 60) in DHCP client packets to forward client traffic to specific DHCP servers. The extended DHCP relay agent compares the option 60 vendor-specific strings received in DHCP client packets against the match criteria that you specify. If there is a match, you can define certain actions for the associated DHCP client packets.

The `vendor-option` statement enables you to specify either an exact, left-to-right match (with the `equals` statement) or a partial match (with the `starts-with` statement), and configure either an ASCII match string (with the `ascii` statement) or a hexadecimal match string (with the `hexadecimal` statement).

You can configure an unlimited number of match strings. Match strings do not support the use of wildcard attributes.

Options

equals—Exact, left-to-right match of the ASCII or hexadecimal match string with the option 60 string.

starts-with—Partial match of the ASCII or hexadecimal match string with the option 60 string. The option 60 string can contain a superset of the ASCII or hexadecimal match string, provided that the leftmost characters of the option 60 string entirely match the characters in the configured match string. When you use the `starts-with` statement, the longest match rule applies; that is, the router matches the string “test123” before it matches the string “test”.

ascii *match-string*—ASCII match string of 1 through 255 alphanumeric characters.

hexadecimal *match-hex*—Hexadecimal match string of 1 through 255 hexadecimal characters (0 through 9, a through f, A through F).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 9.0.

vendor-specific (DHCP Relay Agent)

IN THIS SECTION

- Syntax | 30686
- Hierarchy Level | 30686
- Description | 30686
- Required Privilege Level | 30687
- Release Information | 30687

Syntax

```
vendor-specific {
    host-name;
    location;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay relay-option-82],
[edit forwarding-options dhcp-relay group group-name relay-option-82]
```

Description

Add vendor-specific information to the option-82, suboption 9 field of DHCPv4 control messages on server-facing interfaces. The vendor-specific information can be a hostname, a location (such as a unique connection identifier), or both. The hostname can be a string of characters such as **Juniper-AB-1**. The location should be specified as interface, VLAN ID, and if applicable, stacked VLAN ID. For example, **<ifd-name>:<vlan>** (ae0:100) or **<ifd-name>:<svlan> -<vlan>** (ae0:100-10).

Junos OS automatically adds the remaining vendor-specific information as per RFC 4243. The enterprise ID is 2636. The hostname is option-data 1, and the location is option-data 2. The DHCPv4 relay strips the suboption data from replies from the server before it relays the packets out the client facing interface.

This feature can be useful, in conjunction with operator-developed tools, for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query a single entity to track and troubleshoot subscriber IP information and network attachment points.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 16.2.

vendor-support

IN THIS SECTION

- Syntax | 30687
- Hierarchy Level | 30687
- Description | 30688
- Required Privilege Level | 30688
- Release Information | 30688

Syntax

```
vendor-support ibm;
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Description

Configure support for any vendor-specific template types. Currently, the only vendor-specific template type is `ipflow-tcp-ts`, for which you configure `vendor-specific ibm`.

If you do not configure `vendor-specific ibm`, a warning appears when you commit the configuration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

- *Configuring an LRF Profile for Subscribers*
- Configuring Logging and Reporting for Junos OS Subscriber Aware
- *Configuring Logging and Reporting for Subscriber Management*
- *Logging and Reporting Function for Subscribers*

verify-path

IN THIS SECTION

- Syntax | 30689
- Hierarchy Level | 30689
- Description | 30689
- Options | 30690
- Required Privilege Level | 30690
- Release Information | 30690

Syntax

```
verify-path {
    destination-ip ip-address;
    packet-size bytes;
}
```

Hierarchy Level

```
[edit security ipsec vpn vpn-name vpn-monitor]
```

Description

Verify the IPsec datapath before the secure tunnel (st0) interface is activated and route(s) associated with the interface are installed in the Junos OS forwarding table. This configuration is useful in network topologies where there is a transit firewall located between the VPN tunnel endpoints, and where IPsec data traffic that uses active routes for an established VPN tunnel on the st0 interface might be blocked by the transit firewall.

When this option is configured, the source interface and destination IP addresses that can be configured for VPN monitor operation are not used for IPsec datapath verification. The source for the ICMP requests in the IPsec datapath verification is the local tunnel endpoint.

When IPsec datapath verification is configured, the following actions occur:

1. Upon the establishment of the VPN tunnel, an ICMP request is sent to the peer tunnel endpoint to verify the IPsec datapath.

   The peer tunnel endpoint must be reachable by VPN monitor ICMP requests and must be able to respond to the ICMP request. While the datapath verification is in progress, "`V`" is displayed in the VPN Monitoring field in the `show security ipsec security-association detail` command output.

2. The `st0` interface is activated only when a response is received from the peer.

   The `show interface st0.x` command output shows the st0 interface status during and after the datapath verification: `Link-Layer-Down` before the verification finishes and `Up` after the verification finishes successfully.

3. If no ICMP response is received from the peer, another ICMP request is sent at the configured VPN monitor interval (the default is 10 seconds) until the VPN monitor threshold (the default is 10 times) is reached.

If the verification does not succeed, the KMD_VPN_DOWN_ALARM_USER system log entry indicates the reason as a VPN monitoring verify-path error. The error is logged under tunnel events in the `show security ipsec security-association detail` command output. The `show security ipsec tunnel-events-statistics` command displays the number of times the error occurred.

VPN monitor `interval` and `threshold` values are configured with `vpn-monitor-options` at the [`edit security ipsec`] hierarchy level.

4. If no ICMP response is received from the peer after the VPN monitor threshold is reached, the established VPN tunnel is brought down and the VPN tunnel is renegotiated.

Options

destination-ip *ip-address*

Original, untranslated IP address of the peer tunnel endpoint that is behind a NAT device. This IP address must not be the NAT translated IP address. This option is required if the peer tunnel endpoint is behind a NAT device. The verify-path ICMP request is sent to this IP address so that the peer can generate an ICMP response.

packet-size *bytes*

(Optional) The size of the packet that is used to verify an IPsec datapath before the st0 interface is brought up.

The packet size must be lower than the path maximum transmission unit (PMTU) minus tunnel overhead. The packet used for IPsec datapath verification must not be fragmented.

- **Range:** 64 to 1350 bytes
- **Default:** 64 bytes

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D70.

`packet-size` option added in Junos OS Release 15.1X49-D120.

Support for `verify-path` option with IPsec VPN running IKED process is introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

vpn-monitor | 30943

vpn (Security) | 30946

version-3

IN THIS SECTION

- Syntax | 30691
- Hierarchy Level | 30691
- Description | 30691
- Default | 30692
- Required Privilege Level | 30692
- Release Information | 30692

Syntax

```
version-3;
```

Hierarchy Level

```
[edit protocols vrrp]
```

Description

Enable Virtual Router Redundancy Protocol version 3 (VRRPv3).

NOTE:

- Even though the `version-3` statement can be configured only at the `[edit protocols vrrp]` hierarchy level, VRRPv3 is enabled on all the configured logical systems as well.
- When enabling VRRPv3, you must ensure that VRRPv3 is enabled on all the VRRP routers in the network. This is because VRRPv3 does not interoperate with the previous versions of VRRP.

Default

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Junos OS Support for VRRPv3

version (Flow Aggregation)

IN THIS SECTION

- Syntax | 30693
- Hierarchy Level | 30693
- Description | 30693
- Options | 30693
- Required Privilege Level | 30693
- Release Information | 30693

Syntax

```
version format;
```

Hierarchy Level

```
[edit forwarding-options accounting name output flow-server hostname],
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server
hostname],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server hostname]
```

Description

Specify the version format of the aggregated flows exported to a cflowd server.

Options

format—Format of the flows.

- **Values:** 5 or 8
- **Default:** 5

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

export-format (Flow Monitoring) | 6667

Enabling Flow Aggregation

Directing Traffic Sampling Output to a Server Running the cflowd Application

version-alias (Dynamic Profiles)

IN THIS SECTION

- Syntax | 30694
- Hierarchy Level | 30694
- Description | 30694
- Options | 30695
- Required Privilege Level | 30695
- Release Information | 30695

Syntax

```
version-alias version-alias-string;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Configure an alias to identify a specific configuration variant of a base dynamic client profile. The version alias is a text description that lets you decide how to name different profile variants, so they have an identifier independent of the dynamic version name that is automatically created by the BNG.

The need for a version alias results from the practice of using a given base dynamic profile across multiple BNGs in a network. When you enable dynamic versioning, you can modify the base dynamic profile to provide specific capabilities to subscribers on different BNGs or to new subscribers on a given BNG. Dynamic versioning assigns a version name to each new variation in the base profile.

Consequently, the version name may vary for subscribers on one BNG or across multiple BNGs. In either case, RADIUS cannot determine which version of a profile is in use by any subscriber. This creates an operational challenge because RADIUS is unable to return corresponding attributes and VSAs in a CoA message that are compatible with that version of the profile.

When you configure a version alias for a dynamic client profile, the BNG sends the version alias to the RADIUS server during authentication. It is conveyed in the Juniper Networks client-profile-name VSA (26–4874–174). The version alias is an independent tag that enables you to track which profile variations are in use. Because RADIUS can tell apart the different profile versions, you can normalize the RADIUS back-end configuration for efficient use of CoA messages.

By default, the Client-Profile-Name VSA carries the name of the base dynamic profile. The version alias string is concatenated to the end of the profile name in the VSA, like this:

client-profile-name:*version-alias-string*

Options

version-alias-string	String that identifies a particular version of a dynamic profile. The maximum size of the string is 64 bytes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

Versioning for Dynamic Profiles

version (BFD)

IN THIS SECTION

- Syntax | **30696**
- Hierarchy Level | **30696**

- Description | 30696
- Options | 30696
- Required Privilege Level | 30697
- Release Information | 30697

Syntax

```
version (0 | 1 | automatic);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp oam bfd-liveness-detection],
[edit logical-systems logical-system-name protocols ldp oam fec address bfd-liveness-detection],
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group group-name liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd],
[edit protocols ldp oam bfd-liveness-detection],
[edit protocols ldp oam fec address bfd-liveness-detection]
```

Description

Configure the BFD protocol version to detect.

Options

0	Use BFD protocol version 0.
1	Use BFD protocol version 1.
automatic	Autodetect the BFD protocol version.

- **Default:** `automatic`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients

Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

Configuring BFD for LDP LSPs

version (BFD Liveness Detection)

IN THIS SECTION

- Syntax | 30697
- Hierarchy Level | 30698
- Description | 30698
- Options | 30698
- Required Privilege Level | 30698
- Release Information | 30699

Syntax

```
version (1 | automatic);
```

Hierarchy Level

```
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls)  oam bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)neighbor neighbor-id oam bfd-liveness-detection],
[edit logical-systems name tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-detection],
[edit routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) oam bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls) neighbor neighbor-id oam bfd-liveness-detection],
[edit tenants name routing-instances name protocols (evpn | l2vpn | vpls)mesh-group mesh-group-name neighbor neighbor-id oam bfd-liveness-detection]
```

Description

Specify the BFD version for detection. You can explicitly configure BFD version 1 or the routing device can automatically detect the BFD version. By default, the routing device automatically detects the BFD version.

Options

Configure the BFD version to detect: `1` (BFD version 1) or `automatic` (autodetect the BFD version).

- **Default:** The routing device automatically detects the BFD version.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1

RELATED DOCUMENTATION

Configuring BFD for Layer 2 VPN and VPLS

Example: Configuring BFD Authentication for BGP

Example: Configuring BFD on Internal BGP Peer Sessions

Understanding BFD Authentication for BGP

version (Dynamic IGMP Interface)

IN THIS SECTION

- Syntax | 30699
- Hierarchy Level | 30699
- Description | 30700
- Options | 30700
- Required Privilege Level | 30700
- Release Information | 30700

Syntax

```
version version;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmpinterface interface-name]
```

Description

Specify the version of IGMP.

Options

version—IGMP version number.

- **Range:** 1, 2, or 3
- **Default:** IGMP version 2

> **NOTE**: Routers running different versions of IGMP negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version.
>
> If you have already configured the router to use IGMP version 1 and then configure it to use IGMP version 2, the router continues to use IGMP version 1 for up to 6 minutes and then uses IGMP version 2.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview

Configuring Dynamic DHCP Client Access to a Multicast Network

Changing the IGMP Version

version (Dynamic MLD Interface)

IN THIS SECTION

- Syntax | 30701
- Hierarchy Level | 30701
- Description | 30701
- Options | 30701
- Required Privilege Level | 30702
- Release Information | 30702

Syntax

```
version version;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Description

Configure the MLD version explicitly on the dynamic interface. MLD version 2 (MLDv2) is used only to support source-specific multicast (SSM).

Options

version—MLD version to run on the interface.

- **Range:** 1 or 2
- **Default:** 1 (MLDv1)

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview

Modifying the MLD Version

version (Flow Monitoring Logs for NAT)

IN THIS SECTION

- Syntax | 30702
- Hierarchy Level | 30703
- Description | 30703
- Options | 30703
- Required Privilege Level | 30703
- Release Information | 30703

Syntax

```
version (ipfix | v9);
```

Hierarchy Level

```
[edit services jflow-log template-profile template-profile-name]
```

Description

Specify the flow template format, such as IPFIX or version 9, to be used for generating flow monitoring records for NAT events and for transmitting them to the collector.

Options

`ipfix`—Use the IPFIX flow template format for flow monitoring logs for NAT events.

`v9`—Use the version 9 flow template format for flow monitoring logs for NAT events.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

version (IGMP Snooping)

IN THIS SECTION

- Syntax | 30704
- Hierarchy Level | 30704
- Description | 30704
- Options | 30705
- Default | 30705
- Required Privilege Level | 30705
- Release Information | 30705

Syntax

```
version version;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping],
[edit routing-instances routing-instance-name protocols igmp-snooping],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-
snooping],
[edit routing-instances routing-instance-name protocols igmp-snooping vlan vlan-id],
[edit protocols igmp-snooping vlan vlan-id],
[edit logical-systems logical-system-name protocols igmp-snooping],
[edit protocols igmp-snooping]
```

Description

Specify the IGMP version for the IGMP general query that the snooping device sends to hosts when an interface comes up or a new member is added to a VLAN or bridge domain. If the L2 querier is

configured for a VLAN/bridge-domain, periodic queries are sent to all members with the specified version of IGMP-snooping.

To configure the IGMP snooping version for *all* VLANs in the default instance:

```
set protocols igmp-snooping vlan all version version
```

To configure the IGMP snooping version for *all* VLANs in a specified routing instance:

```
set routing-instances routing-instance name protocols igmp-snooping vlan all version version
```

Options

version—IGMP version number.

Range: 2 and 3.

Default

If you do not configure the version statement, the default is IGMPv3.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1

RELATED DOCUMENTATION

Configuring IGMP or MLD Snooping Version

version (IKE)

IN THIS SECTION

- Syntax | 30706
- Hierarchy Level | 30706
- Description | 30706
- Options | 30706
- Required Privilege Level | 30707
- Release Information | 30707

Syntax

```
version ( 1 | 2);
```

Hierarchy Level

```
[edit services ipsec-vpn ike policy policy-name]
```

Description

Configure the Internet Key Exchange (IKE) version that is used to negotiate dynamic SAs for IPSec.

Options

1—Uses IKEv1.

2—Uses IKEv2.

NOTE: By default, Junos OS uses IKE policy version 1.0. Version 2.0 is supported only in Junos OS Release 11.4 and later. If no version is explicitly configured, Junos OS sets the version to version 1.0.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Configuring IKE Policies

version-ipfix (Forwarding Options)

IN THIS SECTION

- Syntax | 30707
- Hierarchy Level | 30708
- Description | 30708
- Required Privilege Level | 30708
- Release Information | 30708

Syntax

```
version-ipfix {
    template template-name;
}
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family  (inet | inet6 | mpls | vpls)
output flow-server hostname]
```

Description

Specify flow monitoring version IPFIX properties to apply to output sampling records.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices

version-ipfix (Services)

IN THIS SECTION

- Syntax | 30709
- Hierarchy Level | 30710
- Description | 30710
- Required Privilege Level | 30710

- Release Information | **30710**

Syntax

```
version-ipfix {
    template template-name {
        data-record-fields {
            source-prefix-as-path count;
            destination-prefix-as-path count;
            bgp-source-standard-community count;
            bgp-destination-standard-community count;
            bgp-source-extended-community count;
            bgp-destination-extended-community count;
            bgp-source-large-community count;
            bgp-destination-large-community count;
        }
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        flow-key {
            flow-direction;
            vlan-id;
            output-interface;
        }
        (ipv4-template | ipv6-template | mpls-ipv4-template | mpls-ipvx-template | vpls-template);
        multi-bgp-path;
        nexthop-learning (enable |disable);
        observation-domain-id id;
        option-refresh-rate packets packets seconds seconds;
        options-template-id id;
        template-id id;
        template-refresh-rate packets packets seconds seconds;
        tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
    }
}
```

Hierarchy Level

```
[edit services flow-monitoring]
```

Description

Specify the IPFIX output template properties to support flow monitoring.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

`data-record-fields` option introduced in Junos OS Evolved Release 21.4R1.

`multi-bgp-path` option added in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices

Configuring Inline Active Flow Monitoring on PTX Series Routers

version (Liveness Detection)

IN THIS SECTION

Syntax | 30711

- Hierarchy Level | 30711
- Description | 30711
- Options | 30711
- Required Privilege Level | 30711
- Release Information | 30711

Syntax

```
version (1 | automatic);
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection]
```

Description

Configure the Bidirectional Forwarding Detection (BFD) protocol version to detect.

Options

1	Use BFD protocol version 1.
automatic	Autodetect the BFD protocol version.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

version (MLD Snooping)

IN THIS SECTION

- Syntax | 30712
- Hierarchy Level | 30712
- Description | 30712
- Options | 30713
- Default | 30713
- Required Privilege Level | 30713
- Release Information | 30713

Syntax

```
version version;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols mld-snooping],
[edit routing-instances routing-instance-name protocols mld-snooping],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols mld-
snooping],
[edit routing-instances routing-instance-name protocols mld-snooping vlan vlan-id],
[edit protocols mld-snooping vlan vlan-id],
[edit logical-systems logical-system-name protocols mld-snooping],
[edit protocols mld-snooping]
```

Description

Specify the MLD version for the MLD general query that the snooping device sends to hosts when an interface comes up or a new member is added to a VLAN or bridge domain. If the L2 querier is

configured for a VLAN/bridge-domain, periodic queries are sent to all members with the specified version of MLD-snooping.

To configure the MLD snooping version for *all* VLANs in the default instance:

```
set protocols mld-snooping vlan all version version
```

To configure the MLD snooping version for *all* VLANs in a specified routing instance:

```
set routing-instances routing-instance name protocols mld-snooping vlan all version version
```

Options

version—MLD version number.

Range: 1 and 2.

Default

If you do not configure the version statement, the default is MLDv2.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1

RELATED DOCUMENTATION

Configuring IGMP or MLD Snooping Version

version (Protocols BFD)

IN THIS SECTION

- Syntax | 30714
- Hierarchy Level | 30714
- Description | 30714
- Options | 30714
- Required Privilege Level | 30714
- Release Information | 30715

Syntax

```
version (0 | 1 | automatic);
```

Hierarchy Level

```
[edit protocols piminterface (Protocols PIM)  interface-name bfd-liveness-detection],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM)  interface-name bfd-liveness-detection]
```

Description

Specify the bidirectional forwarding detection (BFD) protocol version that you want to detect.

Options

Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version)

- **Default:** automatic

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring BFD for PIM

version (Protocols IGMP)

IN THIS SECTION

- Syntax | 30715
- Hierarchy Level | 30715
- Description | 30716
- Options | 30716
- Required Privilege Level | 30716
- Release Information | 30716

Syntax

```
version version;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp interface interface-name],
[edit protocols igmp interface interface-name]
[edit protocols igmp-snooping vlan                (all | vlan-name)]
```

Description

Specify the version of IGMP.

Options

version—IGMP version number.

- **Range:** 1, 2, or 3
- **Default:** IGMP version 2

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Changing the IGMP Version

version (Protocols IGMP AMT)

IN THIS SECTION

- Syntax | 30717
- Hierarchy Level | 30717
- Description | 30717
- Options | 30717
- Required Privilege Level | 30717
- Release Information | 30717

Syntax

```
version version;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols igmp amt relay defaults],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols igmp
amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances routing-instance-name protocols igmp amt relay defaults]
```

Description

Specify the version of IGMP used through an Automatic Multicast Tunneling (AMT) interface.

Options

version—IGMP version number.

- **Range:** 1, 2, or 3
- **Default:** IGMP version 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

> *Configuring Default IGMP Parameters for AMT Interfaces*

version (Protocols MLD)

IN THIS SECTION

- Syntax | 30718
- Hierarchy Level | 30718
- Description | 30718
- Options | 30718
- Required Privilege Level | 30719
- Release Information | 30719

Syntax

```
version version;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mld interface interface-name],
[edit protocols mld interface interface-name]
```

Description

Configure the MLD version explicitly. MLD version 2 (MLDv2) is used only to support source-specific multicast (SSM).

Options

version—MLD version to run on the interface.

- **Range:** 1 or 2
- **Default:** 1 (MLDv1)

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Modifying the MLD Version

version (Protocols PIM)

IN THIS SECTION

- Syntax | 30719
- Hierarchy Level | 30720
- Description | 30720
- Options | 30720
- Required Privilege Level | 30720
- Release Information | 30720

Syntax

```
version version;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim interface (Protocols PIM) interface-name],
[edit logical-systems logical-system-name protocols pim rp static address address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
interface (Protocols PIM) interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim
rp static address address],
[edit protocols pim interface (Protocols PIM) interface-name],
[edit protocols pim rp static address address],
[edit routing-instances routing-instance-name protocols pim interface (Protocols PIM) interface-
name],
[edit routing-instances routing-instance-name protocols pim rp static address address]
```

Description

Starting in Junos OS Release 16.1, it is no longer necessary to specify a PIM version. PIMv1 is being obsoleted so the version choice is moot.

Options

version—PIM version number.

- **Range:** See the Description, above.
- **Default:** PIMv2 for both rendezvous point (RP) mode (at the `[edit protocols pim rp static address address]` hierarchy level). and interface mode (at the `[edit protocols pim interface interface-name]` hierarchy level).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement deprecated (hidden) in Junos OS Release 16.1 for later removal.

RELATED DOCUMENTATION

Enabling PIM Sparse Mode

Configuring PIM Dense Mode Properties

Configuring PIM Sparse-Dense Mode Properties

version (SNMP)

IN THIS SECTION

- Syntax | 30721
- Hierarchy Level | 30721
- Description | 30721
- Default | 30722
- Options | 30722
- Required Privilege Level | 30722
- Release Information | 30722

Syntax

```
version (all | v1 | v2);
```

Hierarchy Level

```
[edit snmp trap-group group-name]
```

Description

Specify the version number of SNMP traps.

Default

all—Send an SNMPv1 and SNMPv2 trap for every trap condition.

Options

all—Send an SNMPv1 and SNMPv2 trap for every trap condition.

v1—Send SNMPv1 traps only.

v2—Send SNMPv2 traps only.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

version9 (Flow Monitoring)

IN THIS SECTION

- Syntax | 30722
- Hierarchy Level | 30723
- Description | 30723
- Required Privilege Level | 30723
- Release Information | 30723

Syntax

```
version9 {
    template template-name {
```

```
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        flow-key {
            flow-direction;
            vlan-id;
            output-interface;
        }
        (ipv4-template | ipv6-template | mpls-template label-position [ positions ] | mpls-ipv4-
template label-position [ positions ] | mpls-ipvx-template);
        multi-bgp-path;
        option-refresh-rate packets packets seconds seconds;
        options-template-id id;
        peer-as-billing-template;
        source-id id;
        template-id id;
        template-refresh-rate packets packets seconds seconds;
        tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
    }
}
```

Hierarchy Level

```
[edit services flow-monitoring]
```

Description

Specify the version 9 output template properties to support flow monitoring.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Statement introduced in Junos OS Release 19.2R1 for MX Series routers with MPC10E-15C-MRATE line card to define a flow record template suitable for IPv4 or IPv6 traffic only.

`multi-bgp-path` option added in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

> *Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates*
>
> *Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250*

version9 (Forwarding Options)

IN THIS SECTION

- Syntax | 30724
- Hierarchy Level | 30724
- Description | 30725
- Required Privilege Level | 30725
- Release Information | 30725

Syntax

```
version9 {
    template template-name;
}
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (bridge | inet |inet6 |mpls vpls)
output flow-server hostname],
```

```
[edit forwarding-options sampling family (inet |inet6 |mpls| vpls |bridge) output flow-server
hostname]
```

Description

Specify flow monitoring version 9 properties to apply to output sampling records.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the following hierarchy levels introduced in Junos OS Release 18.2R1: `[edit forwarding-options sampling instance instance-name family bridge]`, `[edit forwarding-options sampling instance instance-name family vpls]`, `[edit forwarding-options sampling family bridge]`, and `[edit forwarding-options sampling family vpls]`.

Statement introduced in Junos OS Release 19.2R1 for MX Series routers with MPC10E-15C-MRATE line card to define a flow record template suitable for IPv4 or IPv6 traffic only.

RELATED DOCUMENTATION

> *Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates*

versioning

IN THIS SECTION

- Syntax | 30726
- Hierarchy Level | 30726

- Description | 30726
- Required Privilege Level | 30726
- Release Information | 30726

Syntax

```
versioning;
```

Hierarchy Level

```
[edit system dynamic-profile-options]
```

Description

Enable version support for dynamic profiles on the system. This means that you can create new versions of dynamic profiles that are currently in use by subscribers. Each new version of the profile is stored as a new profile. The profile name is in the format *base-profile-name*$$*number*, where *number* varies sequentially from 01 to 10. There can be up to 10 modified versions of a base profile.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Versioning for Dynamic Profiles

video

IN THIS SECTION

- Syntax | 30727
- Hierarchy Level | 30727
- Description | 30727
- Required Privilege Level | 30727
- Release Information | 30728

Syntax

```
video {
    dscp (Services CoS) (alias | bits);
    forwarding-class (Services PIC Classifiers) class-name;
}
```

Hierarchy Level

```
[edit services (CoS) cos application-profile profile-name sip]
```

Description

Set the appropriate `dscp` and `forwarding-class` values for SIP video traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Application Profiles for Use as CoS Rule Actions

video (Application Profile)

IN THIS SECTION

- Syntax | 30728
- Hierarchy Level | 30728
- Description | 30729
- Default | 30729
- Required Privilege Level | 30729
- Release Information | 30729

Syntax

```
video {
    dscp (alias | bits);
    forwarding-class class-name;
}
```

Hierarchy Level

```
[edit services cos application-profileprofile-name sip]
```

Description

Set the appropriate `dscp` and `forwarding-class` values for SIP video traffic.

Default

By default, the system will not alter the DSCP or forwarding class for SIP video traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

voice (Application Profile) | 30917

video-monitoring

IN THIS SECTION

- Syntax | 30730
- Hierarchy Level | 30732
- Description | 30732
- Required Privilege Level | 30732
- Release Information | 30732

Syntax

```
video-monitoring {
    interfaces {
        interface-name {
            family {
                inet {
                    input-flows {
                        input-flow-name {
                            destination-address [ address ];
                            destination-port [ port ];
                            source-address [ address ];
                            source-port [ port ];
                            template template-name;
                        }
                    }
                    output-flows {
                        output-flow-name {
                            destination-address [ address ];
                            destination-port [ port ];
                            source-address [ address ];
                            source-port [ port ];
                            template template-name;
                        }
                    }
                }
                inet6 {
                    input-flows {
                        input-flow-name {
                            destination-address [ address ];
                            destination-port [ port ];
                            source-address [ address ];
                            source-port [ port ];
                            template template-name;
                        }
                    }
                    output-flows {
                        output-flow-name {
                            destination-address [ address ];
                            destination-port [ port ];
                            source-address [ address ];
                            source-port [ port ];
```

```
                        template template-name;
                    }
                }
            }
            mpls {
                input-flows {
                    input-flow-name {
                        (destination-address [ address ] | source-address [ address ]);
                        destination-port [ port ];
                        payload-type (ipv4 | ipv6);
                        source-port [ port ];
                        template template-name;
                    }
                }
                output-flows {
                    output-flow-name {
                        (destination-address [ address ] | source-address [ address ]);
                        destination-port [ port ];
                        payload-type (ipv4 | ipv6);
                        source-port [ port ];
                        template template-name;
                    }
                }
            }
        }
    }
}
templates {
    template-name {
        interval-duration interval-duration;
        inactive-timeout inactive-timeout;
        rate {
            (layer3 layer3-packets-per-second | media media-bits-per-second);
        }
        delay-factor {
            disable;
            threshold {
                (info | warning | critical) delay-factor-threshold;
            }
        }
        media-loss-rate {
            disable;
            threshold {
```

```
                    (info | warning | critical) percentage mlr-percentage | packet-count mlr-
packet-count);
                }
            }
            media-rate-variation {
                ;
                threshold {
                    (info |warning | critical) mrv-variation;
                }
            }
            media-packets-count-in-layer3 media-packets-count-in-layer3;
            media-packet-size media-packet-size;
        }
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Define the options for video monitoring using media delivery index options for metrics.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring Inline Video Monitoring on MX Series Routers

view (Configuring a MIB View)

IN THIS SECTION

- Syntax | 30733
- Hierarchy Level | 30733
- Description | 30733
- Options | 30734
- Required Privilege Level | 30734
- Release Information | 30734

Syntax

```
view view-name {
    oid object-identifier <include | exclude>;
}
```

Hierarchy Level

```
[edit snmp]
```

Description

Define a MIB view. A MIB view identifies a group of MIB objects. Each MIB object in a view has a common OID prefix. Each object identifier represents a subtree of the MIB object hierarchy. The `view` statement uses a view to specify a group of MIB objects on which to define access. To enable a view, you must associate the view with a community by including the `view` statement at the `[edit snmp community community-name]` hierarchy level.

NOTE: To remove an OID completely, use the `delete view all oid oid-number` command but omit the `include` parameter.

Options

view-name—Name of the view.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

community (SNMP) | 4475

view (SNMP Community)

IN THIS SECTION

- Syntax | 30735
- Hierarchy Level | 30735
- Description | 30735
- Options | 30735
- Required Privilege Level | 30735
- Release Information | 30735

Syntax

```
view view-name;
```

Hierarchy Level

```
[edit snmp community community-name]
```

Description

Associate a view with a community. A view represents a group of MIB objects.

Options

view-name—Name of the view. You must use a view name already configured in the `view` statement at the `[edit snmp]` hierarchy level.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

violation-action (DHCP Local Server and DHCP Relay Agent)

IN THIS SECTION

- Syntax | 30736
- Hierarchy Level | 30736
- Description | 30736

- Options | 30736
- Required Privilege Level | 30737
- Release Information | 30737

Syntax

```
violation-action action;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay lease-time-validation],
[edit forwarding-options dhcp-relay dhcpv6 lease-time-validation],
[edit forwarding-options dhcp-relay group group-name lease-time-validation],
[edit forwarding-options dhcp-relay dhcpv6 group group-name lease-time-validation],
[edit logical-systems logical-system-name ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name...],
[edit routing-instances routing-instance-name ...],
[edit system services dhcp-local-server lease-time-validation],
[edit system services dhcp-local-server dhcpv6 lease-time-validation],
[edit system services dhcp-local-server group group-name lease-time-validation],
[edit system services dhcp-local-server dhcpv6 group group-name lease-time-validation]
```

Description

Configure the action that the router performs when a DHCP lease-time violation occurs. The violation occurs when a third-party DHCP server or address-assignment pool offers a DHCP lease time that is less than the threshold specified by the `lease-time-threshold` statement.

Options

action Action taken by the router when a lease-time violation occurs.

- `drop`—(Optional) For DHCPv4 and DHCPv6 relay agent, the third-party lease is dropped and the client binding fails.
- `override-lease`—(Optional) For DHCPv4 and DHCPv6 local server, the third-party lease is overridden with the value specified by the `lease-time-threshold` statement and binds the client using the new value.
- `strict`—(Optional) For DHCPv4 and DHCPv6 local server, DHCP ignores the third-party lease and the client binding fails.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

DHCP Lease Times for IP Addresses

violation-report-rate (DDoS Flow Detection)

IN THIS SECTION

- Syntax | 30738
- Hierarchy Level | 30738
- Description | 30738
- Options | 30738
- Required Privilege Level | 30738
- Release Information | 30738

Syntax

```
violation-report-rate report-rate;
```

Hierarchy Level

```
[edit system ddos-protection global]
```

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Limit the rate at which bandwidth violations (violation reports) are reported from an FPC to the Routing Engine, for all protocol groups and packet types on all line cards.

Options

report-rate Number of violations per second.

- **Range:** 1 through 50,000
- **Default:** 100

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

> *Configuring the Violation Reporting Rate for All Protocol Groups and Packet Types*
>
> *Setting Up and Using Flow Detection*

virtual-address

IN THIS SECTION

- Syntax | 30739
- Hierarchy Level | 30739
- Description | 30739
- Options | 30739
- Required Privilege Level | 30740
- Release Information | 30740

Syntax

```
virtual-address [ addresses ];
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-
group group-id],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet address address vrrp-group group-id]
```

Description

Configure the addresses of the devices in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group. You can configure up to eight addresses.

Options

addresses Addresses of one or more devices. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the primary device for the group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring Basic VRRP Support
- *Understanding VRRP on SRX Series Devices*
- *Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces*

virtual-channel-group (CoS Interfaces)

IN THIS SECTION

- Syntax | 30740
- Hierarchy Level | 30741
- Description | 30741
- Options | 30741
- Required Privilege Level | 30741
- Release Information | 30741

Syntax

```
virtual-channel-group group-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Assign a virtual channel group to a logical interface.

If you apply a virtual channel group to multiple logical interfaces, separate queues are created on each of the interfaces. The same virtual channel names are used on all the interfaces. You can specify the scheduler and shaping rates in the virtual channels in percentages so that you can apply the same virtual channel group to logical interfaces with different available bandwidths.

Options

group-name—Name of the virtual channel group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

default (CoS) | **5187**

scheduler-map (CoS Virtual Channels) | **16350**

shaping-rate (CoS Virtual Channels) | **16936**

virtual-channels | **30743**

virtual-channel-groups | **30742**

virtual-channel-groups

IN THIS SECTION

- Syntax | 30742
- Hierarchy Level | 30742
- Description | 30742
- Options | 30743
- Required Privilege Level | 30743
- Release Information | 30743

Syntax

```
virtual-channel-groups {
    virtual-channel-group-name {
        virtual-channel-name {
            scheduler-map map-name;
            shaping-rate (percent percentage | rate);
            default;
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Associate a virtual channel with a scheduler map and a shaping rate. Virtual channels and virtual channel groups enable you to direct traffic into a virtual channel and apply bandwidth limits to the channel.

Options

group-name—Name of the virtual channel group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

- default (CoS) | 5187
- scheduler-map (CoS Virtual Channels) | 16350
- shaping-rate (CoS Virtual Channels) | 16936
- virtual-channel-group (CoS Interfaces) | 30740
- virtual-channels | 30743

virtual-channels

IN THIS SECTION

- Syntax | 30744
- Hierarchy Level | 30744
- Description | 30744
- Options | 30744
- Required Privilege Level | 30744
- Release Information | 30744

Syntax

```
virtual-channels {
    virtual-channel-name;
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify a list of virtual channels.

Each virtual channel has eight transmission queues.

Options

virtual-channel-name—Name of the virtual channel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

RELATED DOCUMENTATION

default (CoS) | **5187**

scheduler-map (CoS Virtual Channels) | **16350**

shaping-rate (CoS Virtual Channels) | **16936**

virtual-channel-group (CoS Interfaces) | **30740**

virtual-channel-groups | **30742**

virtual-chassis

IN THIS SECTION

- Syntax | 30745
- Hierarchy Level | 30746
- Description | 30746
- Default | 30746
- Required Privilege Level | 30746
- Release Information | 30747

Syntax

```
virtual-chassis {
    aliases {
        serial-number serial-number {
            alias-name alias-name;
        }
    }
    auto-conversion;
    auto-provisioned;
    auto-sw-update {
        (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
        package-name package-name;
    }
    fast-failover (ge | vcp disable | xe);
    graceful-restart {
        disable;
    }
    id id;
    mac-persistence-timer (minutes | disable);
    member member-id {
        fabric-tree-root;
        location location;
        mastership-priority number;
        no-management-vlan;
        serial-number;
```

```
        role;
    }
    no-auto-conversion;
    no-split-detection;
    preprovisioned;
    traceoptions (Virtual Chassis) {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
regex>;
        flag flag ;
    }
    vc-port {
        lag-hash (packet-based | source-port-based);
    }
    vcp-no-hold-time;
    vcp-snmp-statistics;
}
```

Hierarchy Level

```
[edit]
```

Description

Configure a Virtual Chassis or a Virtual Chassis Fabric (VCF).

The remaining statements are explained separately. See CLI Explorer.

Default

A standalone EX Series switch is a Virtual Chassis by default. It has a default member ID of 0, a default primary-role priority of 128, and a default role as primary.

A QFX Series device configured in standalone mode is a Virtual Chassis by default. It has a default member ID of 0, a default primary-role priority of 128, and a default role as primary.

A standalone XRE200 External Routing Engine or EX8200 switch is not part of an EX8200 Virtual Chassis until a Virtual Chassis configuration is set up.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- *Adding a Device to a Virtual Chassis Fabric*
- Configuring an EX4650 or a QFX Series Virtual Chassis
- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis
- Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis
- Configuring an EX9200 Virtual Chassis

virtual-chassis (MX Series Virtual Chassis)

IN THIS SECTION

- Syntax | 30747
- Hierarchy Level | 30748
- Description | 30748
- Required Privilege Level | 30748
- Release Information | 30749

Syntax

```
virtual-chassis {
    aliases {
        serial-number serial-number;
```

```
    }
    graceful-restart {

        disable;
    }
    heartbeat-address (ip-address | ipv6-address);
    heartbeat-timeout seconds;
    heartbeat-tos seconds;
    heartbeat-tos value;
    locality-bias;
    member member-id {
        role (routing-engine | line-card);
        serial-number serial-number;
    }
    no-split-detection;
    preprovisioned;
    traceoptions {
        file filename <files number> <match regular-expression> <no-stamp> <replace> <size
maximum-file-size> <world-readable | no-world-readable>;
        flag flag <detail> <disable> <receive> <send>;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Create a Virtual Chassis configuration for MX Series routers.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Example: Configuring Interchassis Redundancy for MX Series 5G Universal Routing Platforms Using a Virtual Chassis

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in the Same Subnet

Example: Determining Member Health Using an MX Series Virtual Chassis Heartbeat Connection with Member Routers in Different Subnets

Configuring Preprovisioned Member Information for a Virtual Chassis

Configuring Locality Bias for a Virtual Chassis

virtual-circuit-id

IN THIS SECTION

- Syntax | **30749**
- Hierarchy Level | **30750**
- Description | **30750**
- Options | **30750**
- Required Privilege Level | **30750**
- Release Information | **30750**

Syntax

```
virtual-circuit-id identifier;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit local-switching interface
interface-name backup-neighbor address],
[edit logical-systems logical-system-name protocols l2circuit neighbor address
interface interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface
interface-name backup-neighbor address],
[edit protocols l2circuit local-switching interface interface-name backup-neighbor address],
[edit protocols l2circuit neighbor address interface interface-name],
[edit protocols l2circuit neighbor address interface interface-name backup-neighbor address]
```

Description

Uniquely identify a Layer 2 circuit for either a standard pseudowire or a redundant pseudowire.

Options

identifier—1 through 4,294,967,295

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Hierarchy levels for `backup-neighbor` (pseudowire redundancy) added in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring the Virtual Circuit ID

Configuring Pseudowire Redundancy on the PE Router

Example: Configuring Layer 2 Circuit Switching Protection

virtual-control-channel

IN THIS SECTION

- Syntax | 30751
- Hierarchy Level | 30751
- Description | 30751
- Options | 30751
- Required Privilege Level | 30752
- Release Information | 30752

Syntax

```
virtual-control-channel channel-name {
    west-interface name;
    east-interface name;
}
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring name (east-interface | west-interface)]
```

Description

Specify virtual control channels which are logical interfaces on the east and west interfaces of the major ring.

Options

`west-interface` *`name`*—Logical interface on the major ring's west port.

`east-interface` *`name`*—Logical interface on the major ring's east port.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

- Ethernet Ring Protection Switching Overview
- Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

virtual-cpu

IN THIS SECTION

- Syntax | 30752
- Hierarchy Level | 30753
- Description | 30753
- Options | 30753
- Required Privilege Level | 30753
- Release Information | 30753

Syntax

```
virtual-cpu {
    virtual-cpu-number {
        physical-cpu number | range;
    }
    count number;
    features {
```

```
        hardware-virtualization;
    }
}
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Specify the number of virtual CPUs the VNF can use. By default, a VNF is assigned one virtual CPU, which is independent of any specific physical CPU.

Options

virtual-cpu count *number* Number of virtual CPUs. For example, 2.

virtual-cpu features Features of the virtual cpu. For example, hardware-virtualization.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D40.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | **27723**

virtual-gateway-address

IN THIS SECTION

- Syntax | 30754
- Hierarchy Level | 30754
- Description | 30754
- Options | 30755
- Required Privilege Level | 30755
- Release Information | 30755

Syntax

```
virtual-gateway-address address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family address address]
```

Description

Set the default IPv4 or IPv6 address for the gateway for end hosts. You eliminate the need to proxy for remote gateway IP addresses because you must configure the virtual gateway address using the same IP address as on all of the provider edge (PE) devices. This configuration internally generates the same Virtual Router Redundancy Protocol (VRRP) MAC. You can set a corresponding virtual gateway address for each logical integrated routing and bridging (IRB) interface. A maximum of 64 PEs can have the same virtual gateway address.

To support pinging on the virtual gateway IP address, you must include both the `virtual-gateway-accept-data` statement and the `preferred` statement at the `[edit interfaces irb unit]` hierarchy of the preferred virtual gateway.

Options

address—virtual gateway address. You can't specify the addresses 0.0.0.0 (default route address) or 255.255.255.255 (broadcast IP address).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

EVPN with IRB Solution Overview

Example: Configuring EVPN with IRB Solution

Configuring the Interface Address

virtual-gateway-v4-mac

IN THIS SECTION

- Syntax | 30756
- Hierarchy Level | 30756
- Description | 30756
- Options | 30757
- Required Privilege Level | 30757
- Release Information | 30757

Syntax

```
virtual-gateway-v4-mac ipv4-mac-address
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit logical-unit-number],
[edit dynamic-profiles name logical-systems name interfaces name unit logical-unit-number],
[edit interfaces name unit logical-unit-number]
```

Description

Explicitly configure an IPv4 media access control (MAC) address for a default virtual gateway.

A Juniper Networks device creates a default virtual gateway when you specify a virtual gateway address (VGA) for an integrated routing and bridging (IRB) interface on the device. You use this option on a device that functions as a Layer 3 gateway in an EVPN-VXLAN or an EVPN-MPLS centrally-routed bridging (CRB) overlay fabric.Through its IRB interface, the default virtual gateway enables communication between non-virtualized hosts, virtual machines (VMs), and servers in different VXLANs, MPLS networks, or IP subnetworks.

When you configure a VGA for an IRB interface, the Layer 3 gateway automatically generates IPv4 MAC address 00:00:5e:00:01:01 for that particular virtual gateway. In this topic, we refer to the virtual gateway MAC address as a *virtual MAC*. The automatically generated virtual MAC is not included as the source MAC address in packets generated by the Layer 3 gateway. Instead, the device includes the MAC address for the IRB interface in:

- Data packets
- The source MAC address field in the outer Ethernet header of:
  - Address Resolution Protocol (ARP) replies
  - Neighbor advertisement packets

In this topic, we refer to the MAC address for the IRB interface as the *IRB MAC*. When an ARP reply includes the IRB MAC as the source MAC address instead of the virtual MAC, in CRB overlays you might see unknown unicast packet flooding throughout the domain.

If you explicitly configure a MAC address for a default virtual gateway, the configured virtual MAC overrides the automatically generated virtual MAC. That is, when the Layer 3 gateway sends data packets, ARP replies, and neighbor advertisement packets, the configured virtual MAC is in the outer

Ethernet header of these packets. As a result, explicitly configuring the virtual MAC eliminates the possibility of flooding the domain with unknown unicast packets.

00

> **NOTE**: The MAC address range 02:00:00:00:00:00 through 02:00:00:00:00:FF is used for internal communication. Don't use addresses in this range for a manual virtual MAC assignment.

For more information about the flooding issue and its resolution, see *Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN Overlay Network*.

Options

ipv4-mac-address	IPv4 MAC address for the default virtual gateway.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R5.

RELATED DOCUMENTATION

virtual-gateway-v6-mac | 30757

virtual-gateway-v6-mac

IN THIS SECTION

- Syntax | 30758
- Hierarchy Level | 30758

- Description | 30758
- Options | 30759
- Required Privilege Level | 30759
- Release Information | 30759

Syntax

```
virtual-gateway-v6-mac ipv6-mac-address
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit logical-unit-number],
[edit dynamic-profiles name logical-systems name interfaces name unit logical-unit-number],
[edit interfaces name unit logical-unit-number]
```

Description

Explicitly configure an IPv6 media access control (MAC) address for a default virtual gateway.

A Juniper Networks device creates a default virtual gateway when you specify a virtual gateway address (VGA) for an integrated routing and bridging (IRB) interface on the device. You use this option on a device that functions as a Layer 3 gateway in an EVPN-VXLAN or an EVPN-MPLS centrally-routed bridging (CRB) overlay fabric. Through its IRB interface, the default virtual gateway enables communication between non-virtualized hosts, virtual machines (VMs), and servers in different VXLANs, MPLS networks, or IP subnetworks.

When you configure a VGA for an IRB interface, the Layer 3 gateway automatically generates IPv6 MAC address 00:00:5E:00:02:01 for that particular virtual gateway. In this topic, we refer to the virtual gateway MAC address as a *virtual MAC*. The automatically generated virtual MAC is not included as the source MAC address in packets generated by the Layer 3 gateway. Instead, the device includes the MAC address for the IRB interface in:

- Data packets
- The source MAC address field in the outer Ethernet header of:
  - Address Resolution Protocol (ARP) replies

- Neighbor advertisement packets

In this topic, we refer to the MAC address for the IRB interface as the *IRB MAC*. When an ARP reply includes the IRB MAC as the source MAC address instead of the virtual MAC, in CRB overlays you might see unknown unicast packet flooding throughout the domain.

If you explicitly configure a MAC address for a default virtual gateway, the configured virtual MAC overrides the automatically generated virtual MAC. That is, when the Layer 3 gateway sends data packets, ARP replies, and neighbor advertisement packets, the configured virtual MAC is in the outer Ethernet header of these packets. As a result, explicitly configuring the virtual MAC eliminates the possibility of flooding the domain with unknown-unicast packets.

> **NOTE**: The MAC address range 02:00:00:00:00:00 through 02:00:00:00:00:FF is used for internal communication. Don't use addresses in this range for a manual virtual MAC assignment.

For more information about the flooding issue and its resolution, see *Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN Overlay Network*.

Options

ipv6-mac-address	IPv6 MAC address for the default virtual gateway.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R5.

RELATED DOCUMENTATION

virtual-gateway-v4-mac | 30755

virtual-inet6-address

IN THIS SECTION

- Syntax | 30760
- Hierarchy Level | 30760
- Description | 30760
- Options | 30760
- Required Privilege Level | 30761
- Release Information | 30761

Syntax

```
virtual-inet6-address [ addresses ];
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
```

Description

Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You can configure up to eight addresses.

Options

addresses—Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the primary virtual router for the group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Basic VRRP Support

virtual-ip

IN THIS SECTION

- Syntax | 30761
- Hierarchy Level | 30762
- Description | 30762
- Options | 30762
- Required Privilege Level | 30762
- Release Information | 30762

Syntax

```
virtual-ip name {
 interface interface;
 ip ip;
 use-virtual-mac;
}
```

Hierarchy Level

```
[edit chassis high-availability services-redundancy-group]
```

Description

Specify IP address used for activeness determination and enforcement on the switching side of Multinode High Availability deployments. Required for hybrid and default gateway deployments.

Options

name Virtual IP Identifier.

interface Name of the interface for virtual IP.

ip IPV4/IPV6 prefix, should be in the same subnet as the interface IP

use-virtual-mac Use virtual MAC (vMAC) address for services redundancy group role enforcement. Virtual MAC address dynamically assigned to the interface on active node that faces the switching side.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.3R1.

RELATED DOCUMENTATION

Multinode High Availability

Example: Configure Multinode High Availability in a Default Gateway Deployment

virtual-link

IN THIS SECTION

- Syntax | 30763
- Hierarchy Level | 30763
- Description | 30764
- Options | 30764
- Required Privilege Level | 30764
- Release Information | 30764

Syntax

```
virtual-link neighbor-id router-id transit-area area-id {
    disable;
    authentication key <key-id identifier>;
    dead-interval seconds;
    hello-interval seconds;
    ipsec-sa name;
    retransmit-interval seconds;
    transit-delay seconds
    keychain keychain-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf
area area-id],
[edit protocols (ospf | ospf3) area area-id],
[edit routing-instances routing-instance-name protocols ospf area area-id]
```

Description

For backbone areas only, create a virtual link to use in place of an actual physical link. All area border routers and other routing devices on the backbone must be contiguous. If this is not possible and there is a break in OSPF connectivity, use virtual links to create connectivity to the OSPF backbone. When configuring virtual links, you must configure links on the two routing devices that form the end points of the link, and both of these routing devices must be area border routers. You cannot configure links through stub areas.

Options

neighbor-id ***router-id***	IP address of the routing device at the remote end of the virtual link.
transit-area ***area-id***	Area identifier of the area through which the virtual link transits. Virtual links are not allowed to transit the backbone area.
ipsec-sa ***name***	Apply the named IPsec authentication to the OSPF interface or virtual link or to an OSPFv2 remote sham link.
keychain ***keychain-name—***	For each keychain, you must specify a name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

The option `keychain` *`keychain-name`* introduced in Junos OS and Junos OS Evolved Release 24.2R1.

RELATED DOCUMENTATION

> *Understanding OSPF Areas*
>
> *Example: Configuring OSPF Virtual Links to Connect Noncontiguous Areas*

virtual-link-local-address

IN THIS SECTION

- Syntax | 30765
- Hierarchy Level | 30765
- Description | 30765
- Options | 30766
- Required Privilege Level | 30766
- Release Information | 30766

Syntax

```
virtual-link-local-address ipv6-address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
```

Description

Used to explicitly configure a virtual link-local address for a Virtual Router Redundancy Protocol (VRRP) IPv6 group. The virtual link-local address must be configured to be on a fe80::/64 subnet for proper operation. Early versions of the Junos OS required manual configuration of this parameter. In current Junos releases a virtual link local address is automatically created based on the interface's MAC address.

NOTE: In current Junos releases you do *not* need to configure link-local addresses and virtual link-local addresses when configuring VRRP for IPv6. Junos OS automatically generates link-local addresses and virtual link-local addresses. However, if link local

addresses and virtual link-local addresses are configured, Junos OS honors the explicitly configured addresses.

Options

ipv6-address—virtual link-local IPv6 address for VRRP for an IPv6 group.

- **Range:** 0 through 255

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring Basic VRRP Support

Junos OS Support for VRRPv3

virtual-mac

IN THIS SECTION

- Syntax | 30767
- Hierarchy Level | 30767
- Description | 30767
- Options | 30767
- Required Privilege Level | 30767
- Release Information | 30767

Syntax

```
virtual-mac mac-address/mask;
```

Hierarchy Level

```
[edit protocols l2-learning global-mac-move]
```

Description

Configure MAC addresses that are not to be considered in the VPLS loop prevention algorithm.

> **NOTE**: Starting in Junos OS Release 17.4R1, this statement has been replaced with the exclusive-mac statement.

Options

mac-address/mask MAC address

- **Values:** 00:00:5e:00:01:00/8, 00:00:0c:07:ac:00/8, and 02:bf:00:00:00:00/32

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Example: Configuring Loop Prevention in VPLS Network Due to MAC Moves

global-mac-move | **7716**

virtual-network-functions

IN THIS SECTION

- Syntax | 30768
- Hierarchy Level | 30768
- Description | 30771
- Options | 30771
- Required Privilege Level | 30772
- Release Information | 30772

Syntax

```
virtual-network-functions vnf-name;
```

Hierarchy Level

```
virtual-network-functions {
    ipsec-nm {
        interfaces {
            heth1 | heth2 {
                mapping {
                    vlan {
                        members vlan-name;
                    }
                }
            }
        }
    }
    vjunos0 {
        interfaces {
            em1 {
                mapping {
                    vlan {
                        members vlan-name;
```

```
                    }
                }
            }
        }
    }
    vmx {
        image filepath;
            virtual-cpu {
            count virtual-cpu-count;
            features;
        }
        memory size size;
        storage {
        vdc {
            type type;
            source file-path;
        }
        vdb {
            type type;
            source file-path;
        }
    }
    vsrx {
    vnf-name  {
        type {
            linux-container | virtual-machine;
        }
        image {
            file-path;
            bus-type [ide | virtio];
            image-type [qcow2 | raw];
        }
        init-descriptor file-path;
        memory {
            size size;
            features {
                hugepages ;
            }
        }
        no-autostart;
        storage device-name {
            type {
                cdrom {
```

```
                source {
                    file filename;
                }
            }
            disk {
                bus-type [ide | virtio];
                    file-type [qcow2 | raw];
                    source {
                    file filename;
                }
            }
            usb {
                source {
                    file filename;
                }
            }
        }
    }
    virtual-cpu {
        virtual-cpu-number {
            physical-cpu number | range;
        }
        count number;
        features {
            hardware-virtualization;
        }
    }
    interfaces interface-name {
        pci-address pci-address;
        mapping {
            hsxe0 {
                virtual-function {
                    vlan-id vlan-id;
                }
            }
            hsxe1 {
                virtual-function {
                    vlan-id vlan-id;
                }
            }
            vlan {
                members vlan-name;
                mode [access | trunk];
```

```
                    native-vlan-id vlan-id;
                }
                peer-interfaces peer-interface-name;
            }
            management {
                internal;
                out-of-band;
            }
        }
    }
```

Description

Create an instance of a virtual network function (VNF) on platforms that run disaggregated Junos OS software.

NOTE:

- Creating a VNF instance fails if the resources required by the VNF are not available in the system.
- If you use init-descriptor to define a VNF by specifying and setting different values for the virtual CPU count or memory and later if you delete the virtual cpu count, the system restores the value to a default value of 1 for vCPU and 1GB for memory.
- You can enable the VNF options such virtual-cpu-hardware-virtualization (vmx), hugepages, image-type, and image-bus-type only when you define the VNF initially. You cannot enable or disable the VNF options after committing the VNF configuration. To enable or disable the VNF options, you must delete the VNF configuration and re-configure with the VNF options.

Options

virtual-network-functions *vnf-name*	Name of the VNF instance. It is mandatory to provide one of the options: init-descriptor or image.
file-path	Path of the source file.
number | range	Number or a range of physical CPUs. For example, 2 or 2...5.

interface-name	Name of the VNF interface, which can range from eth0 to eth9. You can configure eth0 and eth1 interfaces and can assign VLAN IDs. To configure eth0 and eth1 interfaces, you must configure `no-default-interfaces` option.
physical-interface-name	Name of the physical interface to which the VNF interface is attached.
vlan-id	SR-IOV virtual function to use to attach a VNF to a physical interface.
native-vlan-id	Virtual LAN identifier for untagged frames. For example, 1...4095
vlan-name	Name of the VLAN members.
size	Amount of memory allocated to a VNF in kilobytes. The default size is 1 GB.
device-name	Name of the storage device.
file-name	Name of the source file of the storage device.
peer-interface-name	Name of the virtual peer interfaces that is mapped to a physical interface.
management	VNF interface management configuration.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D45.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | **27723**

virtual-network-functions (Junos node slicing)

IN THIS SECTION

- Syntax | 30773
- Hierarchy Level | 30773
- Description | 30773
- Options | 30773
- Required Privilege Level | 30774
- Release Information | 30774

Syntax

```
virtual-network-functions vnf-name {
    base-config base-config ;
    chassis-type chassis-type;
    id id;
    resource-template resource-template;
}
```

Hierarchy Level

[edit]

Description

Associate a GNF ID, base configuration, chassis type and resource template with the VNF.

The GNFs that are configured and committed will appear as auto-complete options in operational commands.

Options

vnf-name	Name of the VNF.

base-config *base-config*	Choose a base configuration for the VNF.
chassis-type *chassis-type*	Applicable only to Junos node slicing set up using external server model. Choose the type of the router chassis (for example, MX960) used as the base system (BSYS) in the node slicing setup.
id *id*	Assign a GNF ID. • **Range:** 1–10 • **Range:** 1-4 (for in-chassis model)
resource-template *resource-template*	Assign a resource template to the VNF. The server resource template defines the number of dedicated CPU cores and the size of DRAM to be assigned to a VNF.

Required Privilege Level

root

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

server (Junos node slicing) | **16606**

virtual-service (Traffic Load Balancer)

IN THIS SECTION

- Syntax | **30775**
- Hierarchy Level | **30775**
- Description | **30775**
- Options | **30775**
- Required Privilege Level | **30777**

Release Information | 30777

Syntax

```
virtual-service virtual-service-name {
    address virtual-ip-address;
    group group-name;
    load-balance-method {
        hash {
            hash-key method;
        }
        random;
    }
    mode ( layer2-direct-server-return | direct-server-return | translated );
    <routing-instance routing-instance-name>;
    <routing-metric route-metric>;
    server-interface server-interface;
    service service-name {
        protocol (udp | tcp);
        server-listening-port port;
        virtual-port virtual-port;
    }
}
```

Hierarchy Level

```
[edit services traffic-load-balance instance instance-name]
```

Description

Configure a TLB virtual service.

Options

address *virtual-ip-address*	Address of the virtual service.

group *group-name*

Server group for the virtual service.

load-balance method hash hash-key *method*

Use a combination of these hash-key methods for the session distribution API:

dest-ip Hash on destination IP address.

proto Hash on protocol.

source-ip Hash on source IP address.

load-balance-method random

Use randomizing algorithm for session distribution.

mode (layer2-direct-server-return | direct-server-return | translated)

Traffic load balancer mode of operation:

direct-server-return Transparent mode Layer 3 direct server return.

layer2-direct-server-return Transparent mode Layer 2 direct server return. Load balancing works by changing the Layer 2 MAC of packets; Layer 3 and higher level headers are not modified.

translated The Packet Forwarding Engine performs stateless load balancing.

route-metric

(Optional) Route metric

- **Range:** 1 through 255

routing-instance-name

(Optional) Routing instance for the virtual service. Default is `inet.0`.

server-interface *server-interface*

(Optional) The server-interface specified under the virtual-service, will be used instead of the values provided under the instance level.

service *service-name*

Translated mode details. Packets destined to this virtual ip-address + virtual-port + protocol will be load balanced to the appropriate server. The destination IP address and port are replaced by the real services IP address and the server-listening-port (configured here).

protocol (udp | tcp) Protocol.

server-listening-port *port* Port number.

virtual-port *virtual-port* Virtual port number.

virtual-ip–address

Local address for the virtual service.

virtual-service-name Identifier for the virtual service.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

Traffic Load Balancer Overview

Configuring TLB

virtual-switch

IN THIS SECTION

- Syntax | 30777
- Hierarchy Level | 30778
- Description | 30778
- Required Privilege Level | 30778
- Release Information | 30778

Syntax

```
virtual-switch name bridge-domain name vlan-id [vlan-ids ];
```

Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name default-x]
```

Description

Specify the routing-instance type as a virtual switch, under which bridge-domain MIPs must be enabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring Maintenance Intermediate Points (MIPs)

virus-detection (Security Antivirus)

IN THIS SECTION

- Syntax | 30779
- Hierarchy Level | 30779
- Description | 30779
- Options | 30779
- Required Privilege Level | 30779
- Release Information | 30779

Syntax

```
virus-detection {
    custom-message message;
    custom-message-subject message-subject;
    (notify-mail-sender | no-notify-mail-sender);
     type (message | protocol-only);
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
notification-options]
[edit security utm feature-profile anti-virus avira-engine profile profile-name notification-
options]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name notification-
options]
```

Description

Configure a notification to send when a virus is detected.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

The Express and Kaspersky Antivirus features are not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1X49-D10	The Express and Kaspersky Antivirus features are not supported from Junos OS Release 15.1X49-D10 onwards.

vjunos0

IN THIS SECTION

- Syntax | 30780
- Hierarchy Level | 30781
- Description | 30781
- Options | 30781
- Required Privilege Level | 30781
- Release Information | 30781

Syntax

```
vjunos0 {
    interfaces {
        em1 {
            mapping {
                vlan {
                    members vlan-name;
                }
            }
        }
```

```
    }
}
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

An option that enables or disables the vjunos virtual network function if the vjunos0 option is configured in the system.

Options

interfaces	Name of the interface. For example, em1.
vlan members *vlan-name*	Name of the VLAN members.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D45.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | **27723**

vlan

IN THIS SECTION

- Syntax (Ethernet and 802.1Q Tagging) | 30782
- Syntax (Static MAC-based VLANs) | 30782
- Syntax (Unknown Unicast) | 30783
- Ethernet | 30783
- 802.1Q Tagging | 30783
- Static MAC-based VLANs | 30783
- Unknown Unicast | 30783
- Description | 30783
- Options | 30784
- Required Privilege Level | 30784
- Release Information | 30784

Syntax (Ethernet and 802.1Q Tagging)

```
vlan {
    members [ (all | names | vlan-ids) ];
}
```

Syntax (Static MAC-based VLANs)

```
vlan vlan-name {
    mac  mac-address {
            next-hop interface-name;
    }
}
```

Syntax (Unknown Unicast)

```
vlan  (all | vlan-name) {
    interface interface-name;
}
```

Ethernet

```
[edit interfaces ge-chassis/slot/port                unit                logical-unit-
number                ethernet-switching],
[edit interfaces xe-chassis/slot/port                unit                logical-unit-
number                ethernet-switching]
```

802.1Q Tagging

```
[edit interfaces interface-name                unit                logical-unit-
number                family ethernet-switching]
```

Static MAC-based VLANs

```
[edit ethernet-switching-options static]
```

Unknown Unicast

```
[edit ethernet-switching-options unknown-unicast-forwarding]
```

Description

For both Gigabit Ethernet and aggregated Ethernet interfaces and 802.1Q Tagging, assign an 802.1Q VLAN tag ID to a logical interface.

For Static MAC-based VLANs, specify the name of a VLAN to add to the Ethernet switching table.

For unknown unicast, specify a VLAN from which unknown unicast packets will be forwarded, or specify that the packets should be forwarded from *all* VLANS. Unknown unicast packets are forwarded from a VLAN to a specific trunk interface.

The remaining statements are explained separately. See CLI Explorer.

> **TIP**: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type **?** after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options

all All VLANs.

vlan-name Name of the VLAN to add to the Ethernet switching table.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Bridging and VLANs

MAC Addresses

show ethernet-switching interfaces | **19641**

show ethernet-switching table | **19700**

show ethernet-switching interface | **19636**

Bridging and VLANs

Configuring Routed VLAN Interfaces on Switches (CLI Procedure)

Integrated Routing and Bridging

Bridging and VLANs

show vlans | **27738**

Junos OS Network Interfaces Configuration Guide

vlan (Access Port Security)

IN THIS SECTION

- Syntax | **30785**
- Hierarchy Level | **30786**
- Description | **30786**
- Options | **30787**
- Required Privilege Level | **30787**
- Release Information | **30787**

Syntax

```
vlan (all | vlan-name) {
    (arp-inspection | no-arp-inspection) {
        forwarding-class class-name;
    }
    dhcp-option82 {
        circuit-id {
            prefix hostname;
            use-interface-description;
            use-vlan-id;
        }
         remote-id {
            prefix hostname | mac | none;
            use-interface-description;
            use-string                                  string;
```

```
        }
        vendor-id <string>;
    }
    (examine-dhcp | no-examine-dhcp) {
        forwarding-class class-name;
    }
    (examine-dhcpv6 | no-examine-dhcpv6) {
        forwarding-class class-name;
    }
    examine-fip {
        fc-map fc-map-value;
    }
    (ip-source-guard | no-ip-source-guard);
    (ipv6-source-guard | no-ipv6-source-guard);
    mac-move-limit limit action (drop | log | none | shutdown);
    }
    (neighbor-discovery-inspection | no-neighbor-discovery-inspection);
    no-option37;
}
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port]
```

Description

Apply any of the following security options to a VLAN:

- DHCP snooping
- DHCPv6 snooping with DHCP option 37
- DHCP option 82
- Dynamic ARP inspection (DAI)
- IPv6 neighbor discovery inspection
- FIP snooping
- IP source guard
- IPv6 source guard

- MAC move limiting

The remaining statements are explained separately. See CLI Explorer.

> **TIP**: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type **?** after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options

all—Apply the feature to all VLANs.

vlan-name—Apply the feature to the specified VLAN.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for the `examine-dhcpv6`, `no-option37`, `neighbor-discovery-inspection`, and `ipv6-source-guard` statements introduced in Junos OS Release 14.1x53-D10 for EX Series switches.

RELATED DOCUMENTATION

- Example: Configuring Port Security (non-ELS)
- Example: Configuring IP Source Guard with Other EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces
- Example: Setting Up DHCP Option 82 Using the Same VLAN
- *Example: Configuring an FCoE Transit Switch*

vlan (Bridge Domains)

IN THIS SECTION

- Syntax | 30788
- Hierarchy Level | 30789
- Description | 30789
- Default | 30789
- Options | 30789
- Required Privilege Level | 30789
- Release Information | 30789

Syntax

```
vlan vlan-id {
    all
    immediate-leave;
    interface interface-name {
        group-limit limit;
        host-only-interface;
        multicast-router-interface;
        static {
            group multicast-group-address {
                source ip-address;
            }
        }
    }
    proxy {
        source-address ip-address;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
}
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name protocols igmp-snooping],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping]
```

Description

Configure IGMP snooping parameters for a particular VLAN.

Default

By default, IGMP snooping options apply to all VLANs.

Options

vlan-id—Apply the parameters to this VLAN.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring VLAN-Specific IGMP Snooping Parameters

igmp-snooping | 8387

vlan (DHCP Bindings on Access Ports)

IN THIS SECTION

- Syntax | 30790
- Hierarchy Level | 30790
- Description | 30790
- Options | 30790
- Required Privilege Level | 30790
- Release Information | 30791

Syntax

```
vlan  vlan-name;
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port interface (all | interface-name) static-ip ip-address]
```

Description

Associate the static IP address with the specified VLAN associated with the specified interface.

Options

vlan-name —Name of a specific VLAN associated with the specified interface.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Configuring Static DHCP IP Addresses for DHCP snooping (non-ELS)

vlan-group

IN THIS SECTION

- Syntax | 30791
- Hierarchy Level | 30791
- Description | 30791
- Required Privilege Level | 30792
- Release Information | 30792

Syntax

```
vlan-group group group-name {
vlan (vlan-id |vlan-group |all) {
}
```

Hierarchy Level

```
[edit protocols vstp]
```

Description

Configure VLAN group for Spanning Tree Protocol (VSTP). VSTP is used to prevent loops in Layer 2 networks on a per-VLAN basis.

BEST PRACTICE: Configure RSTP when you configure VSTP. RSTP overhead is minimal and this configuration ensures that a spanning-tree protocol is running on all VLANs on your switch, even when your switch is supporting more than the maximum number of allowed VSTP VLANs.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

vstp | 30983
show spanning-tree bridge | 26548
show spanning-tree interface | 26555

vlan-groups

IN THIS SECTION

- Syntax | 30793
- Hierarchy Level | 30793
- Description | 30793
- Options | 30793
- Required Privilege Level | 30793
- Release Information | 30793

Syntax

```
vlan-groups vlan_group_name{
      vlan-id-list vlan_id_ist
}
```

Hierarchy Level

```
[edit vlans]
```

Description

Configure VLAN group on EX series switches. The 802.1X VLAN group feature uses a single wireless local area network (WLAN) that can support multiple VLANs. The client gets assigned to one of the configured VLANs.

The VLAN group name is added within the Tunnel-Private-Group-ID (defined as RADIUS attribute type 81, RFC 2868) and sent in the RADIUS response instead of a regular VLAN ID or VLAN Name when a supplicant is requesting 802.1X authentication through the RADIUS server. When the switch receives the VLAN group name, the switch assigns the endpoint to one of the VLANs in that group using the round-robin algorithm.

Options

vlan-id-list Specify the associated VLAN id.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

show vlans vlan-groups | 27774

802.1X Authentication

vlan-id (Fast Ethernet)

IN THIS SECTION

- Syntax | 30794
- Hierarchy Level | 30794
- Description | 30794
- Options | 30795
- Required Privilege Level | 30795
- Release Information | 30795

Syntax

```
vlan-id number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For Fast Ethernet and Gigabit Ethernet interfaces only, bind an 802.1Q VLAN tag ID to a logical interface.

Options

number—A valid VLAN identifier.

- **Range:** For 4-port Fast Ethernet PICs configured to handle VPLS traffic, 512 through 1023. For 1-port and 10-port Gigabit Ethernet PICs configured to handle VPLS traffic, 512 through 4094.

> **NOTE**: On Junos OS Evolved, `vlan-id` 0 is not supported.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Interfaces for VPLS Routing

vlan-id

IN THIS SECTION

- Syntax | 30796
- VLANs and Bridge Domain VLANs | 30796
- 802.1Q Tagging | 30797
- VLAN ID to Rewrite | 30797
- VLAN Tagging and Layer 3 Subinterfaces | 30797
- Description | 30797
- Default | 30798

- Options | 30799
- Required Privilege Level | 30800
- Release Information | 30800

Syntax

```
vlan-id (all | none | number);
```

VLANs and Bridge Domain VLANs

For platforms without ELS:

```
[edit vlans vlan-name vlan-range]
```

For platforms without ELS and with ELS:

```
[edit vlans vlan-name]
```

For ELS platforms only:

```
[edit interfaces interface-name unit number]
[edit vlans vlan-name vlan-id-list]
```

```
[edit vlans vlan-name],
[edit logical-systems logical-system-name vlansvlan-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name vlansvlan-
name],
[edit routing-instances routing-instance-name vlans vlan--name]
```

CAUTION: When you want to change a traffic impacting option under a routing instance, use the following procedure.

1. Deactivate the routing instance configuration.
2. Change the traffic impacting option.
3. Reactivate the updated routing instance configuration.

For example, follow this procedure if you need to change settings such as:

- The `vlan-id` in an EVPN routing instance-- Changing the `vlan-id` on the fly, without first deactivating the associated EVPN routing instance would be catastrophic.
- The `service-type` in a MAC-VRF routing instance-- When you change the service type of a running instance, the device might incorrectly change the VLAN ID if it is not deactivated prior to making the change.

802.1Q Tagging

```
[edit vlans vlan-name]
```

VLAN ID to Rewrite

```
[edit interfaces interface-name unit logical-unit-number input-vlan-map],
[edit interfaces interface-name unit logical-unit-number output-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
input-vlan-map],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
output-vlan-map]
```

VLAN Tagging and Layer 3 Subinterfaces

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

For VLANs, specify a VLAN identifier (VID) to include in the packets sent to and from the VLAN, or a VPLS routing instance.

> **NOTE**: When configuring a VLAN identifier for provider backbone bridge (PBB) routing instances, dual-tagged VIDs and the `none` option are not permitted.

For 10-Gigabit Ethernet and aggregated Ethernet interfaces only, bind an 802.1Q VLAN tag ID to a logical interface. Statement introduced in Junos OS Release 9.2 for EX Series switches.

> **NOTE**: The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

For 802.1Q tagging, configure an 802.1Q tag to apply to all traffic that originates on the VLAN.

The number zero is reserved for priority tagging and the number 4095 is also reserved.

For VLAN ID to Rewrite Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface.

You cannot include the `vlan-id` statement with the `swap` statement, `swap-push` statement, `push-push` statement, or `push-swap` statement at the [`edit interfaces` *`interface-name`* `unit` *`logical-unit-number`* `output-vlan-map`] hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the `vlan-id` statement that you include at the [`edit interfaces` *`interface-name`* `unit` *`logical-unit-number`*] hierarchy level.

Default

For 802.1Q Tagging on EX Series and SRX Series, If you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of **1**.

On a EX2300 switch, the maximum number of vlans supported is 2024. The vlans can be in the range of 1-4093.

On a EX3400 switch, the maximum number of vlans that can be created is 4093 including the default vlan with id 1. In a single `vlan-id-list` default `vlan1` is always inherited, so the valid configurable vlan range is 2-4093. You can use `vlan-id` up to and including 4094, but 4093 is the maximum number of vlans that can be configured.

For VLANs on a QFX3500 and QFX3500 switch, if you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of **1**. The number zero is reserved for priority tagging and the number 4093 is also reserved.

On a QFX5100 switch, if you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of **1**. The number zero is reserved for priority tagging and the number 4093 is also reserved.

> **NOTE**: You can only create up 4090 VLANs on a QFX5100 switch. If you create more than 4090 VLANs, the interfaces associated with the extra VLANs are not displayed in the `show vlans` command output. For example, if you create 4094 VLANs, the extra VLANs will not have interfaces associated with the VLANs. The order in which you configure the extra VLANs determines which interfaces are missing from the `show vlans` command output.

For VLAN tagging and Layer 3 subinterfaces, bind an 802.1Q VLAN tag ID to a logical interface.

> **NOTE**: The VLAN tag ID cannot be configured on logical interface unit **0**. The logical unit number must be **1** or higher.

Options

For VLANs:

`number`—A valid VLAN identifier. If you configure multiple VLANs with a valid VLAN identifier, you must specify a unique VLAN identifier for each. However, you can use the same VLAN identifier for VLANs that belong to different virtual switches. Use this option to send single tagged frames with the specified VLAN identifier over VPLS VT interfaces.

> **NOTE**: If you specify a VLAN identifier, you cannot also use the `all` option. They are mutually exclusive.

`all`—Specify that the VLAN spans all the VLAN identifiers configured on the member logical interfaces.

> **NOTE**: You cannot specify the `all` option if you include a routing interface in the VLAN.

`none`—Specify to enable shared VLAN learning or to send untagged frames over VPLS VT interfaces.

> **NOTE**: Multichassis link aggregation (MC-LAG) does not support the `none` option with the `vlan-id` statement with VLANs.

For 802.1Q Tagging:

number—VLAN tag identifier

Range:

-
  - 1 through 4094 (all switches except EX8200 Virtual Chassis)
  - 1 through 4092 (EX8200 Virtual Chassis only)
- **Default:** 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for Layer 2 trunk ports added in Junos OS Release 9.2.

Support for SRX 5600, and SRX 5800 devices added in Junos OS Release 9.6.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Example: Connecting Access Switches with ELS Support to a Distribution Switch with ELS Support

Example: Configuring a Private VLAN on a Single Switch with ELS Support

Creating a Private VLAN on a Single Switch with ELS Support (CLI Procedure)

Creating a Private VLAN Spanning Multiple EX Series Switches with ELS Support (CLI Procedure)

Example: Configuring VLANs on Security Devices

Example: Configuring Interfaces and Routing Instances for a User Logical Systems

Rewriting the VLAN Tag on Tagged Frames

Binding VLAN IDs to Logical Interfaces

vlan-tagging | **30861**

Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support

Configuring Gigabit Ethernet Interfaces (J-Web Procedure)

Configuring a Layer 3 Subinterface (CLI Procedure)

Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support

Junos OS Ethernet Interfaces Configuration Guide

vlan-id (Dynamic Profiles)

IN THIS SECTION

- Syntax | 30801
- Hierarchy Level | 30801
- Description | 30801
- Options | 30802
- Required Privilege Level | 30802
- Release Information | 30802

Syntax

```
vlan-id (number | none);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

Description

For VLAN demux, Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.

Options

number—A valid VLAN identifier. When used in the `dynamic-profiles` hierarchy, specify the `$junos-vlan-id` predefined variable to dynamically obtain the VLAN identifier.

`none`—Enable the use of untagged pseudo-wire frames on dynamic interfaces.

- For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.
- For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.
- VLAN ID 0 is reserved for tagging the priority of frames.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

VLAN demux interface support introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles

vlan-id (Dynamic VLANs)

IN THIS SECTION

- Syntax | 30803
- Hierarchy Level | 30803
- Description | 30803
- Options | 30803

- Required Privilege Level | 30803
- Release Information | 30804

Syntax

```
vlan-id number;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number input-vlan-map],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map]
```

Description

For dynamic VLAN interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface.

You cannot include the `vlan-id` statement with the `swap` statement, `swap-push` statement, `push-push` statement, or `push-swap` statement at the [`edit dynamic-profiles` *`profile-name`* `interfaces` *`interface-name`* `unit` *`logical-unit-number`* `output-vlan-map`] hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the `vlan-id` statement that you include at the [`edit dynamic-profiles` *`profile-name`* `interfaces` *`interface-name`* `unit` *`logical-unit-number`*] hierarchy level.

Options

number A valid VLAN identifier. When used for input VLAN maps, you can specify the `$junos-vlan-map-id` predefined variable to dynamically obtain the VLAN identifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Rewriting the VLAN Tag on Tagged Frames

Binding VLAN IDs to Logical Interfaces

vlan-id inner-all

IN THIS SECTION

- Syntax | 30804
- Hierarchy Level | 30804
- Description | 30804
- Required Privilege Level | 30805
- Release Information | 30805

Syntax

```
vlan-id inner-all;
```

Hierarchy Level

```
[edit routing-instances instance_name]
```

Description

Enable qualified MAC learning on the second (inner) VLAN tag of an ingress 2-tagged packet, without removing the first (outer) tag implicitly. For a single vlan-tagged packet, qualified MAC learning happens

on VLAN 4096. If the ingress packet has more than two tags, all tags beyond the second tag are treated as part of data.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

Understanding Qualified MAC Learning

deep-vlan-qualified-learning | **5176**

vlan-id-list

IN THIS SECTION

- Syntax | **30805**
- Hierarchy Level | **30806**
- Description | **30806**
- Options | **30807**
- Required Privilege Level | **30807**
- Release Information | **30807**

Syntax

```
vlan-id-list [ vlan-id-numbers ];
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name],
[edit interfaces interface-name unit 0],
[edit interfaces interface-name unit logical-unit-number],
[edit vlans vlan-name]
```

Description

Specify a VLAN identifier list to use for a bridge domain or VLAN in trunk mode. VLAN identifier list can be used on C-VLAN interfaces in Q–in–Q tunneling for EX and QFX Series switches.

Specify the `trunk` option in the `interface-mode` statement to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the `vlan-id-list` statement to forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the `access` option to accept packets with no VLAN ID to forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the `vlan-id` statement.

This statement also enables you to bind a logical interface to a list of VLAN IDs, thereby configuring the logical interface to receive and forward a frame with a tag that matches the specified VLAN ID list.

WARNING: On EX2000, EX3000, EX4000, and QFX5000 Series switches, if VLAN identifier list (`vlan-id-list`) is used for Q-in-Q tunnelling, you can apply no more than eight VLAN ID ranges to a physical interface, regardless of the number of logical interfaces or the number of VLAN ID lists configured. When more than eight VLAN ID ranges are configured and committed, only the first eight ranges are considered. Traffic will *not* be forwarded to the vlans specified in VLAN ID ranges nine and above.

In the example below, traffic will not be forwared to the vlans specified in the range `170-180`.

```
user@host# show interfaces xe-0/0/44
flexible-vlan-tagging;
encapsulation flexible-ethernet-services;
unit 1 {
    encapsulation vlan-bridge;
```

```
            vlan-id-list [ 10-20 30-40 50-60 70-80];
            input-vlan-map push;
            output-vlan-map pop;
        }
        unit 2 {
            encapsulation vlan-bridge;
            vlan-id-list [ 90-100 110-120 130-140 150-160];
            input-vlan-map push;
            output-vlan-map pop;
        }
        unit 3 {
            encapsulation vlan-bridge;
            vlan-id-list 170-180;
            input-vlan-map push;
            output-vlan-map pop;
        }
```

Options

vlan-id-numbers—Valid VLAN identifiers. You can combine individual numbers with range lists by including a hyphen.

- **Range:** 0 through 4095

> **NOTE**: On EX Series switches and the QFX Series, the range is 0 through 4094.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring a Bridge Domain

Configuring a VLAN

Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances

Configuring VLAN Identifiers for VLANs and VPLS Routing Instances

Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation

vlan-id-list (Ethernet VLAN Circuit)

IN THIS SECTION

- Syntax | 30808
- Hierarchy Level | 30808
- Description | 30808
- Options | 30809
- Required Privilege Level | 30810
- Release Information | 30810

Syntax

```
vlan-id-list [vlan-id vlan-id-vlan-id];
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Binds a single-tag logical interface to a list of VLAN IDs. Configures a logical interface to receive and forward any tag frame whose VLAN ID tag matches the list of VLAN IDs you specify.

NOTE: When you create a circuit cross-connect (CCC) using VLAN-bundled single-tag logical interfaces on Layer 2 VPN routing instances, the circuit automatically uses `ethernet` encapsulation. For Layer 2 VPN, you need to include the `encapsulation-type` statement and specify the value `ethernet` at either of the following hierarchy levels:

- `[edit routing-instances routing-instance-name protocols l2vpn]`
- `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn]`

For more information about the `encapsulation-type` configuration statement and the Layer 2 encapsulation types `ethernet` and `ethernet-vlan`, see the Junos OS VPNs Library for Routing Devices.

Options

[*vlan-id vlan-id-vlan-id*]—A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

- **Range:** 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

NOTE: Configuring `vlan-id-list` with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

1. **Inefficient**

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id-range 1-4094;
}
```

2. **Best Practice**

```
[edit interfaces interface-name]
unit 0;
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Binding VLAN IDs to Logical Interfaces

encapsulation (Logical Interface) | **6300**

encapsulation (interfaces) | **6289**

Junos OS VPNs Library for Routing Devices

flexible-vlan-tagging | **7235**

vlan-tagging | **30861**

vlan-tags (Dual-Tagged Logical Interface) | **30869**

vlan-id-list (Interface in Bridge Domain)

IN THIS SECTION

- Syntax | **30811**
- Hierarchy Level | **30811**
- Description | **30811**
- Options | **30811**
- Required Privilege Level | **30811**

Release Information | 30811

Syntax

```
vlan-id-list [ number number-number ];
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family bridge],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family bridge]
```

Description

Configure a logical interface to forward packets and learn MAC addresses within each bridge domain configured with a VLAN ID that matches a VLAN ID specified in the list. VLAN IDs can be entered individually using a space to separate each ID, entered as an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.

Options

number number—Individual VLAN IDs separated by a space.

number-number—Starting VLAN ID and ending VLAN ID in an inclusive range.

- **Range:** 1 through 4095

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

- Configuring a Logical Interface for Trunk Mode
- Configuring the VLAN ID List for a Trunk Interface
- Tunnel Services Overview
- Tunnel Interface Configuration on MX Series Routers Overview

vlan-id-list (Interface in VPLS)

IN THIS SECTION

- Syntax | 30812
- Hierarchy Level | 30812
- Description | 30813
- Options | 30813
- Required Privilege Level | 30813
- Release Information | 30813

Syntax

```
vlan-id-list [ numbers number-number ];
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure a logical interface to forward packets and learn MAC addresses within each VPLS routing instance configured with a VLAN ID that matches a VLAN ID specified in the list. VLAN IDs can be entered individually using a space to separate each ID, entered as an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.

Options

number number—Individual VLAN IDs separated by a space.

number-number—Starting VLAN ID and ending VLAN ID in an inclusive range.

- **Range:** 1 through 4095

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced for VPLS in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Interfaces for VPLS Routing

Configuring Interfaces for VPLS Routing

vlan-id (Logical Port in Bridge Domain)

IN THIS SECTION

- Syntax | 30814
- Hierarchy Level | 30814
- Description | 30814

- Options | 30814
- Required Privilege Level | 30814
- Release Information | 30814

Syntax

```
vlan-id number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family bridge],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family bridge]
```

Description

The VLAN ID configured on the logical port. Received packets with no VLAN tags are forwarded within the bridge domain with the matching VLAN ID.

Options

`number`—The VLAN ID.

- **Range:** 1 through 4095

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring Access Mode on a Logical Interface

Tunnel Services Overview

Tunnel Interface Configuration on MX Series Routers Overview

vlan-id-range

IN THIS SECTION

- Syntax | 30815
- Hierarchy Level | 30815
- Description | 30815
- Options | 30816
- Required Privilege Level | 30816
- Release Information | 30817

Syntax

```
vlan-id-range vlan-id-vlan-id
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Bind a range of VLAN IDs to a logical interface.

Options

`number`—The first number is the lowest VLAN ID in the range the second number is the highest VLAN ID in the range.

- **Range:** 1 through 4094

> **NOTE**: On SRX240, SRX300, SRX320, SRX340, SRX345, SRX550M, and SRX650 devices, the VLAN range from 1 to 4094 on inet interfaces and the VLAN range from 1 to 3967 on Ethernet switching interfaces. On Ethernet switching interfaces, the VLAN range from 3968 to 4094 falls under the reserved VLAN address range, and the user is not allowed to configure VLANs in this range.

> **NOTE**: Configuring `vlan-id-range` with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.
>
> The following examples illustrate this further:
>
> 1. **Inefficient**
>
> ```
> [edit interfaces interface-name]
> vlan-tagging;
> unit number {
>  vlan-id-range 1-4094;
> }
> ```
>
> 2. **Best Practice**
>
> ```
> [edit interfaces interface-name]
> unit 0;
> ```

VLAN ID 0 is reserved for tagging the priority of frames.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

RELATED DOCUMENTATION

Understanding VLANs

vlan-id (routing instance)

IN THIS SECTION

- Syntax | 30817
- Hierarchy Level | 30817
- Description | 30818
- Options | 30818
- Required Privilege Level | 30818
- Release Information | 30818

Syntax

```
vlan-id (vlan-id | all | none);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
[edit routing-instances routing-instance-name instance-type]
```

Description

Specify 802.1Q VLAN tag IDs to a routing instance.

Options

vlan-id—A valid VLAN identifier.

- **Range:** For 4-port Fast Ethernet PICs, 512 through 1023. For 1-port and 10-port Gigabit Ethernet PICs configured to handle VPLS traffic, 512 through 4094.

`all`—Include all VLAN identifiers specified on the logical interfaces included in the routing instance.

`none`—Include no VLAN identifiers for the routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

vlan-id-start

IN THIS SECTION

- Syntax | 30819
- Hierarchy Level | 30819
- Description | 30819

- Options | 30819
- Required Privilege Level | 30819
- Release Information | 30819

Syntax

```
vlan-id-start S-VLAN-ID;
```

Hierarchy Level

```
[edit vlans vlan-name interfaces interface-name mapping-range C-VLAN-range (push | swap)]
```

Description

Configure an access interface to map a range of C-VLANs to a range of S-VLANs. Use this statement instead of using multiple `set vlans VLAN-name interface interface-name mapping (push | swap)` statements to configure Q-in-Q tunneling or VLAN translation on a per-VLAN basis. This statement sets the start of the S-VLAN range that the C-VLANs are mapped to. If you omit this option, mapping begins with the first ID in the range of S-VLAN IDs (which you configure using the `set vlans vlan-range`statement).

Options

None

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

vlan-id (VLAN ID to Be Bound to a Logical Interface)

IN THIS SECTION

- Syntax | 30820
- Hierarchy Level | 30820
- Description | 30820
- Options | 30820
- Required Privilege Level | 30821
- Release Information | 30821

Syntax

```
vlan-id number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

For Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.

Options

number—A valid VLAN identifier.

- **Range:** For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023. In Junos OS Evolved `vlan-id 0` is not supported.

For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.

VLAN ID 0 is reserved for tagging the priority of frames used in Junos OS. In Junos OS Evolved `vlan-id 0` is not supported.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Enabling VLAN Tagging

vlan (IGMP Snooping)

IN THIS SECTION

- Syntax (EX Series and SRX-series:SRX 210) | **30822**
- Syntax (EX4600, NFX Series, QFabric Systems, and QFX Series) | **30822**
- Hierarchy Level | **30823**
- Description | **30823**
- Default | **30824**
- Options | **30824**
- Required Privilege Level | **30825**
- Release Information | **30825**

Syntax (EX Series and SRX-series:SRX 210)

```
vlan (all | vlan-name) {
    data-forwarding {
        receiver {
            install;
            mode (proxy | transparent);
            (source-list | source-vlans) vlan-list;
            translate;
        }
        source {
            groups group-prefix;
        }
    }
    disable;
    immediate-leave;
    interface (all | interface-name) {
        multicast-router-interface;
        static {
            group ip-address;
        }
    }
    proxy {
        source-address ip-address;
    }
    robust-count number;
    version number;
}
```

Syntax (EX4600, NFX Series, QFabric Systems, and QFX Series)

```
vlan vlan-name {
    immediate-leave;
    interface interface-name {
        group-limit limit;
        host-only-interface;
        multicast-router-interface;
        static {
            group multicast-group-address {
                source ip-address;
```

```
            }
        }
    }
    (l2-querier | igmp-querier (QFabric Systems only)) {
        source-address ip-address;
    }
    qualified-vlan ;
    proxy {
        source-address ip-address;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
}
```

Hierarchy Level

```
[edit protocols igmp-snooping]
```

Description

Configure IGMP snooping parameters for a VLAN (or all VLANs if you use the `all` option, where supported).

On legacy EX Series switches, which do not support the Enhanced Layer 2 Software (ELS) configuration style, IGMP snooping is enabled by default on all VLANs, and this statement includes a `disable` option if you want to disable IGMP snooping selectively on some VLANs or disable it on all VLANs. Otherwise, IGMP snooping is enabled on the specified VLANs if you configure any statements and options in this hierarchy.

> **NOTE**: You cannot configure IGMP snooping on a secondary (private) VLAN (PVLAN). However, starting in Junos OS Release 18.3R1 on EX4300 switches and EX4300 Virtual Chassis, and Junos OS Release 19.2R1 on EX4300 multigigabit switches, enabling IGMP snooping on a primary VLAN implicitly enables IGMP snooping on its secondary VLANs. See *IGMP Snooping on Private VLANs (PVLANs)* for details.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type `?` after `vlan` or `vlans` on the command line in configuration mode. Note that only one VLAN is displayed for a VLAN range, and for IGMP snooping, secondary private VLANs are not listed.

Default

On devices that support the `all` option, by default, IGMP snooping options apply to all VLANs . For all other devices, you must specify the `vlan` statement with a VLAN name to enable IGMP snooping.

Options

- `all`—All VLANs on the switch. This option is available only on EX Series switches that do not support the ELS configuration style.
- `disable`—Disable IGMP snooping on all or specified VLANs. This option is available only on EX Series switches that do not support the ELS configuration style.
- *`vlan-name`*—Name of a VLAN. A VLAN name must be provided on switches that support ELS to enable IGMP snooping.

TIP: On devices that support the `all` option, when you configure IGMP snooping parameters using the `vlan all` statement, any VLAN that is not individually configured for IGMP snooping inherits the `vlan all` configuration. Any VLAN that is individually configured for IGMP snooping, on the other hand, inherits none of its configuration from `vlan all`. Any parameters that are not explicitly defined for the individual VLAN assume their default values, not the values specified in the `vlan all` configuration.

For example, in the following configuration:

```
protocols {
    igmp-snooping {
        vlan all {
            robust-count 8;
        }
        vlan employee {
            interface ge-0/0/8.0 {
                static {
                    group 239.0.10.3
```

```
                    }
                }
            }
        }
    }
```

all VLANs, except `employee`, have a robust count of 8. Because `employee` has been individually configured, its robust count value is not determined by the value set under `vlan all`. Instead, its robust count is the default value of 2.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Statement updated with enhanced `?` (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.

RELATED DOCUMENTATION

Configuring IGMP Snooping

Configuring IGMP Snooping on Switches

Example: Configuring IGMP Snooping on Switches

Example: Configuring IGMP Snooping on EX Series Switches

Configuring VLAN-Specific IGMP Snooping Parameters

show igmp-snooping vlans | **20198**

vlan-isolation

IN THIS SECTION

- Syntax | 30826
- Hierarchy Level | 30826
- Description | 30826
- Required Privilege Level | 30827
- Release Information | 30827

Syntax

```
vlan-isolation;
```

Hierarchy Level

```
[edit chassis network-slices]
```

Description

Configure VLAN isolation. For the VLAN isolation configuration to take effect, you must configure the `set chassis network-slices vlan-isolation` CLI configuration statement on all the Routing Engines of the BSYS and GNFs, and then reboot the router.

On an MX without Junos Node Slicing enabled, all control traffic between the Routing Engine and the FPCs, as well as between the Routing Engines themselves, occur over an internal network. This network includes both the Routing Engines and all the FPCs. When you configure Junos node slicing, this internal network is shared by all the guest network function (GNF) Routing Engine pairs and their respective set of FPCs. The internal network has sufficient bandwidth to accommodate this traffic without additional configuration. However, if you want to isolate the internal traffic of each GNF from all others, use the `set chassis network-slices vlan-isolation` CLI configuration statement. This configuration assigns a separate internal VLAN for each GNF, resulting in all control traffic within each GNF being VLAN tagged and isolated from those of other GNFs.

NOTE:

- VLAN isolation is disabled by default.
- If you want to configure the sub line card feature, you must ensure that VLAN isolation is disabled. To disable VLAN isolation, delete the `set chassis network-slices vlan-isolation` configuration statement. See also: *Sub Line Card Overview*.

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

Setting Up Junos Node Slicing

vlan members (VLANs)

IN THIS SECTION

- Syntax | 30828
- Hierarchy Level | 30828
- Description | 30828
- Options | 30828
- Required Privilege Level | 30828
- Release Information | 30828

Syntax

```
vlan members [vlan-id];
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Description

Specify multiple VLAN identifiers to create a VLAN for each VLAN identifier.

Options

vlan-id—A list of valid VLAN identifiers. A VLAN is created for each VLAN identifier in the list.

> **NOTE**: If you specify a VLAN identifier list, you cannot configure an IRB interface in the VLAN.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.5.

RELATED DOCUMENTATION

Example: Configuring VLANs on Security Devices

vlan (Mirroring)

IN THIS SECTION

- Syntax | 30829
- Hierarchy Level | 30829
- Description | 30829
- Options | 30829
- Required Privilege Level | 30830
- Release Information | 30830

Syntax

```
vlan (vlan-id | vlan-name);
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input ingress],
[edit forwarding-options analyzer analyzer-name input ingress routing-instance instance-name],
[edit forwarding-options analyzer analyzer-name output],
[edit forwarding-options analyzer analyzer-name output routing-instance instance-name]
```

Description

Configure mirrored traffic to be sent to a VLAN for remote monitoring.

Options

vlan-id	Numeric VLAN identifier.
vlan-name	Name of the VLAN.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10 (ELS).

RELATED DOCUMENTATION

Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX9200 Switches

Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use

vlan (MLD Snooping)

IN THIS SECTION

- Syntax | 30830
- Hierarchy Level | 30831
- Description | 30831
- Default | 30831
- Options | 30832
- Required Privilege Level | 30832
- Release Information | 30833

Syntax

```
vlan (all | vlan-name) {
    disable;
    immediate-leave;
    interface ( all | interface-name) {
```

```
        group-limit  limit;
        host-only-interface;
        immediate-leave;
        multicast-router-interface;
        static {
            group ip-address {
                source ip-address;
            }
        }
    }
    qualified-vlan;
    query-intervalseconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier>;
    }
    version version;
}
```

Hierarchy Level

```
[edit protocols mld-snooping]
[edit routing-instances instance-name protocols mld-snooping]
```

Description

Configure MLD snooping parameters for a VLAN.

When the **vlan** configuration statement is used without the `disable` statement, MLD snooping is enabled on the specified VLAN or on all VLANs.

Default

If the `vlan` statement is not included in the configuration, MLD snooping is disabled.

Options

all (All EX Series switches except EX9200) Configure MLD snooping parameters for all VLANs on the switch.

vlan-name Configure MLD snooping parameters for the specified VLAN.

TIP: When you configure MLD snooping parameters using the `vlan all` statement, any VLAN that is not individually configured for MLD snooping inherits the **vlan all** configuration. Any VLAN that is individually configured for MLD snooping, on the other hand, inherits none of its configuration from **vlan all**. Any parameters that are not explicitly defined for the individual VLAN assume their default values, not the values specified in the **vlan all** configuration.

For example, in the following configuration:

```
protocols {
    mld-snooping {
        vlan all {
            robust-count 8;
        }
        vlan employee {
            interface ge-0/0/8.0 {
                static {
                    group ff1e::1;
                }
            }
        }
    }
}
```

all VLANs, except **employee**, have a robust count of 8. Because **employee** has been individually configured, its robust count value is not determined by the value set under **vlan all**. Instead, its robust count is the default value of 2.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

Support at the `[edit routing-instances instance-name protocols mld-snooping]` hierarchy introduced in Junos OS Release 13.3 for EX Series switches.

Support for the `qualified-vlan`, `query-interval`, `query-last-member-interval`, `query-response-interval`, and `traceoptions` statements introduced in Junos OS Release 13.3 for EX Series switches.

RELATED DOCUMENTATION

> *Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)*

vlan-model

IN THIS SECTION

- Syntax | 30833
- Hierarchy Level | 30834
- Description | 30834
- Options | 30834
- Required Privilege Level | 30834
- Release Information | 30834

Syntax

```
vlan-model one-to-one;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Define the network VLAN model.

Options

`one-to-one`—Specify that any received, dual-tagged VLAN packet triggers the provisioning process in a Layer 2 Wholesale network. Using this option, the router learns VLAN tags for each individual client. The router learns both the outer tag and inner tag of the incoming packets, when the `instance-role` statement is defined as `access`, or the outer VLAN tag only, when the `instance-role` statement is defined as `nni`.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers

Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers

vlan (MSTP)

IN THIS SECTION

- Syntax | 30835
- EX Series | 30835
- Hierarchy Level | 30836
- Description | 30836
- Options | 30836
- Required Privilege Level | 30837
- Release Information | 30837

Syntax

```
vlan vlan-id;
```

EX Series

```
vlan (all | vlan-id | vlan-name) {
     bridge-priority priority;
     forward-delay seconds;
     hello-time seconds;
     interface interface-name {
         bpdu-timeout-action {
             block;
             log;
        }
         cost  cost;
         disable;
         edge;
         mode mode;
         no-root-port;
         priority priority;
    }
     max-age seconds;
```

```
     traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable | no-world-
readable>;
        flag flag;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mstp msti msti-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name
protocols mstp msti msti-id]
```

```
[edit protocols mstp msti msti-id],
[edit protocols mstp msti msti-id]
[edit protocols vstp]
```

```
[edit routing-instances routing-instance-name protocols mstp msti msti-id]
```

Description

Configure the VLANs for a Multiple Spanning Tree Instance (MSTI) or VSTP instance.

NOTE: When you configure VSTP with the `set prototocol vstp vlan all` command, vlan-id 1 is excluded to be compatible with Cisco PVST+. If you want vlan-id 1 to be included in VSTP, you must set it separately with the `set protocol vstp vlan 1` command.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type **?** after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options

vlan-name—Name of the VLAN.

vlan-id—The VLAN identifier associated with the MSTI.

vlan-id-range—Range of VLAN identifiers associated with the MSTI in the form *minimum-vlan-id-maximum-vlan-id*. VLAN identifier ranges are not supported for VSTP.

- **Range:** 1 through 4096

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Statement updated with enhanced **?** (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding MSTP

Understanding VSTP

Example: Configuring Network Regions for VLANs with MSTP on Switches

vlan (PIM Snooping)

IN THIS SECTION

- Syntax | 30838
- Hierarchy Level | 30838
- Description | 30838
- Required Privilege Level | 30838
- Release Information | 30838

Syntax

```
vlan <vlan-id>{
no-dr-flood;
}
```

Hierarchy Level

```
[edit routing-instances <instance-name> protocols pim-snooping],
[edit logical-systems <logical-system-name> routing-instances <instance-name> protocols pim-
snooping]
```

Description

Configure PIM snooping parameters for a VLAN.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

PIM Overview

Configuring Basic PIM Settings

vlan (Port Mirroring)

IN THIS SECTION

- Syntax | 30839
- Hierarchy Level | 30839
- Description | 30839
- Options | 30840
- Required Privilege Level | 30840
- Release Information | 30840

Syntax

```
vlan (vlan-id | vlan-name) {
    No Link Title;
```

Hierarchy Level

For platforms without ELS:

```
[edit ethernet-switching-options analyzer name input ingress],
[edit ethernet-switching-options analyzer name output]
```

For platforms with ELS:

```
[edit forwarding-options analyzer name input (egress | ingress)]
[edit forwarding-options analyzer name output]
[edit forwarding-options port-mirroring [instance name] family ethernet-switching output]
```

Description

When used in an `input` statement, specifies that traffic entering or exiting a VLAN should be mirrored. (You can include this statement in an `ingress` statement or `egress` statement within the `input` statement. It

is not supported in an `egress` statement on all switches) When used in an `output` statement, specifies that mirrored traffic to be sent to a VLAN for remote monitoring.

On some switches, only one interface can be a member of the output (analyzer) VLAN. This limitation does not apply on the QFX10000 switch if traffic is mirrored on ingress. In this case, multiple QFX10000 interfaces can belong to the output VLAN, and traffic is mirrored to all of those interfaces. If traffic is mirrored on egress on a QFX10000 switch, only one interface can be a member of the analyzer VLAN.

Options

vlan-id—Numeric VLAN identifer.

vlan-name—Name of the VLAN.

`no-tag`—Specifies that remote mirrored packets are not tagged with the tag of the output (analyzer) VLAN.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

Option`output vlan` added in Junos OS Release 12.1 for the QFX Series.

Option`no-tag` added in Junos OS Release 15.1X53-D10 for QFX10000 switches.

RELATED DOCUMENTATION

> *Understanding Port Mirroring and Analyzers*
>
> *Configuring Port Mirroring Local and Remote Analysis*

vlan (Port Mirroring)

IN THIS SECTION

- Syntax | 30841
- Hierarchy Level | 30841
- Description | 30841
- Options | 30841
- Required Privilege Level | 30841
- Release Information | 30842

Syntax

```
vlan (vlan-name | vlan-ID);
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family ethernet-switching output]
[edit forwarding-options port-mirroring instance instance-name family ethernet-switching output]
```

Description

Specify the VLAN name or ID for sending copies of packets to an analyzer. This configuration enables remote VLAN port mirroring on EX Series switches.

Options

vlan-name—Name of remote mirroring VLAN.

vlan-ID—ID of the remote mirroring VLAN.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Example: Configuring Layer 2 Port Mirroring to Remote VLAN

Layer 2 Port Mirroring to Remote Destination by Using Destination as VLAN

vlan-prune

IN THIS SECTION

- Syntax | 30842
- Hierarchy Level | 30842
- Description | 30843
- Default | 30843
- Required Privilege Level | 30843
- Release Information | 30843

Syntax

```
vlan-prune;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Description

Prune the Virtual Chassis port (VCP) paths in a Virtual Chassis to ensure received broadcast, multicast, and unknown unicast traffic in a VLAN uses the shortest possible path through the Virtual Chassis to the egress VLAN interface.

By default, all broadcast, multicast, and unknown unicast traffic in a VLAN on an EX Series Virtual Chassis is broadcast to all member switches in the Virtual Chassis. This behavior unnecessarily consumes bandwidth within the Virtual Chassis because unneeded traffic is sent to all Virtual Chassis member switches.

Enabling this option allows you to conserve bandwidth within the Virtual Chassis. Broadcast, multicast, and unknown unicast traffic still enters and exits the Virtual Chassis within the same VLAN, without the added bandwidth consumption that results from broadcasting this traffic to all member switches.

Default

Disabled

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Configuring VLANs for EX Series Switches

vlan-range

IN THIS SECTION

Syntax | 30844

- Hierarchy Level | 30844
- Description | 30844
- Default | 30844
- Options | 30844
- Required Privilege Level | 30844
- Release Information | 30845

Syntax

```
vlan-range vlan-id-low-vlan-id-high;
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Description

Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

Default

None.

Options

vlan-id-low-vlan-id-high—Specify the first and last VLAN ID number for the group of VLANs.

Required Privilege Level

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

- *Configuring VLANs on Switches*
- *Configuring VLANs for EX Series Switches*
- Configuring VLANs for EX Series Switches (J-Web Procedure)
- Configuring Routed VLAN Interfaces on Switches (CLI Procedure)
- *Understanding Bridging and VLANs on Switches*
- *Configuring IRB Interfaces on Switches*

vlan-ranges

IN THIS SECTION

- Syntax | 30845
- Hierarchy Level | 30846
- Description | 30846
- Required Privilege Level | 30846
- Release Information | 30847

Syntax

```
vlan-ranges {
    access-profile profile-name;
    authentication {
        packet-types  [packet-types];
        password password-string;
        username-include {
```

```
            circuit-type;
            circuit-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-18;
            option-37;
            option-82 <circuit-id> <remote-id>;
            radius-realm radius-realm-string;
            remote-id;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | inet);
        accept-out-of-band protocol;
        access-profile vlan-dynamic-profile-name;
        ranges (any | low-tag)-(any | high-tag);
    }
    override;
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure]
```

Description

Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

Configuring Interfaces to Support Both Single and Stacked VLANs

vlan-ranges (RADIUS Options)

IN THIS SECTION

- Syntax | 30847
- Hierarchy Level | 30847
- Description | 30847
- Options | 30848
- Required Privilege Level | 30848
- Release Information | 30848

Syntax

```
vlan-ranges (any | low-tag-high-tag);
```

Hierarchy Level

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
```

Description

Configure the VLAN range of subscribers to which the named NAS-Port options definition applies.

NOTE: You can configure a maximum of 16 NAS-Port options definitions per physical interface. Each definition can include a maximum of 32 VLAN ranges or 32 S-VLAN ranges, but cannot include a combination of VLAN ranges and S-VLAN ranges.

Options

any Entire VLAN range representing all VLAN IDs.

low-tag VLAN ID tag representing the lower limit of the VLAN range.

- **Range:** 1 through 4094

high-tag VLAN ID tag representing the upper limit of the VLAN range.

- **Range:** 1 through 4094

NOTE: To specify a single VLAN ID, set *low-tag* and *high-tag* to the same value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

RADIUS NAS Port Attributes and Options

RADIUS NAS Port Attributes and Options

vlan-rewrite

IN THIS SECTION

- Syntax | 30849
- Hierarchy Level | 30849
- Description | 30849
- Required Privilege Level | 30850
- Release Information | 30850

Syntax

```
vlan-rewrite {
    translate from-vlan-id to-vlan-id;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit number family bridge ]
[edit interfaces interface-name unit number family ethernet-switching]
```

Description

Translate a VLAN ID (the *from-vlan-id*) in an incoming packet to another VLAN ID (the *to-vlan-id*) that you have configured on an interface. The device does the corresponding translation in reverse at egress.

You can configure VLAN translation with the `vlan-rewrite translate` statement as follows:

- Only on trunk mode interfaces.
- The *to-vlan-id* is a VLAN you configure on the device with the specified interface as a member of that VLAN.
- The *from-vlan-id* and *to-vlan-id* VLAN IDs can have values between 1 and 4094.

After you set up a VLAN mapping, the device accepts incoming packets tagged with *from-vlan-id* and translates that VLAN ID into the *to-vlan-id* to process the packet. We refer to the *to-vlan-id* as the *mapped VLAN value*.

The device also accepts packets with VLAN IDs for which you did not configure VLAN translation on an interface, and passes those through to the packet processing pipeline without translation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Rewriting a VLAN Tag and Adding a New Tag

Overlapping VLAN Support Using VLAN Translation in EVPN-VXLAN Networks

vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP)

IN THIS SECTION

- Syntax | 30851
- Hierarchy Level | 30851
- Description | 30851
- Options | 30851
- Required Privilege Level | 30851
- Release Information | 30851

Syntax

```
vlan-rule  (high-low | odd-even);
```

Hierarchy Level

```
[edit chassis fpc slot pic slot forwarding-mode vlan-steering]
```

Description

Configure the interoperation mode of the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) when interoperating with 100 gigabit Ethernet interfaces from other vendors.

If no VLAN rule is configured, all tagged packets are distributed to PFE0.

Options

high-low VLAN IDs 1 through 2047 are distributed to PFE0 and VLAN IDs 2048 through 4096 are distributed to PFE1.

odd-even Odd number VLAN IDs are distributed to PFE1 and even number VLAN IDs are distributed to PFE0.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP

forwarding-mode (100-Gigabit Ethernet) | **7459**

vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) | 30855

vlan (Secure Access Port)

IN THIS SECTION

- Syntax | 30852
- Hierarchy Level | 30853
- Description | 30853
- Options | 30853
- Required Privilege Level | 30854
- Release Information | 30854

Syntax

```
vlan (all | vlan-name) {
    examine-fip {
        examine-vn2vn {
            beacon-period milliseconds;
        }
        fc-map fc-map-value;
        no-fip-snooping-scaling;
    }
    dhcp-option82
        circuit-id {
            prefix (Circuit ID for Option 82) hostname;
            use-interface-description;
            use-vlan-id;
        }
        remote-id {
            prefix (Remote ID for Option 82) hostname | mac | none;
            use-interface-description;
            use-string                          string;
        }
        vendor-id <string>;
```

```
    }
    (arp-inspection | no-arp-inspection);
        circuit-id {
            prefix (Circuit ID for Option 82) hostname;
            use-interface-description;
            use-vlan-id;
        }
        remote-id {
            prefix (Remote ID for Option 82) hostname | mac | none;
            use-interface-description;
             use-string string;
        }
    vendor-id <string>;
    }
    (examine-dhcp | no-examine-dhcp);
    mac-move-limitlimit action action;
}
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port]
```

Description

Apply DHCP snooping, dynamic ARP inspection (DAI), DHCP option 82, and MAC move limiting.

> **TIP**: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type **?** after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

The remaining statements are explained separately. See CLI Explorer.

Options

all—Apply DHCP snooping, DAI, DHCP option 82, and MAC move limiting to all VLANs.

vlan-name —Apply DHCP snooping, DAI, DHCP option 82, and MAC move limiting to the specified VLAN.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

- *Understanding How to Protect Access Ports from Common Attacks*
- Understanding MAC Limiting and MAC Move Limiting for Port Security
- *Understanding and Using Trusted DHCP Servers*
- *Configuring MAC Limiting (QFX Switches)*
- *Understanding and Using Trusted DHCP Servers*

vlan (Static IP)

IN THIS SECTION

- Syntax | 30854
- Hierarchy Level | 30855
- Description | 30855
- Options | 30855
- Required Privilege Level | 30855
- Release Information | 30855

Syntax

```
vlan vlan-name;
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port interface (all | interface-name) static-ip ip-address]
```

Description

Associate a static IP address with the specified VLAN.

Options

vlan-name—Name of a VLAN associated with the specified interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Configuring Static DHCP IP Addresses for DHCP snooping (non-ELS)

vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP)

IN THIS SECTION

- Syntax | 30856
- Hierarchy Level | 30856
- Description | 30856

Required Privilege Level | 30856
Release Information | 30856

Syntax

```
vlan-steering {
    vlan-rule (high-low | odd-even);
}
```

Hierarchy Level

```
[edit chassis fpc slot pic slot forwarding-mode]
```

Description

Configure the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) to interoperate with 100 gigabit Ethernet interfaces from other vendors.

The other statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP

forwarding-mode (100-Gigabit Ethernet) | 7459

sa-multicast (100-Gigabit Ethernet) | 16264
vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) | 30850

vlan-tag

IN THIS SECTION

- Syntax | 30857
- Hierarchy Level | 30857
- Description | 30857
- Default | 30857
- Options | 30858
- Required Privilege Level | 30858
- Release Information | 30858

Syntax

```
vlan-tag (outer | outer-and-inner);
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules
ieee-802.1 (rewrite-name | default)]
```

Description

Apply this IEEE-802.1 rewrite rule to the outer VLAN tag or, if available, both outer and inner VLAN tags.

Default

If you do not include this statement, the rewrite rule applies to the outer VLAN tag only.

Options

`outer`—Apply the rewrite rule to the outer VLAN tag only.

`outer-and-inner`—Apply the rewrite rule to both the outer and inner VLAN tags.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Applying IEEE 802.1p Rewrite Rules to Dual VLAN Tags

vlan-tag (Dynamic Classifiers)

IN THIS SECTION

- Syntax | 30859
- Hierarchy Level | 30859
- Description | 30859
- Default | 30859
- Options | 30859
- Required Privilege Level | 30859
- Release Information | 30859

Syntax

```
vlan-tag (inner | outer);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-number classifiers ieee-802.1]
```

Description

Apply this IEEE-802.1 classifier to the inner or outer VLAN tags in a dynamic profile.

Default

If you do not include this statement, the classifier applies to the outer VLAN tag only.

Options

`inner`—Apply the classifier to the inner VLAN tag only.

`outer`—Apply the classifier to the outer VLAN tag only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying a Classifier to a Subscriber Interface in a Dynamic Profile

classifiers (Definition)

vlan-tag (Dynamic Rewrite Rules)

IN THIS SECTION

- Syntax | 30860
- Hierarchy Level | 30860
- Description | 30860
- Default | 30860
- Options | 30861
- Required Privilege Level | 30861
- Release Information | 30861

Syntax

```
vlan-tag (outer | outer-and-inner);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-number rewrite-rules ieee-802.1]
```

Description

Apply this IEEE-802.1 rewrite rule to the outer or outer and inner VLAN tags in a dynamic profile.

Default

If you do not include this statement, the rewrite rule applies to the outer VLAN tag only.

Options

`outer`—Apply the rewrite rule to the outer VLAN tag only.

`outer-and-inner`—Apply the rewrite rule to both the outer and inner VLAN tags.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Guidelines for Configuring Dynamic CoS for Subscriber Access

Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile

rewrite-rules

vlan-tagging

IN THIS SECTION

- Syntax | **30862**
- Syntax (QFX Series, NFX Series, and EX4600) | **30862**
- Syntax (SRX Series Interfaces) | **30862**
- Hierarchy Level | **30862**
- QFX Series, NFX Series, and EX4600 Interfaces | **30862**
- SRX Series Interfaces | **30862**
- Description | **30863**
- Default | **30863**
- Options | **30863**

Required Privilege Level | 30863
Release Information | 30863

Syntax

```
vlan-tagging;
```

Syntax (QFX Series, NFX Series, and EX4600)

```
vlan-tagging;
```

Syntax (SRX Series Interfaces)

```
vlan-tagging native-vlan-id vlan-id;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit logical-systems logical-system-name interfaces interface-name]
```

QFX Series, NFX Series, and EX4600 Interfaces

```
[edit interfaces (QFX Series) interface-name ]
[edit interfaces (QFX Series) interface-range interface-range-name ]
```

SRX Series Interfaces

```
[edit interfaces interface ]
```

Description

For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

> **NOTE**: For QFX Series configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface. Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.
>
> On EX Series switches except for EX4300 and EX9200 switches, the `vlan-tagging` and `family ethernet-switching` statements cannot be configured on the same interface. Interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to `family ethernet-switching` by the default factory configuration. EX6200 and EX8200 switch interfaces do not have a default `family` setting.

Default

VLAN tagging is disabled by default.

Options

`native-vlan-id`— (SRX Series)Configures a VLAN identifier for untagged packets. Enter a number from 0 through 4094.

> **NOTE**: The `native-vlan-id` can be configured only when either `flexible-vlan-tagging` mode or `interface-mode` trunk is configured.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

- 802.1Q VLANs Overview
- Configuring a Layer 3 Subinterface (CLI Procedure)
- Configuring Tagged Aggregated Ethernet Interfaces
- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch
- vlan-id
- Configuring a Layer 3 Logical Interface
- Configuring VLAN Tagging

vlan-tagging (Dynamic)

IN THIS SECTION

- Syntax | 30864
- Hierarchy Level | 30864
- Description | 30865
- Required Privilege Level | 30865
- Release Information | 30865

Syntax

```
vlan-tagging;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name],
[edit interfaces interface-name]
```

Description

For Fast Ethernet and Gigabit Ethernet interfaces and aggregated Ethernet interfaces configured for VPLS, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

> **NOTE**: For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, the Junos OS supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or bridging domain.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

- *Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs*
- *Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs*
- *Configuring the L2TP LNS Peer Interface*

vlan-tags

IN THIS SECTION

- Syntax | 30866
- Hierarchy Level | 30866
- Description | 30866
- Options | 30866

Required Privilege Level | 30866
Release Information | 30866

Syntax

```
vlan-tags outer number inner number;
```

Hierarchy Level

```
[edit bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains
bridge-domain-name],
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]
[edit vlans vlan-name]
```

Description

Specify dual VLAN identifier tags for a bridge domain, VLAN, or VPLS routing instance.

Options

outer *number*—A valid VLAN identifier.

inner *number*—A valid VLAN identifier.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

Configuring a Bridge Domain

Configuring a VLAN

Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances

Configuring VLAN Identifiers for VLANs and VPLS Routing Instances

Configuring a Layer 2 Virtual Switch

Configuring a Layer 2 Virtual Switch on an EX Series Switch

vlan-tags (Dynamic Profiles)

IN THIS SECTION

- Syntax | 30867
- Hierarchy Level | 30867
- Description | 30868
- Options | 30868
- Required Privilege Level | 30868
- Release Information | 30868

Syntax

```
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
```

Description

For Gigabit Ethernet IQ and IQE interfaces only, binds TPIDs and 802.1Q VLAN tag IDs to a logical interface. You must include the `stacked-vlan-tagging` statement at the `[edit interfaces interface-name]` hierarchy level.

> **NOTE**: The `inner-range vid1-vid2` option is supported on IQE PICs only.

Options

`inner [tpid].vlan-id`

A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the `dynamic-profiles` hierarchy, specify the `$junos-vlan-id` predefined variable to dynamically obtain the VLAN ID.

> **NOTE**: On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the `inner-range tpid. vid1–vid2` option with the `vlan-tags` statement for ISP-facing interfaces.

- **Range:** For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

`outer [tpid].vlan-id`

A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the `dynamic-profiles` hierarchy, specify the `$junos-stacked-vlan-id` predefined variable.

- **Range:** For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

VLAN demux interface support introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring Dual VLAN Tags

vlan-tags (Dual-Tagged Logical Interface)

IN THIS SECTION

- Syntax | 30869
- Hierarchy Level | 30869
- Description | 30869
- Options | 30870
- Required Privilege Level | 30871
- Release Information | 30871

Syntax

```
vlan-tags inner-list [vlan-id vlan-id-vlan-id ] outer <tpid.>vlan-id;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

(MX Series routers only) Binds a dual-tag logical interface to a list of VLAN IDs. Configures the logical interface to receive and forward any dual-tag frame whose inner VLAN ID tag matches the list of VLAN IDs you specify.

NOTE: To create a circuit cross-connect (CCC) using VLAN-bundled dual-tag logical interfaces on Layer 2 VPN routing instances, you must include the `encapsulation-type` statement and specify the value `ethernet-vlan` at the one of the following hierarchy levels:

- `[edit routing-instances routing-instance-name protocols l2vpn]`
- `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn]`

For more information about the `encapsulation-type` configuration statement and the Layer 2 encapsulation types `ethernet` and `ethernet-vlan`, see the Junos OS VPNs Library for Routing Devices.

Options

`inner-list [vlan-id vlan-id vlan-id-vlan-id]`—A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

- **Range:** 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

`outer <tpid.>vlan-id`—An optional Tag Protocol ID (TPID) and a valid VLAN ID.

- **Range:** For TPID, specify a hexadecimal value in the format `0xnnnn`.
- **Range:** For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

NOTE: Configuring `inner-list` with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1 through 4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

1. **Inefficient**

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
```

```
        vlan-tags outer vid inner-list 1-4094;
    }
```

2. **Best Practice**

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id vid;
}
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Binding VLAN IDs to Logical Interfaces

encapsulation (Logical Interface) | **6300**

encapsulation (interfaces) | **6289**

Junos OS VPNs Library for Routing Devices

flexible-vlan-tagging | **7235**

vlan-id-list (Ethernet VLAN Circuit) | **30808**

vlan-tagging | **30861**

vlan-tags-outer

IN THIS SECTION

- Syntax | 30872
- Hierarchy Level | 30872
- Description | 30872
- Required Privilege Level | 30872
- Release Information | 30872

Syntax

```
vlan-tags-outer vlan-tag;
```

Hierarchy Level

```
[edit interfaces interface-set interface-set-name interface interface-name]
```

Description

The S-VLAN outer tag that belongs to a set of interfaces used to configure hierarchical CoS schedulers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Junos OS Class of Service User Guide for Routing Devices

vlan-tags (Stacked VLAN Tags)

IN THIS SECTION

- Syntax | 30873
- Hierarchy Level | 30873
- Description | 30873
- Options | 30874
- Required Privilege Level | 30876
- Release Information | 30876

Syntax

```
vlan-tags inner tpid.vlan-id inner-list value inner-range vid1–vid2 outer tpid.vlan-id;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Bind TPIDs and 802.1Q VLAN tag IDs to a logical interface. TPID fields are used to identify the frame as an IEEE 802.1Q-tagged frame.

Options

inner *tpid.vlan-id*

A TPID and a valid VLAN identifier. TPID is a 16-bit field set to a value of 0x8100 in order to identify the frame as an IEEE 802.1Q-tagged frame.

- **Range:** (most routers) For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported.

inner-list *value*

List or a set of VLAN identifiers. You can configure a maximum of 32 inner-lists for an interface.

> **NOTE**: This is supported on MX Series routers with Trio-based FPCs.

inner-range *tpid. vid1-vid2*

Specify a TPID and a range of VLAN IDs where vid1 is the start of the range and vid2 is the end of the range.

> **NOTE**: On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the `inner-range` *tpid. vid1-vid2* option with the `vlan-tags` statement for ISP-facing interfaces.

- **Range:** For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer *tpid.vlan-id*

A TPID and a valid VLAN identifier.

- **Range:** (most routers) For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported.

> **NOTE**: Configuring `inner-range` with the entire vlan-id range consumes system resources and is not a best practice. The `inner-range` must be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1 through 4094), it has the same result as not specifying a range; however, it consumes Packet Forwarding Engine resources such as VLAN lookup table entries, and so on.
>
> The following examples illustrate this further:

1. **Inefficient**

```
[edit interfaces  interface-name]stacked-vlan-tagging;
unit number {
    vlan-tags outer vid inner-range 1-4094;
}
```

2. **Best Practice**

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id vid;
}
```

NOTE: Configuring `inner-range` with the entire vlan-id range consumes system resources and is not a best practice. The `inner-range` must be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1 through 4094), it has the same result as not specifying a range; however, it consumes Packet Forwarding Engine resources such as VLAN lookup table entries, and so on.

The following examples illustrate this further:

1. **Inefficient**

```
[edit interfaces  interface-name]
stacked-vlan-tagging;
unit number {
    vlan-tags outer vid inner-range 1-4094;
}
```

2. **Best Practice**

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
```

```
                vlan-id vid;
            }
```

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring Dual VLAN Tags
- Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers
- stacked-vlan-tagging

vlan (Unknown Unicast Forwarding)

IN THIS SECTION

- Syntax | 30877
- Hierarchy Level | 30877
- Description | 30877
- Options | 30877
- Required Privilege Level | 30877
- Release Information | 30877

Syntax

```
vlan  (all | vlan-name) {
    interface (Unknown Unicast Forwarding) interface-name;
}
```

Hierarchy Level

```
[edit ethernet-switching-options unknown-unicast-forwarding]
```

Description

Specify a VLAN from which unknown unicast packets will be forwarded or specify that the packets will be forwarded from all VLANS. Unknown unicast packets are forwarded from a VLAN to a specific trunk interface.

The `interface` statement is explained separately.

> **TIP**: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type **?** after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options

all—All VLANs.

vlan-name—Name of a VLAN.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Statement updated with enhanced **?** (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.

RELATED DOCUMENTATION

show vlans

show ethernet-switching table | 19700

Configuring Unknown Unicast Forwarding (CLI Procedure)

vlan (VSTP)

IN THIS SECTION

- Syntax | 30878
- EX Series | 30879
- Hierarchy Level | 30879
- Description | 30879
- Options | 30880
- Required Privilege Level | 30880
- Release Information | 30880

Syntax

```
vlan vlan-id {
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    max-age seconds;
    interface interface-name {
        cost cost;
        edge;
        mode (p2p | shared);
        no-root-port;
        priority interface-priority;
    }
}
```

EX Series

```
                        vlan (all | vlan-id | vlan-name) {
                            bridge-priority                        priority;
                            forward-delay                        seconds;
                            hello-time                        seconds;
                                interface (all | interface-name)  {
        bpdu-timeout-action {
                                        block;
                                        log;
        }
                                cost                                cost;
                                disable;
                                edge;
                                mode                                mode;
                                no-root-port;
                                priority                                priority;
    }
                            max-age                        seconds;
                                traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable | no-world-
readable>;
        flag flag;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols vstp],
[edit protocols vstp]
```

Description

Configure VSTP VLAN parameters.

> **TIP**: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type **?** after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options

`all`	All VLANs.
vlan-id	Numeric VLAN identifier.
vlan-range	Name of the VLAN range.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for logical systems added in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Understanding VSTP
- Configuring VLAN Spanning Tree Protocol

vlans

IN THIS SECTION

- Syntax (QFX Series, QFabric, NFX Series, and EX4600) | **30881**
- Syntax (QFX Series, NFX Series, and EX4600) | **30882**
- Syntax (SRX Series and EX Series) | **30885**
- Syntax (SRX Series) | **30885**
- Syntax (vSRX Virtual Firewall) | **30889**

- Hierarchy Level | 30892
- Description | 30892
- Default | 30893
- Options | 30893
- Required Privilege Level | 30893
- Release Information | 30893

Syntax (QFX Series, QFabric, NFX Series, and EX4600)

```
vlans {
     vlan-name {
        description text-description;
        dot1q-tunneling {
            customer-vlans (id | range);
        }
        filter input filter-name;
         filter output filter-name;{
            isolated;
            mapping  (policy | tag push | native push);
            promiscuous;
        }
        isolation-vlan-id;
        l3-interface vlan.logical-interface-number;
        mac-limit number;
        no-local-switching;
        no-mac-learning;
        primary-vlan vlan-name;
        pvlan extend-secondary-vlan-id vlan-id;
        vlan-id number;
        vlan-range vlan-id-low-vlan-id-high;
    }
}
```

Syntax (QFX Series, NFX Series, and EX4600)

```
vlans {
    vlan-name {
        description text-description;
        domain-type bridge;
        forwarding-options {
            dhcp-security {
                arp-inspection;
                group group-name {
                    interface interface-name {
                        static-ip ip-address {
                            mac mac-address;
                        }
                    }
                    overrides {
                        no-option82;
                        trusted;
                        untrusted;
                    }
                }
                ip-source-guard;
                no-dhcp-snooping;
                option-82 {
                    circuit-id {
                        prefix {
                            host-name;
                            logical-system-name;
                            routing-instance-name;
                        }
                        use-interface-description (device | logical);
                        use-vlan-id;
                    }
                    remote-id {
                        host-name hostname;
                        use-interface-description (device | logical);
                        use-string string;
                    }
                    vendor-id {
                        use-string string;
                    }
                }
```

```
            }
            fip-security {
                examine-vn2vf;
                examine-vn2vn {
                    beacon-period milliseconds;
                }
                fc-map fc-map-value;
                interface interface-name {
                    (fcoe-trusted | no-fcoe-trusted;)
                }
            }
        }
        l3-interface irb.logical-unit-number;
        multicast-snooping-options {
            flood-groups [group-names];
            forwarding-cache {
                threshold {
                    reuse threshold;
                    suppress threshold;
                }
            }
            graceful-restart {
                disable;
                restart-duration duration;
            }
            host-outbound-traffic {
                dot1p bits;
                forwarding-class forwarding-class;
            }
            multichassis-lag-replicate-state;
            nexthop-hold-time time;
            options {
                syslog {
                    level level;
                    mark interval;
                    upto level;
                }
            }
            traceoptions {
                file filename {
                    files number;
                    no-world-readable;
                    size file-size;
```

```
                    world-readable;
                }
                flag flag {
                    disable;
                }
            }
        }
        switch-options {
            interface interface-name {
                interface-mac-limit limit {
                    packet-action action;
                }
                static-mac mac-address;
            }
            interface-mac-limit limit {
                packet-action action;
            }
            mac-move-limit limit {
                packet-action action;
            }
            mac-table-size limit {
                packet-action drop;
            }
            no-mac-learning;
            }
        }
        vlan-id number;
        vlan-id-list [vlan-id | vlan-id-vlan-id];
        vlan-tags
            inner value;
            outer value;
        }
        vxlan {
            ingress-node-replication
            ovsdb-managed
            }
        }
    }
}
```

Syntax (SRX Series and EX Series)

```
vlans {
    vlan-name {
        description text-description;
        dot1q-tunneling  {
            customer-vlans (id | range)
            layer2-protocol-tunneling all | protocol-name {
                drop-threshold number;
                shutdown-threshold number;
            }
        }
        filter input filter-name;
        filter output filter-name;{
            egress;
            ingress;
            mapping (native (push | swap) | policy | tag (push | swap));
            pvlan-trunk;
        }
        isolation-id id-number;
        l3-interface l3-interface-name.logical-interface-number;
        l3-interface-ingress-counting layer-3-interface-name;
        mac-limit limit action action;
        mac-table-aging-time seconds;
        no-local-switching (VLAN);
        no-mac-learning;
        primary-vlan vlan-name;
        vlan-id number;
        vlan-prune;
        vlan-range vlan-id-low-vlan-id-high;
    }
}
```

Syntax (SRX Series)

```
vlans {
    vlan name {
        (vlan-id (1..3967) | vlan-id-list [ vlan-id-numbers]);
        description;
        forwarding-options {
```

```
            dhcp-security {
                arp-inspection;
                dhcpv6-options {
                    option-16 {
                        use-string use-string;
                    }
                    option-18 {
                        prefix {
                            host-name;
                            logical-system-name;
                            routing-instance-name;
                            vlan-id;
                            vlan-name;
                        }
                        use-interface-description (device | logical);
                        use-interface-index (device | logical);
                        use-interface-mac;
                        use-interface-name (device | logical);
                        use-string use-string;
                    }
                    option-37 {
                        prefix {
                            host-name;
                            logical-system-name;
                            routing-instance-name;
                            vlan-id;
                            vlan-name;
                        }
                        use-interface-description (device | logical);
                        use-interface-index (device | logical);
                        use-interface-mac;
                        use-interface-name (device | logical);
                        use-string use-string;
                    }
                }
                group group-name {
                    interface interface-name {
                        static-ip {
                            ip-address {
                            mac-address;
                            }
                        }
                        static-ipv6 {
```

```
                    ip-address {
                    mac-address;
                    }
                }
            }
            overrides {
                no-dhcpv6-options;
                no-option16;
                no-option18;
                no-option37;
                no-option82;
                trusted;
                untrusted;
            }
        }
        ip-source-guard;
        ipv6-source-guard;
        neighbor-discovery-inspection;
        no-dhcp-snooping;
        no-dhcpv6-snooping;
        option-82 {
            circuit-id {
                prefix {
                    host-name;
                    logical-system-name;
                    routing-instance-name;
                }
                use-interface-description (device | logical);
                use-vlan-id;
            }
            remote-id {
                host-name;
                mac;
                use-interface-description (device | logical);
                use-string use-string;
            }
            vendor-id {
                use-string use-string;
            }
        }
    }
    filter {
        input filter-name;
```

```
            }
            flood {
                input filter-name;
            }
        }
        interface interface-name;
        l3-interface l3-interface-name;
        mcae-mac-flush;
        mcae-mac-synchronize;
        service-id service-id;
        switch-options {
            interface name {
                action-priority action-priority;
                encapsulation-type (ethernet | ethernet-vlan);
            ignore-encapsulation-mismatch;
            interface-mac-limit {
                limit;
                packet-action (drop | drop-and-log | log | none | shutdown);
            }
            no-mac-learning;
            pseudowire-status-tlv;
            static-mac mac-address {
                vlan-id value;
            }
            }
            interface-mac-limit {
                limit;
                packet-action (drop | drop-and-log | log | none | shutdown);
            }
            mac-table-aging-time seconds;
            mac-table-size {
                limit;
                packet-action {
                drop;
                }
            }
            no-mac-learning;
            static-rvtep-mac {
                mac mac_addr {
                remote-vtep;
            }
            }
        }
```

```
    }
}
```

Syntax (vSRX Virtual Firewall)

```
vlans {
    vlan name {
        (vlan-id (all | none | number) | vlan-id-list [ vlan-id-numbers] | vlan-tags <inner number>
outer number);
        description;
        forwarding-options {
            dhcp-security {
                arp-inspection;
                dhcpv6-options {
                    option-16 {
                        use-string use-string;
                    }
                    option-18 {
                        prefix {
                            host-name;
                            logical-system-name;
                            routing-instance-name;
                            vlan-id;
                            vlan-name;
                        }
                        use-interface-description (device | logical);
                        use-interface-index (device | logical);
                        use-interface-mac;
                        use-interface-name (device | logical);
                        use-string use-string;
                    }
                    option-37 {
                        prefix {
                            host-name;
                            logical-system-name;
                            routing-instance-name;
                            vlan-id;
                            vlan-name;
                        }
                        use-interface-description (device | logical);
                        use-interface-index (device | logical);
```

```
                use-interface-mac;
                use-interface-name (device | logical);
                use-string use-string;
            }
        }
        group group-name {
            interface interface-name {
                static-ip {
                    ip-address;
                }
                static-ipv6 {
                    ip-address;
                }
            }
            overrides {
                no-dhcpv6-options;
                no-option16;
                no-option18;
                no-option37;
                no-option82;
                trusted;
                untrusted;
            }
        }
        ip-source-guard;
        ipv6-source-guard;
        light-weight-dhcpv6-relay;
        neighbor-discovery-inspection;
        no-dhcp-snooping;
        no-dhcpv6-snooping;
        option-82 {
            circuit-id {
                prefix {
                    host-name;
                    logical-system-name;
                    routing-instance-name;
                }
                use-interface-description (device | logical);
                use-vlan-id;
            }
            remote-id {
                host-name;
                mac;
```

```
                    use-interface-description (device | logical);
                    use-string use-string;
                }
                vendor-id {
                    use-string use-string;
                }
            }
        }
        filter {
            input filter-name;
        }
        flood {
            input filter-name;
        }
    }
    interface interface-name;
    l3-interface l3-interface-name;
    mcae-mac-synchronize;
    no-irb-layer-2-copy;
    service-id service-id;
    switch-options {
        interface name {
            action-priority action-priority;
            encapsulation-type (ethernet | ethernet-vlan);
        ignore-encapsulation-mismatch;
        interface-mac-limit {
            disable;
            limit;
            packet-action (drop | drop-and-log | log | none | shutdown);
        }
        mac-pinning;
        no-mac-learning;
        pseudowire-status-tlv;
        static-mac mac-address {
            vlan-id value;
        }
        }
        interface-mac-limit {
            limit;
            packet-action (drop | drop-and-log | log | none | shutdown);
        }
        mac-statistics;
        mac-table-aging-time seconds;
```

```
            mac-table-size {
                limit;
                packet-action {
                drop;
                }
            }
            no-mac-learning;
            static-rvtep-mac {
                mac mac_addr {
                remote-vtep;
            }
            }
        }
    }
}
```

Hierarchy Level

```
[edit]
```

```
[edit routing-instances routing-instance-name]
```

Description

Configure VLAN properties.

On EX Series switches and SRX Series Firewalls (including vSRX Virtual Firewall), the following configuration guidelines apply:

- Only private VLAN (PVLAN) firewall filters can be used when the VLAN is enabled for Q-in-Q tunneling.
- An S-VLAN tag is added to the packet if the VLAN is Q-in-Q-tunneled and the packet is arriving from an access interface.
- You cannot use a firewall filter to assign an integrated routing and bridging (IRB) interface or a routed VLAN interface (RVI) to a VLAN.
- VLAN assignments performed using a firewall filter override all other VLAN assignments.

Default

If you use the default factory configuration, all switch interfaces become part of the VLAN **default**.

Options

vlan-name—Name of the VLAN. The name can include letters, numbers, hyphens (-), and periods (.) and can contain up to 255 characters long.

The remaining statements are explained separately. See CLI Explorer.

The remaining statements are described separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing–control—To add this statement to the configuration.

system—To view this statement in the configuration.

system–control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Statements for private VLANs and Q-in-Q tunneling introduced in Junos OS Release 12.1 for the QFX Series switches.

RELATED DOCUMENTATION

Integrated Routing and Bridging

Bridging and VLANs

Bridging and VLANs

Bridging and VLANs

Configuring VLANs for EX Series Switches with ELS Support (CLI Procedure)

Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation

Configuring Q-in-Q Tunneling on Security Devices

Integrated Routing and Bridging

Bridging and VLANs

vlans (RA Guard)

IN THIS SECTION

- Syntax | 30894
- Hierarchy Level | 30894
- Description | 30894
- Required Privilege Level | 30895
- Release Information | 30895

Syntax

```
vlans (vlan-name | all) {
    policy policy-name (stateful | stateless);
}
```

Hierarchy Level

```
[edit forwarding-options access-security router-advertisement-guard]
```

Description

Configure IPv6 Router Advertisement (RA) guard on a VLAN. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in a policy.

Before you can configure RA guard on a VLAN, you must first configure a policy at the [`edit forwarding-options access-security router-advertisement-guard`] hierarchy level. The policy is then applied to the VLAN at the [`edit forwarding-options access-security router-advertisement-guard vlans` *`vlan-name`*] hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D55.

RELATED DOCUMENTATION

> *Understanding IPv6 Router Advertisement Guard*
>
> *Configuring Stateful IPv6 Router Advertisement Guard*

vmhost

IN THIS SECTION

- Syntax | 30895
- Hierarchy Level | 30897
- Description | 30897
- Options | 30897
- Required Privilege Level | 30898
- Release Information | 30898

Syntax

```
vmhost {
    interfaces  name {
        family {
            (inet | inet6) {
                address name {
                    master-only;
```

```
                }
                gateway gateway;
            }
        }
    }
    external interface {
                in-band {
                            mapping {
                        interface  interface-name;
                        interface virtual-function;
                                        }
                        family inet {
                                address IPv4-inet-address;
                                gateway IPv4-gateway;
                                        }
                vlan-id vlan-ID;
                }
            out-of-band {
                          family inet {
                                address IPv4-inet-address;
                                gateway IPv4-gateway;
                                        }
                      }
    management-if {
        disable;
        link-mode (automatic | full-duplex | half-duplex);
        speed (100m | 10m | 1g | automatic);
    }
    no-auto-recovery;
    resize {
        vjunos {
            compact;
        }
    }
    services {
        ssh {
            root-login (allow | deny);
        }
    }
    syslog {
        file name {
            vmhost-syslog-object (any | authorization | cron | daemon | kernel | local0 | local1
| local2 | local3 | local4 | local5 | local6 | local7 | lpr | mail | mark | news | privileged |
```

```
syslog | user | uucp) {
                                        level;
            }
        }
        host name {
            transport (tcp | udp);
            vmhost-syslog-object (any | authorization | cron | daemon | kernel | local0 | local1
| local2 | local3 | local4 | local5 | local6 | local7 | lpr | mail | mark | news | privileged |
syslog | user | uucp) {
                level;
            }
        }
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure VM host management properties. Set values in the `edit vmhost` hierarchy of the configuration.

Options

interfaces Configure interface properties of the host.

- management-if0 | management-if1—Configure the host's side management interface0 or interface1.

  family (inet|inet6) address— Configure IPv4 or IPv6 parameters.

  primary-only— Configure the IP address to be used when the Routing Engine is the current primary. The configured IP address is assigned to primary RE host when the Routing Engine is the current primary. It is recommended to set this option for platforms with dual Routing Engine architecture with VM host support.

  gateway— Configure gateway IP address.

in-band Enable external network access through one of the front panel revenue ports.

out-of-band Enable external network access through the front panel management interface.

management-if Configure management interface properties of the host.

- disable—Disable the host interface eth0, which serves as the management port
- link-mode (automatic | half-duplex | full-duplex)—Configure the link mode of the host interface eth0, which serves as the management port as half-duplex or full-duplex. You can also manually select the link mode option as either half-duplex or full-duplex.
- speed (automatic | 10m | 100m | 1g)—Configure the link speed of the host interface eth0, which serves as the management port. If you set the link speed as 10m or 100m, autonegotiation is turned off and the link speed is the speed that you specify.

no-auto-recovery Disable the automatic guest recovery by the host.

services Enable ssh access to the host and enable or disable root-login to the host from guest.

- ssh—Allow ssh access

  root-login—Configure host root access through ssh

  allow | deny—Allow or deny root access through ssh.

syslog Enable the remote syslog configuration from guest to host OS. Based on the severity configured on guest, the syslog information is logged onto the **/etc/syslog.conf** file on the host.

- host <host-name>—Host notified for remote syslog configuration.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F3.

in-band and out-of-band options introduced for NFX devices in Junos OS Release 24.2R1.

NOTE: PTX3000 router supports the Routing and Control Board, RCBPTX.

RELATED DOCUMENTATION

VM Host Overview

Disabling Autorecovery on Routing Engines with VM Host Support

vmhost management-if disable

IN THIS SECTION

- Syntax | 30899
- Hierarchy Level | 30899
- Description | 30899
- Default | 30900
- Required Privilege Level | 30900
- Release Information | 30900

Syntax

```
vmhost management-if disable
```

Hierarchy Level

```
[edit vmhost]
```

Description

Disable the host interface eth0, which serves as the management port. You can the disable the interface if there are any issues associated with security or any hardware failures either at the local end or the remote end of the interface. if you disable the interface, the transmitter is turned off and the link partner experiences a link-down condition.

Default

The host interface eth0 which serves as the management port is enabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F6.

> **NOTE**: The command is supported on the routers with RE-MX-X6, RE-MX-X8, and RE-PTX-X8 Routing Engines only.

RELATED DOCUMENTATION

- vmhost | 30895
- vmhost management-if speed | 30902
- vmhost management-if link-mode | 30900
- show vmhost management-if | 27797

vmhost management-if link-mode

IN THIS SECTION

- Syntax | 30901
- Hierarchy Level | 30901
- Description | 30901
- Default | 30901
- Options | 30901
- Required Privilege Level | 30901

Release Information | 30902

Syntax

```
vmhost management-if link-mode (automatic | half-duplex | full-duplex)
```

Hierarchy Level

```
[edit vmhost]
```

Description

Configure the link mode of the host interface eth0, which serves as the management port as half-duplex or full-duplex. You can also manually select the link mode option as either half-duplex or full-duplex.

Default

The link partners auto-negotiate the speed and duplex link mode and select the highest common capability.

Options

automatic Autonegotiate the link mode of the management interface. if you set the link mode to `automatic`, you must also set the link speed to `automatic`.

half-duplex Set the link mode of the management interface to half-duplex.

full-duplex Set the link mode of the management interface to full-duplex.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F6.

> **NOTE**: The command is supported on the routers with RE-MX-X6, RE-MX-X8, and RE-PTX-X8 Routing Engines only.

RELATED DOCUMENTATION

- vmhost | 30895
- vmhost management-if disable | 30899
- vmhost management-if speed | 30902
- show vmhost management-if | 27797

vmhost management-if speed

IN THIS SECTION

- Syntax | 30902
- Hierarchy Level | 30903
- Description | 30903
- Default | 30903
- Options | 30903
- Required Privilege Level | 30903
- Release Information | 30903

Syntax

```
vmhost management-if speed { automatic | 10m | 100m | 1g }
```

Hierarchy Level

```
[edit vmhost]
```

Description

Configure the link speed of the host interface eth0, which serves as the management port. If you set the link speed as `10m` or `100m`, autonegotiation is turned off and the link speed is the speed that you specify.

Default

The link partners auto-negotiate the speed and duplex link mode and select the highest common capability.

Options

automatic Autonegotiate the link speed of the management interface. If you set the link speed as `automatic`, speed and link mode are auto-negotiated with the link partner.

10m Set the link speed of the management interface to 10Mbps.

100m Set the link speed of the management interface to 100Mbps

1g Set the link speed of the management interface to 1Gbps. If you set link speed to `1Gbps`, autonegotiation is enabled. However, the interface advertises only 1Gbps speed and full-duplex mode.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F6.

NOTE: The command is supported on the routers with RE-MX-X6, RE-MX-X8, and RE-PTX-X8 Routing Engines only.

RELATED DOCUMENTATION

vmhost | 30895
vmhost management-if disable | 30899
vmhost management-if link-mode | 30900
show vmhost management-if | 27797

vmhost resize vjunos compact

IN THIS SECTION

- Syntax | 30904
- Hierarchy Level | 30904
- Description | 30905
- Required Privilege Level | 30905
- Release Information | 30905

Syntax

```
vmhost resize vjunos compact
```

Hierarchy Level

```
[edit]
```

Description

Reduce the routing engine resource footprint on MX Series router to accommodate GNF VMs within the Routing Engine in in-chassis Junos node slicing. Configuring the `vmhost resize vjunos compact` statement enables hyperthreading in Routing Engine CPU. The in-chassis Junos node slicing requires that hyperthreading be enabled.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1R1.

vnf-license-supplement

IN THIS SECTION

- Syntax | 30905
- Hierarchy Level | 30906
- Description | 30906
- Options | 30906
- Required Privilege Level | 30906
- Release Information | 30906

Syntax

```
vnf-license-supplement {
    vnf-id gnf-id
    license-supplement-string augmented-ssrn
}
```

Hierarchy Level

[edit system]

Description

Allocate a MAC address to a licensed guest network function (GNF). For more details, see *Assigning MAC Addresses to GNF*.

Options

vnf-id ***gnf-id***	Specify the GNF ID that will be assigned the licensed MAC address range. Each license can only be assigned to a single GNF.
license-supplement-string ***augmented-ssrn***	Specify the augmented software support reference number (SSRN) string, which contains the MAC address details. Example: `RTU00023003204-01-AABBCCDDEEFF-1100-01-ABCD`.

Required Privilege Level

Release Information

Statement introduced in Junos OS Release 19.2R1.

vnf-name

IN THIS SECTION

- Syntax | 30907
- Hierarchy Level | 30908
- Description | 30908
- Options | 30909
- Required Privilege Level | 30909
- Release Information | 30909

Syntax

```
vnf-name {
    type {
        linux-container | virtual-machine;
    }
    image {
        file-path;
        bus-type [ide | virtio];
        image-type [qcow2 | raw];
    }
    init-descriptor file-path;
    memory {
        size size;
        features {
            hugepages;
        }
    }
    no-autostart;
    storage device-name {
        type {
            cdrom {
                source {
                    file filename;
                }
            }
            disk {
                bus-type [ide | virtio];
                    file-type [qcow2 | raw];
                    source {
                    file filename;
                }
            }
            usb {
                source {
                    file filename;
                }
            }
        }
    }
    virtual-cpu {
        virtual-cpu-number {
```

```
            physical-cpu number | range;
        }
        count number;
        features {
            hardware-virtualization;
        }
    }
    interfaces interface-name {
        pci-address pci-address;
        mapping {
            hsxe0 {
                virtual-function {
                    vlan-id vlan-id;
                }
            }
            hsxe1 {
                virtual-function {
                    vlan-id vlan-id;
                }
            }
            vlan {
                members vlan-name;
                mode [access | trunk];
                native-vlan-id vlan-id;
            }
        }
    }
}
```

Hierarchy Level

```
[edit virtual-network-functions]
```

Description

Name of the virtual network function.

Options

interfaces Name of the interface. For example, em1.

no-autostart An option to disable auto-start of VNF on the VNF configuration commit.

pci-address An option to specify PCI address for the VNF interfaces.

mapping An option to map VNF interfaces on platforms running disaggregated Junos OS.

virtual-cpu An option to specify the number of virtual CPUs the VNF can use. By default, a VNF is assigned one virtual CPU, which is independent of any specific physical CPU.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X53-D45.

RELATED DOCUMENTATION

Understanding Virtual Network Functions

Managing the VNF Life Cycle

show virtual-network-functions | 27723

vni

IN THIS SECTION

- Syntax | 30910
- Hierarchy Level | 30910
- Description | 30910

- Options | **30910**
- Required Privilege Level | **30910**
- Release Information | **30911**

Syntax

```
vni [0-16777214]
```

Hierarchy Level

```
[edit routing-instances routing-instance-name vlans vlan-name vxlan]
[edit vlans vlan-name vxlan]
```

Description

Assign a numeric value to identify a Virtual Extensible LAN (VXLAN). All members of a VXLAN must use the same VNI.

Options

vni Value to specify in the `vni` attribute.

- **Range:** 0 through 16,777,215

> **NOTE**: Starting in Junos OS Release 17.3R3-3, 18.1R3-S3, and 19.1R1, Junos OS supports a VNI value of 0.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R2.

RELATED DOCUMENTATION

Understanding VXLANs

Manually Configuring VXLANs on QFX Series and EX4600 Switches

Examples: Manually Configuring VXLANs on QFX Series and EX4600 Switches

vni(Interfaces)

IN THIS SECTION

- Syntax | 30911
- Hierarchy Level | 30911
- Description | 30912
- Options | 30912
- Required Privilege Level | 30912
- Release Information | 30912

Syntax

```
vni [0-16777214]
```

Hierarchy Level

```
[edit interfaces name unit name tunnel encapsulation vxlan-gpe]
```

Description

Assign a numeric value to identify a Virtual Extensible LAN (VXLAN). All members of a VXLAN must use the same virtual network identifier (VNI).

Options

vni Value to specify in the `vni` attribute.

- **Range:** For flexible tunnel interfaces, the range supported is 0 through 16,777,214 for all MX Series routers (except MX80).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding VXLANs

Flexible Tunnel Interfaces Overview

Configuring Flexible Tunnel Interfaces

Example: Configuring Flexible Tunnel Interfaces on MX Series Routers

show interfaces fti | **20866**

vni-options

IN THIS SECTION

Syntax | **30913**

- Hierarchy Level | 30913
- Description | 30913
- Required Privilege Level | 30913
- Release Information | 30914

Syntax

```
vni-options vni vxlan-network-identifier {
    designated-forwarder-election-hold-time seconds;
    vrf-target {
        community;
        auto;
        import community-name;
        export community-name;
    }
}
```

Hierarchy Level

```
[edit protocols evpn]
[edit routing-instances routing-instance-name protocols evpn]
```

Description

Configure a designated forwarder election hold time and specific route targets (RTs) for each VXLAN network identifier (VNI).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D30 .

RELATED DOCUMENTATION

EVPN Multihoming Overview

Understanding EVPN with VXLAN Data Plane Encapsulation

Example: Configuring an EVPN Control Plane and VXLAN Data Plane

extended-vni-list | 6710

vnid (EVPN)

IN THIS SECTION

- Syntax | 30914
- Hierarchy Level | 30914
- Description | 30915
- Required Privilege Level | 30915
- Release Information | 30915

Syntax

```
vnid;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name]
```

Description

Enables having the same MAC frames across multiple VXLAN segments without traffic crossover. Multicast in VXLAN is implemented as Layer 3 multicast, in which endpoints subscribe to groups. VNID is a 24-bit virtual network identifier that uniquely identifies the VXLAN segment.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

Understanding EVPN with VXLAN Data Plane Encapsulation

Example: Configuring an EVPN Control Plane and VXLAN Data Plane

EVPN Type-5 Route with MPLS encapsulation for EVPN-MPLS

EVPN Type-5 Route with VXLAN encapsulation for EVPN-VXLAN

voice

IN THIS SECTION

- Syntax | 30916
- Hierarchy Level | 30916
- Description | 30916
- Required Privilege Level | 30916
- Release Information | 30916

Syntax

```
voice {
    dscp (Services CoS) (alias | bits);
    forwarding-class (Services PIC Classifiers) class-name;
}
```

Hierarchy Level

```
[edit services (CoS) cos application-profile profile-name sip]
```

Description

Set the appropriate `dscp` and `forwarding-class` values for SIP voice traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring Application Profiles for Use as CoS Rule Actions

voice (Application Profile)

IN THIS SECTION

- Syntax | 30917
- Hierarchy Level | 30917
- Description | 30917
- Default | 30917
- Required Privilege Level | 30917
- Release Information | 30918

Syntax

```
voice {
    dscp (alias | bits);
    forwarding-class class-name;
}
```

Hierarchy Level

```
[edit services cos application-profileprofile-name sip]
```

Description

Set the appropriate `dscp` and `forwarding-class` values for SIP voice traffic.

Default

By default, the system will not alter the DSCP or forwarding class for SIP voice traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

video (Application Profile) | 30728

voip

IN THIS SECTION

- Syntax | 30918
- Hierarchy Level | 30919
- Description | 30919
- Required Privilege Level | 30919
- Release Information | 30919

Syntax

```
voip {
    interface (all | [interface-name | access-ports]) {
        forwarding-class forwarding-class;
        vlan  vlan-name );
        voip-mac-exclusive;
    }
}
```

Hierarchy Level

- For platforms with ELS:

```
[edit switch-options voip]
```

- For platforms without ELS:

```
[edit ethernet-switching-options                voip],
```

Description

Configure voice over IP (VoIP) on interfaces.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Hierarchy level `[edit switch-options]` introduced in Junos OS Release 13.2X50-D10. (See *Using the Enhanced Layer 2 Software CLI* for information about ELS.)

Configuration option `voip-mac-exclusive` introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

> *Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch*
>
> *Example: Configuring VoIP on an EX Series Switch Without Including 802.1X Authentication*
>
> *Example: Configuring VoIP on an EX Series Switch Without Including LLDP-MED Support*

voip-mac-exclusive

IN THIS SECTION

- Syntax | 30920
- Hierarchy Level | 30920
- Description | 30920
- Default | 30921
- Required Privilege Level | 30921
- Release Information | 30921

Syntax

```
voip-mac-exclusive;
```

Hierarchy Level

```
[edit ethernet-switching-options secure-access-port interface (all | interface-name)]
[edit switch-options voip interface (all | interface-name)]
```

Description

NOTE: `[edit switch-options voip interface (all | interface-name)]` is not supported on ELS on switches that use Enhanced Layer 2 Software (ELS).

Restrict a VoIP client MAC address to be learned only in a configured VoIP VLAN.

If the `voip-exclusive-mac` statement is configured at the [`edit ethernet-switching-options secure-access-port interface interface-name`] hierarchy level for an interface in a VoIP VLAN, any MAC address learned on that interface for the VoIP VLAN is not learned on an interface for a data VLAN. If a MAC address has been learned on a data VLAN interface and then later, is learned on a VoIP VLAN with that same interface, the MAC address is removed from the data VLAN interface.

Default

A client MAC address is unrestricted and can be learned on both a VoIP VLAN and a data VLAN.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in JUNOS Release 13.2X50-D10.

RELATED DOCUMENTATION

Overview of Port Security

secure-access-port | **16432**

volume-limit (LRF Rule)

IN THIS SECTION

- Syntax | **30921**
- Hierarchy Level | **30922**
- Description | **30922**
- Options | **30922**
- Required Privilege Level | **30922**
- Release Information | **30922**

Syntax

```
volume-limit volume;
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name then report]
```

Description

Configure the data volume limit to be used for reporting. The template that the LRF rule is using must have `trigger-type volume` configured.

Options

volume Data volume, in megabytes.

- **Range:** 1 through 1024

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

RELATED DOCUMENTATION

Configuring Logging and Reporting for Junos OS Subscriber Aware

Configuring an LRF Profile for Subscribers

Configuring Logging and Reporting for Subscriber Management

vpi (Define Virtual Path)

IN THIS SECTION

- Syntax | 30923
- Hierarchy Level | 30923
- Description | 30923
- Options | 30924
- Required Privilege Level | 30924
- Release Information | 30924

Syntax

```
vpi vpi-identifier {
    maximum-vcs maximum-vcs;
    oam-liveness {
        up-count cells;
        down-count cells;
    }
    oam-period (disable | seconds);
    shaping {
        (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
burst length);
        queue-length number;
    }
}
```

Hierarchy Level

```
[edit interfaces at-fpc/pic/port atm-options]
```

Description

For ATM interfaces, configure the virtual path (VP).

NOTE: Certain options apply only to specific platforms.

Options

vpi-identifier—ATM virtual path identifier. This is one of the VPIs that you define in the "vci" on page 30660 statement. (For a list of hierarchy levels at which you can include the `vci` statement, see "vci" on page 30660.)

- **Range:** 0 through 255

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring the Maximum Number of ATM1 VCs on a VP

vpls-flow-table-size

IN THIS SECTION

- Syntax | 30925
- Hierarchy Level | 30925
- Description | 30925
- Options | 30925

Required Privilege Level | 30926
Release Information | 30926

Syntax

```
vpls-flow-table-size units;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services flow-table-size]
```

Description

Configure the size of the VPLS flow table in units of 256K entries.

NOTE: Any change in the configured size of the flow table size initiates an automatic reboot of the FPC.

NOTE: Starting with Junos OS Release 17.3R1, the maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 256 on MPC5Es and MPC6Es with with 4 GB DDR3 memory or higher. The maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 245 on MPC5Es and MPC6Es with DDR3 memory lower than 4 GB.

NOTE: The recommended flow table size is 4 so that it can scale up to 4x256K flows, which is 1M. You can configure more, however, the system will issue a warning message.

Options

units Number of 256K flow entries available for the VPLS flow table.

- **Range:** 1 through 245
- **Default:** 15 (3840K)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

vpls-flush-on-topology-change

IN THIS SECTION

- Syntax | 30926
- Hierarchy Level | 30927
- Description | 30927
- Required Privilege Level | 30927
- Release Information | 30927

Syntax

```
vpls-flush-on-topology-change;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (mstp | rstp)],
[edit logical-systems logical-system-name protocols vstp vlan vlan-id],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
(mstp | rstp)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vstp
vlan vlan-id],
[edit protocols (mstp | rstp)],
[edit protocols vstp vlan vlan-id],
[edit routing-instances routing-instance-name protocols (mstp | rstp)],
[edit routing-instances routing-instance-name protocols vstp vlan vlan-id]
```

Description

Determine the action the bridge should take when the topology of a multihomed Layer 2 ring with MPLS infrastructure changes: flush the media access control (MAC) cache or not. By default, the bridge does not flush the cache when the topology changes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

vpls-id

IN THIS SECTION

- Syntax | 30928
- Hierarchy Level | 30928
- Description | 30928

- Options | 30928
- Required Privilege Level | 30928
- Release Information | 30929

Syntax

```
vpls-id vpls-id;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances instance-name protocols l2vpn],
[edit logical-systems logical-system-name routing-instances instance-name protocols l2vpn mesh-
group mesh-group-name],
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls],
[edit logical-systems logical-system-name routing-instances instance-name protocols vpls mesh-
group mesh-group-name],
[edit routing-instances instance-name protocols l2vpn],
[edit routing-instances instance-name protocols l2vpn mesh-group mesh-group-name],
[edit routing-instances instance-name protocols vpls],
[edit routing-instances instance-name protocols vpls mesh-group mesh-group-name]
```

Description

Identify the virtual circuit identifier used for the VPLS routing instance or mesh group. This statement is a part of the configuration to enable LDP signaling for VPLS.

Options

vpls-id—Specify a valid identifier for the VPLS routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring LDP Signaling for VPLS

vpls-id-list (protocols vpls mesh-group)

IN THIS SECTION

- Syntax | 30929
- Hierarchy Level | 30929
- Description | 30930
- Options | 30930
- Required Privilege Level | 30930
- Release Information | 30930

Syntax

```
vpls-id-list vc-id-numbers;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name neighbor
address]
```

Description

The `vpls-id-list` option allows you to configure multiple pseudowires under LDP-VPLS neighbor by creating one pseudowire per virtual circuit id. This combination of neighbor and virtual circuit ID must be unique across the l2circuit and all local LDP-VPLS instances. The pseudowire will terminate in the specified mesh-group.

As with other pseudowires, the pseudowires created here will signal the remote PE via LDP label mapping message. For each pseudowire created under a neighbor, VPLS will create a VT/LSI interface and add both it and the and label route to the mpls.0 table. Note that you can create multiple pseudowires between the same pair of PEs in LDP-VPLS for a single routing instances only.

Enable `local-switching` along with `vpls-id-list` to allow local switching of traffic (including BUM traffic) between multiple pseudo wires. This can also eliminate the need to have a dedicated mesh-group for each LDP spoke pseudo wire (which has a limit of 14), and provide support for H-VPLS in EVP-LAN and EP-LAN topologies.

Options

`vc-id-numbers`—The virtual circuit ID for the PE pair. This ID can be any number in the range of 1 to 4294967295.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2 .

RELATED DOCUMENTATION

show vpls connections | **27849**

neighbor (Protocols VPLS) | **11284**

vpls (Interfaces)

IN THIS SECTION

- Syntax | 30931
- Hierarchy Level | 30931
- Description | 30931
- Required Privilege Level | 30931
- Release Information | 30931

Syntax

```
vpls;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family]
```

Description

Specify the VPLS protocol family information for the logical interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Interfaces for VPLS Routing

vpls-mac-move

IN THIS SECTION

- Syntax | 30932
- Hierarchy Level | 30932
- Description | 30932
- Options | 30933
- Required Privilege Level | 30933
- Release Information | 30933

Syntax

```
vpls-mac-move{
    cooloff-time seconds;
    interface-recovery-time seconds;
    statistical-approach-wait-time seconds;
    virtual-mac mac-address /mask;
}
```

Hierarchy Level

```
[edit protocols l2-learning]
```

Description

Configure VPLS MAC move parameters to prevent loop creation in the network.

Options

cooloff-time *seconds*	Time interval, in seconds, during which no additional interface is disabled by the PE after a MAC move is detected. By default, it is 30 seconds.
interface-recovery-time *seconds*	Time interval, in seconds, after which the interface is enabled after a VPLS MAC move is detected. By default, the timer is disabled.
statistical-approach-wait-time *seconds*	Time interval, in seconds, during which the MAC moves are monitored to collect the statistics. By default, it is 30 seconds.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

Statement changed to `global-mac-move` in Junos OS Release 17.4R1

RELATED DOCUMENTATION

MAC Moves Loop Prevention in VPLS Network Overview

Example: Configuring Loop Prevention in VPLS Network Due to MAC Moves

virtual-mac | **30766**

vpls (Routing Instance)

IN THIS SECTION

- Syntax | **30934**
- Hierarchy Level | **30936**

- Description | 30936
- Options | 30936
- Required Privilege Level | 30936
- Release Information | 30936

Syntax

```
vpls {
    mac-pinning;
    active-interface {
        any;
        primary interface-name;
    }
    community COMM;
    connectivity-type (ce | irb);
    control-word;
    encapsulation-type ethernet;
    ignore-encapsulation-mismatch;
    ignore-mtu-mismatch;
    import-labeled-routes [ routing-instance-name ];
    interface interface-name;
    interface-mac-limit limit;
    label-block-size size;
    mac-flush [ explicit-mac-flush-message-options ];
    mac-table-aging-time time;
    mac-table-size size;
    mesh-group mesh-group-name {
        interface interface-name;
        l2vpn-id (as-number:id | ip-address:id);
        local-switching;
        mac-flush [ explicit-mac-flush-message-options ];
        neighbor address {...};
        peer-as all;
        pseudowire-status-tlv hot-standby-vc-on;
        route-distinguisher (as-number:id | ip-address:id);
        vpls-id number;
        vrf-export [ policy-names ];
        vrf-import [ policy-names ];
```

```
        vrf-target {
            community;
            import community-name;
            export community-name;
        }
    }
    mtu mtu;
    no-control-word;
    no-tunnel-services;
    service-type single;
    site site-name {
        active-interface interface-name {
            any;
            primary preference-value;
        }
        best-site;
        interface interface-name {
            interface-mac-limit limit;
        }
        mesh-group mesh-group-name;
        multi-homing;
        site-identifier identifier;
        site-preference preference-value {
            backup;
            primary;
        }
    }
    site-range number;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    tunnel-services {
        devices device-names;
        primary primary-device-name;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
[edit routing-instances routing-instance-name protocols]
```

Description

Configure a virtual private LAN service (VPLS) routing instance.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

`service-type single`—Allows one vlan per routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`mac-flush` option introduced in Junos OS Release 10.0.

`hot-standby-vc-on`, `import-labeled-routes [ routing-instance-name ]`, and `interface interface` options introduced in Junos OS Release 15.1R2.

RELATED DOCUMENTATION

Configuring VPLS Routing Instances

vpls-template

IN THIS SECTION

- Syntax | 30937
- Hierarchy Level | 30937
- Description | 30937
- Required Privilege Level | 30937
- Release Information | 30937

Syntax

```
vpls-template;
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix template]
```

Description

Specify that the IPFIX template is used only for VPLS records. Starting in Junos OS Release 18.2R1, the `vpls-template` option is deprecated; use the `bridge-template` option instead. The `bridge-template` option supports both VPLS and bridge records.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13 .2.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

bridge-template | 3038

vpn-apply-export

IN THIS SECTION

- Syntax | 30938
- Hierarchy Level | 30938
- Description | 30938
- Required Privilege Level | 30939
- Release Information | 30939

Syntax

```
vpn-apply-export;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor neighbor],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor neighbor]
```

Description

Apply both the VRF export and BGP group or neighbor export policies (VRF first, then BGP) before routes from the `vrf` or `l2vpn` routing tables are advertised to other PE routers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Policies for the VRF Table on PE Routers in VPNs

vpn-group-address

IN THIS SECTION

- Syntax | 30939
- Hierarchy Level | 30940
- Description | 30940
- Options | 30940
- Required Privilege Level | 30940
- Release Information | 30940

Syntax

Use `group-address` in place of `vpn-group-address`.

```
vpn-group-address address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim],
[edit routing-instances routing-instance-name protocols pim]
```

Description

Configure the group address for the Layer 3 VPN in the service provider's network.

Options

address—Address for the Layer 3 VPN in the service provider's network.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Starting with Junos OS Release 11.4, to provide consistency with draft-rosen 7 and next-generation BGP-based multicast VPNs, configure the provider tunnels for draft-rosen 6 anysource multicast VPNs at the `[edit routing-instances routing-instance-name provider-tunnel]` hierarchy level. The `mdt`, `vpn-tunnel-source`, and `vpn-group-address` statements are deprecated at the `[edit routing-instances routing-instance-name protocols pim]` hierarchy level.

RELATED DOCUMENTATION

> *Configuring Multicast Layer 3 VPNs*
>
> Junos OS Multicast Protocols User Guide

vpn-label

IN THIS SECTION

- Syntax | 30941
- Hierarchy Level | 30941
- Description | 30941
- Required Privilege Level | 30941
- Release Information | 30941

Syntax

```
vpn-label;
```

Hierarchy Level

```
[edit chassis memory-enhanced]
```

Description

Allocate more jtree memory for Layer 3 VPN labels.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 10.4.

RELATED DOCUMENTATION

Allocating More Memory for Routing Tables, Firewall Filters, and Layer 3 VPN Labels

vpn-localization

IN THIS SECTION

- Syntax | 30942
- Hierarchy Level | 30942
- Description | 30942
- Options | 30943
- Required Privilege Level | 30943
- Release Information | 30943

Syntax

```
vpn-localization {
vpn-core-facing-default;
vpn-core-facing-only;
}
```

Hierarchy Level

```
[edit chassis fpc fpc-slot-number]
```

Description

Configure FPC as either core-facing default or core-facing only in order to localize the routes.

Options

vpn-core-facing-default The core-facing FPC installs all the routes and next hops of the CE-facing interfaces.

vpn-core-facing-only The core-facing FPC installs all the routes and does not store next hops of the CE-facing interfaces.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Understanding VRF Localization in Layer 3 VPNs

Example: Improving Scalability Using VRF Localization for Layer 3 VPNs

localized-fib | 9983

vpn-monitor

IN THIS SECTION

- Syntax | 30944
- Hierarchy Level | 30944
- Description | 30944
- Options | 30944
- Required Privilege Level | 30945
- Release Information | 30945

Syntax

```
vpn-monitor {
    destination-ip ip-address;
    optimized;
    source-interface interface-name;
    verify-path {
        destination-ip ip-address;
        packet-size bytes;
    }
}
```

Hierarchy Level

```
[edit security ipsec vpn vpn-name]
```

Description

Configure settings for VPN monitoring.

Options

destination-ip Specify the destination of the Internet Control Message Protocol (ICMP) pings. If this statement is used, the device uses the peer's gateway address by default.

optimized Enables VPN monitoring optimization for the specific VPN tunnel.

When the administrator enables VPN monitoring optimization, the SRX Series Firewall sends ICMP echo requests (pings) only when there is an outgoing traffic and no incoming traffic from the configured peer, through the VPN tunnel. If there is an incoming traffic through the VPN tunnel, the SRX Series Firewall considers the tunnel to be active and does not send pings to the peer.

VPN monitoring optimization saves resources on the SRX Series Firewall as ICMP echo requests are sent only when needed to determine peer liveliness. Also, ICMP echo requests can activate costly backup links that would otherwise not be used.

- Default: Disabled.

If the administrator does not configure the option explicitly, the firewall sends VPN monitoring packet once every configured interval, also knows as always-send, by default. For more details, see Understanding VPN Monitoring.

source-interface Specify the source interface for ICMP requests (VPN monitoring "hellos"). If no source interface is specified, the device automatically uses the local tunnel endpoint interface.

verification-path Specify the verification path to verify the IPsec datapath before the secure tunnel (st0) interface is activated and route(s) associated with the interface are installed in the Junos OS forwarding table.

- `destination-ip` *`ip-address`*—Original, untranslated IP address of the peer tunnel endpoint that is behind a NAT device. This IP address must not be the NAT translated IP address. This option is required if the peer tunnel endpoint is behind a NAT device. The verify-path ICMP request is sent to this IP address so that the peer can generate an ICMP response.

- `packet-size` *`bytes`*—(Optional) The size of the packet that is used to verify an IPsec datapath before the st0 interface is brought up. The packet size must be lower than the path maximum transmission unit (PMTU) minus tunnel overhead. The packet used for IPsec datapath verification must not be fragmented. The range of the packet size is 64 to 1350 bytes and the default packet size value is 64 bytes

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Support for `verify-path` keyword and `destination-ip` added in Junos OS Release 15.1X49-D70.

Support for `packet-size` option added in Junos OS Release 15.1X49-D120.

Support for `vpn-monitor` and `verify-path` options with IPsec VPN running iked process is introduced in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

IPsec Overview

vpn (Security)

IN THIS SECTION

- Syntax | 30946
- Hierarchy Level | 30948
- Description | 30948
- Options | 30948
- Required Privilege Level | 30951
- Release Information | 30951

Syntax

```
vpn vpn-name {
    bind-interface interface-name;
    df-bit (clear | copy | set);
    distribution-profile (default-spc2-profile | default-spc3-profile | distribution-profile-name);
    copy-outer-dscp;
    establish-tunnels (immediately | on-traffic | responder-only | responder-only-no-rekey);
    match-direction (input | output);
    passive-mode-tunneling;
    tunnel-mtu tunnel-mtu;
    ike {
        anti-replay-window-size anti-replay-window-size;
        gateway gateway-name;
        idle-time seconds;
        install-interval seconds;
        ipsec-policy ipsec-policy-name;
        no-anti-replay;
        proxy-identity {
            local ip-prefix;
            remote ip-prefix;
            service (any | service-name);
```

```
        }
    }
    manual {
        authentication {
            algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
            key (ascii-text key | hexadecimal key);
        }
        encryption {
            algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-256-cbc |
aes-256-gcm | des-cbc);
            key (ascii-text key | hexadecimal key);
        }
        external-interface external-interface-name;
        gateway ip-address;
        protocol (ah | esp);
        spi spi-value;
    }
        multi-sa {
            forwarding-class (expedited-forwarding | assured-forwarding | best-effort | network-
control);
        }
    traffic-selector traffic-selector-name {
        local-ip ip-address/netmask;
        remote-ip ip-address/netmask;

        protocol protocol_name/protocol_id;
        source-port low-high;
        destination-port low-high;
        metric metric_value;
        description description_value;
        term term_name {
        local-ip ip-address/netmask;
        remote-ip ip-address/netmask;
        protocol protocol_name/protocol_id;
        source-port low-high;
        destination-port low-high;
                                    }
    }
    vpn-monitor {
        destination-ip ip-address;
        optimized;
        source-interface interface-name;
        verify-path {
```

```
            destination-ip ip-address;
            packet-size bytes;
        }
    }
}
```

Hierarchy Level

```
[edit security ipsec]
```

Description

Configure an IPsec VPN. A VPN provides a means by which remote computers communicate securely across a public WAN suchas the Internet. A VPN connection can link two LANs (site-to-site VPN) or a remote dial-up user and a LAN. The traffictha t flows between these two points passes through shared resources such as routers, switches, and othernetwork equipment that make up the public WAN. To secure VPN communication while passing throughthe WAN, the two participants create an IP Security (IPsec) tunnel. IPsec is a suite of related protocols for cryptographically securing communications at the IP Packet Layer.

Options

vpn-name Name of the VPN.

bind-interface Configure the tunnel interface to which the route-based virtual private network (VPN) is bound.

copy-outer-dscp Enable copying of Differentiated Services Code Point (DSCP) (outer DSCP+ECN) field from the outer IP header encrypted packet to the inner IP header plain text message on the decryption path. The benefit in enabling this feature is that after IPsec decryption, clear text packets can follow the inner CoS (DSCP+ECN) rules.

distribution-profile Specify a distribution-profile to distribute tunnels. The `distribution-profile` option is introduced to give the administrator an option to select which PICs in the chassis should handle tunnels associated with a certain VPN object. If the default profiles such as `default-spc3-profile` or `default-spc2-profile` are not selected, a new user-defined profile can be selected. In a profile, you need to mention the Flexible PIC Concentrator (FPC) slot and the PIC number. When such a profile is associated with a VPN object, all matching tunnels are distributed across these PIC's.

- Values:
  - default-spc2-profile—Default group for distributing tunnels on SPC2 only
  - default-spc3-profile—Default group for distributing tunnels on SPC3 only
  - distribution-profile-name—Name of the distribution profile.

df-bit

Specify how the device handles the Don't Fragment (DF) bit in the outer header.

On SRX5400, SRX5600, and SRX5800 devices, the DF-bit configuration for VPN only works if the original packet size is smaller than the st0 interface MTU, and larger than the external interface-ipsec overhead.

- Values:
  - `clear`—Clear (disable) the DF bit from the outer header. This is the default.
  - `copy`—Copy the DF bit to the outer header.
  - `set`—Set (enable) the DF bit in the outer header.

establish-tunnels

Specify when IKE is activated: immediately after VPN information is configured and configuration changes are committed, or only when data traffic flows. If this configuration is not specified, IKE is activated only when data traffic flows.

- Values:
  - `immediately`—IKE is activated immediately after VPN configuration changes are committed.

    Starting with Junos OS Release 15.1X49-D70, a warning message is displayed if you configure the `establish-tunnels immediately` option for an IKE gateway with `group-ike-id` or `shared-ike-id` IKE user types (for example, with AutoVPN or a remote access VPN). The `establish-tunnels immediately` option is not appropriate for these VPNs because multiple VPN tunnels may be associated with a single VPN configuration. Committing the configuration will succeed, however the `establish-tunnels immediately` configuration is ignored. The state of the tunnel interface will be up all the time, which was not the case in previous releases when the `establish-tunnels immediately` option was configured.

  - `on-traffic`—IKE is activated only when data traffic flows and must to be negotiated with the peer gateway. This is the default behavior.

- `responder-only`—Responds to IKE negotiations that are initiated by the peer gateway, but does not initiate IKE negotiations from the device. This option is required when another vendor's peer gateway expects the protocol and port values in the traffic selector from the initiating gateway. `responder-only` option added in Junos OS Release 19.1R1.

  This option is supported on unified iked process that is not enabled by default. Administrators must execute the `request system software add optional://junos-ike.tgz` command to load the `junos-ike` package.

- `responder-only-no-rekey`—Option does not establish any VPN tunnel from the device, so the VPN tunnel is initiated from the remote peer. An established tunnel does not start any rekeying from the device and relies on the remote peer to initiate this rekeying. If rekeying does not occur, then the tunnel is brought down after hard-lifetime expires.

  This option is supported on unified iked process that is not enabled by default. Administrators must execute the `request system software add optional://junos-ike.tgz` command to load the `junos-ike` package.

ike Define an IKE-keyed IPsec VPN.

manual Define a manual IPsec security association (SA).

multi-sa Negotiate multiple security association (SAs) based on configuration choice. Multiple SAs negotiates with the same traffic selector on the same IKE SA.

traffic-selector Configure multiple sets of local IP address prefix, remote IP address prefix, source port range, destination port range, and protocol as a traffic selector for an IPsec tunnel.

match-direction Direction for which the rule match is applied

- Values:
  - input—Match on input to interface
  - output—Match on output from interface

passive-mode-tunneling No active IP packet checks before IPSec encapsulation

tunnel-mtu Maximum transmit packet size

- **Range:** 256 through 9192

vpn-monitor Configure settings for VPN monitoring.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Support for IPv6 addresses added in Junos OS Release 11.1.

Support for `copy-outer-dscp` added in Junos OS Release 15.1X49-D30.

`verify-path` keyword and `destination-ip` added in Junos OS Release 15.1X49-D70.

`packet-size` option added in Junos OS Release 15.1X49-D120.

Support for `term`, `protocol`, `source-port`, `destination-port`, `metric`, and `description` options introduced in Junos OS Release 21.1R1.

Support for `vpn-monitor` option with IPsec VPN running iked process is added in Junos OS Release 23.4R1.

RELATED DOCUMENTATION

IPsec Overview

vpn-unequal-cost

IN THIS SECTION

- Syntax | 30952
- Hierarchy Level | 30952
- Description | 30952

- Options | 30952
- Required Privilege Level | 30952
- Release Information | 30953

Syntax

```
vpn-unequal-cost {
    equal-external-internal;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options multipath],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-
options rib routing-table-name multipath],
[edit routing-instances routing-instance-name routing-options multipath],
[edit routing-instances routing-instance-name routing-options rib routing-table-name multipath]
```

Description

When the `vpn-unequal-cost` statement is configured, protocol-independent load balancing is applied to VPN routes that are equal to interior gateway protocol (IGP) metrics. If the `vpn-unequal-cost` statement is not configured, then protocol-independent load balancing is applied to VPN routes that are equal to the router identifier.

Options

`equal-external-internal`—Specifies that both external and internal BGP paths can be selected for multipath.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4. The `equal-external-internal` option was added for Junos OS Release 8.4.

RELATED DOCUMENTATION

Configuring Protocol-Independent Load Balancing in Layer 3 VPNs

Load Balancing and IP Header Filtering for Layer 3 VPNs

vpn-global-import

IN THIS SECTION

- Syntax | 30953
- Hierarchy Level | 30953
- Description | 30954
- Required Privilege Level | 30954
- Release Information | 30954

Syntax

```
vpn-global-import [ vpn-global-import ... ];
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options rib],
[edit logical-systems name routing-options rib],
[edit logical-systems name tenants name routing-instances name routing-options rib],
[edit routing-instances name routing-options rib],
```

```
[edit routing-options rib]
```

Description

Enable leaking of BGP VPN routes to global RIBs to provide service providers the flexibility to allow internet access to VPN customers. Leaking routes of a customer VPN into master global Internet RIBs like `inet.0`and `inet6.0` is for customer route prefixes that come from global Internet namespace and are not private addresses.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 24.2R1.

RELATED DOCUMENTATION

route-distinguisher-id-use-router-id | 15889

vpws-service-id

IN THIS SECTION

- Syntax | 30955
- Hierarchy Level | 30955
- Description | 30955
- Options | 30955
- Required Privilege Level | 30956
- Release Information | 30956

Syntax

```
vpws-service-id {
    local service-id;
    remote serivce-id;
    source-packet-routing {
        srv6;
        locator locator-name {
            end-dx2-sid end-dx2-sid;
            micro-dx2-sid micro-sid-value;
            }
    }
}
```

Hierarchy Level

```
[edit routing-instances instance-name protocols evpn interface interface-name]
```

Description

Specify the local and remote Ethernet VPN (EVPN)-Virtual Private Wire Service (VPWS) service identifiers. These service identifiers are unique to an EVPN and are used to identify the endpoints of the EVPN-VPWS network. These endpoints are auto-discovered by BGP and are used to exchange the service labels (learned from the respective provider edge (PE) routers) that are used by auto-discovered routes per EVI route type. Depending on the mode of operation of the PE routers in the EVPN-VPWS network, these two endpoints of the can be colocated on the same PE router or on different PE routers.

Options

local

Unique local VPWS service identifier of the EVPN-VPWS network. This identifies the logical interface of the PE router forwarding traffic to the PE router with a remote VPWS service identifier in the EVPN-VPWS network.

- **Range:** 1 through 16,777,215

remote

Unique remote VPWS service identifier of the EVPN-VPWS network. This identifies the logical interface of the PE router receiving the traffic from the PE router with a local VPWS service identifier in the EVPN-VPWS network.

- **Range:** 1 through 16,777,215

srv6 Enable segment routing over IPv6.

locator *locator-name* Specify a name to identify the locator.

end-dx2-sid *end-dx2-sid* Specify the SID to enable endpoint decapsulation of EVPN-VPWS Layer 2 cross connect circuit.

micro-dx2-sid *micro-sid-value* Enable micro-sid by specifying the value for micro-SID.

> **NOTE**: The micro-SID value must be in the static SID range of the local micro-SIDs. To display the range of local static SIDs in the local micro-SID, use the `show srv6 block` command.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

`source-packet-routing` statement introduced in Junos OS Release 22.4R1.

`locator` statement introduced in Junos OS Release 22.4R1.

`end-dx2-sid` statement introduced in Junos OS Release 22.4R1.

`micro-dx2-sid` statement introduced in Junos OS Evolved Release 24.2R1

RELATED DOCUMENTATION

Overview of VPWS with EVPN Signaling Mechanisms

EVPN Multihoming Overview

Configuring VPWS with EVPN Signaling Mechanisms

Example: Configuring VPWS with EVPN Signaling Mechanisms

show evpn vpws-instance | **19840**

Configuring EVPN-VPWS over SRv6

vrf-advertise-selective

IN THIS SECTION

- Syntax | 30957
- Hierarchy Level | 30957
- Description | 30957
- Required Privilege Level | 30958
- Release Information | 30958

Syntax

```
vrf-advertise-selective {
    family {
        inet-mvpn;
        inet6-mvpn;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Explicitly enable IPv4 or IPv6 MVPN routes to be advertised from the VRF instance while preventing all other route types from being advertised.

If you configure the `vrf-advertise-selective` statement without any of its options, the router or switch has the same behavior as if you configured the `no-vrf-advertise` statement. All VPN routes are prevented from being advertised from a VRF routing instance to the remote PE routers. This behavior is useful for hub-and-spoke configurations, enabling you to configure a PE router to not advertise VPN routes from the primary (hub) instance. Instead, these routes are advertised from the secondary (downstream) instance.

The options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Limiting Routes to Be Advertised by an MVPN VRF Instance

no-vrf-advertise | 11620

vrf-export

IN THIS SECTION

- Syntax | 30958
- Hierarchy Level | 30959
- Description | 30959
- Default | 30959
- Options | 30959
- Required Privilege Level | 30959
- Release Information | 30959

Syntax

```
vrf-export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name]
[edit protocols evpn interconnect]
[edit routing-instances routing-instance-name]
[edit routing-instances routing-instance-name protocols evpn interconnect]
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]
[edit switch-options]
```

Description

Specify how the device exports routes from the routing table *routing-instance-name*.inet.0 to other devices. If the `instance-type` in the routing instance configuration is `vrf`, you must either:

- Include this statement to apply a routing export policy in the routing instance.
- Use the `vrf-target` statement with a *community* value. In that case, the device advertises all routes with a matching community value.

You can configure multiple export policies on the device.

Default

For instance-type `vrf`, we require you to configure the `vrf-export` or the `vrf-target` statement. The default action is to reject.

Options

policy-names—Names for the export policies.

Required Privilege Level

routing— To view this statement in the configuration.

routing-control— To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the following hierarchy levels introduced in Junos OS Release 20.3R1 on QFX Series switches: `[edit protocols evpn interconnect]` and `[edit routing-instances routing-instance-name protocols evpn interconnect]`.

RELATED DOCUMENTATION

Implementing EVPN-VXLAN for Data Centers

instance-type | **8773**

Configuring Policies for the VRF Table on PE Routers in VPNs

vrf-import

IN THIS SECTION

- Syntax | **30960**
- Hierarchy Level | **30960**
- Description | **30961**
- Options | **30961**
- Required Privilege Level | **30961**
- Release Information | **30961**

Syntax

```
vrf-import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit  logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls mesh-group mesh-group-name]
[edit protocols evpn interconnect]
[edit routing-instances routing-instance-name]
```

```
[edit routing-instances routing-instance-name protocols evpn interconnect]
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]
[edit switch-options]
```

Description

Specify how routes are imported into the routing table (*routing-instance-name*.inet.0) of the local device from the remote device.

You can configure multiple import policies on the device.

One of the following statements is required for importing routes:

- **vrf-target** - When you configure the `vrf-target community` statement without the `vrf-import` statement, by default the device accepts all routes matching the specified target community.
- **vrf-import** - When you configure the `vrf-import` statement instead of the `vrf-target` statement with a community value, there is no default action. The device imports only the routes accepted in the configured `vrf-import` policy statement.

Options

policy-names—Names for the import policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the following hierarchy levels introduced in Junos OS Release 20.3R1 on QFX Series switches: `[edit protocols evpn interconnect]` and `[edit routing-instances routing-instance-name protocols evpn interconnect]`.

RELATED DOCUMENTATION

Implementing EVPN-VXLAN for Data Centers

instance-type | 8773

Configuring Policies for the VRF Table on PE Routers in VPNs

vrf-mtu-check

IN THIS SECTION

- Syntax | 30962
- Hierarchy Level | 30962
- Description | 30962
- Default | 30962
- Required Privilege Level | 30963
- Release Information | 30963

Syntax

```
vrf-mtu-check;
```

Hierarchy Level

```
[edit chassis]
```

Description

On M Series routers (except the M120 and M320 router) and on EX Series 8200 switches, configure path maximum transmission unit (MTU) checks on the outgoing interface for unicast traffic routed on a virtual private network (VPN) routing and forwarding (VRF) instance.

Default

Disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Path MTU Checks for VPN Routing Instances

Configure Path MTU Discovery

vrf-name (Duplicate Accounting)

IN THIS SECTION

- Syntax | 30963
- Hierarchy Level | 30964
- Description | 30964
- Options | 30964
- Required Privilege Level | 30964
- Release Information | 30964

Syntax

```
vrf-name vrf-name;
```

Hierarchy Level

```
[edit access profile profile-name accounting duplication-vrf]
```

Description

Specify a nondefault VRF (LS:RI combination) to which duplicate accounting information is sent. Up to five access profiles can be defined in this VRF; the profiles point to the RADIUS accounting servers that receive the accounting information.

Options

vrf-name Name of a nondefault VRF to receive duplicate accounting reports.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.

Statement supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

RELATED DOCUMENTATION

RADIUS Accounting for Subscriber Access

RADIUS Authentication and Accounting Basic Configuration

vrf-propagate-ttl

IN THIS SECTION

- Syntax | 30965
- Hierarchy Level | 30965
- Description | 30965
- Default | 30965
- Required Privilege Level | 30966
- Release Information | 30966

Syntax

```
vrf-propagate-ttl;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Enable normal time-to-live (TTL) decrementing in a VRF routing instance. You configure this statement once per routing instance, and it affects only RSVP-signaled or LDP-signaled LSPs in the routing instance. When this router acts as an ingress router for an LSP, it pushes an MPLS copies the TTL value from the IP packet header. When the router acts as the penultimate router, it pops the MPLS header and writes the MPLS TTL into the IP packet.

Default

Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP. This statement explicitly configures the default behavior for the VRF routing instance and is useful for overriding the `no-propagate-ttl` configured globally on the router at the `[edit protocols mpls]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

> *Example: Diagnosing Networking Problems Related to Layer 3 VPNs by Disabling TTL Decrementing*
>
> Disabling Normal TTL Decrementing

vrf-table-label

IN THIS SECTION

- Syntax | 30966
- Hierarchy Level | 30967
- Description | 30967
- Options | 30967
- Required Privilege Level | 30968
- Release Information | 30968

Syntax

```
vrf-table-label {
    source-class-usage;
    static;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit routing-instances routing-instance-name]
```

Description

Map the inner label of a packet to a specific VPN routing and forwarding (VRF) instance. This allows the examination of the encapsulated IP header. The first lookup is done on the VPN label to determine which VRF instance to refer to, and the second lookup is done on the IP header to determine how to forward packets to the correct end hosts.

When you include the `vrf-table-label` statement in the configuration of a VRF routing instance, a label-switched interface (LSI) logical interface label is created and mapped to the VRF routing table. Any routes in the VRF routing table are advertised with the LSI logical interface label allocated for the VRF routing table. When packets destined for the VRF routing instance arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the LSI interface and are then forwarded and filtered based on the correct table.

All routes in a VRF routing instance configured with this option are advertised with one label allocated per VRF.

NOTE:

- The `vrf-table-label` statement is supported on PTX5000 and PTX3000 routers only when third-generation FPCs are installed on the router and `enhanced-ip` command is configured on the chassis.
- Starting in Junos OS Release 17.2, you can configure the `enhanced-ip` command, which is supported on platforms using Modular Port Concentrators (MPCs) equipped with Junos Trio chipsets. You can also separate the MPLS labels used for different label spaces which provides more flexibility and scalability. The `vrf-table-label` space is increased to at least 16,000, if the platform can support the scale.

Options

The remaining statements are explained separately.

- **Range:** 16 through 1,048,575 for static label value.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the `source-class-usage` statement added in Junos OS Release 9.3.

Support for the `static` statement added in Junos OS Release 17.2.

RELATED DOCUMENTATION

Filtering Packets in Layer 3 VPNs Based on IP Headers

Configuring EXP-Based Traffic Classification for VPLS

Load Balancing and IP Header Filtering for Layer 3 VPNs

vrf-target

IN THIS SECTION

- Syntax | 30969
- Hierarchy Level | 30969
- Description | 30969
- Options | 30970
- Required Privilege Level | 30970
- Release Information | 30970

Syntax

```
vrf-target {
    community;
    auto;
    import community;
    export community;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
l2vpn mesh-group mesh-group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls
mesh-group mesh-group-name],
[edit protocols evpn interconnect],
[edit routing-instances routing-instance-name protocols evpn interconnect],
[edit routing-instances routing-instance-name protocols evpn vni-options vni vni-value],
[edit routing-instances routing-instance-name protocols l2vpn mesh-group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name],
[edit switch-options]
```

Description

Specify a virtual routing and forwarding (VRF) target community. This value is a route target extended community.

The `vrf-target` statement simplifies configuring a route target by enabling you to configure the statements for route sharing in a routing instance at the `[edit routing-instances]` hierarchy level. In effect, this statement configures a single policy for import and a single policy for export to replace per routing instance policies for every community.

If you configure this statement with only the *community* option, the device generates default VRF import and export policies to import and export routes with the specified route target community. In this case, you can configure matching route targets for corresponding routing instances on different devices. Then those devices can share routes without requiring you to explicitly configure import and export routing policies.

You can explicitly specify import and export route target communities for the routing instance using the `import` or `export` options. Or, if you use the `auto` option, the device automatically derives the route target community value. For example, see *Auto-derived Route targets* for EVPN routing instances.

You can create more complex policies by explicitly configuring routing instance import and export policies using the "vrf-import" on page 30960 and "vrf-export" on page 30958 statements. With those statements, you define the extended community as a routing policy option using the "community (Policy Options)" on page 4465 statement, which includes a *community-name* and one or more community values.

If you don't configure matching route targets with default import and export policies for a routing instance on devices where you want to share routes, you must ensure that for each routing instance, you configure the export target on each device to match the import target on the other devices, and vice versa, as needed.

Options

community—Community name in the format `target:community-id`, where target is a required keyword followed by a colon, and community-id consists of two numbers (each in the range 0 through 65535) separated by another colon in the format *x*:*y*. The first number, *x*, is often (but not required to be) a related autonomous system (AS) number. See "community (Policy Options)" on page 4465 for more on defining communities (including route targets), and *Configuring EVPN Routing Instances* for more on configuring EVPN instances.

auto—Automatically derives the route target. The auto-derived route targets have higher precedence over manually configured route targets in vrf-target, vrf-export policies, and vrf-import policies.

> **NOTE**: Auto-derived route targets are supported only in virtual switch and EVPN routing instances.

`import` *community*—Accept routes from neighbor devices with the specified route target community.

`export` *community*—Advertise routes with the specified route target community to neighbor devices.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`auto option` added in Junos OS Release 19.1R1 for MX series.

Support at the following hierarchy levels introduced in Junos OS Release 20.3R1 on QFX Series switches: `[edit protocols evpn interconnect]` and `[edit routing-instances routing-instance-name protocols evpn interconnect]`.

RELATED DOCUMENTATION

> *Auto-derived Route targets*
>
> *Configuring Policies for the VRF Table on PE Routers in VPNs*
>
> *Example: Configuring FEC 129 BGP Autodiscovery for VPWS*

vrrp

IN THIS SECTION

- Syntax | 30971
- Hierarchy Level | 30972
- Description | 30972
- Options | 30972
- Required Privilege Level | 30972
- Release Information | 30972

Syntax

```
vrrp {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the Virtual Router Redundancy Protocol (VRRP) process.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the VRRP process.
- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take the primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

> **NOTE**: On SRX300 and SRX320 devices, you cannot configure the same VRRP group ID on different interfaces of a single device

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

vrrp-group

IN THIS SECTION

- Syntax | 30973
- Hierarchy Level | 30974
- Description | 30974
- Options | 30974
- Required Privilege Level | 30976
- Release Information | 30976

Syntax

```
vrrp-group group-id {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    advertisements-threshold number;
    authentication-key key;
    authentication-type (md5 | simple);
    fast-interval milliseconds;
    no-accept-data;
    no-preempt;
    preempt {
    hold-time seconds;
    }
    preferred;
    priority number;
    track {
        interface name {
            bandwidth-threshold bits-per-second priority-cost priority {
                priority-cost priority;
            }
        }
        priority-hold-time seconds;
        route route_address {
            routing-instance;
        }
```

```
    }
    virtual-address virtual-link-local-address;
    vrrp-inherit-from {
        active-group active-group;
        active-interface active-interface;
    }
}
```

Hierarchy Level

```
[edit interfaces name unit name family inet address],
[edit interfaces name unit name family inet inet6 address]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet address address]
```

Description

The Routing Engine performs one-by-one state change for subsequent VRRP groups. Every time there is a state change, and the new state for a particular VRRP group is primary, the Routing Engine generates appropriate VRRP announcement messages. When the first VRRP group detected by the Routing Engine changes state, and the new state is primary, the Routing Engine generates appropriate VRRP announcement messages. The Packet Forwarding Engine is informed about the state change, so that hardware filters for that group are reprogrammed without delay. The new primary then sends gratuitous ARP message to the VRRP groups.

Options

group-id

VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from 00:00:5e:00:53:00 through 00:00:5e:00:53:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

> **NOTE**: Under certain circumstances, the group identifier that you enter must be different from any other group identifiers that you configured for logical units of this same physical

interface. See *Configuring Basic VRRP Support* for more information.

- **Range:** `0` through `255`

accept-data

Accept packets destined for virtual IP address

advertise-interval

Advertisement interval (seconds)

- **Range:** `1` through `255`

inet6-advertise-interval

Inet6 advertisement interval (milliseconds)

- **Range:** `100-40000`

advertisements-threshold

Number of vrrp advertisements missed before declaring primary down

- **Range:** `1` through`15`

authentication-key

Authentication key

authentication-type

Authentication type

- Values:
  - `md5`—HMAC-MD5-96
  - `simple`—Simple password

fast-interval

Fast advertisement interval (milliseconds)

- **Range:** `10` through `40950`

no-accept-data

Don't accept packets destined for virtual IP address

no-preempt

Don't allow preemption

preempt

Allow preemption

preferred

Preferred group on subnet

priority

Device election priority

- **Range:** `0` through `255`

track

Interfaces to track for VRRP group

virtual-address	One or more virtual IPv4 addresses
virtual-inet6-address	One or more virtual inet6 addresses
vrrp-inherit-from	VRRP group to follow for the vrrp-group or vrrp-inet6-group

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface

Release Information

Statement introduced before Junos OS Release 18.1R1.

RELATED DOCUMENTATION

Understanding VRRP on SRX Series Devices

Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces

vrrp-inet6-group

IN THIS SECTION

- Syntax | 30977
- Hierarchy Level | 30977
- Description | 30977
- Options | 30978
- Required Privilege Level | 30978
- Release Information | 30978

Syntax

```
vrrp-inet6-group group-id {
    (accept-data | no-accept-data);
    advertisements-threshold number;
    fast-interval milliseconds;
    inet6-advertise-interval seconds;
    (preempt | no-preempt) {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bits-per-second priority-cost priority;
            priority-cost priority;
        }
        priority-hold-time seconds;
        route prefix/prefix-length routing-instance instance-name priority-cost priority;
    }
    virtual-inet6-address [ addresses ];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from vrrp-group;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6 address address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet6 address address]
```

Description

Configure a Virtual Router Redundancy Protocol (VRRP) IPv6 group.

NOTE: The group identifier that you enter must be different from any other group identifiers that you configured for logical units of this same physical interface.

Options

group-id—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the `source-address-filter` statement. MAC addresses ranging from **00:00:5e:00:01:00** through **00:00:5e:00:01:ff** are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

- **Range:** 0 through 255

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Basic VRRP Support

vrrp-inherit-from

IN THIS SECTION

- Syntax | 30979
- Hierarchy Level | 30979
- Description | 30979
- Options | 30979
- Required Privilege Level | 30979
- Release Information | 30979

Syntax

```
vrrp-inherit-from {
    active-group group-index;
    active-interface active-interface-name;
}
```

Hierarchy Level

```
[edit  interfaces interface-name unit logical-unit-number family inet6 vrrp-inet6-group group-id]
[edit  interfaces interface-name unit logical-unit-number family inet No Link Title group-id]
```

Description

VRRP group to follow for the vrrp-group or vrrp-inet6-group.

Options

group-index—Identifier for VRRP active group.

- **Range:** 0 through 255

active-interface-name—Interface name of VRRP active group.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding VRRP

vsi-discovery

IN THIS SECTION

- Syntax | 30980
- Hierarchy Level | 30980
- Description | 30980
- Default | 30980
- Options | 30981
- Required Privilege Level | 30981

Syntax

```
vsi-discovery {
    interface interface-name
    vsi-policy vsi-policy-name
}
```

Hierarchy Level

```
[edit protocols edge-virtual-bridging]
```

Description

Configure Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP). VDP is used to program policies for each individual station interface (VSI).

Default

VDP is disabled by default.

Options

interface-name Name of the interface on which VDP is configured. The remaining statement is explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on an EX Series Switch

Configuring Edge Virtual Bridging on an EX Series Switch

vsi-policy

IN THIS SECTION

- Syntax | 30981
- Hierarchy Level | 30982
- Description | 30982
- Required Privilege Level | 30982

Syntax

```
vsi-policy vsi-policy-namefrom vsi-manager vsi-manager-id vsi-type vsi-type vsi-version vsi-
version vsi-instance instance-number;
```

Hierarchy Level

```
[edit policy-options]
```

Description

Define and apply the named VSI policy to the edge virtual bridging (EVB) configuration. For use with edge virtual bridging, each virtual machine (VM) on the server is uniquely identified by following four parameters, which are contained in a VSI policy:

- vsi-manager-id
- vsi-type
- vsi-version
- vsi-instance-id

The vsi-policy command manually configures these four parameters on the EX switch for the successful association of VM-VSI. VDP protocol helps determine the parameters defined for the virtual machines on the server and configure them on the switch. Use policy options to define the VM-VSI parameters. Configure a firewall filter for each of the VM profiles and use it in this statement.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

> *Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on an EX Series Switch*
>
> *Configuring Edge Virtual Bridging on an EX Series Switch*

vstp

IN THIS SECTION

- MX Series | 30983
- EX Series, QFX Series, QFabric | 30984
- ELS versions: EX Series, QFX Series, NFX Series | 30985
- Hierarchy Level | 30986
- Description | 30986
- Default | 30987
- Required Privilege Level | 30987
- Release Information | 30987

MX Series

```
vstp {
    bpdu-block-on-edge;
    force-version stp;
    interface interface-name {
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode (p2p | shared);
        no-root-port;
        priority interface-priority;
    }
    priority-hold-time seconds;
    vlan vlan-id {
        bridge-priority priority;
        forward-delay seconds;
        hello-time seconds;
        max-age seconds;
        interface interface-name {
```

```
            access-trunk
            bpdu-timeout-action {
                alarm;
                block;
            }
            cost cost;
            edge;
            mode (p2p | shared);
            no-root-port;
            priority interface-priority;
            system-identifier 00:00:01:00:00:01
        }
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

EX Series, QFX Series, QFabric

```
vstp {
                        bpdu-block-on-edge;
                        disable;
                        force-version stp;
                            vlan (all | vlan-id | vlan-name) {
                                bridge-priority                        priority;
                                forward-delay                        seconds;
                                hello-time                        seconds;
                                    interface (all | interface-name)  {
            arp-on-stp;
            bpdu-timeout-action {
                block;
                log;
            }
                                    cost                                cost;
            disable;
                                    edge;
                                    mode                                mode;
                                    no-root-port;
                                    priority                                priority;
```

```
        }
                                max-age                          seconds;
                               system-identifier 00:00:01:00:00:01
        traceoptions {
            file filename <files number > <size size> <no-stamp | no-world-readable | world-
readable>;
            flag flag;
        }
    }
}
```

ELS versions: EX Series, QFX Series, NFX Series

```
vstp {
    bpdu-block-on-edge;
    disable;
    force-version stp;
    interface  (interface-name disable | interface-range-name | all ){
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode (p2p | shared);
        no-root-port;
        priority interface-priority;
    }
priority-hold-time seconds;
    vlan (vlan-id | all){
        bridge-priority priority;
        forward-delay seconds;
        hello-time seconds;
        max-age seconds;
        interface (interface-name disable | interface-range-name | all ){
            bpdu-timeout-action {
                alarm;
                block;
            }
            cost cost;
            edge;
```

```
            mode (p2p | shared);
            no-root-port;
            priority interface-priority;
            system-identifier 00:00:01:00:00:01
        }
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    vlan-group group group-name {
        vlans vlan-name (vlan-id |vlan-range | open-set-of-values) {
            interface all;
            interface interface-name {
                disable;
            }
        }
    }
```

Hierarchy Level

```
[edit protocols]
```

```
[edit logical-systems logical-system-name protocols],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],
```

```
[edit protocols],
```

```
[edit routing-instances routing-instance-name protocols]
```

Description

Configures VLAN Spanning Tree Protocol (VSTP). VSTP is used to prevent loops in Layer 2 networks on a per-VLAN basis.

You can have a maximum of 253 VSTP VLANs per switch.

If the number of VLANs on your switch exceeds the VSTP VLAN limit, you must use the `vlan` statement to specify which VLANs or VLAN groups use VSTP. You also cannot use the **vlan all** option to configure VSTP when your switch has more than the maximum allowed VSTP VLANs. To ensure all VLANs are running a spanning-tree protocol, run RSTP for networks with large numbers of VLANs .

NOTE: When you configure VSTP with the `set protocol vstp vlan all` command, VLAN ID 1 is not set. If you want VLAN ID 1 to be included in the VSTP configuration on your switch, you must set it separately with the `set protocol vstp vlan 1` command.

NOTE: Option `vlan all` is not supported in Junos OS Release 13.2X50.

BEST PRACTICE: Configure RSTP when you configure VSTP. RSTP overhead is minimal and this configuration ensures that some spanning tree protocol is running on all VLANs on your switch, even when your switch has more than the maximum number of allowed VSTP VLANs.

By default, the chassis MAC address is advertised as the bridge address in the BPDU. If you want to change the bridge address advertised in the BPDU, you must configure the system identifier by using the system-identifier option.

The remaining statements are explained separately. See CLI Explorer.

Default

VSTP is not enabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

`bpdu-block-on-edge` statement added in Junos OS Release 9.4.

`bpdu-timeout-action` statement added in Junos OS Release 9.4.

Support for logical systems added in Junos OS Release 9.6.

Statement updated in Junos OS Release 15.1 for EX Series switches to support configuration of spanning tree parameters globally on all interfaces.

`system-identifier` statement added in Junos OS Release 11.3.

NOTE: You cannot disable spanning tree parameters globally on all interfaces.

RELATED DOCUMENTATION

- Understanding VSTP
- Configuring VLAN Spanning Tree Protocol
- Configuring VSTP (CLI Procedure)
- Example: Configuring VSTP on QFX Series Switches and EX4600 Switches
- show spanning-tree bridge | **26548**
- show spanning-tree interface | **26555**

vtep-source-interface

IN THIS SECTION

- Syntax | **30988**
- Hierarchy Level | **30989**
- Description | **30989**
- Options | **30989**
- Required Privilege Level | **30989**
- Release Information | **30989**

Syntax

```
vtep-source-interface
interface-name (inet | inet6 );
```

Hierarchy Level

```
[edit switch-options]
[edit routing-instances routing-instance-name]
```

Description

Configure a source interface for a Virtual Extensible LAN (VXLAN) tunnel. You must provide the name of a logical interface configured on the loopback interface.

If you specify an IPv4 address (`inet` option) as the VTEP source address, the VXLAN tunnel endpoint (VTEP) encapsulates VXLAN packets with an IPv4 outer header and tunnels the packets through an IPv4 underlay network.

If you specify an IPv6 address (`inet6` option) as the VTEP source address, the VTEP encapsulates VXLAN packets with an IPv6 outer header and tunnels the packets through an IPv6 underlay network.

Options

interface-name	Loopback interface name.
inet	IPv4 source address.
inet6	IPv6 source address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

Support at the `[edit routing-instances routing-instance-name]` hierarchy level introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

Understanding VXLANs

EVPN-VXLAN with an IPv6 Underlay

Manually Configuring VXLANs on QFX Series and EX4600 Switches

Examples: Manually Configuring VXLANs on QFX Series and EX4600 Switches

vulnerability-type

IN THIS SECTION

- Syntax | 30990
- Hierarchy Level | 30990
- Description | 30990
- Options | 30991
- Required Privilege Level | 30991
- Release Information | 30991

Syntax

```
vulnerability-type {
    values [ values ];
}
```

Hierarchy Level

```
[edit security idp (Security)              dynamic-attack-group              name
filters]
```

Description

Vulnerability type of attack.

Vulnerabilities are the intersection of three elements: a system susceptibility or flaw, attacker access to the flaw, and attacker capability to exploit the flaw. A security risk is often incorrectly classified as a vulnerability.

Using this field you can perform vulnerability scanning. Vulnerability scanning is an inspection of the potential points of exploit on a network to identify security issues. A vulnerability scan detects and classifies system weaknesses in a networks and predicts the effectiveness of countermeasures.

Options

values Values for vulnerability-type field (for example: buffer overflow, injection, use after free, Cross-site scripting (XSS), Remote Code Execution (RCE), and so on. Specifying the vulnerability type for IDP will indicate which applications are weak and therefore can be manipulated. The type of vulnerability is reported for fixing these vulnerabilities.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 18.2R1.

vxlan

IN THIS SECTION

- Syntax | 30992
- Hierarchy Level | 30992
- Description | 30992
- Options | 30992
- Required Privilege Level | 30992
- Release Information | 30993

Syntax

```
vxlan {
    encapsulate-inner-vlan;
    ingress-node-replication;
    multicast-group;
    ovsdb-managed;
    riot-loopback;
    unreachable-vtep-aging-timer;
    vni;
    multicast-group multicast-group;
    mtu mtu;
}
```

Hierarchy Level

```
[edit vlans name],
[edit bridge-domains bridge-domain-name]
```

Description

Configure support for Virtual Extensible LANs (VXLANs) on a Juniper Networks device.

Options

multicast-group *multicast-group*	Multicast group (IPv4 or IPv6 addresses) registered for VXLAN segment.
mtu *mtu*	(cRPD only) Specify the maximum transmission unit (MTU) size for a VXLAN packet. Range: 100 to 65535 bytes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

`ingress-node-replication` option added for EVPN VXLAN on QFX5100 switches in Junos OS Release 14.1X53-D30.

`multicast-group` *`multicast-group`* option added for MX series routers with MPC and MIC interfaces in Junos OS Release 17.2.

`mtu` option added for cRPD in Junos OS Release 21.2R1.

`riot-loopback` option added for EVPN-VXLAN on QFX5210 switches in Junos OS Release 21.3R1.

RELATED DOCUMENTATION

Understanding VXLANs

Using a RIOT Loopback Port to Route Traffic in an EVPN Network

Manually Configuring VXLANs on QFX Series and EX4600 Switches

Examples: Manually Configuring VXLANs on QFX Series and EX4600 Switches

vxlan-disable-copy-tos-decap

IN THIS SECTION

- Syntax | 30994
- Hierarchy Level | 30994
- Description | 30994
- Default | 30995
- Required Privilege Level | 30995
- Release Information | 30995

Syntax

```
vxlan-disable-copy-tos-decap
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Disable copying the Type of Service (ToS) field from the outer IP header in a Virtual Extensible LAN (VXLAN) packet to the inner IP header when de-encapsulating the original packet. The ToS field includes values for differentiated services code points (DSCP) and explicit congestion notification (ECN).

In networks where you configure quality of service (QoS) options, traffic moving through the network might include a DSCP value in the ToS field of the IP header for classifying and policing packets. Similarly, network nodes might use the ECN bits in the ToS field of the IP header to enable end-to-end notification of network congestion so the sender can reduce the transmission rate before nodes must start dropping packets. If you have VXLAN tunnels configured in your network, when encapsulating the original packet in the VXLAN header, some devices (such as the QFX5000 line of switches) copy the DSCP and ECN bits from the original packet's IP header (the inner IP header) to the VXLAN IP header (the outer IP header) by default. At the other end of the tunnel, the same types of devices de-encapsulate the packet and copy the outer IP DSCP and ECN bits back to the inner IP header.

However, some devices (such as the QFX10000 line of switches) don't have the ability to copy ToS field bits upon VXLAN encapsulation and de-encapsulation. As a result, you might see unexpected results if the devices on both ends of the tunnel don't consistently copy the ToS bits during encapsulation and de-encapsulation. For example, the decapsulating node might overwrite the inner IP header DSCP bits in a packet with 000000, which inadvertently lowers the priority of the packet to best effort.

To account for these differences in VXLAN encapsulation behavior, you can configure the `vxlan-disable-copy-tos-encap` statement to disable copying ToS field values from inner to outer IP headers during encapsulation. Similarly, you can configure the `vxlan-disable-copy-tos-decap` statement to disable copying ToS values from the outer IP header back to the inner IP header during de-encapsulation.

> **CAUTION**: Committing either `vxlan-disable-copy-tos-encap` or `vxlan-disable-copy-tos-decap` restarts the FPC causing a disruption in traffic until the FPC recreates the interfaces and tunnels.

Default

Copying ToS bits during VXLAN tunnel de-encapsulation is enabled on supporting devices.

> **NOTE**: Some platforms support copying and disabling copy of only the DSCP bits in the ToS field, and not the ECN bits. Other platforms (such as EX4650 and QFX5120 switches) support copying and disabling copy of both the DSCP and ECN bits. See Feature Explorer for specific platform and release version support details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to a configuration.

Release Information

Statement introduced in Junos OS Release 18.4R2 to disable copying the DSCP bits in the ToS field during VXLAN de-encapsulation.

Support added in Junos OS Release 21.1R1 for copying and disabling copy of ECN bits in the ToS field during VXLAN de-encapsulation.

RELATED DOCUMENTATION

Understanding VXLANs

vxlan-disable-copy-tos-encap | **30995**

vxlan-tos-copy-filter | **31011**

vxlan-disable-copy-tos-encap

IN THIS SECTION

- Syntax | **30996**
- Hierarchy Level | **30996**

- Description | 30996
- Default | 30997
- Required Privilege Level | 30997
- Release Information | 30997

Syntax

```
vxlan-disable-copy-tos-encap
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Disable copying the Type of Service (ToS) field from the IP header of a packet to the outer IP header of a Virtual Extensible LAN (VXLAN) packet when encapsulating the packet. The ToS field includes values for differentiated services code points (DSCP) and explicit congestion notification (ECN).

In networks where you configure quality of service (QoS) options, traffic moving through the network might include a DSCP value in the ToS field of the IP header for classifying and policing packets. Similarly, network nodes might use the ECN bits in the ToS field of the IP header to enable end-to-end notification of network congestion so the sender can reduce the transmission rate before nodes must start dropping packets. If you have VXLAN tunnels configured in your network, when encapsulating the original packet in the VXLAN header, some devices copy the DSCP and ECN bits from the original packet's IP header (the inner IP header) to the VXLAN IP header (the outer IP header) by default. At the other end of the tunnel, the same types of devices decapsulate the packet and copy the outer IP DSCP and ECN bits back to the inner IP header.

However, some devices (such as the QFX10000 line of switches) don't have the ability to copy ToS field bits upon VXLAN encapsulation and decapsulation. As a result, you might see unexpected results if the devices on both ends of the tunnel don't consistently copy the ToS bits during encapsulation and decapsulation. For example, the decapsulating node might overwrite the inner IP header DSCP bits in a packet with 000000, which inadvertently lowers the priority of the packet to best effort.

To account for these differences in VXLAN encapsulation behavior, you can configure the `vxlan-disable-copy-tos-encap` statement to disable copying ToS field values from inner to outer IP headers during encapsulation. Similarly, you can configure the `vxlan-disable-copy-tos-decap` statement to disable copying ToS values from the outer IP header back to the inner IP header during de-encapsulation.

NOTE: QFX5130 and QFX5700 switches do not require `vxlan-disable-copy-tos-encap`, as this is the default behavior for these switches.

CAUTION: Committing either `vxlan-disable-copy-tos-encap` or `vxlan-disable-copy-tos-decap` restarts the FPC causing a disruption in traffic until the FPC recreates the interfaces and tunnels.

Default

Copying ToS bits during VXLAN tunnel encapsulation is enabled on supporting devices.

NOTE: Some platforms support copying and disabling copy of only the DSCP bits in the ToS field, and not the ECN bits. Other platforms (such as EX4650 and QFX5120 switches) support copying and disabling copy of both the DSCP and ECN bits. See Feature Explorer for specific platform and release version support details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R2 to disable copying the DSCP bits in the ToS field during VXLAN encapsulation.

Support added in Junos OS Release 21.1R1 for copying and disabling copy of ECN bits in the ToS field during VXLAN encapsulation.

RELATED DOCUMENTATION

Understanding VXLANs

vxlan-disable-copy-tos-decap | 30993
vxlan-tos-copy-filter | 31011

vxlan-extended

IN THIS SECTION

- Syntax | 30998
- Hierarchy Level | 30998
- Description | 30998
- Required Privilege Level | 30999
- Release Information | 30999

Syntax

```
vxlan-extended;
```

Hierarchy Level

```
[edit system packet-forwarding-options system-profile]
```

Description

On ACX7100-32C, ACX7100-48L, and ACX7024 devices, you can use either an IPv4 or IPv6 underlay for a VXLAN configuration. To use an IPv6 underlay, enable the `vxlan-extended` system profile.

Before you enable the `vxlan-extended` system profile, remove any existing VXLAN configurations. After you enable the `vxlan-extended` statement, the packet forwarding engine (PFE) reboots. After the PFE comes back up, you can configure an underlay and overlay for the IPv6 VXLAN underlay. Once you have configured the IPv6 VXLAN underlay and overlay, your VXLAN deployment will be using an IPv6 underlay.

If you want to use an IPv4 underlay, enable the default system profile and remove the IPv6 VXLAN underlay configuration.

To enable the default system profile, issue the `delete system packet-forwarding-options system-profile vxlan-extended` command. Apply your underlay and overlay configuration for your IPv4 VXLAN underlay configuration. After the PFE reboots, your VXLAN deployment will be using an IPv4 underlay.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.4R1.

vxlan-gbp-l2-profile

IN THIS SECTION

- Syntax | 30999
- Hierarchy Level | 31000
- Description | 31000
- Required Privilege Level | 31001
- Release Information | 31001

Syntax

```
vxlan-gbp-l2-profile;
```

Hierarchy Level

```
[edit chassis forwarding-options]
```

Description

Enable `vxlan-gbp-l2-profile` on the tunnel termination endpoint in your EVPN-VXLAN deployment to support group-based policies. This setting tells the switch to allocate a share of its resources for L2/L3 group-based policies, whereas otherwise the resources would remain committed for use by all other flows.

The `vxlan-gbp-l2-profile` is suitable when running an edge device in a virtualized network comprised of multiple virtual-severs and VMs. You can also use this profile for direct reachability if the edge device needs more space in the mac-table.

Note that for switches in a virtual chassis, the device must be rebooted for this setting to apply; for stand-alone switches the packet forwarding engine (PFE) restarts.

Table 1835 on page 31000 shows the maximum GBP supported configuration with 1K unique tags.

NOTE: The `vxlan-gbp-l2-profile` is not applicable for the EX4100 switches.

Table 1835: GBP Scales for the vxlan-gbp-l2-profile

Tag Assignment	Maximum GBP Supported Configuration (EX4100 Series)	Maximum GBP Supported Configuration (EX4400 Series)	Maximum GBP Supported Configuration (EX4650 Series and QFX5120 Series)
MAC	NA	32K	96K
IP	NA	20K	26K
Port	NA	2K	2K
VLAN	NA	3K	3K
Port + VLAN	NA	3K	3K

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 for supported EX4400, EX4650, and QFX5120 Series switches.

RELATED DOCUMENTATION

Example: Micro and Macro Segmentation using Group Based Policy in a VXLAN

vxlan-gbp-l3-profile

IN THIS SECTION

- Syntax | 31001
- Hierarchy Level | 31002
- Description | 31002
- Required Privilege Level | 31003
- Release Information | 31003

Syntax

```
vxlan-gbp-l3-profile;
```

Hierarchy Level

```
[edit chassis forwarding-options]
```

Description

Enable `vxlan-gbp-l3-profile` on the tunnel termination endpoint in your EVPN-VXLAN deployment to support group-based policies. This setting tells the switch to allocate a share of its resources for L2/L3 group-based policies, whereas otherwise the resources would remain committed for use by all other flows.

The `vxlan-gbp-l3-profile` profile is suitable for an edge device running in a L3 intensive network where more space in the host-table is required.

Note that for switches in a virtual chassis, the device must be rebooted for this setting to apply; for stand-alone switches the packet forwarding engine (PFE) restarts.

Table 1836 on page 31002 shows the maximum GBP supported configuration with 1K unique tags.

NOTE: The `vxlan-gbp-l3-profile` is not applicable for the EX4100 switches.

Table 1836: GBP Scales for the vxlan-gbp-l3-profile

Tag Assignment	Maximum GBP Supported Configuration (EX4100 Series)	Maximum GBP Supported Configuration (EX4400 Series)	Maximum GBP Supported Configuration (EX4650 Series and QFX5120 Series)
MAC	NA	32K	64K
IPv4	NA	32K	80K
Port	NA	2K	2K
VLAN	NA	3K	3K
Port + VLAN	NA	3K	3K

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.2R1 for supported EX4400, EX4650, and QFX5120 Series switches.

vxlan-gbp-profile

IN THIS SECTION

- Syntax | 31003
- Hierarchy Level | 31003
- Description | 31004
- Default | 31004
- Required Privilege Level | 31005
- Release Information | 31005

Syntax

```
vxlan-gbp-profile
```

Hierarchy Level

```
[edit chassis forwarding-options]
```

Description

Enable `vxlan-gbp-profile` on the tunnel termination endpoint in your EVPN-VXLAN deployment to support group-based policies. This setting tells the switch to allocate a share of its resources for L2/L3 group-based policies, whereas otherwise the resources would remain committed for use by all other flows. Note that for switches in a virtual chassis, the device must be rebooted for this setting to apply; for stand-alone switches the packet forwarding engine (PFE) will be restarted.

Group-based policies (GBP) make use of existing layer 3 VXLAN network identifiers (VNI), in conjunction with firewall filter policies, to provide micro-segmentation at the level of device or tag, independent of the underlying network topology. For example, IoT devices typically only need access to specific applications on the network, so GBP can keep this traffic isolated by automatically applying security policies without the need for L2 or L3 lookups or ACLs. As such, GBP provides a new approach to network access control and security that is especially valuable for enterprise campuses. The `vxlan-gbp-profile` is suitable for a balanced configuration that contains a mix of L2 and L3 networks.

In addition to enabling `vxlan-gbp-profile` on the tunnel termination endpoint, you need to create firewall rules with match conditions for the endpoint devices you want to segregate. Do this on the EX4400 switch in your topology that is deployed in the role of VXLAN gateway for the access layer.

Table 1837 on page 31004 shows the maximum GBP supported configuration with 1K unique tags.

Table 1837: GBP Scales for the vxlan-gbp-profile

Tag Assignment	Maximum GBP Supported Configuration for EX4100 Series	Maximum GBP Supported Configuration for EX4400 Series	Maximum GBP Supported Configuration for EX4650 Series and QFX5120 Series
MAC	32K	32K	32K
IPv4	10K	16K	26K
Port	500	2K	2K
VLAN	Not applicable	3K	3K
Port + VLAN	Not applicable	3K	3K

Default

Not enabled

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1 for EX4400 Series switches.

RELATED DOCUMENTATION

Example: Micro and Macro Segmentation using Group Based Policy in a VXLAN

vxlan-gpe (FTI)

IN THIS SECTION

- Syntax | 31005
- Hierarchy Level | 31006
- Description | 31006
- Options | 31006
- Required Privilege Level | 31007
- Release Information | 31007

Syntax

```
vxlan-gpe {
    destination {
        address address;
    }
    destination-udp-port destination-udp-port;
    source {
        address address;
```

```
        interface interface;
    }
    source-udp-port-range {
        (min min-port-number| max max-port-number)
    }
tunnel-endpoint vxlan;
    vni vni;
}
```

Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit name tunnel encapsulation],
[edit interfaces name unit name tunnel encapsulation]
```

Description

Virtual Extensible LAN (VXLAN) is defined as an encapsulation format that encapsulates Ethernet frames in an outer UDP/IP transport. The `vxlan-gpe` encapsulation refers to the header `gpe` added to the encapsulation `vxlan` to extend the existing VXLAN protocol to provide support for multiprotocol encapsulation.

Options

destination-udp-port

Value to write to the destination-udp-port field.

- **Range:** 1 through 65,535

source-udp-port-range

Value to write to the source-udp-port-range field.

- **min *min-port-number*:** 1 through 65,535.
- **max *max-port-number*:** 1 through 65,535

min-port-number should be less than or equal to *max-port-number*

If you do not specify the source port range, it is randomly set (in the range 49152 – 65535) based on the hash value calculated using various packet headers field.

tunnel-endpoint

Tunnel endpoint type.

- Values:
  - vxlan—To configure VXLAN tunnel endpoint.

vni Value to specify in the `vni` attribute.

- **Range:** 0 through 16,777,214

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Flexible Tunnel Interfaces Overview

Configuring Flexible Tunnel Interfaces

show interfaces fti | **20866**

vni(Interfaces) | **30911**

Example: Configuring Flexible Tunnel Interfaces on MX Series Routers

vxlan-routing

IN THIS SECTION

- Syntax | **31008**
- Hierarchy Level | **31008**
- Description | **31008**
- Required Privilege Level | **31008**
- Release Information | **31008**

Syntax

```
vxlan-routing {
    interface-num integer;
    loopback-port;
    next-hop (VXLAN Routing) integer;
    overlay-ecmp;
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure Virtual Extensible LAN (VXLAN) routing options on a device.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

`overlay-ecmp` statement introduced in Junos OS Release 17.4R1.

`loopback-port` statement introduced in Junos OS Release 21.3R1 for QFX5210 switches.

vxlan-stitching

IN THIS SECTION

- Syntax | 31009
- Hierarchy Level | 31009
- Description | 31009
- Options | 31010
- Default | 31010
- Required Privilege Level | 31010
- Release Information | 31010

Syntax

```
vxlan-stitching {
    control-word;
}
```

Hierarchy Level

```
[edit system packet-forwarding-options system-profile]
```

Description

You must configure `vxlan-stitching` on ACX devices to support Virtual Extensible LAN (VXLAN) data center interconnect (DCI) stitching features.

NOTE: The PFE will restart once this statement is applied.

Options

control-word Enable this option on an ACX device operating as a DCI gateway (GW) when both of the following are true:

- Your MPLS network is using Ethernet VPN (EVPN-MPLS) for the overlay fabric.
- Your ACX DCI GW's peer device is using `control-word`. For example, you may be connecting across the DCI to a PTX, which only supports `control-word`.

Consider the following when configuring control-word support.

- control-word support is disabled by default on ACX devices. Only enable this option when control-word support is required.
- control-word support is enabled by default in MAC-VRF routing-instances.
- In EVPN-VXLAN to EVPN-MPLS networks, when the peer GW **doesn't** support control-word, configure the `vxlan-stitching` knob as described, and configure the `no-control-word` option in the `[edit routing-instances` *`instance-name`* `protocols evpn]` hierarchy.
- In EVPN-VXLAN to EVPN-MPLS networks, when the peer GW **does** support control-word, configure the `control-word` option in the `[edit system packet-forwarding-options system-profile vxlan-stitching` hierarchy.

Default

Stitching features are disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 23.2R1.

`control-word` option introduced in Junos OS Evolved Release 23.4R1.

RELATED DOCUMENTATION

EVPN-VXLAN Data Center Interconnect Through EVPN-MPLS WAN Overview

vxlan-tos-copy-filter

IN THIS SECTION

- Syntax | 31011
- Hierarchy Level | 31011
- Description | 31011
- Default | 31012
- Required Privilege Level | 31012
- Release Information | 31012

Syntax

```
vxlan-tos-copy-filter
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Copy the Type of Service (ToS) field from the IP header of a packet to the outer IP header of a Virtual Extensible LAN (VXLAN) packet when encapsulating the packet. The ToS field includes values for differentiated services code points (DSCP) and explicit congestion notification (ECN).

Default

For QFX5000 Series switches, `vxlan-tos-copy-filter` is only required for QFX5130 and QFX5700 switches. ToS copy from the inner to outer VXLAN header is the default behavior in other QFX5000 Series switches.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved versions 22.1R3, 22.2R2, 22.3R1.

RELATED DOCUMENTATION

Understanding VXLANs

vxlan-disable-copy-tos-decap | 30993

vxlan-disable-copy-tos-encap | 30995

W

IN THIS SECTION

- wait-for-acct-on-ack (Access Profile) | 31014
- wait-to-block-interval | 31016
- wait-to-restore | 31017
- walkup | 31019
- wander-filter-disable | 31021
- warm-standby (Interfaces) | 31022
- warm-standby | 31023
- warning | 31024

- watch-list | **31026**
- watchdog | **31027**
- watchdog (PSM) | **31028**
- watchdog-timeout | **31031**
- wavelength | **31032**
- web-authentication | **31057**
- web-authentication (Access) | **31058**
- web-authentication (Interfaces) | **31060**
- web-filter | **31062**
- web-filter-profile | **31065**
- web-filtering | **31066**
- web-filtering (Logical System Security Feature Profile) | **31072**
- web-filtering (Security Content Security Policy) | **31076**
- web-management (System Processes) | **31078**
- web-management (System Services) | **31079**
- web-redirect | **31084**
- web-redirect-to-https | **31085**
- web-server (Services) | **31086**
- webapi-clear-text (Security) | **31089**
- webapi-ssl (Security) | **31090**
- webapi (System Services) | **31091**
- websense-redirect | **31093**
- weight | **31095**
- weight (Targeted Subscribers) | **31097**
- weight (Security) | **31099**
- weighted-load-balancing (L2TP LAC) | **31100**
- west-interface | **31102**
- when | **31104**
- white-list | **31106**
- white-list (Security Screen) | **31107**
- white-list (Security-Zone) | **31109**
- whitelist | **31111**

- whitelist-url-categories | 31112
- wildcard delete | 31114
- wildcard-group-inet | 31115
- wildcard-group-inet6 | 31117
- wildcard-source (PIM RPF Selection) | 31119
- wildcard-source (Selective Provider Tunnels) | 31120
- window-scale | 31122
- window-size | 31123
- windows-logon (Juniper Secure Connect) | 31124
- winnuke | 31126
- winnuke (IDS Screen Next Gen Services) | 31127
- wins-server (Access) | 31129
- wins-server (System) | 31130
- wired-max-processes | 31131
- wired-process-mem-size | 31133
- withdraw-priority | 31134
- within (Event Policy) | 31136
- wlan | 31137
- world-readable (Forwarding Options) | 31142
- world-readable (System Syslog) | 31144
- world-readable (Next Gen Services Global System Logging) | 31145
- write-interval | 31146
- write-view | 31148

wait-for-acct-on-ack (Access Profile)

IN THIS SECTION

- Syntax | 31015
- Hierarchy Level | 31015

- Description | 31015
- Required Privilege Level | 31015
- Release Information | 31016

Syntax

```
wait-for-acct-on-ack;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router's authd process to wait for an Acct-On-Ack response message from RADIUS before sending new authentication and accounting updates to the RADIUS server. This configuration ensures that when a new subscriber session starts, the authentication and accounting information for the new session does not get deleted when RADIUS clears previously existing session state information.

At subscriber session startup, the Junos OS authd process sends an Acct-On message to the RADIUS server and the new session starts authentication and accounting operations. However, in some service provider environments, upon receipt of the Acct-On message, the RADIUS server cleans up the previous session state and removes accounting statistics. In this scenario, the RADIUS server's cleanup operation can inadvertently delete the new session's authentication and accounting information, which might include customer billing information.

To ensure that the new session's authentication and accounting information is not deleted, you can include the `wait-for-acct-on-ack` statement to configure authd to wait for an Acct-On-Ack response message from the RADIUS accounting server, so the RADIUS cleanup can finish before authd sends any new authentication and accounting updates.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

RADIUS Accounting for Subscriber Access

RADIUS Servers and Parameters for Subscriber Access

wait-to-block-interval

IN THIS SECTION

- Syntax | 31016
- Hierarchy Level | 31016
- Description | 31016
- Options | 31017
- Required Privilege Level | 31017
- Release Information | 31017

Syntax

```
wait-to-block-interval number;
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Enable the Wait to Block (WTB) timer interval when clearing force switch and manual switch commands.

Options

number—Wait-to-block interval, in seconds.

- **Range:** 5 through 10 s
- **Default:** 5 s

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

> *Ethernet Ring Protection Switching Overview*
>
> Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

wait-to-restore

IN THIS SECTION

- Syntax | 31018
- Hierarchy Level | 31018
- Description | 31018
- Options | 31018
- Required Privilege Level | 31018
- Release Information | 31019

Syntax

```
wait-to-restore minutes;
```

Hierarchy Level

```
[edit chassis source interfaces (external-a | external-b | interface interface-name)]
[edit chassis source interfaces (external | interface interface-name)]
[edit chassis source interfaces (external-0/0 | external-1/0 | interface interface-name)]
```

Description

Configure the time in minutes for each port to be up before opening the Ethernet Synchronization Message Channel (ESMC) for messages. When a port's signal transitions out of the signal fail state, it must be fault-free for the wait-to-restore time before it is again considered by the clock selection process.

> **NOTE**: When you perform GRES on MX Series routers, you must execute the `clear synchronous-ethernet wait-to-restore` operational mode command on the new primary Routing Engine to clear the wait-to-restore timer on it. This is because the `clear synchronous-ethernet wait-to-restore` operational mode command clears the wait-to-restore timer only on the local Routing Engine.

Options

minutes Set the time for the port signal to be up before the port is opened to receive and transmit ESMC messages.

- **Range:** 0 through 12 minutes
- **Default:** 5 minutes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

- synchronization (MX Series) | 28795
- synchronization (PTX Series) | 28802
- global-wait-to-restore | 7741

walkup

IN THIS SECTION

- Syntax | 31019
- Hierarchy Level | 31019
- Description | 31020
- Default | 31020
- Required Privilege Level | 31020
- Release Information | 31020

Syntax

```
walkup;
```

Hierarchy Level

```
[edit logical-systems logical-system-name policy-options defaults route-filter],
[edit logical-systems logical-system-name policy-options policy-statement policy-statement-name
defaults route-filter],
```

```
[edit policy-options defaults route-filter],
[edit policy-options policy-statement policy-statement-name defaults route-filter]
```

Description

Enable route filter walkup globally or locally for a particular policy statement. The walkup feature examines more than the longest match route filters in a policy statement term with more than one route filter, allowing consolidation of terms and a potential performance enhancement.

Default

By default, no route filter walkup is performed and only the longest match route filter in a policy statement term is examined.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

- no-walkup | **11623**
- route-filter | **15861**
- *Walkup for Route Filters Overview*
- *Configuring Walkup for Route Filters to Improve Operational Efficiency*
- *Example: Configuring Walkup for Route Filters Globally to Improve Operational Efficiency*
- *Example: Configuring Walkup for Route Filters Locally to Improve Operational Efficiency*

wander-filter-disable

IN THIS SECTION

- Syntax | 31021
- Hierarchy Level | 31021
- Description | 31021
- Required Privilege Level | 31021
- Release Information | 31021

Syntax

```
wander-filter-disable;
```

Hierarchy Level

```
[edit chassis synchronization output interfaces external]
[edit chassis synchronization output interfaces (external-0/0 | external-1/0)]
```

Description

Disable the wander filter on the output interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

synchronization (MX Series) | 28795

warm-standby (Interfaces)

IN THIS SECTION

- Syntax | 31022
- Hierarchy Level | 31022
- Description | 31022
- Required Privilege Level | 31022
- Release Information | 31023

Syntax

```
warm-standby;
```

Hierarchy Level

```
[edit interfaces rlsqnumber redundancy-options]
```

Description

For AS or Multiservices PIC redundancy configurations, specify that the failure detection and recovery involves one backup PIC supporting multiple working PICs. Recovery time is not guaranteed.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces

warm-standby

IN THIS SECTION

- Syntax | 31023
- Hierarchy Level | 31023
- Description | 31023
- Required Privilege Level | 31024
- Release Information | 31024

Syntax

```
warm-standby;
```

Hierarchy Level

```
[edit routing-options]
```

Description

Set the routing protocols process (rpd) mode to warm standby. Warm standby mode helps the backup RE stay synchronized with the primary RE, allowing for faster RE switchover during GRES.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

Statement introduced in Junos OS Evolved Release 21.2R1.

RELATED DOCUMENTATION

Understanding Graceful Routing Engine Switchover

warning

IN THIS SECTION

- Syntax | 31024
- Hierarchy Level | 31025
- Description | 31025
- Options | 31025
- Required Privilege Level | 31025
- Release Information | 31025

Syntax

```
warning low-light-warning {
    (link-down | syslog);
}
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Description

Specifies the action to take if the receiving optics signal is below the optics low-light warning threshold. Starting in Junos OS Release 15.1, for all QSFP-based interfaces, you need not explicitly configure the `syslog` option. The `syslog` option is enabled by default.

Options

`link-down`—Drop the 10-Gigabit Ethernet link and marks link as down.

`syslog`—Write the optics information to the system log.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2R1.

RELATED DOCUMENTATION

Configuring Link Down Notification for Optics Options Alarm or Warning

optics-options | **11847**

100-Gigabit Ethernet OTN Options Configuration Overview

watch-list

IN THIS SECTION

- Syntax | 31026
- Hierarchy Level | 31026
- Description | 31026
- Options | 31026
- Required Privilege Level | 31027
- Release Information | 31027

Syntax

```
watch-list {
    [ routes ];
}
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN interfaces, configure an ISDN list of routes to watch. Used only for dialer watch.

Options

routes—IP prefix of a route. Specify one or more. The primary interface is considered up if there is at least one valid route for any of the addresses in the watch list to an interface other than the backup interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

watchdog

IN THIS SECTION

- Syntax | 31027
- Hierarchy Level | 31027
- Description | 31028
- Options | 31028
- Required Privilege Level | 31028
- Release Information | 31028

Syntax

```
watchdog {
    disable;
    enable;
    timeout value;
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Enable or disable the watchdog timer when Junos OS encounters a problem.

Options

- `disable`—Disable the watchdog timer.
- `enable`—Enable the watchdog timer.
- `timeout` *`value`*—Specify amount of time to wait in seconds.

  Range: 1 through 3600 seconds.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

watchdog (PSM)

IN THIS SECTION

- Syntax | 31029
- Hierarchy Level | 31029
- Description | 31029
- Options | 31029
- Required Privilege Level | 31030
- Platform-Specific Watchdog Behavior | 31030
- Release Information | 31030

Syntax

```
watchdog {
    timeout value;
    pat-frequency value;
}
```

Hierarchy Level

```
[edit chassis psm]
[edit chassis pem]
```

Description

Enable the power supply module (PSM) watchdog feature. Generally, Junos detects the thermal event health events and takes the required action as per user configuration. However, if a thermal health event causes Junos to go down, the PSM watchdog feature detects it and shuts down the router.

> **NOTE**: The PSM watchdog feature works only if all the online PSMs in the router support the feature.

Options

timeout *value*

Specify the watchdog timer in seconds. After the specified duration, the watchdog expires.

- **Range:** 600 through 900 seconds
- **Default:** 900 seconds

On Junos OS, timeout value is not significant since this value is programmed to watchdog reset counter. Instead, you can use the shutdown-timer option while configuring the thermal-health check feature.

pat-frequency *value*

Specify the frequency (in minutes) at which the software resets the watchdog counter. If the watchdog counter doesn't get reset because of reasons such as Routing Engine crash and system reboot, on watchdog timer expiry, the PSM turns off the output power.

- **Range:** 1 through 3 minutes
- **Default:** 1 minute

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Platform-Specific Watchdog Behavior

Use Feature Explorer to confirm platform, and release support for specific features.

Use the following table to review platform-specific behaviors for your platform:

Table 1838:

Platform	Difference
QFX10K devices	Please note that PEM firmware upgrade is required for the PSM watchdog feature. When you configure the watchdog feature on a PSM with outdated firmware, the system will raise a `Watchdog Not Supported` alarm.

Release Information

Statement introduced in Junos OS Release 20.1R1.

RELATED DOCUMENTATION

> *Handling Thermal Health Events Using Thermal Health Check and PSM Watchdog*
>
> thermal-health-check | **29305**

watchdog-timeout

IN THIS SECTION

- Syntax | 31031
- Hierarchy Level | 31031
- Description | 31031
- Options | 31031
- Required Privilege Level | 31032
- Release Information | 31032

Syntax

```
watchdog-timeout seconds;
```

Hierarchy Level

```
[edit access diameter peer peer-name]
```

Description

Configure the amount of time to wait for a Device-Watchdog-Answer message.

Options

seconds

Amount of time to wait.

- **Range:** 1 through 65,535 seconds
- **Default:** 30 seconds

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.1.

RELATED DOCUMENTATION

diameter (Subscriber Aware Policy Control) | 5699

wavelength

IN THIS SECTION

- Syntax | 31032
- Hierarchy Level | 31033
- Description | 31033
- Options | 31033
- Wavelength Options for 400G ZR and ZR+ Optics | 31037
- Required Privilege Level | 31056
- Release Information | 31056

Syntax

```
wavelength nm;
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Description

For 10-Gigabit, 100-Gigabit, 400-Gigabit ZR, and 400-Gigabit ZR+ Ethernet DWDM interfaces only, configure full C-band ITU-Grid tunable optics.

Options

nm—Wavelength value. It can be one of the following:

> **NOTE**: All values are displayed. However, if you configure a value that is not supported by the device, an error message is displayed and the device is not tuned to the specified wavelength.

- `1528.38`—1528.38 nanometers (nm), corresponds to a 50-GHz grid
- `1528.77`—1528.77 nm, corresponds to 50-GHz and 100-GHz grids
- `1529.16`—1529.16 nm, corresponds to a 50-GHz grid
- `1529.55`—1529.55 nm, corresponds to 50-GHz and 100-GHz grids
- `1529.94`—1529.94 nm, corresponds to a 50-GHz grid
- `1530.33`—1530.33 nm, corresponds to 50-GHz and 100-GHz grids
- `1530.72`—1530.72 nm, corresponds to a 50-GHz grid
- `1531.12`—1531.12 nm, corresponds to 50-GHz and 100-GHz grids
- `1531.51`—1531.51 nm, corresponds to a 50-GHz grid
- `1531.90`—1531.90 nm, corresponds to 50-GHz and 100-GHz grids
- `1532.29`—1532.29 nm, corresponds to a 50-GHz grid
- `1532.68`—1532.68 nm, corresponds to 50-GHz and 100-GHz grids
- `1533.07`—1533.07 nm, corresponds to a 50-GHz grid
- `1533.47`—1533.47 nm, corresponds to 50-GHz and 100-GHz grids

- `1533.86`—1533.86 nm, corresponds to a 50-GHz grid
- `1534.25`—1534.25 nm, corresponds to 50-GHz and 100-GHz grids
- `1534.64`—1534.64 nm, corresponds to a 50-GHz grid
- `1535.04`—1535.04 nm, corresponds to 50-GHz and 100-GHz grids
- `1535.43`—1535.43 nm, corresponds to a 50-GHz grid
- `1535.82`—1535.82 nm, corresponds to 50-GHz and 100-GHz grids
- `1536.22`—1536.22 nm, corresponds to a 50-GHz grid
- `1536.61`—1536.61 nm, corresponds to 50-GHz and 100-GHz grids
- `1537.00`—1537.00 nm, corresponds to a 50-GHz grid
- `1537.40`—1537.40 nm, corresponds to 50-GHz and 100-GHz grids
- `1537.79`—1537.79 nm, corresponds to a 50-GHz grid
- `1538.19`—1538.19 nm, corresponds to 50-GHz and 100-GHz grids
- `1538.58`—1538.58 nm, corresponds to a 50-GHz grid
- `1538.98`—1538.98 nm, corresponds to 50-GHz and 100-GHz grids
- `1539.37`—1539.37 nm, corresponds to a 50-GHz grid
- `1539.77`—1539.77 nm, corresponds to 50-GHz and 100-GHz grids
- `1540.16`—1540.16 nm, corresponds to a 50-GHz grid
- `1540.56`—1540.56 nm, corresponds to 50-GHz and 100-GHz grids
- `1540.95`—1540.95 nm, corresponds to a 50-GHz grid
- `1541.35`—1541.35 nm, corresponds to 50-GHz and 100-GHz grids
- `1541.75`—1541.75 nm, corresponds to a 50-GHz grid
- `1542.14`—1542.14 nm, corresponds to 50-GHz and 100-GHz grids
- `1542.54`—1542.54 nm, corresponds to a 50-GHz grid
- `1542.94`—1542.94 nm, corresponds to 50-GHz and 100-GHz grids
- `1543.33`—1543.33 nm, corresponds to a 50-GHz grid
- `1543.73`—1543.73 nm, corresponds to 50-GHz and 100-GHz grids

- 1544.13—1544.13 nm, corresponds to a 50-GHz grid
- 1544.53—1544.53 nm, corresponds to 50-GHz and 100-GHz grids
- 1544.92—1544.92 nm, corresponds to a 50-GHz grid
- 1545.32—1545.32 nm, corresponds to 50-GHz and 100-GHz grids
- 1545.72—1545.72 nm, corresponds to a 50-GHz grid
- 1546.12—1546.12 nm, corresponds to 50-GHz and 100-GHz grids
- 1546.52—1546.52 nm, corresponds to a 50-GHz grid
- 1546.92—1546.92 nm, corresponds to 50-GHz and 100-GHz grids
- 1547.32—1547.32 nm, corresponds to a 50-GHz grid
- 1547.72—1547.72 nm, corresponds to 50-GHz and 100-GHz grids
- 1548.11—1548.11 nm, corresponds to a 50-GHz grid
- 1548.51—1548.51 nm, corresponds to 50-GHz and 100-GHz grids
- 1548.91—1548.91 nm, corresponds to a 50-GHz grid
- 1549.32—1549.32 nm, corresponds to 50-GHz and 100-GHz grids
- 1549.72—1549.72 nm, corresponds to a 50-GHz grid
- 1550.12—1550.12 nm, corresponds to 50-GHz and 100-GHz grids
- 1550.52—1550.52 nm, corresponds to a 50-GHz grid
- 1550.92—1550.92 nm, corresponds to 50-GHz and 100-GHz grids
- 1551.32—1551.32 nm, corresponds to a 50-GHz grid
- 1551.72—1551.72 nm, corresponds to 50-GHz and 100-GHz grids
- 1552.12—1552.12 nm, corresponds to a 50-GHz grid
- 1552.52—1552.52 nm, corresponds to 50-GHz and 100-GHz grids
- 1552.93—1552.93 nm, corresponds to a 50-GHz grid
- 1553.33—1554.33 nm, corresponds to 50-GHz and 100-GHz grids
- 1553.73—1554.73 nm, corresponds to a 50-GHz grid
- 1554.13—1554.13 nm, corresponds to 50-GHz and 100-GHz grids

- `1554.54`—1554.54 nm, corresponds to a 50-GHz grid
- `1554.94`—1554.94 nm, corresponds to 50-GHz and 100-GHz grids
- `1555.34`—1555.34 nm, corresponds to a 50-GHz grid
- `1555.75`—1555.75 nm, corresponds to 50-GHz and 100-GHz grids
- `1556.15`—1556.15 nm, corresponds to a 50-GHz grid
- `1556.55`—1556.55 nm, corresponds to 50-GHz and 100-GHz grids
- `1556.96`—1556.96 nm, corresponds to a 50-GHz grid
- `1557.36`—1557.36 nm, corresponds to 50-GHz and 100-GHz grids
- `1557.77`—1557.77 nm, corresponds to a 50-GHz grid
- `1558.17`—1558.17 nm, corresponds to 50-GHz and 100-GHz grids
- `1558.58`—1558.58 nm, corresponds to a 50-GHz grid
- `1558.98`—1558.98 nm, corresponds to 50-GHz and 100-GHz grids
- `1559.39`—1559.39 nm, corresponds to a 50-GHz grid
- `1559.79`—1559.79 nm, corresponds to 50-GHz and 100-GHz grids
- `1560.20`—1560.20 nm, corresponds to a 50-GHz grid
- `1560.61`—1560.61 nm, corresponds to 50-GHz and 100-GHz grids
- `1561.01`—1561.01 nm, corresponds to a 50-GHz grid
- `1561.42`—1561.42 nm, corresponds to 50-GHz and 100-GHz grids
- `1561.83`—1561.83 nm, corresponds to a 50-GHz grid
- `1562.23`—1562.23 nm, corresponds to 50-GHz and 100-GHz grids
- `1562.64`—1562.64 nm, corresponds to a 50-GHz grid
- `1563.05`—1563.05 nm, corresponds to 50-GHz and 100-GHz grids
- `1563.45`—1563.45 nm, corresponds to a 50-GHz grid
- `1563.86`—1563.86 nm, corresponds to 50-GHz and 100-GHz grids
- `1564.27`—1564.27 nm, corresponds to a 50-GHz grid
- `1564.68`—1564.68 nm, corresponds to 50-GHz and 100-GHz grids

- `1565.09`—1565.09 nm, corresponds to a 50-GHz grid
- `1565.50`—1565.50 nm, corresponds to 50-GHz and 100-GHz grids
- `1565.90`—1565.90 nm, corresponds to a 50-GHz grid
- `1566.31`—1566.31 nm, corresponds to 50-GHz and 100-GHz grids
- `1566.72`—1566.72 nm, corresponds to a 50-GHz grid
- `1567.13`—1567.13 nm, corresponds to 50-GHz and 100-GHz grids
- `1567.54`—1567.54 nm, corresponds to a 50-GHz grid
- `1567.95`—1567.95 nm, corresponds to 50-GHz and 100-GHz grids
- `1568.36`—1568.36 nm, corresponds to a 50-GHz grid
- `1568.77`—1568.77 nm, corresponds to 50-GHz and 100-GHz grids
- **Default:** `1550.12`—1550.12 nm, corresponds to 50-GHz and 100-GHz grids

Wavelength Options for 400G ZR and ZR+ Optics

Supported grid spacing are 100GHz, 75GHz and 6.25GHz currently. The grid selection is based on 100GHz>75GHz>6.25GHZ.

Example 1:

1528.77 1528.77 nm/196.10000 THz 6.25/75/100GHz spacing

1528.77 is part of the 100GHZ, 75GHz, and 6.25GHz grids. So, this wavelength is initialized in 100GHz grid.

Example 2:

1530.53 1530.53 nm/195.87500 THz 6.25/12.5/75GHz spacing

1530.53 is part of the 75GHz and 6.25GHz grids. So, this wavelength is initialized in 75GHz grid.

The wavelength can be one of the following:

```
1528.38              1528.38 nm/196.15040 THz 50Ghz spacing
1528.77              1528.77 nm/196.10000 THz 6.25/12.5/50/75/100GHz spacing
1528.82              1528.82 nm/196.09375 THz 6.25GHz spacing
1528.87              1528.87 nm/196.08750 THz 6.25/12.5GHz spacing
1528.92              1528.92 nm/196.08125 THz 6.25GHz spacing
```

```
1528.97             1528.97 nm/196.07500 THz 6.25/12.5GHz spacing
1529.02             1529.02 nm/196.06875 THz 6.25GHz spacing
1529.07             1529.07 nm/196.06250 THz 6.25/12.5GHz spacing
1529.11             1529.11 nm/196.05625 THz 6.25GHz spacing
1529.16             1529.16 nm/196.05000 THz 6.25/12.5/50GHz spacing
1529.21             1529.21 nm/196.04375 THz 6.25GHz spacing
1529.26             1529.26 nm/196.03750 THz 6.25/12.5GHz spacing
1529.31             1529.31 nm/196.03125 THz 6.25GHz spacing
1529.36             1529.36 nm/196.02500 THz 6.25/12.5/75GHz spacing
1529.41             1529.41 nm/196.01875 THz 6.25GHz spacing
1529.46             1529.46 nm/196.01250 THz 6.25/12.5GHz spacing
1529.50             1529.50 nm/196.00625 THz 6.25GHz spacing
1529.55             1529.55 nm/196.00000 THz 6.25/12.5/50/100GHz spacing
1529.60             1529.60 nm/195.99375 THz 6.25GHz spacing
1529.65             1529.65 nm/195.98750 THz 6.25/12.5GHz spacing
1529.70             1529.70 nm/195.98125 THz 6.25GHz spacing
1529.75             1529.75 nm/195.97500 THz 6.25/12.5GHz spacing
1529.80             1529.80 nm/195.96875 THz 6.25GHz spacing
1529.85             1529.85 nm/195.96250 THz 6.25/12.5GHz spacing
1529.89             1529.89 nm/195.95625 THz 6.25GHz spacing
1529.94             1529.94 nm/195.95000 THz 6.25/12.5/50/75GHz spacing
1529.99             1529.99 nm/195.94375 THz 6.25GHz spacing
1530.04             1530.04 nm/195.93750 THz 6.25/12.5GHz spacing
1530.09             1530.09 nm/195.93125 THz 6.25GHz spacing
1530.14             1530.14 nm/195.92500 THz 6.25/12.5GHz spacing
1530.19             1530.19 nm/195.91875 THz 6.25GHz spacing
1530.24             1530.24 nm/195.91250 THz 6.25/12.5GHz spacing
1530.29             1530.29 nm/195.90625 THz 6.25GHz spacing
1530.33             1530.33 nm/195.90000 THz 6.25/12.5/50/100GHz spacing
1530.38             1530.38 nm/195.89375 THz 6.25GHz spacing
1530.43             1530.43 nm/195.88750 THz 6.25/12.5GHz spacing
1530.48             1530.48 nm/195.88125 THz 6.25GHz spacing
1530.53             1530.53 nm/195.87500 THz 6.25/12.5/75GHz spacing
1530.58             1530.58 nm/195.86875 THz 6.25GHz spacing
1530.63             1530.63 nm/195.86250 THz 6.25/12.5GHz spacing
1530.68             1530.68 nm/195.85625 THz 6.25GHz spacing
1530.72             1530.72 nm/195.85000 THz 6.25/12.5/50GHz spacing
1530.77             1530.77 nm/195.84375 THz 6.25GHz spacing
1530.82             1530.82 nm/195.83750 THz 6.25/12.5GHz spacing
1530.87             1530.87 nm/195.83125 THz 6.25GHz spacing
1530.92             1530.92 nm/195.82500 THz 6.25/12.5GHz spacing
1530.97             1530.97 nm/195.81875 THz 6.25GHz spacing
1531.02             1531.02 nm/195.81250 THz 6.25/12.5GHz spacing
```

```
1531.07             1531.07 nm/195.80625 THz 6.25GHz spacing
1531.12             1531.12 nm/195.80000 THz 6.25/12.5/50/75/100GHz spacing
1531.16             1531.16 nm/195.79375 THz 6.25GHz spacing
1531.21             1531.21 nm/195.78750 THz 6.25/12.5GHz spacing
1531.26             1531.26 nm/195.78125 THz 6.25GHz spacing
1531.31             1531.31 nm/195.77500 THz 6.25/12.5GHz spacing
1531.36             1531.36 nm/195.76875 THz 6.25GHz spacing
1531.41             1531.41 nm/195.76250 THz 6.25/12.5GHz spacing
1531.46             1531.46 nm/195.75625 THz 6.25GHz spacing
1531.51             1531.51 nm/195.75000 THz 6.25/12.5/50GHz spacing
1531.56             1531.56 nm/195.74375 THz 6.25GHz spacing
1531.60             1531.60 nm/195.73750 THz 6.25/12.5GHz spacing
1531.65             1531.65 nm/195.73125 THz 6.25GHz spacing
1531.70             1531.70 nm/195.72500 THz 6.25/12.5/75GHz spacing
1531.75             1531.75 nm/195.71875 THz 6.25GHz spacing
1531.80             1531.80 nm/195.71250 THz 6.25/12.5GHz spacing
1531.85             1531.85 nm/195.70625 THz 6.25GHz spacing
1531.90             1531.90 nm/195.70000 THz 6.25/12.5/50/100GHz spacing
1531.95             1531.95 nm/195.69375 THz 6.25GHz spacing
1532.00             1532.00 nm/195.68750 THz 6.25/12.5GHz spacing
1532.04             1532.04 nm/195.68125 THz 6.25GHz spacing
1532.09             1532.09 nm/195.67500 THz 6.25/12.5GHz spacing
1532.14             1532.14 nm/195.66875 THz 6.25GHz spacing
1532.19             1532.19 nm/195.66250 THz 6.25/12.5GHz spacing
1532.24             1532.24 nm/195.65625 THz 6.25GHz spacing
1532.29             1532.29 nm/195.65000 THz 6.25/12.5/50/75GHz spacing
1532.34             1532.34 nm/195.64375 THz 6.25GHz spacing
1532.39             1532.39 nm/195.63750 THz 6.25/12.5GHz spacing
1532.44             1532.44 nm/195.63125 THz 6.25GHz spacing
1532.49             1532.49 nm/195.62500 THz 6.25/12.5GHz spacing
1532.53             1532.53 nm/195.61875 THz 6.25GHz spacing
1532.58             1532.58 nm/195.61250 THz 6.25/12.5GHz spacing
1532.63             1532.63 nm/195.60625 THz 6.25GHz spacing
1532.68             1532.68 nm/195.60000 THz 6.25/12.5/50/100GHz spacing
1532.73             1532.73 nm/195.59375 THz 6.25GHz spacing
1532.78             1532.78 nm/195.58750 THz 6.25/12.5GHz spacing
1532.83             1532.83 nm/195.58125 THz 6.25GHz spacing
1532.88             1532.88 nm/195.57500 THz 6.25/12.5/75GHz spacing
1532.93             1532.93 nm/195.56875 THz 6.25GHz spacing
1532.98             1532.98 nm/195.56250 THz 6.25/12.5GHz spacing
1533.02             1533.02 nm/195.55625 THz 6.25GHz spacing
1533.07             1533.07 nm/195.55000 THz 6.25/12.5/50GHz spacing
1533.12             1533.12 nm/195.54375 THz 6.25GHz spacing
```

```
1533.17             1533.17 nm/195.53750 THz 6.25/12.5GHz spacing
1533.22             1533.22 nm/195.53125 THz 6.25GHz spacing
1533.27             1533.27 nm/195.52500 THz 6.25/12.5GHz spacing
1533.32             1533.32 nm/195.51875 THz 6.25GHz spacing
1533.37             1533.37 nm/195.51250 THz 6.25/12.5GHz spacing
1533.42             1533.42 nm/195.50625 THz 6.25GHz spacing
1533.47             1533.47 nm/195.50000 THz 6.25/12.5/50/75/100GHz spacing
1533.51             1533.51 nm/195.49375 THz 6.25GHz spacing
1533.56             1533.56 nm/195.48750 THz 6.25/12.5GHz spacing
1533.61             1533.61 nm/195.48125 THz 6.25GHz spacing
1533.66             1533.66 nm/195.47500 THz 6.25/12.5GHz spacing
1533.71             1533.71 nm/195.46875 THz 6.25GHz spacing
1533.76             1533.76 nm/195.46250 THz 6.25/12.5GHz spacing
1533.81             1533.81 nm/195.45625 THz 6.25GHz spacing
1533.86             1533.86 nm/195.45000 THz 6.25/12.5/50GHz spacing
1533.91             1533.91 nm/195.44375 THz 6.25GHz spacing
1533.96             1533.96 nm/195.43750 THz 6.25/12.5GHz spacing
1534.00             1534.00 nm/195.43125 THz 6.25GHz spacing
1534.05             1534.05 nm/195.42500 THz 6.25/12.5/75GHz spacing
1534.10             1534.10 nm/195.41875 THz 6.25GHz spacing
1534.15             1534.15 nm/195.41250 THz 6.25/12.5GHz spacing
1534.20             1534.20 nm/195.40625 THz 6.25GHz spacing
1534.25             1534.25 nm/195.40000 THz 6.25/12.5/50/100GHz spacing
1534.30             1534.30 nm/195.39375 THz 6.25GHz spacing
1534.35             1534.35 nm/195.38750 THz 6.25/12.5GHz spacing
1534.40             1534.40 nm/195.38125 THz 6.25GHz spacing
1534.45             1534.45 nm/195.37500 THz 6.25/12.5GHz spacing
1534.50             1534.50 nm/195.36875 THz 6.25GHz spacing
1534.54             1534.54 nm/195.36250 THz 6.25/12.5GHz spacing
1534.59             1534.59 nm/195.35625 THz 6.25GHz spacing
1534.64             1534.64 nm/195.35000 THz 6.25/12.5/50/75GHz spacing
1534.69             1534.69 nm/195.34375 THz 6.25GHz spacing
1534.74             1534.74 nm/195.33750 THz 6.25/12.5GHz spacing
1534.79             1534.79 nm/195.33125 THz 6.25GHz spacing
1534.84             1534.84 nm/195.32500 THz 6.25/12.5GHz spacing
1534.89             1534.89 nm/195.31875 THz 6.25GHz spacing
1534.94             1534.94 nm/195.31250 THz 6.25/12.5GHz spacing
1534.99             1534.99 nm/195.30625 THz 6.25GHz spacing
1535.04             1535.04 nm/195.30000 THz 6.25/12.5/50/100GHz spacing
1535.08             1535.08 nm/195.29375 THz 6.25GHz spacing
1535.13             1535.13 nm/195.28750 THz 6.25/12.5GHz spacing
1535.18             1535.18 nm/195.28125 THz 6.25GHz spacing
1535.23             1535.23 nm/195.27500 THz 6.25/12.5/75GHz spacing
```

```
1535.28             1535.28 nm/195.26875 THz 6.25GHz spacing
1535.33             1535.33 nm/195.26250 THz 6.25/12.5GHz spacing
1535.38             1535.38 nm/195.25625 THz 6.25GHz spacing
1535.43             1535.43 nm/195.25000 THz 6.25/12.5/50GHz spacing
1535.48             1535.48 nm/195.24375 THz 6.25GHz spacing
1535.53             1535.53 nm/195.23750 THz 6.25/12.5GHz spacing
1535.58             1535.58 nm/195.23125 THz 6.25GHz spacing
1535.63             1535.63 nm/195.22500 THz 6.25/12.5GHz spacing
1535.67             1535.67 nm/195.21875 THz 6.25GHz spacing
1535.72             1535.72 nm/195.21250 THz 6.25/12.5GHz spacing
1535.77             1535.77 nm/195.20625 THz 6.25GHz spacing
1535.82             1535.82 nm/195.20000 THz 6.25/12.5/50/75/100GHz spacing
1535.87             1535.87 nm/195.19375 THz 6.25GHz spacing
1535.92             1535.92 nm/195.18750 THz 6.25/12.5GHz spacing
1535.97             1535.97 nm/195.18125 THz 6.25GHz spacing
1536.02             1536.02 nm/195.17500 THz 6.25/12.5GHz spacing
1536.07             1536.07 nm/195.16875 THz 6.25GHz spacing
1536.12             1536.12 nm/195.16250 THz 6.25/12.5GHz spacing
1536.17             1536.17 nm/195.15625 THz 6.25GHz spacing
1536.22             1536.22 nm/195.15000 THz 6.25/12.5/50GHz spacing
1536.26             1536.26 nm/195.14375 THz 6.25GHz spacing
1536.31             1536.31 nm/195.13750 THz 6.25/12.5GHz spacing
1536.36             1536.36 nm/195.13125 THz 6.25GHz spacing
1536.41             1536.41 nm/195.12500 THz 6.25/12.5/75GHz spacing
1536.46             1536.46 nm/195.11875 THz 6.25GHz spacing
1536.51             1536.51 nm/195.11250 THz 6.25/12.5GHz spacing
1536.56             1536.56 nm/195.10625 THz 6.25GHz spacing
1536.61             1536.61 nm/195.10000 THz 6.25/12.5/50/100GHz spacing
1536.66             1536.66 nm/195.09375 THz 6.25GHz spacing
1536.71             1536.71 nm/195.08750 THz 6.25/12.5GHz spacing
1536.76             1536.76 nm/195.08125 THz 6.25GHz spacing
1536.81             1536.81 nm/195.07500 THz 6.25/12.5GHz spacing
1536.86             1536.86 nm/195.06875 THz 6.25GHz spacing
1536.90             1536.90 nm/195.06250 THz 6.25/12.5GHz spacing
1536.95             1536.95 nm/195.05625 THz 6.25GHz spacing
1537.00             1537.00 nm/195.05000 THz 6.25/12.5/50/75GHz spacing
1537.05             1537.05 nm/195.04375 THz 6.25GHz spacing
1537.10             1537.10 nm/195.03750 THz 6.25/12.5GHz spacing
1537.15             1537.15 nm/195.03125 THz 6.25GHz spacing
1537.20             1537.20 nm/195.02500 THz 6.25/12.5GHz spacing
1537.25             1537.25 nm/195.01875 THz 6.25GHz spacing
1537.30             1537.30 nm/195.01250 THz 6.25/12.5GHz spacing
1537.35             1537.35 nm/195.00625 THz 6.25GHz spacing
```

```
1537.40              1537.40 nm/195.00000 THz 6.25/12.5/50/100GHz spacing
1537.45              1537.45 nm/194.99375 THz 6.25GHz spacing
1537.50              1537.50 nm/194.98750 THz 6.25/12.5GHz spacing
1537.55              1537.55 nm/194.98125 THz 6.25GHz spacing
1537.59              1537.59 nm/194.97500 THz 6.25/12.5/75GHz spacing
1537.64              1537.64 nm/194.96875 THz 6.25GHz spacing
1537.69              1537.69 nm/194.96250 THz 6.25/12.5GHz spacing
1537.74              1537.74 nm/194.95625 THz 6.25GHz spacing
1537.79              1537.79 nm/194.95000 THz 6.25/12.5/50GHz spacing
1537.84              1537.84 nm/194.94375 THz 6.25GHz spacing
1537.89              1537.89 nm/194.93750 THz 6.25/12.5GHz spacing
1537.94              1537.94 nm/194.93125 THz 6.25GHz spacing
1537.99              1537.99 nm/194.92500 THz 6.25/12.5GHz spacing
1538.04              1538.04 nm/194.91875 THz 6.25GHz spacing
1538.09              1538.09 nm/194.91250 THz 6.25/12.5GHz spacing
1538.14              1538.14 nm/194.90625 THz 6.25GHz spacing
1538.19              1538.19 nm/194.90000 THz 6.25/12.5/50/75/100GHz spacing
1538.24              1538.24 nm/194.89375 THz 6.25GHz spacing
1538.28              1538.28 nm/194.88750 THz 6.25/12.5GHz spacing
1538.33              1538.33 nm/194.88125 THz 6.25GHz spacing
1538.38              1538.38 nm/194.87500 THz 6.25/12.5/75GHz spacing
1538.43              1538.43 nm/194.86875 THz 6.25GHz spacing
1538.48              1538.48 nm/194.86250 THz 6.25/12.5GHz spacing
1538.53              1538.53 nm/194.85625 THz 6.25GHz spacing
1538.58              1538.58 nm/194.85000 THz 6.25/12.5/50GHz spacing
1538.63              1538.63 nm/194.84375 THz 6.25GHz spacing
1538.68              1538.68 nm/194.83750 THz 6.25/12.5GHz spacing
1538.73              1538.73 nm/194.83125 THz 6.25GHz spacing
1538.78              1538.78 nm/194.82500 THz 6.25/12.5/75GHz spacing
1538.83              1538.83 nm/194.81875 THz 6.25GHz spacing
1538.88              1538.88 nm/194.81250 THz 6.25/12.5GHz spacing
1538.93              1538.93 nm/194.80625 THz 6.25GHz spacing
1538.98              1538.98 nm/194.80000 THz 6.25/12.5/50/100GHz spacing
1539.03              1539.03 nm/194.79375 THz 6.25GHz spacing
1539.07              1539.07 nm/194.78750 THz 6.25/12.5GHz spacing
1539.12              1539.12 nm/194.78125 THz 6.25GHz spacing
1539.17              1539.17 nm/194.77500 THz 6.25/12.5GHz spacing
1539.22              1539.22 nm/194.76875 THz 6.25GHz spacing
1539.27              1539.27 nm/194.76250 THz 6.25/12.5GHz spacing
1539.32              1539.32 nm/194.75625 THz 6.25GHz spacing
1539.37              1539.37 nm/194.75000 THz 6.25/12.5/50/75GHz spacing
1539.42              1539.42 nm/194.74375 THz 6.25GHz spacing
1539.47              1539.47 nm/194.73750 THz 6.25/12.5GHz spacing
```

```
1539.52             1539.52 nm/194.73125 THz 6.25GHz spacing
1539.57             1539.57 nm/194.72500 THz 6.25/12.5GHz spacing
1539.62             1539.62 nm/194.71875 THz 6.25GHz spacing
1539.67             1539.67 nm/194.71250 THz 6.25/12.5GHz spacing
1539.72             1539.72 nm/194.70625 THz 6.25GHz spacing
1539.77             1539.77 nm/194.70000 THz 6.25/12.5/50/100GHz spacing
1539.82             1539.82 nm/194.69375 THz 6.25GHz spacing
1539.86             1539.86 nm/194.68750 THz 6.25/12.5GHz spacing
1539.91             1539.91 nm/194.68125 THz 6.25GHz spacing
1539.96             1539.96 nm/194.67500 THz 6.25/12.5/50/75GHz spacing
1540.01             1540.01 nm/194.66875 THz 6.25GHz spacing
1540.06             1540.06 nm/194.66250 THz 6.25/12.5GHz spacing
1540.11             1540.11 nm/194.65625 THz 6.25GHz spacing
1540.16             1540.16 nm/194.65000 THz 6.25/12.5/50GHz spacing
1540.21             1540.21 nm/194.64375 THz 6.25GHz spacing
1540.26             1540.26 nm/194.63750 THz 6.25/12.5GHz spacing
1540.31             1540.31 nm/194.63125 THz 6.25GHz spacing
1540.36             1540.36 nm/194.62500 THz 6.25/12.5GHz spacing
1540.41             1540.41 nm/194.61875 THz 6.25GHz spacing
1540.46             1540.46 nm/194.61250 THz 6.25/12.5GHz spacing
1540.51             1540.51 nm/194.60625 THz 6.25GHz spacing
1540.56             1540.56 nm/194.60000 THz 6.25/12.5/50/75/100GHz spacing
1540.61             1540.61 nm/194.59375 THz 6.25GHz spacing
1540.66             1540.66 nm/194.58750 THz 6.25/12.5GHz spacing
1540.71             1540.71 nm/194.58125 THz 6.25GHz spacing
1540.76             1540.76 nm/194.57500 THz 6.25/12.5GHz spacing
1540.80             1540.80 nm/194.56875 THz 6.25GHz spacing
1540.85             1540.85 nm/194.56250 THz 6.25/12.5GHz spacing
1540.90             1540.90 nm/194.55625 THz 6.25GHz spacing
1540.95             1540.95 nm/194.55000 THz 6.25/12.5/50GHz spacing
1541.00             1541.00 nm/194.54375 THz 6.25GHz spacing
1541.05             1541.05 nm/194.53750 THz 6.25/12.5GHz spacing
1541.10             1541.10 nm/194.53125 THz 6.25GHz spacing
1541.15             1541.15 nm/194.52500 THz 6.25/12.5/75GHz spacing
1541.20             1541.20 nm/194.51875 THz 6.25GHz spacing
1541.25             1541.25 nm/194.51250 THz 6.25/12.5GHz spacing
1541.30             1541.30 nm/194.50625 THz 6.25GHz spacing
1541.35             1541.35 nm/194.50000 THz 6.25/12.5/50/100GHz spacing
1541.40             1541.40 nm/194.49375 THz 6.25GHz spacing
1541.45             1541.45 nm/194.48750 THz 6.25/12.5GHz spacing
1541.50             1541.50 nm/194.48125 THz 6.25GHz spacing
1541.55             1541.55 nm/194.47500 THz 6.25/12.5GHz spacing
1541.60             1541.60 nm/194.46875 THz 6.25GHz spacing
```

```
1541.65             1541.65 nm/194.46250 THz 6.25/12.5GHz spacing
1541.70             1541.70 nm/194.45625 THz 6.25GHz spacing
1541.75             1541.75 nm/194.45000 THz 6.25/12.5/50/75GHz spacing
1541.80             1541.80 nm/194.44375 THz 6.25GHz spacing
1541.84             1541.84 nm/194.43750 THz 6.25/12.5GHz spacing
1541.89             1541.89 nm/194.43125 THz 6.25GHz spacing
1541.94             1541.94 nm/194.42500 THz 6.25/12.5GHz spacing
1541.99             1541.99 nm/194.41875 THz 6.25GHz spacing
1542.04             1542.04 nm/194.41250 THz 6.25/12.5GHz spacing
1542.09             1542.09 nm/194.40625 THz 6.25GHz spacing
1542.14             1542.14 nm/194.40000 THz 6.25/12.5/50/100GHz spacing
1542.19             1542.19 nm/194.39375 THz 6.25GHz spacing
1542.24             1542.24 nm/194.38750 THz 6.25/12.5GHz spacing
1542.29             1542.29 nm/194.38125 THz 6.25GHz spacing
1542.34             1542.34 nm/194.37500 THz 6.25/12.5/75GHz spacing
1542.39             1542.39 nm/194.36875 THz 6.25GHz spacing
1542.44             1542.44 nm/194.36250 THz 6.25/12.5GHz spacing
1542.49             1542.49 nm/194.35625 THz 6.25GHz spacing
1542.54             1542.54 nm/194.35000 THz 6.25/12.5/50GHz spacing
1542.59             1542.59 nm/194.34375 THz 6.25GHz spacing
1542.64             1542.64 nm/194.33750 THz 6.25/12.5GHz spacing
1542.69             1542.69 nm/194.33125 THz 6.25GHz spacing
1542.74             1542.74 nm/194.32500 THz 6.25/12.5GHz spacing
1542.79             1542.79 nm/194.31875 THz 6.25GHz spacing
1542.84             1542.84 nm/194.31250 THz 6.25/12.5GHz spacing
1542.89             1542.89 nm/194.30625 THz 6.25GHz spacing
1542.94             1542.94 nm/194.30000 THz 6.25/12.5/50/75/100GHz spacing
1542.99             1542.99 nm/194.29375 THz 6.25GHz spacing
1543.04             1543.04 nm/194.28750 THz 6.25/12.5GHz spacing
1543.08             1543.08 nm/194.28125 THz 6.25GHz spacing
1543.13             1543.13 nm/194.27500 THz 6.25/12.5GHz spacing
1543.18             1543.18 nm/194.26875 THz 6.25GHz spacing
1543.23             1543.23 nm/194.26250 THz 6.25/12.5GHz spacing
1543.28             1543.28 nm/194.25625 THz 6.25GHz spacing
1543.33             1543.33 nm/194.25000 THz 6.25/12.5/50GHz spacing
1543.38             1543.38 nm/194.24375 THz 6.25GHz spacing
1543.43             1543.43 nm/194.23750 THz 6.25/12.5GHz spacing
1543.48             1543.48 nm/194.23125 THz 6.25GHz spacing
1543.53             1543.53 nm/194.22500 THz 6.25/12.5/75GHz spacing
1543.58             1543.58 nm/194.21875 THz 6.25GHz spacing
1543.63             1543.63 nm/194.21250 THz 6.25/12.5GHz spacing
1543.68             1543.68 nm/194.20625 THz 6.25GHz spacing
1543.73             1543.73 nm/194.20000 THz 6.25/12.5/50/100GHz spacing
```

```
1543.78             1543.78 nm/194.19375 THz 6.25GHz spacing
1543.83             1543.83 nm/194.18750 THz 6.25/12.5GHz spacing
1543.88             1543.88 nm/194.18125 THz 6.25GHz spacing
1543.93             1543.93 nm/194.17500 THz 6.25/12.5GHz spacing
1543.98             1543.98 nm/194.16875 THz 6.25GHz spacing
1544.03             1544.03 nm/194.16250 THz 6.25/12.5GHz spacing
1544.08             1544.08 nm/194.15625 THz 6.25GHz spacing
1544.13             1544.13 nm/194.15000 THz 6.25/12.5/50/75GHz spacing
1544.18             1544.18 nm/194.14375 THz 6.25GHz spacing
1544.23             1544.23 nm/194.13750 THz 6.25/12.5GHz spacing
1544.28             1544.28 nm/194.13125 THz 6.25GHz spacing
1544.33             1544.33 nm/194.12500 THz 6.25/12.5GHz spacing
1544.38             1544.38 nm/194.11875 THz 6.25GHz spacing
1544.43             1544.43 nm/194.11250 THz 6.25/12.5GHz spacing
1544.48             1544.48 nm/194.10625 THz 6.25GHz spacing
1544.53             1544.53 nm/194.10000 THz 6.25/12.5/50/100GHz spacing
1544.58             1544.58 nm/194.09375 THz 6.25GHz spacing
1544.63             1544.63 nm/194.08750 THz 6.25/12.5GHz spacing
1544.68             1544.68 nm/194.08125 THz 6.25GHz spacing
1544.72             1544.72 nm/194.07500 THz 6.25/12.5/75GHz spacing
1544.77             1544.77 nm/194.06875 THz 6.25GHz spacing
1544.82             1544.82 nm/194.06250 THz 6.25/12.5GHz spacing
1544.87             1544.87 nm/194.05625 THz 6.25GHz spacing
1544.92             1544.92 nm/194.05000 THz 6.25/12.5/50GHz spacing
1544.97             1544.97 nm/194.04375 THz 6.25GHz spacing
1545.02             1545.02 nm/194.03750 THz 6.25/12.5GHz spacing
1545.07             1545.07 nm/194.03125 THz 6.25GHz spacing
1545.12             1545.12 nm/194.02500 THz 6.25/12.5GHz spacing
1545.17             1545.17 nm/194.01875 THz 6.25GHz spacing
1545.22             1545.22 nm/194.01250 THz 6.25/12.5GHz spacing
1545.27             1545.27 nm/194.00625 THz 6.25GHz spacing
1545.32             1545.32 nm/194.00000 THz 6.25/12.5/50/75/100GHz spacing
1545.37             1545.37 nm/193.99375 THz 6.25GHz spacing
1545.42             1545.42 nm/193.98750 THz 6.25/12.5GHz spacing
1545.47             1545.47 nm/193.98125 THz 6.25GHz spacing
1545.52             1545.52 nm/193.97500 THz 6.25/12.5GHz spacing
1545.57             1545.57 nm/193.96875 THz 6.25GHz spacing
1545.62             1545.62 nm/193.96250 THz 6.25/12.5GHz spacing
1545.67             1545.67 nm/193.95625 THz 6.25GHz spacing
1545.72             1545.72 nm/193.95000 THz 6.25/12.5/50GHz spacing
1545.77             1545.77 nm/193.94375 THz 6.25GHz spacing
1545.82             1545.82 nm/193.93750 THz 6.25/12.5GHz spacing
1545.87             1545.87 nm/193.93125 THz 6.25GHz spacing
```

```
1545.92             1545.92 nm/193.92500 THz 6.25/12.5/75GHz spacing
1545.97             1545.97 nm/193.91875 THz 6.25GHz spacing
1546.02             1546.02 nm/193.91250 THz 6.25/12.5GHz spacing
1546.07             1546.07 nm/193.90625 THz 6.25GHz spacing
1546.12             1546.12 nm/193.90000 THz 6.25/12.5/50/100GHz spacing
1546.17             1546.17 nm/193.89375 THz 6.25GHz spacing
1546.22             1546.22 nm/193.88750 THz 6.25/12.5GHz spacing
1546.27             1546.27 nm/193.88125 THz 6.25GHz spacing
1546.32             1546.32 nm/193.87500 THz 6.25/12.5GHz spacing
1546.37             1546.37 nm/193.86875 THz 6.25GHz spacing
1546.42             1546.42 nm/193.86250 THz 6.25/12.5GHz spacing
1546.47             1546.47 nm/193.85625 THz 6.25GHz spacing
1546.52             1546.52 nm/193.85000 THz 6.25/12.5/50/75GHz spacing
1546.57             1546.57 nm/193.84375 THz 6.25GHz spacing
1546.62             1546.62 nm/193.83750 THz 6.25/12.5GHz spacing
1546.67             1546.67 nm/193.83125 THz 6.25GHz spacing
1546.72             1546.72 nm/193.82500 THz 6.25/12.5GHz spacing
1546.77             1546.77 nm/193.81875 THz 6.25GHz spacing
1546.82             1546.82 nm/193.81250 THz 6.25/12.5GHz spacing
1546.87             1546.87 nm/193.80625 THz 6.25GHz spacing
1546.92             1546.92 nm/193.80000 THz 6.25/12.5/50/100GHz spacing
1546.97             1546.97 nm/193.79375 THz 6.25GHz spacing
1547.02             1547.02 nm/193.78750 THz 6.25/12.5GHz spacing
1547.07             1547.07 nm/193.78125 THz 6.25GHz spacing
1547.12             1547.12 nm/193.77500 THz 6.25/12.5/75GHz spacing
1547.17             1547.17 nm/193.76875 THz 6.25GHz spacing
1547.22             1547.22 nm/193.76250 THz 6.25/12.5GHz spacing
1547.27             1547.27 nm/193.75625 THz 6.25GHz spacing
1547.32             1547.32 nm/193.75000 THz 6.25/12.5/50GHz spacing
1547.37             1547.37 nm/193.74375 THz 6.25GHz spacing
1547.42             1547.42 nm/193.73750 THz 6.25/12.5GHz spacing
1547.47             1547.47 nm/193.73125 THz 6.25GHz spacing
1547.52             1547.52 nm/193.72500 THz 6.25/12.5GHz spacing
1547.57             1547.57 nm/193.71875 THz 6.25GHz spacing
1547.62             1547.62 nm/193.71250 THz 6.25/12.5GHz spacing
1547.67             1547.67 nm/193.70625 THz 6.25GHz spacing
1547.72             1547.72 nm/193.70000 THz 6.25/12.5/50/75/100GHz spacing
1547.77             1547.77 nm/193.69375 THz 6.25GHz spacing
1547.82             1547.82 nm/193.68750 THz 6.25/12.5GHz spacing
1547.87             1547.87 nm/193.68125 THz 6.25GHz spacing
1547.92             1547.92 nm/193.67500 THz 6.25/12.5GHz spacing
1547.97             1547.97 nm/193.66875 THz 6.25GHz spacing
1548.02             1548.02 nm/193.66250 THz 6.25/12.5GHz spacing
```

```
1548.06             1548.06 nm/193.65625 THz 6.25GHz spacing
1548.11             1548.11 nm/193.65000 THz 6.25/12.5/50GHz spacing
1548.16             1548.16 nm/193.64375 THz 6.25GHz spacing
1548.21             1548.21 nm/193.63750 THz 6.25/12.5GHz spacing
1548.26             1548.26 nm/193.63125 THz 6.25GHz spacing
1548.31             1548.31 nm/193.62500 THz 6.25/12.5/75GHz spacing
1548.36             1548.36 nm/193.61875 THz 6.25GHz spacing
1548.41             1548.41 nm/193.61250 THz 6.25/12.5GHz spacing
1548.46             1548.46 nm/193.60625 THz 6.25GHz spacing
1548.51             1548.51 nm/193.60000 THz 6.25/12.5/50/100GHz spacing
1548.56             1548.56 nm/193.59375 THz 6.25GHz spacing
1548.61             1548.61 nm/193.58750 THz 6.25/12.5GHz spacing
1548.66             1548.66 nm/193.58125 THz 6.25GHz spacing
1548.71             1548.71 nm/193.57500 THz 6.25/12.5GHz spacing
1548.76             1548.76 nm/193.56875 THz 6.25GHz spacing
1548.81             1548.81 nm/193.56250 THz 6.25/12.5GHz spacing
1548.86             1548.86 nm/193.55625 THz 6.25GHz spacing
1548.91             1548.91 nm/193.55000 THz 6.25/12.5/50G/75Hz spacing
1548.96             1548.96 nm/193.54375 THz 6.25GHz spacing
1549.01             1549.01 nm/193.53750 THz 6.25/12.5GHz spacing
1549.06             1549.06 nm/193.53125 THz 6.25GHz spacing
1549.11             1549.11 nm/193.52500 THz 6.25/12.5GHz spacing
1549.16             1549.16 nm/193.51875 THz 6.25GHz spacing
1549.21             1549.21 nm/193.51250 THz 6.25/12.5GHz spacing
1549.26             1549.26 nm/193.50625 THz 6.25GHz spacing
1549.32             1549.32 nm/193.50000 THz 6.25/12.5/50/100GHz spacing
1549.37             1549.37 nm/193.49375 THz 6.25GHz spacing
1549.42             1549.42 nm/193.48750 THz 6.25/12.5GHz spacing
1549.47             1549.47 nm/193.48125 THz 6.25GHz spacing
1549.52             1549.52 nm/193.47500 THz 6.25/12.5/75GHz spacing
1549.57             1549.57 nm/193.46875 THz 6.25GHz spacing
1549.62             1549.62 nm/193.46250 THz 6.25/12.5GHz spacing
1549.67             1549.67 nm/193.45625 THz 6.25GHz spacing
1549.72             1549.72 nm/193.45000 THz 6.25/12.5/50GHz spacing
1549.77             1549.77 nm/193.44375 THz 6.25GHz spacing
1549.82             1549.82 nm/193.43750 THz 6.25/12.5GHz spacing
1549.87             1549.87 nm/193.43125 THz 6.25GHz spacing
1549.92             1549.92 nm/193.42500 THz 6.25/12.5GHz spacing
1549.97             1549.97 nm/193.41875 THz 6.25GHz spacing
1550.02             1550.02 nm/193.41250 THz 6.25/12.5GHz spacing
1550.07             1550.07 nm/193.40625 THz 6.25GHz spacing
1550.12             1550.12 nm/193.40000 THz 6.25/12.5/50/75/100GHz spacing
1550.17             1550.17 nm/193.39375 THz 6.25GHz spacing
```

```
1550.22             1550.22 nm/193.38750 THz 6.25/12.5GHz spacing
1550.27             1550.27 nm/193.38125 THz 6.25GHz spacing
1550.32             1550.32 nm/193.37500 THz 6.25/12.5GHz spacing
1550.37             1550.37 nm/193.36875 THz 6.25GHz spacing
1550.42             1550.42 nm/193.36250 THz 6.25/12.5GHz spacing
1550.47             1550.47 nm/193.35625 THz 6.25GHz spacing
1550.52             1550.52 nm/193.35000 THz 6.25/12.5/50GHz spacing
1550.57             1550.57 nm/193.34375 THz 6.25GHz spacing
1550.62             1550.62 nm/193.33750 THz 6.25/12.5GHz spacing
1550.67             1550.67 nm/193.33125 THz 6.25GHz spacing
1550.72             1550.72 nm/193.32500 THz 6.25/12.5/75GHz spacing
1550.77             1550.77 nm/193.31875 THz 6.25GHz spacing
1550.82             1550.82 nm/193.31250 THz 6.25/12.5GHz spacing
1550.87             1550.87 nm/193.30625 THz 6.25GHz spacing
1550.92             1550.92 nm/193.30000 THz 6.25/12.5/50/100GHz spacing
1550.97             1550.97 nm/193.29375 THz 6.25GHz spacing
1551.02             1551.02 nm/193.28750 THz 6.25/12.5GHz spacing
1551.07             1551.07 nm/193.28125 THz 6.25GHz spacing
1551.12             1551.12 nm/193.27500 THz 6.25/12.5GHz spacing
1551.17             1551.17 nm/193.26875 THz 6.25GHz spacing
1551.22             1551.22 nm/193.26250 THz 6.25/12.5GHz spacing
1551.27             1551.27 nm/193.25625 THz 6.25GHz spacing
1551.32             1551.32 nm/193.25000 THz 6.25/12.5/50/75GHz spacing
1551.37             1551.37 nm/193.24375 THz 6.25GHz spacing
1551.42             1551.42 nm/193.23750 THz 6.25/12.5GHz spacing
1551.47             1551.47 nm/193.23125 THz 6.25GHz spacing
1551.52             1551.52 nm/193.22500 THz 6.25/12.5GHz spacing
1551.57             1551.57 nm/193.21875 THz 6.25GHz spacing
1551.62             1551.62 nm/193.21250 THz 6.25/12.5GHz spacing
1551.67             1551.67 nm/193.20625 THz 6.25GHz spacing
1551.72             1551.72 nm/193.20000 THz 6.25/12.5/50/100GHz spacing
1551.77             1551.77 nm/193.19375 THz 6.25GHz spacing
1551.82             1551.82 nm/193.18750 THz 6.25/12.5GHz spacing
1551.87             1551.87 nm/193.18125 THz 6.25GHz spacing
1551.92             1551.92 nm/193.17500 THz 6.25/12.5/75GHz spacing
1551.97             1551.97 nm/193.16875 THz 6.25GHz spacing
1552.02             1552.02 nm/193.16250 THz 6.25/12.5GHz spacing
1552.07             1552.07 nm/193.15625 THz 6.25GHz spacing
1552.12             1552.12 nm/193.15000 THz 6.25/12.5/50GHz spacing
1552.17             1552.17 nm/193.14375 THz 6.25GHz spacing
1552.22             1552.22 nm/193.13750 THz 6.25/12.5GHz spacing
1552.27             1552.27 nm/193.13125 THz 6.25GHz spacing
1552.32             1552.32 nm/193.12500 THz 6.25/12.5GHz spacing
```

```
1552.37             1552.37 nm/193.11875 THz 6.25GHz spacing
1552.42             1552.42 nm/193.11250 THz 6.25/12.5GHz spacing
1552.47             1552.47 nm/193.10625 THz 6.25GHz spacing
1552.52             1552.52 nm/193.10000 THz 6.25/12.5/50/75/100GHz spacing
1552.57             1552.57 nm/193.09375 THz 6.25GHz spacing
1552.62             1552.62 nm/193.08750 THz 6.25/12.5GHz spacing
1552.68             1552.68 nm/193.08125 THz 6.25GHz spacing
1552.73             1552.73 nm/193.07500 THz 6.25/12.5GHz spacing
1552.78             1552.78 nm/193.06875 THz 6.25GHz spacing
1552.83             1552.83 nm/193.06250 THz 6.25/12.5GHz spacing
1552.88             1552.88 nm/193.05625 THz 6.25GHz spacing
1552.93             1552.93 nm/193.05000 THz 6.25/12.5/50GHz spacing
1552.98             1552.98 nm/193.04375 THz 6.25GHz spacing
1553.03             1553.03 nm/193.03750 THz 6.25/12.5GHz spacing
1553.08             1553.08 nm/193.03125 THz 6.25GHz spacing
1553.13             1553.13 nm/193.02500 THz 6.25/12.5/75GHz spacing
1553.18             1553.18 nm/193.01875 THz 6.25GHz spacing
1553.23             1553.23 nm/193.01250 THz 6.25/12.5GHz spacing
1553.28             1553.28 nm/193.00625 THz 6.25GHz spacing
1553.33             1553.33 nm/193.00000 THz 6.25/12.5/50/100GHz spacing
1553.38             1553.38 nm/192.99375 THz 6.25GHz spacing
1553.43             1553.43 nm/192.98750 THz 6.25/12.5GHz spacing
1553.48             1553.48 nm/192.98125 THz 6.25GHz spacing
1553.53             1553.53 nm/192.97500 THz 6.25/12.5GHz spacing
1553.58             1553.58 nm/192.96875 THz 6.25GHz spacing
1553.63             1553.63 nm/192.96250 THz 6.25/12.5GHz spacing
1553.68             1553.68 nm/192.95625 THz 6.25GHz spacing
1553.73             1553.73 nm/192.95000 THz 6.25/12.5/50/75GHz spacing
1553.78             1553.78 nm/192.94375 THz 6.25GHz spacing
1553.83             1553.83 nm/192.93750 THz 6.25/12.5GHz spacing
1553.88             1553.88 nm/192.93125 THz 6.25GHz spacing
1553.93             1553.93 nm/192.92500 THz 6.25/12.5GHz spacing
1553.98             1553.98 nm/192.91875 THz 6.25GHz spacing
1554.03             1554.03 nm/192.91250 THz 6.25/12.5GHz spacing
1554.08             1554.08 nm/192.90625 THz 6.25GHz spacing
1554.13             1554.13 nm/192.90000 THz 6.25/12.5/50/100GHz spacing
1554.18             1554.18 nm/192.89375 THz 6.25GHz spacing
1554.23             1554.23 nm/192.88750 THz 6.25/12.5GHz spacing
1554.29             1554.29 nm/192.88125 THz 6.25GHz spacing
1554.34             1554.34 nm/192.87500 THz 6.25/12.5/75GHz spacing
1554.39             1554.39 nm/192.86875 THz 6.25GHz spacing
1554.44             1554.44 nm/192.86250 THz 6.25/12.5GHz spacing
1554.49             1554.49 nm/192.85625 THz 6.25GHz spacing
```

```
1554.54             1554.54 nm/192.85000 THz 6.25/12.5/50GHz spacing
1554.59             1554.59 nm/192.84375 THz 6.25GHz spacing
1554.64             1554.64 nm/192.83750 THz 6.25/12.5GHz spacing
1554.69             1554.69 nm/192.83125 THz 6.25GHz spacing
1554.74             1554.74 nm/192.82500 THz 6.25/12.5GHz spacing
1554.79             1554.79 nm/192.81875 THz 6.25GHz spacing
1554.84             1554.84 nm/192.81250 THz 6.25/12.5GHz spacing
1554.89             1554.89 nm/192.80625 THz 6.25GHz spacing
1554.94             1554.94 nm/192.80000 THz 6.25/12.5/50/75/100GHz spacing
1554.99             1554.99 nm/192.79375 THz 6.25GHz spacing
1555.04             1555.04 nm/192.78750 THz 6.25/12.5GHz spacing
1555.09             1555.09 nm/192.78125 THz 6.25GHz spacing
1555.14             1555.14 nm/192.77500 THz 6.25/12.5GHz spacing
1555.19             1555.19 nm/192.76875 THz 6.25GHz spacing
1555.24             1555.24 nm/192.76250 THz 6.25/12.5GHz spacing
1555.29             1555.29 nm/192.75625 THz 6.25GHz spacing
1555.34             1555.34 nm/192.75000 THz 6.25/12.5/50GHz spacing
1555.39             1555.39 nm/192.74375 THz 6.25GHz spacing
1555.44             1555.44 nm/192.73750 THz 6.25/12.5GHz spacing
1555.49             1555.49 nm/192.73125 THz 6.25GHz spacing
1555.55             1555.55 nm/192.72500 THz 6.25/12.5/75GHz spacing
1555.60             1555.60 nm/192.71875 THz 6.25GHz spacing
1555.65             1555.65 nm/192.71250 THz 6.25/12.5GHz spacing
1555.70             1555.70 nm/192.70625 THz 6.25GHz spacing
1555.75             1555.75 nm/192.70000 THz 6.25/12.5/50/100GHz spacing
1555.80             1555.80 nm/192.69375 THz 6.25GHz spacing
1555.85             1555.85 nm/192.68750 THz 6.25/12.5GHz spacing
1555.90             1555.90 nm/192.68125 THz 6.25GHz spacing
1555.95             1555.95 nm/192.67500 THz 6.25/12.5GHz spacing
1556.00             1556.00 nm/192.66875 THz 6.25GHz spacing
1556.05             1556.05 nm/192.66250 THz 6.25/12.5GHz spacing
1556.10             1556.10 nm/192.65625 THz 6.25GHz spacing
1556.15             1556.15 nm/192.65000 THz 6.25/12.5/50/75GHz spacing
1556.20             1556.20 nm/192.64375 THz 6.25GHz spacing
1556.25             1556.25 nm/192.63750 THz 6.25/12.5GHz spacing
1556.30             1556.30 nm/192.63125 THz 6.25GHz spacing
1556.35             1556.35 nm/192.62500 THz 6.25/12.5GHz spacing
1556.40             1556.40 nm/192.61875 THz 6.25GHz spacing
1556.45             1556.45 nm/192.61250 THz 6.25/12.5GHz spacing
1556.50             1556.50 nm/192.60625 THz 6.25GHz spacing
1556.55             1556.55 nm/192.60000 THz 6.25/12.5/50/100GHz spacing
1556.61             1556.61 nm/192.59375 THz 6.25GHz spacing
1556.66             1556.66 nm/192.58750 THz 6.25/12.5GHz spacing
```

```
1556.71             1556.71 nm/192.58125 THz 6.25GHz spacing
1556.76             1556.76 nm/192.57500 THz 6.25/12.5/75GHz spacing
1556.81             1556.81 nm/192.56875 THz 6.25GHz spacing
1556.86             1556.86 nm/192.56250 THz 6.25/12.5GHz spacing
1556.91             1556.91 nm/192.55625 THz 6.25GHz spacing
1556.96             1556.96 nm/192.55000 THz 6.25/12.5/50GHz spacing
1557.01             1557.01 nm/192.54375 THz 6.25GHz spacing
1557.06             1557.06 nm/192.53750 THz 6.25/12.5GHz spacing
1557.11             1557.11 nm/192.53125 THz 6.25GHz spacing
1557.16             1557.16 nm/192.52500 THz 6.25/12.5GHz spacing
1557.21             1557.21 nm/192.51875 THz 6.25GHz spacing
1557.26             1557.26 nm/192.51250 THz 6.25/12.5GHz spacing
1557.31             1557.31 nm/192.50625 THz 6.25GHz spacing
1557.36             1557.36 nm/192.50000 THz 6.25/12.5/50/75/100GHz spacing
1557.41             1557.41 nm/192.49375 THz 6.25GHz spacing
1557.46             1557.46 nm/192.48750 THz 6.25/12.5GHz spacing
1557.52             1557.52 nm/192.48125 THz 6.25GHz spacing
1557.57             1557.57 nm/192.47500 THz 6.25/12.5GHz spacing
1557.62             1557.62 nm/192.46875 THz 6.25GHz spacing
1557.67             1557.67 nm/192.46250 THz 6.25/12.5GHz spacing
1557.72             1557.72 nm/192.45625 THz 6.25GHz spacing
1557.77             1557.77 nm/192.45000 THz 6.25/12.5/50GHz spacing
1557.82             1557.82 nm/192.44375 THz 6.25GHz spacing
1557.87             1557.87 nm/192.43750 THz 6.25/12.5GHz spacing
1557.92             1557.92 nm/192.43125 THz 6.25GHz spacing
1557.97             1557.97 nm/192.42500 THz 6.25/12.5/75GHz spacing
1558.02             1558.02 nm/192.41875 THz 6.25GHz spacing
1558.07             1558.07 nm/192.41250 THz 6.25/12.5GHz spacing
1558.12             1558.12 nm/192.40625 THz 6.25GHz spacing
1558.17             1558.17 nm/192.40000 THz 6.25/12.5/50/100GHz spacing
1558.22             1558.22 nm/192.39375 THz 6.25GHz spacing
1558.27             1558.27 nm/192.38750 THz 6.25/12.5GHz spacing
1558.32             1558.32 nm/192.38125 THz 6.25GHz spacing
1558.38             1558.38 nm/192.37500 THz 6.25/12.5GHz spacing
1558.43             1558.43 nm/192.36875 THz 6.25GHz spacing
1558.48             1558.48 nm/192.36250 THz 6.25/12.5GHz spacing
1558.53             1558.53 nm/192.35625 THz 6.25GHz spacing
1558.58             1558.58 nm/192.35000 THz 6.25/12.5/50/75GHz spacing
1558.63             1558.63 nm/192.34375 THz 6.25GHz spacing
1558.68             1558.68 nm/192.33750 THz 6.25/12.5GHz spacing
1558.73             1558.73 nm/192.33125 THz 6.25GHz spacing
1558.78             1558.78 nm/192.32500 THz 6.25/12.5GHz spacing
1558.83             1558.83 nm/192.31875 THz 6.25GHz spacing
```

```
1558.88             1558.88 nm/192.31250 THz 6.25/12.5GHz spacing
1558.93             1558.93 nm/192.30625 THz 6.25GHz spacing
1558.98             1558.98 nm/192.30000 THz 6.25/12.5/50/100GHz spacing
1559.03             1559.03 nm/192.29375 THz 6.25GHz spacing
1559.08             1559.08 nm/192.28750 THz 6.25/12.5GHz spacing
1559.14             1559.14 nm/192.28125 THz 6.25GHz spacing
1559.19             1559.19 nm/192.27500 THz 6.25/12.5/75GHz spacing
1559.24             1559.24 nm/192.26875 THz 6.25GHz spacing
1559.29             1559.29 nm/192.26250 THz 6.25/12.5GHz spacing
1559.34             1559.34 nm/192.25625 THz 6.25GHz spacing
1559.39             1559.39 nm/192.25000 THz 6.25/12.5/50GHz spacing
1559.44             1559.44 nm/192.24375 THz 6.25GHz spacing
1559.49             1559.49 nm/192.23750 THz 6.25/12.5GHz spacing
1559.54             1559.54 nm/192.23125 THz 6.25GHz spacing
1559.59             1559.59 nm/192.22500 THz 6.25/12.5GHz spacing
1559.64             1559.64 nm/192.21875 THz 6.25GHz spacing
1559.69             1559.69 nm/192.21250 THz 6.25/12.5GHz spacing
1559.74             1559.74 nm/192.20625 THz 6.25GHz spacing
1559.79             1559.79 nm/192.20000 THz 6.25/12.5/50/75/100GHz spacing
1559.84             1559.84 nm/192.19375 THz 6.25GHz spacing
1559.90             1559.90 nm/192.18750 THz 6.25/12.5GHz spacing
1559.95             1559.95 nm/192.18125 THz 6.25GHz spacing
1560.00             1560.00 nm/192.17500 THz 6.25/12.5GHz spacing
1560.05             1560.05 nm/192.16875 THz 6.25GHz spacing
1560.10             1560.10 nm/192.16250 THz 6.25/12.5GHz spacing
1560.15             1560.15 nm/192.15625 THz 6.25GHz spacing
1560.20             1560.20 nm/192.15000 THz 6.25/12.5/50GHz spacing
1560.25             1560.25 nm/192.14375 THz 6.25GHz spacing
1560.30             1560.30 nm/192.13750 THz 6.25/12.5GHz spacing
1560.35             1560.35 nm/192.13125 THz 6.25GHz spacing
1560.40             1560.40 nm/192.12500 THz 6.25/12.5/75GHz spacing
1560.45             1560.45 nm/192.11875 THz 6.25GHz spacing
1560.50             1560.50 nm/192.11250 THz 6.25/12.5GHz spacing
1560.56             1560.56 nm/192.10625 THz 6.25GHz spacing
1560.61             1560.61 nm/192.10000 THz 6.25/12.5/50/100GHz spacing
1560.66             1560.66 nm/192.09375 THz 6.25GHz spacing
1560.71             1560.71 nm/192.08750 THz 6.25/12.5GHz spacing
1560.76             1560.76 nm/192.08125 THz 6.25GHz spacing
1560.81             1560.81 nm/192.07500 THz 6.25/12.5GHz spacing
1560.86             1560.86 nm/192.06875 THz 6.25GHz spacing
1560.91             1560.91 nm/192.06250 THz 6.25/12.5GHz spacing
1560.96             1560.96 nm/192.05625 THz 6.25GHz spacing
1561.01             1561.01 nm/192.05000 THz 6.25/12.5/50G/75Hz spacing
```

```
1561.06             1561.06 nm/192.04375 THz 6.25GHz spacing
1561.11             1561.11 nm/192.03750 THz 6.25/12.5GHz spacing
1561.16             1561.16 nm/192.03125 THz 6.25GHz spacing
1561.22             1561.22 nm/192.02500 THz 6.25/12.5GHz spacing
1561.27             1561.27 nm/192.01875 THz 6.25GHz spacing
1561.32             1561.32 nm/192.01250 THz 6.25/12.5GHz spacing
1561.37             1561.37 nm/192.00625 THz 6.25GHz spacing
1561.42             1561.42 nm/192.00000 THz 6.25/12.5/50/100GHz spacing
1561.47             1561.47 nm/191.99375 THz 6.25GHz spacing
1561.52             1561.52 nm/191.98750 THz 6.25/12.5GHz spacing
1561.57             1561.57 nm/191.98125 THz 6.25GHz spacing
1561.62             1561.62 nm/191.97500 THz 6.25/12.5/75GHz spacing
1561.67             1561.67 nm/191.96875 THz 6.25GHz spacing
1561.72             1561.72 nm/191.96250 THz 6.25/12.5GHz spacing
1561.77             1561.77 nm/191.95625 THz 6.25GHz spacing
1561.83             1561.83 nm/191.95000 THz 6.25/12.5/50GHz spacing
1561.88             1561.88 nm/191.94375 THz 6.25GHz spacing
1561.93             1561.93 nm/191.93750 THz 6.25/12.5GHz spacing
1561.98             1561.98 nm/191.93125 THz 6.25GHz spacing
1562.03             1562.03 nm/191.92500 THz 6.25/12.5GHz spacing
1562.08             1562.08 nm/191.91875 THz 6.25GHz spacing
1562.13             1562.13 nm/191.91250 THz 6.25/12.5GHz spacing
1562.18             1562.18 nm/191.90625 THz 6.25GHz spacing
1562.23             1562.23 nm/191.90000 THz 6.25/12.5/50/75/100GHz spacing
1562.28             1562.28 nm/191.89375 THz 6.25GHz spacing
1562.33             1562.33 nm/191.88750 THz 6.25/12.5GHz spacing
1562.39             1562.39 nm/191.88125 THz 6.25GHz spacing
1562.44             1562.44 nm/191.87500 THz 6.25/12.5GHz spacing
1562.49             1562.49 nm/191.86875 THz 6.25GHz spacing
1562.54             1562.54 nm/191.86250 THz 6.25/12.5GHz spacing
1562.59             1562.59 nm/191.85625 THz 6.25GHz spacing
1562.64             1562.64 nm/191.85000 THz 6.25/12.5/50GHz spacing
1562.69             1562.69 nm/191.84375 THz 6.25GHz spacing
1562.74             1562.74 nm/191.83750 THz 6.25/12.5GHz spacing
1562.79             1562.79 nm/191.83125 THz 6.25GHz spacing
1562.84             1562.84 nm/191.82500 THz 6.25/12.5/75GHz spacing
1562.89             1562.89 nm/191.81875 THz 6.25GHz spacing
1562.95             1562.95 nm/191.81250 THz 6.25/12.5GHz spacing
1563.00             1563.00 nm/191.80625 THz 6.25GHz spacing
1563.05             1563.05 nm/191.80000 THz 6.25/12.5/50/100GHz spacing
1563.10             1563.10 nm/191.79375 THz 6.25GHz spacing
1563.15             1563.15 nm/191.78750 THz 6.25/12.5GHz spacing
1563.20             1563.20 nm/191.78125 THz 6.25GHz spacing
```

```
1563.25             1563.25 nm/191.77500 THz 6.25/12.5GHz spacing
1563.30             1563.30 nm/191.76875 THz 6.25GHz spacing
1563.35             1563.35 nm/191.76250 THz 6.25/12.5GHz spacing
1563.40             1563.40 nm/191.75625 THz 6.25GHz spacing
1563.45             1563.45 nm/191.75000 THz 6.25/12.5/50/75GHz spacing
1563.51             1563.51 nm/191.74375 THz 6.25GHz spacing
1563.56             1563.56 nm/191.73750 THz 6.25/12.5GHz spacing
1563.61             1563.61 nm/191.73125 THz 6.25GHz spacing
1563.66             1563.66 nm/191.72500 THz 6.25/12.5GHz spacing
1563.71             1563.71 nm/191.71875 THz 6.25GHz spacing
1563.76             1563.76 nm/191.71250 THz 6.25/12.5GHz spacing
1563.81             1563.81 nm/191.70625 THz 6.25GHz spacing
1563.86             1563.86 nm/191.70000 THz 6.25/12.5/50/100GHz spacing
1563.91             1563.91 nm/191.69375 THz 6.25GHz spacing
1563.96             1563.96 nm/191.68750 THz 6.25/12.5GHz spacing
1564.02             1564.02 nm/191.68125 THz 6.25GHz spacing
1564.07             1564.07 nm/191.67500 THz 6.25/12.5/75GHz spacing
1564.12             1564.12 nm/191.66875 THz 6.25GHz spacing
1564.17             1564.17 nm/191.66250 THz 6.25/12.5GHz spacing
1564.22             1564.22 nm/191.65625 THz 6.25GHz spacing
1564.27             1564.27 nm/191.65000 THz 6.25/12.5/50GHz spacing
1564.32             1564.32 nm/191.64375 THz 6.25GHz spacing
1564.37             1564.37 nm/191.63750 THz 6.25/12.5GHz spacing
1564.42             1564.42 nm/191.63125 THz 6.25GHz spacing
1564.47             1564.47 nm/191.62500 THz 6.25/12.5GHz spacing
1564.53             1564.53 nm/191.61875 THz 6.25GHz spacing
1564.58             1564.58 nm/191.61250 THz 6.25/12.5GHz spacing
1564.63             1564.63 nm/191.60625 THz 6.25GHz spacing
1564.68             1564.68 nm/191.60000 THz 6.25/12.5/50/75/100GHz spacing
1564.73             1564.73 nm/191.59375 THz 6.25GHz spacing
1564.78             1564.78 nm/191.58750 THz 6.25/12.5GHz spacing
1564.83             1564.83 nm/191.58125 THz 6.25GHz spacing
1564.88             1564.88 nm/191.57500 THz 6.25/12.5GHz spacing
1564.93             1564.93 nm/191.56875 THz 6.25GHz spacing
1564.99             1564.99 nm/191.56250 THz 6.25/12.5GHz spacing
1565.04             1565.04 nm/191.55625 THz 6.25GHz spacing
1565.09             1565.09 nm/191.55000 THz 6.25/12.5/50GHz spacing
1565.14             1565.14 nm/191.54375 THz 6.25GHz spacing
1565.19             1565.19 nm/191.53750 THz 6.25/12.5GHz spacing
1565.24             1565.24 nm/191.53125 THz 6.25GHz spacing
1565.29             1565.29 nm/191.52500 THz 6.25/12.5/75GHz spacing
1565.34             1565.34 nm/191.51875 THz 6.25GHz spacing
1565.39             1565.39 nm/191.51250 THz 6.25/12.5GHz spacing
```

```
1565.44          1565.44 nm/191.50625 THz 6.25GHz spacing
1565.50          1565.50 nm/191.50000 THz 6.25/12.5/50/100GHz spacing
1565.55          1565.55 nm/191.49375 THz 6.25GHz spacing
1565.60          1565.60 nm/191.48750 THz 6.25/12.5GHz spacing
1565.65          1565.65 nm/191.48125 THz 6.25GHz spacing
1565.70          1565.70 nm/191.47500 THz 6.25/12.5GHz spacing
1565.75          1565.75 nm/191.46875 THz 6.25GHz spacing
1565.80          1565.80 nm/191.46250 THz 6.25/12.5GHz spacing
1565.85          1565.85 nm/191.45625 THz 6.25GHz spacing
1565.90          1565.90 nm/191.45000 THz 6.25/12.5/50/75GHz spacing
1565.96          1565.96 nm/191.44375 THz 6.25GHz spacing
1566.01          1566.01 nm/191.43750 THz 6.25/12.5GHz spacing
1566.06          1566.06 nm/191.43125 THz 6.25GHz spacing
1566.11          1566.11 nm/191.42500 THz 6.25/12.5GHz spacing
1566.16          1566.16 nm/191.41875 THz 6.25GHz spacing
1566.21          1566.21 nm/191.41250 THz 6.25/12.5GHz spacing
1566.26          1566.26 nm/191.40625 THz 6.25GHz spacing
1566.31          1566.31 nm/191.40000 THz 6.25/12.5/50/100GHz spacing
1566.36          1566.36 nm/191.39375 THz 6.25GHz spacing
1566.42          1566.42 nm/191.38750 THz 6.25/12.5GHz spacing
1566.47          1566.47 nm/191.38125 THz 6.25GHz spacing
1566.52          1566.52 nm/191.37500 THz 6.25/12.5/75GHz spacing
1566.57          1566.57 nm/191.36875 THz 6.25GHz spacing
1566.62          1566.62 nm/191.36250 THz 6.25/12.5GHz spacing
1566.67          1566.67 nm/191.35625 THz 6.25GHz spacing
1566.72          1566.72 nm/191.35000 THz 6.25/12.5/50GHz spacing
1566.77          1566.77 nm/191.34375 THz 6.25GHz spacing
1566.83          1566.83 nm/191.33750 THz 6.25/12.5GHz spacing
1566.88          1566.88 nm/191.33125 THz 6.25GHz spacing
1566.93          1566.93 nm/191.32500 THz 6.25/12.5GHz spacing
1566.98          1566.98 nm/191.31875 THz 6.25GHz spacing
1567.03          1567.03 nm/191.31250 THz 6.25/12.5GHz spacing
1567.08          1567.08 nm/191.30625 THz 6.25GHz spacing
1567.13          1567.13 nm/191.30000 THz 6.25/12.5/50/100GHz spacing
1567.18          1567.18 nm/191.29375 THz 6.25GHz spacing
1567.23          1567.23 nm/191.28750 THz 6.25/12.5GHz spacing
1567.29          1567.29 nm/191.28125 THz 6.25GHz spacing
1567.34          1567.34 nm/191.27500 THz 6.25/12.5GHz spacing
1567.39          1567.39 nm/191.26875 THz 6.25GHz spacing
1567.44          1567.44 nm/191.26250 THz 6.25/12.5GHz spacing
1567.49          1567.49 nm/191.25625 THz 6.25GHz spacing
1567.54          1567.54 nm/191.25000 THz 6.25/12.5/50GHz spacing
1567.59          1567.59 nm/191.24375 THz 6.25GHz spacing
```

```
1567.64             1567.64 nm/191.23750 THz 6.25/12.5GHz spacing
1567.70             1567.70 nm/191.23125 THz 6.25GHz spacing
1567.75             1567.75 nm/191.22500 THz 6.25/12.5GHz spacing
1567.80             1567.80 nm/191.21875 THz 6.25GHz spacing
1567.85             1567.85 nm/191.21250 THz 6.25/12.5GHz spacing
1567.90             1567.90 nm/191.20625 THz 6.25GHz spacing
1567.95             1567.95 nm/191.20000 THz 6.25/12.5/50/100GHz spacing
1568.00             1568.00 nm/191.19375 THz 6.25GHz spacing
1568.05             1568.05 nm/191.18750 THz 6.25/12.5GHz spacing
1568.11             1568.11 nm/191.18125 THz 6.25GHz spacing
1568.16             1568.16 nm/191.17500 THz 6.25/12.5GHz spacing
1568.21             1568.21 nm/191.16875 THz 6.25GHz spacing
1568.26             1568.26 nm/191.16250 THz 6.25/12.5GHz spacing
1568.31             1568.31 nm/191.15625 THz 6.25GHz spacing
1568.36             1568.36 nm/191.15000 THz 6.25/12.5/50GHz spacing
1568.77             1568.77 nm/191.1003 THz 50/100Ghz spacing
```

Required Privilege Level

Interface—To view this statement in the configuration.

Interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port) | **20580**

web-authentication

IN THIS SECTION

- Syntax | 31057
- Hierarchy Level | 31057
- Description | 31057
- Options | 31057
- Required Privilege Level | 31058
- Release Information | 31058

Syntax

```
web-authentication {
    client-match user-or-group-name;
}
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication]
```

Description

Specify that the policy allows access to users who have previously been authenticated by Web authentication. Web authentication must be enabled on one of the addresses on the interface to which the HTTP or HTTPS request is redirected.

Options

`client-match` *`user-or-group`* —(Optional) Username or user group name.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

HTTPS for Web authentication is supported on SRX5400, SRX5600, and SRX5800 devices starting from Junos OS Release 12.1X44-D10 and on vSRX Virtual Firewall, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

Understanding User Role Firewalls

web-authentication (Access)

IN THIS SECTION

- Syntax | 31058
- Hierarchy Level | 31059
- Description | 31059
- Options | 31059
- Required Privilege Level | 31060
- Release Information | 31060

Syntax

```
web-authentication {
   banner {
       success success;
```

```
    }
    default-profile profile-name;
    timeout seconds;
}
```

Hierarchy Level

```
[edit access firewall-authentication
]
[edit logical-systems name tenants name access firewall-authentication],
[edit tenants name access firewall-authentication]
```

Description

Specify that users go through the Web authentication process. The user uses HTTP or HTTPS to access an IP address on the device that is enabled for Web authentication. In this scenario, the user does not use HTTP or HTTPS to access the IP address of the protected resource. The user is prompted for a username and password, which are verified by the device. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the results of this authentication. This method of authentication differs from pass-through authentication in that users need to access the protected resource directly after accessing the Web authentication IP address and being authenticated.

Options

banner *success*;

Configure the banner that appears to users during the Web authentication process. The banner appears during login, after successful authentication, and after failed authentication.

default-profile *profile-name*

Specify the authentication profile to use if no profile is specified in a policy.

timeout *seconds*

Specify the `timeout` option in seconds.

If you do not specify a timeout value, and if the web authentication process takes more than 3 seconds, your browser may display `invalid username and password`, even though the username and password is correct. For example, when you type a username and password in a browser for authentication, SRX Series Firewall checks your account in the database, and after 3 seconds your web browser displays a message `invalid username and password`. However, after 10 seconds, SRX Series Firewall receives a response from the database that

the user authentication is successful, but SRX Series Firewall could not notify you about successful authentication, due to 3 seconds timeout value. If you configure the timeout value from 5 through 60 seconds, then the browser waits for the SRX Series Firewall to respond for the specified time.

- **Default:** 3 seconds
- **Range:** 5 through 60 seconds

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

HTTPS for Web authentication is supported on SRX5400, SRX5600, and SRX5800 devices starting from Junos OS Release 12.1X44-D10 and on vSRX Virtual Firewall, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Option `timeout` introduced in Junos OS Release 15.1X49-D130.

RELATED DOCUMENTATION

Firewall User Authentication Overview

web-authentication (Interfaces)

IN THIS SECTION

- Syntax | **31061**
- Hierarchy Level | **31061**
- Description | **31061**
- Options | **31061**

Required Privilege Level | 31061
Release Information | 31061

Syntax

```
web-authentication {
    http;
    https;
    redirect-to-https;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family-name address]
```

Description

Enable the Web authentication process for firewall user authentication.

Options

`http`—Enable HTTP service.

`https`—Enable authentication through HTTPS.

`redirect-to-https`—Redirect Web authentication to HTTPS.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support for `https` and `redirect-to-https` introduced for SRX5400, SRX5600, and SRX5800 Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX Virtual Firewall, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

Understanding Interfaces

web-filter

IN THIS SECTION

- Syntax | 31062
- Hierarchy Level | 31064
- Description | 31064
- Required Privilege Level | 31064
- Release Information | 31064

Syntax

```
web-filter {
    profile (Web Filter) profile-name {
        dns-filter {
            database-file filename;
            dns-resp-ttl seconds;
            dns-server [ ip-address ];
            hash-key key-string;
            hash-method hash-method-name;
            statistics-log-timer minutes;
            wildcarding-level level;
        }
        dns-filter-template template-name {
            client-interfaces [ client-interface-name ];
```

```
            client-routing-instance client-routing-instance-name;
            dns-filter {
                database-file filename;
                dns-resp-ttl seconds;
                dns-server [ ip-address ];
                hash-key key-string;
                hash-method hash-method-name;
                statistics-log-timer minutes;
                wildcarding-level level;
            }
            server-interfaces  [ server-interface-name ];
            server-routing-instance server-routing-instance-name;
            term term-name {
                from {
                    src-ip-prefix [ source-prefix ];
                }
                then {
                    accept;
                    dns-sinkhole;
                }
            }
        }
        global-dns-stats-log-timer minutes;
        url-filter-database filename;
        url-filter-template template-name {
            client-interfaces [ client-interface-name1 client-interface-name2 ];
            disable-url-filtering;
            dns-resolution-interval minutes;
            dns-resolution-rate seconds;
            dns-retries number;
            dns-routing-instance dns-routing-instance-name;
            dns-server [ ip-address1 ip-address2 ip-address3 ];
            dns-source-interface loopback-interface-name;
            dns-routing-instance dns-routing-instance-name;
            routing-instance routing-instance-name;
            server-interfaces  [ server-interface-name1 server-interface-name2 ];
            term term-name {
                from {
                    src-ip-prefix [prefix1 prefix2];
                    dest-port   [port1 port2];
                }
                then {
                    accept;
```

```
                    custom-page custom-page;
                    http-status-code http-status-code;
                    redirect-url redirect-url;
                    tcp-reset;
                }
            }
            url-filter-database filename
        }
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure filtering of DNS requests for disallowed website domains. Filtering can result in either:

- Blocking access to the site by sending the client a DNS response that includes an IP address or domain name of a sinkhole server instead of the disallowed domain.
- Logging the DNS request and allowing access.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

web-filter-profile

IN THIS SECTION

- Syntax | 31065
- Hierarchy Level | 31065
- Description | 31065
- Options | 31065
- Required Privilege Level | 31066
- Release Information | 31066

Syntax

```
web-filter-profile profile-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the DNS filter profile or the URL filter profile that the service set uses. The filter profile is configured at the `[edit services web-filter]` hierarchy level, and specifies how to filter DNS requests for disallowed website domains or how to filter access to disallowed URLs.

Options

profile-name	Name of the DNS filter profile.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

RELATED DOCUMENTATION

DNS Request Filtering for Disallowed Website Domains

web-filtering

IN THIS SECTION

- Syntax | 31066
- Hierarchy Level | 31069
- Description | 31069
- Options | 31070
- Required Privilege Level | 31071
- Release Information | 31071

Syntax

```
web-filtering {
    http-persist;
    http-reassemble;
    juniper-enhanced {
```

```
        base-filter;
        block-message {
            type custom-redirect-url;
            url;
        }
        cache {
            size kilobytes;
            timeout minutes;
        }
        category name {
            action (block | log-and-permit | permit | quarantine);
            custom-message;
        }
        custom-block-message;
        default (block | log-and-permit | permit | quarantine);
        fallback-settings {
            default (block | log-and-permit);
            server-connectivity (block | log-and-permit);
            timeout (block | log-and-permit);
            too-many-requests (block | log-and-permit);
        }
        no-safe-search;
        quarantine-custom-message;
        quarantine-message {
            type custom-redirect-url;
            url;
        }
        reputation {
            reputation-fairly-safe;
            reputation-moderately-safe;
            reputation-suspicious;
            reputation-very-safe;
        }
        server {
            host;
            port;
            routing-instance;
        }
        site-reputation-action {
            fairly-safe (block | log-and-permit | permit | quarantine);
            harmful (block | log-and-permit | permit | quarantine);
            moderately-safe (block | log-and-permit | permit | quarantine);
            suspicious (block | log-and-permit | permit | quarantine);
```

```
            very-safe (block | log-and-permit | permit | quarantine);
        }
        timeout seconds;
    }
    juniper-local {
        block-message {
            type custom-redirect-url;
            url;
        }
        category name {
            action (block | log-and-permit | permit | quarantine);
            custom-message;
        }
        custom-block-message;
        default (block | log-and-permit | permit);
        fallback-settings {
            default (block | log-and-permit);
            server-connectivity (block | log-and-permit);
            timeout (block | log-and-permit);
            too-many-requests (block | log-and-permit);
        }
        quarantine-custom-message;
        quarantine-message {
            type custom-redirect-url;
            url;
        }
        timeout seconds;
    }
    performance-mode;
    traceoptions {
        flag name;
    }
    url-blacklist;
    url-whitelist;
    websense-redirect {
        account;
        block-message {
            type custom-redirect-url;
            url;
        }
        category name {
            action (block | log-and-permit | permit | quarantine);
            custom-message;
```

```
        }
        custom-block-message;
        fallback-settings {
            default (block | log-and-permit);
            server-connectivity (block | log-and-permit);
            timeout (block | log-and-permit);
            too-many-requests (block | log-and-permit);
        }
        quarantine-custom-message;
        quarantine-message {
            type custom-redirect-url;
            url;
        }
        server {
            host;
            port;
            routing-instance;
        }
        sockets;
        timeout seconds;
    }
}
```

Hierarchy Level

```
[edit security utm feature-profile]
[edit security utm default-configuration]
```

Description

Configure Content Security web filtering features. You can also configure the default Content Security configuration for web filtering feature profile. If you do not configure any option in the web filtering feature profile, the values configured in the default Content Security configuration are applied. The default Content Security Web filtering configuration for HTTP is also applicable for the HTTPS sessions. Web filtering feature's potential policies conflict check is independent of the content filtering, antivirus, and antispam features.

Options

http-persist	Check all HTTP request in a connection. If `http-persist` option is enabled for clear text HTTP traffic, then Web filtering checks every HTTP request packet in the same session.
http-reassemble	Reassemble HTTP request segments. If http-reassemble option is enabled for clear text HTTP traffic, then Enhanced Web Filtering (EWF) reassembles the fragmented HTTP request to avoid evasion instead of packet-based inspection.
juniper-enhanced	Enable enhanced Web filtering on the device.
base-filter	A base filter is an object that contains a category-action pair for all categories defined in the category file.
block-message	Juniper enhanced block message settings.
cache	Set the cache parameters for Surf-Control-Integrated Web filtering and Enhanced Web Filtering.
category	Select a custom URL category list you created (custom objects) for filtering against.
custom-block-message	Enter a custom message to be sent when HTTP requests are blocked.
default	Specify an action for the profile, for requests that experience internal errors in the Web filtering module.
fallback-settings	Fallback settings tell the system how to handle errors.
no-safe-search	Do not perform safe-search for Juniper enhanced protocol. Safe-search redirect supports HTTP only. Therefore it is not possible to generate a redirect response for HTTPS search URLs. Safe-search redirects can be disabled by using the CLI option `no-safe-search`.
quarantine-custom-message	Juniper enhanced quarantine custom message.
quarantine-message	Juniper enhanced quarantine message settings.
reputation	Customize reputation level. The ThreatSeeker Cloud (TSC) provides site reputation information. Based on these reputations, you can choose a block or a permit action.
server	Set server parameters by entering the server name or IP address.
site-reputation-action	Specify the action to be taken depending on the site reputation returned for all types of URLs whether it is categorized or uncategorized.

timeout	Enter a timeout limit for requests. Once this limit is reached, fail mode settings are applied. • **Range:** 1 through 120
juniper-local	Enable Juniper Networks local URL filtering on the device.
block-message	Juniper local block message settings.
traceoptions	Trace options for Web filtering feature.
url-blacklist	This is a global blocklist category, blocking content for Web filtering.
url-whitelist	A URL allowlist is a unique custom list that you define in which all the URLs or IP addresses in that list for a specified category are always bypassed for filtering.
websense-redirect	Web filtering websense redirect engine. Websense occasionally releases new EWF categories. EWF classifies websites into categories according to host, URL, or IP address and performs filtering based on the categories.
type	Type of Web filtering solution or URL filtering solution used by the device.
performance-mode	Improves the performance by only analyzing the traffic requests on ports 80 and 443.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

The surf-control-integrated feature is not supported from Junos OS Release 15.1X49-D10 onwards. For previous releases, statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level introduced in Junos OS Release 18.2R1.

The `performance-mode` statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Understanding Local Web Filtering

Monitoring Web Filtering Configurations

web-filtering (Logical System Security Feature Profile)

IN THIS SECTION

- Syntax | 31072
- Hierarchy Level | 31074
- Description | 31075
- Options | 31075
- Required Privilege Level | 31076
- Release Information | 31076

Syntax

```
web-filtering {
    juniper-enhanced {
        profile name {
            base-filter base-filter;
            block-message {
                type custom-redirect-url;
                url url;
            }
            category name {
                action (block | log-and-permit | permit | quarantine);
                custom-message custom-message;
            }
            custom-block-message custom-block-message;
            default (block | log-and-permit | permit | quarantine);
            fallback-settings {
                default (block | log-and-permit);
                server-connectivity (block | log-and-permit);
```

```
                timeout (block | log-and-permit);
                too-many-requests (block | log-and-permit);
            }
            no-safe-search;
            quarantine-custom-message quarantine-custom-message;
            quarantine-message {
                type custom-redirect-url;
                url url;
            }
            site-reputation-action {
                fairly-safe (block | log-and-permit | permit | quarantine);
                harmful (block | log-and-permit | permit | quarantine);
                moderately-safe (block | log-and-permit | permit | quarantine);
                suspicious (block | log-and-permit | permit | quarantine);
                very-safe (block | log-and-permit | permit | quarantine);
            }
            timeout seconds;
        }
    }
    juniper-local {
        profile name {
            block-message {
                type custom-redirect-url;
                url url;
            }
            category name {
                action (block | log-and-permit | permit | quarantine);
                custom-message custom-message;
            }
            custom-block-message custom-block-message;
            default (block | log-and-permit | permit);
            fallback-settings {
                default (block | log-and-permit);
                server-connectivity (block | log-and-permit);
                timeout (block | log-and-permit);
                too-many-requests (block | log-and-permit);
            }
            quarantine-custom-message quarantine-custom-message;
            quarantine-message {
                type custom-redirect-url;
                url url;
            }
            timeout seconds;
```

```
        }
    }
    websense-redirect {
        profile name {
            account account;
            block-message {
                type custom-redirect-url;
                url url;
            }
            category name {
                action (block | log-and-permit | permit | quarantine);
                custom-message custom-message;
            }
            custom-block-message custom-block-message;
            fallback-settings {
                default (block | log-and-permit);
                server-connectivity (block | log-and-permit);
                timeout (block | log-and-permit);
                too-many-requests (block | log-and-permit);
            }
            quarantine-custom-message quarantine-custom-message;
            quarantine-message {
                type custom-redirect-url;
                url url;
            }
            server {
                host host;
                port port;
                routing-instance routing-instance;
            }
            sockets sockets;
            timeout seconds;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-systems-name security utm feature-profile]
[edit tenants tenant-name security utm feature-profile]
```

Description

Configures the Content Security Web filtering feature for logical systems. The Web filtering allows you to manage Internet usage by preventing access to inappropriate Web content. The potential policies conflict check of the Web filtering feature is independent of the content filtering, antivirus, and antispam features. You can also configure the default Content Security configuration for Web filtering feature profile. If you do not configure any option in the Web filtering feature profile, the values configured in the default Content Security configuration are applied.

Options

juniper-enhanced	Enables Enhanced Web Filtering (EWF) on the device.
base-filter	A base filter is an object that contains a category-action pair for all categories defined in the category file.
block-message	Juniper enhanced block message settings.
category	Select a custom URL category list you created (custom objects) for filtering against.
custom-block-message	Enter a custom message to be sent when HTTP requests are blocked.
default	Specify an action for the profile, for requests that experience internal errors in the Web filtering module.
fallback-settings	Fallback settings helps the system how to handle errors.
no-safe-search	Do not perform safe-search for Juniper enhanced protocol. Safe-search redirect supports HTTP only. Therefore it is not possible to generate a redirect response for HTTPS search URLs. Safe-search redirects can be disabled by using the CLI option `no-safe-search`.
quarantine-custom-message	Juniper enhanced quarantine custom message.
quarantine-message	Juniper enhanced quarantine message settings.
server	Set server parameters by entering the server name or IP address.
site-reputation-action	Specify the action to be taken depending on the site reputation returned for all types of URLs whether it is categorized or uncategorized.
timeout	Enter a timeout limit for requests. Once this limit is reached, fail mode settings are applied.

- **Range:** 1 through 1800

`juniper-local`	Enables Juniper Networks local URL filtering on the device.
`block-message`	Juniper local block message settings.
`websense-redirect`	Web filtering websense redirect engine. Websense occasionally releases new EWF categories. EWF classifies websites into categories according to host, URL, or IP address and performs filtering based on the categories.
`type`	Type of Web filtering solution or URL filtering solution used by the device.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Unified Threat Management Overview

utm default-configuration | **30614**

web-filtering (Security Content Security Policy)

IN THIS SECTION

Syntax | **31077**

Hierarchy Level | **31077**

- Description | 31077
- Required Privilege Level | 31077
- Release Information | 31077

Syntax

```
web-filtering {
    http-profile http-profile;
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm utm-policy policy-name]
[edit logical-systems logical-systems-name security utm utm-policy policy-name]
[edit tenants tenant-name security utm utm-policy policy-name]
```

Description

Configures a Content Security policy for the Web filtering protocols and attach this policy to a security profile to implement it. Web filtering allows you to manage Internet usage by preventing access to inappropriate Web content.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support in default configuration introduced in Junos OS Release 18.2R1.

Support for configuration in logical systems introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Web Filtering Overview

web-management (System Processes)

IN THIS SECTION

- Syntax | 31078
- Hierarchy Level | 31078
- Description | 31078
- Options | 31079
- Required Privilege Level | 31079
- Release Information | 31079

Syntax

```
web-management {
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the Web management process.

Options

- `disable`—Disable the Web management process.
- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

web-management (System Services)

IN THIS SECTION

- Syntax | 31080
- Hierarchy Level | 31080
- Description | 31081
- Options | 31081
- Required Privilege Level | 31083
- Release Information | 31083

Syntax

```
web-management {
    control max-threads max-threads;
    http {
        interface [interface-names] ;
        port port;
    }
    https {
        interface [interface-names];
        ( local-certificate name | pki-local-certificate name | system-generated-certificate );
        port port;
        virtual-domain virtual-domain-name {
            pki-local-certificate pki-local-certificate-name;
          }
    }
    management-url management-url;
    session {
        idle-timeout minutes;
        session-limit number;
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (no-world-readable | world-readable);
        }
        flag flag level level;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure settings for HTTP or HTTPS access. HTTP access allows management of the device using the browser-based J-Web graphical user interface. HTTPS access allows secure management of the device using the J-Web interface. With HTTPS access, communication between the device's Web server and your browser is encrypted.

> **NOTE**: On SRX340, SRX345, and SRX380 devices, the factory-default configuration has a generic HTTP configuration. To use Gigabit Ethernet (ge) and fxp0 ports as management ports, you must use the **set system services web-management http interface** command to configure HTTP access for those interfaces. The Web management HTTP and HTTPS interfaces are changed to fxp0.0 and from ge-0/0/1.0 through ge-0/0/7.0.
>
> vSRX Virtual Firewall 3.0 on Hyper-V does not support the web management https configuration.

Options

control max-threads *max-threads*

Configure the maximum number of simultaneous threads to handle access requests.

- **Range:** 0 through 16

https

Configure the secure version of the HTTP service, HTTPS, which is encrypted.

management-url

Configure the URL path for Web management access.

traceoptions

Set the trace options.

- `file`—Configure the trace file information.
  - *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - `files` *number*— Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the `size` *`maximum file-size`* option.

Range: 2 through 1000 files

Default: 10 files

- `match` *`regular-expression`*—Refine the output to include lines that contain the regular expression.

- `size` *`maximum-file-size`*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

  Range: 10 KB through 1 GB

  Default: 128 KB

  If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` *`number`* option.

- `(world-readable | no-world-readable)`— By default, log files can be accessed only by the user who configures the tracing operation. The `world-readable` option enables any user to read the file. To explicitly set the default behavior, use the `no-world-readable` option.

- `flag` *`flag`*—Specify which tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags.

  - `all`—Trace all areas.

  - `configuration`—Trace configuration.

  - `dynamic-vpn`—Trace dynamic VPN events.

  - `init`—Trace the daemon init process.

  - `mgd`—Trace MGD requests.

  - `webauth`—Trace Web authentication requests.

- `level` *`level`* —Specify the level of debugging output.

  - `all`—Match all levels.

  - `error`—Match error conditions.

- `info`—Match informational messages.
- `notice`—Match conditions that should be handled specially.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.

- `no-remote-trace`—Disable remote tracing.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for `https` introduced for SRX5400, SRX5600, and SRX5800 devices starting from Junos OS Release 12.1X44-D10 and on vSRX Virtual Firewall, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, and SRX1500 devices starting from Junos OS Release 15.1X49-D40.

`virtual-domain` statement is introduced in Junos OS Release 23.1R1

RELATED DOCUMENTATION

Secure Management Access Configuration Summary

Firewall User Authentication Overview

Dynamic VPN Overview

web-redirect

IN THIS SECTION

- Syntax | 31084
- Hierarchy Level | 31084
- Description | 31084
- Required Privilege Level | 31084
- Release Information | 31084

Syntax

```
web-redirect;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name  policy policy-name then permit
firewall-authentication pass-through user-firewall]
```

Description

Optionally, redirect HTTP requests to the device's internal webserver by sending a redirect HTTP response to the client system to reconnect to the webserver for user authentication. The interface on which the client's request arrived is the interface to which the request is redirected.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Starting with Junos OS Release 15.1X49-D70 and Junos OS Release 17.3R1, support for user-firewall added on SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall Services Gateways.

RELATED DOCUMENTATION

Understanding User Role Firewalls

web-redirect-to-https

IN THIS SECTION

- Syntax | 31085
- Hierarchy Level | 31085
- Description | 31085
- Required Privilege Level | 31086
- Release Information | 31086

Syntax

```
web-redirect-to-https;
```

Hierarchy Level

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
firewall-authentication pass-through user-firewall]
```

Description

Redirect unauthenticated HTTP requests to the internal HTTPS webserver of the device.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced on on SRX5600 and SRX5800 Services Gateways starting from Junos OS Release 12.1X44-D10, and on vSRX Virtual Firewall, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Starting with Junos OS Release 15.1X49-D70 and Junos OS Release 17.3R1, support for `user-firewall` added on SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall Services Gateways.

RELATED DOCUMENTATION

Unified Threat Management User Guide

Firewall User Authentication Overview

web-server (Services)

IN THIS SECTION

- Syntax | 31087
- Hierarchy Level | 31087
- Description | 31087
- Options | 31088
- Required Privilege Level | 31088
- Release Information | 31089

Syntax

```
web-server {
    server-name;
    address address;
    connect-method (http | https);
    port port;
}
```

Hierarchy Level

```
[edit services user-identification authentication-source (Services User Identification ClearPass)
aruba-clearpass user-query (Services User Identification)]
```

Description

Specify the name of the webserver configuration on the SRX Series Firewall used for the user query integrated ClearPass authentication and enforcement function. The webserver is the ClearPass server to which the SRX Series Firewall connects to request authentication and identity information for an individual user.

When information for the individual user is not posted to the SRX Series Firewall by ClearPass through Web API POST request messages, the SRX Series Firewall can request this information from the ClearPass Policy Manager (CPPM) under certain circumstances. You must enable the user query function by configuring it.

address- Configure for the integrated ClearPass authentication and enforcement feature the address of the ClearPass webserver that the SRX Series Firewall communicates with. The SRX Series Firewall requests user authentication and identity information for an individual user from the ClearPass webserver whose address is configured. If you configure the user query function, the SRX Series Firewall can obtain this information for a specific user when it does not receive it from the ClearPass Policy Manager through Web API POST requests.

connect-method- Configure the application protocol used for the SRX Series Firewall connection to the ClearPass Policy Manager (CPPM) for user query requests. You identify the connection protocol as part of the configuration that identifies the CPPM server. The user query function allows the SRX Series Firewall to request from the CPPM user authentication and identity information for an individual user.

port- Configure the port on the Juniper Identity Management Service server that the SRX Series Firewall uses to connect to the server.

Options

server-name Specify the Web server name.

address Specify the IP address or hostname of web server.

http Configure HTTP as the connection protocol to use for the SRX Series integrated ClearPass authentication and enforcement feature's connection to the ClearPass Policy Manager (CPPM) webserver for individual user authentication queries. You can identify the connection protocol as part of the configuration that identifies the CPPM webserver (mutually exclusive with HTTPS).

If the SRX Series Firewalls does not find an authentication entry for a user in its local ClearPass authentication table, it can query the Aruba ClearPass webserver for this information.

https Configure HTTPS as the connection protocol used for the SRX Series connection to the ClearPass Policy Manager (CPPM) for user query requests. You identify the connection protocol as part of the configuration that identifies the CPPM webserver.

The integrated ClearPass authentication and enforcement user query function allows the SRX Series Firewall to request from the CPPM user authentication and identity information for an individual when the SRX Series ClearPass authentication table does not contain that information.

The http and https configuration assumes that aruba-clearpass is specified as the authentication source.

The Web API process, acting as an HTTP server, exposes to the Aruba ClearPass Policy Manager (CPPM) an API that allows the CPPM, acting as a client, to send POST request messages to it. The CPPM, which serves as the authentication source, initiates the session to the SRX Series Firewall and sends it user authentication and identity information.

- **Default:** https—The connect-method configuration is optional. If it is not configured, HTTPS is assumed.

port Specify the Web server port number.

- **Default:** 443
- **Range:** 1 through 65535

Required Privilege Level

1. services—To view this statement in the configuration.

2. services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

webapi-clear-text (Security)

IN THIS SECTION

- Syntax | 31089
- Hierarchy Level | 31089
- Description | 31089
- Required Privilege Level | 31089
- Release Information | 31090

Syntax

```
web-api-cleartext
```

Hierarchy Level

```
[edit security zones security-zone zone host-inbound-traffic system-services]
```

Description

Enable the Web API (webapi) service over HTTP host inbound traffic on TCP port 8080 for unencrypted data.

Required Privilege Level

1. security—To view this statement in the configuration.

2. security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

webapi-ssl (Security)

IN THIS SECTION

- Syntax | 31090
- Hierarchy Level | 31090
- Description | 31090
- Required Privilege Level | 31090
- Release Information | 31091

Syntax

```
webapi-ssl
```

Hierarchy Level

```
[edit security zones security-zone zone host-inbound-traffic system-services]
```

Description

Enable the Web API service over HTTPS host inbound traffic on TCP port 8443.

Required Privilege Level

security	To view this statement in the configuration.

security-control To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

webapi (System Services)

IN THIS SECTION

- Syntax | 31091
- Hierarchy Level | 31092
- Description | 31092
- Options | 31092
- Required Privilege Level | 31093
- Release Information | 31093

Syntax

```
webapi {
    client (System Services) name;
    debug-level (System Services) (alert | crit | emerg | error | info | notice | warn);
    debug-log (System Services) file;
    http (System Services) {
        port port;
    }
    https (System Services) {
        certificate certificate;
        certificate-key certificate-key;
        default-certificate;
        pki-local-certificate pki-local-certificate;
        port port;
    }
    user (System Services){
```

```
        user-name;
        password password;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Configure Web API (webapi) to facilitate efficient transmission of user authentication and identity information from the CPPM to the SRX Series Firewall. The CPPM, which is the client in this relationship, initiates a session with the SRX Series Firewall Web API process, which is the server in this relationship. However, the CPPM can do this only if you have configured the Web API function on the SRX Series Firewall. For security reasons, the Web API process is not enabled by default.

The configuring statements are explained separately. See CLI Explorer.

Options

client Address of permitted HTTP/HTTPS request originator.

debug-level Debug level for webapi process.

- Values:
  - alert—Match alert messages
  - crit—Match critical messages
  - emerg—Match emergence messages
  - error—Match error messages
  - info—Match informational messages
  - notice—Match notice messages
  - warn—Match warning messages

debug-log Debug log for webapi process.

http Unencrypted HTTP connection settings.

https Encrypted HTTPS connection settings.

user Specify the User name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X48-D30.

websense-redirect

IN THIS SECTION

- Syntax | 31093
- Hierarchy Level | 31094
- Description | 31094
- Options | 31094
- Required Privilege Level | 31095
- Release Information | 31095

Syntax

```
websense-redirect {
    profile profile-name {
        account value;
```

```
        custom-block-message value;
        fallback-settings {
            default (block | log-and-permit);
            server-connectivity (block | log-and-permit);
            timeout (block | log-and-permit);
            too-many-requests (block | log-and-permit);
        }
        server {
            host host-name;
            port number;
        }
        sockets value;
        timeout value;
        no-safe-search;
    }
}
```

Hierarchy Level

```
[edit security utm default-configuration]
[edit security utm feature-profile web-filtering]
```

Description

Configure the Websense redirect engine features.

Starting with Junos OS Release 17.4R1, you can download and dynamically load new Enhanced Web Filtering (EWF) categories. The downloading and dynamic loading of the new EWF categories do not require a software upgrade. Websense occasionally releases new EWF categories. EWF classifies websites into categories according to host, URL, or IP address and performs filtering based on the categories.

The new categories do not affect the existing configurations. You can modify the existing configurations to make use of the new categories.

Options

no-safe-search Disable the safe search function.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

The `[edit security utm default-configuration]` hierarchy level is introduced in Junos OS Release 18.2R1.

`no-safe-search` option added in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

> *Example: Enhancing Security by Configuring Redirect Web Filtering Using Custom Objects*

weight

IN THIS SECTION

- Syntax | 31095
- Hierarchy Level | 31096
- Description | 31096
- Options | 31096
- Required Privilege Level | 31096
- Release Information | 31096

Syntax

```
weight number;
```

Hierarchy Level

```
[edit chassis cluster redundancy-group group-number interface-monitor interface ]
[edit chassis cluster redundancy-group group-number ip-monitoring IP-address]
```

Description

Specify the relative importance of the object to the operation of the redundancy group. This statement is primarily used with interface monitoring and IP address monitoring objects. The failure of an object—such as an interface—with a greater weight brings the group closer to failover. Every monitored object is assigned a weight.

- interface-monitor objects—If the object fails, its weight is deducted from the threshold of its redundancy group;
- ip-monitoring objects—If a monitored IP address becomes unreachable for any reason, the weight assigned to that monitored IP address is deducted from the redundancy group's global-threshold for IP address monitoring.

Every redundancy group has a default threshold of 255. If the threshold reaches 0, a failover is triggered. Failover is triggered even if the redundancy group is in manual failover mode and preemption is not enabled.

Options

number —Weight assigned to the interface or monitored IP address. A higher weight value indicates a greater importance.

- **Range:** 0 through 255

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 10.1.

RELATED DOCUMENTATION

cluster (Chassis) | 4369

weight (Targeted Subscribers)

IN THIS SECTION

- Syntax | 31097
- Hierarchy Level | 31097
- Description | 31097
- Options | 31098
- Required Privilege Level | 31098
- Release Information | 31098

Syntax

```
weight ($junos-interface-set-target-weight | weight-value);
```

Hierarchy Level

```
[edit interfaces interface-set interface-set-name targeted-options]
```

Description

Configure weight for targeted subscribers like PPPoe, demux, and conventional VLANs. The weight assigned is based on factors such as customer preference, CoS, or bandwidth requirement. The member links are then assigned based on the value of the weight. The value of the weight can range from 1 through 1000.

Options

weight ($junos-interface-set-target-weight | *weight-value*)

Specify the weight for targeted subscribers. Member links for logical interfaces of aggregated Ethernet logical interfaces are assigned based on the value of the weight . When a new VLAN is added to the same aggregated Ethernet bundle, then the primary member link selected for targeting is the one with the minimum primary load and the backup link selected for targeting is the one with the minimum overall load.

When you configure the $junos-interface-set-target-weight predefined variable, the weight value is sourced from VSA 26-214 in the RADIUS Access-Accept message when a dynamic subscriber is authenticated and applies to both the interface set and all its member interfaces; you must configured targeted distribution for both the interface set and its member interfaces. If the interface set is not explicitly configured and RADIUS VSA 26-214 is not received, then the interface set weight derives from the weight assigned to the first member interface that is authorized.

- **Range:** 1 through 1000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

`$junos-interface-set target-weight` option added in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Understanding Support for Targeted Distribution of Logical Interface Sets of Static VLANs over Aggregated Ethernet Logical Interfaces

Using RADIUS-Sourced Weights for Interface and Interface Set Targeted Distribution

RADIUS-Sourced Weights for Interface and Interface Set Targeted Distribution

weight (Security)

IN THIS SECTION

- Syntax | 31099
- Hierarchy Level | 31099
- Description | 31099
- Required Privilege Level | 31100
- Release Information | 31100

Syntax

```
weight (equal | firewall | idp);
```

Hierarchy Level

```
[edit security forwarding-process application-services maximize-idp-sessions]
```

Description

If you are deploying IDP policies, you can tune the device to increase IDP session capacity. By using the provided commands to change the way the system allocates resources, you can achieve a higher IDP session capacity.

Devices ship with an implicit default session capacity setting. This default value gives more weight to firewall sessions. You can manually override the default by using the `maximize-idp-sessions` command. The command allows you to choose between these weight values: `equal`, `firewall`, and `idp`. The following table displays the available session capacity weight and approximate throughput for each.

Table 1839: Session Capacity and Resulting Throughput

Weight Value	Firewall Capacity	IDP Capacity	Firewall Throughput	IDP Throughput
Default	1,000,000	256,000	10 Gbps	2.4 Gbps
`equal`	1,000,000	1,000,000	8.5 Gbps	2 Gbps
`firewall`	1,000,000	1,000,000	10 Gbps	2.4 Gbps
`idp`	1,000,000	1,000,000	5.5 Gbps	1.4 Gbps

This statement is supported on SRX1500, SRX 5800, SRX 5600, and SRX 5400 devices and vSRX Virtual Firewall.

Required Privilege Level

security—To view this in the configuration.

security-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

weighted-load-balancing (L2TP LAC)

IN THIS SECTION

Syntax | **31101**

- Hierarchy Level | 31101
- Description | 31101
- Required Privilege Level | 31101
- Release Information | 31101

Syntax

```
weighted-load-balancing;
```

Hierarchy Level

```
[edit services l2tp]
```

Description

Specify that the router considers tunnel weight when selecting from among multiple tunnels that share the same preference level. A higher maximum session limit on a tunnel corresponds to a higher tunnel weight. A tunnel with a higher weight is more likely to be selected than a tunnel with a lower weight. The distribution of sessions across all tunnels in the preference level, on average, is proportional to the tunnel weight

Disabled by default. By default, tunnel selection within a preference level is strictly random. The **destination-equal-load-balancing** statement must be disabled to successfully enable this statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

L2TP LAC Tunneling for Subscribers

L2TP LAC Tunneling for Subscribers

west-interface

IN THIS SECTION

- Syntax | 31102
- Hierarchy Level | 31103
- Description | 31103
- Required Privilege Level | 31103
- Release Information | 31103

Syntax

```
west-interface {
    node-id mac-address;
    control-channel channel-name {
        vlan number;
        interface name interface-name
    }
    interface-none
    ring-protection-link-end;
    virtual-control-channel {
        west-interface name;
        east-interface name;
    }
}
```

Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

Description

Define one of the two interface ports for Ethernet ring protection, the other being defined by the `east-interface` statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

NOTE: Always configure this port second, after configuring the `east-interface` statement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Ethernet Ring Protection Switching Overview

Ethernet Ring Protection Using Ring Instances for Load Balancing

east-interface | **6186**

ethernet-ring | **6497**

Example: Configuring Ethernet Ring Protection Switching on EX Series Switches

Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS

Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)

when

IN THIS SECTION

- Syntax | 31104
- Description | 31104
- Options | 31104
- Required Privilege Level | 31105
- Sample Output | 31105
- Release Information | 31105

Syntax

```
when
<chassis chassis-id | member member of vc | model model | node node | peer peer | routing-engine
routing-engine | time time>
```

Description

Display inherited data for a specified parameter that is defined as conditions under the configuration "groups" on page 7916.

Options

chassis *chassis-id*	Specify the chassis type of the device.
member *member of vc*	Specify the name of the member of the virtual chassis.
model *model*	Specify the model name of the device.
node *node*	Specify the node of a cluster.
peer *peer*	Specify the peer host.
routing-engine *routing-engine*	Specify the type of the Routing Engine.

time *time*

Specify the start time or time duration. The syntax for specifying the time is: time <start-time> [to <end-time>] using the time format yyyy-mm-dd.hh:mm, hh:mm, or hh.

Required Privilege Level

view

Sample Output

show | display inheritance when

```
user@host# show | display inheritance when
## Last changed: 2023-01-05 03:44:28 PST
version 22.4I-20220921.0.1158;
    login {
        ##
        ## 'wheel' was inherited from group 'global'
        ##
        class wheel {
            ##
            ## 'snmp' was inherited from group 'global'
            ##
            permissions snmp;
        }
        ##
```

Release Information

`when` option introduced in Junos OS Release 11.3.

`peers` option added in Junos OS Release 14.2R6 for the MX Series.

`peers` option added in Junos OS Release 16.1R1 for the EX Series.

RELATED DOCUMENTATION

Example: Configure Conditions for Applying Configuration Groups

show | display inheritance | **19529**

groups | 7916

white-list

IN THIS SECTION

- Syntax | 31106
- Hierarchy Level | 31106
- Description | 31106
- Options | 31107
- Required Privilege Level | 31107
- Release Information | 31107

Syntax

```
white-list {
                IP-address-list IP-address-list-name;
                 file-name file-name;
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name security-intelligence-policy]
```

Description

Configure a global allowlist of IP addresses to allow packets to the destined countries.

Options

- `IP-address-list` *`ip-address-list-name`* —Contains the list of IP addresses that must be allowed. *ip-address-list-name* is the name of the IP address-list specified at the `[edit services web-filter]` hierarchy.
- `file-name` *`file-name`*—Name of the file which contains the IP addresses that must be allowed. The file must be in the **/var/db/url-filterd** location and must have the same name as in the configuration.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1 on MX Series routers with Juniper Advanced Threat Prevention Cloud (Juniper ATP Cloud) .

RELATED DOCUMENTATION

Global Allowlist and Global Blocklist

white-list (Security Screen)

IN THIS SECTION

- Syntax | **31108**
- Hierarchy Level | **31108**
- Description | **31108**
- Options | **31108**
- Required Privilege Level | **31109**
- Release Information | **31109**

Syntax

```
white-list name {
    address [address...];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name security screen],
[edit security screen],
[edit tenants tenant-name security screen]
[edit logical-systems logical-system-name security screen ids-option screen-name udp flood],
[edit security screen ids-option screen-name udp flood],
[edit tenants tenant-name security screen ids-option screen-name udp flood]
```

Description

Configure a list of IP addresses that are exempted from UDP flood detection, which occur during the UDP flood screen protection process. This list of exempted addresses is called an allowlist.

You can use this statement to configure an allowlist of IP addresses that bypass UDP flood detection.

> **NOTE**: This statement is not supported to create UDP flood screen allowlists on SRX5400, SRX5600, and SRX5800 devices.

Both IPv4 and IPv6 allowlists are supported. Addresses in an allowlist must be all IPv4 or all IPv6. In each allowlist, there can be up to 32 IP addresses.

Options

- `name` *White-list name*—The name of the allowlist.
- `address` *address*— The list of IP addresses. You can specify multiple addresses or address prefixes as a sequence of addresses separated by spaces and enclosed in square brackets. You can configure single address or subnet address.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

Support for UDP flood screen allowlist introduced in Junos OS Release 17.4.

`tenant` option added in Junos OS Release 18.3R1.

Support for UDP and TCP flood screen allowlists added in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480 and MX960 routers.

RELATED DOCUMENTATION

- Understanding Allowlists for SYN Flood Screens
- Understanding Allowlist for UDP Flood Screens

white-list (Security-Zone)

IN THIS SECTION

- Syntax | 31110
- Hierarchy Level | 31110
- Description | 31110
- Options | 31110
- Required Privilege Level | 31110
- Release Information | 31110

Syntax

```
white-list [ white-list ... ];
```

Hierarchy Level

```
[edit logical-systems name security zones security-zone],
[edit logical-systems name tenants name security zones security-zone],
[edit security zones security-zone],
[edit tenants name security zones security-zone]
```

Description

Configure allowlist for all IP screen options in a security zone.

You can use this statement to configure allowlist for all IP screen options at zone level to bypass IP packet screening check. Both IPv4 and IPv6 allowlists are supported. Addresses in an allowlist must be all IPv4 or all IPv6. In each allowlist, there can be up to 32 IP addresses.

Options

White-list name—The name of the allowlist.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 23.4R1.

whitelist

IN THIS SECTION

- Syntax | 31111
- Hierarchy Level | 31111
- Description | 31111
- Options | 31112
- Required Privilege Level | 31112
- Release Information | 31112

Syntax

```
whitelist [global-address-book-addresses];
```

Hierarchy Level

```
[edit services ssl proxy profile profile-name]
[edit services ssl termination profile profile-name]
```

Description

Specify the addresses exempted from the SSL proxy. This statement is supported on the SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX Virtual Firewall.

You can selectively bypass SSL proxy processing for some sessions by configuring a allowlist. Typically, you might configure the allowlist to include trusted servers or domains with which you are very familiar. An allowlist include addresses that you want to exempt from undergoing SSL proxy processing.

To configure the allowlist, you need to specify the domain that you want to exempt in an address book and then configure the address in the SSL proxy profile.

Options

- *allowlist-address*—Specify address from the global address book.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Configuring SSL Forward Proxy

Firewall User Authentication Overview

whitelist-url-categories

IN THIS SECTION

- Syntax | 31112
- Hierarchy Level | 31113
- Description | 31113
- Options | 31113
- Required Privilege Level | 31113
- Release Information | 31113

Syntax

```
whitelist-url-categories url-category-list;
```

Hierarchy Level

```
[edit services ssl proxy profile profile-name]
```

Description

Configure the predefined URL categories in SSL proxy profile to exempt from SSL inspection. The URL category identification is leveraged from the Web filtering categories obtained from the Content Security module.

Before you specify URL category list, you must create a web filtering profile with custom objects using custom URL category or use predefined list. Next apply the feature profile to the Content Security policy.

The following example uses a predefined profile, junos-wf-enhanced-default.

```
[edit]
user@host# set security utm feature-profile web-filtering type juniper-enhanced
user@host# set security utm utm-policy policy-name web-filtering http-profile junos-wf-enhanced-
default
```

Starting in Junos OS Release 17.4R1, you can use custom URL categories in SSL proxy profile.

Options

url-category-list	List of predefined or custom URL category.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1X49-D80.

RELATED DOCUMENTATION

Creating an Allowlist of Exempted URL Categories for SSL Proxy

wildcard delete

IN THIS SECTION

- Syntax | 31114
- Description | 31114
- Options | 31114
- Required Privilege Level | 31115
- Release Information | 31115

Syntax

```
wildcard delete <statement-path> <identifier> <regular-expression>
```

Description

Delete a statement or identifier. All subordinate statements and identifiers contained within the specified statement path are deleted with it.

Deleting a statement or an identifier effectively “unconfigures” or disables the functionality associated with that statement or identifier.

If you do not specify *statement-path* or *identifier*, the entire hierarchy starting at the current hierarchy level is removed.

Options

identifier (Optional) Name of the statement or identifier to delete.

regular-expression	(Optional) The pattern based on which you want to delete multiple items. When you use the `wildcard` command to delete related configuration items, the *regular-expression* must be the final statement.
statement-path	(Optional) Path to an existing statement or identifier. Include this if the statement or identifier to be deleted is not at the current hierarchy level.

Required Privilege Level

configure—To enter configuration mode. Other required privilege levels depend on where the statement is located in the configuration hierarchy.

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Using Global Replace in a Device Configuration—Using the upto Option

Example: Use the Wildcard Command with the Range Option

wildcard-group-inet

IN THIS SECTION

- Syntax | 31116
- Hierarchy Level | 31116
- Description | 31116
- Required Privilege Level | 31116
- Release Information | 31117

Syntax

```
wildcard-group-inet {
    wildcard-source {
        inter-region-segmented{
            fan-out fan-out value;
        }
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
        threshold-rate number;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective],
[edit routing-instances routing-instance-name provider-tunnel selective]
```

Description

Configure a wildcard group matching any group IPv4 address.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

The `inter-region-segmented` statement added in Junos OS Release 15.1.

RELATED DOCUMENTATION

wildcard-group-inet6 | 31117

Example: Configuring Selective Provider Tunnels Using Wildcards

Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN

Configuring a Selective Provider Tunnel Using Wildcards

wildcard-group-inet6

IN THIS SECTION

- Syntax | 31117
- Hierarchy Level | 31118
- Description | 31118
- Required Privilege Level | 31118
- Release Information | 31118

Syntax

```
wildcard-group-inet6 {
    wildcard-source {
        inter-region-segmented{
            fan-out fan-out value;
        }
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
```

```
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
        threshold-rate number;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective],
[edit routing-instances routing-instance-name provider-tunnel selective]
```

Description

Configure a wildcard group matching any group IPv6 address.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

The `inter-region-segmented` statement added in Junos OS Release 15.1.

RELATED DOCUMENTATION

wildcard-group-inet | **31115**

Example: Configuring Selective Provider Tunnels Using Wildcards

Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN

Configuring a Selective Provider Tunnel Using Wildcards

wildcard-source (PIM RPF Selection)

IN THIS SECTION

- Syntax | 31119
- Hierarchy Level | 31119
- Description | 31119
- Required Privilege Level | 31119
- Release Information | 31120

Syntax

```
wildcard-source {
    next-hop next-hop-address;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address],
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-
list-addresses]
```

Description

Use a wildcard for the multicast source instead of (or in addition to) a specific multicast source.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring PIM RPF Selection

wildcard-source (Selective Provider Tunnels)

IN THIS SECTION

- Syntax | 31120
- Hierarchy Level | 31121
- Description | 31121
- Required Privilege Level | 31121
- Release Information | 31121

Syntax

```
wildcard-source {
    inter-region-segmented {
        fan-out fan-out value;
    }
ldp-p2mp;
    pim-ssm {
        group-range multicast-prefix;
    }
    rsvp-te {
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
```

```
        static-lsp lsp-name;
    }
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective group group-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet6],
[edit routing-instances routing-instance-name provider-tunnel selective group group-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-
tunnel selective wildcard-group-inet6],
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet],
[edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet6]
```

Description

Configure a selective provider tunnel for a shared tree using a wildcard source.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.0.

The `inter-region-segmented` statement added in Junos OS Release 15.1.

RELATED DOCUMENTATION

wildcard-group-inet | 31115

wildcard-group-inet6 | 31117

Example: Configuring Selective Provider Tunnels Using Wildcards

Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN

Configuring a Selective Provider Tunnel Using Wildcards

window-scale

IN THIS SECTION

- Syntax | 31122
- Hierarchy Level | 31122
- Description | 31122
- Options | 31123
- Required Privilege Level | 31123
- Release Information | 31123

Syntax

```
window-scale {
    match (equal | greater-than | less-than | not-equal);
    value window-scale-factor;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify the scale factor that the session of the attack will use. The window scale extension expands the definition of the TCP window to 32 bits and then uses a scale factor to carry this 32-bit value in the 16-bit window field of the TCP header.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`window-scale-factor`*—Match the number of bytes.
- **Range:** 0 through 255

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

window-size

IN THIS SECTION

- Syntax | 31123
- Hierarchy Level | 31124
- Description | 31124
- Options | 31124
- Required Privilege Level | 31124
- Release Information | 31124

Syntax

```
window-size {
    match (equal | greater-than | less-than | not-equal);
    value window-size;
}
```

Hierarchy Level

```
[edit security idp custom-attack attack-name attack-type signature protocol tcp]
```

Description

Specify the number of bytes in the TCP window size.

Options

- `match` (`equal` | `greater-than` | `less-than` | `not-equal`)—Match an operand.
- `value` *`window-size`*—Match the number of bytes.
- **Range:** 0 through 65,535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

windows-logon (Juniper Secure Connect)

IN THIS SECTION

- Syntax | 31125
- Hierarchy Level | 31125
- Description | 31125
- Options | 31125
- Required Privilege Level | 31126

Release Information | 31126

Syntax

```
windows-logon {
    auto-dialog-open;
    disconnect-at-logoff;
    domain domain;
    eap-auth;
    flush-credential-at-logoff;
    lead-time-duration seconds;
    mode (automatic | manual);
}
```

Hierarchy Level

```
[edit security remote-access client-config]
```

Description

Define windows logon settings for the Juniper Secure Connect remote client device.

Options

auto-dialog-open	Automatically open dialog for connection establishment.
disconnect-at-logoff	Disconnect the session after logoff.
domain	Domain name for automatic windows logon.
eap-auth	EAP authentication method before the profile selection.
flush-credential-at-logoff	Flush cached credentials after logoff.
lead-time-duration	Lead time duration for domain logon in seconds.

- **Default:** 45 seconds
- **Range:** 0 through 120

mode

Set windows logon mode.

- Values:
  - automatic—Automatic Windows logon with configured credentials.
  - manual—Manual Windows logon.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Juniper Secure Connect Administrator Guide

winnuke

IN THIS SECTION

- Syntax | 31127
- Hierarchy Level | 31127
- Description | 31127
- Required Privilege Level | 31127
- Release Information | 31127

Syntax

```
winnuke;
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp]
```

Description

Enable detection of attacks on Windows NetBios communications. Packets are modified as necessary and passed on. Each WinNuke attack triggers an attack log entry in the event alarm log.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

winnuke (IDS Screen Next Gen Services)

IN THIS SECTION

Syntax | 31128

- Hierarchy Level | 31128
- Description | 31128
- Required Privilege Level | 31128
- Release Information | 31128

Syntax

```
winnuke;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Description

Identify and drop TCP segments that are destined for port 139 and have the urgent (URG) flag set, which provides protection against WinNuke attacks.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services

wins-server (Access)

IN THIS SECTION

- Syntax | 31129
- Hierarchy Level | 31129
- Description | 31129
- Required Privilege Level | 31129
- Release Information | 31129

Syntax

```
wins-server address
```

Hierarchy Level

```
[edit access address-assignment pool <name> family (inet | inet6) xauth-attributes]
[edit access profile profile-name]
```

Description

Specify the wins-server IP address.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.4 of Junos OS. Starting with Junos OS 15.1X49-D80 and Junos OS Release 17.3R1, the wins-server option at the [edit access profile] hierarchy level allows you to configure the IPv4 address of a Windows Internet Name Service (WINS) server.

RELATED DOCUMENTATION

Understanding Pass-Through Authentication

wins-server (System)

IN THIS SECTION

- Syntax | 31130
- Hierarchy Level | 31130
- Description | 31130
- Options | 31131
- Required Privilege Level | 31131
- Release Information | 31131

Syntax

```
wins-server {
    address;
}
```

Hierarchy Level

```
[edit system services dhcp],
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]
```

Description

For J Series Services Routers and EX Series switches only. Specify one or more NetBIOS Name Servers. When a DHCP client is added to the network and assigned an IP address, the NetBIOS Name Server

manages the Windows Internet Name Service (WINS) database that matches IP addresses (such as `192.168.1.3`) to Windows NetBIOS names (such as **\\Marketing**). List servers in order of preference.

Options

address—IPv4 address of the NetBIOS Name Server running WINS. To configure multiple servers, include multiple *address* options.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

wired-max-processes

IN THIS SECTION

- Syntax | 31131
- Hierarchy Level | 31132
- Description | 31132
- Options | 31132
- Required Privilege Level | 31132
- Release Information | 31132

Syntax

```
wired-max-processes num-procs;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic slot-number adaptive-services service-package extension-
provider]
```

Description

Configure the number of processes that use wired process memory. Performance can degrade if a process uses memory beyond its Big TLB memory. If this setting is changed, the PIC will reboot.

Options

num-procs Number of processes that use the reserved wired process memory.

- **Range:** 1 through 8

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

wired-process-mem-size | **31133**

forwarding-db-size | **7454**

object-cache-size | **11765**

policy-db-size | **12674**

wired-process-mem-size

IN THIS SECTION

- Syntax | 31133
- Hierarchy Level | 31133
- Description | 31133
- Options | 31133
- Required Privilege Level | 31134
- Release Information | 31134

Syntax

```
wired-process-mem-size mem-size;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-
provider]
```

Description

Configure the size of the reserved wired process memory. You can also configure the object cache. If this setting is changed, the PIC reboots.

Options

megabytes—Size of the reserved wired process memory, in MB. The only size you can set for this statement is 512 MB.

- **Default:** 512 MB
- **Range:** 0 through 512 MB

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

- forwarding-db-size | 7454
- object-cache-size | 11765
- policy-db-size | 12674
- wired-max-processes | 31131

withdraw-priority

IN THIS SECTION

- Syntax | 31134
- Hierarchy Level | 31135
- Description | 31135
- Required Privilege Level | 31135
- Release Information | 31136

Syntax

```
withdraw-priority (expedited | priority priority-queue-number (1-16));
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family family-name sub-family],
[edit logical-systems logical-system-name protocols bgp group group-name family family-name sub-family],
[edit protocols bgp family family-name sub-family],
[edit protocols bgp group group-name family family-name sub-family],
[edit protocols bgp group group-name neighbor neighbor-id family family-name]
```

Description

Within BGP route prioritization, the `withdraw-priority` statement allows you to set specific priority levels for BGP routes that are to be withdrawn. The `withdraw-priority` statement can be configured for BGP neighbors during BGP configuration, or for sub-families within the following address families:

- `evpn`
- `inet`
- `inet-mdt`
- `inet-mvpn`
- `inet-vpn`
- `inet6`
- `inet6-mvpn`
- `inet6-vpn`
- `iso-vpn`
- `l2vpn`
- `route-target`
- `traffic-engineering`

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

output-queue-priority | 12008

Understanding BGP Route Prioritization

within (Event Policy)

IN THIS SECTION

- Syntax | 31136
- Hierarchy Level | 31136
- Description | 31137
- Options | 31137
- Required Privilege Level | 31137
- Release Information | 31137

Syntax

```
within seconds {
    events [ events ];
    not events [ events ];
    trigger (after | on | until) event-count;
}
```

Hierarchy Level

```
[edit event-options policy policy-name]
```

Description

Configure events that must (or must not) occur within a specified time interval before one of the trigger events in order for the event policy to be triggered.

The remaining statements are explained separately. See CLI Explorer.

Options

seconds Time interval between a correlated event and a trigger event.

- **Range:** 60 through 604,800 seconds

Required Privilege Level

maintenance—To view or add this statement in the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Use Correlated Events to Trigger an Event Policy

wlan

IN THIS SECTION

- Syntax (SRX Series) | **31138**
- Hierarchy Level | **31139**
- Description | **31139**
- Options | **31139**
- Required Privilege Level | **31142**
- Release Information | **31142**

Syntax (SRX Series)

```
wlan {
    access-point name {
        description description;
        interface wl interface;
        access-point-options {
            country country-code;
        }
        location location;
        mac-address mac-address;
        radio (1| 2) {
            radio-options {
            channel {
                number (auto | channel-number);
                bandwidth (20 | 40 | 80);
            }
            mode (g | gn | an | acn);
            radio-off;
            transmit-power percent;
            }
        }
    }
    virtual-access-point id {
        description description;
        no-broadcast-ssid;
        maximum-stations number;
        station-isolation;
        upload-limit upload-limit-rate;
        download-limit download-limit-rate;
        ssid ssid;
        vlan vlan-id;
        station-mac-filter (allow-list | deny-list) {
            mac-address addr1 addr2;
        }
    security {
        none;
        wpa-enterprise {
            cipher-suites ccmp;
            radius-server ip-address;
            radius-port port;
            radius-key secret-key;
```

```
            wpa-version v2;
            }
            wpa-personal {
                cipher-suites ccmp;
                key (ascii | hex) key;
                wpa-version v2;
            }
        }
    }
}
```

Hierarchy Level

```
[edit wlan access-point name]
```

Description

Configure WLAN properties on SRX Series Firewalls.

Options

access-point name
Name of the wireless access point.

interface
Wireless LAN interface (wl-*x*/0/0) created for the access-point setting.

To support Mini-PIM in chassis cluster mode, you can configure two WLAN interfaces (wl-*x*/0/0) and (wl-*y*/0/0) on both nodes, then the WLAN configuration for access-point is pushed to the two wireless LAN interface cards on both the nodes.

description
Description of the access point and virtual access point (VAP). The maximum length is 64 characters.

country
The country code.

location
Location of the access point. The maximum number of characters you can use is 64.

channel number (auto | channel number)
Channel number of the radio. If you select `auto`, then the Mini-PIM chooses the channel automatically.

bandwidth

Radio 1 (5 GHz) supports bandwidth of 20MHz, 40MHz, and 80MHz, whereas Radio 2 (2.4 GHz) supports bandwidth of 20MHz and 40MHz. The default value is 20MHz for 2.4GHz and 40MHz for 5GHz.

radio modes (an | acn | gn | g)

Mode for the radio operation.

Radio 1 supports the following modes:

- **an** 802.11a and 802.11n clients operating in 5-GHz frequency can connect to the access point.

- **acn** 802.11a, 802.11b, 802.11n and 802.11ac clients operating in 5-GHz frequency can connect to the access point.

Radio 2 supports the following mode:

- **gn** 802.11g, 802.11b, and 802.11n clients operating in 2.4-GHz frequency can connect to the access point. This is the default mode for this radio.

  g 802.11g clients operating in 2.4-GHz frequency can connect to the access point supported from Junos OS Release 20.4R1.

- **radio-off** Radio is turned off.

- **transmit-power** The percentage of transmit power.

  NOTE: When you configure the transmit power, the Mini-PIM card will fix transmit power to the specified value set, in this case the power by rate functionality does not work. So it is recommended not to set transmit power to a specified value. When you do not configure the transmit power (do not fix the transmit power to a specified value), the power by rate functionality works. If you configure the transmit power percentage to 100, then it chooses the option "`auto`", the behavior is same as no transmit power configuration is done and power by rate functionality will work.

no-broadcast-ssid

Disable broadcast SSID. By default, the broadcast SSID is enabled.

maximum-stations	The number of maximum clients that can be connected to the virtual access point. The range is 1 through 127.
station-isolation	Isolate the clients connected to the same VAP.
security (none | wpa-enterprise | wpa-personal)	Security settings for the VAP. WPA enterprise is a Wi-Fi Alliance standard that uses RADIUS server authentication with AES-CCMP. This mode allows the use of high-security encryption along with centrally managed user authentication. WPA personal is a Wi-Fi Alliance standard that uses preshared key (PSK) authentication with AES-CCMP . Only WPA2 standards are supported on Wi-Fi Mini-PIM.
none	No security. The data transferred between clients and the access point is not encrypted. This method allows clients to associate with the access point without any authentication.
cipher-suites (ccmp)	Select the appropriate WPA cipher algorithm. The value is CCMP algorithms.
radius-server	IP address of the radius server.
radius-port	Port number of the RADIUSs server. The default value is 1812.
radius-key	Secret key of the RADIUS. The maximum number of characters you can use is 64.
key (ascii | hex)	WPA shared key. The range of key length is 8 through 63 for ASCII or 8 through 64 hexadecimal characters.
wpa-version (v2)	Version of the WPA version. Only supported value is WPA2.
upload-limit	Specify the upload rate limit. The range is from 256 Kbps through 1,048,576 Kbps.
download-limit	Specify the download rate limit. The range is from 256 Kbps through 1,048,576 Kbps.
ssid	SSID value for the virtual access point. The range is 2 through 32. SSID value can include only letters, numbers and five special characters—hyphen (-), underscore (_), at (@), hash (#), and period (.) in the value of the SSID.
vlan-id	VLAN ID for the virtual access point. The range is 1 through 4094. The default value is 1.
station-mac-filter (allow-list| deny-list)	Specify the MAC filter. You can set either allow the mac address list or deny it. The MAC address format is like xx:xx:xx:xx:xx:xx. The maximum number of the MAC addresses listed is 16.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

Wi-Fi Mini-Physical Interface Module Overview

Configure Wi-Fi Mini-PIM

world-readable (Forwarding Options)

IN THIS SECTION

- Syntax | 31142
- Hierarchy Level | 31143
- Description | 31143
- Options | 31143
- Required Privilege Level | 31143
- Release Information | 31143

Syntax

```
(world-readable | no-world-readable);
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family (inet |inet6 |mpls) output file],
[edit forwarding-options sampling traceoptionsfile]
```

Description

Enable unrestricted file access.

Options

`no-world-readable`—Restrict file access to owner. This is the default.

`world-readable`—Enable unrestricted file access.

- **Default:** `no-world-readable`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

Configuring Traffic Sampling on MX, M and T Series Routers

Collecting Traffic Sampling Output in a File

Tracing Traffic-Sampling Operations

world-readable (System Syslog)

IN THIS SECTION

- Syntax | 31144
- Hierarchy Level | 31144
- Description | 31144
- Default | 31144
- Required Privilege Level | 31144
- Release Information | 31145

Syntax

```
world-readable | no-world-readable;
```

Hierarchy Level

```
[edit system syslog archive],
[edit system syslog file filename archive]
```

Description

Grant all users permission to read log files, or restrict the permission only to the `root` user and users who have the Junos OS `maintenance` permission.

Default

`no-world-readable`

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Specify Log File Size, Number, and Archiving Properties

world-readable (Next Gen Services Global System Logging)

IN THIS SECTION

- Syntax | 31145
- Hierarchy Level | 31145
- Description | 31145
- Default | 31146
- Options | 31146
- Required Privilege Level | 31146
- Release Information | 31146

Syntax

```
world-readable
```

Hierarchy Level

```
[edit services rtlog traceoptions file]
```

Description

Allow any user to read the log file

Default

By default, the `no-world-readable` option is set. No user is allowed to read the log file.

Options

world-readable Allow any user to read the log file

Required Privilege Level

system

Release Information

Statement introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging

Enabling Global System Logging for Next Gen Services

Configuring System Logging to One or More Remote Servers for Next Gen Services

Configuring Local System Logging for Next Gen Services

write-interval

IN THIS SECTION

- Syntax | 31147
- For Platforms with Enhanced Layer 2 Software (ELS) | 31147
- For Platforms Without ELS | 31147
- For MX Series Platforms | 31147
- Description | 31147
- Options | 31148

Required Privilege Level | 31148
Release Information | 31148

Syntax

```
write-interval seconds;
```

For Platforms with Enhanced Layer 2 Software (ELS)

```
[edit system processes dhcp-service dhcp-snooping-file],
[edit system processes dhcp-service dhcpv6-snooping-file]
```

For Platforms Without ELS

```
[edit ethernet-switching-options secure-access-port dhcp-snooping-file],
[edit ethernet-switching-options secure-access-port dhcpv6-snooping-file]
```

For MX Series Platforms

```
[edit system processes dhcp-service dhcp-snooping-file]
```

Description

Specify how frequently the device writes the database entries from memory into the DHCP snooping database file.

- If you are configuring `write-interval` at the `[edit system processes dhcp-service dhcp-snooping-file]` or the `[edit system processes dhcp-service dhcpv6-snooping-file]` hierarchy level, see *Configuring Persistent Bindings in the DHCP or DHCPv6 (ELS)*.

Options

seconds—Value in seconds.

- **Range:** 60 through 86,400 seconds.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Support at the `[edit system processes dhcp-service dhcp-snooping-file]` hierarchy level introduced in Junos OS Release 13.2X50-D10.

Support at the `[edit system processes dhcp-service dhcpv6-snooping-file]` hierarchy level introduced in Junos OS Release 13.2X51-D20.

Support at the `[edit ethernet-switching-options secure-access-port dhcpv6-snooping-file]` hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

RELATED DOCUMENTATION

Understanding DHCP Snooping (non-ELS)

write-view

IN THIS SECTION

- Syntax | 31149
- Hierarchy Level | 31149
- Description | 31149
- Options | 31149

Required Privilege Level | 31149
Release Information | 31149

Syntax

```
write-view view-name;
```

Hierarchy Level

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix context-
prefix) security-model (any | usm | v1 | v2c) security-level  (authentication | none | privacy)]
```

Description

Associate the write view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).

Options

view-name—Name of the view for which the SNMP user group has write permission.

Required Privilege Level

snmp—To view this statement in the configuration.

snmp-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

X

IN THIS SECTION

- xauth-attributes | 31150
- xe (Port) | 31152
- xlat-source-rule | 31154
- xle (Port) | 31155
- xnm-clear-text | 31156
- xnm-ssl | 31158
- xon (Input Congestion Notification) | 31161

xauth-attributes

IN THIS SECTION

- Syntax (inet) | 31150
- Syntax (inet6) | 31151
- Hierarchy Level | 31151
- Hierarchy Level (inet6) | 31151
- Description | 31151
- Options | 31151
- Required Privilege Level | 31152
- Release Information | 31152

Syntax (inet)

```
xauth-attributes {
    primary-dns IP address;
```

```
    primary-wins IP address;
    secondary-dns IP address;
    secondary-wins IP address;
}
```

Syntax (inet6)

```
xauth-attributes {
    primary-dns-ipv6 IP address;
    secondary-dns-ipv6  IP address;
}
```

Hierarchy Level

```
[edit access address-assignment pool <name> family inet]
```

Hierarchy Level (inet6)

```
[edit access address-assignment pool <name> family inet6]
```

Description

Configure XAuth attributes to use in XAuth authentication.

Options

- `apply-groups`—Groups from which to inherit configuration data.
- `apply-groups-except`—Do not inherit configuration data from these groups.
- `primary-dns`—Specify the primary-dns IP address.
- `secondary-dns`—Specify the secondary-dns IP address.
- `primary-wins`—Specify the primary-wins IP address.
- `secondary-wins`—Specify the secondary-wins IP address.

- `primary-dns-ipv6`—Specify the primary-dns IPv6 address.
- `secondary-dns-ipv6`—Specify the secondary-dns IPv6.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`xauth-attributes` option under `inet6` is introduced in Junos OS Release 20.3R1.

xe (Port)

IN THIS SECTION

- Syntax | 31152
- Hierarchy Level | 31153
- Description | 31153
- Options | 31153
- Required Privilege Level | 31153
- Release Information | 31153

Syntax

```
xe {
    (port port-number | port-range port-range-low port-range-high);
}
```

Hierarchy Level

```
[edit chassis fpc slot pic pic-number]
```

Description

(QFX3600 standalone switch only) Configure a specific port or a range of ports to operate as four 10-Gigabit Ethernet (*xe*) type ports.

CAUTION: The Packet Forwarding Engine on the switch is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the switch.

NOTE: Port Q0 supports only three (not the typical four)10-Gigabit Ethernet ports. Therefore, you can configure up to 63 (not 64) 10-Gigabit Ethernet ports on ports Q0 through Q15.

Options

port-number Port number on which you want to configure the port type. Valid values are 0 through 15.

port-range-low Lowest-numbered port in the range of ports. The lowest possible value is 0.

port-range-high Highest-numbered port in the range of ports. The highest possible value is 15.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2X50-D20.

RELATED DOCUMENTATION

Configuring the Port Type on QFX3600 Standalone Switches

xlat-source-rule

IN THIS SECTION

- Syntax | 31154
- Hierarchy Level | 31154
- Description | 31154
- Required Privilege Level | 31154
- Release Information | 31155

Syntax

```
xlat-source-rule {
                rule-set r1 {
                    rule r1;
                }
            }
```

Hierarchy Level

```
[edit services nat destination rule-set name rule name then destination-nat]
```

Description

Set the source NAT rule to match for NAT464

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

xle (Port)

IN THIS SECTION

- Syntax | 31155
- Hierarchy Level | 31155
- Description | 31155
- Options | 31156
- Required Privilege Level | 31156
- Release Information | 31156

Syntax

```
xle {
    (port port-number | port-range port-range-low port-range-high);
}
```

Hierarchy Level

```
[edit chassis fpc slot pic pic-number]
```

Description

(QFX3500 and QFX3600 standalone switches only) Configure a specific QSFP+ port or a range of QSFP + ports to operate as 40-Gigabit Ethernet (*xle*) type ports.

> **CAUTION**: The Packet Forwarding Engine on the switch is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the switch.

Options

port-number	Port number on which you want to configure the port type. On a QFX3500 standalone switch, specify a value from 0 through 3. On a QFX3600 standalone switch, specify a value from 0 through 15.
port-range-low	Lowest-numbered port in the range of ports. The lowest possible value is 0.
port-range-high	Highest-numbered port in the range of ports. The highest possible value is 3 on QFX3500 standalone switches, and 15 on QFX3600 standalone switches.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2X50-D20.

RELATED DOCUMENTATION

Configuring the QSFP+ Port Type on QFX3500 Standalone Switches

Configuring the Port Type on QFX3600 Standalone Switches

xnm-clear-text

IN THIS SECTION

Syntax | 31157

- Hierarchy Level | 31157
- Description | 31157
- Options | 31157
- Required Privilege Level | 31158
- Release Information | 31158

Syntax

```
xnm-clear-text {
    connection-limit limit;
    rate-limit limit;
}
```

Hierarchy Level

```
[edit system services]
```

Description

Allow Junos XML protocol clear-text requests from remote systems to the local router.

> **NOTE**: Junos OS Evolved supports the `xnm-clear-text` statement starting in Junos OS Evolved release 23.4R1.

Options

connection-limit *limit*

Configure the maximum number of connections sessions for the xnm-clear-text service per protocol (either IPv6 or IPv4).

> **NOTE**: The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

- **Range:** 1 through 250 connections
- **Default:** 75 connections

rate-limit *limit*

Configure the maximum number of connections attempts per minute, per protocol (either IPv6 or IPv4) on an access service. For example, a rate limit of 10 allows 10 IPv6 xnm-clear-text session connection attempts per minute and 10 IPv4 xnm-clear-text session connection attempts per minute.

- **Range:** 1 through 250 connections
- **Default:** 150 connections

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring clear-text or SSL Service for Junos XML Protocol Client Applications

xnm-ssl

IN THIS SECTION

- Syntax | 31159
- Hierarchy Level | 31159
- Description | 31159
- Options | 31160
- Required Privilege Level | 31161

Release Information | 31161

Syntax

```
xnm-ssl {
    connection-limit limit;
    local-certificate name;
    rate-limit limit;
    ssl-renegotiation ;
}
```

Hierarchy Level

```
[edit system services]
```

Description

Allow Junos XML protocol TLS requests from remote systems to the local router.

NOTE: Junos OS Evolved supports the `xnm-ssl` statement starting in Junos OS Evolved release 23.4R1. To use this feature on Junos OS Evolved devices you must first configure the following statement:

```
set interfaces lo0.0 family inet address 127.0.0.1/32
```

NOTE: When FIPS mode is enabled on a device, the `xnm-ssl` service does not support TLS 1.0. Additionally, on a device running Junos OS release 23.1R1, the `xnm-ssl` service only supports TLS 1.2 or later.

For a device in FIPS mode:

- Running Junos OS prior to release 23.1R1, clients can communicate with the `xnm-ssl` service using TLS 1.0 or later.

- Running Junos OS release 23.1R1 or later, clients can communicate with the `xnm-ssl` service using TLS 1.1 or later.

For a device in non-FIPS mode:

- Running Junos OS prior to release 15.1R8, clients can communicate with the `xnm-ssl` service using TLS 1.0 or later.
- Running Junos OS release 15.1R8 or later and prior to release 23.1R1, clients can communicate with the `xnm-ssl` service using TLS 1.1 or later.
- Running Junos OS release 23.1R1, clients can communicate with the `xnm-ssl` service using TLS 1.2 or later.

Options

connection-limit ***limit***

Configure the maximum number of connections sessions for the ftp service per protocol (either IPv6 or IPv4).

> **NOTE**: The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

- **Range:** 1 through 250 connections
- **Default:** 75 connections

local-certificate ***name***

Import or reference a TLS certificate by specifying the name of the local certificate to use.

There is no default. The value for `local-certificate` should be the same as the name provided during the import of the certificate using the CLI configuration statement `local` at the `[edit security certificates]` hierarchy level.

rate-limit ***limit***

Configure the maximum number of connections attempts per minute, per protocol (either IPv6 or IPv4) on an access service. For example, a rate limit of 10 allows 10 IPv6 ftp session connection attempts per minute and 10 IPv4 ftp session connection attempts per minute.

- **Range:** 1 through 250 connections
- **Default:** 150 connections

ssl-renegotiation

Enable TLS re-negotiation for `xnm-ssl` service.

- **Default:** Disabled

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`ssl-renegotiation` introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

Configuring clear-text or SSL Service for Junos XML Protocol Client Applications

Importing SSL Certificates for Junos XML Protocol Support

local

xon (Input Congestion Notification)

IN THIS SECTION

- Syntax | 31162
- Hierarchy Level | 31162
- Description | 31162
- Options | 31162
- Required Privilege Level | 31163
- Release Information | 31163

Syntax

```
xon value;
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name input dscp code-point code-point-bits]
[edit class-of-service congestion-notification-profile profile-name input ieee-802.1 code-point code-point-bits]
```

Description

Specifies XON threshold offset cells for resuming transmission from the dynamic shared buffer.

Each port can have lossy and lossless queues. You assign forwarding class sets (priority groups) to all the lossless queues at ingress. This enables transmitting priority-based flow control (PFC) frames in case of any congestion at the egress queues. Lossless priority groupgs (PGs) get their ingress buffers from the lossless service pool. The PGs dynamically access these buffers based on the alpha value configured. When the ingress shared occupancy reaches its limit, the corresponding PG transmits the PFC frames to the egress peer so that peer can stop transmitting packets. This limit is called XOFF and is not configurable.

The PFC XON threshold is the ingress port's PG shared buffer limit at which the ingress port's egress peer resumes transmitting packets after a brief PAUSE because of the PFC message sent by this ingress port.

WARNING: You cannot configure the PFC XON threshold without configuring PFC for the congestion notification profile.

Options

`xon value` XON threshold offset cells for resuming transmission from the dynamic shared buffer.

Range: 0-100000

Default: 20

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 22.3R1 for QFX5220 devices.

RELATED DOCUMENTATION

buffer-dynamic-threshold | 3047

pfc (Input Congestion Notification) | 12456

Y

IN THIS SECTION

- yang-compliant (NETCONF) | 31163
- yang-modules (NETCONF) | 31166
- yellow-differential-delay | 31168

yang-compliant (NETCONF)

IN THIS SECTION

- Syntax | 31164
- Hierarchy Level | 31164
- Description | 31164

Required Privilege Level | 31165
Release Information | 31166

Syntax

```
yang-compliant;
```

Hierarchy Level

```
[edit system services netconf]
```

Description

Require that the NETCONF server return YANG-compatible configuration data that is consistent with the device's YANG schema and RFC 6020, *YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)*.

> **NOTE**: You must configure the `yang-compliant` statement to enable OpenDaylight (ODL) controllers to manage devices running Junos OS.

When you configure the `yang-compliant` statement, the `<get-config>` and `<get-configuration format="xml">` remote procedure calls (RPCs) include the following changes to their output:

- Suppress unsupported configuration statements that are not defined in the YANG schema for that device (`<undocumented>` elements). This does not suppress deprecated statements.
- Suppress configuration comments (`<junos:comment>` elements).
- Serialize list keys according to the YANG schema for configuration statements that define `choice-ident` and `choice-value` as keys.
- Emit YANG annotations encoded as per RFC 7952.

Table 1840 on page 31165 illustrates the changes in the `<get-config>` and `<get-configuration format="xml">` RPC output.

Table 1840: yang-compliant Changes to Configuration Output

Change	Default Output	yang-compliant Output

Suppresses unsupported statements	`<processes>`
	`  <ntp>`
	`  <undocumented>`
	`  <enable/>`
	`  </undocumented>`
	`  </ntp>`
	`</processes>`
		`<processes>`
		`  <ntp>`
		`  </ntp>`
		`</processes>`

Suppresses comments	`<junos:comment>/* CustA */`
	`</junos:comment>`
	`<route>`
	`  <name>198.51.100.1/24`
	`  </name>`
	`  <next-hop>10.1.1.254`
	`  </next-hop>`
	`  <retain/>`
	`  <no-readvertise/>`
	`</route>`
		`<route>`
		`  <name>198.51.100.1/24`
		`  </name>`
		`  <next-hop>10.1.1.254`
		`  </next-hop>`
		`  <retain/>`
		`  <no-readvertise/>`
		`</route>`

Correctly serializes list keys that include `<choice-ident>` and `<choice-value>`	`<route-filter>`
	`  <address>0.0.0.0/0`
	`  </address>`
	`  <prefix-length-range>/32-/32`
	`  </prefix-length-range>`
	`</route-filter>`
		`<route-filter>`
		`  <address>0.0.0.0/0`
		`  </address>`
		`  <choice-ident>prefix-length-range`
		`  </choice-ident>`
		`  <choice-value>/32-/32`
		`  </choice-value>`
		`</route-filter>`

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3.

RELATED DOCUMENTATION

rfc-compliant (NETCONF) | 15755
ssh (NETCONF) | 28428

yang-modules (NETCONF)

IN THIS SECTION

- Syntax | 31166
- Hierarchy Level | 31167
- Description | 31167
- Default | 31167
- Options | 31167
- Required Privilege Level | 31168
- Release Information | 31168

Syntax

```
yang-modules {
    device-specific;
    emit-anyxml-in-rpc-output;
    emit-extensions;
    emit-family-ns-and-module-name;
}
```

Hierarchy Level

```
[edit system services netconf]
```

Description

Configure how the device running Junos OS serves the native YANG modules.

Default

If you do not include the `yang-modules` statement, the device running Junos OS serves the family-specific YANG data models that are shipped with the device.

Options

device-specific

Instruct the device to generate device-specific YANG data models instead of the family-specific YANG data models that are shipped with the device.

emit-anyxml-in-rpc-output

Instruct the device to generate YANG data models that use the `anyxml` statement in the Junos RPC output schemas. To configure this statement, you must also configure the `device-specific` statement.

> **NOTE**: Starting in Junos OS Release 23.1R1, the YANG data models for Junos RPCs replace the `anyxml` statement with accurate output schemas.

emit-extensions

Instruct the device to generate YANG data models that explicitly include Junos OS extension statements.

> **NOTE**: The device emits the `junos:command` extension statement starting in Junos OS Release 22.4R1 and Junos OS Evolved Release 22.4R1.

emit-family-ns-and-module-name

Instruct the device to generate YANG data models that use the original device family identifier in the module's name, namespace, and filename. To configure this statement, you must also configure the `device-specific` statement.

This statement is only applicable on devices where the default device family identifier changed from one release to the next.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

`emit-anyxml-in-rpc-output` statement added in Junos OS Release 23.1R1.

`emit-family-ns-and-module-name` statement added in Junos OS Evolved Release 23.4R1.

yellow-differential-delay

IN THIS SECTION

- Syntax | 31168
- Hierarchy Level | 31168
- Description | 31169
- Options | 31169
- Required Privilege Level | 31169
- Release Information | 31169

Syntax

```
yellow-differential-delay milliseconds;
```

Hierarchy Level

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
```

Description

For link services interfaces only, configure the yellow differential delay among bundle links to give warning when a link has a differential delay that exceeds the configured threshold.

Options

milliseconds—Yellow differential delay threshold.

- **Range:** 1 through 2000 milliseconds
- **Default:** 72 milliseconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16

action-red-differential-delay | **1782**

red-differential-delay | **14156**

Z

IN THIS SECTION

zone (System Security Profile) | **31170**

zones | **31172**

zone (System Security Profile)

IN THIS SECTION

- Syntax | 31170
- Hierarchy Level | 31170
- Description | 31170
- Options | 31171
- Required Privilege Level | 31171
- Release Information | 31171

Syntax

```
zone {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the zones that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 4000

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only)

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only)

zones

IN THIS SECTION

- Syntax | 31172
- Hierarchy Level | 31174
- Description | 31174
- Options | 31174
- Required Privilege Level | 31174
- Release Information | 31174

Syntax

```
zones {
    functional-zone {
        management {
            description text;
            host-inbound-traffic {
                protocols protocol-name {
                                except;
                }
                system-services service-name {
                                except;
                }
            }
            interfaces interface-name  {
                host-inbound-traffic {
                    protocols protocol-name {
                                    except;
                    }
                    system-services service-name {
                                    except;
                    }
                }
            }
            screen screen-name;
        }
```

```
    }
    security-zone zone-name {
        address-book {
            address address-name {
                ip-prefix {
                    description text;
                }
                description text;
                dns-name domain-name {
                    ipv4-only;
                    ipv6-only;
                }
                range-address lower-limit to upper-limit;
                wildcard-address ipv4-address/wildcard-mask;
            }
            address-set address-set-name {
                address address-name;
                address-set address-set-name;
                description text;
            }
        }
        advance-policy-based-routing;
        application-tracking;
        description text;
        host-inbound-traffic {
            protocols protocol-name {
                                except;
            }
            system-services service-name {
                                except;
            }
        }
        interfaces interface-name {
            host-inbound-traffic {
                protocols protocol-name {
                                    except;
                }
                system-services service-name {
                                    except;
                }
            }
        }
        screen screen-name;
```

```
        tcp-rst;
    }
}
```

Hierarchy Level

```
[edit security]
```

Description

A zone is a collection of interfaces for security purposes. All interfaces in a zone are equivalent from a security point of view. Configure the following zones:

- Functional zone—Special-purpose zone, such as a management zone that can host dedicated management interfaces.
- Security zone—Most common type of zone that is used as a building block in policies.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OS Release 11.1. The `description` option added in Junos OS Release 12.1.

RELATED DOCUMENTATION

Security Zones Overview

Supported System Services for Host Inbound Traffic

Others

IN THIS SECTION

- | (pipe) | 31175

| (pipe)

IN THIS SECTION

- Syntax | 31175
- Description | 31175
- Options | 31176
- Required Privilege Level | 31179
- Release Information | 31179

Syntax

```
 | (compare | count |
display (changed | commit-scripts | detail | inheritance | json | merge | omit | set |
translation-scripts <configured-delta | translated-config | translated-delta> | xml) |
except pattern | find pattern | hold | last lines | match     pattern | no-more | refresh
interval |
request message (all | account@terminal)
resolve <full-names> | save filename | append filename | tee | trim columns )
```

Description

Filter the output of an operational mode or a configuration mode command.

Options

append *filename*
Append the output to a file.

compare (filename | rollback *n*)
Compare configuration changes with another configuration file. In operational mode, use the `show configuration` command. In configuration mode, use the `show` command.

compare | display xml
Compare configuration changes with the active configuration and display them in XML format. In operational mode, use the `show configuration` command. In configuration mode, use the `show` command.

count
Display the number of lines in the output.

display
Display additional information about the configuration contents.

- **changed** — Tag changes with `junos:changed attribute` (XML only).
- **commit-scripts** — (Configuration mode only) Display all statements that are in a configuration, including statements that were generated by transient changes.
- **detail** — (Configuration mode only) Display configuration data detail.
- **inheritance <brief | default | no-comments | groups | terse>** — (Configuration mode only) Display inherited configuration data and source group.
- **json** — Display the output for operational commands and configuration data in JavaScript Object Notation (JSON) format.
- **merge** — Use with the `show ephemeral-configuration` command to display the merged view of the static and ephemeral configuration databases in Junos OS Release 18.1 and earlier releases. Issuing the `show ephemeral-configuration | display merge` command displays the configuration data from all instances of the ephemeral configuration database merged with the complete post-inheritance configuration.

Starting in Junos OS Release 18.2R1, to display the merged view of the static and ephemeral configuration databases, use the `show ephemeral-configuration merge` command instead.

omit (Configuration mode only) Display configuration statements omitted by the `apply-flags omit` configuration statement.

set Display the configuration as a series of configuration mode commands required to re-create the configuration.

translation-scripts Display the configuration with YANG translation scripts applied.

To view the complete post-inheritance configuration with the translated configuration data from all enabled YANG translation scripts included in the output, append the `| display translation-scripts` filter to the `show configuration` command in operational mode or the `show` command in configuration mode.

You can also append one of several keywords to display different views of the configuration data corresponding to the non-native YANG data models:

- `configured-delta`—In configuration mode, compare the candidate and active configurations, and display configuration changes in the statements or hierarchies corresponding to non-native YANG data models before any translation is applied. The XML output displays the deleted content, followed by the new content in the syntax defined by the YANG data model.
- `translated-config`—In operational or configuration mode, display all non-native configuration data present in the committed or candidate configuration, respectively, after processing by all enabled translation scripts into Junos OS syntax.
- `translated-delta`—In configuration mode, compare the candidate and active configurations, and display configuration changes in the statements or hierarchies corresponding to non-native YANG data models after translation is applied. The XML output displays the deleted content, followed by the new content in Junos OS syntax.

xml	(Operational mode only) Display the command output as Junos XML protocol (Extensible Markup Language [XML]) tags.
except *pattern*	Ignore text matching a regular expression when searching the output. If the regular expression contains spaces, operators, or wildcard characters, enclose it in quotation marks.
find *pattern*	Display the output starting at the first occurrence of text matching a regular expression. If the regular expression contains spaces, operators, or wildcard characters, enclose it in quotation marks (" ").
hold	Hold text without exiting the `--More--` prompt.
last *lines*	Display the last number of lines you want to view from the end of the configuration. However, when the number of lines requested is less than the number of lines that the screen length setting permits you to display, Junos returns as many lines as permitted by the screen length setting.
match *pattern*	Search for text matching a regular expression. If the regular expression contains spaces, operators, or wildcard characters, enclose it in quotation marks.
no-more	Display output all at once rather than one screen at a time.
`resolve`	(Operational mode only) Convert IP addresses into Domain Name System (DNS) names. Truncates to fit original size unless `full-names is` specified. To prevent the names from being truncated, use the `full-names` option.
refresh *interval*	Refresh the display of the command according to the interval specified. The screen gets refreshed periodically to show you the current output of the command until you quit the command. The default refresh interval is one second. However, you can also explicitly specify a value from `1` through `604800` for the refresh interval.
request message (all | *account@terminal*)	Display command output on the terminal of a specific user logged in to your router, or on the terminals of all users logged in to your router.
save *filename*	Save the output to a file or URL.
tee	Allows you to both display the command output on screen and write it to a file. Unlike the UNIX `tee` command, if the file cannot be opened, just an error message is displayed.
trim *columns*	Trim specified number of columns from the start line. Only positive values are accepted. An error message appears if a negative value is given.

Required Privilege Level

view

Release Information

Command introduced before Junos OS Release 7.4.

`display commit-scripts` option added in Junos OS Release 7.4.

`tee` option added in Junos OS Release 14.1.

`display json` option added in Junos OS Release 14.2.

`compare | display xml` option added in Junos OS Release 15.1.

`display translation-scripts` option added in Junos OS Release 16.1.

`display merge` option added in Junos OS Release 16.2R2.

`display merge` option deprecated in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Displaying the Current Configuration

Using the Pipe (|) Symbol to Filter Command Output

Comparing Configurations and Displaying the Differences in Text

Understanding the show | compare | display xml Command Output

Junos OS Automation Scripting User Guide

Displaying Output Beginning with the Last Entries

Viewing Files and Directories on a Device Running Junos OS